EUROGRAPHICS 2017 / L. Barthe and B. Benes
(Guest Editors)

Volume 36 (2017), Number 2

Gradient-Domain Photon Density Estimation

Binh-Son Hua[1] Adrien Gruson[2] Derek Nowrouzezahrai[3] Toshiya Hachisuka[2]

[1]Singapore University of Technology and Design [2]The University of Tokyo [3]McGill University

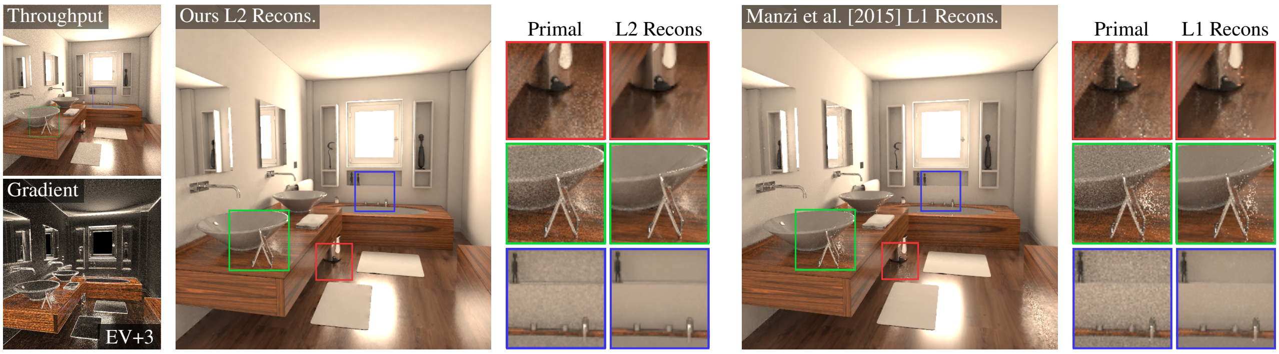

Figure 1: *Equal-time rendering (five minutes) of a bathroom scene with gradient-domain photon density estimation and bi-directional path tracing [MKA*15]. Our method more efficiently handles regions dominated by complex transport paths.*

Abstract

The most common solutions to the light transport problem rely on either Monte Carlo (MC) integration or density estimation methods, such as uni- & bi-directional path tracing or photon mapping. Recent gradient-domain extensions of MC approaches show great promise; here, gradients of the final image are estimated numerically (instead of the image intensities themselves) with coherent paths generated from a deterministic shift mapping. We extend gradient-domain approaches to light transport simulation based on density estimation. As with previous gradient-domain methods, we detail important considerations that arise when moving from a primal- to gradient-domain estimator. We provide an efficient and straightforward solution to these problems. Our solution supports stochastic progressive density estimation, so it is robust to complex transport effects. We show that gradient-domain photon density estimation converges faster than its primal-domain counterpart, as well as being generally more robust than gradient-domain uni- & bi-directional path tracing for scenes dominated by complex transport.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: 3D Graphics & Realism—Raytracing

1. Introduction

Recent adoption of physically-based rendering in feature films, animation, and video game industry has promoted the importance of efficient numerical solutions to light transport simulation. The most common numerical solvers rely on Monte Carlo (MC) integration such as uni- and bi-directional path tracing, or density estimation such as photon mapping. While MC integration works well in many cases, it struggles to resolve effects resulting from *complex light transport paths*, namely those that are localized on small manifolds in a high-dimensional path space. Common examples are so-called *SDS paths* where a diffuse scattering event lies between two or more specular (or near-specular) events. Density estimation approaches are ideal for sampling such types of light transport paths.

Recently, new MC integration approaches based on gradient-domain estimation have improved upon their primal-domain counterparts: here, *image gradients* are estimated numerically (as opposed to *image intensities*), after which a reconstruction post-process is used to compute the final image. The seminal work of Lehtinen et al. [LKL*13] proposed a gradient-domain variant of Metropolis Light Transport [Vea97], and follow-up works similarly extended uni- and bi-directional path tracing to the image gradient-domain [KMA*15, MKA*15].

We propose the first gradient-domain rendering approach based on density estimation. As with prior work, our technique is conceptually simple, but a special care is necessary to ensure robustness and convergence when shifting light paths to compute image-space gradients, computing appropriate Jacobians, and performing density estimation itself. As such, our method cannot be constructed as a trivial variant of existing gradient-domain techniques. Specifically, our technical contributions comprise:

- a procedure to generate offset paths compatible with photon density estimation; namely, we present a *hybrid shift mapping* technique that exploits *half-vector copy* and *manifold exploration*,
- the derivation of the hybrid shift mapping Jacobian necessary for compensating for the change in estimation domains, and
- theory and algorithms for *offset photon density estimation* and *image-space gradient evaluation* which are designed to support stochastic progressive density estimation [HJ09].

Our experiments show that gradient-domain photon density estimation inherits many of the advantages of density estimation *and* gradient-domain rendering, including the ability to sample complex SDS paths more efficiently than MC path-based solvers *and* the ability to leverage primal-domain reconstruction as a means of filtering away noise. Furthermore, gradient-domain photon density estimation performs on par with existing methods on simpler light transport paths (i.e., diffuse-only interreflections), where primal-domain photon density estimation can become inefficient. Figure 1 illustrates these benefits on an example scene where (gradient-domain) uni- and bi-directional path tracing both struggle.

2. Related Work

Gradient-domain rendering Lehtinen et al. [LKL*13] presented a seminal work of gradient-domain light transport which shows how to perform shift mapping using manifold exploration [JM12] for reducing the variance of gradient estimates in the context of Metropolis Light Transport [Vea97]. Their key insight was that focusing sampling budget on pixels with large gradients could lead to significant gains. Kettunen et al. [KMA*15] went on to show that this adaptive sampling of gradients is not fundamentally necessary, proposing a simplified gradient-domain path tracing algorithm that still maintains similar convergence advantages. Their alternative shift mapping approach maintains half-vectors when offsetting base paths.

In a follow-up work, Manzi et al. [MKA*15] extended gradient-domain rendering to the bi-directional setting with a shift mapping based on manifold exploration. In order to achieve practical performance, they consider only diffuse-to-diffuse vertex connections. This limits the applicability of their approach to handle scenes with many specular interactions. We will directly address this limitation by devising path generation and density estimation formulations to enable gradient-domain light transport in the context of photon density estimation. Theses formulations combine the advantages of density estimation (i.e., robust treatment of specular transport) and gradient-domain rendering (i.e., noise filtering through primal-domain reconstruction).

Recently, Manzi et al. [MKD*16] also explored temporal coherence in gradient-domain rendering with a novel adaptive sampling scheme. This work is orthogonal to ours and could be implemented on top of our technique to support integration over time.

Photon density estimation Unlike MC path-based simulation, photon density estimation relies on computing local densities when forming full light transport paths between light- and eye-subpaths. Traditional photon mapping [Jen01] is a two-pass algorithm that stores photons shot from light sources in the first pass. The number of photons in the first pass is inherently limited by the total available memory of the system. Hachisuka et al. [HOJ08, HJ09] proposed a progressive variant that eliminates this limitation by storing photon statistics at gather points in a manner that allows photon data to be discarded and accumulated progressively. Density estimation can also be combined with MC path-based estimators using multiple importance sampling and an extended path-space formulation of light transport that rectifies MC integral estimators with density estimation [HPJ12, GKDS12]. We propose a gradient-domain density estimation and refinement scheme to compute image gradients *progressively*, inheriting the benefits of progressive (primal-domain) photon mapping.

Symbol	Meaning
$\mathbf{z}_k$	vertex k on sensor subpath
$\mathbf{z}_t, \mathbf{z}'_t$	base and offset gather point
$\mathbf{y}_k$	vertex k on light subpath
$\mathbf{y}_s, \mathbf{y}'_s$	base and offset photon
$G(\mathbf{y}_b, \mathbf{y}_{b+1})$	geometry term
$G(\mathbf{y}_b, \dots, \mathbf{y}_s)$	generalized geometry term
$\mathbf{o}_k$	half-vector at a vertex k
E, L, D, S	eye, light, diffuse, and specular vertex

Table 1: *Table of notation.*

Reconstruction In gradient-domain rendering, final images are reconstructed using Poisson reconstruction after image gradients are estimated. Lehtinen et al. [LKL*13] compared an unbiased L2 reconstruction with a more stable, biased L1-norm based reconstruction. The latter does not suffer from dipole artifacts at the cost of some (usually negligible) bias. Manzi et al. [MVZ16] proposed using a G-buffer to better regularize this reconstruction stage. We rely on traditional L2 and L1 norm reconstruction [LKL*13] when generating our final results, thus our algorithm can be used as a drop-in replacement for existing gradient-domain estimators.

3. Gradient-Domain Rendering

Image gradients are estimated using pixel finite differences. Base pixel intensities can be estimated with paths sampled using e.g., MC techniques, and offset pixel intensities (next to the base pixel) can similarly be estimated using offset paths built atop the same set of bases paths. Noise in image-space gradients is reduced by maintaining as much coherence between offset and base paths as possible. We begin by detailing the integral used to estimate the intensity of an offset pixel.

Given a light transport path $\mathbf{x} = \mathbf{x}_0 \dots \mathbf{x}_\ell$ of length ℓ in path space $\mathcal{P}$ (where $\mathbf{x}_0$ lies on an emitter and $\mathbf{x}_\ell$ on a sensor), the integral to compute the intensity I_i of pixel i is

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

where $f(\mathbf{x})$ is the path throughput and $h_i(\mathbf{x})$ is the reconstruction filter at pixel i. In order to estimate image-space gradients, gradient-domain rendering uses finite differences of neighboring pixels, i.e., $\Delta_{ij} = I_j - I_i$ where Δ_{ij} is the finite difference at pixel i due to a neighboring pixel j. While one can apply Equation 1 independently across pixels and subsequently perform finite differencing, we can exploit path coherence between neighboring pixels i and j:

$$I_j^i = \int_{\mathcal{P}} h_i(\mathbf{x}) f(T(\mathbf{x})) \left| \frac{dT(\mathbf{x})}{d\mathbf{x}} \right| d\mathbf{x}, \qquad (2)$$

where T is the shift mapping operator that transforms a path through pixel i (the base path) into a path through pixel j (the offset path), and I_j^i is the intensity at pixel j estimated by shift mapping. To exploit coherence, the operator T is designed such that the base and offset path share as many common vertices as possible. The Jacobian determinant accounts for path density changes due to shift [LKL*13]. Note that the reconstruction filter at the neighboring pixel is equal to that of the base pixel, and so $h_j(T(\mathbf{x})) = h_i(\mathbf{x})$. This formulation assumes that neighboring pixels tend to have similar light transport paths, and so correlated estimation using shifted paths is more efficient than independent estimation using only Equation 1.

We use forward finite differences to estimate horizontal and vertical gradients. Except at image boundaries, we can use either the "primal" integral estimate (Equation 1) or the integral estimate using shift mapping (Equation 2). Consequently, the gradient due to a pair of pixels i and j can be computed using a combination of these two strategies, with equal weighting (i.e., 0.5). Kettunen et al. [KMA*15] propose setting the weights according to a multiple importance sampling scheme, which can further reduce artifacts near depth discontinuities.

Let the estimated horizontal and vertical gradients be G_x and G_y, and the primal-domain image estimate be I_0. Linearizing images to 1D vectors, we can express the final image I as the solution to the following Poisson reconstruction problem [BCCZ08]:

$$I = \arg\min \|D_x I - G_x\|_p + \|D_y I - G_y\|_p + \lambda \|I - I_0\|_p, \qquad (3)$$

where D_x and D_y are 1D convolution operators that compute finite differences along the image's x and y axes, λ controls the similarity between the final image and the primal-domain throughput image, and p (which can take a value 1 or 2) denotes the choice norm used during reconstruction. Specifically, an L2 norm yields an unbiased reconstruction and L1 norm yields a result that is more robust to reconstruction artifacts. For all our results, we solve this optimization problem using an iterative re-weighted least squares solver [LKL*13].

4. Gradient-Domain Photon Density Estimation

Let $\mathbf{x} = \mathbf{y}_0 \ldots \mathbf{y}_s, \mathbf{z}_t \ldots \mathbf{z}_0$ be a light transport path of length $\ell = s + t$ sampled by photon density estimation, where $\mathbf{y} = \mathbf{y}_0 \ldots \mathbf{y}_s$ and $\mathbf{z} = \mathbf{z}_0 \ldots \mathbf{z}_t$ are the light and sensor subpaths, vertex $\mathbf{y}_s$ is a photon, and $\mathbf{z}_t$ a gather point. According to path space extensions [HPJ12], the measurement contribution function of a photon density estimation path includes a kernel $K(\mathbf{y}_s, \mathbf{z}_t)$ to capture the flux transfer between the "disconnected" vertices $\mathbf{y}_s$ and $\mathbf{z}_t$. In our work, we use a simple uniform kernel $K(\mathbf{p}, \mathbf{q}) = K = 1/(\pi r^2)$ with radius r, regardless of the vertices $\mathbf{p}$ and $\mathbf{q}$. We assume that light and sensor subpaths are generated independently and, therefore, the path measure is simply $d\mathbf{x} = d\mathbf{x}_0 \ldots d\mathbf{x}_\ell$. This allows us to treat density estimation and MC integration in a unified manner, and so to apply and extend previous shift mapping methods to our problem [LKL*13, KMA*15, MKA*15]. We use Heckbert's regular expression notation [Hec90] to classify families of paths: each scattering event along a path from eye (E) to light (L) could be diffuse (D) or specular (S).

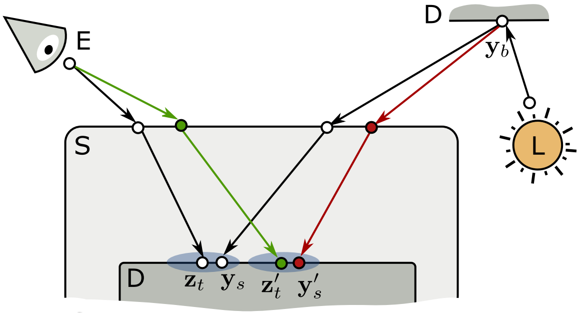

Figure 2: *An example of shifting for photon density estimation. The sensor and light subpaths are first individually shifted, then density estimation is performed on the offset path using the same kernel of the base path (shaded region).*

In order to evaluate gradients, we have to shift each photon density estimation path to generate an offset path. However, doing so with existing shift mapping techniques [LKL*13, KMA*15] is neither straightforward nor feasible, since a photon density estimation path is bi-directional and contains two "disconnected" subpaths. Half-vector copies [KMA*15] are designed to shift unidirectional paths and, while manifold exploration [JM12] is able to shift bidirectionally-sampled paths, the original manifold exploration algorithm is only suitable for complete transport paths (i.e., and not a pair of disconnected subpaths). None of these existing shift mapping approached are appropriate for shifting a photon density estimation path, all while maintaining the kernel.

Our solution is to first shift the base sensor and light subpaths, and then perform kernel evaluation to generate the corresponding offset paths (see Figure 2). We design a shift mapping that restricts the offset photon within the support of the density estimation kernel, keeping the same kernel size as the base path's. The benefit of this shift mapping scheme is two-fold: first, it allows us to avoid many offset paths with zero contribution when the subpaths are shifted independently; second, we can tailor different shift mapping strategies to the sensor and light subpaths.

Specifically, we shift sensor subpaths using a half-vector copy and light subpaths using manifold exploration [JM12]. The rationale is as follows: in photon density estimation, a sensor subpath comprises only ES^*D transport configurations, which are a simpler subset of configurations handled by gradient-domain path tracing (where half-vector copy has been shown to be highly efficient and easier to implement [KMA*15]); since light subpaths can have $L(D|S)^*D$ configurations, a half-vector copy is not suitable as it requires a DD subpath on the base path and a $(D|S)D$ configuration on the offset path in order to perform reconnection. Instead, we employ a modified manifold exploration to shift light subpaths, and we show

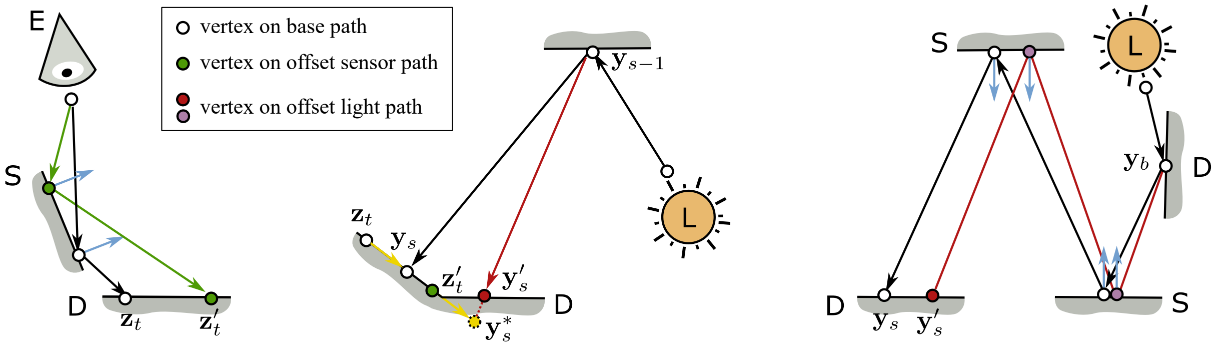

(a) Sensor subpath shift (b) Photon shift (c) Light subpath manifold walk

Figure 3: *Shift mapping strategies: (a) an offset sensor subpath is generated by half-vector copy, (b) an offset photon is generated by preserving the local offset vector (yellow) between the base gather point and the base photon, and (c) the remaining vertices (purple) of the offset light subpath are determined according to a manifold walk.*

how to compute the Jacobian in this case. Figure 3 illustrates the details of our shift mapping techniques.

4.1. Shifting sensor subpaths

In photon density estimation, we assume that a gather point is deposited at the first diffuse vertex on a sensor subpath. The half-vector copy shift mapping can be used without modification to generate the offset sensor subpath here (see Figure 3a). During sensor subpath tracing, each vertex can be classified based on the roughness of the current sampled BSDF component [KMA*15].

Equation 2 is expressed in the surface area measure, and the Jacobian has to remain compatible with this measure. In gradient-domain path tracing, the Jacobian of half-vector copy is expressed in the solid angle measure, and so applying it in our context necessitates a change of measure. Let the last two vertices on the base path be $\mathbf{z}_i$ and $\mathbf{z}_{i+1}$ during eye path tracing, the corresponding vertices on an offset path be $\mathbf{z}'_i$ and $\mathbf{z}'_{i+1}$, and the Jacobian in the solid angle measure be J_ω, the Jacobian in the area measure can be written as:

$$J = J_\omega \left[\frac{G(\mathbf{z}_i, \mathbf{z}_{i+1})}{G(\mathbf{z}'_i, \mathbf{z}'_{i+1})}\right] \left[\frac{\cos(\mathbf{n}'_i, \mathbf{z}'_i \to \mathbf{z}'_{i+1})}{\cos(\mathbf{n}_i, \mathbf{z}_i \to \mathbf{z}_{i+1})}\right], \qquad (4)$$

where G is the geometry term and $\mathbf{n}_i$ is the normal at path vertex i. Note that if $\mathbf{z}_i$ and $\mathbf{z}'_i$ are specular, the Jacobian is simply 1.

4.2. Shifting light subpaths

We generate offset light subpaths in two steps: the position of the last vertex on the offset light subpath (the offset photon) is first determined (red vertex in Figure 3b), before computing the intermediate vertices to connect the offset photon to the base light subpath (purple vertices in Figure 3c).

Step 1. Shift mapping is effective when the base path is perturbed as little as possible. To do so in photon density estimation, we aim to preserve the local photon distribution at the base and offset gather points. First, a candidate location $\mathbf{y}^*_s$ for the offset photon (yellow vertex in Figure 3b) is determined as

$$\mathbf{y}^*_s = \mathbf{z}'_t + M' M^{-1} (\mathbf{y}_s - \mathbf{z}_t), \qquad (5)$$

where $\mathbf{z}_t$ and $\mathbf{z}'_t$ are the base and offset gather points, $\mathbf{y}_s$ is the base photon location, and M and M' are transformation matrices that displace a point from surface tangent space to world space at the base and offset gather locations.

Since the candidate location is not guaranteed to lie on a surface, we trace a ray from the preceeding vertex of the base photon towards the candidate location, and use the intersection $\mathbf{y}'_s = \mathtt{trace}(\mathbf{y}_{s-1} \to \mathbf{y}^*_s)$ as the offset photon location, where `trace` function returns ray intersection with the scene geometry.

The Jacobian of this shift is

$$\left|\frac{\partial \mathbf{y}'_s}{\partial \mathbf{y}_s}\right| = \left|\frac{\partial \mathbf{y}'_s}{\partial \mathbf{y}^*_s}\right| \left|\frac{\partial \mathbf{y}^*_s}{\partial \mathbf{y}_s}\right| = \frac{G(\mathbf{y}_{s-1}, \mathbf{y}^*_s)}{G(\mathbf{y}_{s-1}, \mathbf{y}'_s)}, \qquad (6)$$

where we leverage the fact that the angles from $\mathbf{y}_{s-1}$ to $\mathbf{y}^*_s$ and to $\mathbf{y}'_s$ are equal (due to the ray tracing `trace` function in the first term), and we know the second term $|\partial \mathbf{y}^*_s / \partial \mathbf{y}_s| = 1$ by taking derivatives of both sides of Equation 5.

Step 2. We choose different techniques to generate the intermediate vertices and to form a complete connection according to the type of scattering event at a vertex on the base path. For light subpaths that end with DD configurations, the offset photon can be connected directly to the preceeding vertex of the base photon. Since no new intermediate vertex is introduced, the Jacobian is simply $\left|\frac{\partial \mathbf{y}'_s}{\partial \mathbf{y}_s}\right|$ in this case.

For the other types of subpath configurations, we shift using a variant of manifold exploration [JM12]: let the base light subpath be $\mathbf{y}_0 \dots \mathbf{y}_b \mathbf{y}_{b+1} \dots \mathbf{y}_s$, where $\mathbf{y}_b$ is a diffuse vertex; the offset light subpath will have the form $\mathbf{y}_0 \dots \mathbf{y}_b \mathbf{y}'_{b+1} \dots \mathbf{y}'_s$, where the reconnection to the base light subpath occurs at $\mathbf{y}_b$. The segment $\mathbf{y}'_{b+1} \dots \mathbf{y}'_{s-1}$ contains new vertices generated by manifold exploration, when connecting $\mathbf{y}_b$ to $\mathbf{y}'_s$. Note that the vertices in $\mathbf{y}'_{b+1} \dots \mathbf{y}'_{s-1}$ can be purely specular or glossy, and we treat these two cases differently.

Jacobian for specular chains The Jacobian is only required to account for vertices that change during a shift. When vertices in segment $\mathbf{y}'_{b+1} \dots \mathbf{y}'_{s-1}$ are all purely specular, we only integrate over

$\mathbf{y}'_{b+1}$ in the constrained path space integral; other vertices do not appear due to their Dirac responses. The Jacobian here is

$$J = \left|\frac{\partial \mathbf{y}'_{b+1}}{\partial \mathbf{y}_{b+1}}\right| = \left|\frac{\partial \mathbf{y}'_{b+1}}{\partial \mathbf{y}'_{s}}\right| \left|\frac{\partial \mathbf{y}'_{s}}{\partial \mathbf{y}_{s}}\right| \left|\frac{\partial \mathbf{y}_{s}}{\partial \mathbf{y}_{b+1}}\right|, \tag{7}$$

where we compute the middle term according to Equation 6. Using perturbation theory [JM12], we compute the third term as

$$\left|\frac{\partial \mathbf{y}_{b+1}}{\partial \mathbf{y}_{s}}\right| = \frac{G(\mathbf{y}_b, \mathbf{y}_{b+1}, \ldots, \mathbf{y}_s)}{G(\mathbf{y}_b, \mathbf{y}_{b+1})}, \tag{8}$$

where $G(\mathbf{y}_b, \mathbf{y}_{b+1}, \ldots, \mathbf{y}_s)$ is the generalized geometry term of the specular chain $\mathbf{y}_b \ldots \mathbf{y}_s$. We compute the first term in the same way.

Jacobian for glossy chains For subpaths with glossy vertices, the formulation becomes more involved as we must consider the integration of glossy vertices that are normally classified as specular events during a standard manifold walk. The Jacobian becomes

$$J = \left|\frac{\partial[\mathbf{y}'_{b+1}\cdots\mathbf{y}'_{s}]}{\partial[\mathbf{y}_{b+1}\cdots\mathbf{y}_{s}]}\right| = \left|\frac{\partial[\mathbf{y}'_{b+1}\cdots\mathbf{y}'_{s}]}{\partial[\mathbf{o}'_{b+1}\cdots\mathbf{y}'_{s}]}\right| \left|\frac{\partial[\mathbf{o}'_{b+1}\cdots\mathbf{y}'_{s}]}{\partial[\mathbf{o}_{b+1}\cdots\mathbf{y}_{s}]}\right| \left|\frac{\partial[\mathbf{o}_{b+1}\cdots\mathbf{y}_{s}]}{\partial[\mathbf{y}_{b+1}\cdots\mathbf{y}_{s}]}\right|, \tag{9}$$

where $\mathbf{o}_k$ denotes the half-vector at a vertex k. This determinant has three terms which corresponds to replacing the integration over $\mathbf{y}'_{b+1}\ldots\mathbf{y}'_s$ to one over $\mathbf{o}'_{b+1}\ldots\mathbf{y}'_s$, and then applying the Jacobian in this latter reparameterized space (the middle term), before converting back to the original parameterization. The reformulation of the area measure form of the path integral into an integral over half-vectors was first proposed by Jakob and Marschner in their manifold exploration work [JM12], and was later extended in Kaplanyan et al.'s half-vector space light transport [KHD14]. Lehtinen et al. [LKL*13] also leveraged this idea to calculate the Jacobian in their gradient-domain Metropolis light transport.

The middle term is easy to calculate: since every half-vector is preserved during a manifold walk, the middle term simplifies to $\left|\partial \mathbf{y}'_s/\partial \mathbf{y}_s\right|$, which can be computed as in **Step 1**. We compute the first and third terms as follows: recall that, from perturbation theory [JM12], it is possible to calculate the derivative blocks $\partial \mathbf{o}_k/\partial \mathbf{y}_j$ and $\partial \mathbf{y}_k/\partial \mathbf{y}_s$ from half-vector constraints, for each $k \in [b+1, s-1]$ and $j \in [b+1, s]$. We assemble these blocks to build the Jacobian matrix

$$C = \left(\frac{\partial[\mathbf{o}_{b+1}\cdots\mathbf{y}_{s}]}{\partial[\mathbf{y}_{b+1}\cdots\mathbf{y}_{s}]}\right), \tag{10}$$

for which the determinant $|C|$ is the third term that we require. Similarly, for the first term, let $C' = \left(\partial[\mathbf{o}'_{b+1}\cdots\mathbf{y}'_s]/\partial[\mathbf{y}'_{b+1}\cdots\mathbf{y}'_s]\right)$. If every vertex $\mathbf{y}_{b+1}\ldots\mathbf{y}_s$ is glossy (but treated as specular during manifold walks), the determinant $\left|C'^{-1}\right|$ is the first term we require.

The only remaining issue is when there exists a *combination* of purely specular *and* glossy vertices in the subpath segment $\mathbf{y}_{b+1}\ldots\mathbf{y}_s$. Here, we must selectively discard certain columns and rows in C'^{-1} before computing the determinant [JM12]: specifically, the column k that corresponds to the half-vector of any specular vertex $k \in [b+1, s-1]$ is discarded because the half-vector does not contribute to the integral. Similarly, we discard row $k+1$ since, when vertex k is specular, vertex $k+1$ does not contribute to the integral (see inlined figure). Glossy interactions originally treated as specular ones are denoted "S", and red lines are the columns and rows we discard. Figure 4 demonstrates the shift mapping for two simple scenes with specular and glossy materials. The relative error plots in this figure empirically demonstrate the convergence and confirm that the proposed shift mapping and Jacobian are accurate.

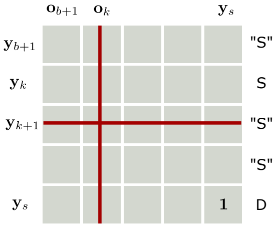

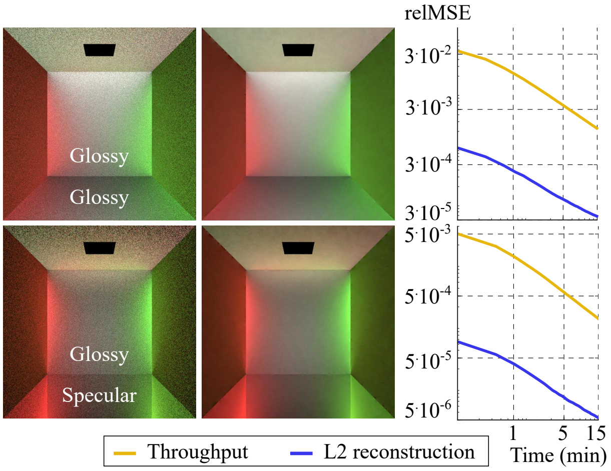

Figure 4: *Modified Cornell Box with modified materials, rendered with our shift mapping. Only paths where the preceeding vertex of a photon is specular and glossy are considered. From left to right: the throughput, L2 reconstruction, and the error plot of the throughput and the reconstruction. Both throughput and L2 reconstruction error converge. The L2 reconstruction error is an order of magnitude lower than the throughput error, which confirms the correctness of our shift mapping and Jacobians.*

4.3. Implementation details

Progressive photon mapping Our method can be easily implemented atop stochastic progressive photon mapping [HJ09] by tracking the total flux of the offset paths. In our implementation, each gather point stores the flux and total number of photons for both the base and offset paths, which can be used to compute the gradients. All flux values in the gather point share the same kernel radius reduction rate. Note that density estimation is only applied to the base gather points. For each base photon, we first compute its flux contribution, then apply shift mapping to generate offset gather points and offset photons. Finally, we compute the contribution of the offset photons, and update the appropriate statistics (i.e., the radii and the total number of photons).

Reversibility To ensure reversibility, we additionally ensure that the line of sight from the preceeding vertex of the offset photon to the base photon is not occluded. If the offset path fails this visibility check, the shift is not reversible and the weight of this offset path (used as a strategy to compute the gradients) is set to 1.

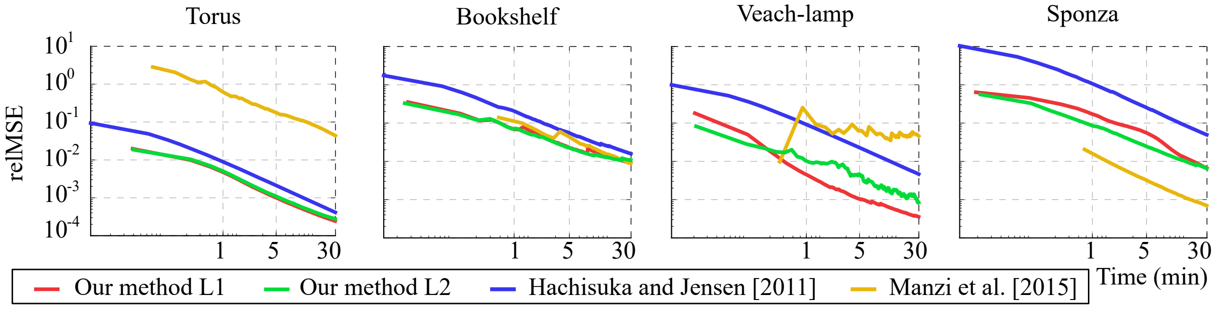

Figure 5: *Convergence plots across scenes, using relative MSE metric. Gradient-domain photon density estimation outperforms the primal-domain algorithm for every scene. Our method also converges faster than GBDPT [MKA*15] in scenes with complex light transport paths (Torus and Veach-lamp). It also performs competitively on Bookshelf and also Sponza, a scene comprising completely diffuse light transport.*

Multiple importance sampling As in previous gradient-domain methods, we apply multiple importance sampling (MIS) to reduce dipole artifacts that could occur during L2 reconstruction, especially around concave corners. For a photon density estimation path, since its sensor and light subpaths are sampled and shifted independently, the balance heuristic for MIS can be expressed as

$$w(\mathbf{x}) = \frac{p(\mathbf{y})p(\mathbf{z})}{p(\mathbf{y})p(\mathbf{z}) + p(T(\mathbf{y}))p(T(\mathbf{z}))\,|T'(\mathbf{y})T'(\mathbf{z})|}, \qquad (11)$$

where the probabilities are in the area measure and T' denotes the Jacobian of the shift for the corresponding subpaths.

Layered BSDFs We follow the same classification scheme for glossy vertices, which could include both diffuse and specular components, as in gradient-domain path tracing [KMA*15]. A vertex is classified as diffuse if its BSDF roughness is greater than an certain threshold, and specular otherwise. For a path vertex that has a multi-component BSDF, only one component is sampled and considered at a time. Offset path vertices are initialized with the same material type as their base path.

5. Results

We implemented our method in Mitsuba [Jak10], and compare against stochastic progressive photon mapping [HJ09], the publicly available implementations of gradient-domain path tracing [KMA*15], and gradient-domain bidirectional path tracing [MKA*15]. Unless other stated, we use the following parameter settings: the λ parameter for Poisson reconstruction is set to a small value, $\lambda = 0.2$, as in previous works [KMA*15, MKA*15]; the density estimation radius is initialized using ray differentials, to approximate the projected pixel size, and we use a (conservative) reduction ratio of $\alpha = 0.9$; we shoot 1 million photons per progressive iteration, and the BSDF roughness threshold (see Section 4.3) is set to 0.05.

We evaluate gradient-domain photon density estimation (GPM) on four scenes: Torus, Bookshelf, Veach-Lamp, and Sponza. Our error metric is relative mean square error $relMSE = (1/n)\sum(R-I)^2/(R^2+0.001)$ where I is the reconstructed image and R is the reference, and n is the number of pixels. In contrast to previous works that use this metric [KMA*15, MKA*15], we found that it is not necessary to discard any high error pixels. We compare our method to stochastic progressive photon mapping (SPPM) and gradient-domain bidirectional path tracing (GBDPT). We do not include gradient-domain path tracing (GPT) as it only performed on par with GBDPT on the teaser scene and the Sponza, but worse elsewhere. More comprehensive results for GPT could be found in our supplementary document.

Figure 5 compares the convergence across algorithms in an equal-time setting with all scenes rendered up to 30 minutes. We output renderings after every five photon passes and collect the images for error measurement. Our technique and GBDPT need to compute gradients and thus take more time than SPPM per photon pass as shown in the plot.

As can be seen, GPM L1 reconstruction usually has lower error than L2 reconstruction, especially at the later photon passes. GBDPT converges for all scenes except the Veach-Lamp; here GBDPT cannot render the glass egg efficiently, which contributes to its slow convergence. We include more comparisons, including relative error maps and interactive comparisons, in our supplemental document.

Figure 6 compares equal-time renderings of our test scenes. As can be seen, in general, GPM consistently outperforms both SPPM and GBDPT in scenes dominated with SDS paths, like Torus and the cupboard area in the Bookshelf scene. Since the gradients are well estimated here, using an L2 reconstruction is sufficient to obtain high-quality reconstructions. In contrast, GBDPT's subpath connection strategy is not able to efficiently sample SDS paths, causing these regions to appear darker than the ground truth. GBDPT's L2 reconstruction also causes dipole artifacts and is generally less robust than L1 reconstruction in these scenarios. Therefore, in Figure 6, we compare our L2 reconstruction results with GBDPT's L1 reconstructions. For scenes with mixed diffuse and glossy materials, our method works comparably to GBDPT, and is still (significantly) more efficient than SPPM.

Compared to SPPM, our method is 1.5 to 3.5× slower at each photon pass. However, since our shift mapping generates coherent light paths, the gradients converges much more quickly, which results in our method's superior *overall performance* compared to SPPM (i.e., using fewer photon passes). In our method, the overhead of each photon pass is mostly due to the manifold exploration we perform, especially when it fails to find valid paths. In these cases,

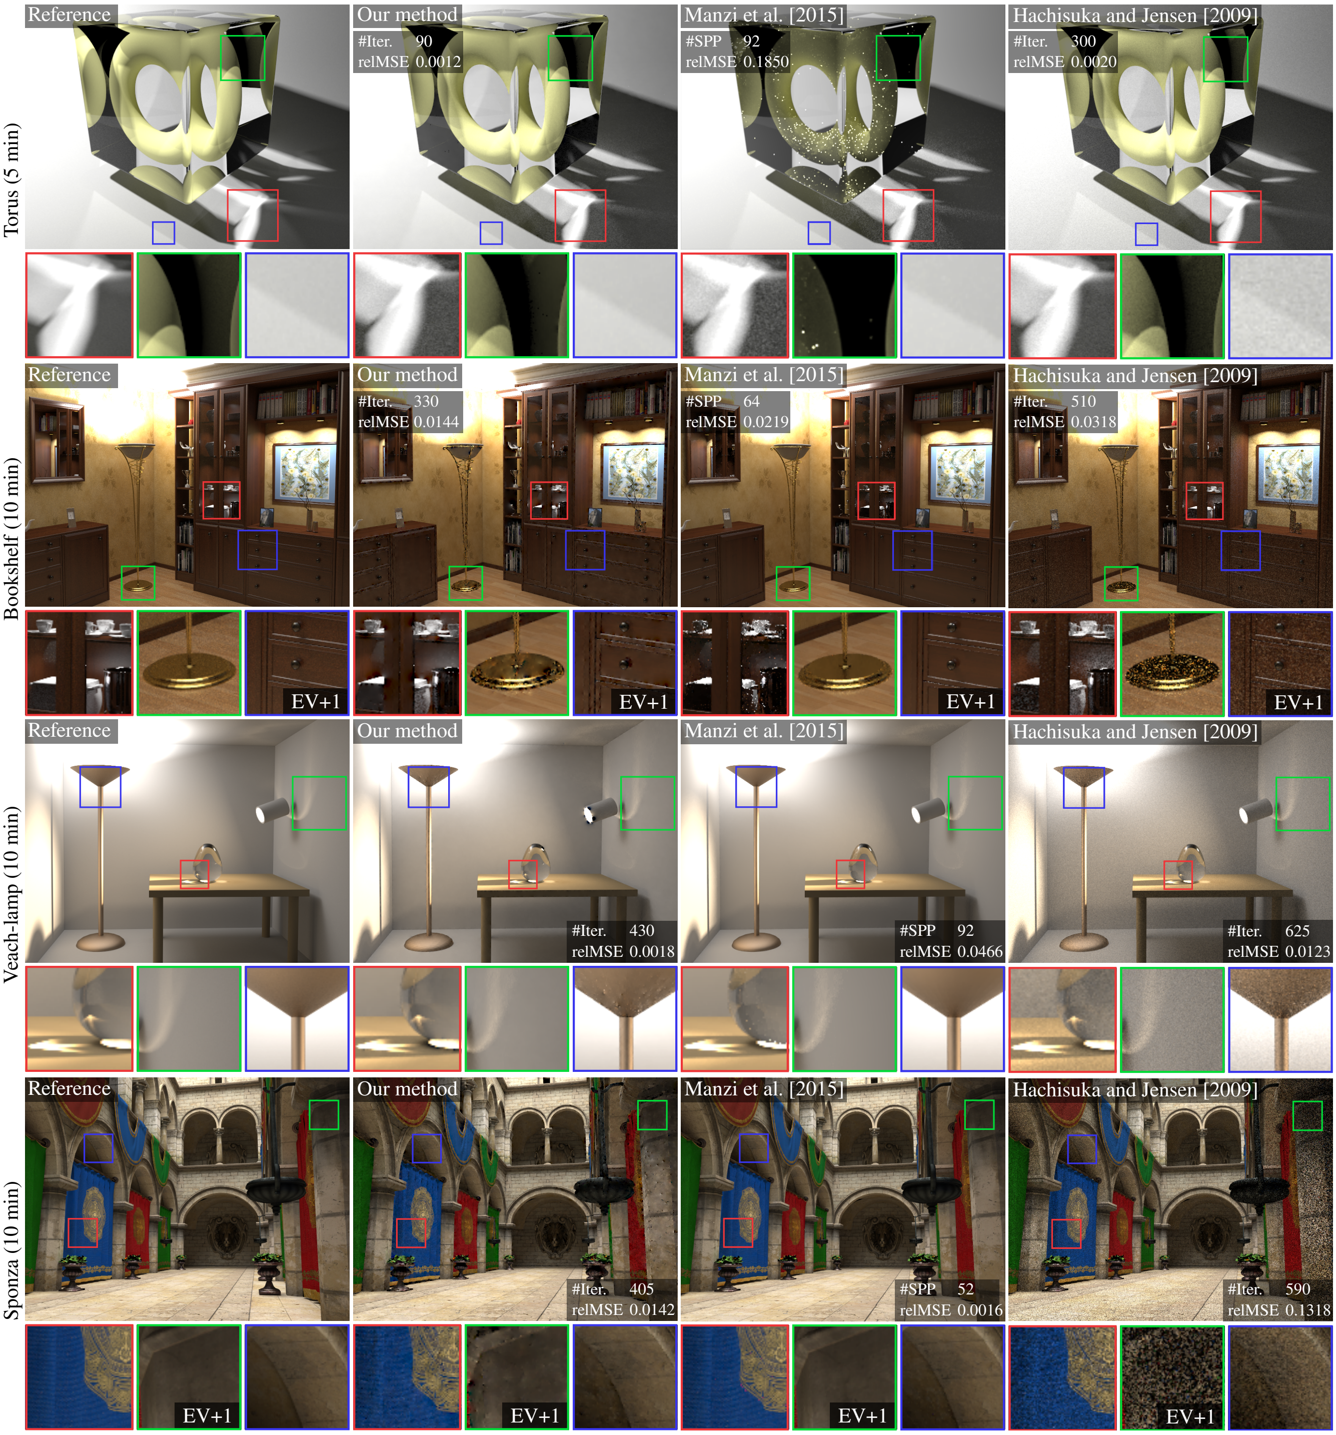

Figure 6: *Equal time comparison between our technique using L2 reconstruction, stochastic progression photon mapping [HJ09] and GBDPT [MKA*15] with L1 reconstruction.*

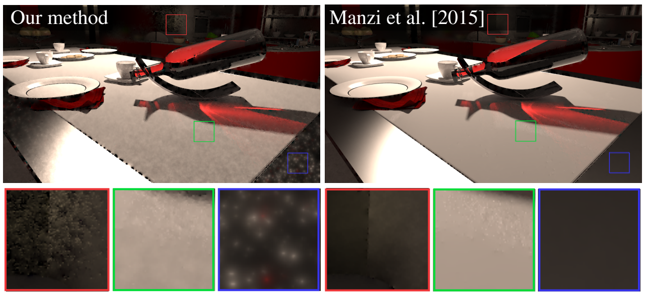

Figure 7: *Failure case: gradient-domain photon density estimation is less efficient at complex glossy regions (red zoom-in) or regions with very few photons visible to the camera (blue zoom-in). Gradient bi-directional path tracing works more effectively in these cases.*

manifold exploration spends all of its computational budget (i.e., maximum number of iterations) to move and project a path in the manifold, but still eventually fails. To limit the impact of this worst-case behavior, we set the maximum number of manifold exploration iterations to 5 (default in Mitsuba is 20), which led to about a 30% improvement in running time per photon pass without sacrificing the quality of the gradients. We use the same setting for GBDPT.

5.1. Limitations

Our method is not without limitations. Most obvious is that gradient-domain photon density estimation inherits the disadvantages of its primal-domain counterpart: essentially, traditional photon mapping tends to perform poorly when photons are gathered on a glossy surface or very few photons are visible from the camera. This adversely affects gradient-domain reconstruction (see Figure 7), and in these scenarios gradient-domain bi-directional path tracing is generally more effective.

6. Conclusions

We presented gradient-domain photon density estimation, a new rendering algorithm that computes image-space gradients using a photon density estimation framework. We show how ideas from existing gradient-domain rendering algorithms need to be adapted and extended in order to support gradient-domain density estimation: shifting photon density estimation paths and treating sensor and light subpaths separately before performing density estimation. We also show how to compute the correct Jacobians during these shifts. Our experiments demonstrate that this gradient-domain density estimation technique is more efficiently than its primal domain counterpart, and it also outperforms gradient-domain path tracing and gradient-domain bi-directional path tracing in scenes with complex transport due to e.g., SDS subpaths.

Several potential extensions of our work would be interesting to pursue. We perform density estimation only once to merge a vertex pair at the moment. It would be interesting to see how this can be extended towards a truly bi-directional photon mapping formulation. Another possible extension is to explore whether a shift mapping technique that does not rely on manifold exploration can still generate shifted paths that are robust to complex transport since manifold exploration is computationally expensive in general.

Acknowledgement We thank the authors of the scenes in our experiments: Benedikt Bitterli (Bathroom), Tiziano Portenier (Bookshelf and Bottle), Eric Veach, Miika Aittala, Samuli Laine, and Jaakko Lehtinen (Veach-Lamp), Frank Meinl (Sponza). We acknowledge the support of the SUTD Digital Manufacturing and Design (DManD) Centre which is supported by the National Research Foundation (NRF) of Singapore. This project is partially funded by JSPS KAKENHI Grant Number 15H05308.

References

[BCCZ08] Bhat P., Curless B., Cohen M., Zitnick C. L.: Fourier analysis of the 2d screened poisson equation for gradient domain problems. In *The European Conference on Computer Vision (ECCV)* (2008). 3

[GKDS12] Georgiev I., Křivánek J., Davidovič T., Slusallek P.: Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics (TOG) 31*, 6 (2012), 192. 2

[Hec90] Heckbert P. S.: Adaptive radiosity textures for bidirectional ray tracing. In *The 17th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)* (1990), ACM, pp. 145–154. 3

[HJ09] Hachisuka T., Jensen H. W.: Stochastic progressive photon mapping. *ACM Transactions on Graphics (TOG) 28*, 5 (2009), 141. 2, 5, 6, 7

[HOJ08] Hachisuka T., Ogaki S., Jensen H. W.: Progressive photon mapping. *ACM Transactions on Graphics (TOG) 27*, 5 (2008), 130. 2

[HPJ12] Hachisuka T., Pantaleoni J., Jensen H. W.: A path space extension for robust light transport simulation. *ACM Transactions on Graphics (TOG) 31*, 6 (2012), 191. 2, 3

[Jak10] Jakob W.: Mitsuba renderer, 2010. 6

[Jen01] Jensen H. W.: *Realistic image synthesis using photon mapping*, vol. 364. Ak Peters Natick, 2001. 2

[JM12] Jakob W., Marschner S.: Manifold exploration: a markov chain monte carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (TOG) 31*, 4 (2012), 58. 2, 3, 4, 5

[KHD14] Kaplanyan A. S., Hanika J., Dachsbacher C.: The natural-constraint representation of the path space for efficient light transport simulation. *ACM Transactions on Graphics (TOG) 33*, 4 (2014). 5

[KMA*15] Kettunen M., Manzi M., Aittala M., Lehtinen J., Durand F., Zwicker M.: Gradient-domain path tracing. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 123. 1, 2, 3, 4, 6

[LKL*13] Lehtinen J., Karras T., Laine S., Aittala M., Durand F., Aila T.: Gradient-domain metropolis light transport. *ACM Transactions on Graphics (TOG) 32*, 4 (2013), 95. 1, 2, 3, 5

[MKA*15] Manzi M., Kettunen M., Aittala M., Lehtinen J., Durand F., Zwicker M.: Gradient-domain bidirectional path tracing. *Proc. Eurographics Symposium on Rendering* (2015). 1, 2, 3, 6, 7

[MKD*16] Manzi M., Kettunen M., Durand F., Zwicker M., Lehtinen J.: Temporal gradient-domain path tracing. *ACM Transactions on Graphics (TOG)* (2016). 2

[MVZ16] Manzi M., Vicini D., Zwicker M.: Regularizing image reconstruction for gradient-domain rendering with feature patches. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 263–273. 2

[Vea97] Veach E.: *Robust monte carlo methods for light transport simulation*. PhD thesis, Stanford University, 1997. 1, 2