Digital Photography and Geometry Capture

NBAY 6120

March 7, 2017

Donald P. Greenberg

Lecture 4

Required Reading

- N. Snavely, S.M. Seitz, and R. Szeliski, “Photo Tourism: Exploring Photo Collections in 3D,” *ACM Trans. Graphics, July 2006, pp. 835-846.* *http://phototour.cs.washington.edu/Photo_Tourism.pdf*

- Bilger, Burkhard. “Auto Correct: Has the Self-Driving Car Arrived at Last?" *The New Yorker*. N.p., 25 Nov. 2013. Web. 10 Sept. 2015.
  - http://www.newyorker.com/magazine/2013/11/25/auto-correct

Required Reading

- Raffi Khatchadourian. "We Know How You Feel, The New Yorker, January 19, 2015. The New Yorker

William F. Strapp, “Technology of the Por

Lincoln’s Daguerreotype

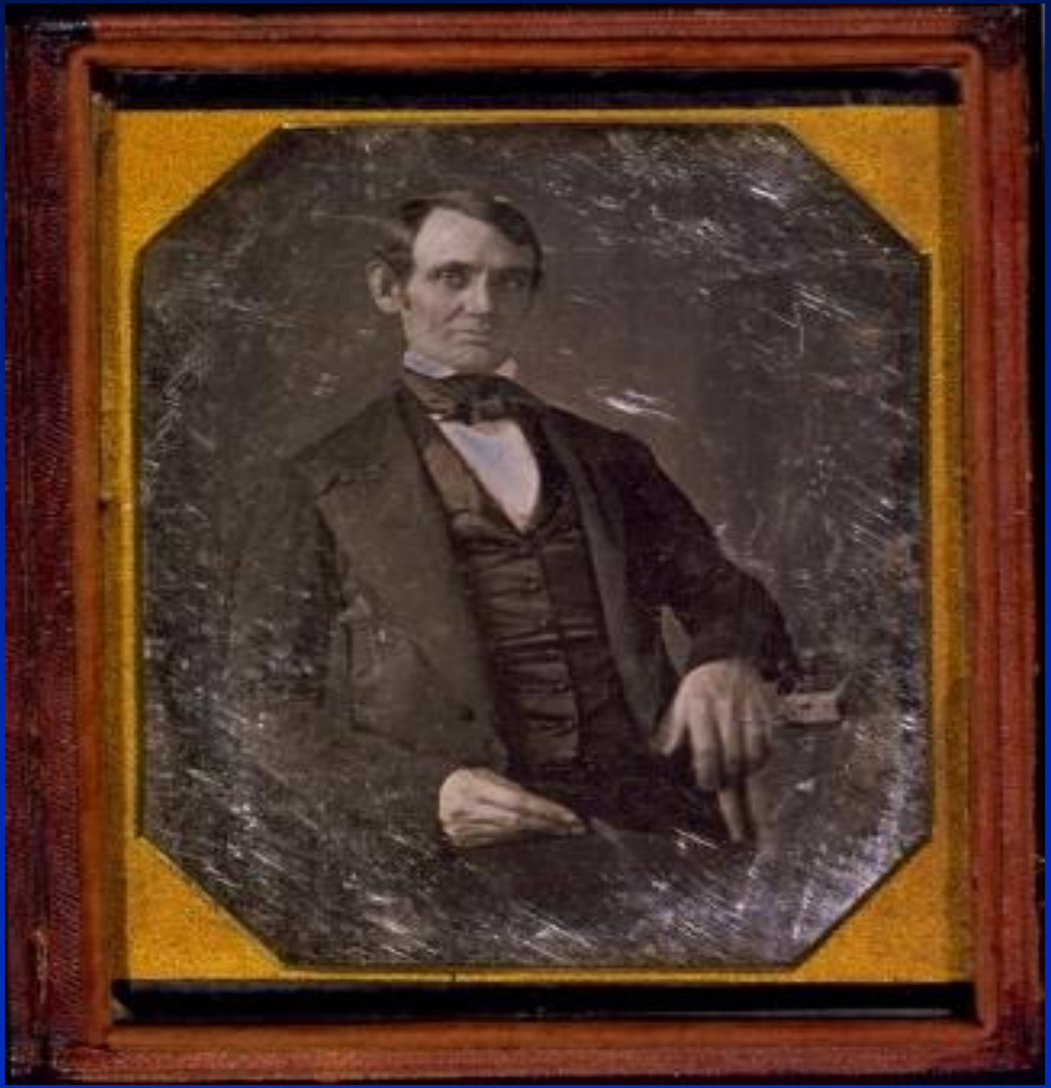

Eadweard Muybridge
- Galloping Horse 1878

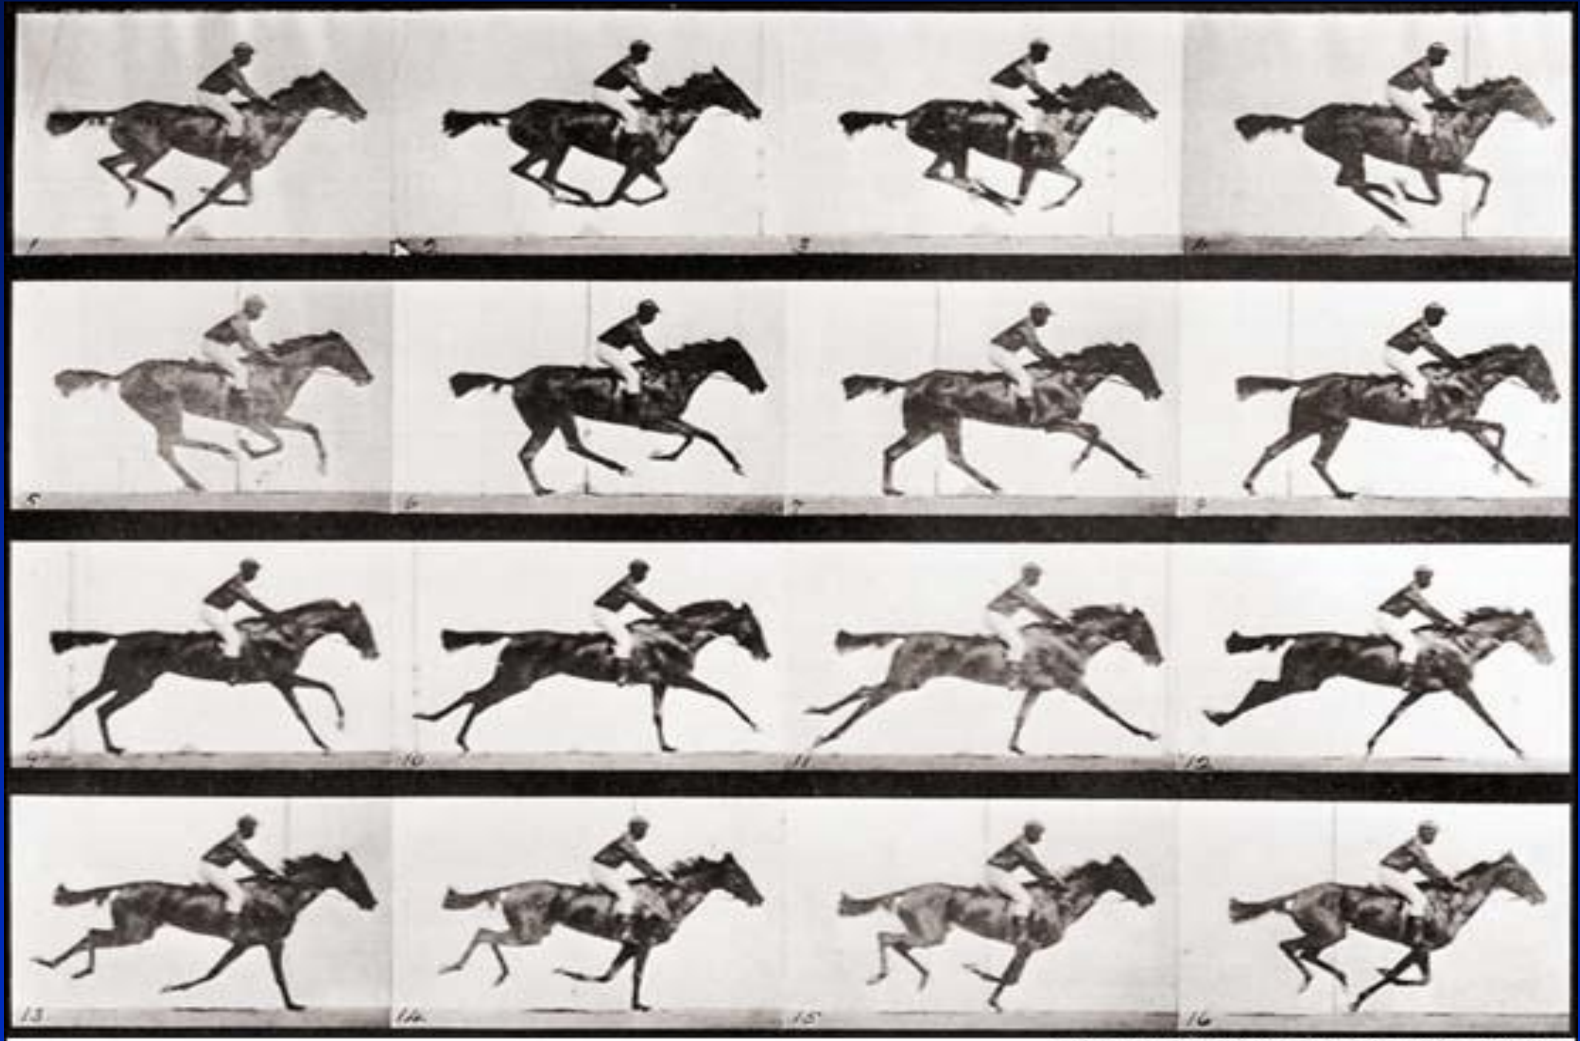

Photographs by Eadweard Muybridge

from *The Story of Kodak*, Douglas Co

Eastman and box - *(from The Story of Kodak, I*

Color Film Paradigm Shift

From multiple lenses or multiple exposures
to multiple layered film

The transition from the optical approach to the
chemical approach formed the new basis for
color photography

Mannes & Godowsky
1920's

(The Story of Kodak, Douglas Collin

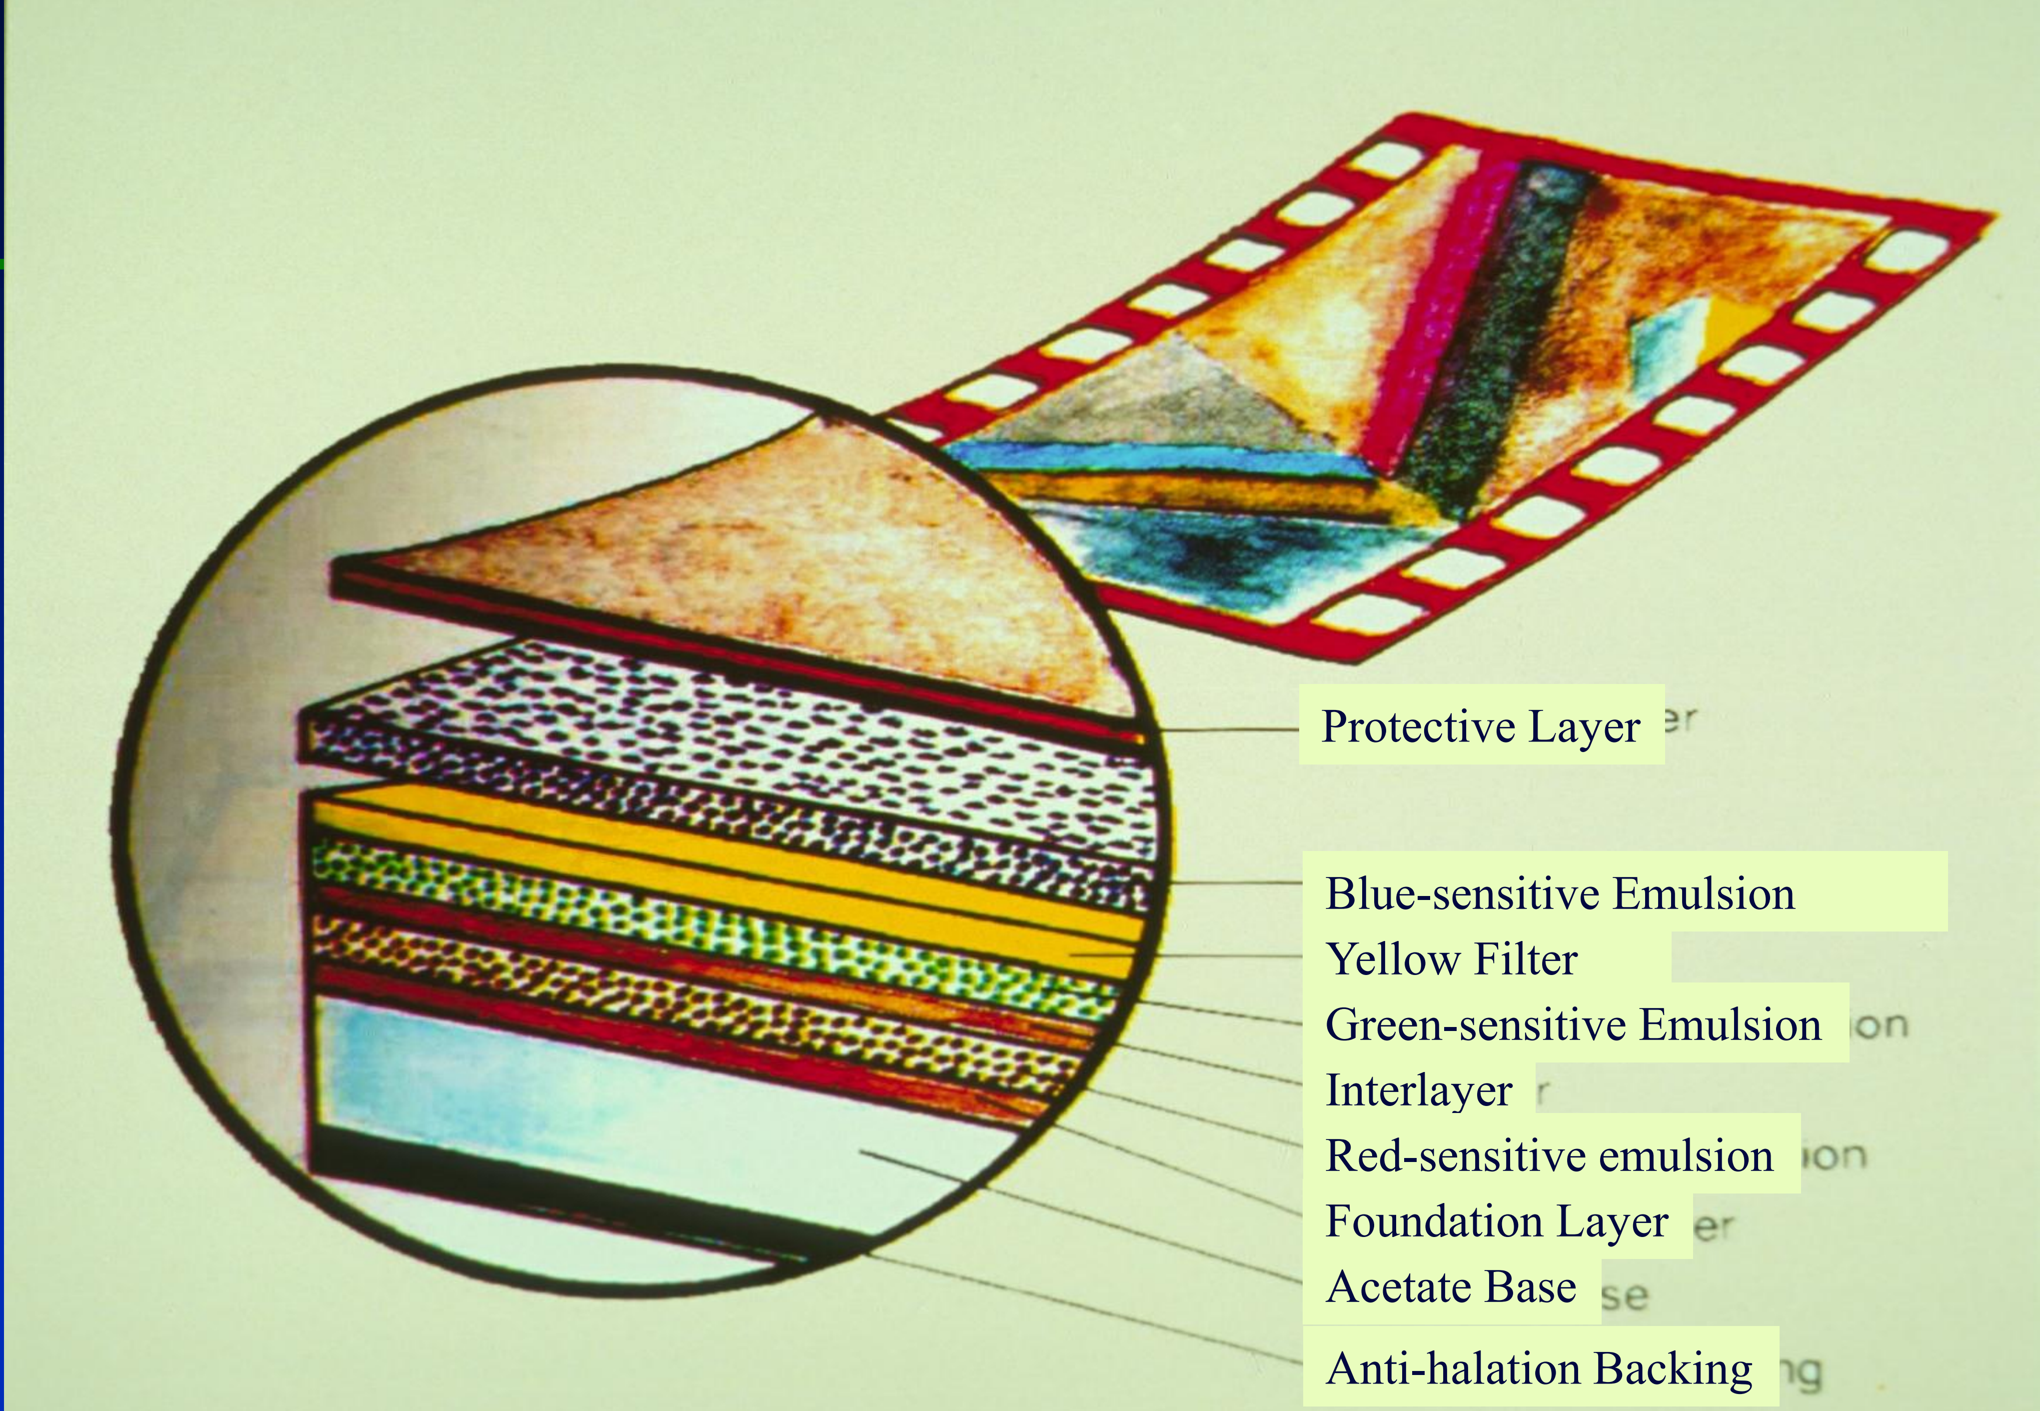

(fig. 1.6, Color Photography, Robert Hirsch

Old Days - You Dropped Off Your Roll And Got Prints Back

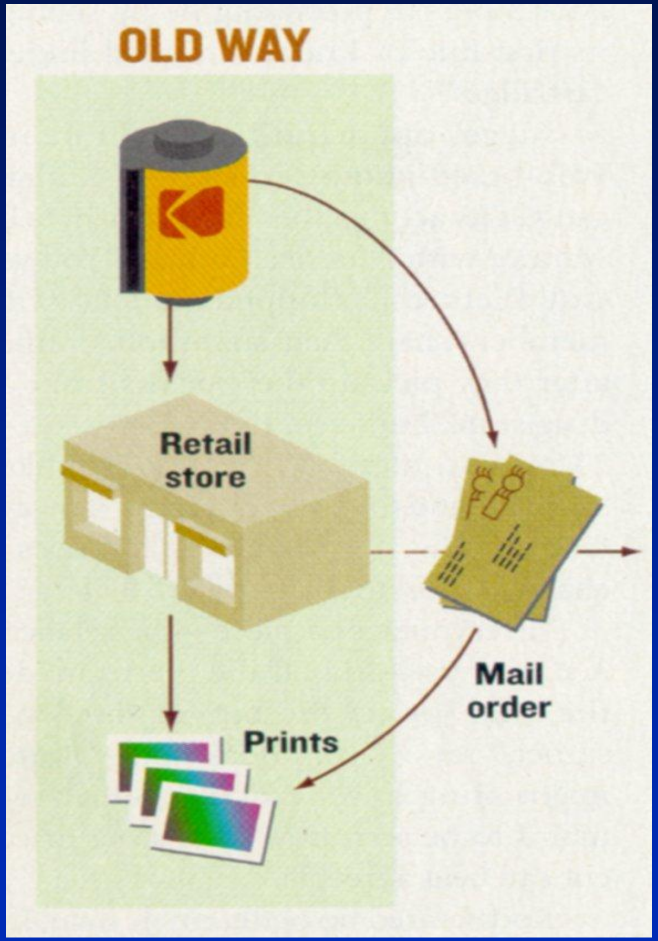

Polaroid Land Camera

Photo-detector Technology

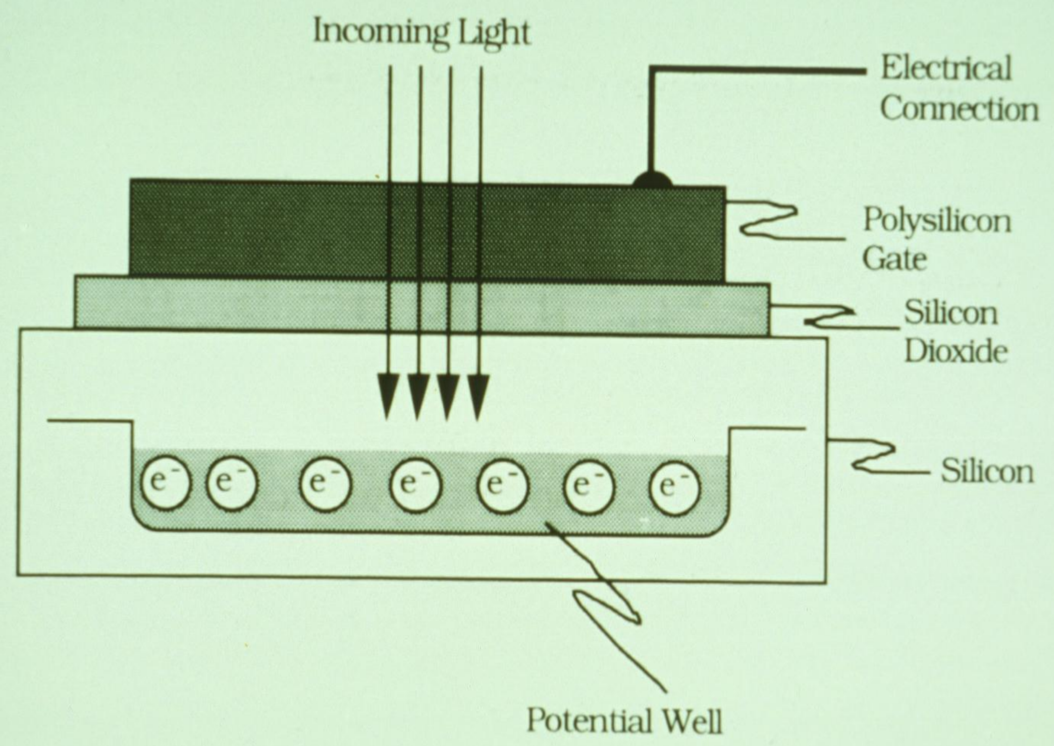

Fig. 2. Simple photodetector

(Charge-Coupled Devices for Quantitative Electronic Imaging, 19

CMOS Technology

- Complementary metal oxide semiconductor
- Cheaper manufacturing technology than CCD’s
  - Follows the semiconductor industry cost curves
  - Reduces the number of chips/camera required
- Processing (which is “free”) can perform calculations on each pixel within frame time (e.g. correct for lighting, motion blur, etc.).

Bayer Pattern

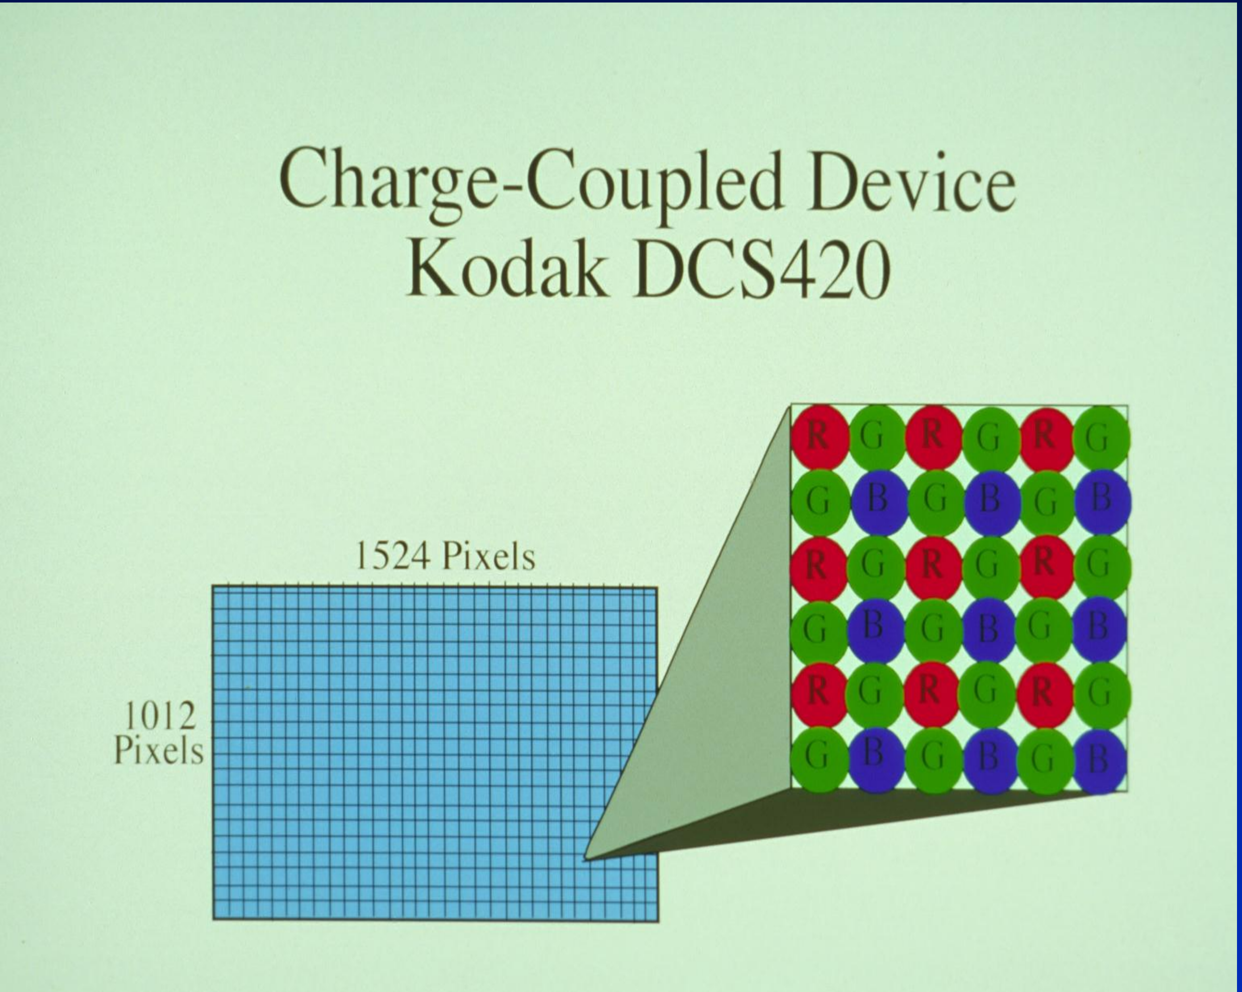

$14,000 approximately, June 1994


SPECIAL REPORT
THE BEST IN EXECUTIVE EDUCATION
BusinessWeek
OCTOBER 20, 1997
A PUBLICATION OF THE McGRAW-HILL COMPANIES
$3.50
Kodak
Royal Gold
200
CAN GEORGE FISHER FIX KODAK?
PAGE 116
Internet: www.businessweek.com America Online: Keyword: BW


GLOBAL 1000 THE TOP COMPANIES
TOBACCO HOW THE DEAL WILL CHANGE
U.S. ECONOMY NEW NUMBERS, ROSIER VIEW
BusinessWeek
JULY 7, 1997
A PUBLICATION OF THE McGRAW-HILL COMPANIES
$3.50
SHOOTOUT!
LEW PLATT, CEO, HEWLETT-PACKARD
HOW HP PLANS TO TAKE ON KODAK— AND REVOLUTIONIZE THE WAY YOU CAPTURE AND PRINT IMAGES
PAGE 100
Internet: www.businessweek.com America Online: Keyword: BW

Requirements For Pervasive Digital Photography

- High resolution, low cost image acquisition devices
- Sufficient computer processing power and memory systems for digital manipulation
- Image enhancement software with easy-to-use interfaces
- High density, low-cost local storage systems

Requirements For Pervasive Digital Photography

- Cheap LCD displays for previewing

- Bandwidth! Bandwidth! Bandwidth!
  - High network bandwidth (wired) for distant transmission
  - Fast throughput (e.g. Firewire) for local transmission
  - Wireless bandwidth (local) for ease of use

- High quality, low cost digital printers

CONSUMER
Digital Cameras 2012

Sony CyberShot
20 MegaPixels
$80

Kodak EASYSHARE Touch
M5370
Cost: $129.95
16 Megapixels

PROFESSIONAL
Digital Cameras 2014

Canon EOS 5DSR
50.6 MegaPixels
$3,899

Nikon Digital SLR
16.2 MegaPixels
$5,999

When is a phone not a phone?

…when it’s a camera...

iPhone 6S Camera – 12 MPixels

Nokia Lumia 1020 – 41 MPixels

Advanced Digital Photography

Eye of a Fly

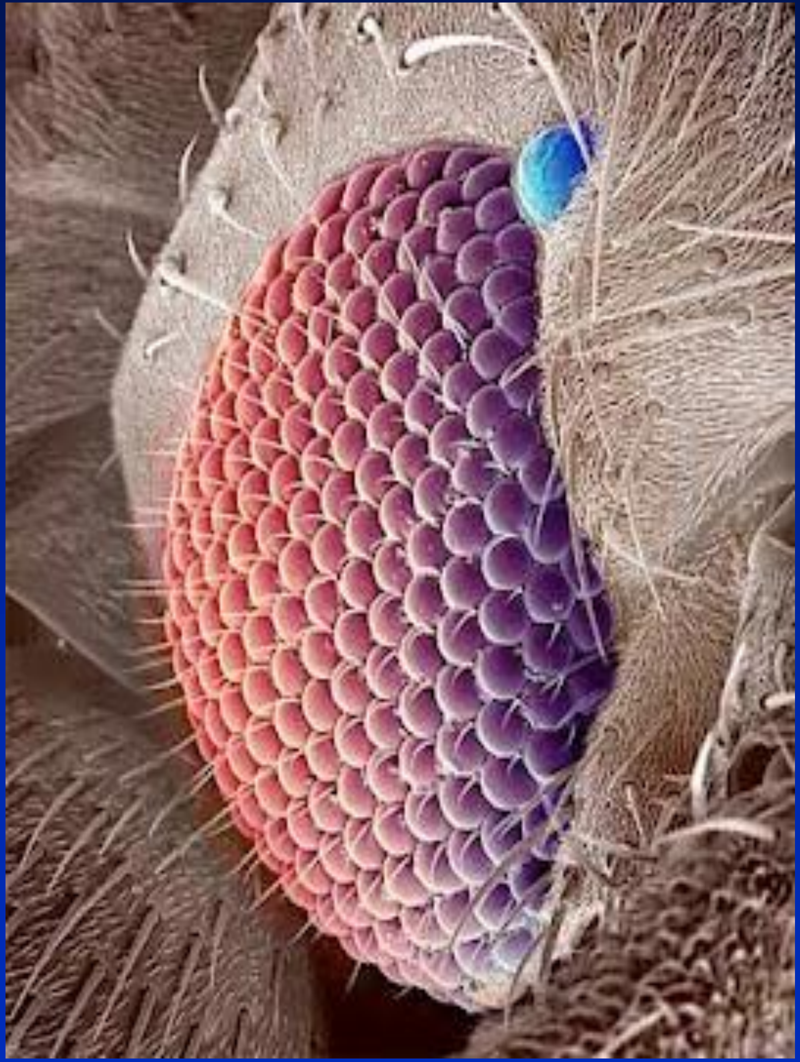

AWARE-2 Duke University

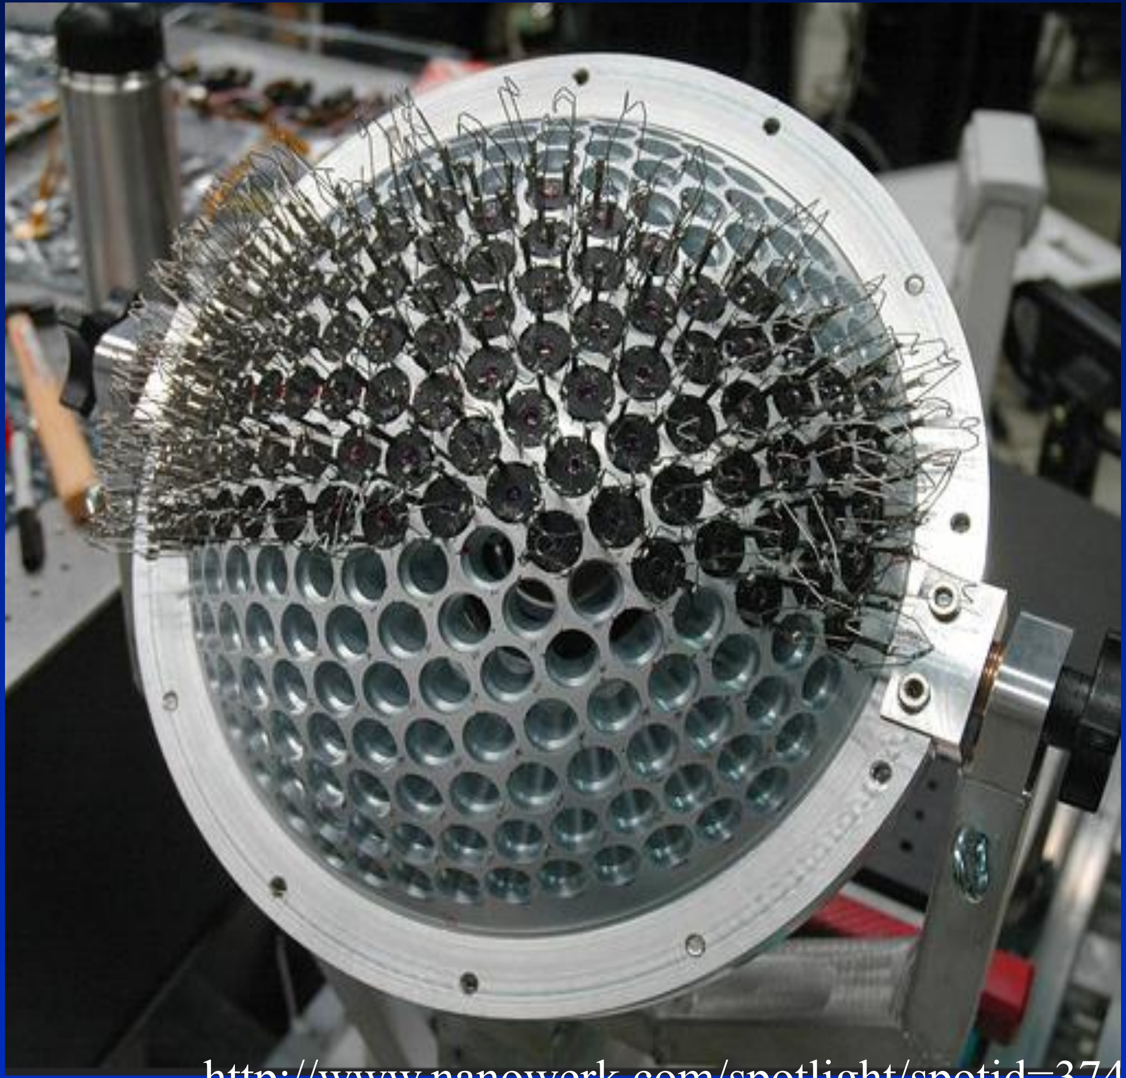

http://www.nanowerk.com/spotlight/spotid=3744.php

AWARE-2

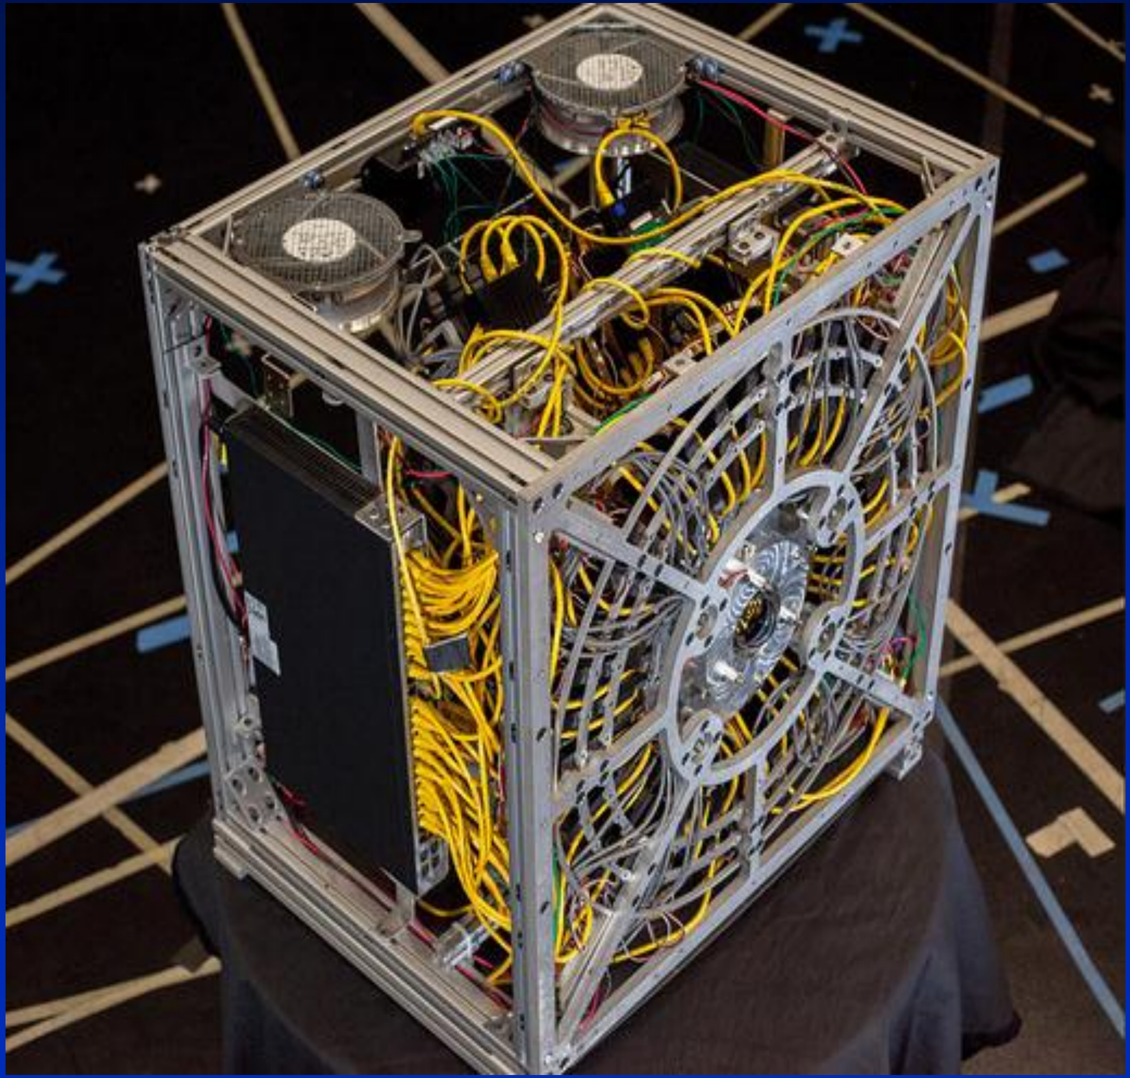

http://mosaic.disp.duke.edu:90/aware/image_list/image_li

AWARE-2

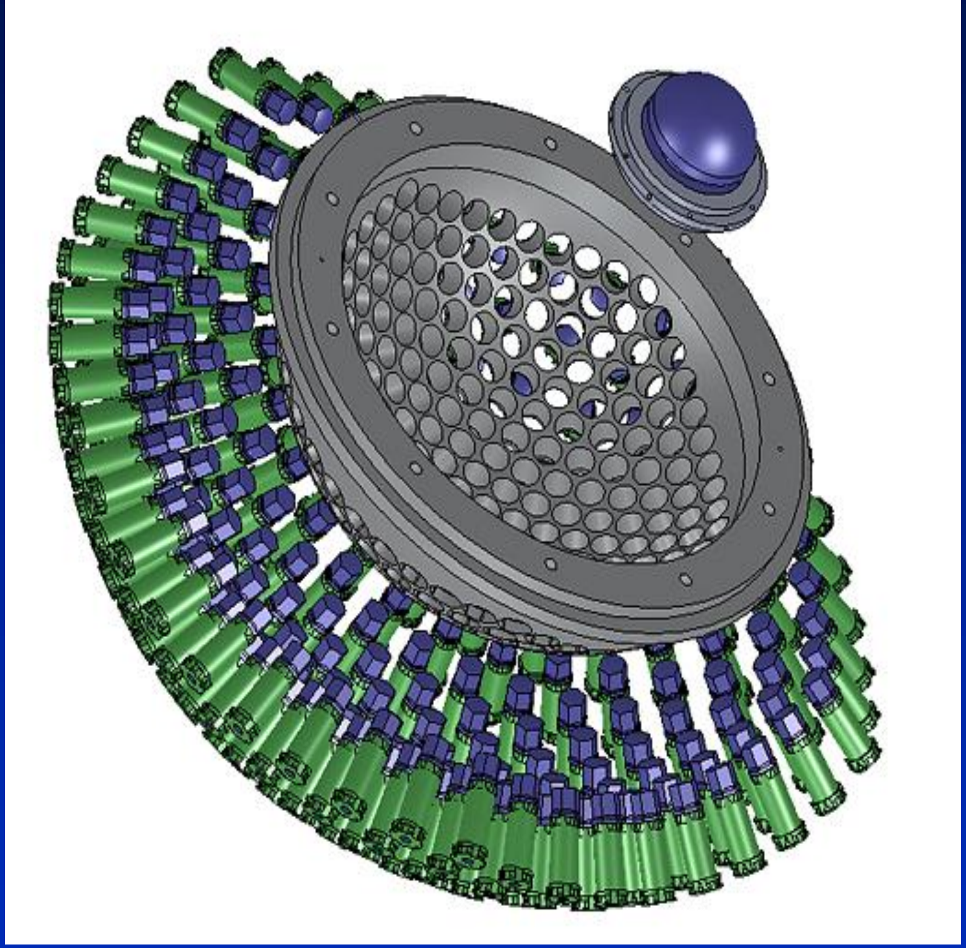

http://mosaic.disp.duke.edu:90/aware/image_list/image_list/public

- http://gigapan.com/galleries/11088/gigapans/146504

Gigapixel Images Prof. Pedro Sander HKUST

- http://gigapan.com/gigapans/58857

Canon’s 250-megapixel camera sensor 09/08/15

- Can read letters 11 miles away!

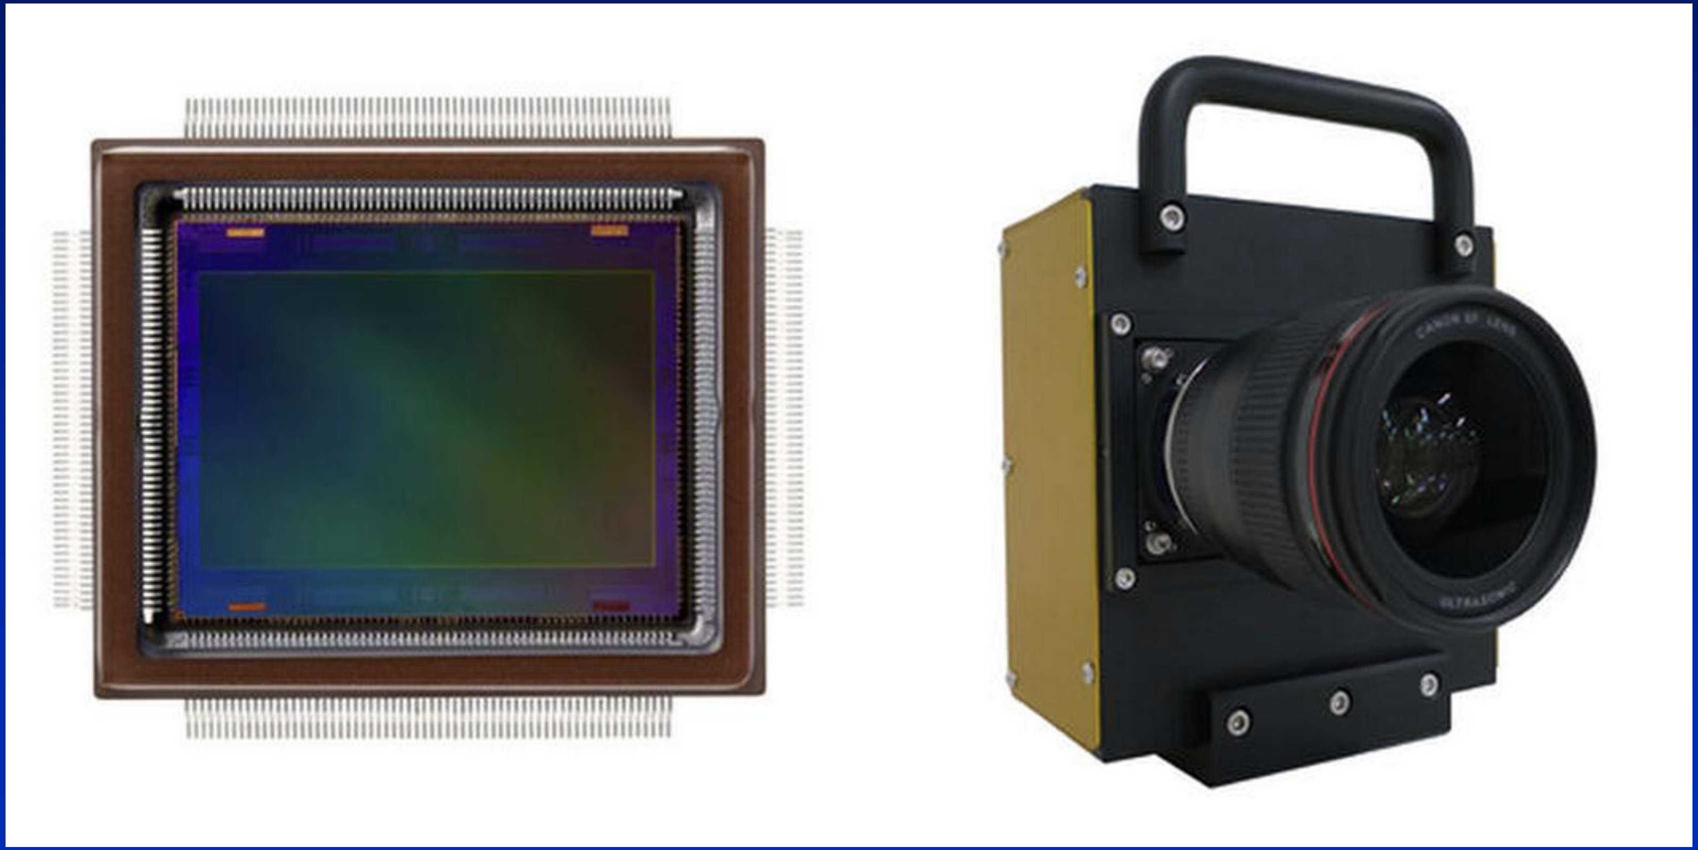

Digital Geometry Capture

- Photographic methods
- Laser scanning
- Pattern projection methods
- Time of Flight

Computer Vision

- The science and technology of machines that see
- Can the machine extract desired information from an image?

What did I see?

Simple case

Known camera positions (x_e, y_e, z_e) , camera optics, known corresponding points each image.

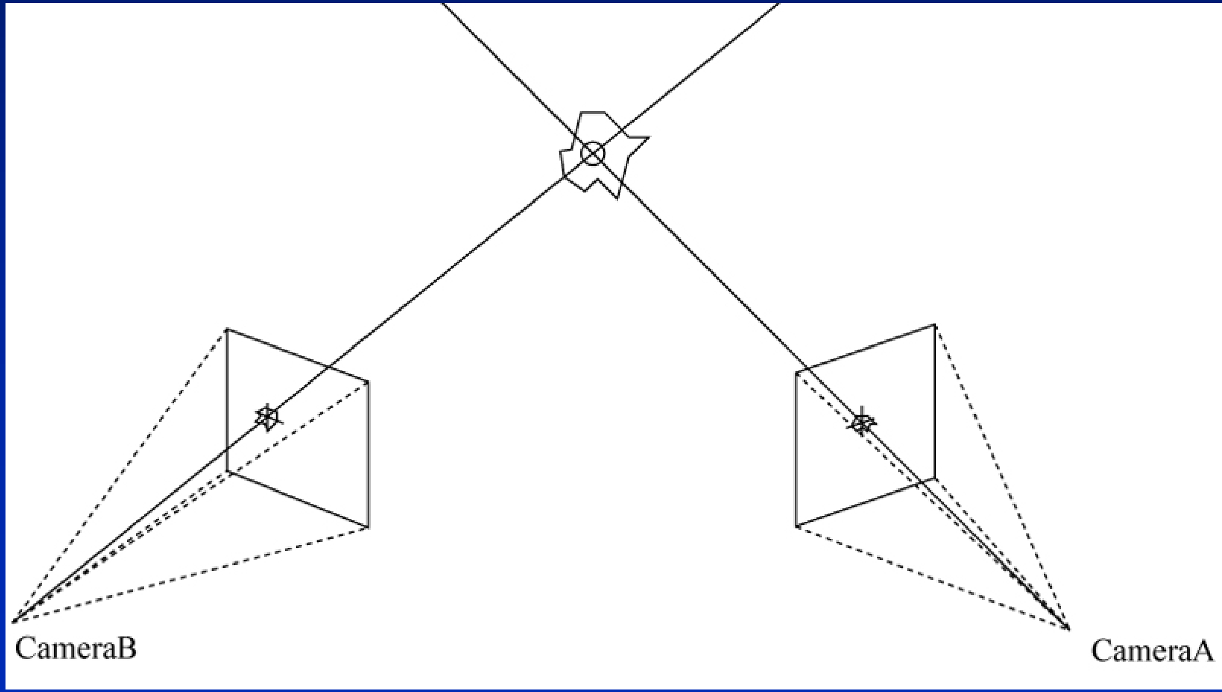

Jeremiah Fairbank. View dependent perspective images. Master's thesis, Cornell University, August 2005.

Early Work - 1975

Sagan House

Sagan House

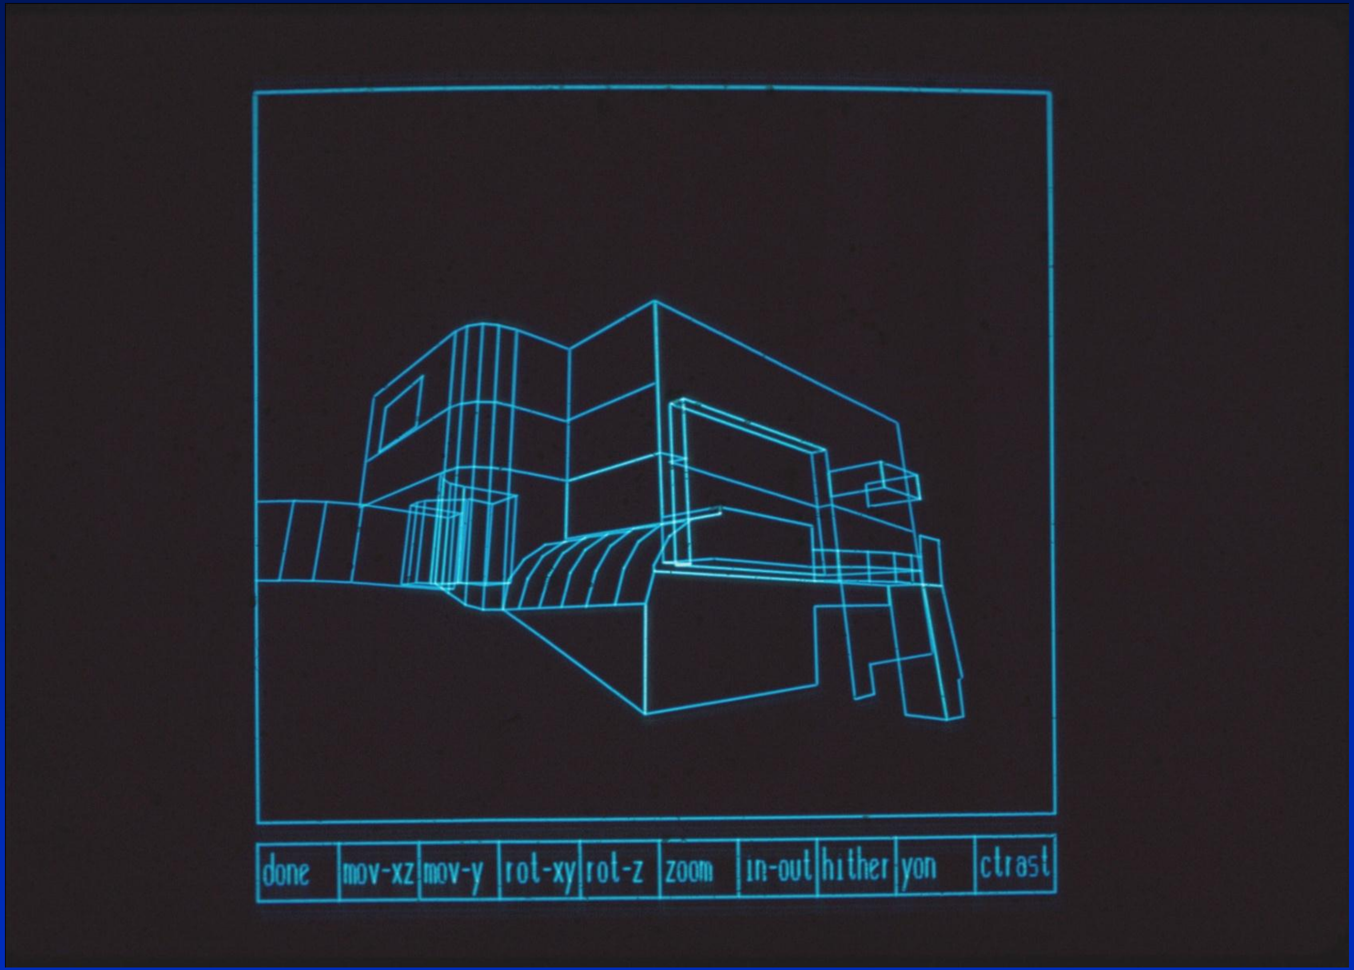

Autodesk 123 Catch

1 2 3 Catch Autodesk

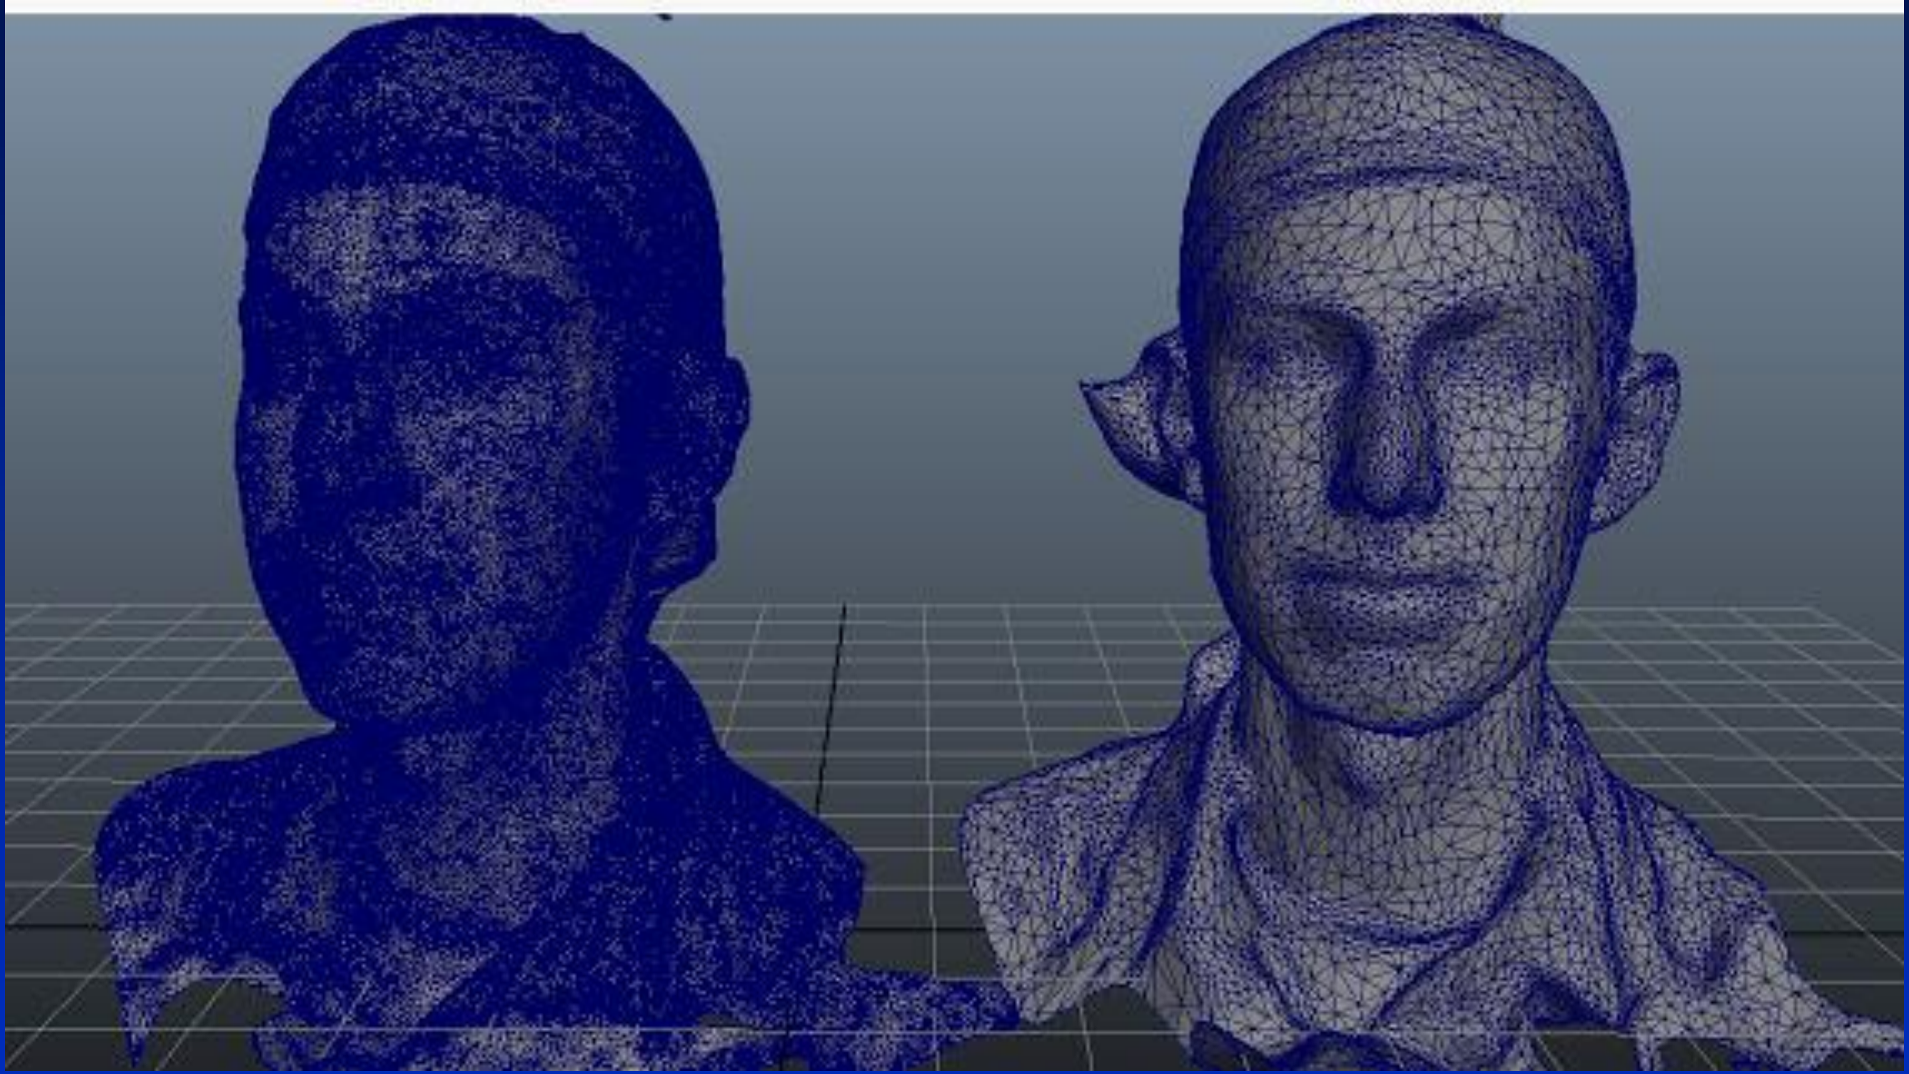

1-2-3D Catch Model from Visual Imaging Course

Credit: Brian Havener

1-2-3D Catch Model from Visual Imaging Course

Credit: Chris Haralampoudis

Autodesk's ReMake 2016

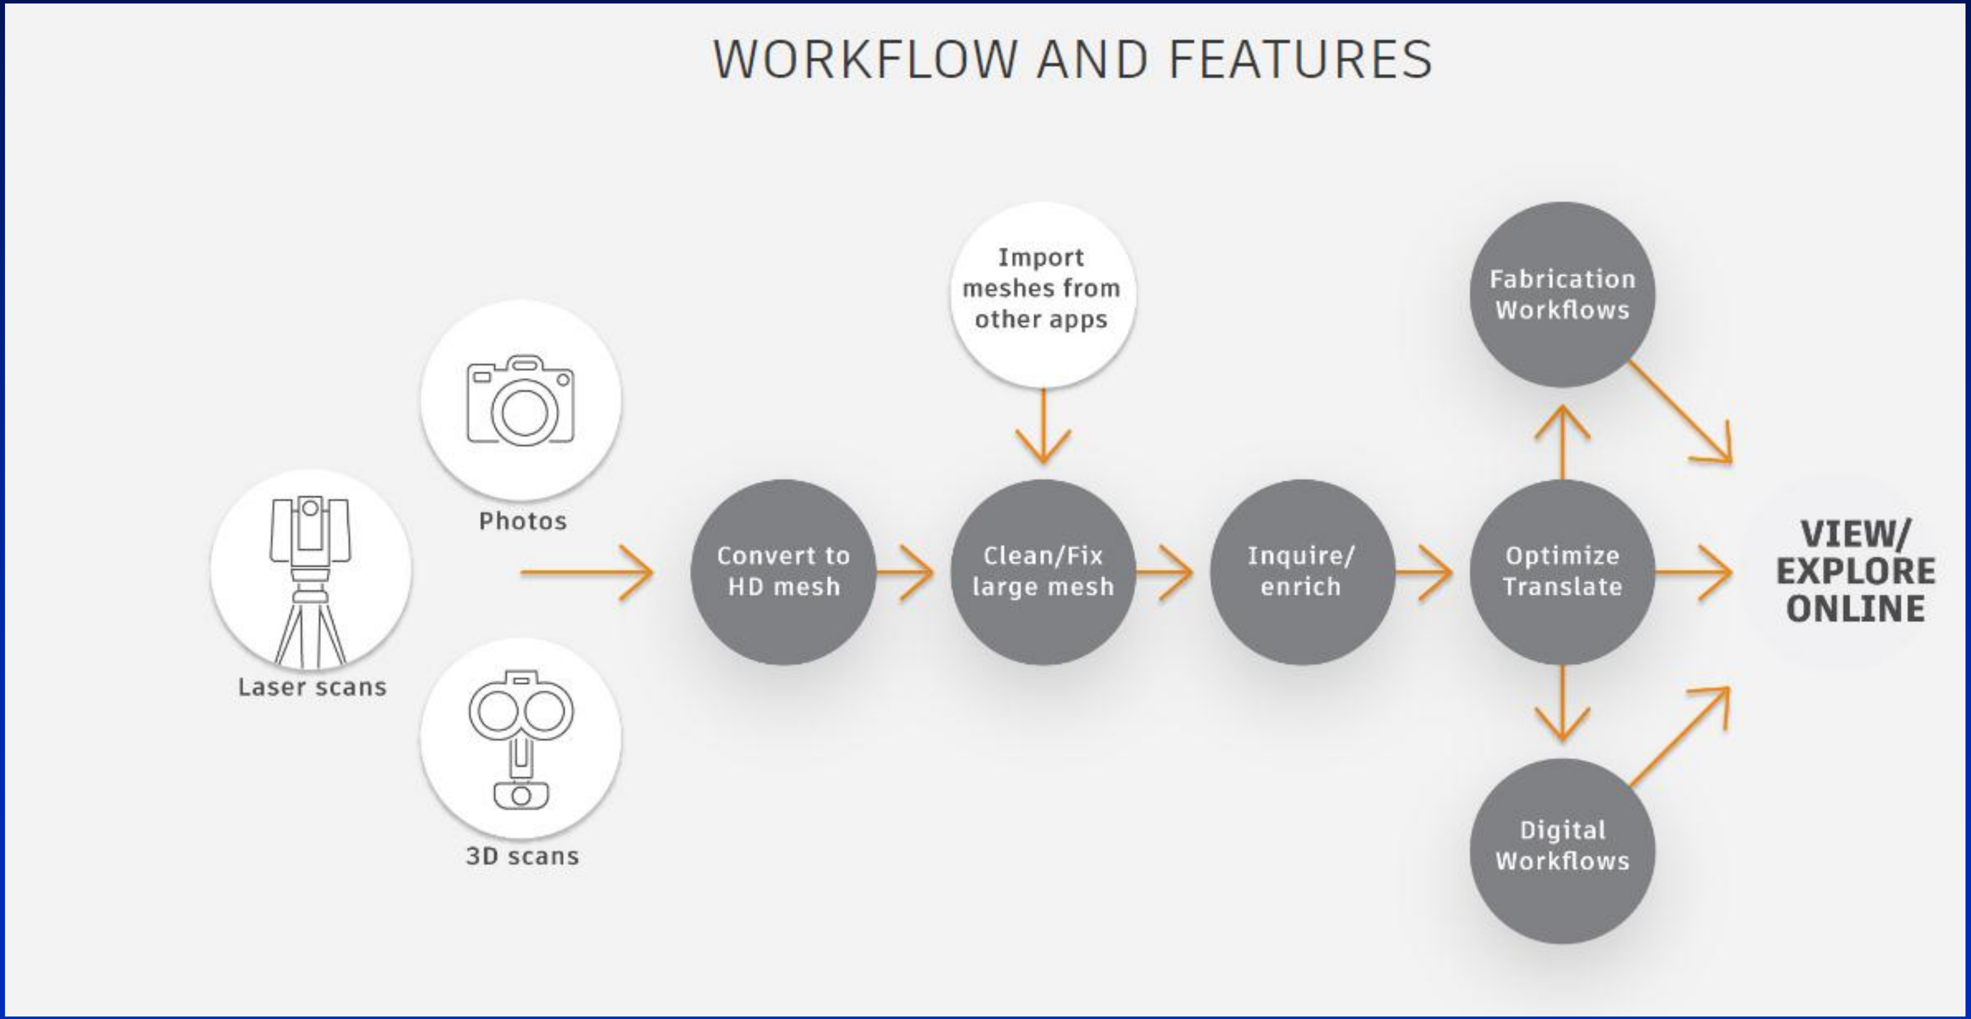

ReMake Model from Visual Imaging Course

Credit: Terrence Vallery

ReMake Model from Visual Imaging Course

Credit: Ashley Yang

Reconstructing Rome[1]

- “The advent of digital photography and the recent growth of photo-sharing websites (flickr) have brought about the seismic change in photography and the use of photo collections.”[1]
- A search for the word “Rome” on flickr returns two million photos.
- This collection, or others like it, capture every popular site, facade, statue, fountain, interior, café, etc.

Characteristics of Typical Photo Sets

- The photos are unstructured
  - No particular order or distribution of camera viewpoints
- The photos are uncalibrated
  - Nothing is known about the camera settings (exposure, focal length, etc.)
- The scale is enormous
  - (millions, not thousands of photos)

and

We need to do this fast!

Correspondence and 3D Structure from Different Camera Positions

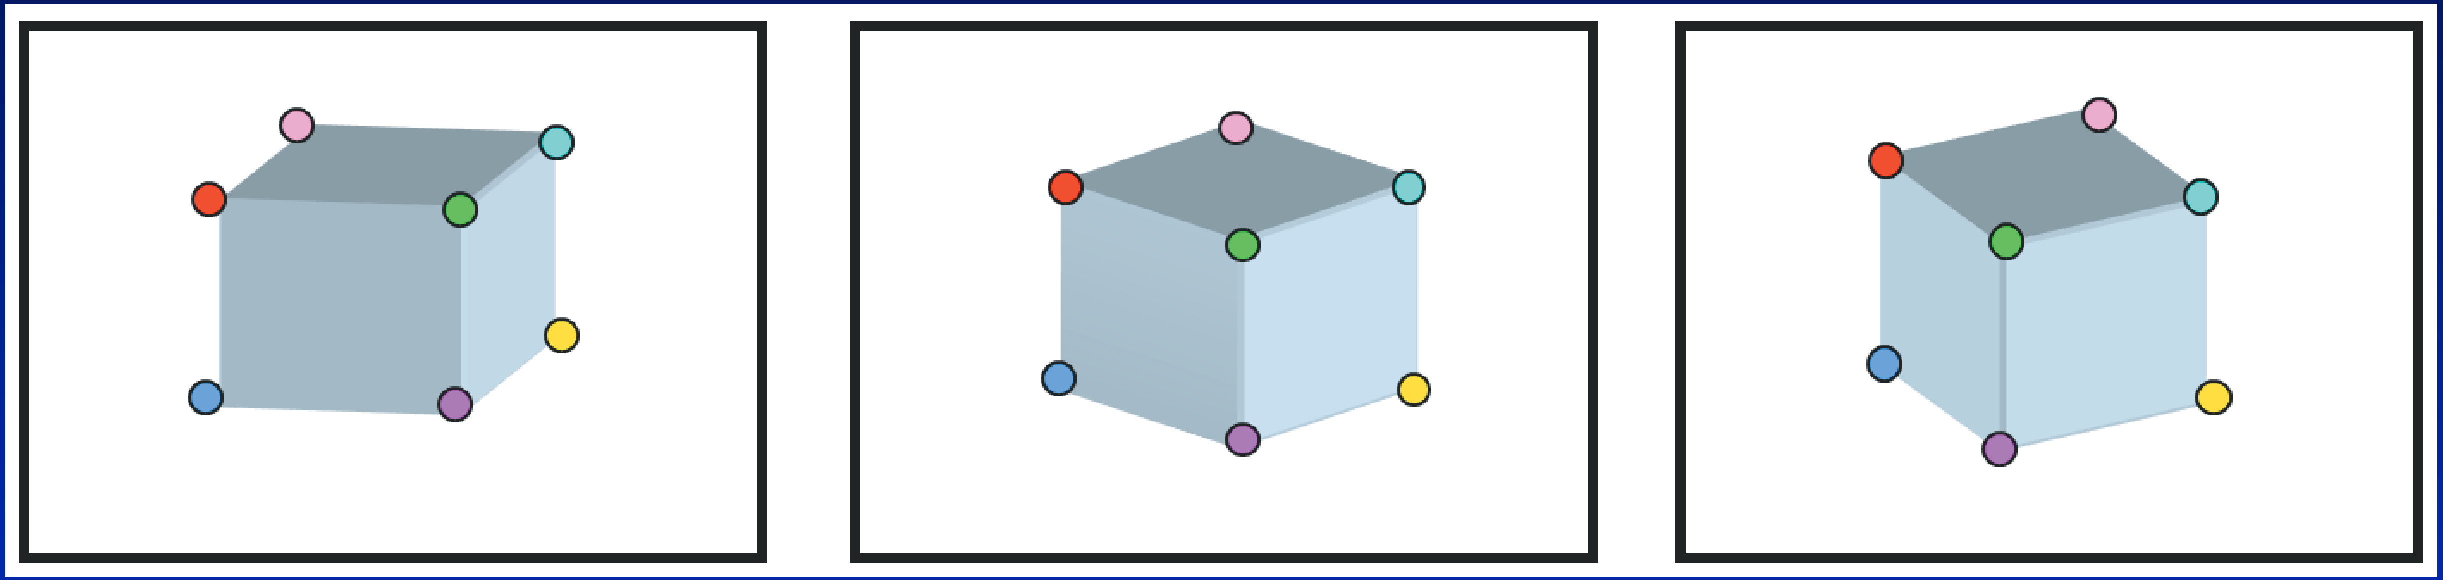

Note: The pictures are in correspondence

2D dots with same color correspond to the same 3D points.

3D Structure from Different Camera Positions

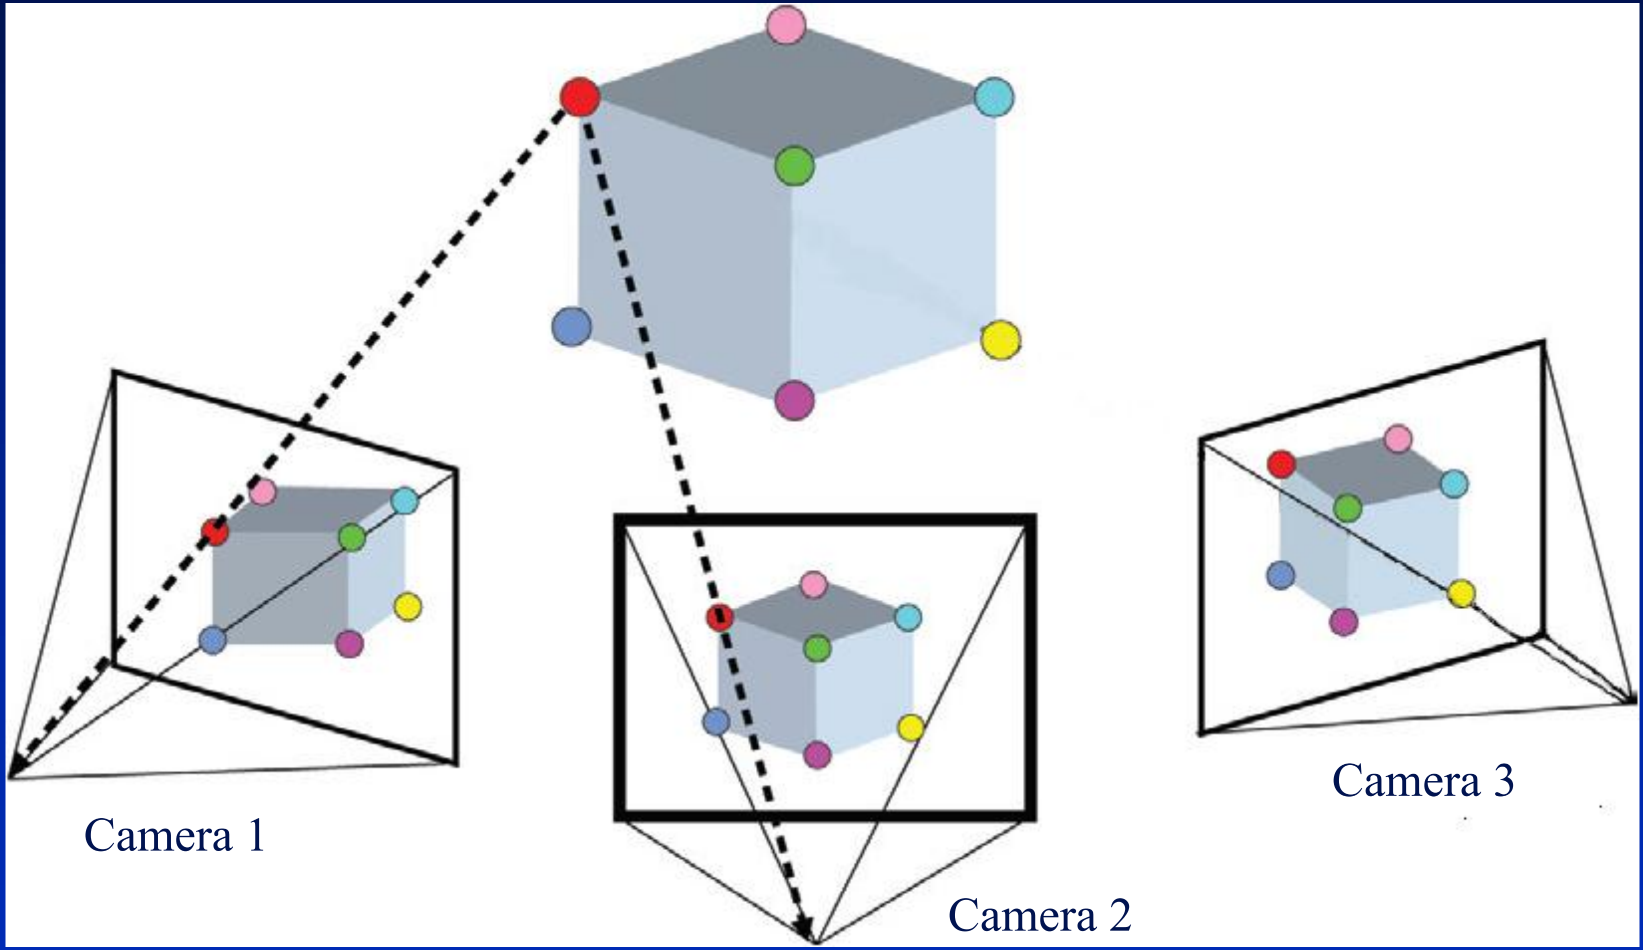

3D Structure from Different Camera Positions

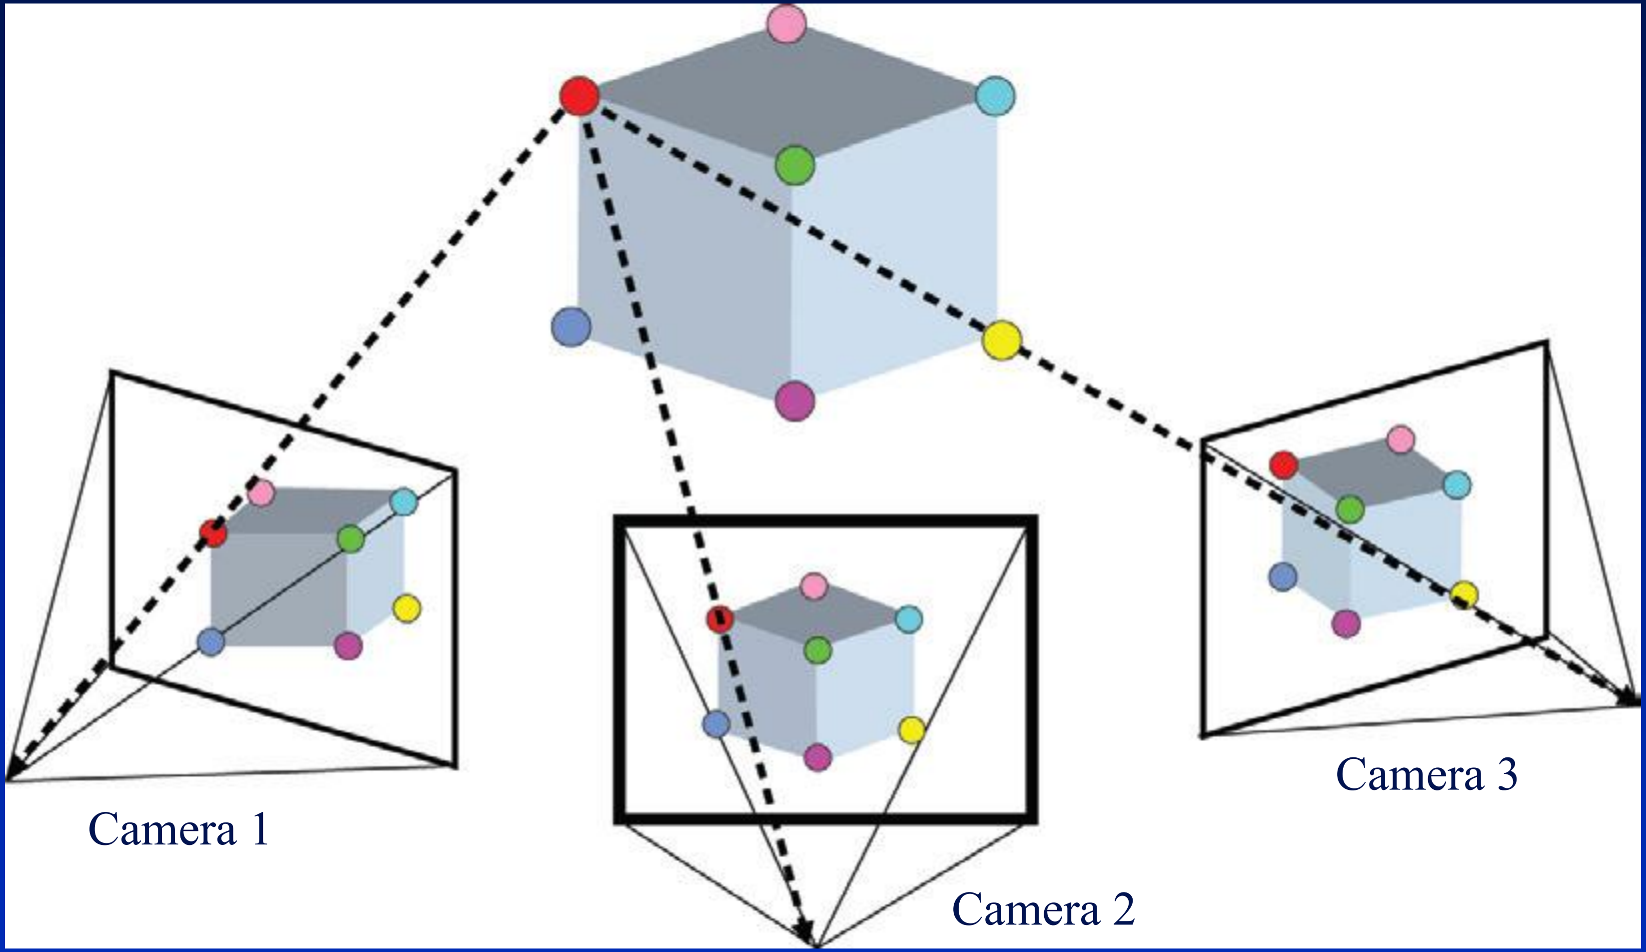

Assuming the position of the red dot is known, there is reprojection error in Camera 3.

Change the Problem to an optimization problem

- Minimize the sum of the squares of the reprojection errors.
- This non-linear least squares problem is difficult to solve due to local minima and maxima.

Trevi Fountain, Rome Italy

http://en.wikipedia.org/wiki/Trevi_Fountain

Feature Detection and Matching

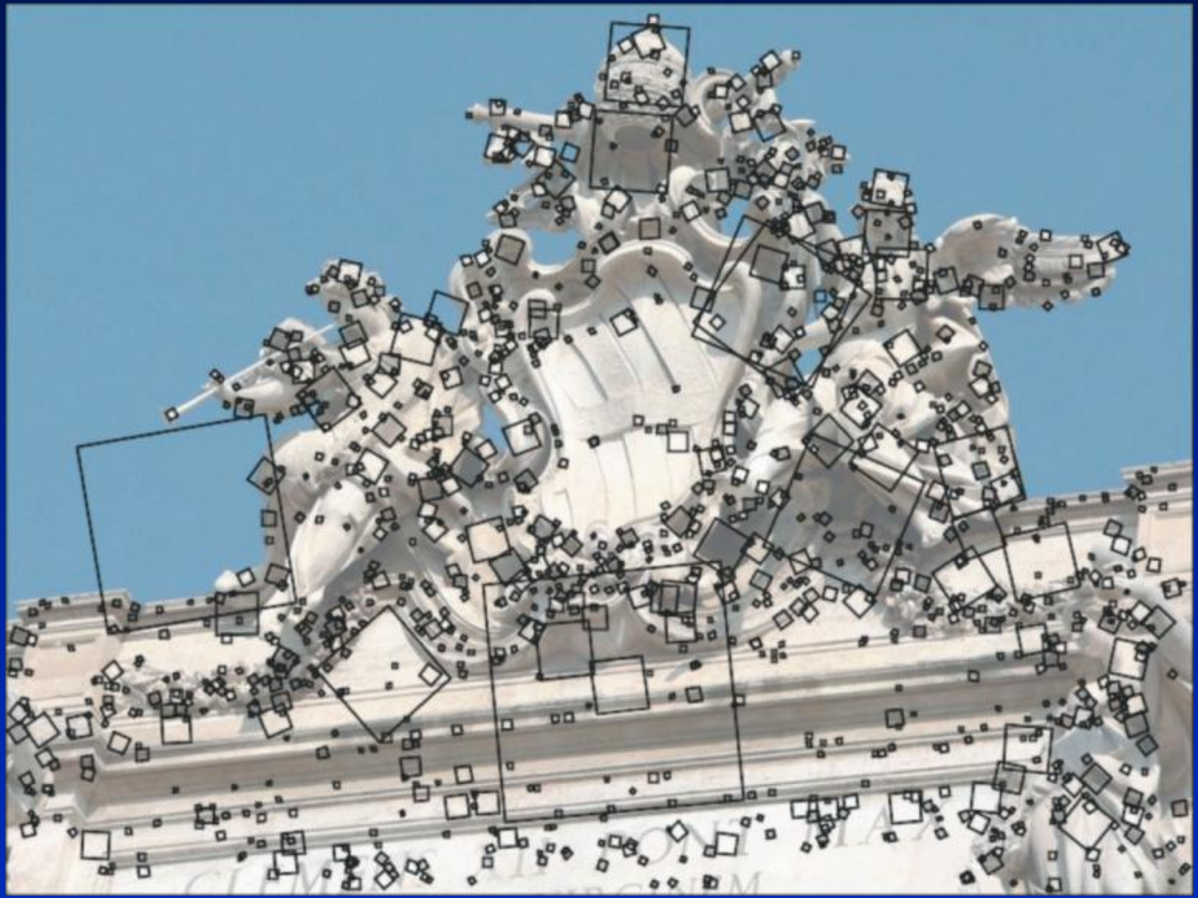

The position and orientation of scale-invariant feature transform (SIFT) features on an image of the Trevi Fountain.

Sameer Agarwal, Yasutaka Furukawa, Naoh Snavely, Brian Curless, Steve M. Seitz, Richard Szeliski. “Reconstructing Rome”, IEEE Computer, June 2010.

Feature Detection and Matching

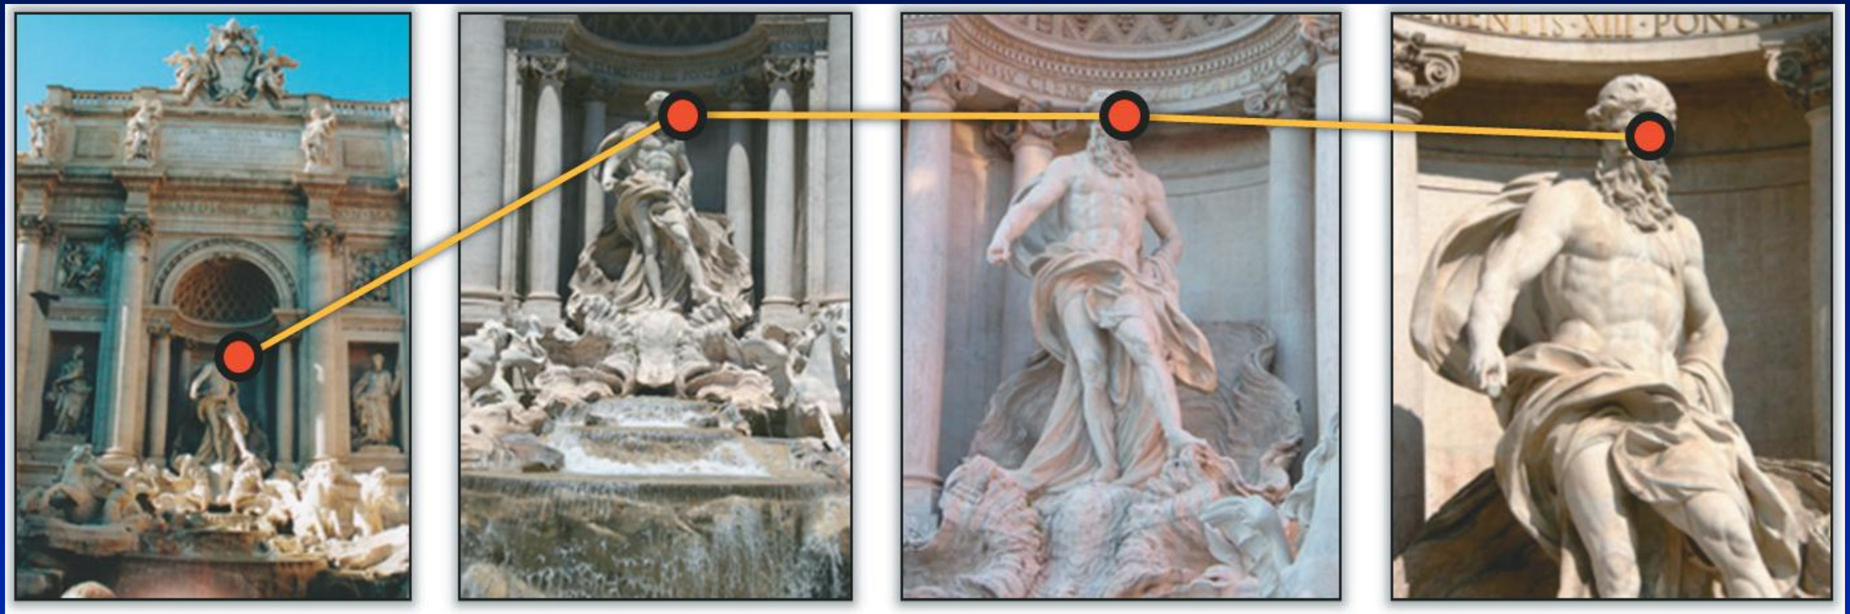

A track corresponding to a point on the face of the central statue of Oceanus at the Trevi Fountain, the embodiment of a river encircling the world in Greek mythology.

Sameer Agarwal, Yasutaka Furukawa, Naoh Snavely, Brian Curless, Steve M. Seitz, Richard Szeliski. “Reconstructing Rome”, IEEE Computer, June 2010.

Colosseum

The Colosseum (Rome)

Reconstructed dense 3D point models. For places with many available images, reconstruction quality is very high.

Sameer Agarwal, Yasutaka Furukawa, Naoh Snavely, Brian Curless, Steve M. Seitz, Richard Szeliski. “Reconstructing Rome”, IEEE Computer, June 2010.

Cornell Campus, McGraw Hall - Noah Snavely

Cyberware Scanner Diagram

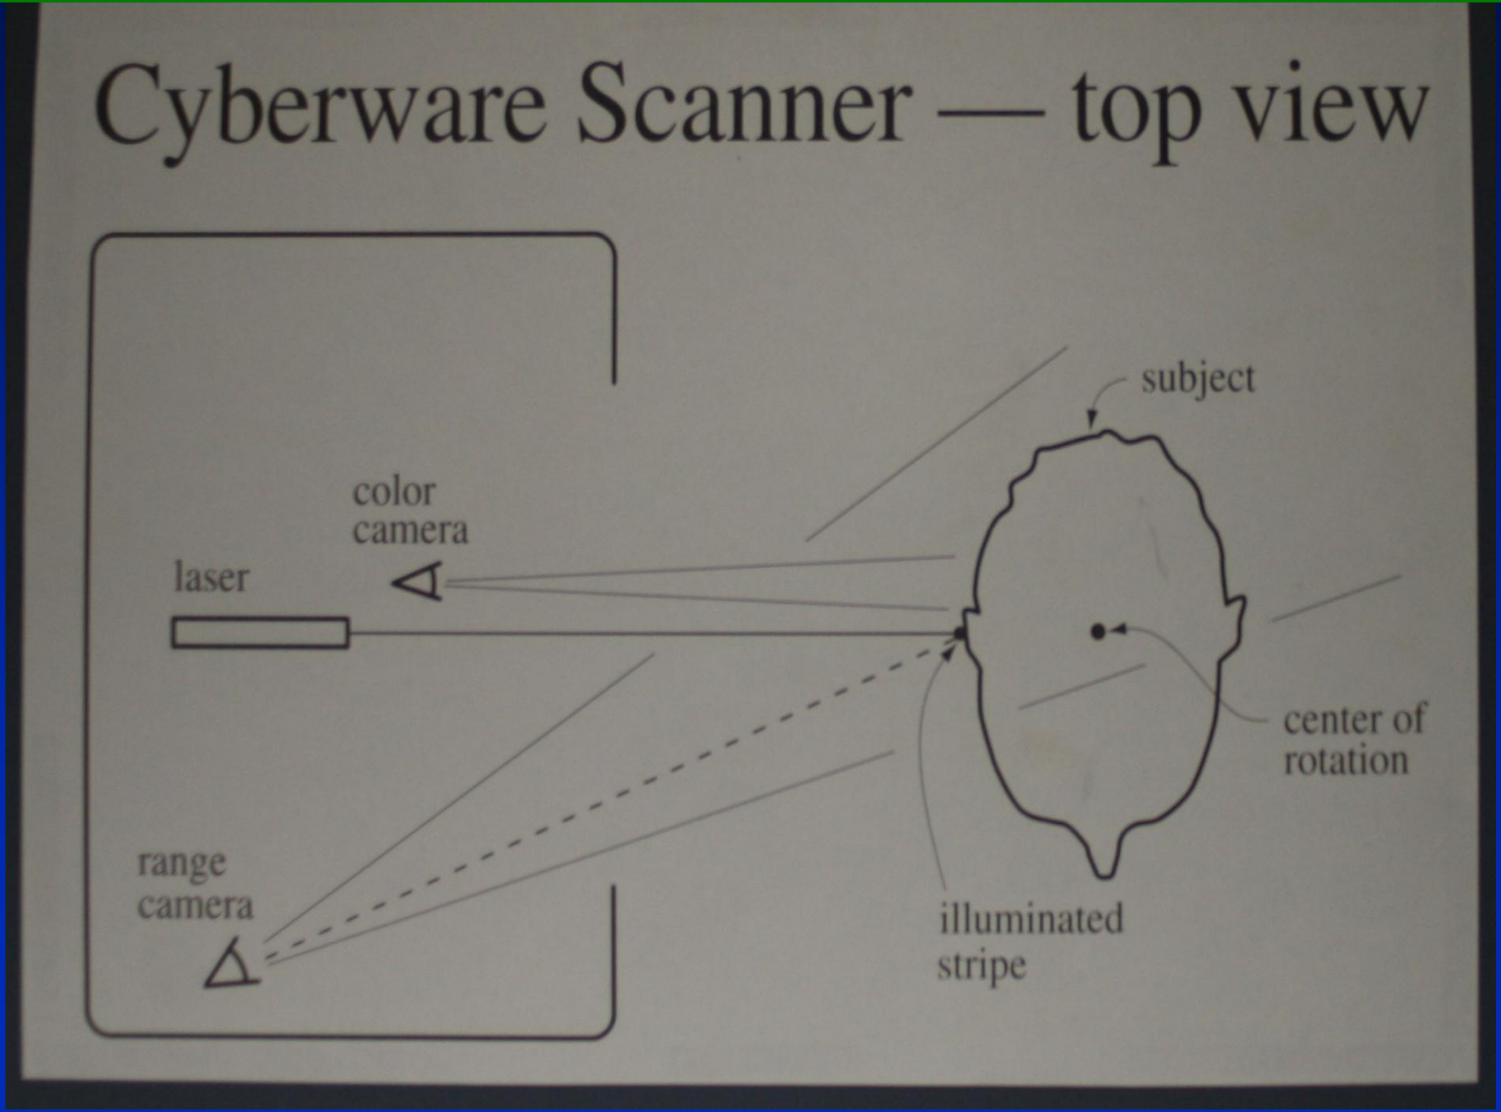

Cyberware Scanner

Uncle Don

Microsoft's Kinect

Microsoft’s Kinect

Making Things See, Greg Borenstein

The Kinect uses a pattern projection and machine learning

- Inferring body position is a two-stage process: First Compute a depth map (using projected pattern), then infer body position (using machine learning)

Structured Light Imaging (Kinect)

- Kinect uses a spatial pseudo random neighborhood pattern with unique coding with different sized dots.

Kinect: Depth Image and Real Image

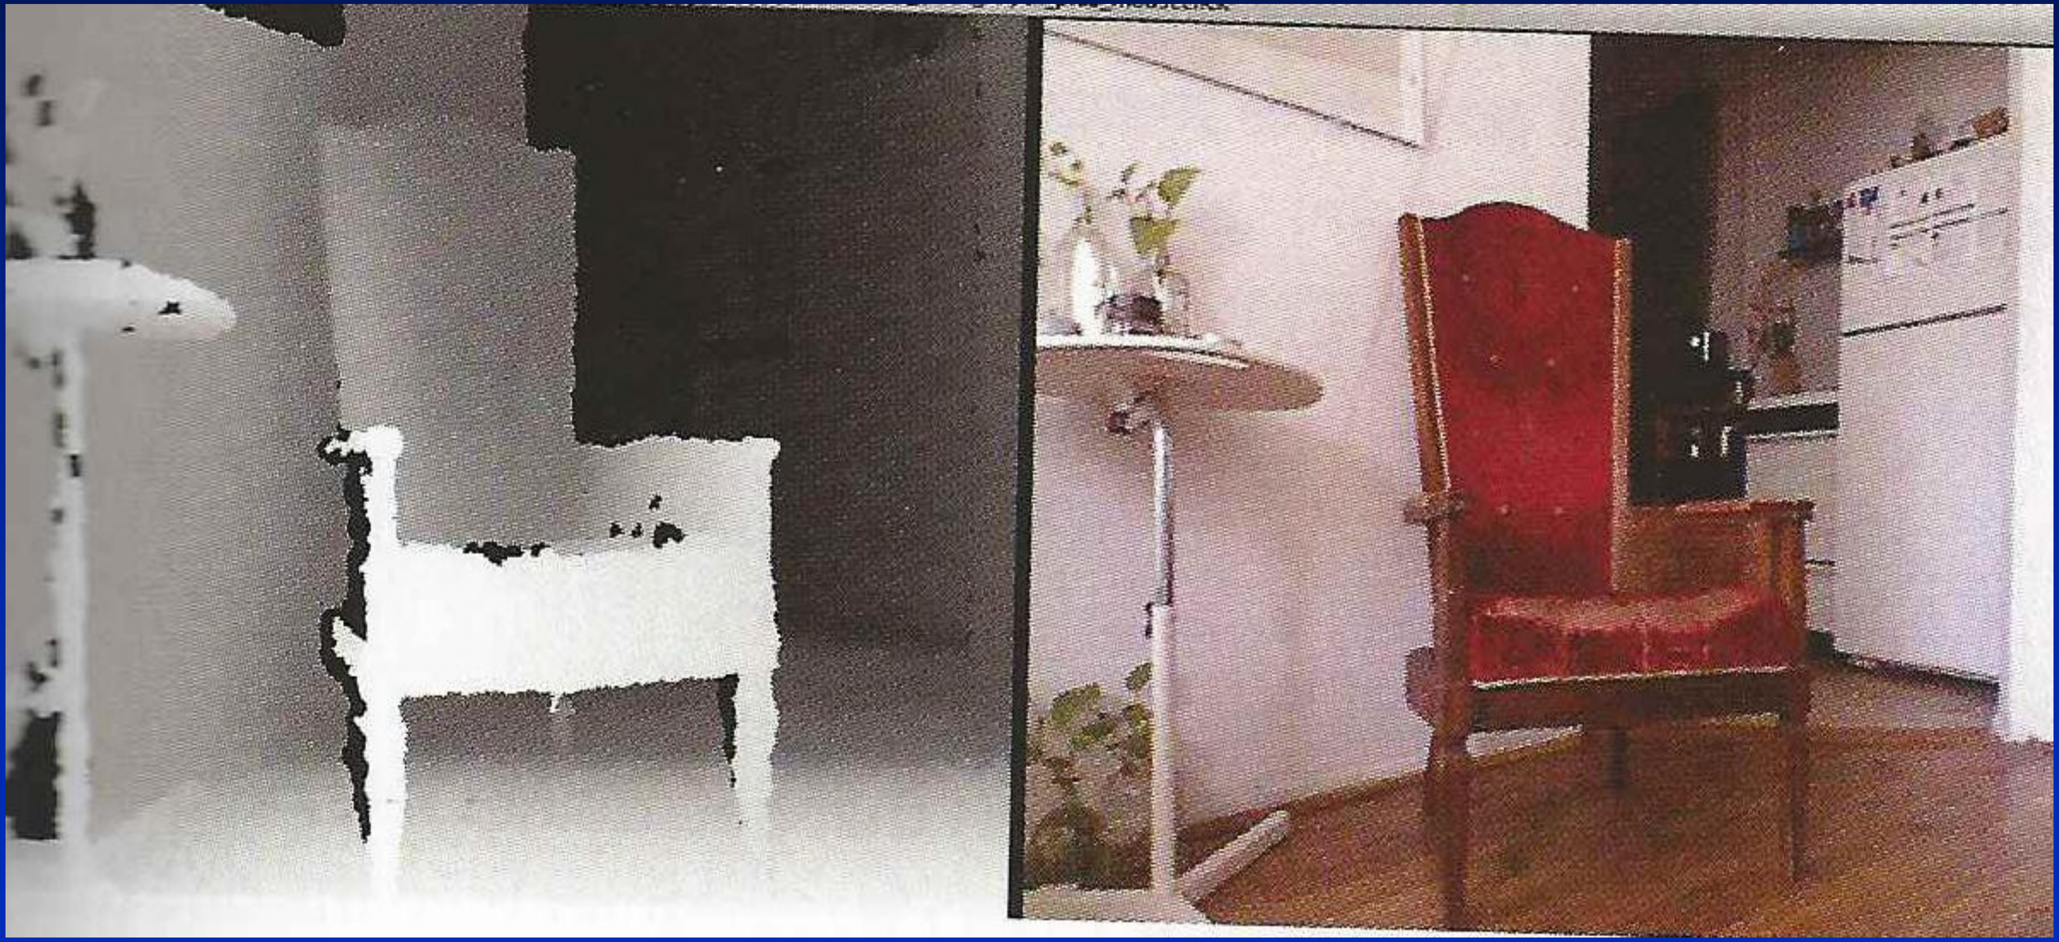

Step 1: Compute a Depth Map

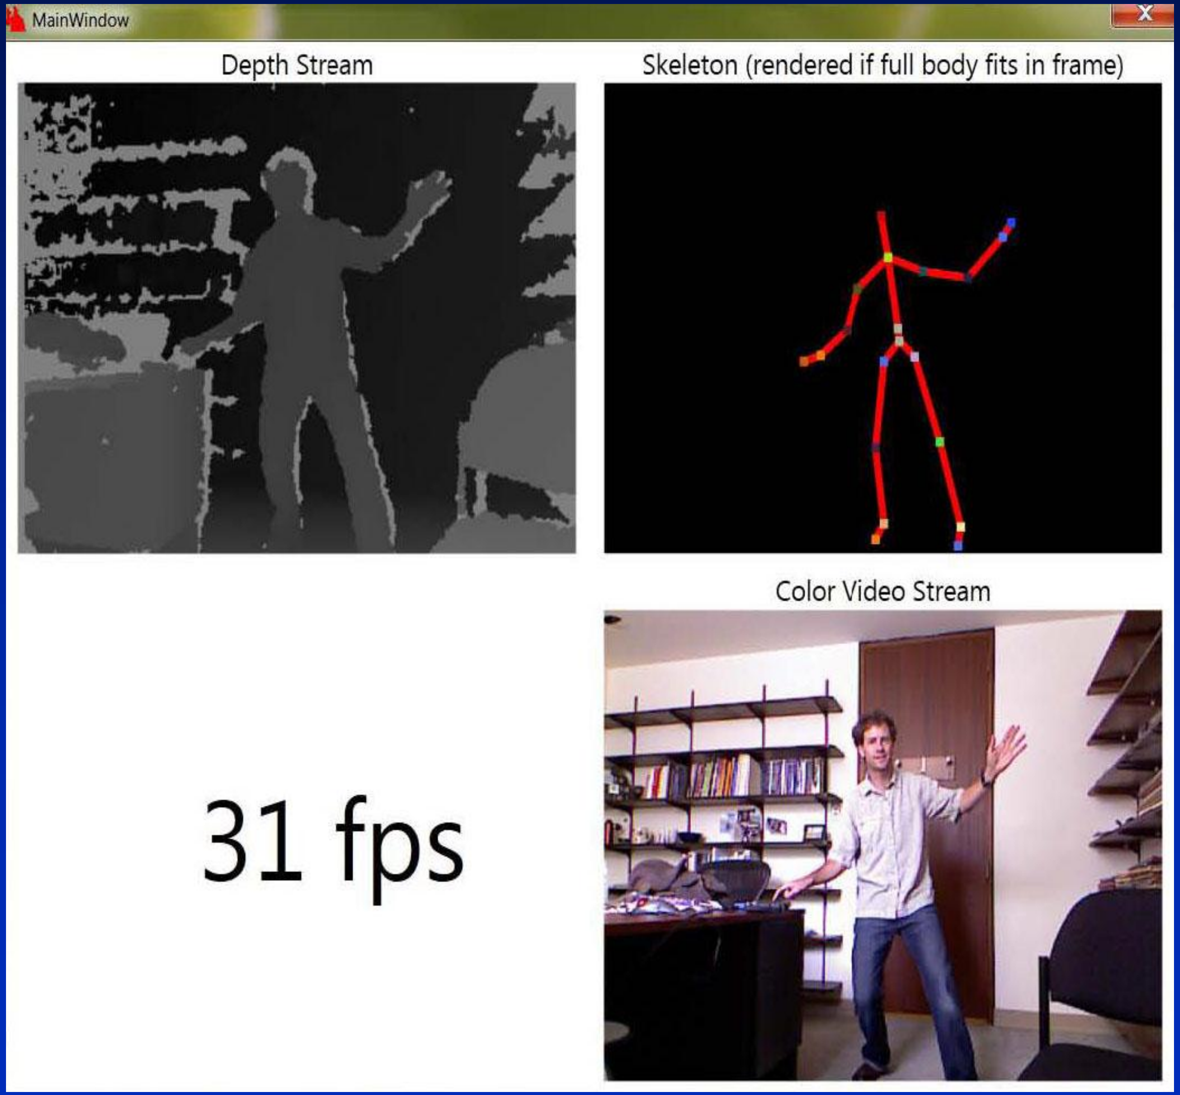

Step 2: Infer a Body Position

Skeleton Manipulation

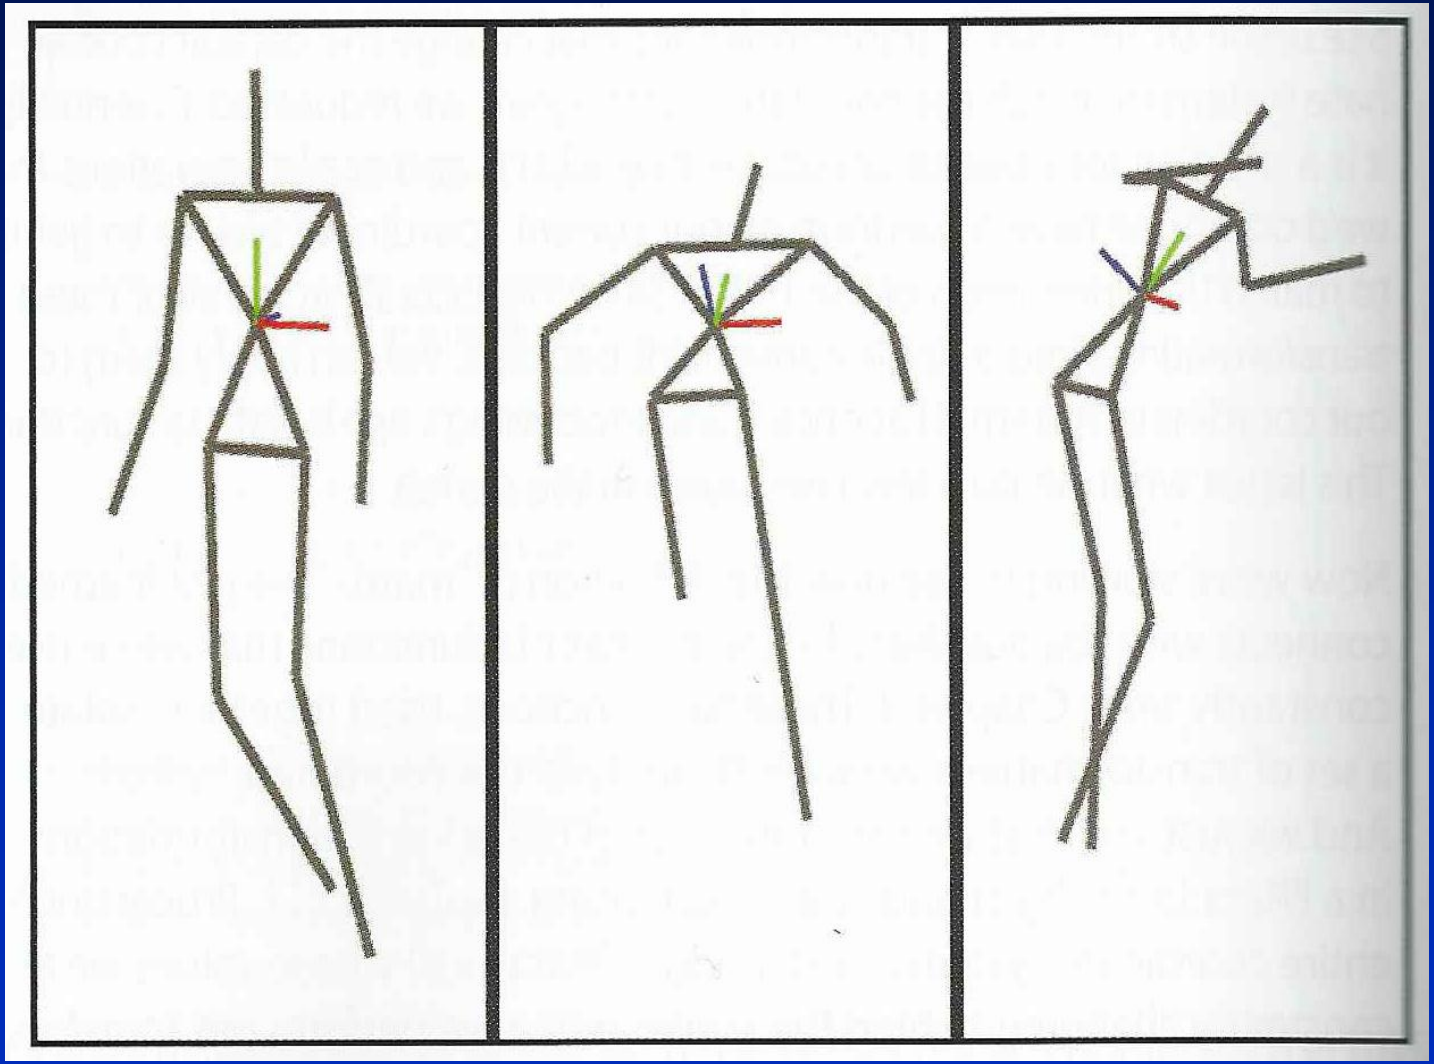

Extracted Skeleton Kinect

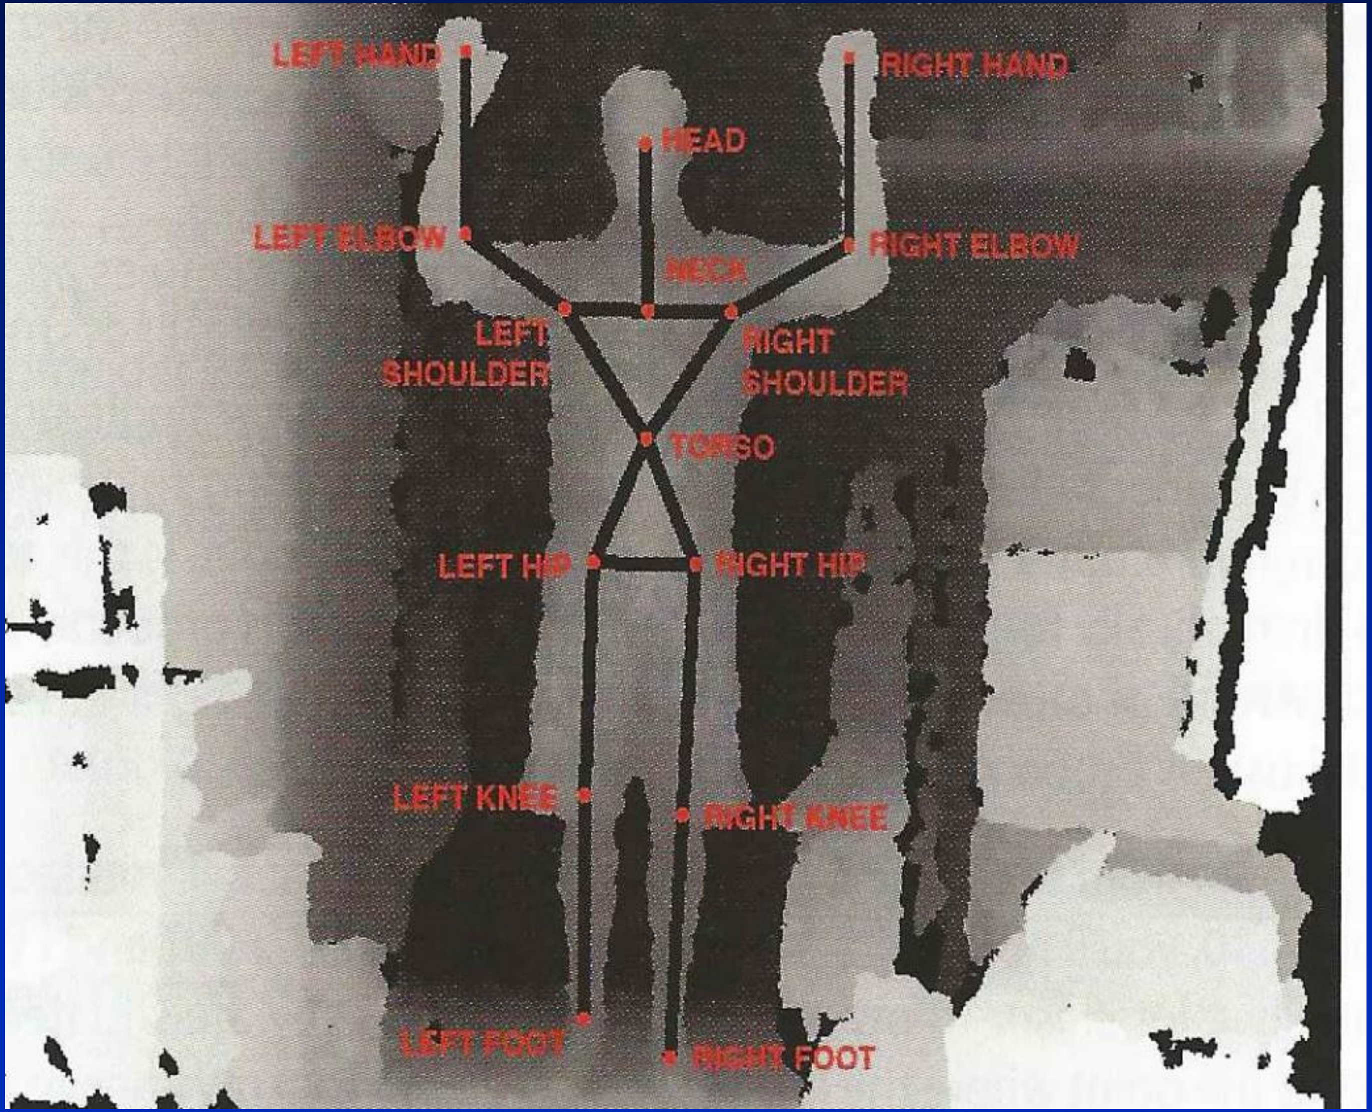

Tracking

Pulsed Modulation

- Measure distance to a 3D object by measuring the absolute time a light pulse needs to travel from a source into the 3D scene and back, after reflection
- Speed of light is constant and known, $c = 3 \cdot 10^8$ m/s

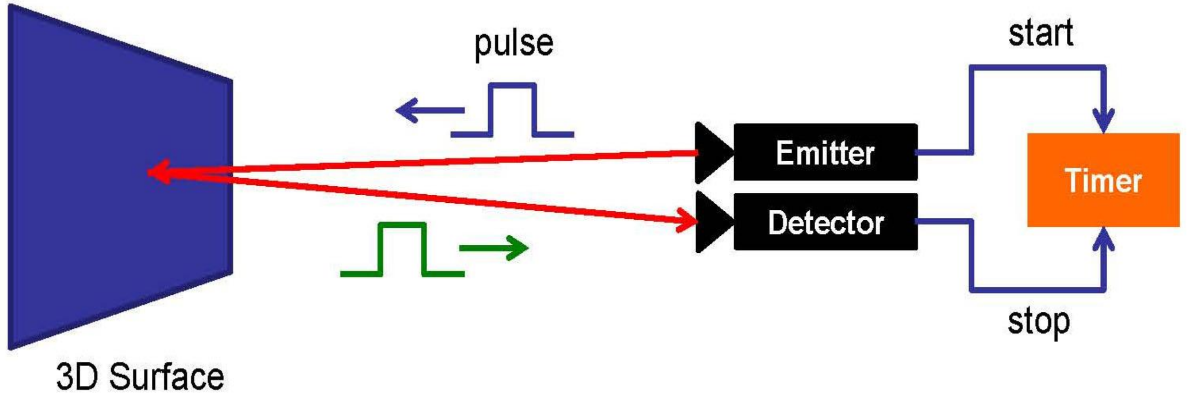

Kinect 2

Kinect 2


Kinect For Windows 1
Processed Image From Kinect
Kinect For Windows 2
Image via http://blogs.msdn.com

Floored

- http://labs.floored.com/clients/mrp-realty/900-g-roof/

Matterport

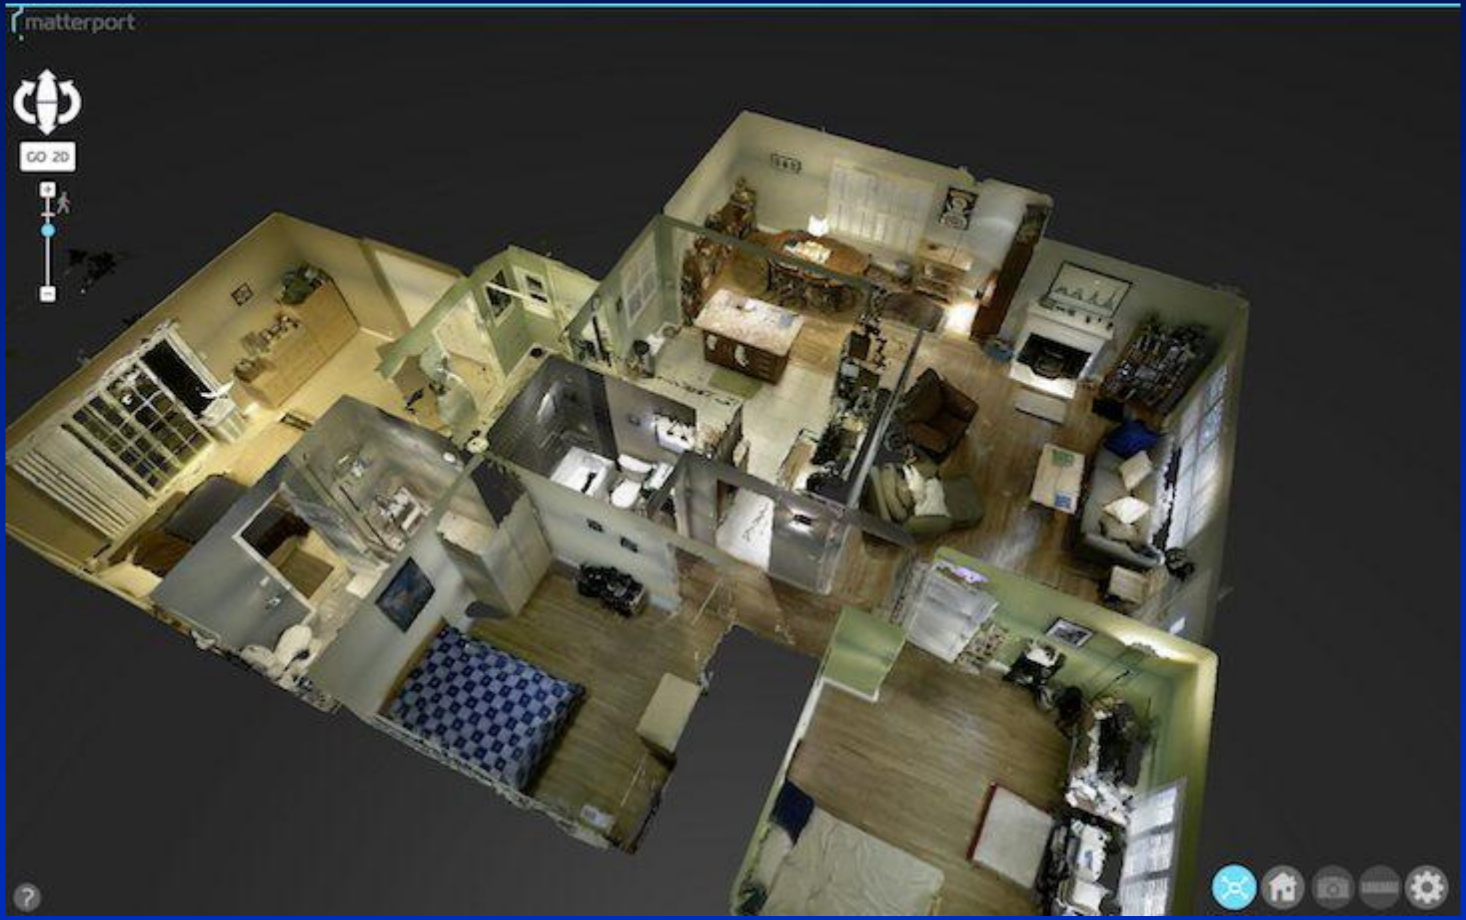

Matterport

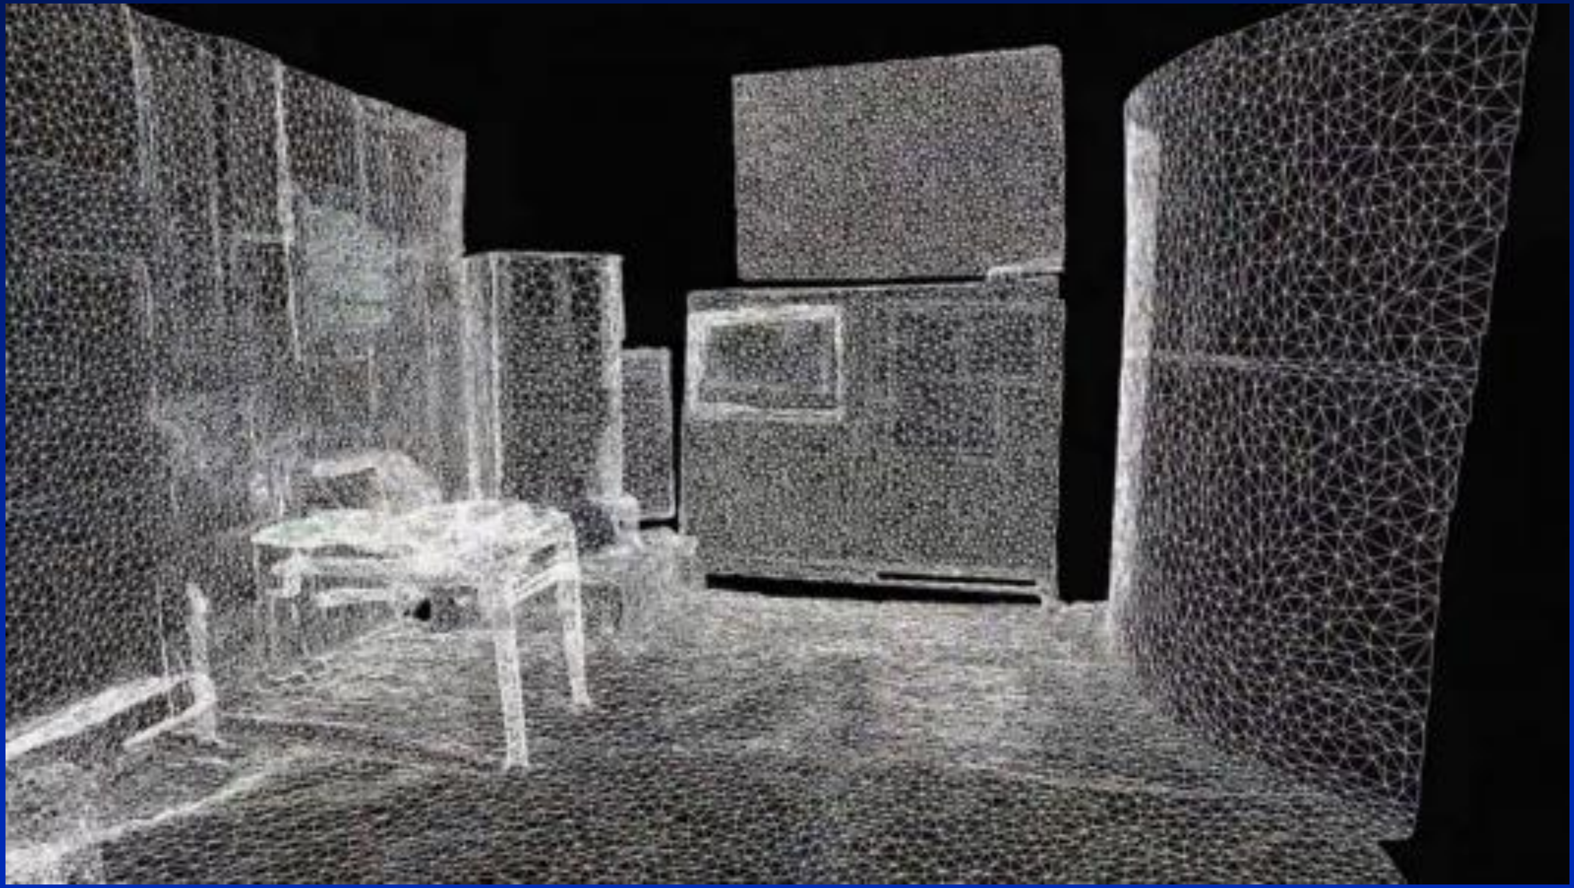

Matterport

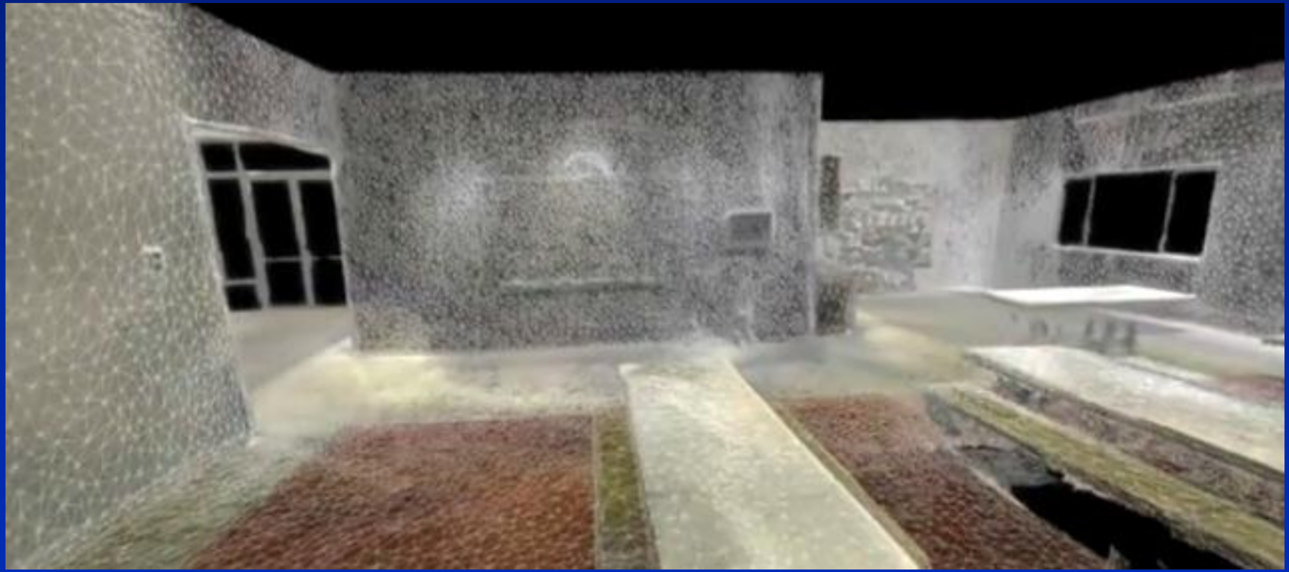

Leica and Autodesk

2017- Not Yet Available

- Calibrated Full spherical image, HDR, LED Flash Support.
- 360,000 Laser Scan PTS/Sec. with selectable resolution settings.
- Less than 3 minutes for full 360° capture

Time of Flight Point Cloud

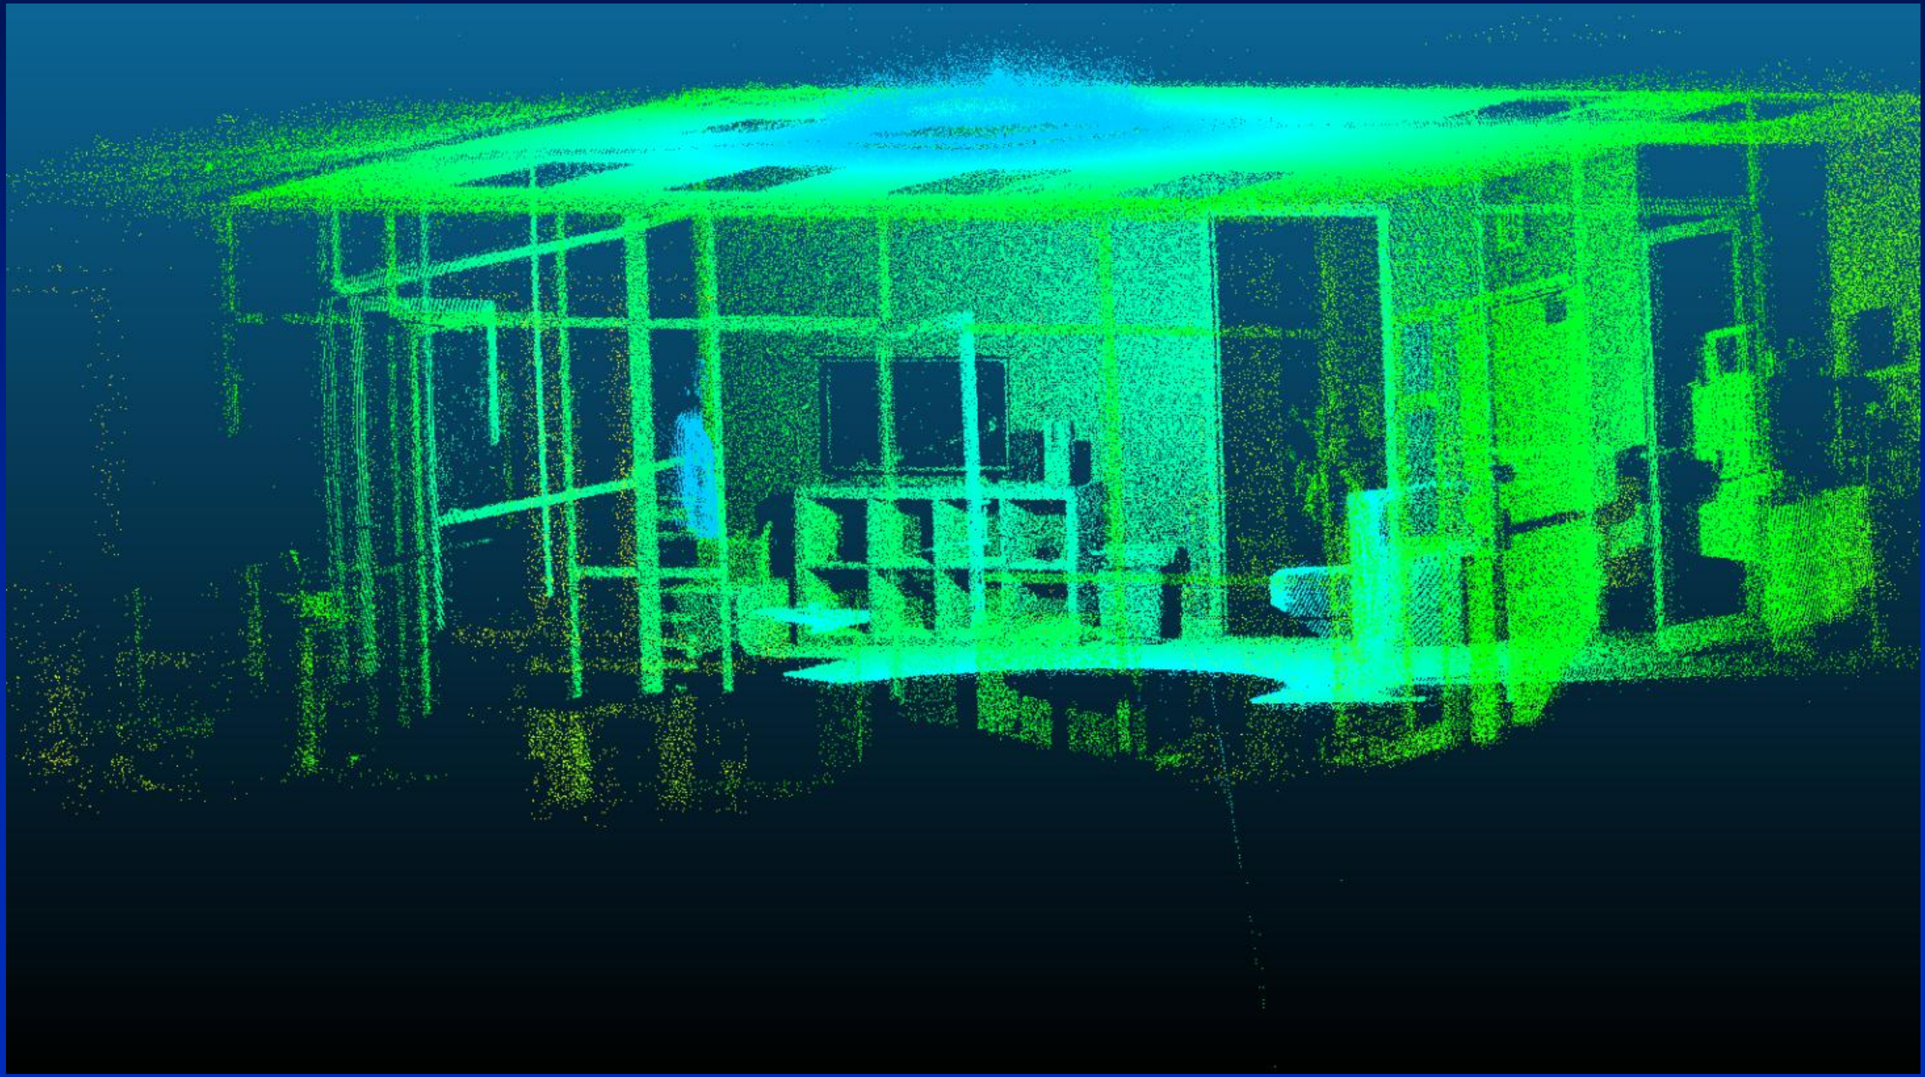

Digital Geometry Capture

- Photographic methods
- Time of Flight
- Radar
- Sonar
- All of the Above

Google Street View and Google Maps 2007-2012

- In 2007, Larry Page requests Thrun and Levandowski to create a virtual map of the U.S.
- Engineers jury-rigged some vans with GPS and rooftop cameras which shot 360° panoramas for any address. They equipped 100 cars which were sent around the U.S.
- Data was put together with a program written by Marc Levoy.
- In 2011, Google announced it would start charging (large) commercial sites
- In 2012, Google allows users to post photos and reviews of locations

By October 2012, Google will have updated 250,000 miles of U.S. roads

Note: They have also added Google Moon and Google Mars

R7 Street View Camera System - 2009

The system is a rosette ® of 15 small, outward-looking cameras using 5-megapixel CMOS image sensors and custom, low-flare, controlled-distortion lenses.

Drafomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stepheane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, Josh Weaver. “Google Street View: Capturing The World At Street Level,” IEEE Computer, June 2010.

Street View Vehicular Platforms

Second-(right) and Third- (left)

Drafomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stepheane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, Josh Weaver. “Google Street View: Capturing The World At Street Level,” IEEE Computer, June 2010.

Google Street View Car Fleet

Google Street View Acquisition Map 2012

Google Street View

- The world contains roughly 50 million miles of roads, paved and unpaved, across 219 countries (ref.)
- This is equivalent to circumnavigating the globe 1250 times.
- To date, hundreds of cities in many countries across four continents have been captured.
- Google has developed several vehicular platforms and texture information in the project's seven year history.

Drafomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stepheane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, Josh Weaver. "Google Street View: Capturing The World At Street Level," IEEE Computer, June 2010.

Imagery from new Street View Vehicle is accompanied by laser range data

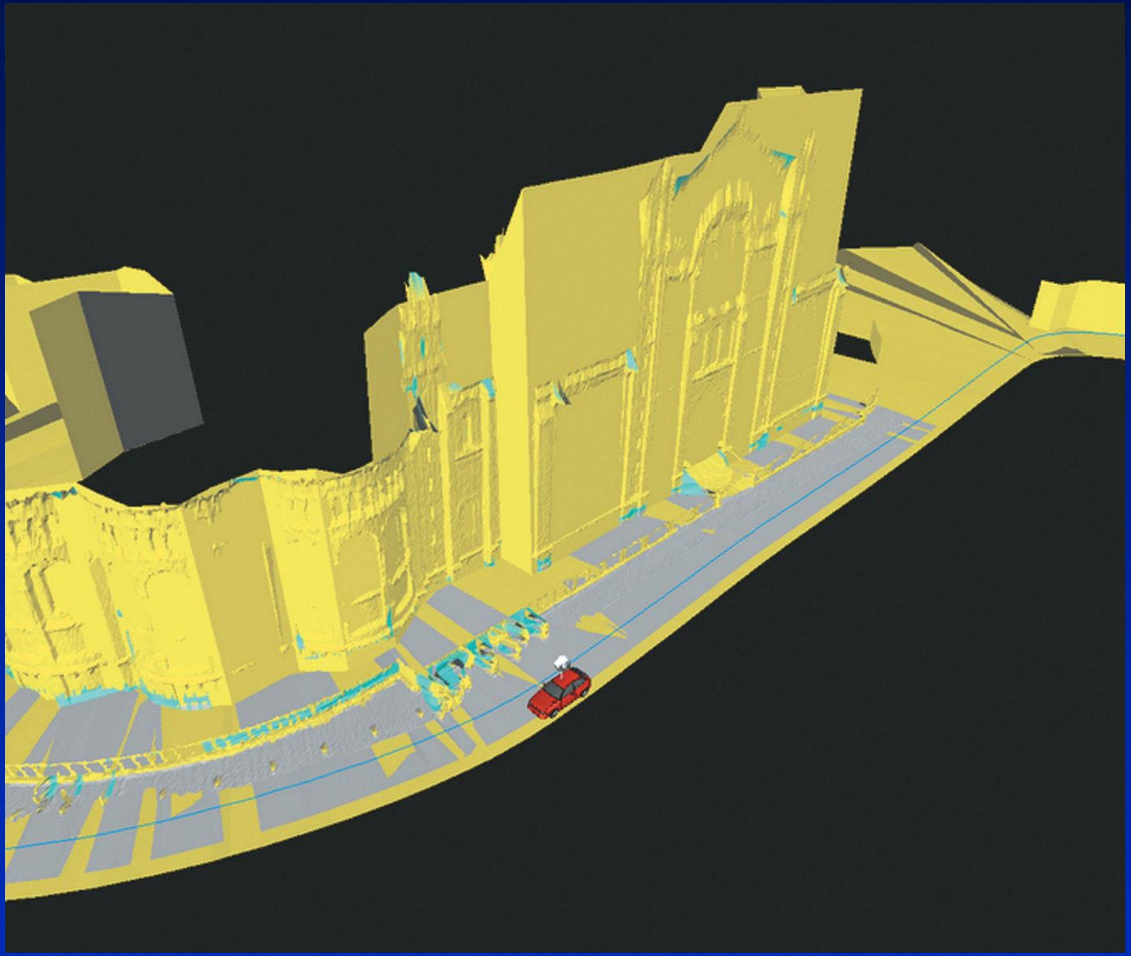

- which is aggregated and simplified by robustly fitting it in a coarse mesh that models the dominant scene surfaces.

Drafomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stepheane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, Josh Weaver. “Google Street View: Capturing The World At Street Level,” IEEE Computer, June 2010.

Using Street View data to enhance user walk-through experiences in Google Earth.

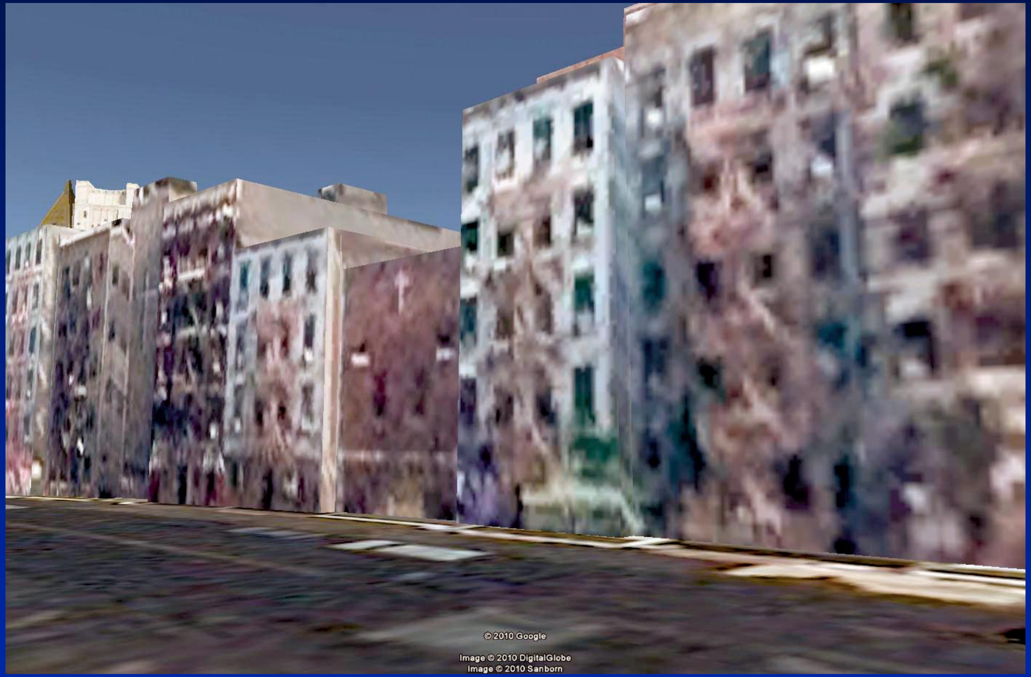

Original 3D models of a New York City scene from airborne data only.

Fused 3D model with high-resolution facades. The visual quality is considerably higher, and many storefronts and signs can now be easily identified and recognized.

Drafomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stepheane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, Josh Weaver. "Google Street View: Capturing The World At Street Level," IEEE Computer, June 2010.

Autonomous Driving Vehicles Pre-2000

- “There was no way, before 2000, to make something interesting”
- “The sensors weren’t there
- The computers weren’t there
- The mapping wasn’t there”
- “Radar was a device on a hilltop that cost $200M”

Sebastian Thrun

Founder of the Google Car Project

Google's Autonomous Driving Vehicle 2013

- Uses multiple sensors, each with a different view of the world
- Laser
  - 64 beams @ 10 revolutions/second scanning 1.3 million points in concentric waves starting 8 feet from the car
  - It can spot a 14" object at a distance of 160 feet
- Radar
  - Has twice the range of the Laser, but much less precision
- Photography
  - Excellent at identifying road signs, turn signals, colors and lights

Google’s Autonomous Driving Vehicle

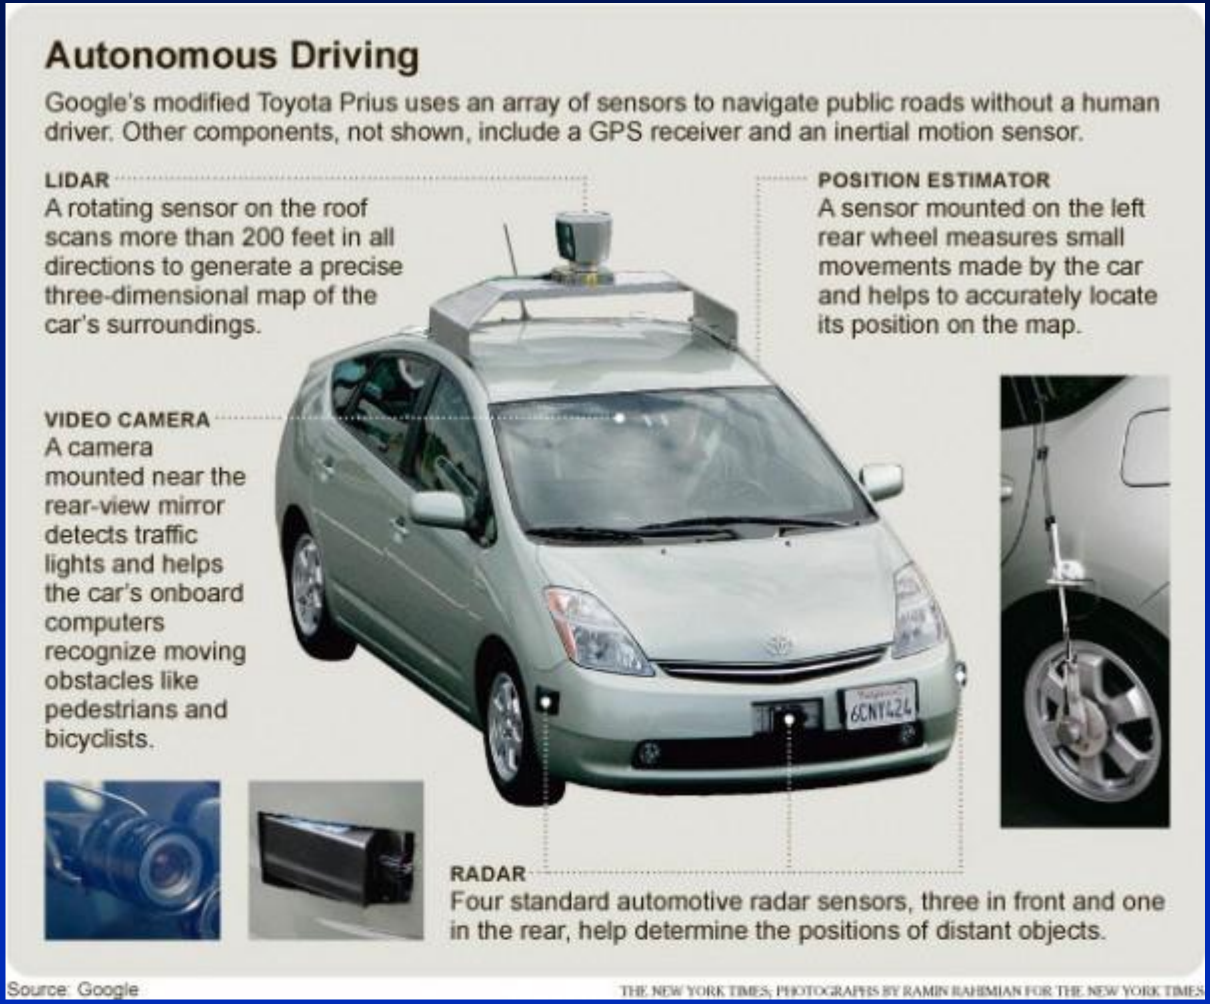

http://erictric.com/2010/10/09/google-secretly-develops-tests-automated-cars/

Lombard Street, San Francisco

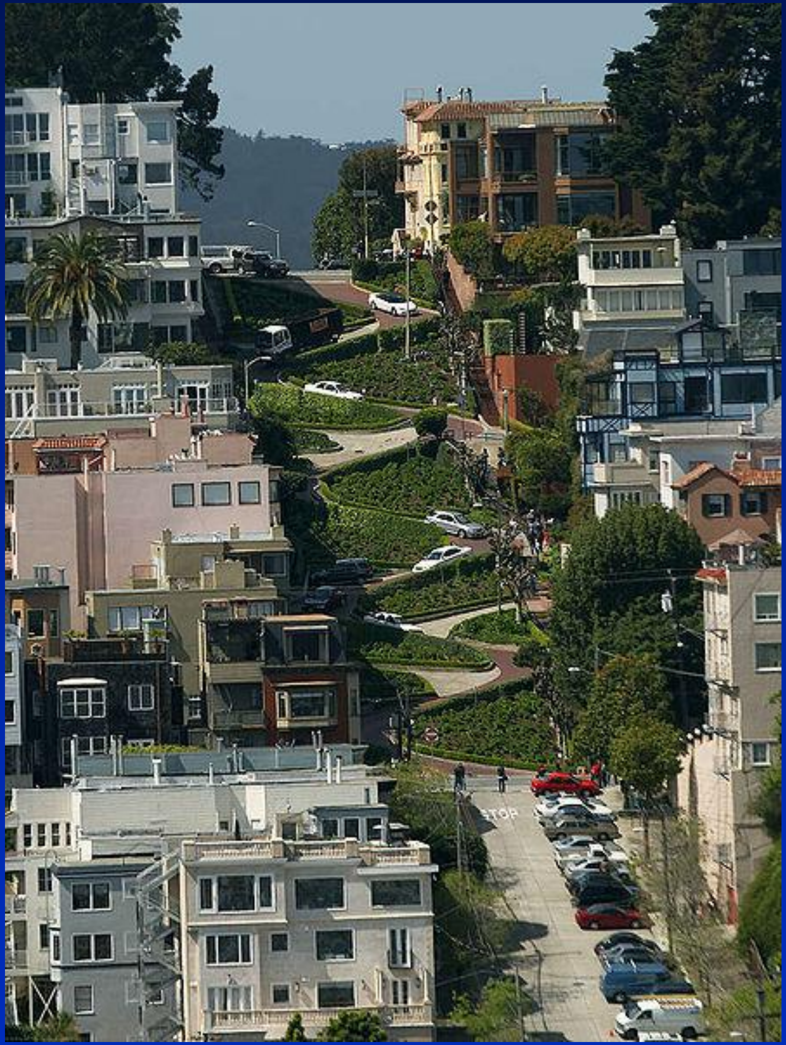

Google’s Autonomous Driving Vehicle 2014-2015

- New laser sensors
  - 2 X range
  - 30 X resolution
  - @ 300’ can spot a metal plate <2” thick
  - Size of a coffee mug
  - Cost ≈ $10,000 (less than current model @ $80,000)

Google's Recording Rig 2015

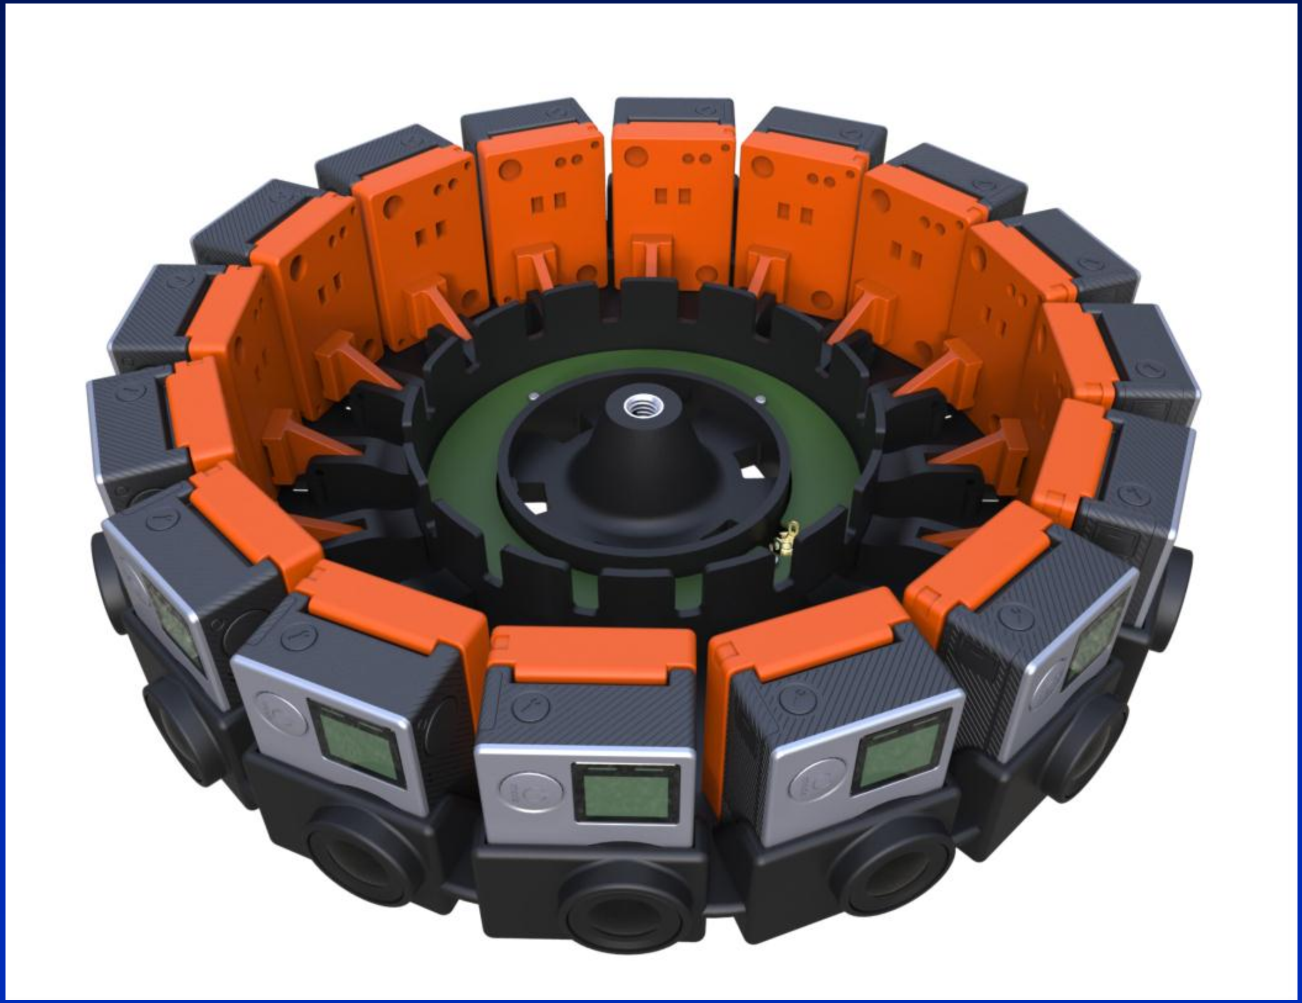

Google Earth, New York City 2016

Motion Capture Markers

Motion Capture

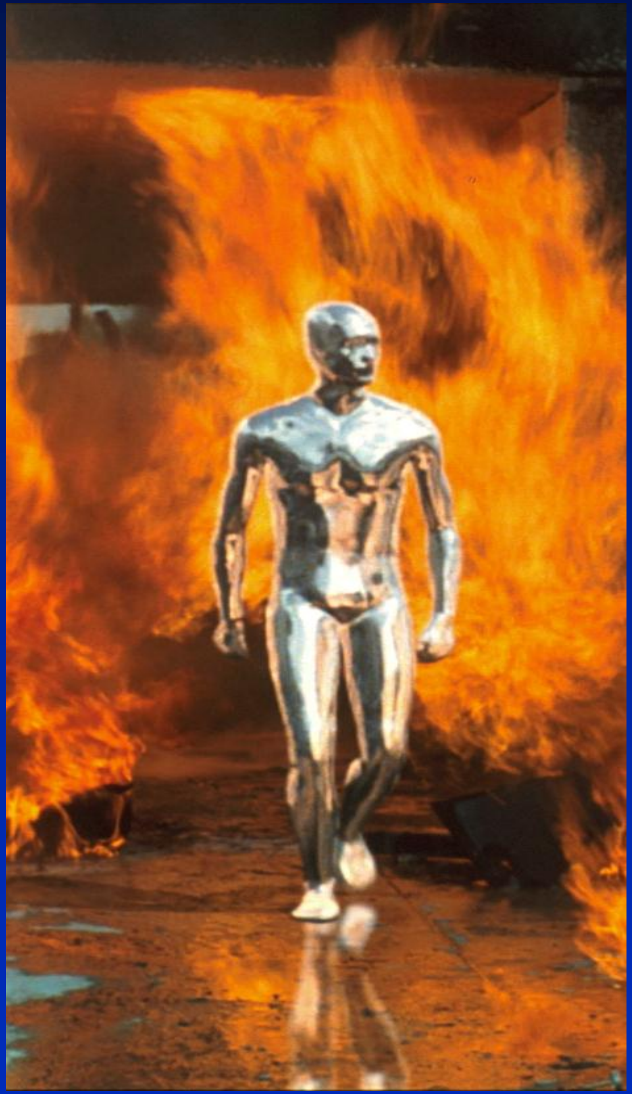

Markerless Motion Capture

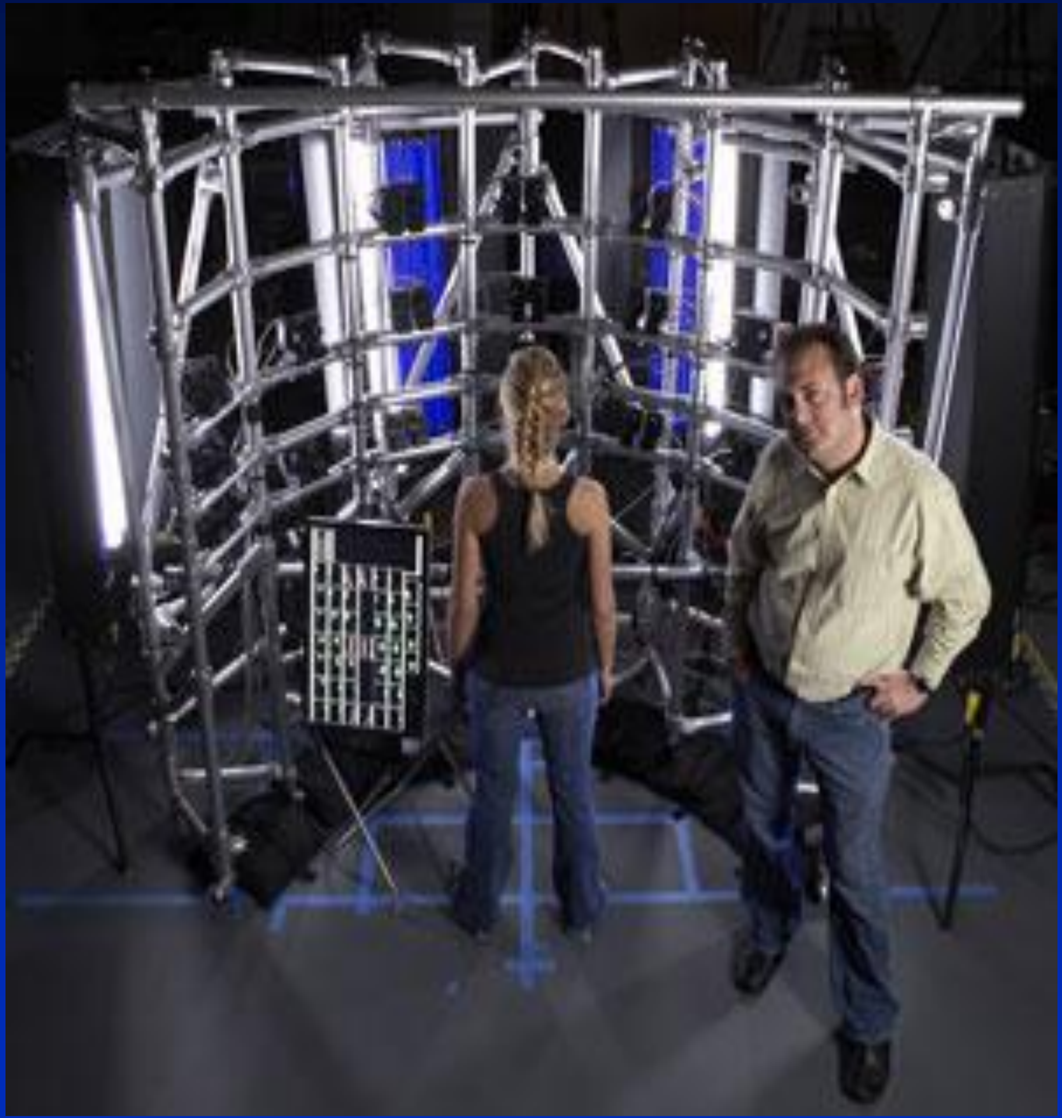

http://macanimationpro.digitalmedianet.com/articles/viewarticle.jsp?id=54623

The Curious Case of Benjamin Button 2008

Cinefex 116, January 2009

The Curious Case of Benjamin Button 2008

Cinefex 116, January 2009

The Curious Case of Benjamin Button 2008

Cinefex 116, January 2009

The Curious Case of Benjamin Button 2008

Ed Ulbrich: How Benjamin Button got his face

BenButton2.wmv

http://www.ted.com/index.php/talks/ed_ulbrich_shows_how_benjamin_button_got_his_face.html

Affective Computing

Facial Recognition

Science Magazine, 30.Jan.2015 – The End of Privacy

Mona Lisa

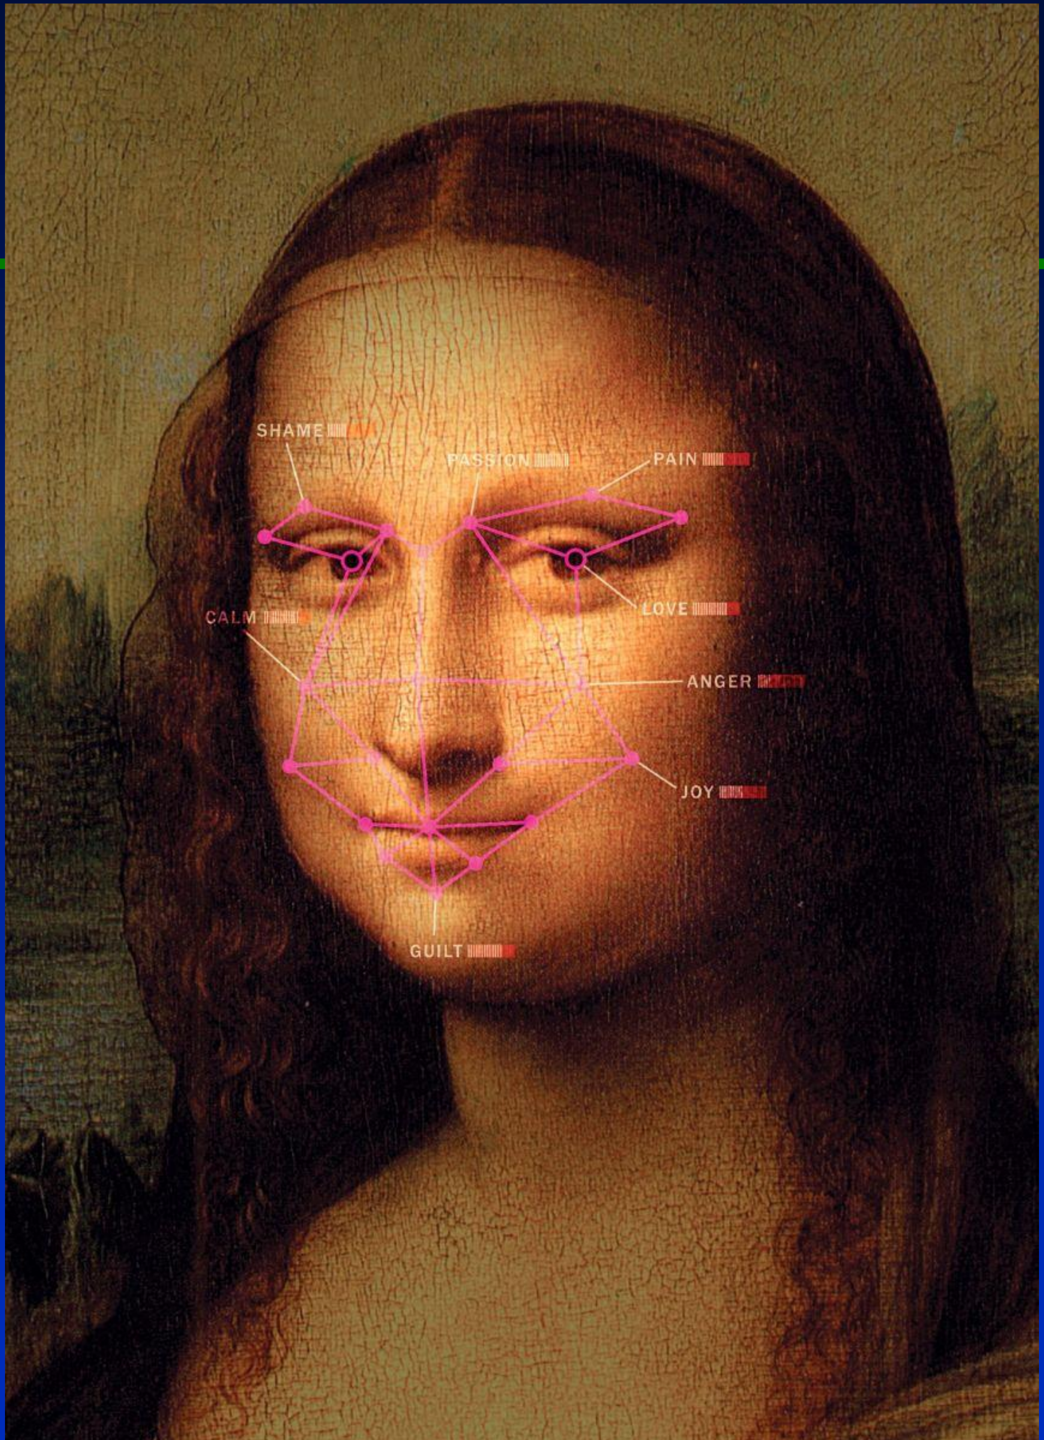

Eckman

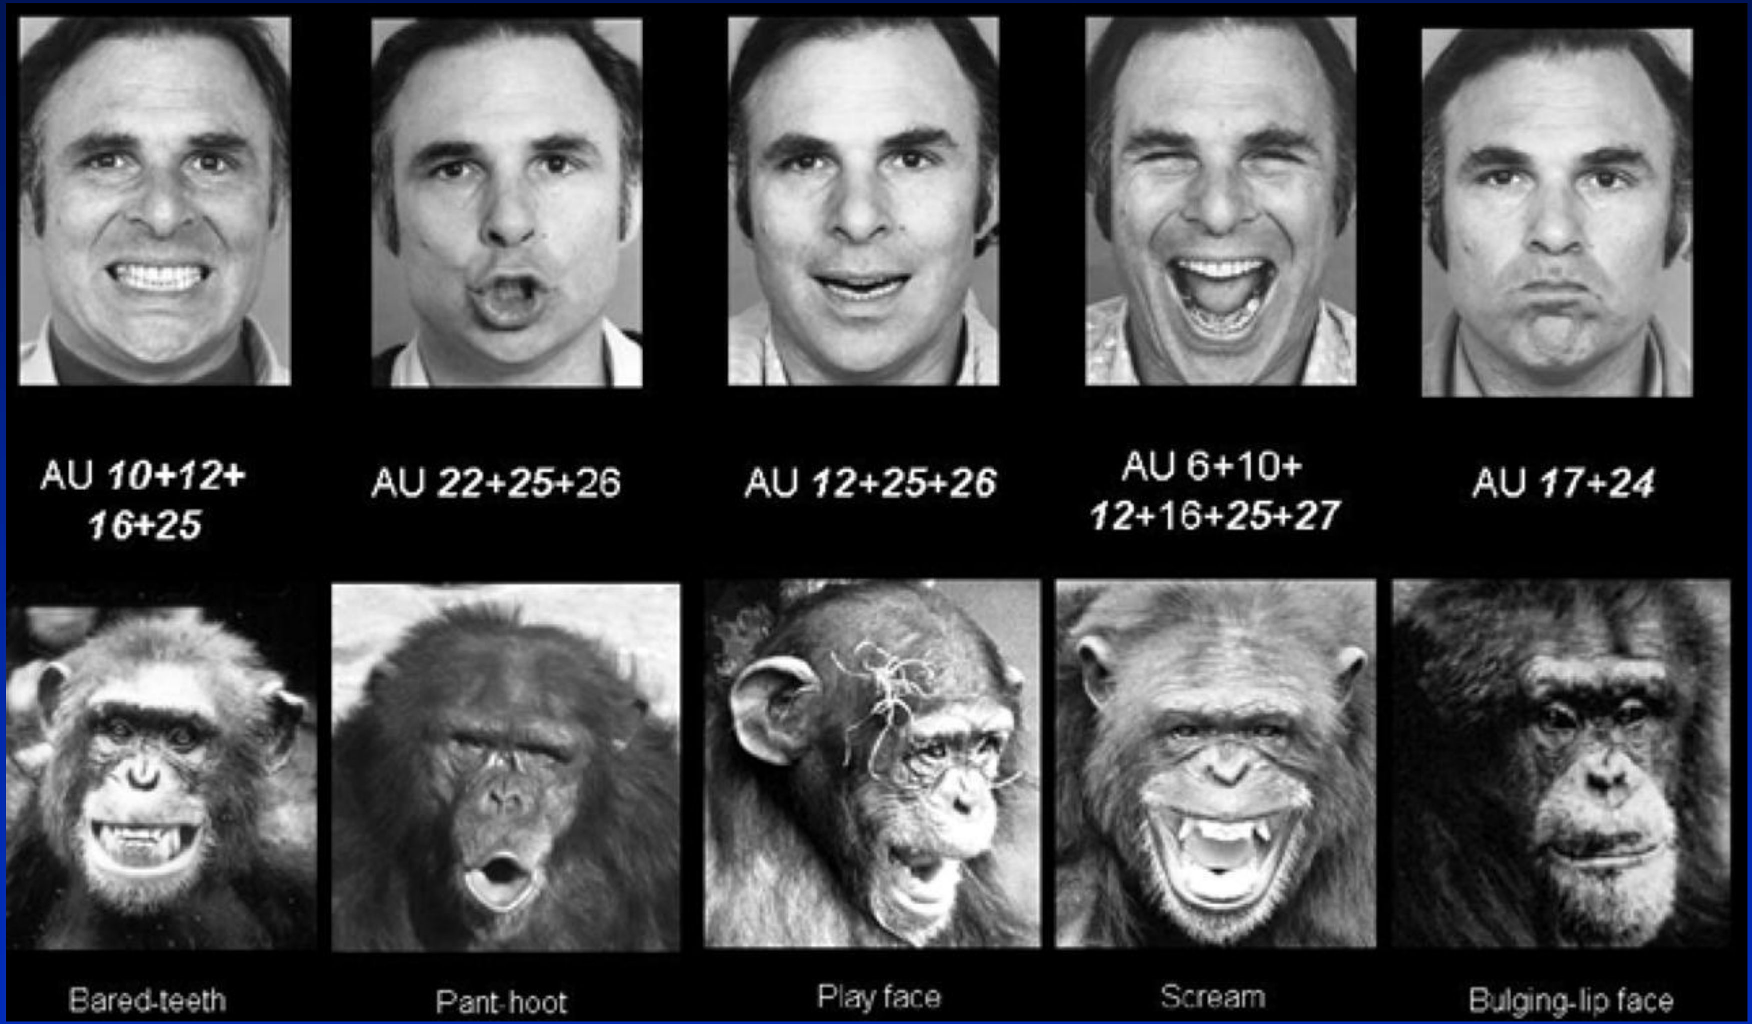
AU 10+12+
16+25
AU 22+25+26
AU 12+25+26
AU 6+10+
12+16+25+27
AU 17+24
Bared-teeth
Pant-hoot
Play face
Scream
Bulging-lip face

Eckman

anger
1 eyebrows down and together
2 eyes glare
3 narrowing of the lips
disgust
1 nose wrinkling
2 upper lip raised
fear
1 eyebrows raised and pulled together
2 raised upper eyelids
3 tensed lower eyelids
4 lips slightly stretched horizontally back to ears
happiness
A real smile always includes:
1 crow's feet wrinkles
2 pushed up cheeks
3 movement from muscle that orbits the eye
sadness
1 drooping upper eyelids
2 losing focus in eyes
3 slight pulling down of lip corners
surprise
Lasts for only one second:
1 eyebrows raised
2 eyes widened
3 mouth open
contempt
1 lip corner tightened and raised on only one side of face

Inside Out

DISGUST
FEAR
JOY
SADNESS
ANGER
INSIDE OUT
INSIDE OUT
INSIDE OUT
INSIDE OUT
INSIDE OUT

Inside Out

End