Introduction

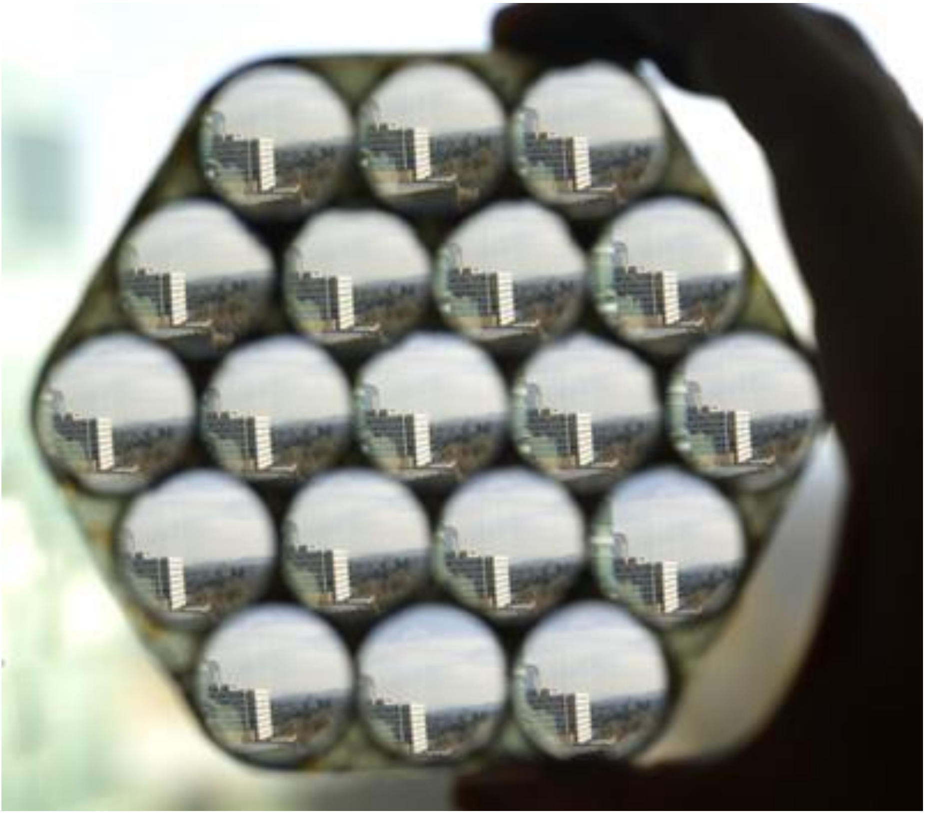

15-463, 15-663, 15-862
Computational Photography
Fall 2019, Lecture 1

http://graphics.cs.cmu.edu/courses/15-463

Overview of today’s lecture

- Teaching staff introductions
- What is computational photography?
- 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/

Looking around corners

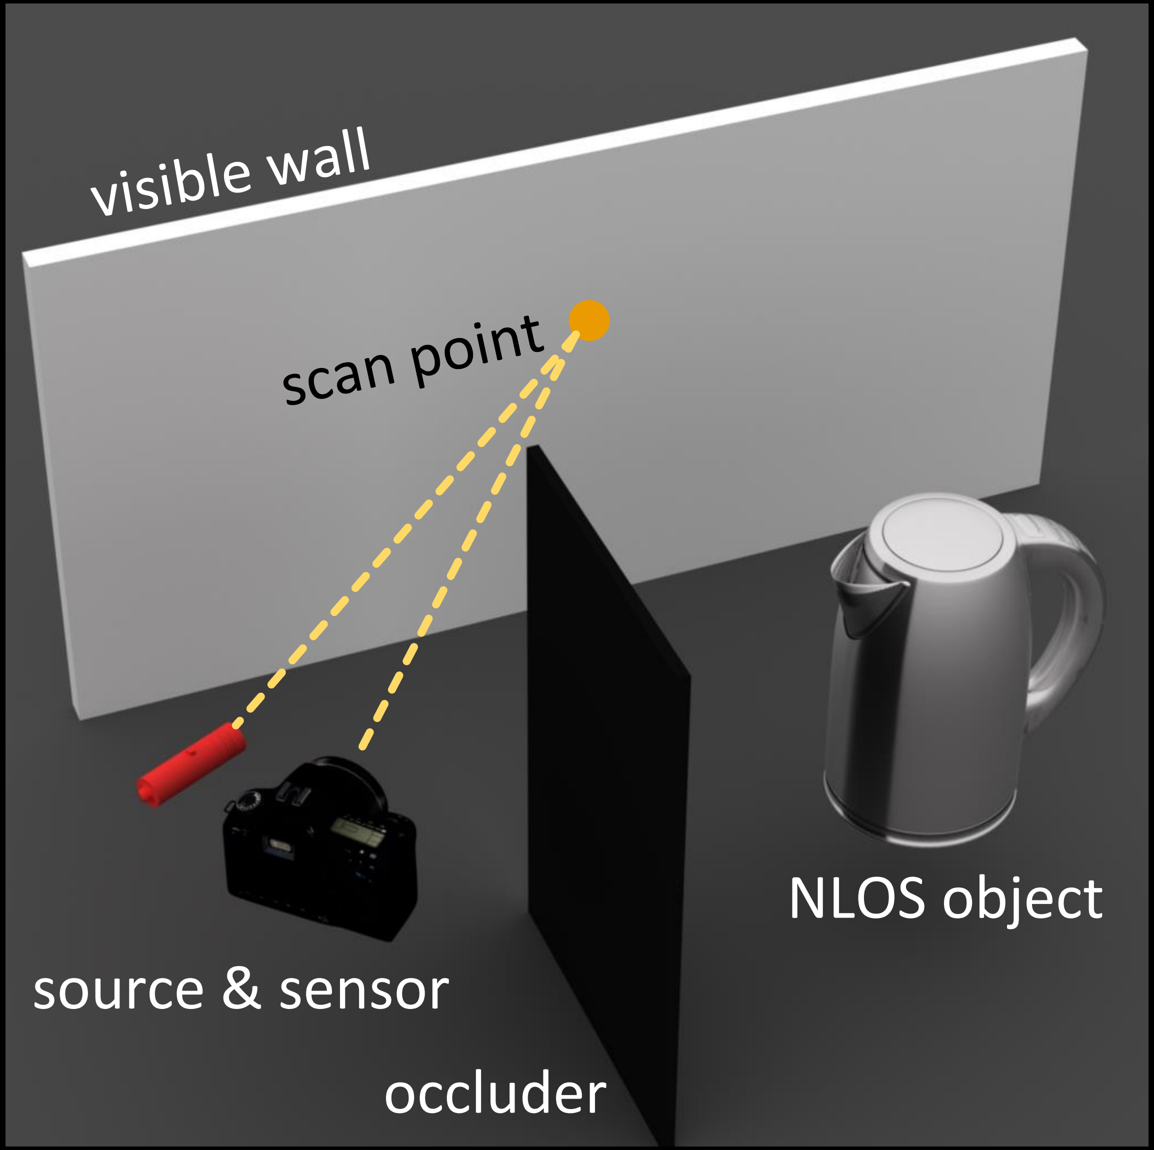

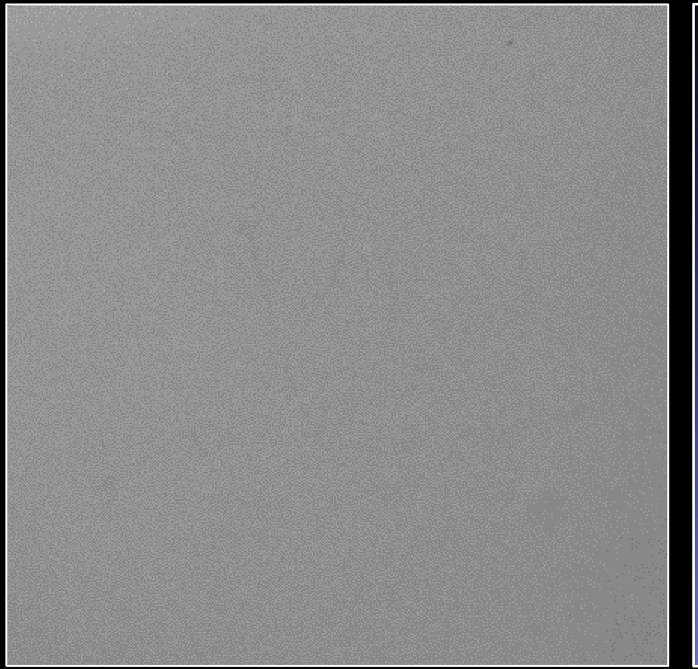

what a regular camera sees

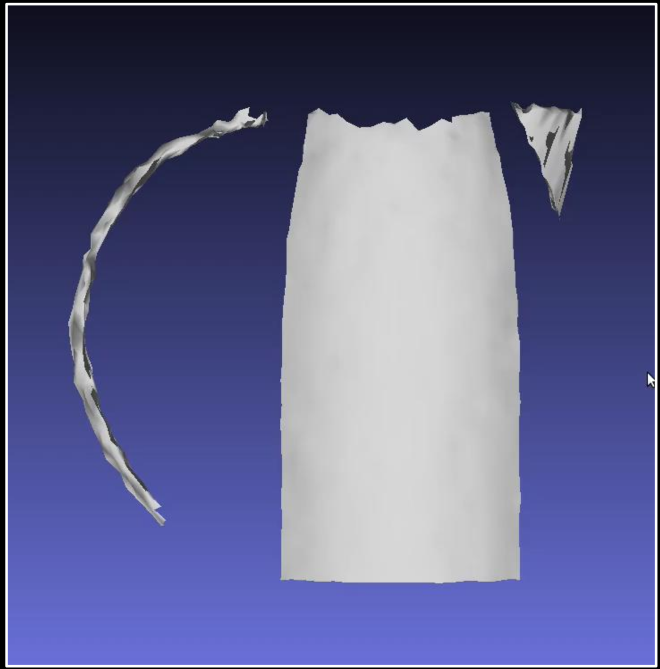

what we can reconstruct

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

Looking inside deep scattering objects

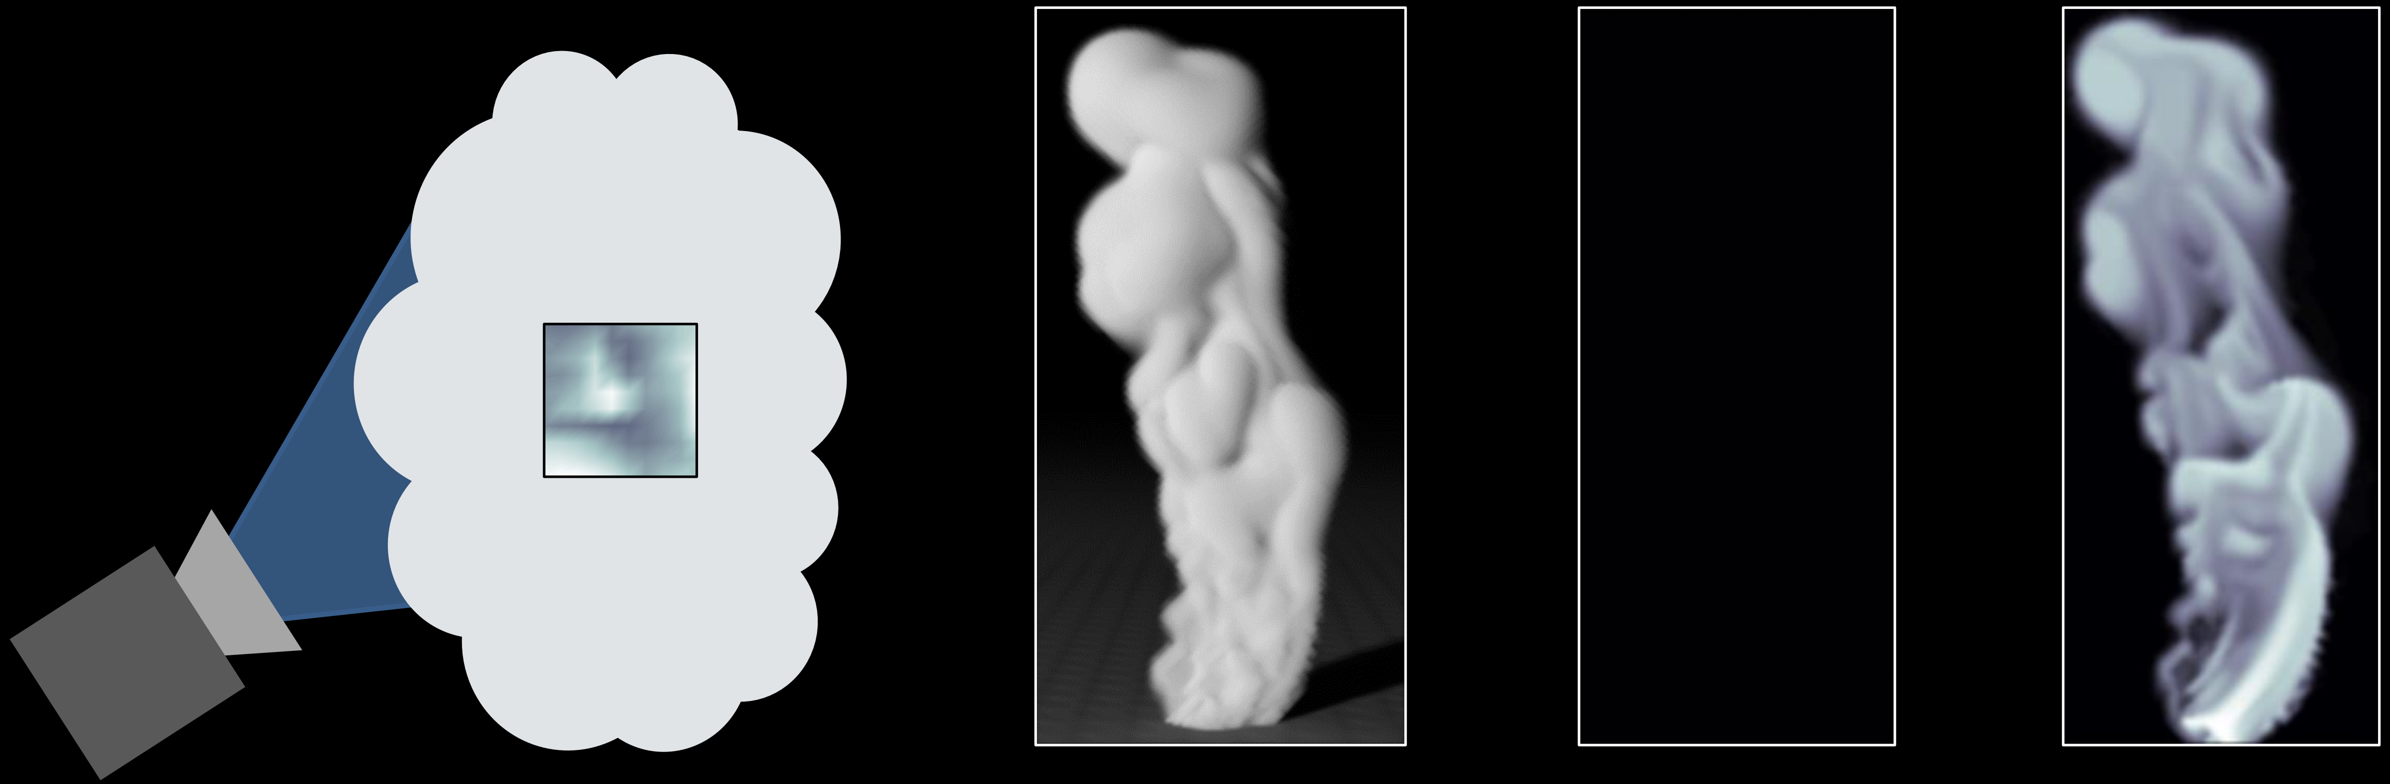

camera

thick smoke cloud

simulated camera measurements

reconstructed cloud volume

slice through the cloud

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

Seeing light in flight

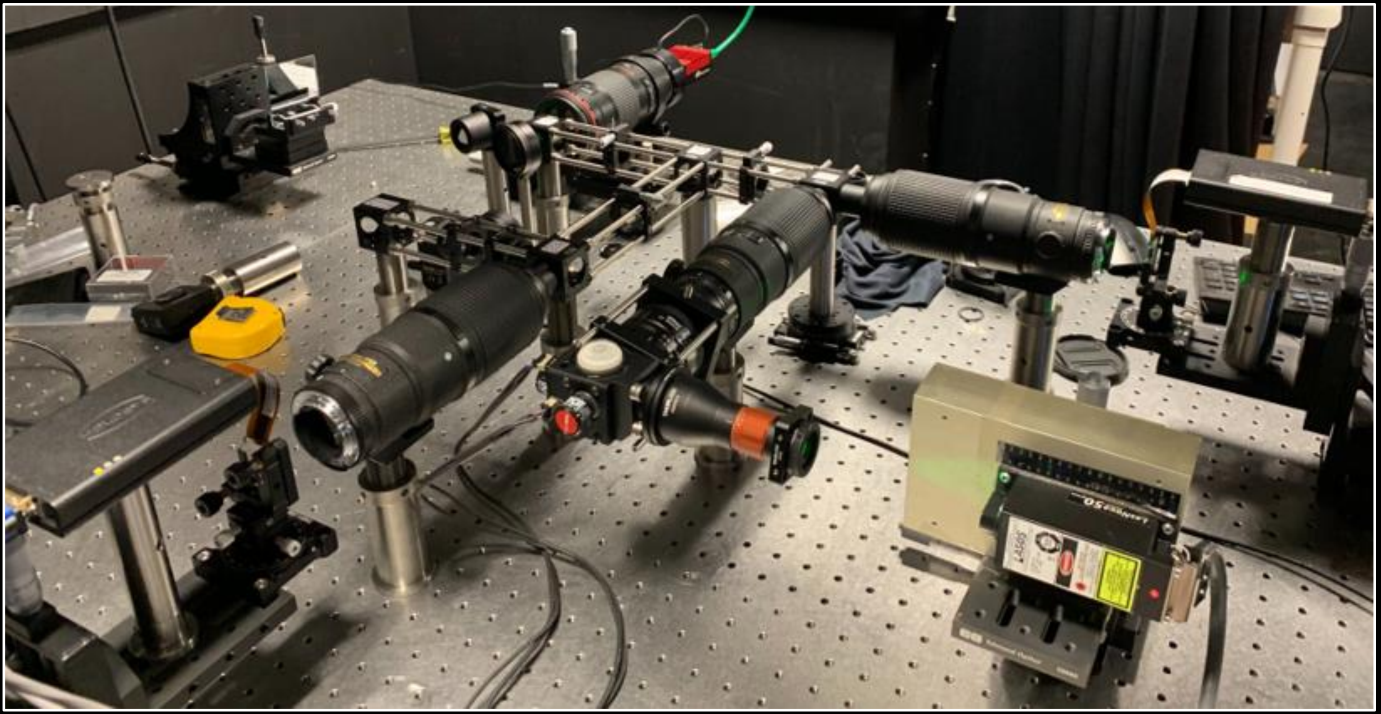

camera for capturing video at 10^{15} frames per second

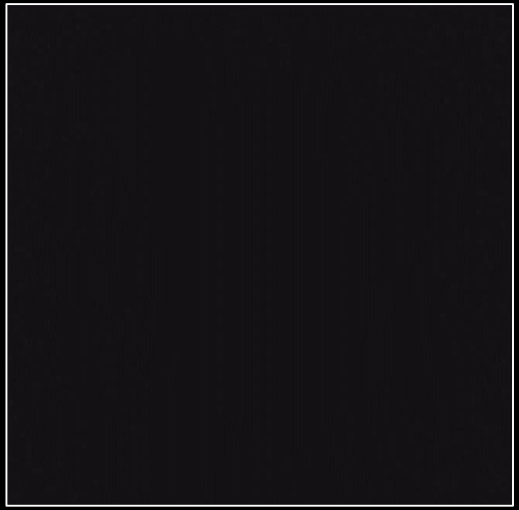

http://imaging.cs.cmu.edu/

Rendering wave effects

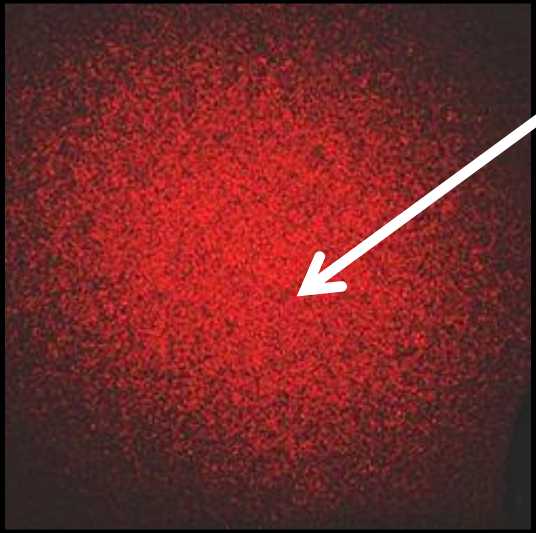

speckle: noise-like pattern

what real laser images look like

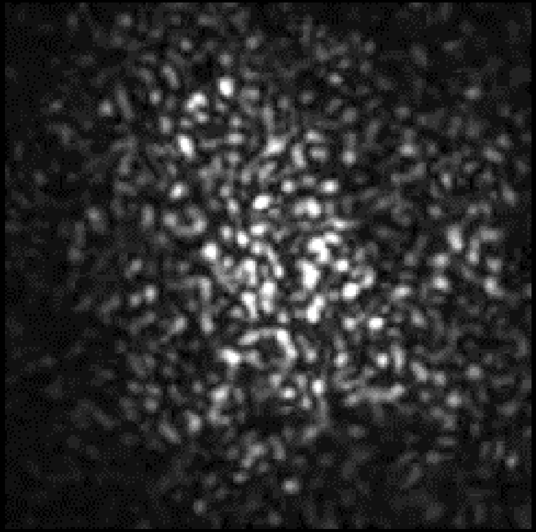

what real laser videos look like

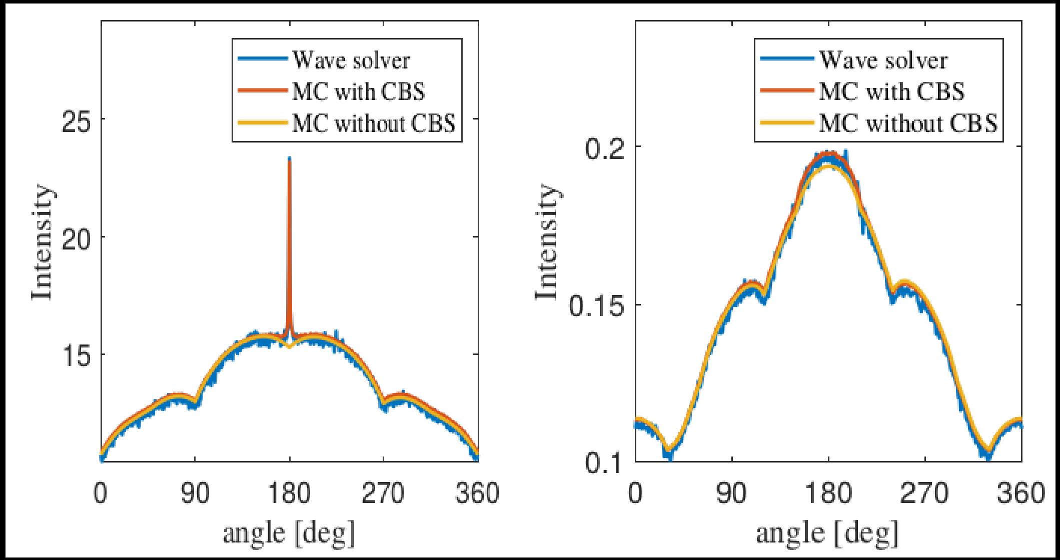

match wave equation solvers, **10^5x faster**

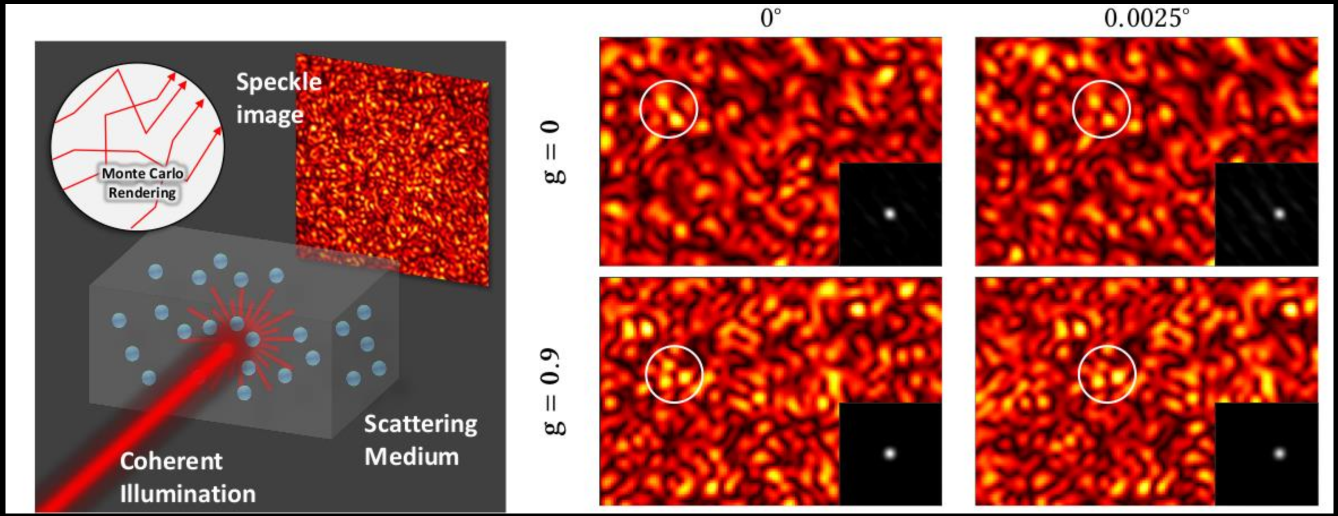

reproduce physical effects like memory effect

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

Differentiable rendering

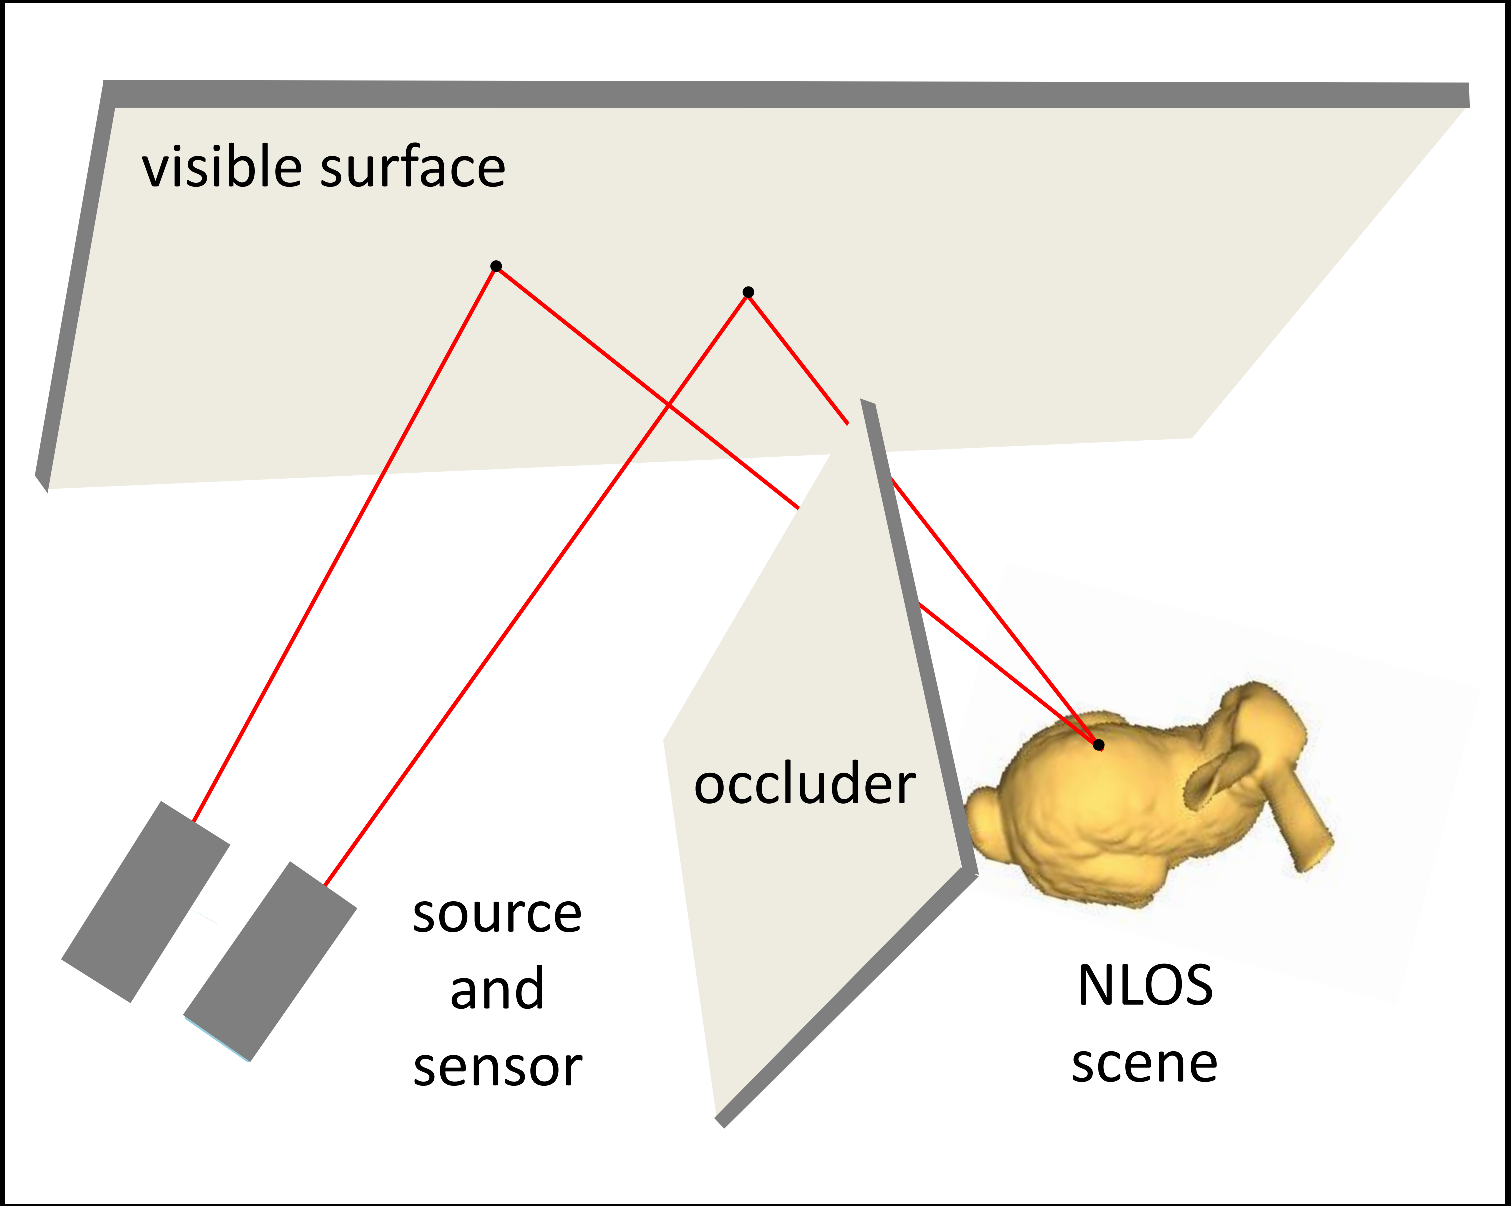

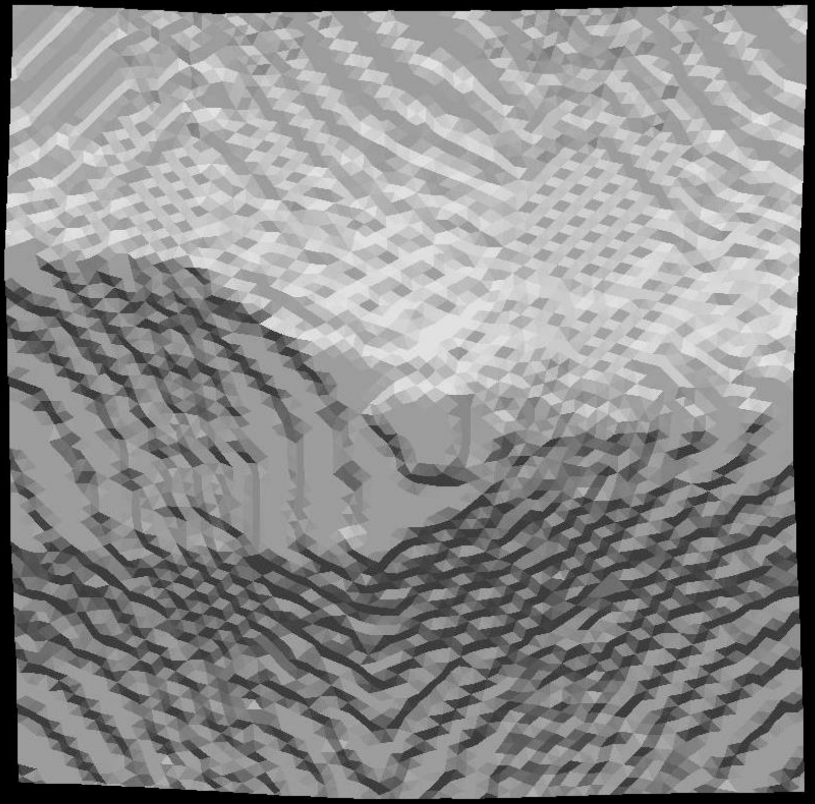

reconstruction evolution

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

TA: Tiancheng Zhi

- CSD PhD Student
- Advisors: Srinivasa Narasimhan and Martial Hebert
- Research Interests: Multispectral Imaging and Material Recognition

- Education:
  - Undergraduate student, EECS, Peking University, 2012-2016
  - PhD student, SCS, Carnegie Mellon University, 2016-

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

Material-aware Cross-spectral Stereo Matching

Difficult materials in Stereo Matching

(a) Left RGB (b) Right NIR

(c) Difficult regions for matching (d) Predicted disparity

Fixing Unreliable Regions with Material Awareness

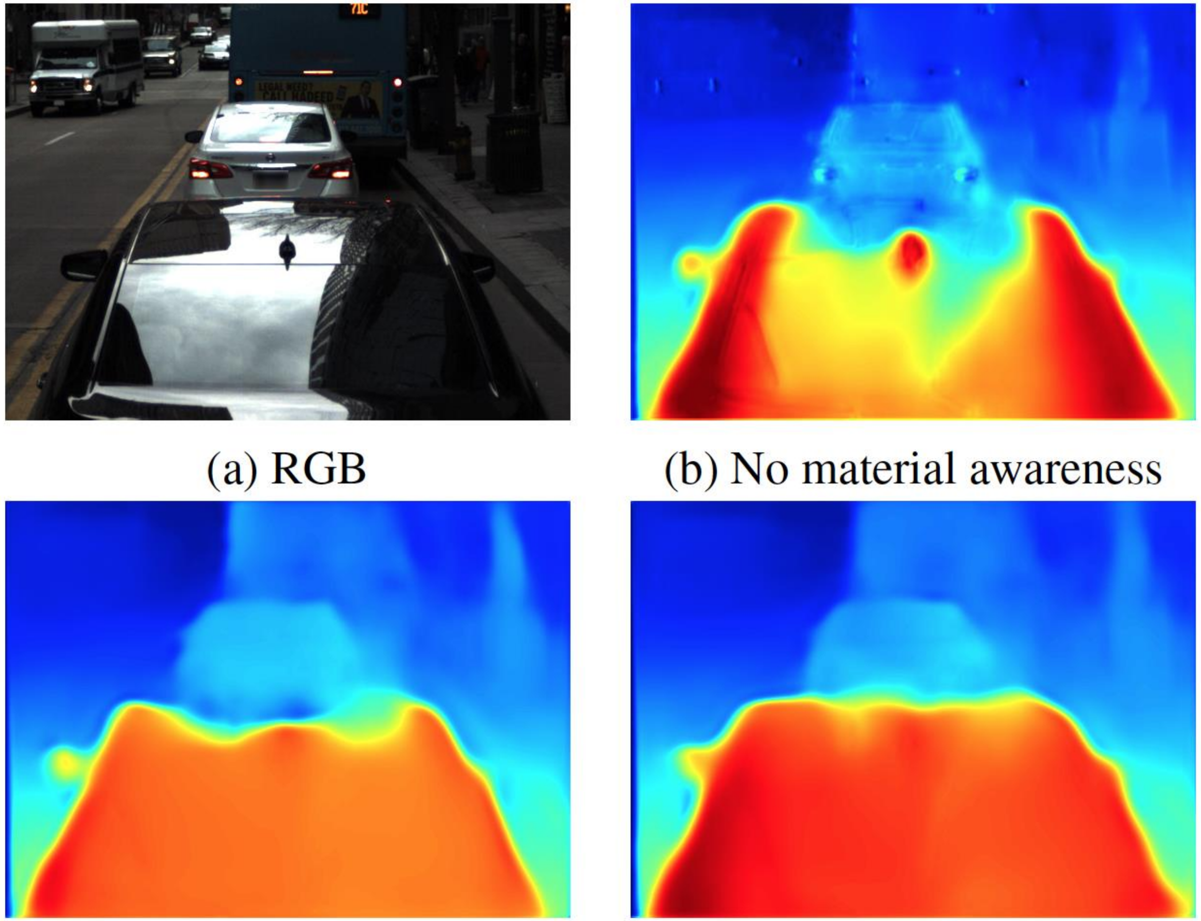

(a) RGB (b) No material awareness

(c) Smoothing w/o confidence (d) Smoothing w/ confidence

Multispectral Imaging for Powder Recognition

20 Different Powders in Different Spectra

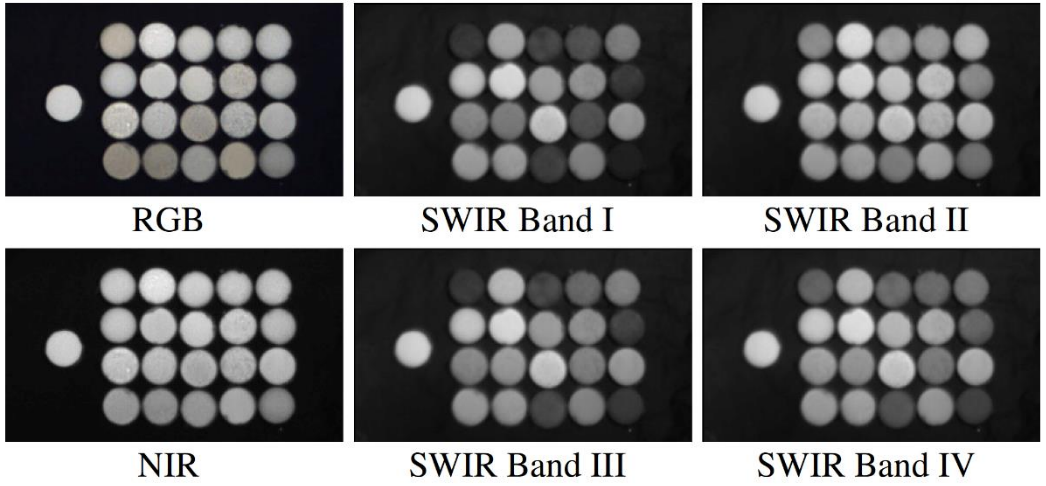

Imaging System

Detection of Powders on Human

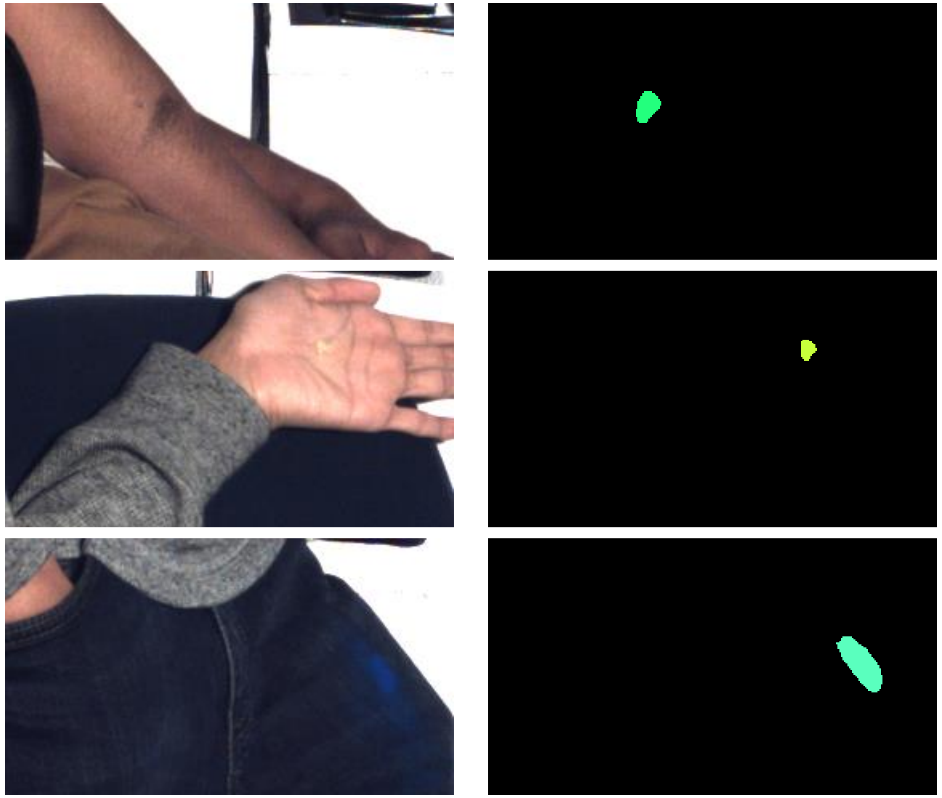

What is computational photography?

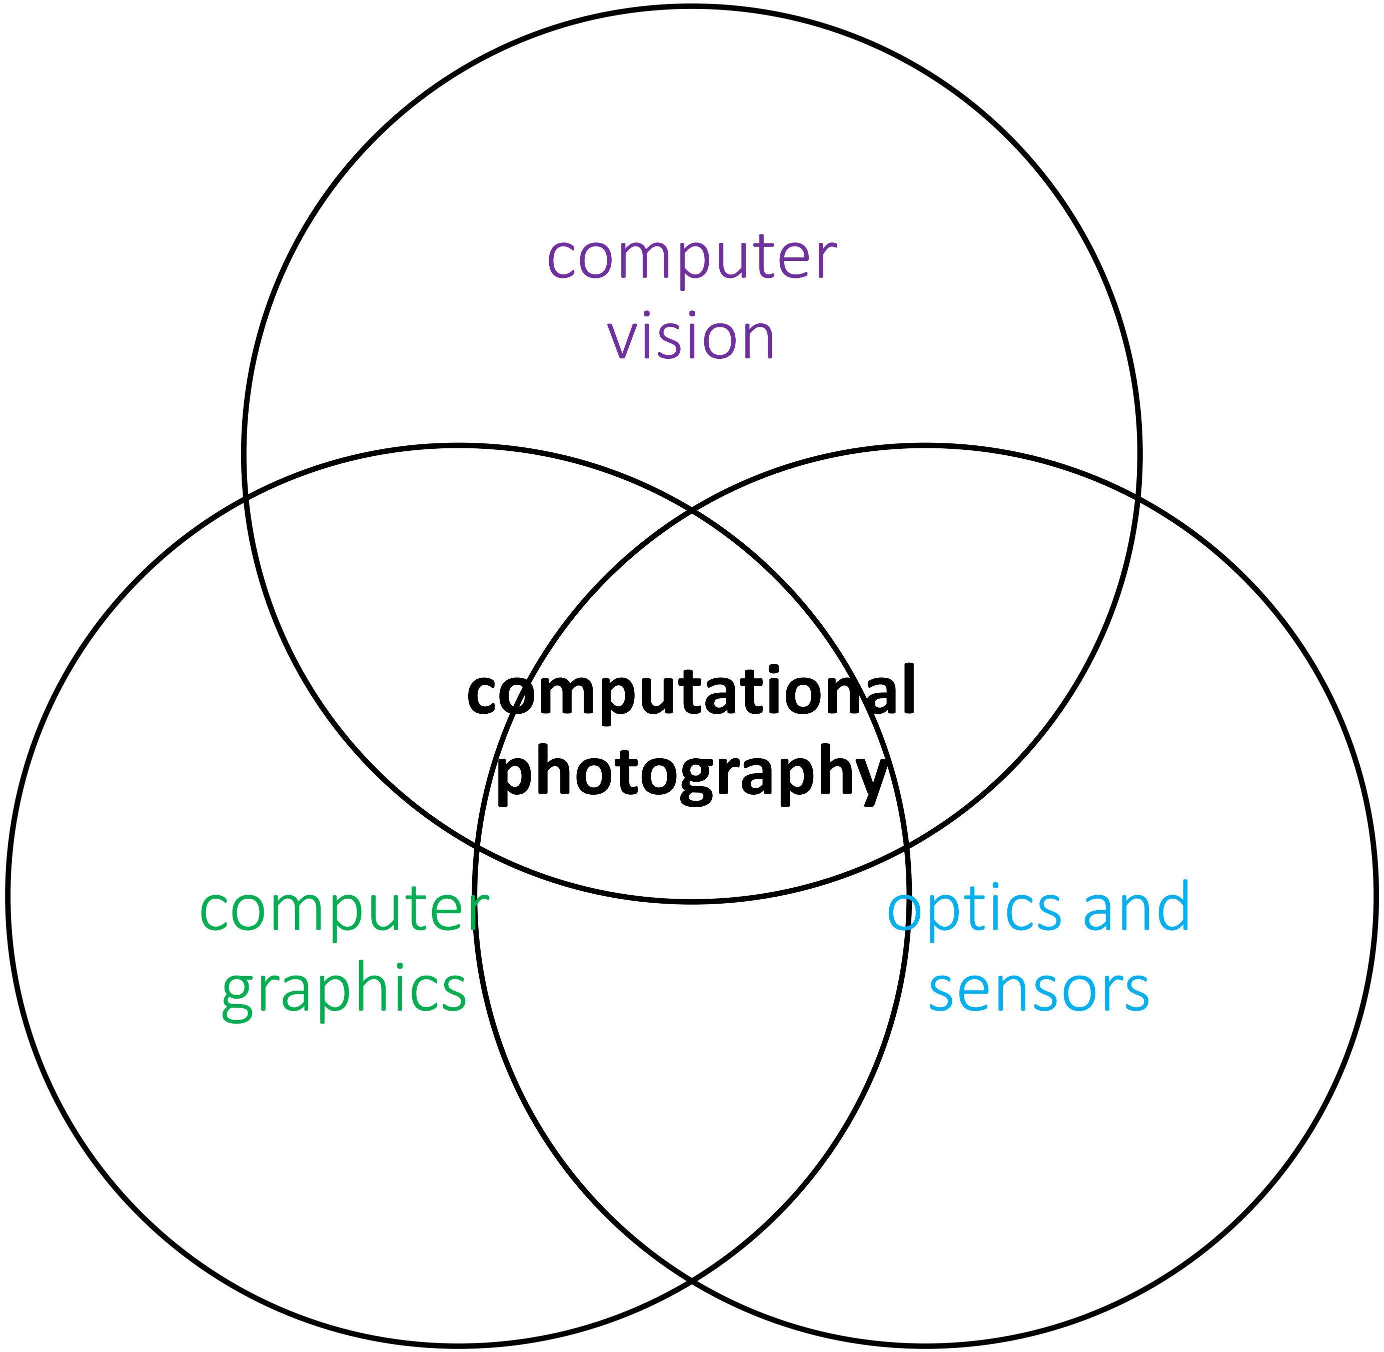

[Slide credit: Kris Kitani]

Analog photography

optics to focus light on an image plane

film to capture focused light (chemical process)

dark room for limited post-processing (chemical process)

Digital photography

optics to focus light on an image plane

digital sensor to capture focused light (electrical process)

on-board processor for post-processing (digital process)

Computational photography

optics to focus light on
an image plane

digital sensor to capture focused
light (electrical process)

arbitrary computation
between sensor and image

Overcome limitations of digital photography

Image enhancement and photographic look

camera output

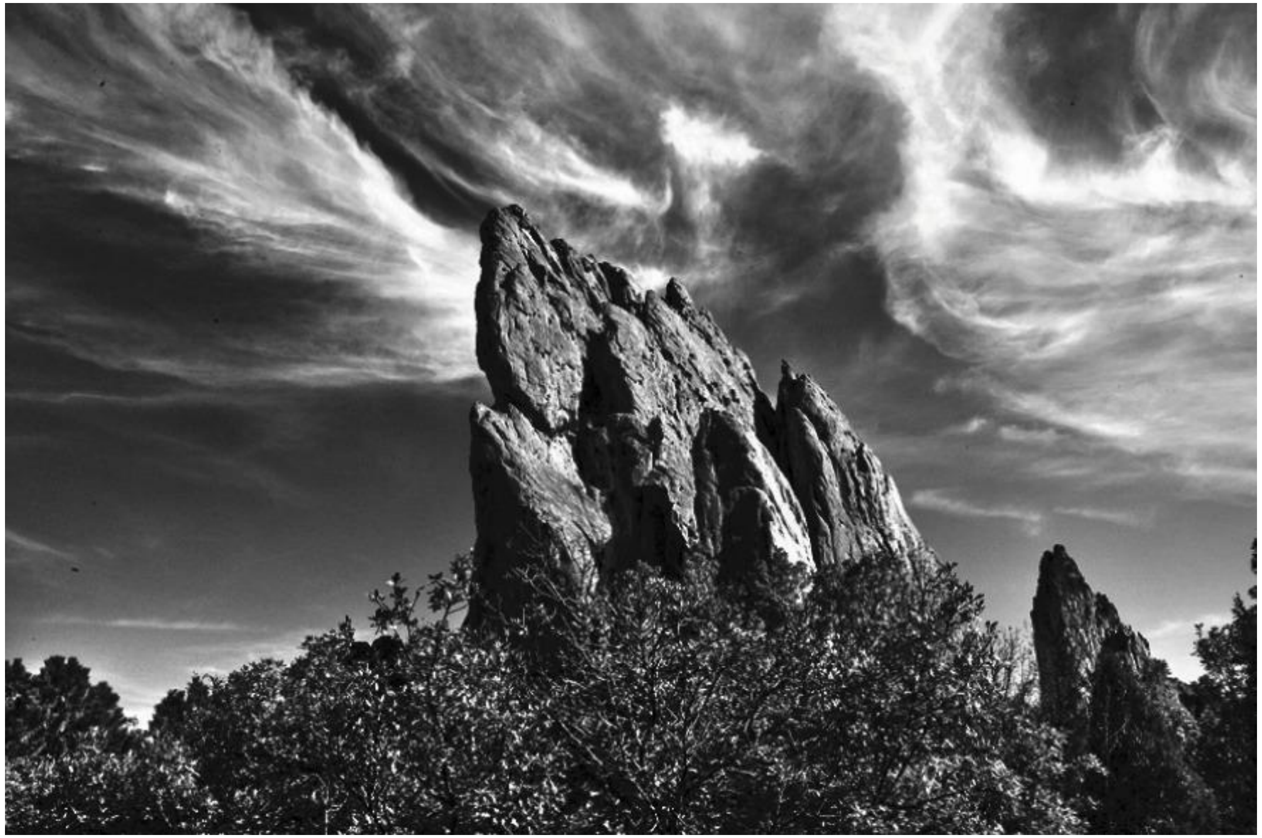

image after stylistic tonemapping

[Bae et al., SIGGRAPH 2006]

Overcome limitations of digital photography

High dynamic range (HDR) imaging

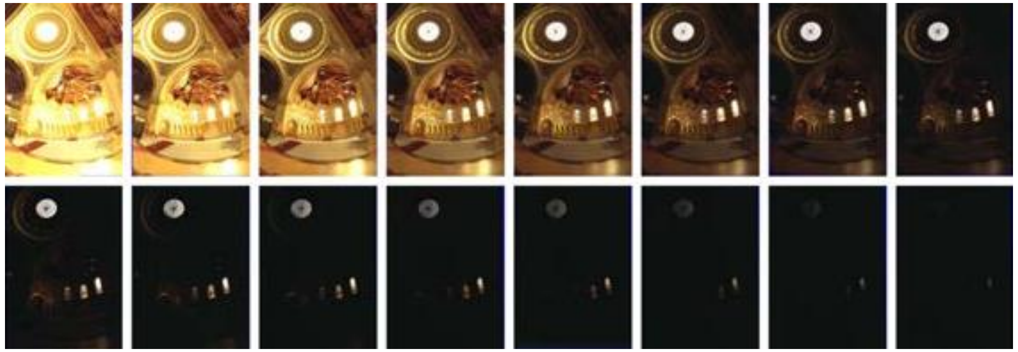

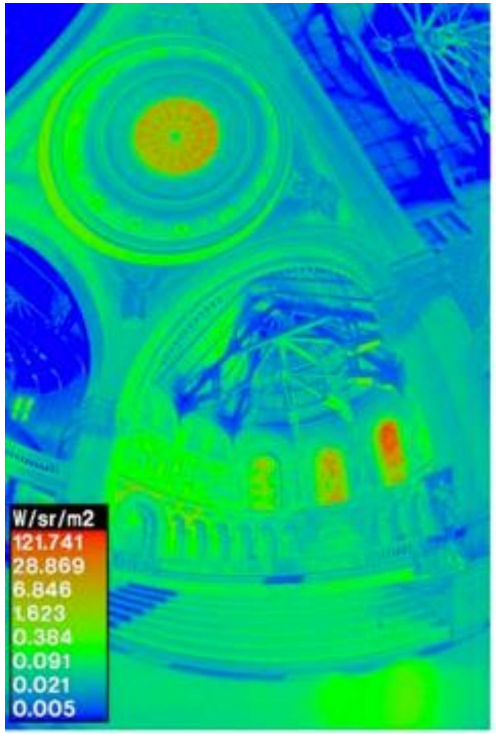

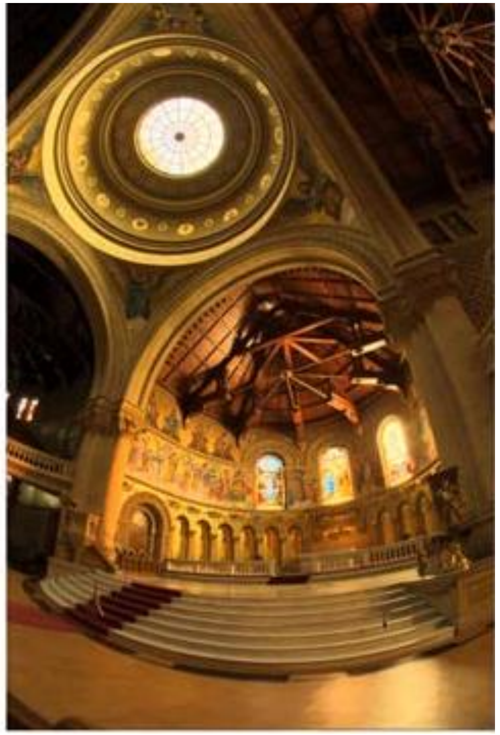

One of your homeworks!

[example from www.dpreview.com] [Debevec and Malik, SIGGRAPH 1997]

Create realistic new imagery

Image blending and harmonization

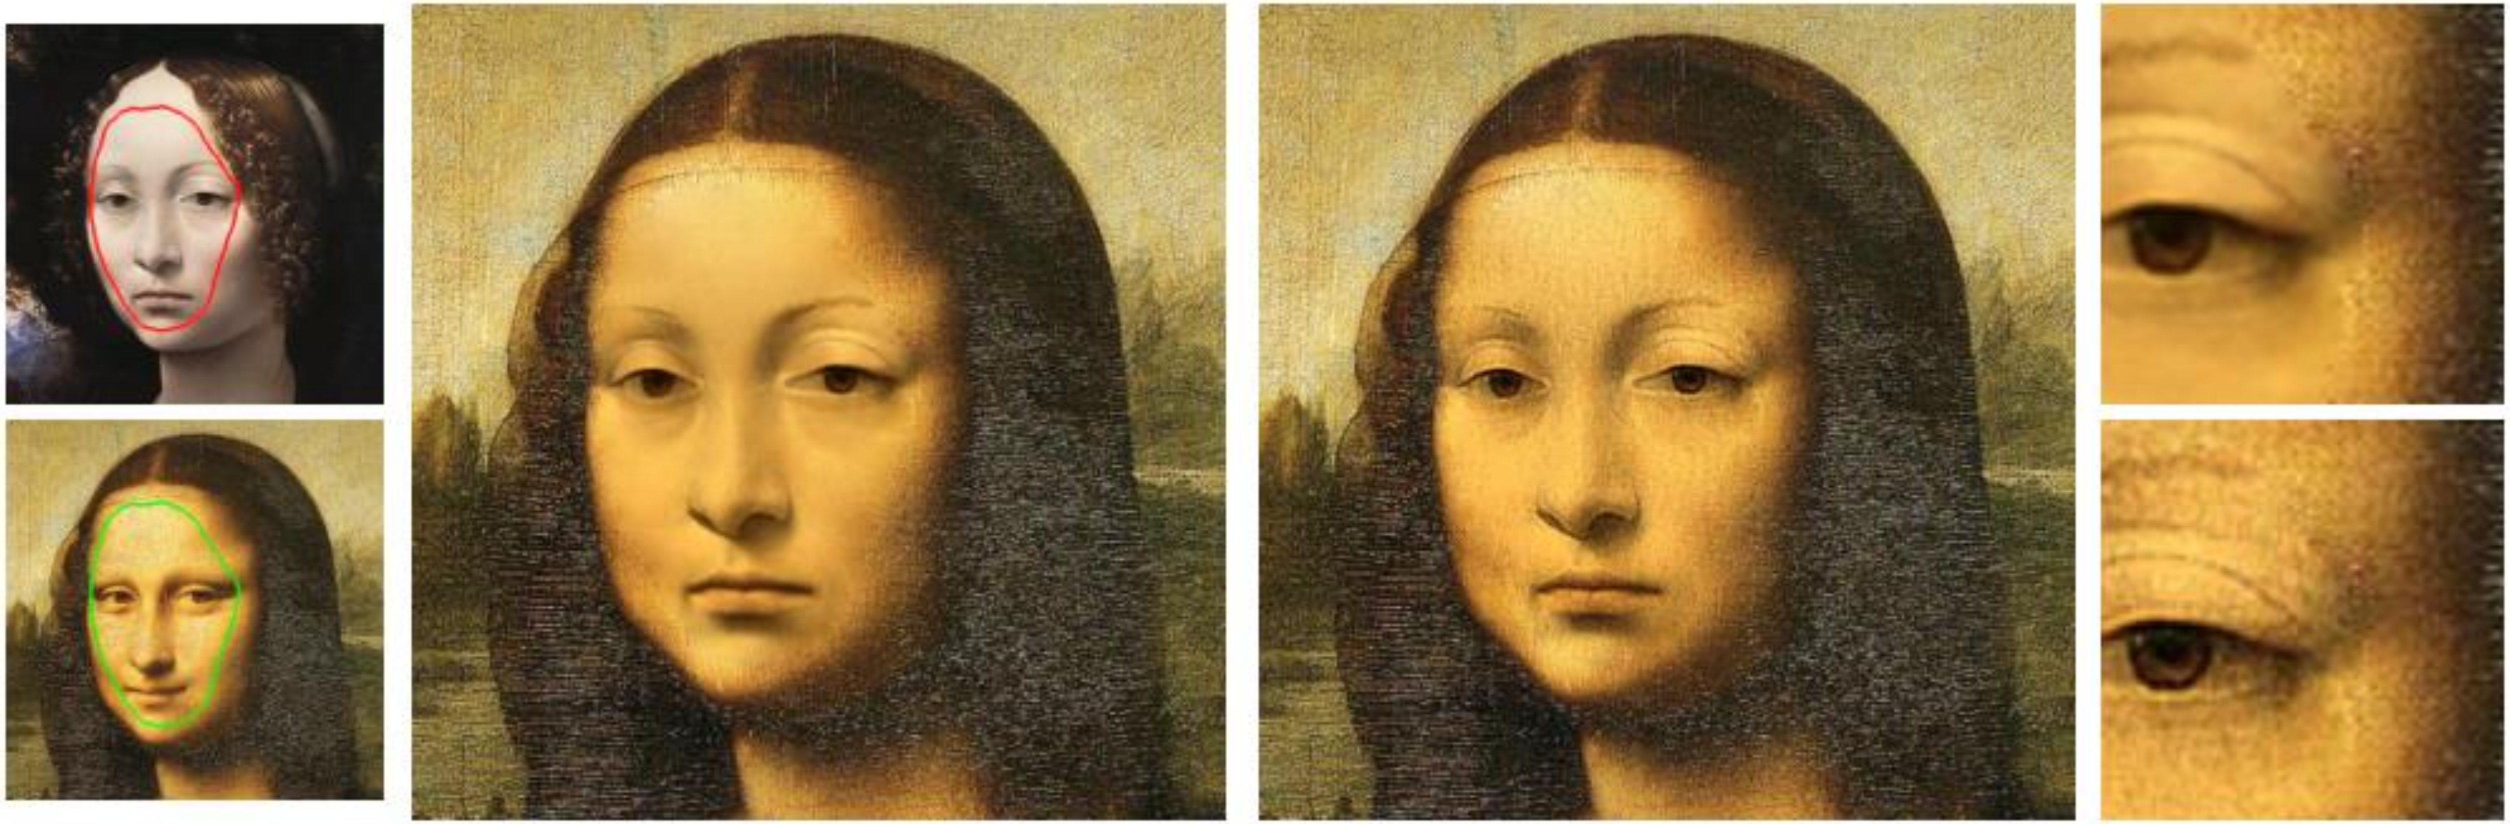

One of your homeworks!

[Sunkavalli et al., SIGGRAPH 2010]

Post-capture image compositing

Computational zoom

images captured at three zoom settings

post-capture synthesis of new zoom views

[Badki et al., SIGGRAPH 2017]

Process image collections

Auto-stitching images into panoramas

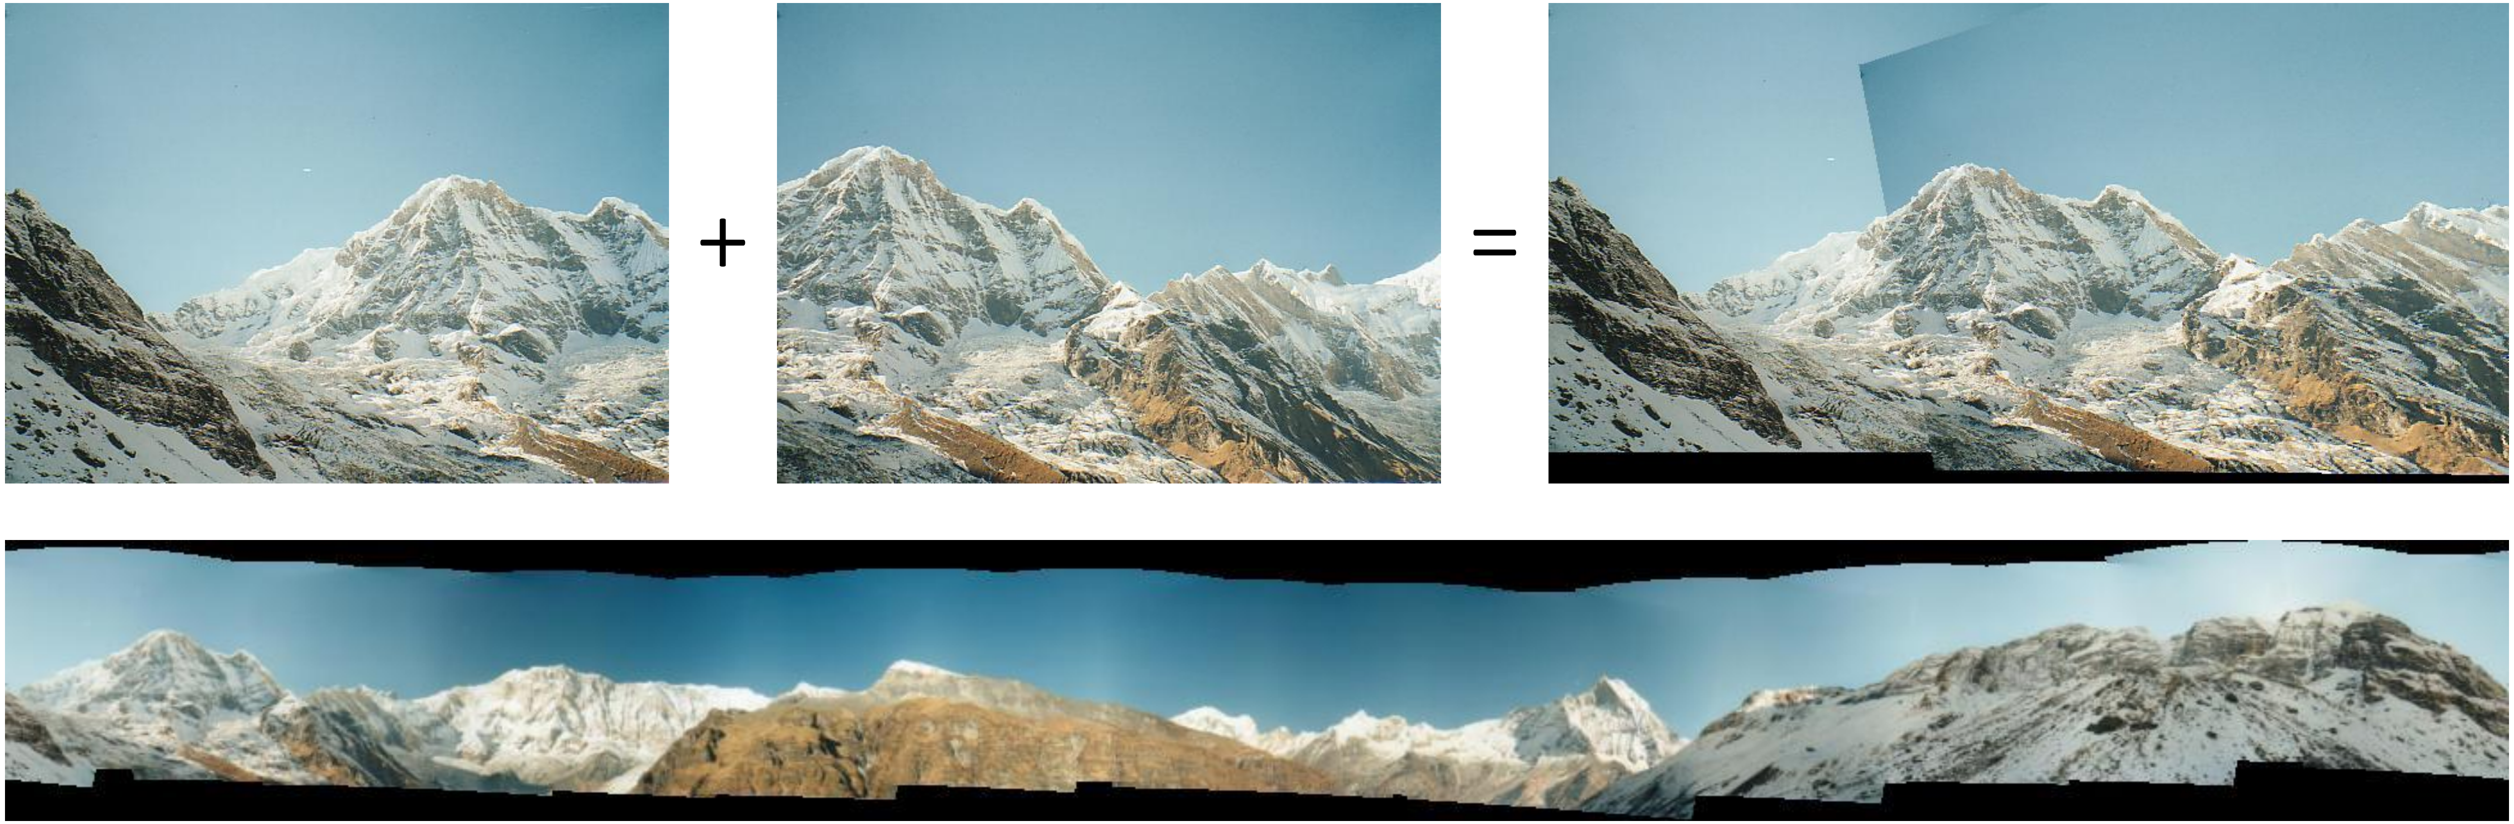

[Brown and Lowe, IJCV 2007]

Process (very) large image collections

Using the Internet as your camera

reconstructing cities from Internet photos

time-lapse from Internet photos

[Agarwal et al., ICCV 2009] [Martin-Brualla et al., SIGGRAPH 2015]

Computational photography

optics to focus light on
an image plane

digital sensor to capture focused
light (electrical process)

arbitrary computation
between sensor and image

Computational photography

generalized optics between scene and sensor

digital sensor to capture focused light (electrical process)

arbitrary computation between sensor and image

*Sometimes people discriminate between *computational photography* and *computational imaging*. We use them interchangeably.

Capture more than 2D images

Lightfield cameras for plenoptic imaging

post-capture refocusing

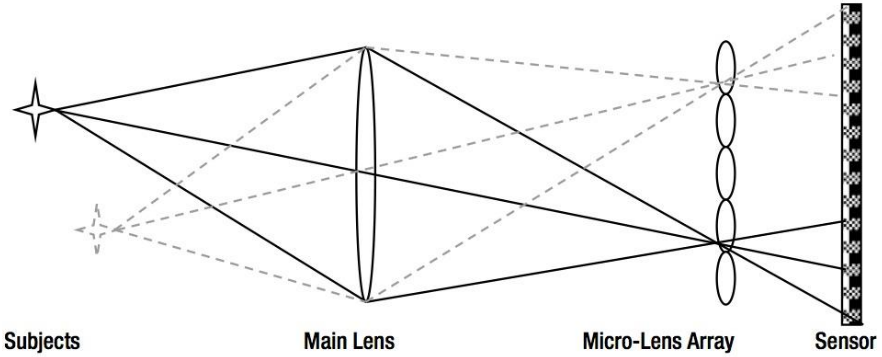

One of your homeworks!

[Ng et al., SIGGRAPH 2005] [Lytro Inc.]

Capture more than 2D images

Lightfield cameras for plenoptic imaging

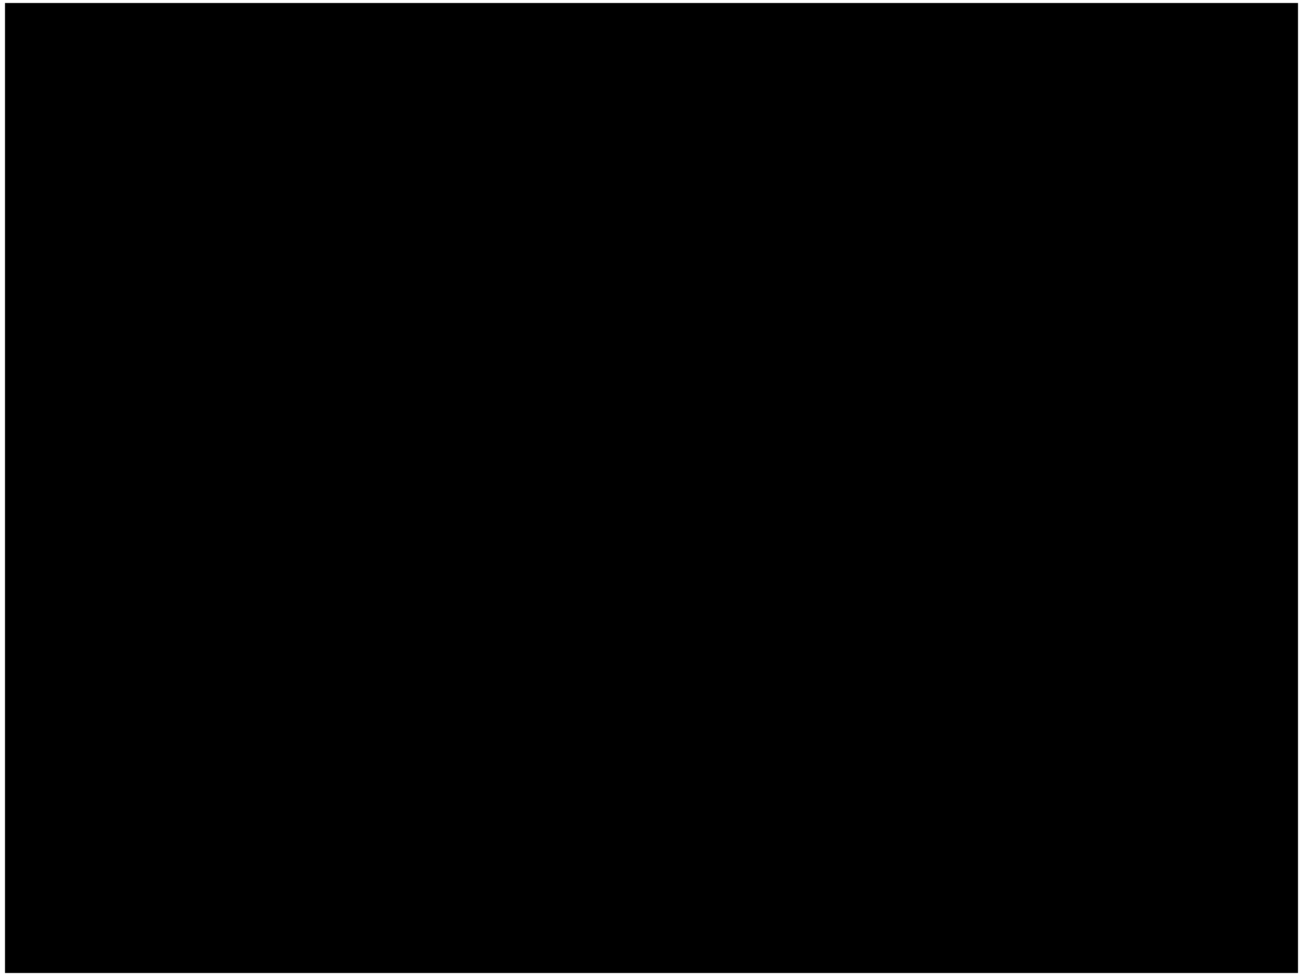

[Ng et al., SIGGRAPH 2005] [Lytro Inc.]

Measure 3D from a single 2D image

Coded aperture for single-image depth and refocusing

conventional vs
coded lens

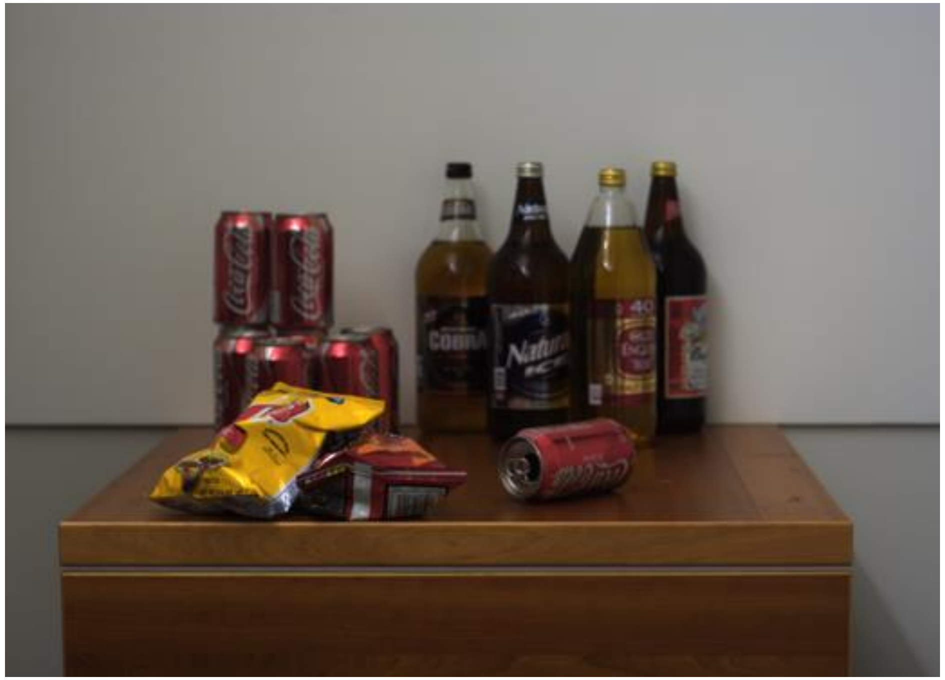

input image

inferred depth

[Levin et al., SIGGRAPH 2007]

Measure 3D from a single 2D image

Coded aperture for single-image depth and refocusing

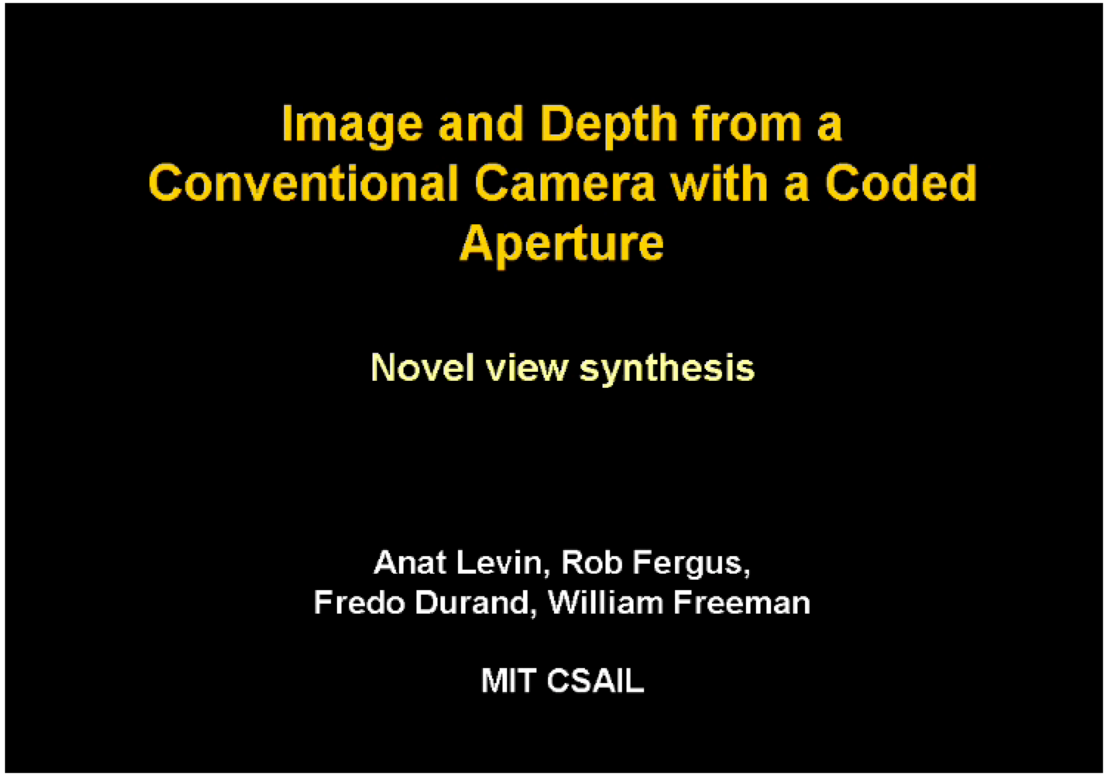

[Levin et al., SIGGRAPH 2007]

Remove lenses altogether

FlatCam: replacing lenses with masks

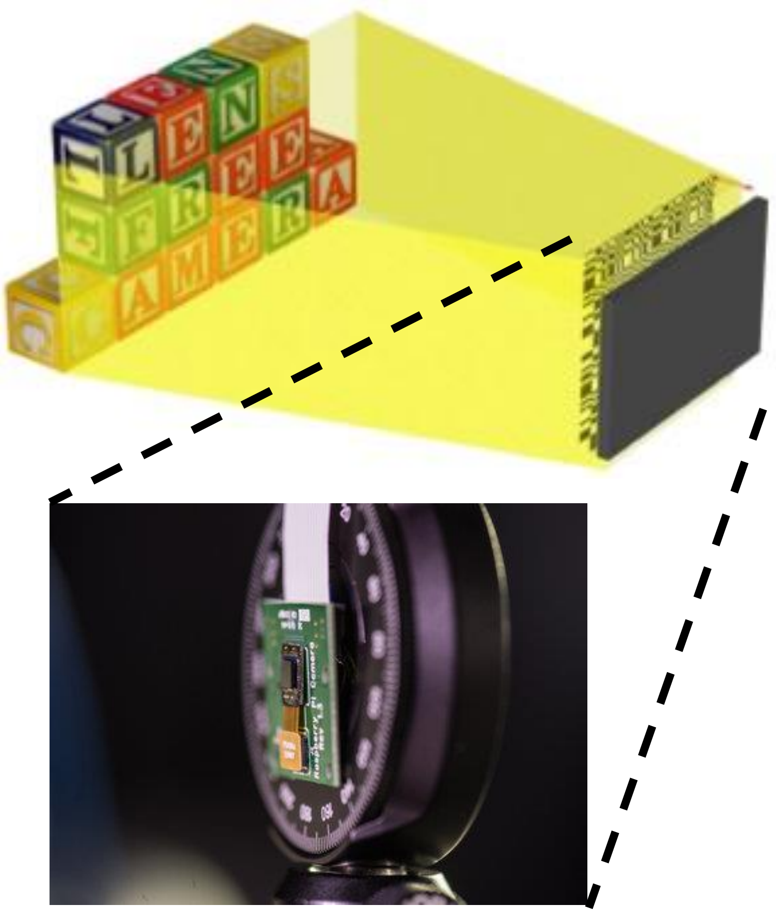

prototype

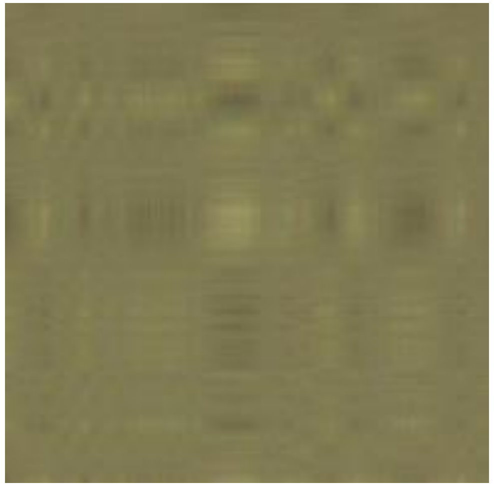

sensor measurements

reconstructed image

[Asif et al. 2015]

Computational photography

generalized optics
between scene and sensor

digital sensor to capture focused
light (electrical process)

arbitrary computation
between sensor and image

Computational photography

generalized optics
between scene and sensor

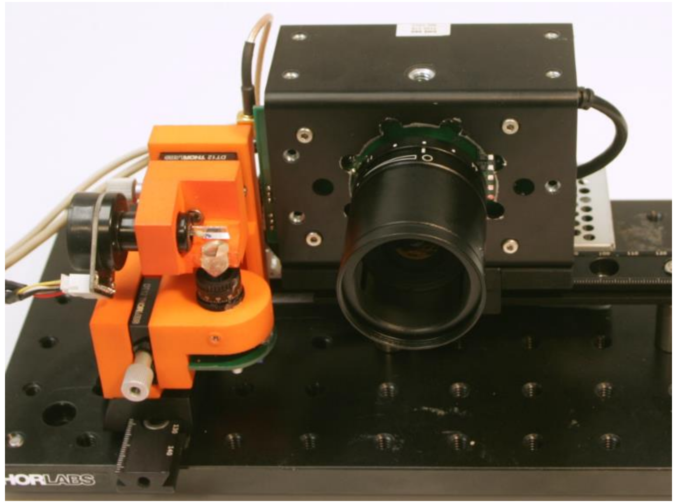

unconventional sensing
and illumination

arbitrary computation
between sensor and image

Measure depth

Time-of-flight sensors for real-time depth sensing

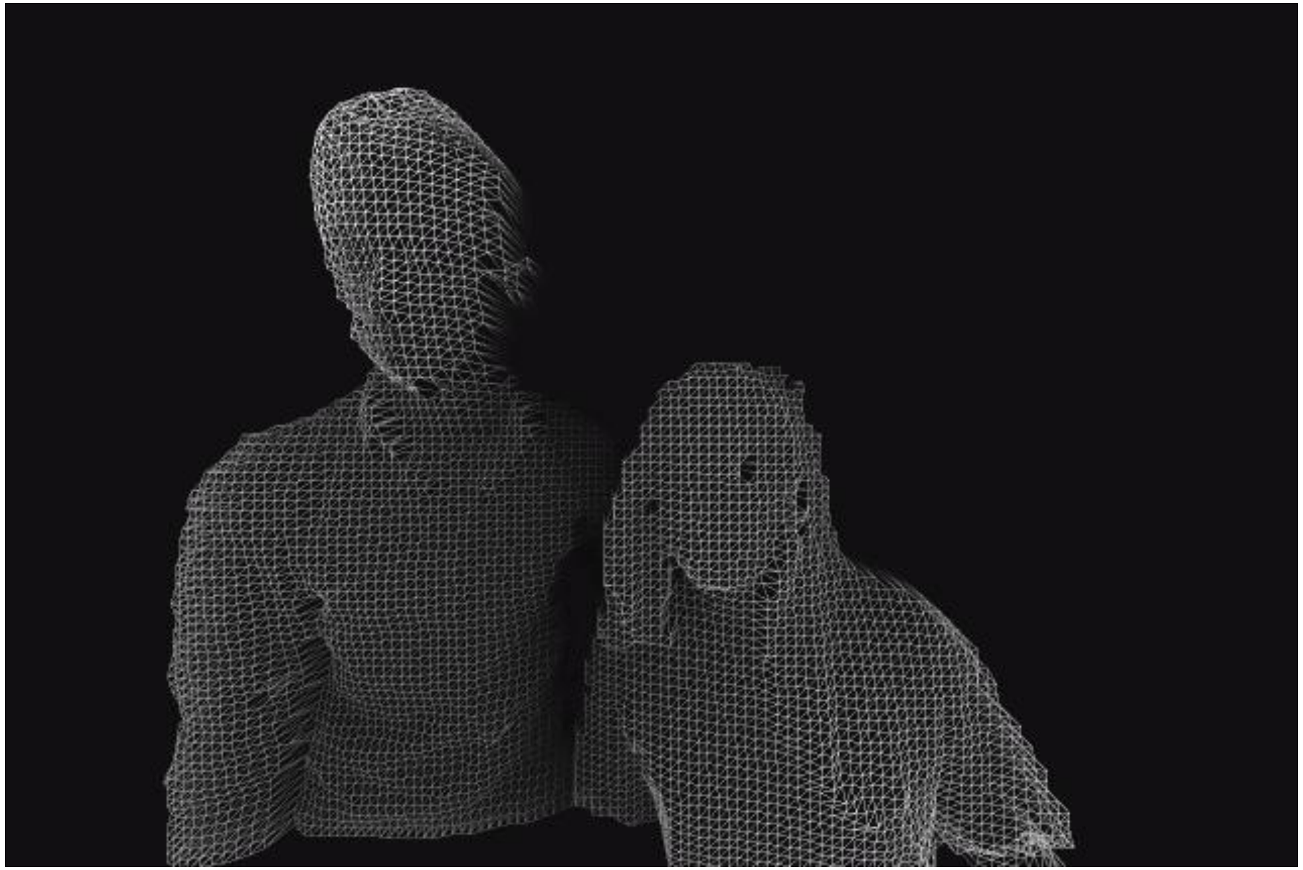

[Microsoft Inc.]

Measure light in flight

Streak camera for femtophotography

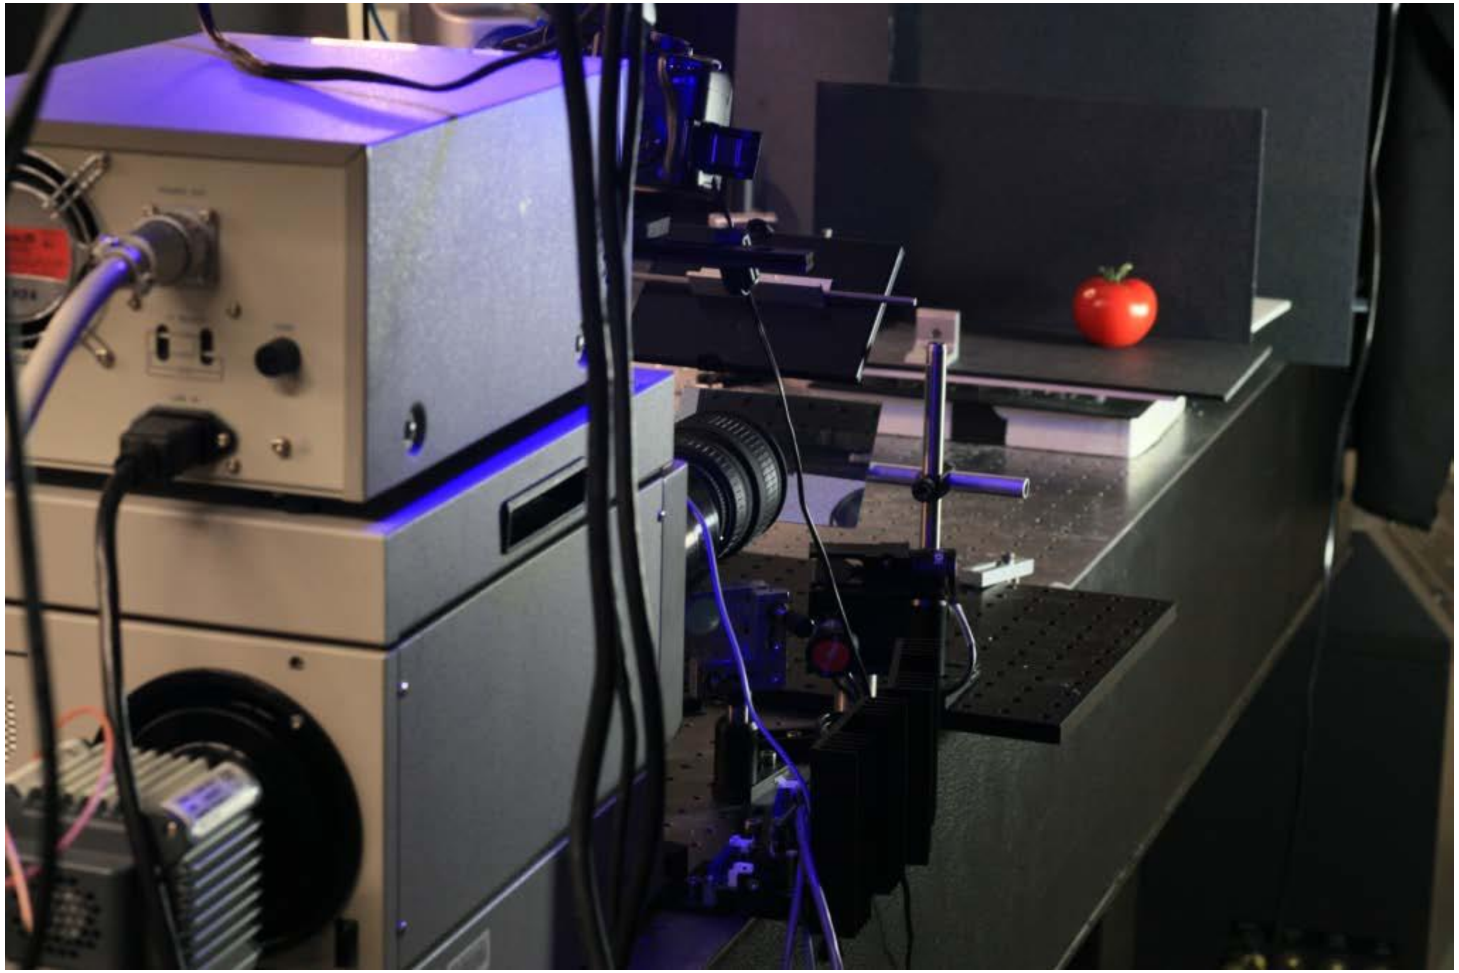

[Velten et al., SIGGRAPH 2013]

Measure light in flight

Streak camera for femtophotography

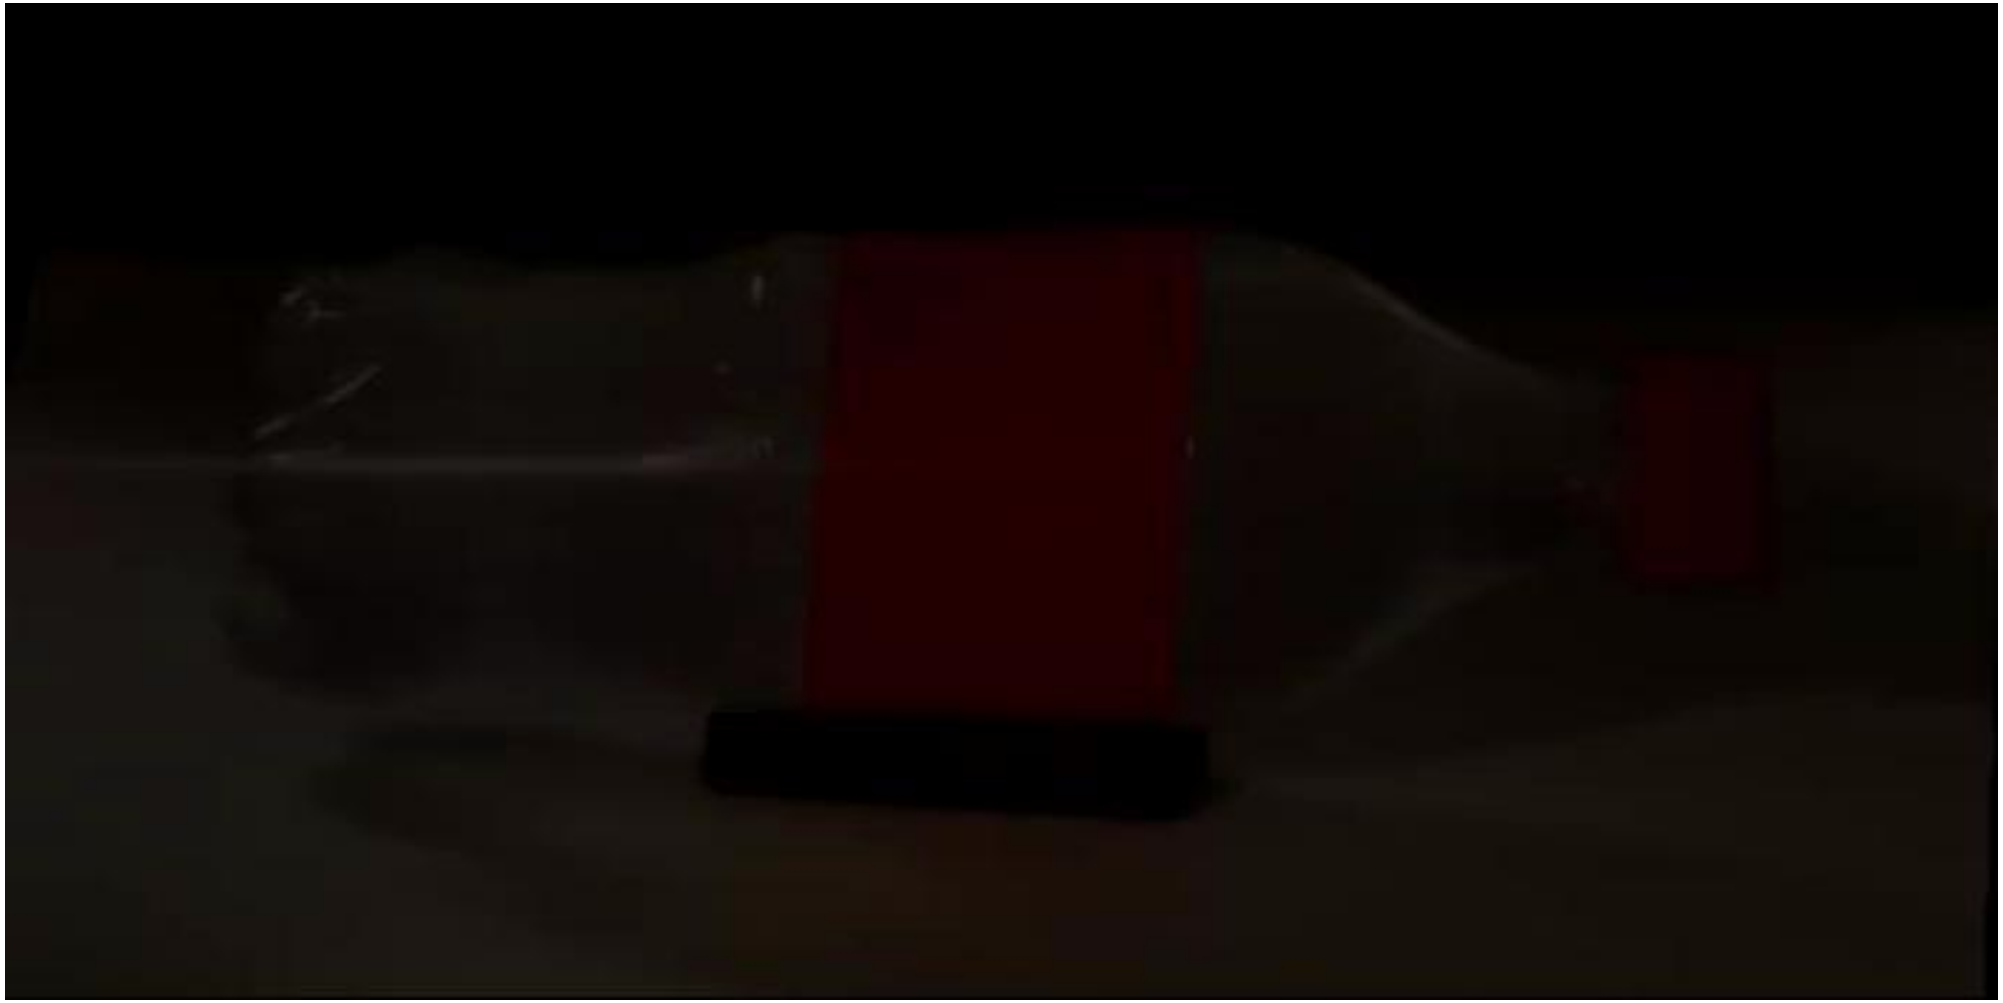

[Velten et al., SIGGRAPH 2013]

Measure photons selectively

Structured light for epipolar imaging

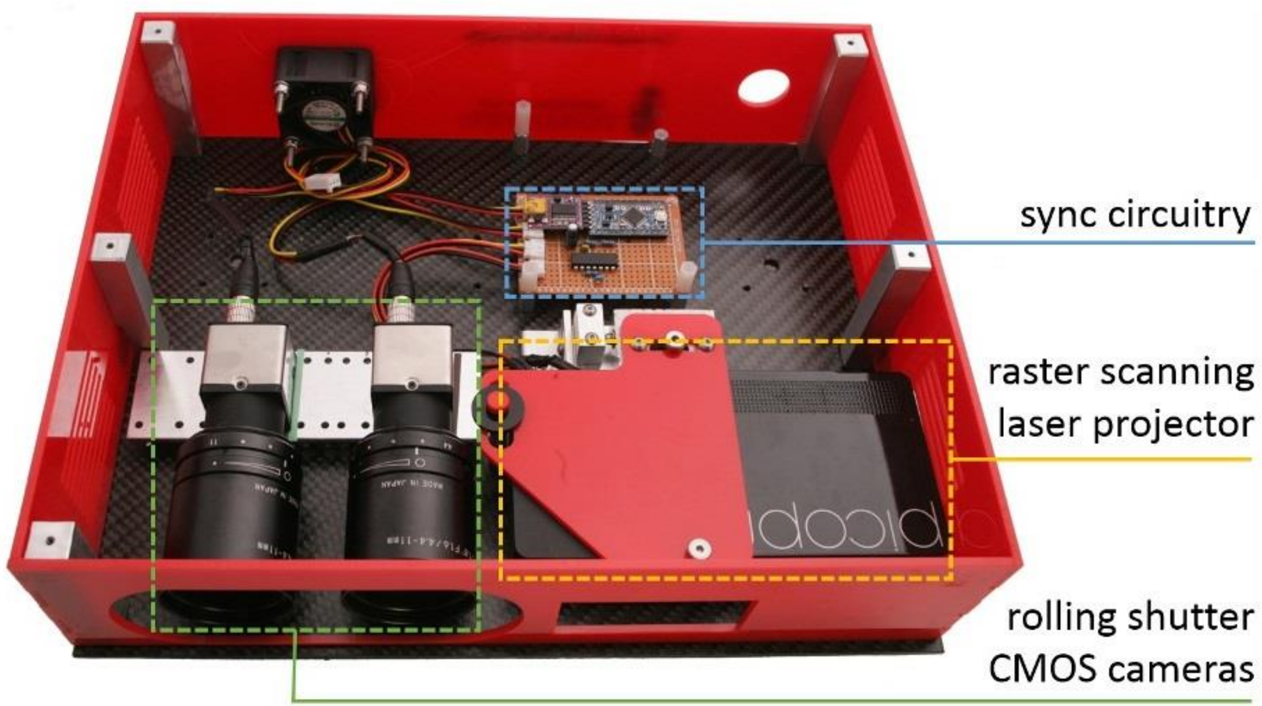

[O'Toole et al., SIGGRAPH 2015]

Measure photons selectively

One of your homeworks!

Structured light for epipolar imaging

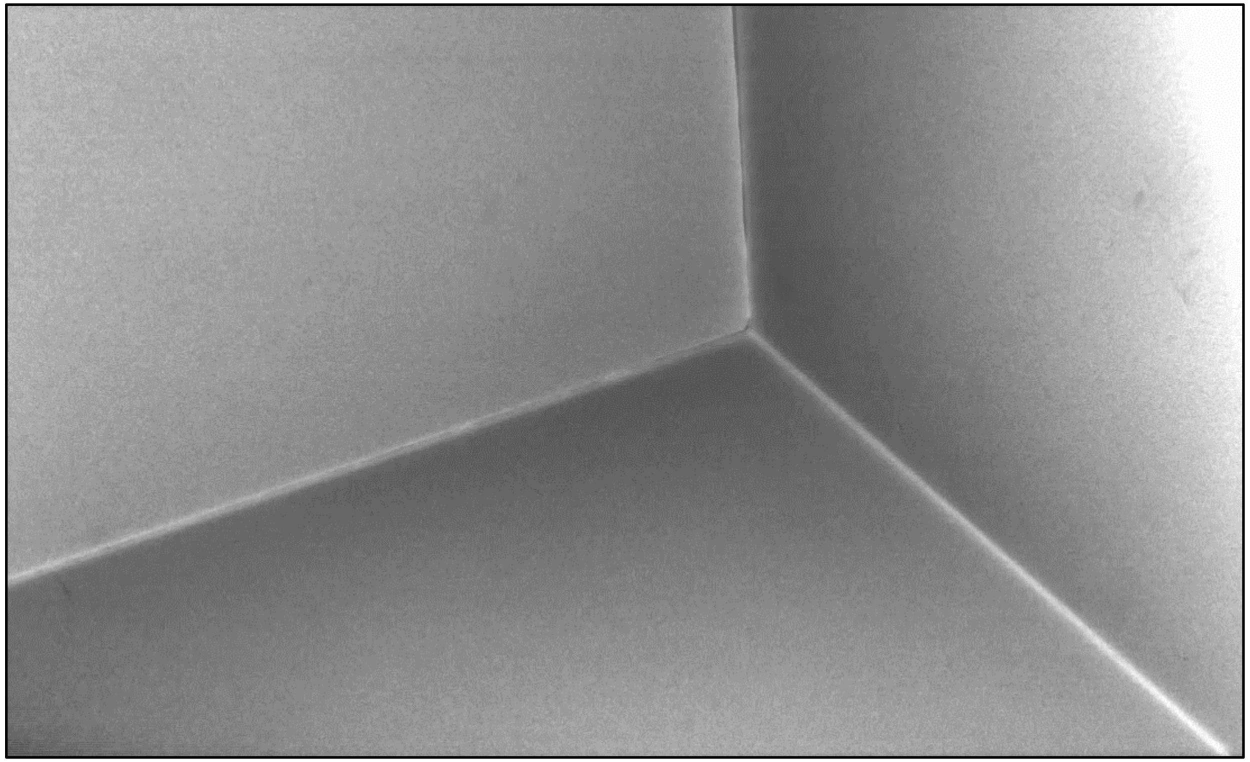

direct photons

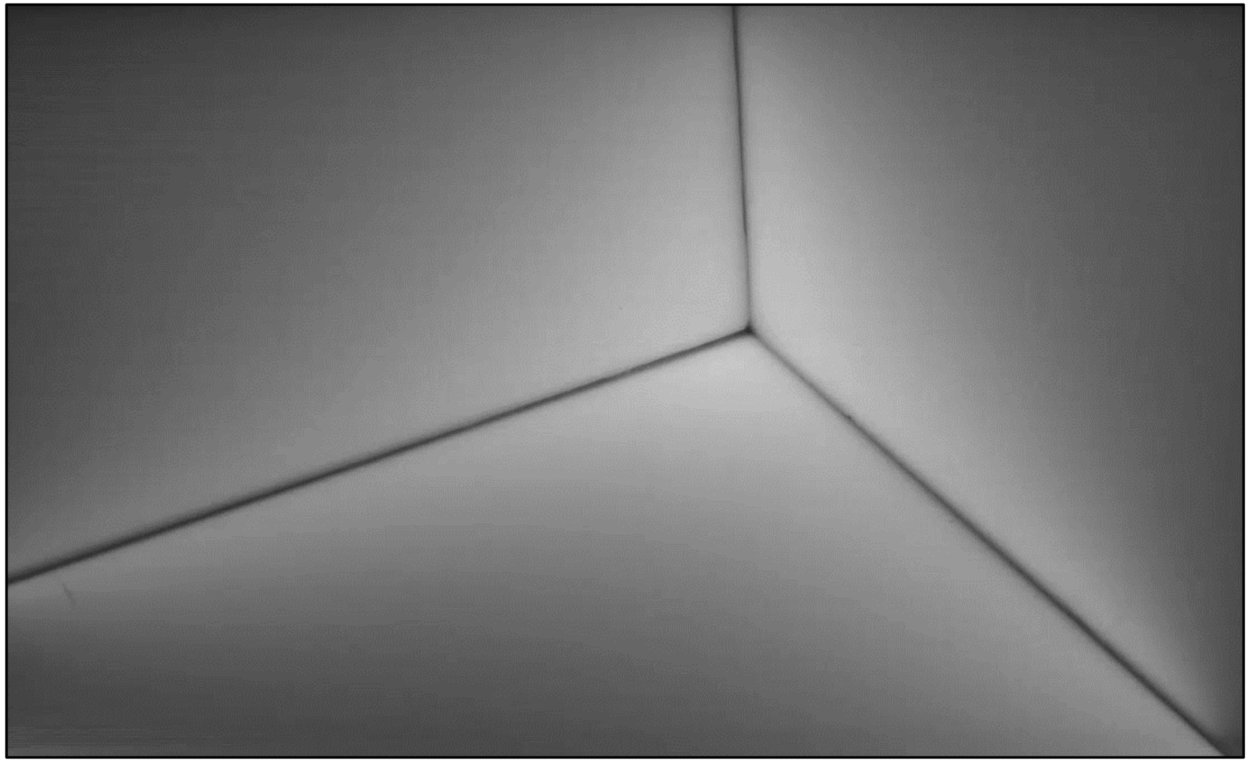

indirect photons

[O'Toole et al., SIGGRAPH 2015]

Computational photography

generalized optics
between scene and sensor

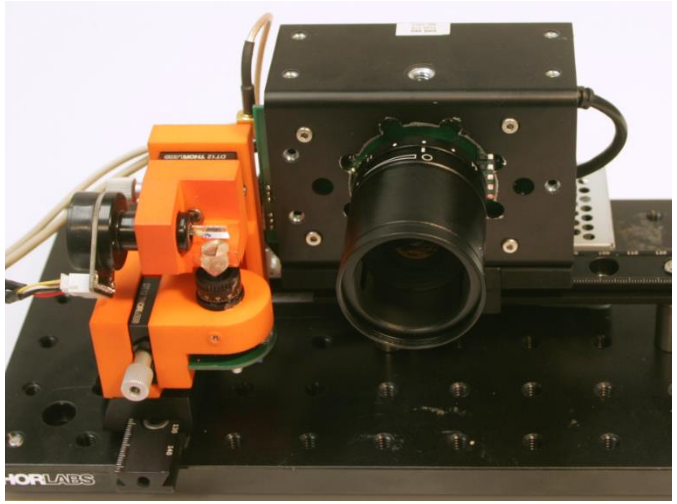

unconventional sensing
and illumination

arbitrary computation
between sensor and image

Computational photography

generalized optics
between scene and sensor

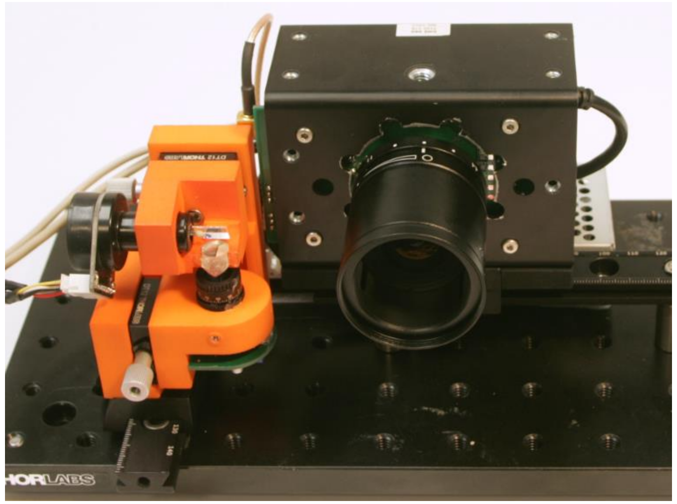

unconventional sensing
and illumination

arbitrary computation
between sensor and image

joint design of optics, illumination, sensors, and computation

Putting it all together

Looking around corners

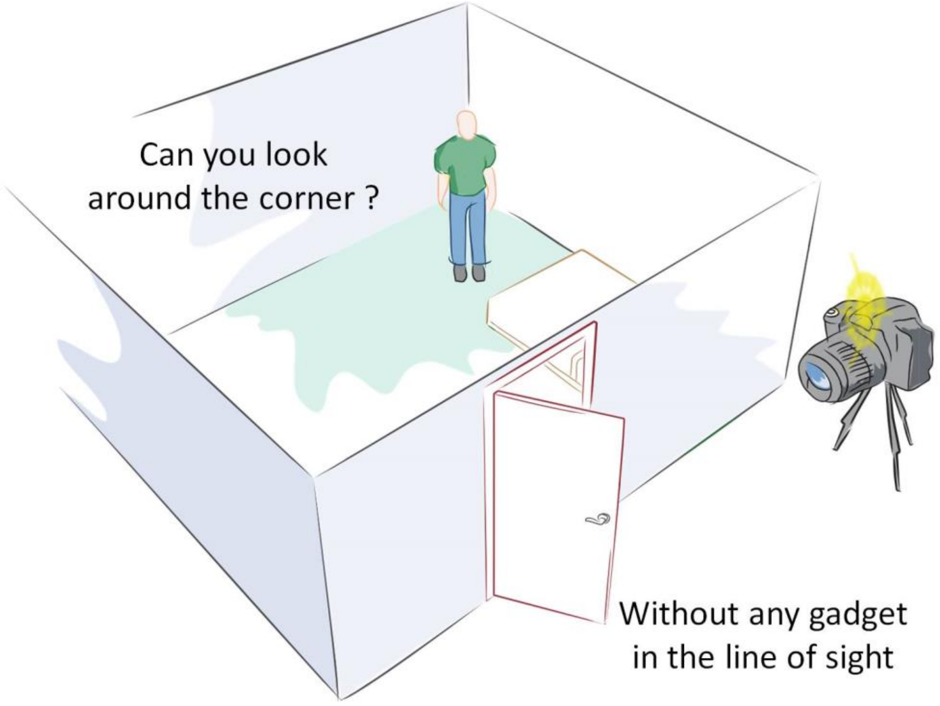

One of your homeworks!

[MIT Media Lab, DARPA REVEAL]

Putting it all together

Looking through tissue

Opportunity

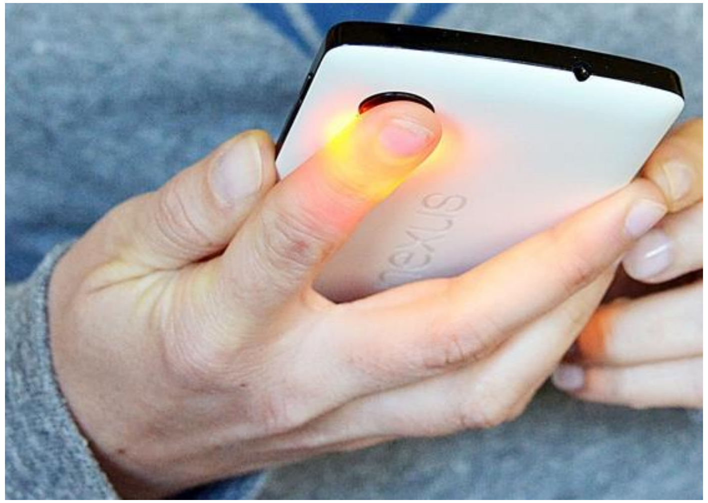

+ Light travels deep inside the body
+ It is non-ionizing (400-1100nm)
+ Cheap to produce and control

Scattering Barrier

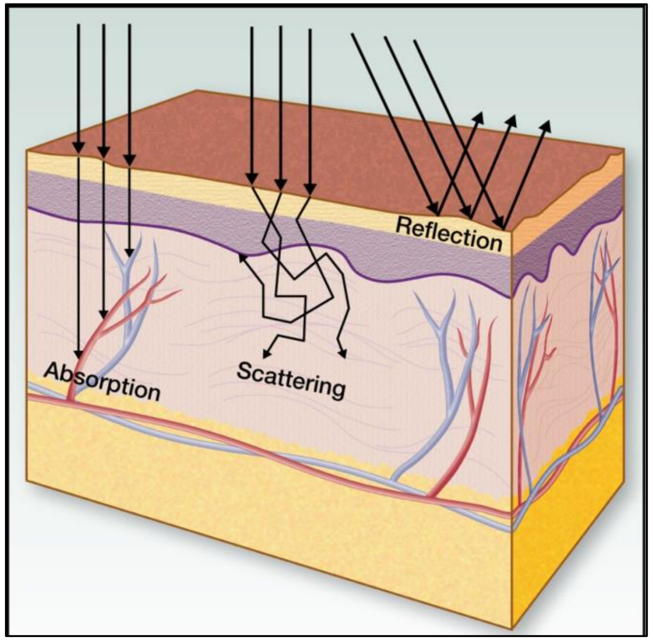

- Most pass-through photons are scattered
- Avg 10 scattering events per mm
- By 50mm, avg 500 scattering events !
- Large-scale inverse problem with low SNR

Practical imaging up to 50mm

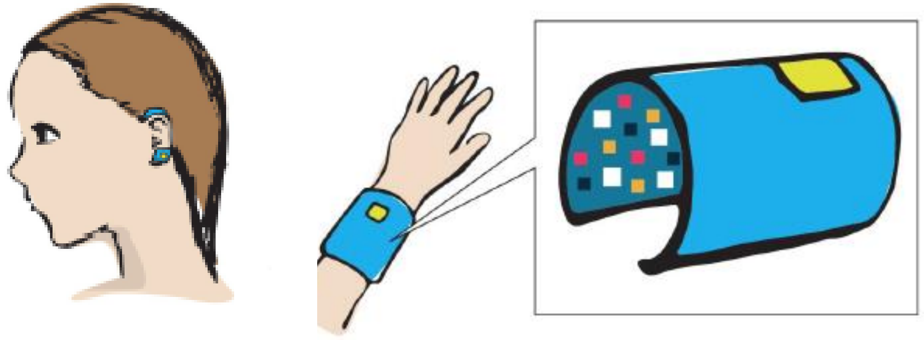

Wearables (1-10mm)

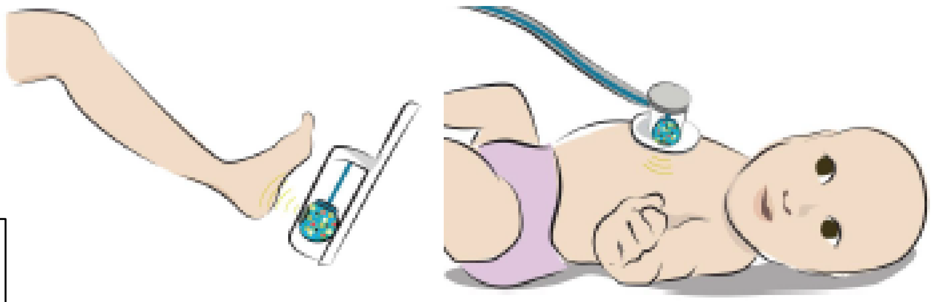

Non-invasive point of care devices (10-50mm)

[NSF Expedition]

Computational photography

generalized optics
between scene and sensor

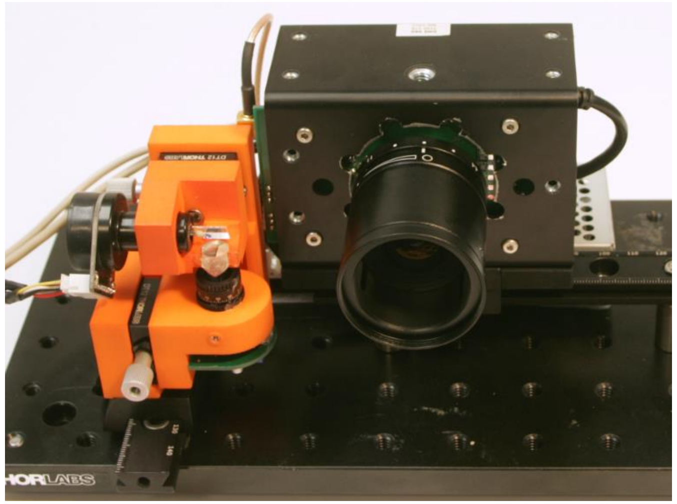

unconventional sensing
and illumination

arbitrary computation
between sensor and image

joint design of optics, illumination, sensors, and computation

Course fast-forward and logistics

Course fast-forward

Tentative syllabus at:

http://graphics.cs.cmu.edu/courses/15-463

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

Topics to be covered

Digital photography:

- optics and lenses
- color
- exposure
- aperture
- focus and depth of field
- image processing pipeline

[Photo from Gordon Wetzstein]

Topics to be covered

Image manipulation and merging:

- image filtering
- image compositing
- image blending
- image warping
- morphing
- high-performance image processing

[Banerjee et al., SIGGRAPH 2014]

Topics to be covered

Types of cameras:

- geometric camera models
- light-field cameras
- coded cameras
- lensless cameras
- compressive cameras
- hyperspectral cameras

Topics to be covered

Active illumination and sensing:

- time-of-flight sensors
- structured light
- computational light transport
- transient imaging
- non-line-of-sight imaging
- optical computing

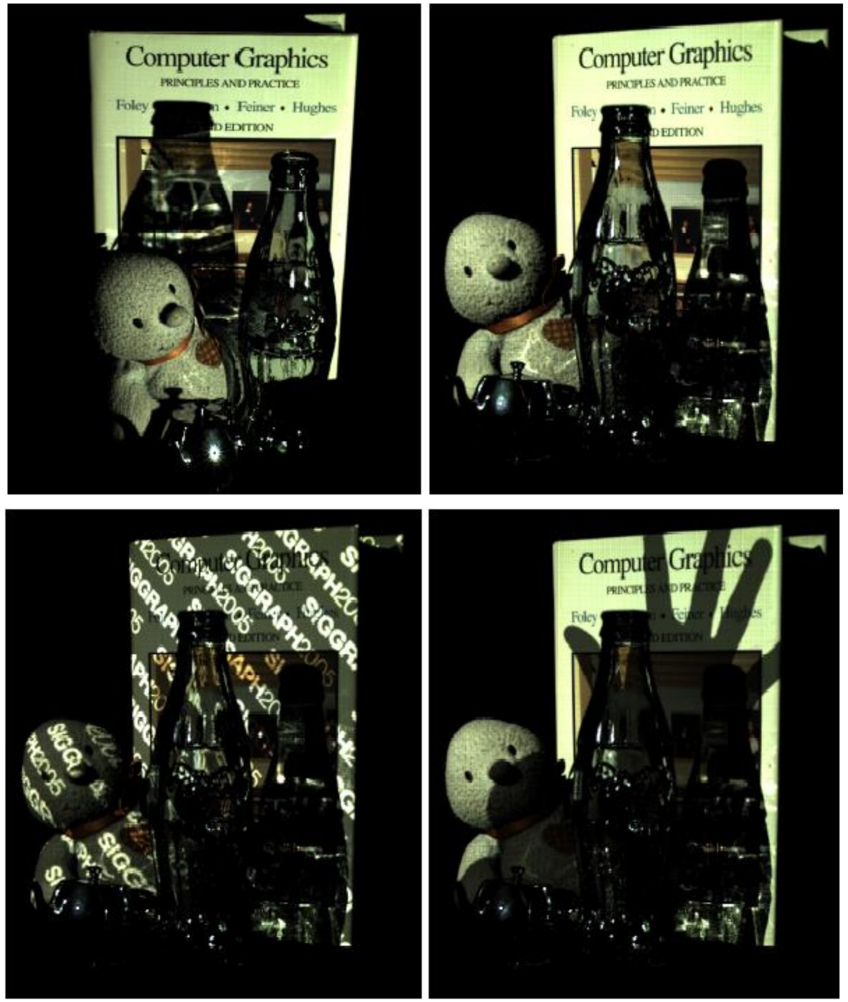

[Sen et al., SIGGRAPH 2005]

Course logistics

- Course website:

  http://graphics.cs.cmu.edu/courses/15-463

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

  https://piazza.com/class/jzoctlm269oe0

- Canvas for homework submissions:

  https://canvas.cmu.edu/courses/12162

Please take the start-of-semester survey!

- Posted on Piazza as well:

  https://docs.google.com/forms/d/e/1FAIpQLScBsG58jLZ-B6krqS5Jd4QnP5fDljz7uC-ZrDyKdARcZYDsZQ/viewform

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 15-420.
- An image processing course at the level of 18-793.

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.
- Laplacian pyramid.
- Poisson blending.
- Homogeneous coordinates.
- Homography.
- RANSAC.
- Epipolar geometry.
- XYZ space.
- Radiance and radiometry.
- Lambertian, diffuse, and specular reflectance.
- n-dot-l lighting.
- Monte Carlo rendering.
- Thin lens, prime lens, and zoom lens.
- Demosaicing.
- Refraction and diffraction.

Evaluation

- Six-plus-one homework assignments (60% + 10%):
  - programming and capturing your own photographs.
  - all programming will be in Matlab.
  - first assignment will serve as a gentle introduction to Matlab.
  - five late days, you can use them as you want.

- Final project (35%):
  - we will provide more information near the end of September.
  - 15-663, 15-862 require more substantive project.
  - if your ideas require imaging equipment, talk to us in advance.
  - **no exam, but final project presentations are during the exam period.**

- Class and Piazza participation (5%):
  - be around for lectures.
  - participate in Piazza discussions.
  - ask questions.

Do I need a camera?

- You will need to take your own photographs for assignments 1-7 (all of them):
    - Assignment 1: pinhole camera – you need a high-sensitivity camera.
    - Assignment 2: HDR – you need a camera with manual exposure controls.
    - Assignment 3: image filtering – you can use your phone camera.
    - Assignment 4: lightfields – you need a camera with manual focus control..
    - Assignment 5: photometric stereo – you need a camera with RAW support.
    - Assignment 6: structured light – you can use your phone camera.
    - Assignment 7: corner cameras – you need a high-sensitivity camera.

- We have 20 Nikon D3300/3400 kits (camera + lens + tripod) for students.
    - If you have your own camera, please use that!

Contact information and office hours

- Feel free to email us about administrative questions.
  - please use [15463] in email title!
- Technical questions should be asked on Piazza.
  - we won’t answer technical questions through email.
  - you can post anonymously if you prefer.
- Office hours will be determined by poll.
  - feel free to email Yannis about additional office hours.
  - you can also just drop by Yannis’ office (Smith Hall (EDSH) Rm 225).

We are looking for undergraduate research assistants

- Projects relating to imaging, rendering, and graphics in general.
- Funding available if you are eligible for NSF REU.
- Please email if interested.

Please take the course survey (posted on Piazza) before the next lecture!