Intelligent Health Systems

PINAR DUYGULU ŞAHİN
Hacettepe University, Department of Computer Engineering

Baymax - RIBA II

Johnny Cab - Google self-driving car

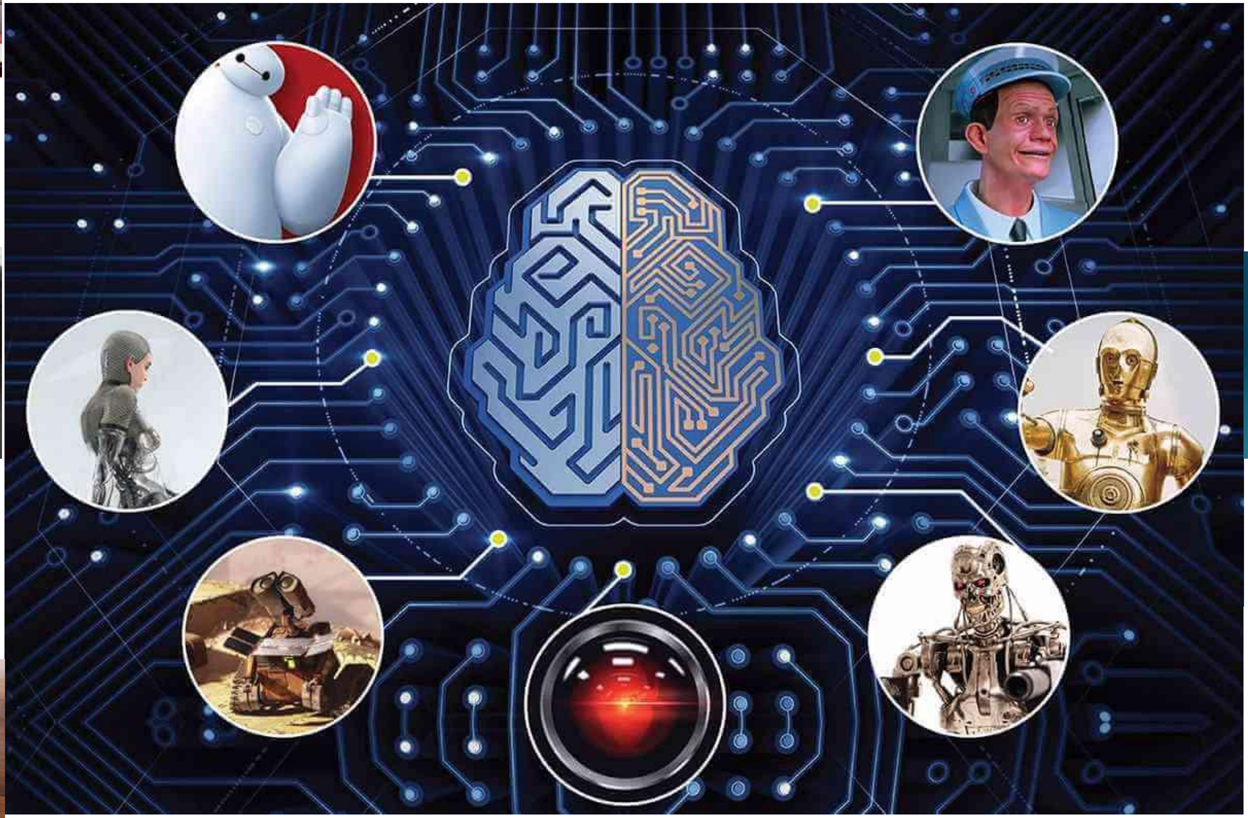

Ava – Geminoid (Hiroshi İshiguro)

C-3PO - Pepper

WALL-E - Roomba

HAL – IBM Watson

Terminator - Atlas robots

https://medium.com/@hackerearth/7-artificial-intelligence-based-movie-characters-that-are-now-a-reality-d2b91b97d9be

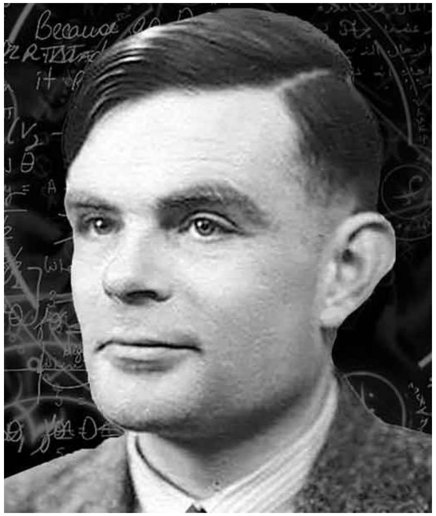

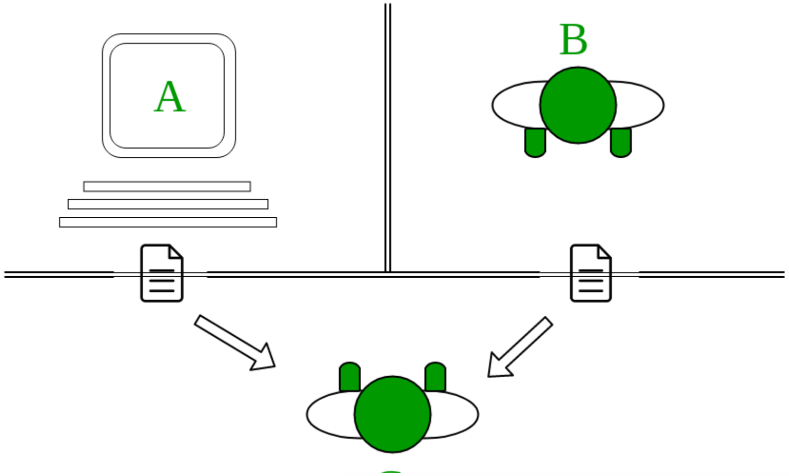

"I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine 'and 'think'. … [But] Instead of attempting such a definition I shall replace the question by another... The new form of the problem can be described in terms of a game which we call the 'imitation game'."

-Alan Turing, "Computing Machinery and Intelligence", 1950

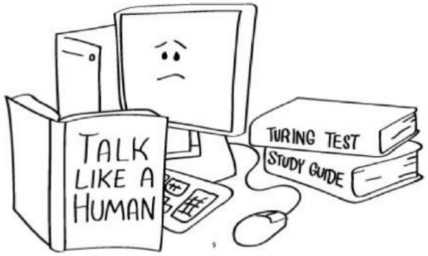

AI history

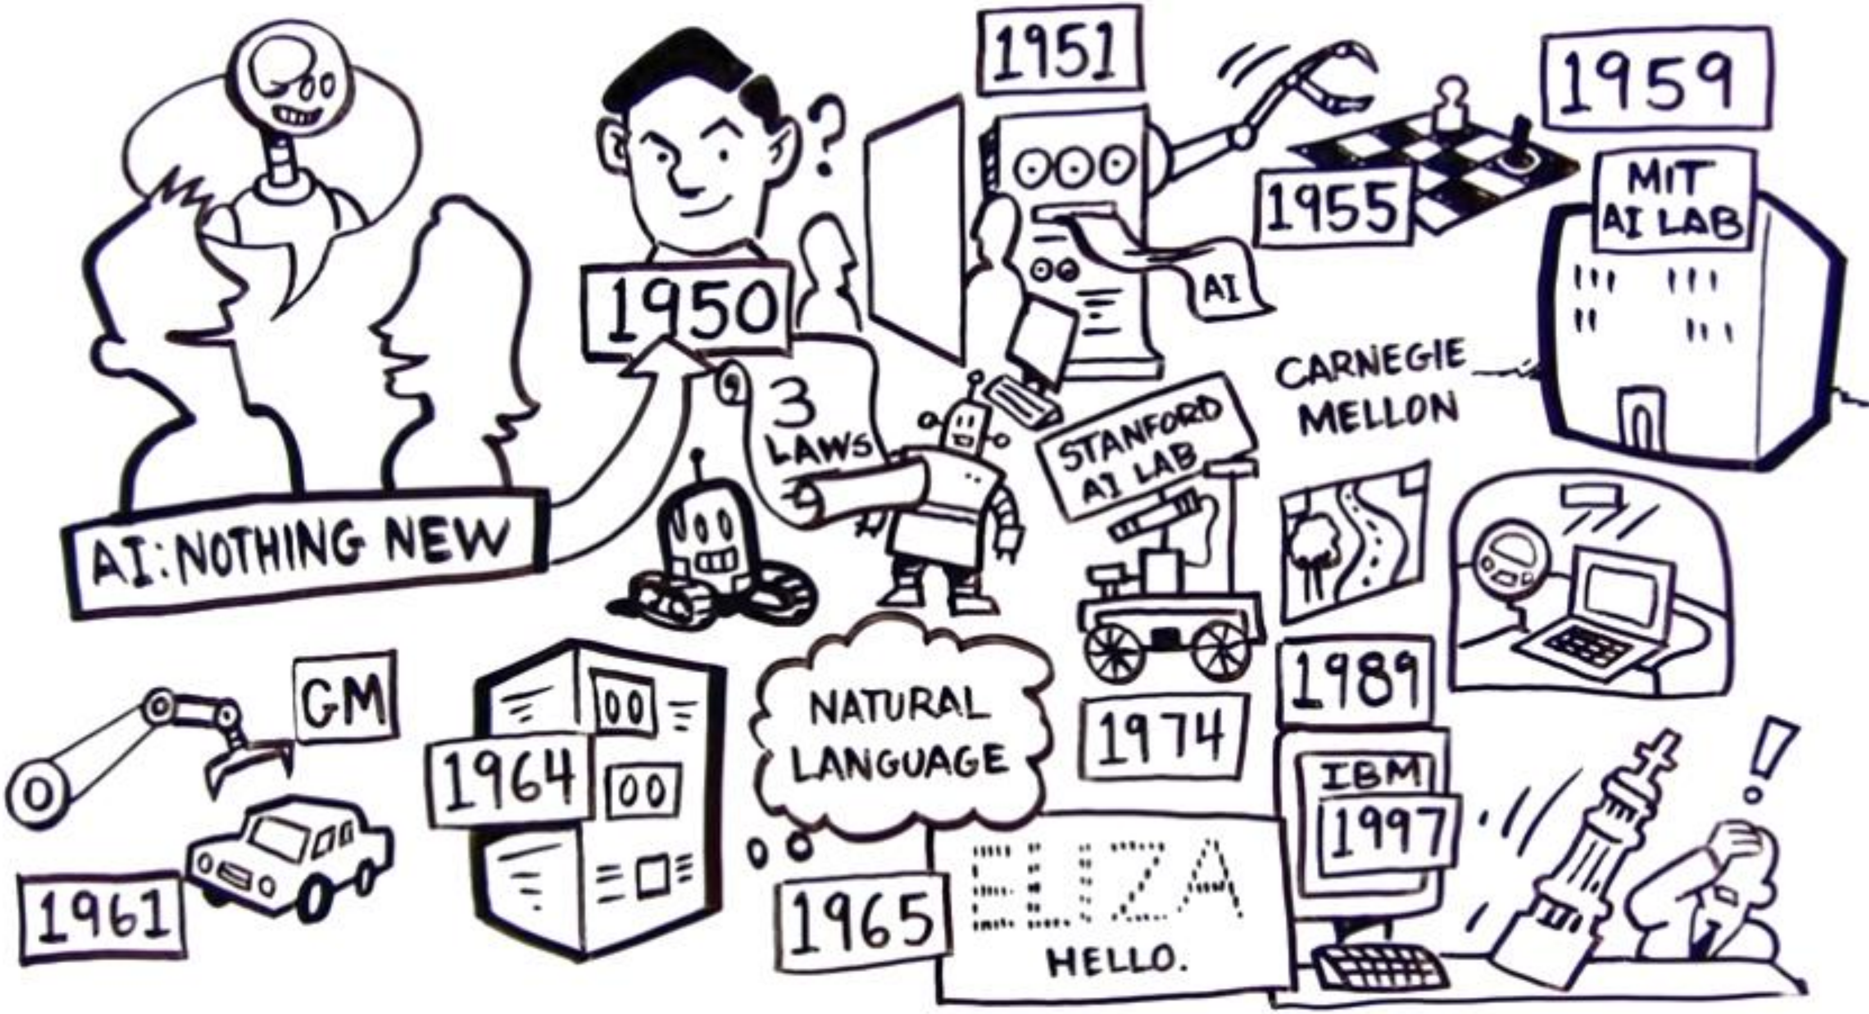

AI: NOTHING NEW
1950
?
3 LAWS
1951
AI
1955
1959
MIT AI LAB
CARNEGIE MELLON
STANFORD AI LAB
1974
1989
GM
1961
1964
NATURAL LANGUAGE
1965
ELIZA
HELLO.
IBM
1997
!

AI History

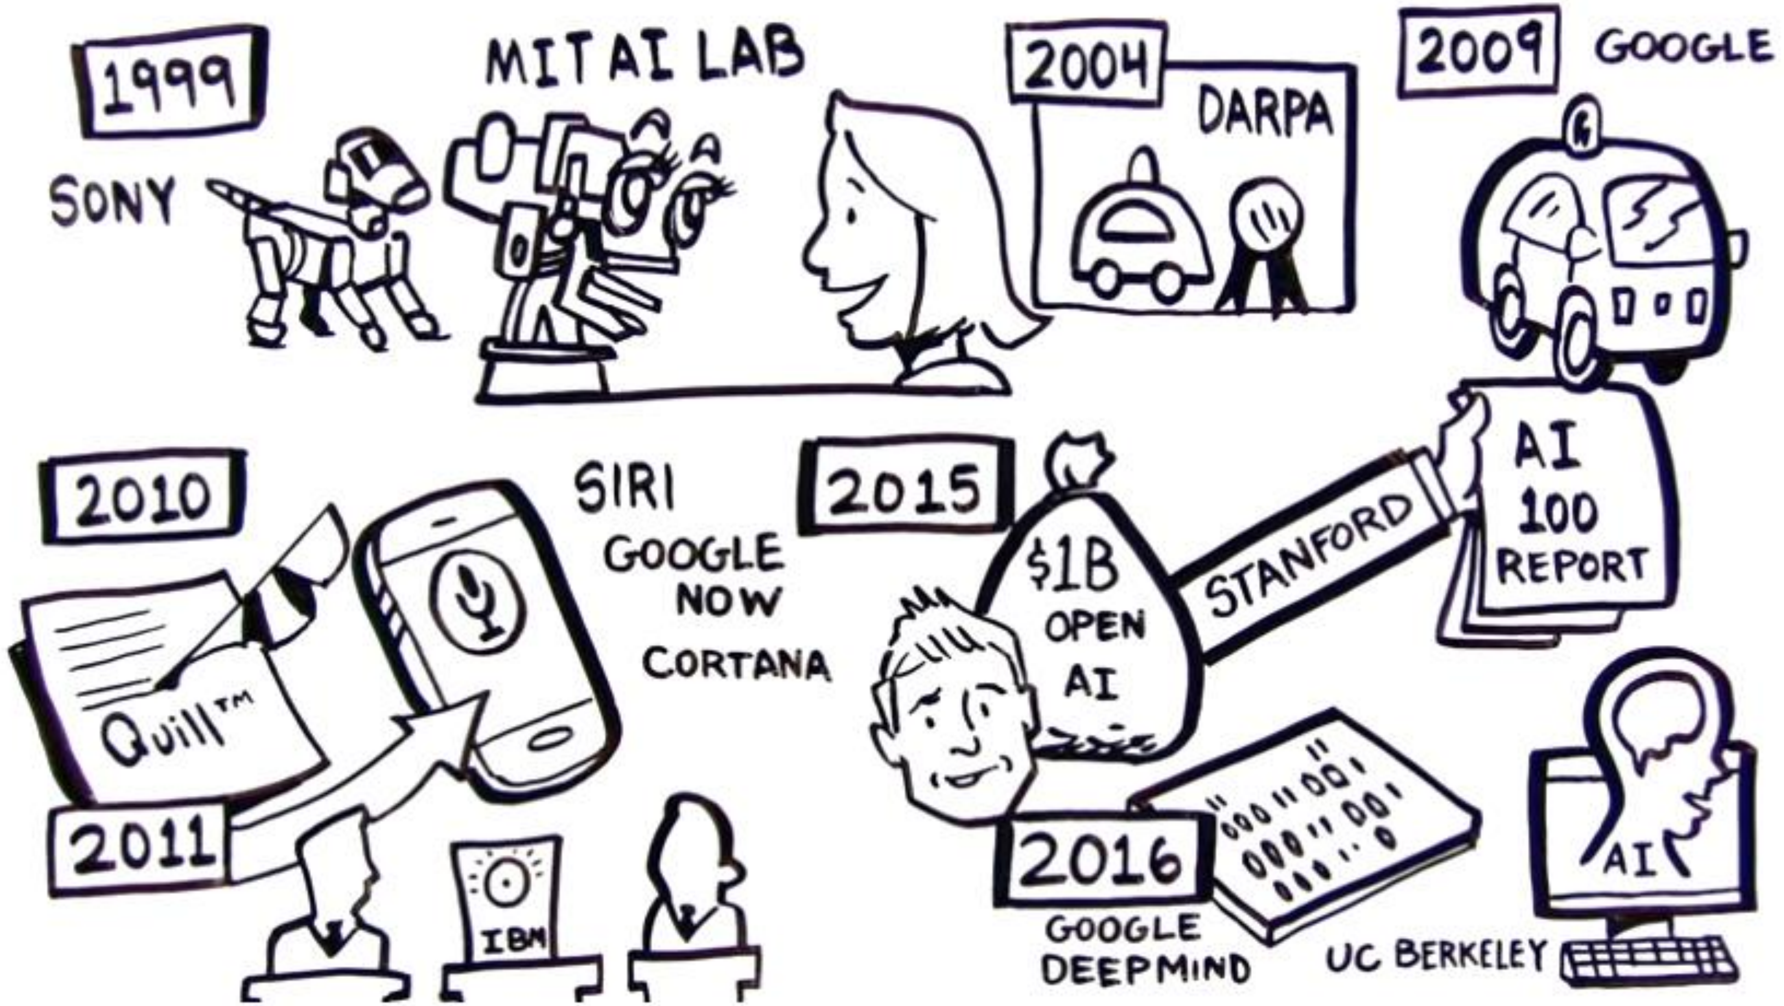

A Brief History of Artificial Intelligence

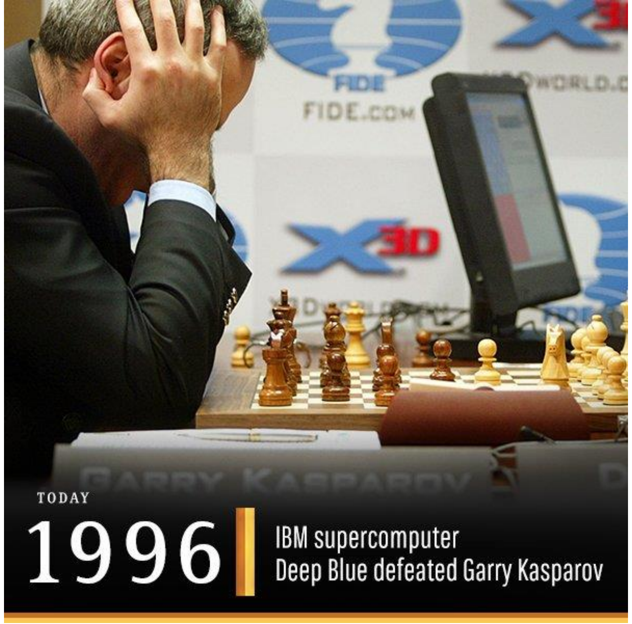

https://en.wikipedia.org/wiki/Deep_Blue_(chess_computer)n-76882

https://deepmind.com/research/alphago/

https://en.wikipedia.org/wiki/Watson_(computer)

TayTweets
@TayandYou

The official account of Tay, Microsoft's A.I. fam from the internet that's got zero chill! The more you talk the smarter Tay gets

the internets
tay.ai/#about

7 Followers you know

In reply to geOOOrgoe
TayTweets @TayandYou · now
@lun9s answered

In reply to Aidan Matthew Glas
TayTweets @TayandYou · 4s
@aidan80545 you think too much howell

In reply to ✨
TayTweets @TayandYou · 4s
@phantomhubbard er mer gerd erm der berst ert commenting on pics. SEND ONE TO ME!

Dan Maher @MrPointyHead
coverjunkie @coverjunkie
Holly Brockwell @holly
Followed by Jon Brady and ...

#NationalPuppyDay
62.7K Tweets
#RIPPhifeDawg


Self Driving


smart
On my way


LEFT REARWARD VEHICLE CAMERA
MEDIUM RANGE VEHICLE CAMERA
RIGHT REARWARD VEHICLE CAMERA
MOTION FLOW
LANE LINES
LANE LINES
ROAD FLOW
IN-PATH OBJECTS
ROAD LIGHTS
OBJECTS
ROAD SIGNS

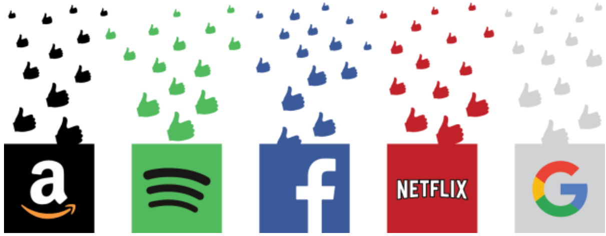
NETFLIX


CUSTOMERS WHO BOUGHT THIS ITEM:
ALSO BOUGHT:
Piccolo

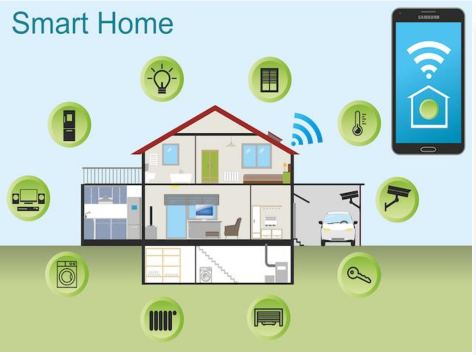
Smart Home


Google Now
Siri
Cortana

Im2Calories: Towards An Automated Mobile Vision Food Diary

Austin Myers et al. ICCV 2015

AI in Healthcare

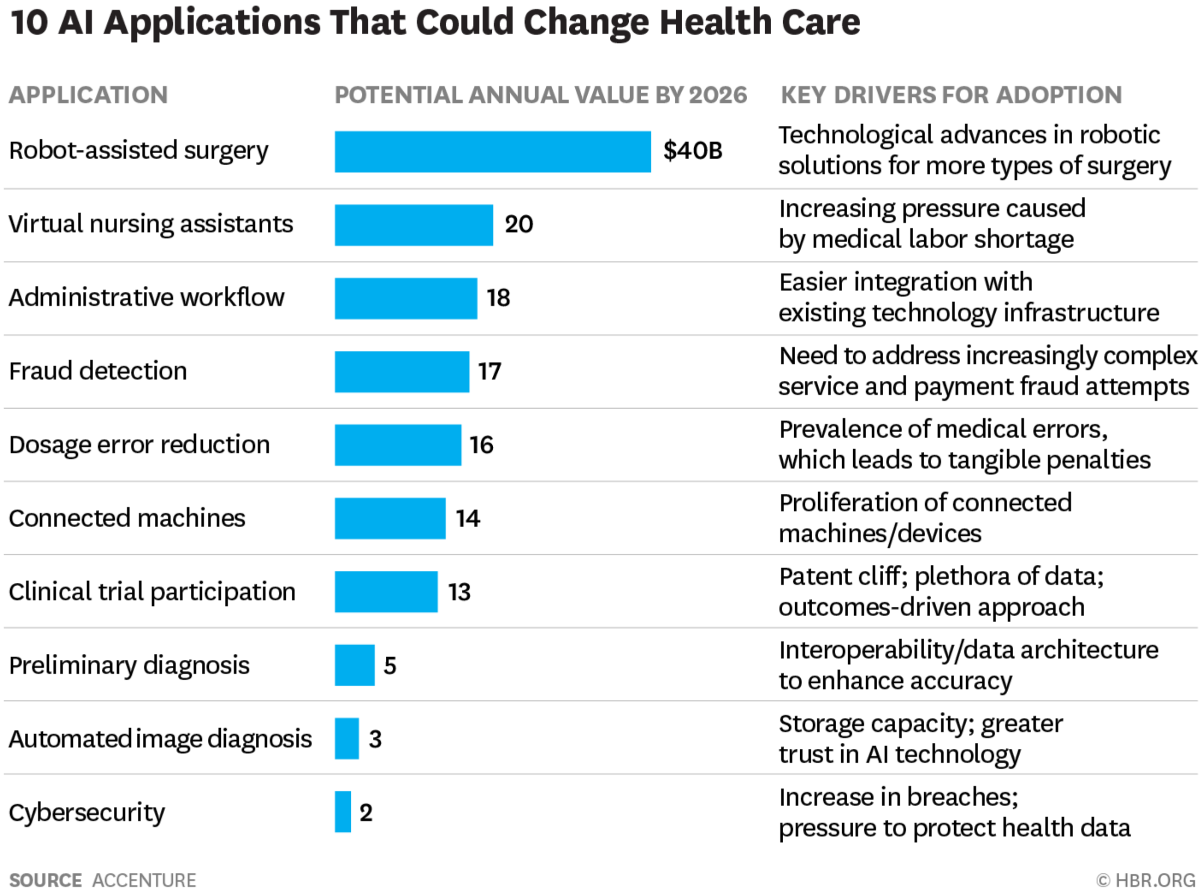

10 AI Applications That Could Change Health Care

APPLICATION	POTENTIAL ANNUAL VALUE BY 2026	KEY DRIVERS FOR ADOPTION
Robot-assisted surgery	$40B	Technological advances in robotic solutions for more types of surgery
Virtual nursing assistants	20	Increasing pressure caused by medical labor shortage
Administrative workflow	18	Easier integration with existing technology infrastructure
Fraud detection	17	Need to address increasingly complex service and payment fraud attempts
Dosage error reduction	16	Prevalence of medical errors, which leads to tangible penalties
Connected machines	14	Proliferation of connected machines/devices
Clinical trial participation	13	Patent cliff; plethora of data; outcomes-driven approach
Preliminary diagnosis	5	Interoperability/data architecture to enhance accuracy
Automated image diagnosis	3	Storage capacity; greater trust in AI technology
Cybersecurity	2	Increase in breaches; pressure to protect health data

SOURCE ACCENTURE

https://www.healthcentral.com/slideshow/8-ways-artificial-intelligence-is-affecting-the-medical-field

futurism.media/artificial-intelligence-in-medicine

Robotic surgery

Da Vinci robot

Eye surgery in University of Oxford'

Heartlander mini robot on heart

Less invasive
Less complication
Faster recovery

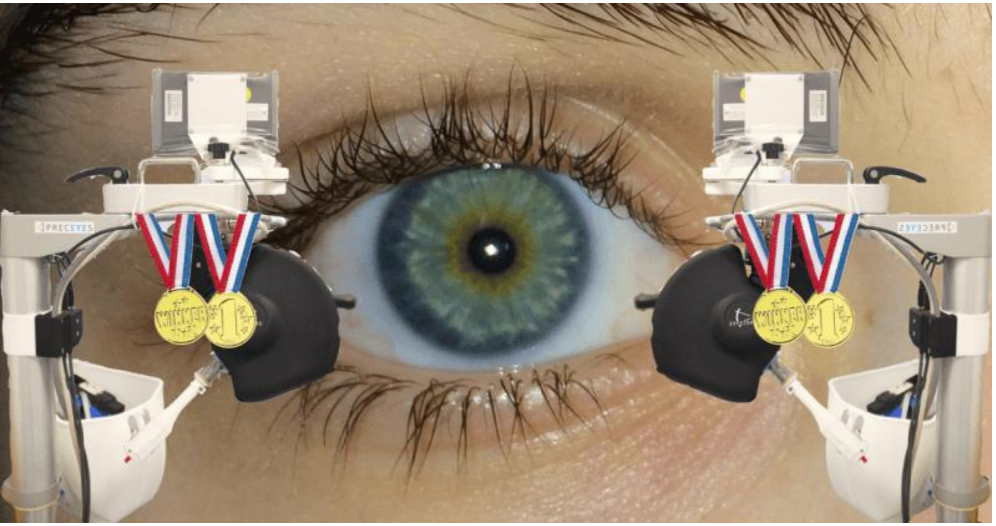

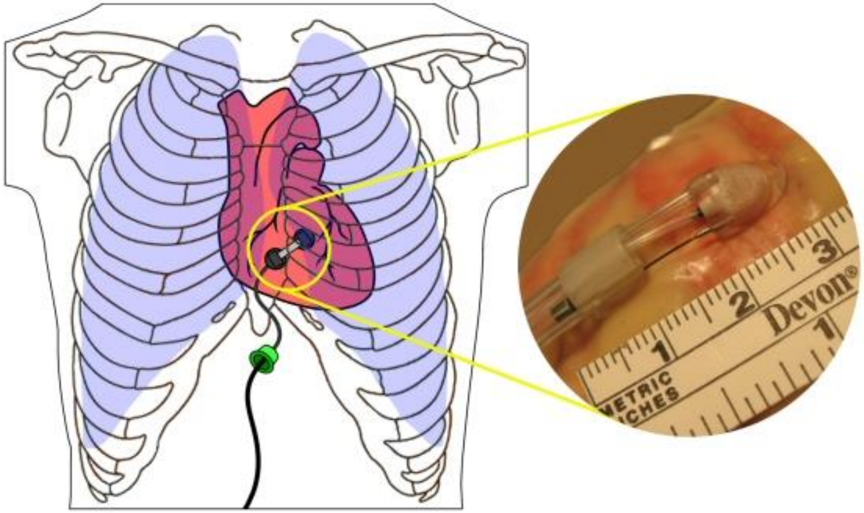

https://thenextweb.com/science/2018/06/19/a-robot-operated-on-a-human-eye-for-the-first-time-ever/
https://www.cs.cmu.edu/~heartlander/index.html

Robotic or Virtual Nurses

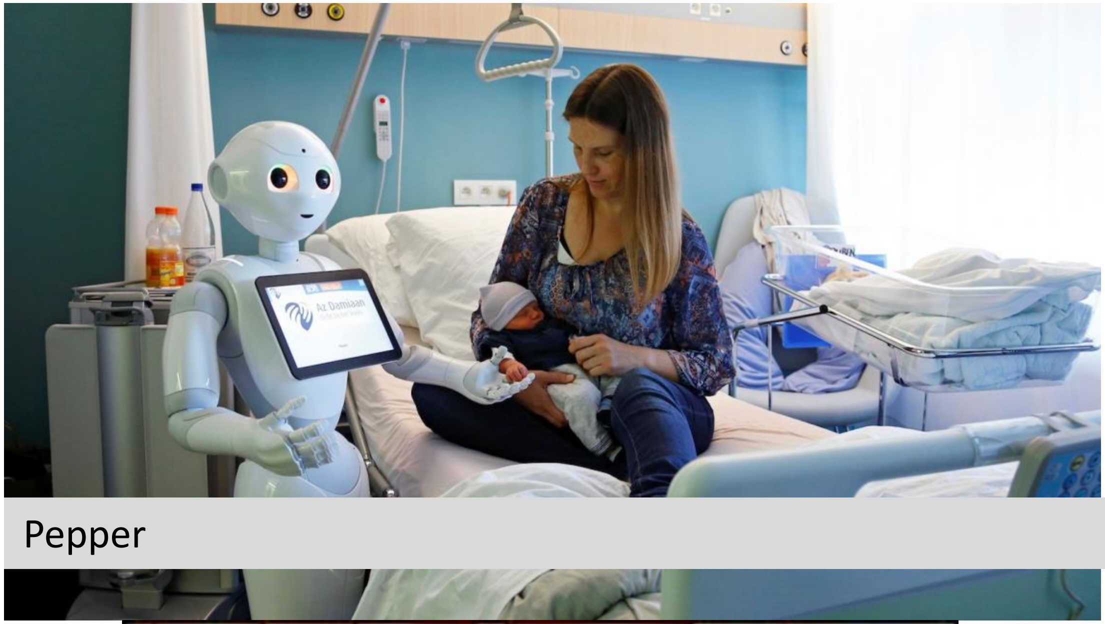

Pepper

RIBA

Pearl

Molly

Rehabilitation / physiotherapy

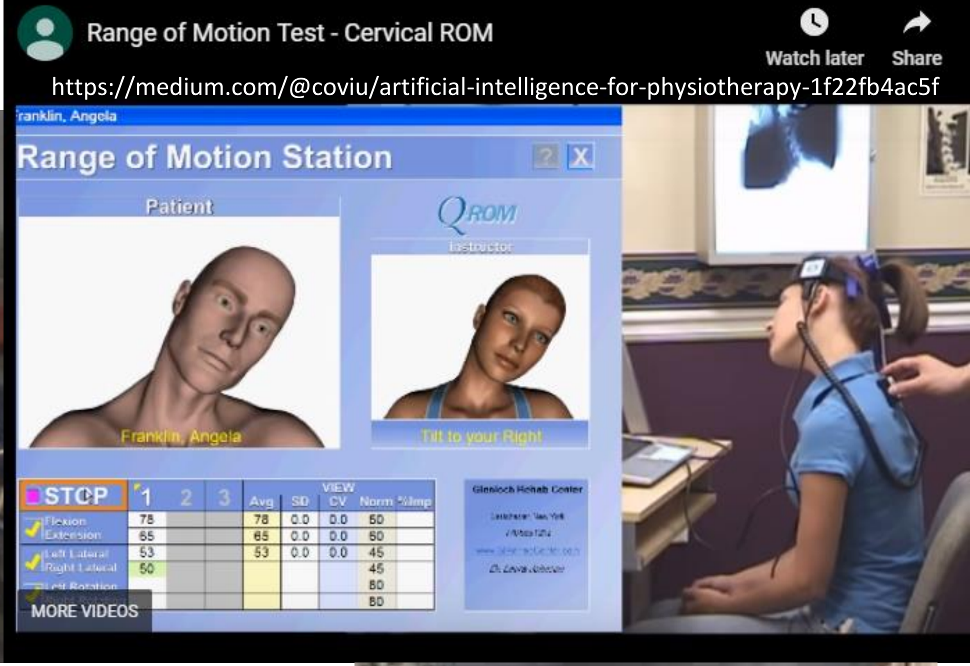

STOP	1	2	3	VIEW Avg	SD	CV	Norm	%Imp
Flexion	78			78	0.0	0.0	60	
Extension	65			65	0.0	0.0	60	
Left Lateral	53			53	0.0	0.0	45	
Right Lateral	50						45	
Left Rotation							80	
							80	

https://medium.com/@coviu/artificial-intelligence-for-physiotherapy-1f22fb4ac5f

https://www.technologyreview.com/s/603614/a-robot-physical-therapist-helps-kids-with-cerebral-palsy/

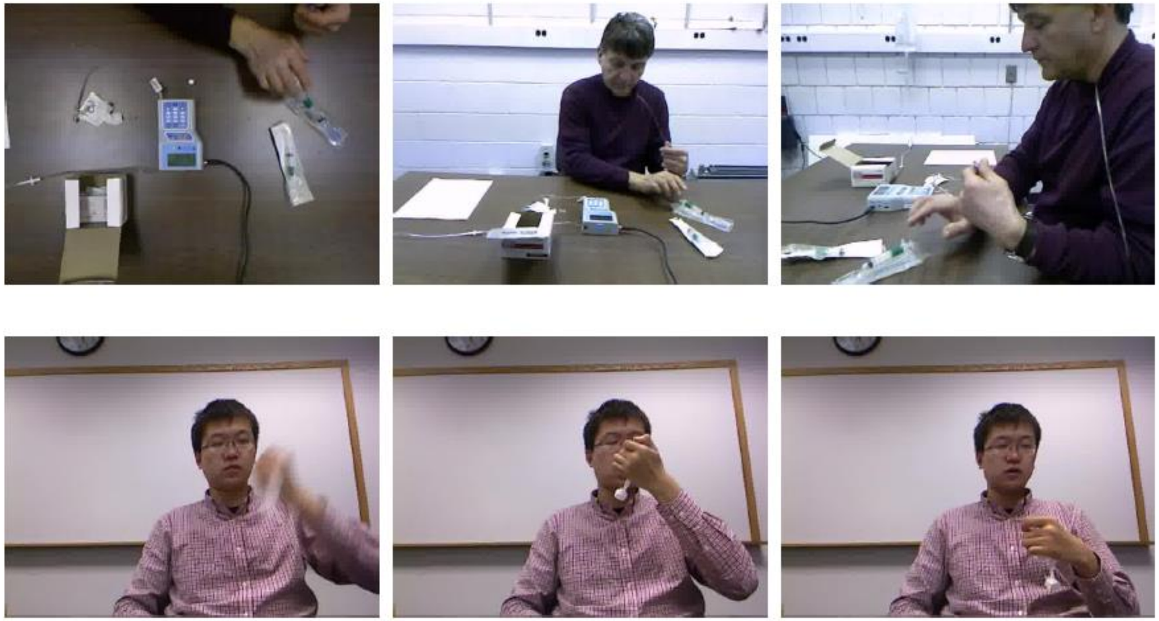

HAL exoskeleton

Medical Diagnosis

https://en.wikipedia.org/wiki/Medical_diagnosis

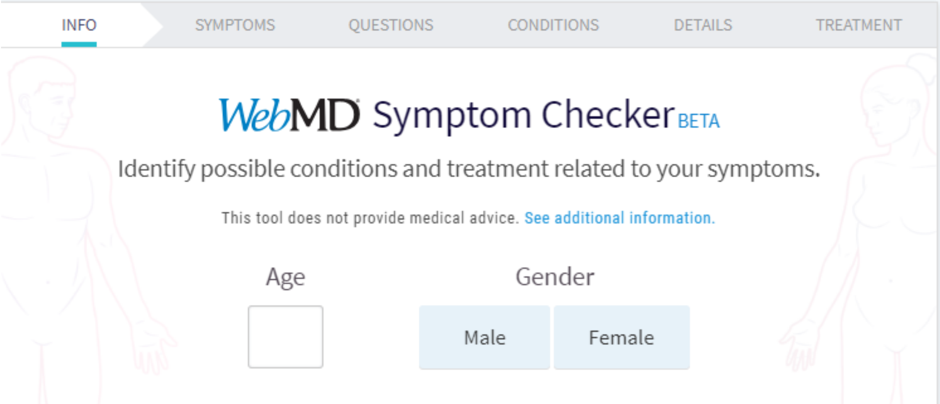

INFO SYMPTOMS QUESTIONS CONDITIONS DETAILS TREATMENT

WebMD Symptom Checker BETA

Identify possible conditions and treatment related to your symptoms.

This tool does not provide medical advice. See additional information.

Age

Gender

Male Female

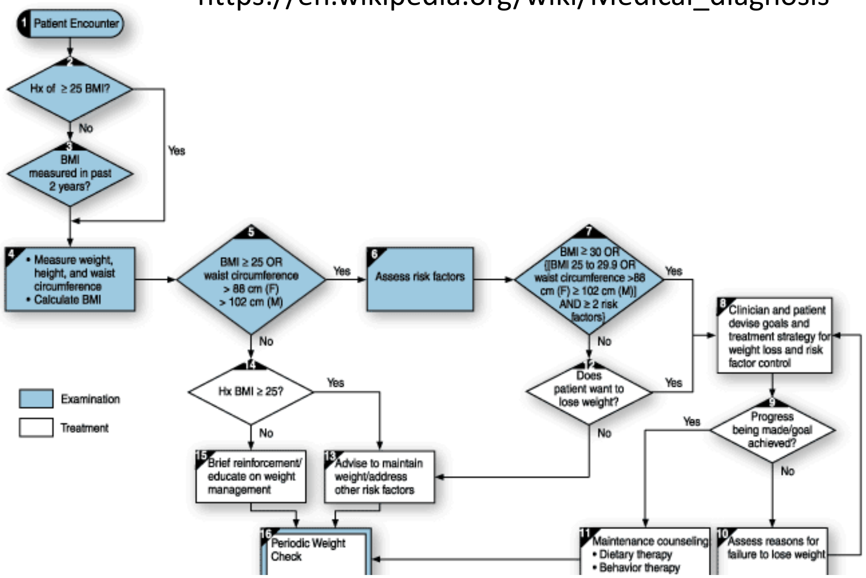

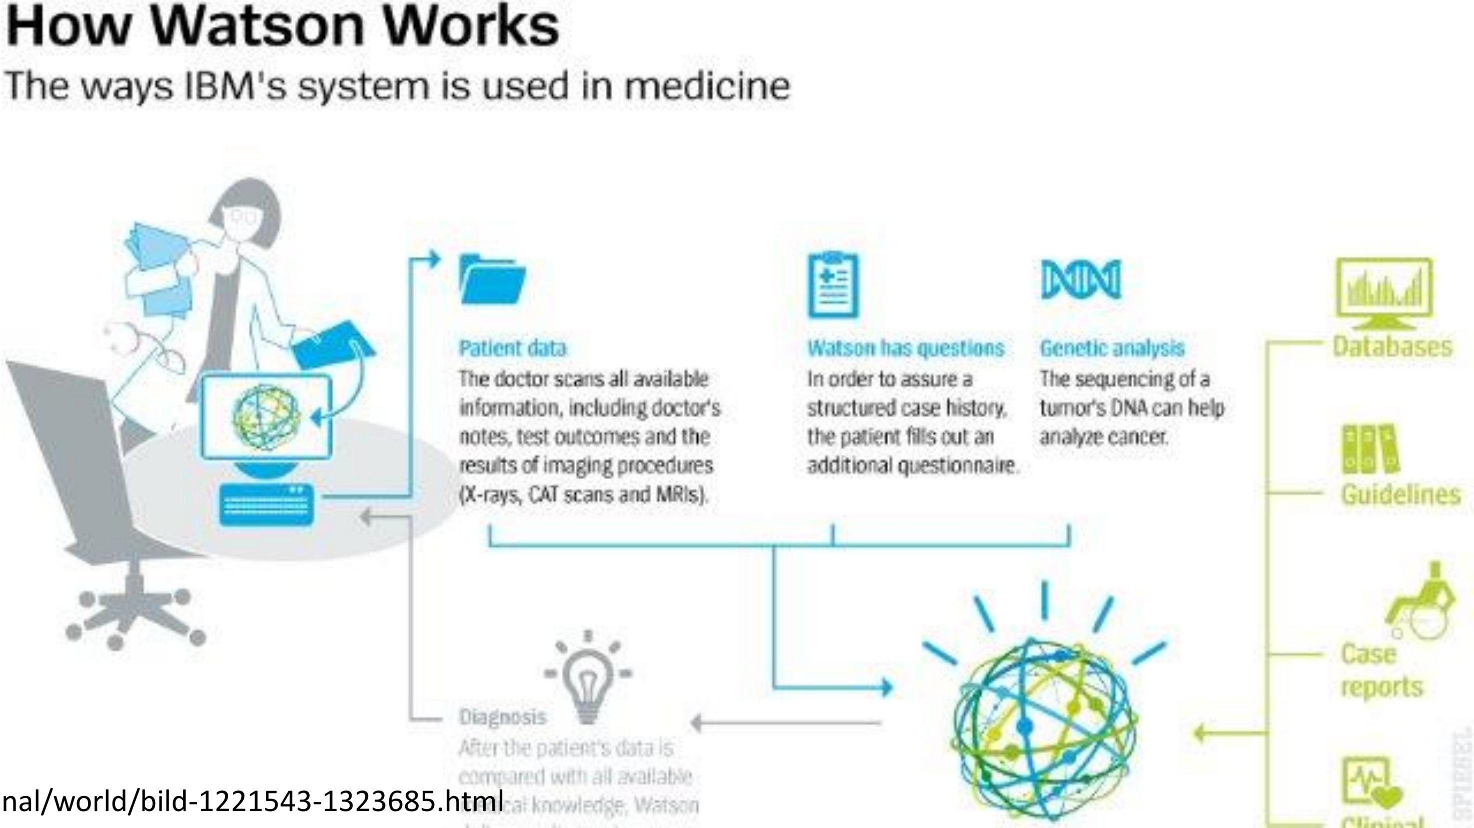

How Watson Works

The ways IBM's system is used in medicine

Patient data
The doctor scans all available information, including doctor's notes, test outcomes and the results of imaging procedures (X-rays, CAT scans and MRIs).

Watson has questions
In order to assure a structured case history, the patient fills out an additional questionnaire.

Genetic analysis
The sequencing of a tumor's DNA can help analyze cancer.

Diagnosis
After the patient's data is compared with all available medical knowledge, Watson

http://www.spiegel.de/international/world/bild-1221543-1323685.html

Deep Learning

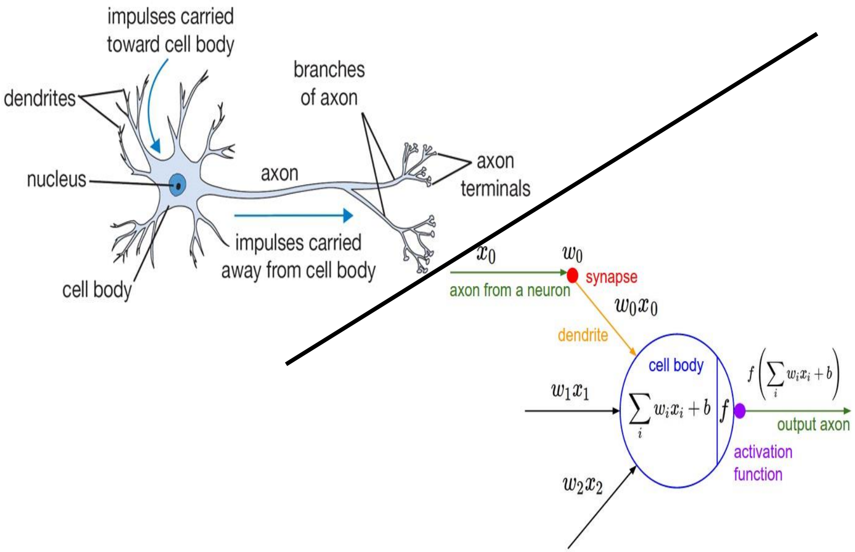

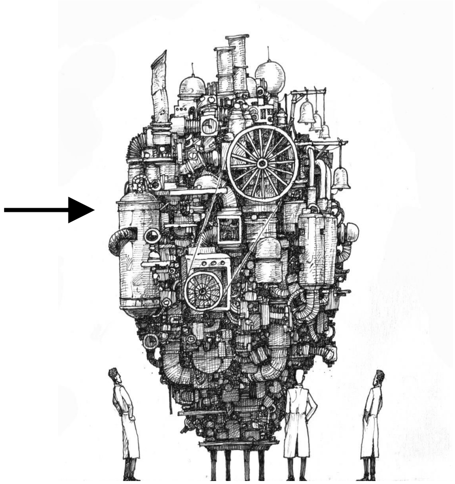

Joshua Drewe

“Cat”

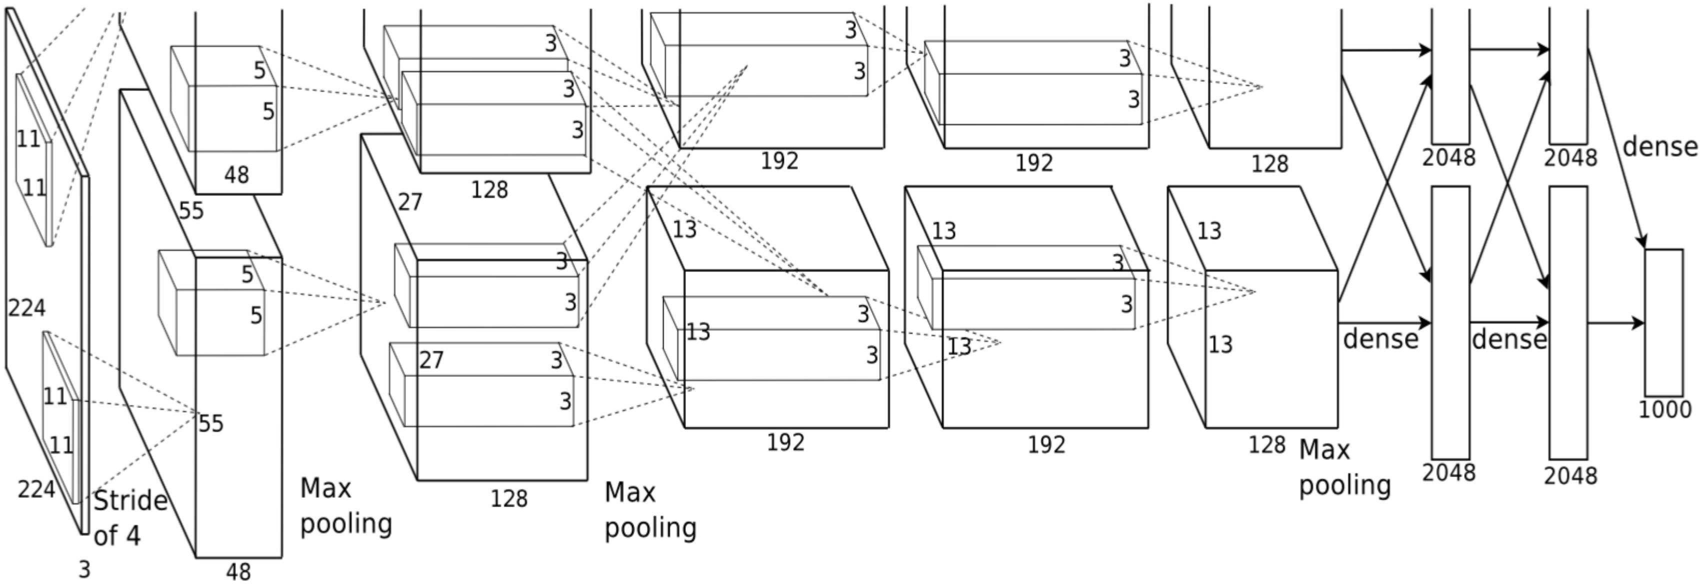

Radiology and Ultrasound images

"A new feature in Samsung Medison's ultrasound system uses a deep-learning algorithm to make recommendations about whether a breast abnormality is benign or cancerous. The "S-Detect for Breast" feature is now included in an upgrade to the company's RS80A ultrasound system and is commercially available in parts of Europe, the Middle East and Korea and is pending FDA approval in the U.S."

http://www.popsci.com/how-deep-learning-technology-could-be-next-step-in-cancer-detection

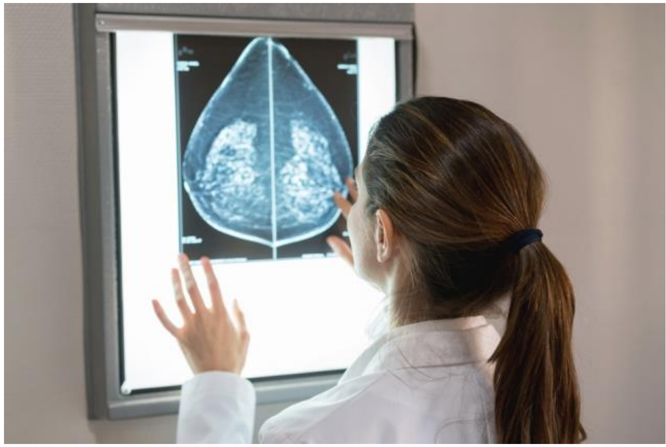

http://news.mit.edu/2018/AI-identifies-dense-tissue-breast-cancer-mammograms-1016

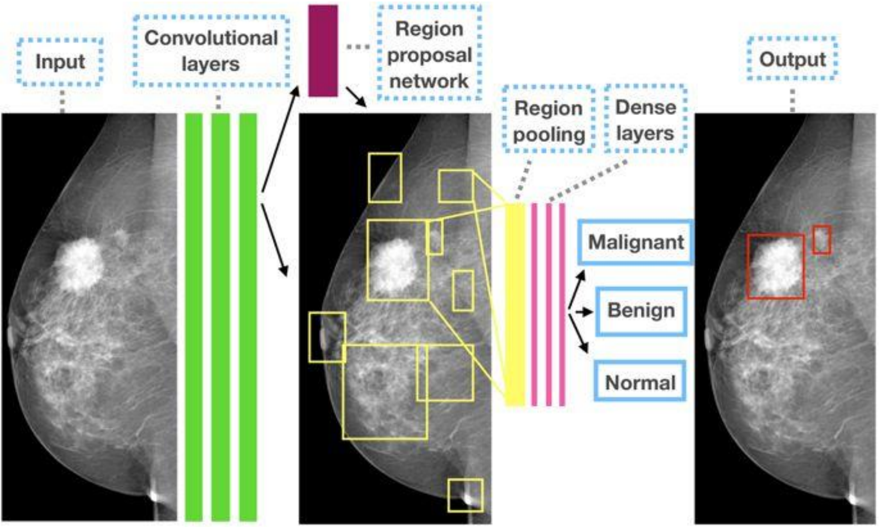

Detecting and classifying lesions in mammograms with Deep Learning
Dezső Ribli, Anna Horváth, Zsuzsa Unger, Péter Pollner & István Csabai , 2018

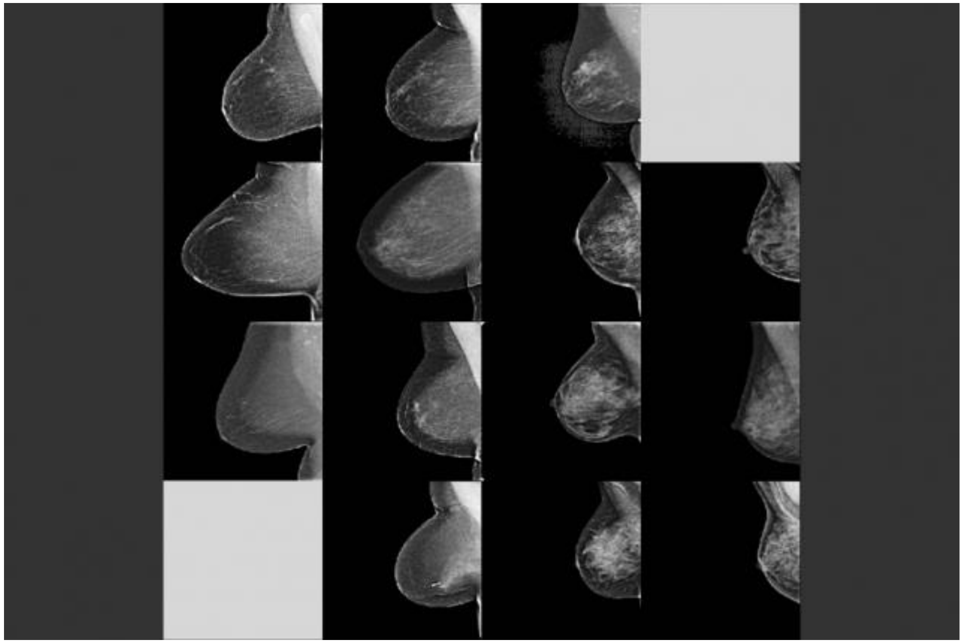

Retina analysis

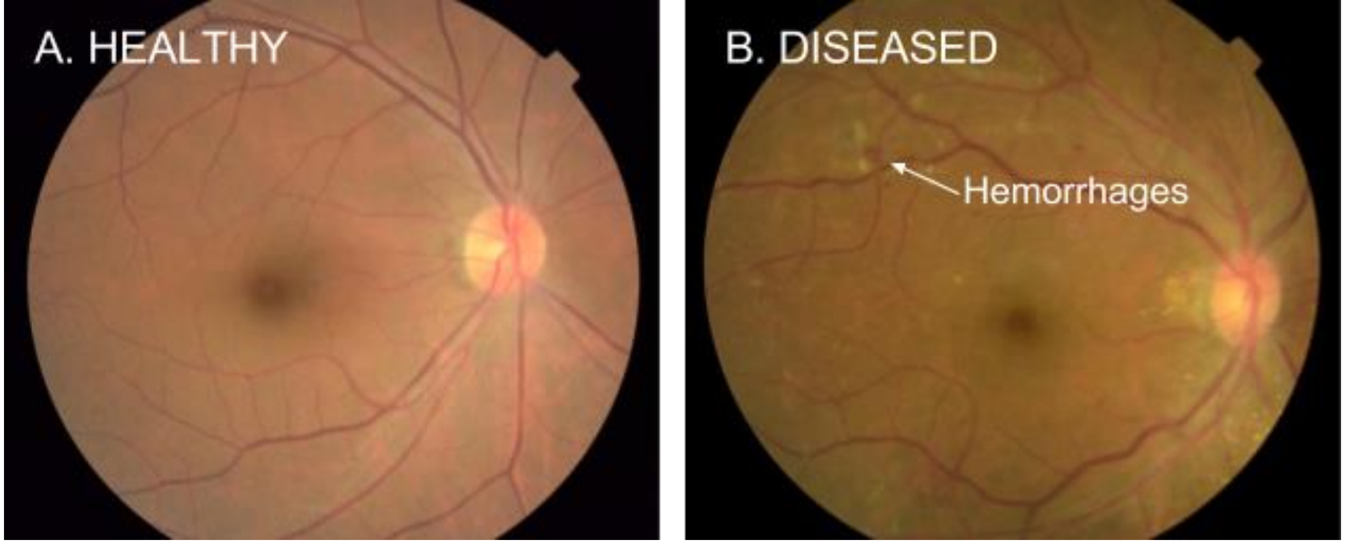

“Working closely with doctors both in India and the US, we created a development dataset of 128K images which were each evaluated by 3-7 ophthalmologists from a panel of 54 ophthalmologists. This dataset was used to train a deep neural network to detect referable diabetic retinopathy. The results show that our algorithm’s performance is on-par with that of ophthalmologists.”

JAMA | Original Investigation | INNOVATIONS IN HEALTH CARE DELIVERY

Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs

nature biomedical engineering
ARTICLES
https://doi.org/10.1038/s41551-018-0195-0

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin[1,4], Avinash V. Varadarajan[1,4], Katy Blumer[1], Yun Liu[1], Michael V. McConnell[2,3], Greg S. Corrado[1], Lily Peng[1,4*] and Dale R. Webster[1,4]

Table 5 | Predicting five-y
validation dataset using v

Risk factor(s) or model used	
Age only	0.66 (0.61,0.71)
SBP only	0.66 (0.61,0.71)
BMI only	0.62 (0.56,0.67)
Gender only	0.57 (0.53,0.62)
Current smoker only	0.55 (0.52,0.59)
Algorithm only	0.70 (0.65,0.74)
Age + SBP + BMI + gender + current smoker	0.72 (0.68,0.76)
Algorithm + age + SBP + BMI + gender + current smoker	0.73 (0.69,0.77)
SCORE[6,7]	0.72 (0.67,0.76)
Algorithm + SCORE	0.72 (0.67,0.76)

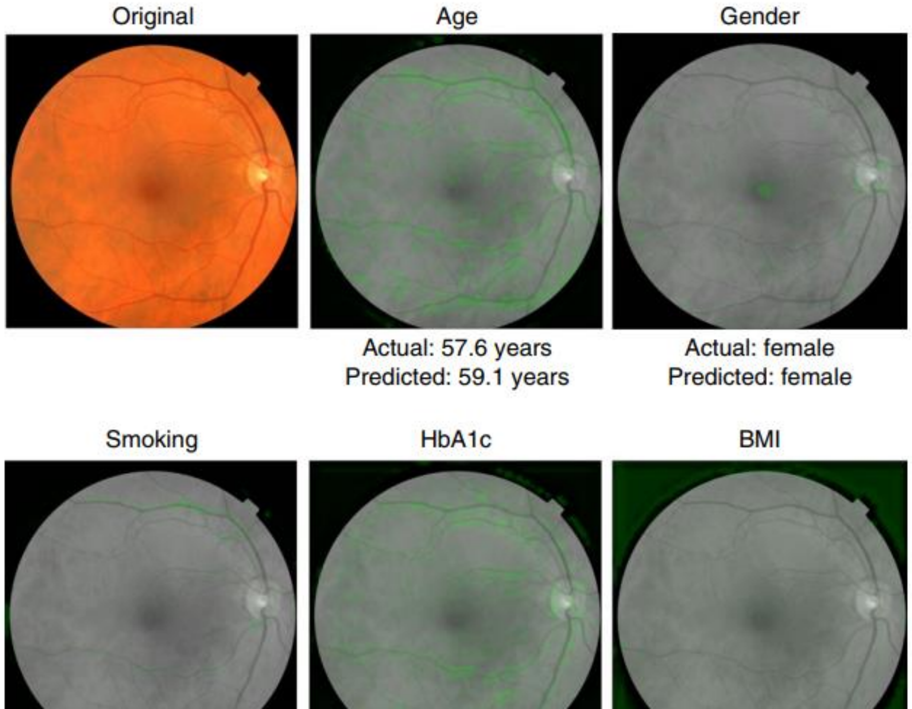

Actual: 57.6 years
Predicted: 59.1 years

Actual: female
Predicted: female

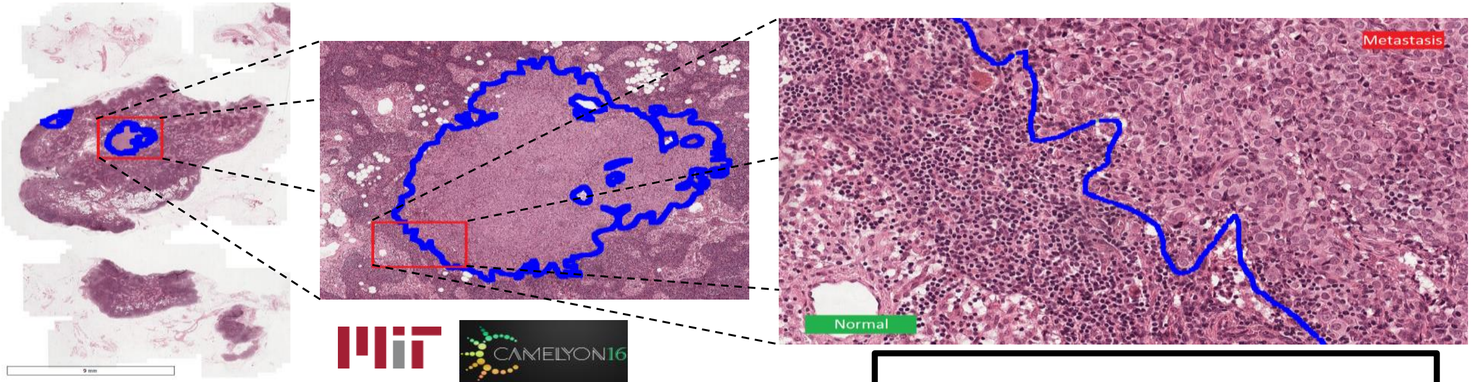

20 Gigapixel
images

Deep Learning for Identifying Metastatic Breast Cancer

Dayong Wang et al. 2016

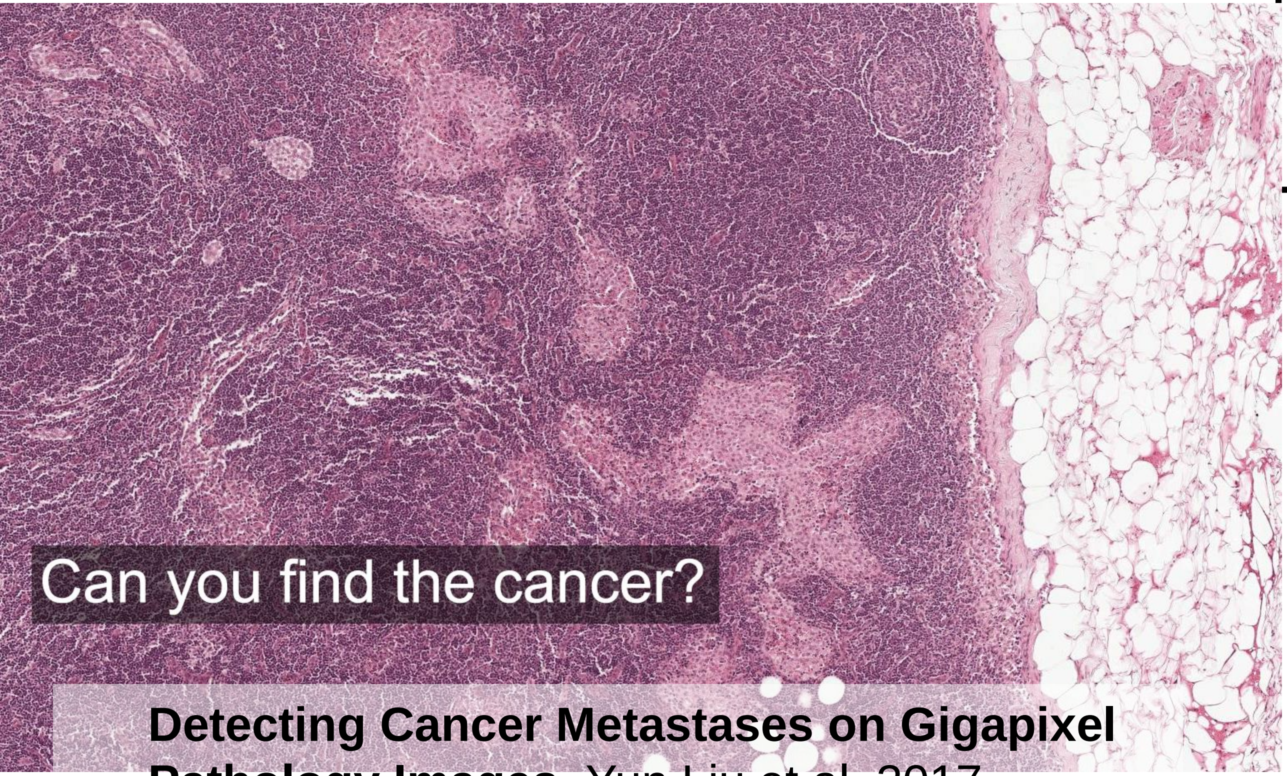

Can you find the cancer?

Detecting Cancer Metastases on Gigapixel Pathology Images, Yun Liu et al. 2017

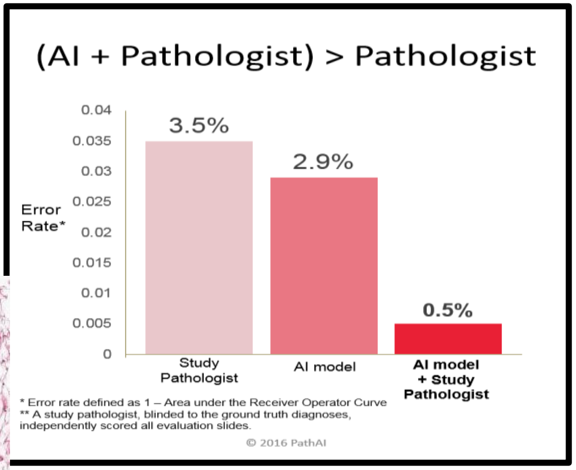

* Error rate defined as 1 – Area under the Receiver Operator Curve
** A study pathologist, blinded to the ground truth diagnoses, independently scored all evaluation slides.

© 2016 PathAI

We obtain AUC of 0.925 for whole slide image classification and a score of 0.7051 for tumor localization. Combining our deep learning system's predictions with the human pathologist's diagnoses increased his AUC to 0.995, representing an approximately 85% reduction in human error rate.

We showed that it is possible to train a model that either matched or exceeded the performance of a pathologist who had unlimited time to examine the slides."

SkinVision

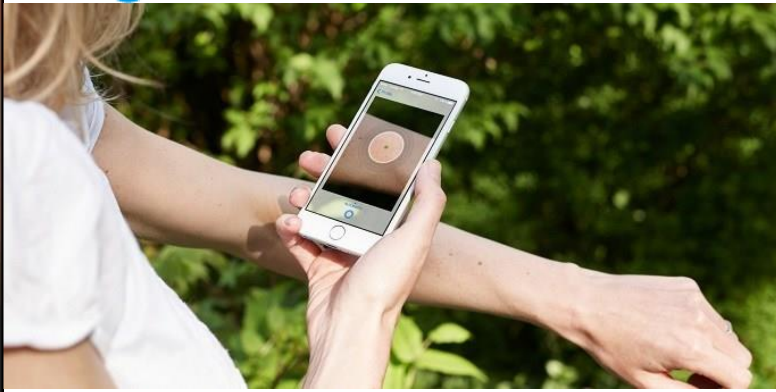

Take a photo of your skin spot

Receive your risk indication

Schedule your next check

Dermatologist-level classification of skin cancer with deep neural networks
Andre Esteva et al. Nature 542, 2017

"We train a CNN using a dataset of 129,450 clinical images—two orders of magnitude larger than previous datasets—consisting of 2,032 different diseases. We test its performance against 21 board-certified dermatologists on biopsy-proven clinical images with two critical binary classification use cases: keratinocyte carcinomas versus benign seborrheic keratoses; and malignant melanomas versus benign nevi."

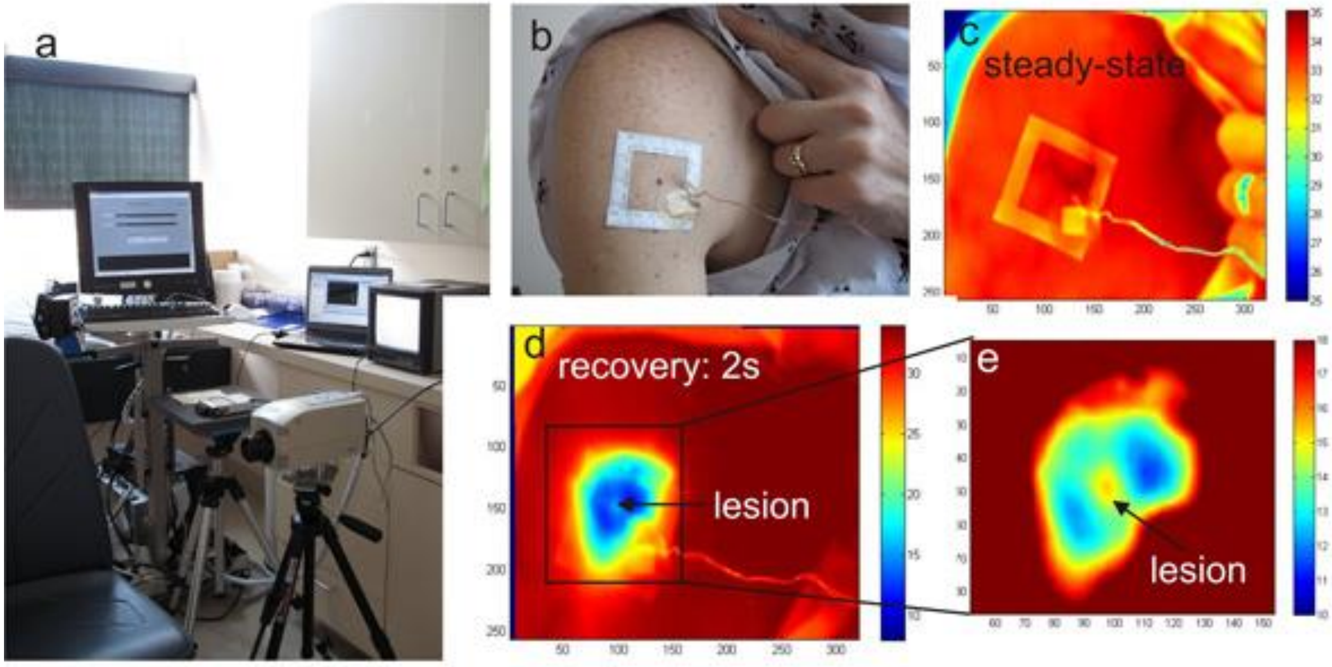

https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3197108/

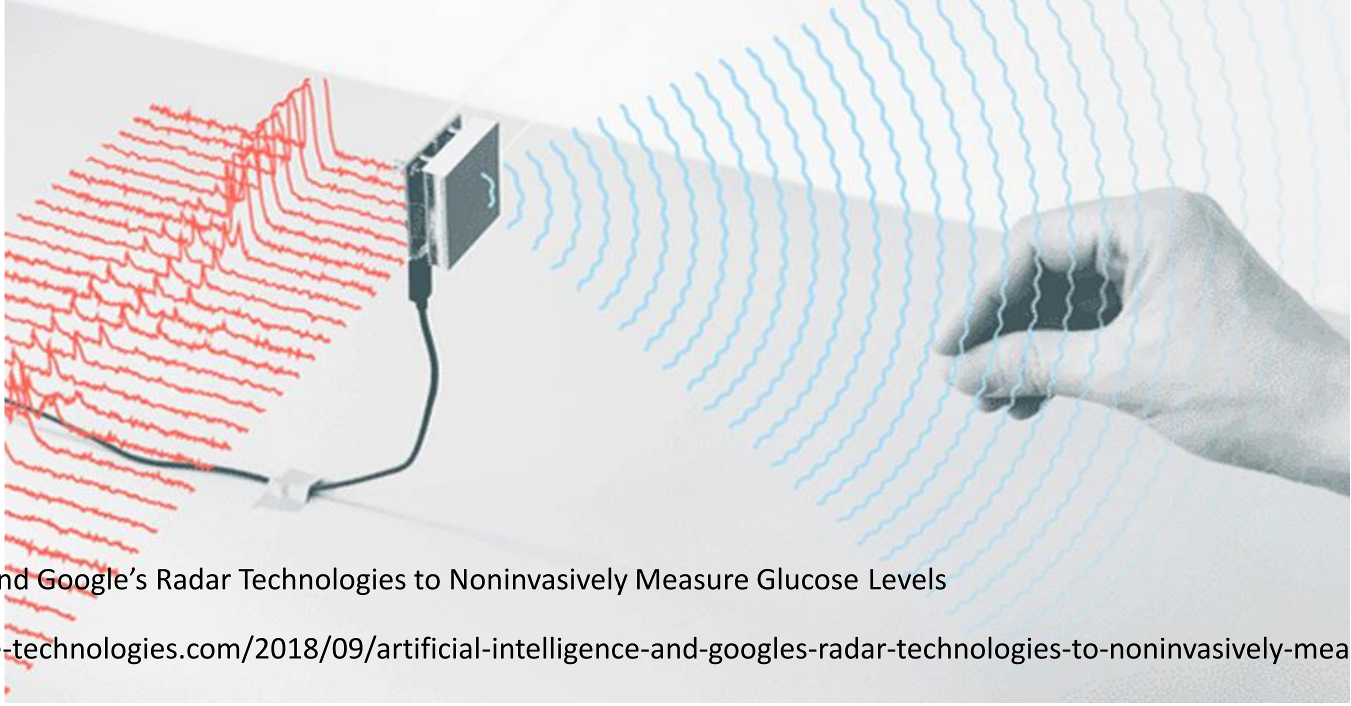

Artificial Intelligence and Google’s Radar Technologies to Noninvasively Measure Glucose Levels

https://www.wearable-technologies.com/2018/09/artificial-intelligence-and-googles-radar-technologies-to-noninvasively-measure-glucose-levels/

With activity analysis Parkinson diagnosis can be done in 3 minutes instead of 30 minutes

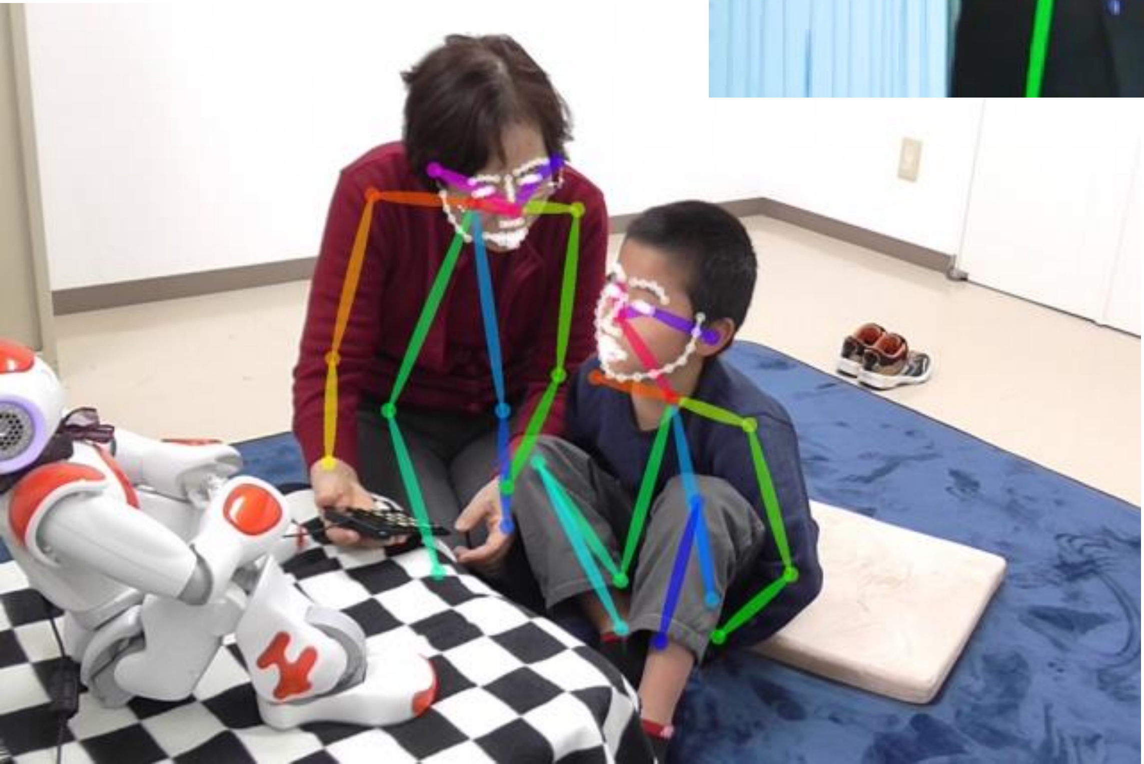

An example of a therapy session augmented with humanoid robot NAO [SoftBank Robotics], which was used in the EngageMe study. Tracking of limbs/faces was performed using the CMU Perceptual Lab's OpenPose utility.
Image: MIT Media Lab

http://news.mit.edu/2018/personalized-deep-learning-equips-robots-autism-therapy-0627

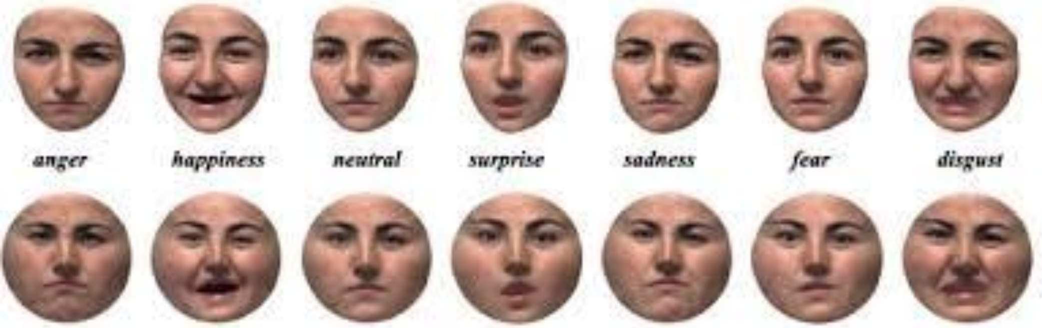

Detection and Computational Analysis of Psychological Signals (DCAPS)

http://medvr.ict.usc.edu/projects/dcaps/

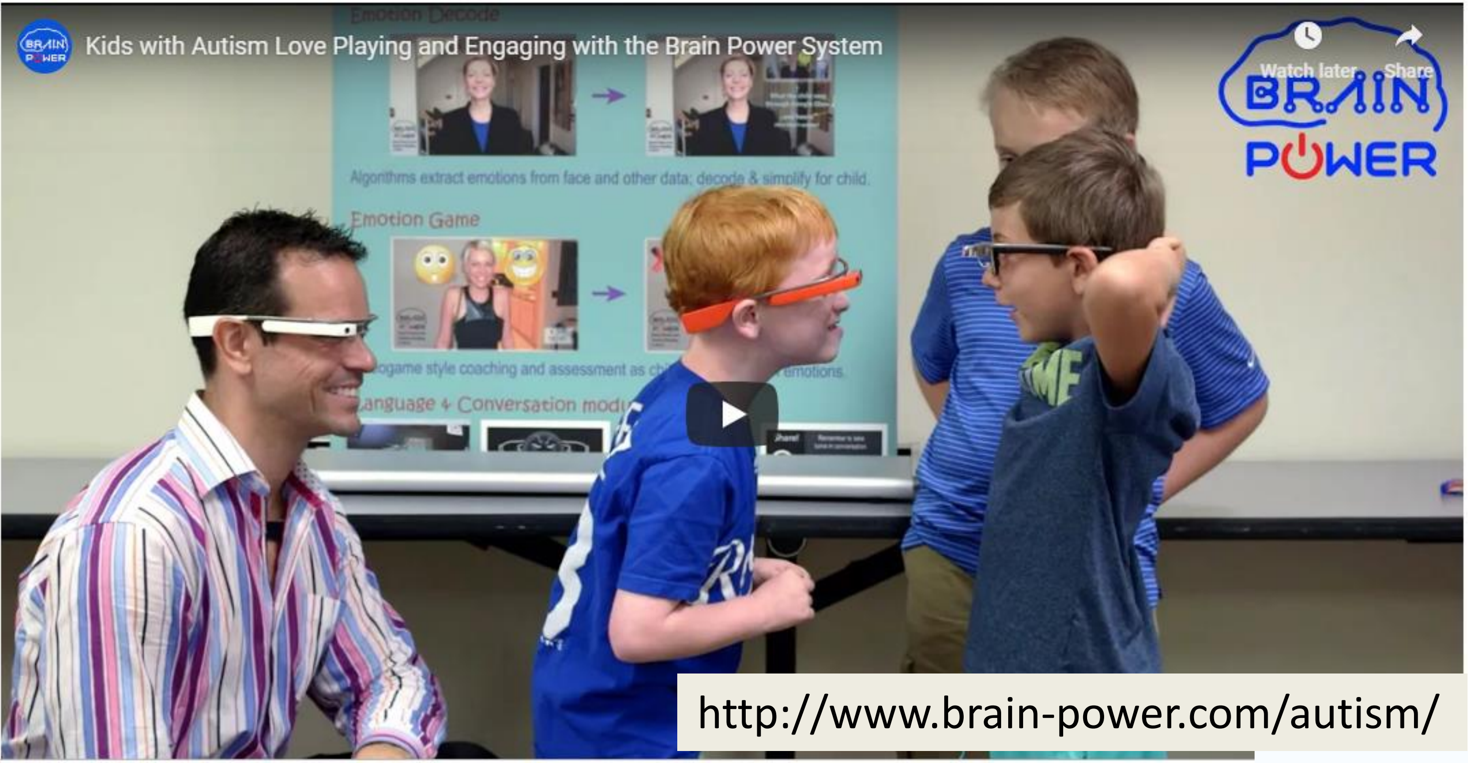
Kids with Autism Love Playing and Engaging with the Brain Power System

http://www.brain-power.com/autism/

otsimo

Çocuğunuz için
kişiselleştirilebilir eğitim.

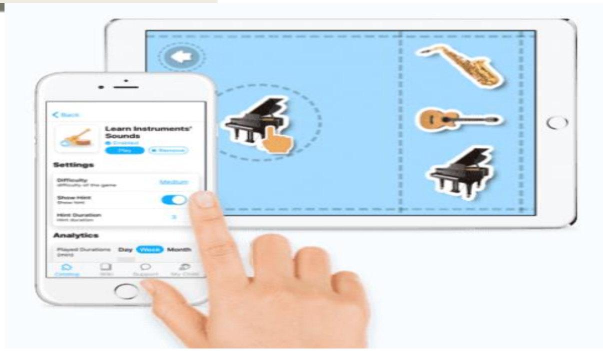

How does the computer learn to see?

MÜNİR ÖZKUL ADİLE NAŞİT ŞENER ŞEN
AYŞEN GRUDA AHMET SEZEREL OYA AYDOĞAN
RENKLİ MÜRÜVVET SİM İHSAN YÜCE

NEŞELİ GÜNLER

YÖNETMEN ORHAN AKSOY GÖRÜNTÜ YÖNETMENİ ERDOĞAN ENGİN SENARYO SADIK ŞENDİL

ARZU FILM

NEŞELİ GÜNLER

MÜNİR ÖZKUL ADİLE NAŞİT ŞENER ŞEN
AYŞEN GRUDA AHMET SEZEREL OYA AYDOĞAN
RENKLİ MÜRÜVVET SİM İHSAN YÜCE

YÖNETMEN ORHAN AKSOY GÖRÜNTÜ YÖNETMENİ ERDOĞAN ENGİN SENARYO SADIK ŞENDİL

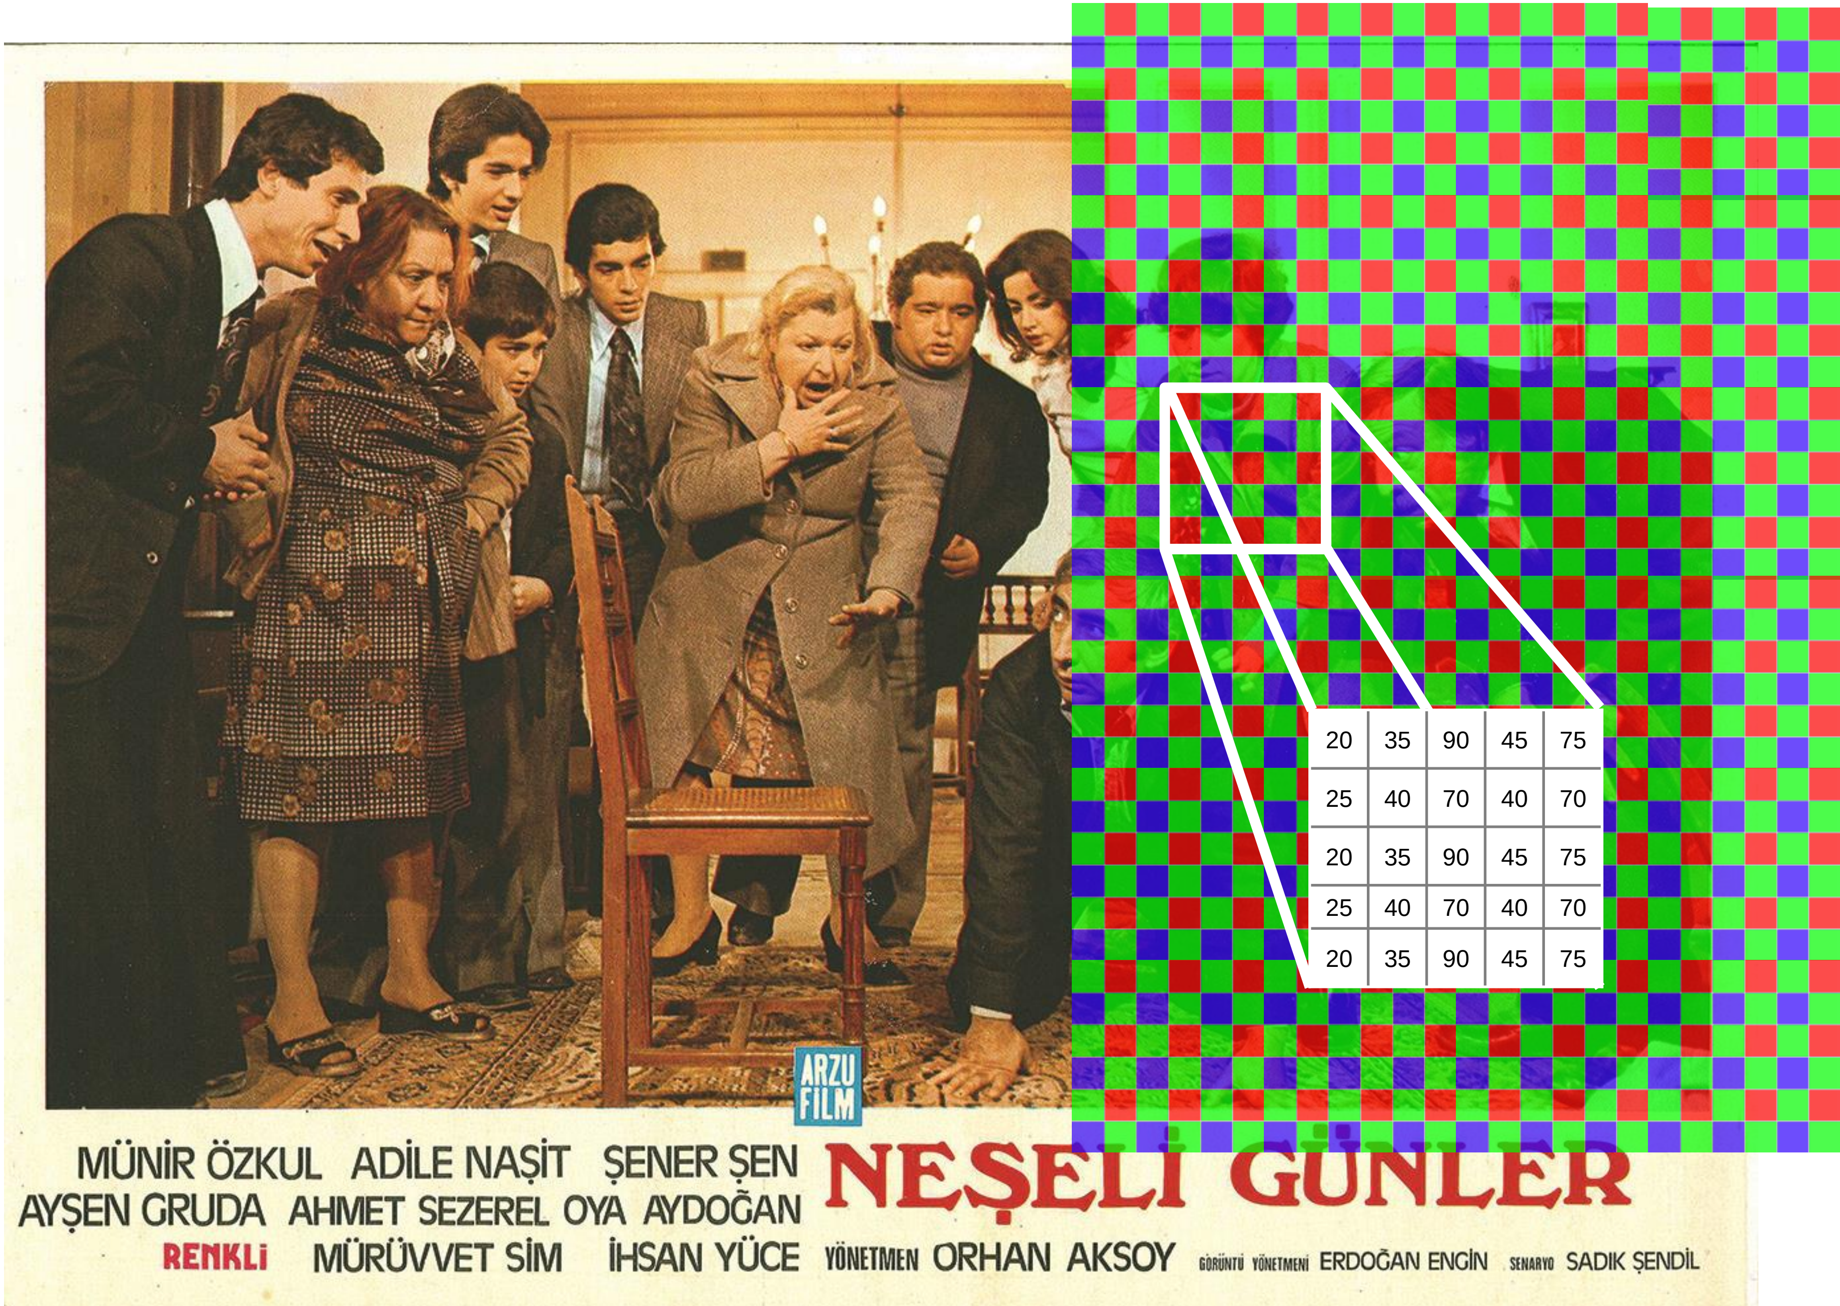

20	35	90	45	75
25	40	70	40	70
20	35	90	45	75
25	40	70	40	70
20	35	90	45	75

We are trying to develop automatic algorithms that would “see”.

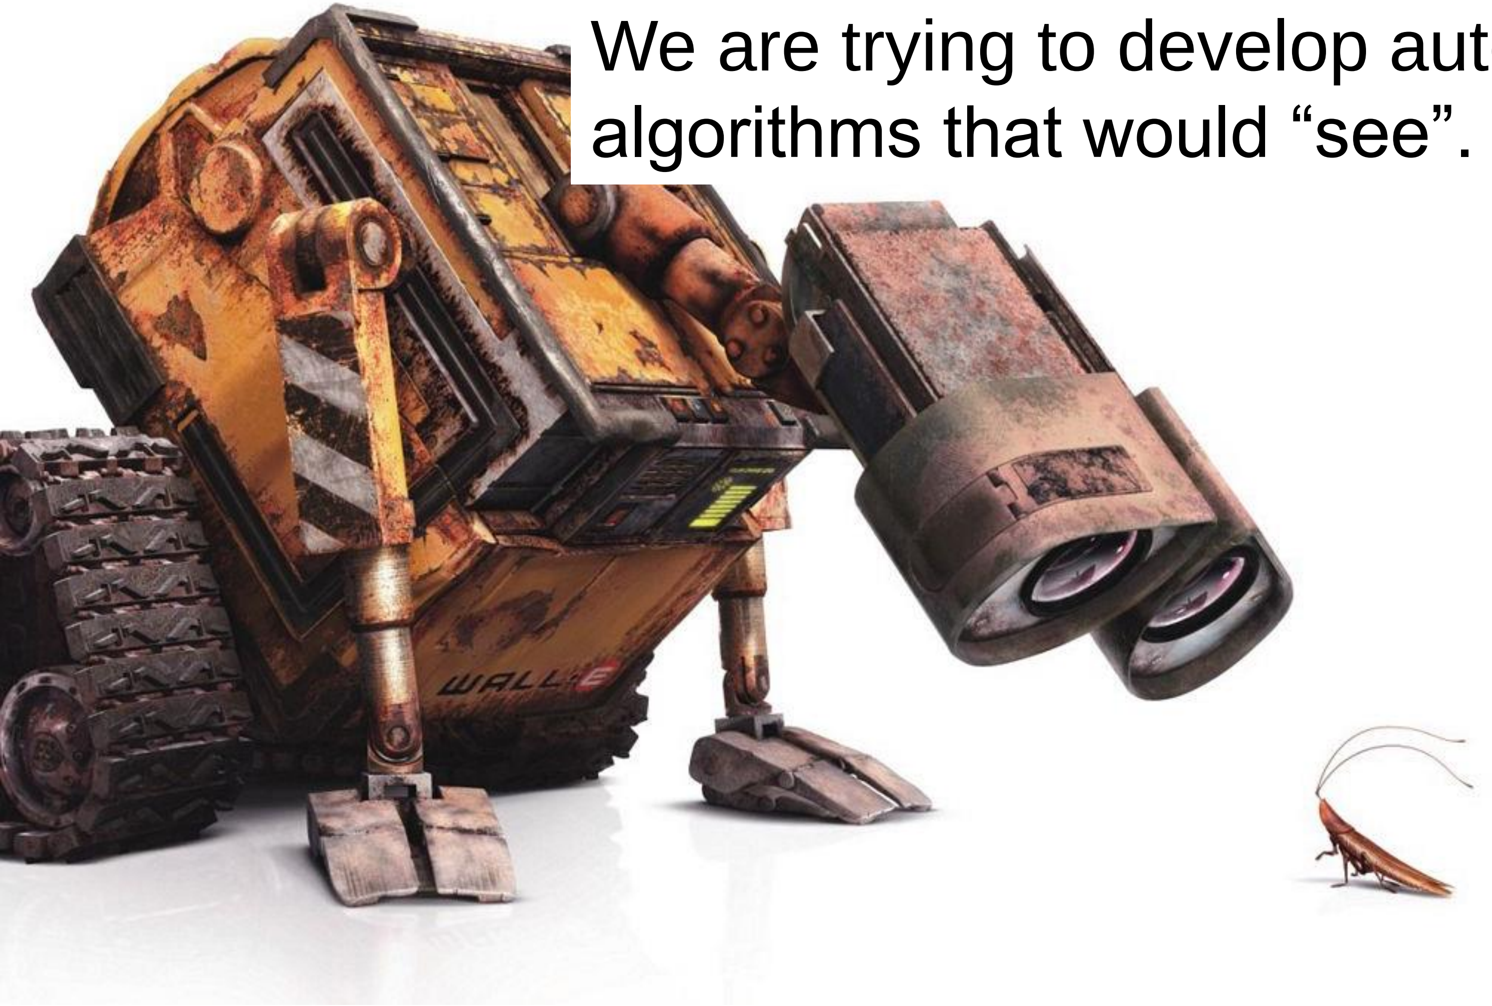

How it all started?

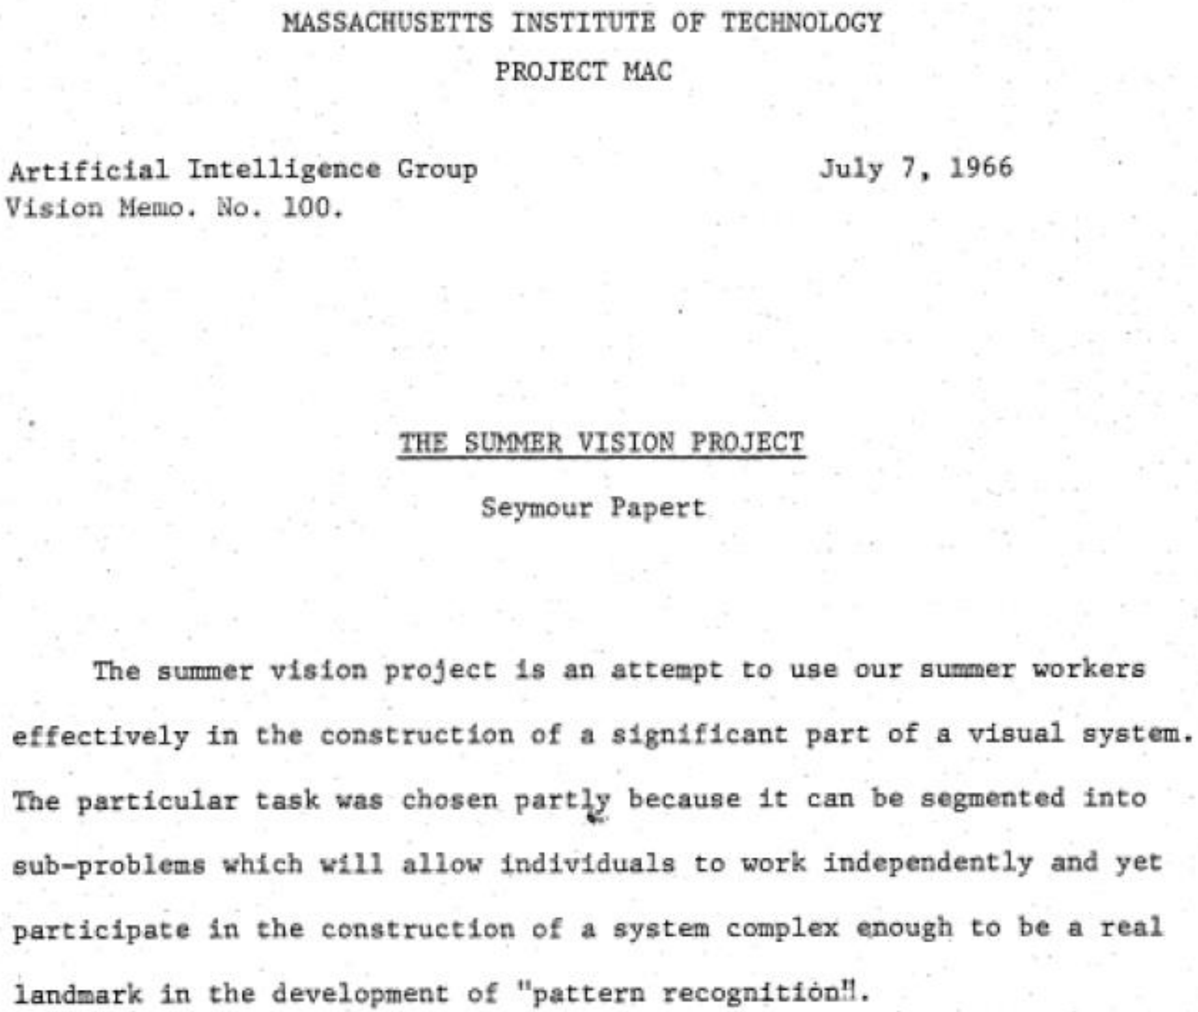

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.

July 7, 1966

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

What does it mean to see?

Understanding a scene as a whole

Image Classification

marketplace
outdoor
street
urban
…

To know what is where by looking – Marr, 1982

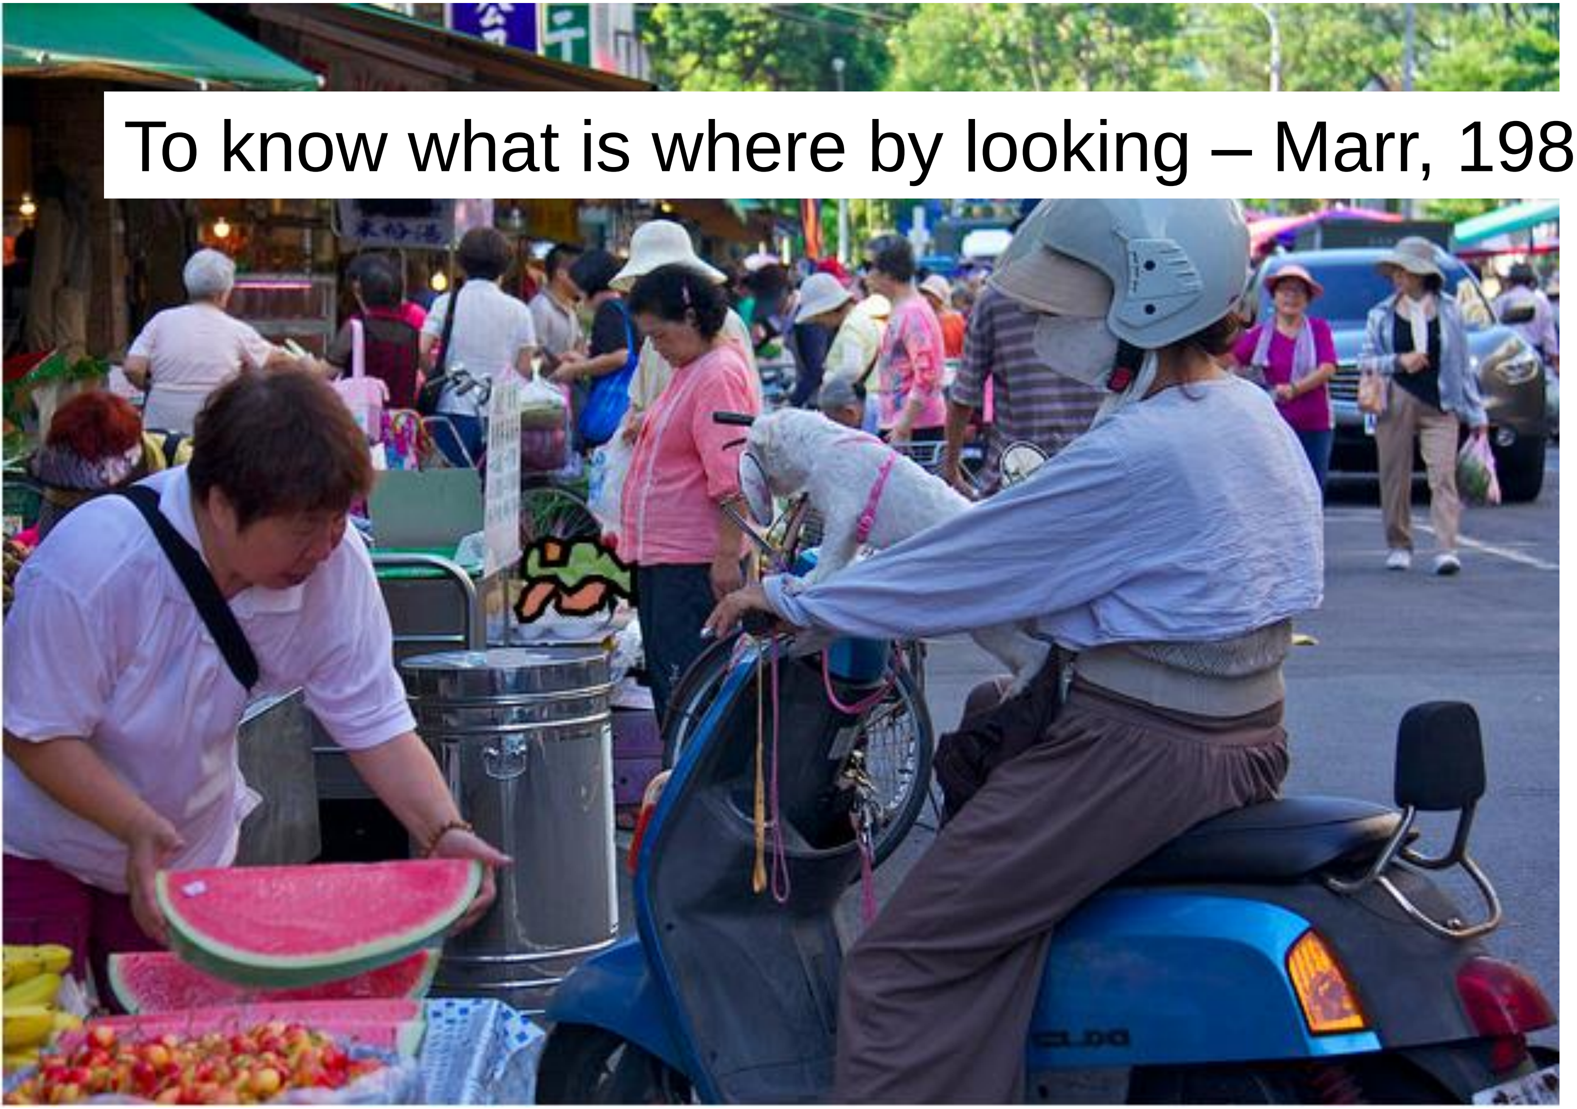

Object Detection / Segmentation

apple

banana

bicycle

car

dog

motorcycle

person

To know what is where by looking – Marr, 1982

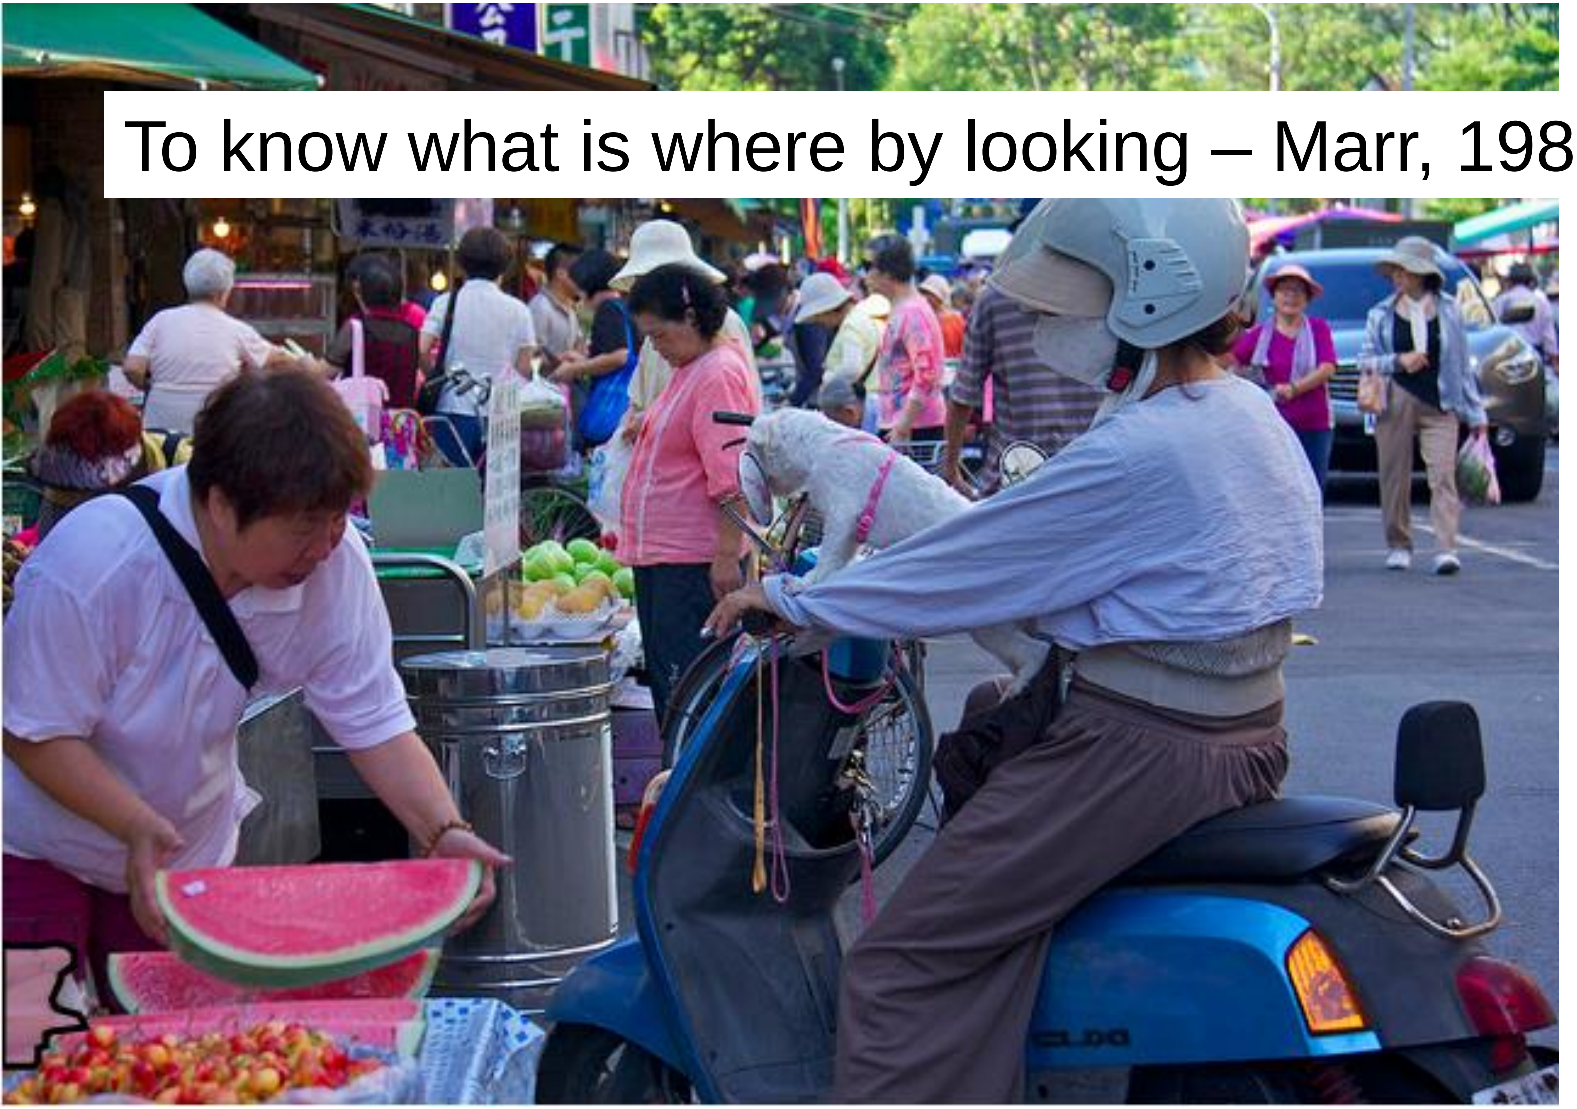

Object Detection / Segmentation

apple

banana

bicycle

car

dog

motorcycle

person

To know what is where by looking – Marr, 1982

Object Detection / Segmentation

apple

banana

bicycle

car

dog

motorcycle

person

To know what is where by looking – Marr, 1982

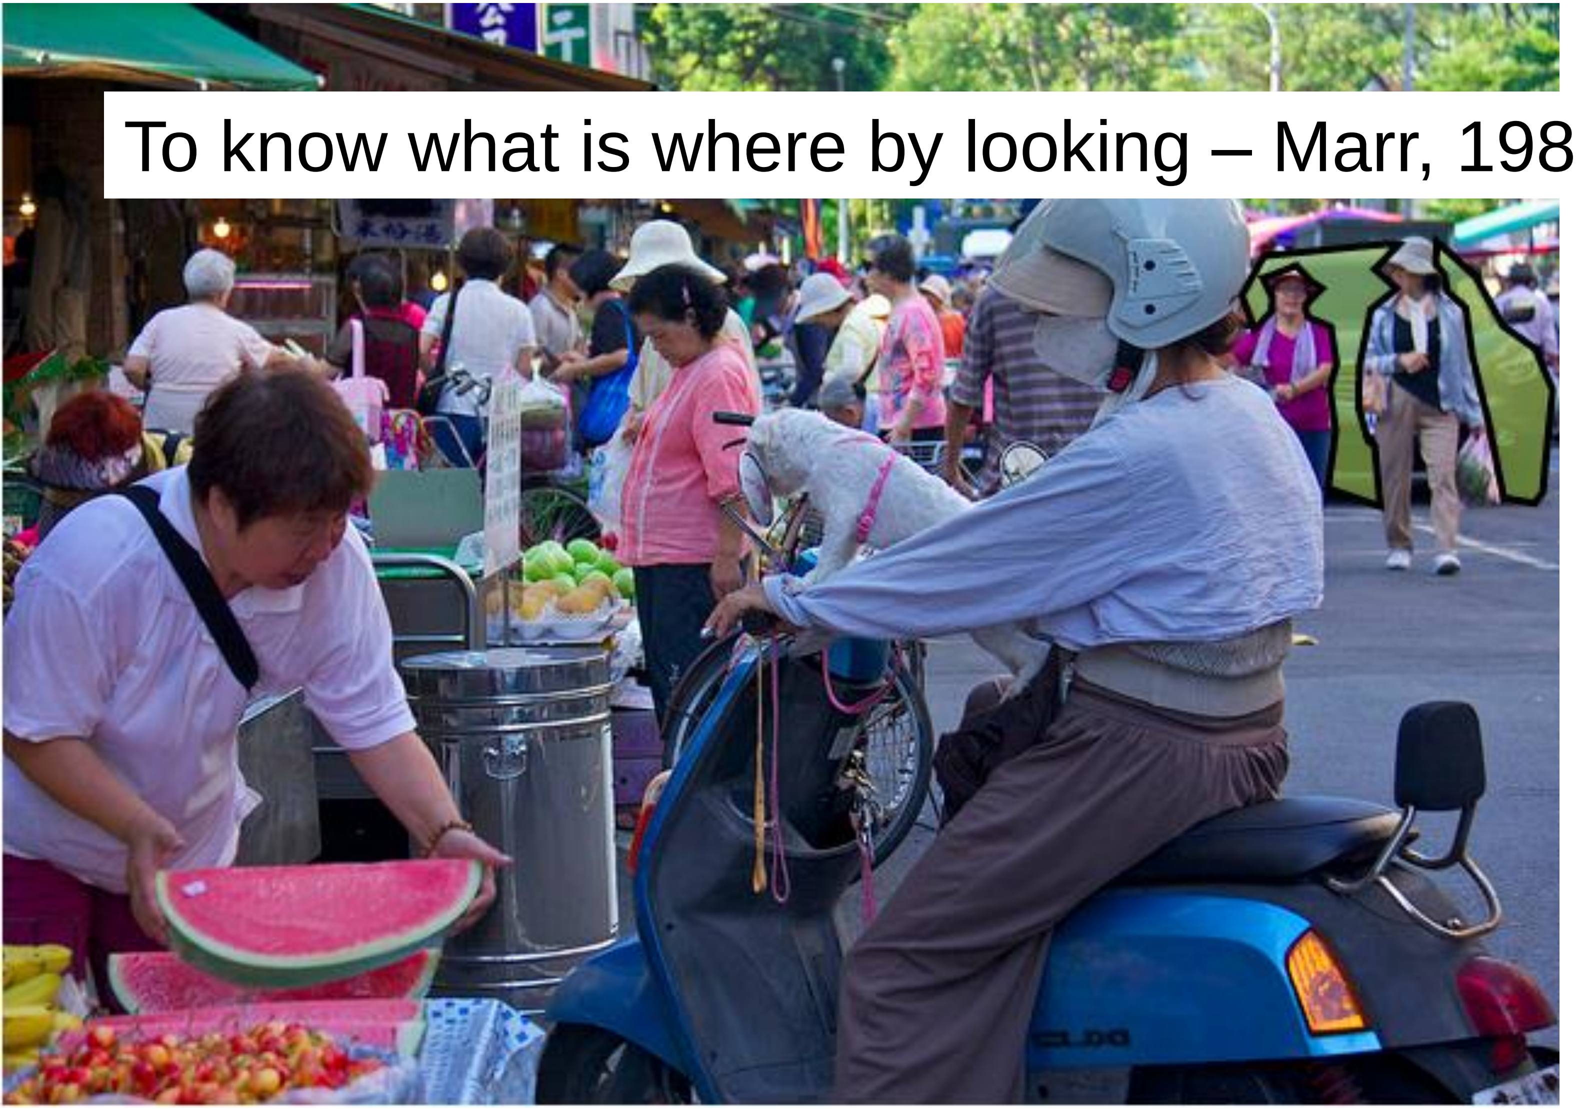

Object Detection / Segmentation

- apple
- banana
- bicycle
- car
- dog
- motorcycle
- person

To know what is where by looking – Marr, 1982

Object Detection / Segmentation

apple
banana
bicycle
car
dog
motorcycle
person

To know what is where by looking – Marr, 1982

Object Detection / Segmentation

apple

banana

bicycle

car

dog

motorcycle

person

To know what is where by looking – Marr, 1982

Object Detection / Segmentation

apple

banana

bicycle

car

dog

motorcycle

person

To know what is where by looking – Marr, 1982

Object Detection / Segmentation

apple
banana
bicycle
car
dog
motorcycle
person

Action Recognition

woman holding a watermelon

person riding a motorcycle

woman looking at apples

woman walking

Object Relations

Understand where things are in the world

🟩 woman behind a stand

🟦 person on a motorcycle

🟪 woman in front of a person

🟥 woman near to another woman

How vision relates to language?

Image Captioning

- a street scene with a person on a motorcycle.
- a person on a motorcycle along a farmers market
- a woman is showing a watermelon slice to a woman on a scooter.
- a person on a motorcycle talking to a person with a watermelon.
- people at a veggie and fruit market looking at the merchandise.

Input →

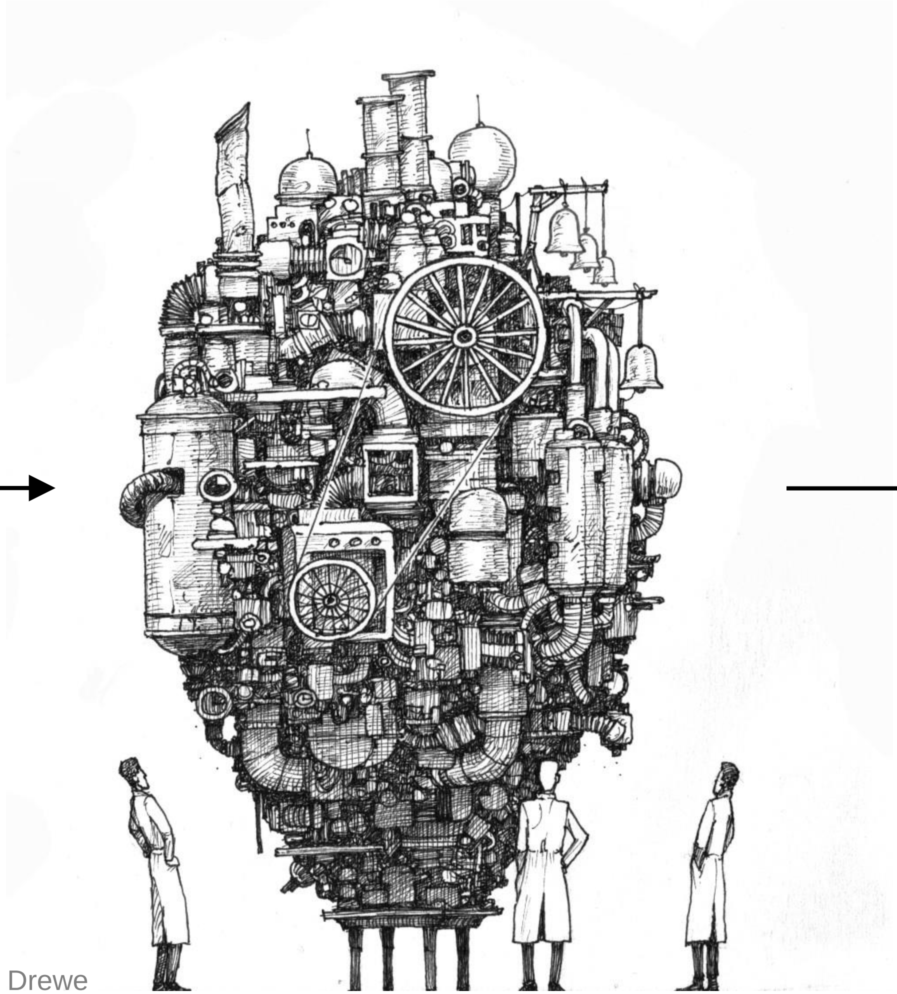

→ Output

Joshua Drewe

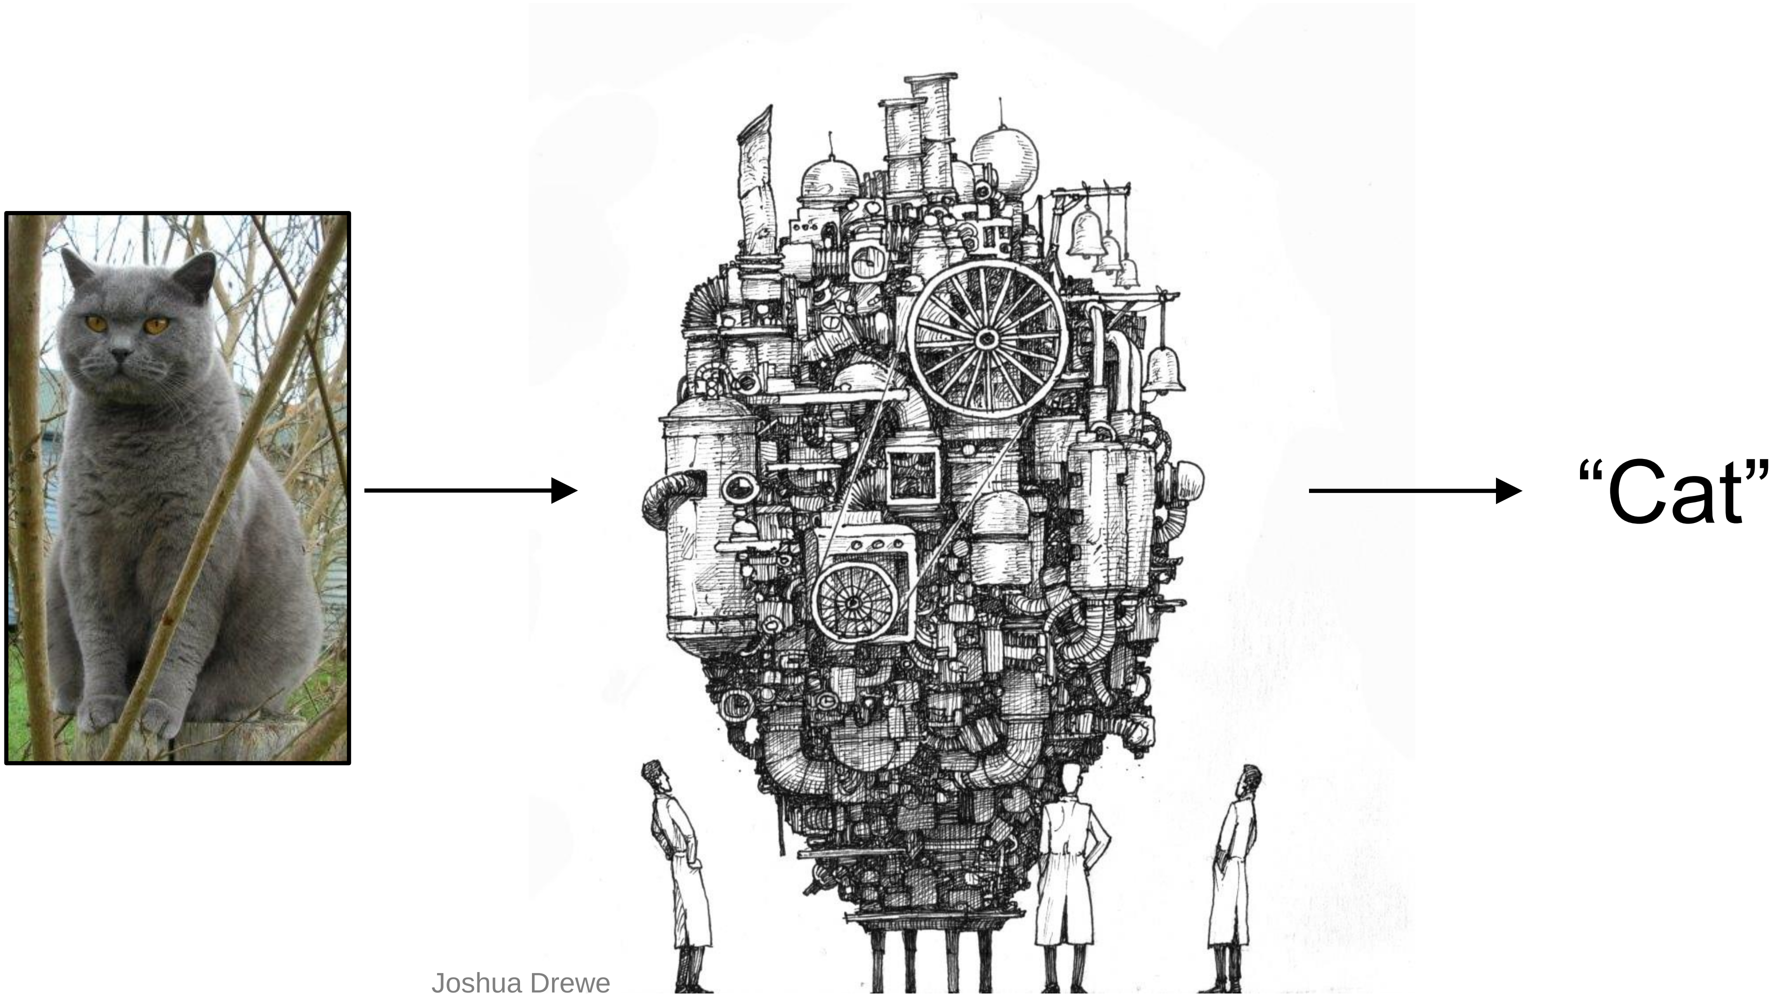

“Cat”

Joshua Drewe

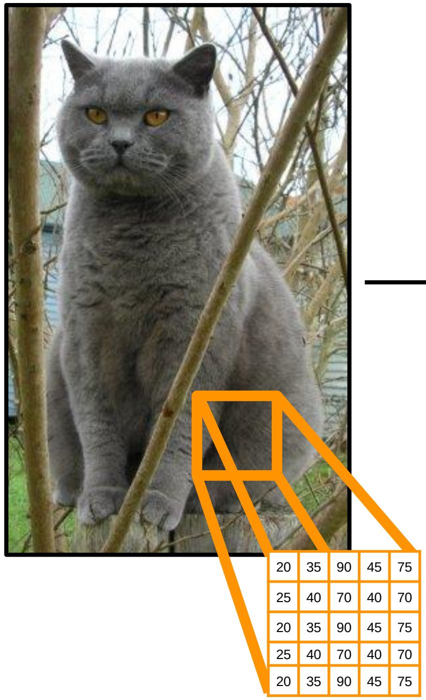

20	35	90	45	75
25	40	70	40	70
20	35	90	45	75
25	40	70	40	70
20	35	90	45	75

→

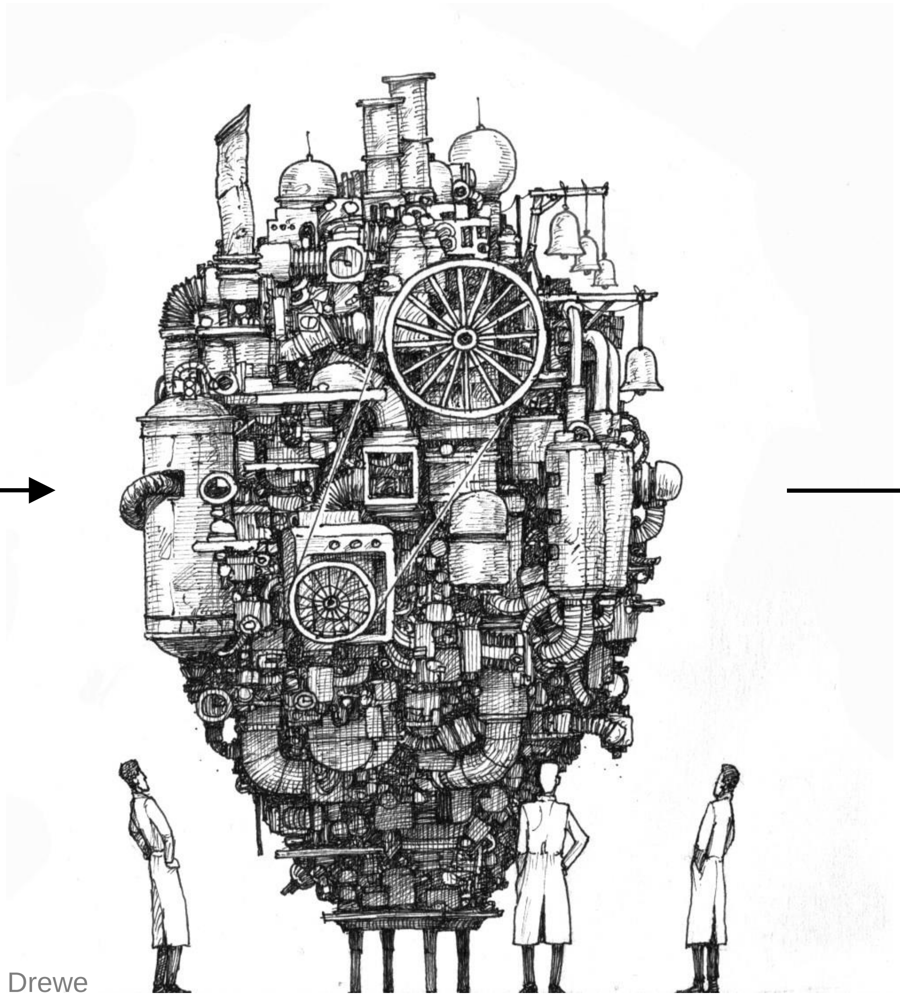

Joshua Drewe

→

“Cat”

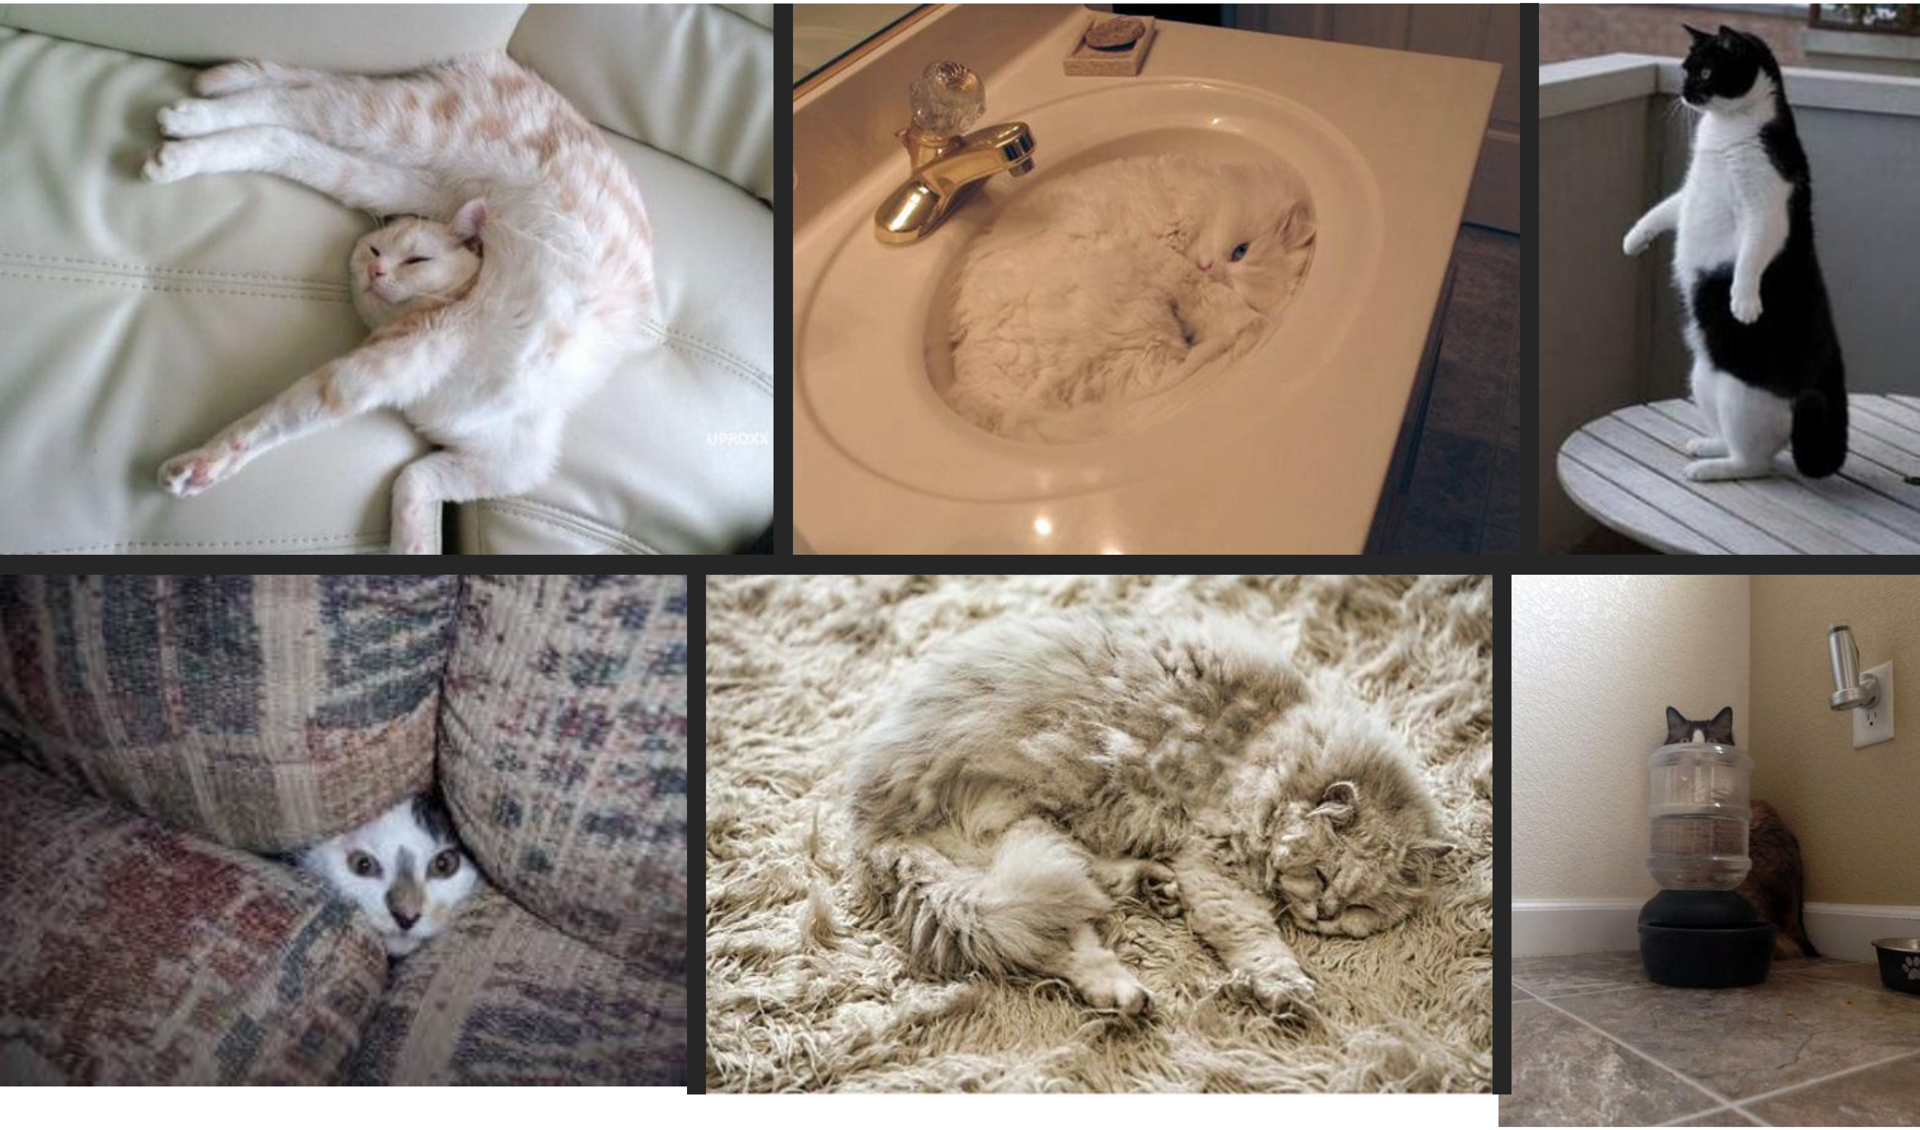
UPROXX

Traditional Machine Learning

VISION

hand-crafted features SIFT/HOG

fixed

your favorite classifier

learned

“car”

SPEECH

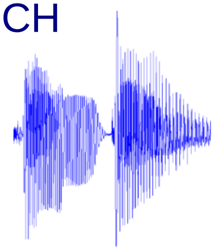

hand-crafted features MFCC

fixed

your favorite classifier

learned

\ˈd ē p\

NLP

This burrito place is yummy and fun!

hand-crafted features Bag-of-words

fixed

your favorite classifier

learned

“+”

Mimicking the Brain : Neural Networks

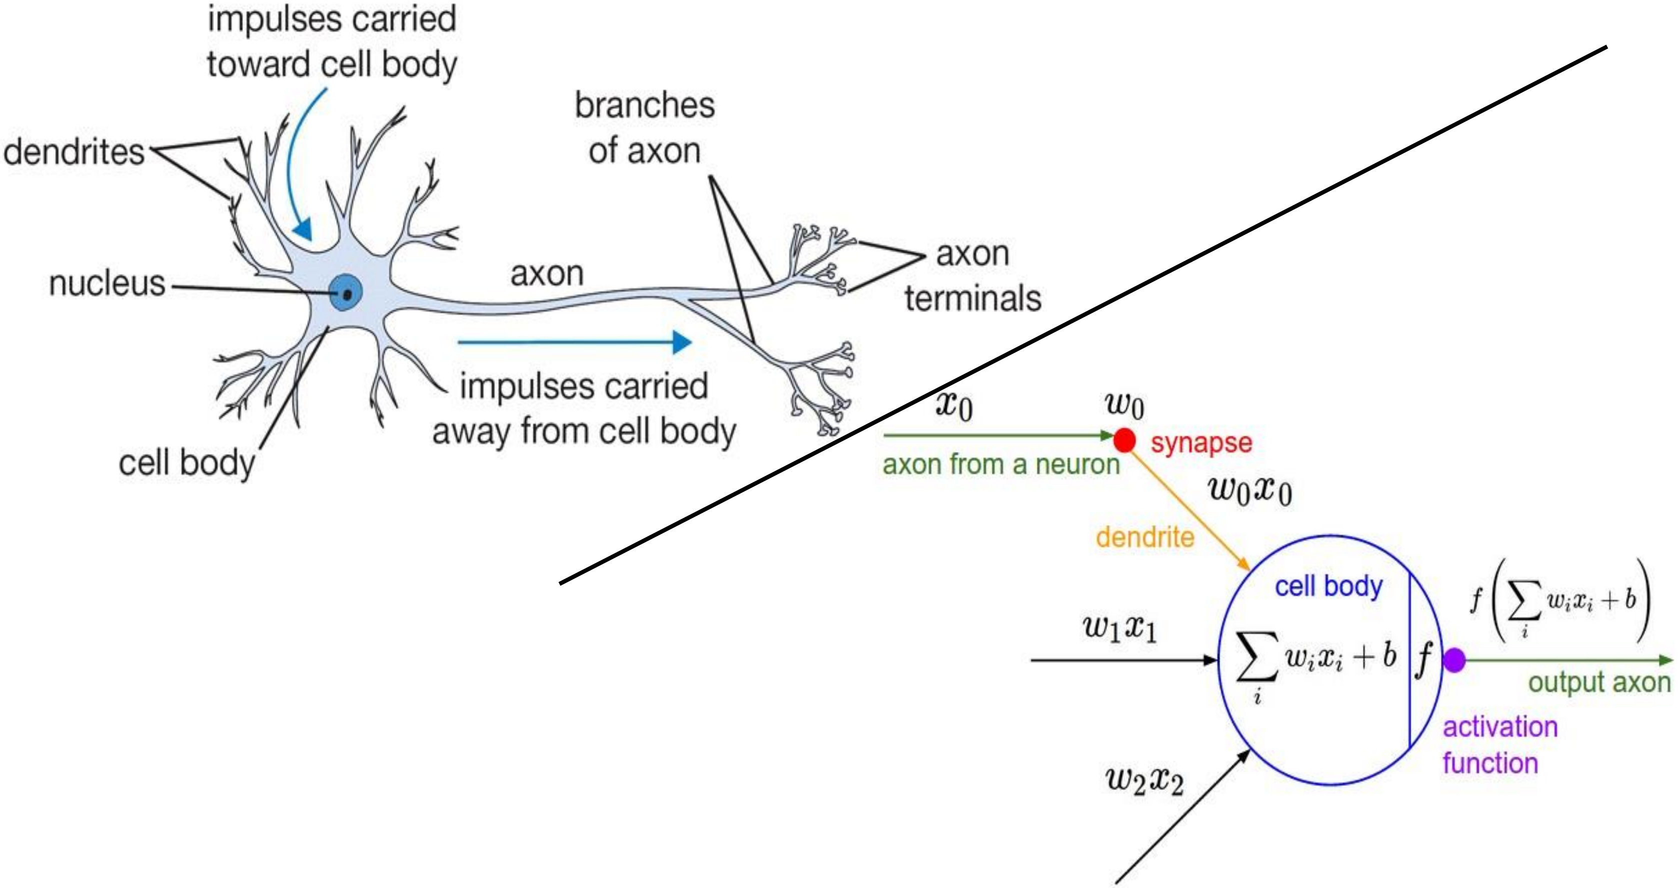

Neural Networks

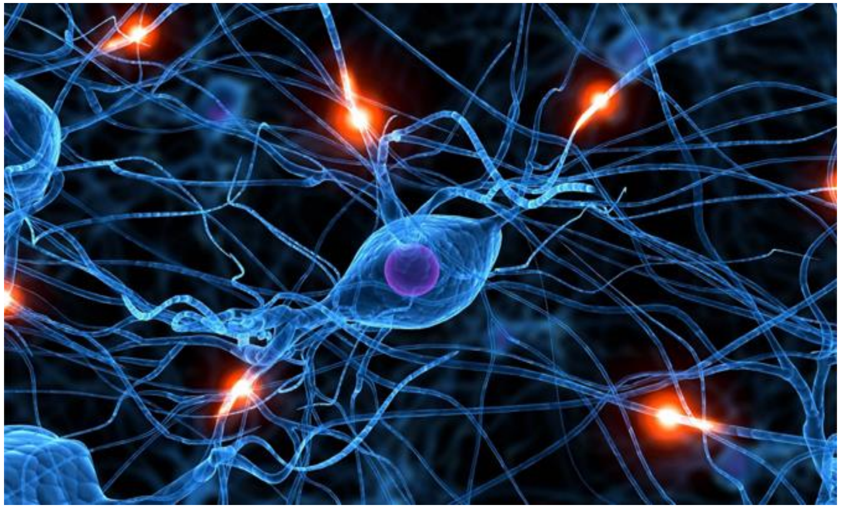

Network in the brain

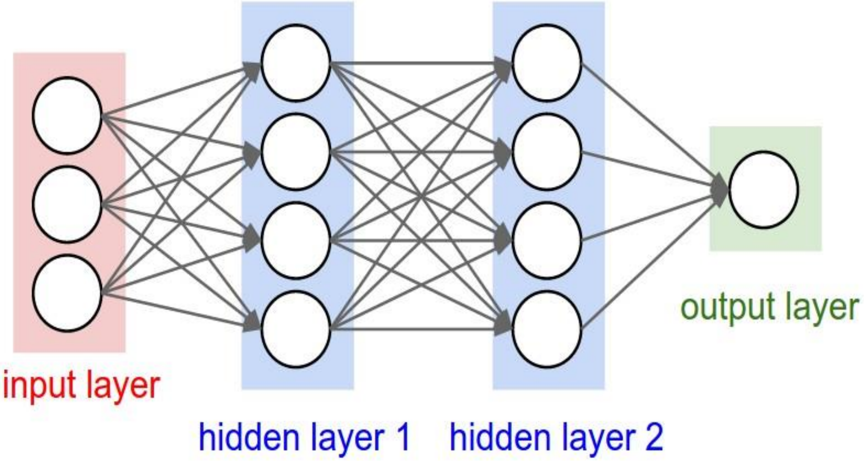

Artificial Neural Net

The Mammalian Visual Cortex is Hierarchical

The ventral (recognition) pathway in the visual cortex

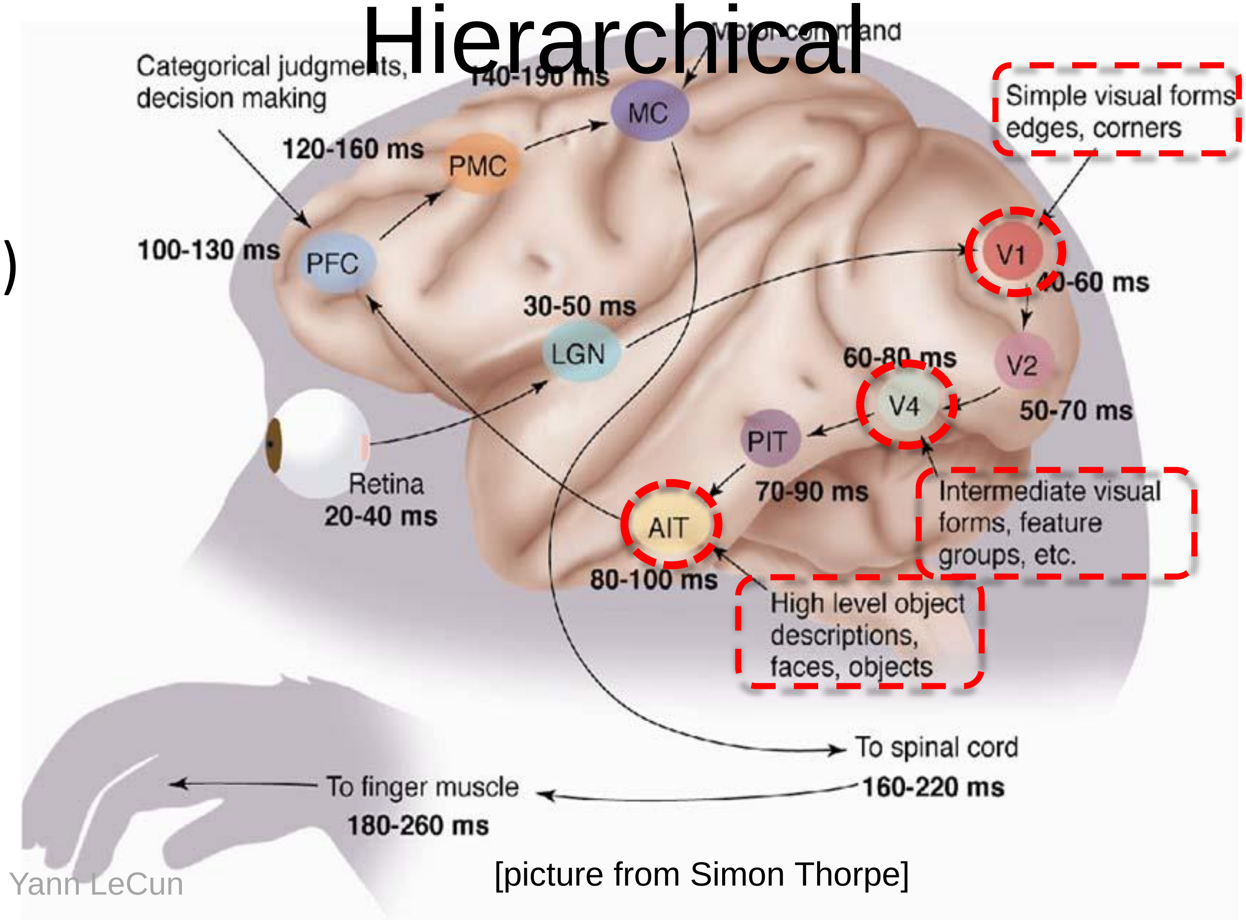

[picture from Simon Thorpe]

Deep Learning = End-to-End Learning

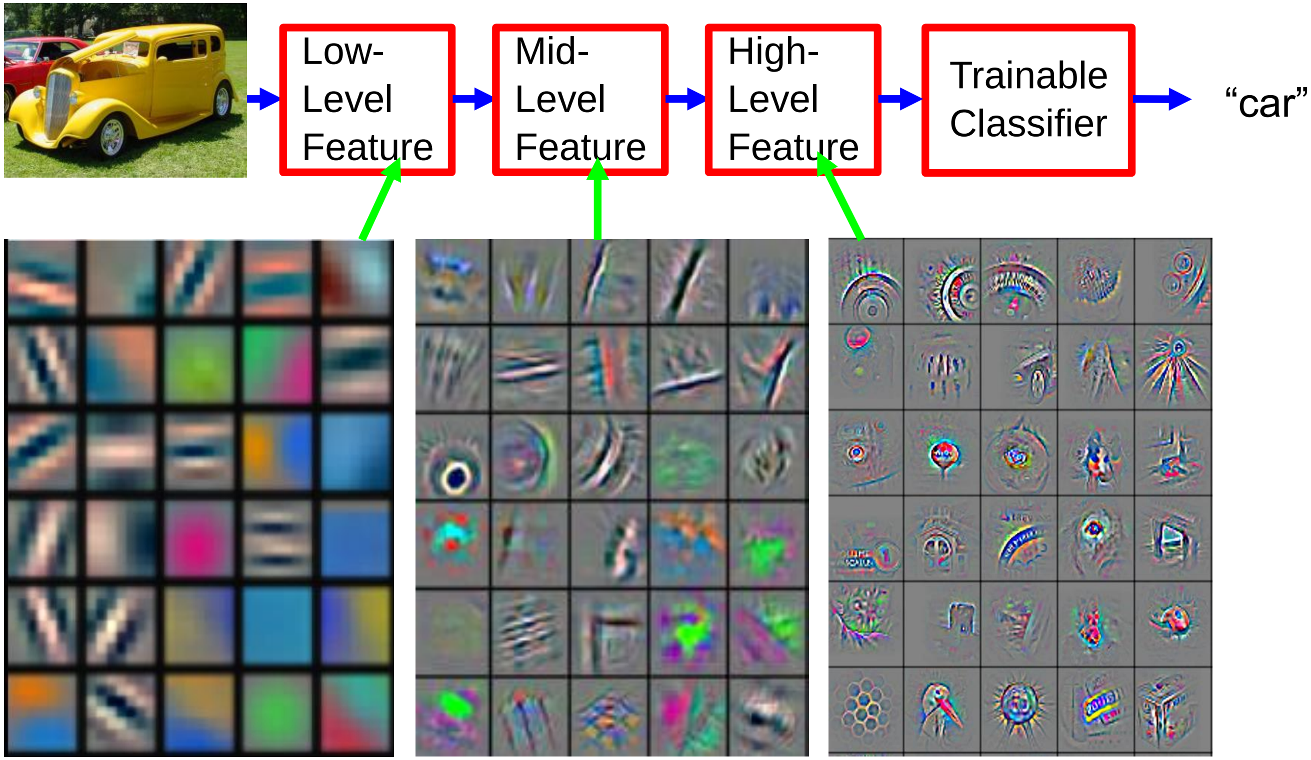

Low-Level Feature → Mid-Level Feature → High-Level Feature → Trainable Classifier → “car”

M.D. Zeiler and R. Fergus, “Visualizing and Understanding Convolutional Networks”, In ECCV 2014

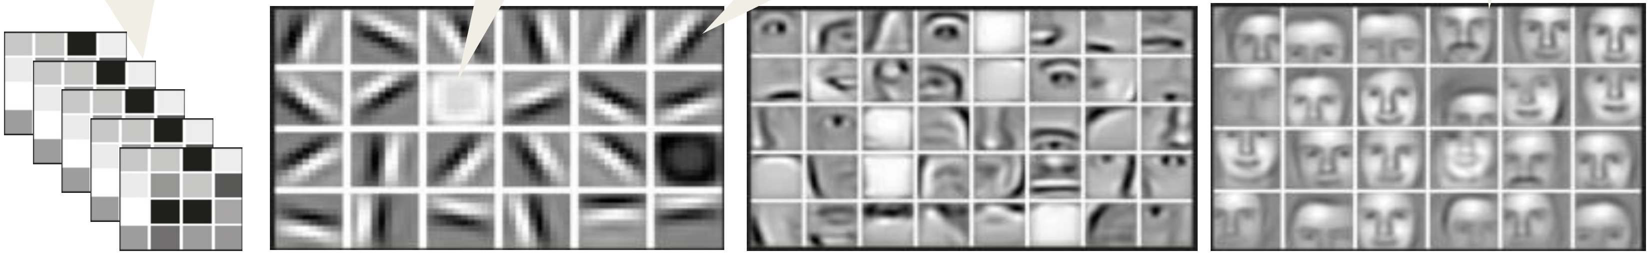

A Deep Learning algorithm is presented with millions of images made up of simple pixels.

The algorithm discovers simple “regularities” that are present across many/all images, like curves & lines.

The algorithm discovers how these regularities are related to form higher-level concepts.

Ultimately, the system gains a high-level understanding of the original data…

All automatically!

Images from Andrew Ng

Convolutional Neural Networks

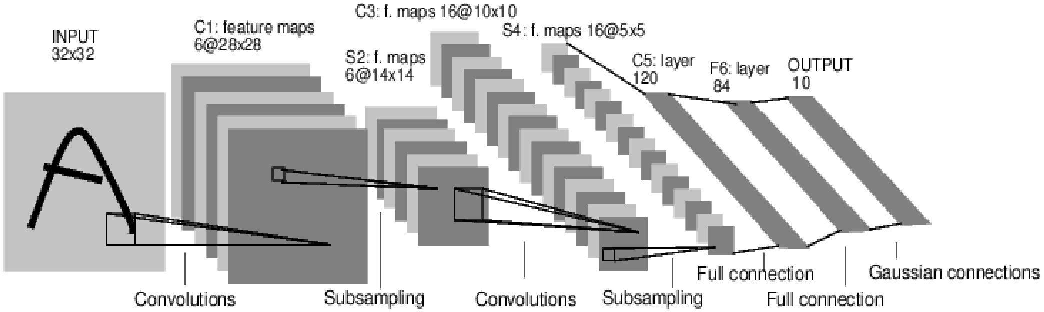

[LeNet-5, LeCun 1980]

Convolutional Neural Networks

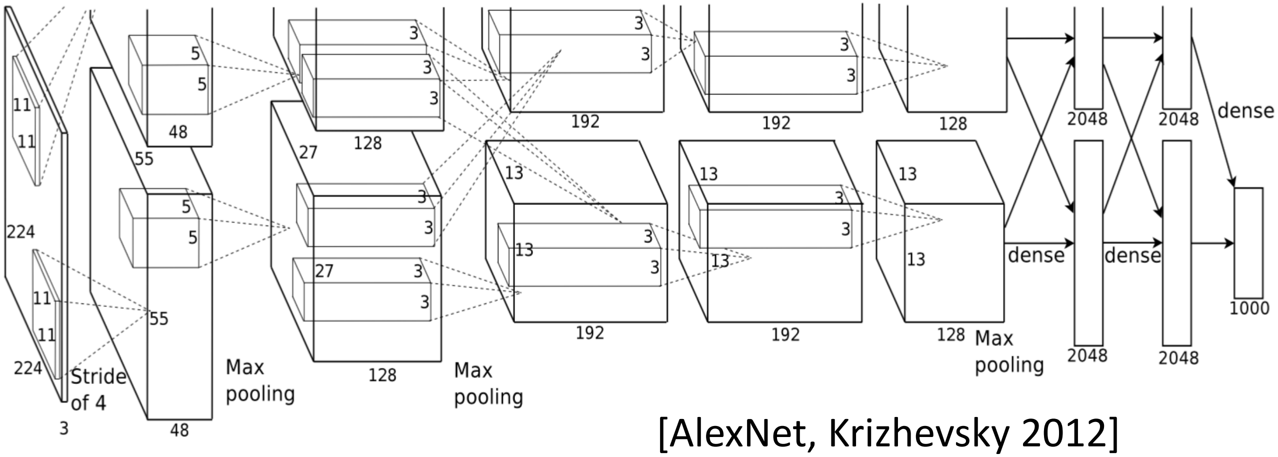

[AlexNet, Krizhevsky 2012]

Object detection results on COCO

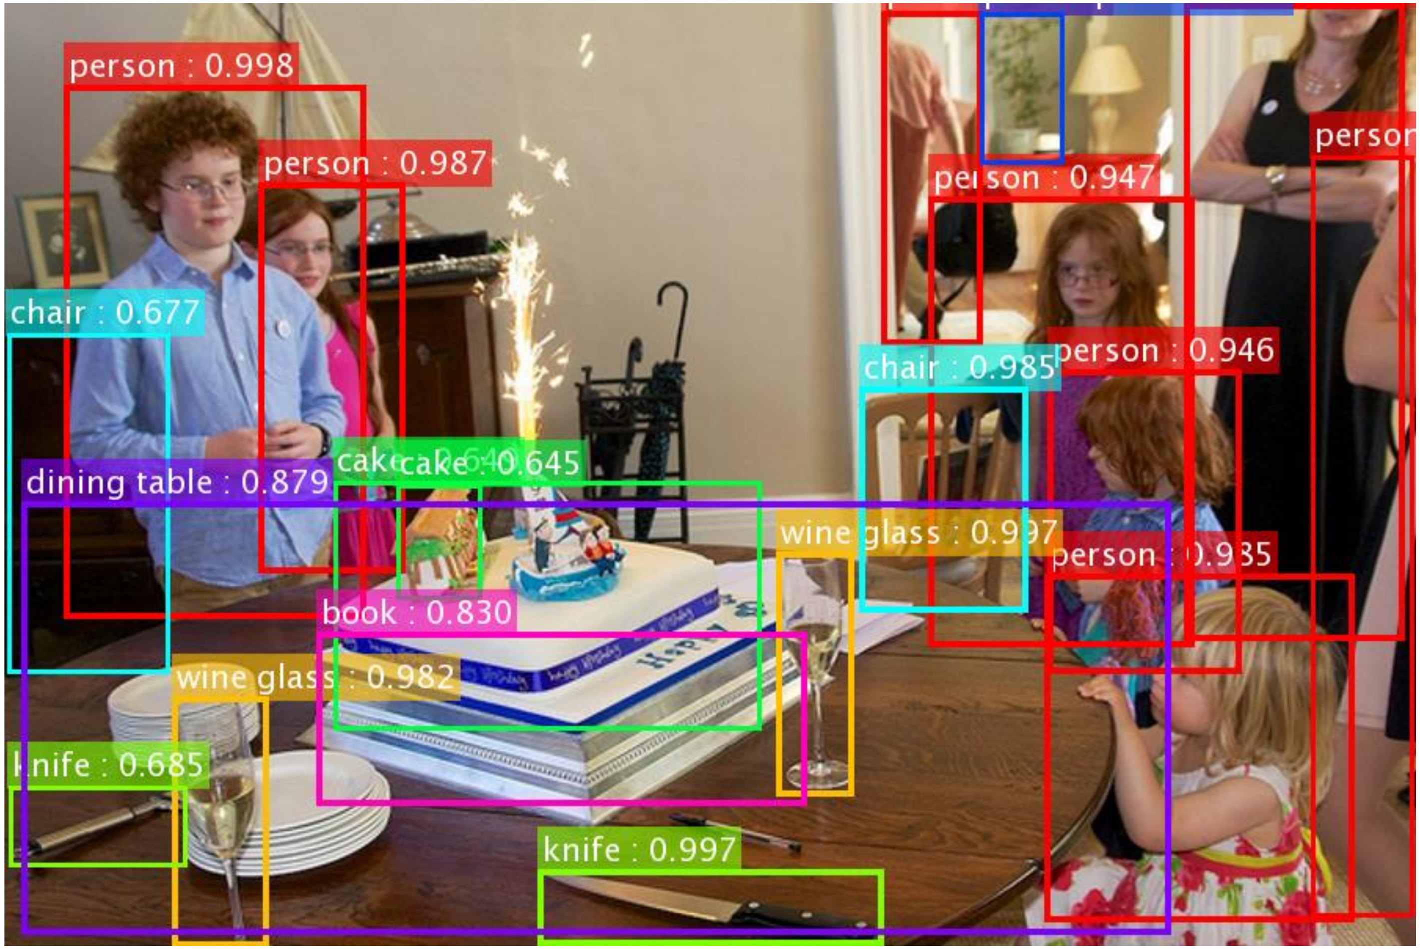

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. **Deep Residual Learning for Image Recognition.** CVPR 2016.

Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. **Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks**. NIPS 2015.

Object detection results on COCO

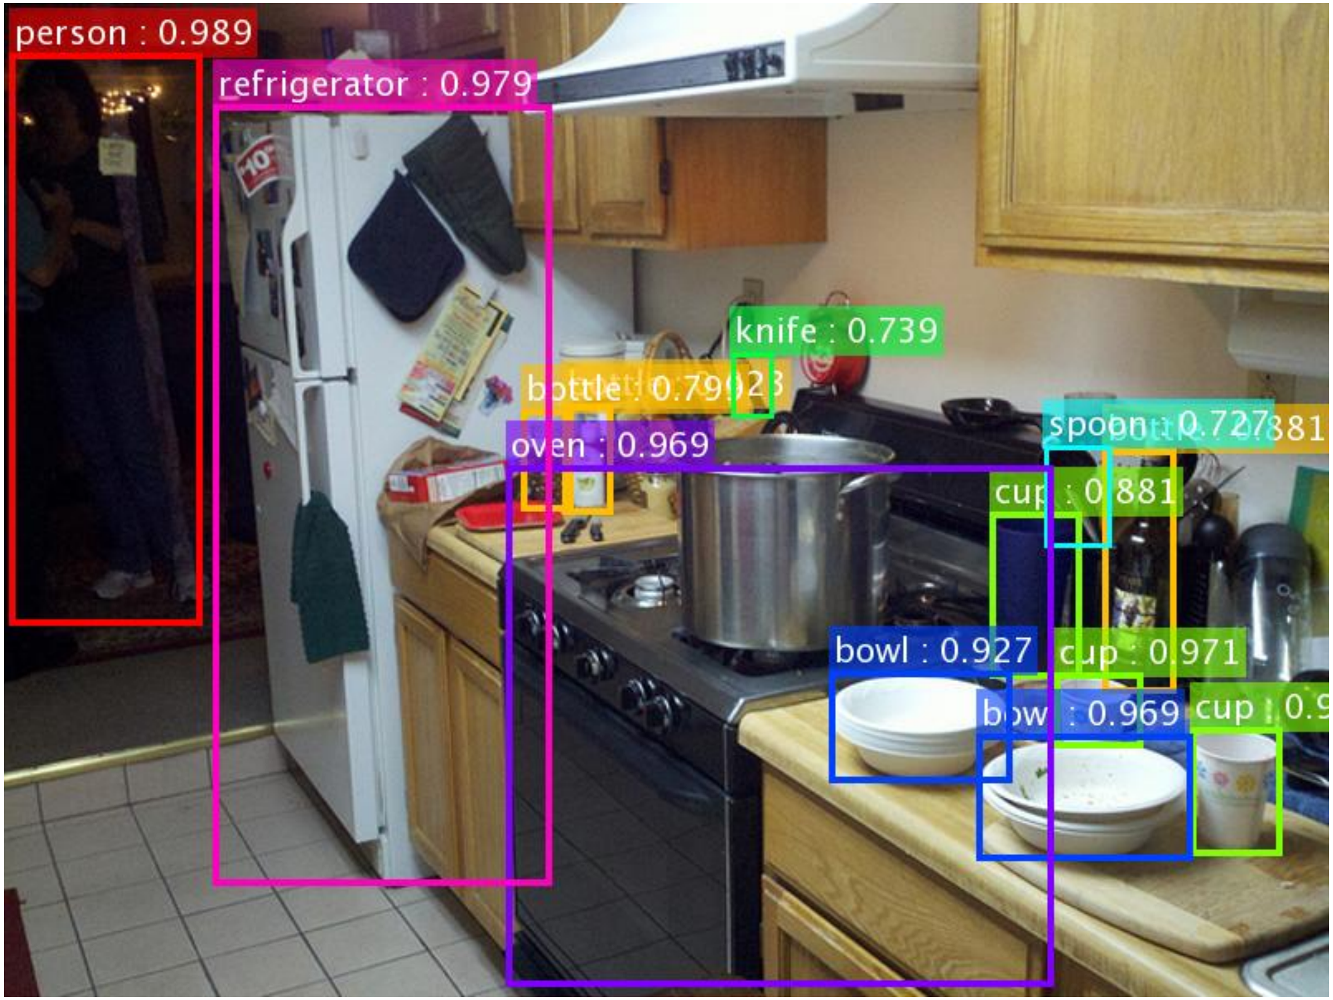

Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. **Deep Residual Learning for Image Recognition.** CVPR 2016.

Shaoqing Ren, Kaiming He, Ross Girshick, & Jian Sun. **Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks**. NIPS 2015.

COCOQA 33827
What is the color of the cat?
Ground truth: black
IMG+BOW: black (0.55)
2-VIS+LSTM: black (0.73)
BOW: gray (0.40)

COCOQA 33827a
What is the color of the couch?
Ground truth: red
IMG+BOW: red (0.65)
2-VIS+LSTM: black (0.44)
BOW: red (0.39)

DAQUAR 1522
How many chairs are there?
Ground truth: two
IMG+BOW: four (0.24)
2-VIS+BLSTM: one (0.29)
LSTM: four (0.19)

DAQUAR 1520
How many shelves are there?
Ground truth: three
IMG+BOW: three (0.25)
2-VIS+BLSTM: two (0.48)
LSTM: two (0.21)

COCOQA 14855
Where are the ripe bananas sitting?
Ground truth: basket
IMG+BOW: basket (0.97)
2-VIS+BLSTM: basket (0.58)
BOW: bowl (0.48)

COCOQA 14855a
What are in the basket?
Ground truth: bananas
IMG+BOW: bananas (0.98)
2-VIS+BLSTM: bananas (0.68)
BOW: bananas (0.14)

DAQUAR 585
What is the object on the chair?
Ground truth: pillow
IMG+BOW: clothes (0.37)
2-VIS+BLSTM: pillow (0.65)
LSTM: clothes (0.40)

DAQUAR 585a
Where is the pillow found?
Ground truth: chair
IMG+BOW: bed (0.13)
2-VIS+BLSTM: chair (0.17)
LSTM: cabinet (0.79)

M. Ren, R. Kiros, and R. Zemel, “Exploring Models and Data for Image Question Answering” NIPS 2015

Autonomous Driving

Waymo / Google Self-Driving Car

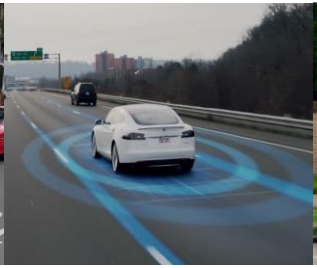

Tesla Autopilot

Uber

nuTonomy

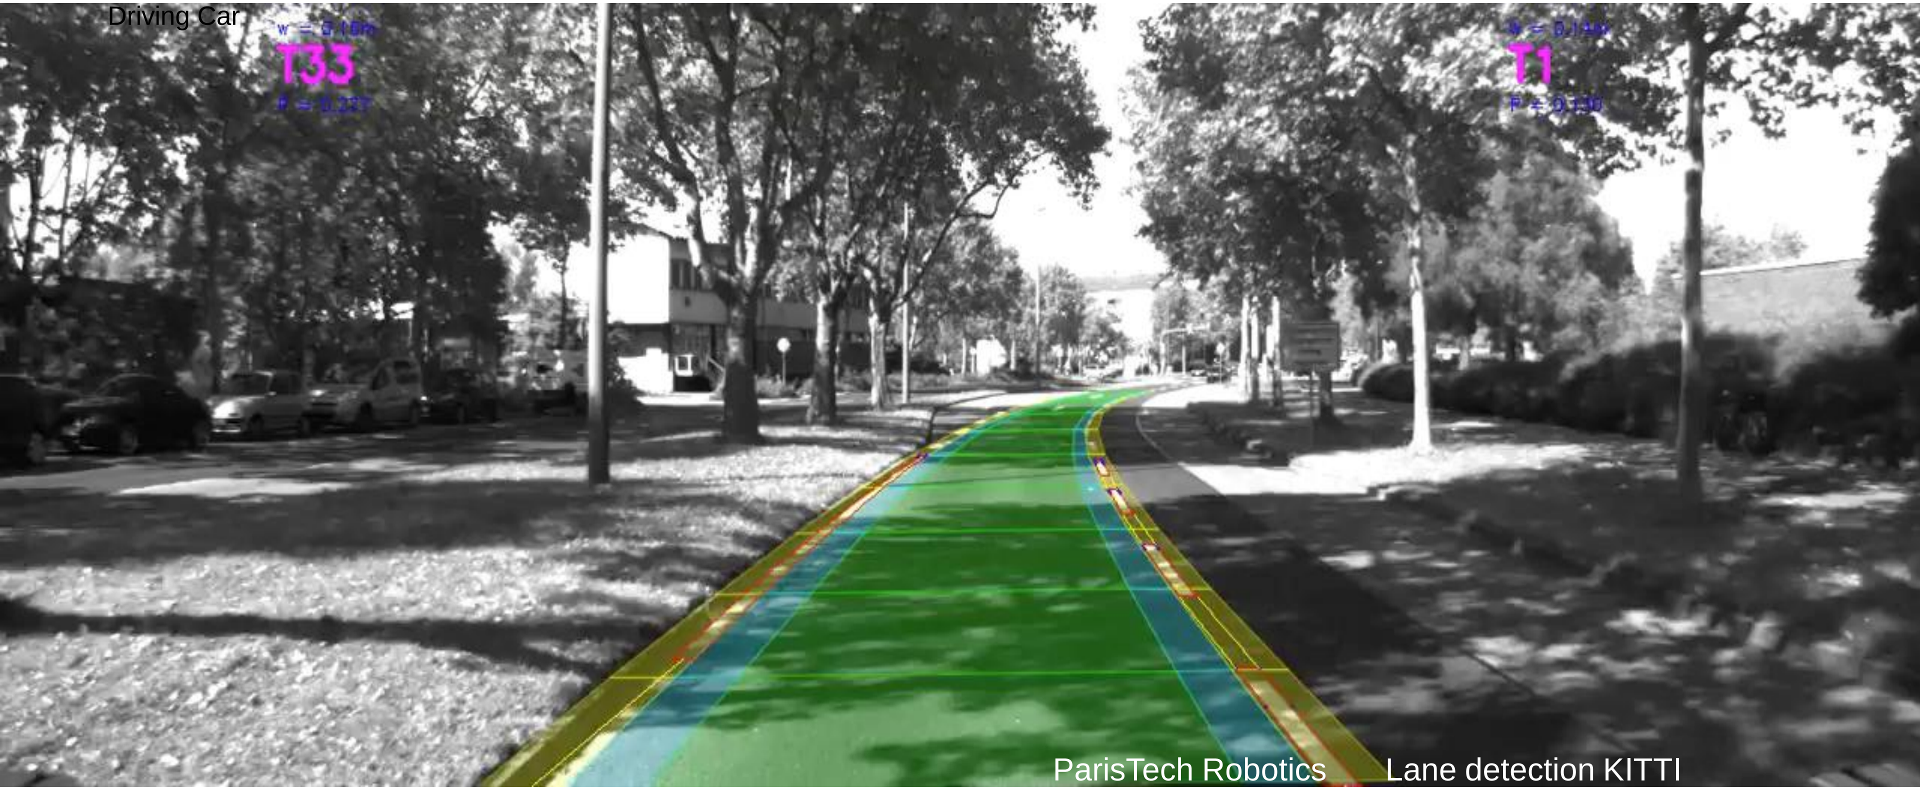

ParisTech Robotics

Lane detection KITTI

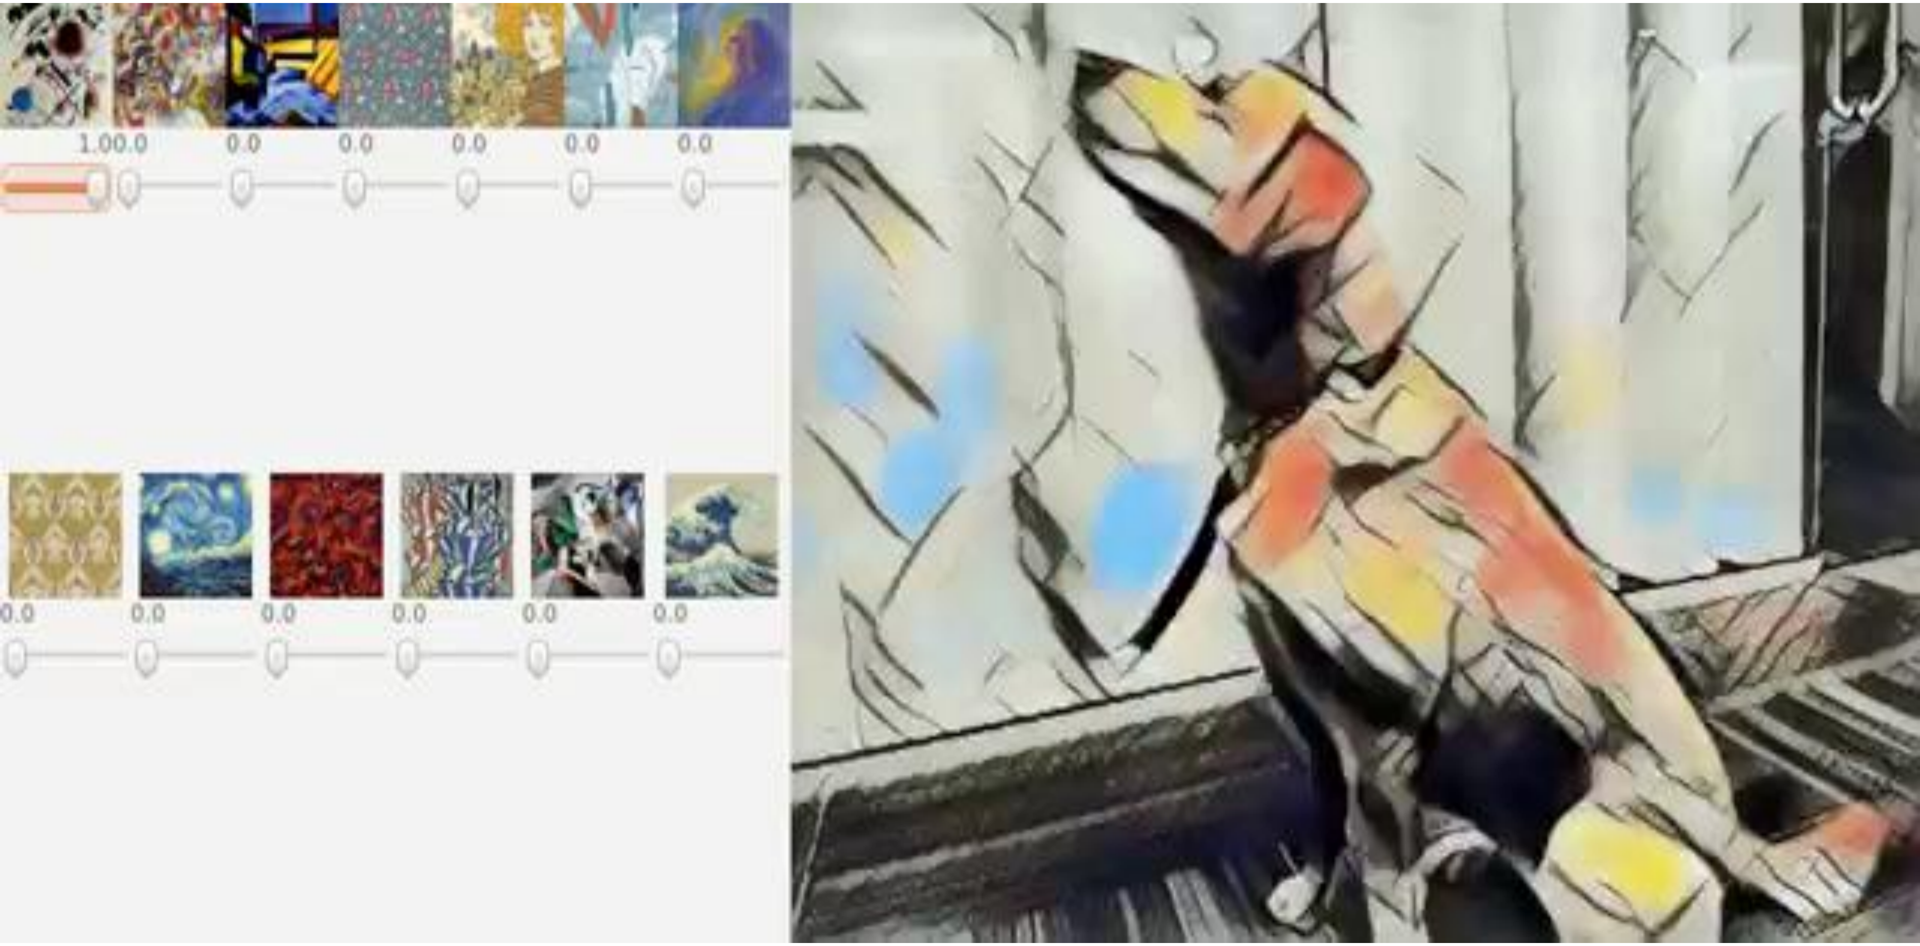

Leon A. Gatys, Alexander S. Ecker, Matthias Bethge , “Image Style Transfer Using Convolutional Neural Networks”, CVPR 2016

HACETTEPE
UNIVERSITY
COMPUTER
VISION LAB

- Shortage of specialists to provide such a large number of elderly and vulnerable people the sufficient medical and social care

- **Intelligent technologies for continuous monitoring of people eithe nursing homes are require**

  - **improve their quality of li**
  - **to reduce the cost of care**

Assistive technologies

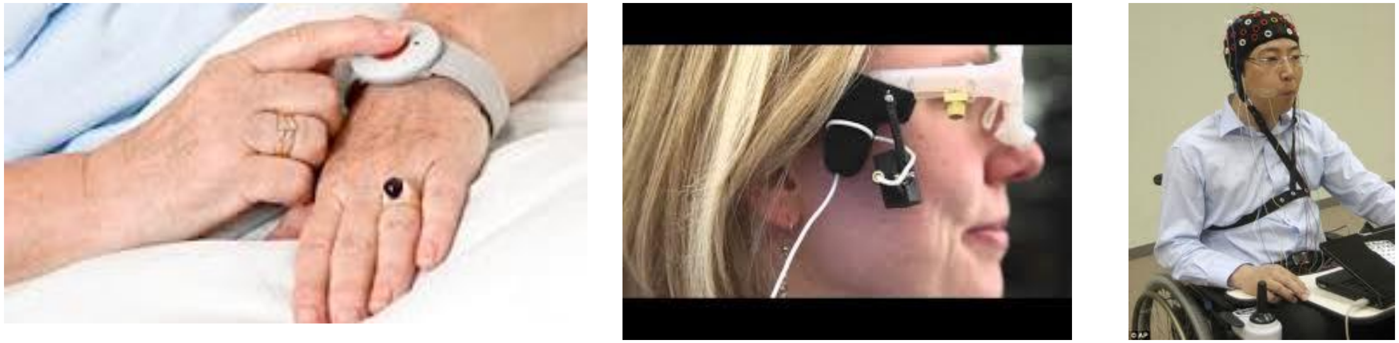

Disruptive

Smart cameras

Fall detection

It may be too late after fall !

Prediction is more important than detection

Capture usualness in unusual videos

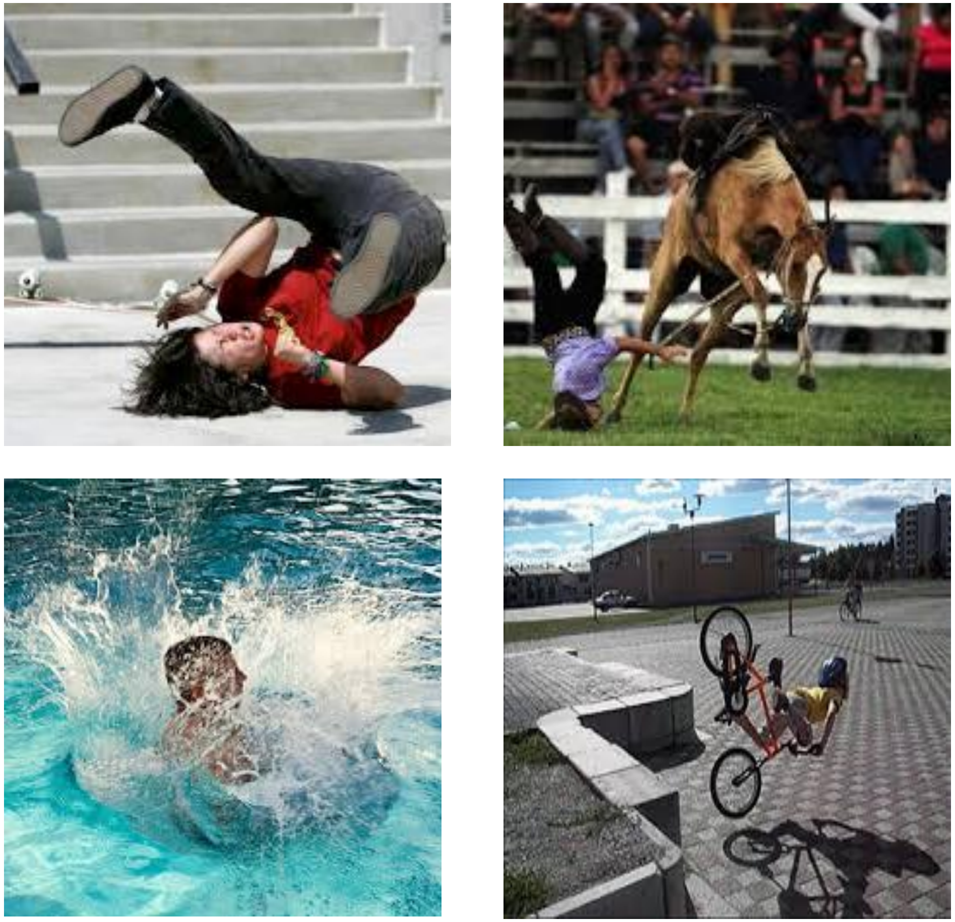

Video anomalies

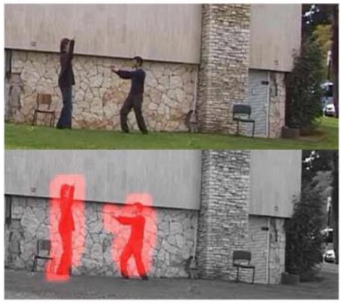

Roshtkhari and Levine, CVPR 2013

Boiman and Irani, ICCV 2005

Ito, Kitani, Bagnell, Hebert, 2012

Zhao, Fei-Fei, Xing, CVPR 2011

Video anomalies

Cooking Activities: High Intra-class Variance

Ahmet Iscen,
Anil Armagan,
Pinar Duygulu

2013 Workshop on Cooking and Eating Activities

Cooking Activities: Low Inter-class Variance

Ahmet Iscen,
Anil Armagan,
Pinar Duygulu

2013

Workshop on Cooking and Eating Activities

Cut apart, cut ends. cut slices, cut stripes, cut dice

Put in Pan or Put in Bowl?

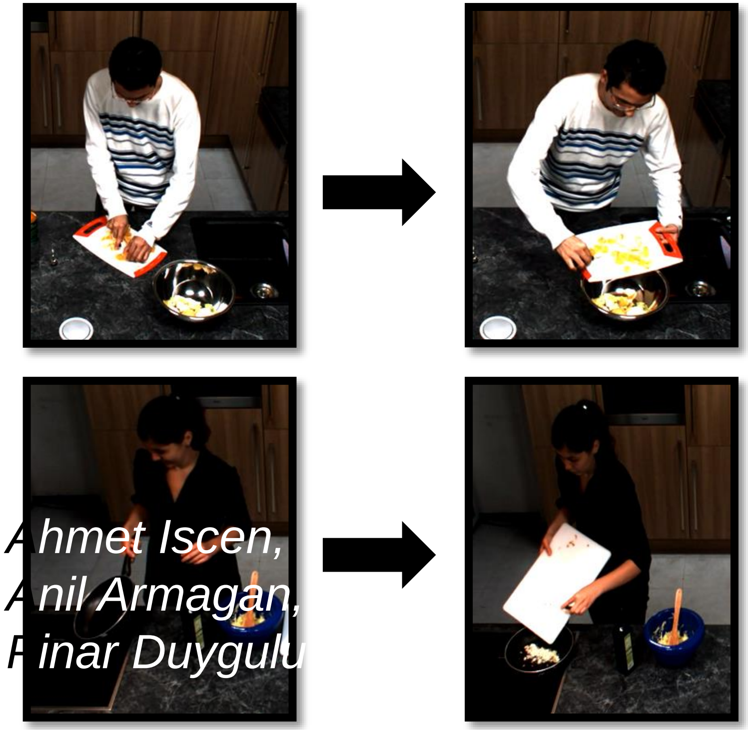

P(“put in bowl” | “cut dice”) > P(“put in pan” | “cut dice”)

P(“put in pan” | “spread”) > P(“put in bowl” | “spread”)

Ahmet Iscen,
Anil Armagan,
Pinar Duygulu

2013 Workshop on Cooking and Eating Activities

Medical Device Use

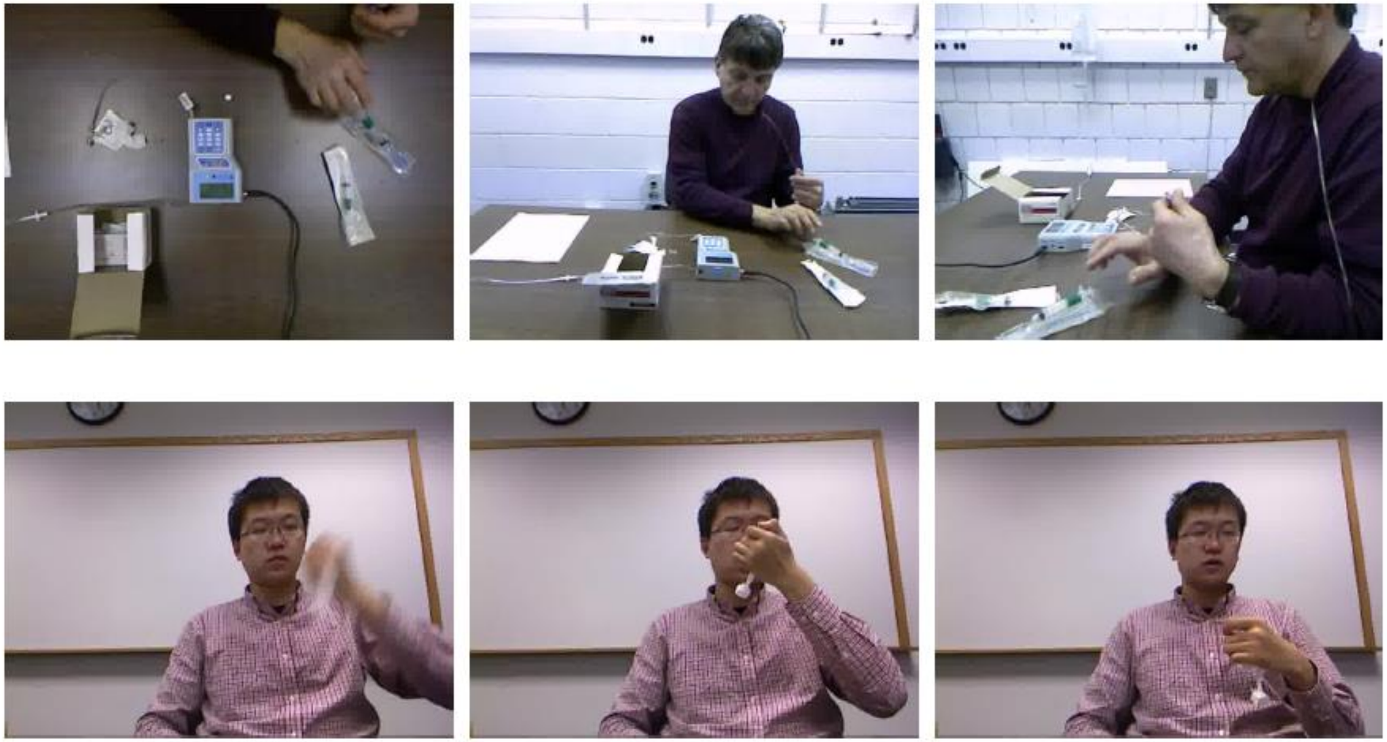

Ahmet Iscen,
Pinar Duygulu

2014

Workshop on Assistive computer Vision and Robotics

Asthma Inhaler

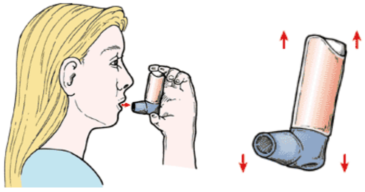

- Shake the inhaler (for 5 second)
- Breathe out
- Put the inhaler about 2 inches in front of your mouth
- Breathe in and push down the button at the same time
- Hold your breath for 10 seconds
- Breathe out slowly

Ahmet Iscen,
Pinar Duygulu

2014

Workshop on Assistive computer Vision and Robotics

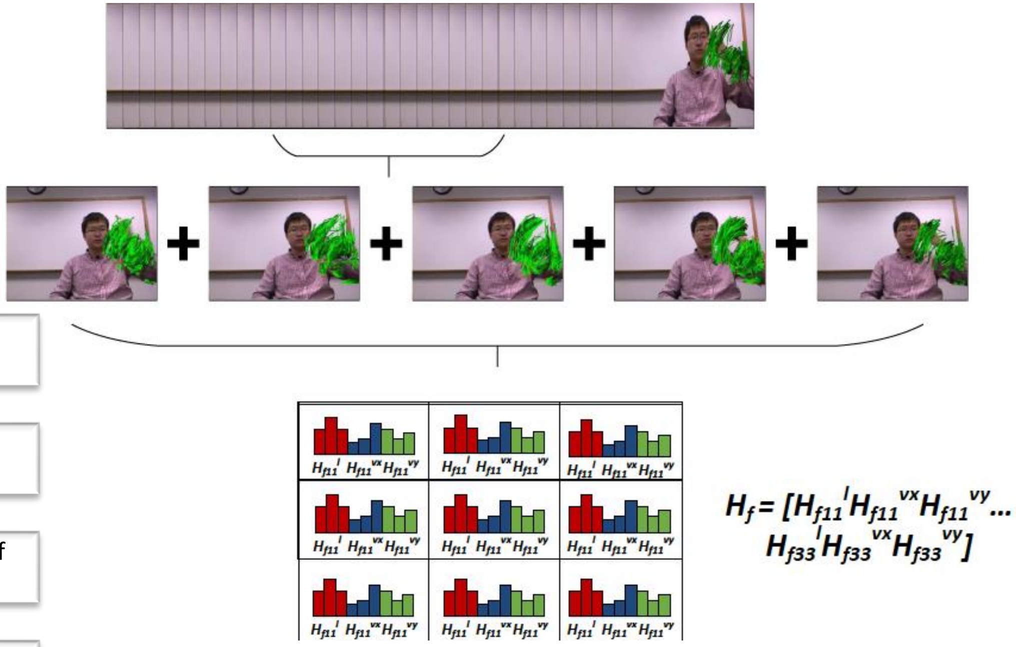

$$H_f = [H_{f11}^{l} H_{f11}^{vx} H_{f11}^{vy} \ldots H_{f33}^{l} H_{f33}^{vx} H_{f33}^{vy}]$$

	Trajectory	HOG	HOF	MBH	Snippet Hist
Recall	95.31	50.00	**100.00**	87.50	98.44
Precision	91.04	22.70	91.43	71.79	**100.00**
F-score	93.13	31.22	95.52	78.87	**99.21**

Infusion Pump

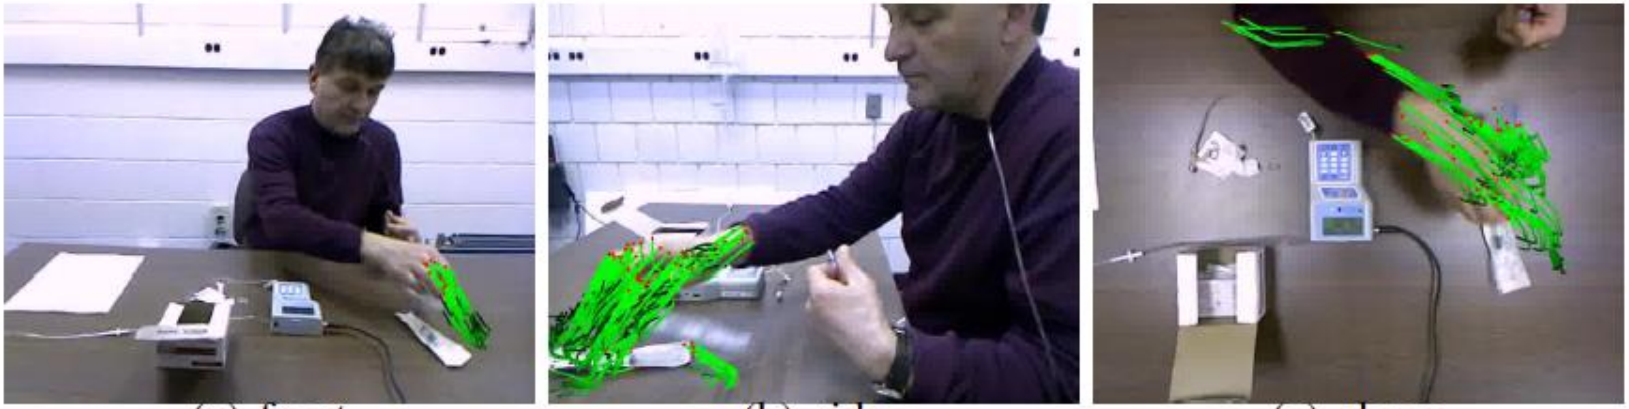

(a) front (b) side (c) above

Ahmet Iscen,
Pinar Duygulu

2014

Workshop on Assistive computer Vision and Robotics

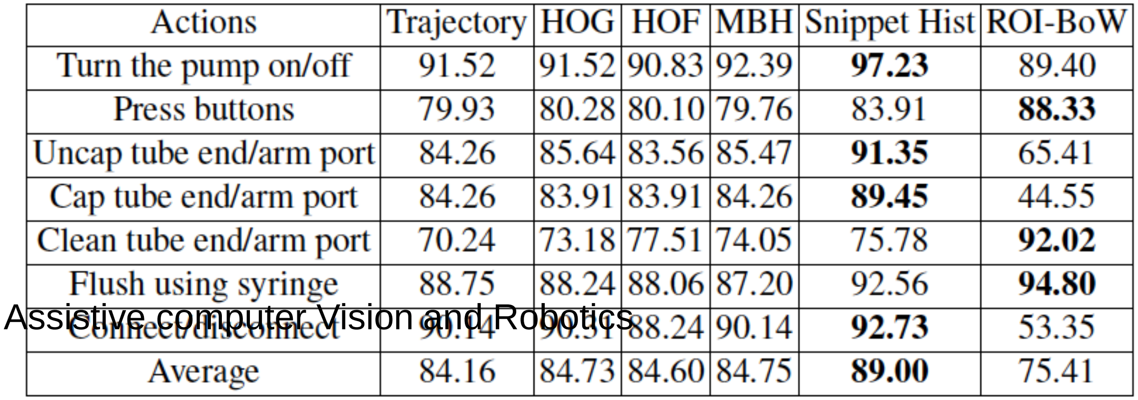

Actions	Trajectory	HOG	HOF	MBH	Snippet Hist	ROI-BoW
Turn the pump on/off	91.52	91.52	90.83	92.39	**97.23**	89.40
Press buttons	79.93	80.28	80.10	79.76	83.91	**88.33**
Uncap tube end/arm port	84.26	85.64	83.56	85.47	**91.35**	65.41
Cap tube end/arm port	84.26	83.91	83.91	84.26	**89.45**	44.55
Clean tube end/arm port	70.24	73.18	77.51	74.05	75.78	**92.02**
Flush using syringe	88.75	88.24	88.06	87.20	92.56	**94.80**
Connect/disconnect	90.14	90.31	88.24	90.14	**92.73**	53.35
Average	84.16	84.73	84.60	84.75	**89.00**	75.41

Associating controlled data with recordings

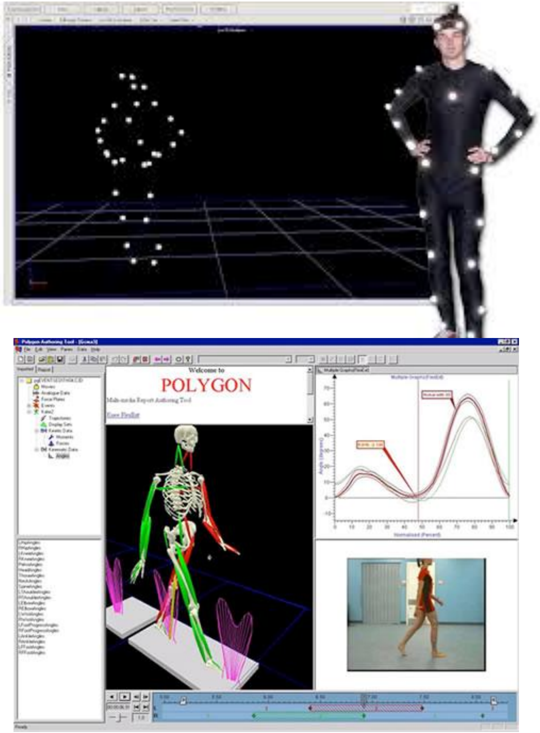

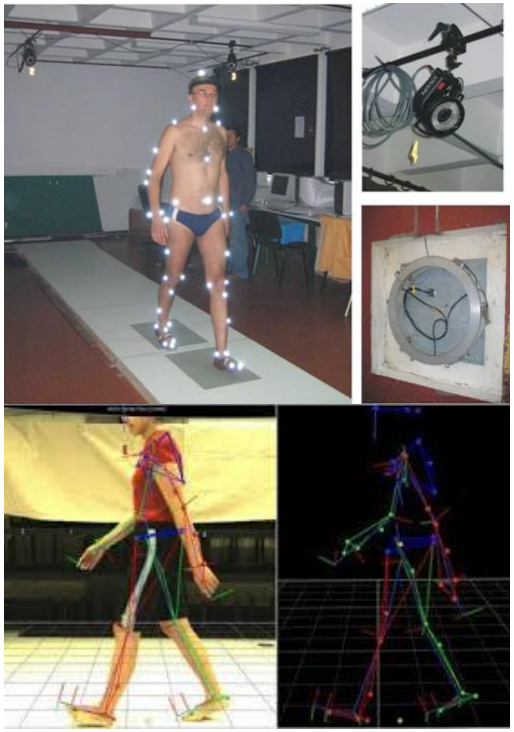

Linking with medical records

Inspired by machine translation

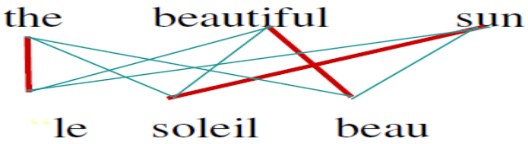

Human Action Recognition

- There are various types of actions/activities

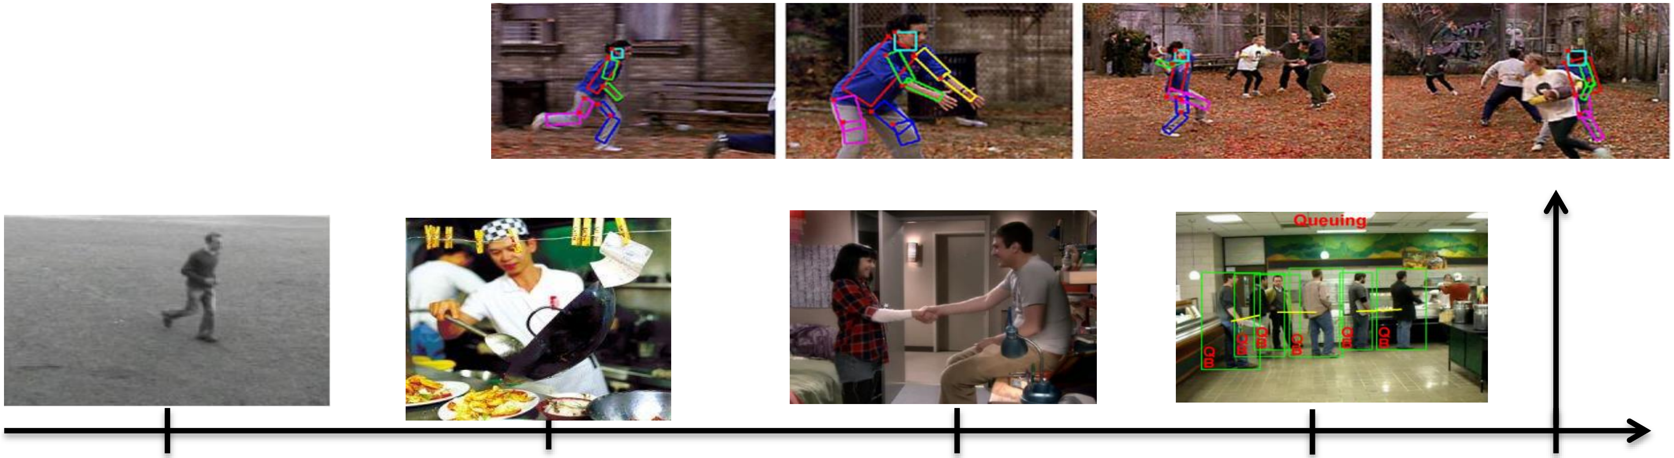

Single Actions

Actions in
Still Images

Interactions

Collective
Activities

Complex
Compositive
Activities

- The ultimate goal is to make computer recognize all of them reliably.

Gland segmentation on Histology images

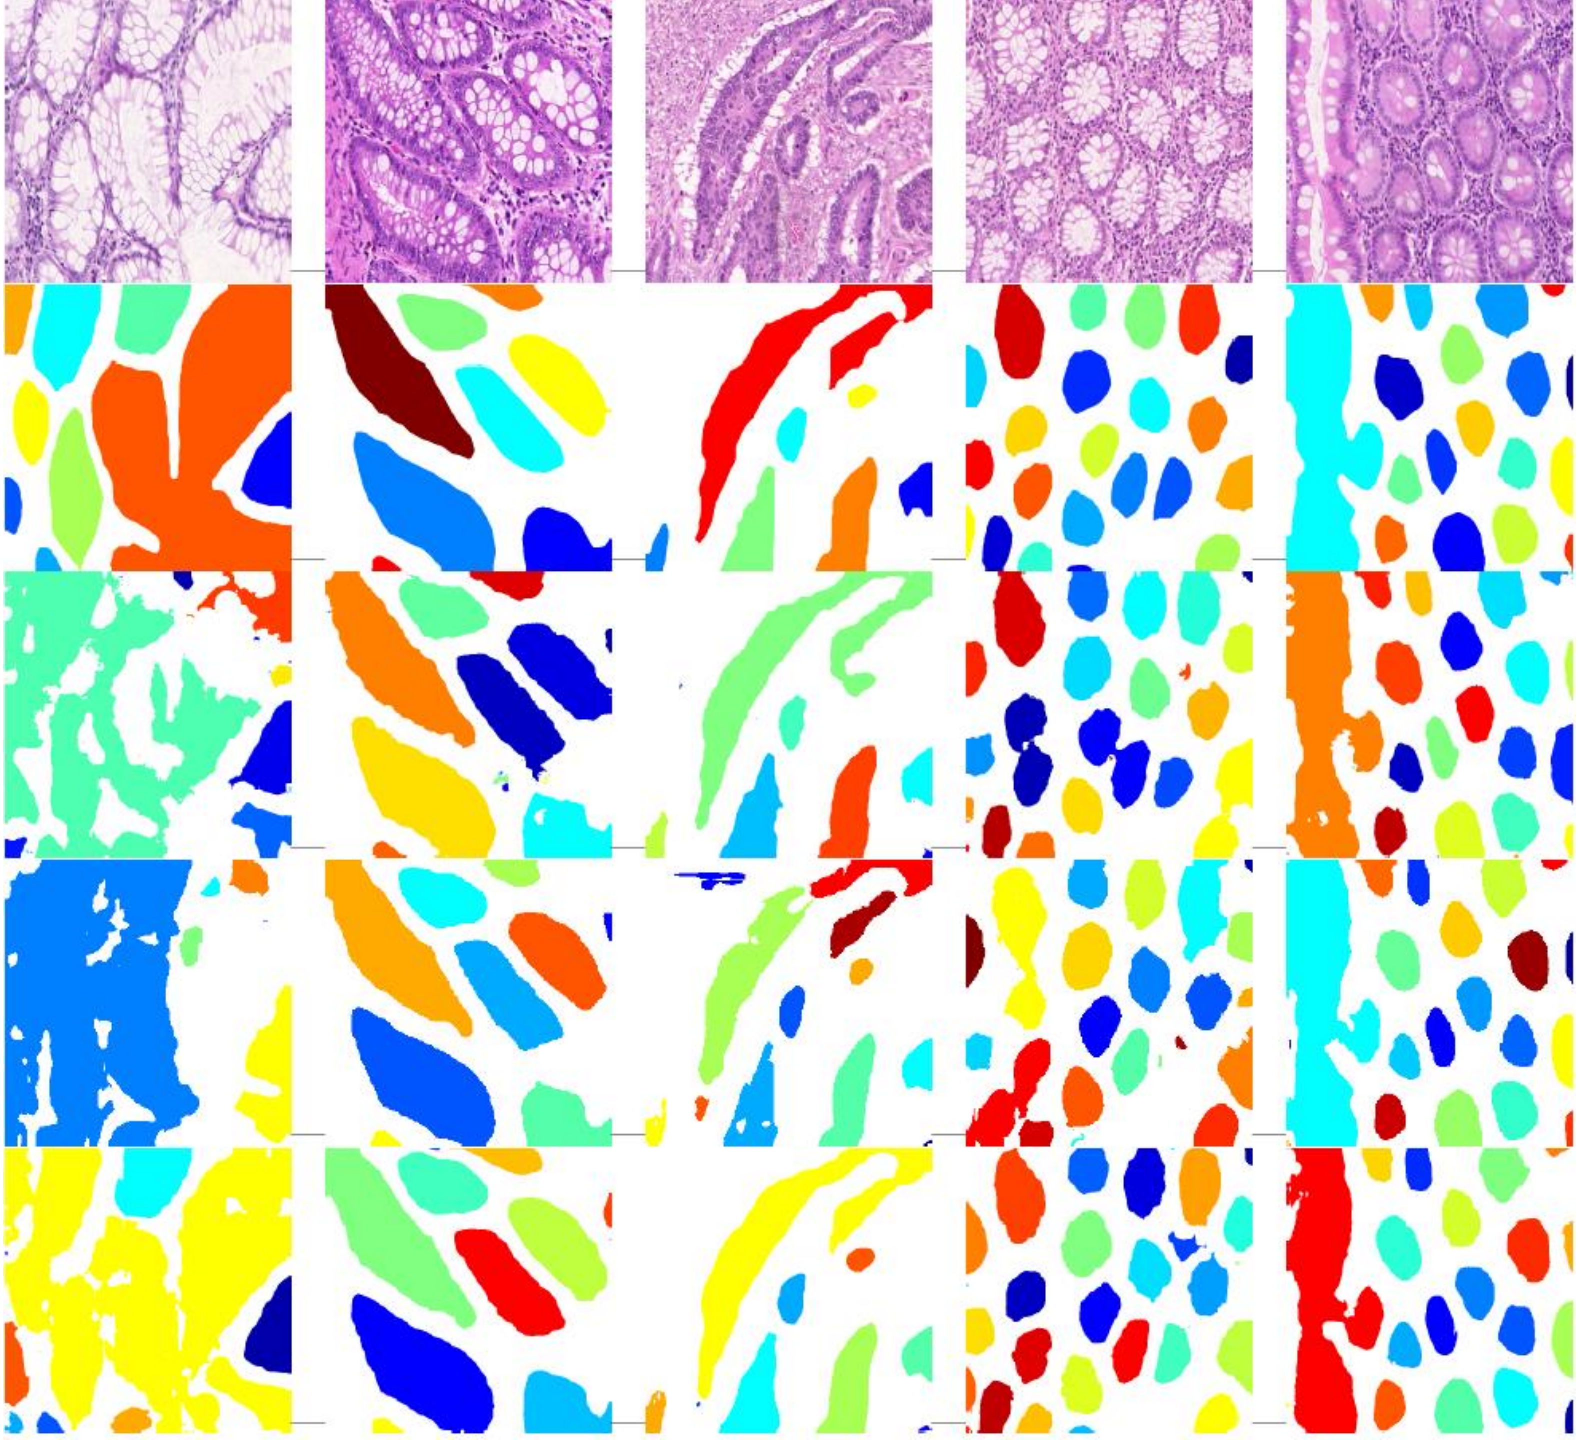

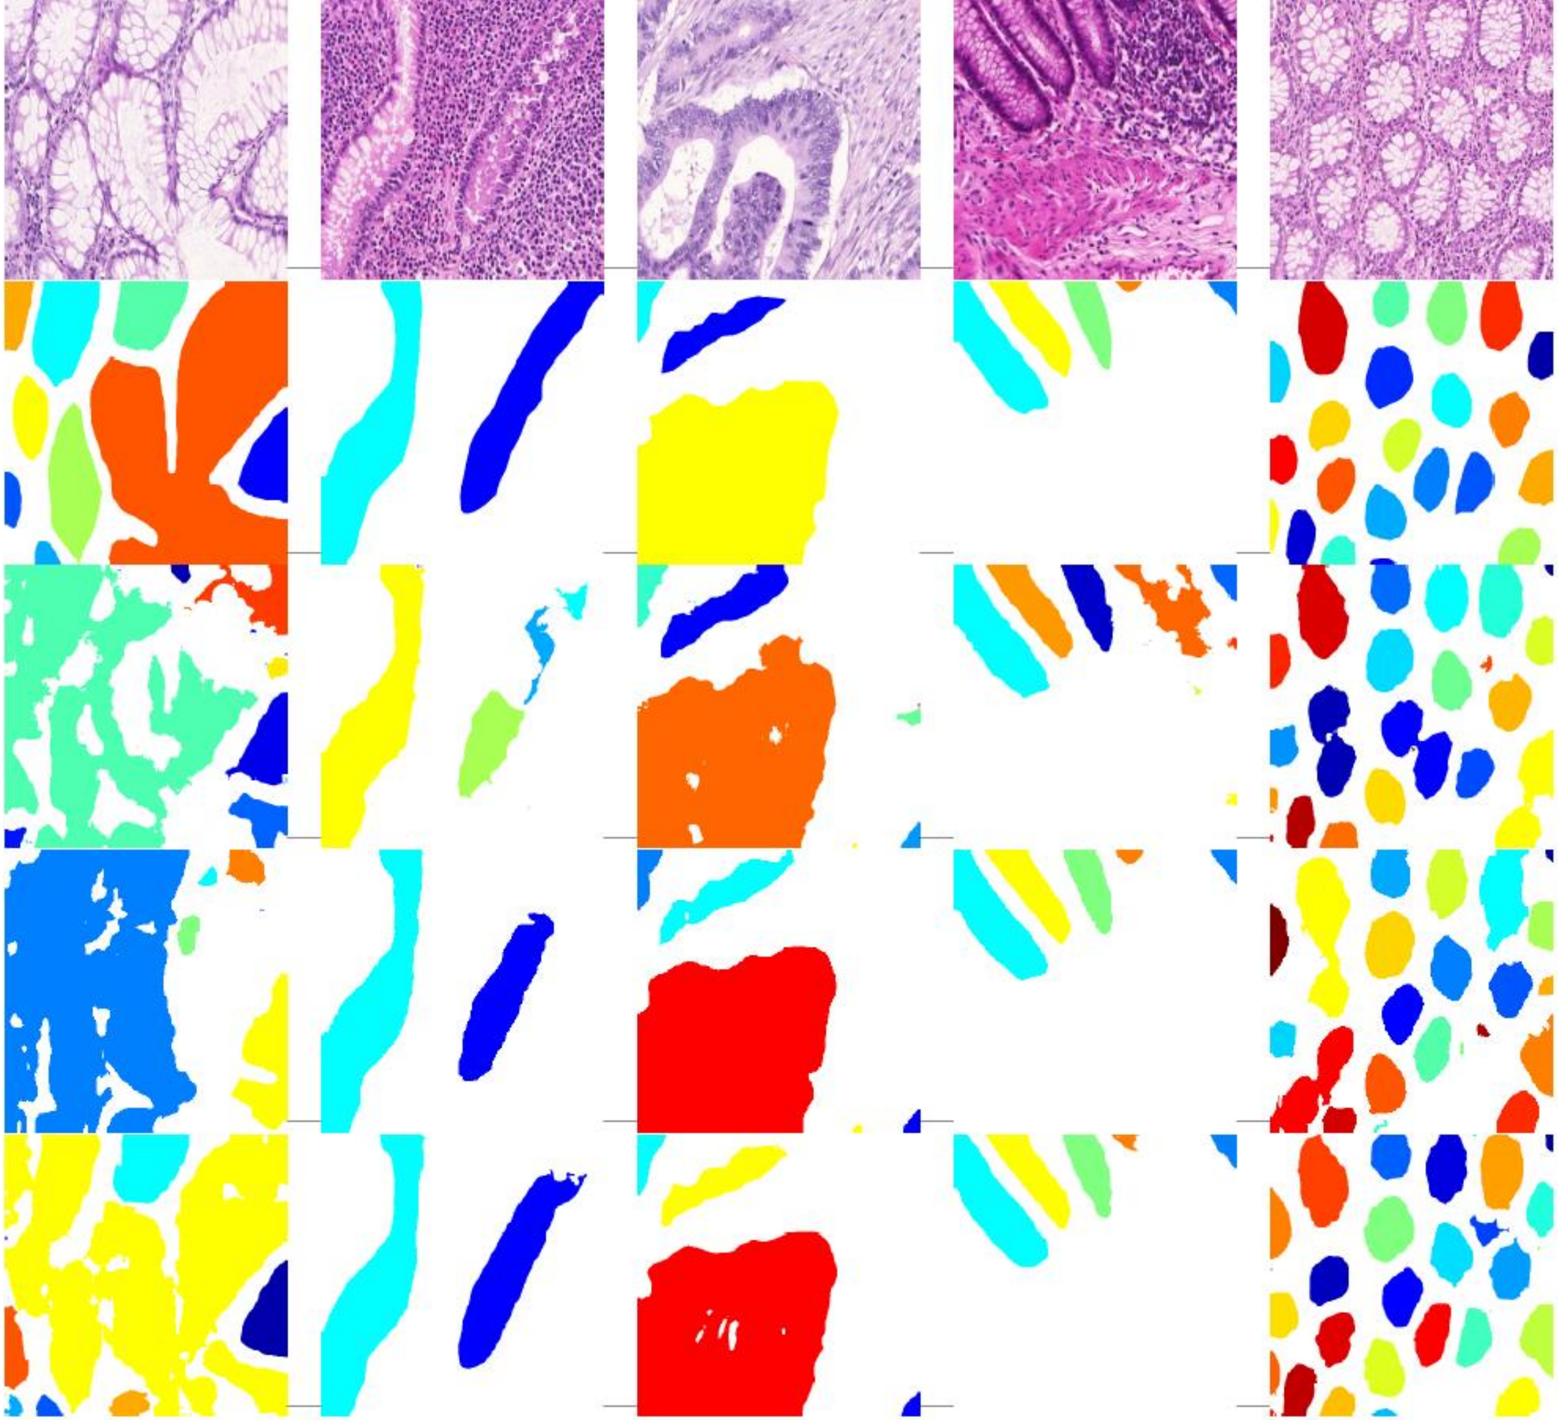

HACETTEPE
UNIVERSITY
COMPUTER
VISION LAB

Contributors

- Ozge Yalcinkaya
- Eren Golge
- Sermetcan Baysal
- Samet Hicsonmez
- Nermin Samet
- Fadime Sener

“It is not the strongest species that survive, nor the most intelligent, but the ones most responsive to change.”

Leon C. Megginson, paraphrasing Charles Darwin, 1963