

Photon mapping in image-based visual comfort assessments with BSDF models of high resolution

Lars O. Grobe ^{a,b}

^aLucerne University of Applied Sciences and Arts, Technikumstrasse 21, 6048 Horw, Switzerland; ^bIzmir Institute of Technology, Urla, 35430 Izmir, Turkey

ARTICLE HISTORY

Compiled June 5, 2020

ABSTRACT

Data-driven models replicate the irregular Bidirectional Scattering Distribution Functions (BSDFs) of optically Complex Fenestration Systems in daylight simulation. RADIANCE employs the *tensor tree* to store the BSDF at high directional resolution. Its application in backward ray-tracing is however challenging, since the density of stochastic samples must match the model resolution. *BSDF proxy* and *peak extraction* address this problem, but are limited to cases when either the fenestration geometry, or the shape and direction of the transmission peak are known. Photon Mapping is proposed to efficiently sample arbitrary BSDFs from the known sun direction. The existing implementation in RADIANCE is extended to account for light sources and their reflections in the field of view, that are of particular importance for visual comfort assessments. The method achieves a high degree of accordance with ray-tracing, and reduces simulation times by $\approx 95\%$ with data-driven models of high resolution.

Abbreviations:

BRT	Backward Ray-Tracing
BSDF	Bidirectional Scattering Distribution Function
CBDM	Climate-Based Daylight Modelling
CFS	Complex Fenestration System
DGI	Daylight Glare Index
DGP	Daylight Glare Probability
FPM	Five Phase Method
GPGPU	General Purpose Graphics Processing Unit
LCP	Laser Cut Panel
OoC	Out-of-Core
PM	Photon Mapping
RMSE	Root Mean Squared Error

KEYWORDS

Photon mapping; BSDF; Radiance; complex fenestration; visual comfort; daylight simulation

This is an Accepted Manuscript of an article published by Taylor & Francis in *Journal of Building Performance Simulation* on 03 Sep 2019, available online: <https://doi.org/10.1080/19401493.2019.1653994>.

CONTACT Lars O. Grobe. Email: larsoliver.grobe@hslu.ch

1. Introduction

1.1. *Daylight simulation with data-driven BSDF models in Radiance*

The daylight simulation suite RADIANCE combines deterministic and stochastic algorithms into a hybrid implementation of Backward Ray-Tracing (BRT) (Ward 1994). The deterministic testing of concentrated light-sources, such as the sun, as well as regular transmission and reflection are solved by deterministic ray-tracing (Whitted 1980). Stochastic sampling by randomly distributed rays accounts for diffuse-indirect illumination (Cook, Porter, and Carpenter 1984). RADIANCE provides physically plausible models for transmission and reflection, has been thoroughly validated (Ward and Shakespeare 1998; Grynberg 1989; Schregle and Wienold 2004; Geisler-Moroder and Dür 2008; Jones and Reinhart 2017), and drives numerous front-ends (e.g. DAYSIM / DIVA, the DIAL+SUITE, IDA ICE, ESP-R, and OPENSTUDIO (Jakica 2018)) for applications in building simulation. ACCELERAD is a variant of RADIANCE that accelerates simulations by the massive parallelism of modern General Purpose Graphics Processing Unit (GPGPU) architectures, and allows for interactive visual comfort assessments (Jones and Reinhart 2017, 2019).

To account for irregular light scattering, a data-driven model approximates arbitrary Bidirectional Scattering Distribution Functions (BSDFs) (Heckbert 1991, p. 26) by a set of coefficients (Ward et al. 2011; McNeil, Lee, and Jonsson 2017). Since assessments of visual comfort, and glare in particular, ask for image-based evaluation techniques (Wymelenberg and Inanici 2014, 2016) that are sensitive to the capability of the fenestration model to replicate directionality (McNeil 2011; Lee, Geisler-Moroder, and Ward 2018), a data-structure of adaptive resolution is implemented in RADIANCE (Ward, Kurt, and Bonneel 2014). The *tensor tree* is compact by merging regions of low variance, yet it resolves features of the BSDF such as peaks caused by directional transmission and reflection. Models can be generated from measurements, and lend themselves in particular to micro-structures featuring complex light scattering properties, such as daylight redirecting films or coatings (Ward, Kurt, and Bonneel 2014; Kazanasmaz et al. 2016; Grobe, Wittkopf, and Kazanasmaz 2017; Grobe 2018).

The data-driven model also allows to model the irregular transmission characteristics of entire Complex Fenestration Systems (CFSs). These are otherwise not supported by the deterministic ray-tracing algorithm, which requires to know the sample directions leading toward the sun a priori, nor by stochastic backward sampling due to the impractically high amount of random rays required to sample a small source such as the sun. The software `genBSDF`, distributed with RADIANCE, pre-computes the BSDFs of such systems and compiles them into data-driven models supported by the backward algorithm (Molina et al. 2015; Mainini et al. 2019). To account for the geometric detail of macro-structured CFSs, e.g. Venetian blinds, the computationally generated, data-driven model can be evaluated only in the indirect-diffuse calculation, while an embedded geometric representation maintains visual detail and shadow patterns caused by the fenestration (Ward, Kurt, and Bonneel 2012).

One of the main challenges in the application of data-driven models even of moderate resolution in RADIANCE is their adequate sampling in the stochastic, indirect-diffuse calculation. The maximum resolution that can be achieved when modelling anisotropic reflection and transmission is currently limited to $2^7 \times 2^7 = 16\,384$ outgoing (and, equally, incident) directions. To account for directional transmission through such models in the ambient calculation pass, an equal or higher number of random rays need to be spawned at each inter-reflection step in the building interior (Rogers

2013). This drastically increases simulation times, since only a fraction of these typically collide with a data-driven model, and even fewer can be expected to reach the BSDF under an incident direction that leads to a directional light source (Figure 2a). Acceleration by GPGPUs is currently not possible due to the lacking support for the data-driven model in ACCELERAD.

A recently added modification of the data-driven model, `aBSDF`, interprets distinct peaks in the BSDF as ideal direct transmission. Its implementation concentrates all light transmitted through a region defined by the model’s resolution in its centre, and is therefore capable to model highly directional transmission of direct sun-light e.g. through fabric even with data-driven models of low or moderate resolution. This not only reduces the required number of random rays to adequately sample directional scattering BSDFs, but also overcomes limitations in the measurable resolution of CFSs that are caused by the interdependency of sampling aperture and apparent beam diameter in far-field gonio-photometry (Lee et al. 2018). Peak extraction has been demonstrated to achieve good results to model the visibility of the sun through shades, but effectively eliminates all information about the peak shape – given that such information is available e.g. from refined measurements – by reducing it to one direction. This may be acceptable for cases where directional transmission is limited to one sharp peak, but is problematic with complex distributions comprising multiple peaks and other distinct features.

1.2. Daylight simulation with the Radiance Photon Map

The recursive simulation of light propagation in ray-tracing leads to a tree of rays. These are commonly classified by a formalized ray notation as listed in Table 1 (Heckbert 1990; Veach 1997). Hybrid BRT as implemented in RADIANCE replicates numerous mechanisms of light transport occurring in buildings, lending itself to applications in lighting design, daylighting and building design, in the form $\boxed{\mathbf{E}(\mathbf{S}^*) (\mathbf{D}|\mathbf{G})*\mathbf{L}}$ for deterministic, and $\boxed{\mathbf{E}(\mathbf{S}^*) [\mathbf{D}|\mathbf{G}] (\mathbf{D}|\mathbf{G})+\mathbf{L}}$ for stochastic ray-tracing. However, both algorithms are not capable to account for primary or secondary caustics $\boxed{\mathbf{E}\mathbf{D}([\mathbf{S}|\mathbf{G}]+)(\mathbf{D}^*)\mathbf{L}}$ (Arvo et al. 2001). Photon Mapping (PM) is a bi-directional algorithm that addresses this limitation and allows to simulate light transport in optically complex scenes (Jensen 2001). Its integration in RADIANCE allows to model light redirection by non-planar reflectors and refracting structures that is not properly accounted for by BRT (Schregle, Grobe, and Wittkopf 2015). Light redirecting elements can be geometrically modelled as any other parts of the scene and, unlike the utilization of pre-computed BSDFs, do not require any pre-processing. Recent enhancements of the RADIANCE PHOTON MAP, such as its Out-of-Core (OoC) data-structure to store large amounts of photons, and the introduction of the CONTRIBUTION PHOTON MAP allow to employ the module in illuminance-based Climate-Based Daylight Modelling (CBDM) techniques (Schregle, Grobe, and Wittkopf 2016; Schregle et al. 2015; Schregle 2015; Bauer and Wittkopf 2016).

The PHOTON MAP module reflects the bidirectional nature of the algorithm by splitting the simulation into two separate passes (Schregle 2004). The forward *photon distribution* from the light source is implemented by the program `mkpmap`. *Photon gathering* is integrated into the core simulation tools of RADIANCE, namely `rtrace` for computing single pixel values or sensor response, `rpict` for image generation, and `rcontrib` for the generation of contribution coefficients as required in CBDM.

In the distribution pass, light-sources emit photons that are reflected or transmitted

Table 1.: Scattering phenomena defining the direction of subsequent rays with corresponding `rtype` values in `RADIANCE`, and the general notation of rays (Arvo et al. 2001).

Scattering	Direction	RADIANCE ray type	Symbol
none (primary)	from eye	PRIMARY	$\boxed{\text{E}}$
any	toward light source	SHADOW	$\boxed{\text{L}}$
regular reflection	mirrored incidence	REFLECTED	$\boxed{\text{S}}$
regular transmission	unchanged	TRANS	$\boxed{\text{S}}$
refraction	by Snell's law	REFRACTED	$\boxed{\text{S}}$
glossy reflection	perturbation of mirror	REFLECTED SPECULAR	$\boxed{\text{G}}$
forward scatter	perturbation of incident	TRANS SPECULAR	$\boxed{\text{G}}$
diffuse reflection	random	REFLECTED AMBIENT	$\boxed{\text{D}}$
diffuse transmission	random	TRANS AMBIENT	$\boxed{\text{D}}$

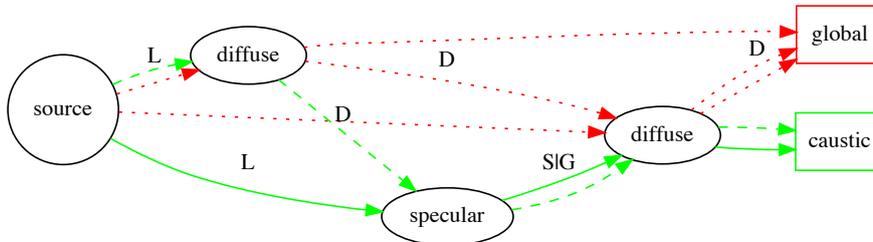
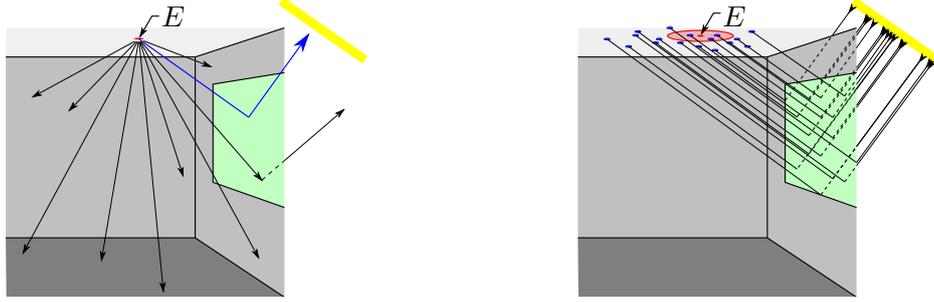


Figure 1.: Exemplary paths of global (dotted) and caustic (continuous and dashed) photons. Note that directionally scattering ($\boxed{\text{G}}$) and regular transmission or reflection ($\boxed{\text{S}}$) are handled identically as specular in the original PHOTON MAP, but not in BRT by `RADIANCE`.

by the scene geometry, and eventually deposited on diffusely scattering surfaces (blue in Figure 2b). After each collision with surfaces featuring a diffuse scattering component, photons are stored in the *global photon map* (dotted in Figure 1), or, for CBDM, in the contribution map for CBDM. Photons that collide with a diffuse surface after having been scattered by specular reflection or transmission are additionally stored in an optional *caustic photon map* (dashed and continuous lines in Figure 1). This photon map refines the representation of the direct caustic path $\boxed{\text{L}([\text{G}|\text{S}]+)\text{D}}$. Note that the path is reversed, when compared to BRT, and starts with $\boxed{\text{L}}$. The primary eye ray $\boxed{\text{E}}$ is not included in the distribution pass. Caustic photons are not only recorded after passing specular surfaces that are directly exposed to light sources, but also - as secondary caustics - if diffusely scattered photons are further scattered by specular reflection or transmission (e.g. $\boxed{\text{L}(\text{D}+)([\text{G}|\text{S}]+)\text{D}}$, dashed path in Figure 1).

In the photon gathering pass, the photon density is evaluated within an adaptive search radius (red in Figure 2b). Other than caustic photons, the global photon density is by default not directly visualized. Instead, to reduce noise, local ambient illuminance is evaluated as the integral of the photon densities reached by one indirect-diffuse re-



(a) BRT: Random sampling, only few rays (blue) that do not miss the light source contribute to the integral E .

(b) PM: After passing the BSDF, caustic photons (blue) are deposited. Photon density within a search radius (red) evaluates to E .

Figure 2.: Calculation of illuminance E on diffuse surfaces under directional illumination.

flection by stochastic sampling. An alternative visualization mode directly computes illuminance from the local density. While this direct visualization is faster, it significantly increases noise and bias, and thereby impacts the appearance of the generated image.

With the local diffuse illumination being solved by gathering of global and caustic photons, the view ray toward diffusely reflecting surfaces and subsequent aimed shadow rays $\overline{\mathbf{E}(\mathbf{S}^*)\mathbf{DL}}$, as well as the deterministic path $\overline{\mathbf{E}(\mathbf{S}^*)\mathbf{L}}$ are traced backward during image generation. To avoid double-counting of rays already accounted for by the photon map, all ray sequences $\overline{\mathbf{E}\mathbf{D}(\mathbf{S}^+)\mathbf{L}}$ must be eliminated in the stochastic backward sampling. This is implemented by the macro `srcRayInPmap(r)`. Unfortunately, since rays of types $\overline{\mathbf{D}}$ and $\overline{\mathbf{G}}$ are both handled by the stochastic scattering routines in RADIANCE, this effectively also suppresses the view-dependent ray sequence $\overline{\mathbf{E}\mathbf{G}(\mathbf{S}^+)\mathbf{L}}$, which is not represented by the photon map since photons are deposited only at $\overline{\mathbf{D}}$ path segments. While this behaviour achieves correct results for local illuminance, it introduces an error in image generation when surface exhibit directional scattering, e.g. are specular-glossy or translucent (Schregle 2016). Typical examples for this are the reflection of the sun on a glossy desk behind clear glazing (Figure 5), or forward-scattering by translucent objects toward the observer (Figure 6 and Figure 7). Consequently, the current implementation of the PHOTON MAP, while efficient in the sampling of data-driven models of high resolution, cannot be applied just to the problem where it could be most beneficial in current visual comfort research - the generation of imagery under sunny sky conditions.

1.3. Objectives

This research aims to enhance the applicability of the RADIANCE PHOTON MAP as a means to increase the efficiency of daylight simulation employing BSDF models of high directional resolution to image generation.

- The implementation is modified to account for directional transmission and reflection of scattered sun-light in image generation.
- Validity of the modified implementation is tested by comparison to BRT.
- The impact of the modification is evaluated by comparing glare ratings based on imagery, as generated by the original and the modified PHOTON MAP.
- A first benchmark compares the efficiency of the PHOTON MAP with that of

BRT when data-driven models are employed.

The command-line interface of RADIANCE is not affected by the proposed modifications. Therefore, the presented method to increase efficiency and accuracy in image generation with data-driven models, in particular of CFSs, can be immediately applied by expert users familiar with RADIANCE, and could be easily implemented in future releases of front-ends.

2. Method

2.1. Modification of the Photon Map implementation in Radiance

The modified implementation aims to employ deterministic BRT rather than PM for sources of known direction, e.g. when light sources are seen directly or through transparent layers, or by mirror-like reflection ($\overline{\mathbf{E}(\mathbf{S}^*)\mathbf{L}}$). Deterministic BRT is also applied when diffusely reflecting surfaces are lit directly, or through transparent surfaces without any intermediate scattering ($\overline{\mathbf{ED}(\mathbf{S}^*)\mathbf{L}}$) by non-extended light sources (types `light` or `spotlight` in RADIANCE). Note that $\overline{\mathbf{S}}$ applies only to regular transmission here. The shadow testing algorithm in RADIANCE would require *virtual light-sources* to account for regular reflection toward diffuse surfaces, therefore, the path ($\overline{\mathbf{ED}(\mathbf{S}^*)\mathbf{L}}$) leads to the deposition of a caustic photon in cases when $\overline{\mathbf{S}}$ stands for regular reflection. BRT toward known source directions, but with randomly jittered rays, is employed to account for forward-scattering of light-sources within the field of view $\overline{\mathbf{EGS}^*\mathbf{L}}$, which is actively suppressed in the original implementation of the PHOTON MAP for RADIANCE.

In all other cases, PM is employed to solve for local illuminance on diffusely scattering surfaces. Furthermore, PM is extended to the diffuse reflection of light emerging from extended sources of types `glow` and `illum`. The contribution of such sources was systematically underestimated in the original implementation of the PHOTON MAP.

To account for forward-scattering toward the observer, the macro `srcRayInPmap(r)` in `src/rt/pmapmat.h` is disabled. Its functionality is replaced by refined criteria when photons are deposited, so that double-counting is avoided.

Other than in the original implementation, regular and forward-scattered transmission are distinguished. In the case of regular transmission ($\overline{\mathbf{LS}^*\mathbf{D}}$), e.g. transparency, indicated by `r->rtype&TRANS`) without any preceding deflection, no photons are deposited (Figure 3). This case is accounted for by sending rays directly toward the light source, as in deterministic ray-tracing.

All directional scattering ($\overline{\mathbf{LG}^*\mathbf{D}}$) as `r->rtype&SPECULAR`, e.g. translucency, and deflection by refraction ($\overline{\mathbf{LS}^*\mathbf{D}}$) if `r->rtype&REFRACTED`), mirror-like reflection ($\overline{\mathbf{LS}^*\mathbf{D}}$) if `r->rtype&REFLECTED`), or regular transmission either following prior deflection or emission from an extended source is accounted for by caustic photons. Consequently, the modified macro `ambRayInPmap(r)` disables Gaussian sampling when ambient rays hit directionally scattering surfaces in the backward pass to avoid double-counting.

All described criteria for deposition of caustic photons are bundled in the new macro `CAUSTICFLAGS(r)` in `src/rt/pmapmat.h`. The same macro is called to decide on the deposition of a caustic photon, the blocking of backward rays to avoid double-counting, and to mark caustic contribution photons.

To account for the illuminance by extended sources of types `glow` or `illum`, e.g. the sky, *primary photons* are introduced. These are deposited in the case $\overline{\mathbf{LS}^*\mathbf{D}}$, e.g. on diffusely transmitting glazing, or on diffusely reflecting surfaces that are exposed to the source either directly ($\overline{\mathbf{LD}}$), or by regular transmission ($\overline{\mathbf{LSD}}$).

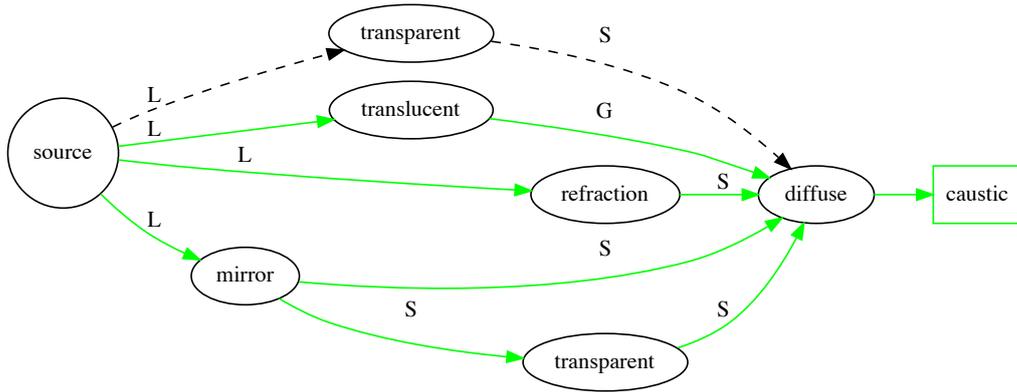


Figure 3.: Exemplary paths leading to deposition of caustic photons in the modified PHOTON MAP (continuous lines). No photon is deposited after regular transmission without prior deflection (dashed lines).

A patch to apply the proposed modification to a current release of RADIANCE is provided with the supplementary materials.

2.2. Testing validity for image generation

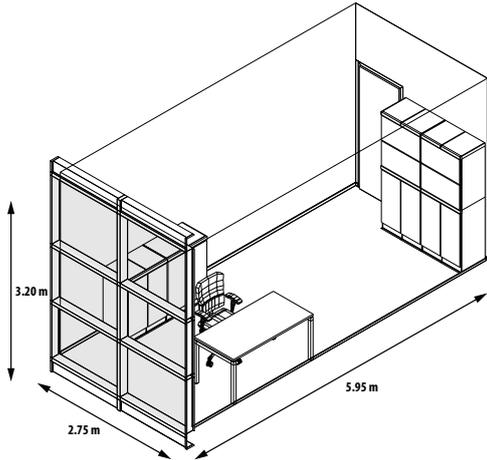
Visual comfort is evaluated in architectural context, which is typically characterized by a high degree of geometric detail, as well as different optical properties of materials and finishes of interior room surfaces, furniture, glazing, and shading systems. To account for this complexity, a detailed model of an exemplary, South-oriented office is employed to test the modifications (Figure 4a).

The façade of the office features a CFS, comprising a Laser Cut Panel (LCP), which is embedded in the glazing of the upper window zone. Clear glazing is applied to the remaining, lower window zones.

The LCP is expected to deflect incident sun-light toward the ceiling, and thereby to increase the depth of the daylighted zone. Figure 4b, right shows a sample provided for characterization and modelling. A data-driven model of the LCP is compiled from measurements on a scanning goniophotometer. The tensor-tree with 16 384 incident and 16 384 outgoing directions before data-reduction by 98% achieves a maximum resolution of $\approx 1.4^\circ$. The model is available as part of the supplementary material. The lower window areas are covered with clear glazing and are modelled by the `glass` model in RADIANCE.

Imagery for two exemplary views is computed. View *v1* takes the perspective of a standing occupant facing the façade. View *v2* corresponds to an occupant seated at a desk facing the Eastern side-wall of the office. The modified and the original implementation of the PHOTON MAP are employed for image generation, as well as BRT as a reference. For each simulation technique, approaches to accelerate the calculation are investigated.

For BRT, the CFS is modelled with and without *peak extraction*. This technique (enabled by the use of the `aBSDF` transmission model) partially replaces the compu-



(a) Model of the exemplary cellular office (Grobe 2017).



(b) Sample of the LCP.

Figure 4.: Office model and CFS as an exemplary test-case.

tationally demanding stochastic sampling of the transmission peak by deterministic shadow-testing. This may reduce artefacts of the ambient cache without increasing the stochastic sampling density. However, since the technique assumes regular transmission without any forward-scattering, it is expected that peak extraction will introduce an error in the image generation. BRT employing the data-driven model without peak-extraction (type BSDF) is therefore considered ground-truth in the scope of this work.

The PHOTON MAP supports two visualization modes of global photons. *Indirect photon visualization* estimates local illuminance by the evaluation of the global photon density at the first diffuse reflection of randomly distributed backward rays. This helps to reduce photon noise and bias as two typical artefacts of PM. *Direct photon visualization* skips the intermediary stochastic sampling and is therefore faster, but at the expense of more pronounced artefacts in the resulting imagery. Both visualization modes are tested for their impact on the appearance of the generated images, and on the results of visual comfort assessments.

To accelerate the simulations, the parallel image-generation program `rt pict` is extended to support PM, and is employed with PM as well as with BRT.

The parameters for the generation of the reference imagery are listed in Table 2 (column BRT). To achieve comparable results, the length of light paths is restricted to 4 reflections with BRT and PM. For the latter, this comprises three forward and one backward segment. Ambient accuracy and divisions are chosen to achieve acceptable image quality for a scene of moderate complexity. For BRT, a higher density of ambient rays is configured by refined settings of `-ad` and `-ad`, so that the data-driven fenestration model is sampled sufficiently. This is not necessary with the PM, when only the diffuse inter-reflection between the internal surfaces needs to be accounted for. The maximum ray weight was set to approximately $M^{-1} \cdot N^{-1} * 0.1$, with M being the number of specular super-samples and N being the number of ambient divisions.

Light simulation with the PHOTON MAP comprises two passes. Parameters for the photon distribution with `mkpmap`, and the subsequent photon-gathering and image generation by `rt pict`, are listed in Table 2 (column PM). Note that, compared to BRT, the parameters of the indirect-diffuse, or ambient, calculation are relaxed, and that at most one indirect-diffuse scattering event (`-ab 1`) is accounted for in the backward-

Table 2.: Parameters for image generation by BRT and PM.

Description	Parameter	BRT	PM
mkpmap:			
photon port modifier	<i>-apo</i> $\langle s \rangle$		portMat
global photon map, photon count target	<i>-apg</i> $\langle s \rangle \langle N \rangle$		g.pm, 1M
caustic photon map, photon count target	<i>-apc</i> $\langle s \rangle \langle N \rangle$		c.pm, 1M
inter-reflections	<i>-lr</i> $\langle N \rangle$		3
rtpict:			
reflections	<i>-lr</i> $\langle N \rangle$	4	4
ambient reflections	<i>-ab</i> $\langle N \rangle$	4	1 or -1
ambient accuracy	<i>-aa</i> $\langle k \rangle$	0.15	0.15
ambient divisions	<i>-ad</i> $\langle N \rangle$	1024	256
ambient subdivisions	<i>-as</i> $\langle N \rangle$	512	64
ambient resolution	<i>-ar</i> $\langle N \rangle$	64	48
maximum ray weight	<i>-lw</i> $\langle k \rangle$	2×10^{-5}	8×10^{-4}
specular super-samples	<i>-ss</i> $\langle N \rangle$	4	4
specularity threshold	<i>-st</i> $\langle k \rangle$	0.0	0
photon map, bandwidth	<i>-ap</i> $\langle s \rangle \langle N \rangle$		g.pm, 80
photon map, bandwidth min max	<i>-ap</i> $\langle s \rangle \langle M \rangle \langle N \rangle$		c.pm, 40, 400
pixel resolution	<i>-x</i> $\langle M \rangle$ <i>-y</i> $\langle N \rangle$	2048, 2048	2048, 2048

Table 3.: Glare classification based on thresholds applied to DGP and DGI.

	Classification	DGP	DGI
•	imperceptible	< 0.35	< 0.18
•	Perceptible	< 0.40	< 0.24
•	Disturbing	< 0.45	< 0.31
•	Intolerable	> 0.45	> 0.31

Table 4.: Photometric quantities computed from the imagery by `evalglare`.

Symbol	Unit	Description
\bar{L}	$\text{cd m}^{-2} \text{sr}^{-1}$	Average luminance of all pixels.
\tilde{L}	$\text{cd m}^{-2} \text{sr}^{-1}$	Median of the pixel values.
E_v	cd m^{-2}	Eye illuminance, cosine-weighted integral of pixel values.
\bar{L}_{src}	$\text{cd m}^{-2} \text{sr}^{-1}$	Average luminance of the detected glare sources.
Ω_{src}	sr	Solid angle of the glare sources.
\bar{L}_b	$\text{cd m}^{-2} \text{sr}^{-1}$	Average luminance of the background.

pass. Setting this parameter to a negative value (-ab -1) effectively suppresses the indirect-diffuse calculation entirely and triggers the direct visualization of global photons. A variable bandwidth of 40 to 400 photons is applied in the gathering of caustic photons to reduce bias in image regions where the illuminance gradient is steep.

For the quantitative comparison of the results of the PHOTON MAP calculations with the reference images, relative luminance difference is computed pixel-wise:

$$\Delta_r = \frac{|L_1 - L_2|}{L_1 + L_2} \cdot 2 \quad (1)$$

2.3. Exemplary glare assessments

To assess the accordance of the modified PHOTON MAP with the reference, and to estimate the impact of the limitations in the original PM implementation when employed in image-based visual comfort assessments, glare metrics are calculated by `evalglare` (Wienold and Christoffersen 2006) from the results of all three implementations. The glare metrics Daylight Glare Probability (DGP) and Daylight Glare Index (DGI), as well as the underlying photometric quantities (Table 4), are reported (Bellia et al. 2008). Glare sources are detected according to a fixed luminance threshold of $2000 \text{ cd m}^{-2} \text{sr}^{-1}$ (Pierson, Wienold, and Bodart 2018), and with the default peak detection. The metrics are classified according to a set of thresholds listed in Table 3 (Jakubiec and Reinhart 2012).

Due to the chosen solar geometry, the view to the sun is obstructed by the LCP. Sunlight cannot directly reach the view-point by transmission through the clear glazing, but by forward-scattering through the LCP and on the glossy surface of an office desk.

It is expected that the resulting highlights are accounted for in the glare evaluations based on BRT as well as by the modified PHOTON MAP, but not by its original implementation.

2.4. *Testing performance and sensitivity to model parameters*

The sensitivity of the performance impact to selected parameters is evaluated for two models.

The detailed office model, that is also used to test the validity of the method, provides high realism, since it represents a typical application of light simulation supporting visual comfort assessments. A direct comparison of simulation times is performed. Since both BRT and PM inevitably introduce artefacts, it is not possible to compare the simulation times for a defined optimum parametrization. Rather, a parametrization is proposed that can be justified as meaningful, and that is shown to achieve results what can be assumed to be acceptable.

A modified *Cornell Box* (Cornell University, Program of Computer Graphics 2002) is prepared as a simplified model to test the method’s performance with reduced complexity, and its sensitivity to selected model variables. The model does not provide the realism of a detailed building model, but supports the understanding of isolated phenomena and allows to relate artefacts to parameters. For this research, the light source of the original box model is eliminated. Instead, a South-oriented window featuring a data-driven model of the LCP as described before is introduced.

The simplified model allows to test the effect of the two variables controlling sampling of the data-driven BSDF with BRT and PM:

Ambient rays: With BRT, the accuracy of sampling, and its computational expense, are governed by the density of indirect-diffuse, or ambient, rays. This density is set by the `-ad N` parameter. To test the impact on image quality and rendering times, the parameter is varied $N = 128$ to $32\,768$ ¹.

Number of photons: Bias and noise as introduced by the PM algorithm decrease with increasing photon counts. To measure the impact on image quality and efficiency, the target counts for both global and caustic photons was varied $N = 10\,000$ to $2\,560\,000$ in the photon distribution pass. The photon bandwidth in the image synthesis pass was scaled accordingly²

The impact of these variables on image quality is measured by pixel-wise comparison of the resulting imagery to a reference in terms of Root Mean Squared Error (RMSE). The parametrization of the PHOTON MAP that is assumed to achieve highest accuracy is defined as reference. The elapsed simulation times are measured.

The performance gain by bidirectional PM is due to the concentration of samples on few, narrow light sources, where BRT would have to sample the entire hemisphere over the light-facing side of the BSDF to warrant for each light source’s contribution. This advantage fades away with increased solid angle - or number - of light sources. An extreme case would be the distribution of sun disks in CBDM, since each solar position would have to be accounted for by a sufficient number of photons. This effect is evaluated by varying the number of light sources, each with an angular diameter of 0.5° ,

¹The maximum ray weight, set by `-lw k`, was scaled by the inverse of the varied ambient rays N and the constant specular super-samples $M = 4$, so that $k = N^{-1} \cdot 0.2 \cdot M^{-1}$. Ambient accuracy was set to 0.1, and ambient super-samples up to half of the initial ambient rays allow to refine the calculation.

²The bandwidth of global and caustic photons was set to $N = \sqrt{0.008 \cdot M}$ with M being the target photon count according to (Schregle, Grobe, and Wittkopf 2016, p. 12, Table A.1).

from 1 to 64. First, one source is located at an elevation of 45° in front of the window. In the subsequent simulations, the sources are distributed over the South hemisphere (below and above the horizon) based on an equal-area subdivision algorithm (Shirley and Chiu 1997).

3. Results and discussion

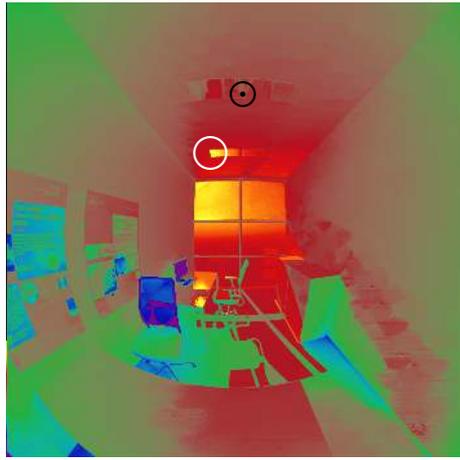
3.1. Visual inspection and comparison of imagery

Luminance maps corresponding to views $v1$ and $v2$, are shown in Figure 5 and Figure 6. A logarithmic scale is applied to account for the high dynamic range of the images. The position of the pixel of highest luminance is marked by a black dot surrounded by a circle in each image. For $v1$, the highest luminance after application of a simple Gaussian filter by `pfilt -r .7` is shown by a white circle, to eliminate effects of image resolution leading to individual pixels of very high luminance due to specular reflections of the sun.

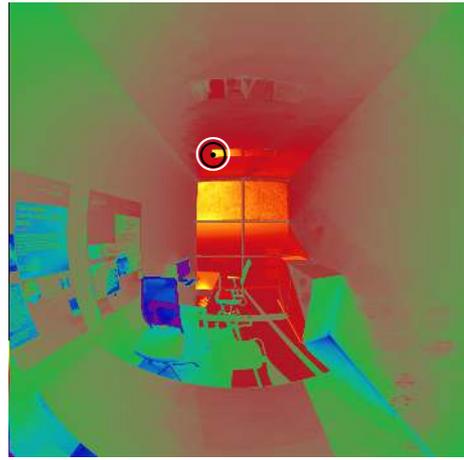
The results by BRT with and without peak-extraction achieve high accordance. The brightest pixel in $v1$ is located on the specular luminaire in the foreground in imagery by BRT (Figure 5a), but on the distant luminaire according to BRT with peak extraction (Figure 5b). After application of the Gaussian filter, the pixel positions agree. For $v2$, the brightest values are identified in the image area of the LCP (Figure 6a and Figure 6a). The stochastic sampling of the indirect-diffuse calculation introduces noise, and leads to the visible cloud-like artefacts of the ambient cache along the right side-wall in $v1$ (Figure 5a). This artefact is less apparent in $v2$ (Figure 6a). The effect of the concentration of the peak region in one singular direction becomes obvious in Figure 7, which shows a part of the upper window zone in $v2$. The forward-scattered light in a circular region centred at the direction toward the sun is assumed to be non-scattered and concentrated in few, very bright pixels, leading to a visible artefact.

The PHOTON MAP introduces visible noise in all generated images. Noise is more apparent when global photons are directly visualized, and in the case of the original PHOTON MAP, in image regions receiving sun-light through the clear glazing (e.g. parts of the desk in $v2$, Figure 6d and Figure 6d). Even with the indirect visualization of global photons, noise on the ceiling and the side-wall in the right half of the image is visible in $v1$ (Figure 5c Figure 5e). This is due to the high contribution of caustic photons triggered by directional transmission through the LCP. These photons are always directly visualized, and the resulting noise is not modulated by the reflective pass in the gathering of global photons.

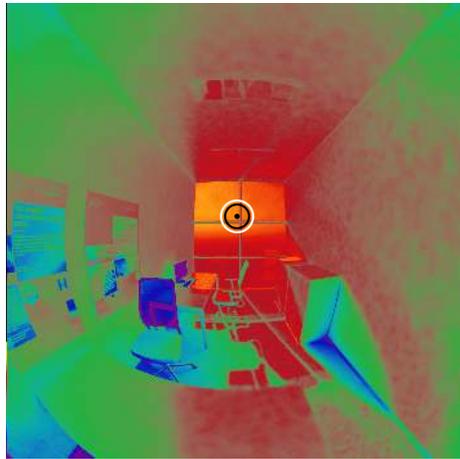
As expected, Figure 5c, Figure 5d, Figure 6c and Figure 6d show that the original PHOTON MAP does not account for the forward-scattering toward the observer, and thereby misses the most pronounced highlights. This applies to both views $v1$ and $v2$, and to both modes of photon visualization (indirect and direct). The missing visibility of the sun through the LCP is shown in detail in Figure 7. Due to the lack of pronounced highlights, the coordinates of pixels with highest L_{max} are inconsistent between Figure 5c and Figure 5d in $v1$, and between Figure 6c and Figure 6d in view $v2$. When global photons are directly visualized with the original implementations (Figure 5d and Figure 6d), regions not receiving sun-light are darker than in all other results. This can be explained by the missing contribution of diffuse sky-light, which is modelled as an extended source of type `glow` and is not accounted for in the original implementation, when indirect-diffuse sampling is suppressed in favour of the direct



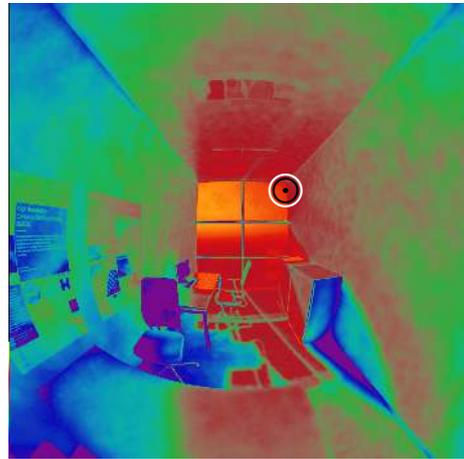
(a) BRT as reference.



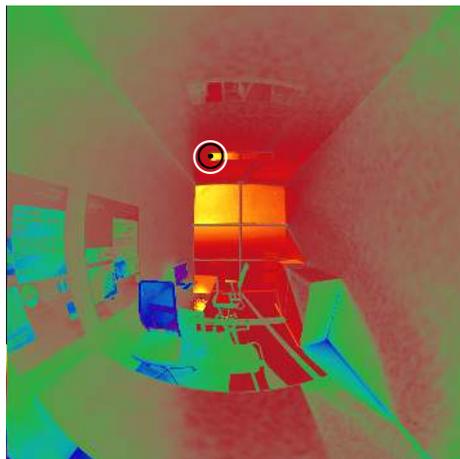
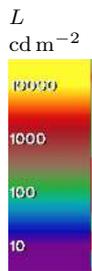
(b) BRT with peak-extraction.



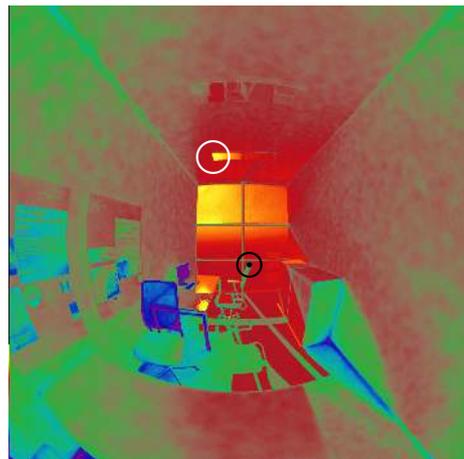
(c) Original PM, indirect photon visualization.



(d) Original PM, direct photon visualization.

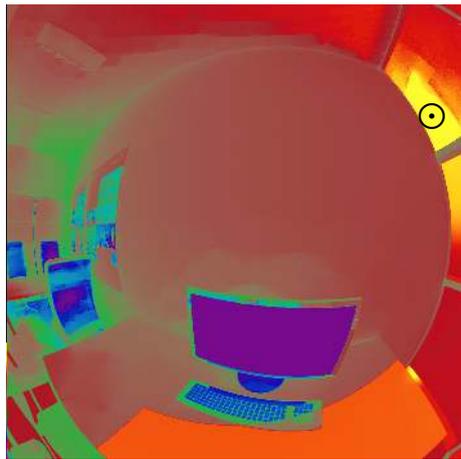


(e) Modified PM, indirect photon visualization.

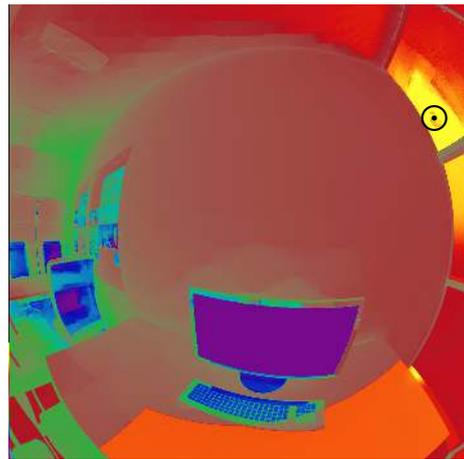


(f) Modified PM, direct photon visualization.

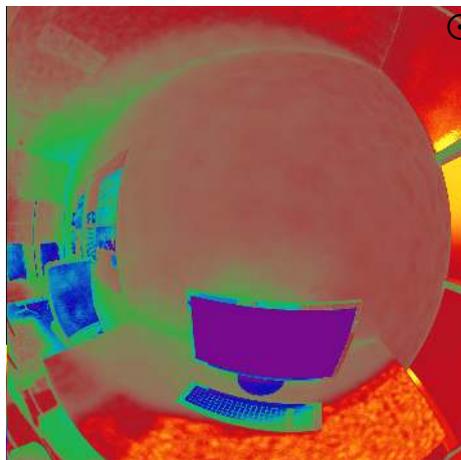
Figure 5.: Luminance maps for view $v1$.



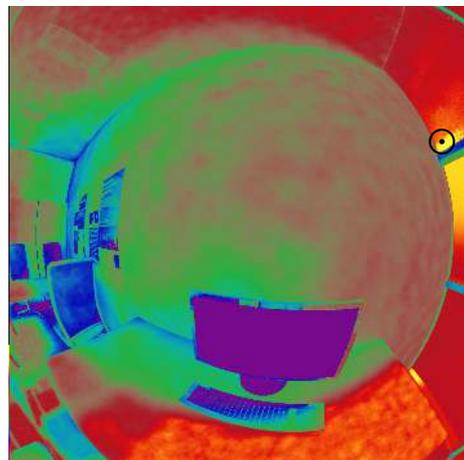
(a) BRT as reference.



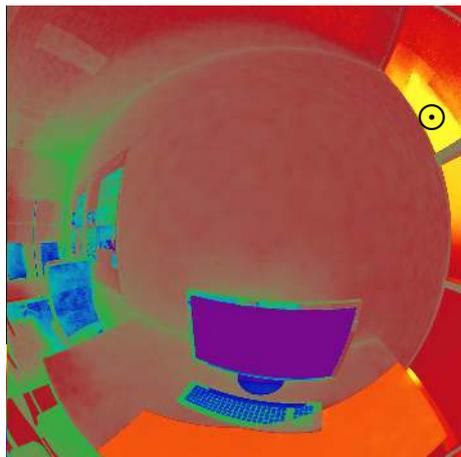
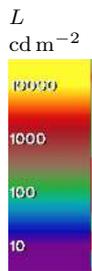
(b) BRT with peak-extraction.



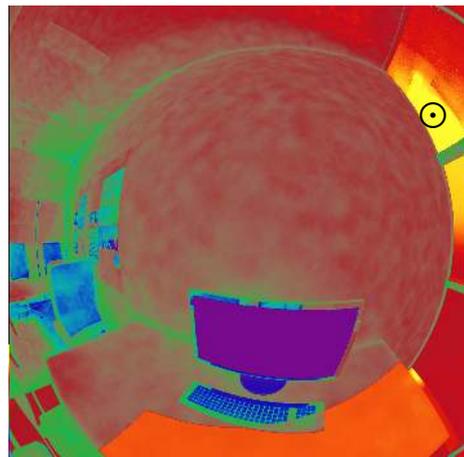
(c) Original PM, indirect photon visualization.



(d) Original PM, direct photon visualization.



(e) Modified PM, indirect photon visualization.



(f) Modified PM, direct photon visualization.

Figure 6.: Luminance maps for view v_2 .

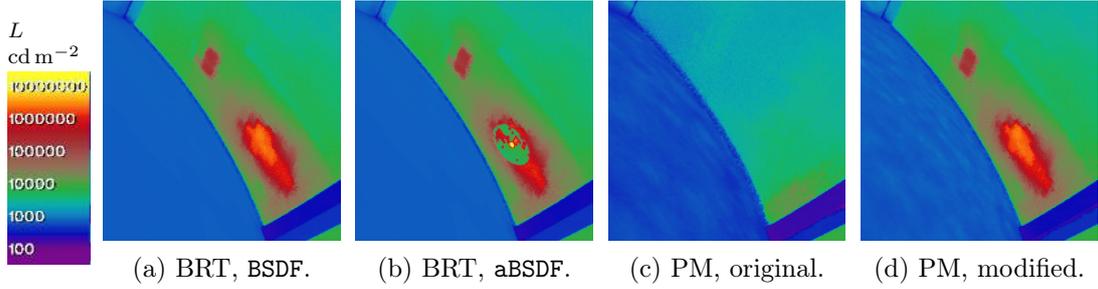


Figure 7.: Sun seen through LCP in view $v2$. Note the concentration of forward scattered transmission to one direction due to peak extraction (b), and its absence with the original PHOTON MAP implementation (c).

visualization of global photons.

The modified PHOTON MAP generally achieves good agreement with BRT. Figure 5f shows that the highest luminance in $v1$ is found in a highlight due to reflection on the furniture when global photons are directly visualized. This is inconsistent with BRT and indirect photon visualization, but compensated by Gaussian filtering. Other than the original implementation, the modified PHOTON MAP maintains sharp boundaries of the brightly illuminated regions of the desk in $v2$ (Figure 6e and Figure 6f). Due to the introduction of primary photons, the diffuse sky-light is accounted for. The highlight caused by the sun seen through the LCP agrees with the result of backward ray-tracing (Figure 7).

3.2. Quantitative comparison of imagery

The relative luminance (pixel) differences between imagery generated by the modified PHOTON MAP and BRT are shown in Figure 8 and Figure 9. Low-frequency photon noise is present in both views. If global photons are visualized indirectly, noise is reduced in regions with mostly diffuse illuminance (e.g. left side-wall in $v1$, Figure 8, which covers most of the background in view $v2$). Noise by caustic photons (e.g. right side-wall in $v1$) is not affected by the mode of photon visualization. Besides noise, bias is introduced by direct visualization of global photons, leading to linear edge-artefacts along the junctions of the inner surfaces of the room (Figure 8b and Figure 9b).

The differences between the results of the two PHOTON MAP implementations are reflected by the photometric quantities calculated from the imagery (Table 5).

For $v1$, Table 5 shows a high degree of accordance between the photometric characteristics calculated from imagery as generated by BRT and the modified PHOTON MAP. L_{max} is identical for BRT with peak extraction and the PHOTON MAP, but lower than that by BRT without peak extraction. The reason for this inconsistency is not the peak extraction mechanism, which does not effect the value of the identified pixel, but the attribution of L_{max} to a different pixel position in Figure 5a. The original implementation of the PHOTON MAP underestimates L_{max} , which is in fact attributed to arbitrary locations on the LCP due to the omission of specular reflections of the sun, as well as all other values. Yet, the deviation in terms of eye illuminance, average luminance and solid angle of glare sources, and background illuminance in $v1$ is only moderate.

The results for $v2$ in Table 5 confirm the good agreement of the modified PHOTON MAP and BRT without peak extraction. Since the L_{max} is attributed to identical

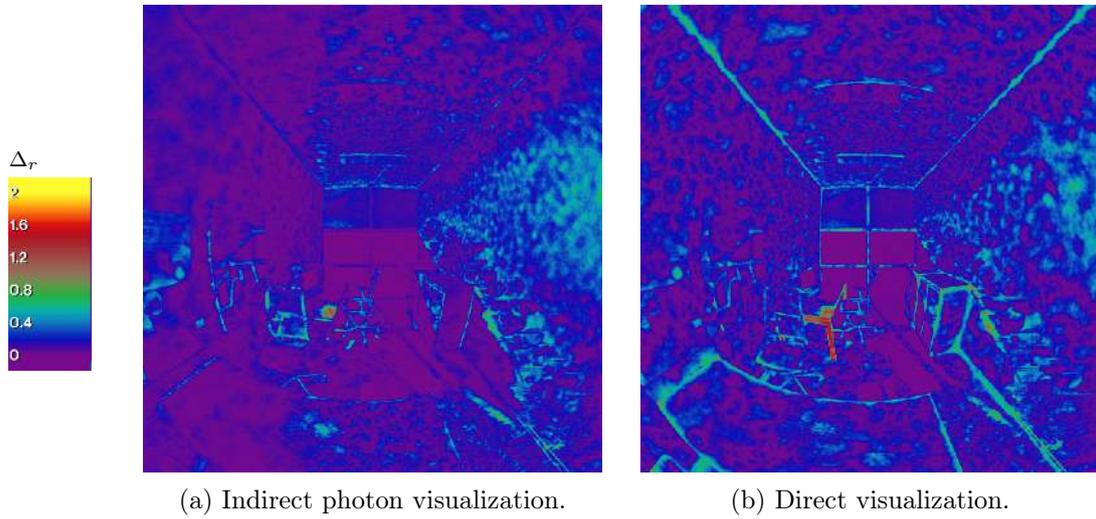


Figure 8.: View $v1$: Relative differences of results by modified PHOTON MAP and BRT.

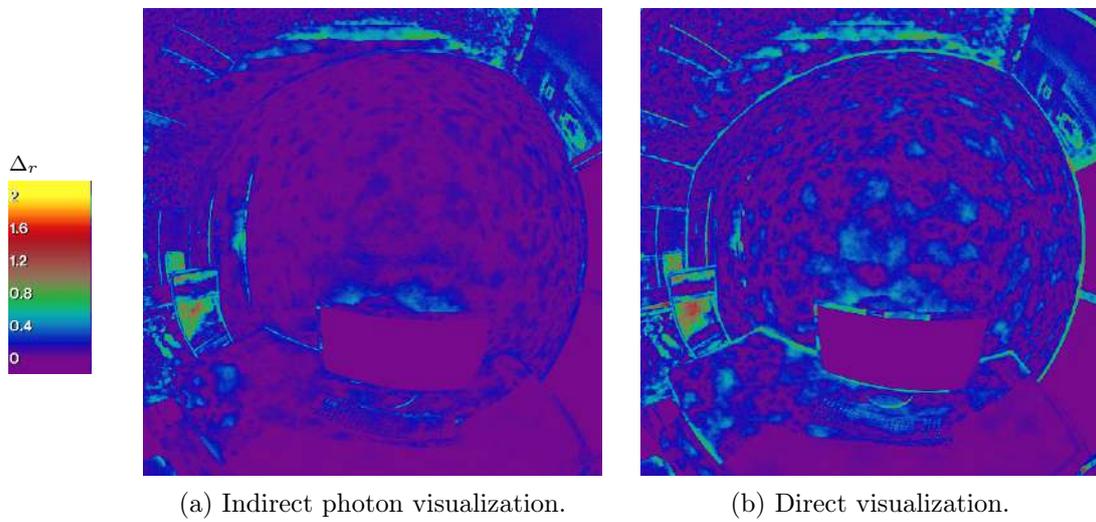


Figure 9.: View $v2$: Relative differences of results by modified PHOTON MAP and BRT.

Table 5.: Photometric quantities computed for views $v1$ and $v2$. Results employing peak extraction are given in brackets. The reference (BRT) is highlighted.

Implementation	\bar{L} [cd m ⁻²]	\tilde{L} [cd m ⁻²]	E_v [lx]	\bar{L}_{src} [cd m ⁻²]	Ω_{src} [sr]	\bar{L}_b [cd m ⁻²]	L_{max} [cd m ⁻²]
View $v1$:							
BRT	774	368	3471	3393	0.622	476	137114
	(716)	(336)	(3273)	(3229)	(0.620)	(445)	(107042)
PM, original	644	327	2872	2844	0.539	453	90395
PM, modified	786	413	3482	3387	0.607	495	107042
View $v2$:							
BRT	1786	547	1693	17530	0.478	455	3168300
	(1196)	(538)	(1625)	(9862)	(0.476)	(453)	(15358200)
PM, original	591	427	1236	3408	0.408	355	71421
PM, modified	1764	530	1634	19074	0.436	438	3168300

coordinates here, and the visibility of the sun through the LCP is exclusively covered by BRT in both implementations, identical values for L_{max} are returned. The effect of the concentration of forward-scattered sun-light by the peak extraction algorithm leads to an extremely high L_{max} . Since the peak extraction models a resolution beyond that achieved by the measurement, this value is, although plausible, impossible to verify by the means of the presented research. The unmodified PHOTON MAP fails short to predict any of the listed photometric quantities. This can be explained by the predominant role of forward-scattering of sun-light in $v2$, which is known to be not accounted for by this implementation.

The similar values of Ω_{src} predicted by BRT with and without peak extraction, and by the modified PHOTON MAP, indicate that the size of glare sources, which is – due to the moderate luminance threshold of 2000 cd m⁻² – not limited to the highlights caused by specular transmission and reflection, is only moderately affected by the rendering techniques. The original PHOTON MAP underestimates the size of the glare sources in particular in $v1$, since it does not account for glossy reflections of the sun in $v1$.

3.3. Results of exemplary glare assessments

The chosen glare metrics DGP and DGI extend the photometric quantities listed in Table 5 to models of human response. Results for views $v1$ and $v2$ are listed in Table 6.

For $v1$, the modified PHOTON MAP and BRT show a very high degree of accordance. This applies to both glare metrics, and includes the faster direct visualization of global photons. Peak extraction has a minor impact. The unmodified PHOTON MAP significantly underestimates DGP, but achieves good agreement for DGI.

The presence of the forward-scattered image of the sun in the field of view in $v2$ leads to almost identical results, when the glare metrics are calculated from imager by BRT and the modified PHOTON MAP. The mode of global photon visualization has no effect. The original PHOTON MAP fails to predict glare by omitting the predominant source. The impact on the luminance-driven DGI is higher than on DGP, a consequence

Table 6.: Glare evaluation, views $v1$ and $v2$. Results with acceleration techniques (aBSDF model, direct visualization of global photons) are listed in brackets.

Implementation	$v1$				$v2$			
	DGP		DGI		DGP		DGI	
BRT	●	0.40	●	22.68	●	0.38	●	18.76
(peak-extraction)	●	(0.38)	●	(22.37)	●	(0.33)	●	(17.34)
PM, original	●	0.36	●	21.85	●	0.24	●	11.87
(direct visualization)	●	(0.35)	●	(22.08)	●	(0.22)	●	(13.12)
PM, modified	●	0.40	●	22.51	●	0.38	●	18.82
(direct visualization)	●	(0.39)	●	(22.67)	●	(0.38)	●	(18.92)

of the drastic underestimation of the average luminance of glare sources as indicated by Table 5.

The application of a luminance threshold of $2000 \text{ cd m}^{-2} \text{ sr}^{-1}$ to the images identifies approximately identical image areas as glare sources, marked red in Figure 10 and Figure 11. Note that this spatial attribution is solely based on the threshold of 2000 cd m^{-2} , and does not quantify the luminance further. The exposure to direct sunlight causes the entire upper window areas in $v1$, as well as the lit area of the desk to exceed the threshold, so that specular scatter is not further distinguished from the diffuse transmission and reflection background. The deflection of light toward the ceiling is above the threshold in the results of backward ray-tracing and the modified PHOTON MAP, but not the original PHOTON MAP.

The exemplary application of glare metrics gives differentiated ratings for DGP and DGI in the view toward the façade $v1$. According to DGP, when based on the reference as well as the modified PHOTON MAP with indirect visualization of global photons, disturbing glare is expected. Based on DGI and identical imagery, on the other hand, glare is just perceptible. When peak extraction (with BRT) or direct visualization of global photons (with the modified PHOTON MAP) are employed, both metrics indicate perceptible but not disturbing glare. This inconsistency illustrates a problem of the application of thresholds – the seemingly contradicting ratings are based on values for DGP that are almost identical, but just in the range of the threshold. The original implementation, not accounting for the highlights, gives results at the lower threshold of the perceptible but not disturbing range, thereby giving the same rating as BRT with peak-extraction and the modified PHOTON MAP with direct photon visualization, although the absolute difference of the predicted DGP values is much higher. The rating for DGI is consistent for all generated imagery, indicating a low impact of the highlights in this particular view configuration.

The wall-facing view $v2$, with the forward-scattered sun covering an extended image region, shows a distinct difference between the reference and the modified PHOTON MAP on the one, and the original PHOTON MAP as well as BRT with peak extraction on the other hand. The former give practically identical results and predict perceptible, but non-disturbing glare. Peak-extraction leads to a clear underestimate of glare by both DGP and DGI – a surprising result, since the artefact due to the concentration of the peak leads to very high pixel values toward the sun, but reduces the average luminance of the detected light source. The original PHOTON MAP clearly underestimates glare

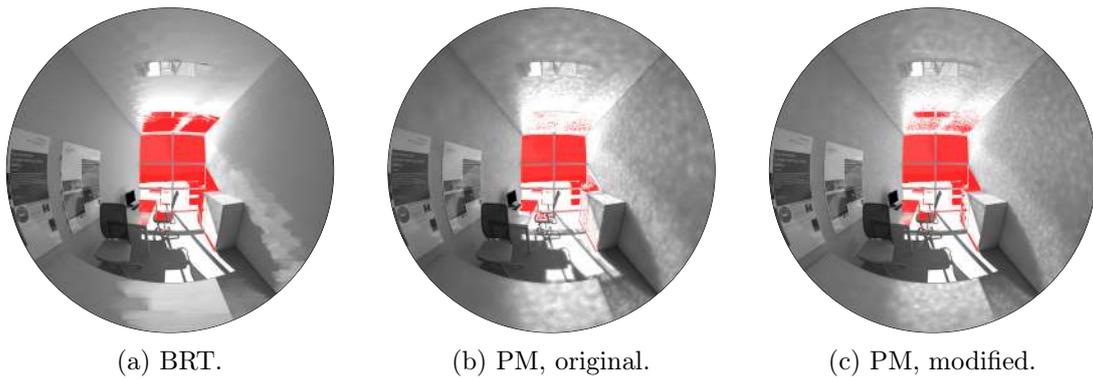


Figure 10.: Glare sources in $v1$ as detected by `evalglare`.

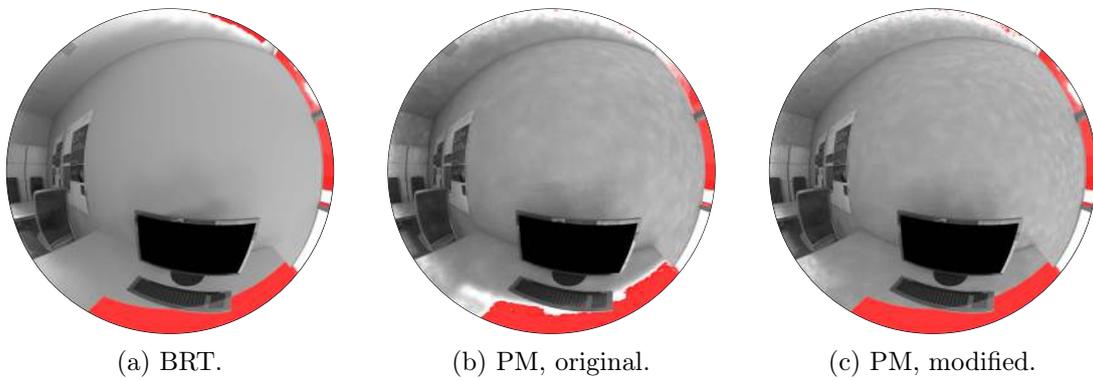


Figure 11.: Glare sources in $v2$ as detected by `evalglare`.

Table 7.: Duration of the simulation passes for the different simulation techniques, view $v1$. Results with acceleration techniques (aBSDF model, direct visualization of global photons) are listed in brackets. Pass 1: Forward distribution of photons with `mkpmap` (only with PM). Pass 2: Image synthesis with `rt pict`.

Implementation	Pass 1	Pass 2	Total
BRT	NA	52 800 s	52 800 s
(with peak extraction)	NA	(49 320 s)	(49 320 s)
PM, original	552 s	1692 s	2244 s
(direct photon visualization)	(552 s)	(1284 s)	(1836 s)
PM, modified	696 s	1824 s	2520 s
(direct photon visualization)	(696 s)	(1332 s)	(2028 s)

in $v2$, as expected, since it does not account for the predominant glare source by forward-scattering from the sun in $v2$.

3.4. Initial benchmark

Table 7 shows the simulation times for the imagery shown in Figure 5, with the parametrization as reported in Table 2.

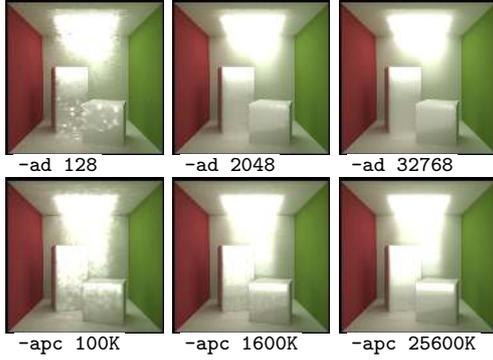
BRT without peak-extraction requires a very high amount of sample rays in the stochastic computation to sample the transmission peaks. Consequently, the computation time is ≈ 20 times higher than that of the modified PHOTON MAP. The purely stochastic sampling through the BSDF results in strong artefacts caused by the interpolation from the ambient cache (Figure 5a). Peak extraction, while incapable to replicate forward-scattering, reduces these artefacts and yet achieves a notable acceleration by 3480 s, or $\approx 7\%$ (Figure 5b).

The additional code triggering the disposal of caustic photons only after scattering events, that was introduced by the modification of the PHOTON MAP, only slightly increases the duration of the photon distribution pass when compared to the original implementation. The acceleration by direct visualization of global photons is limited to the image synthesis pass and reduces the simulation time by 492 s, or $\approx 20\%$.

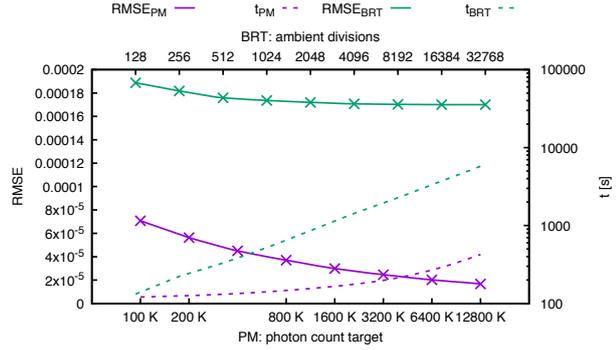
The effect of varying ambient divisions with BRT, and the target count of photons with PM, on the simulation results with the simplified Cornell Box model are shown in Figure 12 and Figure 13.

With one light source, RMSE moderately decreases with both simulation techniques, but is consistently lower for the evaluated parameter range with PM (Figure 12b). According to the tested parameter range, it cannot be expected that the two technique’s results converge. The imagery shown in Figure 12a reflects this result. Increasing the number of ambient samples as well as the target photon count clearly improves image quality by reducing artefacts of the ambient cache (BRT) and photon noise and bias (PM). However, BRT is not capable to replicate the sharp projection of the window onto the ceiling and the back wall, as simulated by PM.

With 64 sources, RMSE by BRT reaches a minimum at ≈ 2048 ambient divisions, and then starts to slowly increase (Figure 13b). This indicates that BRT converges not exactly with PM, which was employed to generate the reference image. With PM,

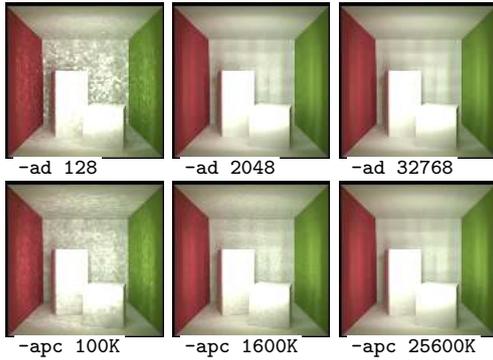


(a) Results by BRT (top), PM (bottom).

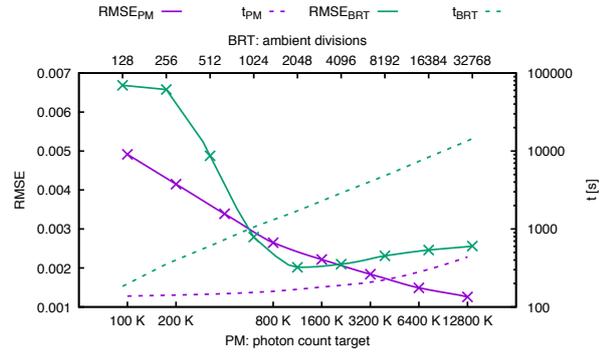


(b) RMSE, simulation times.

Figure 12.: Simplified model, varied simulation parameters and one light source.



(a) Results by BRT (top), PM (bottom).



(b) RMSE, simulation times.

Figure 13.: Simplified model, varied simulation parameters and 64 light sources.

on the other hand, RMSE decreases over the entire tested parameter range. With moderately refined parameters, the imagery by BRT and PM replicates the overlaid subtle projections of the light sources on the back wall, and the brightly illuminated boxes in the foreground.

Simulation times (Figure 12b, Figure 13b) increase approximately linearly with the number of ambient rays with BRT. With PM, exhibit an increasing slope with photon map size. This is due to the decreasing weight of the time spent on the backward ray-tracing computation, which is independent from the photon count. However, the measured simulation times are by far lower than for BRT over the entire tested parameter range.

4. Conclusion

In view of the increasing interest in luminance-based metrics for visual comfort, as well as the need to visualize the effect of CFSs beyond their functional aspects, the PHOTON MAP is a promising means to leverage the benefits of data-driven models in research as well as design and planning practice. The modification of the PHOTON MAP implementation for RADIANCE allows to employ the data-driven BSDF model in the rapid synthesis of imagery that is not only visually appealing – and lacking the typical artefacts due to indirect-diffuse sampling – but also valid for the evaluated case. This is considered a crucial step toward the wider acceptance of data-driven models as

a general means to replicate the irregular optical properties of CFSs.

Forward-scattering, that is not properly reflected by the original implementation of the PHOTON MAP but can contribute significantly to glare, is accurately replicated, e.g. when the sun or its specular reflections are in the field of view. This allows to employ the bidirectional algorithm to generate imagery for visual comfort and glare assessments. The original implementation, on the other hand, should not be applied in image-based visual comfort assessments, if any forward scattering may occur. Peak extraction to accelerate BRT or PM should be applied only with great care, and if the underlying assumption of ideal regular transmission holds true.

The modified PHOTON MAP achieves a high degree of accordance with BRT in glare assessments. In the case of the evaluated office model, simulation times were reduced to $\approx 5\%$ in an exemplary single point-in-time simulation. This result is not representative, since the efficiency depends on the model's geometry as well as the reflection and scattering models employed.

In particular, the required number of photons, that depends on the size and detail of the geometric model, has been shown to impact the efficiency of the technique. However, for the case of a moderately sized model and directional illumination, the possible relaxation of the ambient calculation offers an increase in efficiency that makes the application of high resolution, data-driven models more applicable. Accounting for diffuse sky-light by BRT, and employing PM only for direct sun-light, may be an optimization to constrain the size of the photon maps. In general, the use of the PHOTON MAP is beneficial in scenes characterized by directional light sources. Photon emission from extended sources, such as the sky hemisphere, does not promise any performance gain.

The direct visualization of global photons promises only minor acceleration compared to the indirect visualization by one reflection, but causes artefacts that may affect not the results of visual comfort assessments, but the perceived image quality and evaluations of aesthetic aspects of CFSs. Progressive PM (Schregle, Grobe, and Wittkopf 2015) can leverage the presented modification of the PHOTON MAP and may further improve image quality with data-driven models.

The PHOTON MAP lends itself not only to the simulation of light transport through geometrically modelled CFSs, but also to the efficient sampling of data-driven BSDF models showing strong directional behaviour. With the BRT algorithm of RADIANCE, such models can be only efficiently modelled if peak extraction reduces the directional peak to one direction. Peak extraction successfully reduces computation times, but effectively loses information of the close-to-peak distribution and can lead to wrong results if non-ideal forward-scattering occurs. However, in practice, the measurement of CFSs typically achieves a directional resolution that is too low to back such exact modelling of the peak shape. In such cases, the elimination of the peak shape from the model does not comprise a loss of information, but replaces one synthetic peak shape (defined by the resolution of the data-driven model) by another, the ideal regular transmission.

A potential future application of the PHOTON MAP is to replace the calculation of the direct sun component in the Five Phase Method (FPM) to reduce the complexity of the method. This would also allow to account for non-Lambertian reflection, that may contribute to visual discomfort but is currently omitted in typical CBDM evaluations. A manuscript presenting such a CBDM method is currently under preparation. Another potential application of the PHOTON MAP is the generation of imagery of models with static illumination, but for varying view positions and directions, since the result of the forward-distribution of photons is view-independent and can be reused.

The scalability of PM is limited since the required number of photons increases with the spatial dimensions of the model, when average photon density is kept constant. The OoC data-structure addresses this, but is not efficient with the indirect visualization of global photons. Progressive PM appears to be fields for future research to widen the scope of potential applications in daylight simulation.

Acknowledgement(s)

Thanks to my colleague Roland Schregle for his advice and guidance through the PHOTON MAP code, and to Greg Ward for discussing its modifications. Andreas Noback prepared Figure 4a. I am grateful to Tuçe Kazanasmaz and Stephen Wittkopf for their valuable feed-back on my PhD research at Izmir Institute of Technology. The LCP was kindly provided by Ian Edmonds and Chantal Basurto-Dàvila.

Disclosure statement

The author declares no conflict of interest.

Funding

This research was supported by the Swiss Innovation Agency Innosuisse under grant #1155000149 as part of the Swiss Competence Center for Energy Research SCCER FEEB&D.

References

- Arvo, James, Marcos Fajardo, Pat Hanrahan, Henrik Wann Jensen, Don Mitchell, Matt Pharr, and Peter Shirley. 2001. "State of the art in Monte Carlo ray tracing for realistic image synthesis." Los Angeles, US, Aug.
- Bauer, Carsten, and Stephen Wittkopf. 2016. "Annual daylight simulations with EvalDRC: Assessing the performance of daylight redirecting components." *Journal of Facade Design and Engineering* 3 (3-4): 253–272.
- Bellia, Laura, Arcangelo Cesarano, Giuseppe Fabio Iuliano, and Gennaro Spada. 2008. "Daylight glare: a review of discomfort indexes." In *Visual quality and energy efficiency in indoor lighting: today for tomorrow*, Rome, IT, Mar.
- Cook, Robert L, Thomas Porter, and Loren Carpenter. 1984. "Distributed ray-tracing." In *SIGGRAPH Computer Graphics*, Vol. 18, Minneapolis, US, Jul., 137–145. ACM.
- Cornell University, Program of Computer Graphics. 2002. "Cornell Box Data." Last accessed on 2019-07-15, <https://www.graphics.cornell.edu/online/box/data.html>.
- Geisler-Moroder, David, and Arne Dür. 2008. "Validation of Radiance against CIE 171:2006 and improved adaptive subdivision of circular light sources." 7th International Radiance Workshop, Fribourg, CH, Oct.
- Grobe, Lars Oliver. 2017. "Computational combination of the optical properties of fenestration layers at high directional resolution." *Buildings* (1): 22.
- Grobe, Lars Oliver. 2018. "Characterization and data-driven modeling of a retro-reflective coating in Radiance." *Energy and Buildings* 162: 121–133.
- Grobe, Lars Oliver, Stephen Wittkopf, and Zehra Tuçe Kazanasmaz. 2017. "High-resolution

- data-driven models of Daylight Redirection Components.” *Journal of Facade Design and Engineering* 5 (2): 101–113.
- Grynberg, Anat. 1989. *Validation of Radiance*. Technical Report. Lawrence Berkeley Laboratory.
- Heckbert, Paul Seagrave. 1990. “Adaptive radiosity textures for bidirectional ray tracing.” *SIGGRAPH Computer Graphics* 24 (4): 145–154.
- Heckbert, Paul Seagrave. 1991. “Simulating global illumination using adaptive meshing.” PhD diss., University of California, Berkeley, CA, USA.
- Jakica, Nebojsa. 2018. “State-of-the-art review of solar design tools and methods for assessing daylighting and solar potential for building-integrated photovoltaics.” *Renewable and Sustainable Energy Reviews* 81: 1296 – 1328.
- Jakubiec, J Alstan, and Christoph F Reinhart. 2012. “The adaptive zone – A concept for assessing discomfort glare throughout daylit spaces.” *Lighting Research and Technology* 44 (2): 149–170.
- Jensen, Henrik Wann. 2001. *Realistic image synthesis using photon mapping*. New York, US: AK Peters / CRC Press.
- Jones, Nathaniel L., and Christoph F. Reinhart. 2017. “Experimental validation of ray tracing as a means of image-based visual discomfort prediction.” *Building and Environment* 113: 131–150.
- Jones, Nathaniel L, and Christoph F Reinhart. 2019. “Effects of real-time simulation feedback on design for visual comfort.” *Journal of Building Performance Simulation* 12 (3): 343–361.
- Kazanasmaz, Tuçe, Lars Oliver Grobe, Carsten Bauer, Marek Krehel, and Stephen Wittkopf. 2016. “Three approaches to optimize optical properties and size of a South-facing window for spatial Daylight Autonomy.” *Building and Environment* 102: 243–256.
- Lee, Eleanor, Taoning Wang, Jacob C Jonsson, Greg Ward, Lars Oliver Grobe, Jan Wienold, and David Geisler-Moroder. 2018. “Generating high-resolution BSDFs for the direct beam component.” 17th International Radiance Workshop, Loughborough, UK, Sep.
- Lee, Eleanor S, David Geisler-Moroder, and Greg Ward. 2018. “Modeling the direct sun component in buildings using matrix algebraic approaches: methods and validation.” *Solar Energy* 160: 380–395.
- Mainini, Andrea Giovanni, Andrea Zani, Giuseppe De Michele, Alberto Speroni, Tiziana Poli, Michele Zinzi, and Andrea Gasparella. 2019. “Daylighting performance of three-dimensional textiles.” *Energy and Buildings* 190: 202 – 215.
- McNeil, Andrew. 2011. *On the sensitivity of daylight simulations to the resolution of the hemispherical basis used to define bidirectional scattering distribution functions*. Technical Report. DOE Technical Memo.
- McNeil, Andrew, Eleanor S Lee, and Jacob C Jonsson. 2017. “Daylight performance of a microstructured prismatic window film in deep open plan offices.” *Building and Environment* 113: 280–297.
- Molina, Germán, Waldo Bustamante, Jiwu Rao, Paul Fazio, and Sergio Vera. 2015. “Evaluation of Radiance’s genBSDF capability to assess solar bidirectional properties of complex fenestration systems.” *Journal of Building Performance Simulation* 8 (4): 216–225.
- Pierson, Clotilde, Jan Wienold, and Magali Bodart. 2018. “Daylight discomfort glare evaluation with Evalglare: Influence of parameters and methods on the accuracy of discomfort glare prediction.” *Buildings* 8 (8): 94.
- Rogers, Zack. 2013. *Enhanced skylight modeling and validation*. Technical Report. California Energy Commission.
- Schregle, Roland. 2004. “Daylight simulation with photon maps.” PhD diss., Universität des Saarlandes, Saarbrücken, DE.
- Schregle, Roland. 2015. *Development and integration of the Radiance photon map extension*. Technical Report. Lucerne University of Applied Sciences and Arts.
- Schregle, Roland. 2016. “Memorandum: Known bug in the Radiance photon map.” Unpublished Technical Note.
- Schregle, Roland, Carsten Bauer, Lars Oliver Grobe, and Stephen K Wittkopf. 2015. “Eval-

- DRC: A tool for annual characterisation of daylight redirecting components with photon mapping.” In *CISBAT 2015 Future Buildings and Districts Sustainability from Nano to Urban Scale*, Vol. 1, Lausanne, Sep., 217–222. LESO-PB, EPFL.
- Schregle, Roland, Lars Oliver Grobe, and Stephen Wittkopf. 2016. “An out-of-core photon mapping approach to daylight coefficients.” *Journal of Building Performance Simulation* 9 (6): 620–632.
- Schregle, Roland, Lars Oliver Grobe, and Stephen K Wittkopf. 2015. “Progressive photon mapping for daylight redirecting components.” *Solar Energy* 114: 327–336.
- Schregle, Roland, and Jan Wienold. 2004. “Physical validation of global illumination methods: measurement and error analysis.” In *Computer Graphics Forum*, Vol. 23, 761–781.
- Shirley, Peter, and Kenneth Chiu. 1997. “A low distortion map between disk and square.” *Journal of Graphics Tools* 2: 45–52.
- Veach, Eric. 1997. “Robust Monte Carlo methods for light transport simulation.” PhD diss., Stanford University, Stanford, CA, US.
- Ward, Greg, Murat Kurt, and Nicolas Bonneel. 2012. *A practical framework for sharing and rendering real-world Bidirectional Scattering Distribution Functions*. Technical Report. Berkeley, CA, US: Lawrence Berkeley National Laboratory.
- Ward, Greg, Murat Kurt, and Nicolas Bonneel. 2014. “Reducing anisotropic BSDF measurement to common practice.” In *Eurographics Workshop on Material Appearance Modeling*, edited by Reinhard Klein and Holly Rushmeier, Lyon, FR, Jun. The Eurographics Association.
- Ward, Greg, and Rob Shakespeare. 1998. *Rendering with Radiance*. San Francisco, CA: Morgan Kaufmann Publishers.
- Ward, Gregory, Richard Mistrick, ES Lee, Andrew McNeil, and Jacob Jonsson. 2011. “Simulating the daylight performance of complex fenestration systems using bidirectional scattering distribution functions within Radiance.” *Leukos* 7 (4): 241–261.
- Ward, Gregory J. 1994. “The Radiance lighting simulation and rendering system.” In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, New York, NY, USA, 459–472. ACM.
- Whitted, Turner. 1980. “An improved illumination model for shaded display.” *Communications of the ACM* 23: 343–349.
- Wienold, Jan, and Jens Christoffersen. 2006. “Evaluation methods and development of a new glare prediction model for daylight environments with the use of CCD cameras.” *Energy and Buildings* 38 (7): 743–757.
- Wymelenberg, Kevin Van Den, and Mehlika Inanici. 2014. “A critical investigation of common lighting design metrics for predicting human visual comfort in offices with daylight.” *Leukos* 10: 145–164.
- Wymelenberg, Kevin Van Den, and Mehlika Inanici. 2016. “Evaluating a new suite of luminance-based design metrics for predicting human visual comfort in offices with daylight.” *Leukos* 12: 113–138.