



Warped-Area Reparameterization of Differential Path Integrals

PEIYU XU, University of California, Irvine, USA

SAI BANGARU, MIT CSAIL, USA

TZU-MAO LI, University of California, San Diego, USA

SHUANG ZHAO, University of California, Irvine, USA




Fig. 1. We introduce the formulation of *reparameterized differential path integrals* for physics-based differentiable rendering. Our formulation can be efficiently estimated using advanced methods like bidirectional path tracing without requiring explicit sampling of discontinuity boundaries. In this example, we show several glass and metal chess pieces lit by an area light. The derivatives (obtained with our *bidirectional* estimator) are w.r.t. the position of the light.

Physics-based differentiable rendering is becoming increasingly crucial for tasks in inverse rendering and machine learning pipelines. To address discontinuities caused by geometric boundaries and occlusion, two classes of methods have been proposed: 1) the edge-sampling methods that directly sample light paths at the scene discontinuity boundaries, which require nontrivial data structures and precomputation to select the edges, and 2) the reparameterization methods that avoid discontinuity sampling but are currently limited to hemispherical integrals and unidirectional path tracing.

We introduce a new mathematical formulation that enjoys the benefits of both classes of methods. Unlike previous reparameterization work that focused on hemispherical integral, we derive the reparameterization in the path space. As a result, to estimate derivatives using our formulation, we can apply advanced Monte Carlo rendering methods, such as bidirectional path tracing, while avoiding explicit sampling of discontinuity boundaries. We show differentiable rendering and inverse rendering results to demonstrate the effectiveness of our method.

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: Differentiable rendering, differential path integral, warped-area reparameterization

ACM Reference Format:

Peiyu Xu, Sai Bangaru, Tzu-Mao Li, and Shuang Zhao. 2023. Warped-Area Reparameterization of Differential Path Integrals. *ACM Trans. Graph.* 42, 6, Article 213 (December 2023), 18 pages. <https://doi.org/10.1145/3618330>

Authors' addresses: Peiyu Xu, peiyux3@uci.edu, University of California, Irvine, USA; Sai Bangaru, sbangaru@mit.edu, MIT CSAIL, USA; Tzu-Mao Li, tzli@ucsd.edu, University of California, San Diego, USA; Shuang Zhao, shz@ics.uci.edu, University of California, Irvine, USA.



This work is licensed under a Creative Commons Attribution International 4.0 License. © 2023 Copyright held by the owner/author(s). 0730-0301/2023/12-ART213

<https://doi.org/10.1145/3618330>

1 INTRODUCTION

Physics-based differentiable rendering is the task of numerically computing derivatives of radiometric measurements with respect to arbitrary scene parameters such as object shapes and optical properties. Such scene derivatives not only can enable gradient-based optimization for solving inverse rendering problems (e.g., [Azinović et al. 2019; Luan et al. 2021]), but also are a key ingredient for integrating physics-based rendering into probabilistic-inferences and machine-learning pipelines (e.g., [Che et al. 2020]).

A key challenge for developing general-purpose differentiable rendering techniques is the differentiation with respect to scene geometries (such as the pose of an object or the position of a mesh vertex). This is because such geometries affect visibility and, if not handled properly, can lead to severely biased derivative estimates—which has been demonstrated by many prior works (e.g., [Li et al. 2018; Loubet et al. 2019; Zhang et al. 2019]).

To address this problem, two categories of techniques have been introduced. The first category *directly samples discontinuity boundaries* [Li et al. 2018; Zhang et al. 2019, 2020, 2021b], and the state of the art is Zhang et al.'s [2020] *differential path integral* formulation which tracks and handles discontinuities at the path level. The second category, on the other hand, *reparameterizes rendering integrals* to avoid explicit handling of discontinuities altogether [Loubet et al. 2019; Bangaru et al. 2020], with the state of the art being Bangaru et al.'s [2020] *warped-area reparameterization*.

In practice, Zhang et al.'s differential path integrals offer the flexibility to develop advanced Monte Carlo estimators, such as bidirectional path tracing, but require nontrivial data structures and precomputation [Yan et al. 2022] to efficiently sample discontinuity boundaries. Bangaru et al.'s warped-area reparameterization, on the other hand, enjoys the convenience of not needing explicit tracking

of discontinuities but has been limited to hemispherical integrals and unidirectional path tracing.

In this paper, we introduce a new formulation that enjoys the benefits of both these methods. Specifically, we adopt Bangaru et al.'s [2020] technique to reparameterize Zhang et al.'s [2020] differential path integrals. Our resulting formulation can be estimated using advanced methods (e.g., bidirectional path tracing) without explicit sampling of discontinuity boundaries.

Concretely, our contributions include:

- Introducing a new formulation for *differential one-bounce light transport* by applying Bangaru et al.'s [2020] warped-area method on top of Zhang et al.'s [2020] material-form reparameterization (§3).
- Establishing the formulation of *reparameterized differential path integrals* by generalizing the one-bounce result in §3 to handle full path integrals (§4).
- Developing Monte Carlo estimators leveraging unidirectional and bidirectional path sampling methods for our reparameterized differential path integrals (§5).

We validate our technique by comparing derivative estimates obtained using finite difference (FD) and our methods with high sample counts (Figure 11). Further, we show differentiable rendering (Figures 12, 13, 15 and 18) and inverse rendering (Figures 12–17) results to demonstrate the effectiveness of our technique.

2 RELATED WORK

Recently, great progresses have been made in the field of physics-based differentiable rendering. In what follows, we discuss most relevant techniques. For a more comprehensive overview, please refer to online course materials, e.g., [Zhao et al. 2020].

Handling discontinuities. A main challenge toward developing general-purpose differentiable rendering engines has been the differentiation with respect to scene geometry, which generally requires calculating additional boundary integrals.

To address this problem, two classes of techniques have been introduced. The first class directly samples discontinuity boundaries. Specifically, Li et al. [2018] introduced Monte Carlo edge sampling—the first technique in this class—by differentiating Kajiya's [1986] rendering equation. Zhang et al. [2019] later generalized this technique to differentiate the radiative transfer equation [Chandrasekhar 1960] for volumetric light transport. Further, Zhang et al. [2020; 2021b] have recently introduced the formulation of *differential path integrals* where discontinuities are tracked and handled at the path level, leading to significantly better performance.

The second class of methods reparameterize boundary integrals and avoid explicit handling of discontinuity boundaries altogether. The first approach of this category was biased and introduced by Loubet et al. [2019]. This method was later extended by Bangaru et al.'s [2020] *warped-area reparameterization*—which is capable of offering unbiased and consistent derivative estimates.

In §3, we will derive a new formulation by applying Bangaru et al.'s warped-area technique on top of Zhang et al.'s material-form reparameterization to differentiate one-bounce light transport. This

result will then be generalized in §4 to produce *reparameterized differential path integrals*—the main result of this paper.

Efficient sampling and differentiation. Previously, differentiable rendering algorithms have been largely repurposing Monte Carlo methods developed for forward rendering to sample light paths. Recently, several sampling methods have been developed specifically for differentiable rendering [Zeltner et al. 2021; Zhang et al. 2021a; Nimier-David et al. 2022; Yan et al. 2022; Yu et al. 2022]. Our technique is largely complementary to these methods, although we consider developing new Monte Carlo techniques specific to our formulation an important topic for future research.

Additionally, several approaches have been proposed to allow differentiable renderers scaling out to complex scenes with large numbers of parameters [Nimier-David et al. 2020; Vicini et al. 2021]. Our mathematical formulation is largely orthogonal to these techniques. Practically, our *unidirectional* estimator is compatible with these techniques. On the other hand, how these methods can be used with bidirectional path-sampling techniques (including our *bidirectional* estimator) remains an open problem.

3 DIFFERENTIATING ONE-BOUNCE LIGHT TRANSPORT WITH EVOLVING EMITTERS

We introduce in this section a simple form of our technique that differentiates *one-bounce light transport with evolving area lights*.

Specifically, we revisit the *material-form* reparameterization introduced by Zhang et al. [2020] in §3.1 before introducing our method that further reparameterizes the material-form result using the *warped-area* technique developed by Bangaru et al. [2020] in §3.2.

Problem specification. Under one-bounce light transport (aka. *direct illumination*), the scattered (i.e., reflected or refracted) radiance leaving a surface point $\mathbf{y}$ toward some point $\mathbf{y}'$ can be expressed as an integral over the surface $\mathcal{M}_e$ of all emitters:

$$I_{\text{dir}}(\mathbf{y} \rightarrow \mathbf{y}') = \int_{\mathcal{M}_e} L_e(\mathbf{x} \rightarrow \mathbf{y}) f_s(\mathbf{x} \rightarrow \mathbf{y} \rightarrow \mathbf{y}') G(\mathbf{x} \leftrightarrow \mathbf{y}) dA(\mathbf{x}), \quad (1)$$

where L_e denotes the **source emission**, f_s indicates the **bidirectional scattering distribution function** (BSDF), G is the (visibility-aware) **geometric term**, and A is the area measure.

In what follows, we focus on the configuration where the emitter surface $\mathcal{M}_e$ evolves with some parameter $\theta \in \mathbb{R}$ while all other surfaces—as well as $\mathbf{y}$ and $\mathbf{y}'$ on the right-hand side of Eq. (1)—are static (i.e., independent of θ). Under this setting, we consider the problem of differentiating the outgoing radiance I_{dir} with respect to θ . We summarize commonly used symbols and their definitions in Table 1.

3.1 Preliminary: Material-Form Reparameterization

To facilitate the differentiation of I_{dir} , Zhang et al. [2020] propose to rewrite the integral as one over some fixed **reference surface** $\mathcal{B}_e$. This can be achieved using a differentiable one-to-one mapping (aka. a **deformation**) $X(\cdot, \theta)$ that transforms the fixed reference surface $\mathcal{B}_e$ to the evolving one $\mathcal{M}_e(\theta)$. To distinguish points on the fixed $\mathcal{B}_e$ and the evolving $\mathcal{M}_e(\theta)$, we call any $\mathbf{p} \in \mathcal{B}_e$ a **material point** and $\mathbf{x} \in \mathcal{M}_e(\theta)$ a **spatial point**.




Fig. 2. **Material-form reparameterization:** This example involves a simple configuration with a fixed occluder and an emitter over an evolving surface $\mathcal{M}_e(\theta)$. Leveraging Zhang et al.’s [2020] material-form reparameterization, one can use a constant reference surface $\mathcal{B}_e$ coupled with a differentiable motion X such that, for any θ , $X(\cdot, \theta)$ is a differentiable one-to-one mapping (i.e., a deformation) from $\mathcal{B}_e$ to $\mathcal{M}_e(\theta)$. We note that, although the reference surface $\mathcal{B}_e$ is fixed, the discontinuity curves $\Delta\mathcal{B}_e$ capturing visibility boundaries with respect to the surface point $\mathbf{y}$ generally depends on the parameter θ .

Applying a change of variable $\mathbf{x} := X(\mathbf{p}, \theta)$ to the original surface integral (1) produces a *material-form*¹ variant:

$$I_{\text{dir}} = \int_{\mathcal{B}_e} \underbrace{L_e(\mathbf{x} \rightarrow \mathbf{y}) f_s(\mathbf{x} \rightarrow \mathbf{y}') G(\mathbf{x} \leftrightarrow \mathbf{y}) J(\mathbf{p}, \theta)}_{=: F_{\text{dir}}(\mathbf{p})} dA(\mathbf{p}), \quad (2)$$

where

$$J(\mathbf{p}, \theta) := \left\| \frac{dA(\mathbf{x})}{dA(\mathbf{p})} \right\|, \quad (3)$$

is the Jacobian determinant resulting from the change of variable. In Eqs. (2) and (3), we consider $\mathbf{x} = X(\mathbf{p}, \theta)$ a function of the material point $\mathbf{p} \in \mathcal{B}_e$ and the parameter θ . Also, we omit the dependency of the integrand F_{dir} on the parameter θ for notational simplicity.

Fundamentally, the material-form reparameterization has moved the derivative contribution from the evolution of the domain $\mathcal{M}_e$ into the new integrand F_{dir} —whose derivative is easier to compute due to fewer types of discontinuities that requires handling.

Choice of reference surfaces. When estimating derivatives at $\theta = \theta_0$ (for some fixed θ_0), the reference surface is usually set to $\mathcal{B}_e = \mathcal{M}_e(\theta_0)$. Under this configuration, the mapping $X(\cdot, \theta_0)$ becomes the identity map, causing the corresponding Jacobian determinant $J(\mathbf{p}, \theta_0)$ to reduce to one. We note that, the derivative of $J(\mathbf{p}, \theta)$ with respect to θ is generally nonzero—even at $\theta = \theta_0$.

Differential direct-illumination integral. In general, differentiating Eq. (2) using Reynolds transport theorem [1903; 1973] produces *material-form differential integrals*:

$$\frac{dI_{\text{dir}}}{d\theta} = \underbrace{\int_{\mathcal{B}_e} \frac{dF_{\text{dir}}(\mathbf{p})}{d\theta} dA(\mathbf{p})}_{\text{interior}} + \underbrace{\int_{\Delta\mathcal{B}_e} \Delta F_{\text{dir}}(\mathbf{p}) V(\mathbf{p}) d\ell(\mathbf{p})}_{\text{boundary}}, \quad (4)$$

where ℓ denotes the curve-length measure, and the *interior* component is a surface integral identical to Eq. (1) except for having the

¹The term “material form” originates in continuum mechanics where reparameterizations like Eq. (2) are common.

Table 1. Commonly used symbols in §3. The right-most column indicates θ -dependency.

Symbol	Definition	θ -dep.
$\mathcal{M}_e$	Emitter surface evolving with some parameter θ	Yes
$\mathcal{B}_e$	Reference surface	No
$X(\cdot, \theta)$	Differentiable one-to-one mapping (aka. deformation) from $\mathcal{B}_e$ to $\mathcal{M}_e(\theta)$	Yes
I_{dir}	Radiance $L(\mathbf{y} \rightarrow \mathbf{y}')$ with fixed $\mathbf{y}$ and $\mathbf{y}'$	Yes
F_{dir}	Integrand of the material-form integral (2)	Yes
$\Delta\mathcal{B}_e$	Curves comprising jump discontinuities of F_{dir}	Yes
$\mathcal{B}_e^{\text{wa}}$	Regions on the reference surface $\mathcal{B}_e$ comprising points $\mathbf{p}$ such that $X(\mathbf{p}, \theta)$ is visible to $\mathbf{y}$	Yes
$\partial\mathcal{B}_e^{\text{wa}}$	Boundary of $\mathcal{B}_e^{\text{wa}}$ (and a superset of $\Delta\mathcal{B}_e$)	Yes
$\mathbf{n}^\partial$	Unit-normal field associated with $\Delta\mathcal{B}_e$	Yes
$\mathbf{v}^\partial$	Velocity field over $\Delta\mathcal{B}_e$ (and later extended to $\partial\mathcal{B}_e^{\text{wa}}$)	Yes
$\mathbf{v}^{\text{dis}}$	Discontinuous velocity field over $\mathcal{B}_e$	Yes
$\mathbf{v}$	Continuous velocity field obtained by convolving $\mathbf{v}^{\text{dis}}$	Yes
$\mathbf{w}$	Spatially varying kernel for convolving $\mathbf{v}^{\text{dis}}$	No

integrand F_{dir} differentiated. The *boundary* component in Eq. (4)—which we will define and discuss in §3.1.1 and §3.1.2 below—is a line integral that is unique to differentiable rendering.

3.1.1 Discontinuities. The *boundary* integral in Eq. (4) is over a set of curves $\Delta\mathcal{B}_e$ comprising jump-discontinuity points of the integrand F_{dir} (with respect to $\mathbf{p}$). We note that, although the reference surface $\mathcal{B}_e$ is fixed, the discontinuity curves $\Delta\mathcal{B}_e$ typically depend on the parameter θ . Also, thanks to the material-form reparameterization, the discontinuity curves do not contain the topological boundary $\partial\mathcal{B}_e$ of the reference surface $\mathcal{B}_e$ (when $\mathcal{B}_e$ is open) since $\partial\mathcal{B}_e$ does not depend on the parameter θ . Additionally, for any discontinuity point $\mathbf{p} \in \Delta\mathcal{B}_e$, ΔF_{dir} captures the difference in F_{dir} across the boundary.

In this paper, we assume the emitted radiance L_e and BSDF f_s to be C^0 -continuous. Under this assumption, the discontinuities emerge solely from the mutual visibility between $\mathbf{x}$ and $\mathbf{y}$ —a factor of the geometric term G in Eqs. (1) and (2). Precisely, for any $\theta \in \mathbb{R}$, the discontinuity curves $\Delta\mathcal{B}_e$ consist of material points $\mathbf{p}$ whose spatial representations $\mathbf{x} = X(\mathbf{p}, \theta)$ are jump discontinuities of the mutual visibility $\mathbb{V}(\mathbf{x} \leftrightarrow \mathbf{y})$ (with $\mathbf{y}$ fixed), as illustrated in Figure 2. Further, we assume without loss of generality that the curve normal $\mathbf{n}^\partial(\mathbf{p})$ points toward the occluded side for all $\mathbf{p}$ (see Figure 3). Then, it holds that

$$\Delta F_{\text{dir}}(\mathbf{p}) = F_{\text{dir}}(\mathbf{p}), \quad (5)$$

where ΔF_{dir} denotes the difference in F_{dir} across a discontinuities boundary

3.1.2 Scalar normal velocity. In Eq. (4), another key component of the *boundary* integral is the **scalar normal velocity** V capturing how fast the discontinuity boundary $\Delta\mathcal{B}_e$ evolves (with respect to the parameter θ) along the curve normal $\mathbf{n}^\partial$. Precisely, this term is defined as

$$V(\mathbf{p}) = \mathbf{n}^\partial(\mathbf{p}) \cdot \mathbf{v}^\partial(\mathbf{p}), \quad (6)$$

ACM Trans. Graph., Vol. 42, No. 6, Article 213. Publication date: December 2023.




Fig. 5. We define $\mathcal{B}_e^{wa} \subset \mathcal{B}_e$ as the set of all material points $\mathbf{p}$ whose spatial representations $\mathbf{x} = \mathbf{X}(\mathbf{p}, \theta)$ are visible to $\mathbf{y}$. This example shows a spherical emitter with no topological boundary. In this case, the boundary $\partial\mathcal{B}_e^{wa}$ (illustrated in magenta) equals the union of visibility boundary $\Delta\mathcal{B}_e$ (illustrated in black) caused by the small occluder and the silhouette (illustrated in gray) of the sphere with respect to $\mathbf{y}$.

representations $\mathbf{x} = \mathbf{X}(\mathbf{p}, \theta)$ belong to one of the following categories:

- C.1 Topological boundary of the emitter surface $\mathcal{M}_e(\theta)$;
- C.2 Silhouette of the emitter (with respect to $\mathbf{y}$) comprising spatial points $\mathbf{x}$ satisfying $\mathbf{n}(\mathbf{x}) \cdot \mathbf{x}\bar{\mathbf{y}} = 0$;
- C.3 Visibility boundaries resulting from occlusion.

It follows that the discontinuity boundary $\Delta\mathcal{B}_e$ —which corresponds only points from the last category (C.3)—is a subset of $\partial\mathcal{B}_e^{wa}$. This allows us to rewrite the left-hand side of Eq. (10) as

$$\int_{\partial\mathcal{B}_e^{wa}} \left(F_{\text{dir}}(\mathbf{p}) \mathbf{v}^\partial(\mathbf{p}) \right) \cdot \mathbf{n}^\partial(\mathbf{p}) d\ell(\mathbf{p}), \quad (13)$$

where the normal velocity $\mathbf{v}^\partial(\mathbf{p})$ is set to zero for all $\mathbf{p}$ beyond the ordinary discontinuity boundaries $\Delta\mathcal{B}_e$.

3.2.2 Interpolating $\mathbf{v}^\partial$. Our objective of this stage is to specify a continuous velocity field $\mathbf{v}$ in the interior of the domain $\mathcal{B}_e^{wa}$ such that $\mathbf{v}$ agrees with $\mathbf{v}^\partial$ on the domain boundary $\partial\mathcal{B}_e^{wa}$. To this end, we adopt Bangaru et al.'s [2020] warped-area approach and employ a two-step process: we first define a discontinuous velocity field $\mathbf{v}^{\text{dis}}$ and then smooth it to make it continuous. In the following, we detail each of the two steps.

Step 1. In the first step, we define a discontinuous velocity field $\mathbf{v}^{\text{dis}}$ over the entire reference surface $\mathcal{B}_e$ (see Figure 6).

For any material point $\mathbf{p}$, when its spatial representation $\mathbf{x} = \mathbf{X}(\mathbf{p}, \theta)$ is invisible to $\mathbf{y}$ or resides on a visibility boundary (i.e., $\mathbf{p} \notin \mathcal{B}_e^{wa}$), we define $\mathbf{v}^{\text{dis}}(\mathbf{p})$ in a similar fashion as the boundary velocity $\mathbf{v}^\partial(\mathbf{p})$ described in §3.1.2. Precisely, let $\mathbf{y}^O$ be an intersection of the open line segment $\mathbf{y}\bar{\mathbf{x}}$ and the scene geometry. Then, by substituting $\mathbf{y}^B$ with $\mathbf{y}^O$ in Eq. (7), we define $\mathbf{v}^{\text{dis}}(\mathbf{p})$ using Eq. (8). Based on this construction, it is easy to verify that $\mathbf{v}^{\text{dis}}(\mathbf{p}) = \mathbf{v}^\partial(\mathbf{p})$ for any point $\mathbf{p}$ on the boundary $\partial\mathcal{B}_e^{wa}$.

On the other hand, when $\mathbf{x}$ is visible to $\mathbf{y}$ (i.e., $\mathbf{p} \in \mathcal{B}_e^{wa}$), there is no occluder between $\mathbf{x}$ and $\mathbf{y}$. In this case, we simply set $\mathbf{v}^{\text{dis}}(\mathbf{p}) = \mathbf{0}$.




Fig. 6. **Discontinuous velocity $\mathbf{v}^{\text{dis}}$.** For each material point $\mathbf{p} \in \mathcal{B}_e$, if its spatial representation $\mathbf{x} = \mathbf{X}(\mathbf{p}, \theta)$ is visible to $\mathbf{y}$, we set $\mathbf{v}^{\text{dis}}(\mathbf{p}) = \mathbf{0}$. Otherwise, there must exist an intersection $\mathbf{y}^O$ (illustrated as the red point) between the open line segment $\mathbf{y}\bar{\mathbf{x}}$ and some occluder. We compute $\mathbf{z} = \text{rayIntersect}(\mathbf{y} \rightarrow \mathbf{y}^O, \mathcal{M}_e)$ and set the velocity $\mathbf{v}^{\text{dis}}(\mathbf{p}) = \frac{d}{d\theta} \mathbf{X}^{-1}(\mathbf{z}, \theta)$.

Step 2. With the discontinuous velocity field $\mathbf{v}^{\text{dis}}$ defined, we construct the continuous velocity field $\mathbf{v}$ by smoothing $\mathbf{v}^{\text{dis}}$. Specifically, for any $\mathbf{p}$, we set

$$\mathbf{v}(\mathbf{p}) := \frac{\int_{\mathcal{B}_e} \mathbf{w}(\mathbf{q}; \mathbf{p}) \mathbf{v}^{\text{dis}}(\mathbf{q}) dA(\mathbf{q})}{\int_{\mathcal{B}_e} \mathbf{w}(\mathbf{q}; \mathbf{p}) dA(\mathbf{q})}, \quad (14)$$

where $\mathbf{w}$ is a **spatially varying kernel**.

To ensure that the resulting $\mathbf{v}$ agrees with $\mathbf{v}^{\text{dis}}$ (and, in turn, $\mathbf{v}^\partial$) at the boundary $\partial\mathcal{B}_e^{wa}$, we need $\mathbf{w}(\cdot, \mathbf{p})$ to behave like *Dirac delta distributions* for all $\mathbf{p} \in \partial\mathcal{B}_e^{wa}$ on the boundary, while having a smooth falloff when $\mathbf{p}$ moves away from the boundary. We follow Bangaru et al. [2020] and use:

$$\mathbf{w}(\mathbf{q}; \mathbf{p}) := (\mathbf{D}(\mathbf{q}; \mathbf{p}) + \mathbf{B}(\mathbf{q}))^{-1}, \quad (15)$$

where $\mathbf{D}(\mathbf{q}; \mathbf{p})$ is the **distance function** that must approach zero when $\mathbf{q}$ approaches $\mathbf{p}$, and $\mathbf{B}(\mathbf{q})$ is the **boundary-test function** that should approach zero when $\mathbf{q}$ approach the boundary $\partial\mathcal{B}_e^{wa}$. We note that, since the kernel $\mathbf{w}$ is used solely for smoothing $\mathbf{v}^{\text{dis}}$, we consider $\mathbf{w}$ independent of the parameter θ (i.e., “detached”).

In the following, we express our choices of the distance and the boundary-test functions.

Distance function. Bangaru et al. [2020] have introduced a distance function, but we found it non-robust—which we will demonstrate in §6.2. To address this problem, we choose a different distance function by letting

$$\mathbf{D}(\mathbf{q}; \mathbf{p}) := \left(\frac{1}{\sigma} \left(1 - \exp \left(-\frac{r^2}{\sigma} \right) \right) \right)^a, \quad (16)$$

where

$$r := \|\mathbf{q} - \mathbf{p}\|/r_0, \quad \sigma := \sigma_0/r_0, \quad (17)$$

with $r_0 := \|\mathbf{x} - \mathbf{y}\|$ being a normalization factor that makes r and σ invariant of global scaling of the scene. In Eqs. (16) and (17), a and σ_0 are user-specified hyperparameters—for which we use $a = 3$ and $\sigma_0 = 0.006$ in practice which work well in all our experiments.




Fig. 7. **Boundary-test function B :** For any material point q , we set $B(q)$ based on its spatial representation $\mathbf{x}' = \chi(q, \theta)$. In (a), we set $B(q)$ to $r_\theta(x'')$ —the shortest distance from x'' to the edge of the occluder. In (b), x' resides on the emitter surface that has no topological boundary. In this case, we set $B(q)$ to $|\vec{x'y} \cdot \mathbf{n}(x')|$.

Boundary-test function. The last remaining ingredient for completing the kernel (15) and, in turn, the convolution in Eq. (14) is the boundary-test function B . For all q on the boundary $\partial\mathcal{B}_e^{\text{wa}}$, we set $B(q) = 0$. For each $q \in \mathcal{B}_e^{\text{wa}}$, we have $v^{\text{dis}}(q) \equiv 0$ (as discussed in Step 1 above), and the value of $w(q; p)$ does not matter (as long as it is finite). Thus, we set $B(q) = 1$.

We now focus on the last case where $q \in \mathcal{B}_e \setminus (\mathcal{B}_e^{\text{wa}} \cup \partial\mathcal{B}_e^{\text{wa}})$. According to the definition of $\mathcal{B}_e^{\text{wa}}$ in Eq. (12), this implies that the spatial representation $\mathbf{x}' = \chi(q, \theta)$ is (strictly) invisible to y (i.e., $\mathbb{V}(\mathbf{x}' \leftrightarrow y) = 0$). In this case, to facilitate the specification of $B(q)$, we first define a **boundary-distance function** B^{dist} for any spatial point x as

$$B^{\text{dist}}(x) := \min(|\vec{xy} \cdot \mathbf{n}(x)|, r^\partial(x)), \quad (18)$$

where $\vec{xy}$ denotes the unit vector pointing from x to y , and $r^\partial(x)$ equals the shortest distance from x to an open boundary of the surface that contains x and $+\infty$ if the surface containing x is closed (and, thus, has no topological boundary). Based on this construction, it is easy to verify that, when x approaches a visibility boundary with respect to y , $B^{\text{dist}}(x)$ approaches zero.

With the boundary-distance function B^{dist} specified, we define the boundary-test function $B(q)$ based on the surface normal $\mathbf{n}(x')$ at the spatial representation $\mathbf{x}' = \chi(q, \theta)$ as follows.

- When $\mathbf{n}(x') \cdot \vec{x'y} > 0$, x' lies on a “front-facing” surface with respect to y (see Figure 7-a), and the open line segment $\vec{yx'}$ must intersect an occluder at x'' . We set $B(q) = B^{\text{dist}}(x'')$.
- When $\mathbf{n}(x') \cdot \vec{x'y} < 0$, x' is located on a “back-facing” surface with respect to y (see Figure 7-b). We set $B(q) = B^{\text{dist}}(x')$.

3.2.3 Completing the derivation. After constructing the continuous vector field v in §3.2.2, the right-hand side of Eq. (10) becomes well defined.

Lastly, according to the definition of the domain $\mathcal{B}_e^{\text{wa}}$ in Eq. (12), the spatial representation $\mathbf{x} = \chi(p, \theta)$ of any material point $p \notin \mathcal{B}_e^{\text{wa}}$ must be invisible to y (i.e., $\mathbb{V}(\mathbf{x} \leftrightarrow y) = 0$). Since the mutual visibility $\mathbb{V}(\mathbf{x} \leftrightarrow y)$ is a factor of $F_{\text{dir}}(p)$, it follows that that $F_{\text{dir}}(p) = 0$ and, in turn,

$$[\nabla \cdot (F_{\text{dir}} v)](p) = 0, \quad \text{for all } p \notin \mathcal{B}_e^{\text{wa}}. \quad (19)$$

This allows us to extend the domain of the integral on the right-hand side of Eq. (10) from $\mathcal{B}_e^{\text{wa}}$ to the full reference surface $\mathcal{B}_e$:

$$\int_{\mathcal{B}_e^{\text{wa}}} \nabla \cdot (F_{\text{dir}} v) dA = \int_{\mathcal{B}_e} \nabla \cdot (F_{\text{dir}} v) dA, \quad (20)$$

which, when substituted into Eq. (4), yields Eq. (11)—our main result for this section.

Discussion. The original warped-area method has been applied to the *boundary* integral over spherical curves (obtained by differentiating spherical rendering integrals). We will discuss the relation of our result and Bangaru et al.’s [2020] warped-area method in Appendix A.

4 WARPED-AREA REPARAMETERIZATION OF DIFFERENTIAL PATH INTEGRALS

In what follows, we first provide a brief recap of *differential path integrals* [Zhang et al. 2020, 2021b] in §4.1 and then introduce the general form of our technique—the main result of this paper—in §4.2.

We summarize the important symbols used in this section in Table 2.

4.1 Preliminary: Differential Path Integrals

4.1.1 Path integrals for forward rendering. At the foundation of many, if not most, advanced forward rendering techniques is the formulation that expresses the response I of a radiometric detector as a **path integral** [Veach 1997]:

$$I = \int_{\Omega} f(\bar{x}) d\mu(\bar{x}), \quad (21)$$

where $\bar{x} = (x_0, \dots, x_N)$ denotes a **light transport path** (with x_0 on an emitter and x_N on a detector), $\Omega = \bigcup_{N=1}^{\infty} \mathcal{M}^{N+1}$ is the **path space** (with $\mathcal{M}$ is the union of all object surfaces), and μ is the corresponding area-product measure.

The integrand of the path integral is the **measurement contribution function** f given by

$$f(\bar{x}) := L_e(x_0 \rightarrow x_1) W_e(x_{N-1} \rightarrow x_N) \left[\prod_{n=1}^{N-1} f_s(x_{n-1} \rightarrow x_n \rightarrow x_{n+1}) \right] \left[\prod_{n=1}^N G(x_{n-1} \leftrightarrow x_n) \right], \quad (22)$$

where L_e and W_e are the source emission and **detector importance** (or response), f_s is the bidirectional scattering distribution function (BSDF), and G is the geometric term.

Material-form path integrals. When the scene geometry $\mathcal{M}$ evolves with some parameter θ , using the material-form reparameterization described in §3.1, one can capture the evolution of $\mathcal{M}$ using a fixed reference surface $\mathcal{B}$ coupled with a deformation $\chi(\cdot, \theta)$ that maps $\mathcal{B}$ to $\mathcal{M}(\theta)$ for any θ .

Let $\bar{\mathbf{p}} = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ be a **material light path** with each vertex $\mathbf{p}_n$ on a reference surface $\mathcal{B}$. Applying the mapping $\chi(\cdot, \theta)$ to each vertex of $\bar{\mathbf{p}}$ produces a spatial light path $\bar{\mathbf{x}} = (\mathbf{x}_0, \dots, \mathbf{x}_N)$ where $\mathbf{x}_n = \chi(\mathbf{p}_n, \theta)$ for all $n = 0, 1, \dots, N$.

Then, by changing the variable of integration from spatial light paths $\bar{\mathbf{x}}$ to material ones $\bar{\mathbf{p}}$, the path integral of Eq. (21) can be rewritten in material form as:

$$I = \int_{\hat{\Omega}} \hat{f}(\bar{\mathbf{p}}) d\mu(\bar{\mathbf{p}}), \quad (23)$$

where the domain of integration is the **material path space**

$$\hat{\Omega} := \bigcup_{N=1}^{\infty} \mathcal{B}^{N+1}, \quad (24)$$

comprising material paths (of finite lengths). Additionally, the integrand of Eq. (23) is the **material measurement contribution** $\hat{f}$ that equals the product of the original measurement contribution f defined in Eq. (22) and the ratio of the two integral elements $d\mu(\bar{\mathbf{x}})/d\mu(\bar{\mathbf{p}}) = \prod_{n=0}^N J(\mathbf{p}_n, \theta)$:

$$\hat{f}(\bar{\mathbf{p}}) := f(\bar{\mathbf{x}}) \prod_{n=0}^N J(\mathbf{p}_n, \theta), \quad (25)$$

where J follows the definition in Eq. (3).

4.1.2 Differential path integrals. Zhang et al. [2020] have derived using Reynolds transport theorem [1903] derivatives of the material-form path integral (23) with respect to arbitrary parameter θ . The result can generally be expressed as **material-form differential path integrals**:

$$\frac{dI}{d\theta} = \underbrace{\int_{\hat{\Omega}} \frac{d\hat{f}(\bar{\mathbf{p}})}{d\theta} d\mu(\bar{\mathbf{p}})}_{\text{interior}} + \underbrace{\int_{\partial\hat{\Omega}} \Delta\hat{f}_K(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\mu(\bar{\mathbf{p}})}_{\text{boundary}}, \quad (26)$$

where the *interior* and *boundary* components are, respectively, over continuous regions of $\hat{f}$ and the discontinuous boundaries separating these regions. In what follows, we explain the *boundary* component in more details.

Domain of integration, discontinuity curves. The *boundary* integral in Eq. (26) is over the **material boundary path space** $\partial\hat{\Omega}$. They contain paths where one of the path segment goes through the scene boundary (see Figure 8 left). The space equals the union of disjoint subspaces $\partial\hat{\Omega}_{N,K}$:

$$\partial\hat{\Omega} = \bigcup_{N=1}^{\infty} \bigcup_{K=0}^{N-1} \partial\hat{\Omega}_{N,K}. \quad (27)$$

Each subspace $\partial\hat{\Omega}_{N,K}$ comprises **material boundary paths** $\bar{\mathbf{p}} = (\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_N)$ where $\mathbf{p}_n \in \mathcal{B}$ for all $n \neq K$ and $\mathbf{p}_K$ is constrained over a set of curves $\Delta\mathcal{B}_K \subset \mathcal{B}$ such that $\mathbf{p}_K$'s spatial representation $\mathbf{x}_K = \chi(\mathbf{p}_K, \theta)$ resides on a visibility boundary with respect to $\mathbf{x}_{K+1}$. We call the spatial line segment $\bar{\mathbf{x}}_K \mathbf{x}_{K+1}$ —which intersects the scene geometry $\mathcal{M}(\theta)$ at a single point $\mathbf{x}^B$ —a **boundary segment** (see Figure 8).

Similar to $\Delta\mathcal{B}_e$ for one-bounce light transport (see §3), the discontinuity curves $\Delta\mathcal{B}_K$ generally depend on the parameter θ . In fact, $\Delta\mathcal{B}_e$ are essentially $\Delta\mathcal{B}_0$ (with $N = 2$).




Fig. 8. The spatial representation of a material boundary path that is an element of the subspace $\partial\hat{\Omega}_{3,1}$ and has the boundary segment $\bar{\mathbf{x}}_1 \mathbf{x}_2$ that intersects the scene at exactly one point $\mathbf{x}^B$. Every path from this subspace contain four vertices $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$ with $\mathbf{p}_1$ constrained to a set of curves $\Delta\mathcal{B}_1$ comprising points whose spatial representations beside on visibility boundaries with respect to $\mathbf{x}_2$. To simplify the illustration, we only show part of the reference surface $\mathcal{B}$ that corresponds to the vertical plane containing the constrained $\mathbf{x}_1$.

Lastly, the measure $\hat{\mu}$ associated with the material boundary path space $\partial\hat{\Omega}$ satisfies that

$$d\hat{\mu}(\bar{\mathbf{p}}) = d\ell(\mathbf{p}_K) \prod_{n \neq K} dA(\mathbf{p}_n). \quad (28)$$

Difference in measurement contribution. The term $\Delta\hat{f}_K(\bar{\mathbf{p}})$ captures the difference in material measurement contribution $\hat{f}$ when the vertex $\mathbf{p}_K$ crosses the discontinuity boundary $\Delta\mathcal{B}_K$ (with all the other vertices fixed).

As stated in §3.1.1, we assume without loss of generality that the boundary $\Delta\mathcal{B}_K$'s normal $\mathbf{n}_K^\partial$ points toward the occluded side (as shown on the right of Figure 8). Then, it holds that

$$\Delta\hat{f}_K(\bar{\mathbf{p}}) = \hat{f}(\bar{\mathbf{p}}). \quad (29)$$

Boundary velocity. The term $V_K(\mathbf{p}_K)$ is the scalar normal velocity defined as

$$V_K(\mathbf{p}_K) = \mathbf{v}_K^\partial(\mathbf{p}_K) \cdot \mathbf{n}_K^\partial(\mathbf{p}_K), \quad (30)$$

where $\mathbf{v}_K^\partial(\mathbf{p}_K)$ is the boundary velocity of the discontinuity point $\mathbf{p}_K$ and can be computed using the process described by Eqs. (7) and (8) in §3.1.2.

In this case, the spatial point $\mathbf{z}$ (that coincides $\mathbf{x}_K$) in Eq. (7) should be computed by intersecting the ray $\mathbf{x}_{K+1} \rightarrow \mathbf{x}^B$ and the scene geometry $\mathcal{M}(\theta)$. When differentiating this ray-intersection process, the derivatives $d\mathbf{x}_{K+1}/d\theta$ and $d\mathbf{x}^B/d\theta$ are computed by differentiating the mappings $\chi(\mathbf{p}_{K+1}, \theta)$ and $\chi(\mathbf{p}^B, \theta)$, respectively. For the latter, $\mathbf{p}^B := \text{detach}[\chi^{-1}(\mathbf{x}^B, \theta)]$ is the “detached” material representation of the point $\mathbf{x}^B$.

4.2 Warped-Area Reparameterization of Boundary Path Integrals

We now generalize our one-bounce result derived in §3.2 to reparameterize the *boundary* component of the material-form differential path integrals in Eq. (26).

Derivation outline. To this end, our derivation will involve the following three main steps:

Table 2. Commonly used symbols in §4. The right-most column indicates θ -dependency.

Symbol	Definition	θ -dep.
$\mathcal{M}$	The union of all surfaces in the scene	Yes
$\mathcal{B}$	Reference surface	No
$\chi(\cdot, \theta)$	Differentiable one-to-one mapping (aka. deformation) transforming $\mathcal{B}$ to $\mathcal{M}(\theta)$	Yes
$\hat{\Omega}$	material path space defined in Eq. (24)	No
$\partial\hat{\Omega}$	material boundary path space defined in Eq. (27)	Yes
$\hat{f}$	material measurement distribution defined in Eq. (25)	Yes
$\hat{f}_K$	Equals $\hat{f}$ but treats the vertex $\mathbf{p}_K$ as the only variable (and considers all the other vertices fixed)	Yes
$\Delta\mathcal{B}_K$	Jump discontinuities of $\hat{f}_K$ with respect to $\mathbf{p}_K$	Yes
$\mathcal{B}_K^{\text{wa}}$	Regions on the reference surface $\mathcal{B}$ comprising points $\mathbf{p}_K$ such that $\chi(\mathbf{p}_K, \theta)$ is visible to $\mathbf{x}_{K+1}$	Yes
$\mathbf{n}_K^\partial$	Unit-normal field associated with $\Delta\mathcal{B}_K$	Yes
$\mathbf{v}_K^\partial$	Boundary velocity over $\Delta\mathcal{B}_K$	Yes
$\mathbf{v}_K$	Continuous velocity field over $\mathcal{B}$ that agrees with $\mathbf{v}_K^\partial$ on the boundary $\Delta\mathcal{B}_K$	Yes

- (1) Expanding the *boundary* path integral in Eq. (26) as nested integrals of individual path vertices $\{\mathbf{p}_n : n = 0, 1, \dots, N\}$. This allows us to isolate the integral with respect to $\mathbf{p}_K$ over discontinuity curves $\Delta\mathcal{B}_K$.
- (2) Reparameterizing the isolated line integral as an integral over the reference surface $\mathcal{B}$ using the warped-area method discussed in §3.2.
- (3) Rewriting the nested integrals with the reparameterized surface integral with respect to $\mathbf{p}_K$ as a new *interior* path integral which, in turn, is merged with the *interior* component of Eq. (26) to produce our *reparameterized differential path integral* in Eq. (38).

In the following, we detail each of these steps.

Step 1. We first express the *boundary* path integral as a sum of integrals over the subspaces $\partial\hat{\Omega}_{N,K}$ defined below Eq. (27):

$$\int_{\partial\hat{\Omega}} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\dot{\mu}(\bar{\mathbf{p}}) = \sum_{N=1}^{\infty} \sum_{K=0}^{N-1} \int_{\partial\hat{\Omega}_{N,K}} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\dot{\mu}(\bar{\mathbf{p}}). \quad (31)$$

For all N, K , the integral on the right-hand side of Eq. (31) can be further rewritten as nested integrals of individual vertices of the material path $\bar{\mathbf{p}} = (\mathbf{p}_0, \dots, \mathbf{p}_N)$:

$$\begin{aligned} & \int_{\partial\hat{\Omega}_{N,K}} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\dot{\mu}(\bar{\mathbf{p}}) \\ &= \int_{\mathcal{B}^N} \left(\int_{\Delta\mathcal{B}_K} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\ell(\mathbf{p}_K) \right) \prod_{n \neq K} dA(\mathbf{p}_n). \end{aligned} \quad (32)$$

On the right-hand side of this equation, the inner integral—which depends on $\mathbf{p}_n$ for all $n \neq K$ —is with respect to the vertex $\mathbf{p}_K$ and

over the discontinuity curves $\Delta\mathcal{B}_K$:

$$\begin{aligned} I(\mathbf{p}_n : n \neq K) &:= \int_{\Delta\mathcal{B}_K} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\ell(\mathbf{p}_K) \\ &= \int_{\Delta\mathcal{B}_K} \left(\hat{f}(\bar{\mathbf{p}}) \mathbf{v}_K^\partial(\mathbf{p}_K) \right) \cdot \mathbf{n}_K^\partial(\mathbf{p}_K) d\ell(\mathbf{p}_K). \end{aligned} \quad (33)$$

Step 2. Eq. (33) is essentially identical to the one-bounce *boundary* integral on the left-hand side of Eq. (10). Therefore, we can apply our warped-area reparameterization given by Eq. (10) to rewrite Eq. (33) as a surface integral:

$$I(\mathbf{p}_n : n \neq K) = \int_{\mathcal{B}_K^{\text{wa}}} \left[\nabla \cdot \left(\hat{f}_K \mathbf{v}_K \right) \right] (\mathbf{p}_K) dA(\mathbf{p}_K). \quad (34)$$

In this equation:

- The domain of integration $\mathcal{B}_K^{\text{wa}}$ follows the definition in Eq. (12) and comprises all material points $\mathbf{p}$ whose spatial representations $\mathbf{x}_K = \chi(\mathbf{p}_K, \theta)$ are visible to the (spatial) vertex $\mathbf{x}_{K+1} = \chi(\mathbf{p}_{K+1}, \theta)$. That is, $\mathcal{B}_K^{\text{wa}} = \{\mathbf{p}_K \in \mathcal{B} : \mathbb{V}(\mathbf{x}_K \leftrightarrow \mathbf{x}_{K+1}) = 1\}$.
- The vector field $\mathbf{v}_K$ is obtained using the process presented in §3.2.2 (that smoothly extends the velocity $\mathbf{v}_K^\partial(\mathbf{p}_K)$ defined on the boundary $\Delta\mathcal{B}_K$ into the interior of $\mathcal{B}_K^{\text{wa}}$).

Additionally, the term $\hat{f}_K$ in Eq. (34) equals the material measurement contribution $\hat{f}$ but considers $\mathbf{p}_K$ the only variable and all the other vertices $\{\mathbf{p}_n : n \neq K\}$ constants. Specifically, let $\hat{g}_K(\mathbf{p}_K)$ be the product of all factors of $\hat{f}$ that depend on $\mathbf{p}_K$:²

$$\begin{aligned} \hat{g}_K(\mathbf{p}_K) &:= G(\mathbf{x}_{K-1} \leftrightarrow \mathbf{x}_K) G(\mathbf{x}_K \leftrightarrow \mathbf{x}_{K+1}) \\ & f_s(\mathbf{x}_{K-2} \rightarrow \mathbf{x}_{K-1} \rightarrow \mathbf{x}_K) f_s(\mathbf{x}_{K-1} \rightarrow \mathbf{x}_K \rightarrow \mathbf{x}_{K+1}) \\ & f_s(\mathbf{x}_K \rightarrow \mathbf{x}_{K+1} \rightarrow \mathbf{x}_{K+2}) J(\mathbf{p}_K), \end{aligned} \quad (35)$$

where $\mathbf{x}_n = \chi(\mathbf{p}_n, \theta)$ is spatial representation of material point $\mathbf{p}_n$ for $n = K-1, K, K+1$. Then, it holds that

$$\hat{f}_K(\bar{\mathbf{p}}) = \hat{g}_K(\mathbf{p}_K) \text{detach} \left(\frac{\hat{f}(\bar{\mathbf{p}})}{\hat{g}_K(\mathbf{p}_K)} \right). \quad (36)$$

We note that $\hat{f}_K(\bar{\mathbf{p}})$ will be differentiated only with respect to the point $\mathbf{p}_K$ (and not the scene parameter θ) when evaluating the divergence $\nabla \cdot (\hat{f}_K \mathbf{v}_K)$ —which we will discuss in §5.1.

Step 3. As shown in §3.2.3, the divergence $[\nabla \cdot (\hat{f}_K \mathbf{v}_K)](\mathbf{p}_K) = 0$ for all $\mathbf{p}_K \notin \mathcal{B}_K^{\text{wa}}$ because the mutual visibility $\mathbb{V}(\mathbf{x}_K \leftrightarrow \mathbf{x}_{K+1}) = 0$ by definition. This allows us to extend the domain of the reparameterized surface integral in Eq. (34) from $\mathcal{B}_K^{\text{wa}}$ to the full reference surface $\mathcal{B}$. Then, substituting the extended Eq. (34) into Eqs. (31)

²We omit the dependency of $\hat{g}_K$ on the points $\mathbf{p}_{K-1}$ and $\mathbf{p}_{K+1}$ (that are considered constants) for notational simplicity.

and (32) produces

$$\begin{aligned} & \int_{\partial\Omega} \hat{f}(\bar{\mathbf{p}}) V_K(\mathbf{p}_K) d\mu(\bar{\mathbf{p}}) \\ &= \sum_{N=1}^{\infty} \sum_{K=0}^{N-1} \int_{\mathcal{B}^{N+1}} \left[\nabla \cdot (\hat{f}_K \mathbf{v}_K) \right] (\mathbf{p}_K) d\mu(\bar{\mathbf{p}}) \\ &= \int_{\hat{\Omega}} \left[\sum_{K=0}^{N-1} \left[\nabla \cdot (\hat{f}_K \mathbf{v}_K) \right] (\mathbf{p}_K) \right] d\mu(\bar{\mathbf{p}}), \end{aligned} \quad (37)$$

where the second equality is obtained by: (i) exchanging inner summation and the integration; and (ii) rewriting the sum of integrals over $\mathcal{B}^{N+1}$ for all N as one integral over the material path space $\partial\hat{\Omega}$.

Lastly, substituting Eq. (37) into the material-form differential path integrals (26) yields the main result of this paper—the **reparameterized differential path integral**:

$$\frac{dI}{d\theta} = \int_{\hat{\Omega}} \left(\frac{d\hat{f}(\bar{\mathbf{p}})}{d\theta} + \sum_{K=0}^{N-1} \left[\nabla \cdot (\hat{f}_K \mathbf{v}_K) \right] (\mathbf{p}_K) \right) d\mu(\bar{\mathbf{p}}), \quad (38)$$

where N indicates the length (i.e., number of segments) of the material light path $\bar{\mathbf{p}}$.

5 OUR MONTE CARLO ESTIMATORS

We now introduce Monte Carlo estimators for our reparameterized differential path integral expressed in Eq. (38). Specifically, we first discuss in §5.1 the estimation of the divergence term $\nabla \cdot (\hat{f}_K \mathbf{v}_K)$. Then, in §5.2, we introduce our path-space Monte Carlo estimators and how do we cache values along paths to increase computational efficiency.

5.1 Estimating Divergence

A key component of our reparameterized differential path integral of Eq. (38) is the divergence

$$\nabla \cdot (\hat{f}_K \mathbf{v}_K) = (\nabla \hat{f}_K) \cdot \mathbf{v}_K + \hat{f}_K (\nabla \cdot \mathbf{v}_K), \quad (39)$$

at some vertex $\mathbf{p}_K$ of a given material path $\bar{\mathbf{p}} = (\mathbf{p}_0, \dots, \mathbf{p}_N)$. In the following, we discuss the estimation of $\hat{f}_K$ and $\nabla \hat{f}_K$ in §5.1.1 as well as $\mathbf{v}_K$ and $\nabla \cdot \mathbf{v}_K$ in §5.1.2.

5.1.1 Measurement contribution. The right-hand side of Eq. (39) involves the term $\hat{f}_K$ and its gradient $\nabla \hat{f}_K$. We recall that, as defined in Eq. (36), $\hat{f}_K(\bar{\mathbf{p}})$ equals the material measurement contribution $\hat{f}(\bar{\mathbf{p}})$ except for treating $\mathbf{p}_K$ as the only variable and all other vertices $\{\mathbf{p}_n : n \neq K\}$ as constants. Therefore, we have $\hat{f}_K(\bar{\mathbf{p}}) = \hat{f}(\bar{\mathbf{p}})$, and:

$$\left[\nabla \hat{f}_K \right] (\bar{\mathbf{p}}) = \left[\frac{\partial \hat{f}_K}{\partial \mathbf{p}_K} \right] (\bar{\mathbf{p}}) = [\nabla \hat{g}_K] (\mathbf{p}_K) \frac{\hat{f}(\bar{\mathbf{p}})}{\hat{g}_K(\mathbf{p}_K)}, \quad (40)$$

where $\hat{g}_K$ is defined in Eq. (35) and the gradient $\nabla \hat{g}_K := \partial \hat{g}_K / \partial \mathbf{p}_K$ can be computed using automatic differentiation.

5.1.2 Continuous velocity. The remaining terms on the right-hand side of Eq. (39) are the continuous velocity $\mathbf{v}_K(\mathbf{p}_K)$ and its divergence $[\nabla \cdot \mathbf{v}_K](\mathbf{p}_K)$ (given the neighboring vertex $\mathbf{p}_{K+1}$ that is considered fixed). We outline this process in Algorithm 1.




Fig. 9. **Discontinuous velocity** $\mathbf{v}_1^{\text{dis}}$: Let $\mathbf{p}_2 \in \mathcal{B}$ be a fixed material point $\mathbf{p}_2 \in \mathcal{B}$ with the spatial representation $\mathbf{x}_2 = X(\mathbf{p}_2, \theta)$. For each point $\mathbf{p}_1$, we set the discontinuous velocity $\mathbf{v}_1^{\text{dis}}(\mathbf{p}_1)$ to zero if $\mathbf{p}_1$'s spatial counterpart $\mathbf{x}_1 = X(\mathbf{p}_1, \theta)$ is visible to $\mathbf{x}_2$. Otherwise, there must exist an intersection $\mathbf{x}^0$ (illustrated as the red square) between the open line segment $\overline{\mathbf{x}_1 \mathbf{x}_2}$ and some occluder. We compute $\mathbf{z}_1 = \text{rayIntersect}(\mathbf{x}_2 \rightarrow \mathbf{x}^0, \mathcal{M})$ and set the velocity as $\mathbf{v}_1^{\text{dis}}(\mathbf{p}_1) = \frac{d}{d\theta} X^{-1}(\mathbf{z}_1, \theta)$.




Fig. 10. We sample points $\{q_{ij}\}$ near a point $\mathbf{p}_K$ by drawing points q'_i in the tangent plane of $\mathbf{p}_K$ (from a Gaussian distribution centered at $\mathbf{p}_K$) and finding all projections of q'_i on the reference $\mathcal{B}$ along the surface normal $\mathbf{n}(\mathbf{p}_K)$.

According to the definition in Eq. (14), we express the continuous velocity $\mathbf{v}_K(\mathbf{p}_K)$ as

$$\mathbf{v}_K(\mathbf{p}_K) = \tilde{\mathbf{v}}_K(\mathbf{p}_K) / W(\mathbf{p}_K). \quad (41)$$

On the right-hand of this equation, the denominator is the *normalization term* W given by

$$W(\mathbf{p}_K) := \int_{\mathcal{B}} w(\mathbf{q}; \mathbf{p}_K) dA(\mathbf{q}), \quad (42)$$

and the numerator is the *convolved velocity* $\tilde{\mathbf{v}}_K$ defined as

$$\tilde{\mathbf{v}}_K(\mathbf{p}_K) := \int_{\mathcal{B}} w(\mathbf{q}; \mathbf{p}_K) \mathbf{v}_K^{\text{dis}}(\mathbf{q}) dA(\mathbf{q}). \quad (43)$$

In Eq. (43), $\mathbf{v}_K^{\text{dis}}$ is the discontinuous vector field that agrees with the boundary velocity $\mathbf{v}_K^{\partial}$ and can be constructed (conditioned on the vertex $\mathbf{p}_{K+1}$) using the process described in Step 1 of §3.2.2 (see Figure 9).

Sampling $\mathbf{q}$. Estimating the continuous velocity $\mathbf{v}_K$ and its divergence $(\nabla \cdot \mathbf{v}_K)$ based on Eqs. (42) and (43) requires sampling material points $\mathbf{q}$ from the reference surface $\mathcal{B}$.

ALGORITHM 1: Consistent estimator for the velocity $\mathbf{v}_K$ and its divergence $\nabla \cdot \mathbf{v}_K$

```

1 Estimate_Velocity( $\mathbf{p}_K, \mathbf{p}_{K+1}$ )
   Input: Two vertices  $\mathbf{p}_K$  and  $\mathbf{p}_{K+1}$  of a material path
   Output:  $\mathbf{v}_K(\mathbf{p}_K)$  and  $[\nabla \cdot \mathbf{v}_K](\mathbf{p}_K)$ 
2 begin
   /* Estimate  $W$  and  $\tilde{\mathbf{v}}_K$  (unbiased and consistent) */
3    $W \leftarrow 0; \nabla W \leftarrow \mathbf{0};$ 
4    $\tilde{\mathbf{v}}_K \leftarrow \mathbf{0}; \nabla \cdot \tilde{\mathbf{v}}_K \leftarrow 0;$ 
5   for  $i = 1$  to  $N_q$  do
6     Sample point  $\mathbf{q}'_i$  from the tangent plane of  $\mathbf{p}_K$  with
       probability pdf( $\mathbf{q}'_i$ );
7     Compute projections  $\{q_{ij} : j = 1, 2, \dots\}$  of  $\mathbf{q}'_i$  on the
       reference surface  $\mathcal{B}$ ;
8     foreach projection  $q_{ij}$  of  $\mathbf{q}'_i$  do
9        $w \leftarrow w(\mathbf{q}_{ij}; \mathbf{p}_K);$ 
10       $\nabla w \leftarrow [\nabla w(\cdot; \mathbf{p}_K)](\mathbf{q}_{ij});$ 
11       $\text{pdf} \leftarrow \text{pdf}(\mathbf{q}'_i) |\mathbf{n}(\mathbf{p}_K) \cdot \mathbf{n}(\mathbf{q}_{ij})|;$ 
12      Compute  $\mathbf{v}^{\text{dis}} = \mathbf{v}_K^{\text{dis}}(\mathbf{q}_{ij})$  and  $\nabla \cdot \mathbf{v}^{\text{dis}}$  (using AD)
        conditioned on  $\mathbf{p}_{K+1}$ ; // §3.2.2, Step 1
13       $W += w/\text{pdf};$  // Eq. (44)
14       $\nabla W += \nabla w/\text{pdf};$  // Eq. (60)
15       $\tilde{\mathbf{v}}_K += (w \mathbf{v}^{\text{dis}})/\text{pdf};$  // Eq. (45)
16       $\nabla \cdot \tilde{\mathbf{v}}_K += (\nabla w \cdot \mathbf{v}^{\text{dis}} + w (\nabla \cdot \mathbf{v}^{\text{dis}}))/\text{pdf};$  // Eq. (61)
8     end
9   end
10  /* Estimate  $\mathbf{v}_K$  and  $\nabla \cdot \mathbf{v}_K$  (biased but consistent) */
11   $\mathbf{v}_K \leftarrow \tilde{\mathbf{v}}_K/W;$  // Eq. (47)
12   $\nabla \cdot \mathbf{v}_K \leftarrow (\nabla \cdot \tilde{\mathbf{v}}_K)/W - (\nabla W \cdot \tilde{\mathbf{v}}_K)/W^2;$  // Eq. (62)
13  return  $\mathbf{v}_K, \nabla \cdot \mathbf{v}_K;$ 
14 end

```

To this end, we employ a process similar to the sampling of outgoing locations for *bidirectional subsurface scattering distribution functions* (BSSRDFs) as follows. As illustrated in Figure 10, we first draw a point $\mathbf{q}'$ inside the tangent plane of $\mathbf{p}_K$ from a 2D (isotropic) Gaussian distribution centered at $\mathbf{p}_K$. Then, we project $\mathbf{q}'$ onto the reference surface $\mathcal{B}$ along the surface normal $\mathbf{n}(\mathbf{p}_K)$ and return all resulting projections $\mathcal{B} \cap \{\mathbf{q}' + t\mathbf{n}(\mathbf{p}_K) : t \in \mathbb{R}\}$ as samples of $\mathbf{q}$.

Estimating $\mathbf{v}_K$. Leveraging the same point sampling method described above, we estimate the *normalization term* W and the *convolved velocity* $\tilde{\mathbf{v}}_K$ in an unbiased and consistent fashion as follows.

First, we draw N_q samples $\{\mathbf{q}'_i : i = 1, 2, \dots, N_q\}$ of $\mathbf{q}$ from the tangent plane of $\mathbf{p}_K$. Let the probability for drawing $\mathbf{q}'_i$ be $\text{pdf}(\mathbf{q}'_i)$. Then, unbiased estimators of W and $\tilde{\mathbf{v}}_K$ can be obtained via

$$\langle W(\mathbf{p}_K) \rangle = \sum_{i=1}^{N_q} \sum_j \frac{w(\mathbf{q}_{ij}; \mathbf{p}_K)}{\text{pdf}(\mathbf{q}_{ij})}, \quad (44)$$

$$\langle \tilde{\mathbf{v}}_K(\mathbf{p}_K) \rangle = \sum_{i=1}^{N_q} \sum_j \frac{w(\mathbf{q}_{ij}; \mathbf{p}_K) \mathbf{v}_K^{\text{dis}}(\mathbf{q}_{ij})}{\text{pdf}(\mathbf{q}_{ij})}, \quad (45)$$

where $\{q_{ij} : j = 1, 2, \dots\}$ are the projections of $\mathbf{q}'_i$ on $\mathcal{B}$, and

$$\text{pdf}(\mathbf{q}_{ij}) = \text{pdf}(\mathbf{q}'_i) |\mathbf{n}(\mathbf{p}_K) \cdot \mathbf{n}(\mathbf{q}_{ij})|. \quad (46)$$

ALGORITHM 2: Our *unidirectional* estimator for the reparameterized differential path integral of Eq. (38).

```

1 Path_Tracing()
   // We use  $\mathbf{x}_n^D$  and  $\mathbf{x}_n^S$  to denote, respectively, the
   // spatial representations of  $\mathbf{p}_n^D$  and  $\mathbf{p}_n^S$  for all  $n$ 
2 begin
   /* Initialize detector subpath */
3   Sample camera ray  $\mathbf{p}_0^D \rightarrow \mathbf{p}_1^D$  with probability pdf;
4    $\hat{f} \leftarrow W_e(\mathbf{x}_1^D \rightarrow \mathbf{x}_0^D) G(\mathbf{x}_1^D \leftrightarrow \mathbf{x}_0^D) J(\mathbf{p}_0^D) J(\mathbf{p}_1^D);$ 
5    $H^D \leftarrow 0; dI \leftarrow 0;$ 
6   for  $n = 1, 2, \dots$  do
7     /* Sample light vertex  $\mathbf{p}_n^S$  */
8     Sample  $\mathbf{p}_n^S$  on the emitter with probability  $\text{pdf}_n^{\text{NEE}}$ ;
9     /* Compute  $\hat{f}^{\text{NEE}} = \hat{f}(\mathbf{p}_n^S, \mathbf{p}_n^D, \dots, \mathbf{p}_0^D)$  and  $d\hat{f}^{\text{NEE}}/d\theta$  */
10     $\hat{f}^{\text{NEE}} \leftarrow \hat{f} \cdot f_s(\mathbf{x}_n^S \rightarrow \mathbf{x}_n^D \rightarrow \mathbf{x}_{n-1}^D) G(\mathbf{x}_n^S \leftrightarrow \mathbf{x}_n^D) J(\mathbf{p}_n^S);$ 
11     $\hat{f}^{\text{NEE}} *= L_e(\mathbf{x}_n^S \rightarrow \mathbf{x}_n^D);$ 
12    Compute  $d\hat{f}^{\text{NEE}} := d\hat{f}^{\text{NEE}}/d\theta$  using AD;
13    /* Compute the sum of divergences via Eq. (49) */
14     $\Sigma_{\text{div}} \leftarrow \hat{f}^{\text{NEE}} (H^D + h(\mathbf{p}_n^D; \mathbf{p}_n^S, \mathbf{p}_{n-1}^D) + h(\mathbf{p}_n^S; \mathbf{p}_n^D));$ 
15    /* Accumulate path contribution */
16     $dI += (d\hat{f}^{\text{NEE}} + \Sigma_{\text{div}})/(\text{pdf} \cdot \text{pdf}_n^{\text{NEE}});$ 
17    /* Sample next vertex  $\mathbf{p}_{n+1}^D$  */
18    Sample  $\mathbf{p}_{n+1}^D$  with probability  $\text{pdf}_{n+1}$ ;
19     $\text{pdf} *= \text{pdf}_{n+1};$ 
20    /* Update measurement contribution */
21     $\hat{f} *= f_s(\mathbf{x}_{n+1}^D \rightarrow \mathbf{x}_n^D \rightarrow \mathbf{x}_{n-1}^D) G(\mathbf{x}_{n+1}^D \leftrightarrow \mathbf{x}_n^D) J(\mathbf{p}_{n+1}^D);$ 
22    /* Update pre-fix sum based on Eq. (50) */
23     $H^D += h(\mathbf{p}_n^D; \mathbf{p}_{n+1}^D, \mathbf{p}_{n-1}^D);$ 
24  end
25  return  $dI;$ 
26 end

```

Lastly, given Eq. (41), a biased-but-consistent estimator of the continuous velocity $\mathbf{v}_K$ can be obtained using the unbiased estimations of the *normalization term* W and the *convolved velocity* $\tilde{\mathbf{v}}_K$ given by Eqs. (44) and (45), respectively:

$$\langle \mathbf{v}_K(\mathbf{p}_K) \rangle = \langle \tilde{\mathbf{v}}_K(\mathbf{p}_K) \rangle / \langle W(\mathbf{p}_K) \rangle. \quad (47)$$

Further, the estimation of the divergence $\nabla \cdot \mathbf{v}_K$ can be accomplished in a similar fashion using unbiased estimates of the gradient ∇w of the *normalization term* W and the divergence $\nabla \cdot \tilde{\mathbf{v}}_K$ of the *convolved velocity* $\tilde{\mathbf{v}}_K$. We provide more details of this process in Appendix B.

Debiasing. Bangaru et al. [2020] have introduced a process based on Russian roulette to debias estimators of $\mathbf{v}_K(\mathbf{p}_K)$ and $[\nabla \cdot \mathbf{v}_K](\mathbf{p}_K)$. On the other hand, achieving full unbiasedness requires the distance and boundary-test functions to satisfy specific conditions, which can be nontrivial to obtain (please see their work for more details). We consider the debiasing of our estimators a future research topic.

5.2 Path-Space Estimators

Our reparameterized differential path integral of Eq. (38) allows the development of path-space Monte Carlo estimators of the form:

$$\left\langle \frac{dI}{d\theta} \right\rangle = \frac{1}{\text{pdf}(\bar{\mathbf{p}})} \left(\frac{d\hat{f}(\bar{\mathbf{p}})}{d\theta} + \sum_{K=0}^{N-1} \left[\nabla \cdot (\hat{f}_K \mathbf{v}_K) \right] (\mathbf{p}_K) \right), \quad (48)$$

where $\text{pdf}(\bar{\mathbf{p}})$ denotes any probability density for sampling the material light path $\bar{\mathbf{p}} = (\mathbf{p}_0, \dots, \mathbf{p}_N)$, and the divergence term $\nabla \cdot (\hat{f}_K \mathbf{v}_K)$ is computed using estimators presented in §5.1.

We introduce an *unidirectional* (§5.2.1) and a *bidirectional* estimator (§5.2.2) that, respectively, apply unidirectional path tracing (PT) with next-event estimation (NEE) and bidirectional path tracing (BDPT) to sample the material path $\bar{\mathbf{p}}$.

Time complexity of naïve implementation. Given a material path $\bar{\mathbf{p}}$ with N segments, evaluating Eq. (48) naïvely takes $O(N^2)$ time since computing the divergence $\nabla \cdot (\hat{f}_K \mathbf{v}_K)$ —which involves computing and differentiating material measurement contribution $\hat{f}_K(\bar{\mathbf{p}})$ (with respect to $\mathbf{p}_K$)—takes $O(N)$ time for each K .

The time complexity becomes even higher when using PT with NEE and BDPT: For the former, each NEE effectively produces one full path, causing the cost for tracing one path to be $O(N^3)$; For the latter, since $O(N^2)$ paths are created³ using one pair of source and detector subpaths, the total complexity becomes $O(N^4)$.

In what follows, we address this challenge by developing new techniques to estimate the divergence $\nabla \cdot (\hat{f}_K \mathbf{v}_K)$ efficiently for our *unidirectional* and *bidirectional* estimators.

5.2.1 Our unidirectional estimator. We now detail our *unidirectional* estimator—which we outline in Algorithm 2.

For any material path $\bar{\mathbf{p}} = (\mathbf{p}_0, \dots, \mathbf{p}_N)$, according to Eq. (40) and the fact that $\hat{f}_K(\bar{\mathbf{p}}) = \hat{f}(\bar{\mathbf{p}})$, we can rewrite the sum of divergence $\Sigma_{\text{div}}(\bar{\mathbf{p}}) := \sum_K [\nabla \cdot (\hat{f}_K \mathbf{v}_K)](\mathbf{p}_K)$ by factoring out the material measurement contribution $\hat{f}(\bar{\mathbf{p}})$. That is,

$$\begin{aligned} \Sigma_{\text{div}}(\bar{\mathbf{p}}) &:= \sum_{K=0}^{N-1} \left[\nabla \cdot (\hat{f}_K \mathbf{v}_K) \right] (\mathbf{p}_K) \\ &= \hat{f}(\bar{\mathbf{p}}) \sum_{K=0}^{N-1} \underbrace{\left(\frac{[\nabla \hat{g}_K](\mathbf{p}_K)}{\hat{g}_K(\mathbf{p}_K)} \cdot \mathbf{v}_K(\mathbf{p}_K) + [\nabla \cdot \mathbf{v}_K](\mathbf{p}_K) \right)}_{=: h(\mathbf{p}_K; \mathbf{p}_{K-1}, \mathbf{p}_{K+1})}. \end{aligned} \quad (49)$$

In this equation, the term $\hat{g}_K$ —which is a function of $\mathbf{p}_K$ conditioned on $\mathbf{p}_{K-1}$ and $\mathbf{p}_{K+1}$ —is defined in Eq. (35), and the velocity $\mathbf{v}_K(\mathbf{p}_K)$ and its divergence $[\nabla \cdot \mathbf{v}_K](\mathbf{p}_K)$ can be estimated using Algorithm 1 from §5.1.2.

We now make an important observation that Eq. (49) can be computed in $O(N)$ time since the term $h(\mathbf{p}_K; \mathbf{p}_{K-1}, \mathbf{p}_{K+1})$ on the right-hand side is essentially a function of one point $\mathbf{p}_K$ and, thus, can be computed in $O(1)$ time.

³In this paper, we focus on standard bidirectional path tracing (BDPT) that connects all pairs of vertices between the source and the detector subpaths.




Fig. 11. We **validate** our techniques by comparing derivative images (visualized using the same color mapping as Figure 1) estimated by our method (c) to finite-difference (FD) references (b). Our “teapot” and “bunny” results are generated, respectively, using our *unidirectional* and *bidirectional* estimators. The derivatives are visualized using the same color map as Figure 1.

Based on Eq. (49), we introduce a *unidirectional* estimator that uses next-event estimation (NEE) that works as follows. When tracing a detector subpath $(\mathbf{p}_0^D, \mathbf{p}_1^D, \dots)$, we maintain the prefix sum

$$H_n^D := \sum_{m=1}^{n-1} h(\mathbf{p}_m^D; \mathbf{p}_{m+1}^D, \mathbf{p}_{m-1}^D) = H_{n-1}^D + h(\mathbf{p}_{n-1}^D; \mathbf{p}_n^D, \mathbf{p}_{n-2}^D), \quad (50)$$

for each vertex $\mathbf{p}_n^D$. After performing NEE at this vertex using a light sample point $\mathbf{p}_n^S$, we obtain a full path $\bar{\mathbf{p}}_n := (\mathbf{p}_n^S, \mathbf{p}_n^D, \dots, \mathbf{p}_1^D, \mathbf{p}_0^D)$. Leveraging the prefix sum H_n , we can compute the sum of divergences $\Sigma_{\text{div}}(\bar{\mathbf{p}}_n)$ defined in Eq. (49) using $O(1)$ time via:

$$\Sigma_{\text{div}}(\bar{\mathbf{p}}_n) = \hat{f}(\bar{\mathbf{p}}_n) \left(H_n^D + h(\mathbf{p}_n^D; \mathbf{p}_n^S, \mathbf{p}_{n-1}^D) + h(\mathbf{p}_n^S; \mathbf{p}_n^D) \right), \quad (51)$$

where the material measurement contribution $\hat{f}(\bar{\mathbf{p}})$ is computed in $O(1)$ the same way as (forward-rendering) unidirectional path tracing (see Algorithm 2).

5.2.2 Our bidirectional estimator. In addition to the *unidirectional* estimator, we also introduce a *bidirectional* one capable of handling challenging light-transport effects such as caustics.

Our *bidirectional* estimator samples paths the same way as forward-rendering bidirectional path tracing (BDPT) [Veach and Guibas 1995; Lafortune and Willems 1996]. Specifically, it generates a *source subpath* $\bar{\mathbf{p}}^S = (\mathbf{p}_0^S, \mathbf{p}_1^S, \dots)$ and a *detector subpath* $\bar{\mathbf{p}}^D = (\mathbf{p}_0^D, \mathbf{p}_1^D, \dots)$.




Fig. 12. **Equal-sample comparisons** between our *unidirectional* estimator with Redner_WAS [Bangaru et al. 2020]. This example involves a Lucy lit by a star-shaped area light. (top) We compare derivatives with respect to the horizontal displacement of the light; (bottom) We compare inverse-rendering results where we infer the location and the pose of the light by minimizing the L^2 loss between the target and rendered images shown in the row marked as “forward”. The configuration visualizations and parameter RMSE plots in the bottom row are used only for evaluation (and not for optimization).

Table 3. Performance statistics for our inverse-rendering results where “time” measures per-iteration optimization time (in seconds) on an Amazon EC2 c6a.8xlarge instance with 32 VCPUs. Additionally, “# img.” indicates the number of images used for inverse rendering, and “# param.” the number of scene parameters being optimized.

Scene		Time	# img.	# param.
Star emitter	(Fig. 12)	3.8	1	3
Chess	(Fig. 13)	4.5	40	30 000
Bunny in glass	(Fig. 14)	22.5	50	20 000
Lamp	(Fig. 15)	7.0	40	30 000
Dodoco	(Fig. 16)	8.4	40	30 000
Caustics	(Fig. 17)	21.3	1	12

Then, by connecting each $\mathbf{p}_m^S$ and $\mathbf{p}_n^D$ (for any m and n), we obtain a complete path sample $\bar{\mathbf{p}}_{m,n} := (\mathbf{p}_0^S, \dots, \mathbf{p}_m^S, \mathbf{p}_n^D, \dots, \mathbf{p}_0^D)$.

For efficient computation of the sum $\sum_K \nabla \cdot (\hat{f}_K \mathbf{v}_K)$ for each path $\bar{\mathbf{p}}_{m,n}$, we extend the process described in §5.2.1 by maintaining two prefix sums ($H_m^S : m = 0, 1, \dots$) and ($H_n^D : n = 0, 1, \dots$) where the former equals

$$H_m^S = H_{m-1}^S + h(\mathbf{p}_{m-1}^S; \mathbf{p}_{m-2}^S, \mathbf{p}_m^S), \quad (52)$$

and the latter follows Eq. (50).

With both pre-fix sums, we can compute the sum of divergences in Eq. (49) in $O(1)$ time using

$$\Sigma_{\text{div}}(\bar{\mathbf{p}}_{m,n}) = \hat{f}(\bar{\mathbf{p}}_{m,n}) \cdot \left(H_n^D + H_m^S + h(\mathbf{p}_m^S; \mathbf{p}_{m-1}^S, \mathbf{p}_n^D) + h(\mathbf{p}_n^D; \mathbf{p}_m^S, \mathbf{p}_{n-1}^D) \right), \quad (53)$$

resulting in the following bidirectional estimator

$$\left\langle \frac{dI}{d\theta} \right\rangle_{\text{bidir}} = \sum_{m,n} \frac{w_{\text{MIS}}(\bar{\mathbf{p}}_{m,n})}{\text{pdf}(\bar{\mathbf{p}}_{m,n})} \left(\frac{d}{d\theta} \hat{f}(\bar{\mathbf{p}}_{m,n}) + \Sigma_{\text{div}}(\bar{\mathbf{p}}_{m,n}) \right). \quad (54)$$

In this equation, for each m and n , the material measurement contribution $\hat{f}(\bar{\mathbf{p}}_{m,n})$, the probability density pdf($\bar{\mathbf{p}}_{m,n}$), and the multiple importance sampling (MIS) weight $w_{\text{MIS}}(\bar{\mathbf{p}}_{m,n})$ can all be computed in $O(1)$ time (with $O(N)$ precomputations) the same way as forward-rendering BDPT.

6 RESULTS

We implement our *unidirectional* and *bidirectional* estimators (introduced in §5.2.1 and §5.2.2, respectively) based on the differentiable renderer released by Yu et al. [2022]. This system is CPU-based and uses the Enzyme automatic differentiation framework [Moses and Churavy 2020].

In what follows, we first validate our estimators using several differentiable-rendering experiments in §6.1. Then, we compare




Fig. 13. We compare **differentiable- and inverse-rendering** results generated with: (a1) our *unidirectional* estimator with the new distance function in Eq. (16); (a2) our *unidirectional* estimator with Redner_WAS’s distance function in Eq. (57); (b1, b2) Redner_WAS with the two distance functions, respectively. The differentiable-rendering results (top) are computed with respect to the displacement of the chess piece. For inverse rendering (bottom), we optimize the shape of the chess piece by only looking at its cast shadow on the ground. We use 40 images among which the object has varying known poses and show one in the row marked “forward”. The mesh error information (plotted on the bottom right) is used for evaluation only.

our estimators to two main baselines—Bangaru et al.’s [2020] (indicated as Redner_WAS) and Zhang et al.’s [2020] (indicated as PSDR) methods—in §6.2 and §6.3, respectively. Lastly, we show additional inverse-rendering results in §6.4.

Please refer to Table 3 for performance statistics and the supplement for animated versions of our inverse-rendering results.

6.1 Validation

In Figure 11, we validate our technique by comparing derivative estimates generated using our method and finite differences (FD). The “Teapot” example shows the cast shadow of a teapot lit by an area light (with the configuration shown as inset), and the derivatives are with respect to the displacement of the teapot. The “Bunny” scene involves a diffuse bunny inside a Cornell box with an area light facing the ceiling, creating an indirect-illumination dominated configuration. For this example, we differentiate the ordinary rendering (a) with respect to the displacement of the bunny. In both examples, the derivative estimates produced by our estimators (c) closely match the FD references (b).

6.2 Comparisons with Redner_WAS

As discussed in Appendix A, although both our *unidirectional* estimator and Bangaru et al.’s [2020] method (indicated as Redner_WAS) apply unidirectional path tracing, the two methods differ significantly. We now demonstrate the practical advantages of our *unidirectional* estimator over Redner_WAS using several differentiable and inverse-rendering examples.

We conduct all comparisons with Redner_WAS using *equal sample* (instead of equal time) due to the significant differences between the two codebases.

Material form. The first advantage of our estimator is the use of the material-form reparameterization (described in §3.1). This allows our warped-area reparameterization to handle fewer types of discontinuities compared with Redner_WAS, leading to overall less variance and bias.

To demonstrate this, we compare derivatives estimated with our *unidirectional* estimator and Redner_WAS on the top of Figure 12. This example uses a “Star emitter” scene involving a Lucy lit by a star-shaped area emitter under direct illumination. We estimate




Fig. 14. We compare **inverse-rendering** performance of *bidirectional* estimators from our technique (c), PSDR [Zhang et al. 2020] (d), and PSDR_aq [Yan et al. 2022] (e). In this example, we optimize the shape of a diffuse object inside a rough-glass cube using 50 multi-view images (with one shown on the top). We use a multi-stage setup for the optimization, where each stage uses varying sample counts and learning rates. We visualize the initial, target, and optimized shapes on the bottom. Without guiding, our method outperforms PSDR and offers a similar level of performance as PSDR_aq (at equal time). The mesh error information (plotted on the bottom right) is used for evaluation only.

derivatives with respect to the horizontal displacement of the emitter. At equal sample, our *unidirectional* estimator produces derivatives with significantly less variance. This is because topological boundaries of the emitter—namely the edges of the star shape—do not need to be reparameterized by our method since they are fixed under the reference configuration. Redner_WAS, on the other hand, has to reparameterize these boundaries and, thus, produces higher variance.

Based on this configuration, we further show an inverse-rendering comparison on the bottom of Figure 12. For this example, we optimize the position and orientation of the star-shaped emitter without directly looking at it. Under identical optimization settings (including initializations and learning rates), our *unidirectional* estimator allows the optimization to converge nicely to the groundtruth. At equal sample, Redner_WAS suffers from much higher variance and causes the optimization to stuck.

Distance function. Another difference between our *unidirectional* and Redner_WAS—as detailed in Appendix A—is the choice of distance function D expressed in Eq. (16).

We demonstrate the effectiveness of our new distance function using differentiable- and inverse-rendering comparisons in Figure 13. This figure uses a “Chess” scene where a chess piece is lit by a small area light, resulting in a cast shadow on the ground.

The differentiable-rendering results indicate that, at equal time, our new distance function is crucial for fast convergence of derivative estimates.

For the inverse-rendering comparison, we use multiple shadow images with the object having varying known poses (with one shown in the figure) and optimize the shape (expressed as per-vertex positions) of the object. With our “new” distance function defined in Eq. (16), both our *unidirectional* estimator and Redner_WAS lead to

reasonable reconstruction results. Our method slightly outperforms Redner_WAS in this case due to the use of material-form reparameterization. In contrast, with the “old” distance function introduced by Bangaru et al. [2020], both methods produce severely biased gradient estimates, yielding low-quality shape reconstructions.

6.3 Comparisons with PSDR

We now compare our *bidirectional* estimator with the *bidirectional* one introduced by Zhang et al. [2020] (indicated as PSDR) as well as an improved variant (indicated as PSDR_aq) by Yan et al. [2022]. When estimating the material-form differential path integrals, both methods handle the *boundary* component by directly sampling material boundary paths. To improve efficiency, both PSDR and PSDR_aq perform primary-sample-space guiding where the former relies on regular grids and the latter leverages adaptive gridding. For all comparisons, we set the sample counts for all methods so that they take approximately equal time to estimate the gradients.

In Figure 14, we show a “bunny in glass” scene that was originally modeled by Yan et al. [2022] and contains a diffuse bunny inside a rough-glass cube. Using 50 multi-view images, we optimize the shape of the bunny. Further, we configure all optimizations to use identical initializations (i.e., a sphere) and learning rates as well as approximately equal time per iteration.

The complexity of this scene causes the optimization using PSDR to converge slowly (as shown in Figure 14-d). On the other hand, PSDR_aq’s guiding improves the rate of convergence considerably (see Figure 14-e). *Without using any precomputation or guiding*, our method is capable of offering a similar level of performance as PSDR_aq (see Figure 14-b).

Figure 15 shows another comparison with PSDR_aq. This example uses a “Lamp” scene where an object is lit by a desk lamp and




Fig. 15. We compare **differentiable- and inverse-rendering** performance of *bidirectional* estimators given by our technique and PSDR_aq [Yan et al. 2022]. The differentiable-rendering results (top) are computed with respect to the horizontal displacement of the object. For inverse rendering (bottom), we optimize the shape of the occluder by only looking at its cast shadow. We use 40 shadow images with the occluder having varying known poses (with one image shown). Without guiding, our method outperforms PSDR_aq at equal time. The mesh error information (plotted on the bottom right) is used for evaluation only.

casts a shadow on the desktop. Further, the area emitter is encapsulated within a glass bulb with low surface roughness, creating a very challenging situation for primary-sample-space guiding (since the high-value region is extremely small). This greatly reduces the effectiveness of guiding, causing PSDR_aq to produce more noisy derivative estimates, as demonstrated by the differentiable-rendering results (shown on the top of the figure).

For the inverse-rendering comparison (shown on the bottom), we use 40 images of the cast shadow with the object having multiple known poses (with one shown in the figure) and optimize the shape of the object. Using identical optimization configurations and equal time, our method outperforms PSDR_aq without the need of precomputation or guiding.

6.4 Additional Inverse-Rendering Results

We now show additional synthetic inverse-rendering results to demonstrate the usefulness of our *bidirectional* estimator.

Figure 16 uses a “Dodoco” scene that contains a diffuse object inside the Cornell box. The object is lit indirectly by a small area light facing the ceiling (similar to the “Bunny” scene in Figure 11). Using 40 images (with one shown in the figure), we optimize the shape of the object.

Additionally, Figure 17 use a “Caustics” scene with a glass bunny inside the Cornell box. We optimize the vertex positions of a cubical control cage by only looking at the ground.

Since both scenes involve complex light transport effects, Redner_WAS and our *unidirectional* estimator—both of which rely on unidirectional path tracing—perform poorly. On the contrary, our *bidirectional* estimator—which is enabled by our reparameterized differential path integral of Eq. (38)—offers the robustness and efficiency to simulate light transport in both scenes, allowing the optimizations to converge nicely.

7 DISCUSSION AND CONCLUSION

Limitations and future work. Our reparameterized differential path integral neglects volumetric light transport. Generalizing it to reparameterize Zhang et al.’s [2021b] generalized differential path integrals is an important future topic.

Also, our Monte Carlo estimators still mostly rely on unidirectional and bidirectional path tracing—techniques developed for forward rendering—to sample light paths. In the future, more efficient sampling schemes using, for instance, guiding, can be helpful.




Fig. 16. We compare **inverse-rendering** performance of our *bidirectional* and *unidirectional* estimators as well as Redner_WAS [Bangaru et al. 2020]. To ensure fairness, we use equal time between our *bidirectional* and *unidirectional* estimators, and equal sample between our *unidirectional* estimator and Redner_WAS.

Lastly, adapting our method to support differentiable rendering of implicit geometries [Vicini et al. 2022; Bangaru et al. 2022] is worth investigating.

Conclusion. In this paper, we bridge the gap in physics-based differentiable rendering techniques between the warped-area sampling methods and the path-space methods, by introducing the formulation of reparameterized differential path integral to reparameterize boundary components of material-form differential path integrals. Our resulting formulation enjoys the advantages of both differential path integrals—by enabling advanced Monte Carlo estimators (such as bidirectional path tracing)—and warped-area reparameterization—by not requiring explicit handling of discontinuity boundaries.

We validated our technique and evaluated its effectiveness using several differentiable rendering and inverse rendering examples.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive comments and Guangyan Cai, Wesley Chang, and Cheng Zhang for proofreading the paper. This project was partially funded by NSF grants 1900927 and 2105806.

REFERENCES

- Dejan Azinović, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. 2019. Inverse Path Tracing for Joint Material and Lighting Estimation. In *Computer Vision and Pattern Recognition*.
- Sai Bangaru, Michael Gharbi, Tzu-Mao Li, Fujun Luan, Kalyan Sunkavalli, Milos Hasan, Sai Bi, Zexiang Xu, Gilbert Bernstein, and Fredo Durand. 2022. Differentiable Rendering of Neural SDFs through Reparameterization. In *ACM SIGGRAPH Asia 2022 Conference Proceedings* (Daegu, Republic of Korea) (SIGGRAPH Asia '22). Association for Computing Machinery, New York, NY, USA, Article 22, 9 pages. <https://doi.org/10.1145/3550469.3555397>
- Sai Praveen Bangaru, Tzu-Mao Li, and Fredo Durand. 2020. Unbiased Warped-Area Sampling for Differentiable Rendering. *ACM Trans. Graph.* 39, 6 (2020), 245:1–245:18.
- Subrahmanyan Chandrasekhar. 1960. *Radiative transfer*. Courier Corporation.
- Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas. 2020. Towards Learning-based Inverse Subsurface Scattering. In *2020 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.

- Harley Flanders. 1973. Differentiation under the integral sign. *The American Mathematical Monthly* 80, 6 (1973), 615–627.
- James T. Kajiya. 1986. The Rendering Equation. In *SIGGRAPH '86*. 143–150.
- Eric P LaFortune and Yves D Willems. 1996. Rendering participating media with bidirectional path tracing. In *Rendering techniques '96*. Springer, 91–100.
- Tzu-Mao Li. 2023. redner: Differentiable rendering without approximation. <https://github.com/BachiLi/redner>.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable Monte Carlo ray tracing through edge sampling. *ACM Trans. Graph.* 37, 6 (2018), 222:1–222:11.
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans. Graph.* 38, 6 (2019), 228:1–228:14.
- Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. 2021. Unified Shape and SVBRDF Recovery using Differentiable Monte Carlo Rendering. *Computer Graphics Forum* 40, 4 (2021), 101–113.
- William Moses and Valentin Churavy. 2020. Instead of Rewriting Foreign Code for Machine Learning, Automatically Synthesize Fast Gradients. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12472–12485.
- Merlin Nimier-David, Thomas Müller, Alexander Keller, and Wenzel Jakob. 2022. Unbiased Inverse Volume Rendering with Differential Trackers. *ACM Trans. Graph.* 41, 4 (2022), 44:1–44:20.
- Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. 2020. Radiative backpropagation: an adjoint method for lightning-fast differentiable rendering. *ACM Trans. Graph.* 39, 4 (2020), 146:1–146:15.
- Osborne Reynolds. 1903. *Papers on mechanical and physical subjects: the sub-mechanics of the universe*. Vol. 3. The University Press.
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Vol. 1610. Stanford University PhD thesis.
- Eric Veach and Leonidas Guibas. 1995. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*. Springer, 145–167.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2021. Path Replay Backpropagation: Differentiating Light Paths Using Constant Memory and Linear Time. *ACM Trans. Graph.* 40, 4, Article 108 (2021), 108:1–108:14 pages.
- Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2022. Differentiable Signed Distance Function Rendering. *ACM Trans. Graph.* 41, 4 (2022), 125:1–125:18.
- Kai Yan, Christoph Lassner, Brian Budge, Zhao Dong, and Shuang Zhao. 2022. Efficient estimation of boundary integrals for path-space differentiable rendering. *ACM Trans. Graph.* 41, 4 (2022), 123:1–123:13.
- Zihan Yu, Cheng Zhang, Derek Nowrouzezahrai, Zhao Dong, and Shuang Zhao. 2022. Efficient Differentiation of Pixel Reconstruction Filters for Path-Space Differentiable Rendering. *ACM Trans. Graph.* 41, 6 (2022), 191:1–191:16.
- Tizian Zeltner, Sébastien Speierer, Iliyan Georgiev, and Wenzel Jakob. 2021. Monte Carlo estimators for differential light transport. *ACM Trans. Graph.* 40, 4 (2021), 78:1–78:16.




Fig. 17. We compare **inverse-rendering** performance between our *unidirectional* and *bidirectional* estimators. We adjust sample counts so that both estimators take approximately equal time per iteration.

- Cheng Zhang, Zhao Dong, Michael Doggett, and Shuang Zhao. 2021a. Antithetic sampling for Monte Carlo differentiable rendering. *ACM Trans. Graph.* 40, 4 (2021), 77:1–77:12.
- Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. 2020. Path-space differentiable rendering. *ACM Trans. Graph.* 39, 4 (2020), 143:1–143:19.
- Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A differential theory of radiative transfer. *ACM Trans. Graph.* 38, 6 (2019), 227:1–227:16.
- Cheng Zhang, Zihan Yu, and Shuang Zhao. 2021b. Path-space differentiable rendering of participating media. *ACM Trans. Graph.* 40, 4 (2021), 76:1–76:15.
- Shuang Zhao, Wenzel Jakob, and Tzu-Mao Li. 2020. Physics-Based Differentiable Rendering: A Comprehensive Introduction. In *ACM SIGGRAPH 2020 Courses*. Article 14, 14:1–14:30 pages.

A RELATION WITH ORIGINAL WARPED-AREA REPARAMETERIZATION

In the following, we discuss the relation between our *unidirectional* method and Redner_WAS—the original warped-area technique introduced by Bangaru et al. [2020]. For simplicity, we use the one-bounce setting discussed in §3.2.

When expressing one-bounce scattered radiance leaving a surface point $\mathbf{y}$ in the direction ω_o , Redner_WAS uses the spherical-integral formulation:

$$I_{\text{dir}} = \int_{\mathbb{S}^2} L_{e,i}(\mathbf{y}, \omega_i) f_s(\mathbf{y}, \omega_i, \omega_o) |\mathbf{n}(\mathbf{y}) \cdot \omega_i| d\sigma(\omega_i), \quad (55)$$

$\underbrace{\hspace{10em}}_{=: F_{\text{dir}}^\sigma(\omega_i)}$

where $L_{e,i}(\mathbf{y}, \omega_i)$ indicates emitted radiance entering $\mathbf{y}$ from the incident direction ω_i , $\mathbf{n}(\mathbf{y})$ denotes the (unit-length) surface normal at $\mathbf{y}$, and σ is the solid-angle measure.

Differentiating Eq. (55) with respect to some scene parameter θ yields

$$\frac{dI_{\text{dir}}}{d\theta} = \underbrace{\int_{\mathbb{S}^2} \frac{dF_{\text{dir}}^\sigma(\omega_i)}{d\theta} d\sigma(\omega_i)}_{\text{interior}} + \underbrace{\int_{\Delta\mathbb{S}^2} \Delta F_{\text{dir}}^\sigma(\omega_i) V(\omega_i) d\ell(\omega_i)}_{\text{boundary}}, \quad (56)$$

where:

- F_{dir}^σ is the integrand defined in Eq. (55);
- $\Delta\mathbb{S}^2$ are *spherical curves* comprising jump discontinuities of F_{dir}^σ ;
- $V(\omega_i) = \mathbf{n}^\theta(\omega_i) \cdot \frac{d\omega_i}{d\theta}$ is the scalar velocity capturing how fast the curves $\Delta\mathbb{S}^2$ evolves with respect to the parameter θ .

Mathematically, Redner_WAS reparameterizes the *boundary* component of Eq. (56) that differs from our *boundary* integral in Eq. (4) in several significant ways as follows.

First, since the material-form reparameterization (§3.1) is not used in Eq. (56), the discontinuity curves $\Delta\mathbb{S}^2$ need to include topological boundaries and discontinuities of surface normals (see Figure 18) of the emitter surface. As demonstrated in Fig. 12, this makes their method much less efficient than ours when handling emitters with complex geometries.

Second, the scalar velocity $V(\omega_i)$ is computed differently: (i) the curve normal $\mathbf{n}^\theta(\omega_i)$ is defined with respect to spherical curves over $\mathbb{S}^2$ (as opposed to general curves over the reference surface $\mathcal{B}_e$); and (ii) the velocity $d\omega_i/d\theta$ is computed with respect to direction ω_i and without differentiating the deformation’s inverse $X^{-1}(\cdot, \theta)$ as in Eq. (8).

At a more technical level, Redner_WAS uses a different distance function given by:

$$D^{\text{was}}(\omega'_i; \omega_i) := \exp(1 - \langle \omega_i, \omega'_i \rangle) - 1. \quad (57)$$

When handling challenging scenes containing complicated geometries, small area lights, or intersecting discontinuity boundaries (with different velocities), the convergence of the consistent estimator $\langle v_K \rangle$ of the continuous velocity v_K in Eq. (47) is known to be sensitive to the choice of distance functions [Bangaru et al. 2022]. In practice, our distance function expressed in Eq. (16) behaves well in all our experiments (including the example shown in Figure 13) with $N_q = 8$ auxiliary samples. The one in Eq. (57) used by Redner_WAS, on the other hand, can lead to convergence issues even with significantly greater N_q .




Fig. 18. **Equal-sample comparison** between our *unidirectional* estimator and Redner_WAS [Bangaru et al. 2020]. This example uses a scene with a lucy model lit by a cube-shaped area light. The formulation used by Redner_WAS requires the cube edges (where surface normals are discontinuous) to be handled by warped-area reparameterization. However, the public implementation of Redner_WAS [Li 2023] neglects this case and, thus, produces noisy and *highly biased* gradients (c). Our technique uses the material-form reparameterization (§3.1) and does not need to reparameterize the cube edges (as they are fixed on the reference surface). This allows our estimator to produce clean gradients (b) that closely match the reference (a).

B MONTE CARLO ESTIMATION OF PER-VERTEX DIVERGENCE

Provided a material path $\bar{p} = (p_0, \dots, p_N)$ and an index $0 \leq K < N$, we now describe the estimation of the divergence $[\nabla \cdot v_K](p_K)$ of the continuous velocity v_K defined in Eq. (41).

Estimating ∇W and $\nabla \cdot \tilde{v}_K$. It is easy to verify that our kernel w defined in Eq. (15) and the discontinuous vector field v_K^{dis} are differentiable everywhere except at their zero-measure jump discontinuities (that are fixed with respect to p_K). Thus, provided Eqs. (42) and (43), we have

$$[\nabla W](p_K) = \int_{\mathcal{B}} \nabla w \, dA, \quad (58)$$

$$[\nabla \cdot \tilde{v}_K](p_K) = \int_{\mathcal{B}} \left((\nabla w) \cdot v_K^{\text{dis}} + w (\nabla \cdot v_K^{\text{dis}}) \right) dA. \quad (59)$$

Using the point-sampling process described in §5.1.2, we estimate ∇W and $\nabla \cdot \tilde{v}_K$ in an unbiased and consistent fashion via

$$\langle [\nabla W](p_K) \rangle = \sum_{i=1}^{N_q} \sum_j \frac{\nabla w(q_{ij}; p_K)}{\text{pdf}(q_{ij})}, \quad (60)$$

and

$$\begin{aligned} \langle [\nabla \cdot \tilde{v}_K](p_K) \rangle &= \\ \sum_{i=1}^{N_q} \sum_j \frac{\nabla w(q_{ij}; p_K) \cdot v_K^{\text{dis}}(q_{ij}) + w(q_{ij}; p_K) [\nabla \cdot v_K^{\text{dis}}](q_{ij})}{\text{pdf}(q_{ij})}, \end{aligned} \quad (61)$$

where the gradient ∇w of the kernel w and the divergence $\nabla \cdot v_K^{\text{dis}}$ of the discontinuous velocity v_K^{dis} can both be computed using automatic differentiation.

Estimating $\nabla \cdot v_K$. Finally, because

$$\nabla \cdot v_K = (\nabla \cdot \tilde{v}_K)/W - ((\nabla W) \cdot \tilde{v}_K)/W^2,$$

we obtain a biased-but-consistent estimator of the divergence $\nabla \cdot v_K$ using Eqs. (60) and (61):

$$\begin{aligned} \langle [\nabla \cdot v_K](p_K) \rangle &= \\ \frac{\langle [\nabla \cdot \tilde{v}_K](p_K) \rangle}{\langle W(p_K) \rangle} - \frac{\langle [\nabla W](p_K) \rangle \cdot \langle \tilde{v}_K(p_K) \rangle}{\langle W(p_K) \rangle^2}, \end{aligned} \quad (62)$$

where $\{q_{ij}\}$ are samples of q drawn with the procedure explained in §5.1.2.