

# Visualization of Coherent Structures of Light Transport

Tobias Zirr<sup>1</sup>, Marco Ament<sup>1</sup>, and Carsten Dachsbacher<sup>1</sup>

<sup>1</sup>Karlsruhe Institute of Technology

---

## Abstract

*Inspired by vector field topology, an established tool for the extraction and identification of important features of flows and vector fields, we develop means for the analysis of the structure of light transport. For that, we derive an analogy to vector field topology that defines coherent structures in light transport. We also introduce Finite-Time Path Deflection (FTPD), a scalar quantity that represents the deflection characteristic of all light transport paths passing through a given point in space. For virtual scenes, the FTPD can be computed directly using path-space Monte Carlo integration. We visualize the FTPD field for several example scenes and discuss the revealed structures. Lastly, we show that the coherent regions visualized by the FTPD are closely related to the coherent regions in our new topologically-motivated analysis of light transport. FTPD visualizations are thus also visualizations of the structure of light transport.*

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications, I.6.6 [Computer Graphics]: Simulation And Modeling—Simulation Output Analysis

---

## 1. Introduction

The simulation of light transport has been subject to extensive research [RDGK12, KKG\*14]. While the main focus is on the development of efficient rendering algorithms, some more recent works visualize and analyze light transport to gain insight into the complex nature of light propagation. Previous approaches employed different domains for analyzing light transport, like ray space [Rus99, GFE\*12], frequency domain [DHS\*05], light field domain [RKR12], and path space [SNM\*13]. However, the high-dimensional nature of light transport makes it difficult to convey a meaningful visual representation on a two-dimensional screen.

Although high-dimensional data is often explored with information visualization techniques [TM04], light transport inherently takes place in the spatio-angular domain and it is difficult to build a mental picture of a scene from abstract visualizations. Therefore, we opt for a spatial visualization that can complement the geometry of the scene. However, a common problem of spatial visualizations is the lack of dimensions, which is also an important aspect for research in vector and tensor field visualization. A general approach to this issue is a reduction of dimensions, for example, to scalar fields, like the approach by Kindlmann et al. [KW99] for visualizing tensor fields with volume rendering or the technique by Haller [Hal01] for visualizing coherent structures of vector fields with the finite-time Lyapunov exponent (FTLE) field.

However, unlike vector and tensor fields, light transport not only depends on the position, but also on the direction, which poses additional challenges for visualization. Moreover, light transport usually exhibits strongly varying characteristics with respect to position and direction: it is constant along a ray in vacuum, it can be smooth under diffuse lighting on a surface, but it also contains discontinuities at shadow edges, mirrors, or refractive objects. Due to these rather unique properties, it is not feasible to use existing techniques from vector and tensor field visualization. Therefore, we present a novel method that is specifically tailored for visualizing light transport, but that is also inspired by well-known approaches to vector field visualization.

Our contributions are the introduction of a light transport analogy to vector field topology, as well as a novel technique for visualizing coherent structures of light transport. In our analogy, we discuss the applicability of topological aspects such as critical points and separatrices to light transport. Moreover, we introduce the *finite-time path deflection* (FTPD), a scalar-valued quantity that represents the deflection characteristic of all light transport paths passing through a given point in space. The FTPD may be computed using path-space Monte Carlo Integration, requiring only minor modifications of well-established algorithms of light transport [KKG\*14]. The absolute value of the FTPD hints at the amount of light scattering in the neighborhood of a given point. Discontinuous change in the FTPD hints

at changes in visibility, surrounding geometry, and material parameters. The FTPD may be computed with or without the weighting of actual incoming light, allowing for the analysis of both actual and potential light transport.

Color-coding the FTPD and its derivative yields visualizations of the structure of light transport in a given scene. The FTPD is computed from the variance of slightly diverging path samples, which makes it conceptually somewhat similar to the FTLE used for the extraction of Lagrangian coherent structures (LCS) in vector fields [Hal01]. Like for the FTLE field, ridges in the FTPD field trivially hint at repellers, such as walls in a scene, and its magnitude indicates the amount of light scattering. However, in contrast to the FTLE, many interesting phenomena are found in the FTPD's first and second derivatives. Since human perception is sensitive to discontinuities in color and first-order discontinuities in brightness these structures are visible in the visualizations as well.

## 2. Related Work

In this section, we discuss previous work related to the visual analysis of light transport and vector fields. The latter is important for our work, because it is a well-studied subject and serves as an inspiration in terms of methodology.

### 2.1. Visual Analysis of Light Transport

Global light transport in the absence of participating media is described by Kajiya's rendering equation [Kaj86] and the complex nature of light interaction has led to a large body of solution techniques. However, we focus on the visual analysis of light transport and we refer the reader to the works by Ritschel et al. [RDGK12] and Krivánek et al. [KKG\*14] for a detailed discussion on rendering algorithms.

A common approach to studying light transport is a frequency analysis of the plenoptic function. For example, Ramamoorthi et al. [RH04] showed that the reflection equation is a convolution with the BRDF in the angular domain. Subsequently, Ramamoorthi et al. [RMB07] demonstrated how the convolving nature of reflection affects gradients in the spatial domain. Durand et al. [DHS\*05] visualized and analyzed the frequency content of radiance and how it is altered by shading, occlusion, and transport, which can be exploited for adaptive sampling. Similar sampling optimizations, based on frequency analyses, were achieved by Egan et al. [EHDR11] and by Mehta et al. [MYRD14]. For Monte Carlo integration, Durand [Dur11] employed a frequency analysis to study the link between properties of the sampling pattern and the integrand. Furthermore, Ramamoorthi et al. [RAMN12] presented a Fourier approach to analyze and optimize the sampling of area light sources. However, although frequency analyses help understand certain phenomena of light transport or optimize algorithms, it is hard to visually gain insight in light transport.

In contrast, our approach visualizes changes in visibility,

surrounding geometry, and material parameters. Computing visibility is a major topic in computer graphics and we refer the reader to the surveys by Bittner and Wonka [BW03] and Cohen-Or et al. [COCSD03]; however, we study coherent structures in visibility. Durand et al. [DDP97] presented the visibility skeleton, a structural representation of coherent visibility in a scene. However, the focus on direct visibility limits its use for a holistic analysis of light transport. Discontinuities in the light field also play an important role for irradiance caches [SJJ12]. Chajdas et al. [CWW11] visualized the irradiance vector field, which indicates the dominant light direction. However, this strongly limits insight due to the high loss of information.

Visualizing light transport is an emerging topic in visualization and computer graphics, and can be inspired by real-world experiments [HFI\*08]. One of the most simple approaches is the visualization of light rays as geometric primitives [Rus99, GFE\*12]. However, insight is strongly limited, either due to an overly sparse representation or due to visual cluttering. On a higher level of abstraction, Reiner et al. [RKR12] presented a set of visualization techniques like spherical plots and particle flow tools for a selective inspection of light transport, which supports digital artists in understanding where certain illumination features originate. Schmidt et al. [SNM\*13] visualized light paths with edge bundling [HvW09] to reduce visual cluttering and to support artists in path re-targeting. Spencer et al. [SJJ14] employed parallel coordinates and the focus-plus-context paradigm [CMS99] to visualize photon distributions. However, all these methods focus on only local regions of a scene and require that a user manually explores the domain of light transport. In contrast, we provide a dense visualization technique that automatically highlights coherent structures.

### 2.2. Topology-based Visualization of Vector Fields

Our approach of visualizing coherent structures of light transport is inspired by topology-based flow visualization. However, since the definitions, methodology, and algorithms of flow visualization cannot be directly adopted for light transport, we focus only on the common basic principles on a conceptual level. For a more detailed discussion, we refer to the book chapter by Scheuermann and Tricoche [ST05] as well as the surveys by Lamee et al. [LHZP07] and Pobitzer et al. [PPF\*11]. The topological skeleton of a steady vector field consists of its critical points and their lower-dimensional invariant manifolds, the separatrices, partitioning regions of different flow behaviour [HH91]. In a similar spirit, we extract the so-called visibility skeleton of light transport that partitions regions of different direct and indirect visibility; however, therefore, we introduce our own definitions of interaction points and separatrices.

For time-dependent vector fields, LCS [HY00] can provide information on flow separation similar to separatrices for steady vector fields. In particular, Haller [Hal01] com-

puted the FTLE of velocity fields and identified the ridges of the FTLE field as LCS. Moreover, Shadden et al. [SLM05] showed that particles seeded near the FTLE ridges do not cross them, demonstrating the physical significance of LCS. In a similar spirit, we introduce a novel quantity, the FTPD, which is inspired by the idea of FTLE to measure flow separation. However, our definition of the FTPD is tailored for light transport, which is significantly different from fluid flow and FTLE. Nonetheless, the FTPD allows us to extract coherent structures of light transport in a similar manner to the ridge extraction of the FTLE field.

### 3. Background

Unlike flows and similar dynamical systems, traveling light cannot be described by a simple vector field. Unless interacting with matter, light travels freely along straight lines in space, through any point in any number of directions at the same time without interfering. Consequently, no single vector can fully describe all light passing through a given point.

#### 3.1. Plenoptic Function

The *plenoptic function*  $L(\mathbf{x}, \omega)$  describes the amount of light passing through every point in space in every direction. It thus fully describes the *light field*. In 3D space,  $L$  is a 5D function that defines the radiance propagated along a ray starting at a given point and extending in a given direction.

In order to conveniently define constraints resulting from interactions such as reflection, transmission or refraction, the plenoptic function is rarely used directly, but rather split into functions of incident radiance  $L_i$  and exitant radiance  $L_o$  [Vea98]. While incident radiance  $L_i(\mathbf{p}, \omega)$  describes the radiance arriving at  $\mathbf{p}$  from direction  $\omega$ , exitant radiance  $L_o(\mathbf{p}, \omega)$  describes the radiance leaving  $\mathbf{p}$  in direction  $\omega$ . In particular, the radiance along an unobstructed line segment in empty space remains constant:

$$L(\mathbf{p}, \omega) = L_o(\mathbf{p}, \omega) = L_i(\mathbf{p}, -\omega). \quad (1)$$

#### 3.2. Rendering Equation

In this paper, we limit ourselves to radiance transfer between interacting surfaces. The reflective behavior of opaque scene geometry is typically described by the rendering equation [Kaj86] that relates the outgoing radiance at surface points with the radiance incoming from above the surface:

$$L_o(\mathbf{x}, \omega) = \int_{\Omega(\mathbf{x})} f_r(\mathbf{x}, \mathbf{v} \rightarrow \omega) L_i(\mathbf{x}, \mathbf{v}) \langle \mathbf{n}(\mathbf{x}), \mathbf{v} \rangle + d\mathbf{v}. \quad (2)$$

Here,  $\Omega(\mathbf{x})$  denotes the unit hemisphere of directions, aligned with the surface normal  $\mathbf{n}(\mathbf{x})$ . The bidirectional reflectance distribution function (BRDF)  $f_r$  specifies the ratio between differential outgoing radiance  $L_o(\mathbf{x}, \omega)$  and differential incoming irradiance  $L_i(\mathbf{x}, \mathbf{v}) \langle \mathbf{n}(\mathbf{x}), \mathbf{v} \rangle + d\mathbf{v}$ .

While the equation in its presented form only accounts for

opaque surface reflection, similar equations can be defined to describe transmittance, refraction or even scattering of light in participating media.

#### 3.3. Radiance Transfer Function

Together, propagation equation (1) and rendering equation (2) make for a recursive definition of light transport. Expansion yields the infinitely-dimensional integral equations solved by light transport algorithms. In order to allow for a generic analysis, we introduce the global *radiance transfer function* that quantifies the amount of radiance exitant at each point  $\mathbf{y}$  in each direction  $\mathbf{v}$  when radiance passes through the differential cone  $(dA(\mathbf{x}), d\omega)$  at point  $\mathbf{x}$  in direction  $\omega$  (Fig. 1(a)). Thus, it fully describes all light transport:

$$T_o(\mathbf{x}, \omega \rightarrow \mathbf{y}, \mathbf{v}) = \frac{d^3}{dL_o(\mathbf{x}, \omega) dA(\mathbf{x}) d\omega} L_o(\mathbf{y}, \mathbf{v}). \quad (3)$$

The radiance transfer function  $T_o$  basically specifies the influence of light at one point on other points. It resolves all constraints imposed on light transport, allowing for an abstract analysis without reference of individual rendering equations. In 3D space,  $T_o$  is a 10-dimensional function.

### 4. A Vector Field Topology Analogy

Although light transport cannot be described by a vector field, we can borrow some aspects of vector field topology to build a terminology for its structure. This section defines such a structure in an abstract way that does not allow direct computation. In the next section we will develop computable means for its visualization.

#### 4.1. Interaction Points

In light transport, there is no clear match for critical points as defined in vector field topology. Still, there are points that change light transport and thus dictate its structure where the equation of unobstructed light propagation (1) is violated:

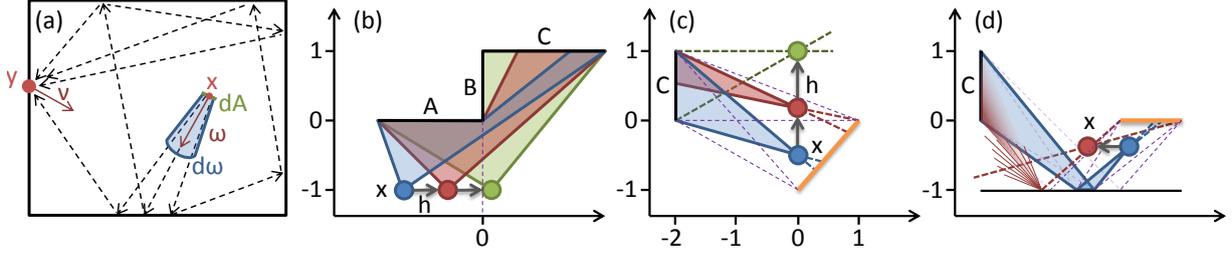
$$\mathcal{I} = \{ \mathbf{x} \mid \exists \omega: \frac{d}{dL_i(\mathbf{p}, -\omega)} L_o(\mathbf{p}, \omega) \neq 1 \}. \quad (4)$$

In the case of scene geometry,  $\mathcal{I}$  includes all surfaces. In general,  $\mathcal{I}$  contains all points of light interaction, henceforth called *interaction points*. In participating media,  $\mathcal{I}$  contains all points. However, light sources that do not interfere with light passing through, only adding to it, are not included.

**Light Sources** To include such non-interfering light sources, we can add the violation of Eqn. (1) in its non-differential form as a criterion for the enhanced set of interaction points  $\mathcal{I}_e$ :

$$\mathcal{I}_e = \mathcal{I} \cup \{ \mathbf{x} \mid \exists \omega: L_o(\mathbf{p}, \omega) \neq L_i(\mathbf{p}, -\omega) \}. \quad (5)$$

**Sensors** For convenience, we define sensors to absorb incoming light. Thus, sensors are also included in  $\mathcal{I}$  and the set of  $k$ -vertex light transport paths  $\Omega_k$  can be written as  $\mathcal{I}_e^k$ .



**Figure 1:** (a): The radiance transfer function  $T_o(\mathbf{x}, \omega \rightarrow \mathbf{y}, \nu)$  yields the radiance reflected at point  $\mathbf{y}$  in direction  $\nu$  per radiance passing through the differential cone ( $dA(\mathbf{x}), d\omega$ ) located at point  $\mathbf{x}$  with direction  $\omega$ . It respects all possible light paths. (b): Visibility separatrix, (c): light separatrix and (d): indirect light separatrixes, marked in purple. Area lights are marked orange. Influence is marked in the color of the corresponding point of interest. Case (d) depicts indirect light separatrixes caused by a mirroring ground plane. The reflection lobe illustrates the effect of a glossy ground instead of a mirror.

## 4.2. Coherent Structures

In light transport, it is common for light to be scattered in all directions, easily reaching any point in space after only a few interactions. It is therefore futile to partition the light field by separation of particle trajectories. To obtain meaningful regions, we consider two adjacent points *coherent* if their *influence* on almost all interaction points is similar. We group such coherent points into *coherent regions*, separated by the set  $\mathcal{S}$  of *separating points* forming *separatrixes*.

### 4.2.1. Influence

We define the influence  $I(\mathbf{x} \rightarrow \mathbf{y})$  of light passing through the *point of interest*  $\mathbf{x}$  on outgoing light at point  $\mathbf{y}$  as:

$$I(\mathbf{x} \rightarrow \mathbf{y}) = \int_{S^2} J(\mathbf{x}, \omega) \int_{S^2} T_o(\mathbf{x}, \omega \rightarrow \mathbf{y}, \nu) d\nu d\omega. \quad (6)$$

Essentially, this function computes all radiance transfer from point  $\mathbf{x}$  to point  $\mathbf{y}$ . The *influence weight function*  $J$  allows for a weighted consideration of radiance transfer depending on the outgoing light direction at  $\mathbf{x}$ . We will later set  $J(\mathbf{x}, \omega) = 1$  to analyze abstract light transport, whereas we will set  $J(\mathbf{x}, \omega) = L_o(\mathbf{x}, \omega)$  for concrete light transport.

In order to separate regions of coherent influence, we quantify the global change in influence at point  $\mathbf{x}$  along directions  $\mathbf{h}$ , by integrating the local change in influence on all interaction points with respect to  $\mathbf{h}$ :

$$D_{\mathbf{h}}(\mathbf{x}) := \int_{\mathcal{I}} \left| \mathbf{h} \cdot \frac{\partial}{\partial \mathbf{x}} I(\mathbf{x} \rightarrow \mathbf{y}) \right| dA(\mathbf{y}). \quad (7)$$

### 4.2.2. Separatrixes

Unfortunately, the absolute value of  $D_{\mathbf{h}}(\mathbf{x})$  is not a good criterion to locally identify separating points: Figure 1(b) shows how the influence on interaction points constantly changes as we move the point of interest, even in regions that we would intuitively call coherent. Instead, discontinuities in  $D_{\mathbf{h}}$  indicate large-scale change in influence (see Section 4.3). Consequently, we look for separatrixes where the

derivative of  $D_{\mathbf{h}}$  exceeds a given threshold  $\tau$ :

$$\mathcal{S}_{\tau} = \{ \mathbf{x} \mid \exists \omega: |\nabla_{\omega} D_{\omega}(\mathbf{x})| \geq \tau \}. \quad (8)$$

Discontinuities of  $D_{\mathbf{h}}$  are included in that set, since the magnitude of its derivative at these points is unbounded.

## 4.3. Examples

For brevity, we first consider only *direct* influence. We can compute the direct influence for point  $\mathbf{x}$  in Figure 1(b, c) as:

$$\begin{aligned} I(\mathbf{x} \rightarrow \mathbf{y}) &= \int_{S^2} J(\mathbf{x}, \vec{\mathbf{x}}\vec{\mathbf{y}}) V(\mathbf{x} \leftrightarrow \mathbf{y}) G(\mathbf{x} \leftrightarrow \mathbf{y}) f_r(\mathbf{y}, \vec{\mathbf{y}}\vec{\mathbf{x}} \rightarrow \nu) d\nu \\ &= J(\mathbf{x}, \vec{\mathbf{x}}\vec{\mathbf{y}}) V(\mathbf{x} \leftrightarrow \mathbf{y}) 2k_d \frac{\langle \mathbf{n}(\mathbf{y}), \vec{\mathbf{y}}\vec{\mathbf{x}} \rangle_+}{\|\mathbf{x} - \mathbf{y}\|^2} =: J V C_d. \end{aligned}$$

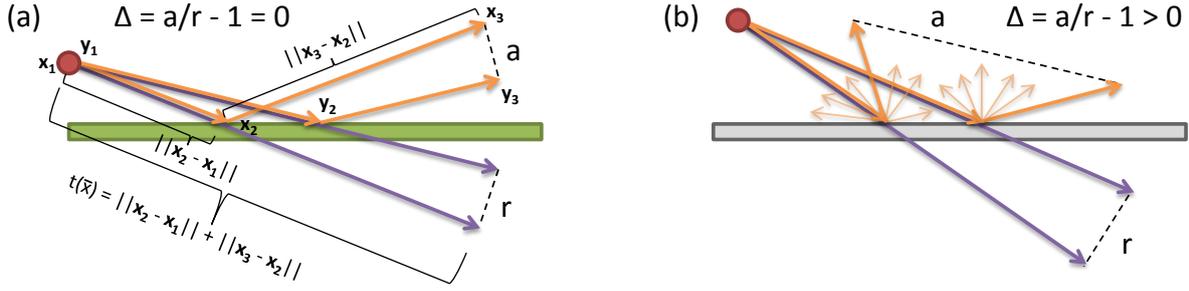
The arrow above vector pairs indicates unit vectors. Here,  $G$  is the geometry term,  $V$  denotes visibility, and  $f_r$  denotes the BRDF, which, for simplicity, we assume to be diffuse, i.e.  $\frac{k_d}{\pi}$ , collapsing continuous terms to  $C_d$ . Note that in Eq. (6), the inner integral  $\int_{S^2} T_o d\nu = V C_d$ .

### 4.3.1. Visibility Separatrixes

Consider the example shown in Figure 1(b). We set  $J = 1$  since there is no radiance weighting,  $V(\mathbf{x} \leftrightarrow (\mathbf{y} \in A)) = 1$ , and  $V((x_1, x_2) \leftrightarrow ((y_1, y_2) \in B \cup C)) = \sigma(y_1 + x_1)$ , where  $\sigma$  is the step function. For the change in influence, we get:

$$D_{\mathbf{h}}(\mathbf{x}) = \int_A |C'_d| dA(\mathbf{y}) + \int_{\substack{\mathbf{y} \in B \cup C, \\ y_1 \geq -x_1}} |C'_d| dA(\mathbf{y}) + \sum_{\substack{\mathbf{y} \in B \cup C, \\ y_1 = -x_1}} |C_d|,$$

with  $C'_d = h \cdot \frac{\partial}{\partial \mathbf{x}} C_d$ . The summation and the constrained integral result from the product rule yielding  $\delta$  and  $\sigma$ , respectively. For points  $\mathbf{x}$  on the left of the separatrix (marked blue and red), the sum is reduced to a single summand. Due to continuity of the BRDF and the geometry term as well as continuous expansion of the second integral domain,  $D_{\mathbf{h}}$  is continuous at these points. On the right of the separatrix (green point), the summation is empty and both integral domains are constant, ergo  $D_{\mathbf{h}}$  is also continuous. For points on



**Figure 2:** Computation of the pairwise path deviation  $\Delta$  from two paths  $\bar{x}, \bar{y}$  deflected by a mirror (a) and a diffuse surface (b)

the separatrix, however, the domain expanse for the second integral is discontinuous, abruptly adding  $C'_d$  integrated over the entire surface  $B$ . Since  $C'_d > 0$ ,  $D_h$  is discontinuous, indicating the separatrix. Note that while there are also infinitely many summands in this case,  $C_d = 0$  in all of these.

#### 4.3.2. Light Separatrices

In example (c), we set  $J(\mathbf{x}, \omega)$  to 1 if there is a ray in direction  $\omega$  from the light source through point  $\mathbf{x}$ , otherwise we set it to 0. For the depicted points, this corresponds to  $J(\mathbf{x} \rightarrow \mathbf{y}) = \sigma(y_2 - 3x_2)$ . Computing the change in influence, we get:

$$D_h(\mathbf{x}) = \int_{y_2 \geq 3x_2} |C'_d| dA(\mathbf{y}) + \left| C_d \right|_{y_2 = 3x_2}.$$

Clearly, the domain of the constrained integral shrinks continuously as the point of interest  $\mathbf{x}$  moves upwards. The second term, however, only exists for points in-between the two upper separatrices, causing discontinuities in  $D_h$  on both. These mark the additional light-induced separatrices.

#### 4.3.3. Fuzzy Separatrices

Figure 1(d) shows the indirect light-induced separatrices that would be detected if the ground plane was a perfect mirror. The example works in much the same way as example (c), in this case the masking  $\sigma$  terms are simply passed through the rendering equation by the Dirac  $\delta$  in the mirror BRDF.

In case of a glossy ground plane (as indicated by the reflection lobe for the red point), the masking  $\sigma$  terms undergo a convolution with the BRDF in the rendering equation. In this case,  $D_h$  is no longer discontinuous on separatrices: As indicated by the red gradient, points on the left side of the leftmost separatrix still influence points of  $C$ , there is no abrupt cut-off of a differential influence term in the change of influence. However, for a moderately glossy ground plane, the change in  $D_h$  would still be high near such ‘fuzzy’ separatrices. Their inclusion in  $\mathcal{S}_\tau$  depends on the threshold  $\tau$ .

### 5. Finite-time Path Deflection

In general, the separatrix criterion in Eqn. (8) cannot be computed directly: Since the integrand in Eqn. (7) depends on

exact solutions of the pointwise influence derivatives, progressive numerical methods such as Monte Carlo integration cannot be employed.

Motivated by the attempts at vector field topology based on the divergence of advected particles trajectories, such as the FTLE, we try to take a similar approach to the structural analysis of light transport. In light transport, however, there are no unique particle trajectories. We therefore need to take into account distributions of trajectories rather than comparing distinct trajectories.

In fact, when computing light transport numerically, we sample light paths, which are essentially distributions of particle trajectories. By associating each path with a probability that equals its relative contribution, we obtain a distribution of ‘trajectories’ that fully describes all light transport that led to a certain measurement.

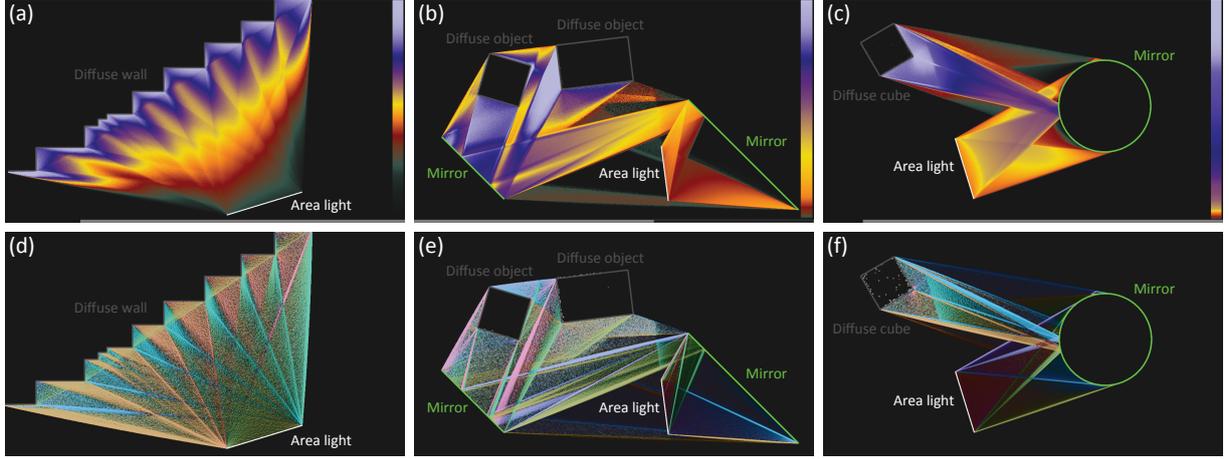
#### 5.1. Pairwise Path Deviation

To quantify the divergence of a pair of paths, we define the *pairwise path deviation* (PWPD)  $\Delta$  that characterizes the deviation of the end points  $\mathbf{x}_k, \mathbf{y}_l$  of two paths  $\bar{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k), \bar{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_l)$  of Euclidean run lengths  $t(\bar{x}) = \sum \|\mathbf{x}_{i+1} - \mathbf{x}_i\|, t(\bar{y})$  from what would be expected in empty space when starting at  $\mathbf{x}_1$  towards  $\mathbf{x}_2$  and at  $\mathbf{y}_1$  towards  $\mathbf{y}_2$ , respectively. Figure 2 illustrates the basic properties of  $\Delta$ :

$$\Delta(\bar{x}, \bar{y}) = \frac{\|\mathbf{x}_k - \mathbf{y}_l\|}{\|(\mathbf{x}_1 + t(\bar{x})\overrightarrow{\mathbf{x}_1\mathbf{x}_2}) - (\mathbf{y}_1 + t(\bar{y})\overrightarrow{\mathbf{y}_1\mathbf{y}_2})\|} - 1. \quad (9)$$

As it is impossible to come up with a meaningful and computable one-to-one mapping between ‘similar’ paths of two distributions to compute the PWPD for, we take the alternative approach of computing  $\Delta$  for *all* pairs of paths. Applying  $\Delta$  to a product distribution of path distributions yields a PWPD distribution that can be characterized using standard statistical measures such as mean value and variance.

In order to obtain well-defined results,  $\Delta$  needs to be regularized, as discussed in the supplementary materials to an (almost)  $C^1$ -function. There are several options for that. We opt to multiply by  $-2r^3 + 3r^2$  where  $r$  is proportional to the tangent of the half angle between the starting segments.



**Figure 3:** FTPD (a),(b),(c) and FTPD gradient (d),(e),(f) for three scenes. (a),(d): A diffuse ‘stair wall’ (gray) illuminated by an area light (white). (b),(e): An area light (white) that indirectly illuminates two diffuse objects (gray) through a two-piece system of mirrors (green). (c),(f): An area-lit mirroring sphere (green) that illuminates a diffuse cube (gray).

## 5.2. PYPD Moments

Based on the distributions of paths starting near a point  $\mathbf{x}$  and direction  $\omega$ , we define the  $i$ -th directional PYPD moment as:

$$m_{\Delta_i}^{[t_0, t_1]}(\mathbf{x}, \omega) = \int_{[t_0, t_1] \times \tilde{\Omega}'_* \times \tilde{\Omega}''_*} \Delta(\bar{x}, \bar{y})^i p_{\times}^{(\mathbf{x}, \omega)}(\bar{x}, \bar{y}) d\mu_{\times}(t, \bar{x}, \bar{y}). \quad (10)$$

The product probability density  $p_{\times}^{(\mathbf{x}, \omega)}(\bar{x}, \bar{y})$  equals the product of the relative contributions of paths  $\bar{x}, \bar{y}$ . The product measure  $d\mu_{\times}(t, \bar{x}, \bar{y}) = \frac{dt}{t_1 - t_0} d\mu(\bar{x}) d\mu(\bar{y})$  simply averages paths of all lengths  $t \in [t_0, t_1]$ . We denote the set of all paths of run length  $t$  starting at any point in space by  $\tilde{\Omega}'_*$ . In the supplementary material we discuss how  $d\mu(\bar{x})$  needs to be defined to properly measure the extended path space  $\tilde{\Omega}''_*$ .

## 5.3. FTPD Field

The *finite-time path deflection* (FTPD), which we will denote by  $\sigma^{[t_0, t_1]}(\mathbf{x})$ , is a scalar value that characterizes the amount of light scattering that occurs on all paths passing through point  $\mathbf{x}$  in all directions. It is computed from a weighted average of the directional PYPD moments:

$$m_{\Delta_i}^{[t_0, t_1]}(\mathbf{x}) = \int_{\Omega} w(\mathbf{x}, \omega) m_{\Delta_i}^{[t_0, t_1]}(\mathbf{x}, \omega) d\omega, \quad (11)$$

$$\sigma^{[t_0, t_1]}(\mathbf{x}) = \sqrt{m_{\Delta^2}^{[t_0, t_1]}(\mathbf{x})} + \sqrt{m_{\Delta^2}^{[t_0, t_1]}(\mathbf{x}) - m_{\Delta^1}^{[t_0, t_1]}(\mathbf{x})^2}. \quad (12)$$

Just like in the theoretical topology analogy, the weight function  $w$  allows for the incorporation of concrete lighting conditions by setting  $w(\mathbf{x}, \omega)$  proportional to the actual radiance as opposed to constantly setting it  $\frac{1}{2\pi}$ .

The FTPD  $\sigma^{[t_0, t_1]}(\mathbf{x})$  itself is defined as the sum of two ‘standard deviations’, the first characterizing the scattering of light in terms of deviation from unobstructed light prop-

agation, the second characterizing the variance of pairwise deviation values.

## 5.4. Coherent Structures

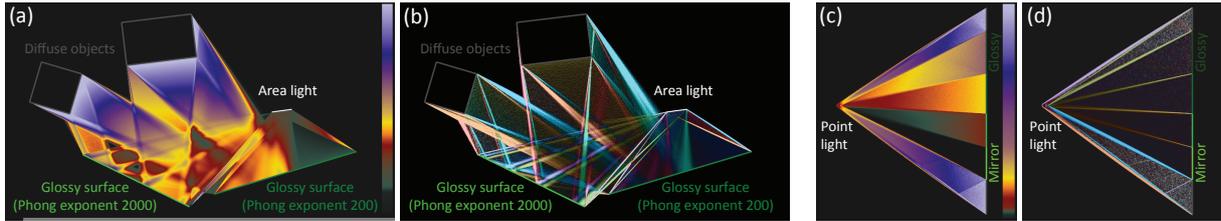
As we will shortly see in the visualizations, there is a direct relationship between the FTPD and the influence of passing light as discussed in the previous section. The mathematical relationship is also visible in our analogous definition of coherent structures (compare with Eqn. (8)), and further discussed in Sections 3.1-3.4 in the supplementary materials. The threshold-parameterized set of FTPD separatrices is:

$$\mathcal{S}_{\tau}^{\sigma^{[t_0, t_1]}} = \left\{ \mathbf{x} \mid \exists \omega: |\nabla_{\omega}^2 \sigma^{[t_0, t_1]}(\mathbf{x})| \geq \tau \right\}. \quad (13)$$

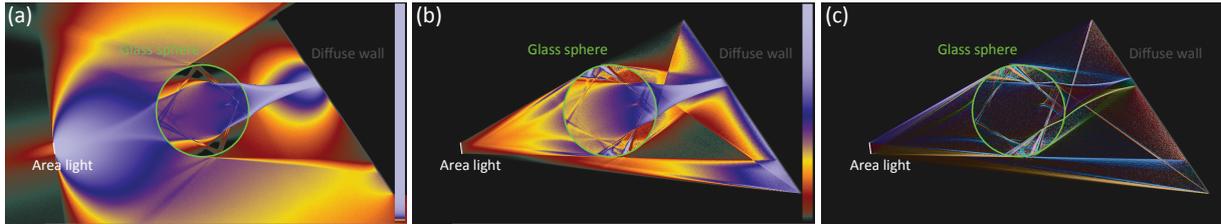
Here,  $\nabla_{\omega}^2$  denotes the second directional derivative. As for interaction points, there is not much to be gained from the FTPD, since they simply coincide with the scene geometry known as part of the simulation of light transport.

## 5.5. Visual Approach to Understanding the FTPD

This section provides interpretations of visualizations obtained from the radiance-weighted FTPD. We summarize the most important phenomena visible in the results, more detail is provided in the supplementary materials. Figure 3 shows the color-coded radiance-weighted FTPD and FTPD gradients computed for various diffuse scenes in ‘flatland’ (radiance transfer in 2D worlds). The FTPD field encodes the amount of deflection of the light transport paths passing through each point in the field. Unsurprisingly, the FTPD increases from the light source towards the geometry. The stair steps (left column) introduce discontinuities into the FTPD field where the topology analogy would find separatrices. While the gradient magnitude (luminance) is rather



**Figure 4:** FTPD (a) and FTPD gradient (b) fields for two glossy surfaces lit by a small area light, showing band-limited (fuzzy) separatrices caused by Phong-exponents 200 (right) and 2000 (left). FTPD (c) and FTPD gradient (d) fields for point-lit surface patches with different glossiness (green, descending from moderately glossy to a perfect mirror). The borders on either side of the glossy surfaces are diffuse.



**Figure 5:** 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.

noisy due to the sample variance of the FTPD, the texture that results from the color-coded gradient directions still reveals coherent regions.

**Propagation of Separatrices** Figure 3 also shows how separatrices propagate through mirrors. This permits to track the influence of light on the scene at every point of light transport paths. Even for the curved mirror, the indirectly lit cube clearly partitions the light field between the area light and the sphere. Also note that while not zero, the FTPD value is still low near curved reflection.

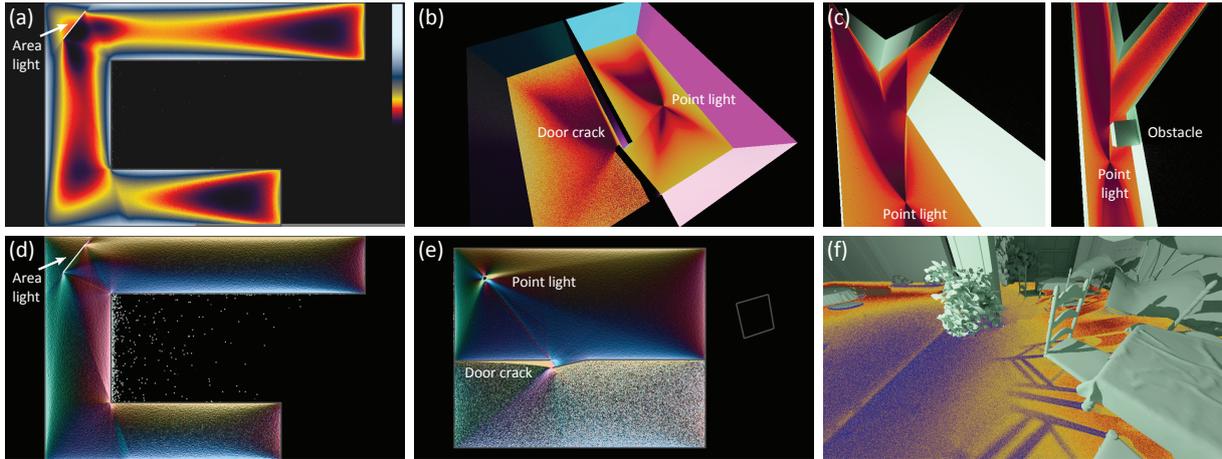
**Effects of Materials / BRDFs** Figures 4(a) and (b) show a similar setup where the perfect mirrors have been replaced by glossy surfaces (Phong-BRDF with exponents 200 and 2000). We can actually observe the fuzzy separatrices resulting from the band-limiting nature of glossy reflections as discussed in Section 4.3.3. Note that this affects separatrices on both sides of the reflection due to the glossy spreading of deflection pairs and due to a fuzzy radiance weighting caused by the filtered incoming light. Higher Phong-exponents result in sharper separatrices. Figures 4(c) and (d) compare the FTPD computed for patches with different glossiness values (Phong-BRDF with exponents 37, 112, 1029, and a perfect mirror). Higher Phong-exponents result in lower path deflection. This corresponds to the intuition of the FTPD being a variance-like value.

**Coherent Structures in Forward-Time** Figures 3 and 4 also illustrate how the FTPD can be interpreted as a kind of ‘reverse shadow’ that allows to foresee interactions with

objects on the routes of light transport paths in forward-time. We can observe how the resulting separatrices at the light source indicate on which side of the responsible objects or corners the light passing on either side of the separatrices will end up. In contrast, structure found in the light field as defined by the plenoptic function solely results from backward-time events. This gives further intuition on the meaning of coherent structures in light transport: They analyze coherency with respect to where light is headed and what geometry and materials it encounters on the way. The FTPD maps the response of the scene to light impulses going out from points in the field to a single scalar value.

**Refraction** Figure 5 shows the FTPD field for a glass sphere featuring curved refraction and reflection including the evaluation of Fresnel-terms. The complex light transport resulting from the many possible light paths result in many different separatrices caused by different phenomena. In this special case of highly structured light transport, the radiance field is already highly structured in itself. Note how the radiance-weighted FTPD retains the structure of the radiance field. The FTPD field additionally puts its forward-directed structure on top of the radiance structure. A detailed discussion is provided in the supplementary materials.

**2D and 3D Example Scenes** In contrast to the previous 2D-flatland examples, Figure 6 shows the FTPD computed for a planar slice through diffuse 3D scenes. Separatrices induced by light sources follow the rays leading away from the source and are best captured if the light source, and thus light rays emanating from the source, are contained in the



**Figure 6:** 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.

slice plane. Otherwise light rays only penetrate the plane (not forming distinct linear structures), and hit geometry on the other side which results in visible projections of this geometry onto the slice.

**Discontinuity Analysis** Many separatrices in the FTPD fields in Figure 6 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, affirming the analogous definition of FTPD separatrices. At this point we note that it should in fact be unexpected to find that many separatrices marked by zero-order discontinuities in the FTPD fields of previous examples. These result from the fact that in practice, we cannot spawn pairs of paths with infinitesimally close directions, causing discontinuities to occupy more than an infinitesimally small solid angle in the FTPD integral.

### 5.6. Unweighted FTPD

So far, we have only looked at radiance-weighted FTPD. Setting  $w = 1$ , the FTPD equally respects all potential light directions. Thus, there are no more light-induced separatrices, leaving the change of visibility as the only cause (mostly direct, indirect only near singular BRDFs). Figure 7 shows such an FTPD field with gradients. Without radiance weighting, the FTPD still increases towards surfaces and decreases in larger open spaces. The mirror and the prism open additional ‘virtual’ empty spaces via (total internal) reflection. We can even observe an indirect separatrix leading off the upper end of the mirror, caused by change in visibility of the back part of the central arrow head, blocked by the contained square on one side but influenced by light on the other side.

### 5.7. Importance-weighted FTPD

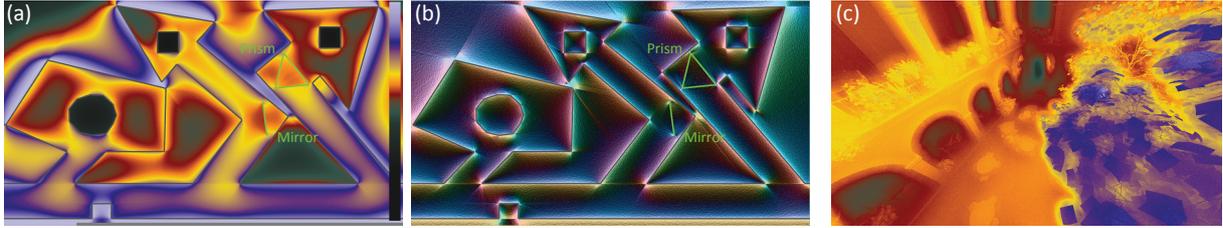
The FTPD field in Figure 7(c) was computed from deflection pairs started along camera rays. The FTPD naturally decreases as the distance to the first interaction increases. Its appearance resembles ambient occlusion in the details since it encodes not only the distance to the first hit point, but also the subsequent path geometry and thus the shape of the surrounding scene. Paths constrained by the coves amount to lower FTPD values than paths leading to equidistant points in the open. This reillustrates the FTPD’s meaning as a representative value of the scene’s response to light impulses and thus path geometry, scene geometry and materials.

## 6. Implementation

In this section, we describe how to compute the FTPD and how it is employed for visualization. For illustration purposes and to provide a proof of concept, we implemented a flatland version of FTPD with a simple path tracer and next event estimation using C++ and GLSL compute shaders. Here, we build on the theoretical investigation of 2D light transport by Jarosz et al. [JSKJ12]. Our 3D implementation is based on the OptiX ray tracing framework.

### 6.1. Computing the Finite-time Path Deflection

The FTPD is straightforward to compute using Monte Carlo integration (MCI), much in the same way as any light transport algorithm: For each point, we start by choosing a random Euclidean run length  $T \in [t_0, t_1]$  and a random direction for which to compute the directional PWP moments. For radiance-weighted FTPD, we then sample a light subpath in the opposite direction, trying to connect the current point with a light source. Independently of the weight function,



**Figure 7:** FTPD (a) and FTPD gradient (b) fields computed without any directional weighting by radiance/light sources. (c): Visualization of a 3D scene [Lla10] with ‘Importance-weighted’ FTPD, traced along primary rays.

we randomly perturb the first direction to obtain a pair of directions. We then sample two paths of run length  $T$  in one of the directions each, starting at the current point. From these two paths, we compute the pairwise path deviation.

We compute the first and second moments of the PWPD with MCI. Therefore we accumulate both the sum of the PWPD  $\Sigma\Delta$  and the sum of its squares  $\Sigma\Delta^2$ , each element multiplied by the relative contribution of both paths and the weight function. Generally, we cannot normalize with respect to weight and contribution in advance, therefore we also integrate the total of these cofactors  $\Sigma P$  to allow for normalization of the final FTPD values. Furthermore, for correct MCI, we divide all computed values by the respective probability densities. Finally, we compute the FTPD values:

$$\sigma_{[t_0, t_1]} = \sqrt{\frac{\Sigma\Delta^2}{\Sigma P}} + \sqrt{\frac{\Sigma\Delta^2}{\Sigma P} - \left(\frac{\Sigma\Delta}{\Sigma P}\right)^2}. \quad (14)$$

All our results were visualized with 100K to 500K samples and computation times ranging from one to five hours.

## 6.2. Visualization

In 3D, the FTPD is a volumetric scalar field that we visualize on a planar slice through the scene. For visualization, we use the CIELAB color space to minimize fake non-linearities introduced by perceptually non-linear interpolation. Since local changes are small compared to the high range of FTPD values, we evenly redistribute colors on the occurring values using a smooth histogram-based mapping. For the FTPD gradient visualizations, we apply the Sobel operator to the FTPD field, mapping 2D gradient direction to chromaticity  $a^*$  and  $b^*$ , and gradient magnitude to luminance  $L^*$ .

## 7. Conclusion

We have presented a technique to extract meaningful structures from global light transport in virtual scenes. For that, we introduced an analogy to vector field topology that replaces the analysis of deterministic trajectories by an analysis of distributed radiance transfer, thus allowing application to light transport. Since the separation criteria of the analogy cannot be computed directly, we introduced the FTPD as a quantity suitable for Monte Carlo integration. Based on the FTPD field, we defined analogous separation criteria and demonstrated that these reveal similar structures.

We discussed the meaning of the resulting coherent regions, demonstrating that the separatrices indicate non-local changes in the global influence on the scene. The coherent structures subsume what could be called the visibility skeleton, separating regions of continuous visibility *change*. However, the structures also reflect indirect visibility, as for example caused by specular interactions and refractions, even via moderately glossy interactions. In this regard, the FTPD field can assist in targeted light placement, the regions revealing the major options for different effect.

Analysis of concrete instances of light transport reveals regions where the passing light has a coherent effect on radiance transfer. Here, separatrices reveal splits in where the passing light is headed. In contrast to ordinary flux visualizations, our visualizations focus on light that is actually influencing surface points, and they complete post-interaction structures with pre-interaction structures, visualizing distinct groups of complete light paths up to the light source. As such, our visualizations can assist in occluder placement: Here, each region indicates a different effect of occlusion. We also visualize material properties. While more complex, future work could try to extract information on the effect of using specific materials in contexts such as interior design.

Besides these applications with utility in digital design and architecture, the structures visualized by our approach are closely related to correlated integrals in light transport (see supplementary Secs. 3.1.2, 3.4). Finding, understanding and exploiting such correlations is subject to active research.

Our approach can be limited by the inherently high dimensionality of light transport. As a consequence, the number of separatrices can become rather high even for moderately complex scenes. In addition, the dimensionality reduction causes separatrices to cross each other, which is different from the separatrices of vector fields. Therefore, it is usually more difficult to interpret our results compared to FTLE visualizations, especially for complex scenes. One possibility to reduce clutter is to limit the number of light bounces.

In future work, however, more sophisticated filtering techniques could be developed to extract only relevant structures — automatically or by means of explorative tools and focus-plus-context visualizations [CMS99]. We also plan to adapt importance sampling for the computation of the pairwise path deflection to decrease variance.

## References

- [BW03] BITTNER J., WONKA P.: Visibility in computer graphics. *Environment and Planning B: Planning and Design* 30, 5 (2003), 729–755. 2
- [CMS99] CARD S. K., MACKINLAY J. D., SHNEIDERMAN B. (Eds.): *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc., 1999. 2, 9
- [COCS03] COHEN-OR D., CHRYSANTHOU Y. L., SILVA C. T., DURAND F.: A survey of visibility for walkthrough applications. *IEEE Transactions on Visualization and Computer Graphics* 9, 3 (2003), 412–431. 2
- [CWW11] CHAJDAS M. G., WEIS A., WESTERMANN R.: Assisted environment probe placement. In *Proceedings of SIGRAD* (2011), pp. 1–9. 2
- [DDP97] DURAND F., DRETTAKIS G., PUECH C.: The visibility skeleton: A powerful and efficient multi-purpose global visibility tool. In *Proc. ACM SIGGRAPH* (1997), pp. 89–100. 2
- [DHS\*05] DURAND F., HOLZSCHUCH N., SOLER C., CHAN E., SILLION F. X.: A frequency analysis of light transport. *ACM Transactions on Graphics* 24, 3 (2005), 1115–1126. 1, 2
- [Dur11] DURAND F.: *A frequency analysis of Monte-Carlo and other numerical integration schemes*. Tech. Rep. MIT-CSAIL-TR-2011-052, MIT, 2011. 2
- [EHDR11] EGAN K., HECHT F., DURAND F., RAMAMOORTHY R.: Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Transactions on Graphics* 30, 2 (2011), 9:1–9:13. 2
- [GFE\*12] GRIBBLE C., FISHER J., EBY D., QUIGLEY E., LUDWIG G.: Ray tracing visualization toolkit. In *Proceedings of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2012), pp. 71–78. 1, 2
- [Hal01] HALLER G.: Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D* 149, 4 (2001), 248–277. 1, 2
- [HFI\*08] HULLIN M. B., FUCHS M., IHRKE I., AJDIN B., SEIDEL H.-P., LENSCH H.: Direct visualization of real-world light transport. In *Proceedings of the Workshop on Vision, Modeling, and Visualization* (2008), pp. 363–371. 2
- [HH91] HELMAN J., HESSELINK L.: Visualizing vector field topology in fluid flows. *IEEE Computer Graphics and Applications* 11, 3 (1991), 36–46. 2
- [HvW09] HOLTEN D., VAN WIJK J. J.: Force-directed edge bundling for graph visualization. *Computer Graphics Forum* 28, 3 (2009), 983–990. 2
- [HY00] HALLER G., YUAN G.: Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D* 147, 3–4 (2000), 352–370. 2
- [JSKJ12] JAROSZ W., SCHÖNEFELD V., KOBBELT L., JENSEN H. W.: Theory, analysis and applications of 2d global illumination. *ACM Transactions on Graphics* 31, 5 (2012), 125:1–125:21. 8
- [Kaj86] KAJIYA J. T.: The rendering equation. *Computer Graphics (Proc. SIGGRAPH)* 20, 4 (1986), 143–150. 2, 3
- [KKG\*14] KŘIVÁNEK J., KELLER A., GEORGIEV I., KAPLANYAN A., FAJARDO M., MEYER M., NAHMAS J.-D., KARLÍK O., CANADA J.: Recent advances in light transport simulation: Theory and practice. In *ACM SIGGRAPH Courses* (2014). 1, 2
- [KW99] KINDLMANN G., WEINSTEIN D.: Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *Proc. of IEEE Visualization* (1999), pp. 183–189. 1
- [LHZP07] LARAMEE R. S., HAUSER H., ZHAO L., POST F. H.: Topology-based flow visualization, the state of the art. In *Topology-based Methods in Visualization*, Hauser H., Hagen H., Theisel H., (Eds.). Springer Berlin Heidelberg, 2007, pp. 1–19. 2
- [Lla10] LLAGUNO G. M. L.: “San Miguel” scene. Published in *Physically Based Rendering, Second Edition*, 2010. 8, 9
- [MYRD14] MEHTA S. U., YAO J., RAMAMOORTHY R., DURAND F.: Factored axis-aligned filtering for rendering multiple distribution effects. *ACM Transactions on Graphics* 33, 4 (2014), 57:1–57:12. 2
- [PPF\*11] POBITZER A., PEIKERT R., FUCHS R., SCHINDLER B., KUHN A., THEISEL H., MATKOVIĆ K., HAUSER H.: The state of the art in topology-based visualization of unsteady flow. *Computer Graphics Forum* 30, 6 (2011), 1789–1811. 2
- [RAMN12] RAMAMOORTHY R., ANDERSON J., MEYER M., NOWROUZEZAHRAI D.: A theory of Monte Carlo visibility sampling. *ACM Transactions on Graphics* 31, 5 (2012), 121:1–121:16. 2
- [RDGK12] RITSCHEL T., DACHSBACHER C., GROSCH T., KAUTZ J.: The state of the art in interactive global illumination. *Computer Graphics Forum* 31, 1 (2012), 160–188. 1, 2
- [RH04] RAMAMOORTHY R., HANRAHAN P.: A signal-processing framework for reflection. *ACM Transactions on Graphics* 23, 4 (2004), 1004–1042. 2
- [RKRD12] REINER T., KAPLANYAN A., REINHARD M., DACHSBACHER C.: Selective inspection and interactive visualization of light transport in virtual scenes. *Computer Graphics Forum* 31, 2 (2012), 711–718. 1, 2
- [RMB07] RAMAMOORTHY R., MAHAJAN D., BELHUMEUR P.: A first-order analysis of lighting, shading, and shadows. *ACM Transactions on Graphics* 26, 1 (2007), 2:1–2:21. 2
- [Rus99] RUSSEL J. A.: *An interactive web-based ray tracing visualization tool*. Master’s thesis, University of Washington, 1999. 1, 2
- [SJJ12] SCHWARZHAUPT J., JENSEN H. W., JAROSZ W.: Practical Hessian-based error control for irradiance caching. *ACM Transactions on Graphics* 31, 6 (2012), 193:1–193:10. 2
- [SJJ14] SPENCER B., JONES M. W., LIM I. S.: A visualization tool used to develop new photon mapping techniques. *Computer Graphics Forum* 34, 1 (2014), 127–140. 2
- [SLM05] SHADDEN S. C., LEKIEN F., MARSDEN J. E.: Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena* 212, 3–4 (2005), 271–304. 3
- [SNM\*13] SCHMIDT T.-W., NOVAK J., MENG J., KAPLANYAN A. S., REINER T., NOWROUZEZAHRAI D., DACHSBACHER C.: Path-space manipulation of physically-based light transport. *ACM Transactions on Graphics* 32, 4 (2013), 129:1–129:11. 1, 2
- [ST05] SCHEUERMANN G., TRICOCHE X.: Topological methods for flow visualization. In *Visualization Handbook*, Johnson C. R., Hansen C. D., (Eds.). Academic Press, Inc., 2005, pp. 341–356. 2
- [TM04] TORY M., MÖLLER T.: Rethinking visualization: A high-level taxonomy. In *Proceedings of the IEEE Symposium on Information Visualization* (2004), pp. 151–158. 1
- [Vea98] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1998. 3, 8