An Error Estimation Framework for Photon Density Estimation

Toshiya Hachisuka* Wojciech Jarosz†,* Henrik Wann Jensen*

*UC San Diego †Disney Research Zürich

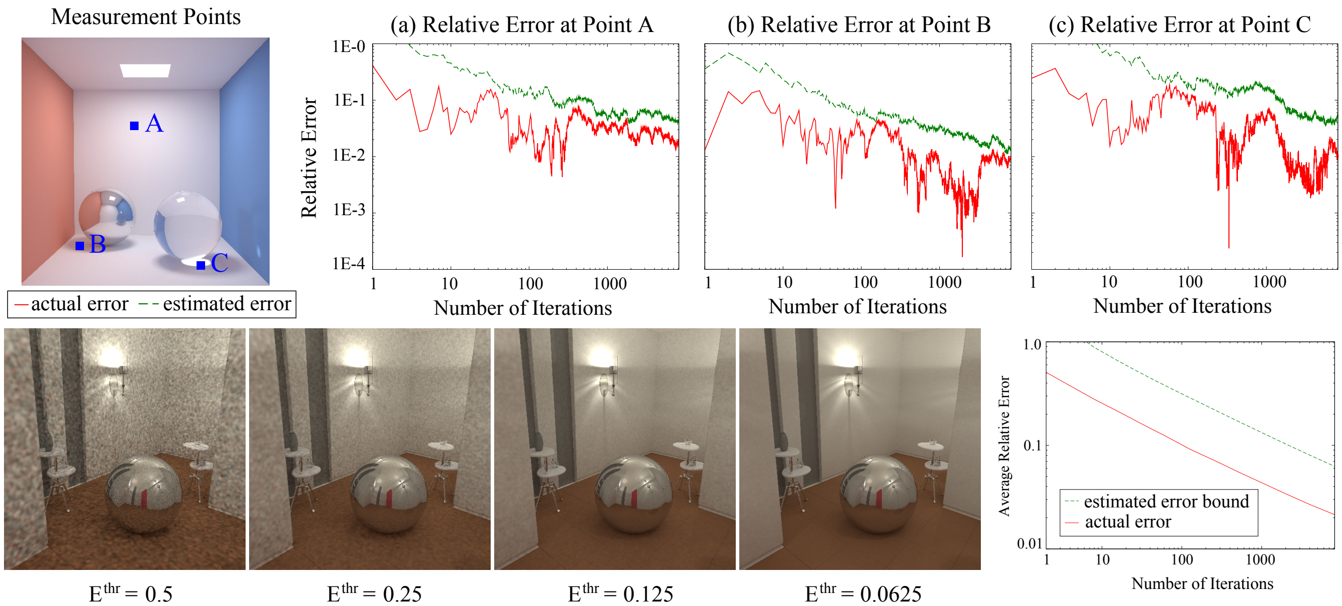

Figure 1: *Top: error bound estimation on a test scene. We show the reference rendering (left), as well as the actual error (red) and estimated error bound (green) at the three points (a,b,c) shown in the reference images. The specified confidence is 90%. Each iteration uses 15K photons. Bottom: rendering with specified error thresholds. The rendering process is terminated once the error estimate reaches the specified threshold. Our conservative error bound predicts the rate of convergence of the true average error automatically (log-log plot).*

Introduction Estimating error is an important task in rendering. For many predictive rendering applications such as simulation of car headlights, lighting design, or architectural design it is import to provide an estimate of the actual error to ensure confidence and accuracy of the results. Even for applications where accuracy is not critical, error estimation is still useful for improving aspects of the rendering algorithm. Examples include terminating the rendering algorithm automatically, adaptive sampling where the parameters of the rendering algorithm are adjusted dynamically to minimize the error, and interpolating sparsely sampled radiance within a given error bound.

We present a general error bound estimation framework for global illumination rendering using photon density estimation. Our method estimates bias due to density estimation of photons, and variance due to photon tracing. Our error bound estimation is robust for any light transport configuration since it is based on progressive photon mapping (PPM). We demonstrate that our estimated error bound captures error under complex light transport. Figure 1 shows that our error bound estimation works well under complex illumination including caustics. As a proof of concept of our error bound estimation, we demonstrate that it can be used to automatically terminate rendering without any subjective trial and error by a user. Our framework is the first general error estimation framework for photon based rendering methods that can handle complex global light transport with arbitrary light paths and materials. Existing work is either based on heuristics that do not capture error, or are limited to specific light paths or materials. We believe that our work is the first step towards answering the important question: "How many photons are needed to render this scene?".

Our Framework Unbiased Monte Carlo ray tracing algorithms are often preferred for predictive rendering since the error bound can be estimated based on variance. However, unbiased methods are not robust in the presence of specular reflections or refractions of caustics from small light sources, which can be often seen in applications of lighting simulation (light bulbs, headlights etc). We therefore use progressive photon mapping (PPM) [Hachisuka et al. 2008], which is a biased Monte Carlo algorithm that is robust under these lighting conditions. Our error estimation framework estimates the following stochastic error bound: $P(-E_i < L_i - L - B_i < E_i) = 1-\beta$, where $E_i = t(i, 1-\frac{\beta}{2})\sqrt{\frac{V_i}{i}}$, $t(i,x)$ is the x percentile of the t-distribution with degree i, L_i is estimated radiance, B_i is estimated bias, V_i is estimated variance, L is the correct radiance, and $1-\beta$ is user-defined confidence of this stochastic bound. Using this stochastic error bound, the user can simply specify their desired confidence as $1-\beta$, and our framework can tell how far the current estimated radiance is from the correct radiance without knowing L. We estimate the bias as $B_i \approx \frac{1}{2}R_i^2 k_2 \Delta L$ [Silverman 1986], where k_2 is a constant derived from the kernel and R_i is the radius for density estimation. We show how to apply this technique to the progressive density estimation used in progressive photon mapping. This equation uses the Laplacian of the unknown, correct radiance ΔL. We estimate this value using the Laplacian of the kernel, which has been used in standard density estimation techniques. Although the original PPM does not support smooth kernels that have a Laplacian, we derive the necessary conditions for incorporating these kernel functions within PPM. Using our kernel-based PPM, we can estimate any order of derivatives including the Laplacian of radiance in a consistent way, which has not been done in existing work and would be useful for analysis of illumination.

Unfortunately, the standard procedure to estimate variance cannot be used in biased methods. We therefore propose estimating variance using bias-corrected radiance $L_i - B_i$, which also removes the dependency between samples in PPM. The same framework can be applied in a straight forward way for photon mapping or grid-based photon-density estimation by simply estimating the Laplacian. The advantage of using PPM is that the estimation of the Laplacian converges to the correct Laplacian in the limit, not just an approximation. That means that the entire framework is consistent except for the approximation of bias.

References

Hachisuka, T., Ogaki, S., and Jensen, H. W. 2008. Progressive photon mapping. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2008) 27*, 5, Article 130.

Silverman, B. 1986. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, New York, NY.