Introduction

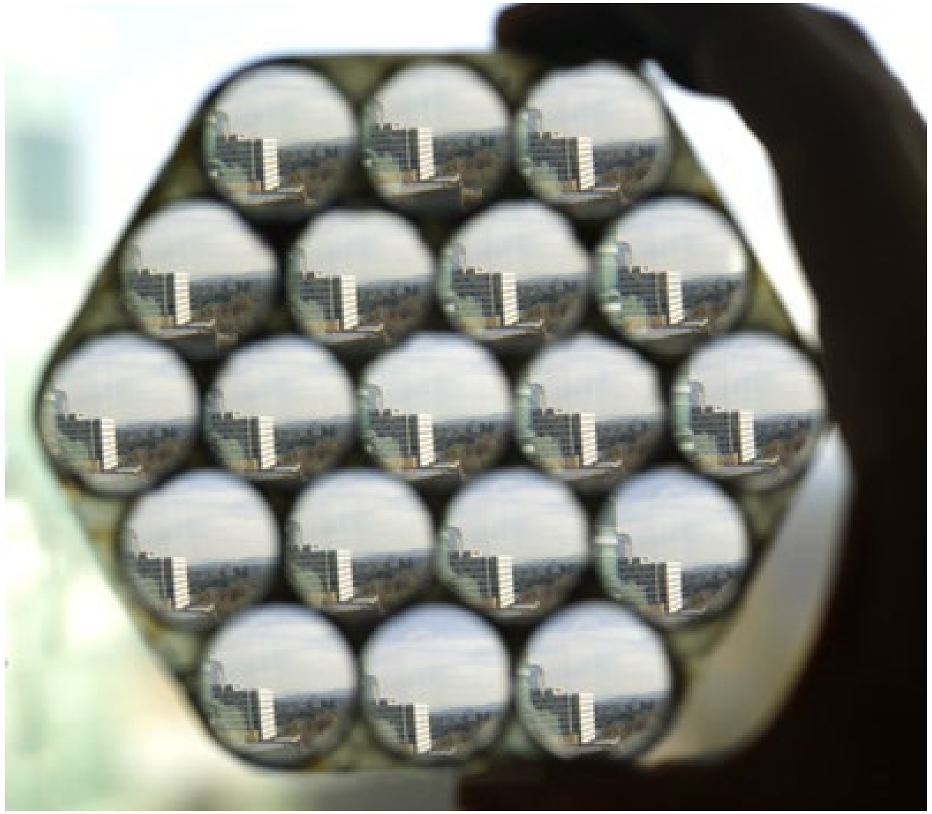

15-463, 15-663, 15-862
Computational Photography
Fall 2024, Lecture 1

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

Lecture etiquette

- Lecture slides (PDF) are posted on the course website before each lecture.
- You are expected to attend lectures in person.
- Feel free to ask questions! Please make sure to raise your hand to both ask your own questions and answer mine.

Overview of today’s lecture

- Teaching staff introductions
- What is computational photography?
- Course fast-forward and logistics

Teaching staff introductions

Instructor: Ioannis (Yannis) Gkioulekas

You can call me Yannis.

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/

3D sensing

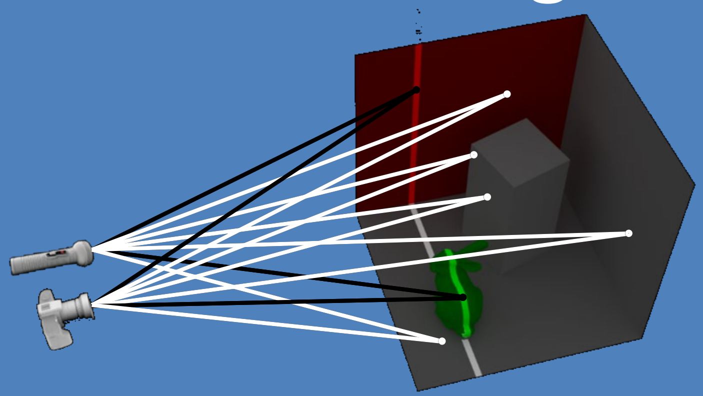

- micrometer 3D sensing using interferometry
- ultrafast 3D sensing using ultrasonic lenses

speckle imaging

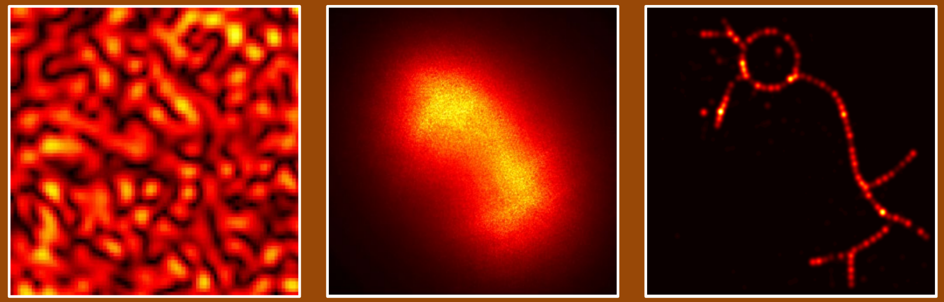

- Monte Carlo rendering of wave optics
- fluorescence microscopy

acousto-optics

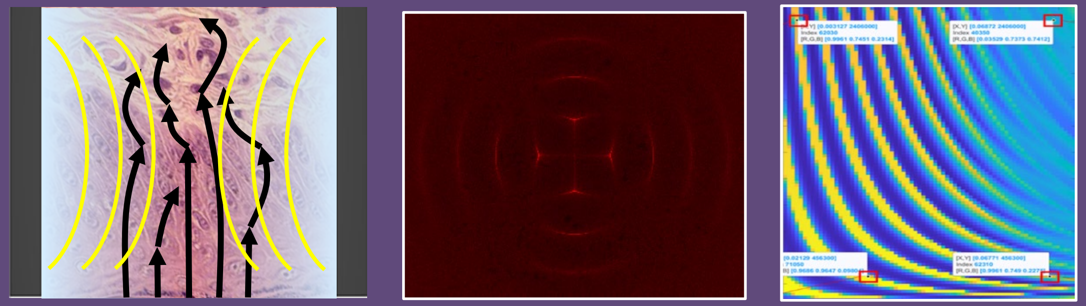

- rendering refractive radiative transfer
- steering light inside tissue

tissue imaging

- differentiable rendering for tissue tomography
- blood and vein imaging

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

Looking around corners

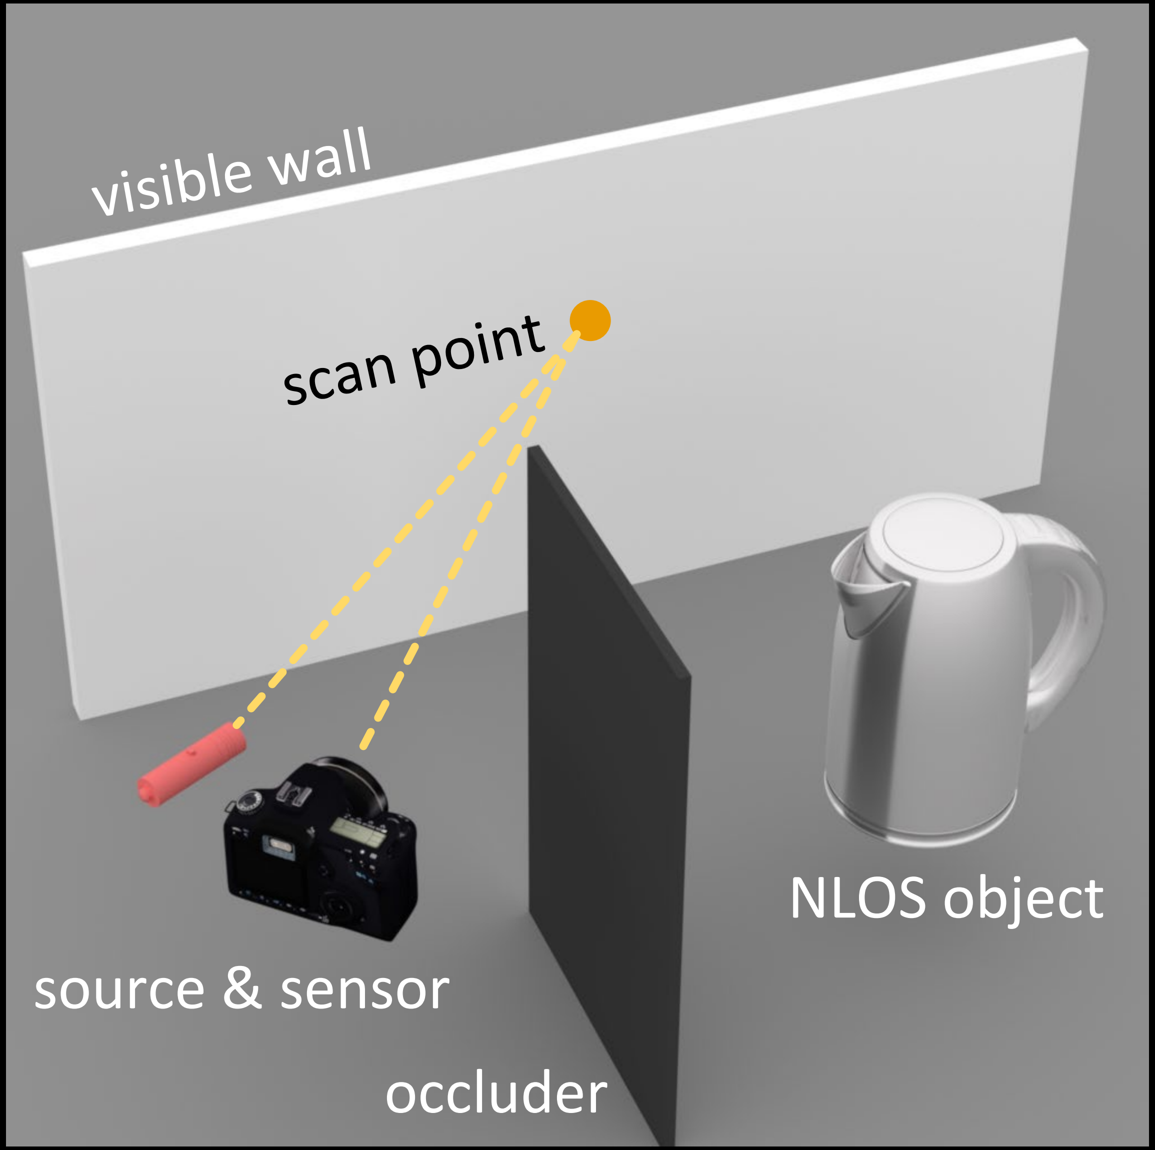

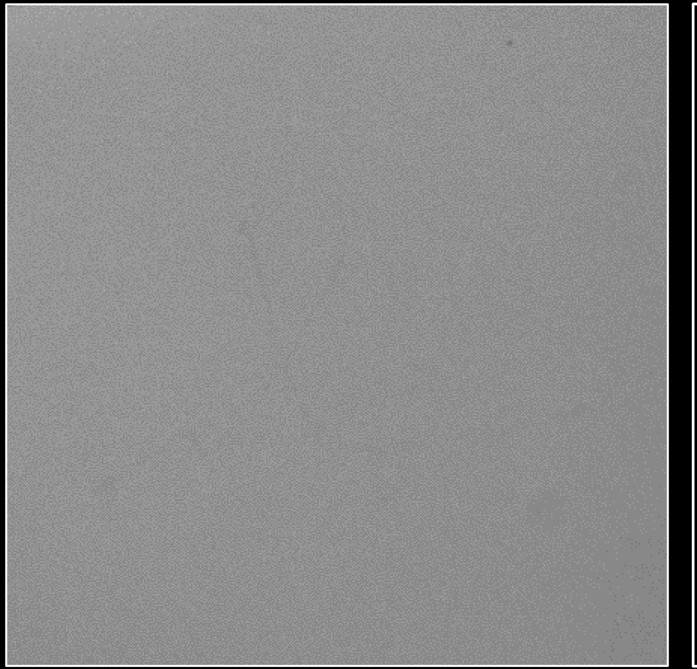

what a regular camera sees

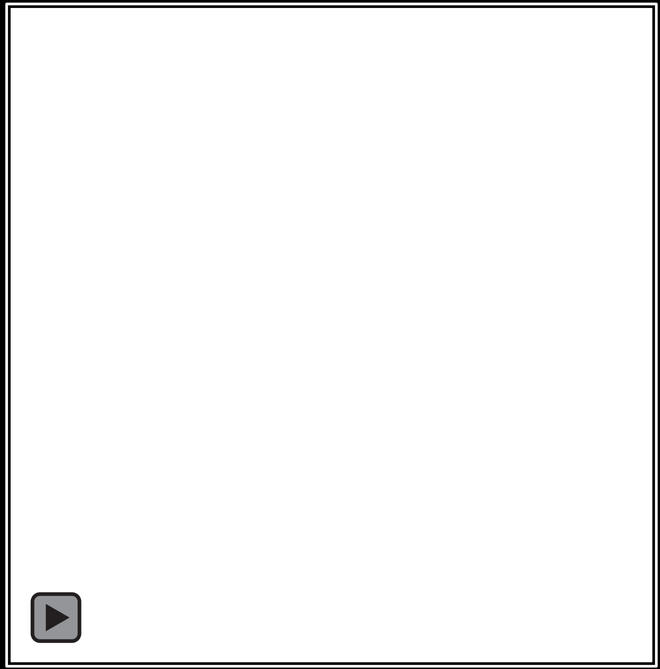

what we can reconstruct

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

Seeing light in flight

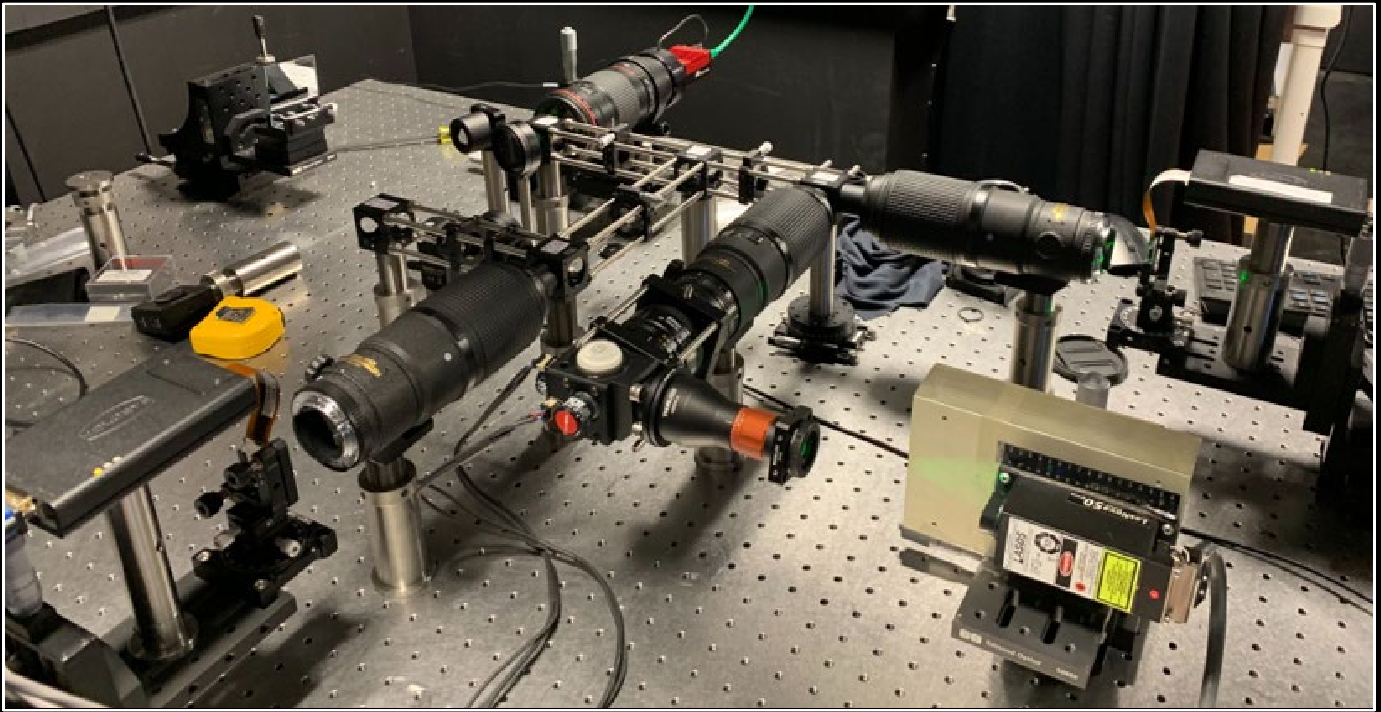

camera for capturing video at 10^{15} frames per second

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

Ultrafast lidar

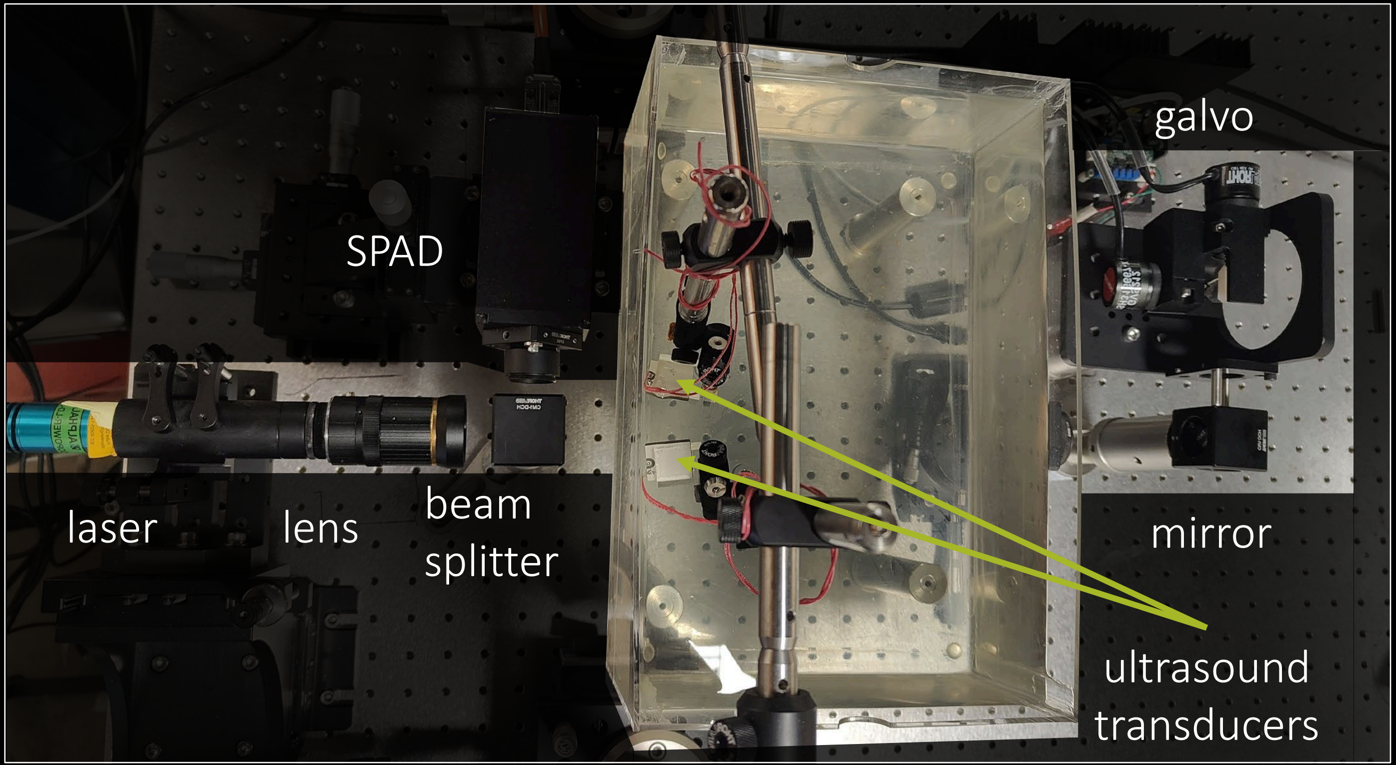

1000 × faster than a
commercial Velodyne lidar

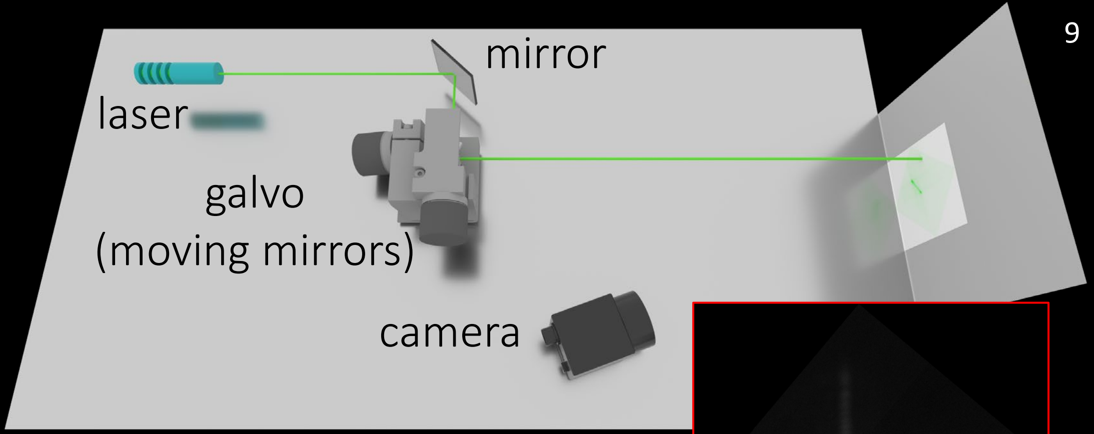

galvo steering (1 kHz)

1 ms exposure

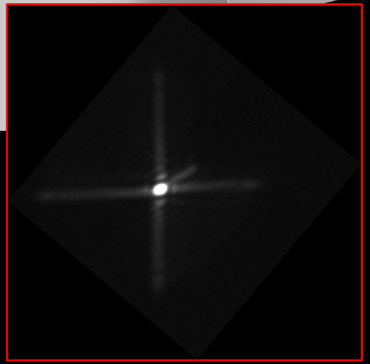

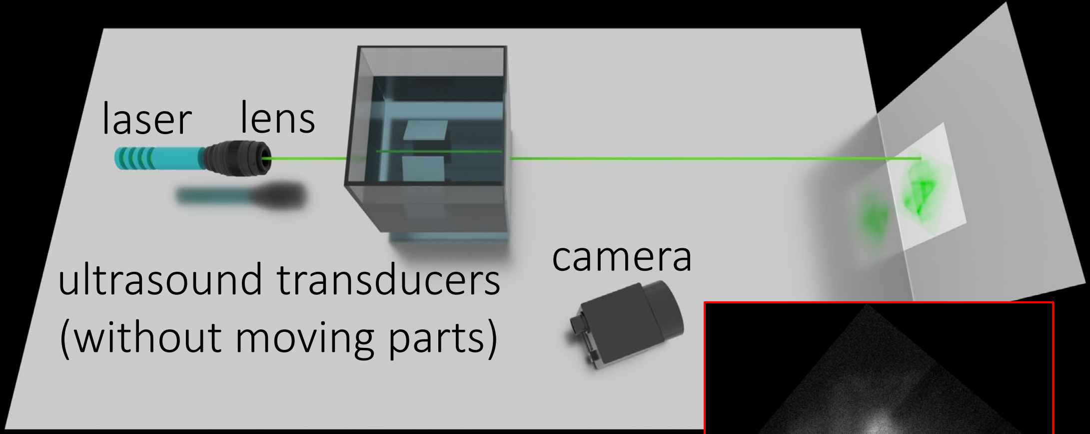

our acousto-optic steering (1 MHz)

1 ms exposure

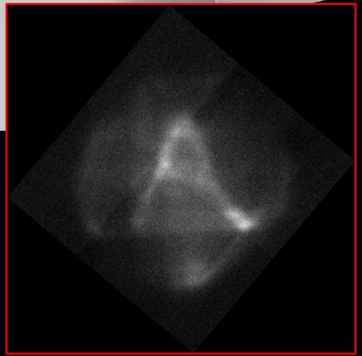

Sunlight micro-3D scanning

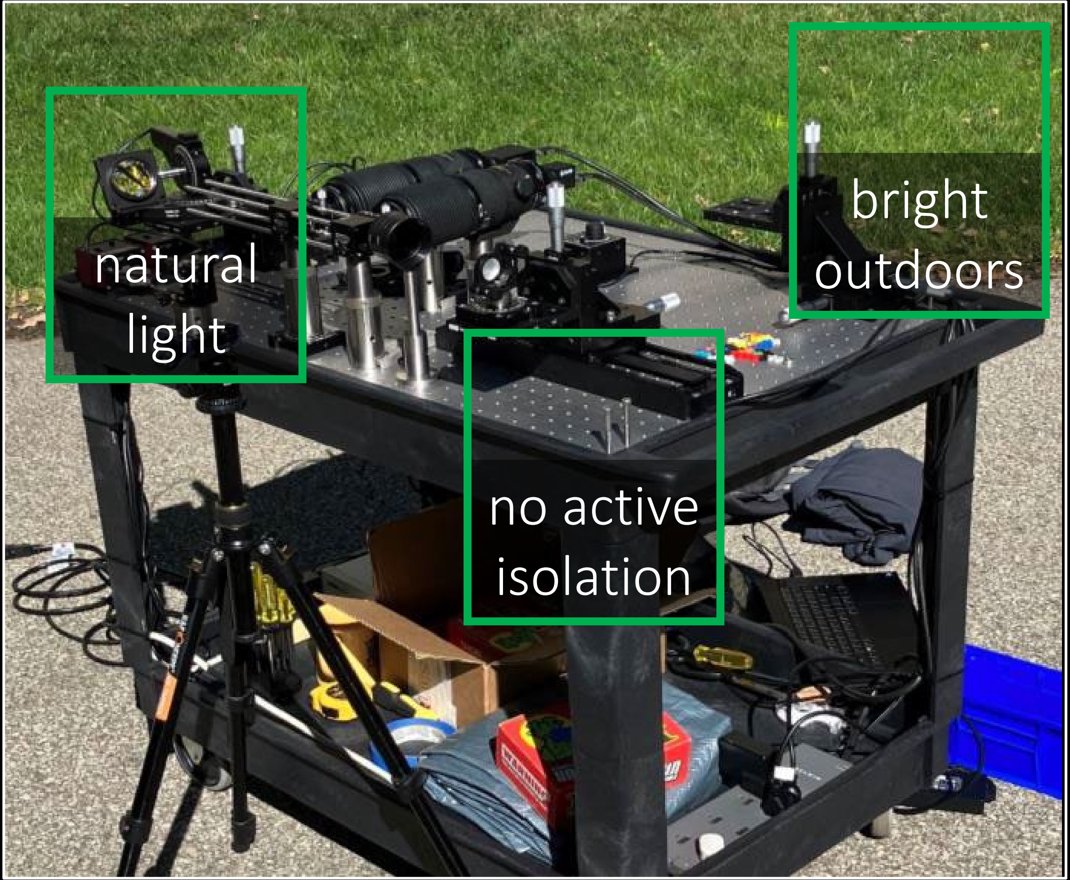

scene

depth map

input images

scene transient

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

What is computational photography?

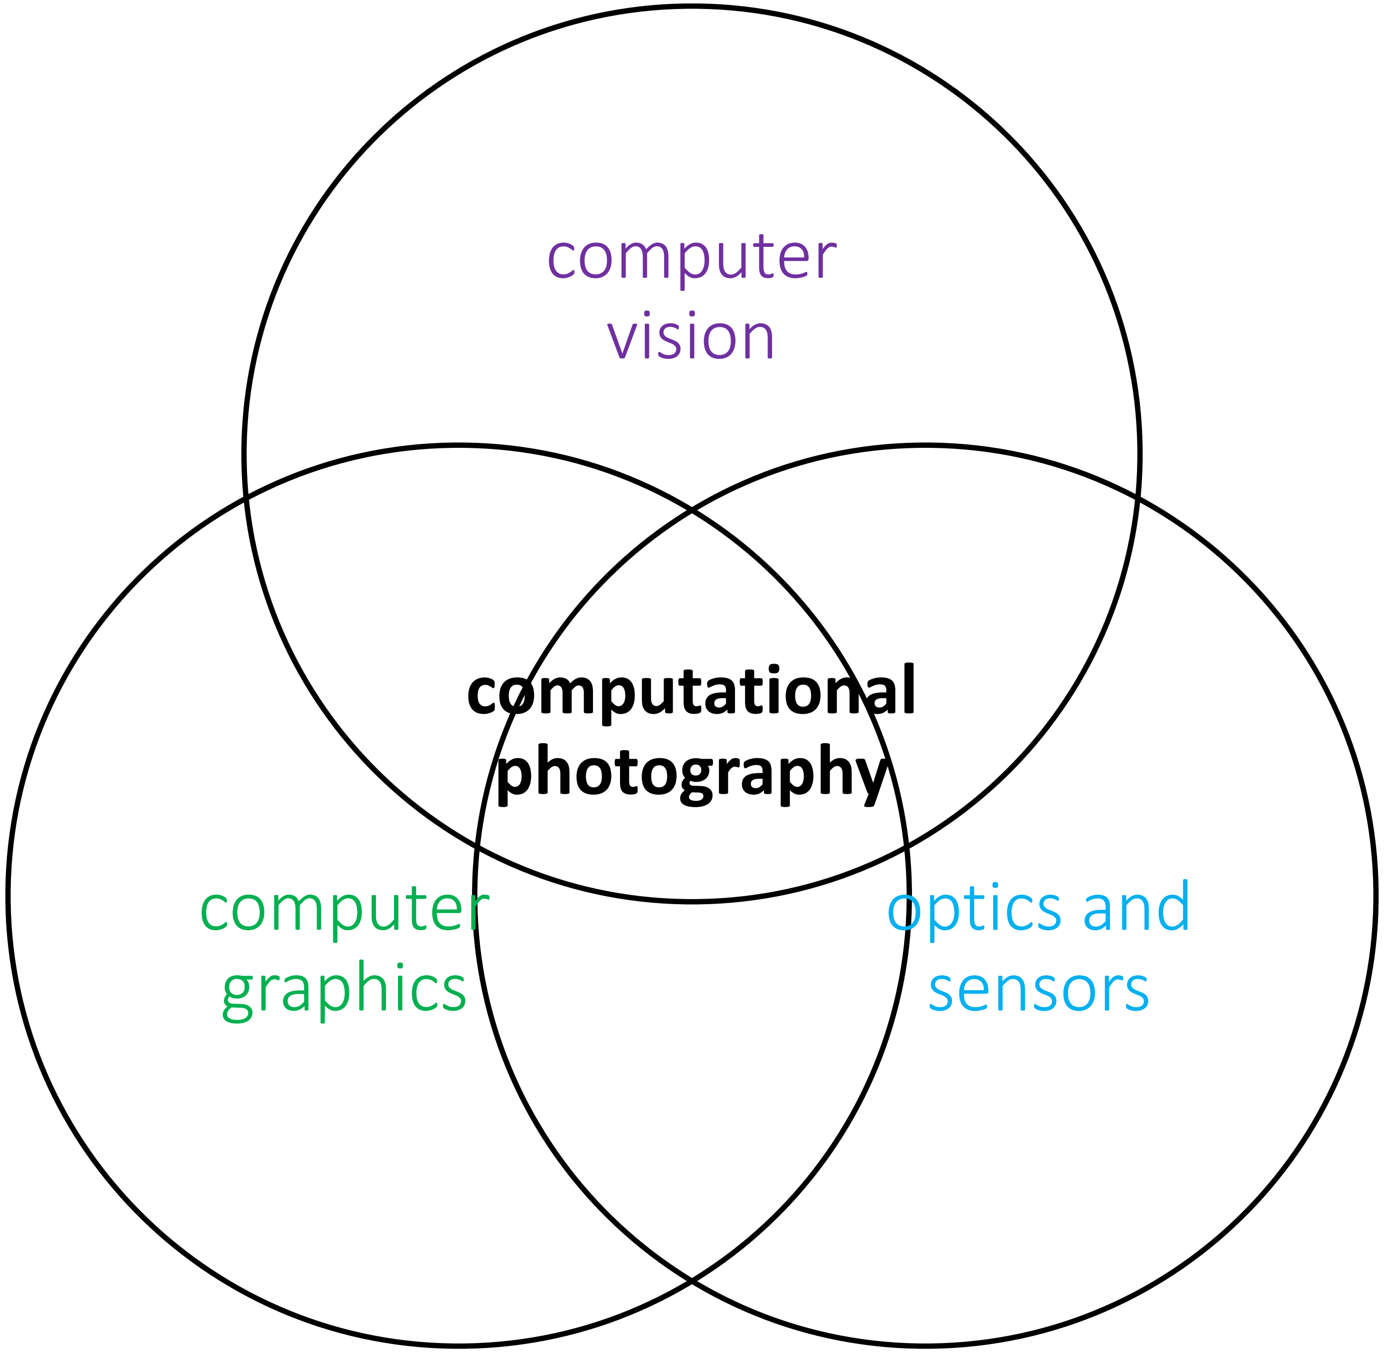

[Slide credit: Kris Kitani]

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

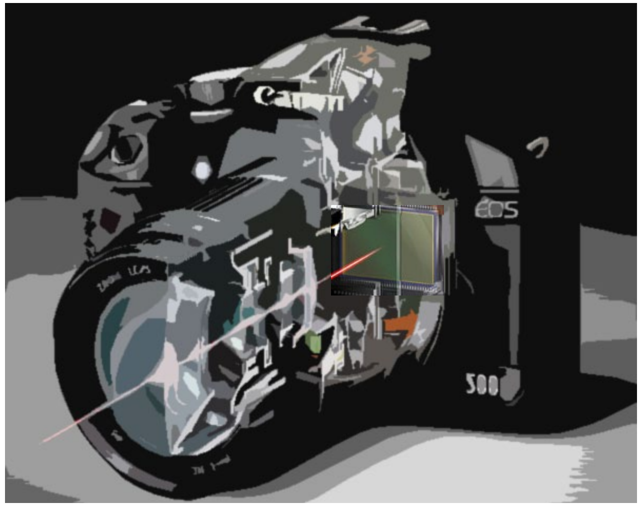

digital sensor to capture focused light (electrical process)

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

Computational photography

optics to focus light on an image plane

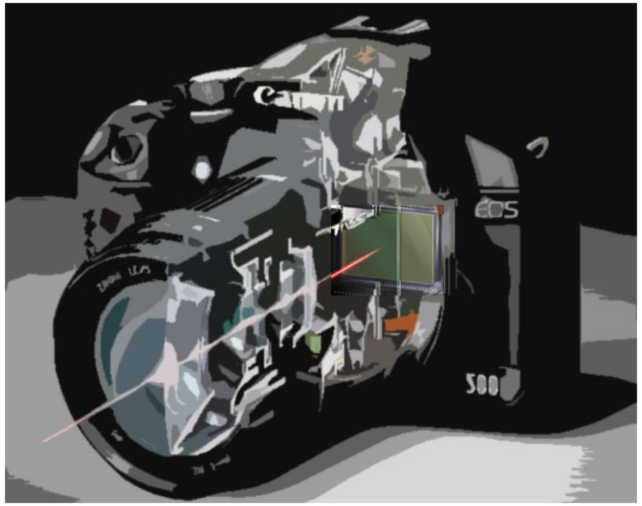

digital sensor to capture focused light (electrical process)

arbitrary computation between sensor and image

Overcome limitations of digital photography

Image enhancement and photographic look

camera output

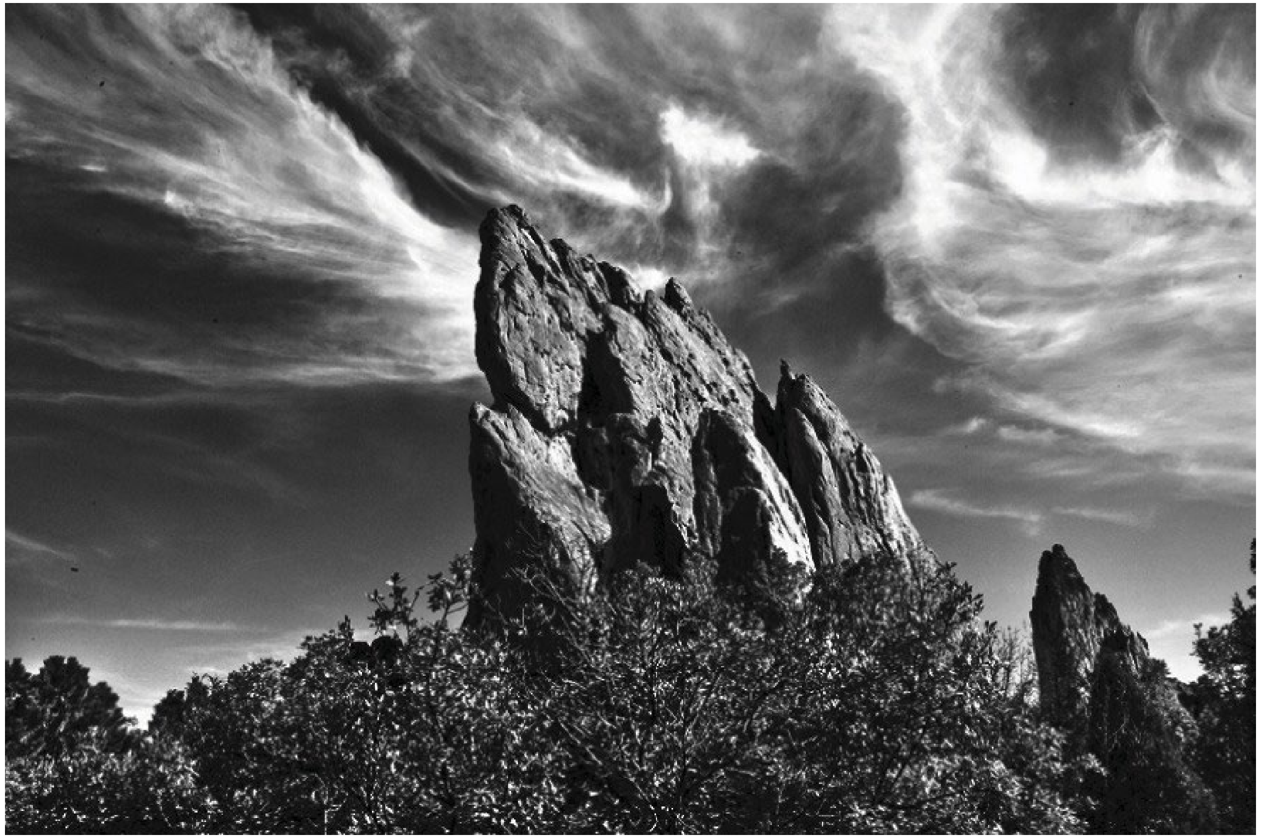

image after stylistic tonemapping

[Bae et al., SIGGRAPH 2006]

Overcome limitations of digital photography

High dynamic range (HDR) imaging

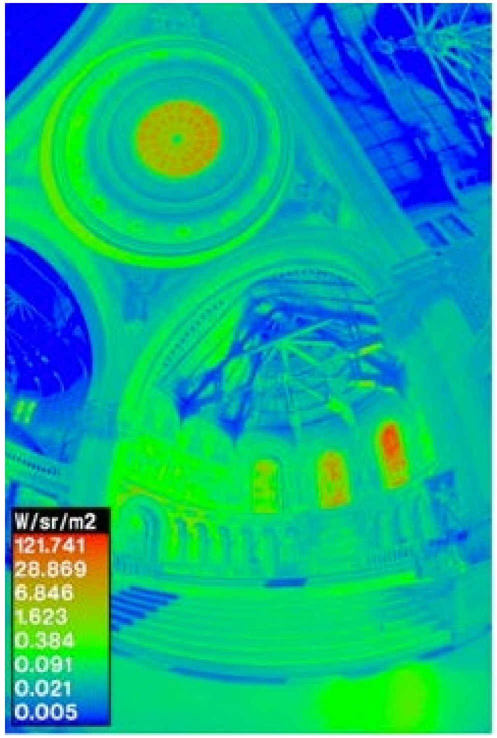

One of your homeworks!

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

Create realistic new imagery

Image blending and harmonization

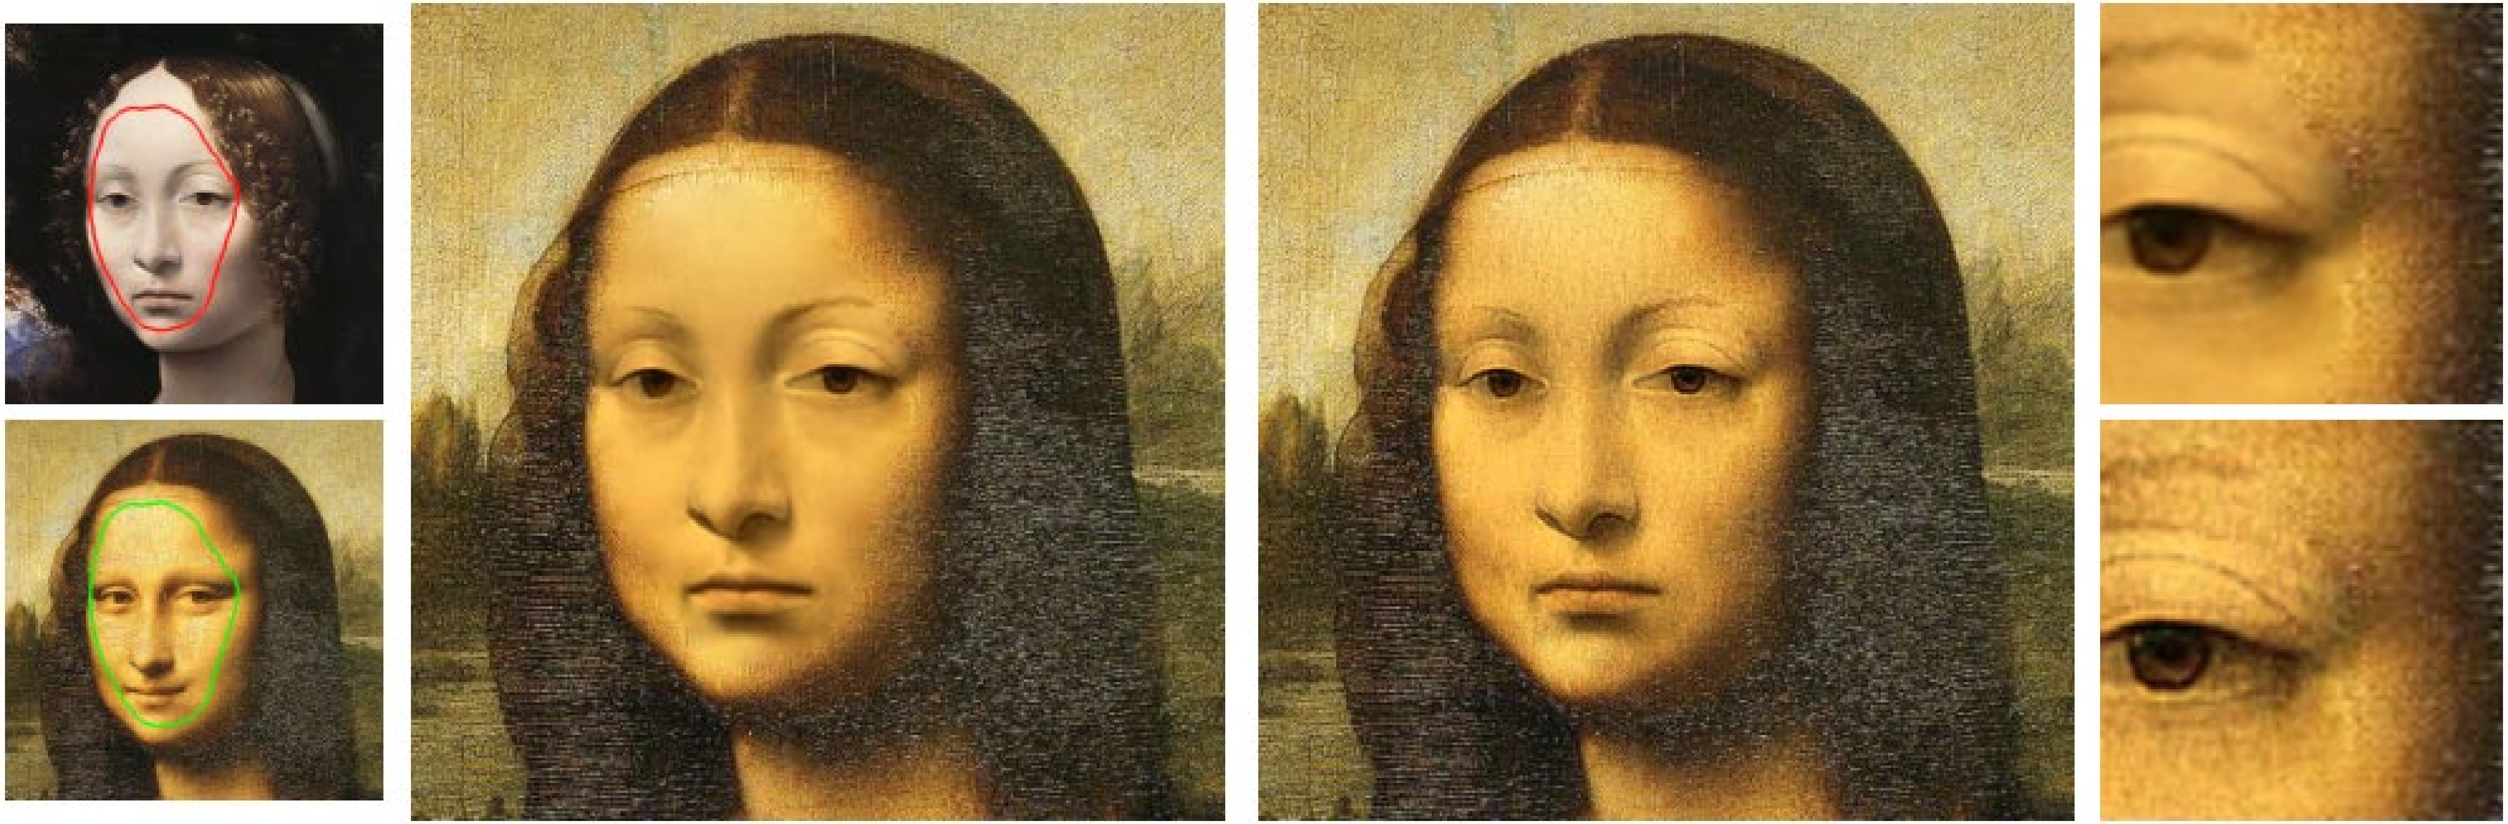

One of your homeworks!

[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

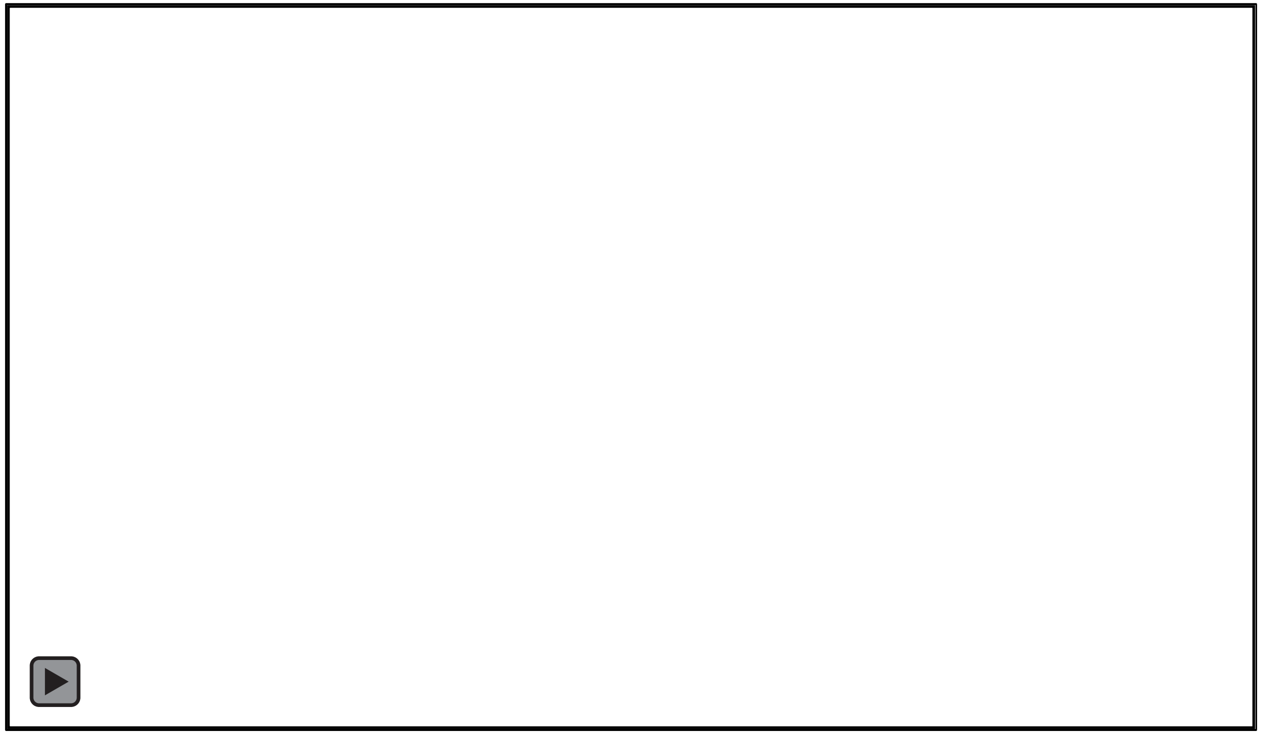

reconstructing cities from Internet photos

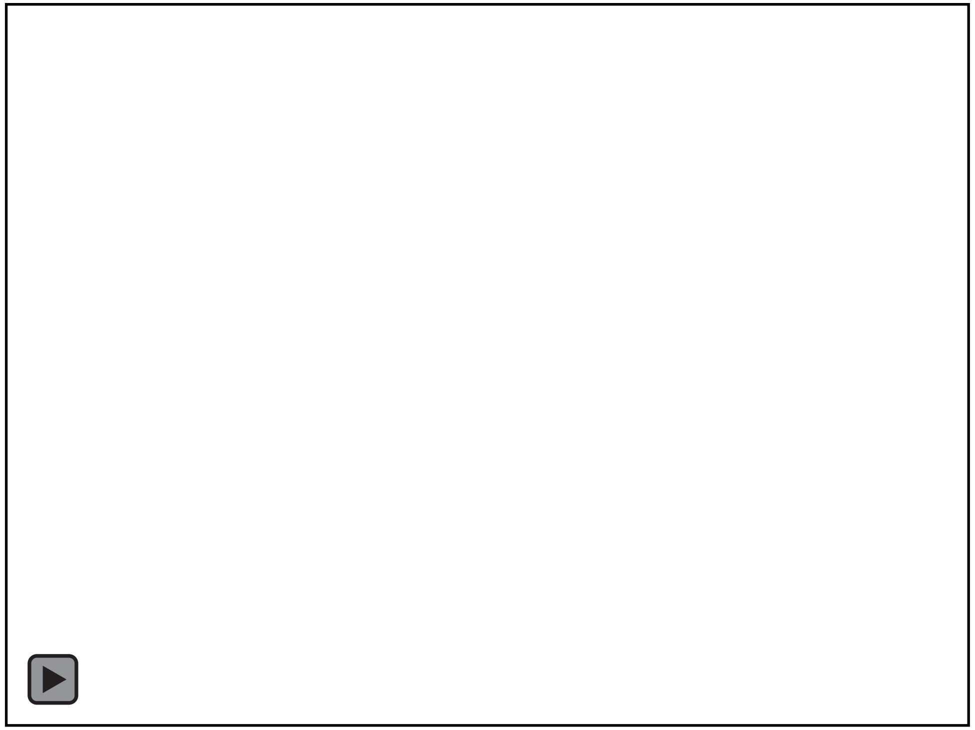

time-lapse from Internet photos

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

Computational photography

optics to focus light on an image plane

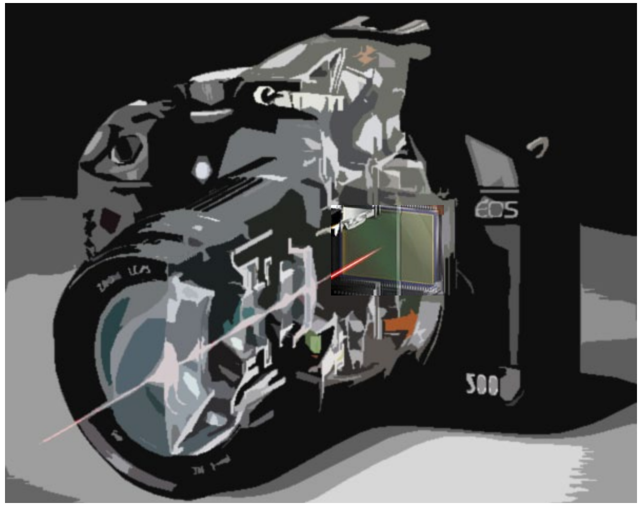

digital sensor to capture focused light (electrical process)

arbitrary computation between sensor and image

Computational photography

generalized optics
between scene and sensor

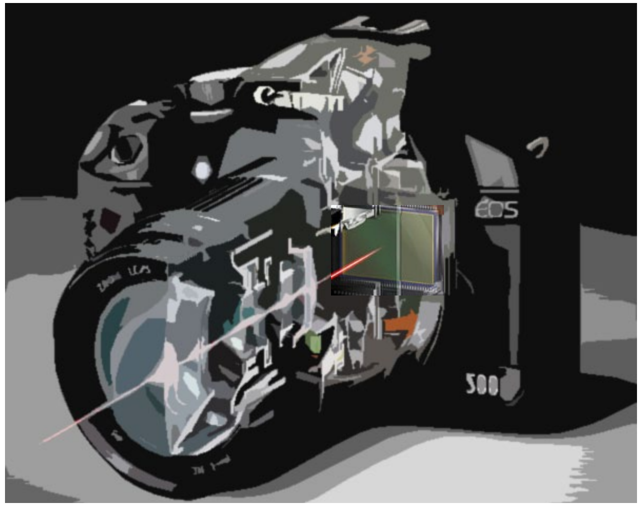

digital sensor to capture focused
light (electrical process)

arbitrary computation
between sensor and image

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

Capture more than 2D images

Lightfield cameras for plenoptic imaging

post-capture refocusing

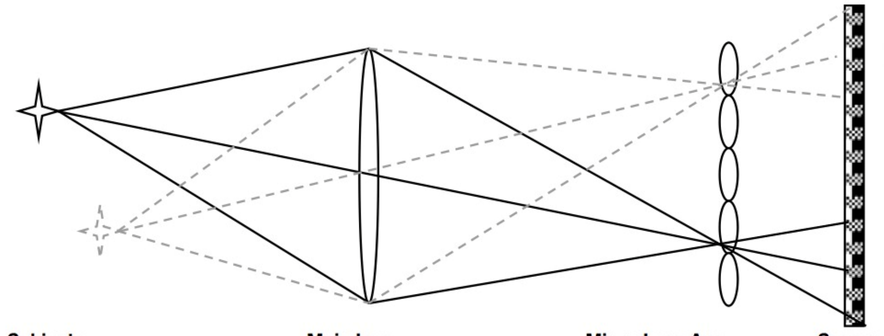

One of your homeworks!

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

Capture more than 2D images

Lightfield cameras for plenoptic imaging

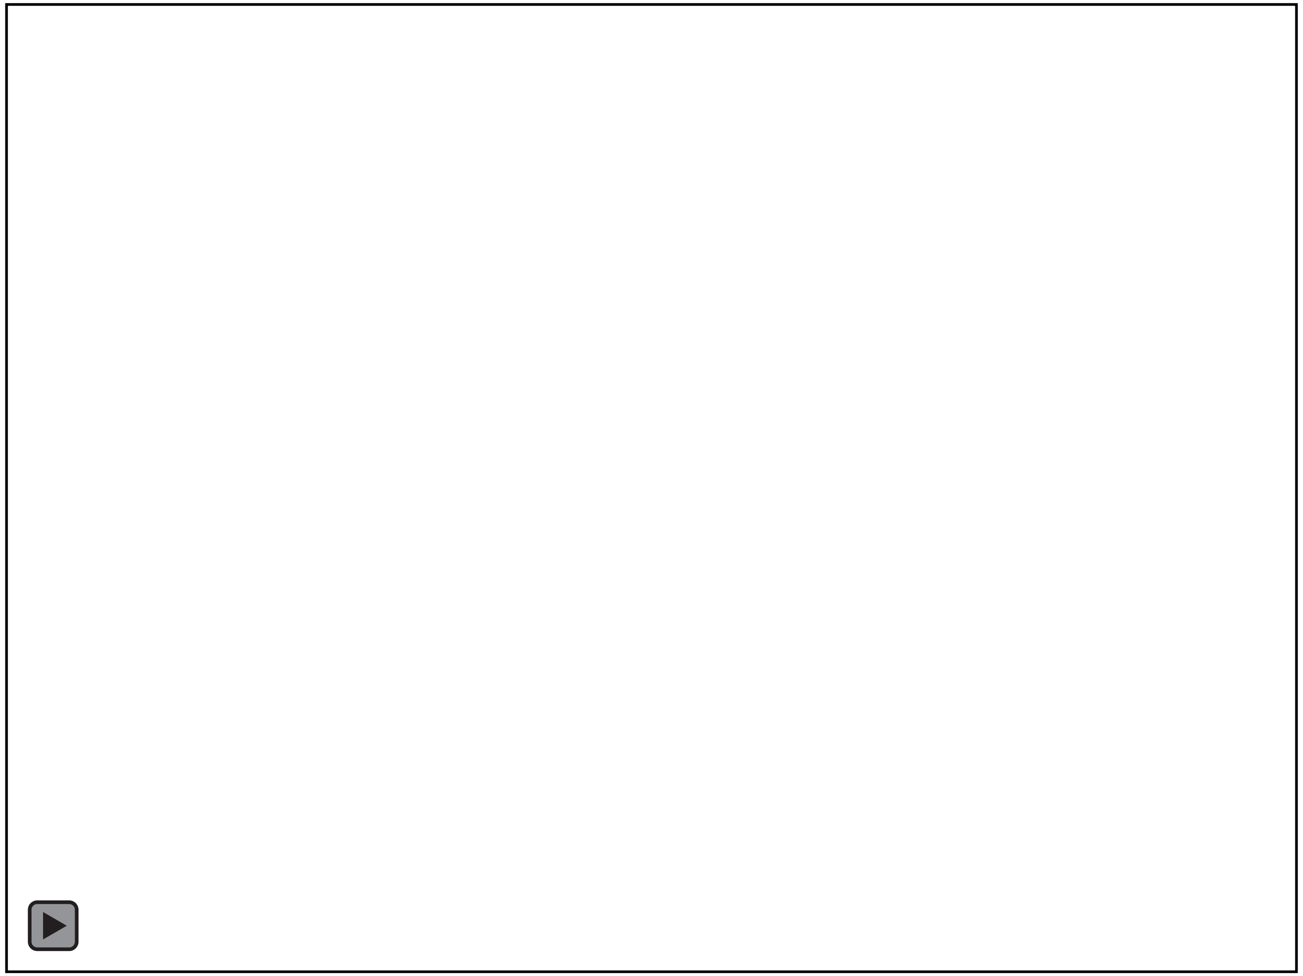

[Ng et al., SIGGRAPH 2005] [Lytro Inc.]

Measure 3D from a single 2D image

Coded aperture for single-image depth and refocusing

conventional vs
coded lens

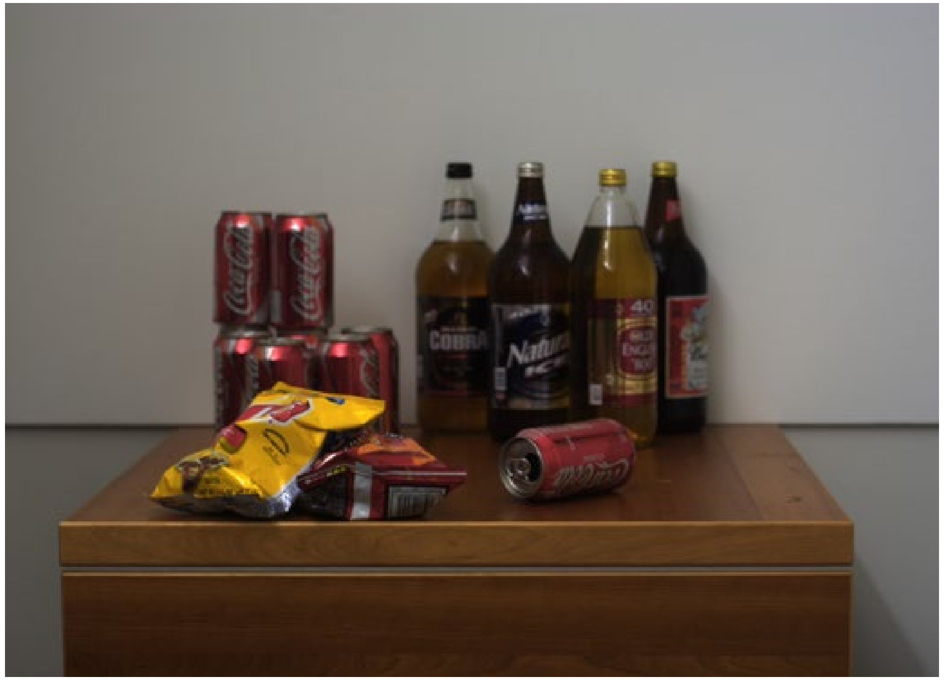

input image

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

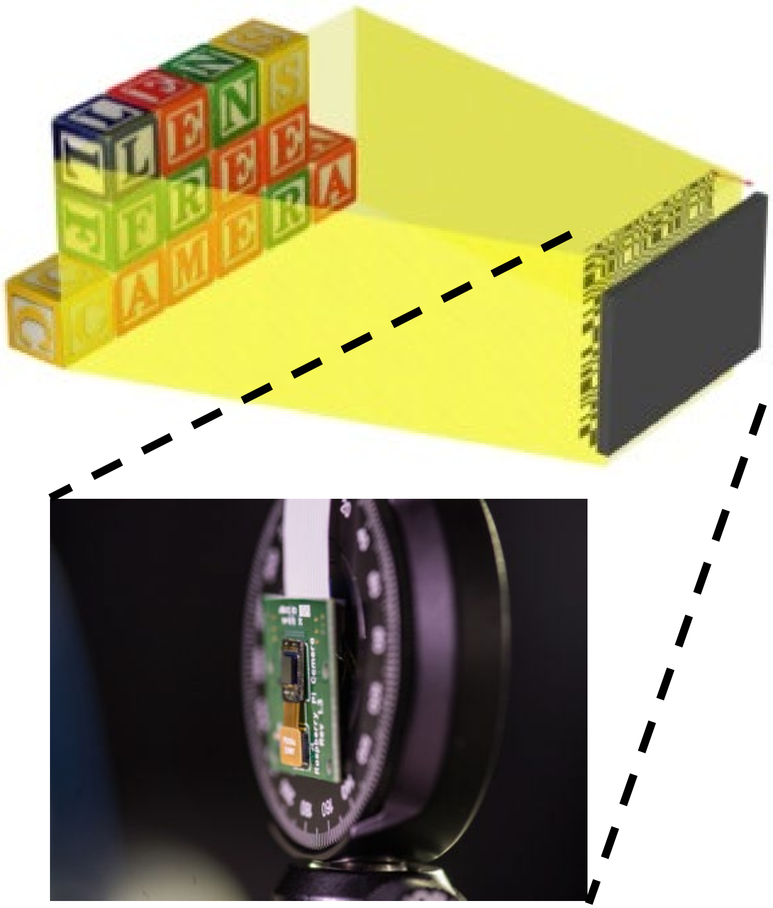

prototype

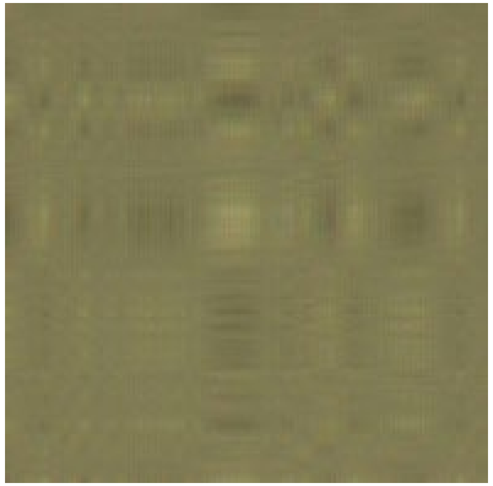

sensor measurements

reconstructed image

[Asif et al. 2015]

Computational photography

generalized optics between scene and sensor

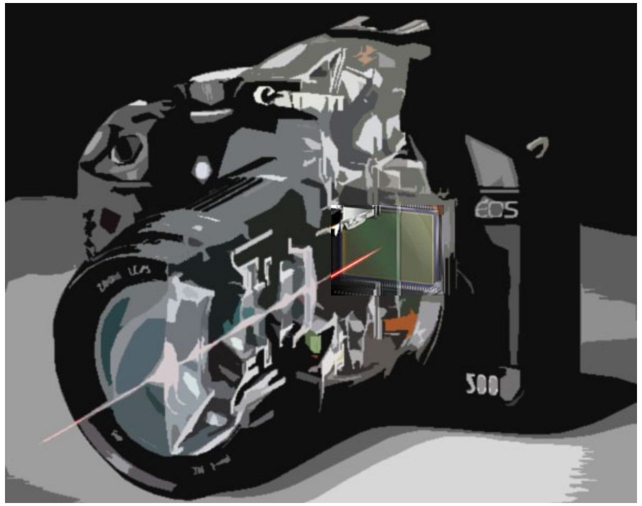

digital sensor to capture focused light (electrical process)

arbitrary computation between sensor and image

Computational photography

generalized optics
between scene and sensor

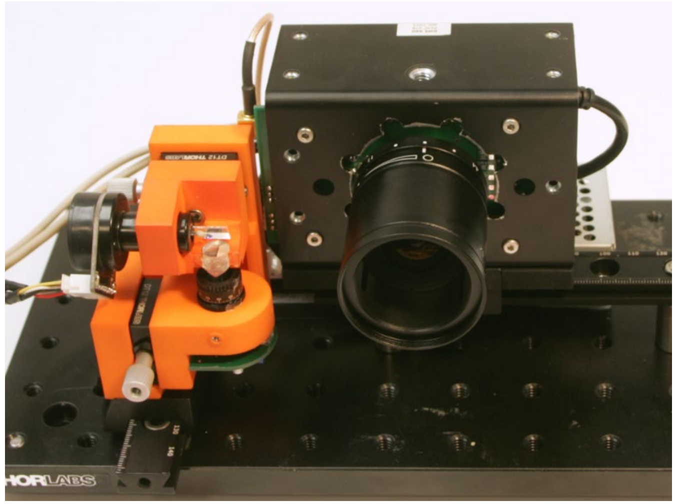

unconventional sensing
and illumination

arbitrary computation
between sensor and image

Measure depth

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

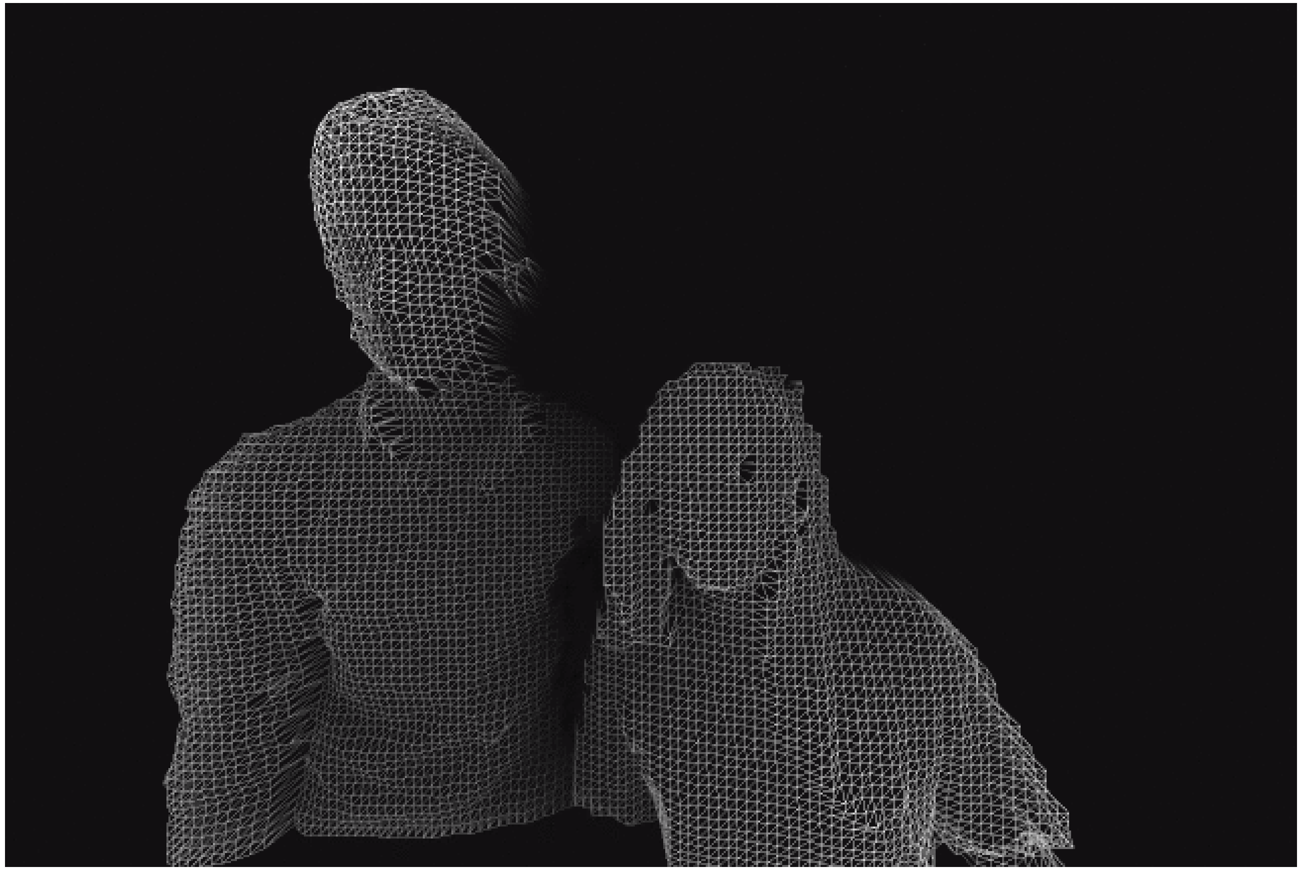

[Microsoft Inc.]

Measure light in flight

Streak camera for femtophotography

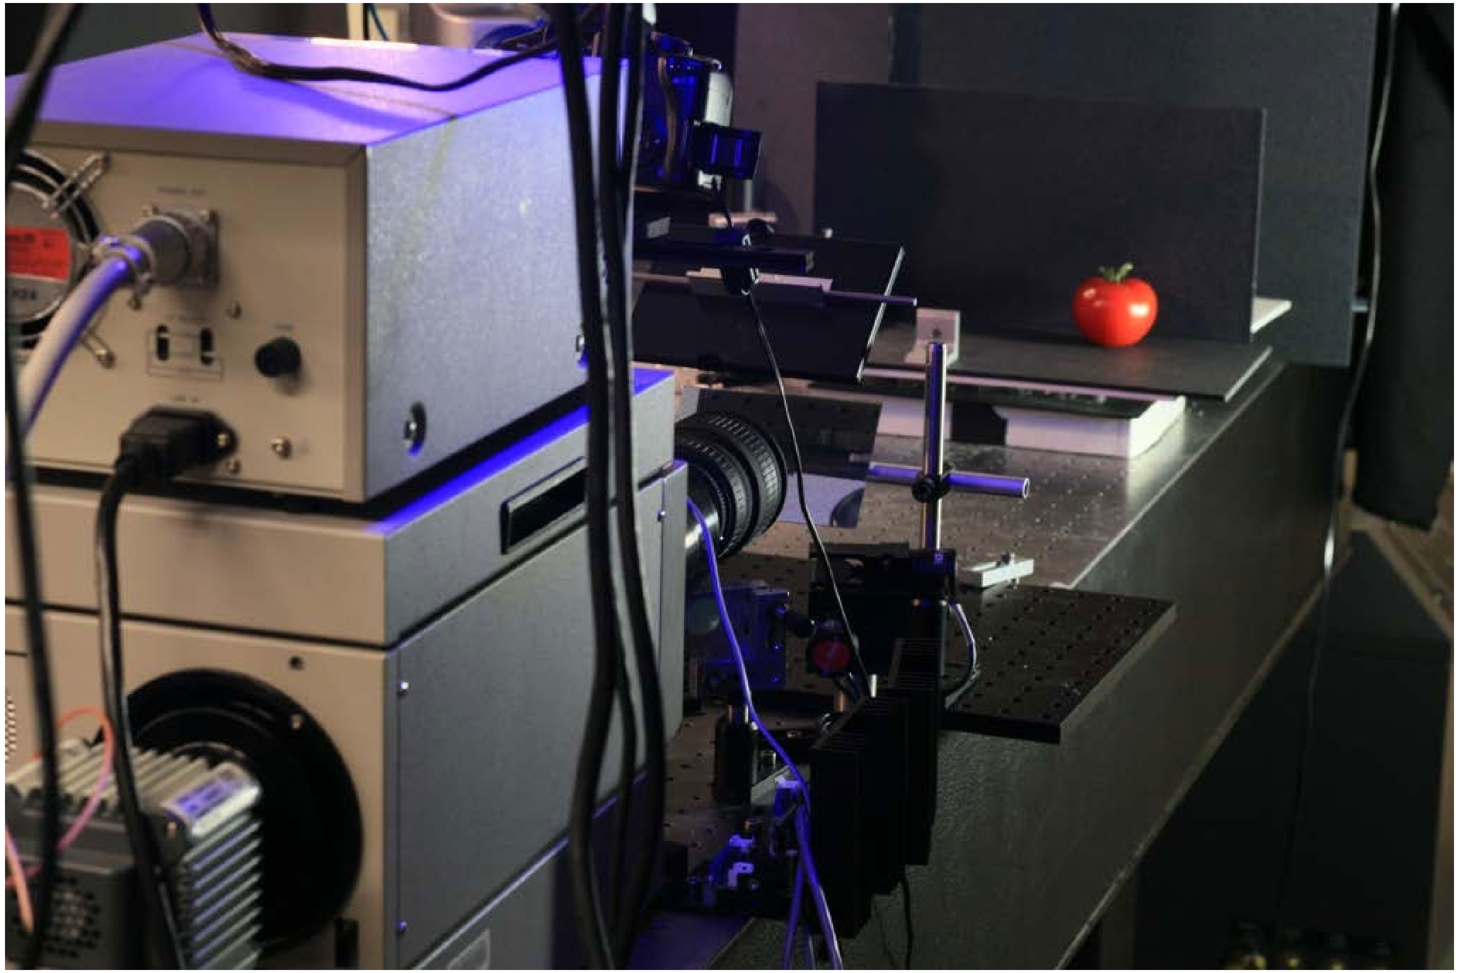

[Velten et al., SIGGRAPH 2013]

Measure light in flight

Streak camera for femtophotography

[Velten et al., SIGGRAPH 2013]

Measure photons selectively

Structured light for epipolar imaging

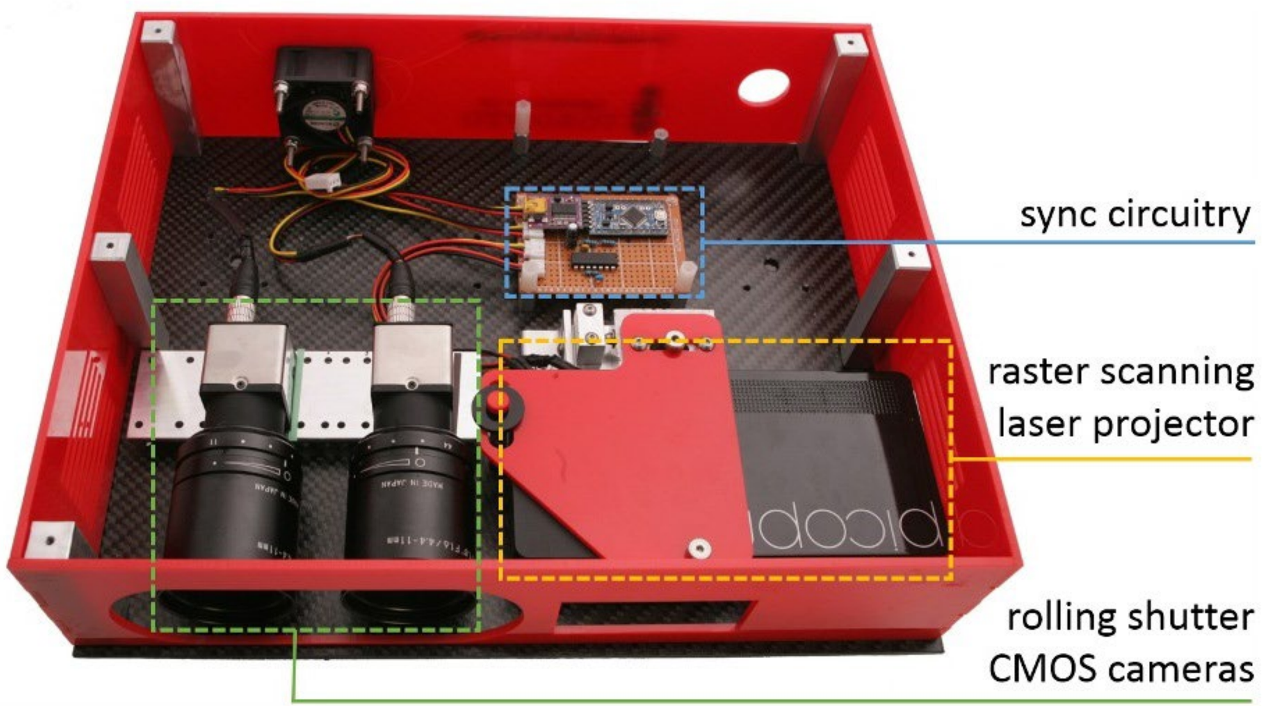

[O'Toole et al., SIGGRAPH 2015]

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

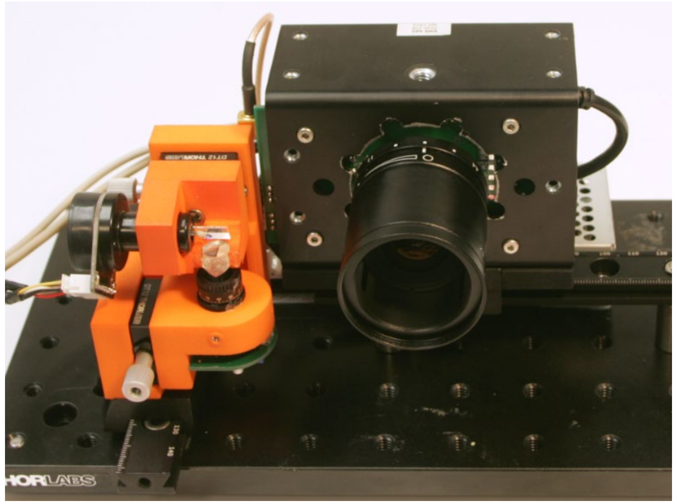

unconventional sensing
and illumination

arbitrary computation
between sensor and image

Computational photography

generalized optics
between scene and sensor

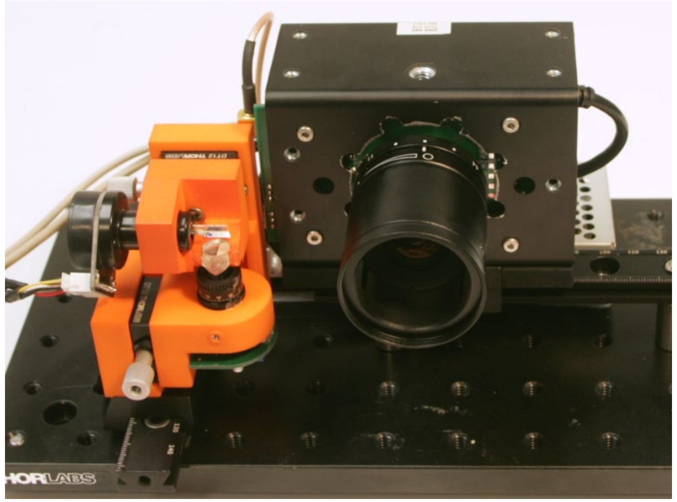

unconventional sensing
and illumination

arbitrary computation
between sensor and image

joint design of optics, illumination, sensors, and computation

Putting it all together

Looking around corners

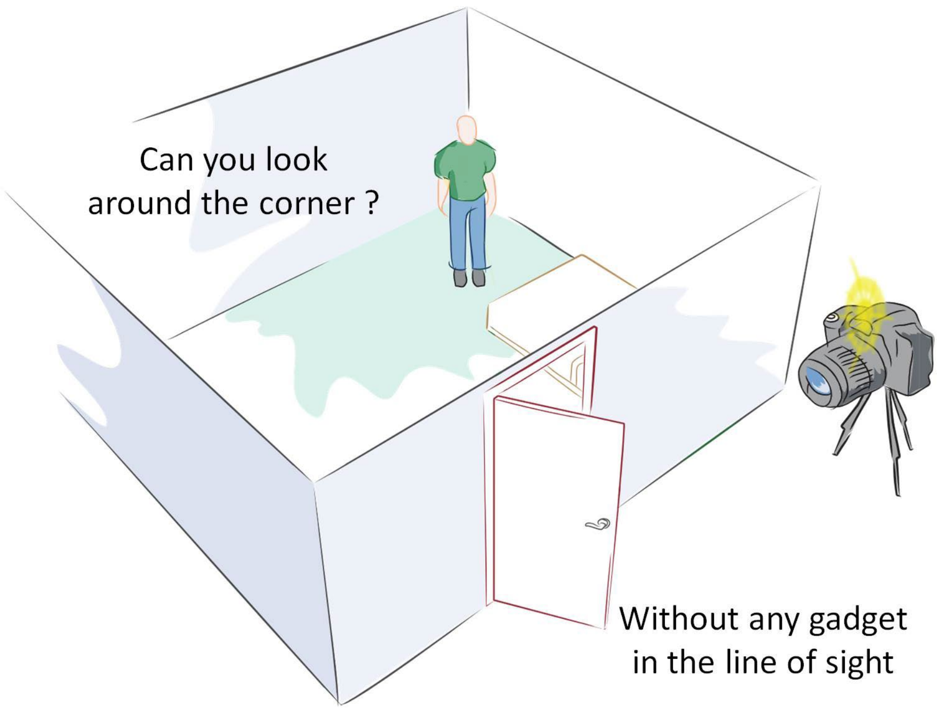

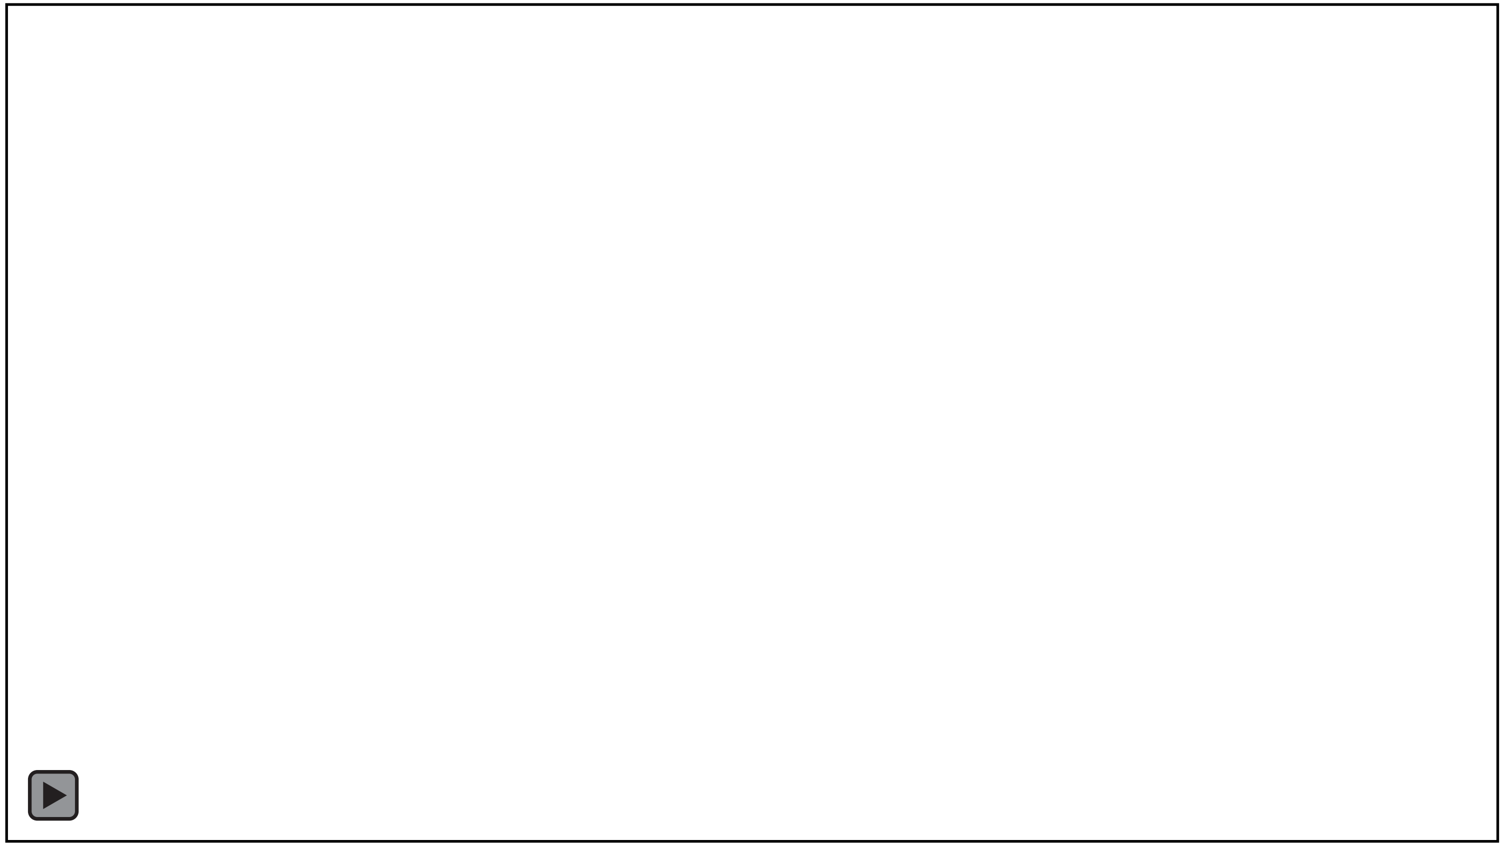

[MIT Media Lab, DARPA REVEAL]

Putting it all together

Looking through tissue

Opportunity

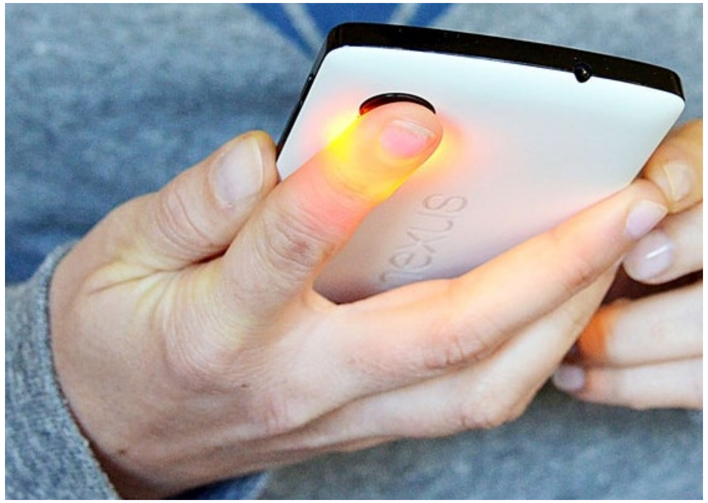

+ Light travels deep inside the body
+ It is non-ionizing (400-1100nm)
+ Cheap to produce and control

Scattering Barrier

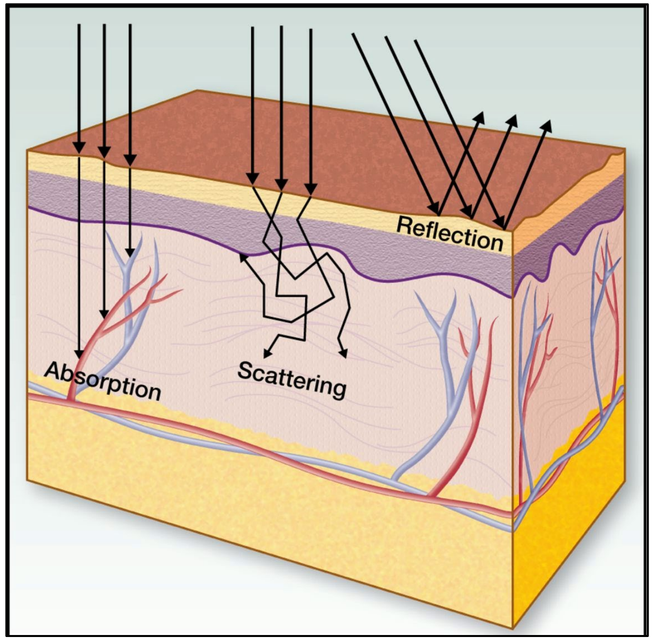

- Most pass-through photons are scattered
- Avg 10 scattering events per mm
- By 50mm, avg 500 scattering events !
- Large-scale inverse problem with low SNR

Practical imaging up to 50mm

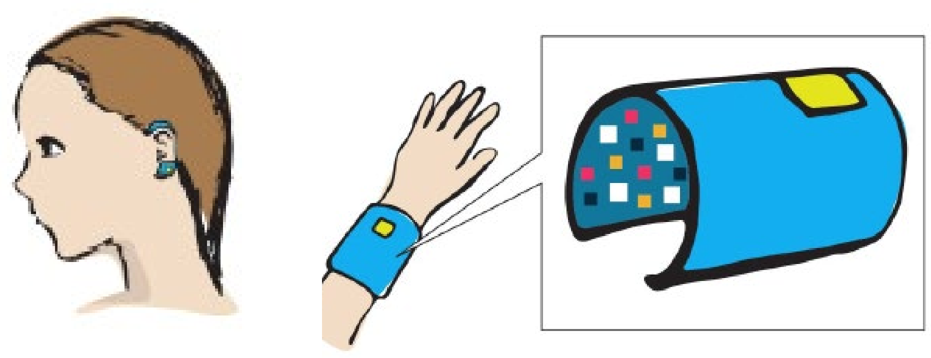

Wearables (1-10mm)

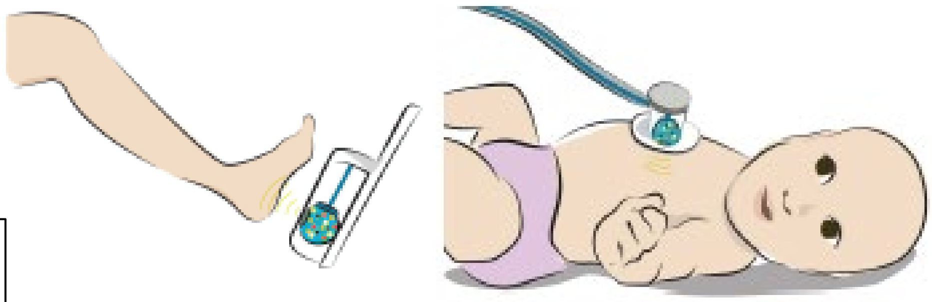

Non-invasive point of care devices (10-50mm)

[NSF Expedition]

Computational photography

generalized optics
between scene and sensor

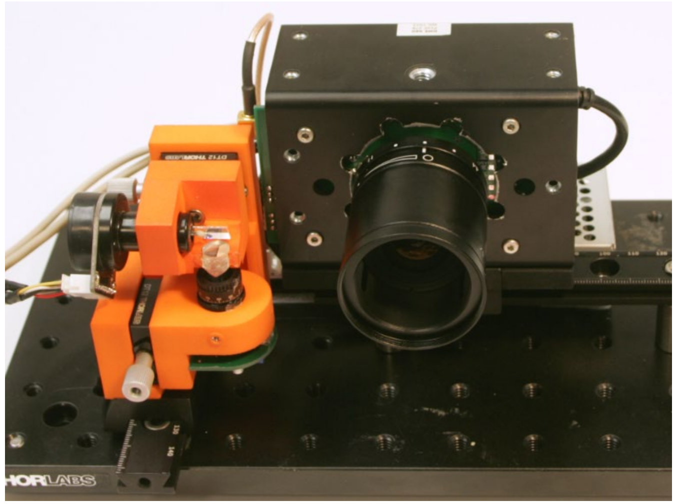

unconventional sensing
and illumination

arbitrary computation
between sensor and image

joint design of optics, illumination, sensors, and computation

Course fast-forward and logistics

Course fast-forward

Tentative syllabus on the course website:

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
- no required textbook, see suggestions on the website, and references at the end of lectures

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 fusion:

- high-dynamic-range imaging
- bilateral filtering
- edge-aware filtering
- gradient-domain image processing
- flash/no-flash photography
- high-performance image processing

[Banerjee et al., SIGGRAPH 2014]

Topics to be covered

Types of cameras:

- geometric camera models
- lightfield 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

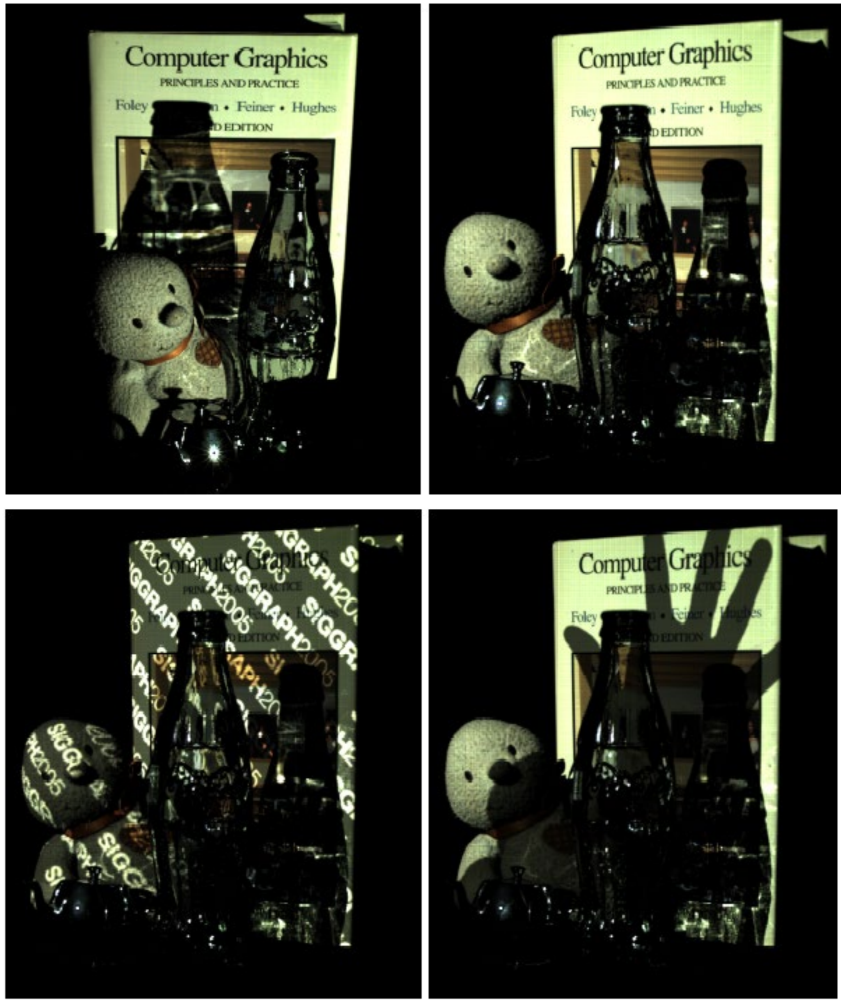

[Sen et al., SIGGRAPH 2005]

Course online platforms

- Course website for all resources:

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

- Canvas and Gradescope for homework submissions and grading.
- Slack for announcements and discussion.
- Links available on the course website.

Please take the start-of-semester survey!

- Posted on Slack as well:

  https://docs.google.com/forms/d/e/1FAIpQLScY5gtWcuSZ4X1n7MUtT5DPto921t5A80jFlB1mmq43oetVrA/viewform

- We use the survey to:
  - Get a better idea of students' background.
  - Get a better idea of equipment needs.
  - Decide on day and time of office hours.

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.
- 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.
- Multi-view stereo.
- Radiance and radiometry.
- Lambertian, diffuse, and specular reflectance.
- n-dot-l lighting.
- Thin lens, prime lens, and zoom lens.
- Demosaicing.
- Refraction and diffraction.

Evaluation

- Six two-week homework assignments (75%):
  - Two parts, programming in **Python**, capturing your own photographs.
  - Generous extra credit components to help you catch up on missed credit.
  - Released and due every second Friday.
  - Five free late days, you can use them as you want. Penalty 10%/day after that.
  - No assignment may be submitted more than three days after deadline.
  - See course website for schedule and submission guidelines.
  - **Strict submission deadlines. No extra late days except for medical reasons.**
- Final project (20%):
  - See final project page on course website for detailed logistics (updated this year).
  - Start thinking about final project early, especially if you require imaging equipment.
  - **No exam or final presentations. Final project videos will appear on YouTube.**
- Class participation (5%):
  - Be around during lectures.
  - Participate in Slack discussions.
  - Ask and answer questions.

Do I need a camera?

- You will need to take your own photographs for assignments 1-6 (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 if it has flash.
  - 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.

- We have 60 Nikon D3X00 kits (camera + lens + tripod) for students.
  - If you have your own camera, please use that!
  - Tutorial available on course website.

- Sign up for a camera (distributed in the second week of classes):
  https://docs.google.com/spreadsheets/d/1aWtAWuxstwshd5eq7BIW-idBc4RJLVh1OMoL_8FbucY

Tips

- Capturing your own photos is *very* time consuming and tedious. Lots of handiwork and trial and error.
- Do not spend all your late days on assignment 1. Assignments 2, 5, 6 take up a *lot* more time.
- Plan how to use late days to accommodate CVPR or other paper deadlines, job interviews, friends' weddings, other coursework, research, etc.
- Start early! Do at least a cursory read of each homework assignment the day it is released, and plan programming and data capture components.
- Uploading solutions with many images can be slow. Plan to submit early to avoid wasting late days.

On average, how did the homework assignments in this course compare to homework assignments in other courses you have taken in terms of difficulty?

Copy

6 responses

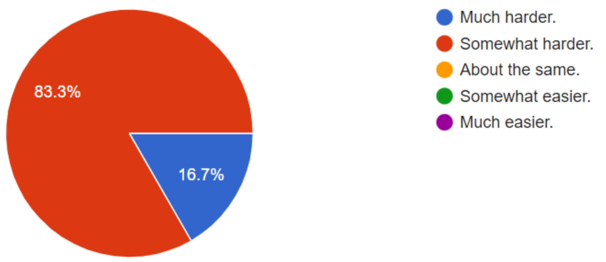

Explain your previous answer.

3 responses

In the parts where we needed to capture our own data, there were more practical aspects to consider such as finding the right materials or finding a dark room. It can be very time-consuming since they required trial-and-error.

Definitely takes more time than assignments in the other course as we have both theoretical and hands-on task for each assignment. There are some challenges to generate my own data set (e.g., at the first place I did not know whether my light condition is enough, and afterward I realize that I need more light, so I took pictures again etc). However, the assignment instructions are clear so it is easy to follow step by step.

I think for me, a lot of the math was already hard to fully understand. As an undergrad, we typically don't read research papers in our classes, so that was also new to me. The research papers often detailed harder math and theory than was discussed in class, which made it difficult to figure out what to do. Most of the homework assignments also took a lot of trial and error to get the code completely right, and I often found it near impossible to debug if I had a minor conceptual understanding, since my code would "look" correct to me, and the visual outputs would usually not provide any clue as to where the bug might be. There was also a lot of documentation reading and figuring out what was possible with numpy / how to actually do tasks properly using it. I think this more systems-related aspect combined with the need to really understand the math/theory behind the assignments made the assignments harder than other class assignments, which would often either focus on systems-related challenges or theory/algorithm-related

Tips

- Capturing your own photos is *very* time consuming and tedious. Lots of handiwork and trial and error.
- Do not spend all your late days on assignment 1. Assignments 2, 5, 6 take up a *lot* more time.
- Plan how to use late days to accommodate CVPR or other paper deadlines, job interviews, friends' weddings, other coursework, research, etc.
- Start early! Do at least a cursory read of each homework assignment the day it is released, and plan programming and data capture components.
- Uploading solutions with many images can be slow. Plan to submit early to avoid wasting late days.

On average, how did the homework assignments in this course compare to homework assignments in other courses you have taken in terms of how much you enjoyed them?

Copy

6 responses

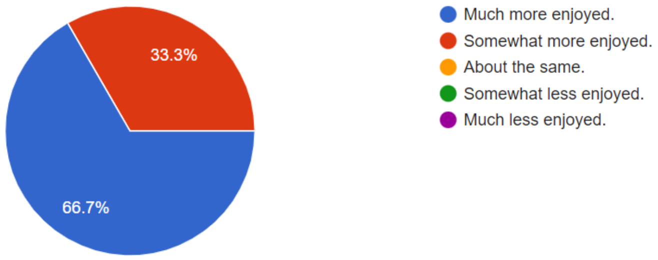

What are the strongest features of this course?

6 responses

- Homeworks, lectures
- Homework assignments that really forced the students to know how to implement the core ideas in class.
- Getting to do advanced image processing on images we took ourselves is so amazing
- Very interesting
- both theoretical contents and hands-on implementation
- The rewarding/interactive/photography homeworks, and the lectures

Don't make the homeworks easier!

Final project competition

- At the end of the semester, we will ask other computational photography faculty at CMU (Srinivasa Narasimhan, Matthew O’Toole, Aswin Sankaranarayanan, Jun-Yan Zhu) to watch final project videos and vote on the two best final projects.

- The two winning students will receive a **free DSLR camera kit** (same as the one provided for homework).

- Previous years’ projects for inspiration: Fall 2020, Fall 2021, Fall 2022, Fall 2023.

Final project competition

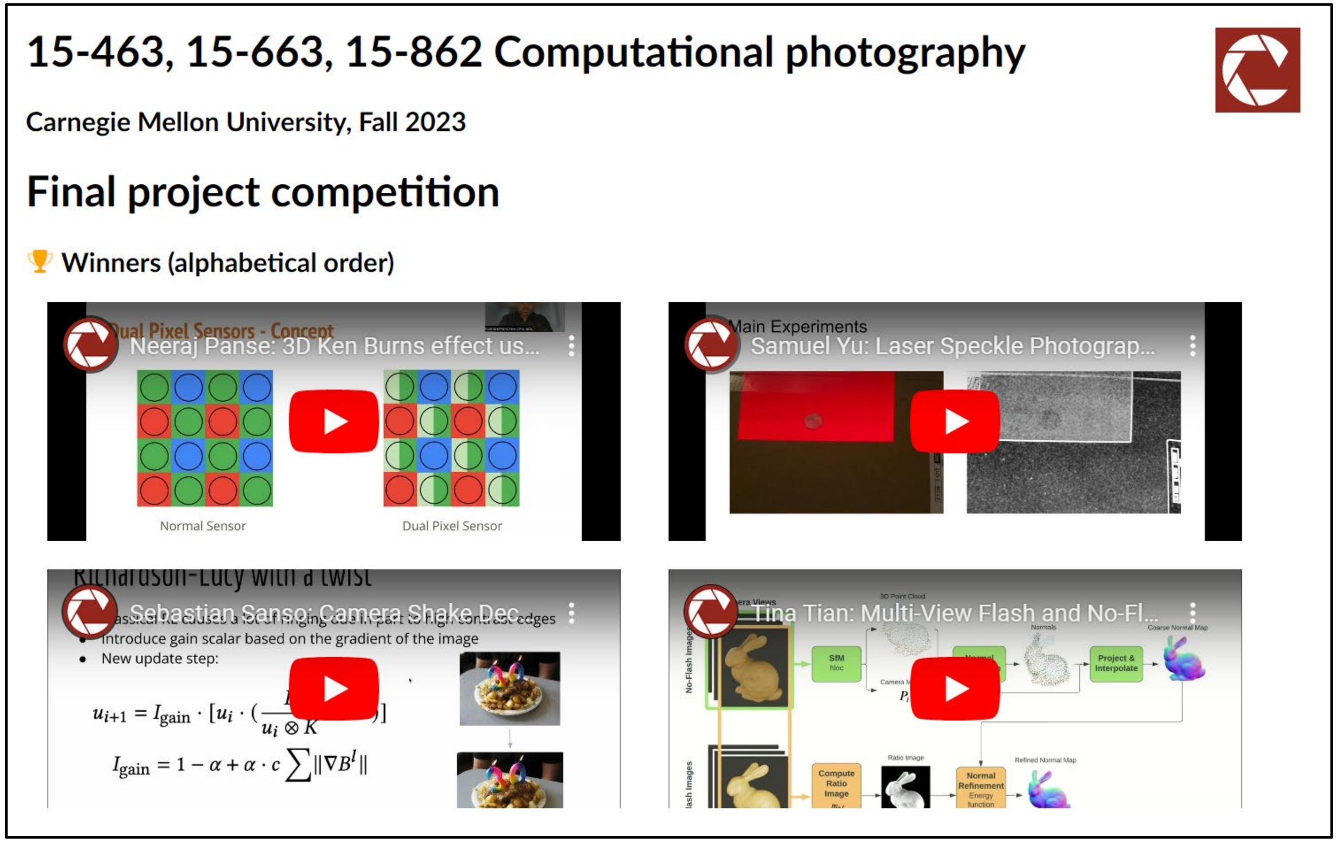

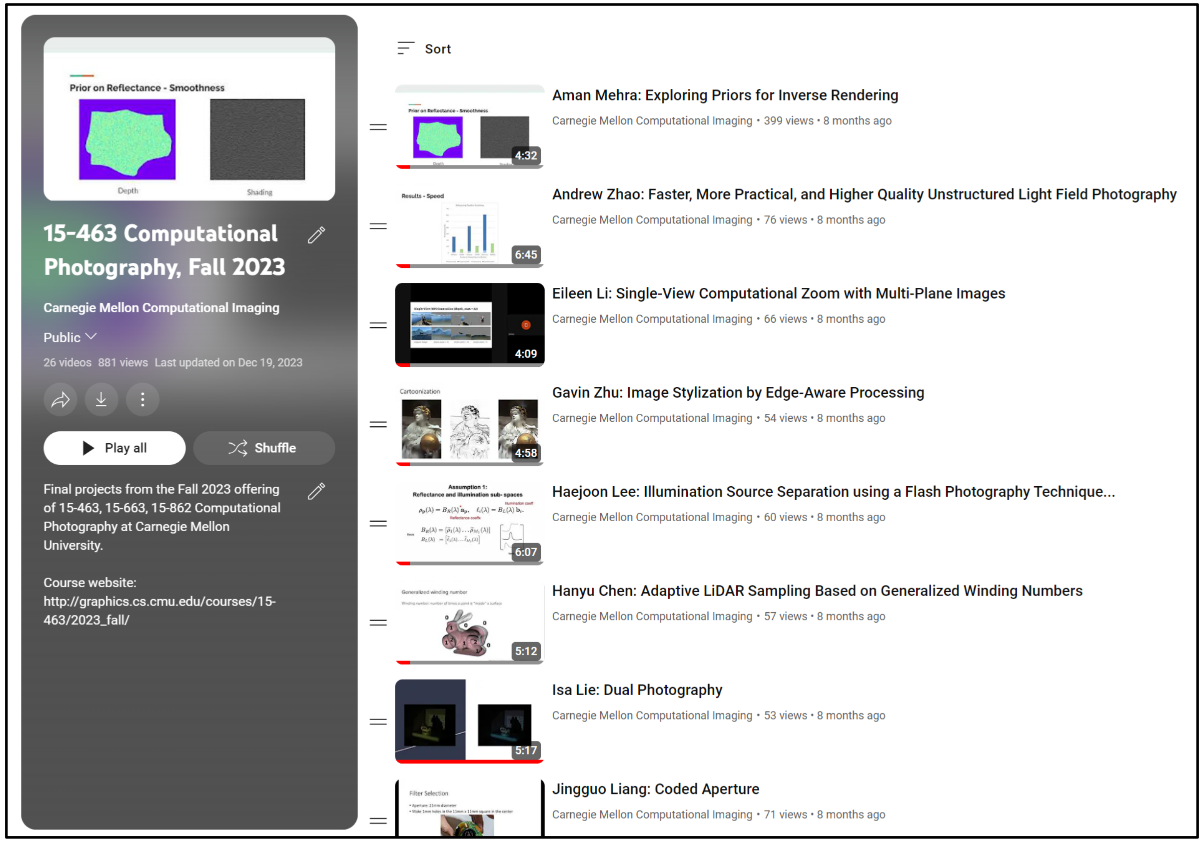

All final project videos will be available on course website and YouTube.

Homework assignment competitions

- After each homework assignment, the teaching staff will select one of the submissions that produced the most compelling result in the “capture your own images” part.

- The winning student will receive a **free camera-related gift.** Tentative list:
  - Assignment 1: Thingify pinhole "lens“.
  - Assignment 2: Colorchecker passport.
  - Assignment 3: flash.
  - Assignment 4: Lytro camera.
  - Assignment 5: telecentric lens.
  - Assignment 6: pocket projector.

- Previous years’ results for inspiration: Fall 2022, Fall 2023.

Contact information and office hours

- Feel free to email us about administrative questions.
  - Please use [15463] in email title!
- Ask technical questions on Slack.
  - We won't answer technical questions through email.
- We will decide office hours using votes in the start-of-semester survey.
  - Office hours will be in person in Smith Hall (EDSH) Rm 236 (graphics lounge).
  - Feel free to email Yannis about additional office hours.
  - You can also just drop by Yannis' office (Smith Hall (EDSH) Rm 225).
  - You can also post on Slack for additional office hours.
  - Office hours for this week only: **Friday, noon – 1 pm ET.**

Interested in research?

- Check out the Carnegie Mellon computational imaging group:

  Website: https://imaging.cs.cmu.edu/

  YouTube: https://www.youtube.com/@cmu-computational-imaging

- Email Yannis if you want to be added to the imaging group mailing list and attend our weekly meetings (day and time for the semester TBD).

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

Carnegie Mellon computational imaging group

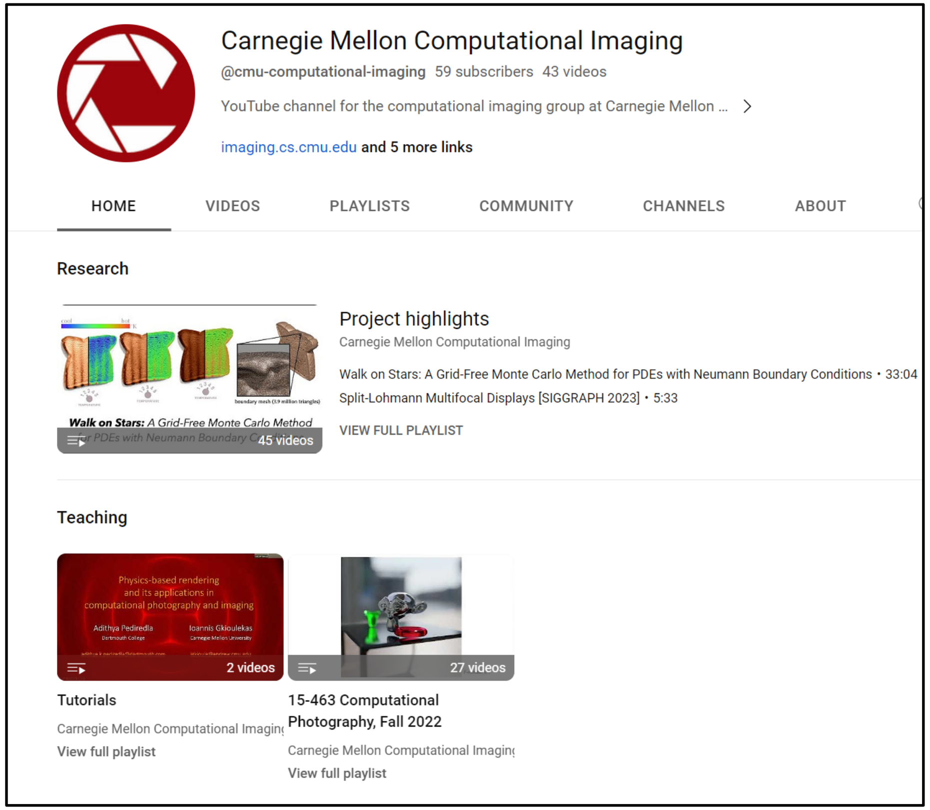

Resources on both research and teaching (CVPR tutorials, SIGGRAPH courses)

International Conference on Computational Photography YouTube channel

https://www.youtube.com/@iccp-conference

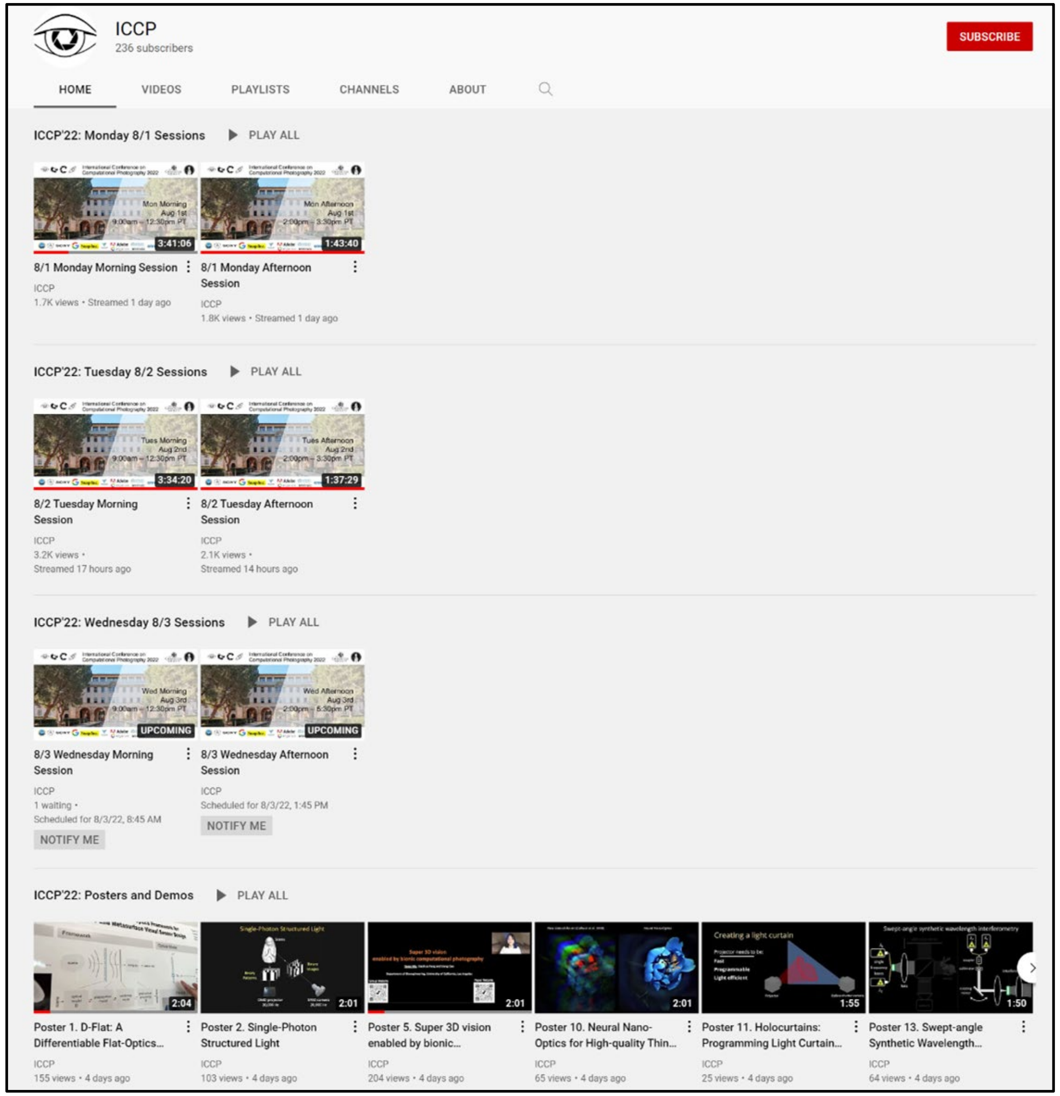

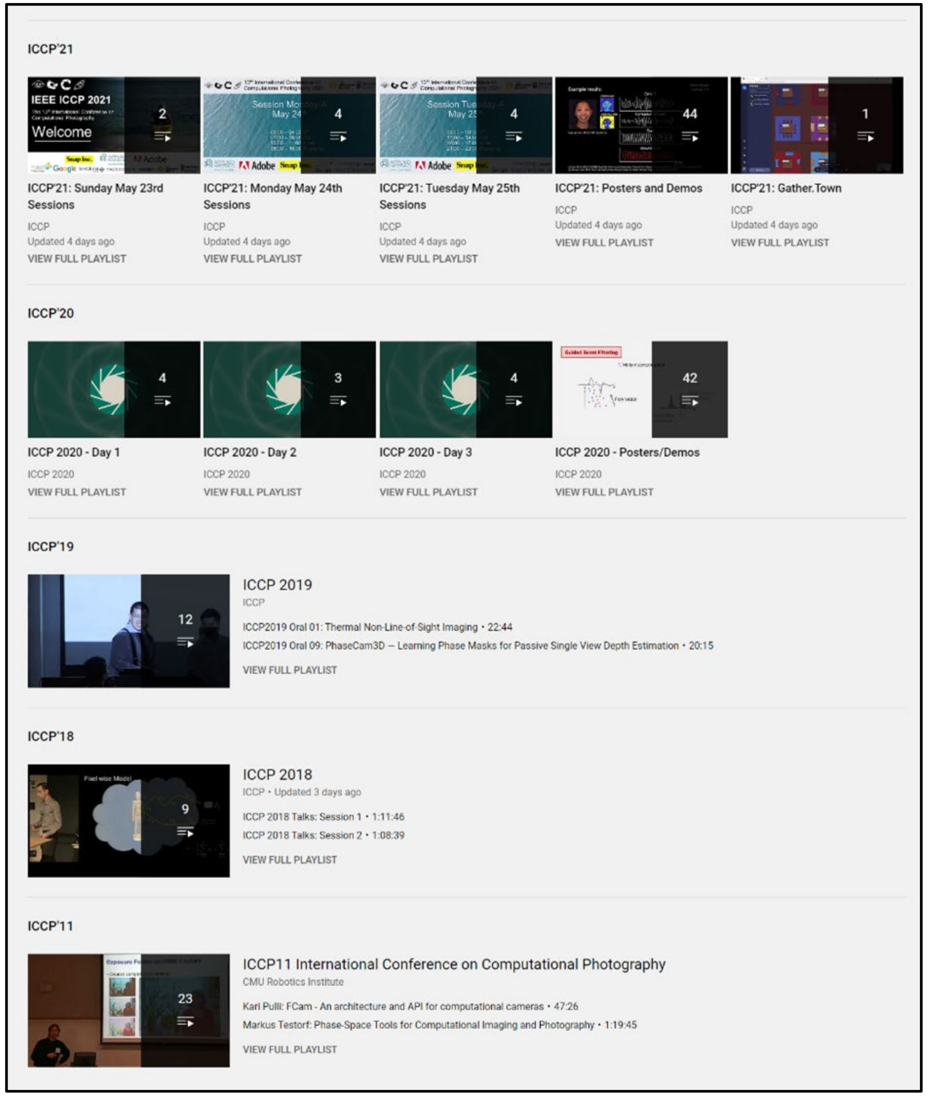

CMU has a strong presence at ICCP

ICCP 2011

ICCP 2018

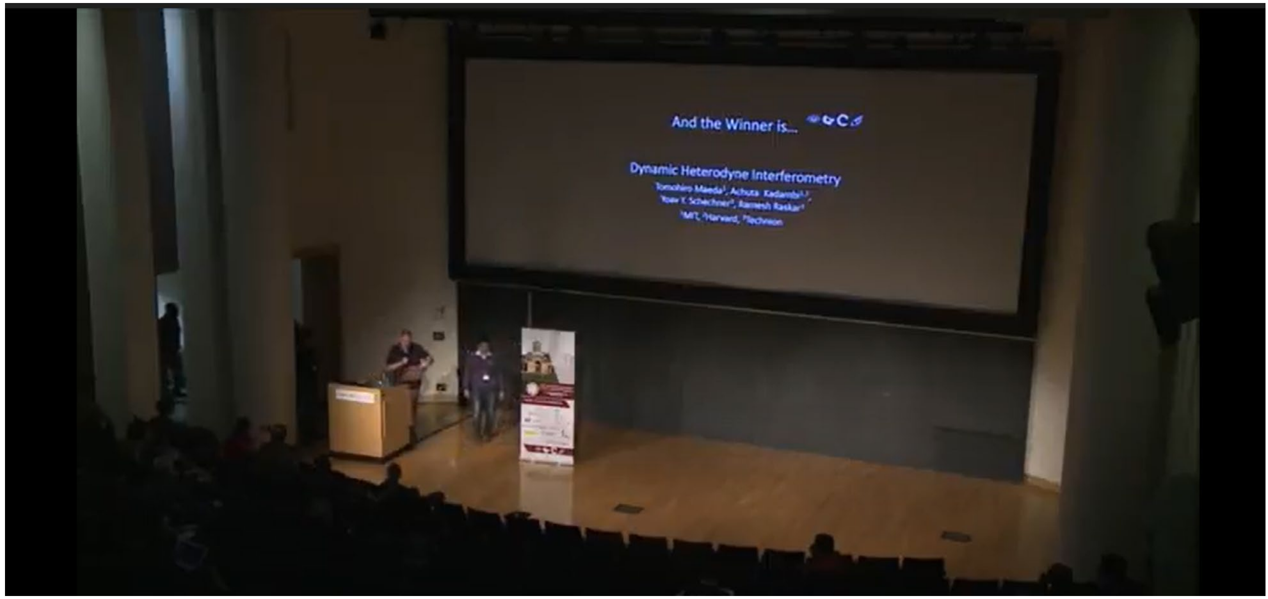

Please take the start-of-semester survey and sign up for a camera before the next lecture!

Survey link:

https://docs.google.com/forms/d/e/1FAIpQLScY5gtWcuSZ4X1n7MUtT5DPto921t5A80jFlB1mmq43oetVrA/viewform

Camera sign up:

https://docs.google.com/spreadsheets/d/1aWtAWuxstwshd5eq7BIW-idBc4RJLVh1OMoL_8FbucY/

Both links available on Slack.