Raytraverse: Navigating the Lightfield to Enhance Climate-Based Daylight Modeling

Stephen Wasilewski
stephen.wasilewski@hslu.ch
HSLU and EPFL
Switzerland

Lars O. Grobe
HSLU
Switzerland

Roland Schregle
HSLU
Switzerland

Jan Wienold
EPFL
Switzerland

Marilyne Andersen
EPFL
Switzerland

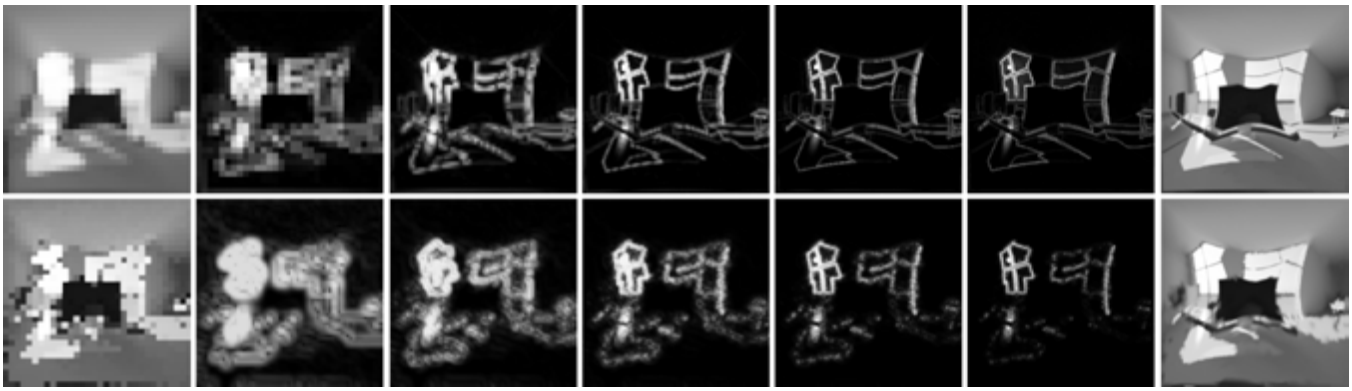

Figure 1: Wavelet decomposition (top row) and matching detail estimation generated during the sampling process (bottom row). Image on left is the lowest level estimation / starting resolution. Image on right is the original image / sampled reconstruction

ABSTRACT

We propose a new method for climate-based daylight modeling (CBDM) based on simulating and evaluating only the most important features. By adaptively sampling the temporal lightfield that describes daylight in buildings, our method escapes the curse of dimensionality that binds typical approaches. The method is centered around an iteratively guided sampling strategy that produces unordered sky and sun coefficient samples throughout an architectural space. This data is organized into a lightfield data structure that can be evaluated under any sky condition. Outputs include photometric quantities, visual comfort metrics, and high dynamic range images. Unlike existing CBDM methods, our approach is not limited by the rendering algorithm, material model, level of detail, or metric output. Our method is accurate and fast. As such it can serve as a general interface for any view or point-based daylighting question, including those answered by illuminance and luminance based metrics.

SimAUD 2021, April 15–17, 2021, Online

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

KEYWORDS

Daylight, Visual Comfort, Glare Analysis, CBDM, Climate Based Daylight Modeling, Lightfield

ACM Reference Format:
Stephen Wasilewski, Lars O. Grobe, Roland Schregle, Jan Wienold, and Marilyne Andersen. 2021. Raytraverse: Navigating the Lightfield to Enhance Climate-Based Daylight Modeling. In *2021 Proceedings of the Symposium on Simulation in Architecture and Urban Design.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Simulating daylight for visual comfort typically serves one of two purposes: image-based evaluations and sensor-based evaluations (usually as a cosine weighted irradiance). Naturally, this follows from how physical space is measured, with high dynamic range images or sensors behind diffusers. For analyses spanning a large spatial-temporal scope, an image based analysis is often impractical due to computation and storage concerns. A recent review of simulation methods for annual glare assessments [19] found that this practicality has led to a gap between subject base research, which suggests that the distribution of light in the visual field is important, and modeling, which is dominated by simplified illuminance based calculations. Additionally, a sensor based analysis offers limited information and, without careful consideration of the rendering techniques employed, can fail to capture important contributions to the scene (such as specular reflections from nearby

buildings). In this paper, we propose a new method, called *raytraverse*, for sampling and evaluating daylight rendering data that bridges this gap between the demands of photorealistic rendering and the limitations of sensor-based analysis. Based on an adaptive sampling technique using wavelets and reconstruction with view origin based kd-trees, *raytraverse* provides a general solution for daylight simulations capable of evaluating large spatial temporal domains through both subjective image assessment and calculation of objective daylight metrics.

A simulation and analysis workflow can be described in four parts: *modeling*: how geometry, material, and light source are represented; *sampling*: how the analysis dimensions are subdivided into discrete points to simulate; *rendering*: how light is propagated through the model; and *evaluating*: how the data are represented for interpretation. Advances in simulation methods can tackle any of these parts, singularly or in combination. Our method is primarily focused on sampling, which has also necessitated the development of a new means of evaluating the data. The following paragraphs outline the important implications sampling has for daylight simulations, and, briefly, the impact existing methods focused on modeling, rendering, and evaluating have on sampling.

Because the sky is always changing, daylight is never a fixed quantity. Building inhabitation is not fixed either. In total there are seven dimensions needed to describe daylight in a building. These are: position (3 dimensions), direction (2 dimensions), wavelength, and time. The description of this space is called a lightfield, or since it is also changing over time, a temporal lightfield. In this paper, we refer to it as either a daylight field, or if we are dealing with a particular source component, a sky lightfield or a sun lightfield. In order to make the problem tractable, any daylight simulation, whether for informing design or documenting performance, includes modeling a representative sky or skies and selecting representative locations within a building. The range of approaches for selecting sky conditions include: single source approaches like the daylight factor, a few site specific sky conditions taken across seasons, annual sky conditions from hourly typical meteorological year data, and site specific data captured at sub-hourly intervals. While increases in computing power have tended to lead analyses toward including more sky conditions, full resolution image generation, whether for subjective evaluation, presentation and/or incorporation in VR studies [3], still typically selects representative sky conditions due to time constraints, disk storage, and the practical consideration of looking at so many images. These different models of incident daylight represent the temporal dimension, the magnitude of which varies from one timestep to thousands of timesteps.

Similarly, for some daylight simulation questions, a carefully considered single point is taken as representative of a particular space or situation. In other cases, a grid of points at workplace or eye height across an entire building may be warranted. Using the IES LM-83 standard for daylight autonomy [5] as an example, the enumeration of the daylight field would be as follows:

- 25 points per 100ft^2 (10m^2): 2' x 2' (0.61 m) grid spacing
- 1 direction: an integrated horizontal illuminance
- 1 band: photopic luminance
- 3,650 sky conditions: one year of hourly data for ten hours per day
- 91,250 total measurements per 100ft^2 or 365 kilobytes of data

In this type of analysis, good positional and temporal resolution is preserved at the expense of directional detail. On the other hand, an annual glare analysis for a single point, such as the enhanced simplified daylight glare probability (eDGPs) [20]) favors directional resolution to measure contrast, but does not extend across positions:

- 1 point: typically a workstation view
- 786,000 directions: a 1000x1000 pixel angular fisheye (excluding corners)
- 1 band: photopic luminance
- 3,650 sky conditions: one year of hourly data for ten hours per day
- 2.9 million total measurements per point or 11 gigabytes of data

Even a single point in the daylight field requires a significant amount of data when rendering at a sufficient resolution to accurately evaluate visual comfort in all conditions. Combining the resolution of both approaches, an uncompressed high-resolution daylight field would require 550 gigabytes of data per 100ft^2 stored either as a single luminance float or in the RADIANCE HDR file format (which provides three color channels).

Existing CBDM approaches handle this data constraint in a number of ways. Key among these is evaluating the contribution from the sky as a set of discrete patches. Originally proposed by Tregenza and Waters [15], the sky-patch based model has been widely adopted and further developed in methods proposed by Mardaljevic [9] and by Reinhart and Herkel [12]. This model transforms the temporal dimension of the daylight field into a pair of angular dimensions corresponding to coordinates on the sky-dome. These coordinates are discretized into a fixed number of sky patches. Regardless of how many time-steps are evaluated, the simulation data is bound by this number of sky patches.

Building on the sky coefficient technique, there are also a series of multi-phase approaches to daylight coefficient simulations. The 3-phase method [10] separates the simulation between a sensor and the window and between the window and the sky. The primary limitation of this approach is the required use of a discretized material model, which does not represent specular transmission well. The 5-phase method [8] circumvents this requirement by simulating the direct sun as a separate component. The sun is transmitted through the facade using the normal rendering process, accounting for specular transmission and reflection. Two sources of bias from the 3-phase method remain: it does not include light reflected off of the window, and the indirect contribution of the sun is still based on the low resolution window model. While not a problem in simple scenarios, this may not be suitable in all cases.

Daylight coefficient and matrix/phase approaches are changes to the model within a simulation workflow that have direct implications on the sampling. Advancements in rendering techniques, including utilizing GPU hardware [7] or changing the algorithms used for light propagation, such as the extension of photon-mapping to daylight coefficient methods [4], do not directly impact sampling, but can change the calculus of how much can be sampled. In the case of photon mapping, sampling at a higher spatial resolution is less costly than increasing the number of timesteps. Other methods change the means of evaluation to limit the required data, but

have limited applicability. The imageless daylight glare probability (iDGP) method [6] uses a matrix based light propagation to directly assign sky patches as glare sources, but it is only valid for direct specular transmission in absence of scattering and glare from reflected sources. Santos and Caldas [13] introduce a heuristic approach based on illuminance that cannot account for discomfort glare in low-light high contrast scenarios.

Another way to manage the data is to only store computed metrics and discard the full resolution at an intermediate step. However, without the full data (and the ability to visualize it), it is difficult to find mistakes and make reproductions. These are important considerations for research and consulting. To respond to them, *raytraverse* aims to satisfy the following conditions:

(1) provide an interface for virtual photometry [1]. In other words, enable the measurement of any lighting based quantity, including the formation of an image, from anywhere within a space.
(2) be compatible with any rendering technique that accurately captures light transport within an architectural space, regardless of material model and light distribution.
(3) minimize computation time without sacrificing accuracy or introducing additional biases not present in the underlying rendering engine.
(4) minimize storage requirements without sacrificing the efficiency of evaluation.
(5) keep all simulation data available for exploration and iterative evaluations without additional computation.
(6) limit the introduction of additional assumptions regarding light transport above those set by the rendering algorithm.

The next section provides an overview of the *raytraverse* method before demonstrating the accuracy of the method compared to reference simulations under a range of typical and challenging conditions.

2 METHOD

Compression works by transforming a signal into an alternate basis where the signal can be represented sparsely: many of the coefficients are zero, or in the case of lossy-compression, many of the coefficients can be treated as zero without losing meaningful information. Once encoded in a sparse domain, the zeros can be discarded, and only the meaningful coefficients (plus some way of locating them) are stored. The signal can then be recovered through the inverse transform. Compressive sensing asserts that, subject to certain formal conditions, if a signal can be represented sparsely then there is a way to accurately measure the signal with far fewer samples than would be necessary with conventional sampling [2]. While our method does not assert the formal conditions required by compressive sensing, the motivation and relationship to compression is similar. Our method attempts to characterize the daylight field with the least number of samples possible to ensure a desired level of accuracy.

Commonly used in image compression, the wavelet transform also has demonstrated effectiveness for sampling applications in computer graphics [11, 17]. The top row of Figure 1 illustrates part of a two dimensional wavelet transformation. A multilevel wavelet transformation separates the signal into a low resolution average of the signal, which in the case of a full decomposition of a view could be described as the average luminance, and levels of detail coefficients. These detail coefficients describe, at higher and higher resolutions, the variance across the view, where higher levels correspond to higher frequency changes. Figure 2 provides an overview of the method, outlining how this iterative sampling approach is used to produce recoverable luminance information.

Wrapped around this general sampling approach, our method separates the sky and sun coefficients, for both the sampling and integration steps. This is done to leverage the same efficiencies gained by existing CBDM methods, utilizing sky patch contributions. But in our case there is an added benefit: the sky sampling can be used to guide the sun sampling to discover specular reflections that the iterative wavelet sampling may miss, which is explained in more detail in the next section.

2.1 Iteratively Guided Sampling

In simulation, we do not have access to these coefficients *a priori*, so the inverse transform is replaced by additional sampling. Arranged like a multi-level wavelet transformation, we iteratively sample through the levels of the decomposition beginning with a low resolution estimate of the average coefficients (see Figure 2). Based on this estimate, a wavelet transformation estimates the detail coefficients shown in the bottom row of Figure 1, which are used as a probability distribution to choose additional samples. The projection used in Figure 1 is the Shirley-Chiu square mapping from a sphere [14]. This mapping is used because it preserves solid angle, a requirement for unbiased sampling; and preserves adjacency, a requirement for generating meaningful detail coefficients. For full spherical sampling, samples are drawn from a double square, mapping to two paired hemispheres. Subsequent figures are shown in the more familiar angular fisheye projection. At each wavelet frequency the sampling density is highest wherever the variation in that frequency is greatest, which will closely match the distribution of the "true" detail coefficients. Figures 3a and 3b confirm that the sampling can recover a similar set of details to what the compression saves. Note that in the sampled image, the sample positions are jittered within the grid at each decomposition level, while the reference decomposition shows coefficients at grid centers.

The sampling is executed in three steps and the results are combined at integration time. The first step samples the sky coefficients, where sampling densities are driven by the scene geometry, with most samples taken at edges and material boundaries, particularly around windows.

The second step samples the ambient and reflected contribution from the direct sun. Unlike the 5-phase method, this is a full depth simulation including scattered sunlight. This is tractable because most samples are aimed at direct paths to the sun and sky, leading to early ray termination. The details uncovered in the sun sampling step include the direct patches of sun within a space and semi-specular reflections and transmission. The wavelet guided sampling will fail wherever a subsequent level of detail is not well estimated from its antecedent. This includes views to the direct sun and perfectly specular reflections of the sun. To check for these specular reflections, we use the sky sampling to guide a level of the sun sampling, as illustrated in Figure 3c.

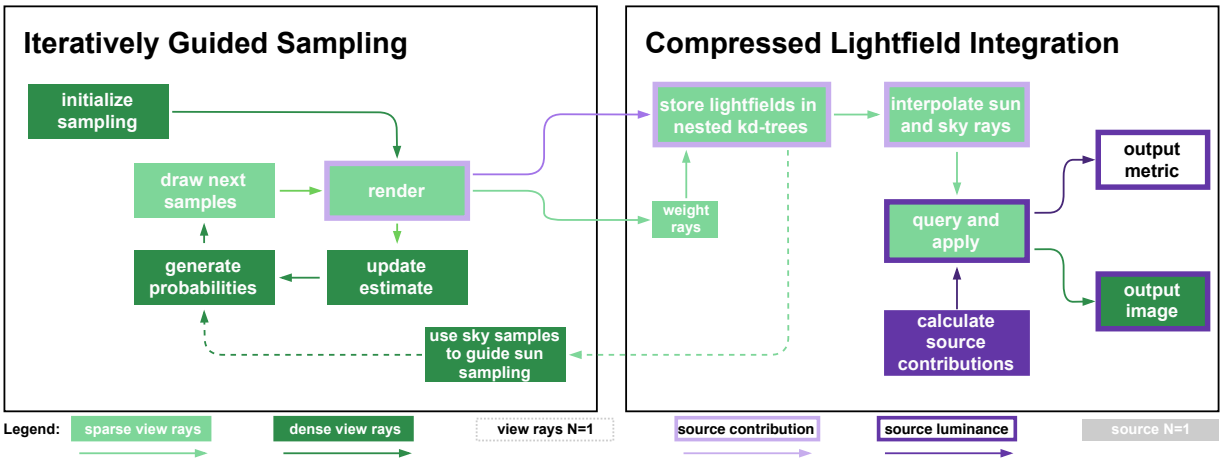

Figure 2: Overview of the *raytraverse* process. Box fills indicate whether the step operates on a sparse or dense set of view rays. The box *output metric* has no fill color because the data has been reduced over the view rays axes. Box outlines indicate whether sources are by contribution or luminance. Boxes without an outline either do not track source information or have source data that is a single value (such as an average or maximum).

The final step is to check for direct visibility of the sun. Directly visible sun has a huge influence on visual comfort and daylight levels and these values can be extremely sensitive in cases where the sampling resolution is not high enough, especially if the sun is partially occluded. Because we know the view direction of the sun we are simulating, we can run the same iterative sampling process used for the entire view for a small view cone aimed directly at the sun. The results of these sun coefficients can then be added to the sky and reflected sun contribution at integration.

The result of this process is a compressed representation of a high resolution lightfield. Sampling density is highest where contrast is high and lowest where contrast is low. This can substantially reduce simulation time, but a new means of evaluating the sampled data is necessary. The next section provides a brief overview of how we prepare this data for fast integration.

2.2 Compressed Lightfield Integration

The output from the sampling/rendering process is unordered data containing view rays and source contributions from points throughout the analysis area. While the rendered data was produced iteratively, the output data is all read in together at the end of the sampling process. Figure 2 outlines the lightfield construction and evaluation. Sample vectors are sorted into kd-trees [16], with one tree for each sensor point in the analysis area. A kd-tree provides

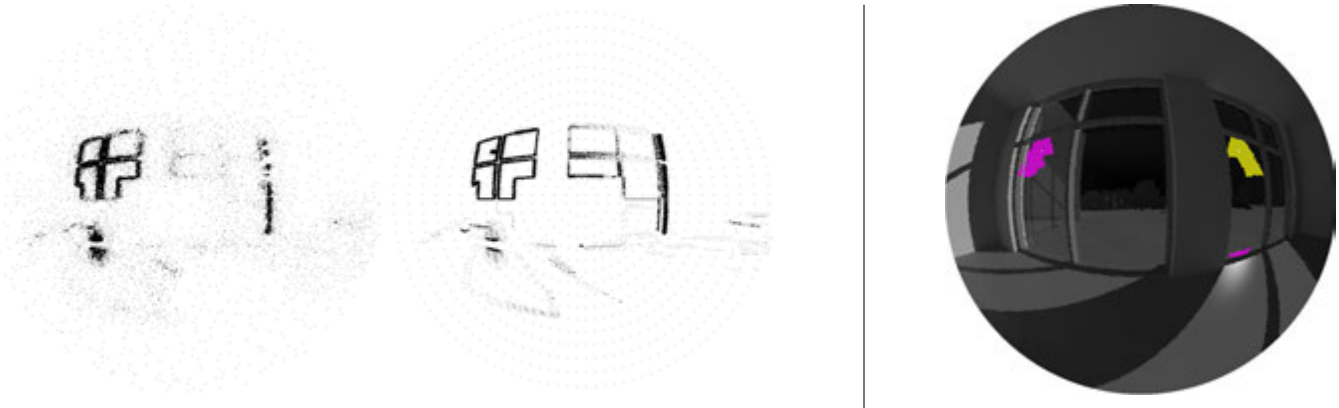

(a) (b) (c)

Figure 3: (a) The sample positions of the 11,800 largest wavelet coefficients (matching the number of samples in (b)). (b) Sample positions drawn by the complete simulation process applied to a full sky description (sun and sky). (c) Specular source sampling is guided by the initial sky patch sampling. The contributions from the sky patch in which the sun resides (yellow) are cropped to isolate purely specular reflections (magenta). These regions are sampled at a resolution high enough to catch the small direct source. Note how even the small area of the mullion is captured by the pre-sampling because the target source is large.

fast queries for spatial data. Each kd-tree represents the 360 degree view from a point. Each node in the tree has a direction vector, daylight coefficients for a sun or sky-patches, and an angular subtense, calculated from the Voronoi region of the node.

Similar to existing daylight coefficient methods, the coefficients are multiplied by a sky matrix that represents each time-step in an annual evaluation. Because the number of vectors is a small subset of a typical image resolution, this step requires much less computation time than operating on a dense matrix.

3 ACCURACY AND EFFICIENCY

To check the performance of our method, a range of lighting conditions were rendered using *raytraverse* and compared against high resolution reference renderings, low resolution renderings that match the sampling rate of the *raytraverse* sun sampling, and illuminance based simulations. Reference renderings made with RADIANCE [18] are used instead of physical measurements because what we wish to validate is the sampling and integration workflow, not the rendering engine itself. For this reason, where possible, the simulation parameters used in the test runs match those of the reference images.

The scenes include typical building simulation scenarios, such as translucent skylights, clear windows, and fabric rollershades. In addition, there are conditions that were specifically selected because they are typically quite difficult for low resolution sampling to properly characterize. These include views to partially obscured direct sun, tightly spaced louvers, frit, and specular reflections. The difficult conditions are in general beyond what can be accurately simulated by existing CBDM methods and would thus require single point in time renderings. One of the objectives of this accuracy test is to see whether *raytraverse* is able to adequately circumvent this limitation.

Figure 4 describes the scene. In each condition, only one set of apertures is present and there are a variety of materials and treatments applied to the apertures. Apertures are grouped by: South, East/West, and Skylights. Table 1 enumerates the configurations. For each of the 56 simulated views and three simulation methods, images and visual comfort metrics are produced. The calculated metrics are: vertical illuminance (E_v), unified glare rating (UGR), and daylight glare probability (DGP). E_v is useful to measure the accuracy of the energy preservation between simulation methods. UGR, which measures contrast, is useful for characterizing the accuracy of the distribution of intensities. DGP combines both E_v and contrast.

4 RESULTS AND DISCUSSION

Figure 5 shows image outputs for our method (*RYT*), the reference simulation (*REF*), and the low resolution simulations (*LOW*). The false color scale used on these images spans four decades of luminance values to show relative detail across the full luminance range. In general, the *RYT* images offer a striking amount of detail when considering that there are on average 0.95% of the number of samples (split between sun and sky sampling) in the *RYT* images compared to *REF*. Subjectively, the *RYT* images appear more similar to the *REF* images than the *LOW* images. Even so, there are some notable visual errors in the *RYT* images:

Table 1: The scene configurations. All four views are simulated for each option. Runs with Sun 1 include furniture and mullions. Runs with Sun 2 use simplified geometry and no furniture.

Material	Sun	Apertures	# Results
Clear Glass	1 and 2	South and East/West	16
Fabric Shades	1	South and East/West	8
Ext. Louvers	1	South and East/West	8
Linear Frit	2	South, East/West, and Skylight	12
Dot Frit	2	South	4
Translucent	1 and 2	Skylight	8
		Total:	**56**

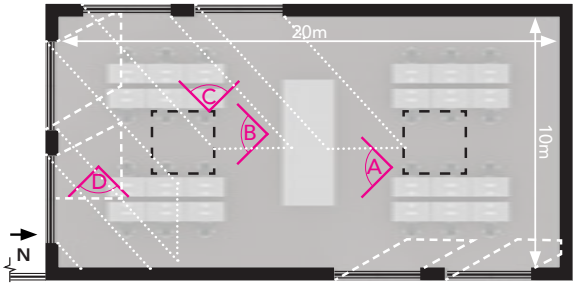

Figure 4: The test scene. Magenta view cones indicate view positions and directions. Space has a 4m ceiling height. Black dashed lines are skylights above. White dashed and dotted lines show sun penetration of Sun 1 and Sun 2, respectively.

(1) darker regions with lots of small detail, such as the workstations in images A and C, appear blurry.
(2) areas of low relative contrast but high frequency, such as the louvers in image B and dot frit in front of the adjacent wall in image D, are inconsistently rendered and full of artifacts.
(3) some small details, like pieces of the horizontal mullion in image A, are missing.
(4) medium contrast edges, like edges between sunlight and shadow in images A and B or the skylight well in image C, appear jagged.

Higher image fidelity could be achieved both by increasing the sampling rate (while still maintaining substantial compression) and/or through image reconstruction and interpolation techniques. However, the primary goal of this method is to calculate high accuracy metrics. The intended applications will produce a quantity of data for which examining individual images will be impractical. Instead, the image reconstruction is meant to be descriptive enough for useful subjective assessment and validation of the lighting condition, which these sample images appear to offer. Compared with other rapid feedback images, which often limit the number of light

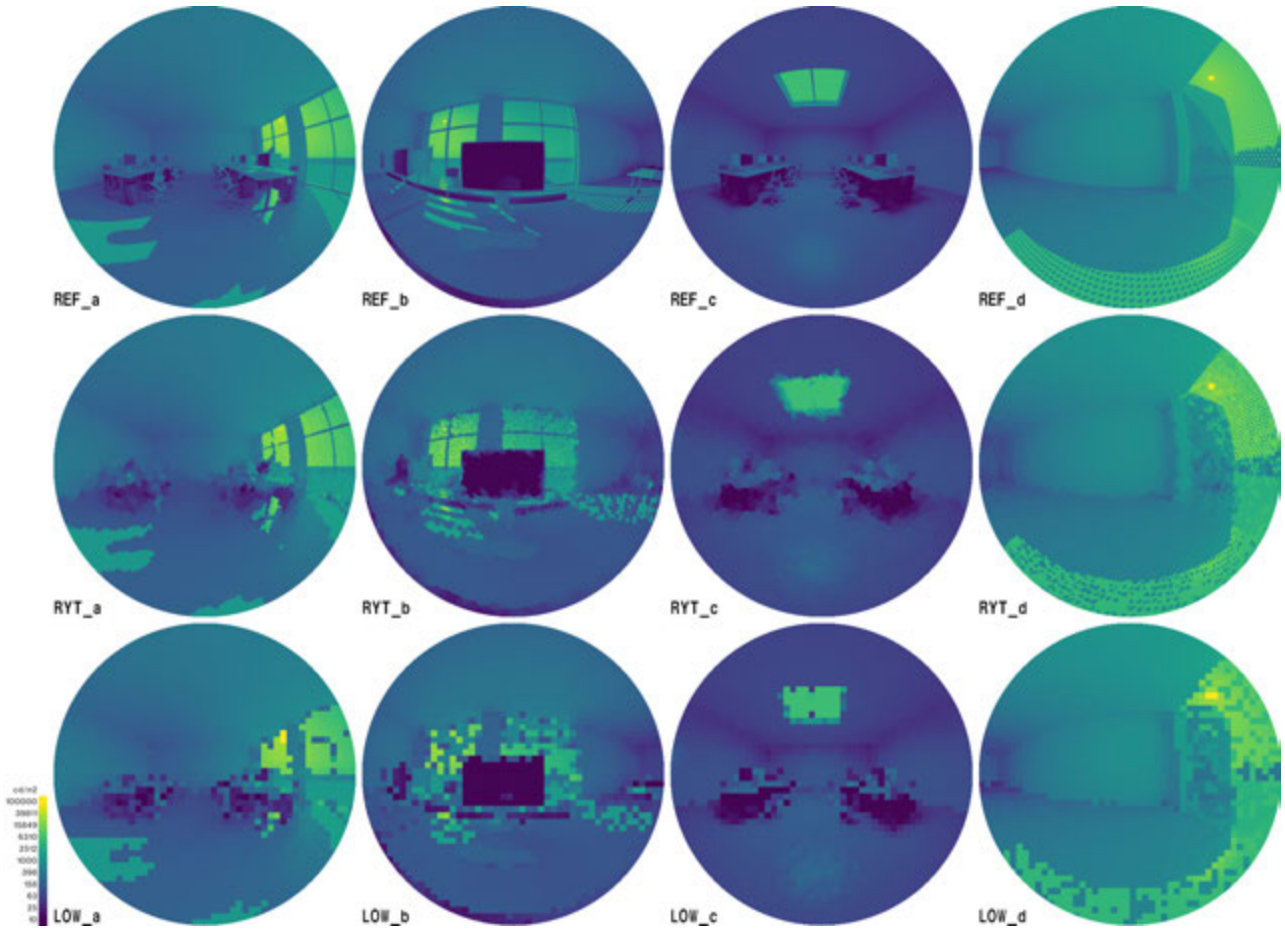

Figure 5: Four out of the 56 scenes/views analyzed for errors. Images labeled *REF* are the reference simulations, *RYT* are generated with *raytraverse*, and *LOW* are 48x48 pixel *low-res* images. View (a) has clear glazing; (b) has exterior louvers, modeled as geometry; (c) is a trans material; (d) is a dot shaped frit. *REF* simulations were rendered at 3000 x 3000 pixels and then filtered to 1000 x 1000 pixels. *RYT* images were interpolated at 1000 x 1000 pixels. *RYT* and *REF* images are further filtered to fit on the page. *LOW* images are resized to match.

bounces beyond what the full simulation will consider, these images include the full propagation of light in the space. Future work and feedback from researchers and practitioners will be needed to assess the suitability of these images for this purpose. The remainder of the discussion focuses on an analysis of the calculated metrics.

4.1 Accuracy

Table 2 lists the summary error statistics for comparing *RYT* to *REF* as well as two other simulation methods to serve as comparables. The *LOW* option uses the same uniform resolution process as the reference simulation (including all parameters) except the images are rendered at 48x48 pixels. This resolution was chosen to match the average sample density of the *RYT* samples visible in the 56 simulated views. A third method, *ILL*, simulates illuminance directly using RTRACE with the -I parameter. Here, samples are distributed according to the ambient division parameter in the same manner that subsequent surface interactions are handled. All three methods include a direct sampling component to account for the sun's contribution.

The *RYT* results have by far the smallest error relative to the comparison methods, both in absolute and relative terms. Figure 6 shows the distribution of errors for *RYT*, *LOW*, and *ILL*. For illuminance, the relative mean absolute deviation (MAD) is 2.72% or six times less than *LOW* and 8.3 times less than *ILL*. Successive *REF* runs with the same parameters for this set of scenes have a relative MAD of 1.7%, so a fair amount of the error in *RYT* could be attributable to expected random noise. *RYT* also performs well for DGP, with a MAD of less than 0.01 (in absolute units). As glare probability categories are spaced by 0.05, this could in some cases lead to a categorical shift. Among the cases between DGP 0.3 - 0.5 (N=10), there was only one case with a deviation greater than 0.02 (.31 for REF vs. 0.20 for RYT). UGR is the most sensitive metric to simulation and evaluation differences, but *RYT* still shows a

Table 2: Error metrics for the three reduced sampling techniques. rMSE: root mean squared error (exposes outliers), MAD: mean absolute deviation (keeps units of metric), MSD: mean signed deviation (shows bias). * *illuminance based methods cannot calculate a true DGP, but because the DGPs method is widely (and often mistakenly) used in annual glare evaluations [19], it is included here in comparison to the reference DGP value.*

E_v (Lux)	*RYT*	*LOW*	*ILL*
rMSE	281	2,012	2,047
MAD	110	557	533
MAD (relative)	2.72%	16.54%	22.64%
MSD	7.23	-208.21	-312.41
DGP			**DGPs***
rMSE	0.017	0.119	0.145
MAD	0.007	0.035	0.055
MAD (relative)	2.91%	9.65%	20.18%
MSD	-0.001	-0.030	-0.038
UGR			
rMSE	4.08	7.77	
MAD	1.38	3.64	
MAD (relative)	7.27%	20.76%	
MSD	-0.069	-2.815	

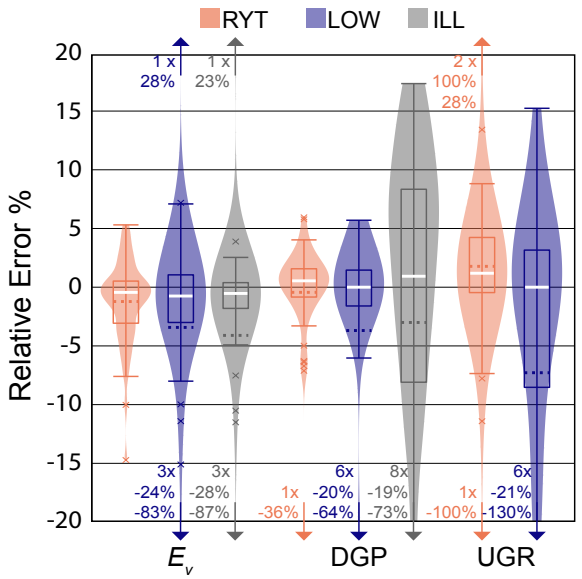

Figure 6: The relative error between reduced sampling methods and the reference simulation. White lines show median value and dashed horizontal line is the mean. Relative error is calculated as $X_{test}/X_{ref} - 1$, so positive values are over predictions relative to the reference and negative values are under predictions.

significant improvement over *LOW*, perhaps most importantly by avoiding outliers.

Outliers can be interpreted as misses or wrong answers. The *RYT* outliers all occur in low E_v and low contrast settings where the glare source threshold becomes especially sensitive. The problem with the outliers in *LOW* and *ILL* are two-fold. First, many annual metrics rely on binary condition counting statistics, whether a quantity is above or below a threshold. Small errors with uniform distributions will tend to evenly fall above and below the threshold, averaging away the error. But large errors will frequently flip the threshold, which, given the wrong distribution of annual conditions, will lead to a substantial over or under prediction of a metric. Second, these misses tend to be systematic. *ILL* never accounts for specular refections or direct sampling through certain material models, and *LOW* does not preserve the contrast of directly visible sun and misses specular reflections.

4.2 Storage Efficiency

The sampling density of a 1000x1000 pixel image representing a 180 degree angular fisheye is nearly constant and equals 125,000 rays/steradian. For *RYT*, the density ranged from 320-5,200 with a mean/median/mode of 1,190/790/350 rays/steradian across the 56 test cases. In relative terms this is a mean/median/mode of 1.0/0.6/0.3% sample density relative to the reference simulation. Because the data is unstructured, additional positional data and the tree index values need to be stored as well. Table 3 describes the file space usage for each of the options (excluding *ILL*). While *RYT* is not efficient for single time-steps because of the additional vector information and sky components, on an assumed annual output, *RYT* requires 25% less storage than the *LOW* option and 99.7% less than *REF*. Returning to the enumerations in the introduction, we estimate that the *raytraverse* lightfield can describe high resolution annual results (full directional description on a 2'x2' grid) in, on average, 965 megabytes of storage per 100ft^2 or 0.33% of the reference resolution. For simplified scenes more typically used in CBDM, this storage requirement will be even smaller.

Table 3: Storage requirements (in megabytes) for the three directional sampling methods. *RYT* in the first column already includes most of the data required for annual integration. The annual *RYT* value is estimated based on an estimate of the number of suns positions (100) required for accurate results, which requires future validation. Values for *LOW* and *REF* assume run line encoded HDR images, the standard image format used by RADIANCE.

	total (N=56)	per point	annual per 100ft^2
sky	596	11	266
sun	16	0.3	699
RYT	612	11	**965**
LOW	1	0.01	**1,280**
REF	177	3.2	**288,700**

4.3 Computation Efficiency

Because of the large variation in render times due to both scene properties and hardware, simulations for all methods were completed on the same machine, with the same version of RADIANCE.

The machine is a 2.9 GHz Intel Core i9 macbook pro (6 cores) with 16 gigabytes of RAM. Table 4 compares the computation times of the four methods, including both simulation and calculation of the metrics.

Table 4: Computation times for the four simulation methods given in per thread time. Per surface area times are estimations, described in the main text.

	average (h:m/thr)	annual 100ft^2 (/thr)
sky sampling	0:19	8h, 1m
sun sampling	0:22	0y, 1d, 17h
integrating	0:00	0h, 22m
RYT	**0:41**	**0y, 2d, 2h**
rendering	2:08	1y, 0d, 8h
evaluation	0:00	0h, 54m
LOW	**2:08**	**1y, 0d, 8h**
ILL	**0:09**	**0y, 210d, 21h**
rendering	3:34	21y, 119d, 13h
evaluation	0:00	18h, 33m
REF	**3:34**	**21y, 120d, 8h**

Average times are calculated directly as the total time divided by the 56 views divided by 12 (the number of threads used to render all methods). These times are not directly comparable between RYT and REF/LOW because the RYT simulations represent a full 360 degree view (12.56 steradians) from each point, whereas the angular fisheye projection used for REF/LOW cover only 7.38 steradians (the hemisphere plus the corners). Additionally, the sky sampling of RYT potentially represents a whole year of simulation data. To make a more level comparison, an estimate of the annual simulation time was also made. For REF and LOW, which do not include sky coefficients, this involved first completing the population of the ambient cache by rendering the off-views from each point. This time is used as an estimate for each of the other 24 points in a 100 sq. ft. simulation area. The total of the render times for the primary view and off-views is multiplied by 3,650 to represent the annual hourly rendering time. The total rendering time is added to the average evaluation time multiplied by 25 and 3,650. To estimate the annual time for RYT, the sky sampling time is multiplied by 25 (to get the 100 sq. ft.), but no additional simulation is needed. The sun sampling time is multiplied by 100 (the number of suns assumed to match the resolution of an hourly timestep) and again by 25. The bulk of the evaluation time only needs to be done once per sun position, and then the sky vectors can all be applied together for each sun position since the metric evaluations are computed directly on the raw lightfield data and no image is synthesized.

While *RYT* and *LOW* sampled a similar number of rays, *RYT* took less than half the time of *LOW*. One possible explanations for this speed up is that the *RYT* samples are concentrated in brighter more direct paths to light sources, requiring fewer bounces to converge, but this requires further investigation. In any case, the additional time required to generate the sample positions is a negligble fraction of the total time. The test cases in this analysis, with large numbers of specular surfaces, complex geometry, and procedurally defined materials, are more computationally intensive than typical daylight models. Even so, using *raytraverse* on a modern laptop, highly detailed annual results for similar scenes should be attainable in around 10 minutes per point.

5 CONCLUSION

This study has shown that the *raytraverse* method can accurately duplicate the results of our reference simulation across difficult conditions, far exceeding the accuracy of a uniform resolution method. Most importantly, it avoids complete misses, which manifest as out of bounds outliers in Figure 6. *Raytaverse* achieves insignificant differences from high resolution reference simulations with 1% of the total sampling rate. Additionally, *raytraverse* scales to evaluating multiple time-steps and view directions without additional sampling. This study was limited to an accuracy study of our method only, with a particular emphasis on difficult to simulate cases rather than a typical scene and the typical full annual simulation. As such, we did not include performance comparisons to other existing methods, which is a necessary piece of work currently in process.

Compared to the reference simulation method, which can be considered best practice for single sky condition rendering, *raytraverse* reduces simulation time by a factor of 350. *Raytraverse* should also compare favorably to illuminance based calculation when considering all of the additional detail that is recoverable from the result. The extension of the method to full annual and complete room scope, given that the method leverages the already established daylight coefficient texchnique, will yield further gains in efficiency. An annual simulation validation will reveal an appropriate relaxation for the *raytraverse* accuracy parameter, RADIANCE parameters, level of sky discretization, and suitable direct sun spacing. All of which should contribute to faster analysis times while maintaining accurate and unbiased results for any daylighting scenario. The proposed method will enable the practical extension of glare and other luminance-based evaluations to building zones. The *Raytraverse* software developed for this research can be used either as a standalone command line tool, similar to RADIANCE, and/or by interfacing directly with the Python library. It is the hope of the authors that simulation tool developers working in Python will consider *Raytraverse* to help expand the capabilities and performance of their tools.

6 SUPPLEMENTAL MATERIAL

The *raytraverse* source code used to run the simulations in this paper is archived at: https://doi.org/10.5281/zenodo.4278918. Source code documentation for the version used for this paper, and link to the current repository, can be found at: https://raytraverse.readthedocs.io/en/v1.0.4.

ACKNOWLEDGMENTS

This research was supported by the Swiss National Science Foundation as part of the ongoing research project “Light fields in climate-based daylight modeling for spatio-temporal glare assessment” (SNSF #179067). We thank Stephen Wittkopf the project’s principal investigator, for accompanying the work with his constructive and critical feed-back.

REFERENCES

[1] Ian Ashdown. 1998. The Virtual Photometer: Modeling the Flow of Light. In *IESNA Annual Conference*. Illuminating Engineering Society of North America.

[2] E. J. Candes and M. B. Wakin. 2008. An Introduction To Compressive Sampling. *IEEE Signal Processing Magazine* 25, 2 (2008), 21–30. https://doi.org/10.1109/MSP.2007.914731 https://doi.org/10.1109/MSP.2007.914731.

[3] K. Chamilothori, G. Chinazzo, J. Rodrigues, E.S. Dan-Glauser, J. Wienold, and M. Andersen. 2019. Subjective and physiological responses to façade and sunlight pattern geometry in virtual reality. *Building and Environment* 150 (2019), 144–155. https://doi.org/10.1016/j.buildenv.2019.01.009

[4] Lars Oliver Grobe. 2019. Photon-mapping in Climate-Based Daylight Modelling with High-resolution BSDFs. *Energy and Buildings* 205 (2019), 109524. https://doi.org/10.1016/j.enbuild.2019.109524

[5] IESNA. 2012. IES LM- 83-12, Approved Method: IES Spatial Daylight Autonomy (sDA) and Annual Sunlight Exposure (ASE).

[6] Nathaniel L. Jones. 2020. Fast Climate-Based Glare Analysis and Spatial Mapping. In *Proceedings of Building Simulation 2019: 16th Conference of IBPSA*, V. Corrado, E. Fabrizio, A. Fasparella, and F. Patuzzi (Eds.). International Building Performance Simulation Association, Rome, Italy. https://doi.org/10.26868/25222708.2019.210267

[7] Nathaniel L Jones and Christoph F Reinhart. 2017. Experimental validation of ray tracing as a means of image-based visual discomfort prediction. *Building and Environment* 113 (2017), 131–150. https://doi.org/10.1016/j.buildenv.2016.08.023 bibtex*[publisher=Elsevier].

[8] Eleanor S. Lee, David Geisler-Moroder, and Gregory Ward. 2018. Modeling the direct sun component in buildings using matrix algebraic approaches: Methods and validation. *Solar Energy* 160 (2018), 380–395. http://www.sciencedirect.com/science/article/pii/S0038092X17311118 https://doi.org/10.1016/j.solener.2017.12.029.

[9] J. Mardaljevic. 2000. Simulation of annual daylighting profiles for internal illuminance. *International Journal of Lighting Research and Technology* 32, 3 (2000). https://journals.sagepub.com/doi/abs/10.1177/096032710003200302 https://doi.org/10.1177/096032710003200302.

[10] A. McNeil and E. S. Lee. 2013. A validation of the Radiance three-phase simulation method for modelling annual daylight performance of optically complex fenestration systems. *Journal of Building Performance Simulation* 6, 1 (2013), 24–37. https://doi.org/10.1080/19401493.2012.671852 https://doi.org/10.1080/19401493.2012.671852.

[11] Ryan S. Overbeck, Craig Donner, and Ravi Ramamoorthi. 2009. Adaptive Wavelet Rendering. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 1–12. https://doi.org/10.1145/1618452.1618486 https://doi.org/10.1145/1618452.1618486.

[12] Christoph F. Reinhart and Sebastian Herkel. 2000. The simulation of annual daylight illuminance distributions — a state-of-the-art comparison of six RADIANCE-based methods. *Energy and Buildings* 32, 2 (July 2000), 167–187. http://www.sciencedirect.com/science/article/pii/S0378778800000426 https://doi.org/10.1016/S0378-7788(00)00042-6.

[13] Luis Santos and Luisa Caldas. 2018. Assessing the Glare Potential of Complex Fenestration Systems: A Heuristic Approach Based on Spatial and Time Sampling. In *PLEA 2018 - Smart and Healthy within the Two-Degree Limit: Proceedings of the 34th International Conference on Passive and Low Energy Architecture*, Vol. 1.

[14] Peter Shirley and Kenneth Chiu. 1997. A Low Distortion Map Between Disk and Square. *Journal of Graphics Tools* 2, 3 (Jan. 1997), 45–52. https://doi.org/10.1080/10867651.1997.10487479

[15] P. R. Tregenza and I. M. Waters. 1983. Daylight coefficients. *Lighting Research and Technology* 15, 2 (June 1983), 65–71. http://dx.doi.org/10.1177/096032718301500201 https://doi.org/10.1177/096032718301500201.

[16] Pauli Virtanen et al. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. https://doi.org/10.1038/s41592-019-0686-2.

[17] Beibei Wang, Xiangxu Meng, and Tamy Boubekeur. 2015. Wavelet Point-Based Global Illumination. *Computer Graphics Forum* 34, 4 (2015), 143–153. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12686 https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12686 https://doi.org/10.1111/cgf.12686.

[18] Gregory J. Ward. 1994. The RADIANCE lighting simulation and rendering system. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques (SIGGRAPH '94)*. Association for Computing Machinery, New York, NY, USA, 459–472. https://doi.org/10.1145/192161.192286 https://doi.org/10.1145/192161.192286.

[19] Stephen Wasilewski, Lars Oliver Grobe, Jan Wienold, and Marilyne Andersen. 2019. A Critical Literature Review of Spatio-temporal Simulation Methods for Daylight Glare Assessment. *SDAR* Journal of Sustainable Design & Applied Research* 7, 1 (Jan. 2019). https://arrow.tudublin.ie/sdar/vol7/iss1/4 https://doi.org/10.21427/87r7-kn41.

[20] Jan Wienold. 2009. Dynamic daylight glare evaluation. In *11th International IBPSA Conference, Glasgow, Scotland*. 944–951. http://www.ibpsa.org/proceedings/BS2009/BS09_0944_951.pdf.