Editor: Mike Potel

Rapid Dissemination of Light Transport Models on the Web

Gladimir V.G. Baranoski, Thomas Dimson, Tenn F. Chen, Bradley Kimmel, Daniel Yim, and Erik Miranda
University of Waterloo

The pursuit of realism has inspired many computer graphics advances, especially regarding the modeling of material appearance. Progress in this area has been directly connected to the development of light transport models—that is, how different materials absorb and scatter light. Such research, however, isn't restricted to realistic image synthesis. Many models developed for rendering applications have employed approaches developed in other fields. Moreover, a number of these models have been evaluated through methodologies that compare the modeled results with actual measured data, a basic requirement in the physical and life sciences. So, these models are applicable to a range of scientific investigations outside the computer graphics domain. (For some examples, see the sidebar.)

Model developers in these various domains must overcome similar scientific and technical hurdles to increase their simulation algorithms' effectiveness. For instance, to predict the amount of radiation required to initiate photosynthesis in plant tissue or melanogenesis (melanin production) in human tissue, researchers must simulate similar light scattering and absorption processes and rely on the scarce availability of measured data for these processes.[1,2]

We believe that interdisciplinary collaboration supported by openly accessible research resources holds the key for timely advances in these areas. These resources include knowledge, data, and models. Viewed in this context, the online availability of light transport models could not only improve the code used in the simulations (for example, by allowing for the detection and fixing of implementation errors) but also strengthen cross-fertilization among different scientific domains.

However, few light transport models have their source code openly available for download. Moreover, simply making the code available might not be enough. The relative complexity of light transport models, especially for users with limited programming experience, usually prevents their use beyond the research groups that develop and maintain them. Clearly, if a model's computer code is accessible to only its developers, the model becomes a black box with limited applicability.

To address these code availability issues and to make light transport models broadly useful, we developed the Natural Phenomena Simulation Group Distributed (NPSGD) framework. NPSGD is modeling-language independent. We designed and implemented it to make these models easily accessible for online use, thereby facilitating global interdisciplinary collaboration. Its source code and documentation are also available for download (www.npsg.uwaterloo.ca/models.php), along with example implementations of validated models. So, other research groups can install, modify, and use the framework to disseminate their own models.

Framework Overview

NPSGD acts as an interface between model implementation and usage by providing a webpage to configure and view simulations (see Figure 1). On the webpage, users can manipulate simulation parameters, such as experimental conditions (for example, the angle of incidence and spectral range) and material characterization data, and receive customized results. NPSGD provides email feedback of job status and model results. By design, this allows time-intensive simulations to run without explicit interaction.

Currently, three light transport models are available through NPSGD. ABM-U is an algorithmic BSDF (bidirectional surface-scattering distri-

 Published by the IEEE Computer Society

bution function) model for unifacial plant leaves. ABM-B is an algorithmic BSDF model for bifacial plant leaves. BioSpec is a biophysically based spectral model of light interaction with human skin.

When a model run finishes, NPSGD attaches the results to an email as a PDF document. For user convenience, this document includes the values assigned to the input parameters. Additional attachments include the raw spectral output data, provided as Excel-readable spreadsheets, and the corresponding output graphs (see Figure 2), provided as PNG image files. Whereas the spectral output data have the same resolution (5 nm) for all models, the simulation's wavelength range varies from 400 nm to 2,500 nm for ABM-U and ABM-B and from 400 nm to 700 nm for BioSpec. Accordingly, the spreadsheets might have up to 421 and 51 rows of data for the leaf and skin models, respectively. For ABM-B and ABM-U, the document provides three radiometric quantities (reflectance, transmittance, and absorptance) as separate columns of data. For BioSpec, it provides only a column with reflectance data.

We took special care to make NPSGD easy to use for both model developers and users. To integrate existing models into NPSGD, researchers need only write a small wrapper around their implementations. Each wrapper is a class, written in Python, specifying the model's name, a list of typed input parameters, and a "run" method that executes the simulation. A wrapper can also contain methods for image generation and LaTeX markup. NPSGD uses the declarations and methods to serve the interactive webpage, produce emails, and create worker processes for model execution. After users submit input data through the website, NPSGD automatically maps the data to member variables in an instantiation of the wrapper class. Model implementations can be in any language.

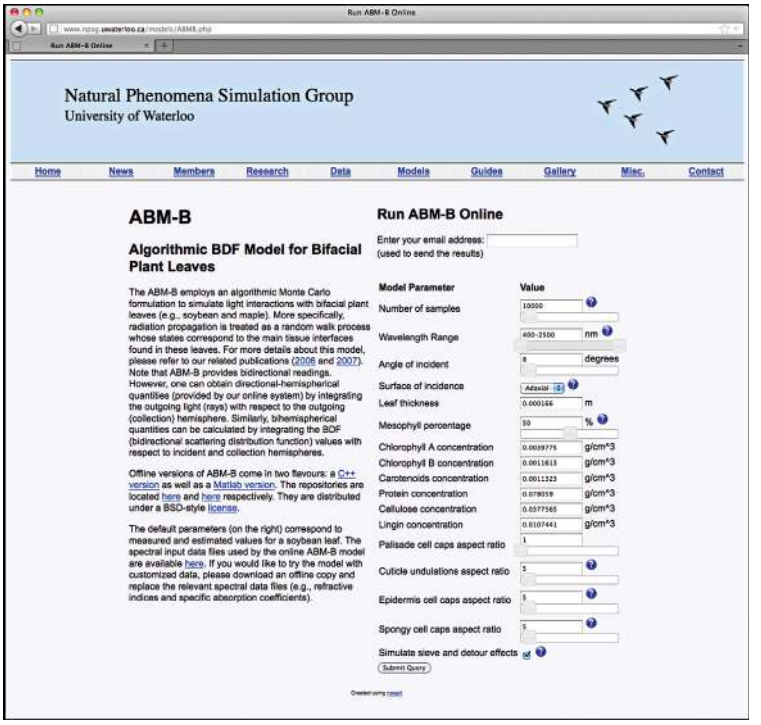

Figure 1. The Web interface for one of the models available through the Natural Phenomena Simulation Group Distributed (NPSGD) framework. Through the webpage, users can configure and execute light transport simulations.

The NPSGD source code includes an example that executes a compiled C++ binary and an example that manipulates parameters using Matlab. NPSGD also includes three categories of templates to create the webpages' HTML, define the email's content, and format LaTeX-based attachments. Using a high-level syntax, researchers can customize

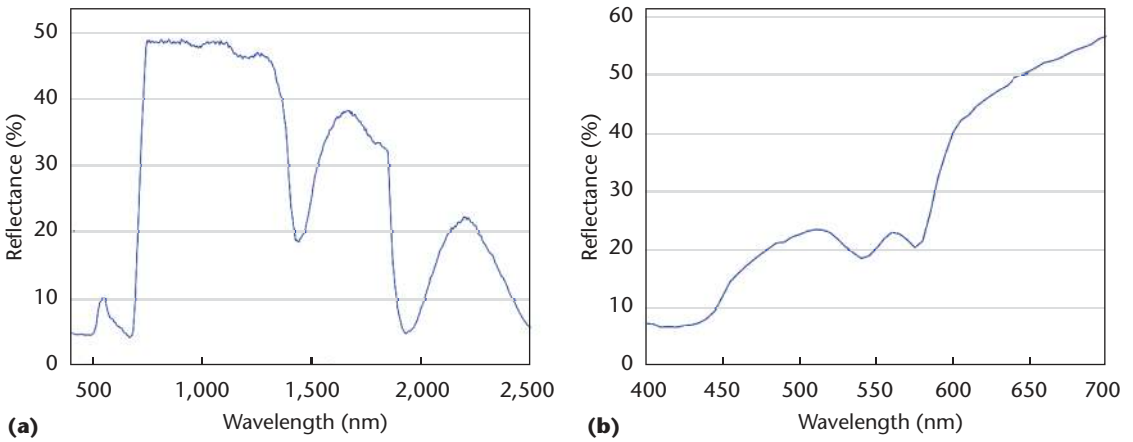

Figure 2. NPSGD output graphs showing data generated by (a) ABM-B, an algorithmic BSDF (bidirectional surface-scattering distribution function) model for bifacial plant leaves, and (b) BioSpec, a biophysically based spectral model of light interaction with human skin. The simulations involved 100,000 samples.

Using Light Transport Models

Light transport models can be used in the forward or inverse mode. In the forward mode, the model usually generates spectral radiometric data (for example, reflectance and transmittance) that depend on the wavelength and direction of the light impinging on the target material. Scientists can obtain comprehensive spectral radiometric datasets by varying the input parameters describing the material's biochemical and biophysical properties. They can then use such datasets, for example, to help predict changes in the concentration and distribution of important material constituents triggered by environmental stimuli.

In the inverse mode, a model can take as input reflectance and transmittance data measured for the target material. It can then quantify the material's constituents or provide information about how the material absorbs and scatters light.

Two typical uses for light transport models involve simulating the spectral appearance of human skin and plants.

Skin Spectral Attributes

In recent years, researchers have devoted substantial effort to developing noninvasive biomedical procedures based on modeling skin's spectral-appearance attributes. These procedures aim to assist the diagnosis, treatment, and prevention of a variety of medical conditions, including skin cancer associated with the abnormal production of melanin.[1]

For example, some medical therapies, such as phototherapy for treating jaundice or photodynamic therapy for treating malignant skin lesions,[2] require in vivo light dosimetry. Dosimetry involves measuring the amount (dose) of radiant energy required to achieve a specific photobiological effect. Medical professionals perform in vivo light dosimetry by either inserting a probe into the tissue or placing the probe some distance from it. In the latter, noninvasive, case, the probe provides reflectance measurements that serve as input to a model-inversion procedure to estimate the tissue absorption and scattering coefficients. These, in turn, can be used to estimate the depth and distribution of light in the tissue by serving as input to a suitable light transport model in the forward mode.

The cosmetics industry also employs computer simulations of biophysical processes affecting skin appearance to assist development of skin protection products such as lotions and sunscreens.[2] Predictive simulation is becoming more important in this area because legislation regulating cosmetics development is becoming stricter worldwide.

Plant Spectral Attributes

Plants are among the most investigated remote sensing targets owing to their importance in sustaining human and animal life. From a remote sensing perspective, the main connection between environmental changes and the light transferred from a landscape is through variations in the natural materials' appearance, or spectral signature.

Regarding vegetation areas, the data collected by multispectral remote sensing systems, such as those on Landsat satellites, serve as input to inversion procedures on light transport models operating at the canopy level. Combining these procedures with inversion procedures on models operating at the leaf level can make quantifying essential parameters, such as chlorophyll, nutrient, and water content, more generic and effective. Accordingly, one of the focal points of remote sensing research has been the development of predictive light transport models for plant leaves.[3] Additionally, forward-mode use of these models has been instrumental in investigating plant photobiology.[3] Disruptions to leaves' photosynthetic apparatus, normally associated with variations in chlorophyll, nutrient, and water content, significantly affect not only food and biofuel production but also the increasingly scarce freshwater supplies used in agriculture.

So, predictive simulation of plant spectral responses to environmental stimuli is essential for formulating effective procedures for managing and conserving valuable agricultural and forestry resources. Such simulation is also essential for accurately assessing ecological disasters involving land or marine vegetation. From an agricultural viewpoint, one of the most important benefits of such in silico experiments is that researchers can quickly perform them to evaluate the trade-offs of strategies for increasing crops' nutrient content and water-use efficiency, given a set of environmental conditions.[4]

References

1. T. Binzoni et al., "Detection Limits of Multi-spectral Optical Imaging under the Skin Surface," *Physics in Medicine and Biology*, vol. 53, no. 3, 2008, pp. 617–636.
2. V. Tuchin, *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*, 2nd ed., Int'l Soc. Optical Eng., 2007.
3. S. Jacquemoud and S. Ustin, "Modeling Leaf Optical Properties," *Photobiological Sciences Online*, 2008; www.photobiology.info/Jacq_Ustin.html.
4. P. Minorsky, "Achieving the In Silico Plant: Systems Biology and the Future of Plant Biological Research," *Plant Physiology*, vol. 132, no. 2, 2003, pp. 404–409.

these templates to change the webpages' appearance and enrich emails and attached PDF documents with specific information about their affiliated university or research group.

NPSGD removes many of the traditional stumbling blocks preventing model distribution: lack of usability, unreliability, and unscalability. Accordingly, NPSGD has three layers (see Figure 3):

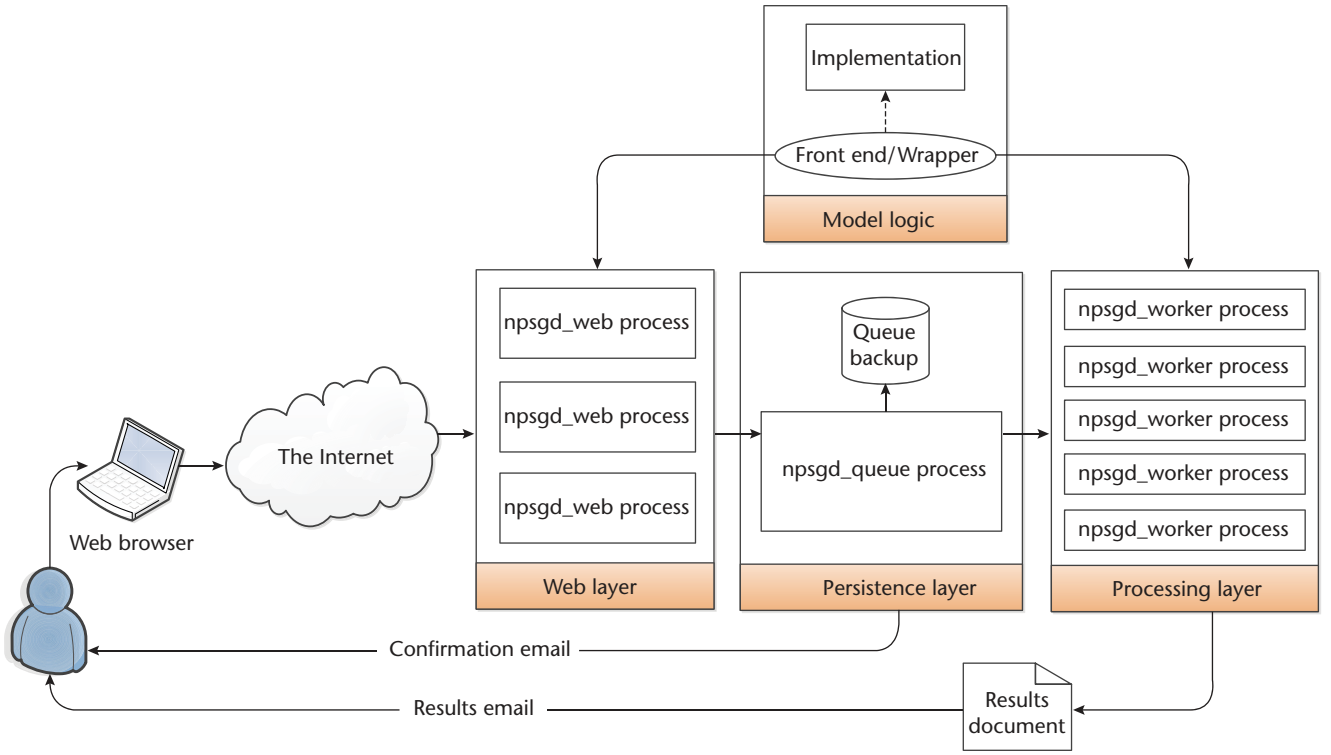

Figure 3. The information flow during a simulation request. The *NPSGD Web* process serves a model webpage by reading the corresponding description provided by the model-handling unit. After the user submits input data via a browser, NPSGD sends a confirmation email and pushes the submission onto the queue provided by the *NPSGD queue* process. When available, an *NPSGD worker* process acquires a request from the queue and runs the simulation by providing the simulation parameters to the underlying model implementation. Upon completion, NPSGD sends the results to the user via email.

- The Web layer consists of instances of *NPSGD Web*, an HTTP server that serves the model webpages to browsers.
- The persistence layer consists of *NPSGD queue*, a disk-backed queue that ensures the results' delivery even after machine failures.
- The processing layer consists of many instances of *NPSGD worker*, processes that compute results using model implementations.

These layers act as a distributed system communicating minimally over HTTP. NPSGD Web and worker instances can be partitioned across multiple servers to facilitate higher loads and reduce the risk of system failure. The system is communicative and fault tolerant. If the queue process or worker processes are offline, the website displays an error message. Previously submitted requests persist on the queue's disk and are reattempted once the process recovers. During execution, the worker processes use a heartbeat signal to communicate to other system members. If a worker process dies during execution, the heartbeat quietly times out, and the queue retries the request on a different worker. If the request fails repeatedly, an email informs the requester.

We programmed NPSGD in Python using the Tornado Web server (www.tornadoweb.org); it runs on Linux and other Unix-like OSs. To cope with long-running processes, the system can employ an asynchronous programming style using a queue that defers computation until worker processes become available. NPSGD has modest hardware requirements. Currently, NPSGD Web and queue processes run on a 3-GHz single-core Pentium 4 CPU with 512 Mbytes of RAM. Simulations run across a dual six-core 2.66-GHz Intel Xeon machine with 24 Gbytes of RAM. Considering the broadest spectral ranges and the maximum number of samples allowed by NPSGD (see Figure 2), these simulations take on average 80 seconds for ABM-U and ABM-B and 30 seconds for BioSpec.

Framework Utilization

The NPSG research program focuses on developing predictive light transport models for realistic image synthesis applications involving natural materials—for example, plant leaves and human skin (see Figure 4). However, because this research program is interdisciplinary, its outcomes have also been published and referenced outside the graphics domain. For example, ABM-U and ABM-B have been

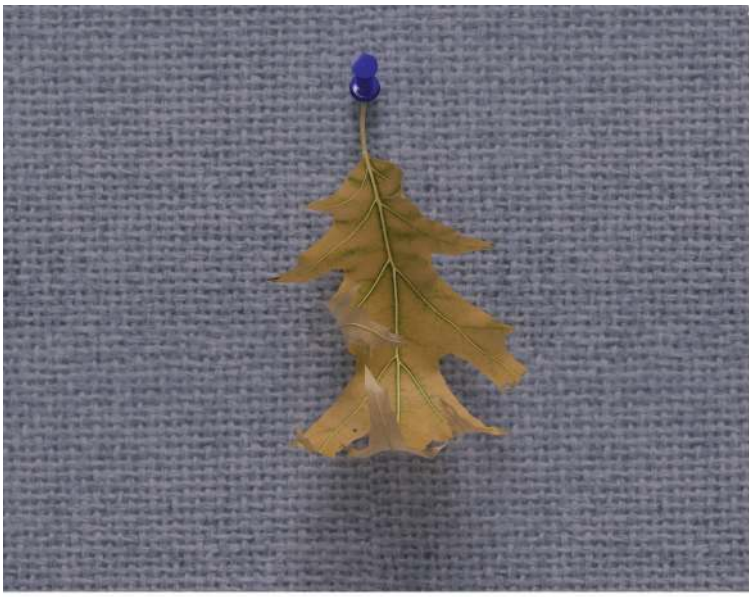

(a)

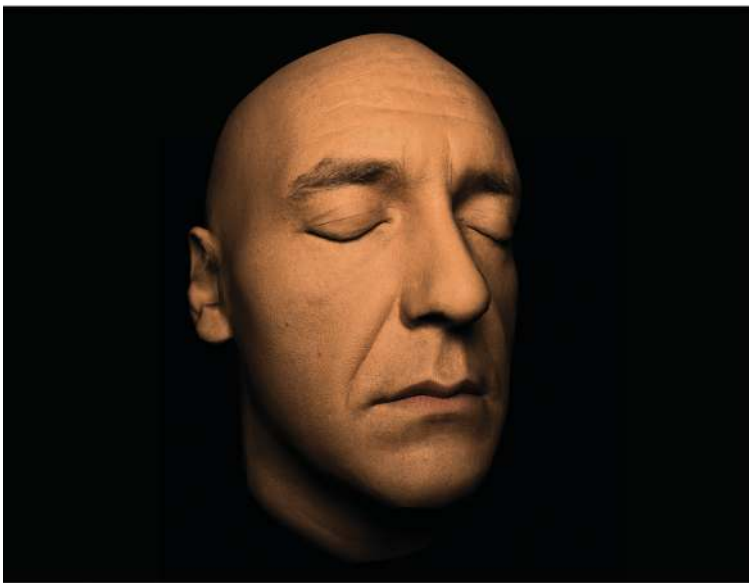

(b)

Figure 4. Two uses of NPSG light transport models. (a) A dried leaf rendered using spectral-appearance attributes provided by ABM-B. (Leaf mesh courtesy of Sung Ming Hong.) (b) A human head rendered using spectral-appearance attributes provided by BioSpec. (Polygonal mesh and textures courtesy of XYZ RGB.)

examined not only in books on modeling material appearance for computer graphics applications[3] but also in state-of-the-art surveys on modeling leaf optical properties for agriculture, remote sensing, and photobiology applications.[2,4]

Accordingly, researchers from these fields often ask to use our models. These requests motivated us to make ABM-U and ABM-B the first models available online through NPSGD. Offline versions of these models and the supporting simulation data (for example, refractive indices and specific absorption coefficients) are also available at the NPSGD website so that researchers can download the code for these models and perform the simulations on their own computers. Since we made these models available, we've received positive feedback from researchers working on modeling and detecting plant spectral signatures, including Stephane Jacquemoud (Institut de Physique du Globe de Paris), a leading scientist in this field.

Similarly, BioSpec has attracted the attention of researchers in such fields as medical imaging[5] and biomedical optics.[6] So, we chose BioSpec as the next model to be available online. Since we made it available, researchers from these fields have used it to obtain modeled spectral data for skin specimens. For example, Jacob Scharcanski's medical-image-processing group at Universidade Federal do Rio Grande do Sul have employed BioSpec-modeled spectral data to classify pigmented skin lesions using standard color-camera images. Because these lesions can be benign or malignant, their reliable classification can be instrumental in the early diagnosis of life-threatening diseases such as melanoma.

One of the next models we'll make available through NPSGD will be CLBlood, a novel cell-based model of light interaction with human blood.[7] With CLBlood, researchers can predictively control the rendering of blood samples with distinct appearances (see Figure 5). For this, CLBlood requires only a few biophysically meaningful parameters, such as the volume fraction occupied by the blood cells (called the *hematocrit*, or simply HCT) and the type of pigment in these cells. Like ABM-U, ABM-B, and BioSpec, CLBlood employs a first-principles approach. However, it probabilistically positions each cell in the material volume to enable the simulation of that cell's interaction with light. After the simulation, CLBlood discards the cell, and the process repeats for the next cell. CLBlood provides spectral and scattering responses for samples over a range of physiological and rheological conditions, making it the most comprehensive model of its kind.

Although the models currently available through NPSGD have been previously published, the online and offline versions of their source code are "fresh." That is, different programmers reimplemented the models. This approach let us filter out possible bugs and improve model running performance by using more efficient software and hardware features.

We plan to make other models of light interaction with natural materials (for example, sand and the human iris) available through NPSGD. We also plan to extend its applicability to different

areas of biomedical and remote sensing research. In addition, we plan to expand its output scope by including spectral-appearance data associated with the spatial distribution of light propagated by different target materials.

Despite the benefits of making data and models publicly available, such an approach's technical hurdles often discourage researchers from following it. NPSGD lets researchers make their light transport models openly accessible in a user-friendly way. So, it provides a sound compromise between code availability and due diligence in the dissemination of research results. More important, as a proof of concept, it demonstrates that the reproducibility of research results through complete model transparency (code and implementation) is feasible. Such reproducibility can lead to fruitful collaborations between model developers and users, regardless of their field of expertise.

Figure 5. Blood samples on microscope slides rendered using spectral-appearance attributes provided by the CLBlood model. From front to back, the samples contain the pigments oxyhemoglobin, deoxyhemoglobin, methemoglobin, and sulfhemoglobin. The last two pigments occur in large amounts in human blood only under abnormal conditions.

References

1. V. Tuchin, *Tissue Optics: Light Scattering Methods and Instruments for Medical Diagnosis*, 2nd ed., Int'l Soc. Optical Eng., 2007.
2. S. Jacquemoud and S. Ustin, "Modeling Leaf Optical Properties," *Photobiological Sciences Online*, 2008; www.photobiology.info/Jacq_Ustin.html.
3. J. Dorsey, H. Rushmeier, and F. Sillion, *Digital Modeling of Material Appearance*, Morgan Kaufmann, 2008.
4. W. Dorigo et al., "A Review of Reflective Remote Sensing and Data Assimilation Techniques for Enhanced Agroecosystem Modeling," *Int'l J. Applied Earth Observation and Geoinformation*, vol. 9, no. 2, 2007, pp. 165–193.
5. T. Binzoni et al., "Detection Limits of Multi-spectral Optical Imaging under the Skin Surface," *Physics in Medicine and Biology*, vol. 53, no. 3, 2008, pp. 617–636.
6. D. Yudovsky and L. Pilon, "Rapid and Accurate Estimation of Blood Saturation, Melanin Content, and Epidermis Thickness from Spectral Diffuse Reflectance," *Applied Optics*, vol. 49, no. 10, 2010, pp. 1707–1719.
7. D. Yim et al., "A Cell-Based Light Interaction Model for Human Blood," *Computer Graphics Forum*, vol. 31, no. 2, 2012, pp. 845–854.

Gladimir V.G. Baranoski *is an associate professor of computer science at the University of Waterloo. Contact him at gvgbaran@uwaterloo.ca.*

Thomas Dimson *was an undergraduate research student in the University of Waterloo's Natural Phenomena Simulation Group during the development of this project. He currently works in the software industry. Contact him at tdimson@gmail.com.*

Tenn F. Chen *is a PhD student in the University of Waterloo's Natural Phenomena Simulation Group. Contact him at t4chen@cs.uwaterloo.ca.*

Bradley Kimmel *is a PhD student in the University of Waterloo's Natural Phenomena Simulation Group. Contact him at bwkimmel@uwaterloo.ca.*

Daniel Yim *is a master's student in the University of Waterloo's Natural Phenomena Simulation Group. Contact him at dyim@uwaterloo.ca.*

Erik Miranda *is a master's student in the University of Waterloo's Natural Phenomena Simulation Group. Contact him at emiranda@uwaterloo.ca.*

Contact department editor Mike Potel at potel@wildcrest.com.