Geometric Antithetic Sampling for Rendering Spatiotemporally Modulated Light Supplementary Material

JUHYEON KIM, Dartmouth College, USA
JACK CUI, Dartmouth College, USA
WOJCIECH JAROSZ, Dartmouth College, USA
ADITHYA PEDIREDLA, Dartmouth College, USA

CCS Concepts: • **Computing methodologies** → **Ray tracing**; **Computational photography**.

Additional Key Words and Phrases: Structured Light, Continuous-wave Time-of-Flight Imaging, Modulated Light Rendering

ACM Reference Format:
Juhyeon Kim, Jack Cui, Wojciech Jarosz, and Adithya Pediredla. 2026. Geometric Antithetic Sampling for Rendering Spatiotemporally Modulated Light Supplementary Material. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26), December 01–04, 2026, Kuala Lumpur, Malaysia.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3829340.3842310

1 Unbiasedness with Paired Inverse Mappings

We prove the property discussed in Section 3.2 of the main manuscript. Assume that the mapping T is chosen randomly from two mappings T_1, T_2 that are inverses of each other:

$$T_2 = T_1^{-1}. \tag{1}$$

We select either mapping with equal probability. Let $p(x)$ denote the density of the primal sample, and let $q_i(x)$ denote the density induced by applying T_i to a primal sample:

$$q_i(x) = p\big(T_i^{-1}(x)\big)\left|\det\left(\frac{\partial T_i^{-1}(x)}{\partial x}\right)\right|. \tag{2}$$

Suppose that T_1 is selected, producing the antithetic sample

$$x_a = T_1(x_p). \tag{3}$$

Instead of evaluating an additional mapped point for the MIS weight of x_p, we use the reverse density $q_2(x_p)$ and define

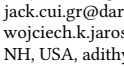
$$\widehat{I}_1 = \frac{p(x_p)}{p(x_p)+q_2(x_p)}\frac{f(x_p)}{p(x_p)} + \frac{q_1(x_a)}{p(x_a)+q_1(x_a)}\frac{f(x_a)}{q_1(x_a)}. \tag{4}$$

Taking the expectation over $x_p \sim p$, with $x_a \sim q_1$, gives

$$\mathbb{E}[\widehat{I}_1] = \int f(x)\left[\frac{p(x)}{p(x)+q_2(x)} + \frac{q_1(x)}{p(x)+q_1(x)}\right]\mathrm{d}x. \tag{5}$$

This conditional estimator is not generally unbiased by itself.

Authors' Contact Information: Juhyeon Kim, Dartmouth College, Hanover, NH, USA, juhyeon.kim.gr@dartmouth.edu; Jack Cui, Dartmouth College, Hanover, NH, USA, jack.cui.gr@dartmouth.edu; Wojciech Jarosz, Dartmouth College, Hanover, NH, USA, wojciech.k.jarosz@dartmouth.edu; Adithya Pediredla, Dartmouth College, Hanover, NH, USA, adithya.k.pediredla@dartmouth.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.
SA Conference Papers '26, Kuala Lumpur, Malaysia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2842-6/2026/12
https://doi.org/10.1145/3829340.3842310

Similarly, if T_2 is selected, we define

$$\widehat{I}_2 = \frac{p(x_p)}{p(x_p)+q_1(x_p)}\frac{f(x_p)}{p(x_p)} + \frac{q_2(x_a)}{p(x_a)+q_2(x_a)}\frac{f(x_a)}{q_2(x_a)}, \tag{6}$$

which has expectation

$$\mathbb{E}[\widehat{I}_2] = \int f(x)\left[\frac{p(x)}{p(x)+q_1(x)} + \frac{q_2(x)}{p(x)+q_2(x)}\right]\mathrm{d}x. \tag{7}$$

Because T_1 and T_2 are selected with equal probability, the expectation of the overall estimator is

$$\begin{aligned}
\mathbb{E}[\widehat{I}] &= \frac{1}{2}\mathbb{E}[\widehat{I}_1] + \frac{1}{2}\mathbb{E}[\widehat{I}_2] && (8)\\
&= \frac{1}{2}\int f(x)\left[\frac{p(x)}{p(x)+q_2(x)} + \frac{q_2(x)}{p(x)+q_2(x)}\right]\mathrm{d}x \\
&\quad + \frac{1}{2}\int f(x)\left[\frac{p(x)}{p(x)+q_1(x)} + \frac{q_1(x)}{p(x)+q_1(x)}\right]\mathrm{d}x && (9)\\
&= \frac{1}{2}\int f(x)\mathrm{d}x + \frac{1}{2}\int f(x)\mathrm{d}x && (10)\\
&= \int f(x)\mathrm{d}x. && (11)
\end{aligned}$$

In the second equality, we reorder the second integrand terms from $\mathbb{E}[\widehat{I}_1]$ and $\mathbb{E}[\widehat{I}_2]$ so that terms with the same denominator are grouped together. Each pair then sums to one. Therefore, although the conditional estimators $\widehat{I}_1$ and $\widehat{I}_2$ are not necessarily unbiased individually, the overall estimator is unbiased after averaging over the equally likely choices of T_1 and T_2.

We present the derivation for two mappings, but the same argument extends to any collection of inverse mapping pairs selected with equal probability. This extension is useful for handling 2D patterns, such as the top, bottom, left, and right translations of a checkerboard pattern.

2 Implementation Details

This section provides implementation details and derivations omitted from the main manuscript.

2.1 CW-ToF Coordinate-Specific Derivatives

Newton iterations and mapping Jacobian determinant evaluation require partial derivatives of F with respect to ξ and ξ', which depend on the path-length gradient $\nabla\ell$ and Hessian $\nabla^2\ell$, as shown in Kim et al. [2026]. We first evaluate these derivatives with respect to the vertex position $\mathbf{p} \in \mathbb{R}^3$. For $\ell(\mathbf{p}) = \|\mathbf{p}_1 - \mathbf{p}\| + \|\mathbf{p}_2 - \mathbf{p}\|$, the gradient and Hessian are

$$\nabla_{\mathbf{p}}\ell = -(\mathbf{d}_1 + \mathbf{d}_2) = -\mathbf{s}, \qquad \nabla^2_{\mathbf{p}}\ell = \sum_{i=1}^{2}\frac{I - \mathbf{d}_i\mathbf{d}_i^T}{\|\mathbf{p}_i - \mathbf{p}\|} = A, \tag{12}$$

where $\mathbf{d}_i = (\mathbf{p}_i - \mathbf{p})/\|\mathbf{p}_i - \mathbf{p}\|$ are unit direction vectors and $\mathbf{s}$ is the unnormalized half-vector. We then transform these derivatives to the coordinates $\xi \in \mathbb{R}^2$ using the chain rule:

$$\nabla_\xi \ell = J^T(-\mathbf{s}), \qquad \nabla^2_\xi \ell = J^T A J + Q, \tag{13}$$

where $J = \partial\mathbf{p}/\partial\xi \in \mathbb{R}^{3\times 2}$ is the coordinate Jacobian, whose columns span the local tangent plane, and $Q \in \mathbb{R}^{2\times 2}$ accounts for the curvature of the parameterization:

$$Q_{ij} = (-\mathbf{s})^T \left(\frac{\partial^2 \mathbf{p}}{\partial\xi_i \partial\xi_j}\right), \qquad i, j \in \{1, 2\}. \tag{14}$$

The same expressions apply to ξ' at the mapped point. Following the general formulation of Kim et al. [2026], we derive J and Q for the parameterizations considered below.

We broadly divide these parameterizations into two groups: local coordinates, which avoid ray tracing during Newton iterations but may require a final visibility or validity check, and global coordinates, which require ray tracing at every iteration. Local coordinates are computationally inexpensive but cannot follow the surface across primitive boundaries. Global coordinates support such crossings at the cost of additional ray tracing.

(Local) Barycentric Triangle Coordinates. For an intersection on a triangle with world-space vertices $\mathbf{v}_0$, $\mathbf{v}_1$, and $\mathbf{v}_2$, we parameterize the surface point using barycentric coordinates:

$$\xi = (\beta_1, \beta_2), \qquad \beta_0 = 1 - \beta_1 - \beta_2. \tag{15}$$

The surface position is

$$\mathbf{p}(\beta_1, \beta_2) = \beta_0 \mathbf{v}_0 + \beta_1 \mathbf{v}_1 + \beta_2 \mathbf{v}_2 = \mathbf{v}_0 + \beta_1 \mathbf{e}_1 + \beta_2 \mathbf{e}_2, \tag{16}$$

where $\mathbf{e}_1 = \mathbf{v}_1 - \mathbf{v}_0$ and $\mathbf{e}_2 = \mathbf{v}_2 - \mathbf{v}_0$ are the triangle edge vectors. The coordinate Jacobian is

$$J_{\text{bary}} = \frac{\partial \mathbf{p}}{\partial \xi} = \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 \end{bmatrix}. \tag{17}$$

Because this parameterization is affine, all second derivatives of the surface position vanish:

$$\frac{\partial^2 \mathbf{p}}{\partial\xi_i \partial\xi_j} = \mathbf{0}, \qquad Q_{\text{bary}} = \mathbf{0}. \tag{18}$$

A point lies within the triangle if and only if

$$\beta_1 \geq 0, \qquad \beta_2 \geq 0, \qquad \beta_1 + \beta_2 \leq 1. \tag{19}$$

These conditions allow us to verify triangle membership without ray tracing, although shadow rays are still required to check visibility.

(Global) Hemispherical Coordinates. Alternatively, we parameterize $\mathbf{p}$ by the direction of a ray originating from $\mathbf{o}$, which we set to $\mathbf{p}_1$. Let $\{\mathbf{t}, \mathbf{b}, \mathbf{n}_o\}$ be an orthonormal frame at $\mathbf{o}$ and let $\xi = (\mu, \psi)$, where $\mu = \cos\theta \in [0, 1]$ and $\psi = \phi/(2\pi) \in [0, 1)$. The ray direction is

$$\mathbf{d}(\mu, \psi) = \sqrt{1-\mu^2}\cos(2\pi\psi)\mathbf{t} + \sqrt{1-\mu^2}\sin(2\pi\psi)\mathbf{b} + \mu\mathbf{n}_o. \tag{20}$$

The surface point is

$$\mathbf{p}(\mu, \psi) = \mathbf{o} + \rho(\mu, \psi)\mathbf{d}(\mu, \psi), \tag{21}$$

where the intersection distance ρ is obtained by ray tracing at each iteration.

To evaluate the derivatives, we locally approximate the intersected surface by a plane with normal $\mathbf{n}$, neglecting surface curvature. For compactness, define

$$\mathbf{d}_i = \frac{\partial \mathbf{d}}{\partial \xi_i}, \qquad \mathbf{d}_{ij} = \frac{\partial^2 \mathbf{d}}{\partial\xi_i \partial\xi_j}, \tag{22}$$

and

$$a_0 = \mathbf{n}^T \mathbf{d}, \qquad a_i = \mathbf{n}^T \mathbf{d}_i, \qquad a_{ij} = \mathbf{n}^T \mathbf{d}_{ij}. \tag{23}$$

Implicit differentiation of the ray–plane intersection yields the Jacobian columns

$$J_i = \frac{\partial \mathbf{p}}{\partial \xi_i} = \rho\left(\mathbf{d}_i - \frac{a_i}{a_0}\mathbf{d}\right), \tag{24}$$

giving

$$J_{\text{hemi}} = \begin{bmatrix} J_\mu & J_\psi \end{bmatrix}. \tag{25}$$

The second derivatives are

$$\mathbf{p}_{ij} \coloneqq \frac{\partial^2 \mathbf{p}}{\partial\xi_i \partial\xi_j} = \rho\left(\mathbf{d}_{ij} - \frac{a_i}{a_0}\mathbf{d}_j - \frac{a_j}{a_0}\mathbf{d}_i - \frac{a_{ij}}{a_0}\mathbf{d} + 2\frac{a_i a_j}{a_0^2}\mathbf{d}\right). \tag{26}$$

Thus, the coordinate-dependent Hessian term is

$$(Q_{\text{hemi}})_{ij} = -\mathbf{s}^T \mathbf{p}_{ij}, \qquad i, j \in \{\mu, \psi\}. \tag{27}$$

Unlike the local tangent-plane parameterizations, Q_{hemi} is generally nonzero even for planar surfaces because the intersection point depends nonlinearly on the hemispherical coordinates.

(Global) Initial-Ray-Centered Hemispherical Chart. The conventional (μ, ψ) parameterization has a coordinate singularity at the pole and a discontinuity in the coordinate representation at the azimuthal seam. These features can cause numerical difficulties during Newton iterations. We empirically find that a fixed chart centered on the initial ray direction often improves robustness.

Let $\mathbf{d}_0$ denote the initial ray direction and let $\{\mathbf{t}_0, \mathbf{b}_0, \mathbf{d}_0\}$ be an orthonormal frame. For $\xi = (u, v) \in \mathbb{R}^2$, define

$$\mathbf{p}(u, v) = \mathbf{d}_0 + u\mathbf{t}_0 + v\mathbf{b}_0, \qquad q(u, v) = 1 + u^2 + v^2. \tag{28}$$

The unit ray direction is

$$\mathbf{d}(u, v) = \frac{\mathbf{p}(u, v)}{\sqrt{q(u, v)}}. \tag{29}$$

This chart covers the open hemisphere centered on $\mathbf{d}_0$ and places the initial ray at the nonsingular coordinate $\xi = (0, 0)$. The frame and chart center remain fixed throughout the Newton iterations.

The first derivatives of the direction are

$$\mathbf{d}_u = q^{-1/2}\mathbf{t}_0 - uq^{-3/2}\mathbf{p}, \tag{30}$$

$$\mathbf{d}_v = q^{-1/2}\mathbf{b}_0 - vq^{-3/2}\mathbf{p}. \tag{31}$$

The second derivatives are

$$\mathbf{d}_{uu} = -\left(2u\mathbf{t}_0 + \mathbf{p}\right) q^{-3/2} + 3u^2\mathbf{p}q^{-5/2}, \tag{32}$$

$$\mathbf{d}_{uv} = -\left(v\mathbf{t}_0 + u\mathbf{b}_0\right) q^{-3/2} + 3uv\mathbf{p}q^{-5/2}, \tag{33}$$

$$\mathbf{d}_{vv} = -\left(2v\mathbf{b}_0 + \mathbf{p}\right) q^{-3/2} + 3v^2\mathbf{p}q^{-5/2}. \tag{34}$$

The surface point is again obtained by ray tracing:

$$\mathbf{p}(u, v) = \mathbf{o} + \rho(u, v)\mathbf{d}(u, v). \tag{35}$$

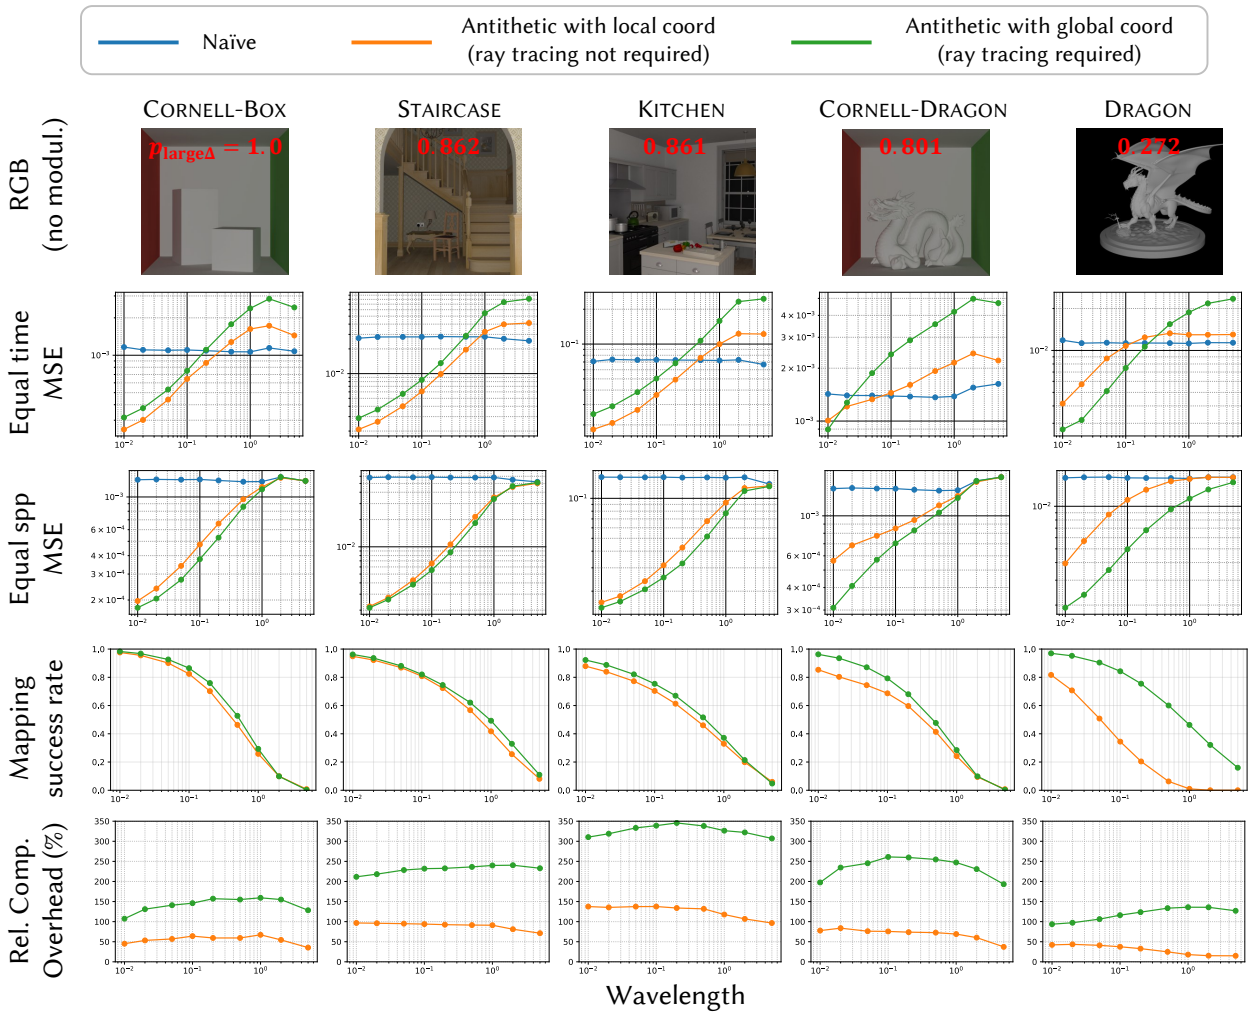


Fig. 1. Comparison of local barycentric coordinates and the global initial-ray-centered hemispherical chart for CW-ToF rendering with different wavelengths. Rows show scene appearance, equal-time MSE, equal-SPP MSE, mapping success rate, and computational overhead. Using global coordinates improves mapping success and equal-SPP accuracy, particularly for finely tessellated geometry, but requires ray tracing at each iteration, resulting in higher cost.

The ray–plane intersection derivatives in Eqs. (24) and (26) therefore apply directly using $i, j \in \{u, v\}$ and the direction derivatives above. In particular,

$$J_i = \rho \left(\mathbf{d}_i - \frac{a_i}{a_0} \mathbf{d} \right), \qquad a_0 = \mathbf{n}^T \mathbf{d}, \qquad a_i = \mathbf{n}^T \mathbf{d}_i, \tag{36}$$

and

$$(Q_{\text{chart}})_{ij} = -\mathbf{s}^T \mathbf{p}_{ij}. \tag{37}$$

The chart induces a nonuniform solid-angle measure:

$$\mathrm{d}\omega = \left| \mathbf{d}^T (\mathbf{d}_u \times \mathbf{d}_v) \right| \mathrm{d}u \, \mathrm{d}v = q^{-3/2} \, \mathrm{d}u \, \mathrm{d}v. \tag{38}$$

Converting the coordinate-space shift Jacobian to solid-angle measure requires multiplying by the ratio of this density at the mapped and original coordinates. Since the original path lies at $(u, v) = (0, 0)$, its density is one, and the correction reduces to $(1 + u'^2 + v'^2)^{-3/2}$, evaluated at the mapped coordinates (u', v').

2.2 Local vs. Global Coordinate Comparison

We compare local barycentric coordinates with the global initial-ray-centered hemispherical chart for the 2D Newton search in Fig. 1. Barycentric coordinates avoid ray tracing during Newton iterations but restrict the mapped point to the original triangle. The hemispherical chart traces each updated direction, allowing iterations to cross triangle boundaries. We evaluate the three scenes from the main manuscript and two additional scenes, Cornell-Dragon and Dragon, which contain increasingly prevalent finely tessellated geometry. Alongside the RGB images in the first row, we report $p_{\text{large}\Delta}$, the probability that a randomly sampled surface point lies on a large triangle, where a triangle is classified as large if its area exceeds 0.1. We compare several quantities across different wavelengths; the x-axis of every plot in Fig. 1 represents wavelength.

Under equal SPP (third row), global coordinates achieve lower MSE across all scenes thanks to their higher mapping success rate (fourth row). This advantage is most pronounced in Dragon, where fine triangles frequently cause local mappings to fail. Under equal time (second row), local coordinates generally perform better due to their lower computational overhead (fifth row). The exception is Dragon, which consists entirely of fine geometry, where the advantage of global coordinates outweighs their additional computational cost. These results favor local coordinates for typical synthetic scenes with moderately sized primitives and global coordinates for scenes dominated by finely tessellated geometry.

Further Discussion on Computational Overhead. Our antithetic mapping incurs greater overall overhead than that of the method of Kim et al. [2026], as we apply it at every NEE vertex within a single shader, roughly doubling the associated ray-tracing work (although this does not exactly double the total runtime, as other shader operations also contribute). In contrast, their method applies the mapping only to the sample selected by RIS, with most of the computational overhead arising from other ReSTIR operations.

2.3 Structured Light Jacobian Derivation

Here, we evaluate the Jacobian of the transformation for structured light. Since the BRDF PDF is often defined with respect to solid angle, we evaluate the Jacobian in the directional domain ω.

We consider the structured-light path $\mathbf{o} \to \mathbf{y} \to \mathbf{x}$, where $\mathbf{o}$ is the projector origin and $\mathbf{y}$ and $\mathbf{x}$ are path vertices. The direction $\omega_{\mathbf{o}\to\mathbf{y}}$ is parameterized by the normalized projector coordinate $\xi = (\xi_1, \xi_2) \in [0,1]^2$. Antithetic sampling is performed in this normalized projector space. Our goal is to evaluate the Jacobian of $\omega_{\mathbf{x}\to\mathbf{y}}$ with respect to the projector coordinate ξ.

Using the chain rule, we have

$$\left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}{\mathrm{d}\xi}\right| = \left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}{\mathrm{d}A(\mathbf{y})}\right| \left|\frac{\mathrm{d}A(\mathbf{y})}{\mathrm{d}\omega_{\mathbf{o}\to\mathbf{y}}}\right| \left|\frac{\mathrm{d}\omega_{\mathbf{o}\to\mathbf{y}}}{\mathrm{d}\xi}\right|. \tag{39}$$

Here, $\theta_{\mathbf{x}_1\mathbf{x}_2}$ denotes the angle between the surface normal at $\mathbf{x}_1$ and the direction $\omega_{\mathbf{x}_1\to\mathbf{x}_2}$. When $\mathbf{x}_1 = \mathbf{o}$, the corresponding normal is the projector's forward direction. For the first two terms, we have

$$\left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}{\mathrm{d}A(\mathbf{y})}\right| = \frac{\left|\cos\theta_{\mathbf{yx}}\right|}{|\mathbf{x}-\mathbf{y}|^2}, \tag{40}$$

and

$$\left|\frac{\mathrm{d}A(\mathbf{y})}{\mathrm{d}\omega_{\mathbf{o}\to\mathbf{y}}}\right| = \frac{|\mathbf{o}-\mathbf{y}|^2}{\left|\cos\theta_{\mathbf{yo}}\right|}. \tag{41}$$

For the third term, let the normalized projector coordinate ξ map to the following point on the projector image plane:

$$\mathbf{z}(\xi) = \begin{bmatrix} s_1(\xi_1 - c_1) \\ s_2(\xi_2 - c_2) \\ f \end{bmatrix}, \tag{42}$$

where f is the focal length, (c_1, c_2) is the principal point in normalized projector coordinates, and s_1 and s_2 convert normalized projector coordinates to image-plane coordinates. The corresponding projector direction is

$$\omega_{\mathbf{o}\to\mathbf{y}} = \frac{\mathbf{z}(\xi)}{|\mathbf{z}(\xi)|}. \tag{43}$$

The differential image-plane area corresponding to $\mathrm{d}\xi$ is

$$\mathrm{d}A_{\mathrm{proj}} = s_1 s_2 \, \mathrm{d}\xi. \tag{44}$$

The conversion from image-plane area to solid angle is

$$\left|\frac{\mathrm{d}\omega_{\mathbf{o}\to\mathbf{y}}}{\mathrm{d}A_{\mathrm{proj}}}\right| = \frac{\cos\theta_{\mathbf{oy}}}{|\mathbf{z}(\xi)|^2}. \tag{45}$$

Since

$$\cos\theta_{\mathbf{oy}} = \frac{f}{|\mathbf{z}(\xi)|}, \tag{46}$$

we have

$$|\mathbf{z}(\xi)| = \frac{f}{\cos\theta_{\mathbf{oy}}}. \tag{47}$$

Therefore,

$$\left|\frac{\mathrm{d}\omega_{\mathbf{o}\to\mathbf{y}}}{\mathrm{d}\xi}\right| = \frac{s_1 s_2}{f^2}\cos^3\theta_{\mathbf{oy}}. \tag{48}$$

Combining the three terms gives

$$\left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}{\mathrm{d}\xi}\right| = \frac{\left|\cos\theta_{\mathbf{yx}}\right|}{|\mathbf{x}-\mathbf{y}|^2}\frac{|\mathbf{o}-\mathbf{y}|^2}{\left|\cos\theta_{\mathbf{yo}}\right|}\frac{s_1 s_2}{f^2}\cos^3\theta_{\mathbf{oy}}. \tag{49}$$

Similarly, we evaluate the Jacobian of $\omega_{\mathbf{x}\to\mathbf{y}'}$ with respect to ξ', where the new normalized projector coordinate is obtained through the antithetic mapping

$$\xi' = T(\xi). \tag{50}$$

The Jacobian between the two directions is then

$$J_{T,\omega} = \left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}'}}{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}\right| = \frac{\left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}'}}{\mathrm{d}\xi'}\right|}{\left|\frac{\mathrm{d}\omega_{\mathbf{x}\to\mathbf{y}}}{\mathrm{d}\xi}\right|}\left|\frac{\mathrm{d}\xi'}{\mathrm{d}\xi}\right|. \tag{51}$$

We use this Jacobian to evaluate the PDF of the antithetic sample:

$$p_a(\mathbf{x}\to\mathbf{y}') = \frac{p_p(\mathbf{x}\to\mathbf{y})}{J_{T,\omega}}, \tag{52}$$

where $p_p(\mathbf{x}\to\mathbf{y})$ is obtained from standard BRDF importance sampling.

3 CW-ToF Phase and Depth Reconstruction

Here, we present additional CW-ToF phase and depth reconstruction results omitted from the main manuscript. We render the Cornell-Bunny scene using CW-ToF with high-frequency modulation of $\lambda = 0.02$ and four phase shifts. We then perform standard phase reconstruction from these measurements, as shown in the first row of Fig. 2.

We find that naive sampling produces noisy phase reconstructions, particularly near the corner of the back wall, due to noisy estimates of indirect illumination. In contrast, antithetic sampling produces significantly cleaner phase reconstructions.

To better visualize the effect of global illumination, we subtract the reconstructed phase from that obtained using direct-illumination-only rendering, as shown in the second row of Fig. 2. Even the ground-truth result exhibits a nonzero phase difference because of global illumination. Our antithetic sampling accurately reproduces this deviation, enabling the analysis of subtle artifacts caused by indirect transport.

Finally, we perform depth reconstruction using standard spatial phase unwrapping, implemented using the `unwrap_phase` function

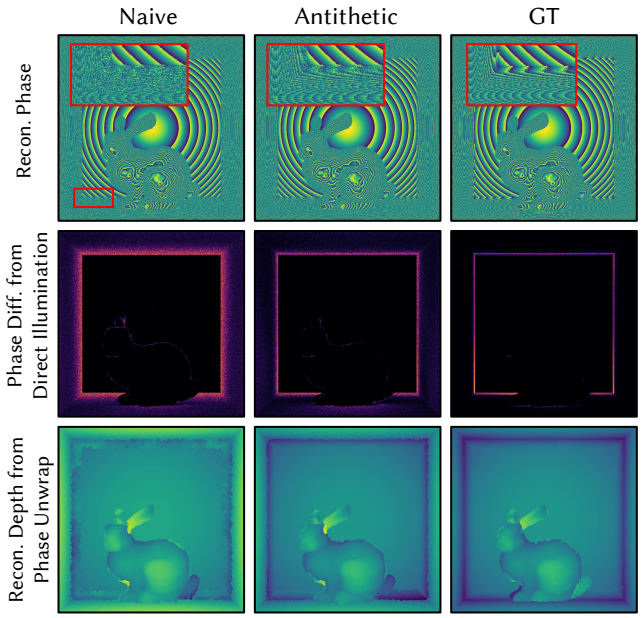


Fig. 2. The first row shows phase reconstruction using four phase shifts. The second row shows phase differences relative to direct-illumination-only rendering, visualizing errors induced by global illumination. The last row shows depth reconstruction obtained through phase unwrapping.

in scikit-image [Van der Walt et al. 2014]. Even the ground-truth rendering exhibits reconstruction artifacts caused by global illumination, particularly near the corners. Our antithetic sampling successfully reproduces these effects and produces results visually similar to the ground truth, whereas naive sampling produces substantially different artifacts near the corners.

Because we employ a basic phase-unwrapping algorithm, the bunny geometry is not perfectly reconstructed, with discontinuities appearing near the ears. More advanced unwrapping strategies using multiple modulation frequencies [Baek et al. 2023; Gupta et al. 2015] could further improve reconstruction quality. Nevertheless, our antithetic sampling framework remains effective for accurately simulating subtle global-illumination effects in these imaging systems.

References

Seung-Hwan Baek, Noah Walsh, Ilya Chugunov, Zheng Shi, and Felix Heide. 2023. Centimeter-Wave Free-Space Neural Time-of-Flight Imaging. *ACM Transactions on Graphics* 42, 1 (March 2023), 3:1–3:18. doi:10/gspq49

Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. 2015. Phasor Imaging: A Generalization of Correlation-Based Time-of-Flight Imaging. *ACM Transactions on Graphics* 34, 5 (2015). doi:10/gfz5k7

Juhyeon Kim, Wojciech Jarosz, and Adithya Pediredla. 2026. ToF ReSTIR: Time-of-Flight Rendering with Spatio-temporal Reservoir Resampling. (2026). arXiv:2605.11536 [cs.GR]

Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. 2014. scikit-image: image processing in Python. *PeerJ* 2 (2014), e453.