

Into the Blue: Better Caustics through Photon Relaxation

B. Spencer and M. W. Jones

Visual and Interactive Computing Group, Swansea University, UK

Abstract

The photon mapping method is one of the most popular algorithms employed in computer graphics today. However, obtaining good results is dependent on several variables including kernel shape and bandwidth, as well as the properties of the initial photon distribution. While the photon density estimation problem has been the target of extensive research, most algorithms focus on new methods of optimising the kernel to minimise noise and bias. In this paper we break from convention and propose a new approach that directly redistributes the underlying photons. We show that by relaxing the initial distribution into one with a blue noise spectral signature we can dramatically reduce background noise, particularly in areas of uniform illumination. In addition, we propose an efficient heuristic to detect and preserve features and discontinuities. We then go on to demonstrate how reconfiguration also permits the use of very low bandwidth kernels, greatly improving render times whilst reducing bias.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism - Global Illumination—

1. Introduction

Photon mapping [Jen96a] is one of the most effective methods for synthesising and rendering realistic caustics. The decoupling of radiant flux from scene geometry makes the technique highly suited to computing illumination over a variety of surfaces and BRDFs. In addition, view-independence and temporal coherency make it ideal for animation. By far the most successful method of computing incident radiance from the photon map has been that of k-nearest neighbour (k-NN) density estimation. In order to be effective, however, this approach requires careful tuning so as to minimise the effects of signal noise and bias. At the most fundamental level, noise and bias are opposing quantities on a continuum defined by kernel bandwidth. By increasing the bandwidth, high-frequency components of the signal are filtered out while error in the form of bias is introduced. For both relatively high and relatively low bandwidths, visually objectionable artefacts begin to appear, evident as bias and high-frequency noise respectively.

In order to reduce errors in the radiance estimate, a large body of prior research has focussed on optimal bandwidth selection and kernel filtering. Broadly stated, the problem is one of signal noise reduction and is well-known throughout

computer science and engineering. Many methods rely on adapting the size and shape of the kernel so as to achieve better results.

The main contributions of this paper are a novel approach to noise and bias reduction based upon reconfiguration of the photon map, and a novel heuristic to rapidly identify and preserve edges and discontinuities. Our new approach relaxes the photons into a configuration with a blue noise spectral signature. This allows for the use of compact kernels with only tens of k-nearest neighbours (k-NN). In addition we apply a heuristic in order to preserve important features within the map. This dual strategy has several important advantages (figure 1):

- A blue noise distribution yields high-quality radiance estimates with very few photons. This greatly improves render time (by an average factor of 5 in our test scenes) and offsets the precomputation required for relaxation.
- Topology, proximity and boundary bias are all minimised by the compact kernel size.
- Edges and discontinuities are effectively preserved using photon migration constraints.

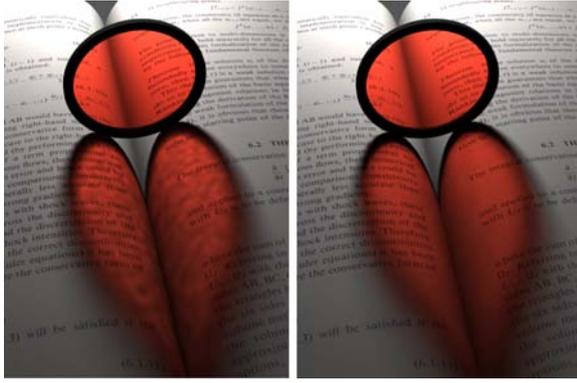


Figure 1: A caustic from a red camera lens filter projected onto a textbook. Upper frame: using an unrelaxed photon map rendered with 200 photons in the radiance estimate. Lower frame: relaxed by 20 iterations and rendered with 20 photons in the radiance estimate. Image inspired by Yusuke Okaue.

2. Previous Work

In this section we review research carried out into the three distinct areas of computer rendering applicable to the scope of this paper. First, we list developments made on bias and noise reduction in photon density estimation. We then go on to cover work published in the field of sampling and its link to photon distribution. Finally, we briefly examine relevant literature in the field of point repulsion.

2.1. Bias and Noise Reduction in Photon Density Estimation

The naïve approach to k -nearest neighbour (k -NN) density estimation involves solving the problem for a given point within the photon map. The irradiance at point x can be estimated as the sum of the power of the k nearest photons to x , divided by the area of the disc enclosing the points. Since the introduction of photon density estimation in global illumination, expansions to the algorithm have largely been divided between two categories: selecting an optimum kernel bandwidth and photon weighting and filtering.

Jensen and Christensen [JC95] proposed differential checking as a method of selecting an optimal bandwidth. Given an initial estimate, the search radius is expanded incrementally by adding progressively more neighbouring photons. If the irradiance differential between the estimates exceeds some given threshold then the search is halted. This approach maximises the search radius and hence minimises noise, while at the same time avoiding penalties from bias.

Schregle [Sch03] proposed a bias compensation operator by using a binary search to hone in upon the optimal bandwidth. An estimate of the expected value of the combined

bias and variance for the current bandwidth is calculated, and the variance contribution is differentiated from the bias by using knowledge that it matches a Gaussian distribution. This bias estimate can govern the binary search, and thus the optimum bandwidth can be discovered. The method is effective; adapting bandwidth across an image in order to preserve both peak irradiance and sharp edges in irradiance, at the same time allowing higher bandwidths, and therefore reduced variance, in smoothly varying regions.

Hey and Purgathofer [HP02] approach the problems of boundary and topological bias by using the underlying model geometry. During radiance estimation, polygons in the nearest neighbourhood to the sample position are determined, their area is calculated and photon power is distributed according to their incoming direction. Although very effective, Hey's method is restrictive in that it extensively refers to the surface polygonal mesh.

Myszkowski [Mys97] introduced a bias-reduction strategy for the kernel-based density estimation framework established by Shirley *et al* [SWH*95]. Myszkowski proposed an enhanced nearest-neighbour (ENN) method guided by a statistical bias-error metric. By calculating the error from an array of density estimates adjacent to the sample point, the ENN algorithm reduces bias by minimisation across the domain of sample radii.

Walter *et al* [WHS97] addressed the issue of bias within the density estimation framework using polynomial regression. Given that boundary bias is also an issue with regression techniques, Walter used a locally-weighted, least squares variant augmented with a complex system to handle boundary regions. Introduced in the paper and further elaborated upon in his PhD thesis [Wal98], Walter also proposed a perceptually-driven bandwidth selector that chooses kernel sizes based upon the limits of visually objectionable noise.

A second class of bias-reduction algorithms focusses on adapting the filter support of the kernel to help reduce unwanted blurring. Isotropic kernels are used in both k -NN- and kernel-based methods of density estimation as well as with techniques based on splatting [Col94]. Jensen [Jen96b] first integrated the concept into the photon mapping paradigm primarily as a means of reducing proximity bias in caustics.

Later, Schjøth *et al* [SOS06] introduced a technique based around anisotropic diffusion filtering. Using a structure tensor derived from the photon distribution, they are able to define a shape-adapting kernel aligned to important visual details such as boundaries and discontinuities. The results are shown to be superior to isotropic kernels at the expense of a small precomputation and rendering overhead.

More recently, Schjøth *et al* [SFES07] proposed adapting ray differential tracing [Ige99] to guide the shape of the density estimation kernel. At render time, ray differential information is used to create a unique kernel based upon the

shape of the footprints of the photons within the gather radius. This approach is especially effective at rendering high-frequency caustics. The authors also demonstrate the power of the technique in realising fine details using only a small number of photons.

Spencer *et al* [SJ09] use a hierarchical data structure to cluster photon flux data, thus allowing estimate areas of an arbitrary size at a near-constant query cost. Relative photon density per unit area is adjusted by controlling the depth to which the balanced kd-tree is traversed. A useful side-effect of this approach is that of better photon stratification at shallower cut-off depths. The result is that noise is rapidly reduced as the kernel radius is increased.

2.2. Sample Distribution

The problem of optimal sampling within a given domain is of great significance in many areas of computer graphics. Monte Carlo integration, dithering, importance sampling and a host of other techniques all benefit from well-contrived sample distributions. Integrators which rely upon purely random functions to generate samples are known to converge more slowly than those which employ stratified or quasi-random methods. The quality of a sample set is sometimes measured by its discrepancy [Shi91] or the measure of equidistribution between points. Another useful tool for determining distribution quality is the Fourier transform. It is recognised that a set of points with a blue noise power spectrum and low angular anisotropy are most suitable for convolution operations such as anti-aliasing (see [DW85, Coo86]).

Optimal distributions are highly favourable in photon mapping since low photon discrepancy intrinsically means less noise in the radiance estimate. Jensen [JC95] proposed stratifying the photons as they are emitted. This can be accomplished by jittering, n-rooks, Poisson-disk and other sample schemes. Alternatively, a quasi-random sequence [Kel95] produces well-ordered sets with the added property of good stratification across multiple dimensions.

Point light sources benefit highly from good photon stratification at the emission phase. At a singularity, the domain over which the samples are distributed exists as the closed two-dimensional space $\mathbb{R}^2 = (\theta, \phi)$ where θ and ϕ are zenith and azimuth angles respectively. This can be mapped directly onto a planar geometric surface parameterised by $\mathbb{R}^2 = (u, v)$, allowing the low-discrepancy properties of the photon distribution to be effectively preserved. Unfortunately, such a mapping fails as extra dimensions are introduced. An area light source (for example, a quad) requires a 4-dimensional sampling domain $\mathbb{R}^4 = (u, v, \theta, \phi)$ resulting in an immersion of the distribution into two dimensions. Additionally, scattering from non-specular surfaces, arbitrary geometry and participating media results in rapid degeneration of the stratification into random noise.

2.3. Point Relaxation and Diffusion

The concept of diffusing a poorly-distributed point set into a more optimal configuration has been successfully adapted to solve many problems. McCool and Fiume [MF92] used Lloyd's method [Llo83] to generate correctly-spaced Poisson-disk samples. Ostromoukhov *et al* [ODJ04] also employed Lloyd's method to improve the spatial distribution of samples generated using a hierarchical method based on Penrose tiling. Kopf *et al* [KCODL06] used a similar scheme to enforce subdivision rules when generating blue noise.

Turk [Tur91] introduced a relaxation technique based on point repulsion to facilitate organic texture synthesis. This idea was later elaborated to define a method for re-tiling polygonal surfaces at arbitrary levels of detail [Tur92]. More recently, Jensen [JB02] utilised the same concept to ensure equidistribution of irradiance samples.

3. Our Method - Photon Relaxation

In section 2 we reviewed some of the prior research carried out into bias and noise reduction in photon density estimation. We also described how a good primal photon distribution cannot always guarantee a low-discrepancy sequence on intersecting geometry. In this section we outline our new approach to the problem based upon photon redistribution. By performing an additional pass after photon tracing, we aim to redistribute photons into an arrangement with a blue noise spectral signature. Our approach can be broken down into two distinct steps:

- Systematically search through the photon map for features and discontinuities, storing the inferred information in the photon data structure (section 3.2).
- Iteratively relax each photon according to a repulsion vector derived from the k-nearest neighbours. The data obtained during the previous step can then be used to constrain point migration and preserve important details (section 3.1).

3.1. Relaxation

In order to redistribute photons over the surface of intersecting geometry we employ a point repulsion method similar to that described in [Tur91]. This approach is simple and intuitive and does not require a Voronoi tessellation of the sample distribution.

For a given photon p at point \vec{x} , we calculate the force of repulsion \vec{f} as being the weighted sum of the offset k-nearest neighbours:

$$\vec{f} = \frac{1}{K} \sum_{k=1}^K (\vec{x} - \vec{x}_k) \left(\frac{r\tau(i)}{d_k + \epsilon} - \frac{d_k}{r\tau(i)} \right) \quad (1)$$

Here, d_k represents the distance from x_k to x . ϵ is an arbitrarily small constant. We choose a value of $K = 6$ given

that points relaxed using Lloyd’s method naturally converge to a hexagonal lattice-like distribution with 6 adjacent neighbours per point. r represents the radius of the disc enclosing the $K + 1$ neighbouring photons. This is necessary to prevent the K^{th} photon always lying on the periphery of the disc and thus having zero weight. $\tau(i)$ amplifies the value of r and hence the magnitude of \vec{f} . We call this the over-relaxation coefficient and define it as:

$$\tau(i) = \lambda_{min} + (\lambda_{max} - \lambda_{min}) \exp\left(-\frac{6i^2}{I^2}\right) \quad (2)$$

Where λ_{max} and λ_{min} represent the higher and lower relaxation bounds, blended together by a Gaussian falloff. Here, i is the current relaxation iteration and I the total number of iterations (see section 5.2 for more detailed discussion on the effects of varying λ).

One of the most notable advantages of photon mapping is its paradigm of decoupling irradiance from geometry. While there are established methods of tracking particle migration over surfaces [Tur92], it is simpler and more versatile to keep our algorithm geometry-independent. This can be achieved by projecting \vec{f} into the plane tangent to the surface at x . We derive an orthonormal basis from the surface normal \vec{n} and using it to define the new, projected vector, \vec{f}' :

$$\begin{aligned} \vec{u} &= \vec{n} \times (-n_x, n_y, -n_z) \\ \vec{v} &= \vec{n} \times \vec{u} \\ \vec{f}' &= \vec{u}(\vec{u} \cdot \vec{f}) + \vec{v}(\vec{v} \cdot \vec{f}) \end{aligned} \quad (3)$$

We apply the relaxation step to each photon in succession, computing the repulsion vector \vec{f}' then adding it to the position of the photon.

We find that the number of relaxation steps required to remove all-frequency noise from a sample distribution depends on the method used to cast photons. Purely random point distributions exhibit noise across the entire spectrum of spatial frequencies. In this case, 30 or more relaxation steps may be required to satisfactorily remove all objectionable noise. Conversely, quasi-random distributions confine noise to higher frequencies. In these instances we find that 20 relaxation steps are usually sufficient (see section 4 for examples).

3.2. Structure Preservation

Point repulsion can be effectively expressed as a diffusion operation. The direction of migration is a function of the derived gradient of photon density meaning particles naturally flow from more dense to less dense regions. In exchange for removing objectionable noise from the sample distribution, photon relaxation introduces a systematic error that we call

diffusion bias. This results in blurring and loss of high frequency detail and is especially apparent along edges and discontinuities.

To solve this problem it is necessary to acquire data on the structure of the photon map and use it to constrain movement relative to an axis perpendicular to the irradiance gradient. Schjøth *et al* [SOS06] use this approach to define a shape-adapting kernel based upon a derived structure tensor. While their solution could be adapted to constrain particle migration, relying purely on the distribution local to each photon does not give us enough precision. Figure 2 demonstrates the ghosting artefacts that appear when we use a purely gradient-based approach.

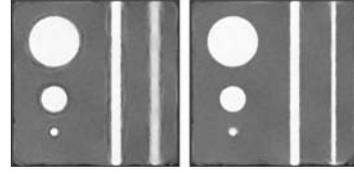


Figure 2: Comparison of edge detection methods. Left: using a gradient-based edge detection algorithm (note that a structure tensor-based variant would correct the unwanted diffusion in the right-most vertical bar). Right: our heuristic-based approach.

3.2.1. Feature Detection

To address these problems we introduce a novel method of controlling migration by assigning each photon a constraining vector \vec{g} and a weighting coefficient w . The vector \vec{g} defines an axis which lies in the plane of the photon. The weight w defines the extent to which the repulsive force \vec{f}' is constrained by \vec{g} . Thus, a maximally constrained photon with weight 1 can only migrate along \vec{g} . An unconstrained photon with weight 0 can move freely in its plane.

In order to find these values we begin by computing the irradiance differential at a given origin x_o from the distribution of the k -nearest neighbouring photons. The vector from x_o to the average photon position defines the gradient vector $\vec{\delta}$. We project this vector into the plane tangent to the surface (using the method described in equation 3) and then normalise it. If a discontinuity passes through the disc containing the k -NN then it is likely that it lies perpendicular to $\vec{\delta}$.

We compute the signed distance ψ_n from each photon n to the plane through the origin lying perpendicular to $\vec{\delta}$ (figure 3 (a)). The photons are then sorted in ascending order of ψ . Using this information we apply a heuristic, χ , that analyses each neighbour and assesses the uniformity of the distribution lying on each side:

$$\chi = \frac{\sigma_L}{\Psi_L} + \frac{\sigma_R}{\Psi_R} \quad (4)$$

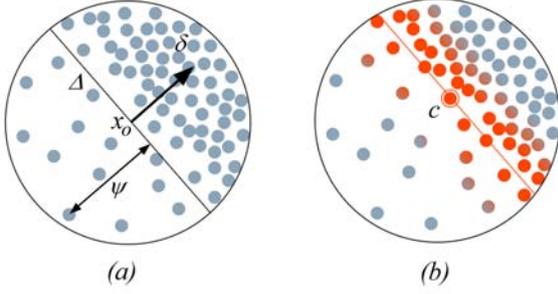


Figure 3: Feature detection. (a) The photon gradient δ is calculated and the perpendicular distance ψ to the plane Δ centred at the origin x_o is calculated for each photon. (b) The heuristic determines the photon c to have the minimum value of χ and hence to be an optimal candidate. Each photon in the estimate is constrained according to its proximity to the gradient plane at c .

The value of σ is derived from the standard deviation of the perpendicular distance between photons. We define it as:

$$\sigma = \sqrt{\frac{1}{N} \sum_{n=1}^N (\psi_n - \bar{\psi}_{n-1} - \bar{\psi})^2 \left(1 - \frac{\psi_n^2}{r^2}\right)} \quad (5)$$

The weighting function on the right-hand side of the equation helps to alleviate false positives which may arise from signal noise. $\bar{\psi}$ is the mean distance, defined as:

$$\bar{\psi} = \frac{1}{N} \sum_{n=1}^N (\psi_n - \psi_{n-1}) \quad (6)$$

Here, σ_L and σ_R are found by applying σ to the two subsets of K that lie to the left and to the right of a given photon, k . These ranges are defined as $L = [0, k)$ and $R = (k, K + 1]$. For the cases of ψ_0 and ψ_{K+1} , the extremes of the gather radius, $-r$ and r , are used respectively.

Our goal is to find the photon in K with the smallest corresponding value of χ , since any discontinuity is most likely to cross this point (figure 3 (b)). In practise it is not necessary to test all the photons, merely those at discrete intervals along $\bar{\delta}$ equivalent to the mean distance between the photons in the estimate: $\bar{\psi}_K \sqrt{\frac{K}{\pi}}$.

Once a suitable candidate discontinuity has been found we

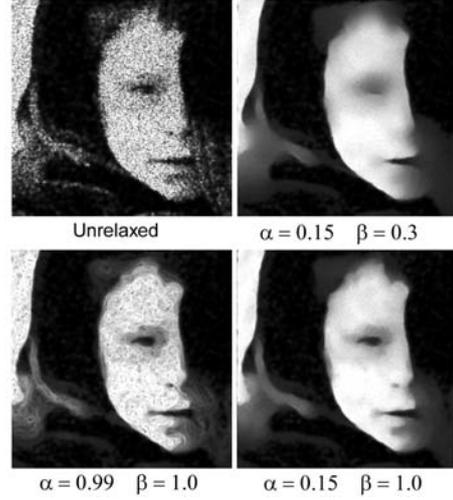


Figure 4: Diffusion of a pattern of photons with density proportional to the pixel intensity of an underlying guide image.

compute a homogeneity metric, ι , that represents the ratio between the means of the left and right partitions:

$$\iota(\bar{\psi}_L, \bar{\psi}_R) = \begin{cases} \frac{\min(\bar{\psi}_R, \bar{\psi}_L)}{\max(\bar{\psi}_R, \bar{\psi}_L)} & \text{when } \max(\bar{\psi}_R, \bar{\psi}_L) \neq 0 \\ 1 & \text{otherwise} \end{cases} \quad (7)$$

The ι function returns a value in the range $[0, 1]$ where 1 implies an entirely homogeneous distribution and 0 a maximal discontinuity. This value is used to determine the magnitude of the constraint to apply to each photon. We re-map a sub-range of ι between the user-specified limits α and β to tune the sensitivity of the heuristic:

$$\iota' = 1 - \frac{\iota - \alpha}{\beta - \alpha} \text{ where } \alpha < \beta \quad (8)$$

ι' is then clamped to the range $[0, 1]$. We found that for scenes containing high-frequency caustics (for example, figure 8), a value of α between 0.1 and 0.2 was optimal to highlight discontinuities. Conversely, in scenes containing low-frequency caustics (for example, figure 1), an α value of 0.0 was more appropriate. Figure 4 demonstrates the effects of different values of the two parameters on a sample distribution.

Whenever a photon p is found to lie within the feature

gradient (figure 3 (b)), we update its migration constraints as follows:

$$\begin{aligned}\vec{g}_p &= \vec{g}_p + \nu'(\vec{\delta} \times \vec{n}_o) \\ w_p &= \max(w_p, \nu')\end{aligned}\quad (9)$$

Where n_o is the surface normal at the origin x_o . We use the maximum of the two values of w_p and ν' because the weight of multiple gradients on each photon would most likely not average to 1, permitting unwanted migration across edges. Once the entire photon map has been evaluated we normalise all values of \vec{g} . For photons where w_p is zero, we derive an arbitrary value of \vec{g} from the surface normal using the first line of equation 3.

Given the normalised constraining vectors and associated weights, we can redefine equation 3 to modify the repulsive force \vec{f}_p' :

$$\begin{aligned}\vec{v}_p &= \vec{g}_p \times \vec{n}_p \\ \vec{f}_p'' &= \vec{v}_p(\vec{v}_p(1 - w_p) \cdot \vec{f}_p') + \vec{g}_p(\vec{g}_p \cdot \vec{f}_p')\end{aligned}\quad (10)$$

3.2.2. Searching the Photon Map

While the structure preservation heuristic outlined in the previous section is effective at constraining photons which lie on important features, the cost of applying the heuristic at every photon is relatively high. One of the benefits of iteratively sweeping through the k -nearest neighbours is that the heuristic can locate the most likely feature edge candidate regardless of where it lies in the local neighbourhood. Therefore, we need only ensure that every photon is captured and analysed at most once during feature detection. Choosing an optimum bandwidth for feature searching depends on the density and distribution of the photon map. We found that a value of $K = 150$ worked well in all our examples.

The most effective way of preventing unnecessary gradient searches is to mark photons captured within each estimate radius as "touched". Covering the map involves moving sequentially along a list of pointers to each of the photons. If a given photon is untouched, it is chosen as the center of a region to be searched. Otherwise, it is skipped. This approach effectively excludes photons that have already been swept and greatly decreases precomputation time. In our implementation we mark all photons within a distance of $0.6r$ from x_o as being touched at each search.

Unfortunately, the penalty of minimising the number of feature searches is that holes or weak migration constraints may appear along prominent discontinuities. This occurs when a discontinuity lies at the very edge of the search radius and is disregarded as noise by the heuristic. In addition, faceting artefacts may appear as a result of a straight gradient vector being used to constrain a curved edge. To address these problems we perform additional searches on photons

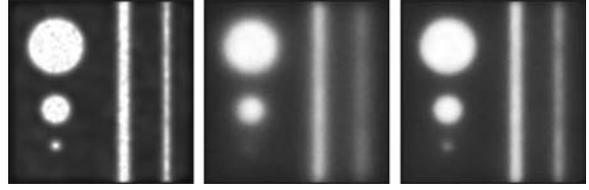


Figure 5: Diffusion bias: Left: unrelaxed test pattern rendered with 200 photons in the radiance estimate. Centre: relaxed by 30 iterations **without** diffusion limiting. Right: relaxed by 30 iterations **with** diffusion limiting. Note that the constraining vectors defined in section 3.2 are not applied in these examples.

lying along identified feature gradients (figure 3 (b)). This means that once a feature or discontinuity is detected, the local neighbourhood is explored laterally across the distribution gradient, further identifying and reinforcing it.

3.3. Limiting Diffusion in Unconstrained Photons

In section 3.2 we showed how the feature detection and migration constraints are effective at preserving sharp edges. However, blurring and loss of fidelity may still occur due to photon migration across steep gradients that do not necessarily qualify as discontinuities. Figure 5 (centre) demonstrates the effect of excessive diffusion on fine details.

To compensate, we introduce another constraint which applies a Gaussian falloff to the magnitude of each repulsion vector \vec{f}_p'' :

$$\vec{f}_p''' = \vec{f}_p'' \exp\left(-\left(\frac{\|\vec{f}_p''\|}{r_p}\right)^2 - \frac{\ln(0.5)}{\rho^2}\right)\quad (11)$$

The net effect is that the motions of photons with high migration pressure are damped while those with a low pressure are free to move. This results in the effective diffusion of low-frequency noise while still preserving higher-frequency details. In equation 11, ρ is a scaling parameter that controls the extent of the constraint. In all our examples, we set this parameter to 0.3. We found that using the value of r_p buffered from the previous relaxation iteration prevented unwanted artefacts being introduced into the distribution.

3.4. Flux Diffusion

One of the drawbacks of using very compact kernel sizes is the potential for noise to appear as a result of variance in chromaticity and intensity between photons. This may occur when caustics from two different coloured emitters mix or when a spectrally-based photon model is used (see [Wal98] for an example). We can address this problem by applying a

diffusion operation to the flux of the photons at each relaxation step.

Given a photon p , its new flux Φ'_p can be calculated as the weighted mean of itself and k -nearest adjacent neighbours.

$$\Phi'_p = \frac{\Phi_p + \sum_{k=1}^K \Phi_k W(k)}{1 + \sum_{k=1}^K W(k)} \quad (12)$$

Where $W(k)$ is an Epanechnikov kernel weighting function. Note that the migration constraining vectors of the photon k are used to prevent flux diffusion across discontinuities:

$$W(k) = \left(1 - \frac{d_k^2}{r^2}\right) (w_k \|\vec{g}_k \cdot (\vec{x}_p - \vec{x}_k)\| + (1 - w_k)) \quad (13)$$

4. Results

To test the efficacy of our algorithm we compare images of various scenes rendered using photon relaxation and standard k -NN, static bandwidth photon mapping. In all cases we use an isotropic kernel with an Epanechnikov weighted filter. In addition, photons are cast from quad area light sources. Sampling over the domain of (u, v, θ, ϕ) is achieved using a low-discrepancy Halton sequence using the first 4 primes with a Faure permutation [Fau92] applied to improve the distribution. This method effectively reduces very low-frequency noise negating the need for high numbers of relaxation iterations. All images rendered with our technique use 20 photons in the radiance estimate unless otherwise stated. Below this threshold we find that kernel artefacts (which are typically negligible at larger bandwidths) begin to appear.

All the images in our examples (including test patterns) are tone mapped using the operator outlined by Reinhard *et al* [RSSF02]. We apply our relaxation technique to photon maps from scenes containing a mixture of high- and low-frequency caustics generated as a result of both reflection and refraction. In all examples, $\lambda_{max} = 2.0$ and $\lambda_{min} = 1.2$. Table 1 provides further detailed statistics for all of the example images.

Figure 1 demonstrates the effectiveness of photon relaxation at removing noise from caustics exhibiting large areas of uniform illumination. Such scenes are typically difficult to render efficiently using existing methods given there is little high-frequency detail to mask the irregularities in the caustic illumination, thus necessitating the use of a large kernel. Notice how all low-frequency noise has been removed using our technique while preserving sharp discontinuities around the shadow penumbra.

Figure 6 contains a focussed cardioid caustic from an tinted metal ring. This scene highlights the effectiveness of the feature detection algorithm at locating and preserving discontinuities in illumination. Note how the compact kernel

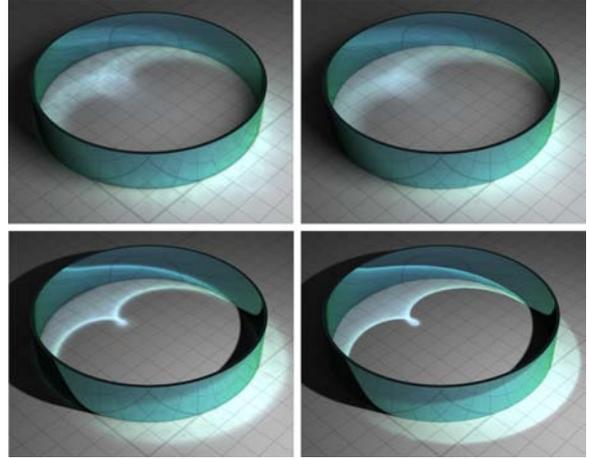


Figure 6: A cardioid caustic from a coloured metal ring. Upper and lower left: unrelaxed photon maps rendered with 100 photons in the radiance estimate. Upper and lower right: relaxed by 20 iterations and rendered with 20 photons in the radiance estimate. Upper left and right: low-frequency caustics from a large area light. Lower left and right: high-frequency caustics from a point light.

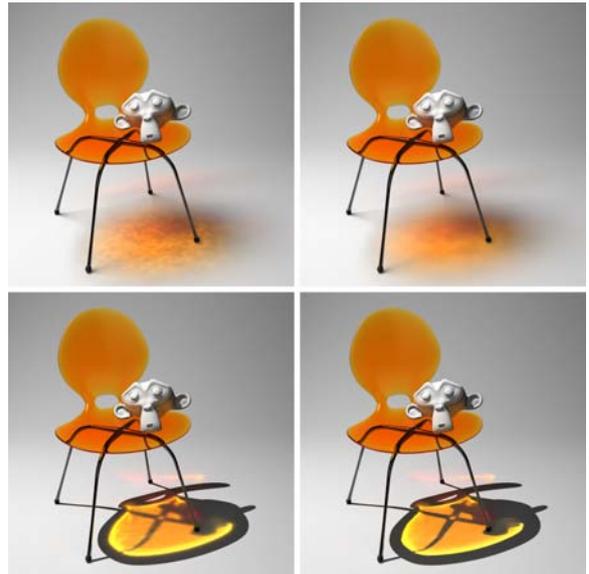


Figure 7: Caustics from light passing through a transparent plastic chair. The frames in this figure correspond to those in figure 6. Unrelaxed photon maps are rendered with 200 photons in the radiance estimate. Using our method, photons are relaxed by 20 iterations and rendered with 20 in the radiance estimate.

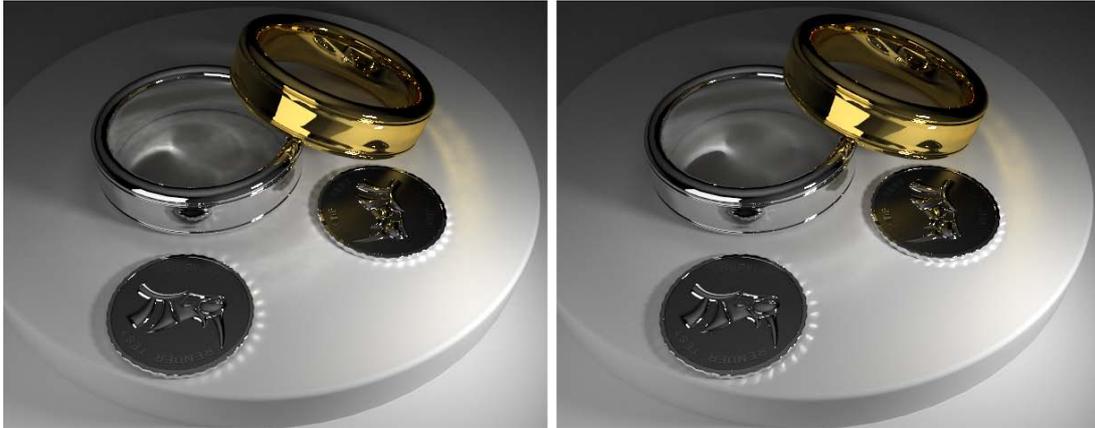


Figure 8: Rings and Coins. Left frame: unrelaxed photon map rendered with 350 photons in the radiance estimate. Right frame: relaxed by 20 iterations and rendered with 20 photons in the radiance estimate.



Figure 9: A close-up of the coins in figure 8. Left frame: Relaxed by 20 iterations; 20 photons in the radiance estimate. Middle frame: Unrelaxed; 350 photons in the radiance estimate. Right frame: Unrelaxed; 20 photons in the radiance estimate.

size effectively eliminates proximity bias along the bright edges of the caustic.

In figure 7, caustics caused by light passing through the thin plastic of the chair are projected onto the ground plane. The fidelity of the diffuse caustics on the underside of the monkey demonstrates how our meshless relaxation approach successfully preserves photon map integrity even on complex geometry. Also note how the detail is preserved around sharp discontinuities while simultaneously removing low-frequency noise in more uniform areas.

Figure 8 is an example of a scene containing intricate caustics. The edges of the coins create a serrated pattern which is blurred out when a high-bandwidth kernel is applied. Notice how our method successfully identifies edge details which are correctly preserved and enhanced during the relaxation step.

5. Analysis and Discussion

Table 1 provides statistics for photon map pre-processing and render time speed-up. One of the most prominent advantages demonstrated by these data is the decreased render time due to fewer photons being required during density estimation. We also analyse the effectiveness of our ap-

proach when compared to existing methods and demonstrate the blue noise properties of the photon distribution by means of Fourier decomposition.

5.1. Bias Tests

Previous literature on the subject of bias and noise reduction provides experimental validation of their techniques with the use of photon test patterns. In particular, Schjøth *et al* [SOS06] use a purely random distribution of 50,000 photons superimposed with an ellipse and two vertical bars at a density 10 times that of the background. The flux of each photon is a constant. The authors present this pattern in order to demonstrate the edge-preserving qualities of their algorithm. In figure 10 we reproduce their examples to demonstrate and compare the effectiveness of our relaxation-based approach.

5.2. Fourier Decomposition

To demonstrate the blue noise properties of the relaxed photon sample distribution we apply a discrete Fourier transform to a 256^2 pixel sample region with an absolute point density of 5%. From these data we derive the radially-averaged power spectrum and angular anisotropy. In all tests we use

	Red Filter		Coins		Chair (point)		Chair (area)		Ring (point)		Ring (area)	
Image Resolution	480 x 640		840 x 525		600 x 600		600 x 600		600 x 478		600 x 478	
Polygons	13,002		126,797		26,570		26,570		6,062		6,062	
Photons	58,000		433,000		74,000		73,000		60,000		56,000	
Render time μ	5.09		10.05		6.25		5.5		5.54		4.86	
Feature search t	1.26s		10.78s		4.24s		2.8s		2.67s		1.74s	
Relaxation t	6.67s		47.01s		12.99s		12.33s		5.83s		5.42s	
Iterations	20		20		20		20		20		20	
$[\alpha, \beta]$	[0, 0.35]		[0.2, 0.35]		[0.2, 0.35]		[0, 0.35]		[0.2, 0.35]		[0, 0.35]	
Radiance estimate	Std.	Rlx.	Std.	Rlx.	Std.	Rlx.	Std.	Rlx.	Std.	Rlx.	Std.	Rlx.
	200	20	350	20	200	20	200	20	100	20	100	20

Table 1: Render statistics. Note that **render time** μ represents the speed-up in calculation of illumination from the caustic component only. **Std.** and **Rlx.** represent values as applicable to the standard method and our relaxed method respectively.

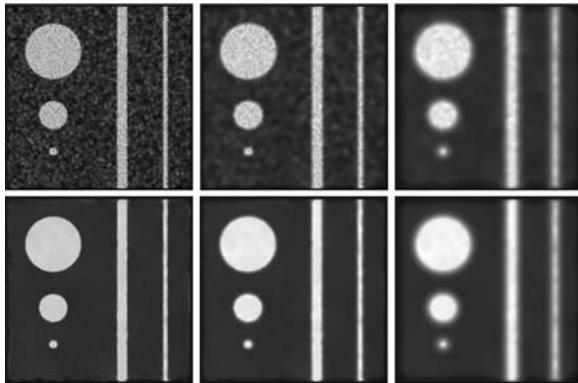


Figure 10: Bias and noise tests. Top row: unrelaxed photon map. Bottom row: relaxed by 20 iterations. Columns from left to right: 10, 50 and 250 photons in the radiance estimate respectively.

the measured anisotropy of white noise (0.26) as a baseline of $-10dB$ (see [Uli88]).

Figure 11 shows analytical data obtained from a random point distribution after 50 relaxation steps. In this example over-relaxation is disabled (by setting both λ_{min} and λ_{max} to 1.0). After a sufficient number of iterations this configuration will eventually yield an ideal, hexagonal lattice-like distribution. Unfortunately, the time required to reach this state is prohibitive. We found that by increasing λ_{min} to a value of 1.2 and λ_{max} to 2.0, the photons relax into a configuration with distinct cellular patterns similar to those in [Tur91]. Significantly, the resulting distribution exhibits a blue noise spectral signature with stronger oscillations and a corresponding lower render variance similar to that of the ideal state obtained when over-relaxation is disabled. Analytical data obtained using these parameters can be found in figure 12.

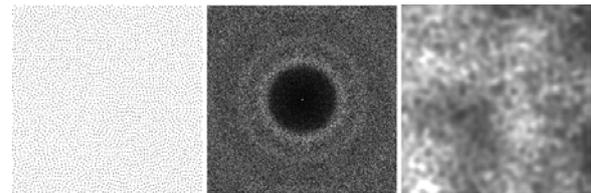
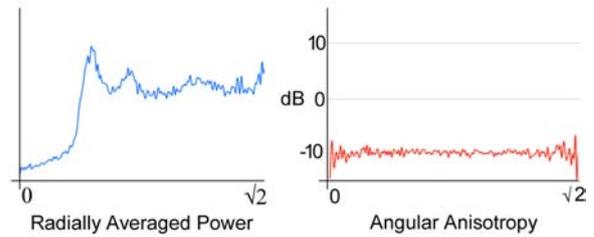


Figure 11: Top left and right: The radial power spectrum and angular anisotropy of an initially random distribution of photons after 50 relaxation iterations. Over-relaxation is disabled. Bottom left: The relaxed point distribution. Bottom centre: The Fourier transform of the previous frame. Note the blue noise spectral signature. Bottom right: A render of the photon map with high gamma correction to highlight noise.

6. Conclusion

In this paper we have introduced a new method of improving the quality of caustics generated using photon mapping. We have shown the technique to be effective in a range of test scenarios and demonstrated the enhanced rendering times (table 1) and lower noise as a result.

For future work we would like to explore the potential of the technique when rendering volume photon maps. We would also like to extend our feature detection algorithm to handle flux intensity and chromaticity boundaries, not just discontinuities in photon density.

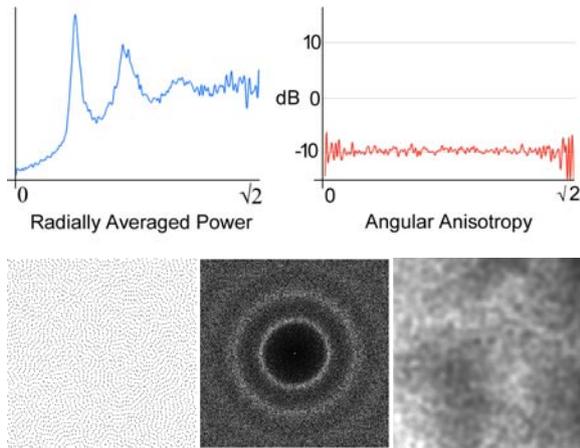


Figure 12: As figure 11 except with the over-relaxation parameters λ_{max} and λ_{min} set to 2.0 and 1.2 respectively. Note the resulting cellular patterns in the photon distribution and the slightly reduced variance in the final render.

References

- [Col94] COLLINS S.: Adaptive splatting for specular to diffuse light transport. In *Proceedings of the 5th Eurographics Workshop on Rendering* (1994), pp. 119–135.
- [Coo86] COOK R. L.: Stochastic sampling in computer graphics. *ACM Transactions on Graphics* 5, 1 (1986), 51–72.
- [DW85] DIPPÉ M. A. Z., WOLD E. H.: Antialiasing through stochastic sampling. In *SIGGRAPH '85: Proceedings of the 12th annual conference on Computer graphics and interactive techniques* (1985), ACM, pp. 69–78.
- [Fau92] FAURE H.: Good permutations for extreme discrepancy. In *J. Number Theory* (1992), vol. 42, pp. 47–56.
- [HP02] HEY H., PURGATHOFER W.: Advanced radiance estimation for photon map global illumination. vol. 21, pp. 541–545.
- [Ige99] IGEHY H.: Tracing ray differentials. In *SIGGRAPH '99: Proceedings of the 26th annual conference on computer graphics and interactive techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 179–186.
- [JB02] JENSEN H. W., BUHLER J.: A rapid hierarchical rendering technique for translucent materials. In *ACM Transactions on Graphics* (2002), vol. 21, ACM, pp. 576–581.
- [JC95] JENSEN H. W., CHRISTENSEN N. J.: Photon maps in bidirectional Monte Carlo ray tracing of complex objects. In *Computers and Graphics* (1995), vol. 19, pp. 215–224.
- [Jen96a] JENSEN H. W.: Global illumination using photon maps. In *Proceedings of the Eurographics workshop on Rendering techniques '96* (1996), Springer-Verlag, pp. 21–30.
- [Jen96b] JENSEN H. W.: *The Photon Map in Global Illumination - PhD Thesis*. Technical University of Denmark, Lyngby, 1996.
- [KCODL06] KOPF J., COHEN-OR D., DEUSSEN O., LISCHINSKI D.: Recursive Wang tiles for real-time blue noise. *ACM Transactions on Graphics* 25, 3 (2006), 509–518.
- [Kel95] KELLER A.: A quasi-Monte Carlo algorithm for the global illumination in the radiosity setting. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing* (1995), Springer, pp. 239–251.
- [Llo83] LLOYD S. A.: An optimization approach to relaxation labeling algorithms. *Image and Vision Computing* 1, 2 (May 1983), 85–91.
- [MF92] MCCOOL M., FIUME E.: Hierarchical Poisson disk sampling distributions. In *Proceedings of the conference on Graphics interface '92* (1992), Morgan Kaufmann Publishers Inc., pp. 94–105.
- [Mys97] MYSZKOWSKI K.: Lighting reconstruction using fast and adaptive density estimation techniques. In *Proceedings of the Eurographics Workshop on Rendering Techniques '97* (1997), pp. 251–262.
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics* 23, 3 (2004), 488–495.
- [RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. In *SIGGRAPH '02* (2002), ACM, pp. 267–276.
- [Sch03] SCHREGLE R.: Bias compensation for photon maps. *Computer Graphics Forum* 22, 4 (2003), 729–742.
- [SFES07] SCHJØTH L., FRISVAD J. R., ERLEBEN K., SPORRING J.: Photon differentials. In *GRAPHITE '07* (2007), ACM Press, pp. 179–186.
- [Shi91] SHIRLEY P.: Discrepancy as a quality measure for sample distributions. In *Computer Graphics Forum: Proceedings of Eurographics '91* (1991), Elsevier Science Publishers, pp. 183–194.
- [SJ09] SPENCER B., JONES M. W.: Hierarchical photon mapping. *IEEE Transactions on Visualization and Computer Graphics* 15, 1 (Jan-Feb 2009), 49–61.
- [SOS06] SCHJØTH L., OLSEN O. F., SPORRING J.: Diffusion based photon mapping. In *International conference on Computer Graphics Theory and Applications* (2006), NSTICC Press, Setubal, Portugal, pp. 168–175.
- [SWH*95] SHIRLEY P., WADE B., HUBBARD P. M., ZARESKI D., WALTER B., GREENBERG D. P.: Global illumination via density estimation. In *Proceedings of 6th Eurographics Workshop on Rendering* (1995), Springer, pp. 219–230.
- [Tur91] TURK G.: Generating textures on arbitrary surfaces using reaction-diffusion. In *SIGGRAPH '91* (1991), ACM, pp. 289–298.
- [Tur92] TURK G.: Re-tiling polygonal surfaces. In *SIGGRAPH '92* (1992), vol. 26, ACM, pp. 55–64.
- [Uli88] ULICHNEY R.: Dithering with blue noise. *Proceedings of the IEEE* 76, 1 (Jan 1988), 56–79.
- [Wal98] WALTER B.: *Density Estimation Techniques for Global Illumination - PhD Thesis*. Cornell University, 1998.
- [WHSG97] WALTER B., HUBBARD P. M., SHIRLEY P., GREENBERG D. P.: Global illumination using local linear density estimation. *ACM Transactions on Graphics* 16, 3 (1997), 217–259.