Toward Evaluating the Usefulness of Global Illumination for Novices in Lighting Design Tasks – Supplemental Material

Ondřej Karlík, Martin Růžička, Václav Gassenbauer, Fabio Pellacini, and Jaroslav Křivánek

1 Interactive Global Illumination Solution

Our interactive global illumination solution is based on the Direct-to-Indirect Transfer (DTIT) algorithm [1]. In this section, we describe the differences of our solution from the original algorithm. Please refer to the original paper for more details.

Hašan et al. encode the light transfer from gather samples to view samples using three matrices (the multi-bounce matrix, and the final gather matrices for the diffuse and glossy components). Instead of this separation, our solution uses the 'one-pass formulation' of the DTIT algorithm as described in Section 3.3 of the original paper. That is to say, we encode the light transfer using a single matrix $\mathbf{T}$:

$$\mathbf{v} = \mathbf{T}\mathbf{g} = (\mathbf{T}\mathbf{W}^T) \cdot (\mathbf{W}\mathbf{g}) = \mathbf{T}^w \mathbf{g}^w, \qquad (1)$$

where $\mathbf{v}$ is the vector of indirect illumination for individual image pixels, $\mathbf{g}$ is the vector of diffuse direct illumination on gather samples distributed on all scene surfaces, $\mathbf{T}$ is the transfer matrix, and $\mathbf{W}$ is the Haar wavelet basis matrix (i.e. $\mathbf{T}^w$ denotes the transfer matrix where all rows are projected onto the Haar wavelet basis, and $\mathbf{g}^w$ is the Haar wavelet basis projection of $\mathbf{g}$).

Please note that while in the original algorithm the $\mathbf{v}$ vector represents illumination at view samples located in the scene, in our approach it contains the actual pixel values. This formulation has a big advantage because it allows us to encode *antialiasing* (in addition to light transfer) in the $\mathbf{T}$ matrix. Being able to bake antialiasing into the transfer matrix was the primary reason for dropping the diffuse and glossy component separation in the final-gather matrix, and, as a consequence, also the idea of separating $\mathbf{T}$ into the multi-bounce and final-gather matrices.

An important advantage of *not* separating $\mathbf{T}$ into the multi-bounce and final-gather matrices is that we can now encode more light paths. While in Hašan et al.'s formulation the first bounce from the light source as well as the second bounce from the camera is limited to Lambertian reflection, in our solution only the first bounce from the light source has this limitation. This allows us to encode some indirect glossy and specular paths that would be ignored by the original formulation.

To precompute the $\mathbf{T}$ matrix, we path trace the image using 50 000 paths per pixel. At each path vertex we find the nearest gather sample, and add the accumulated path throughput to the matrix element corresponding to the current pixel and the nearest gather sample.

To obtain high-quality, artifact-free indirect illumination, we use 1 million gather samples and between 250 and 330 wavelet coefficients for each transfer matrix row (depending on how much GPU memory is left after loading the scene data).

- *O. Karlík, M. Růžička, and J. Křivánek are with Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic. E-mails: karlik@cgg.mff.cuni.cz, martinruzickatm@seznam.cz, jaroslav.krivanek@mff.cuni.cz.*
- *Václav Gassenbauer is with IRISA / INRIA, Rennes. E-mail: vgassenb@irisa.fr.*
- *F. Pellacini is with Dartmouth College and Sapienza University of Rome. E-mail: fabio@cs.dartmouth.edu*

2 Subject eligibility for testing

Subjects participating in the study were checked to meet the following requirements:

- basic familiarity with using a PC,
- normal or corrected-to-normal vision, and
- no substantial previous knowledge of rendering algorithms nor experience with 3D content creation (which includes inability to answer the questions "What is global illumination?" and "What are fill lights in computer graphics?").

3 Test procedure

At the beginning of the test procedure each subject is given a high-level description of the study and test procedure, asked to turn off her/his mobile phone to avoid any distractions, instructed about the payment, asked to adjust the workplace according to her/his ergonomic preferences, and offered snacks and non-alcoholic beverages (which are available during the entire course of the experiment).

3.1 Training procedure

The instructor goes through the training with the test subject. The training consists of the following 11 trials:

1) Learning the basics of the relighting application and learning viewport manipulation. The goal is to match the camera view.
2) Learning the basics of light manipulation and manipulating a light to match a screenshot of the interactive viewport.
3) Same as previous, only with different scene.
4) Matching the illumination of a single (key) light (similarly to Experiment 1 in the actual study).
5) Matching the illumination of a single key light with fixed position through light attribute manipulation (intensity, color, and directionality).
6) Introduction of fill lights – fill lights are described, and the subject is asked to switch a light between key and fill mode, observing the results.
7) Matching the illumination of a single fill light (similarly to Experiment 2 in the actual study).
8) Introduction of global illumination – GI is described, and the user is asked to turn it on and off, observing the results.
9) Matching the illumination of a single key light facing a wall with GI enabled (bounce light) – similarly to Experiment 3 in the actual study.
10) Introducing multiple lights – the user is asked to create additional lights and manipulate their position/orientation/scale to match a screenshot of the interactive viewport.
11) Free training – the user is given a scene with all features enabled and is encouraged to experiment on her/his own, and ask questions about the interface or study, should she/he have any.

4 Missing Data

We have recorded three incidents during the testing causing partial loss of measured data. All results and analysis presented in the paper and supplemental materials correctly account for these losses.

- The testing application crashed due to overheating when Subject 4 was matching lighting with global illumination in Cartoon scene in Experiment 1; we have discarded his result in this scene along with the paired direct lighting matching.
- Subject 10 incorrectly ranked lighting features in Experiment 5 by using indirect lighting twice and not using fill lighting. We have discarded this ranking.
- Subject 25 was mistakenly given the opportunity to turn GI on or off in Experiment 4, as discovered by inspecting the recording of the testing. We have discarded his entire results in this experiment.

5 Statistical Analysis

5.1 Paired differences

Since all subjects in our study do all pairs of tasks we want to compare (e.g. matching with GI on and off), we use the paired differences to reduce the overall variance. This means that instead of computing statistics for each task independently, and then comparing them, we work with the inter-subject paired differences. Suppose, that for i-th of N subjects we have measured values a_i and b_i (e.g matching time without GI and with GI). Then, the difference d_i for this subject is

$$d_i = b_i - a_i.$$

The sample *mean difference* is the basic statistic we use. It is the estimate of the true difference in the entire population. It is computed as the average of all subjects' differences:

$$\bar{d} = \frac{1}{N}\sum_{i=1}^{N} d_i = \frac{1}{N}\sum_{i=1}^{N}(b_i - a_i).$$

The associated sample standard deviation is

$$s = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}\left(\bar{d} - d_i\right)^2}.$$

The standard error of the sample mean is

$$SE = \frac{s}{\sqrt{N}}.$$

Using *mean of differences* instead of *difference of means* efficiently reduces variance by eliminating inter-subject variance. Consider a hypothetical case where Subject A finishes matching without GI in 50 seconds and with GI in 100, but Subject B takes 200 and 250 seconds. The paired difference is then the same for both subjects (50 seconds) with zero variance. If we were to compare the means for GI off and GI on, we would get the same difference of means (50 seconds), but both means would have huge variance, making the results inconclusive.

5.2 Paired t-test confidence interval construction

We construct the intervals for data with normal distribution according to Gardner and Altman [2]. An α-confidence interval is constructed as

$$\left[\bar{x} - (t_{1-\alpha/2} \cdot SE), \quad \bar{x} + (t_{1-\alpha/2} \cdot SE)\right],$$

where t is the value of Student's t-distribution with $N-1$ degrees of freedom.

5.3 Cohen's *d*

This effect size indicator, given by Cohen [3], is simply a difference of (unpaired) means $\bar{a}$, $\bar{b}$ divided by the standard deviation s:

$$d = \frac{\bar{a} - \bar{b}}{s}.$$

Since we have two populations, we compute the standard deviation s as the pooled standard deviation of both group standard deviations s_1, s_2:

$$s = \sqrt{\frac{(N-1)\cdot(s_1^2 + s_2^2)}{2N}}$$

5.4 Pearson's *r*

This effect size indicator is identical to the Pearson correlation coefficient, which is defined as

$$r = \frac{\sum_{i=1}^{N}(X_i - \bar{X}) \cdot (Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{N}(X_i - \bar{X})^2} \cdot \sqrt{\sum_{i=1}^{N}(Y_i - \bar{Y})^2}}.$$

We get the vectors X, Y from the two vectors of measured data $\mathbf{a}$, $\mathbf{b}$ as follows: X is simply concatenation of the vectors, and Y_i equals 0 if X_i was taken from the vector $\mathbf{a}$, and 1 otherwise.

REFERENCES

[1] M. Hašan, F. Pellacini, and K. Bala, "Direct-to-indirect transfer for cinematic relighting," *ACM Trans. Graph. (SIGGRAPH 2006)*, vol. 25, no. 3, pp. 1089–1097, 2006.

[2] M. J. Gardner and D. G. Altman, "Confidence intervals rather than p values: estimation rather than hypothesis testing," *British Medical Journal (Clinical Research Edition)*, vol. 292, no. 6522, pp. 746–750, March 1986.

[3] J. Cohen, *Statistical power analysis for the behavioral sciences : Jacob Cohen.*, 2nd ed. Lawrence Erlbaum, Jan. 1988.