


SPECIAL ISSUE PAPER

High-fidelity iridal light transport simulations at interactive rates

Boris Kravchenko | Gladimir V. G. Baranoski  | Tenn Francis Chen | Erik Miranda |
Spencer R. Van Leeuwen

Natural Phenomena Simulation Group,
University of Waterloo, 200 University
Avenue West Waterloo, Ontario, Canada

Correspondence

Gladimir V.G. Baranoski, Natural
Phenomena Simulation Group, University of
Waterloo, 200 University Avenue West,
Waterloo, Ontario N2L 3G1, Canada.
Email: gvgbaran@gmail.com

Abstract

Predictive light transport models based on first-principles simulation approaches have been proposed for complex organic materials. The driving force behind these efforts has been the high-fidelity reproduction of material appearance attributes without one having to rely on the manipulation of ad hoc parameters. These models, however, are usually considered excessively time consuming for rendering and visualization applications requiring interactive rates. In this paper, we propose a strategy to address this open problem with respect to one of the most challenging of these organic materials, namely the human iris. More specifically, starting with the configuration of a predictive iridal light transport model on a parallel-computing platform, we analyze the sensitivity of iridal appearance attributes to key model running parameters in order to achieve an optimal balance between fidelity and performance. We believe that the proposed strategy represents a step toward the real-time and predictive synthesis of high-fidelity iridal images for rendering and visualization applications, and it can be extended to other organic materials.

KEYWORDS

biophysical phenomena, human iris, simulation, visualization

1 | INTRODUCTION

The generation of realistic images has always been a challenge for the computer graphics community, especially when these images depict organic materials. One of the main research avenues in this area involves the correct modeling of appearance attributes.¹ Another avenue involves the use of specialized graphics processing units (GPUs) to enable the fast generation of believable images for interactive applications (e.g., Dorsey, Rushmeier, and Sillion and Chiang and Fyffe^{2,3}). As the advances in this area continue, the demand for higher correctness to cost ratios in the image synthesis process has been increasing accordingly. Recently, this trend has been particularly intensified by initiatives aimed at the development of hyperspectral visualization and rendering systems.^{4,5}

As stated by Greenberg et al.,⁶ predictability is instrumental to increase the fidelity⁷ of the image synthesis process.⁶ It is

also the key for making it less dependent on manual tweaks.⁸ To achieve predictable rendering results controlled by biophysically meaningful parameters, one often needs to resort to light transport models developed using first-principles simulation approaches. Due to their performance overhead, however, these models are usually not considered for use in rendering and visualization applications with turnaround times on the order of milliseconds.

Our research addresses these issues while focusing on the human iris, arguably one of the most noticeable organic materials employed in the creation of virtual characters. As pointed out by Lee et al.⁹ and Walt Disney himself,¹⁰ natural-looking eyes are quite desirable for entertainment applications depicting close-ups of human faces. From a life sciences perspective, there is also a strong interest in the visualization of changes in the human iris's chromatic attributes determined by variations on the attenuation of light within the iridal tissues. For example, in the medical field, a number of studies

have related the incidence of several ocular diseases, such as the degeneration of ocular tissues¹¹ and melanoma,¹² to these attributes. The model known as *iridal light transport* (ILIT) was specifically designed to support these computer graphics and interdisciplinary applications involving the predictive simulation of light interactions with the human iris.¹³ Its ray optics formulation based on Monte Carlo methods, however, is bound to incur relatively high processing times.

In this paper, we demonstrate that predictive light transport simulations leading to realistic depictions of the human iris can be executed at interactive rates. Initially, to make inroads toward this objective, we have configured the ILIT simulation algorithms on the GPU using compute unified device architecture (CUDA).¹⁴ Although substantial performance gains can be obtained using this platform, they are insufficient for achieving our objective. For this reason, we have focused our efforts on analyzing the sensitivity of iridal appearance attributes to changes in key model running parameters, namely spectral resolution and number of sample rays. Our findings indicate that the selection of appropriate values for these parameters is essential for obtaining a high level of appearance fidelity while taking full advantage of the performance gains resulting from the configuration of ILIT on the GPU. We believe that such a strategy can be used to obtain similar results for other complex organic materials.

The remainder of this paper is organized as follows. In the next section, we provide a concise review of related work. We then describe the experimental framework employed in the analysis of the effects of key ILIT running parameters on the fidelity of modeled appearance attributes. Afterwards, we present the results of our research and discuss their practical implications. Finally, we conclude the paper and indicate directions for future work.

2 | RELATED WORK

Different approaches have been employed by the computer graphics community to generate realistic images of ocular structures, in particular the human iris. For example, Sager et al.¹⁵ proposed the rendering of an anatomically detailed eye model to be employed in a surgical simulator. In their work, the iris was represented using a Gouraud-shaded polygon with colors specified by a color ramp. Lefohn et al.¹⁶ presented the first biologically motivated algorithm specifically designed for the rendering of realistic-looking irides. Their modeling approach was based on an ocular prosthetics methodology where several multi-transparent layers are combined to create an artificial iris. Wecker et al.¹⁷ proposed an image-based technique that consists in decomposing photographs of the human iris into several components. These are subsequently recombined to generate distinct syn-

thetic iridal images. Lam and Baranoski¹³ then presented ILIT as the first biophysically-based appearance model for the human iris. We remark that we employ this model as the instrument of our investigation. However, to conserve space, we refer the reader interested in its detailed description to the publications where it was originally described.^{13,18}

After the development of ILIT, several works, mostly targeting the generation of believable iridal images in real time, followed. For example, Francois et al.^{19,20} introduced an image-based method for estimating both iridal morphology and optical parameters from iridal photographs. The resulting iridal layered model was then employed in the real-time rendering of believable eye images using the GPU. Pamplona et al.²¹ proposed an image-based model for iridal pattern deformation to be used in conjunction with a physiologically-based model for pupil light reflex. Chiang and Fyffe² used a GPU-adapted normal mapping approach to generate believable eye representations for real-time facial rendering purposes. The appearance of the human iris was modeled employing iridal chromatic attributes obtained from photographs along with artistic tools to emulate effects like caustics and refraction darkening. Jimenez et al.³ presented a real-time shader-based eye rendering solution for use in interactive applications. More recently, Bérard et al.²² presented a system to capture the geometry and texture of ocular structures.

Although not directly related to the generation of realistic images of ocular structures, for completeness, it is worth mentioning the model of individual photoreceptors found in the human eye proposed by Deering.²³ This model was developed as part of a framework for investigating how changes in rendering algorithms and in image display devices are related to artifacts visible to human observers.

3 | EXPERIMENTAL FRAMEWORK

In this section, we present the data and procedures employed in our search for a practical strategy leading to the generation of high-fidelity modeled appearance attributes at interactive rates. We also concisely outline the GPU-specific optimizations performed during the implementation of the GPU version of ILIT. For the sake of conciseness, the reader interested in more details about the steps involved in the reconfiguration of the ILIT algorithms on the GPU using CUDA is referred to a supporting document.²⁴

3.1 | Data and procedures

In our experiments, we considered three iris specimens with distinct levels of pigmentation. These specimens, henceforth

TABLE 1 Three sets of ILIT parameters employed to characterize the distinct iridal specimens considered in this work

Parameters	Darkly pigmented	Moderately pigmented	Lightly pigmented
Eumelanin mass (mg)	2.53E-2	6.33E-3	1.00E-4
Pheomelanin mass (mg)	6.90E-3	1.75E-3	3.00E-5
ABL eumelanin (%)	90	50	10
ABL pheomelanin (%)	90	50	10
SL blood (%)	7	4	1

Note. ABL = anterior border layer; SL = stromal layer; ILIT = iridal light transport.

referred to as lightly pigmented, moderately pigmented, and darkly pigmented, are characterized by the biophysical parameters depicted in Table 1. The values assigned for these parameters correspond to actual biophysical data used in the characterization of the main iridal layers, namely the anterior border layer and the stromal layer. The reader interested in more detailed information about their sources is referred to the original publications describing the ILIT model.^{13,18} Similarly, for the spectral absorption coefficients of the pigments considered in our implementations of ILIT, we used measured data²⁵ also reported in the scientific literature. We note that the parameters depicted in Table 1 were those effectively modified to characterize the three selected specimens, while the remaining model parameters (fully listed on ILIT website²⁶) were kept fixed during our simulations.

The ILIT model can provide radiometric readings for distinct illumination and collection geometries. For example, one can obtain directional-hemispherical reflectance quantities by integrating the outgoing rays with respect to the collection hemisphere using a virtual spectrophotometer.²⁷ These quantities were used in our analyses of the sensitivity of modeled readings to changes in key model running parameters, namely spectral resolution (denoted by s_r) and the number of sample rays (denoted by n_r). More specifically, we have computed iridal directional-hemispherical reflectance curves within the spectral region between 400–700 nm for the selected specimens (Figure 1). In the computation of these curves, we considered an angle of incidence of 0° and distinct values for s_r and n_r . We note that the spectral iridal reflectances computed for a given specimen tend to converge to specific values as n_r increases. We employed the curves associated with these values as the fidelity reference curves in our analyses. To obtain these reference curves, we used s_r equal to 1 nm (301 sample wavelengths) and n_r equal to 10^6 rays. It is also worth mentioning that the reflectance curves depicted in this work, which were obtained through the online version of the GPU-based implementation of the ILIT model,²⁸ match those that can be obtained through the online version of its original implementation (central processing unit [CPU]-based)²⁶ using the same parameter values.

We also generated iris swatches to complement our analyses (Figures 2 and 3). The swatches' chromatic attributes




FIGURE 1 Comparisons of modeled reflectance curves computed for the selected iris specimens using the central processing unit (CPU)- and graphics processing unit (GPU)-based implementations of iridal light transport (ILIT). From top to bottom: darkly pigmented, moderately pigmented, and lightly pigmented specimen, respectively. The curves were computed using distinct numbers of rays. From left to right: 10^3 , 10^4 , 10^5 , and 10^6 , respectively. All curves were computed considering a spectral resolution of 1 nm



FIGURE 2 Iris swatches generated using modeled reflectance data computed considering different spectral resolutions. From top to bottom: swatches generated for the darkly, moderately, and lightly pigmented specimens, respectively. Leftmost column: employing data sampled at three wavelengths (463, 504, and 600 nm). Subsequent columns, from left to right: employing data sampled from 400 to 700 nm at every 50, 30, 10, and 1 nm, respectively. The simulations employed to obtain the modeled data were carried out using 10^6 rays



FIGURE 3 Iris swatches generated using modeled reflectance data computed considering different numbers of sample rays. From top to bottom: swatches generated for the darkly, moderately, and lightly pigmented specimens, respectively. From the leftmost column to the fourth column: swatches generated using modeled data obtained considering a spectral resolution of 30 nm and 10^3 , 10^4 , 10^5 , and 10^6 sample rays, respectively. Rightmost column: reference solutions obtained considering a spectral resolution of 1 nm and 10^6 rays

were obtained from the convolution of the illuminant spectral power distribution, the modeled reflectance data, and the broad spectral response of the human photoreceptors.²⁹ This last step was performed employing a standard XYZ to sRGB conversion procedure.³⁰ Unless otherwise stated,

the iris swatches and the additional images presented in this work were rendered considering the CIE standard D65 illuminant.³¹

Surely, the spectral responses of a given iridal specimen can vary from one measurement point to another. Although

one could modulate the values of the biophysical parameters affecting its chromatic attributes according to the characteristics of its underlying iridal structure, this process would require detailed information about the specimen's morphology (e.g., tissue thickness variations) that is not readily available. Hence, for practical purposes, we have elected to use iridal texture variations to modulate the lightness of the colors derived from the modeled reflectance data. For consistency, we employed the same iridal texture for all swatches.

To assess appearance fidelity, we initially visually compared swatches (Figure 2) generated using modeled reflectance data computed considering distinct spectral resolutions and n_r equal to 10^6 with their respective swatch reference, which was generated using the specimen's fidelity reference curve. We assigned to s_r integer factors of 300 nm (10, 30, and 50 nm) to ensure a uniform sampling of the visible spectral region of interest (from 400 to 700 nm).

Besides visual inspection, we also employed a device-independent CIE-based metric to compare modeled results obtained using different combinations of simulation running parameters. More specifically, we computed the CIELAB differences³² between modeled chromatic attributes and their corresponding reference values. These differences are defined as

$$\Delta E_{ab}^* = \sqrt{(d_L^2 + d_a^2 + d_b^2)}, \quad (1)$$

where d_L , d_a , and d_b represent the differences $L_1^* - L_2^*$, $a_1^* - a_2^*$, and $b_1^* - b_2^*$, respectively, in which L^* , a^* , and b^* correspond the CIELAB color space dimensions. These are calculated for the modeled chromatic attributes (indicated by subscript 1) and their respective reference values (indicated by subscript 2). We performed these calculations using standard colorimetric formulas³³ and setting the white point to the CIE D65 standard illuminant.³¹

We then repeated the comparisons for swatches (Figure 3) generated using modeled reflectance data computed considering different numbers of sample rays and the spectral resolution determined in the previous set of comparisons. The purpose of these two sets of comparisons was to find an optimal balance between appearance fidelity and performance (reduced computational time). By performing the comparisons in this order, we have attributed more importance to s_r . This choice was motivated by the fact that this parameter has a larger impact on the qualitative reproduction of appearance attributes.

To assess the performance gains associated with the execution of iridal light transport simulations on the CUDA platform, we have compared the framework's GPU-based implementation of ILIT with a multithreaded (C++)

CPU-based implementation of this model. These two implementations are henceforth referred to as ILIT-GPU and ILIT-CPU, respectively. Both implementations were tested on a server running Ubuntu 14.04 LTS with two hyper-threaded Intel Xeon E5-2630 2.4 GHz CPUs (comprising a total of 16 cores, 32 threads), 64 GB of RAM and a GTX 980 GPU. For completeness, we outline the GPU-specific optimizations that were performed as part of this work in the next subsection.

3.2 | GPU-specific optimizations

The GTX 980 is able to execute blocks with up to 1,024 threads in each block, and the number of threads per block should be a multiple of the thread warp size (32) to maximize hardware utilization.³⁴ Through empirical testing,²⁴ we elected to employ 3,200 blocks, with each block containing 32 threads. This combination yields roughly 10^5 threads, which is the number of sample rays normally used to obtain asymptotically convergent simulated spectrophotometric readings.²⁷ For simulations employing such a higher number of sample rays (usually a multiple, denoted by k , of 10^5), we run these simulations k times on each thread.

We have also replaced all double precision floats in the CPU-based implementation of ILIT by single-precision floats. This was motivated by the fact that consumer GPUs tend to have much faster single-precision performance compared to double precision. A GTX 980 is capable of executing $32\times$ more single-precision floating point instructions than double-precision floating point instructions in a single clock cycle.³⁵ In addition, all math functions (cos, sin, log ...) were replaced by intrinsic single-precision functions provided by the CUDA API. Although these functions are less precise than their C++ counterparts, they use fewer instructions. As we demonstrate in the next section, the use of single-precision floating point functions is sufficient for obtaining simulation results with the same level of convergence and fidelity observed in the results provided by the (C++) CPU-based implementation of ILIT.

Finally, we optimized thread occupancy (the ratio between the number of currently executing thread warps per streaming multiprocessor unit to the maximum number of thread warps). This optimization was attained by selecting an appropriate number of registers to be allocated for a thread. A GTX 980 can have up to 2,048 threads in-flight on a single streaming multiprocessor unit, assuming every thread uses no more than 32 registers. Setting the number of registers to 32, however, caused excessive register spillover to occur. Again through empirical testing,²⁴ we determined that, by using 64 registers, we can improve performance at the expense of having a smaller number of active threads.

4 | RESULTS AND DISCUSSION

Initially, we compared the modeled reflectance curves provided by the ILIT-GPU and ILIT-CPU implementations in order to assess the correctness of the former. As stated in the previous section, ILIT-GPU employs single-precision floating point arithmetic to maximize performance. Hence, one might expect that it would provide results with a higher level of error than those provided by ILIT-CPU. The graphs presented in Figure 1 demonstrate, however, that the effects of precision differences are negligible in comparison with random fluctuations associated with the stochastic nature of the ILIT model. We note that such random fluctuations are expected, especially when one sets a relatively low value for n_r . As n_r is increased, they are reduced accordingly, and the curves computed by both ILIT implementations tend to show a close agreement.

We then proceeded to assess the computational gains that can be obtained using the reconfigured ILIT algorithms. As illustrated by the data presented in Table 2, ILIT-GPU outperformed the multithreaded CPU-based implementation of ILIT. We note that ILIT-GPU has an initialization time of 200 ms, which was included in Table 2. It is also worth observing that the computational times increase for both

TABLE 2 Running times recorded for the light transport simulations performed for the selected specimens using the GPU-based (ILIT-GPU) and the multithreaded CPU-based (ILIT-CPU) implementations of the ILIT model. The simulations were carried out considering distinct numbers of rays (n_r) and a spectral resolution of 1 nm

	ILIT-GPU (s)	ILIT-CPU (s)
	Darkly pigmented	
n_r		
10^3	0.39	1.80
10^4	0.40	2.87
10^5	0.75	12.78
10^6	4.04	79.33
	Moderately pigmented	
n_r		
10^3	0.41	1.71
10^4	0.46	3.70
10^5	1.12	19.62
10^6	7.54	110.29
	Lightly pigmented	
n_r		
10^3	0.46	1.81
10^4	0.54	4.49
10^5	1.62	25.04
10^6	12.18	149.27

Note. ILIT = iridal light transport; GPU = graphics processing unit; CPU = central processing unit.

implementations as the level of pigmentation is reduced. This can be explained by the fact that the random walks tend to take more steps to terminate when the probability of light absorption is reduced.

For completeness, we also provide the corresponding speedup factors in Table 3. It is worth noting that the recorded speedup factors of the ILIT-GPU with respect to a single threaded CPU-based implementation of ILIT were on the order of 100×.

For applications aimed at providing predictable results at interactive rates, one should strive to maintain high levels of both fidelity and performance. With this goal in mind, we investigated the sensitivity of modeled chromatic attributes to changes in the key model running parameters, namely s_r and n_r . We started with the former as we describe next.

Traditionally, image synthesis systems employ three sample wavelengths, whose selection may vary from one user to another. For comparison purposes, we considered three wavelengths, namely 463, 504, and 600 nm, representative of the blue, green, and red regions of the light visible spectrum, respectively. As it can be verified by visually inspecting the iris swatches presented in Figure 2, the results obtained when one considers only three sample wavelengths may differ markedly from more comprehensive spectral solutions. They also indicate that the results tend to converge to a reference solution as we increase the number of sample wavelengths (from 3 to 7, 11, 31, and 301). However, in order to obtain an image with a high appearance fidelity with respect to this reference solution, one does not necessarily need to resort to an extremely fine, and consequently costly, spectral resolution. For example, as it can also be observed in Figure 2, swatches obtained using a sampling interval of 30 nm (11 wavelengths) are visually indistinguishable from their respective reference swatches obtained considering an interval of 1 nm (301 wavelengths).

In order to quantitatively assess the fidelity of the swatches generated using modeled data obtained considering s_r equal

TABLE 3 Speedup factors of the ILIT-GPU with respect to the multithreaded CPU-based implementation of the ILIT model. These factors refer to the iridal light transport simulations performed for the selected specimens, which were carried out considering distinct numbers of rays (n_r) and a spectral resolution of 1 nm

n_r	Darkly pigmented	Moderately pigmented	Lightly pigmented
10^3	4.57×	4.15×	3.94×
10^4	7.10×	7.90×	8.19×
10^5	16.95×	17.50×	15.44×
10^6	19.61×	14.62×	12.24×

Note. ILIT = iridal light transport; GPU = graphics processing unit; CPU = central processing unit.

TABLE 4 CIELAB ΔE_{ab}^* differences computed for the selected specimens' swatches (depicted in Figure 2), which were generated using modeled data computed considering three distinct spectral resolutions (s_r , given in nanometers) and 10^6 rays. The color differences were calculated with respect to the specimens' specific reference swatches (also depicted in Figure 2, rightmost column), which were generated using modeled data computed considering a spectral resolution of 1 nm

s_r	Darkly pigmented	Moderately pigmented	Lightly pigmented
50	1.59	5.75	6.77
30	0.27	0.33	0.38
10	0.05	0.11	0.05



FIGURE 4 Image rendered using modeled iridal reflectance data

to 10, 30, and 50 nm, we computed their CIELAB ΔE_{ab}^* differences Equation (1) with respect to the corresponding reference solutions. As expected, these differences (Table 4) decreased as we reduced the spectral sampling intervals. Moreover, for s_r equal to 30 nm, the values are smaller than 2.3, the experimentally determined perceptibility threshold for CIELAB chromatic differences.^{32,36} This quantitatively demonstrates that the swatches generated considering a sampling interval of 30 nm show no noticeable differences with respect to their reference swatches.

After determining a suitable value for s_r , we examined the impact of n_r . Again, as it can be inferred by visually inspecting the swatches presented in Figure 3, a combination considering a spectral resolution of 30 nm and 10^4 rays (second column from the left) is sufficient to obtain results which closely agree with their corresponding reference solutions (rightmost column). This aspect is also supported by the CIELAB ΔE_{ab}^* differences Equation (1) computed for the swatches generated using these parameter values. More specifically, when compared to the experimentally determined threshold of 2.3,^{32,36} the magnitude of these differences, namely 1.16, 0.63, and 0.57 for the darkly, moderately, and lightly pigmented specimens, respectively, indicates that this combination of parameter values can be used to obtain chromatic attributes that are indistinguishable from their corresponding reference solutions.

In terms of performance, setting s_r equal to 30 nm and n_r equal to 10^4 rays enables us to reduce the time spent in the simulations to 3.72, 6.30, and 9.71 ms, respectively, which are values below a latency time of 16 ms (associated with an update rate of 60 Hz) appropriate for interactive applications.³⁷ We note that these values do not include the GPU initialization time mentioned earlier because this step needs to be performed only once for interactive applications. In other words, it does not need to be repeated for each frame rendered using modeled iridal reflectance data.

We fully recognize that the synthesis of realistic-looking images of human eyes is a challenging process. As such, it requires due attention to be given to all stages of the rendering pipeline, starting with the selection of geometrical models and textures representing the distinct ocular structures. Clearly, these aspects are beyond the scope of this work, which focuses on the iris's spectral responses. However, we remark that these responses play a central role in this process. Hence, we believe that the proposed strategy can effectively contribute to the interactive generation of images that are both believable and predictable depictions of the human eye. This aspect is illustrated by the images presented in Figures 4, 5, and 6,




FIGURE 5 Images rendered using modeled iridal reflectance data computed for the darkly (left), moderately (center), and lightly (right) pigmented iris specimens




FIGURE 6 Images rendered using modeled iridal reflectance data computed for the darkly (left) and moderately (right) pigmented iris specimens and considering two distinct CIE standard illuminants: A (left) and D50 (right)

which were rendered using a standard ray tracer and modeled iridal reflectance data obtained considering the selected combination of simulation running parameters ($s_r = 30$ nm and $n_r = 10^4$ rays).

With respect to usability of the accelerated version of ILIT developed for this work, we note that it can be seamlessly integrated into rendering systems that support GPU-based implementations. Alternatively, it can be employed to provide iridal appearance data on demand for image synthesis applications running on different computational environments.

Finally, it is worth stressing that only a small subset of the ILIT parameters (Table 1) needs to be modified so that one can obtain markedly distinct chromatic attributes for iridal specimens. This is particularly convenient when one wants to streamline the image synthesis process. If necessary, however, its detailed parameter space allows experimentation with a wide range of biophysical factors affecting iridal appearance. Hence, besides the support to the efficient rendering of realistic eye images for artistic and entertainment applications, the ILIT-GPU can potentially be employed in educational and scientific applications involving rapid visualizations of iridal appearance variations as illustrated in the supplementary video.³⁸

5 | CONCLUSION AND FUTURE WORK

Distinct trends can be observed in image synthesis research. At one end of the spectrum, efforts have been directed toward the fast generation of believable images for a large variety of artistic and entertainment applications, from movie effects to video games. At the other end of the spectrum, different initiatives have been aimed at the generation of images with the highest possible level of biophysical correctness. Among these initiatives, one can include the development of predictive appearance models for a large variety of natural and

man-made materials. The simultaneous realization of high levels of appearance fidelity and performance is usually problematic, however, because researchers often have to sacrifice one to get the other.


In this paper, we presented a strategy aimed at conciliating these seemingly conflicting goals. We used an experimental framework that incorporates advances at both ends of the rendering research spectrum in order to demonstrate its feasibility. More specifically, after reconfiguring the ILIT algorithms on the GPU using CUDA, we examined the trade-offs involving different combinations of key model running parameters. Our findings indicate that high-fidelity appearance attributes can be obtained through first-principles light transport algorithms executed at interactive rates. Accordingly, these attributes can be employed in the efficient synthesis of iridal images that are not only realistically looking, but also predictable.

We believe that our experimental investigation can be viewed as a proof of concept for a broad range of interactive applications involving the rendering and visualization of different organic materials. Accordingly, as future work, we intend to demonstrate that the proposed strategy for balancing fidelity and performance can lead to similar results for other organic tissues like human skin and blood.

REFERENCES

1. Dorsey J, Rushmeier H, Sillion F. Digital modeling of material appearance, Burlington, MA, USA: Morgan Kaufmann/Elsevier; 2007.
2. Chiang M, Fyffe G. Realistic real-time rendering of eyes and teeth: University of Southern California, Los Angeles, CA, USA; 2010. Technical Report ICT-TR-01-2010.
3. Jimenez J, Danvoye E, Pahlen J. Separable subsurface scattering and photorealistic eyes rendering. Advances in Real-Time Rendering in Games, SIGGRAPH Course, Los Angeles, CA, USA; 2012.
4. Kim MH, Harvey TA, Kittle DS, et al. 3D omaging spectroscopy for measuring hyperspectral patterns on solid objects. ACM

- Trans Graph. 2012;31(4):38:1–38:11. presented at the SIGGRAPH 2012.
5. Baranoski GVG. Hyperspectral modeling of material appearance: General framework, challenges and prospects. *RITA*. 2015;22(2):203–232.
 6. Greenberg DP, Arvo J, Lafortune E, et al. A framework for realistic image synthesis. *SIGGRAPH, Annual Conference Series*, Los Angeles, CA, USA; 1997. p. 477–494.
 7. Gross DC. Report from the fidelity implementation study group. *Simulation Interoperability Workshop, Simulation Interoperability and Standards Organization*; Orlando, FL, USA; 1999. Paper 99S-SIW-167.
 8. Fournier A. The tiger experience. *Workshop on Rendering, Perception and Measurement*, Cornell University; USA. p. 1999.
 9. Lee S, Badler J, Badler N. Eyes alive. *ACM Trans Graph*. 2002;21(3):637–644.
 10. Watt A, Watt M. *Advanced animation and rendering techniques*. New York, N.Y., USA: Addison-Wesley; 1992.
 11. Delori FC, Dorey CK, Fitch KA. Characterization of ocular melanin by iris reflectometry. *Invest Ophthalmol Vis Sci*. 1991;32(1-4):1141.
 12. Regan S, Egan KM, Gragoudas ES. Iris color as a risk factor in uveal melanoma. *Invest Ophthalmol Vis Sci*. 1997;38(4):3771–3771.
 13. Lam MWY, Baranoski GVG. A predictive light transport model for the human iris. *Comp Graph Forum*. 2006;25(3):359–368.
 14. Chei S, Boyer M, Meng J, Tarjan D, Sheaffer JW, Skadron K. A performance study of general-purpose applications on graphics processors using CUDA. *J Parallel Distrib Comput*. 2008;68:1370–1380.
 15. Sagar MA, Bullivant D, Mallinson GD, Hunter PJ. A virtual environment and model of the eye for surgical simulation. *SIGGRAPH, Annual Conference Series, ACM*, Orlando, FL, USA; 1994. p. 205–212.
 16. Lefohn A, Budge B, Shirley P, Caruso R, Reinhard E. An ophthalmologist's approach to human iris synthesis. *IEEE Comput Graph*. 2003;23(6):70–75.
 17. Wecker L, Samavati F, Gavrilova M. Iris synthesis: A reverse subdivision technique. *3rd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*, Dunedin, New Zealand; 2005. p. 121–125.
 18. Lam MWY. *ILIT: A predictive light transport model for the human iris*. Master's thesis, University of Waterloo, Canada, August 2006.
 19. Francois GP, Breton G, Bouatouch K. *Anatomically accurate modeling and rendering of the human eye*. Rennes, France: INRIA; 2007. Technical Report.
 20. Francois GP, Breton G, Bouatouch K. Image-based modeling of the human eye. *IEEE T Vis Comput Gr*. 2009;15(5):815–827.
 21. Pamplona VF, Oliveira MM, Baranoski GVG. Photorealistic models for pupil light reflex and iridal pattern deformation. *ACM Trans Graph*. 2009;28(4):106:1–106:12.
 22. Bérard P, Bradley D, Nitti M, Beeler T, Gross MH. High-quality capture of eyes. *ACM Trans Graph*. 2014;33(6):223–1.
 23. Deering MF. A photon accurate model of the human eye. *ACM Trans Graph*. 2005;24(3):649–658.
 24. Kravchenko B. Balancing fidelity and performance in iridal light transport simulations aimed at interactive applications. Master's thesis, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2016.
 25. Natural Phenomena Simulation Group (NPSG). *Human iris data*. Ontario, Canada: School of Computer Science, University of Waterloo; 2016. Available from: <http://www.npsg.uwaterloo.ca/data.php>.
 26. Natural Phenomena Simulation Group (NPSG). *Run ILIT online*. Ontario, Canada: School of Computer Science, University of Waterloo; 2013. Available from: <http://www.npsg.uwaterloo.ca/models/ilit.php>.
 27. Baranoski GVG, Rokne JG, Xu G. Virtual spectrophotometric measurements for biologically and physically-based rendering. *Visual Comput*. 2001;17(8):506–518.
 28. Natural Phenomena Simulation Group (NPSG). *Run ILIT interactive*. Ontario, Canada: School of Computer Science, University of Waterloo; 2016. <http://www.npsg.uwaterloo.ca/models/ilitInteractive.php>.
 29. Hopfield JJ. Olfaction and color vision: More in simpler.. In: Ong NP, Bhatt RN, editors. *More is different: Fifty years of condensed matter physics*. Princeton, NJ, USA: Princeton University Press, 2001. p. 269–284.
 30. Stone MC. *A field guide to digital color*. Natick, MA, USA: AK Peters; 2003.
 31. Hunt RWG. *Measuring colour*. 2nd edn. Chichester, England: Ellis Horwood Limited; 1991.
 32. Stokes M, Fairchild MD, Berns RS. Precision requirements for digital color reproduction. *ACM Trans Graph*. 1992;11(4):406–422.
 33. Brainard DH. Color appearance and color difference specification. *Sci Color*. 2003;2: 191–216.
 34. NVIDIA. *CUDA C best practices guide*. Santa Clara, CA, USA: NVIDIA Corporation; 2015. Technical report.
 35. NVIDIA. *Cuda C programming guide*. Santa Clara, CA, USA: NVIDIA Corporation; 2015. Technical report.
 36. Mahy M, Van Eycken L, Oosterlinck A. Evaluation of uniform color spaces developed after the adoption of CIELAB and CIELUV. *Color Res Appl*. 1994;19(2):105–121.
 37. Ellis SR, Mania K, Adelstein BD, Hill MI. Generalizeability of latency detection in a variety of virtual environments. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*; 2004. p. 2632–2636.
 38. Kravchenko B, Baranoski GVG, Chen TF, Miranda E, Van Leeuwen SR. High-fidelity iridal light transport simulations at interactive rates. Available from: <https://youtu.be/ZdqPW8hKSmc>, Supplementary video; 2017.



Boris Kravchenko completed his Master of Mathematics degree at the University of Waterloo in 2016. He received a Bachelor of Computer Science from the same institution in 2014. His research interests include real-time 3D graphics, material appearance simulation, and general-purpose computation on the GPU. He is a former member and currently a collaborator to the Natural Phenomena Simulation Group at the University of Waterloo.



Gladimir V. G. Baranoski obtained his PhD in computer science at the University of Calgary (Canada) in 1998. Currently, he is a faculty member of the David R. Cheriton School of Computer Science at the University of Waterloo (Canada),

where he has established the Natural Phenomena Simulation Group. His research interests include primarily the predictive simulation of light interactions with natural materials aimed at interdisciplinary applications.



Tenn Francis Chen received his PhD degree in computer science from the University of Waterloo (Canada) in 2016. While there, he conducted research related to material appearance with a focus on human skin and relevant temporal phenomena. He is

currently working on technologies related to virtual reality at Google Inc. He is a former member and currently a collaborator to the Natural Phenomena Simulation Group at the University of Waterloo.



Erik Miranda received the BSE (Bachelor of Software Engineering) degree and the MMath degree in computer science from the University of Waterloo (Canada) in 2010 and 2014, respectively. His primary research interests focus on the

application of computer graphics techniques for acous-

tic propagation. He is a former member and currently a collaborator to the Natural Phenomena Simulation Group at the University of Waterloo. former member and currently a collaborator to the Natural Phenomena Simulation Group at the University of Waterloo.



Spencer R. Van Leeuwen is currently pursuing a Master of Mathematics in computer science at the David R. Cheriton School of Computer Science at the University of Waterloo (Canada). He holds a Bachelor of Mathematics in computer science and combinatorics and optimization with a minor in music from the same institution. He is also a member of the Natural Phenomena Simulation Group at the

University of Waterloo.

How to cite this article: Kravchenko B, Baranoski GVG, Chen TF, Miranda E, Van Leeuwen SR. High-fidelity iridal light transport simulations at interactive rates. *Comput Anim Virtual Worlds*. 2017;28:e1755. <https://doi.org/10.1002/cav.1755>