Gradient-domain Volumetric Photon Density Estimation

ADRIEN GRUSON, The University of Tokyo and JFLI CNRS UMI 3527
BINH-SON HUA, The University of Tokyo and Singapore University of Technology and Design
NICOLAS VIBERT, McGill University
DEREK NOWROUZEZAHRAI, McGill University
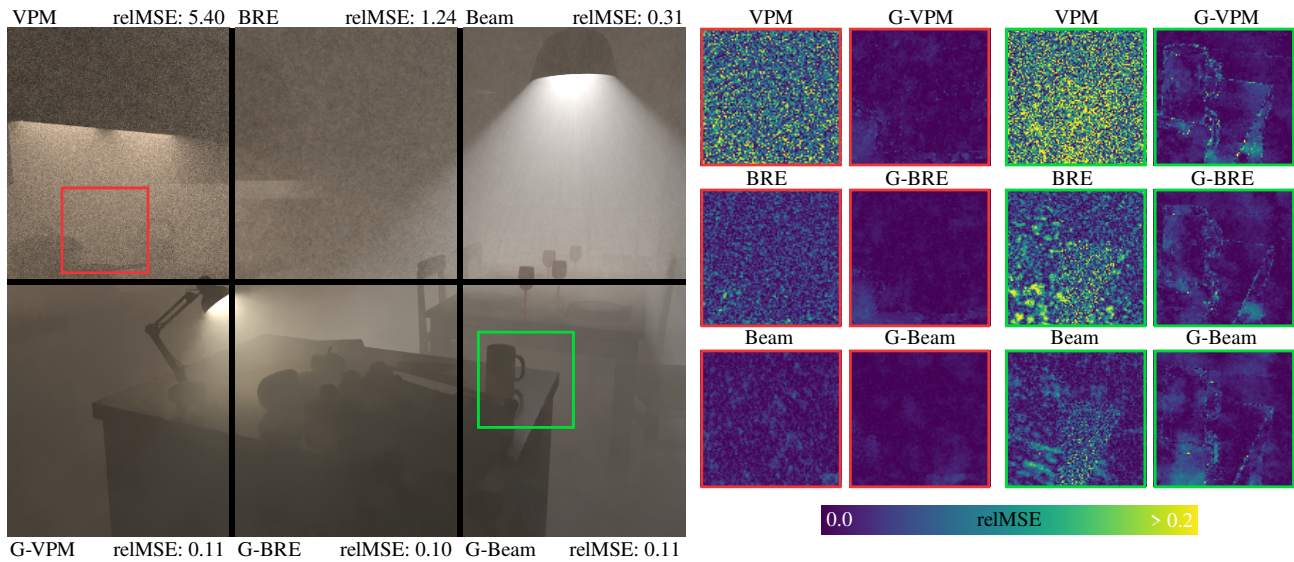TOSHIYA HACHISUKA, The University of Tokyo

Fig. 1. Equal-time (5 minutes) renderings of a smoky KITCHEN scene. Gradient-domain volumetric rendering techniques with L1 reconstruction converge faster than primal-domain volumetric rendering technique. The relMSE error metric has a unitless scale of 10^{-2}.

Gradient-domain rendering can improve the convergence of surface-based light transport by exploiting smoothness in image space. Scenes with participating media exhibit similar smoothness and could potentially benefit from gradient-domain techniques. We introduce the first gradient-domain formulation of image synthesis with homogeneous participating media, including four novel and efficient gradient-domain volumetric density estimation algorithms. We show that naïve extensions of gradient domain path-space and density estimation methods to volumetric media, while functional, can result in inefficient estimators. Focussing on point-, beam- and plane-based gradient-domain estimators, we introduce a novel shift mapping that eliminates redundancies in the naïve formulations using spatial relaxation within the volume. We show that gradient-domain volumetric rendering improve convergence compared to primal domain state-of-the-art, across a suite of scenes. Our formulation and algorithms support progressive estimation and are easy to incorporate atop existing renderers.

CCS Concepts: • **Computing methodologies → Ray tracing**;

Additional Key Words and Phrases: gradient rendering, participating media

1 INTRODUCTION

Accurately simulating light transport in participating media remains a longstanding challenge in computer graphics. While Monte Carlo methods robustly render a range of surface-only scenes, density estimation techniques tend towards more desirable convergence behavior in scenes with volumetric media. Recent advances in volumetric density estimation, such as the beam radiance estimate [Jarosz et al. 2008], photon point and beam estimators [Jarosz et al. 2011a], virtual ray and beam lights [Novák et al. 2012a,b], and other higher-dimensional representation such as photon planes and volumes [Bitterli and Jarosz 2017], further improve this behavior.

Several recent extensions of conventional rendering methods to the gradient domain have greatly improved visual convergence [Kettunen et al. 2015; Lehtinen et al. 2013; Manzi et al. 2015]. The seminal work on gradient-domain Metropolis light transport [Lehtinen et al. 2013] led to gradient-domain variants of uni- and bi-directional path

tracing [Kettunen et al. 2015; Manzi et al. 2015], as well as gradient-domain density estimation on surfaces [Hua et al. 2017; Sun et al. 2017]. These methods directly estimate image *gradients*, instead of pixel intensities, and apply Poisson reconstruction to generate the final image. Gradient-domain approaches tend to be more efficient than their conventional counterparts since images tend to be piecewise smooth in many scenes. Existing gradient-domain methods, however, only model and treat *surface light transport*. Scenes with participating media also often exhibit image-space smoothness that could potentially benefit from gradient-domain rendering.

We present the first gradient-domain formulation of volumetric rendering, including several progressive algorithm for rendering scenes with homogeneous participating media. As in surfaces light transport, we exploit the coherence of light paths by introducing a novel shift mapping within the volume. A key difference is that we carefully take distance sampling, an operation unique to volumetric rendering, into account in shift mapping to form an efficient gradient estimator in volumes. Our gradient estimators are *progressive* and generalize several primal-domain volumetric density estimators to the gradient-domain. We show that our gradient-domain volumetric rendering algorithms improve the efficiency compared to the primal-domain estimators. Concretely, our contributions are:

- the first image-space gradient formulation for rendering in the presence of homogeneous participating media,
- novel shift mapping techniques that account for the additional distance dimension in volumetric light transport, and
- several new gradient-domain volume rendering algorithms using photons, beams and planes: specifically, point-point, point-beam, beam-beam and beam-plane gradient-domain estimators.

For comprehensive evaluation, we also extend uni- and bi-directional gradient-domain path tracing to homogeneous participating media; while these new gradient-domain path-based methods are not core contributions, we are the first to prototype them, to our knowledge.

Figure 1 illustrates representative results in a scene with diffuse and specular volumetric transport. We compare *volumetric photon mapping* (VPM) [Jensen and Christensen 1998], *beam radiance estimate* (BRE) [Jarosz et al. 2008] and *photon beams* (Beam) [Jarosz et al. 2011b] with our gradient-domain methods. In the experiments, we show that our gradient-domain methods can consistently produce smoother results compared to primal-domain rendering methods.

2 RELATED WORK

Monte Carlo Integration. Monte Carlo integration can be used to solve the radiative transfer equation (RTE) [Chandrasekhar 2013], and its application to rendering leads to the uni- and bi-directional volumetric path tracing algorithms [Lafortune and Willems 1996] that generate paths to connect sensors to lights. Kulla and Fajardo [2012] improve light sampling techniques for single-scattering, and Georgiev et al. [2013] extended this idea to lower-order multiple scattering. Monte Carlo integration is very general but its convergence can be slow, even in geometrically (and visually) simple scenes, e.g., a Cornell Box with homogeneous media. We build atop density estimation methods that often converge faster in volumetric media.

Density Estimation. Jensen and Christensen [1998] proposed a density estimation algorithm for volumetric light transport that traces light paths to deposit *volume photons* in a scene, before integrating their density along camera rays to synthesize a final image. While more expensive than surface photon mapping [Jensen 2001], this *volumetric photon mapping* is competitive with Monte Carlo counterparts. Jarosz et al.'s *beam radiance estimate* (BRE) [2008] performs density estimation directly along *an entire camera ray*, leading to significant performance improvements. A recent extension of photon point primitives to beam primitives, including a theory of density estimation using these *photon beams*, has led to improvements when rendering scenes in certain scenarios [Jarosz et al. 2011a]: Krivanek et al.'s [2014] variance analysis of various point- and beam-based estimators concludes that point-based estimators are best-suited to optically thick media, whereas beam primitive-based estimators are beneficial in optically thin media. More recently, Bitterli and Jarosz [2017] proposed to use higher order primitives (like plane or volume) to get smoother results of multiple-scattering inside the participating media.

Progressive density estimation [Hachisuka and Jensen 2009; Knaus and Zwicker 2011] effectively eliminates the memory constraints of traditional density estimators, and their extensions to volumes using points [Hachisuka and Jensen 2009] and beams [Jarosz et al. 2011b] result in similar gains. We present the first gradient-domain progressive volumetric density estimators using photon points, the BRE, photon beams and photon planes.

Gradient-domain Rendering. Gradient-domain methods directly estimate image-space gradients using coherent paths, and then reconstruct a final image with Poisson reconstruction [Bhat et al. 2008]. The seminal work of gradient-domain Metropolis light transport [Lehtinen et al. 2013] inspired several relevant follow-up works. Kettunen et al. [2015] formulated the image gradient computation as a Monte Carlo estimator by extending path tracing to the gradient-domain. Manzi et al. [2015] similarly extended bi-directional path tracing and keeping the gradient computation overhead low by ignoring the connection between non-diffuse vertices. A temporal extension of gradient-domain path tracing was also recently proposed by Manzi et al. [2016a] in order to render image *sequences* more efficiently. Bauszat et al. [2017] also extended gradient-domain path tracing to leverage path reuse. All such techniques estimate gradients by using paths constructed by Monte Carlo path tracing.

Hua et al. [2017] proposed gradient-domain photon density estimation on surfaces. They use a two-step shift mapping to shift both the camera and photon subpaths before performing density estimation. Sun et al. [2017] combined gradient-domain bidirectional path tracing and density estimation into the same framework.

We extend gradient-domain photon density estimation to volumetric light transport. We show that a trivial extension of gradient-domain surface volumetric density estimation does not result in an efficient rendering technique, and that special care needs to be taken both when formulating the problem and devising generalized volumetric shift mappings. We propose a novel shift mapping variant well-suited for gradient computation in the presence of volumetric transport that takes the advantage of photon density kernel support to further increase performance.

Table 1. Table of notation.

Symbol	Meaning
$\mathbf{z}_k, \mathbf{y}_k$	vertex k on a sensor and light subpath
$\mathbf{z}_t, \mathbf{z}'_t$	base and offset gather points
$\mathbf{y}_s, \mathbf{y}'_s$	base and offset photons
$f(\bar{x})$	measurement contribution function for a path $\bar{x}$
$\rho(\mathbf{x}_k)$	abstract vertex interaction at $\mathbf{x}_k$
$\tau(\mathbf{x}_k, \mathbf{x}_{k+1})$	abstract edge transmittance from $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$
$\sigma_t, \sigma_s, \sigma_a$	extinction, scattering and absorption coefficients
E, L, D, S, M	camera, light, diffuse, specular & media interactions
$T(\bar{x})$	shift mapping operator on a path $\bar{x}$

3 BACKGROUND

We propose new algorithms to render scenes with participating media using progressive gradient-domain volumetric photon density estimation. While conceptually simple, extending gradient-domain rendering formulations to participating media *and* to density estimation requires addressing two important challenges.

First, since existing gradient-domain rendering formulations rely on *Monte Carlo* integration and *surface-only transport*, we must devise a new gradient-domain formulation that is suitable for *density estimation* and *volumetric transport*. Second, a trivial generalization of surface-based shift mapping to volumes would result in inefficient estimators. We identify that this is mainly due to the additional distance sampling in volumes. To alleviate this, we propose a new shift operator based on spatial relaxation of volume photons that allows us to skip a large percentage of expensive shift mappings during gradient estimation, resulting in significantly faster convergence.

3.1 Volume Light Transport with Density Estimation

We begin with the unified path integral formulation for surfaces and volumes [Jakob 2013] which states that:

$$I_i = \int_{\mathcal{P}} h_i(\bar{x}) f(\bar{x}) \mathrm{d}\mu(\bar{x}), \tag{1}$$

where $\bar{x}$ is a light path, f is the measurement contribution function, and μ is the measure of paths in path space $\mathcal{P}$. The filter h_i determines the contribution of path $\bar{x}$ to pixel i. To solve this path integral numerically, uni- and bi-directional path tracing build paths by sampling vertices sequentially. Specifically, for a length-k path with $k+1$ vertices $\mathbf{x}_0 \ldots \mathbf{x}_k$, the measurement contribution function can be written as $f(\bar{x}) = \rho(\bar{x}) G(\bar{x})$ with

$$\rho(\bar{x}) = \prod_{i=0}^{k} \rho(\mathbf{x}_i) \prod_{i=0}^{k-1} \tau(\mathbf{x}_i, \mathbf{x}_{i+1}) \text{ and } G(\bar{x}) = \prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}), \tag{2}$$

where $\rho(\mathbf{x})$ models the surface or medium interaction at a vertex, τ models the transmittance value between two vertex, and G models the geometry factors. Specifically, in a volume, the vertex interaction function ρ from Equation 2 is

$$\rho(\mathbf{x}_i) = \begin{cases} L_e(\mathbf{x}_0 \to \mathbf{x}_1) & \text{if } i = 0 \\ W_e(\mathbf{x}_{k-1} \to \mathbf{x}_k) & \text{if } i = k \\ \rho_s(\mathbf{x}_{i-1} \to \mathbf{x}_i, \mathbf{x}_i \to \mathbf{x}_{i+1}) & \text{if } \mathbf{x}_i \text{ is on a surface} \\ \sigma_s(\mathbf{x}_i)\rho_p(\mathbf{x}_{i-1} \to \mathbf{x}_i, \mathbf{x}_i \to \mathbf{x}_{i+1}) & \text{otherwise,} \end{cases}$$
$$\tag{3}$$

where L_e and W_e are light emission and camera importance, σ_s the scattering coefficient, ρ_s the BSDF and ρ_p the phase function. The transmittance along an edge is $\tau(\mathbf{x}_i, \mathbf{x}_{i+1})$. The geometric term is

$$G(\mathbf{x}_i, \mathbf{x}_{i+1}) = \cos_\perp(\mathbf{x}_i) \cos_\perp(\mathbf{x}_{i+1}) \Big/ \|\mathbf{x}_{i+1} - \mathbf{x}_i\|^2, \tag{4}$$

where $\cos_\perp(\cdot)$ is the cosine of the angle between the incident light direction and the surface normal for surface scattering events, and simply 1 at volumetric scattering events.

Given recent advances, photon density estimation becomes an attractive and efficient alternative to Monte Carlo integration, especially for scenes dominated by volumetric transport and paths that remain on a localized manifold in the path space (i.e., SDS events). A path $\bar{x}$ can be constructed by connecting a camera subpath $\bar{z}$ and light subpath $\bar{y}$ with a blurring kernel. We follow the comprehensive theory of radiance estimation in volumes using photons points and beams [Jarosz et al. 2011a] and a recent extension to photon planes [Bitterli and Jarosz 2017] and consider four typical estimators (in camera ray query × photon data order): point-point, beam-point, beam-beam, and beam-plane estimators. In general, the major difference among these estimators are how the last vertices of both subpaths are sampled, resulting in different strategies to generate the last two vertices on the camera and light subpath [Hachisuka et al. 2017]. For brevity, here we let the last two vertices on the light subpath be $\mathbf{a} = \mathbf{y}_{s-1}, \mathbf{b} = \mathbf{y}_s$, and those on the camera subpath be $\mathbf{d} = \mathbf{z}_t, \mathbf{c} = \mathbf{z}_{t-1}$. Path vertices can therefore be grouped into three parts: $\bar{x} = (\mathbf{y}_0 \ldots \mathbf{abdc} \ldots \mathbf{z}_0)$, where $\mathbf{y}_0 \ldots \mathbf{ab}$ is the light subpath, $\mathbf{dc} \ldots \mathbf{z}_0$ the camera subpath, and $\mathbf{abdc}$ the vertices participating in kernel sampling (See Figure 6 for an illustration). Applying Monte Carlo estimation and density estimation, the path contribution is:

$$\frac{f(\bar{x})}{p(\bar{x})} = \frac{f(\mathbf{y}_0 \ldots \mathbf{a})}{p(\mathbf{y}_0 \ldots \mathbf{a})} \frac{f(\mathbf{z}_0 \ldots \mathbf{c})}{p(\mathbf{z}_0 \ldots \mathbf{c})} D(\mathbf{a}, \mathbf{b}, \mathbf{d}, \mathbf{c}) \tag{5}$$

where the subpath contributions $f(\mathbf{y}_0 \ldots \mathbf{a})$ and $f(\mathbf{z}_0 \ldots \mathbf{c})$ can be defined as in Equation 2. Note that vertex interactions at $\mathbf{a}$ and $\mathbf{c}$ are included in the subpath contributions. The density estimator D computes the contribution of the last light edge $\mathbf{ab}$ and the last camera edge $\mathbf{dc}$:

$$D(\mathbf{a}, \mathbf{b}, \mathbf{d}, \mathbf{c}) = \frac{\tau(\mathbf{a}, \mathbf{b}) G(\mathbf{a}, \mathbf{b}) \rho(\mathbf{a}, \mathbf{b}, \mathbf{d}, \mathbf{c}) G(\mathbf{c}, \mathbf{d}) \tau(\mathbf{d}, \mathbf{c}) K(\mathbf{a}, \mathbf{b}, \mathbf{d}, \mathbf{c})}{p(\mathbf{b}, \mathbf{d} \mid \mathbf{y}0, \ldots, \mathbf{a}, \mathbf{c}, \ldots, \mathbf{z}_0)}$$

With this convention, we adapt the extended path integral for volumes [Hachisuka et al. 2017] to evaluate path contributions with different kernels, which include photon points and photon beams.

3.2 Gradient-domain Rendering

Without loss of generality, let the image-space gradient Δ for pixel i be estimated using finite differences with its neighboring pixel $i+1$, as $\Delta_i = I_{i+1} - I_i$. Separately estimating I_i and I_{i+1} with Monte Carlo integration will cause gradients to have the same level of variance as the image itself, and gradient-domain rendering will not yield significant improvement after reconstruction. Instead *shift mapping* is used to correlate the paths passing through pixels i and $i+1$, resulting in an estimate of the gradients that improves the reconstructed image quality.

(a) Generate the offset camera path
and sample the offset gather point

(b) Generate the offset photon position
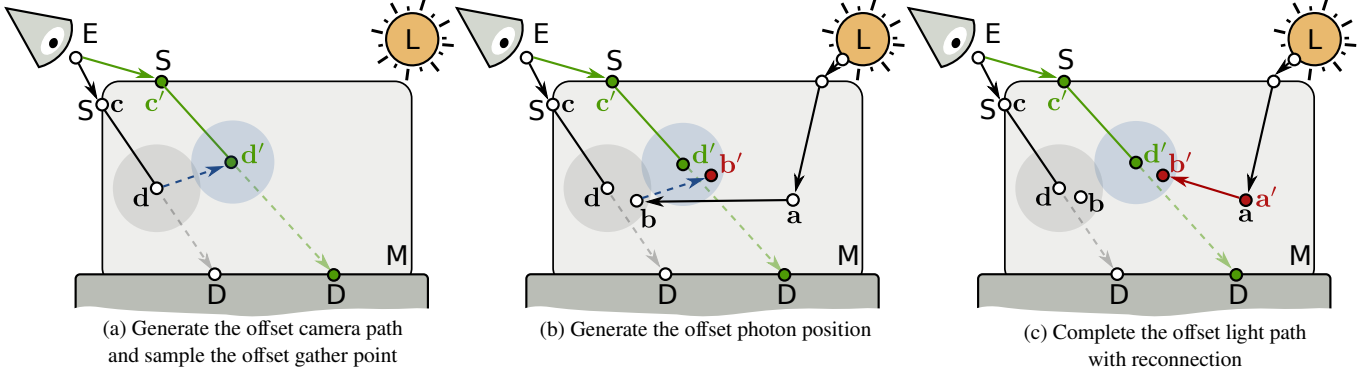
(c) Complete the offset light path
with reconnection

Fig. 2. A photon density estimation path is shifted in three steps: first, the camera subpath is shifted and the offset gather point on the offset path is generated; next, for each photon collected around the base gather point, we generate a offset photon within the kernel support of the offset gather point; finally, we complete the offset photon path by performing reconnection. Blue arrows indicate a shift and red arrows a reconnection.

Shift Mapping. The goal of *shift mapping* [Lehtinen et al. 2013] is thus to construct a new path, *through* a neighboring pixel, by modifying the original path as little as possible (see Figure 2).

We can express the intensity of pixel $i + 1$ using paths of pixel i and a shift mapping function T as

$$I_{i+1} = \int_{\mathcal{P}} h_{i+1}(T(\bar{x})) f(T(\bar{x})) \left| \frac{\partial T(\bar{x})}{\partial \bar{x}} \right| d\mu(\bar{x}) , \qquad (6)$$

where T is the shift mapping that transforms paths of pixel i into paths of pixel $i + 1$, and the Jacobian accounts for the change in the path density that occurs during the shift. Since the reconstruction filter value remains unchanged in the shift, we have $h_{i+1}(T(\bar{x})) = h_i(\bar{x})$, and the gradient can be computed as

$$\Delta_i = \int_{\mathcal{P}} h_i(\bar{x}) \left(f(T(\bar{x})) \left| \frac{\partial T(\bar{x})}{\partial \bar{x}} \right| - f(\bar{x}) \right) d\mu(\bar{x}) . \qquad (7)$$

In general, we want to keep the variance of the gradient estimator as small as possible by (1) keeping base and offset paths as similar as possible, (2) avoiding non-invertible offset paths, and (3) avoiding singularities in the Jacobian. For the last point, we apply the multiple importance sampling (MIS) solution proposed by Kettunen et al. [2015], where the gradient at pixel i is estimated using a two-way shift from pixel i to j and vice versa as

$$\Delta_i = \int_{\mathcal{P}} w(\bar{x}) h_i(\bar{x}) \left(f(T(\bar{x})) \left| \frac{\partial T(\bar{x})}{\partial \bar{x}} \right| - f(\bar{x}) \right) d\mu(\bar{x}) \qquad (8)$$

$$+ \int_{\mathcal{P}} w^{-1}(\bar{x}) h_{i+1}(\bar{x}) \left(f(\bar{x}) - f(T^{-1}(\bar{x})) \left| \frac{\partial T^{-1}(\bar{x})}{\partial \bar{x}} \right| \right) d\mu(\bar{x}) ,$$

where T^{-1} is the inverse shift mapping that transforms paths from pixel $i + 1$ to pixel i, w and w^{-1} are MIS weights of the forward and inverse shift mapping such that $w(\bar{x}) + w^{-1}(T(\bar{x})) = 1$. Note that to apply MIS (Section 6), the shift must be invertible, otherwise we need to devise MIS weights that maintain a unit sum.

After estimating gradients, we perform a Poisson reconstruction:

$$I = \arg\min \left\| \begin{matrix} D_x I - G_x \\ D_y I - G_y \end{matrix} \right\|_p^p + \|\lambda(I - I_0)\|_p^p , \qquad (9)$$

where D_x and D_y are convolution filters that compute finite differences in the image I, and G_x and G_y are the estimated gradients.

The λ parameter is a regularization parameter that controls the contribution of the primal-domain image, and the parameter p defines the Lp-norm in the error metric. Lehtinen et al. [2013] showed that the L2 metric yields an unbiased reconstruction but is sensitive to discrepancies between the primal and gradient images, whereas an L1 metric is robust to these discrepancies but introduces a bias.

Gradient-domain Photon Density Estimation. To apply photon density estimation to gradient-domain rendering, we have to consider shifting the light and camera subpath and the kernel vertices (**abdc**) that fall within the support of a density estimation kernel K. Similar to the primal domain, the offset path has the configuration

$$\bar{x}' = (\mathbf{y}_0' \dots \mathbf{a}' \mathbf{b}' \mathbf{d}' \mathbf{c}' \dots \mathbf{z}_0'), \qquad (10)$$

Estimating Equation 6 with Monte Carlo, we have

$$\frac{f(\bar{x}')}{p(\bar{x})} \left| \frac{\partial \bar{x}'}{\partial \bar{x}} \right| = \frac{f(\mathbf{y}_0' \dots \mathbf{a}')}{p(\mathbf{y}_0 \dots \mathbf{a})} \frac{f(\mathbf{c}' \dots \mathbf{z}_0')}{p(\mathbf{c} \dots \mathbf{z}_0)} \left| \frac{\partial(\mathbf{y}_0' \dots \mathbf{a}')}{\partial(\mathbf{y}_0 \dots \mathbf{a})} \right| \left| \frac{\partial(\mathbf{c}' \dots \mathbf{z}_0')}{\partial(\mathbf{c} \dots \mathbf{z}_0)} \right|$$

$$\cdot D(\mathbf{a}', \mathbf{b}', \mathbf{d}', \mathbf{c}') \qquad (11)$$

Here, we evaluate the offset path contribution weighted by the probability density of the *base* path and the shift's Jacobian.

For efficiency, we need to ensure that $f(\bar{z}') > 0$ and $f(\bar{y}') > 0$ and avoid zero kernel contribution. Hua et al. [2017] proposed a shift operator to resolve this challenge in the case of surface light transport, as they shift the relative configuration of the base path endpoints (**b** and **d**) while enforcing $K(\mathbf{b}', \mathbf{d}') \neq 0$. In order to construct a full path with the shifted endpoints (**b**' and **d**'), they apply a half-vector copy [Kettunen et al. 2015] to shift the camera subpath and then use a diffuse reconnection or manifold exploration [Jakob and Marschner 2012] to shift the light subpath.

Our approach also uses the concept of shifting an entire configuration of endpoints, but we need to additionally address complications that arise due to distance integration and importance sampling in volumetric light transport, as well as different density estimation kernel dimensionalities (estimator D in Equation 11). The density estimator of the offset path is

$$D(\mathbf{a}', \mathbf{b}', \mathbf{d}', \mathbf{c}') = \tau(\mathbf{a}', \mathbf{b}')G(\mathbf{a}', \mathbf{b}')\rho(\mathbf{a}', \mathbf{b}', \mathbf{d}', \mathbf{c}')G(\mathbf{c}', \mathbf{d}')\tau(\mathbf{d}', \mathbf{c}')$$

$$\cdot \frac{K(\mathbf{a}', \mathbf{b}', \mathbf{d}', \mathbf{c}')}{p(\mathbf{b}, \mathbf{d} \mid \mathbf{y}0, \dots, \mathbf{a}, \mathbf{c}, \dots, \mathbf{z}_0)} \left| \frac{\partial(\mathbf{b}' \mathbf{d}')}{\partial(\mathbf{b} \mathbf{d})} \right| \qquad (12)$$

The probability of kernel sampling $p(\mathbf{b}, \mathbf{d})$ and its Jacobian could be derived depending on which density estimation scheme is used, e.g., point-point, point-beam, beam-beam, plane-beam.

Equation 12 is a theoretical formula for estimating the contribution at the last two vertices of the offset camera and light subpaths, where shift mapping jointly shifts the gather point and photon. In practice, our shift mapping is applied in three steps: first shifting the camera subpath, then shifting the gathered photon location and finally shifting the light subpath. By shifting the camera subpath without considering the light subpath, we can estimate the Jacobian in Equation 12 by separating the gather point and the photon as:

$$\left|\partial(\mathbf{b}'\mathbf{d}')/\partial(\mathbf{b}\mathbf{d})\right| = \left|\partial\mathbf{b}'/\partial\mathbf{b}\right|\left|\partial\mathbf{d}'/\partial\mathbf{d}\right|. \tag{13}$$

In the following sections, we will detail the shift mapping for volumetric density estimation techniques, starting with volumetric photon mapping (VPM) (Section 4), and then extending to higher dimensional photon representations including beam radiance estimate (BRE), photon beam (Beam), and photon plane (Plane) (Section 5). For each technique, we will also show how to compute the Jacobian for the gradient-domain rendering integral. Along with the discussion of shift mapping, we will also introduce *mixed shift*, a new shift mapping that combines traditional shift mapping with photon reuse, which leads to more efficient shift mapping particularly when the density estimation kernels has great overlaps.

4 GRADIENT-DOMAIN VOLUME PHOTON MAPPING

We first extend *volumetric photon mapping* (VPM) to the gradient-domain, before deriving our shift mapping: we extend surface-based gradient-domain photon mapping [Hua et al. 2017] to volumes by considering eye-segment integration due to transmittance. We estimate this distance integral with a density estimator D that performs distance sampling to determine the last vertice $\mathbf{d}$ and $\mathbf{b}$ on the camera and light subpaths. In VPM, $\mathbf{d}$ and $\mathbf{b}$ are sampled independently.

In the primal domain, we can estimate this integral with ray marching or Monte Carlo integration. We choose the latter to more closely match the path tracing formalism. Here, we first importance sample the k^{th} camera edge $\mathbf{z}_{k-1}\mathbf{z}_k$ based on its relative throughput, before distance sampling t_c proportional to edge transmittance as

$$p(t_c) = \frac{\sigma_t}{1 - \exp(-\sigma_t s)} \exp(-\sigma_t t_c) p(k), \tag{14}$$

where s is the distance between $\mathbf{z}_{k-1}$ and $\mathbf{z}_k$. Once we sample t_c, we have $\mathbf{d} = \mathbf{c} + t_c \omega_c$ where $\mathbf{c} = z_{k-1}$ and $\omega_c = \mathbf{z}_{k-1}\mathbf{z}_k$. We perform density estimation at vertex $\mathbf{d}$ to find photons $\mathbf{b}$ that overlap the kernel support centered at $\mathbf{d}$, forming complete paths. We shift this (completed) path and estimate its image-space gradients according to the procedure we detail in Sections 4.1 and 4.2, and in Figure 2.

4.1 Shifting the Camera Subpath

In surface density estimation, we assume that a base camera path ends on the first diffuse surface encountered during tracing. In volume density estimation, however, this gather point is not necessarily the base camera path's last vertex: instead, it can result from distance sampling. The base camera path has the form $\mathbf{z}_0 \ldots \mathbf{c}\mathbf{d}$, where $\mathbf{z}_0 \ldots \mathbf{c}$ is the prefix and $\mathbf{d}$ the gather point. Note that all base camera
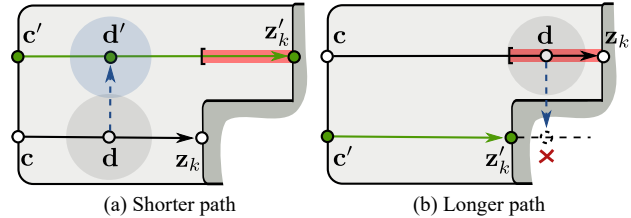

(a) Shorter path   (b) Longer path

Fig. 3. Shifting the gather point on the base camera subpath. We copy the transmittance distance from the base, to the offset camera subpath. If the base camera subpath edge is shorter, we attribute zero contribution to the omitted distance integration region; if it is longer, the shift is unsuccessful.

path prefix vertices are surface interactions, but the gather point could lie either on a surface or in a volume.

Therefore, our camera path shift mapping requires two steps: first, we shift the base camera path prefix and then the base gather point. For the first shift step, since the prefix only comprises surface interactions, we can apply path tracing and half-vector copy shift mappings proposed by Kettunen et al. [2015] and Hua et al. [2017] to generate the offset prefix $\mathbf{z}_1' \ldots \mathbf{c}'$. After copying the half-vectors, we need to generate an offset gather location by determining a corresponding distance t_c' on the offset path based on the sampled distance t_c on the base camera path. With this strategy, the Jacobian due to the camera subpath shift is:

$$\left|\frac{\partial(\mathbf{z}_0' \ldots \mathbf{c}'\mathbf{d}')}{\partial(\mathbf{z}_0 \ldots \mathbf{c}\mathbf{d})}\right| = \left|\frac{\partial(\omega_0' \ldots \omega_c')}{\partial(\omega_0 \ldots \omega_c)}\right|\left|\frac{\partial t_c'}{\partial t_c}\right| \tag{15}$$

$$\cdot \frac{G(\mathbf{z}_0, \mathbf{z}_1)}{G(\mathbf{z}_0', \mathbf{z}_1')} \frac{G(\mathbf{z}_1, \mathbf{z}_2)}{G(\mathbf{z}_1', \mathbf{z}_2')} \cdots \frac{G(\mathbf{c}, \mathbf{d})}{G(\mathbf{c}', \mathbf{d}')}$$

where we use the relation between the volume measure at a point $\mathbf{z}$ and the subtended solid angle measure $d\mathbf{z}\, G(\tilde{\mathbf{z}}, \mathbf{z}) = d\omega\, dt$ and $\mathbf{z} = \tilde{\mathbf{z}} + \omega t$ to compute the Jacobian in the volume measure from the half-vector copy Jacobian [Kettunen et al. 2015]. Note that, for the Jacobian due to changes in distance, we need only consider the distance on the last camera subpath edge (from t_c to t_c') since the preceding vertices are generated without any distance sampling. Only a ray-surface intersection is required to trace these vertices.

To determine the offset gather location, we copy the distance $t_c' \equiv t_c$, but only after ensuring that the shift distance t_c' generates a valid gather point (see Figure 3). Here, the shift Jacobian is simply 1. With this offset distance, we now define our offset gather point $\mathbf{d}'$.

Discussion. We experimented with two additional shift maps, including linearly scaling $t_c' = k t_c$ (where k is the ratio of the offset and the base camera edge length) and reusing the random number that generates t_c on the base camera path. We found that our simpler copying of t_c works best, in general. This behavior results from first recognizing that any explicit shift guarantee will undoubtedly displace gather points more, and so result in a less coherent shift. This is especially true in the presence of sharp lighting features, like caustics and visibility changes. Secondly, even though the simpler copy strategy can lead to under-sampled regions in the integration domain (red line region in Figure 3), these regions are very small in practice and do not cause noticeable convergence errors in practice. We can address this (admittedly minor) inconvenience by marking the offset path as irreversible.

(a) Simple shift

(b) Null shift

(c) Mixed shift

- ● Base gather point
- ● Offset gather point
- ○ Base photon
- ● Offset photon
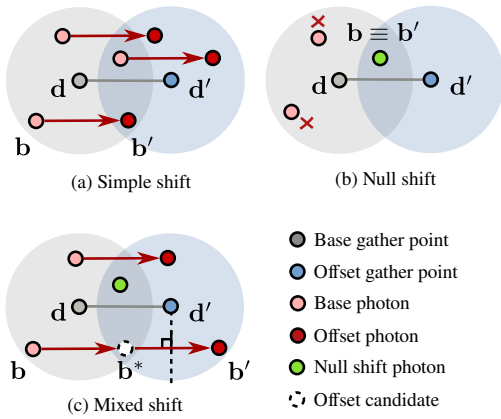- ● Null shift photon
- ⊙ Offset candidate

Fig. 4. Shift mapping to generate the photon. Mixed shifts (c) reuse photons in the overlapped region (darker blue) while also guaranteeing reversibility by reflecting the photons in the non-overlapped region in the kernel support across the plane perpendicular to $\mathbf{b}^*\mathbf{b}'$ (dashed line). In (b), photons marked with a cross cannot contribute to the offset gather point.

4.2 Shifting the Photon

After shifting the base camera subpath with the base gather point, we construct a kernel at the offset gather point to prepare for a complete density-estimated path. We always apply the same kernel at both the base and offset gather points. Hua et al. [2017] apply a shift mapping guided by the distance between the offset and the base gather points, which we refer to as a *simple shift*, here. The simple shift is general, but could be expensive in volumetric rendering for two reasons: first, the kernel size in volumetric rendering is generally larger than in surface rendering, and so more photons will overlap a kernel; second, many density estimations are performed along a camera path, resulting in many more shift mappings per rendering pass.

To reduce the cost of shift mapping in gradient-domain volumetric light transport, we exploit this kernel overlap to design a new shift mapping we call the *mixed shift*. Here, we discuss simple shift and introduce an additional intermediate concept called the *null shift*, before combining these concepts together to form our mixed shift.

Simple shift. Here, the offset photon position is computed as $\mathbf{b}' = \mathbf{b} + (\mathbf{d}' - \mathbf{d})$ (see Figure 4a). This simple transformation is reversible and guarantees that every region in the kernel can be explored. This always shifts a base photon, regardless of where it falls in the kernel, which is suboptimal since photons that lie in the overlap of the base and offset kernel supports need not be shifted.

Null shift. One can also simply treat each base photon as a offset photon (Figure 4b): this shift mapping is always successful and effectively free for photons that overlap both kernel supports. On the other hand, photons outside the overlap contribute zero energy to the offset gather point and yield an ineffective shift map.

Mixed shift. Motivated by the pros and cons of the simple and null shifts, we propose a more efficient method (Figure 4c): photons in the kernel overlap region are not shifted (null shift); (2) otherwise, the photons are shifted using a simple shift. To ensure that the shift mapping is reversible, we verify whether photons shifted using the
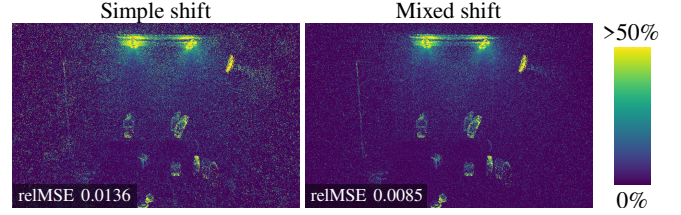
Fig. 5. Failure rate of simple and mixed shift in the SPOTLIGHT scene dominated by specular transport: mixed shift decreases the shift failure rate and lower the relative MSE by more than 30% for equal render time (5 mins) using gradient-domain volumetric photon mapping.

simple shift lie in the overlap region. If so, they are re-shifted outside the overlap. As such, the final offset photon position is

$$\mathbf{b}' = \mathbf{b}^* + 2\mathbf{u}^\top(\mathbf{d}' - \mathbf{b}^*)\mathbf{u}, \tag{16}$$

where $\mathbf{b}^*$ is the initial offset position after one application of a simple shift, and $\mathbf{u} = (\mathbf{d}' - \mathbf{d})/\|\mathbf{d}' - \mathbf{d}\|$ the unit vector pointing towards the offset gather point. The dot product term is the projected distance from the initial photon offset position to the offset gather point onto unit vector $\mathbf{u}$. This term makes sure that the final photon offset position $\mathbf{b}'$ is out of the overlapping region. Figure 5 illustrates the benefits of mixed shift regarding the percentage of successful shifts in the SPOTLIGHT scene with complex shift configurations.

4.3 Reconnection

After computing the offset photon, we reconnect it to the base path. To increase coherence, it is preferable to connect the photon to the base light path as early as possible, all while avoiding zero contribution scenarios. To do so, after classifying the preceding light path vertex of the base photon as diffuse or specular, we apply specialized strategies: for diffuse, we directly connect the offset photon to the preceding vertex; otherwise, we rely on manifold exploration for reconnection [Hua et al. 2017; Lehtinen et al. 2013].

5 GRADIENT-DOMAIN VOLUME DENSITY ESTIMATION WITH PHOTON BEAMS AND PLANES

Equipped with insights from gradient-domain volumetric photon mapping, we will detail extensions of higher-order primitives to gradient-domain density estimation, including *beam radiance estimate* (BRE) [Jarosz et al. 2008], *photon beams* (Beams) [Jarosz et al. 2011b] and *photon planes* (Planes) [Bitterli and Jarosz 2017].

5.1 Gradient-domain Beam Radiance Estimate

The *beam radiance estimate* (BRE) often outperforms standard volumetric photon mapping in the primal domain. Having generalized the *point-to-point* volumetric photon mapping estimator to the gradient domain (i.e., G-VPM), and motivated by the relationship between the primal-domain VPM and BRE algorithms, we will extend the *beam-to-point* BRE estimator to the gradient domain (i.e., G-BRE).

Point-Beam 2D. The BRE gathers all photons along a camera path segment using a beam query, which can be interpreted as a splatting technique where each photon-aligned kernel's energy is projected onto the screen. This formulation is more efficient than VPM and enables the use of 2D kernels. To enable gradient-domain BRE with a 2D kernel (Figure 6c), we need to ensure that $\mathbf{d}'\mathbf{b}'$ are perpendicular
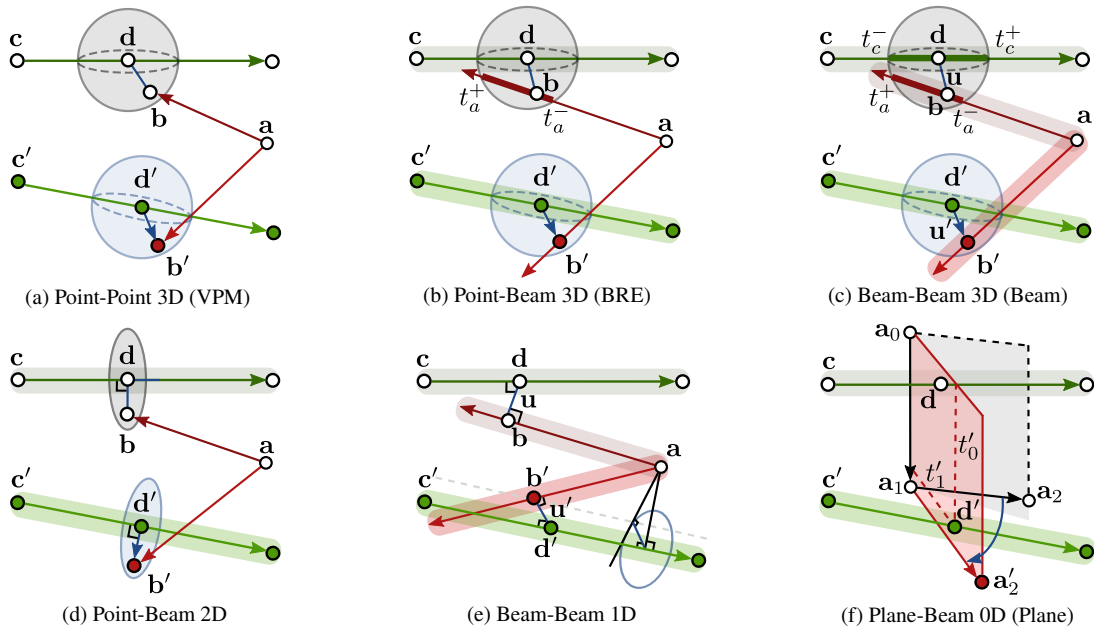
Fig. 6. Shift mapping for volumetric density estimation techniques with various kernel dimensions. In general, shift mapping with lower kernel dimensions is more restricted as it requires preserving more geometric constraints to guarantee reversibility. Technique naming convention: data (Point, Beam, Plane) - query (Point, Beam) and kernel dimension (3D, 2D, 1D, 0D). (For conceptual illustration, beam size and kernel radii are not drawn to have the same scale.)

to the offset camera beam. We do so by expressing the vector $\mathbf{db}$ in the base camera beam coordinate frame, and by copying its local coordinates to the offset camera beam. The associated shift mapping only applies rigid transformations, and so its Jacobian is 1.

However, since the overlap of two 2D kernels is either null nor degenerate, the application of a *mixed shift* becomes more challenging. To overcome this issue, we rely on a 3D kernel variant of the BRE, which we call Point-Beam 3D.

Point-Beam 3D. We rely on an extended volumetric path integral formulation [Hachisuka et al. 2017] to apply a 3D kernel in BRE and devise our new shift mapping (Figure 6b). This shift mapping procedure is related to the mixed shift for G-VPM (Section 4.2), with two notable differences. First, we position a kernel at a photon location instead of a gather point, for G-BRE. This, however, is only a conceptual difference and does not impose any algorithmic changes. To evaluate the contribution, in the primal domain, we sample the base gather point in $[t_c^-, t_c^+]$ in the kernel, and compute the contribution of the photon to this gather point. Shift mapping in the gradient domain can then be applied as in Section 4.2. Second, and more importantly, we have to ensure that the camera beam segment and its length remain constant during a shift, i.e., $[t_c^{-\prime}, t_c^{+\prime}] = [t_c^-, t_c^+]$. We achieve this by keeping the distance between the photon and the gather point constant $|\mathbf{d'b'}| = |\mathbf{db}|$.

5.2 Gradient-domain Volume Density Estimation with Photon Beams

After lifting camera subpath gather points to beams for G-BRE, we continue with a generalization of photon points in gradient-domain volumetric density estimation to higher dimensional photon

primitives: first, from photon points to photon beams. As with G-BRE, we assume that a camera beam is always used for gathering. The camera beam could intersect a photon beam constructed from a light subpath edge.

In general, kernel evaluation with beams involves sampling vertex pairs $(\mathbf{d}, \mathbf{b})$ [Bitterli and Jarosz 2017; Hachisuka et al. 2017]. These vertices can be treated as a gather point on the camera beam, and a photon along the photon beam. In the gradient domain, in addition to tracking the probability density of the sampled vertex pair $(\mathbf{d}, \mathbf{b})$, we need to track the Jacobian associated to this sampling, as in Equation 12. For efficiency, our sampling takes the geometric intersection of the camera beam and the photon beam into account. We provide details for density estimators with photon beams using 1D and 3D kernels, below.

Beam-Beam 1D. Here we opt to preserve the distance between the photon and the gather point (Figure 6e): we first project the beam origin $\mathbf{a}$ onto the offset camera ray starting from $\mathbf{c}$, and fit a disk S of radius r at the projected location so that the disk is perpendicular to the offset camera ray. We then compute a tangent vector from the beam origin $\mathbf{a}$ to the disk, which yields the offset distance vector $\mathbf{u'}$. We slide $\mathbf{u'}$ along the offset camera ray, until it reaches the offset gather point $\mathbf{d'}$, to determine the offset photon at $\mathbf{b'}$. Note that there are two possible tangent vectors during this shift. The offset photon from both cases will have the same distance to the beam origin, due to symmetry, and so the Jacobian is equal for both cases. We estimate the Jacobian of this shift mapping numerically with finite differences, and find that a value of one produces sufficiently accurate reconstructions.

In theory, the shift mapping we present is sound, preserving geometric constraints to guarantee reversibility. To be concrete, after

the shift, the distance between the photon and the gather point remains unchanged, and the distance vector is perpendicular to the camera and photon beam. To generalize this shift mapping to specular and glossy path reconnections, manifold exploration [Jakob and Marschner 2012] needs to be extended to preserve such geometric constraints.

In practice, we approximate this shift mapping with a simpler strategy. We first find the candidate photon $\mathbf{b}_* = \mathbf{b} - \mathbf{d} + \mathbf{d}_*$ to determine the offset photon beam, where the candidate gather point $\mathbf{d}_*$ is computed by copying the same distance from the primal camera subpath. Then, we perform an additional beam-beam intersection here to find the actual photon $\mathbf{b}'$ and the actual gather point $\mathbf{d}'$. The Jacobian of the photon shift here is

$$\left|\frac{\partial \mathbf{d}'}{\partial \mathbf{d}}\right|\left|\frac{\partial \mathbf{b}'}{\partial \mathbf{b}}\right| = \left|\frac{\partial \mathbf{d}'}{\partial \mathbf{d}_*}\right|\left|\frac{\partial \mathbf{d}_*}{\partial \mathbf{d}}\right|\left|\frac{\partial \mathbf{b}'}{\partial \mathbf{b}_*}\right|\left|\frac{\partial \mathbf{b}_*}{\partial \mathbf{b}}\right| = \left|\frac{\partial \mathbf{d}'}{\partial \mathbf{d}_*}\right|\left|\frac{\partial \mathbf{b}'}{\partial \mathbf{b}_*}\right| \qquad (17)$$

which models the shift from the candidate to the actual intersection. In practice, this value is very close to 1. Empirically, we find that the approximate shift is effective and integration converges without issues. Moreover, there is negligible impact on convergence between the accurate and approximate shift mappings, across our test scenes.

Beam-Beam 3D. For Beam-Beam 3D, we find the photon beam segment that intersects the kernel support, and then sample a photon along the segment (Figure 6c). Given the sampled photon, the gather point can be sampled as in the Point-Beam 3D (BRE) case. We then apply the same kernel shift as G-VPM (Section 4) to generate the offset gather point $\mathbf{d}'$ and offset photon $\mathbf{b}'$. Diffuse reconnection or manifold exploration can be used to connect the offset photon at $\mathbf{b}'$ to the base light subpath, as usual. Note that, in this case, the geometry of the offset photon beam (its start and end points) is undetermined until successful reconnection. While the 3D kernel sampling (to estimate primal-domain contribution) is slightly different from Point-Beam 3D, the shift mapping and its Jacobian remain unchanged in the gradient domain.

5.3 Gradient-domain Volume Density Estimation with Photon Planes

Finally, we show that gradient-domain rendering can be extended beyond point and beam primitives. Namely, we extend shift mapping to photon planes with a zero dimensional kernels (Figure 6f). Here, density estimation involves a camera beam and a photon plane, and we shift this configuration as follows.

As before, we first fix the distance on the camera subpath, setting $t'_c = t_c$ to compute the offset gather point $\mathbf{d}'$. We then proceed to establish the offset photon. Since we use a 0D kernel, the offset photon is expected to be at the same location as the offset gather point, and so $\mathbf{b}' = \mathbf{d}'$. The offset photon plane is then constructed by preserving $\mathbf{a}_0$ and $\mathbf{a}_1$, rotating the base plane by moving $\mathbf{a}_2$ to $\mathbf{a}'_2$, where the offset plane contains $\mathbf{d}'$. We mark this shift as invalid if the offset camera beam does not intersect the offset photon plane, or when t'_0 is negative because we opt not to change $\mathbf{a}_0$ and $\mathbf{a}_1$ and only do a rotation to generate offset plane.

The Jacobian determinant of this shift is

$$\left|\frac{\partial \mathbf{b}'}{\partial \mathbf{b}}\right| = \left|\frac{\partial(t'_0 t'_1)}{\partial(t_0 t_1)}\right|\left|\frac{\partial \mathbf{b}'}{\partial(t'_0 t'_1)}\right|\left|\frac{\partial(t_0 t_1)}{\partial \mathbf{b}}\right| \qquad (18)$$

$$= \left|\frac{\partial(t'_0 t'_1)}{\partial(t_0 t_1)}\right|\frac{\omega_0 \cdot (\omega'_1 \times \omega'_c)}{\omega_0 \cdot (\omega_1 \times \omega_c)} = \frac{t'_0}{t_0}\frac{t'_1}{t_1}\frac{\omega_0 \cdot (\omega'_1 \times \omega'_c)}{\omega_0 \cdot (\omega_1 \times \omega_c)},$$

where $\mathbf{d}' = \mathbf{b}' = \mathbf{a}_0 + \omega'_0 t'_0 + \omega'_1 t'_1$ where $\omega'_0 = \mathbf{a}_0 \mathbf{a}_1$ and $\omega'_1 = \mathbf{a}_1 \mathbf{a}'_2$. We operate in a local coordinate frame parameterized by (t_0, t_1), compute the Jacobian, and transform back to the original frame. The local parameterization yields a Jacobian determinant that is similar to photon plane variable coupling or ray-triangle intersection [Möller and Trumbore 1997]. In this local space, we need only consider the scaling of t_0 and t_1 during a shift, hence the formula above. Contrary to our previous shift mapping, we do not need to reconnect the light path (Section 4.3) as neither $\mathbf{a}_0$ nor the first edge change during the shift mapping.

6 EXPERIMENTS

We employ five test scenes: Kitchen, Staircase, Spotlight, Bathroom, and Laser. The first four scenes were rendered with surface-volume and volume-volume light transport, to evaluate all the gradient techniques except gradient-domain photon planes. We use the Laser scene to evaluate gradient-domain photon planes, since photon planes [Bitterli and Jarosz 2017] are best suited to volume-volume light transport and multiple scattering. Before performing density estimation, we trace camera subpaths until they hit diffuse surfaces. We use the vertex classification convention of previous work [Hua et al. 2017; Kettunen et al. 2015], classifying a surface vertex as specular if its roughness is < 0.05. We use a relative mean square error metric *relMSE*$= (1/n) \sum (R - I)^2/(R^2 + 0.001)$, where R is the reference and I our rendered output. All gradient-domain techniques use Poisson reconstruction with L1-norm (Equation 9) to produce the final images.

We generate all our results on a dual Intel Westmere EP Xeon X5650 workstation clocked at 2.66GHz with 12 cores and 24 GB of memory. The reference images are rendered using BRE with a 3D kernels, except the Laser scene which uses bidirectional path tracing. We provide an open source implementation of our methods.

6.1 Implementation Details

We implement our work in Mitsuba [Jakob 2010], evaluating our gradient-domain volumetric photon mapping (G-VPM), gradient-domain beam radiance estimate (G-BRE), gradient-domain photon beams (G-Beam) and gradient-domain photon planes (G-Plane) against their four primal-domain variants.

All our techniques support progressive photon density estimation [Knaus and Zwicker 2011], we use the standard reduction ratio of $\alpha = 0.7$ and deposit 1 million photons or 20×10^3 photons beams and planes per iteration.

We clamp the number of manifold exploration [Jakob and Marschner 2012] iterations to 5 when shifting a specular or glossy light path. To achieve even faster manifold exploration and increase the number of successful shifts, we relax the acceptance condition of these walks by accepting a proposal path so long as its last vertex (towards the camera) falls inside the offset photon's kernel. This relaxation biases
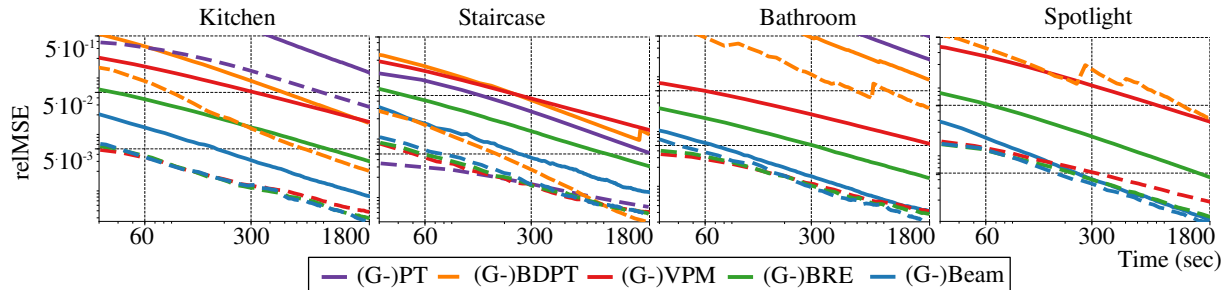
Fig. 7. Convergence plots for our test scenes using a relative MSE metric. Solid and dash lines represent primal- and gradient-domain methods. G-VPM and G-BRE outperform their primal counterparts by a large margin. G-Beam outperforms Beam marginally for scenes dominated with specular light transport.

manifold exploration but, due to progressive kernel size reduction, the resulting estimator is still consistent.

We set the λ parameter for Poisson reconstruction to 0.2, the kernel radius scale to 0.10% of the scene bounding box diagonal (for every method except VPM, where we use a slightly larger scale of 0.15% to reduce variance.) We use a uniform kernel for all our results. For (G-)VPM, we use 40 stratified samples per pixel for distance sampling. For beam and plane density estimation, we use a BVH to accelerate beam-beam and beam-plane intersection tests.

With the exception of photon planes, we use 3D kernels for all volumetric photon density estimators. Note that, while Beam-Beam 1D estimators may appear to be more efficient (as the kernel is of lower dimension), we found empirically that there is no clear benefit when compared to Beam-Beam 3D: in fact, it is not fair to compare Beam-Beam 1D and Beam-Beam 3D in a progressive density estimation setting, since kernel reduction rates differ across kernel dimensions.

Gradient Estimation with Photon Subsets. We propose to use only subsets of all available photons during gradient estimation to reduce rendering time. Our observation is that, in density estimation, neighboring pixels gather from mostly the same set of photons, resulting in similar pixel values and small gradients. This is especially true for volumetric rendering with camera beam queries. We therefore propose to split the photons into two equal subsets that are gathered by each pixel in an interleaved manner. This balance between introducing artifacts in gradient estimation and improving computation time when gathering photons. We employ this optimization for all our rendering techniques that rely on camera beam queries, including G-BRE, G-Beam and G-Plane. To further reduce the impact of noise from the primal domain, we reuse offset paths to estimate the primal-domain contribution [Lehtinen et al. 2013].

Multiple Importance Sampling. As with previous gradient-domain methods, we apply multiple importance sampling to reduce the variance of Jacobian determinant estimates. As light and camera subpaths are generated independently, we can compute the weight using a balance-like heuristic, as

$$w(\mathbf{x}) = \frac{p(\mathbf{y})p(\mathbf{z})}{p(\mathbf{y})p(\mathbf{z}) + p(T(\mathbf{y}))p(T(\mathbf{z}))|T'(\mathbf{y})||T'(\mathbf{z})|} , \qquad (19)$$

where all probabilities are expressed in the area measure and T' is the Jacobian of the shift mapping. For mixed shifts, when a base light path is reused, we simply have $T(\mathbf{y}) = \mathbf{y}$ and the weight only treats the base and offset camera paths.

Gradient-domain Volumetric Path Tracing. We additionally implement novel extensions of gradient-domain path tracing (G-PT) [Kettunen et al. 2015] and gradient-domain bidirectional path tracing (G-BDPT) [Manzi et al. 2015] to volumetric light transport. In G-PT, we use a half-vector copy shift mapping and also copy the distance sampled on the base path to the offset path (similarly to Section 4.1). We use a balance heuristic for MIS in G-PT instead of a power heuristic [Kettunen et al. 2015] since, in a volume, the path probability from distance sampling is very small and may suffer from floating point precision issues. Our G-BDPT extension to volumes is much more straightforward, incorporating the transmittance evaluation into the existing implementation [Manzi et al. 2015].

6.2 Experimental Results

Convergence. We select scenes with various types of light transport to evaluate our techniques. Particularly, for scenes with mostly diffuse light transport, we choose STAIRCASE and KITCHEN, and for scenes with moderate and dominant specular-diffuse-specular light transport, we choose BATHROOM and SPOTLIGHT. Figure 7 shows the convergence plots of the different algorithms, for equal rendering times. The result images are shown in Figure 8. Figure 12 illustrates the proportions of different shift mapping used in each gradient-domain technique.

Gradient-domain techniques consistently outperform their primal-domain counterpart, with the widest margin in scenes dominated by diffuse light transport (STAIRCASE and KITCHEN), since the shift mapping of diffuse light subpaths only requires direct reconnection (which is both cheap to perform and more likely to succeed.) This applies to both gradient-domain path-based and photon-based approaches. Figure 8 (first row) illustrates the rendering of the STAIRCASE scene. G-PT and G-BDPT also perform as well as our gradient-domain photon density estimation techniques here since their primal-domain counterpart can effectively sample the path space.

The BATHROOM and SPOTLIGHT scenes are dominated by specular-diffuse-specular transport. Here, we only show the results of volume photon density techniques as we found (G-)PT and (G-)BDPT yield very high variance estimators. Please refer to our supplemental document for comprehensive results, including relative error maps and interactive comparisons.

In the BATHROOM scene (Figure 8, second row), the main light source has a high frequency occluder that reduces the efficiency
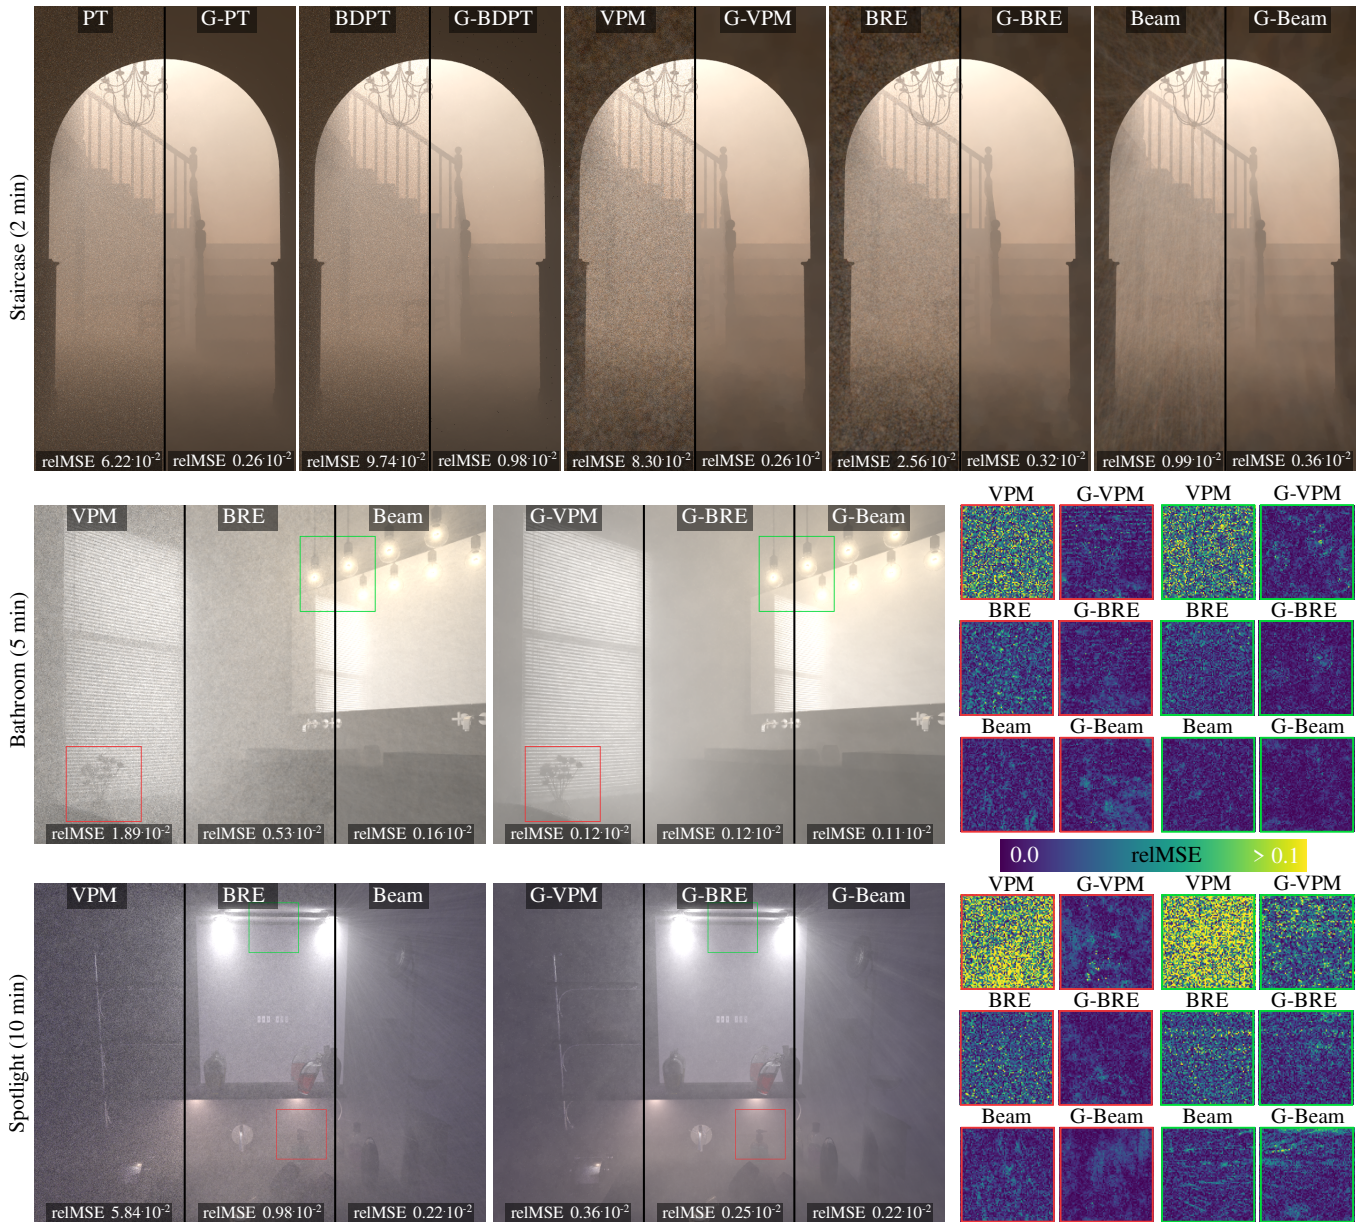
Fig. 8. Equal-time rendering of primal- and gradient-domain methods for the STAIRCASE, BATHROOM and SPOTLIGHT scenes. For scenes dominated by specular-diffuse-specular light transport, such as BATHROOM and SPOTLIGHT, gradient-domain volumetric photon density estimation (G-VPM, G-BRE, G-Beam) significantly outperforms gradient-domain path-based techniques (G-PT and G-BDPT). For full rendering results, please refer to our supplemental document.

of direct reconnection. Here, our gradient-domain techniques still outperform all primal-domain techniques.

SPOTLIGHT (Figure 8, last row) is our most challenging scene, as most of the illumination comes from complex specular light transport, since luminaire lenses are modelled realistically. Our techniques are very robust, presenting at least a 2× reduction in relMSE for G-VPM and G-BRE, compared to VPM and BRE. G-Beam also improves upon Beam, albeit only marginally. This small gain

can be attributed to the fact that the primal-domain Beam technique can already efficiently treat specular light transport.

Anisotropic Phase Function. Figure 9 shows relative gain between primal- and gradient-domain photon plane in LASER scene along with other volumetric density estimation techniques. This scene is challenging because it has a point light source with a single fixed lighting direction for laser. We use direct reconnection for shift mapping here because there does not exist a diffuse vertex on the light source for manifold exploration to work. Gradient-domain
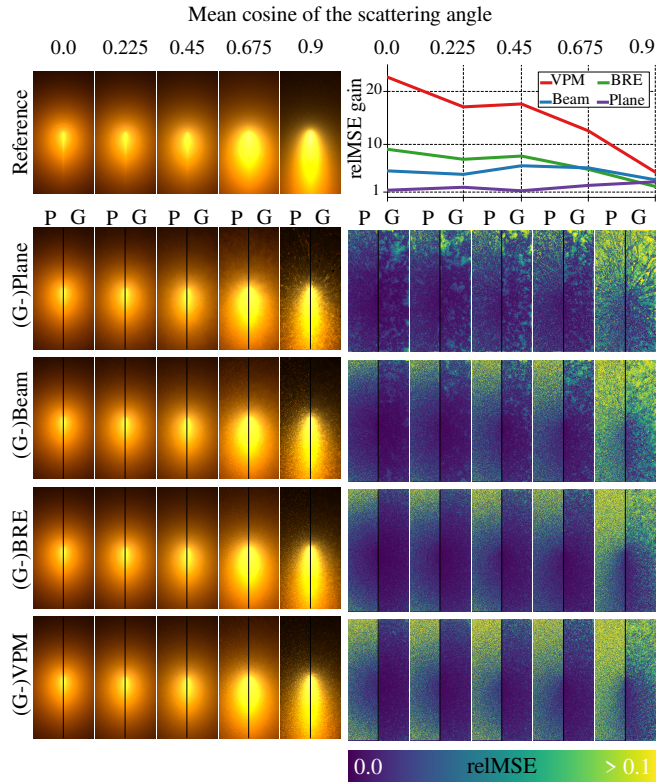
Fig. 9. The LASER scene rendered with primal- and gradient-domain volumetric density estimation. For isotropic and low anisotropic phase functions, gradient-domain methods outperform primal-domain variants, including photon planes. Top left: the relMSE ratio between primal- and gradient-domain volumetric density estimation methods (10 minutes). Gradient-domain methods are less effective with increasing phase function anisotropy.
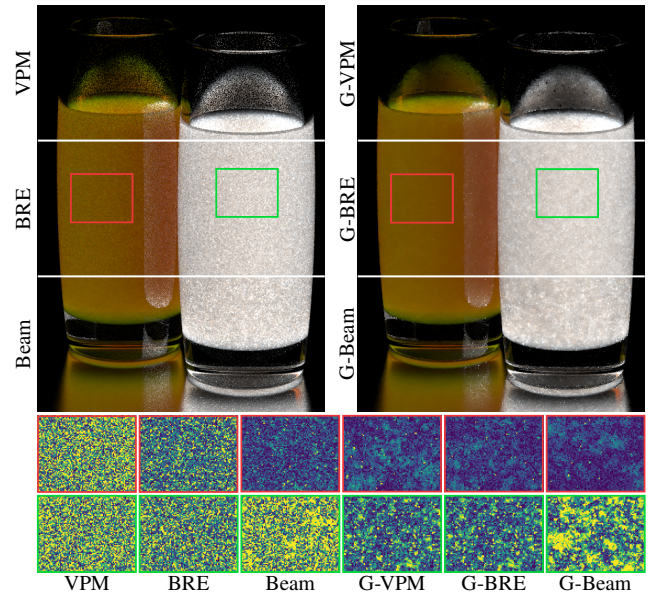

Fig. 10. Equal-time renderings (five minutes) of the GLASS scene. Our gradient-domain methods improve upon primal-domain counterparts for both the milk and orange juice media. The rendering performance for the glass of milk drops due to noisier primal-domain estimates in dense media.
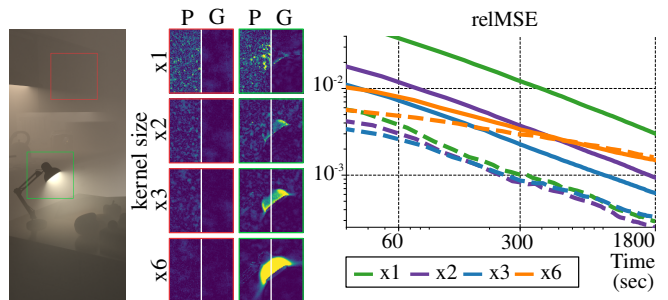

Fig. 11. G-BRE with various kernel sizes: as we increase kernel size (and, so, bias), the performance gap between gradient- and primal-domain narrows (x1 uses the same kernel size as in Figure 1.) For extremely large initial radii (x6), both primal- and gradient-domain rendering have large bias but the gradient domain method still outperforms, especially in early iterations.

techniques still work effectively for this scene, but the gain reduces when the data representation becomes higher dimensional such as beams and planes as the primal-domain method is already efficient.

In this experiment, we also investigate the influence of anisotropic phase function in gradient-domain rendering. We excluded single scattering and varies the phase function from 0.0 to 0.9. We found that photon plane is more effective data representation for primal-domain density estimation in scenes with near-isotropic phase functions. However, when the phase function is strongly directional, all primal-domain techniques result in noisy estimations. Gradient-domain techniques, by contrast, can still perform better and result in less noisy estimations.

Dense Participating Media. Figure 10 shows a scene with a glass of milk and a glass of orange juice. Manifold exploration can handle the dielectric interface of the glasses and successfully generates offset paths for gradient-domain methods. In this experiment, we set the radius to be 0.5% of the bounding box diagonal of the participating medium for all the techniques. Similar to previous experiments, our gradient-domain techniques consistently outperform their primal-domain counterparts. We also observe that, with dense media, gradient-domain methods inherit the noise that primal-domain rendering results also exhibit in these scenarios.

Bias vs. Variance. In volumetric photon density estimation, there is trade-off between bias and variance, which is controlled by the kernel size. In general, using smaller kernel size results in higher variance and lower bias in the primal images, which favors the de-noising capability of gradient-domain techniques. Larger kernel size results in larger bias, in which cases gradient-domain techniques become less effective. Figure 11 demonstrates the trade-off with various kernel size with (G-)BRE. As can be seen from the convergence plot in this figure, it could be just safe to set a small kernel size for gradient-domain volumetric photon density estimation techniques.

Performance Gain. To justify the performance of our techniques, Figure 12 visualizes statistics of shift mapping in our test scenes, and Figure 13 shows the relative gain between two classes of gradient-domain techniques, one with a naïve shift mapping that only uses simple shift, and the other with our approach with optimizations

(including mixed shift and kernel relaxation in manifold exploration). As can be seen in the Spotlight scene in Figure 13, the naïve approach exhibits noisy gradient estimation because of the overhead caused by expensive manifold exploration due to specular light transport. From the statistics in Figure 12, mixed shift can account for about 30% of the total number shifts, and so with mixed shift and all optimizations applied, we show that a significant amount of manifold walks can be skipped, resulting in overall gains in the convergence plot in Figure 13. Visual comparison of the gradients estimated by the naïve and our approach is also shown in this figure. In addition, we note that these gains are more significant for G-VPM as this methods uses a bigger radius.

6.3 Limitations

Despite generating smoother results, each iteration of a gradient-domain iteration is slower than its primal-domain variants' iteration, due to the overhead of shift mapping and gradient computation. In our experiments, G-VPM is about 30% slower per iteration, G-BRE twice as slow, and G-Beam up to four times slower. We use equal-time comparisons for all error comparisons, however. Please refer to our supplemental material for more detailed timing statistics. The additional transformation in mixed shift can cause large light subpath perturbations, compared to a simple shift. This increases variance in the gradient estimation for mixed shifts however, in practice, we compensate for this by reducing the cost when kernels overlap. Mixed shift also generally results in a higher number of successful shifts, which contributes to variance reduction. While kernel overlap reduces during progressive density estimation, this poses no issues in practice as the simple shift's cost also reduces.

G-VPM, G-BRE, G-Beam, and G-Plane inherit the limitations of their primal-domain counterparts: for example, radiance estimation with VPM and BRE could be too noisy to generate a good throughput image, resulting in noisy gradient and Poisson reconstruction. This could happen in scenes where photon visibility is complex to sample. These issues could be addressed by incorporating, for example, a Markov chain Monte Carlo sampler.

Gradient-domain methods are much less sensitive to the initial setting of the kernel radii compared to primal-domain techniques,

Fig. 12. Shift statistics after 30 minutes of rendering. For scenes dominated by specular and glossy transport (e.g., as in Bathroom and Spotlight), shift mapping is more costly due to additional manifold exploration steps. Mixed shift reduces the need to perform manifold exploration by reusing photons for shift mapping, thus retaining the efficiency of gradient-domain methods.
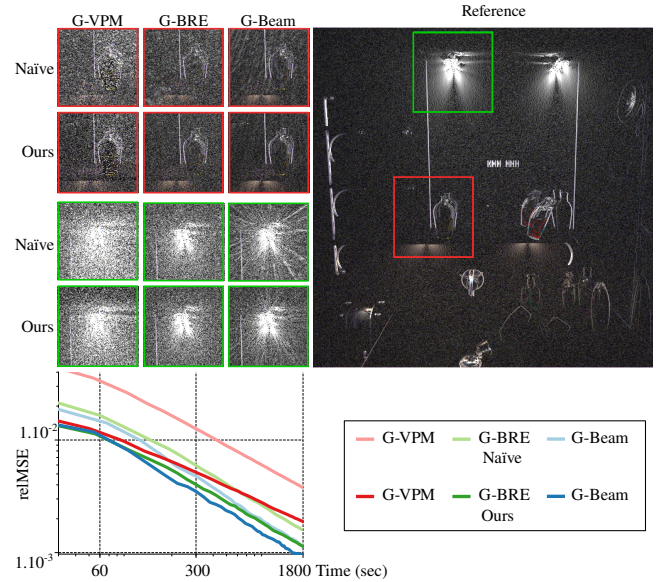
Fig. 13. Comparing gradient estimation in the Spotlight scene with a naïve approach with simple shifts and our approach. Except at object boundaries, we expect smooth near-zero gradients. Our approach generates remarkable low-noise gradient estimates, resulting in faster convergence.

however rendering error is dominated by bias and not variance when employing extremely large initial radii (see x6 in Figure 11). Gradient-domain reconstruction is not able to reduce this bias, and so designing new reconstruction schemes that take this bias into account can be an interesting avenue of future work.

7 CONCLUSION

We proposed the first gradient-domain formulation of light transport in participating media, resulting in four new rendering algorithms: gradient-domain volumetric photon mapping (G-VPM), gradient-domain beam radiance estimate (G-BRE), gradient-domain photon beams (G-Beam) and gradient-domain photon planes (G-Plane). To do so, we introduced novel shift mappings tailored to each rendering algorithm that apply to volumetric transport and photon density estimation formalisms. Our experiments validate that G-VPM and G-BRE *both* consistently outperform their primal-domain GPM and BRE variants. G-Beam and G-Plane also significantly outperform the primal-domain Beam and Plane algorithms in diffuse scenes.

There are two key theoretical insights we gain from our investigation of the various gradient-domain volumetric photon density estimation (and path-based) techniques. First, while it is generally accepted that gradient-domain methods for surface rendering improve upon their primal-domain counterparts (and, there, usually by a significant margin), we show that gain in the volumetric setting scales differently: while gradient-domain methods still outperform their primal-domain counterparts, the benefits of moving to the gradient-domain depends on the *dimensionality* of the density estimation primitives being used. The efficiency gap between gradient- and primal-domain volumetric density estimators narrows for higher-order primitives, e.g., when moving from using beams

to using planes. This a somewhat counter-intuitive result: in images rendered with photon beams and planes, the noise distribution in the throughput and gradient images has lower frequency. This, however, can lead to non-optimal Poisson reconstruction when using gradients. Designing a method that takes care to consider, e.g., the frequency in the throughput and gradient images may be an interesting direction of future work.

Second, we showed how to exploit the unique aspect of biased rendering methods in the gradient domain. As with primal-domain progressive photon density estimation, we found it perfectly reasonable to employ biased but consistent shift mappings during progressive gradient-domain photon density estimation: e.g., with density kernel overlaps that reduce the computational cost of the shift mapping (mixed shift), or by employing a manifold exploration relaxation scheme. In fact, such biased (but consistent) shift mappings have never been presented in the prior art.

We expect these insights to open new avenues of exploration for more efficient shift mapping schemes. This is particular interesting in the context of shift mappings that currently employ manifold exploration as, while being useful, the manifold exploration operation can form a computational bottleneck as it was not originally designed for gradient-domain shift mappings. Our work also points to several promising open questions that deserve future investigation. For example, can Monte Carlo-based shift mappings be devised to enable (efficient) volumetric rendering in gradient-domain uni- and bi-directional path tracing? Can Monte Carlo and (volumetric) density estimation formalism be combined, much like in the primal-domain [Křivánek et al. 2014], to form a unified model of image synthesis in the gradient domain.

ACKNOWLEDGMENTS

REFERENCES

Pablo Bauszat, Victor Petitjean, and Elmar Eisemann. 2017. Gradient-Domain Path Reusing. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)* (2017).

Pravin Bhat, Brian Curless, Michael Cohen, and C. Lawrence Zitnick. 2008. Fourier Analysis of the 2D Screened Poisson Equation for Gradient Domain Problems. In *The European Conference on Computer Vision (ECCV)*.

Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond Points and Beams: Higher-Dimensional Photon Samples for Volumetric Light Transport. *ACM Transactions on Graphics (Proc. of SIGGRAPH)* (2017).

Eva Cerezo, Frederic Perez-Cazorla, Xavier Pueyo, Francisco Seron, and François Sillion. 2005. A Survey on Participating Media Rendering Techniques. *The Visual Computer* (2005).

Subrahmanyan Chandrasekhar. 2013. *Radiative Transfer*. Courier Corporation.

Carsten Dachsbacher, Jaroslav Křivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable Realistic Rendering with Many-Light Methods. *Computer Graphics Forum* (2014).

Iliyan Georgiev, Jaroslav Krivanek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint importance sampling of low-order volumetric scattering. *ACM Transactions on Graphics* (2013).

Toshiya Hachisuka, Iliyan Georgiev, Wojciech Jarosz, Jaroslav Křivánek, and Derek Nowrouzezahrai. 2017. Extended Path Integral Formulation for Volumetric Transport. In *Proc. of EGSR (Experimental Ideas & Implementations)*.

Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic progressive photon mapping. *ACM Transactions on Graphics* (2009).

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics* (2012).

Miloš Hašan, Jaroslav Křivánek, Bruce Walter, and Kavita Bala. 2009. Virtual Spherical Lights for Many-light Rendering of Glossy Scenes. *ACM Transactions on Graphics* (2009).

Louis G Henyey and Jesse Leonard Greenstein. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93 (1941).

Binh-Son Hua, Adrien Gruson, Derek Nowrouzezahrai, and Toshiya Hachisuka. 2017. Gradient-Domain Photon Density Estimation. *Computer Graphics Forum* (2017).

Wenzel Jakob. 2010. Mitsuba Renderer. (2010).

Wenzel Jakob. 2013. *Light Transport on Path-Space Manifolds*. Ph.D. Dissertation. Cornell University.

Wenzel Jakob and Steve Marschner. 2012. Manifold exploration: a Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics* (2012).

Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011a. A Comprehensive Theory of Volumetric Radiance Estimation Using Photon Points and Beams. *ACM Transactions on Graphics* (2011).

Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011b. Progressive photon beams. *ACM Transactions on Graphics* (2011).

Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008. The Beam Radiance Estimate for Volumetric Photon Mapping. *Computer Graphics Forum* (2008).

Henrik Wann Jensen. 2001. *Realistic image synthesis using photon mapping*.

Henrik Wann Jensen and Per H Christensen. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH*.

James T. Kajiya. 1986. The Rendering Equation. *ACM Transactions on Graphics* (1986).

Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain path tracing. *ACM Transactions on Graphics* (2015).

Claude Knaus and Matthias Zwicker. 2011. Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics* (2011).

Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics* (2014).

Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. *Computer Graphics Forum* (2012).

Eric P Lafortune and Yves D Willems. 1996. Rendering participating media with bidirectional path tracing. In *Rendering Techniques*.

Jaakko Lehtinen, Tero Karras, Samuli Laine, Miika Aittala, Frédo Durand, and Timo Aila. 2013. Gradient-domain metropolis light transport. *ACM Transactions on Graphics* (2013).

Marco Manzi, Markus Kettunen, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain bidirectional path tracing. *Eurographics Symposium on Rendering* (2015).

Marco Manzi, Markus Kettunen, Frédo Durand, Matthias Zwicker, and Jaakko Lehtinen. 2016a. Temporal gradient-domain path tracing. *ACM Transactions on Graphics* (2016).

M. Manzi, D. Vicini, and M. Zwicker. 2016b. Regularizing Image Reconstruction for Gradient-Domain Rendering with Feature Patches. *Computer Graphics Forum* (2016).

Tomas Möller and Ben Trumbore. 1997. Fast, Minimum Storage Ray-triangle Intersection. *Journal of Graphics Tools* (1997).

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012a. Progressive virtual beam lights. In *Computer Graphics Forum*.

Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012b. Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics* (2012).

Fabrice Rousselle, Wojciech Jarosz, and Jan Novák. 2016. Image-space Control Variates for Rendering. *ACM Transactions on Graphics* (2016).

Weilun Sun, Xin Sun, Nathan A Carr, Derek Nowrouzezahrai, and Ravi Ramamoorthi. 2017. Gradient-Domain Vertex Connection and Merging. (2017).