Supplemental Material: Adaptive Progressive Photon Mapping

ANTON S. KAPLANYAN and CARSTEN DACHSBACHER
Karlsruhe Institute of Technology

List of Figures

1. SIMULATION OF 2D RADIANCE ESTIMATION

In order to validate the accuracy and performance of the derived optimal values, the derivative estimator, and the adaptive bandwidth selector, we performed a statistical simulation reproducing the estimation process of Progressive Photon Mapping (PPM). This experiment allows us to measure and analyze the estimation error and study the asymptotic properties and convergence of the estimators. All simulations are performed in 2D, as the optimal parameters and conditions have been derived for on-surface radiance estimation. The estimation is computed on the domain $(-1;1) \times (-1;1)$ on a regular grid Ω of 400×400 estimation points $\mathbf{y}_i$. We draw samples from the "radiance distribution", which is, for the sake of this experiment, a design function as described in the following text. We estimate the radiance at all points $\mathbf{y}_i$ and compute the Mean Integrated Squared Error (MISE). The MISE is computed as the squared difference between the estimated values at $\mathbf{y}_i$ and the exact value of the radiance design function (or its derivatives, respectively). In every experiment we draw 10^{11} samples in order to study the asymptotic behavior of the Mean Squared Error (MSE) for different estimators.

The samples are drawn according to a density function using tabulated CDFs (using 150 million entries). The MSE is computed for all estimation points based on the ground truth values of the actual function $f(\mathbf{y})$. All computation is done in double precision. For the kernel function k we use the Perlin kernel $k(t) = 1 + t^3(-6t^2 + 15t - 10)$ [Perlin 2002]. The simulation is using the recursive PPM estimator; every step N comprises the following operations.

—We draw $J = 2000$ sample pairs $(\mathbf{x}_j, \gamma_j)$ of the estimated function $f(\mathbf{y})$ according to the density function: $\mathbf{x}_j \in (-1;1) \times (-1;1)$ is the position of the sample, and $\gamma_j \in \mathbb{R}_+$ is the contribution of the j-th sample which is drawn from the normal distribution $\mathcal{N}(\mu, \sigma^2)$. Here, μ is the expected value of γ and the variance σ^2 mimics the Monte-Carlo integration of the path integral in the pixel measurement equation. We use $\mu = 10$ and $\sigma^2 = 3$ for our simulations.

—Using these samples, the radiance values $L_N(\mathbf{y}_i)$ are updated for all estimation points $\mathbf{y}_i$. The values are recursively estimated using the equations we use in the paper: $L_N(\mathbf{y}_i) = \frac{N-1}{N} L_{N-1}(\mathbf{y}_i) + \frac{1}{N} \left(\frac{1}{J} \sum_{j=1}^{J} k(\mathbf{y}_i - \mathbf{x}_j) \gamma_j \right)$

—The mean integrated squared error is computed.

We use the following design functions $f(t)$, where $t \in (0;1)$ is the Euclidean length of an input vector $\mathbf{y}$:

—A RADIANCE function $f(t) = \sin \arctan \frac{h}{t}$ which resembles illumination of a point light source onto a plane below it (in spirit of Hachisuka et al. [2010], see Figure 1.1). Here, h is the elevation of the light source above the plane ($h = 0.2$ in our experiments). We take this function as a representative for the class of continuous functions.

—The DOPPLER function $f(t) = 0.5 + \sqrt{t(1-t)} \sin \frac{2.1\pi}{t+0.05}$ representing the class of highly discontinuous functions. It is shown in Figure 1.2.

We use the complex DOPPLER design function in all experiments unless stated otherwise.

For the original PPM estimator, we always use $\alpha_{opt} = 2/3$ for the bandwidth shrinkage (unless we examine different values of α, where these are explicitly given), and the initial bandwidth is determined by a k-NN search. As the value of k is crucial for the comparison of the efficiency, we always determine the globally optimal k by empirically minimizing the MISE after 1000 iterations.

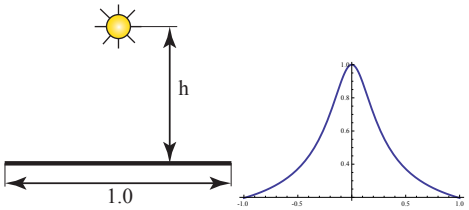

Fig. 1.1: RADIANCE function. Scene with a plane and a point light source in the middle.

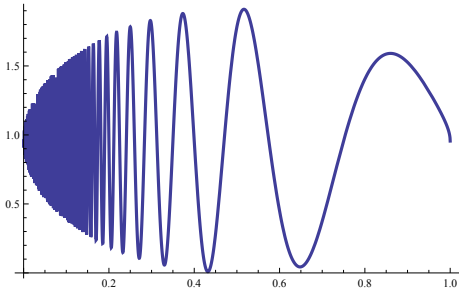

Fig. 1.2: DOPPLER function.

1.1 Different Bandwidth Shrinkage Rates for Asymptotic Estimation

We show the asymptotic behavior of the radiance estimation of the original PPM and its MISE depending on different values of α (and, respectively, the asymptotic rate of bandwidth reduction) for the DOPPLER design function.

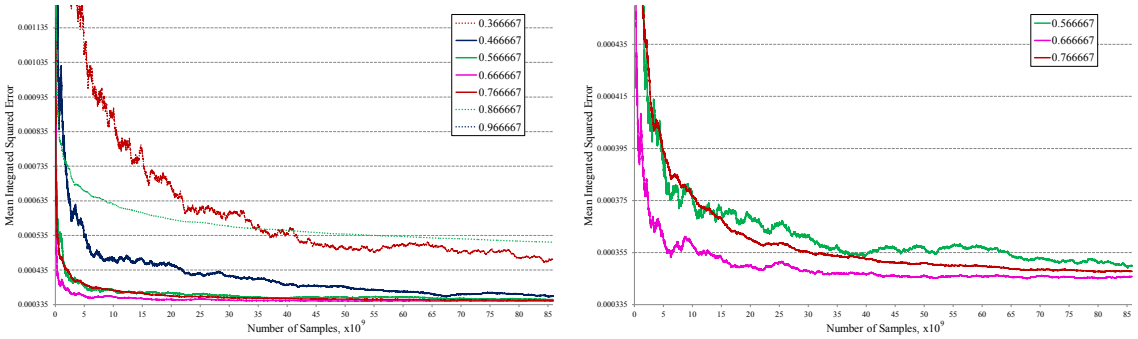

Fig. 1.3: Different α for the recursive estimation. We determine the initial bandwidth using a k-NN search with $k = 7$.

1.2 Different Bandwidth Shrinkage Rates for the Second Derivative Estimation

Similar to Section 1.1, we study the asymptotically optimal value of α' for the estimation of the second derivative of the DOPPLER design function. In addition, the simulation confirms that the estimation of the second derivative diverges if $\alpha' \notin (2/3; 1)$.

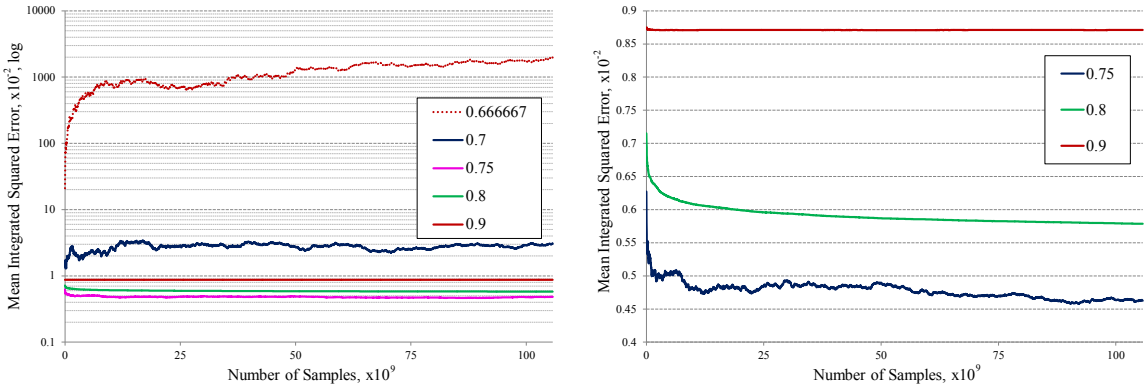

Fig. 1.4: Different α' for the estimation of the second derivative. We determine the initial bandwidth using a k-NN search with $k = 8$.

1.3 Comparison of the Adaptive Estimator to the PPM Estimator without Bandwidth Selection

Next we compare the two recursive estimators discussed in the article: the original PPM estimator with a k-NN bandwidth selection and our new adaptive approach.

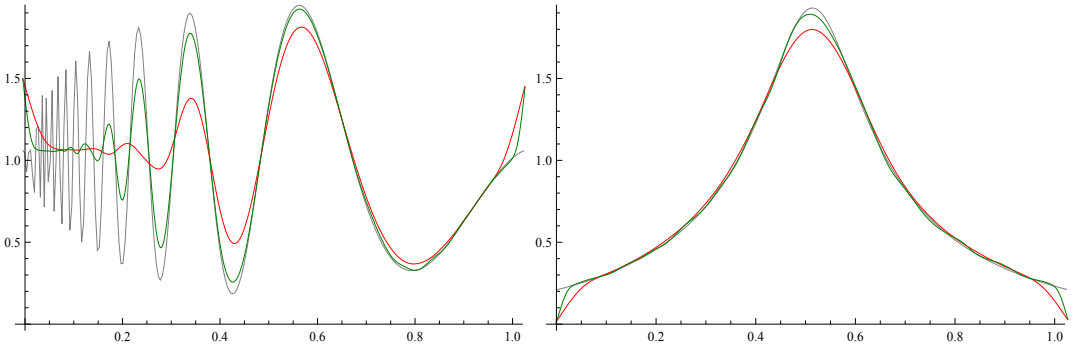

Fig. 1.5: Estimation results after drawing 100.000 samples with DOPPLER (left) and RADIANCE (right) design functions for the recursive estimation with constant bandwidth (red) and plugin-bandwidth (our method, green). Note that these plots show a 1D-slice of the 2D estimation. In PPM we determine the initial bandwidth using a k-NN search with $k = 7$ and $k = 14$ for DOPPLER and RADIANCE design functions, respectively.

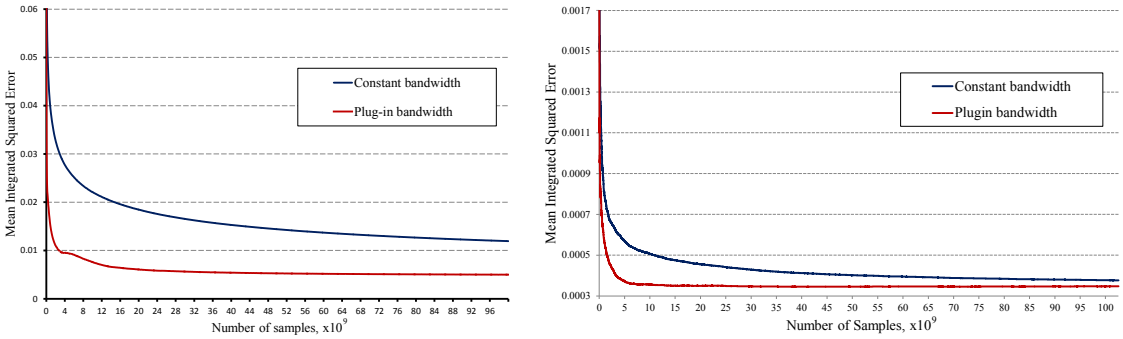

Fig. 1.6: The MISE for the recursive estimation with constant bandwidth and plugin bandwidth (our method) for the DOPPLER (left) and RADIANCE (right) design functions.

1.4 Comparison of the Bandwidth Selection for the Estimation of the Second Derivative

Here, we compare the accuracy of the estimation of the second derivative of the radiance along the x-axis. We use different bandwidth selection strategies: (1) initial bandwidth with k-NN search and (2) the local rule of thumb bandwidth selection using the equation given in the article.

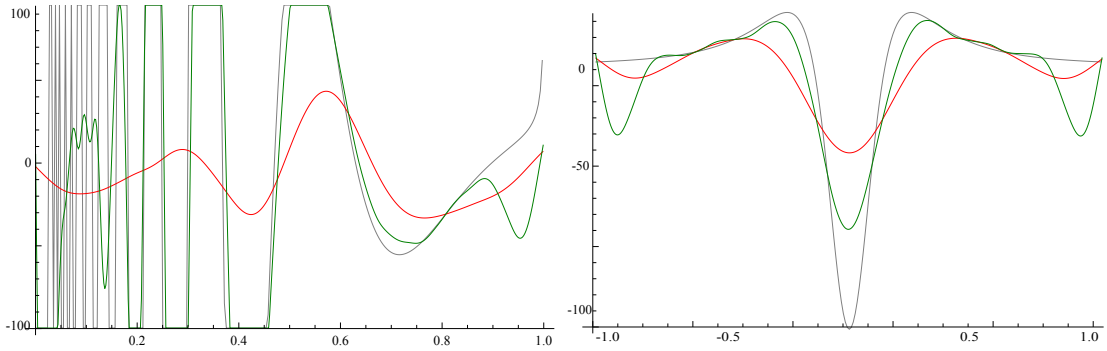

Fig. 1.7: Estimation of the second derivative after drawing 100.000 samples with the DOPPLER (left) and RADIANCE (right) design functions. We compare the recursive estimation with an initial bandwidth selected using a k-NN search (red, $k = 7$ and $k = 14$ for the left and right plot, respectively), and using the rule of thumb (green). Note that this is a slice of the 2D estimation along x-axis through the origin.

2. CONVERGENCE RATES FOR DIFFERENT TEST SCENES

In this section we show different scenes where we compare APPM (our method) to original PPM using plots of the MISE. The MISE is computed to the reference solution obtained from the original PPM with 10^{11} photons where the global bandwidth was smaller than the pixel footprint [Hachisuka and Jensen 2009].

In all examples the photon map size is 256^2. For the original PPM method we use $\alpha = 2/3$ and select the initial bandwidth according to a k-NN search with $k = 10$. We show renderings after having emitted 5, 10, 20, 50, and 100 million photons. The bandwidth range is specified in per mille with respect to the bounding radius of the corresponding scene.

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons (f) top: bandwidth after $100 \cdot 10^6$ photons, bottom: converged result

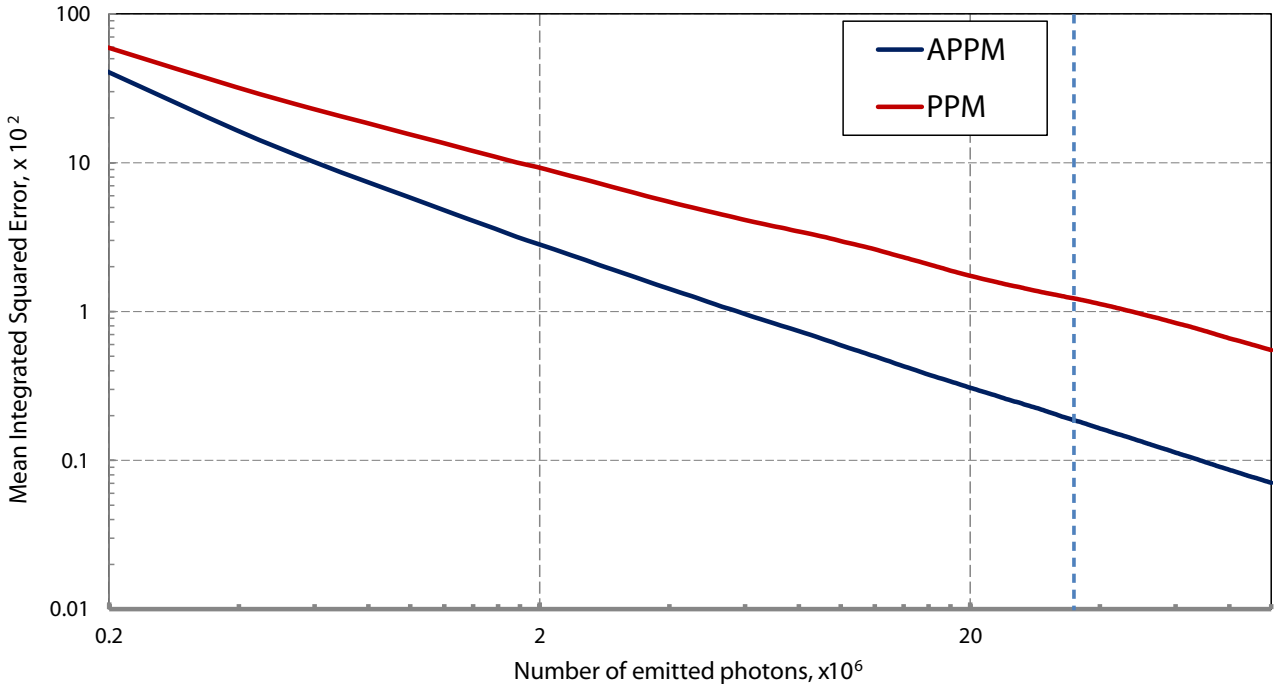

Fig. 2.1: Progressive rendering for the scene ROOM. The dashed vertical blue line roughly shows the stabilization of the bandwidth estimator.

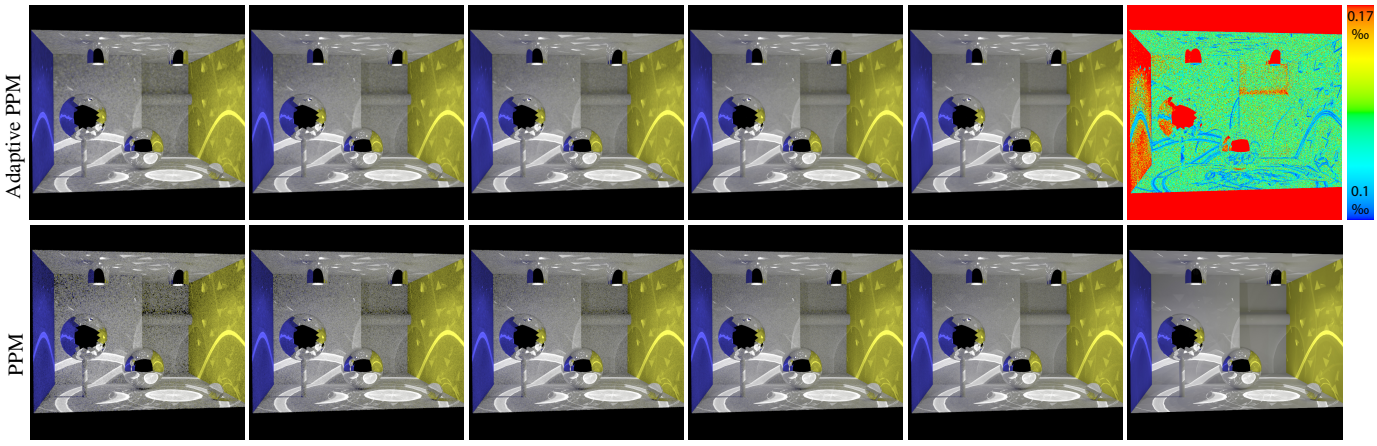

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons (f) top: bandwidth after $100 \cdot 10^6$ photons, bottom: converged result

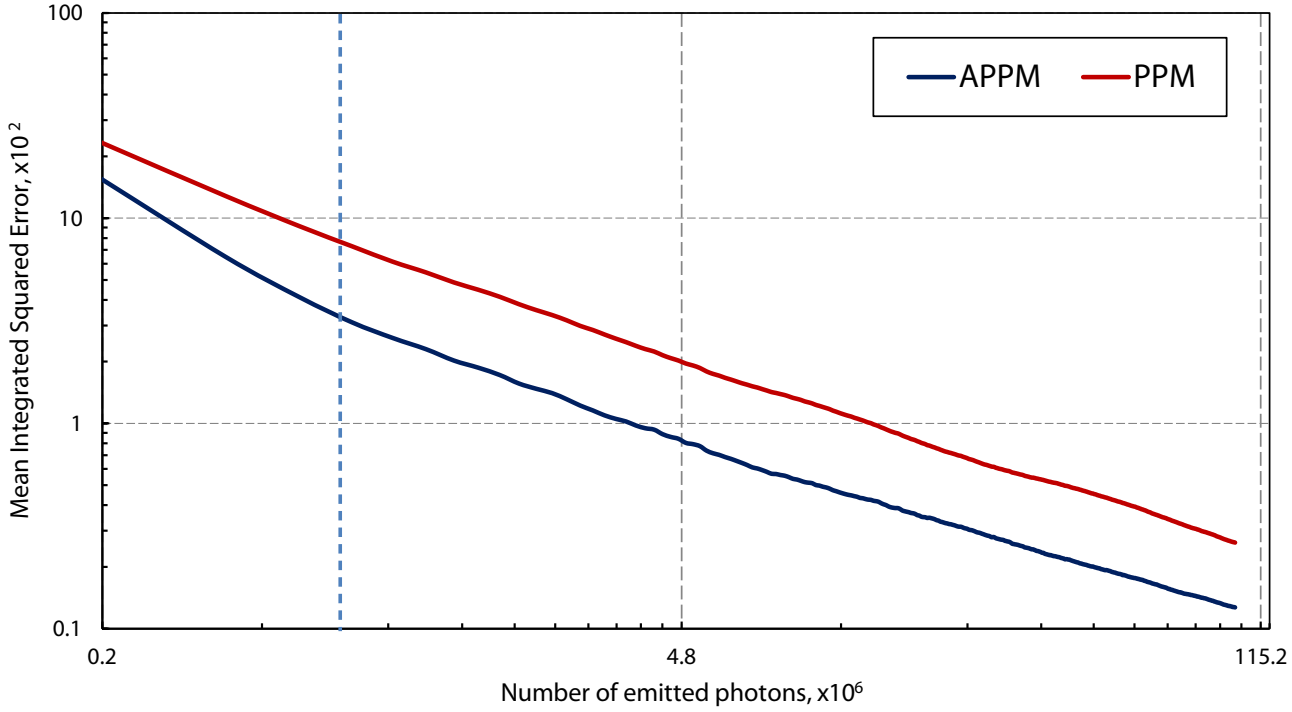

Fig. 2.2: Progressive rendering for the scene BOX. The dashed vertical blue line roughly shows the stabilization of the bandwidth estimator.

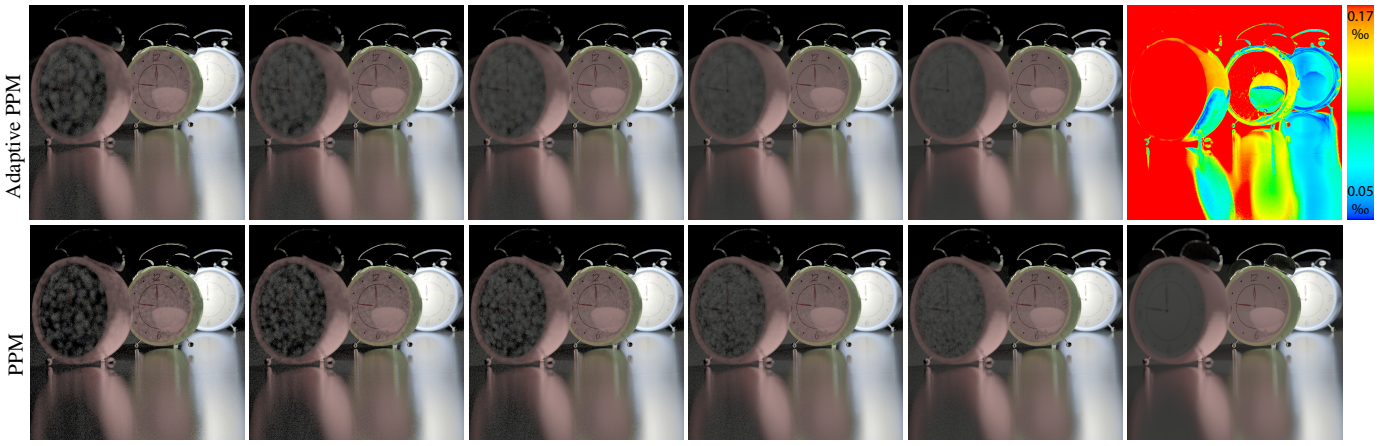

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons (f) top: bandwidth after $100 \cdot 10^6$ photons, bottom: converged result

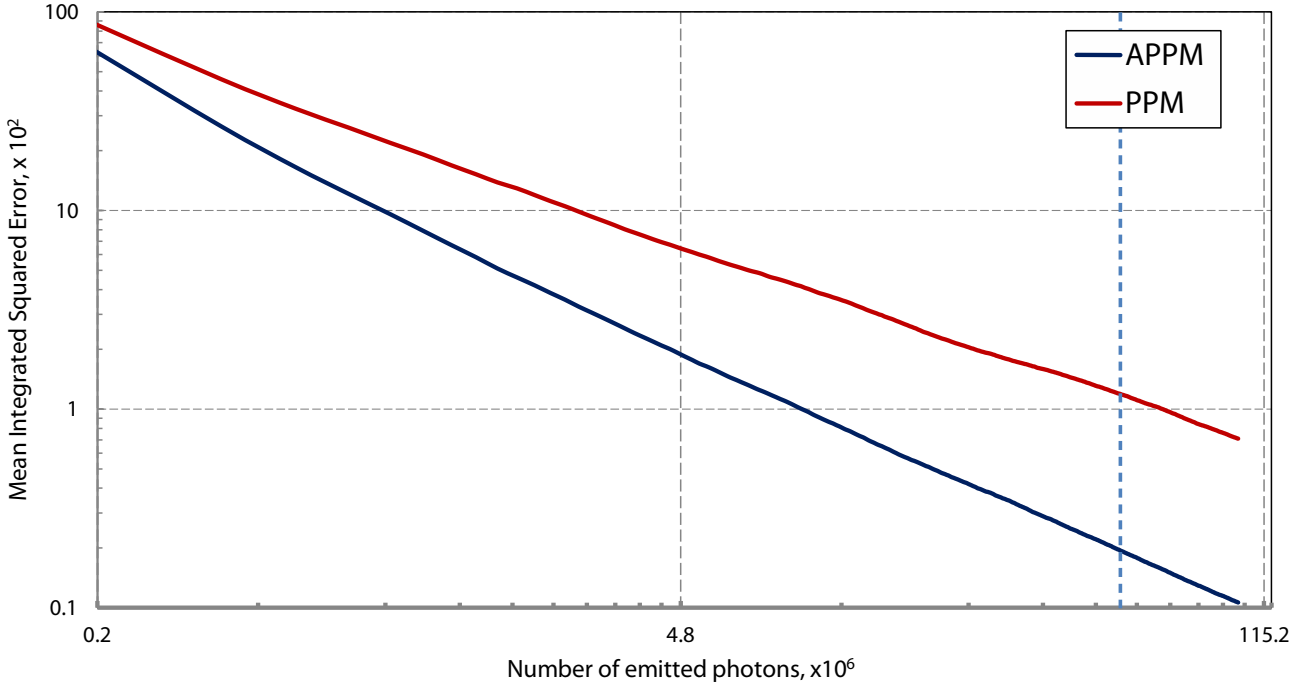

Fig. 2.3: Progressive rendering for the scene CLOCKS. The dashed vertical blue line roughly shows the stabilization of the bandwidth estimator.

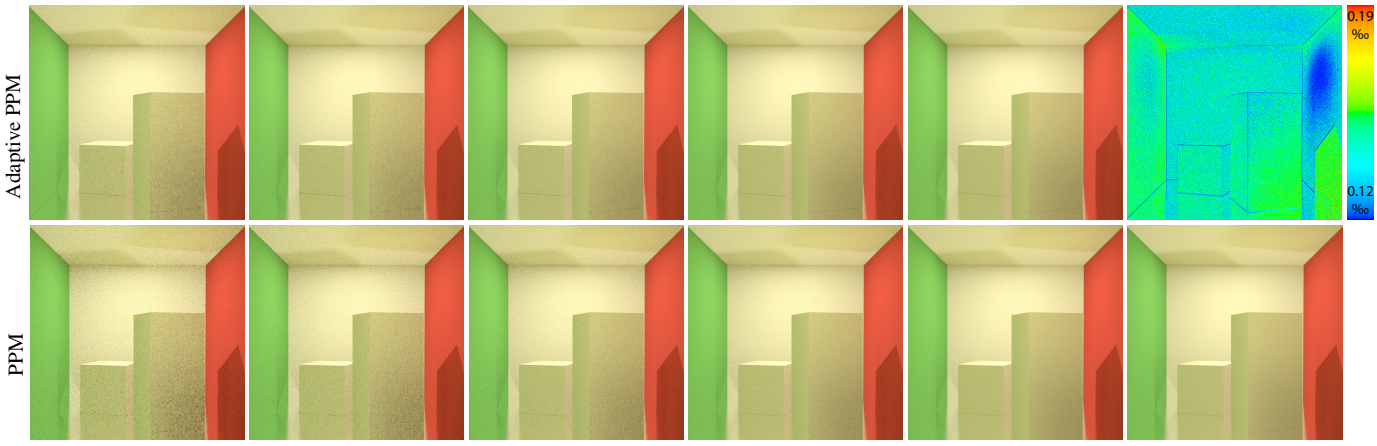

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons (f) top: bandwidth after $100 \cdot 10^6$ photons, bottom: converged result

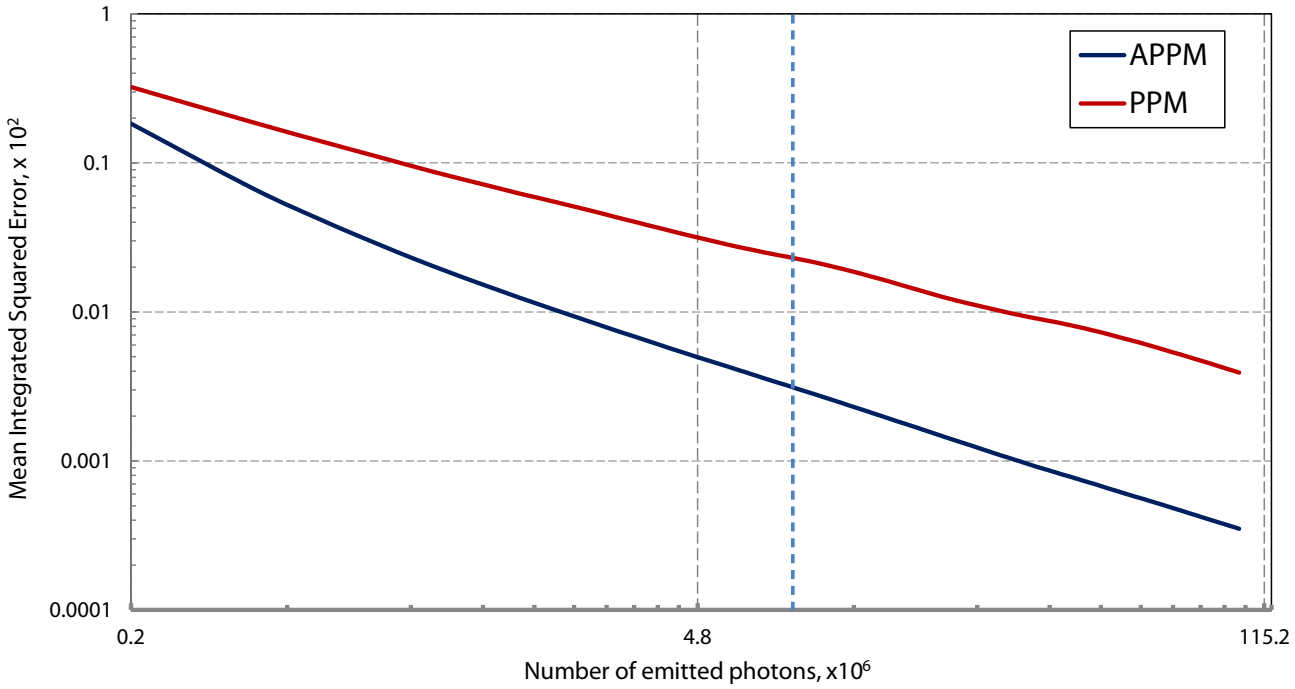

Fig. 2.4: Progressive rendering for the scene CORNELL. The dashed vertical blue line roughly shows the stabilization of the bandwidth estimator.

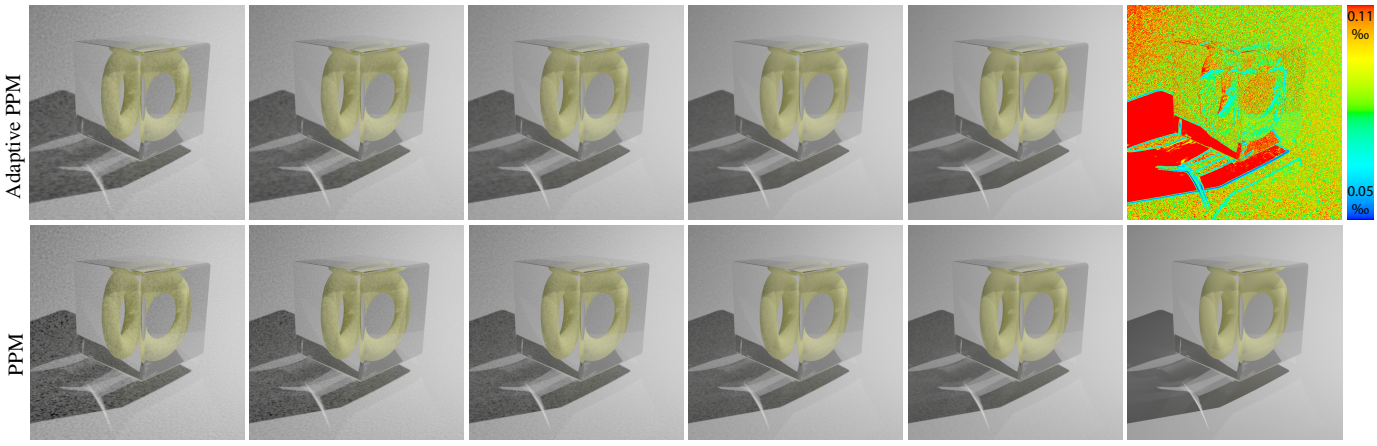

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons (f) top: bandwidth after $100 \cdot 10^6$ photons, bottom: converged result

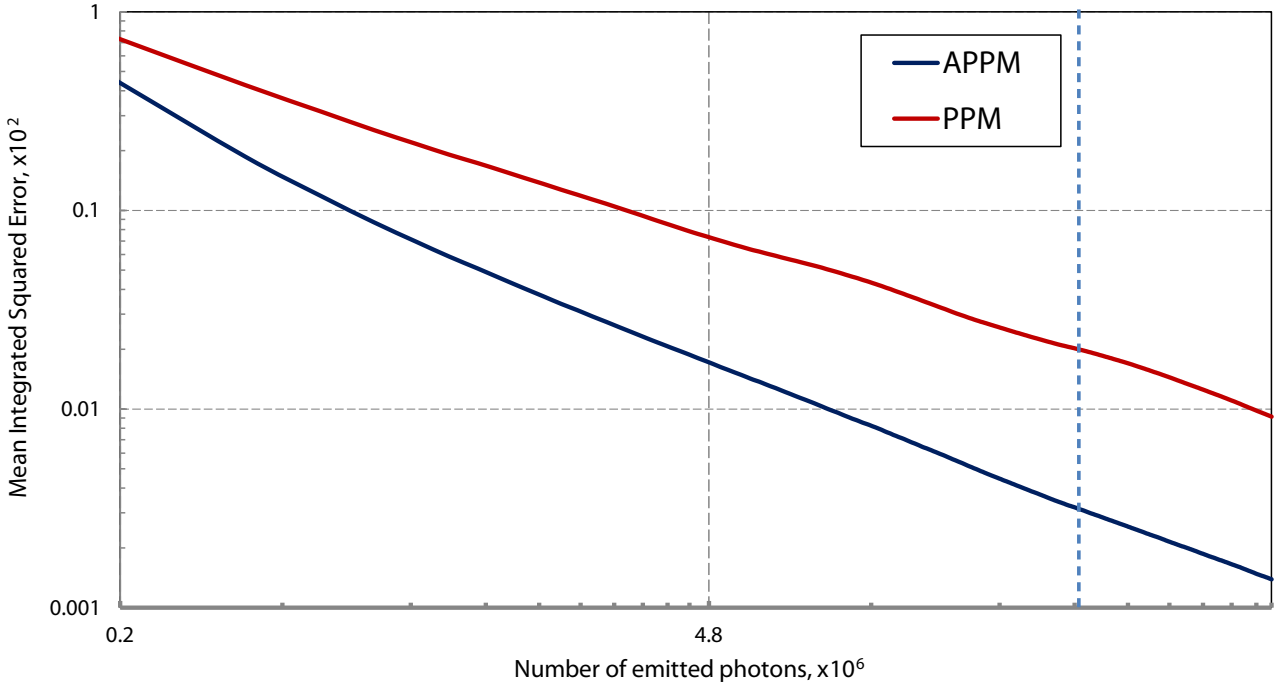

Fig. 2.5: Progressive rendering for the scene TORUS. The dashed vertical blue line roughly shows the stabilization of the bandwidth estimator.

3. CONVERGED RESULTS FOR DIFFERENT SCENES

In this section we show quasi-converged images for different scenes after emitting 100 billion photons (10^{11} photons). Note that bias remains in standard PPM if the initial bandwidth is not chosen close to the optimal value. The bandwidth range is specified in per mille with respect to the bounding radius of the corresponding scene.

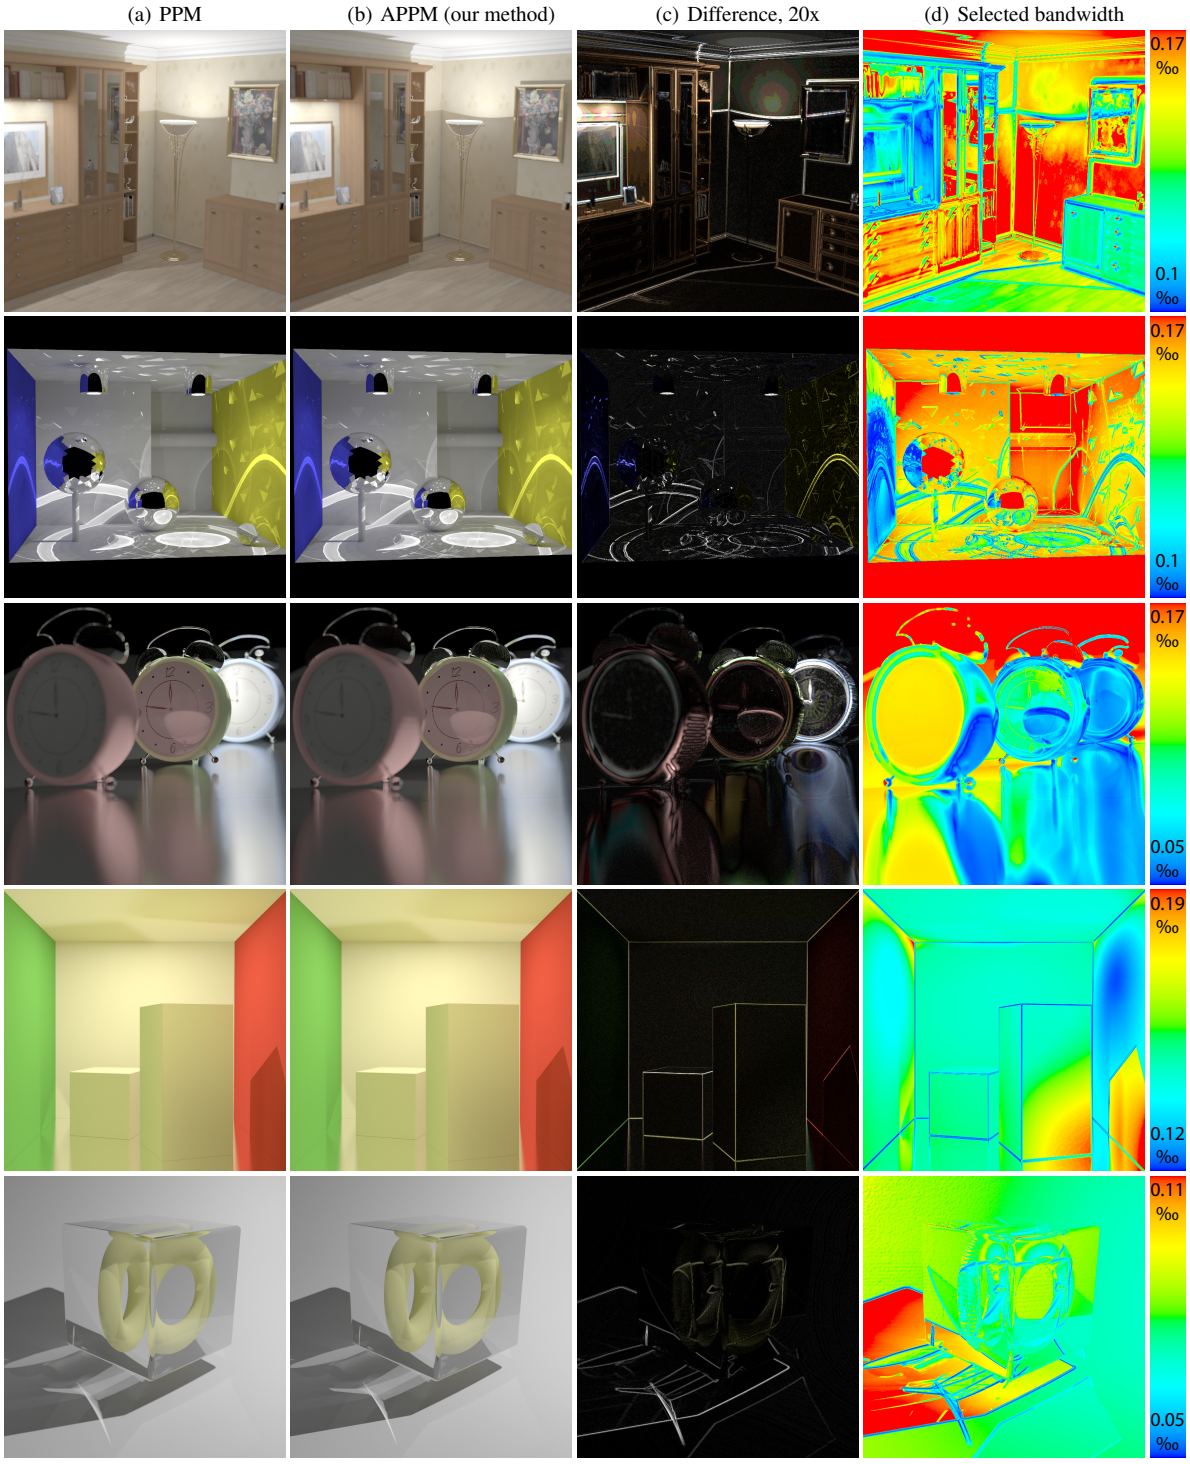

Fig. 3.1: Results for our test scenes rendering with standard PPM where the initial radius is chosen using k-NN search with $k = 10$ (first column), and with our APPM method which does not require user-specified parameters (second column). The third column shows the absolute differences; the fourth column visualizes the selected bandwidth of the APPM method.

4. EXACT SOLUTION TO THE AMSE

In Section 5.1 of the article we derived the asymptotic mean squared error of the pixel measurement. As it is a bicubic equation, it has a single proper root for r_N. Let us denote the following intermediate terms:

$$\Theta_N = 27J^2\pi^2\Delta\mathrm{I}^4k_2^4p_l^2\mathrm{Var}[\psi] - J^3\pi^3\Delta\mathrm{I}^6k_2^6p_l^3\Sigma_N^3 + 3\sqrt{3}\pi^2\sqrt{-J^4\Delta\mathrm{I}^8k_2^8p_l^4\left(2J\pi\Delta\mathrm{I}^2k_2^2p_l\Sigma_N^3 - 27\mathrm{Var}[\psi]\right)\mathrm{Var}[\psi]},$$

where $\Sigma_N = \sum_{i=1}^{N-1} r_i^2$.

Then the optimal bandwidth is

$$r_N = \frac{1}{\sqrt{3}}\sqrt{\frac{J\pi\Delta\mathrm{I}^2k_2^2p_l\Sigma_N^2}{\Theta_N^{1/3}} + \frac{\Theta_N^{1/3}}{J\pi\Delta\mathrm{I}^2k_2^2p_l} - \Sigma_N}.$$

This bandwidth is not only more expensive to compute, but it also requires storing additional values, such as Σ_N. Thus we propose to use the simplified formula as described in the article.

5. APPROXIMATE SOLUTION TO THE AMSE

In Section 5.1 of the article, we also introduced a simplified term $AMSE'$ and propose to use it for the bandwidth selection strategy. Here we study the impact of this simplification on the estimation.

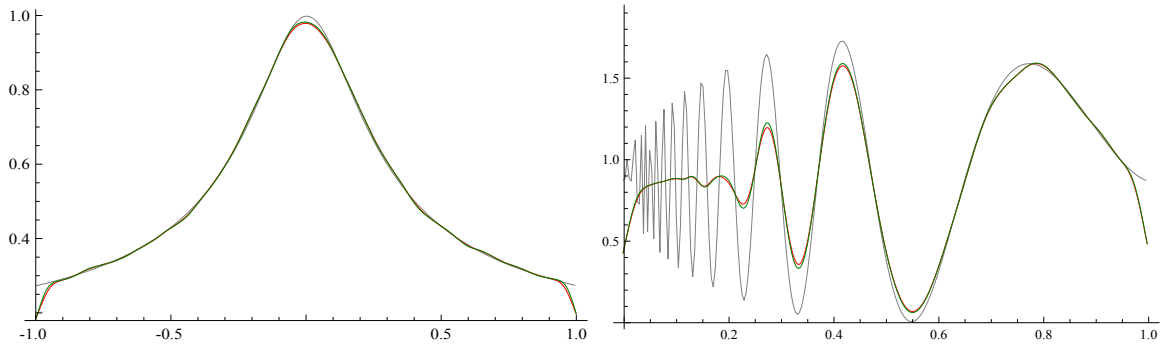

Fig. 5.1: Estimation plots for the RADIANCE (left) and the DOPPLER (right) design functions after drawing 10.000 samples. We compare the bandwidth selection strategy using the exact AMSE (green) and the approximate AMSE' (red). Note that both yield almost identical performance.

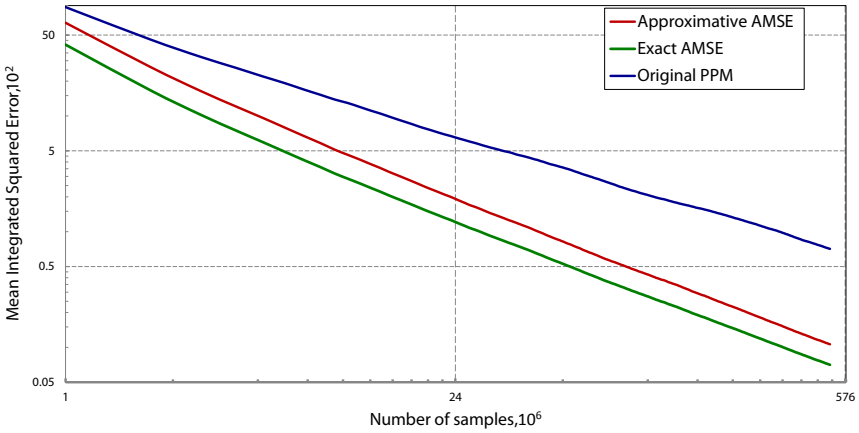

Fig. 5.2: The MISE computed for the design function DOPPLER using the bandwidth obtained from the solution for the exact AMSE (green) and approximate AMSE' (red). The MISE of an estimator with constant bandwidth (blue) is shown for comparison.

6. COMPARISON OF ADAPTIVE BANDWIDTH SELECTION TO K-NEAREST-NEIGHBORS SEARCH

In this section we provide a comparison sequence of bandwidth selection (radii shown color-coded) using two methods: our adaptive selector and the classic k-nearest-neighbors (k-NN) search (range query), with different number of photons. For the k-NN bandwidth selector, we empirically determined the globally optimal parameter $k = 7$. It is updated with every new photon map and is reduced based on the optimal shrinkage rate $O(N^{-1/6})$. It can be observed from the sequence that the k-NN bandwidth selector relies on the local density of samples to compute the bandwidth. Unfortunately this leads to higher bandwidths at all types of radiance discontinuities (notice the smooth bandwidth gradient on shadow edges) and underestimated bandwidth on surfaces with relatively flat illumination (e.g., on the floor). Consequently, this leads to increased bias at hard shadow edges and, at the same time, to lower convergence in flat regions.

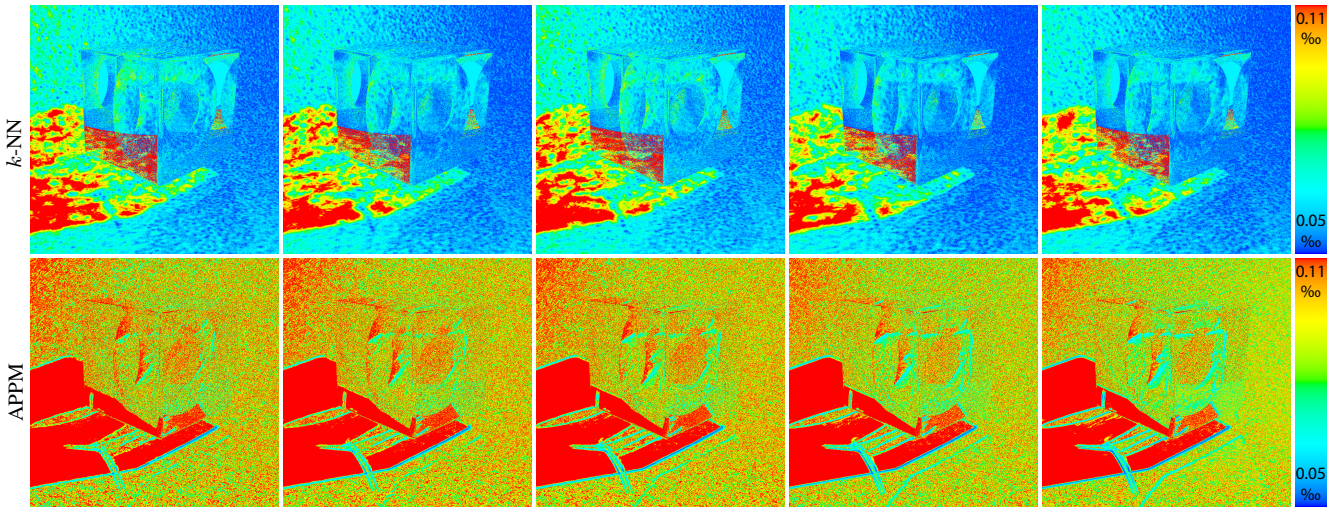

(a) $5 \cdot 10^6$ photons (b) $10 \cdot 10^6$ photons (c) $20 \cdot 10^6$ photons (d) $50 \cdot 10^6$ photons (e) $100 \cdot 10^6$ photons

Fig. 6.1: Progressive rendering for the scene TORUS. The image shows the false-colored sequence of bandwidth evolution depending on the bandwidth selection algorithm.

7. COMPARISON OF PPM WITH APPM USING PERCEPTUAL ERROR PREDICTION

In this section we show difference images of the BOX scene rendered with PPM and APPM generated using the High-Dynamic Range Visual Difference Predictor (HDR-VDP 2.1.1) [Mantiuk et al. 2011]. We used the following display settings for the VDP: resolution 768×768 on a diagonal of 7 inches (which corresponds to an image of that resolution displayed on a 26-inch monitor with Full-HD resolution), viewed from a distance of 80 cm. The images produced by APPM and PPM are compared to the reference image shown on the bottom (computed with APPM). The false-color image was overlaid onto the reference. The color scheme depicts the probability of a human perceiving the difference: red is 1 (definitely recognizable), blue is 0. We can observe that the artifacts in an early intermediate image (5 million photons) are slightly more perceivable with APPM (on the floor and the right wall) due to more low-frequency patterns produced by our method. However, this changes quickly: with 20 million photons APPM has already clearly less noticeable differences than PPM. For PPM we use optimal parameter $\alpha = 2/3$ and k-NN bandwidth selector with the globally optimal $k = 7$ (empirically determined).

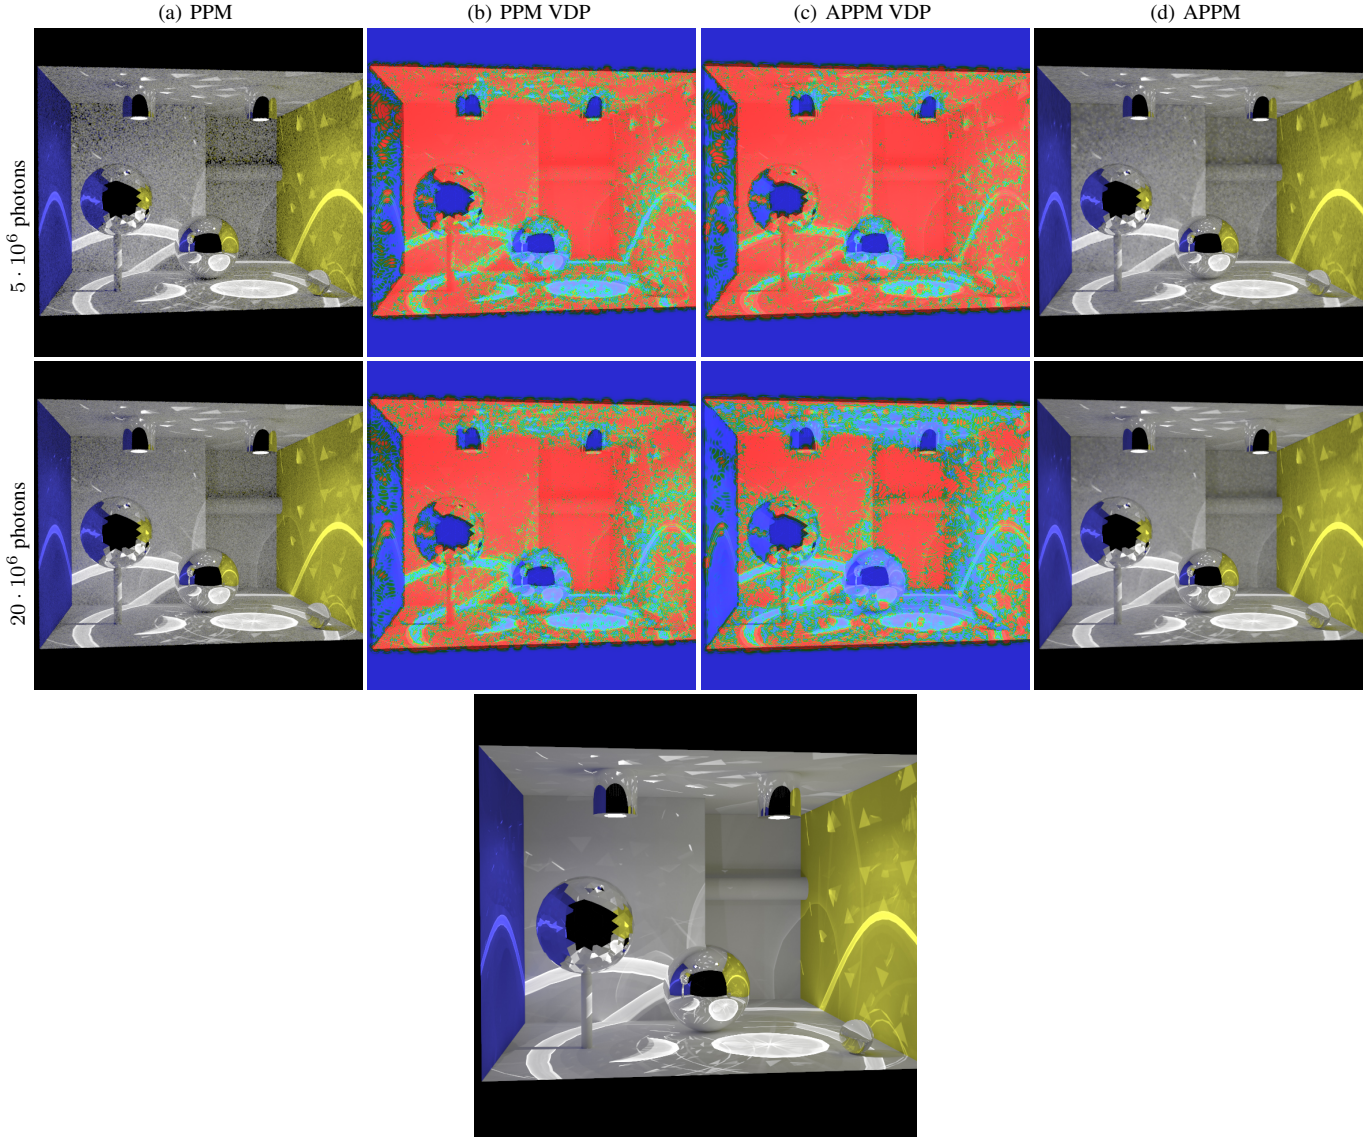

Fig. 7.1: PPM and APPM compared in the BOX scene using visual difference predictor.

8. COMPARISON TO BIDIRECTIONAL PATH TRACING

In this section we provide several comparison images with PPM and BiDirectional Path Tracing (BDPT) [Lafortune and Willems 1993; Veach and Guibas 1994] to illustrate that BDPT does not handle all transport paths. We used explicit connections to the camera to capture caustics with point light sources in BDPT. One can clearly see the lack of illumination in parts containing reflected caustics (specular-diffuse-specular paths). Such paths are impossible to explicitly construct with an unbiased method. Moreover the construction is an undecidable process, which means that there is no algorithm that can certainly determine such a path. As a consequence, the images of the BOX scene appear brighter than in the previous section, as we had to make the maximum number of ray-surface interactions equal for PPM and BDPT.

(a) BDPT (b) PPM

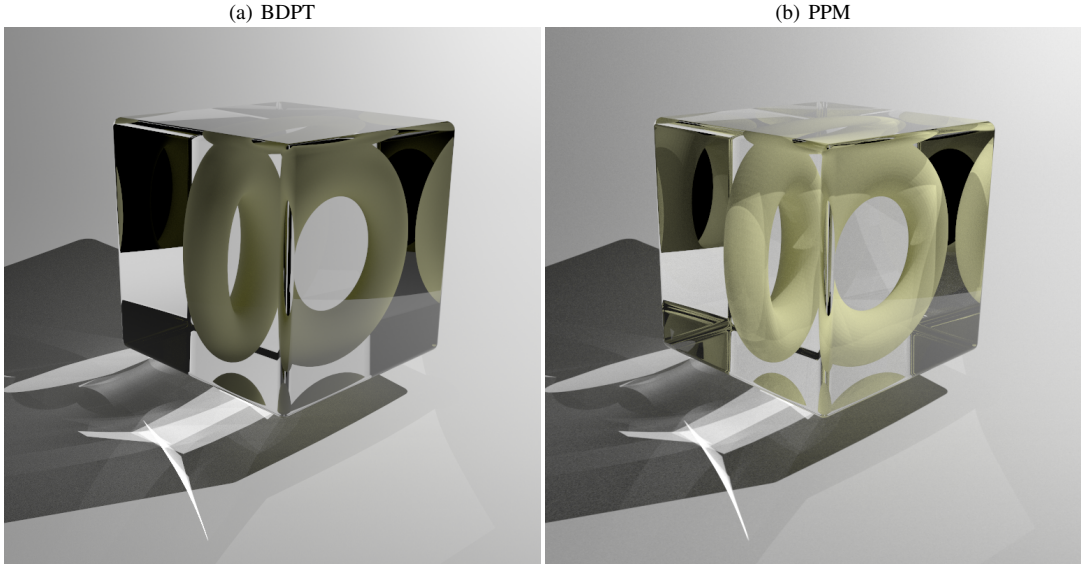

Fig. 8.1: The TORUS scene rendered until convergence with the unbiased BDPT method and PPM. The scene is rendered using a pinhole camera and a single point light source. The lack of illumination inside the glass cube is due to the impossibility of sampling these paths in an unbiased way.

(a) BDPT (b) PPM

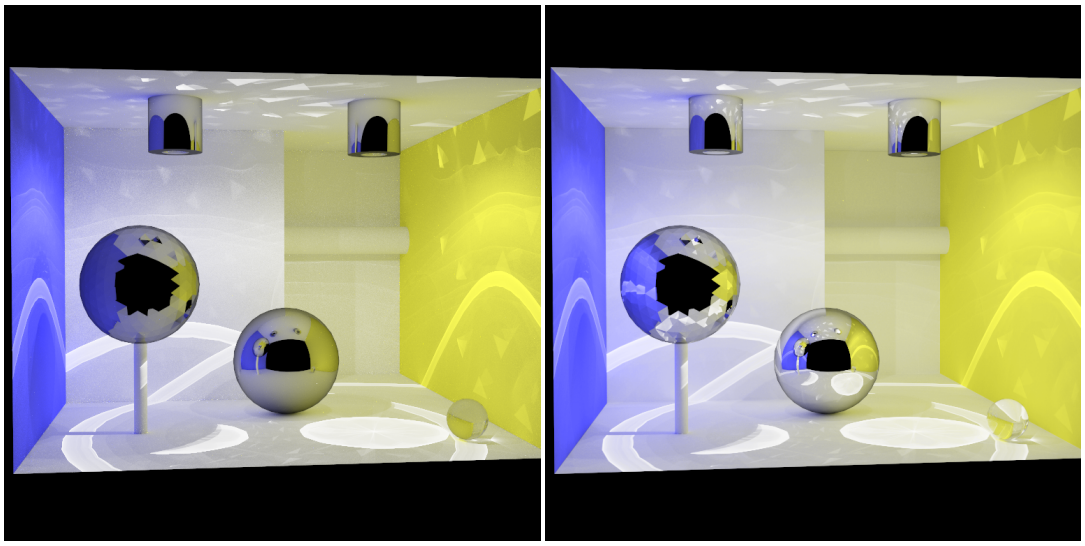

Fig. 8.2: The BOX scene rendered until convergence with the unbiased BDPT method and PPM. The scene is again rendered using a pinhole camera and a single point light source. Here the lack of illumination is noticeable on the pure glass spheres. Also the illumination of the ceiling is slightly different, as S-D-S paths are not taken into account in BDPT.

REFERENCES

HACHISUKA, T., JAROSZ, W., AND JENSEN, H. W. 2010. A progressive error estimation framework for photon density estimation. *ACM Transactions on Graphics 29,* 6, 144:1–144:12.

LAFORTUNE, E. P. AND WILLEMS, Y. D. 1993. Bi-directional path tracing. In *Proc. of 3rd int. Conference on Computational Graphics and Visualization Techniques (COMPUGRAPHICS)*. 145–153.

MANTIUK, R., KIM, K. J., REMPEL, A. G., AND HEIDRICH, W. 2011. Hdr-vdp-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. *Transactions on Graphics (Proc. of ACM SIGGRAPH) 30,* 4, 40:1–40:14.

PERLIN, K. 2002. Improving noise. *ACM Transactions on Graphics 21,* 3, 681–682.

VEACH, E. AND GUIBAS, L. 1994. Bidirectional estimators for light transport. In *Proc. of Eurographics Rendering Workshop*. 147–162.