

Adaptive Lighting for Data-Driven Non-Line-of-Sight 3D Localization and Object Identification

Sreenithy Chandran
schand56@asu.edu

Arizona State University
Tempe, USA

Suren Jayasuriya
sjayasur@asu.edu

Abstract

Non-line-of-sight (NLOS) imaging of objects not visible to either the camera or illumination source is a challenging task with vital applications including surveillance and robotics. Recent NLOS reconstruction advances have been achieved using time-resolved measurements which requires expensive and specialized detectors and laser sources. In contrast, we propose a data-driven approach for NLOS 3D localization and object identification requiring only a conventional camera and projector. To generalize to complex line-of-sight (LOS) scenes with non-planar surfaces and occlusions, we introduce an adaptive lighting algorithm. This algorithm, based on radiosity, identifies and illuminates scene patches in the LOS which most contribute to the NLOS light paths, and can factor in system power constraints. We achieve an average identification of 87.1% object identification for four classes of objects, and average localization of the NLOS object's centroid with a mean-squared error (MSE) of 1.97 cm in the occluded region for real data taken from a hardware prototype. These results demonstrate the advantage of combining the physics of light transport with active illumination for data-driven NLOS imaging.

1 Introduction

Non-line-of-sight (NLOS) imaging is an emerging field of research with applications for autonomous vehicle collision avoidance, search and rescue operations, and industrial safety and inspection. This problem is challenging because NLOS objects are outside of the line-of-sight (LOS) of both the camera and the illumination source(s). Recent success in NLOS reconstruction utilizes time-of-flight to help separate NLOS from LOS path contributions [1, 2, 3]. Using time-resolved measurements, backpropagation [4] and/or optimization [5] can reconstruct the NLOS scene, typically under the assumption of a flat LOS wall to avoid LOS indirect light. These time-resolved detectors, while achieving superior results, are costly, consume higher power than conventional CMOS image sensors, and require specialized aligned optics and lasers. To alleviate these issues, research has also focused on NLOS imaging using conventional cameras [6, 7, 8, 9, 10].

We propose a data-driven approach to the problem of NLOS 3D localization as well as object identification that leverages modern existing convolutional neural network architectures. This approach requires only a conventional 2D camera and a projector, and no optical

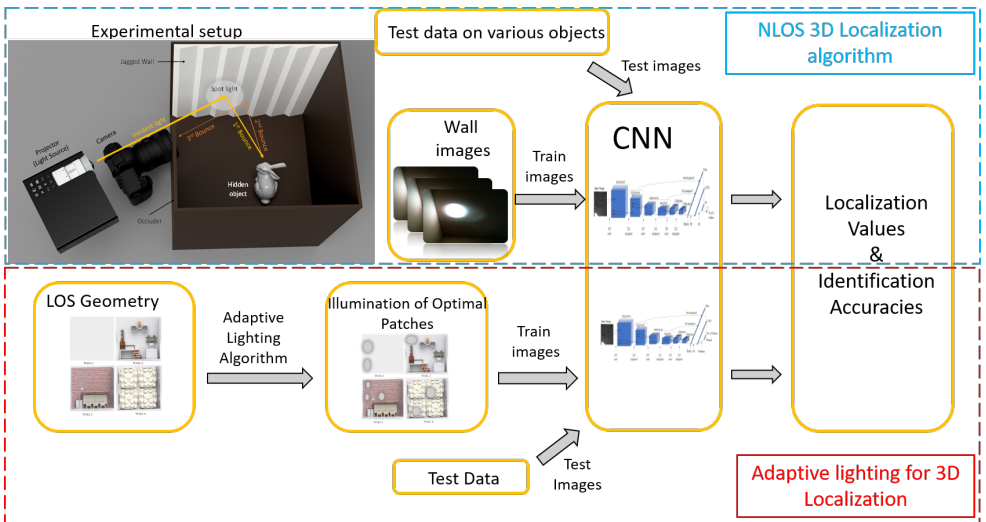


Figure 1: Our research uses deep learning to perform 3D localization of an object in the non-line-of-sight (NLOS) of a camera and projector. Our complete system pipeline consists of a CNN trained on diffuse wall images, and the adaptive lighting procedure to improve performance. We use two different network architectures for localization and identification.

alignment or projector-camera synchronization. To handle LOS scenes with non-planar geometries and occlusions, we develop an adaptive lighting algorithm based on radiosity that improves our learning performance. To the best of our knowledge, we are the first paper to perform NLOS imaging on non-planar, complex LOS scenes.

In particular, our specific contributions are (1) solving NLOS 3D localization as well as NLOS object identification using convolutional neural networks with a conventional camera and projector, (2) an adaptive lighting algorithm for identifying optimal LOS patches to maximize NLOS radiosity captured, and distribute illumination power among these patches given power constraints, and (3) robust performance for non-planar complex LOS geometries and improved NLOS localization and identification using our adaptive lighting algorithm. To validate these contributions, we demonstrate results using a physically-based renderer as well as a real working experimental setup in the lab. We publicly release our code, trained network models, and datasets at https://github.com/sreenithy/AdaptiveLighting_NLOS for use by the research community. This paper combines the advantages of physics-based vision with deep learning, and shows how the choice of illumination can play a major role in the performance of NLOS imaging algorithms.

2 Related Work

NLOS Imaging with Time-Resolved Detectors: Kirmani et al. introduced “looking around corners” using 5D time-resolved light transport [23]. This was later realized using streak cameras and a femto-second laser to perform elliptic backpropagation [15, 24]. Researchers demonstrated NLOS reconstruction using continuous-wave time-of-flight sensors via optimization [16] and NLOS tracking using array signal processing [20]. Single-photon avalanche diodes (SPADs) [11] have been used for NLOS reconstruction [9, 30]. Advances in reconstruction algorithms for SPADs include accounting for partial occluders and surface

normals [17], confocal NLOS reconstruction [10], geometric modeling [32], backpropagation [2, 27], space-carving [41], temporal focusing [60], wave-based models in the frequency domain [28], and the identification of Fermat paths [46]. Higher level applications such as detection and tracking have been shown [7, 12]. Recently, neural networks have been applied to time-resolved measurements for NLOS imaging [6, 33].

NLOS Imaging without Time-Resolved Detectors: Sen et al. transposed the light transport matrix to see an object in the camera’s NLOS, although it is still in the projector’s LOS [54]. Accidental pinholes and pinspeck cameras can visualize the NLOS scene [40]. Bouman et al. visualize 1D and 2D slices of the NLOS scene using a physical wall/corner [9]. Known occluders in the NLOS can help reconstruct the light field [9] and hidden scene [39]. Smith et al. used laser speckle for tracking NLOS objects [36]. Similar to this paper, Klein et al. demonstrate a real-time tracking algorithm using radiosity [24], Tancik et al. perform deep learning for conventional cameras for NLOS imaging [37, 38], and Chen et al. recently show high quality reconstruction from RGB cameras using a convolutional neural network [8]. However, our work differs from the previous works by introducing the idea of adaptive lighting to improve the performance of data-driven NLOS localization, and show this for non-planar LOS scenes.

3 NLOS Localization and Identification with Deep Learning

Our NLOS imaging scenario consists of a projector illuminating a spot on a LOS scene (not necessarily planar), and a camera captures an image of this spot. An illuminated spot on the wall undergoes at least three diffuse reflections or bounces as it travels from a point on the wall to hidden NLOS object and back to the wall, before being captured by the focused camera. This setup is present in the top left corner of Fig. 1. The tasks we are interested in include localization, i.e. determining the (x, y, z) location of an object’s centroid, and identification of that object is present in the NLOS. In this paper, we assume a fixed surveillance scenario where information about the line-of-sight (LOS) is known, and images can be captured for training of neural networks. While these assumptions limit the practical application of this research (i.e. does not generalize to unseen locations or moving scenes), it allowed us tractability in which to investigate questions about localization, identification, and adaptive lighting.

3.1 Dataset Rendering

For deep learning, large datasets of images are typically needed to train the networks. However, collecting such large numbers of image data for a real NLOS scene would be prohibitive. Thus we utilize synthetic data to train our algorithms, and then fine-tune the networks on the real experimental dataset which is smaller in scale. We use dataset augmentation including flipping, cropping, translation, and rotation, to help the network be robust to changes in camera viewpoint and improve generalisation performance. For synthetic data generation, we use the physically-based renderer Mitsuba [18]. We render scenes similar to those shown in Figure 3(A), that include complex geometries, occlusions, colors and textures. In our experimental results, we show how our adaptive lighting method robustly improves performance on all these scenes.

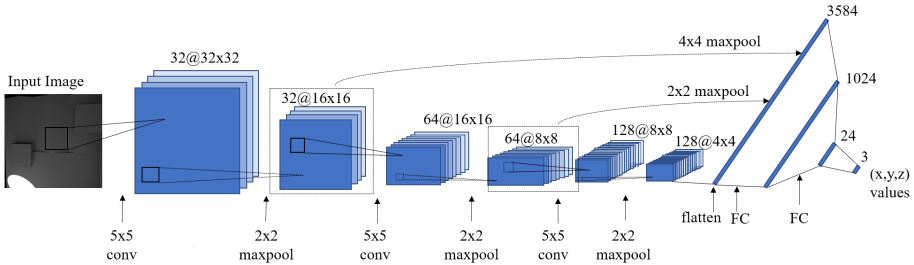


Figure 2: Proposed network architecture for NLOS localization which uses convolutional layers followed by fully connected layers to perform the end task.

We used four types of NLOS objects: a sphere (diameter 5cm), a human model (width 5.5 cm and height 17.5 cm), a cylinder (diameter 6 cm and height 8 cm), and the Stanford bunny (width 7.4 cm and height 5 cm). The illuminated spot was randomly positioned on the LOS wall. For the adaptive lightings described later in Section 4, we position the spot(s) to illuminate the patch(es) given by the optimization algorithm.

To quickly model image formation, we use the Instant Radiosity algorithm [27] available in Mitsuba. This does not generate physically-realistic images as compared to using Monte Carlo integration for solving the rendering equation [21, 43], but it enables extremely fast rendering speeds. Rendering one image of 64×64 resolution with 20,000 samples per pixel (spp) takes 18 seconds on a GeForce GTX1080 Ti GPU. This enabled dataset creation up to 100,000 images suitable for deep learning in a short amount of time, while rendering a full Monte Carlo path tracer at 20,000 spp takes 22 minutes per image. Previous work has shown that physically-realistic rendering using Monte Carlo compared to radiosity does lead to improved performance for computer vision tasks [44]. However, this paper localizes an object of width 5.5 cm with an MSE of 1.55 cm using radiosity images. We leave it to future work to show NLOS scenes where physically-realistic rendering is critical for learning performance.

3.2 Network Architectures

Our network architectures are similar for localization and identification with a few changes to the last layers. The network architecture consists of five layers: three convolutional layers for feature extraction and either two fully connected layers for localization, or a fully connected layer and soft-max classifier for object identification. We use multi-scale features via skip connections similar to [35]. Convolution layers are of size 5×5 with a stride of 1. Pooling layers are of size 2×2 . A ReLU is performed after each pooling layer and before each fully connected layer. The localization network has three outputs corresponding to the predicted (x, y, z) centroid location. The identification network uses softmax for classification into object classes. Each network is trained with images with multiple NLOS objects but only one type of LOS scene/wall. The network architecture for localization is in Fig 2, a similar network architecture for classification is in the supplementary material. Note that we do not perform background subtraction or calibration with knowledge that there is no NLOS object in the scene for our images at testing/inference as compared to other works [24, 57]. For implementation details including training times, epochs, learning rate, please see the supplemental material.

4 Adaptive Lighting for Improved NLOS Localization and Identification

A central question we explore in this paper is what are the best lighting patterns in the LOS which maximize NLOS localization and identification? In other words, given a LOS scene with N patches, what patches maximize photons that travel along NLOS light paths? To approach this problem, we use physics-based constraints on light transport given by radiosity [10, 13, 22]. Radiosity describes the transfer of radiant energy between surfaces in the scene, and is calculated based on surface geometry and reflectances/emittances. We make the following modeling assumptions: (1) we have knowledge of the LOS geometry (i.e. via structured light or 3D scanning), (2) diffuse reflection and emittance only in our scene, and (3) only modeling up to third bounce light similar to assumptions made by Klein et al. [24]. We thus formulate an optimization to determine lighting that will maximize signal from the NLOS, and can handle complex LOS scenes with occlusions and non-planar geometries. This is a first step to having NLOS imaging work for complex LOS scenes in the real world.

The total radiosity leaving a single patch will be its emission (i.e. if its a light source) along with the summation of all reflected radiosity from other patches [13]:

$$B_i = E_i + \rho_i \sum_{j=1}^N F_{ij} B_j, \quad F_{ij} = \left(\frac{\cos\theta_i \cos\theta_j}{\pi r^2} \right) V(x_i, x_j) A_i \quad (1)$$

where B_i is the outgoing radiosity of patch i , E_i is the energy emitted out of the patch i , ρ_i is the diffuse reflectance of the patch i , F_{ij} is the form factor which determines how much light is transferred from one patch to another, and B_j is the outgoing radiosity of patch j . F_{ij} depends on x_i as the central point on patch i , x_j as a central point on patch j , $r = \|x_i - x_j\|$ as the distance between x_i and x_j , θ_i as the angle between normal N_i and vector $x_j - x_i$, θ_j as the angle between normal N_j and vector $x_i - x_j$, and $V(x_i, x_j)$ as the visibility function of both patches. Note that this is an approximation of radiosity, and that radiosity form factors can be better calculated using hemicubes [9]. However, this is a common approximation used in practice for NLOS imaging [17, 24].

This visibility function accounts for occlusion, and is 1 if patch i is visible from patch j , and 0 otherwise. We utilize basic raycasting to determine if two patches are visible to one another, which allows us to handle complex LOS scenes with occlusion. However, we are not robust to occlusion in the NLOS as this is difficult to determine from only LOS information.

In the LOS scenes we consider, we include objects that have varying colors and material reflectance that include mixed diffuse (k_d) and specular (k_s) terms. However, for our algorithm, we only consider the k_d term to formulate a radiosity optimization. Of course accounting for specularly may improve the performance of our adaptive lighting algorithm, but we found that we achieved significant improvement using diffuse modeling assumptions alone.

4.1 Maximizing NLOS Radiosity

We now present our adaptive lighting optimization algorithm based on radiosity. Please refer the supplemental material for a more complete derivation for the light paths that are chosen for optimizing over. Using the above radiosity equation, we can calculate the radiosity for LOS light paths and NLOS light paths for three bounce light. If we seek to illuminate m

patches out of N total LOS patches, we wish to maximize the following:

$$\operatorname{argmax}_{[1,\dots,m] \subseteq [1,\dots,N]} \sum_{i=1}^m B_i^{NLOS} \quad (2)$$

where B_i^{NLOS} is the NLOS radiosity from patch i . We can determine B_i^{NLOS} by summing all three-bounce light paths that hit both the patch i and the NLOS voxel (see supplemental material for full derivation and algorithm). We return the top $1, \dots, m$ patches which maximize the NLOS radiosity¹. In the supplemental material, we discuss why adaptive lighting focuses on patches as opposed to spatially-varying lighting, which we show is not optimal because it spreads out power over a large region rather than concentrating photons on a few key patches.

Distributed Light Algorithm: The real advantage of adaptive lighting is when the imaging system has tight power constraints for its illumination. This is common for embedded and mobile computer vision platforms such as robotics and unmanned aerial vehicles. New energy-efficient active illumination also exploits illuminating and capturing only salient light paths to extract signals of interest [0, 26, 29, 42, 45]. In this vein, we develop a *distributive* algorithm which tells us not only what are the optimal patches to shine light, but what fraction of power from our power budget should we shine on each patch.

We formulate a joint optimization over the m optimal patches and the illumination radiant exitances $I_i, i = 1, \dots, M$ for each patch. This is given by the following equation:

$$\operatorname{argmax}_{\substack{[1,\dots,m] \subseteq [1,\dots,N] \\ I_1, \dots, I_m}} \sum_{i=1}^m B_i^{NLOS}(I_i) \text{ such that } \sum_{i=1}^m I_i = T, \forall I_i \leq I_o. \quad (3)$$

Here, $B_i^{NLOS}(I_i)$ is the radiosity for patch i given illumination radiant exitance I_i , T is the total illumination power constraint, and each radiant exitance cannot exceed a maximum I_o .

We use the Sequential Least Squares Programming algorithm available in SciPy library [49] to solve our optimization problems. The optimization treats the indices as continuous variables to utilize conventional solvers (as opposed to combinatorial optimization), and we find the solutions converged to valid integer values nevertheless. The full optimization for adaptive lighting takes between 1-2 minutes to run on a conventional CPU.

Training and Inference using Adaptive Lighting: We describe the training and inference procedure for our adaptive lighting algorithm. This is summarized in Figure 1. To determine adaptive lighting patterns for the NLOS, we divide the NLOS region into voxels in order to calculate the optimization that maximizes radiosity back from this voxel region. Then for our dataset, we train the network with an object contained in NLOS voxel i with the associated adaptive lighting given by our algorithm for the same voxel. We are thus training the network to recognize these adaptive lightings per voxel class to maximally extract NLOS information. At inference, we present N adaptive lightings, each corresponding to an associated possible N NLOS voxels where the object could be located. We modify the final layer of the network to have a softmax output to give probabilities per voxel. We run all these N images through forward propagation of the network and save this softmax output of each image. We then take the maximum confident probability across all N images as the location of the NLOS object.

¹Here, the indices of these patches are relabeled to be $1, \dots, m$ for notation. The optimization returns m indices corresponding to the m patches that is a random subset of the N total patches.

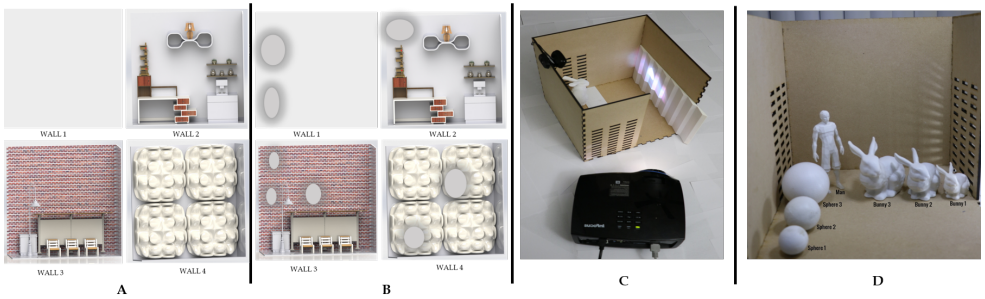


Figure 3: (A) The various LOS scenes used for experiments; (B) The LOS scenes with optimal adaptive lighting patterns visualized; (C) Real experimental prototype in the lab; (D) The 3D printed objects used in real experiments.

5 Experimental Results

Simulated Results: In Table 1, we display the results for simulated data for four types of LOS walls for the localization and identification of the 4 objects. For each task, we perform adaptive lighting to select the optimal patch ($m = 1$) to illuminate, and the 2nd and 3rd choice correspond to the subsequent patches returned by the adaptive lighting optimization. As can be determined from the table, the adaptive optimal patch had an average mean-squared error (MSE) of 2.09 cm, while the second and third chosen patches had an average MSE of 3.60 cm and 5.08 cm respectively. Thus adaptive lighting improved localization on average by 1.51 cm across all experiments. In addition, we tested object identification in the NLOS. This was a four class object identification with an additional fifth class of “No object” if the softmax score of the classifier was below 50%. Average identification of the NLOS object calculated across the four walls is 91.02% for the adaptive lighting, while the next best illuminated patch averaged 82.28%.

Network Generalisation: Our network generalizes across changes in camera viewpoint at testing, since we performed data augmentation during training of this nature. We tested rotating a camera a full 360 degrees in increments of 45 degrees at test time (for a NLOS sphere using LOS Wall 2). This incurred localization errors that varied up to 0.75 cm. Similarly, for translation of the camera position along the x,y,z axis, we only saw variation up to 0.4 cm. Thus our dataset augmentation yielded generalisation to camera motion at test time.

However, our network does not generalize as well to NLOS objects its never seen before. Using new NLOS object of the Stanford dragon (length of 6 cm, width of 3.5 cm and height of 5 cm) resulted in 8.02 cm localization error, and the Happy Buddha statue (width of 5 cm and height of 15 cm) resulted in 7.18 cm localization error. It remains future research to make the networks more robust to unknown objects or simply to add enough data for the network to generalize.

Distributed Lighting: Also in Table 1, we present the distributed lighting algorithm results that jointly optimizes over patch selection and illumination power (Equation 3). To compare, we present two baseline methods: Method 1 only optimizes patch selection using the adaptive lighting algorithm and distributes power evenly among these patches, while Method 2 neither optimizes the distribution of power nor the selection of patches. Our simulated experiment is for two objects (Man and Bunny) with LOS Wall 3, for two ($m = 2$) illumination patches. Note how we get improvements of 40% (3.25 cm \rightarrow 2.17 cm average) improvement of jointly optimizing over Method 1, and as high as 65% (5.22 cm \rightarrow 2.17 cm average) improvement over Method 2.

Wall	Object	Localization with Adaptive Lighting			Identification with Adaptive Lighting	
		Optimal Patch	2nd Patch	3rd Patch	Optimal Patch	2nd Patch
1	SPHERE	2.19 cm	3.26 cm	4.89 cm	92.6%	81.7%
1	BUNNY	1.89 cm	2.54 cm	5.93 cm	92.8%	82.3%
1	MAN	1.50 cm	2.30 cm	3.91 cm	93.3%	82.1%
1	CYLINDER	1.52 cm	2.94 cm	3.75 cm	93.6%	84.2%
1	NO OBJECT	-	-	-	98.5%	96.3%
2	SPHERE	2.34 cm	3.81 cm	5.23 cm	89.9%	79.1%
2	BUNNY	2.31 cm	3.37 cm	4.91 cm	90.3%	84.9%
2	MAN	1.56 cm	3.09 cm	4.87 cm	91.4%	81.3%
2	CYLINDER	1.98 cm	3.92 cm	4.11 cm	91.0%	83.5%
2	NO OBJECT	-	-	-	98.1%	94.6%
3	SPHERE	2.90 cm	4.9 cm	6.15 cm	86.4%	76.1%
3	BUNNY	2.81 cm	4.17 cm	5.85 cm	87.8%	78.3%
3	MAN	1.92 cm	3.98 cm	5.02 cm	88.1%	79.2%
3	CYLINDER	2.30 cm	3.26 cm	5.68 cm	89.3%	80.6%
3	NO OBJECT	-	-	-	97.6%	93.0%
4	SPHERE	2.98 cm	4.51 cm	5.81 cm	84.6%	71.8%
4	BUNNY	1.87 cm	4.20 cm	5.96 cm	85.9%	73.4%
4	MAN	1.63 cm	3.43 cm	4.56 cm	86.3%	75.6%
4	CYLINDER	1.71 cm	3.87 cm	4.69 cm	87.1%	78.3%
4	NO OBJECT	-	-	-	95.8%	89.3%
Distributed Lighting						
Object	Optimal lighting using distributed lighting			Method 1	Method 2	
MAN	1.56 cm			2.45 cm	4.32 cm	
BUNNY	2.36 cm			3.19 cm	5.78 cm	

Table 1: Simulated Results for 4 different walls and 4 objects, including the performance improvement due to adaptive lighting for the optimal patch versus 2nd and 3rd patches returned by the optimization. Distributed lighting optimizes both patch selection and power distribution for $m = 2$ patches, while Method 1 only performs adaptive patch selection with equal distribution of power and Method 2 does not optimize either the illumination or patches.

Tracking and Ablation Studies: In Figure 4, we show tracking results for our NLOS localization, as well as ablation studies for the effects of object size and noise. Figure 4(a) depicts the tracking results of a sphere, with the black trajectory denoting the ground truth and red denoting the network localization. When the sphere is within approximately 10 cm from the LOS wall (one-third of the LOS wall length), the localization error is smaller and it tends to increase progressively as the object moves farther away from the LOS wall. This is due to the loss of light due to the radiometric falloff as the NLOS object moves farther away, conforming to our intuitions.

For object size, we varied the diameter of a simulated sphere in the NLOS, and noticed how localization performance improves as the object gets bigger. This is intuitive since a larger object yields more reflected light back to the LOS which can be extracted by the network. Noise is usually added to make the model more robust to changes [14]. We additionally study the effect of adding noise to the training examples. In Figure 4(c), we add Gaussian noise to all training examples, and report the MSE performance of the resulting network. Note that the performance stays relatively flat, showing that the model is relatively robust to noise (this may be because the renderer itself has noise that helps regularize the network). In the supplemental material, we also perform saliency analysis to determine what parts of the LOS scene is salient for the NLOS localization task.

Real Experimental Results: For our hardware setup, we constructed a physical NLOS imaging prototype (35.6 cm \times 35.6 cm \times 35.6 cm dimensions) to validate our proposed method in the real world, as shown in Figure 3(C). The LOS scene used is a variation of




Figure 4: a) Tracking results for a sphere of diameter 5 cm for simulated data. The black trajectory denotes the actual path taken by object and red denotes the predicted position of the object by using the network proposed in Fig 2; b) Effect of shape of object on the MSE. As the diameter of the sphere increases, the MSE reduces, giving better accuracy in localization; c) Effect of adding noise to the dataset while training.

Object	Localization with Adaptive Lighting		Identification with Adaptive Lighting	
	Optimal Patch	2nd Patch	Optimal Patch	2nd Patch
MAN	1.55 cm	2.41 cm	90.1%	81.9%
CYLINDER	2.07 cm	3.12 cm	85.0%	75.8%
BUNNY	1.89 cm	2.82 cm	87.5%	78.6%
BALL	2.38 cm	3.45 cm	85.8%	77.2%
Distributed Lighting				
Object	Optimal Lighting		Method 1	Method 2
MAN	1.78 cm		2.97 cm	5.01 cm
CYLINDER	2.41 cm		3.28 cm	6.16 cm
BUNNY	2.55 cm		3.53 cm	5.43 cm
BALL	2.63 cm		4.10 cm	5.69 cm

Table 2: Real data results using our hardware prototype for four objects trained on the complex LOS wall 2. Note that our optimization algorithms consistently give better localization and identification performance as compared to the next-best patch illumination.

LOS Wall 2 in that figure, and consisted of black, white, and grey colors and coarse texture given by the 3D printer. Objects were also 3D printed. Further implementation details are provided in the supplemental material.

In Table 2, we present the results from our real hardware prototype. The adaptive lighting for a single patch resulted in an average MSE of 1.97 cm in localization, while the next best-patch had average error of 2.95 cm. The values reported in Table 2 is for a sphere (diameter of 5 cm), cylinder(radius of 3 cm and height of 8 cm), a bunny (width of 7.4 cm and height of 5 cm), and a man silhouette (width of 5.5cm and height of 17.5cm). For identification, we have similar improvements in accuracy due to optimal choice of lighting (87.1%) as opposed to the next-best choice (78.3%).

For the distributed method, we run our algorithm to optimize for two patches simultaneously in terms of location and illumination power. The adaptive algorithm achieves a MSE of 1.78 cm for the man and 2.55 cm for the bunny, 2.63 cm for the ball, and 2.41 cm for the cylinder; all of which outperform both Method 1 and Method 2. Our real results help validate the idea that adaptive lighting can help improve NLOS localization and identification in a real world scene.

6 Discussion

Our results provide several points of discussion, as well as further implications for adaptive lighting methods. We note the effectiveness of deep learning methods, even with fairly conventional CNN architectures, at being able to extract information from the NLOS light paths and perform localization. In our best real data result from our hardware prototype, we localize the centroid of an object of width 5.5 cm and height 17.5 cm to an MSE of 1.55 cm using our adaptive lighting method. This outperforms state-of-the-art results reported in the literature including 1.7 cm localization error for a 12 cm wide object (in 2D) from [84] and 6.1 cm localization error for a 10 cm wide object (in 3D) from [24]. Further, to the best of our knowledge, our paper is the first method to handle non-planar LOS scenes for NLOS imaging. Yet we caution the reader that it is difficult to make a fair comparison between different algorithms/methods as the objects, simulators, and training data used are not standardized. Recently, a quantitative benchmark for NLOS imaging was introduced that could potentially help in this regard [25]. To further these efforts, we publicly release our code and dataset at this Github link: https://github.com/sreenithy/AdaptiveLighting_NLOS for the research community.

There are several limitations for our research approach. It is unclear whether our method can perform full 3D NLOS object reconstruction with a conventional camera and projector. We showed that while our networks generalize with respect to camera motion, they are not tolerant to unknown NLOS objects or if the LOS scene changes. Finally, we have not considered occlusions in the NLOS that limits our modeling techniques for adaptive lighting.

There are several avenues for future work. More physically-realistic light transport modeling, including specular light paths, would enable better adaptive lighting and better data sets for training. We chose to optimize radiosity to boost the number of photons returning from the NLOS as a first step towards adaptive lighting. However, a full end-to-end network which jointly optimizes the lighting patterns with the inverse NLOS task can probably yield superior results. New architecture designs for our CNNs could potentially improve their performance and generalizability, including incorporating generative adversarial networks or performing joint localization and identification. Finally, the method could be engineered for real-time localization/tracking at inference by accelerating the radiosity optimization with techniques such as cached shadow maps implemented on a GPU. However, we hope that this paper is the first step towards intelligent lighting for NLOS imaging in the future.

Acknowledgments: We would like to thank Srinivasa Narasimhan, Kyros Kutulakos, and Ioannis Gkioulekas for initial discussions. Shenbagaraj Kannapiran from ASU helped render illustrations used in this paper. The AME Fabrication lab also assisted with 3D printing, scene setup, and lending the projector for our hardware prototype. This work was partially supported by the Defense Advanced Research Projects Agency (REVEAL Grant HR00111620021) to S. Jayasuriya, as well as joint support from both the Herberger Research Initiative in the Herberger Institute for Design and the Arts (HIDA) and the Fulton Schools of Engineering (FSE) at ASU to S. Chandran and S. Jayasuriya. This work was also supported by a hardware GPU donation by NVIDIA.

References

- [1] Supreeth Achar, Joseph R Bartels, William L Whittaker, Kiriakos N Kutulakos, and Srinivasa G Narasimhan. Epipolar time-of-flight imaging. *ACM Transactions on*

- Graphics*, 36(4):37, 2017.
- [2] Victor Arellano, Diego Gutierrez, and Adrian Jarabo. Fast back-projection for non-line of sight reconstruction. *Optics Express*, 25(10):11574–11583, 2017.
- [3] Manel Baradad, Vickie Ye, Adam B Yedidia, Frédo Durand, William T Freeman, Gregory W Wornell, and Antonio Torralba. Inferring light fields from shadows. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6267–6275, 2018.
- [4] Katherine L Bouman, Vickie Ye, Adam B Yedidia, Fredo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. Turning corners into cameras: Principles and methods. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2289–2297. IEEE, 2017.
- [5] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Eliceiri, and Andreas Velten. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics Express*, 23(16):20997–21011, 2015.
- [6] Piergiorgio Caramazza, Alessandro Boccolini, Daniel Buschek, Matthias Hullin, Catherine Higham, Robert Henderson, Roderick Murray-Smith, and Daniele Faccio. Neural network identification of people hidden from view with a single-pixel, single-photon detector. *Scientific Reports*, (8), 2018. ISSN 2045-2322.
- [7] Susan Chan, Ryan E Warburton, Genevieve Garipey, Jonathan Leach, and Daniele Faccio. Non-line-of-sight tracking of people at long range. *Optics Express*, 25(9):10109–10117, 2017.
- [8] Wenzheng Chen, Simon Daneau, Fahim Mannan, and Felix Heide. Steady-state non-line-of-sight imaging. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [9] Michael F Cohen and Donald P Greenberg. The hemi-cube: A radiosity solution for complex environments. In *ACM SIGGRAPH Computer Graphics*, volume 19, pages 31–40. ACM, 1985.
- [10] Michael F Cohen and John R Wallace. *Radiosity and realistic image synthesis*. Elsevier, 2012.
- [11] S. Cova, M. Ghioni, A. Lacaíta, C. Samori, and F. Zappa. Avalanche photodiodes and quenching circuits for single-photon detection. *Applied Optics*, 35(12):1956–1976, Apr 1996.
- [12] Genevieve Garipey, Francesco Tonolini, Robert Henderson, Jonathan Leach, and Daniele Faccio. Detection and tracking of moving objects hidden from view. *Nature Photonics*, 10(1):23–26, 2016.
- [13] Cindy M Goral, Kenneth E Torrance, Donald P Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. In *ACM SIGGRAPH Computer Graphics*, volume 18, pages 213–222. ACM, 1984.

- [14] Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2222–2232, 2017.
- [15] Otkrist Gupta, Thomas Willwacher, Andreas Velten, Ashok Veeraraghavan, and Ramesh Raskar. Reconstruction of hidden 3d shapes using diffuse reflections. *Optics Express*, 20(17):19096–19108, 2012.
- [16] Felix Heide, Lei Xiao, Wolfgang Heidrich, and Matthias B Hullin. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3222–3229, 2014.
- [17] Felix Heide, Matthew O’Toole, Kai Zhang, David Lindell, Steven Diamond, and Gordon Wetzstein. Non-line-of-sight imaging with partial occluders and surface normals. *ACM Transactions on Graphics*, 38, 2019.
- [18] Wenzel Jakob. Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.
- [19] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL <http://www.scipy.org/>. [Online; accessed 2019-04-29].
- [20] Achuta Kadambi, Hang Zhao, Boxin Shi, and Ramesh Raskar. Occluded imaging with time-of-flight sensors. *ACM Transactions on Graphics*, 35(2):15, 2016.
- [21] James T Kajiya. The rendering equation. In *ACM Siggraph Computer Graphics*, volume 20, pages 143–150. ACM, 1986.
- [22] Alexander Keller. Instant radiosity. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 49–56. ACM Press/Addison-Wesley Publishing Co., 1997.
- [23] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar. Looking around the corner using transient imaging. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 159–166. IEEE, 2009.
- [24] Jonathan Klein, Christoph Peters, Jaime Martín, Martin Laurenzis, and Matthias B Hullin. Tracking objects outside the line of sight using 2d intensity images. *Scientific Reports*, 6:32491, 2016.
- [25] Jonathan Klein, Martin Laurenzis, Dominik L. Michels, and Matthias B. Hullin. A quantitative platform for non-line-of-sight imaging problems. In *British Machine Vision Conference 2018, BMVC 2018, Northumbria University, Newcastle, UK, September 3-6, 2018*, page 104, 2018.
- [26] Hiroyuki Kubo, Suren Jayasuriya, Takafumi Iwaguchi, Takuya Funatomi, Yasuhiro Mukaigawa, and Srinivasa G Narasimhan. Acquiring and characterizing plane-to-ray indirect light transport. In *Computational Photography (ICCP), 2018 IEEE International Conference on*, pages 1–10. IEEE, 2018.

- [27] Marco La Manna, Fiona Kine, Eric Breitbart, Jonathan Jackson, Talha Sultan, and Andreas Velten. Error backprojection algorithms for non-line-of-sight imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [28] David B. Lindell, Gordon Wetzstein, and Matthew O’Toole. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM Transactions on Graphics*, 38(4):116:1–116:13, July 2019. ISSN 0730-0301.
- [29] Matthew O’Toole, Supreeth Achar, Srinivasa G Narasimhan, and Kiriakos N Kutulakos. Homogeneous codes for energy-efficient illumination and imaging. *ACM Transactions on Graphics*, 34(4):35, 2015.
- [30] Matthew O’Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338, 2018.
- [31] A. Pediredla, A. Dave, and A. Veeraraghavan. Snlos: Non-line-of-sight scanning through temporal focusing. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–13, May 2019.
- [32] Adithya Kumar Pediredla, Mauro Buttafava, Alberto Tosi, Oliver Cossairt, and Ashok Veeraraghavan. Reconstructing rooms using photon echoes: A plane based model and reconstruction algorithm for looking around the corner. In *2017 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 2017.
- [33] Guy Satat, Matthew Tancik, Otkrist Gupta, Barmak Heshmat, and Ramesh Raskar. Object classification through scattering media with deep learning on time resolved measurement. *Optics Express*, 25(15):17466–17479, 2017.
- [34] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R Marschner, Mark Horowitz, Marc Levoy, and Hendrik Lensch. Dual photography. *ACM Transactions on Graphics*, 24(3):745–755, 2005.
- [35] Pierre Sermanet and Yann LeCun. Traffic sign recognition with multi-scale convolutional networks. In *International Joint Conference on Neural Networks*, pages 2809–2813, 2011.
- [36] Brandon M Smith, Matthew O’Toole, and Mohit Gupta. Tracking multiple objects outside the line of sight using speckle imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6258–6266, 2018.
- [37] Matthew Tancik, Guy Satat, and Ramesh Raskar. Flash photography for data-driven hidden scene recovery. *arXiv preprint arXiv:1810.11710*, 2018.
- [38] Matthew Tancik, Tristan Swedish, Guy Satat, and Ramesh Raskar. Data-driven non-line-of-sight imaging with a traditional camera. In *Imaging Systems and Applications*, pages IW2B–6. Optical Society of America, 2018.
- [39] Christos Thrampoulidis, Gal Shulkind, Feihu Xu, William T Freeman, Jeffrey Shapiro, Antonio Torralba, Franco Wong, and Gregory Wornell. Exploiting occlusion in non-line-of-sight active imaging. *IEEE Transactions on Computational Imaging*, 2018.
- [40] Antonio Torralba and William T Freeman. Accidental pinhole and pinspeck cameras. *International Journal of Computer Vision*, 110(2):92–112, 2014.

- [41] 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 *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2336–2344. IEEE, 2017.
- [42] Tomoki Ueda, Hiroyuki Kubo, Suren Jayasuriya, Takuya Funatomi, and Yasuhiro Mukaigawa. Slope disparity gating using a synchronized projector-camera system. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–9. IEEE, 2019.
- [43] Eric Veach. *Robust monte carlo methods for light transport simulation*. Number 1610. Stanford University PhD thesis, 1997.
- [44] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounsi G Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3:745, 2012.
- [45] Jian Wang, Joseph Bartels, William Whittaker, Aswin C Sankaranarayanan, and Srinivasa G Narasimhan. Programmable triangulation light curtains. In *European Conference on Computer Vision*, pages 20–35. Springer, 2018.
- [46] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6800–6809, 2019.
- [47] Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, and Thomas Funkhouser. Physically-based rendering for indoor scene understanding using convolutional neural networks. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.