Computer Vision

CSE P576

Dr. Matthew Brown

The Course

- People
  - Matthew Brown + Guest(s) TBA
  - TAs: Nishat Khan, Dianqi Li
- Time and location
  - Lectures: Tuesdays 6:30-9:20pm
  - Office hours: Thursday 5:30-6:30pm (or by appointment)
- Evaluation
  - 4 projects, equally weighted
- Resources
  - https://courses.cs.washington.edu/courses/csep576/20sp
  - Piazza=Discussion board, Canvas=Assignments
  - Book 1: "Computer Vision", Szeliski,
  - Book 2: "Deep Learning", Goodfellow et al.
  - Stanford CS231N (CNNs for Vision)

Face Detection

[Motorola]

Camera Tracking

[Boujou -- Vicon/OMG]

3D Reconstruction

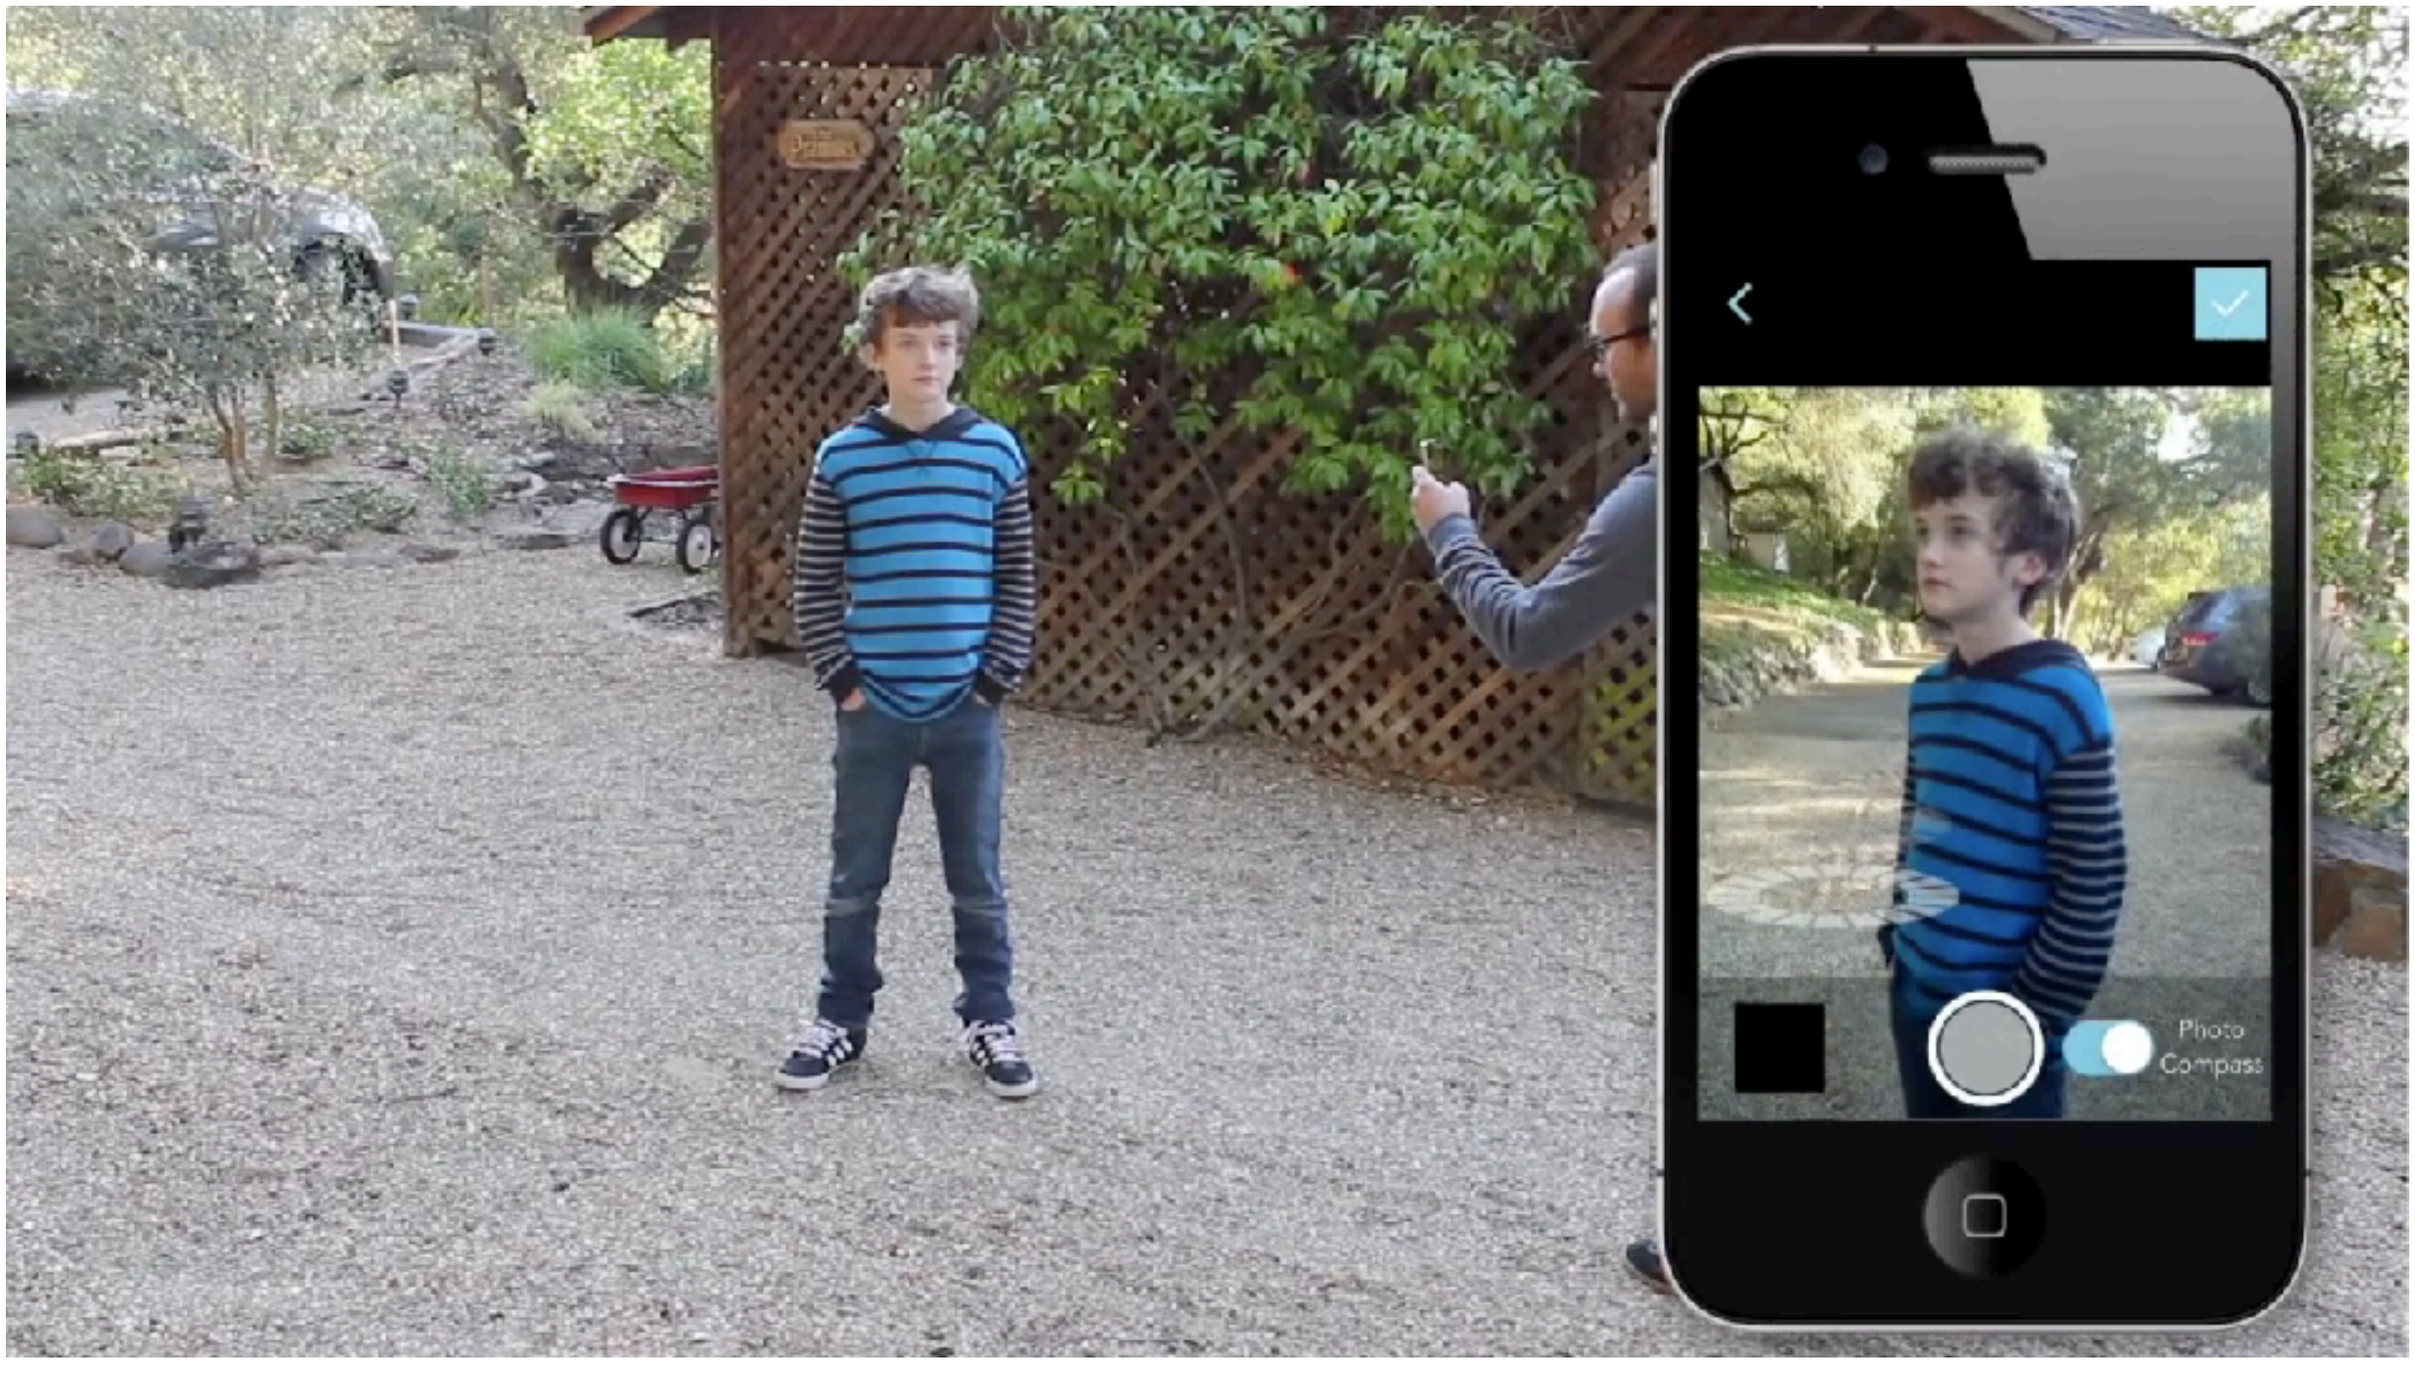

[Autodesk 123D Catch]

Body Pose Tracking

[Microsoft Xbox Kinect]

Body Pose Tracking

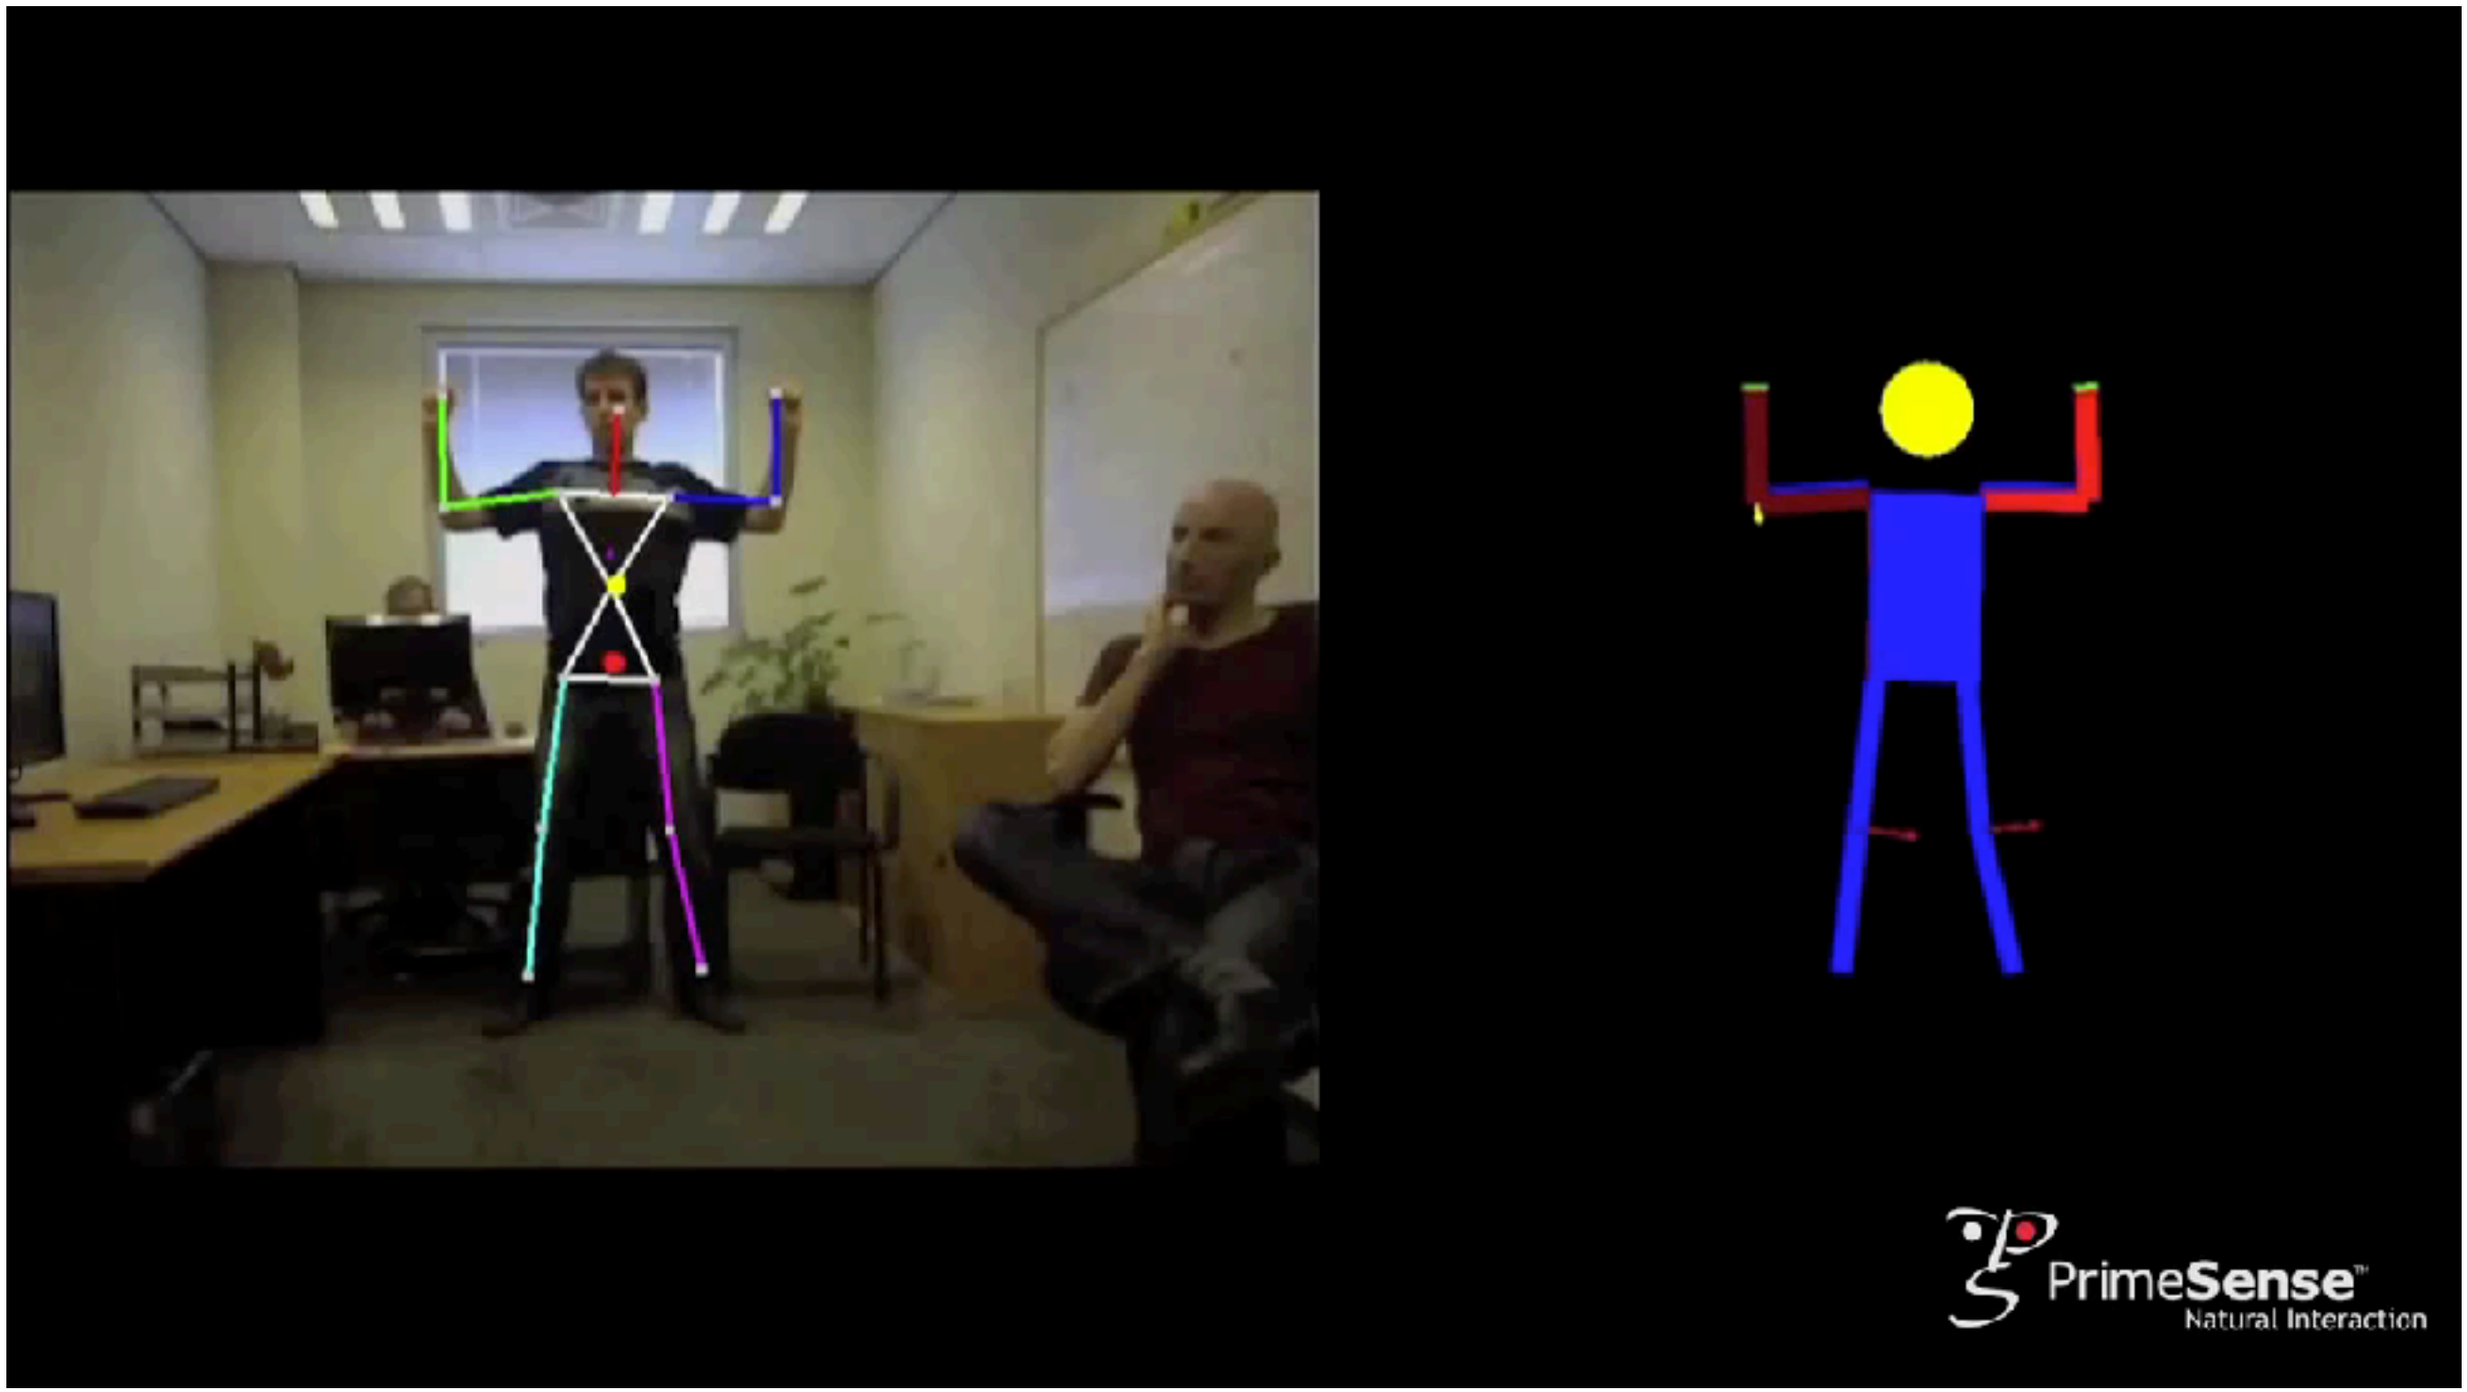

[PrimeSense]

Image Recognition and Search

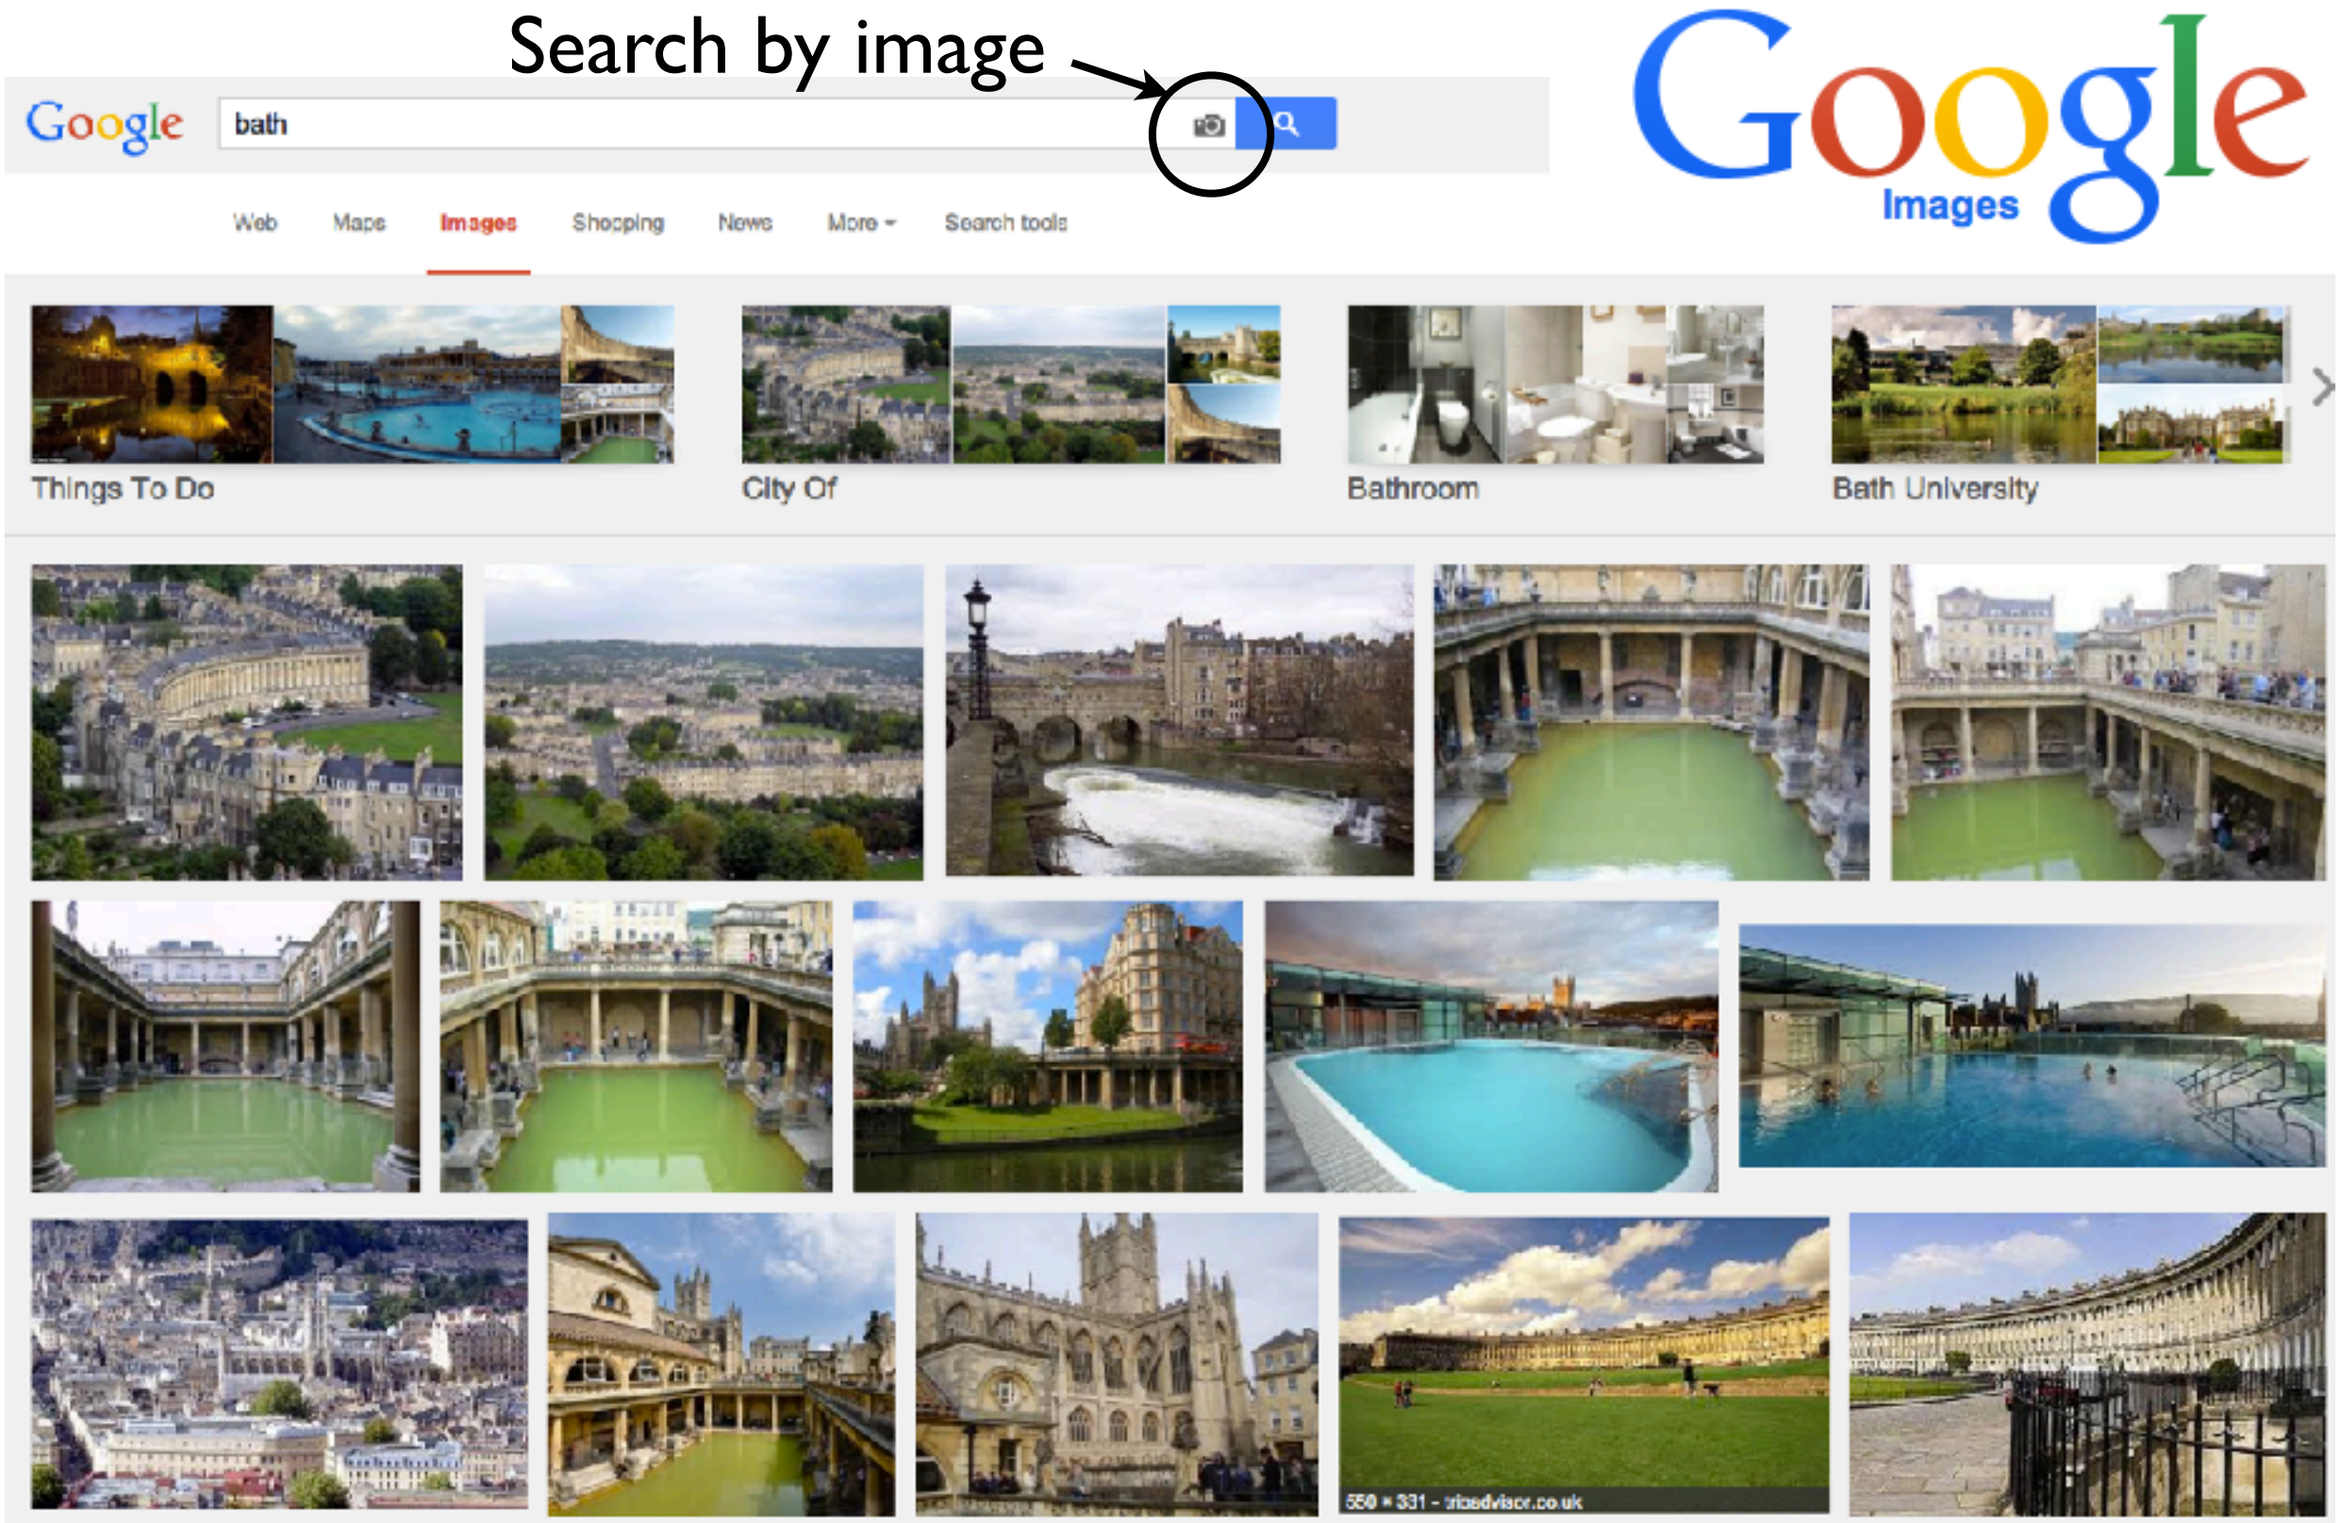

Self Driving Cars

[Google]

Flying Vehicles

www.skydio.com

AR/VR

[Microsoft Hololens]

Mobile Apps

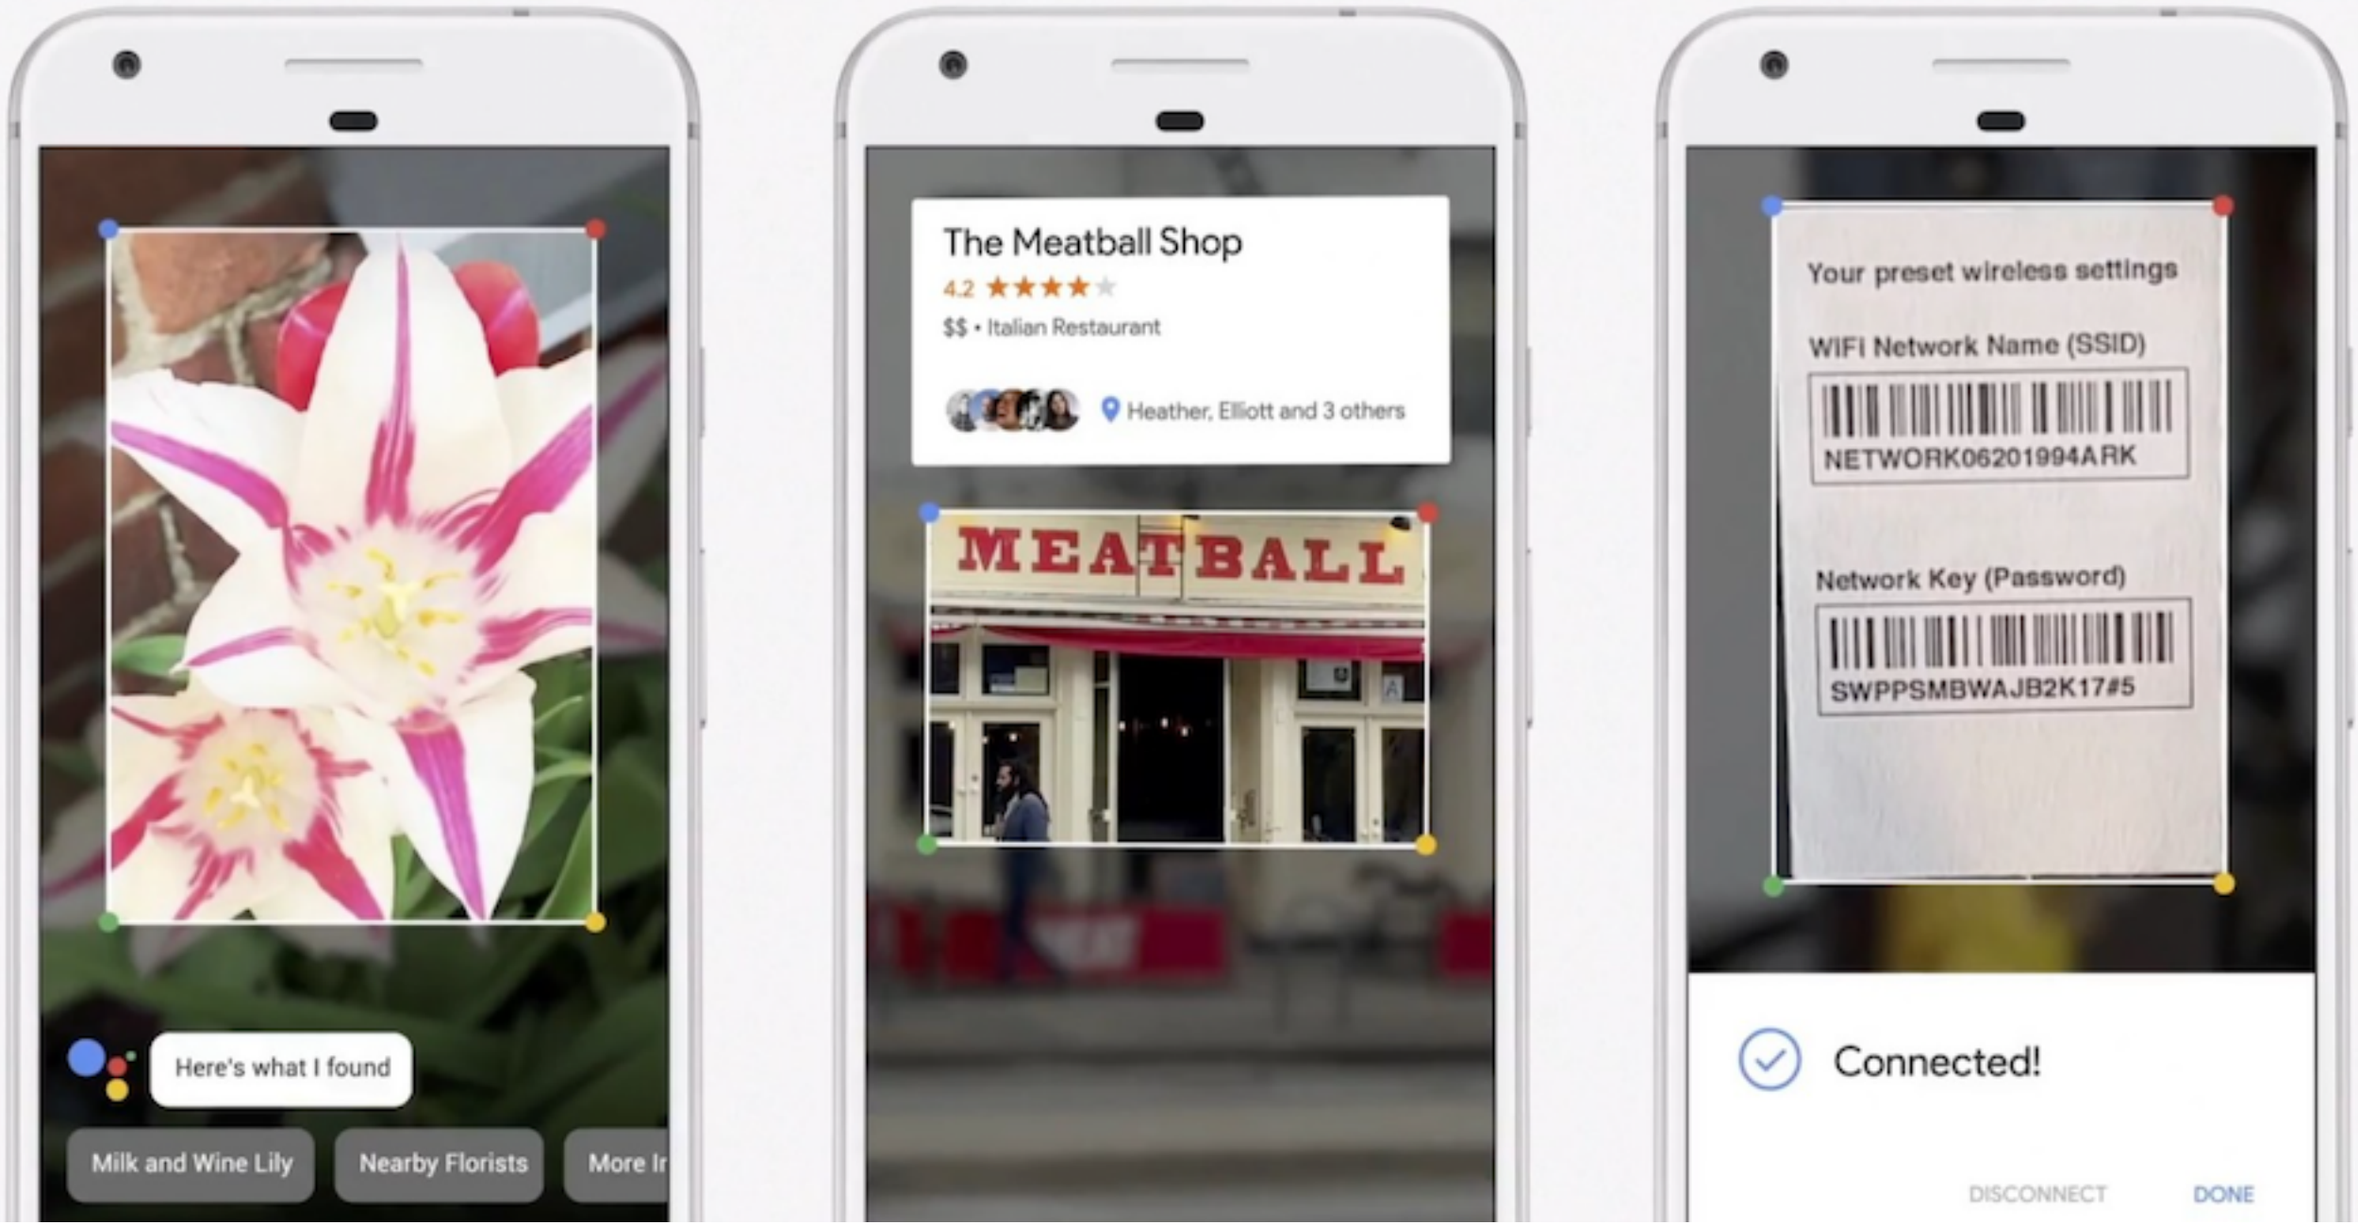

[Google Lens]

Art

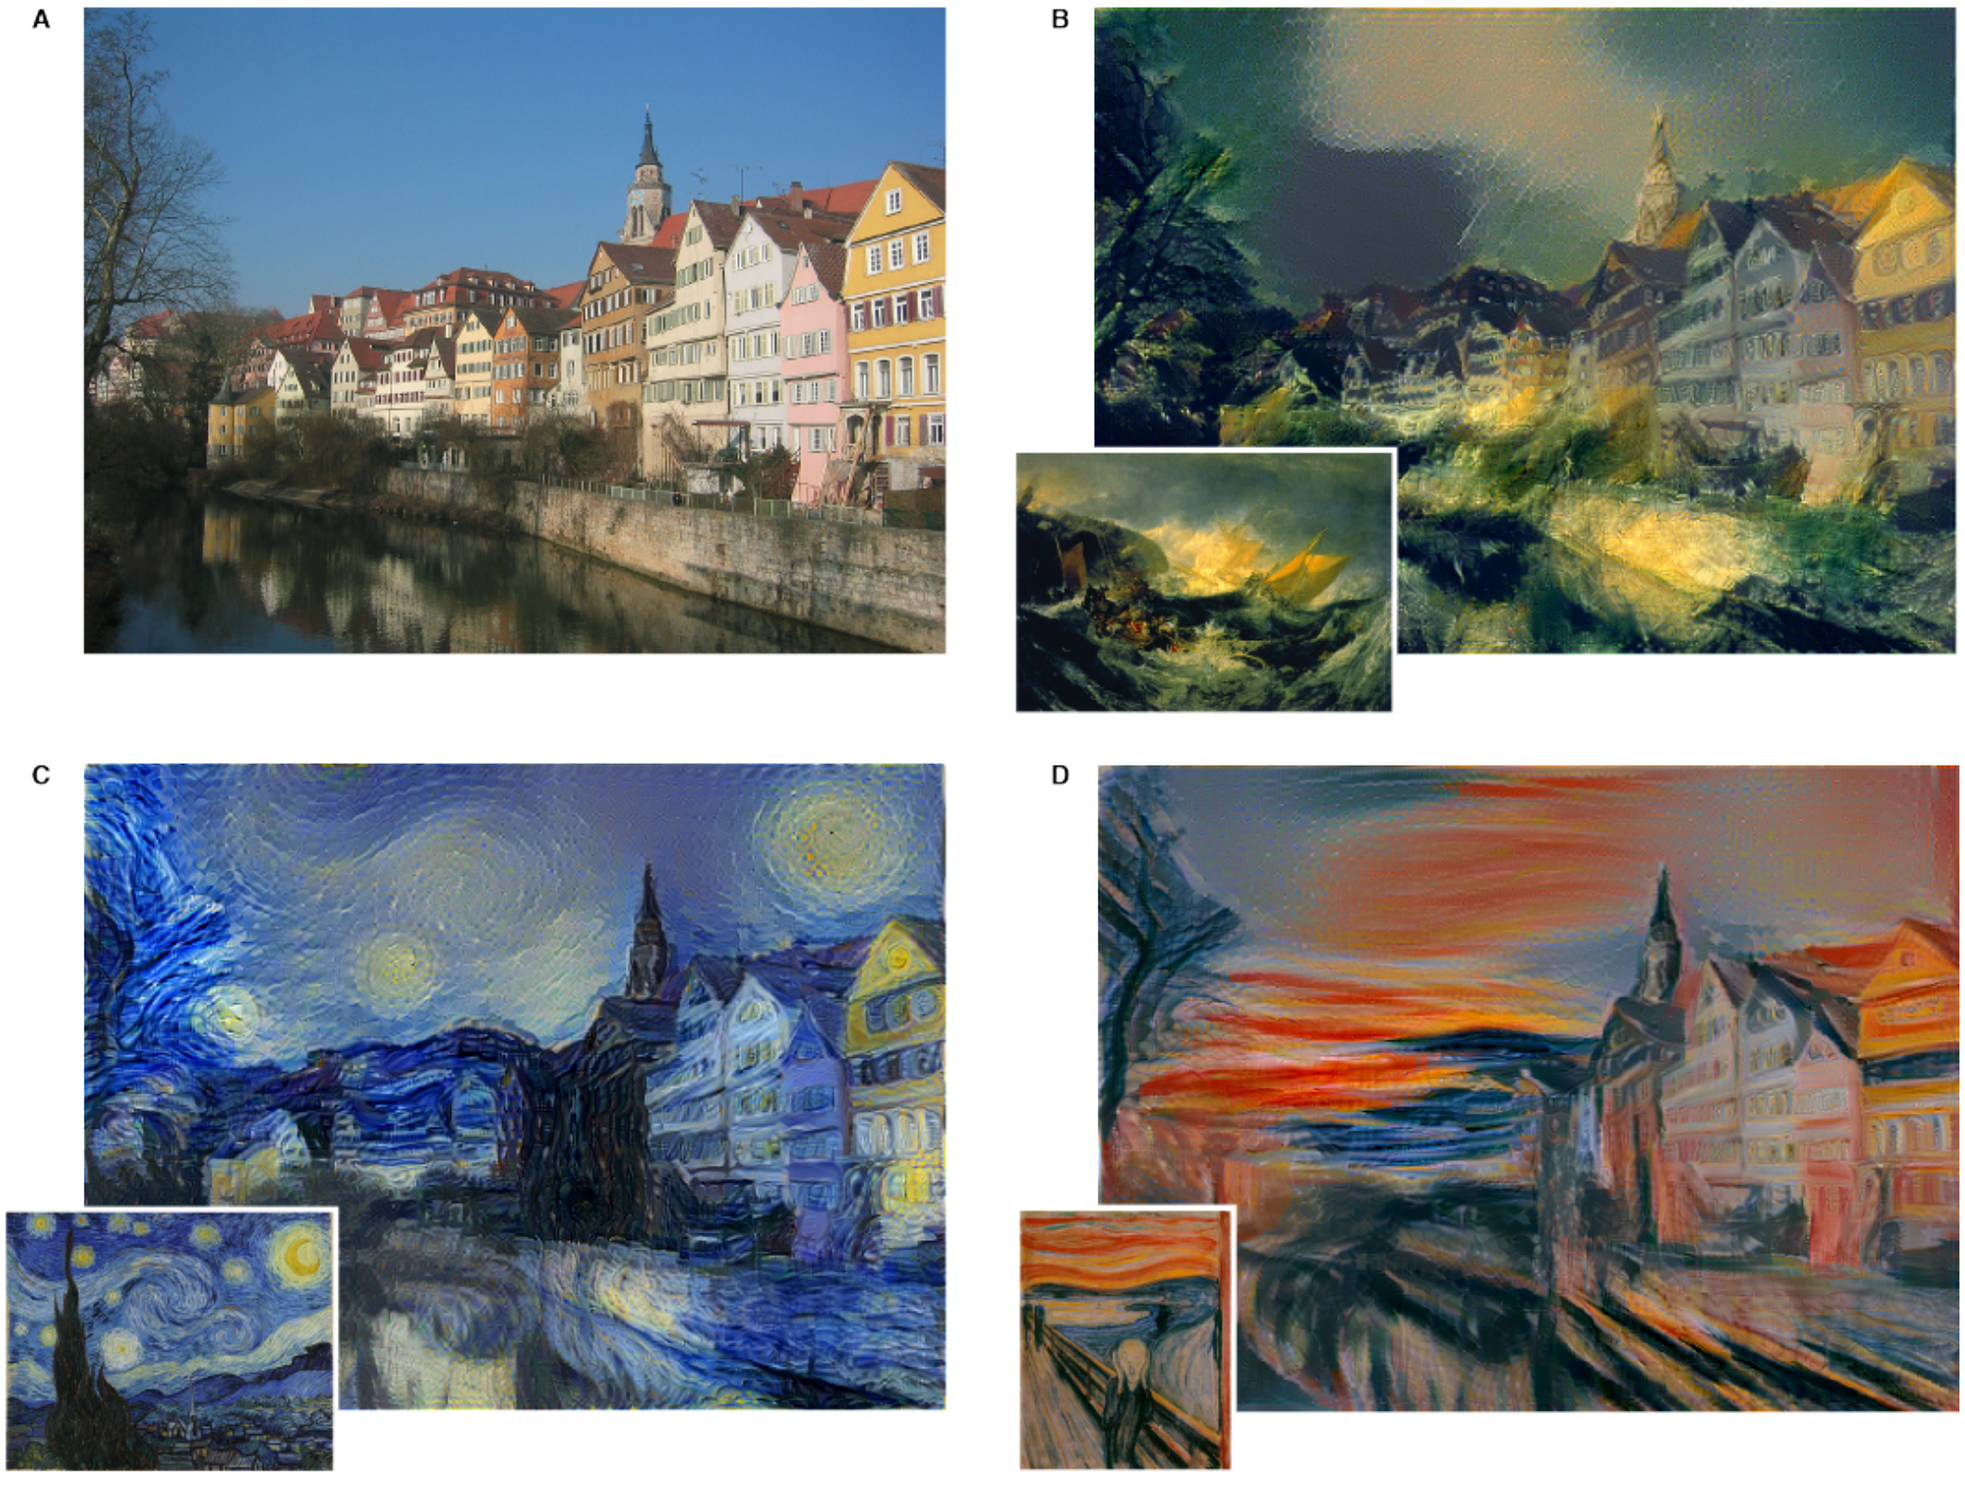

[Gatys, Ecker, Bethge 2015]

Applications of Computer Vision

- Digital Entertainment + Consumer
  - Camera tracking, 3D reconstruction, visual effects, virtual reality, augmented reality, product recognition
- Science and Medicine
  - Visual data analytics, anatomical measurement/analysis, tumour detection
- Engineering and Industry
  - Robotics, self driving cars, reverse engineering, visual servoing, industrial part inspection, OCR, precision agriculture
- Photography/Videography and Editing
  - Face detection, scene recognition, video stabilisation, drone camera, gap filling, image blending, panorama stitching, high dynamic range
- Mapping and Environmental
  - Image registration, 3D building modelling, streetview, numberplate recognition, landmark recognition, species identification

Definitions of Computer Vision #1

“Inverse Computer Graphics”

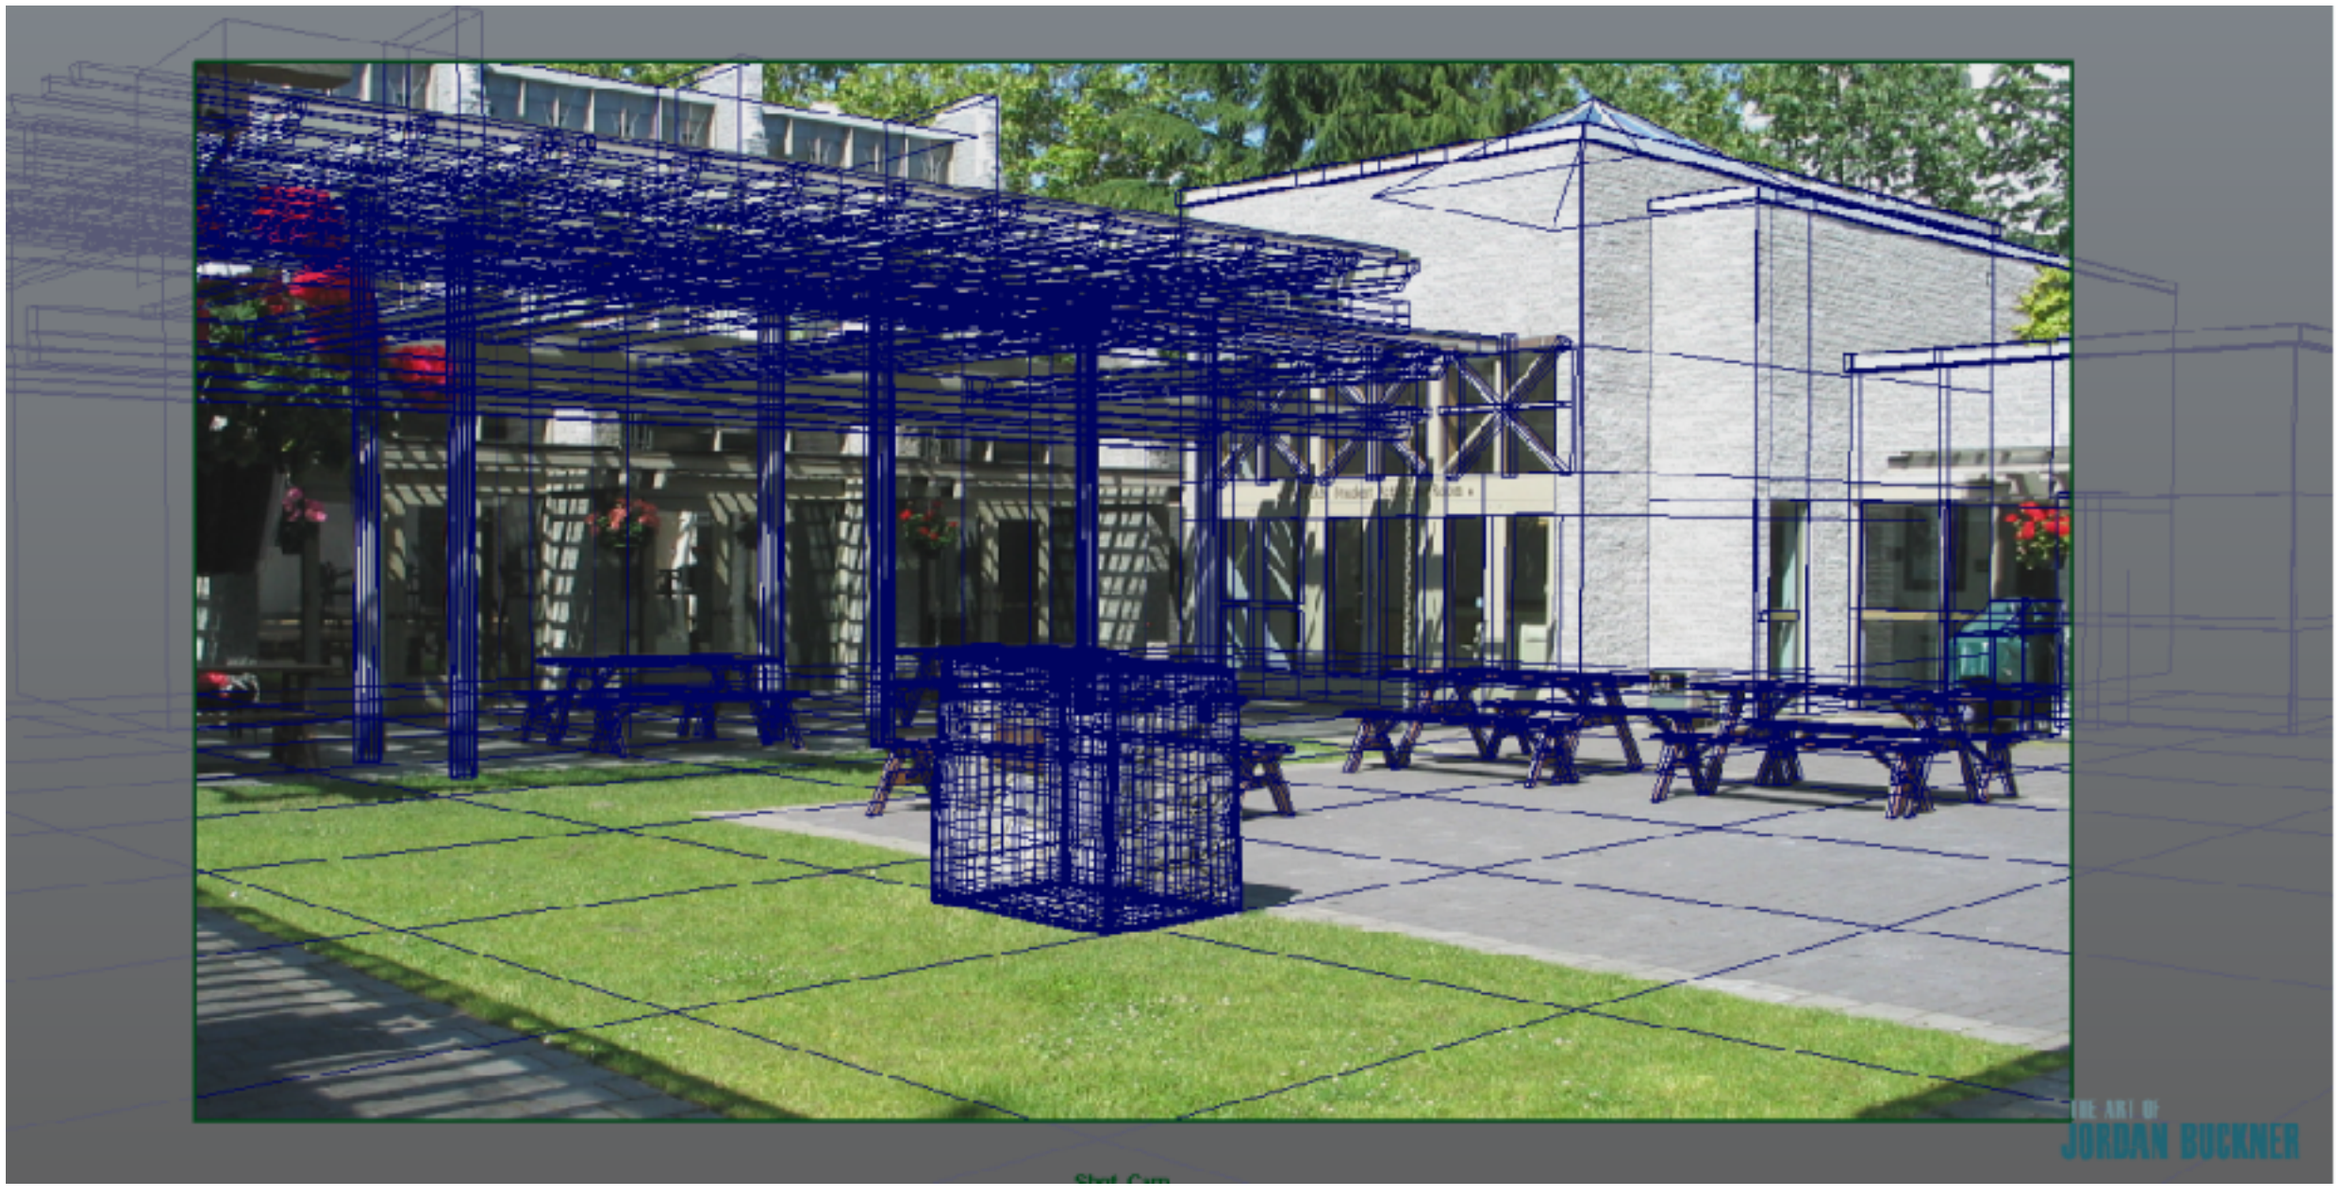

[J. Buckner]

Definitions of Computer Vision #1

"Inverse Computer Graphics"

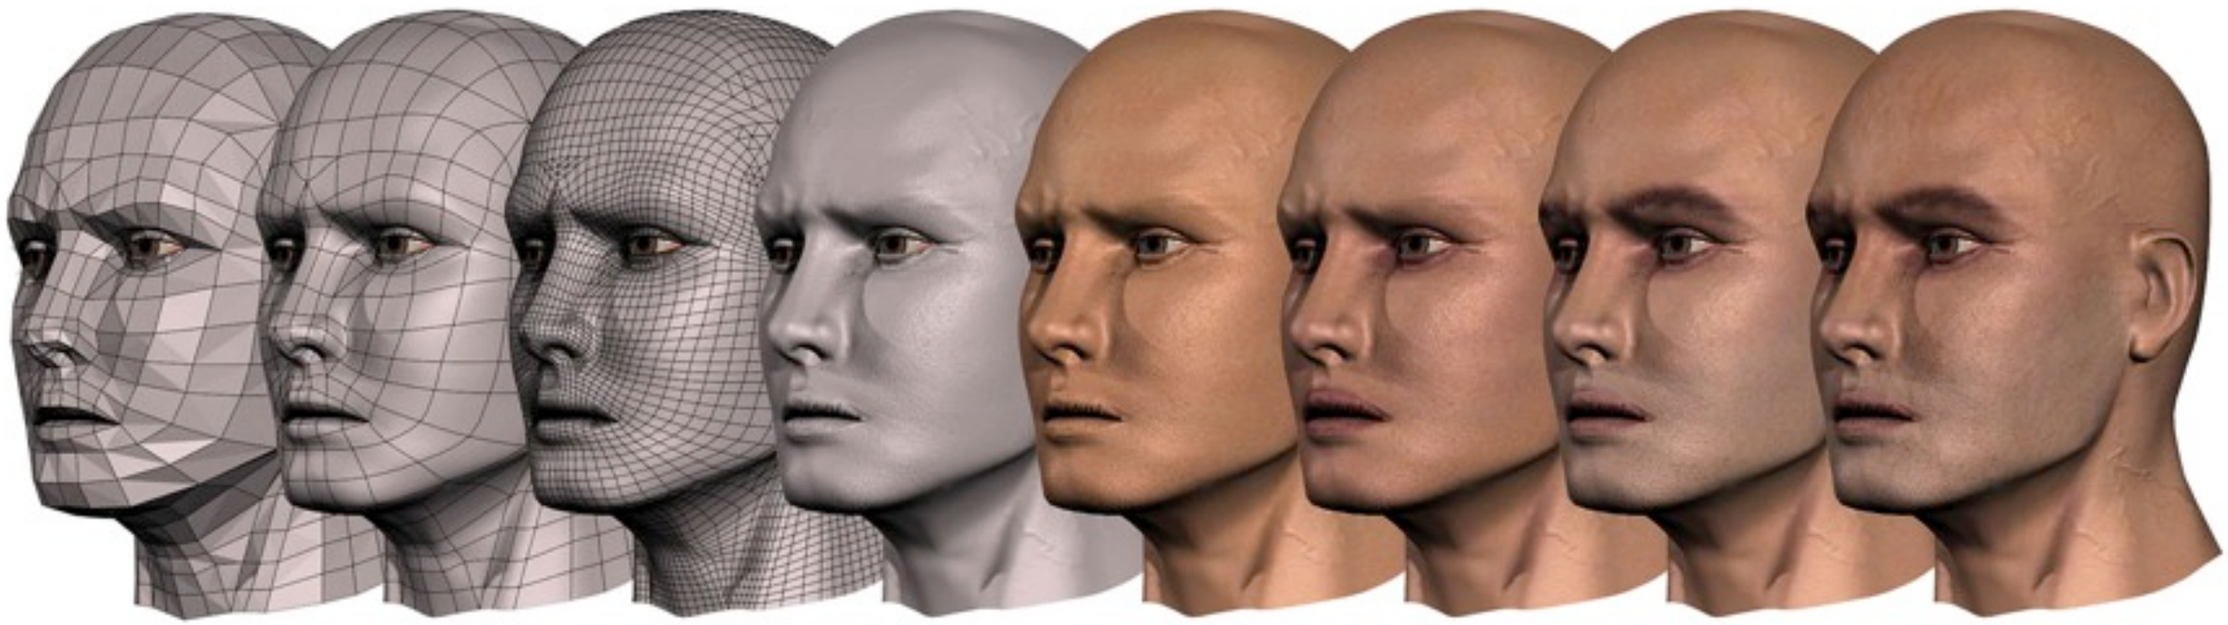

Graphics

Vision

Photometric Capture

- Capture reflectance as well as geometry ("Light Stage")

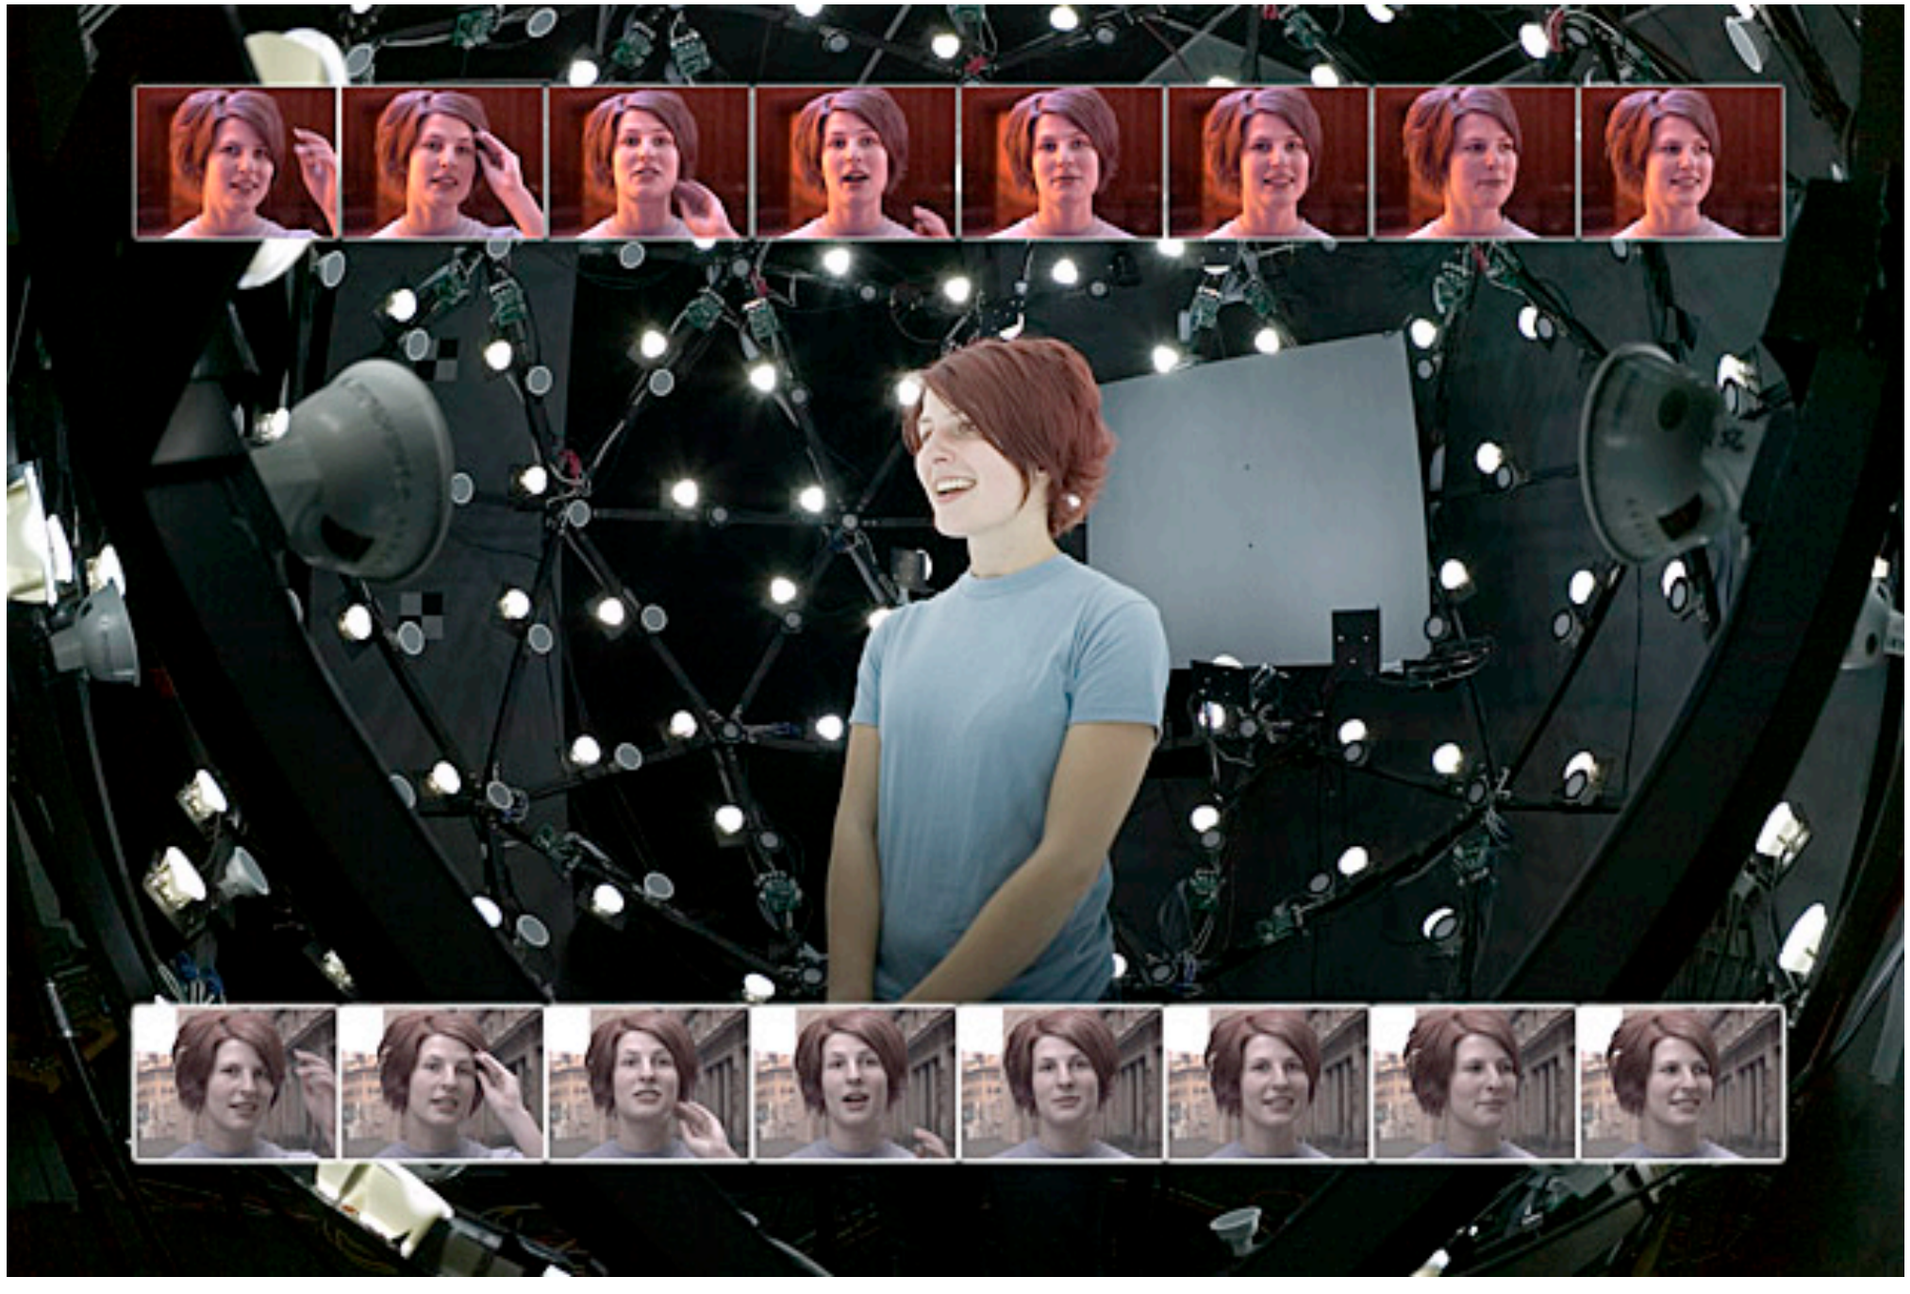

[Paul Debevec, USC]

Definitions of Computer Vision #2

“Replicate Human Vision”

=

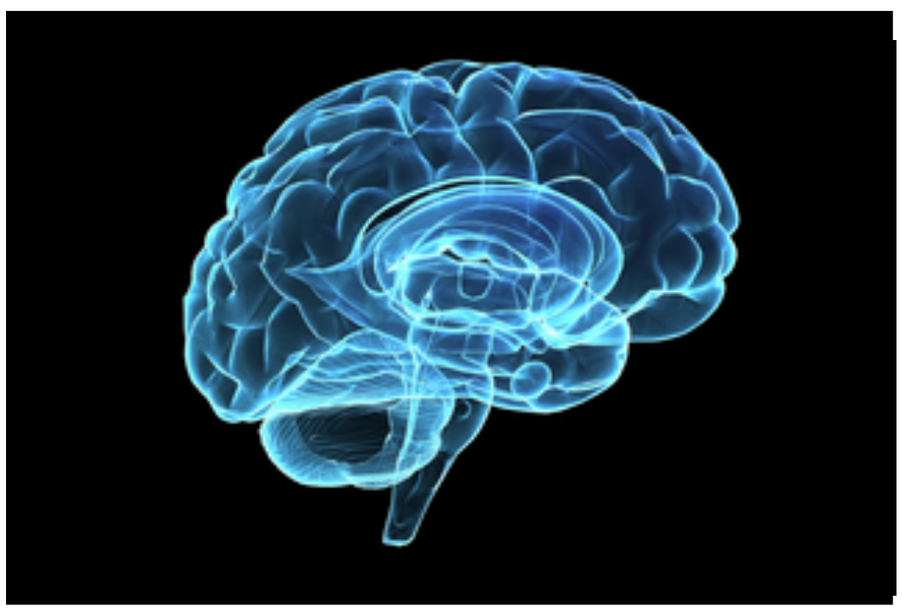

=

Definitions of Computer Vision #2

“Replicate Human Vision”

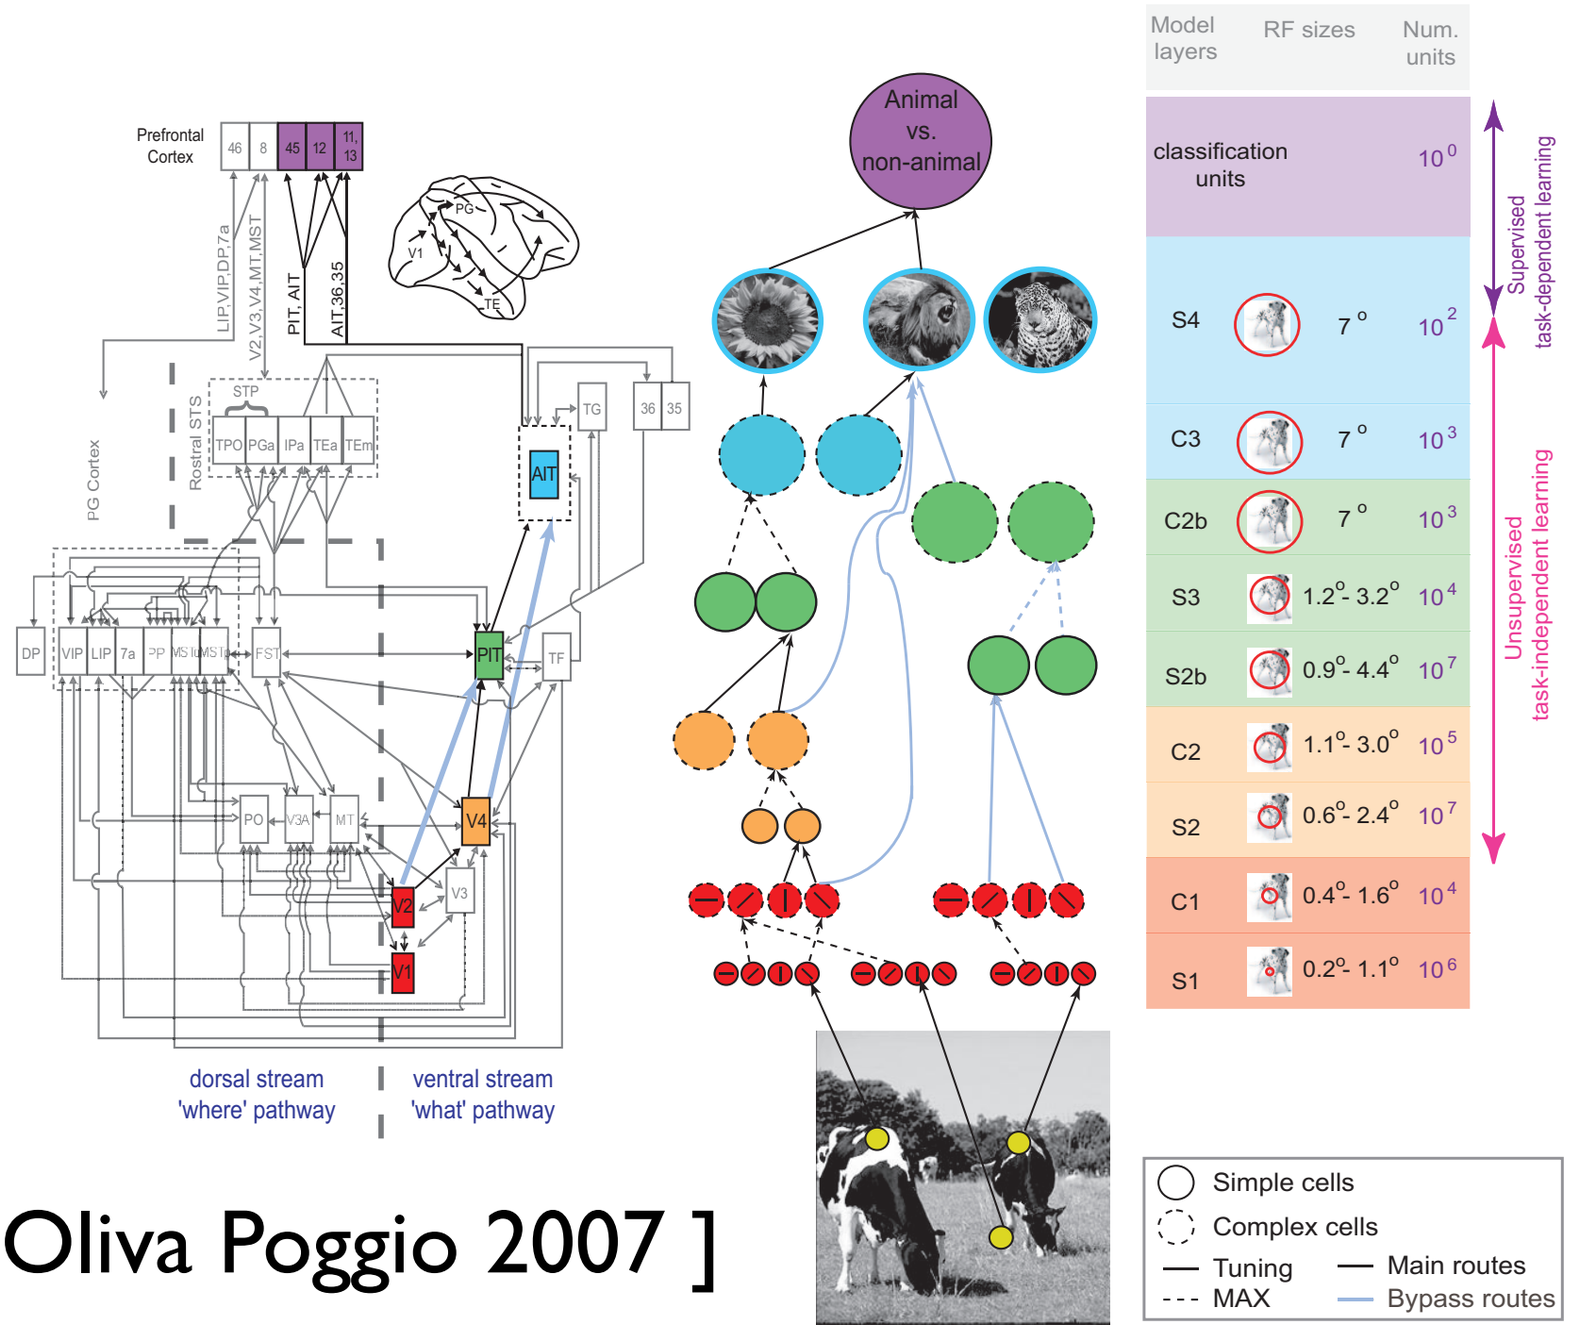

[Serre Oliva Poggio 2007]

ImageNet

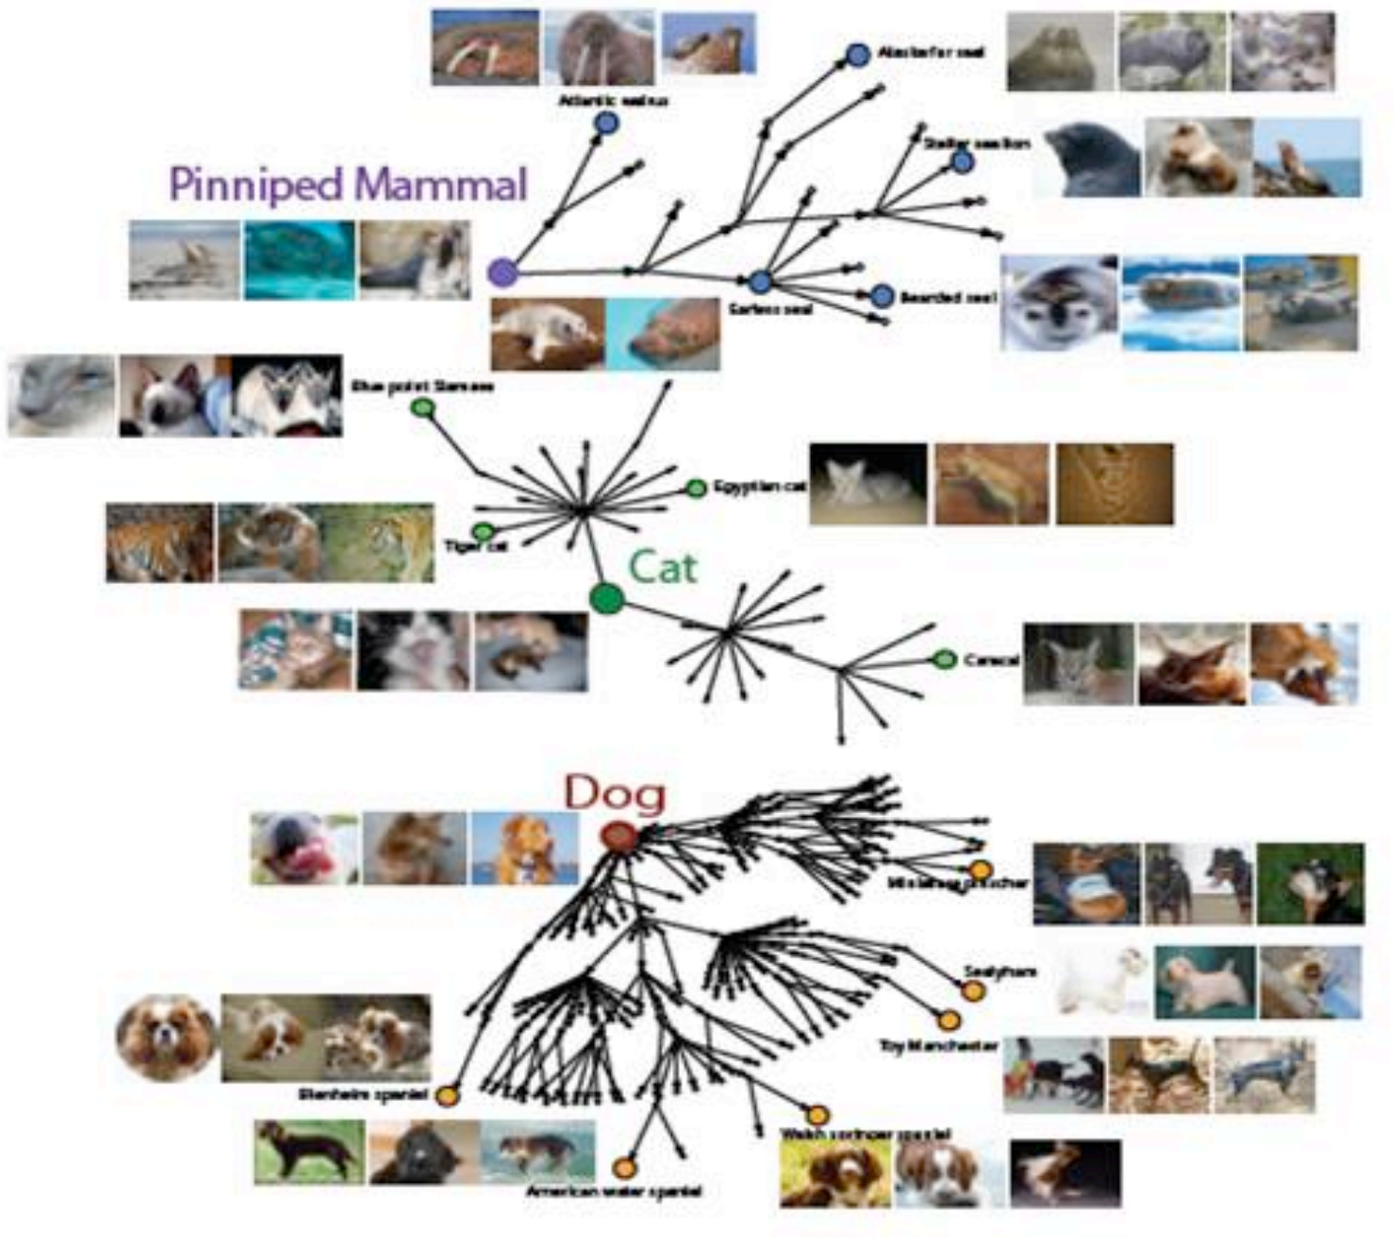

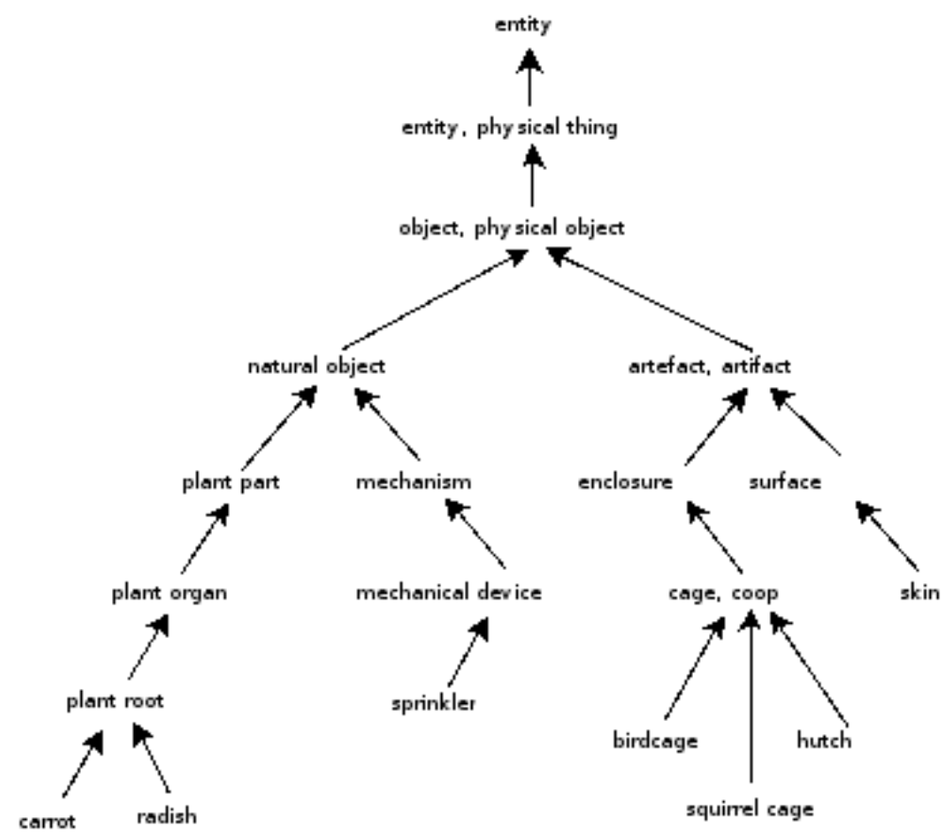

Figure 1. "is a" relation example

15 million images in 22,000 categories

[F. F. Li et al]

ImageNet Classification via CNN

- "Alexnet" gave breakthrough results on the ImageNet 2012 Large Scale Visual Recognition Challenge (ILSVRC 2012)

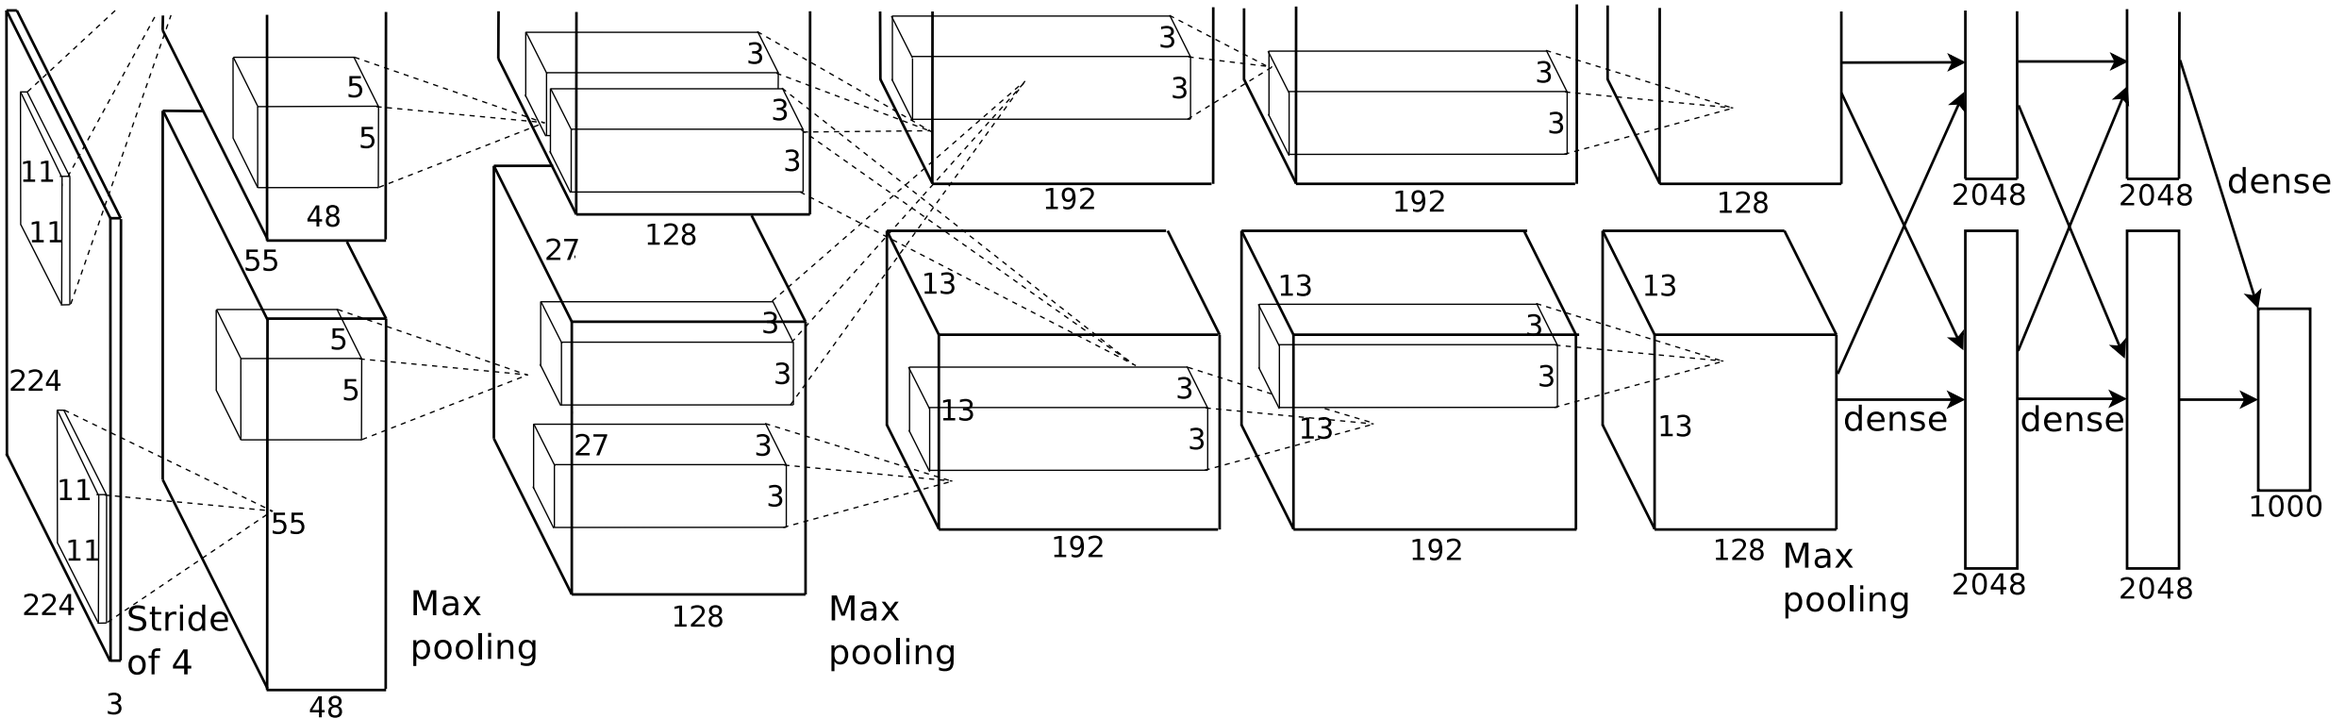

[Krizhevsky, Sutskever, Hinton 2012]

Definitions of Computer Vision #3

"Image/Video Understanding"

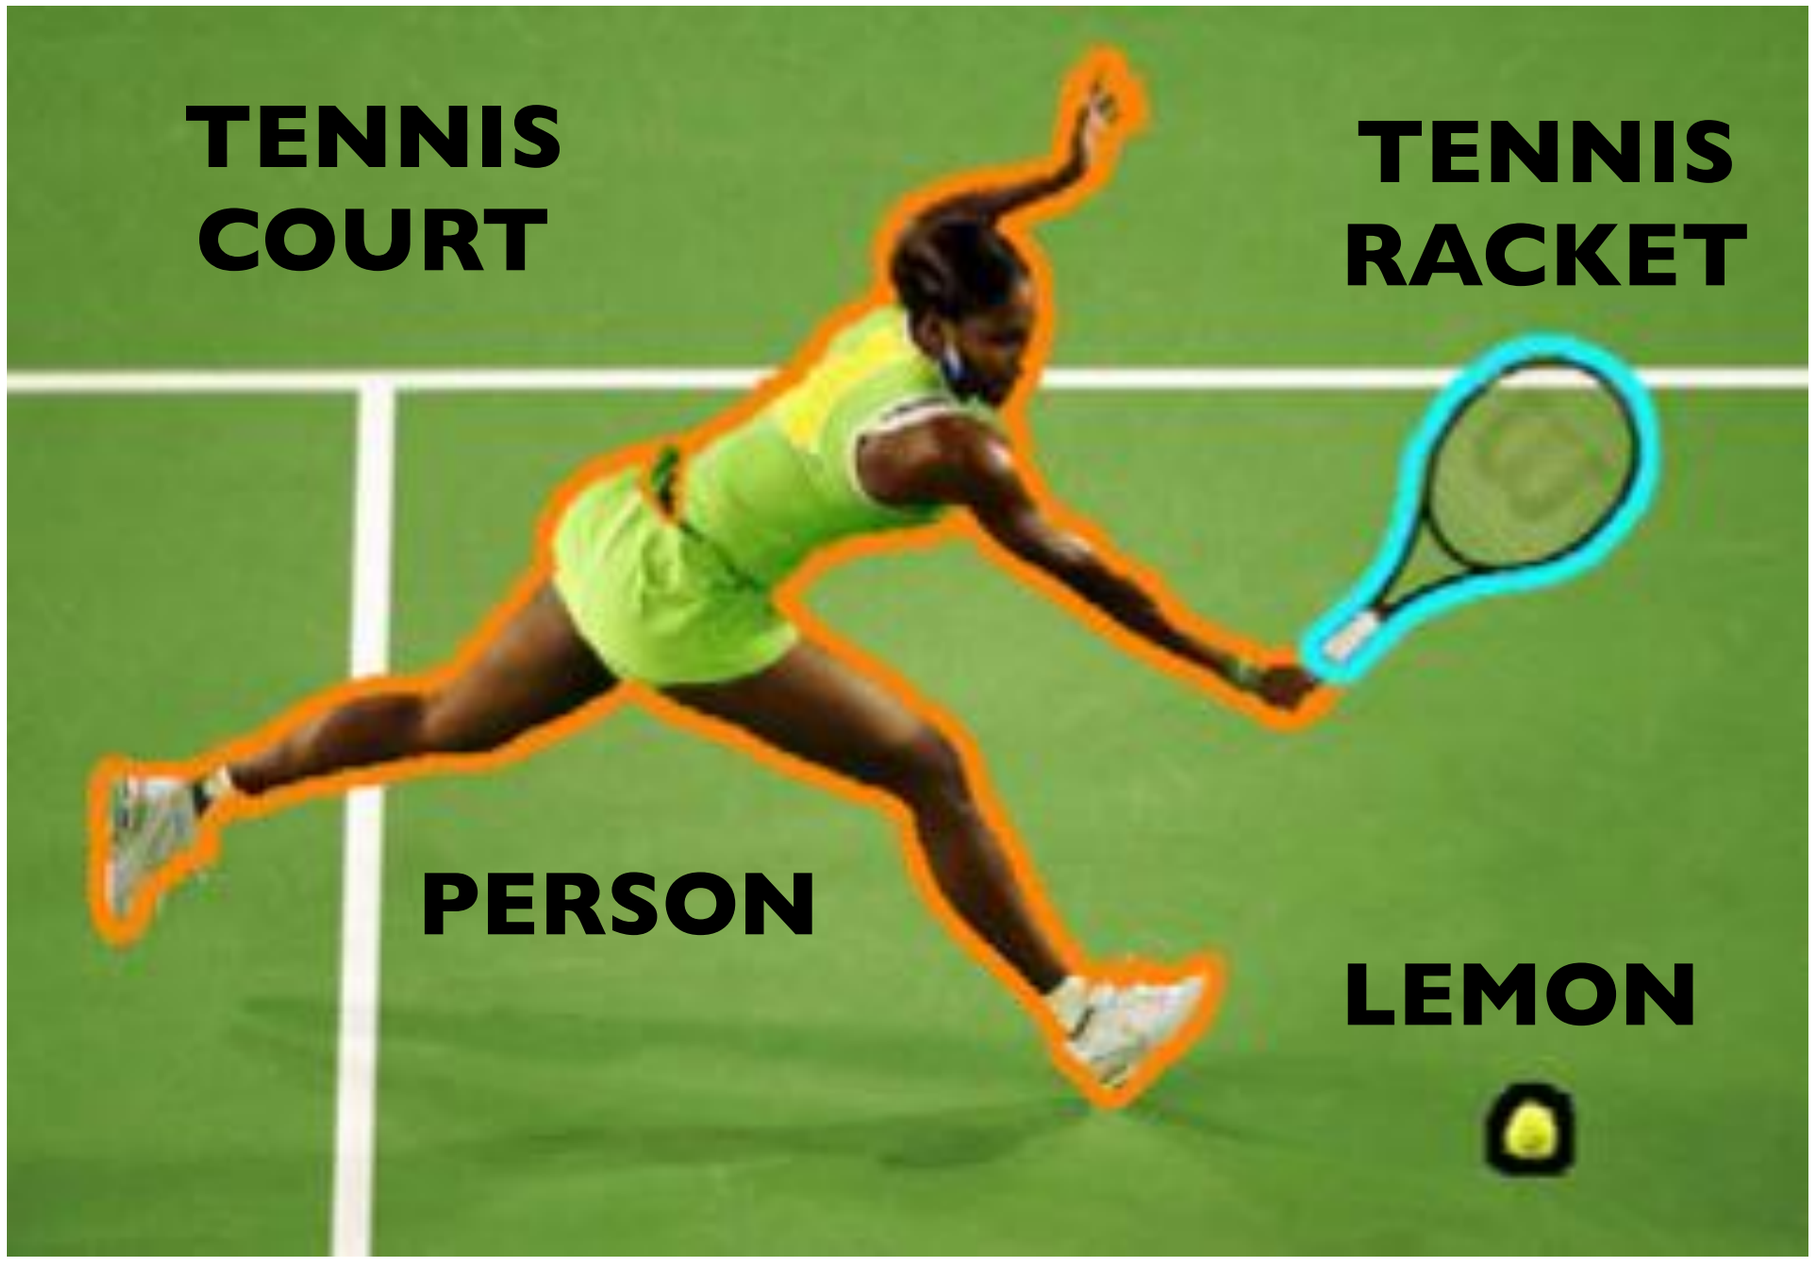

[Rabinovich, Galleguillos, Wiewiora, Belongie 2007]

Definitions of Computer Vision #3

“Image/Video Understanding”

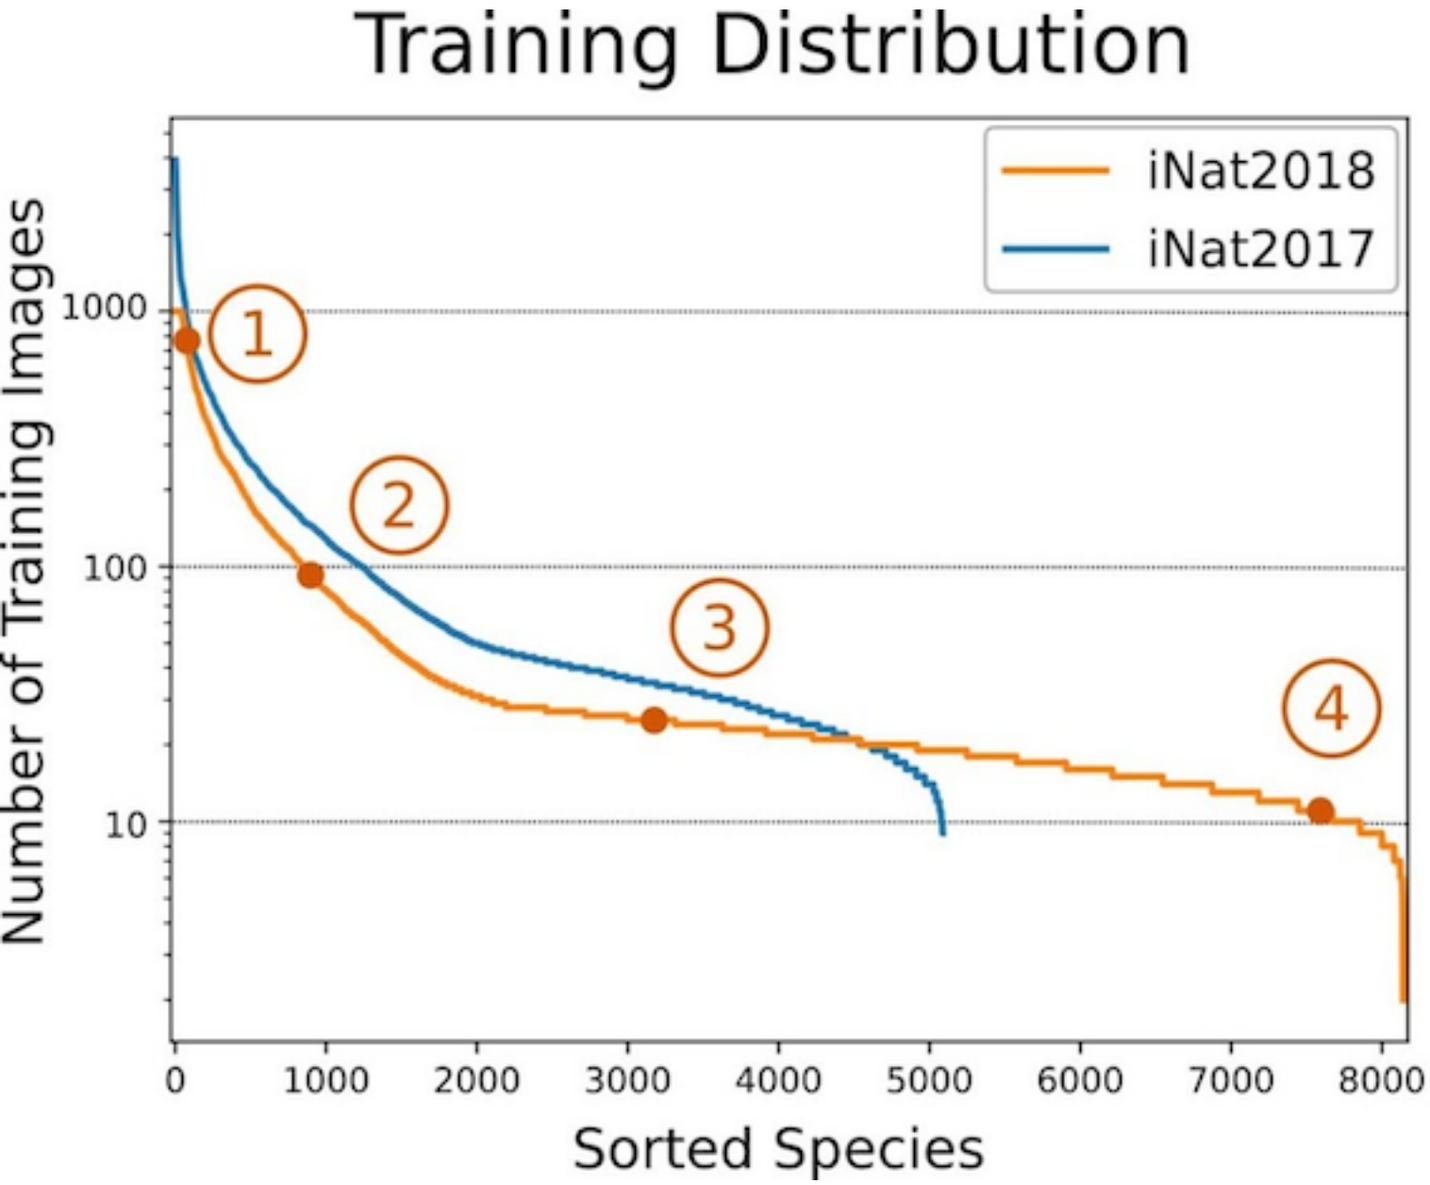

iNaturalist Challenge 2018

1 Cooper's Hawk

2 American Bison

3 Mallow Bindweed

4 Island Fox

[Van Horn, Mac Aodha]

Definitions of Computer Vision #3

"Image/Video Understanding"

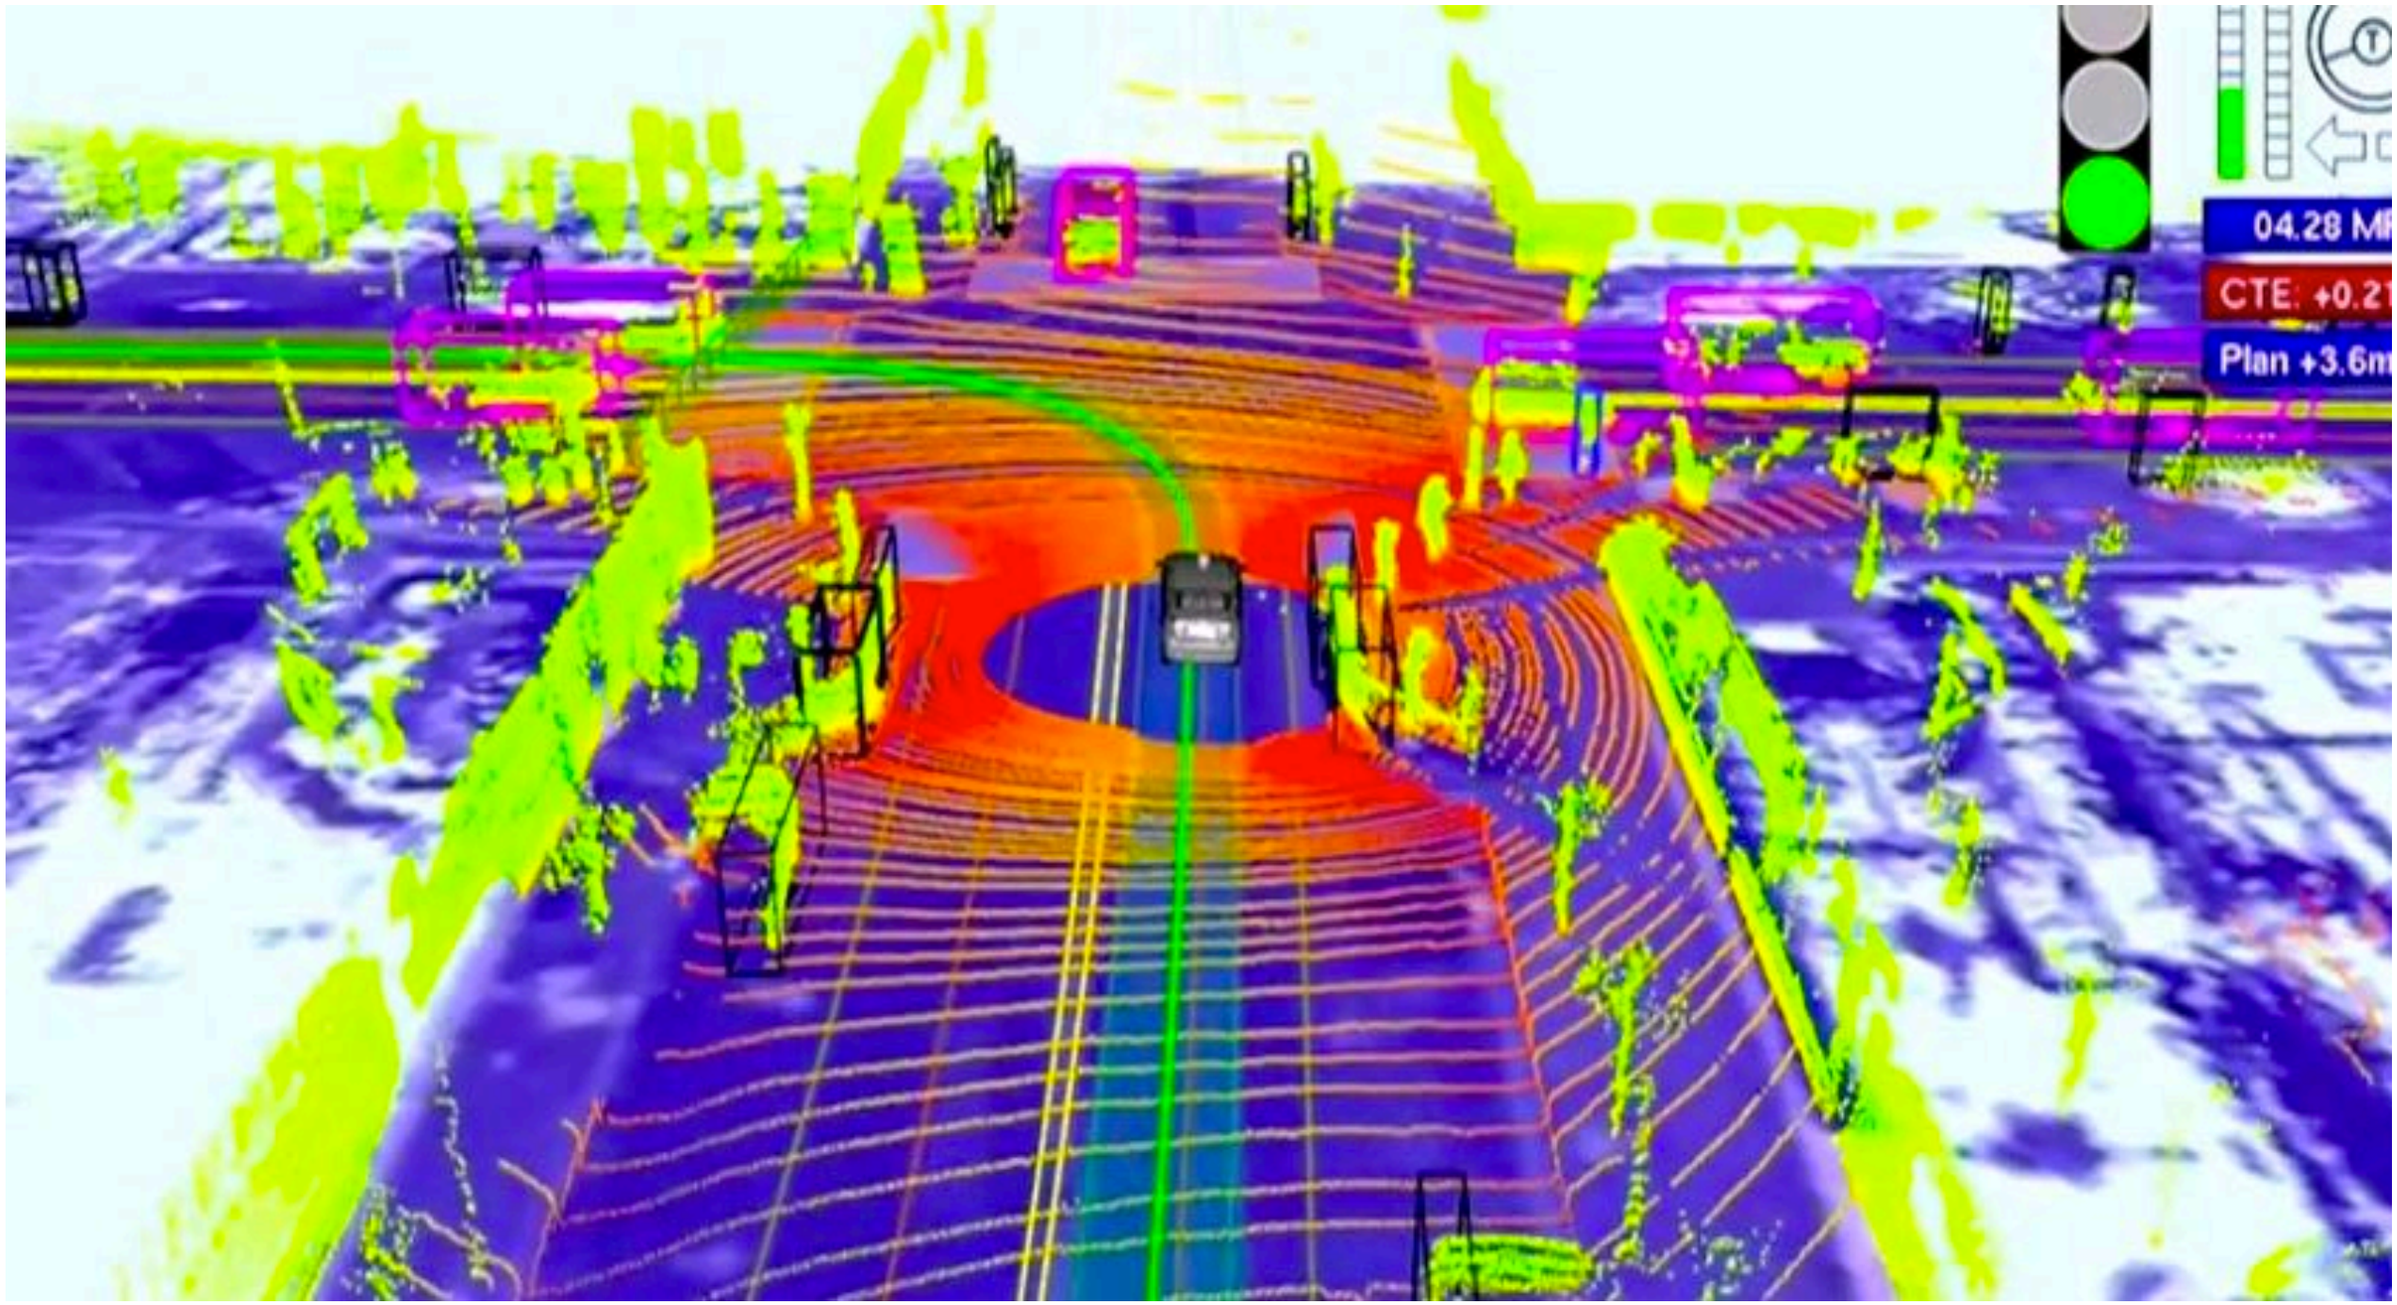

[Thrun Urmson Google]

This Course

- Computer Vision, with emphasis on **visual geometry + learning** (roughly 50-50 split between the two)
- 10 lectures, + office hours
- 4 projects, equally weighted
- **Project 1**: Feature Extraction and Matching
- **Project 2**: Panoramic Image Stitching
- **Project 3**: Image Classifi cation using CIFAR10
- **Project 4**: Pixel Labelling Project
- Projects will use iPython notebooks (e.g., Jupyter, Colab)
- Numpy for numerics
- Tensorflow for machine learning

Schedule

Date	**Lecture Topics**	**Project**
03/31	Introduction, Image Formation	
04/07	Filtering and Pyramids, Features and Matching	P1 assigned
04/14	Planar Geometry. 2-view Alignment, RANSAC	
04/21	Epipolar + Multiview Geometry, SFM/SLAM	P2 assigned P1 due 4/24
04/28	Dense correspondence, Stereo, Flow	

Schedule

Date	Lecture Topics	Project
05/05	Machine Learning, NN, SVM, Decision Trees, Boosting	P3 assigned P2 due 5/8
05/12	Linear/Logistic Regression, NNets, CNNs, Backprop	
05/19	Per-Pixel Labelling, Depth, Flow, Segmentation	P4 assigned P3 due 5/22
05/26	Object Detection, Applications + Architectures	
6/2	Computational Photography, Past Present and Future	P4 due

Recommended Text 1

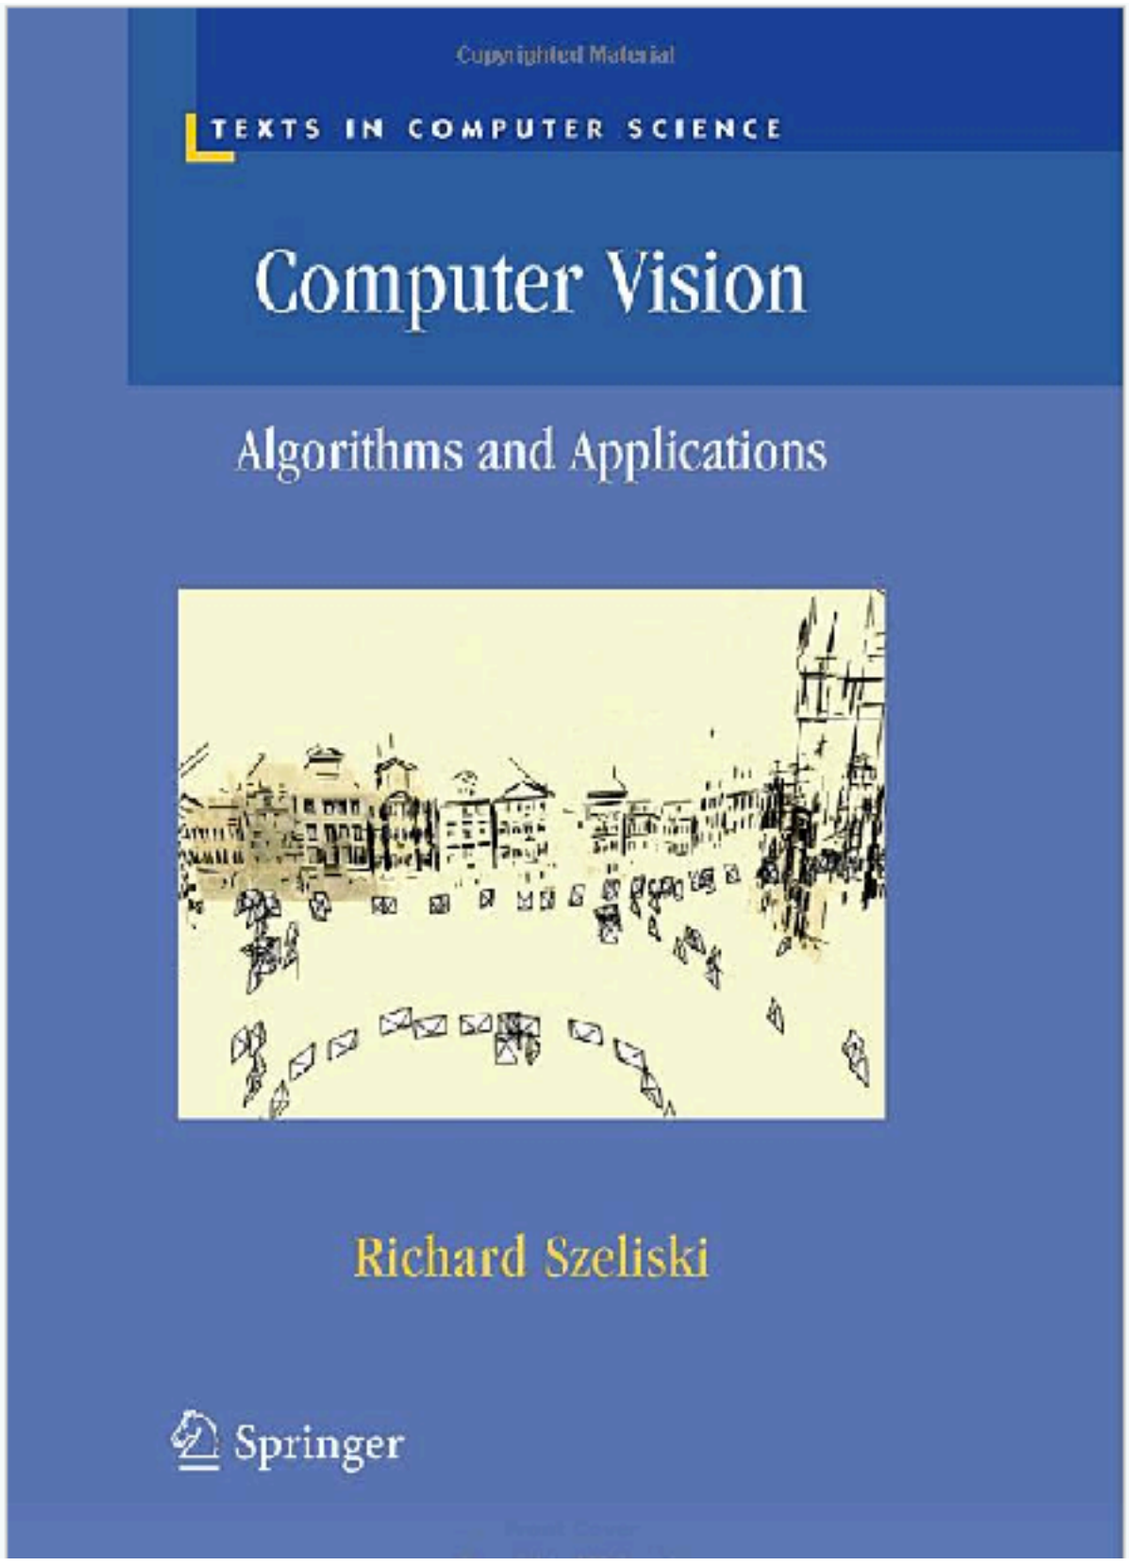

Computer Vision: Algorithms and Applications

Richard Szeliski

http://szeliski.org/Book

Core textbook for the course. Good coverage of most topics, oriented around practical applications

Computer Vision: Szeliski

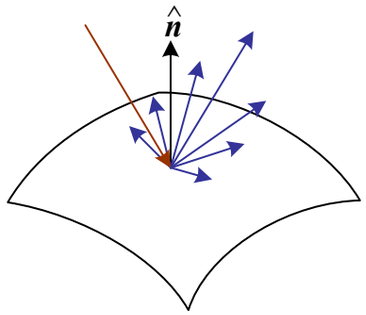

2. Image Formation

3. Image Processing

4. Features

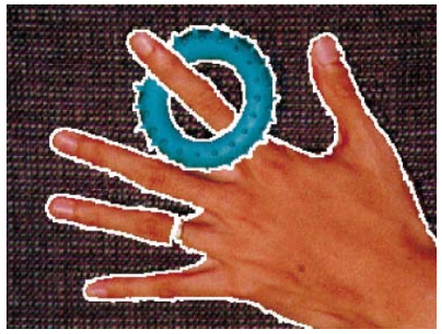

5. Segmentation

6-7. Structure from Motion

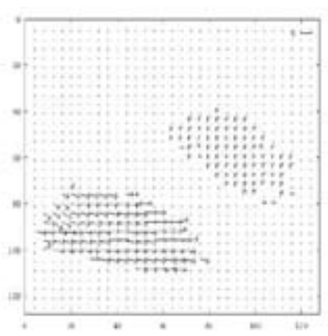

8. Motion

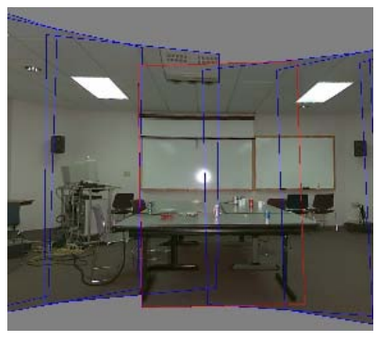

9. Stitching

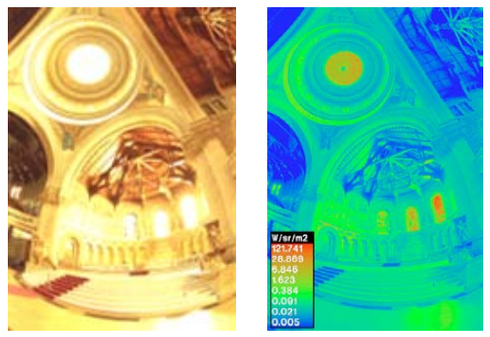

10. Computational Photography

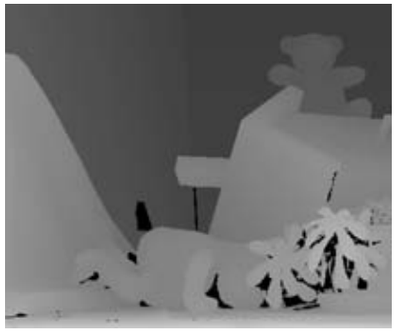

11. Stereo

12. 3D Shape

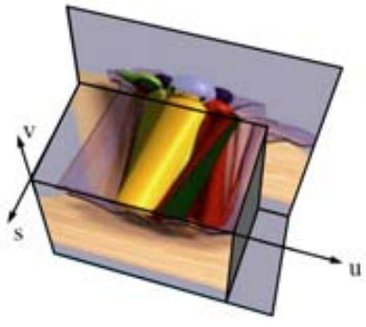

13. Image-based Rendering

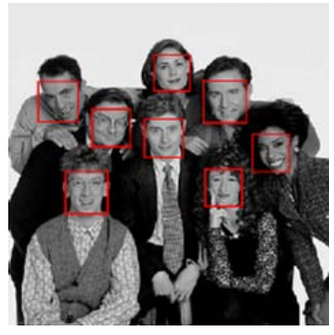

14. Recognition

Recommended Text 2

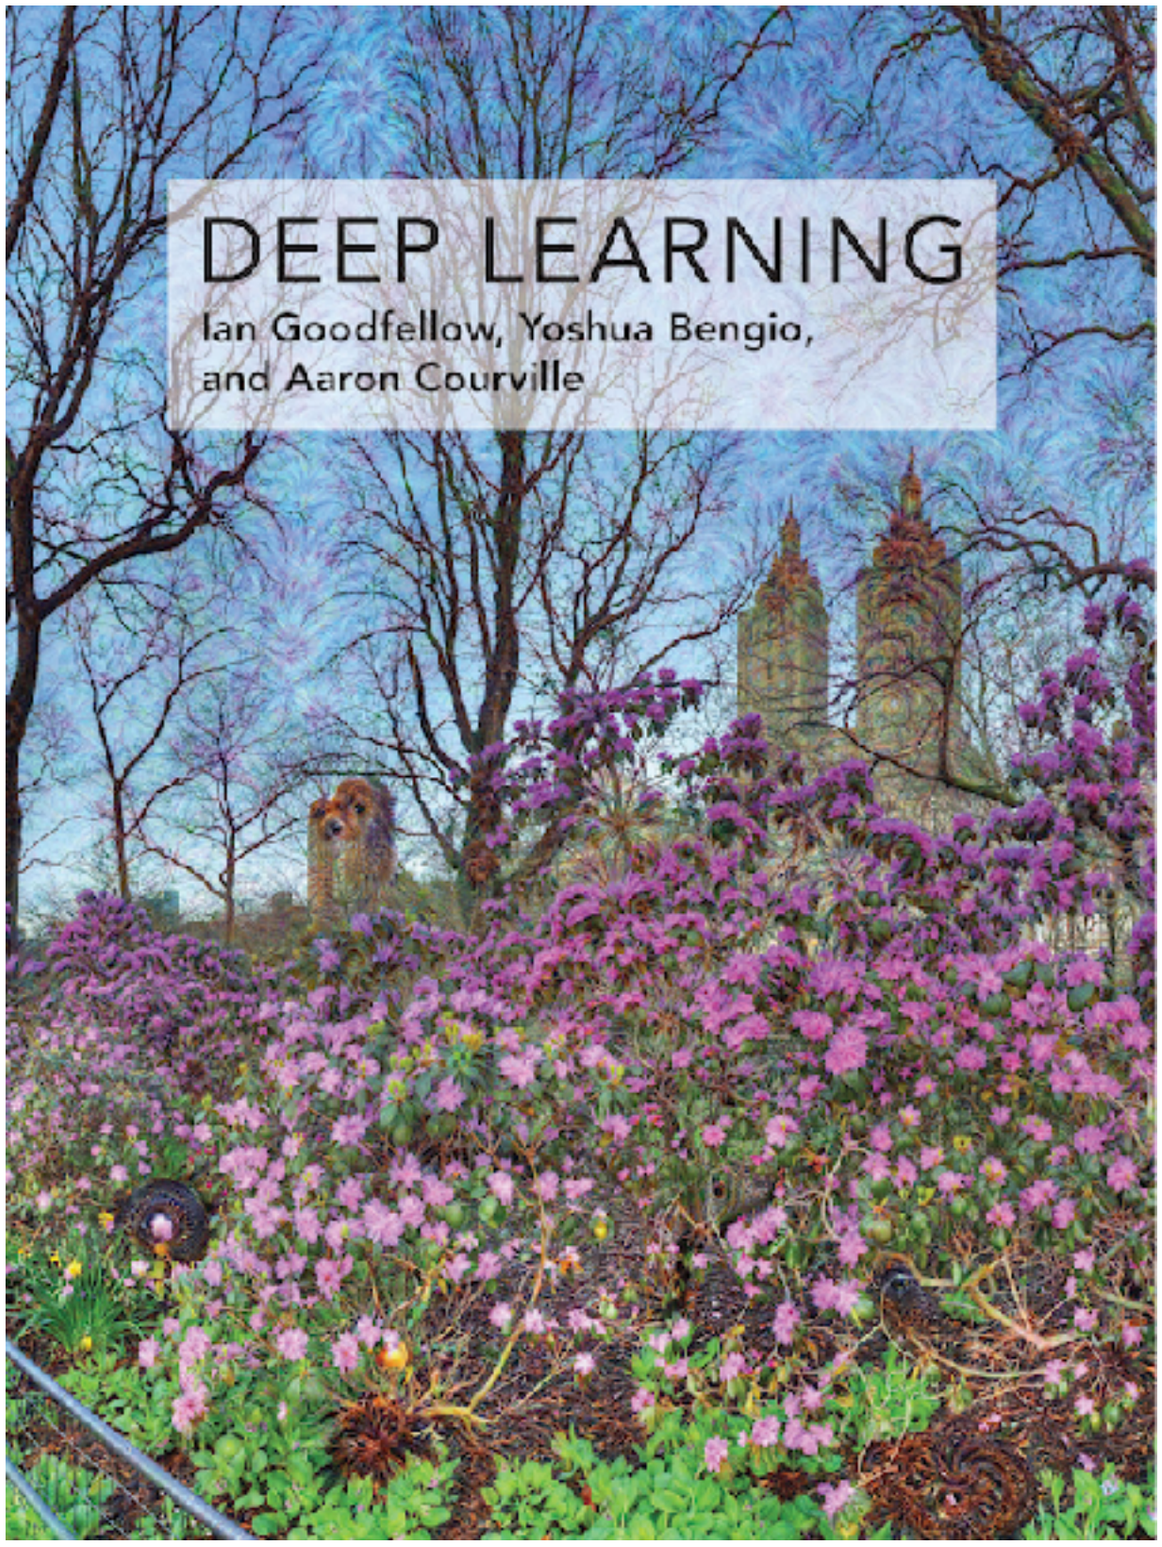

Deep Learning: Goodfellow, Bengio, Courville

deeplearningbook.org

Background maths + probability, practical deep nets, deep learning research

Also cs231n.stanford.edu
— CNNs for Vision

Inceptionism

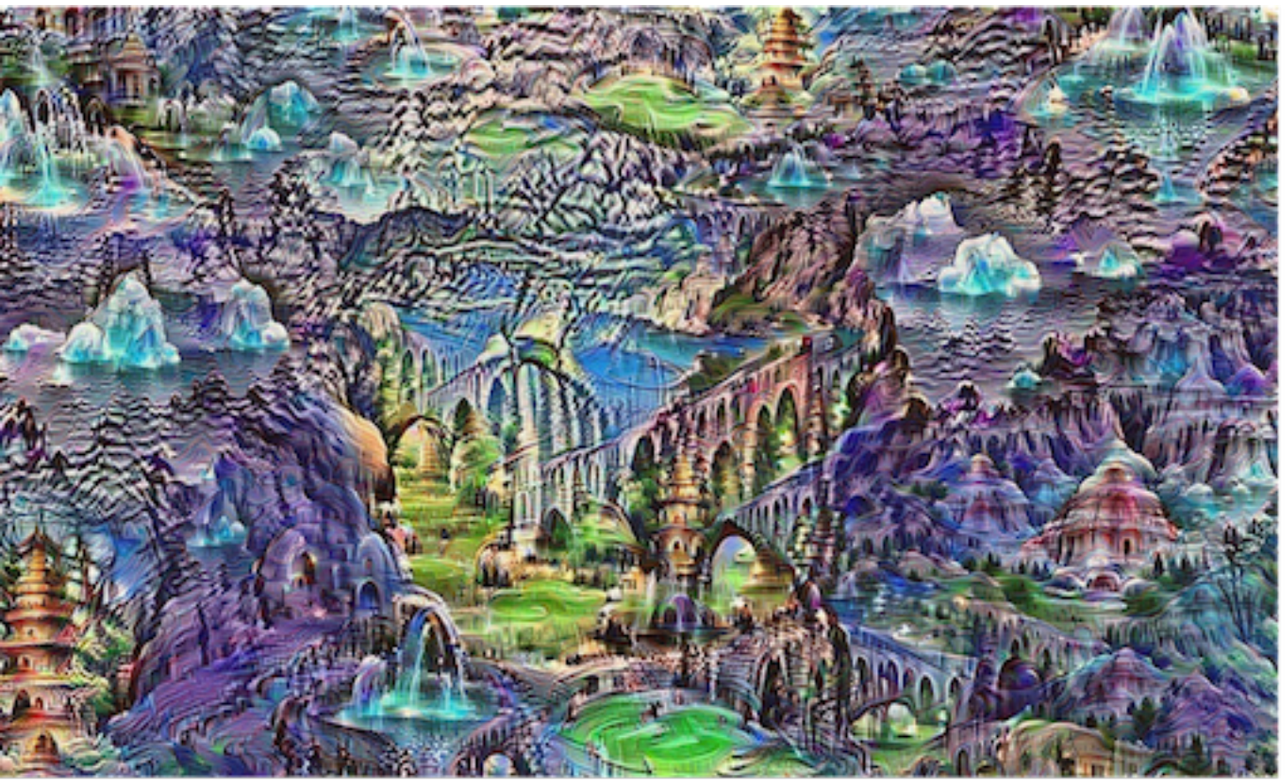

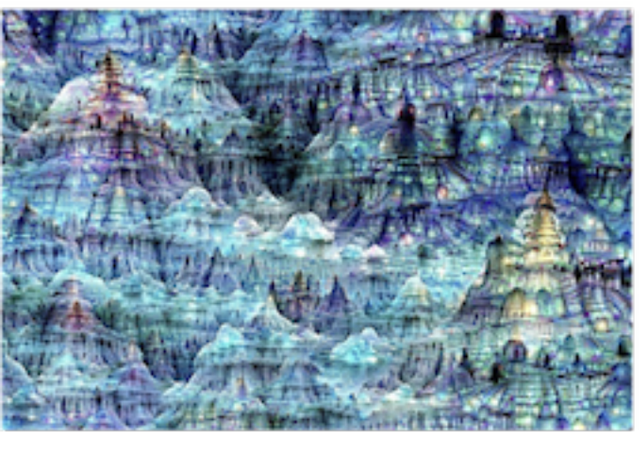
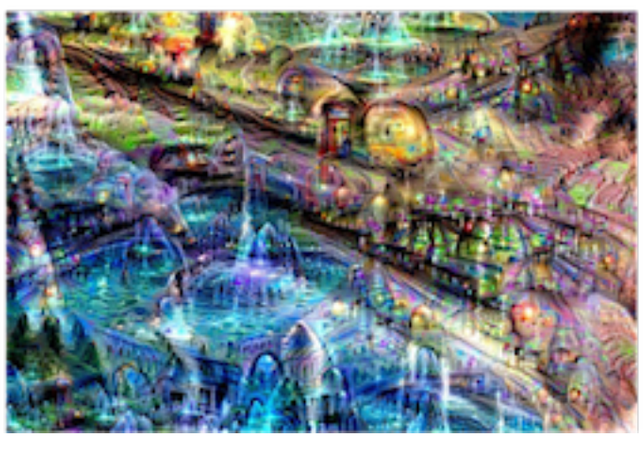

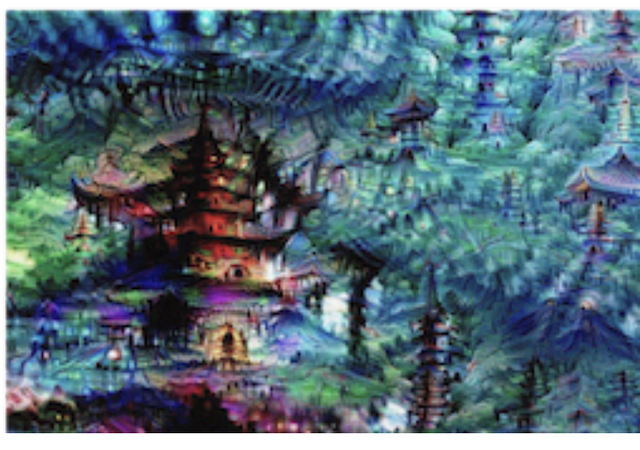

[Mordvintsev, Olah, Tyka 2015]

Next Lecture

- Cameras + Image Formation

Try getting Jupyter/Colab up and running,
and work through Justin Johnson's Python intro:
http://cs231n.github.io/python-numpy-tutorial