Real-Time Visual Comfort Feedback for Architectural Design

NATHANIEL L. JONES, CHRISTOPH F. REINHART

Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT: Today's predictions of visual comfort are based on high-quality physically-based visualization renderings. Unfortunately, designers and practitioners rarely realize the full benefit of physically-based lighting simulation due to the amount of time required for these simulations. Visual comfort analysis is generally performed late in the design process as a form of validation, if at all. We propose a design workflow wherein certain quantitative visual comfort metrics can be displayed immediately to the designer as the scene changes, often before the physically-based visualization reaches a finished quality. In our prototype software, live-updating predictions of daylight glare probability, task luminance, and contrast are presented alongside a progressively rendered image of the scene so that the user may decide when to accept the values and move on with the design process. In most cases, sufficiently accurate results are available within seconds after rendering only a few frames.
Keywords: daylighting, visual comfort, glare, real-time, simulation, rendering

INTRODUCTION

In creative disciplines, flow is a highly focused state of consciousness in which mental tasks seem automatic and effortless. Csikszentmihalyi (1996) identifies nine characteristics of flow, including that it requires immediate feedback toward clear goals and reduces awareness of passing time. He suggests that the sensation of flow causes the enjoyment that people find in creative disciplines. Designers should favor tools that foster flow and therefore lead to the enjoyment of creative tasks.

The conventional approach to predicting visual comfort in buildings is to run a ray tracing simulation with high accuracy settings, wait while the simulation processes, and then repeat as necessary with modifications to the scene and settings. This violates the principles of flow. High-accuracy simulations are too time-consuming to produce immediate feedback, and waiting produces an awareness of time that distracts from the design task. As a result, visual comfort prediction is generally carried out late in the design process in order to validate a completed design, if it occurs at all. For architecture to benefit from daylight as a practical alternative to electric lighting, visual comfort prediction must be available in real time during the design process.

In this paper, we propose a set of predictive visual comfort metrics that can be calculated without interrupting flow and present a prototype software tool to calculate them. We show that several image-based metrics, including vertical eye illuminance (E_v), daylight glare probability (DGP), task area luminance (TAL), and contrast ratio (CR), can be obtained from lower-quality renderings and do not vary significantly as rendering quality increases. Our software prototype initially produces a noisy rendering and progressively refines it through iterative path tracing. Rather than choosing accuracy settings first and then waiting for the simulation to finish before viewing the rendering, the user sees a continually updating image of the rendering alongside plots of the visual comfort metrics. The user decides when the simulation has achieved sufficient accuracy, which often takes only a few seconds.

BACKGROUND

For the past twenty years, RADIANCE has stood out as the most validated simulation software for physically-based rendering and irradiance calculation in building design (Larson & Shakespeare, 1998). Its ability to predict daylit sensor readings has been validated in controlled environments (Mardaljevic, 1995; Reinhart & Herkel, 2000; Mardaljevic, 2001; Reinhart & Walkenhorst, 2001) and buildings (Grynberg, 1989; Ng, et al., 2001; Galasiu & Atif, 2002). Recently, Jones & Reinhart (2015) showed that RADIANCE predicts real scene luminance within the 20% error bound recommended set by the Illuminating Engineering Society of North America (Rea, 2000).

RADIANCE performs light backward ray tracing, where rays emanate from a virtual camera or sensor and bounce around the scene to sample the environment. Wherever a ray intersects a surface, it recursively spawns one or more new rays, depending on the surface material, and gathers their results into a single value that is returned as the parent ray's result (Whitted, 1980). By far the most time consuming portion of this calculation is the diffuse contribution, which requires a large number of sampling rays at each intersection. In an effort to improve rendering quality and speed, the computer graphics community has produced many alternative methods for computing global illumination. One alternative is path tracing, which traces only a single ray from each intersection but does so iteratively in order to build up a complete sampling of the scene (Lafortune & Willems, 1993). Path tracing and many of its extensions offer the benefit that intermediate results may be displayed before the rendering is finished.

While path tracing and other new global illumination methods have received little attention in the architectural community, some effort has been made to achieve faster rendering speeds through RADIANCE. Accelerad is a RADIANCE derivative that performs faster simulations in parallel on the graphics processing unit (GPU) instead of using a single central processing unit (CPU) thread (Jones & Reinhart, 2014; Jones & Reinhart, 2015). This strategy allows hours-long RADIANCE simulations to be run in minutes, but it does not achieve real-time speeds.

Visual Comfort Metrics

Visual comfort prediction in architecture depends on accurately rendering the view of a building occupant. Ray tracing and path tracing produce physically-based renderings with accurate light levels from which visual comfort metrics can be derived. We consider five metrics:

Vertical eye illuminance (E_v) represents the illuminance value of a sensor placed at the observer's eye position and serves as a measure of overall scene brightness. It is calculated as:

$$E_v = \sum_{i=1}^{n} L_{p,i}\omega_{p,i}\cos\theta_{p,i} \tag{1}$$

where $L_{p,i}$ is the luminance of the ith pixel in the field of view, $\omega_{p,i}$ is the solid angle occupied by that pixel, and $\theta_{p,i}$ is the angle to that pixel from the view direction.

Daylight glare probability (DGP) represents the probability that an occupant will experience glare in the given view (Wienold & Christoffersen, 2006). Values greater than 45% correspond to intolerable glare, while those under 35% predict imperceptible glare. Recently, adaptive visual comfort metrics using images taken from multiple vantage points have been proposed to more accurately model building occupant behavior (Jakubiec & Reinhart, 2012). DGP is calculated as:

$$\begin{aligned} DGP = 0.16 + 5.87 \times 10^{-5} E_v + 0.0918 \\ \times \log_{10}\left(1 + \sum_{i=1}^{n} \frac{L_{s,i}^2 \omega_{s,i}}{E_v^{1.87} P_i^2}\right) \end{aligned} \tag{2}$$

where $L_{s,i}$ and $\omega_{s,i}$ are the luminance and solid angle of the ith source, and P_i is the Guth position index representing the eye's sensitivity to the source direction.

Task area luminance (TAL) is the luminance of a user-defined region of the image—typically a work surface for which visibility of a task is important. It is the solid-angle-weighted average of pixel luminance within a region R, calculated as:

$$TAL_R = \sum_{i \in R} L_{p,i}\omega_{p,i} \Big/ \sum_{i \in R} \omega_{p,i} \tag{3}$$

Contrast ratio (CR) measures the contrast between two regions. It is typically used to measure contrast on a computer monitor, where regions H and L represent areas of high and low pixel states. Ratios above 25%, resulting from bright reflections on the monitor, are considered intolerable (ISO, 2008). It is calculated as:

$$CR = \frac{TAL_L}{TAL_H} = \frac{\sum_{i \in L} L_{p,i}\omega_{p,i} / \sum_{i \in L} \omega_{p,i}}{\sum_{j \in H} L_{p,j}\omega_{p,j} / \sum_{j \in H} \omega_{p,j}} \tag{4}$$

We consider one more contrast metric, **RAMMG** (named for its inventors' initials) (Rizzi, et al., 2004). RAMMG computes mean local pixel variation over a subsampled image pyramid structure, or MIP-map, as follows:

$$RAMMG = \frac{1}{m} \sum_{\forall level} \left(\frac{1}{8n} \sum_{i=0}^{n} \sum_{j=0}^{8} \alpha_j \left| L_{p,i} - L_{p,j} \right| \right) \tag{5}$$

where m is the number of pyramid levels, n is the number of pixels in the current level, and α_j is a weight applied to the jth neighboring pixel:

$$\alpha_j = \frac{1}{4 + 2\sqrt{2}} \begin{bmatrix} \sqrt{2}/2 & 1 & \sqrt{2}/2 \\ 1 & & 1 \\ \sqrt{2}/2 & 1 & \sqrt{2}/2 \end{bmatrix} \tag{6}$$

A preliminary study showed high correlation between RAMMG and subjective ratings of images of daylit spaces (Rockcastle & Andersen, 2015). In our study, we use RAMMG as a measure of image quality because noise in low-quality path-traced images takes the form of high local contrast.

METHOD AND IMPLEMENTATION

We modified the RADIANCE source code to perform path tracing instead of ray tracing. By default, RADIANCE allows the user to specify the number of ambient divisions for diffuse sampling. In our implementation, the number of ambient divisions is held at one, and instead of achieving better accuracy through sampling density, we achieve it using multiple rendering passes. The first pass (frame zero) traces only direct and specular paths. Subsequent passes calculate the diffuse component at low accuracy. The results are aggregated and progressively refined for each pixel as follows:

$$L'_{p,n} = L_{p,0} + \frac{1}{n}\sum_{i=1}^{n} L_{p,i} \quad (7)$$

where $L'_{p,n}$ is the luminance of the pixel after n frames.

In order to speed up the rendering of individual frames, we perform path tracing in parallel on a GPU. We implemented path tracing with the OptiX™ ray tracing engine (Parker, et al., 2010). The ray intersection routines, though carried out on the GPU, remain identical to those in RADIANCE, so the results after aggregating a large number of frames are also expected to be identical to RADIANCE. We and ran simulations using NVIDIA® Tesla® K40 graphics accelerators, each with 2880 compute cores.

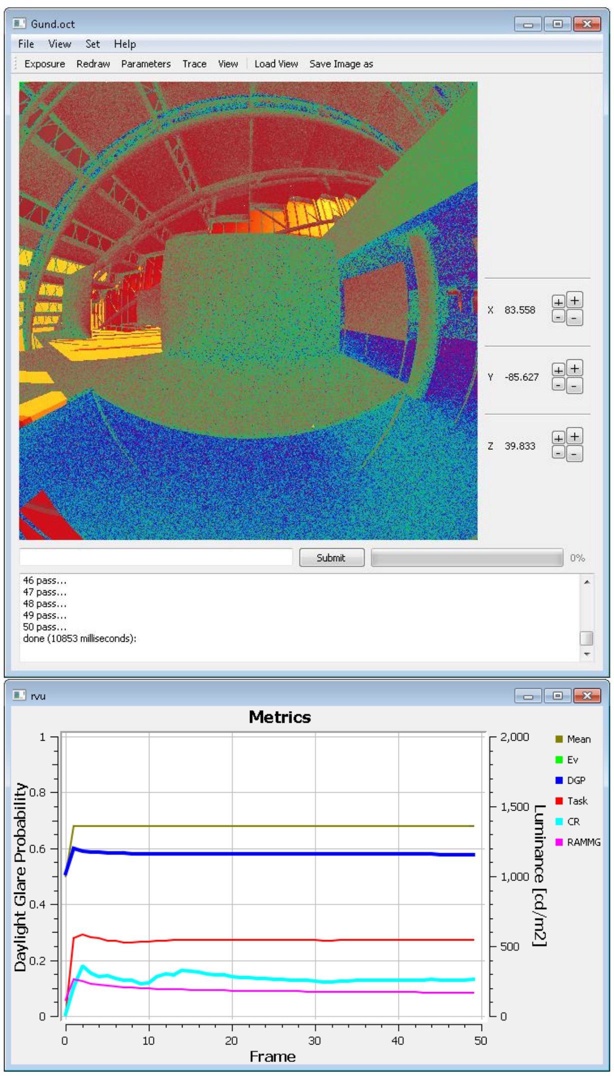

Figure 1: The prototype's user interface presents live views of the rendering in progress and predicted visual comfort values.

Post-processing and calculation of metrics is split between the GPU and CPU. Pixel-level calculations including luminance calculation occur in parallel on the GPU, while the summations required to calculate the visual comfort metrics occur on the CPU. Tone-mapping for scene visualization also occurs on the GPU, and while the high-dynamic range (HDR) pixel values are stored on the GPU and may be saved to a RADIANCE HDR file upon request, only the tone-mapped image is returned to the CPU for display at every frame.

The graphic user interface for our prototype is modified from the *rvu* program included with RADIANCE (Figure 1). The user may choose between photoreal and falsecolor tone-mapping using the same options that are available in Thomas Bleicher's *wxfalsecolor* program. We recommend the use of falsecolor tone-mapping because the HDR extents of many daylit scenes exceed the viewable range on most monitors. In its current implementation, the interface allows a single task area and pair of contrast regions to be monitored by the TAL and CR metrics. However, nothing prevents further development from allowing an unlimited number of regions within the image to be monitored simultaneously. A separate window displays a frame-by-frame history of visual comfort metric values.

RESULTS

We tested our prototype in ten scenes. The scenes were modeled in either Rhinoceros or SketchUp; the Rhinoceros models were exported to RADIANCE format using DIVA-for-Rhino (Jakubiec & Reinhart, 2011), and the SketchUp models were exported using Thomas Bleicher's *su2rad*. Our prototype makes all RADIANCE projections available, but we used a 180° angular fisheye view in all cases because it approximates the human field of vision. The 512×512 pixel images rendered in between 0.15 and 2 seconds per frame, depending on scene complexity. Each simulation was allowed to run through 10,000 frames in order to reach a stable value, much longer than turned out to be necessary. We compare intermediate results to the final value to show how quickly the visual comfort predictions converge on a stable value.

Our method has a clear speed advantage over RADIANCE. The scene shown in Figure 1 rendered its first ten frames in 2 seconds and reached its 100th frame in 22 seconds. On a 3.4 GHz workstation, renderings of comparable quality made in *rvu* by setting the number of ambient divisions to 10 and 100 took 42 and 238 seconds, respectively. Furthermore, our progressive rendering technique makes useful results available at intermediate frames, while *rvu*'s does not. We stress the importance of this speedup as a means to enable flow in the creative process.

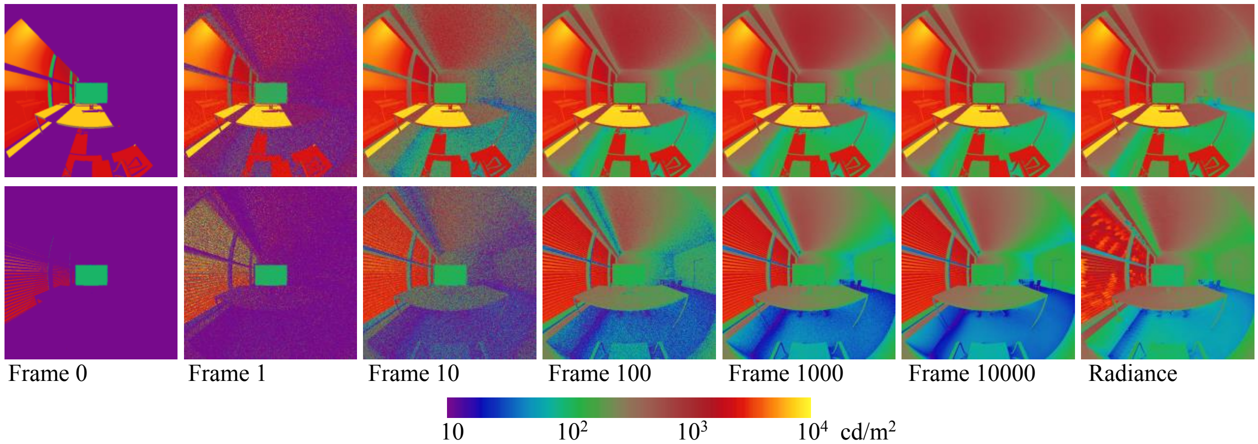

Figure 2: Falsecolor renderings of an office with unshaded windows (top) and venetian blinds (bottom) show a progression from direct contribution only (frame 0) to well-sampled diffuse contribution (frame 10000). Comparison to RADIANCE rendering (right) shows the accuracy of our method.

Image Quality

Progressive path tracing allows the user to observe the scene as it renders. Figure 2 shows intermediate frames from two scenes rendered with our software prototype. The first frames are significantly noisy, but general patterns of illumination are visible by the tenth frame, and little change in illumination is apparent between the 100th and 10,000th frames. Comparison to RADIANCE rendering shows that our method produces accurate luminance distributions; however, the irradiance caching algorithm (Ward & Heckbert, 1992) used by RADIANCE gives the venetian blinds a mottled appearance that is likely to result in inaccurate visual comfort predictions. Path tracing does not suffer from this problem.

The evolution of the RAMMG contrast metric provides and indicator of rendering quality. RAMMG is sensitive to pixel-level noise and decreases as the image quality improves (Figure 3). In the scenes we tested, RAMMG is accurate to within 10-12% of its final value after 100 frames and within 2-3% after 1000 frames. The eventual values reached by RAMMG is related to brightness of light sources in the scene.

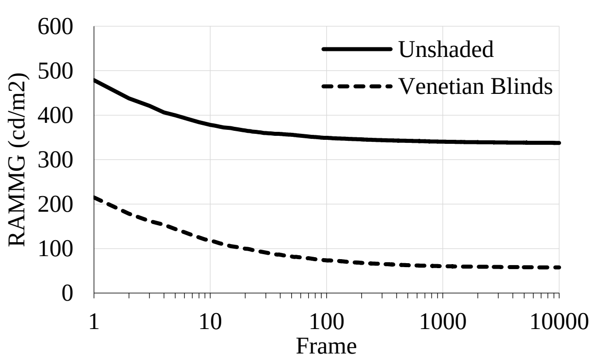

Figure 3: RAMMG for the scenes in Figure 2 converges toward a minimum as rendering progresses.

Vertical Eye Illuminance

In contrast to image quality, E_v changes very little during progressive rendering. Random pixel-level noise tends to cancel itself, resulting in near-constant E_v predictions (Figure 4). In the scenes we tested, E_v was correct within 0.2% of its final value after the first frame and within 0.01% by the tenth frame. Immediate availability of accurate results makes E_v the most compatible metric with flow.

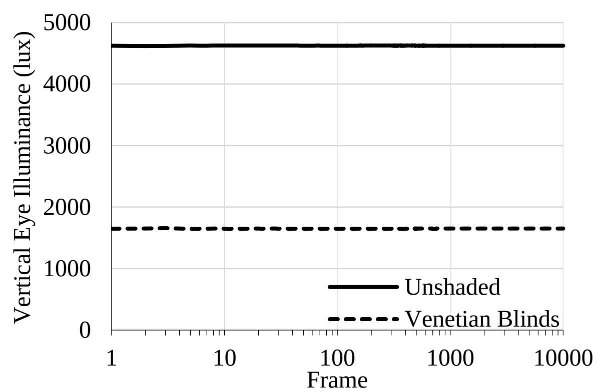

Figure 4: Vertical eye illuminance for the scenes in Figure 2 is constant as rendering progresses.

Daylight Glare Probability

DGP is sensitive to both global light levels and local pixel variation, so its behavior should be between those of RAMMG and E_v. Figure 5 shows renderings created by our prototype using the same model geometry under different sky conditions. Movement of the sun into and out of the field of view results in differing DGP values (Figure 6). DGP reaches its highest value at 11 AM when the sun is directly visible through the window and is lowest under overcast sky conditions.

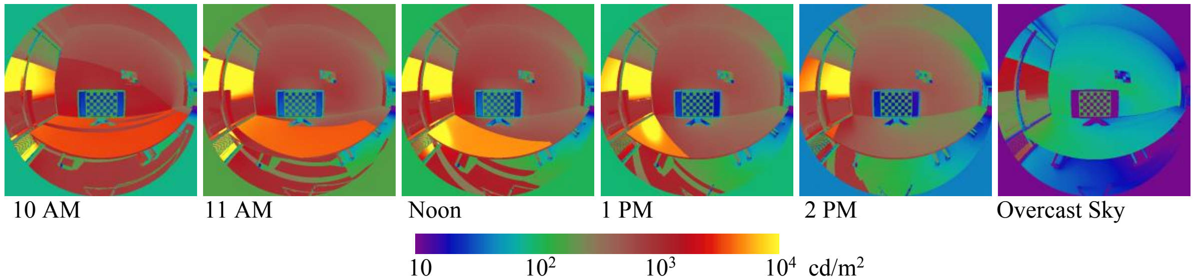

Figure 5: Falsecolor renderings of an office under multiple sky conditions show time-dependent changes in luminance distribution.

Initial predictions from our prototype overestimate DGP because bright random noise is interpreted as a glare source. The initial error is reduced in luminous scenes where real glare sources are more severe. Because our method never underpredicts glare, it will not report false negatives. For scenes with actual DGP in the intolerable zone above 45%, our method produced very little variation as rendering progressed. In the worst case, the predicted DGP was off by 5% after ten frames and by less than 1% after 100 frames. We propose that only when the initial DGP value is in the perceptible glare range between 35% and 45% is it necessary to run the simulation for multiple frames using our method, and even then the number of frames required to reach steady state or drop below the glare threshold is small.

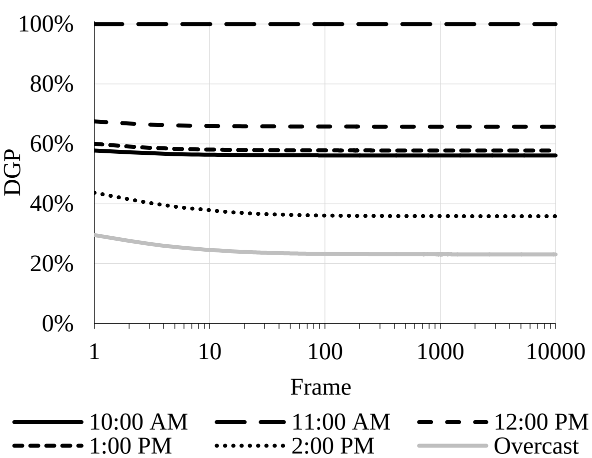

Figure 6: DGP for the scenes in Figure 5 converges downward as rendering progresses. Greater values undergo less change.

Task Area Luminance and Contrast Ratio

As with E_v, random noise is likely to cancel itself when calculating TAL and CR. However, while E_v is calculated over a large area, the user-selected regions for TAL and CR may be quite small (Figure 7). As a result, noise may persist through more frames (Figure 8). For the view used in Figure 5, the task region on the desk occupies 5646 pixels. Error in TAL was under 5% for the first frame and reduced to less than 1% by the tenth frame. The high and low regions of the monitor used for CR calculation each occupy only 32 pixels. Due to the small sample size, error in CR was not reliably below 1% until the 1000th frame. We suggest that the time required to reach convergence should be inversely proportional to the size of the region.

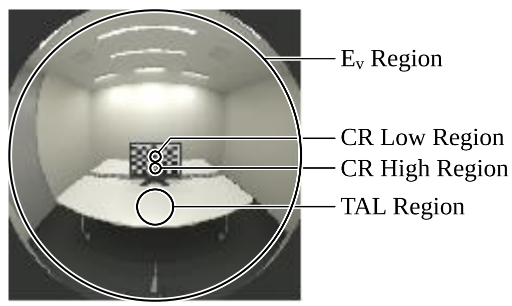

Figure 7: The user-specified regions for calculating TAL and CR are small relative to the entire rendering.

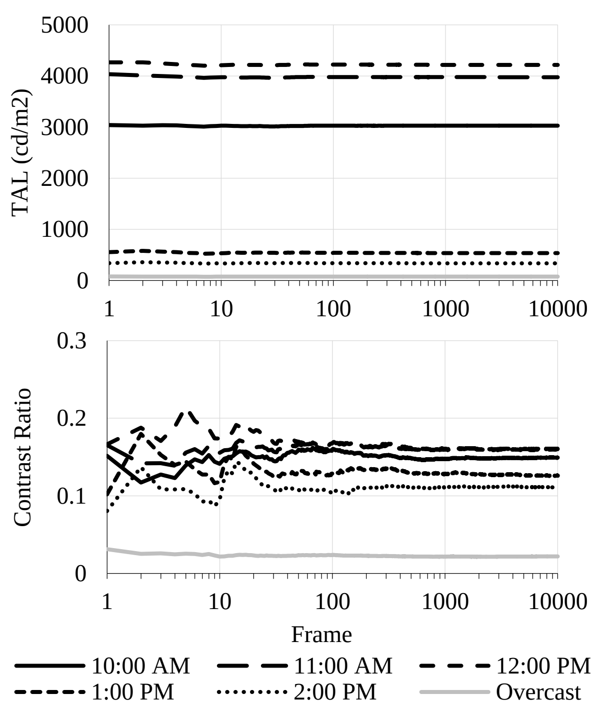

Figure 8: TAL (above) and CR (below) for the scenes in Figure 5 show varying random noise depending on the number of pixels in the region.

CONCLUSION

In this paper, we have demonstrated a software prototype and proof-of-concept to provide architects with visual comfort feedback in real time. Our method uses progressive path tracing to display the current rendering state and calculates visual comfort metrics for each frame. Contrary to conventional wisdom, reasonably accurate visual comfort metrics can be obtained from fast, noise-filled renderings. Given that lighting simulation is expected to produce luminance results within 20% of as-built values (Rea, 2000), the additional error introduced by using unconverged renderings is negligible.

Our prototype opens new avenues for investigation and tool design. It can serve as a platform for evaluating visual comfort and perceptual metrics. The techniques used by our path tracer could lead to progressively updating visualizations of illuminance distribution or climate-based daylighting metrics over large floorplans. In the future, we hope to make validated progressive renderings and visual comfort feedback directly accessible through computer aided design software and to study their effect on real design processes.

Most significantly, we have opened the possibility for immediate visual comfort feedback compatible with the flow of a creative design process. Progressive rendering and graphic display of visual comfort metrics allow users to detect errors and make informed design decisions without interrupting their train of thought. By making these metrics easily accessible to non-expert users, we hope to expand their use in architectural design.

ACKNOWLEDGEMENTS

This research was funded through the Kuwait-MIT Center for Natural Resources and the Environment by the Kuwait Foundation for the Advancement of Sciences. The Tesla K40 accelerators used for this research were donated by the NVIDIA Corporation. J. Alstan Jakubiec provided several of the models and RADIANCE input files.

REFERENCES

Csikszentmihalyi, M., 1996. *Creativity: Flow and the Psychology of Discovery and Invention.* 1st ed. New York: HarperCollins.

Galasiu, A. D. & Atif, M. R., 2002. Applicability of daylighting computer modeling in real case studies: comparison between measured and simulated daylight availability and lighting consumption. *Building and Environment,* 37(4), pp. 363-377.

Grynberg, A., 1989. *Validation of Radiance,* Berkeley, CA: Lawrence Berkeley Laboratories. Document ID 1575.

ISO 9241-303, 2008. *Ergonomics of human-system interaction – Part 303: Requirements for electronic visual displays,* Geneva: International Organization for Standardization.

Jakubiec, J. A. & Reinhart, C. F., 2011. DIVA 2.0: Integrating daylight and thermal simulations using Rhinoceros 3D and EnergyPlus. *Proceedings of Building Simulation 2011: 12th Conference of International Building Performance Simulation Association, Sydney, 14-16 November,* pp. 2202-2209.

Jakubiec, J. A. & Reinhart, C. F., 2012. The 'adaptive zone' – A concept for assessing discomfort glare throughout daylit spaces. *Lighting Research and Technology,* 44(2), pp. 149-170.

Jones, N. L. & Reinhart, C. F., 2014. Physically based global illumination calculation using graphics hardware. *Proceedings of eSim 2014: The Canadian Conference on Building Simulation,* pp. 474-487.

Jones, N. L. & Reinhart, C. F., 2015. Validation of GPU lighting simulation in naturally and artificially lit spaces. *Proceedings of BS2015: 14th Conference of International Building Performance Simulation Association, Hyderabad, India, December 7-9,* pp. 1229-1236.

Lafortune, E. P. & Willems, Y. D., 1993. Bi-directional path tracing. *Proceedings of CompuGraphics,* Volume 93, pp. 145-153.

Larson, G. W. & Shakespeare, R., 1998. *Rendering with Radiance: The Art and Science of Lighting Visualization.* San Francisco: Morgan Kaufmann Publishers, Inc.

Mardaljevic, J., 1995. Validation of a lighting simulation program under real sky conditions. *Lighting Research and Technology,* 27(4), pp. 181-188.

Mardaljevic, J., 2001. The BRE-IDMP dataset: a new benchmark for the validation of illuminance prediction techniques. *Lighting Research and Technology,* 33(2), pp. 117-136.

Ng, E. Y.-Y., Poh, L. K., Wei, W. & Nagakura, T., 2001. Advanced lighting simulation in architectural design in the tropics. *Automation in Construction,* 10(3), pp. 365-379.

Parker, S. G. et al., 2010. OptiX: A general purpose ray tracing engine. *ACM Transactions on Graphics – Proceedings of ACM SIGGRAPH 2010,* 29(4).

Rea, M. S., 2000. *The IESNA lighting handbook: reference & application.* 10th ed. New York: Illuminating Engineering Society of North America.

Reinhart, C. F. & Herkel, S., 2000. The simulation of annual daylight illuminance distributions – a state-of-the-art comparison of six RADIANCE-based methods. *Energy and Buildings,* 32(2), pp. 167-187.

Reinhart, C. F. & Walkenhorst, O., 2001. Validation of dynamic RADIANCE-based daylight simulations for a test office with external blinds. *Energy and Buildings,* 33(7), pp. 683-697.

Rizzi, A., Algeri, T., Medeghini, G. & Marini, D., 2004. A proposal for Contrast Measure in Digital Images. *Conference: CGIV 2004 – Second European Conference on Color in Graphics, Imaging and Vision,* pp. 187-192.

Rockcastle, S. & Andersen, M., 2015. Human perceptions of daylight composition in architecture: A preliminary study to compare quantitative contrast measures with subjective user assessments in HDR renderings. *Proceedings of the 14th International Conference of the International Building Performance Simulation Association,* pp. 1205-1212.

Ward, G. J. & Heckbert, P. S., 1992. Irradiance Gradients. *Third Eurographics Workshop on Rendering,* Volume 8598.

Whitted, T., 1980. An improved illumination model for shaded display. *Communications of the ACM,* 23(6), pp. 343-349.

Wienold, J. & Christoffersen, J., 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), p. 743–757.