

# Non-Line-of-Sight Reconstruction using Efficient Transient Rendering

JULIAN ISERINGHAUSEN and MATTHIAS B. HULLIN, University of Bonn, Germany

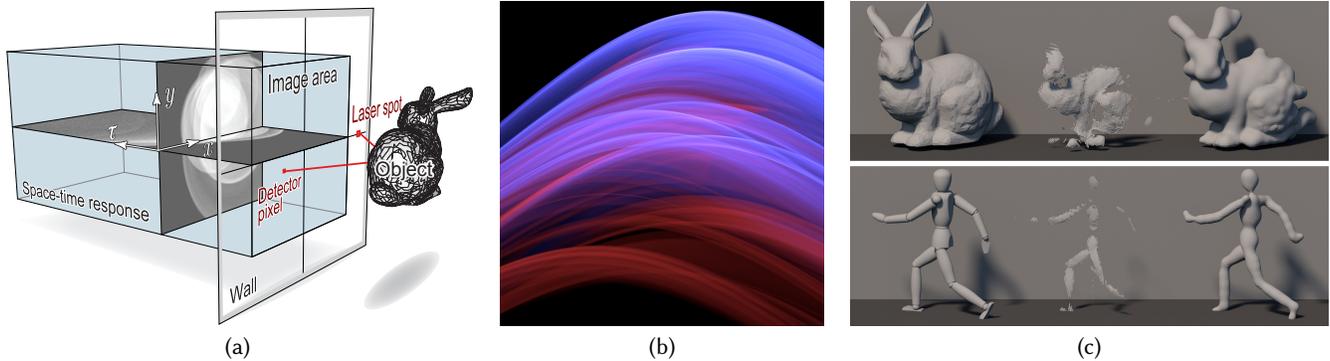


Fig. 1. (a) The challenge of looking around the corner deals with the recovery of information about objects beyond the direct line of sight. In this illustration of a setting proposed by Velten et al. [2012], an unknown object is located in front of a wall, but additional obstacles occlude the object from any optical devices like light sources or cameras. Our only source of information are therefore indirect reflections off other surfaces (here, a planar “wall”). A point on the wall that is illuminated by an ultrashort laser pulse turns into an omnidirectional source of indirect light (“laser spot”). After scattering off the unknown object, some of that light arrives back at the wall, where it forms an optical “echo” or space-time response (shown are 2D slices) that can be picked up by a suitable camera. Locations on the wall can be interpreted as omnidirectional detector pixels that receive different mixtures of backscattered light contributions at different times. We assume that neither camera nor laser can directly illuminate or observe the object, leaving us with the indirect optical space-time response as the only source of information. Note that for the sake of clarity, laser source, camera, and occluder are not shown here. The complete setup is illustrated in Figure 2. (b) We propose a novel transient renderer to simulate such indirectly scattered light transport efficiently enough for use as a forward model in inverse problems. In this artistic visualization, light contributions removed by the shadow test are marked in red, and the net intensity in blue. Together with an optimization algorithm, the renderer can be used to reconstruct the geometry of objects outside the line of sight. (c) Left to right: ground-truth object geometry; reconstruction using a state-of-the-art method (ellipsoidal backprojection); reconstruction using the technique presented in this paper. Top row: BunnyGI dataset; bottom row: Mannequin1Laser dataset. Our method relies on highly efficient and near-physical forward simulation, and it exemplifies the use of computer graphics as a technical tool to solve inverse problems in other fields.

Being able to see beyond the direct line of sight is an intriguing prospective and could benefit a wide variety of important applications. Recent work has demonstrated that time-resolved measurements of indirect diffuse light contain valuable information for reconstructing shape and reflectance properties of objects located around a corner. In this paper, we introduce a novel reconstruction scheme that, by design, produces solutions that are consistent with state-of-the-art physically-based rendering. Our method combines an efficient forward model (a custom renderer for time-resolved three-bounce indirect light transport) with an optimization framework to reconstruct object geometry in an analysis-by-synthesis sense. We evaluate our algorithm on a variety of synthetic and experimental input data, and show that it gracefully handles uncooperative scenes with high levels of noise or non-diffuse material reflectance.

This work was supported by the German Research Foundation (DFG) under grant (HU-2273/2-1), the X-Rite Chair for Digital Material Appearance, and the European Research Council under ERC Starting Grant ECHO. We thank Hendrik Lensch, Michael Wand, Johannes Hanika, Felix Heide, Wolfgang Heidrich, Ivo Ihrke and Ramesh Raskar for valuable discussion over the course of this project. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan Xp GPU that was used in this research.

Authors’ address: Julian Iseringhausen, iseringhausen@cs.uni-bonn.de; Matthias B. Hullin, hullin@cs.uni-bonn.de, University of Bonn, Institute of Computer Science II, Bonn, 53115, Germany.

© 2020 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3368314>.

CCS Concepts: • **Computing methodologies** → **Computational photography**; **3D imaging**; **Reconstruction**; **Physical simulation**.

**ACM Reference Format:**

Julian Iseringhausen and Matthias B. Hullin. 2020. Non-Line-of-Sight Reconstruction using Efficient Transient Rendering. *ACM Trans. Graph.* 39, 1, Article 8 (February 2020), 14 pages. <https://doi.org/10.1145/3368314>

## 1 MOTIVATION

Every imaging modality from ultrasound to x-ray knows situations where the target is partially or entirely occluded by other objects and therefore cannot be directly observed. In a recent strand of work, researchers have aimed to overcome this limitation, developing a variety of approaches to extend the line of sight of imaging systems, for instance using wave optics [Boger-Lombard and Katz 2018; Katz et al. 2014] or by using the occluder itself as an accidental imager [Bouman et al. 2017]. Among all the techniques proposed, a class of methods has received particular attention within the computer vision and imaging communities. The main source of information for these methods are indirect reflections of light within the scene, represented by time-resolved impulse responses. From such responses, it has been shown that the presence and position of objects “around a corner” [Kirmani et al. 2009], or even their shape [Velten et al. 2012] and/or reflectance [Naik et al. 2011] can be reconstructed. In this paper, we focus on the archetypal challenge of reconstructing

the shape of an unknown object from 3-bounce indirect and (more or less) diffuse reflections off a planar wall (Figure 1(a)) [Kirmani et al. 2009]. The overwhelming majority of approaches to this class of problem rely on ellipsoidal *backprojection*, where intensity measurements are smeared out over the loci in space (ellipsoidal shells) that correspond to plausible scattering locations under the given geometric constraints [Arellano et al. 2017; Buttafava et al. 2015; Garipey et al. 2016; Kadambi et al. 2016; Velten et al. 2012]. Ellipsoidal backprojection implicitly assumes that the object is a volumetric scatterer, and it does not take into account surface orientation and self-occlusion of the object. More importantly, unlike linear backprojection used in standard emission or absorption tomography, ellipsoidal backprojection is not the adjoint of a physically plausible forward light transport operator. Where such operators have been identified [La Manna et al. 2018], they are typically constrained to rudimentary volumetric, non-opaque, isotropic scattering models. This necessitates filtering, and the reconstructed shapes are typically flat and low in detail. On the other hand, algorithms based on ellipsoidal backprojection generally have much shorter runtimes than our approach, since they do not require a global optimization scheme.

Here, we propose an alternative approach that mitigates some of the problems of backprojection by formulating the non-line-of-sight sensing problem in an analysis-by-synthesis sense. In other words, we develop a physically plausible and efficient forward simulation of light transport (transient renderer) and combine it with a nonlinear optimizer to determine the scene hypothesis that best agrees with the observed data. The method is enabled by a number of technical innovations, which we consider the key contributions of this work:

- a scene representation based on level sets and a surface-oriented scattering model for time-resolved light transport around a corner (wall to object to wall) based on time-resolved radiative transfer,
- an extremely efficient GPU-based custom renderer for three-bounce backscatter that features near-physical handling of occlusion effects and a novel temporal filtering scheme for triangular surfaces, and
- a global, self-refining optimization strategy to minimize the reconstruction error.

We evaluate our method on a number of synthetic and experimental datasets and find that it is capable of achieving significantly higher object coverage and detail than ellipsoidal backprojection, even on greatly reduced and degraded input data. Our renderer not only naturally accommodates surface BRDFs, but is also open to extensions like higher-order light bounces or advanced background models that will be needed in order to tackle future non-line-of-sight sensing problems. The method, as proposed here, is not capable of delivering high reconstruction rates in this first implementation. However, we believe that being able to generate transient renderings for the around-the-corner setting very efficiently will enable novel approaches to the problem, for instance based on machine learning.

## 2 RELATED WORK

The research areas of transient imaging and non-line-of-sight reconstruction have recently received tremendous attention from the computer vision, graphics, imaging and optics communities. For a structured overview on the state of the art, we refer the interested reader to a recent survey [Jarabo et al. 2017].

### 2.1 Transient imaging

Imaging light itself as it propagates through space and time poses the ultimate challenge to any imaging system. To obtain an idea of the frame rate required, consider that in vacuum, light only takes about 3 picoseconds ( $3 \cdot 10^{-12}$  s) per millimeter of distance traversed. The typical transient imaging system consists of an ultrashort (typically, sub-picosecond) light source and an ultrafast detector. Oddly, three of the highest-performing detection technologies are over 40 years old: streak tubes [Velten et al. 2011] wherein a single image scanline is “smeared out” over time on a phosphor screen; holography using ultrashort pulses [Abramson 1978], and gated image intensifiers [Laurenzis and Velten 2014b]. More common nowadays, however, are semiconductor devices that achieve comparable temporal resolution without the need for extreme light intensities or voltages. Among the technologies reported in literature are regular reverse-biased photodiodes [Kirmani et al. 2009], as well as time-correlated single-photon counters which conveniently map to standard CMOS technology [Garipey et al. 2015]. On the low end, it has also been shown that transient images can be computationally reconstructed from multi-frequency correlation time-of-flight measurements [Heide et al. 2013], although data thus obtained typically suffers from the low temporal bandwidth of these devices, which necessitates heavy regularization.

### 2.2 Transient rendering

The simulation of transient light transport, when done naïvely, is no different from regular physically-based rendering, except that for each light path that contributes to the image, its optical length must be calculated and its contribution stored in a time-of-flight histogram [Smith et al. 2008]. A number of offline transient renderers have been made available to the public [Jarabo et al. 2014; Slaney and Chou 2014]. Even with advanced temporal sampling [Jarabo et al. 2014] and efficiency-increasing filtering strategies such as photon beams [Marco et al. 2017], such renderers still take on the order of hours to days to produce converged results. In contrast, the special-purpose renderer introduced in this paper is capable of producing close-to-physical renderings of around-the-corner settings in a matter of milliseconds. Finally, there have been efforts to simulate the particular characteristics of single-photon counters [Hernandez et al. 2017], an emerging type of sensor that can be expected to assume a major role in transient imaging.

### 2.3 Analysis of transient light transport and looking around corners

The information carried by transient images has been the subject of several investigations. Wu et al. laid out the geometry of space-time streak images for lensless imaging [2012], and discussed the influence of light transport phenomena such as subsurface scattering

on the shape of the temporal response [2014]. Economically, the most important use of transient light transport analysis today is likely in multi-path backscatter removal for correlation-based time-of-flight ranging [Fuchs 2010, and many others].

In this paper, we direct our main attention to the idea of exploiting time-resolved measurements of indirect reflections for the purpose of extending the direct line of sight and, in effect, looking around corners [Kirmani et al. 2009; Velten et al. 2012]. While a variety of geometric settings have been investigated, the bulk of work in this area relies on the arrangement illustrated in Figure 1(c) and Figure 2 and further introduced in the following Section 3.

The reconstruction strategies can be roughly grouped in two classes. One major group is formed by backprojection approaches where each input measurement casts votes on those locations in the scene where the light could have been scattered [Arellano et al. 2017; Buttafava et al. 2015; Garipey et al. 2016; Kadambi et al. 2016; Laurenzis and Velten 2014b; Velten et al. 2012]. A smaller but more diverse group of work relies on the use of forward models to arrive at a scene hypothesis that best agrees with the measured data. Here, reported approaches fall into several categories. A combinatorial labeling scheme was developed by Kirmani et al. [2009]. If the capture geometry is sufficiently constrained, frequency-domain inverse filtering [O’Toole et al. 2018a] can be employed. Variational methods using simple linearized light transport tensors [Heide et al. 2014; Naik et al. 2011] and simplistic models based on radiative transfer [Klein et al. 2016; Pediredla et al. 2017] are (in principle) capable of expressing opacity effects like shadowing and occlusion, and physically plausible shading. These approaches are closest to our proposed method. In concurrent work, Heide et al. [2017] added such extra factors as additional weights into their least-squares data term, achieving non-line-of-sight reconstructions of significantly improved robustness. Thrampoulidis et al. [2017] applied a similar idea on the reconstruction of 2D albedo maps on known geometry that are further obscured by known occluders between object and wall. For homogeneous volumetric media in direct sight, Gkioulekas et al. [2013] extensively relied on physically-based rendering to recover their scattering parameters and phase function. With the proposed method, we demonstrate what we believe is the first reconstruction scheme for non-line-of-sight object geometry that is based on a near-physical yet extremely efficient special-purpose renderer and, by design, produces solutions that are self-consistent. We believe that our work can serve as an example for other uses of computer graphics methodology as a technical tool for solving inverse problems in imaging and vision.

### 3 PROBLEM STATEMENT

Here we introduce the geometry of the non-line-of-sight reconstruction problem as used in the remainder of the paper. For simplicity, we neglect the constant factor  $c$  (the speed of light) connecting *time* and (*optical*) *path length*. Thus, time and distance can be used synonymously and all discussions become independent of the absolute scale.

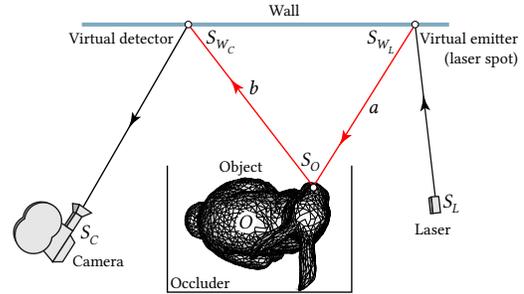


Fig. 2. Schematic top view of the scene arrangement, where the unknown object is occluded from direct observation. We assume that the temporal response has been “unwarped” (e.g., [Kadambi et al. 2016]), so only the occluded segments  $a$  and  $b$  contribute to the total time of flight and to the shading in Equation 4.

#### 3.1 Problem geometry and transient images

We model our setting after the most common scenario from literature (Figure 2), where the unknown object is observed indirectly by illuminating a wall with a laser beam and measuring light reflected back to the wall. Following Kadambi et al. [2016], the laser spot on the wall acts as an area light source, and observed locations on the wall are equivalent to omnidirectional detectors that produce an “unwarped” transient image [Velten et al. 2013] (Figure 1(a)). The extent of the observed wall, the size of the object and its distance to the wall are usually on the same order of magnitude. The *transient image* or *space-time response*  $\mathbf{I} \in \mathbb{R}^{n_x \times n_\tau}$  is the entirety of measurements taken using this setup,  $n_x$  being the number of combinations of detector pixels and illuminated spots and  $n_\tau$  the number of bins in a time-of-flight histogram recorded per location. For a two-dimensional array of observed locations (for instance, when using a time-gated imager), the space-time response can be interpreted as a three-dimensional data cube similar to a video.

#### 3.2 Problem formulation

The idea underlying ellipsoidal backprojection is that any entry in the transient image, or the response of a pair of emitter and detector positions for a given travel time, corresponds to an ellipsoidal locus of possible candidate scattering locations. If no further information is available, any measured quantity of light therefore “votes” for all locations on its ellipsoid. Finally, the sum or product of all such votes is interpreted as occupancy measure, or probability of there being an object at any point in space. We refer to a recent study [La Manna et al. 2018] that discusses the design options for such algorithms in great detail.

In contrast, we formulate the reconstruction task as a non-linear least-squares minimization problem

$$\min_{\mathbf{P}} \|\mathbf{I}_{\text{ref}} - \mathbf{I}(G(\mathbf{P}))\|_2^2, \quad (1)$$

where  $\mathbf{P}$  is a parameter vector describing the scene geometry,  $G(\cdot)$  is a function that generates explicit scene geometry (a triangle mesh),  $\mathbf{I}_{\text{ref}}$  is the measured space-time scene response, and  $\mathbf{I}(\cdot)$  is a forward model (renderer) that predicts the response under the scene hypothesis passed as argument. The purpose of the optimization is

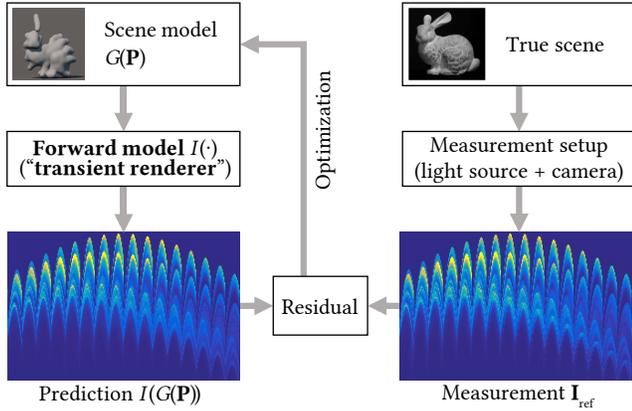


Fig. 3. Overview of our analysis-by-synthesis scheme for looking around a corner. Our pipeline heavily relies on custom-made components (scene representation, renderer, residual function, optimizer) to make this approach viable.

to find the scene geometry  $G(\mathbf{P})$  that minimizes the sum of squared pixel differences between the predicted and the observed responses. Figure 3 illustrates this principle.

A key feature of this formulation is that the solution by its very definition is optimally consistent with the chosen physical model of light transport, and that ongoing improvements in forward modeling will also benefit the reconstruction. Furthermore, our approach naturally handles opaque, oriented surfaces, whereas in backprojection, surface geometry is implicitly defined and needs to be derived using additional filtering steps. Furthermore, our method is able to handle arbitrary surface BRDFs, where current backprojection methods implicitly assume diffuse cloud-like scattering [La Manna et al. 2018]. A downside of our approach is that it requires a full model of the scene, and that any unknowns (such as background or noise) can distort the solution in ways that are hard to predict. On the other hand, we believe that our approach lends itself for future extensions like higher-order light bounces.

## 4 METHOD

In the following, we introduce the components of our reconstruction algorithm in detail.

### 4.1 Geometry representation

We seek to parameterize the scene geometry in terms of a vector  $\mathbf{P}$  that has a small number of degrees of freedom to make the optimization in Equation 1 tractable. Rather than using  $\mathbf{P}$  to directly store a mesh representation with vertices and faces, we express the geometry implicitly as an isosurface of a scalar field  $M_{\mathbf{P}}(\mathbf{x})$  composed of globally supported basis functions. This approach is also common in surface reconstruction from point clouds [Carr et al. 2001]. In our case, the vector  $\mathbf{P}$ ,

$$\begin{aligned} \mathbf{P} &= (\mathbf{p}_1, \dots, \mathbf{p}_m) \\ &= ((\mathbf{x}_1, \sigma_1), \dots, (\mathbf{x}_m, \sigma_m)), \end{aligned} \quad (2)$$

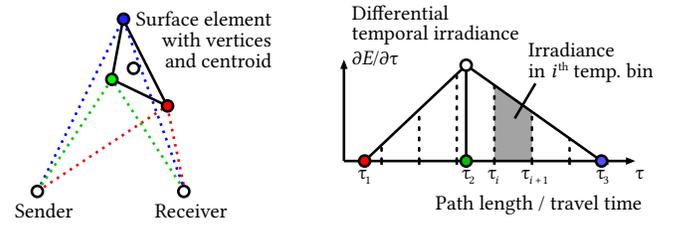


Fig. 4. To compute the total irradiance  $\alpha_t$  contributed by a surface triangle to a given detector pixel, we evaluate the radiative transfer using the element’s centroid. We then use a first-order filter to distribute this irradiance over the temporal bins that are affected by the triangle. To this end, we compute the three optical path lengths, or travel times,  $\tau_1, \dots, \tau_3$  belonging to the triangle’s three vertices. The irradiance ending up in any temporal bin is then obtained by constructing a triangular function of total area  $\alpha_t$  using the three arrival times as illustrated, then geometrically integrating over the time interval that corresponds to the bin. The true temporal distribution depends on the position and orientation of the triangle. However, the effectiveness of the temporal filter can be seen in Table 1 and Figures 5 and 8.

lists the centers  $\mathbf{x}_i$  and standard deviations  $\sigma_i$  of  $m$  isotropic Gaussian blobs. From the scalar field

$$M_{\mathbf{P}}(\mathbf{x}) = \sum_{i=1}^m e^{-\|\mathbf{x}-\mathbf{x}_i\|_2^2 / (2\sigma_i^2)} \quad (3)$$

we extract the triangle mesh  $G(\mathbf{P})$  using a GPU implementation of Marching Cubes [Lorensen and Cline 1987]. For all our reconstructions, we used a fixed resolution of  $128^3$  voxels for the reconstruction volume, and a fixed threshold of  $\frac{3}{4}$  for the isosurface. The extension to other implicit functions, such as anisotropic Gaussians or general radial basis functions, is trivial.

### 4.2 Rendering (synthesis)

We propose a custom renderer that is suitable for use as forward model  $I(\cdot)$  inside the objective function, Equation 1. In order to be suited for this purpose, the renderer must be sufficiently close to physical reality. At the same time, it has to be very efficient because hundreds of thousands of renderings may be required over the course of the optimization run. We achieve this efficiency by restricting the renderer to a single type of light path and rendering only light bounces from the wall to the object and back to the wall. Following the notation of [Pharr and Humphreys 2010] and by dropping any constant terms, we can write the incoming radiance for each camera pixel as

$$L = \int_O f(S_{W_L} \rightarrow S_O \rightarrow S_{W_C}) \eta(S_O \leftrightarrow S_{W_C}) \eta(S_{W_L} \leftrightarrow S_O) dS_O, \quad (4)$$

where  $O = G(\mathbf{P})$  denotes the object,  $f$  the object’s BRDF and  $S_{\_}$  surface points as shown in Figure 2. The geometric coupling term  $\eta$  is defined as

$$\eta(S_1 \leftrightarrow S_2) = V(S_1 \leftrightarrow S_2) \frac{|\cos(\theta_1)| |\cos(\theta_2)|}{\|S_1 - S_2\|_2^2}, \quad (5)$$

with  $V$  being the binary visibility function and  $\theta_i$  the angle of the ray connecting  $S_1$  and  $S_2$  to the respective surface normal. Since our object is already represented as a triangle mesh, we are able to

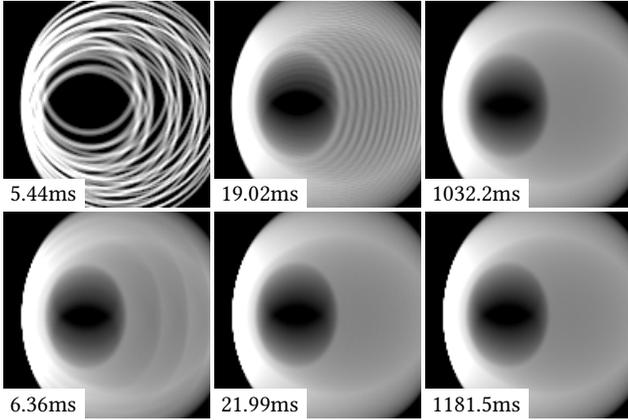


Fig. 5. The temporal filter also results in overall smoother spatial slices of the space-time response. Here we verify the performance of the filter by rendering the response generated by a planar square using different levels of detail. Shown is a single time slice without (top row) and with temporal filtering (bottom row). From left to right: coarse tessellation ( $4 \times 4$  quads), medium tessellation ( $16 \times 16$  quads), fine tessellation ( $128 \times 128$  quads). Numbers indicate the rendering time for the entire transient data cube ( $128 \times 128$  pixels, 192 time bins) on an NVIDIA GTX 980. Note the significant quality improvement at only 14–17% increased computational cost.

approximate Equation 4 by assuming a constant radiance over each triangles' surface,

$$L \approx \sum_{t \in T} f(S_{W_L} \rightarrow S_t \rightarrow S_{W_C}) \eta(S_t \leftrightarrow S_{W_C}) \eta(S_{W_L} \leftrightarrow S_t) A_t \quad (6)$$

$$=: \sum_{t \in T} \alpha_t.$$

Here,  $T$  is the set of all triangles of our object,  $P_t$  is the centroid, and  $A_t$  the area. We denote the total irradiance contributed by triangle  $t$  as  $\alpha_t$ . In our experiments, we use Lambertian and metal BRDFs, but other reflectance functions can be used as well. This approximation can be seen as an extension of the one found in [Klein et al. 2016]. We further add two important features to increase physical realism and generate a smooth transient image.

Our first addition are visibility tests ( $V$ ) for both segments of the light path, which is necessary for handling non-convex objects. We first connect the laser point and the triangle centroid by a straight line, and test whether this segment intersects with any of the other triangles of the object mesh. For all visible triangles for which no intersection is found, we test the visibility of the second path segment (return of scattered light to the wall) in the same way. This shadow test avoids overestimation of backscatter from self-occluding object surfaces. We note, however, that our way of performing the test only for the triangle centroid leads to a binary decision (triangle entirely visible or entirely shadowed) and therefore potentially makes the objective non-continuous. This can be reduced by using a triangle grid of sufficiently high resolution.

To render a transient image, we extend the pixels of the steady-state renderer to record time-of-flight histograms. The light contribution  $\alpha_t$  enters into this histogram according to the geometric length of the corresponding light path; this length is simply the sum

of the two Euclidean distances from laser point to point on triangle and back to the receiving point on the wall (see Figure 2). We found that the temporal response is prone to artifacts if only the centroid of the triangle is taken into account for the path length. Instead, we use the path lengths for the triangle's three corner vertices to determine the temporal footprint of the surface element. Using a linear filter, we then distribute the contribution  $\alpha_t$  over the temporal domain (Figure 4). This procedure ensures that the rendered outcome is smooth in the temporal and spatial domains even when a single surface element covers dozens of temporal bins (Figure 5).

### 4.3 Optimization (analysis)

The optimization problem in Equation 1 is non-convex and non-linear, so special care has to be taken to find a solution (a set of blobs) that, when rendered, minimizes the cost function globally. While it would be desirable to optimize over the whole parameter vector  $\mathbf{P}$  simultaneously, this is computationally prohibitive. To address this problem, we developed the iterative optimization scheme summarized in Algorithm 1, with subroutines provided in Algorithms 2 and 3. Figure 6 shows several intermediate results during execution of the optimization scheme.

---

#### Algorithm 1 Global optimization scheme

---

**Input:** Reference image  $I_{\text{ref}}$ , Threshold  $c_{\text{thresh}}$

**Output:** Parameter vector  $\mathbf{P}$ , Cost  $c$

```

1:  $\mathbf{x} \leftarrow \text{SAMPLE}(\emptyset)$ 
2:  $\mathbf{P}, c \leftarrow \text{ADD\_BLOB}(\emptyset, \mathbf{x})$ 
3: while  $c > c_{\text{thresh}}$  do
4:    $\mathbf{x} \leftarrow \text{SAMPLE}(\mathbf{P})$ 
5:    $\mathbf{P}_1, c_1 \leftarrow \text{ADD\_BLOB}(\mathbf{P}, \mathbf{x})$ 
6:    $\mathbf{P}_2, c_2 \leftarrow \text{DUPLICATE\_BLOB}(\mathbf{P}, \mathbf{x})$ 
7:    $\mathbf{P}_3, c_3 \leftarrow \text{DELETE\_BLOB}(\mathbf{P}, \mathbf{x})$ 
8:    $i \leftarrow \arg \min_x c_x$ 
9:   if  $c_i < c$  then
10:     $\mathbf{P}, c \leftarrow \mathbf{P}_i, c_i$ 
11:   end if
12:    $\mathbf{P}_r, c_r \leftarrow \text{REITERATE}(\mathbf{P})$ 
13:   if  $c_r < c$  then
14:     $\mathbf{P}, c \leftarrow \mathbf{P}_r, c_r$ 
15:   end if
16:    $\mathbf{P}, c \leftarrow \text{CHECK\_DELETE}(\mathbf{P}, c)$ 
17: end while

```

---

The heart of our optimization algorithm is the inner optimization loop  $\text{ITERATE}(\mathbf{p}, \mathbf{P})$ , which determines the  $k = 10$  nearest neighbors of a given pivot blob  $\mathbf{p}$  using the routine  $\text{FIND\_NEIGHBORS}(\mathbf{p}, \mathbf{P})$ . It then optimizes the *positions* of those blobs using the Levenberg-Marquardt algorithm,  $\text{LEVENBERG\_MARQUARDT}(\mathbf{P})$  [Levenberg 1944; Marquardt 1963]. The function  $\text{SET\_VARIABLE}(\mathbf{x})$  is used to label these parameters as variable to the solver, while all other blobs are kept fixed during the optimization run using  $\text{SET\_FIXED}(\mathbf{x})$ . Derivatives for the Jacobian matrix are computed numerically using finite differences (by repeatedly executing our forward renderer with the perturbed parameter vector). In a subsequent step, the sizes of the selected blobs are also included in a second optimization run,

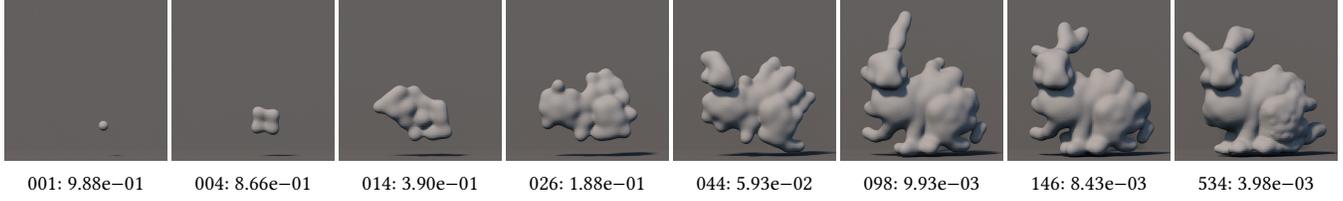


Fig. 6. Convergence of reconstructed geometry for the Bunny dataset over the course of the optimization. Number pairs denote iteration number and value of cost function (relative to start value).

---

**Algorithm 2** Inner optimization scheme
 

---

```

1: function ITERATE( $\mathbf{p}, \mathbf{P}$ )
2:    $\mathbf{P}_{\text{opt}} \leftarrow \text{FIND\_NEIGHBORS}(\mathbf{p}, \mathbf{P}, 10)$ 
3:   SET_FIXED( $\mathbf{P}$ )
4:   for all  $(\tilde{\mathbf{p}}, \tilde{\sigma}) \in \mathbf{P}_{\text{opt}}$  do
5:     SET_VARIABLE( $\tilde{\mathbf{p}}$ )
6:   end for
7:    $\mathbf{P} \leftarrow \text{LEVENBERG\_MARQUARDT}(\mathbf{P})$ 
8:   for all  $(\tilde{\mathbf{p}}, \tilde{\sigma}) \in \mathbf{P}_{\text{opt}}$  do
9:     SET_VARIABLE( $\tilde{\mathbf{p}}$ )
10:    SET_VARIABLE( $\tilde{\sigma}$ )
11:  end for
12:   $\mathbf{P} \leftarrow \text{LEVENBERG\_MARQUARDT}(\mathbf{P})$ 
13:   $c \leftarrow \text{COMPUTE\_COST}(\mathbf{P})$ 
14:  return  $\mathbf{P}, c$ 

```

---



---

**Algorithm 3** Subroutines to Algorithm 1.
 

---

```

1: function ADD_BLOB( $\mathbf{P}, \mathbf{x}$ )
2:    $\mathbf{p} \leftarrow (\mathbf{x}, \sigma_0)$ 
3:   return ITERATE( $\mathbf{p}, \mathbf{P} \cup \mathbf{p}$ )
1: function CHECK_DELETE( $\mathbf{P}$ )
2:   for all  $\mathbf{p} \in \mathbf{P}$  do
3:     if  $\text{COMPUTE\_COST}(\mathbf{P} \setminus \mathbf{p}) < \eta \cdot c$  then
4:        $\mathbf{P} \leftarrow \mathbf{P} \setminus \mathbf{p}$ 
5:     end if
6:   end for
7:    $c \leftarrow \text{COMPUTE\_COST}(\mathbf{P})$ 
8:   return  $\mathbf{P}, c$ 
1: function DUPLICATE_BLOB( $\mathbf{P}, \mathbf{x}$ )
2:    $\mathbf{p} \leftarrow \text{FIND\_NEAREST}(\mathbf{P}, \mathbf{x})$ 
3:    $\mathbf{p}_1, \mathbf{p}_2 \leftarrow \text{SPLIT}(\mathbf{p})$ 
4:   return ITERATE( $\mathbf{p}, \mathbf{P} \setminus \mathbf{p} \cup \mathbf{p}_1 \cup \mathbf{p}_2$ )
1: function REITERATE( $\mathbf{P}$ )
2:    $\mathbf{p} \leftarrow \text{CHOOSE\_RANDOM}(\mathbf{P})$ 
3:   return ITERATE( $\mathbf{p}, \mathbf{P}$ )
1: function REMOVE_BLOB( $\mathbf{P}, \mathbf{x}$ )
2:    $\mathbf{p} \leftarrow \text{FIND\_NEAREST}(\mathbf{P}, \mathbf{x})$ 
3:   return ITERATE( $\mathbf{p}, \mathbf{P} \setminus \mathbf{p}$ )

```

---

with a parameter  $\sigma_{\text{max}}$  defining an upper limit for the blob size. We found that this two-stage approach is necessitated by the strong

non-convexity of the objective function. By optimizing over multiple blobs simultaneously, we allow the optimizer to recover complex geometry features that are influenced by more than a single blob.

The algorithm starts with a single blob as initial solution, then performs an outer loop over four phases: sampling, mutation, re-iteration, and regularization. In the following, we provide a full description of the individual phases and explain our design choices. The parameters used in our reconstructions are shown in Table 2.

*Sampling.* Our algorithm pivots around locations in the reconstruction volume that are randomly chosen according to a distribution (PDF) that aims to give problematic regions a higher probability of being sampled. We obtain the PDF by backprojecting the absolute value of the current residual image into the working volume. For locations  $\mathbf{x}$  that are sampled by the function `SAMPLE()`, our working hypothesis is that *something* about the solution should change there; we address this by selecting the nearest blob to this location (`FIND_NEAREST( $\mathbf{P}, \mathbf{x}$ )`) and applying and testing our three mutation strategies on it. Since each mutation probably increases the cost function, it is followed by a relaxation of the neighborhood of the pivot blob.

*Mutation.* We employ three mutation strategies to generate variations of the current solution. `ADD_BLOB( $\mathbf{P}, \mathbf{x}$ )` adds a new blob  $(\mathbf{x}, \sigma_0)$  to  $\mathbf{P}$ . `DELETE_BLOB( $\mathbf{P}, \mathbf{x}$ )` deletes the blob  $\mathbf{p} \in \mathbf{P}$  that is closest to  $\mathbf{x}$ . `DUPLICATE_BLOB( $\mathbf{P}, \mathbf{x}$ )` replaces the blob  $\mathbf{p} \in \mathbf{P}$  by two new blobs that are displaced by a vector  $\pm \mathbf{d}$  from the original position so they can be separated by the optimizer. Out of the three solutions (each one after performing an inner optimization `ITERATE( $\mathbf{p}, \mathbf{P}$ )` on the neighborhood), the one with the lowest cost  $c_i$  is chosen to be the new solution. A call to `ITERATE` consists of two non-linear optimizations, one solely over the blob positions, followed by an optimization over both blob positions and sizes. This procedure is essential due to the non-convexity of the cost function, initial experiments have shown that skipping the first optimization generally results in unwanted, strong local minima, where a single blob spans large parts of the reconstruction volume.

*Reiteration.* As the next step, another call to `ITERATE` is performed on a random group of neighboring blobs. This re-evaluation of previously relaxed blobs is necessary to avoid being stuck in local minima during early iterations, when the hypothesis does not yet contain enough blobs to properly describe the transient response.

*Regularization.* Finally, the algorithm first checks each blob for its significance to the solution (`CHECK_DELETE`), and deletes it if

doing so does not worsen the total cost by more than a small factor  $\eta$ . This regularizing step prevents the build-up of excess geometry in hidden regions that is not supported by the data. It is the only step that can lead to an increase in the cost  $c$ ; all other heuristics ensure that the cost falls monotonically.

#### 4.4 Implementation details

Our reconstruction software is written in C++. Geometry generation and rendering are implemented on the GPU, using NVIDIA CUDA and the Thrust parallel template library for the bulk of the tasks and the NVIDIA OptiX prime ray-tracing engine for the shadow tests. The optimization algorithm is implemented using the Ceres solver [Agarwal et al. 2015]. Intermediate results are visualized on-the-fly using the VTK library [Schroeder et al. 2006]. We used various workstations in our experiments, with Intel Core i7 CPUs and NVIDIA GeForce GPUs ranging from GTX 780 to Titan Xp.

### 5 EVALUATION

In this section, we verify the correctness of our renderer, and use it to reconstruct geometry from simulations and experimental measurements of around-the-corner scattered light. Input data, relay wall parameters, scene geometry, as well as output volumes and meshes of our proposed method and the state-of-the-art ellipsoidal backprojection method of [Arellano et al. 2017] can be found in the supplemental material.

#### 5.1 Scene geometry

All our synthetic experiments were conducted in a consistent scene setup and our models use the same arbitrary unit for length and time. The standard temporal resolution (size of histogram bin) of our virtual detectors is 0.4 units. Typical time resolutions of real-world devices are 10ps for streak cameras or 100ps for SPAD detectors. Equating the bin size with these time constants results in a conversion factor to real-world distances of 7.5mm and 75mm per world unit, respectively. We arranged the scene such that the relay wall is a diffuse plane at  $z = 0$  with normal in positive  $z$  direction. The object, with a typical size of 50 units, was located on the  $z$  axis at  $z = 45$ . The laser spot was modeled as a cosine-lobe light source pointing in positive  $z$  direction at one of four wall locations  $(45, 0, 0)$ ,  $(-45, 0, 0)$ ,  $(0, 45, 0)$  and  $(0, -45, 0)$ . The range of observed points on the wall was represented by an area of  $80 \times 80$  units<sup>2</sup> which was observed by an orthographic camera centered at  $(x, y) = (0, 0)$ . The synthetic scene parameters can also be found in Table 2 and the capture geometry of the measured scenes can be found in the respective publications [Buttafava et al. 2015; O’Toole et al. 2018a].

#### 5.2 Correctness of renderer

Before we evaluate the performance of our overall reconstruction system, we test correctness and performance of the forward model that is at its heart, our custom renderer. To this end, we prepare test scenes and render reference images using Microsoft’s Time of Flight Tracer [Slaney and Chou 2014], a transient renderer based on pbrt version 2 [Pharr and Humphreys 2010].

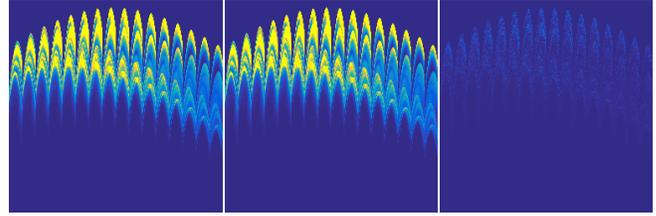


Fig. 7. The physically-based renderings with and without global illumination are virtually indistinguishable. From left to right: Rendering with global illumination; rendering without global illumination; difference of the two renderings.

Using a 30% reflective triangle mesh model of the Stanford Bunny, we generated two reference renderings of  $16 \times 16 \times 256$  spatio-temporal resolution using the physically-based renderer, one with full global illumination and one with a maximum path length of 2 reflections. With the cosine light source representing the spot lit by the laser, a path length of 2 includes light scattering from the wall to the object and back to the wall, but not light that has been interreflected at the object or that has bounced between object and wall multiple times. In Figure 7, both versions are shown along with the difference. At least for our around-the-corner setting, we found that the error caused by truncating the path length to 2 is not very significant, with 69.809dB peak signal-to-noise ratio (PSNR) or a relative  $L_2$  difference of 0.486%.

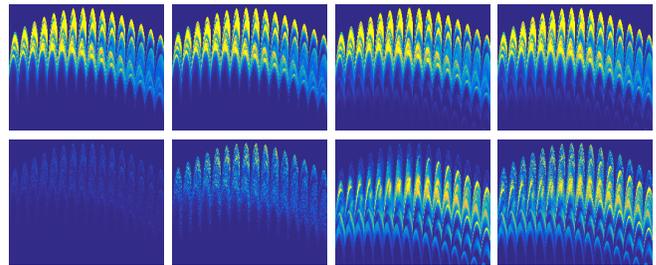


Fig. 8. The effect of our augmentations on the rendering error. The top row shows transient renderings made with our renderer, the bottom row shows the respective difference to the ground truth toftracer rendering (range scaled for print). From left to right: Our renderer with all features turned on; temporal filtering turned off; shadow tests turned off; temporal filtering and shadow tests turned off. Error metrics for these renderings are provided in Table 1.

We then used the truncated rendering as reference for our own renderer, and tested the effect of temporal filtering and shadow testing on the difference (Figure 8). A naïve version of our renderer, with all refinements disabled, reached the reference up to an error of a little under 10%. After activating the temporal filtering and the shadow tests, our fast renderer delivered a close approximation to the ray-traced reference with 69.796dB peak signal-to-noise ratio (PSNR) or a relative  $L_2$  difference of 0.489%. All error values are provided at a glance in Table 1. The main result from this investigation is that both features are essential to our renderer. The gain in accuracy comes at the expense of significantly increased runtime

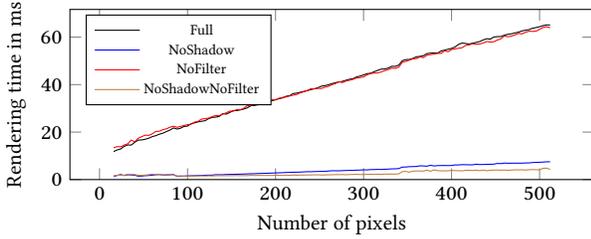


Fig. 9. Rendering performance of four versions of our algorithm (with/without filtering, with/without shadow test) as a function of output pixel count.

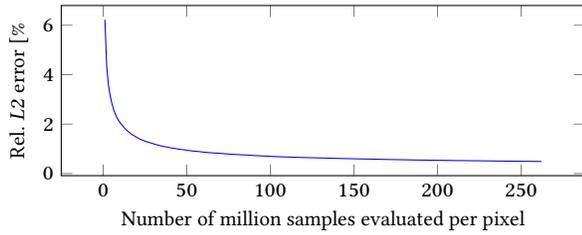


Fig. 10. Difference between our fast renderer and the ray-traced reference solution with a varying number of samples per pixels.

Comparison	PSNR [dB]	Rel. L2 error [%]
RTFull / RTTrunc	69.809	0.486
RTTrunc / OursFull	69.796	0.489
RTTrunc / OursNoFilter	53.379	3.237
RTTrunc / OursNoShadow	45.638	7.892
RTTrunc / OursNoShadowNoFilter	44.942	8.550

Table 1. Using the Stanford Bunny as test object, we compare our renderer to ray-traced renderings with maximum path lengths of 2 (RTTrunc) and  $\infty$  (RTFull). With all the features enabled (OursFull), our renderer matches the ray-traced solution for the 3-bounce setting (wall-object-wall) to 0.49%, which is on the same order as the influence of global illumination (RTFull) on this scene. Omission of shadow tests and temporal filtering result in significantly higher error values.

when using the shadow test (Figure 9). For small numbers of pixels, a significant part of that runtime is caused by the construction of acceleration structures—here, about 10ms for an object with approximately 55,000 triangles. Another noteworthy observation is that the Monte-Carlo rendering used as reference was likely not fully converged (Figure 10) even after evaluating 250 million samples per pixel. We expect that more exhaustive sampling would likely have further reduced the error.

### 5.3 Geometry reconstruction

We used various types of input data to test our algorithm: synthetic data generated using a path tracer or our own fast renderer, as well as experimental data obtained from other sources. The results from these reconstructions are scattered throughout the paper, referencing the datasets from Table 2 by their respective names. Meshes are rendered in a daylight environment using Mitsuba [Jakob 2010],

with a back wall and ground plane added as shadow receivers for better visualization of the 3D shapes. Note that these planes are not part of the experimental setup.

*Synthetic datasets.* After establishing in Section 5.2 that our fast renderer produces outcomes that are almost identical to the ray-traced reference, we used both the path tracer and our fast renderer to generate a variety of around-the-corner input data. In particular, we prepared several variations of the Mannequin scene, reducing the number of pixels, the number of laser spots, as well as the temporal resolution. An overview of our datasets, as well as the parameters used for reconstruction, can be found in Table 2. Like the backprojection method, our method has a small number of parameters: the blob size upper bound  $\sigma_0$  and the regularization parameter  $\eta$ .

We show renderings of the reconstructed meshes alongside the backprojected solutions, obtained using the Fast Backprojection code provided by Arellano et al. [2017], and ground truth (Figure 1(c)). They show that the quality delivered by our algorithm, in general, outperforms the state-of-the-art method on the synthetic datasets examined in this study. The meshes produced by our method tend to be more complete, smoother, and overall closer to the true surface. We also performed more quantitative evaluations. Figures 11 and 12 show the error of the recovered surface in  $z$ -direction for three datasets. We present backprojection results using two different filters: the Laplacian filter, since it is the most popular choice in literature, and the modified Difference of Gaussian filter [Laurenzis and Velten 2014a], which produced the most accurate and low-noise results from a wide range of tested filters. In general, meshes generated using the backprojection method tend to lie in front of the true surface. This is due to the way surface geometry is reconstructed from the density volumes obtained by the backprojection algorithm. Even if the peak of the density distribution lies exactly on the object geometry, extracting an isosurface will displace it by a certain distance. We were able to reduce this effect for the modified DoG filter by suppressing non-maximum density values along the  $z$ -axis. However, using this filter, we were still not able to produce reconstructions that achieved the same high coverage and low error as our method. It is the combination of a surface-oriented scattering model, in combination with our analysis-by-synthesis scheme, which vitally enables a reconstruction with systematic errors as low as the ones exhibited by our method.

*Degradation experiments.* To put the robustness of our method to the test, we performed a series of experiments that deliberately deviate from an idealized, noise-free, Lambertian and global-illumination-free light transport model, or reduce the amount of input data used for the reconstruction. In a first series of experiments, we sub-sampled the Mannequin dataset both spatially and temporally, and observed the degradation in reconstructed outcome (Figure 15). In a second series, we added increasing amounts of Poisson noise (Figure 16). Next, we replaced the diffuse reflectance of the BunnyGI model by a metal BRDF (Blinn model as implemented by pbrt) and decreased the roughness value (Figure 13). Our fast renderer used during reconstruction was set to the same BRDF parameters that were used to generate the input data. Finally, we constructed a strongly concave synthetic scene (Bowl) and used

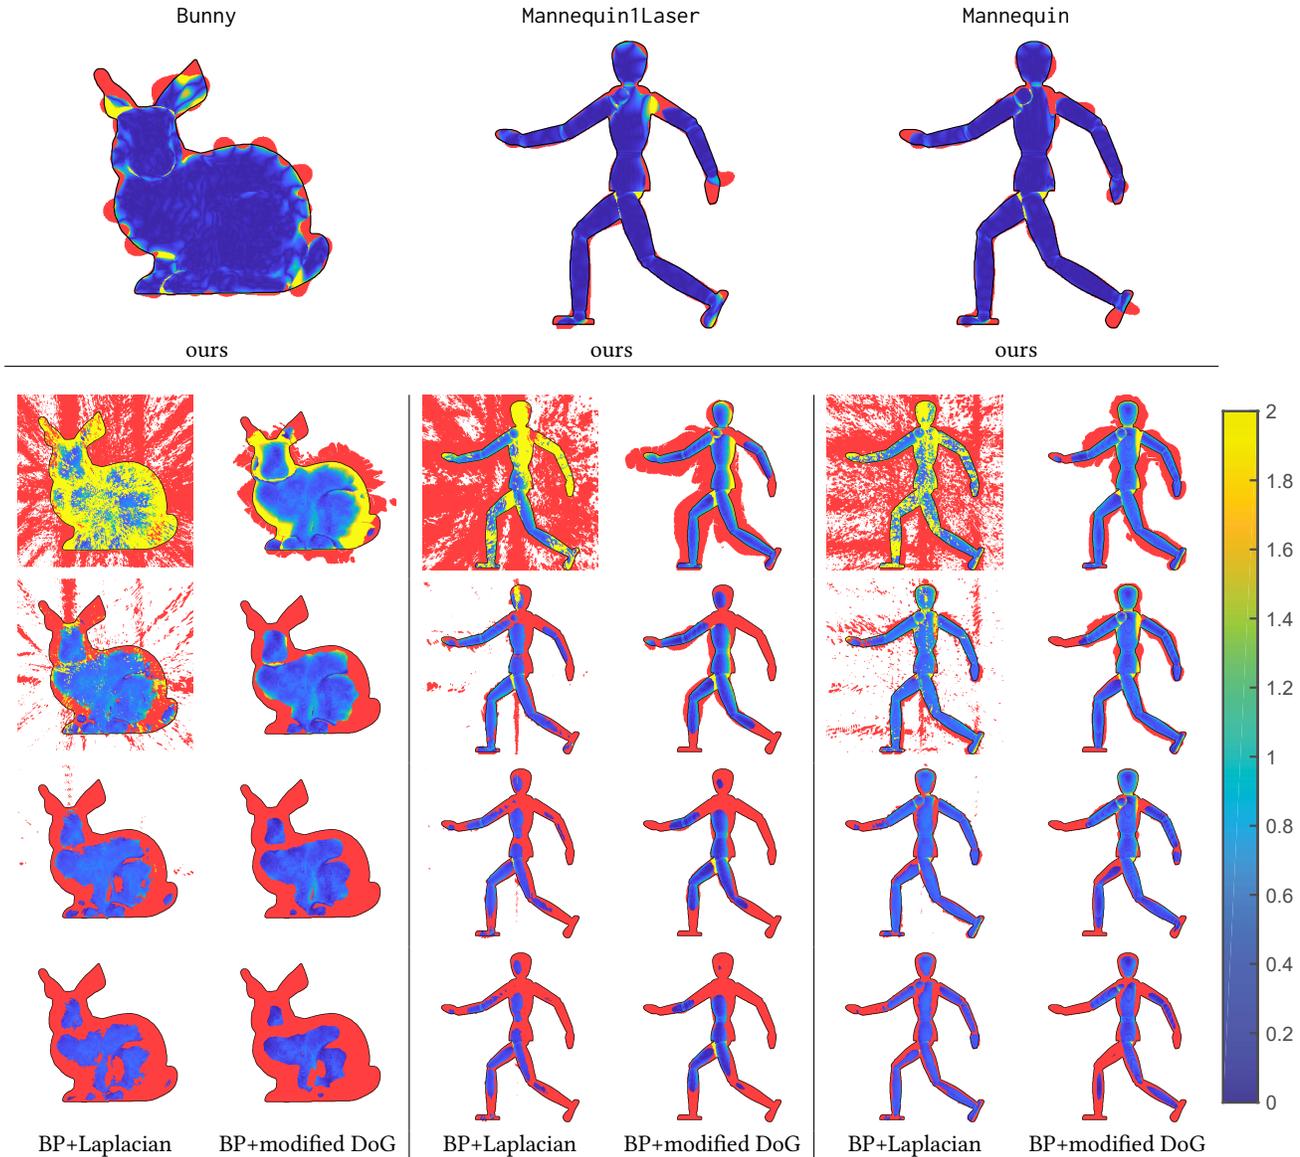


Fig. 11. Absolute depth error (in world units) in the reconstructions obtained from the synthetic Bunny (left), Mannequin1Laser (middle), and Mannequin (right) datasets. The top row shows error plots for our method, the bottom part shows depth errors for backprojection (BP) using the Laplacian filter and the modified Difference of Gaussian filter [Laurenzis and Velten 2014a] respectively, with increasing isovalues from top to bottom. The black line indicates the ground-truth object silhouette. Red color inside the silhouette indicates a missing (false-negative) surface and outside a silhouette it indicates excess (false-positive) geometry. Note that the range is clamped to  $[0, 2]$  for visualization; values plotted in yellow can be significantly higher. See Figure 12 for a quantitative analysis.

high albedo values in order to test the influence of unaccounted-for global illumination on the reconstructed geometry (Figure 14).

As expected, in all these examples, the further the data deviates from the ideal case, the more the reconstruction quality decreases. While backprojection tends to be more robust with respect to low-frequency bias (Bowl experiment), our method quite gracefully deals with high-frequency noise by fitting a low-frequent rendering to it.

For highly specular materials, the discretization of the surface mesh and the sensing locations on the wall may lead to sampling issues: specular glints that are missed by the forward simulation cannot contribute to the solution.

*Experimental datasets.* We show reconstructions of two experimental datasets obtained using SPAD sensors.

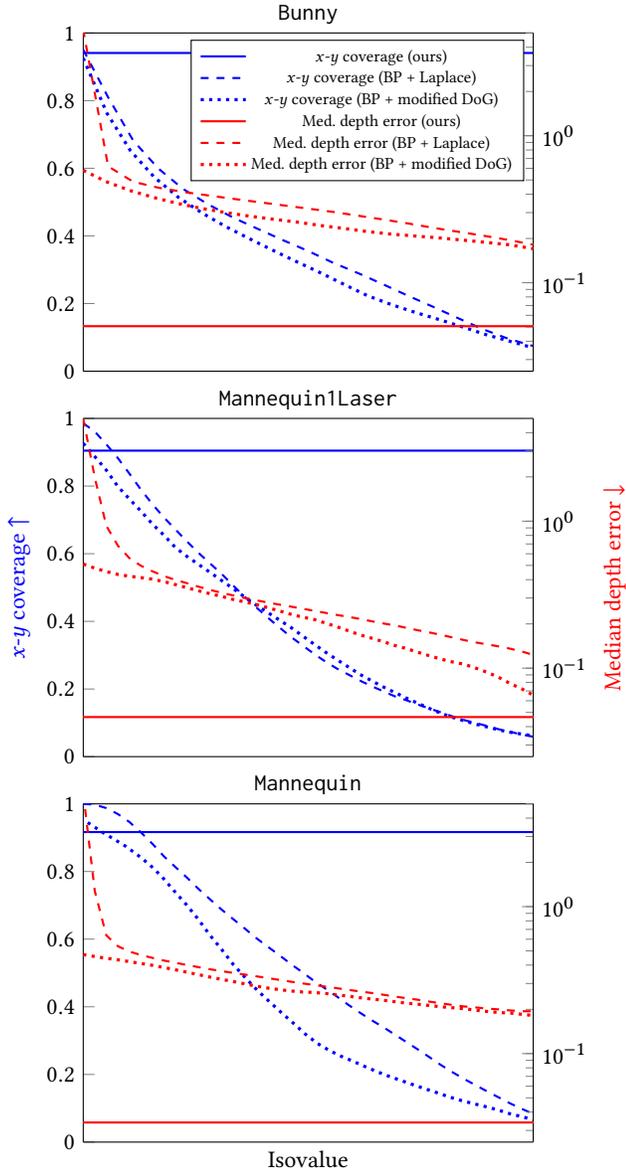


Fig. 12. Evaluation of the depth map coverage in the  $x$ - $y$  plane (higher is better) and the median absolute depth error in  $z$  direction (lower is better) for the Bunny, Mannequin1Laser, and Mannequin datasets. The proposed method achieves coverage values above 90% with a median depth error as low as 0.03 to 0.05 world units. For the state-of-the-art method (evaluated using the Laplacian filter and the modified Difference of Gaussian filter [Laurenzis and Velten 2014a]), no isovalue is capable of simultaneously achieving high coverage and low depth error. We note that the modified DoG filter generates less noise than the Laplacian filter for low isovalues. A qualitative visualization of this study can be found in Figure 11.

The first dataset (SPADScene) was measured by Buttafava et al. [2015], by observing a single location on the wall with a SPAD

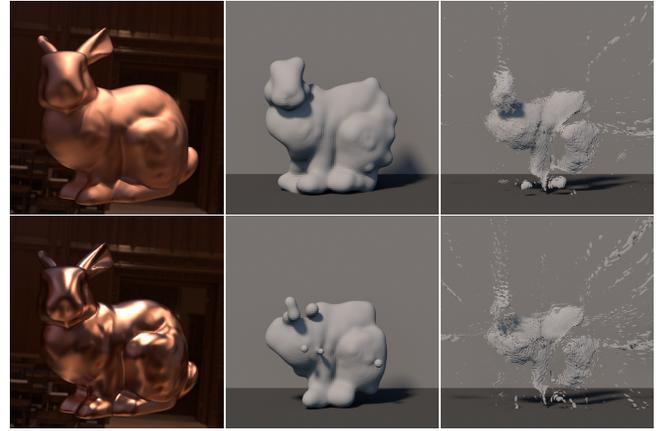


Fig. 13. Reconstruction of the BunnyMetal\* scenes with pbrt's metal BRDF applied to the object (top row: Blinn roughness 0.05; bottom row: Blinn roughness 0.01). From left to right: reference rendering in Grace Cathedral environment [Debevec 1998]; our proposed method; backprojection.

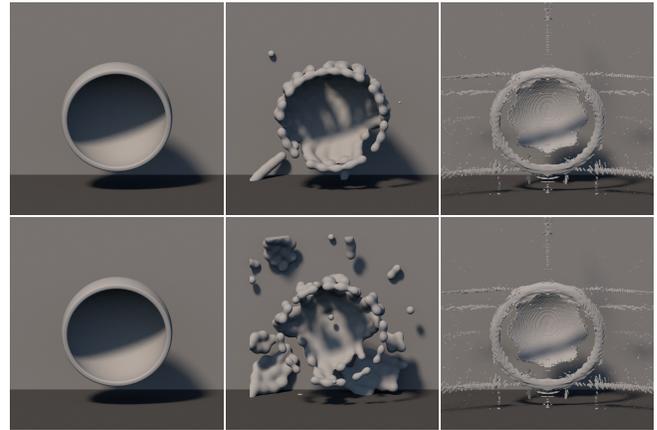


Fig. 14. Bowl scene. A strongly concave shape with high albedo (top row: 30%; bottom row: 100%) features large amounts of interreflected light in the input data, which leads to spurious features in the reconstructed geometry. From left to right: reference geometry; our proposed method; backprojection.

detector, and scanning a pulsed laser to a rectangular grid of locations. We note that this setup is dual, and hence equivalent for our purpose, to illuminating the single spot and scanning the detector to the grid of different locations. The dataset came included with the Fast Backprojection code provided by Arellano et al. [2017]. To apply our algorithm on the SPADScene dataset, we first subtracted a lowpass-filtered version (with  $\sigma = 1000$  bins) of the signal to reduce noise and background, then downsampled the dataset from its original temporal resolution by a factor of 25.

Like in the original work, the reconstruction remains vague and precise details are hard to make out (Figure 17). The reconstructed blobby objects appear to be in roughly the right places, but their shapes are poorly defined. We note that our method quite clearly

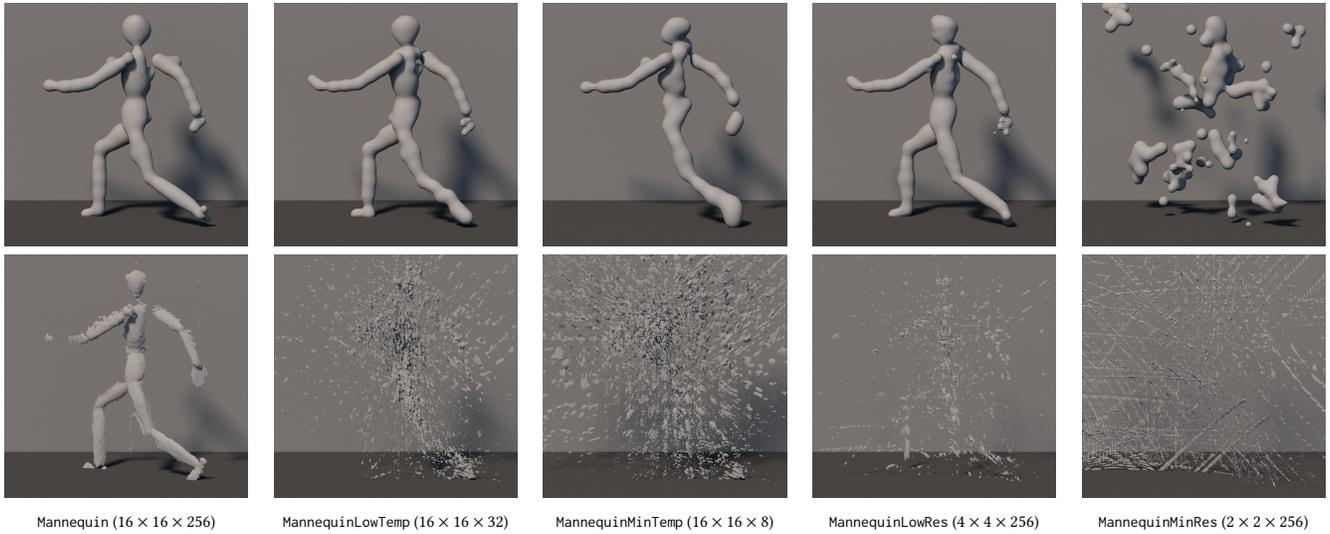


Fig. 15. Reconstruction of the Mannequin\* dataset using different levels of degradation. From left to right: Mannequin, MannequinLowTemp, MannequinMinTemp, MannequinLowRes, MannequinMinRes. Top row: Our reconstruction, bottom row: backprojection. Unlike backprojection, our reconstruction method handles degradations in the input data quite gracefully. Even an extremely low spatial resolution of  $2 \times 2$  pixels or a temporal resolution of only 8 bins still produces roughly identifiable results.

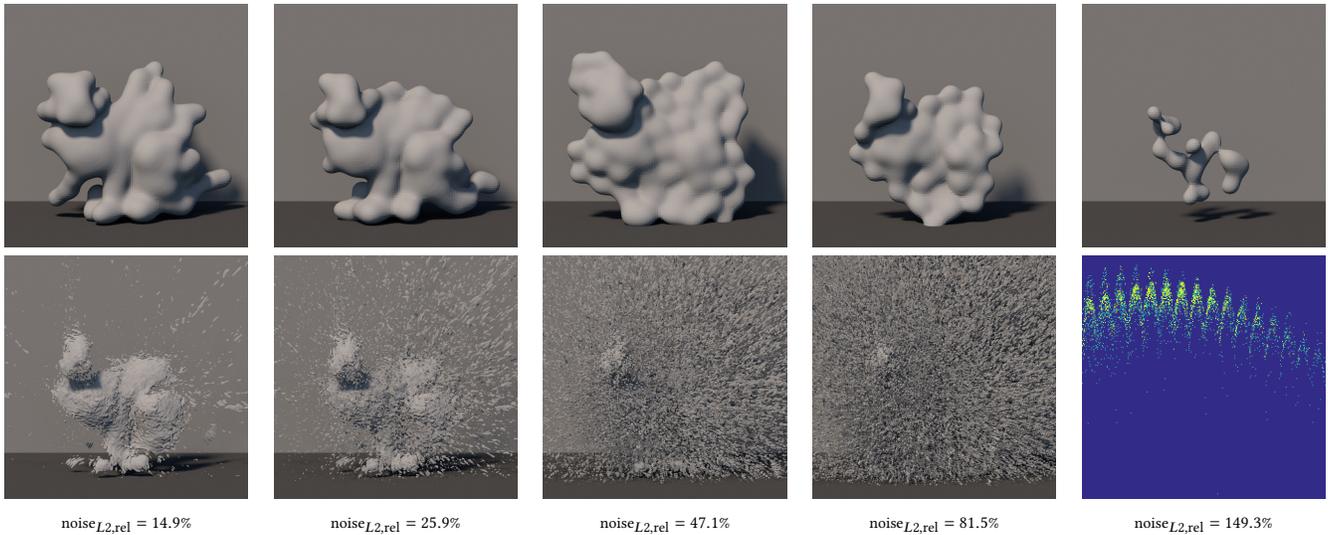


Fig. 16. Reconstruction of the BunnyGI dataset with different levels of Poisson noise applied to the input data. Relative  $L_2$  error from left to right: 14.9%, 25.9%, 47.1%, 81.5%, 149.3%. Top row: Our reconstruction, bottom row: backprojection. Our algorithm is based on a noise-free forward model. It therefore manages to localize the object reliably even under very noisy conditions (albeit at reduced reconstruction quality). In the rightmost example (streak plot), at most two photons have been counted per pixel, resulting in data that contains 50% more noise than signal.

carves out the letter “T” where backprojection delivers a less clearly defined shape (Figure 18).

The second dataset (O’TooleDiffuseS) is a measurement of a letter “S” cut from white cardboard, which O’Toole et al. measured via a diffuse wall using their confocal setup [O’Toole et al. 2018a]. In this setup, illumination and observation share the same optical path and are scanned across the surface. We downsampled the input

data by a factor of  $4 \times 4 \times 4$  in the spatial and temporal domains. Although the inclusion of the direct reflection in the data allowed for a better background subtraction and white point correction than in the case of the previous dataset, it becomes clear that there must be more sources of bias. In particular, we identified a temporal blur of roughly 3 time bins. Adding a similar blur to our renderer (a box

filter of width 3 bins), made the reconstructed “S” shape much more clearly recognizable as such (Figure 19).

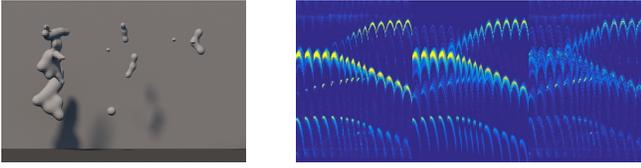


Fig. 17. Reconstruction of the experimental SPADScene dataset [Buttafava et al. 2015]. Shown is the output mesh and the transient data (from left to right: observation, prediction, residual).

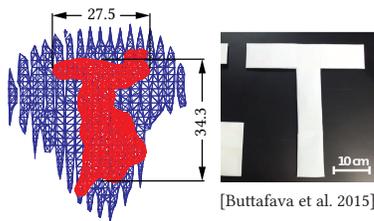


Fig. 18. The “T” object from the experimental SPADScene dataset published by Buttafava et al. [2015]. Shown are reconstructions obtained using back-projection (blue) and the proposed method (red), along with approximate dimensions using the scale provided in the original work (right).

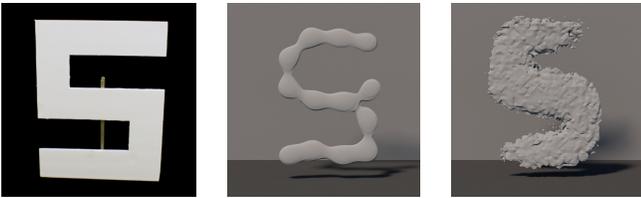


Fig. 19. OToolDiffuseS dataset [O’Toole et al. 2018a]. From left to right: photo of diffuse “S”-shaped cutout; surface mesh reconstructed using our method; mesh reconstructed using method described in [O’Toole et al. 2018a].

## 6 DISCUSSION

In the proposed approach, we develop computer graphics methodology (a near-physical, extremely efficient rendering scheme) to reconstruct occluded 3D shape from three-bounce indirect reflections. To our knowledge, this marks the first instance of a non-line-of-sight reconstruction algorithm that is consistent with a physical forward model. This solid theoretical foundation leads to results that, under favorable conditions, show higher object coverage and detail than the de-facto state of the art, error backprojection. In extreme situations, like very low spatial / temporal resolutions or high noise levels, we have shown that our method breaks down significantly later than the current state of the art (Figures 15 and 16). Under conditions that are not covered by the forward model (noise, bias/background,

global illumination) the results are on par or slightly inferior to existing methods. In terms of runtime, our method typically takes several hours or even days for a reconstruction run (Table 2) and therefore cannot compete with recent optimized versions of error backprojection [Arellano et al. 2017] or GPU-based deconvolvers [O’Toole et al. 2018b], which are typically on the order of 10s to 100s and 1s respectively. However, we consider this a soft hindrance that has to be considered together with the fact that the capture of suitable input data, too, is far from being instantaneous. This latter factor is governed by the physics of light and therefore may turn out, in the long run, to impose more severe limitations to the practicality of non-line-of-sight sensing solutions.

We noted that the reconstruction quality of the SPAD datasets stays behind the quality of the synthetic datasets (whether path-traced or using our own renderer). Our image formation model approximates the physical light transport up to very high accuracy (as shown in Section 5.2), but does not explicitly model the SPAD sensor response to the incoming light. The SPAD data is biased due to background noise and dark counts, and the temporal impulse response is asymmetric and smeared out due to time jitter and afterpulsing [Gulinatti et al. 2011; Hernandez et al. 2017]. While these effects could easily be incorporated into our forward model, doing so would require either a careful calibration of the imaging setup (which was not provided with the public datasets) or an estimation of the noise parameters from input data. In this light, we find the presented results very promising for this line of research, and consider the explicit application of measured noise profiles and the modeling of additional imaging setups as future work.

A key feature of our method is that, within the limitations of the forward model (opaque, but not necessarily diffuse, light transport without further interreflections) good solutions can be immediately identified by a low residual error. However, the non-convex objective and possibly unknown noise and background terms may make it challenging to reach this point. Our optimization scheme, while delivering good results in the provided examples, offers no guarantee of global convergence. As of today, it is unclear which of the two factors will prove more important in practice, the physical correctness of the forward model or the minimizability of the objective derived from it.

## 7 FUTURE WORK

We imagine that extended versions of our method could be used to jointly estimate geometry and material. Advanced global optimization heuristics could further improve the convergence behavior and the overall quality of the outcome. We imagine that hierarchical approaches or hybrid solutions might bring further improvement, for instance by using the (physically inaccurate but global) solution of one reconstruction scheme to warm-start another local optimization run using a more accurate model like ours.

The extrinsic and intrinsic calibration of traditional 2D imaging setups is well understood [Zhang 2000]. However, this problem has not been satisfyingly solved by the NLOS reconstruction community so far. The current best practice is to manually estimate the positions and normals of the projected camera pixels, potentially

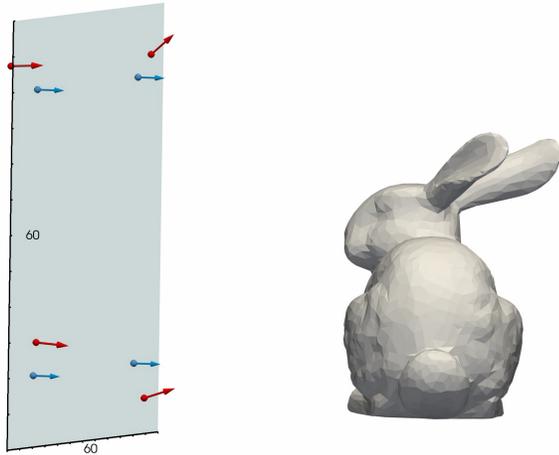


Fig. 20. An illustration of our preliminary NLOS camera calibration experiment. A transient image of the Bunny scene has been acquired using our transient renderer. The blue arrows denote the ground truth positions and normals of the projected camera pixels. The red arrows show an initial guess before optimization with a positional RMSE of 2.7 units and an average angular error of  $16.6^\circ$ . After optimization using Equation 1, the optimized pixel positions and normals coincide with the ground truth up to floating point precision. Due to the greatly reduced number of variables compared to the geometry reconstruction problem, the optimization concluded in less than one minute.

leading to a systematic bias in the (typically non-metric) reconstructions. Our proposed method presents not only an alternative solution to NLOS reconstruction, but also lays out a foundation for solving related problems. Here, we presented a method for recovering the scene geometry, where the acquisition geometry was assumed to be known. In future work, we would like to study the *dual problem*, where the scene geometry is known (a calibration target), but the acquisition geometry is unknown. We conducted initial experiments with our synthetic Bunny dataset and were able to recover the positions and normals of four projected pixels up to a very high precision, regardless of an overly imprecise initial guess, see Figure 20. Again, we utilize Equation 1 as the objective function, but the parameter vector  $\mathbf{P}$  consists of the positions and normals of the projected pixels. Challenges will include the generalization to real-world data, the design of an optimum calibration target, and the validation against measured data. We could also imagine utilizing our forward model to estimate the parameters of a SPAD sensor response model [Hernandez et al. 2017].

Finally, our renderer is not constrained to use in a costly iterative solver. Just as well, we can imagine using it to enable new machine learning approaches to the problem. A suitably trained feedforward neural network, for example, would deliver instant results. Whereas existing renderers are too slow for generating large amounts of training data, our renderer would be fast enough to obtain millions of datasets in a single day. Together with a suitable signal degradation model [Hernandez et al. 2017], we expect that it will be possible to closely approximate the most relevant real-world scenarios.

## REFERENCES

- Nils Abramson. 1978. Light-in-flight recording by holography. *Optics Letters* 3, 4 (Oct 1978), 121–123. <https://doi.org/10.1364/OL.3.000121>
- Sameer Agarwal, Keir Mierle, and Others [sic]. 2015. Ceres Solver. <http://ceres-solver.org>.
- Victor Arellano, Diego Gutierrez, and Adrian Jarabo. 2017. Fast back-projection for non-line of sight reconstruction. *Optics Express* 25, 10 (2017).
- Jeremy Boger-Lombard and Ori Katz. 2018. Non line-of-sight localization by passive optical time-of-flight. *arXiv preprint arXiv:1808.01000* (2018).
- Katherine L Bouman, Vickie Ye, Adam B Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. 2017. Turning Corners into Cameras: Principles and Methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2270–2278.
- Mauro Buttaviva, Jessica Zeman, Alberto Tosi, Kevin Eliceiri, and Andreas Velten. 2015. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics Express* 23, 16 (2015), 20997–21011.
- Jonathan Carr, Richard Beatson, Jon Cherrie, Tim Mitchell, W Fright, Bruce McCallum, and Tim Evans. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *Proc. 28th Annual Conf. on Computer Graphics and Interactive Techniques*. ACM, 67–76.
- Paul Debevec. 1998. Light Probe Image Gallery. <http://www.pauldebevec.com/Probes/>.
- Stefan Fuchs. 2010. Multipath interference compensation in time-of-flight camera images. In *Pattern Recognition (ICPR), 2010 20th International Conference on*. IEEE, 3583–3586.
- G Garipey, N Krstajic, R Henderson, C Li, RR Thomson, GS Buller, B Heshmat, R Raskar, J Leach, and D Faccio. 2015. Single-photon sensitive light-in-flight imaging. *Nature Communications* 6 (2015).
- Genevieve Garipey, Francesco Tonolini, Robert Henderson, Jonathan Leach, and Daniele Faccio. 2016. Detection and tracking of moving objects hidden from view. *Nature Photonics* 10, 1 (2016).
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 162.
- A. Gulinatti, I. Rech, M. Assanelli, M. Ghioni, and S. Cova. 2011. A physically based model for evaluating the photon detection efficiency and the temporal response of SPAD detectors. *Journal of Modern Optics* 58, 3-4 (2011), 210–224. <https://doi.org/10.1080/09500340.2010.536590>
- Felix Heide, Matthias B. Hullin, James Gregson, and Wolfgang Heidrich. 2013. Low-budget Transient Imaging using Photonic Mixer Devices. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 32, 4 (2013), 45:1–45:10.
- Felix Heide, Matthew O’Toole, Kai Zhang, David B. Lindell, Steven Diamond, and Gordon Wetzstein. 2017. Robust Non-line-of-sight Imaging with Single Photon Detectors. *CoRR* abs/1711.07134 (2017). [arXiv:1711.07134](http://arxiv.org/abs/1711.07134) <http://arxiv.org/abs/1711.07134>
- Felix Heide, Lei Xiao, Wolfgang Heidrich, and Matthias B. Hullin. 2014. Diffuse Mirrors: 3D Reconstruction from Diffuse Indirect Illumination Using Inexpensive Time-of-Flight Sensors. *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (2014).
- Quercus Hernandez, Diego Gutierrez, and Adrian Jarabo. 2017. A Computational Model of a Single-Photon Avalanche Diode Sensor for Transient Imaging. *arXiv preprint arXiv:1703.02635* (2017).
- Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics (TOG)* 33, 6 (2014), 177.
- Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017), 65–79.
- Achuta Kadambi, Hang Zhao, Boxin Shi, and Ramesh Raskar. 2016. Occluded Imaging with Time-of-Flight Sensors. *ACM Transactions on Graphics* 35, 2, Article 15 (March 2016), 12 pages. <https://doi.org/10.1145/2836164>
- Ori Katz, Pierre Heidmann, Mathias Fink, and Sylvain Gigan. 2014. Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations. *Nature Photonics* 8, 10 (2014), 784–790.
- A. Kirmani, T. Hutchison, J. Davis, and R. Raskar. 2009. Looking around the corner using transient imaging. In *Proc. ICCV*. 159–166.
- Jonathan Klein, Christoph Peters, Jaime Martín, Martin Laurenzis, and Matthias B. Hullin. 2016. Tracking objects outside the line of sight using 2D intensity images. *Scientific Reports* 6, 32491 (2016).
- M. La Manna, F. Kine, E. Breitbach, J. Jackson, T. Sultan, and A. Velten. 2018. Error Backprojection Algorithms for Non-Line-of-Sight Imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018), 1–1. <https://doi.org/10.1109/TPAMI.2018.2843363>
- Martin Laurenzis and Andreas Velten. 2014a. Feature selection and back-projection algorithms for nonline-of-sight laser-gated viewing. *Journal of Electronic Imaging* 23, 6 (2014), 1 – 6. <https://doi.org/10.1117/1.JEI.23.6.063003>

Name	Reference	Resolution	# Lasers	$s_{\text{geom}}$	$s_{\text{camera}}$	$\eta$	$\sigma_0$	$t_0$	$\delta_t$	$c_n/c_0$ [%]	T [min]	$n_{\text{iters}}$	$n_{\text{blobs}}$
Bunny	Ours	16 × 16 × 256	4	40 × 40	80 × 80	1.01	1.5	80	0.4	0.32	5096	660	156
BunnyGI	pbrt	16 × 16 × 256	4	40 × 40	80 × 80	1.01	1.5	80	0.4	0.59	5611	181	109
BunnyMetal0.05	pbrt	16 × 16 × 256	4	40 × 40	80 × 80	1.01	1.5	80	0.4	2.02	3419	259	101
BunnyMetal0.01	pbrt	16 × 16 × 256	4	40 × 40	80 × 80	1.01	1.5	80	0.4	11.75	3005	361	87
Bow1Albedo0.3	pbrt	16 × 16 × 256	4	26 × 26	80 × 80	1.001	0.4	80	0.4	4.36	5579	167	155
Bow1Albedo1	pbrt	16 × 16 × 256	4	26 × 26	80 × 80	1.001	0.4	80	0.4	31.53	4280	267	197
Mannequin	Ours	16 × 16 × 256	4	40 × 494	480 × 80	1.005	1.5	90	0.4	1.46	2326	505	69
MannequinLowRes	Ours	4 × 4 × 256	4	40 × 49	80 × 80	1.005	1.5	90	0.4	1.41	1251	252	76
MannequinMinRes	Ours	2 × 2 × 256	4	40 × 49	80 × 80	1.005	1.5	90	0.4	4.44	931	350	101
MannequinLowTemp	Ours	16 × 16 × 32	4	40 × 49	80 × 80	1.005	1.5	90	3.2	1.21	1322	166	67
MannequinMinTemp	Ours	16 × 16 × 8	4	40 × 49	80 × 80	1.005	1.5	90	12.8	7.95	420	102	23
Mannequin1Laser	Ours	16 × 16 × 256	1	40 × 49	80 × 80	1.005	1.5	90	0.4	0.59	1419	243	57
SPADScene	Measured	185 × 1 × 256	1	—	—	1.005	4.5	373	0.748	20.31	1280	328	43
OTooleDiffuseS	Measured	64 × 64 × 2048	1	—	—	1.01	0.015	0.756	0.0012	33.43	67	13	13

Table 2. Parameters of our reconstructed scenes, where  $s_{\text{geom}}$  is the size of the ground truth object projected onto the diffuse camera wall in world units,  $s_{\text{camera}}$  is the area covered by the camera in world units,  $\eta$  is the drop deletion factor in Algorithm 3,  $\sigma_0$  is the initial blob standard deviation,  $t_0$  is the time stamp of the first time bin,  $\delta_t$  is the size of a time bin, and  $c_n/c_0$  is the residual cost after optimization (relative to the initial cost). The total reconstruction times  $T$  are taken from file timestamps and vary due to manual termination of the reconstruction procedure, execution on different GPU models, overhead through parallel execution of multiple jobs, as well as debugging output. The optimizations terminated after  $n_{\text{iters}}$  iterations and consist of  $n_{\text{blobs}}$  Gaussian blobs. Please note that the exact scene geometry is only known for the synthetic experiments.

Martin Laurenzis and Andreas Velten. 2014b. Nonline-of-sight laser gated viewing of scattered photons. *Optical Engineering* 53, 2 (2014), 023102–023102.

Kenneth Levenberg. 1944. A Method for the Solution of Certain Non-linear Problems in Least Squares. *Quart. Appl. Math.* 2, 2 (1944), 164–168.

William Lorensen and Harvey Cline. 1987. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In *Proc. 14th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. ACM, 163–169. <https://doi.org/10.1145/37401.37422>

Julio Marco, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2017. Transient Photon Beams. In *Spanish Computer Graphics Conference (CEIG)*. The Eurographics Association. <https://doi.org/10.2312/ceig.20171216>

Donald W. Marquardt. 1963. An algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM J. Appl. Math.* 11, 2 (1963), 431–441. <https://doi.org/10.1137/0111030>

N. Naik, S. Zhao, A. Velten, R. Raskar, and K. Bala. 2011. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graph.* 30, 6 (2011), 171.

Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. 2018a. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* 555, 25489 (2018), 338–341.

Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. 2018b. Real-time Non-line-of-sight Imaging. In *ACM SIGGRAPH 2018 Emerging Technologies* (Vancouver, British Columbia, Canada) (SIGGRAPH '18). ACM, New York, NY, USA, Article 14, 2 pages. <https://doi.org/10.1145/3214907.3214920>

Adithya Kumar Pediredla, Mauro Buttafava, Alberto Tosi, Oliver Cossairt, and Ashok Veeraraghavan. 2017. Reconstructing rooms using photon echoes: A plane based model and reconstruction algorithm for looking around the corner. In *Computational Photography (ICCP), 2017 IEEE International Conference on*. IEEE, 1–12.

Matt Pharr and Greg Humphreys. 2010. *Physically Based Rendering, Second Edition: From Theory To Implementation* (2nd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Will Schroeder, Ken Martin, and Bill Lorensen. 2006. *The Visualization Toolkit—An Object-Oriented Approach To 3D Graphics* (fourth ed.). Kitware, Inc.

Malcolm Slaney and Philip A. Chou. 2014. *Time of Flight Tracer*. Technical Report. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/time-of-flight-tracer/>

Adam Smith, James Skorupski, and James Davis. 2008. *Transient Rendering*. Technical Report UCSC-SOE-08-26. School of Engineering, University of California, Santa Cruz.

Christos Thrapoulidis, Gal Shulkind, Feihu Xu, William T Freeman, Jeffrey H Shapiro, Antonio Torralba, Franco NC Wong, and Gregory W Wornell. 2017. Exploiting Occlusion in Non-Line-of-Sight Active Imaging. *arXiv preprint arXiv:1711.06297* (2017).

Andreas Velten, Ramesh Raskar, and Mounqi Bawendi. 2011. Picosecond Camera for Time-of-Flight Imaging. In *Imaging and Applied Optics*. IMB4. <https://doi.org/10.1364/ISA.2011.IMB4>

A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M.G. Bawendi, and R. Raskar. 2012. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications* 3 (2012), 745.

Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounqi Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femto-photography: Capturing and Visualizing the Propagation of Light. *ACM Trans. Graph.* 32, 4, Article 44 (July 2013), 8 pages. <https://doi.org/10.1145/2461912.2461928>

Di Wu, Andreas Velten, Matthew O’Toole, Belen Masia, Amit Agrawal, Qionghai Dai, and Ramesh Raskar. 2014. Decomposing Global Light Transport Using Time of Flight Imaging. *International Journal of Computer Vision* 107, 2 (01 Apr 2014), 123–138. <https://doi.org/10.1007/s11263-013-0668-2>

Di Wu, Gordon Wetzstein, Christopher Barsi, Thomas Willwacher, Matthew O’Toole, Nikhil Naik, Qionghai Dai, Kyros Kutulakos, and Ramesh Raskar. 2012. Frequency analysis of transient light transport with applications in bare sensor imaging. *Computer Vision—ECCV 2012* (2012), 542–555.

Zengyou Zhang. 2000. A Flexible New Technique for Camera Calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (December 2000), 1330–1334. <https://www.microsoft.com/en-us/research/publication/a-flexible-new-technique-for-camera-calibration/> MSR-TR-98-71, Updated March 25, 1999.