

Beyond Volumetric Albedo

— A Surface Optimization Framework for Non-Line-of-Sight Imaging: Supplementary Material

Chia-Yin Tsai, Aswin C. Sankaranarayanan, and Ioannis Gkioulekas
Carnegie Mellon University

1. Introduction

In this supplement, we cover the following topics:

- In Section 2, we prove and discuss Proposition 1. In particular, we detail the GGX reflectance model, and derive the associated derivative expressions. We then derive the expression for the derivative with respect to mesh vertices. Finally, we discuss computing derivatives while correctly accounting for visibility terms.
- In Section 3, we discuss in detail the rendering algorithm used for estimating the transient and its derivatives. Also, we discuss extensively our optimization procedure and related engineering issues.
- In Section 4, we show additional experimental results. In particular, we provide quantitative metrics for the shape reconstruction experiments in the main paper. We additionally evaluate performance for different numbers of measurements and levels of noise.

2. Differentiating transients

2.1. Reflectance model

We assume that the NLOS surface has a spatially-uniform BRDF, which we model using the GGX microfacet BRDF [29]. We first define the surface normal $\hat{\mathbf{n}}$, incoming and outgoing directions ω_i and ω_o , and *half-vector* $\mathbf{h} = (\omega_i + \omega_o) / \|\omega_i + \omega_o\|$. Then, the BRDF model is:

$$\begin{aligned} f_s[\pi](\hat{\mathbf{n}}, \omega_i, \omega_o) &= f_{\text{GGX}}[\alpha](\hat{\mathbf{n}}, \omega_i, \omega_o) \\ &= \frac{F_{\text{Cook-Torrance}}(\omega_o, \mathbf{h}) \cdot D_{\text{GGX}}(\hat{\mathbf{n}}, \mathbf{h}) \cdot G_{\text{Smith}}(\hat{\mathbf{n}}, \omega_i, \omega_o)}{4}, \end{aligned} \quad (14)$$

where

$$D_{\text{GGX}}(\hat{\mathbf{n}}, \mathbf{h}) = \frac{\alpha^2}{\pi \left[\langle \hat{\mathbf{n}}, \mathbf{h} \rangle^2 (\alpha^2 - 1) + 1 \right]^2}, \quad (15)$$

$$G_{\text{Smith}}(\hat{\mathbf{n}}, \omega_i, \omega_o) = G_1(\hat{\mathbf{n}}, \omega_i) \cdot G_1(\hat{\mathbf{n}}, \omega_o), \quad (16)$$

$$G_1(\hat{\mathbf{n}}, \omega) = \frac{2}{\langle \hat{\mathbf{n}}, \omega \rangle + \sqrt{\alpha^2 + (1 - \alpha^2) \langle \hat{\mathbf{n}}, \omega \rangle^2}}, \quad (17)$$

and $F_{\text{Cook-Torrance}}(\omega_o, \mathbf{h})$ is the Fresnel reflection of an ideal reflector, independent of α . This leads to a one-dimensional reflectance parameterization, $\pi = \{\alpha\}$, where the parameter $\alpha \in [0, 1]$ controls surface roughness.

2.2. Proposition 1

Proposition 1 of the main paper is a direct consequence of the Equations (18) and (21) we derive below.

Reflectance derivatives. Differentiating the surface integral of Equation (1) of the main paper with respect to reflectance is straightforward. By observing that the integration measure is independent of the reflectance parameters π , and under weak continuity conditions that are known to be satisfied for radiometric integrals [2], we can simply exchange the order of differentiation and integration to obtain

$$\frac{\partial I(t; \mathbf{l}, \mathbf{s})}{\partial \pi} \Big|_{v, \pi} = \int_{S_{\text{NLOS}}} \overbrace{W(\mathbf{x}; t) \frac{\partial f(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}))}{\partial \pi}}^{g_r(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}))} \Big|_{v, \pi} \cdot v(\mathbf{x}, \mathbf{l}) v(\mathbf{x}, \mathbf{s}) \, dA(\mathbf{x}), \quad (18)$$

where

$$\begin{aligned} \frac{\partial f(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}))}{\partial \pi} \Big|_{v, \pi} &= \frac{\partial f_s[\pi](\hat{\mathbf{n}}(\mathbf{x}), \hat{\omega}_l(\mathbf{x}), \hat{\omega}_s(\mathbf{x}))}{\partial \pi} \Big|_{v, \pi} \\ &\quad \cdot \frac{\langle -\hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \langle \hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle}{\|\mathbf{x} - \mathbf{l}\|^2} \\ &\quad \cdot \frac{\langle -\hat{\omega}_s(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{s}) \rangle \langle \hat{\omega}_s(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle}{\|\mathbf{x} - \mathbf{s}\|^2}, \end{aligned} \quad (19)$$

and the derivatives $\partial f_s / \partial \pi$ can be computed analytically from Equations (14)-(17),

$$\frac{\partial f_s}{\partial \alpha} = \frac{\partial f_{\text{GGX}}[\alpha](\hat{\mathbf{n}}, \omega_i, \omega_o)}{\partial \alpha}. \quad (20)$$

We used symbolic differentiation to compute this derivative, and then implemented the resulting expression in our rendering code.

Surface derivatives. In contrast to the reflectance case, differentiating the surface integral of Equation (1) of the main paper with respect to mesh vertices is complicated by the fact that the integration measure $A(\mathbf{x})$ is now a function of these same mesh vertices. An additional complicating factor is the presence of the binary and discontinuous visibility terms v in the integrand. We can ameliorate the second problem by making the approximation that the visibility terms are independent of the mesh geometry. This approximation is justified by the fact that the visibility terms have non-zero derivatives only on a zero-measure part of the surface (that is, on occluding contours [13]), and is common in computer vision and graphics problems [1, 22, 14].

Under this assumption, we can directly apply the derivation in Section 3.2 of Delaunoy and Prados [4]. Then, the derivative of Equation (1) of the main paper with respect to each vertex \mathbf{v}_i becomes:

$$\begin{aligned} \left. \frac{\partial I(t; \mathbf{l}, \mathbf{s})}{\partial \mathbf{v}_i} \right|_{\mathbf{v}, \pi} &= \sum_{k \in J_i} \int_{\mathcal{T}_k} \overbrace{[g_{s1}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}); i, k) + g_{s2}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}); i, k)]}^{g_s(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}); i, k)} \\ &\quad \cdot v(\mathbf{x}, \mathbf{l}) v(\mathbf{x}, \mathbf{s}) \, dA(\mathbf{x}), \end{aligned} \quad (21)$$

where

$$\begin{aligned} g_{s1}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}); i, k) &= \nabla_{\mathbf{x}} g(\mathbf{x}, \hat{\mathbf{n}}_k) \phi_i(\mathbf{x}), \\ g_{s2}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x}); i, k) &= -\frac{\mathbf{e}_{k,i}}{2A_k} \wedge [g(\mathbf{x}, \hat{\mathbf{n}}_k) \hat{\mathbf{n}}_k + g_{\hat{\mathbf{n}}}(\mathbf{x}, \hat{\mathbf{n}}_k)]. \end{aligned} \quad (22)$$

J_i corresponds to the set of triangles that contain vertex i . $\nabla_{\mathbf{x}}$ is the gradient of g with respect to \mathbf{x} , $\phi_i(\mathbf{x})$ is the linear interpolating basis function. $\mathbf{e}_{k,i}$ is the opposite edge of vertex i in the face k pointing counterclockwise. Operator \wedge denotes the cross product. Finally,

$$g_{\hat{\mathbf{n}}} = \nabla_{\hat{\mathbf{n}}} g(\mathbf{x}, \hat{\mathbf{n}}_k) - \langle \nabla_{\hat{\mathbf{n}}} g(\mathbf{x}, \hat{\mathbf{n}}_k), \hat{\mathbf{n}}_k \rangle \hat{\mathbf{n}}_k. \quad (24)$$

In the following, we detail some quantities that are used in the above expressions. For simplicity, we show the expression for the case of confocal imaging ($\mathbf{l} = \mathbf{s}$). The

derivation for non-confocal case is straightforward, only involves longer expressions. As defined in Equations (1) and (2) of the main paper,

$$\begin{aligned} g(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) &= W(\mathbf{x}; t) f(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) \\ &= W(\mathbf{x}; t) f_s(\hat{\mathbf{n}}(\mathbf{x}), \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\omega}_{\mathbf{s}}(\mathbf{x})) \\ &\quad \cdot \frac{\langle -\hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle}{\|\mathbf{x} - \mathbf{l}\|^2} \\ &\quad \cdot \frac{\langle -\hat{\omega}_{\mathbf{s}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{s}) \rangle \langle \hat{\omega}_{\mathbf{s}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle}{\|\mathbf{x} - \mathbf{s}\|^2}, \\ &= W(\mathbf{x}; t) f_s(\hat{\mathbf{n}}(\mathbf{x}), \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\omega}_{\mathbf{s}}(\mathbf{x})) g_0(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) \end{aligned} \quad (25)$$

The temporal importance W function is an unit rectangular function. This means that the derivatation will only be non-zero on the boundary. To make g defferentiable, we substitute W with a Gaussian that has full width at half max equal to the temporal resolution of the sensor. Then,

$$\begin{aligned} W(\mathbf{x}; t) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\tau(\mathbf{x}) - t)^2}{2\sigma^2}\right) \\ \frac{\partial W(\mathbf{x}; t)}{\partial \mathbf{x}} &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\tau(\mathbf{x}) - t)^2}{2\sigma^2}\right) \frac{2[\tau(\mathbf{x}) - t]}{\sigma^2} \hat{\omega}_{\mathbf{l}} \\ \frac{\partial W(\mathbf{x}; t)}{\partial \hat{\mathbf{n}}} &= \mathbf{0} \end{aligned} \quad (26)$$

In addition, the terms $\nabla_{\mathbf{x}} g(\mathbf{x}, \hat{\mathbf{n}}_j)$ and $\nabla_{\hat{\mathbf{n}}} g(\mathbf{x}, \hat{\mathbf{n}}_j)$ can be computed through chain rule, as follows

$$\begin{aligned} g_0(\mathbf{x}, \hat{\mathbf{n}}) &= \frac{\langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}} \rangle^2 \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle^2}{\|\mathbf{x} - \mathbf{l}\|^4} \\ \frac{\partial g_0(\mathbf{x}, \hat{\mathbf{n}})}{\partial \mathbf{x}} &= \frac{2 \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}} \rangle}{\|\mathbf{x} - \mathbf{l}\|^5} \\ &\quad \cdot [\hat{\mathbf{n}}(\mathbf{l}) \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}} \rangle + \hat{\mathbf{n}} \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \\ &\quad + 4 \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}} \rangle \hat{\omega}_{\mathbf{l}}(\mathbf{x})] \\ \frac{\partial g_0(\mathbf{x}, \hat{\mathbf{n}})}{\partial \hat{\mathbf{n}}} &= \frac{2 \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle^2 \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}} \rangle}{\|\mathbf{x} - \mathbf{l}\|^4} \hat{\omega}_{\mathbf{l}}(\mathbf{x}) \end{aligned} \quad (27)$$

Finally, the reflectance function is also a function of \mathbf{x} and $\hat{\mathbf{n}}$. Notice that ω_i and ω_o are function of \mathbf{x} . We use the symbolic differentiation to compute the derivative with respect to \mathbf{x} and $\hat{\mathbf{n}}$, respectively.

2.3. Accounting for visibility changes

Delaunoy and Prados [4] derived the gradient of a function that contains visibily terms for the case of surfaces represented as triangular meshes. In particular, they show how to differentiate energy functionals of the form,

$$E(\mathcal{S}) = \int_{\mathcal{S}} \langle \mathbf{g}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})), \hat{\mathbf{n}}(\mathbf{x}) \rangle v(\mathbf{x}, \mathbf{s}_0) \, dA(\mathbf{x}), \quad (28)$$

where \mathbf{s}_0 is the camera center for the line-of-sight imaging case. The visibility function checks whether a point is visible to the camera. Their derivative expression can be directly adapted to our setting in the case of confocal imaging $\mathbf{l} = \mathbf{s}$. In that case, we observe that the image formation model of Equation (1) of the main paper simplifies to

$$I(t; \mathbf{l}) = \int_{\mathcal{S}_{\text{NLOS}}} g(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) v(\mathbf{x}, \mathbf{l}) dA(\mathbf{x}), \quad (29)$$

where

$$g(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) = f_s(\hat{\mathbf{n}}(\mathbf{x}), \hat{\omega}_l(\mathbf{x}), \hat{\omega}_l(\mathbf{x})) \cdot \frac{(\langle -\hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle)^2 (\langle \hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle)^2}{\|\mathbf{x} - \mathbf{l}\|^4}. \quad (30)$$

If we define

$$\mathbf{g}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) = f_s(\hat{\mathbf{n}}(\mathbf{x}), \hat{\omega}_l(\mathbf{x}), \hat{\omega}_l(\mathbf{x})) \cdot \frac{(\langle -\hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle)^2 \langle \hat{\omega}_l(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle}{\|\mathbf{x} - \mathbf{l}\|^4} \cdot \hat{\omega}_l(\mathbf{x}), \quad (31)$$

then we can rewrite Equation (29) as

$$I(t; \mathbf{l}) = \int_{\mathcal{S}_{\text{NLOS}}} \langle \mathbf{g}(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})), \hat{\mathbf{n}}(\mathbf{x}) \rangle v(\mathbf{x}, \mathbf{l}) dA(\mathbf{x}), \quad (32)$$

which now matches the form of Equation (28).

We briefly describe the derivative expression derived by Delaunoy and Prados [4] for functionals of the form of Equation (28). The derivative contains three terms, terms due to the variation of the normal, term due to the tetrahedra of the visible adjacent triangles, and the term due to the movement of the crepuscular cone. Please refer to [4] for the details.

$$\frac{\partial E(\mathcal{S})}{\partial \mathbf{v}_i} = \mathbf{G}_i^{\text{norm}} + \mathbf{G}_i^{\text{int}} + \mathbf{G}_i^{\text{horiz}} \quad (33)$$

$$\mathbf{G}_i^{\text{norm}} = - \sum_{k \in J_i} \frac{\mathbf{e}_{k,i}}{2A_k} \wedge \int_{\mathcal{T}_k} P_{\hat{\mathbf{n}}_k}(D_{\hat{\mathbf{n}}_k} \mathbf{g}(\mathbf{x}, \mathbf{n}_k)^T \hat{\mathbf{n}}_k) v(\mathbf{x}, \mathbf{s}_0) ds \quad (34)$$

$$\mathbf{G}_i^{\text{int}} = \sum_{k \in J_i} \hat{\mathbf{n}}_k \int_{\mathcal{T}_k} \nabla \cdot \mathbf{g}(\mathbf{x}, \hat{\mathbf{n}}_k) \phi_i(\mathbf{x}) ds \quad (35)$$

$$\mathbf{G}_i^{\text{horiz}} = \sum_{\mathbf{H}_{i,k}} \frac{1}{2} \int_0^1 \{ [p(T(\mathbf{y}(u))) - p(\mathbf{y}(u))] \left[\frac{\mathbf{y}(u)}{\|\mathbf{y}(u)\|^4} \wedge \mathbf{H}_{i,k} \right] (1-u) \} du \quad (36)$$

$P_{\hat{\mathbf{n}}_k}(\cdot)$ is the projection on the orthogonal plane to $\hat{\mathbf{n}}_k$. A_k is the area of face k . $\mathbf{H}_{i,k}$ is the vector such that $[\mathbf{v}_i,$




Figure 8. **Normal smoothness.** For surface smoothness regularization, we optimize vertices around a face so that the normal is aligned with the weighted normal around the face.

$\mathbf{v}_i + \mathbf{H}_{i,k}$ is the edge of triangle k generating the horizon. \mathbf{y} corresponds to points sampled on the edge $\mathbf{H}_{i,k}$. $T(\mathbf{x})$ corresponds to the point located behind \mathbf{x} in the direction of the viewpoint. $\mathbf{g}(\mathbf{x}) = p(\mathbf{x}) \frac{\mathbf{x}}{x_z^3}$

2.4. Surface Regularization

When optimizing geometry, we follow Delaunoy and Prados [4] and augment the loss function $E(\mathbf{v}, \boldsymbol{\pi})$ of Equation (7) of the main paper, with a normal smoothing regularization term

$$\mathcal{R}(\mathbf{v}) = \sum_{k=1}^T A_k (1 - \hat{\mathbf{h}}_k \cdot \hat{\mathbf{n}}_k), \quad (37)$$

where A_k and $\hat{\mathbf{n}}_k$ are the area and face normal, respectively, of the k -th mesh triangle, and $\hat{\mathbf{h}}_k$ is the weighted average of the face normals of all triangles in its neighborhood N_k ,

$$\hat{\mathbf{h}}_k = \frac{\sum_{i \in N_k} A_i \hat{\mathbf{n}}_i}{\|\sum_{i \in N_k} A_i \hat{\mathbf{n}}_i\|}. \quad (38)$$

A visualization of this regularization term is shown in Figure 8. Its derivative with respect to mesh vertices is provided by Delaunoy and Prados [4].

3. Stochastic estimation and optimization

In this section, we provide details relating to the rendering and optimization algorithms discussed in Section 4.2 in the main paper.

3.1. Monte Carlo rendering with stratified area sampling

We first discuss the rendering algorithm we use to estimate transients $I(t; \mathbf{l}, \mathbf{s})$, and their derivatives with respect to surface, $\frac{\partial I}{\partial \mathbf{v}}$, and with respect to reflectance, $\frac{\partial I}{\partial \boldsymbol{\pi}}$ (Equations (1), (9), and (10) in the main paper, respectively). We present the algorithm in the context of estimating $I(t; \mathbf{l}, \mathbf{s})$, but the discussion applies exactly for the $\frac{\partial I}{\partial \mathbf{v}}$ and $\frac{\partial I}{\partial \boldsymbol{\pi}}$ cases. Algorithm 1 shows an overview of our rendering procedure.

We use Monte Carlo integration to approximate the surface integral of Equation (1): We first use any probability distribution μ on $\mathcal{S}_{\text{NLOS}}$ to sample a set of points $\{\mathbf{x}_j \in \mathcal{S}_{\text{NLOS}}, j = 1, \dots, J\}$. Then, we form an estimate as in Equation (11) of the main paper, reproduced here for convenience:

$$\langle I \rangle = \sum_{j=1}^J \frac{g(\mathbf{x}_j, \hat{\mathbf{n}}(\mathbf{x}_j)) v(\mathbf{x}_j, \mathbf{l}) v(\mathbf{x}_j, \mathbf{s})}{\mu(\mathbf{x}_j)}, \quad (39)$$

In the context of light transport simulation, this is referred to as *Monte Carlo rendering* [23, 19, 5]. We note that Monte Carlo integration by sampling points \mathbf{x} on $\mathcal{S}_{\text{NLOS}}$ is equivalent to the *area sampling* strategy described by Veach and Guibas [25]. The fact that we use area sampling is also the reason for the presence of the shading terms $\langle -\hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{l}) \rangle \langle \hat{\omega}_{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle / \|\mathbf{x} - \mathbf{l}\|^2$ and $\langle -\hat{\omega}_{\mathbf{s}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{s}) \rangle \langle \hat{\omega}_{\mathbf{s}}(\mathbf{x}), \hat{\mathbf{n}}(\mathbf{x}) \rangle / \|\mathbf{x} - \mathbf{s}\|^2$ in the estimate of Equation (39) (see Equation (9) in Veach and Guibas [25]).

We note that standard path sampling algorithms (e.g., path tracing [9], bidirectional path tracing [24], Metropolis light transport [26]) typically use *directional sampling* instead of area sampling. However, area sampling is advantageous in our setting for two reasons: First, when the NLOS surface $\mathcal{S}_{\text{NLOS}}$ is small compared to the LOS surface \mathcal{S}_{LOS} , directional sampling will result in a large number of missed rays, greatly reducing rendering efficiency. This has also been observed by Klein et al. [12], who report an order of magnitude efficiency improvement when using area sampling instead of directional sampling.

Second, area sampling facilitates *stratified sampling* [6]. We start by noting that, when $\mathcal{S}_{\text{NLOS}}$ is a triangular mesh, Equation (1) can be decomposed into a sum of per-triangle integrals,

$$I = \sum_{k=1}^T \int_{\mathcal{T}_k} g(\mathbf{x}, \hat{\mathbf{n}}(\mathbf{x})) v(\mathbf{x}, \mathbf{l}) v(\mathbf{x}, \mathbf{s}) dA(\mathbf{x}), \quad (40)$$

where \mathcal{T}_k is the k -th mesh triangle, and $\mathcal{S}_{\text{NLOS}} = \bigcup_{k=1}^T \mathcal{T}_k$. We can then estimate I as

$$\langle I \rangle = \sum_{k=1}^T \sum_{j=1}^{\lceil J/T \rceil} \frac{g(\mathbf{x}_j^k, \hat{\mathbf{n}}(\mathbf{x}_j^k)) v(\mathbf{x}_j^k, \mathbf{l}) v(\mathbf{x}_j^k, \mathbf{s})}{\mu(\mathbf{x}_j^k)}, \quad (41)$$

where for each k , $\{\mathbf{x}_j^k \in \mathcal{T}_k, j = 1, \dots, \lceil J/T \rceil\}$. When μ is the uniform distribution, Equation (39) uses J points \mathbf{x}_j uniformly sampled from the entire mesh $\mathcal{S}_{\text{NLOS}}$. By contrast, Equation (41) splits the J points into T equal-size sets, with points in the k -th set uniformly sampled from a *stratum* corresponding to the triangle \mathcal{T}_k . This stratification procedure can significantly reduce the variance of $\langle I \rangle$ [15, 21], by up to a factor of $1/T$ if the integrand of Equation (40) is

approximately constant within each triangle—as is the case for meshes with very fine triangulation. Empirically, we observed that stratified sampling reduces variance by an order of magnitude compared to uniform sampling.

An additional advantage of stratified sampling becomes evident when we consider the visibility terms v in Equations (39) and (41). For large meshes, visibility tests can account for the bulk of the rendering computational cost. Therefore, it is critical that these tests be performed using highly-optimized libraries specifically designed for this task [28, 17]. When using stratified sampling, the ray bundles $\{\mathbf{s} \rightarrow \mathbf{x}_j^k\}$ and $\{\mathbf{l} \rightarrow \mathbf{x}_j^k\}$ for all points of the same k can be treated as *coherent ray packets* for the purposes of visibility testing [20, 27]. Empirically, we found that this results in a fourfold rendering acceleration, highlighting another advantage of our triangle-based stratification procedure.

We conclude this discussion by noting one important disadvantage of our area sampling procedure (either stratified or otherwise): When the reflectance f_s of the underlying surface becomes very specular, then the estimate of Equation (39) becomes very inefficient. This is because, when \mathbf{x}_j is sampled uniformly on $\mathcal{S}_{\text{NLOS}}$ or a triangle \mathcal{T}_k , the half-vector corresponding to the path $\mathbf{l} \rightarrow \mathbf{x}_j \rightarrow \mathbf{s}$ will, with high probability, deviate significantly from the normal $\hat{\mathbf{n}}(\mathbf{x}_j)$. As a result, most of the sampled paths will have very low integrand values $g(\mathbf{x}_j, \hat{\mathbf{n}}(\mathbf{x}_j))$. The effect of this can be observed in Section 6.1 of the main paper, when optimizing for the NLOS case under very specular BRDF (GGX with $\alpha = 0.1$). In the future, we plan to overcome this by incorporating into our framework area sampling techniques that can account for BRDF effects through half-vector importance sampling [10]. An alternative direction is to combine area and directional sampling techniques using multiple importance sampling [25].

3.2. Optimization pipeline

In this section, we describe in detail our optimization pipeline. We show an overview in Algorithm 2, and discuss below the various components and strategies we use to make convergence faster and more robust.

Stochastic gradient descent. We alternately optimize for reflectance and surface vertices using stochastic gradient descent (SGD). In particular, we use the Adam [11] variant of SGD, in order to take advantage of the adaptive, per-parameter, learning-rate scheduling.

When optimizing for surface vertices, Adam ordinarily would maintain $3 \cdot V$ independent learning rates, where V is the number of vertices, and each vertex is a three-dimensional vector. Empirically, we found it beneficial to constrain all 3 coordinates of the same vertex to have the same learning rate, reducing the independent learning rates to V —one per 3D vertex vector. To achieve this, we modi-

Algorithm 1 Monte Carlo rendering for estimating transient I , or its derivatives $\frac{\partial I}{\partial v}$, $\frac{\partial I}{\partial \pi}$.

Require: Integrand function $G \in \{g, g_s, g_r\}$.

Require: Mesh $\mathcal{S}_{\text{NLOS}} = \bigcup_{k=1}^K \mathcal{T}_k$.

Require: Reflectance parameters π .

Require: Virtual source l and detector s .

Require: Number of rendering samples J .

```

1: function RENDER( $G, \mathcal{S}_{\text{NLOS}}, \pi, l, s, J$ )
2:    $\langle I \rangle \leftarrow 0$ . ▷ Initialize estimate.
3:   for  $k \in \{1, \dots, K\}$  do ▷ Select triangle  $\mathcal{T}_k$ .
4:     ▷ Compute triangle quantities.
5:      $A \leftarrow \text{ComputeTriangleArea}(\mathcal{T}_k)$ .
6:      $\hat{n} \leftarrow \text{ComputeTriangleNormal}(\mathcal{T}_k)$ .
7:     for  $j \in \{1, \dots, \lceil J/K \rceil\}$  do
8:       ▷ Uniformly sample a point on the triangle.
9:        $x \leftarrow \text{SamplePointUniformly}(\mathcal{T}_k)$ .
10:      ▷ Evaluate the function to be integrated.
11:       $y \leftarrow \text{EvaluateIntegrand}(G; \pi, x, \hat{n}, l, s)$ .
12:      ▷ Perform visibility tests.
13:       $v_l \leftarrow \text{EvaluateVisibility}(\mathcal{S}_{\text{NLOS}}, x, l)$ .
14:       $v_s \leftarrow \text{EvaluateVisibility}(\mathcal{S}_{\text{NLOS}}, x, s)$ .
15:       $\langle I \rangle \leftarrow \langle I \rangle + y \cdot v_l \cdot v_s \cdot A / \lceil J/K \rceil$ .
16:     end for
17:   end for
18:   return  $\langle I \rangle$ .
19: end function

```

fied Adam to update the learning rate for each vertex using the total magnitude of the 3D gradient for that vertex’s coordinates. For the learning rate update rules, we refer to the original Adam publication [11].

Coarse-to-fine surface optimization. We use a coarse-to-fine scheme to accelerate the surface optimization. We begin with a mesh of a resolution equivalent to the scanning resolution on the visible wall \mathcal{S}_{LOS} , as produced by a volumetric reconstruction procedure. We then iterate between evolving the mesh for N gradient-descent iterations, and increasing the number of vertices by 1.25. As discussed in the main paper, to increase the number of vertices, we first use El Topo [3] to create a non-intersecting version of the evolved mesh, then perform isotropic remeshing with an increased number of target vertices.

Continuation. As discussed in the main paper and in Section 2.4 of this supplement, we use a regularized loss function for surface optimization,

$$E(v, \pi) + \lambda \mathcal{R}(v). \quad (42)$$

Following common practice in regularized optimization [7], we adopt a *continuation* scheme and gradually decrease the weight λ of the regularization term: We start with an initial value λ_0 , and decrease λ by a factor of 1.25 every time we

increase the surface resolution. This gives more and more emphasis on the data term in stead of the regularization.

Rendering sample budgeting. We follow the so-called *increasing precision* strategy proposed by Pfeiffer and Sato [18]: As the optimization proceeds, we increase the number of samples J used for rendering the gradients with respect to surface and reflectance. Intuitively, as the optimization gets closer to a local minimum, and therefore the true gradient becomes smaller, the gradient-descent procedure becomes more sensitive to variance in the gradient estimates; therefore, reducing this variance by increasing the rendering samples can help convergence. In our implementation, we use an initial number of J_0 samples, and we increase this by a factor of 1.25 every time we alternate between surface and reflectance optimization.

4. Additional Experiments

We show additional results from experiments using synthetic data.

Quantitative evaluation of shape reconstructions. To quantify the improvement in surface reconstruction in Figure 5 of the main paper, we compare in Table 9 initialization and final reconstruction to ground-truth using two metrics: 1) the loss function of Equation (7) in the main paper; 2) the mean vertex-to-surface error. In all cases, our method provides significant improvement, often by one or more orders of magnitude. We note that mean vertex-to-surface error is not representative of the difference in resolution between the two meshes, because of the vastly different number of vertices in the initialization and final reconstructions. The following extreme example illustrates the issue: A mesh consisting of just one triangle whose three vertices happened to be exactly on the ground-truth surface would have zero error, despite not capturing the underlying shape.

Number of measurements. We compare reconstructions obtained by the technique of O’Toole et al. [16] and our procedure, using measurements from confocal scans ($l = s$) at 5 different scanning resolutions within the same scanning area on the visible wall \mathcal{S}_{LOS} . In particular, we consider resolutions: 16×16 , 32×32 , 64×64 , 128×128 , and 256×256 scan points.

Figure 10 shows the results. We note that, for the 128×128 and 256×256 cases, we found it beneficial to initialize our optimization procedure using a reconstruction obtained by applying the algorithm of O’Toole et al. [16] on only 64×64 measurements. Despite the loss in detail, this initialization provided better coverage on the NLOS surface.

We observe that, even in the case of 16×16 measured transients, our algorithm can recover some of the detail on the surface of the ground-truth mesh (e.g., texture on the

		bunny	armadillo	bear	bust	einstein	skull	soap
Loss function	init	$5.01 \cdot 10^{-4}$	$1.4 \cdot 10^{-3}$	$1.7 \cdot 10^{-3}$	$7.17 \cdot 10^{-4}$	$3.8 \cdot 10^{-3}$	$1.3 \cdot 10^{-3}$	$2.4 \cdot 10^{-3}$
	result	$1.13 \cdot 10^{-5}$	$2.3 \cdot 10^{-5}$	$2.57 \cdot 10^{-5}$	$6.49 \cdot 10^{-5}$	$9.65 \cdot 10^{-5}$	$4.31 \cdot 10^{-5}$	$2.08 \cdot 10^{-4}$
Vertex distance	init	$3.1 \cdot 10^{-3}$	$6.6 \cdot 10^{-3}$	$6 \cdot 10^{-3}$	$3.4 \cdot 10^{-3}$	$8.1 \cdot 10^{-3}$	$6.8 \cdot 10^{-3}$	$2.3 \cdot 10^{-3}$
	result	$9.95 \cdot 10^{-4}$	$3.4 \cdot 10^{-3}$	$1.2 \cdot 10^{-3}$	$1.9 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$3.55 \cdot 10^{-4}$

Figure 9. **Surface reconstruction metrics:** For each shape in Figure 5 of the main paper, we compare the initialization and final reconstruction to ground-truth using two metrics: 1) the loss function of Equation (7) in the main paper (*loss function*); 2) the mean vertex-to-surface error (*vertex distance*).

bunny’s leg), despite the fact that the initialization has no visible details. We also see that, with just 32×32 measured transients, our algorithm can reconstruct more detail than what is possible using the volumetric reconstruction from O’Toole et al. [16] with 256×256 measured transients. That is, we can obtain a more detailed reconstruction even though we are using 64 times fewer measurements.

Noise levels. In this experiment, we evaluate the performance of our method as a function of the amount of noise contaminating our transient measurements. We use synthetic transients, and simulate SPAD noise using the method of Hernandez et al. [8]. The simulated noise contains three components, ambient noise, Poisson noise, and SPAD jitter.

Figure 11 shows reconstructions under three different noise levels, corresponding to different number of laser pulses M . Even though performance deteriorates as noise increases, we see that our proposed method can take SPAD jitter into account and is robust to ambient and Poisson noise.

We note that, as the noise level increases, the level of detail we can recover decreases. In this case, continuing to increase the spatial resolution of the mesh can be counter-productive, as the amount of noise means that gradient estimates have very high variance: We simply have not measured enough photons in order for our measurements to sufficiently regularize the surface S_{NLOS} . This is shown in Figure 12, where we show the evolution of the mesh as we continue to upsample it at higher spatial resolutions. In practice, this implies that, for higher noise levels, we need to stop the optimization procedure at earlier resolution levels.

References

- [1] Ben Appleton and Hugues Talbot. Globally minimal surfaces by continuous maximal flows. *IEEE transactions on pattern analysis and machine intelligence*, 28(1):106–118, 2006. 2
- [2] James Arvo. *Analytic methods for simulated light transport*. PhD thesis, PhD thesis, Yale University, 1995. 1
- [3] Tyson Brochu and Robert Bridson. El topo, 2009. <https://www.cs.ubc.ca/labs/imager/tr/2009/eltopo/eltopo.html>. 5
- [4] Amaël Delaunoy and Emmanuel Prados. Gradient flows for optimizing triangular mesh-based surfaces: Applications to 3d reconstruction problems dealing with visibility. *IJCV*, 95(2):100–123, 2011. 2, 3
- [5] Philip Dutré, Kavita Bala, and Philippe Bekaert. *Advanced global illumination*. AK Peters, Ltd., 2006. 4
- [6] George Fishman. *Monte Carlo: concepts, algorithms, and applications*. Springer Science & Business Media, 1996. 4
- [7] Elaine T Hale, Wotao Yin, and Yin Zhang. Fixed-point continuation for l_1 -minimization: Methodology and convergence. *SIAM Journal on Optimization*, 19(3):1107–1130, 2008. 5
- [8] Quercus Hernandez, Diego Gutierrez, and Adrian Jarabo. A computational model of a single-photon avalanche diode sensor for transient imaging. *arXiv*, preprint arXiv:1703.02635, 2017. 6
- [9] James T Kajiya. The rendering equation. In *ACM Siggraph Computer Graphics*, volume 20, pages 143–150, 1986. 4
- [10] Anton S Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. The natural-constraint representation of the path space for efficient light transport simulation. *ACM Transactions on Graphics (TOG)*, 33(4):102, 2014. 4
- [11] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 4, 5
- [12] Jonathan Klein, Martin Laurenzis, Dominik L Michels, and Matthias B Hullin. A quantitative platform for non-line-of-sight imaging problems. In *Proceedings BMVC*, pages 1–12, 2018. 4
- [13] Jan J Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13(3):321–330, 1984. 2
- [14] Hsueh-Ti Derek Liu, Michael Tao, and Alec Jacobson. Paparazzi: Surface editing by way of multi-view image processing. *ACM TOG*, 2018. 2
- [15] Don P Mitchell. Consequences of stratified sampling in graphics. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 277–280. ACM, 1996. 4

Algorithm 2 Surface and reflectance optimization pipeline.

Require: Initial NLOS vertices v_0 .
Require: Initial number of NLOS vertices V_0 .
Require: Initial NLOS triangles T_0 .
Require: Initial reflectance parameters π_0 .
Require: Initial learning rates $\eta_{s,0}, \eta_{r,0}$.
Require: Measured transient \tilde{I} .
Require: Virtual source l and detector s .
Require: Initial regularization weight λ_0 .
Require: Initial number of rendering samples J_0 .
Require: Number of iterations N .

Initialization.

```
1:  $v \leftarrow v_0$ .  
2:  $V \leftarrow V_0$ .  
3:  $T \leftarrow T_0$ .  
4:  $\pi \leftarrow \pi_0$ .  
5:  $\lambda \leftarrow \lambda_0$ .  
6:  $\eta_s \leftarrow \eta_{s,0}$ .  
7:  $\eta_r \leftarrow \eta_{r,0}$ .  
8:  $J \leftarrow J_0$ .  
9: while not converged do
```

Gradient-descent optimization.

```
10:  $\mathcal{S}_{\text{NLOS}} \leftarrow \text{CreateMesh}(v, T)$ .  
11:  $\triangleright$  Update reflectance.  
12:  $\pi \leftarrow \text{OptReflectance}(\pi, \eta_r, \mathcal{S}_{\text{NLOS}}, \tilde{I}, l, s, N, J)$ .  
13:  $\triangleright$  Update surface.  
14:  $v' \leftarrow \text{OptSurface}(v, \eta_s, \pi, T, \tilde{I}, l, s, \lambda, N, J)$ .
```

Geometry processing.

```
15:  $\triangleright$  Deform mesh.  
16:  $(v, T) \leftarrow \text{ElTopo}(v, v', T)$ .  
17:  $\triangleright$  Increase mesh resolution.  
18:  $V \leftarrow V \cdot 1.25$ .  
19:  $(v, T) \leftarrow \text{IsoRemesh}(v, T, V)$ .
```

Update parameters.

```
20:  $\lambda \leftarrow \lambda / 1.25$ .  $\triangleright$  Continuation.  
21:  $J \leftarrow J \cdot 1.25$ .  $\triangleright$  Increasing precision.  
22: end while
```

- [16] Matthew O'Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696):338, 2018. [5](#), [6](#)
- [17] Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. Optix: a general purpose ray tracing engine. *ACM TOG*, 29(4):66, 2010. [4](#)
- [18] Gustavo T Pfeiffer and Yoichi Sato. On stochastic optimization methods for monte carlo least-squares problems. *arXiv preprint arXiv:1804.10079*, 2018. [5](#)
- [19] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically*

Algorithm 3 Adam routines for optimizing NLOS mesh vertices and reflectance.

Require: Initial NLOS vertices v_0 .
Require: Initial learning rates η_0 .
Require: Reflectance parameters π .
Require: NLOS triangles T .
Require: Measured transient \tilde{I} .
Require: Virtual source l and detector s .
Require: Regularization weight λ .
Require: Number of iterations N .
Require: Number of rendering samples J .

```
1: function OPTSURFACE( $v_0, \eta_0, \pi, T, \tilde{I}, l, s, \lambda, N, J$ )  
2:    $v \leftarrow v_0$ .  $\triangleright$  Initialize estimate.  
3:    $\eta \leftarrow \eta_0$ .  $\triangleright$  Initialize learning rates.  
4:   for  $n \in \{1, \dots, N\}$  do  
5:      $\mathcal{S}_{\text{NLOS}} \leftarrow \text{CreateMesh}(v, T)$ .  
6:      $\triangleright$  Render required quantities.  
7:      $I \leftarrow \text{Render}(g; \mathcal{S}_{\text{NLOS}}, \pi, l, s, J)$ .  
8:      $I_v \leftarrow \text{Render}(g_s; \mathcal{S}_{\text{NLOS}}, \pi, l, s, J)$ .  
9:      $R_v \leftarrow \text{ComputeReguGradient}(\mathcal{S}_{\text{NLOS}})$ .  
10:     $\triangleright$  Compute gradient.  
11:     $g \leftarrow (I - \tilde{I}) I_v + \lambda R_v$   
12:     $\triangleright$  Perform Adam updates.  
13:     $v \leftarrow \text{AdamUpdateParameters}(v, \eta, g)$   
14:     $\eta \leftarrow \text{AdamUpdateLearningRates}(v, \eta, g)$   
15:  end for  
16:  return  $v$ .  
17: end function
```

Require: Initial reflectance parameters π_0 .
Require: Initial learning rates η_0 .
Require: NLOS mesh $\mathcal{S}_{\text{NLOS}}$.
Require: Measured transient \tilde{I} .
Require: Virtual source l and detector s .
Require: Number of iterations N .
Require: Number of rendering samples J .

```
18: function OPTREFLECTANCE( $\pi_0, \eta_0, \mathcal{S}_{\text{NLOS}}, \tilde{I}, l, s, N, J$ )  
19:    $\pi \leftarrow \pi_0$ .  $\triangleright$  Initialize estimate.  
20:    $\eta \leftarrow \eta_0$ .  $\triangleright$  Initialize learning rates.  
21:   for  $n \in \{1, \dots, N\}$  do  
22:     $\triangleright$  Render required quantities.  
23:     $I \leftarrow \text{Render}(g; \mathcal{S}_{\text{NLOS}}, \pi, l, s, J)$ .  
24:     $I_\pi \leftarrow \text{Render}(g_r; \mathcal{S}_{\text{NLOS}}, \pi, l, s, J)$ .  
25:     $\triangleright$  Compute gradient.  
26:     $g \leftarrow (I - \tilde{I}) I_\pi$   
27:     $\triangleright$  Perform Adam updates.  
28:     $\pi \leftarrow \text{AdamUpdateParameters}(\pi, \eta, g)$   
29:     $\eta \leftarrow \text{AdamUpdateLearningRates}(\pi, \eta, g)$   
30:  end for  
31:  return  $\pi$ .  
32: end function
```

based rendering: From theory to implementation. Morgan Kaufmann, 2016. 4

- [20] Matt Pharr, Craig Kolb, Reid Gershbein, and Pat Hanrahan. Rendering complex scenes with memory-coherent ray tracing. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 101–108, 1997. 4
- [21] Adrien Pilleboue, Gurprit Singh, David Coeurjolly, Michael Kazhdan, and Victor Ostromoukhov. Variance analysis for monte carlo integration. *ACM Transactions on Graphics (TOG)*, 34(4):124, 2015. 4
- [22] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR*, 2006. 2
- [23] Eric Veach. *Robust monte carlo methods for light transport simulation*. Stanford University Stanford, 1998. 4
- [24] Eric Veach and Leonidas Guibas. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*, pages 145–167. Springer, 1995. 4
- [25] Eric Veach and Leonidas J Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 419–428. ACM, 1995. 4
- [26] Eric Veach and Leonidas J Guibas. Metropolis light transport. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 65–76. ACM Press/Addison-Wesley Publishing Co., 1997. 4
- [27] Ingo Wald, Philipp Slusallek, Carsten Benthin, and Markus Wagner. Interactive rendering with coherent ray tracing. In *Computer graphics forum*, volume 20, pages 153–165. Wiley Online Library, 2001. 4
- [28] Ingo Wald, Sven Woop, Carsten Benthin, Gregory S Johnson, and Manfred Ernst. Embree: a kernel framework for efficient cpu ray tracing. *ACM TOG*, 33(4):143, 2014. 4
- [29] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206. Eurographics Association, 2007. 1

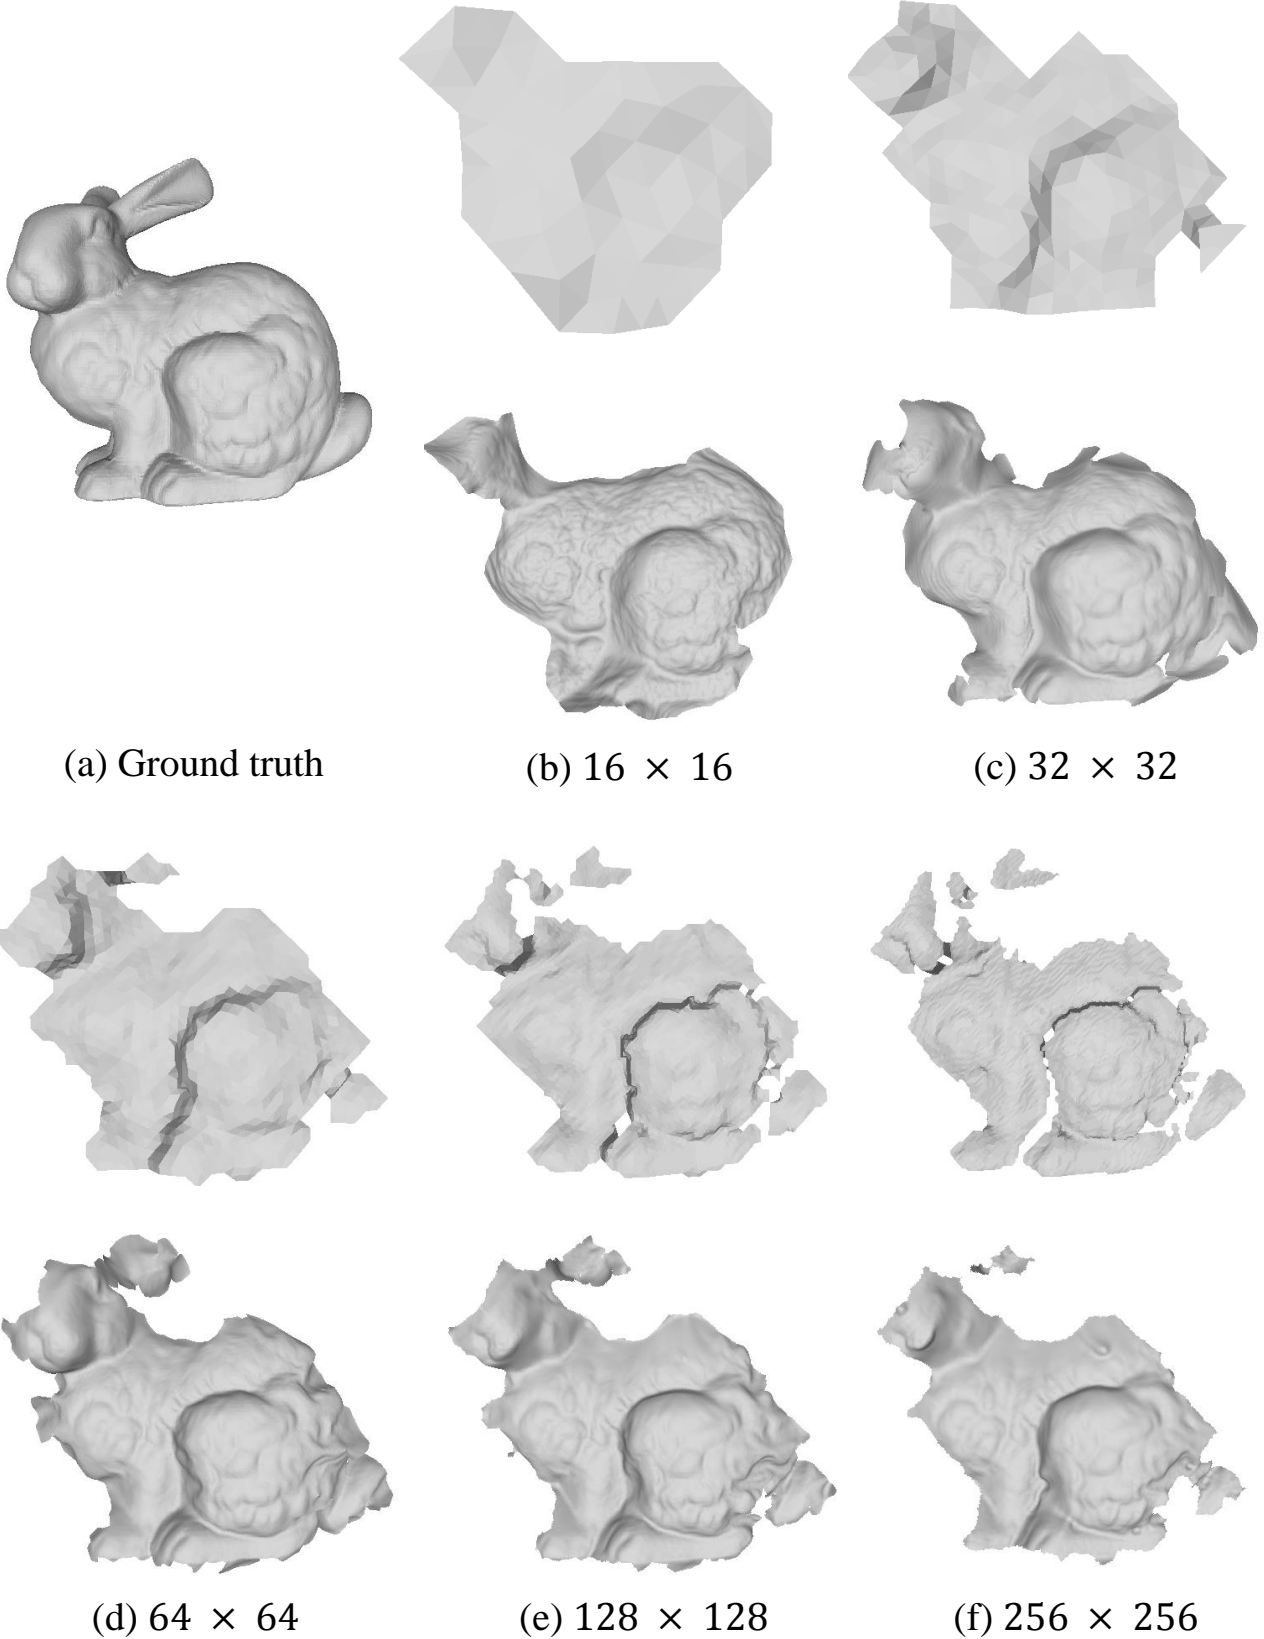


Figure 10. **Surface optimization using different number of measurements:** We perform experiments for different numbers of measured transients (i.e., scanned points on the LOS surface \mathcal{S}_{LOS}). When the number of measurements is very small, the initialization does not recover any discernible shape, whereas our surface optimization framework still recovers details of the ground-truth mesh. As the number of measurements increases, the level of detail of both the initialization and our reconstruction increases; in all cases our method significantly improves the final surface reconstruction.

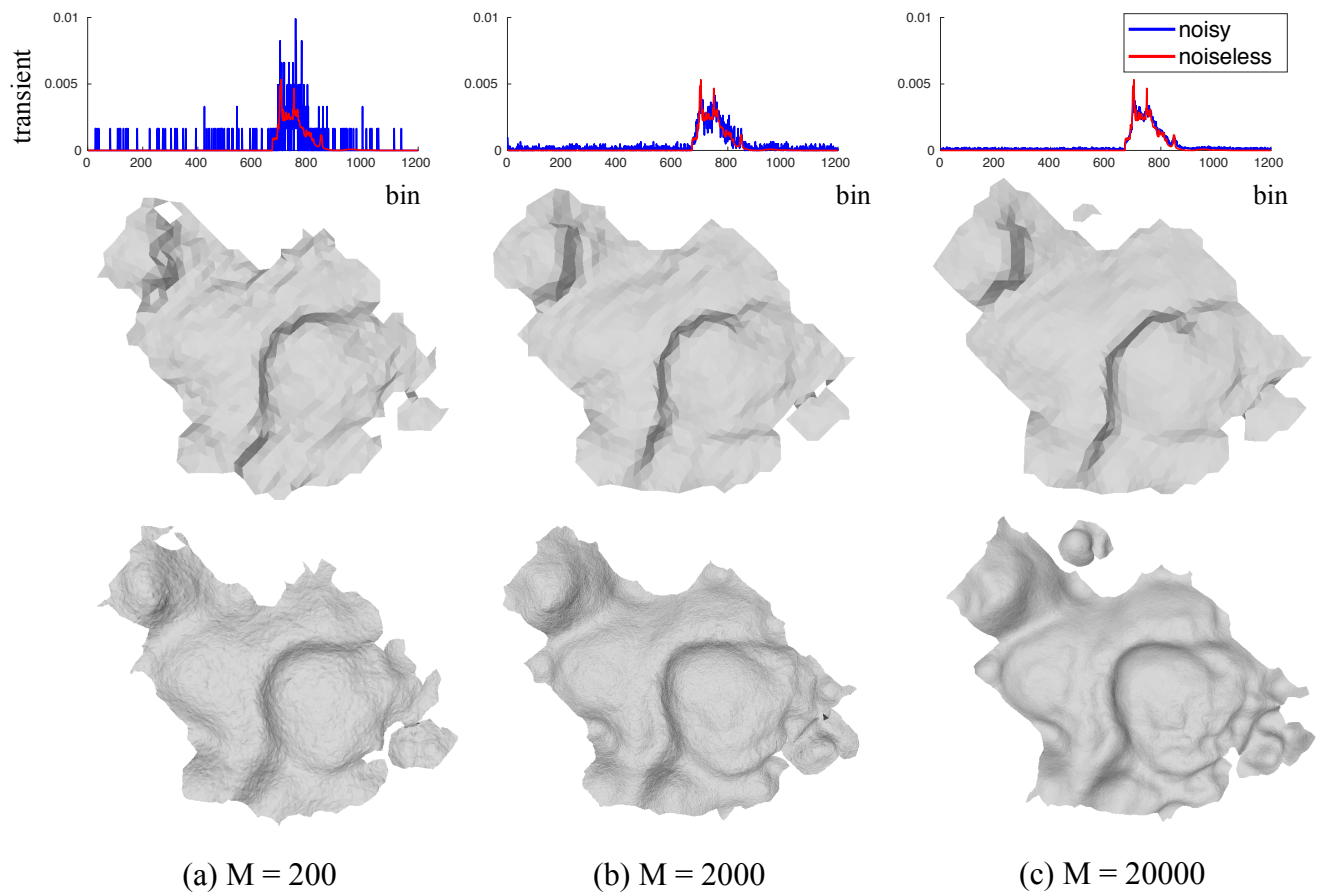


Figure 11. **Noise experiment:** We simulate measurement noise by including ambient light, Poisson noise, and SPAD jitter. (Top) Sample of the noisy transients. (Middle) Initialization (Bottom) Our recovered results.




Figure 12. **Noise experiment:** We show the progression of the surface optimization for the case of the highest noise ($M = 200$). The number of mesh vertices increases with iteration. As the spatial resolution of the surface increases, the effect of noise is to reduce local surface quality, even if the global surface shape remains acceptable.