

A Framework for Developing and Benchmarking Sampling and Denoising Algorithms for Monte Carlo Rendering

Jonas Deyson B. Santos · Pradeep Sen · Manuel M. Oliveira

Abstract Although many adaptive sampling and reconstruction techniques have been proposed in the last few years, the case for which one should be used for a specific scene is still to be made. Moreover, developing a new technique has required selecting a particular rendering system, which makes the technique tightly coupled to the chosen renderer and limits the amount of scenes it can be tested on. In this paper, we propose a renderer-agnostic framework for testing and benchmarking sampling and denoising techniques for Monte Carlo rendering. It decouples techniques from rendering systems by hiding the renderer details behind an API. This improves productivity and allows for direct comparisons among techniques originally developed for different rendering systems. We demonstrate its effectiveness by using our API to instrument four rendering systems, most state-of-the-art Monte Carlo denoising techniques and three sampling strategies, and by performing a benchmark across rendering systems.

Keywords Monte Carlo Rendering, Adaptive Sampling and Reconstruction, Denoising, Benchmarking.

This work was sponsored by CAPES and CNPq-Brazil (fellowships and grants 306196/2014-0 and 423673/2016-5), as well as US National Science Foundation grants IIS-1321168 and IIS-1619376.

Jonas Deyson B. Santos
Instituto de Informática, UFRGS, Porto Alegre, RS, Brazil
E-mail: jdbsantos at inf.ufrgs.br

Pradeep Sen
University of California, Santa Barbara
E-mail: psen at ece.ucsb.edu

Manuel M. Oliveira
Instituto de Informática, UFRGS, Porto Alegre, RS, Brazil
E-mail: oliveira at inf.ufrgs.br

1 Introduction

Rendering is one of the most important problems in computer graphics and has been the subject of over half a century of research. In particular, there has been a tremendous amount of exploration on Monte Carlo (MC) physically-based rendering systems [8] such as path-tracing [15] and its various extensions [38]. To address shortcomings in Monte Carlo rendering, more than three decades of research has explored a wide variety of different ideas, including adaptive sampling and reconstruction algorithms [12], faster acceleration structures and intersection algorithms [33], improved sampling patterns [13], and Monte Carlo denoisers [36, 31, 16], to name a few broad categories.

Although developing a new algorithm that successfully improves Monte Carlo rendering in some way is a challenging task in itself, researchers face two further challenges. First, *they must integrate their algorithm into an actual rendering system* so they can test it on complex scenes. After all, renderers have several key components required to produce high-quality images (*e.g.*, scene I/O, samplers, ray-traversal acceleration data structures, primitive-ray intersectors, shading systems, and reconstruction filters), and many of these components are often orthogonal to the algorithm being explored. Therefore, rendering researchers often leverage the infrastructure provided by existing rendering systems. However, integrating a new algorithm into a rendering system is often a time-consuming task, precluding its deployment on multiple renderers to properly test the technique.

Second, *researchers must find several high-quality scenes to test their algorithm and demonstrate their performance*. Since most rendering researchers are not digital artists, constructing complex aesthetically pleas-

ing scenes is often a non-trivial, time-consuming task, and “programmer-art” scenes do not tend to be of the same quality as those constructed by professional artists. Moreover, rendering systems tend to adopt proprietary scene-description formats. Thus, researchers tend to stick to a handful of test scenes available for the selected rendering system.

The consequence of these two challenges is that most researchers are often only able to demonstrate their algorithm on a single rendering system using a small number of test scenes. This significantly limits their ability to thoroughly test and explore the proposed method, and for reviewers to properly evaluate its performance. Furthermore, it is often difficult to compare against existing methods, since they often have been implemented in other rendering systems and tested on different scenes. Having good “apples-to-apples” comparisons is important when trying to gauge the benefits of a new method.

Finally, porting a recently published algorithm to a new rendering system is not easy, since the developers performing the port are usually not experts on the new algorithm, even though they may be very familiar with the target rendering system. Therefore, they usually have to *translate* the available implementation (or the algorithm described in the paper) to the rendering system. This can introduce bugs in the process and may not produce ideal results, since the algorithmic parameters that worked successfully for one rendering system might not work for the new one. Trying to determine the optimal parameters for an algorithm that one did not develop can be a very time consuming task.

To address these problems, *we present a novel framework that allows researchers to develop, test, compare, and even deploy sampling and denoising algorithms for Monte Carlo rendering.* Specifically, we propose an application program interface (API) that allows developers to easily port their algorithms to different rendering systems by providing the necessary communication between such algorithms and the other components of an existing rendering system. In other words, instead of having the researchers port their algorithms to multiple rendering systems, we have done the leg work for them by instrumenting rendering systems to provide the necessary services through our API.

Therefore, a researcher only needs to implement an algorithm once, and can immediately use it with all rendering systems that support our framework. This allows researchers to rapidly test and deploy their algorithms on a range of rendering systems, and test them on a wide variety of scenes. This allows for automatic independent benchmarking of algorithms, which is quite useful when submitting new techniques for publication.

As a proof of concept, we have initially instrumented four rendering systems (PBRT-v3, PBRT-v2, Mitsuba, and a procedural renderer), pretty much all state-of-the-art MC denoising algorithms, and three sampling techniques. We plan to open-source our framework so that other Monte Carlo renderers can support the API directly themselves. This will also allow third-party rendering systems to rapidly adopt recently proposed algorithms that conform to our API.

To demonstrate the effectiveness of our framework, we conduct a case study involving Monte Carlo (MC) denoising algorithms. Such a study illustrates key aspects of our system: provide easy integration of algorithms and rendering systems (by means of just a few calls to the API); provide an independent benchmark for MC techniques that works across various rendering systems; and, allow developers to evaluate the performance of rendering systems with various algorithms, and vice versa. These are desirable features for algorithm and rendering system developers, as well as for the academic, industry, and end-user communities, who should be able to make better informed decisions when choosing a technique and/or a rendering system to render a given scene.

For our study, we have instrumented essentially all state-of-the-art MC denoising algorithms (*e.g.*, NFOR [6], LBF [16], RHF [9], LWR [21], RDFC [31], RPF [36], SBF [19], NLM [30], and GEM [29]), allowing them to be used with the four rendering systems, even though most of these algorithms have originally been developed for a single renderer. Furthermore, our system’s ability to automatically generate benchmark reports allows for the comparison of the different methods on an even playing field. In our study, we compare the performance of different MC denoising methods and discuss some of their identified potential limitations.

Although this paper does not propose a new MC rendering algorithm *per se*, this kind of *meta-research* system (*i.e.*, a system designed to aid the research process) is not new to the graphics and vision communities. Successful examples include the Middlebury benchmark [34, 35], which has transformed the way two-frame dense stereo correspondence algorithms are developed and compared, as well as the benchmarks on Alpha Matting [28, 27], optical flow [4, 3, 2], and video matting [11, 10]. More recently, Anderson et al. [1] proposed a framework to compile PDF sampling patterns for Monte Carlo.

Inspired by these works, our system provides test scenes intended to stress the limits of Monte Carlo techniques and reveal their potential limitations. It is extensible, allowing for easy support of new rendering systems, as well as sampling and denoising strategies. The community should be able to contribute new scenes

and techniques in a simple way. Our system is publicly available through our project website, providing valuable feedback to the research and industry communities.

In summary, the **contributions** of our work include:

- A framework for developing, testing, and benchmarking sampling and denoising Monte Carlo algorithms (Section 3). Our framework decouples the algorithms from rendering systems by means of an API, allowing researchers to implement their techniques once and run them on any rendering system supporting our framework. It easily incorporates new algorithms, rendering systems, and testing datasets;
- An automatic and independent benchmarking system for comparing Monte Carlo algorithms across multiple rendering systems and supporting a large number of scenes (Section 3). This should be a useful tool for assessing the quality of new Monte Carlo algorithms against established ones, especially for submission purposes;
- A detailed evaluation of the state-of-the-art Monte Carlo denoising algorithms using our framework and a discussion of their performance and limitations (Section 4).

While the use of an API might reduce the performance of an application, a careful design of the API minimizes such an impact. Nevertheless, the benefits provided by our framework highly supersede a potential performance reduction, specially in off-line rendering environments. Once tested on different rendering systems and on a variety of scenes, one can decide to provide native implementations for specific rendering systems.

2 Related Work

We begin by discussing meta-research systems in both graphics and vision which, like our own framework, have been developed to facilitate/improve the research process. Afterwards, we focus on previous work on Monte Carlo denoising, which is the application that we use in our case study to illustrate the benefits of our framework.

2.1 Meta-Research in Graphics

Several systems have been proposed over the years to facilitate research development in graphics. Some of the most popular ones include Cg [20], Brook [7], and Halide [25]. Cg is a general-purpose, hardware-oriented, programming language and supporting system designed for the development of efficient GPU applications, and providing easy integration with the two major 3D graphics APIs (OpenGL and Direct 3D). Brook [7] is also

a system for general-purpose computation that allows developers to use programmable GPUs as streaming co-processors, while abstracting GPU architectural details. Halide [25] tries to optimize image-processing algorithms by decoupling the algorithm’s description from its schedule. This allows for an algorithm to be described once, while specific schedules are provided for different target platforms (*e.g.*, CPUs, GPUs, mobile devices, etc.). Automatic generation of optimized schedules in Halide has been addressed in a follow-up work [22].

While the primary goal of these systems is to generate efficient code while abstracting hardware details from developers, our focus is on decoupling Monte Carlo algorithms from rendering systems. This greatly simplifies the task of porting algorithms to multiple rendering systems, freeing developers from the burden of knowing implementation details of specific renderers to be able to perform integration. Our system also makes a wider range of scenes available for testing, providing a comprehensive, multi-rendering system benchmark for Monte Carlo algorithms. Recently, Anderson et al. [1] proposed an approach to compile sampling BRDFs for MC applications. Their method complements our work.

2.2 Benchmarking Systems in Computer Vision

Quantitative benchmarks have been proposed for several computer vision areas, including optical flow [4, 3], dense two-frame stereo correspondence [34], and alpha matting [28]. These initiatives have provided independent tools for assessing the quality of the results produced by existing and new algorithms, and have led to significant progress in these areas.

Optical Flow – Barron et al. [4] compared accuracy, reliability, and density of velocity measurements for several established optical flow algorithms, and showed that their performance could vary significantly from one technique to another. Baker et al. [3] proposed another benchmark for optical-flow algorithms that considers aspects not covered by Barron et al. These include sequences containing non-rigid motion, realistic synthetic images, high frame-rate video to study interpolation errors, and modified stereo sequences of static scenes. The authors have made their datasets and evaluation results publicly available, and provide the option for one to submit his own results for evaluation [2]

Stereo Correspondence – The Middlebury benchmark [34] provided a taxonomy and evaluation for dense two-frame stereo correspondence algorithms. The datasets and evaluation are publicly available on the web, and anyone can submit results for evaluation [35].

Alpha Matting – Rhemann et al. [28] introduced a benchmark system for alpha matting techniques. The authors provide some training data and use a test dataset for which the corresponding ground truth has not been disclosed. Similarly to the optical-flow and dense stereo correspondence benchmarks mentioned before, the results are available on-line, and anyone can submit results for evaluation [27].

Video Matting – Erofeev et al. [11] extended the alpha matting benchmark to videos, supporting both objective and subjective evaluations of video matting techniques. Training and test datasets are provided, with results and submissions being available through the web [10].

Unlike such systems, ours goes beyond rating submitted results computed off-line. It provides an API that allows Monte Carlo algorithms to be tested with different rendering systems using a variety of scenes. Thus, it can compare different techniques across multiple rendering systems, something that was not previously possible without requiring the developer to create multiple implementations tailored to individual rendering systems.

2.3 Monte Carlo Denoising Algorithms

Although there has been a significant amount of work on reducing the variance of MC rendered images through sampling/reconstruction (see [24, 39]), for brevity we shall only focus on previous post-processing approaches that filter *general* Monte Carlo noise (*i.e.*, noise from any and all distributed effects, path tracing, and so on).

Soon after the seminal paper by Cook et al. [8] raised the problem of MC noise, there was some early work in general MC filtering, including approaches using nonlinear median and alpha-trimmed mean filters for edge-aware spike removal [18] and variable-width filter kernels to preserve energy and salient details [32]. However, in the years that followed, researchers largely ignored general MC filtering algorithms in favor of other variance reduction techniques, due to the inability of these filters to successfully remove the MC noise while preserving scene detail.

Recently, interest in general MC filtering algorithms has enjoyed a significant revival. For example, Sen and Darabi [36] demonstrated that filters could effectively distinguish between noisy scene detail and MC noise. To do this, they used mutual information to determine dependencies between random parameters and scene features, and combined these dependencies to weight a cross-bilateral filter at each pixel in the image. Rousselle et al. [30] proposed to use a non-local means filter to remove general MC noise. Kalantari and Sen [17] applied median absolute deviation to estimate the noise

level at every pixel to use any image denoising technique for filtering the MC noise. Finally, Delbracio et al. [9] modified the non-local means filter to use the color histograms of patches, rather than the noisy color patches, in the distance function.

Other approaches have effectively used error estimation for filtering general distributed effects. For example, Rousselle et al. [29] used error estimates to select different filter scales for every pixel to minimize reconstruction error. Furthermore, Li et al. [19] proposed to use Stein’s unbiased risk estimator (SURE) [37] to select the best parameter for the spatial term of a cross-bilateral filter. Rousselle et al. [31] extended this idea to apply the SURE metric to choose the best of three candidate filters. Moon et al. [21] estimated the error for discrete sets of filter parameters using a weighted local regression. Bauszat et al. [5] posed the filter selection problem as an optimization and solved it with graph cuts. More recently, Kalantari et al. [16] introduced a machine learning approach in which a neural network is used to drive the MC filter.

All of these techniques have strengths and weaknesses in terms of the scene features they can satisfactorily handle, memory costs, execution time, etc. All these variables make a direct comparison of the various algorithms difficult. Our framework is intended to fill-in this gap. Hopefully, it will help developers better understand the interplay among the various involved elements and available metrics, shedding some light on the occasional situations in which publications seem to disagree about the quality rank of different techniques.

3 System Design

A physically-based rendering system has to perform several tasks in order to generate an image. These include, for instance, read the scene description file, build the internal scene representation data structure, generate well-distributed sampling positions in a high-dimensional space, compute shading, compute global illumination, reconstruct the final image, and save the result. Several of those tasks hide a significant amount of complexity.

A good renderer implementation usually employs design practices that allow some level of extensibility. For example, it is common practice to facilitate adding new materials, shapes, cameras, samplers, reconstruction filters, and integrators, given they obey pre-defined interfaces. However, although different renderers use similar abstractions, implementing a new technique (*e.g.*, a new denoising filter) still requires choosing a particular renderer (*e.g.*, PBRT-v3 or Mitsuba). *This makes it hard to compare techniques implemented in different*

systems, and time consuming to implement a technique in multiple rendering systems.

Our framework seeks to avoid these limitations by decoupling the implementation of a new technique from any specific rendering system. For this, it hides sample value computation details associated to individual rendering systems behind a general sampling evaluation interface. Thus, it allows for any technique to be seamlessly integrated with different rendering systems, and provides a direct and simple mechanism for comparing techniques' results on multiple (rendering) systems.

Denoising techniques are responsible for reconstructing images from sample values computed by the renderer. As such, they are a key component of all Monte Carlo rendering systems. Thus, we have chosen MC denoising algorithms to demonstrate the effectiveness of our framework. One should note, however, that our system can be used to implement/evaluate other MC techniques, such as general sampling and reconstruction algorithms. In our case study, we consider both adaptive and non-adaptive MC denoising techniques. Next, we describe the main components of our system and discuss how they are used to support the integration of techniques and rendering systems, and to perform benchmarks.

3.1 Main Components

The architecture of the proposed system has three main components (Figure 1): the *client process*, the *benchmark process*, and the *renderer process*. The client process implements the *technique* being integrated. The renderer process interfaces with the actual rendering system to evaluate the samples requested by the client process. The benchmark process controls the overall execution, and saves the final image along with useful metadata information, such as reconstruction and rendering times.

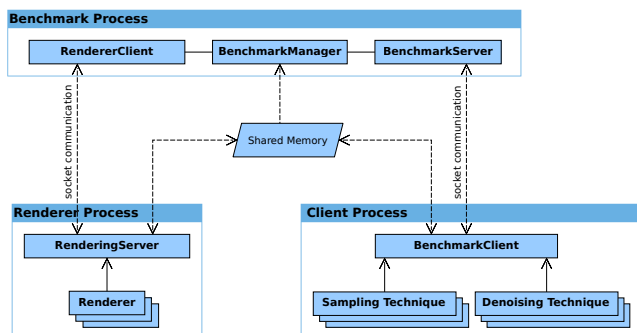


Fig. 1 Main components of the system.

Figure 2 shows a sequence diagram of a typical execution of the system. The benchmark process is executed,

receiving a list of scenes to be rendered, each scene with a list of different sample budgets. For a given scene, the renderer process is executed. The benchmark process forwards the client's request to the renderer process and keeps track of the execution time and sample budget limits. Once the client process receives the requested samples, it reconstructs the final image and sends it to the benchmark process. The cycle can start again with another request (*e.g.*, for a different sample density for the same scene, a different scene, a different technique, etc.). Note that we use the expression “benchmark process” to refer to this intermediate layer regardless of the system being used to locally evaluate a single or multiple techniques, or to perform or complement an actual full benchmark. The system also provides a web-based graphical user interface (GUI) for visual exploration of the results.

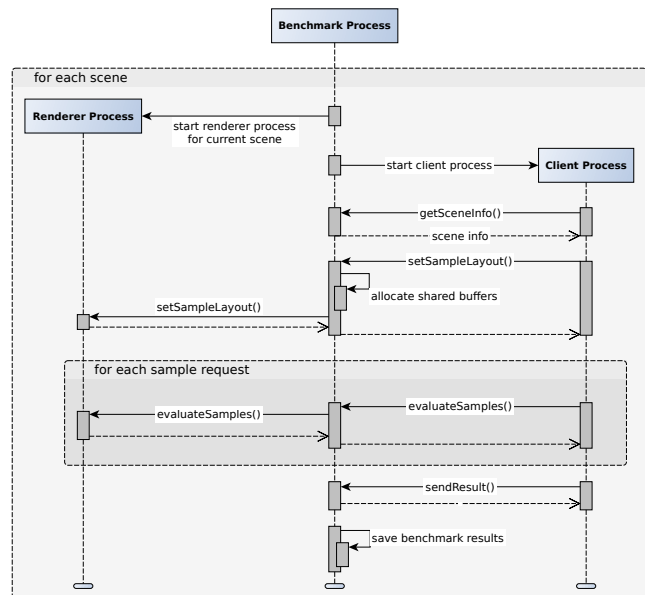


Fig. 2 A typical execution of the system.

The separation between client and renderer processes allows us to provide a clean API to the client process, simplifying the task of implementing a new technique. Once a new technique is implemented using this renderer-agnostic API, it can be readily tested on a variety of scenes and compared against other techniques.

On the renderer side, this separation allows us to provide different renderers as back-ends to the system. When rendering a scene, the client process does not need to know the specific renderer being used. This also tests the robustness of the technique to variations in sample values computed by different renderers.

3.1.1 Client Process

The system expects techniques to follow the template shown in Figure 3. Such a flow is general enough to cover a large variety of techniques, including MC denoising — adaptive, non-adaptive, a priori and a posteriori [39] — as well as sampling techniques. If a particular technique does not provide sampling positions, the renderer transparently supplies them.




Fig. 3 Template for techniques supported by our system.

When the client process starts, it is given a sample budget. In the *initial sampling* step (Figure 3), the technique decides what portion of the sample budget to spend initially. If the technique is non-adaptive, the entire budget is spent in this step. Otherwise, one or more iterations of *sample analysis* and *adaptive sampling* are performed, until the sample budget is completely consumed. After the final image is reconstructed, the client process finishes. Besides the sample budget, the client has access to more information about the scene through a scene information querying API. This information allows the technique to adjust its parameters depending on the characteristics of the scene.

Our framework is general enough to support advanced techniques with adaptive sampling, allowing them to generate sample positions based on information from previous iterations. If the technique does not perform adaptive sampling, the renderer itself generates the sample positions. The API also allows the technique to specify which features it needs for each sample. Some may require only color information, while others may require geometric features like normals, depth, etc. The technique also specifies the exact layout of sample components in memory.

3.1.2 Benchmark Process

The benchmark process manages the system execution and mediates the communication between client and renderer processes (Figure 2). It is responsible for starting the renderer process, providing information about the sample layout and additional rendering parameters, and later collecting the computed samples to be forwarded to the client technique. The benchmark process also keeps track of the current sample budget, client process

execution time, and saves the image reconstructed by the client process along with an execution log.

3.1.3 Renderer Process

It consists of a common rendering system that has been instrumented to communicate with the benchmark process. It is responsible for computing the samples needed by the client process, as well as providing information about the current loaded scene. To help instrumenting existing rendering systems, we provide a few auxiliary classes that implement the necessary API and help collecting the sample data throughout the system.

3.2 Scenes

Our system includes two general categories of scenes: *production*, and *experimental*. The first category includes scenes one would usually find in a production environment (*e.g.*, most scenes shown in Figure 10). They usually contain more detailed geometry and textures, a bigger variety of illumination settings, and aesthetically pleasing results. The second category includes scenes designed specifically to stress certain aspects of the filters. Figure 6 shows examples of experimental scenes. By including a variety of scenes in both categories, we hope to avoid biases when comparing different techniques.

When evaluating a scene, we consider two main aspects: *features* and *noise sources*. Features are legitimate details that denoising techniques must preserve, like textures and materials, geometric details, shading highlights, etc. Noise sources are elements that introduce undesired noise artifacts, like camera effects (motion blur and depth-of-field), glossy materials, area lights, and indirect illumination.

3.3 Implementation Details

Instrumenting additional rendering systems for use with our framework only requires implementing the endpoints needed to communicate with the benchmark process. We provide an auxiliary class called `RenderingServer`, which implements the communication protocol and exposes a higher level API using a signal-slot mechanism. The `RenderingServer` class and a few other auxiliary classes make it easy to instrument a renderer.

Synchronization, control messages, and small messages are implemented using TCP socket messages on a predefined local port. Large buffers use shared memory, saving memory and avoiding transfer overhead. They are used to transfer samples to the client, and the reconstructed image to the benchmark process (Figure 1).

4 Results

We have implemented our framework in C++. As a proof of concept, we have ported three well-known renderers, PBRT-v2 [23], PBRT-v3 [24], and Mitsuba [14], plus a procedural renderer to work as back-ends of our system. We have also adapted essentially all state-of-the-art MC denoising methods: LWR [26], NFOR [6], LBF [16], RPF [36], SBF [19], RHF [9], NLM [30], RDFC [31], and GEM [29]. For this, we have instrumented the original source code provided by the rendering systems’ developers and by the authors of these techniques with calls to our API. In the case of NFOR, we could not get the source code and implemented it from scratch. All results shown in the paper were generated on a 4 GHz i7-4790K CPU with 32 GB of RAM. This section provides several examples illustrating the use of the four rendering systems and eight state-of-the-art MC denoising algorithms. We also demonstrate the support for sampling techniques by adapting three commonly used ones: stratified, Sobol, and low discrepancy.

Some techniques use geometric features from the first intersection point to help them preserve scene details. This strategy tends to perform poorly on scenes with transparent glass and mirrors, as shown in Figure 4 (center). To make comparisons among techniques fairer, we implemented modified versions of these techniques using the first non-specular intersection point instead. We indicate the modified versions by a suffix “-mf” (modified features) — Figure 4 (right).



Fig. 4 Rendering using geometric features. Reference image (left). Overblurring on transmitted scene details caused by relying on features at the first intersection point (center). Using features from the first non-specular intersection allows the denoiser to preserve those details (right).

Figure 10 shows results of a benchmark created with seven MC denoising techniques and nine scenes from our scene pool. The scenes were selected as to form a representative set of situations that can challenge a denoiser. *Measure One* contains several glossy highlights, and *Measure One Moving* adds motion blur on top of that. *Crown in Glass* contains intricate bumpy textures with sources of caustics, all behind a layer of glass. *Furry Bunny* and *Curly Hair* contain fine geometric features that can easily be overblurred. *Bathroom* is a typical

interior scene with several fine textures reflected by mirrors. *Country Kitchen Night* is a challenging global illumination scene with hidden light sources, being prone to fireflies (artifacts consisting of bright single pixels scattered over the image). Finally, *Glass of Water* is a mostly specular scene with many specular highlights.

The first row of Figure 10 shows thumbnails of the ground truth images for the selected scenes. Although the image resolutions vary, their typical size is about $1,024 \times 1,024$ pixels. A small square highlights a challenging region in each scene. The corresponding regions for the noisy result, for the outputs generated by the various denoising techniques, and for the reference images are shown in the subsequent rows. From the scenes shown in Figure 10, *Bathroom*, *Glass of Water*, and *Country Kitchen Night* were rendered using Mitsuba; the remaining six were rendered using PBRT-v3. Figures 5 and 6 show examples of images generated with our framework using PBRT-v2 and a procedural renderer, respectively.

Our system can be used with and provides support for testing and comparing different sampling strategies. Figure 7 shows a scene rendered using our framework with three sampling techniques: stratified, Sobol, and low discrepancy. The images were generated using PBRT-v2, with 64 samples per pixel. The reference image was rendered using 811,008 spp.




Fig. 5 Images generated with our system using PBRT-v2 and the techniques NFOR, RHF, and GEM, respectively.




Fig. 6 Examples of experimental scenes rendered with our system using a procedural renderer. (left) Mandelbrot set. (right) Increasing sinusoidal bands ($\sin(x^2)$).

The results in Figure 10 show that all techniques have some degree of trouble with glossy highlights, as shown in the scene *Measure One*. The glossy highlights are often overblurred or contain patchy artifacts. Glossy




Fig. 7 Examples of images produced by different sampling techniques using our system with the PBRT-v2 (top). The reference image was generated with 811,008 spp, while the stratified, Sobol, and low-discrepancy images used 64 spp. For comparison, the bottom row shows zoomed-in versions of the highlighted regions shown on top.

highlights are troublesome because the extra features used by the denoisers to tell legitimate scene details from noise do not help detecting the highlights. Another instance of this problem can be seen in Figure 8. The subtle checker patterns seen on the reference image (Figure 8 (bottom right)) come from a texture applied to the specular component of the material. This specular component is not part of the albedo feature used by denoisers, causing them to remove the detail.

Back to Figure 10, scene *Measure One Moving* is a motion blur version of the previous scene. The strong motion blur effect makes the overblurring of the glossy highlights less visible, but it may also lead to other situations that may cause denoisers to produce overblur. All techniques have trouble preserving the fine motion blur details over the noisy glossy background.

The *Crown in Glass* scene contains bump-mapping details behind a layer of glass. Some techniques do a good job at preserving these details on the less noisy areas (e.g., NLM and RDFC). In darker, noisier regions, all techniques introduce some degree of overblurring.

Very fine geometry details, as commonly found in hair (*Curly Hair*) and fur (*Furry Bunny*) is also a frequent source of problems. Notice that even denoisers that rely on geometric features, as in the case with LBF, can overblur these details — although hair and fur are being captured by the geometric features, the sub-pixel-level detail in the presence of noise constitutes a challenge.

Scenes with very challenging illumination conditions — which translates to high levels of noise — are also problematic. High-energy spikes (fireflies) are very difficult to spread out while preserving energy, causing blob artifacts. As the *Country Kitchen Night* scene example shows, some techniques like *RDFC* do a good job at

spreading fireflies, but small variations in the geometry of the scene can cause artifacts.

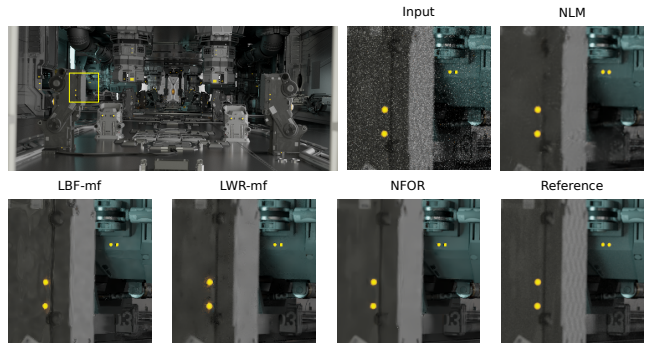


Fig. 8 Texture details in the specular component of some materials (see the checker pattern on the light gray rectangle in the reference image) are not part of the “albedo” feature, making the denoisers to remove such details.

The results shown in Figure 10 and the previous discussion illustrate the potential of our framework to provide qualitative assessments of MC denoising techniques, as well as to identify potential limitations of current approaches. As such, our system provides relevant information for guiding future research in the area. Our framework also contains a GUI for interactive exploration of benchmark results, which include quantitative assessments based on several metrics (MSE, rMSE, PSNR, SSIM, and execution time – Figure 9). Table 1 shows the values of rMSE, PSNR, and SSIM for all examples in Figure 10.

The proposed system is available for download in our project website¹. The interactive version of the benchmark results corresponding to Figure 10, using our web-based GUI, is also available. We would like to encourage the reader to explore such material. A video providing a brief tutorial on how to interactively explore the benchmark results can be found in the project website. Our GUI can be used with the results obtained for any technique that uses our framework. Figure 9 shows a snapshot of some of the quantitative information obtained when using our framework to render the *Bathroom* scene using the seven denoising techniques shown in Figure 10. The graphs compare the performance of the techniques according to PSNR, rMSE, SSIM, and execution time for 16, 32, 64, and 128 spp.

4.1 Discussion

Communication Overhead – It can be significant depending on how a technique requests samples. If the

¹ http://www.inf.ufrgs.br/~oliveira/projects/FBKSD/FBKSD_page.html

5 Conclusion and Future Work

We have presented a novel framework for easy evaluation and development of sampling and denoising MC algorithms on multiple rendering systems. Conversely, it also allows for rendering systems to quickly incorporate new algorithm that conform to our API. This makes it straightforward to perform benchmarks involving various algorithms across different renderers. We have demonstrated the effectiveness of our system by instrumenting four rendering systems (PBRT-v3, PBRT-v2, Mitsuba, and a procedural renderer), eight state-of-the-art MC denoising algorithms, three sampling techniques, and by benchmarking these denoising algorithms on multiple renderers.

We have provided a qualitative assessment of the evaluated MC denoising techniques, identifying potential limitations of existing approaches. This information might guide future research in the area. The visual exploration of the quantitative data collected during the benchmark also provides valuable feedback for researchers and users, helping them to address the practical question of identifying the most effective techniques for rendering scenes with a given set of features.

5.1 Future Work

We plan on releasing an on-line benchmarking service to allow researchers to submit their techniques for evaluation and ranking. This would be similar to other benchmark services, such as the Middlebury [34].

To solve the current memory overhead problem, we plan to group samples into smaller-sized chunks and send them to the client process as soon as they are computed by the renderer. This should increase the communication overhead but will make the memory complexity independent of the amount of requested samples, which is a fair tradeoff. We will make such an implementation available at the project website.

Finally, we intend to provide automatic conversion of scene file formats among renderers. This would complete our vision of making the choices of technique, rendering system, and scene all orthogonal to each other.

Acknowledgements We would like to thank Matt Pharr, Greg Humphreys, and Wenzel Jakob for making the source code of PBRT-v2, PBRT-v3, and Mitsuba publicly available. We also thank the authors of the following techniques for kindly providing their source code: LBF, RHF, LWR, RDFC, RPF, SBF, NLM, GEM. The following individuals and institutions provided the scenes used in the paper: Martin Lubich (Crown), Cem Yuksel (Curly Hair), Jesper Lloyd (Toy Gyro), Wojciech Jarosz (Chess), Andrew Kensler (Toasters), Bernhard Vogl and Stanford CG Lab (Furry Bunny), Marek (Bathroom), aXel (Glass of Water), Jay-Artist (Country Kitchen),

Beeple (Measure One), Anat Grynberg and GregWard (Conference), Duc Nguyen, Ron Fedkiw, and Nolan Goodnight (Smoke).

References

1. Anderson, L., Li, T.M., Lehtinen, J., Durand, F.: A Domain-Specific Language for Monte Carlo Sampling. *ACM Trans. Graph.* **36**(4) (2017)
2. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M., Szeliski, R.: Middlebury Flow Accuracy and Interpolation Evaluation (2011). URL <http://vision.middlebury.edu/flow/eval/>
3. Baker, S., Scharstein, D., Lewis, J.P., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. *IJCV* **92**(1), 1–31 (2011)
4. Barron, J.L., Fleet, D.J., Beauchemin, S.S.: Performance of optical flow techniques. *IJCV* **12**(1), 43–77 (1994)
5. Bauszat, P., Eisemann, M., Eisemann, E., Magnor, M.: General and robust error estimation and reconstruction for monte carlo rendering. *Computer Graphics Forum* **34**(2), 597–608 (2015)
6. Bitterli, B., Rousselle, F., Moon, B., Iglesias-Gutián, J.A., Adler, D., Mitchell, K., Jarosz, W., Novák, J.: Nonlinearly Weighted First-order Regression for Denoising Monte Carlo Renderings. *Computer Graphics Forum* **35**(4), 107–117 (2016)
7. Buck, I., Foley, T., Horn, D., Sugerman, J., Fatahalian, K., Houston, M., Hanrahan, P.: Brook for gpus: Stream computing on graphics hardware. *ACM Trans. Graph.* **23**(3), 777–786 (2004)
8. Cook, R.L., Porter, T., Carpenter, L.: Distributed ray tracing. In: *Proc. SIGGRAPH '84*, pp. 137–145
9. 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), 8:1–8:15 (2014)
10. Erofeev, M., Gitman, Y., Vatolin, D., Fedorov, A., Wang, J.: Videomattng (2014). URL <http://videomattng.com/>
11. Erofeev, M., Gitman, Y., Vatolin, D., Fedorov, A., Wang, J.: Perceptually motivated benchmark for video matting. In: X. Xie, M.W. Jones, G.K.L. Tam (eds.) *BMVA Press Proc.*, pp. 99.1–99.12 (2015)
12. 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 TOG* **27**(212), 1 (2008)
13. Heck, D., Schlömer, T., Deussen, O.: Blue noise sampling with controlled aliasing. *ACM Trans. Graph.* **32**(3), 25:1–25:12 (2013)

14. Jakob, W.: Mitsuba renderer (2010). [Http://www.mitsuba-renderer.org](http://www.mitsuba-renderer.org)
15. Kajiya, J.T.: The rendering equation. *SIGGRAPH'86* **20**(4), 143–150 (1986)
16. Kalantari, N.K., Bako, S., Sen, P.: A machine learning approach for filtering Monte Carlo noise. *ACM Trans. Graph.* **34**(4), 122:1–122:12 (2015)
17. Kalantari, N.K., Sen, P.: Removing the noise in monte carlo rendering with general image denoising algorithms. *Computer Graphics Forum* **32**(2pt1), 93–102 (2013)
18. Lee, M.E., Redner, R.A.: Filtering: A note on the use of nonlinear filtering in computer graphics. *IEEE Comput. Graph. Appl.* **10**(3), 23–29 (1990)
19. Li, T.M., Wu, Y.t., Chuang, Y.y.: SURE-based optimization for adaptive sampling and reconstruction. *ACM Trans. Graph.* **31**, 1 (2012)
20. Mark, W.R., Glanville, R.S., Akeley, K., Kilgard, M.J.: Cg: A system for programming graphics hardware in a c-like language. *ACM Trans. Graph.* **22**(3), 896–907 (2003)
21. Moon, B., Carr, N., Yoon, S.E.: Adaptive rendering based on weighted local regression. *ACM Trans. Graph.* **33**(5), 170:1–170:14 (2014)
22. Mullapudi, R.T., Adams, A., Sharlet, D., Ragan-Kelley, J., Fatahalian, K.: Automatically scheduling halide image processing pipelines. *ACM Trans. Graph.* **35**(4), 83:1–83:11 (2016)
23. Pharr, M., Humphreys, G.: *Physically Based Rendering, from Theory to Implementation*, 2 edn. Morgan Kaufmann (2010)
24. Pharr, M., Jakob, W., Humphreys, G.: *Physically Based Rendering, from Theory to Implementation*, 3 edn. Morgan Kaufmann (2016)
25. Ragan-Kelley, J., Adams, A., Paris, S., Levoy, M., Amarasinghe, S., Durand, F.: Decoupling algorithms from schedules for easy optimization of image processing pipelines. *ACM Trans. Graph.* **31**(4), 32:1–32:12 (2012)
26. Ren, P., Wang, J., Gong, M., Lin, S., Tong, X., Guo, B.: Global illumination with radiance regression functions. *ACM Trans. Graph.* **32**, 1 (2013)
27. Rhemann, C., Rother, C., Wang, J., Gelautz, M., Kohli, P., Rott, P.: Alpha matting evaluation website (2009). URL http://www.alphamatting.com/eval_25.php
28. Rhemann, C., Rother, C., Wang, J., Gelautz, M., Kohli, P., Rott, P.: A perceptually motivated online benchmark for image matting. In: *CVPR*, pp. 1826–1833 (2009)
29. Rousselle, F., Knaus, C., Zwicker, M.: Adaptive sampling and reconstruction using greedy error minimization. *ACM Trans. Graph.* **30**(6) (2011)
30. Rousselle, F., Knaus, C., Zwicker, M.: Adaptive rendering with non-local means filtering. *ACM Trans. Graph.* **31**(6), 195:1–195:11 (2012)
31. Rousselle, F., Manzi, M., Zwicker, M.: Robust denoising using feature and color information. *Comput. Graph. Forum* **32**(7), 121–130 (2013)
32. Rushmeier, H.E., Ward, G.J.: Energy preserving non-linear filters. In: *Proc. SIGGRAPH '94*, pp. 131–138 (1994)
33. Samet, H.: Sorting in space: Multidimensional, spatial, and metric data structures for computer graphics applications. In: *ACM SIGGRAPH ASIA 2010 Courses, SA '10*, pp. 3:1–3:52 (2010)
34. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV* **47**(1-3), 7–42 (2002)
35. Scharstein, D., Szeliski, R., Hirschmüller, H.: Middlebury Stereo Vision Page (2002). URL <http://vision.middlebury.edu/stereo/>
36. Sen, P., Darabi, S.: On filtering the noise from the random parameters in Monte Carlo rendering. *ACM Trans. Graph.* **31**(3), 1–15 (2012)
37. Stein, C.M.: Estimation of the mean of a multivariate normal distribution. *Ann. Statist.* **9**(6), 1135–1151 (1981)
38. Veach, E., Guibas, L.J.: Metropolis light transport. In: *Proc. SIGGRAPH '97*, pp. 65–76 (1997)
39. Zwicker, M., Jarosz, W., Lehtinen, J., Moon, B., Ramamoorthi, R., Rousselle, F., Sen, P., Soler, C., Yoon, S.E.: Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering. *Computer Graphics Forum* **34**(2), 667–681 (2015)

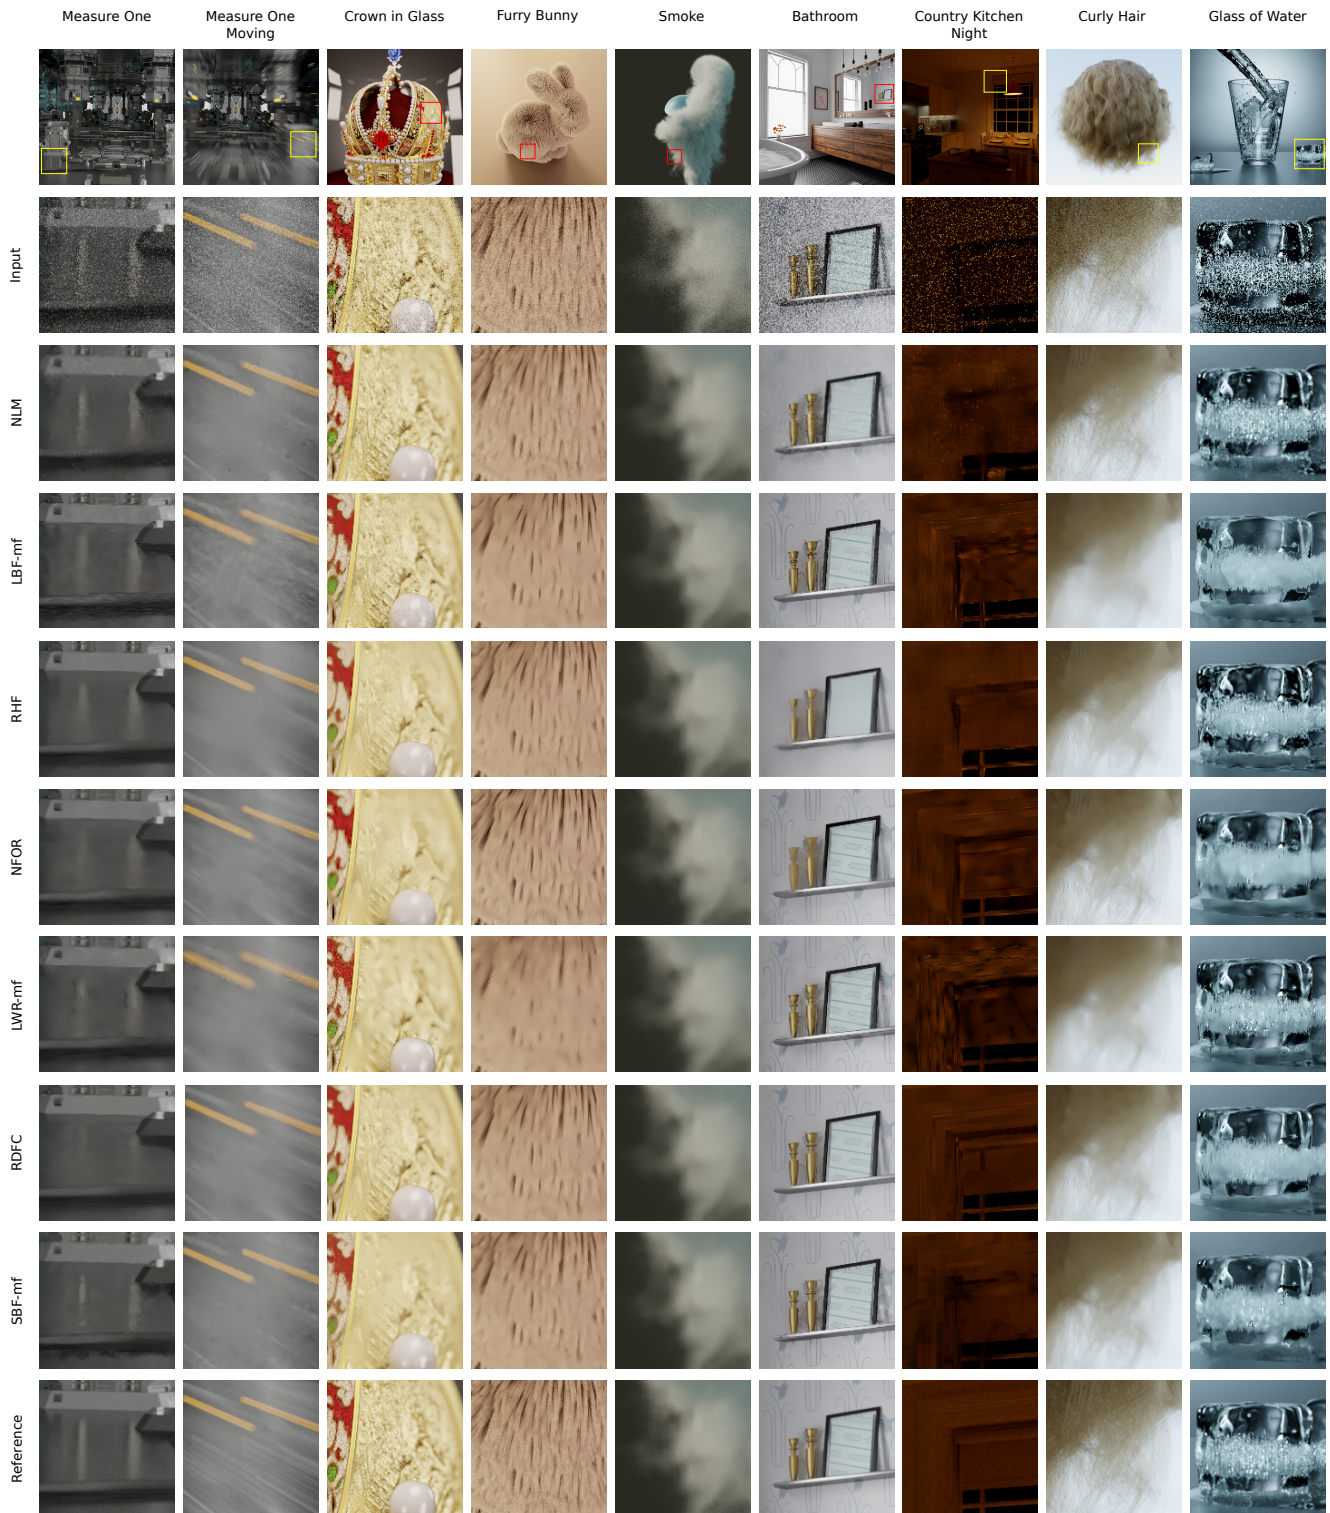


Fig. 10 Results from a benchmark including seven MC denoising techniques and nine scenes (from our scene pool) that pose challenges to denoising methods. All results were generated with 128 samples per pixel.