Introduction

http://www.cs.cmu.edu/~16385/

16-385 Computer Vision
Spring 2020, Lecture 1

Overview of today's lecture

- Teaching staff introductions
- What is computer vision?
- 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 (2004-2009)

Harvard University (2009-2017)

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

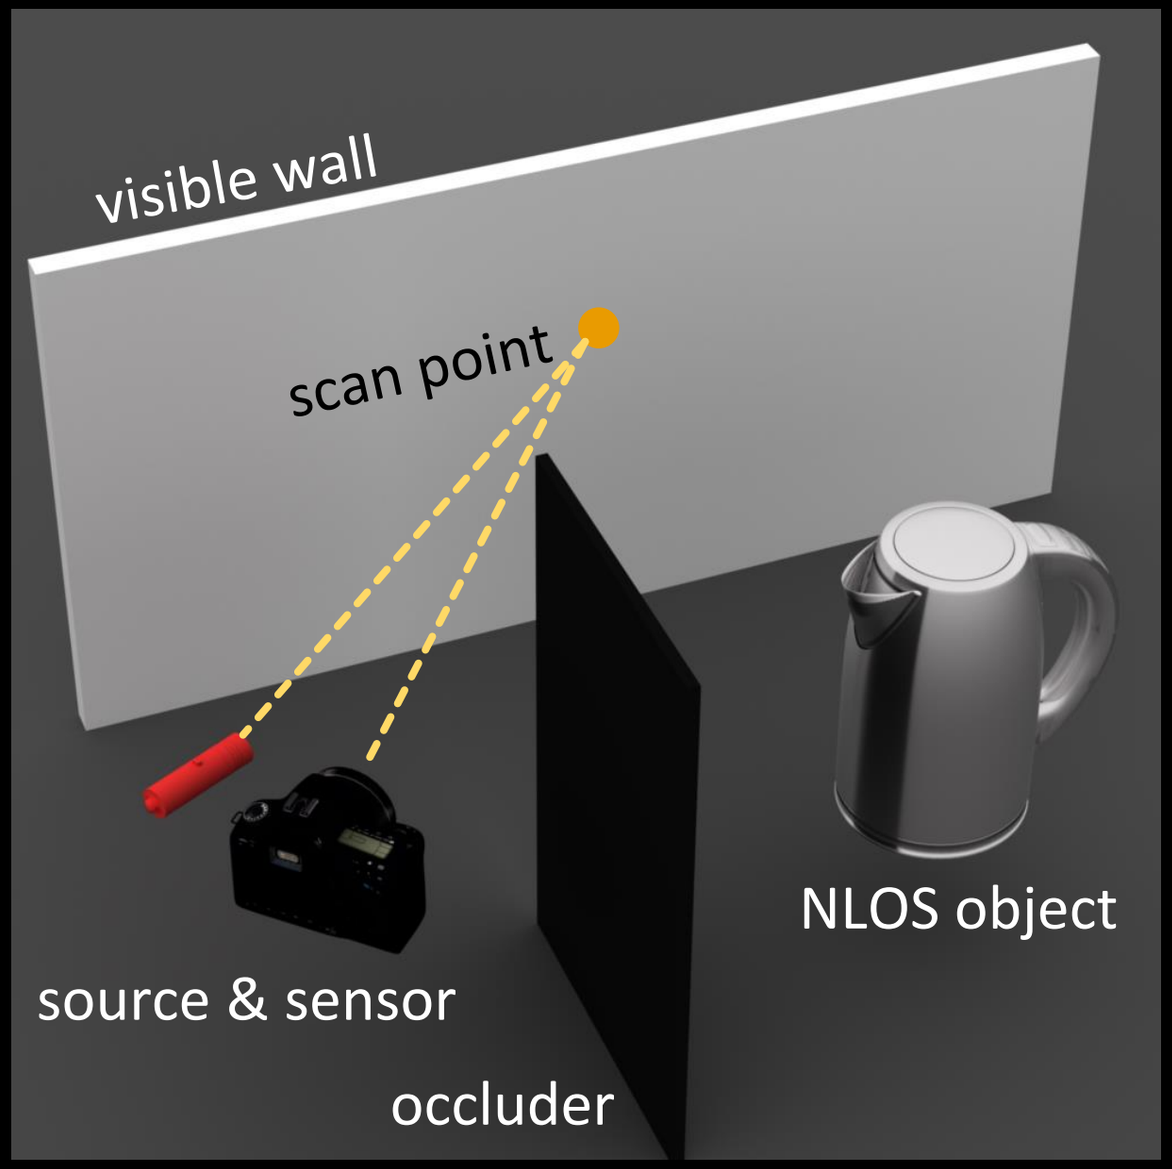

what a regular camera sees

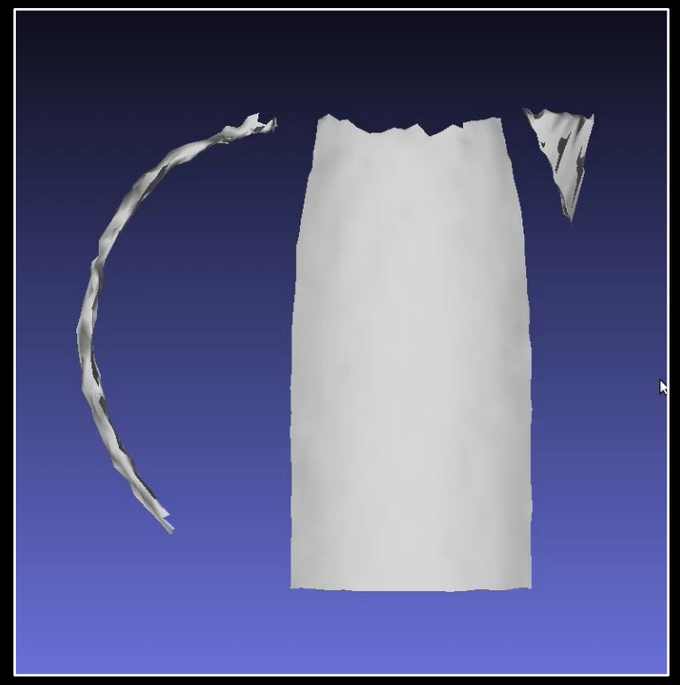

what we can reconstruct

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

Looking inside deep scattering objects

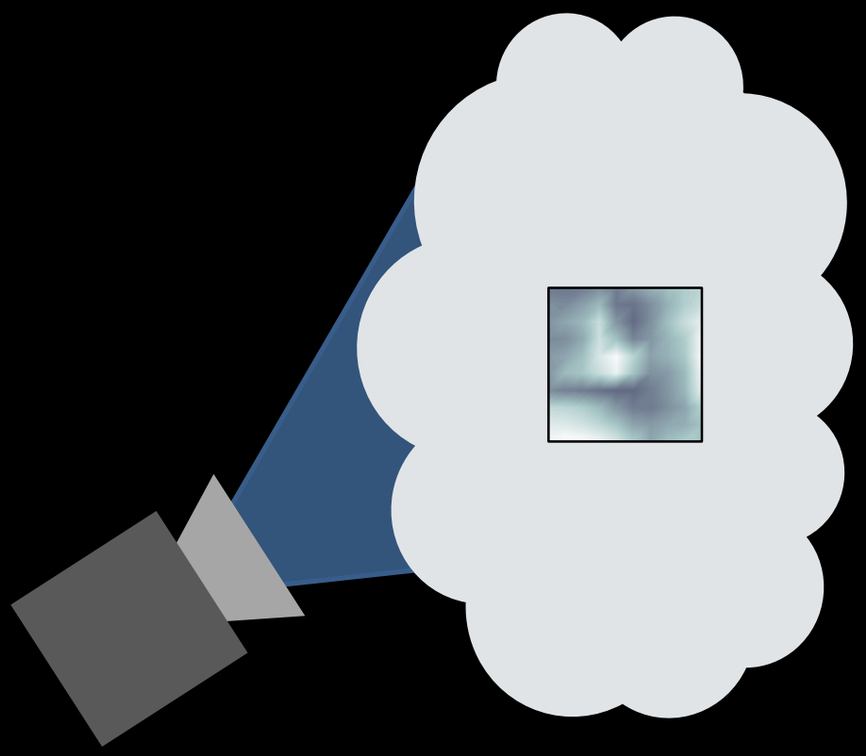

camera

thick smoke cloud

simulated camera measurements

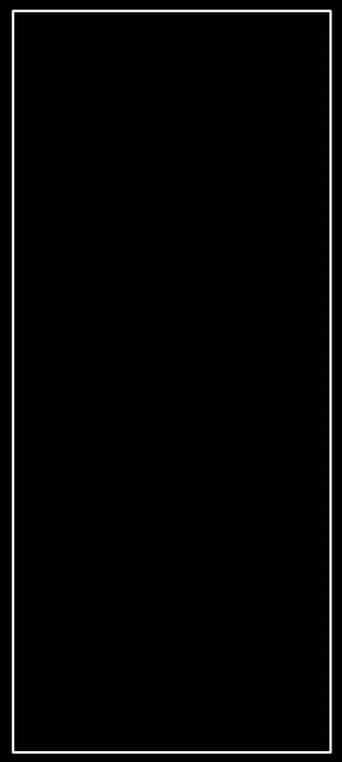

reconstructed cloud volume

slice through the cloud

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

Seeing light in flight

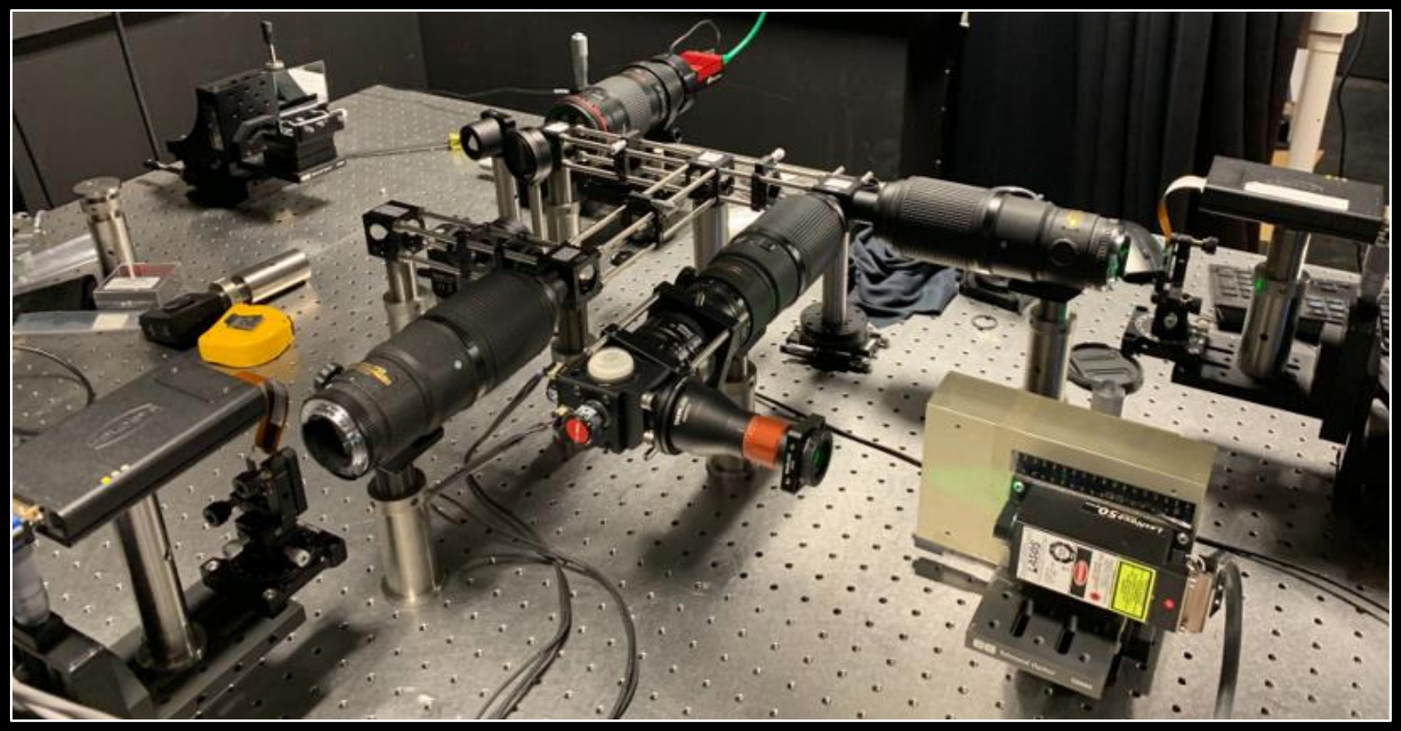

camera for capturing video at 10^{15} frames per second

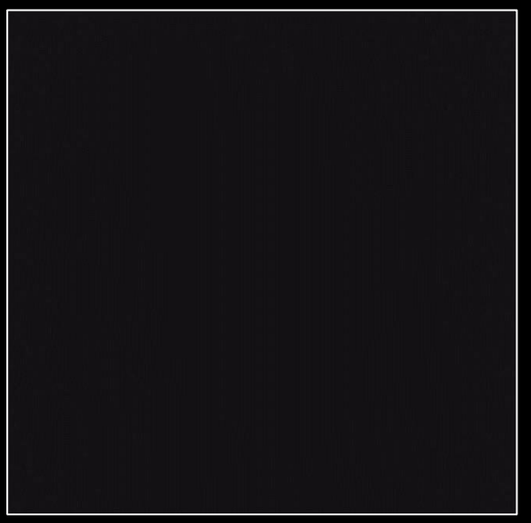

http://imaging.cs.cmu.edu/

Rendering wave effects

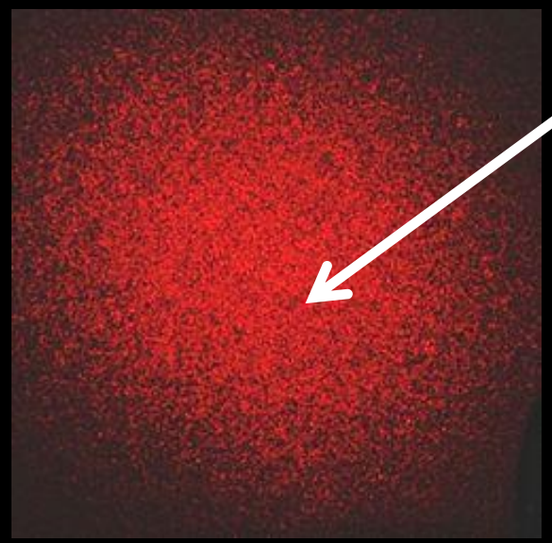

speckle: noise-like pattern

what real laser images look like

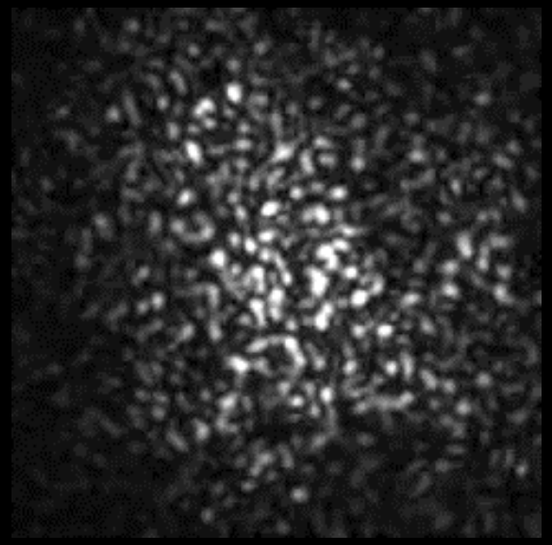

what real laser videos look like

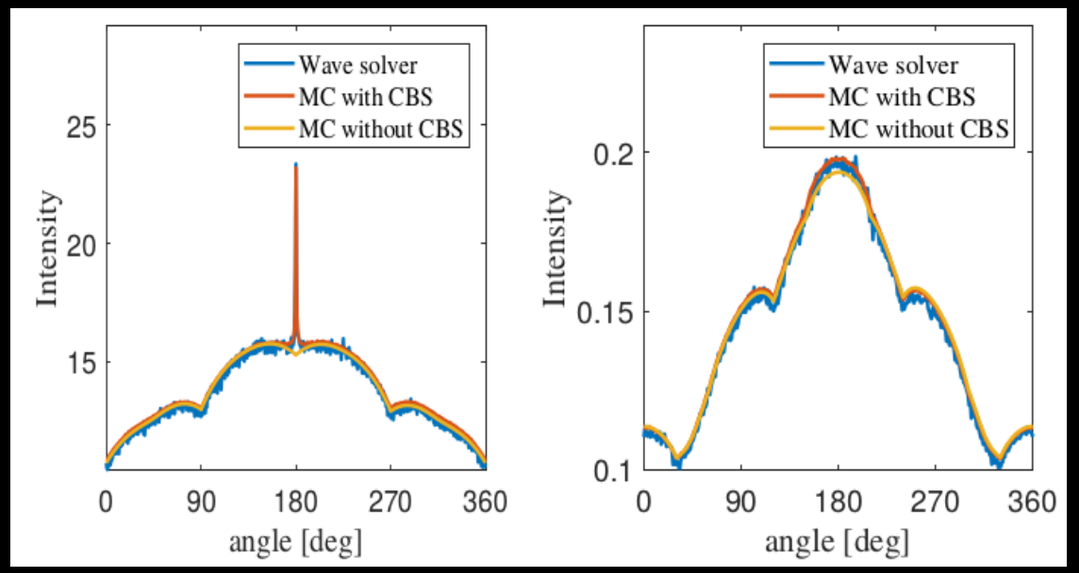

match wave equation solvers, **<u>10^5x faster</u>**

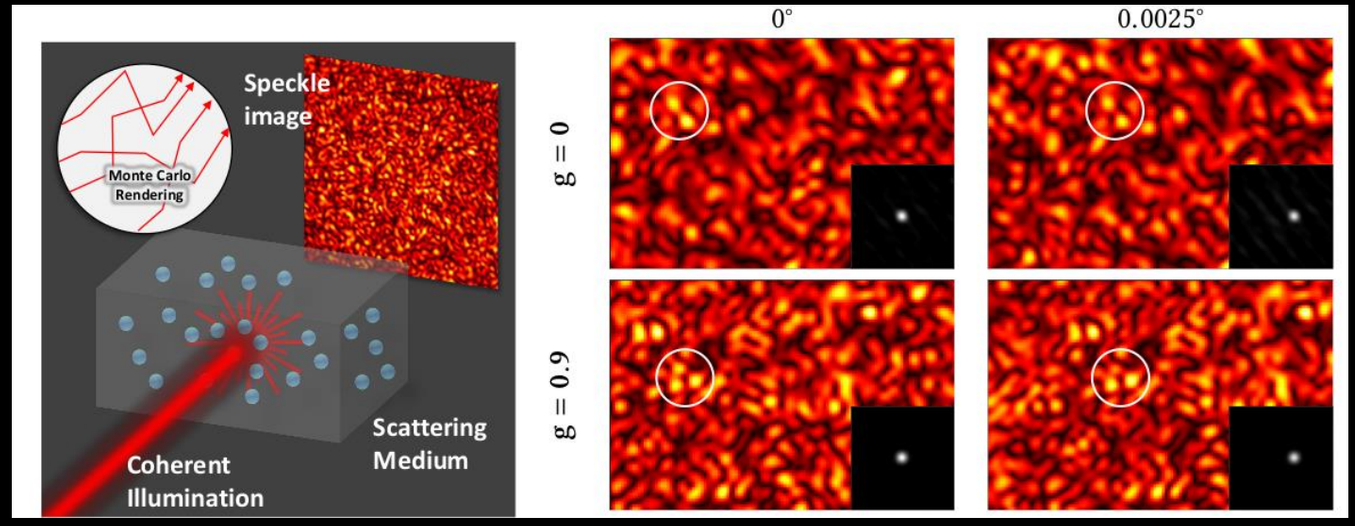

reproduce physical effects like memory effect

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

Differentiable rendering

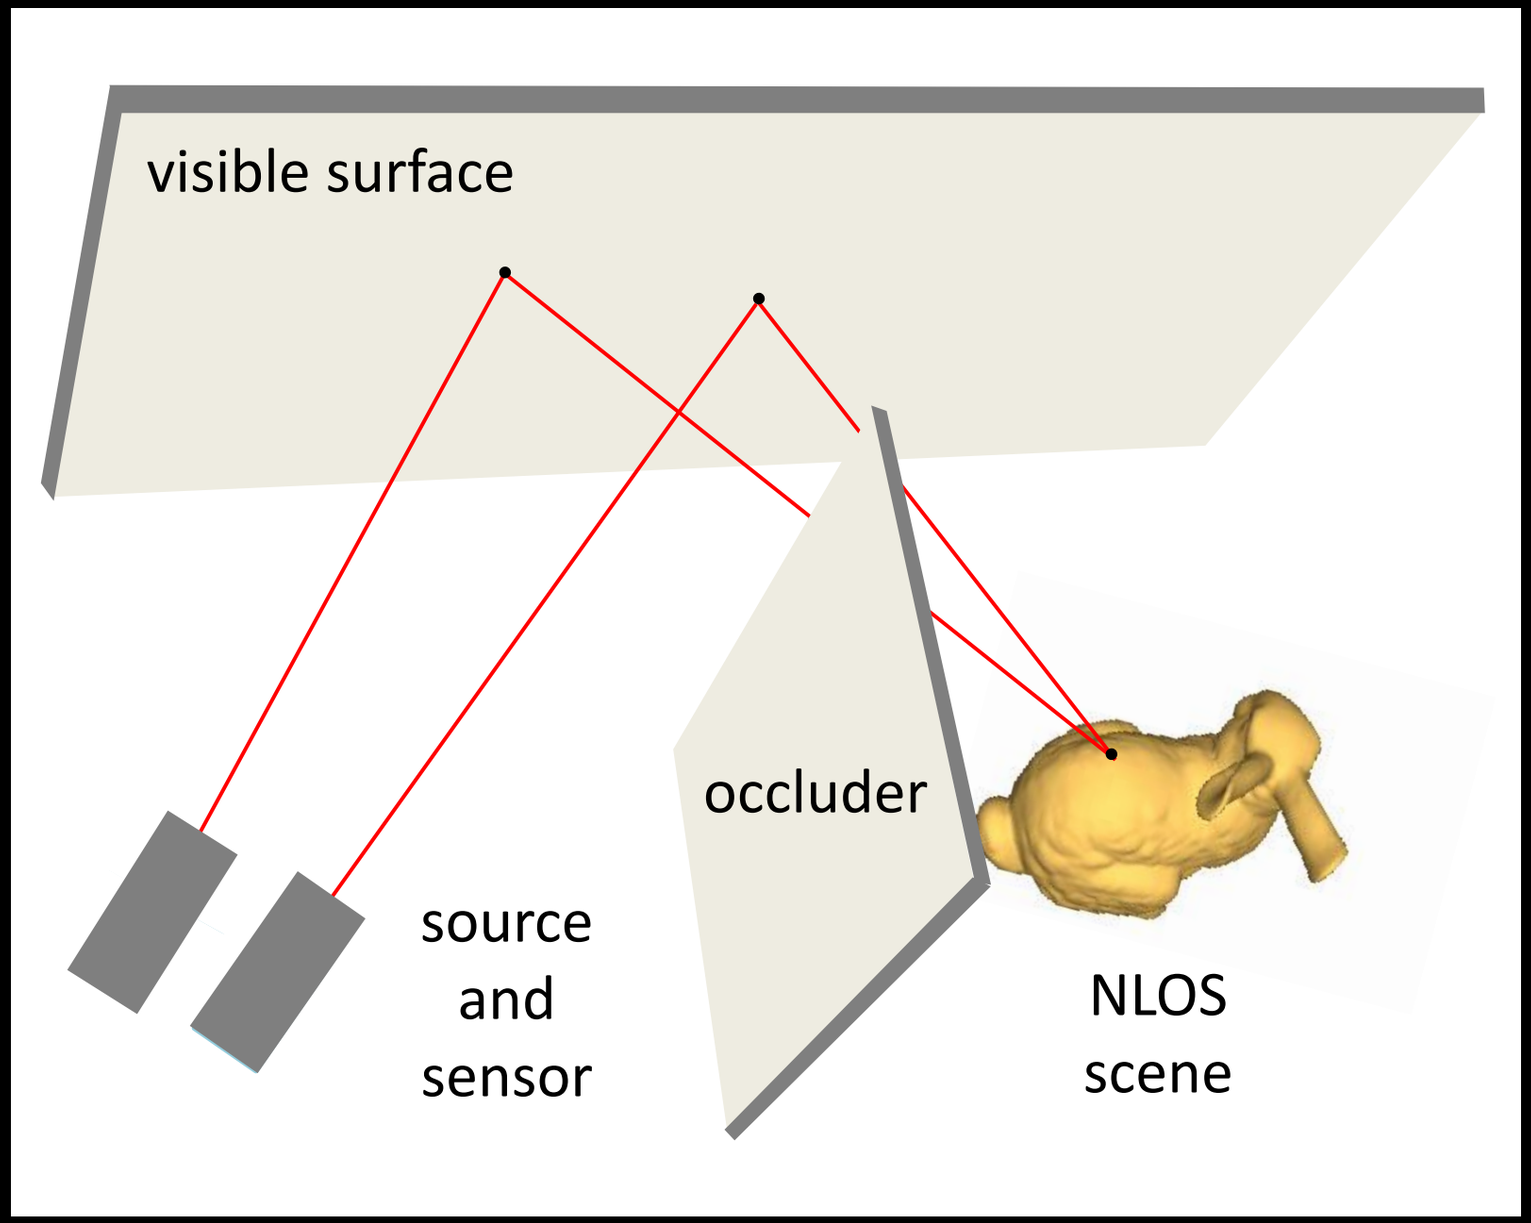

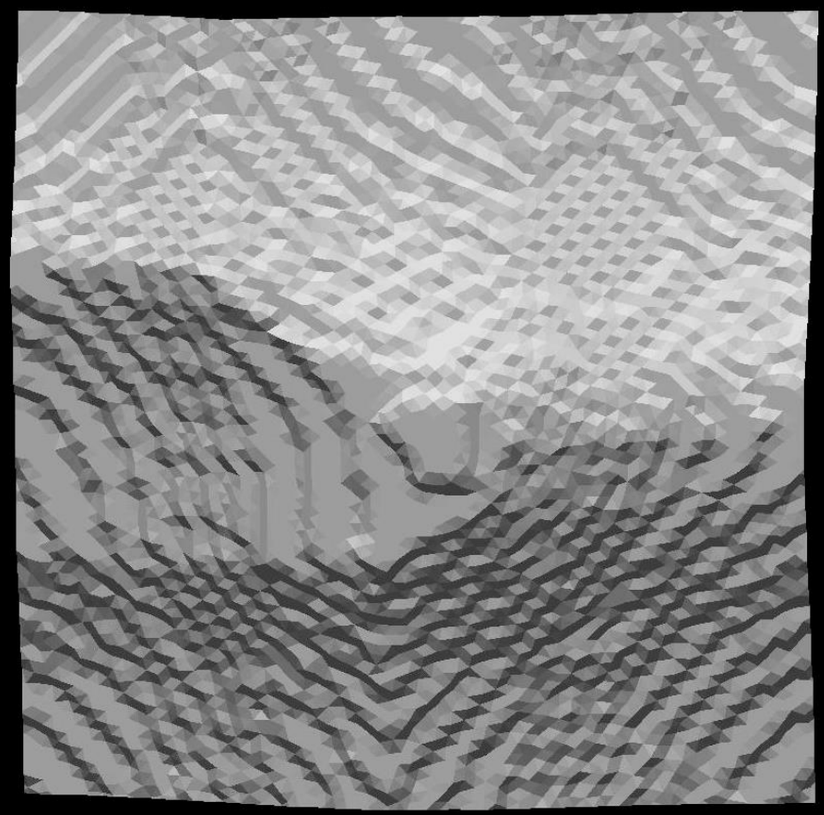

reconstruction evolution

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

TA: Anand Bhoraskar

Master of Science in Computer Vision (MSCV)

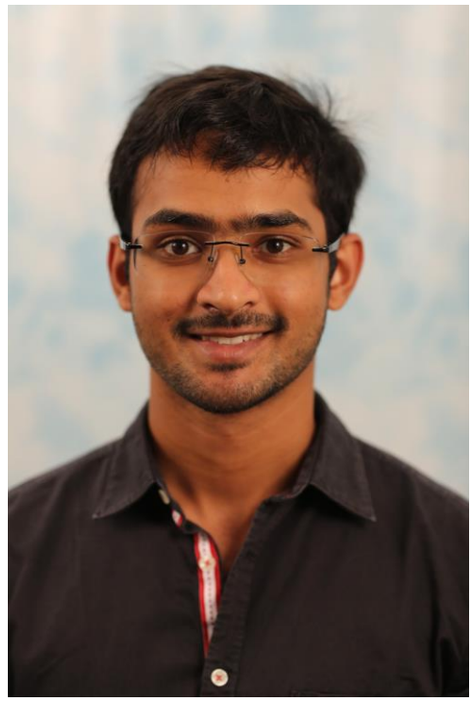

Research Interests:
Simultaneous Localization and Mapping (SLAM), Deep Learning, Object Detection/Tracking

Current Area of Research:
Long term mapping for SLAM for dynamic environments

Past Research:
Video Stabilization, Object Tracking

TA: Prakhar Kulshreshtha (PK)

- **Master of Science in Computer Vision (MSCV)**

- **Research Interests**:
  - Deep Learning for Detection and Instance Segmentation
  - SLAM
  - Online Learning in Computer Vision

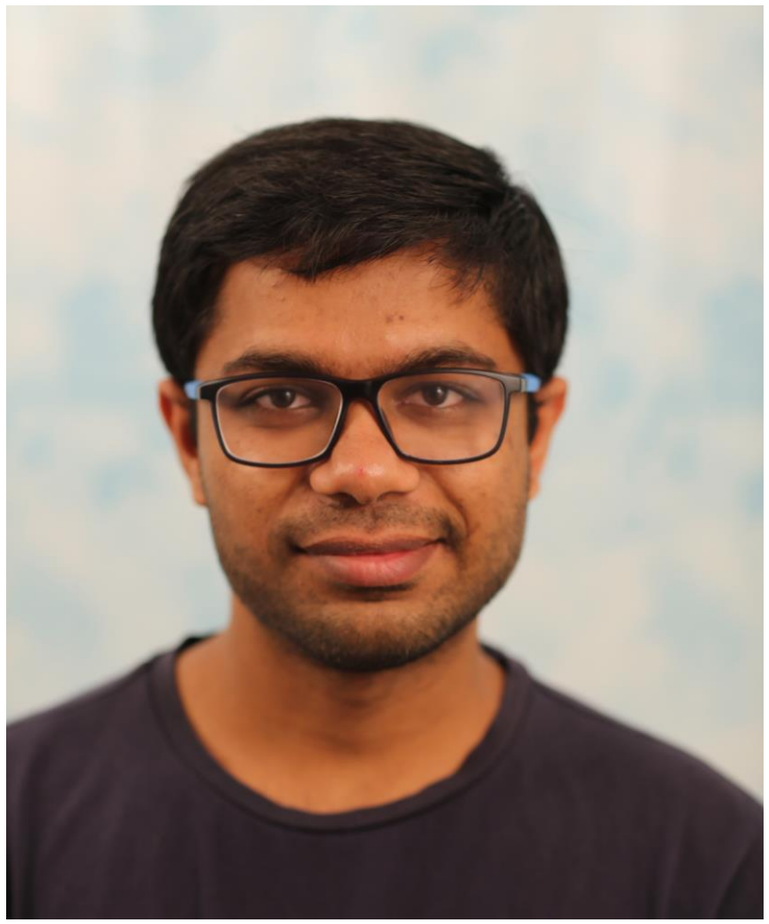

- **Current area of research**:
  - Long-term SLAM for Dynamic environments (under Prof. Michael Kaess)

- **Past Reseach**:
  - Instance segmentation for quality estimation of food grains on a smartphone
  - Text Intelligence in smartphone keyboard apps
  - Online Face Clustering

What is computer vision?

What a person sees

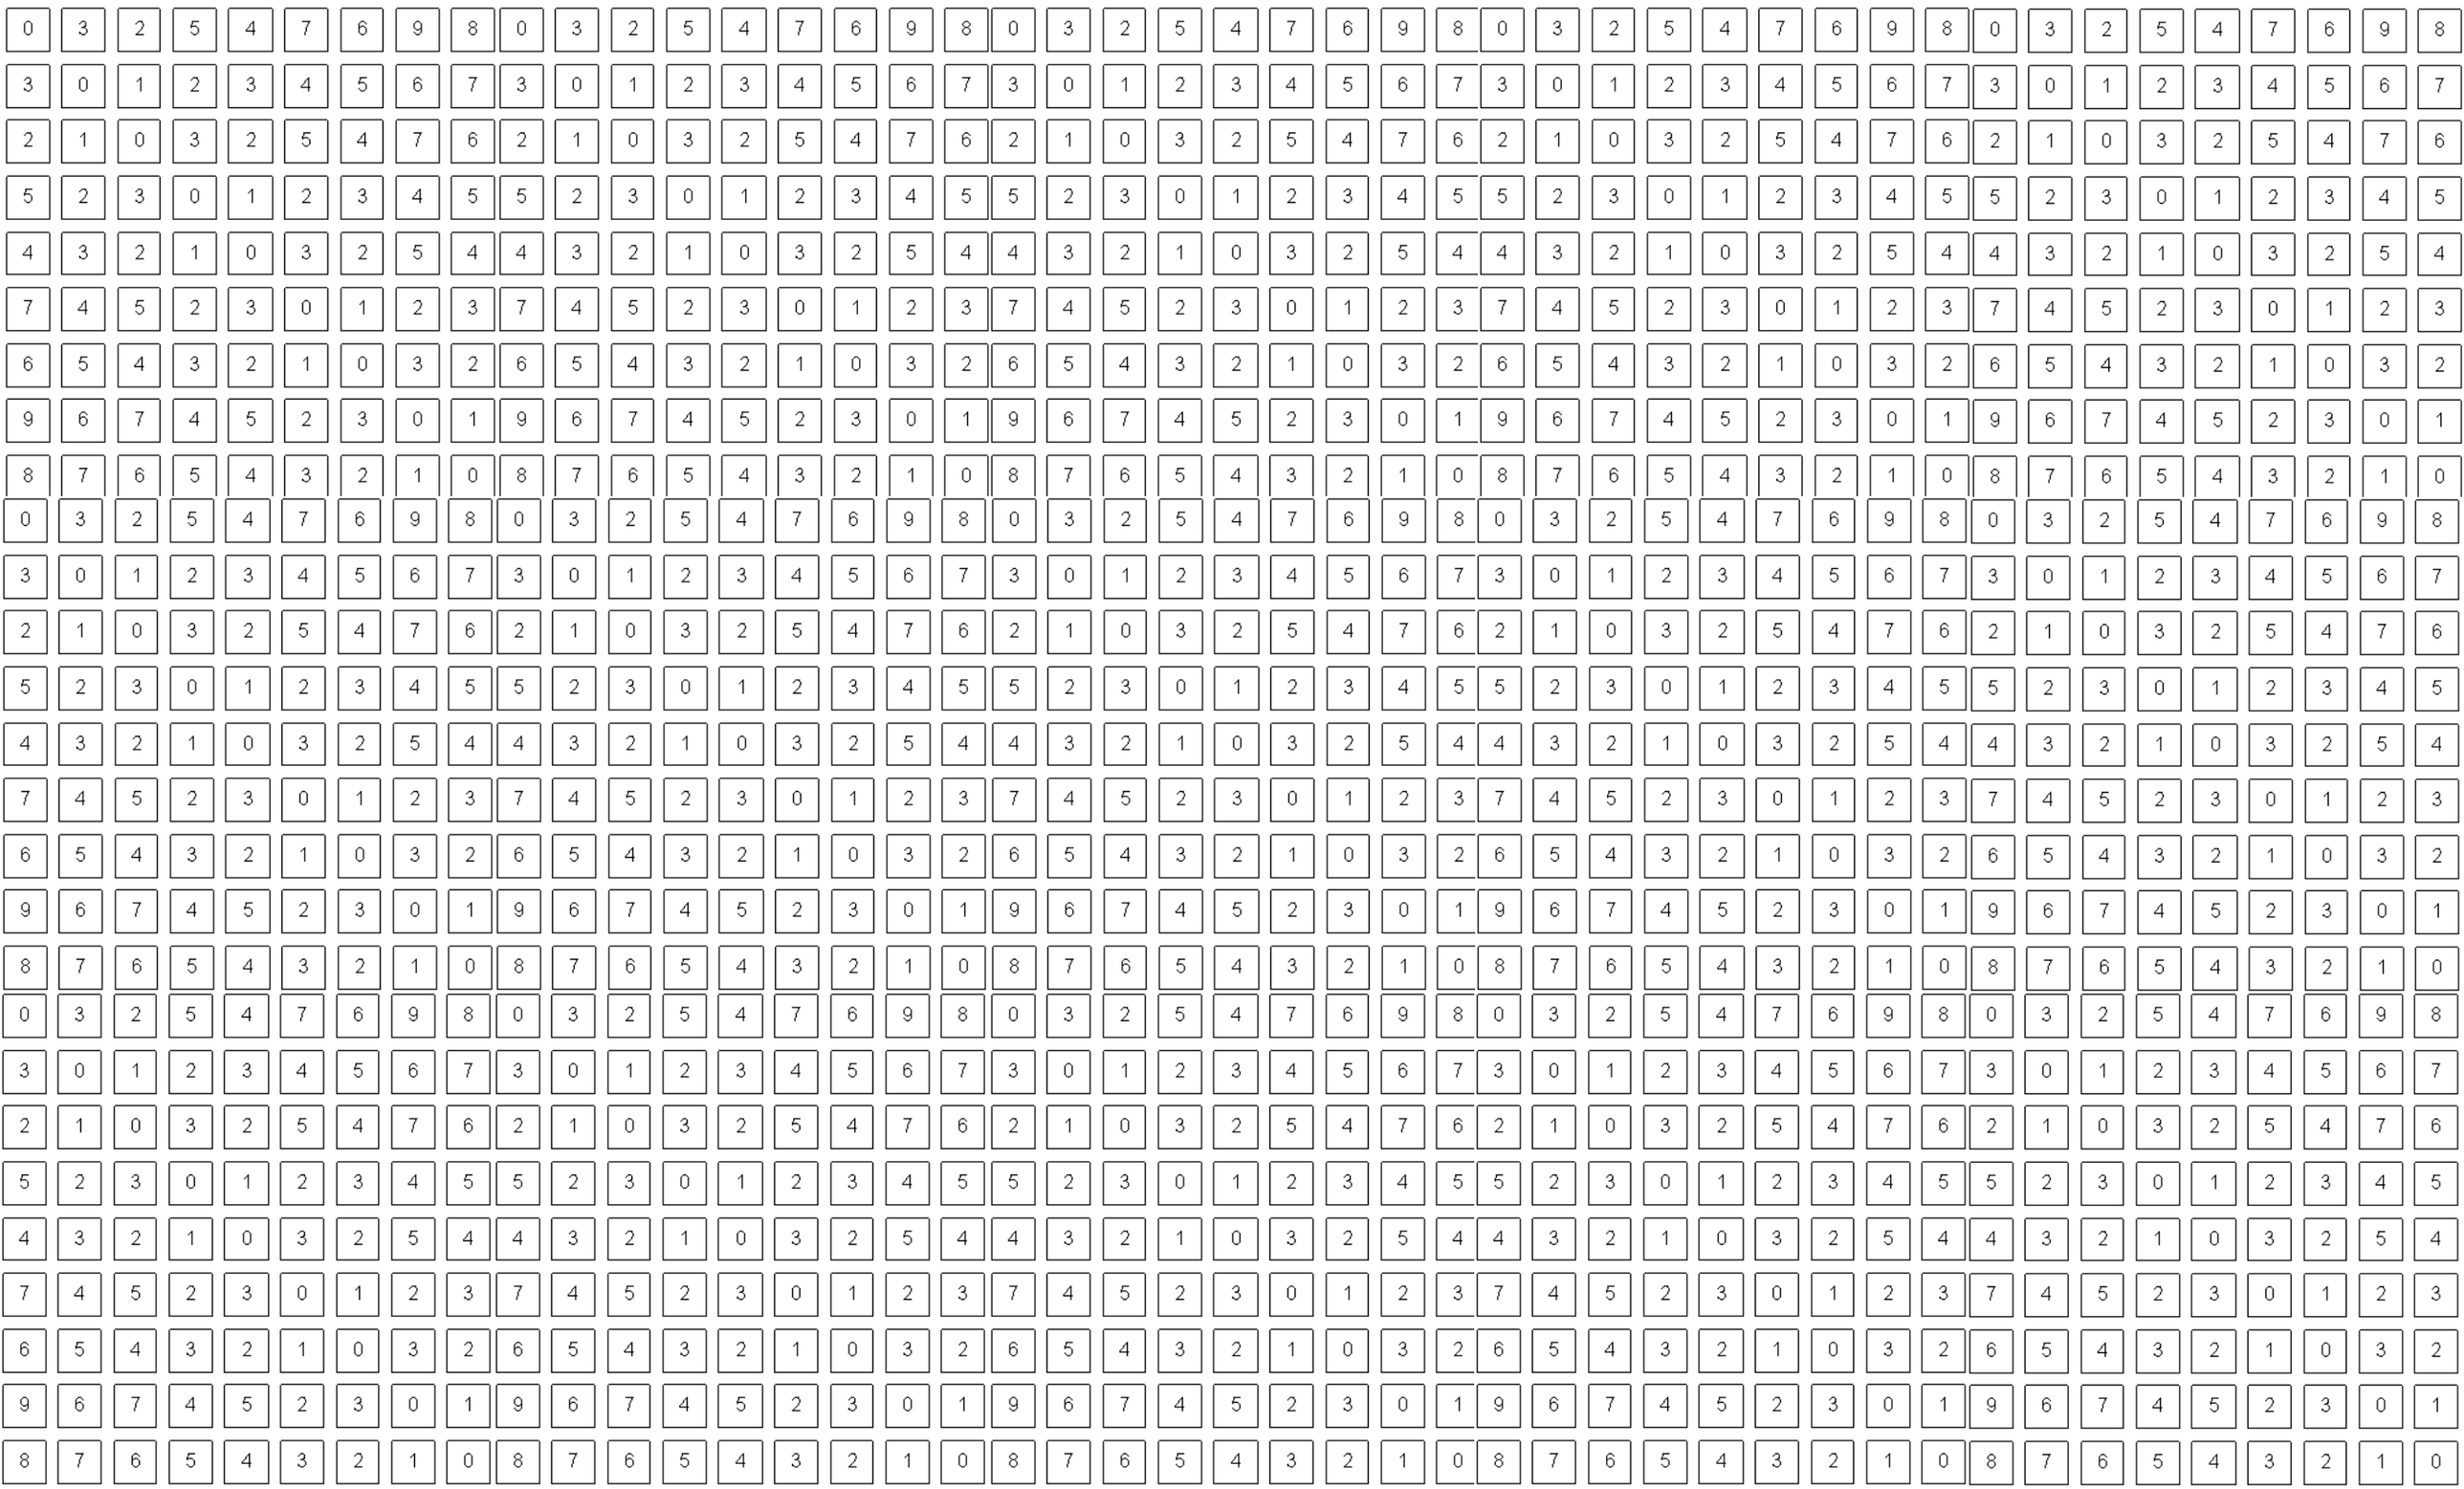

What a computer sees

Why are we able to interpret this image?

The goal of computer vision is
to give computers
(super) human-level perception

typical perception pipeline

representation

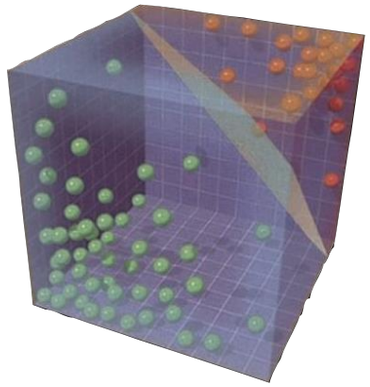

‘fancy math’

output

typical perception pipeline

representation

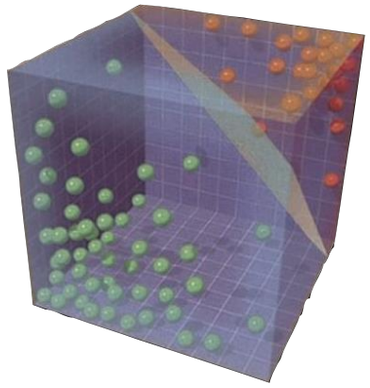

‘fancy math’

what should we look at?
(image features)

output

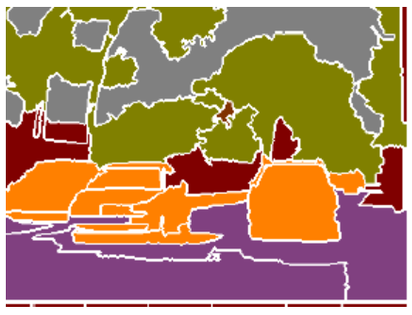

what can we understand?
(semantic segmentation)

typical perception pipeline

representation

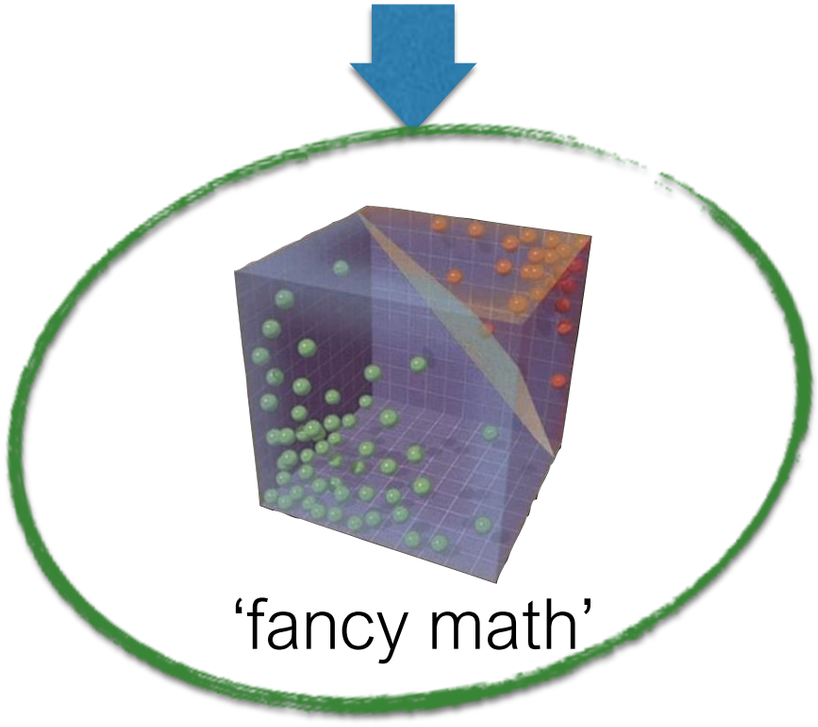

output

what should we look at?
(image features)

easy to get lost in
the techniques

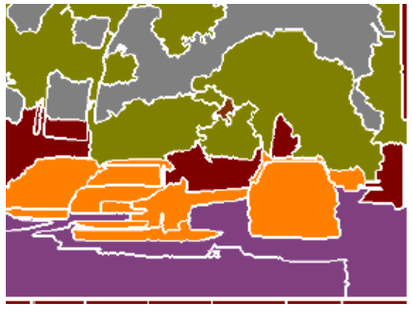

what can we understand?
(semantic segmentation)

typical perception pipeline

representation

what should we look at?
(image features)

The parts that we are most interested in

'fancy math'

output

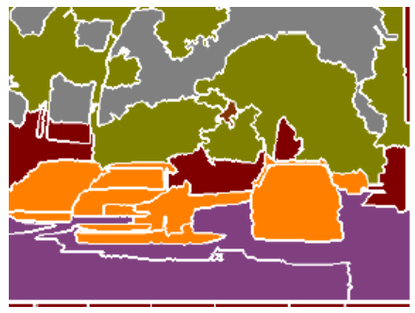

what can we understand?
(semantic segmentation)

Important note:

In general, computer vision does not work

Important note:

In general, computer vision does not work

(except in certain situations/conditions)

Applications of computer vision

Machine vision

Automated visual inspection

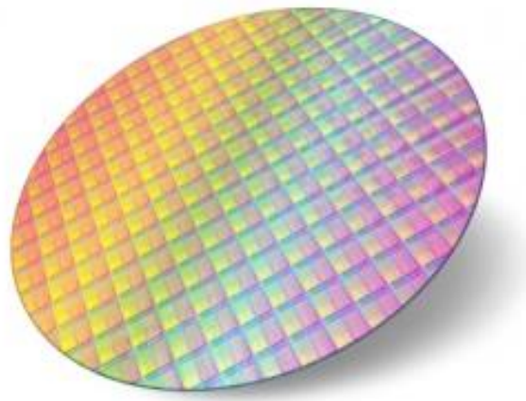

Object Recognition

Toshiba Tech IS-910T 2013

DataLogic LaneHawk LH4000 2012

Face detection

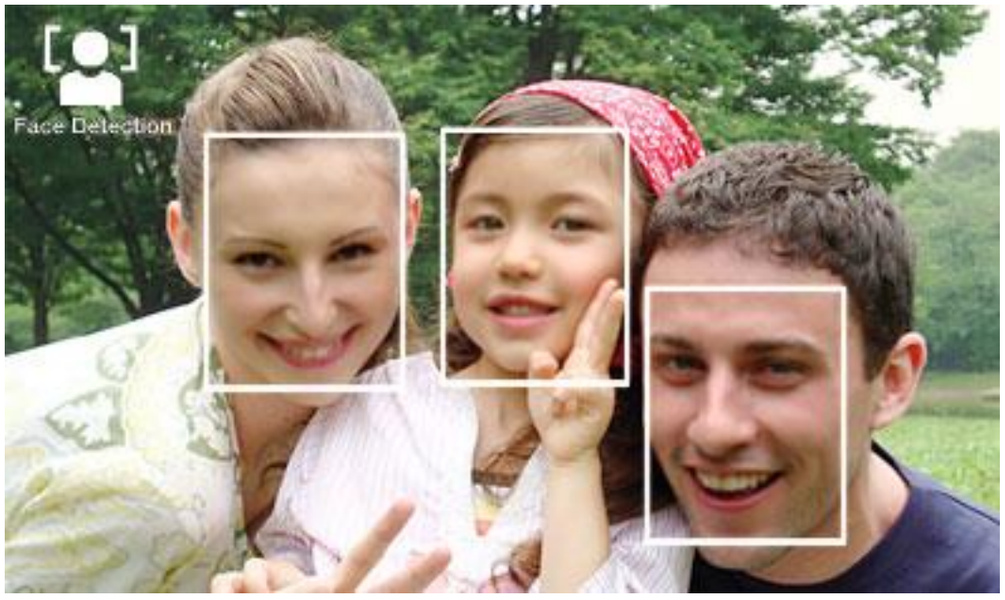

Sony Cyber-shot

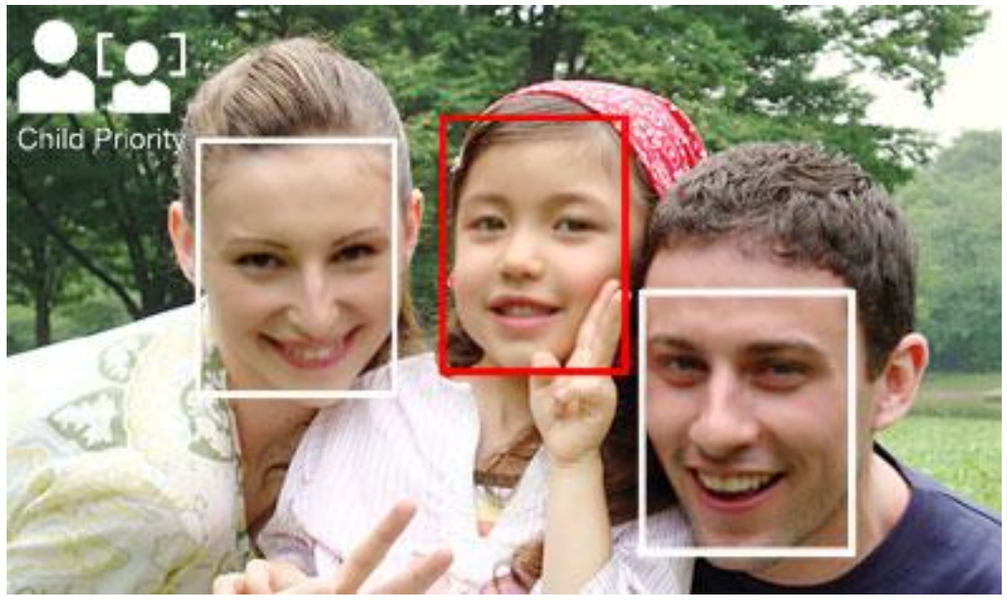

Age recognition

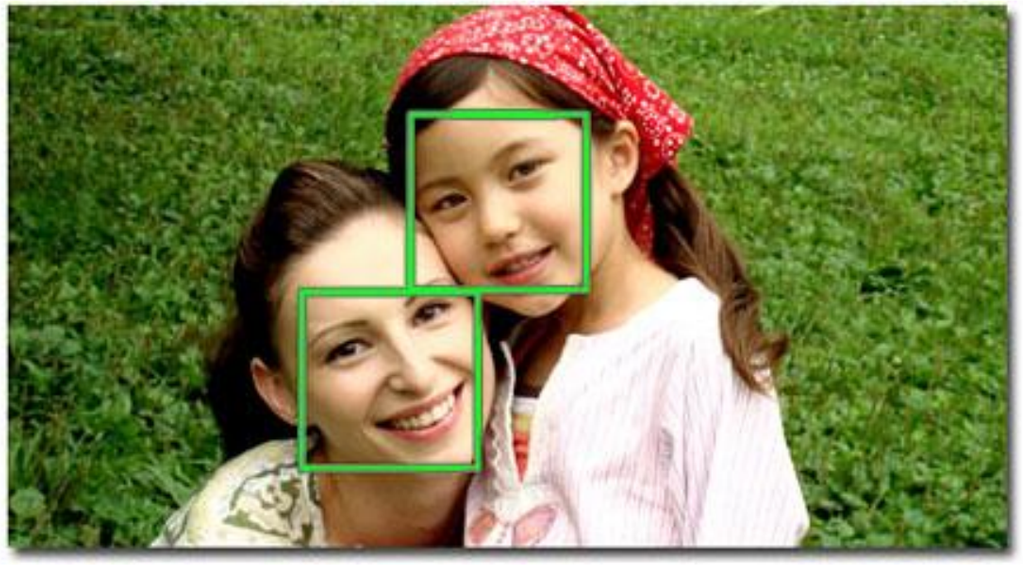

Smile recognition

Face makeovers

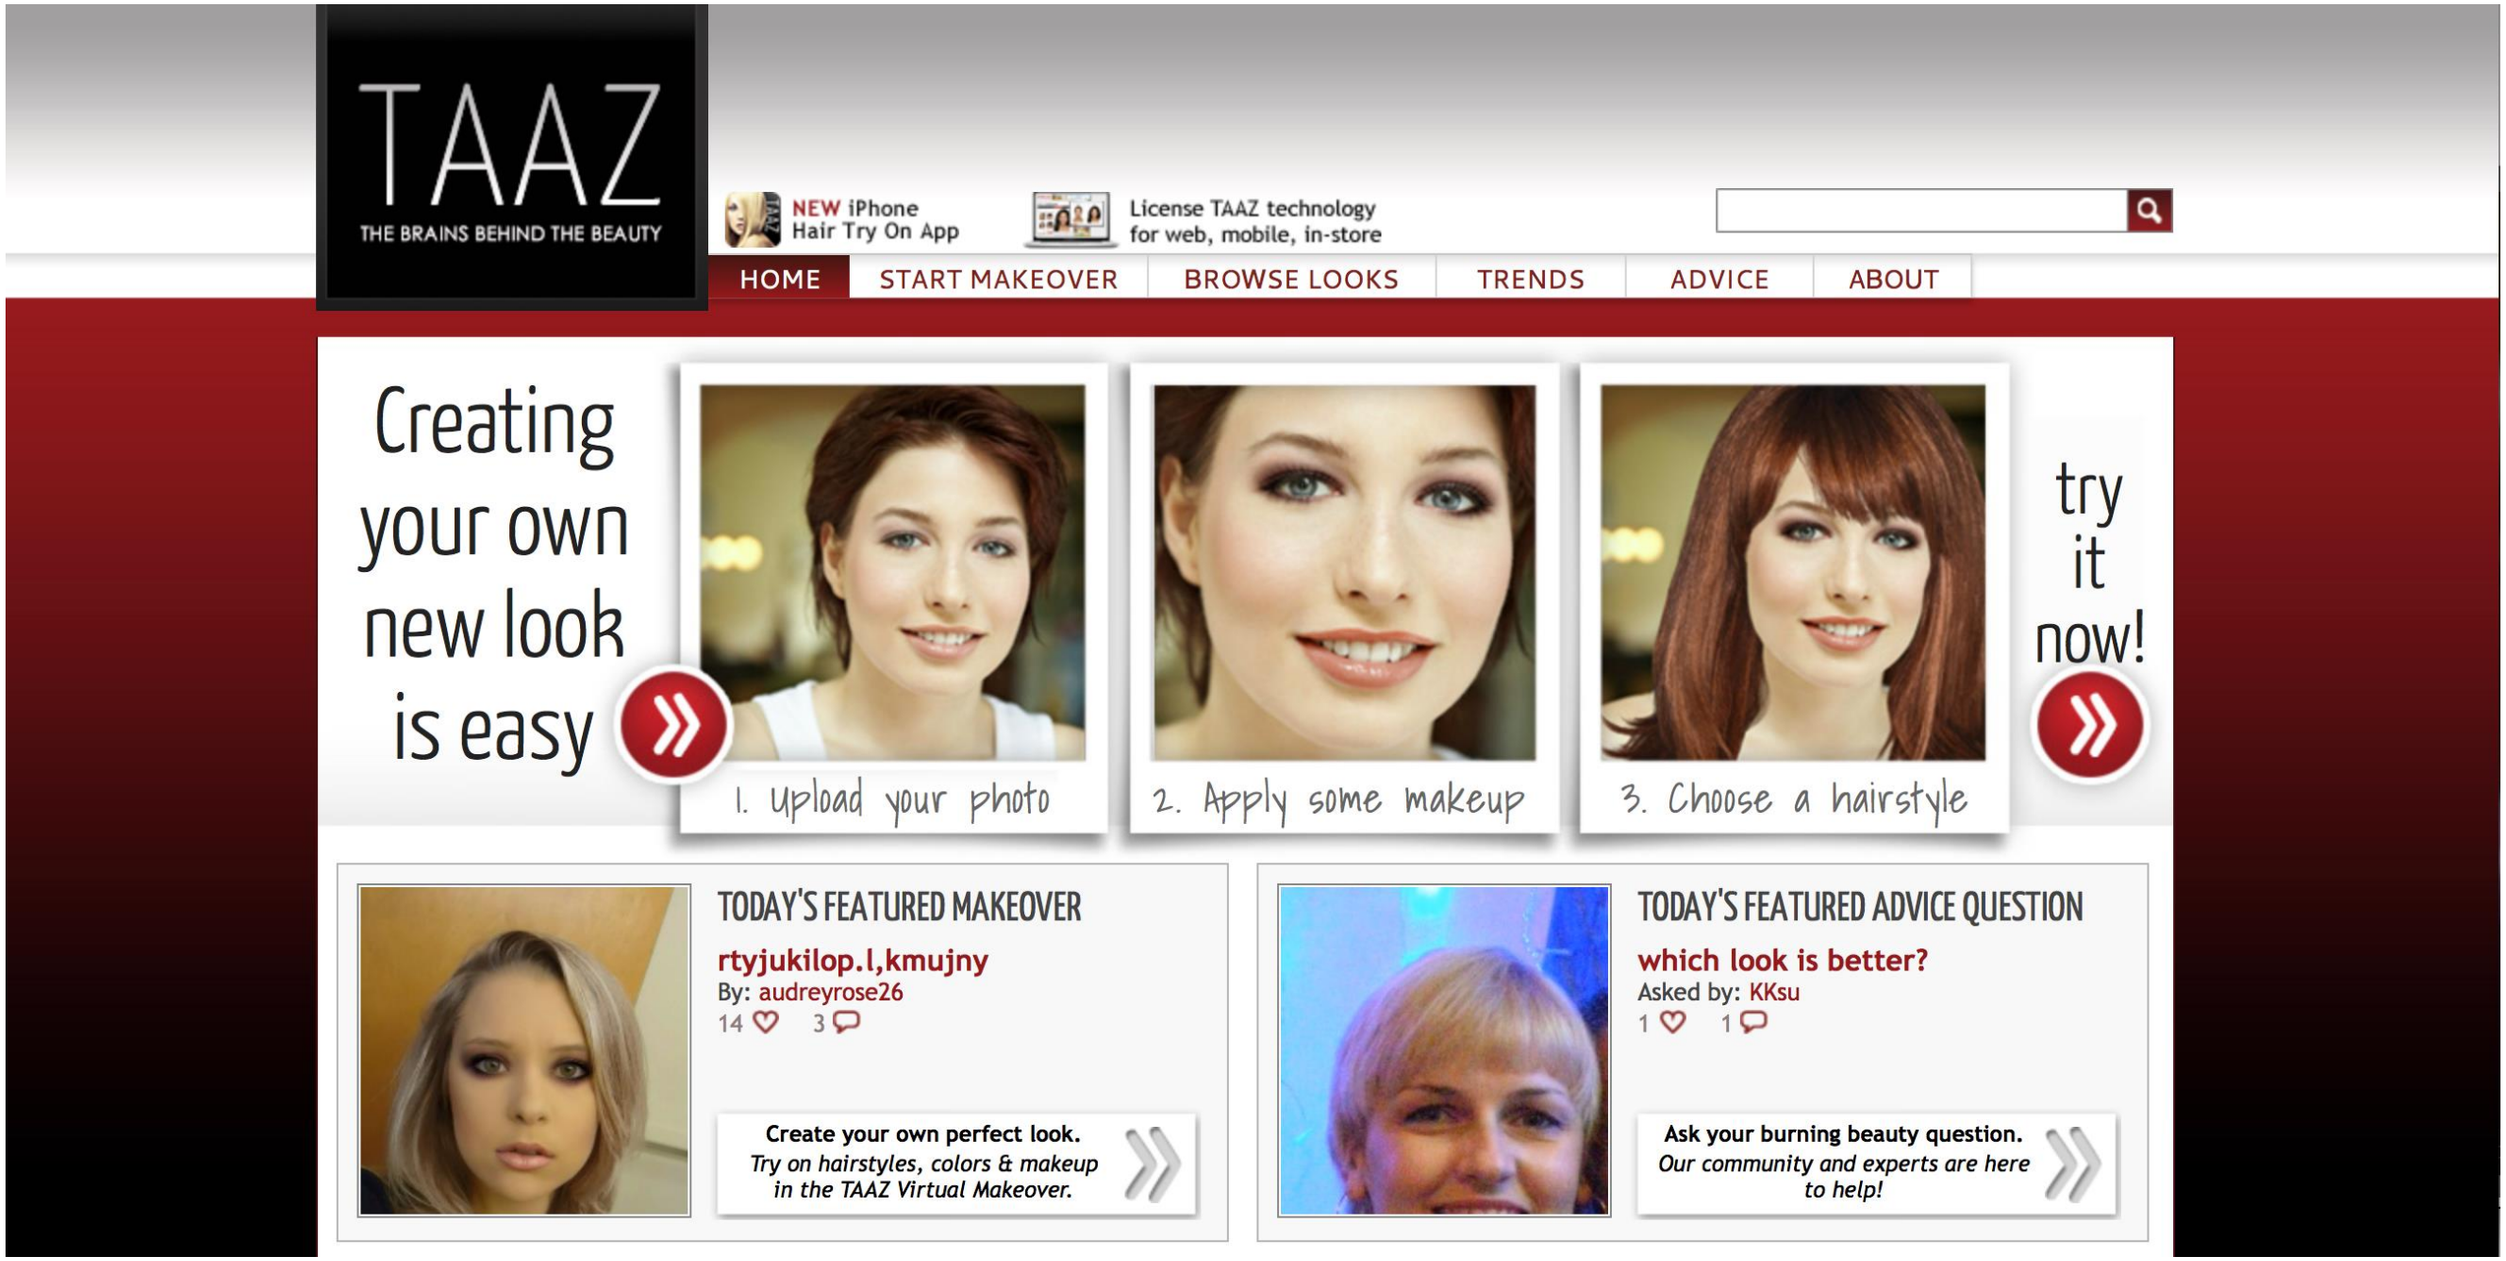

leafsnap


Verizon
12:38 PM
Back
Results
Map
Snap It! Results
1
Red Maple
Acer rubrum
2
Striped Maple
Acer pensylvanicum
3
Sycamore Maple
Acer pseudoplatanus
12:20 PM
100%
First
Last
Scientific
Ilex opaca
American Hornbeam
Carpinus caroliniana
American Linden
Tilia americana
American Sycamore
Platanus occidentalis
Amur Corktree
Phellodendron amurense
Q A B C D E F G H J K L M N O P Q R S T U W Y
Home
Browse
Collection
Options
Snap It!

Word Lens

Word Lens

www.QuestVisual.com

First-down line

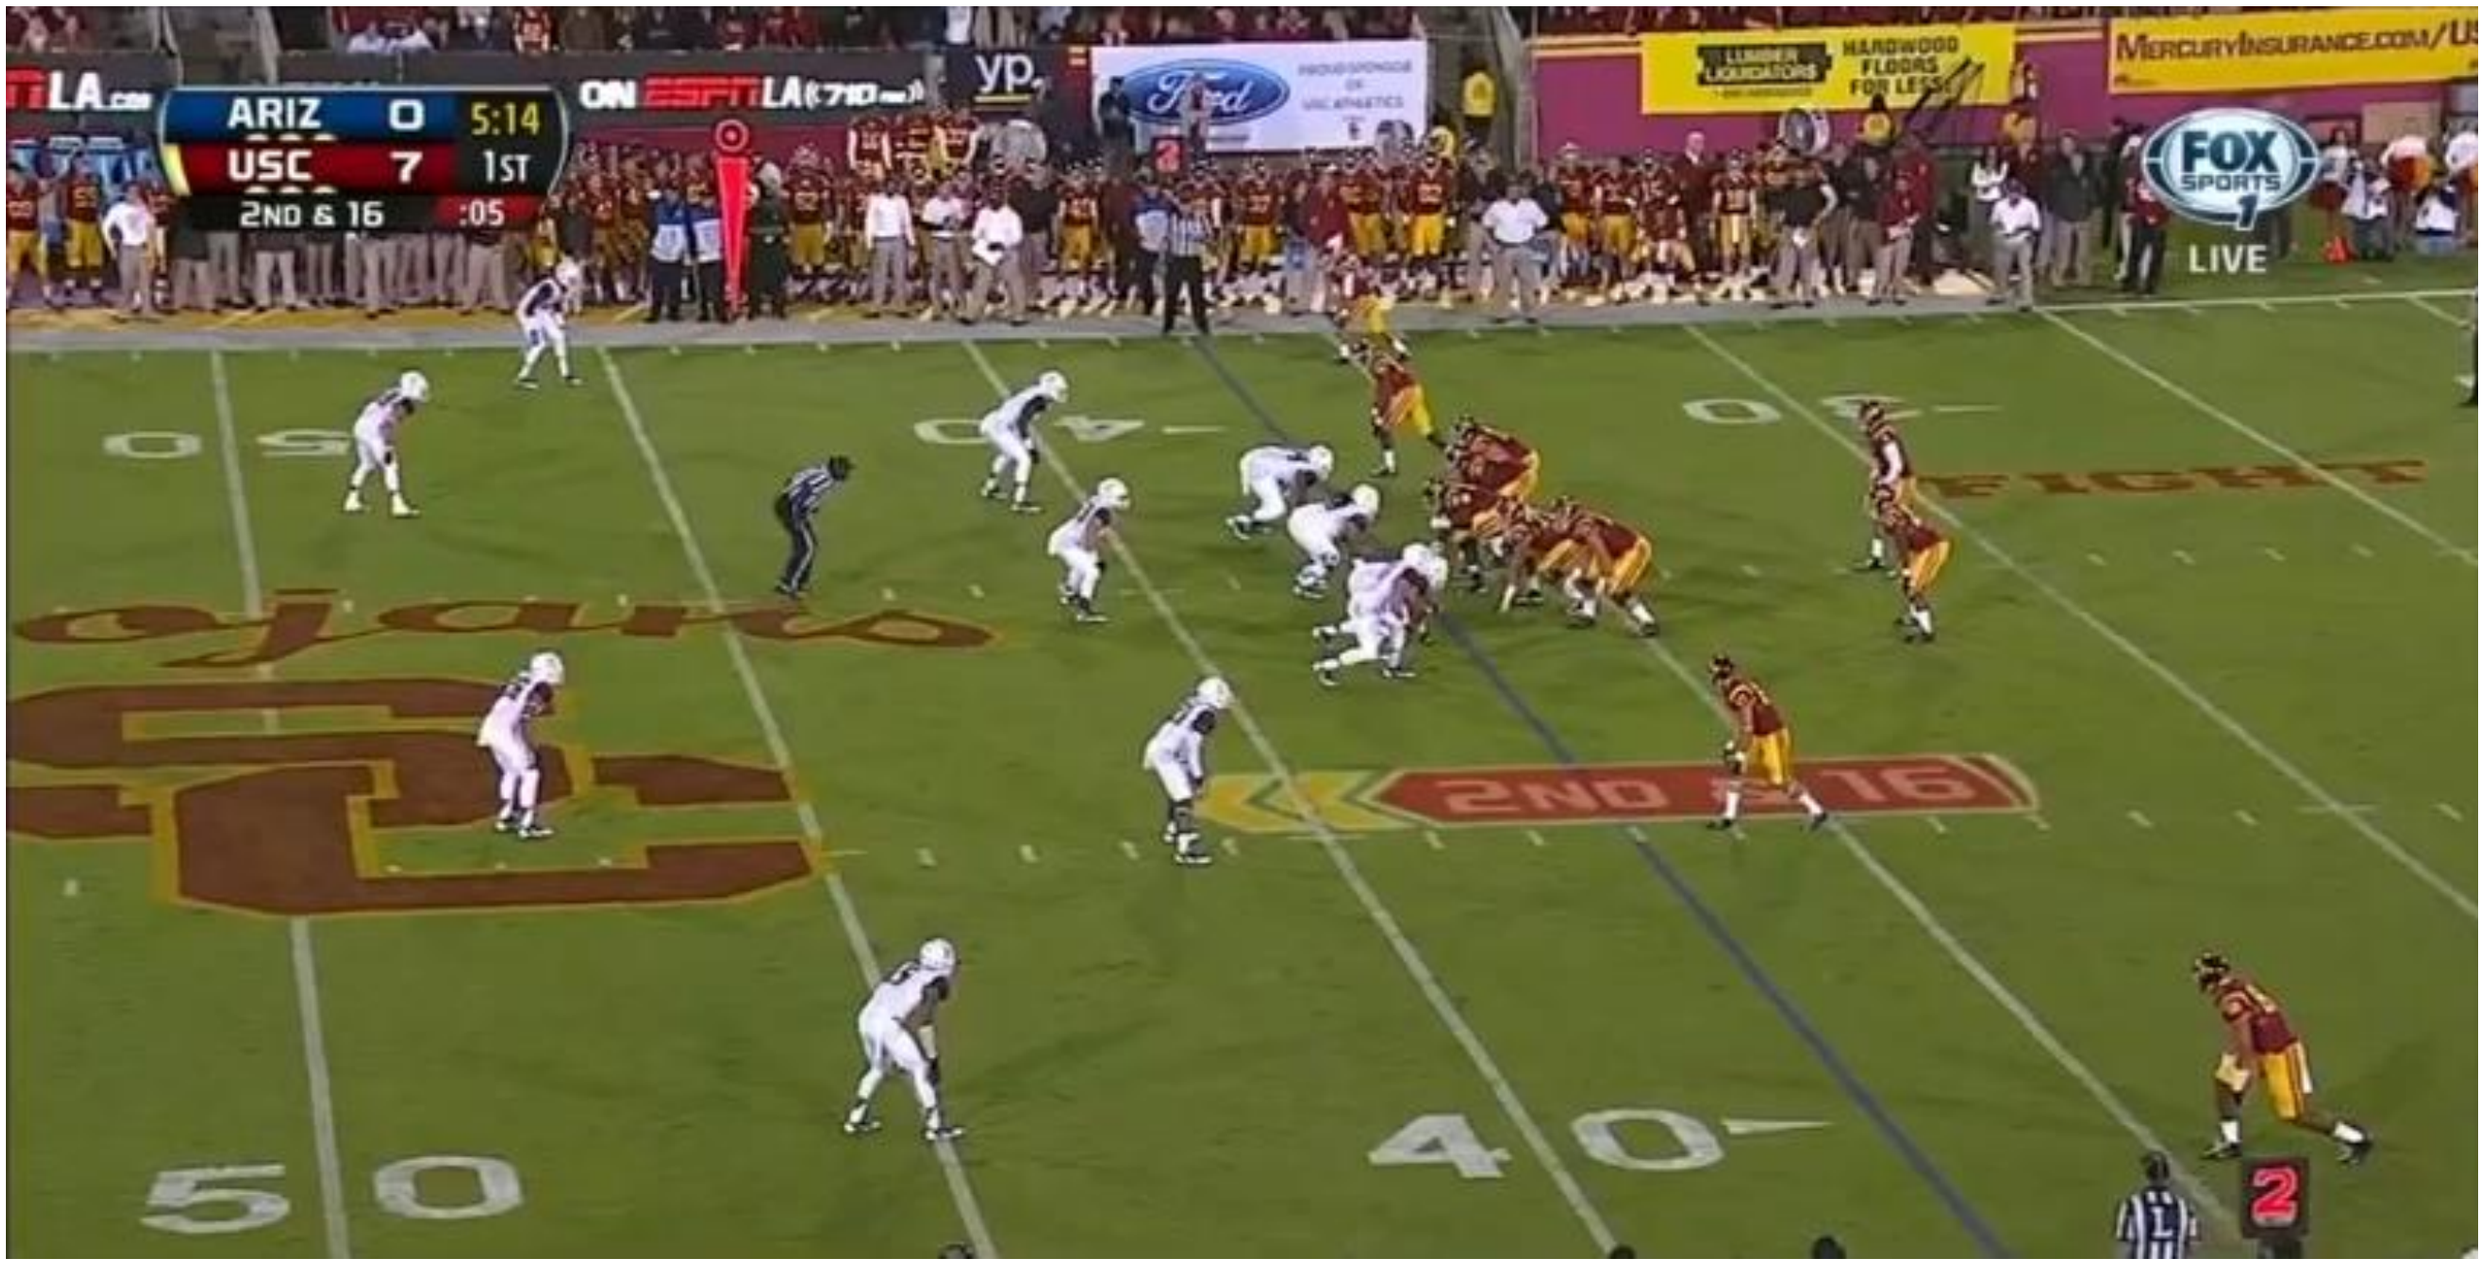

BMW 5 series

BMW night vision

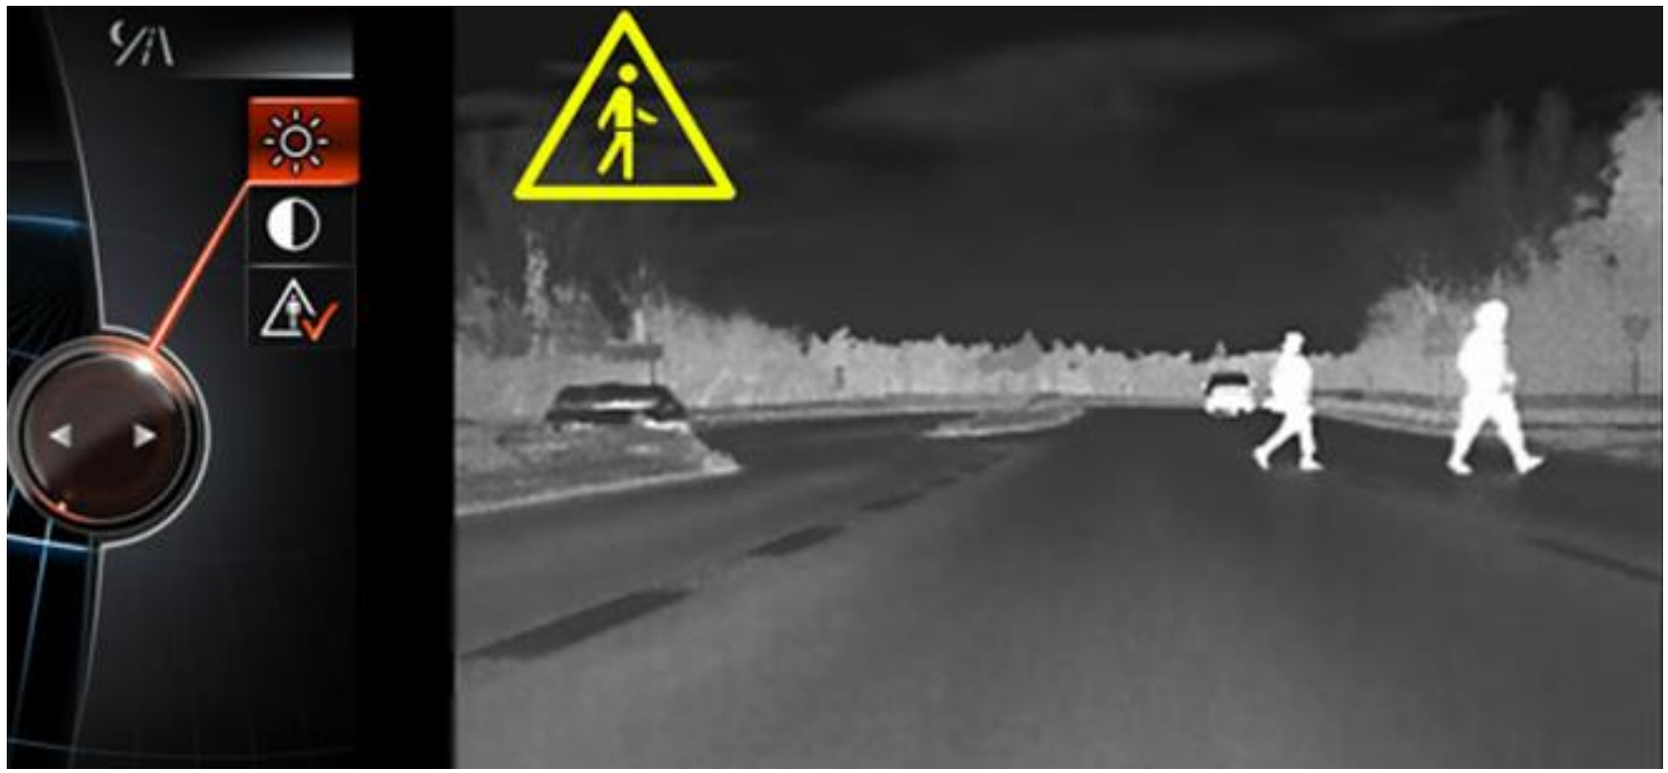

Infinity EX

“Around view” camera

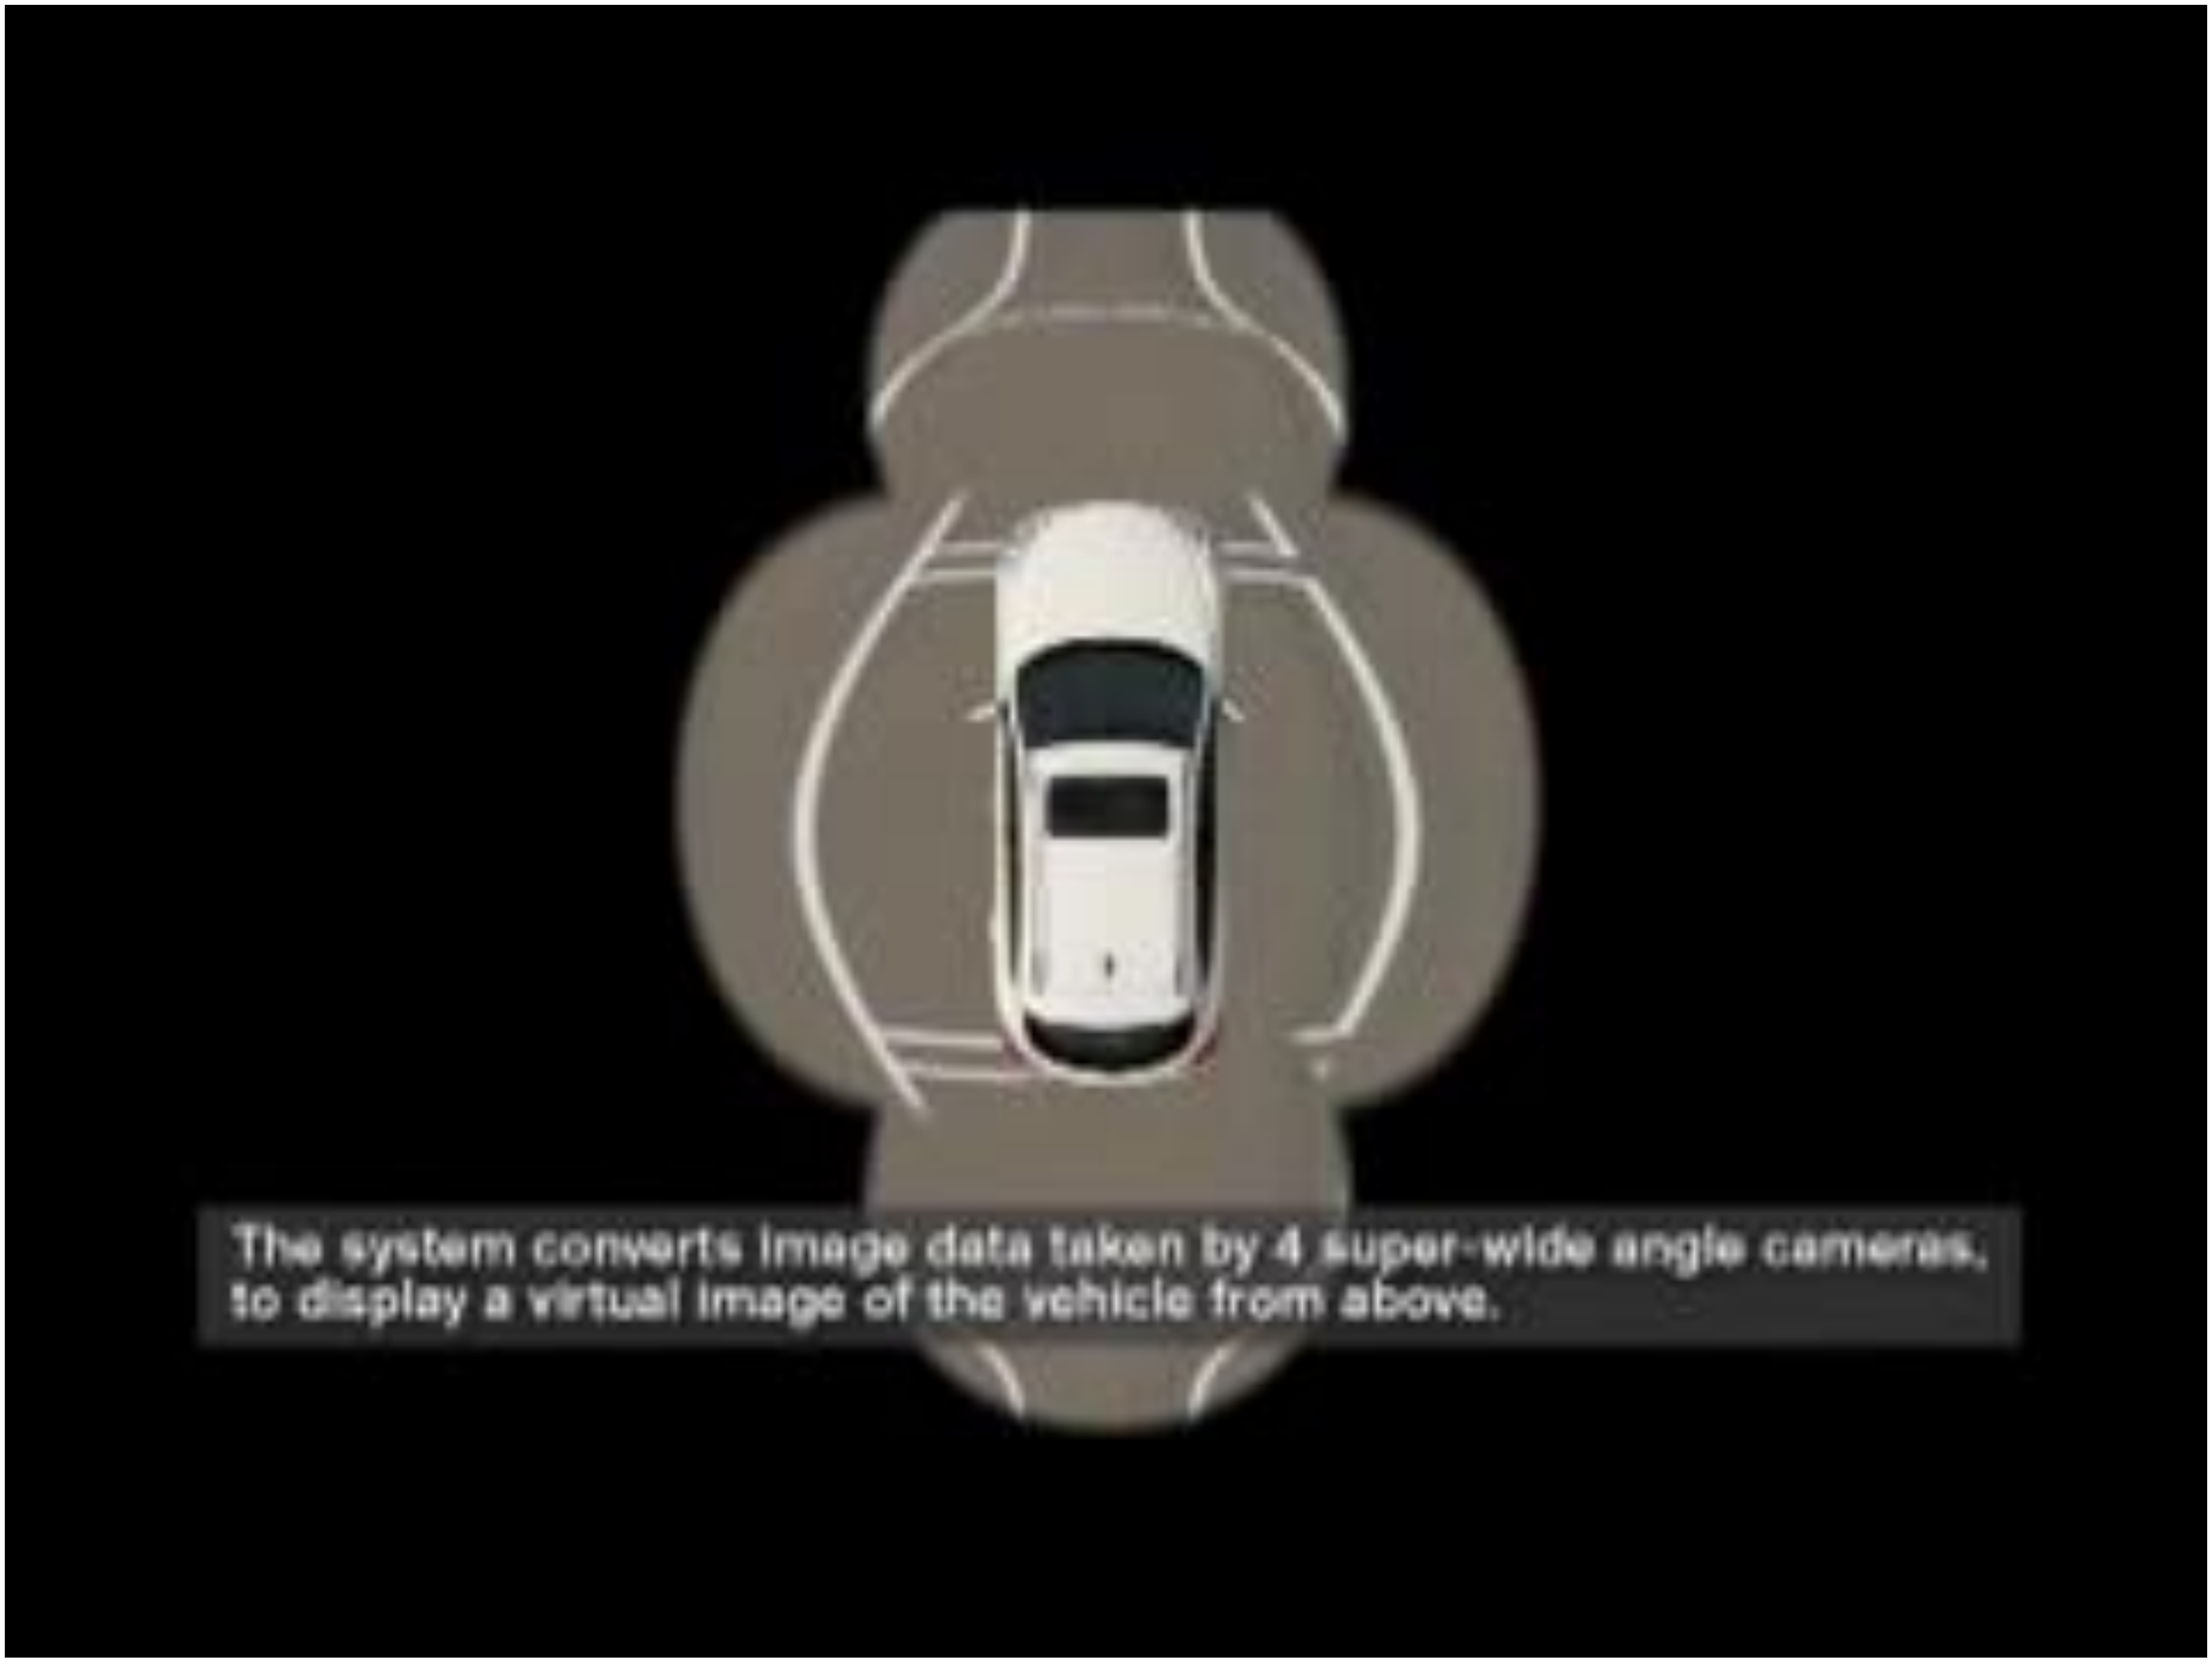
The system converts image data taken by 4 super-wide angle cameras,
to display a virtual image of the vehicle from above.

2015

Vision in Cars

Image stitching

Photosynth

Tango

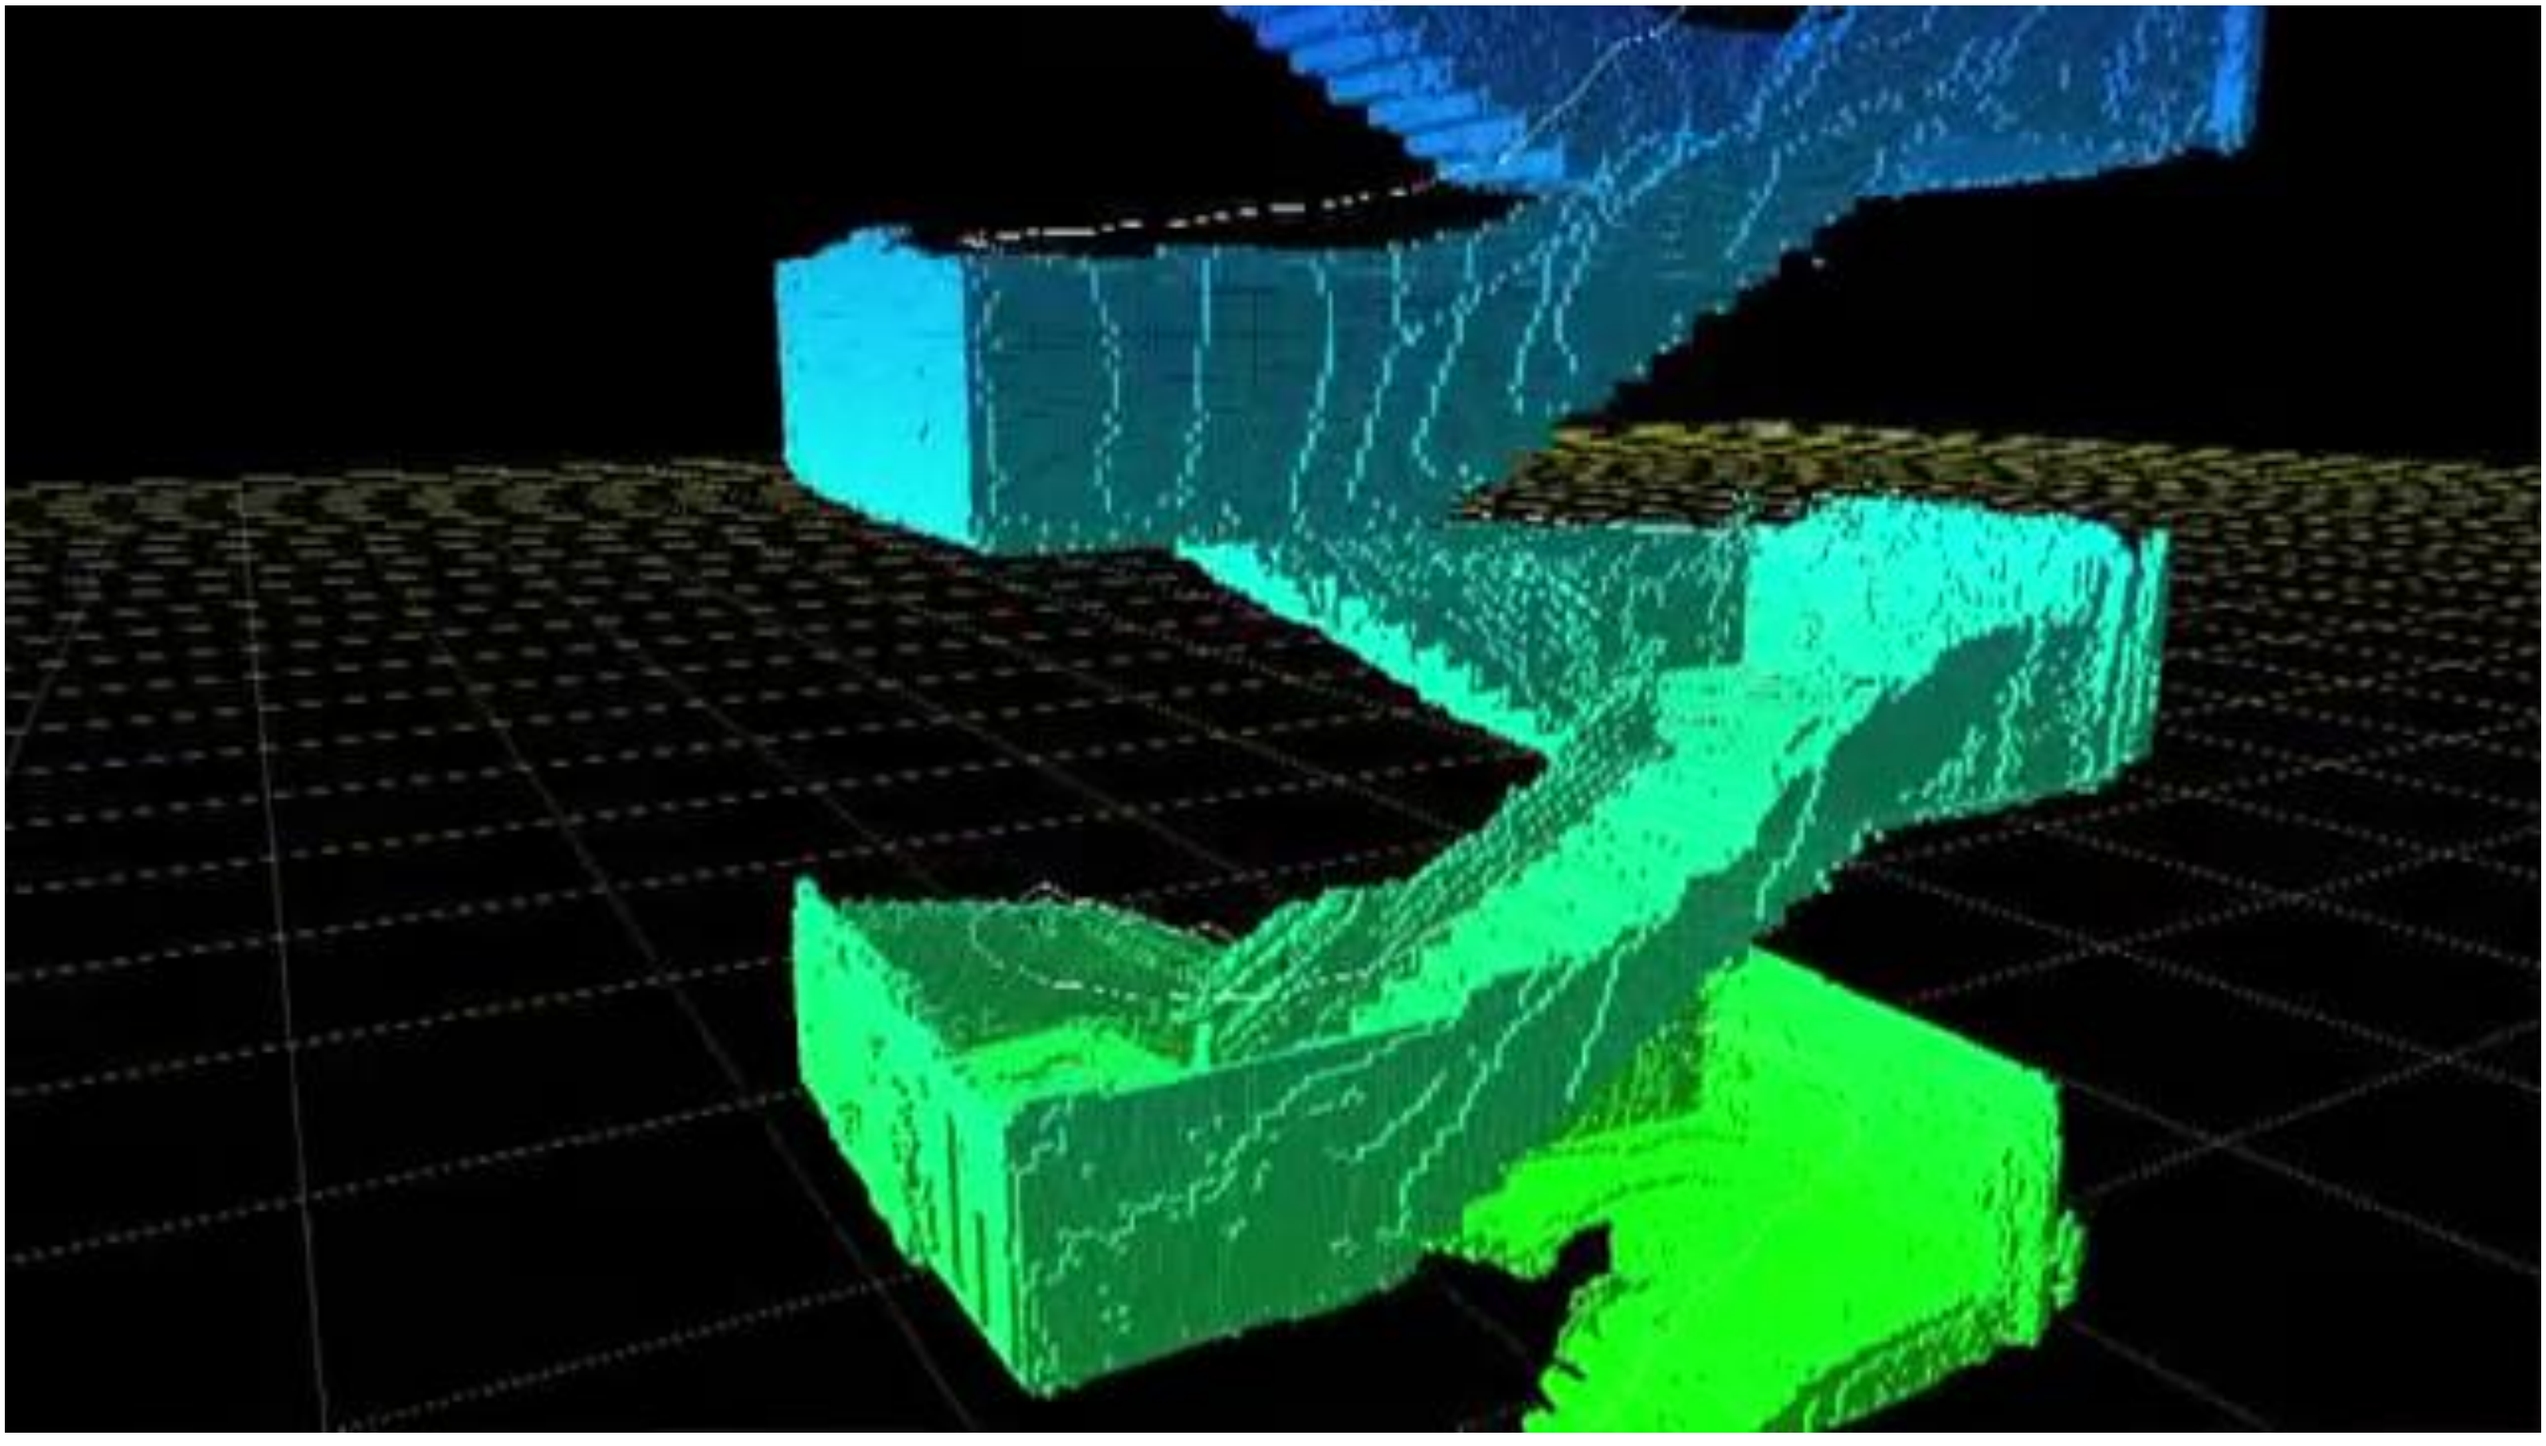

Virtual Fitting

2015

Computer Vision for VR

Deep Face

Deep Dream

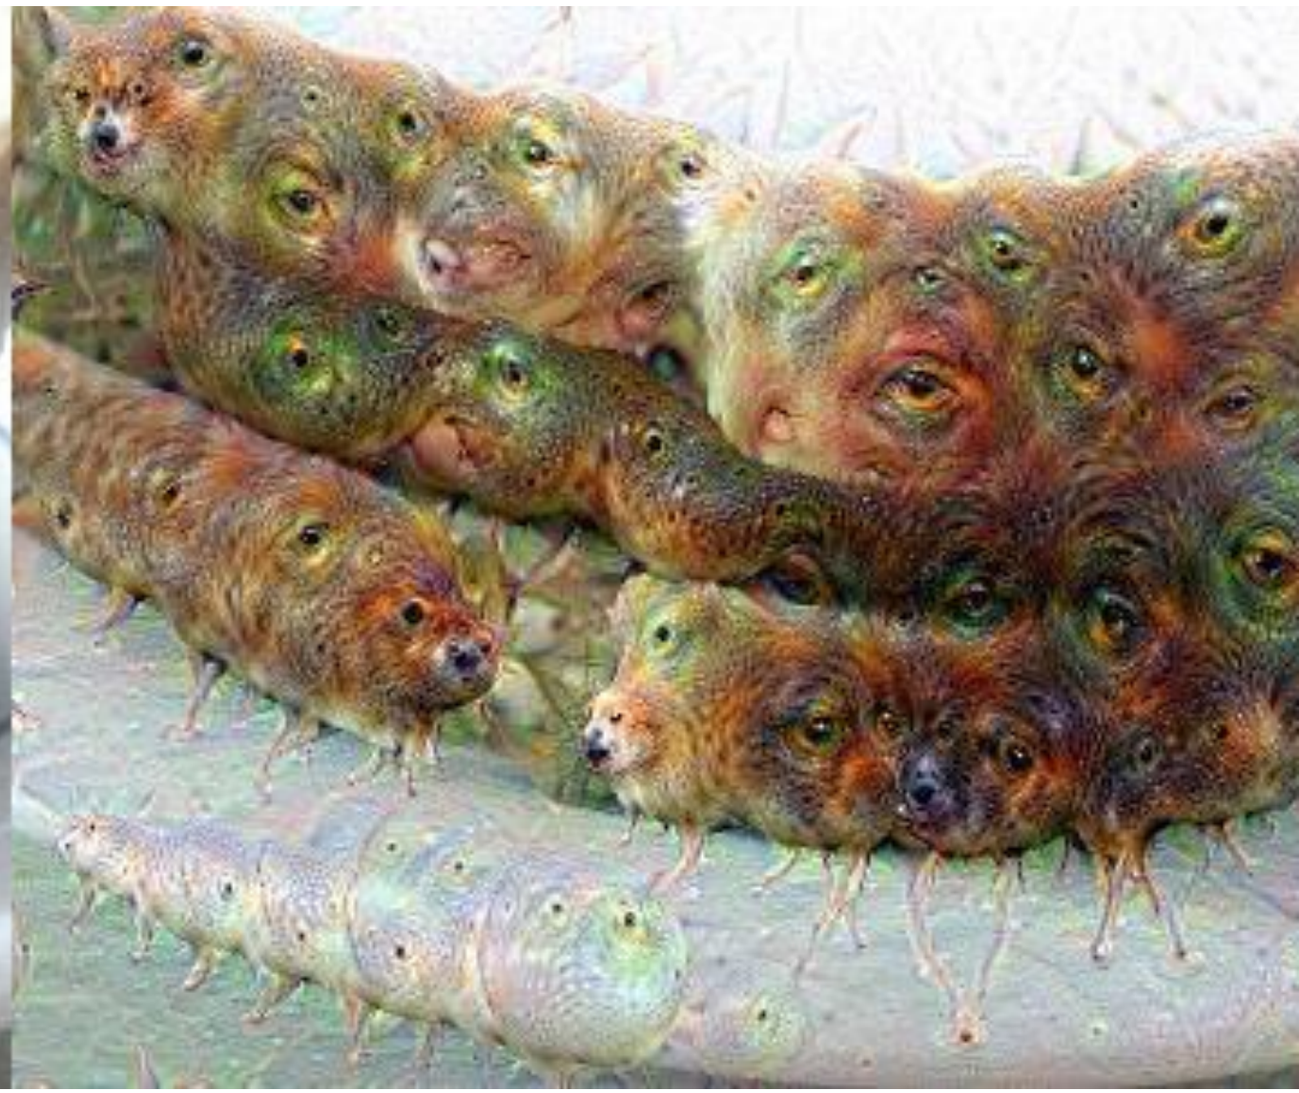

Facebook video style transfer 2016

Face2Face: Real-time Face Capture and Reenactment of RGB Videos

Justus Thies[1], *Michael Zollhöfer*[2],
Marc Stamminger[1], *Christian Theobalt*[2],
Matthias Nießner[3]

[1]University of Erlangen-Nuremberg
[2]Max-Planck-Institute for Informatics
[3]Stanford University

CVPR 2016 (Oral)

It’s a good time to do computer vision

Industry aggressively hiring CV faculty from universities

UW

Berkeley

CMU

Toronto

UBER

UCLA

USC

Dropbox

UCSD

Columbia

NYU

CMU CMU

NVIDIA

Toronto

amazon.com

CalTech

UW

Stanford

CMU

Stanford

NYU

GTech

GTech

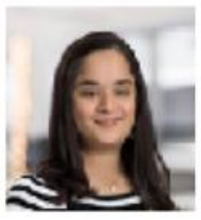

UT Austin

NYU

Google

Zillow

Oculus VR

SFU

CMU

ARGO AI

CMU CMU GTech

CMU

IBM

MIT

MIT Toronto UW

Industry aggressively hiring CV graduates, or even students!

(~~strong~~ dominant industrial presence at conferences for recruitment)

facebook research | Research Areas | Publications | People | Programs | Downloads | Careers | Blog

JULY 21, 2017

Advancing computer vision technologies at CVPR 2017

By: Facebook Research

Google Research Blog

The latest news from Research at Google

Google at CVPR 2017

Friday, July 21, 2017

CVPR GROWTH

Original slide courtesy of CVPR 2016

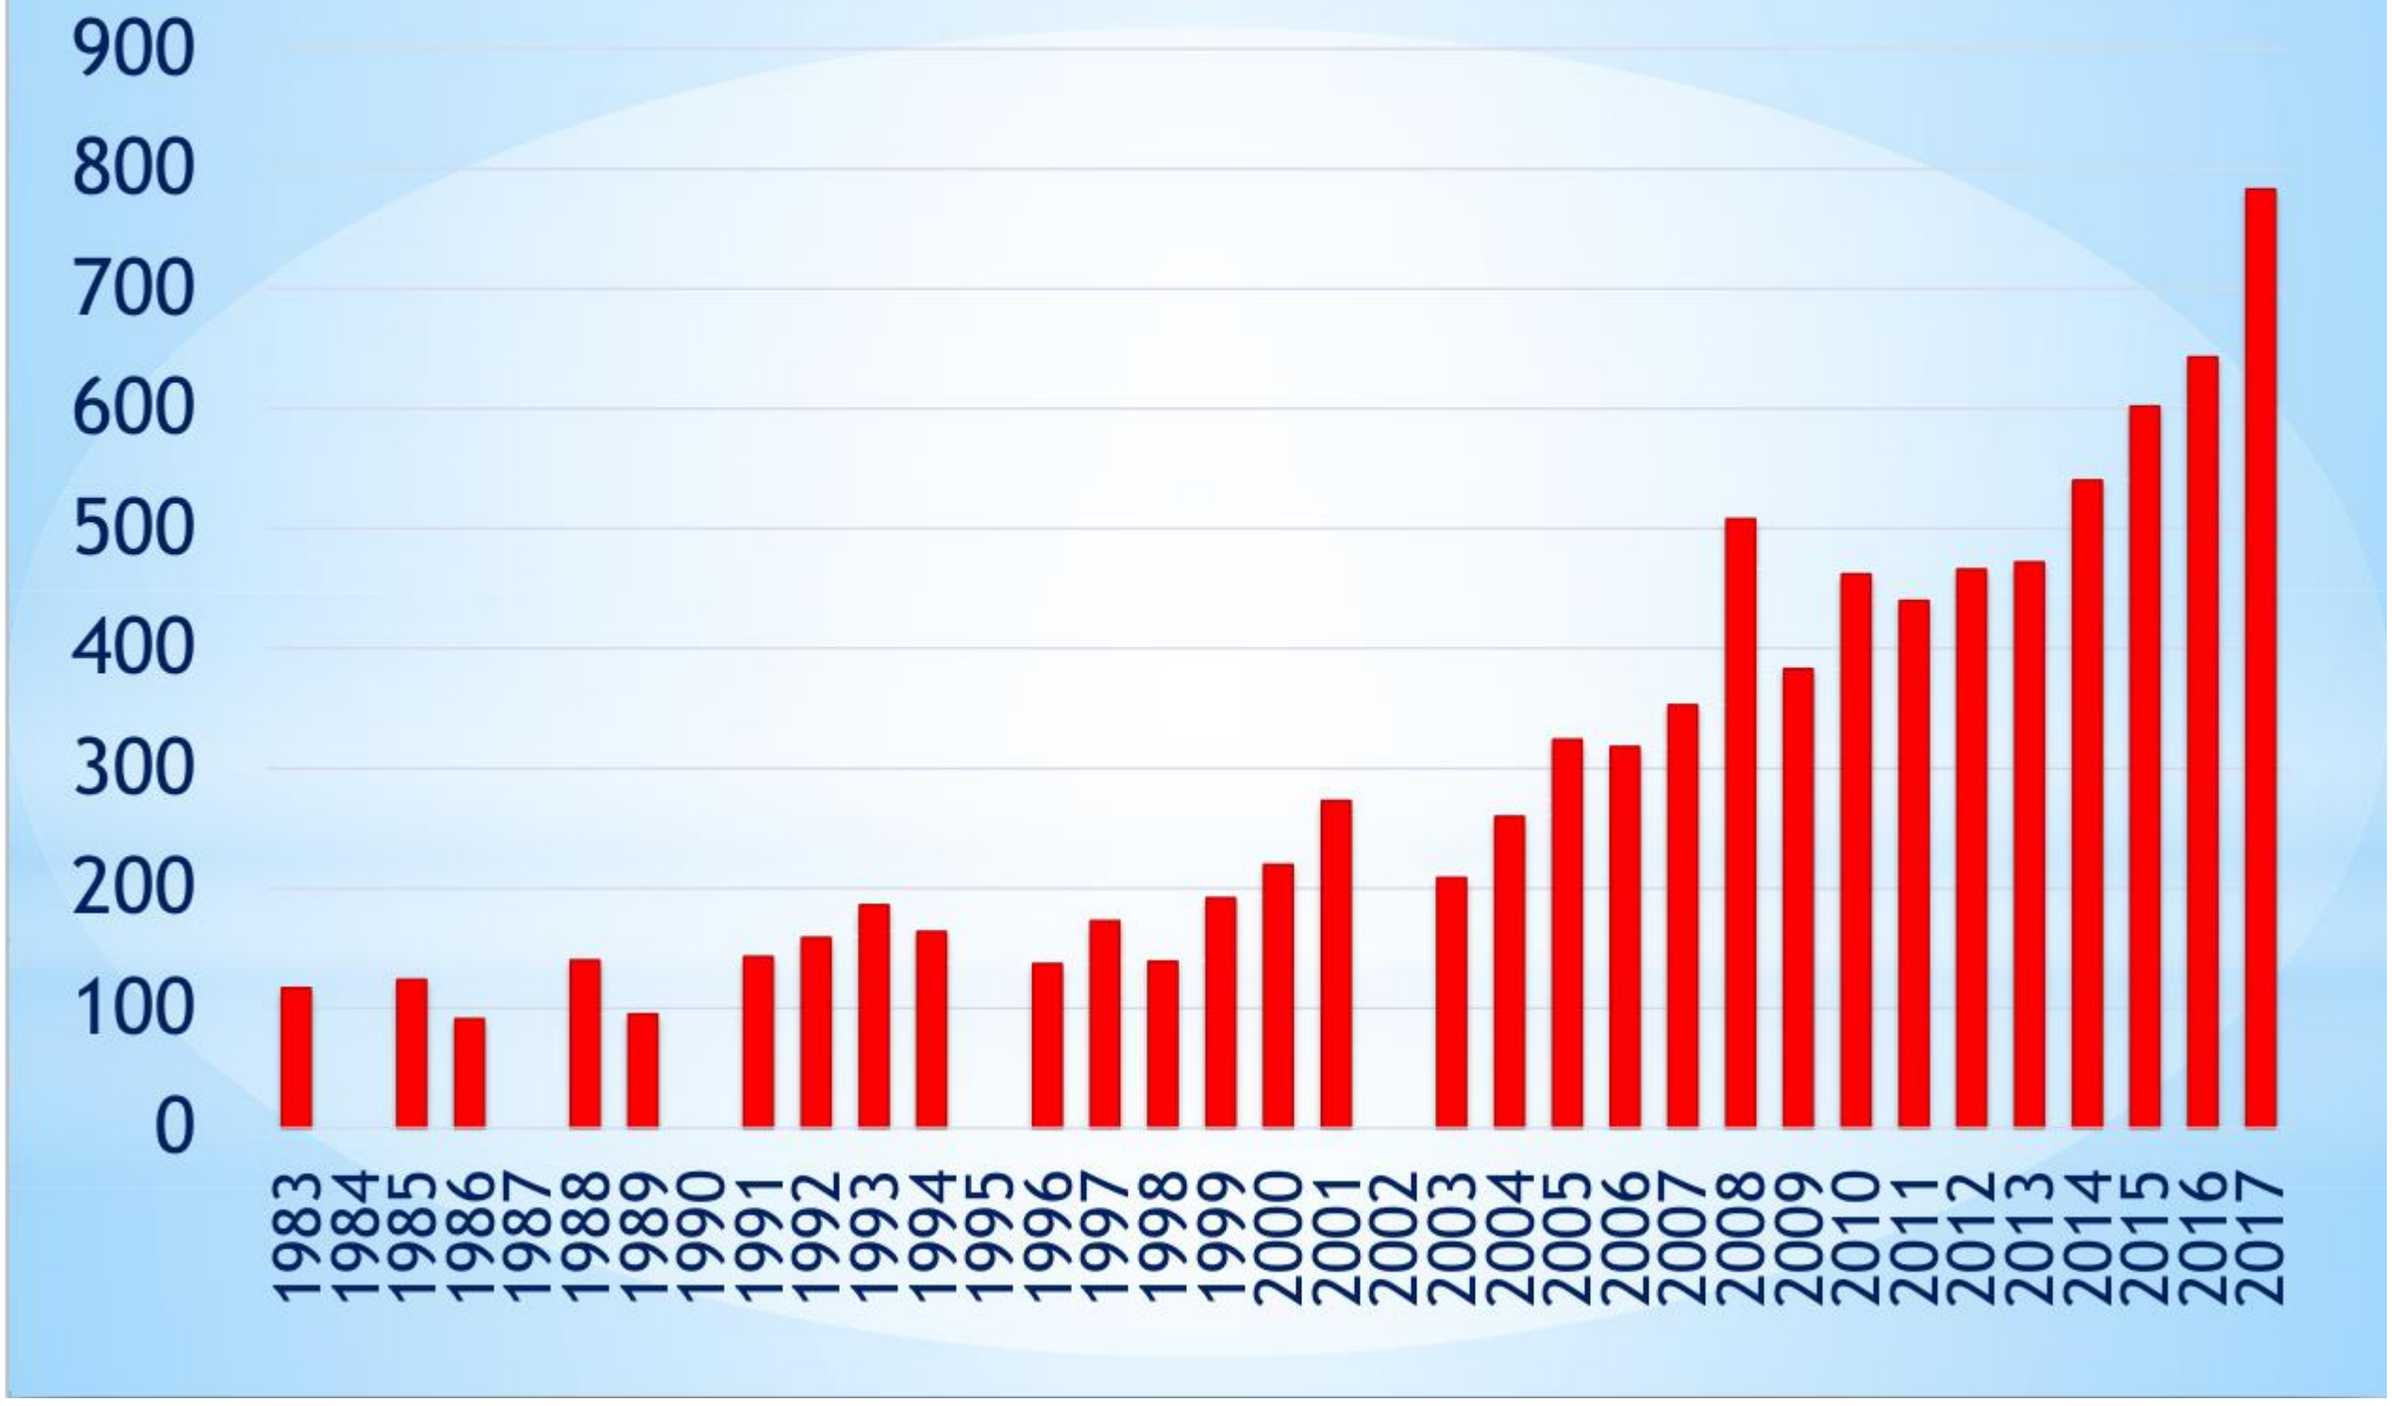

CVPR GROWTH

Original slide courtesy of CVPR 2016

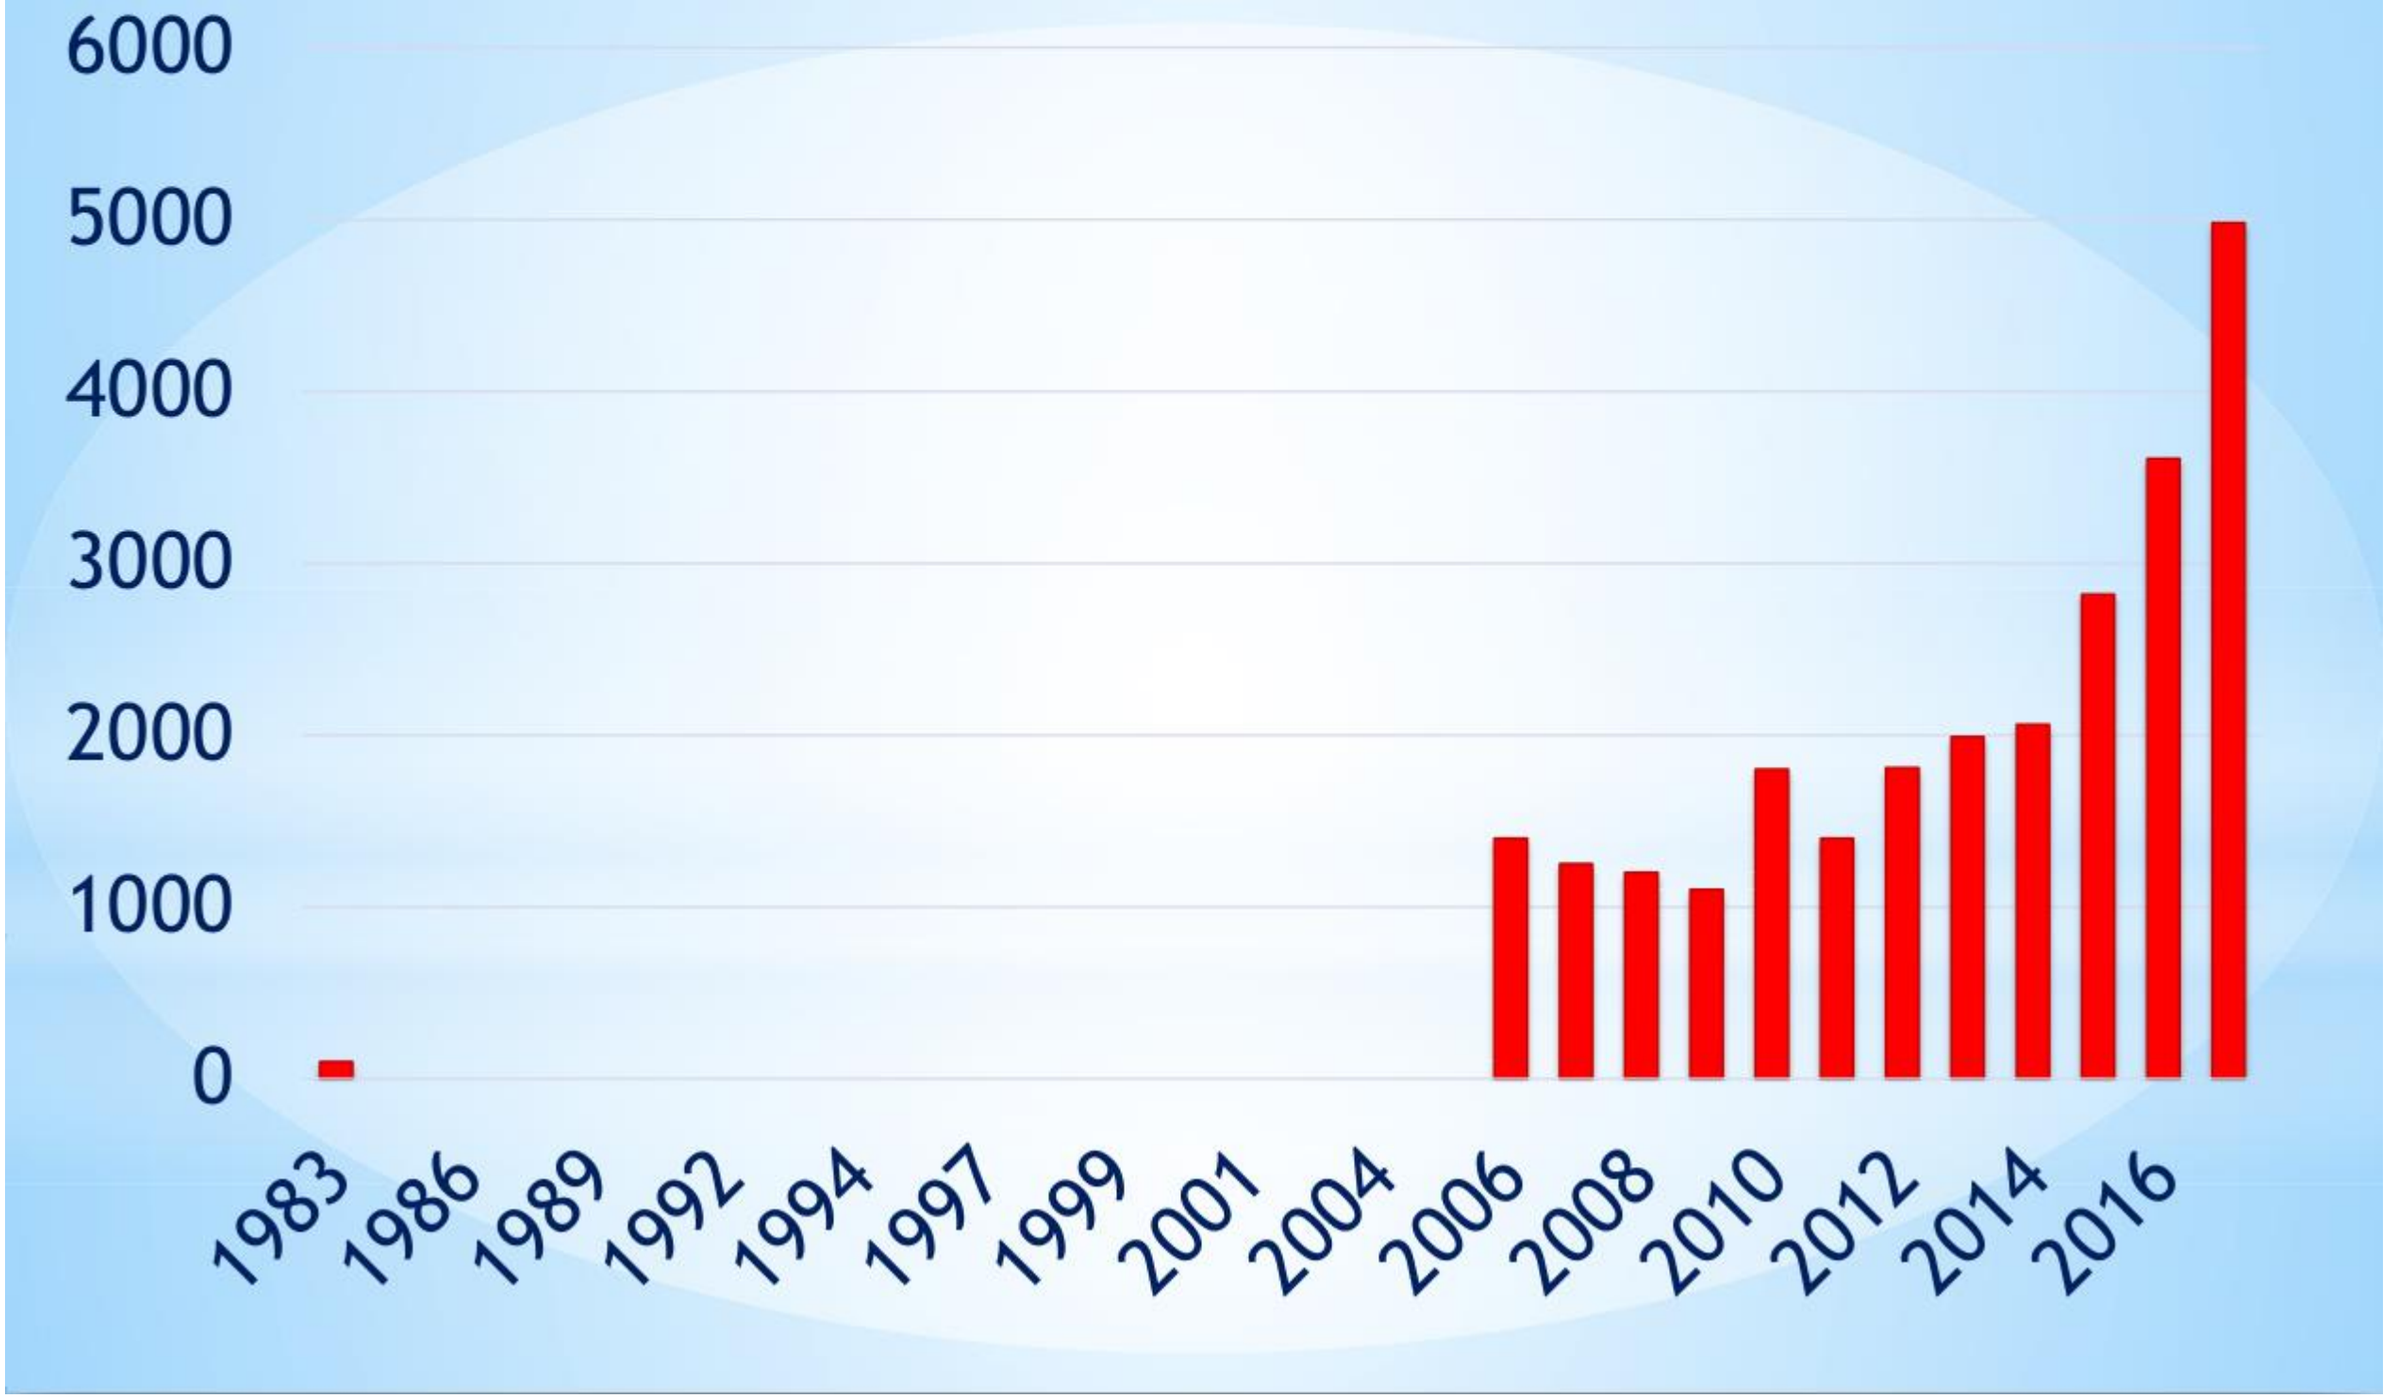

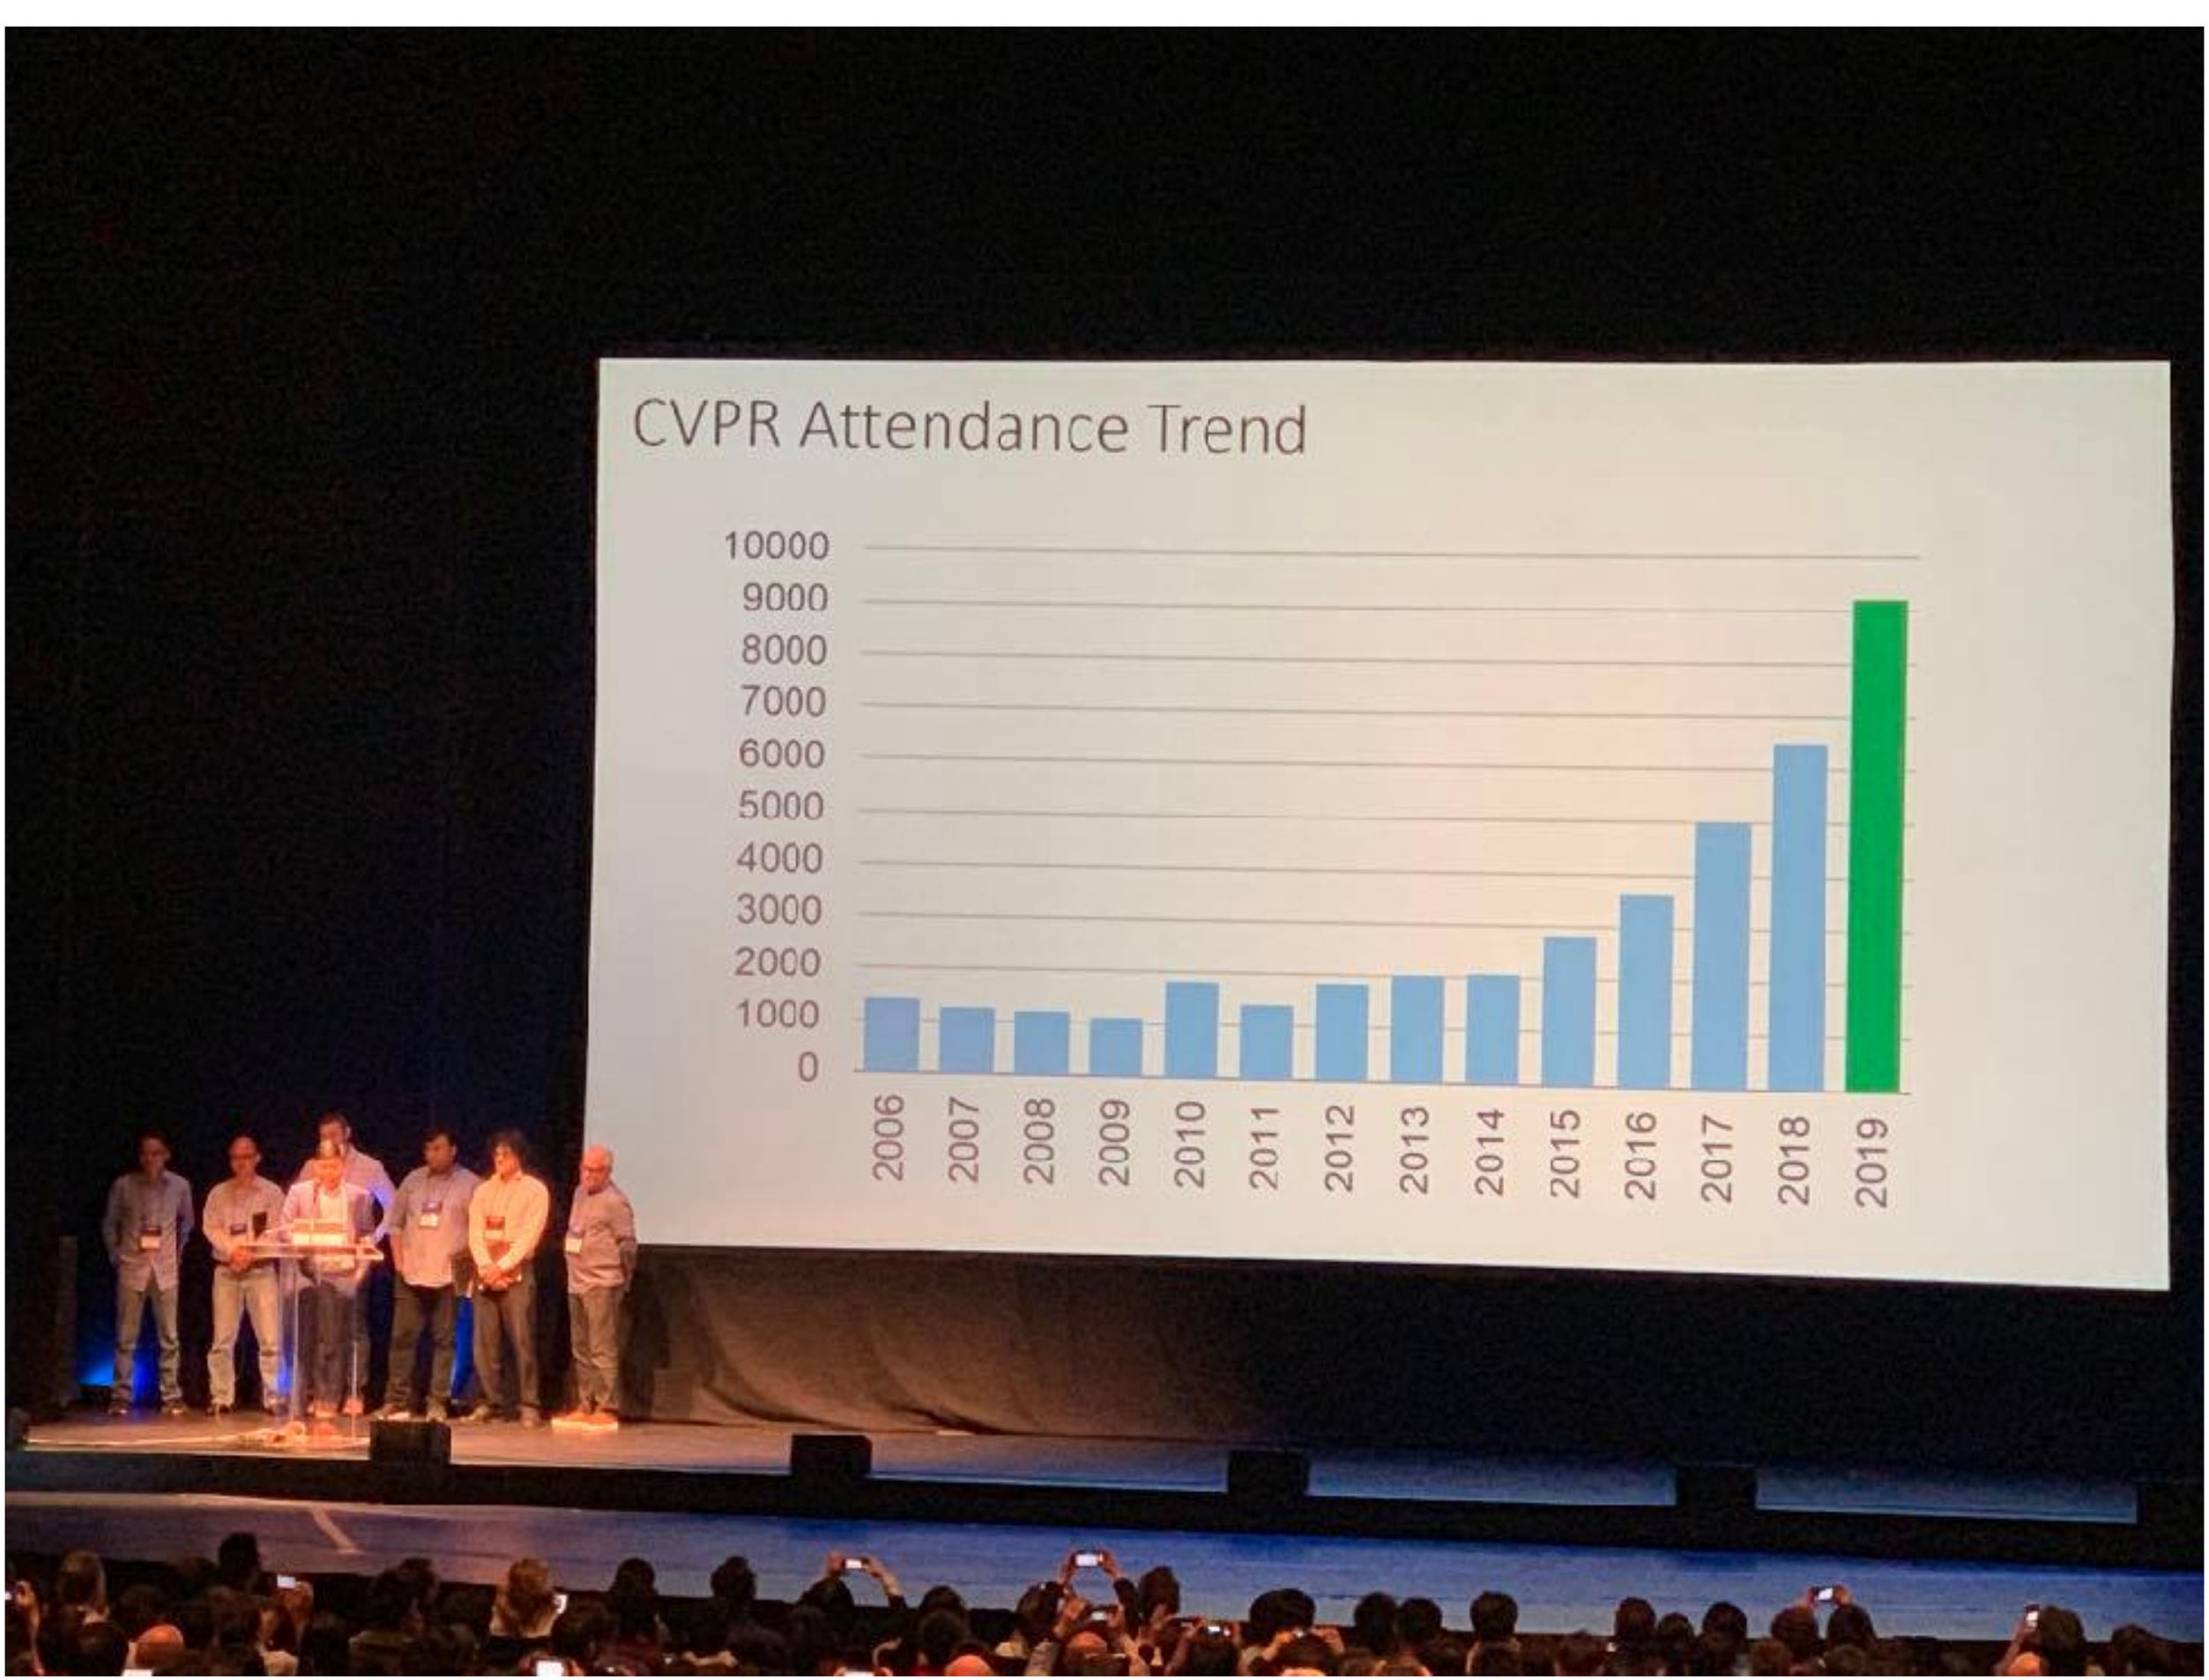
CVPR Attendance Trend
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019

Computer vision at CMU

Dedicated courses for each subject we cover in this class:

- Physics-based Methods in Vision
- Geometry-based Methods in Computer Vision
- Computational Photography
- Visual Learning and Recognition
- Statistical Techniques in Robotics
- Sensors and sensing

… plus an entire department's worth of ML courses.

CVPR 2019: CMU was the second most common academic affiliation among authors

(can you guess the first?)

Master in Computer Vision at CMU

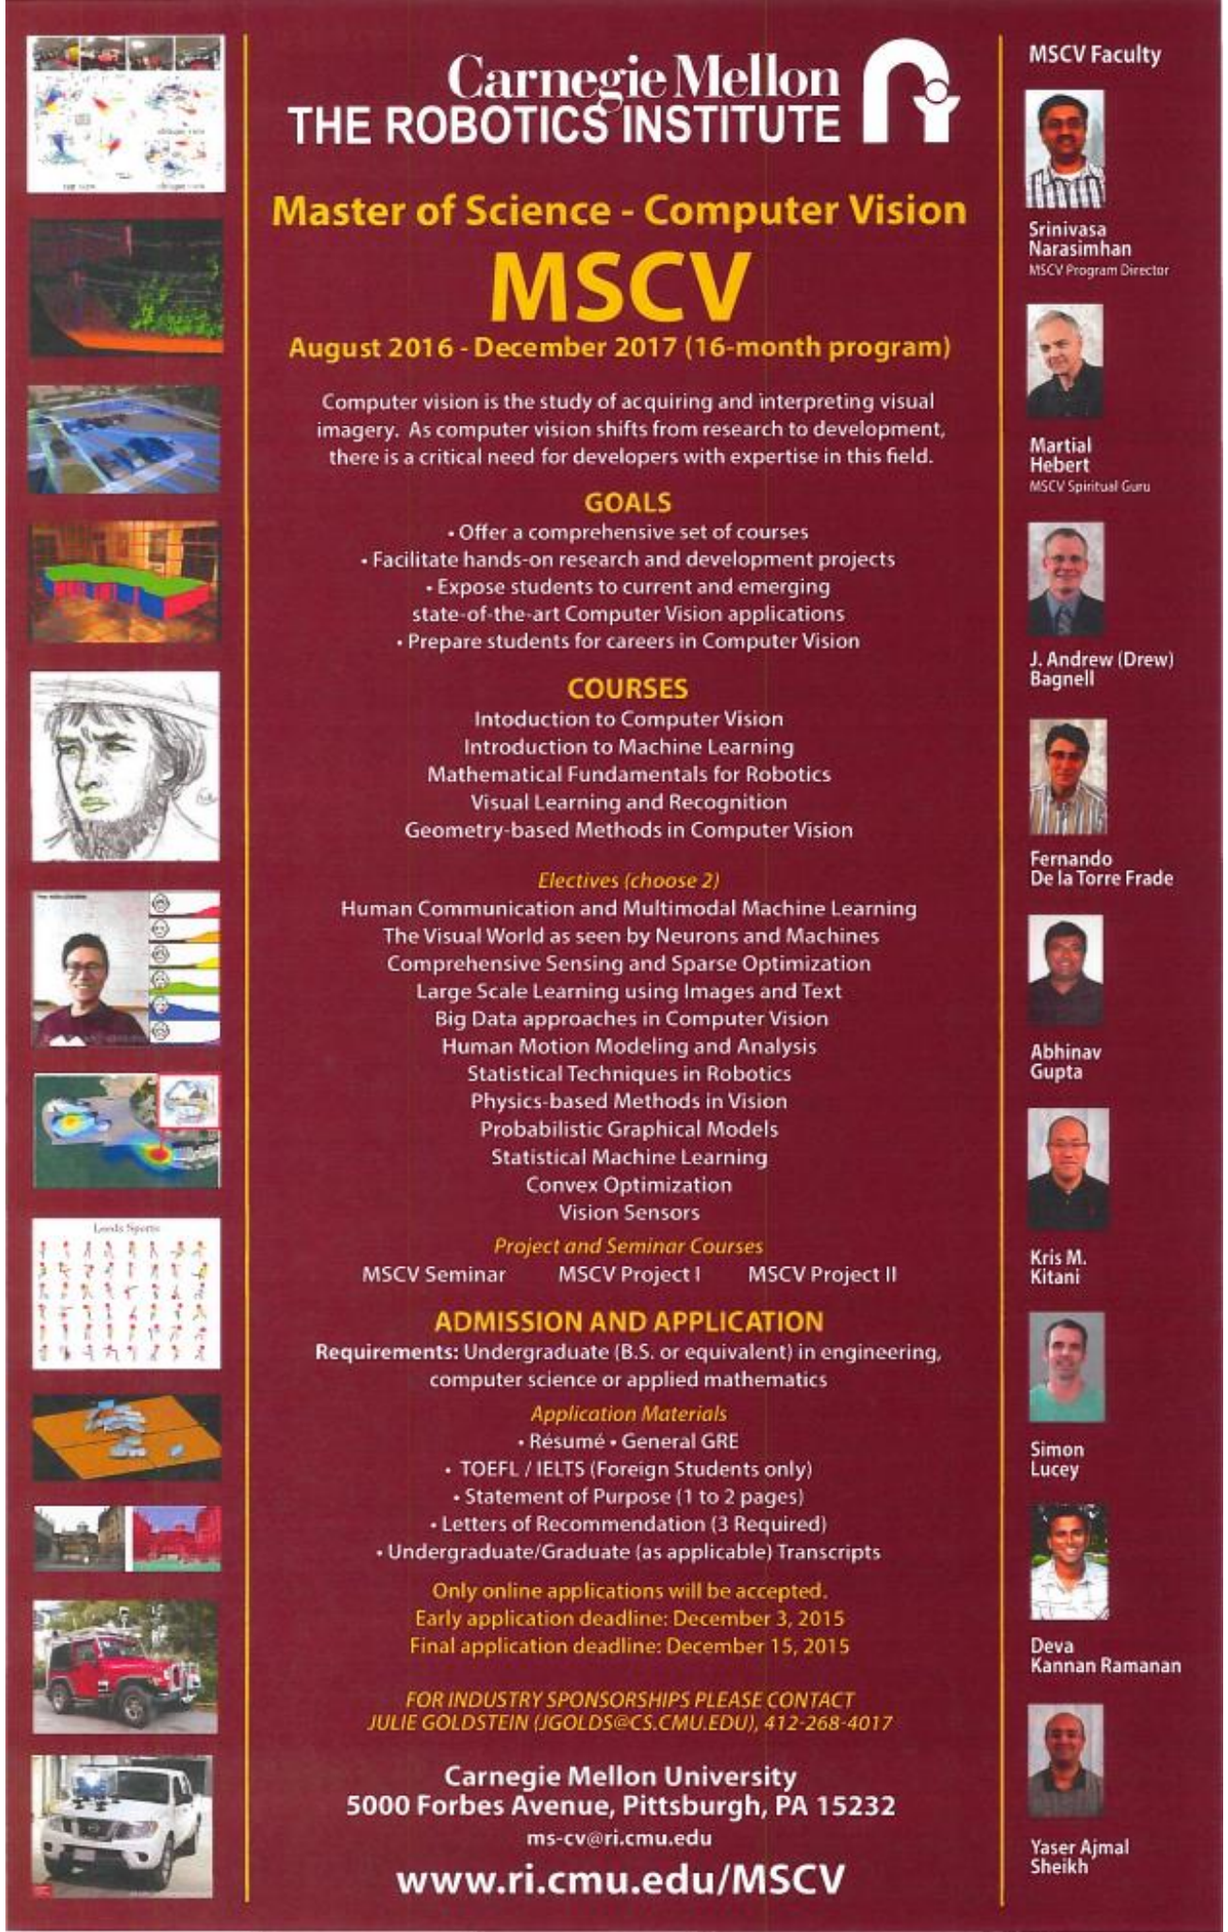

Course logistics

Website

http://www.cs.cmu.edu/~16385/

(includes links to Canvas and Piazza)

Assignments Canvas

https://canvas.cmu.edu/courses/14118

Discussion&notes piazza

https://piazza.com/class/k53x5h48my264d

(you should sign up here on your own)

Topics to be covered

Image processing:

- Basics of filtering.
- Image pyramids.
- Gradients and lines.
- Hough transforms.

Topics to be covered

Feature detection and correspondences:

- Corner detection.
- SIFT et al.
- Feature descriptors.
- RANSAC.

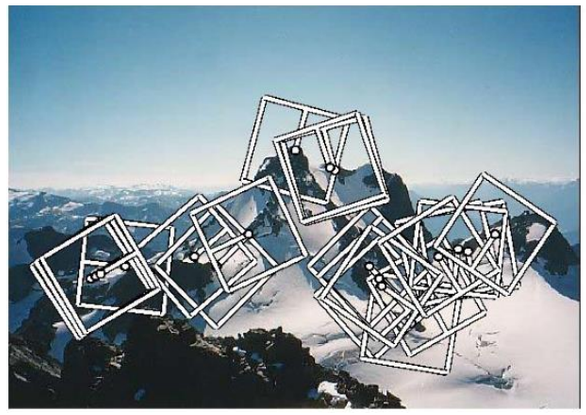
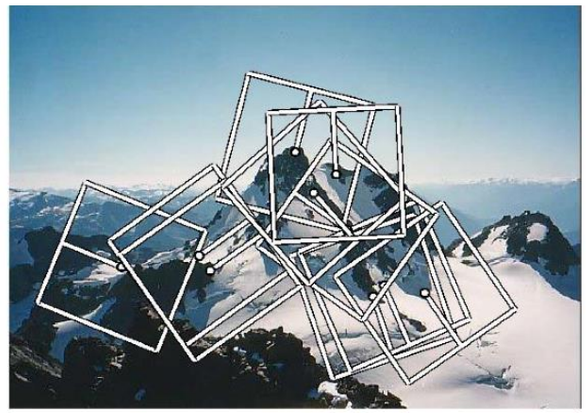
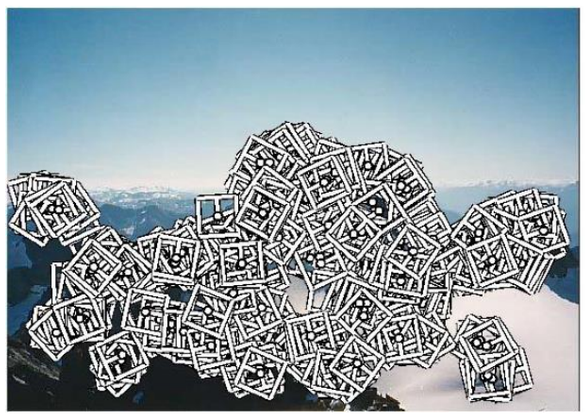
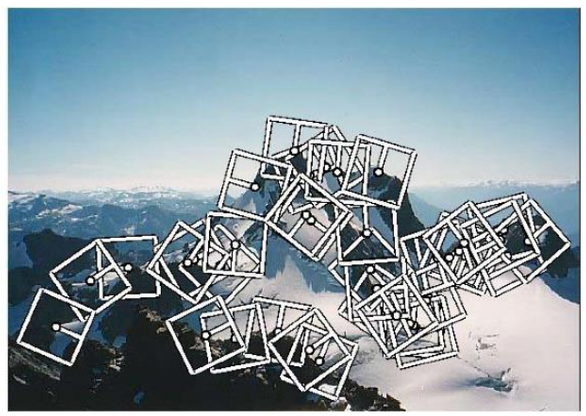

Topics to be covered

Transformations and geometry:

- Homographies and image alignment.
- Camera models.
- Fundamental matrix.
- Epipolar geometry and stereo.
- Structure from motion.

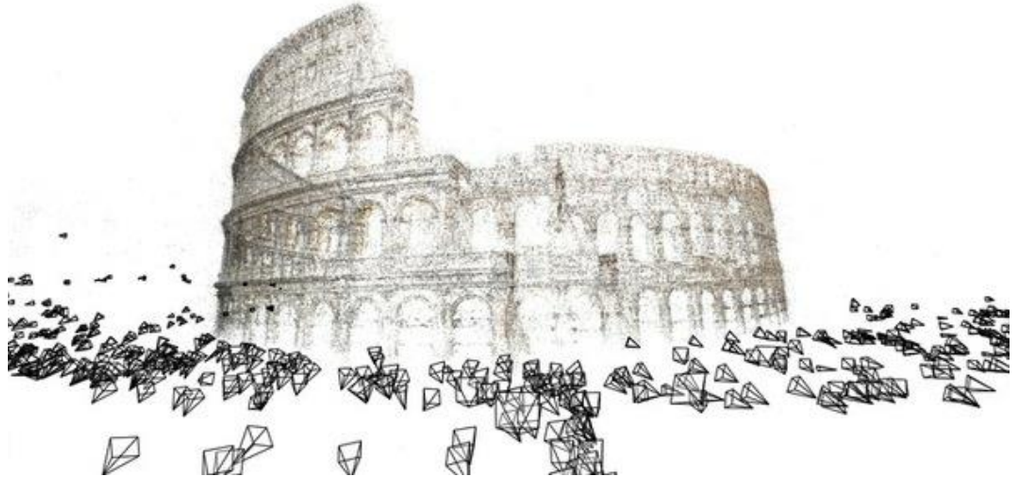

Topics to be covered

Physics-based vision:

- Reflectance and image formation.
- Radiometry.
- Shape from shading.
- Photometric stereo.
- Color.

Topics to be covered

Objects, faces, and learning:

- Basics of probability.
- K-means, KNN, PCA, SVM.
- Bag of words.
- Viola-Jones face detection.
- Perceptron, backpropagation.
- Convolutional neural networks.

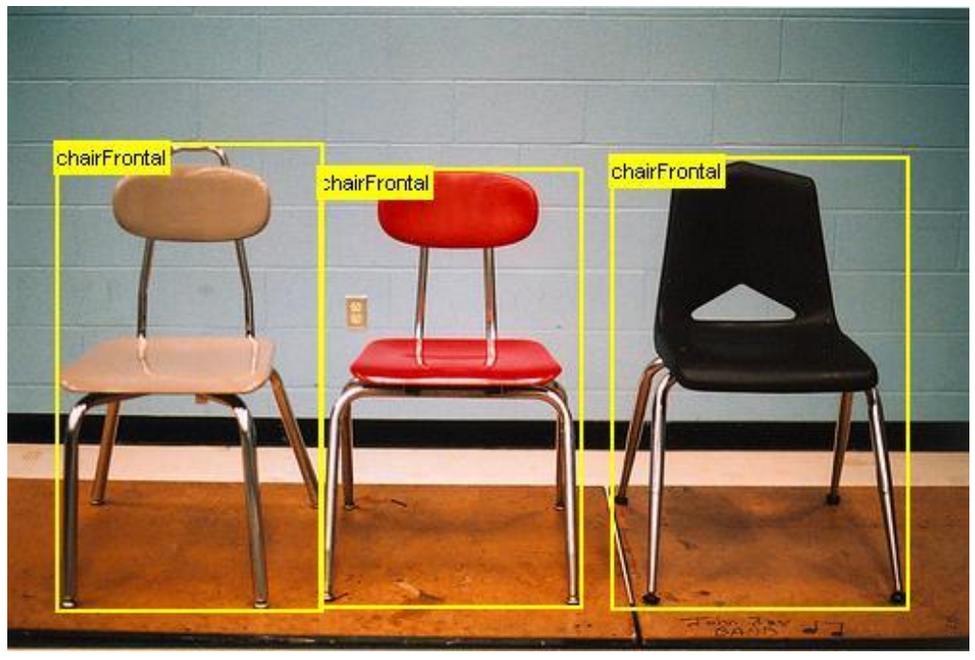

Topics to be covered

Dealing with motion:

- Optical flow (LK, HS).
- Image registration.
- Kalman Filtering.
- Tracking (KLT, Mean-Shift).

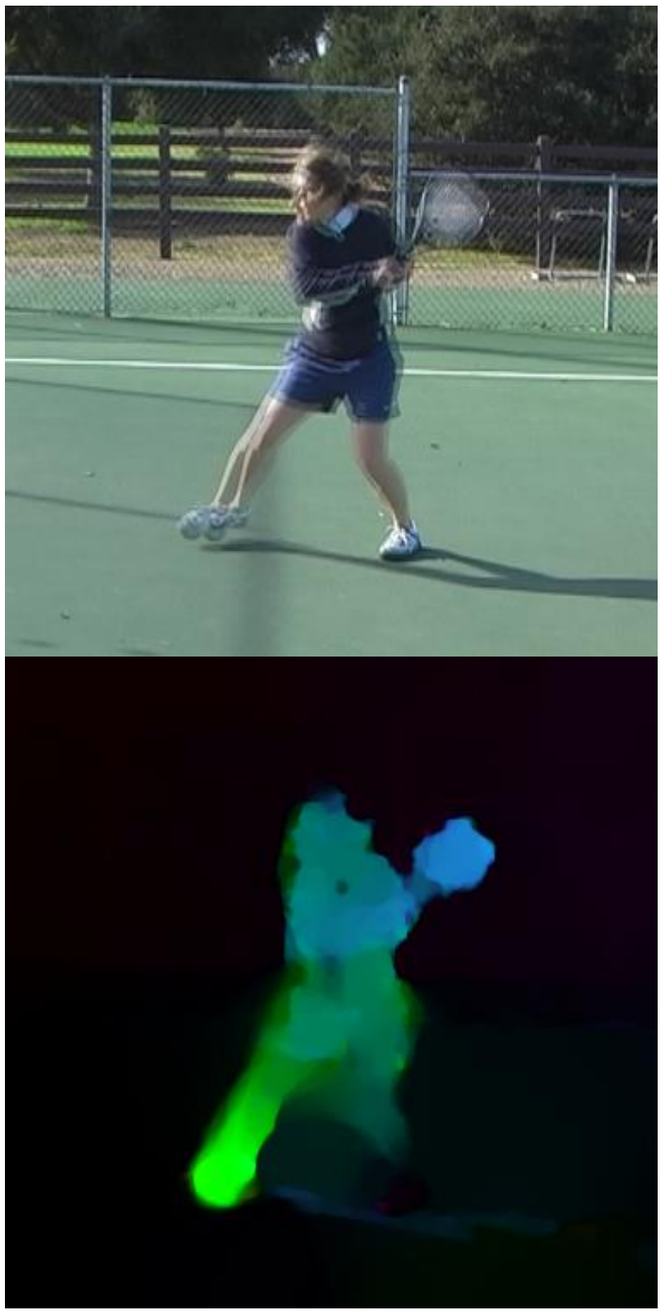

Topics to be covered

Special topics:

- Computational photography.
- ???

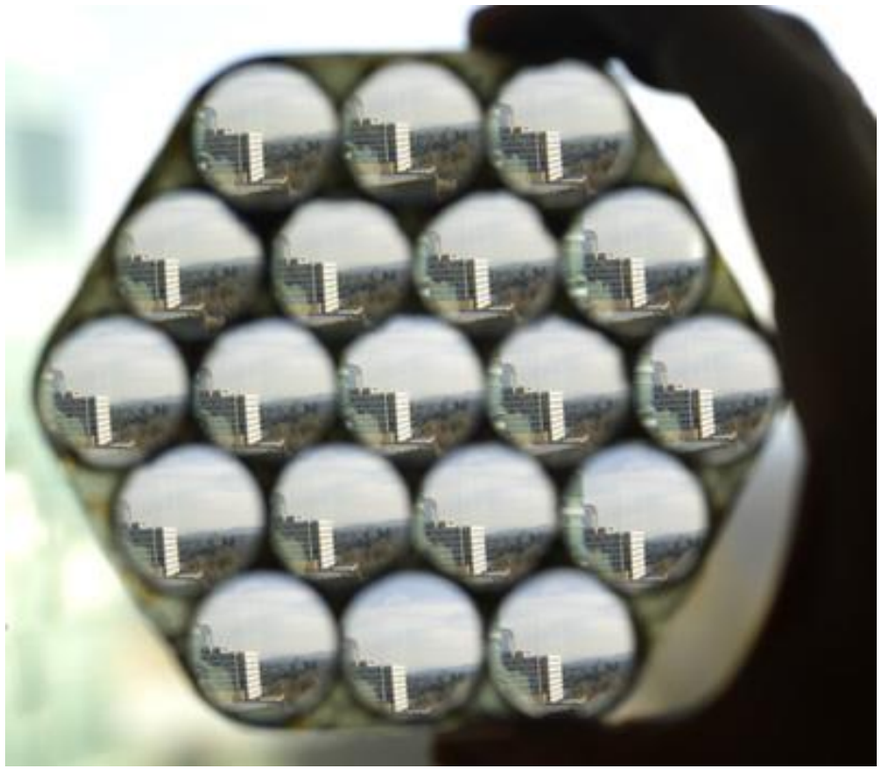

Grading

- Seven two-week programming assignments: 70%
- Twelve weekly take-home quizzes: 27%
- Class and Piazza participation: 3%

Take-home quizzes:

- New this year.
- Two-three theory questions.
- Replace mid-term and final.

Participation:

- Be around for lectures.
- Post on Piazza discussions.
- Ask and answer questions.

Programming assignments

Assignment 1 Hough Transform
Assignment 2 Homography
Assignment 3 Stereo
Assignment 4 Photometric Stereo
Assignment 5 Bag of Words
Assignment 6 Convolutional Neural Nets
Assignment 7 Lucas-Kanade Tracking

- a lot of programming in Matlab and Python.
- hours and hours of programming.
- days and days of debugging.
- generous grading policy (like grad school)
- take advantage of extra credit

Programming assignments

Assignment 1 Hough Transform
Assignment 2 Homography
Assignment 3 Stereo
Assignment 4 Photometric Stereo
Assignment 5 Bag of Words
Assignment 6 Convolutional Neural Nets
Assignment 7 Lucas-Kanade Tracking

seriously, a lot of

programming, so start early!

- a lot of programming in Matlab and Python.
- hours and hours of programming.
- days and days of debugging.
- generous grading policy (like grad school)
- take advantage of extra credit

Schedule

- Tentative schedule on course website.
- Likely to change.
- Always check course website and Piazza for updates!

Date	Topics	Slides	Assignments
M, Jan 13	Introduction		
W, Jan 15	Image filtering		
M, Jan 20	**No class (Martin Luther King day)**		
W, Jan 22	Image pyramids and Fourier transform		PA1 out
M, Jan 27	Hough transform		TQ1 out
W, Jan 28	Feature and corner detection		
M, Feb 3	Feature descriptors and matching		TQ1 due, TQ2 out
W, Feb 5	2D transformations		PA1 due, PA2 out
M, Feb 10	Image homographies		TQ2 due, TQ3 out
W, Feb 12	Camera models		
M, Feb 17	Two-view geometry		TQ3 due, TQ4 out
W, Feb 19	Stereo		PA2 due, PA3 out
M, Feb 24	Structure from motion		TQ4 due, TQ5 out
W, Feb 26	Radiometry and reflectance		
M, Mar 2	Photometric stereo and shape from shading		TQ5 due, TQ6 out
W, Mar 4	Color		PA3 due, PA4 out
M, Mar 9	**No class (spring break)**		
W, Mar 11	**No class (spring break)**		
M, Mar 16	Image processing pipeline		TQ6 due, TQ7 out
W, Mar 18	Introduction to recognition		PA4 due, PA5 out
M, Mar 23	Bag of works		TQ7 due, TQ8 out
W, Mar 25	Neural networks		
M, Mar 30	Convolutional neural networks		TQ8 due, TQ9 out
W, Apr 1	Optimization		PA5 due, PA6 out
M, Apr 6	Faces		TQ9 due, TQ10 out
W, Apr 8	Optical flow		
M, Apr 13	Alignment		TQ10 due, TQ11 out
W, Apr 15	Tracking		PA6 due, PA7 out
M, Apr 20	Temporal models and SLAM		TQ11 due, TQ12 out
W, Apr 22	Graph-based methods		
M, Apr 27	Segmentation		TQ12 due
W, Apr 29	Wrap-up and discussion		PA7 due

Leniency

Late days for programming assignments:

- 10% reduction of points per late day
- 6 free late days total
- use them wisely… save for later (harder) assignments!

Option to skip take-home quizzes:

- you only need to submit 9 out of 12 quizzes
- late quizzes will not be graded

Book

We will be posting readings after each lecture

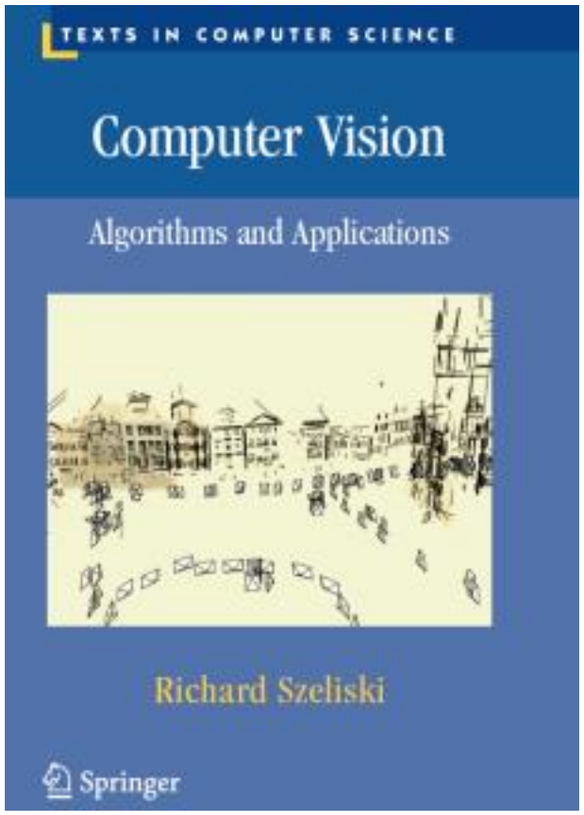

PDF online

http://szeliski.org/Book/

Prerequisites

We assume familiarity with calculus, linear algebra, basic probability, and programming.

Formal prerequisites:

- "Mathematical Foundations of Electrical Engineering" (18-202) and "Principles of Imperative Computation" (15-122)

OR

- "Matrix Algebra with Applications" (21-240) and "Matrices and Linear Transformations" (21-241) and "Calculus in Three Dimensions" (21-259) and "Principles of Imperative Computation" (15-122)

If you are missing a prerequisite but still want to enroll, let me know and we'll discuss it.

Contact information and office hours

- Feel free to email us about administrative questions.
  - please use [16385] 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).

<u>Yannis will announce office hours for this week.</u>

Please take the course survey before the next lecture!

(also posted on Piazza)