

Supplementary Materials: Visualization of Coherent Structures of Light Transport

Tobias Zirr¹, Marco Ament¹, and Carsten Dachsbacher¹

¹Karlsruhe Institute of Technology




Figure 1: *Omni-directionally integrated radiance (a), FTPD (b) and FTPD gradient (c) fields for a refracting glass sphere lit by a small area light. Light paths used for radiance integration and radiance weighting were limited to a length of 6 segments.*

1. Discussion of Details in the Results

This section provides additional details on some of the visualizations shown in the paper.

1.1. Refraction

The radiance structure inside the glass sphere (Figure 1) results from total internal reflection that allows light to reach points that are otherwise in shadow due to refraction perturbing light paths towards the center. The hard boundaries of these structures mark the narrow stripes where light travels around the sphere several times by total internal reflection, after it just managed to enter the sphere at extremely shallow angles. The stripes leading of the sides of the sphere are also caused by light traveling on such multiple total internal reflection paths before barely managing to exit the sphere. The rest of the radiance field conforms with what is intuitively expected of the caustics caused by glass spheres.

The FTPD field additionally puts its forward-directed structure on top of the radiance structure. In particular, we can see several separatrices on the back side of the sphere induced by the transmitted light scattered off the diffuse wall. The caustic in the center basically acts like another small area light and introduces separatrices where its light is split into light hitting the sphere and light missing the sphere. The entire diffuse wall acts like another bigger, but less intense area light that introduces similar separatrices. However, these are not tangential to the sphere as in this case, total internal reflection sets in even before paths diverge due

to hitting or missing the sphere, already causing a high degree of divergence before the tangent angle is reached.

1.2. 2D and 3D Example Scenes

Figures 2(b, e) depict the classic door crack Metropolis Light Transport [Vea98] example scene, where light emitted by a point light source in one room indirectly lights the other room through the crack of an almost closed door. The walls reflect the actual radiance. While the soft shadows do not reveal much, the FTPD field reveals structure in the light transport. As would be expected, the door crack introduces a separatrix that splits points with higher energy flow through the crack from points with lower energy flow through the crack.

The gradient visualization in Figure 2(e) reveals additional separatrices. One is introduced by the change in visibility where view of the other room is completely blocked by the door, co-linearly extending the door crack upwards. Two more very subtle ones mark the changes in influence of the diffuse indirect light on the dark room, scattered by different walls adjacent to the upper right corner on either side of the first-order discontinuities. In the dark room, corresponding separatrices can be found resulting from the actual soft shadow boundaries in the radiance field. Here, one more separatrix is introduced by the geometric discontinuity of the lower left corner, hit by the rather directed indirect light as focused by the small crack.

The canal examples in Figure 2(c) show two more FTPD

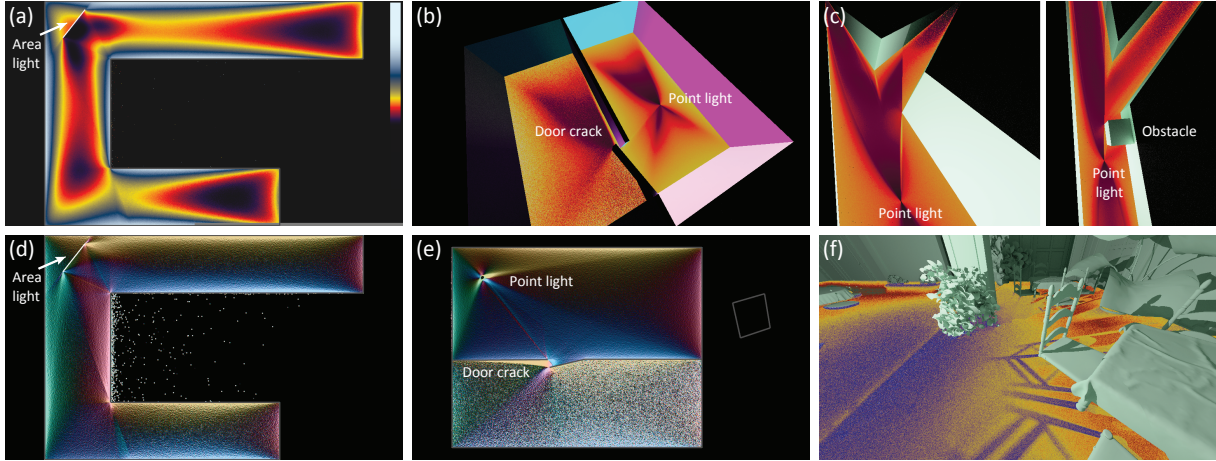


Figure 2: FTPD (a),(b),(c),(f) and FTPD gradient (d),(e) fields for more scenes, including 3D examples. (a),(b): U-shaped room lit by an area light. (b),(d): Slice through the 3D ‘Metropolis Light Transport [Vea98] Door Scene’ lit by a point light. The walls reflect the actual radiance. (c): FTPD for a slice of a 3D canal lit by a point light, unblocked (left) and blocked by an obstacle (right). (f): 3D scene [Lla10] where the slice does not contain the point light above the pot plant in front of the column.




Figure 3: FTPD on a slice of a 3D scene (red): If the light source is not contained in the slice plane (a), light rays hit different points for each point on the slice. This creates unintuitive projections (blue) of the scene rather than revealing separatrices in empty space. Moving the plane to contain the light (b), we can observe the separatrices known from our 2D analysis.

slices through 3D scenes. The FTPD field pretty much follows the intuition established by now, separatrices are introduced by discontinuities in geometry and visibility.

1.2.1. Discontinuity Analysis

Many separatrices in the FTPD fields in Figure 2 are rather subtle. Discontinuities in the gradient visualization separate coherent regions more clearly. This complies with our exemplary analytic consideration of light transport in Section 4.3 of the paper that demonstrates how separatrices are marked by first-order discontinuities. In the gradient visualization of the U-shaped room, the subtle light-independent discontinuities leading off the inner corners are caused by change in the visibility of the respective corridors. Another first-order discontinuity can be observed along the light-induced separatrix leading towards the lower inner corner. This separatrix follows the illustration in Figure 1(b) of the paper. The

remaining light-induced separatrices in the example of the U-shaped room follow the established intuition. Besides, the gradient field reflects the soft shadow of the area light below the lower inner corner.

1.2.2. 3D Slice Projection Issues

As briefly discussed in the paper and shown in Figure 2(f), inconsiderate slice placement can lead to unintuitive planar projections. Figure 3 illustrates how if the slices do not contain the light sources, light rays only penetrate the plane and hit geometry on the other side, which results in visible projections of this geometry onto the slice. The resulting visualizations reveal less to little information about the light transport, as they do not contain the linear structures known from the 2D case. In contrast, if the light source is contained by the slice, light rays through points of the slice are also contained in the slice plane, restoring the expected structures.

2. Measuring the Enhanced Path Space $\bar{\Omega}_*^t$

The PYPD moments and thus the FTPD are defined in terms of an integration over an *enhanced* set of paths. In this section we highlight the important properties and differences to the path space used for standard light transport simulation in computer graphics [Vea98]. Normally, the path space only contains complete paths that connect light sources to a sensor. In contrast, our enhanced path space $\bar{\Omega}_*^t$ contains paths that may start anywhere in space and have a fixed Euclidean run length of t .

2.1. Fixed Run Length

Since $\bar{\Omega}_*^t$ enforces a particular length, paths are not required to end on surfaces, but rather are cut off anywhere in space. This does not break the classic path-space area-product measure [Vea98], since it corresponds to ‘virtual’ spherical surfaces around the respective penultimate points. These ‘virtual’ surfaces are measured by the differential area corresponding to the endpoint of each path.

2.2. Arbitrary Starting Points

However, note that the arbitrary starting point breaks the classic path-space area-product measure $d\mu$, as it introduces a whole new unmeasured dimension that is orthogonal to surface patches, and thus breaks convergence of the integral. We must therefore use a modified measure $d\mu_*(\bar{x}) = d\mu(\bar{x})dl$ that includes a differential distance dl in-between parallel planes of starting points. Each of these planes then corresponds to a ‘virtual’ surface that can be measured by the differential area corresponding to the starting point in the classic area-product measure.

Thus, $\bar{\Omega}_*^t$ is actually a union of all path spaces where paths start on parallel planes \perp (point, normal):

$$\bar{\Omega}_*^t = \bigcup_{l=-\infty}^{+\infty} \bar{\Omega}_{\perp(l\omega, \omega)}^t. \quad (1)$$

3. Relationship between the FTPD and Light Transport

We have already used the path integral formulation, introduced by Veach [Vea98] for the problem of light transport, to define the FTPD in terms of a probability distribution of light transport paths. The product probability density $p_{\times}^{(\mathbf{x}, \omega)}(\bar{x}, \bar{y}) = p_{\mathbf{x}, \omega}(\bar{x}) p_{\mathbf{x}, \omega}(\bar{y})$ used in Eqn. (10) in the paper weights all pairwise path deviations according to their paths’ share in the light transport going through a given point \mathbf{x} ($\approx \mathbf{x}_1$), i.e., their contribution:

$$\begin{aligned} p_{\mathbf{x}, \omega}(\bar{x}) &\propto W_{-e}^{(\mathbf{x}, \omega)}(\mathbf{x}_1 \rightarrow \mathbf{x}_2) G(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2) \\ &\quad f_s(\mathbf{x}_1 \rightarrow \mathbf{x}_2 \rightarrow \mathbf{x}_3) G(\mathbf{x}_2 \leftrightarrow \mathbf{x}_3) \dots \\ &\quad \dots f_s(\mathbf{x}_{k-2} \rightarrow \mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) G(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_k). \end{aligned} \quad (2)$$

Since in our context, paths do not end on surfaces, we always define the differential patches at the end points to be orthogonal to the last segment in the corresponding geometry term. Note that the reversed emitted importance function W_{-e} emits importance *in the direction* of light propagation, i.e., reversed to what one is accustomed to in rendering.

The function $W_{-e}^{(\mathbf{x}, \omega)}$ is close to a direct delta centered at a given point \mathbf{x} and a given direction ω . As we intend to compute pairwise path deviations for pairs of slightly different paths, in practice $W_{-e}^{(\mathbf{x}, \omega)}$ needs to leave room for some small variation and may therefore not be singular.

Intuitively, rather than computing the radiance transported by paths as in the simulation of light transport, we compute the pairwise deviation transported by pairs of paths. Still, computation of the FTPD is closely related to the computation of radiance in that they share the same integral rendering framework. Particularly, both have the rendering equation at their core, imposing the same constraints on their results.

3.1. Mathematical Relationship

The remainder of this section mathematically relates the FTPD to light transport. These paragraphs provide a short summary of what is to follow.

3.1.1. Relating Influence to PYPD Moments

While the definitions of the influence-based topology analogy and the FTPD-based structural analysis appear to be quite different, Section 3.3 demonstrates that the influence I and the PYPD moments $m_{\Delta i}$ can be brought into a form where the only difference is the pairwise path deviation Δ . Based on that, a comparison of the directional change in influence $D_{\mathbf{h}}$ with the directional derivative of the PYPD moments reveals the similarity between the separatrix criteria of light transport and the FTPD field.

3.1.2. Interpretation of Coherent Structures as Continuous Contribution Transfer

Intuitively, we expect adjacent points to be coherent when the paths starting at their respective locations are similar in geometry and contribution. In particular, we expect the sets of these paths and their contribution functions to be identical in the limit as we move from one point infinitesimally close to another. Mathematically, this means that we expect contribution to be transferred *continuously* from the paths starting at one point to the paths starting at a coherent point.

Section 3.4 shows that the existence of such a mapping F that continuously maps paths at one point to paths at adjacent points while also maintaining continuity of their contribution is sufficient for these points to fall into the same coherent region with respect to some separatrix threshold τ .

3.2. Radiance Transfer Path Integral

We start by rewriting the radiance transfer function that lies at the heart of our light field analysis as a path integral. Introducing a (non-reversed) emitted importance factor $W_e^{(\mathbf{y}, \mathbf{v})}$, we can measure the flux incoming at point \mathbf{y} from direction \mathbf{v} using the measurement equation. From that, we can compute the corresponding radiance as follows:

$$I^{(\mathbf{y}, \mathbf{v})} = \int_{\tilde{\Omega}} L_o(\mathbf{x}_1, \overrightarrow{\mathbf{x}_1 \mathbf{x}_2}) \frac{p_{\mathbf{x}, \omega}(\bar{\mathbf{x}})}{W_{-e}^{(\mathbf{x}, \omega)}} W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}), \quad (3)$$

$$L_i(\mathbf{y}, \mathbf{v}) = \frac{d^2}{dA(\mathbf{y})d\mathbf{v}} I^{(\mathbf{y}, \mathbf{v})} = \lim_{W_e \rightarrow \delta} \int_{\tilde{\Omega}} L_o(\mathbf{x}_1, \overrightarrow{\mathbf{x}_1 \mathbf{x}_2}) \frac{p_{\mathbf{x}, \omega}(\bar{\mathbf{x}})}{W_{-e}^{(\mathbf{x}, \omega)}} W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}). \quad (4)$$

As $W_e^{(\mathbf{y}, \mathbf{v})}$ approaches the delta function centered around point \mathbf{y} and direction \mathbf{v} , the measured flux $I^{(\mathbf{y}, \mathbf{v})}$ per area and per solid angle incoming at point \mathbf{y} from direction \mathbf{v} approaches the corresponding incoming radiance L_i in the light field. Now, fixing source point \mathbf{x} and source direction ω by reintroducing $W_{-e}^{(\mathbf{x}, \omega)}$, differentiation yields the *incoming* radiance transfer function:

$$\frac{d^2}{dA(\mathbf{x})d\omega} L_i(\mathbf{y}, \mathbf{v}) = \lim_{W_{\pm e} \rightarrow \delta} \int_{\tilde{\Omega}} L_o(\mathbf{x}_1, \overrightarrow{\mathbf{x}_1 \mathbf{x}_2}) p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}) \quad (5)$$

$$= L_o(\mathbf{x}, \omega) \lim_{W_{\pm e} \rightarrow \delta} \int_{\tilde{\Omega}} p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}), \quad (6)$$

$$\frac{d^3}{dL_o(\mathbf{x}, \omega)dA(\mathbf{x})d\omega} L_i(\mathbf{y}, \mathbf{v}) = \frac{d}{dL_o(\mathbf{x}, \omega)} L_o(\mathbf{x}, \omega) \lim_{W_{\pm e} \rightarrow \delta} \int_{\tilde{\Omega}} p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}) \quad (7)$$

$$= \lim_{W_{\pm e} \rightarrow \delta} \int_{\tilde{\Omega}} p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_e^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}). \quad (8)$$

Note that in Section 3.3 of the paper, we defined the radiance transfer function in terms of outgoing radiance L_o rather than incoming radiance L_i . To account for that, we introduce an enhanced set of paths $\tilde{\Omega}^+$ that respectively appends end points on an infinitesimal sphere around each end point of the original paths in $\tilde{\Omega}$. Integrating over this enhanced set of paths, we can adjust the radiance transfer function to yield outgoing radiance:

$$T_o(\mathbf{x}, \omega \rightarrow \mathbf{y}, \mathbf{v}) = \frac{d^3}{dL_o(\mathbf{x}, \omega)dA(\mathbf{x})d\omega} L_o(\mathbf{y}, \mathbf{v}) = \lim_{W_{-e} \rightarrow \delta} \int_{\tilde{\Omega}^+} p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}). \quad (9)$$

3.3. Relating Influence to PWPD moments

Now that we have rewritten the radiance transfer function in terms of a path integral, we can directly compare the light transport formalism to the formalism of the FTPD. Writing the influence I in terms of the derived radiance transfer function, we get:

$$I(\mathbf{x} \rightarrow \mathbf{y}) = \int_{S^2} J(\mathbf{x}, \omega) \int_{\tilde{\Omega}^+} p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mu(\bar{\mathbf{x}}) d\mathbf{v} d\omega \quad (10)$$

$$= \int_{S^2} J(\mathbf{x}, \omega) \int_{S^2} \int_{\tilde{\Omega}^+} W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) p_{\mathbf{x}, \omega}(\bar{\mathbf{x}}) \left[\int_{\tilde{\Omega}} p_{\mathbf{x}, \omega}(\bar{\mathbf{y}}) d\mu(\bar{\mathbf{y}}) \right] d\mu(\bar{\mathbf{x}}) d\mathbf{v} d\omega \quad (11)$$

$$= \int_{S^2} J(\mathbf{x}, \omega) \int_{S^2} \int_{\tilde{\Omega}^+ \times \tilde{\Omega}^+} W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) p_{\times}^{(\mathbf{x}, \omega)}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mu_{\times}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mathbf{v} d\omega. \quad (12)$$

Note that the integral inserted in the second line simply evaluates to 1. We can rearrange the PWPD moment path integral to obtain a comparable integrand over \mathcal{I} :

$$m_{\Delta^i}^{[t_0, t_1]}(\mathbf{x}) = \int_{S^2} w(\mathbf{x}, \omega) m_{\Delta^i}^{[t_0, t_1]}(\mathbf{x}, \omega) d\omega = \int_{S^2} w(\mathbf{x}, \omega) \int_{[t_0, t_1] \times \tilde{\Omega}^+ \times \tilde{\Omega}^+} \Delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})^i p_{\times}^{(\mathbf{x}, \omega)}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mu_{\times}(t, \bar{\mathbf{x}}, \bar{\mathbf{y}}) d\omega \quad (13)$$

$$= \int_{\mathcal{I}} \int_{S^2} w(\mathbf{x}, \omega) \int_{[t_0, t_1] \times \tilde{\Omega}^+ \times \tilde{\Omega}^+} \Delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})^i \left[\int_{S^2} W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) d\mathbf{v} \right] p_{\times}^{(\mathbf{x}, \omega)}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mu_{\times}(t, \bar{\mathbf{x}}, \bar{\mathbf{y}}) d\omega dA(\mathbf{y}) \quad (14)$$

$$= \int_{\mathcal{I}} \int_{S^2} w(\mathbf{x}, \omega) \int_{S^2} \int_{[t_0, t_1] \times \tilde{\Omega}^+ \times \tilde{\Omega}^+} \Delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})^i W_{-e}^{(\mathbf{y}, \mathbf{v})}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k) p_{\times}^{(\mathbf{x}, \omega)}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mu_{\times}(t, \bar{\mathbf{x}}, \bar{\mathbf{y}}) d\mathbf{v} d\omega dA(\mathbf{y}) \quad (15)$$

$$=: \int_{\mathcal{I}} M_{\Delta^i}^{[t_0, t_1]}(\mathbf{x} \rightarrow \mathbf{y}) dA(\mathbf{y}). \quad (16)$$

Here, we locked the penultimate point of one path in each deviation pair to one interaction point \mathbf{y} and collapsed the integrand

into a fraction of the PWP moment $M_{\Delta^i}^{[t_0, t_1]}(\mathbf{x} \rightarrow \mathbf{y})$ specific to that point. Associating the paths in $\bar{\Omega}^+$ with the paths in $\bar{\Omega}_*^i$ by matching points $\mathbf{x}_1 \dots \mathbf{x}_{k-1}$ (all but the last) and associating the weight functions $J = w$, we can now see a direct correspondence between the definition of influence I and fractional PWP moment M_{Δ^i} . Thus, $M_{\Delta^i}(\mathbf{x} \rightarrow \mathbf{y})$ directly accounts for the influence $I(\mathbf{x} \rightarrow \mathbf{y})$ of point \mathbf{x} on point \mathbf{y} .

In contrast to I , M_{Δ^i} additionally convolves the pairwise path deviation Δ with path contribution. We can see why this is necessary by comparing the way the influence and the fractional PWP moment are integrated over interaction points. We defined the change in influence on interaction points as:

$$D_{\mathbf{h}}(\mathbf{x}) = \int_{\mathcal{I}} \left| \mathbf{h} \cdot \frac{\partial}{\partial \mathbf{x}} I(\mathbf{x} \rightarrow \mathbf{y}) \right| dA(\mathbf{y}). \quad (17)$$

In an actual computation using a progressive numerical method such as Monte Carlo integration, however, we have no way of differentiating the pointwise integrands that are likewise integrals that cannot be solved analytically. Therefore, we are restricted to differentiating the integral values that are the PWP moments:

$$|\nabla_{\mathbf{h}} m_{\Delta^i}^{[t_0, t_1]}(\mathbf{x})| = |\nabla_{\mathbf{h}} M_{\Delta^i}^{[t_0, t_1]}(\mathbf{x})| \quad (18)$$

$$= \left| \mathbf{h} \cdot \frac{\partial}{\partial \mathbf{x}} \int_{\mathcal{I}} M_{\Delta^i}^{[t_0, t_1]}(\mathbf{x}_1 \rightarrow \mathbf{y}) dA(\mathbf{y}) \right| \quad (19)$$

$$= \left| \int_{\mathcal{I}} \mathbf{h} \cdot \frac{\partial}{\partial \mathbf{x}} M_{\Delta^i}^{[t_0, t_1]}(\mathbf{x}_1 \rightarrow \mathbf{y}) dA(\mathbf{y}) \right|. \quad (20)$$

As we see now, the fundamental difference between the theoretical light field formalism and the FTPD formalism is the point of taking the absolute value. Whereas the former relies on the ability to differentiate the *integrand*, the latter is based on a differentiation of *integrals* — relying on the FTPD to ‘always’ yield different values for different scenarios while retaining the same properties with respect to continuity and curvature. It is this difference that makes the FTPD directly computable, whereas it would be impossible to compute the point-wise differences in influence without first settling on either a discretization over interaction points or a fixed-time approximation of point-to-point radiance transfer.

Note that while in practice, we differentiate the FTPD values $\sigma^{[t_0, t_1]}$ rather than the PWP moments $m_{\Delta^i}^{[t_0, t_1]}$, the FTPD values inherit the same relevant behavior as a composition of continuous elementary operations.

Coming back to the convolution with the PWP Δ in $M_{\Delta^i}^{[t_0, t_1]}$, we note that without Δ , integrating M_{Δ^i} over all interaction points would always evaluate to the integral of w . In this case, it is obvious that any FTPD derivative would be constantly 0.

3.4. Interpretation of Coherent Structures as Continuous Contribution Transfer

The function of $\Delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ in that convolution is to assign each path contribution a different weight. Thus, it creates a gradient between different paths, making changes in the contribution of different paths detectable in different changes of the FTPD, while continuous transfer of contribution between similar paths of adjacent points still translates to continuous change of the total FTPD value. In that way, the pairwise path deviation can be seen as a continuous hashing mechanism of path geometry, that in turn represents the constraints imposed by scene geometry. The resulting contribution-weighted fractional and total PWP moments can be seen as a continuous hash of the influence on distinct and all interaction points, respectively, representing all constraints imposed on radiance transfer by BRDFs, rendering equations, and geometry.

Intuitively, we expect adjacent points to be coherent when the paths starting at their respective locations are similar in geometry and contribution. In particular, we expect the sets of these paths and their contribution functions to be identical in the limit as we move one point infinitesimally close to another. Mathematically, this means that we expect contribution to be transferred *continuously* from the paths starting at one point to the paths starting at adjacent coherent points.

In order to analyze transfer of contribution between similar paths as a point of interest \mathbf{p} is moved, we now merge contributions and weights into a weighted contribution value $p_{\mathbf{p}, \mathbf{v}}^w(\bar{\mathbf{x}}) = \sqrt{w(\mathbf{p}, \mathbf{v})} p_{\mathbf{p}, \mathbf{v}}(\bar{\mathbf{x}})$. In terms of this contribution, the PWP moment becomes:

$$m_{\Delta^i}^{[t_0, t_1]}(\mathbf{x}) = \int_{\bar{\Omega}_*^{[t_0, t_1]}} \int_{\bar{\Omega}_*^{[t_0, t_1]}} \Delta(\bar{\mathbf{x}}, \bar{\mathbf{y}})^i \int_{\mathcal{S}^2} p_{\mathbf{x}, \mathbf{v}}^w(\bar{\mathbf{x}}) p_{\mathbf{x}, \mathbf{v}}^w(\bar{\mathbf{y}}) d\mathbf{v} d\mu_*(\bar{\mathbf{x}}) d\mu_*^{[t_0, t_1]}(\bar{\mathbf{y}}). \quad (21)$$

As we move the point of interest (POI) \mathbf{p} , we can describe continuous transfer of contribution between adjacent paths by a continuous mapping from paths starting at point \mathbf{p} to similarly contributing paths starting at the moved point $\mathbf{p} + h\omega$. We denote such a mapping by $F(h, \bar{\mathbf{x}})$. To formalize the idea of continuous contribution transfer, we require the mappings $F(h, *)$ to be injective




Figure 4: The mapping F C^1 -continuously maps paths \bar{x}_0 at $h = 0$ to paths \bar{x}_h such that paths associated by $F(*, \bar{x})$ have C^1 -continuous contribution $p_{\mathbf{p}, \mathbf{v}_h}^w$ in h . (a) The mapping F can handle continuous visibility changes by changing the path density. Continuity of the geometry term ensures that new paths to newly visible points occur with continuously changing contribution so long as BRDF and lighting conditions are continuous. The mapping F simply maps corner paths to new corner paths and spreads out the paths in-between. (b) Abrupt changes in visible area cannot be handled by F since the vanishing paths cannot be mapped to new paths while maintaining C^1 -continuity of path density and contribution. The same is true the other way round, in this case there is no set of paths that can be continuously mapped to the newly appearing paths. (c) Singular BRDFs are also handled by F , since reflection segments can be continuously mapped to adjusted reflection segments, so long as continuity on the remaining subpaths following the segment can be maintained. Thus, singular BRDFs propagate indirect visibility change.

and the relative path density $\frac{d\mu_*(F(h, \bar{x}))}{d\mu_*(\bar{x})}$ to be C^1 in h and continuous in \bar{x} almost everywhere. Moreover, we require the contribution $p_{\mathbf{p}+h\omega, F(h, \bar{x})}^w(F(h, \bar{x}))$ of h -associated paths to be C^1 in h for all $\mathbf{v} \in S^2$ and an injective $F(h, \mathbf{v})$ that is C^1 in \mathbf{v} a.e.

Figure 4 gives an intuition for the meaning of the mapping F . In fact, this mapping generically describes the geometric constraints that arise from BRDFs and scene geometry, such as imposed by specular reflection and refraction: For a perfect mirror reflecting a small cube (Figure 4,c), it would simply map ‘valid’ (contributing) reflection paths connecting the POI \mathbf{p} with the cube to valid reflection paths connecting the moved POI $\mathbf{p} + h\omega$ with the cube. As such, it is comparable to the mappings that algorithms such as Manifold Exploration [JM12] and the Half Vector Space Mutation [KHD14] try to predict based on local assumptions on BRDFs and geometry in order to find paths with similar contribution to previously constructed paths in scenarios constrained by specular interactions. However, we will not require knowledge of a concrete mapping. Rather, we will now show that the existence of such a mapping implies C^1 -continuity of the FTPD.

Given a set of paths $\bar{\Omega}_0$ and an interval $[0, \varepsilon]$ for which we can find such a continuous path contribution transfer mapping F to $\bar{\Omega}_h = F(h, \bar{\Omega}_0)$, we can see that:

$$m_{\Delta'}^{[t_0, t_1]}|_{\bar{\Omega}_0, F}(\mathbf{p} + h\omega) = \int_{F(h, \bar{\Omega}_0^{[t_0, t_1]})} \int_{F(h, \bar{\Omega}_0^{(\bar{y})})} \Delta(\bar{x}, \bar{y})^i \int_{S^2} p_{\mathbf{p}+h\omega, \mathbf{v}}^w(\bar{x}) p_{\mathbf{p}+h\omega, \mathbf{v}}^w(\bar{y}) d\mathbf{v} d\mu_*(\bar{x}) d\mu_*^{[t_0, t_1]}(\bar{y}) \quad (22)$$

$$= \int_{\bar{\Omega}_0^{[t_0, t_1]}} \int_{\bar{\Omega}_0^{(\bar{y})}} \Delta(F(h, \bar{x}), F(h, \bar{y}))^i \int_{S^2} p_{\mathbf{p}+h\omega, F(h, \mathbf{v})}^w(F(h, \bar{x})) p_{\mathbf{p}+h\omega, F(h, \mathbf{v})}^w(F(h, \bar{y})) \left| \frac{\partial F(h, \mathbf{v})}{\partial \mathbf{v}} \right| d\mathbf{v} \quad (23)$$

$$\frac{d\mu_*(F(h, \bar{x}))}{d\mu_*(\bar{x})} d\mu_*(\bar{x}) \frac{d\mu_*(F(h, \bar{y}))}{d\mu_*(\bar{y})} d\mu_*^{[t_0, t_1]}(\bar{y}).$$

Now assume the pairwise path deviation $\Delta(\bar{x}, \bar{y})$ to be C^1 -continuous in all points of both paths \bar{x}, \bar{y} as a C^1 -composition of C^1 -continuous operations. Then, deriving $m_{\Delta'}^{[t_0, t_1]}|_{\bar{\Omega}_0, F}(\mathbf{p} + h\omega)$ with respect to h clearly yields a continuous derivative (chain rule, composition of continuous functions and derivatives). In particular, if we find such a mapping F on $\bar{\Omega}_0 = \bar{\Omega}_h = \bar{\Omega}$, then all PWP momenta are C^1 -continuous functions on $[h = 0, \varepsilon]$ and it follows that the FTPD is also C^1 -continuous on that interval. Consequently, adjacent points of interest \mathbf{p} for which we can find such mappings over the entire path space for any differentiation direction ω are part of the same coherent region (for some threshold τ , we ignore second-order curvature).

In practice, our regularized pairwise path deviation $\Delta(\bar{x}, \bar{y})$ is C^1 -continuous *almost everywhere*, therefore after integration, the argumentation still holds.

4. Discussion of the FTPD

While we have now seen that the FTPD fulfills our expectations in many cases, there are a few caveats. Its hash-like function principally makes it prone to collisions, which may result in undetectably subtle separatrices or even loss of separatrices. In some of the examples, we have seen that such subtlety — as compared to the noise resulting from the high variance of FTPD fields, even more amplified by differentiation — can be a practical problem.

4.1. Run Length Variation

The question may arise why we sample pairwise path deflection for paths of a large interval of run lengths rather than of one fixed run length. Reformulation of the PWP moment integral in terms of interaction points showed that the PWP moments for paths of specific lengths more or less correspond to the influence on the interaction point that lies at their penultimate vertex. Therefore, fixed-length paths would insufficiently capture interactions with closer interaction points. In practice, this can be seen in that limiting sampling to very long paths, discontinuities fade out, whereas short paths naturally fail to incorporate the influence on farther interaction points.

4.2. Many Samples

As to sampling many pairs of paths per outgoing light direction, it is clear that these are necessary in order to capture non-singular scattering of light in all directions such as by diffuse surfaces.

4.3. Angle vs. Space Variation

Since our pairwise sampling is naturally related to partial differentiation with finite differences, the question may arise as to why we always perturb ray directions rather than ray origins in order to generate pairs, when we could perturb (“differentiate with respect to”) ray origins as well. For flow maps, Weinkauff and Theisel [WT10] found that differentiation with respect to all variables allows to extract streak lines and surfaces in the same way as stream and path lines, thus opening these up to the same kinds of analysis methods. However, so far, we did not find a corresponding utility of full differentiation in our analysis of light transport. There are several reasons why we chose directions:

For one, we can expect very similar results when perturbing ray origins rather than ray directions: The effect of perturbing directions and perturbing origins are basically the same, both result in pairwise hit points that are slightly apart. However, since we have to work with finite differences, perturbing origins near surfaces breaks meaningful results when they end up on different sides of scene geometry.

Moreover, in its provided intuitive definition, the pairwise

path deviation works better with respect to directional perturbation, since the expected distance always grows with run length. Thus, the ratio smoothly handles deviation from the original directions on interaction with curved surfaces. For origin perturbation, however, the expected distance is constant with respect to run length, easily resulting in large PWP values on the slightest divergence after the first interaction. Of course, the definition can be adapted to additionally divide by run length to get back to the original behavior. Still, we did not find an advantage in perturbing origins rather than directions warranting the additional error cases.

4.4. Alternative Operations

While it might seem interesting to look at the structures revealed by a continuous minimum and maximum of the directional FTPD over all outgoing light directions, this would require a very fine discretization of the sphere of outgoing light directions in order to first accumulate the directional FTPD by separate Monte Carlo integration for each direction and then apply the minimum/maximum. We have not tried that since it would likely lead to impracticable computation times. Additionally, we do not expect the resulting visualizations to be more informative. Changes would still be continuous, making structures similarly hard to detect in first-order discontinuities. Besides, such selective operators naturally limit the data at each point to correspond to one direction only, which might even reveal less structure. In contrast, the weighted integration of the FTPD is straightforward to compute using one single and fully progressive Monte Carlo integration per point of interest.

References

- [JM12] JAKOB W., MARSCHNER S.: Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics* 31, 4 (2012), 58:1–58:13. [6](#)
- [KHD14] KAPLANYAN A. S., HANIK A. J., DACHSBACHER C.: The natural-constraint representation of the path space for efficient light transport simulation. *ACM Transactions on Graphics* 33, 4 (2014), 102:1–102:13. [6](#)
- [LLa10] LLAGUNO G. M. L.: “San Miguel” scene. Published in *Physically Based Rendering, Second Edition*, 2010. [2](#)
- [Vea98] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1998. [1](#), [2](#), [3](#)
- [WT10] WEINKAUF T., THEISEL H.: Streak lines as tangent curves of a derived vector field. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1225–1234. [7](#)

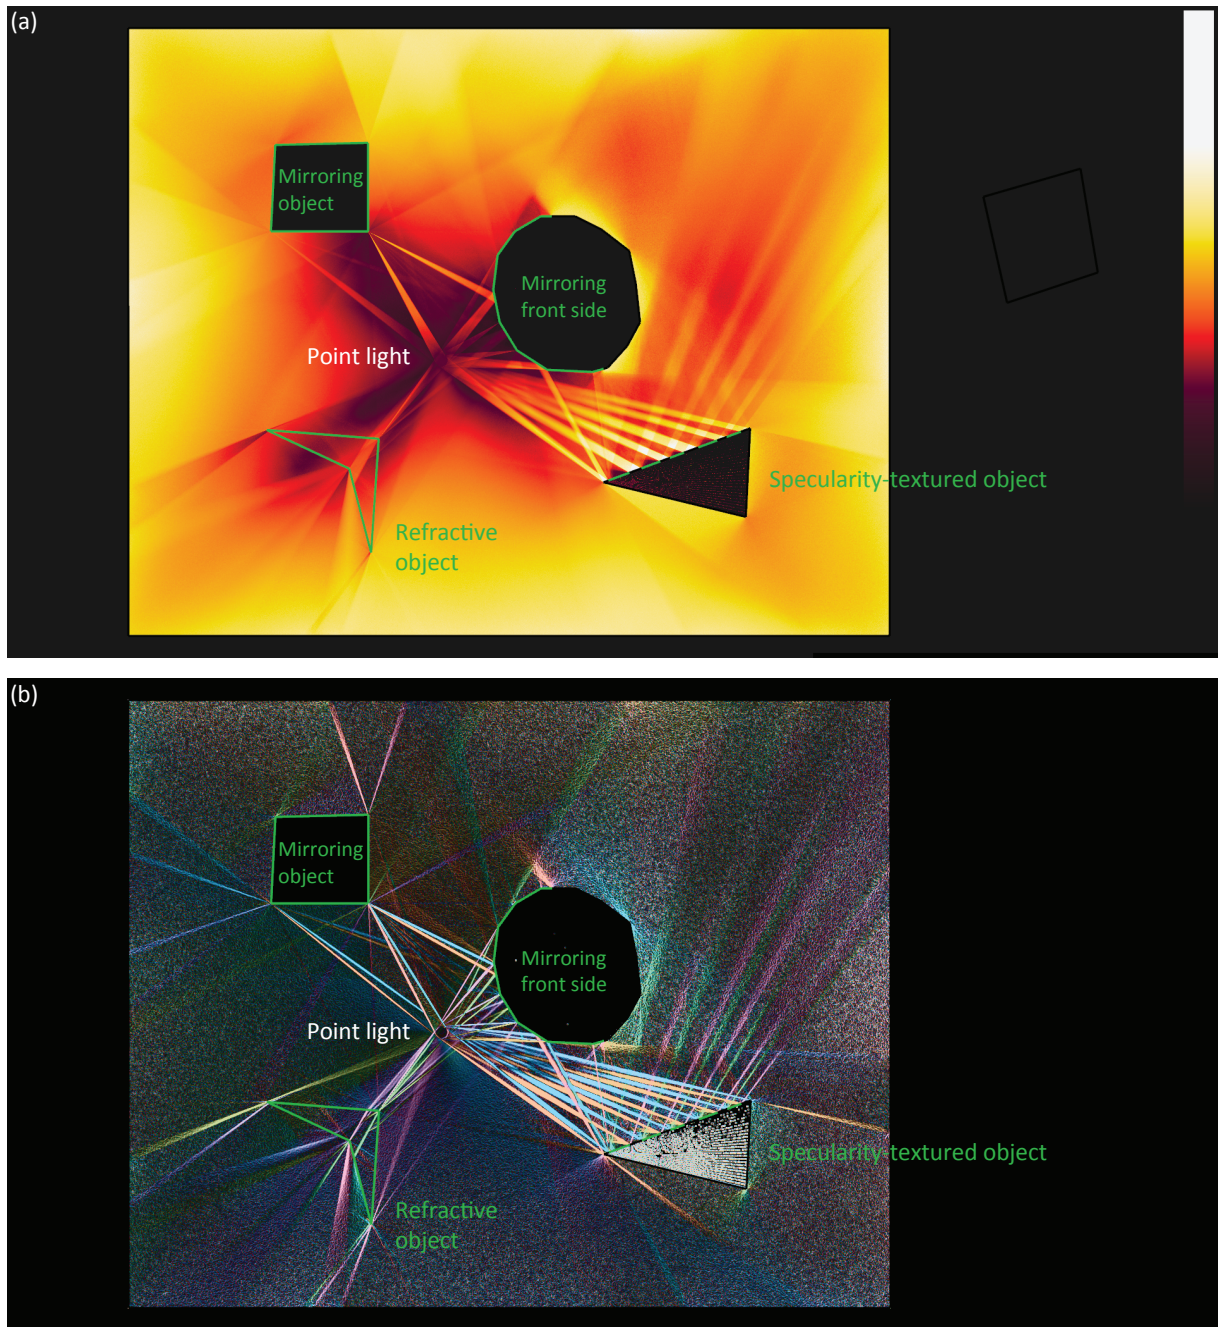


Figure 5: *FTPD (a) and FTPD gradient (b) fields for a moderately complex point-lit scene with several mirroring objects, a refractive object (lower left) and specularity-textured object (lower right).*