The global structure of the visual light field and its relation to the physical light field

Kartashova, Tatiana; Sekulovski, D; de Ridder, Huib; te Pas, SF; Pont, Sylvia

DOI
10.1167/16.10.9

Publication date
2016

Document Version
Final published version

Published in
Journal of vision

Citation (APA)
Kartashova, T., Sekulovski, D., de Ridder, H., te Pas, SF., & Pont, S. (2016). The global structure of the visual light field and its relation to the physical light field. *Journal of vision*, *16*(10), 1-16. https://doi.org/10.1167/16.10.9

The global structure of the visual light field and its relation to the physical light field

Tatiana Kartashova Perceptual Intelligence Laboratory, Delft University of Technology, Delft, the Netherlands

Dragan Sekulovski Visual Experiences Group, Philips Research Europe, Eindhoven, the Netherlands

Huib de Ridder Perceptual Intelligence Laboratory, Delft University of Technology, Delft, the Netherlands

Susan F. te Pas Experimental Psychology, Utrecht University, Utrecht, the Netherlands

Sylvia C. Pont Perceptual Intelligence Laboratory, Delft University of Technology, Delft, the Netherlands

Human observers have been demonstrated to be sensitive to the local (physical) light field, or more precisely, to the primary direction, intensity, and diffuseness of the light at a point in a space. In the present study we focused on the question of whether it is possible to reconstruct the global visual light field, based on observers' inferences of the local light properties. Observers adjusted the illumination on a probe in order to visually fit it in three diversely lit scenes. For each scene they made 36 settings on a regular grid. The global structure of the first order properties of the light field could then indeed be reconstructed by interpolation of light vectors coefficients representing the local settings. We demonstrate that the resulting visual light fields (individual and averaged) can be visualized and we show how they can be compared to physical measurements in the same scenes. Our findings suggest that human observers have a robust impression of the light field that is simplified with respect to the physical light field. In particular, the subtle spatial variations of the physical light fields are largely neglected and the visual light fields were more similar to simple diverging fields than to the actual physical light fields.

Introduction

It is fascinating how light manipulations can change the appearance of an object or a scene (Cuttle, 1973; Ganslandt & Hofmann, 1992; Hunter, Biver, & Fuqua, 2007). Many professionals—artists, photographers, light designers, and architects—use practical knowledge about light in their work. They put things in a spotlight to put attention on them, they use diffuse lighting to make surfaces look smoother, or play more sophisticated tricks with our visual system to create illusions. So far, the theory behind such lighting techniques is not particularly extensive (Cuttle, 2003; Gilchrist & Radonjic, 2009). Moreover, we are not aware of any lighting design books that address how complicated optical interactions between lighting, spatial geometry, materials, and objects result in the light distribution or light field in a space. Most literature in this field focuses on providing enough light (intensity) on working surfaces, people, and objects (Boyce, 1981). Only a few lighting designers address the issue that human-centered lighting design should be based on the observer's experience of light arriving at the eye, which is determined not only by the light sources, but also by a scene geometry, furnishing objects, and materials as well as the position and motion of the observer and the people in the scene (Cuttle, 2003; Ganslandt & Hofmann, 1992). These

Citation: Kartashova, T., Sekulovski, D., de Ridder, H., te Pas, S. F., & Pont, S. C. (2016). The global structure of the visual light field and its relation to the physical light field. *Journal of Vision*, *16*(10):9, 1–16, doi:10.1167/16.10.9.

doi: 10.1167/16.10.9 Received November 5, 2015; published August 18, 2016 ISSN 1534-7362

lighting designers provide some ideas on how to progress toward working with the light field, but none were able to describe the global structure of light distributions in three-dimensional spaces. Thus, a deeper understanding of visual qualities of light might help to improve multiple fields of knowledge.

How an image is perceived depends on the light in the image's scene but at the same time, the light can be judged through the objects in that scene (Koenderink, Pont, van Doorn, Kappers, & Todd, 2007). There are many studies on perceptual interrelations between light and object shape (Berbaum, Bever, & Chung, 1983; Koenderink & van Doorn, 2006; O'Shea, Agrawala, & Banks, 2010), light and surfaces properties (Doerschner, Boyaci, & Maloney, 2007; Fleming, Dror, & Adelson, 2003; Ho, Landy, & Maloney, 2006; Marlow, Kim, & Anderson, 2012), and light and spatial geometry (Madsen, 2007; Yamauchi, Ikeda, & Shinoda, 2003). However, there are only few studies with the light field as the main focus in physics (Dror, Willsky, & Adelson, 2004; Gershun, 1939; Moon & Spencer, 1981; Mury, Pont, & Koenderink, 2007, 2009) and in visual perception (Adelson & Bergen, 1991; Gerhard & Maloney, 2010; Koenderink et al., 2007; Maloney, Gerhard, Boyaci, & Doerschner, 2010; Morgenstern, Geisler, & Murray, 2014; Pentland, 1982; Pont, 2013; Pont & Koenderink, 2007; Pont, van Doorn, de Ridder, & Koenderink, 2010; Xia, Pont, & Heynderickx, 2014). Of those studies, we list the most relevant findings, on which we based our study of the reconstruction of the spatial structure of a perceived light field and its relation to its physical counterpart.

Gershun (1939) was one of the pioneers working on the definition of light in a space. He introduced the term *light field*, which is a function that describes the amount of light traveling in every direction through every point. It describes light as a function of position in space and direction resulting in a five-dimensional spherical function, describing the radiance arriving to a point x, y, z from direction θ, φ. One may interpret this description as a huge collection of panoramic images for all positions in a space. Gershun (1939) introduced *radiant flux density* — the net flux that passes through any given surface element from either side; and the *light vector* — the magnitude and direction of the net flux. He also introduced *light tubes* as a manner to describe the net light transport through a space. Adelson and Bergen (1991) described an equivalent of the physical light field, the *plenoptic* function (from *plenus*, meaning "all," plus optic), which specifies the structure of light as a function of position, direction, wavelength, and time. In our work, we ignore the variables wavelength and time and focus on the luminance distribution — as in Gershun's (1939) definition. This concerns a five-dimensional function, which is quite complex for most real scenes.

The development of mathematical and computer modeling tools provided the means for further elaboration of the topic. Mury et al. (2007, 2009) studied the spatial structure of physical light fields. They developed methods to describe the physical light field by extending Gershun's (1939) definition of light's components, as well as methods to measure and visualize its (global) structure. They found experimentally that the low-order components of the light field vary smoothly over scenes, implying that it is possible to make a reconstruction of the physical light field of a scene based on a relatively small number of measurement points. Mury et al. (2007, 2009) also demonstrated a strong relationship between the low-order components and the geometrical layouts of the scenes, concluding that, in some way, the physical light field can be thought of as a property of the geometry.

The visual light field was measured by Koenderink et al. (2007). In order to test human sensitivity to various parameters of the physical light field, they introduced a white sphere as a *gauge object*. Their setup consisted of stereoscopic photographs of a scene under three light conditions and a probe (a white matte sphere). The task was to set the lighting (light direction, intensity, diffuseness) on the probe so that it appeared to belong to the scene. The resulting probe images were compared to the photographs of a real sphere in the same positions. Koenderink et al. (2007) concluded that human observers have expectations of how a given object would appear if it was introduced in a scene at some arbitrary position and named this awareness the *visual light field*.

Xia et al. (2014) tested the findings of Koenderink et al. (2007) for real scenes. They created a real setup in which the lighting on a scene and on a probe could be manipulated separately, and the scene and probe could be fused using a semitransparent mirror. Twenty participants were asked to judge whether the probe fitted the scene with regard to illumination intensity, direction, and diffuseness. The observers were found to be sensitive to variations in light intensity, direction, and diffuseness.

In reviewing studies on illumination, lightness and brightness perception, Schirillo (2013) argued the necessity "to develop a broader perceptual theory, including the crucial variable of how light is inferred in open space" (p. 905). Throughout the paper, he provided several examples of experiments in which the same test patch was perceived significantly lighter or darker depending on its apparent position in space — under dim or bright illumination, respectively. Schirillo (2013) discussed the possibility that we infer light itself, raising questions such as "why does the object appear colored, but not the space in front of it, in that this space contains light of the same wavelengths as that at the surface of the object" (p. 908). In the end, he draws

Figure 1. Light conditions: visible light source (LAMP), diffused light sources in the ceiling on the right side of the scene (DIFFUSE), two collimated light sources in the ceiling, one on the left and the other on the right side of the scene (SPOTLIGHTS).

the conclusion that there exists a mental representation of the light in three-dimensional space. He notes that such awareness must be derived from reflecting surfaces in a scene, but that the quality is not brightness. He concluded that the awareness of light extending into the space between the surfaces and the eye is phenomenologically real and measurable.

Schirillo (2013) also noted that our awareness of the light in empty space does not necessarily mirror physical illumination. For example, neutral light in a room furnished in red may be perceived as reddish. Ostrovsky, Cavanagh, and Sinha (2005) found that the structure of the physical light field can also be misperceived. They demonstrated that humans often neglect inconsistencies in illumination, which would not be the case if the light field is perceived in a physically veridical way. This "deficiency" was explained by the visual system not trying to verify the global consistency of the local estimates because, they speculated, human evolution in a single light source environment yields a fast local analysis sufficient for obtaining the illumination information needed. Likewise, in Koenderink et al.'s (2007) experiment there was a condition in which the judgments of all but one observer were inconsistent with the physical truth. It was the condition in which a probe was in the volume of a cast shadow of an object, implying that observers were not aware of certain details of the light field's structure.

Van Doorn, Koenderink, Todd, and Wagemans (2012) recently studied the perception of several global structures in natural light fields. They showed that human observers can perceive uniform, diverging, and converging light fields (which can be represented by sunlight, candlelight, and a ring or sphere of light surrounding a scene or object of interest, respectively) but are insensitive to rotational and deformation light flow patterns (Cuttle, 1973), which are less common and may be formed by complicated lighting/geometry.

In conclusion, there is growing evidence that although the visual and physical light fields show similarities, they are not identical to each other. However, we are not aware of any studies in which the global structure of the visual light field was actually measured and compared to the physical light field. One reason for this gap certainly concerns the lack of a proper method to measure the visual light field.

The aim of the present study is to fill this gap in measuring and comparing the global structures of corresponding physical and visual light fields. In doing so, an important question is whether measuring many instances of a local visual light field in a scene allows the reconstruction of its global structure by interpolating its lower order features in a similar way as Mury et al. (2009) did for the physical light field. We found that our data was sufficiently smooth to accomplish this.

The measurements were done on a real scene, a living room, in which the type of illumination was varied to create three different light fields. We took measurements over nearly regular grids, sufficiently capturing the variations in the light field over the three illumination types. In this study, we focused mostly on the directional component of the light field, the light vector, but we also measured the ambient component in order to make diffuseness comparisons. The paper is organized as follows: First, we present the methods we used to measure and reconstruct the visual light fields, demonstrate approaches for their visualization, and discuss the measurements results. Second, we explain the physical measurements and their processing. Finally, we compare the visual and physical light fields and discuss the overall results of the study.

The visual light field

Methods

Scene

The scene was part of a laboratory room, furnished as a common living room (see Figure 1). The scene width was 340 cm, the depth (from the back wall to the front-most line of measurements) was 255 cm, and the height was 300 cm. Most of the objects in the scene had matte surfaces varying in roughness. Some surfaces were shiny, including a smooth specular tabletop, a metal fruit plate and lamp stand, and a transparent glass flower vase. We used the following light conditions: a lamp on the table (LAMP); two sets of fluorescent lamps, one behind another, on the right side of the ceiling creating a diffusely lit scene (DIFFUSE); and two spotlights, one on the left and the other on the right side of the room with their center of symmetry slightly to the left of the room center, to create a more focused lighted scene (SPOTLIGHTS).

We photographed the scene under each of the three light conditions. The camera was standing 5 m away from the back wall of the scene with the following camera settings: *f*-number equal to 7.1 and exposure times equal to 1/2, 1/15, and 1/8 s for LAMP, DIFFUSE, and SPOTLIGHTS, respectively, because of the differences between illumination levels. We then converted the photographs to grayscale, with each picture ranging in pixel brightness from 0 to 255, and cropped them to hide a part of the ceiling. Thus, the images were made so that the light source was visible for the first scene and invisible for the other two scenes. The resulting pictures can be seen in Figure 1.

Probe

To measure the visual light field for each light condition, we used a grid of 36 positions in the photographs in three depth layers (see Figure 2A). At each position, a computer-generated rendering of a probe was superposed on the photographed image to assess the perceived light in that position. The probe, similar to the one introduced by Koenderink et al. (2007), is a white matte sphere on a black monopod (the "pole") superimposed on predetermined locations in the image. Observers could control both the direction and the intensity of a collimated light beam on the sphere and the intensity of the ambient light. The direction was controlled by mouse movements and the intensities by keyboard buttons. In this experiment, the diffuseness was defined as the ratio of collimated and ambient luminances (Xia, Pont, & Heynderickx, 2016b), ranging from fully collimated light (e.g., spotlight) to fully diffuse or Ganzfeld illumination (e.g., as in mist or a snowy field on a cloudy day). The diffuseness could be controlled by adjusting the ratio between the collimated and ambient intensities instead of by an explicit extra parameter, as in Koenderink et al. (2007). This adjustment simplified the interface and is based on studies into diffuseness characterization (Cuttle, 2003; Xia et al., 2016a, 2016b, 2016c, 2016d).

The poles indicated the probe position in the scene. This method can be safely used to replace the stereoscopic representation of the probe location, because in Koenderink et al. (2007) the stereoscopic representation served solely for defining the probe location. Moreover,

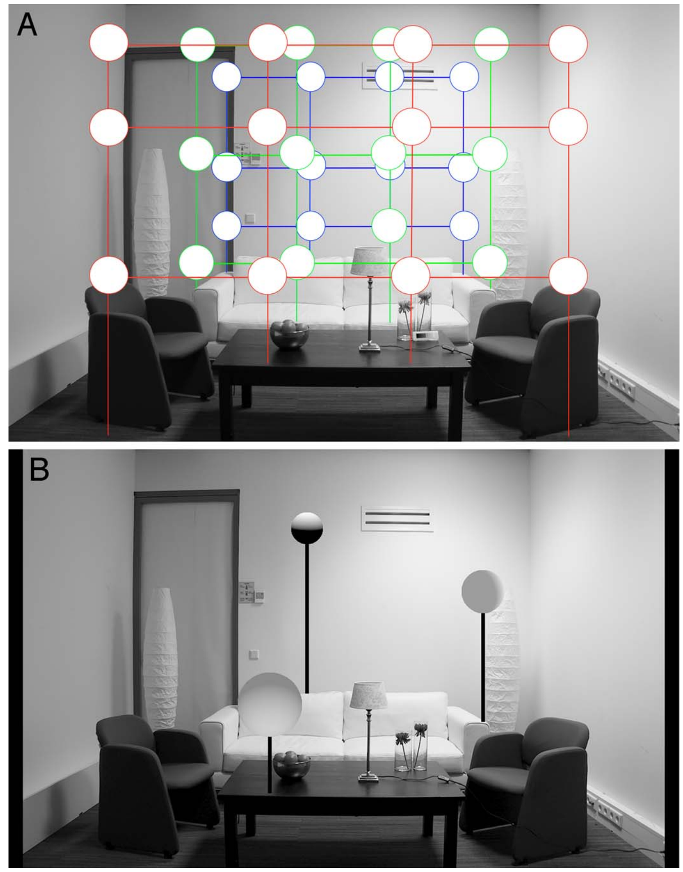

Figure 2. (A) All probe positions in the scene. Lines of the same color connect the positions at the same depth and the vertical lines denote the ends of the poles. (B) Examples of the probes for each depth. The shading on the probes demonstrates a few possible variations of settings. These are not observers' results.

on Pont's (2011) poster, the methods of Wijntjes and Pont (2010) were used to test whether the representation of the probe's position using a pole resulted in robust spatial percepts. The relative depth structure was found to be robust up to condition- and observer-dependent depth range scaling. The poles always had to be entirely visible to the observer, which restricted the available space for probe placement (e.g., one probe on the sofa was moved up in order to not occlude the lamp; see Figure 2A). We estimated the probe coordinates in the space of the scene (a) in depth and width by placing the pole end on objects with known positions; (b) in height, by relating to a reference object (photographs of the cubic light meter; see Methods of the section on the physical light field measurements). The sphere size was scaled according to the perspective in the photographs depending on its position (see Figure 2B).

Setup

The experiment was performed on a high-resolution 15-in. computer screen (2880 × 1800 pixels, Retina

Figure 3. Subjective light source(s) position(s), where red dots represent single sources, and color dots represent pairs of sources. The actual positions of the light sources are shown using dotted lines. For (A) LAMP, all observers pointed at the lamp; for (B) DIFFUSE, the points were more spread out; and for (C) SPOTLIGHTS, most observers drew single sources, red dots and a transparent line for a single elongated source. The source next to the table lamp "is standing in front of the scene." Four observers drew a pair of sources.

Display, luminance range from 0.4 cd/m^2 to 330.8 cd/m^2), with software developed using the Psychtoolbox library (Brainard, 1997). The light in the room was switched off. The images and the screen were calibrated linearly; in other words, the monitor luminances in the image were linearly related to the original luminances in the scene. The viewing distance of the observer was fixed with a chin rest at 27 cm from the monitor in order to keep the same viewing angle as the camera (62°).

Participants

Ten observers (five men, five women) participated in this experiment. The participants were naive with respect to the setup and purpose of this experiment. All participants had normal or corrected-to-normal vision. They all gave written, informed consent. All experiments were done in agreement with the Declaration of Helsinki, Dutch Law, local ethical guidelines, and approved by the TUDelft Human Research Ethics Committee.

Procedure

Before the start of the experiment, we explained the experiment procedure, the probe, and its controls to the observers. In both training and trial sessions, the task was to set the parameters of the probe's lighting (direction, intensities of directed and ambient light) to make it appear like it fit in the scene. In order to prevent misunderstanding of the position of the probe in the scene, it was explicitly stated that the end of the pole was always standing on a visible object. We then showed photographs of a real white sphere under each light condition and ensured that the participant understood how to control the probe by doing three training trials. The following main part of the experiment consisted of 108 trials, including 36 probe positions for each light condition. We ceded the trial repetitions to be able to test a larger number of probe locations, taking into account that in the Koenderink et al. (2007) experiment, the reproducibility over sessions was stated as "fair" (medians of the quartile deviations over all observers of 5.7° were found for the slant, 4.9° for the tilt, 0.023 for the intensity, which was defined on a range from 0 to 1, and 0.093 for the directedness, which was defined on a range from −1 to 1). The order of the 108 trials was randomized per observer. After the first 54 trials, the observers were given a break. Altogether, the experiment took between 1 and 2 hr, with an average of 1 hr and 20 min. After the experiment, we asked the participants to draw where they thought the light source(s) were positioned on pictures of the three scenes, and to describe the position of the source(s) in words.

Results

Source estimation

The results of the survey on the inferred light sources positions are presented in Figure 3. For the LAMP condition with the visible light source, all observers pointed at the lamp. For the DIFFUSE condition, there was more variation but most of the light sources were "placed" in the upper right part of the ceiling. The observer who pointed a light source on the left part of the picture still stated that it is directed to the right. Another observer positioned the light source above the standing lamp next to the ventilation hole and commented that it is an invisible light source floating next to the wall. Finally, only 4 of 10 observers noticed that there were two light sources in the SPOTLIGHTS condition. Of the others, three stated that the light source was in the middle of the ceiling, one claimed that there was only one elongated source in the middle, one placed the light source on the right, and one placed it in front of the scene.

Local settings

For each light parameter and each condition, we evaluated the intersubjective spread in the parameter settings by calculating the medians of the quartile deviations over all observers per point using the same statistics that were used in Koenderink et al.'s (2007) experiment. Overall, the intersubjective spreads in the settings (called *spread* in the rest of this section) seem to stay within reasonable percentages of the full ranges. The directional settings were decomposed into two angles, *polar angle* and *azimuth*. Polar angle indicates an angle perpendicular to the picture plane, varying from 0° (frontal illumination from the point of view of the observer) to 180° (backward illumination). Azimuth indicates an angle in the picture plane, varying from 0° (right side of the probe illuminated) to 360°, counter-clockwise. For these angles, we used circular statistics. There does not seem to be a strong dependency of polar angle spread on light condition (medians of 11.3°, 14.5°, and 15.0° for LAMP, DIFFUSE, and SPOTLIGHTS, respectively). However, there was a strong dependency of azimuth spread on light condition, increasing from a median of 7.1° in the LAMP light condition to 14.6° in the DIFFUSE condition, and 18.3° in the SPOTLIGHTS condition. Please note that these intersubjective spreads of the probe settings correlate with the apparent spread of the subjective light sources' positions in Figure 3.

The photographs were taken with varying exposure times, so there is no meaning in direct comparisons between light intensities of the settings. However, the ratio of the directed and ambient intensity, or the diffuseness, is relative and will be addressed extensively in the comparison of the visual with the physical qualities. The spread of the diffuseness (on a range 0–1) was 0.19 for the LAMP condition, 0.14 for the DIFFUSE condition, and 0.18 for the SPOTLIGHTS condition. Compared to spreads for diffuseness found by others (reviewed in Xia et al., 2016a), these values are typical and also stay within reasonable percentages of the full range.

Thus, the local settings show rather robust behavior, allowing comparison of the average structure of the light field over observers with the physical light field (see Comparison of visual and physical light fields section). Now we will address the question whether it is possible to actually do these reconstructions in combination with demonstrating possible ways to visualize them.

Reconstruction and visualization of the visual light field

Light field visualization is difficult due to the high dimensionality of the data. As the light field is a function of location (x, y, z) and direction (θ, φ), we need to project five dimensions into two for its representation in a flat image. The task is even more challenging because flat images inherently contain ambiguities. Mury et al. (2009) demonstrated several methods for visualizations of (certain properties of) the physical light field, such as contour plots of its components' strength distributions, collections of panoramic projections, fields of projected light vectors, and light tubes.

One way to visualize the visual light field is to simply superpose the adjusted probe objects on the scene (see Figure 4, top three rows). The white spheres represent the observer's fits in straightforward manner. Using this method, we can get an impression of the settings that is highly visual. This method, unlike others, allows display of the ambient component. However, the direction and diffuseness of the light on a smooth white sphere have been shown to interact perceptually (Pont & Koenderink, 2007). Additionally, such representation is discrete; it needs integration and interpolation to infer the global structure of the visual light field.

The light vector representation (see Figure 4, fourth row) consists of arrows depicting adjusted directions and relative intensities of the directed light. These vectors are taken to be the perceptual equivalent of the physical light vector. The directional components are perhaps more clear than in the sphere images, but still it is difficult to see the global structure.

The last method we demonstrate is the light tubes visualization method (see Figure 4 bottom row; Gershun, 1939; Mury et al., 2009). The light tubes are locally tangential to the light vector and their width is inversely proportional to its strength. One can think about the tube as an enclosure of a part of the light flow: the amount of light passing through the cross-section is constant over the length of the tube. In physical light fields the tubes usually diverge from a light source and end on light absorbing surfaces. Such representations give an impression of the so-called light flow (Cuttle, 1973) through a space at first glance, and concern the global structure of the light field.

We calculated the light tubes for the visual light fields using interpolation of the light vectors parameters (direction, intensity) between neighboring measurement points. In their reconstructions of physical light fields, Mury et al. (2009) started the tubes from the top, since all the light sources were on the ceiling. Our algorithm for the tubes visualization started from the light absorbing surfaces because we had a condition in which the light source was located within the volume of the scene. The recursive algorithm created a matrix of tubes components: points constituting the tube path through the volume plus the tube widths. At initialization, the algorithm calculated the starting coordinates such that the tubes origins were evenly distributed over the outer

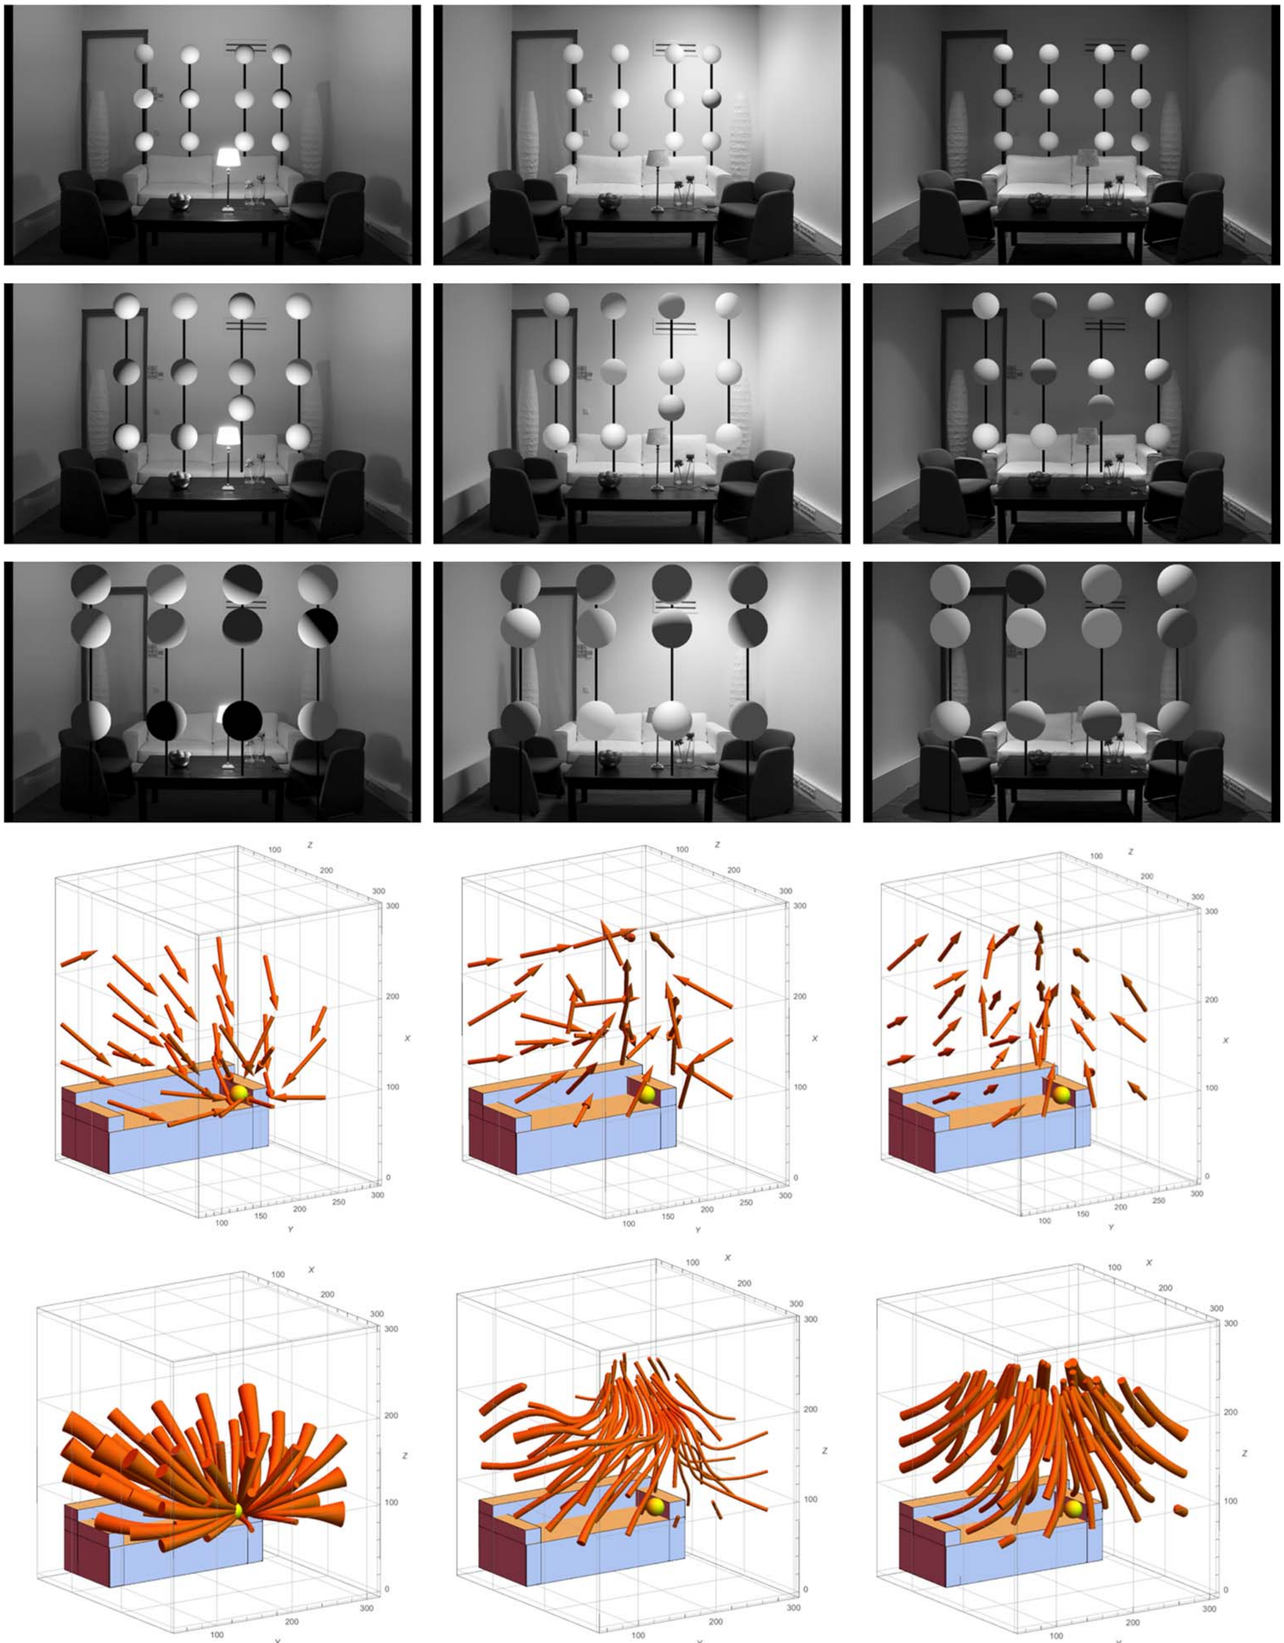

Figure 4. Visual light field visualization methods: spheres, arrows, and tubes. For one of the observers, each column illustrates the settings in one of the light conditions: from left to right, LAMP, DIFFUSE, and SPOTLIGHTS. The spheres show the appearance of the adjusted probes per depth plane. The arrows show the adjusted direction of the light pointing in the direction from which the light arrives with the arrows' lengths corresponding to the relative strengths of the directed light. The thickness of the tubes is inversely proportional to the relative strength of the directed light and their direction is locally tangential to the adjusted light direction.

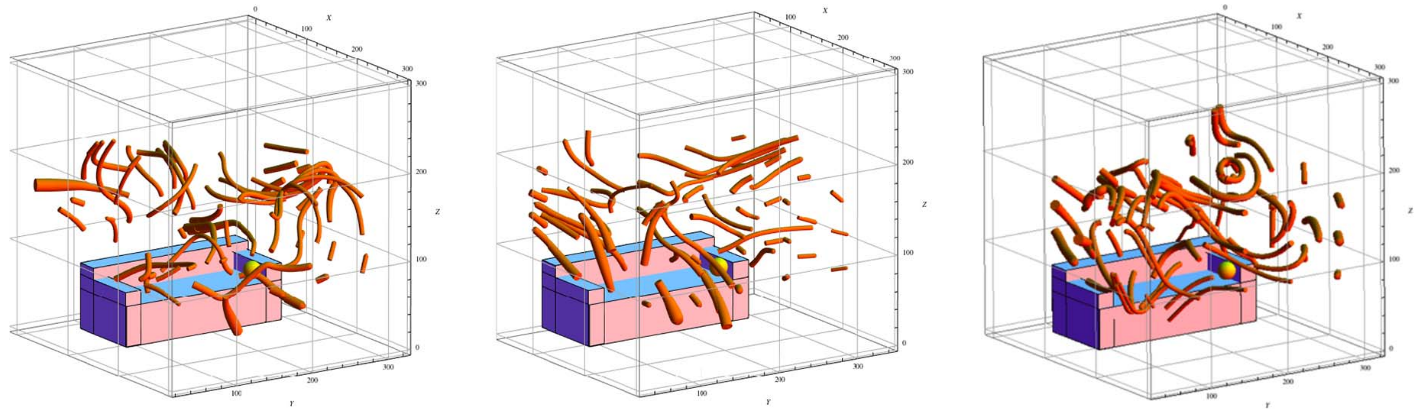

Figure 5. Examples of tubes created on the basis of randomly oriented vectors, which we made in order to check how coherent the interpolated flow patterns would be for locally random settings. It is obvious that these light flows do not represent a coherent "field."

bounds of the visualization volume. The tubes' initial widths and light vector directions were then interpolated from these coordinates using linear interpolation functions for the visualized light condition (Mury et al., 2009). For each tube, the algorithm made a step in the interpolated light vector direction on the next iterations. The stopping condition for each tube was either reaching the predefined limit of steps, the tube leaving the volume of the scene, or the tube fluctuating in a small area (which would mean that the tube reached a light source in the visualized volume). The initial width, number of steps, and number of tubes were adjusted manually to optimize the imaged light flow and avoid cluttering.

It is important to note that the light fields presented in Figure 4 are examples of the individual visual light fields. We provided visualizations of all individual visual light fields in the supplementary material with this paper. Although the intersubjective spreads were quite small, some of the reconstructed individual visual light fields for the DIFFUSE and SPOTLIGHTS look quite different. This is probably due to global differences in the settings, which have a minor influence on the intersubjective spreads because those were based on local settings, but clearly change the overall reconstruction's appearance. We found that our experimental method provides sufficiently robust data to reconstruct the individual visual light fields for our scenes. Mury et al. (2009) tested how good their reconstructions were by comparing several interpolated values in between initially measured points with extra measurements at those points, concluding that, although the local values were not exactly identical, the global structure could be measured robustly. In order to avoid increasing the number of measurements, we used the existing points only. First, we excluded a point from the grid, then ran the interpolation on this grid, and finally calculated the angular and intensity differences between the excluded point values (measured vector) and interpolated ones (approximated vector on the basis of the other points of the grid). We repeated this for each point in each individual light field reconstruction and found that the median values for the angular differences were fairly small, 19°, 23°, and 23° for LAMP, DIFFUSE, and SPOTLIGHTS conditions, respectively.

We also ran the light flow reconstruction algorithm using randomly directed vectors with randomly appointed intensities. The angular differences were then found to be around 90° after running the algorithm on 100 sets of random vectors. To obtain the difference in intensities, we took the median of the absolute difference between the values, and it constituted 0.15, 0.12, 0.13, and 0.28 for LAMP, DIFFUSE, SPOTLIGHTS, and random conditions, respectively. To illustrate the contrast between the visual light fields and the models based on random values, Figure 5 presents examples of the outcome when running the interpolation and visualization algorithm on sets of random vectors. We can see that the reconstruction method breaks down for random data. Thus, our reconstructions of the global visual light field do represent its structure and are not an artifact of the interpolation method.

The physical light field

Measurements

In order to compare the visual light fields with the corresponding physical light fields, we first had to reconstruct the global structure of the physical light fields for the same scene and light conditions as in the psychophysical experiment. Mury et al. (2009) captured physical light fields for the same light lab as we used, though it was empty at that time. They measured over a grid of $3 \times 5 \times 3$ points with a step size of 1 m using a plenopter, a custom-made illuminance meter with 12 measuring heads, to be able to reconstruct the light field up to its second order spherical harmonics representation, which can be described by nine coefficients, making the 12 measured values sufficient for its estimation. We took this approach and tuned it to our purposes.

The measurements were done over 49 points: a grid of three in height, five in width and three in depth of the scene, and four additional points (Figure 6). Unfortunately, the furniture disposition in the scene did not allow making the grid perfectly regular. The positions of the measurements can be seen in Figure 7. The heights of the measurements were 80, 145, and 210 cm, except for the lowest row behind the sofa, which was measured at the height of 122 cm. The heights of the four additional points are all 145 cm.

In our study we limited ourselves to the ambient and light vector components, which correspond to the zero and first order components in the spherical harmonics representation. Consequently, we did not need 12 measurements at each position as Mury et al. (2009) needed to include the second order spherical harmonics, so a substantially lower number of measurements were included, which were just enough to be able to estimate the four coefficients describing the zero and first order spherical harmonics representation. A cube was the closest approximation of a sphere having a regular shape with four or more faces.

We used a custom-made device (for details, see Xia et al., 2016b) for cubic illuminance measurements (Cuttle, 2013) on the basis of a Konica Minolta T-10MA illuminance meter (Konica Minolta, Inc., Tokyo, Japan), plus extra heads. The main part of the device is a cube, covered by black velvety paper, with a sensor on each surface. The cube was fixed on a stick and tilted such that its long corner-to-corner diagonal was vertical, and the stick was fixed on a tripod. Our construction allowed the lowest measurement to be 50 cm from the ground and the highest about 2.5 m. On the top of the cube there was a bubble level and on the tripod there were protractors, which allowed adjusting the cube orientation horizontally. All six sensors were connected to a laptop through the luminance meter's main body. In this way, we could make simultaneous illuminance measurements from all six sensors. Following Cuttle's (2013) procedure, we could then derive the local light field parameters from these six records and finally reconstruct the global structure of the light field by interpolation. The cubic illuminance measurements satisfied us in almost all respects. Yet, for one point close to the light source in the LAMP light condition, the resulting vector did not point to the lamp, due to known limitations of the cubic measurements whenever individual sensors record (close to) zero illuminance, as discussed by Xia et al. (2016b). This point was excluded from the analysis.

The procedure of the measurements was the following: set the tripod to the position, check the orientation of the cube, then make measurement and photograph (for reference) for each light condition. The minimum values of the sensors measurements were 3.97, 45.4, and 19.22 lux and the maximum values were 173.9, 2040, and 1493 lux for LAMP, DIFFUSE, and SPOTLIGHTS conditions, respectively.

Data processing and visualization

The resulting measurement data constituted the six measurements per position and light condition. We translated the measurements to light vectors (a) via the method introduced by Cuttle (2013) and (b) via spherical harmonics approximations as introduced by Mury et al. (2007, 2009). These two methods are based on the same concepts, but framed within different mathematical approaches. In Xia et al. (2016b) an extensive comparison of these two methods is given, which leads to the conclusion that they give very similar results under natural circumstances.

Cuttle's (2013) method has a straightforward approach. The components of the resulting light vector are calculated by subtraction of the measurements of opposing faces, and rotation of the results to align with standard axes, taking into account that, initially, measurements were done with a tilted cube (see above).

In a spherical harmonics representation of the light field, the zero order component corresponds to the

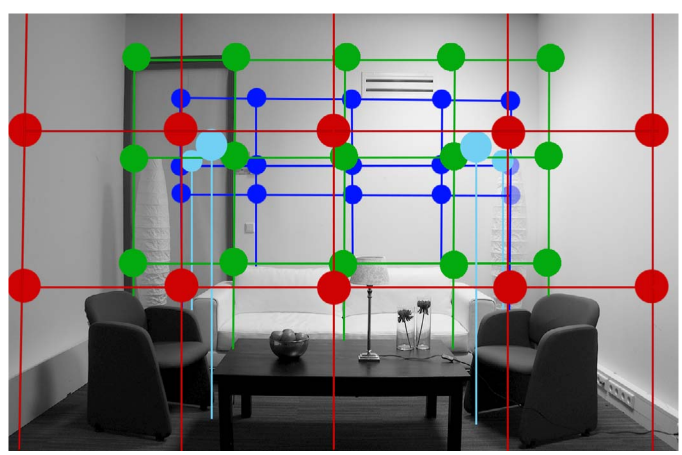

Figure 6. Scheme of measurement positions for the physical light measurements. Colored lines connect measurement positions at the same depth and denote where the cubical light meter was standing. The top row of the closest measurements is above the field of view of the camera.

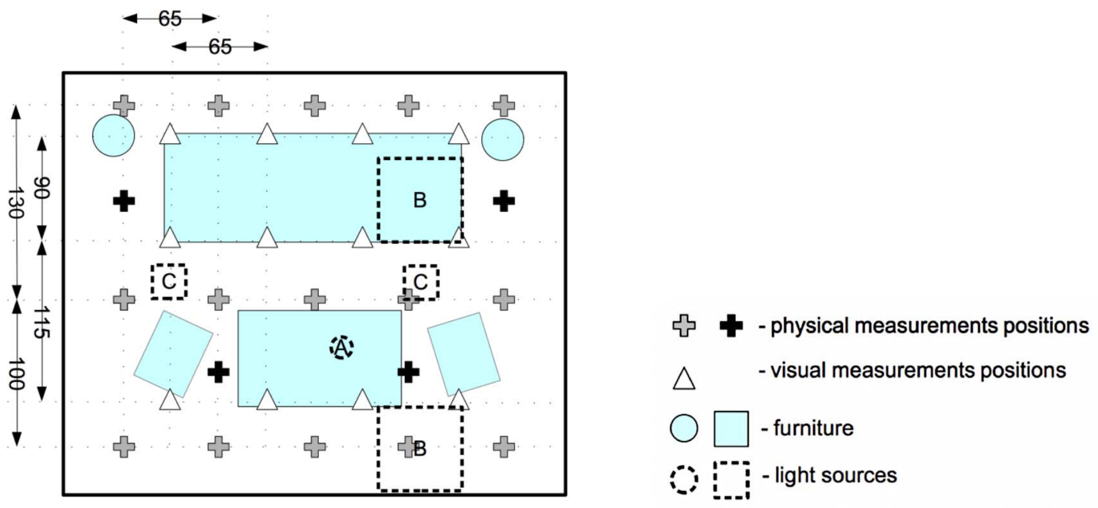

Figure 7. Scene scheme. Blue rectangles and circles define furniture positions; gray crosses represent the physical measurement positions and black crosses show four additional measurement positions. Triangles represent the psychophysical measurement positions, and dashed shapes define the light sources positions: A = lamp on the table, B = diffused light sources in the ceiling, and C = directed light sources in the ceiling.

ambient component of the light (Cuttle's "density of light") and the first order component to the light vector. In order to calculate the coefficients, we used the system of equations proposed by Mury et al. (2009) and adapted it for 6 measurements instead of 12, since we were interested in the zero and first orders only. To use this system we had to define the sensitivity profile of the sensors, which we retrieved from the documentation of the sensors. The difference between the resulting vectors from the two methods was negligible so we used Cuttle's (2013) method for the calculations.

We applied the same interpolation and visualization algorithm for the physical light fields as for the visual light fields (see the section, Reconstruction and visualization of the visual light field).

Diffuseness was calculated according to Cuttle's (2003, 2013) method using a cubic illuminance meter:

$$E_{(x)} = E_{x+} - E_{x-} \quad (1)$$

$$E_{\text{vector}} = \sqrt{E_{(x)}^2 + E_{(y)}^2 + E_{(z)}^2} \quad (2)$$

$$\sim E_x = \frac{E_{x+} + E_{x-} - E_{x\,\text{vector}}}{2} \quad (3)$$

$$E_{\text{symmetric}} = \frac{\sim E_x + \sim E_y + \sim E_z}{3} \quad (4)$$

$$E_{\text{scalar}} = E_{\text{vector}}/4 + E_{\text{symmetric}} \quad (5)$$

$$(D_{\text{Cuttle}})_{\text{Normalized}} = 1 - (E_{\text{vector}}/E_{\text{scalar}})/4 \quad (6)$$

E_+, E_- are the measurements of opposite sides of the cube; $E_{(x)}$, $E_{(y)}$, $E_{(z)}$ are the light vector components on each axis; E_{vector} is the light vector magnitude; $E_{\text{symmetric}}$ is the symmetric illuminance; and E_{scalar} is the mean illuminance in a point. For the observers' settings, E_{vector} and $E_{\text{symmetric}}$ were taken as the magnitudes of directed and ambient light, respectively. The normalized diffuseness $(D_{\text{Cuttle}})_{\text{Normalized}}$ ranges from 0 (fully collimated light) to 1 (fully diffuse light). See Xia et al. (2016a) for extensive theoretical and empirical comparisons of this method with other methods to define and measure diffuseness.

Comparison of visual and physical light fields

As was already stated, light fields have a complex structure, which makes the visualization, but also the quantitative analysis, difficult. To compare the light fields we chose to analyze the light vectors' directions first. Figure 8A, top row, shows three-dimensional illustrations of the physical and observers' averaged vectors in the LAMP, DIFFUSE, and SPOTLIGHTS conditions. We analyzed these data by calculating for each position, observer, and light condition the angular difference between the visual and physical light vectors in three dimensions, as well as between their projections on the picture plane. We took into account the projected vectors comparisons (two-dimensional; i.e., taking into account the azimuthal angle settings and neglecting the polar angle settings) because polar angle settings of lighting on a sphere were proven to suffer from the bas-relief ambiguity (Belhumeur, Kriegman, & Yuille, 1999; Koenderink, van Doorn, Kappers, te

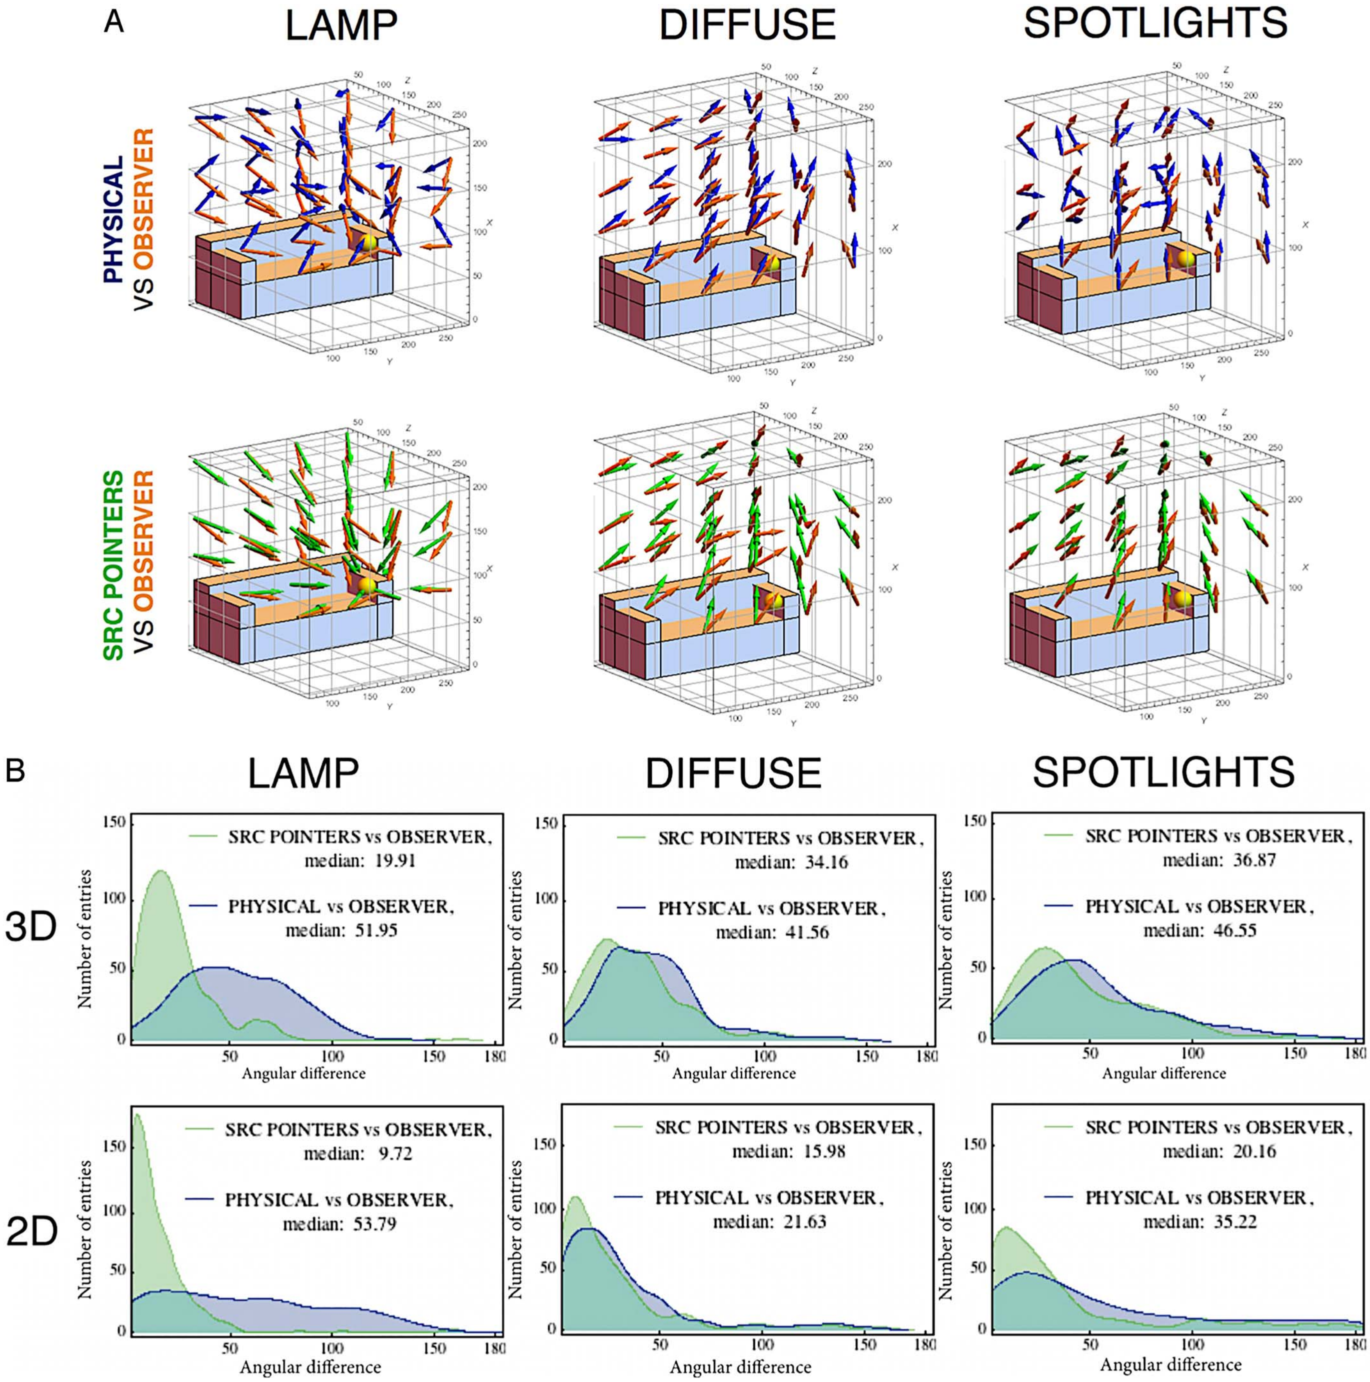

Figure 8. Light direction comparisons. (A) Three-dimensional illustrations of (normalized) light vectors for each light condition, with the blue arrows representing the PHYSICAL vectors, the orange representing (averaged) OBSERVER vectors, and green the SRC POINTERS vectors. (B) Histograms of angular differences between physical light vectors and observers' directional components (PHYSICAL vs. OBSERVER = blue histograms), and between vectors pointing at the subjective light sources and observers' directional components (SRC POINTERS vs. OBSERVER = green histograms) in three-dimensional representations (top row) and two-dimensional representations (bottom row).

Pas, & Pont, 2003; Koenderink et al., 2007). The resulting angular differences for all positions and all observers were summarized in smoothed histograms (see Figure 8B). First, we took the unsigned differences of the observers' settings and the physical measurements results (PHYSICAL vs. OBSERVER). For each light condition, the two-dimensional results showed a maximum around 10° difference, which is rather small for this probing method but typical for azimuth settings in illuminance flow inferences (Koenderink et al., 2007;

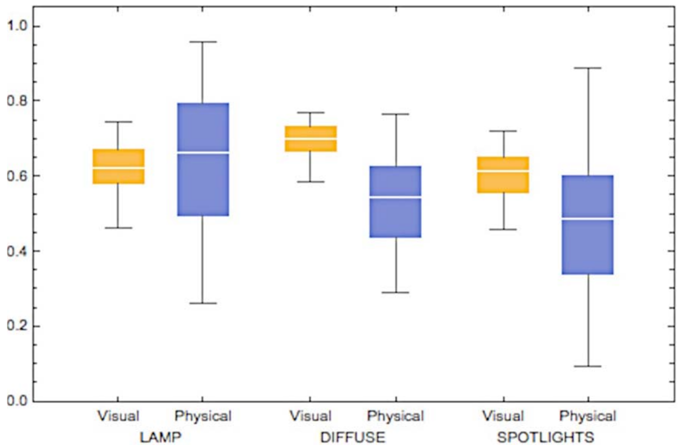

Figure 9. Box-and-whisker chart of normalized visual and physical light diffuseness as a function of light condition, ranging between 0 (fully collimated) and 1 (fully diffused light). The boxes represent the range between 25% and 75% (first and third quartile) and the bars represent the full range.

Xia et al., 2016d). The distributions are broad, however, resulting in medians of 54°, 22°, and 35°. The three-dimensional results had maxima, as expected, at higher angular differences, namely between 40° and 60°, and medians of 52°, 42°, and 47°. Next, we used the same approach to compare the observers' settings and a model (SRC POINTERS vs. OBSERVER), based on the physical light source positions. The SRC POINTERS model represents essentially a simple divergent field, which would occur if observers would be simply pointing to light sources, and in physics, if there would be a light source in empty space. We obtained the positions of the light sources from the scene measurements for each light condition and pointed all the vectors to these positions (see Figure 8a, second row). For the DIFFUSE and SPOTLIGHTS conditions with two light sources, the directions of the vectors were calculated using linear superposition and the inverse square law of the distances to the sources. The histograms of comparisons between the observers' settings and the SRC POINTERS model moved noticeably to lower values, with respect to those for the PHYSICAL vs. OBSERVER histograms, resulting in lower medians for all three lighting conditions: 10°, 16°, and 20° for the two-dimensional analysis and 20°, 34°, and 37° for the three-dimensional analysis. Pairwise Kolmogorov-Smirnov tests on the cumulative raw data confirmed this shift by showing that the shapes of the PHYSICAL vs. OBSERVER histograms are significantly different from the SRC POINTERS vs. OBSERVER for all two-dimensional and three-dimensional comparisons (for $p < 0.05$). Thus, human observers' settings were closer to the predictions of the diverging field models than to the physical values.

In Figure 9 we show the averaged data for the diffuseness values. Perhaps surprisingly, the spreads of the diffuseness settings of the observers were always smaller than the spreads of the physical diffuseness values. The overall means were different for physical and visual light fields. Specifically, human observers considered the DIFFUSE light condition the most diffuse (as did the authors, judging on the visual appearance), whereas physically the overall most diffuse condition turned out to be the LAMP condition. The ranges of the diffuseness values showed the same relation between conditions: for both physical and perceptual measurements, the DIFFUSE condition showed the smallest range, and the SPOTLIGHTS condition showed the biggest range.

We present the three-dimensional visualizations of the physical light fields and (averaged between observers) visual light fields in Figure 10. The physical light field for the LAMP light condition shows curved tubes at the top, due to the scattering of the light from the ceiling and interreflections in the corners. Towards the back of the scene, the light is more diffuse due to reflections from the ceiling and walls and dimmer because of the distance to the lamp. This causes a change of direction over the tube's length and thus, its curvy shape. None of the observers took these changes in the local average light direction into account for the LAMP condition, so their averaged settings form rather straight tubes diverging away from the light. The results of the physical and visual light field reconstructions were found to be most similar for the DIFFUSE light condition: the tubes seem to have the same origin, in the top left part of the ceiling, and spread out from there. In this condition, the physical tubes are also rather straight. In the SPOTLIGHTS condition, the physical reconstructions show slim, almost vertical, tubes under the lamps, where the light is highly directed. In other parts of the scene, the light is primarily due to scattering, causing the tubes to be thick and odd-shaped. The visual light field for this light condition shows tubes converging to the ceiling, and no odd-shaped tubes.

Discussion

Measuring visual light fields is a novel technique that reveals how human observers make inferences about the structure of the physical light field. We developed this method by merging existing approaches of measuring the physical light field (Mury et al., 2009) and the local visual light field (Koenderink et al., 2007). We have shown that the data obtained by this method is sufficient to reconstruct and visualize individual visual light fields. In addition to the psychophysical

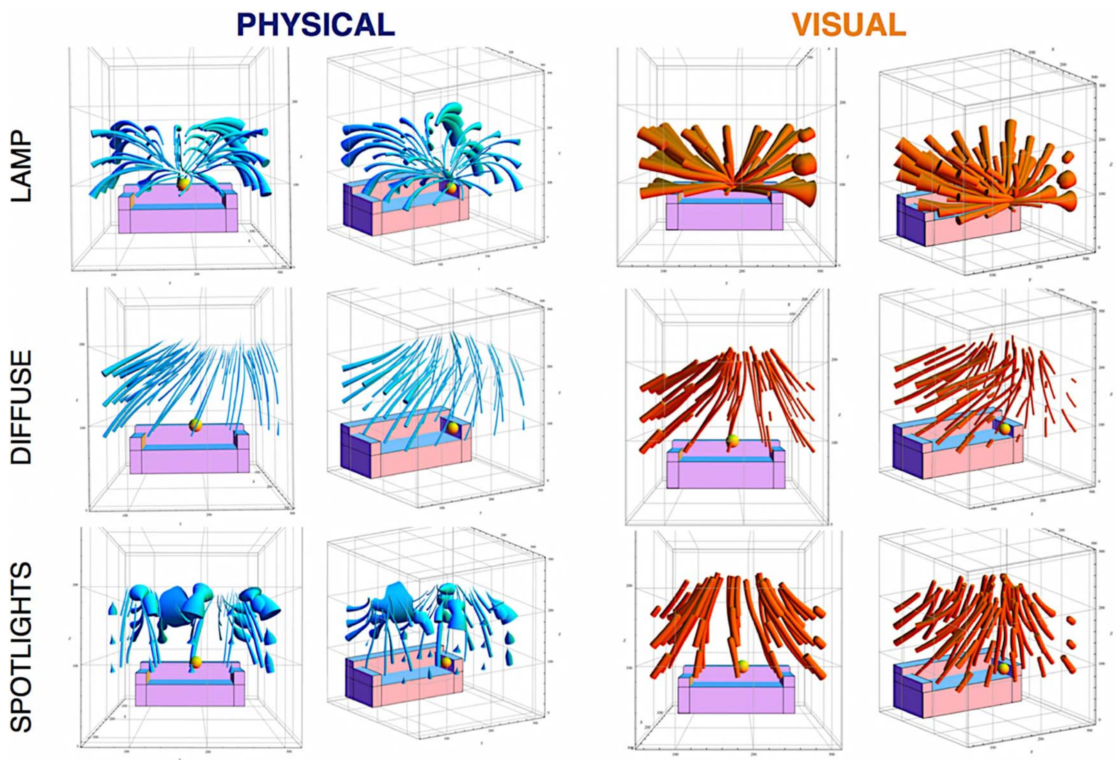

Figure 10. Physical and visual light fields. Blue shows reconstructions via light tubes of the physical measurements and orange shows reconstructions of the averaged observers' settings. Two different perspectives are shown for each condition.

measurements, we did physical measurements and compared the visual and physical light fields for three light conditions.

The goal of this study was to examine the structure of the visual light field. One way to assess it is to make multiple local measurements and interpolate them in order to provide values of the light properties in an arbitrary point within the measured volume. For the local measurements we used Koenderink et al.'s (2007) method of visual fit for illumination probing, which proved its reliability already in their and others' studies (Pont & Koenderink, 2007; Pont et al., 2010; Xia et al., 2014). Additionally, we took into account that Mury et al. (2009) demonstrated that the grid of 45 ($3 \times 5 \times 3$) points was sufficient to reconstruct a physical light field in the light lab where we also did our measurements. We did our measurements in only part (about half) of the room, but with a similar number of measurement points in the grid, so that the sampling density was about a factor of two higher. This was done because the living room scene is more complex than Mury et al.'s (2009) empty room.

Our first result was the finding that observers' settings on a probe are indeed reliable enough to reconstruct the visual light field for real scenes. In Koenderink et al.'s (2007) experiment, quartile deviations between repetitions were within 4° and 12° for the polar angle ("slant" in their terms) and 1.6° and 13° for azimuth ("tilt"). Since we did not do repetitions of trials, it is not possible to make direct comparisons of the spreads with their study. Additionally, it was stated that all observers reproduced azimuth (tilt) within 5° and 10°, which is more accurate than our results. We connect it to the fact that the light fields in our experiment were more complex than theirs. The striking difference between the interpolations of the visual measurements (Figures 4 and 10) and of the random vector grids (Figure 5) demonstrates that the data was regular enough to reconstruct visual light fields.

We discovered that the structure of the visual light field can be quite different from the physical one. The analysis of the light flows (see illustration via light tubes in Figure 10) shows that the participants' settings seem to grasp the basic structure of the physical light field and converge at light sources, but ignore subtle changes due to (inter)reflections. Van Doorn, Koenderink, and Wagemans (2011) and van Doorn et al. (2012) revealed that human observers are able to infer convergent and divergent two-dimensional light flows but not more complex ones, such as rotational and deformation two-dimensional flows. Our findings confirm their conclusions also for three-dimensional global structures of light fields. More specifically, the observers' results correlated better with the models containing almost straight tubes diverging out from the sources, than with the physical, more complicated truth. This was particularly evident in the LAMP condition with its strong reflections from the ceiling and walls. Figure 3 shows that observers were actually not able to infer where the sources were unless the lamp was visible. Altogether, these results suggest that what the observers do is far from "inverse optics" and that instead, a diverging field is a template for the visual light field.

An important remaining question is: is the inferred global visual light field something that is represented perceptually? In other words, do observers only make local settings on the basis of object appearance and are the reconstructed visual fields mere mathematical inferences by the authors, or do the observers have a fairly good internal representation of the light field? The consistencies within and between the findings in Koenderink et al. (2007), van Doorn et al., (2011, 2012), Schirillo (2013), and our study perhaps suggest that observers have such representations, though they concern strongly simplified (convergent/divergent light flow structures with fairly straight flow lines/tubes) compared to natural light flows. However, further research is needed to provide conclusive research on the internal representation of the luminous environment.

Our results revealed some dependencies of the visual light field on the features of the scenes, which is especially apparent in the similarities between the settings on the spheres and light location estimates. Specifically, in the scene with a single visible light source in the image, the LAMP condition, the observers seemed to simply have pointed at the source both in the probe settings and in the light source position estimations. For this condition, the settings on the spheres showed the highest consistency between subjects. In the DIFFUSE light condition, the light source was not visible but the physical light field had a relatively simple structure and observers inferred it fairly well. In the SPOTLIGHTS condition, both the settings on the probe and the source estimations were the least consistent overall. The dependencies between the complexity of the stimulus scene and lighting need further investigation. The approach demonstrated in this paper could be used in studies answering questions such as whether perception of the luminous environment is dependent on the scene being empty or full of objects, and how different scene geometries and materials affect the structure of the perceived light field. On the basis of earlier results — plus the basic fact that the optical structure entering the eye confounds geometry, material, and light — we expect interactions between shape/space, material, and light perception (Anderson, 2011; Fleming, 2014; Pont & te Pas, 2006; Zhang, de Ridder, & Pont, 2015).

Currently there is a rise of interest into the topic of light diffuseness (Koenderink et al., 2007; Pont et al., 2010; Morgenstern et al., 2014; Xia et al., 2014, 2015, 2016a, 2016b). As expected on the basis of these studies, we found that the physical diffuseness varied in a much wider range over the scenes than the visual diffuseness. In the LAMP light condition, a lampshade scattered the light in all directions except for the light exiting in the direction of the ceiling. The white ceiling and walls function as big diffusers. In the DIFFUSE light condition, the light sources themselves are rather diffuse, but they were directed downward to the dark floor, which functioned as a light absorber. This result, namely that the scene with physically the most diffuse light turned out not to be the DIFFUSE condition as the authors and observers thought but, instead, the LAMP condition, is perhaps surprising. The resultant light (field) is thus determined by the relation between the light source positions and the scene geometry and materials, and not primarily by the illuminants. Koenderink and van Doorn (1983), Mury et al. (2007), and Xia et al. (2016d) confirm and give insights into such optical mechanisms.

We found that human observers have a robust impression of the light field that is simplified with respect to the physical light field and that corresponds rather well with a model based on simple divergent fields from the light sources. These results have high practical interest. For example, the understanding of observers' inferences of light propagation through spaces can be used in lighting design for architecture or computer graphics. Moreover, the experiment setup itself might be used as a tool for visual light probing, exempting a designer or other lighting professional from physical measurements, if he or she is only interested in perceived light qualities. The method allows us to obtain the spatial structure of a visual light field from a single experimental session. Having such a tool, it is possible to construct the inferences that observers make about the light propagation through a scene and the variation of its qualities along the flow. We think that this fast and cheap tool has a big potential in perception studies, lighting, and computer graphics industries.

Conclusions

In this study, we focused on the global structure of the visual light field. Our main questions were whether the global visual light field can be measured, and if so, how similar it is to the physical light field. Our method for constructing the visual light field and its visualization via interpolation of regularized local measurements is shown to be robust even for individual light fields. Additionally, our comparisons of the visual and physical measurements results suggest that human observers have consistent impressions of the light field,

though not exactly corresponding to the physical truth; specifically, they tend to neglect subtle spatial variations in the physical light fields.

Keywords: light, light field, visual perception, probe, illumination, light qualities, visualization, flow

Acknowledgments

The authors would like to thank Ingrid Heynderickx for her help on initial stages of the study and the anonymous reviewers for their valuable comments and suggestions to improve the manuscript. We thank the Visual Experiences Group, Philips Research Europe, for the use of their light lab. This work has been funded by the EU FP7 Marie Curie Initial Training Networks (ITN) project PRISM, Perceptual Representation of Illumination, Shape and Material (PITN-GA-2012-316746).

Commercial relationships: none.
Corresponding author: Tatiana Kartashova.
Email: t.kartashova@tudelft.nl.
Address: Landbergstraat 15, 2628CE, Delft, The Netherlands.

References

Adelson, E., & Bergen, J. (1991). The plenoptic function and the elements of early vision. In M. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 3–20). Cambridge, MA: MIT Press.

Anderson, B. L. (2011). Visual perception of materials and surfaces. *Current Biology*, *21*, R978–R983.

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. L. (1999). The bas-relief ambiguity. *International Journal of Computer Vision*, *35*, 33–44.

Berbaum, K., Bever, T., & Chung, C. S. (1983). Light source position in the perception of object shape. *Perception*, *12*, 411–416.

Boyce, P. (1981). *Human factors in lighting*. London, UK: Applied Science Publishers.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Cuttle, C., & Ilium, M. (1973). *The sharpness and the flow of light*. In R. Küller (Ed.), *Architectural psychology—Proceedings of the conference held at Lund University* (pp. 12–22). Lund, Sweden: Lund University.

Cuttle, C. (2003). *Lighting by design*. Oxford, UK: Architectural Press.

Cuttle, C. (2013). Research note: A practical approach to cubic illuminance measurement. *Lighting Research & Technology*, *46*(1), 31–34.

Doerschner, K., Boyaci, H., & Maloney, L. T. (2007). Testing limits on matte surface color perception in three-dimensional scenes with complex light fields. *Vision Research*, *47*, 3409–3423.

Dror, R. O., Willsky, A. S., & Adelson, E. H. (2004). Statistical characterization of real-world illumination. *Journal of Vision*, *4*(9):11, 821–837, doi:10.1167/4.9.11. [PubMed] [Article]

Fleming, R. W., Dror, R. O., & Adelson, E. H. (2003). Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, *3*(5):3, 347–368, doi:10.1167/3.5.3. [PubMed] [Article]

Fleming, R. W. (2014). Visual perception of materials and their properties. *Vision Research*, *94*, 62–75.

Ganslandt, R., & Hofmann, H. (1992). *Handbook of lighting design*. Lüdenscheid, Germany: Vieweg, ERCO Edition.

Gerhard, H., & Maloney, L. (2010). Estimating changes in lighting direction in binocularly viewed three-dimensional scenes. *Journal of Vision*, *10*(9): 14, 1–22, doi:10.1167/10.9.14. [PubMed] [Article]

Gershun, A. (1939). The light field [Translated by P. Moon & G. Timoshenko]. *Journal of Mathematics and Physics, 18*, 51–151.

Gilchrist, A. L., & Radonjic, A. (2009). Functional frameworks of illumination revealed by probe disk technique. *Journal of Vision*, *10*(5):6, 1–12, doi:10.1167/10.5.6. [PubMed] [Article]

Ho, Y.-X., Landy, M. S., & Maloney, L. T. (2006). How direction of illumination affects visually perceived surface roughness. *Journal of Vision*, *6*(5): 8, 634–648, doi:10.1167/6.5.8. [PubMed] [Article]

Hunter, F., Biver, S., & Fuqua, P. (2007). *Light science and magic*. Oxford, UK: Elsevier.

Koenderink, J. J., Pont, S. C., van Doorn, A. J., Kappers, A. M. L., & Todd, J. T. (2007). The visual light field. *Perception*, *36*, 1595–1610. doi:10.1068/p5672.

Koenderink, J. J., & van Doorn, A. J. (1983). Geometrical modes as a general method to treat diffuse interreflections in radiometry. *Journal of the Optical Society of America*, *73*, 843–850.

Koenderink, J. J., & van Doorn, A. J. (2006). Shape from shading. In N. Paragios, Y. Chen, & O. D. Faugeras (Eds.), *Handbook of mathematical models in computer vision* (pp. 375–388). New York: Springer US.

Koenderink, J. J., van Doorn, A. J., Kappers, A. M. L., te Pas, S. F., & Pont, S. C. (2003). Illumination direction from texture shading. *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, *20*, 987–995.

Madsen, M. (2007). Light-zones(s): As concept and tool. *ARCC Journal*, *4*(1), 50–59.

Maloney, L. T., Gerhard, H. E., Boyaci, H., & Doerschner, K. (2010). Surface color perception and light field estimation in 3D scenes. In L. Harris & M. Jenkin (Eds.), *Vision in 3D environments* (pp. 65–88). New York: Cambridge University Press.

Marlow, P. J., Kim, J., & Anderson, B. L. (2012). The perception and misperception of specular surface reflectance. *Current Biology*, *22*, 1909–1913. doi:10.1016/j.cub.2012.08.009.

Moon, P., & Spencer, D. E. (1981). *The photic field.* Cambridge, MA: MIT Press.

Morgenstern, Y., Geisler, W. S., & Murray, R. F. (2014). Human vision is attuned to the diffuseness of natural light. *Journal of Vision*, *14*(9):15, 1–18, doi:10.1167/14.9.15. [PubMed] [Article]

Mury, A., Pont, S. C., & Koenderink, J. J. (2007). Light field constancy within natural scenes. *Applied Optics*, *46*, 7308–7316.

Mury, A., Pont, S. C., & Koenderink, J. J. (2009). Representing the light field in finite three-dimensional spaces from sparse discrete samples. *Applied Optics*, *48*, 450–457.

O'Shea, J., Agrawala, M., & Banks, M. S. (2010). The influence of shape cues on the perception of lighting direction. *Journal of Vision*, *10*(12):21, 1–21, doi:10.1167/10.12.21. [PubMed] [Article]

Ostrovsky, Y., Cavanagh, P., & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *Perception*, *34*, 1301–1314. doi:10.1068/p5418.

Pentland, A. P. (1982). Finding the illuminant direction. *Journal of Optical Society of America*, *72*(4), 448–455.

Pont, S. (2011). An ecologically valid description of the light field. *Journal of Vision*, *11*(11):, 345, doi:10.1167/11.11.345. [Abstract]

Pont, S. C. (2013). Spatial and form-giving qualities of light. *Handbook of Experimental Phenomenology: Visual Perception of Shape, Space, and Appearance* (pp. 205–222). New York: Wiley.

Pont, S. C., & Koenderink, J. J. (2007). Matching illumination of solid objects. *Perception & Psychophysics*, *69*, 459–468.

Pont, S. C., & te Pas, S. F. (2006). Material—Illumination ambiguities and the perception of solid objects. *Perception*, *35*, 1331–1350.

Pont, S. C., van Doorn, A., de Ridder, H., & Koenderink, J. J. (2010). The visual light field in the in- and outside. *Perception, 39*, ECVP Abstract Supplement, 104.

Schirillo, J. A. (2013). We infer light in space. *Psychonomic Bulletin & Review*, *20*, 905–915. doi:10.3758/s13423-013-0408-1.

van Doorn, A. J., Koenderink, J. J., Todd, J. T., & Wagemans, J. (2012). Awareness of the light field: The case of deformation. *I-Perception*, *3*, 467–480.

van Doorn, A. J., Koenderink, J. J., & Wagemans, J. (2011). Light fields and shape from shading. *Journal of Vision*, *11*(3):21, 1–21, doi:10.1167/11.3.21. [PubMed] [Article]

Wijntjes, M. W. A., & Pont, S. C. (2010). Pointing in pictorial space: Quantifying the perceived relative depth structure in mono and stereo images of natural scenes. *ACM Transactions on Applied Perception*, *7*(4), 1–8.

Xia, L., Pont, S. C., & Heynderickx, I. (2014). The visual light field in real scenes. *i-Perception, 5*(7), 613–629. doi:10.1068/i0654.

Xia, L., Pont, S. C., & Heynderickx, I. (2015). Simultaneous measurement and visualization of light flow and diffuseness in 3D spaces. In X. Name (Ed.), *Proceedings of 28th CIE Session* (pp. 556–563). City, ST: Publisher.

Xia, L., Pont, S. C., & Heynderickx, I. (2016a). Light diffuseness metric Part 1: Theory. *Lighting Research & Technology,* E-pub ahead of print. Available at http://lrt.sagepub.com/content/early/recent

Xia, L., Pont, S. C., & Heynderickx, I. (2016b). Light diffuseness metric II: describing measuring and visualizing the light flow and diffuseness in 3D spaces. *Lighting Research & Technology, XX,* 1–17.

Xia, L., Pont, S. C., & Heynderickx, I. (2016c). Probing the sensitivity of observers for light qualities in real scenes. Manuscript in preparation.

Xia, L., Pont, S. C., & Heynderickx, I. (2016d). Effects of scene content and layout on the perceived light direction in 3D spaces. Manuscript in preparation.

Yamauchi, R., Ikeda, M., & Shinoda, H. (2003). Walls surrounding a space work more efficiently construct a recognized visual space of illumination than do scattered objects. *Optical Review, 10,* 166–173.

Zhang, F., de Ridder, H., & Pont, S. C. (2015). The influence of lighting on visual perception of material qualities. In *Proceedings of SPIE 9394, Human Vision and Electronic Imaging XX,* 93940Q, doi:10.1117/12.2085021.