

Sampling and Reconstruction of Visual Appearance: From Denoising to View Synthesis

CSE 274 [Fall 2021], 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)
- **View Synthesis:** Predict new views or light field (4D)
- **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

Instructor

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

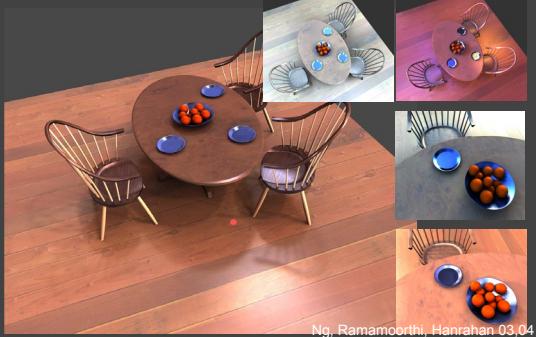
- PhD Stanford, 2002 [with Pat Hanrahan, 2020 Turing Award]
“*Spherical Harmonic Lighting*” widely used in games (e.g. Halo series), movies (e.g. Avatar), etc. (Adobe, ...)
- At Columbia 2002-2008, UC Berkeley 2009-2014
- “*Monte Carlo denoising*” inspired raytracing offline, real-time
- At UCSD since Jul 2014: Director, Center for Visual Computing
- Awards for research: White House PECASE (2008), SIGGRAPH Significant New Researcher (2007), ACM Fellow
- <https://www.youtube.com/watch?v=qpyCXqXGe7I>
- *Key contributions in view synthesis (NeRFs ECCV 20)*
- *Key contributions in Monte Carlo denoising, consult NVIDIA*

Fast Motion Blur Rendering (3D)



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

Precomputed Real-Time Rendering (6D)

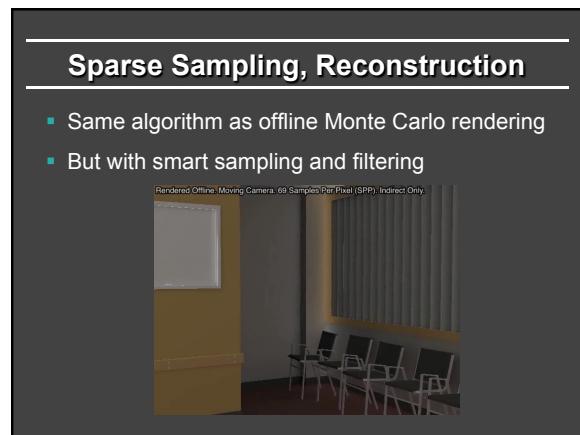
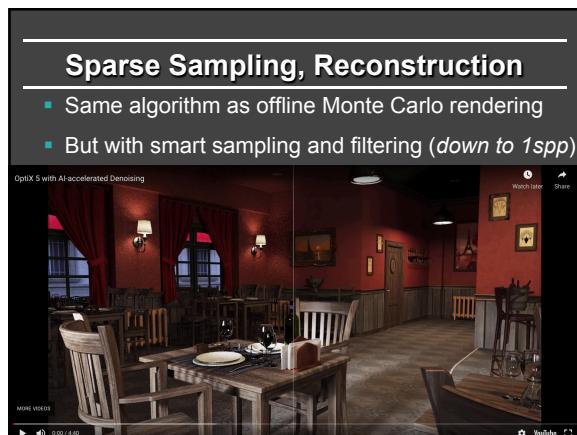
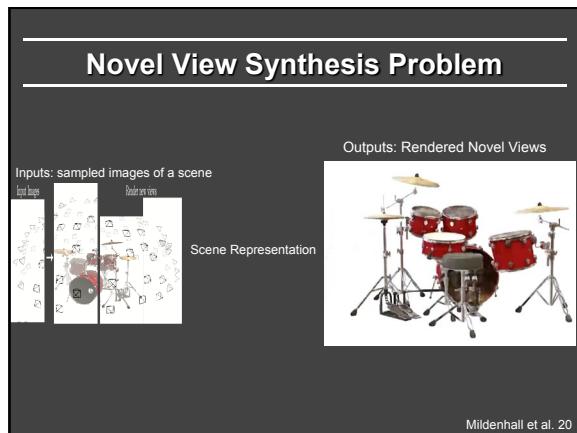
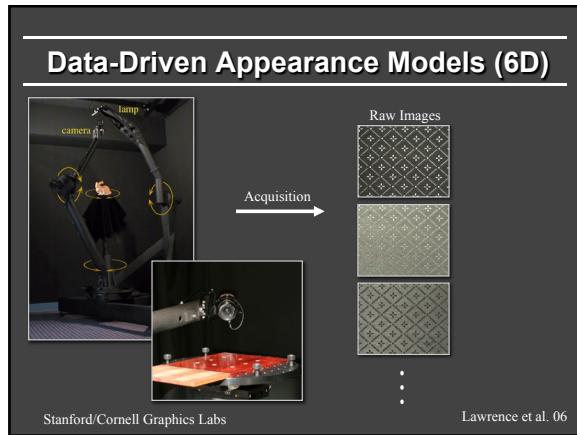


Ng, Ramamoorthi, Hanrahan 03,04

Real-Time Global Illumination (4D)

Interactive Physically-Based Indirect Illumination Using Axis-Aligned Filtering

[Mehta, Wang, Ramamoorthi, Durand13]





Sparse Sampling, Reconstruction

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

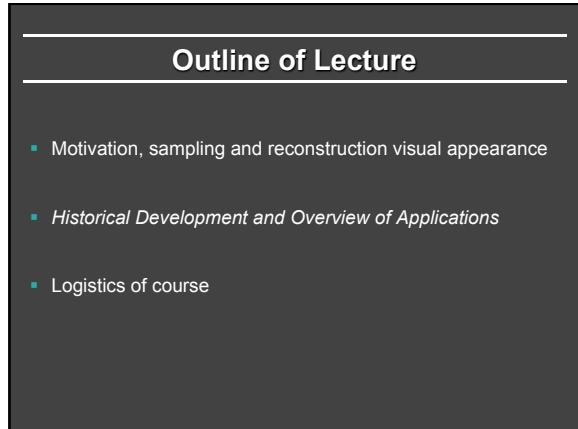
Mathematical and signal-processing approach

- 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

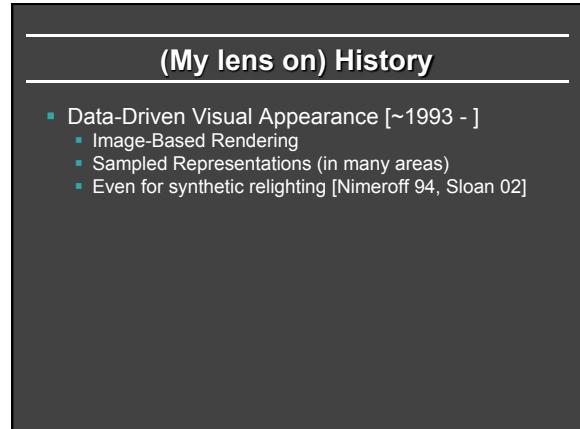
- *Monte Carlo Rendering* (order-of-mag. advances classical problem)
- *View synthesis* (best results in decades)
- Imaging and Computational Photography
- Precomputed Rendering, Physics-Based Vision, Animation, ...

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



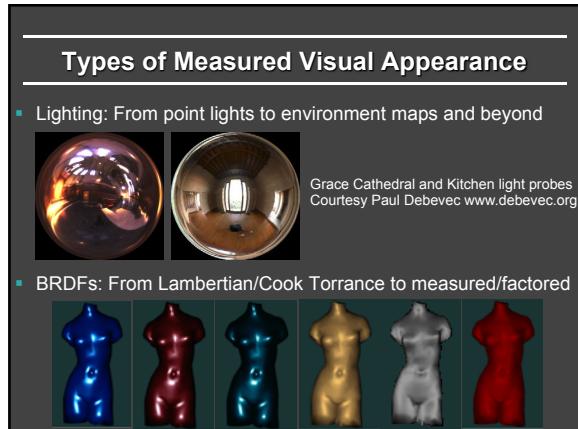
Outline of Lecture

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

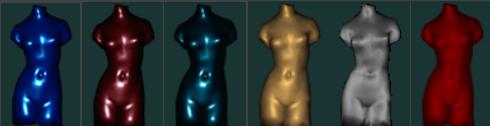


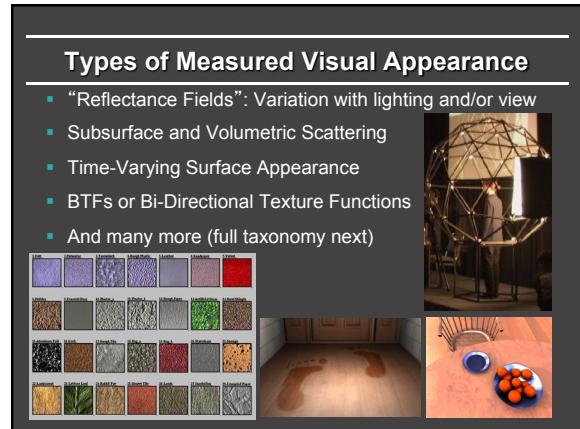
(My lens on) History

- 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
 
- 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

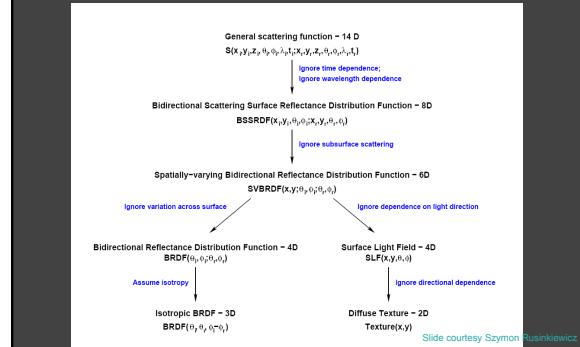
- All knowledge of light in scene [Adelson 91] (x, y, z)
- Anywhere in space (θ, ϕ)
- In any direction (t)
- At any time instant (λ)
- For any wavelength of light (λ)
- Function of 7 variables, therefore 7D function
- We care about taxonomy of *scattering functions*
 - General Scattering Function is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

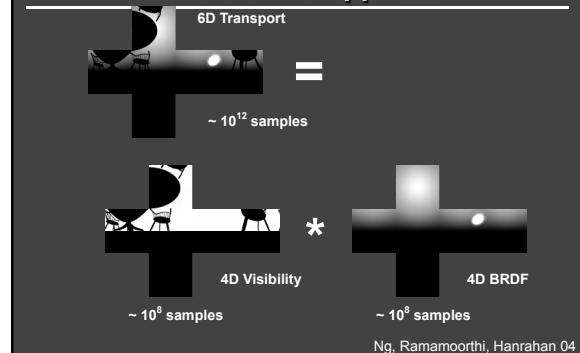
Taxonomy of Appearance

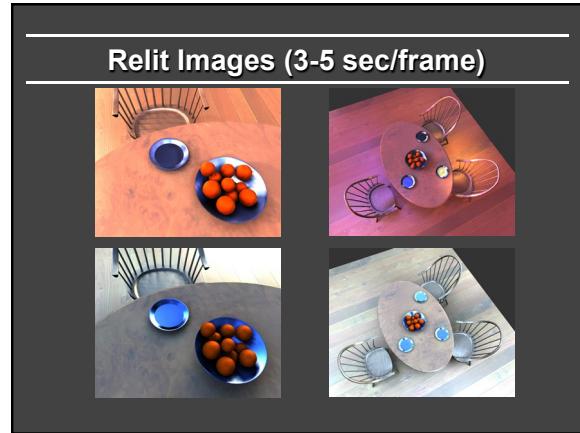
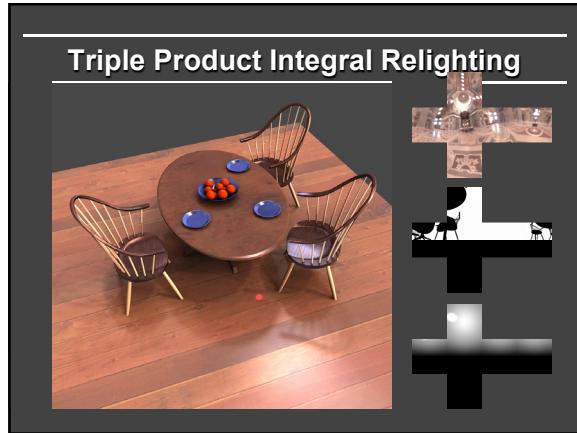


(My lens on) History

- 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 \rightarrow 4D]
 - Efficient Mathematical Representations [spherical harmonics, wavelets, radial basis functions]
 - New Computational Methods [analytic formulae, convolution, triple products, Nystrom, ...]
 - New machine learning representations (recent)

Factorization Approach



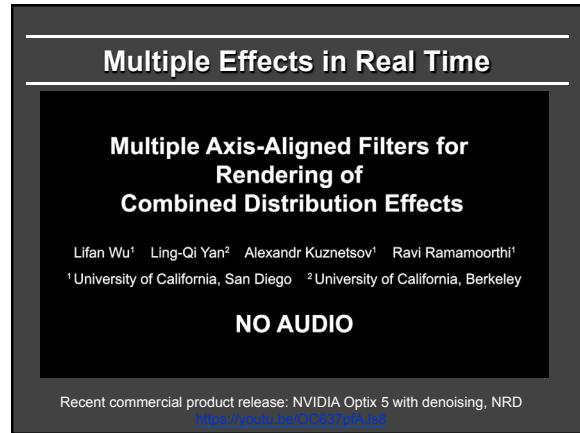
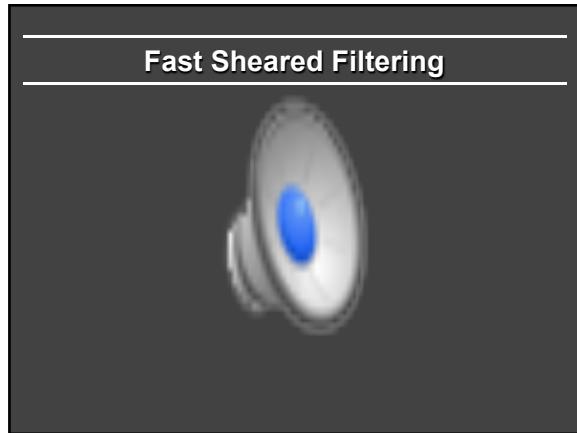


(My lens on) History and Future

- 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 \rightarrow 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 (old CSE168)
 - One of main application areas 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)*



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](#) ; TuringRT 10G rays/second: [Video](#)
- Fast precomputation of light transport matrices
 - Precomputed rendering widely used in games etc.
- Sparse acquisition of Light Transport, View Synthesis
- Light Field Imaging, Interpolation

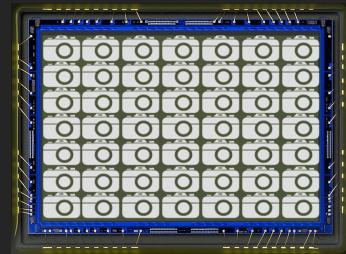


Resolution trade-off

Limited resolution

High angular

Low spatial



UCSan Diego Kalantari et al.



Solution: angular super-resolution



UCSan Diego Kalantari et al.

Our result



UCSan Diego Kalantari et al.

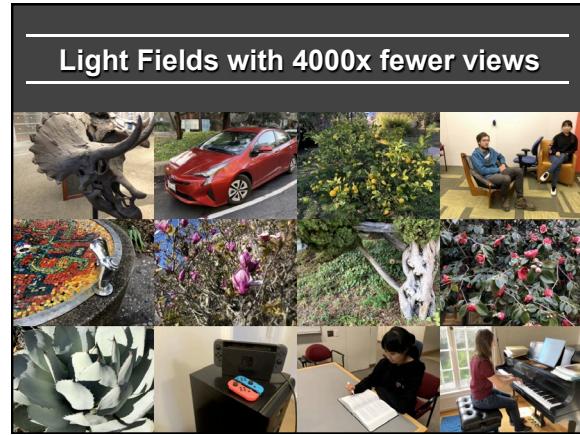


Virtual Experiences of Real-World Scenes



Input Images





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, view synthesis methods leverage CNNs
 - *Convergence real-time, offline rendering, machine learning*
- 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: *From Denoising to View Synthesis*
 - Need to cover a lot of background research papers
 - Then discuss current frontiers in the field
- Also basics of rendering to cover for CSE 168
- UCSD is the best place for this!!
- Format: Alternate lectures, student presentations of papers
- Website: <http://viscomp.ucsd.edu/classes/cse274/fa21/274.html>

Course Logistics

- No textbooks. Required readings are papers available online (and some handouts for books)
 - Handouts at <http://viscomp.ucsd.edu/classes/cse274/fa21/readings>
- Office hours: after class or email. My contact info is on my webpage: <http://www.cs.ucsd.edu/~ravir>
- Zoom: <https://ucsd.zoom.us/my/ravir11> (if we need to move to remote instruction temporarily or permanently)
- TA: Alex Kuznetsov (atkuznet@engr.ucsd.edu) Office hours in CSE 4150 MTh 12-1pm or email for another time.
- 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)
 - If class is large, groups of 2 can present 1 paper
 - 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, or for image-based rendering or view synthesis
 - 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 not required; first few weeks will cover basic background (CSE 168). But also doesn't hurt, consider UCSD online CSE 168
- 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
 - Intended as advanced graphics course for those with CSE 167/168/equivalent and very interested in area
- 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) [only if not done already]
 - 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")*

Questions?