Soft Shadow Diffusion (SSD): Physics-inspired Learning for 3D Computational Periscopy

Fadlullah Raji[®] and John Murray Bruce[®]

Computer Science and Engineering, University of South Florida 4202 E. Fowler Avenue, Tampa, FL, USA 33620. {fraji, murraybruce}@usf.edu

Abstract. Conventional imaging requires a line of sight to create accurate visual representations of a scene. In certain circumstances, however, obtaining a suitable line of sight may be impractical, dangerous, or even impossible. Non-line-of-sight (NLOS) imaging addresses this challenge by reconstructing the scene from indirect measurements. Recently, passive NLOS methods that use an ordinary photograph of the subtle shadow cast onto a visible wall by the hidden scene have gained interest. These methods are currently limited to 1D or low-resolution 2D color imaging or to localizing a hidden object whose shape is approximately known. Here, we generalize this class of methods and demonstrate a 3D reconstruction of a hidden scene from an ordinary NLOS photograph. To achieve this, we propose a novel reformulation of the light transport model that conveniently decomposes the hidden scene into light-occluding and non*light-occluding* components to yield a separable non-linear least squares (SNLLS) inverse problem. We develop two solutions: A gradient-based optimization method and a physics-inspired neural network approach, which we call Soft Shadow diffusion (SSD). Despite the challenging illconditioned inverse problem encountered here, our approaches are effective on numerous 3D scenes in real experimental scenarios. Moreover, SSD is trained in simulation but generalizes well to unseen classes in simulation and real-world NLOS scenes. SSD also shows surprising robustness to noise and ambient illumination.

Keywords: Computational imaging \cdot Machine learning \cdot 3D generative models \cdot Diffusion models \cdot Separable non-linear least squares

1 Introduction

Collision avoidance systems in autonomous navigation, detection of hidden adversaries in military operations, danger assessment in search-and-rescue operations, and arterial blockage assessment in biomedical imaging are examples of applications that will benefit from the ability to form images of scenes beyond an observer's line of sight. Because of its breadth of potential applications, this field called non-line-of-sight (NLOS) imaging continues its rapid growth. The fundamental underlying goal is to infer information about the NLOS scene from



Fig. 1: NLOS imaging configuration. Non-occluding objects diffusely reflect light toward the visible wall; the reflected light is partially occluded by an occluding object in the hidden scene to create a soft shadow on the visible wall.

measurements of signals, such as light [15,25,51,52], sound [4,28], heat [22,37], or other modality [1], that reach the observer after being diffusely reflected or emitted from the NLOS scene.

In this work, we demonstrate 3D NLOS imaging from measurements of the light reaching a surface that is visible to the observer and the hidden scene. Prior methods that similarly rely on measurements of light fall into either of two categories: *active methods*, which exploit controlled illumination of the hidden scene [15, 17, 25, 29, 32–34, 39, 42, 43, 47, 53, 59, 62]; or *passive methods*, which exploit already existing light without any control over the hidden scene illumination [3, 5, 13, 16, 40, 52, 54, 55, 65, 70]. Active methods have demonstrated remarkable success in achieving high-resolution 3D NLOS imaging [42, 43, 62], by using ultrafast pulsed-laser illumination, and time-resolved single-photon detection. Despite these advances, the complexity and cost of the equipment required by active methods can be a limiting factor in several applications. In contrast, passive methods offer a simpler alternative but are currently limited to low-resolution 1D [5,54] or 2D [51,52,55,70] reconstructions of hidden scenes. A recent method [13] exploits two orthogonal edges of a door frame to achieve remarkable 3D imaging of non-occluding, hidden-scene objects. However, the exploited shadows are cast by the visible, known door frame occluder. Here, we generalize occluder-aided passive methods to include 3D imaging of occluders that are hidden or otherwise unknown.

Contributions: In this paper, we propose a reformulation of the computational periscopy as a separable non-linear least squares (SNLLS) inverse problem that enables 3D reconstructions from a single penumbra photograph, and we offer a gradient-based optimization to solve it. In addition, we introduce *Soft Shadow*

Diffusion (SSD) a novel generative model based on denoising probabilistic diffusion for reconstructing 3D shapes, as high-resolution 3D point clouds, from a single photograph of their soft shadows. Finally, we experimentally demonstrate the first joint reconstructions of the 3D occluding and the 2D non-occluding objects in a real experimental scene.

1.1 Related work

Computational Periscopy: Figure 1(a) shows a photograph of the NLOS imaging scenario considered here and by prior related works (e.g., [3,51,52,70]). In this scenario, a preexisting ambient light source illuminates the hidden scene area (and the *visible wall* in the observer's LOS). Light reflected, toward the visible wall, from some parts of the hidden scene (here, the chair and basketballs) is partially occluded by other hidden scene objects (here, the spade-like shape) producing a penumbra (or soft shadow) on the visible wall.

From simple photographs of penumbra, Saunders et al. [51,52] and Baradad et al. [3] produce reconstructions of the light-reflecting (non-light-occluding) portions of the hidden scene by assuming that the light-occluding hidden scene structures have a known shape but unknown location, or that projections of the occluder on the visible wall are available from a prior calibration step. Furthermore, by assuming that the occluder is approximately planar, Yedidia et al. [70] reconstruct its 2D projection and a 2D video of the non-light-occluding portion of the hidden scene from a video of the visible penumbra. In this paper, we develop a new computational strategy to reconstruct the 3D structure of the light-occluding objects from a 2D photograph of penumbra. Hence, the output computed by our method can augment these prior approaches by providing the necessary prior knowledge about the hidden occluders.

Instead of reconstructing its image, directly inferring the contextual properties of a hidden scene may be important in various applications. Along those lines, Sharma et al. proposed learning-based methods for accurately determining human activity or classifying the number of people in an unknown room from a video recording of penumbra [56]. Medin et al. further showed that a single penumbra encodes enough information for biometric identification of an individual without a line of sight [38]. Although direct photographs of sharp human silhouette shadows have long been explored for biometrics [20, 63], Medin et al. [38] consider a challenging variation: The observed shadow is instead of an object that occludes light reflected by the face of the hidden individual. Rather impressively, both methods [38,56] demonstrate accurate inference without prior knowledge of the hidden occluder or scene geometry.

Shape from S: Shape from shading [19, 46] and shape-from-shadow [7] are areas of longstanding and broad interest to the computer vision community. Shape from shading reconstructs a 3D shape from its 2D image by analyzing variations in brightness (shading) across the image, to deduce the orientation of each point on the object's surface. In contrast, shape from shadow reconstructs a 3D shape

from several (binary) shadow images, created by varying the illumination direction, by analyzing the darkest parts across the image stacks to deduce the depth and surface orientations. Recent approaches have exploited data-driven approaches to yield impressive reconstructions [23, 30, 66], even in a monocular setting [9]. These prior works, exploit cues from shading due to direct illumination and sharp shadows caused by mutual occlusion under known or possibly unknown illuminations. Here, we use a single image of the soft shadow created on a nearby surface, by hidden 3D occluding objects partially blocking a complex, diffuse, and hidden illumination both lying outside the image's field of view.

Diffusion Probabilistic Models: Generative diffusion probabilistic models [18,58,60]. These models have demonstrated high-quality image generation [49, 67], rivaling generative adversarial networks (GANs) [14]. Their application has broadened to include tasks such as time-series forecasting [48], audio generation [26], text-to-speech [27, 45]. Several works have explored these models for point cloud reconstructions via a two-stage generation using flow models [68] or diffusion priors [6, 72]. To reduce generation latency, Zhou et al. [74] proposed PVD, a single-stage diffusion prior for generating point clouds. Beyond point clouds, text-conditional 3D generation has advanced with methods like Dream-Fields [21] and CLIP-guided mesh optimization [24]. Other approaches focus on reconstructing 3D from multi-view images using diffusion-based priors [31,44,69].

This work is a novel extension of diffusion models to generate 3D point clouds from a 2D soft shadow photograph. We subsequently use a trained regression model to generate a 3D mesh representation of the object casting the shadow. Building on prior work that performs diffusion over individual points, our approach conditions the diffusion process on soft shadow images.

2 Light Propagation Model

Consider the configuration shown in Figure 1(a) and let \mathbf{x} denote any point in the hidden scene volume. A hidden scene point whose light paths to the visible wall portion within the camera's FOV are (un)occluded along some directions, but not all, is well-conditioned for recovery. The set \mathcal{X} of all such hidden scene points forms the *computational FOV* (CFOV) of the imaging system [52]. Let $f(\mathbf{x})$ denote the unknown radiosity distribution at a hidden scene point \mathbf{x} , then the irradiance $i(\mathbf{p})$ of a visible wall point \mathbf{p} is [52]:

$$i(\mathbf{p}) = \int_{\mathbf{x}\in\mathcal{X}} \frac{g(\mathbf{p},\mathbf{x})}{\|\mathbf{x}-\mathbf{p}\|_2^2} v(\mathbf{x},\mathbf{p};\boldsymbol{\theta}_{\mathrm{o}}) f(\mathbf{x}) \,\mathrm{d}\mathbf{x} + b(\mathbf{p}), \tag{1}$$

where $g(\mathbf{p}, \mathbf{x}) = \cos(\measuredangle(\mathbf{x} - \mathbf{p}, \mathbf{n}_{\mathbf{x}})) \cos(\measuredangle(\mathbf{p} - \mathbf{x}, \mathbf{n}_{\mathbf{p}}))$ models Lambertian reflection and foreshortening with $\mathbf{n}_{\mathbf{p}}$ and $\mathbf{n}_{\mathbf{x}}$ representing the surface normals at \mathbf{p} and \mathbf{x} , respectively. The function $v(\mathbf{x}, \mathbf{p}; \boldsymbol{\theta}_{o})$ models the visibility between \mathbf{p} and \mathbf{x} for an unknown occluding object whose shape is implicitly parameterized by $\boldsymbol{\theta}_{o}$; $v(\mathbf{x}, \mathbf{p}; \boldsymbol{\theta}_{o}) = 1$ when the ray from \mathbf{x} to \mathbf{p} is unoccluded, otherwise it is 0. Finally, $b(\mathbf{p})$ models contributions due to noise, and other sources of illumination—predominantly any existing visible side illumination sources and illumination

from hidden scene points outside the CFOV, i.e. $\int_{\mathbf{x}\notin\mathcal{X}} \frac{g(\mathbf{p},\mathbf{x})}{\|\mathbf{x}-\mathbf{p}\|_2^2} f(\mathbf{x}) d\mathbf{x}$ —that do not form penumbra. The recoverable hidden scene \mathcal{X} is typically assumed to be a single plane at a known depth [2,52,70], or a few planes [3,51] at known depths, from the visible wall. Discretizing the measurement plane and the hidden scene plane(s) into pixels gives the discrete forward model:

$$\mathbf{y} = \mathbf{A}(\boldsymbol{\theta}_{\mathrm{o}})\mathbf{f} + \mathbf{b},\tag{2}$$

where $\mathbf{y} \in \mathbb{R}^M$ is a vectorization of the measured photograph, $\mathbf{f} \in \mathbb{R}^N$ is the vector containing the unknown radiosities of the hidden scene pixels, and $\mathbf{A}(\boldsymbol{\theta}_{\mathrm{o}}) \in \mathbb{R}^{M \times N}$ is the light transport matrix given $\boldsymbol{\theta}_{\mathrm{o}}$ (which implicitly represents the shape and position of the occluding objects).

Saunders et al. [52] assume a known occluder shape but use θ_{o} to represent its location. This work considers the case where the occluder's shape is also unknown and uses θ_{o} as an implicit representation for the 3D shape of an unknown occluder. Estimating the unknowns θ_{o} , **f** and **b** using (2) is an ill-conditioned *separable nonlinear inverse problem*. In Section 4 we present two solutions for this inverse problem by building on our model and explicitly delineating the hidden scene as comprising two components: those that act as occluders of light, and those that act only as emitters/reflectors of light.



Fig. 2: Computational FOV and two simple examples of occluders. (a) The camera's FOV is projected towards the non-occluding scene plane. The occluding region is the volume created by joining the boundaries of the measurement FOV to the boundaries of the desired hidden scene plane. (b) An ideal pinhole projects an image of the scene. (c) An ideal pinspeck projects a negative image of the scene.

2.1 Hidden Scene Decomposition and Computational FOV

We aim to reconstruct simultaneously a 3D image of the hidden occluders and a 2D image of the hidden non-occluders (i.e., emitters/reflectors of light). To obtain a feasible inverse problem and facilitate accurate reconstructions, we must restrict the reconstruction volume to include only hidden scene portions that can cast penumbra somewhere within the photographed portion of the visible wall. The 2D image of the non-occluders is assumed to be the radiosity of a plane of fixed size $W_x \times H_x$ at a depth D from the visible wall (see Figure 2a). Then, as illustrated in Figure 2a, the occluding volume is then defined by the frustum

whose edges coincide with the lines formed by connecting each vertex of the camera FOV to each vertex of the non-occluding hidden scene plane.

For our problem setting, the 3D *occluding volume* and the chosen 2D *non-occluding plane*, are collectively referred to as the computational FOV (CFOV). This represents a natural extension of the CFOV concept [52] from 2D to 3D.

3 3D Computational Periscopy: Models & Representation

To facilitate 3D reconstruction using our new reformulation, we consider two different representations of 3D structures. The first, presented in Section 3.2, is an explicit representation that uses a uniform voxelization of the occluding frustum portion of the CFOV. The second, presented in Section 3.4, is an implicit neural point cloud-based representation, which imposes a *deep prior* for 3D shapes. Each representation inspires a corresponding reconstruction method presented in Section 4 and builds upon a simple pinspeck occluder.

3.1 A Simple Pinspeck Occluder

A pinspeck (or anti-pinhole) occluder [11, 61] is the optical complement of a pinhole. Thus, a single *pinspeck* at $\mathbf{u} = (u_x, u_y, u_z)$ blocks all rays that reach the point \mathbf{u} such that the visibility between a point $\mathbf{p} = (p_x, 0, p_z)$ on the visible wall, and a point $\mathbf{x} = (x, y, z)$ on the non-occluding hidden scene plane is

$$v_{\rm ps}(\mathbf{p}, \mathbf{x}; \mathbf{u}) = 1 - \delta \left(p_x - mu_x - \frac{xy}{u_y} \right) \delta \left(p_z - mu_z - \frac{yz}{u_y} \right) \tag{3}$$

where $\delta(\cdot)$ is the Dirac delta function, and $m = 1 - y/u_y$. Imposing the optical complementarity property [11, 61], gives the following expression for the (negative) image formed on the visible plane:

$$i(\mathbf{p}) = \int_{\mathbf{x}\in\mathcal{X}} \frac{g(\mathbf{p},\mathbf{x})}{\|\mathbf{x}-\mathbf{p}\|_2^2} \left(1 - v_{\rm ph}(\mathbf{p},\mathbf{x};\mathbf{u})\right) f(\mathbf{x}) \,\mathrm{d}\mathbf{x},\tag{4}$$

where the pinhole visibility $v_{\rm ph}(\mathbf{p}, \mathbf{x}; \mathbf{u}) = 1 - v_{\rm ps}(\mathbf{p}, \mathbf{x}; \mathbf{u})$.

3.2 Uniform Voxels Representation

Assume the light-occluding frustum is uniformly discretized into K voxels whose centers are located at $\{\mathbf{u}_k\}_{k=1}^K$, then with a dense discretization, arbitrary 3D shapes can be well-approximated by a subselection of voxels identified by their indices. In this scenario, $v(\mathbf{x}, \mathbf{p}; \boldsymbol{\theta}_0) = 1 - \sum_{k=1}^{K} \alpha_k v_{\text{ph}}(\mathbf{p}, \mathbf{x}; \mathbf{u}_k)$ where the binary-valued scalar α_k indicates the presence (1) or absence (0) of a pinspeck at \mathbf{u}_k , such that the visible wall irradiance is approximately:

$$i(\mathbf{p}) = \int_{\mathbf{x}\in\mathcal{X}} \frac{g(\mathbf{p},\mathbf{x})}{\|\mathbf{x}-\mathbf{p}\|_2^2} \left(1 - \sum_{k=1}^K \alpha_k v_{\mathrm{ph}}(\mathbf{p},\mathbf{x};\mathbf{u}_k) \right) f(\mathbf{x}) \,\mathrm{d}\mathbf{x}.$$
(5)

We discretize the non-occluding plane into $N = (N_{\rm x} \times N_{\rm z})$ pixels, and represent the unknown radiosity of the *n*-th pixel by f_n and the collection of all N radiosities by $\mathbf{f} \in \mathbb{R}^N$. Additionally, we define $\{\mathbf{p}_m\}_{m=1}^M$ as the locations of pixels on the visible wall that are captured by the $M = (M_{\rm x} \times M_{\rm z})$ camera pixels. Thus,

$$\mathbf{y} = \mathbf{A} \left(1 - \sum_{k=1}^{K} \alpha_k \mathbf{V}_k \right) \mathbf{f} = \left(\mathbf{A} - \sum_{k=1}^{K} \alpha_k \mathbf{A} \mathbf{V}_k \right) \mathbf{f}, \tag{6}$$

where $\mathbf{y} \in \mathbb{R}^M$ is the vectorized penumbra photograph, $\mathbf{V}_k \in \mathbb{R}^{M \times N}$ is the matrix corresponding to visibility of a pinhole at \mathbf{u}_k . Reconstructing the 3D shape is, thus, equivalent to computing $\boldsymbol{\theta}_0 \stackrel{\text{def}}{=} \boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_K)^{\mathsf{T}}$ from \mathbf{y} .

3.3 Point Cloud Representation

While a uniform voxelization has the potential to provide accurate scene representations with finer discretizations; it, however, leads to increased computational complexity, memory requirements, and poor conditioning. Furthermore, a binary constraint on $\boldsymbol{\alpha}$ yields a nonconvex optimization. In contrast, a point cloud representation deviates from the uniform grid-based, directly optimizing the vector locations \mathbf{u}_k instead of the binary vector $\boldsymbol{\alpha}$. It also provides a lower-dimensional representation that effectively captures a wide range of 3D shapes while being compatible with our pinspeck-based model (4). Thus, we aim to solve $\mathbf{y} = \mathbf{A}(\boldsymbol{\theta}_0)\mathbf{f}$, where $\boldsymbol{\theta}_0 \stackrel{\text{def}}{=} {\{\mathbf{u}_k\}_{k=1}^K}$. Directly optimizing the locations of K points is challenging. Nonetheless, to facilitate accurate and high-resolution reconstructions we develop and exploit a learned point cloud shape prior which we call **soft shadow diffusion**.

3.4 Soft Shadow Diffusion Model

Our proposed soft shadow diffusion (SSD) model adapts the state-of-the-art denoising diffusion models [36, 57, 58, 73] to conditional point cloud generation. Here, the denoising diffusion process is conditioned on a latent representation $\hat{\mathbf{y}}$ of a measured soft shadow photograph \mathbf{y} . Our motivation is to leverage SSD as a powerful learned prior for high-resolution 3D point clouds that explain the soft shadows in a penumbra photograph.

Once trained, our SSD model $G_w(\mathbf{y})$ is used to generate a denoised point cloud $\{\mathbf{u}_k^0\}_{k=1}^K$ by iteratively reversing a Markov forward process, as depicted in Figure 3b. Specifically, conditioned on a soft shadow photograph latent $\hat{\mathbf{y}}$, the model iteratively reduces the noise from an initial point cloud $\{\mathbf{u}_k^T\}_{k=1}^K$ that is a random realization of a Gaussian noise process.

The training process (see Figure 3a) begins by setting $\{\mathbf{u}_k^0\}_{k=1}^K$ to be a random sample from our dataset of noise-free point clouds. Then, a series of progressively noisier point clouds $\{\mathbf{u}_k^t\}_{k=1}^K$ for t = 1, 2, ..., T are computed from $\{\mathbf{u}_k^0\}_{k=1}^K$ using $\mathbf{u}_k^t = \sqrt{\overline{\alpha}_t}\mathbf{u}_k^0 + \sqrt{1 - \overline{\alpha}_t}\epsilon$ for k = 1, 2, ..., K, where ϵ is zero-mean



Fig. 3: Soft Shadow Diffusion Model Pipeline. (a) In training, the model maps a clean input point cloud to noise using the forward diffusion process. (b) At inference, the model starts from an initial noise and transforms this noise into a point cloud representation of light-occluding structure by conditioning on the encoded soft shadow photograph $\hat{\mathbf{y}}$. The point cloud is then converted into a mesh representation of the occluding structure.

white Gaussian noise with unit variance, and the coefficient α_t gradually decreases as the time step t increments. With these noisy samples, the weights w of the underlying neural network are updated by minimizing the following loss function [60]:

$$L = \mathbb{E}_{t, \{\mathbf{u}_k^0\}_{k=1}^K, \epsilon} \left\| \epsilon_{\theta}(\{\mathbf{u}_k^t\}_{k=1}^K, \hat{\mathbf{y}}, t) - \epsilon \right\|_2^2.$$

We provide additional details on training and sampling from our SSD model in the supplement.

4 3D Computational Periscopy: Inversion

Assuming background contributions to the term \mathbf{b} is small, a solution to (2) can be obtained by solving the Tikhonov-regularized minimization

$$\underset{(\mathbf{f},\boldsymbol{\theta}_{o})}{\arg\min} \|\mathbf{A}(\boldsymbol{\theta}_{o})\mathbf{f} - \mathbf{y}\|_{2}^{2} + \lambda \|\mathbf{f}\|_{2}^{2},$$
(7)

where $\|\cdot\|_2$ is the 2-norm, and $\lambda > 0$ is the regularization parameter. An alternating approach for the minimization (7) is,

$$\widehat{\mathbf{f}}^{k} = \arg\min_{\mathbf{f}} \left\| \mathbf{A}(\widehat{\boldsymbol{\theta}}_{o}^{k-1})\mathbf{f} - \mathbf{y} \right\|_{2}^{2} + \lambda \|\mathbf{f}\|_{2}^{2}$$

$$= \left(\mathbf{A}^{\mathsf{T}}(\widehat{\boldsymbol{\theta}}_{o}^{k-1})\mathbf{A}(\widehat{\boldsymbol{\theta}}_{o}^{k-1}) + \lambda \mathbf{I} \right)^{-1} \mathbf{A}^{\mathsf{T}}(\widehat{\boldsymbol{\theta}}_{o}^{k-1})\mathbf{y}$$

$$(8)$$

$$\widehat{\boldsymbol{\theta}}_{o}^{k} = \underset{\boldsymbol{\theta}_{o}}{\arg\min} \left\| \mathbf{A}(\boldsymbol{\theta}_{o}) \widehat{\mathbf{f}}^{k} - \mathbf{y} \right\|_{2}^{2}.$$
(9)

Moreover, by replacing f in (7) with its closed-form estimate, we obtain the minimization problem:

$$\widehat{\boldsymbol{\theta}}_{o} = \underset{\boldsymbol{\theta}_{o}}{\arg\min} \left\| \left[\left(\mathbf{A}^{\mathsf{T}}(\boldsymbol{\theta}_{o}) \mathbf{A}(\boldsymbol{\theta}_{o}) + \lambda \mathbf{I} \right)^{-1} \mathbf{A}^{\mathsf{T}}(\boldsymbol{\theta}_{o}) - \mathbf{I} \right] \mathbf{y} \right\|_{2}^{2}.$$
(10)

In principle, reconstructing θ_{o} from penumbra measurement **y** is feasible without explicitly forming estimates of **f** as an intermediate step. We use this observation in a physics-inspired neural network inversion model outlined in Section 4.2.

4.1 Inversion via Gradient-based Optimization

Using the uniform voxels representation (described in Section 3.2), we wish to estimate $\boldsymbol{\theta}_{o} \stackrel{\text{def}}{=} \{\alpha_{k}\}_{k=1}^{K}$ and **f** from **y** such that (2) holds. To that end, we allow the background to be non-negligible and solve the equivalent Tikhonov-regularized minimization problem:

$$\underset{(\mathbf{f},\boldsymbol{\theta}_{o},\mathbf{b},\lambda)}{\arg\min} \|\mathbf{A}(\boldsymbol{\theta}_{o})\mathbf{f} + \mathbf{b} - \mathbf{y}\|_{2}^{2} + \lambda \|\mathbf{f}\|_{2}^{2}.$$
 (11)

To impose convexity, the binary constraint is relaxed by using $\sigma(\mathbf{z}_k)$, where $\sigma(\cdot)$ is the sigmoid function, as a proxy for the binary-valued α_k . Hence, we instead optimize \mathbf{z}_k using gradient descent. This approach demonstrates robust performance on simulated datasets. In real experiments, we found that incorporating a positivity constraint when background and noise contributions are not negligible is more effective.

Algorithm 1 summarizes the steps of our voxelized optimization-based approach. Empirical results (see supplementary material) show superior convergence properties when directly optimizing for $\mathbf{f}, \mathbf{z}, \mathbf{b}$ and λ in an alternating fashion using Algorithm 1, compared to assuming a negligible background.

| Algorithm 1 Alternating Minimization Method | |
|------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Require: $\mathbf{A}, [\mathbf{V}_1, \mathbf{V}_2, \mathbf{V}_3,, \mathbf{V}_K], \mathbf{y}, \texttt{numIte}$ | r |
| 1: Initialize \mathbf{z}_0 randomly, $\mathbf{b}_0 = 0$, and $\lambda_0 =$ | 1 |
| 2: Initialize step sizes η_b , η_z , and η_λ | |
| 3: for $i = 1$ to numIter do | |
| 4: $\mathbf{A}_v = A \odot \left(1 - \frac{1}{K} \sum_{k=1}^K \mathbf{V}_k \cdot \sigma(\mathbf{z}_{i_k})\right)$ | Impose binary constraint |
| 5: $\mathbf{f}^* = (\mathbf{A}_v^{T} \mathbf{A}_v + \lambda_{i-1} I)^{-1} \mathbf{A}_v^{T} \mathbf{y}$ | Estimate 2D non-occluding scene |
| 6: $\mathbf{f}_i = \max(\mathbf{f}^* + \mathbf{b}_{i-1}, 0)$ | |
| 7: $\mathcal{L}(\mathbf{y}, \mathbf{A}_v \mathbf{f}_i) = \frac{1}{M} \sum_{j=1}^M \ \mathbf{y}_j - (\mathbf{A}_v \mathbf{f}_i)_j\ $ | 22 |
| 8: $\mathbf{b}_i \leftarrow \mathbf{b}_{i-1} - \eta_b \frac{\partial \mathcal{L}(\mathbf{y}, \mathbf{A}_v \mathbf{f}_i)}{\partial \mathbf{b}_{i-1}}$ | Update background contribution estimate |
| 9: $\mathbf{z}_i \leftarrow \mathbf{z}_{i-1} - \eta_z \frac{\partial \mathcal{L}(\mathbf{y}, \mathbf{A}_v \mathbf{f}_i)}{\partial \mathbf{z}_{i-1}}$ | Update occluder estimate |
| 10: $\lambda_i \leftarrow \lambda_{i-1} - \eta_\lambda \frac{\partial \mathcal{L}(\mathbf{y}, \mathbf{A}_v \mathbf{f}_i)}{\partial \lambda_{i-1}}$ U | Jpdate estimation of the regularization constant |
| 11: end for | |
| 12: Return $\widehat{\mathbf{f}} = \mathbf{f}_i, \widehat{\boldsymbol{\alpha}} = \sigma(\mathbf{z}_i) > 0.5$ | |

Implementation details: The alternating minimization algorithm (Algorithm 1) is implemented in PyTorch. The occluder region is discretized into $N_x \times N_y \times N_z$ uniform voxels. A high discretization is required for occluders. The required computation and memory are high, even for coarse discretizations like $10 \times 5 \times 10$.

To alleviate the high complexity and memory needed to precompute the matrix \mathbf{V}_k for all possible pinspeck contributions, we employ a sparse matrix representation for \mathbf{V}_k . This approach stores only the specific pixel locations to which each pinspeck contributes. The required gradient updates (for all pinspeck proxies $\{\mathbf{z}_k\}_k$) are computed in this new representation. (See supplementary material for more details and simulated examples.)

4.2 Inversion via Physics-inspired Neural Network

Here, we propose a reconstruction approach that unifies our physics-based separable model with SSD as a learned shape prior. Specifically, because SSD generates plausible point clouds of 3D shapes alone and not their locations in the NLOS scene, we augment our estimation pipeline with an adaptation of the location estimation algorithm in [52]. Our entire reconstruction pipeline, which includes estimation of \mathbf{f} , proceeds as follows:

 Stage 1: Conditional generation of 3D shapes. Sample the 3D shape from a conditional diffusion model given the soft shadow photograph y.

$$\{\widehat{\mathbf{u}}_k^0\}_{k=1}^K = G_w(\mathbf{y}),\tag{12}$$

- Stage 2: Localization of generated 3D shapes. Estimate the occluding object's spatial location by finding the 3D translation vector, $\widehat{\delta_{\mathbf{u}}}$, with: $\widehat{\delta_{\mathbf{u}}} = \arg \max_{\delta_{\mathbf{u}}} \left\| \mathbf{H}(\{\widehat{\mathbf{u}}_{k}^{0} + \delta_{\mathbf{u}}\}_{k=1}^{K}) \mathbf{y} \right\|_{2}^{2}$, where $\mathbf{H}(\cdot) = \mathbf{A}(\cdot) \left(\mathbf{A}(\cdot)^{T} \mathbf{A}(\cdot)\right)^{-1} \mathbf{A}(\cdot)^{T}$. - Stage 3: Reconstruction of the non-occluding component. We con-
- Stage 3: Reconstruction of the non-occluding component. We convert the properly translated occluder point cloud $\{\widehat{\mathbf{u}}_{k}^{0} + \widehat{\boldsymbol{\delta}_{\mathbf{u}}}\}_{k=1}^{K}$, and compute its mesh representation to obtain $\widehat{\boldsymbol{\theta}}_{o}$. We then formulate and solve the following TV regularized problem: $\widehat{\mathbf{f}} = \arg\min_{\mathbf{f}} ||\mathbf{y} \mathbf{A}(\widehat{\boldsymbol{\theta}}_{o})\mathbf{f}||_{2}^{2} + \lambda ||\mathbf{f}||_{TV}^{2}$.

Implementation details: We implement our physics-inspired neural network model, which includes the SSD network, in PyTorch.

Datasets. We train our SSD model on thousands of 3D models. We rendered all 3D models in the ShapeNet [8] dataset from 3 random camera angles as RGB images using Blender [12]. These images depict the non-occluding objects confined on the 2D plane, while the 3D models are the light-occluding structures. Our Blender script normalizes each 3D model to a bounding unit cube, configures a standard lighting setup, and exports RGB images using Blender's built-in realtime rendering engine. We then utilized the rendering equation in (2) on all possible pairs of rendered images and 3D models; and randomly selected pairs to simulate a diverse set of training examples. To produce point clouds of a 3D model, we first constructed a dense point cloud for each object by sampling 50,000 points along the 3D object's surface. The dense point cloud of unevenly spaced points is then resampled to produce a point cloud of 2,048 evenly spaced points. Our dataset is roughly 260,000 simulated instances of various object categories.

Soft Shadow Image Encoder. To encode the soft shadow photographs into latent representation, $\hat{\mathbf{y}}$, we employed the CrossFormer architecture [64] and trained it jointly with the diffusion process. CrossFormer introduces a cross-scale attention mechanism crucial for processing visual inputs. In addition, unlike traditional vision transformers, CrossFormer effectively captures the correlation between features of different scales, which is vital for understanding complex visual inputs like our penumbra photographs.

Model Implementation. Deviating from previous work on point cloud reconstructions that often employed specialized 3D architectures [74] or transformerbased [41] models, we take advantage of a simpler yet effective UNet-based architecture [50] with 1D convolutions. Our model allows fast sampling of more points than the model was trained on, and is designed to predict the mean μ of the noise in (3.4), from the inputs $\hat{\mathbf{y}}$, t, and $\{\mathbf{u}_k^t\}_{k=1}^K$, i.e., the encoded representation of a soft shadow 2D photograph, current time step t during the diffusion process, and the intermediate noisy point cloud, respectively. The entire model architecture is illustrated in Figure 3 and additional details are provided in the supplementary document.

3D mesh representation. Reconstructing the non-occluding portion of the hidden scene **f** requires accurate rendering of soft shadow images of the estimated occluding objects. Sparse point clouds may lead to inaccuracies that yield poor approximations of the soft shadows. Hence, we convert the generated point clouds from soft shadow diffusion into meshes. Using the entire ShapeNet [8] dataset, we train a signed distance field (SDF) neural network model to predict the SDF of the reconstructed point cloud. The SDF model follows the autoencoder architecture [10]. We then apply the marching cubes algorithm [35] to the computed SDF to generate a mesh representation of the light-occluding object. The mesh representations facilitate higher fidelity renders of the soft shadows.

5 Experiments

Experimental results for real and synthetic data are presented to demonstrate the efficacy of our proposed methods. Because prior methods focus on different acquisition scenarios (such as access to video frames [70, 71], pre-calibrated occluder visibility functions [3], known occluder shapes [52], or even fully known occluders [16,65]) direct and fair comparisons are not possible. These prior techniques focus on reconstructing the light-emitting portion of the hidden scene alone, whereas our approach is the first to seek, in addition, a 3D reconstruction of hidden occluders. We are considering a much harder problem and as such the 2D reconstructions achievable will, in expectation, be no better than existing methods.



(b) SSD-based reconstruction on Large 3D Objects.

Fig. 4: Real Experimental Reconstruction Results from Soft Shadow Diffusion Model. The model takes the soft shadow photographs (column 3) as input to condition the diffusion process and generates a 3D mesh representation (column 6) of the 3D light-occluding structure (column 2) in the hidden area. The mesh representation is then translated for the reconstruction of the light-emitting objects (column 4) as a 2D RGB image.

5.1 Qualitative Results

Real experimental examples of the *soft shadow diffusion* are presented in Figures 1 and 4, and the real experimental examples of the *gradient-based* method are shown in Figure 5. Figures 4 and 5 depicts the non-occluding portion of the hidden scene on a monitor at a distance of 1.08 meters to the visible wall, and 3D occluders placed between them. Measurements were obtained by photographing the visible wall with a 4-megapixel camera. The measurements were inputs to our physics-inspired neural network inversion method or Algorithm 1 for the gradient-based inversion method. In the final stage of SSD, we reconstruct each channel of the 2D non-occluding portion using TV regularization. The reconstructions show good qualitative agreement with the ground truth





Fig. 5: Comparing reconstructions of SSD and gradient-based optimization. SSD is trained on 3D objects and outputs a (3D) four-legged table for the planner occluder. The gradient-based reconstructions use $10 \times 5 \times 10$ discretization of the occluding region.

hidden scenes. However, the gradient-based method does not scale well for more complex occluder shapes requiring finer discretizations (as shown in the supplementary document).

Even when the non-occluding hidden scene portion reflects light (rather than being self-luminous) and spans a large depth (instead of being confined to a plane), the reconstructed 3D image of the occluder and 2D image of the nonoccluders are remarkably accurate as shown in Figure 1. This experiment further highlights the surprising generalizability of the proposed SSD-based method, which was trained using only simulated data.

Additional qualitative results with real experiments and simulations are given in the supplement.



Fig. 6: Robustness of SSD-based approach to background illumination. The insights show the measured photograph and reconstructions at each background level. With the monitor brightness fixed, a larger illuminance value indicates increased ambient light source contributions in (b).

5.2 Quantitative Results

Here, we investigate the robustness of our approach to increasing levels of background illumination (Figure 6), while reconstructions obtained under varying levels of measurement noise are presented in the supplementary document (Supplementary Section 4.1). In the real experimental result, we introduce an additional external light source that increases the overall brightness of the captured photograph but does not cast penumbra. We report the mean squared error (MSE) for the 2D reconstructions, and Chamfer Distance (CD) for the 3D occluding object reconstructions. Our approach degrades slowly for reasonable SBRs in simulated experiments, with the 3D reconstructions remaining accurate even in SBRs of around 10 dB, 2D reconstructions degrade gradually for SBRs below 20 dB. Similar trends are observed for real data reconstructions. The surprising stability of the 3D reconstructions even at very low SBRs demonstrates the utility of the strong shape prior imposed by the proposed SSD network.

SSD demonstrates similar robustness to measurement noise, with 3D reconstruction quality relatively unaffected even at signal-to-noise ratios as low as 10 dB. Additional experiments, shown in Supplementary Section 4.2, suggest that SSD also generalizes well to reconstructing unseen 3D shape classes.

6 Discussion & Conclusion

We introduced a formulation of computational periscopy that enables 3D reconstructions. Our formulation inspired a physics-based algorithm, and a physicsinspired neural network for computing 3D reconstructions of light-occluders, in addition to 2D reconstructions of the light-reflecting structures in the hidden scene. This demonstrates an advance over prior techniques in occluder-aided passive NLOS imaging that relied on *a priori* known information about the lightoccluders in the scene. Central to our physics-inspired neural network model is a novel soft shadow diffusion (SSD) network, which reconstructs point clouds from soft shadows and extends denoising diffusion models to interpret soft shadows. Although SSD was trained in simulations, both approaches were successful in real experiments. SSD was further shown to degrade gracefully with noise and additional background.

Notwithstanding, some aspects of our work warrant further investigation: First, our gradient-based method is limited to reconstructing simple shapes due to the computational demand of higher discretization. Although a sparse matrix approach mitigated this limitation in simulations, it was less effective in realworld scenarios (as shown in the supplementary material). Second, our SSD model predicts a view-consistent 3D point cloud without information about the true location of the points in the NLOS configuration. To facilitate recovery of the non-occluding object, we partition the occluder space into multiple 3D bounding boxes of the occluder size. The subsequent step involves a grid search over these bounding boxes to position the occluder accurately within the hidden area. However, the size of the occluding object is typically unknown in real-world uncontrolled scenarios. These shortcomings are avenues for future research.

References

- 1. Adib, F., Katabi, D.: See through walls with wifi! In: Proc. ACM SIGCOMM. p. 75–86 (2013) 2
- Aittala, M., Sharma, P., Murmann, L., Yedidia, A., Wornell, G., Freeman, B., Durand, F.: Computational mirrors: Blind inverse light transport by deep matrix factorization. Advances in Neural Information Processing Systems (NeurIPS) 32 (2019) 5
- Baradad, M., Ye, V., Yedidia, A.B., Durand, F., Freeman, W.T., Wornell, G.W., Torralba, A.: Inferring light fields from shadows. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 6267–6275 (2018) 2, 3, 5, 11
- Boger-Lombard, J., Slobodkin, Y., Katz, O.: Towards passive non-line-of-sight acoustic localization around corners using uncontrolled random noise sources. Scientific Reports 13(1), 4952 (2023) 2
- Bouman, K.L., Ye, V., Yedidia, A.B., Durand, F., Wornell, G.W., Torralba, A., Freeman, W.T.: Turning corners into cameras: Principles and methods. In: Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV) (Oct 2017) 2
- Cai, R., Lai, H., Chami, I., Guibas, L.J.: Learning gradient fields for shape generation. In: Proc. Eur. Conf. Computer Vision (ECCV). pp. 10751–10760 (2020)
 4
- 7. Cavanagh, P., Leclerc, Y.G.: Shape from shadows. Journal of Experimental Psychology: Human Perception and Performance 15(1), 3 (1989) 3
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository (2015) 10, 11
- Chaudhury, A.N., Keselman, L., Atkeson, C.G.: Shape from shading for robotic manipulation. In: Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV). pp. 8389–8398. IEEE (2024) 4
- Chou, G., Bahat, Y., Heide, F.: Diffusion-sdf: Conditional generative modeling of signed distance functions. In: Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV) (2023) 11
- 11. Cohen, A.L.: Anti-pinhole imaging. Optica Acta 29(1), 63-67 (1982) 6
- Community, B.O.: Blender a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), http://www.blender.org 10
- Czajkowski, R., Murray-Bruce, J.: Two-edge-resolved three-dimensional non-lineof-sight imaging with an ordinary camera. Nat. Commun. 15(1162) (2024) 2
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) 4
- Faccio, D., Velten, A., Wetzstein, G.: Non-line-of-sight imaging. Nature Reviews Physics 2(6), 318–327 (2020) 2
- Geng, R., Hu, Y., Lu, Z., Yu, C., Li, H., Zhang, H., Chen, Y.: Passive non-line-ofsight imaging using optimal transport. IEEE Trans. Image Process. 31, 110–124 (2022) 2, 11
- Heide, F., O'Toole, M., Zang, K., Lindell, D.B., Diamond, S., Wetzstein, G.: Nonline-of-sight imaging with partial occluders and surface normals. ACM Trans. Graphics (2019) 2
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 33, pp. 6840–6851 (2020)

- 16 Raji and Murray-Bruce
- Horn, B.K.P., Brooks, M.J. (eds.): Shape from shading. MIT Press, Cambridge, MA, USA (1989) 3
- Iwashita, Y., Stoica, A., Kurazume, R.: Gait identification using shadow biometrics. Pattern Recognition Letters 33(16), 2148–2155 (2012) 3
- Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR) (December 2021) 4
- 22. Kaga, M., Kushida, T., Takatani, T., Tanaka, K., Funatomi, T., Mukaigawa, Y.: Thermal non-line-of-sight imaging from specular and diffuse reflections. IPSJ Transactions on Computer Vision and Applications 11 (12 2019) 2
- Karnieli, A., Fried, O., Hel-Or, Y.: Deepshadow: Neural shape from shadow. In: Proc. Eur. Conf. Computer Vision (ECCV). p. 415–430 (2022) 4
- Khalid, N.M., Xie, T., Belilovsky, E., Tiberiu, P.: Clip-mesh: Generating textured meshes from text using pretrained image-text models. In: SIGGRAPH Asia 2022 Conference Papers (December 2022) 4
- Kirmani, A., Jeelani, H., Montazerhodjat, V., Goyal, V.K.: Diffuse imaging: Creating optical images with unfocused time-resolved illumination and sensing. IEEE Signal Processing Letters 19(1), 31–34 (2012) 2
- Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. In: Proc. Int. Conf. on Learning Representations (ICLR) (2021) 4
- gil Lee, S., Kim, H., Shin, C., Tan, X., Liu, C., Meng, Q., Qin, T., Chen, W., Yoon, S.H., Liu, T.Y.: Priorgrad: Improving conditional denoising diffusion models with data-dependent adaptive prior. In: Proc. Int. Conf. on Learning Representations (ICLR) (2022) 4
- Lindell, D.B., Wetzstein, G., Koltun, V.: Acoustic non-line-of-sight imaging. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 6773–6782 (2019) 2
- Lindell, D.B., Wetzstein, G., O'Toole, M.: Wave-based non-line-of-sight imaging using fast fk migration. ACM Trans. Graphics 38(4), 1–13 (2019) 2
- Ling, J., Wang, Z., Xu, F.: Shadowneus: Neural sdf reconstruction by shadow ray supervision. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 175–185 (2023)
- Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. arXiv preprint arXiv:2309.16653 (2023) 4
- Liu, X., Bauer, S., Velten, A.: Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems. Nat. Commun. 11(1), 1645 (2020) 2
- 33. Liu, X., Guillén, I., La Manna, M., Nam, J.H., Reza, S.A., Huu Le, T., Jarabo, A., Gutierrez, D., Velten, A.: Non-line-of-sight imaging using phasor-field virtual wave optics. Nature 572(7771), 620–623 (2019) 2
- Liu, X., Wang, J., Xiao, L., Fu, X., Qiu, L., Shi, Z.: Few-shot non-line-of-sight imaging with signal-surface collaborative regularization. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 13303–13312 (2022) 2
- Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: SIGGRAPH. pp. 163–169. ACM (1987) 11
- Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR) (2021)
 7

- Maeda, T., Wang, Y., Raskar, R., Kadambi, A.: Thermal non-line-of-sight imaging. In: Proc. IEEE Int. Conf. Computational Photography (ICCP). p. 1–11. IEEE (2019) 2
- Medin, S.C., Weiss, A., Durand, F., Freeman, W.T., Wornell, G.W.: Can shadows reveal biometric information? In: Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV). pp. 869–879 (January 2023) 3
- Metzler, C.A., Heide, F., Rangarajan, P., Balaji, M.M., Viswanath, A., Veeraraghavan, A., Baraniuk, R.G.: Deep-inverse correlography: towards real-time high-resolution non-line-of-sight imaging. Optica 7(1), 63–71 (2020) 2
- Murray-Bruce, J., Saunders, C., Goyal, V.K.: Occlusion-based computational periscopy with consumer cameras. In: Wavelets and Sparsity XVIII. vol. 11138, pp. 286–297. SPIE (2019) 2
- Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts (2022), https://arxiv.org/ abs/2212.08751 11
- O'Toole, M., Lindell, D.B., Wetzstein, G.: Confocal non-line-of-sight imaging based on the light-cone transform. Nature 555(7696), 338–341 (2018) 2
- Pawlikowska, A.M., Halimi, A., Lamb, R.A., Buller, G.S.: Single-photon threedimensional imaging at up to 10 kilometers range. Optics express 25(10), 11919– 11931 (2017) 2
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2309.16653 (2022) 4
- 45. Popov, V., Vovk, I., Gogoryan, V., Sadekova, T., Kudinov, M.: Grad-tts: A diffusion probabilistic model for text-to-speech. In: ICML (2021) 4
- Ramachandran, V.S.: Perception of shape from shading. Nature 331, 163–166 (1988) 3
- 47. Rapp, J., Saunders, C., Tachella, J., Murray-Bruce, J., Altmann, Y., Tourneret, J.Y., McLaughlin, S., Dawson, R.M., Wong, F.N., Goyal, V.K.: Seeing around corners with edge-resolved transient imaging. Nat. Commun. 11(1), 5929 (2020) 2
- Rasul, K., Seward, C., Schuster, I., Vollgraf, R.: Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In: Proc. Int. Conf. on Machine Learning (ICML) (2021) 4
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 10684–10695 (June 2022) 4
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 234–241. Lecture Notes in Computer Science, Springer International Publishing (2015) 11
- Saunders, C., Bose, R., Murray-Bruce, J., Goyal, V.K.: Multi-depth computational periscopy with an ordinary camera. In: Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP). pp. 9299–9305 (2020) 2, 3, 5
- 52. Saunders, C., Murray-Bruce, J., Goyal, V.K.: Computational periscopy with an ordinary digital camera. Nature 565(7740), 472–475 (2019) 2, 3, 4, 5, 6, 10, 11
- 53. Seidel, S., Rueda-Chacón, H., Cusini, I., Villa, F., Zappa, F., Yu, C., Goyal, V.K.: Non-line-of-sight snapshots and background mapping with an active corner camera. Nat. Commun. **14**(1), 3677 (2023) 2
- 54. Seidel, S.W., Ma, Y., Murray-Bruce, J., Saunders, C., Freeman, W.T., Christopher, C.Y., Goyal, V.K.: Corner occluder computational periscopy: Estimating a

hidden scene from a single photograph. In: Proc. IEEE Int. Conf. Computational Photography (ICCP). pp. 1–9. IEEE (2019) 2

- Seidel, S.W., Murray-Bruce, J., Ma, Y., Yu, C., Freeman, W.T., Goyal, V.K.: Twodimensional non-line-of-sight scene estimation from a single edge occluder. IEEE Trans. Comput. Imaging 7, 58–72 (2021) 2
- 56. Sharma, P., Aittala, M., Schechner, Y.Y., Torralba, A., Wornell, G.W., Freeman, W.T., Durand, F.: What you can learn by staring at a blank wall. In: Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV). pp. 2330–2339 (October 2021) 3
- Shim, J., Kang, C., Joo, K.: Diffusion-based signed distance fields for 3d shape generation. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 20887–20897 (June 2023) 7
- Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Proc. Int. Conf. on Machine Learning (ICML). p. 2256–2265 (2015) 4, 7
- Somasundaram, S., Dave, A., Henley, C., Veeraraghavan, A., Raskar, R.: Role of transients in two-bounce non-line-of-sight imaging. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 9192–9201 (2023) 2
- 60. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. Advances in neural information processing systems 32 (2019) 4, 8
- 61. Torralba, A., Freeman, W.T.: Accidental pinhole and pinspeck cameras: Revealing the scene outside the picture. Int. J. Computer Vision **110**(2), 92–112 (2014) **6**
- Velten, A., Willwacher, T., Gupta, O., Veeraraghavan, A., Bawendi, M.G., Raskar, R.: Recovering three-dimensional shape around a corner using ultrafast time-offlight imaging. Nat. Commun. 3(1), 745 (2012) 2
- 63. Verlekar, T.T., Soares, L.D., Correia, P.L.: Gait recognition in the wild using shadow silhouettes. Image and Vision Computing **76**, 1–13 (2018) **3**
- 64. Wang, W., Yao, L., Chen, L., Lin, B., Cai, D., He, X., Liu, W.: Crossformer: A versatile vision transformer hinging on cross-scale attention. In: Proc. Int. Conf. on Learning Representations (ICLR) (2022) 11
- Wang, Y., Zhang, Y., Huang, M., Chen, Z., Jia, Y., Weng, Y., Xiao, L., Xiang, X.: Accurate but fragile passive non-line-of-sight recognition. Communications Physics 4(1), 88 (Dec 2021) 2, 11
- Yamashita, Y., Sakaue, F., Sato, J.: Recovering 3d shape and light source positions from non-planar shadows. In: Proc. IEEE/CVF Int. Conf. Pattern Recognition (ICPR). pp. 1775–1779 (2010) 4
- Yang, B., Gu, S., Zhang, B., Zhang, T., Chen, X., Sun, X., Chen, D., Wen, F.: Paint by example: Exemplar-based image editing with diffusion models. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 18381– 18391 (2023) 4
- Yang, G., Huang, Y., Hao, Z., Liu, M., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV) (2019) 4
- Ye, J., Wang, P., Li, K., Shi, Y., Wang, H.: Consistent-1-to-3: Consistent image to 3d view synthesis via geometry-aware diffusion models. arXiv preprint arXiv:2309.16653 (2024) 4
- Yedidia, A.B., Baradad, M., Thrampoulidis, C., Freeman, W.T., Wornell, G.W.: Using unknown occluders to recover hidden scenes. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 12231–12239 (2019) 2, 3, 5, 11

- Yedidia, A.B., Baradad, M., Thrampoulidis, C., Freeman, W.T., Wornell, G.W.: Using unknown occluders to recover hidden scenes. In: Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR). pp. 12231–12239 (2019) 11
- 72. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. In: Advances in Neural Information Processing Systems (NeurIPS) (2022) 4
- 73. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 10021–10039 (2022) 7
- Zhou, L., Du, Y., Wu, J.: 3d shape generation and completion through point-voxel diffusion. In: Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV). pp. 5806–5815 (2021) 4, 11