Computer Vision

CSE P576

Dr. Matthew Brown

The Course

- People
  - Matthew Brown, Richard Newcombe
  - Rob Gens, David Rosen
  - Tanner Schmidt, Yun Hsuan Su (Melody)
- Time and location
  - Lectures: Thursdays 6:30-9:20pm JHN 075
  - Office hours: Wednesday 5:00-6:00pm CSE 674 (or by appointment)
- Evaluation
  - 4 projects, equally weighted
- Resources
  - https://courses.cs.washington.edu/courses/csep576/18sp
  - Google group: csep576-18sp-discussion
  - Book 1: "Computer Vision", Szeliski,
  - Book 2: "Deep Learning", Goodfellow et al.

Face Detection

[Motorola]

Camera Tracking

[Boujou -- Vicon/OMG]

3D Reconstruction

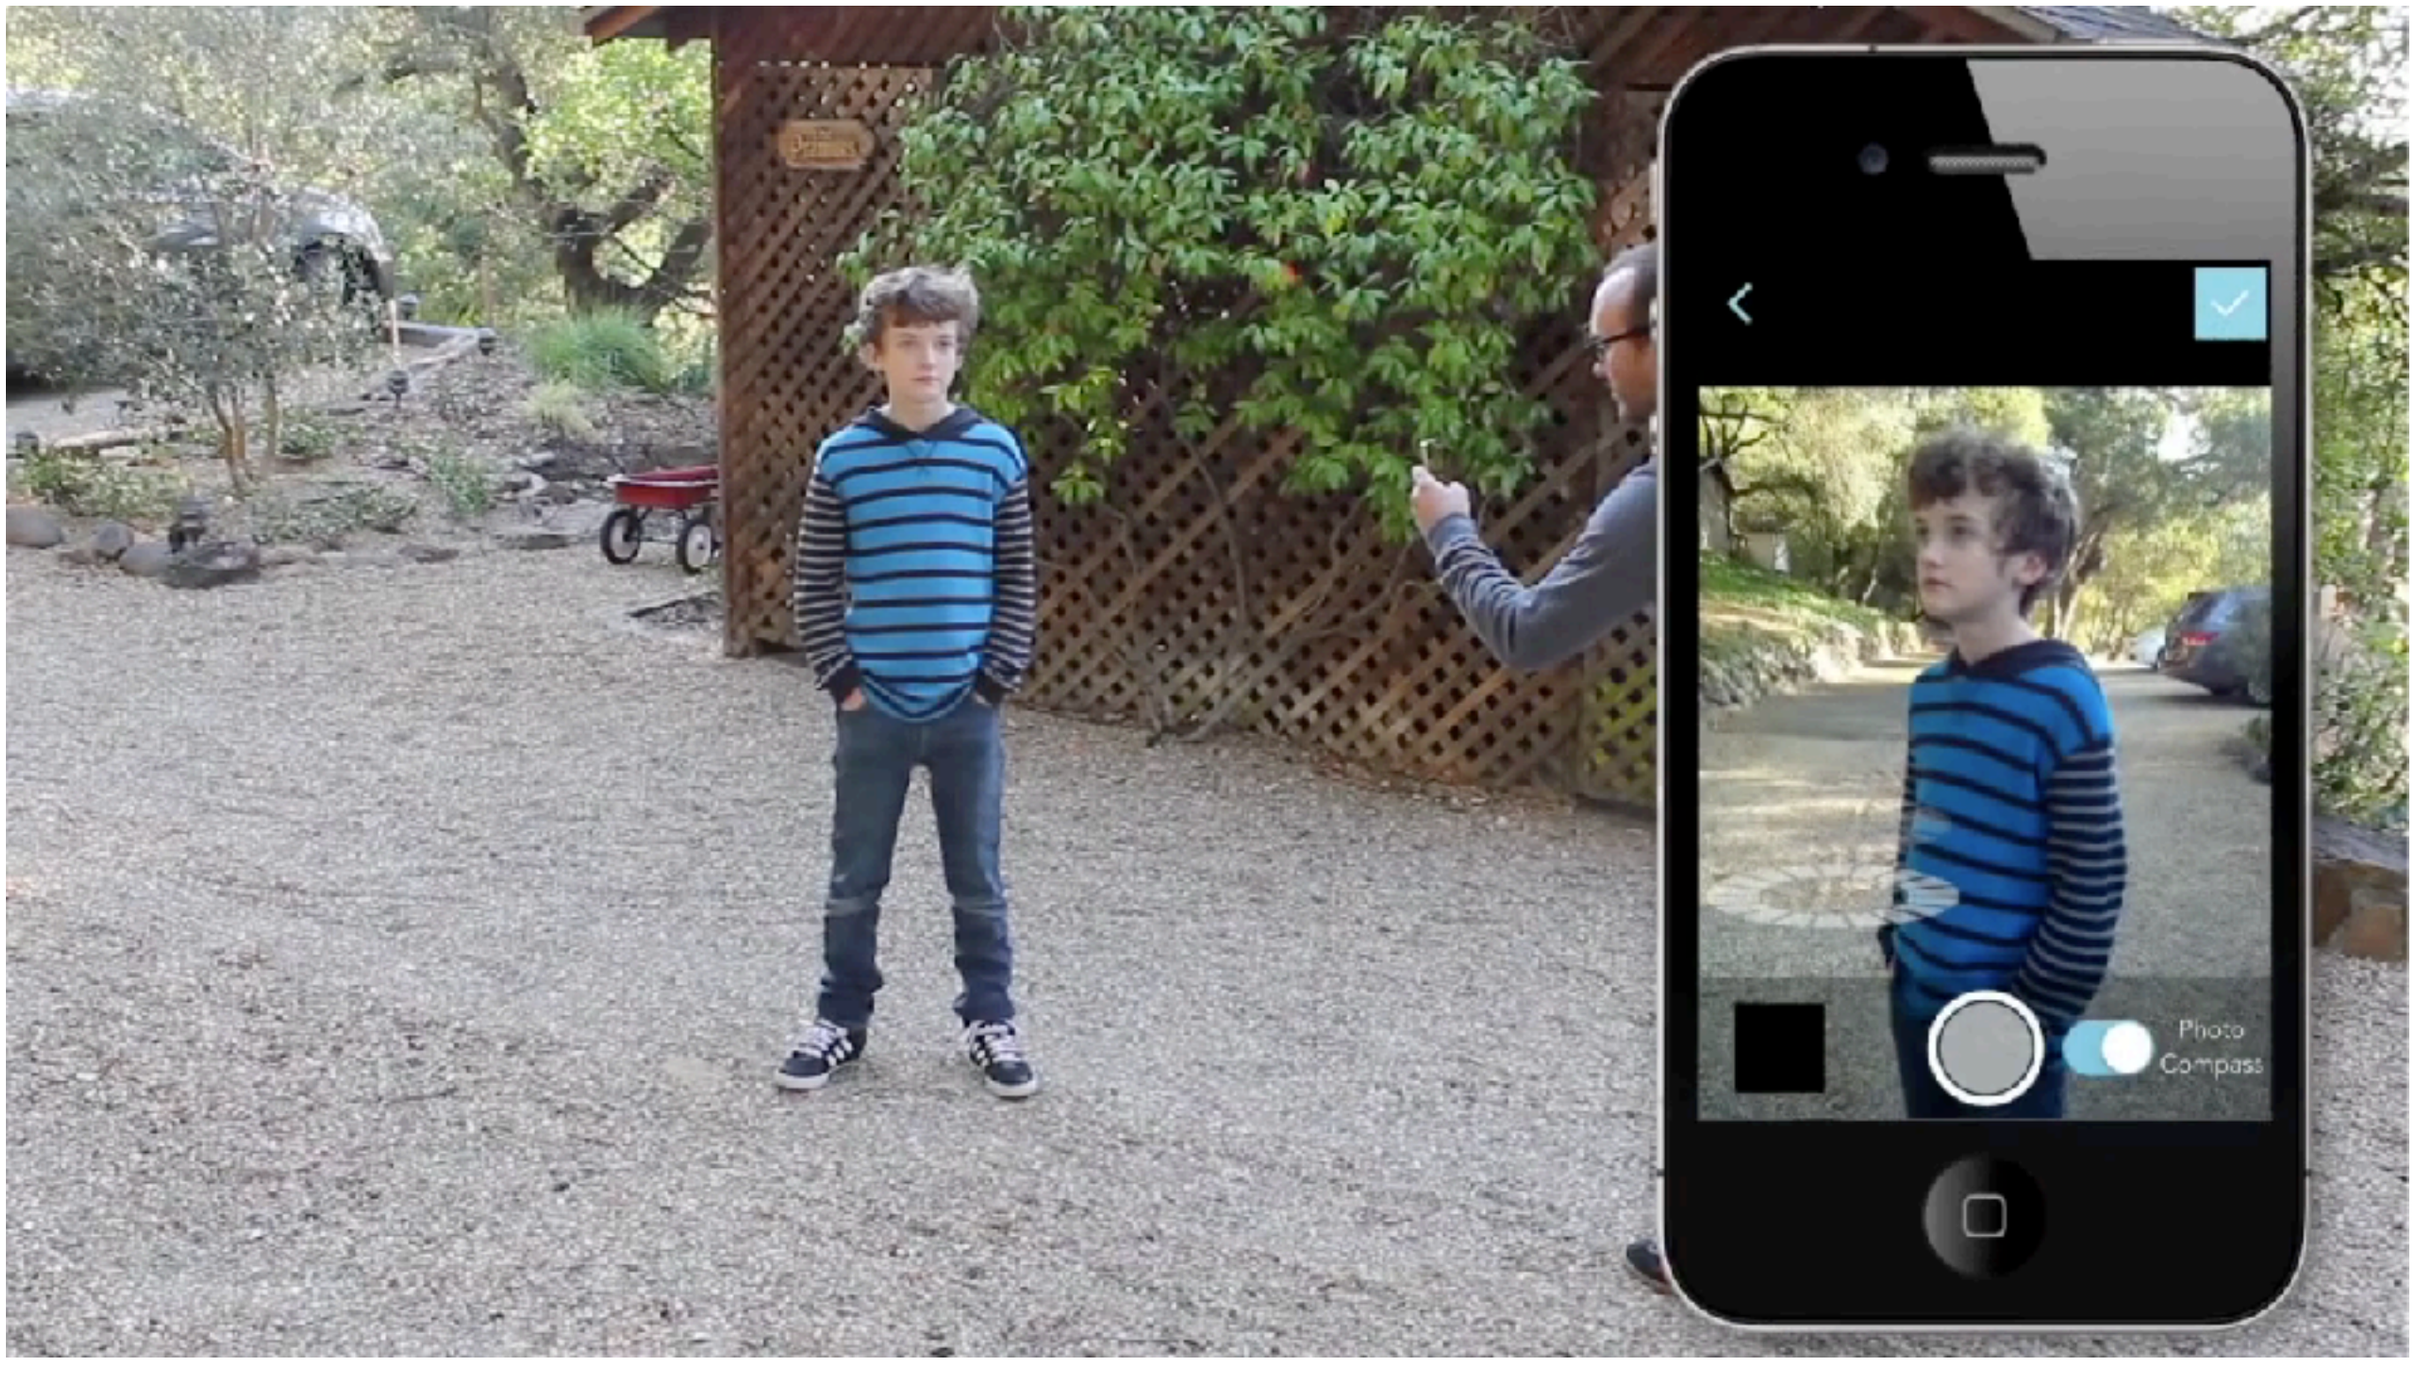

[Autodesk 123D Catch]

Body Pose Tracking

[Microsoft Xbox Kinect]

Body Pose Tracking

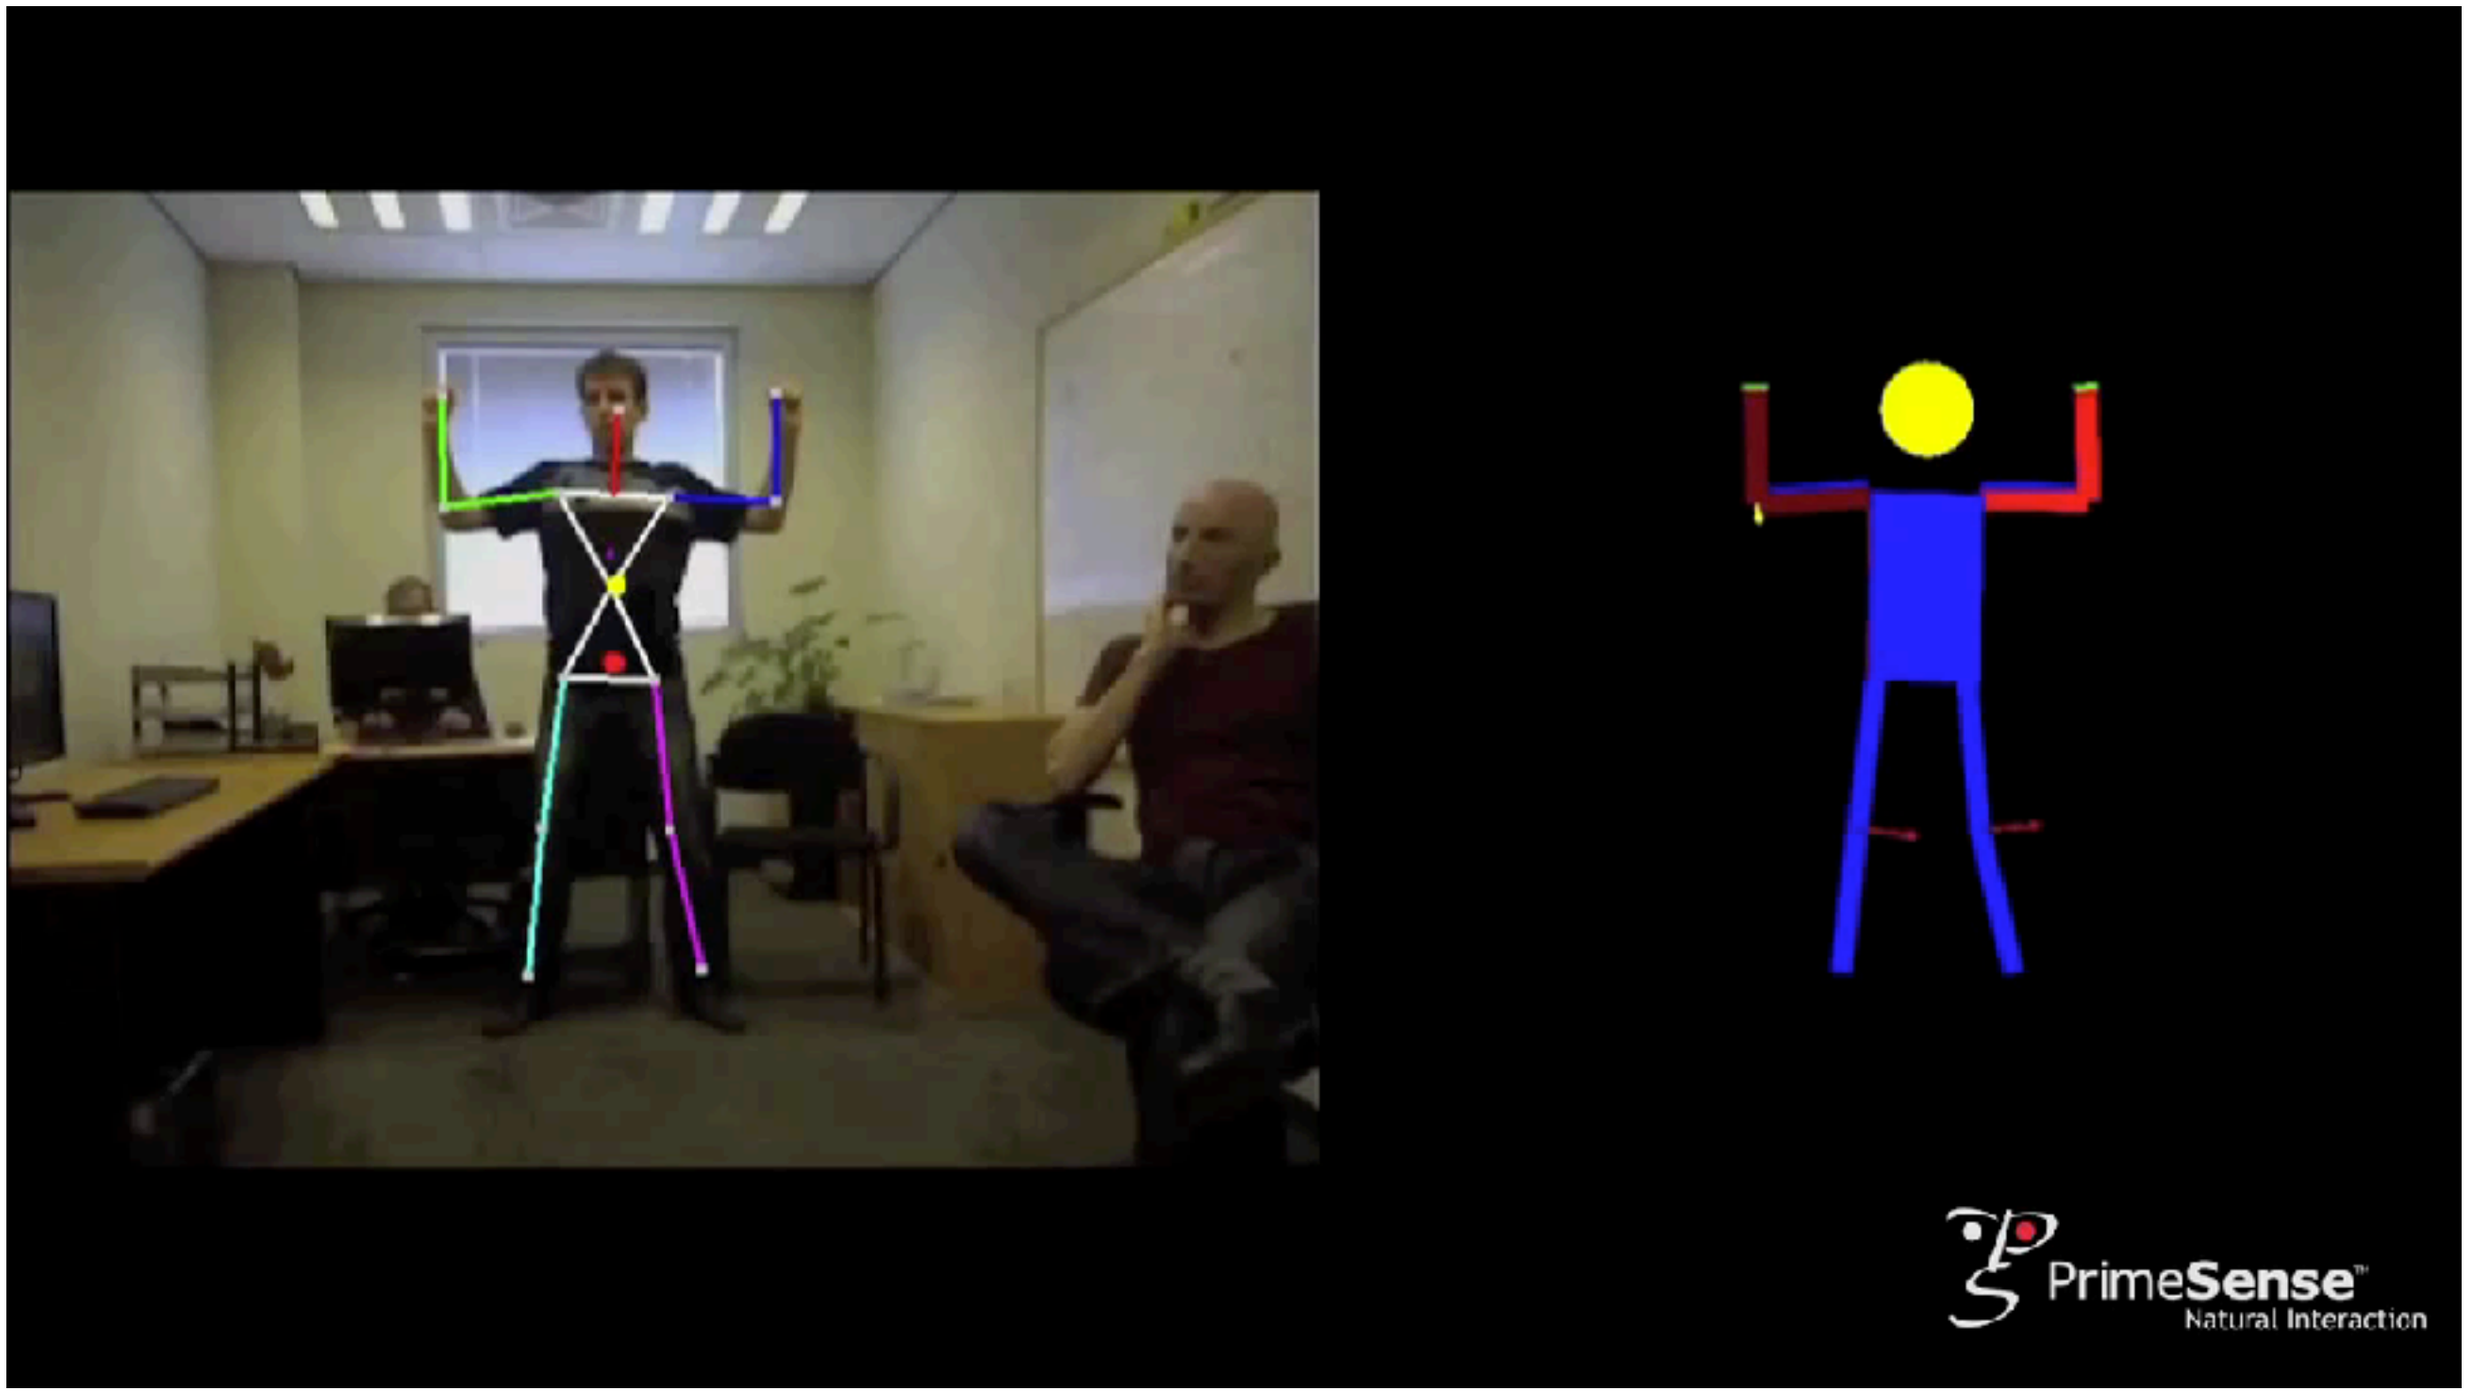

[PrimeSense]

Image Recognition and Search

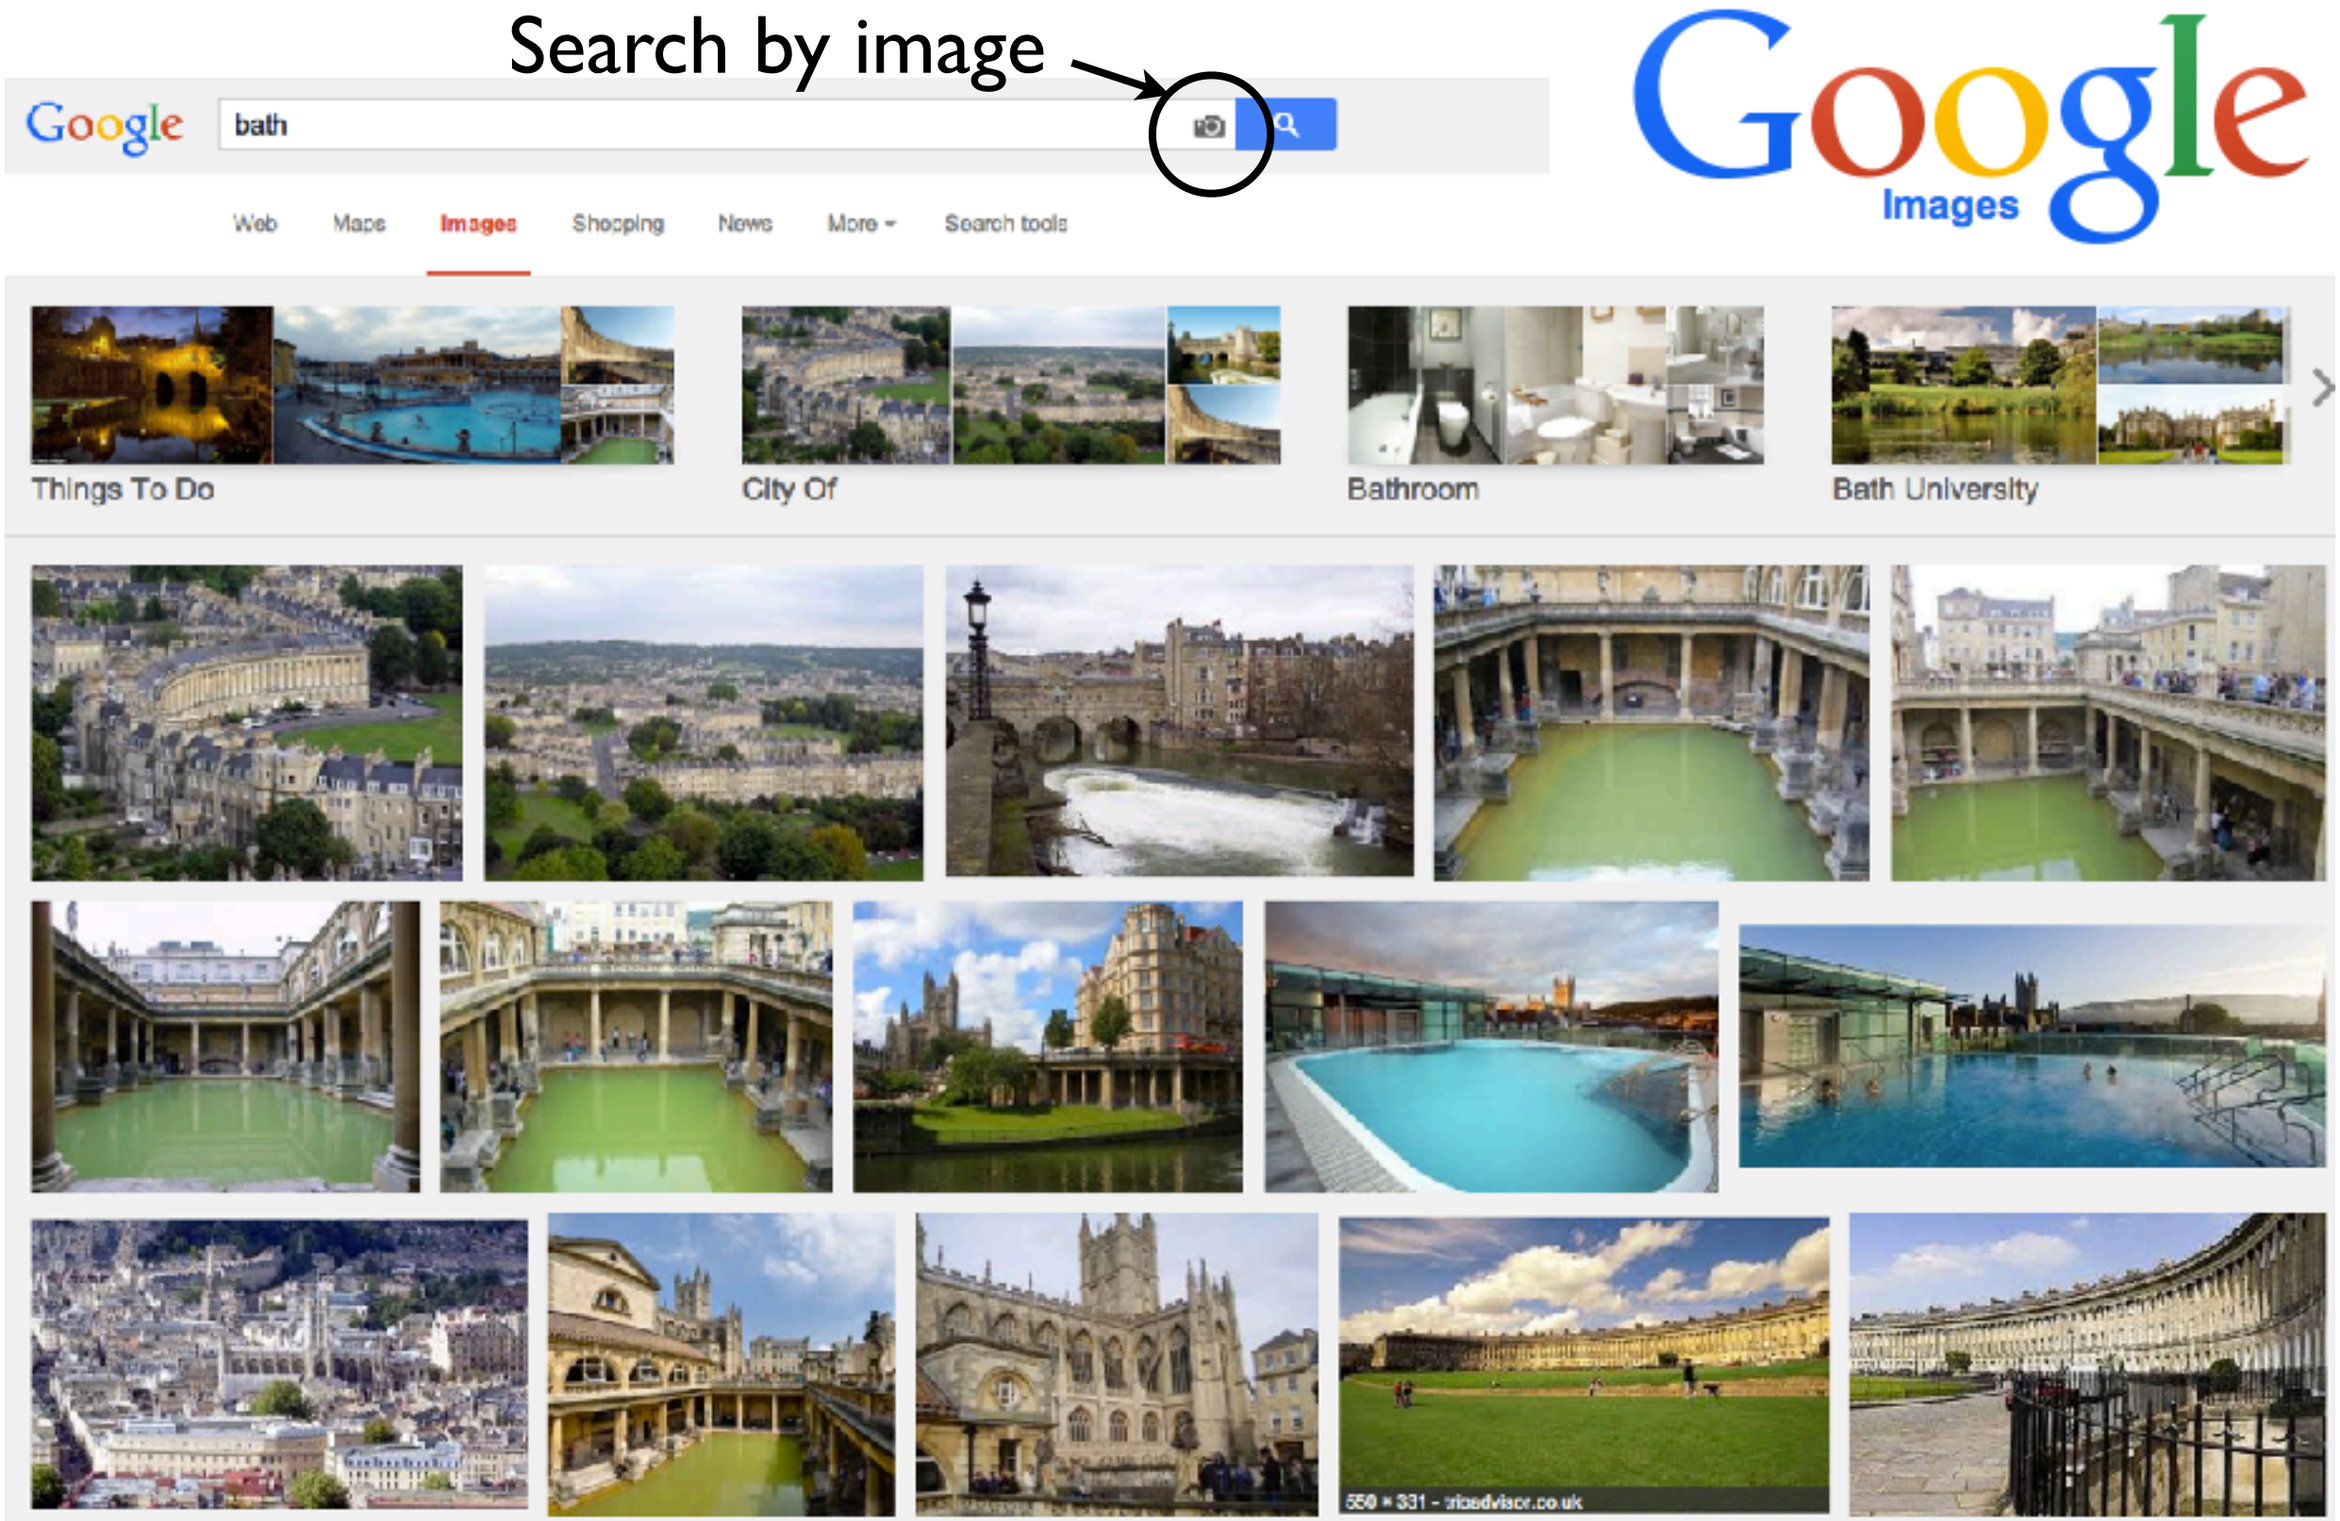

Self Driving Cars

[Google]

Flying Vehicles

www.skydio.com

Live Slides web content

To view

Download the add-in.

liveslides.com/download

Start the presentation.

AR/VR

[Microsoft Hololens]

Mobile Apps

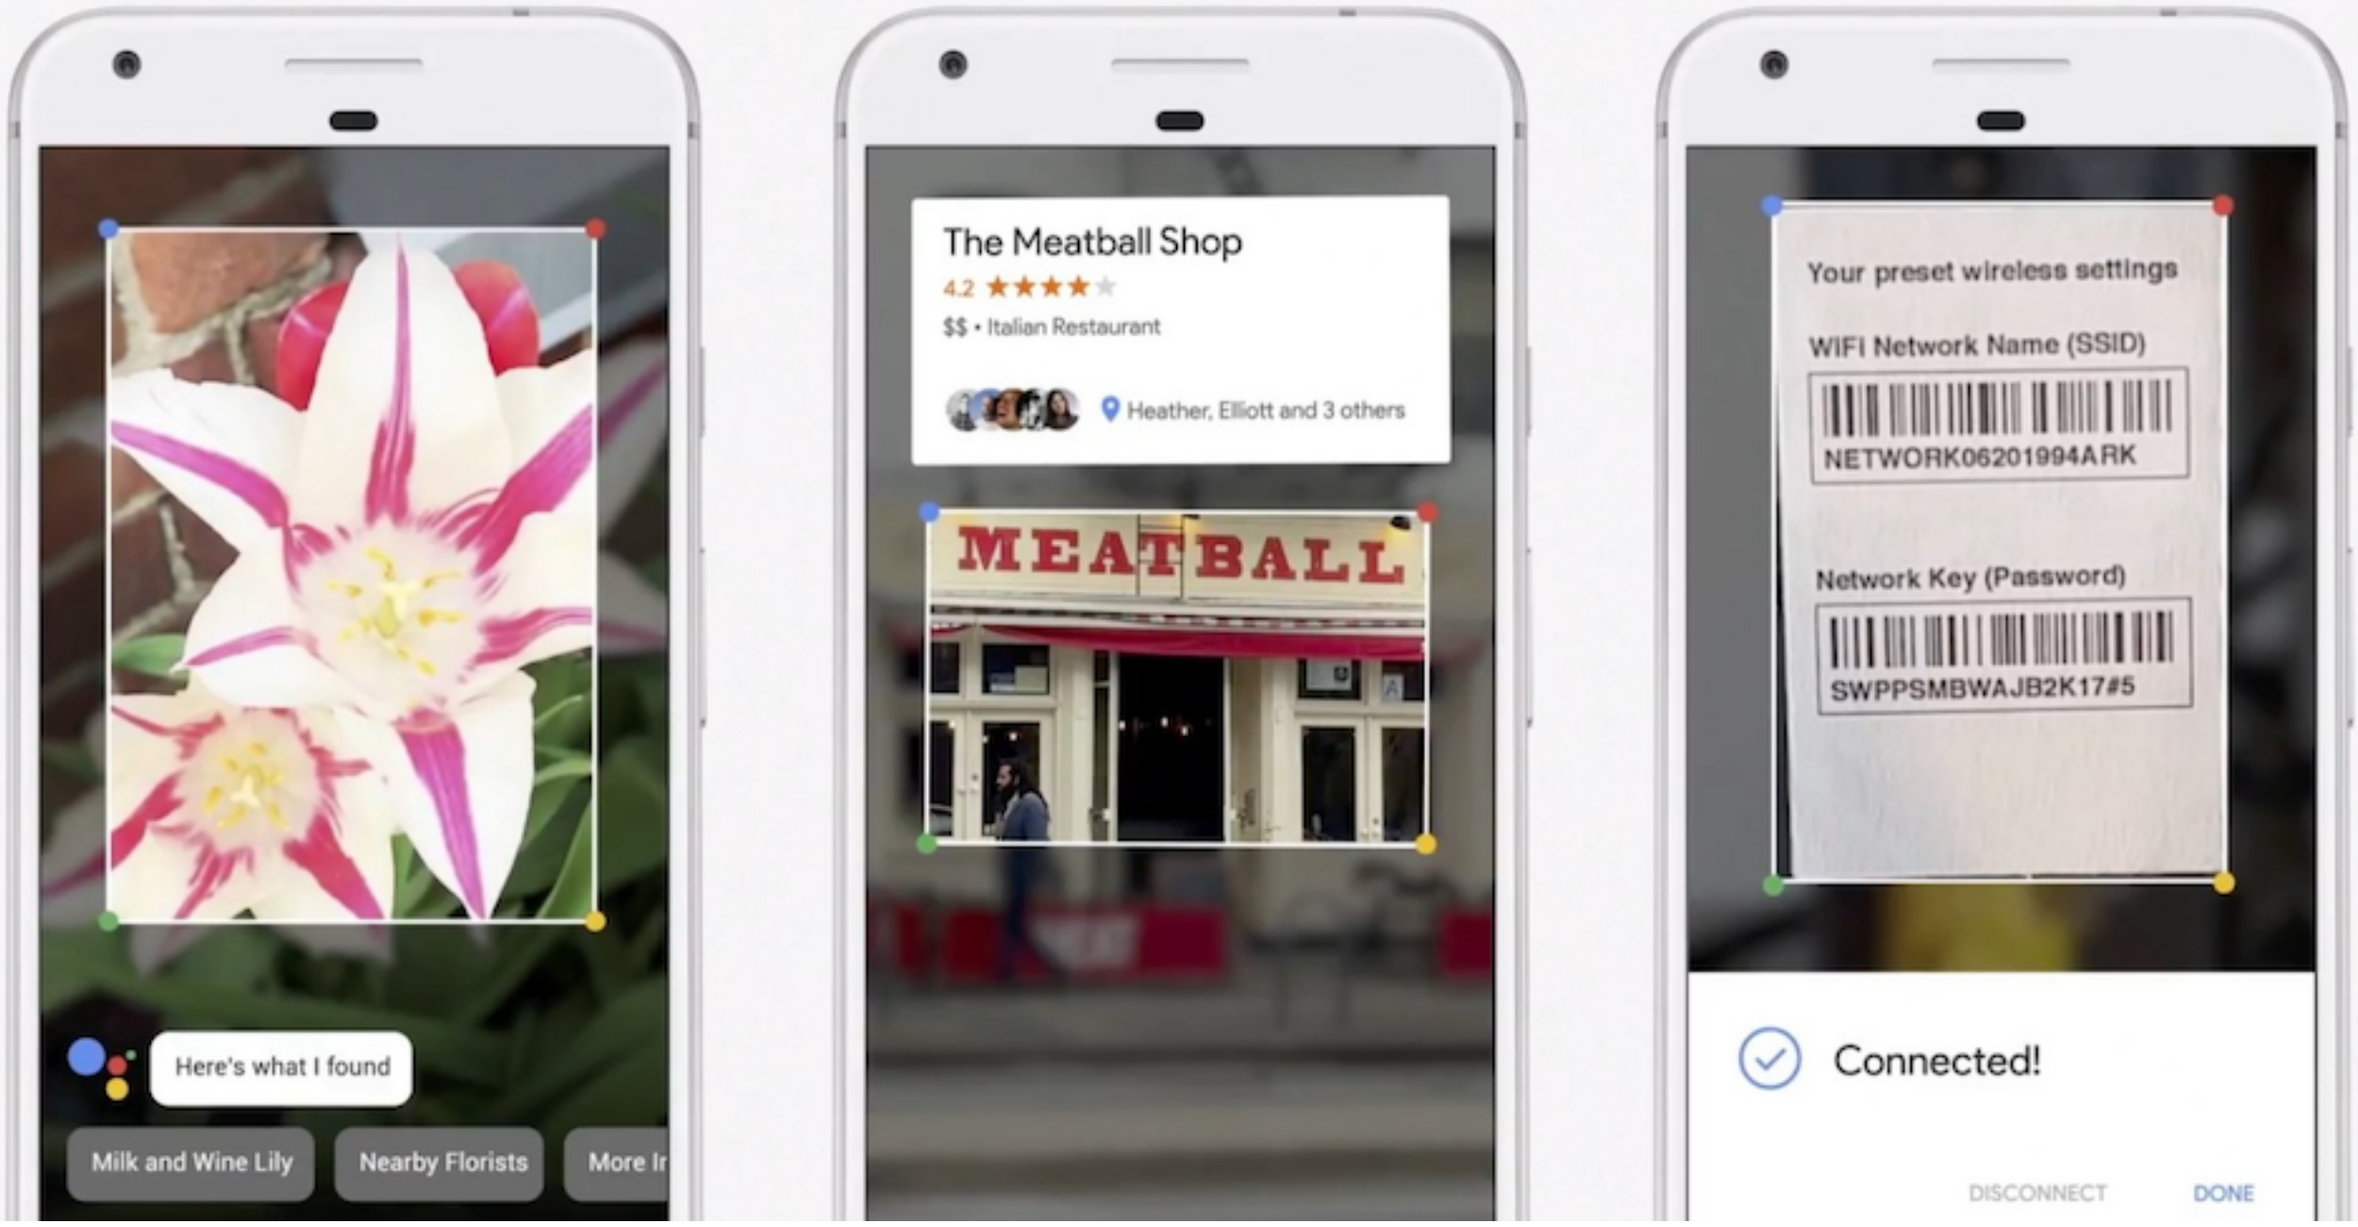

[Google Lens]

Art

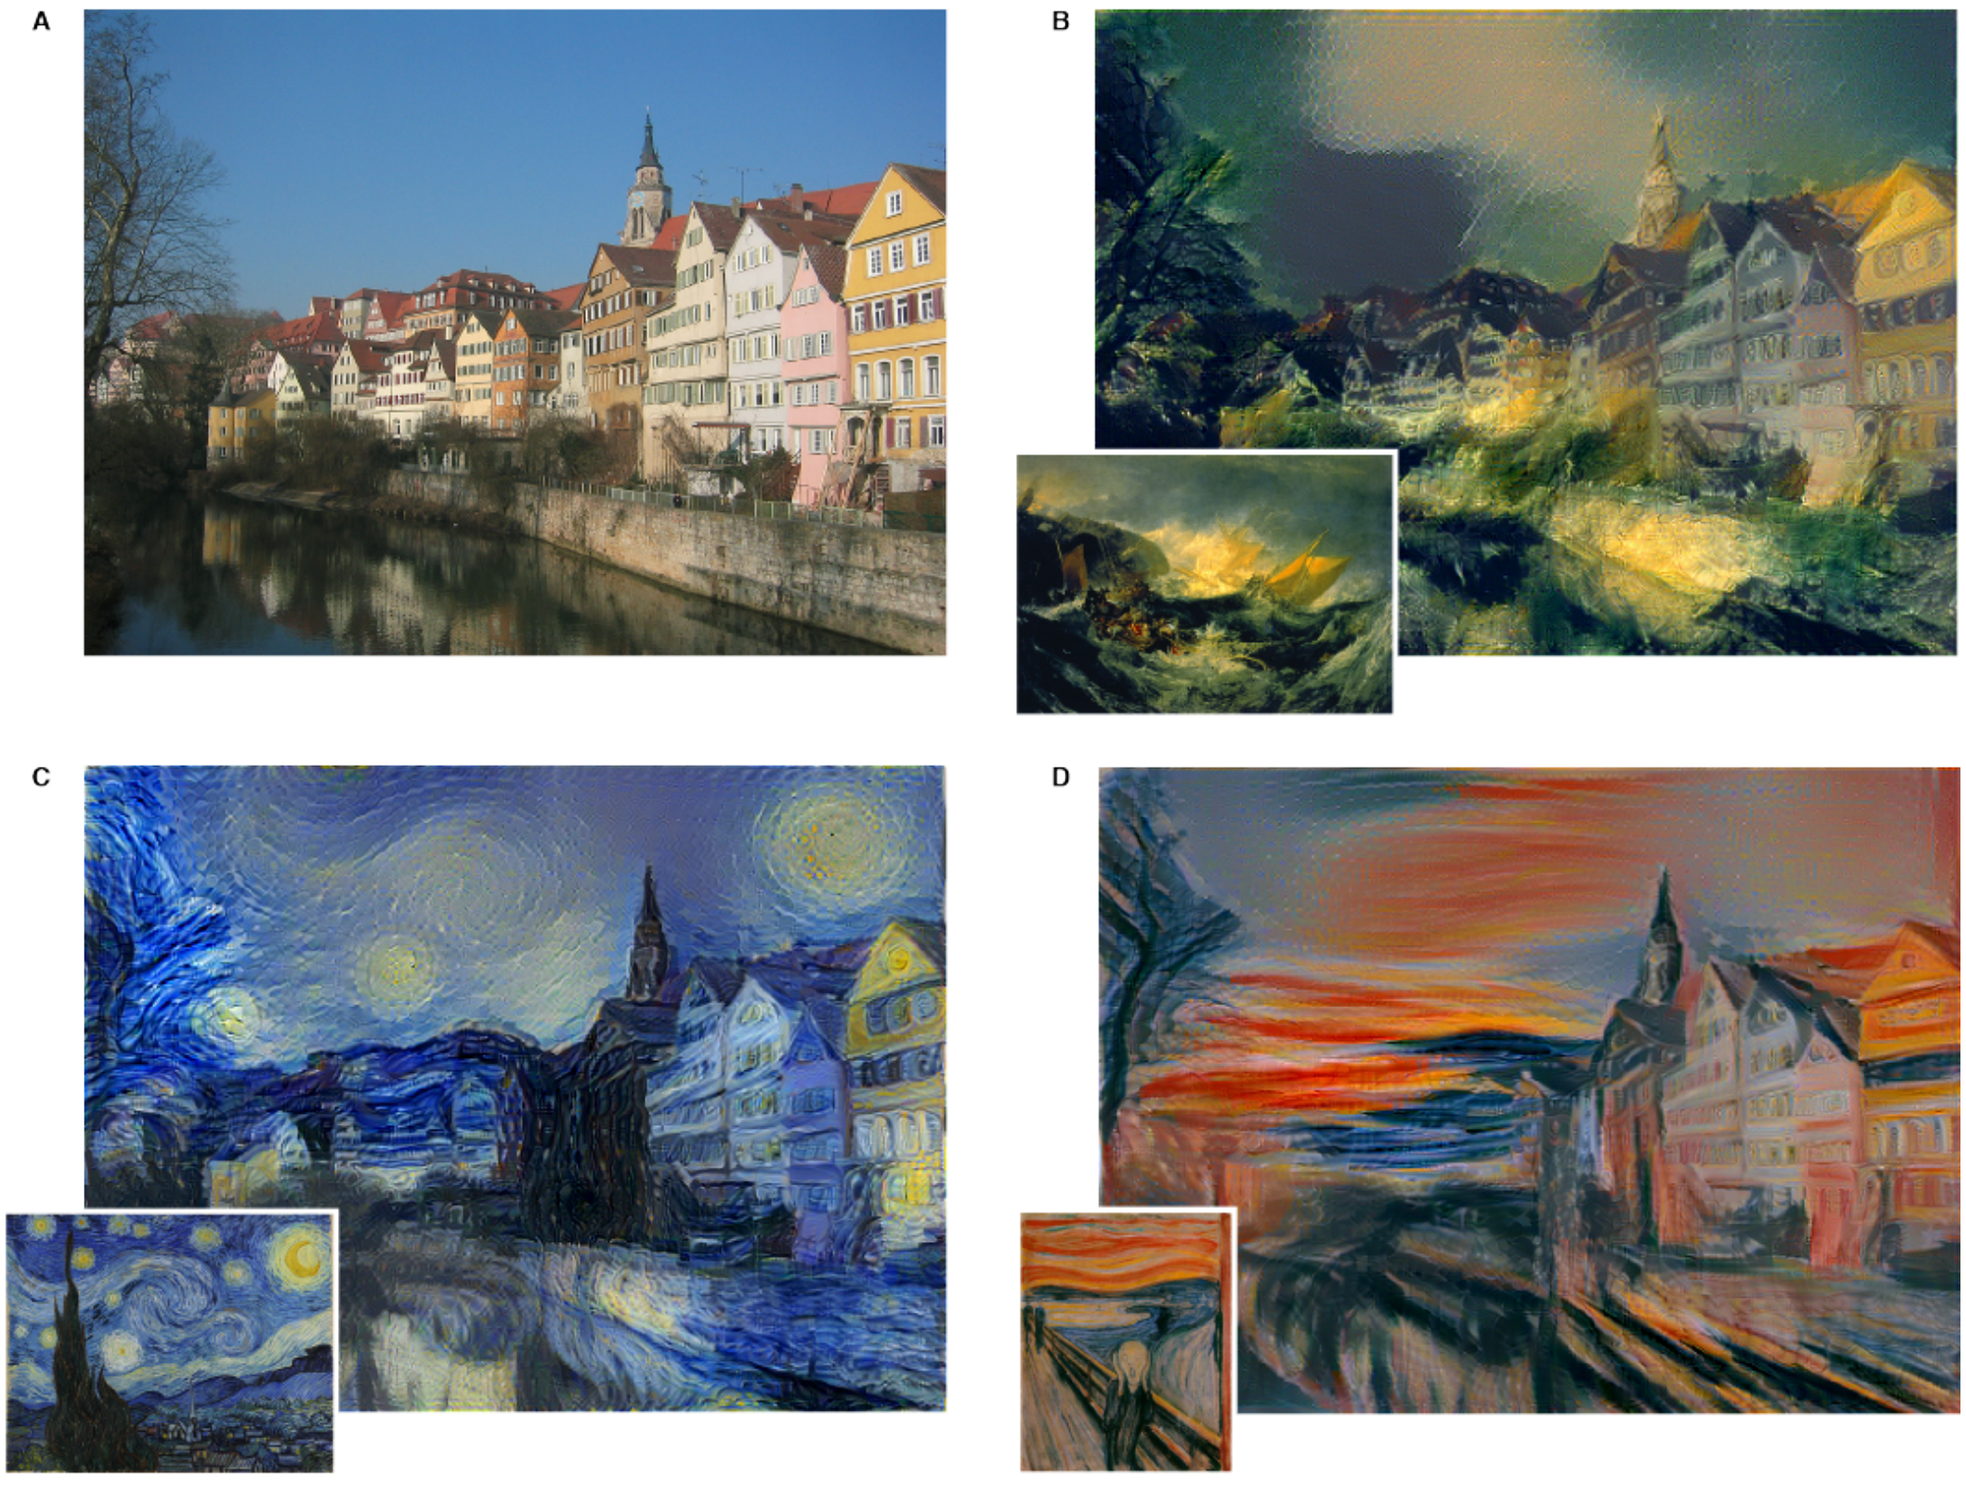

[Gatys, Ecker, Bethge 2015]

Applications of Computer Vision

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

Definitions of Computer Vision #1

“Inverse Computer Graphics”

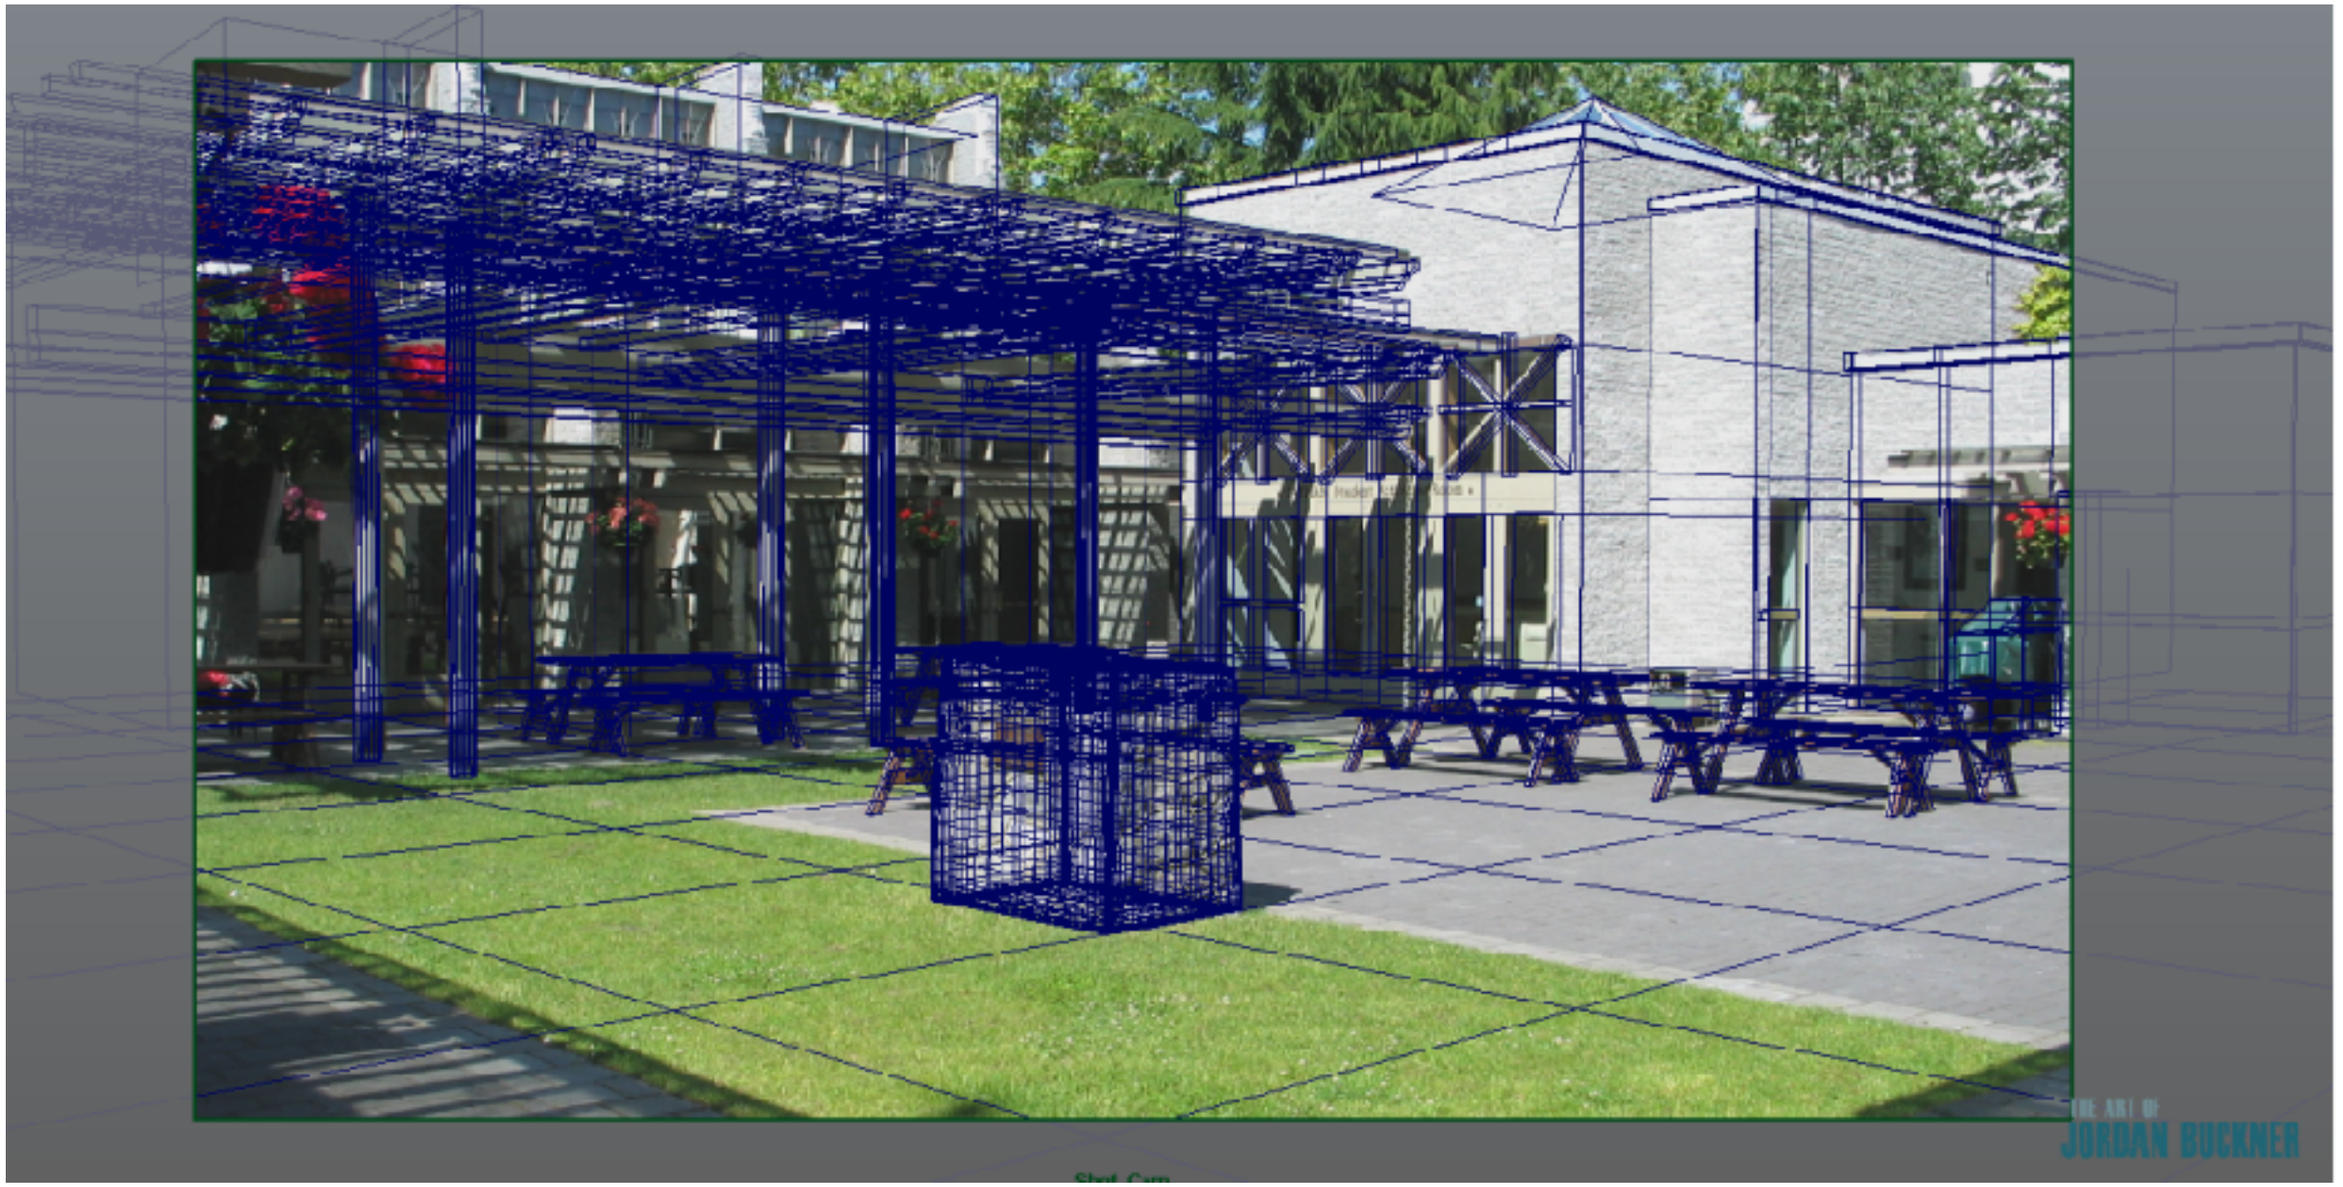

[J. Buckner]

Definitions of Computer Vision #1

"Inverse Computer Graphics"

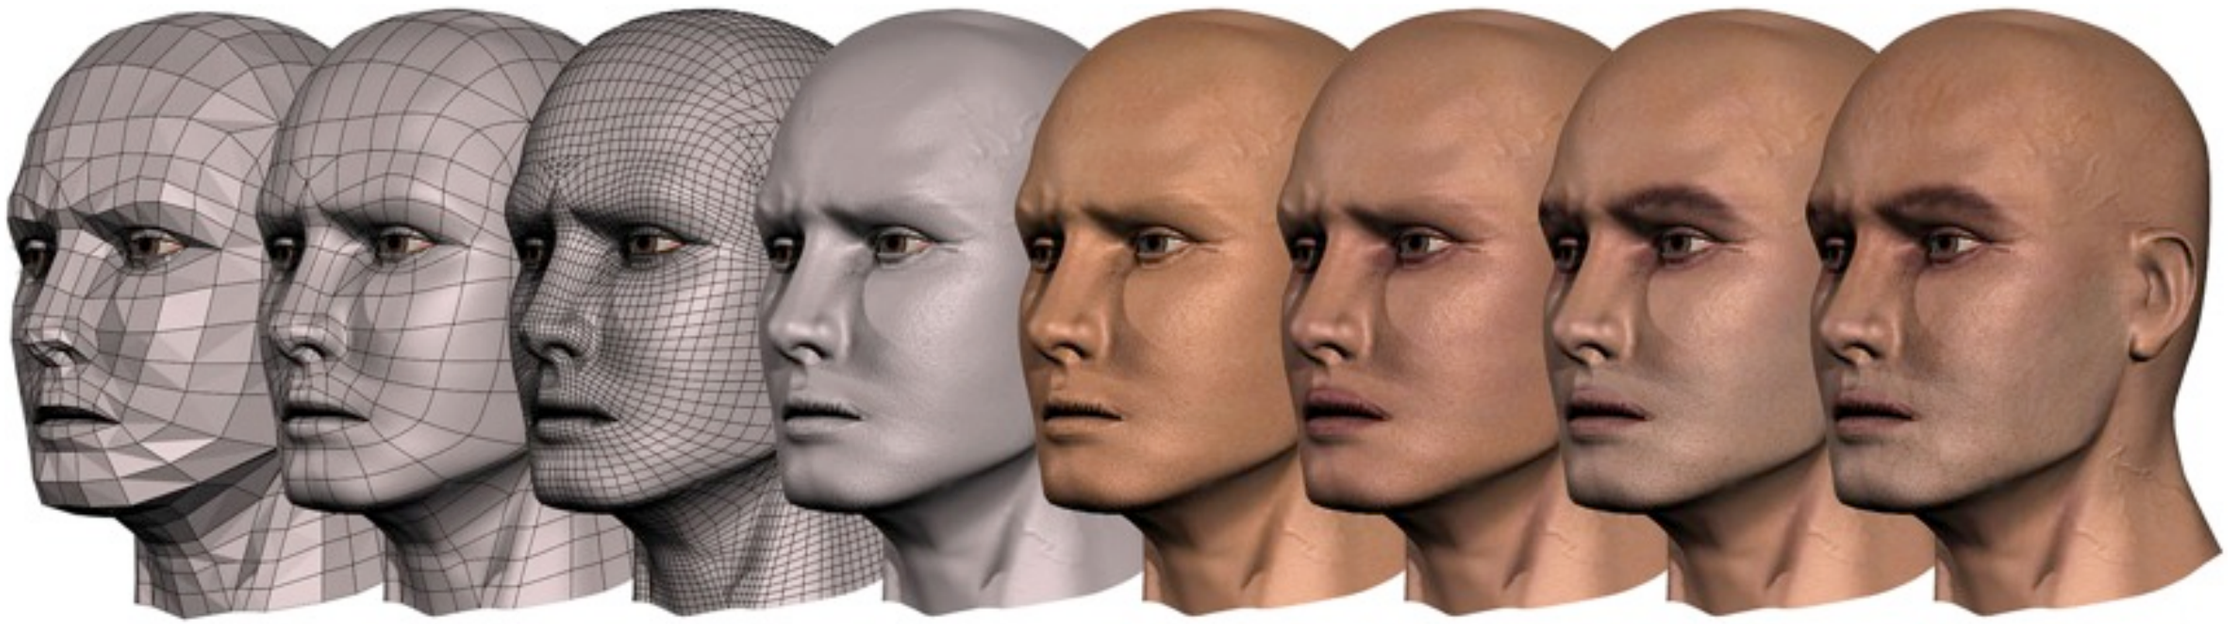

Graphics

Vision

Photometric Capture

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

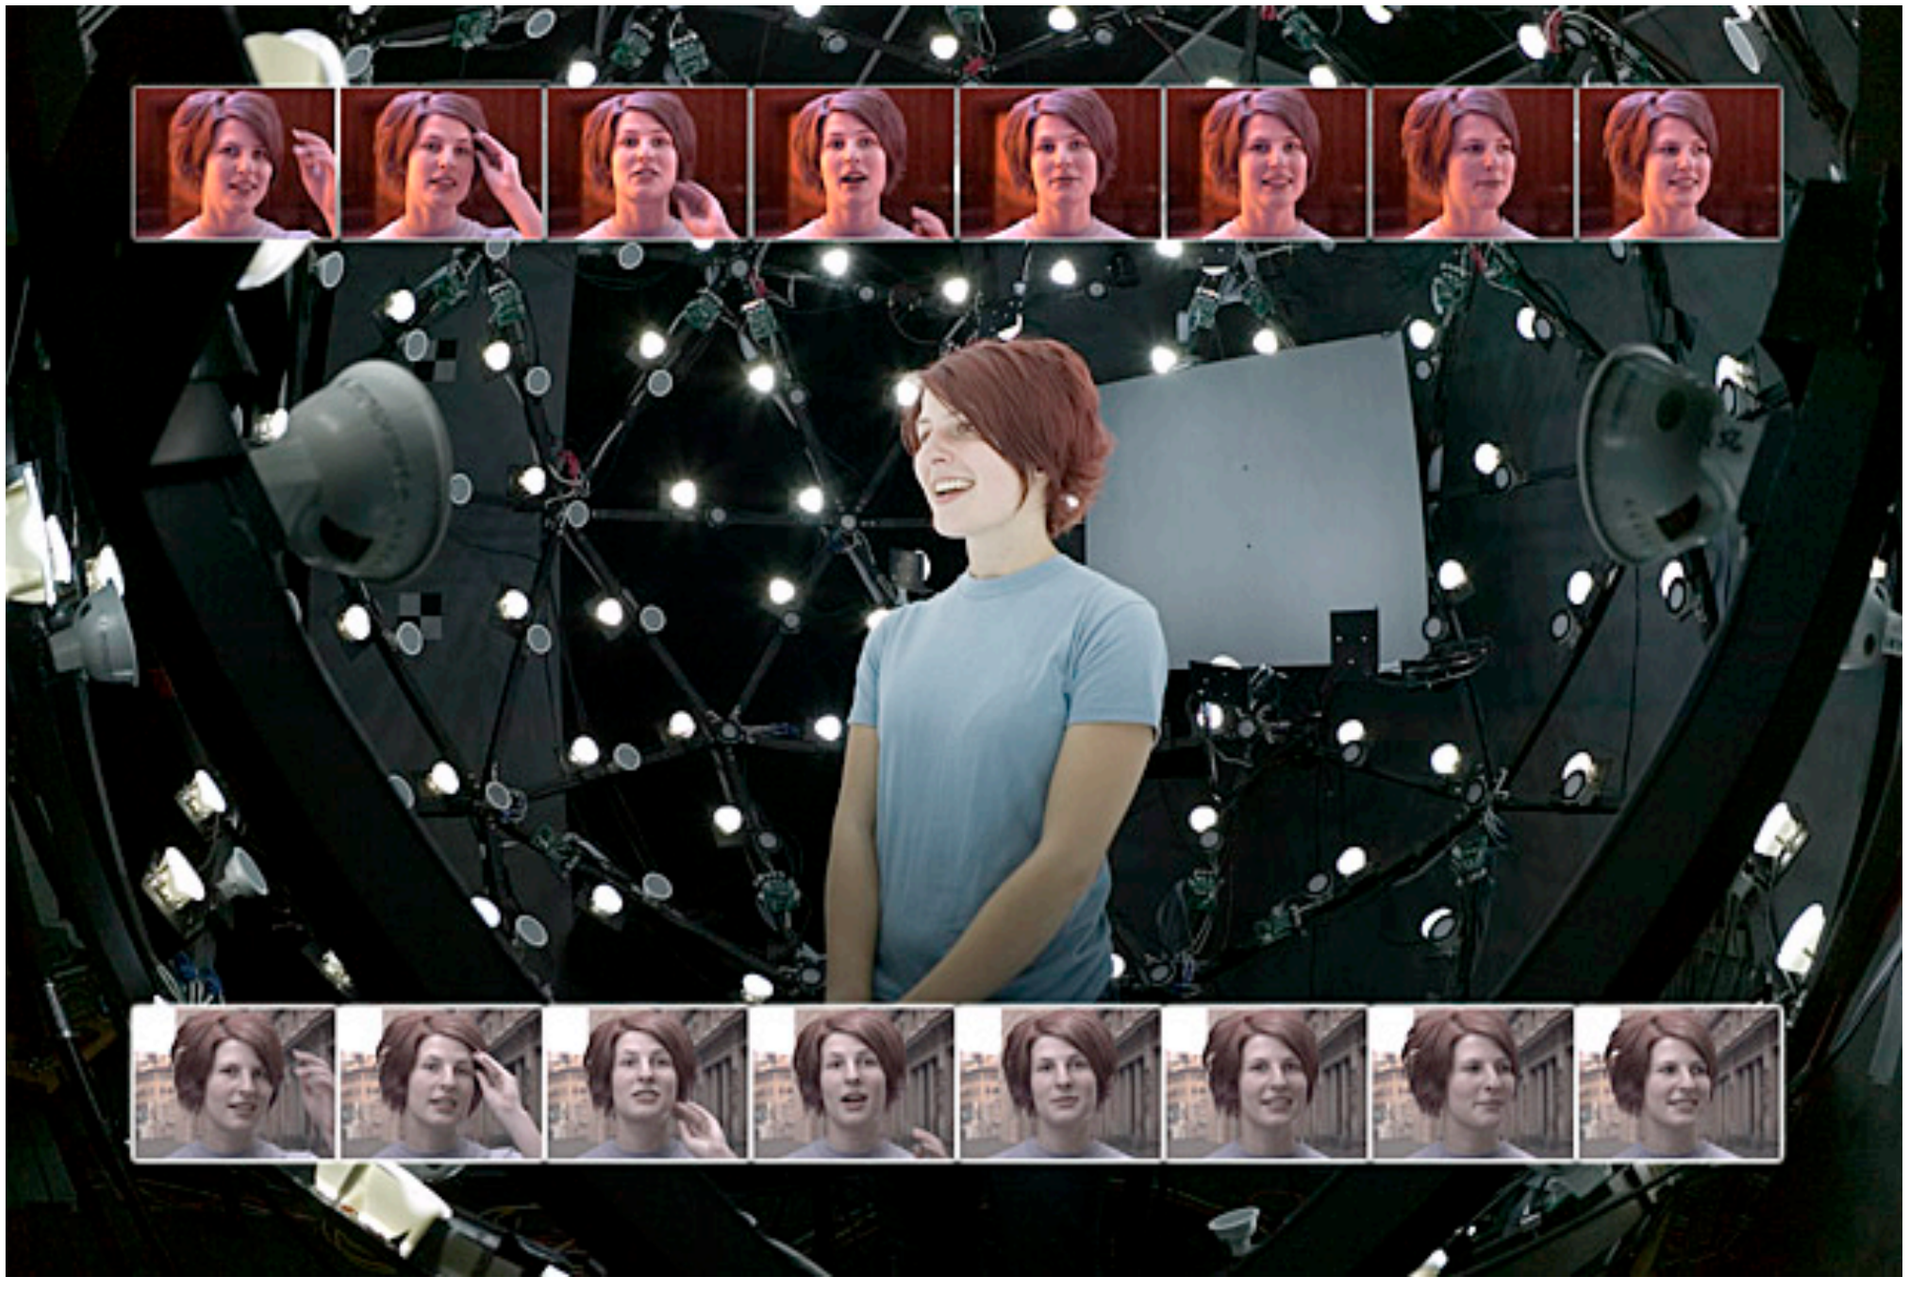

[Paul Debevec, USC]

Definitions of Computer Vision #2

“Replicate Human Vision”

=

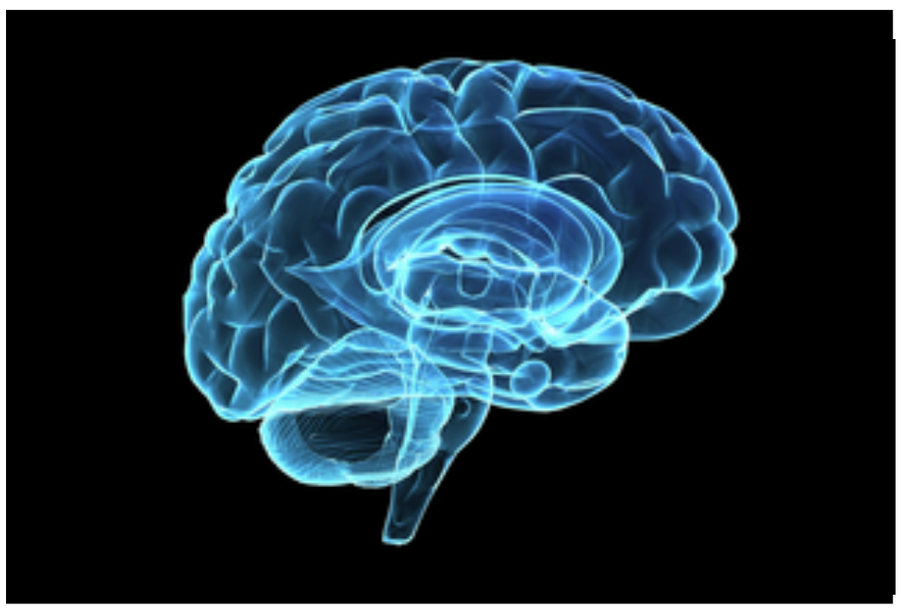

=

Definitions of Computer Vision #2

“Replicate Human Vision”

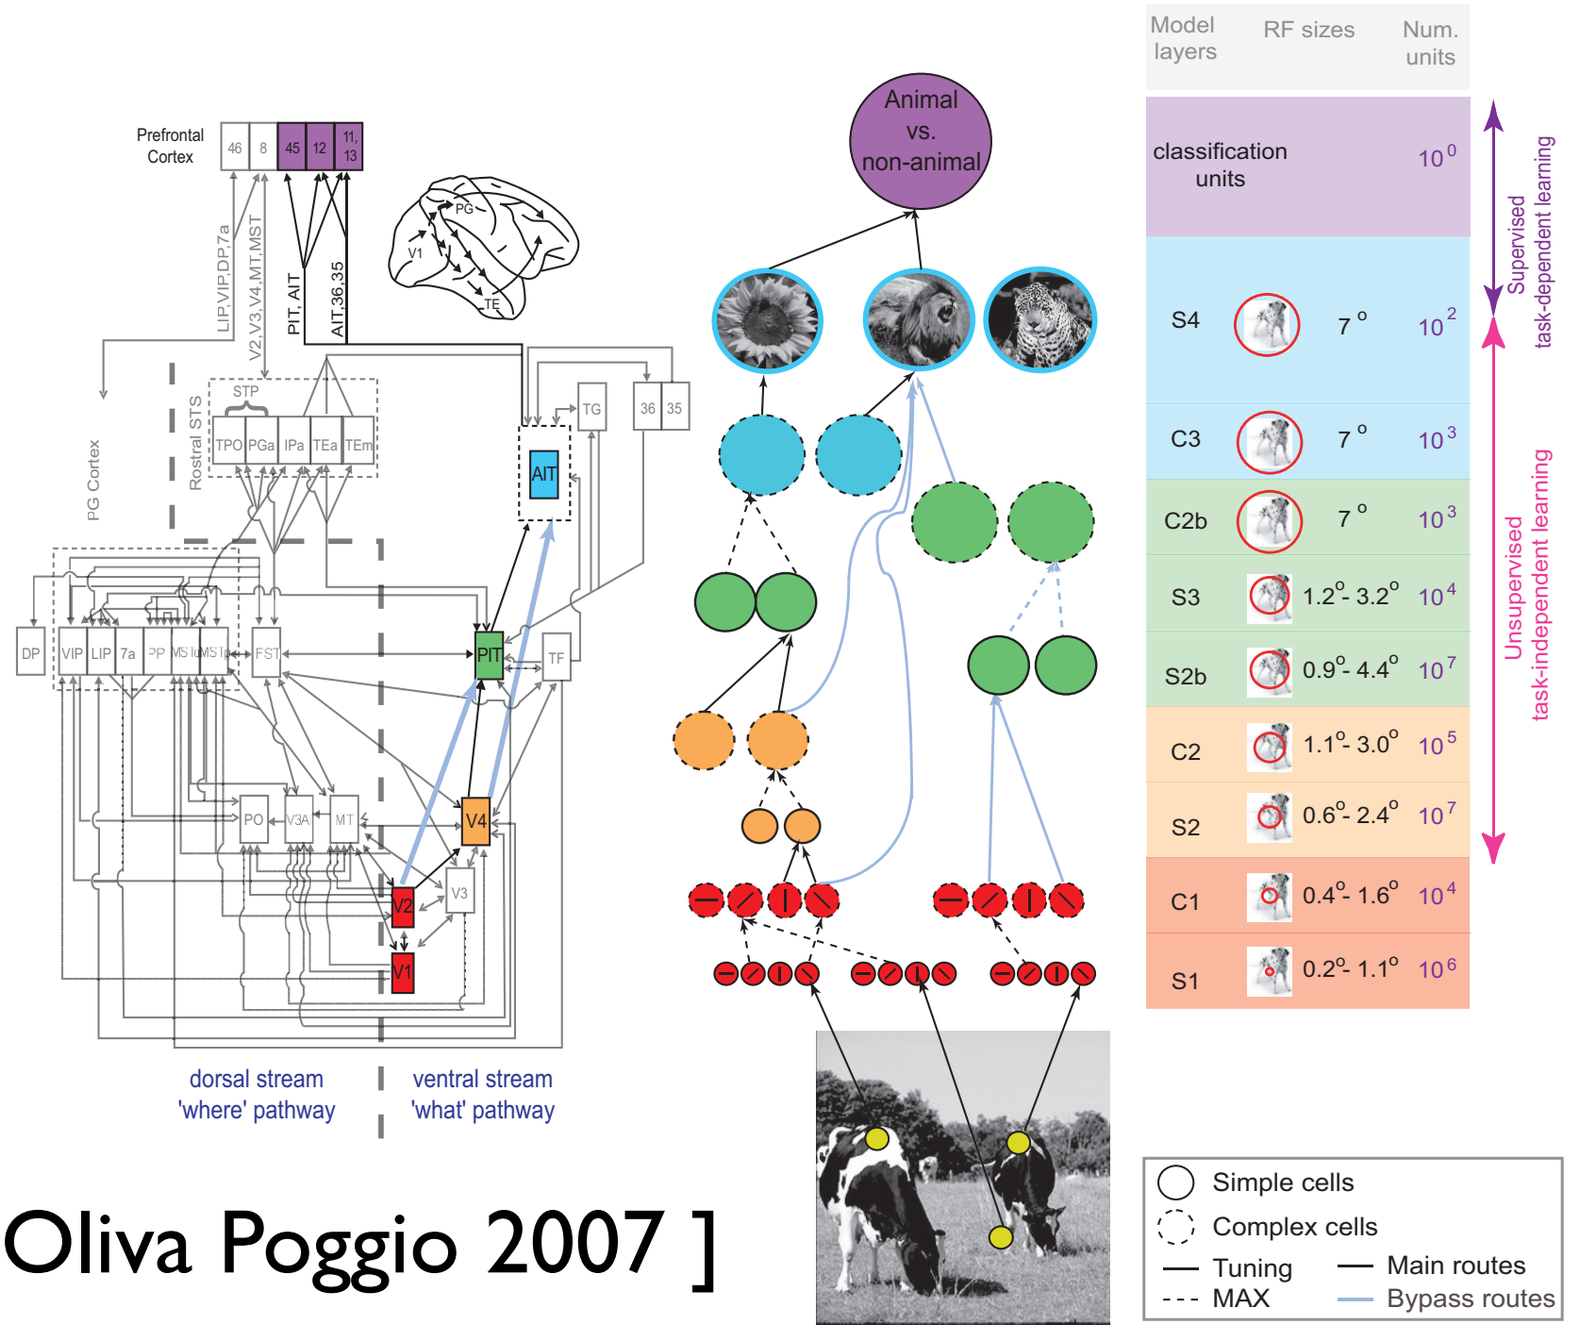

[Serre Oliva Poggio 2007]

ImageNet

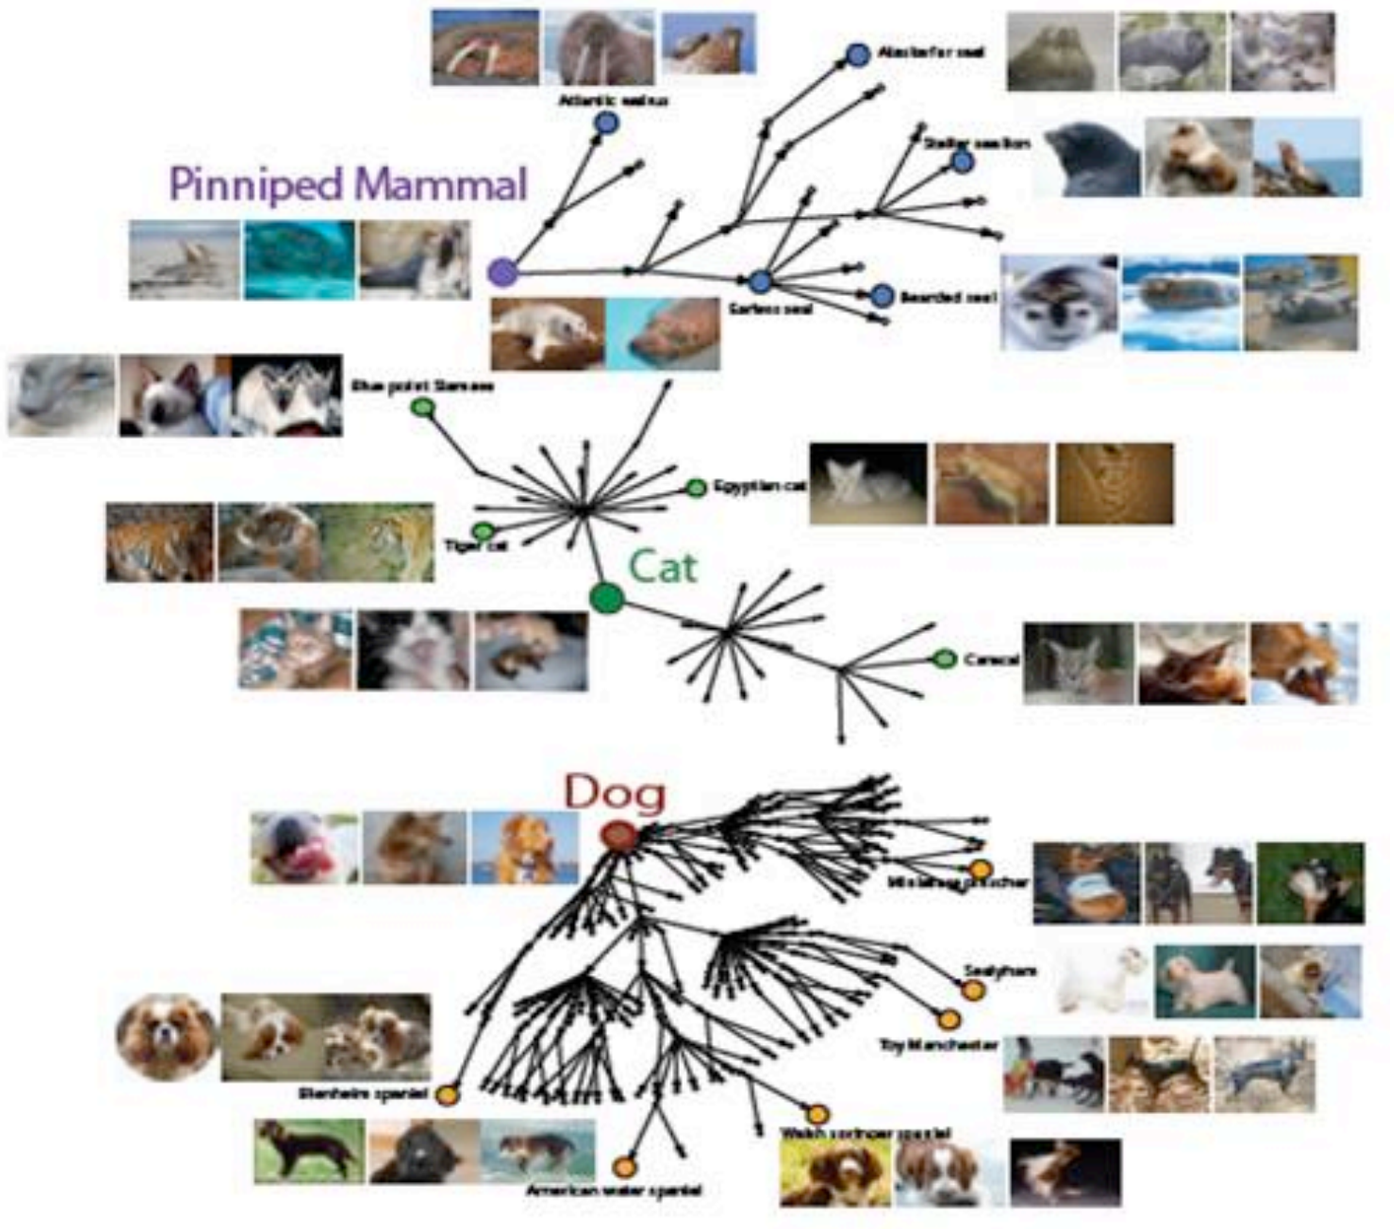

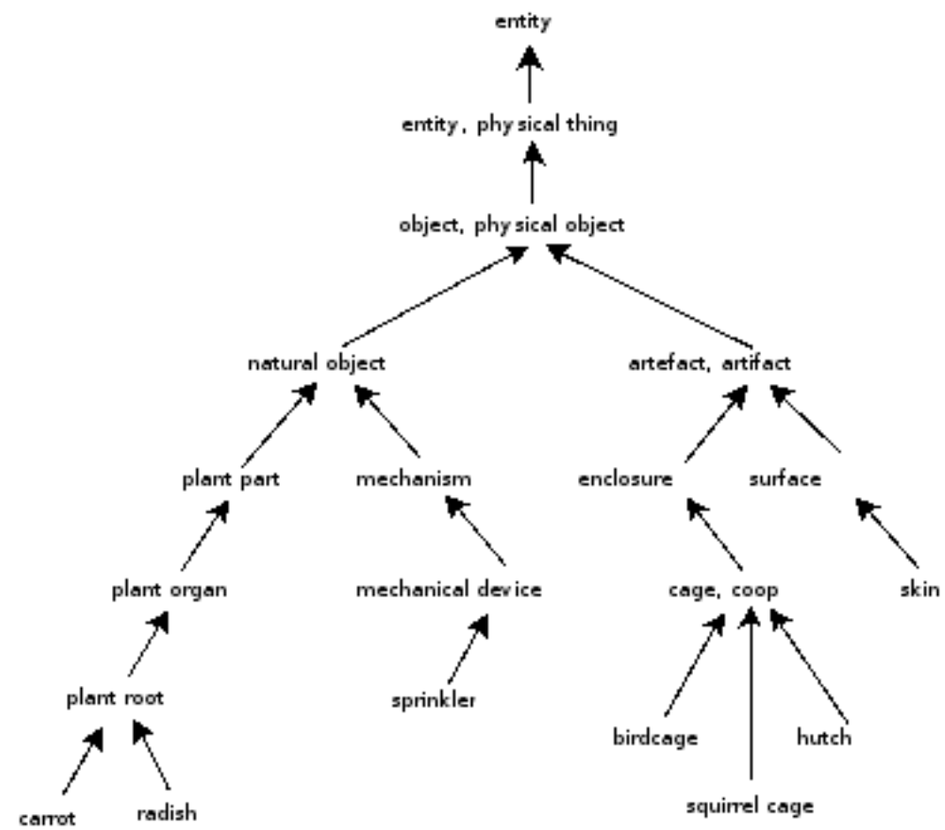

Figure 1. "is a" relation example

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)

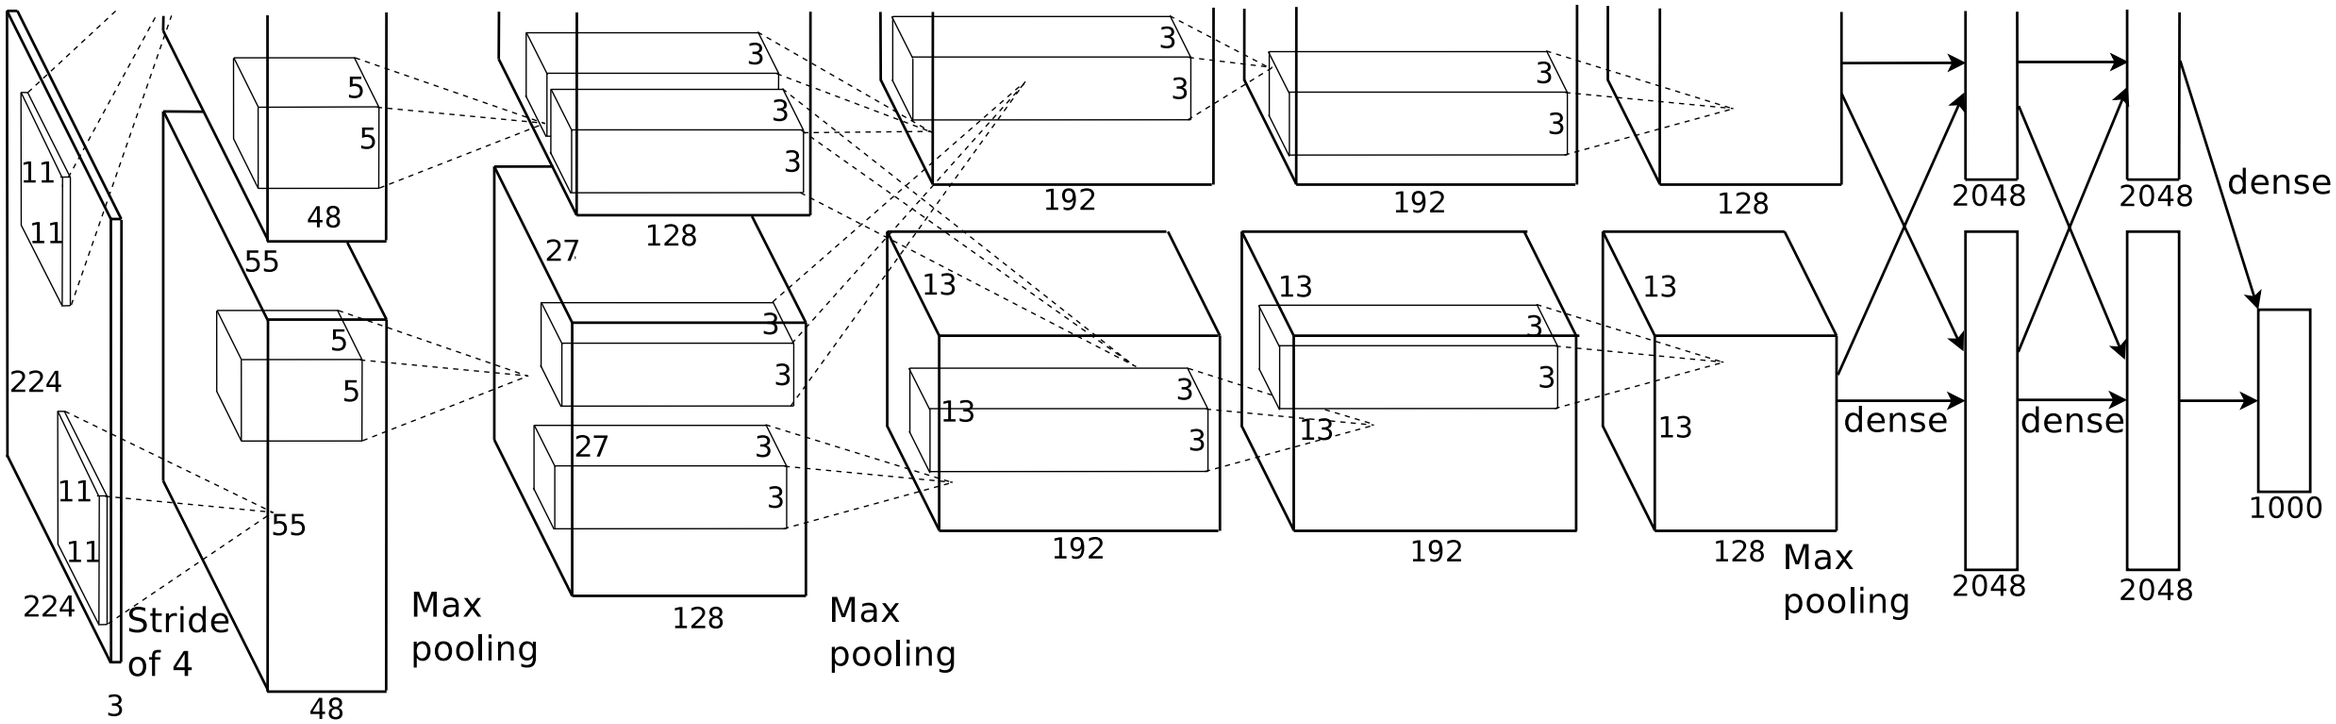

[Krizhevsky, Sutskever, Hinton 2012]

Definitions of Computer Vision #3

“Image/Video Understanding”

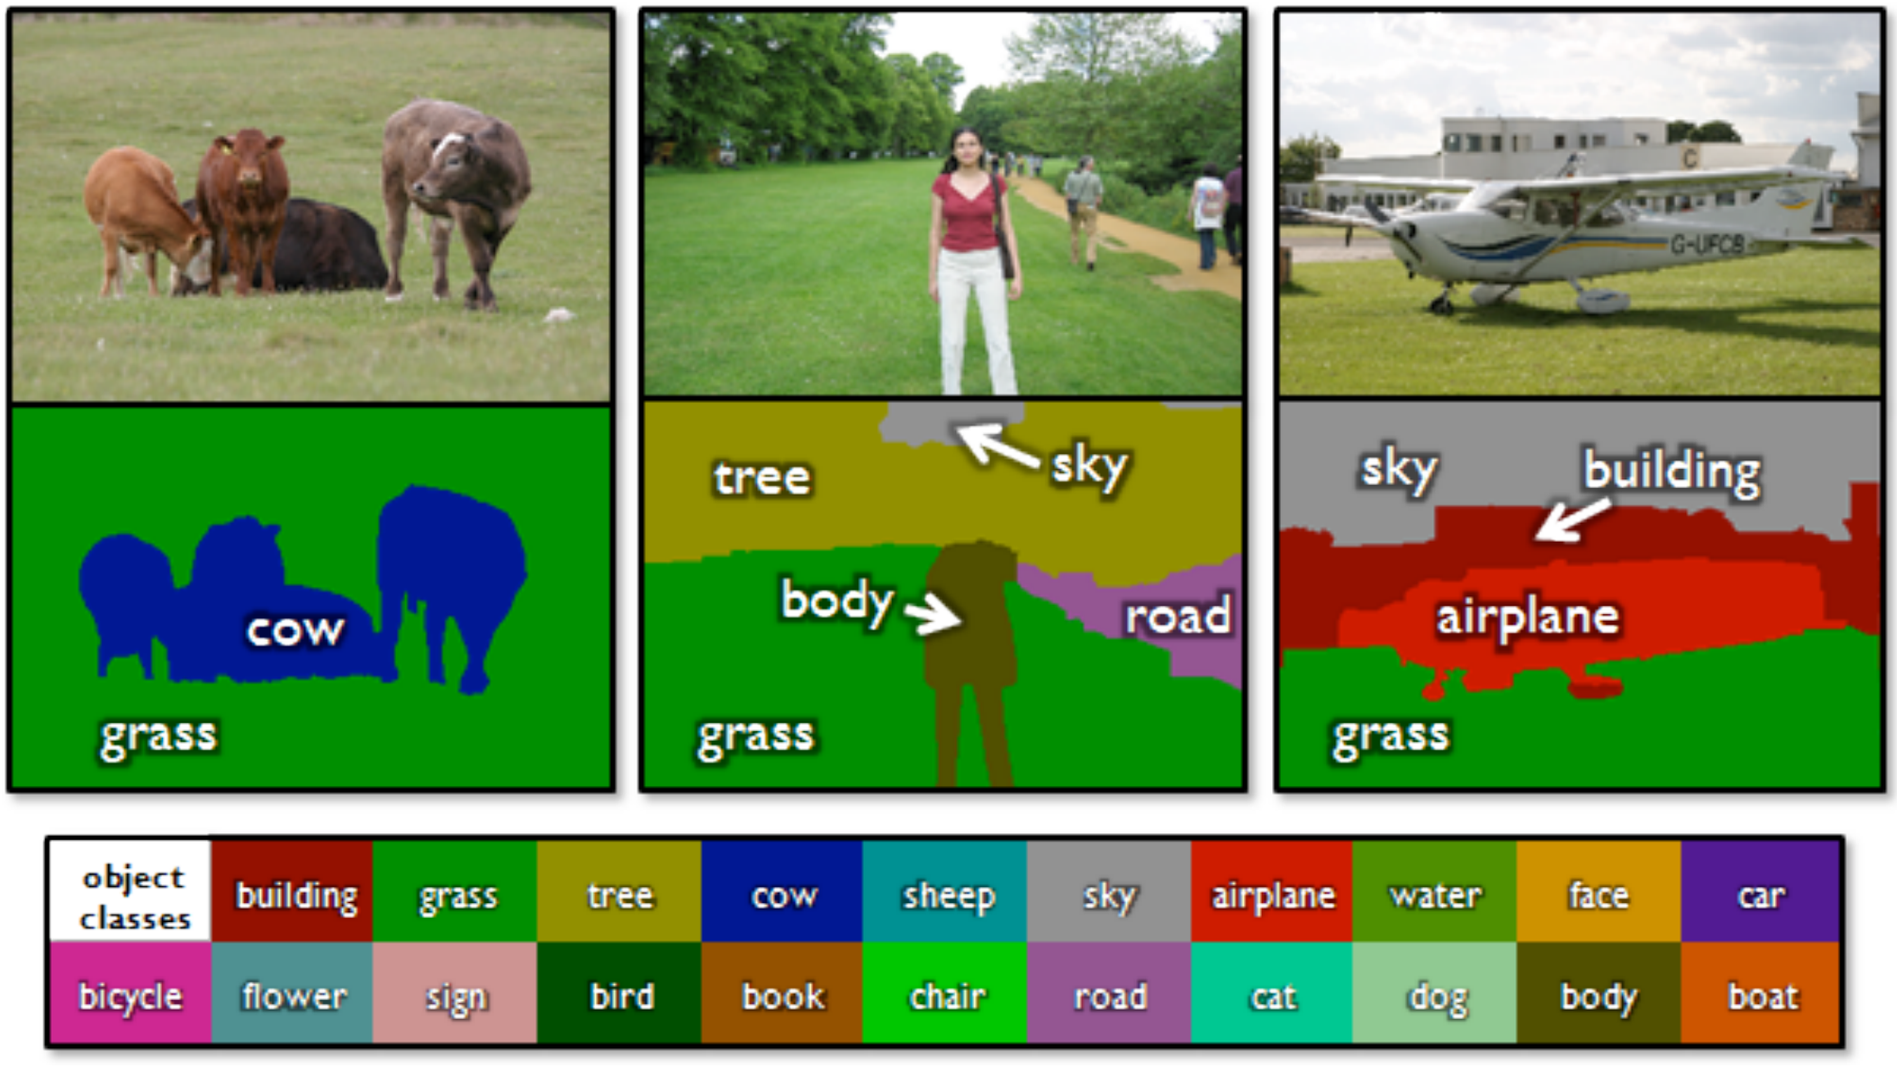

[Shotton Winn Rother Criminisi 2006]

Definitions of Computer Vision #3

"Image/Video Understanding"

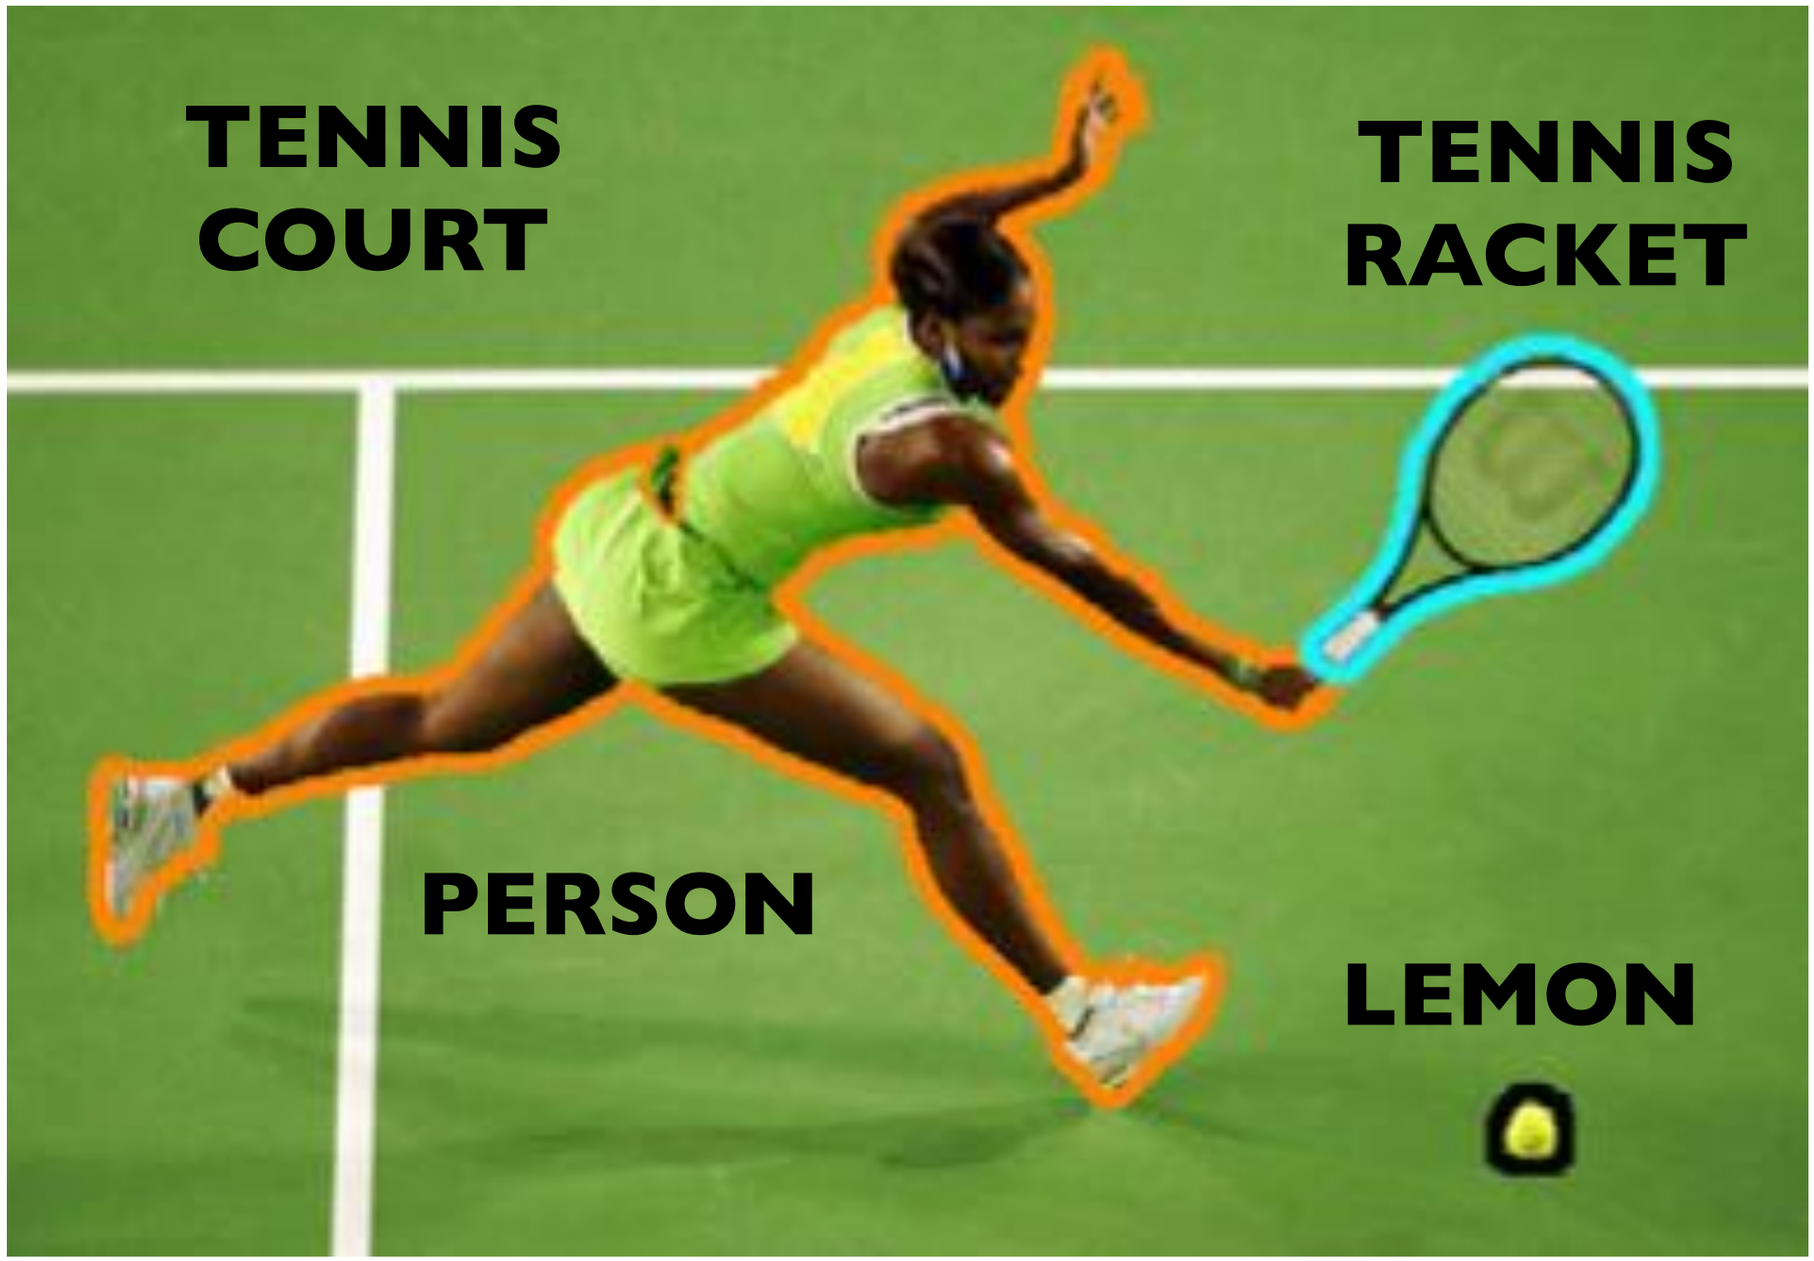

[Rabinovich, Galleguillos, Wiewiora, Belongie 2007]

Definitions of Computer Vision #3

“Image/Video Understanding”

[cityscapes-dataset.com]

3R's of Computer Vision #1

"Registration"

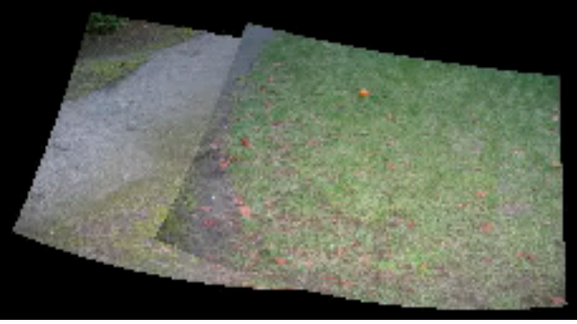

[Brown Lowe 2003, 2007]

3R's of Computer Vision #1

"Registration"

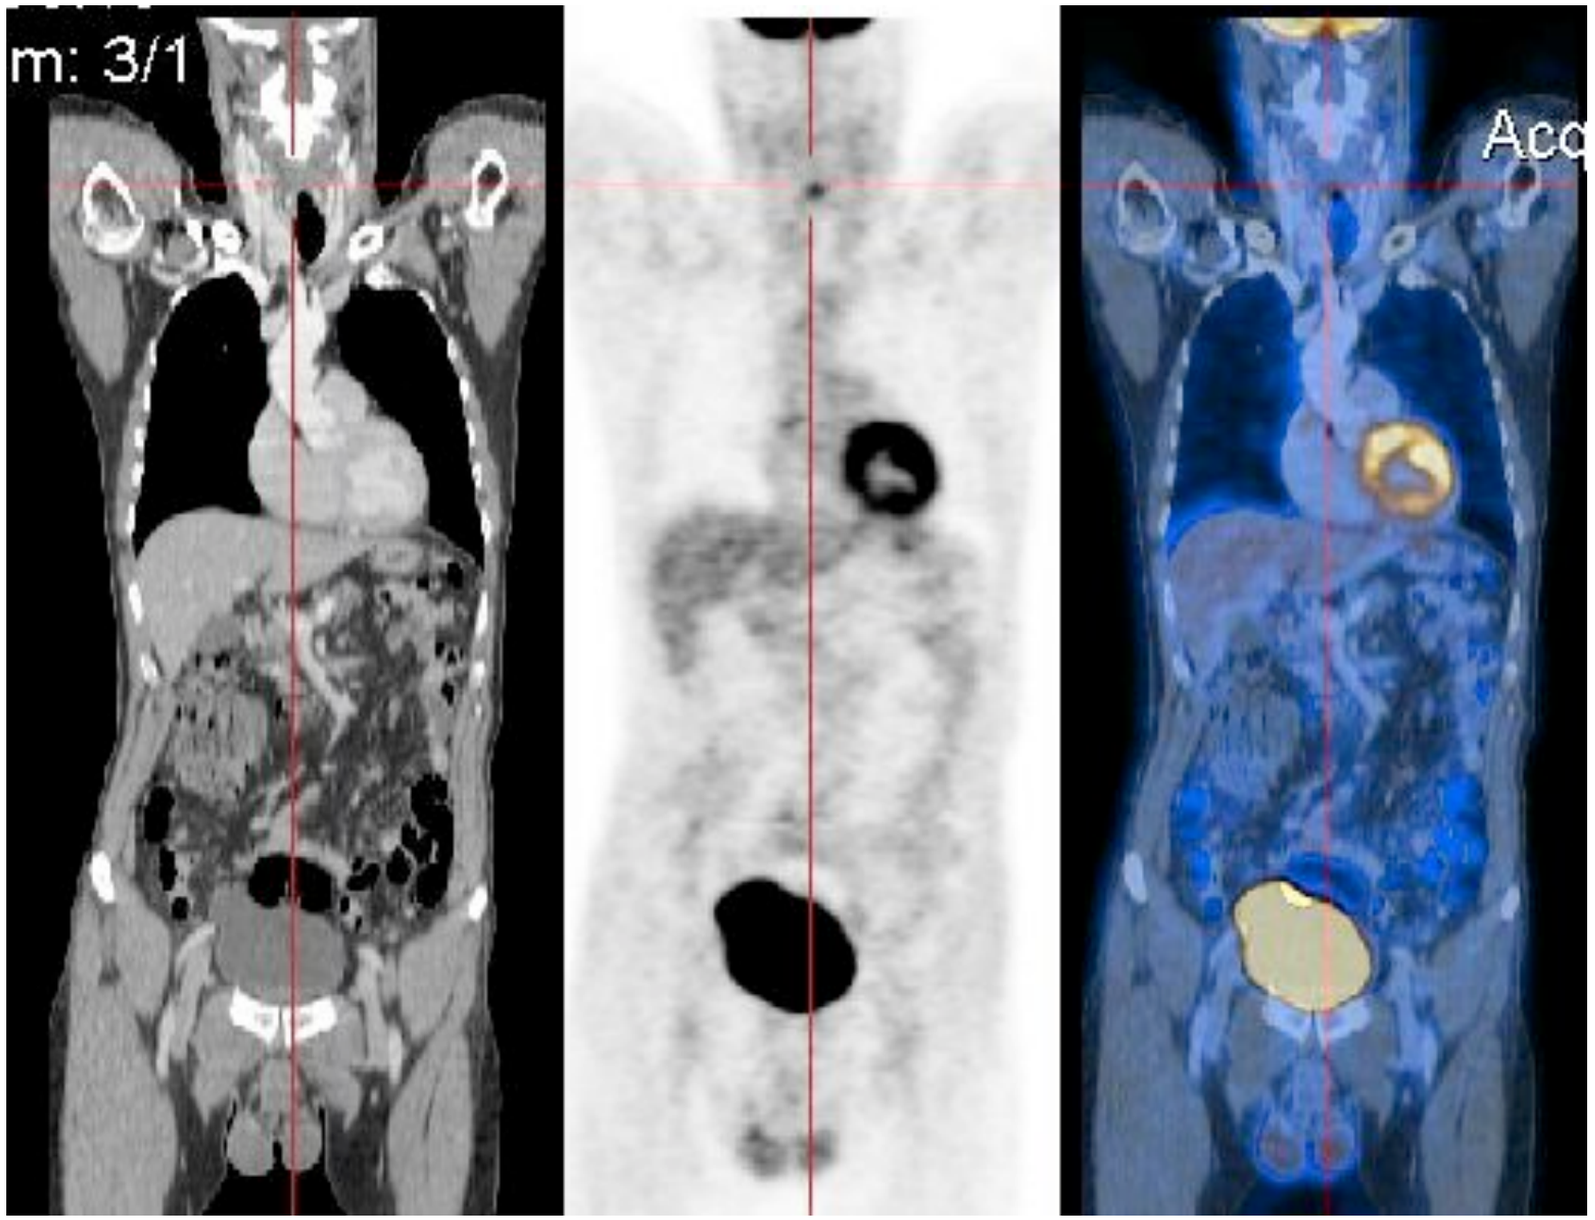

CT PET [Wikipedia]

3R's of Computer Vision #2

"Reconstruction"

Structure from Motion
[Noah Snavely]

Multi-view Stereo
[Y. Furukawa]

3R's of Computer Vision #2

"Reconstruction"

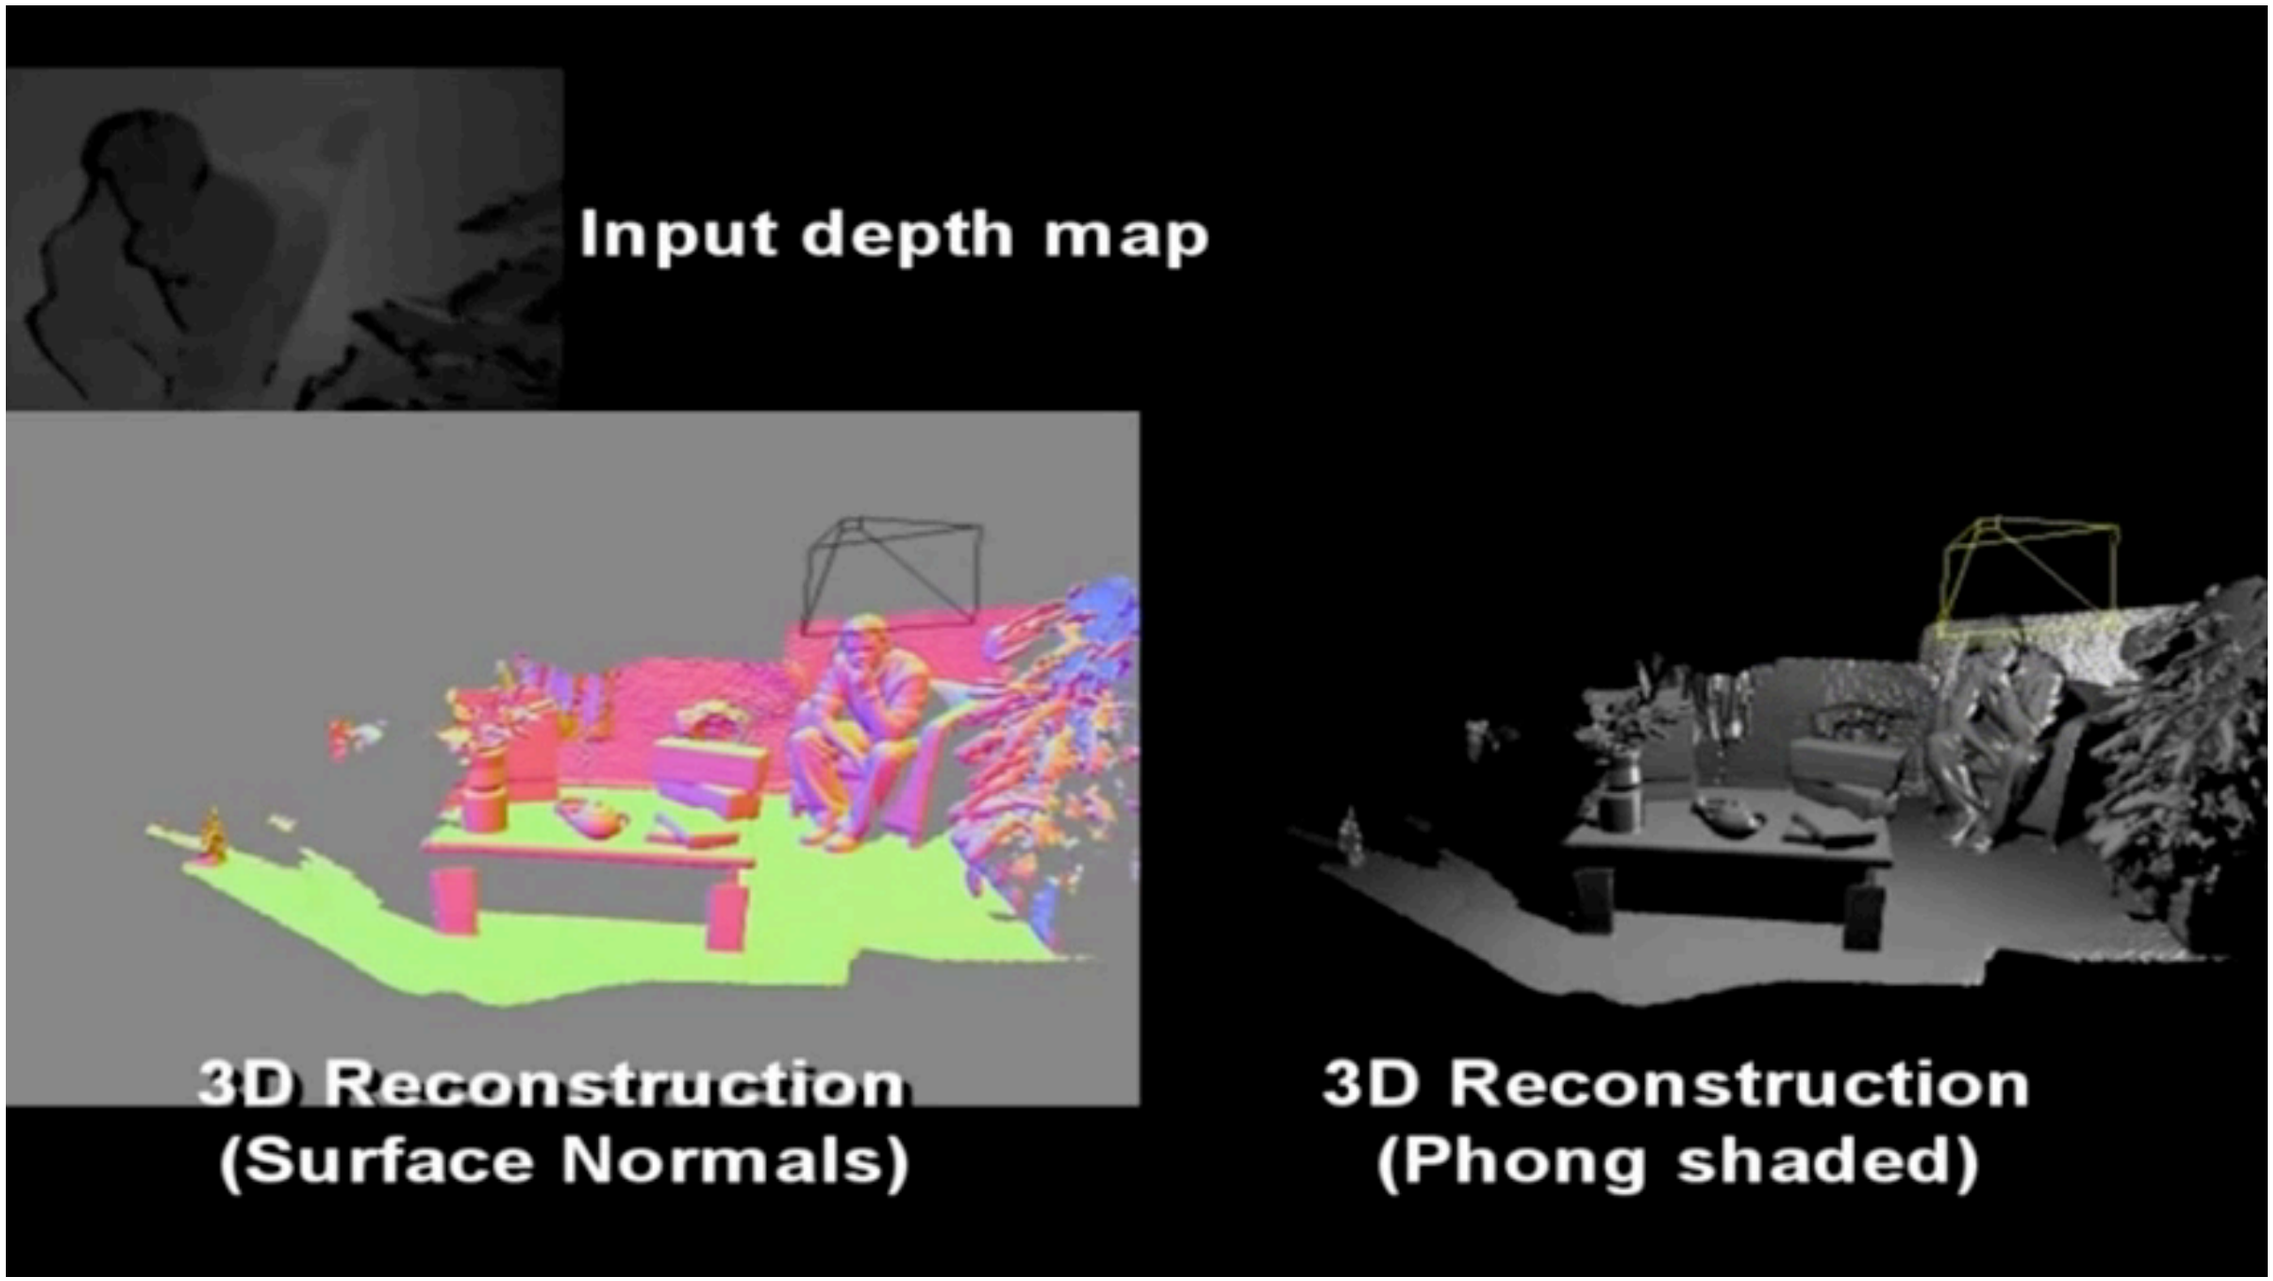

[KinectFusion Izadi et al]

3R's of Computer Vision #3

"Recognition"

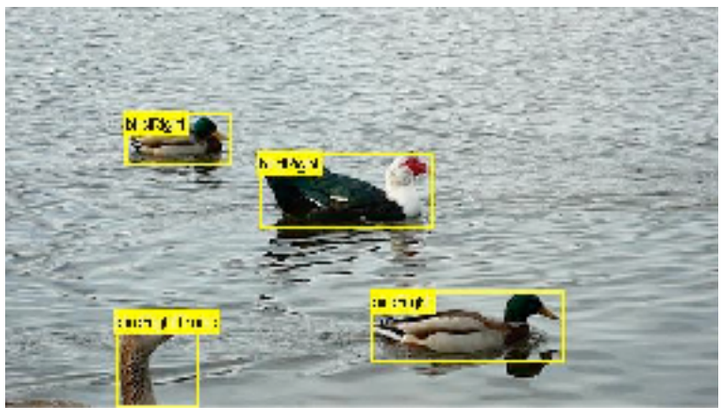

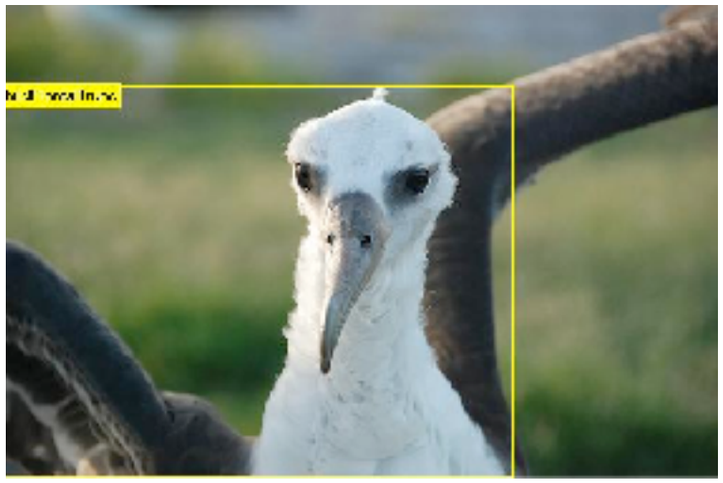

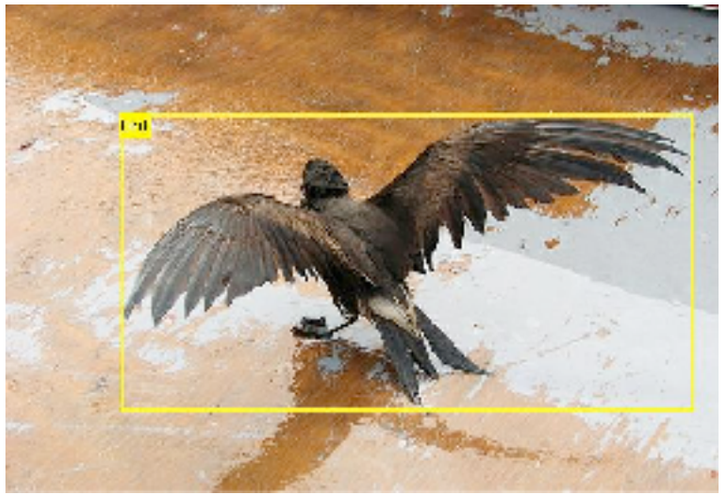

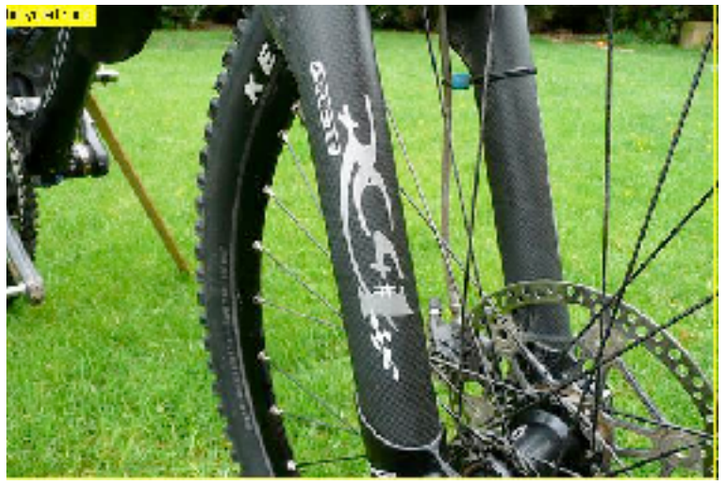

3R's of Computer Vision #3

"Recognition"

Instance Recognition

Category Recognition

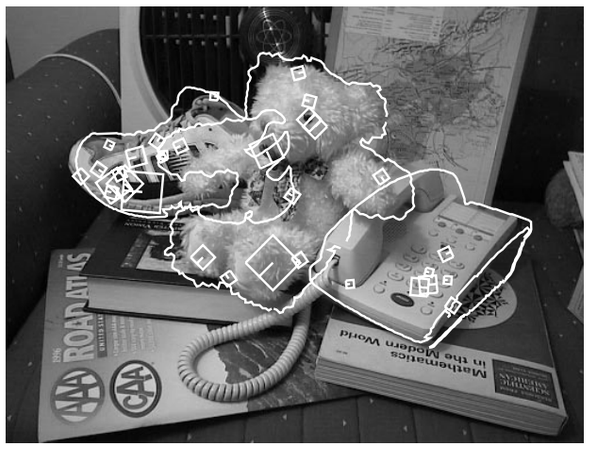

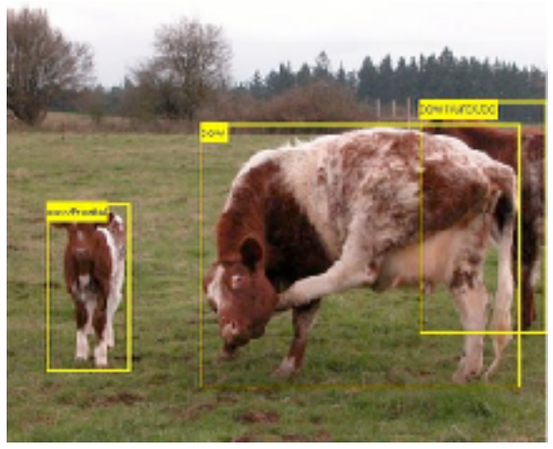

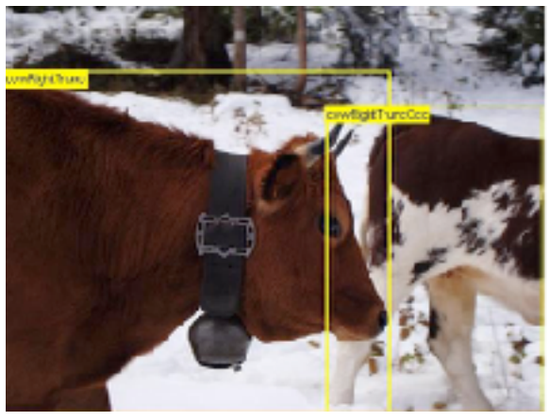

[Lowe 1999, 2004]

[Pascal Challenge]

S = Segmentation

Input Image

Instance Segmentation

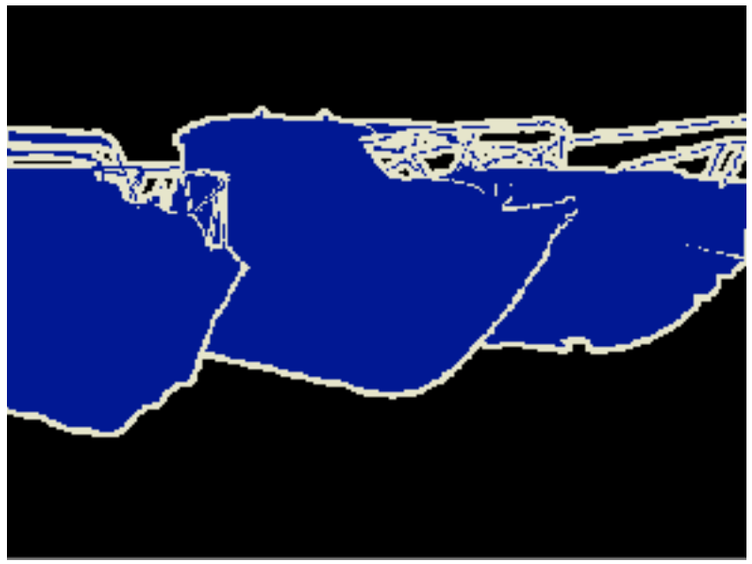

Category Segmentation

T = Tracking

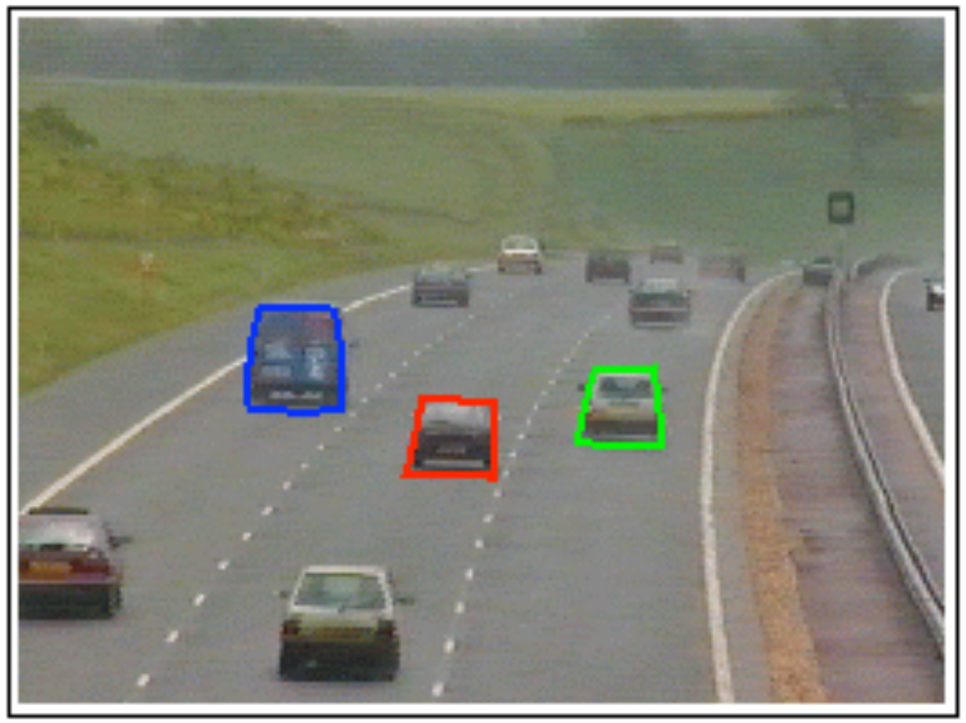

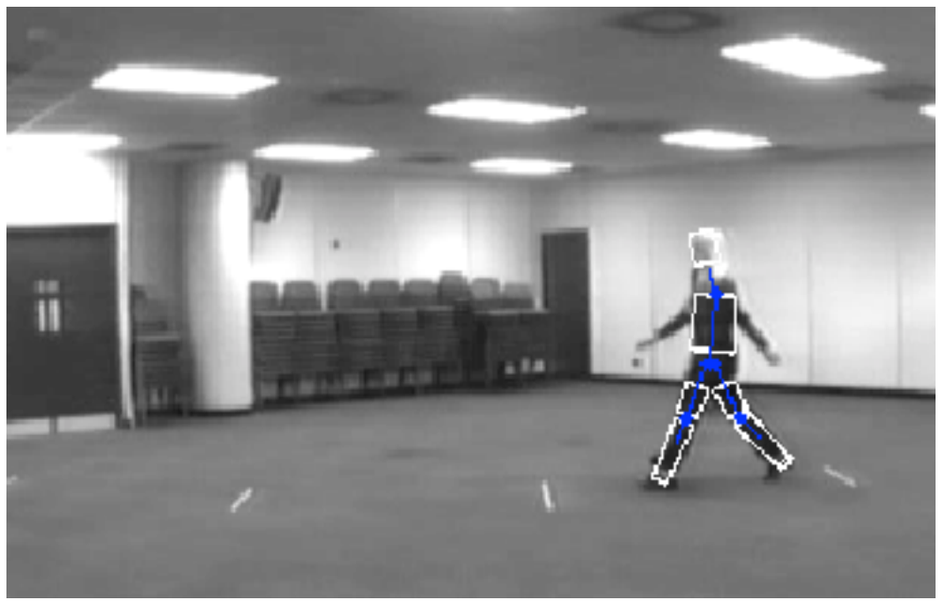

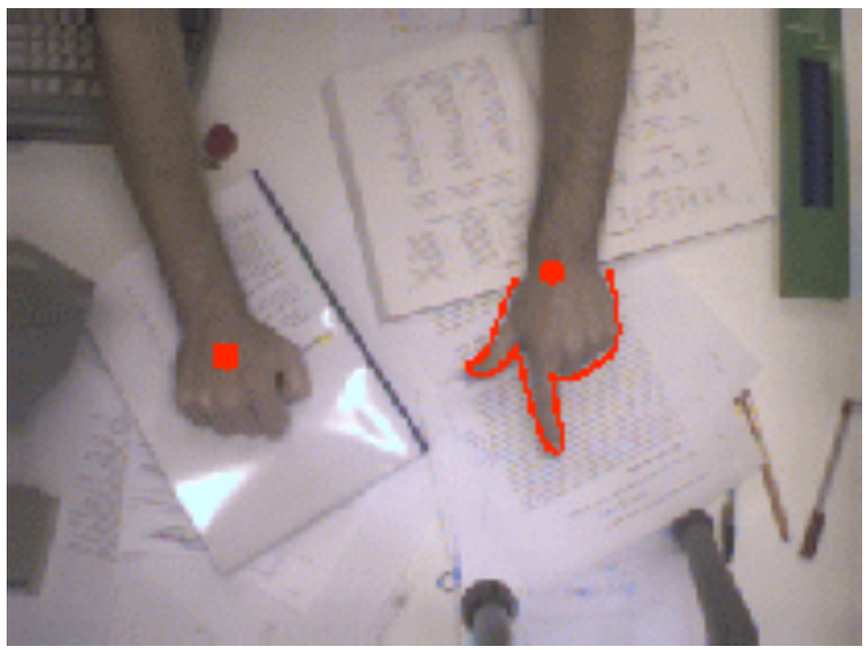

[Active Contours, Blake Isard 1998]

This Course

- Computer Vision, with emphasis on **visual geometry + deep learning**
  - Image Formation, Low-level processing, Camera models, 2D/3D geometry, Stereo, Optical Flow, Multi-view methods, Deep Learning for vision, CNNs, Regression/Classification, Applications
- 10 lectures, +office hours
- 4 projects, equally weighted
- **Project 1**: Feature Extraction and Matching
- **Project 2**: Panoramic Image Stitching
- **Project 3**: Image Classification using CNNs
- **Project 4**: Deep Learning for Stereo Matching
- Projects will use iPython notebooks (e.g., Jupyter, Colab)
- Numpy for numerics
- Tensorflow for machine learning

Recommended Text 1

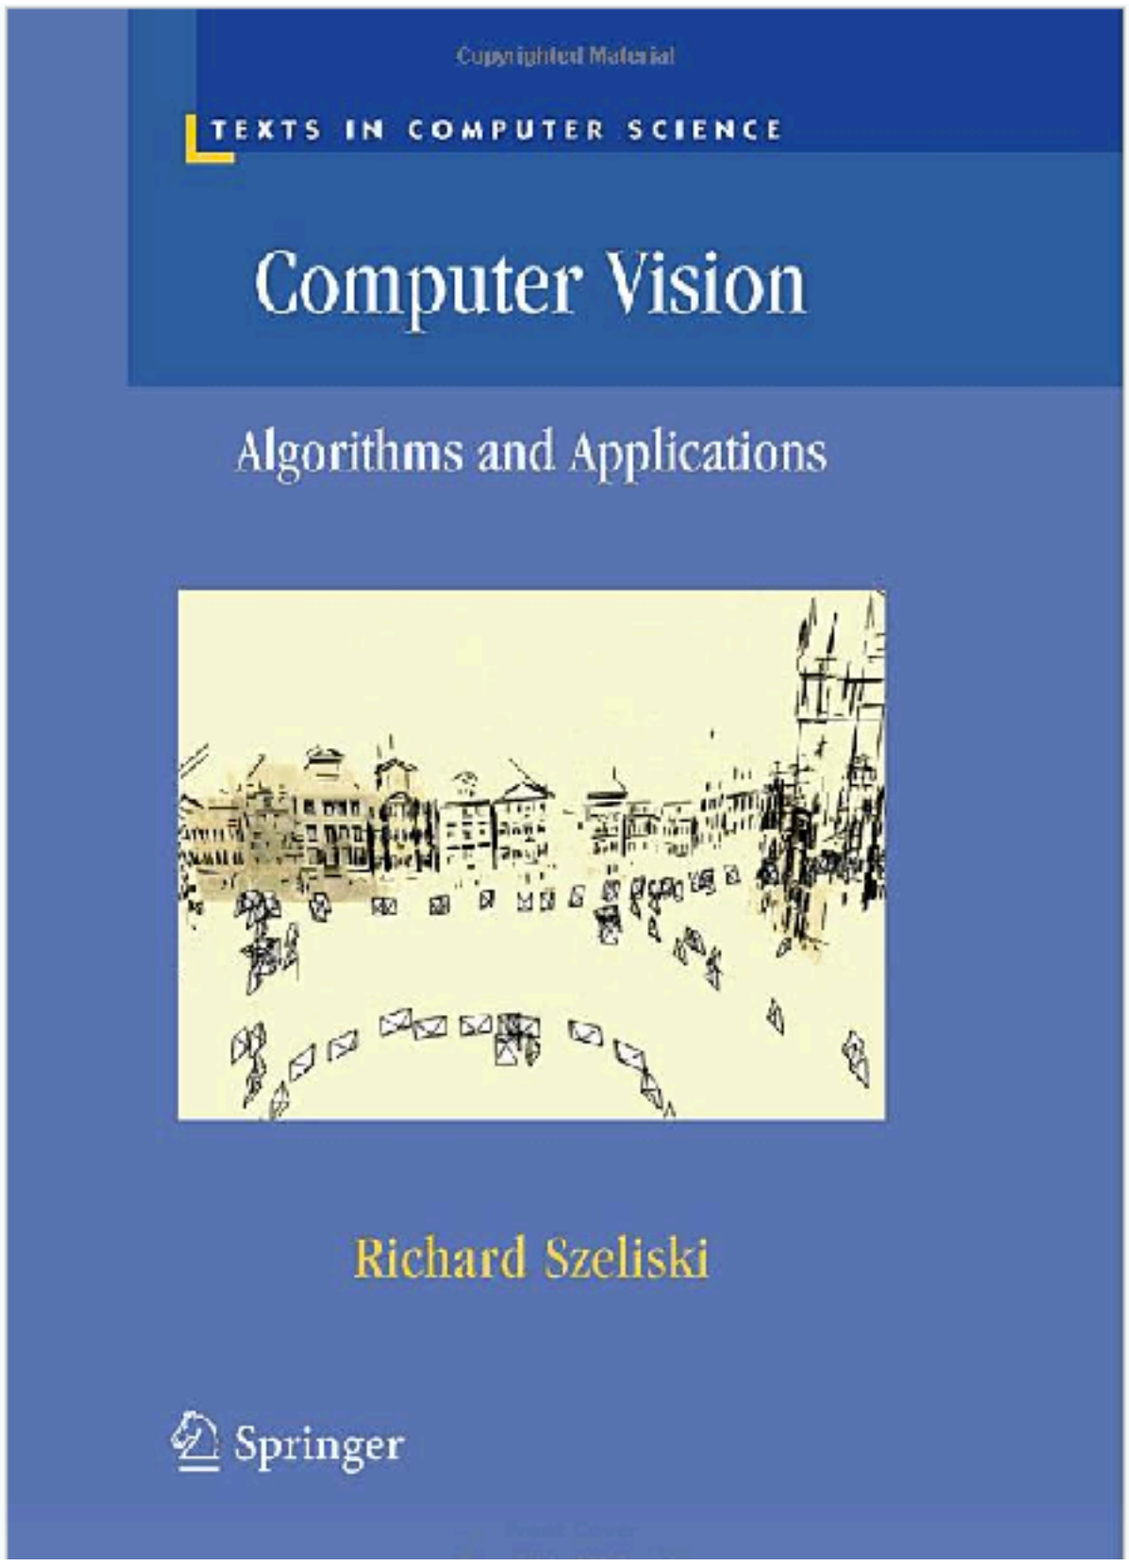

Computer Vision: Algorithms and Applications

Richard Szeliski

http://szeliski.org/Book

Core textbook for the course. Very good coverage of all topics, oriented around practical applications

Computer Vision: Szeliski

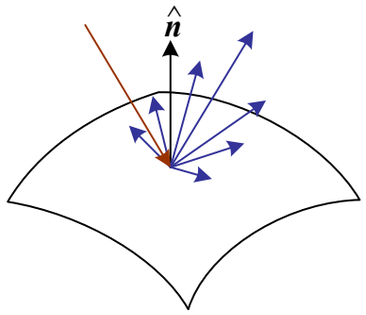

2. Image Formation

3. Image Processing

4. Features

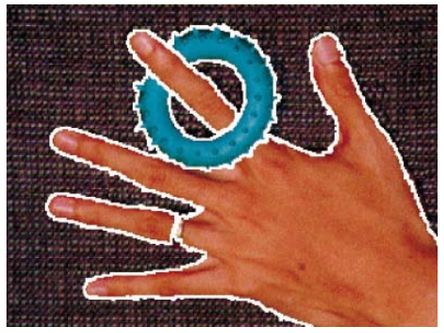

5. Segmentation

6-7. Structure from Motion

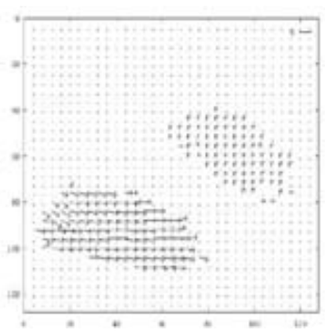

8. Motion

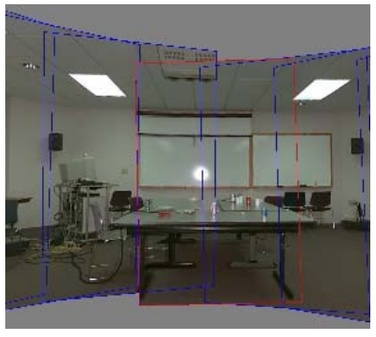

9. Stitching

10. Computational Photography

11. Stereo

12. 3D Shape

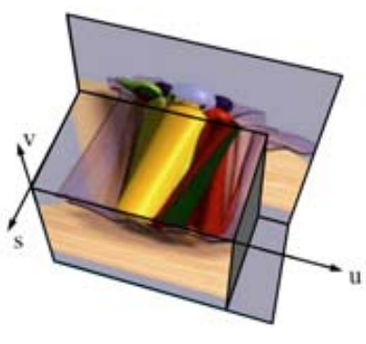

13. Image-based Rendering

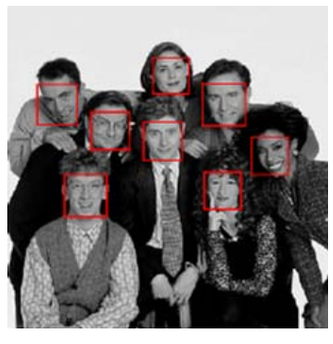

14. Recognition

+ Learning

Recommended Text 2

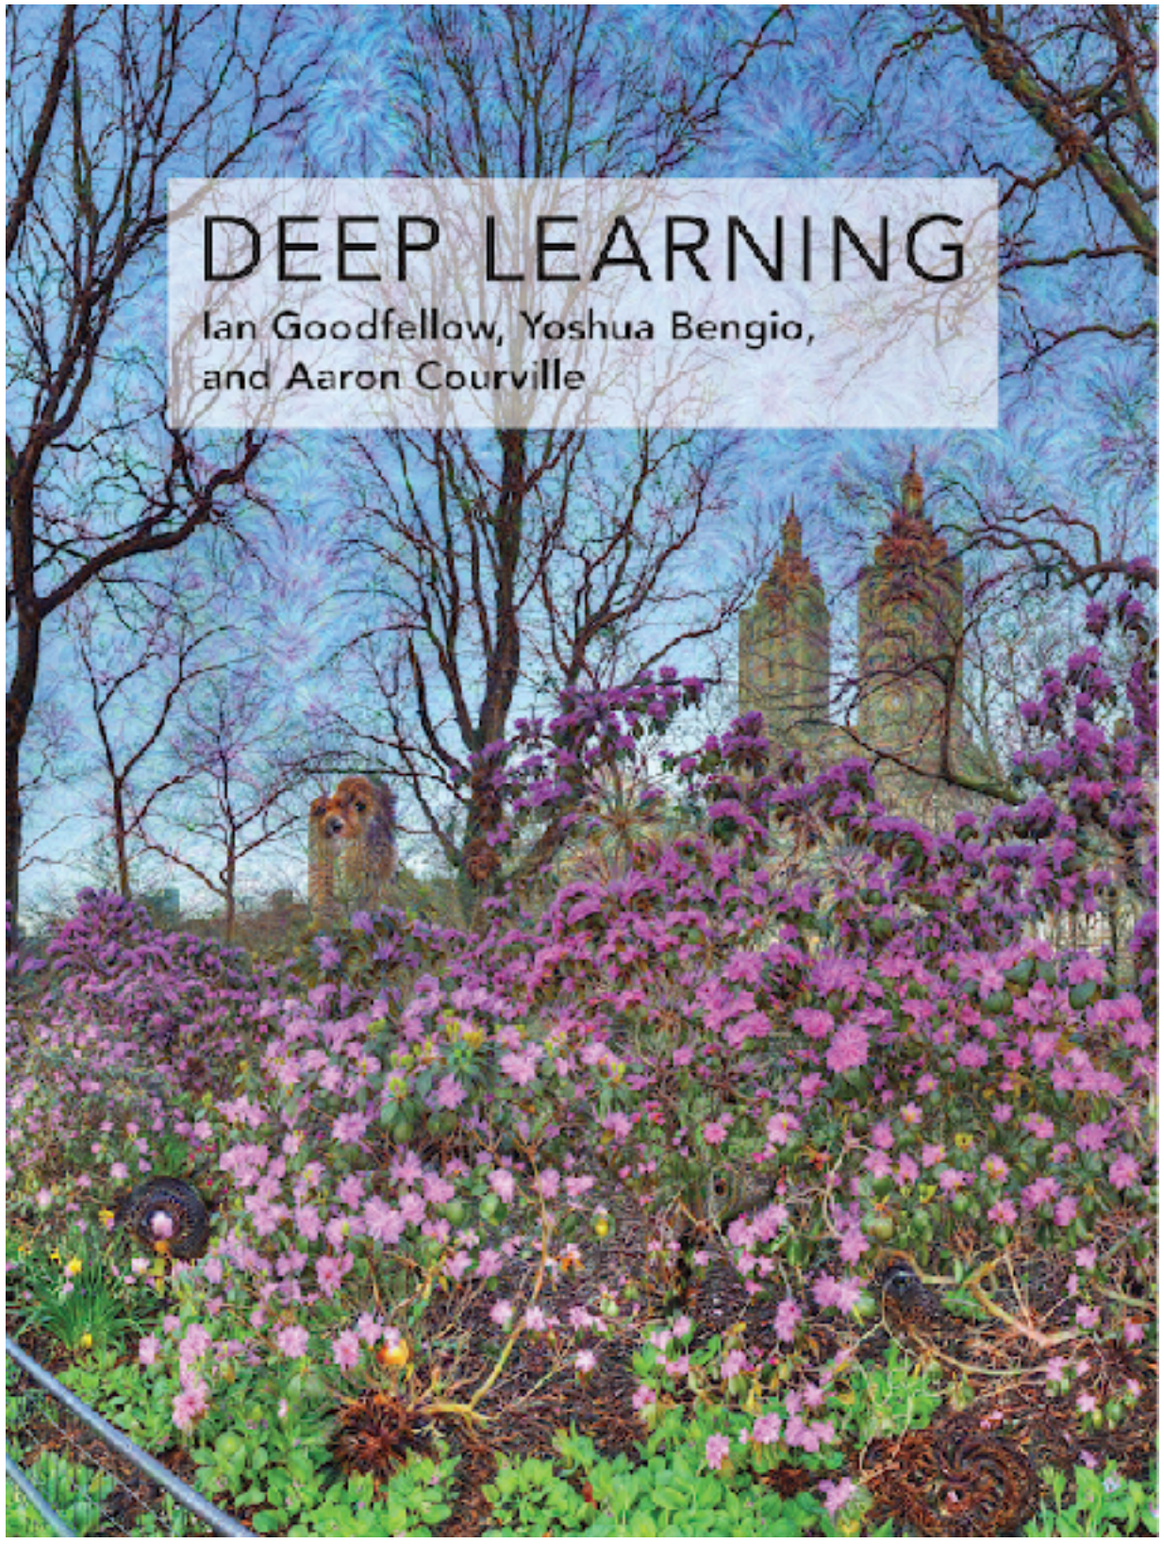

Deep Learning: Goodfellow, Bengio, Courville

deeplearningbook.org

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

Inceptionism

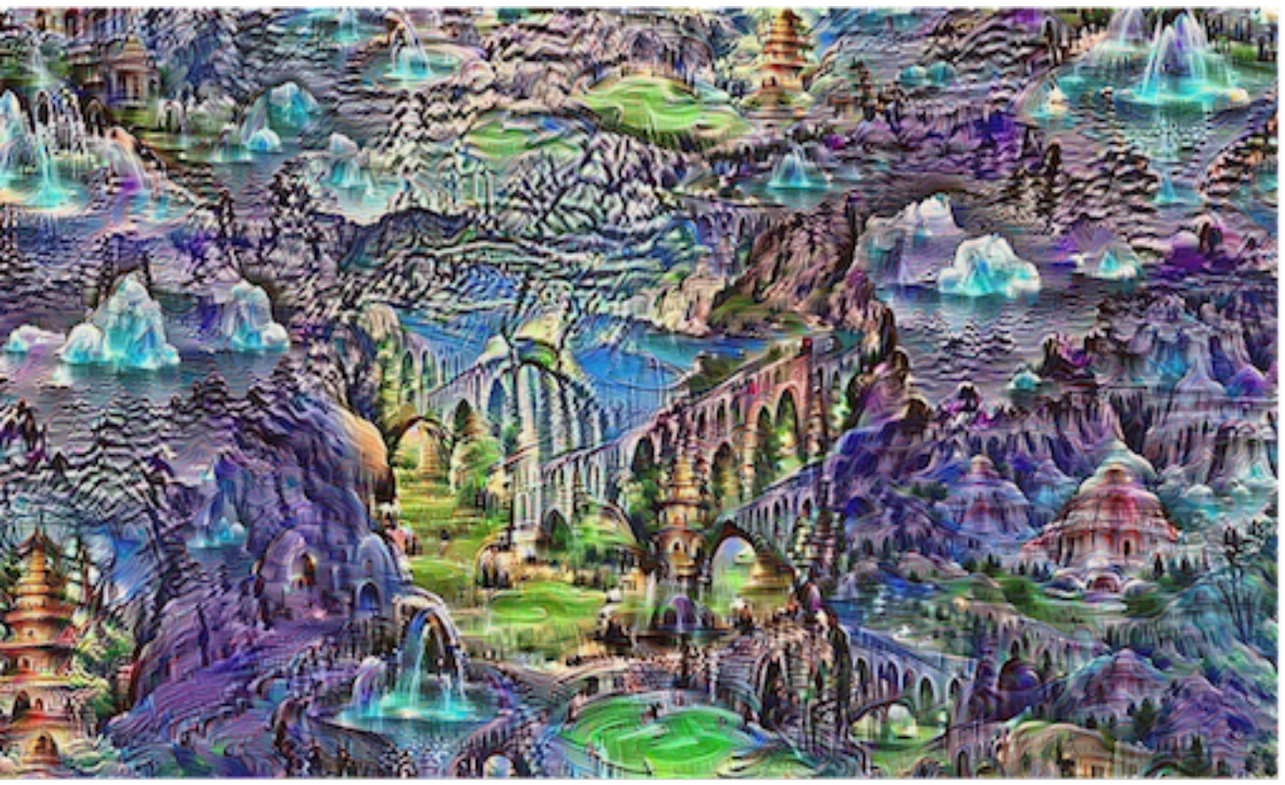

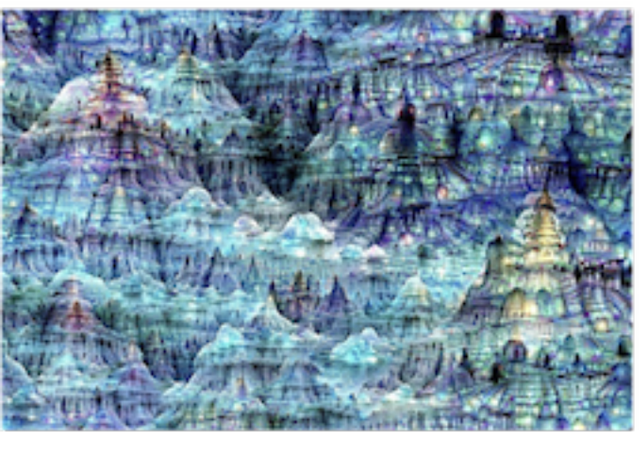
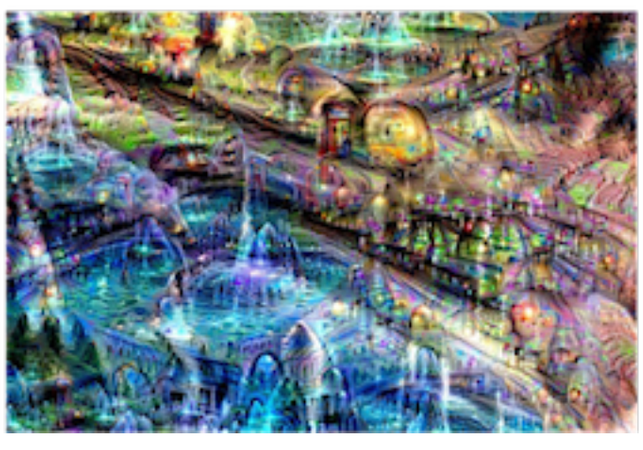

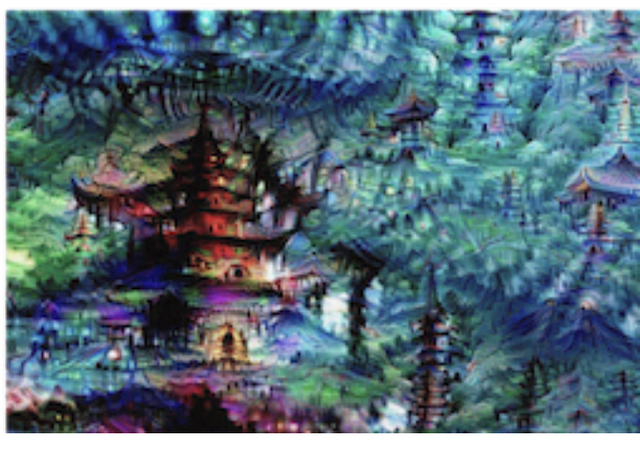

[Mordvintsev, Olah, Tyka 2015]

Schedule

Date	Lecture Topics	Project
03/29	Introduction, Image Formation	
04/05	Filtering and Pyramids, Features and Matching	P1 assigned
04/12	Planar + Epipolar Geometry. 2-view Alignment, RANSAC	
04/19	Multiview Geometry, SFM/ SLAM, Optimization	P1 due P2 assigned
04/26	Dense correspondence, Stereo, Flow, Depth cams	

Schedule

Week Begin	Lecture Topics	Lab
05/03	Machine Learning, NN, SVM, Decision Trees, Boosting	P2 due P3 assigned
05/10	Linear/Logistic Regression, NNets, CNNs, Backprop	
05/17	Object Detection, Instance Segmentation	P3 due P4 assigned
05/24	Tracking, SLAM	
05/31	Depth Estimation, History and Futurology	P4 due

Break

- Next: 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