Introduction

16-385 Computer Vision
Spring 2022, Lecture 1

Overview of today's lecture

- Teaching staff introductions
- What is computer vision?
- Course fast-forward and logistics

Teaching staff introductions

Hi!

Matthew O’Toole (Instructor)

Harry Zhang

Arpit Agarwal

What is computer vision?

What a person sees

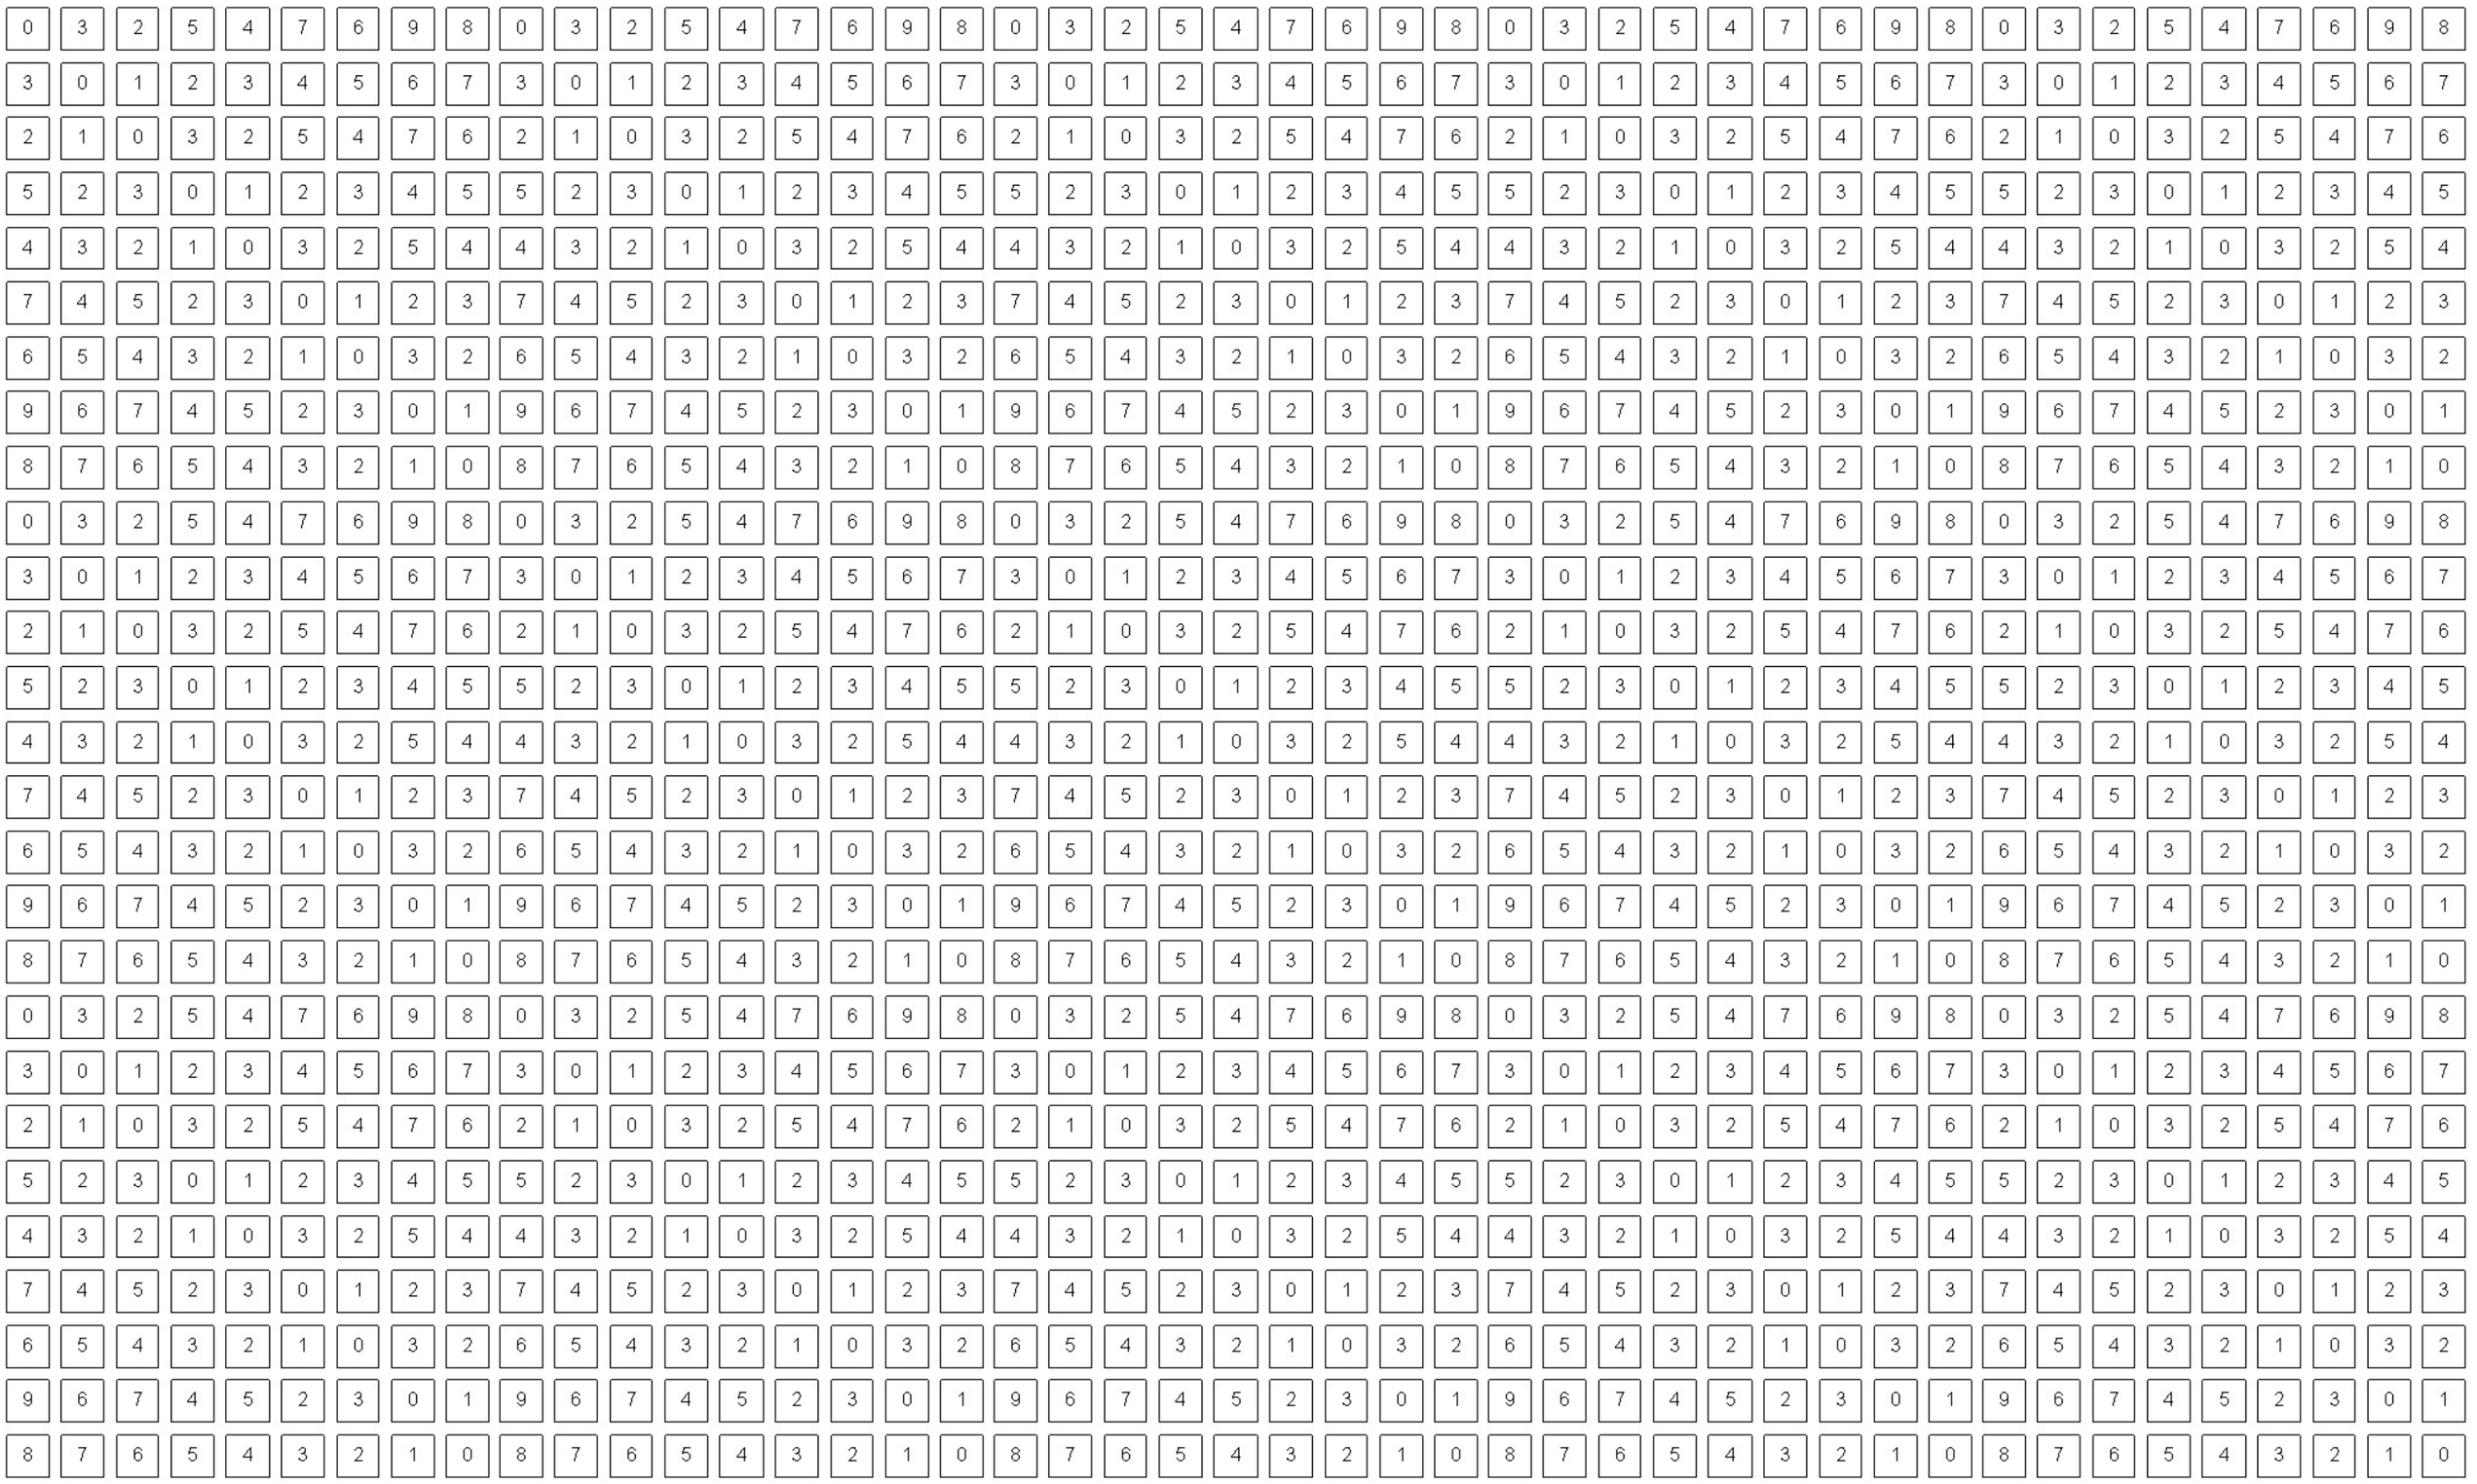

What a computer sees

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

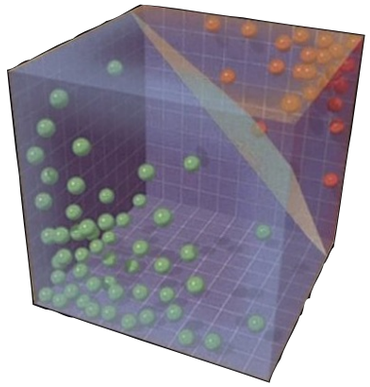

‘fancy math’

output

typical perception pipeline

representation

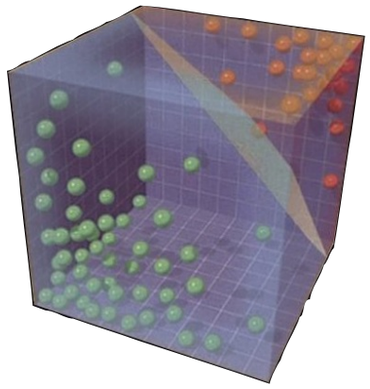

‘fancy math’

output

what should we look at?
(image features)

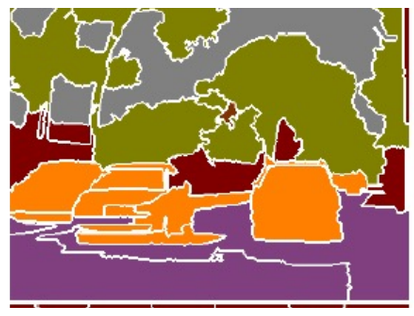

what can we understand?
(semantic segmentation)

typical perception pipeline

representation

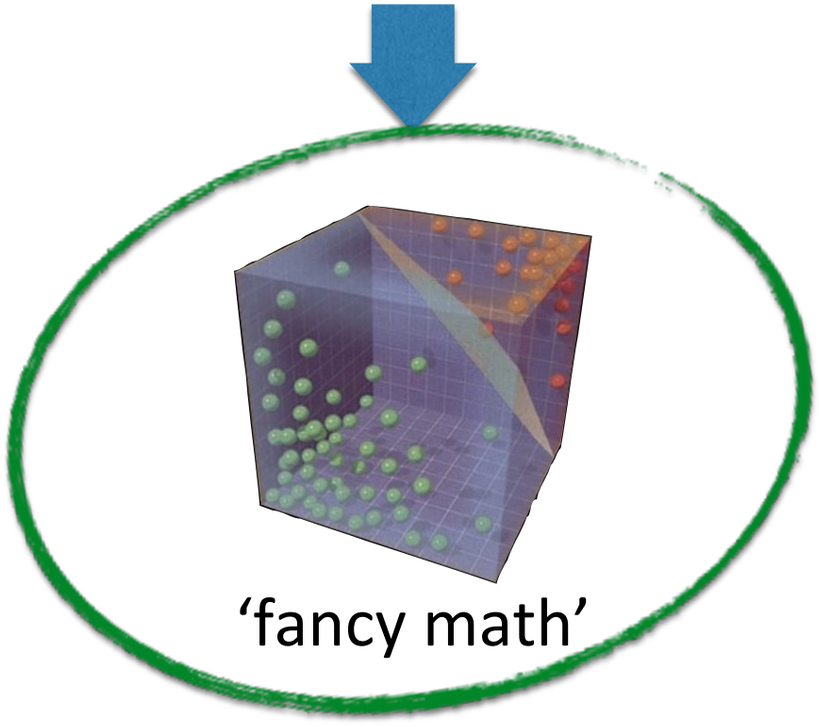

output

easy to get lost in the techniques

what should we look at?
(image features)

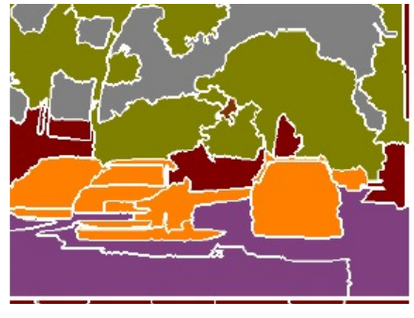

what can we understand?
(semantic segmentation)

typical perception pipeline

representation

what should we look at?
(image features)

The parts that we are most interested in

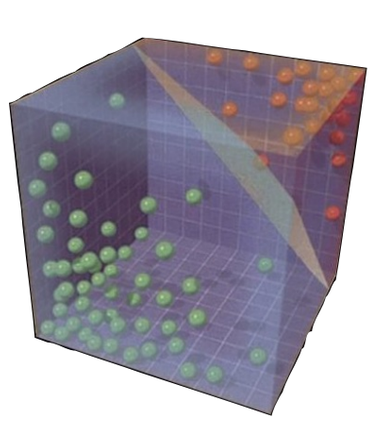

‘fancy math’

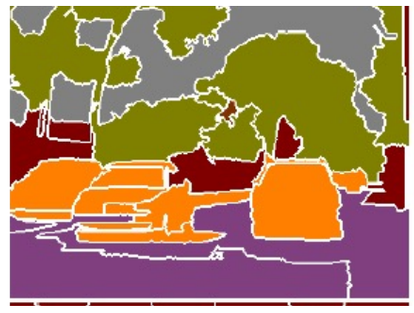

what can we understand?
(semantic segmentation)

output

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

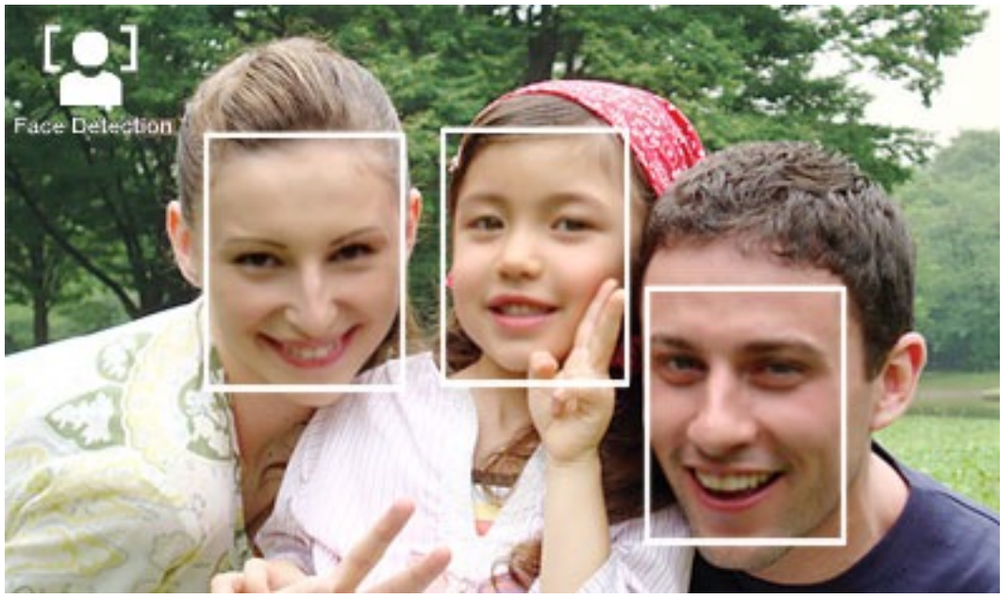

Sony Cyber-shot

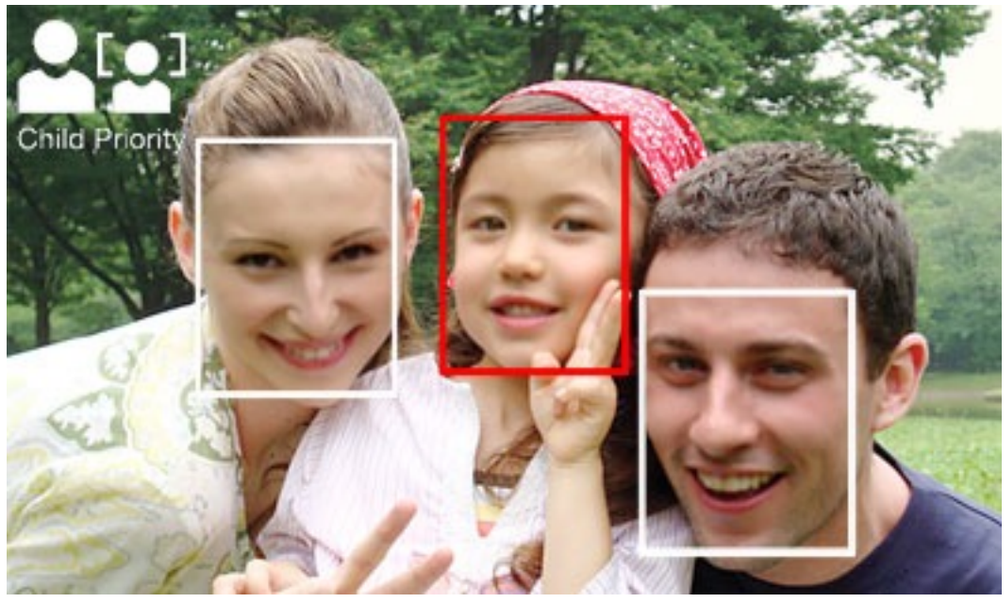

Age recognition

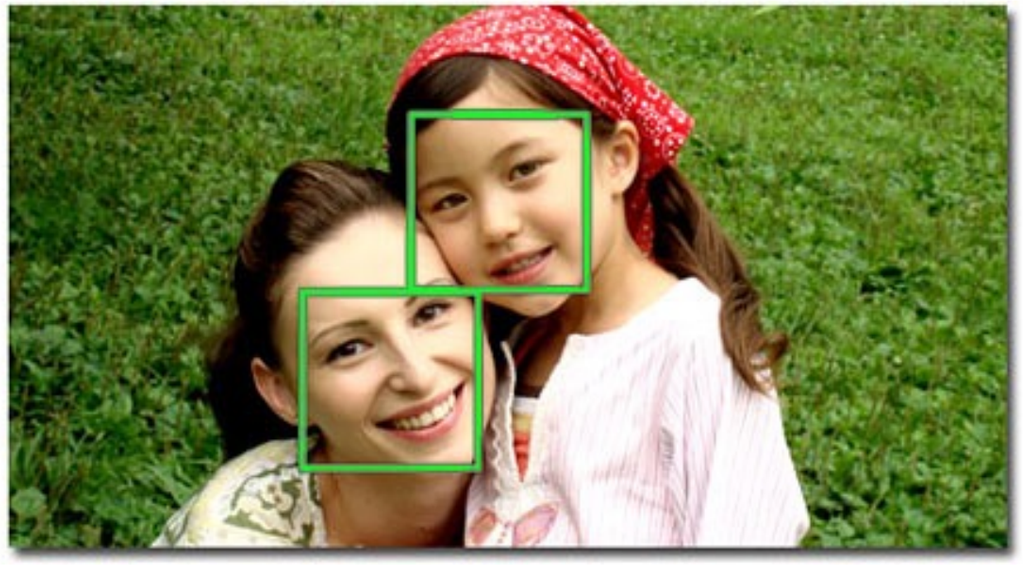

Smile recognition

Face ID

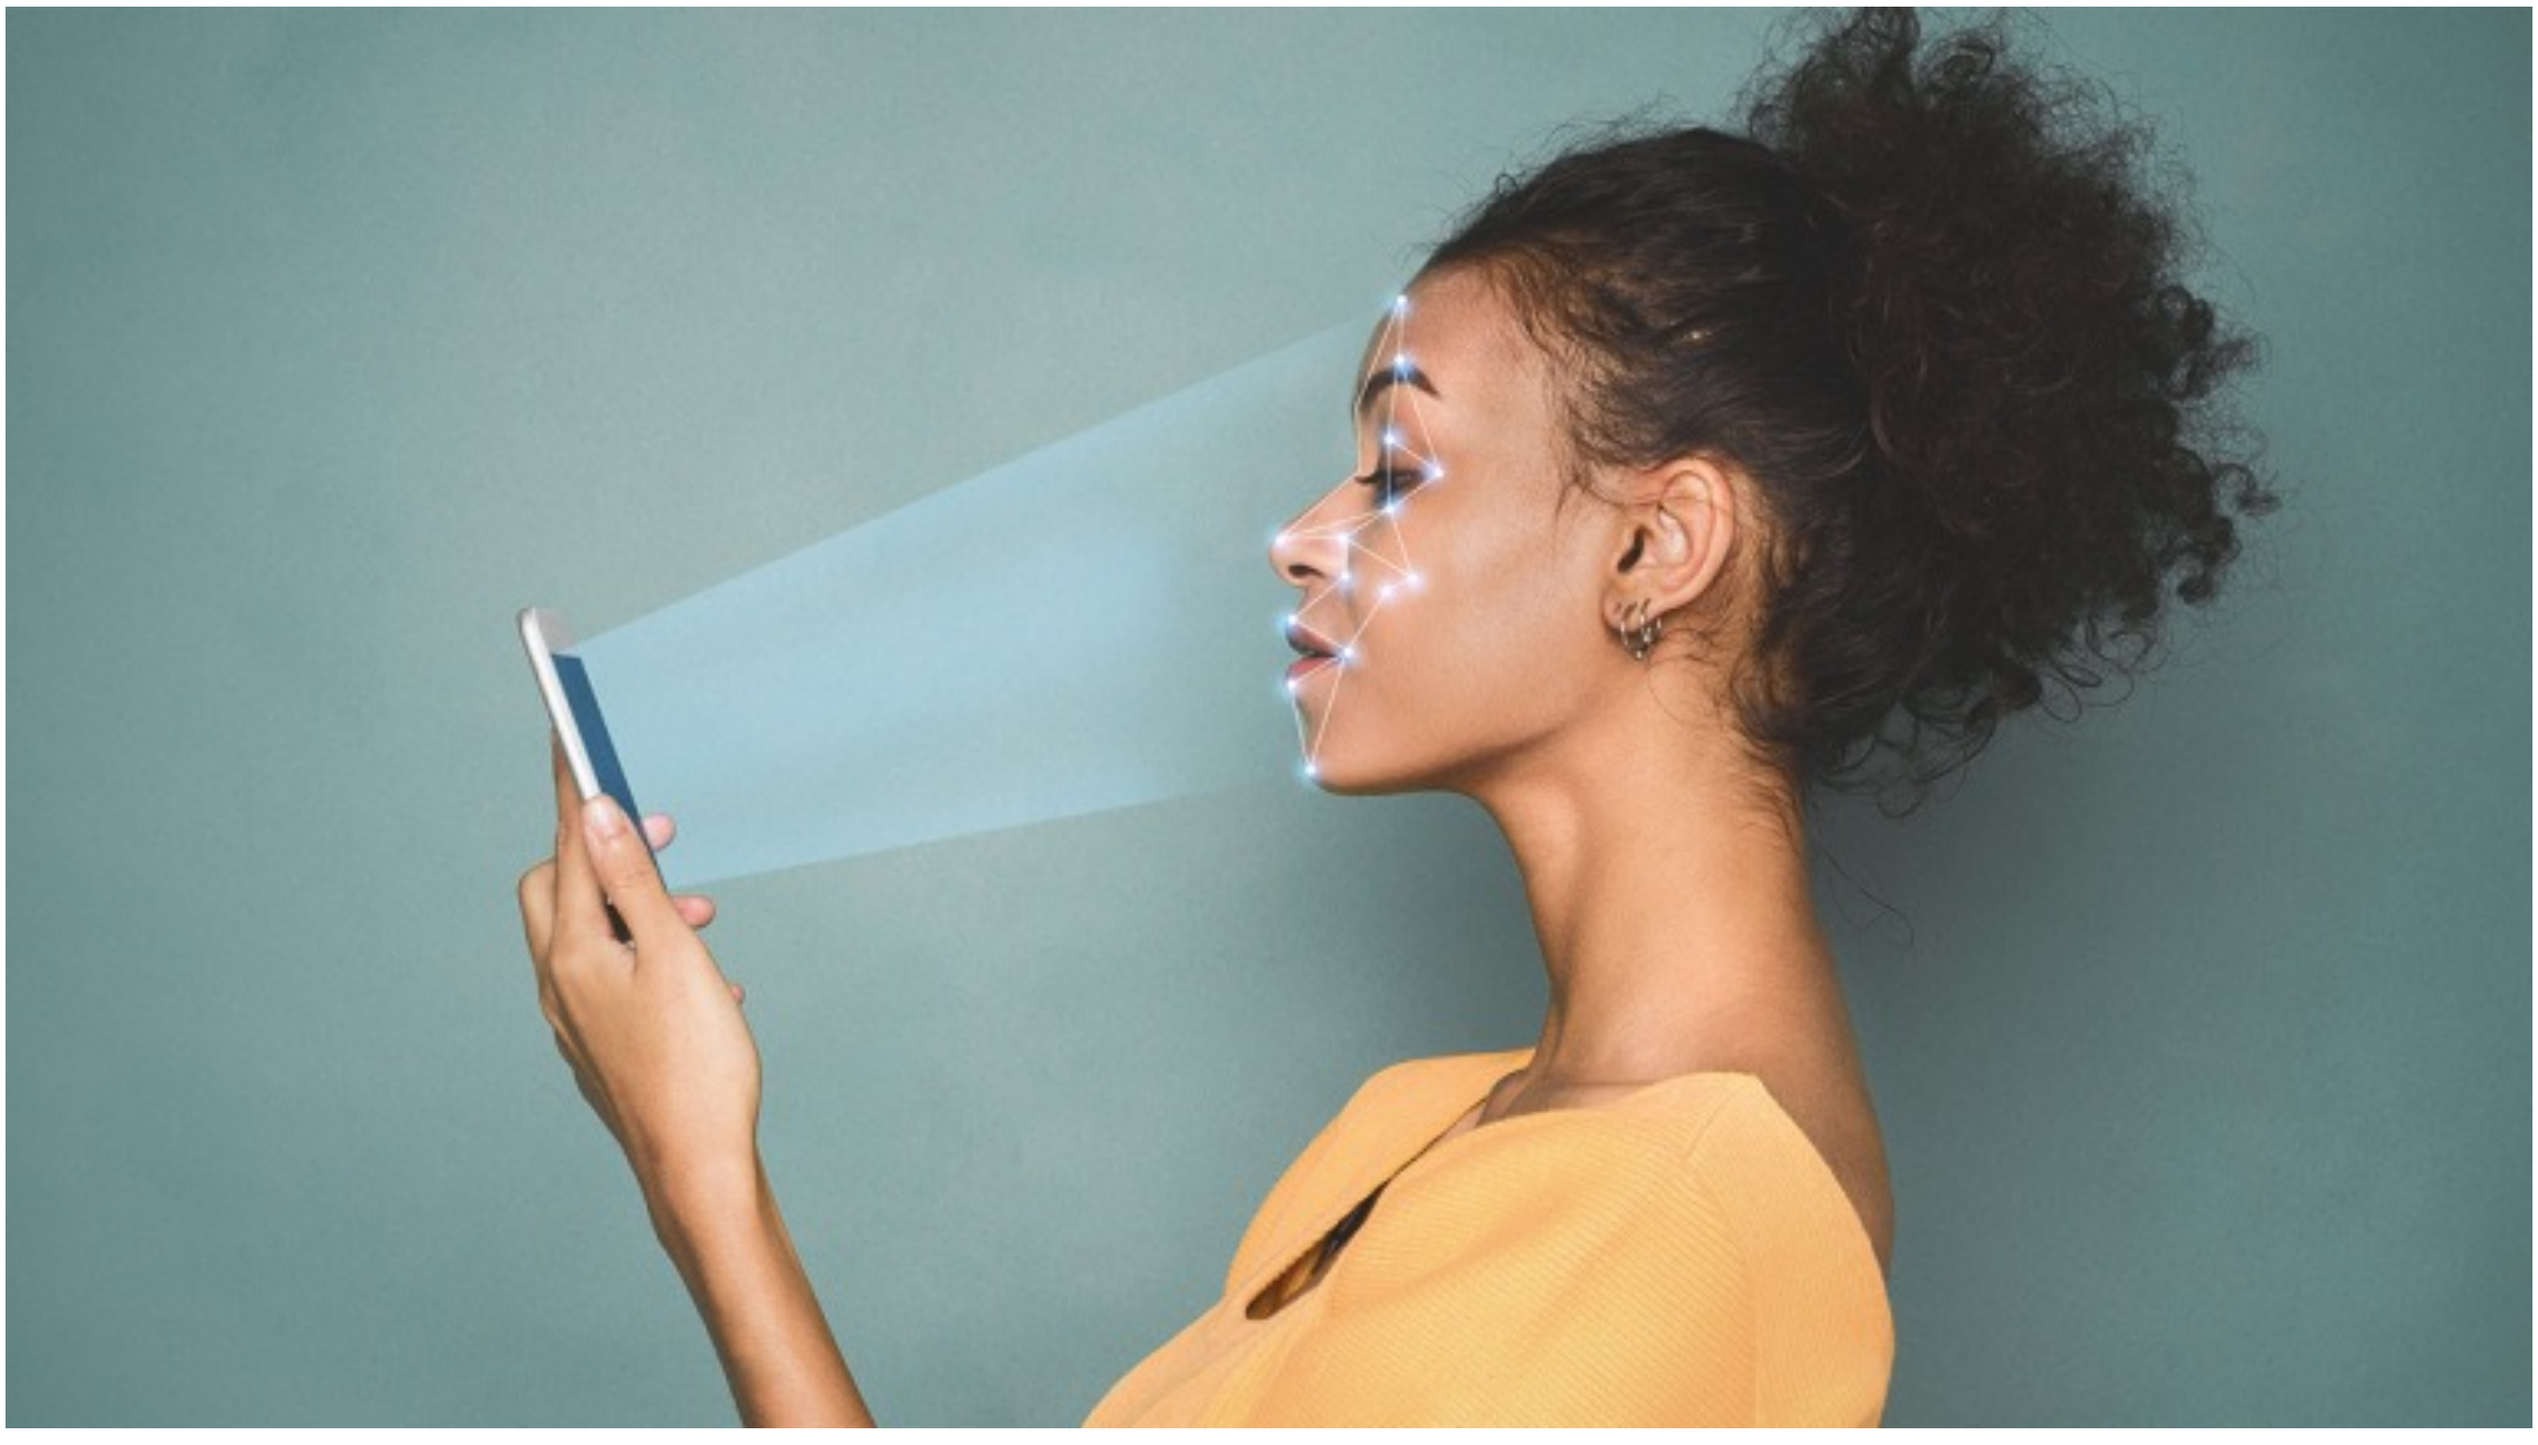

Face ID

Identifying plants

Google translate

First-down line

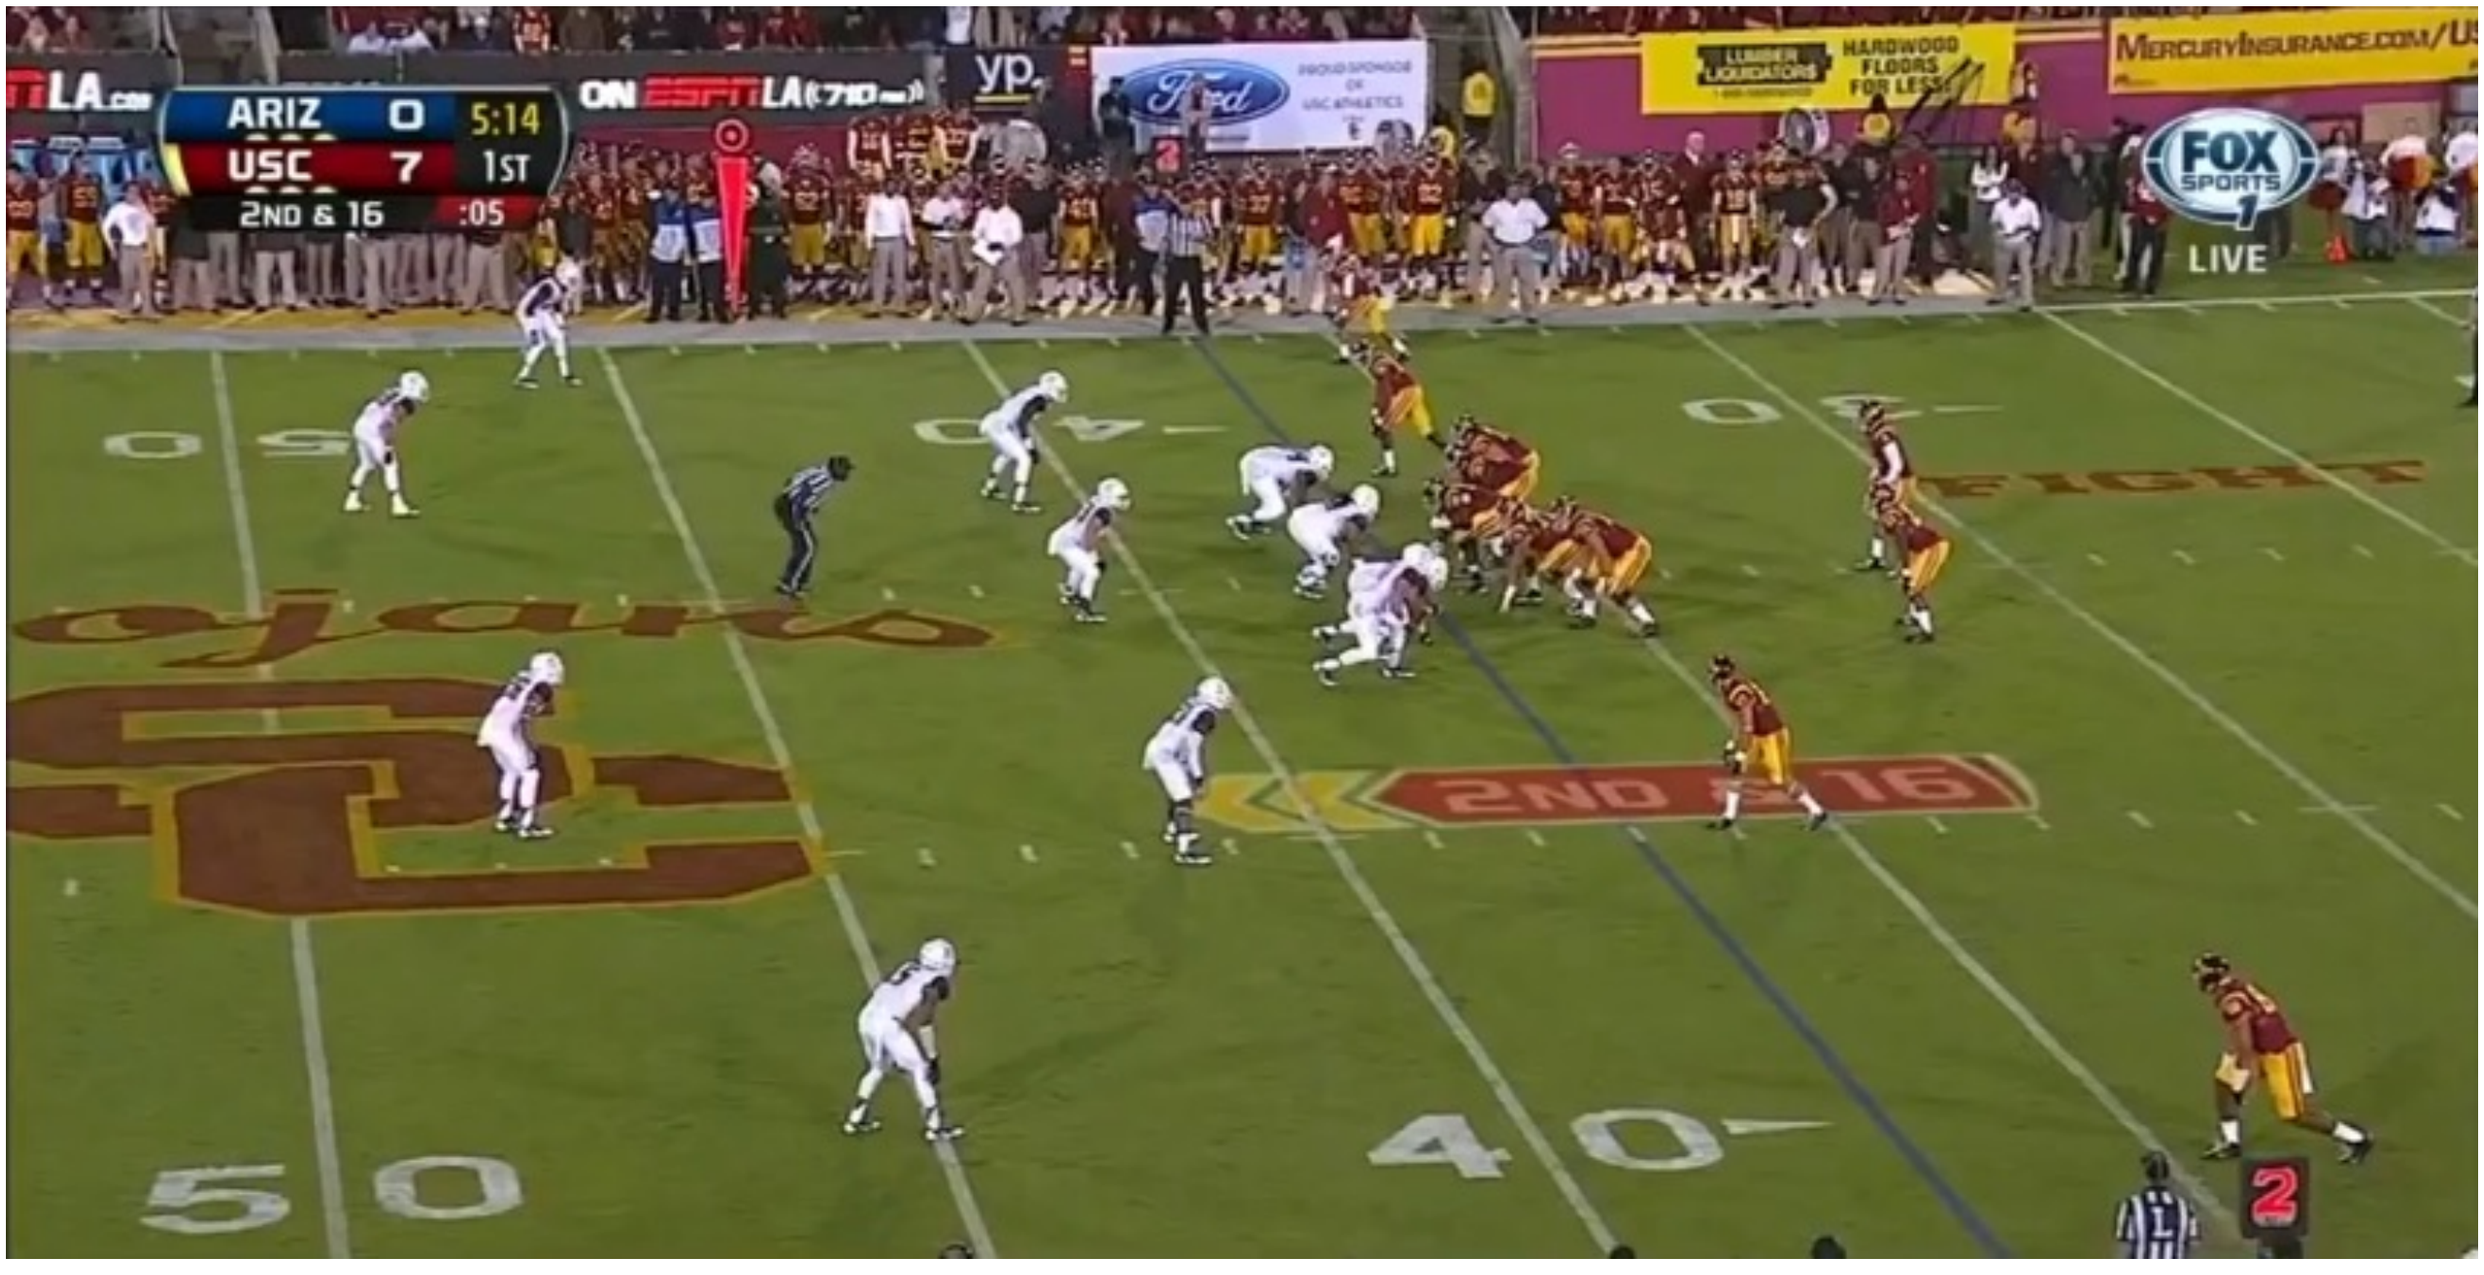

BMW 5 series

BMW night vision

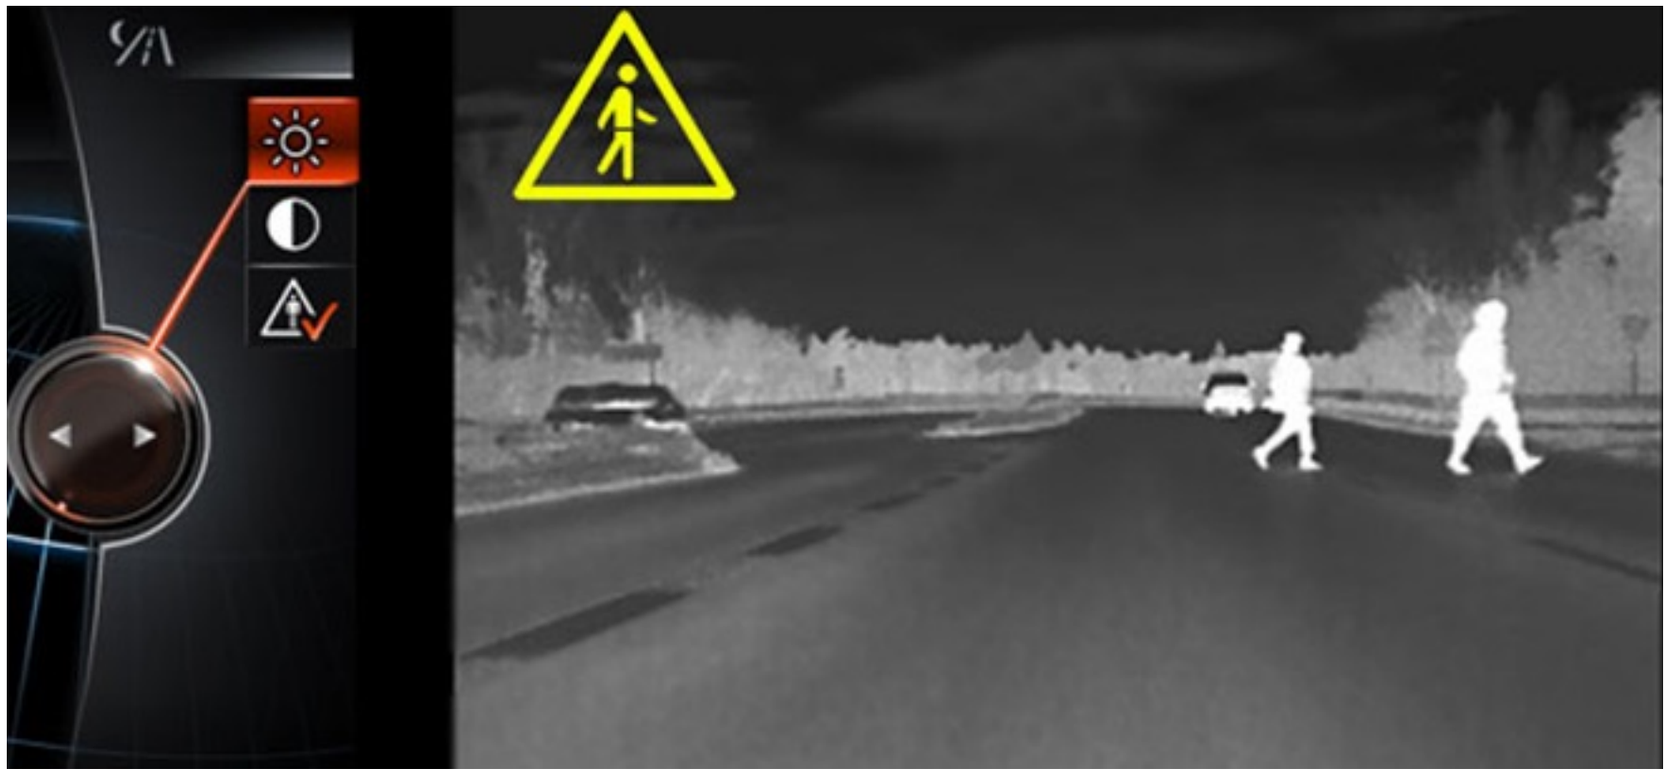

Vision in Cars

Image stitching

Photosynth

Tango

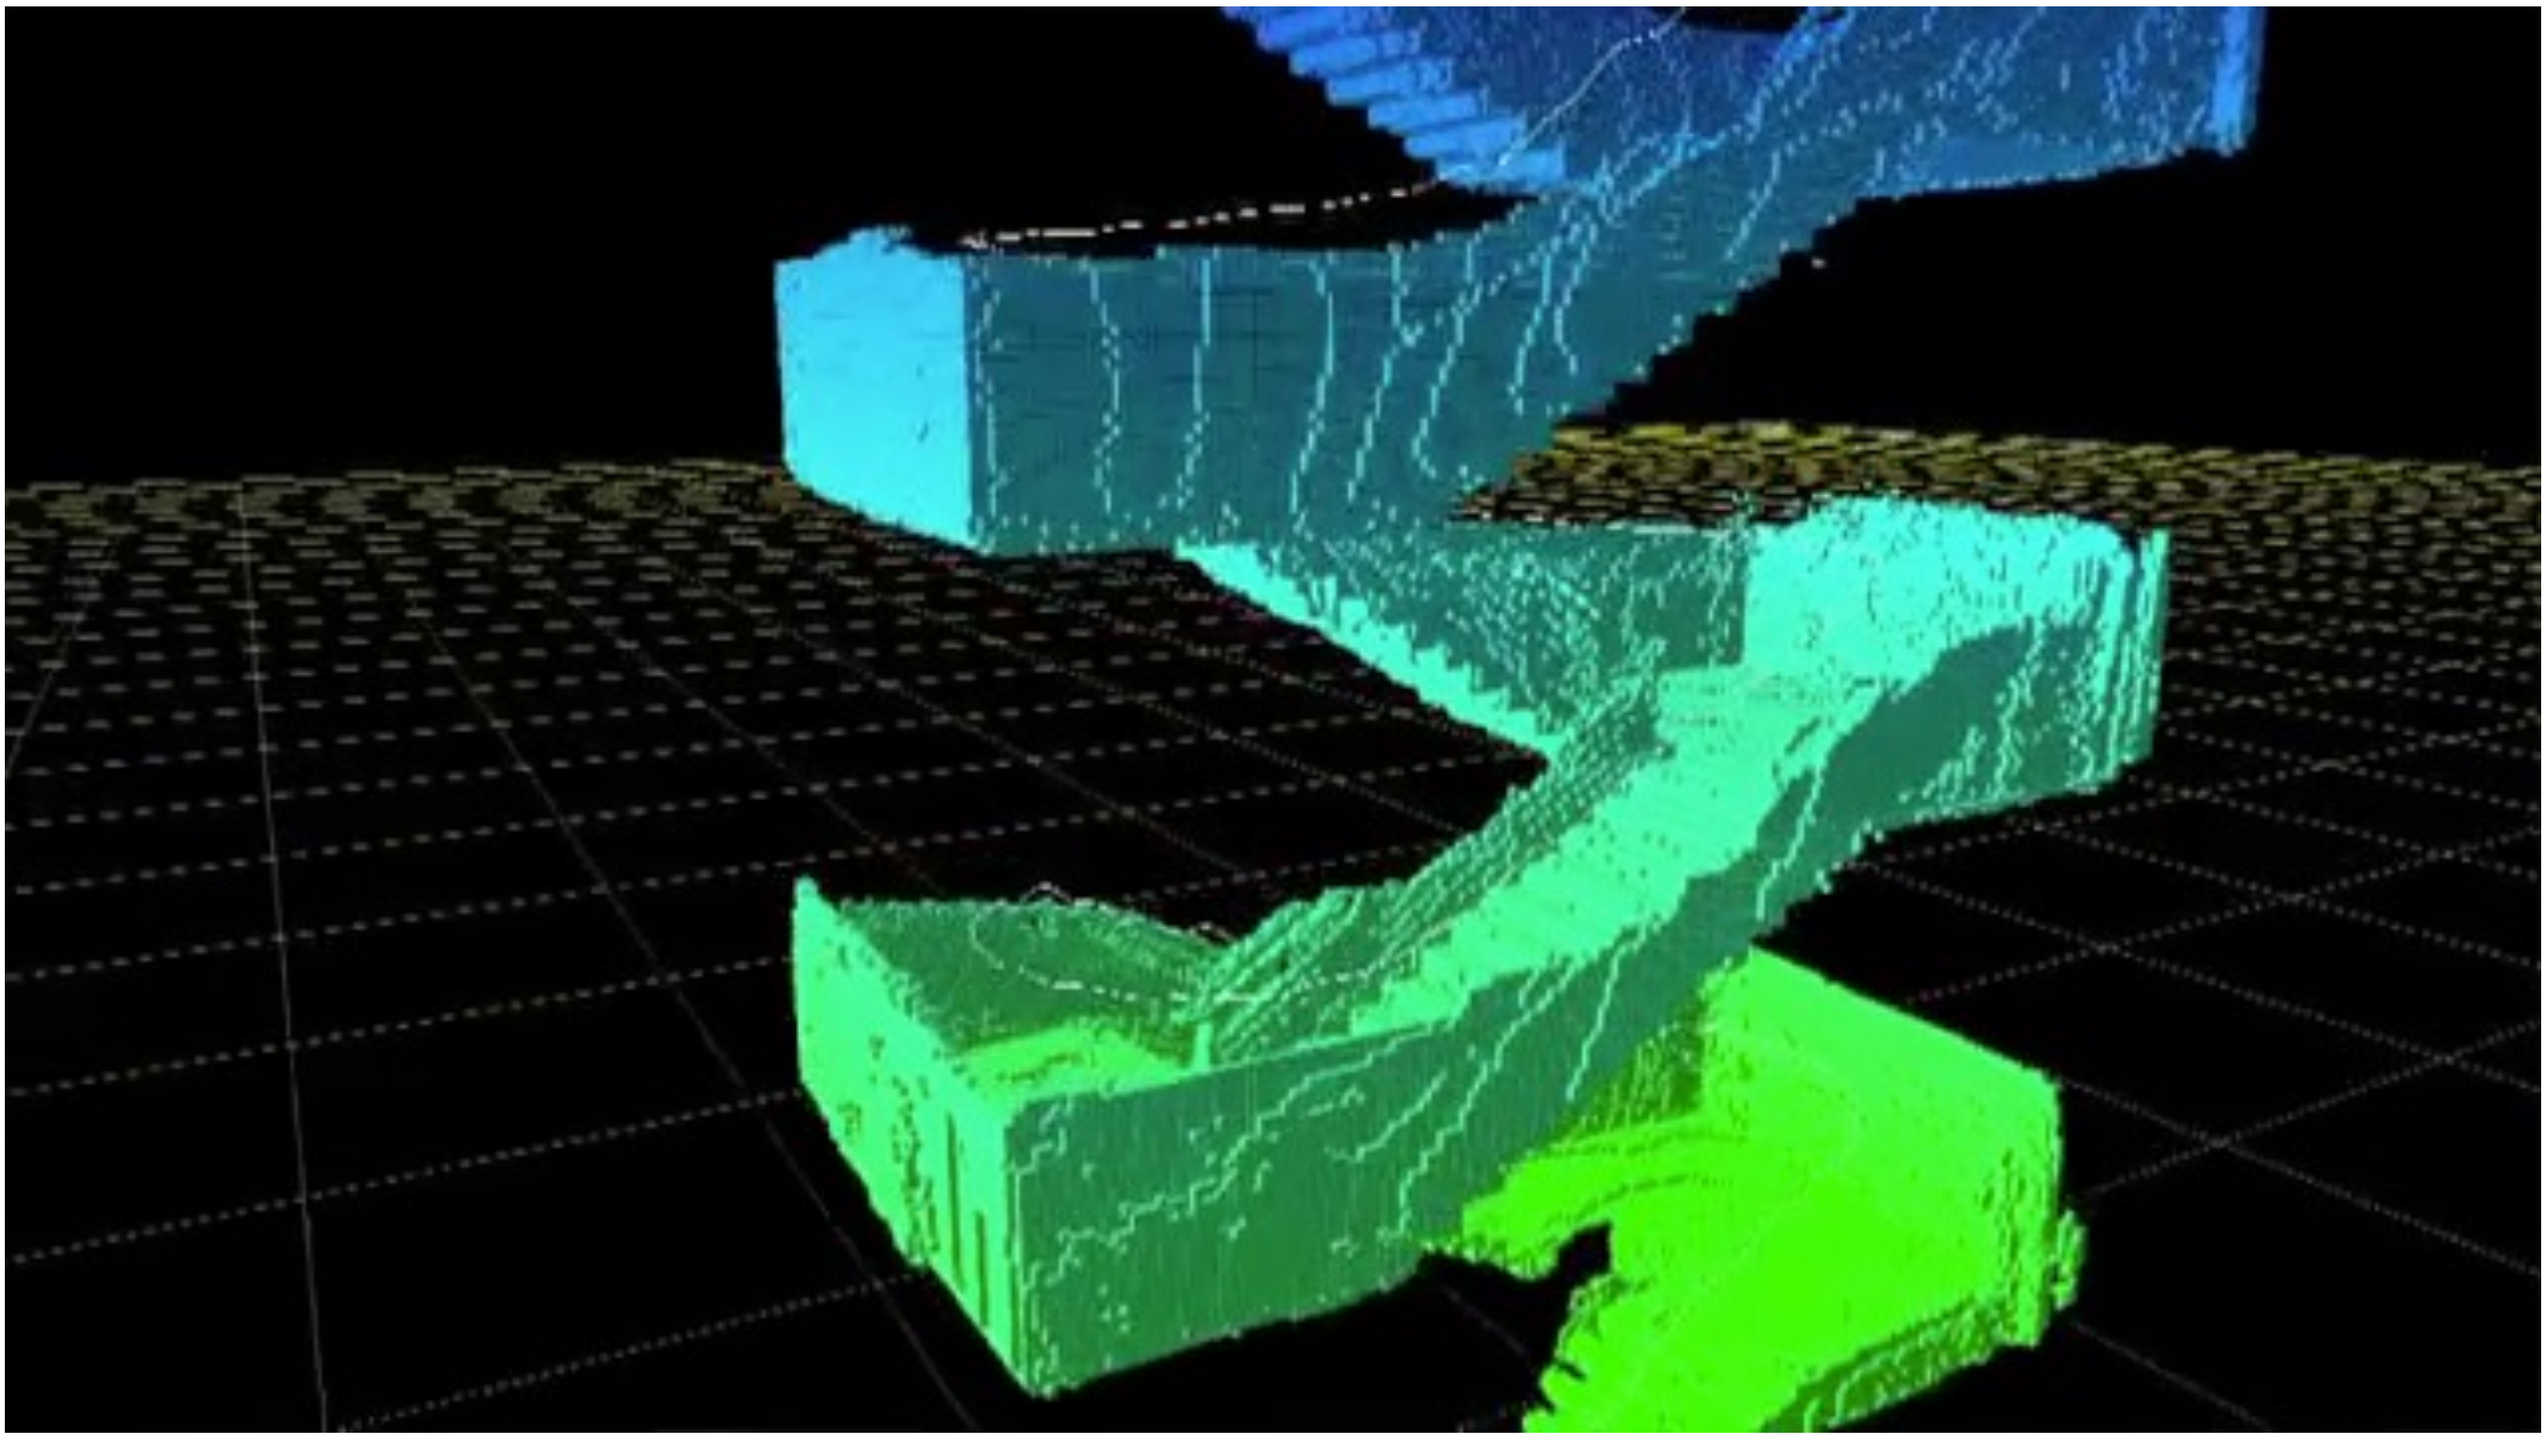

Computer Vision for VR

Virtual Fitting

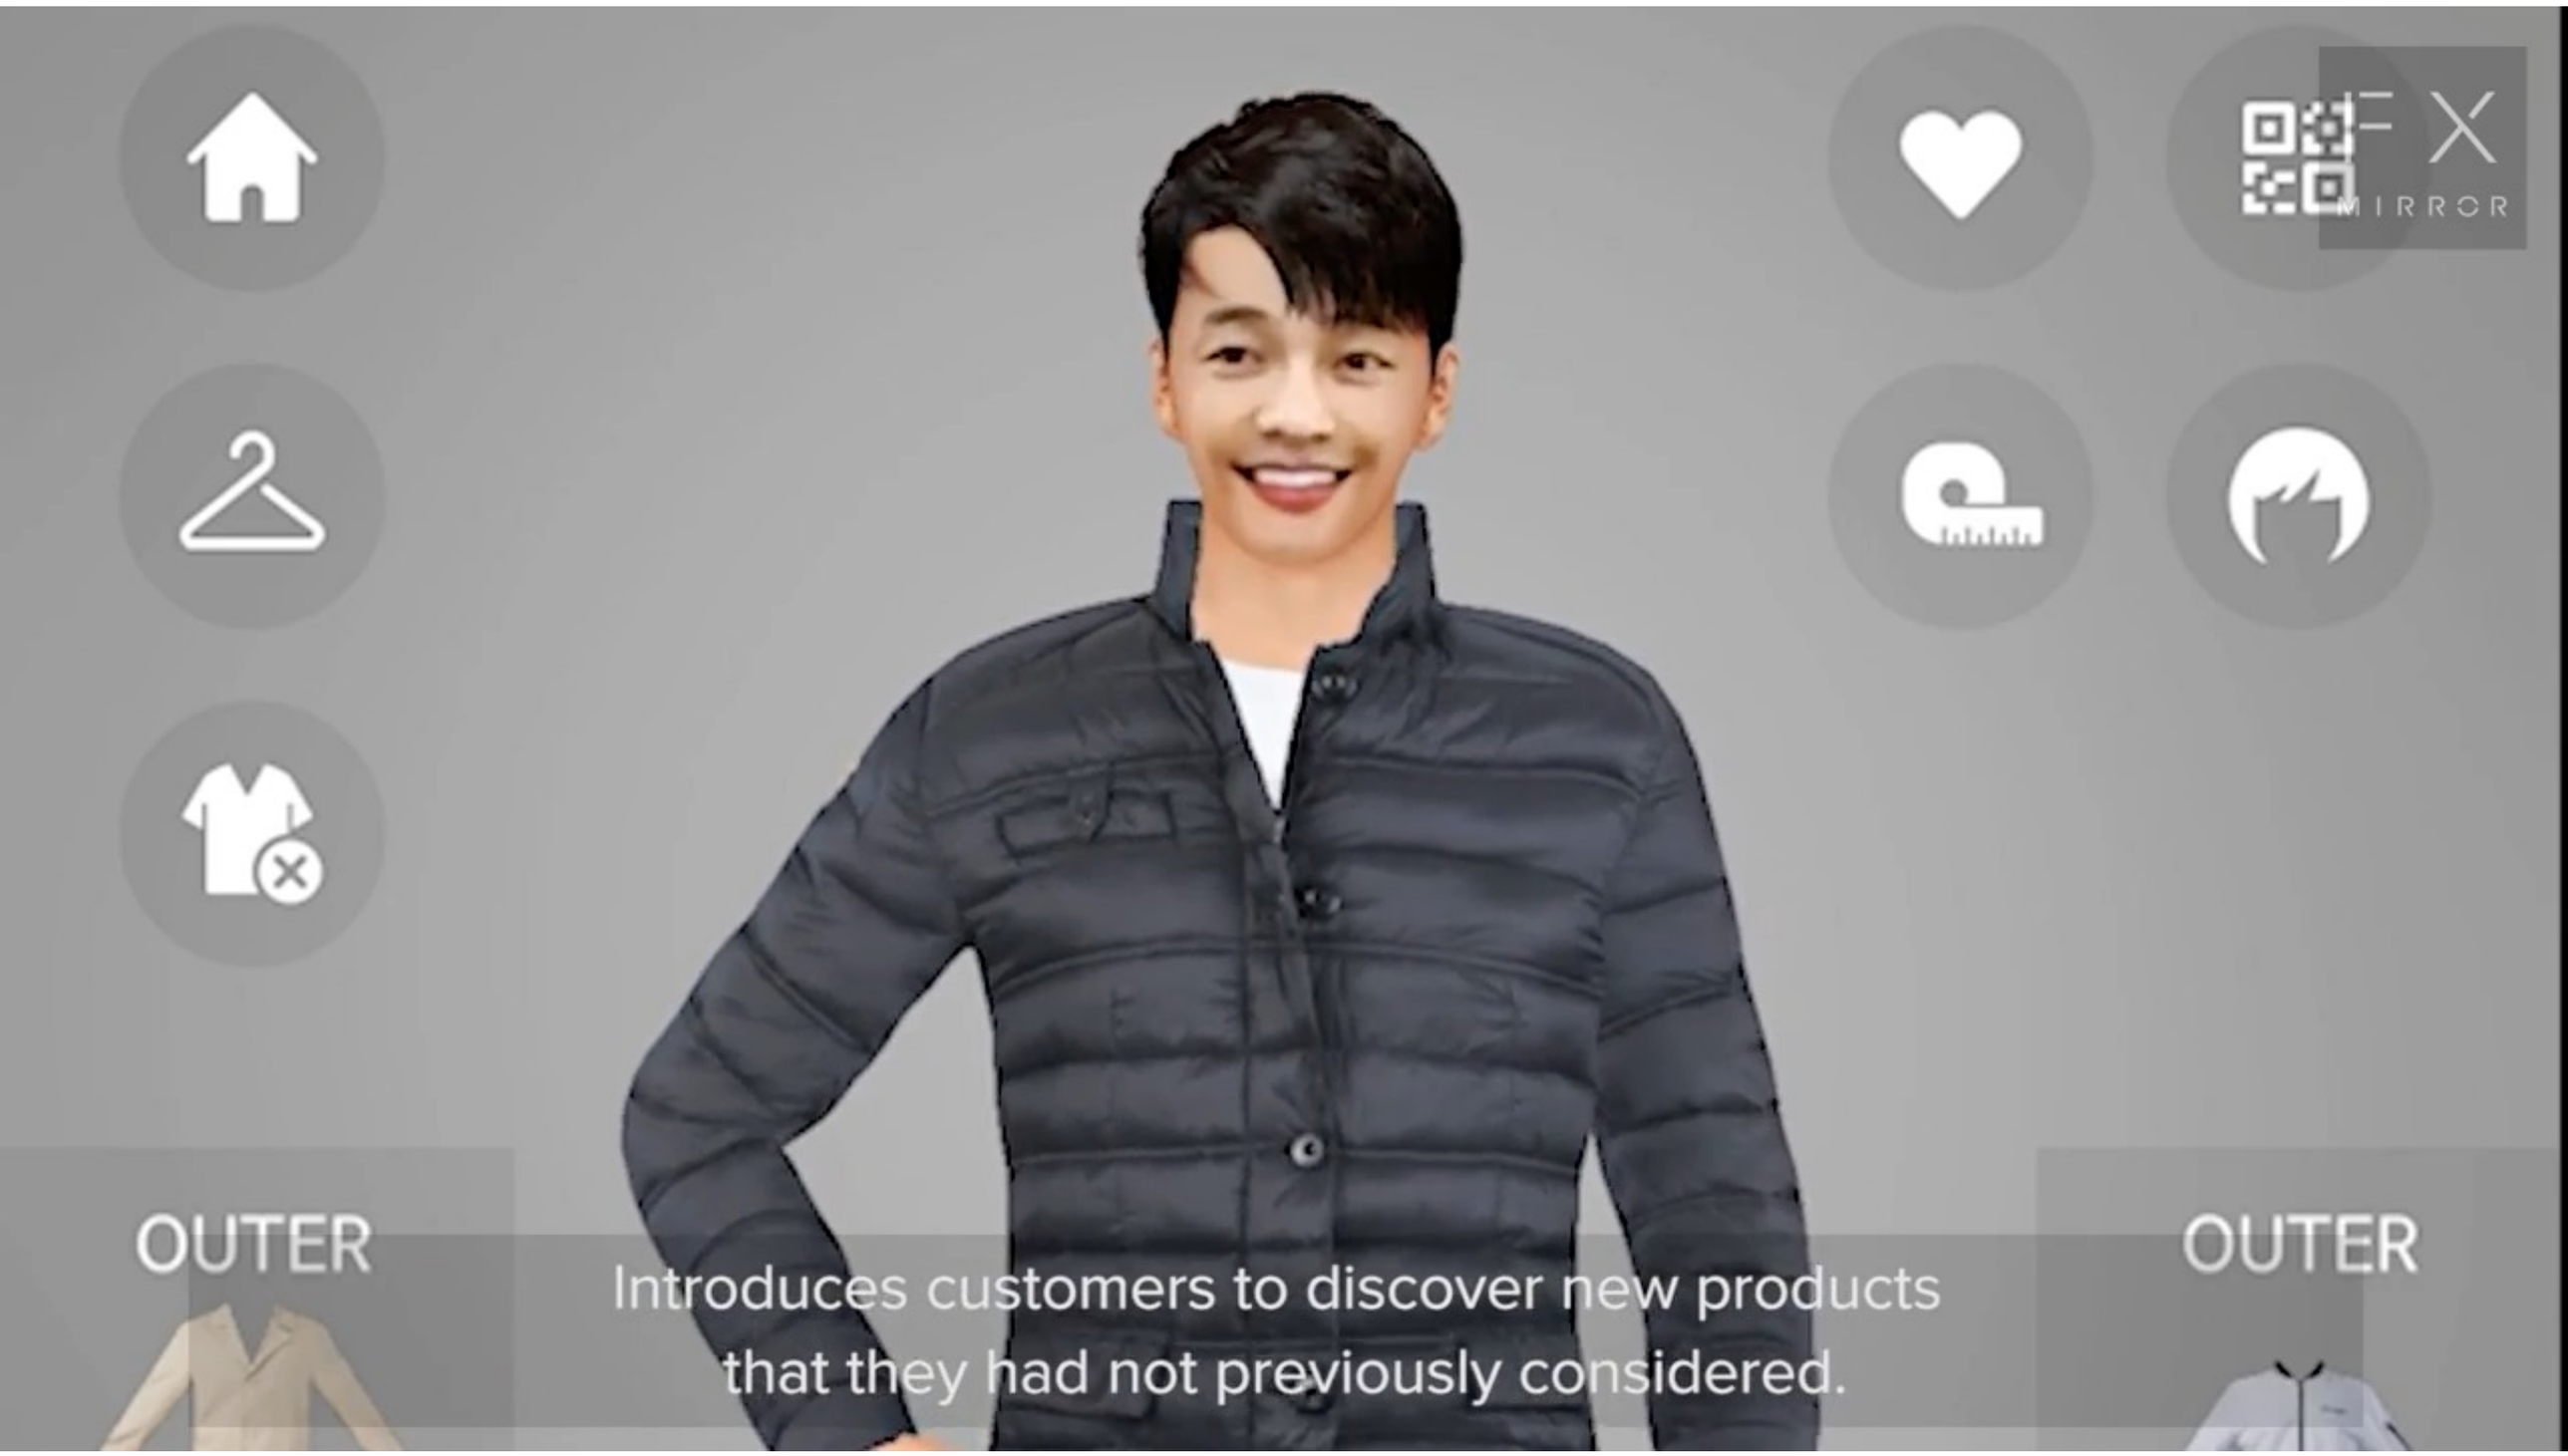

Style Transfer

Source: [Jamriška et al. 2019]

Reenactment Pipeline

Input Source

Tracking Source

Expr. Transfer

Input Target

Tracking Target

Deep Fake

Deep Fake

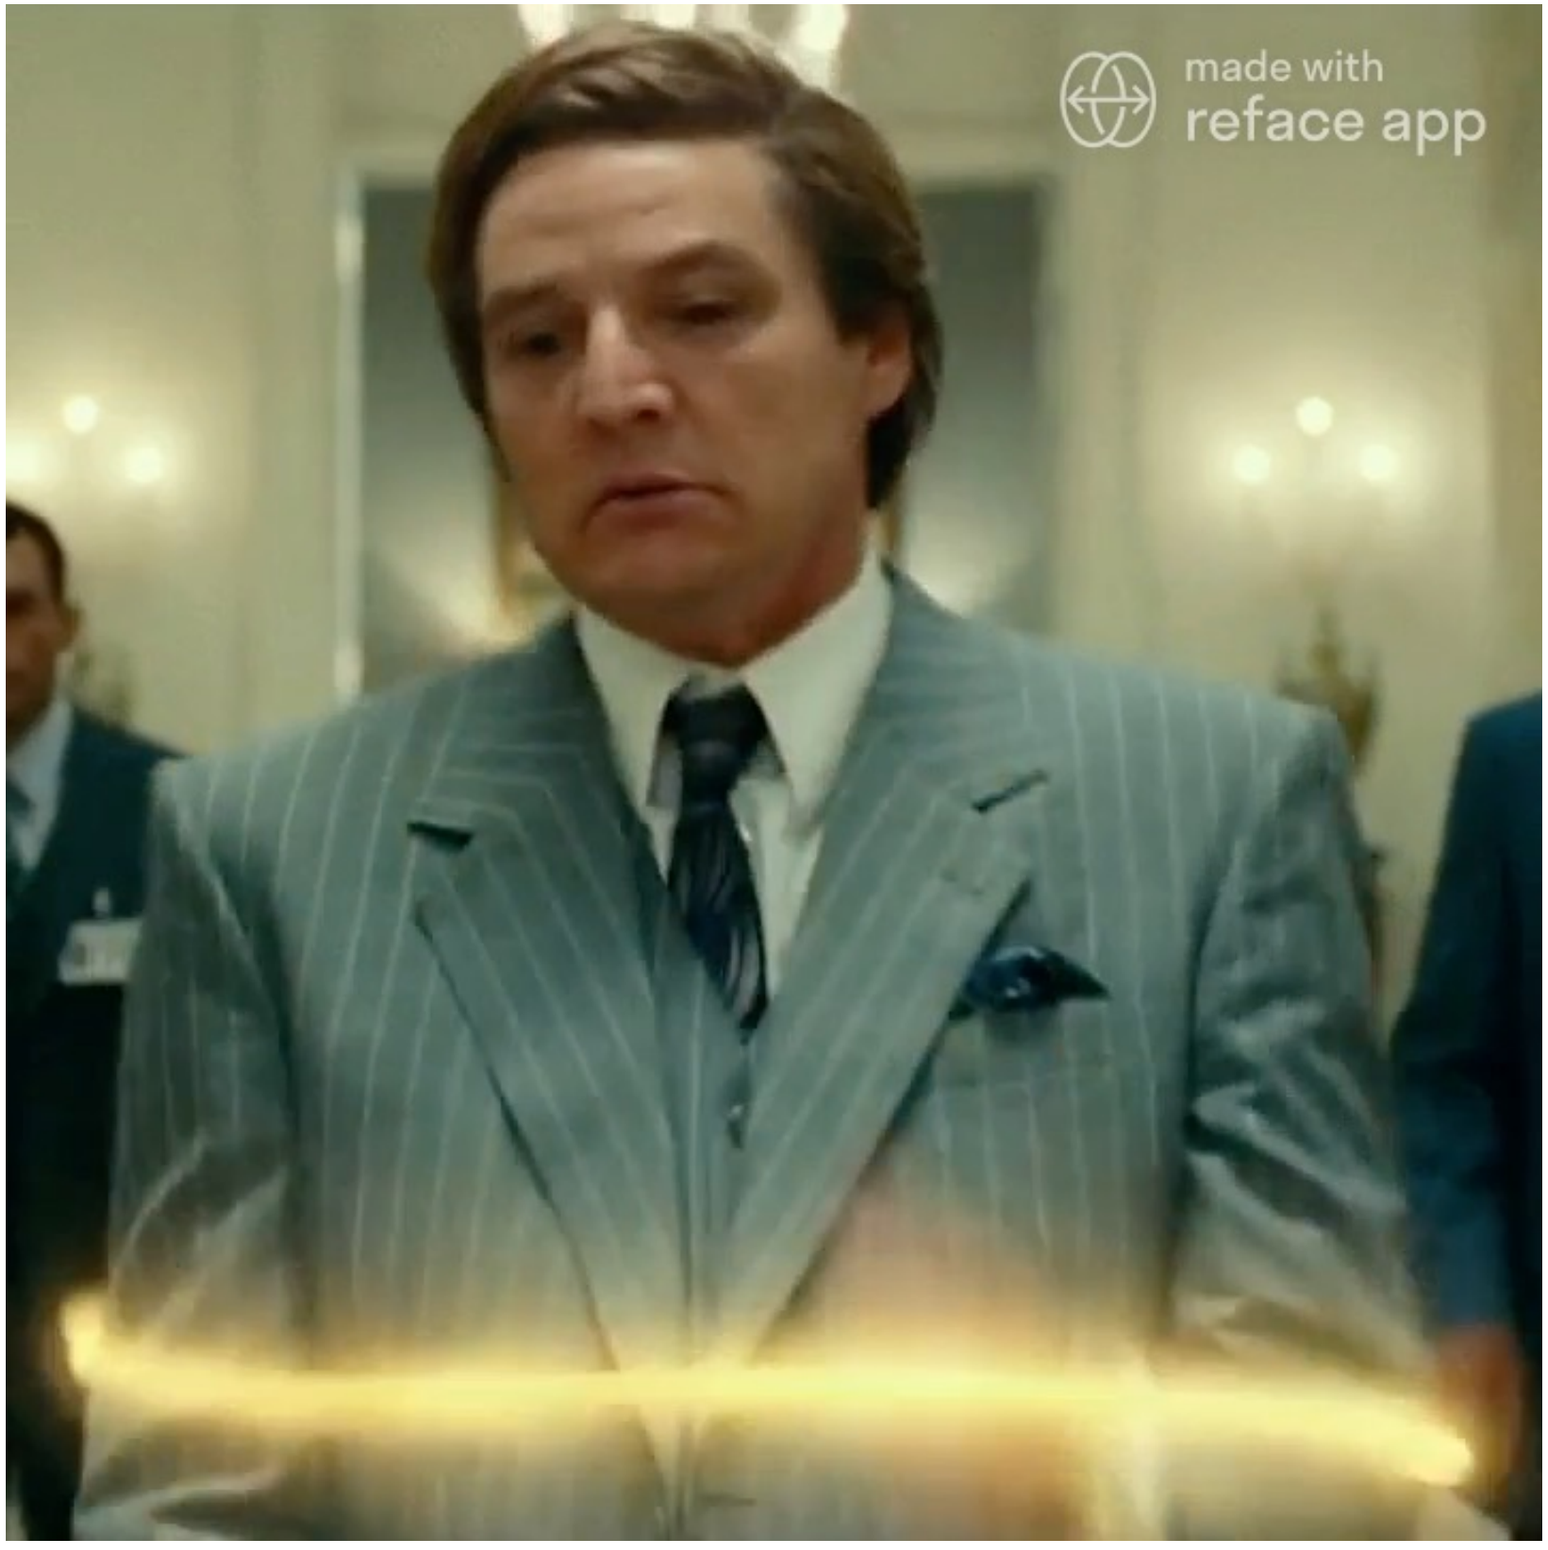

It’s a good time to do computer vision

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

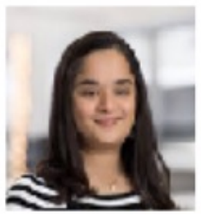

UT Austin

NYU

Google

Zillow

Oculus VR

ARGO AI

SFU

CMU

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)

Stats for CVPR (Computer Vision and Pattern Recognition)

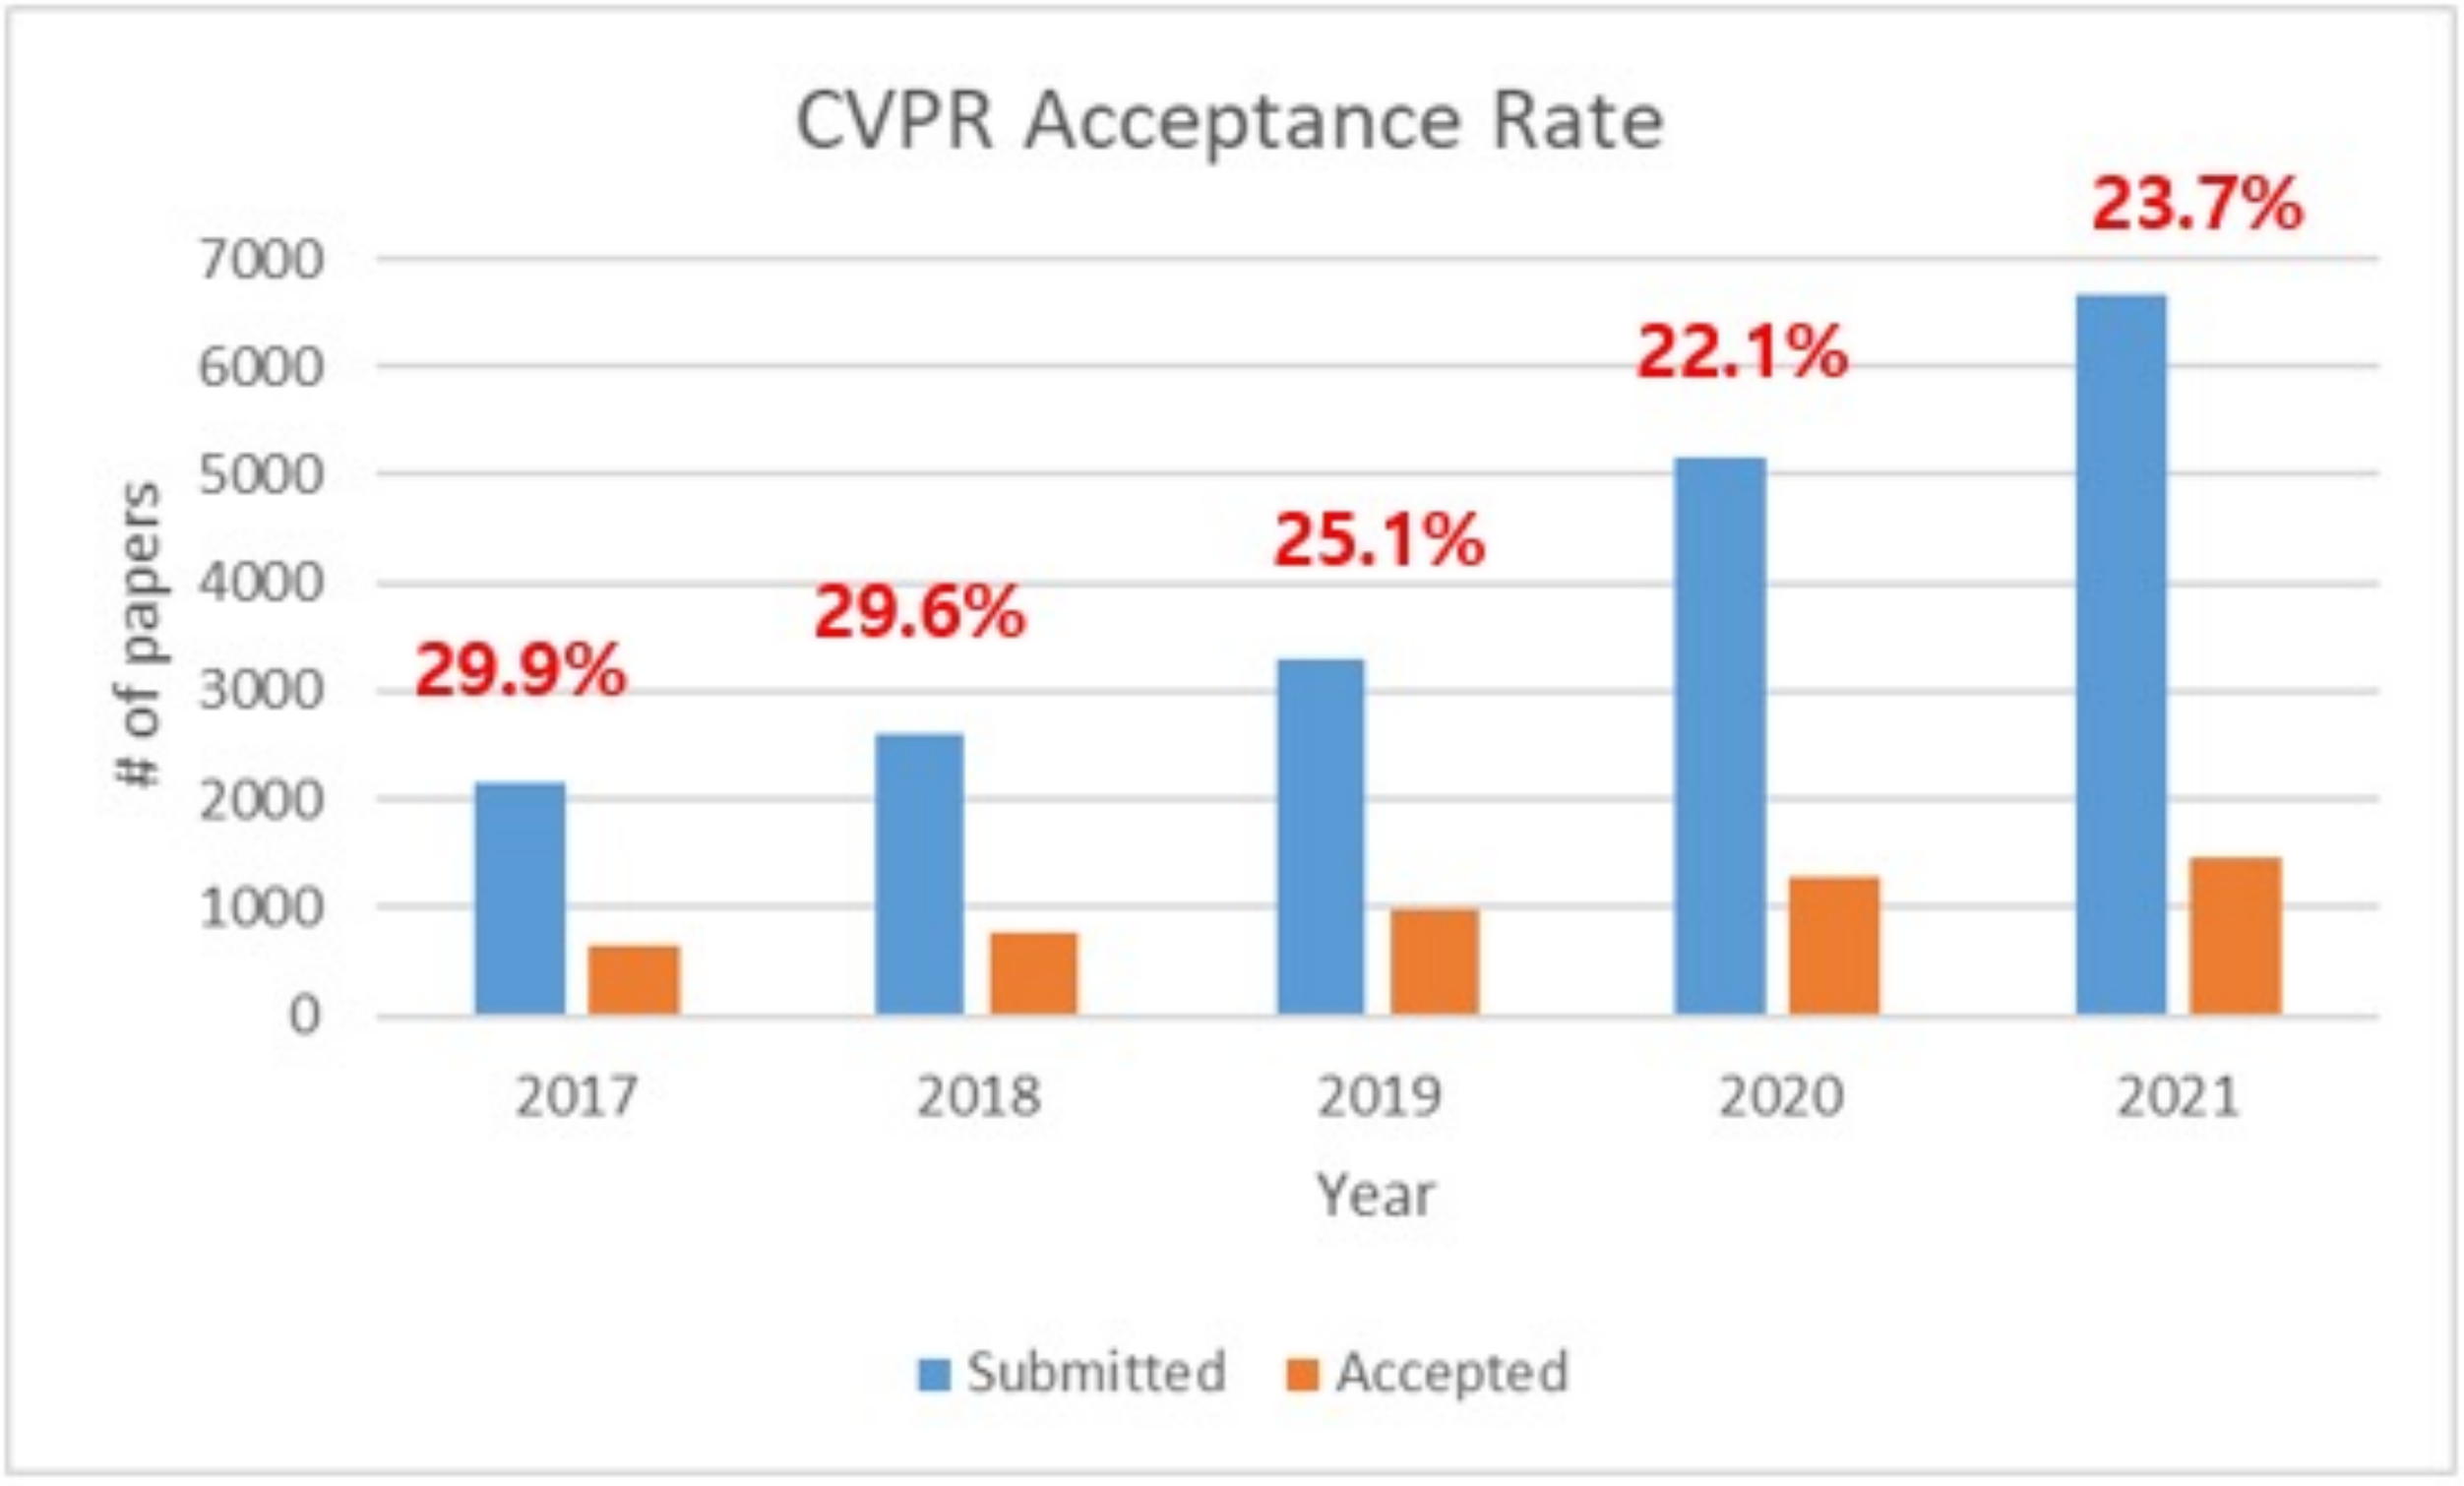

Stats for CVPR (Computer Vision and Pattern Recognition)

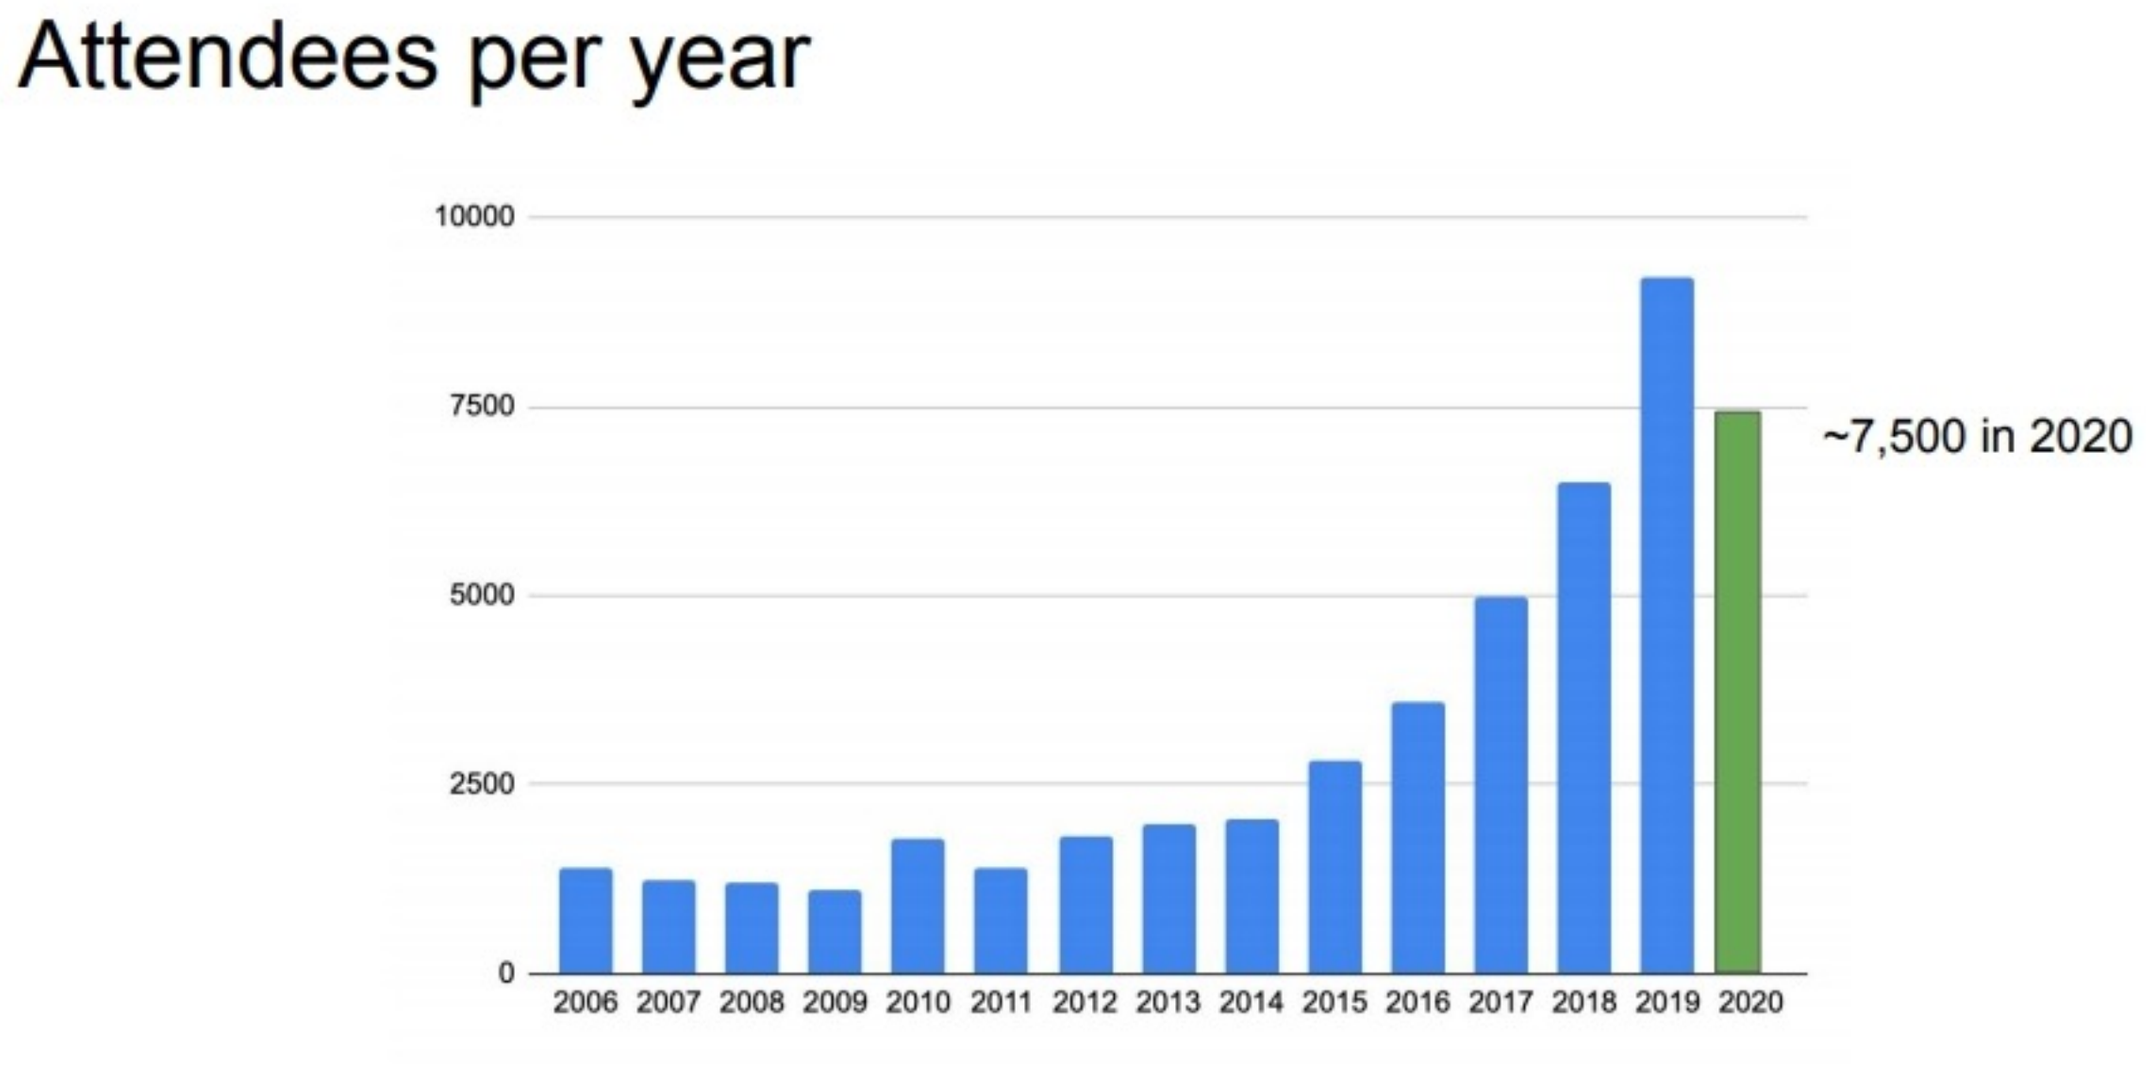

Computer vision at CMU

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.

Master in Computer Vision at CMU

Course logistics

Website

http://16385.courses.cs.cmu.edu/

(includes links to Canvas and Piazza)

Assignments Canvas

https://canvas.cmu.edu/courses/28379

Discussion & Notes piazza

https://piazza.com/cmu/spring2022/16385

Topics to be covered

Image processing:

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

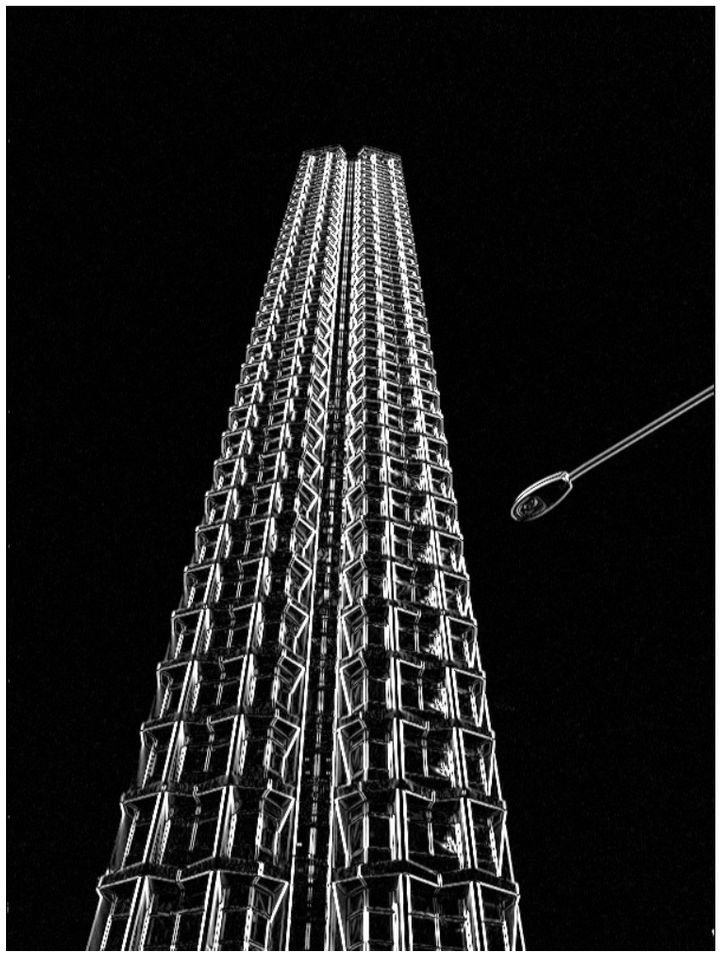

Topics to be covered

Feature detection and correspondences:

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

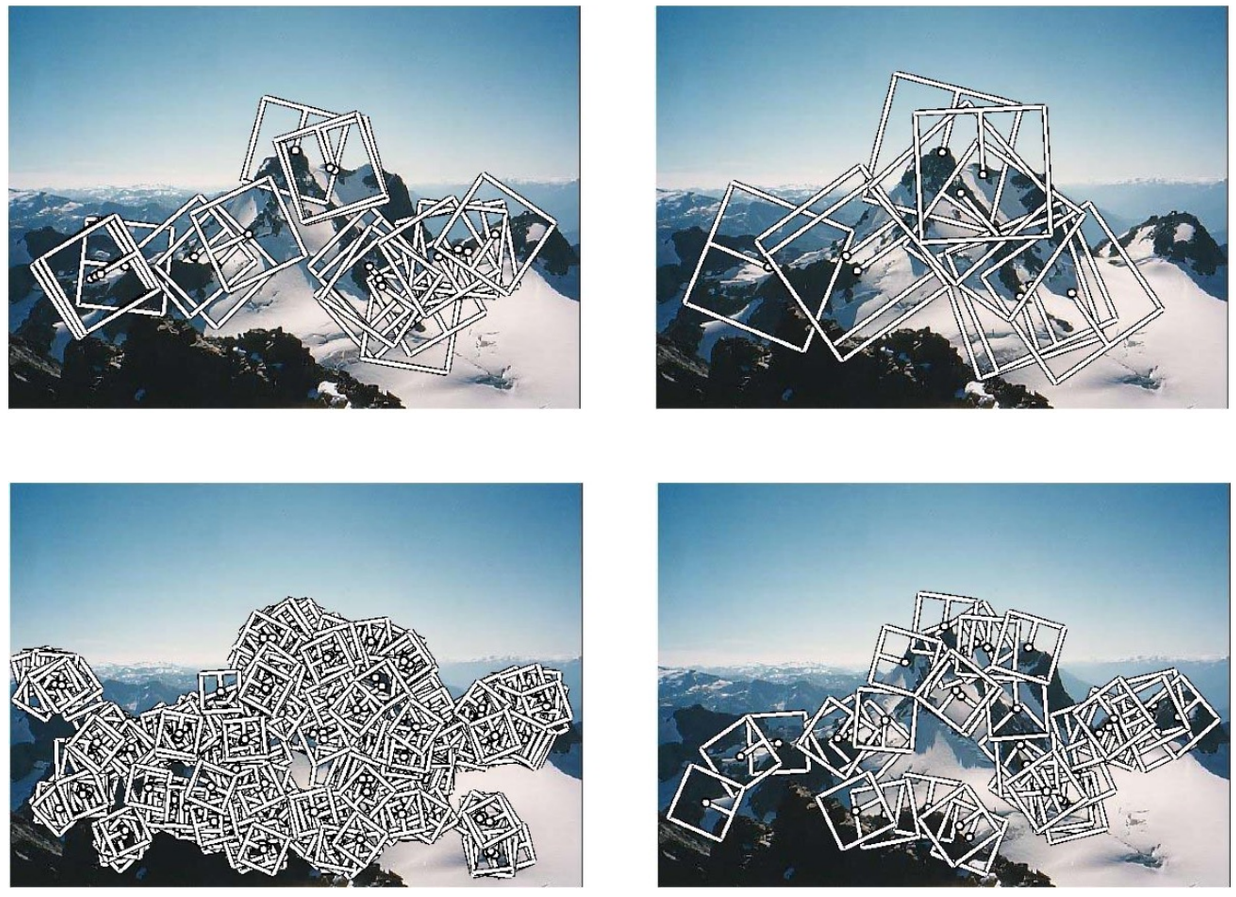

Topics to be covered

Transformations and geometry:

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

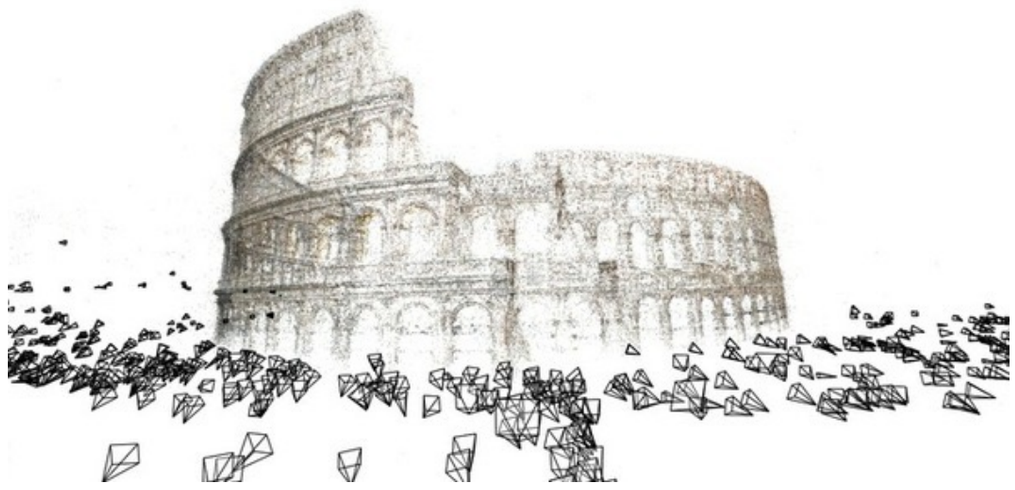

Topics to be covered

Physics-based vision:

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

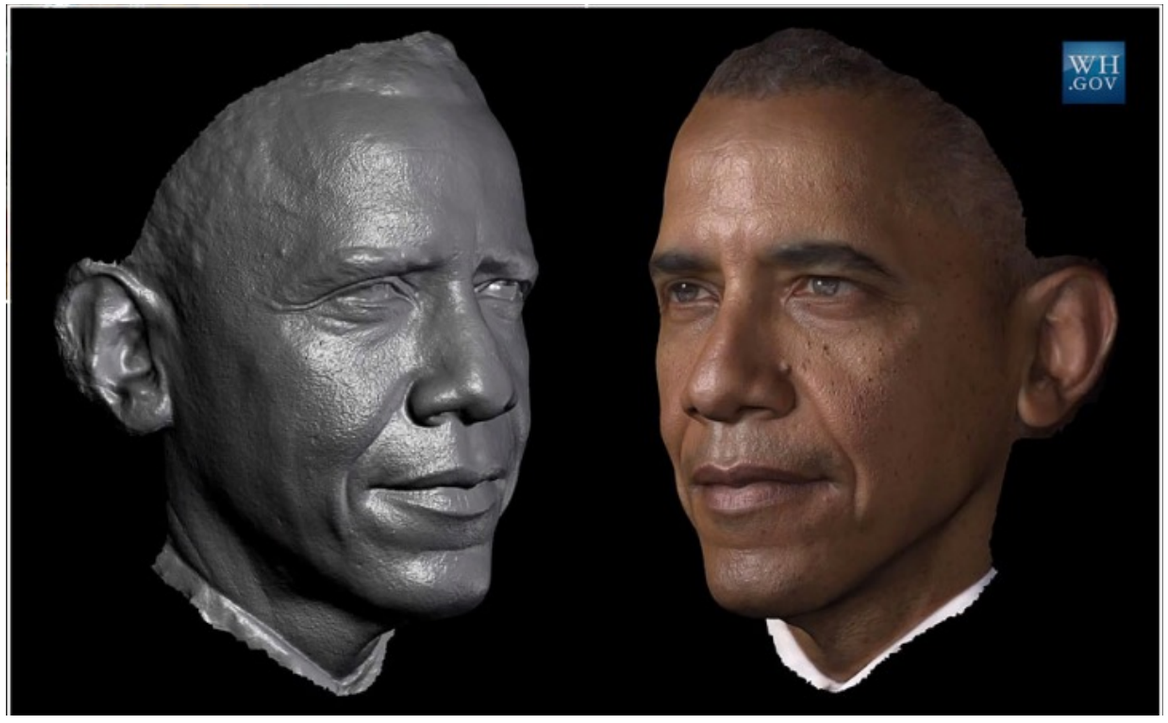

Topics to be covered

Objects, faces, and learning:

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

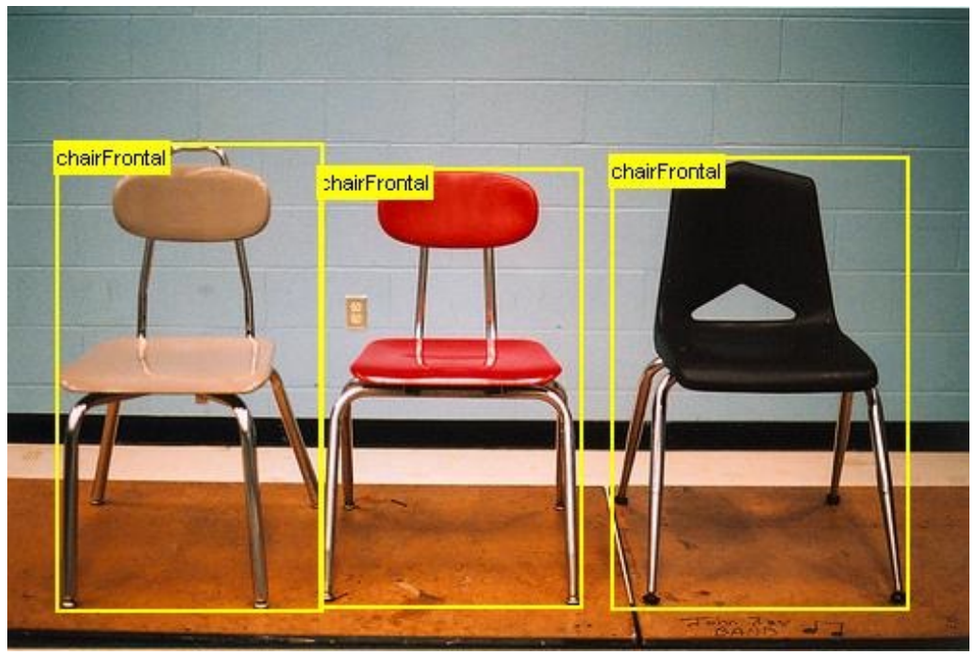

Topics to be covered

Dealing with motion:

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

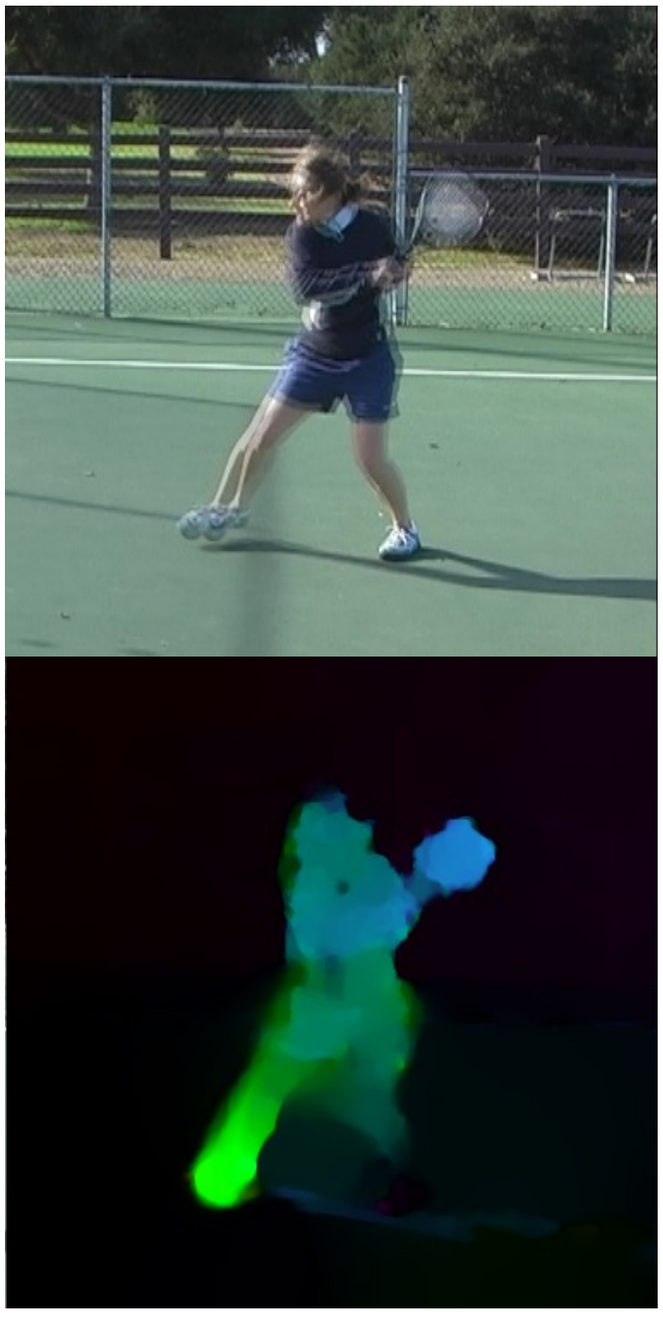

Grading

- Six two-week programming assignments: 70%
- Eleven weekly take-home quizzes: 27%
- Class, Website, and Piazza participation: 3%

Participation:

- Be active! Ask questions.
- Post on Piazza and course website.

Programming Assignments

- a lot of programming in Python
- hours and hours of programming
- days and days of debugging
- generous grading policy
- take advantage of extra credit

Assignment 1 Hough Transform
Assignment 2 Homography
Assignment 3 Stereo
Assignment 4 Bag of Words
Assignment 5 Convolutional Neural Nets
Assignment 6 Image Alignment

Programming Assignments

- a lot of programming in Python
- hours and hours of programming
- days and days of debugging
- generous grading policy
- take advantage of extra credit

Assignment 1 Hough Transform
Assignment 2 Homography
Assignment 3 Stereo
Assignment 4 Bag of Words
Assignment 5 Convolutional Neural Nets
Assignment 6 Image Alignment

Seriously.. a lot of programming, so start early!

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 8 out of 11 quizzes
- late quizzes will not be graded

Book

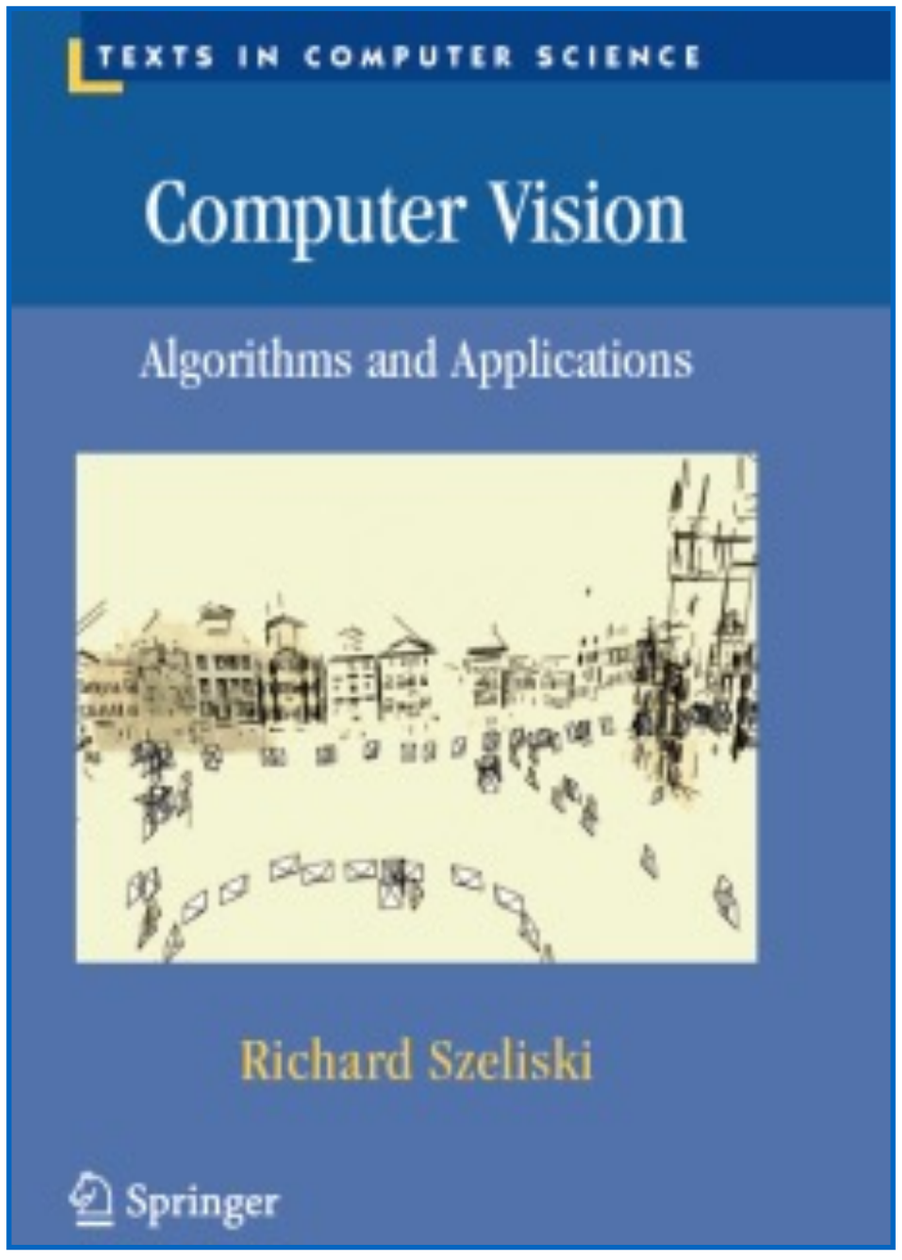

PDF online

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

- Feel free to email us about administrative questions.
    - please use [16385] in email title!

- Lecture questions should be asked on course website (or in lecture), and assignment/quiz/logistic 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 me about additional office hours.

<u>I will announce office hours at the end of this week.</u>