Introduction

15-468, 15-668, 15-868
Physics-based Rendering
Spring 2022, Lecture 1

http://graphics.cs.cmu.edu/courses/15-468

Online lecture etiquette

- Lectures **are recorded**, including all discussions. This is to facilitate students that cannot attend the lectures live.

- Recordings become available on **Canvas** a few hours (usually <= 3) after the lecture. Please note that you are **not** allowed to share these recordings with anyone outside this class. This is to protect your FERPA rights and those of your fellow students.

- Please keep your Zoom window muted when you are not speaking.

- You are welcome to keep your camera on or off.

- Feel free to ask questions! Either use the “raise hand” option (preferable), or post in the chat. If I miss you, please repeat. And if I keep missing you, please unmute yourselves and mention that you have a question.

- I’ll be staying around for another 30 minutes after the lecture for additional Q&A.

Overview of today's lecture

- Teaching staff introductions
- What is this course about?
- Course fast-forward and logistics

Teaching staff introductions

Instructor: Ioannis (Yannis) Gkioulekas

I won’t hold it against you if you mispronounce my last name

Originally from Greece

National Technical University of Athens (2004-09)

Harvard University (2009-17)

Carnegie Mellon University (2017-now)

Yannis at Harvard in 2011

My website: http://www.cs.cmu.edu/~igkioule

See also: http://imaging.cs.cmu.edu/

Broadly interested in computational imaging, physics-based vision, and physics-based rendering

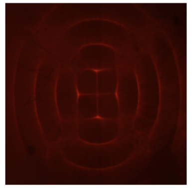

Rendering Near-Field Speckle Statistics in Scattering Media
Adithya Pediredla, Yasin Karimi Chalmiani, Matteo Giuseppe Scopelliti, Maysam Chamanzar, Srinivasa Narasimhan, **Ioannis Gkioulekas**
ACM Transactions on Graphics (SIGGRAPH Asia), 2020
paper | project page

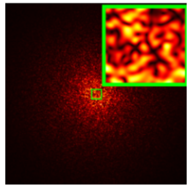

Rendering Near-Field Speckle Statistics in Scattering Media
Chen Bar, **Ioannis Gkioulekas**, Anat Levin
ACM Transactions on Graphics (SIGGRAPH Asia), 2020
paper | project page

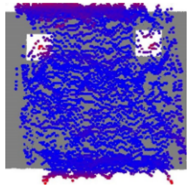

A Theory of Fermat Paths for 3D Imaging Sonar Reconstruction
Eric Westman, **Ioannis Gkioulekas**, Michael Kaess
IEEE International Conference on Intelligent Robots and Systems (IROS), 2020
paper | project page

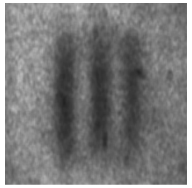

Interferometric Transmission Probing with Coded Mutual Intensity
Alankar Kotwal, Anat Levin, **Ioannis Gkioulekas**
ACM Transactions on Graphics (SIGGRAPH), 2020
paper | project page

Langevin Monte Carlo Rendering with Gradient-based Adaptation
Fujun Luan, Shuang Zhao, Kavita Bala, **Ioannis Gkioulekas**
ACM Transactions on Graphics (SIGGRAPH), 2020
paper | project page

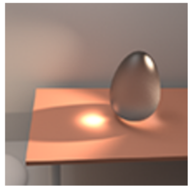

Path-Space Differentiable Rendering
Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, Shuang Zhao
ACM Transactions on Graphics (SIGGRAPH), 2020
paper | project page

Effect of Geometric Sharpness on Translucent Material Perception
Bei Xiao, Shuang Zhao, **Ioannis Gkioulekas**, Wenyan Bi, Kavita Bala Journal of Vision (JOV), 2020
paper | code

Use rendering to study human perception

A Volumetric Albedo Framework for 3D Imaging Sonar Reconstruction
Eric Westman, **Ioannis Gkioulekas**, Michael Kaess
IEEE International Conference on Robotics and Automation (ICRA), 2020
paper | project page

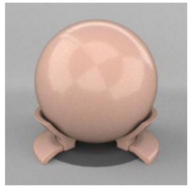

Towards Reflectometry from Interreflections
Kfir Shem-Tov, Sai Praveen Bangaru, Anat Levin, **Ioannis Gkioulekas**
IEEE International Conference on Computationa Photography (ICCP), 2020
paper | project page

Use rendering to make reflectometry easier

Towards Learning-based Inverse Subsurface Scattering
Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, **Ioannis Gkioulekas**
IEEE International Conference on Computationa Photography (ICCP), 2020
paper | project page

Use rendering to make neural networks better

Zhang et al. '20

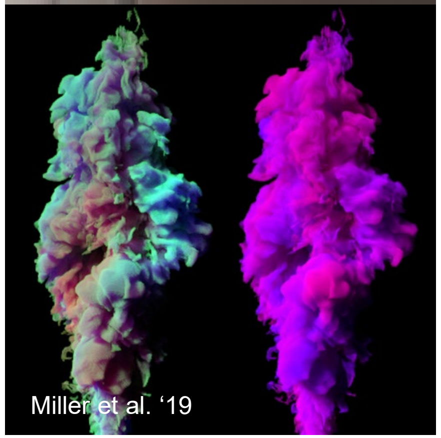

Miller et al. '19

some prior works

Bailey Miller advised by Prof. Ioannis Gkioulekas

Teaching Assistant for 15-468 and 2nd year PhD in CSD

Monte Carlo methods for PDEs

(aka physical simulations that work with complex geometry)

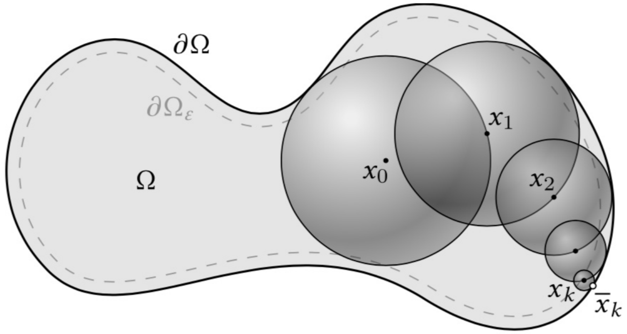

Figure from Sawhney et al '20

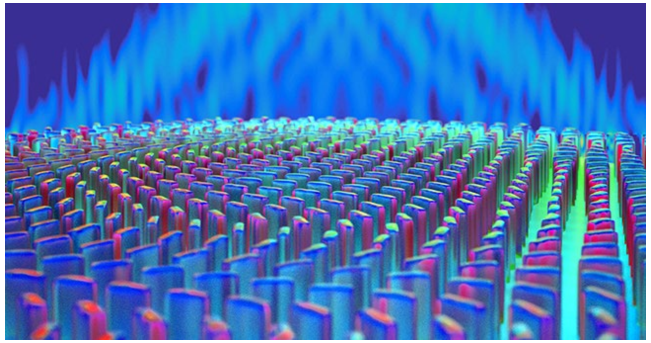

Image from startup Metalenz

neural rendering for reconstruction

(aka computer vision methods that use physics-based rendering)

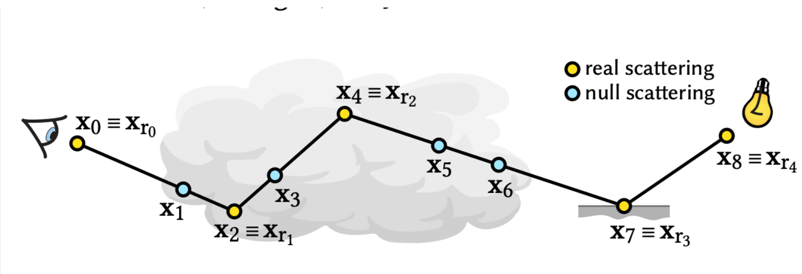

Figure from Miller et al '19

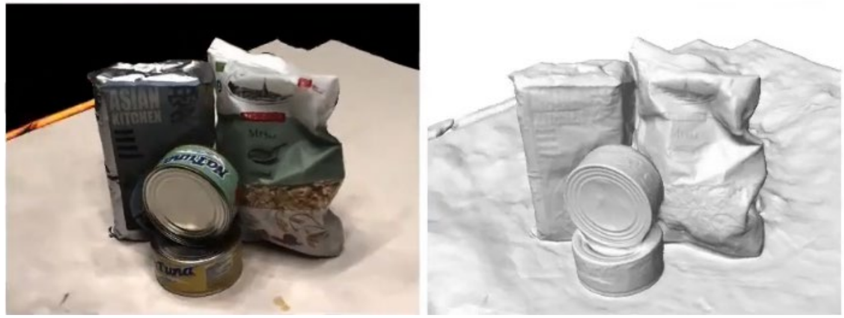

Image from Oechsle et al '21

Ryan Po
15-468 TA and Senior in CS

Research: imaging with single-photon sensors

SPAD setup

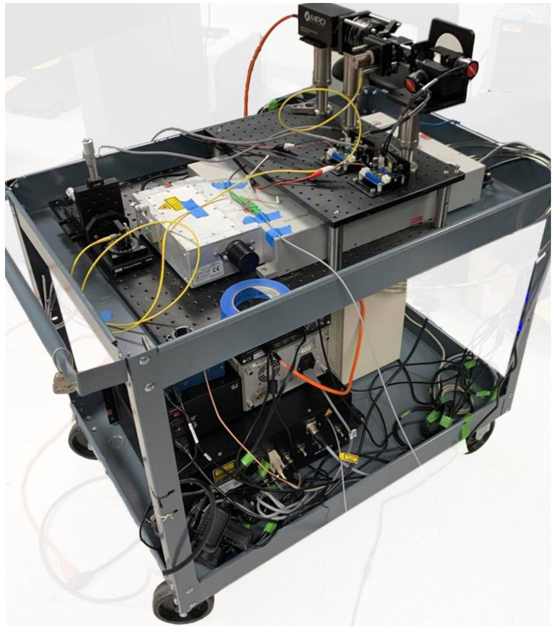

3D scan outside CFA

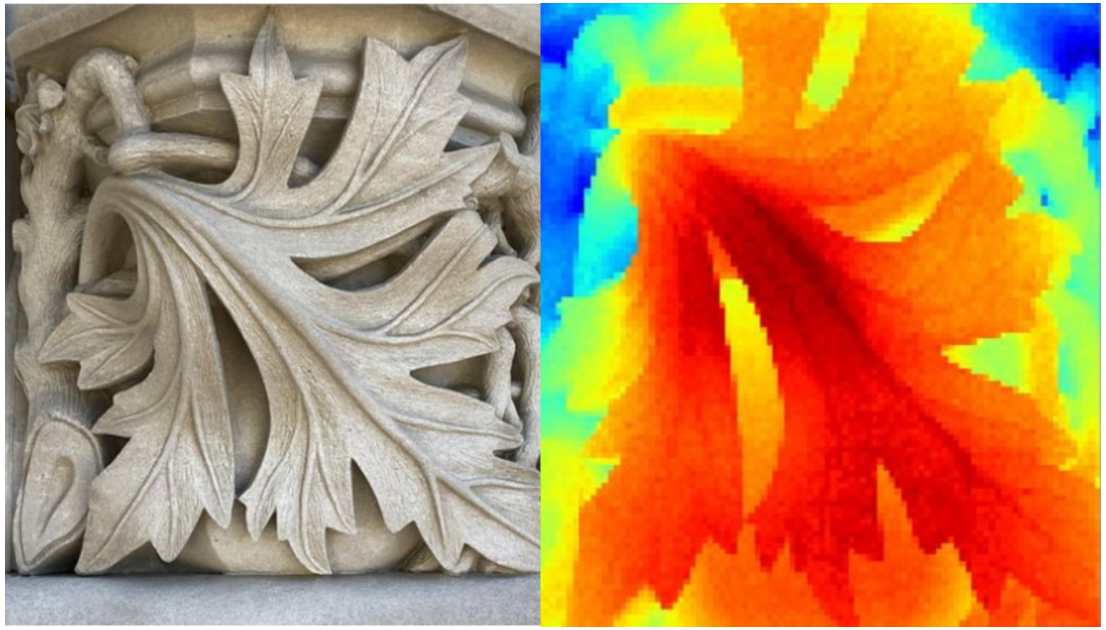

3D scan of Yannis' office

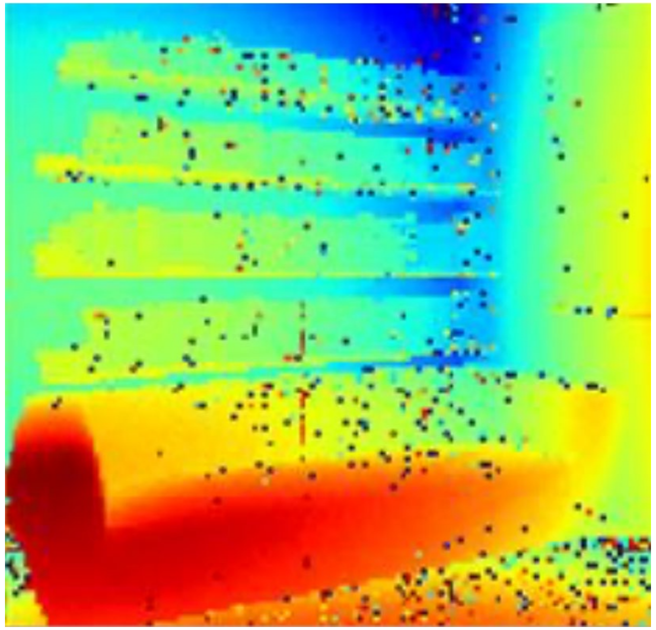

What is this course about?

Forward rendering

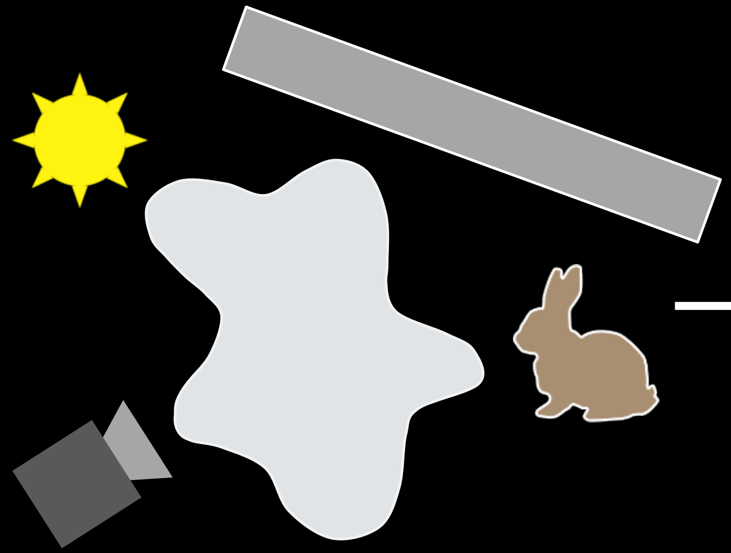

digital scene specification
(geometry, materials,
optics, light sources)

physically-accurate
rendering

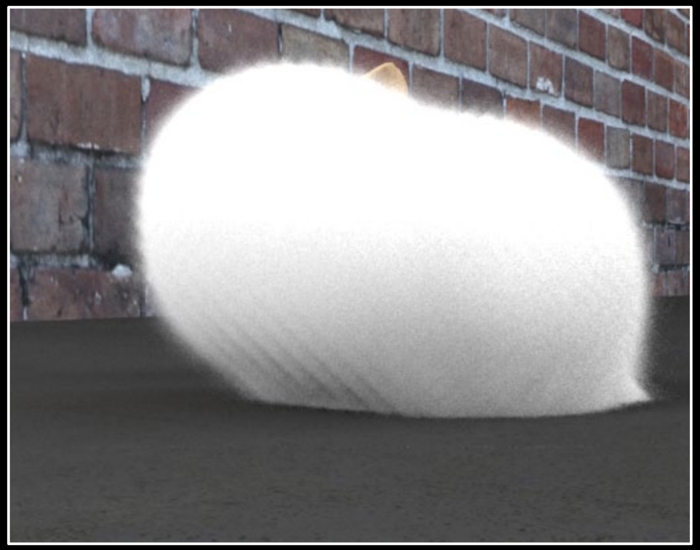

photorealistic
simulated image

What is this class about?

- How can we generate realistic images?
- Why do things look the way they do?

Motivation

Wojciech Jarosz

Light transport in the real world

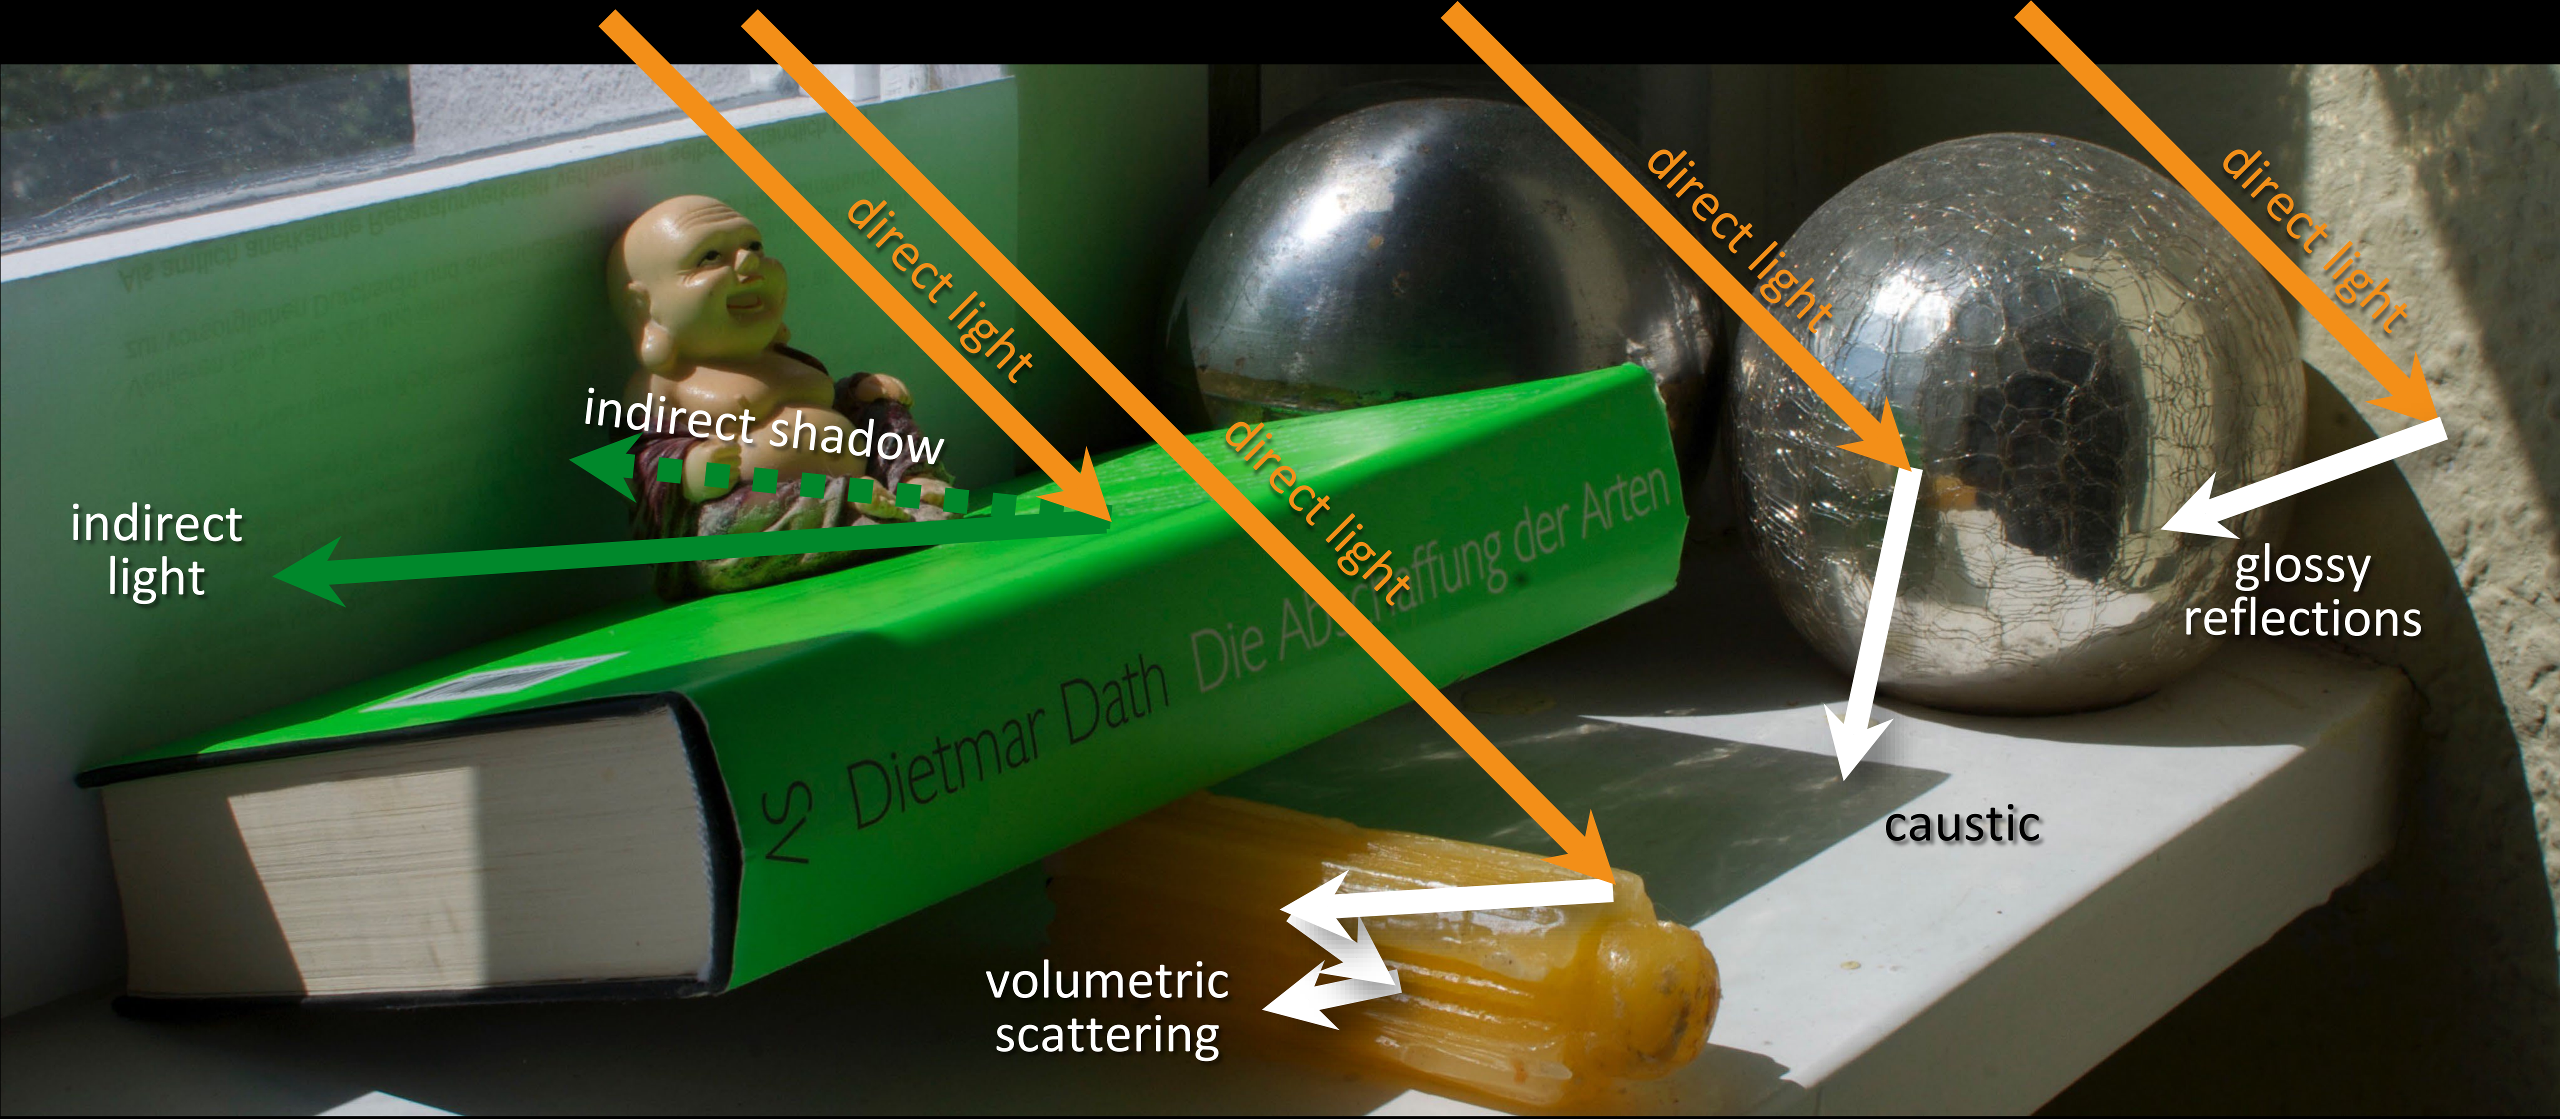

After [Ritschel et al 2011]

Physics-based rendering

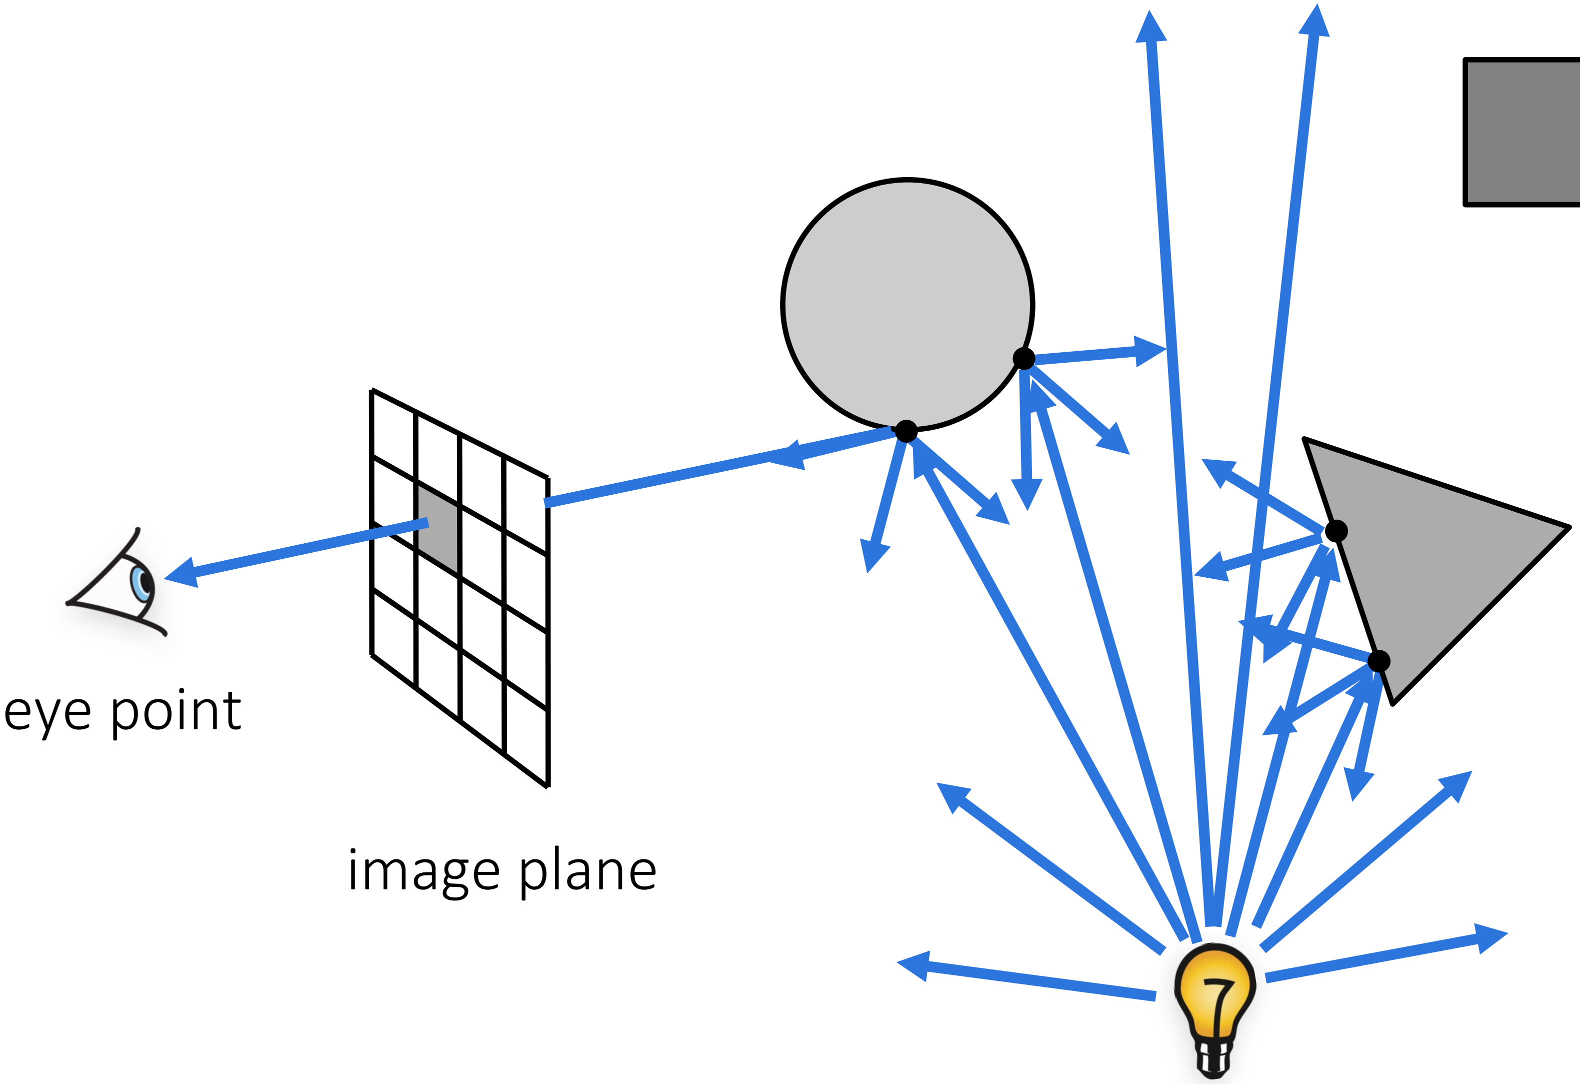

Mimic the physics of light transport using ray tracing

Ray tracing in production

Arnold Renderer

SOLIDANGLE

PIXAR's RenderMan

Hyperion

Wojciech Jarosz

Visual effects

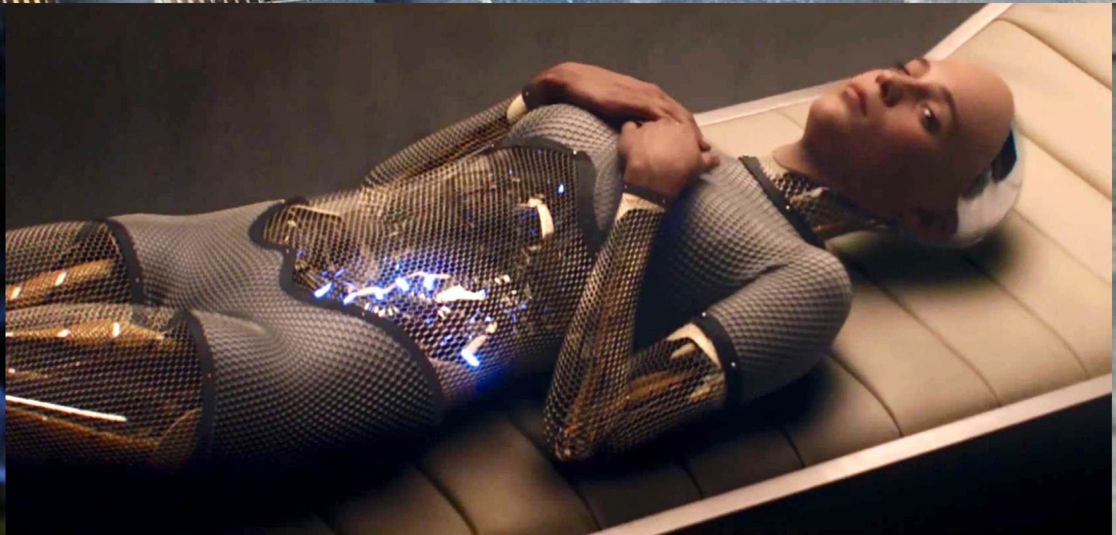

Wojciech Jarosz

Animated films

Wojciech Jarosz

Video games




Wojciech Jarosz

Architectural visualization

Wojciech Jarosz

Architectural visualization

Wojciech Jarosz

Advertising/product visualization

Wojciech Jarosz

Advertising & E-commerce

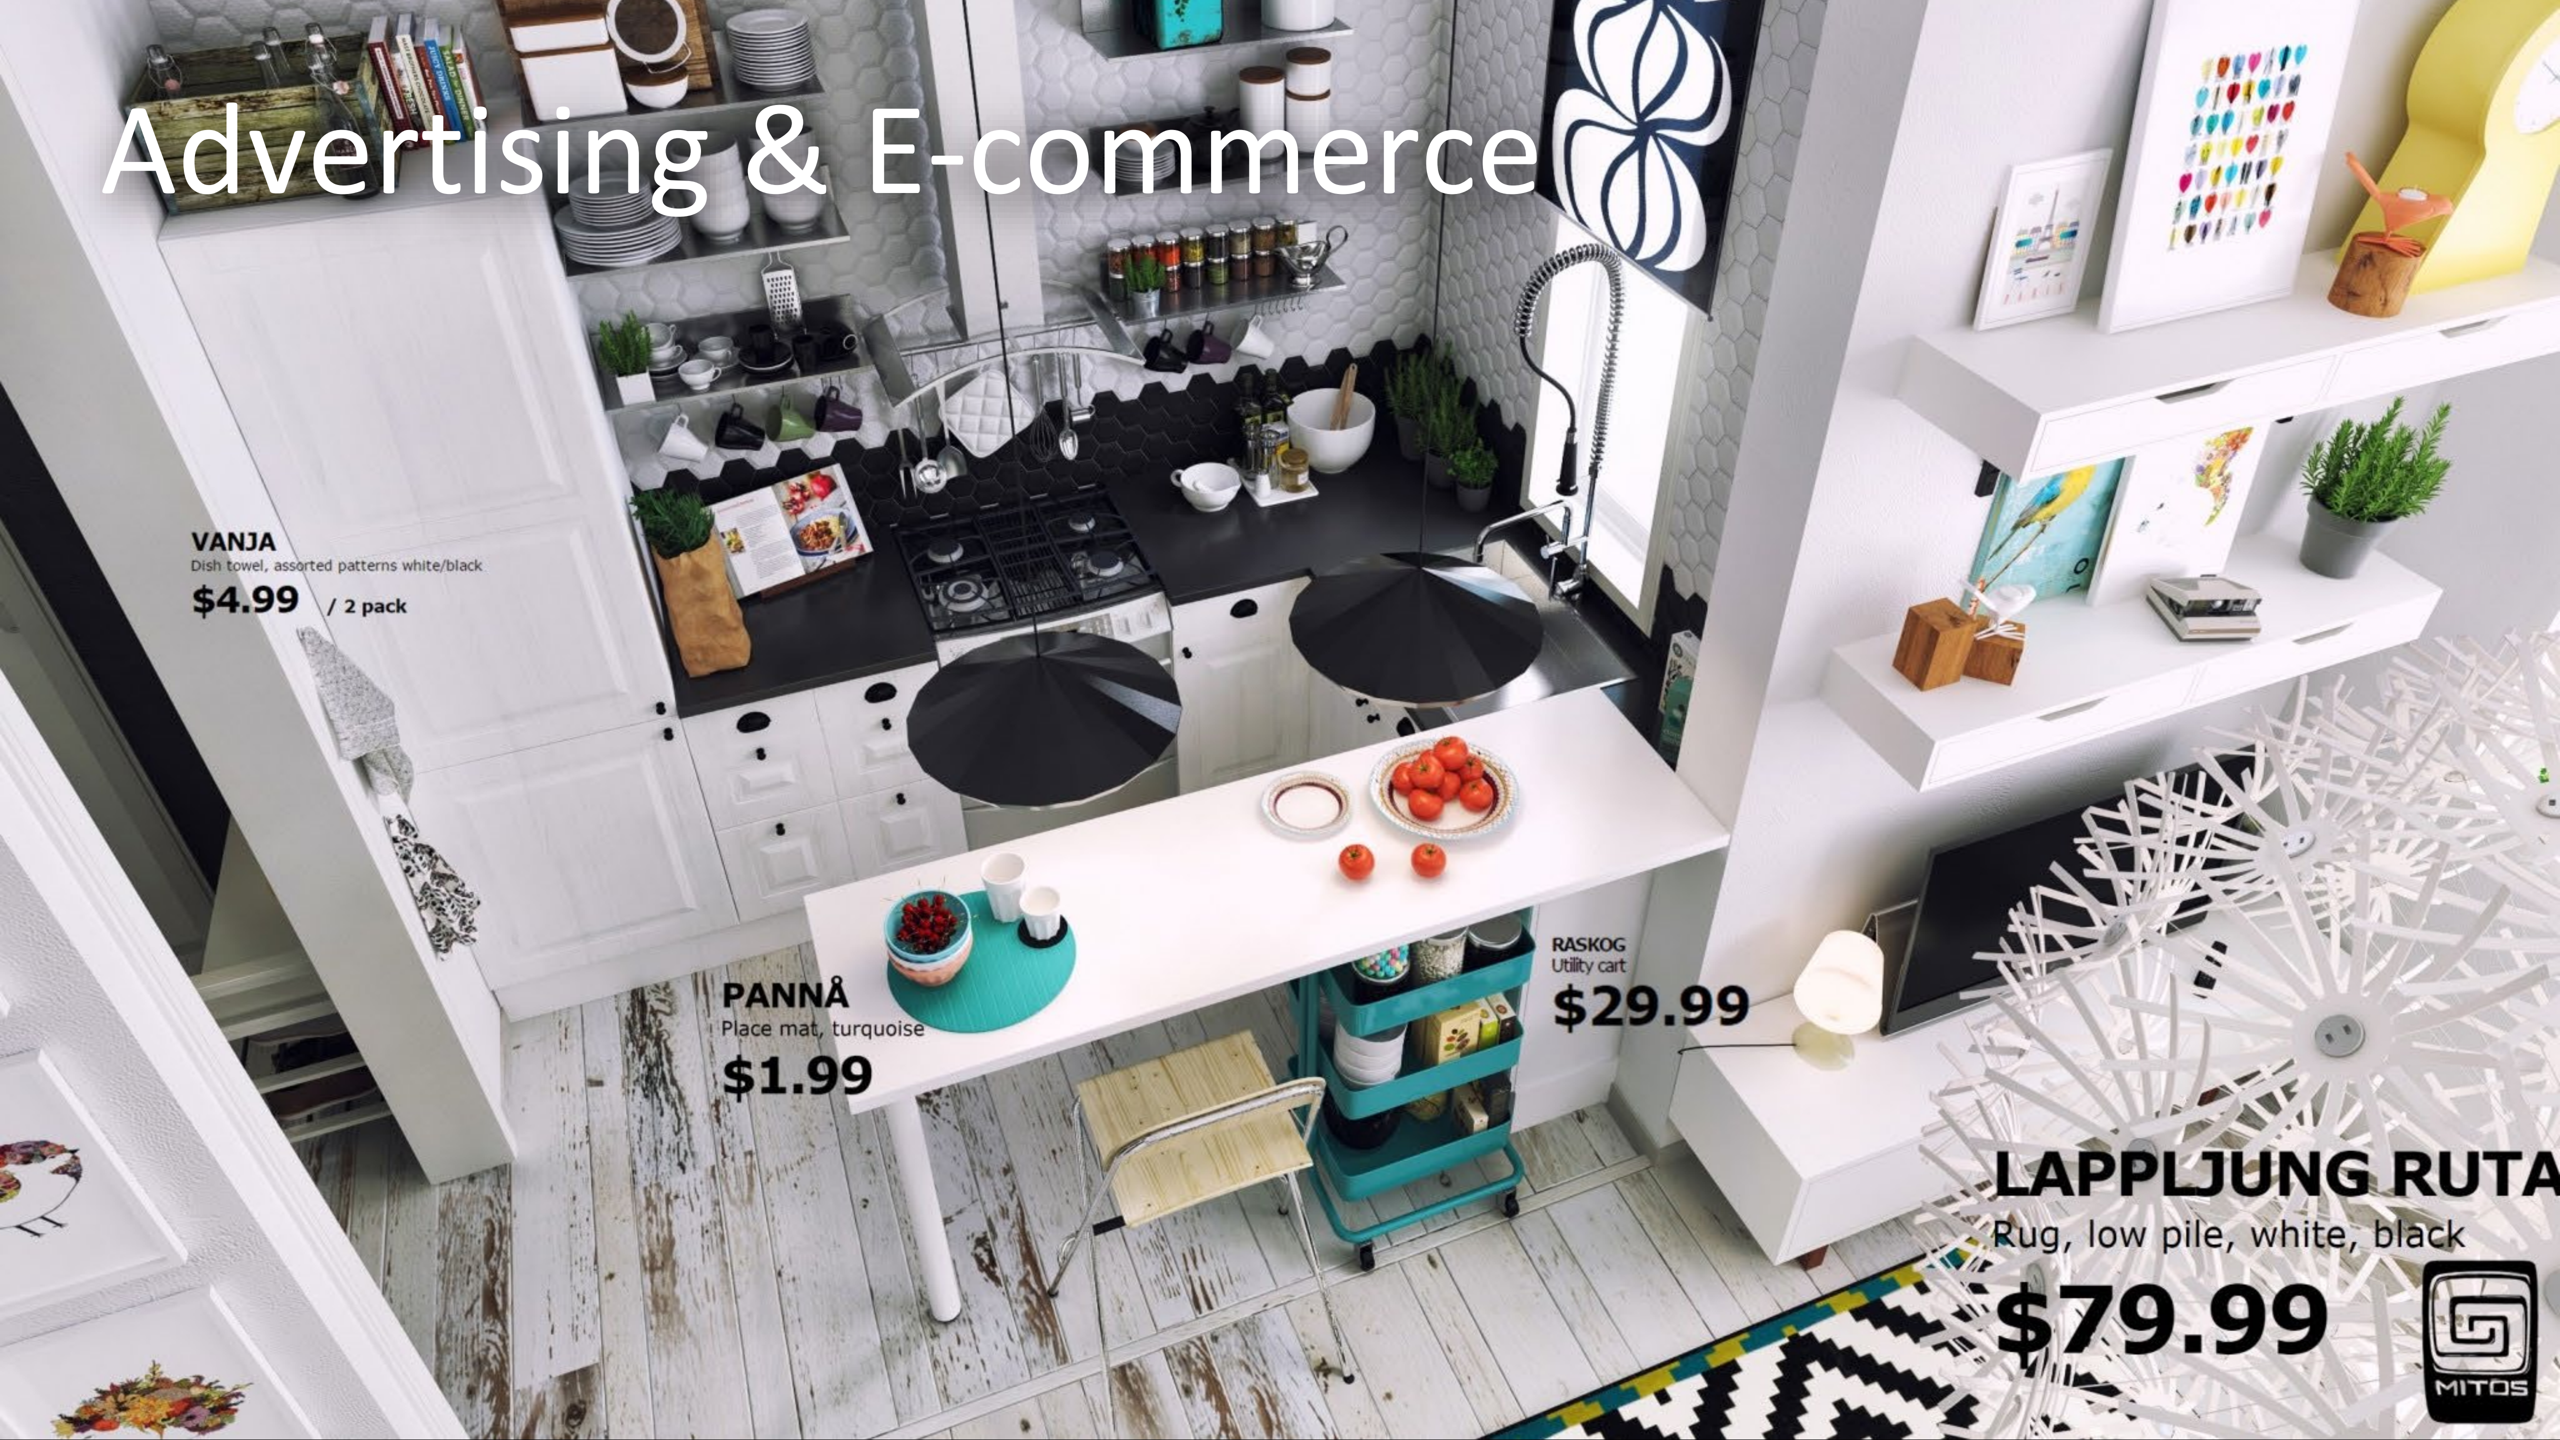

Cultural heritage

Wojciech Jarosz

Digital fabrication

Wojciech Jarosz

Scientific visualization

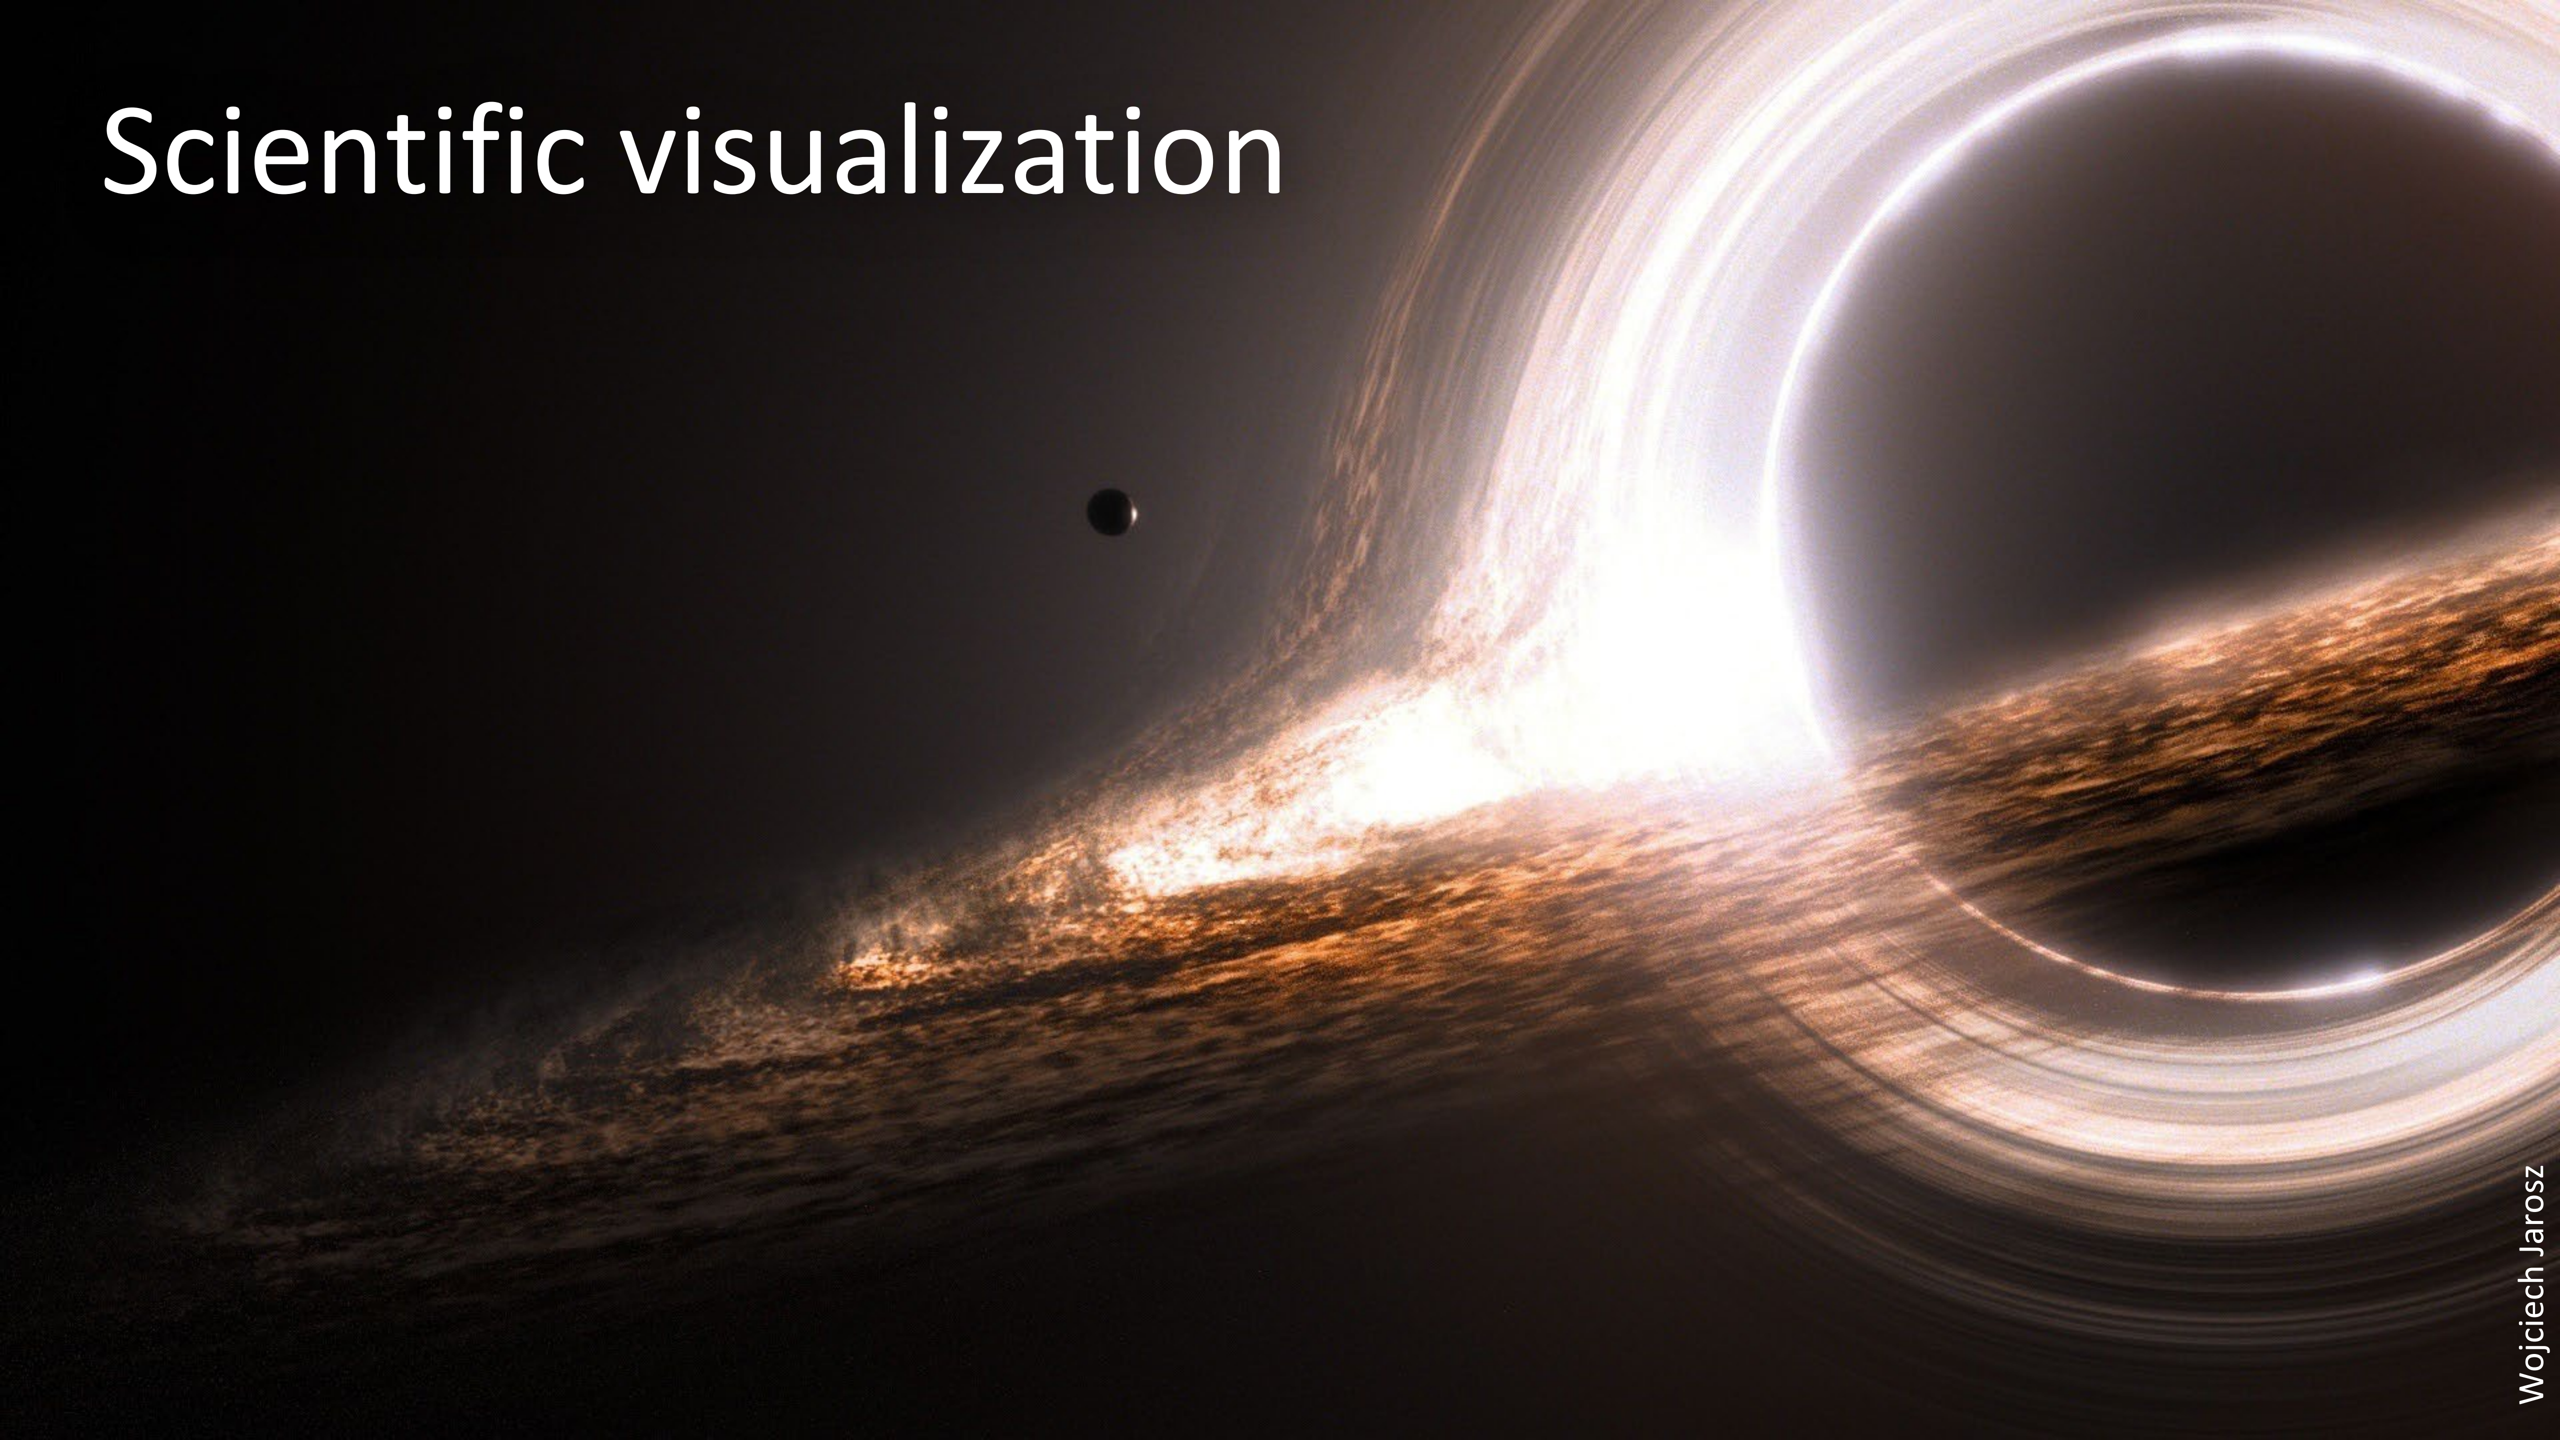

Wojciech Jarosz

Scientific imaging

rendering computational light transport

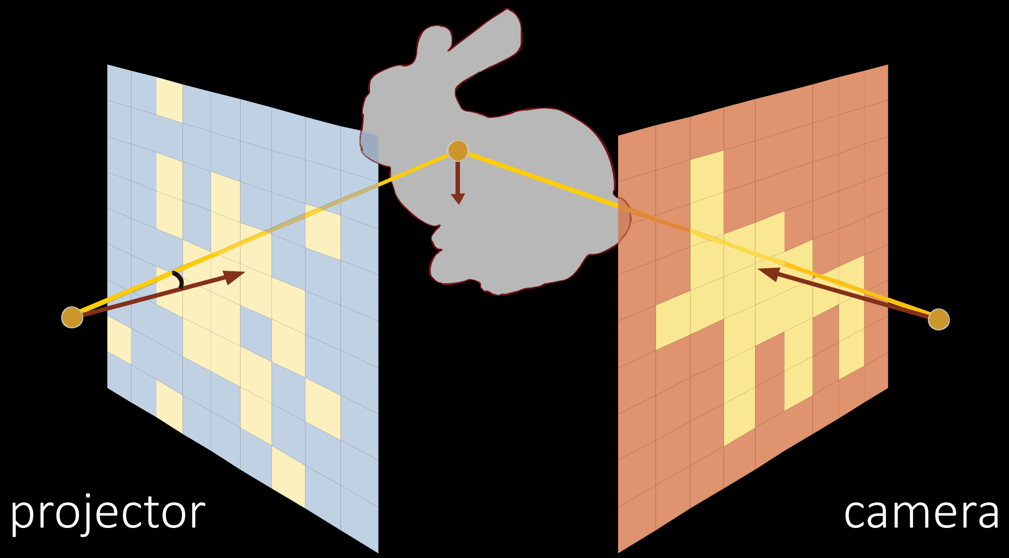

Used by CMU imaging projects:

convolutional DOT

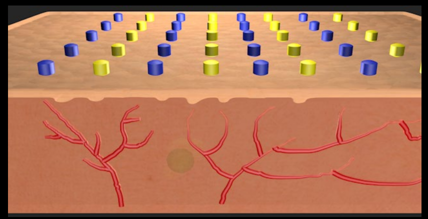

coded coherence

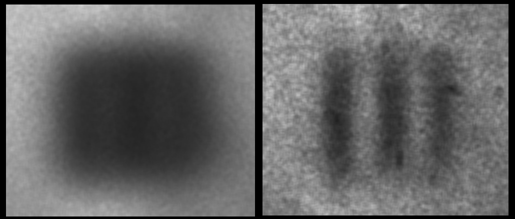

coded spectrum

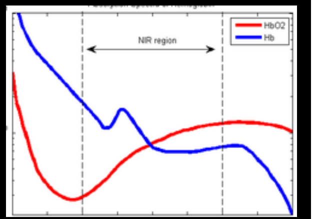

rendering time-of-flight sensors

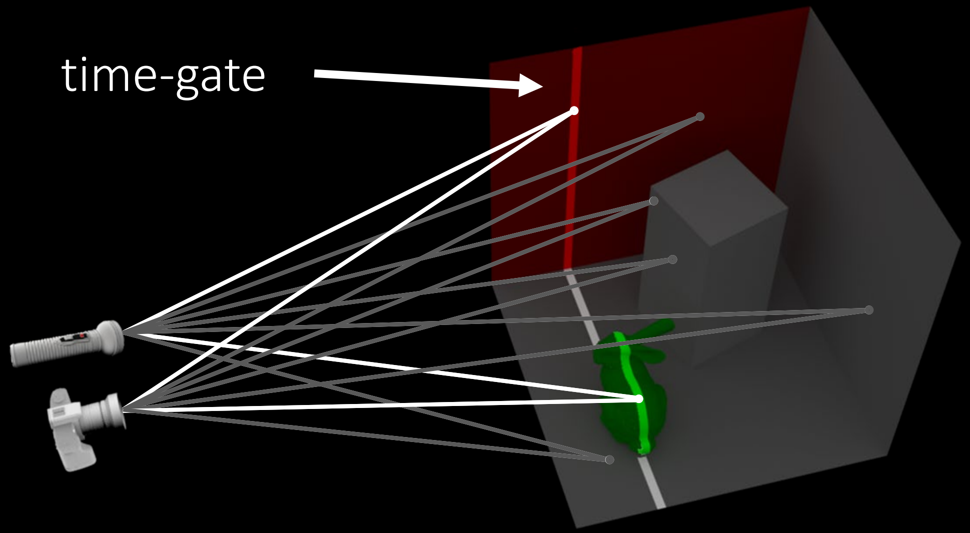

Used by CMU imaging projects:

ToF DOT

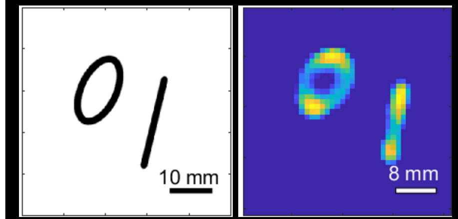

all-photon imag.

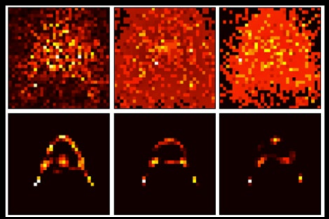

differential SPAD

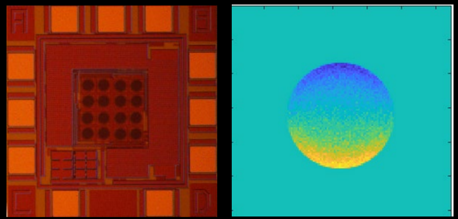

Rendering wave effects

speckle: noise-like pattern

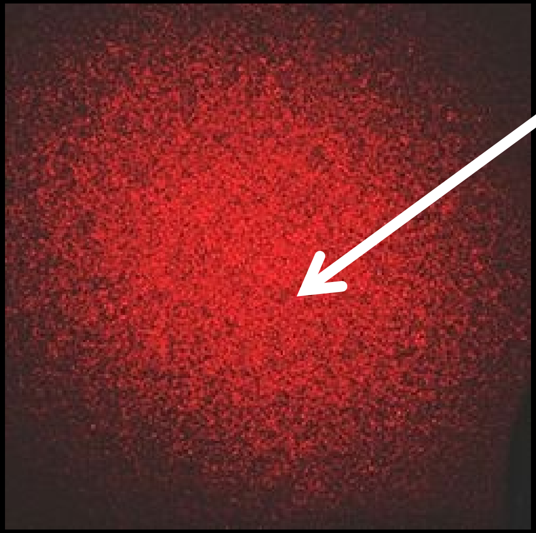

what real laser images look like

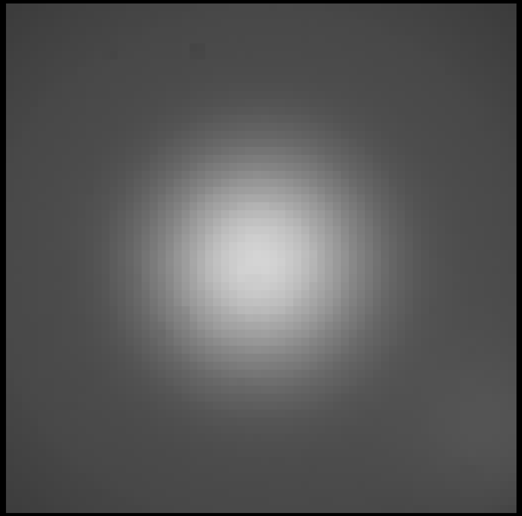

what standard rendered images look like

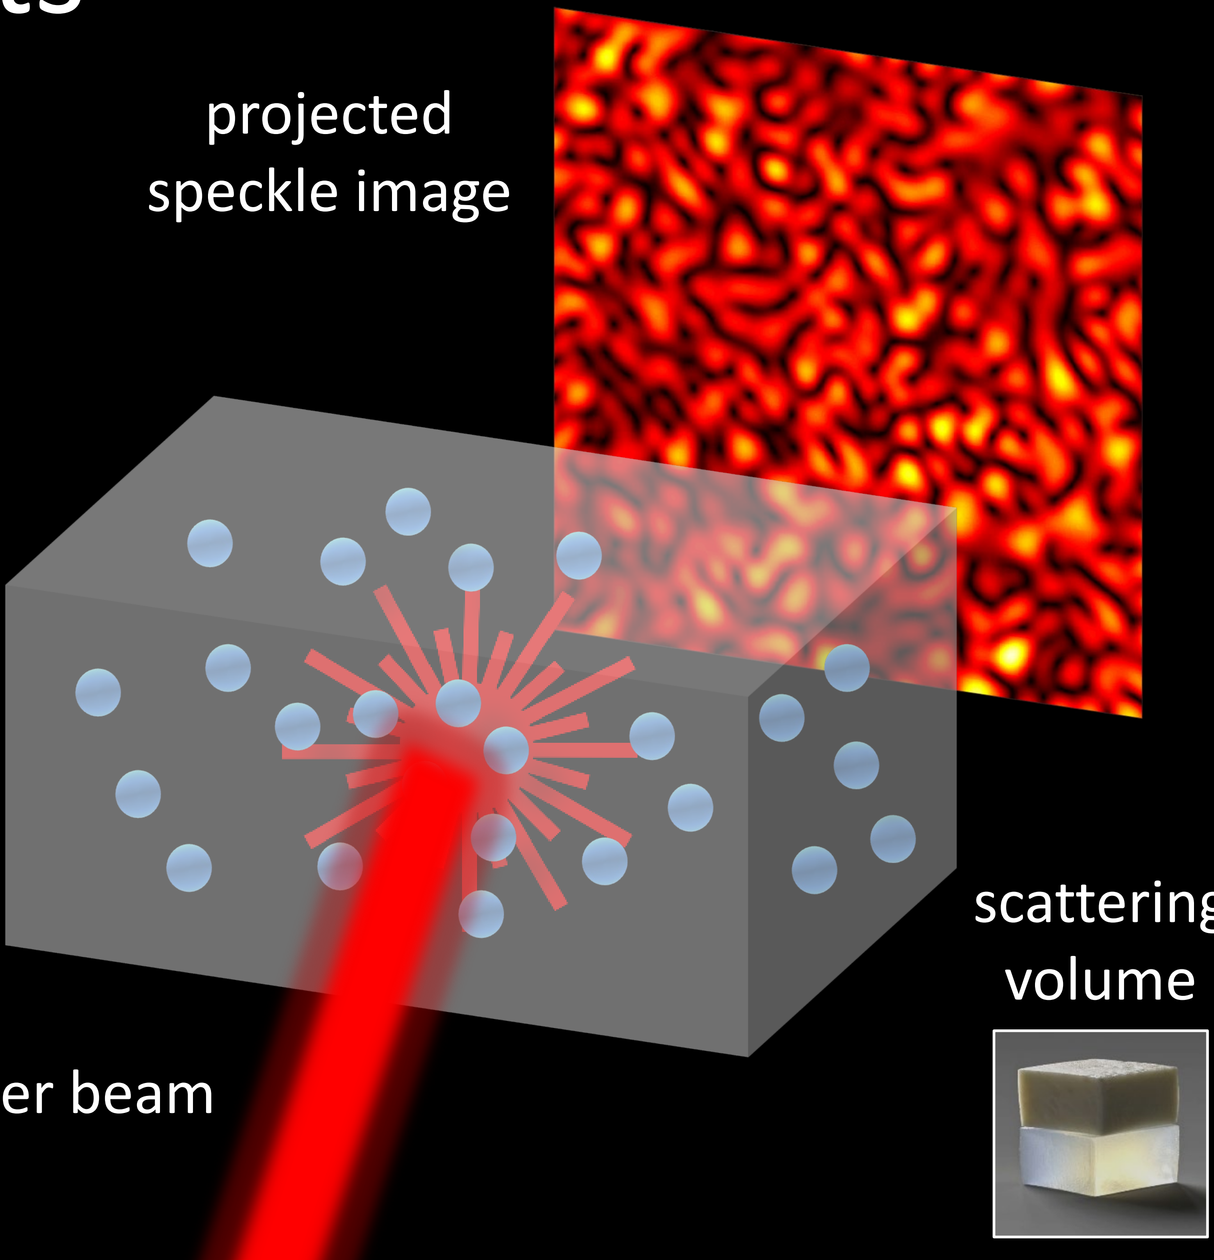

Application: fluorescence Microscopy

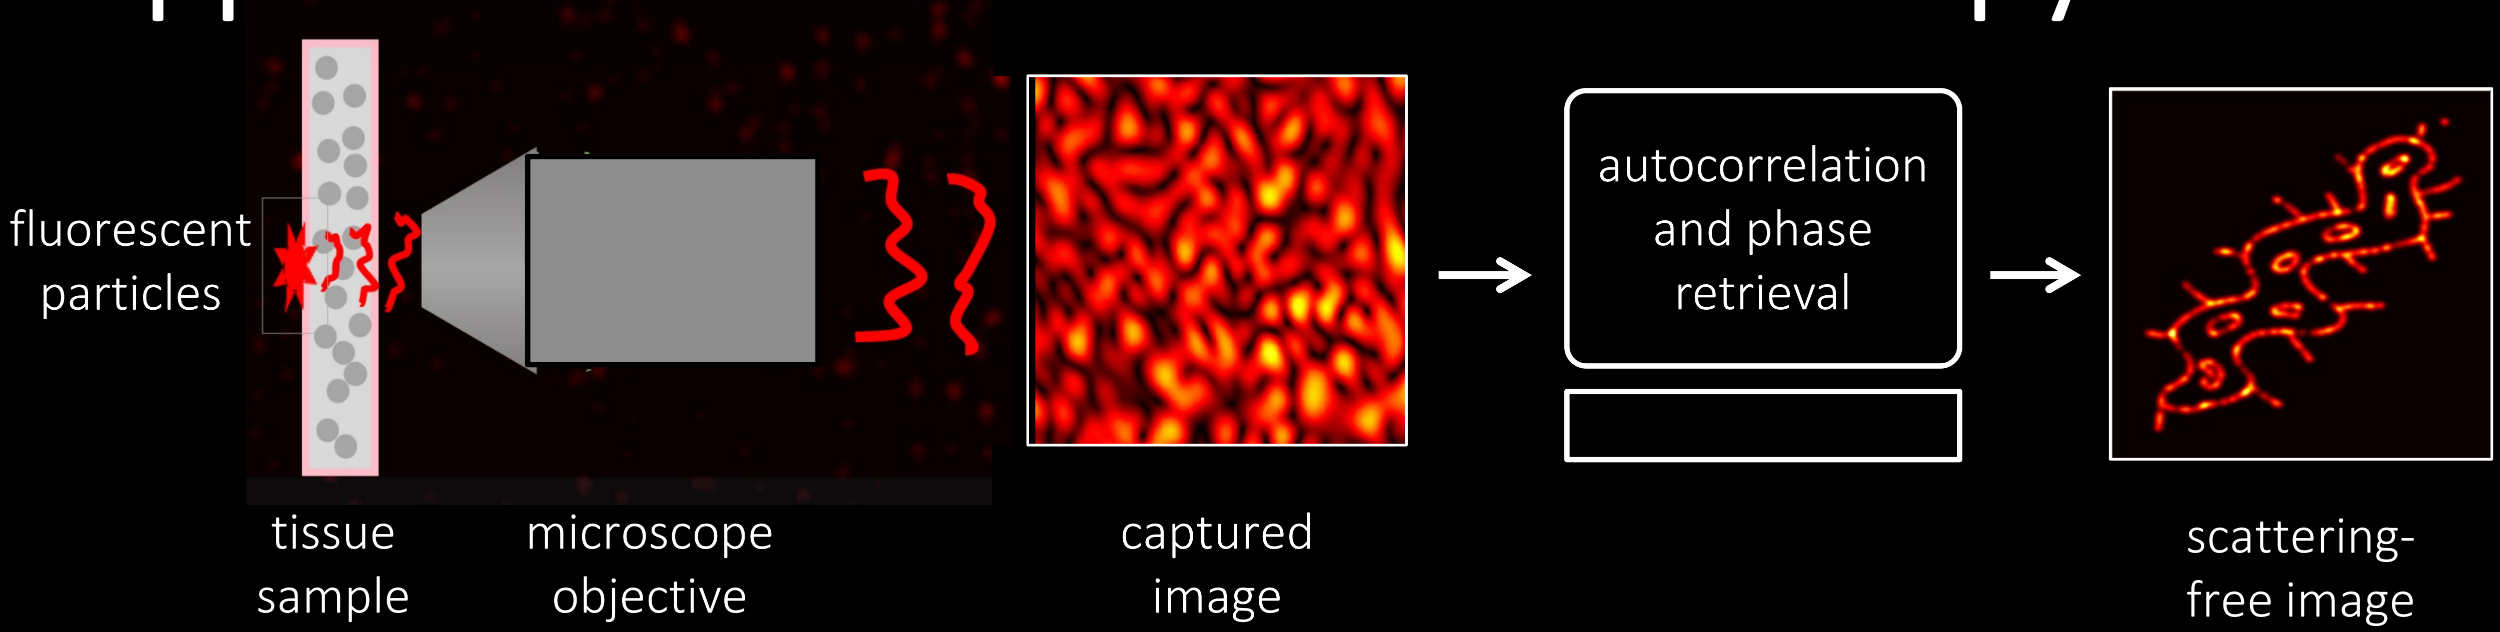

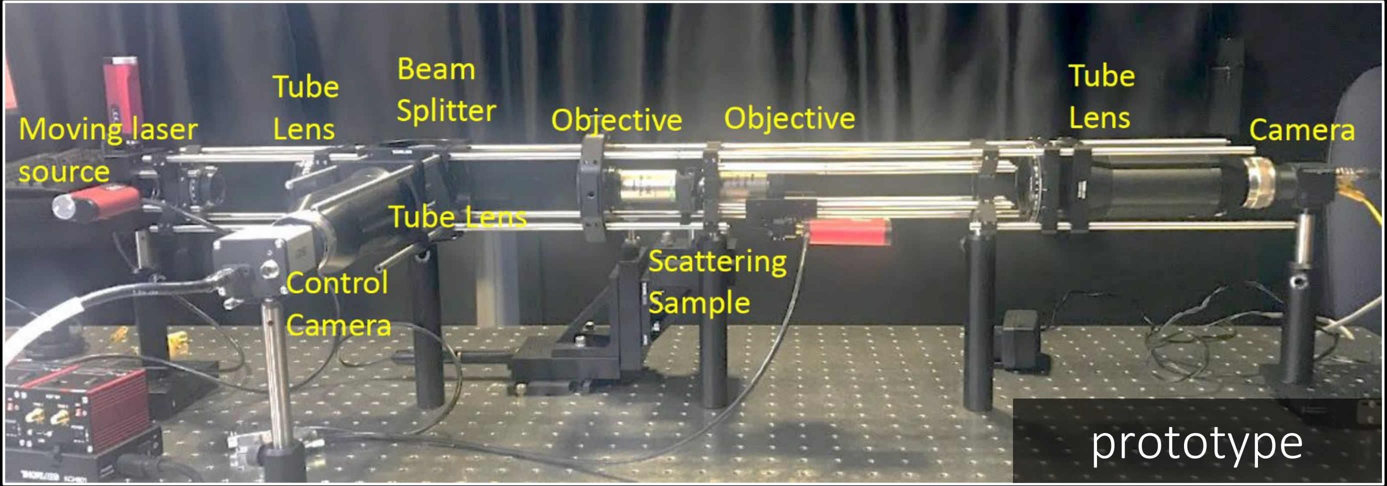

Performance strongly depends on:

- speckle statistics
- image priors
- tissue parameters

Rendering-assisted exploration and new algorithms!

Rendering eikonal transport

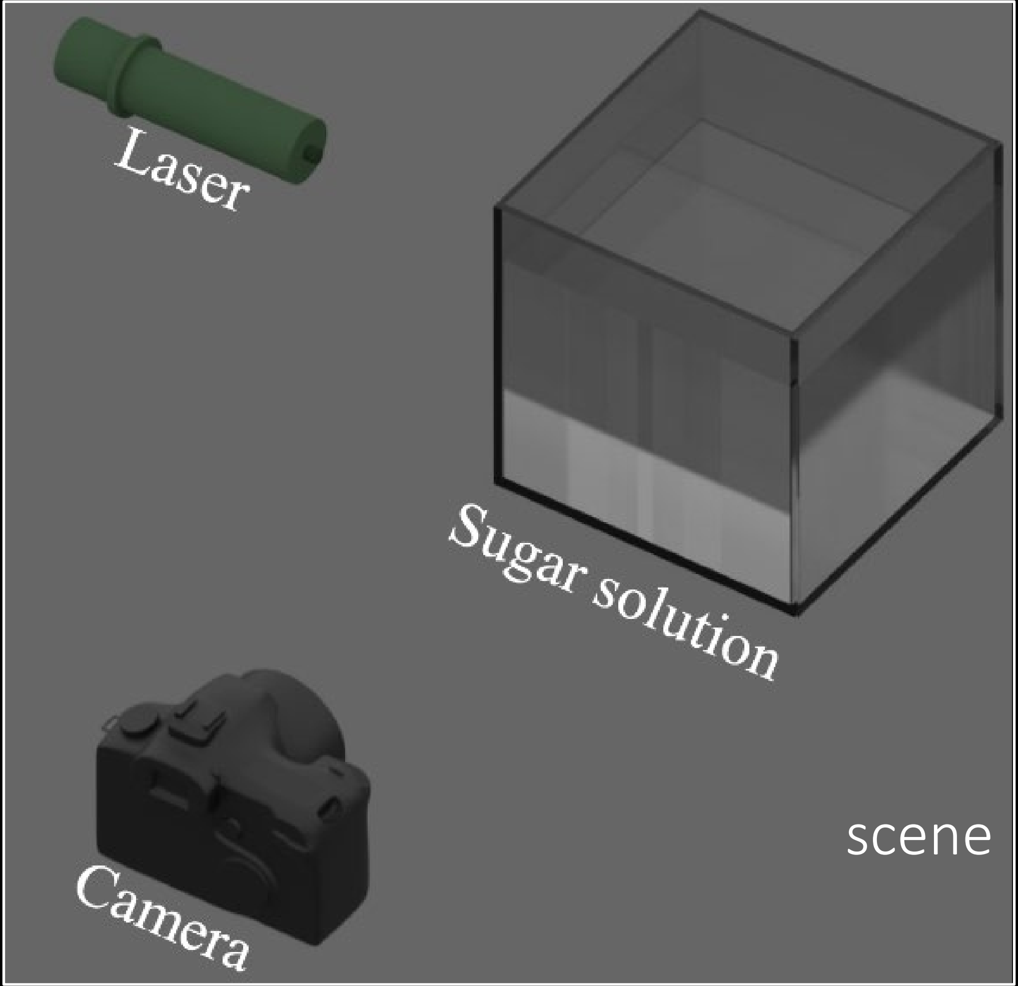

Application: acousto-optics

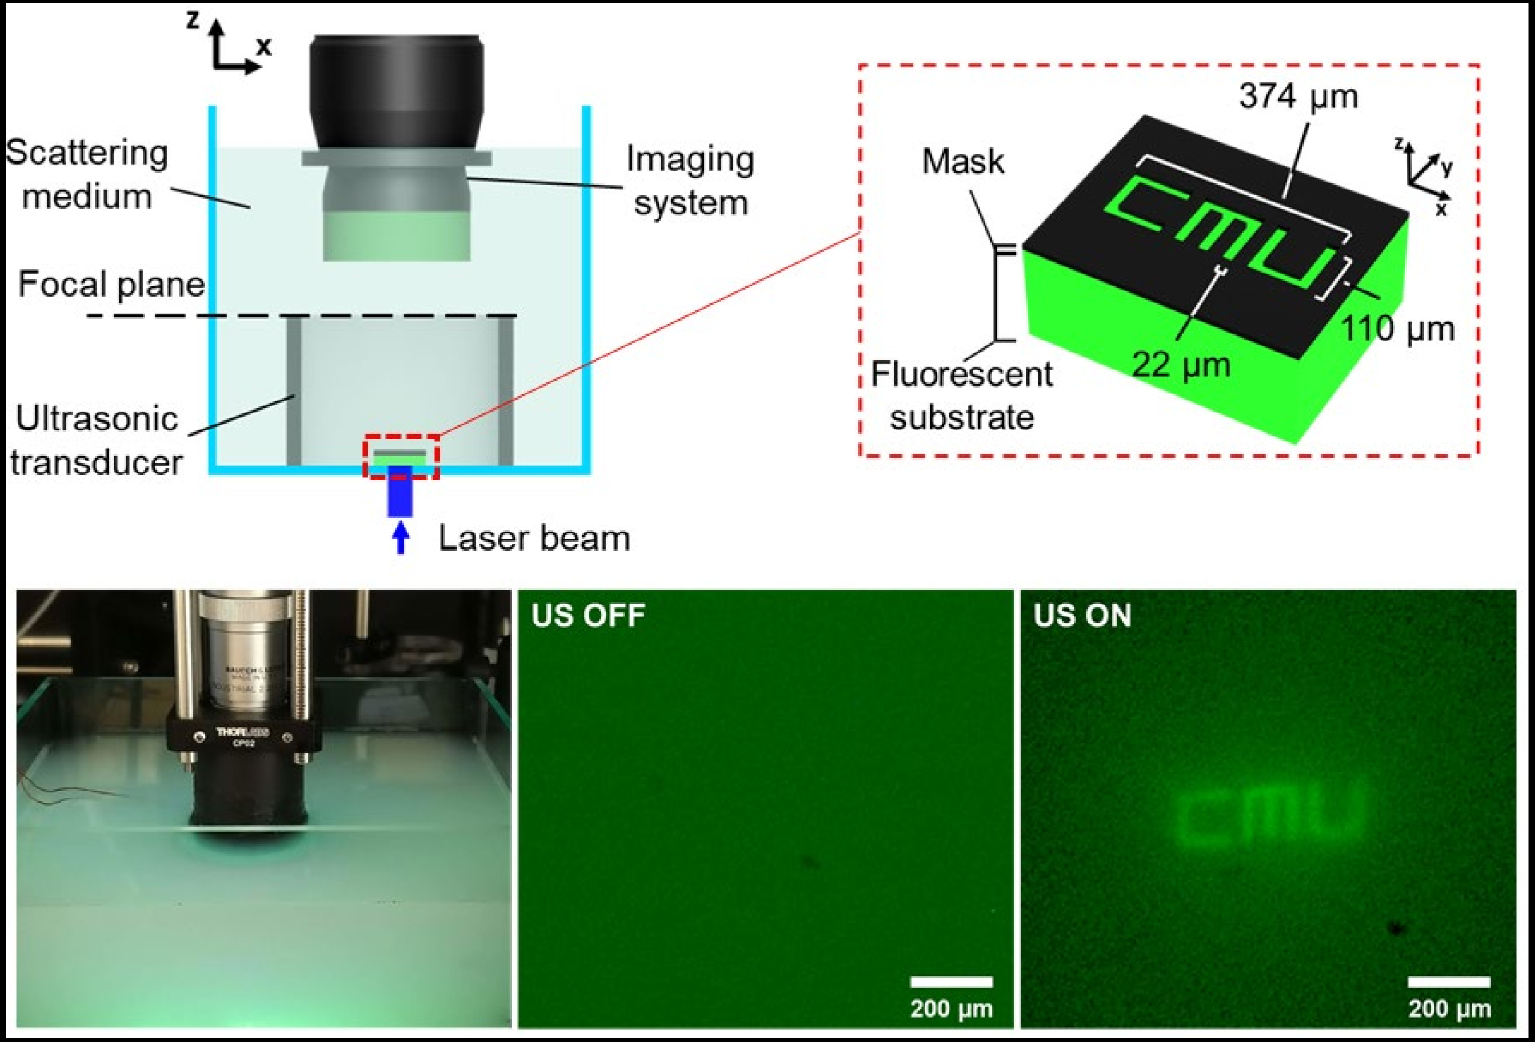

real capture

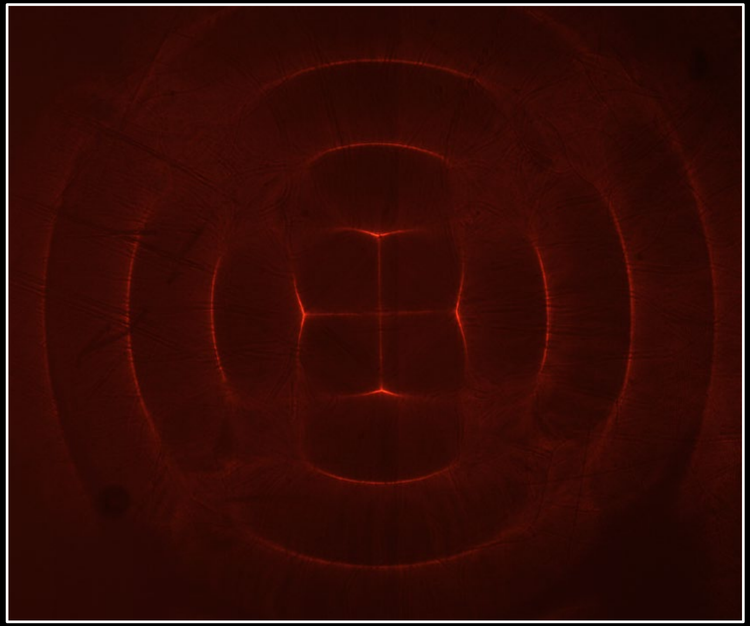

our algorithm

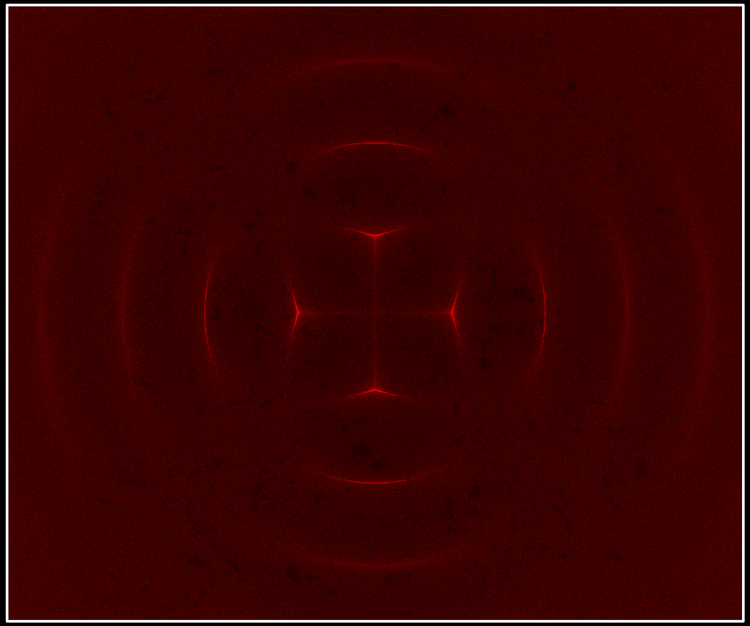

Forward rendering

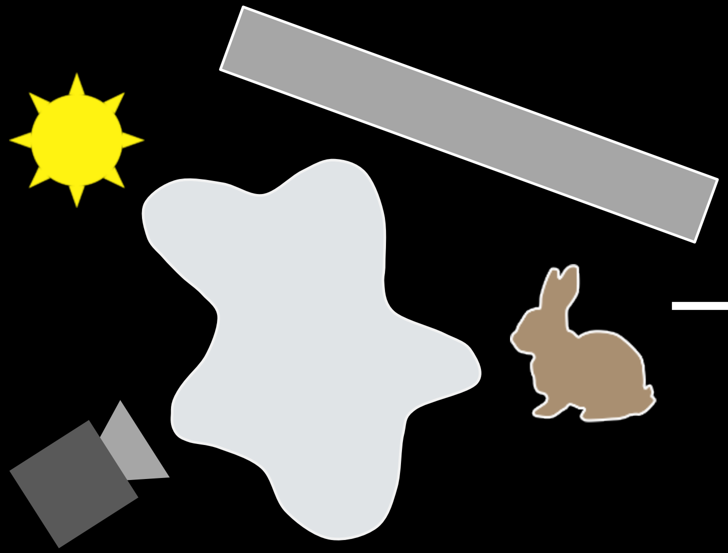

digital scene specification
(geometry, materials,
optics, light sources)

physically-accurate
rendering

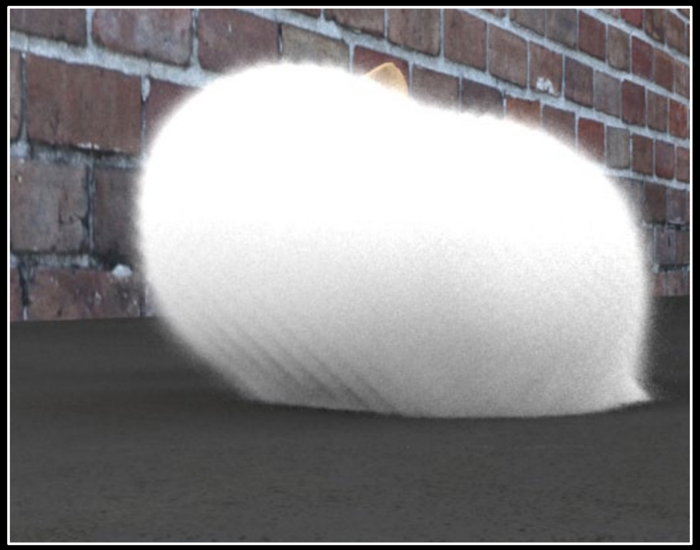

photorealistic
simulated image

Inverse rendering

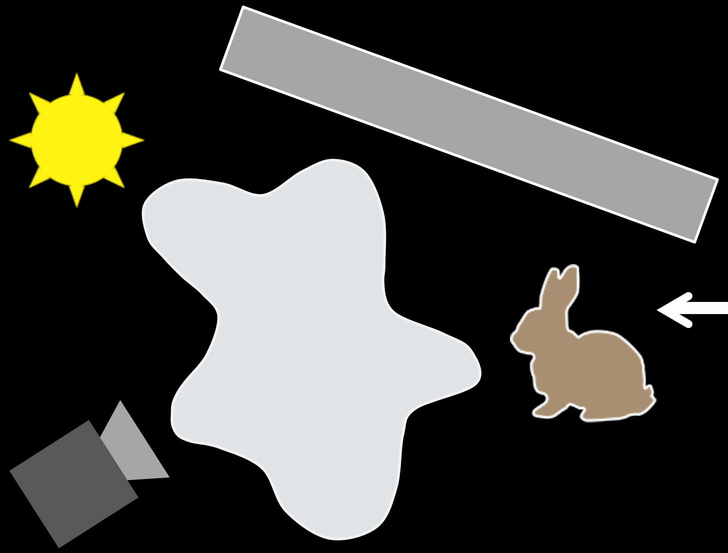

physically-accurate inverse rendering

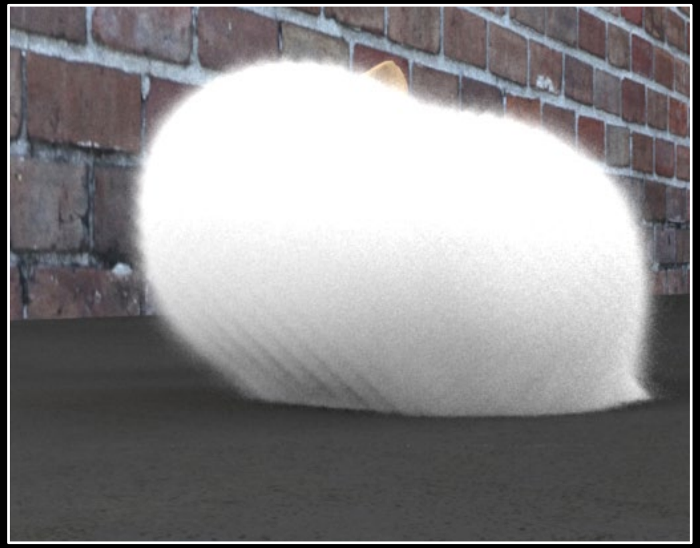

digital scene specification (geometry, materials, camera, light sources)

photorealistic synthetic image

image measurement

Differentiable rendering

Original image | Derivative image

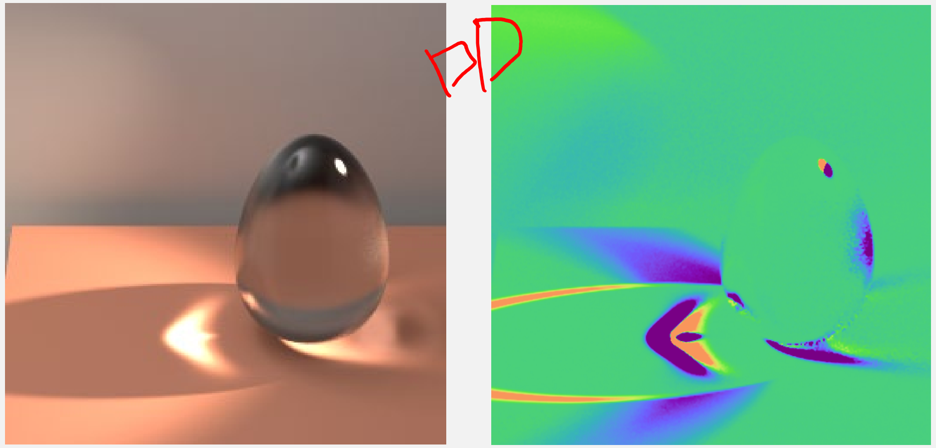

Original image | Derivative image

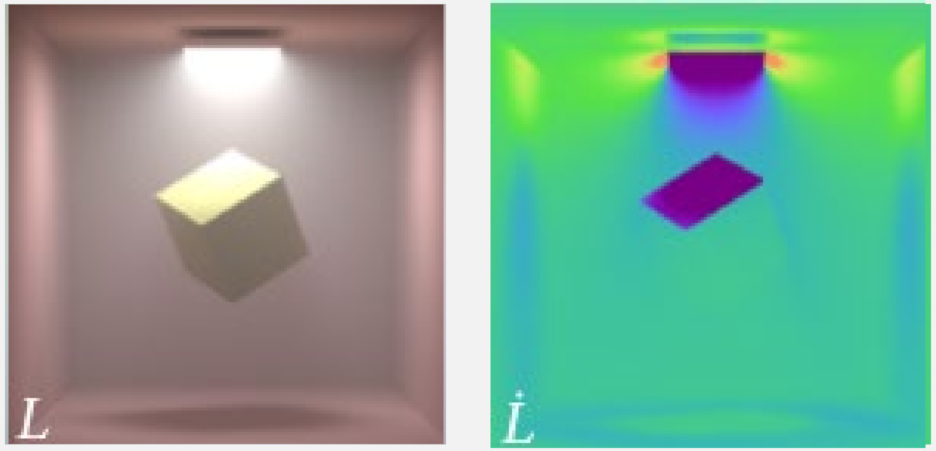

Original image | Derivative image

Original image | Derivative image

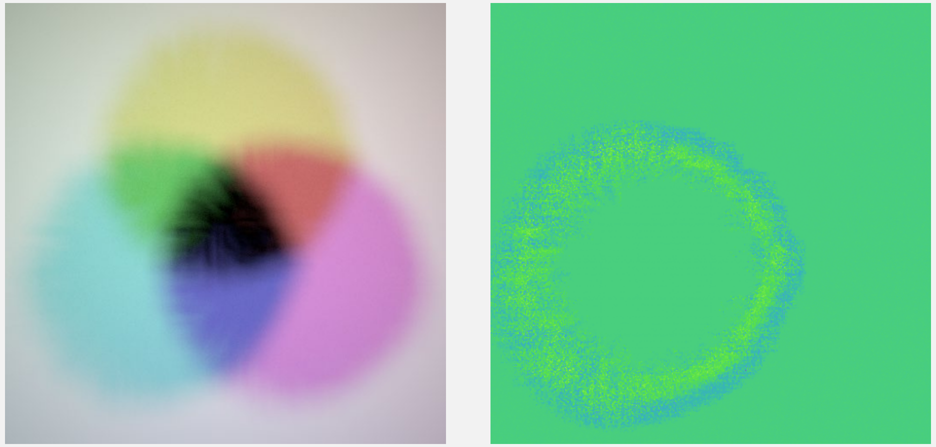

Application: shape optimization

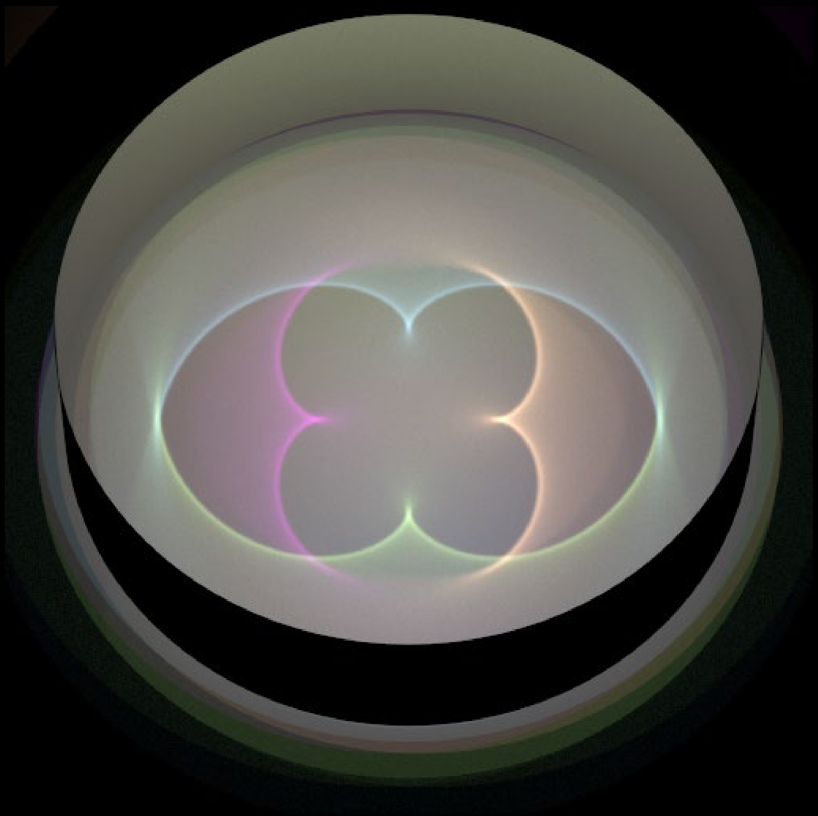

Initial

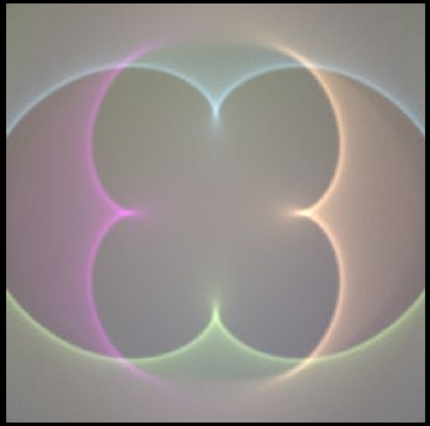

Target image

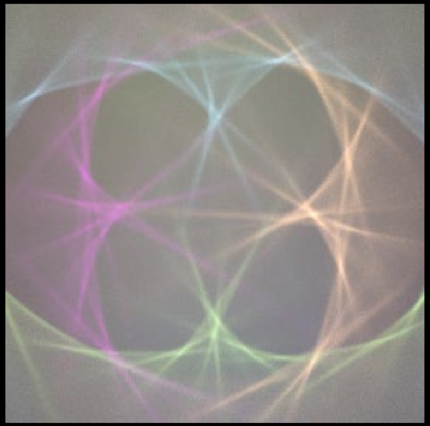

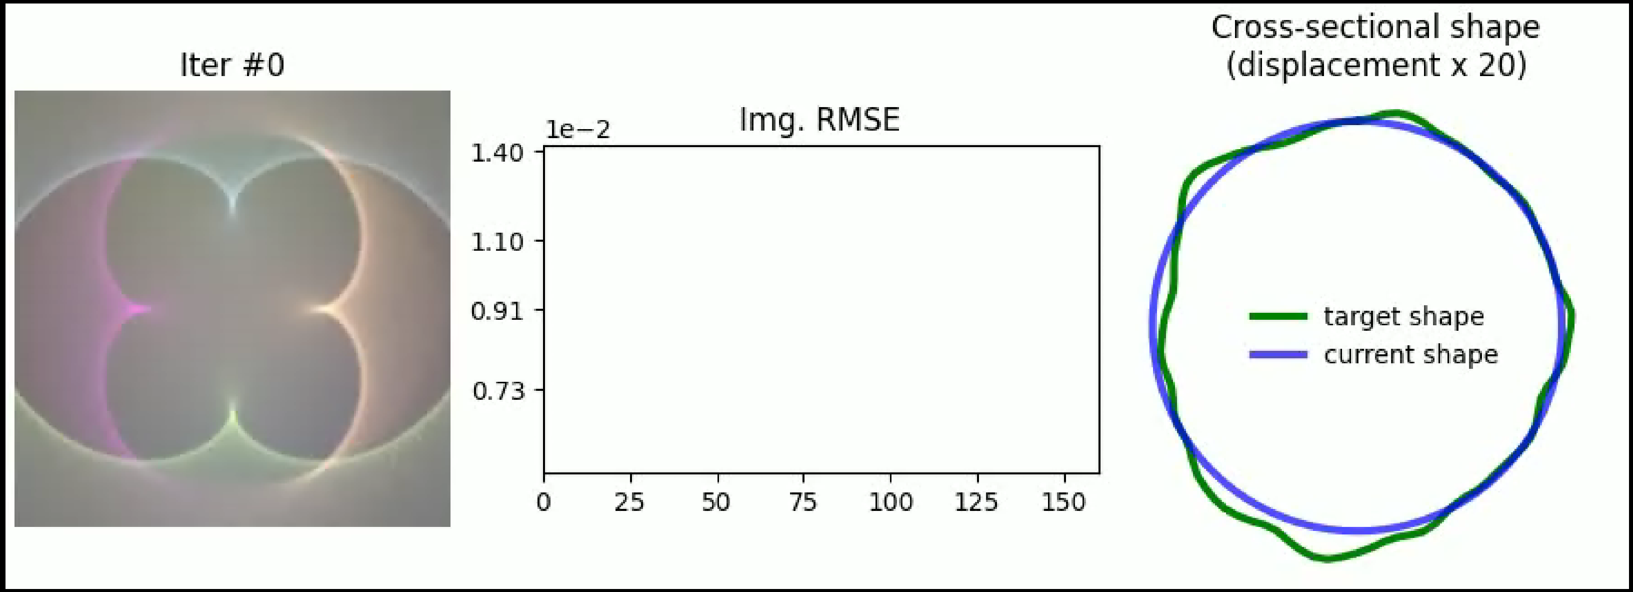

Application: non-line-of-sight imaging

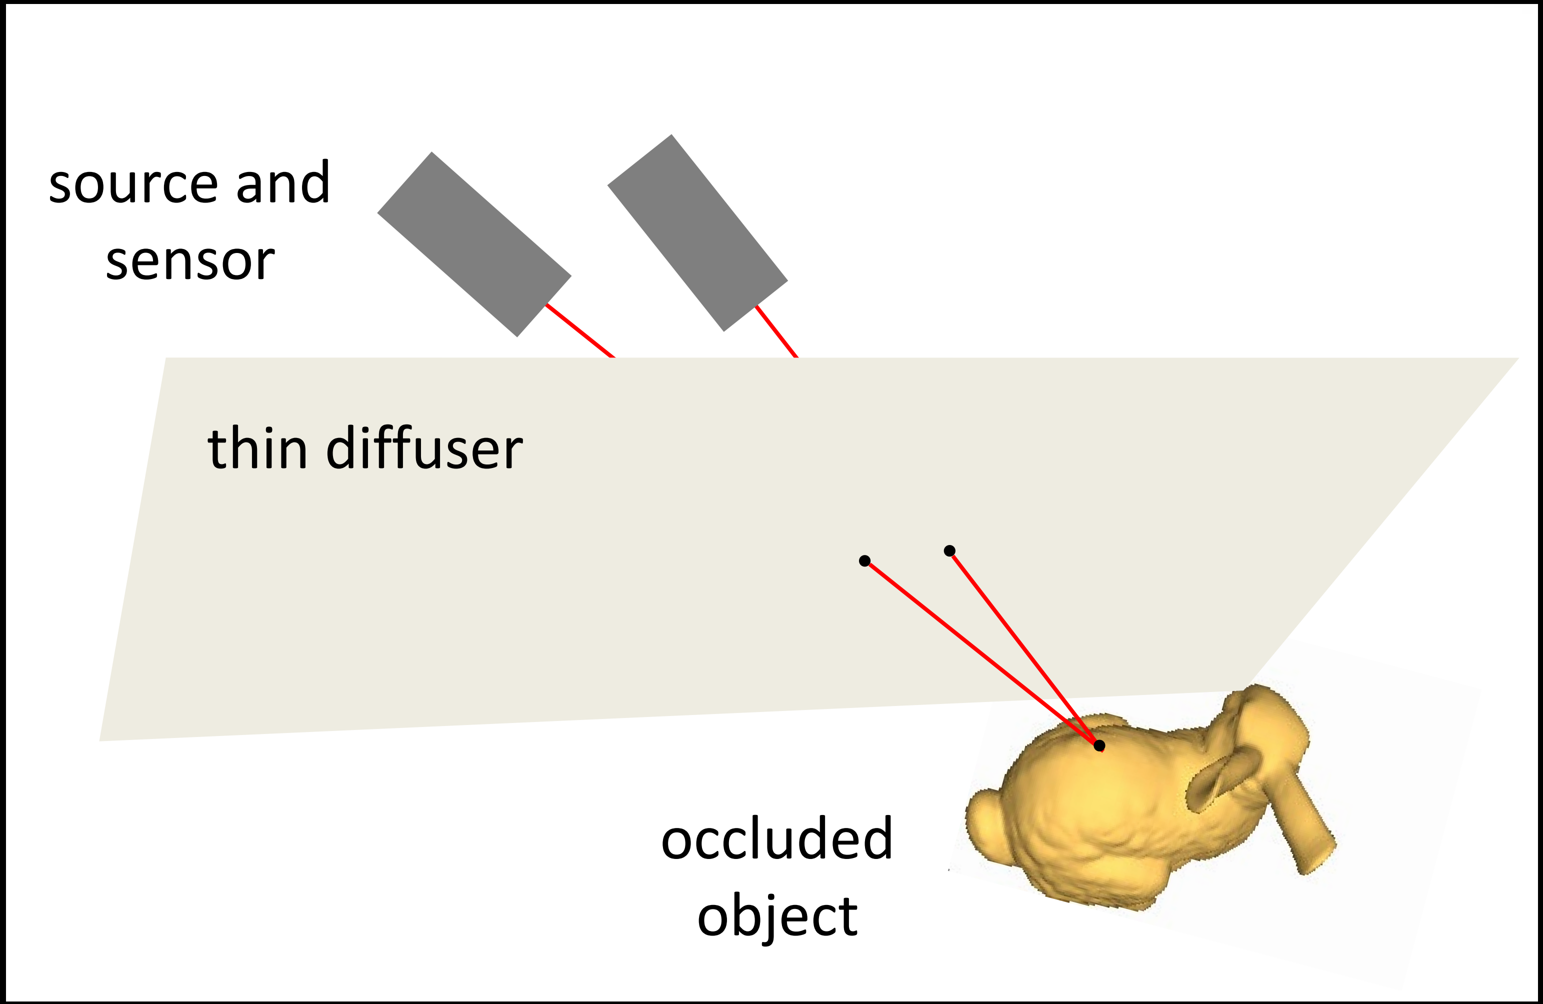

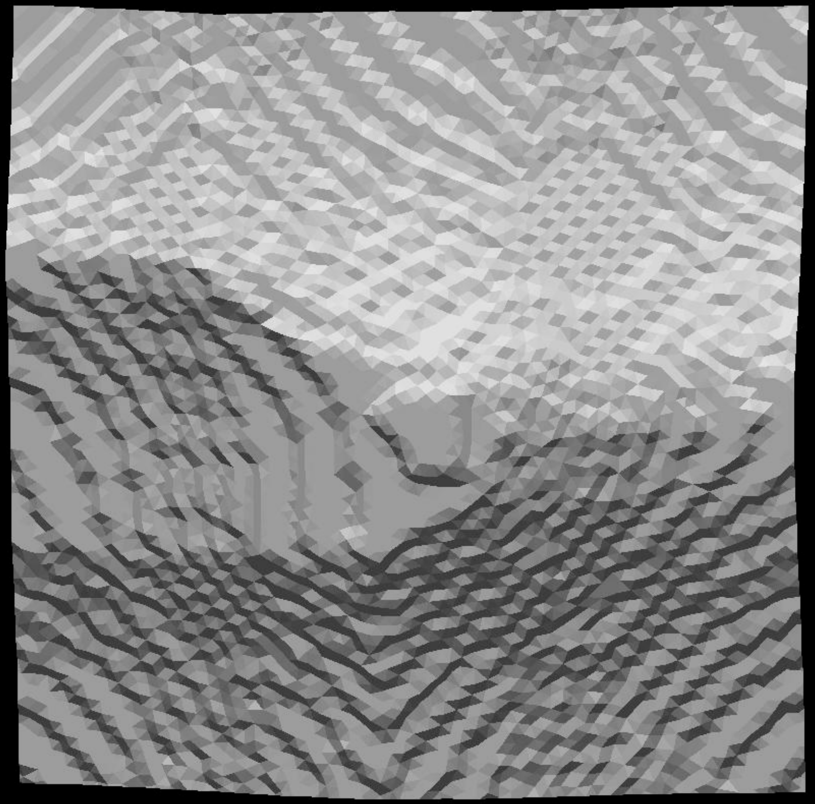

reconstruction evolution

Application: non-invasive tomography

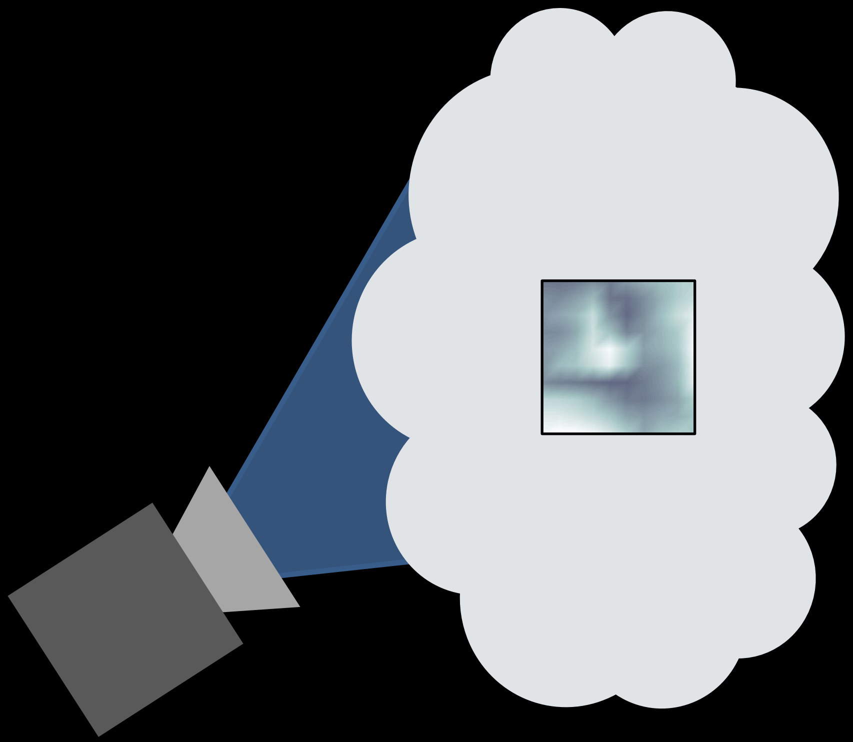

camera

thick smoke cloud

simulated camera measurements

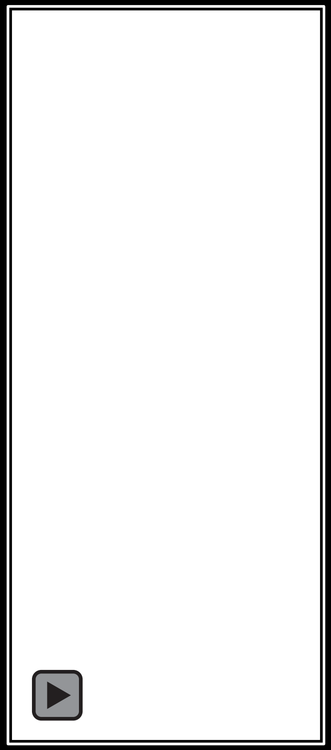

reconstructed cloud volume

slice through the cloud

Application: vision and machine learning

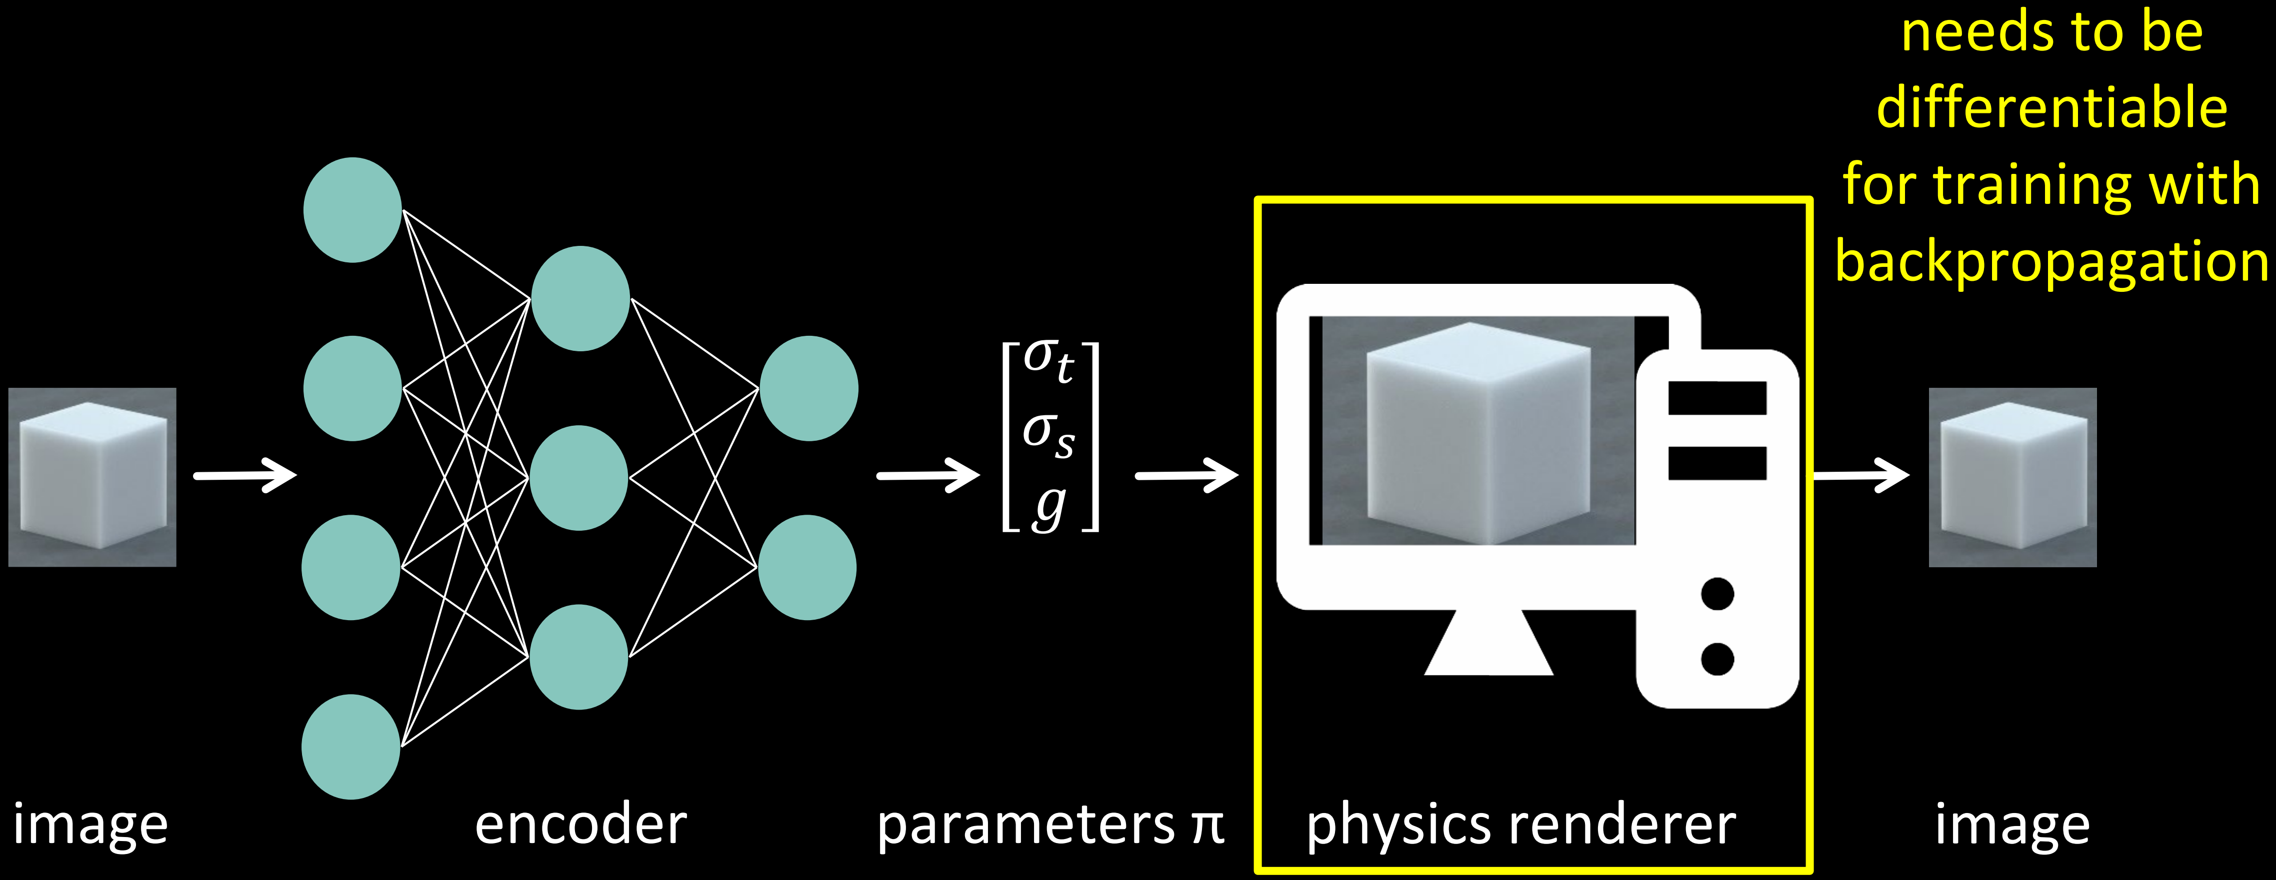

Application: vision and machine learning

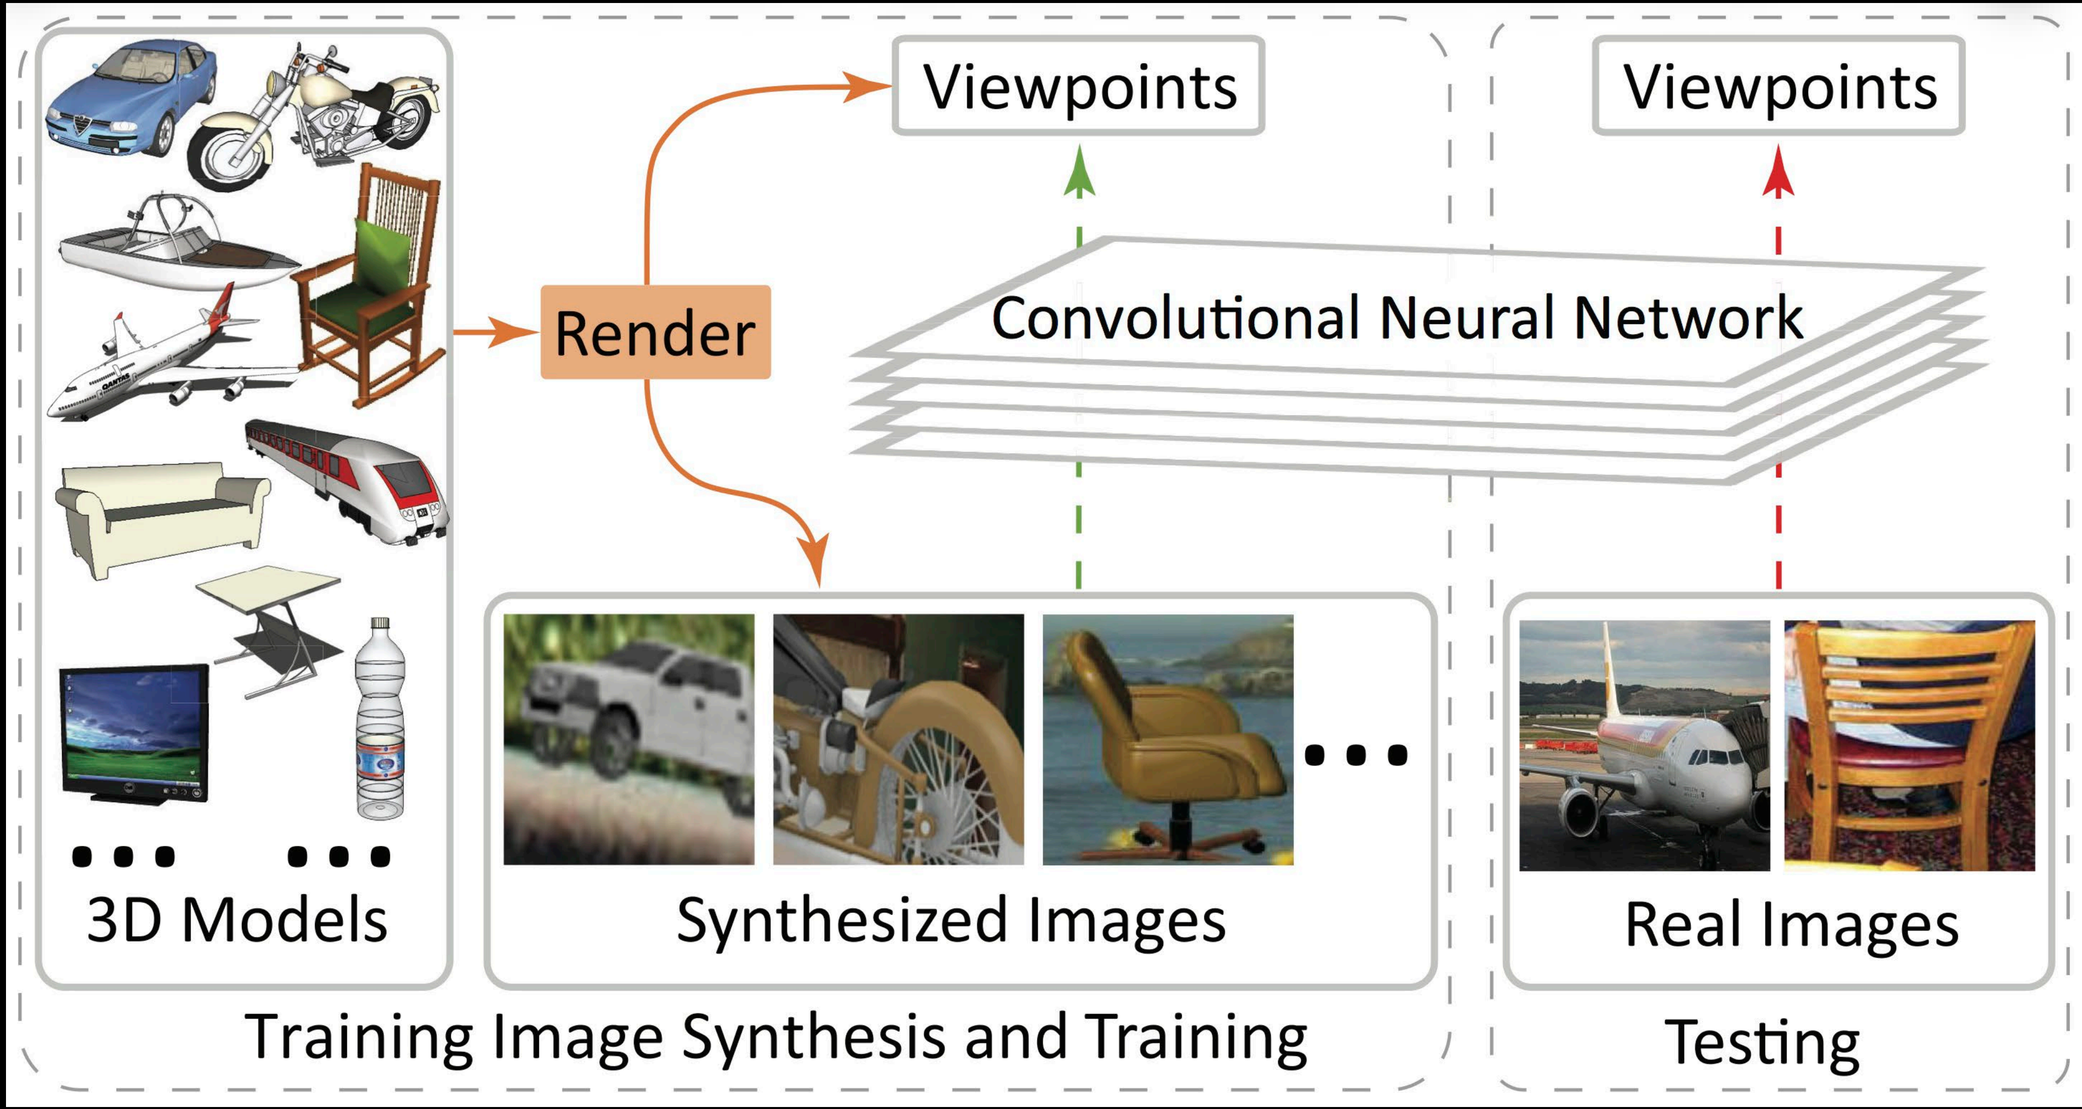

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views [Su et al. ICCV 2015]

Application: neural rendering

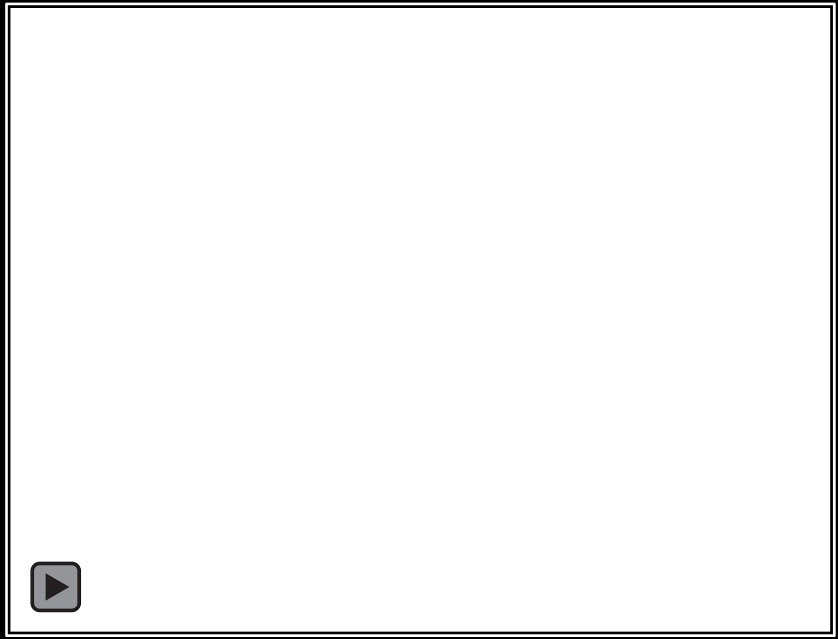
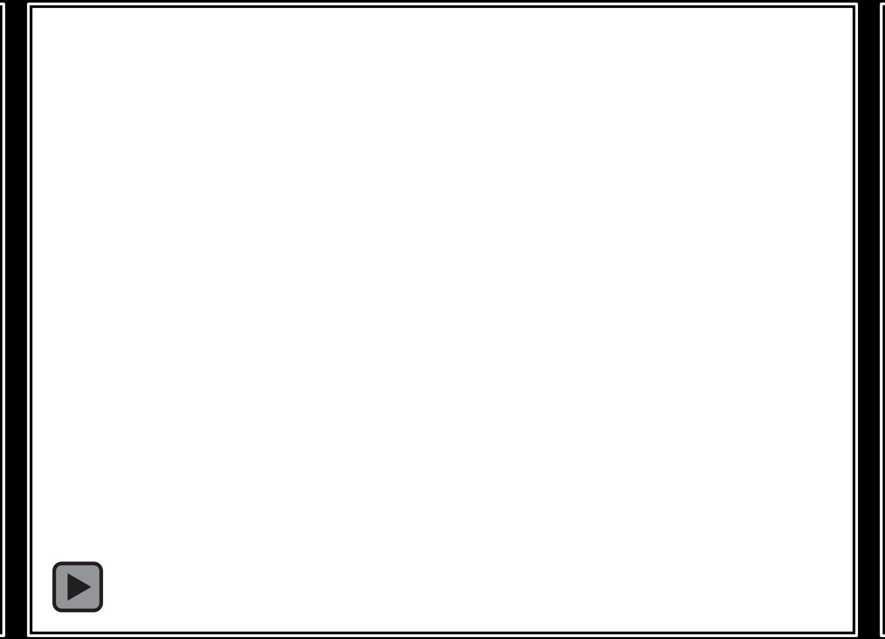
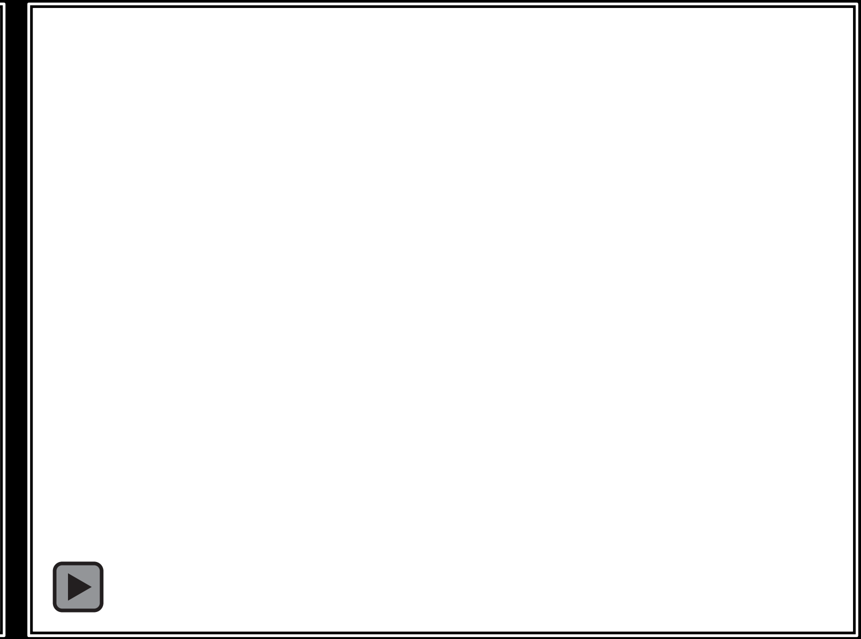

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis [MildenHall et al. ECCV 2020]

What is this class about?

Producing realistic images by:

- *simulating* light transport (global illumination)
- *simulating* light-material interactions (appearance modeling)

Understanding why things look the way they do:

- Why is the sky blue?
- Why is the grass green?
- Why does metal look different than marble?

Course fast-forward and logistics

Please take the start-of-semester survey!

- Posted on Piazza as well:

https://docs.google.com/forms/d/e/1FAIpQLSdZONUsOzpfkLGTcpuc8nlHRx7EkVIL_yGzXZt5xdwTsDlzxw/viewform

Course logistics

- Course website:

  http://graphics.cs.cmu.edu/courses/15-468

- Piazza for discussion and announcements (sign up!):

  https://piazza.com/class/ky96bnus9u54ul

- Canvas for homework submissions, Zoom links, and recordings:

  https://canvas.cmu.edu/courses/27795

- Slack server for real-time discussion:

  See Piazza for the invite link

Course fast-forward

Tentative syllabus at:

http://graphics.cs.cmu.edu/courses/15-468

- schedule and exact topics will almost certainly change during semester
- keep an eye out on the website for updates

Topics to be covered

Basics of ray tracing:

- trace-intersect recursions
- basic camera and illumination models
- shading
- intersection queries
- texture mapping

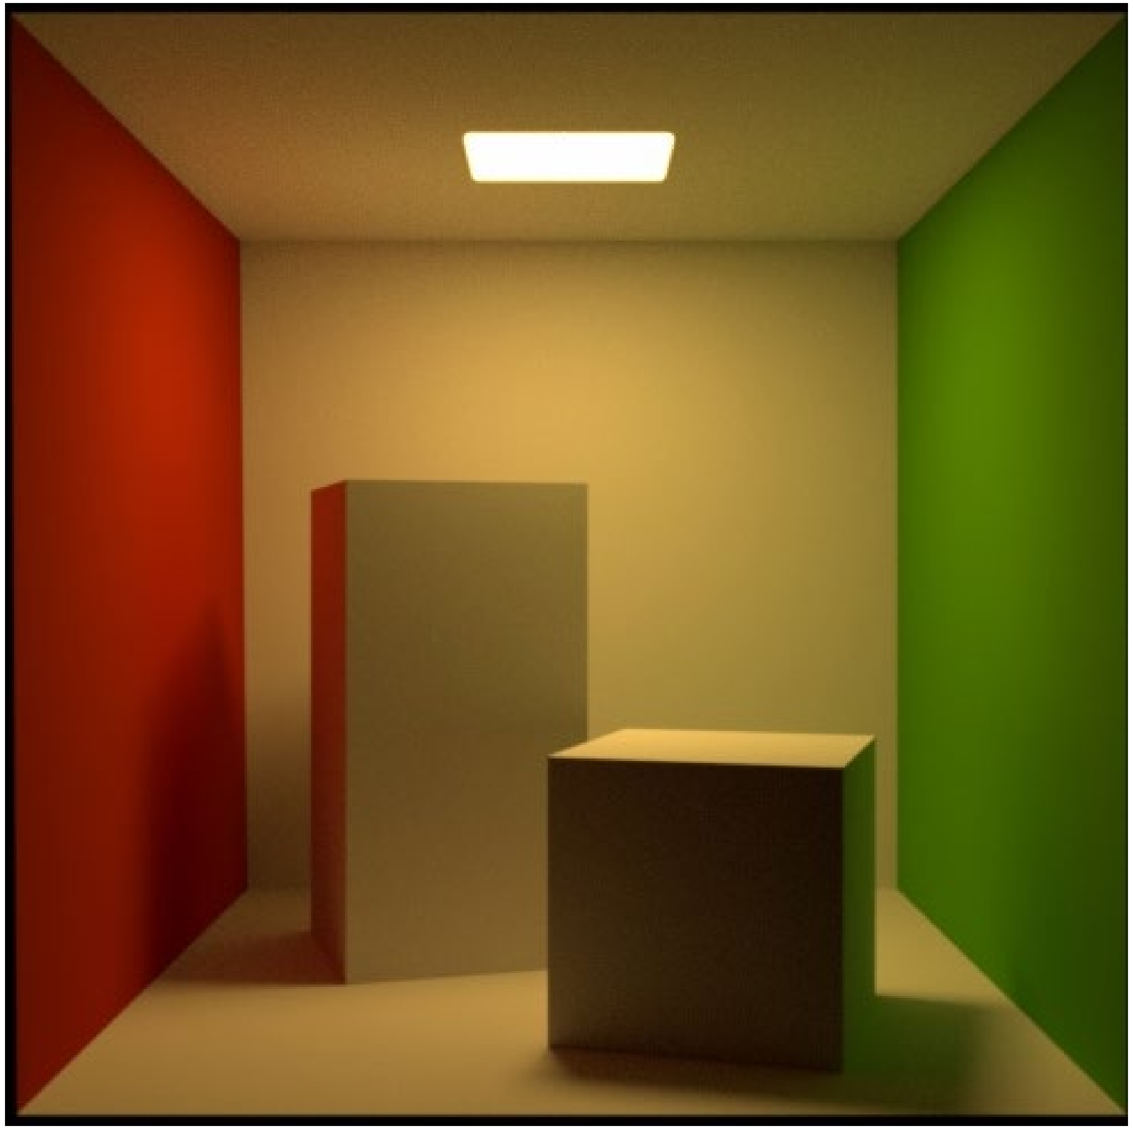

Topics to be covered

Theory of light transport and materials:

- rendering equation
- radiative transfer equation
- path integral formulations
- microfacet reflectance models
- statistical scattering models

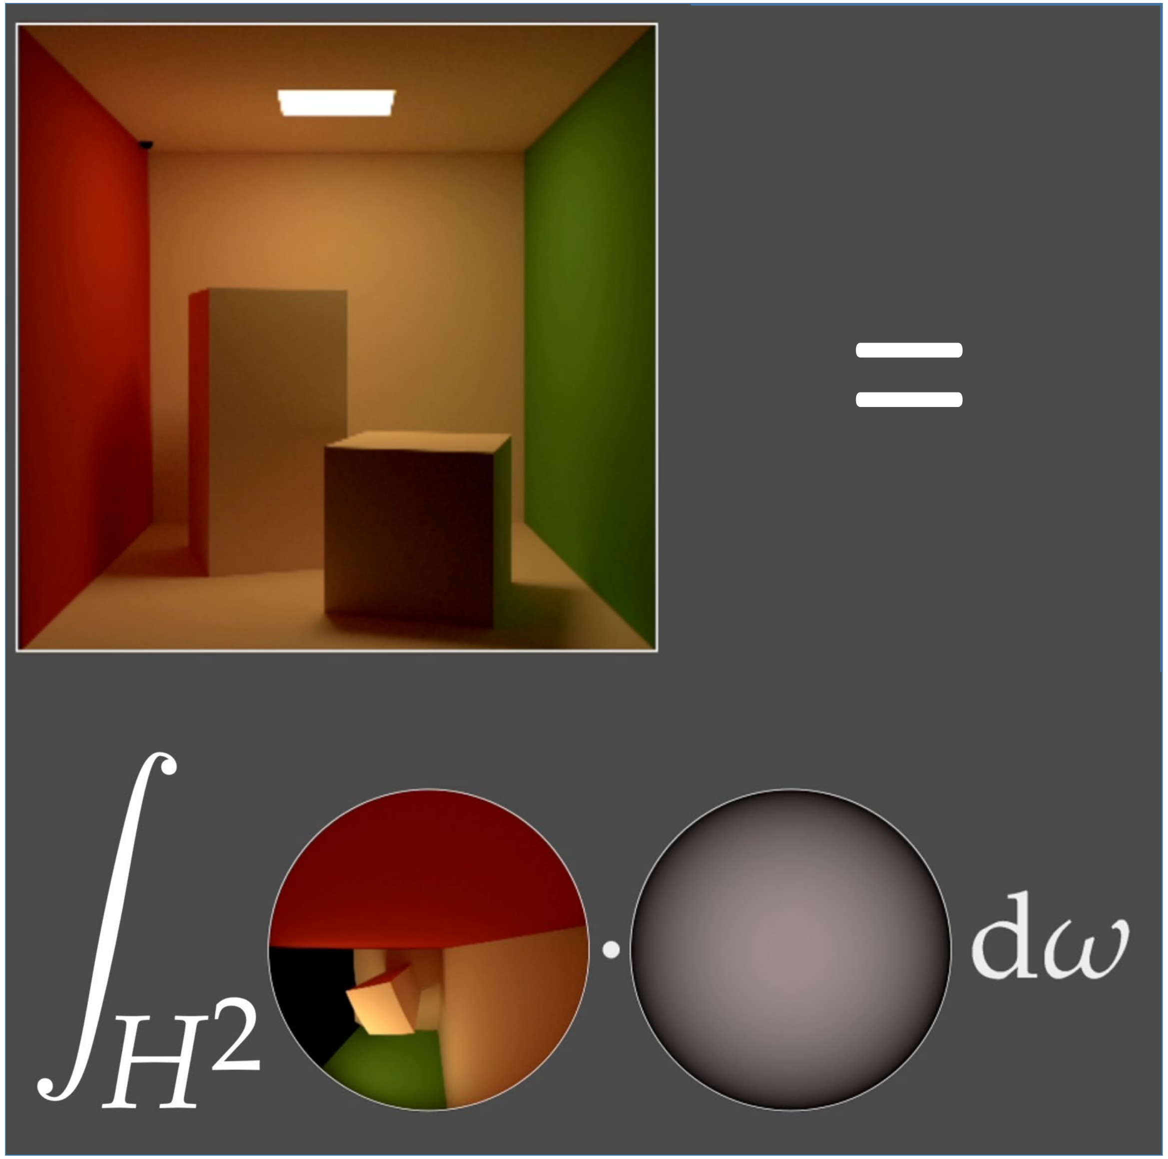

Topics to be covered

Monte Carlo rendering algorithms:

- unidirectional and bidirectional estimators
- Markov chain Monte Carlo techniques
- volumetric rendering
- importance sampling techniques
- quasi-Monte Carlo techniques

Topics to be covered

Advanced topics:

- differentiable and neural rendering
- neural rendering
- rendering wave-optics effects
- rendering specular transport effects
- rendering eikonal transport effects

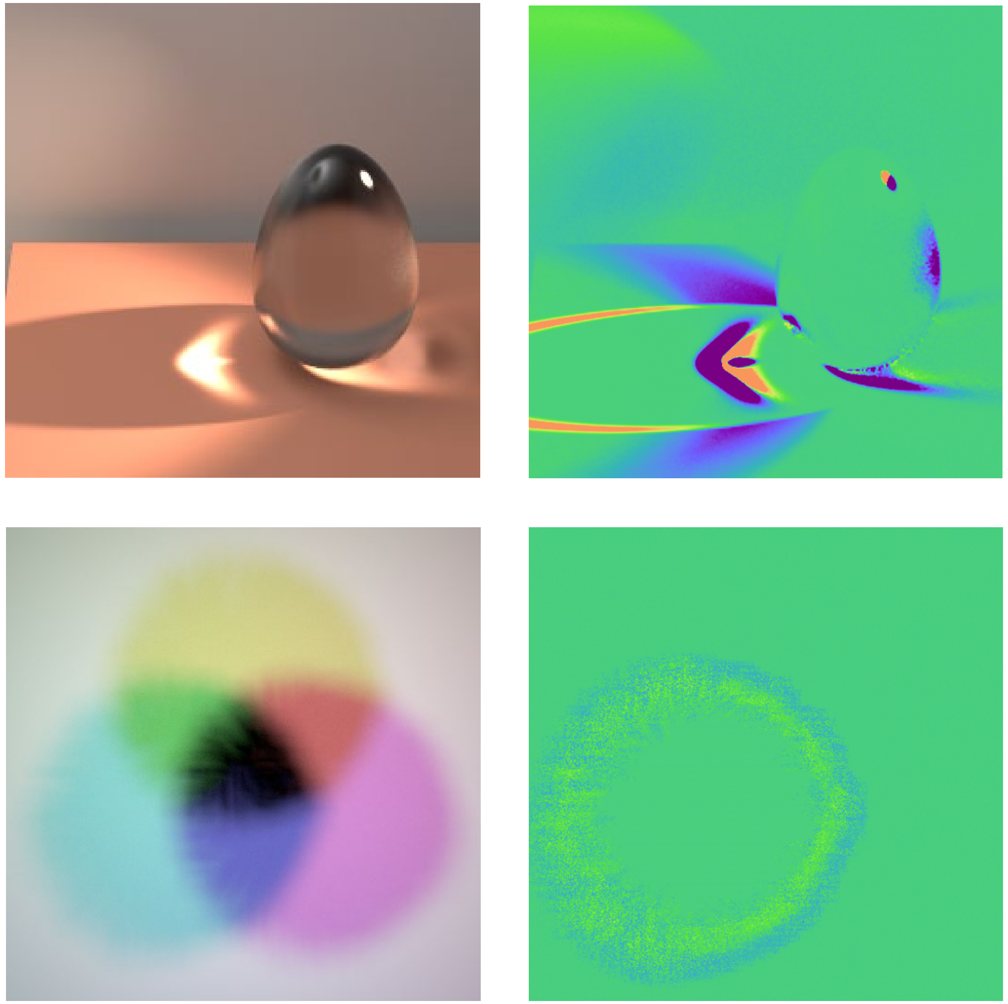

Books

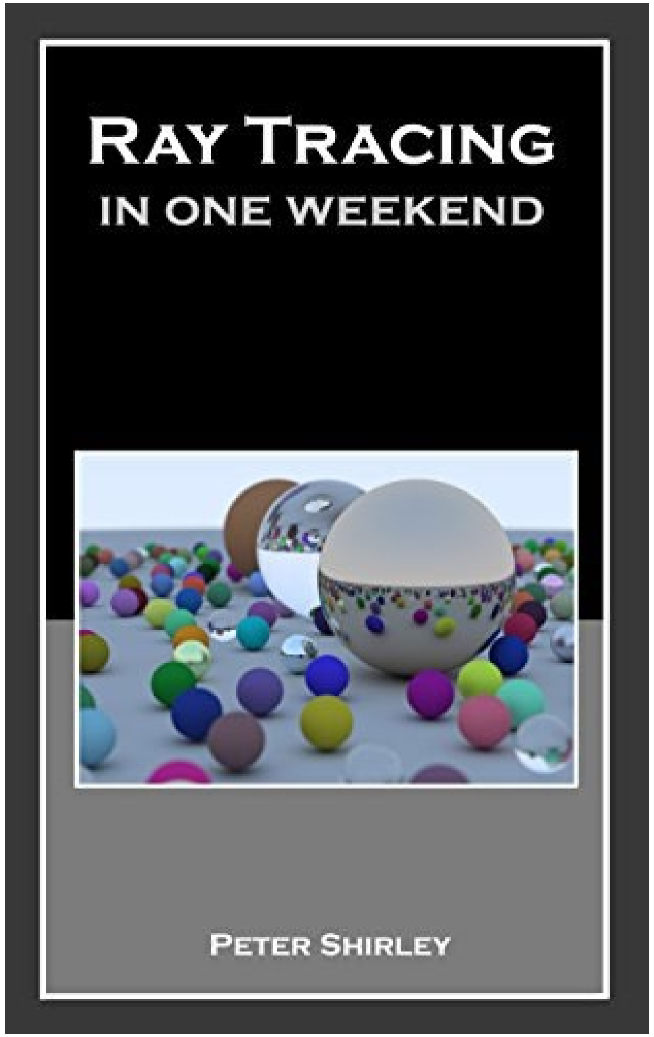

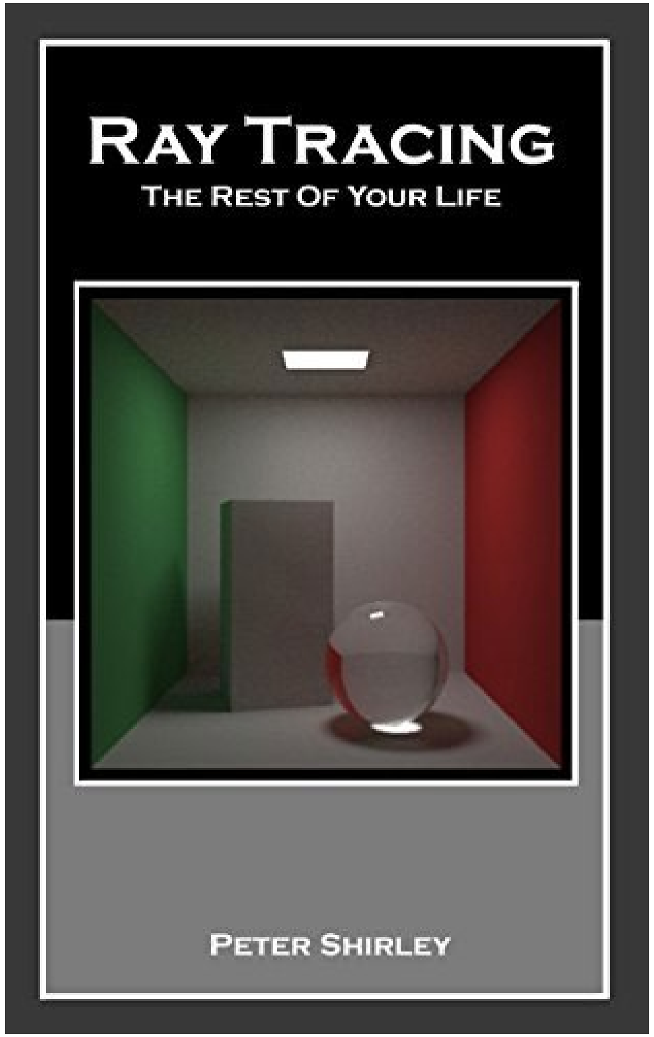

Peter Shirley's "Ray Tracing" series.

- Great reference material for first programming assignment.

Books

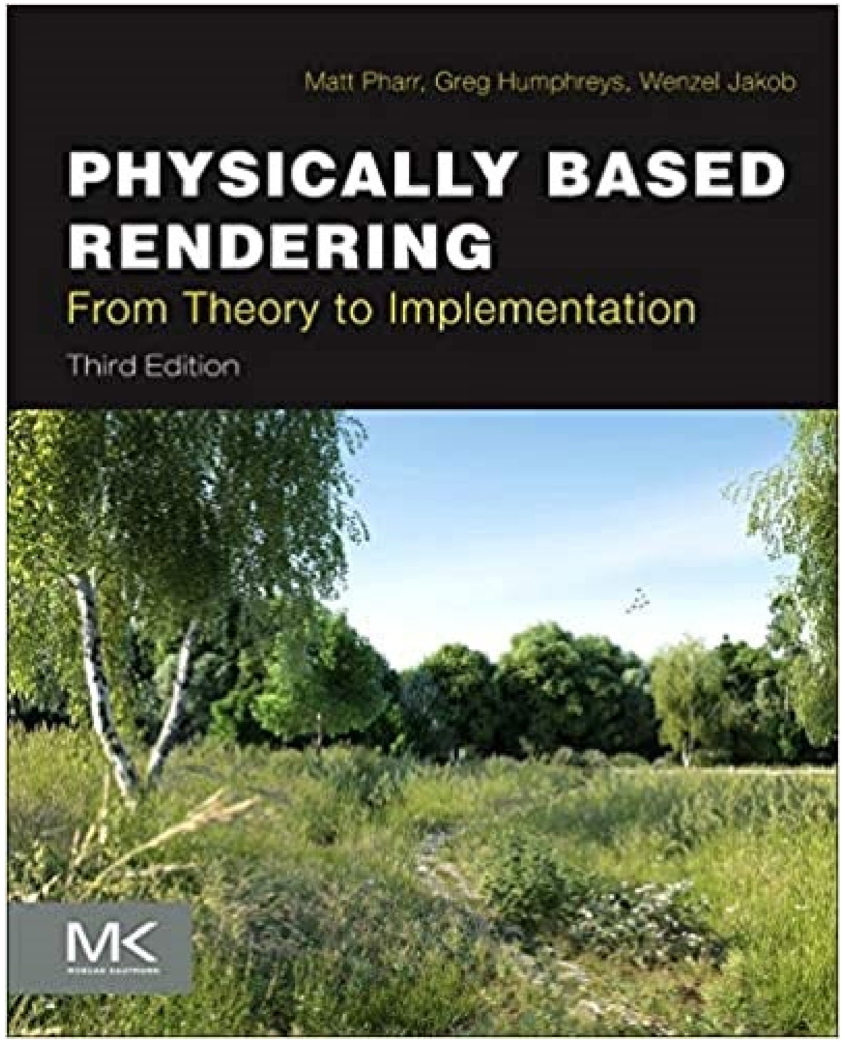

“PBR(T)”, great reference for later programming assignments.

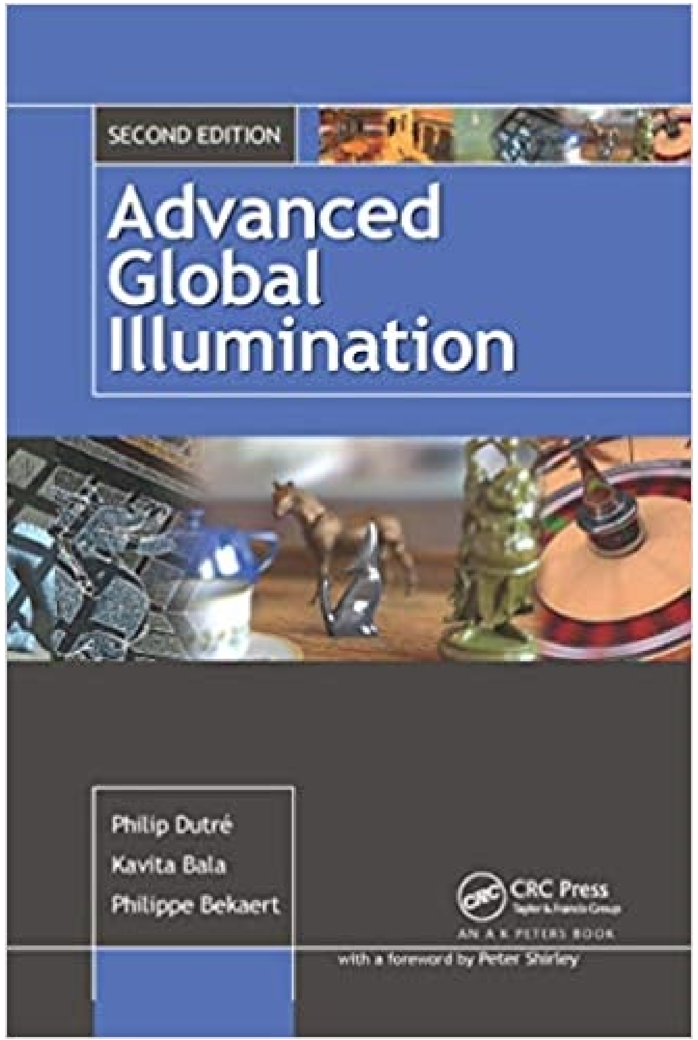

“AGI”, great reference for theoretical aspects of the course.

Prerequisites

At least one of the following:

- A computer vision course at the level of 16-385 or 16-720.
- A computer graphics course at the level of 15-462/15-662.
- A computational photography course at the level of 15-463/15-663/15-862.

Pop quiz

How many of you know or have heard of the following terms:

- Gaussian and box filtering.
- Convolution and Fourier transform.
- Aliasing and anti-aliasing.
- Homogeneous coordinates.
- Affine transforms and homographies.
- Pinhole, perspective, and orthographic camera.
- Triangular mesh.
- Ray-mesh intersections.
- Texture mapping.
- Radiometry and radiance.
- Lambertian, diffuse, and specular BRDFs.
- n-dot-l lighting.
- Environment map.
- Point and directional light sources.
- Ray tracing.
- Monte Carlo estimation.
- Refraction and diffraction.

Evaluation

- Four programming assignments (50%):
  - implement progressively more advanced features within an existing barebones rendering framework.
  - all programming will be in **C++**.
  - 0-th assignment will serve as a gentle introduction to our simplified version of Nori educational renderer.
  - five late days, no more than three per assignment, 10% penalty per additional late day.
- Ten take-home quizzes (20%):
  - solve 2-3 simple math problems related to each week's lectures.
  - **no late days, we will do solutions in recitations.**
  - you can skip two out of ten quizzes without penalty.
- Final project and rendering competition (25%):
  - implement rendering features of your choice and produce compelling imagery.
  - compete for **two free SIGGRAPH registrations (technical award and artistic award)!**
  - we will provide more information towards the end of February.
  - 15-668, 15-868 require more substantive project.
  - **no exam, but final project presentations are during the exam period.**
- Class, Piazza, and Slack participation (5%):
  - be around for lectures, office hours, reading groups (lenient this semester).
  - participate in Piazza and Slack discussions.
  - ask questions and answer other people's questions.

Submission deadlines will be enforced strictly!

Final project competition, Spring 2021

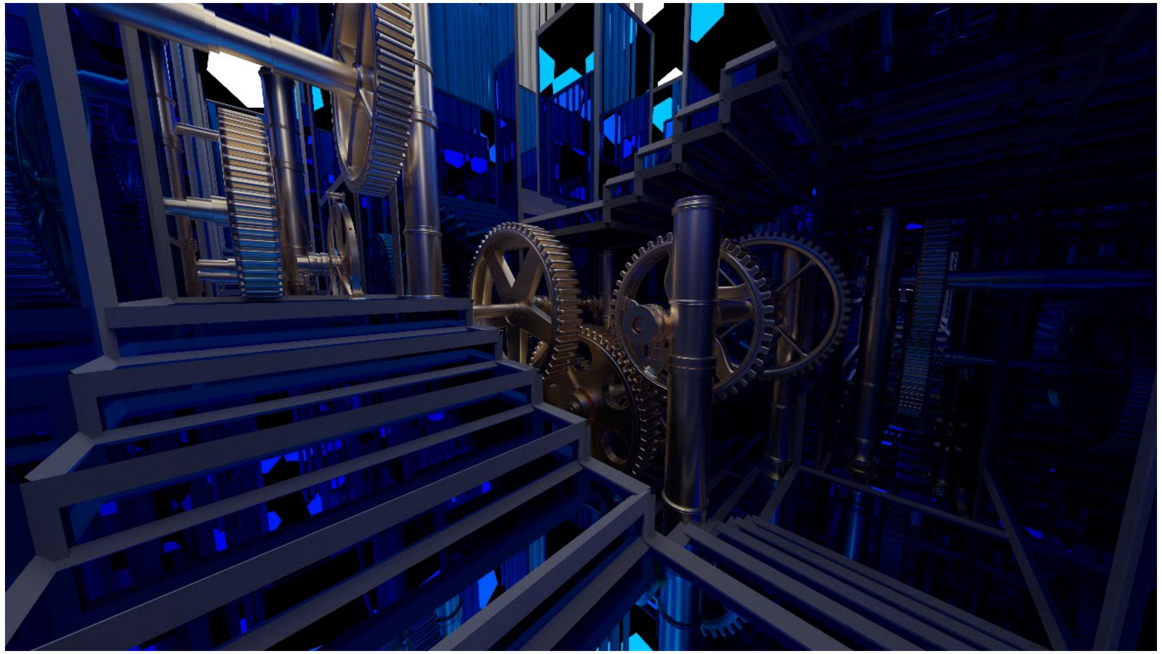

Technical award winner: Max Slater

Art award winner: Arpit Agarwal

- All of last year's final projects:
  - presentations - https://docs.google.com/presentation/d/1qeFYNXn3Z_pbmvTCtEUOtU8JGy1v8zeaQ9MlUJgCP-8/edit
  - renderings - http://graphics.cs.cmu.edu/courses/15-468/2021_spring/rendering_competition.html

Rendering competitions elsewhere

Look at rendering competitions for similar courses at other universities for inspiration!

- Dartmouth (2019, 2017, 2016)
- EPFL (2019, 2018, 2017)
- ETH Zurich (2017, 2016, Fall 2015, Spring 2015, 2014, 2013, 2012)
- UC San Diego (2011, 2010, 2008, 2007, 2006, 2005, 2004, 2003)
- Stanford.

Wednesday recitations

- Every Wednesday, there will be a recitation, where we go over the solutions to that week's take-home quiz.
- Typically, recitations take the form of whiteboard derivations, and free-form discussion.
- **Participation is completely optional.** Recitations will be recorded.
- Time will be decided by vote in the start-of-semester survey.

Friday reading groups

- Every second Friday, there will be a reading group to cover in detail an advanced topic selected by students and instructors.
- Typically, reading groups take the form of a review of a group of papers, whiteboard derivations, and free-form discussion.
- **Participation is completely optional.** Reading groups will be recorded.
- Time will be decided by vote in the start-of-semester survey.
- Topics covered last year: specular next-event estimation, Metropolis light transport, optimal multiple importance sampling, ReSTIR, BRDF acquisition, path guiding, transmittance, neural field representations.

Contact information, office hours, and discussion

- Feel free to email us about administrative questions.
  - please use [15468] in email title!
- Technical questions should be asked on Piazza or Slack.
  - we won't answer technical questions through email.
  - you can post anonymously if you prefer.
- Office hours will be determined by vote in the start-of-semester survey.
  - office hours will be in person at the Smith Hall (EDSH) graphics lounge (by through Zoom until January 31st).
  - feel free to email Yannis about additional office hours.
  - you can also just drop by Yannis' office (Smith Hall (EDSH) Rm 225) after January 31st.
  - you can also post on Piazza or DM on Slack for additional office hours.
  - office hours for this week will be announced on Piazza.
- Post-lecture Q&A for 30 minutes.

This course is still highly experimental!

- Second time this course is offered, so expect things to inevitably change throughout the semester.

Interested in research?

- Visit the graphics lab and imaging group websites:

  http://graphics.cs.cmu.edu/
  https://imaging.cs.cmu.edu/

- Email Yannis if you want to be added to the graphics lab mailing list and attend our weekly meetings (**Tuesdays, 1 – 2 pm ET**).

- We are actively recruiting research assistants for projects relating to **rendering**, imaging, and graphics in general. Please email Yannis if interested.

Please take the start-of-semester survey!

- Posted on Piazza as well:

https://docs.google.com/forms/d/e/1FAIpQLSdZONUsOzpfkLGTcpuc8nlHRx7EkVIL_yGzXZt5xdwTsDlzxw/viewform