


Computer Graphics II: Rendering

CSE 168 [Spr 20], Lecture 13: Monte Carlo Denoising
Ravi Ramamoorthi



<http://viscomp.ucsd.edu/classes/cse168/sp20>



To Do

- Homework 4 (importance sampling) due May 18
- These lectures cover more advanced topics
 - May be relevant for your final project
 - Or curiosity in terms of frontiers of modern rendering
- This lecture on Monte Carlo denoising summary of whole CSE 274 class I taught last academic year
- Topic of great current interest in both research and production offline and real-time (OptiX, RTX GPUs)
- Good idea for final project, or simply leverage denoiser built into modern OptiX implementations
- Can get down to 1-4 samples per pixel. Amazing!**
- Lecture is high level, ask if need detailed pointers

Cook et al. [1984] results




Motivation

- Distribution effects (depth of field, motion blur, global illumination, soft shadows) are slow. Many dimensions sample



- Ray Tracing physically accurate but slow, not real-time
- Can we adaptively sample and filter for fast, real-time?

Monte Carlo Path Tracing




1000 paths/pixel

Jensen

Sampling and Reconstruction

- Monte Carlo is noisy at low sample counts
- Can we reduce time/samples by smart adaptive sampling and smart filtering/reconstruction?
- General area of Monte Carlo denoising
- Long history [Mitchell 91, Guo 98]

Sample result



Adaptive sampling + reconstruction

- These algorithms use 2 kinds of noise reduction strategies, sometimes combined:



1. Adaptive sampling algorithms

- Use information from renderer to position new samples better to reduce noise

2. Reconstruction (filtering) algorithms


- Use information from renderer to remove MC noise directly

- Both methods have been explored in the past, but new algorithms make remarkable advances



History

- Adaptive sampling old technique Mitchell et al. 87, 91,...



- But not very widely used... artifacts, can miss features
- After seminal papers in 87-91, not much follow on

Directional Coherence Maps

- Allocate samples to edges (Guo 98) Most of variance at those edges in the image




Figure 5: Comparison of the sampling patterns of adaptive stochastic sampling (top row) and the DCM (bottom row). The patterns in the left column are taken from RADANCE work image described in Section 5. Some of the fine features are shown in zoomed views of the sampling patterns in the right column. These zoomed views correspond to the same region as the zoomed views in Fig. 4.

Directional Coherence Maps (Guo 98)




Figure 2: Progressive renderings of an office scene lit by sunlight transferred through a light shelf. (a) The approximate image at the end of the regular subdivision, with 1.6% evaluated pixels located at the corners of the 8×8 blocks in the work image. (b) The approximate image after boundary evaluations for all 8×8 blocks in the work image, with 5% of pixels evaluated. (c) The approximate image after evaluating about 6% of the pixels, whose locations are shown in Fig. 5 (bottom left). (d) The final image as rendered by the baseline RADANCE system. The scene model was supplied courtesy of Greg W. Lawson.

Guo 98

A Frequency Analysis of Light Transport

F. Durand, MIT CSAIL

N. Holzschuch, C. Soler, ARTIS/GRAVIR-IMAG INRIA

E. Chan, MIT CSAIL


F. Sillion, ARTIS/GRAVIR-IMAG INRIA

Fourier analysis 101

- Spectrum corresponds to blurriness:
 - Sharpest feature has size $\sim 1/F_{\max}$
- Convolution theorem:
 - Multiplication of functions: spectrum is convolved
 - Convolution of functions: spectrum is multiplied
- Classical spectra:
 - Box becomes sinc
 - Dirac becomes constant


Transport

- Shear: $x' = x - v d$




Transport in Fourier space

- Shear in primal: $x' = x - v d$
- Shear in Fourier, along the other dimension



Transport becomes Shear

- This is consistent with light field spectra [Chai et al. 00, Isaksen et al. 00]




BRDF integration

- Ray-space: **convolution**
 - Outgoing light: convolution of incoming light and BRDF
 - For rotationally-invariant BRDFs
- Fourier domain: **multiplication**
 - Outgoing spectrum: multiplication of incoming spectrum and BRDF spectrum



Adaptive shading sampling


- Per-pixel prediction of max. frequency (bandwidth)
 - Based on curvature, BRDF, distance to occluder, etc.
 - No spectrum computed, just estimate max frequency




Per-pixel bandwidth criterion

Adaptive shading sampling


- Per-pixel prediction of max. frequency (bandwidth)
 - Based on curvature, BRDF, distance to occluder, etc.
 - No spectrum computed, just estimate max frequency



Uniform sampling




Adaptive sampling




Resurgence (2008 -)

- Eurographics 2015 STAR report by Zwicker et al. [former UCSD faculty]
- [Durand et al. 2005] Frequency analysis light transport. Proposed use for adaptive sampling. Not very practical




Multi-Dimensional Adaptive Sampling

- Hachisuka, Jarosz, ... Zwicker, Jensen [MDAS 2008]
- Scenes with motion blur, depth of field, soft shadows
- Involves high-dimensional integral, converges slowly
- Exploit high-dimensional info to sample adaptively
- Sampling in 2D image plane or other dims inadequate
 - Need to consider full joint high-dimensional space



Multidimensional Adaptive Sampling



Multidimensional Adaptive Sampling




Figure 10: Visualizations of projected sample distributions using our method for the chess scene from Figure 8 and the pool scene from Figure 7. Our adaptive sampler places samples both around high frequency image discontinuities (in focus chess piece and stationary pool ball) as well as in regions which exhibit significant motion blur or depth of field effects.


Multi-Dimensional Adaptive Sampling



Motion Blur and Depth of Field 32 samples per pixel

A-Priori Methods

- Egan et al. 2009: Frequency Analysis and Sheared Filtering for Motion Blur; first deep use frequency anal.




Fast Motion Blur Rendering



Garfield: A Tail of Two Kitties
Rhythm & Hues Studios
Twentieth Century-Fox Film Corporation

A Simple Approach

- For each pixel
 - Sample many different moments in time
 - Very expensive. Can we do better sampling, filtering?









$t = 0.00$
 $t = 0.25$
 $t = [0.50, 1.0]$
 $t = 0.75$
 $t = 1.00$

Observation 1

- Motion-blurred images have low spatial frequency




Egan, Tseng, Holzschuch, Durand, Ramamoorthi 09




Basic Example – Fourier Domain

- Low velocity, small shear in both domains




Basic Example – Fourier Domain

- Large shear




Basic Example – Fourier Domain

- Non-linear motion, wedge shaped spectra




Standard Reconstruction Filter

- Standard anti-aliasing and reconstruction filter is axis-aligned



Sheared Reconstruction Filter

- Our sheared filter allows for much tighter packing of replicas (ie sparse sampling)





Car Scene

Our Method,
4 samples per pixel



Static Render





Fourier Analysis, Sheared Filtering

Previous Work

- Shinya 93: Spatio-temporal filtering of uniform velocity
- Chai et al. 00: Plenoptic Sampling: wedge spectra
- Hachisuka et al. 08: Multidimensional Adaptive Sampling

Our Subsequent Work


- Adaptive Wavelet Rendering [Overbeck et al 09]
- Area Light Soft Shadows [Egan et al 11a]
- Spherical Harmonic Directional Occlusion [Egan et al 11b]
- Fast (real-time) Sheared Filtering [Yan et al 15]**

Real-Time Axis-Aligned Filtering

- Soft Shadows [Mehta Wang Ramamoorthi 12]
- Global Illumination [Mehta Wang Ramamoorthi Durand 13]
- Multiple Effects [Mehta Yao Ramamoorthi Durand 14]
- Multiple Axis-Aligned Filtering [Wu et al. 17]**

Fast Sheared Filtering (FSF)

- Separable sheared filter




Fourier: sheared shape

Fourier: compact pack replicas

Primal: **separable** filter, still **hard** to extend to higher dimensions (can be done by approximation)

Primal: **low** sampling rate








MAAF: Fourier

- Motivation: Cover the spectrum **compactly**

$$W(\Omega_x, \Omega_y) = \sum_{p=-(N-1)}^{N-1} \mu_p G(\Omega_x - C_x^p, \sigma_x^p) G(\Omega_y - C_y^p, \sigma_y^p)$$

- Multiple Axis-Aligned Filter (MAAF)



(b) 3 components (c) 5 components (d) 7 components (e) 25 components

Video




Sparse Sampling and Reconstruction

A Priori Methods



[Chai et al. 2000] [Durand et al. 2005] [Egan et al. 2009]

A Posteriori Methods



[Hachisuka et al. 2008] [Li et al. 2012]



[Rousselle et al. 2012] [Moon et al. 2016]

Adaptive Wavelet Rendering

The figure displays a comparison of rendering results for three different scenes, comparing four methods: 'Our Method', 'Multires (PSTT)', 'Reference Monte Carlo', and 'MCAS'. Each method's results are shown in a column, with the first column also including a reference image. The rows represent different scenes: a sphere, a fractal, and a landscape. Each image is accompanied by a caption indicating the number of samples per pixel and the error metric (RMSE or L1 error).

Scene	Our Method	Multires (PSTT)	Our Method	Reference Monte Carlo	MCAS
Sphere	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 40 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE
Fractal	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 40 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE
Landscape	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 40 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE	100 Accelerated DOF + Motion Blur average 32 samples per pixel 0.001 RMSE


Overbeck et al 09
General high-D
effects. Simple
and fast (renders
into wavelet domain)

Feature-Space Methods

- General practical denoising (no frequency) [2012-]
- General effects (Sec 2.3 of EG STAR Report)
- General image-space denoising framework
- *But use auxiliary features (depth, normals, etc.)*
- Basis for methods deployed in industry today

Random Parameter Filtering

- Sen Darabi 12, importance of each feature
 - Addresses noisy features (e.g. depth of field)
 - Notion of mutual information
- Weighted bilateral filter, very good at low samples
 - Parameters determined by feature importance
 - Auxiliary features are key to beat image denoising
 - Has led to newer methods, commercialization



(a) Screen position (b) Random parameters (c) World space coords. (d) Surface normals (e) Texture value (f) Sample color



Subsequent Work

- SURE (Stein's unbiased risk estimator: general kernels, adaptive sampling, general effects)




Figure 1: Comparisons between greedy error minimization (GEM) [Rousset et al. 2011] and our SURE-based filtering. With SURE, we are able to use kernels (cross bilateral filters in this case) that are more effective than GEM's isotropic Gaussians. Thus, our approach better adapts to anisotropic features (such as the motion blur pattern due to the motion of the airplane) and preserves scene details (such as the textures on the floor and curtains). The kernels of both methods are visualized for comparison.


Subsequent Work

- Moon et al. local linear or polynomial models, treat as regression. Many other methods
- APR: Polynomial order chosen to minimize error
- Newest methods use deep learning instead




Impact: Offline

- Handle general effects. Sample and denoise (builds on AWR, AAF, FSF, MAAF. Predict general filter kernel)
- Many more sophisticated methods available now; used in almost every major production rendering software
- Based on Deep Learning for Monte Carlo Denoising



Impact: Real-Time



Impact: Real-Time

- Extend AAF, FSF, MAAF: Predict Filter based on Deep Learning (sample and AI-based denoising)
- NVIDIA software (OptiX 2017), hardware (RTX 2018)
- 40-year journey: ray tracing curiosity to every pixel



From SIGGRAPH 18



Real Photo: Speaker and Turner Whitted at SIGGRAPH 18