

Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering

M. Zwicker¹, W. Jarosz², J. Lehtinen³, B. Moon⁴, R. Ramamoorthi⁵, F. Rousselle², P. Sen⁶, C. Soler⁷, and S.-E. Yoon⁴

¹University of Bern

²Disney Research Zürich

³Aalto University and NVIDIA

⁴KAIST (Korea Advanced Institute of Science and Technology)

⁵University of California, San Diego

⁶University of California, Santa Barbara

⁷INRIA Rhône-Alpes

Abstract

Monte Carlo integration is firmly established as the basis for most practical realistic image synthesis algorithms because of its flexibility and generality. However, the visual quality of rendered images often suffers from estimator variance, which appears as visually distracting noise. Adaptive sampling and reconstruction algorithms reduce variance by controlling the sampling density and aggregating samples in a reconstruction step, possibly over large image regions. In this paper we survey recent advances in this area. We distinguish between “a priori” methods that analyze the light transport equations and derive sampling rates and reconstruction filters from this analysis, and “a posteriori” methods that apply statistical techniques to sets of samples to drive the adaptive sampling and reconstruction process. They typically estimate the errors of several reconstruction filters, and select the best filter locally to minimize error. We discuss advantages and disadvantages of recent state-of-the-art techniques, and provide visual and quantitative comparisons. Some of these techniques are proving useful in real-world applications, and we aim to provide an overview for practitioners and researchers to assess these approaches. In addition, we discuss directions for potential further improvements.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms

Today, Monte Carlo methods are widely accepted as the most practical methods for realistic image synthesis. The rendering equation [Kaj86] formulates this problem as an integral over all light paths that connect any point on a light source to a point on an image sensor. Monte Carlo methods estimate this integral by randomly sampling light paths and accumulating their image contributions. Even simple Monte Carlo rendering algorithms come with a number of very desirable properties: they are consistent, which means that as the number of sampled paths increases, the estimated image converges to the correct solution; some algorithms, like (bidirectional) path tracing, are also unbiased, that is, the expected value of the estimated image corresponds to the correct solution and the error consists only of variance; and finally, they are applicable to most scene configurations that are relevant in practice.

On the other hand, because only a limited number of random light paths can be sampled to compute each image, all Monte Carlo methods suffer from variance in the estimated pixel values, which appears as image noise. Unfortunately, computation times to obtain visually satisfactory results without noticeable noise are often in the minutes and hours. Therefore, researchers have proposed a wide variety of noise or variance reduction strategies over the years, from different path sampling strategies (importance sampling, bidirectional techniques, Metropolis sampling) to statistical techniques (quasi-Monte Carlo sampling using low-discrepancy sequences, density estimation, control variates), or signal processing methods (frequency analysis, non-linear filtering), to name the most prominent ones. In this paper, we survey recent advances in adaptive sampling and reconstruction,

which have proven very effective at reducing the computational cost of Monte Carlo techniques in practice.

The amount of variance generally varies greatly between local regions of rendered images. Adaptive sampling refers to techniques that control sampling densities based on previously acquired samples, in contrast to sampling predetermined target densities, to distribute samples according to the local amount of variance. Adaptive reconstruction computes output pixel values using locally defined reconstruction filters. These filters reduce variance by sharing information between pixels, and by aggregating samples over larger image regions, while trying to avoid blurriness. Adaptive sampling generally requires an adaptive reconstruction step, while adaptive reconstruction may also be performed without adaptive sampling. A common strategy is to sample and reconstruct iteratively, where the estimated reconstruction error determines sampling densities in the next step.

Pioneering efforts in adaptive sampling and reconstruction occurred almost simultaneously with the development of the first Monte Carlo algorithms. For example, Mitchell [Mit87] proposed a two-step approach to adaptively sample the image plane considering a contrast metric inspired by human perception, and he developed a reconstruction filter to deal with the nonuniform sample distributions. Parker and Sloan [PS89] and Guo [Guo98] sample the image plane using progressive refinement, and both apply polynomial reconstruction filters. Ward et al.'s irradiance caching algorithm [WRC88] sparsely and adaptively samples irradiance in the image plane in a greedy fashion, and uses a custom tailored reconstruction strategy to interpolate irradiance at each pixel. Inspired by these works, other researchers have considered more advanced perceptual error estimates [BM98, RPG99], or alternative reconstruction strategies such as splatting [RW94] and anisotropic diffusion [McC99]. Despite these early successes, recent advances in adaptive sampling and reconstruction have been significant, reducing the number of samples often by orders of magnitudes without sacrificing quality compared to these earlier methods.

As a starting point that sparked these advances, we would like to single out two papers that represent trends in what we call “a priori” and “a posteriori” techniques emerging in their aftermath. Durand et al. [DHS*05] embarked on the project to develop a local frequency analysis of light transport. This revived the idea to enhance samples with information obtained from a local, analytic analysis of the light transport equations, and use this to control sampling and reconstruction. A single sample enhanced with such information could a priori, without considering other samples, determine the local sampling density or support of a reconstruction filter. Hence, for the purpose of organizing our survey, we use the term “a priori” to characterize methods following this pattern. A priori methods are not restricted to frequency analysis, but may rely on other information such as derivatives. A key idea in a priori methods is also to perform the analysis in

a higher dimensional domain, which was first proposed by Mitchell [Mit91].

On the other hand, Overbeck et al. [ODR09] estimate the error of a 2D wavelet approximation of the rendered image after distributing an initial set of samples, and then they iteratively add more samples. Here, sample densities and reconstruction filters are derived a posteriori from the statistics of a set of samples. Hence, we call such methods “a posteriori”. While this strategy goes back to the earliest adaptive sampling techniques [Mit87], Overbeck’s approach combines it with more powerful reconstruction filters, which was a crucial step paving the way for further advances.

We review developments in adaptive sampling and reconstruction roughly starting with the aforementioned works that arguably initiated a revival of this area. We broadly group approaches into “a priori” (Section 1) and “a posteriori” (Section 2), although some techniques employ a combination of both strategies. Finally, we discuss strengths, weaknesses, and practicability of these techniques, for example in production rendering (Section 3).

1. A Priori Methods

A priori methods enhance Monte Carlo samples with information derived from an analytical analysis of the light transport equations. To evaluate this analysis in practice, they often require access to additional scene information including local 3D geometry or analytic BRDF models. They then construct adaptive reconstruction filters based on this information. To ensure that the analysis remains tractable, most *a priori* methods focus on a restricted set of effects, such as depth of field, motion blur, soft shadows, or diffuse indirect illumination, and some combinations thereof. We organize our review according to the type of analysis of each method, where we distinguish three main types: local frequency analysis (Section 1.1), light field structure analysis (Section 1.2), and derivative analysis (Section 1.3).

1.1. Local Frequency Analysis

Adaptive techniques based on frequency analysis leverage classical signal processing theory [OS09] to determine local sampling rates and reconstruction filters. Nyquist’s sampling theorem states that, for a bandlimited signal, a sampling rate of at least twice the signal bandwidth ensures that the signal can be perfectly reconstructed from the samples without aliasing. A key idea is to determine bandlimits locally and apply the sampling theorem to obtain nonuniform, locally optimal sampling densities. Nyquist’s sampling theorem can also be leveraged to solve the adaptive reconstruction part: reconstruction filters should simply match the local sampling densities, that is, their spectral support should extract exactly the central replica of the sampled signal. The most common application of this theory in computer graphics is in texture

filtering, and the work reviewed here was also inspired by the pioneering contribution by Heckbert [Hec86] in this area.

In geometrical optics, the light energy in a scene is described by its radiance distribution. Restricted to rays incident or reflected from surfaces, radiance distributions can be represented by 4D functions in the spatio-angular domain called light fields [LH96]. The frequency analysis of light fields has been studied for image based rendering [CTCS00], computational cameras [LHG*09], and light field displays [ZMDP06, WLHR11]. Durand et al. [DHS*05] embarked on the project to conduct a local frequency analysis of light fields in the neighborhood of individual rays, and study the effects of light transport phenomena such as transport along the ray, reflection on curved surfaces, and occlusion. We show part of this analysis schematically in 2D flatland in Figure 1, where each ray is given by a position x and direction v . Determining local frequency spectra under various light transport effects would make it possible to derive adaptive sampling rates of rays.

A key question for practical applications is how to represent the local spectra. One option is to estimate only the bandwidths of the spectra, since this is sufficient to determine sampling rates via the Nyquist theorem. Below we will also discuss additional strategies that have been explored. Note that light field spectra are 4D functions in the spatio-angular domain, hence sampling densities and reconstruction filters are most naturally defined in a 4D space. For practical implementation, however, it can be more efficient to work in 2D image space by projecting from 4D to 2D.

Durand et al. [DHS*05] demonstrate a proof-of-concept algorithm for adaptive sampling of glossy reflection under environment illumination and partial occlusion. They derive approximate, local 2D bandwidths of the image signal and distribute samples accordingly. They reconstruct the image by adjusting the spatial support of a cross-bilateral filter according to the bandwidth estimates. Bagher et al. [BSS*13a] extend this approach to measured BRDFs using wavelets to estimate local bandwidths of illumination, textures, and BRDFs.

A key insight of Durand et al.’s work is that effects including diffuse emission and reflection, transport in free space, and occlusion lead to light fields composed of locally linear structures as can be seen in Figure 1. The slopes of these structures are linked to distances to diffuse emitters and reflectors, distances of transport, and distances to occluders, respectively. The corresponding frequency spectra have wedge-like shapes with slopes perpendicular to the slopes in the spatio-angular domain. A series of consecutive works exploits this observation by deriving bandwidths of sheared filters that cover these wedges, as shown in Figure 2.

Egan et al. [ETH*09] added motion blur to the frequency analysis, while including the effects of shading, occlusion, and cast shadows. They show that these effects locally determine an “effective velocity”, and the light field energy is

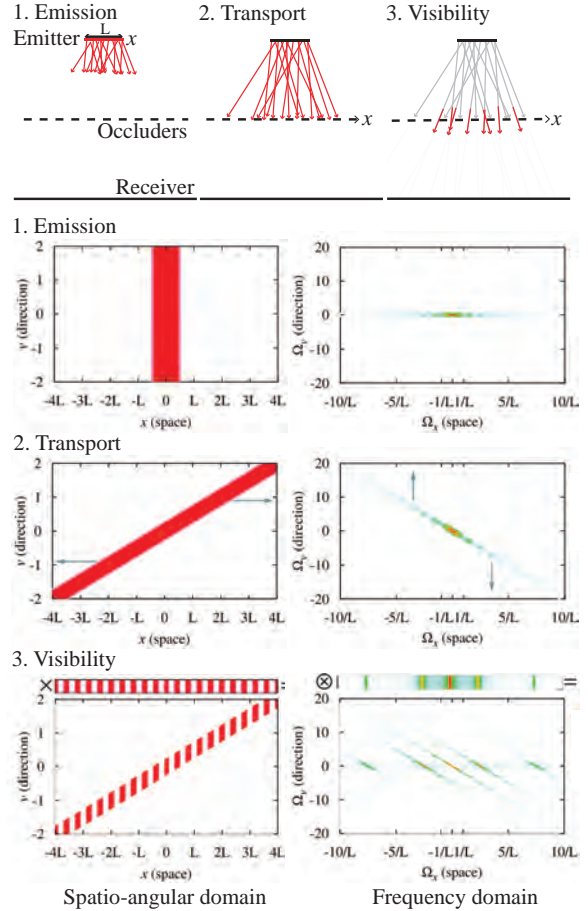


Figure 1: Frequency analysis of light fields under diffuse emission, transport, and occlusion. A source (Step 1) of size L emits light uniformly over directions, hence the light field is constant over v , and non-zero only for positions $|x| < L/2$. In the frequency domain, this is a sinc over spatial frequencies Ω_x . Transport to the occluder (Step 2) corresponds to moving the x axis to the occluder plane. This results in a shear along x in the spatio-angular domain, because the directions of rays (their v coordinates) stay constant. Their new x coordinates, on the other hand, change proportionally to their direction v . In the frequency domain, this causes a related shear along the other axis (angular frequencies Ω_v). Finally, occlusion (Step 3) is a multiplication with a binary function in the spatial domain, and a convolution in the frequency domain. Figure adapted from Durand et al. [DHS*05].

mostly contained in a wedge in the frequency domain defined by the minimum and maximum velocities. This analysis is similar to the work by Chai et al. [CTCS00] who first described such wedge spectra in the context of light field rendering. It is also related to early work by Shinya et al. [Shi93] who derived spatio-temporal filters using motion vectors at




Figure 2: Linear structures in light fields, caused by diffuse emission and reflection, transport, occlusion, and linear motion, lead to wedge-like frequency spectra. They can be packed densely in the frequency domain (left), corresponding to sparse sampling. Reconstruction employs sheared filters (middle) that extract the central wedge (right). This figure is adapted from Egan et al.'s work on motion blur [ETH*09], illustrating spectra over a spatio-temporal domain (x, t) .

the sub-pixel level to anti-alias animation sequences. Egan et al. derive sparse sampling rates for the wedge spectra that are guaranteed to avoid aliasing, and construct matching sheared 3D spatio-temporal filters to reconstruct the signal as shown in Figure 2. In practice, they approximate the ideal sheared box filters with truncated Gaussians.

Subsequently, they also studied soft shadows from planar area light sources [EHDR11] following a similar approach. They derive wedge spectra for occluders, where the shape of the wedge depends on the range of occluder depths, transform the spectra to the frequency space of the receiver, and derive a sheared reconstruction filter on the receiver that avoids aliasing. In a practical algorithm, they uniformly but sparsely sample occlusions across the image and determine 4D sheared filters at each sample in a first pass. In a second pass, they use the 4D sheared filters to reconstruct soft shadows at each pixel from the sparse samples. Finally, Egan et al. [EDR11] analyzed distant illumination under occlusion and including low-frequency BRDFs, and they derive local bandwidths and 4D sheared filters similarly as above. Also the algorithmic implementation follows the same two step approach with sparse, uniform initial sampling followed by reconstruction using the sheared filters.

Mehta et al. [MWR12] observed that the 4D sheared filtering approach developed in previous methods [ETH*09, EHDR11, EDR11] adds minutes of overhead to the rendering time. In addition, the two step algorithm requires additional storage for the initial samples. Instead, they propose a method suitable for real-time rendering of soft shadows based on axis-aligned 2D filtering in image space as shown in Figure 3, and without additional storage requirements. Using similar analysis as above, they derive sampling rates and filter sizes to adaptively sample and reconstruct, but with axis aligned filters directly in 2D image space. Their method also converges to reference Monte Carlo solutions as sampling rates increase beyond the minimum required by the analytical analysis. Later, they extended this approach to include diffuse indi-




Figure 3: Mehta et al. [MWR12, MWRD13, MYRD14] use axis aligned filters with bandwidths (Ω_x^f, Ω_y^f) in (x, y) image space (left, red rectangle) to reconstruct the wedge spectra (light blue). Axis aligned filters can be packed less tightly (right) than sheared filters (Figure 2), implying higher sampling rates, but they can be implemented very efficiently.

rect illumination [MWRD13]. Most recently they developed a method to combine primary (depth of field) and secondary effects (soft shadows, diffuse and moderately glossy indirect illumination) in a single framework [MYRD14], again using image space adaptive sampling and axis-aligned filtering. Their method requires little computation and incurs no storage overhead, yet it achieves impressive results as shown in Figure 4. In their paper, they demonstrate up to $30\times$ sampling rate reduction compared to equal quality noise-free MC.

Soler et al. also analyzed depth of field [SSD*09] using the same frequency domain machinery [DHS*05] to model transport, occlusion, and aperture effects of local light field spectra around a ray. In contrast to the methods described above, they do not represent the spectra using bandlimits and sheared filters, but by sampling the local power spectra and transforming the samples according to the effects of transport, occlusion, and lens aperture. They then derive bandwidths in the image and on the lens aperture to determine image and aperture sampling densities. They reconstruct the image by collecting the k -nearest samples at each pixel.

Beyond bandlimits in combination with sheared filters, and sampled spectra, Belcour et al. [BSS*13b] propose a third alternative to represent local light field spectra. They use covariance matrices that define multi-dimensional Gaussians, which may be degenerate if the matrices are rank deficient. Various light transport effects (transport in free space, occlusion, reflection, etc.) are modeled as operators applied to the multi-dimensional Gaussians, which boil down to simple manipulations of the covariance matrices. The unique feature of this technique is that different light transport and lens effects (motion blur, depth of field) are easily combined by concatenating these operators. In a proof-of-concept implementation, they propose to propagate covariance spectra along sampled paths, project and accumulate them in image space, and deter-




Figure 4: This figure is adapted from Mehta et al. [MYRD14], who handle primary (depth of field) and secondary effects (soft shadows, indirect illumination) in a single framework. The approach relies on fast, axis aligned filtering in image space. We report the number of rays (that is, ray segments) per pixel (rpp) and show an equal time comparison (on the GPU) to stratified Monte Carlo rendering.

mine sampling rates and reconstruction filters based on this information. Subsequently, Munkberg et al. [MVH*14] built on Belcour et al.’s light field operators to develop real-time capable filters for motion blur and depth of field. The key idea in their work is to split samples into image plane tiles and partition each tile into depth layers. Then they filter each layer separately with an appropriate filter derived using the light field operators.

We summarize the methods based on local frequency analysis of light fields in Table 1. A limitation of most methods is that they analyze only a single or a limited combination of effects, and they require specialized algorithms for each. The work by Mehta et al.’s [MYRD14] and Belcour et al. [BSS*13b] are the most general ones, supporting depth of field, soft shadows, and diffuse and slightly glossy indirect illumination. Mehta et al. provides lower computational overhead and higher quality at similar sampling rates than Belcour et al. On the other hand, Belcour et al.’s approach is more easily extendable to additional effects including motion blur and participating media [BBS14] by deriving appropriate operators on covariance matrices for these effects. The covariance matrix operators have also been adapted to real-time filtering as well [MVH*14].

1.2. Light Field Structure Analysis

It is not surprising that light field structures as discussed in the previous section and shown in Figure 1 may be exploited for rendering without resorting to a frequency analysis. Lehtinen et al. [LAC*11] proposed a method that starts with a sparse set of “fat” samples that contain, in addition to radiance, auxiliary information such as positions, normals, veloc-

Table 1: Summary of the main characteristics of methods based on local frequency analysis of light fields. In the “Effects” column, “D” stands for direct, “II” for indirect illumination (an additional “(G)” identifies methods that also support glossy surfaces), “MB” for motion blur, “SS” for soft shadows, “AO” for ambient occlusion, “DOF” for depth of field, and “PM” for participating media. In the “Rep.” column we note the main representation of the local spectra that are used for analysis, where “Bandlim.” stands for bandlimits of ideal (sheared) filters, “Sampled” indicates a sampled representation of the spectra, and “Covar.” denotes a multidimensional Gaussian approximation using covariance matrices. In the “Rec.” column we list reconstruction strategies, where “IS” stands for image space, “ISL” for image space with depth layers, “HD” for high dimensional space, “xb” for cross-bilateral, “AA” for axis-aligned, “aniso.” for anisotropic, “knn” for averaging the k -nearest samples, and “V” for 3D volume. For high dimensional reconstructions, the dimensionality depends on the effects being considered.

Method	Effects	Rep.	Rec.
[DHS*05]	D(G)	Bandlim.	xb IS
[BSS*13a]	D(G)	Bandlim.	xb IS
[ETH*09]	MB	Bandlim.	aniso. HD
[EHDR11]	SS	Bandlim.	aniso. HD
[EDR11]	AO	Bandlim.	aniso. HD
[MWR12]	SS	Bandlim.	AA IS
[MWRD13]	II(G)	Bandlim.	AA IS
[MYRD14]	II(G),SS,DOF	Bandlim.	AA IS
[SSD*09]	DOF	Sampled	knn IS
[BSS*13b]	II(G),SS,DOF,MB	Covar.	aniso. IS
[BBS14]	PM	Covar.	aniso. IS & V
[MVH*14]	DOF, MB	Covar.	aniso. ISL

ities, etc., and uses these samples for densely reconstructing the temporal light field. The upsampled light field is then integrated by brute force, yielding pictures of far higher quality than those obtained from naive reconstruction.

The key idea in the technique is treating the samples as 3D scene points, and reprojecting them for novel viewpoints across the lens, sub-frame times in the shutter interval, and light source positions, using the auxiliary information gathered during the initial sampling. In effect, they employ an analytic geometric model that accurately describes the local light field structures along which the integrand is almost constant, an effect that is crucial also in the algorithms based on frequency analysis (Section 1.1). Reprojection allows the reuse of the initial samples along the predicted trajectories as shown in Figure 5, effectively multiplying the per-pixel sample rates by a large factor. An important component is a technique to detect occlusions and robustly determine the visibility of reprojected samples.

The algorithm handles combinations of motion blur, depth of field, and soft shadows with excellent results at very low

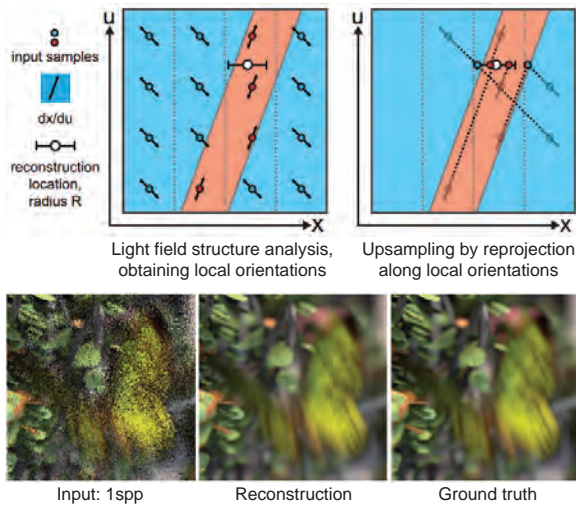


Figure 5: Lehtinen et al. [LAC*11] analyze light field structures at a sparse set of initial samples. They upsample by reprojecting along the light field structures, achieving high quality reconstruction using very sparse initial sampling. The example here uses only one initial sample per pixel (1spp).

initial sampling densities. The authors later extended this approach to the reconstruction of diffuse indirect illumination [LALD12] following a similar strategy. They also generalize their method to glossy materials, where the reprojection is restricted according to the bandwidth of the material reflectance; however, the computational cost of this extension is significant.

1.3. Derivative Analysis

A basic intuition why derivative analysis is useful for adaptive sampling and reconstruction is that small gradients indicate locally constant functions, allowing for low sample rates and large reconstruction filters, and vice versa. Gradient information may be used not only to determine the support size of reconstruction filters, but also their anisotropic shapes. Irradiance caching [WRC88] is a highly practical realization of this basic intuition, designed specifically for evaluating irradiance due to indirect illumination. It computes high quality irradiance samples at sparse locations in the image plane, and estimates an upper bound on the irradiance gradients to control the density of irradiance samples as well as the support of the reconstruction filters.

Since its introduction irradiance caching has been continuously refined. Ward and Heckbert derived a more accurate estimate for irradiance gradients that they used to get smoother and more accurate reconstruction using linear extrapolation with Taylor series [WH92]. Křivánek et al. extended it to handle low-frequency glossy materials [KGPB05] by storing, and differentiating [KGPB05], incident radiance using




Figure 6: Illustration of the benefit of occlusion Hessians (from Schwarzhaupt et al. [SJJ12]). We visualize indirect irradiance on the top left. Occlusion Hessians (top right) adapt the distribution of irradiance samples on the ground plane to indirect shadows cast by an occluder. They lead to more accurate reconstructions in particular with few irradiance samples (“records”). The Cornell box compares results obtained by a first-derivative approach (bounded split-sphere [WRC88], middle row) to using occlusion-aware Hessians (bottom row) at varying sample counts.

spherical harmonics to preserve some directional information. Jarosz et al. extended the method to handle participating media [JDZJ08, JZJ08].

While most irradiance caching variants were developed with an eye towards practical applicability, Ramamoorthi et al. [RMB07] conducted a more comprehensive and thorough first-order analysis of lighting, shading, and shadows in direct illumination. As a proof-of-concept, they also describe an adaptive sampling algorithm based on gradient magnitudes. Jarosz et al. [JSKJ12] recently extended Ramamoorthi et al.’s work by performing a second-order analysis of indirect illumination. Interestingly, there is a close connection between second order derivatives and Gaussian approximations of local frequency spectra as discussed by Belcour et al. [BBS14]. By leveraging the resulting second-order derivatives, Jarosz et al. proposed a Hessian-based irradiance extrapolation for increased accuracy in the Taylor expansion. Schwarzhaupt et al. [SJJ12] improved upon this further. A key contribution is their second-order analysis of occlusions, visualized in Figure 6. Their approach leads to significant improvements in particular with sparse irradiance samples. Despite these improvements, irradiance caching remains a specialized method targeted at handling indirect illumination constrained to diffuse or moderately glossy materials.




Figure 7: The generic template followed by many techniques discussed in Section 2. Typically, the iteration terminates when a given sample budget is exhausted. This figure is adapted from Rousselle et al. [RKZ11].

2. A Posteriori Methods

Instead of analyzing the light transport equations to derive local signal properties analytically, like *a priori* methods, *a posteriori* techniques statistically analyze sets of samples acquired using a Monte Carlo renderer with little additional information. At the minimum these methods require local, usually per-pixel, estimates of variance. We illustrate the generic template followed by many methods in this section in Figure 7. A key idea in most techniques is to use a family of reconstruction filters and develop error estimates for the filter outputs, usually using a mean squared error metric (MSE). In general, MSE can be expressed as a sum of squared bias and variance, where bias corresponds to blurriness, and variance is residual noise from the input Monte Carlo samples. A family of filters should provide a trade-off between bias and variance, and the goal is to locally select the best filter that optimally balances bias and variance to minimize MSE. In addition, the local error estimates also make it possible to control sampling densities. We organize our survey according to the types of filters that are used, distinguishing multiscale filters (Section 2.1), filters designed for generic image denoising (Section 2.2), and filters derived from auxiliary features (Section 2.3). Most approaches rely on image space filtering because of its simplicity and efficiency.

A notable exception is the work by Hachisuka et al. [HJW*08], where they store samples in multidimensional path space, which may include effects such as motion blur, depth of field, and soft shadows. They estimate the local integration error incurred by an initial set of samples, and adaptively distribute more samples where the error is highest in the multidimensional space. This work inspired many of the techniques from Section 1 that also work in a higher dimensional space. In contrast, however, Hachisuka et al.’s method is agnostic of the underlying rendering effects. As a consequence, it cannot match the quality of specialized techniques. On the other hand, this should in principle make it easier to apply the approach to arbitrary combinations of effects, but developing efficient algorithms with higher dimensional path spaces is challenging.

2.1. Multiscale Filters

A family of filters obtained by scaling the support of a basis filter provides a natural trade-off between bias and variance. Overbeck et al. [ODR09] proposed an algorithm following the strategy outlined in Figure 7 using Daubechies wavelets. They perform wavelet analysis of the rendered image and analyze its error using Mitchell’s contrast metric [Mit87]. They perform wavelet shrinkage [DJ94], that is, they subtract a noise estimate from individual wavelet coefficients, to minimize the error metric. In addition, they use the error metric to adaptively distribute samples in the image plane focusing on regions with high error. Rousselle et al. [RKZ11] observed that critically sampled wavelets lead to ringing artifacts when used for denoising. They obtain higher filtering quality using Gaussian filters at several scales. They perform an MSE analysis by estimating bias and variance separately, and they locally select the scale that minimizes the error. We visualize their approach in Figure 8.

The main advantage of such multiscale filters is that they are very efficient to compute, facilitating potential real-time applications. For example, Dammertz et al. [DSHL10] use wavelets for real-time filtering. On the other hand, families of filters parameterized by a single scale parameter require small scales to filter complex image structures without introducing blurriness. While including parameters to describe anisotropy should be possible, this may not be enough to reach the quality of methods that enable even more complex filter shapes, as described in the following sections.

2.2. Leveraging Image Denoising Filters

Image denoising filters developed by the image processing community to restore images corrupted by spatially uniform noise have recently achieved impressive results [DFKE07, BCM05, EA06]. Hence, it seems promising to leverage these methods for adaptive sampling and reconstruction in Monte Carlo rendering.

Patch-based techniques compute denoising filters by as-




Figure 8: Visualization of Gaussian filter scales (top left), brighter levels indicating larger filters. Adaptive sampling (top right, brighter means more samples) places more samples where filters are small or where there is more noise due to defocus or soft shadow effects. The bottom row shows an equal time quality comparison to naive random sampling.

sessing the similarity of local image patches. They are currently among the most successful filters for generic image denoising. A key idea is that similar patches may be used for denoising within large neighborhoods, even globally over the whole image, without introducing blurriness. This justifies the term “non-local”. The non-local means (NL-means) filter [BCM05] embodies the simplest implementation of the strategy. It generalizes bilateral filtering [TM98], which Xu and Pattanaik [XP05] used for denoising in Monte Carlo rendering. Instead of weighting a neighboring pixel based on the difference of its value to the center pixel, NL-means determines the weight based on the similarity of small patches around the neighbor pixel and the center. Compared to bilateral filtering, NL-means is a much more effective denoising filter, since the filter weights are more robust to noise in the input. Rousselle et al. [RKZ12] exploit NL-means filtering for adaptive rendering, and modify it to cope with non-uniform noise levels. Unfortunately, a full error analysis of NL-means is challenging because the filter itself depends on the noisy input. Instead, Rousselle et al. only consider variance and neglect bias. They estimate output variance simply using the per-pixel difference of two independent filtered images. The per-pixel variance estimate then drives adaptive sampling.

Delbracio et al. [DMB*14] propose a similar non-local denoising approach. Instead of comparing patch similarity based on pixel values, however, they gather histograms of




Figure 9: A scene with motion blur and soft shadows (top left: reference, to right: 4spp). The figure compares the sheared filter for motion blur derived from a frequency analysis [ETH*09] (bottom left) to adaptive reconstruction using generic image denoising filters [KS13] (RNGD, bottom right). Both use four samples per pixel (top right) as input.

sampled values at each pixel, and assess patch similarity based on the differences of these histograms measured by the χ^2 distance. In addition, they propose a multiscale approach to remove low frequency noise. They demonstrate that their approach improves upon NL-means filtering.

Kalantari and Sen [KS13] propose a generic method (Removing Noise With General Denoising, “RNGD” in the following, see Table 2 for a list of acronyms) that is designed to operate with any denoising method for globally uniform noise, such as BM3D [DFKE07], assuming that it has a parameter to adjust the filter to the global noise level. First, they propose a novel per-pixel variance estimate using the median absolute deviation [DJ94]. Then, to deal with spatially nonuniform variance in Monte Carlo rendering, they propose to denoise the complete rendered images several times, each time setting a different noise level. Finally, they compose the result image by selecting pixels from the filter outputs with appropriate noise levels.

The techniques discussed here significantly improve the output quality compared to simple multiscale filters. We show an example result in Figure 9. The methods mentioned here and in the previous subsection all come with the advantage that they require minimal changes to an existing renderer. Besides noisy renderings they only need per pixel statistics such as contrast [ODR09] or variance [RKZ12, KS13] estimates, or pixel histograms [DMB*14] as inputs.

2.3. Filters using Auxiliary Features

Besides noisy images, it is easy to save auxiliary image features such as per-pixel normals, depth, or diffuse reflectance from most Monte Carlo rendering systems. The feature images contain rich information about image structures that can be exploited to design adaptive filters that preserve these structures. An early approach by McCool [McC99] used




Figure 10: Comparison of RPF [SD12] (main image) with SURE [LWC12] and RDFC [RMZ13] (in close-ups) at a low, equal sampling rate (8spp). RPF is most effective at such low sampling rates.

depth and normal information to control anisotropic diffusion. More recently, researchers explored the cross-bilateral filter [ED04, PSA*04]. Here, the core idea is to derive the weight with which a neighbor contributes to the filter output at a center pixel from the pairwise feature differences.

Sen and Darabi [SD12] observed that the relation between the feature differences and filter weights should depend on the degree to which the feature determines the true pixel value. After all, certain features are highly dependent on the random parameters of the Monte Carlo process and are therefore not reliable for filtering. They propose to measure this dependency by computing the mutual information between sample features and the random parameters used to compute them. Their approach (Random Parameter Filtering, or “RPF” in the following, see also Table 2) naturally handles noisy features, which occur when rendering effects such as motion blur or depth of field.

A disadvantage of this approach is that the complexity of the algorithm depends on the number of samples and can become expensive for more than a few dozen samples. In

addition, while RPF is based on useful intuition about the influence that features should have on the filter, it does not try to minimize the output error directly, and methods based on explicit error estimation have proven superior at high sampling rates. However, at low sampling rates (e.g., less than 16spp), the ability of RPF to detect noisy features can yield improvements over other methods (see Figure 10). Park et al. [PMKY13] subsequently proposed a modification of RPF whose complexity is independent of the number of samples without significant quality degradation.

Li et al. [LWC12] first introduced Stein’s unbiased risk estimator (SURE) to adaptive sampling and reconstruction in rendering. SURE is a general technique to estimate the accuracy of a statistical estimator. Van De Ville and Kocher [VDVK09] applied it to estimate global parameters in NL-means image denoising, and Li et al. showed how to leverage it to obtain local, unbiased MSE estimates of cross-bilateral filters. They use a set of cross-bilateral filters with spatial support windows of varying sizes, and select the best size at each pixel by picking the filter leading to the smallest estimate of MSE using SURE. Their approach (“SURE” in the following, see also Table 2) typically outperforms RPF presumably because it selects filters according to an actual error estimate. In Figure 11 we show an equal-sample comparison between SURE, the axis-aligned filtering technique based on *a priori* frequency analysis by Mehta et al. [MYRD14], and the approach by Kalantari and Sen (RNGD) using BM3D. While Mehta et al. use GPU raytracing, the other methods are implemented on top of PBRT [PH10]. The filtering overhead of SURE and RNGD is on the order of 30-40 seconds. The methods perform similarly at a fixed sampling rate, where SURE and RNGD tend to overblur, and Mehta has some residual noise.

Rousselle et al. [RMZ13] also build on SURE error estimation and cross-bilateral filtering. Instead of minimizing error by varying the spatial filter support, they carefully design three filters with the same support, but differing in their other parameters. One filter is tuned to be most sensitive to the noisy color information and disregards the features, one disregards color and is determined purely by the features, and the third is in between. These three filters lead to fewer outliers in the error estimate, since error estimation of small filters, such as in Li et al. [LWC12], is very sensitive to noise. In addition, they propose a feature prefiltering step to deal with noisy features and introduce two new features to better represent soft shadows due to direct illumination and caustics. Figure 12 shows the improvements of their approach (Robust Denoising using Feature and Color Information, “RDFC” in the following, see also 2) over Li et al. (SURE).

Similar to the bilateral filter, the guided filter [HST10] can be applied for edge preserving image denoising. Given a denoising window, the ground truth image is locally approximated as a linear function of a guide image. The two coefficients of the linear function are estimated by minimizing




Figure 11: An equal-sample comparison in terms of ray segments per pixel (rpp) of Mehta et al.’s approach [MYRD14] based on a priori frequency analysis for soft shadows, indirect illumination, and depth of field (main image), with a posteriori techniques (in close-ups) by Li et al. [LWC12] (SURE) and Kalantari and Sen [KS13] (RNGD).

the L_2 error between the linearly transformed guide and the noisy input image. Bauszat et al. [BEM11] applied the guided filter to reduce noise generated by Monte Carlo rendering while preserving edges by using geometric features as guide images. Recently, Moon et al. [MCY14] also used a linear model to approximate the ground truth function locally, but in addition they weigh the error of each pixel individually based on the features (“LWR” for Local Weighted Regression, see also Table 2). They introduce an MSE estimate by analyzing bias and variance separately, and they optimize the parameters used to determine the weights to minimize the MSE at each pixel. Finally, they deal with noisy features by performing dimensionality reduction via a truncated SVD (Singular Value Decomposition), and their optimization is solved on the local features with a reduced dimensionality. As shown in Figure 13, their approach achieves high-quality rendering results even in nonlinear functions such as glossy highlights and defocused areas thanks to the automatic parameter selection and the truncated SVD, respectively.



Figure 12: Equal-time comparison of Rousselle et al. [RMZ13] (RDFC, main image) and Li et al. [LWC12] (SURE, in close-ups). RDFC is more robust to noisy error estimates, and it includes a feature to capture soft shadows from direct illumination, which leads to significant quality improvements.

The methods discussed so far in this section use features that are obtained directly as byproducts of usual Monte Carlo rendering. As an interesting alternative, Moon et al. [MJL*13] proposed a new feature, a virtual flash image, which serves as an edge-stopping function. The virtual flash image is motivated by flash photography [ED04, PSA*04] where a flash image provides high-frequency information of the noisy input image taken without a flash. To emulate the flash image in rendering, they put an additional light source on the viewing position and construct the virtual image through an actual shading. This simple idea provides a means to capture a wide set of edges such as reflected and refracted edges with little computational overhead since it is performed by reusing a subset of light paths. Furthermore, they define homogeneous pixels as neighboring pixels whose sample means are within a confidence interval of the true mean at the center pixel. By restricting the denoising to homogeneous pixels, they preserve additional edges (e.g., caustics) that the flash image does not include, and they obtain consistent denoising results that converge to the ground truth.

Table 2 summarizes the main characteristics of the *a posteriori* methods cited in this section. This table does not list the light transport effects supported, since all cited methods are generic (though some, like Dammertz et al. [DSHL10], give bad results if their assumption of noise free features is violated). A commonality of all methods discussed here is that they reconstruct pixel values by building (locally

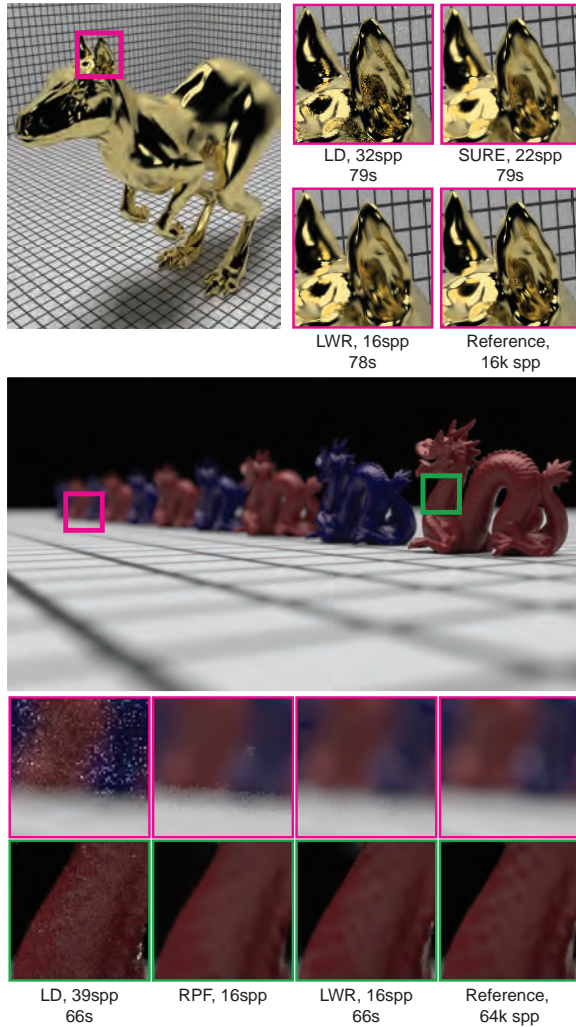


Figure 13: Moon et al. [MCY14] (LWR, main images) use weighted local regression to estimate the denoised output. At equal render times, LWR shows better results on a variety of rendering effects compared to naive low-discrepancy sampling (LD), Li et al. [LWC12] (SURE), and Sen and Darabi [SD12] (RPF, all in close-ups). This figure is modified from Moon et al. [MCY14].

adaptive) weighted averages of complete Monte Carlo path samples. While this is particularly simple, it could be further generalized to aggregating partial paths that do not immediately connect a light source and the image sensor. For example, photon mapping [Jen01] and its generalizations [KDB14, HPJ12, GKDS12] aggregate several “light subpaths” (partial paths connected to light sources) when connecting to individual “eye subpaths” (partial paths connected to the image sensor). The lightcuts algorithm builds adaptive hierarchies of light subpaths, and connects these to individ-

Table 2: Summary of the main characteristics of the a posteriori methods cited. In the “ASR” column, “S” stands for adaptive sampling, and “R” for adaptive reconstruction. Methods performing both (“SR”) also support iteratively adjusting the sampling density and adding more samples. The “Metric” column indicates the error metric used to guide adaptive sampling or reconstruction. If it is empty the method does not adapt to any error estimate. In the “Reconstruction” column, an additional “*” identifies methods that leverage information contained in feature buffers. The last column “Acr.” lists acronyms that we use in the comparison figures this paper.

Method	ASR	Metric	Reconstr.	Acr.
[Mit87]	S	Contrast	uniform	-
[RW94]	R	Variance	splatting	-
[BM98]	S	Perceptual	uniform	-
[McC99]	R	Variance	aniso. diff.*	-
[RPG99]	S	Perceptual	uniform	-
[FP04]	S	Perceptual	uniform	-
[XP05]	R	-	bilateral	-
[ODR09]	SR	Contrast	wavelet	-
[DSHL10]	R	-	wavelet*	-
[BEM11]	R	-	local regr.*	-
[RKZ11]	SR	MSE (bias, var.)	Gaussian	-
[RKZ12]	SR	Variance	NL-means	-
[SD12]	R	Mutual inf.	bilateral*	RPF
[LWC12]	SR	MSE (SURE)	bilateral*	SURE
[KS13]	SR	Variance	gen. denoising	RNGD
[RMZ13]	SR	MSE (SURE)	bilateral*	RDFC
[DMB*14]	R	χ^2	NL-means	-
[MCY14]	SR	MSE (bias, var.)	local regr.*	LWR

ual [WFA*05] or groups of eye subpaths [WABG06] while respecting a desired error bound. It may be interesting to investigate extensions of the algorithms summarized here to operate on light subpaths also, instead of requiring complete path samples.

There is also room to explore between the extremes of working in 2D image space, like most methods discussed here, and the full high-dimensional path space [HJW*08]. For example, deCoro et al. [DWR10] represent samples in a joint image and color space to detect and reject outlier samples in this 5D space. Finally, while all reconstruction methods discussed here are driven (implicitly) by (weighted) L_2 errors, one could generalize this to formulations inspired by total variation methods or compressive sensing. They rely on sparse error norms that have proven to provide visually more convincing results, for example, in image restoration [PF14]. While some methods based on compressive sensing have been proposed for rendering [SD11, HWRH13], it is likely that the full potential of these techniques has not been realized yet.

3. Discussion and Conclusions

We have discussed recent advances in adaptive sampling and reconstruction for Monte Carlo, broadly categorizing approaches into two groups, *a priori* and *a posteriori*. *A priori* methods attempt to leverage information obtained from an analytic analysis of the light transport equations. We have distinguished three strategies, local frequency analysis, light field structure analysis, and derivative analysis. Frequency analysis methods are attractive because they build on the solid foundation of the Nyquist sampling theorem to derive sampling rates and reconstruction filters. The derived filters have relatively simple shapes, however, such as sheared anisotropic filters approximated by Gaussians. Hence high frequency effects may require small filters (in the spatial domain) and high sampling rates. On the other hand, fast axis-aligned filters are showing promise for real-time applications. Instead of filtering, the approaches using light field structure analysis that we discussed upsample the light field by extrapolating new samples from a sparse initial set. Because of the explicit visibility determination, this leads to more complex support shapes of the extrapolated samples, and excellent results at very low initial sampling densities. As a downside, the methods require the storage of all initial samples, and reconstruction overhead is higher. A challenge for all these techniques is to combine the analysis of all relevant effects, from motion blur to glossy indirect illumination, into one framework. Currently, Belcour et al.'s light transport operators on covariance matrices achieve this, but this approach is also limited by simple Gaussian filter shapes. Finally, derivative analysis is mostly used in irradiance caching techniques, and recent advances using higher order analysis have further advanced the approximation quality for a given computational budget. While irradiance caching is highly successful in practice, it remains limited to sampling and reconstructing diffuse or moderately glossy indirect illumination.

A posteriori methods forego analytic analysis for statistical error estimation based on acquired samples. They are typically ignorant of the underlying rendering effects and applicable to arbitrary combinations of light transport phenomena. A key property of these methods is also that they rely on heuristically defined filters, such as cross-bilateral filters using auxiliary features, which can have arbitrary support shapes and extend over large image areas. Hence, these methods can be effective at removing noise even in image regions with intricate structures. In addition, many methods are consistent because their filters are eventually restricted by the empirical variance of rendered pixel means, which vanishes with increasing sample counts. It is also relatively straightforward in practice to modify existing renderers to support *a posteriori* adaptive sampling and reconstruction, since the required information such as feature images or pixel variances are available as a byproduct of conventional rendering. As a disadvantage, the most sophisticated reconstruction filters, like local weighted regression or methods requiring several

passes of cross-bilateral filtering, are still too expensive for real-time applications even using GPU implementations.

The techniques discussed in this report may well have an impact on CG movie production. Until recently, rasterization-based pipelines, such as Pixar's REYES architecture, have been the norm, with the notable exception of Blue Sky Studios, which used ray tracing on all of its productions. However, we are currently seeing an industry-wide shift to physically based Monte Carlo rendering, spearheaded by Solid Angle's Arnold renderer that is now used in a wide array of productions. Pixar is introducing a new Monte Carlo rendering pipeline, RIS, in its RenderMan software. Walt Disney Animation Studios (WDAS) developed its own Monte Carlo renderer, Hyperion, and Weta Digital recently introduced its renderer, Manuka. These renderers come in addition to existing ones, such as Mental Ray from Mental Images (now NVIDIA) and V-Ray from Chaos Group, both of which are readily integrated in professional production tools. With this shift to Monte Carlo rendering comes a strong interest for robust adaptive rendering solutions, and denoising in particular. For example, Anderson and Meyer from Pixar proposed a statistical technique based on PCA to denoise animation sequences [MA06]. Goddard, from Image Engine, recently presented the denoising technique they developed for the production of Elysium [God14]. WDAS also used denoising in the production of Big Hero 6, which proved to be crucial for reaching production deadlines, and Pixar is currently investigating various denoising alternatives for its own productions. These production environments bring new constraints and demands that stretch the capabilities of existing methods and will hopefully foster new research.

Although some of the methods described in this survey have been extended to filter animation sequences, we believe this is the area that provides biggest potential for further quality improvements. It would also be tempting to try to combine the advantages of *a priori* methods, like taking advantage of sheared structures in light fields, and *a posteriori* techniques, like the complex filter shapes. Methods for real-time rendering still lag the quality of off-line methods, and more effort should be spent on trying to close this gap. Finally, there is a variety of techniques and methodologies that remain largely unexplored in the context covered by this survey, including the use of data-driven and learning-based approaches, or general frameworks for sparse sampling and reconstruction based on work in compressive sensing and matrix completion. From a more theoretical perspective, an important open question is also whether lower bounds on the number of samples can formally be proved.

Acknowledgments

We would like to thank the following persons and institutes for making 3D objects and scenes publicly available: headus/Rezard for the Killeroo model, the Stanford 3D Scanning Repository for the dragon model, the Cornell Univer-

sity Program of Computer Graphics for the Cornell box, Guillermo M. Leal Llaguno of Evolución Visual for the San Miguel scene, Andrew Kensler and the Utah 3D Animation Repository for the toasters scene, Marko Dabrovic and Mihovil Odak for the Sibenik cathedral, and Anat Grynberg and Greg Ward for the conference room. R. Ramamoorthi acknowledges NSF grants 1115242, 1451830, the Intel Science and Technology Center for Visual Computing, and equipment donations from NVIDIA. P. Sen was funded by NSF grant 1342931. S. E. Yoon was supported in part by NRF-2013R1A1A2058052. M. Zwicker acknowledges support by the Swiss National Science foundation under grant no. 143886. Finally, we are grateful to the Monte Carlo rendering research community, in particular those working in adaptive sampling and reconstruction whose work we discuss here. They, too, have contributed to this report.

Appendix

Code for the following methods is available online (retrieved February 2015):

- Hachisuka et al. [HJW*08]: <http://www.ci.i.u-tokyo.ac.jp/~hachisuka/mdas.zip>
- Egan et al. [ETH*09]: <http://www.cs.columbia.edu/cg/mb/>
- Overbeck et al. [ODR09]: <http://www1.cs.columbia.edu/~rso2102/AWR/>
- Lehtinen et al. [LAC*11]: <http://groups.csail.mit.edu/graphics/tlfr/>
- Rousselle et al. [RKZ11]: http://www.cgg.unibe.ch/downloads/asr-auxiliary.zip/at_download/file
- Lehtinen et al. [LALD12]: <http://groups.csail.mit.edu/graphics/ilfr/>
- Li et al. [LWC12]: <http://www.cmlab.csie.ntu.edu.tw/project/sbf/>
- Mehta et al. [MWR12]: <http://graphics.berkeley.edu/papers/UdayMehta-AAF-2012-12/index.html>
- Rousselle et al. [RKZ12]: http://www.cgg.unibe.ch/downloads/nlm-code-data.zip/at_download/file
- Kalantari and Sen [KS13]: http://www.ece.ucsb.edu/~psen/PaperPages/RemovingMCNoiseStuff/RemovingMCNoise_v1.0.zip
- Moon et al. [MJL*13]: <http://sglab.kaist.ac.kr/VFL/>
- Rousselle et al. [RMZ13]: http://www.cgg.unibe.ch/downloads/pg2013_code_data.zip/at_download/file
- Moon et al. [MCY14]: <http://sglab.kaist.ac.kr/WLR/>
- Munkberg et al. [MVH*14]: <http://fileadmin.cs.lth.se/graphics/research/papers/2014/recon5d/>

References

- [BBS14] BELCOUR L., BALA K., SOLER C.: A local frequency analysis of light scattering and absorption. *ACM Trans. Graph.* 33, 5 (Aug. 2014). 5, 6
- [BCM05] BUADES A., COLL B., MOREL J.-M.: A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation* 4, 2 (2005), 490–530. 7, 8
- [BEM11] BAUSZAT P., EISEMANN M., MAGNOR M.: Guided image filtering for interactive high-quality global illumination. *Computer Graphics Forum* 30, 4 (2011), 1361–1368. 10, 11
- [BM98] BOLIN M. R., MEYER G. W.: A perceptually based adaptive sampling algorithm. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques* (1998), SIGGRAPH '98, ACM, pp. 299–309. 2, 11
- [BSS*13a] BAGHER M., SOLER C., SUBR K., BELCOUR L., HOLZSCHUCH N.: Interactive rendering of acquired materials on dynamic geometry using frequency analysis. *Visualization and Computer Graphics, IEEE Transactions on* 19, 5 (May 2013), 749–761. 3, 5
- [BSS*13b] BELCOUR L., SOLER C., SUBR K., HOLZSCHUCH N., DURAND F.: 5d covariance tracing for efficient defocus and motion blur. *ACM Trans. Graph.* 32, 3 (July 2013), 31:1–31:18. 4, 5
- [CTCS00] CHAI J.-X., TONG X., CHAN S.-C., SHUM H.-Y.: Plenoptic sampling. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (2000), SIGGRAPH '00, ACM, pp. 307–318. 3
- [DFKE07] DABOV K., FOI A., KATKOVNIK V., EGIAZARIAN K.: Image denoising by sparse 3-d transform-domain collaborative filtering. *Image Processing, IEEE Transactions on* 16, 8 (2007), 2080–2095. 7, 8
- [DHS*05] DURAND F., HOLZSCHUCH N., SOLER C., CHAN E., SILLION F. X.: A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3 (July 2005), 1115–1126. 2, 3, 4, 5
- [DJ94] DONOHO D. L., JOHNSTONE I. M.: Ideal spatial adaptation by wavelet shrinkage. *Biometrika* 81, 3 (1994), 425–455. 7, 8
- [DMB*14] DELBRACIO M., MUSÉ P., BUADES A., CHAUVIER J., PHELPS N., MOREL J.-M.: Boosting Monte Carlo rendering by ray histogram fusion. *ACM Trans. Graph.* 33, 1 (Feb. 2014), 8:1–8:15. 8, 11
- [DSHL10] DAMMERTZ H., SEWITZ D., HANIKA J., LENSCH H. P. A.: Edge-avoiding à-trous wavelet transform for fast global illumination filtering. In *Proceedings of the Conference on High Performance Graphics* (2010), HPG '10, Eurographics Association, pp. 67–75. 7, 10, 11
- [DWR10] DECORO C., WEYRICH T., RUSINKIEWICZ S.: Density-based outlier rejection in monte carlo rendering. *Computer Graphics Forum* 29, 7 (2010), 2119–2125. 11
- [EA06] ELAD M., AHARON M.: Image denoising via sparse and redundant representations over learned dictionaries. *Image Processing, IEEE Transactions on* 15, 12 (Dec 2006), 3736–3745. 7
- [ED04] EISEMANN E., DURAND F.: Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 673–678. 9, 10
- [EDR11] EGAN K., DURAND F., RAMAMOORTHI R.: Practical filtering for efficient ray-traced directional occlusion. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 180:1–180:10. 4, 5

- [EHDR11] EGAN K., HECHT F., DURAND F., RAMAMOORTHY R.: Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Trans. Graph.* 30, 2 (Apr. 2011), 9:1–9:13. 4, 5
- [ETH*09] EGAN K., TSENG Y.-T., HOLZSCHUCH N., DURAND F., RAMAMOORTHY R.: Frequency analysis and sheared reconstruction for rendering motion blur. *ACM Trans. Graph.* 28, 3 (July 2009), 93:1–93:13. 3, 4, 5, 8, 13
- [FP04] FARRUGIA J.-P., PÉROCHE B.: A progressive rendering algorithm using an adaptive perceptually based image metric. *Computer Graphics Forum* 23, 3 (Sept. 2004), 605–614. 11
- [GKDS12] GEORGIEV I., KŘIVÁNEK J., DAVIDOVIČ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 192:1–192:10. 11
- [God14] GODDARD L.: Silencing the noise on Elysium. In *ACM SIGGRAPH 2014 Talks* (New York, NY, USA, 2014), SIGGRAPH '14, ACM, pp. 38:1–38:1. 12
- [Guo98] GUO B.: Progressive radiance evaluation using directional coherence maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques* (1998), SIGGRAPH '98, ACM, pp. 255–266. 2
- [Hec86] HECKBERT P.: *Fundamentals of Texture Mapping and Image Warping*. Master's thesis, UC Berkeley, June 1986. 3
- [HJW*08] HACHISUKA T., JAROSZ W., WEISTROFFER R. P., DALE K., HUMPHREYS G., ZWICKER M., JENSEN H. W.: Multidimensional adaptive sampling and reconstruction for ray tracing. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 33:1–33:10. 7, 11, 13
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 191:1–191:10. 11
- [HST10] HE K., SUN J., TANG X.: Guided image filtering. In *Proceedings of the 11th European Conference on Computer Vision: Part I* (Berlin, Heidelberg, 2010), ECCV'10, Springer-Verlag, pp. 1–14. 9
- [HWRH13] HEIDE F., WETZSTEIN G., RASKAR R., HEIDRICH W.: Adaptive image synthesis for compressive displays. *ACM Trans. Graph.* 32, 4 (July 2013), 132:1–132:12. 11
- [JDZJ08] JAROSZ W., DONNER C., ZWICKER M., JENSEN H. W.: Radiance caching for participating media. *ACM Trans. Graph.* 27, 1 (Mar. 2008). 6
- [Jen01] JENSEN H. W.: *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001. 11
- [JSKJ12] JAROSZ W., SCHÖNEFELD V., KOBBELT L., JENSEN H. W.: Theory, analysis and applications of 2D global illumination. *ACM Trans. Graph.* 31, 5 (Sept. 2012), 125:1–125:21. 6
- [JZJ08] JAROSZ W., ZWICKER M., JENSEN H. W.: Irradiance gradients in the presence of participating media and occlusions. *Computer Graphics Forum (Proc. EGSR)* 27, 4 (June 2008), 1087–1096. 6
- [Kaj86] KAJIYA J. T.: The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug. 1986), 143–150. 1
- [KDB14] KELLER A., DAHM K., BINDER N.: Path space filtering. In *ACM SIGGRAPH 2014 Talks* (2014), SIGGRAPH '14, pp. 68:1–68:1. 11
- [KGBP05] KŘIVÁNEK J., GAUTRON P., BOUATOUCH K., PATTANAİK S.: Improved radiance gradient computation. In *Proceedings of the 21st Spring Conference on Computer Graphics* (2005), SCCG '05, pp. 155–159. 6
- [KGPB05] KŘIVÁNEK J., GAUTRON P., PATTANAİK S., BOUATOUCH K.: Radiance caching for efficient global illumination computation. *Visualization and Computer Graphics, IEEE Transactions on* 11, 5 (2005), 550–561. 6
- [KS13] KALANTARI N. K., SEN P.: Removing the noise in Monte Carlo rendering with general image denoising algorithms. *Computer Graphics Forum* 32, 2 (2013), 93–102. 8, 10, 11, 13
- [LAC*11] LEHTINEN J., AILA T., CHEN J., LAINE S., DURAND F.: Temporal light field reconstruction for rendering distribution effects. *ACM Trans. Graph.* 30, 4 (July 2011), 55:1–55:12. 5, 6, 13
- [LALD12] LEHTINEN J., AILA T., LAINE S., DURAND F.: Reconstructing the indirect light field for global illumination. *ACM Trans. Graph.* 31, 4 (July 2012), 51:1–51:10. 6, 13
- [LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (1996), SIGGRAPH '96, ACM, pp. 31–42. 3
- [LHG*09] LEVIN A., HASINOFF S. W., GREEN P., DURAND F., FREEMAN W. T.: 4d frequency analysis of computational cameras for depth of field extension. *ACM Trans. Graph.* 28, 3 (July 2009), 97:1–97:14. 3
- [LWC12] LI T.-M., WU Y.-T., CHUANG Y.-Y.: Sure-based optimization for adaptive sampling and reconstruction. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 194:1–194:9. 9, 10, 11, 13
- [MA06] MEYER M., ANDERSON J.: Statistical acceleration for animated global illumination. *ACM Trans. Graph.* 25, 3 (July 2006), 1075–1080. 12
- [McC99] MCCOOL M. D.: Anisotropic diffusion for Monte Carlo noise reduction. *ACM Trans. Graph.* 18, 2 (Apr. 1999), 171–194. 2, 8, 11
- [MCY14] MOON B., CARR N., YOON S.-E.: Adaptive rendering based on weighted local regression. *ACM Trans. Graph.* 33, 5 (Sept. 2014), 170:1–170:14. 10, 11, 13
- [Mit87] MITCHELL D. P.: Generating antialiased images at low sampling densities. *SIGGRAPH Comput. Graph.* 21, 4 (Aug. 1987), 65–72. 2, 7, 11
- [Mit91] MITCHELL D. P.: Spectrally optimal sampling for distribution ray tracing. *SIGGRAPH Comput. Graph.* 25, 4 (July 1991), 157–164. 2
- [MJL*13] MOON B., JUN J. Y., LEE J., KIM K., HACHISUKA T., YOON S.-E.: Robust image denoising using a virtual flash image for Monte Carlo ray tracing. *Computer Graphics Forum* 32, 1 (2013), 139–151. 10, 13
- [MVH*14] MUNKBERG J., VAIDYANATHAN K., HASSELGREN J., CLARBERG P., AKENINE-MÖLLER T.: Layered reconstruction for defocus and motion blur. *Computer Graphics Forum* 33, 4 (2014), 81–92. 5, 13
- [MWR12] MEHTA S. U., WANG B., RAMAMOORTHY R.: Axis-aligned filtering for interactive sampled soft shadows. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 163:1–163:10. 4, 5, 13
- [MWRD13] MEHTA S. U., WANG B., RAMAMOORTHY R., DURAND F.: Axis-aligned filtering for interactive physically-based diffuse indirect lighting. *ACM Trans. Graph.* 32, 4 (July 2013), 96:1–96:12. 4, 5
- [MYRD14] MEHTA S. U., YAO J., RAMAMOORTHY R., DURAND F.: Factored axis-aligned filtering for rendering multiple distribution effects. *ACM Trans. Graph.* 33, 4 (July 2014), 57:1–57:12. 4, 5, 9, 10
- [ODR09] OVERBECK R. S., DONNER C., RAMAMOORTHY R.: Adaptive wavelet rendering. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 140:1–140:12. 2, 7, 8, 11, 13

- [OS09] OPPENHEIM A. V., SCHAFER R. W.: *Discrete-Time Signal Processing*, 3rd ed. Prentice Hall, 2009. 2
- [PF14] PERRONE D., FAVARO P.: Total variation blind deconvolution: The devil is in the details. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on* (June 2014), pp. 2909–2916. 11
- [PH10] PHARR M., HUMPHREYS G.: *Physically Based Rendering: From Theory To Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2010. 9
- [PMKY13] PARK H., MOON B., KIM S., YOON S.-E.: P-RPF: Pixel-based random parameter filtering for Monte Carlo rendering. In *Computer-Aided Design and Computer Graphics (CAD/Graphics), 2013 International Conference on* (Nov 2013), pp. 123–130. 9
- [PS89] PAINTER J., SLOAN K.: Antialiased ray tracing by adaptive progressive refinement. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 281–288. 2
- [PSA*04] PETSCHNIG G., SZELISKI R., AGRAWALA M., COHEN M., HOPPE H., TOYAMA K.: Digital photography with flash and no-flash image pairs. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 664–672. 9, 10
- [RKZ11] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive sampling and reconstruction using greedy error minimization. *ACM Trans. Graph.* 30, 6 (Dec. 2011), 159:1–159:12. 7, 11, 13
- [RKZ12] ROUSSELLE F., KNAUS C., ZWICKER M.: Adaptive rendering with non-local means filtering. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 195:1–195:11. 8, 11, 13
- [RMB07] RAMAMOORTHI R., MAHAJAN D., BELHUMEUR P.: A first-order analysis of lighting, shading, and shadows. *ACM Trans. Graph.* 26, 1 (Jan. 2007). 6
- [RMZ13] ROUSSELLE F., MANZI M., ZWICKER M.: Robust denoising using feature and color information. *Computer Graphics Forum* 32, 7 (2013), 121–130. 9, 10, 11, 13
- [RPG99] RAMASUBRAMANIAN M., PATTANAİK S. N., GREENBERG D. P.: A perceptually based physical error metric for realistic image synthesis. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (1999), SIGGRAPH '99, ACM, pp. 73–82. 2, 11
- [RW94] RUSHMEIER H. E., WARD G. J.: Energy preserving non-linear filters. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), SIGGRAPH '94, ACM, pp. 131–138. 2, 11
- [SD11] SEN P., DARABI S.: Compressive rendering: A rendering application of compressed sensing. *Visualization and Computer Graphics, IEEE Transactions on* 17, 4 (April 2011), 487–499. 11
- [SD12] SEN P., DARABI S.: On filtering the noise from the random parameters in Monte Carlo rendering. *ACM Trans. Graph.* 31, 3 (June 2012), 18:1–18:15. 9, 11
- [Shi93] SHINYA M.: Spatial anti-aliasing for animation sequences with spatio-temporal filtering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, ACM, pp. 289–296. 3
- [SJJ12] SCHWARZHAUPT J., JENSEN H. W., JAROSZ W.: Practical Hessian-based error control for irradiance caching. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 193:1–193:10. 6
- [SSD*09] SOLER C., SUBR K., DURAND F., HOLZSCHUCH N., SILLION F.: Fourier depth of field. *ACM Trans. Graph.* 28, 2 (May 2009), 18:1–18:12. 4, 5
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Computer Vision, 1998. Sixth International Conference on* (1998), pp. 839–846. 8
- [VDVK09] VAN DE VILLE D., KOCHER M.: Sure-based non-local means. *Signal Processing Letters, IEEE* 16, 11 (2009), 973–976. 9
- [WABG06] WALTER B., ARBREE A., BALA K., GREENBERG D. P.: Multidimensional lightcuts. *ACM Trans. Graph.* 25, 3 (2006), 1081–1088. 11
- [WFA*05] WALTER B., FERNANDEZ S., ARBREE A., BALA K., DONIKIAN M., GREENBERG D. P.: Lightcuts: a scalable approach to illumination. *ACM Trans. Graph.* 24, 3 (2005), 1098–1107. 11
- [WH92] WARD G. J., HECKBERT P.: Irradiance gradients. In *Third Eurographics Workshop on Rendering* (1992), vol. 8598, Alan Chalmers and Derek Paddon. 6
- [WLHR11] WETZSTEIN G., LANMAN D., HEIDRICH W., RASKAR R.: Layered 3d: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Trans. Graph.* 30, 4 (July 2011), 95:1–95:12. 3
- [WRC88] WARD G. J., RUBINSTEIN F. M., CLEAR R. D.: A ray tracing solution for diffuse interreflection. *SIGGRAPH Comput. Graph.* 22, 4 (June 1988), 85–92. 2, 6
- [XP05] XU R., PATTANAİK S.: A novel Monte Carlo noise reduction operator. *Computer Graphics and Applications, IEEE* 25, 2 (2005), 31–35. 8, 11
- [ZMDP06] ZWICKER M., MATUSIK W., DURAND F., PFISTER H.: Antialiasing for Automultiscopic 3D Displays. In *Proceedings of Eurographics Symposium on Rendering* (2006), pp. 73–82. 3