

Towards 3D Vision with Low-Cost Single-Photon Cameras

Fangzhou Mu[†], Carter Sifferman[†], Sacha Jungerman, Yiquan Li[‡], Mark Han[‡], Michael Gleicher, Mohit Gupta, Yin Li

[†] co-first author, [‡] equal contribution

University of Wisconsin-Madison

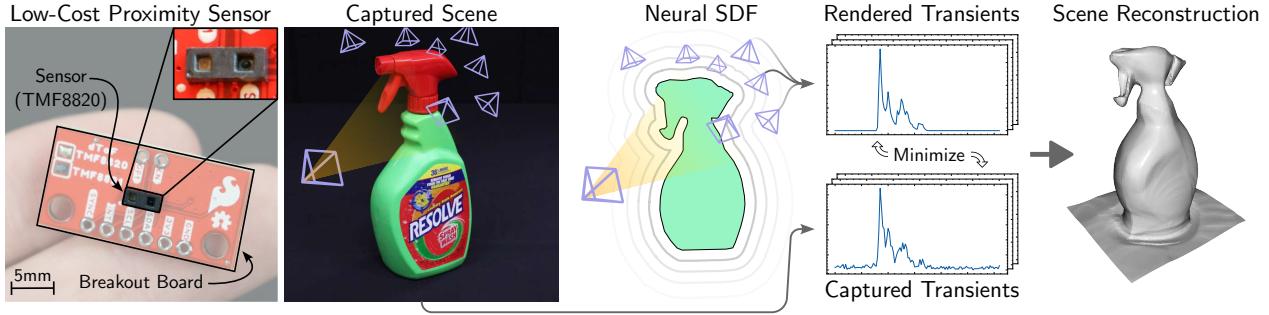


Figure 1. We demonstrate that measurements from spatially distributed low-cost single-photon proximity sensors (left) can be used to reconstruct 3D shape of real world objects (right). Our method combines a differentiable image formation model and neural rendering to recover 3D geometry based on measurements (transient histograms) from sensors with known poses. This is done by minimizing the difference between the observed and rendered sensor measurements. For clarity, a subset of sensor poses and measurements are shown.

Abstract

We present a method for reconstructing 3D shape of arbitrary Lambertian objects based on measurements by miniature, energy-efficient, low-cost single-photon cameras. These cameras, operating as time resolved image sensors, illuminate the scene with a very fast pulse of diffuse light and record the shape of that pulse as it returns back from the scene at a high temporal resolution. We propose to model this image formation process, account for its non-idealities, and adapt neural rendering to reconstruct 3D geometry from a set of spatially distributed sensors with known poses. We show that our approach can successfully recover complex 3D shapes from simulated data. We further demonstrate 3D object reconstruction from real-world captures, utilizing measurements from a commodity proximity sensor. Our work draws a connection between image-based modeling and active range scanning, and offers a step towards 3D vision with single-photon cameras. Our project webpage is at https://cpsiff.github.io/towards_3d_vision/.

1. Introduction

Reconstructing 3D shape of real objects remains a central problem in vision, solutions to which have evolved into two parallel branches. *Image-based modeling* [45] leverages a plethora of visual cues from multiple photographs

(e.g., stereo, motion, shading), leading to problems including multi-view stereo [42], photometric stereo [1] and the more recent neural radiance fields (NeRF) [33]. Conversely, *active range scanning* [22] combines an active light source with an imaging sensor, giving rise to imaging techniques such as structured light [13], and time-of-flight [16]. Conventional wisdom suggests that range scanning yields more precise 3D geometry than image-based modeling at the cost of using specialized, expensive hardware.

An emerging approach for range scanning is direct time-of-flight imaging with active *single-photon cameras*, a form of time-resolved image sensor. This approach couples a pico-to-nanosecond detector with a fast coherent light source, illuminates the scene with a very short pulse of light, and measures the intensity of the light over time as it reflects back from the scene. The resulting incident wavefront is recorded and quantized, forming a *transient histogram*. A special case of this approach is single-photon LiDAR, in which the light source (laser) is highly focused, the detector finds the peak in the histogram, and the sensor reports a single distance value per detector pixel. When using a diffuse light source, these time-resolved sensors capture visual information beyond the distance measurements extracted in LiDAR. Transient histograms in this case record distributions of times-of-flight, encoding the product of scene geometry and reflectance over each imaged scene patch [23].

The entirety of information in the transient histogram has been previously utilized in applications such as fluorescence

lifetime imaging [26] and non-line-of-sight imaging [48]. Recently, time-resolved measurements from such sensors have been explored for line-of-sight 3D reconstruction and novel view synthesis [30]. Unfortunately, these systems, due to their prototype-like nature, exist in a bulky benchtop form factor, require significant power, and are costly (over \$10,000 USD). They are thus not suited for many applications including autonomous drones, wearable computing, and augmented reality, in which low-cost, energy efficient, miniature sensors are required.

Low-cost single-photon cameras have recently become available in the form of active single photon avalanche diodes (SPADs). They include one or more SPADs paired with an eye-safe diffuse light source (*e.g.*, an infrared VCSEL laser), are very small (<20 mm³), inexpensive ($<\$5$ USD), and power efficient (<10 milliwatts per measurement) [2, 32]. These are sold as proximity sensors and some can be configured to report transient histograms. Compared to laboratory-grade systems, however, they lack precise optics, calibration, and timing characteristics and have an order of magnitude lower spatial and temporal resolution. Still, these sensors have proven successful for material classification [4], human pose recognition [41], and simple shape recovery (*i.e.*, a planar surface) [23, 43].

In this work, we address the problem of reconstructing 3D shape of arbitrary Lambertian objects from a set of spatially distributed low-cost single-photon cameras with known poses. We present an approach that combines a neural signed distance field surface representation, a differentiable transient formation model for practical active single-photon cameras, and an optimization scheme following the analysis-by-synthesis pipeline. Fig. 1 illustrates our sensor, imaging setting, and approach. We show that our approach can successfully recover complex 3D shapes from simulated data. We further demonstrate 3D object reconstruction from real-world captures, utilizing measurements from a low-cost, off-the-shelf proximity sensor.

Our approach draws a connection between image-based modeling and active range scanning while avoiding some of their pitfalls. By using an active light source and SPADs, our approach illuminates and images an object from multiple angles, resembling the imaging setup of multiview photometric stereo. Our model also echoes the design of NeRF. The key difference is that our model considers an input of low spatial resolution (*i.e.*, hundreds of pixels per scene), and leverages rich temporal information encoded in the transient histograms. Unlike image-based modeling, our approach can operate under low light conditions, and reconstruct objects without texture. Unlike most range scanning methods, our approach performs favorably in the presence of strong ambient flux and is robust to mildly specular surfaces (see top of the spray bottle in Fig. 1).

Practical Applications. Our approach presents a compact,

energy-efficient, and low-cost solution for 3D sensing when relative sensor poses are known. We envision applying our method to settings where relative sensor poses are fixed, and 3D geometry can be estimated on a per-frame basis by combining measurements from a distributed set of sensors. This includes applications like wearable hand tracking [10, 20], and collision avoidance and mapping for drones, mobile robots, and robot manipulators. This work provides a first step towards enabling such applications.

Scope and limitation. Similarly to other methods for active range scanning, the proposed method fails on highly specular objects and requires hundreds of views to reconstruct a complex 3D shape. Our imaging system is tailored for low-cost sensors with limited range (4-6 meters) and temporal resolution (over 100 picoseconds), and thus it is unproven at imaging larger space or finer details. While methods exist for fast NeRF [35], our approach is not optimized for speed and takes a few hours to reconstruct a single object. Instead, we focus on demonstrating the feasibility of reconstructing real-world 3D objects using commodity sensors.

2. Related Work

Time-resolved imaging. Time-resolved sensors have long been used in applications such as non-line-of-sight (NLOS) imaging [11, 50], where a scene is recovered from around the corner, fluorescence lifetime imaging [27], a microscopy technique for characterizing a biological sample, and to measure distance using direct time-of-flight [15, 17]. Recent techniques [30] used high-end time-resolved sensors to directly recover scene geometry. However, these methods rely on hardware prototypes which are often prohibitively expensive and not accessible to consumers. For example, Transient NeRF [30] performs view synthesis and 3D reconstruction using a lab prototype with no ambient light and 2-5 views of 512×512 transient histograms, each with 1500 bins. In contrast, in this work, we use commodity hardware with ambient light and 128-240 total transients, each comprised of 128 bins.

3D Imaging with low-cost SPADs. With cheap time-resolved proximity sensors becoming commonplace, recent works have investigated their use for 3D reconstruction. Callenberg *et al.* [6] demonstrate that, with some additional hardware, high-resolution depth imaging from a single viewpoint is possible. A low-cost SPAD has been used to augment an RGB SLAM system [51]. Other works utilize supervised machine learning to recover geometric information from single low-cost SPAD measurements, such as 3D human pose [41] or high-resolution depth images [23, 53]. Jungerman *et al.* [23] use differentiable rendering to recover two degrees of freedom of a planar surface from a single low-cost SPAD transient histogram. Sifferman *et al.* [43] extend this method to fully recover a planar surface

using a low-cost SPAD with multiple detector pixels.

Neural implicit representations. Neural representations, as popularized by NeRF [33], enable novel view synthesis and 3D reconstruction by representing the scene as a neural network. While the original NeRF representation encoded view-dependent volumetric effects, alternative encodings have been proposed to better model geometry and reconstruct surfaces. NeuS [49] represents the scene as a level set, allowing for better modeling of surfaces at the expense of not being able to represent volumetric effects. Many works extend these ideas to work with different sensing modalities and external supervision, such as depth queues from structure-from-motion [9], RGB images plus continuous-wave time-of-flight sensors [56], only depth information [29, 37], or more recently using only transients [30, 60]. In this work, we perform 3D reconstruction using only transient histograms as captured using commodity hardware by adapting the implicit surface representation introduced by NeuS to render transients.

Multiview photometric stereo. Our approach illuminates and images an object from multiple views in order to reconstruct its 3D shape, having conceptual similarity to image-based modeling methods. This concept has previously been explored as multiview photometric stereo (MVPS) [7, 18]: reconstructing 3D geometry given distributed views of the scene and various lighting conditions. More recent approaches resort to deep learning models [55] and consider NeRF-based representations [24, 52].

3. 3D Vision with Single-Photon Cameras

We propose to reconstruct 3D scene geometry using a sparse set of measurements from low-cost time-resolved SPAD sensors. Our approach assumes a distributed set of N SPAD sensors with known pose, each comprising a single-pixel detector co-located with a diffuse laser. The time of flight, t , of returning photons, received over the detector FoV ¹ is recorded and binned into *transient histograms* [23, 43], $\{\mathbf{h}_j\}_{j=1}^N$: time-resolved measurements that encode rich information about the scene geometry and reflectance. Our goal is to recover scene geometry from these measurements.

Traditional depth ranging. Our setup differs from the conventional use of SPADs for depth ranging [15, 17, 30], which assumes a collimated laser and focused detector. There, the depths of individual scene points can be independently estimated from separate transients, and a dense scan of thousands of scene points is needed to obtain full scene geometry. These systems require high-quality beam lasers and detectors with high temporal resolution. In contrast, our imaging configuration is indicative of the real-world use

¹We assume that the illumination and detector FoV coincide and refer to both as the sensor FoV in the rest of the paper.

cases of low-cost SPAD sensors with *diffuse lasers*, wide FoVs, and low-resolution detectors. In our case, the transient encodes rich information beyond its peak location.

Low-cost SPAD sensing. Our setup extends recent systems [23, 43] built around low-cost SPAD sensors for 3D sensing in two ways. First, our system consists of multiple posed sensors as opposed to a single sensor. Second, it is capable of capturing and reconstructing complex, non-parametric scenes as opposed to the parametric geometry of a single plane. With these upgrades, our system is reminiscent of multi-view stereo and NeRF-like systems for conventional RGB cameras, and is a substantial step towards practical 3D vision using commodity proximity sensors.

3.1. Transient Formation Model

We model the transient formation process in two steps. We first derive the *transient waveform* of a scene by modeling the interaction of light with scene geometry. The next step uses this waveform as the input to a sensor model that accounts for non-linearities that occur when capturing *transient histograms* with real hardware.

Transient waveform. As in [23, 43], we ignore high-order light paths and only model direct reflection from the object surface. The transient waveform $\tau_s(t)$ of a sensor s located at the origin can be written in path integral form [47] as

$$\tau_s(t) = \int_S \rho(\mathbf{x}) \frac{G(\mathbf{x}, \omega)}{\|\mathbf{x}\|^4} \delta\left(\|\mathbf{x}\|_2 - \frac{ct}{2}\right) dA(\mathbf{x}), \quad (1)$$

$$G(\mathbf{x}, \omega) = f_r(\mathbf{x}, \mathbf{n}_x, \omega) V(\mathbf{x}) \langle \omega, \mathbf{n}_x \rangle^2,$$

where $dA(\mathbf{x})$ is an infinitesimal area around point \mathbf{x} on the object surface S , ρ is the albedo, δ is the Dirac delta function, c the speed of light, and ω is a ray direction from \mathbf{x} to s . Time-independent effects such as V , the visibility function governed by scene geometry and sensor FoV, f_r the bidirectional reflectance distribution function (BRDF), and foreshortening effects where \mathbf{n}_x is the surface normal at \mathbf{x} , are modeled as $G(\mathbf{x}, \omega)$.

Sensor model. The sensor model accounts for laser and detector characteristics when converting a transient waveform into a transient histogram [19]. Our sensor model considers the laser pulse, laser power, detector quantum efficiency, ambient photon flux, internal detector noise, pile-up effect, and time jitter. In practice, the *laser pulse* is not a perfect impulse and, despite bandpass filters, the measured transient also captures some constant ambient light. To model this, we convolve $\tau(t)$ ² with the laser's impulse response $g(t)$, scaled by ϕ^{scale} which absorbs laser power and quantum efficiency of the detector, and then offset its intensity by ϕ^{bkgd} which encapsulates ambient photon flux and internal detector noise:

$$\tilde{\tau}(t) = \phi^{scale} (\tau * g)(t) + \phi^{bkgd}. \quad (2)$$

²We drop the subscript s hereafter for clarity.

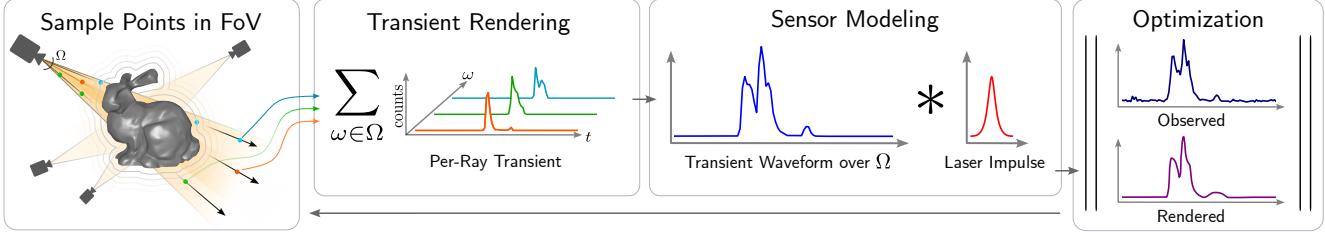


Figure 2. Method Overview: The scene is modeled as a neural implicit surface in the form of an SDF. To render a transient, we approximate Eq. 8 by sampling rays within each pixel’s FoV, and subsequently points on those rays. This idealized transient waveform is then convolved with the sensor’s laser impulse response to model the transient histogram formation. Finally, we optimize the scene representation by minimizing a loss between the rendered transients and the observations.

$\tilde{\tau}(t)$ is subsequently discretized into a histogram of Poisson rates $\mathbf{r} = [r_1, \dots, r_B]$ with B bins. The probability q_i of at least one photon falling inside the i^{th} bin is given by [8]

$$q_i = 1 - \exp(-r_i). \quad (3)$$

In practice, SPADs aggregate photon counts over C laser cycles, with only the first incident photon being detected in each cycle. This results in *pile-up*, a nonlinear distortion of transients, leaving photons arriving at a later timestamp less likely to be detected [39]. Specifically, the probability p_i of detecting a photon in the i^{th} bin in a cycle is given by [39]

$$p_i = q_i \prod_{k=1}^{i-1} (1 - q_k). \quad (4)$$

The photon counts $[h_1, \dots, h_B]$ in a transient histogram $\tilde{\mathbf{h}}$ follow a multinomial distribution:

$$[h_1, \dots, h_{B+1}] \sim \text{Multinomial}(C, (p_1, \dots, p_{B+1})), \quad (5)$$

where $p_{B+1} = 1 - \sum_{i=1}^B p_i$, and h_{B+1} counts the number of cycles without detected photons. $\tilde{\mathbf{h}}$ is subsequently convolved with a discretized *time jitter* kernel \mathbf{s} to account for the temporal uncertainty of photon detection events, yielding the final histogram \mathbf{h} measured by a SPAD detector:

$$\mathbf{h}[b] = \sum_k \mathbf{h}[k] \mathbf{s}[b - k]. \quad (6)$$

We now present a reconstruction algorithm based on the analysis-by-synthesis principle to recover scene geometry from the transients. The input to this system is a distributed set of posed transients $\{\mathbf{h}_j\}_{j=1}^N$ captured of the scene.

3.2. Neural Scene Reconstruction from SPADs

The reconstruction problem that we attempt to solve is extremely challenging. Unlike in conventional depth ranging where the depth of a scene point can be directly determined from the histogram via peak finding, a transient from our system represents the superposition of light reflected from numerous scene points, as illuminated by a diffuse laser, and is further contaminated by non-idealities of the detector (*e.g.*, pile-up). The direct inversion of the signal is thus a

highly ill-posed problem. Further, we cannot adapt methods from the NLOS imaging literature [28, 38, 50] as they only support dense 2D scans, whereas our system uses a distributed, sparse and unstructured set of measurements.

To overcome these challenges, we resort to an analysis-by-synthesis approach based on differentiable rendering. Our approach allows flexible positioning of sensors and accurate modeling of histogram formation, thereby enabling high-quality reconstruction of scene geometry. We now describe our reconstruction algorithm in detail.

Neural scene representation. Following NeuS [49], we represent the scene geometry as a signed distance function (SDF), parameterized as a multi-layer perceptron (MLP) $f_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$. f_θ maps the position-encoded (PE) xyz coordinates of a point \mathbf{x} to its signed distance d :

$$d = f_\theta(\text{PE}(\mathbf{x})). \quad (7)$$

Compared to [23, 43], this neural SDF allows our method to represent scene geometry beyond simple parametric shapes as the level set $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 | f_\theta(\mathbf{x}) = 0\}$.

Transient volume rendering. The key idea behind our analysis-by-synthesis approach is to render f_θ into transients and compare them with those captured by our system. To adapt Equation 1 for the rendering of f_θ , we first rewrite it in angular integral form as

$$\tau(t) = \int_{\Omega} \frac{\rho}{\pi} \frac{V(\mathbf{x}) \langle -\omega, \mathbf{n}_x \rangle}{\|\mathbf{x}\|^2} \delta \left(\|\mathbf{x}\|_2 - \frac{ct}{2} \right) d\omega, \quad (8)$$

where ω are ray directions in the sensor FoV Ω , and \mathbf{x} the point where ω intersects with the object surface \mathcal{S} (∞ if no intersection). For simplicity, we assume a learned spatially uniform albedo ρ and Lambertian BRDF $f_r = 1/\pi$ ³.

Inspired by NeRF [33] and NeuS [49], we approximate Equation 8 via volume rendering to resolve surface discontinuities, enabling the optimization of θ via gradient descent:

$$\hat{\tau}(t) = \int_{\Omega} \frac{\rho}{\pi} \frac{T^2(t) \sigma(\mathbf{p}(\omega, t)) \langle -\omega, \mathbf{n}_p \rangle}{\|\mathbf{p}(\omega, t)\|^2} d\omega. \quad (9)$$

³This assumption may be relaxed to allow more expressive BRDF models such as the Phong reflection model [5].

Here, $\mathbf{p}(\omega, t) = ct/2\omega$ are points along ω , the volume density σ is a function of f_θ as in NeuS [49], and the transmittance T is given by

$$T(t) = \exp \left(- \int_0^t \sigma(\mathbf{p}(u)) du \right). \quad (10)$$

In practice, we discretize $\hat{\tau}(t)$ over the transient bin intervals $\{[t_i, t_{i+1}]\}_{i=1}^B$ and work with the histogram $\hat{\tau} = [\hat{\tau}_1, \dots, \hat{\tau}_B]$, where

$$\hat{\tau}_i = \int_{\Omega} \frac{\rho}{\pi} \int_{t_i}^{t_{i+1}} \frac{T^2(t) \sigma(\mathbf{p}(\omega, t)) \langle -\omega, \mathbf{n}_p \rangle}{\|\mathbf{p}(\omega, t)\|^2} dt d\omega. \quad (11)$$

We estimate the intractable Equation 11 via Monte Carlo sampling of ω and subsequently of $\mathbf{p}(\omega, t)$.

Bilevel importance sampling. Similar to [33, 49], the sampling of $\mathbf{p}(\omega, t)$ is weighted by a probability density function (PDF) over the equally sized bin intervals. This PDF is proportional to the per-bin weights w_i given by

$$w_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \Delta) (1 - \exp(-\sigma_i \Delta)), \quad (12)$$

where σ_i is evaluated at the mid-point of the i^{th} bin, and Δ is the bin size in distance.

We extend this idea to the importance sampling of ω . Specifically, the sampling PDF over a uniform partitioning of FoV Ω is proportional to the cumulative weights $w^{(k)}$ over rays $\omega^{(k)}$ drawn from each partition k :

$$w^{(k)} = \sum_{i=1}^B w_i^{(k)}. \quad (13)$$

Intuitively, this allows us to point more rays at high-density regions occupied by the object surface.

Differentiable sensor modeling. Modeling sensor behavior is particularly important for our analysis-by-synthesis approach. This is because the synthesis targets τ are not determined by the scene geometry alone but reflect the complex interplay of geometry with sensor non-idealities including pulse shape, pile-up and time jitter. To this end, we cascade $\hat{\tau}$ to a *differentiable sensor model* Γ to simulate the transformation applied by the sensor to raw waveforms.

Specifically, Γ closely follows the sensor model in Section 3.1; Equations 2-4 are differentiable and applied sequentially on $\hat{\tau}$, yielding per-bin photon detection probabilities $\hat{\mathbf{p}} = [\hat{p}_1, \dots, \hat{p}_B]$. Instead of sampling photon counts using Equation 5, we directly convolve $\hat{\mathbf{p}}$ with the jitter kernel as in Equation 6. This allows us to sidestep the non-differentiable sampling step while producing an unbiased estimate of the transient $\hat{\mathbf{h}} = \Gamma(\hat{\tau})$ for loss evaluation.

Loss functions. The optimization of θ is driven by three loss terms: a histogram reconstruction loss $\mathcal{L}_{\text{hist}}$ that minimizes the L1 distance between $\hat{\mathbf{h}}$ and \mathbf{h} , an Eikonal loss [14]

$\mathcal{L}_{\text{Eikonal}}$ that regularizes the SDF, and a total variation regularizer [34] \mathcal{L}_{TV} that penalizes floaters in empty space. The combined loss function \mathcal{L} is thus given by

$$\mathcal{L} = \mathcal{L}_{\text{hist}} + \lambda_{\text{Eikonal}} \mathcal{L}_{\text{Eikonal}} + \lambda_{\text{TV}} \mathcal{L}_{\text{TV}}, \quad (14)$$

where λ_{Eikonal} and λ_{TV} are the respective loss weights.

Comparison to NLOS-NeuS [12]. NLOS-NeuS is a concurrent work for scene reconstruction using SPAD sensors. Despite similarities in the imaging model and reconstruction algorithm, the two works differ substantially in application (direct line-of-sight *vs.* non-line-of-sight), laser characteristics (eye-safe diffuse laser *vs.* high-energy beam laser), sensor quality and cost (commodity, low-cost *vs.* laboratory-grade, high-cost), transient modeling (with *vs.* without sensor non-idealities) and scan pattern (distributed and sparse *vs.* dense and structured 2D grid). We believe the two works complement each other and together unveil an exciting avenue toward single-photon 3D vision systems.

4. Experiments

We demonstrate the effectiveness of our method for 3D geometric reconstruction of various objects in simulation, and in the real world with a low-cost SPAD. We provide qualitative and quantitative results for both settings. See the supplement for implementation details (*e.g.* learning rates).

Baselines. We compare our method to two baselines: reprojection [15, 17] and space carving [25, 46].

Reprojection, also known as back-projection, reconstructs a scene as a point cloud and is the *de facto* standard for depth ranging. We compare to two forms of reprojection. The peak method finds the distance d corresponding to the histogram bin with the highest intensity. For a sensor at position \mathbf{s} with an outwards pointing optical axis \mathbf{u} , a point is placed in the scene at position $\mathbf{s} + d\mathbf{u}$. The threshold method works in the same way but finds d by locating the lowest-index bin with intensity above a threshold t_p . If no bin passes the threshold, no point is projected. We do not apply surface reconstruction to the generated point clouds, as it is often unreliable for poorly imaged scenes. See the supplement for further discussion.

Space carving reconstructs a scene as a voxel grid. Like thresholded reprojection, it finds the distance d corresponding to the lowest-index bin with intensity above a threshold t_s . All voxels in the sensor's FoV and nearer than d are marked empty, along with voxels outside the FoV. Voxels in the FoV and further than d are marked as occupied. The carved scene is the union of the occupied set for all sensors.

To ensure strong baselines for real-world experiments, we perform a brute-force search over t_p and t_s and choose values that minimize Chamfer distance over the entire real-world dataset. Space carving voxel size was set to 1.0cm.

Method	Chamfer Distance (mm) ↓							
	Armadillo	Bear	Bunny	Digit	Einstein	Skull	Soap	Sphere
Reprojection (Peak)	54.29	40.36	34.95	55.85	43.25	48.90	51.71	51.07
Reprojection (Threshold)	65.43	60.72	54.05	60.64	61.31	65.14	68.74	63.16
Space Carving	34.78	24.53	22.29	45.44	26.60	25.49	21.44	25.47
Ours	3.93	5.95	3.84	3.27	3.51	3.22	3.23	3.77

Table 1. **Quantitative results on simulated data.** Our method more accurately recovers 3D shapes than baselines across 8 objects.

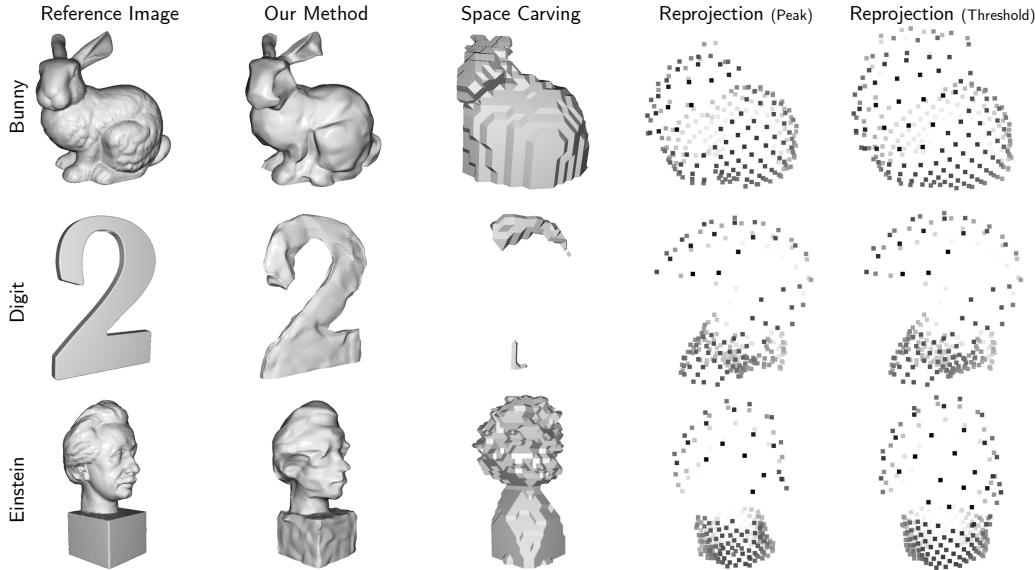


Figure 3. **Qualitative results on simulated data.** Our method reconstructs dense and detailed 3D shapes. Space carving provides only hulls of a target shape, and is prone to carving away extra space when thin structures are present. Reprojection yields sparse points.

Evaluation protocol. Following NeuS [49], we evaluate all methods using Chamfer distance. We report standard (two-way) Chamfer on simulated data. For real-world captures, we report Chamfer in both directions to evaluate the quality of reconstruction. Prior to Chamfer calculation, we convert ground-truth meshes and reconstructions from our method to point clouds by drawing 5 million points uniformly at random on the mesh surface. For space carving, occupied voxels are converted to points if they touch unoccupied space, excluding the edge of the grid.

4.1. Simulated Experiments

Experiment setup. We simulate transients for eight scenes of varying complexity using the image formation model in Section 3.1. The objects are centered on the ground plane ($z = 0$) with the largest dimension $\approx 0.3\text{m}$. Sensors with a conical FoV are uniformly distributed on a hemisphere at the origin with a radius of 0.5m, and are all pointed at the origin. In our simulation, $N = 256$, $B = 256$, $\Delta = 5\text{mm}$, $\text{FoV} = 30^\circ$, $\phi^{scale} = 1$, $\phi^{bkgd} = 0.001$, $C = 5000$ and $\rho = 0.8$. The laser pulse, g , has a full-width-at-half-maximum (FWHM) of 50ps, and s is a tabulated PDF obtained from experiments [19]. The sensor parameters are deliberately chosen to reflect the characteristics of low-cost

sensors. See the supplement for a sensitivity study on how these parameters affect the quality of reconstruction.

Results. Table 1 summarizes the quantitative results of all methods. Our method achieves an average Chamfer distance of $< 5\text{mm}$, an order of magnitude lower than all baselines. A key reason is that the baselines only use depth information from a single histogram bin, whereas our method makes effective use of the entire waveform, which contains rich geometry cues about a large scene patch.

We provide visualizations of our results in Figure 3. Our method recovers global scene structure as well as local geometry details. In contrast, reprojection yields sparse point clouds without sufficient coverage of the scene. While space carving produces dense reconstructions, the occupancy grid only represents an envelope of the scene, leaving it difficult to recognize the precise shape of an object.

4.2. Hardware Prototype

We use the SPAD-based AMS TMF8820 proximity sensor [2], which retails for \$10 USD. We connect the sensor to a microcontroller via I²C and use the AMS-provided driver to extract transient histograms.

The sensor contains a total of 216 SPADs, which are pooled onboard the sensor into 3×3 zones, each of which

Method	Chamfer Distance (mm) ↓									
	Big Box*		Block*		Pyramid*		Toy Container†		Cereal Box†	
Reprojection (Peak)	77.4	24.9/52.5	51.8	12.8/39.0	94.7	17.5/77.1	71.0	24.9/46.0	49.3	17.3/31.9
Reprojection (Threshold)	67.5	14.8/52.7	52.3	8.5/43.8	75.4	5.9/69.5	52.4	8.9/43.5	51.6	19.2/32.4
Space Carving	67.9	35.1/32.8	69.2	33.4/35.8	80.1	39.5/40.6	98.9	52.8/46.0	44.1	24.4/19.6
Ours	12.5	6.1/ 6.4	9.8	5.6/ 4.2	18.4	9.0/ 9.3	11.5	5.8/ 5.6	16.3	8.3/ 8.0

Table 2. **Quantitative results on real-world captures.** Our method more accurately reconstructs real-world objects with homogeneous (*) and rich (†) texture. Reprojection yields sparse and unevenly distributed points, harming one-way Chamfer from GT to reconstruction.

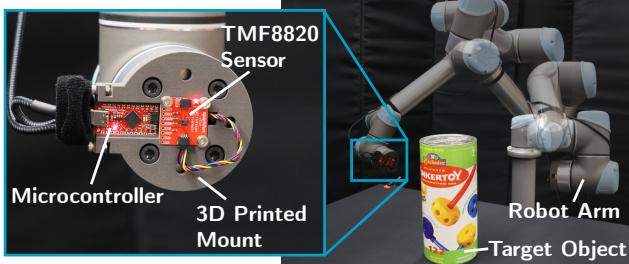


Figure 4. To capture real-world data from a wide set of viewpoints, we mount the TMF8820 proximity sensor to a robot arm. Forward kinematics of the robot are used to gather sensor pose.

images a different FoV. The sensor captures one transient histogram for each zone. We pool histograms from all zones, which is equivalent to capturing one wide-FoV histogram per-measurement as SPADs do not suffer from readout noise [54]. In doing so, we avoid inter-histogram interference previously observed by [43] and avoid the need to model individual fields-of-view of the sensor, which we empirically observed to have soft and poorly specified boundaries. We slightly modify our method to accommodate the AMS TMF8820 sensor used in real-world experiments.

Laser Impulse. The laser impulse response of the TMF8820 is not Gaussian and varies slightly between measurements. Fortunately, the sensor captures the shape of its laser impulse for each measurement in a “reference histogram”. We record this histogram for each measurement and incorporate it into our forward model by cross-correlating the idealized scene response with this recorded reference histogram. We observe that the bin size Δ_r of the reference histogram is smaller than the bin size Δ of the transient histograms captured by the sensor. To account for this, we scale the reference histogram in the temporal dimension by a factor Δ_r/Δ before cross-correlation. Further, we find that it is necessary to temporally shift the reference histogram by a fixed amount ϕ^{delay} before correlation.

To calibrate the parameters Δ , Δ_r , and ϕ^{delay} , we perform the one-off intrinsic calibration procedure separately introduced by [43]. The TMF8820 sensor is pointed at a planar surface from a range of known distances and angles-of-incidence. A differentiable render-and-compare method is used to optimize for the unknown sensor intrinsic parameters given known planar geometry.

Pile-up Correction. While our forward model assumes that the target transients exhibit nonlinear distortion due to pile-up, the TMF8820 sensor performs pile-up correction on-sensor, and it cannot be disabled. To accommodate this, we incorporate the differentiable Coates’ correction [8] as a final step in the forward model.

Other Sensors. Our design can be easily applied to any SPAD with a co-located diffuse illumination source, including other low-cost sensors [31], high-end setups [23, 36], and things in between. The only modification necessary is the calibration of sensor intrinsics (*e.g.*, using [43]).

4.3. Real-world Experiments

Experiment Setup. We capture a real-world tabletop dataset of eight objects⁴ of varying geometry and texture. To capture many posed views of the target object, we attach the sensor to a Universal Robots UR5 robot arm. We program the arm to automatically move to a set of poses and record sensor measurements at each pose. To obtain sensor poses, we use the forward kinematics of the robot, which are accurate to $\pm 0.5\text{mm}$ [40]. Each object is captured from between 128 and 240 viewpoints. Five of the objects are simple geometric primitives, for which we manually generate ground-truth meshes based on the dimensions of the target object and measurements of its position from the robot’s forward kinematics. Meshes are trimmed to an axis-aligned bounding box 16cm larger than the target object in each dimension before the Chamfer distance calculation.

Results. As seen in Table 2, our method outperforms all baselines by a wide margin as measured by two-way Chamfer distance. While reprojection is at times competitive in one-way distance from reconstruction to ground truth, it performs poorly in the opposite direction due to the sparse and unevenly distributed point cloud generated, as visualized in Figure 5. While space carving outperforms reprojection on simulated data under a highly structured sensor pose distribution (*i.e.* all sensors are facing the center of the object), it yields poor results on real-world scenes, in which we vary sensor orientation by $\pm 10^\circ$ to emulate real-world capture conditions and increase coverage. By contrast, our method benefits from the more varied sensor poses as is shown in Figure 6.

⁴See supplement for full results over the entire dataset.

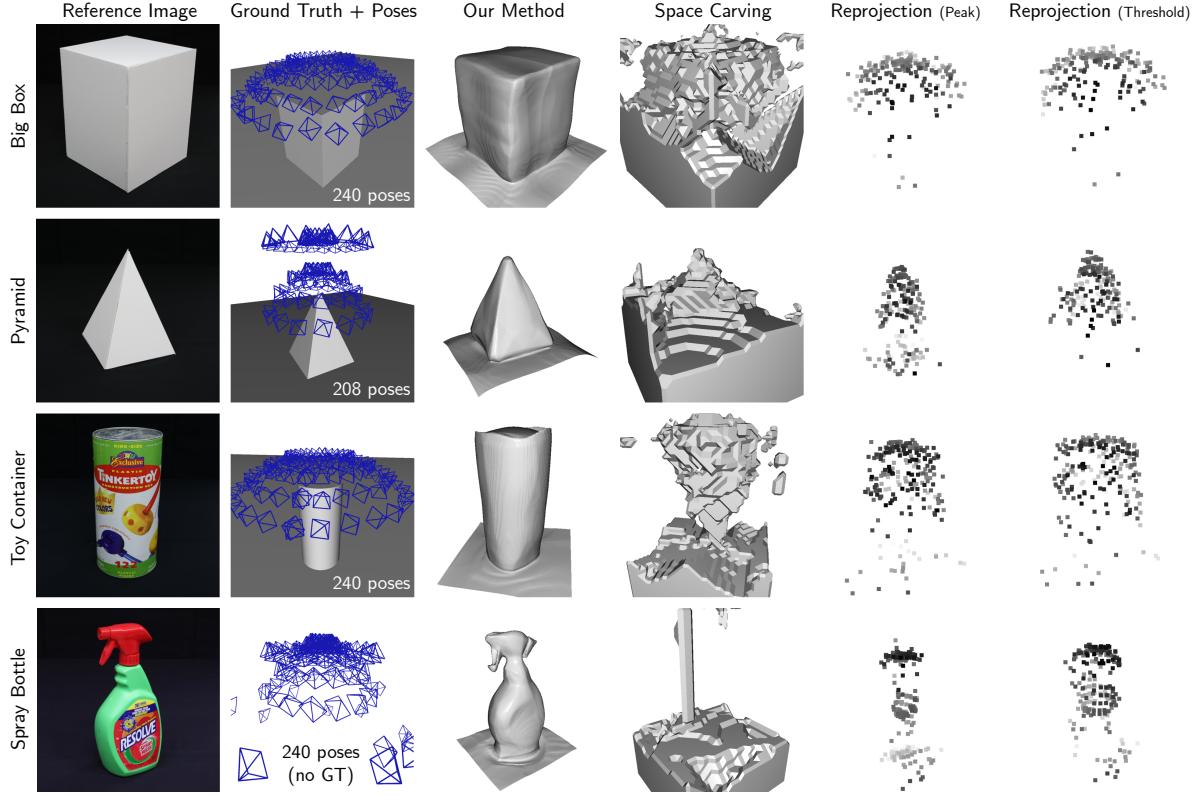


Figure 5. **Qualitative results on real-world captures.** Our method again attains the highest reconstruction quality. Poses in column two are subsampled by a factor of two for clarity. See supplement for additional qualitative results.

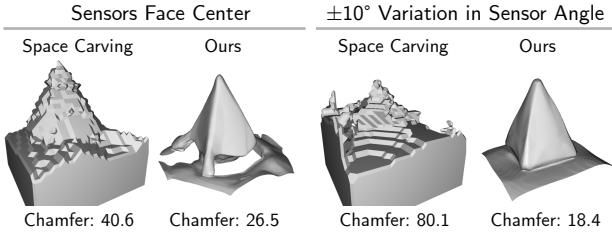


Figure 6. Space carving performs more poorly when sensor poses include some variation in target point, while our system takes advantage of the increased view diversity and coverage.



Figure 7. Because our reconstruction method assumes a Lambertian surface, it fails to reconstruct highly specular scenes, such as a glossy white bust (left) or mirror-finish kettle (right).

Further, our method is surprisingly robust to violation of assumptions made about surface reflectance; it successfully reconstructs non-Lambertian objects with rich texture despite assuming Lambertian BRDF with a spatially uniform albedo. These include both simple shapes (Toy Container in Figure 5 and Cereal Box) and challenging objects with complex geometry (Spray Bottle in Figure 5). We hypothesize

that the overlapping FoVs of distributed sensors help constrain the optimization of our model and encourage a plausible reconstruction that best explains all transients. Our strong results on simulated and real-world data validate our modeling approach and demonstrate a single-photon 3D vision system for real-world scene reconstruction.

5. Discussion and Future Work

Despite assuming spatial uniform reflectance and albedo, our method is robust to rich textures (Fig. 5) and compares favorably to baseline methods for reconstructing challenging scenes with high specularities (Fig. 7). Future work will investigate recovering spatially varying reflectance or incremental learning of geometry as new measurements become available [37, 44], enabling applications like real-time mapping and SLAM. Our method may be particularly relevant in applications such as robotics and wearable computing, where the small size, low power requirements, and robust hardware of proximity sensors are very valuable.

Acknowledgements: This work was supported by Los Alamos National Lab and the Department of Energy, Wisconsin Alumni Research Foundation, NSF CAREER Award 1943149, NSF grant CNS-2107060, and a SONY Faculty Innovation Award. Gleicher and Gupta hold concurrent appointments at Amazon and Cruise, respectively. This work is not associated with either company.

References

- [1] Jens Ackermann and Michael Goesele. A survey of photometric stereo techniques. *Foundations and Trends® in Computer Graphics and Vision*, 2015.
- [2] AMS OSRAM AG. *TMF882X Datasheet*. AMS OSRAM AG.
- [56] 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. *NeurIPS*, 2021.
- [4] Cienna N Becker and Lucas J Koerner. Plastic classification using optical parameter features measured with the tmf8801 direct time-of-flight depth sensor. *Sensors*, 2023.
- [5] Phong Bui-Tuong. Illumination for computer generated pictures. *CACM*, 1975.
- [6] Clara Callenberg, Zheng Shi, Felix Heide, and Matthias B Hullin. Low-cost spad sensing for non-line-of-sight tracking, material classification and depth imaging. *ACM Transactions on Graphics (ToG)*, 2021.
- [7] Ziang Cheng, Hongdong Li, Yuta Asano, Yinqiang Zheng, and Imari Sato. Multi-view 3D reconstruction of a texture-less smooth surface of unknown generic reflectance. In *CVPR*, 2021.
- [8] PB Coates. The correction for photonpile-up’ in the measurement of radiative lifetimes. *Journal of Physics E: Scientific Instruments*, 1968.
- [9] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *CVPR*, 2022.
- [10] Nathan Devrio and Chris Harrison. DiscoBand: Multiview depth-sensing smartwatch strap for hand, body and environment tracking. In *UIST*, pages 1–13, 2022.
- [11] Daniele Faccio, Andreas Velten, and Gordon Wetzstein. Non-line-of-sight imaging. *Nature Reviews Physics*, 2020.
- [12] Yuki Fujimura, Takahiro Kushida, Takuya Funatomi, and Yasuhiro Mukaigawa. NLOS-NeuS: Non-line-of-sight neural implicit surface. In *ICCV*, 2023.
- [13] Jason Geng. Structured-light 3D surface imaging: a tutorial. *Advances in optics and photonics*, 2011.
- [14] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, 2020.
- [15] Anant Gupta, Atul Ingle, Andreas Velten, and Mohit Gupta. Photon-flooded single-photon 3D cameras. In *CVPR*, 2019.
- [16] Miles Hansard, Seungkyu Lee, Ouk Choi, and Radu Patrice Horaud. *Time-of-flight cameras: principles, methods and applications*. Springer Science & Business Media, 2012.
- [17] Felix Heide, Steven Diamond, David B Lindell, and Gordon Wetzstein. Sub-picosecond photon-efficient 3D imaging using single-photon sensors. *Scientific reports*, 2018.
- [18] Carlos Hernandez, George Vogiatzis, and Roberto Cipolla. Multiview photometric stereo. *TPAMI*, 2008.
- [19] Quercus Hernandez, Diego Gutierrez, and Adrian Jarabo. A computational model of a single-photon avalanche diode sensor for transient imaging. *arXiv preprint arXiv:1703.02635*, 2017.
- [20] Fang Hu, Peng He, Songlin Xu, Yin Li, and Cheng Zhang. FingerTrak: Continuous 3d hand pose tracking by deep learning hand silhouettes captured by miniature thermal cameras on wrist. *IMWUT*, 4(2):1–24, 2020.
- [60] Shengyu Huang, Zan Gojcic, Zian Wang, Francis Williams, Yoni Kasten, Sanja Fidler, Konrad Schindler, and Or Litany. Neural LiDAR fields for novel view synthesis. In *CVPR*, 2023.
- [22] Ray A Jarvis. A perspective on range finding techniques for computer vision. *TPAMI*, 1983.
- [23] Sacha Jungerman, Atul Ingle, Yin Li, and Mohit Gupta. 3D scene inference from transient histograms. In *ECCV*, 2022.
- [24] Berk Kaya, Suryansh Kumar, Carlos Oliveira, Vittorio Ferrari, and Luc Van Gool. Multi-view photometric stereo revisited. In *WACV*, 2023.
- [25] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *IJCV*, 2000.
- [26] João L Lagarto, Federica Villa, Simone Tisa, Franco Zappa, Vladislav Shcheslavskiy, Francesco S Pavone, and Riccardo Cicchi. Real-time multispectral fluorescence lifetime imaging using single photon avalanche diode arrays. *Scientific Reports*, 2020.
- [27] Jongho Lee, Atul Ingle, Jenu V Chacko, Kevin W Eliceiri, and Mohit Gupta. CASPI: collaborative photon processing for active single-photon imaging. *Nature Communications*, 2023.
- [28] David B Lindell, Gordon Wetzstein, and Matthew O’Toole. Wave-based non-line-of-sight imaging using fast fk migration. *ACM Transactions on Graphics (ToG)*, 2019.
- [29] Xinyang Liu, Yijin Li, Yanbin Teng, Hujun Bao, Guofeng Zhang, Yinda Zhang, and Zhaopeng Cui. Multi-modal neural radiance field for monocular dense slam with a light-weight tof sensor. In *ICCV*, 2023.
- [30] Anagh Malik, Parsa Mirdehghan, Sotiris Nousias, Kiriakos N. Kutulakos, and David B. Lindell. Transient neural radiance fields for LiDAR view synthesis and 3D reconstruction. *NeurIPS*, 2023.
- [31] ST Microelectronics. *VL53L8CH Datasheet*, 2024.

[32] ST Microelectronics. *VL6180X Proximity and Ambient Light Sensing Module Datasheet*, 2024.

[33] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.

[34] Fangzhou Mu, Sicheng Mo, Jiayong Peng, Xiaochun Liu, Ji Hyun Nam, Siddeshwar Raghavan, Andreas Velten, and Yin Li. Physics to the rescue: Deep non-line-of-sight reconstruction for high-speed imaging. *TPAMI*, 2022.

[35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 2022.

[36] Mark Nishimura, David B. Lindell, Christopher Metzler, and Gordon Wetzstein. Disambiguating monocular depth estimation with a single transient. In *ECCV*, 2020.

[37] Joseph Ortiz, Alexander Clegg, Jing Dong, Edgar Sucar, David Novotny, Michael Zollhoefer, and Mustafa Mukadam. isDF: Real-time neural signed distance fields for robot perception. In *RSS*, 2022.

[38] Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 2018.

[39] Adithya K Pediredla, Aswin C Sankaranarayanan, Mauro Buttafava, Alberto Tosi, and Ashok Veeraraghavan. Signal processing based pile-up compensation for gated single-photon avalanche diodes. *arXiv preprint arXiv:1806.07437*, 2018.

[40] Martin Pollák, Marek Kočíško, Dušan Paulišin, and Petr Baron. Measurement of unidirectional pose accuracy and repeatability of the collaborative robot UR5. *Advances in Mechanical Engineering*, 12(12): 1687814020972893, 2020.

[41] Alice Ruget, Max Tyler, Germán Mora Martín, Sterling Scholes, Feng Zhu, Istvan Gyongy, Brent Hearn, Steve McLaughlin, Abderrahim Halimi, and Jonathan Leach. Pixels2Pose: Super-resolution time-of-flight imaging for 3D pose estimation. *Science Advances*, 2022.

[42] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR*, 2006.

[43] Carter Sifferman, Yiping Wang, Mohit Gupta, and Michael Gleicher. Unlocking the performance of proximity sensors by utilizing transient histograms. *RA-L*, 2023.

[44] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew Davison. iMAP: Implicit mapping and positioning in real-time. In *ICCV*, 2021.

[45] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Nature, 2022.

[46] Chia-Yin Tsai, Kiriakos N Kutulakos, Srinivasa G Narasimhan, and Aswin C Sankaranarayanan. The geometry of first-returning photons for non-line-of-sight imaging. In *CVPR*, 2017.

[47] Eric Veach. *Robust Monte Carlo methods for light transport simulation*. Stanford University, 1998.

[48] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mouni G Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications*, 2012.

[49] 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. *NeurIPS*, 2021.

[50] 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 *CVPR*, 2019.

[51] Liu Xinyang, Li Yijin, Teng Yanbin, Bao Hujun, Zhang Guofeng, Zhang Yinda, and Cui Zhaopeng. Multi-modal neural radiance field for monocular dense slam with a light-weight tof sensor. In *ICCV*, 2023.

[52] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K. Wong. PS-NeRF: Neural inverse rendering for multi-view photometric stereo. In *ECCV*, 2022.

[53] Li Yijin, Liu Xinyang, Dong Wenqi, Zhou han, Bao Hujun, Zhang Guofeng, Zhang Yinda, and Cui Zhaopeng. DELTAR: Depth estimation from a light-weight tof sensor and rgb image. In *ECCV*, 2022.

[54] Franco Zappa, Simone Tisa, Alberto Tosi, and Sergio Cova. Principles and features of single-photon avalanche diode arrays. *Sensors and Actuators A: Physical*, 2007.

[55] Dongxu Zhao, Daniel Lichy, Pierre-Nicolas Perrier, Jan-Michael Frahm, and Soumyadip Sengupta. MVPSNet: Fast generalizable multi-view photometric stereo. In *ICCV*, 2023.

Towards 3D Vision with Low-Cost Single-Photon Cameras

Supplementary Material

In this supplement, we describe (1) The implementation details of our reconstruction pipeline (Section A); (2) a sensitivity analysis to understand the robustness of our method w.r.t. imaging parameters (Section B); (3) A discussion of the difficulty of point cloud surface reconstruction (Section C); (4) additional qualitative results on simulated data and real world captures (Section D); and (5) further discussion of our work (Section E). We also include three short videos, which provide animated 360° views of qualitative reconstruction results. The videos can also be viewed on the project web page: https://cpsiff.github.io/towards_3d_vision/.

For sections, figures and equations, we use numbers (e.g., Sec. 1) to refer to the main paper and capital letters (e.g., Sec. A) to refer to this supplement.

A. Implementation Details

For both simulated and real world experiments, we use an 8-layer MLP with 256 hidden units as our SDF, f_θ , and initialize it as a sphere, centered at the origin with radius 0.3m, using geometric initialization [57]. For each transient, we sample 256 rays ω over Ω and sample 256 points per ray. We set $\lambda_{Eikonal}$ to 0.1 across all experiments and set λ_{TV} to 0 and 0.01 respectively for the simulated and real-world experiments. We train f_θ for 300K steps using Adam [62] with a mini-batch size 2, a learning rate 0.0005, and cosine decay. The learned SDFs are converted to meshes using Marching Cubes [63].

B. Sensitivity Analysis

We perform extensive experiments to understand the robustness of our method in comparison to baselines under varying sensor parameters in simulation. All experiments are based on the Bunny scene and the parameters are varied one at a time while other parameters remain fixed at the base condition (as described in Section 4 of the main paper). To ensure strong baselines for every sensor configuration, we calibrate the thresholds t_p and t_s for the projection (threshold) and space carving baselines respectively per reconstruction. We perform a brute force search over possible thresholds and report the best Chamfer achieved. As this amounts to calibrating on the test set, the numbers reported represent the best possible performance of the baseline methods on the given data. The results of this sensitivity analysis are presented in Figure A. In what follows we discuss some of the main findings.

Sensor Placement. We study two key parameters that control sensor placement: the number of views and the mini-

mum elevation angle at which the sensors are placed. Our method consistently outperforms all baselines in Chamfer distance by an order of magnitude across a broad range of parameter choices. In particular, our method readily supports as few as 128 views above a considerably large elevation angle of 30° without harming reconstruction quality. This robust gain in performance confirms that our method takes advantage of broad-band signal in transients not exploited by the baseline methods.

Temporal Resolution. Our system takes advantage of the temporal information in transient histograms, and therefore benefits when that information is present at a high resolution. Because of this, our method outperforms baselines by a very wide margin at a small bin size, but the margin vanishes as bins become wider than 2cm (equivalently 66ps), because decomposing the temporal signal becomes impractical beyond this limit. Fortunately, today’s commodity SPADs operate at a smaller bin size (~ 40 ps). Baseline methods show no performance gain at small bin sizes, as they do not take advantage of the temporal resolution.

Angular Resolution. Our system resolves spatial resolution from wide-FoV sensors by taking advantage of the time dimension. In this regime, the optimal sensor field-of-view size is not obvious: a smaller FoV means more highly constrained geometry, as each histogram images a smaller region, but too small of a field-of-view means a lack of coverage and under-constrained geometry. We find that an angular resolution in the 30° to 60° range is optimal for reconstructing 3D geometry with our method on the bunny scene. Reprojection based methods benefit more from a smaller field-of-view, while space carving performs best with a wider field-of-view so that space is sufficiently carved away. In every case, our method outperforms baselines by a wide margin.

Signal-to-noise ratio (SNR). We consider three parameters that jointly impact SNR: illumination power, ambient flux, and number of illumination cycles. Our method again outperforms all baselines by a significant margin across all test conditions. Notably, the baselines fail or perform considerably worse under high ambient flux, as signal photons are blocked by background photons due to pile-up. By contrast, our method is robust against a broad range of ambient flux levels, as we model the effects of ambient flux directly.

C. Point Cloud Surface Reconstruction

We do not apply surface reconstruction to the point clouds computed by reprojection because off-the-shelf reconstruc-

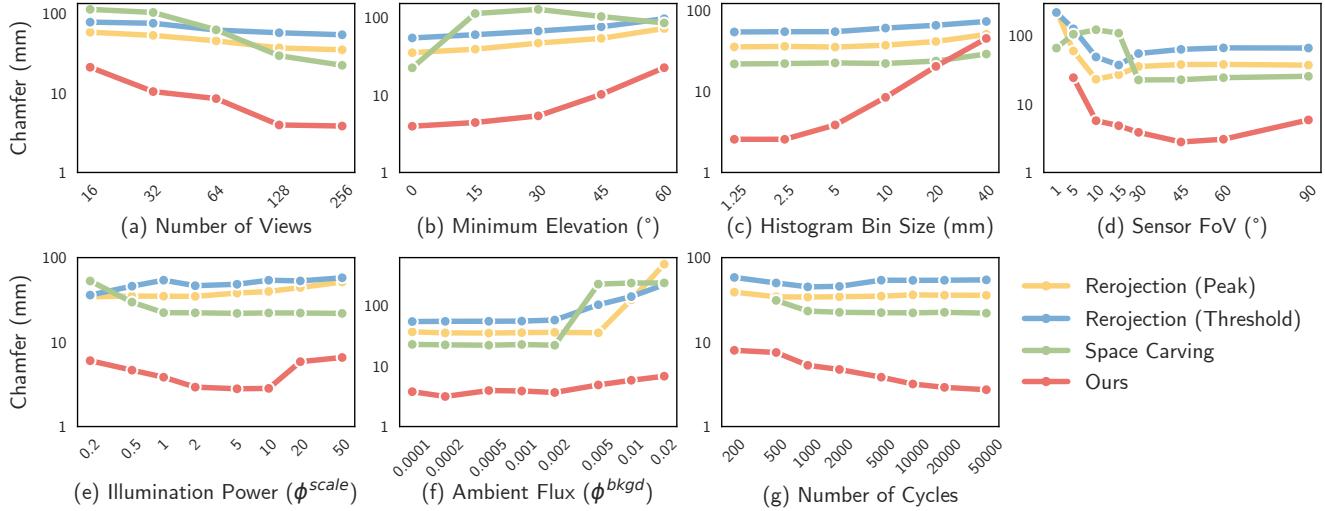


Figure A. **Sensitivity analysis of our method compared to baselines across a range of imaging parameters.** In almost every case, our method outperforms baseline methods on Chamfer distance. Missing datapoint in (d) indicates that our method failed to converge. Illumination power (e) is unit-less as it also absorbs factors like quantum efficiency and does not map directly to any real world parameter.

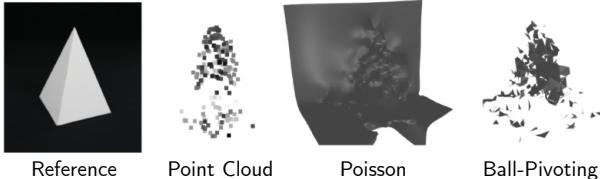


Figure B. Off-the-shelf algorithms for surface reconstruction (Poisson Surface Reconstruction [61] and The Ball-Pivoting Algorithm [58]) do not perform well on point clouds generated by reprojection.

tion techniques do not perform reliably on the generated point clouds, as shown in Figure B. Additionally, calculating Chamfer distance to the computed point cloud ensures that we are capturing the efficacy of reprojection rather than a given surface reconstruction method. We include one-way reconstruction to ground truth Chamfer distance in Table 2 to provide a metric which does not penalize the sparsity of point clouds produced by reprojection.

D. Additional Qualitative Results

More reconstruction results. We present additional qualitative results on simulated data (Figure D) and real-world captures (Figure C). These results were omitted from the main paper due to lack of space.

Surface normal visualization. Moving beyond the 3D shapes, we further examine the surface normal of our reconstructed 3D objects. The surface normal of a point \mathbf{x} on the reconstructed mesh is estimated as

$$\tilde{\mathbf{n}}_{\mathbf{x}} = \frac{\nabla_{\mathbf{x}}(f_{\theta}(\text{PE}(\mathbf{x})))}{\|\nabla_{\mathbf{x}}(f_{\theta}(\text{PE}(\mathbf{x})))\|}, \quad (\text{A})$$

where f_{θ} is the learned SDF and PE denotes the positional encoding function. The error $e_{\mathbf{x}}$ w.r.t. the ground-truth normal $\mathbf{n}_{\mathbf{x}}$ is given by

$$e_{\mathbf{x}} = |\langle \mathbf{n}_{\mathbf{x}}, \tilde{\mathbf{n}}_{\mathbf{x}} \rangle|. \quad (\text{B})$$

Surface normal results. We provide visualizations of surface normals for simulated data in Figure E. Our method can successfully recover smoothly varying normals. Error typically occurs at edges and depth discontinuities with fast-changing normals. We hypothesize that sensors with higher temporal and spatial resolution are needed for more accurate surface normal reconstruction.

E. Further Discussion

Beyond Lambertian objects. Our method assumes a spatially uniform Lambertian BRDF, but in practice can effectively reconstruct objects with spatially varying albedo and slightly glossy appearance (e.g. the spray bottle). In theory, our method can easily be adapted to incorporate a parametric lighting model. Recovery of the parameters of such a model are likely possible because, by sharing information among many observations, the BRDF is effectively sampled at many incident and exitant angles. An intriguing direction for future work is investigating which BRDF parameterizations can be recovered with our imaging setup, and the effect of the reflectance model on reconstruction quality. We suspect that a non-parametric NeRF-like BRDF would not be suitable as it does not sufficiently constrain the optimization. A parametric lighting model, e.g. Phong [59] or Oren-Nayar [66] may appropriately constrain the opti-

mization while allowing the model to learn a more accurate scene representation.

Runtime efficiency. Our method takes on the order of hours to reconstruct a scene, making it unsuitable for real-time applications in its current state. Future work should investigate ways to speed up forward rendering and model training. Improved importance sampling would likely yield modest improvements in convergence time. Another option is to render only summary statistics of the histogram (*e.g.* mean, peak locations or widths) rather than the entire histogram, which would likely be faster to render at the expense of yielding a lower-quality reconstruction.

Sensor pose. In this work, we used an industrial robot arm to gather posed sensor measurements. We chose this modality as it is guaranteed to provide highly accurate sensor poses, and allows control over precise sensor placement. For applications like wearable computing and distributed sensing for robotics, camera poses might be pre-calibrated and remain fixed relative to each other during operation. Alternatively, the low-cost single-photon camera could be combined with a sensor-based localization system (*e.g.*, an IMU based [68] or a camera based [65] system) to recover camera pose, a setup which is standard in related works [64, 67]. Such a capture setup would allow capture of more organic and large scale scenes, which more closely mimic the potential use cases of the sensor (*e.g.* on mobile robots and drones).

Comparison to other 3D imaging modalities. Our work provides a low-cost 3D imaging system using single-photon cameras. We provide detailed comparisons between our method and baseline methods, but do not compare our reconstructions to those gathered from other 3D modalities, such as continuous wave time-of-flight [56] or LiDAR [60]. Future work should provide a comparison to these other modalities to provide insights into the niche (in terms of accuracy, size, power, etc.) filled by each.

Comparison to a NeRF with a miniature RGB camera. Using multiple calibrated RGB cameras, NeRF-like approaches can leverage photometric cues (*e.g.* correspondence and shading) to recover 3D shapes from a distributed set of cameras. *Conceptually*, NeRF-like approaches fall short under non-ideal lighting (*e.g.* low-light) or insufficient correspondence (*e.g.* textureless objects), while our method with active lighting and depth remains effective. With sufficient light and distinct texture patterns, NeRF-like approaches on RGB images will yield a higher quality reconstruction than our method, due to significantly increased data rates. Without compression, a VGA-resolution image with 8-bit color channels contains 7Mbits of information, while an image (*i.e.*, transient histogram) from our sensor contains 3Kbits of information. With $2400 \times$ the information per-view, it is not surprising that NeRF with RGB cam-

eras could outperform our method given the same number of input views. *Practically*, existing low-cost RGB cameras are larger and less power efficient than the SPAD sensors that we utilize. While miniature RGB cameras do exist (*e.g.*, those used for endoscopy), they are $>20 \times$ the cost of a SPAD.

Commodity sensors. One challenge for future work is a lack of hardware support for measurement and use of transient histograms. Very few low-cost sensors allow access to transient histograms, and those that do often perform pre-processing that is proprietary or undocumented. Additionally, most sensors are equipped with very low-bandwidth I²C interfaces, limiting their effective FPS. We hope that manufacturers will see value in users having access to transient histogram data and support the use of this data with documentation, low-level access, and high-bandwidth interfaces in the future.

Ethical concerns. Our work presents a new method for imaging 3D objects with low-cost single-photon cameras. We do not anticipate major ethical concerns.

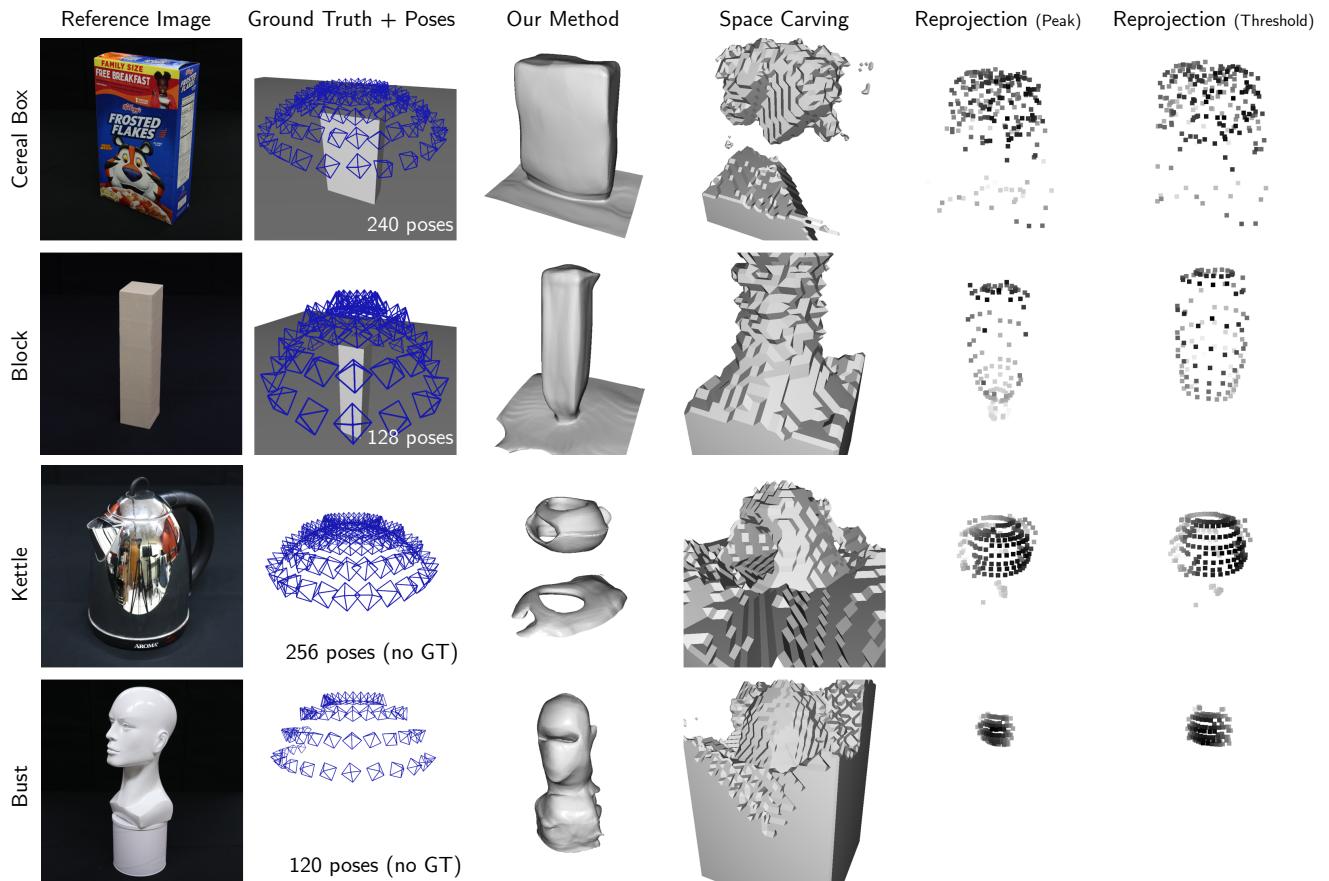


Figure C. **Additional qualitative results on real-world captures.** Our method achieves the highest reconstruction quality. Poses in column two are subsampled by a factor of two for clarity.

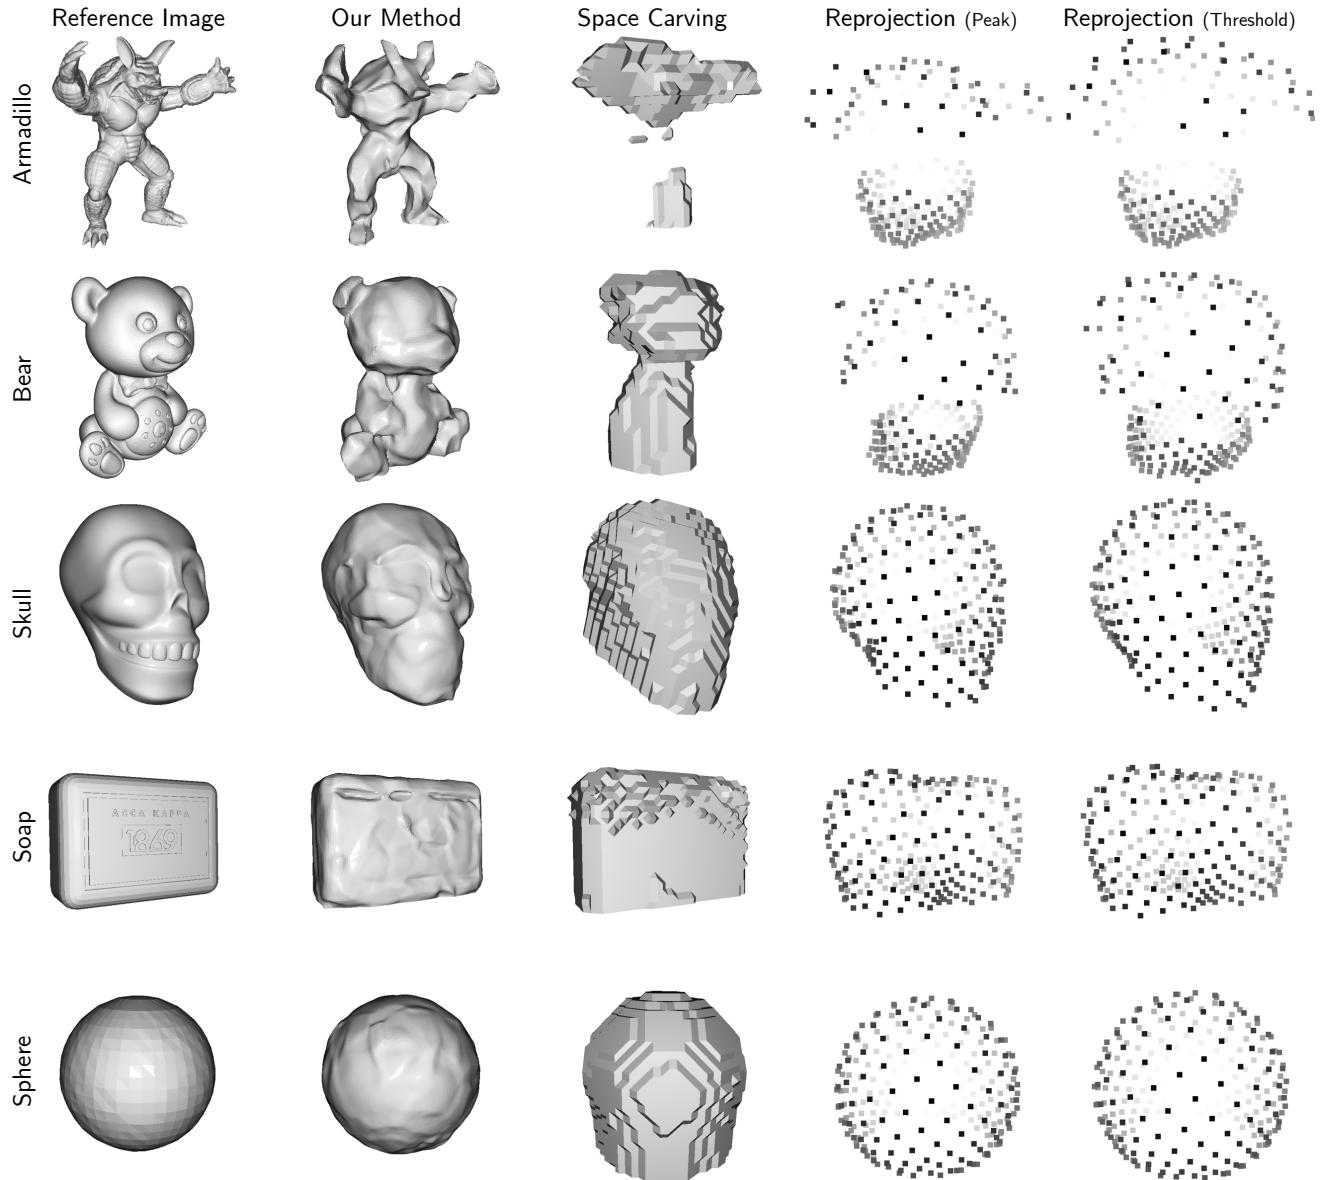


Figure D. Additional qualitative results on simulated data. Our method achieves the highest reconstruction quality. Space carving captures an envelope of the shape, and may carve away occupied areas in concave shapes (e.g. Armadillo). Reprojection gives a sparse reconstruction of convex shapes, (e.g. skull, soap, sphere), the scale of which may be distorted due to biases introduced by the wide field-of-view of the sensor.

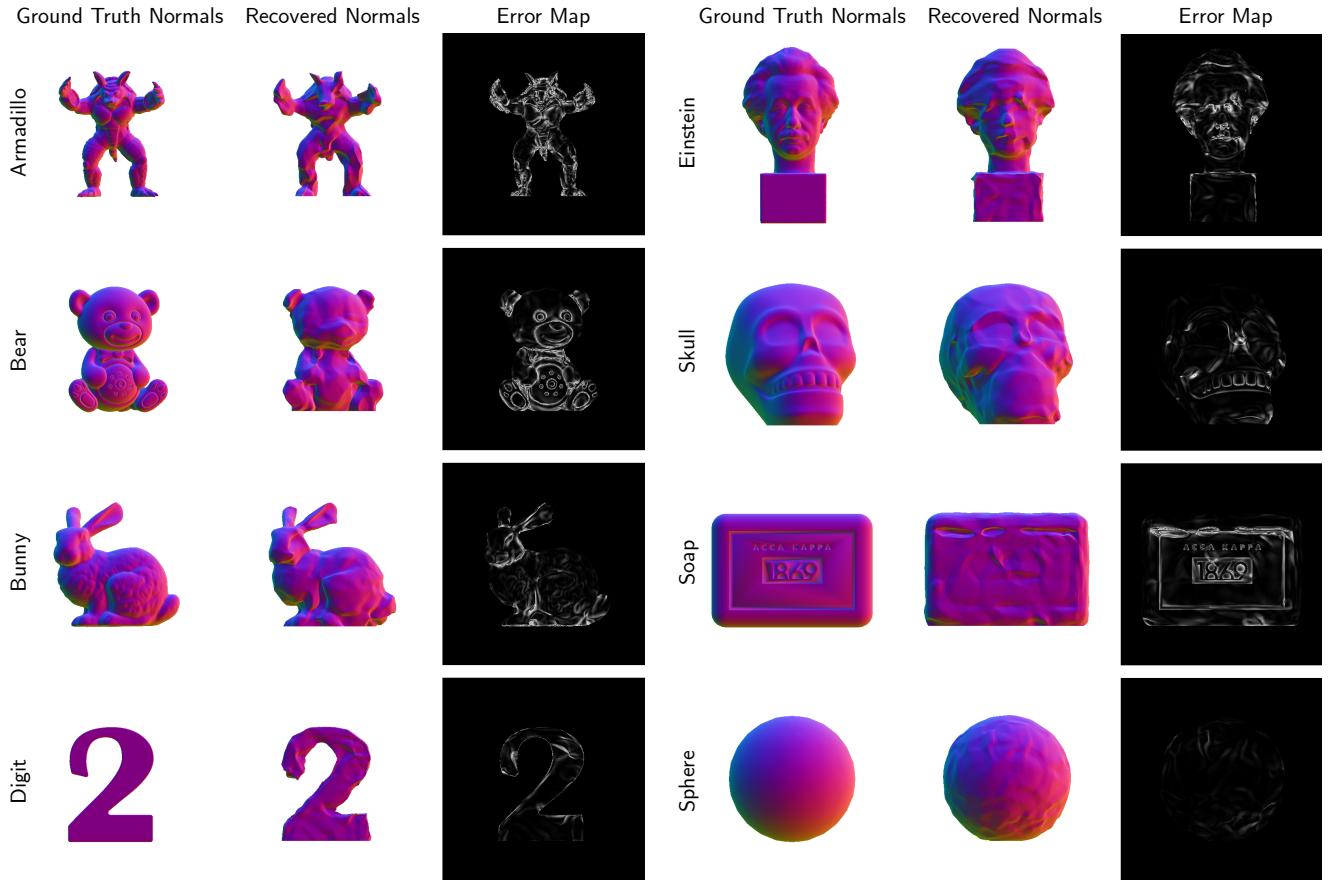


Figure E. **Visualizations of surface normals for simulated data.** Our method correctly estimates surface normals in flat regions. Error mainly occurs at edges and depth discontinuities. We hypothesize that sensors with higher temporal and spatial resolution are needed to detect rapid changes in surface normals.

Supplementary References

- [56] 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. *NeurIPS*, 2021.
- [57] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *CVPR*, 2020.
- [58] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, 1999.
- [59] Phong Bui-Tuong. Illumination for computer generated pictures. *CACM*, 1975.
- [60] Shengyu Huang, Zan Gojcic, Zian Wang, Francis Williams, Yoni Kasten, Sanja Fidler, Konrad Schindler, and Or Litany. Neural LiDAR fields for novel view synthesis. In *CVPR*, 2023.
- [61] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, 2006.
- [62] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [63] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, 1987.
- [64] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [65] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 2015.
- [66] Michael Oren and Shree K Nayar. Generalization of lambert’s reflectance model. In *SIGGRAPH*, 1994.
- [67] Joseph Ortiz, Alexander Clegg, Jing Dong, Edgar Sucar, David Novotny, Michael Zollhoefer, and Mustafa Mukadam. isdf: Real-time neural signed distance fields for robot perception. In *RSS*, 2022.
- [68] Jingang Yi, Junjie Zhang, Dezhen Song, and Suhada Jayasuriya. Imu-based localization and slip estimation for skid-steered mobile robots. In *IROS*, 2007.