

A Critical Comparison of Annual Glare Simulation Methods

Stephen Wasilewski^{1,2*}, Jan Wienold², Marilyne Andersen²

¹ Lucerne University of Applied Sciences and Arts, Horw, Switzerland

² École Polytechnique Fédérale de Lausanne, Switzerland

* corresponding author: stephen.wasilewski@hslu.ch

Abstract

Recently, there have been multiple proposals for faster methods to calculate glare metrics, daylight glare probability (DGP) in particular. This is driven simultaneously by the lengthy times required to simulate DGP with a conventional image-based approach and accumulating evidence from subjective glare evaluation experiments showing that accounting for both the saturation and contrast in a view improves the accuracy of a glare prediction. While some of these methods have been presented with their own validations, comparisons of accuracy across methods are limited by the differences in tested scenarios and resulting distributions of daylight conditions. This study compares six point-based methods and two zonal estimations for quickly calculating hourly DGP values from three viewpoints with different relationships to the window across a range of scenarios. These include scenarios with direct and semi-specular transmission and others with specular and semi-specular reflection. We find that while some of these fast methods closely align with results from a reference simulation, others introduce large and consequential errors.

Introduction

Several methods exist for the efficient calculation of annual glare conditions evaluated by the daylight glare probability metric:

$$\text{DGP} = 5.87 \cdot 10^{-5} \cdot E_v +$$

$$9.18 \cdot 10^{-2} \log_{10} \left(1 + \sum_{i=1}^n \frac{L_{s,i}^2 \omega_{s,i}}{E_v^{1.87} P_{s,i}^2} \right) + 0.16.$$

Where E_v is vertical illuminance (in *lux*), n is the number of glare sources, $L_{s,i}$ is the luminance of glare source i (in cd/m^2), $\omega_{s,i}$ is the solid angle of glare source i (in *steradians*), and $P_{s,i}$ is the position index. The first term, which is dependent only on E_v , accounts for glare due to saturation (the brighter it is, the higher the chance of discomfort). The second term accounts for glare due to contrast (the brighter potential glare sources are relative to the observer's adaptation, the higher chance of discomfort). Each of these methods offers trade-offs between accuracy, simulation time, and sensitivity to different scenarios which may cause glare. Such trade-offs are necessary because of the large domain of the evaluation space and the lengthy simulation times required to yield accurate

photometric results with a conventional image-based approach.

The point-based methods evaluated are: simplified DGP (DGPs) (Wienold, 2009), enhanced simplified DGP (eDGPs) (Wienold, 2009), imageless DGP (Jones, 2019), *ClimateStudio* annual glare (www.solemma.com), *ray-traverse* (Wasilewski et al. 2021), and adaptive glare coefficients (AGC) (Wienold, 2022). The two zonal estimation methods are: glare annual classes evaluation (GLANCE) (Giovannini et al., 2020) and a heuristic based approach proposed by Santos and Caldas (2021). All the tested implementations of these methods use *Radiance* (Ward, 1994) for computing ray values, including luminance, illuminance, and daylight-coefficients. The reference simulation is also completed using *Radiance*, so there is a perfect match between modelled geometry, material properties, viewpoints and directions, and sky model across all methods. As such the errors presented here are those introduced by the tested methods on top of whatever unknown error would exist between the reference simulation and the real-world equivalent of our test scenes. Before describing the reference simulations and test procedures in the next section, each of the tested methods is briefly introduced below.

DGPs

Because DGPs is entirely based on a linear function of E_v , its accuracy and simulation time is dependent on the method used to simulate E_v . DGPs ignores individual glare sources and can predict glare only if saturation glare is present. DGPs fails in low-light conditions with high contrast, as in the case of electrochromic windows when the sun is visible from the observer viewpoint (Wienold et al., 2019). Results are presented for this method using high-accuracy climate-based daylight modelling (CBDM) method detailed in the *Methods* section.

eDGPs

The eDGPs method's core assumption is that glare sources can be detected by using simplified images that do not consider diffuse interreflections (i.e., using $-ab$ 0 in *Radiance*). To determine the DGP values for the required timesteps, a simplified image is rendered and glare sources detected for every timestep when the direct irradiance is larger than 50W/m^2 . The original implementation of eDGPs uses E_v values calculated with *daysim*. In our study, results are presented with CBDM sensor point

data matching the other methods using pre-computed values.

Imageless DGP

Imageless DGP is calculated using the *dcglare* executable distributed with *Radiance*. As input, it takes a sky matrix with sky vectors for each timestep and two daylight coefficient matrices: DC_{total} , which includes the total contribution (direct and indirect) of a sky patch at a cosine weighted sensor point and view direction, and DC_{direct} , which includes only the direct part of the same contribution. DGP is calculated by deriving the glare sources and E_v directly from these matrices. Therefore, unlike an illuminance-only approach, it can have some knowledge of the angular distribution of light. Imageless DGP was previously validated against 2-phase daylight coefficient images and eDGPs simulations (Jones, 2019). The scene was a corner open office with two fenestrations: clear glazing and geometrically modelled venetian blinds. Reference data was calculated annually at a single East facing window adjacent point and spatially (6552 views) for two time-steps. Because the 2-phase reference images do not model the direct sun with an accurate source size, we consider the eDGPs reference to be more accurate. Compared to eDGPs, the validation found a root mean square error (RMSE) of 0.06 to 0.10 and a mean signed deviation (MSD) of +0.04 to +0.07 across spatial and temporal reference sets for both window scenarios.

ClimateStudio Annual Glare

ClimateStudio is a plug-in for the 3D-modelling software Rhinoceros. It implements a progressive path tracer using *Radiance* and GPU based hardware acceleration. According to the documentation (<https://climatestudio-docs.com/docs/annualGlare.html> (accessed 3.17.22)), “[f]or annual DGP simulations, *ClimateStudio* relies on the vertical illuminance portion of the DGP formula, plus a contrast measurement from the solar disc.” Because of the way it sends samples from a point, individual traced rays can update multiple view directions simultaneously, further reducing simulation time when conducting a multi-view direction analysis. While *ClimateStudio* is widely used academically and in industry, we do not know of a peer reviewed validation of its annual glare calculation against high quality reference data (whether simulated or measured).

Raytraverse

Raytraverse applies a novel sampling approach to leverage the results produced by a *Radiance*-based renderer. Using information garnered from the model and iterative rounds of simulation, *raytraverse* requires a small fraction of the samples needed by full resolution images to describe light at a point with the same effective resolution. These results are directly evaluated, yielding additional time savings while maintaining high accuracy. *Raytraverse* has been previously validated for two different implementations (Wasilewski et. al., 2021; 2022).

First, Wasilewski et. al. (2021) compared the directional sampling algorithm applied to sky coefficients and full contributions from the direct sun to high quality reference simulations across four views and 14 different scenes varying in window location, transmitting material, and geometric detail. That study found a RMSE for DGP of 0.017 and MSD of 0.00.

Second, Wasilewski et. al. (2022) implemented a 2-phase method with direct sun (DDS, for dynamic daylight simulation (Walkenhorst, 2002)). This method was simulated with a range of quality parameters and compared to high quality point-in-time reference simulations for eight views across a mix of interior and perimeter open areas and private offices. No RMSE was reported, but the distribution normalized (accounting for the frequency of different conditions in the reference data) mean absolute error (MAE), which will always be less than or equal to RMSE, for the lowest quality setting was 0.01, and the MSD was 0.00.

The implementation tested here is slightly different from the second validation. In terms of representing the direct sun by true source and sky patch, this is equivalent to a sensor point 2-phase DDS approach, with one less bounce/intersection represented by the true source than a 2-phase DDS image.

AGC

This new approach (Wienold, 2022) optimizes the image generation part of the eDGPs method and replaces the “brute force” (=at every timestep) simulation of a simplified image by a method that calculates only the relevant parts of an image for sun positions significantly different from those which have already been calculated (“adaptive”). Potential glare sources originating from the sky luminance distribution are determined in a separate simulation step and added to the glare equation. This method, like eDGPs, uses pre-calculated E_v values, so results are presented with CBDM sensor point data.

Zonal Estimations

GLANCE is an illuminance-based method like DGPs (Giovannini et al., 2020). The novelty of the method is that the thresholds for the different glare categories are calibrated to the scene by means of an eDGPs analysis at a single point with the same view direction across the space. The heuristic zonal approach (Santos and Caldas, 2021) posits that E_v can be used for an initial search of glare conditions zonally and then, based on a fixed threshold, highlight those conditions for further evaluation. To assess the reliability of these two methods, we can evaluate how well both the fixed and calibrated thresholds predict glare according to DGP when applied across the range of scenarios included in our analysis. For GLANCE, this is limited to the two points with matching view directions.

Methods

With a focus on the sensitivity of the methods to detect possible glare scenarios, this study looks at several fixed

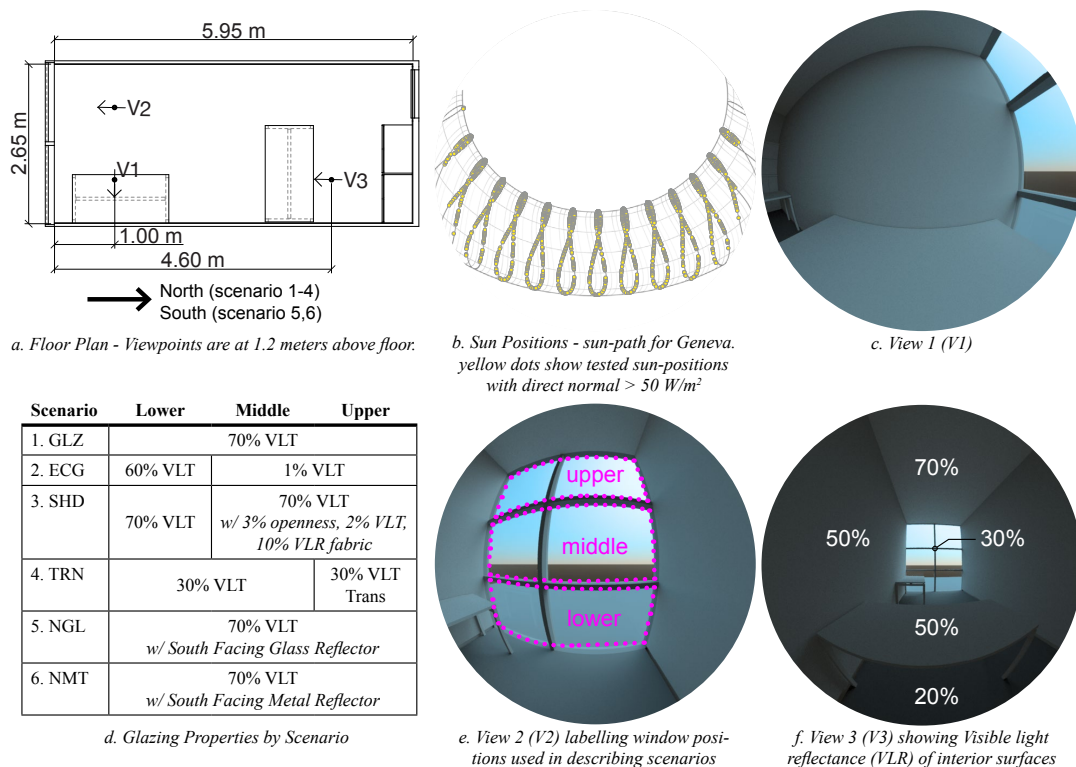


Figure 1: Test Scene and Evaluated Views (rendered with synthetic clear sky and no sun)

state conditions for a single occupancy office with a range of transmitting materials and reflection scenarios that can cause glare in buildings. We evaluate the office across the occupied area from three points, two facing the window wall and one facing perpendicular to the window. To verify the accuracy, we use a reference high resolution image simulation using high quality point-in-time *Radiance* simulations.

Test Scenarios

The tested scene is a simplified cellular office model published in the public domain (Grobe et al., 2020). The model was simplified by removing the desk chairs and monitors, and by simplifying the materials. See Figure 1 for plan, base material properties, solar positions, and viewpoints.

Sky conditions are derived from the typical meteorological year IWEK dataset for Geneva, Switzerland (<https://www.ladybug.tools/epwmap/> (accessed 3.17.22)). Scenarios use direct normal irradiance and diffuse horizontal irradiance as parameters to the Perez all-weather sky model, as implemented by the *gendaylit* executable of *Radiance*. The file *perezlum.cal* was modified to exclude the sky/ground blending to better match the output of *gendaymtx*, which is used by some of the methods. All timesteps with direct normal irradiance less than 50 watts/m² are excluded. Sun positions are calculated using *gendaylit* assuming the half hour for each hourly entry in the energyplus weather file. Note that the radiance version used is *Radiance* 5.4a (2021-11-21), and the solar position algorithm was last updated Oct. 2019.

Using this model and set of sky conditions, we simulated six different scenarios, shown in Figure 2, with glazing material properties in Figure 1e. Four scenarios face South. The other two scenarios face North with a reflective surface 20 meters high and 12 meters wide located 6 meters North of the window.

Reference Simulation

The 36,396 conditions (six scenarios, three views, and 2,022 skies) were simulated with point-in-time simulations using the *Radiance rpict* executable, as managed by the *rad* program. The final command (cleaned to remove redundant settings) with oversampling and filtering is:

```
rpict -x 3000 -y 3000 -vta -vh 180 -vv 180\
-dp 4096 -dt .01 -dc 1 -ds .2 -dr 3\
-ms 0.025 -ss 16 -st .01 -lr 12 -lw 1e-5\
-af afile -av 0 0 0 -aa .075 -ar 600 -ab 6\
-ad 1500 -as 750 -ps 3 -pt .04 octree > img.unf
pfilt -1 -e 1 -m .25 -x /3 -y /3 img.unf > img.hdr
```

Reference DGP values were calculated from these reference images using *evalglare* v3.02 with the default settings (glare source threshold 2000 cd/m² and search radius 0.2 radians).

CBDM Illuminance

Several of the test methods require an externally calculated illuminance value. While we can use the illuminance calculated from the reference images to understand something about the error intrinsic to or added by the method, this would not represent a practical workflow. Instead, we calculate illuminance using the 2-phase DDS CBDM method described as described by Subramaniam (2017). We focused on achieving a high accuracy result and did

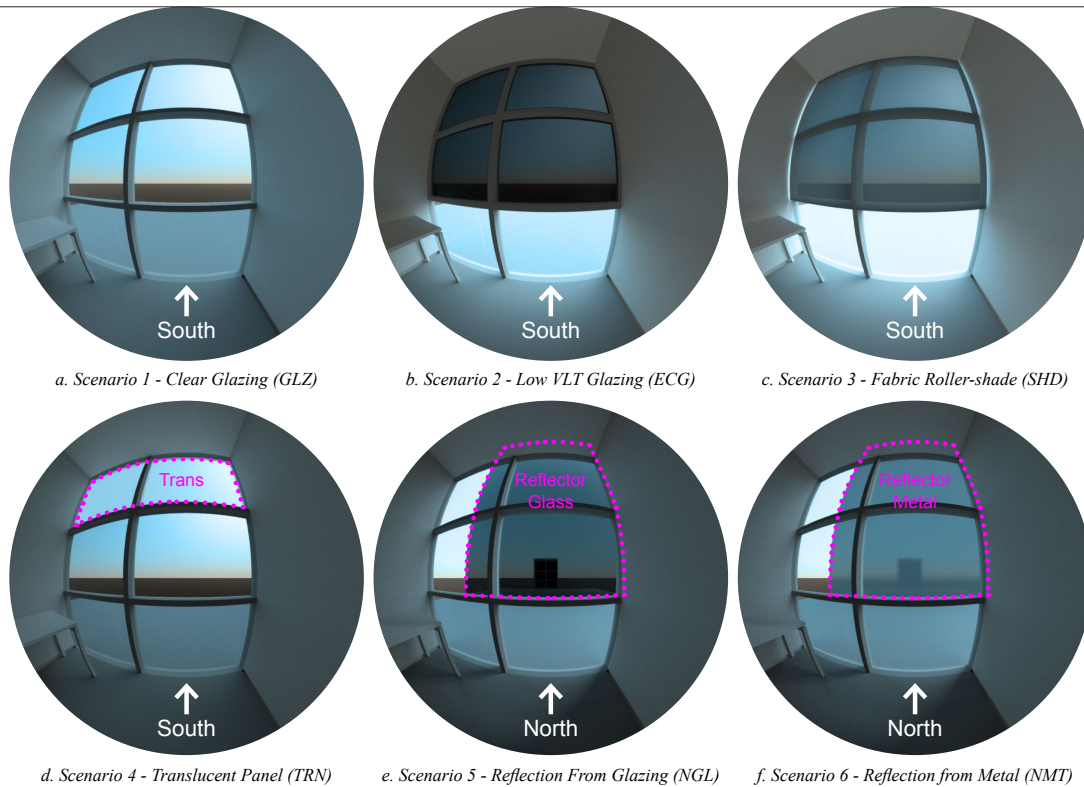


Figure 2: Transmission and Reflection Scenarios (rendered with synthetic clear sky and no sun)

not attempt to optimize the simulation parameters for speed. Daylight Coefficients were calculated for Reinhart sky divisions with MF:4 (2,305 sky patches + ground) and sun coefficients with MF:6 (5,185 positions). *Rfluxmtx* settings for the daylight coefficient and direct daylight coefficient (with *-ab 1* and non-transmitting materials black) matrices are: *-lw 1e-8 -ab 7 -lr -12 -ad 200000*. *Rcontrib* settings for the direct sun matrix (with non-transmitting materials black): *-ab 1 -n 16 -ad 10000 -lw 1e-5 -dc 1 -dt 0 -dj 0 -e MF:6*.

Parameters for Tested Methods

Beyond those for generating illuminance results, DGPs requires no additional parameters, as the methods only entails a linear fit of E_v to approximate DGP. For calculating eDGPs, the image resolution is set to 1200x1200. To calculate imageless DGP, the default parameters of *dclglare* are used. For *ClimateStudio*, the default simulation parameters used are *-ad 1 -ab 6 -lw 0.01* with 100 passes and 64 samples per pass. For *raytraverse*, the simulation parameters used are *-ad 7250 -as 0 -lw 5.5e-5 -ab 6*, with 144 sky patches and a final sampling resolution of 256x256 per hemisphere. For AGC, the parameters are: Maximum distance between sun positions: 2.5° (0.27° for the north facing variants), Image resolution: 1200x1200, no. of luminance bins for glare source detection: 8.

Results and Discussion

Figure 3 plots the distribution of glare conditions as categorized by the reference simulation. Of these 36,396 conditions, 17,277 fall between DGP 0.25 and 0.75. We choose to limit the calculation of error metrics to this range where errors due to simulation method could intro-

duce meaningful differences in the interpretation of results. The thresholds for glare categories for DGP are be-




Figure 3: Distribution of reference conditions by scene. See Figure 2 for scene descriptions. Dashed box indicates range used for calculating error statistics

tween 0.35-0.45, although recent research suggests that under certain lighting conditions the thresholds between comfort and discomfort could be higher (McNeil and Burrell, 2016). Among these filtered conditions, only one case for DGPs and seven cases for *ClimateStudio* fail to predict any glare ($DGP > 0.35$) when the reference DGP is greater than 0.75.

Time-Step Error

Calculating DGP including a contrast term is somewhat more complicated. Table 1 shows RMSE, MAE, and MSD for the six tested methods as well as a reference error. The reference error is a result of changing the way the source data is grouped, the threshold for including pixels as part of the direct solar source and normalizing the solar source area. Given our reference data, this error, on a per

scene basis, is likely the lowest achievable by any method. Table 2 highlights those methods that are within 100% of this reference error. Only eDGPs, *raytraverse*, and AGC meet this benchmark for any of the scenes for RMSE and MAE. For MSD, unbiased results are also observed for DGPs and imageless DGP for some scenes, although the average bias across all scenes is still significantly larger than the reference magnitude.

Table 1: Error in DGP Calculation. White/bolded cells are within 100% of reference error (by column). Colors range from blue to white to red from -0.08 to 0.08.

RMSE	GLZ	ECG	SHD	TRN	NGL	NMT	AVG
reference	0.01	0.02	0.02	0.01	0.03	0.01	0.02
DGPs	0.03	0.07	0.07	0.05	0.11	0.03	0.06
eDGPs	0.01	0.03	0.03	0.02	0.04	0.02	0.03
imageless DGP	0.04	0.07	0.07	0.05	0.09	0.03	0.06
climtestudio	0.04	0.06	0.06	0.06	0.14	0.07	0.07
raytraverse	0.02	0.02	0.02	0.01	0.04	0.03	0.02
AGC	0.02	0.03	0.03	0.03	0.04	0.02	0.03

MAE	GLZ	ECG	SHD	TRN	NGL	NMT	AVG
reference	0.00	0.01	0.01	0.00	0.02	0.01	0.01
DGPs	0.03	0.05	0.05	0.04	0.08	0.02	0.04
eDGPs	0.01	0.02	0.02	0.01	0.03	0.01	0.02
imageless DGP	0.01	0.07	0.06	0.02	0.06	0.02	0.04
climtestudio	0.03	0.06	0.05	0.04	0.09	0.05	0.05
raytraverse	0.01	0.01	0.01	0.01	0.03	0.02	0.01
AGC	0.01	0.03	0.03	0.02	0.03	0.01	0.02

MSD	GLZ	ECG	SHD	TRN	NGL	NMT	AVG
reference	0.00	0.01	0.01	0.00	0.01	0.00	0.01
DGPs	0.01	-0.05	-0.04	0.02	-0.05	0.01	-0.02
eDGPs	0.00	0.00	0.00	0.01	0.02	-0.01	0.00
imageless DGP	0.01	-0.07	-0.06	0.00	-0.05	-0.01	-0.03

<i>climtestudio</i>	-0.01	-0.05	-0.04	-0.03	-0.08	-0.04	-0.04
<i>raytraverse</i>	0.00	0.01	0.01	0.00	0.02	0.00	0.01
<i>AGC</i>	0.00	-0.01	0.00	0.01	0.02	-0.01	0.00

Figure 4 shows scatter plots for each of the six methods compared to the reference simulation. Here the origin of the error and bias in Table 2 is readily visible by scene and view. DGPs, *ClimateStudio*, and to a lesser extent imageless DGP, all appear to systematically underpredict glare for certain scenarios. *ClimateStudio* underpredicts glare for TRN, NMT, and NGL at view 2. DGPs underpredicts glare for NGL at view 2 and ECG at views 2 and 3. Imageless DGP underpredicts glare for ECG, SHD, and NGL at views 2 and 3. *Raytraverse* consistently matches the reference, except scenario NMT, where predictions at view 2 have a wide spread quantified by the higher RMSE values. AGC closely matches the results of eDGPs, both in terms of the magnitude and pattern of deviation from the reference.

Quality of Illuminance Data

An important lighting quantity for accurately measuring DGP is E_v , both as the saturation term and the adaptation component of the contrast term. Table 1 approximates the error introduced by the deviation in E_v , by the RMSE, MAE, and MSD of DGPs calculated with the reference illuminance compared to the illuminance used by each test method for cases with reference DGP between 0.25 and 0.75. None of the methods demonstrate bias greater in magnitude than 0.01, except for *ClimateStudio* in the clear glazing scenario. *ClimateStudio* also has significant

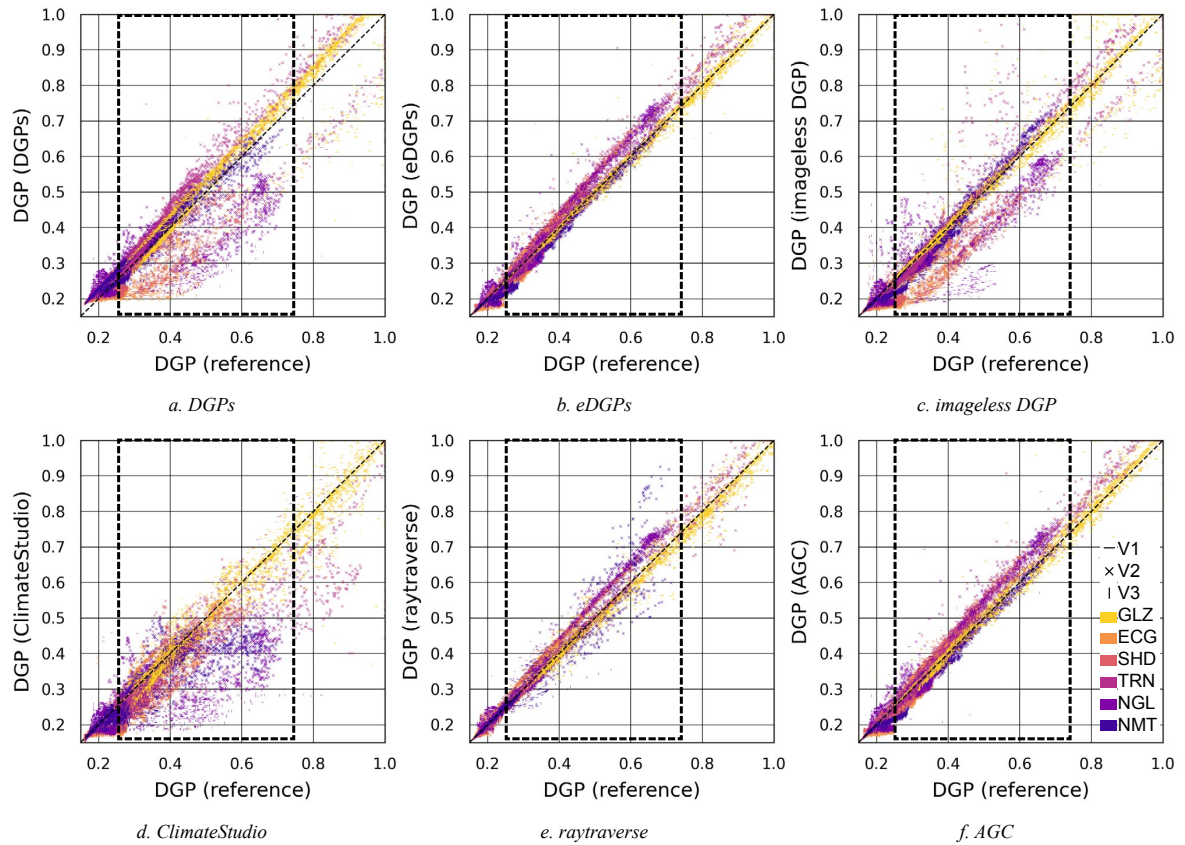


Figure 4: Scatter plots comparing reference DGP (x-axis) to test methods (y-axis). Dashed box indicates range used for calculating error statistics in Table 2.

RMSE across all scenarios, but smaller MAE, indicating a large spread in deviations. Both 2-phase methods have MAE less than or equal to 0.01 across all methods, but the 2-phase DDS performs better for clear glazing and the translucent panel. *Raytraverse* has low MAE for all methods but with a RMSE of 0.03 for the north metal reflection.

Table 2: Error in Illuminance Calculation scaled as DGPs. AGC, DGPs, eDGPs use 2-phase DDS. Imageless DGP uses 2-phase matrices.

RMSE	GLZ	ECG	SHD	TRN	NGL	NMT	AVG
2-phase DDS	0.01	0.01	0.01	0.02	0.02	0.01	0.01
2-phase	0.03	0.01	0.01	0.03	0.02	0.01	0.02
ClimateStudio	0.04	0.02	0.02	0.04	0.04	0.05	0.03
raytraverse	0.02	0.00	0.00	0.01	0.01	0.03	0.01
MAE							
2-phase DDS	0.01	0.00	0.01	0.01	0.01	0.01	0.01
2-phase	0.01	0.00	0.01	0.01	0.01	0.01	0.01
ClimateStudio	0.03	0.01	0.01	0.02	0.02	0.03	0.02
raytraverse	0.01	0.00	0.00	0.01	0.01	0.01	0.01
MSD							
2-phase DDS	0.00	0.00	0.01	0.01	0.01	0.00	0.01
2-phase	0.01	0.00	0.00	0.01	0.01	0.00	0.00
ClimateStudio	0.02	0.00	0.00	-0.01	0.01	-0.01	0.00
raytraverse	0.00	0.00	0.00	0.00	0.01	0.00	0.00

Annual Distribution Error

Often, the goal of generating hourly simulations is to quantify the annual performance, instead of precisely understanding the lighting condition at a given hour. Typically, this is done by measuring the number of hours above (or below) a set threshold or measuring a given percentile condition from the distribution. To understand how the observed timestep errors for the tested methods propagate to such annual results, the 75th, 85th, and 95th percentile DGP values for each of the eighteen view/scene combinations are calculated. To account for the imprecision of the reference values described in Table 1, deviations are taken from the nearer of the two reference calculation methods, handling cases where the tested method is as accurate as the uncertainty in the reference.

Figure 5 shows the distribution of the error across these eighteen scenarios for the six methods. Percentiles are determined based on working hours (8AM-6PM), 365 days per year. For these percentiles to be accurate, we are assuming that the 1,822 calculated values that occur during working hours (recall we only simulated sky conditions with direct normal greater than 50 watts/m²) contain the top 912 (25%) conditions.

Here we see that the time step errors tend to propagate to the annual results. The negative bias observed for imageless DGP and *ClimateStudio* is stronger than the mean for the 95th percentile results, but the percentile errors are smaller than the average MSD when observing higher frequency events. DGPs has a mix of positive and negative MSD across the scenes, so it is not entirely surprising that the percentile results would also range from positive to negative, with the higher end of the distribution tending towards a negative bias and the lower end towards a positive one.




Figure 5: Distribution of error in annual percentiles for 18 view/scene combinations. Method abbreviations: imageless DGP (iDGP), *ClimateStudio* (CS), *raytraverse* (RYT)

Raytraverse produces annual results with a very close distribution to the reference data, only one scenario (north reflected glazing at view 2, where the 95th percentile DGP is 0.64) is off by more than 0.01. eDGPs and AGC have percentile errors of less than 0.05, but all deviations less than -0.02 occur for percentile values less than 0.30, well below the threshold for perceptible glare. The larger positive biases are introduced by the E_v simulation. For cases where the reference percentile value falls between 0.3 and 0.5, The average deviations in the annual percentiles are: DGPs (0.021), eDGPs (0.007), imageless DGP (0.020), *ClimateStudio* (0.025), *raytraverse* (0.002), and AGC (0.010).

Convergence Properties

To understand whether the systematic errors observed for DGPs, imageless DGP, and *ClimateStudio* are intrinsic to the method or are the result of our chosen simulation parameters, we went on to test whether, given enough simulation time, the results begin to converge towards the correct simulation. We look at these three methods not only because of their larger errors, but also because the chosen simulation parameters for these methods resulted in faster simulation times than *raytraverse* and AGC (see next section for caveats).

For DGPs and imageless DGP, we increased the *rcontrib* simulation parameters to: *-lw 1e-8 -ab 7 -lr -12 -ad 800000* and increased the Reinhart sky divisions to MF:6 (5,185 sky patches + ground). For *ClimateStudio*, we ran 500 passes with 640 samples per pass. For DGPs, the average MSD improves from -0.02 to -0.01, but RMSE does not improve. For imageless DGP, the average MSD remains at -0.03, and RMSE improves from 0.06 to 0.05. For *ClimateStudio*, neither RMSE or MSD improves.

For DGPs, imageless DGP, and *ClimateStudio*, the same number of glare prediction misses (as described in the last section) are observed with these high-quality settings.

Repeating the check for the quality of the illuminance data using a reference DGPs value, both the 2-phase results (affecting DGPs and imageless DGP) and *ClimateStudio*

results are improved with the increased simulation accuracy. For *ClimateStudio*, this is particularly true for GLZ, ECG, and SHD, where the average RMSE is reduced from 0.03 to 0.01. This indicates that there is an intrinsic source of error and bias in how these three methods account for glare sources that is not solved for by an improved estimation of E_v .

Simulation Time for Context

This study is primarily focused on understanding the accuracy of the tested methods but given that a goal of all of these methods is to complete faster analyses of annual glare metrics, some discussion of computation time is relevant. Unfortunately, because these methods are implemented for different operating systems and have different simulation scopes in mind, comparing practical analysis times is even less reliable than usual. However, given the range of observed times between methods, it is still useful context. Table 3 summarizes the total time needed to complete analysis for each test method and information about the computer used for the simulation.

Table 3: Simulation times for the test methods. Base time is with initial settings. High Settings are those specified in the Convergence Properties Section.

	time*(base)	time*(high)	Computer
DGP _s	9.8	19.0	A
eDGP _s ²	6170	-	A, D
imageless DGP	3.8	12.5	A
<i>ClimateStudio</i> ¹	2.3	16.2	B
<i>Raytraverse</i>	15.5	-	A
AGC ²	31.8	-	A, C
*times given are real time in average seconds per point and scene.			
A. 2018 MacBook Pro with Intel 2.9 GHz Core i9 processor, 16 GB RAM, and a solid-state hard drive. Using 12 processes.			
B. Windows 10 Virtual Machine with Intel Xeon Gold 6248R CPU @ 3.00GHz, 8 GB RAM, Nvidia GRID RTX8000P-2Q GPU			
C. 2022 MacBook Pro M1max. Using 3 processes.			
D. Dell PowerEdge R6515 Server AMD EPYC 7413 with Ubuntu 20.04.3 LTS, using 1 process.			
1. <i>ClimateStudio</i> requires simulating a grid of points, so six points were simulated instead of three. Reported times are pro-rated.			
2. Illuminance calculations done on computer A.			

It is notable that the times with high settings for DGPs, imageless DGP, and *ClimateStudio* are roughly the same as the times for *raytraverse* and AGC. Although, given the hardware differences it cannot be assumed that these results are replicable. Given the increased accuracy and broader applicability of AGC and *raytraverse*, there is little performance advantage to using methods that make more assumptions about the scene and likely glare conditions. An advantage of *ClimateStudio* is its ease of use for people more comfortable in a 3D modelling GUI environment than on a command line. However, there is a danger that these users may be less aware of where and when their simulations are valid, which based on this study appears to only be for high transmission clear glazing without reflections.

Extending to Zonal Analyses

Among the methods tested in detail, *ClimateStudio* and *raytraverse* have some advantages when extending to a zonal analysis where the time to evaluate N points and views is not simply N times the time to evaluate one view-point. Both methods, without additional simulation time, can evaluate additional view directions. *Raytraverse*, through its adaptive sampling approach, can also adaptively sample a zone, reducing the number of points evaluated in areas of low variance. This should offer particularly large time savings in deeper spaces with much of the evaluated floor area far from the façade.

As mentioned in the introduction, there are also approaches to zonal analysis based on extrapolating fewer high information samples (image based) to a broader low information sampling (illuminance sensor based). For the GLANCE method, thresholds are determined through a minimization, as described in Giovannini et al. (2020). Because no exact minimization algorithm is specified in that paper, we used a brute-force approach in 50 lux increments followed by a golden section minimization bounded by 100 lux on either side of the brute-force result. GLANCE is intended to be used with eDGPs calculated at a single point to then predict glare categories across a space for views with the same direction. The other heuristic based zonal approach, by Santos and Caldas (2021), uses a fixed threshold (2300 lux) to identify conditions with perceptible glare. Figure 6 shows the correct prediction rate of each method for the tested set of conditions between DGP 0.25 and 0.75. We use a threshold of 0.35 to detect the border of perceptible glare because this is the only level for which a threshold is proposed by the fixed threshold method. Like the rest of the methods requiring input E_v , both GLANCE and the fixed threshold methods are calculated with the 2-phase DDS values.

Across our scenes, GLANCE does not appear to offer a benefit to the fixed threshold method either using a point near the façade (V2) to predict glare at a deeper point (V3) or vice versa. Additionally, a simple DGPs prediction outperforms both methods. Compared to the other tested




Figure 6: Prediction rates averaged across the six scenes for DGP with a detection threshold of 0.35..

measures, the illuminance thresholds only outperform *ClimateStudio*. Given that AGC, *raytraverse*, and imageless DGP appear to offer higher quality predictions without taking significantly more time than generating a high-quality illuminance value (even less in the case of imageless DGP), there is unlikely to be a use case where either of these methods is preferred.

Conclusion

This paper collected several recently developed methods for calculating DGP values faster than conventional image-based approaches and assessed their accuracy across a range of common daylight conditions. While eDGPs have long been the benchmark for faster but still reliable data generation, the computational time required makes multi-point or multi-scenario analyses impractical. Of the five other methods for calculating hourly DGP values evaluated here, only *raytraverse* and AGC have similar errors as eDGPs when compared to data from high-quality reference images across all assessed conditions, both in terms of hourly errors, annual distributions, and threshold prediction accuracy. Additionally, we showed that there is a risk in mischaracterizing zonal glare conditions with using illuminance-based proxy measurements. Imageless DGP takes less time to simulate than an equivalent 2-phase DDS illuminance calculation and offers sufficient accuracy for predicting glare category and annual distributions for scenarios limited to clear glazing and rough specular transmission (where DGP is well correlated with illuminance). Confirming the limitations mentioned by Jones (2019), imageless DGP should not be used with low transmission glazing, fabric roller shades, or specular reflections. *ClimateStudio* offers tremendous potential both for its graphical interface, progressive results, and utilization of hardware-acceleration. We found that the default settings are not able to generate accurate hourly illuminance values and that increasing the sampling does not improve the DGP values sufficiently to accurately categorize glare conditions. Given the progress made in terms of simulation speed, both through novel methods and advancing hardware, we see little reason to continue to rely on poorly performing illuminance-based proxies for simulating and predicting glare with DGP, when methods like AGC and *raytraverse* show that being fast does not mean sacrificing accuracy.

References

- Grobe, L.O., Plöner, D., Noback, A., 2020. Architectural cases for lighting simulation. Digital dataset: <https://c4science.ch/source/AC4LS>
- Giovannini, L., Favoino, F., Lo Verso, V.R.M., Serra, V., Pellegrino, A., 2020. GLANCE (GLare ANnual Classes Evaluation): An approach for a simplified spatial glare evaluation. Building and Environment 186, <https://doi.org/10.1016/j.buildenv.2020.107375>
- Grobe, L.O., Plöner, D., Noback, A., 2020. Architectural cases for lighting simulation. Digital dataset: <https://c4science.ch/source/AC4LS>
- Jones, N.L., 2019. Fast Climate-Based Glare Analysis and Spatial Mapping. in 16th International IBPSA Conference.
- McNeil, A., Burrell, G., 2016. Applicability of DGP and DGI for evaluating glare in a brightly daylit space (2016) ASHRAE and IBPSA-USA Building Simulation Conference, pp. 57-64.
- Pierson, C., Wienold, J., Bodart, M., 2018. Daylight Discomfort Glare Evaluation with Evalglare: Influence of Parameters and Methods on the Accuracy of Discomfort Glare Prediction. Buildings 8, 94. <https://doi.org/10.3390/buildings8080094>
- Santos, L., Caldas, L., 2021. Assessing the glare potential of side-lit indoor spaces: a simulation-based approach. Architectural Science Review 64, 139–152. <https://doi.org/10.1080/00038628.2020.1758622>
- Subramaniam, S., 2017. Daylighting Simulations with Radiance using Matrix-based Methods. LBNL. <https://www.radiance-online.org/learning/tutorials>
- Walkenhorst, O., Luther, J., Reinhart, C., Timmer, J., 2002. Dynamic annual daylight simulations based on one-hour and one-minute means of irradiance data. Solar Energy 72, 385–395. [https://doi.org/10.1016/S0038-092X\(02\)00019-1](https://doi.org/10.1016/S0038-092X(02)00019-1)
- Ward, G.J., 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, pp. 459–472. <https://doi.org/10.1145/192161.192286>
- Wasilewski, S., 2021. Raytraverse: Navigating the Lightfield to Enhance Climate-Based Daylight Modeling 9. SimAUD.
- Wasilewski, S., Grobe, L.O., Wienold, J., Andersen, M., 2022. Efficient Simulation for Visual Comfort Evaluations. Energy and Buildings <https://doi.org/10.1016/j.enbuild.2022.112141>
- Wienold, J., Andersen, M., 2022. Adaptive glare coefficient method for climate-based daylight glare analyses. To be submitted in summer 2022.
- Wienold, J., 2009. Dynamic daylight glare evaluation, in: 11th International IBPSA Conference. pp. 944–951.
- Wienold, J., Iwata, T., Sarey Khanie, M., Errell, E., Kaftan, E., Rodriguez, R., Yamin Garreton, J., Tzempelikos, T., Konstantzos, I., Christoffersen, J., Kuhn, T., Pierson, C., Andersen, M., 2019. Cross-validation and robustness of daylight glare metrics. Lighting Research & Technology <https://doi.org/10.1177/1477153519826003>