

CPPM: Chi-squared Progressive Photon Mapping

ZEHUI LIN, SHENG LI*, XINLU ZENG, and CONGYI ZHANG, Dept. of Computer Science and Technology, Peking University

JINZHU JIA, Dept. of Biostatistics and Center for Statistical Science, Peking University

GUOPING WANG, Dept. of Computer Science and Technology, Peking University

DINESH MANOCHA, University of Maryland at College Park




Fig. 1. The Glass scene rendered by our approach (CPPM) and APPM [Kaplanyan and Dachsbacher 2013]. This scene includes transparent objects, a glossy floor, and complex caustics. Both methods render the scene using 10,000 iterations (655M photons), and the bandwidth is initialized by $k_{NN} = 10$ for k-NN search. For APPM, MSE=1.369, while for CPPM, MSE=1.047. We illustrate zoomed-in results and their heat map visualization, which show the absolute difference w.r.t. the reference image, and the blue to red color scale corresponds to a small to large error range. CPPM exhibits better rendering quality.

We present a novel chi-squared progressive photon mapping algorithm (CPPM) that constructs an estimator by controlling the bandwidth to obtain superior image quality. Our estimator has parametric statistical advantages over prior nonparametric methods. First, we show that when a probability density function of the photon distribution is subject to uniform distribution, the radiance estimation is unbiased under certain assumptions. Next, the local photon distribution is evaluated via a chi-squared test to determine whether the photons follow the hypothesized distribution (uniform distribution) or not. If the statistical test deems that the photons inside the bandwidth are uniformly distributed, bandwidth reduction should be suspended. Finally, we present a pipeline with a bandwidth retention and conditional reduction scheme according to the test results. This pipeline not only accumulates sufficient photons for a reliable chi-squared test, but also guarantees that the estimate converges to the correct solution under our assumptions. We evaluate our method on various benchmarks and observe significant improvement in the running time and rendering quality in terms of mean squared error over prior progressive photon mapping methods.

*corresponding author: lisheng@pku.edu.cn
Project URL: [bactlink.github.io/CPPM](https://github.com/bactlink/CPPM)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.
0730-0301/2020/12-ART1 \$15.00
<https://doi.org/10.1145/3414685.3417822>

CCS Concepts: • **Computing methodologies** → **Rendering**; **Ray tracing**.

Additional Key Words and Phrases: progressive photon mapping, chi-squared test, probability density function, radiance estimate, mean squared error

ACM Reference Format:

Zehui Lin, Sheng Li, Xinlu Zeng, Congyi Zhang, Jinzhu Jia, Guoping Wang, and Dinesh Manocha. 2020. CPPM: Chi-squared Progressive Photon Mapping. *ACM Trans. Graph.* 39, 6, Article 1 (December 2020), 12 pages. <https://doi.org/10.1145/3414685.3417822>

1 INTRODUCTION

Monte Carlo path tracing is widely used to solve the rendering equation without introducing bias in photo-realistic rendering [Dutre et al. 2006]. However, this method is not the optimal choice for accurate lighting simulation, especially for caustics emanating from small light sources. Photon mapping (PM) can solve the rendering equation by simulating the propagation and distribution of energy through a large number of photons [Jensen 1996]. Compared to path tracing, PM has many advantages in terms of rendering caustics and other phenomena caused by light transport in Specular-Diffuse-Specular (SDS) paths, including lower-frequency glossy and diffuse inter-reflections [Jensen 2001]. However, the variance and bias problems that arise due to kernel estimation in PM are regarded as major issues. One way to alleviate these problems is to use as many photons as possible. However, the memory overhead can be high, and these artifacts may still exist.

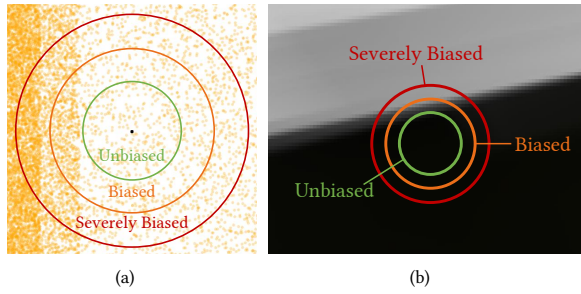


Fig. 2. Illustration of conditions for biased/unbiased estimation. (a) photon distribution within different bandwidths may lead to unbiased/biased results; (b) a rendering instance corresponding to (a). Uniformly distributed samples tend to make an unbiased estimation.

Progressive photon mapping (PPM) theoretically solves these problems by gathering infinite photons with limited memory overhead, progressively alleviating the bias [Hachisuka et al. 2008]. PPM generally launches multi-passes of photons and estimates the radiance using progressive bandwidth reduction. There has been considerable work on improving the performance of PPM, including Stochastic Progressive Photon Mapping (SPPM) [Hachisuka and Jensen 2009], Adaptive Progressive Photon Mapping (APPM) [Kaplanyan and Dachsbacher 2013], and other variants [Fu and Jensen 2012; Knaus and Zwicker 2011; Weiss and Grosch 2012].

As the state of the art, APPM obtains an optimal convergence of bandwidth, which is based on the adaptive estimation of parameters of local distributions [Kaplanyan and Dachsbacher 2013]. From a statistical view, this is a nonparametric method, which does not hypothesize about the photon distribution. However, it is also well-known that nonparametric methods are usually less efficient than parametric methods when useful prior knowledge is utilized [Györfi et al. 2006]. We find that a uniform distribution of photons within the bandwidth may generally lead to unbiased estimation, as illustrated in Figure 2. Our goal is to design an algorithm that utilizes this finding by introducing the hypothesis about photon distribution and statistical analysis of the distribution.

Main Results: We present a chi-squared progressive photon mapping algorithm (CPPM) that uses the chi-squared test [Bagdonavicius and Nikulin 2011] to find a bandwidth such that the photon distribution is consistent with a uniform distribution hypothesis and leads to a feasible estimation of radiance. If such a bandwidth for a pixel is determined, its bandwidth can remain unchanged while the radiance estimation still theoretically converges to the true value. The main contributions of our work include:

- We prove that the probability density function of the photon distribution subject to a uniform distribution would theoretically result in an unbiased estimator under our location-independent assumption.
- We propose a progressive approach that can find a fixed bandwidth to achieve high-quality rendering when the photon distribution inside is consistent with a statistical hypothesis, i.e. uniform distribution. The local photon distribution is evaluated via a chi-squared test to determine whether it is consistent with the hypothesis or not.

- We present a pipeline based on our novel scheme to handle bandwidth reduction and preservation according to the chi-squared test result. The bandwidth reduction scheme helps reduce bias while maintaining low variance.
- We provide a theoretical analysis of our algorithm's fast convergence and show that it degenerates to SPPM in the worst case, while it can theoretically converge to the correct value under our assumptions.

2 RELATED WORK

In this section, we give a brief overview of prior works in photon mapping and related areas.

2.1 Photon Mapping and Progressive Schemes

Photon mapping (PM) has been widely used to render different scenarios and phenomena, including participating media [Jarosz et al. 2011], time-varying volume data [Jönsson and Ynnerman 2017], and indirect illumination in diffuse or glossy scenes [Herzog et al. 2007]. It calculates a *hit point* for each pixel and sets a *bandwidth* to collect photons located inside it. The region within this bandwidth is a *searching area*. Progressive photon mapping (PPM) [Hachisuka et al. 2013, 2008] is a popular progressive scheme that uses a parameter α to progressively reduce the bandwidth. Stochastic PPM (SPPM) [Hachisuka and Jensen 2009] makes it possible to compute an accurate average radiance value in a region (pixel's measurement) using multiple distributed ray tracing passes.

In general, reducing the bias (visual blur) or variance (visual noise) to yield a high-quality image can be challenging in photon mapping. There is a considerable amount of work on improving rendering performance on photon density estimates [Hachisuka et al. 2013]. Hachisuka et al. [2010] presented a quantitative error estimation framework for photon density. Fu and Jensen [2012] described an adaptive noise reduction technique. Schregle [2003] presented a bias compensating operator that adapts the bandwidth to the photon distribution. Knaus and Zwicker [2011] presented a probabilistic derivation that includes convergence analysis for variance and bias. Kaplanyan and Dachsbacher [2013] deduced an optimal parameter of $\alpha = 2/3$ for pixel measurement estimation to minimize asymptotic mean squared error (AMSE) and proposed adaptive PPM that optimally balances bias and variance to minimize overall error based on approximations. These methods can analyze the bias and variance and improve the accuracy of the radiance estimate. We also try to improve convergence and performance of PPM by reducing bias and variance.

2.2 Other Improvements on Photon Mapping

PPM has also been extended to handle dynamic scenes [Weiss and Grosch 2012], integrated with bidirectional path tracing into a framework using vertex connection and merging (VCM) [Georgiev et al. 2012] or unified path sampling (UPS) [Hachisuka et al. 2012], integrated with Markov Chain Monte Carlo methods [Šik et al. 2016], used with out-of-core distributed processing [Günther and Grosch 2014], and used as a temporal estimator for transient rendering [Jarabo et al. 2014].

Some considerations of photon gathering include a hierarchical algorithm used to compute the radiance from a photon map [Spencer and Jones 2009a], GPU-based spatial data caching and interpolation [Wang et al. 2009], and unbiased photon gathering [Qin et al. 2015]. Photon tracing techniques can result in uniform relative error in the radiance estimate [Chen et al. 2011; Gruson et al. 2016; Hachisuka and Jensen 2011]. Optimized emission strategies can automatically guide photons towards regions and provide substantial variance reduction [Grittmann et al. 2018]. Photon relaxation methods are used to remove noise from photon maps and to minimize the residual bias [Spencer and Jones 2009b, 2013a,b]. Different filter kernels [Hernández et al. 2014], varying bandwidths [Schregle 2003], adaptive kernel shapes [Schjøth 2009], and boundary bias removal [Havran et al. 2005] have also been proposed for radiance estimates. Jakob et al. [2011] used a Gaussian mixture model to efficiently create noise-free renderings. Zhu et al. [2020] applied deep neural network to predict a kernel function to aggregate photon contributions for radiance estimation. In contrast to these methods, we aim to improve the rendering performance by finding a bandwidth with low bias and variance.

3 MOTIVATION AND RATIONALE

Conventional progressive photon mapping formulations assume that the estimation can only converge to the true pixel value when the bandwidth converges to zero and the cumulative number of photons is unlimited [Hachisuka et al. 2008]. APPM [Kaplanyan and Dachsbacher 2013] presented the optimal bandwidth convergence rate $O(N^{-1/6})$ in terms of asymptotic MSE (AMSE), as well as the computation of the optimal bandwidth by estimating variance and bias during each iteration with approximations. However, the estimated variance and bias may not provide sufficient accuracy when the sample size is not big enough, which limits its performance.

If the radiance estimation is unbiased within a fixed bandwidth, i.e. the expected value of the estimate equals the true value, the radiance estimated using an infinite number of photons can still converge to the true value according to the Big Number Theory, even if this bandwidth does not converge to zero. Based on this principle, a key aspect of our approach is to determine whether a bias-free region with a bandwidth can be found (illustrated in Figure 2). Our solution to find such bandwidth is to progressively reduce the bandwidth, and suspend the reduction based on the *chi-squared test*, which can be used to evaluate the property of photon distribution. In the field of statistics, this test can evaluate the goodness of fit between an observed distribution and a hypothetical probability distribution [Cochran 1952].

We organize the exposition of our approach as follows:

- We first set up the theoretical basis of the radiance estimation of our algorithm (Section 4). We argue that if the photon distribution within the bandwidth is consistent with uniform distribution, the pixel measurement estimation is unbiased under our assumptions. The uniform distribution serves as our hypothesized distribution. Under this condition, this bandwidth should not reduce.
- Next, the key issue is to find an effective solution for identifying whether the photon distribution is consistent with

the hypothesized distribution or not. To this end, we use a chi-squared test to perform hypothesis testing on the photon distribution within the bandwidth (Section 5).

- The bandwidth should reduce when the chi-squared test deems the photons are not consistent with hypothesized uniform distribution. However, prior reduction schemes can not work well with this test. Therefore, we propose a novel reduction scheme that guarantees that the estimate converges to the correct value under our assumptions. (Section 6).
- Based on these principles, we present a novel pipeline of our progressive photon mapping algorithm that takes into account the pixel measurement estimation, the chi-squared test, and the bandwidth reduction scheme (Section 7).

4 THEORETICAL FOUNDATION

In this section, we start from a formulation of the pixel measurement estimation and its unbiased condition within a bandwidth; then we present the photon density function for multiple hit points based on a unified space; finally, we propose a sufficient unbiased condition for radiance estimation under our assumptions.

4.1 Pixel Measurement Estimation

The pixel measurement I in photon mapping can be estimated by kernel estimation, as proposed in APPM [2013] :

$$\hat{I} = \frac{1}{NJ} \sum_{i=1}^N \sum_{j=1}^J k_{R_i}(\vec{x}_i - \vec{p}_{i,j}) \psi_{i,j}, \quad (1)$$

where \hat{I} is the estimate of the pixel measurement I ; N is the number of hit points of a pixel; J is the size of the photon map per iteration; R_i is the bandwidth at the i -th iteration; k_{R_i} is a normalized function with bandwidth R_i ; \vec{x}_i is the location of the hit point; $\vec{p}_{i,j}$ is the location of the j -th photon in the i -th photon map; and $\psi_{i,j}$ is the contribution of the full path constructed between the eye sub-path ending at \vec{x}_i and the light sub-path ending at $\vec{p}_{i,j}$. $\psi_{i,j}$ includes various factors such as BRDF value, contributions of the sub-paths, and other possible factors [Kaplanyan and Dachsbacher 2013].

We make a location-independent assumption to regulate $\psi_{i,j}$, i.e. we assume that $\psi_{i,j}$ is independent of the location of the photon $\vec{p}_{i,j}$ for any hit point \vec{x}_i . This assumption can help to obtain a theoretically unbiased condition, which will be analyzed and discussed in Subsection 4.3. Similar to conventional methods, an isotropic and circular on-surface kernel function and a planar assumption of the searching area [Hey and Purgathofer 2002] are used. In addition, we assume that any two hit points are independent, and any two photons are also independent.

From Eq. (1) and according to the principle in Section 3, having a bias-free bandwidth R for the pixel measurement I should satisfy

$$I = E \left[\frac{1}{NJ} \sum_{i=1}^N \sum_{j=1}^J k_R(\vec{x}_i - \vec{p}_{i,j}) \psi_{i,j} \right], \quad (2)$$

where $E[X]$ indicates the expectation of expression X .

Eq. (2) refers to the estimation of a probability density function based on statistical sampling. To take advantage of the uniform distribution of photons as illustrated in Figure 2, we transform the




Fig. 3. Illustration of different hit points' searching areas being aligned to a unified domain Ω_R . All the photons inside the corresponding searching areas are also mapped to this unified space.

sampling of photons in Eq. (2) to an integration of the photon density function as

$$I = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_{\mathcal{M}} k_R(\vec{x}_i - \vec{x}) p(\vec{x}) \Psi(\vec{x}_i, \vec{x}) d\vec{x} \right], \quad (3)$$

where \mathcal{M} is the manifold of the all scene surfaces in \mathbb{R}^3 ; p is the photon density function corresponding to $\psi_{i,j}$ in Eq. (2), and $\Psi(\vec{x}_i, \vec{x})$ indicates the contribution of all possible full paths constructed by the eye path ending at \vec{x}_i and the light paths ending at \vec{x} .

4.2 Density Function in Unified Space

Given that the N hit points in Eq. (3) are located at different positions, their associated searching areas also do not coincide. Therefore, we construct a unified space Ω with a unified domain Ω_R on it, and we then align different hit points to the origin of this 2D space Ω and their searching areas to this unified domain Ω_R . This process can be treated as a conversion from the scene space to a unified space, as illustrated in Figure 3. A position \vec{x} within the searching area of \vec{x}_i can be mapped to $\vec{x}' \in \Omega_R$ through

$$\vec{x}' = (\langle \vec{x} - \vec{x}_i, \vec{u}_i \rangle, \langle \vec{x} - \vec{x}_i, \vec{v}_i \rangle), \quad (4)$$

where \vec{u}_i and \vec{v}_i are two orthonormal bases specified on the tangent plane at hit point \vec{x}_i , and $\langle \cdot, \cdot \rangle$ indicates the operator of the inner product.

Then the photon density function over \vec{x}_i 's searching area is mapped to be a function \tilde{p}_i over Ω_R :

$$\tilde{p}_i(\vec{x}'(x'_u, x'_v)) = p(x'_u \vec{u}_i + x'_v \vec{v}_i + \vec{x}_i). \quad (5)$$

Specifically, $\tilde{p}_i(\vec{0}) = p(\vec{x}_i)$. Function \tilde{p}_i over Ω_R is an aligned 2D form of the PDF centered at \vec{x}_i .

Similarly, $\Psi(\vec{x}_i, \vec{x})$ is also mapped to the unified space as

$$\tilde{\Psi}_i(\vec{x}'(x'_u, x'_v)) = \Psi(\vec{x}_i, x'_u \vec{u}_i + x'_v \vec{v}_i + \vec{x}_i). \quad (6)$$

After the N photon density functions are aligned, we define \bar{p} , the *average photon density function (APDF)* over Ω_R as

$$\bar{p}(\vec{x}') = \frac{1}{N} \sum_{i=1}^N \tilde{p}_i(\vec{x}'). \quad (7)$$

4.3 Unbiased Condition

The condition for a bias-free bandwidth in Eq. (3) can be transformed to an integration over the unified domain Ω_R by Eqs. (5), (6) as

$$I = \mathbb{E} \left[\int_{\Omega_R} k_R(\vec{x}') \frac{1}{N} \sum_{i=1}^N \tilde{p}_i(\vec{x}') \tilde{\Psi}_i(\vec{x}') d\vec{x}' \right]. \quad (8)$$

Under the location-independent assumption, $\tilde{\Psi}_i(\vec{x}')$ becomes a constant function for a given hit point \vec{x}_i because $\tilde{\Psi}_i(\vec{x}')$ is independent to the location \vec{x} . We can then take $\tilde{\Psi}_i(\vec{x}')$ outside of the integration, and by applying Eq. (7), we have

$$I = \mathbb{E} \left[\int_{\Omega_R} k_R(\vec{x}') \bar{p}(\vec{x}') d\vec{x}' \right] \mathbb{E} [\tilde{\Psi}]. \quad (9)$$

If APDF is a constant function, i.e. $\bar{p}(\vec{x}') \equiv \bar{p}(\vec{0})$, then

$$I = \mathbb{E} \left[\int_{\Omega_R} k_R(\vec{x}') d\vec{x}' \right] \mathbb{E} [\bar{p}] \mathbb{E} [\tilde{\Psi}]. \quad (10)$$

Obviously, the first term on the right side equals one. Note that since the expected value of a radiance estimate with an infinitely small kernel bandwidth corresponds to the exact radiance [Knaus and Zwicker 2011], we can express I as

$$I = \mathbb{E} [\tilde{\Psi}] \mathbb{E} [\delta(\vec{x}')] = \mathbb{E} [\tilde{\Psi}] \mathbb{E} [\bar{p}(\vec{0})], \quad (11)$$

where δ is the Dirac delta function. Eq. (11) explains that Eq. (10) is true. In summary, the pixel measurement estimation is unbiased if APDF is a constant function (uniform distribution of photons) under our assumptions. We provide a detailed proof in Appendix A.

In general, radiance over a surface is mainly determined by the photon density. Therefore, when the photons are uniformly distributed within a relatively small region, the contributions of the constructed paths are usually close and can be approximately treated as irrelevant to the photons' locations. Consequently, when the APDF is a constant function, the location-independent assumption is usually tenable. However, there still exist some special cases in practice in which this location-independent assumption may be violated. As a result, bias may be introduced in these cases. We will discuss these cases in Subsection 9.4.

5 PHOTON DISTRIBUTION EVALUATION

In this section, we present techniques to determine whether the photons satisfy the constant condition of APDF.

5.1 Chi-squared Test Overview

We use the chi-squared test to evaluate the property of photon distribution through an evaluation of the goodness of fit. The chi-squared test can evaluate the goodness of fit between an observed distribution and a hypothetical probability distribution [Cochran 1952]. It first sets up a *null hypothesis*, i.e. there is *no* difference between the observed distribution and the hypothetical probability distribution. Next, it calculates a statistic (chi-squared statistic) based on the observed distribution to decide whether to reject the null hypothesis. Compared to other goodness of fit tests, including the KS-test (Kolmogorov-Smirnov test) and the Cramér-von Mises criterion

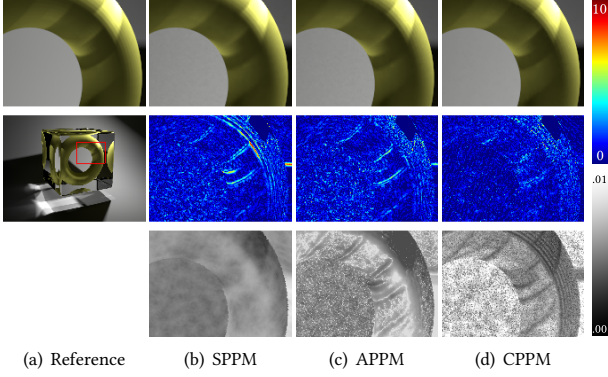


Fig. 4. Validity of photon alignment on a curved surface (Torus). We highlight the adaptive bandwidth (bottom row) computed by CPPM with better results (error visualization in the middle row). The black to white color scale corresponds to a small to large bandwidth range.

[Stephens 1974], the chi-squared test is lightweight because it does not require an empirical distribution function of the samples.

5.2 Photon Projection and Alignment

To perform the chi-squared test on the accumulated photons within a bandwidth, all the photons in multiple searching areas should be aligned to the unified domain Ω_R according to the conversion in Subsection 4.2 as illustrated in Figure 3.

Given that the photons associated with a hit point can be located on curved surfaces, the photons should first be projected to the tangent plane of this hit point \vec{x}_i according to the planar assumption. Each projected photon maintains its attributes, including orientation and distance, with respect to hit point, as well as its contribution to the radiance. We then align these projected photons to the unified domain Ω_R by Eq. (4). The above process can handle the photon at any position in the scene space regardless of whether the surface is flat or there is occlusion due to the projection. Taking a Torus scene as the example, we show our strategy is workable for the curved surface and can compute an appropriate bandwidth with better results, as shown in Figure 4. We provide a formulation of this process that combines the projection and alignment of the photons in Appendix B.

5.3 Chi-squared Test over Photon Distribution

In this section, we explain how to perform the chi-squared test, including how to calculate the chi-squared test statistic and when to reject the null hypothesis. The null hypothesis of the chi-squared test in CPPM is that there is no difference between the aligned photons and uniform distribution, i.e. the APDF is a constant function.

The chi-squared test requires grouping the samples into categories. To this end, we first partition the unified domain Ω_R into n_a concentric annuli of equal areas. Next, we further partition each annulus into n_s sectors of equal areas, as illustrated in Figure 5. The effect of different values of n_a and n_s will be discussed in Subsection 9.3. Using equally sized sectors makes it easier to compute the sector to which each photon belongs. Other partitioning schemes can also be used if the sector size is not too small.




Fig. 5. Disc shaped domain Ω_R with bandwidth R_i (left); partitioning of this disc into annuli with equal areas, and then further partitioning of annuli into sectors with equal areas (middle). When reducing the bandwidth, our method tests the distribution in every sub-disc (concentric discs with different radius), finds one (green region) not rejected by the chi-squared test, and assigns its radius to the bandwidth R_{i+1} (right).

The value of the chi-squared test statistic is

$$\chi^2 = \sum_{a=1}^{n_a} \sum_{s=1}^{n_s} \frac{(O_{a,s} - Mp_{a,s})^2}{Mp_{a,s}} = \sum_{a=1}^{n_a} \sum_{s=1}^{n_s} \frac{O_{a,s}^2}{Mp_{a,s}} - M, \quad (12)$$

where M is the total number of photons collected in Ω_R ; $p_{a,s}$ is the expected probability of the s -th sector of the a -th annulus under the null hypothesis, and $O_{a,s}$ is the number of photons in the s -th sector of the a -th annulus. M , $p_{a,s}$ and $O_{a,s}$ are shared by multiple hit points of a pixel. When all sectors have the same area and the null hypothesis is that the average photon density function is a constant function, $p_{a,s}$ is simply $\frac{1}{n_a n_s}$. The chi-squared test statistic approximately follows the chi-squared distribution with $(n_a n_s - 1)$ degrees of freedom if the null hypothesis is true.

A predetermined significance level α_χ (typically 0.05) is used to decide whether to reject the null hypothesis. The quantile $(1 - \alpha_\chi)$ of the chi-squared distribution with $(n_a n_s - 1)$ degrees of freedom is taken as a critical value c_χ , i.e. the probability that X larger than c_χ is α_χ , where X is a random variable following the chi-squared distribution with $(n_a n_s - 1)$ degrees of freedom. The chi-squared test statistic χ^2 is compared to the critical value c_χ . If χ^2 is larger than c_χ , the null hypothesis is rejected. Otherwise, the chi-squared test does not reject the null hypothesis, and the distribution within this bandwidth is deemed to make an unbiased estimation under our assumptions.

6 PROGRESSIVE BANDWIDTH SCHEME

In this section, we introduce a novel bandwidth strategy based on the results of the chi-squared test in Subsection 6.1. Accordingly, we set a photon count lower bound for the bandwidth in Subsection 6.2, which can guarantee infinite photons being collected progressively for radiance estimation as well as sufficient photons for the chi-squared test.

6.1 Bandwidth Reduction

The result of the chi-squared test will be either "rejecting the null hypothesis" or "not rejecting the null hypothesis," as described in Subsection 5.3. If the test fails to reject the null hypothesis and deems that the estimation is unbiased, the bandwidth should remain unchanged. Otherwise, the bandwidth should be reduced to eliminate bias.




Fig. 6. Illustration of bias variation along with changing of the bandwidth in kernel estimation, taking the green point in the left figure as the central point. Compared to SPPM, CPPM can reduce its bandwidth to the largest possible bandwidth that can make bias-free estimation at the i -th iteration.

To determine a new bandwidth for the next iteration, an additional chi-squared test is performed on each *sub-disc* in Ω_R using the statistics of the previously collected photons. In order to find an ideal bandwidth as soon as possible, the bandwidth is reduced using the following scheme.

- If any of the distribution within a sub-disc is not rejected by the chi-squared test, we select the largest of the sub-discs not rejected by the test and assign its radius as the new bandwidth for the next iteration (right of Figure 5). The benefit of this approach is illustrated in Figure 6, and the reduced bandwidth can obtain a bias-free result while helping reduce the variance due to a relatively large bandwidth when compared to SPPM. Our reduction rate R_{i+1}/R_i is at least $\sqrt{\frac{1}{n_a}}$.
- If none of these distributions passes the test, we reduce the bandwidth in a novel way. Here, we introduce a new parameter k (constant) to regulate the next bandwidth as:

$$R_{i+1}^2 = kR_i^2, \quad 0 < k < 1. \quad (13)$$

Generally, k controls the rate of bandwidth reduction with a reduction rate $R_{i+1}/R_i = \sqrt{k}$. A smaller k is appropriate for reducing the bandwidth quickly but may also lead to an increase in variance due to an insufficient number of photons. We will theoretically analyze the convergence of our approach as a function of k in Section 8, and discuss the selection of k empirically in Subsection 9.3.

After the new bandwidth is determined, we discard the statistics of previously collected photons and partition the new domain $\Omega_{R_{i+1}}$.

6.2 Constraint on Photon Accumulation

Our method should follow the general rule that *the number of accumulated photons should increase* to guarantee convergence. Furthermore, the chi-squared test also requires sufficient photons to ensure its accuracy. Therefore, a certain number of photons needs to be accumulated before the chi-squared test is performed. However, when CPPM decides to reduce the bandwidth, the reduction rate R_{i+1}/R_i is at least $\min(\sqrt{k}, \sqrt{\frac{1}{n_a}})$ according to Subsection 6.1. This is very aggressive if the bandwidth keeps reducing continuously.

In order to satisfy the convergence rule mentioned above, we set a progressively increasing lower bound on photon accumulation. It is also used as a necessary condition for the chi-squared test, i.e. the

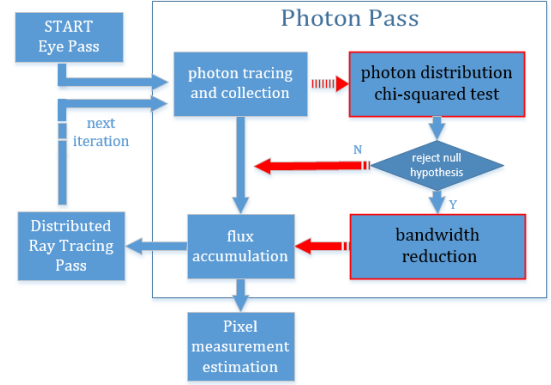


Fig. 7. Our pipeline with the novel modules highlighted. The chi-squared test is activated only if enough photons are collected, and the bandwidth reduction occurs only if the chi-squared test rejects the null hypothesis.

test will be performed only when the number of photons collected within the current bandwidth exceeds this lower bound.

We denote the lower bound as B and introduce a parameter β to control the increment of B to guarantee convergence. For a hit point, the lower bound at the i -th iteration is B_i , and successive B_{i+1} is given as

$$B_{i+1} = \begin{cases} B_i & \text{if } R_i \text{ is fixed in the } i\text{-th iteration;} \\ \beta B_i & \text{if } R_i \text{ is reduced in the } i\text{-th iteration,} \end{cases} \quad (14)$$

where β is a user-predefined parameter. $\beta > 1$ can guarantee infinite photons being collected for radiance estimation. B_1 is set to k_{NN} search to trigger the chi-squared test at the very beginning.

We will analyze the convergence of our algorithm due to β theoretically in Section 8 and discuss the selection of β empirically in Subsection 9.3.

7 ALGORITHM AND IMPLEMENTATION

Based on the aforementioned chi-squared test on the photon distribution and bandwidth strategy, we present our chi-squared PPM (CPPM) algorithm in this section.

7.1 Pipeline

The overall pipeline of the CPPM method is illustrated in Figure 7. Each iteration starts with an *eye pass* (i.e. ray tracing) to obtain a hit point for each pixel, followed by a *photon tracing and collection* module, where a new photon map is generated; photon collection is then performed. The two highlighted modules, *photon distribution chi-squared test* (see Section 5) and *bandwidth reduction* (see Section 6), are the two novel components in this pipeline. The chi-squared test will be activated only if a certain number of photons has been accumulated within the same bandwidth, as discussed in Subsection 6.2. Otherwise, the algorithm will jump from photon collection directly to a *flux accumulation* module while keeping the bandwidth unchanged. Meanwhile, if the chi-squared distribution does not reject the null hypothesis, the pipeline will also skip over the *bandwidth reduction* to a *flux accumulation* module without bandwidth reduction. Each pixel will be re-traced regardless of

the bandwidth reduction in the *distributed ray tracing pass*. Consequently, a new hit point associated with the corresponding pixel will be generated in the next iteration. During each iteration, the *pixel measurement estimation* module can synthesize the rendering results using Eq. (1).

In prior SPPM algorithms, the bandwidth is updated during each iteration. In contrast, the bandwidth reduction in CPPM happens only if the chi-squared test rejects the null hypothesis. Our method continues to collect photons and then checks for photon distribution even though the test failed to reject the null hypothesis in last iteration. Since any false null hypothesis will be rejected using infinite samples [Cochran 1952], our method can progressively reduce the probability of not rejecting a false null hypothesis.

7.2 Photon Collection

An extra data module is required to perform the chi-squared test for each pixel other than the conventional setting in SPPM. According to the method discussed in Subsection 5.3, the number of photons accumulated within the bandwidth and within each sector should be stored. Particularly, we discard those photons with zero contribution.

Specifically, we use two types of photon counters for each pixel in CPPM: one is to record the total number of photons in the current bandwidth, which is used to determine whether a sufficient number of photons has been collected; the other one is to record the number of photons in each sector for the chi-squared test statistic (Eq. 12). Therefore, only additional $(n_a n_s + 1)$ integers are required for each pixel, and the additional memory overhead (generally 10% ~ 20% in different scenes over SPPM) is constant during all these iterations.

8 CONVERGENCE ANALYSIS

Generally, the chi-squared test is assumed to produce a correct result, i.e. a true null hypothesis is never rejected by the chi-squared test, while a false null hypothesis is always rejected. We first analyze the worst case: the null hypothesis is always rejected, i.e. CPPM fails to find a desired bandwidth and the bandwidth keeps reducing. Next, we analyze the opposite case in which we can find the bandwidth that satisfies the hypothesis.

If the null hypothesis is always rejected, i.e., CPPM fails to find a bandwidth that satisfies the hypothetical model, the bandwidth will keep reducing and converge to 0. According to our progressive bandwidth scheme, the bandwidth convergence can be given as

$$R_N = O\left(N^{-\frac{1}{2} \log \frac{\beta}{k} \frac{1}{k}}\right). \quad (15)$$

We give a detailed analysis in Appendix C. If $\beta^{\alpha-1} = k^\alpha$, this bandwidth convergence is asymptotically equivalent to that of SPPM, which is $O\left(N^{(\alpha-1)/2}\right)$ according to [Knaus and Zwicker 2011].

We adopt the definition of the average radiance estimation error $\bar{\epsilon}_N$ from Knaus and Zwicker [2011]. Based on their probabilistic framework, we can derive the convergence rate of bias and variance

from the convergence rate of bandwidth:

$$E[\bar{\epsilon}_N] = \begin{cases} O(R_N^2) = O\left(N^{-\log \frac{\beta}{k} \frac{1}{k}}\right), & \beta > 1 \\ O\left(\frac{\log N}{N}\right), & \beta = 1 \end{cases} \quad (16)$$

$$\text{Var}[\bar{\epsilon}_N] = O\left(N^{-1} R_N^{-2}\right) = O\left(N^{-1 + \log \frac{\beta}{k} \frac{1}{k}}\right). \quad (17)$$

Specifically, the variance does not converge if $\beta = 1$. Based on the analysis in [Kaplanyan and Dachsbacher 2013] and the formulas in Eqs.(16,17), the optimal AMSE $[\hat{I}_N]$ can be obtained if $\text{Var}[\bar{\epsilon}_N] = E[\bar{\epsilon}_N]^2$, which requires $\beta = \frac{1}{k^2}$,

$$\text{AMSE}[\hat{I}_N] = O\left(N^{-\frac{2}{3}}\right). \quad (18)$$

We get the same optimal AMSE convergence rate as SPPM/APPM, when we have the same optimal bandwidth convergence rate.

If the chi-squared test does not reject the null hypothesis, based on our assumptions given in Subsection 4.1, the pixel measurement can be estimated in an unbiased way, while the bandwidth does not reduce. Therefore, the estimator becomes a pure Monte Carlo estimator, and we have $\text{Var}[\bar{\epsilon}_N] = O(N^{-1})$ and $E[\bar{\epsilon}_N] = 0$. Based on the analysis in [Kaplanyan and Dachsbacher 2013], $\text{AMSE}[\hat{I}_N]$ can be approximated as

$$\begin{aligned} \text{AMSE}[\hat{I}_N] &\approx \frac{1}{N} \text{Var}[\phi L] + E[\phi]^2 (\text{Var}[\bar{\epsilon}_N] + E[\bar{\epsilon}_N]^2) \\ &= O(N^{-1}). \end{aligned} \quad (19)$$

This indicates that CPPM can theoretically converge and achieve very fast convergence even if the bandwidth does not reduce.

Although CPPM degenerates to SPPM in the worst case, we still benefit from the situation when the samples are consistent with the null hypothesis. Although the chi-squared test cannot be expected to always produce correct results, our experimental results (Section 9) demonstrate that this test is quite reliable in practice.

9 EXPERIMENT AND DISCUSSION

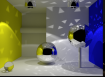


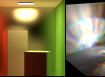
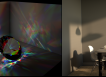




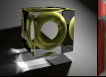
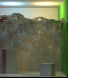
In this section, we provide experiments and results on various benchmarks and make comparisons between CPPM, SPPM [Hachisuka and Jensen 2009] (baseline), and APPM [Kaplanyan and Dachsbacher 2013] (state of the art). In addition, we discuss the effect of the parameters related to CPPM in the experiment, including initial value (k_{NN}), convergence rate related parameters (β and k), domain partition (n_a , n_s), and test significant level (α_χ).

9.1 Settings

In total, we use 11 benchmarks to evaluate the performance. The Box, Dining Room, Glass, and Water scenes are rendered based on the Mitsuba renderer [Jakob 2010], while others are rendered based on NVIDIA OptiX 6.0 [Parker et al. 2010], which demonstrates the compatibility and availability of our method. Mitsuba runs on an Intel(R) Xeon(R) Silver 4210 CPU on a Linux System, while OptiX runs on an NVIDIA GeForce RTX 2080 GPU with an Intel(R) Core(TM) i7-8700 CPU on a Windows system.

The initial bandwidths for the SPPM, APPM, and CPPM are determined by setting $k_{NN} = 10$ using k -NN search [Hachisuka and

Table 1. Benchmarks (in alphabetical order) for performance comparison on the iterations and time (sec.) between CPPM and APPM. MSE value obtained from SPPM after 1,000 iterations is used as the baseline. CPPM shows significant improvement over APPM on all scenes.

Scene	Box	Clocks	Conference	Cornell	Diamond	Dining	Glass	Sibenik	Sponza	Torus	Water
Overview											
MSE \leq	4.825	2.288	12.663	11.494	4.233	19.145	10.525	18.068	15.442	3.500	11.870
SPPM	Iters	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
	Time	128.685	78.802	108.082	232.403	120.371	166.856	371.223	123.665	70.701	205.176
APPM	Iters	765	655	538	763	678	961	778	788	823	556
	Time	101.415	48.391	62.782	195.307	81.137	162.215	300.411	88.289	56.283	204.589
CPPM	Iters	376	327	239	547	362	541	582	562	499	252
	Time	63.505	24.483	27.614	151.350	59.506	101.389	255.023	69.435	35.861	176.873




Fig. 8. Close-up images of the Torus scene for comparison. The heat map visualizes the absolute difference w.r.t. the reference image.

Jensen 2009]. The optimal setting $\alpha = 2/3$ is used for SPPM [Kaplanian and Dachsbaeh 2013]. In our implementation, we use $n_a = 2$, $n_s = 6$, $k = 0.8$ and $\beta = 1.2$ for CPPM unless otherwise stated. The same number of photons (2^{16}) per iteration is used for all the methods.

9.2 Results

Performance comparisons between different methods for all the benchmarks are listed in alphabetical order in Table 1. The MSE value obtained after running SPPM with 1,000 iterations is used as the baseline. Then we record the number of iterations and corresponding time cost (in seconds) required to reach that MSE. These benchmarks include scenes with caustics (Box and Torus), scenes with glossy surfaces (Cornell Box), more challenging scenes with complicated caustics (Diamond, Glass, and Water), and other scenes of different scales without caustics. Overall, CPPM outperforms prior algorithms on all the scenes with the fewest iterations and the lowest time costs. We also show their plots over 1,000 iterations in the supplemental material.

We show the difference between each algorithm in Figure 8. In a local region, SPPM has noticeable noise and severe blur problems at discontinuities; APPM has some visual noise near discontinuities, as pointed out by Kaplanian and Dachsbaeh [2013]; CPPM has less blur or noise.

To evaluate CPPM's performance, we plot the bandwidth and absolute error over iterations in Figure 9, picking two points, A and B (illustrated in Figure 8), as representatives. We provide these algorithms with the same photon maps for fair comparisons. Point

A is a typical point with smooth illumination and Point B is at the boundary of a caustic.

At Point A (see Figure 9 (a)(b)), where the initial bandwidth brings lower bias, CPPM's bandwidth is fixed for 1,000 iterations and is generally larger than the bandwidth of SPPM and APPM. It is apt to collect more photons, so CPPM's error quickly converges to 0. APPM's bandwidth varies drastically, which may introduce extra bias and variance. This is because its bandwidth is obtained by estimating bias and variance, which may have some fluctuations.

At Point B (see Figure 9 (c)(d)), where the initial bandwidth brings severe bias, CPPM quickly reduces its bandwidth to find a distribution close to uniform distribution using our bandwidth selection scheme. The bandwidth reduces like a step function. APPM also quickly reduces its bandwidth to lower the bias. However, many approximations used in computing the optimal bandwidth make APPM underestimate the bias.

These diagrams exhibit different behaviors of CPPM, SPPM, and APPM in terms of bandwidth reduction. CPPM's bandwidth behavior exhibits a more robust evaluation of the local photon distribution than APPM. The error caused by CPPM during these iterations tends to be consistently less than that generated by SPPM or APPM, regardless of whether the bandwidth reduces or not. This also indicates that CPPM's faster convergence to a correct value is due to its adaptive bandwidth reduction. In fact, CPPM can obtain a bandwidth consistent with the null hypothesis under different conditions. This strategy helps to reduce the bias and variance to achieve better performance compared to other methods. We also provide bandwidth visualizations of different methods in the supplementary video.

We design an artware scene including a translucent object, a glossy material, a mirror, and complex caustics for comparison, as shown in Figure 10. CPPM demonstrates better results than APPM. We provide more results with comparisons between different methods to show the advantages of our method in the supplementary material.

9.3 Parameters and Discussions

We investigate the reliability of CPPM by tuning the parameters used in the experiments that may affect the performance.

• Initial Bandwidth k_{NN}

Generally, the initial bandwidth determined by k_{NN} is set to 10 in prior progressive methods. We vary this value on the Cornell Box




Fig. 9. Visualization of bandwidth and error at two representative points: A and B. These diagrams are in log-log plot except (b). The highlighted bandwidth (green) of CPPM in (c) indicates that the chi-squared test has been performed at that iteration. CPPM exhibits better performance than SPPM and APPM.




Fig. 10. Artware scene rendered by our approach (CPPM) and APPM [Kaplanyan and Dachsbacher 2013] (10,000 iterations). We illustrate zoomed-in results with absolute error visualization. The overall MSE of APPM and CPPM are 1.670 and 0.840, respectively. CPPM exhibits better rendering quality.

(simple flat surfaces with smooth radiance) and the Water (complex caustic effects) scene and show the effect of k_{NN} in Figure 11. For further comparisons, we also use a global bandwidth (10^{-1} scene scale) for APPM, which is also used in [Kaplanyan and Dachsbacher 2013]. We can see that CPPM always outperforms APPM using different k_{NN} on different types of scenes, which demonstrates the reliability of our method regardless of the variation in the initial value. In fact, the plots of APPM using different initial bandwidth almost coincident in (b).

The optimal initial bandwidth is generally related to the type of scene. If a scene has smooth radiance like Cornell Box, there are many pixels where the photon distribution within a large bandwidth is consistent with the null hypothesis, then a larger initial bandwidth will help to collect more photons and reduce the variance. In contrast, in a scene like Water, which has more complicated radiance, the photon distribution within a large bandwidth has a high probability of being rejected by the test. Therefore, using a smaller initial bandwidth will help quickly find the bandwidth where the photon distribution is consistent with the null hypothesis, reducing the bias in some early iterations.

• Convergence Control k, β

As described in Subsection 5.3 and analyzed in Section 8, β and k theoretically affect the convergence of CPPM. We vary values of β and k and show the results in Table 2, where the first, the fourth, and the last group of parameters satisfy $\beta \approx 1/k^2$.

If the null hypothesis is always rejected, CPPM will degenerate to SPPM, and using $\beta = 1/k^2$ should have the best performance according to Section 8. If $\beta < 1/k^2$, the bandwidth reduces more




Fig. 11. Various k_{NN} for CPPM over iterations on the Cornell Box (no caustics) and the Water (caustics) scenes. APPM-GR represents a 10^{-1} scene scale is used as the initial bandwidth.

Table 2. Different k and β for the Sibenik (no caustics) and the Diamond and Torus (caustics) scenes after 1,000 iterations.

k, β		0.7, 2.0	0.8, 1.0	0.8, 1.2	0.8, 1.6	0.9, 1.2
Sibenik	MSE	15.929	14.887	14.907	15.445	15.074
	Time	110.722	108.216	110.969	108.582	110.500
Diamond	MSE	2.907	2.347	2.350	2.692	2.467
	Time	125.911	127.206	126.063	126.989	132.470
Torus	MSE	1.758	1.452	1.435	1.572	1.446
	Time	42.096	41.566	41.482	41.752	42.365

aggressively and is apt to quickly find a desired bandwidth consistent with the null hypothesis. However, if the null hypothesis is always rejected, aggressively reducing the bandwidth will result in high variance. In Table 2, we can see that $\beta < 1/k^2$ have better performance, which implies that most pixels can achieve a bandwidth where the null hypothesis is not rejected by the chi-squared

Table 3. Water scene rendered using different n_a and n_s by CPPM. $n_a = 2$ and $n_s = 6$ achieves the best performance.

MSE	$n_a = 1$	$n_a = 2$	$n_a = 3$
$n_s = 4$	12.315	10.326	10.555
$n_s = 6$	12.006	10.237	10.385
$n_s = 8$	11.981	10.248	10.448
$n_s = 10$	12.051	10.271	10.452

Table 4. Cornell Box and Sponza rendered with different α_χ with 1,000 iterations.

MSE	$\alpha_\chi = 0.01$	$\alpha_\chi = 0.05$	$\alpha_\chi = 0.10$
Cornell	0.871	1.095	1.651
Sponza	11.371	10.359	10.842

test. Even if $\beta = 1.0$, where the bandwidth reduction is extremely aggressive and the convergence cannot be guaranteed, CPPM still exhibits a very good performance. This indicates that our bandwidth preserving strategy works. We conservatively use $k = 0.8$ and $\beta = 1.2$ to guarantee MSE convergence in the benchmarks.

• Domain Partition n_a, n_s

A larger n_a can provide more candidates for bandwidth selection, but when the chi-squared test is not reliable enough, the bandwidth may reduce inaccurately and thereby result in a noisy image. A larger n_s can reduce the impact of the parametric partition, but it increases the number of sectors and the sample size should also be increased accordingly, otherwise the result from the chi-squared test may become unreliable, also resulting in a noisy image. We list the MSE on the Water scene using different n_a and n_s , as shown in Table 3. We also get similar results on other scenes, which shows $(n_a, n_s) = (2, 6)$ is the best setting for CPPM. Using $n_a = 1$ always produce worse results, which shows the benefits of our bandwidth selection (based on $n_a \geq 2$) described in Subsection 6.1.

• Significance Level α_χ

From a statistical view, when a smaller α_χ is used, the chi-squared test is more tolerant of the null hypothesis and is less likely to reject the null hypothesis; when a larger α_χ is used, the chi-square test is stricter and more likely to reject the null hypothesis. In other words, using a small α_χ can reduce the probability of incorrectly rejecting the null hypothesis when the null hypothesis is true; using a large α_χ can reduce the probability of not rejecting the null hypothesis when the null hypothesis is false. We vary α_χ for different type of scenes and show the results in Table 4.

For scenes (like the Cornell Box) with overall smooth radiance, where a vast majority of pixels have large bandwidths consistent with the null hypothesis, a small α_χ is apt to reduce the probability of rejecting a true null hypothesis and then maintain the bandwidth. So the Cornell Box scene gets the best performance using $\alpha_\chi = 0.01$. For scenes (like the Sponza) with relatively complex radiance, where many pixels may not have large bandwidths consistent with the null hypothesis when the bandwidth is large, a relatively large α_χ is apt to reduce the probability of not rejecting a false null hypothesis and then reduce the bandwidth. Therefore, the Sponza scene gets the best performance using $\alpha_\chi = 0.05$. Overall, $\alpha_\chi = 0.01 \sim 0.05$ is a good trade-off in practice.

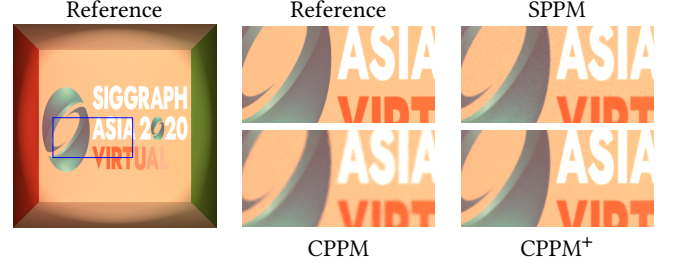


Fig. 12. A special case using a textured spotlight, where the location-independent assumption is violated and CPPM exhibits blurrier results than SPPM. CPPM⁺ improves the quality by specific processing based on CPPM.

9.4 Discussion

It is possible that the chi-squared test fails to reject the null hypothesis for several iterations when the null hypothesis is false. At the point B in Figure 8, the estimation is always biased in the first 1,000 iterations of CPPM, as shown in Figure 9(d). We can see in Figure 9(c) that the chi-squared test is performed in many iterations. In most of these iterations, it does not reject the null hypothesis and decides not to reduce the bandwidth, which implies that the null hypothesis is not correctly rejected. However, CPPM continues to collect photons and check the photon distribution even though the test in the last iteration failed to reject the null hypothesis; the null hypothesis is rejected later and the bandwidth continues to reduce. This shows that our method is robust in this case.

On the other hand, it is also possible that the null hypothesis is true while the estimation is still biased due to a violation of the location-independent assumption. Spatially-varying BRDF, textured spotlight, complex geometry with occlusions, or other special cases may cause the violation of assumption. Under such conditions, the contributions of the constructed paths are dependent on their locations even though the photons are uniformly distributed. Taking textured spotlight as an example, CPPM suspends the bandwidth reduction incorrectly only based on the uniformity of the photons while wrongly treating the contribution as identical and performs biased estimation with an over blurred result, as shown in Figure 12.

To deal with this case, an add-on to correct the photon contribution and make it identical would make CPPM workable. This add-on combines several techniques together including rendering RGB channels separately, emitting the light path proportional to the light source luminance, and bouncing the light path by Russian Roulette. Using this specific processing in CPPM, we can tackle the assumption-violated problem to some extent, the new result (CPPM⁺) is shown in Figure 12 with better quality. In addition, a diffuse textured surface, which is a particular case of spatially-varying BRDF, we believe that using the hit point's BRDF for all photons can handle this situation. However, these assumption-violated cases generally need specific case-by-case processing with careful implementation, and these processes may consume more time overhead. Theoretically, identical photon contribution is not strictly equivalent to identical path contribution, there may be some situations that make difference between them. Therefore, the proposed add-on can not cover all these special cases.

We have evaluated CPPM in various scenes that include caustics, glossy surfaces, complex geometries, and different kinds of light sources. The experimental results suggest that extreme cases rarely occur in general settings and CPPM outperforms prior methods.

10 CONCLUSION, LIMITATION, AND FUTURE WORK

We present an algorithm that consistently evaluates the distribution of photons based on the chi-squared test and a novel bandwidth reduction strategy. It suspends bandwidth reduction if the test does not reject the null hypothesis and deems that the distribution within this bandwidth permits unbiased estimation. We analyze the asymptotic average error of our CPPM algorithm, which degrades to SPPM in the worst case. We demonstrate through experiments that CPPM performs better than SPPM and APPM. In practice, our approach can be easily implemented in any progressive framework. In our framework, other hypothetical models such as linear functions or generalized odd functions can also be used instead of the constant function; we provide more results in supplementary material.

However, CPPM has some limitations. For example, CPPM can be more sensitive to bias than APPM, meaning that it reduces the bandwidth in a more dramatic way in some local regions, including wall corners and the edges of shadows, resulting in more noise (variance) than APPM. The overall performance of CPPM is better, but APPM may get better performance than CPPM in some local regions. It would be an interesting future work to combine APPM and CPPM in a hybrid manner, e.g., automatically choosing one of them for each pixel. If the null hypothesis is always false and being rejected, the convergence rate will degenerate to the same as that in SPPM. Fortunately, CPPM performs well when the actual photon distribution is slightly different from the uniform distribution. Although there always needs to be specific processing for different assumption-violated cases, as discussed in Subsection 9.4, a better solution is to design a hypothetical model that takes the contribution of photons at different location into account, which we leave for future work. CPPM's bandwidth cannot become larger than the initial bandwidth. It would be interesting to modify the bandwidth reduction scheme of CPPM to allow the bandwidth to grow. Finally, CPPM has the potential to work in combination with other sampling-based techniques like VCM [Georgiev et al. 2012] / UPS [Hachisuka et al. 2012] for realistic rendering. However, a method for integrating them seamlessly needs further investigation.

There are many areas for future work in addition to overcoming the limitations. Our CPPM is designed for on-surface estimation. Applying CPPM to volumetric radiance estimation requires modifying the partition scheme, i.e. making a spatial partition instead of a planar one. The optimal bandwidth for volumetric radiance estimation results in an MSE convergence rate $O(N^{-4/7})$ [Kaplanyan and Dachsbacher 2013] that is slower than on-surface estimation. Finding a low bias bandwidth in this context would be an interesting future work. With hypothetical testing, finding some solutions to reduce the bias and variance of the plug-in variables in APPM is also an interesting idea.

ACKNOWLEDGEMENTS

We would like to thank all the anonymous reviewers for helpful suggestions. This project was supported by the National Key R&D Program of China (No.2017YFB0203002, No.2017YFB1002700) and NSFC of China (No.61632003, No.6166114600).

APPENDIX

A Unbiased Pixel Measurement Estimation

According to [Kaplanyan and Dachsbacher 2013],

$$I = \int_{\mathcal{M} \times \mathcal{M}} \delta(\vec{x} - \vec{x}') \Psi(\vec{x}, \vec{x}') d\vec{x} d\vec{x}' \quad (20)$$

where δ is the Dirac delta function, and $\Psi(\vec{x}, \vec{x}')$ is the contribution of all possible full paths constructed by the hit point \vec{x} and photon \vec{x}' . Expressing the Dirac delta function in Eq. (20) with limiting, we obtain

$$I = \lim_{r \rightarrow 0} \int_{\mathcal{M} \times \mathcal{M}} k_r(\vec{x} - \vec{x}') \Psi(\vec{x}, \vec{x}') d\vec{x} d\vec{x}'. \quad (21)$$

Expressing the integral in Eq. (21) with expectation of sampling hit point \vec{x}_i , we have

$$I = \lim_{r \rightarrow 0} \mathbb{E} \left[\int_{\mathcal{M}} k_r(\vec{x}_i - \vec{x}') p(\vec{x}') \Psi(\vec{x}_i, \vec{x}') d\vec{x}' \right]. \quad (22)$$

We then transform \mathcal{M} to Ω_r and substitute p and Ψ by Eqs. (5) and (6); Eq. (22) becomes

$$I = \lim_{r \rightarrow 0} \mathbb{E} \left[\int_{\Omega_r} k_r(\vec{x}') \tilde{p}_i(\vec{x}') \tilde{\Psi}_i(\vec{x}') d\vec{x}' \right]. \quad (23)$$

We assume that $\tilde{\Psi}_i$ is a constant function. Therefore, we can factor the expectation into a product of expectations:

$$I = \lim_{r \rightarrow 0} \mathbb{E} \left[\int_{\Omega_r} k_r(\vec{x}') \tilde{p}_i(\vec{x}') d\vec{x}' \right] \mathbb{E} [\tilde{\Psi}]. \quad (24)$$

Substitute the first expectation of Eq. (24) by the expectation of the average of N samples and apply Fubini's theorem; we then have

$$I = \lim_{r \rightarrow 0} \mathbb{E} \left[\int_{\Omega_r} k_r(\vec{x}') \frac{1}{N} \sum_{i=1}^N \tilde{p}_i(\vec{x}') d\vec{x}' \right] \mathbb{E} [\tilde{\Psi}]. \quad (25)$$

When APDF $\bar{p}(\vec{x}') = \frac{1}{N} \sum_{i=1}^N p_i(\vec{x}')$ is a constant function within bandwidth R , we can extract it outside of the integral as

$$I = \lim_{r \rightarrow 0} \mathbb{E} \left[\int_{\Omega_r} k_r(\vec{x}') d\vec{x}' \right] \mathbb{E} [\bar{p}] \mathbb{E} [\tilde{\Psi}], \quad (26)$$

where the first expectation equals one. Therefore, we can substitute the bandwidth r with R , remove the limit symbol, and go back along with the above deductions to Eq. (23). Then, we will get Eq. (8), which implies the pixel measurement estimation is unbiased.

B Photon Projection and Alignment

We calculate the corresponding 2D coordinate of the j -th photon in the searching area of the i -th hit point as:

$$\vec{p}'_{i,j} = \frac{\|\vec{p}_{i,j} - \vec{x}_i\| \cdot (\langle \vec{p}_{i,j} - \vec{x}_i, \vec{u}_i \rangle, \langle \vec{p}_{i,j} - \vec{x}_i, \vec{v}_i \rangle)}{\|(\langle \vec{p}_{i,j} - \vec{x}_i, \vec{u}_i \rangle, \langle \vec{p}_{i,j} - \vec{x}_i, \vec{v}_i \rangle)\|}, \quad (27)$$

where \vec{u}_i and \vec{v}_i are two orthonormal bases vectors on the tangent plane at hit point \vec{x}_i ; $\langle \cdot, \cdot \rangle$ indicates inner product; and $\|\cdot\|$ indicates Euclidean length.

C Bandwidth Convergence Analysis

We give an analysis of the bandwidth convergence when the chi-squared test always rejects the null hypothesis; we then show that the bandwidth convergence in CPPM will degenerate to the same as that in SPPM in this case.

Assume the duration between two successive bandwidth reductions is P_i . When the chi-squared test always rejects the null hypothesis, from Eq. (13), we can derive that P_i is proportional to the required number of photons B_i , and inversely proportional to the area of the searching area. Therefore, we can obtain P_i as

$$P_i = \left(\frac{\beta}{k}\right)^{i-1} P_1. \quad (28)$$

For the N hit points in the first N iterations, let t be the number of different bandwidths. We can infer bounds of N by P_i as

$$\sum_{i=1}^{t-1} P_i \leq N \leq \sum_{i=1}^t P_i. \quad (29)$$

Substitute P_i with Eq. (28) and we have the relation between t and N :

$$\sum_{i=1}^{t-1} \left(\frac{\beta}{k}\right)^{i-1} P_1 \leq N \leq \sum_{i=1}^t \left(\frac{\beta}{k}\right)^{i-1} P_1 \Rightarrow \left(\frac{\beta}{k}\right)^t = O(N). \quad (30)$$

Since R_N is the t -th bandwidth, based on Eq. (30), we assess the bandwidth convergence as

$$R_N = k^{\frac{1}{2}}(t-1) R_1 = O\left(N^{-\frac{1}{2} \log_{\beta} \frac{1}{k}}\right). \quad (31)$$

REFERENCES

- VB Bagdonavicius and MS Nikulin. 2011. Chi-squared goodness-of-fit test for right censored data. *International Journal of Applied Mathematics and Statistics* 24, SI-11A (2011), 30–50.
- Jiating Chen, Bin Wang, and Jun-Hai Yong. 2011. Improved stochastic progressive photon mapping with metropolis sampling. *Computer Graphics Forum* 30, 4 (2011), 1205–1213.
- William G. Cochran. 1952. The Chi-square Test of Goodness of Fit. *The Annals of Mathematical Statistics* 23, 3 (1952), 315–345.
- Philip Dutre, Philippe Bekaert, and Kavita Bala. 2006. *Advanced global illumination*. A K Peters Ltd.
- Zhe Fu and Henrik Wann Jensen. 2012. Noise reduction for progressive photon mapping. In *ACM SIGGRAPH 2012 Talks*. ACM Siggraph, 29.
- Iliyan Georgiev, Jaroslav Krivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (2012), 192–1.
- Pascal Grittmann, Arsène Pérard-Gayot, Philipp Slusallek, and Jaroslav Krivánek. 2018. Efficient Caustic Rendering with Lightweight Photon Mapping. *Computer Graphics Forum* 37, 4 (2018), 133–142.
- Adrien Gruson, Mickaël Ribardière, Martin Šik, Jiří Vorba, Rémi Cozot, Kadi Bouatouch, and Jaroslav Krivánek. 2016. A Spatial Target Function for Metropolis Photon Tracing. *ACM Transactions on Graphics (TOG)* 36, 1 (2016), 4.
- Tobias Günther and Thorsten Grosch. 2014. Distributed Out-of-Core Stochastic Progressive Photon Mapping. *Computer Graphics Forum* 33, 6 (2014), 154–166.
- László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. 2006. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media.
- Toshiya Hachisuka, Wojciech Jarosz, Iliyan Georgiev, Anton Kaplanyan, Derek Nowrouzezahrai, and Ben Spencer. 2013. State of the art in photon density estimation. In *SIGGRAPH Asia 2013 Courses*. ACM Siggraph, 15.
- Toshiya Hachisuka, Wojciech Jarosz, and Henrik Wann Jensen. 2010. A progressive error estimation framework for photon density estimation. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 144.
- Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic progressive photon mapping. *ACM Transactions on Graphics (TOG)* 28, 5 (2009), 141.
- Toshiya Hachisuka and Henrik Wann Jensen. 2011. Robust adaptive photon tracing using photon path visibility. *ACM Transactions on Graphics (TOG)* 30, 5 (2011), 114.
- Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive photon mapping. *ACM Transactions on Graphics (TOG)* 27, 5 (2008), 130.
- Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A path space extension for robust light transport simulation. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 191.
- Vlastimil Havran, Jiri Bittner, Robert Herzog, and Hans-Peter Seidel. 2005. Ray maps for global illumination. In *Proceedings of the Sixteenth Eurographics conference on Rendering Techniques*. Eurographics Association, 43–54.
- Rubén Jesus García Hernández, Carlos Urena, Jordi Poch, and Mateu Sbert. 2014. Overestimation and underestimation biases in photon mapping with non-constant kernels. *IEEE transactions on visualization and computer graphics* 20, 10 (2014), 1441–1450.
- Robert Herzog, Vlastimil Havran, Shinichi Kinuwaki, Karol Myszkowski, and Hans-Peter Seidel. 2007. Global illumination using photon ray splatting. *Computer Graphics Forum* 26, 3 (2007), 503–513.
- Heinrich Hey and Werner Purgathofer. 2002. Advanced radiance estimation for photon map global illumination. *Computer Graphics Forum* 21, 3 (2002), 541–545.
- Wenzel Jakob. 2010. Mitsuba renderer.
- Wenzel Jakob, Christian Regg, and Wojciech Jarosz. 2011. Progressive Expectation–Maximization for Hierarchical Volumetric Photon Mapping. *Computer Graphics Forum (Proceedings of EGSR)* 30, 4 (June 2011). <https://doi.org/10/dtwcjj>
- Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A Framework for Transient Rendering. *ACM Transactions on Graphics (SIGGRAPH Asia 2014)* 33, 6, Article 177 (2014).
- Wojciech Jarosz, Derek Nowrouzezahrai, Robert Thomas, Peter-Pike Sloan, and Matthias Zwicker. 2011. Progressive photon beams. *ACM Transactions on Graphics (TOG)* 30, 6 (2011), 181.
- Henrik Wann Jensen. 1996. Global illumination using photon maps. In *Rendering Techniques '96*. Springer, 21–30.
- Henrik Wann Jensen. 2001. *Realistic image synthesis using photon mapping*. Vol. 364. Ak Peters Natick.
- Daniel Jönsson and Anders Ynnerman. 2017. Correlated Photon Mapping for Interactive Global Illumination of Time-Varying Volumetric Data. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 901–910.
- Anton S Kaplanyan and Carsten Dachsbacher. 2013. Adaptive progressive photon mapping. *ACM Transactions on Graphics (TOG)* 32, 2 (2013), 16.
- Claude Knaus and Matthias Zwicker. 2011. Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics (TOG)* 30, 3 (2011), 25.
- Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. 2010. Optix: a general purpose ray tracing engine. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 66.
- Hao Qin, Xin Sun, Qiming Hou, Baining Guo, and Kun Zhou. 2015. Unbiased photon gathering for light transport simulation. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 208.
- Lars Schjøth. 2009. *Anisotropic density estimation in global illumination: a journey through time and space*. Ph.D. Dissertation. Ph. D. thesis, University of Copenhagen.
- Roland Schrengle. 2003. Bias compensation for photon maps. *Computer Graphics Forum* 22, 4 (2003), 729–742.
- Martin Šik, Hisanari Otsu, Toshiya Hachisuka, and Jaroslav Krivánek. 2016. Robust light transport simulation via metropolised bidirectional estimators. *ACM Trans. Graph.* 35, 6 (2016), 245.
- Ben Spencer and Mark W Jones. 2009a. Hierarchical photon mapping. *IEEE Transactions on visualization and computer graphics* 15, 1 (2009), 49–61.
- Ben Spencer and Mark W Jones. 2009b. Into the blue: Better caustics through photon relaxation. *Computer Graphics Forum* 28, 2 (2009), 319–328.
- Ben Spencer and Mark W Jones. 2013a. Photon parameterisation for robust relaxation constraints. *Computer Graphics Forum* 32, 2pt1 (2013), 83–92.
- Ben Spencer and Mark W Jones. 2013b. Progressive photon relaxation. *ACM Transactions on Graphics (TOG)* 32, 1 (2013), 7.
- Michael A Stephens. 1974. EDF statistics for goodness of fit and some comparisons. *Journal of the American statistical Association* 69, 347 (1974), 730–737.
- Rui Wang, Rui Wang, Kun Zhou, Minghao Pan, and Hujun Bao. 2009. An efficient GPU-based approach for interactive global illumination. *ACM Transactions on Graphics (TOG)* 28, 3 (2009), 91.
- Maayan Weiss and Thorsten Grosch. 2012. Stochastic progressive photon mapping for dynamic scenes. *Computer Graphics Forum* 31, 2pt3 (2012), 719–726.
- Shilin Zhu, Zexiang Xu, Henrik Wann Jensen, Hao Su, and Ravi Ramamoorthi. 2020. Deep Kernel Density Estimation for Photon Mapping. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 35–45.