

Sampling and Reconstruction of Visual Appearance

CSE 274 [Winter 2018], Lecture 1

Ravi Ramamoorthi

<http://www.cs.ucsd.edu/~ravir>



Motivation: High-Dim. Appearance

Appearance in vision/graphics involves high-dimensional data

- **(Precomputed) Real-Time Rendering:** Light transport matrix stores variation across surface, light, view (4D-6D)
- **Monte Carlo Rendering:** Need to sample across time, light source, depth of field for each pixel (3D-7D)
- **Appearance Acquisition:** Acquire reflectance functions and light transport (4D-8D)
- **Computer Vision:** Effect of lighting on images (4D)

Consider real-time rendering (6D): With ~ 100 samples/dimension
 $\sim 10^{12}$ samples total!! : Intractable computation, rendering

Fast Motion Blur Rendering (3D)



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

Precomputed Real-Time Rendering (6D)




Ng, Ramamoorthi, Hanrahan 03.04

Real-Time Global Illumination (4D)

Interactive Physically-Based Indirect Illumination Using Axis-Aligned Filtering

[Mehta, Wang, Ramamoorthi, Durand13]

Data-Driven Appearance Models (6D)



Stanford/Cornell Graphics Labs

Lawrence et al. 06


Volumetric Materials (Gigavoxels)



[Hasan and Ramamoorthi 13]

Lighting-Insensitive Recognition

Illuminate subject from many incident directions
Understand space of images as lighting is varied
Low Dimensional Subspace [Ramamoorthi 02]



Debevec et al. 00

Sparse Sampling, Reconstruction

- Same algorithm as offline Monte Carlo rendering
- But with smart sampling and filtering (current work)



Sparse Sampling, Reconstruction

Rendered Offline, 40 Samples Per Pixel (SPP), Moving Geometry, 2-Bounce, Indirect



Overview

Many problems in graphics, vision are hard to solve directly

- Visual appearance** is high-dimensional
- High sampling rate, computational cost
- Often prohibitive to even sample before compression

Mathematical and signal-processing approach

- Treat illumination and reflection functions as signals
- Exploit "sparsity" for acquisition of light transport and appearance
- Reduces to efficient **sampling and high quality reconstruction**

Large range of applications in graphics, vision

- Monte Carlo Rendering* (order-of-mag. advances classical problem)
- Appearance Acquisition and Editing
- Imaging and Computational Photography
- Precomputed Rendering, Physics-Based Vision, Animation, ...

Much recent work (esp at UCSD): CSE 274 Topics in Graphics:
Sampling and Reconstruction of Visual Appearance

Outline of Lecture


- Motivation, sampling and reconstruction visual appearance
- Historical Development and Overview of Applications*
- Logistics of course

(My lens on) History

- Data-Driven Visual Appearance [~1993 -]
 - Image-Based Rendering
 - Sampled Representations (in many areas)
 - Even for synthetic relighting [Nimeroff 94, Sloan 02]


Types of Measured Visual Appearance

- Lighting: From point lights to environment maps and beyond




Grace Cathedral and Kitchen light probes
Courtesy Paul Debevec www.debevec.org

- BRDFs: From Lambertian/Cook Torrance to measured/factored



Types of Measured Visual Appearance

- "Reflectance Fields": Variation with lighting and/or view
- Subsurface and Volumetric Scattering
- Time-Varying Surface Appearance
- BTFs or Bi-Directional Texture Functions
- And many more (full taxonomy next)



Time-Varying Appearance: [Video](#)



General Plenoptic Function


- All knowledge of light in scene [Adelson 91]
- Anywhere in space (x, y, z)
- In any direction (θ, ϕ)
- At any time instant (t)
- For any wavelength of light (λ)
- Function of 7 variables, therefore 7D function
- We care about taxonomy of *scattering functions*
 - General Scattering Function is 14D (bet. two plenoptics)
$$f(x_i, y_i, z_i, \theta_i, \phi_i, \lambda_i, t_i; x_o, y_o, z_o, \theta_o, \phi_o, \lambda_o, t_o)$$

Common Assumptions

- Ignore time dependence (no phosphorescence or time-varying BRDF properties)
- Ignore wavelength (no fluorescence, assume RGB)
- Travel in free space, parameterize on surfaces (no z)
 - Alternative for light fields: 4D space of rays (intersections in 2 planes)
- Each of these removes 1D of plenoptic, 2D of scattering
- Left with 8D function of greatest importance for class
- 8D Bi-Directional Surface Scattering Distribution Function (BSSRDF)

$$f(x_i, y_i, \theta_i, \phi_i; x_o, y_o, \theta_o, \phi_o)$$


Taxonomy of Appearance



(My lens on) History

- Data-Driven Visual Appearance [~1993 -]
 - Image-Based Rendering
 - Sampled Representations (in many areas)
 - Even for synthetic relighting [Nimeroff 94, Sloan 02]
- Mathematical and Computational Fundamentals [~2000 -] (my term, topic of my Career Award)
 - Lower-Dimensional Factorizations [e.g. 6D -> 4D]
 - Efficient Mathematical Representations [spherical harmonics, wavelets, radial basis functions]
 - New Computational Methods [analytic formulae, convolution, triple products, Nystrom, ...]


Factorization Approach



Triple Product Integral Relighting



Relit Images (3-5 sec/frame)



(My lens on) History and Future

- Data-Driven Visual Appearance [~1993 -]
 - Sampled representations and real-world data
 - How do we make the data practical and usable?*
- Mathematical and Computational Fundamentals [~2000 -] (my term, topic of my Career Award)
 - Lower-Dimensional Factorizations [e.g. 6D -> 4D]
 - Data still is acquired or precomputed by brute force*
- Sampling Theory of Appearance [~2007 -]
 - How many samples do we need for final representation?
 - Theory and efficient acquisition/computation strategies
 - Really a core signal-processing problem (Chai et al 00)
 - Newest approaches based on machine learning
 - Sampling and Reconstruction of Visual Appearance*

Applications and Context

- Monte Carlo global illumination rendering (CSE 168)
 - Main application area discussed in course
 - Often perceived as mature: "rendering is a solved problem"
 - But not widely used in production rendering until 2011
- Production Rendering is now Physically-Based
 - Sea change since 2011. Ad-hoc methods gone. MC used.
 - Sampling and Reconstruction (denoising) is key*
 - Leads to 1-2 order of magnitude speedup in mature area*
 - Essential part of today's production renderers (since ~2014)*



Fast Sheared Filtering



Multiple Effects in Real Time

Multiple Axis-Aligned Filters for Rendering of Combined Distribution Effects

Lifan Wu¹ Ling-Qi Yan² Alexandr Kuznetsov¹ Ravi Ramamoorthi¹
¹University of California, San Diego ²University of California, Berkeley

NO AUDIO

Applications and Context

- Fast Enough for Real-Time Rendering
 - Sample and reconstruct for real-time started ~2012
 - Latest papers machine learning + 1 sample per pixel
 - [Video](#)
- Fast precomputation of light transport matrices
 - Precomputed rendering widely used in games etc.
- Sparse acquisition of Light Transport
- Light Field Imaging, Interpolation




Resolution trade-off

Limited resolution

High angular

Low spatial



UC San Diego Kalantari et al.

Solution: angular super-resolution



UC San Diego Kalantari et al.

Our result



UC San Diego Kalantari et al.

Relation to Signal Processing, Learning

- Signal processing itself undergoing revolution
 - From Nyquist to compressive to sparse low rank
 - Will cover these topics briefly where relevant
 - Exciting time to work in this area
 - *Unified Sampling Theory of Appearance?*
- Newest advances in machine learning
 - Deep Convolutional Neural Networks (CNNs)
 - Introduced for computer vision but many exciting applications for image synthesis
 - Latest denoising methods leverage CNNs
- This is a graphics course, but we will touch on above methods as needed. Exciting convergence.

Outline of Lecture

- Motivation, sampling and reconstruction visual appearance
- Historical Development and Overview of Applications
- *Logistics of course*

Course Goals, Format

- Goal: Background and current graphics research
 - Topic: Sampling and Reconstruction of Visual Appearance
 - Need to cover a lot of background research papers
 - Then discuss current frontiers in the field
- UCSD is the best place for this!!
- Format: Alternate lectures, student presentations of papers
- Website: <http://www.cs.ucsd.edu/~ravir/274/18/274.html>

Course Logistics

- No textbooks. Required readings are papers available online (and some handouts for books)
 - Handouts at <http://www.cs.ucsd.edu/~ravir/274/18/readings>
- Office hours: after class or email. My contact info is on my webpage: <http://www.cs.ucsd.edu/~ravir>
- Should count for PhD, MS, BS electives in graphics and vision, see me if there is a problem or you need a certification

Requirements

- Pass-Fail (2 units)
 - Show up to class regularly
 - Present 1 or 2 paper(s) if needed
 - Prefer you do this rather than just sit in
- Grades (4 units)
 - Attend class, participate in discussions
 - Present 1 or 2 papers (please do this well)
 - Project (key part of grade)

Project

- Wide flexibility if related to course. Can be done groups of 2
 - Default: Implement (part of) one of papers and produce an impressive demo for real-time or offline rendering
 - See/e-mail me re ideas
 - Best projects will go beyond simple implementation (try something new, some extensions)
- Alternative (less desirable): Summary of 3+ papers in an area
 - Best projects will explore links/framework not discussed by authors, and suggest future research directions

Prerequisites

- Strong interest in graphics, rendering
- Computer graphics experience (167 or equivalent)
 - What if lacking prerequisites? Next slide
 - Experience with rendering (CSE 168) not required
- Course will move quickly
 - Covering recent and current active research
 - Some material quite technical
 - Considerable background material is covered
 - Assume some basic knowledge
 - Many topics. Needn't fully follow each one, but doing so will be most rewarding.

If in doubt/Lack prerequisites

- Material is deep, not broad
 - May be able to pick up background quickly
 - Course requirements need you to really fully understand only one/two areas (topics)
 - But if completely lost, won't be much fun
- If in doubt, see if you can more or less follow some of papers after background reading
- Ultimately, your call

Assignment this week

- E-mail me (ravir@cs.ucsd.edu)
 - Name, e-mail, status (Senior, PhD etc.)
 - Will you be taking course grades or P/F
 - Background in graphics/any special comments
 - Optional: Papers you'd like to present FCFS (only those that say "presented by students")
- Paper presenters for [first set of papers]

Questions?
