Graphics Hardware (cont'd)
Geometry Capture
Motion Capture

Visual Imaging in the Electronic Age

Lecture # 18

November 5, 2020

Donald P. Greenberg

Ford Assembly Line 1913

Graphics Pipeline Hardware

“Moore's Law is for wimps.”

Interesting Trends

- Moore’s Law is approximately 1.7x increase in speed a year; graphics accelerators are improving at 2x to 4x a year

Moore’s Law – GPU Transistor Counts

Processor	Transistor count	Date of introduction	Manufacturer	Process	Area
R520	321,000,000	2005	AMD	90 nm	288 mm^2
R580	384,000,000	2006	AMD	90 nm	352 mm^2
G80	681,000,000	2006	NVIDIA	90 nm	480 mm^2
R600 Pele	700,000,000	2007	AMD	80 nm	420 mm^2
G92	754,000,000	2007	NVIDIA	65 nm	324 mm^2
RV790XT Spartan	959,000,000	2008	AMD	55 nm	282 mm^2
GT200 Tesla	1,400,000,000	2008	NVIDIA	65 nm	576 mm^2
Cypress RV870	2,154,000,000	2009	AMD	40 nm	334 mm^2
Cayman RV970	2,640,000,000	2010	AMD	40 nm	389 mm^2
GF100 Fermi	3,200,000,000	Mar 2010	NVIDIA	40 nm	526 mm^2
GF110 Fermi	3,000,000,000	Nov 2010	NVIDIA	40 nm	520 mm^2
GK104 Kepler	3,540,000,000	2012	NVIDIA	28 nm	294 mm^2
Tahiti RV1070	4,312,711,873	2011	AMD	28 nm	365 mm^2
GK110 Kepler	7,080,000,000	2012	NVIDIA	28 nm	561 mm^2
RV1090 Hawaii	6,300,000,000	2013	AMD	28 nm	438 mm^2
GM204 Maxwell	5,200,000,000	2014	NVIDIA	28 nm	398 mm^2
GM200 Maxwell	8,100,000,000	2015	NVIDIA	28 nm	601 mm^2
Fiji	8,900,000,000	2015	AMD	28 nm	596 mm^2

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

GPU Transistor Counts

Processor	Transistor count	Date of introduction	Manufacturer	Process	Area
GP102 Pascal	12,000,000,000	2016	Nvidia	16 nm	471 mm^2
GP100 Pascal	15,300,000,000	2016	Nvidia	16 nm	610 mm^2
GV100 Volta	21,100,000,000	2017	Nvidia	12 nm	815 mm^2

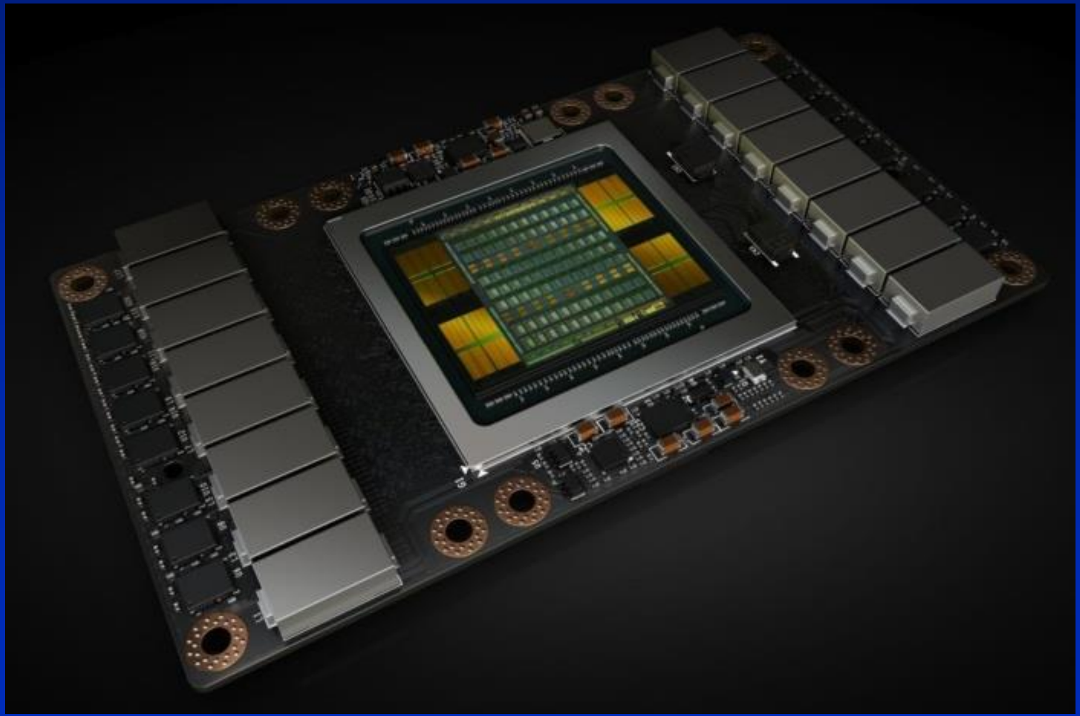

GV100 Volta

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

Nomenclature

- FLOPS- floating point operations per second.
- MIPS- millions of instructions per second
- MEGAFLOPS- million floating point instructions per second
- GIGAFLOPS- billion floating points per second
- TERAFLOPS- trillions of floating point instructions per second.

Instructions Per Second (IPS)

- IPS = No. of cores x (No. of cycles/sec.) x (Instructions/cycle)
- Example: Digital Equipment Corporation
  - VAX 11/780, 5Mhz
  - Cost $250,000 → 1977
- IPS = $1x5x10^6$(cycles/sec)x0.2(Instructions/cycle)
  - =$1x10^6$(Instructions)
  - □ The first 1.0 MIPS machine

Instructions Per Clock-cycle per second (IPS) (Microprocessor CPU’S)

Processor / System	MIPS and Frequency	(IPS / clock cycles per second)	(IPS / clock cycles per second / cores)	Year
Intel 4004	0.092 MIPS at 0.740 MHz	0.124	0.124	1971
VAX-11/780	1.000 MIPS at 5.000 MHz	0.2	0.2	1977
Motorola 68030	9 MIPS at 25 MHz	0.36	0.36	1987
Intel Pentium	188 MIPS at 100 MHz	1.88	1.88	1994
Graphics Processor / System	**MIPS and Frequency**	**(IPS / clock cycles per second)**	**(IPS / clock cycles per second / cores)**	**Year***
ARM Cortex A7	2,850 MIPS at 1.5 GHz	1.9	1.9	2011
Nvidia Tegra 3	13,800 MIPS at 1.5 GHz	9.2	2.5	2011
ARM Cortex A73 (4-core)	71,120 MIPS at 2.8 GHz	25.4	6.35	2016
Graphics Boards/ System	**MIPS and Frequency**	**(IPS / clock cycles per second)**	**(IPS / clock cycles per second / cores)**	**Year***
*Nvidia Ge-Force RTX-2080T.I	76,000,000 MIPs at 1.5 GHz	50,666.7	11.6	2019
*Nvidia Quadro RTX-2080T.I	78,000,000 MIPS at 1.6 GHz	48750.0	11.2	2019
Graphics Servers/ System	**MIPS and Frequency**	**(IPS / clock cycles per second)**	**(IPS / clock cycles per second / cores)**	**Year***
AMD Ryzen 71800X	304,510 MIPS at 3.6 GHz	84.6	10.6	2017
*Nvidia DGX -1	500,000,000 MFLOPs at 1.46 GHz	357142.9	17.4	2017

* These measurements are now a mix of various types of standard operations

Ford 2020


RH SEALER TAPE #1AND #2
LOOK FOR LAMP CAN TAPE RH
LOOK FOR LAMP CAN TAPE LH
LH SEALER TAPE #1AND#2

BMW 2020

Nvidia Titan Xp August 2017

3840 CUDA Cores
12,000,000 transistors
224 texture mapping units
16 nm process technology

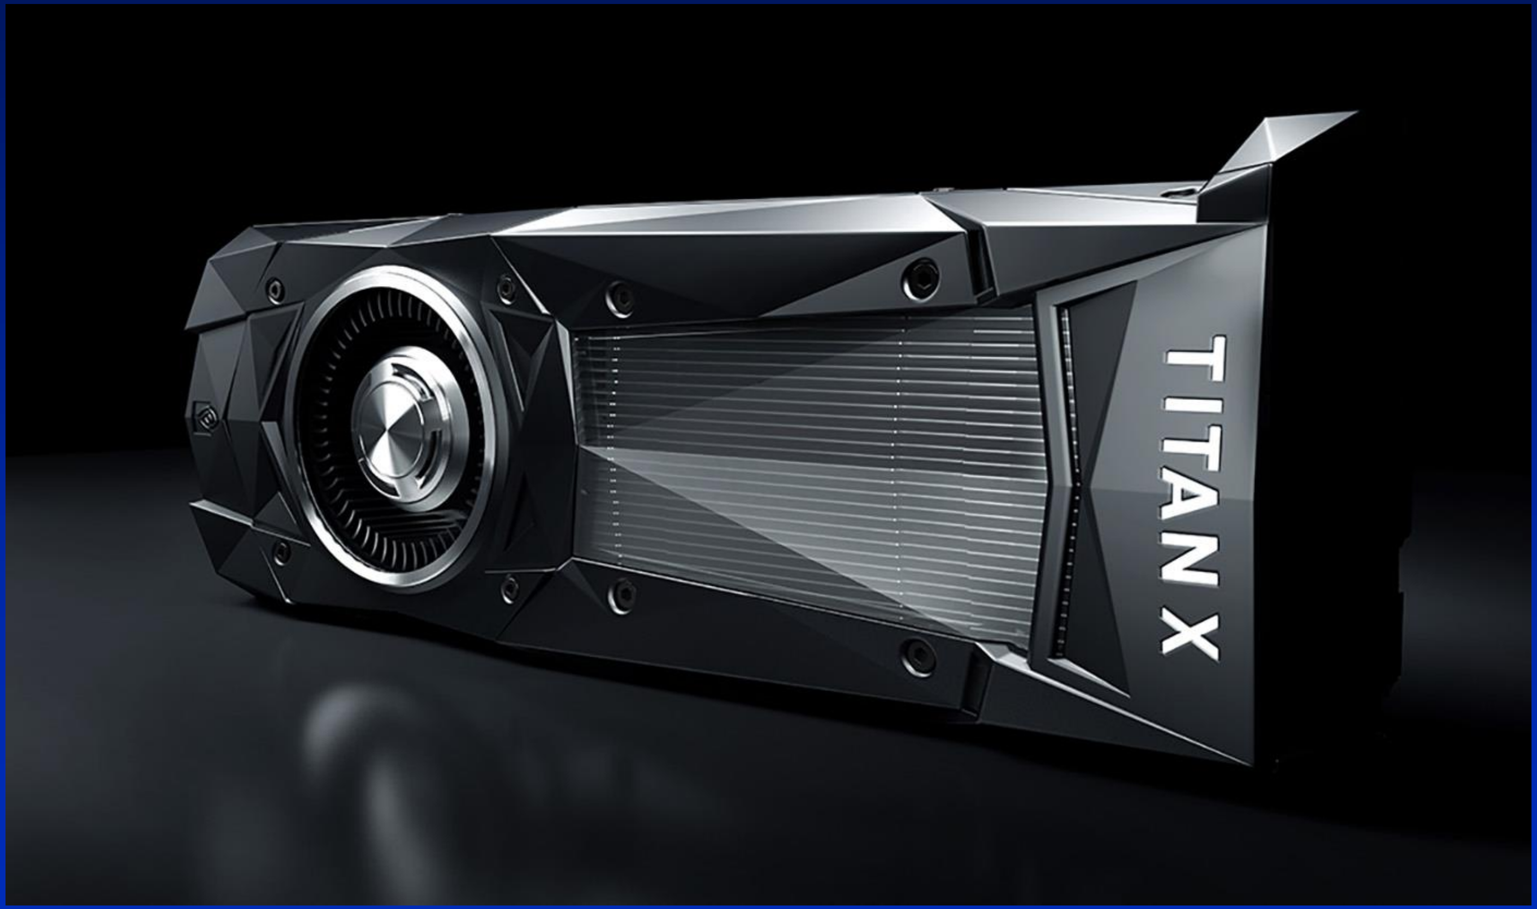

Nvidia RTX 2080 Ti 2018

Nvidia DGX 2018

4X Tesla V100
500 TFLOPS
20,480 CUDA Cores
256 GB memory

AMD Threadripper 2020

64 Cores, 4.35 GHz, 8.46 Inst/cycle=

2,356,000 MIPS

Interesting Trends

- Moore's Law is approximately 1.7x increase in speed a year; graphics accelerators are improving at 2x to 4x a year

Is reducing energy (watts) is more important than increasing processing speed (flops)?

Nvidia tries to buy ARM

Nvidia Ampere RTX 3090 2020

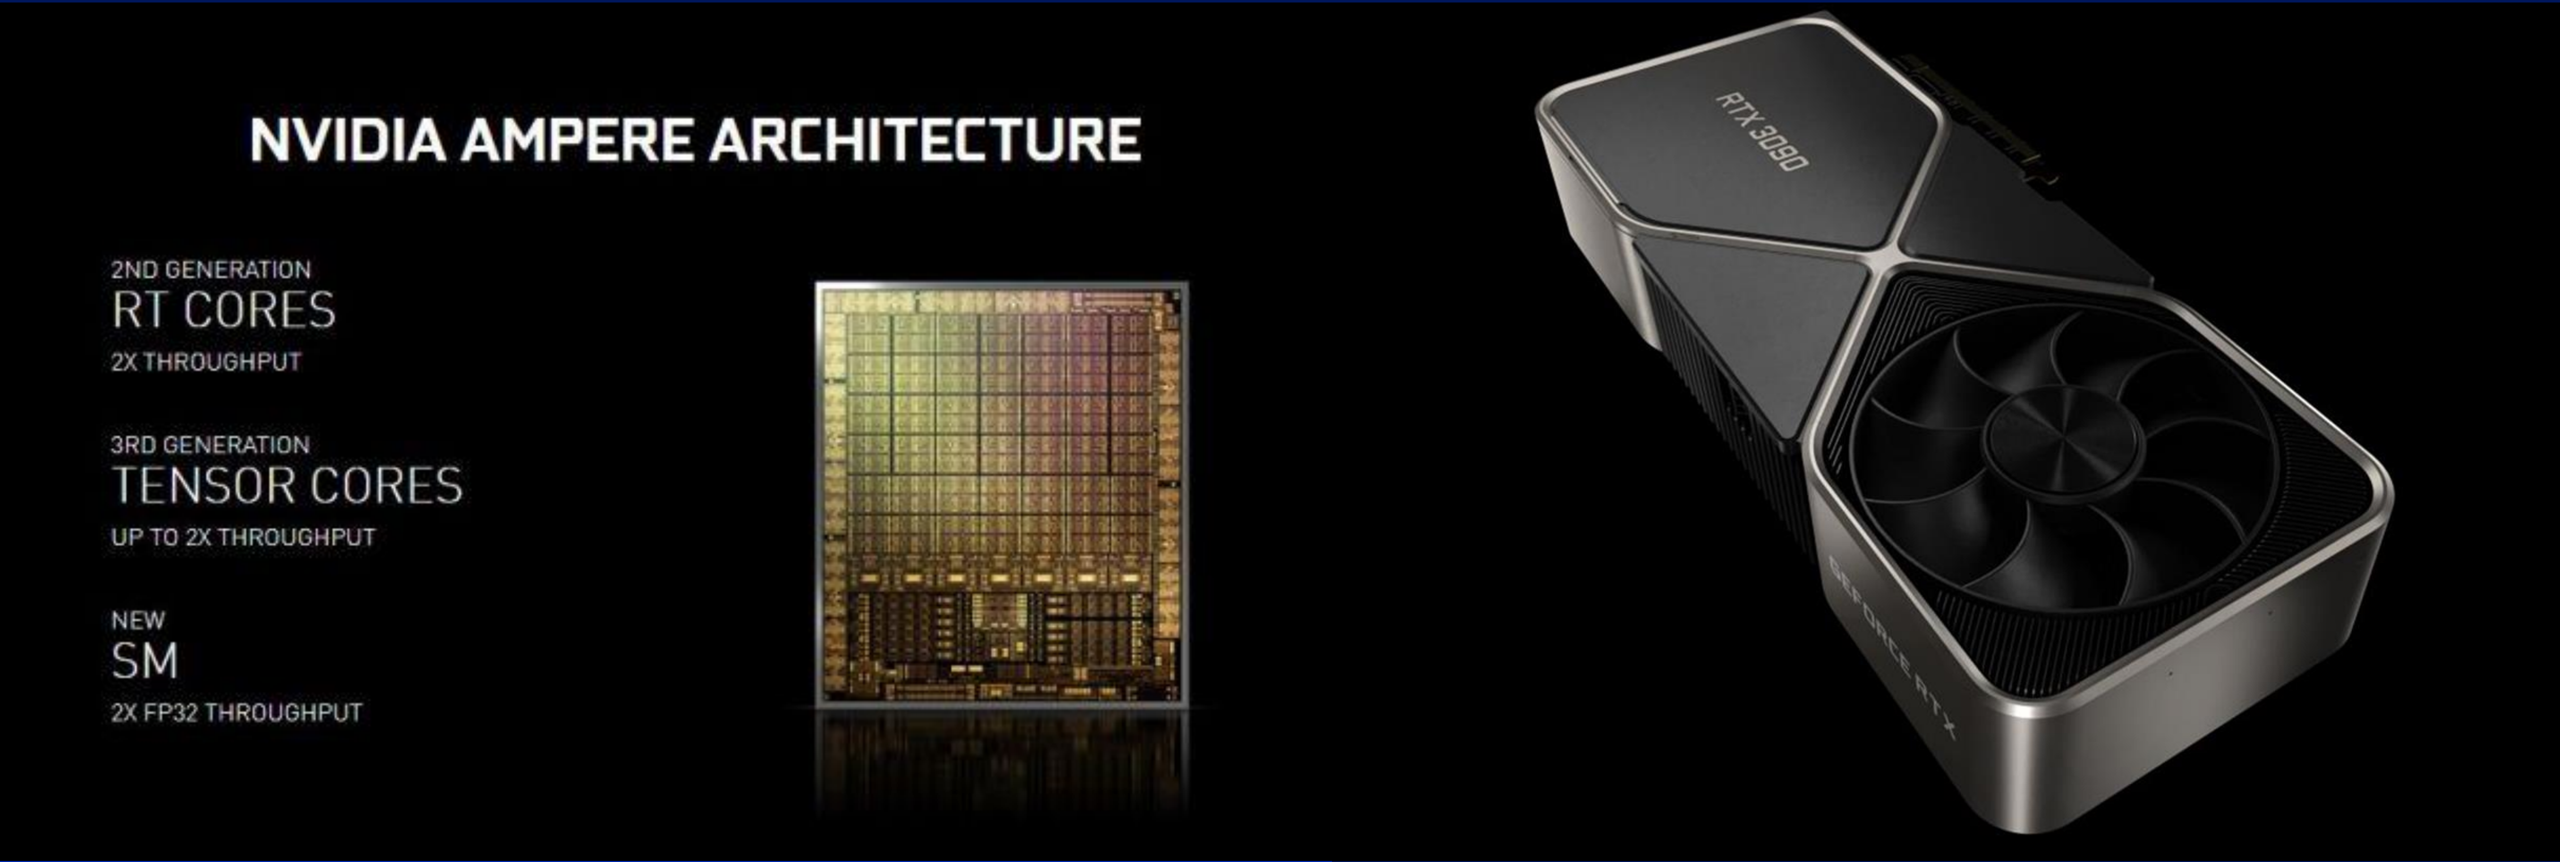

Digital Geometry Capture

- Photographic methods (cont’d)
- Laser scanning
- Time of Flight Sensors

ReMake Autodesk

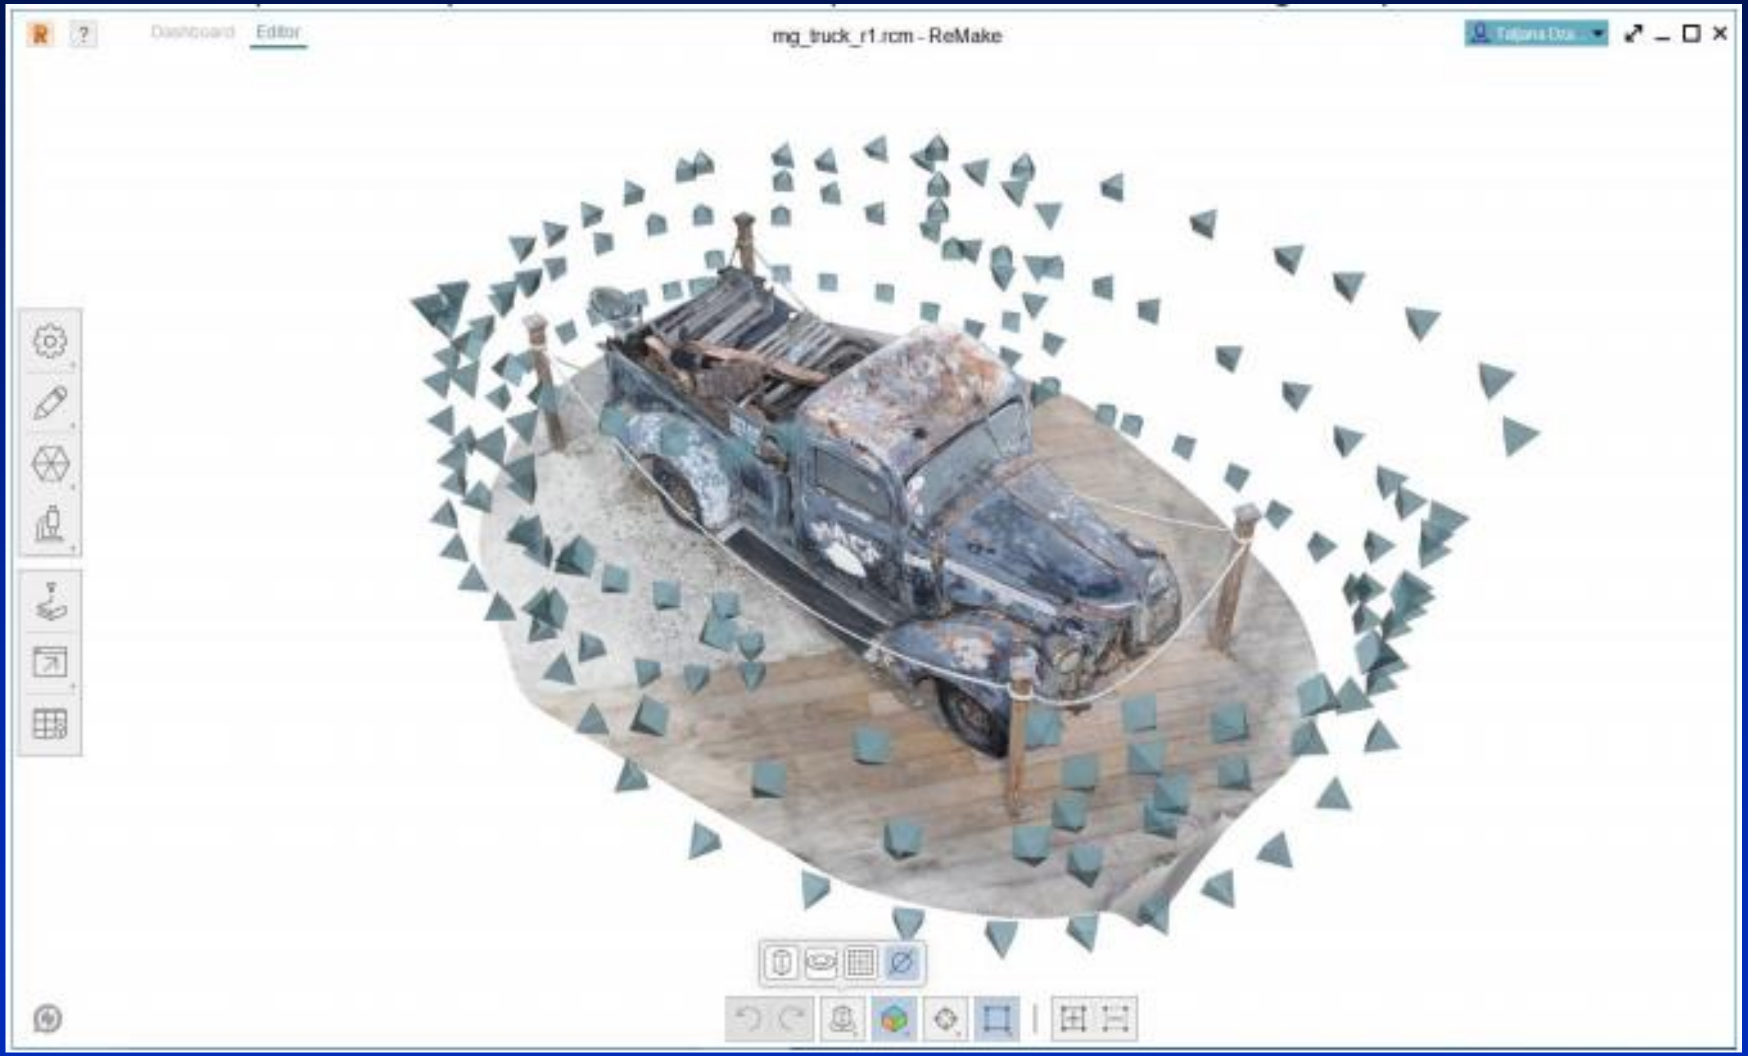

Cornell Campus, McGraw Hall - Noah Snavely

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.

Kinect for xBox 360

The original Kinect used 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)

Kinect speckle pattern

- Near region (0.8 – 1.2m)
  Small dots
- Middle region (1.2 – 2.0m)
- Far region (2.0 – 3.5 m)
  Large dots

Structured Light Imaging (Kinect One)

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

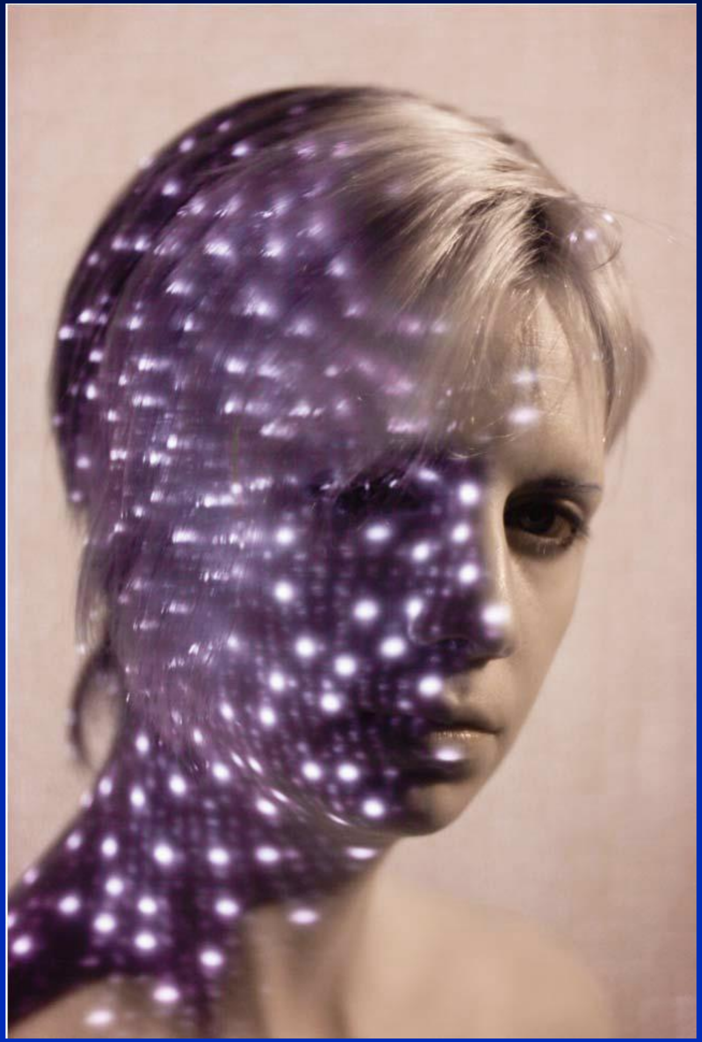

Point Cloud Drawn from the Kinect's 3D data

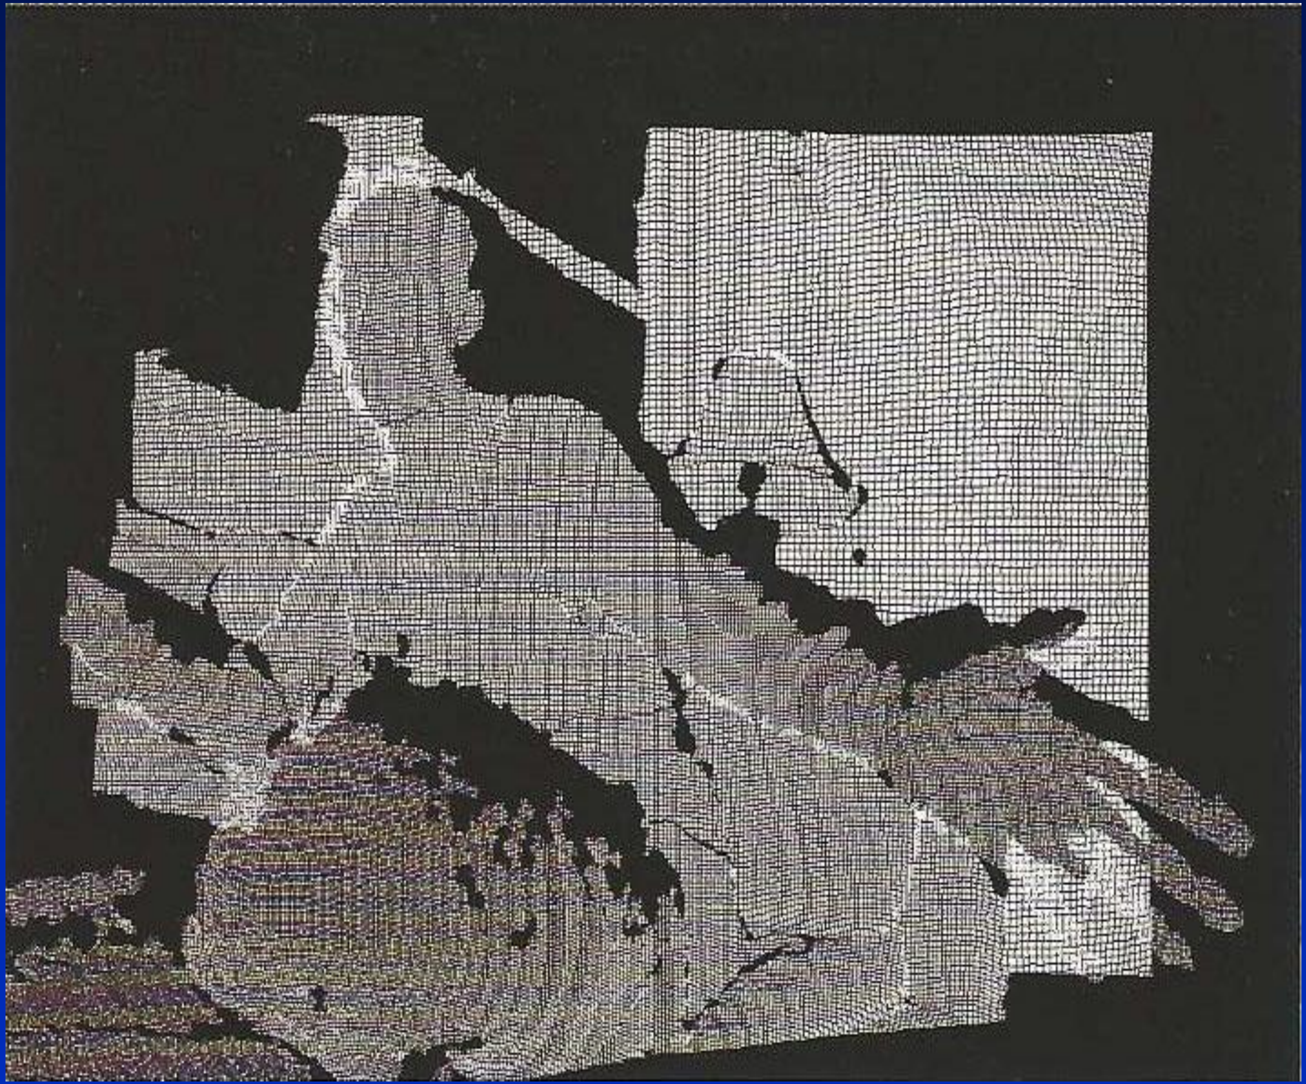

Kinect: Depth Image and Real Image

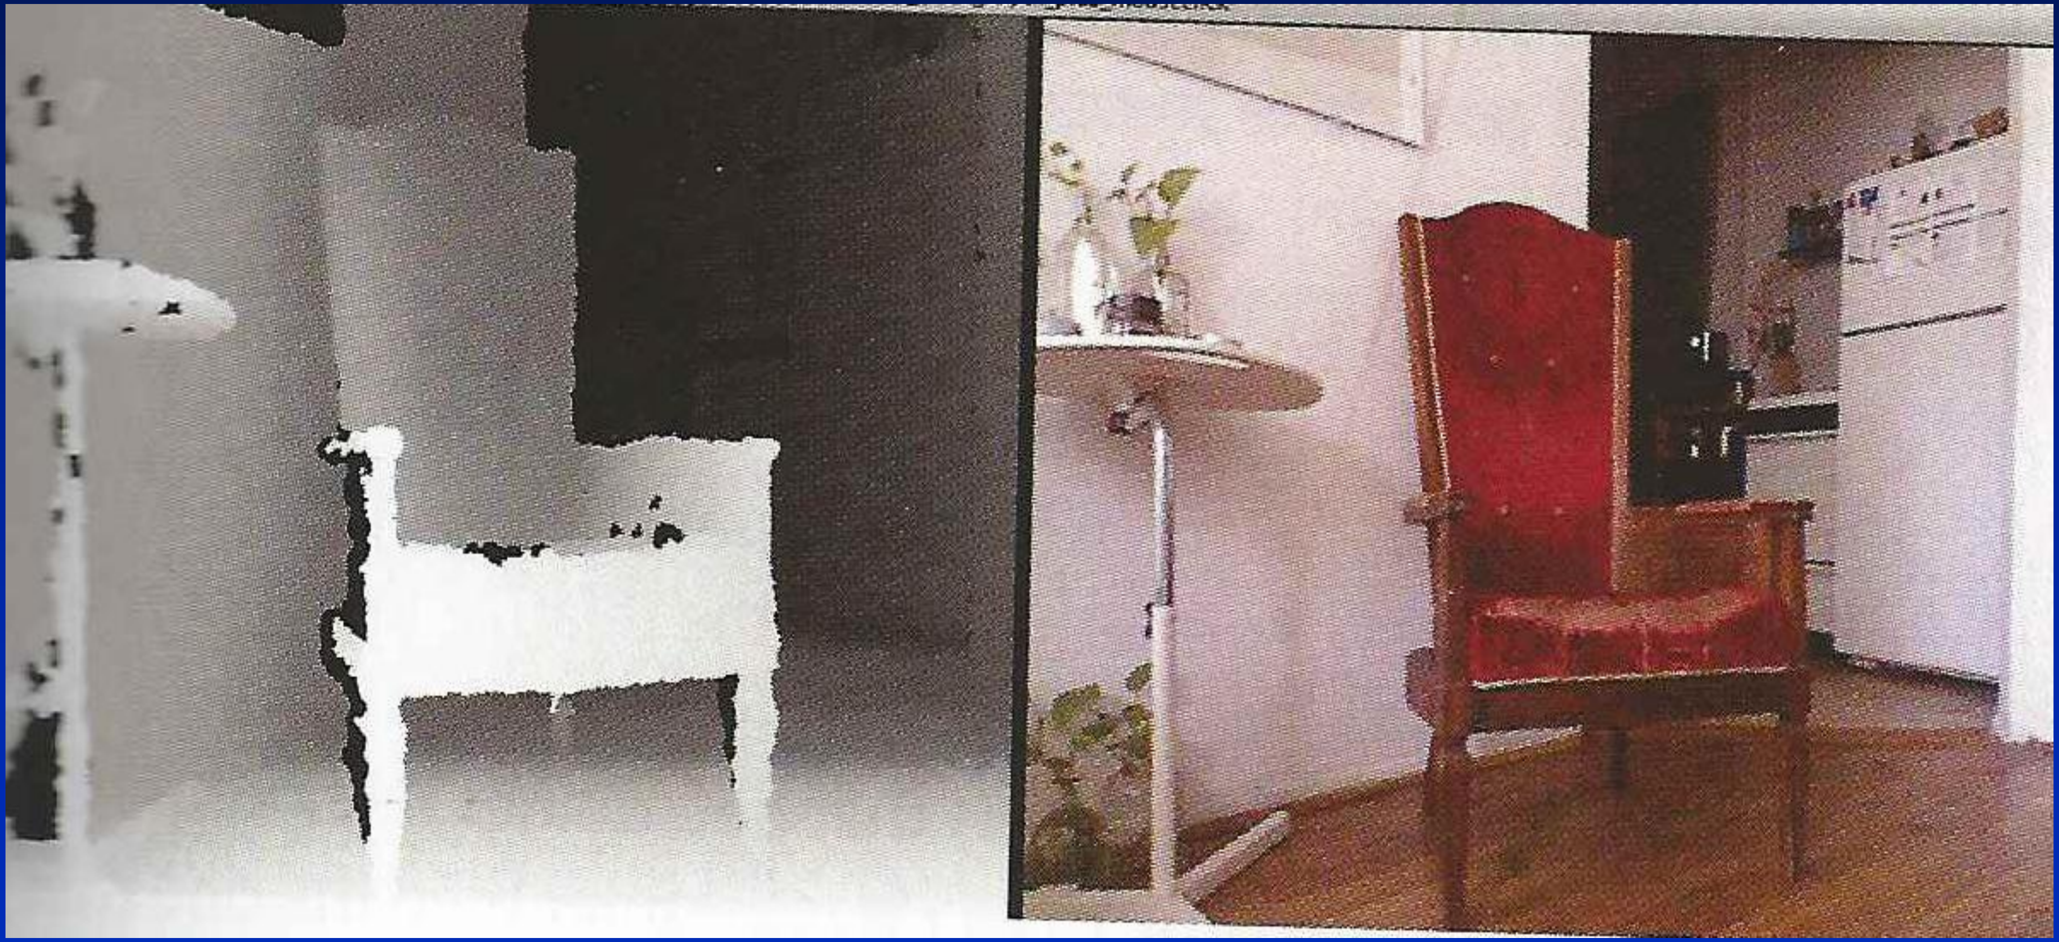

Digital Geometry Capture

- Photographic methods
- Laser scanning
- Time of Flight Sensors

Cyberware Scanner

Cyberware Scanner Diagram

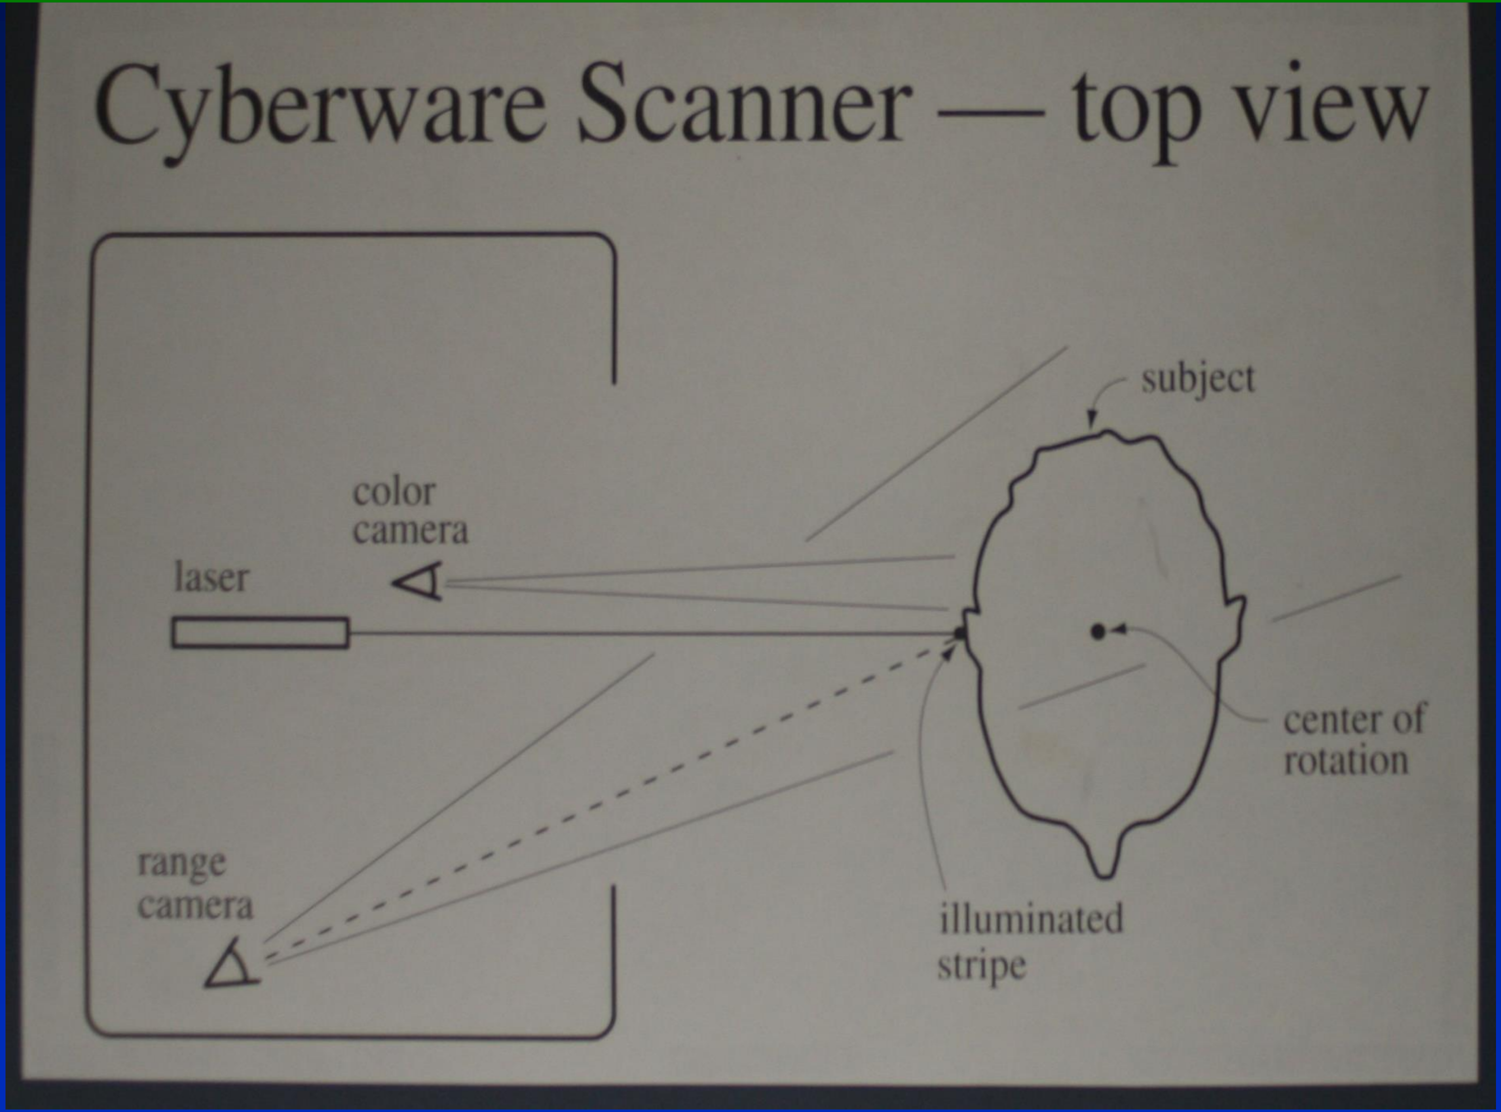

Cyberware Scanner

Uncle Don

Cyberware vs. 1 2 3 Catch

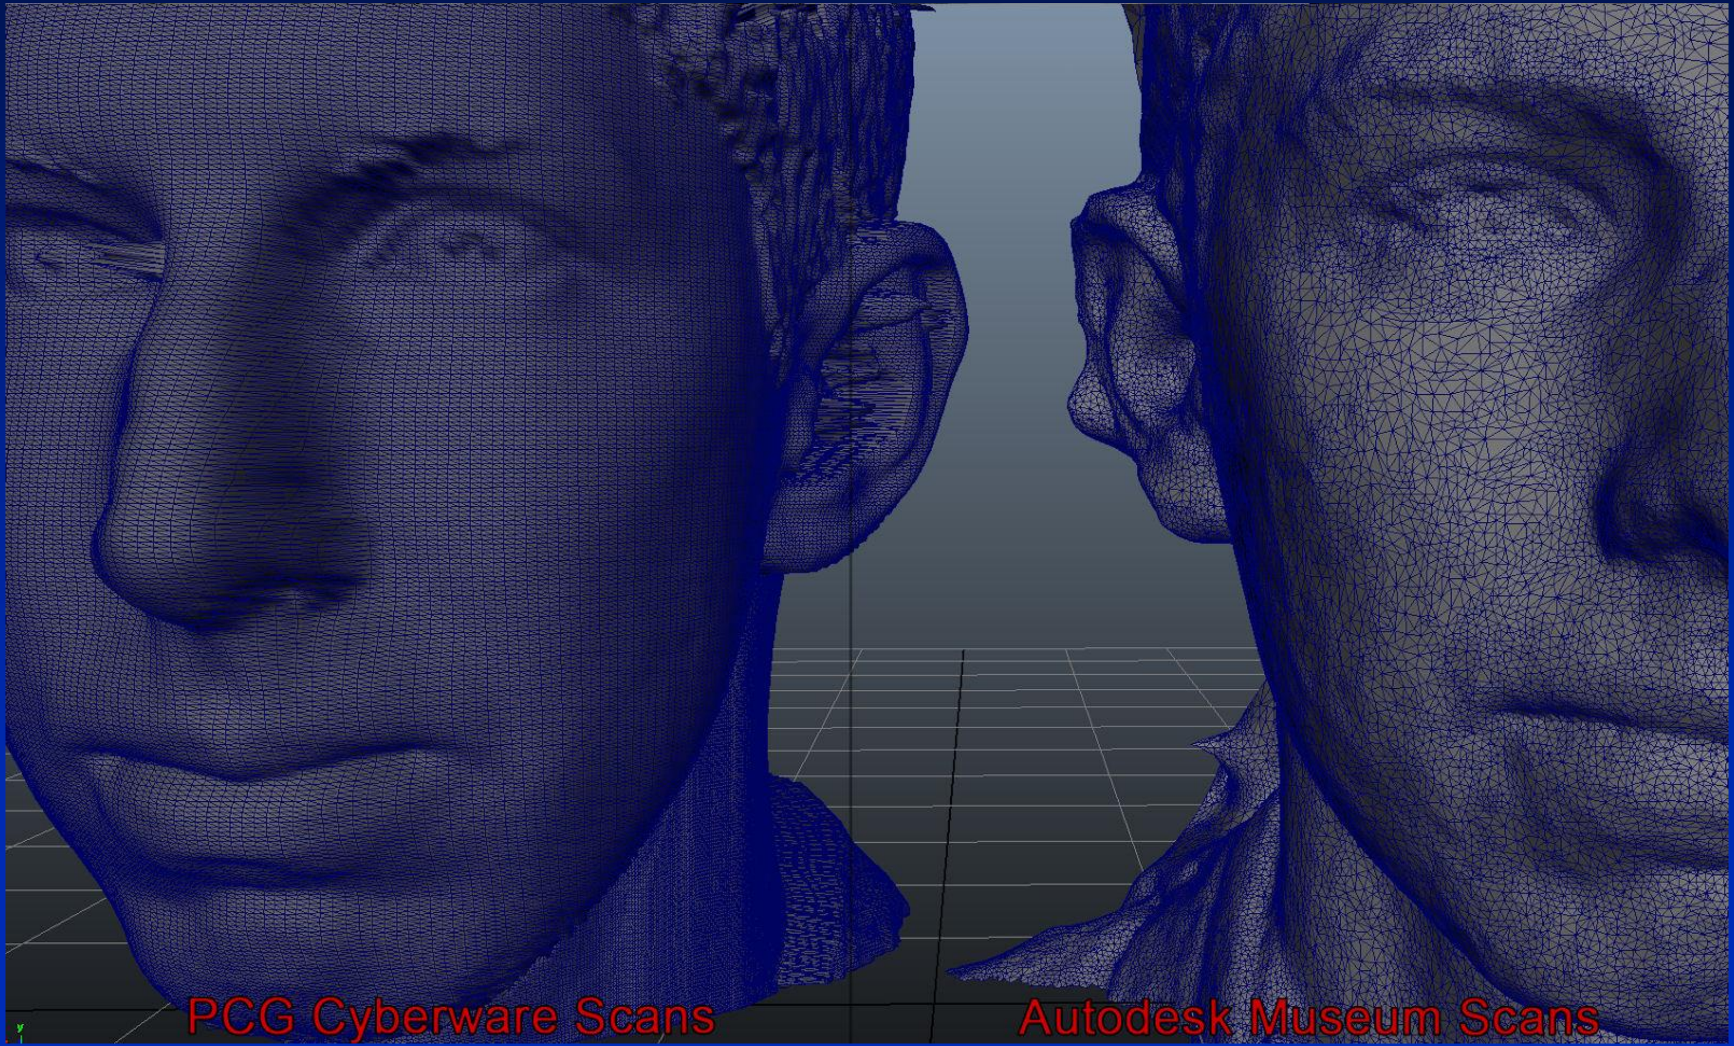

Digital Geometry Capture

- Photographic methods
- Laser scanning
- Time of Flight Sensors

Time of Flight Limitations

Limited by distance
Limited by duration of pulse
Ability to maximize reflection

Measurement Metrics

- Speed of light (300 meters $x10^8$/sec)
- Number of samples/sec ~200 Thz
- Thz = 200 x 10^{12} cycles/sec

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, $c = 3 \cdot 10^8 m/s$

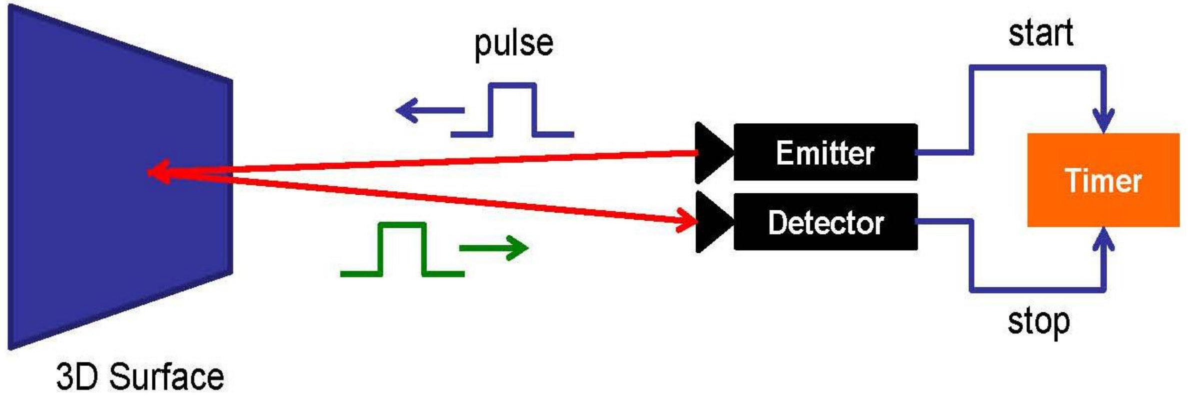

Matterport 2016

Seeing Behind Walls 2019

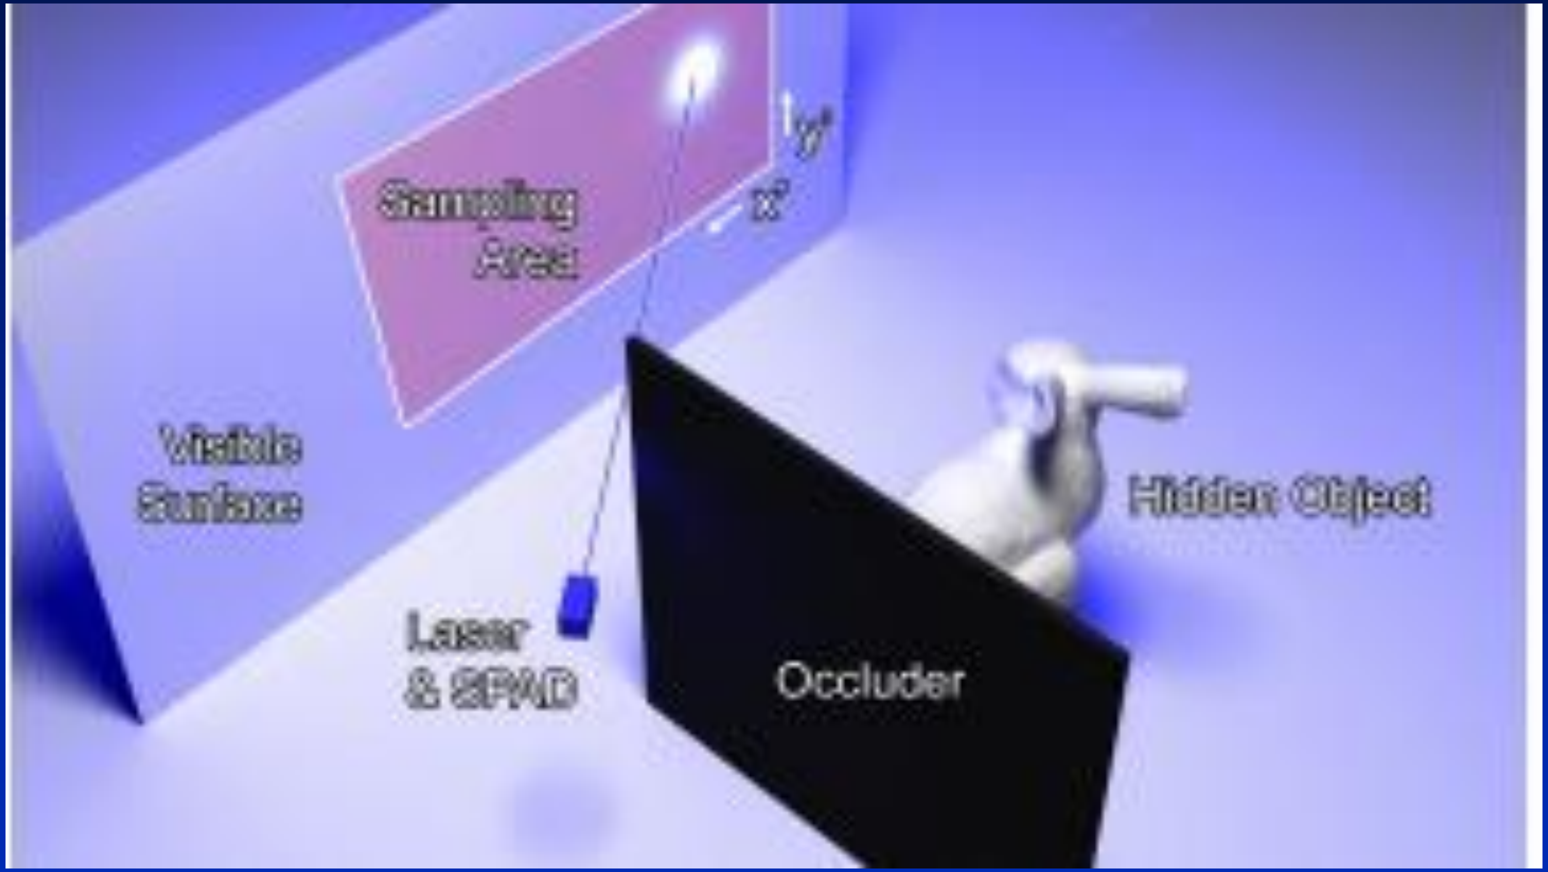

Wetzstein

Stanford

Confocal Non-Line-of-Sight Imaging

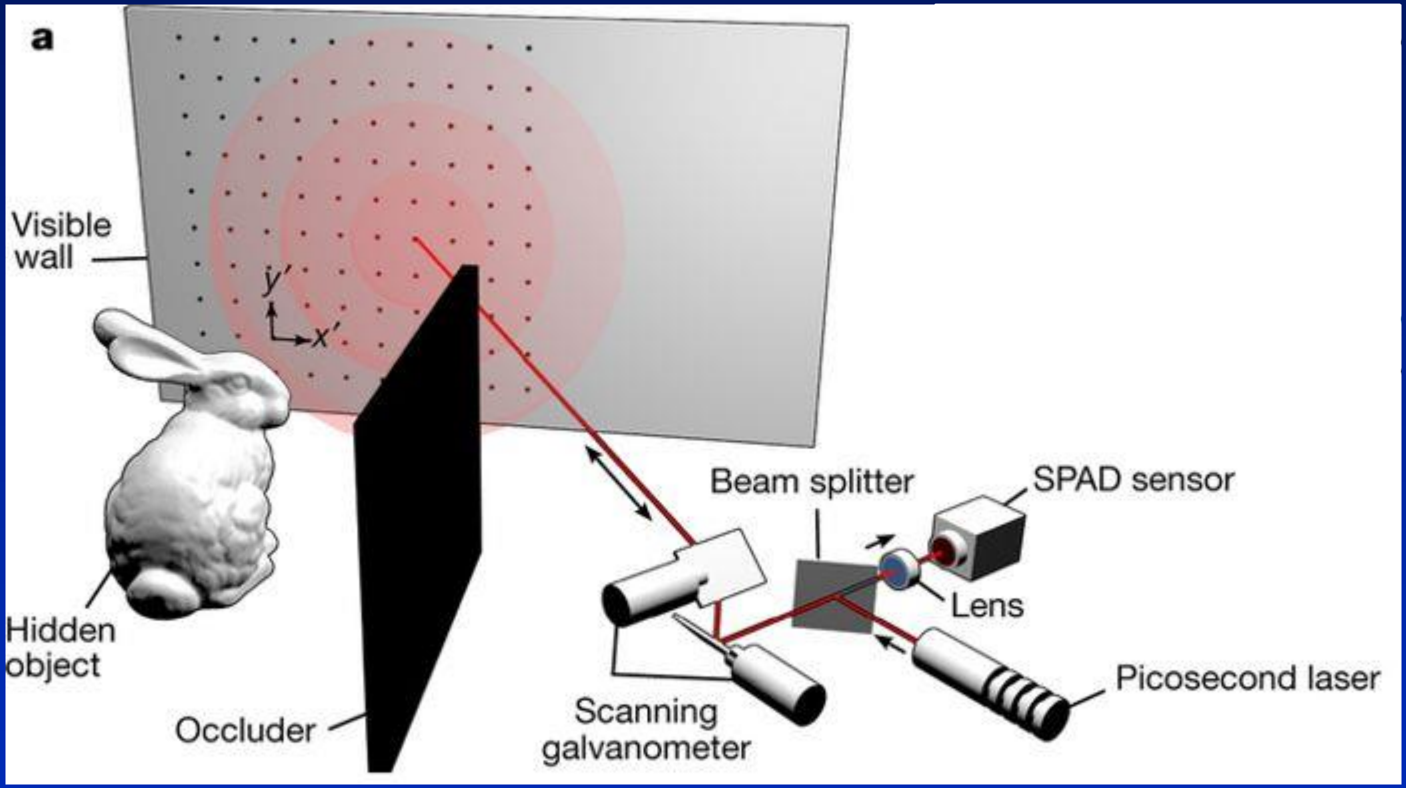

Nature, March 2018

C-NLOS Reconstruction and True Geometry

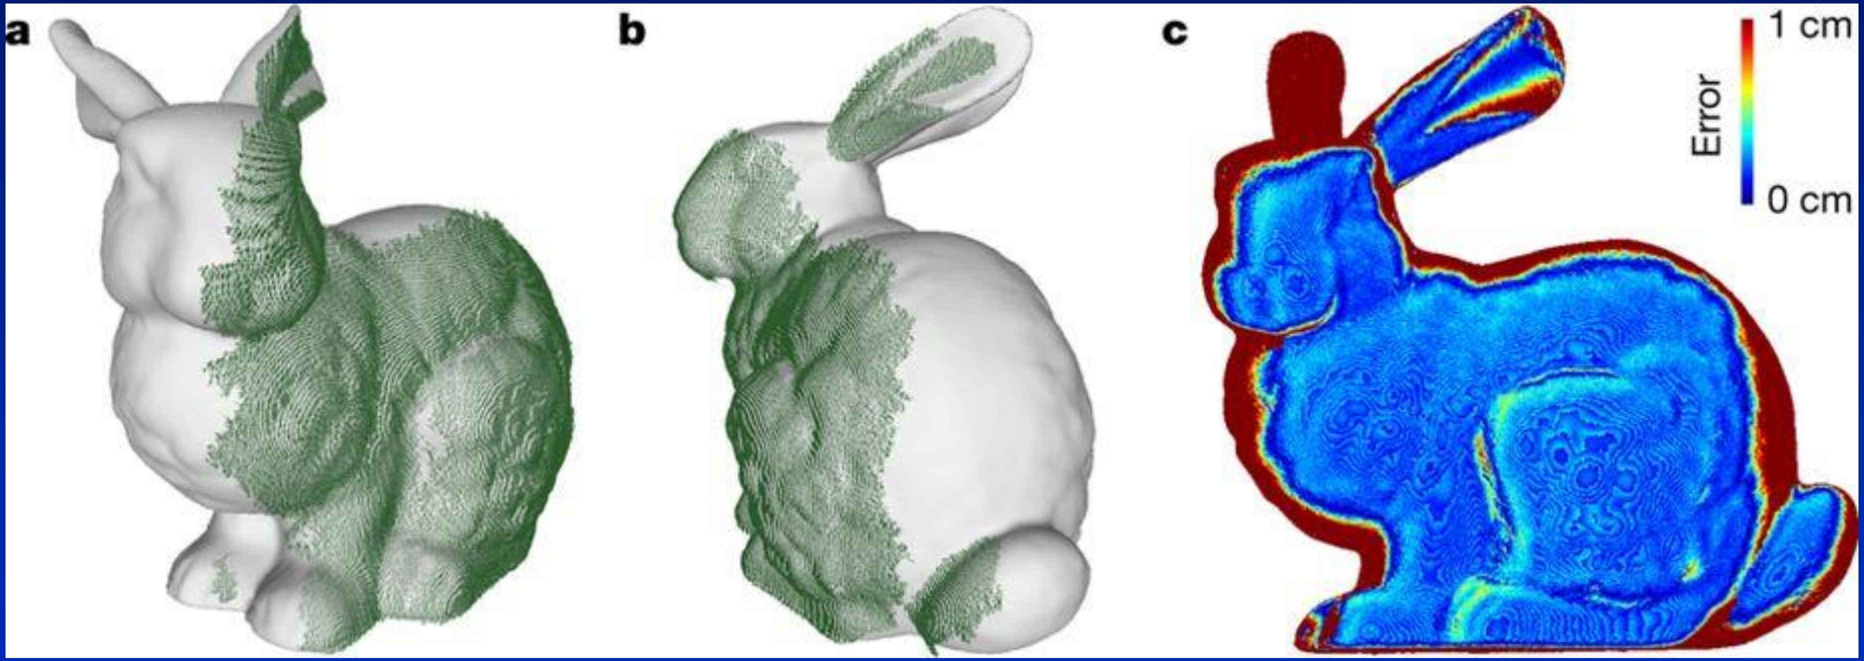

Nature, March 2018

Digital Geometry and Motion Capture

- Photographic methods
- Laser scanning
- Sonar
- Time of Flight
- All of the Above

Johansson mo cap video 1973

Motion Capture Markers

Motion Capture

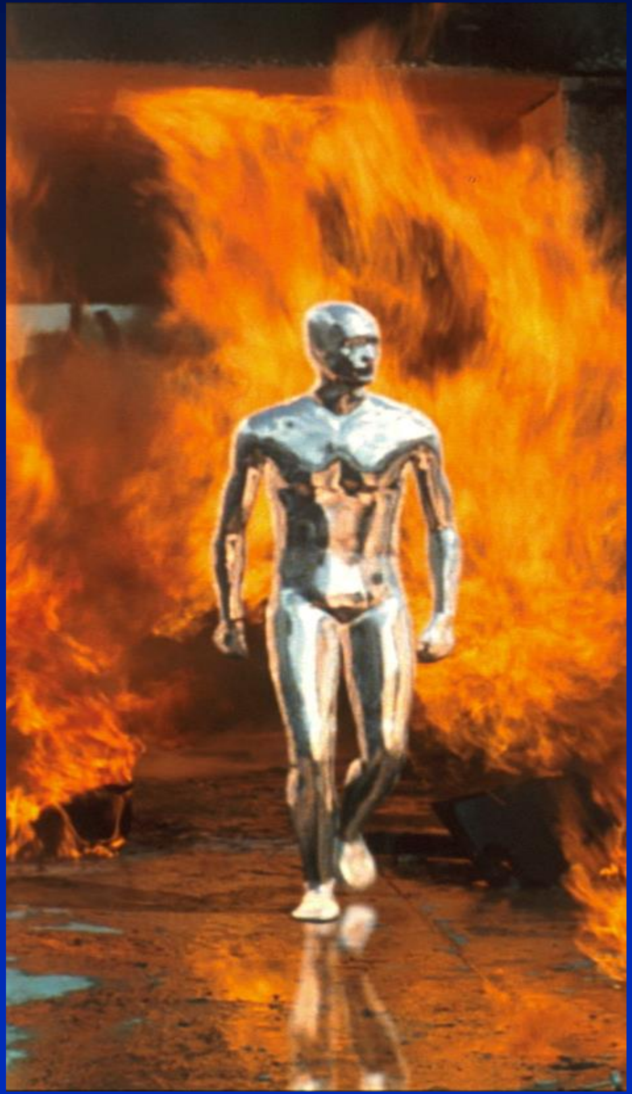

Kinect for xBox 360

Step 1: Compute a Depth Map

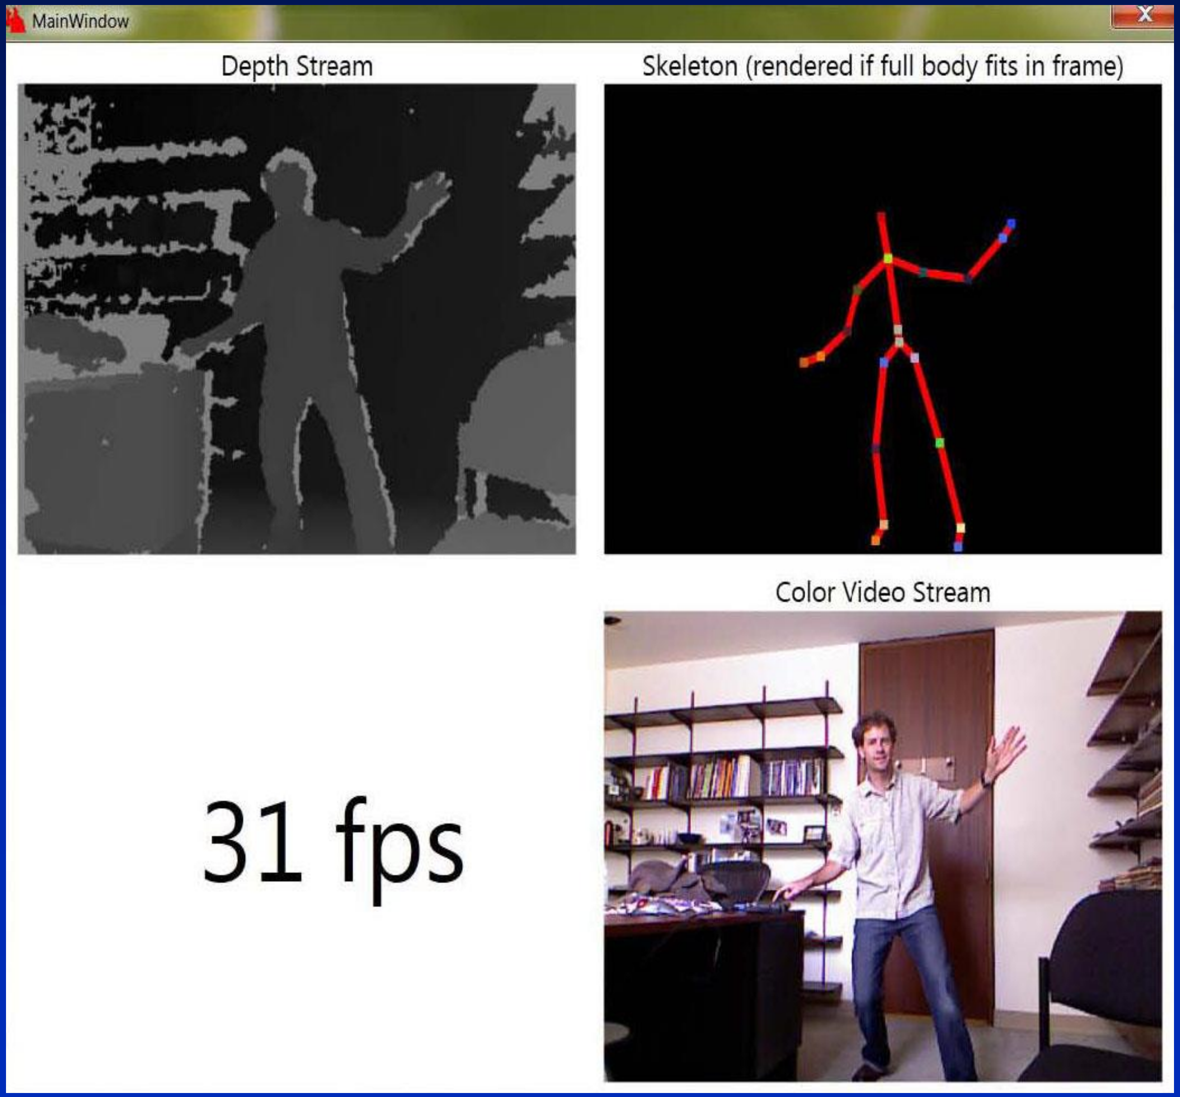

Extracted Skeleton Kinect

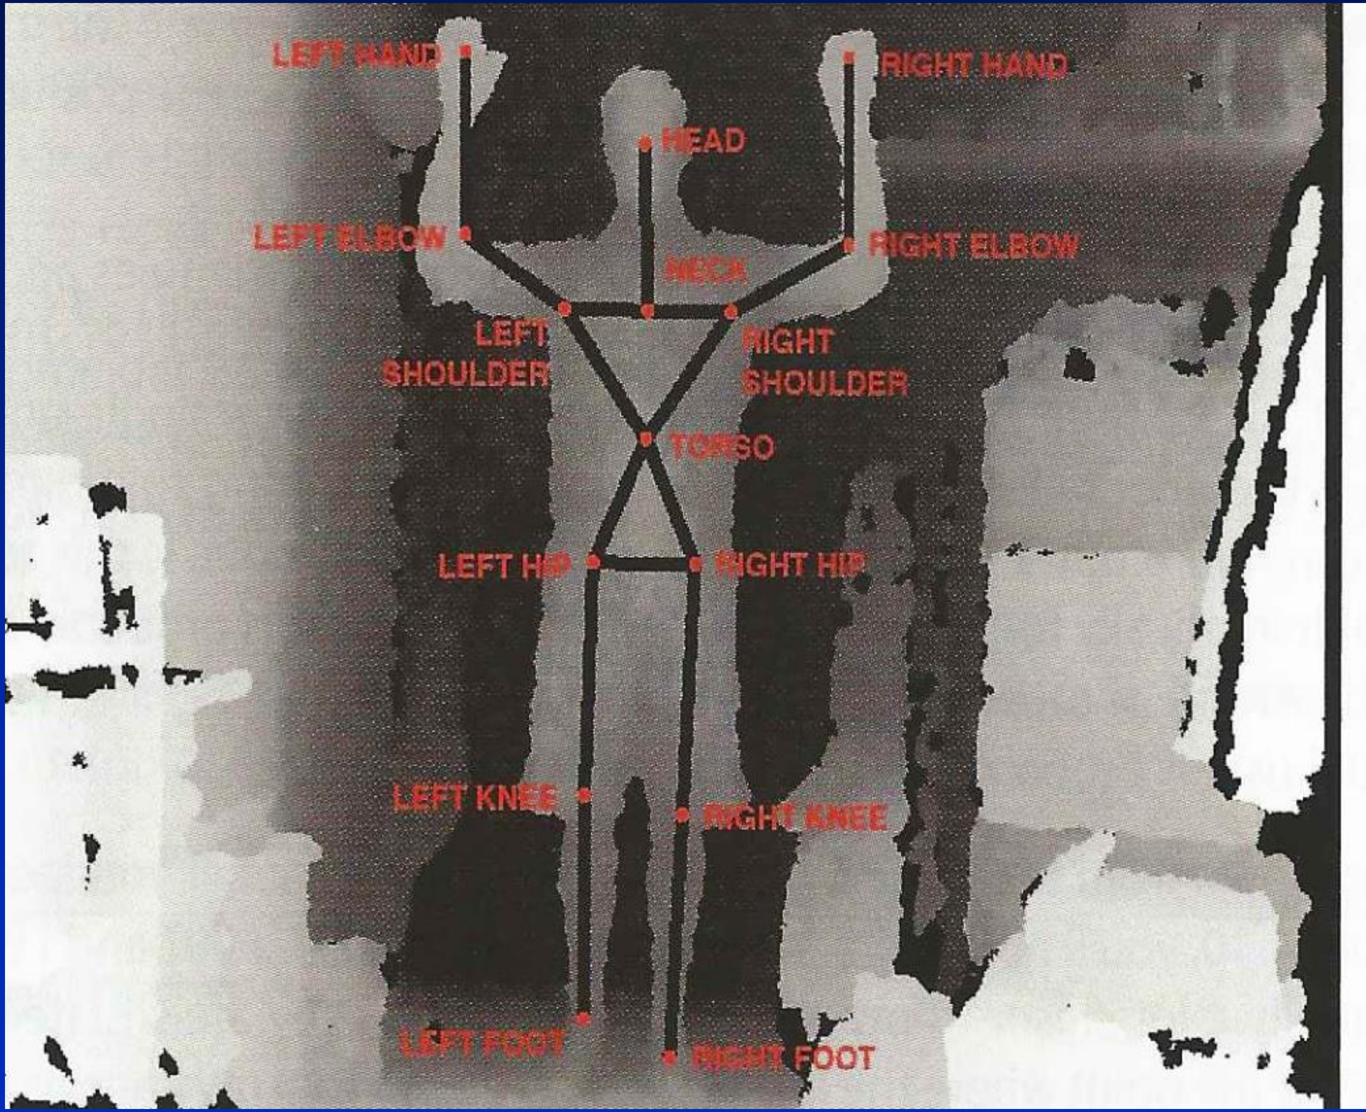

Step 2: Infer a Body Position

Skeleton Manipulation

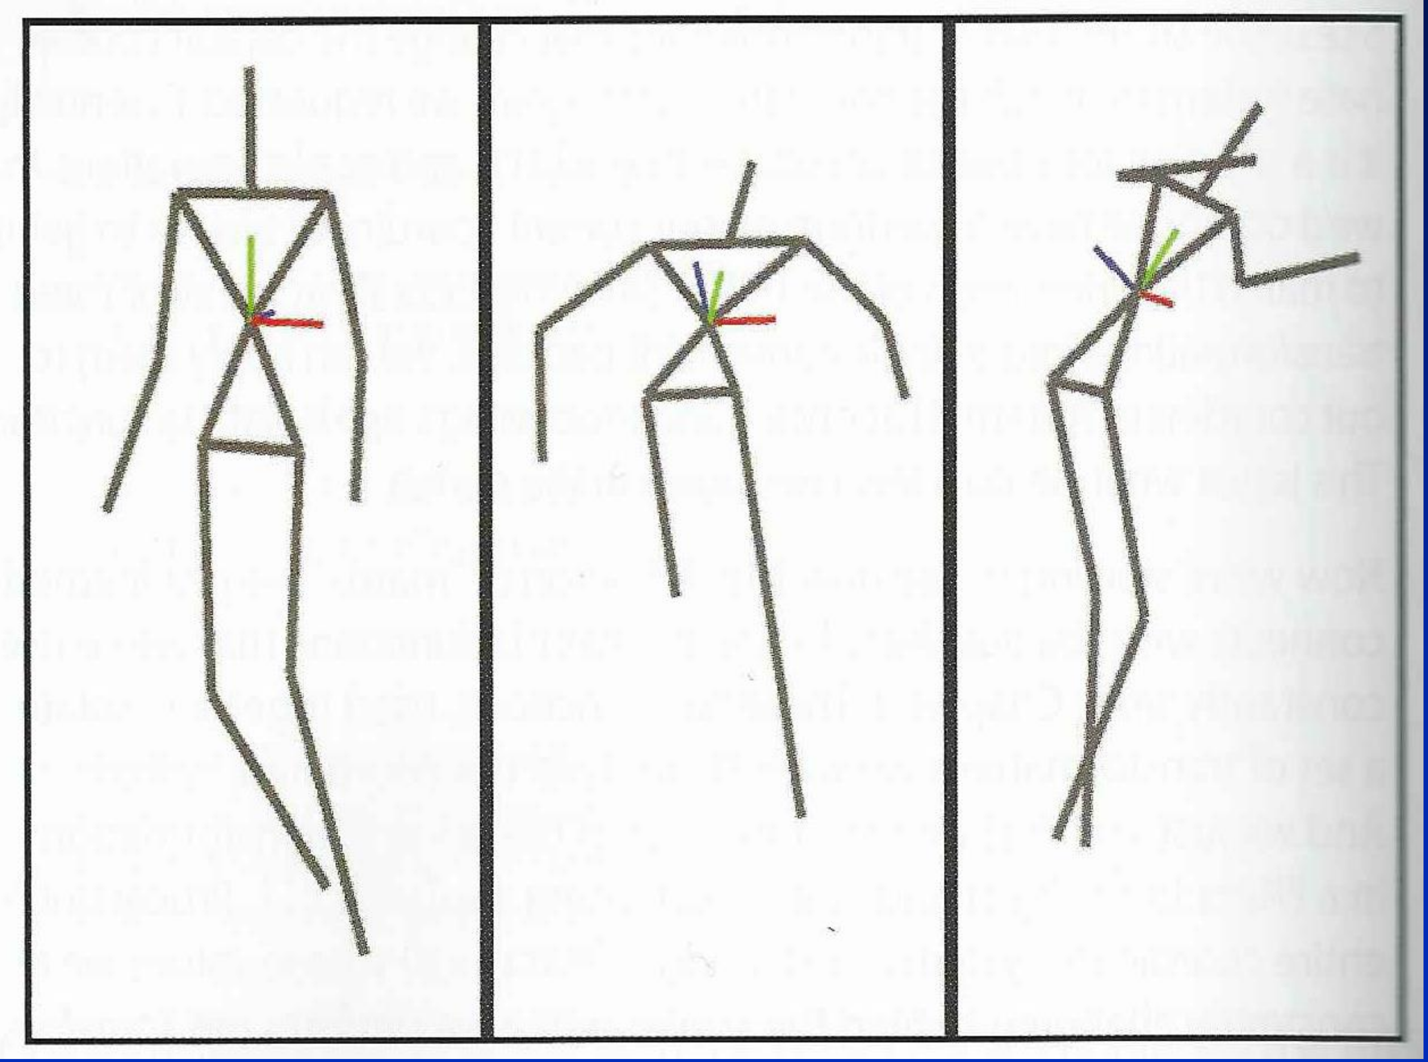

Tracking

The Curious Case of Benjamin Button 2008

Cinefex 116, January 2009

Markerless Motion Capture

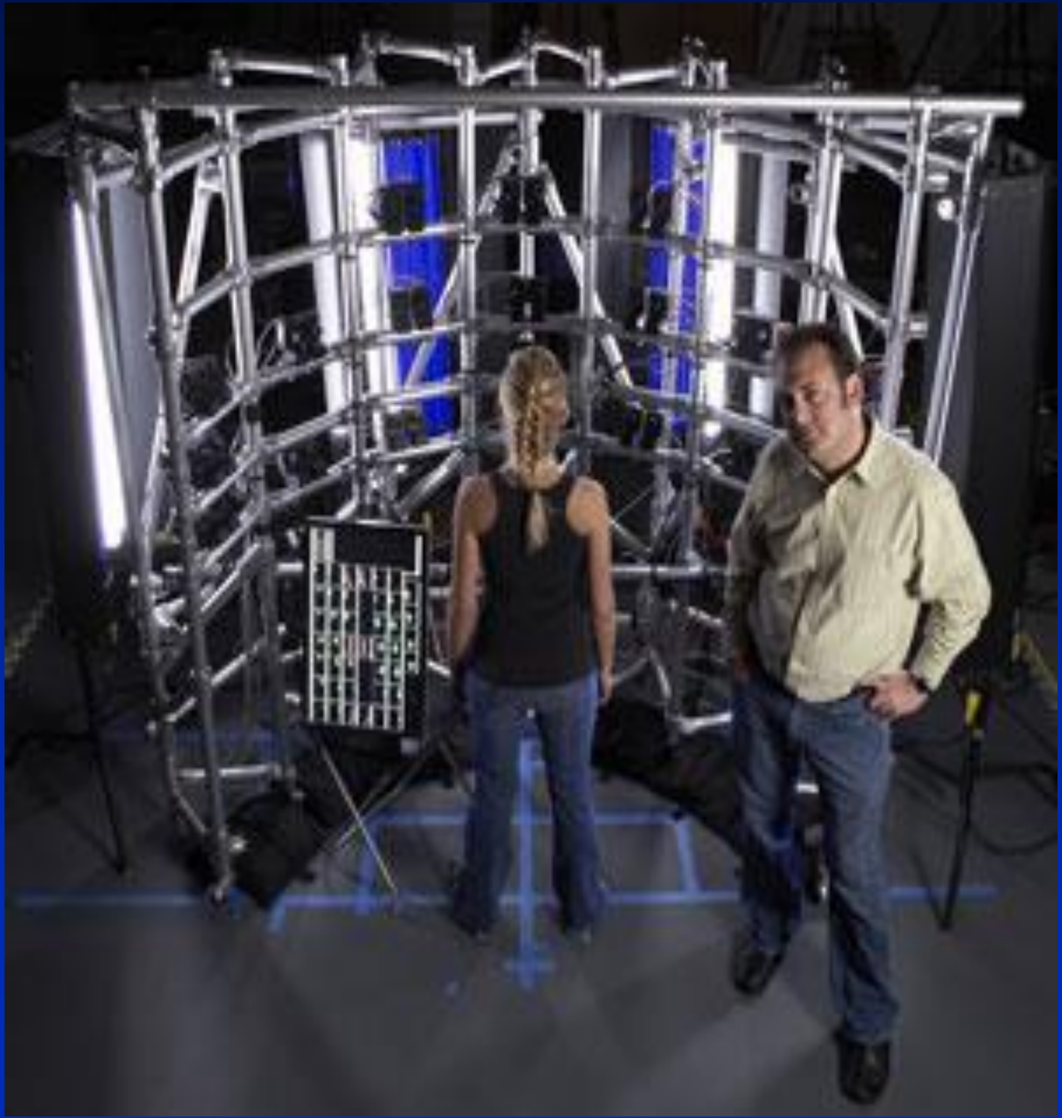

http://macanimationpro.digitalmedianet.com/articles/viewarticle.jsp?id=54623

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

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

Can we detect emotions?

Six Universal Expressions of Emotion

- Paul Ekman

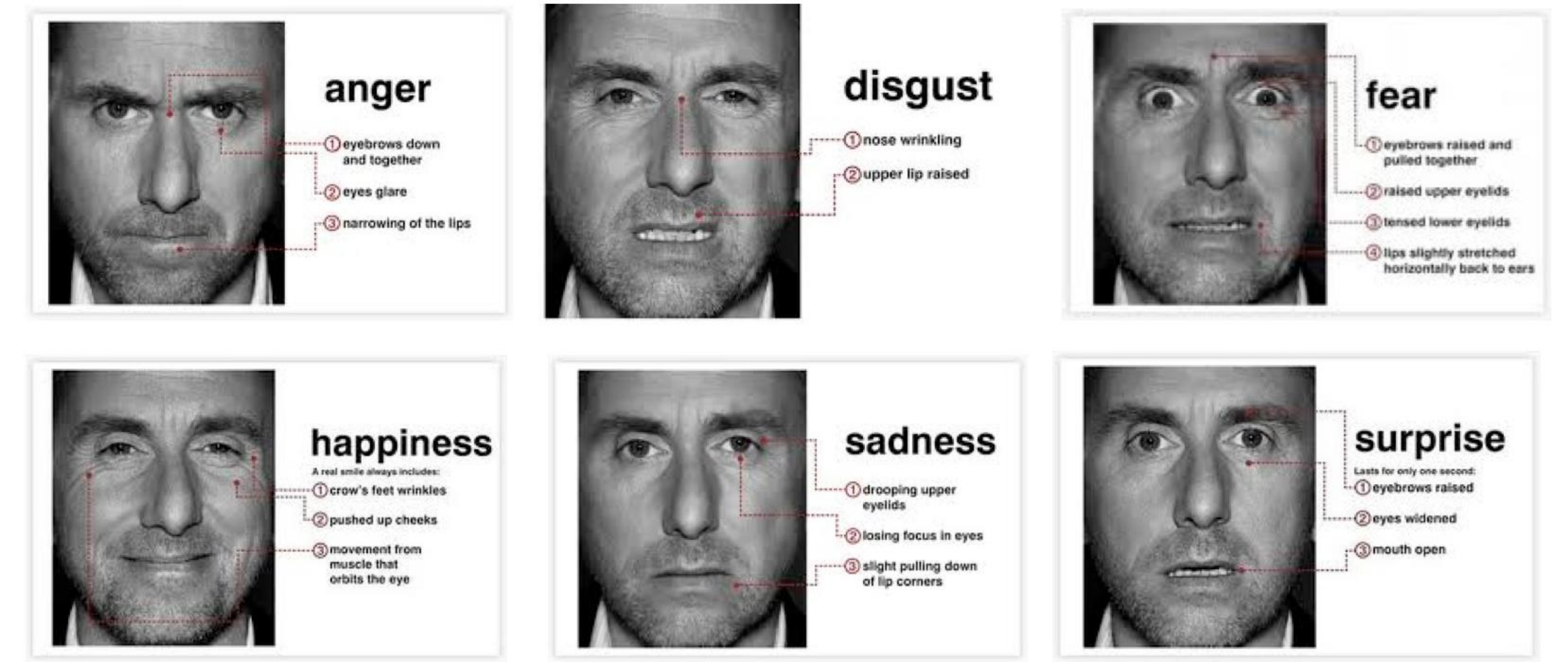

Ekman's Emotion

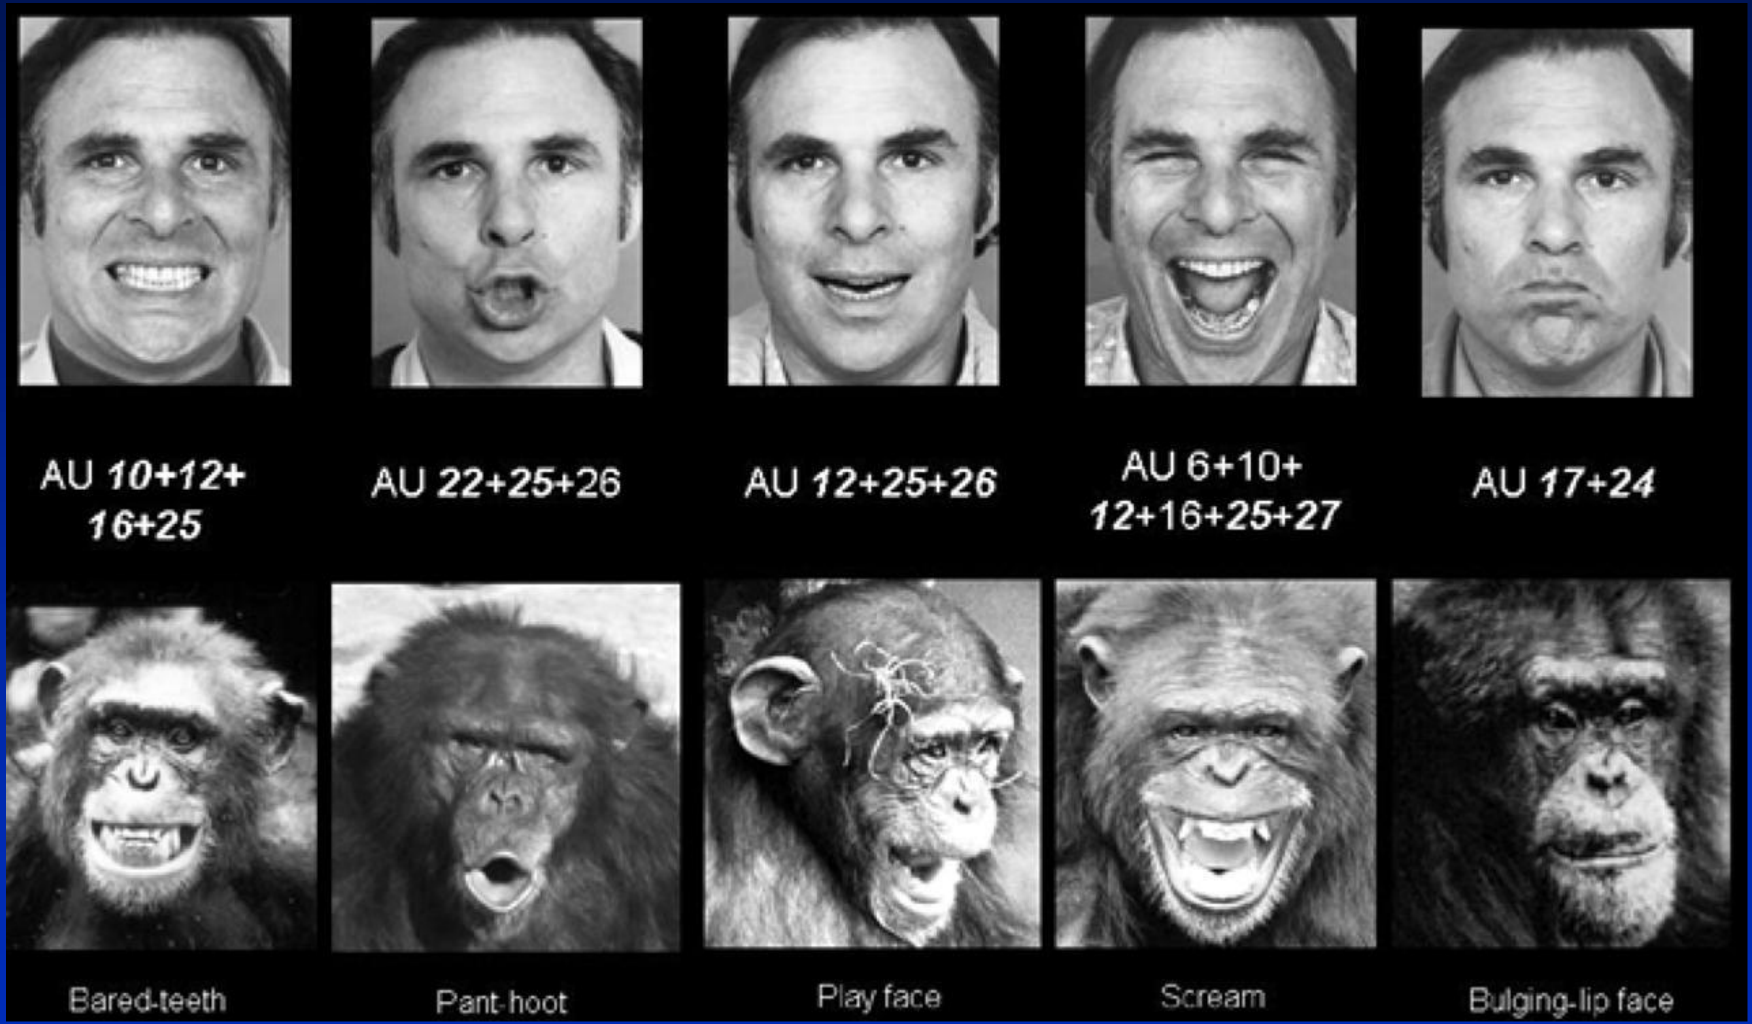

Inside Out

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

Affidex Software

- Scan the image for a face(s) and isolate one.
- Using feature detection algorithms, identify the face's main regions (mouth, nose, eyes, eyebrows, etc.) and ascribe dots to each.
- Separate the dots into "deformable" and "non-deformable" points.
- Deformable points serve as anchors to estimate the magnitude of movement.

Affectiva Computing

- "I think that, ten years down the line, we won't remember what it was like when we couldn't just frown at our device, and our device would say, "Oh, you didn't like that, did you?"
  - Rana el Kaliouby
    Affectiva

End

- “The Art of Studying Moods” (for military strategy)
  - Sun Tzu, *The Art of War*, 600 BC
- “Emotions are physiological and universal in nature”
  - *The Expression of the Emotions in Men and Animals*, Charles Darwin, 1872

Mona Lisa

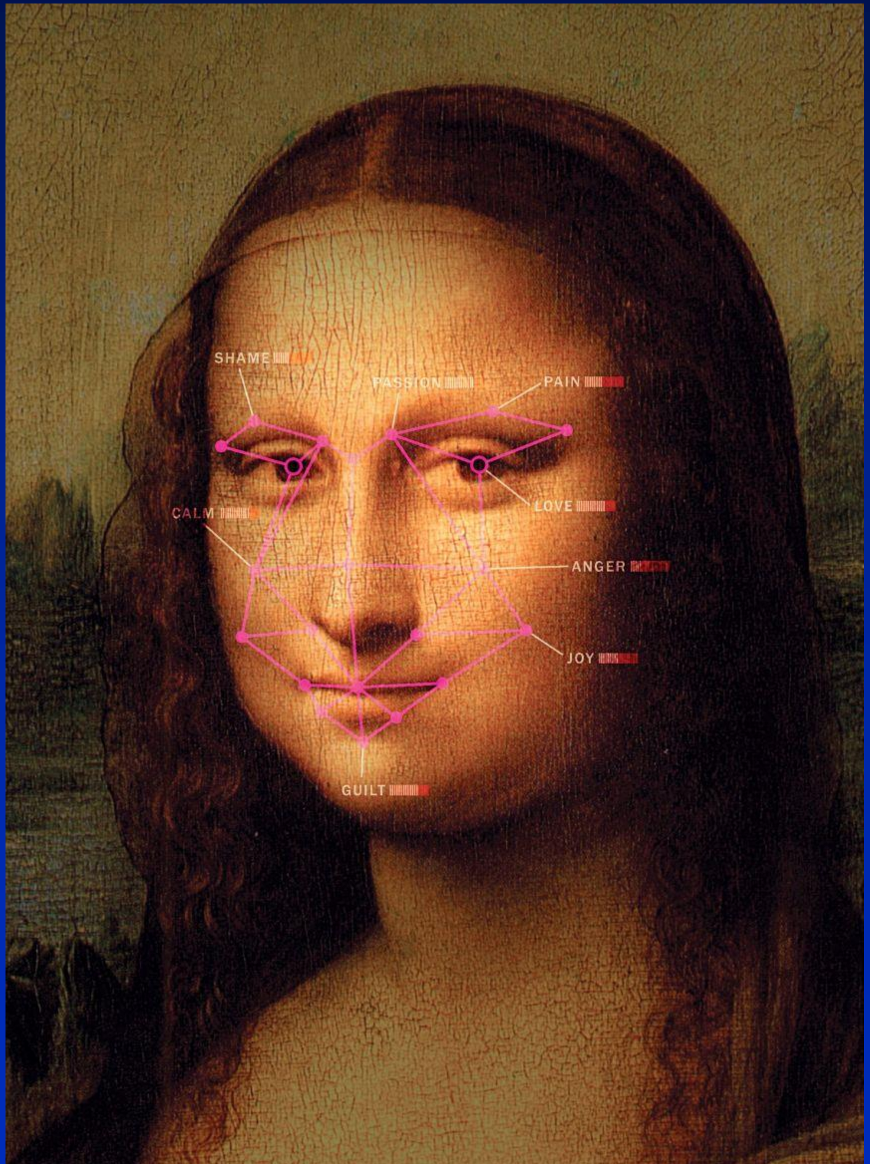

Lidar System

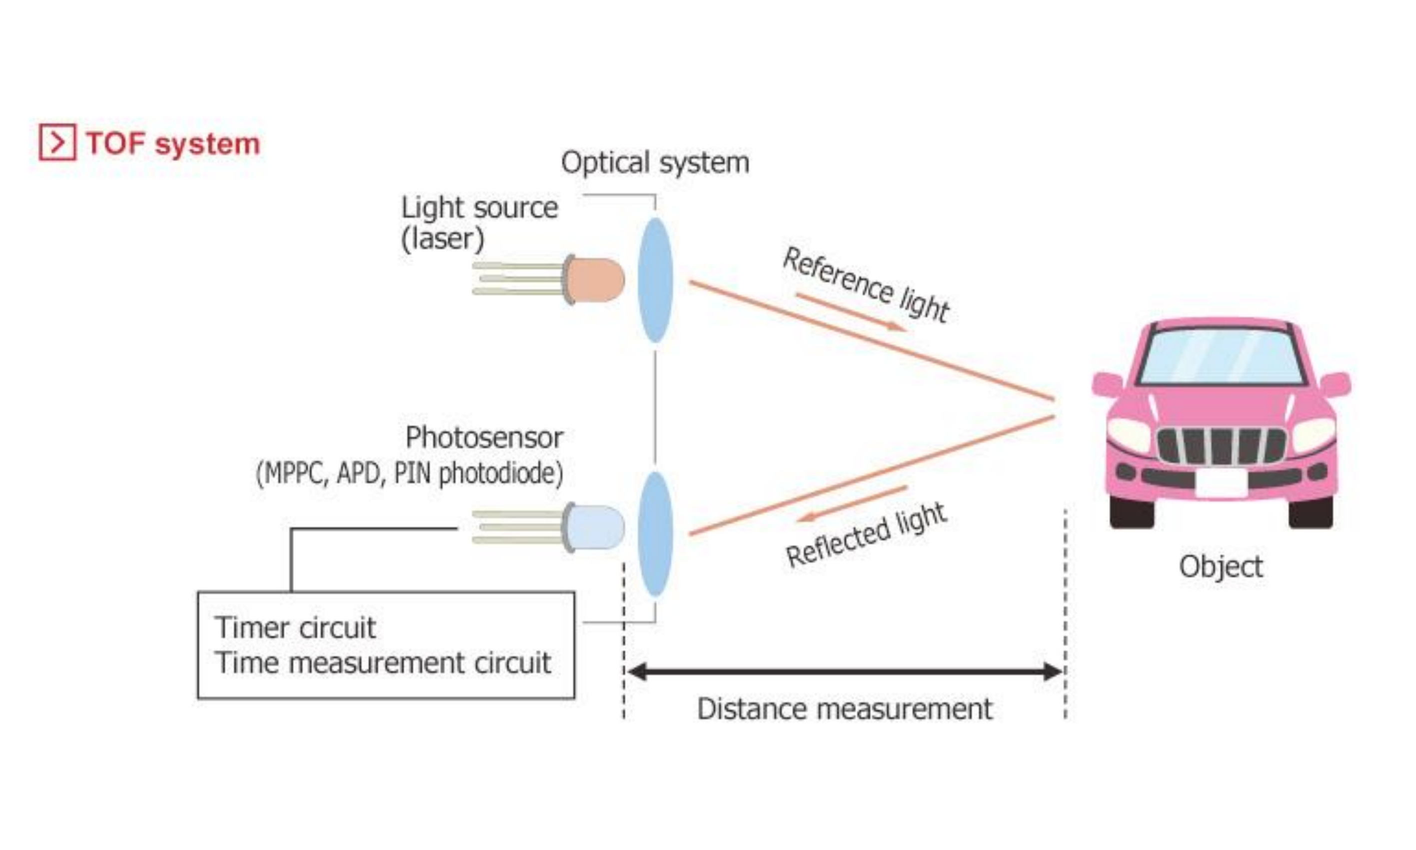

The LIDAR system uses the time of flight a reflection of a of light beam (pulsed or continuous) to sense the surroundings

https://www.microcontrollertips.com/lidar-and-time-of-flight-part-2-operation/

Multiple Sensing Technologies for ADV

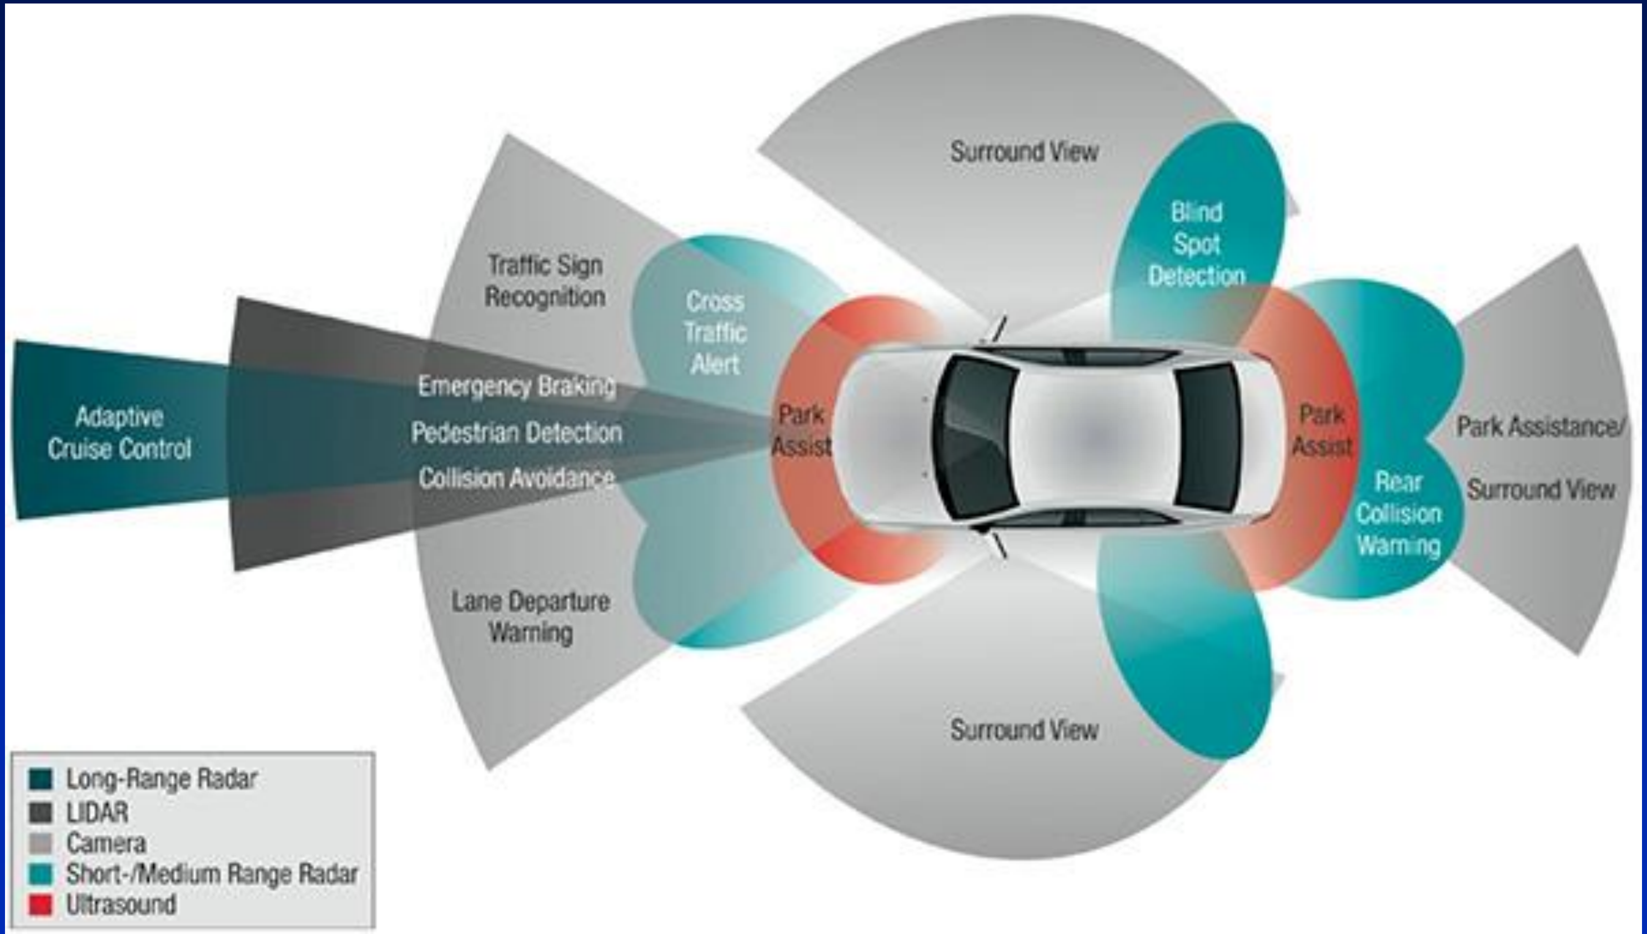

Fully autonomous vehicle will need surrounding- and ambient-sensing and imaging based on multiple technologies

https://www.microcontrollertips.com/lidar-and-time-of-flight-part-1-introduction-faq/

Inside Out

References

- “A basis for deconstructing affect in facial expressions”
  - *Facial Action Coding System*, Paul Eckman, 1978
- “Competitors are beginning to acquire the ability to express and recognize affect.”
  - *Affective Computing*, Rosalind Picard, 1995

Can we compute fast enough for Virtual Reality?

Examples of Computations

Motorola 68030

Intel Pentium