POLITECNICO
MILANO 1863

Artificial Neural Networks and Deep Learning

- Machine Learning vs Deep Learning-

Matteo Matteucci, PhD (matteo.matteucci@polimi.it)
Artificial Intelligence and Robotics Laboratory
Politecnico di Milano

AIRLAB
ARTIFICIAL INTELLIGENCE AND ROBOTICS LAB

«Me, Myself, and I»

Matteo Matteucci, PhD
Associate Professor
Dept. of Electronics, Information & Bioengineering
Politecnico di Milano
matteo.matteucci@polimi.it

My research interests

- Robotics & Autonomous Systems
- Machine Learning
- Pattern Recognition
- Computer Vision & Perception

Courses I teach

- Robotics (BS+MS)
- Machine Learning (MS)
- Deep Learning (MS+PhD)
- Cognitive Robotics (MS)

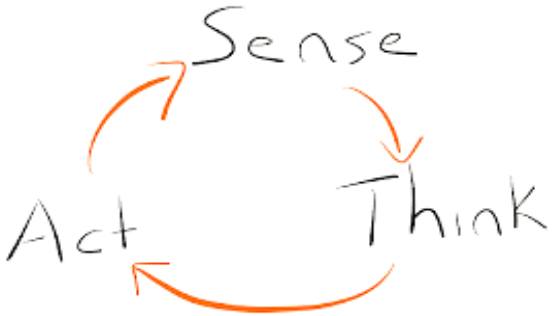

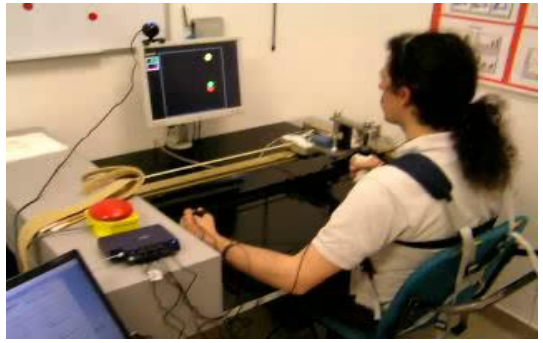

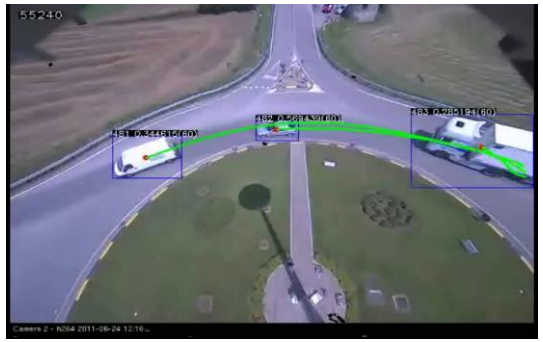

Enable physical and software autonomous systems to perceive, plan, and act without human intervention in the real world

About Machine Learning & De

Over the top of Emerging Technologies Hype Cycle!!

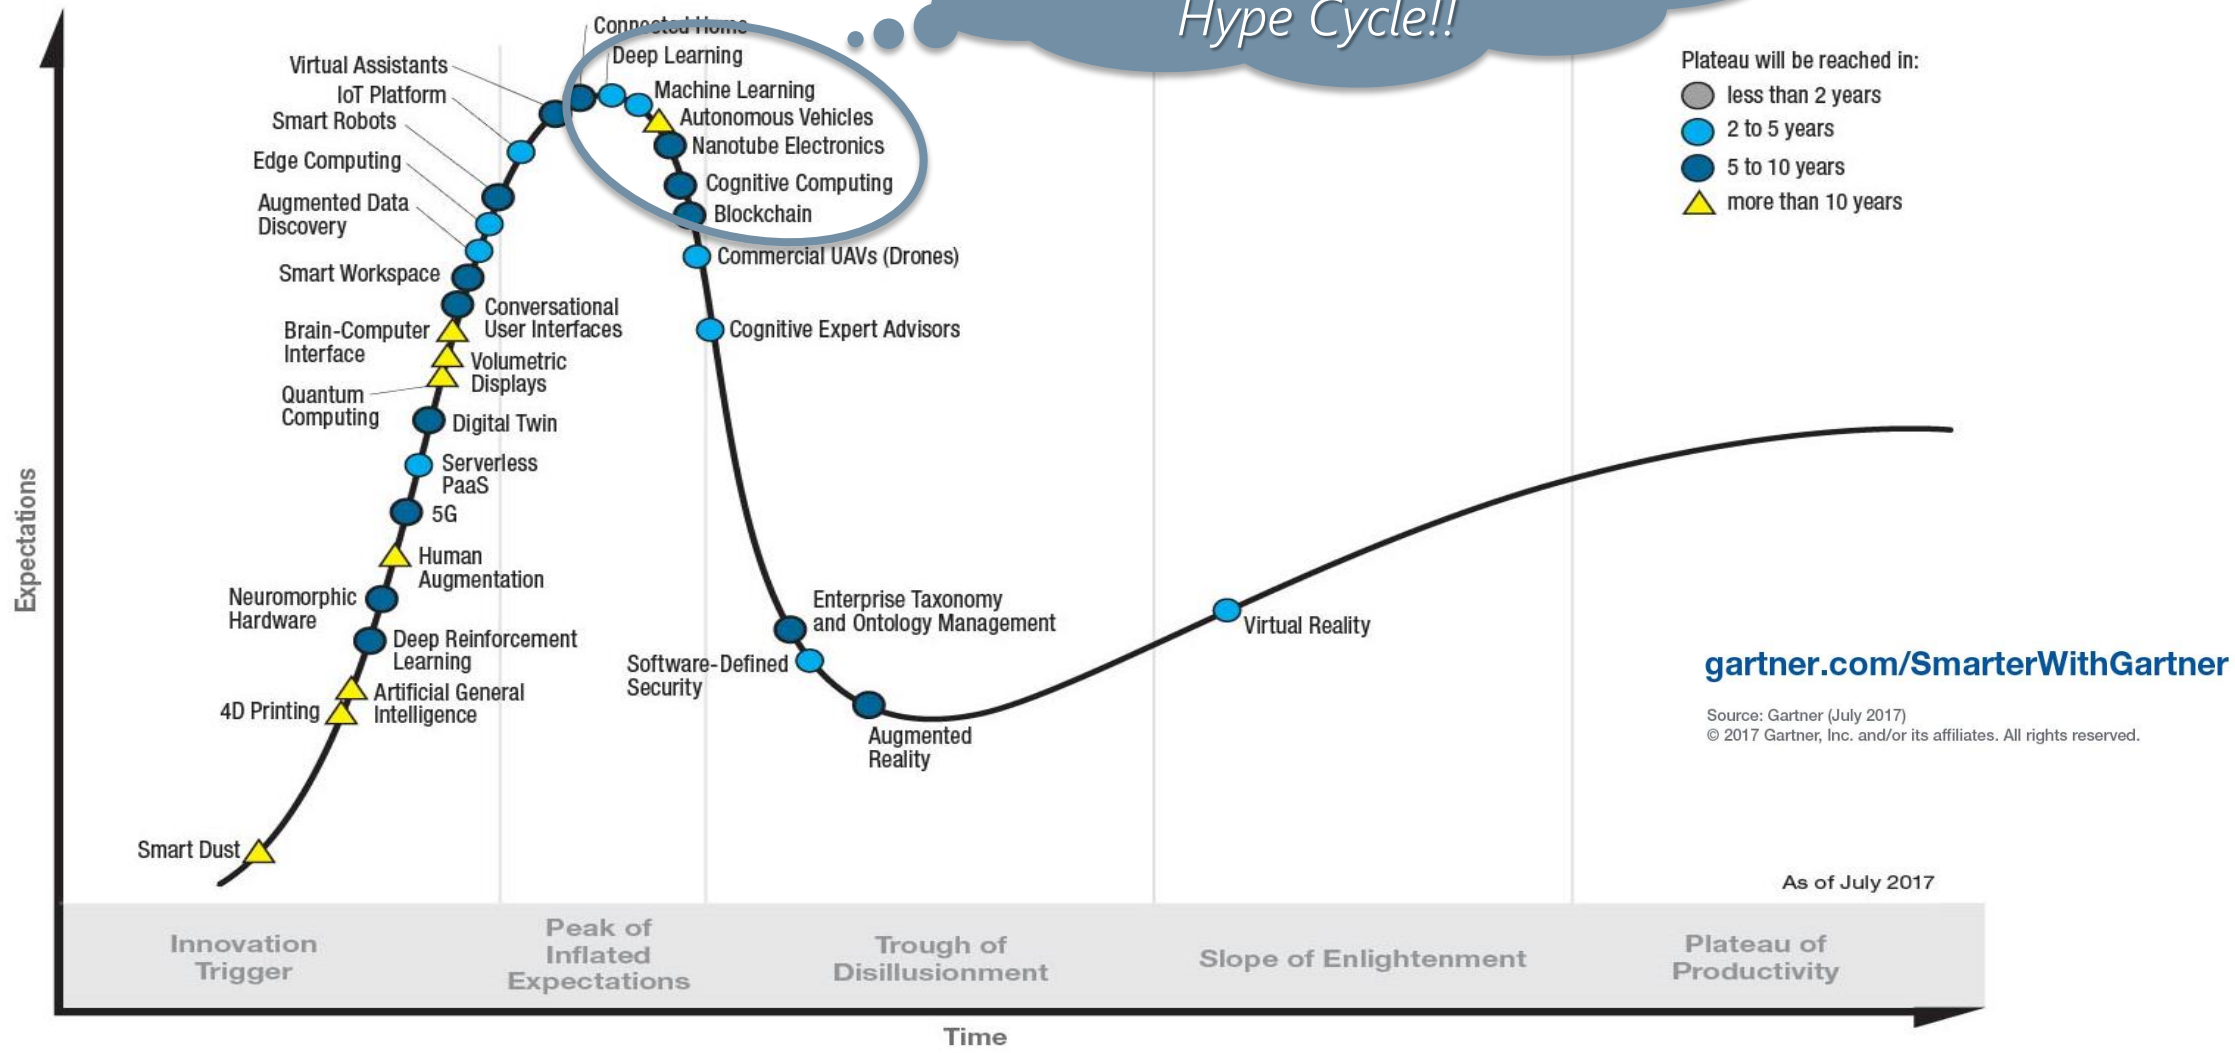

«Deep Learning is not AI, nor Machine Learning»

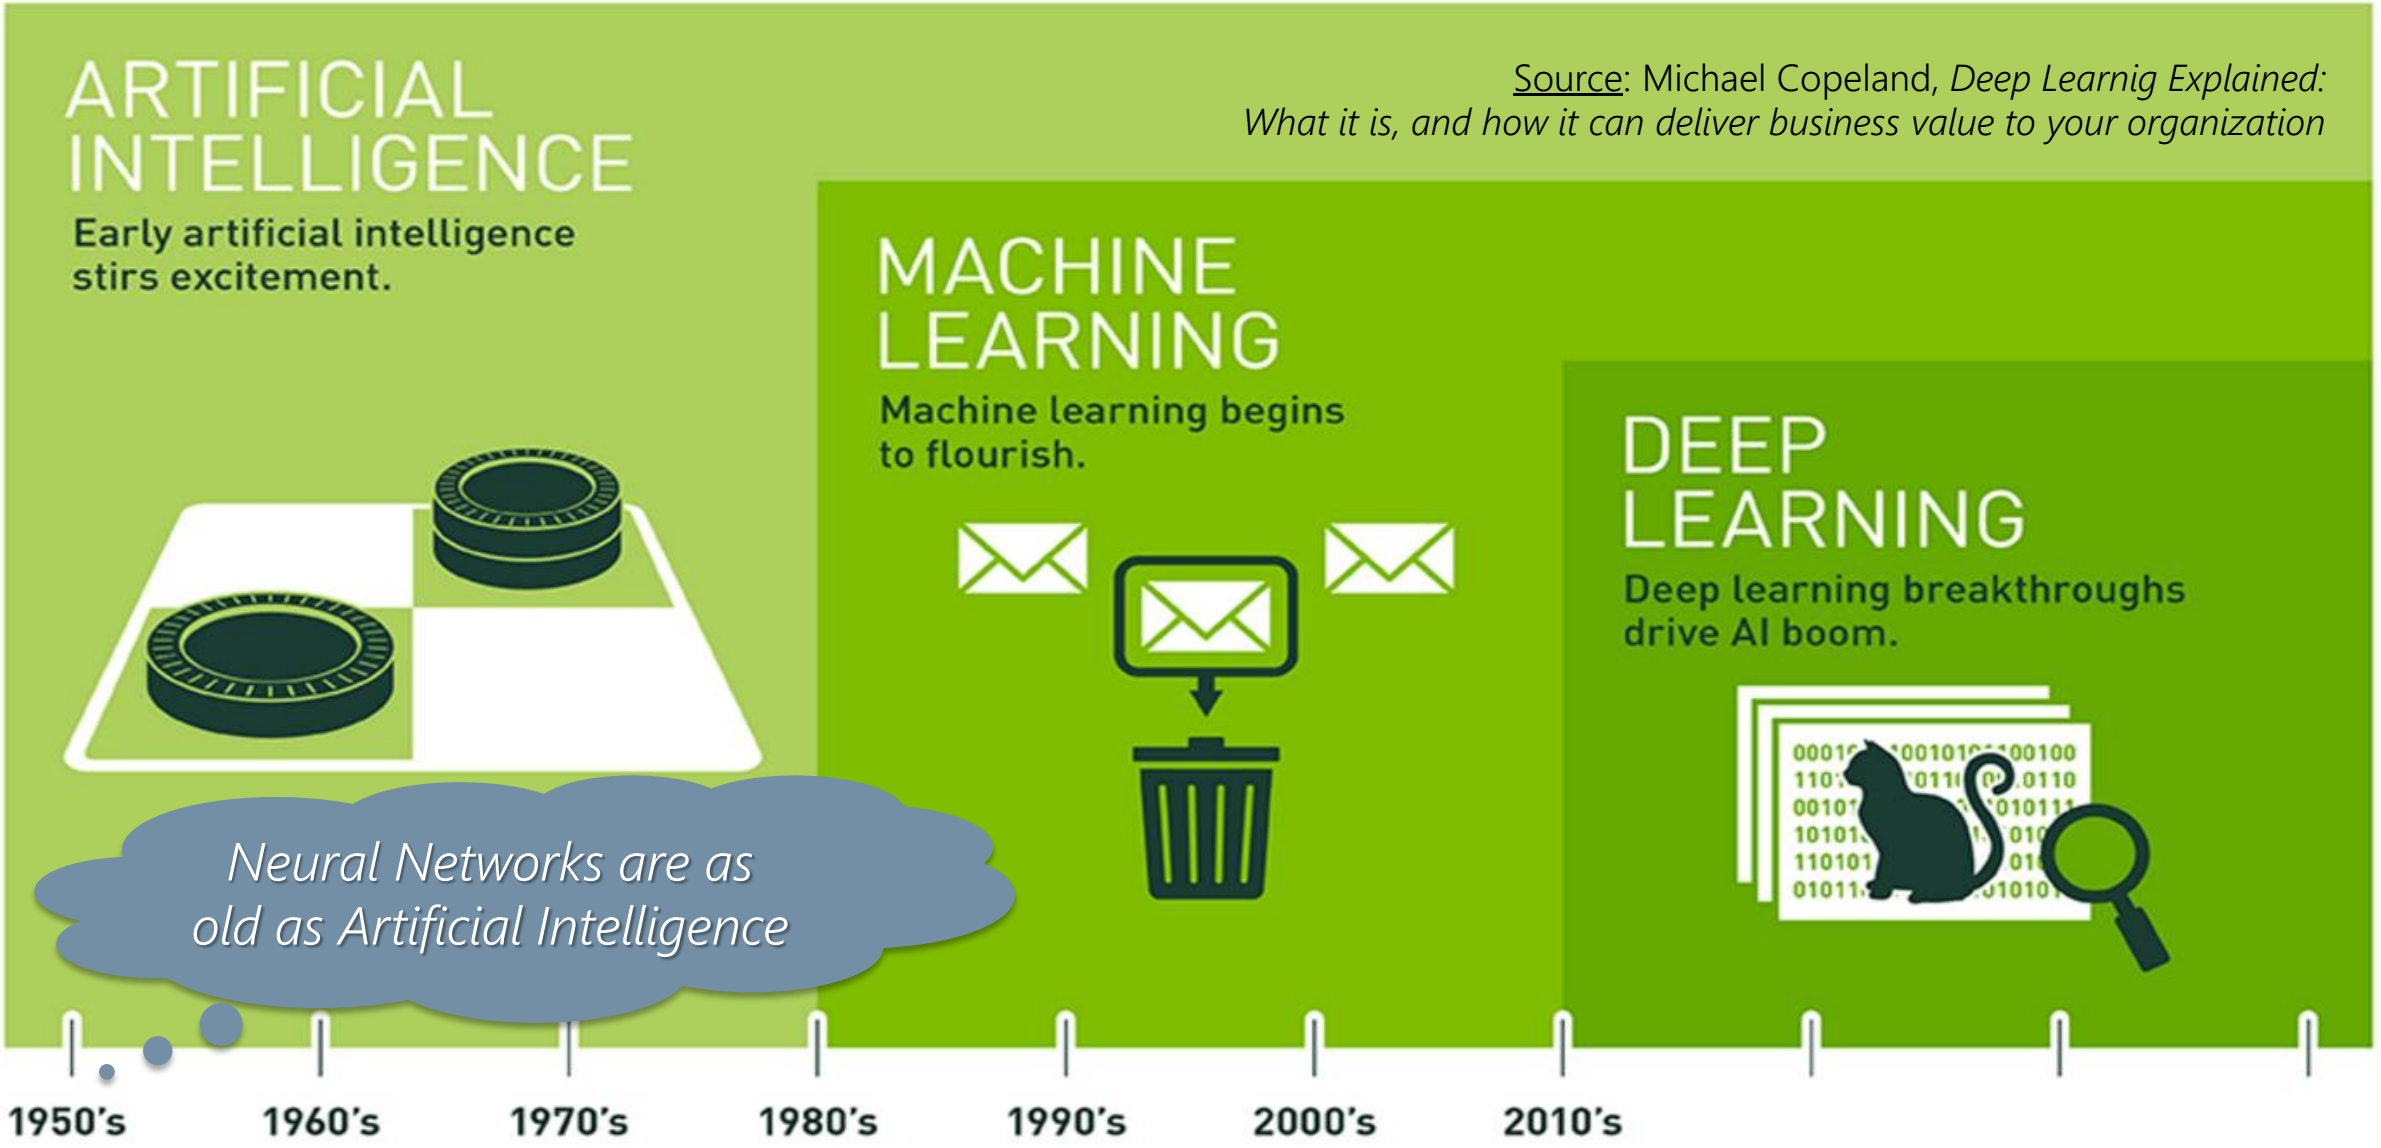

Source: Michael Copeland, *Deep Learnig Explained: What it is, and how it can deliver business value to your organization*

Machine Learning

Machine learning is a category of research and algorithms focused on finding patterns in data and using those patterns to make predictions. Machine learning falls within the artificial intelligence (AI) umbrella, which in turn intersects with the broader field of knowledge discovery and data mining.

Source: SAS, 2014 and PwC, 2016 *and myself, 2017*

Machine Learning

Machine Learning (Tom Mitchell – 1997)

T = Regression/Classification/…
E = Data
P = Errors/Loss

"A computer program is said to learn from experience E with respect to some class of task T and a performance measure P, if its performance at tasks in T, as measured by P, improves because of experience E."

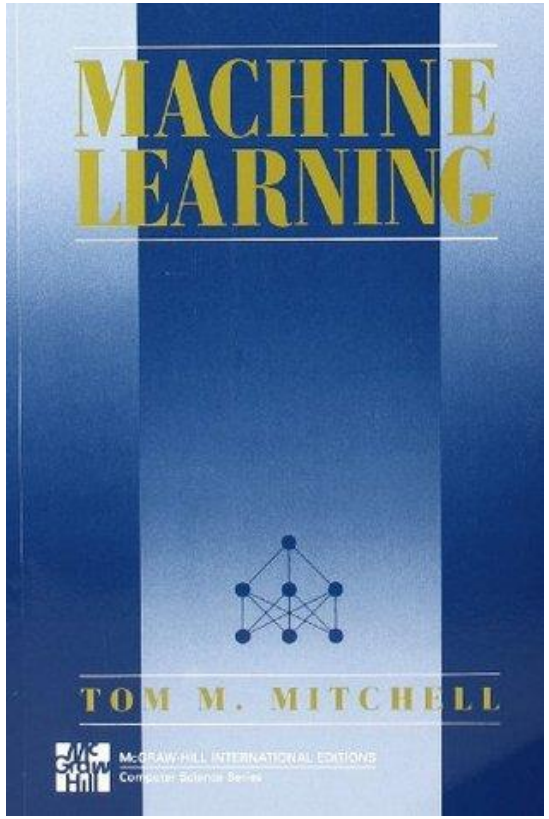

Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- ***Supervised learning***: given the desired outputs $t_1, t_2, t_3, \dots, t_N$ learn to produce the correct output given a new set of input
- ***Unsupervised learning***: exploit regularities in D to build a representation to be used for reasoning or prediction
- ***Reinforcement learning***: producing actions $a_1, a_2, a_3, \dots, a_N$ which affect the environment, and receiving rewards $r_1, r_2, r_3, \dots, r_N$ learn to act in order to maximize rewards in the long term

Reinforcement Learning is Wellknown

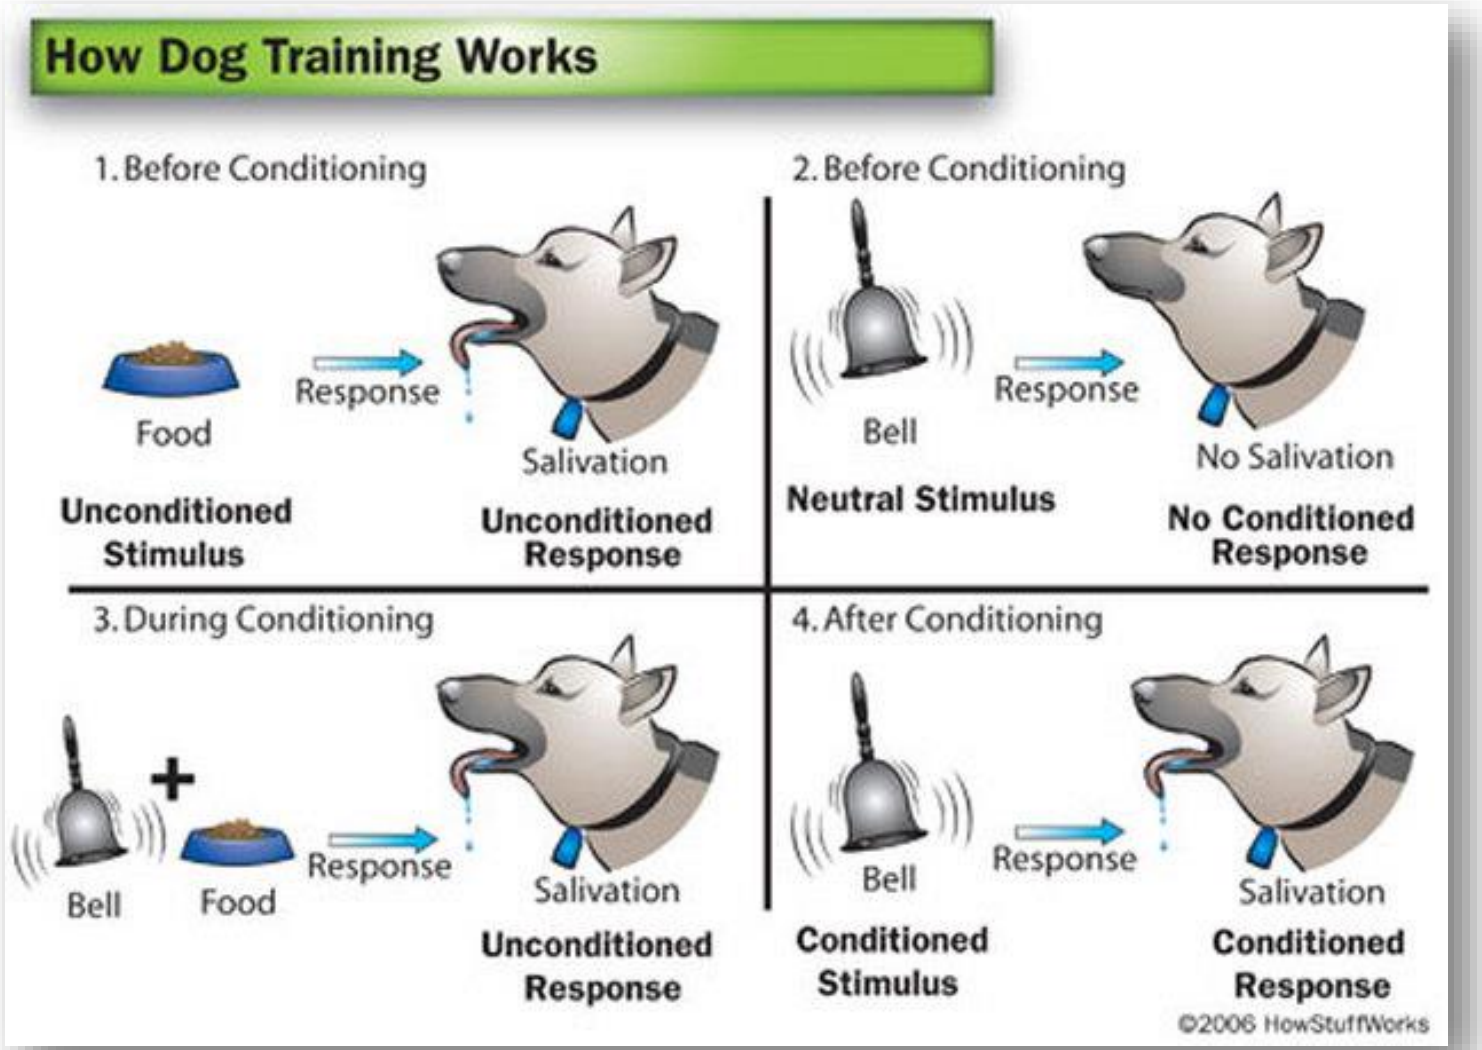

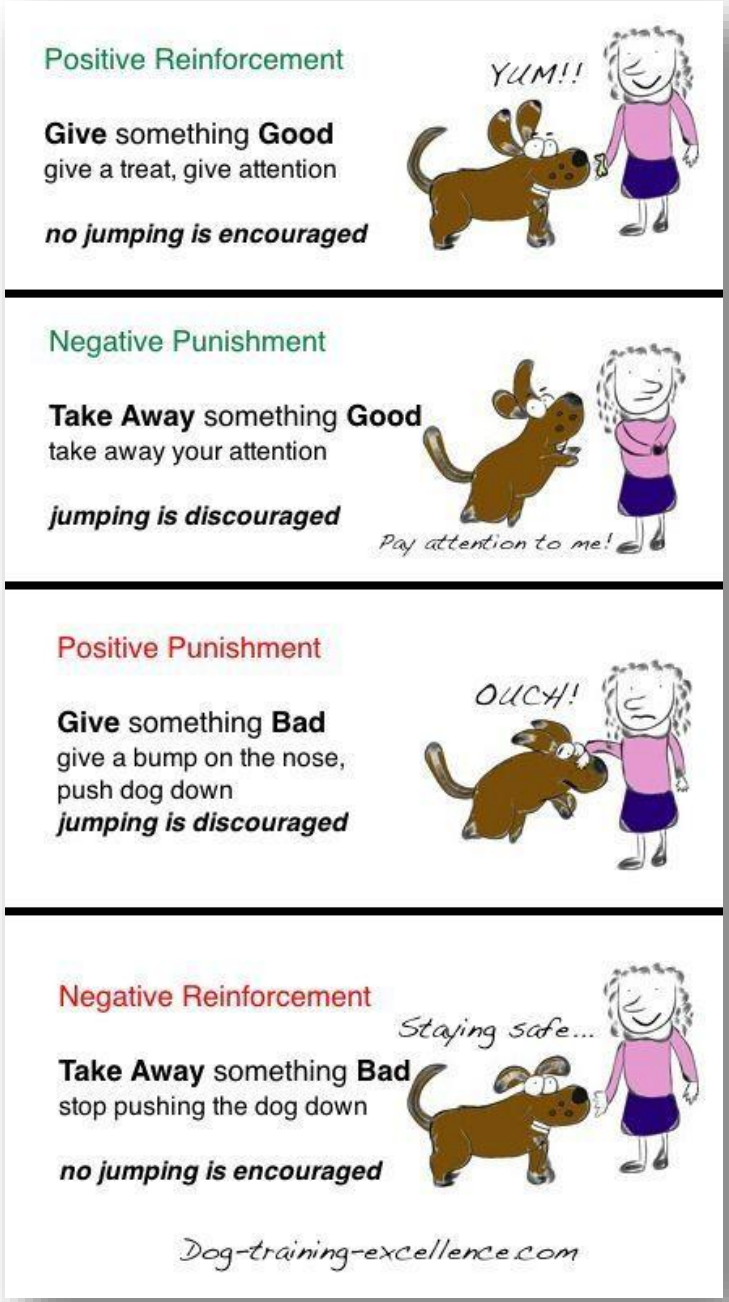

Reinforcement Learning

Let's our machine be an agent interacting with an unknown environment

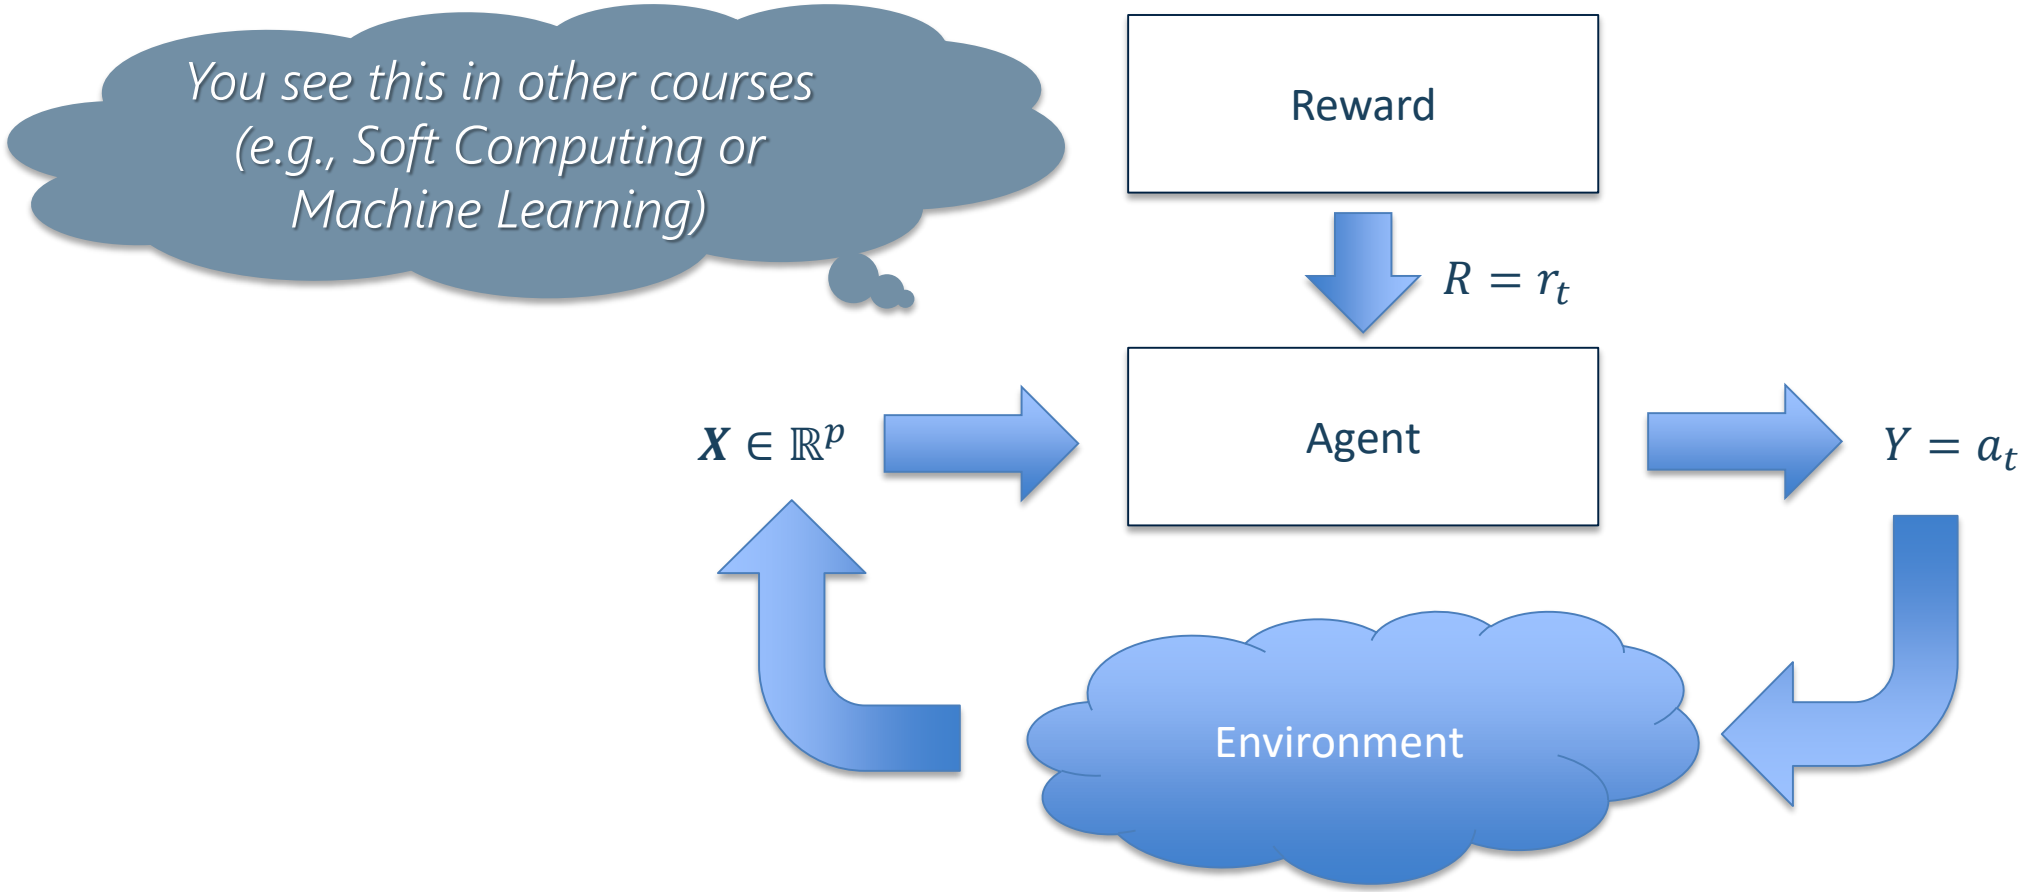

Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- ***Supervised learning***: given the desired outputs $t_1, t_2, t_3, \dots, t_N$ learn to produce the correct output given a new set of input
- ***Unsupervised learning***: exploit regularities in D to build a representation to be used for reasoning or prediction
- ***Reinforcement learning***: producing actions $a_1, a_2, a_3, \dots, a_N$ which affect the environment, and receiving rewards $r_1, r_2, r_3, \dots, r_N$ learn to act in order to maximize rewards in the long term

Supervised learning: Classification

Cars

Motorcy

Learning is about modeling …

Hand-crafted Features

Learned Classifier

Motorcycle

Terminology in Classification

- Input
- Features
- Observations
- Independent Variables

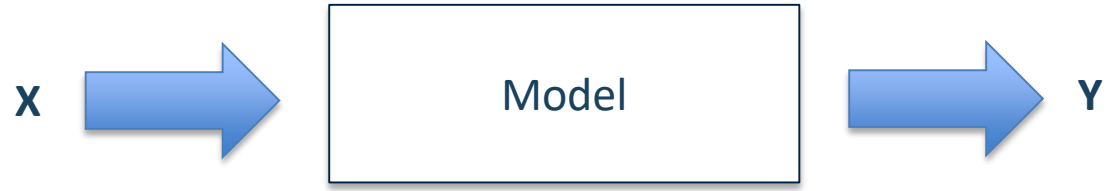

- Output
- Class
- Dependent Variable

- Classifier
- Inductive Hypothesis
- Learning Machine
- …

Supervised learning: Regression


12000 $


15000 $

6000 $


2000 $


8000 $


22000 $


4000 $


28000 $


6000 $


35000 $

Hand-crafted Features

Learned Regressor

3800 $

Terminology in Regression

- Input
- Predictor
- Observations
- Independent Variable

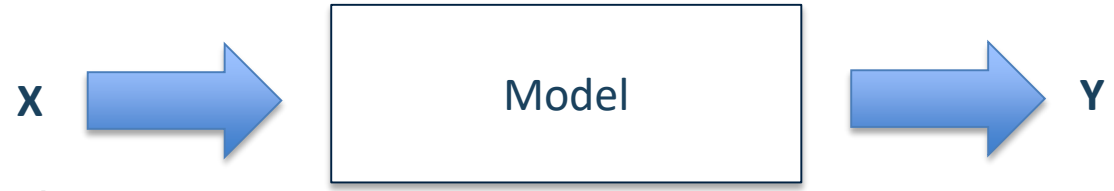

- Output
- Prediction
- Response
- Dependent Variable

- Model
- Function
- Inductive Hypothesis
- Learning Machine
- …

Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- ***Supervised learning***: given the desired outputs $t_1, t_2, t_3, \dots, t_N$ learn to produce the correct output given a new set of input
- ***Unsupervised learning***: exploit regularities in D to build a representation to be used for reasoning or prediction
- ***Reinforcement learning***: producing actions $a_1, a_2, a_3, \dots, a_N$ which affect the environment, and receiving rewards $r_1, r_2, r_3, \dots, r_N$ learn to act in order to maximize rewards in the long term

Unsupervised learning: Clustering



















Unsupervised learning: Clustering



















Unsupervised learning: Clustering




















Unsupervised learning: Clustering




















Unsupervised learning: Clustering




















Unsupervised learning: Clustering



















Unsupervised learning: Clustering



















Notation in Brief

You see this in other courses (e.g., Datamining)

In this case the training dataset is given by a set of input records

$$D =< x_1 >< x_2 >< x_3 >< \cdots >< x_N >$$

The task is to produce a representation of the data which highlights some knowledge about its organization.

Sometimes this knwoledge is named «patterns» …

Machine Learning Paradigms

Immagine you have a certain experience E, i.e., data, and let's name it

$$D = x_1, x_2, x_3, \dots, x_N$$

- *Supervised learning*: given the desired outputs $t_1, t_2, t_3, \dots, t_N$ learn to produce the correct output given a new set of input
- *Unsupervised learning*: exploit regularities in D to build a representation to be used for reasoning or prediction
- *Reinforcement learning*: producing actions $a_1, a_2, a_3, \dots, a_N$ which affect the environment, and receiving rewards $r_1, r_2, \dots$ to maximize rewards in the long term

This course focuses most on Supervised Learning (with some unsupervised spots)

The Master Algorithm (Pedro Domingos, 2015)

"The master algorithm is the ultimate learning algorithm. It's an algorithm that can learn anything from data and it's the holy grail of machine learning …"

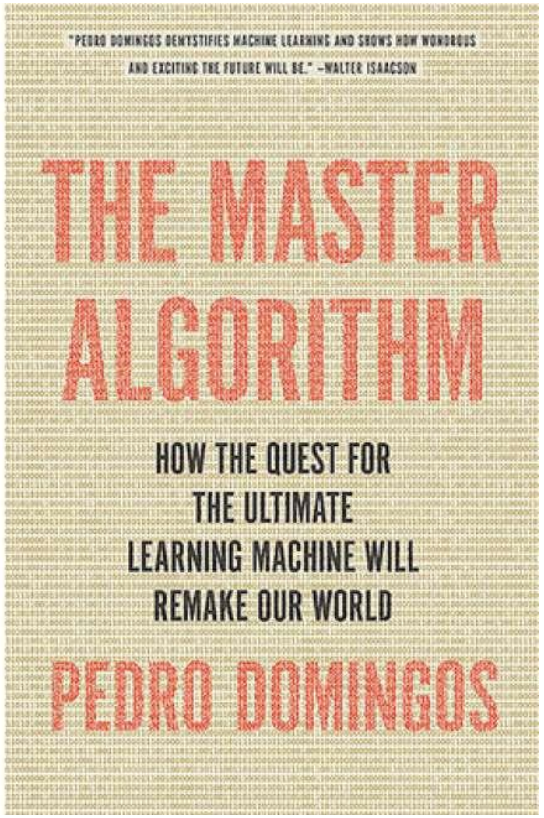

The Master Algorithm (Pedro Domingos, 2015)

Symbolists	Bayesians	Connectionists	Evolutionaries	Analogizers
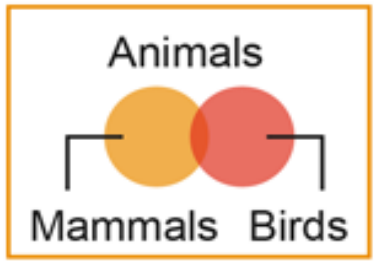		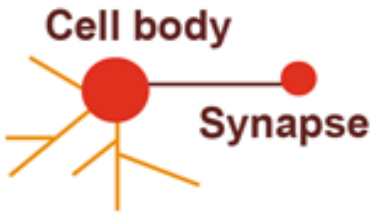		
Use symbols, rules, and logic to represent knowledge and draw logical inference	Assess the likelihood of occurrence for probabilistic inference	Recognize and generalize patterns dynamically with matrices of probabilistic, weighted neurons	Generate variations and then assess the fitness of each for a given purpose	Optimize a function in light of constraints ("going as high as you can while staying on the road")
Favored algorithm Rules and decision trees	**Favored algorithm** Naive Bayes or Markov	**Favored algorithm** Neural networks	**Favored algorithm** Genetic programs	**Favored algorithm** Support vectors

Source: Pedro Domingos, *The Master Algorithm*, 2015

Deep Learning: The Master Algorithm?

MIT Technology Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Introduction | The 10 Technologies | Past Years

Deep Learning
With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.

Temporary Social Media
Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous.

Prenatal DNA Sequencing
Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child?

Additive Manufacturing
Skeptical about 3-D printing? GE, the world's largest manufacturer, is on the verge of using the technology to make jet parts.

Baxter: The Blue-Collar Robot
Rodney Brooks's newest creation is easy to interact with, but the complex innovations behind the robot show just how hard it is to get along with people.

Memory Implants
A maverick neuroscientist believes he has deciphered the code by which the brain forms long-term memories. Next: testing a prosthetic implant for people suffering from long-term memory loss.

Smart Watches
The designers of the Pebble watch realized that a mobile phone is more useful if you don't have to take it out of your pocket.

Ultra-Efficient Solar Power
Doubling the efficiency of a solar cell would completely change the economics of renewable energy. Nanotechnology just might make it possible.

Big Data from Cheap Phones
Collecting and analyzing information from simple cell phones can provide surprising insights into how people move about and behave – and even help us understand the spread of diseases.

Supergrids
A new high-power circuit breaker could finally make highly efficient DC power grids practical.

Enabling Cross-Lingual Conversations in Real Time

Microsoft Research
May 27, 2014 5:58 PM PT

The success of the team's progress to date was on display May 27, in a talk by Microsoft CEO Satya Nadella in Rancho Palos Verdes, Calif., during the Code Conference. During Nadella's conversation
Kara Swisher and Walt Mossberg of
Re/code tech website relating to a new
of personal computing, he asked
deep Pall to join him on stage. Pall, the
rosoft corporate vice president of
pe, demonstrated for the first time
licly the Skype Translator app, with Pall
versing in English with German-

View milestones on the path to Skype Translator

#speech2speech

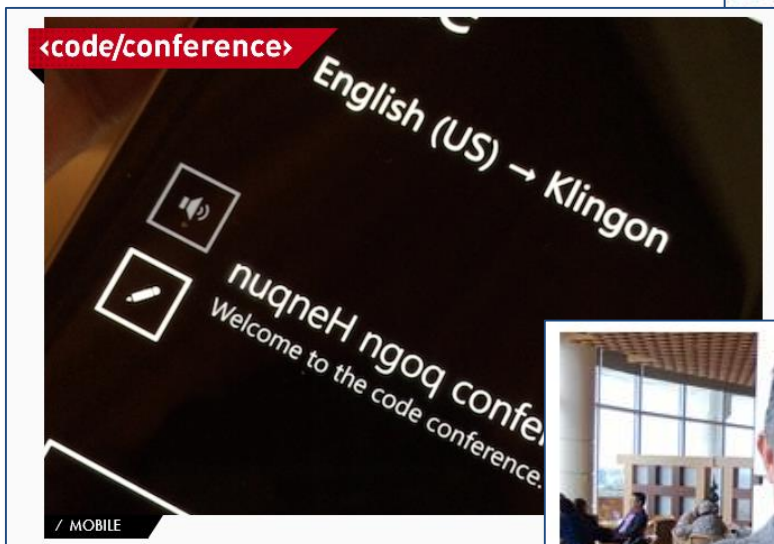

Microsoft's Skype "Star Trek" Language Translator Takes on Tower of Babel

May 27, 2014, 5:48 PM PDT

Remember the universal translator on Star Trek? The gadget th
to aliens?

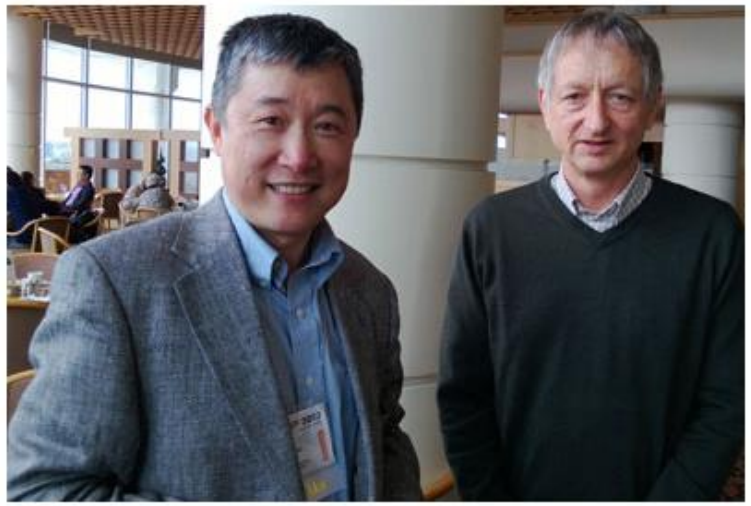

Li Deng (*left*) and Geoff Hinton.

A core development that enables Skype translation came from Redmond researcher Li Deng. He invited Geoff Hinton, a professor at the University of Toronto, to visit Redmond in 2009 to work on new neural-network learning methods, based on a couple of seminal papers from Hinton and his collaborators in 2006 that had brought new

The path to the Skype Translator gained momentum with an encounter in the autumn of 2010. Seide and colleague Kit Thambiratnam had developed a system they called The Translating! Telephone for live speech-to-text and speech-to-speech translation of phone calls.

IMAGENET

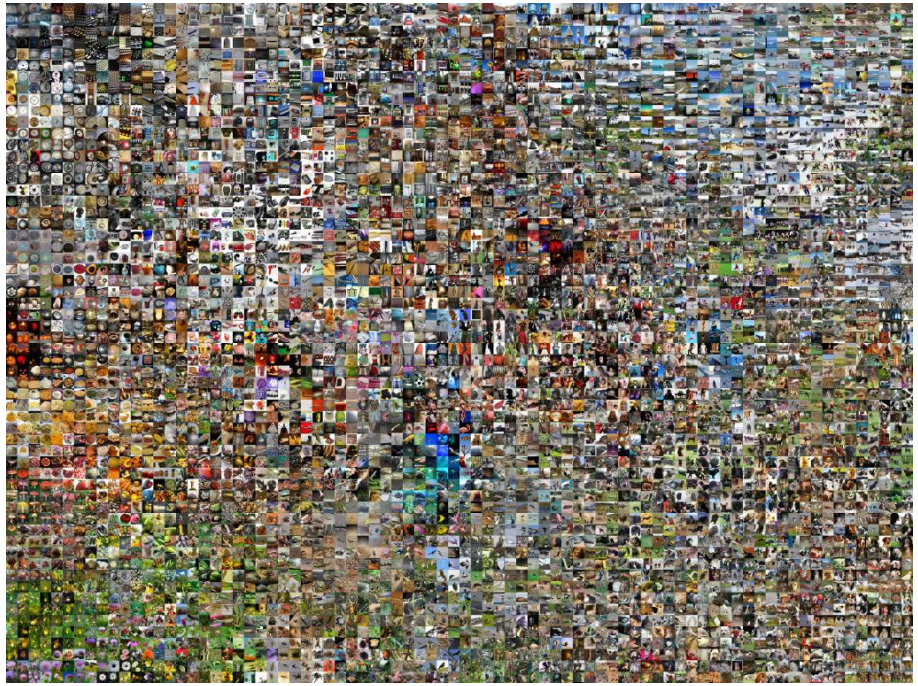

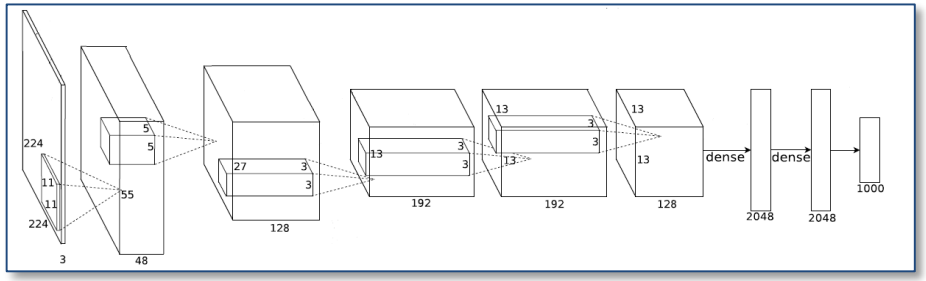
224
11
11
224
3
5
5
55
48
27
3
3
128
13
3
3
13
192
13
3
3
13
192
13
13
128
dense
2048
dense
2048
1000

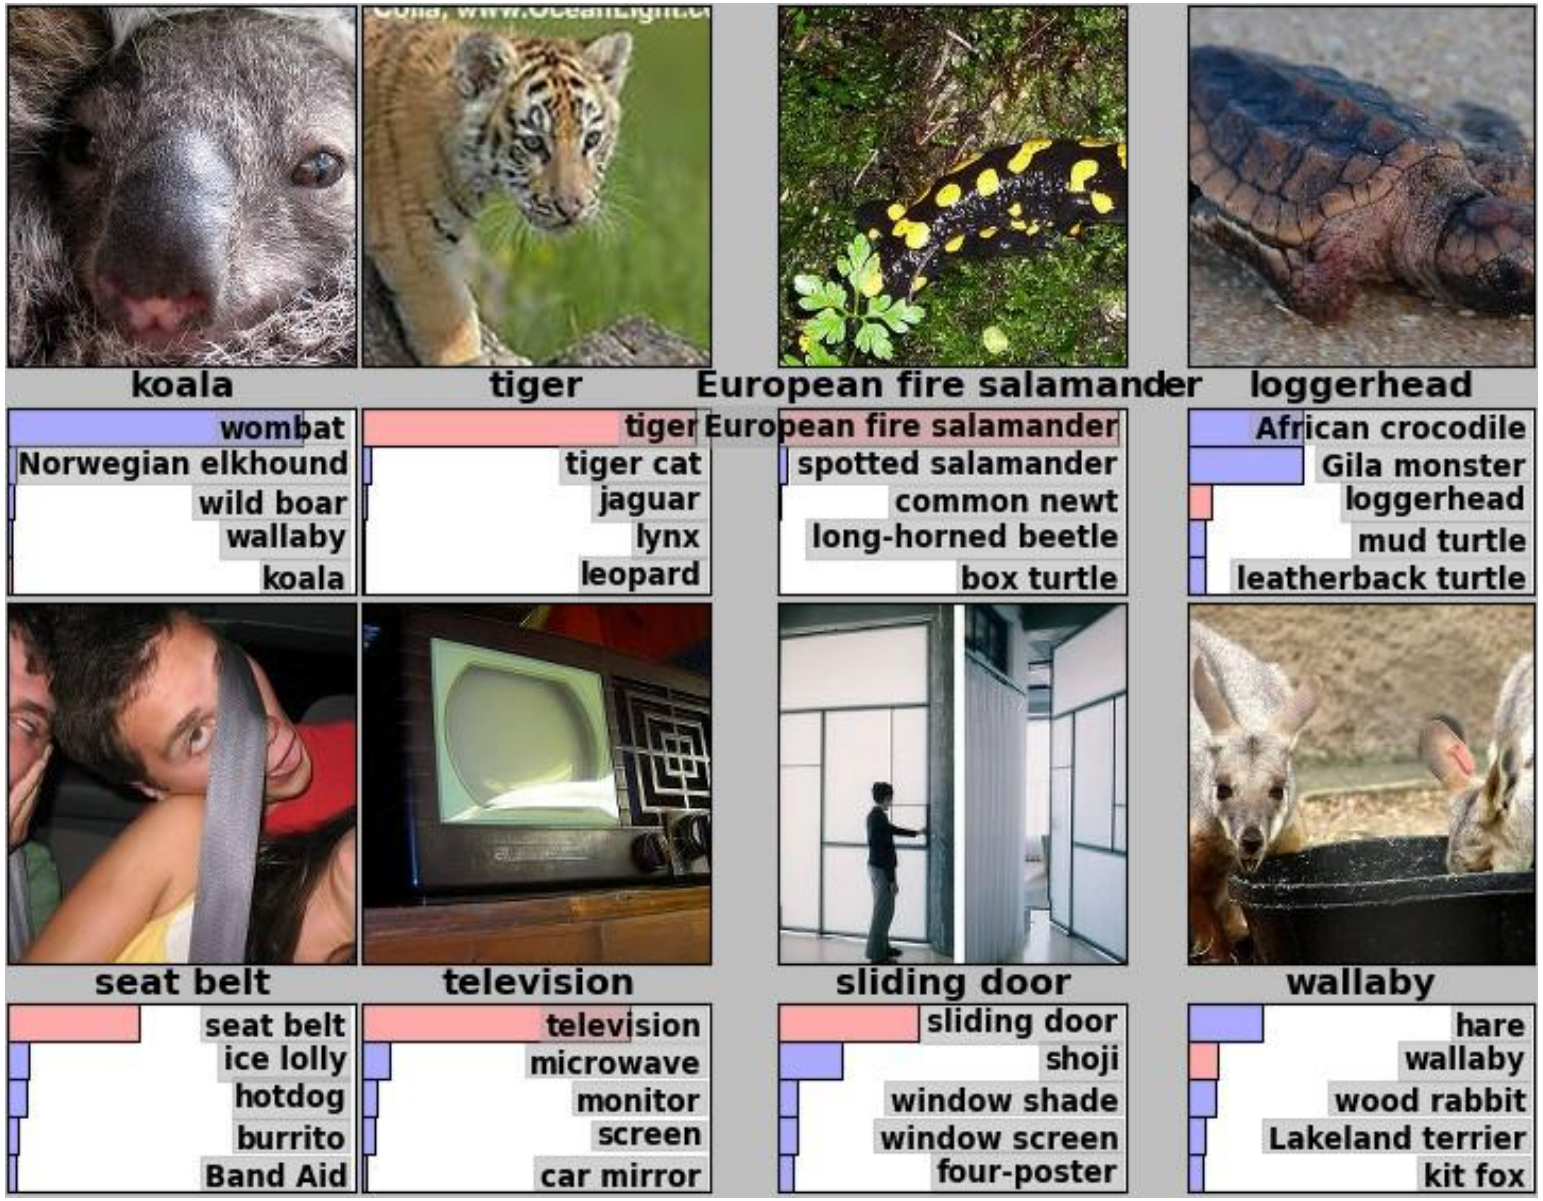
koala
wombat
Norwegian elkhound
wild boar
wallaby
koala
tiger
tiger
tiger cat
jaguar
lynx
leopard
European fire salamander
European fire salamander
spotted salamander
common newt
long-horned beetle
box turtle
loggerhead
African crocodile
Gila monster
loggerhead
mud turtle
leatherback turtle
seat belt
seat belt
ice lolly
hotdog
burrito
Band Aid
television
television
microwave
monitor
screen
car mirror
sliding door
sliding door
shoji
window shade
window screen
four-poster
wallaby
hare
wallaby
wood rabbit
Lakeland terrier
kit fox

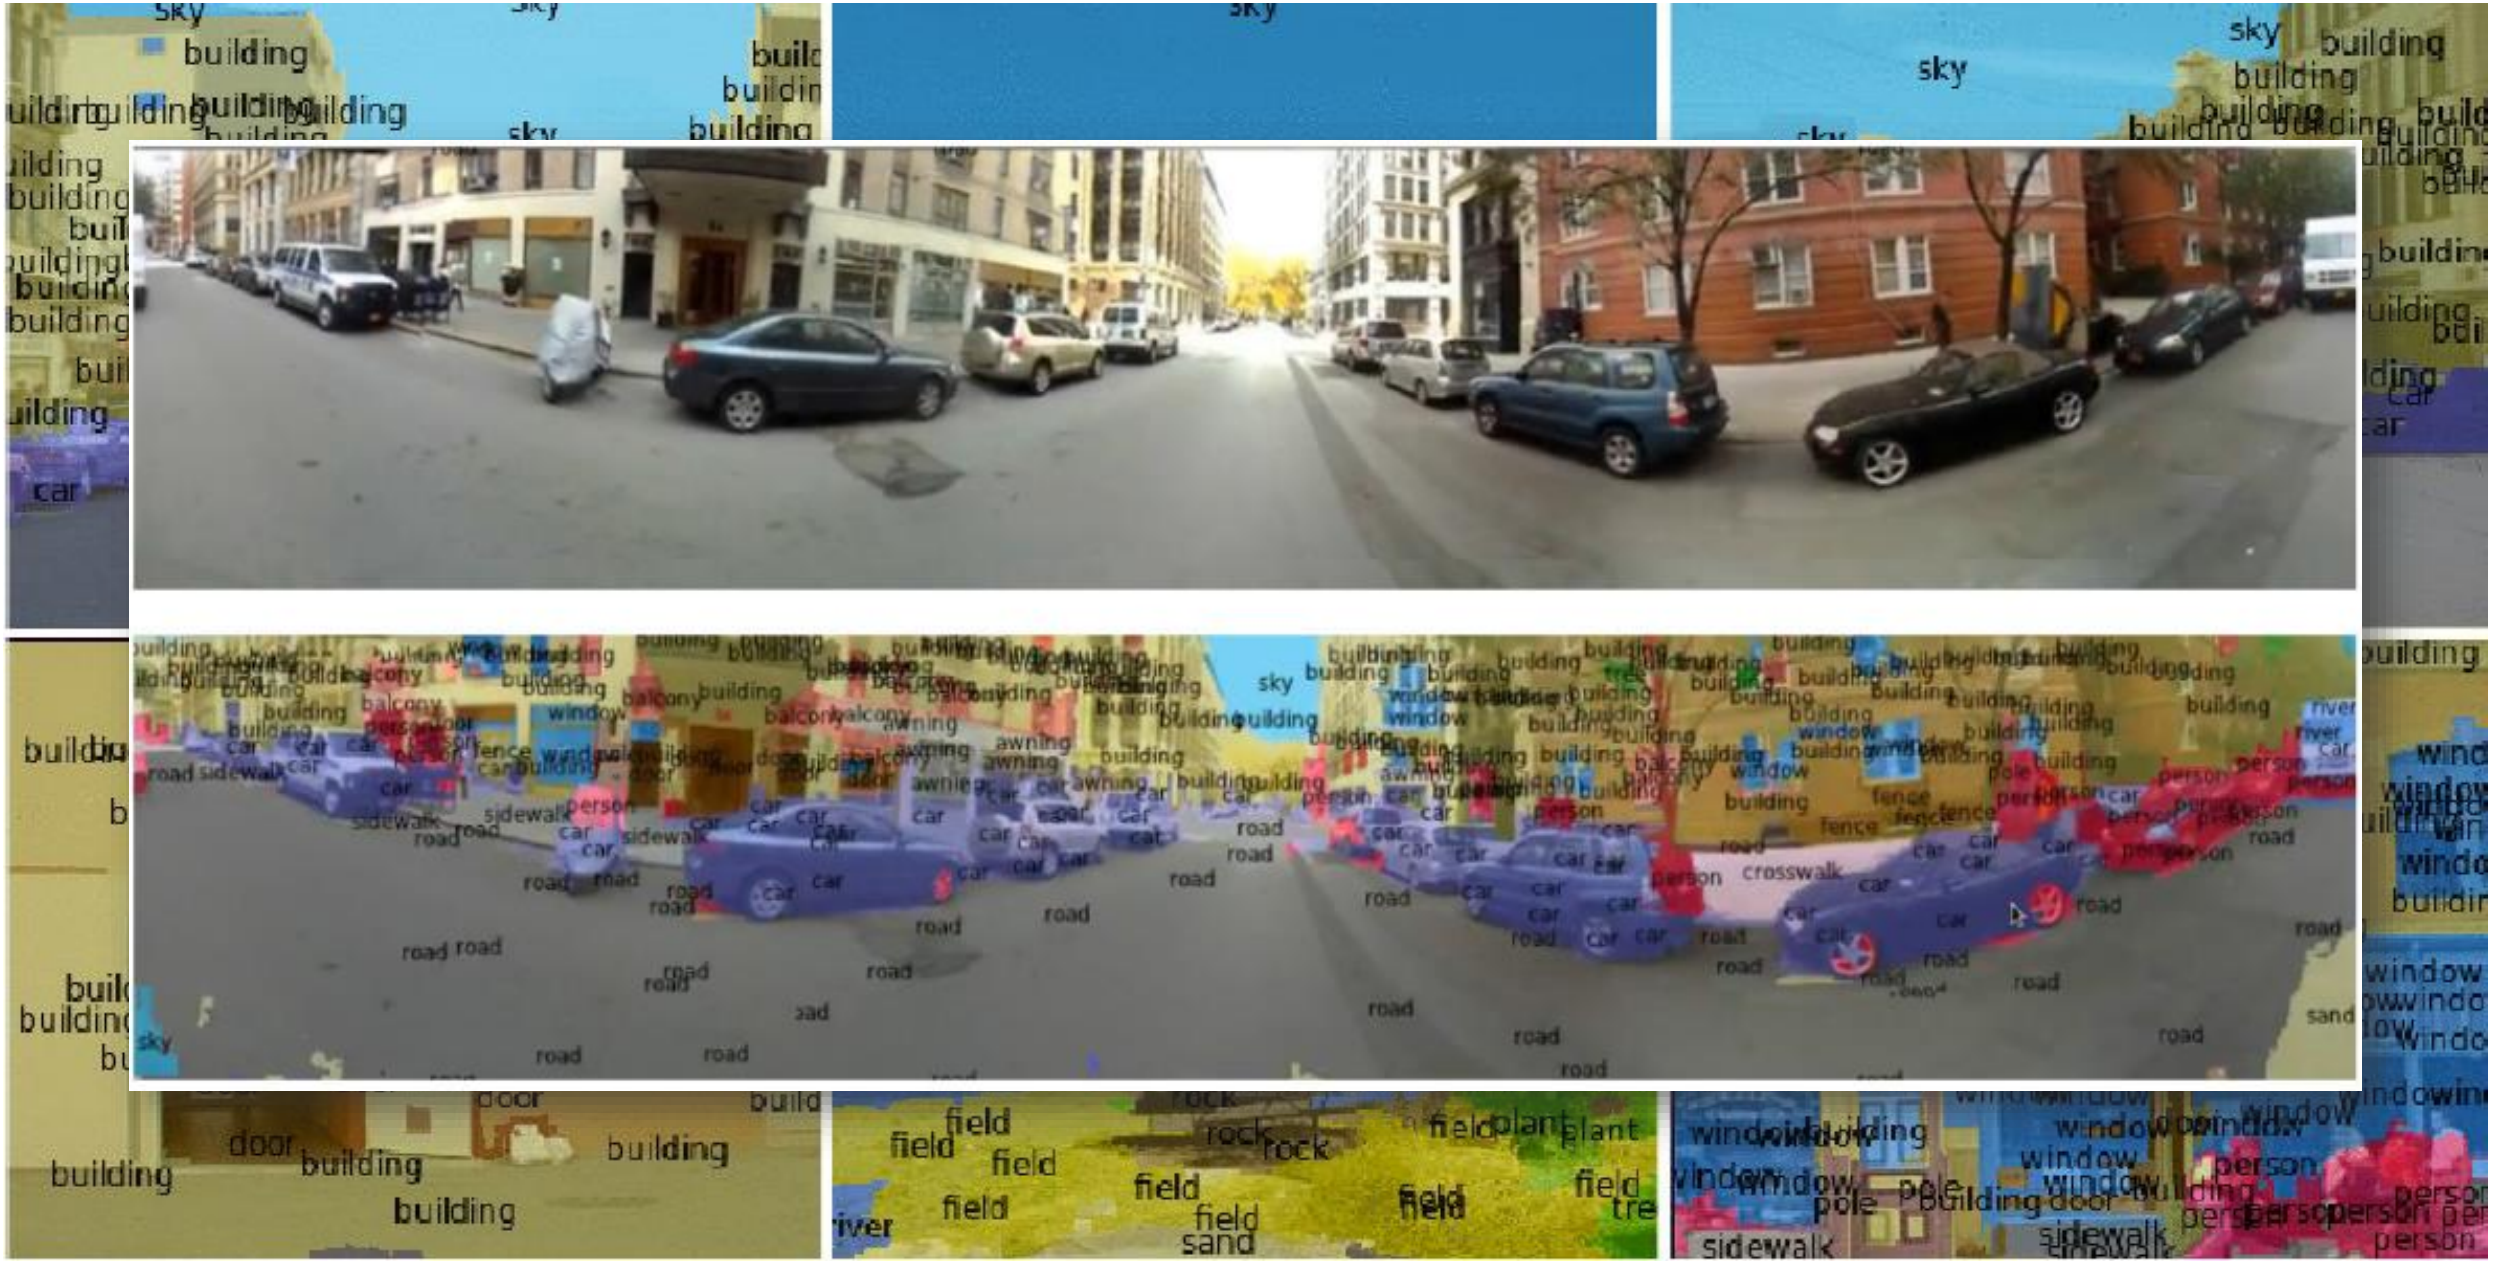


umbrella.99
person1.00
person1.00
person1.00
car.99
car.93

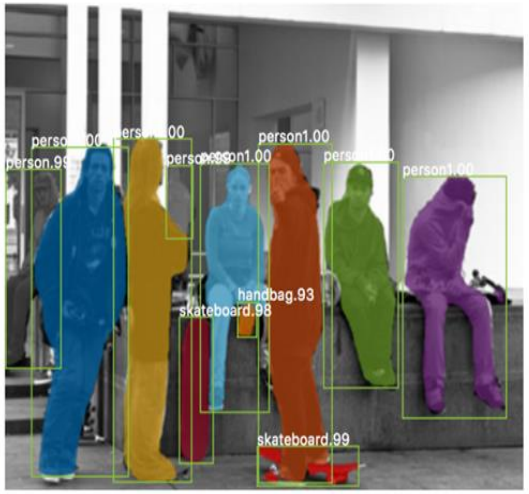
person1.00
handbag.93
skateboard.98
skateboard.99

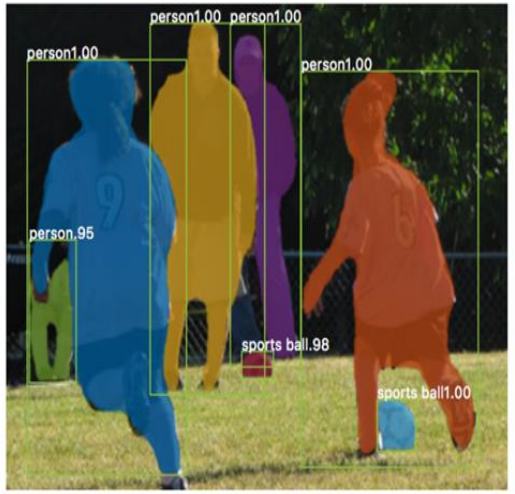
person1.00
person1.00
person1.00
person1.00
person.95
sports ball.98
sports ball1.00

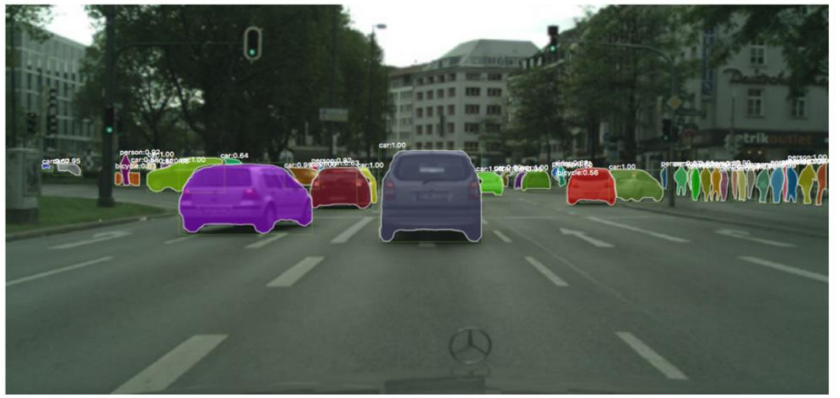


imax PHOTOGRAPHY

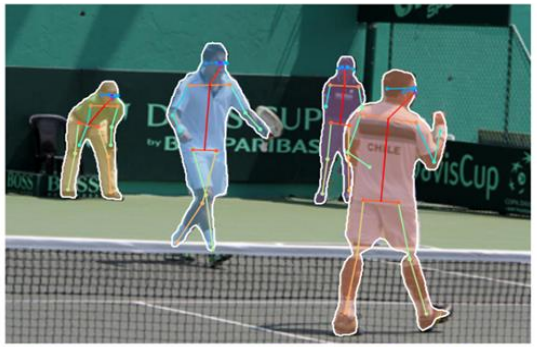

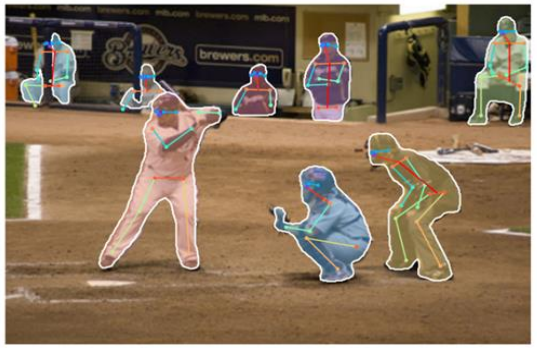
brewers.com

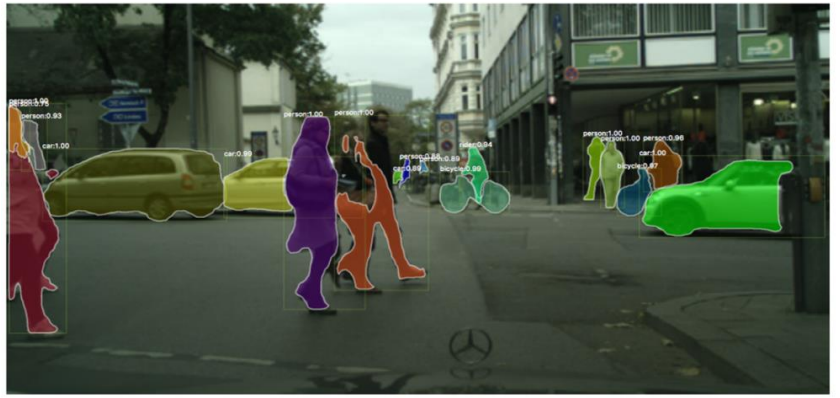

https://github.com/luanfujun/deep-photo-styletransfer

https://github.com/jcjohnson/neural-style

https://github.com/jcjohnson/fast-neural-style

https://ml4a.github.io/ml4a/style_transfer/

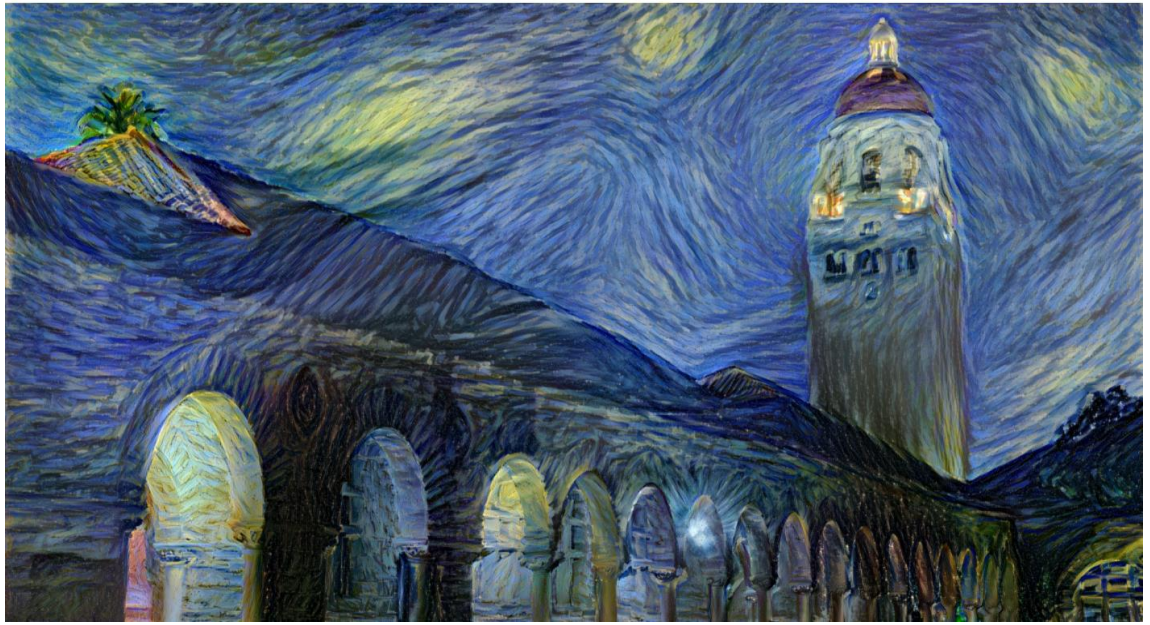

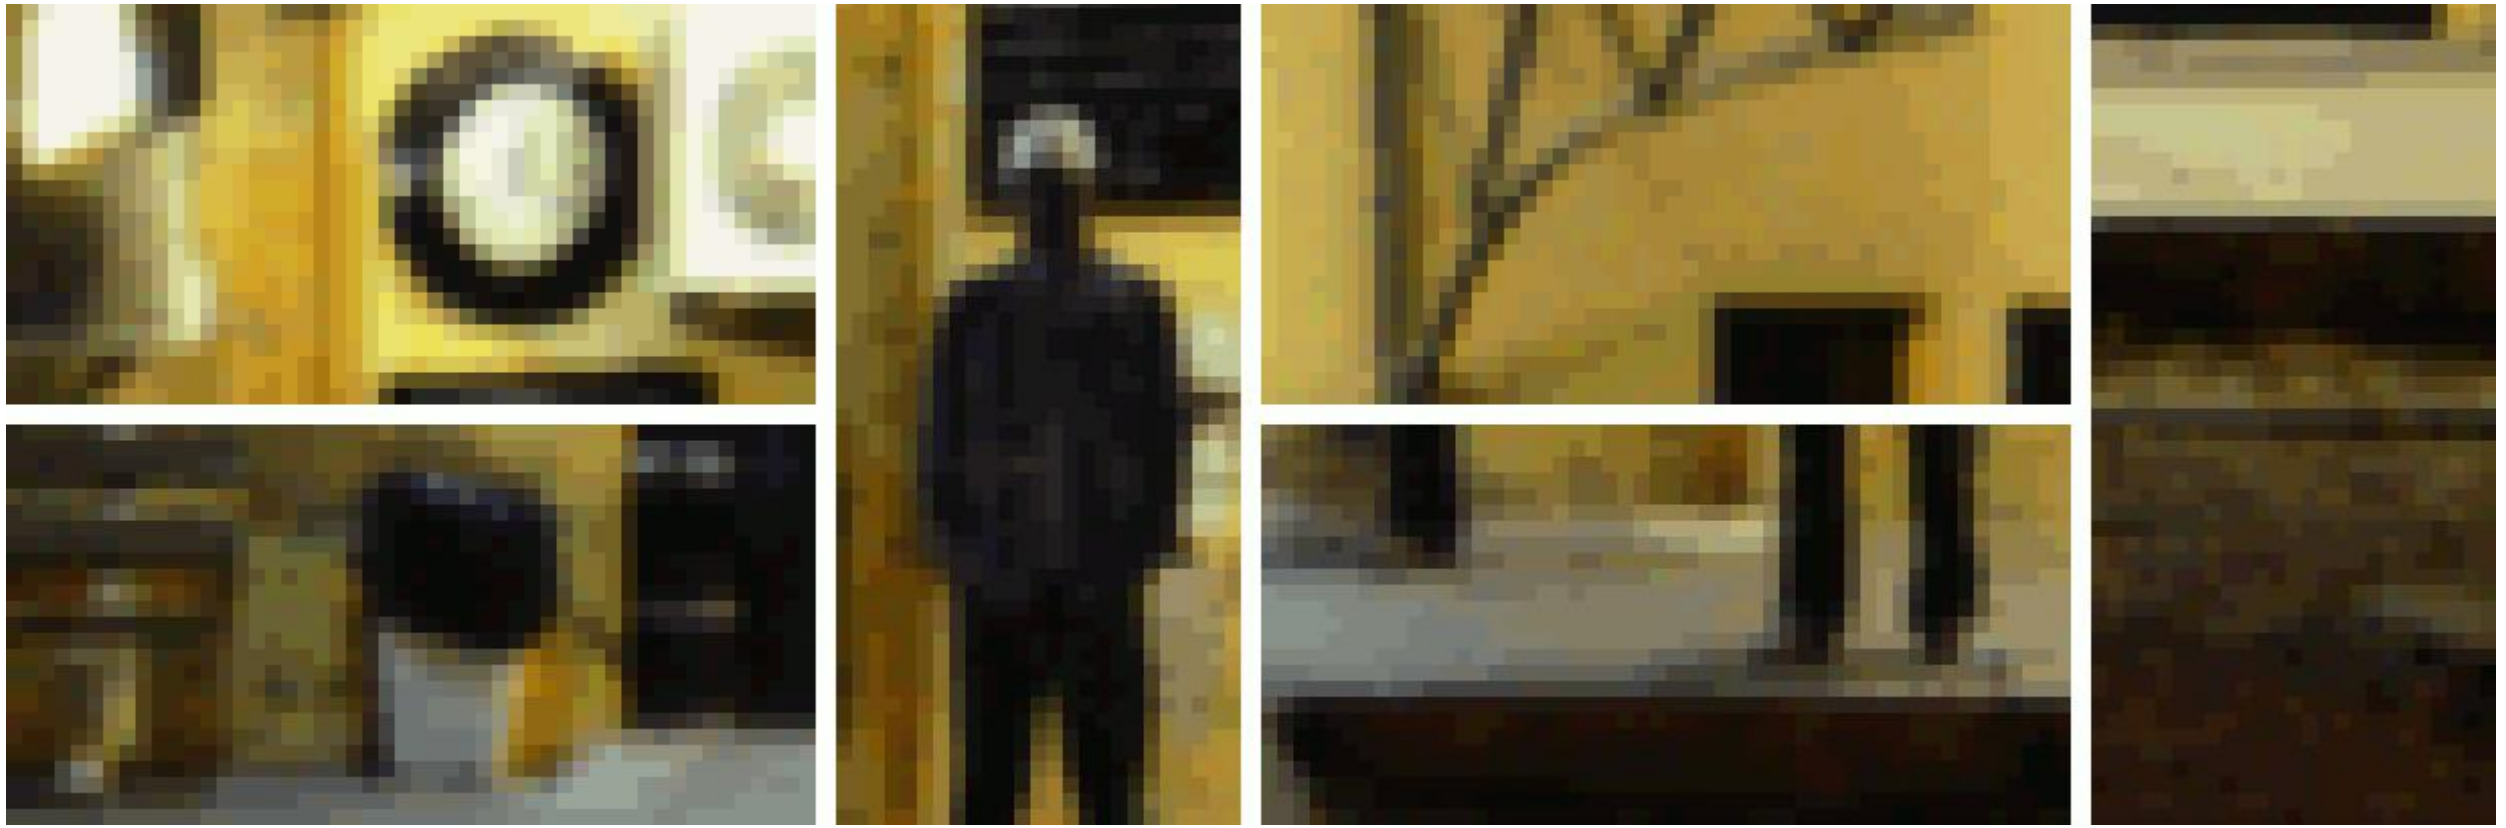

https://github.com/alexjc/neural-enhance

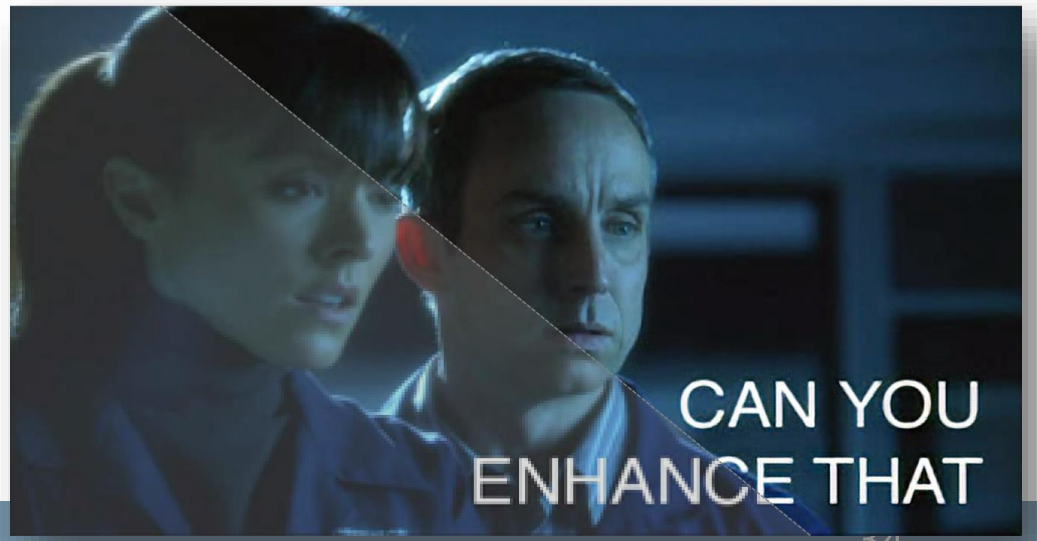

Text description	This flower has petals that are white and has pink shading	This flower has a lot of small purple petals in a dome-like configuration	This flower has long thin yellow petals and a lot of yellow anthers in the center	This flower is pink, white, and yellow in color, and has petals that are striped	This flower is white and yellow in color, with petals that are wavy and smooth	This flower has upturned petals which are thin and orange with rounded edges	This flower has petals that are dark pink with white edges and pink stamen
256x256 StackGAN							

Text description	This bird is red and brown in color, with a stubby beak	The bird is short and stubby with yellow on its body	A bird with a medium orange bill white body gray wings and webbed feet	This small black bird has a short, slightly curved bill and long legs	A small bird with varying shades of brown with white under the eyes	A small yellow bird with a black crown and a short black pointed beak	This small bird has a white breast, light grey head, and black wings and tail
256x256 StackGAN							

'Go is implicit. It's all pattern matching. But that's what deep learning does very well.'

—DEMIS HASSABIS, DEEPMIND

The win is more than a novelty. Online services like Google, Facebook, and Microsoft, already use deep learning to identify images, recognize spoken words, and understand natural

with a technology called reinfor
methods, point the way to a futu
can learn to perform physical ta
environment. "It's a natural fit f

It's incredibly difficult to build a machine that duplicates the kind of intuition that makes the top human players so good at

In the mid-'90s, a computer program called Chinook beat the world's top player at the game of checkers. A few years later, IBM's Deep Blue supercomputer shocked the chess world when it wiped the proverbial floor with world champion Gary Kasparov. And more

BM machine, Watson, topped the best
dy!, the venerable TV trivia game.
o mastered Othello, Scrabble,
poker. But in the wake of Crazy Stone's
Coulom predicted that another ten years
a machine could beat a grandmaster
rt.

IN A HUGE BREAKTHROUGH, GOOGLE'S AI BEATS A TOP PLAYER AT THE GAME OF GO

WIRED

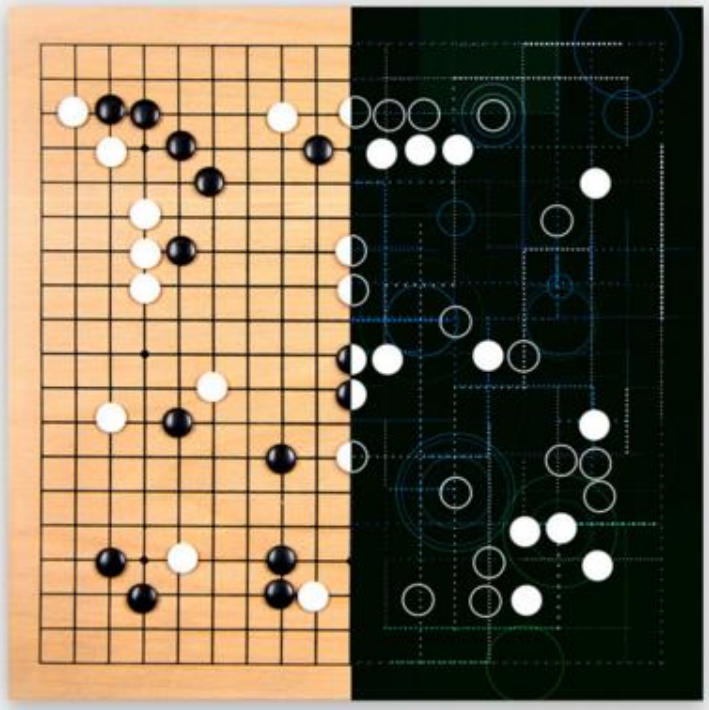

The Master Algorithm (Pedro Domingos, 2015)

"The master algorithm is the ultimate learning algorithm. It's an algorithm that can learn anything from data and it's the holy grail of machine learning …"

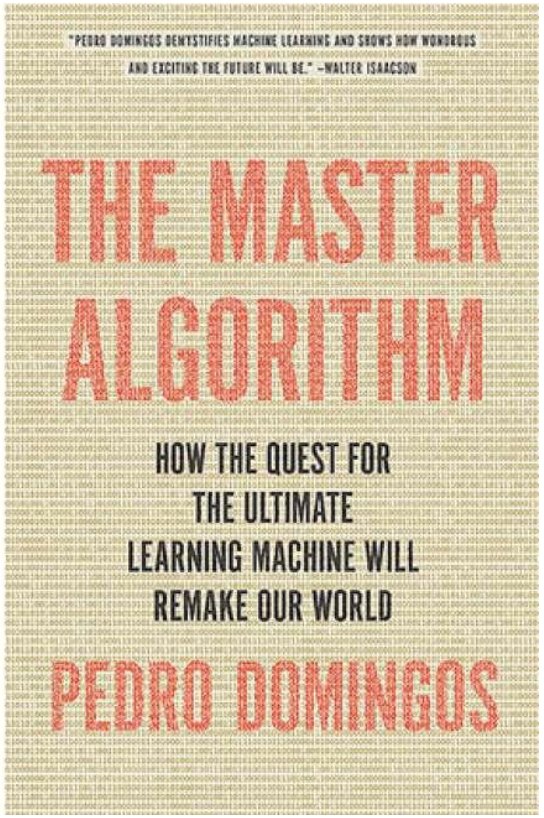

Deep Learning: The Master Algorithm?

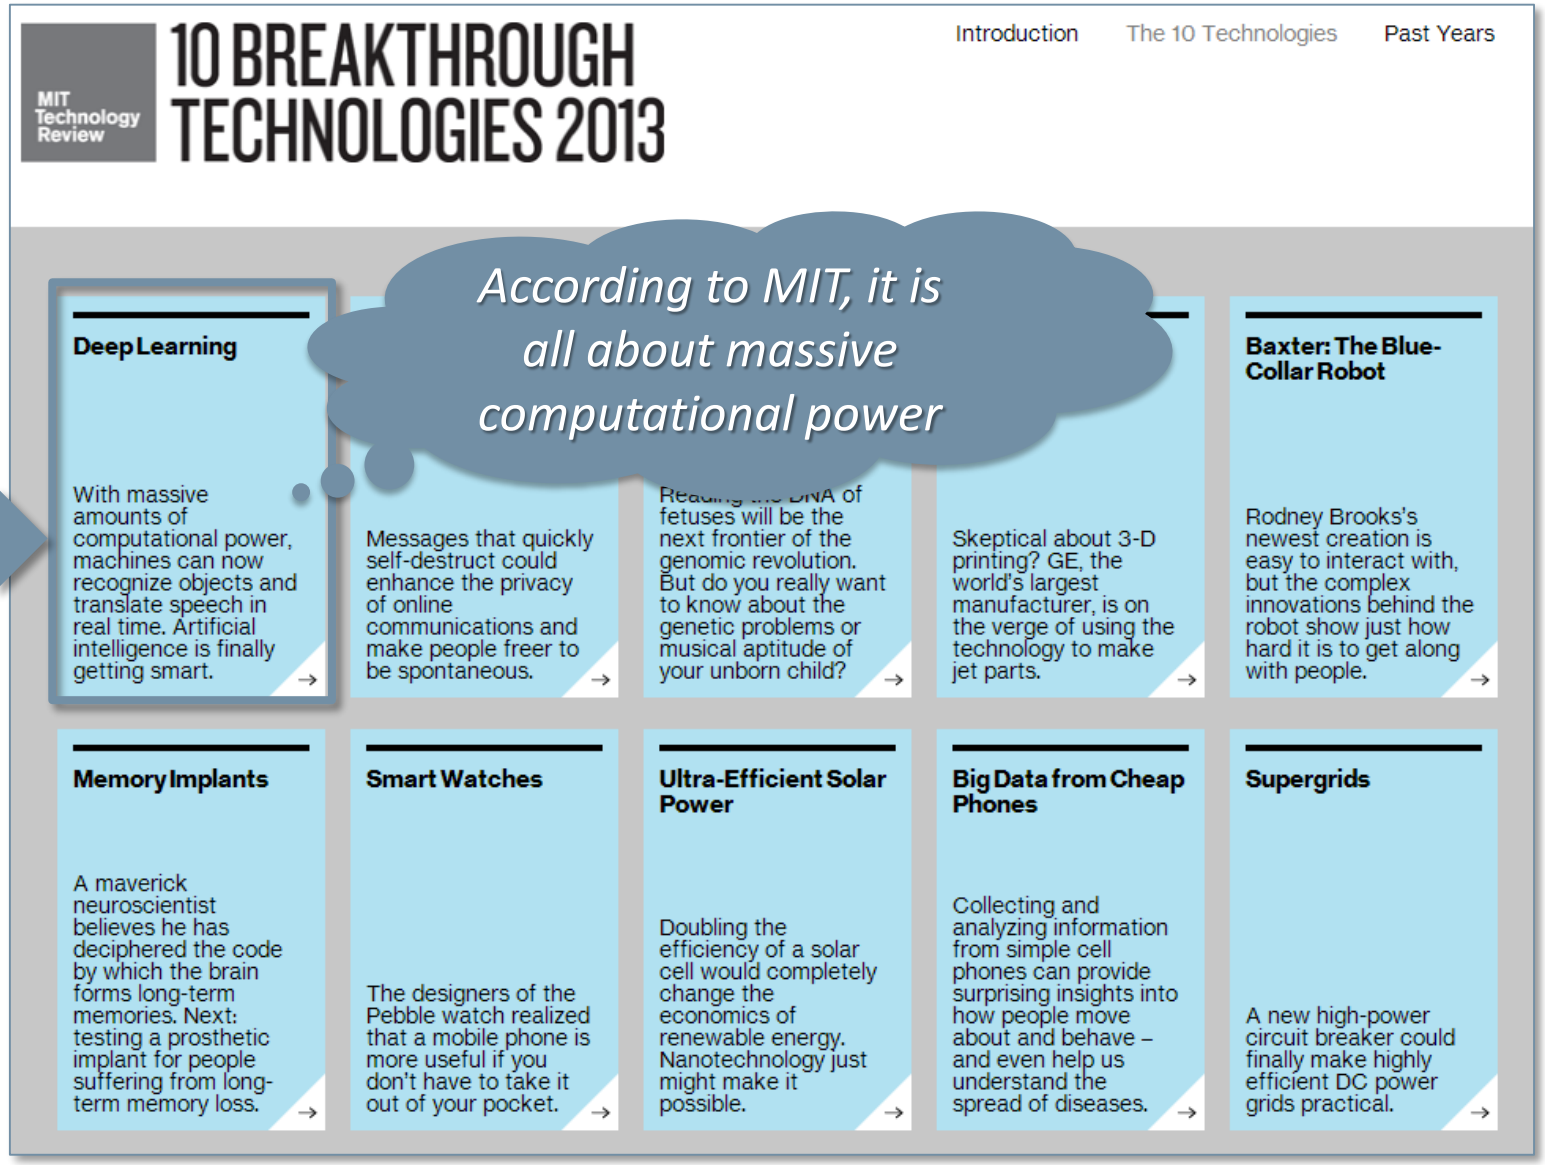

Deep Learning: The Master Algorithm?

MIT Technology Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Introduction | The 10 Technologies | Past Years

Deep Learning
With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart.

Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous.

Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child?

Skeptical about 3-D printing? GE, the world's largest manufacturer, is on the verge of using the technology to make jet parts.

Baxter: The Blue-Collar Robot
Rodney Brooks's newest creation is easy to interact with, but the complex innovations behind the robot show just how hard it is to get along with people.

Memory Implants
A maverick neuroscientist believes he has

Smart Watches

Ultra-Efficient Solar Power
Doubling the efficiency of a solar cell would completely change the economics of renewable energy. Nanotechnology just might make it possible.

Big Data from Cheap Phones
Collecting and analyzing information from simple cell phones can provide surprising insights into how people move about and behave – and even help us understand the spread of diseases.

Supergrids
A new high-power circuit breaker could finally make highly efficient DC power grids practical.

According to MIT, it is all about massive computational power

The Economist got it right! It is all about (Big) Data

Recall about Supervised Learning

Cars

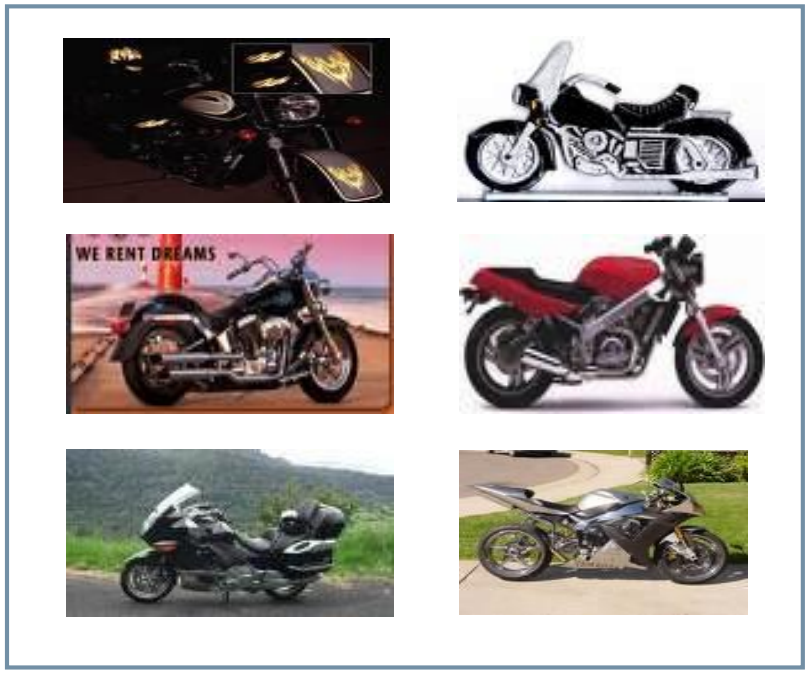

Motorcycles

Hand-crafted Features

Learned Classifier

CAR

Recall about Supervised Learning

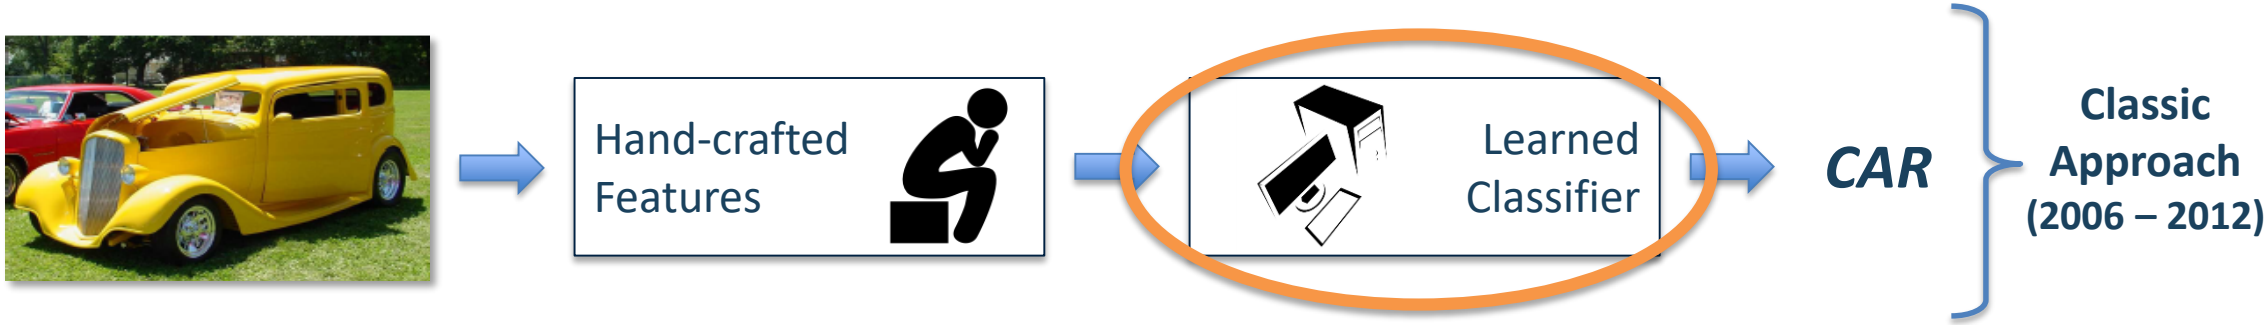

Features are based on domain knowledge or heuristics:

- Words in a Dictionary for text classification
- MFCC for Speech Recognition
- SIFT, HoG, BRIEF in Visual Tasks

How Machine Learning can help with this?

However …

- They need to be carefully designed depending on the task
- They are fixed and sometimes they do not generalize between datasets

Beyond Supervised Learning

Lots of labeled examples required!

Cars

Motorcycles

Hand-crafted Features

Learned Classifier

CAR

Semi-supervised learning

Unlabeled images (all cars/motorcycles)

Hand-crafted Features

Unsupervised Learnng

$\phi_1, \phi_2, \ldots, \phi_k$

Car

Motorcycle

Few labeled images

$\phi_1, \phi_2, \ldots, \phi_k$

Learned Classifier

CAR / BYKE

Semi-supervised learning

Unlabeled images (all cars/motorcycles)

Hand-crafted Features

Unsupervised Learnng

$\phi_1, \phi_2, \ldots, \phi_k$

Hand-crafted Features

Learned Feature Projection

Learned Classifier

CAR

Modern Pattern Recogniton

Speech recognition (early 90's – 2011)

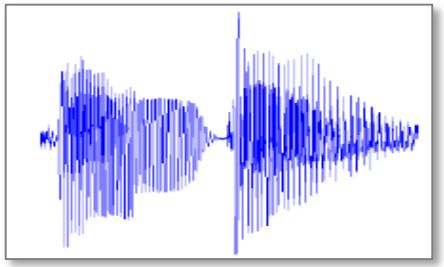

Low level features

MFCC

Fixed Features

Mid level features

Mixture of Gaussian

Unsupervised Learning

Learned Classifier

Supervised Learning

Horn

Object recognition (2006 – 2012)

Low level features

SIFT/HoG

Fixed Features

Mid level features

K-means Sparse Coding

Unsupervised Learning

Pooling

Learned Classifier

Supervised Learning

Car

Transfer Learning

Unlabeled images (all cars/motorcycles)

Hand-crafted Features

Unsupervised Learnng

$\phi_1, \phi_2, \ldots, \phi_k$

Hand-crafted Features

Learned Feature Projection

Learned Classifier

CAR

Transfer Learning

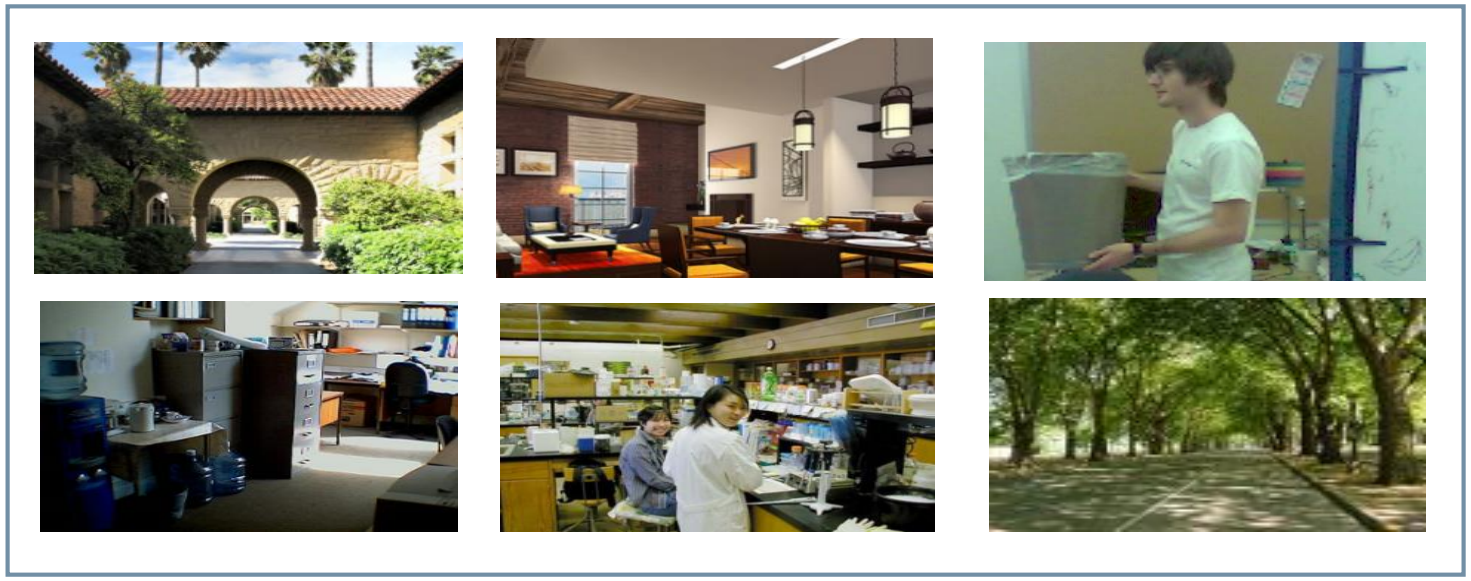

Unlabeled images (random images from the web)

$\phi_1, \phi_2, \ldots, \phi_k$

Transfer Learning

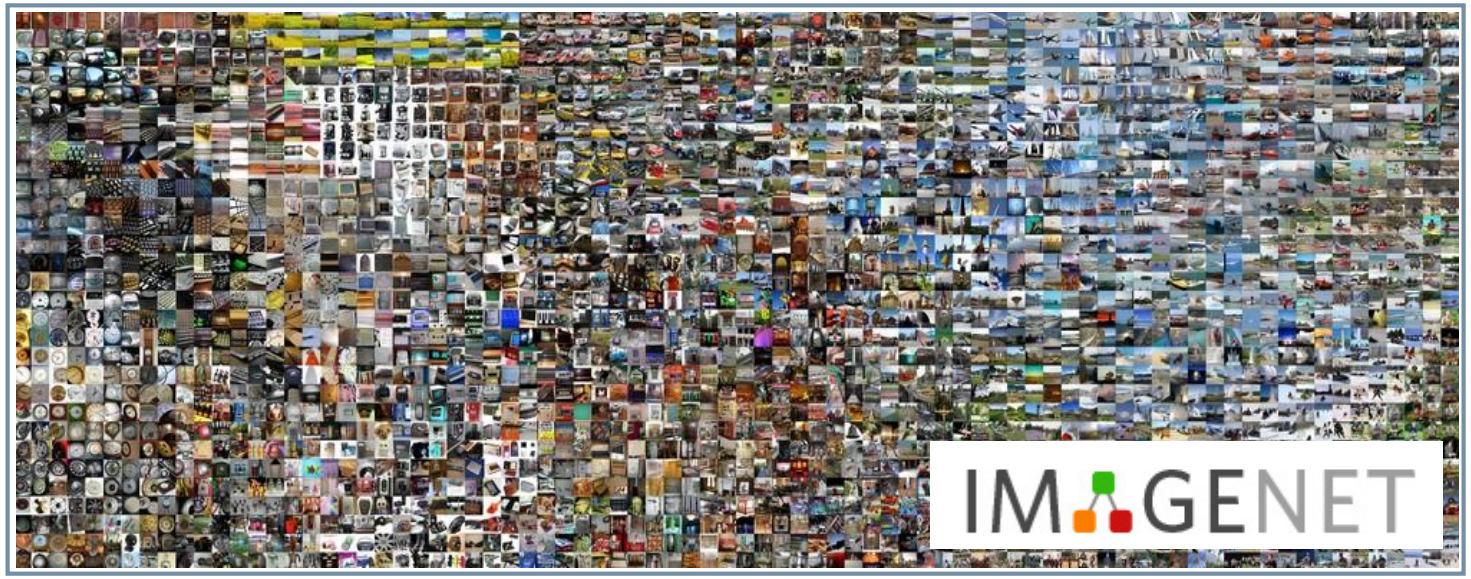

Unlabeled images (random images from the web)

Hand-crafted Features

Unsupervised Learnng

$\phi_1, \phi_2, \ldots, \phi_k$

Hand-crafted Features

Learned Feature Projection

Learned Classifier

CAR

It's all about features …

What if we do not get these right?

Hand-crafted Features → Learned Feature Projection → Learned Classifier → ***CAR***

Classic approach (2006 – 2012)

It's all about features …

Learned Features → Learned Features → Learned Features → Learned Classifier → **CAR**

Deep Learning (2012 – …)

Learn from data!

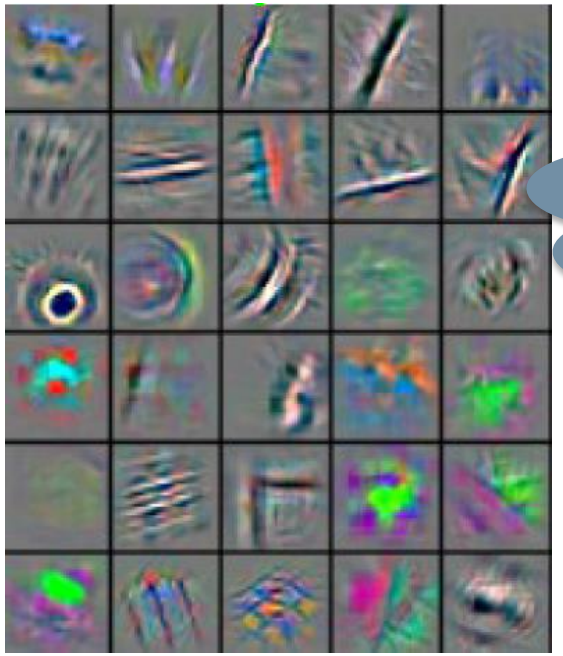

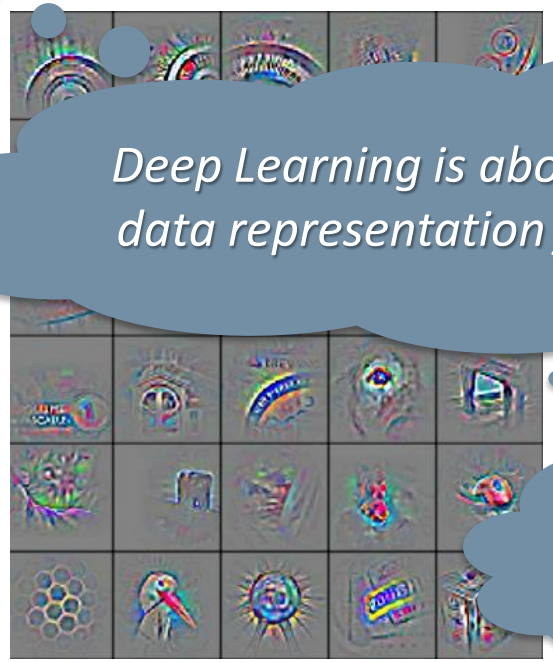

Deep Learning is about learning data representation from data!

But which data?


IN 60 SECONDS...
1 NEW DEFINITION IS ADDED ON urban
1,600+ READS ON Scribd.
13,000+ HOURS MUSIC STREAMING ON PANDORA
12,000+ NEW ADS POSTED ON craigslist
370,000+ MINUTES VOICE CALLS ON skype
98,000+ TWEETS
20,000+ NEW POSTS ON tumblr.
THE LARGEST SOCIAL READING PUBLISHING COMPANY!!
New Craigslist Ads
320+ NEW twitter ACCOUNTS
13,000+ iPhone APPLICATIONS DOWNLOADED
100+ NEW Linked in ACCOUNTS
1 associatedcontent NEW ARTICLE IS PUBLISHED
THE WORLD'S LARGEST COMMUNITY CREATED CONTENT!!
QUESTIONS ASKED ON THE INTERNET...
100+ Answers.com
40+ YAHOO! ANSWERS
600+ NEW VIDEOS
6,600+ NEW PICTURES ARE UPLOADED ON flickr
50+ WordPress DOWNLOADS
25+ HOURS TOTAL DURATION
70+ DOMAINS REGISTERED
695,000+ facebook STATUS UPDATES
125+ PLUGIN DOWNLOADS
60+ NEW BLOGS
1,500+ BLOG POSTS
168 MILLION EMAILS ARE SENT
694,445 SEARCH QUERIES
1,700+ Firefox DOWNLOADS
79,364 WALL POSTS
510,040 COMMENTS
Google
Google Search

It's all about features …

Learn from data!

Learned Features → Learned Features → Learned Features → Learned Classifier → ***CAR***

Deep Learning (2012 – …)

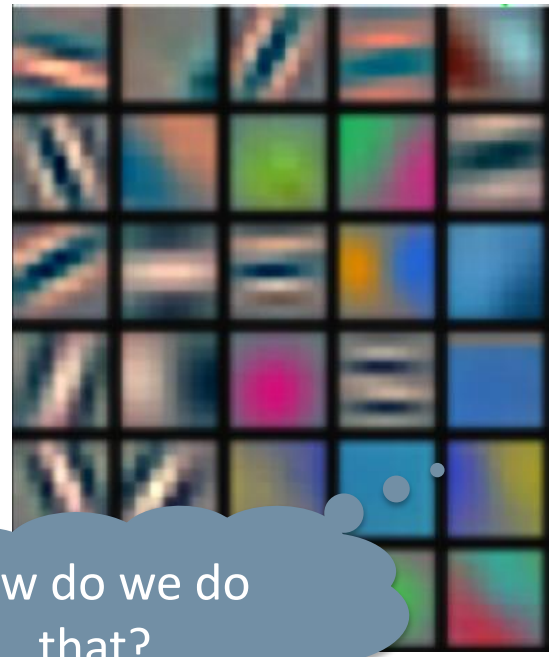

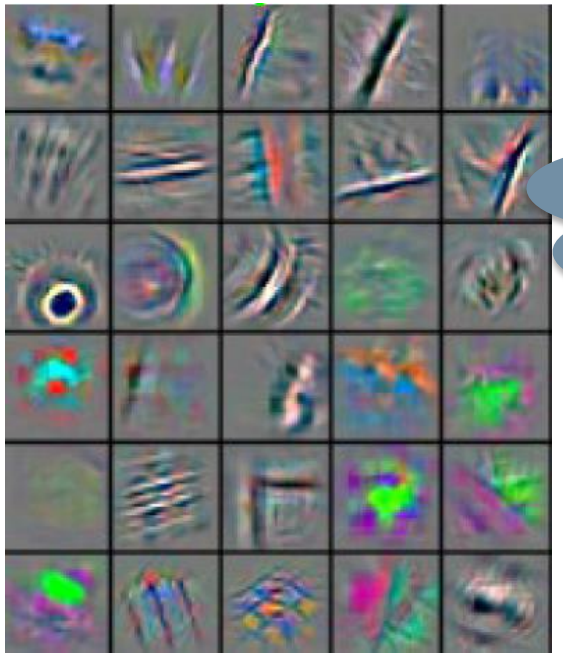

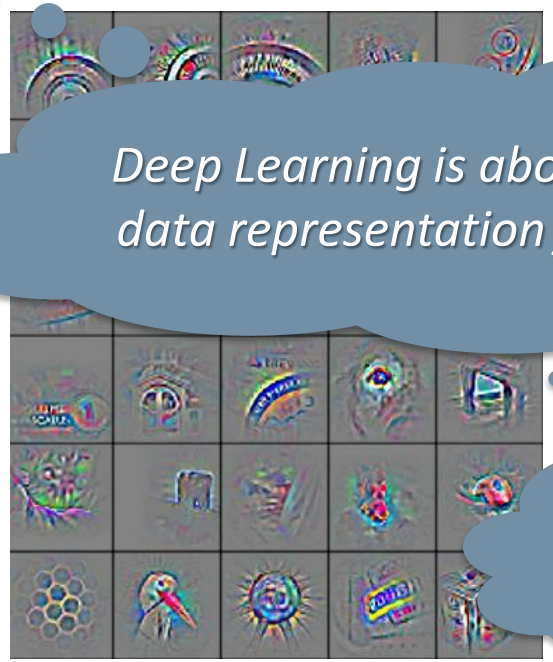

Deep Learning is about learning data representation from data!

How do we do that?

But which data?

Some people call it the «Fourth Paradigm»

"Scientific breakthroughs powered by advanced computing capabilities that help researcher manipulate and explore massive datasets"

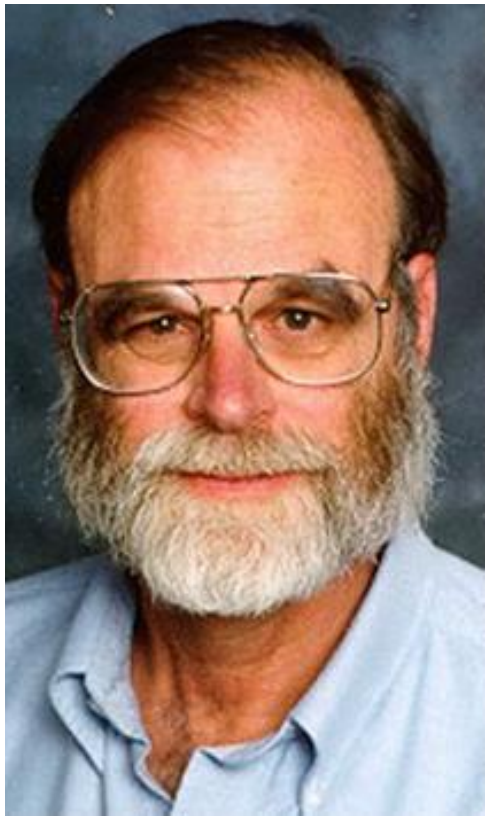

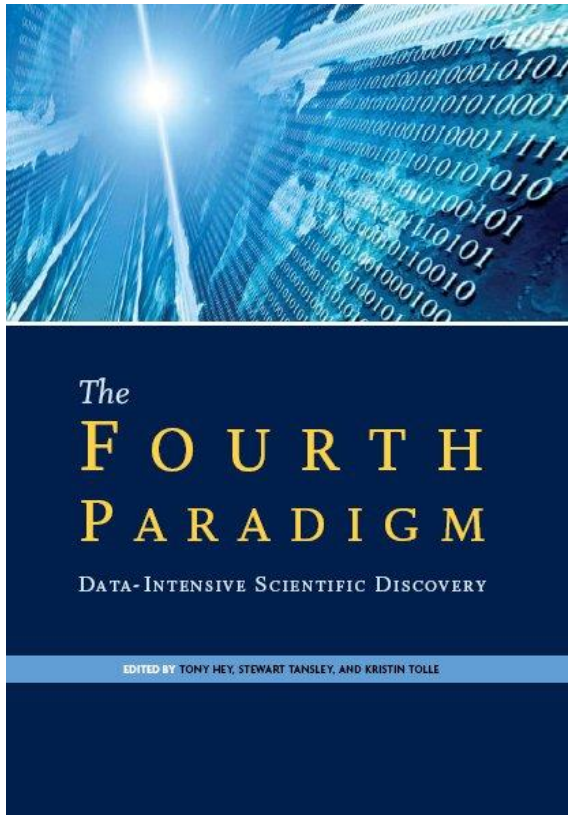

The Fourth Paradigm explain

Deep Learning, i.e., representation learning from data, is the fourth paradigm for AI!

Empirical science

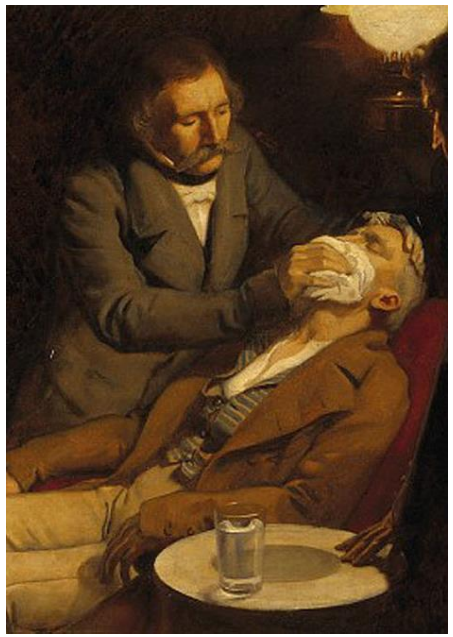

Morton – 1846
(Anesthesia)

Theoretical science

Pasteur – 1870
(Germ Theory)

Computational science

Bradford Hill – 1920
(Randomised Trials)

Data-intensive science

Next Generation Sequencing – 2000
(Towards personalized medicine)

Representation Learning in Context

Learning the representation is a challenging problem for Machine Learning, Computer Vision, Artificial Intelligence, Neuroscience, …

Cognitive perspective

- How can a perceptual system build itself looking at the external world?
- How much prior structure is necessary?

Neuroscience perspective

- Does the cortex «run» a single, general learning algorithm?

Representation Learning in Context

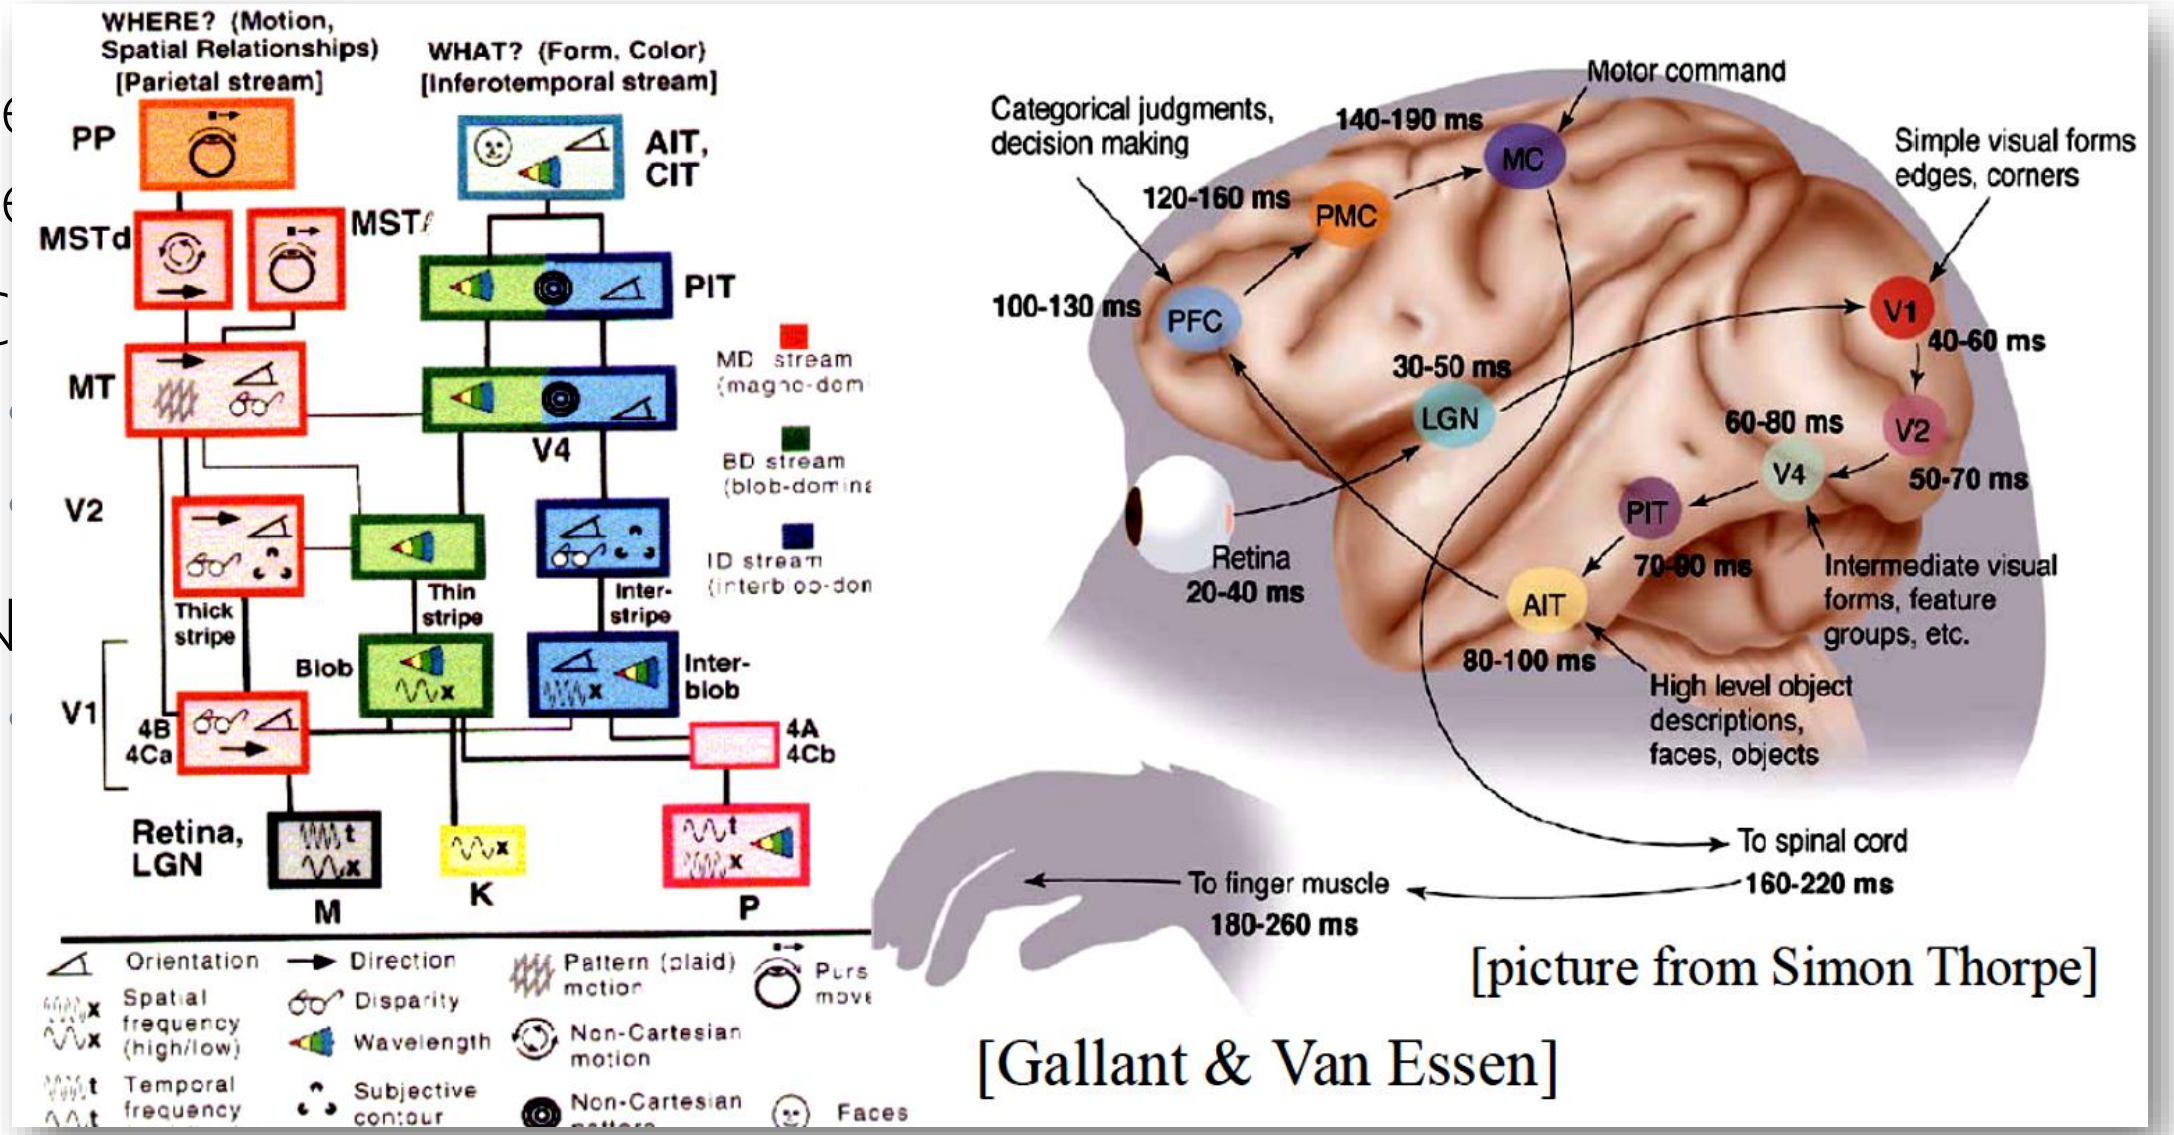

Representation Learning in Context

Learning the representation is a challenging problem for Machi
Learning, Computer Vision, Artificial Intelligence, Neuroscience,

Cognitive perspective

- How can a perceptual system build itself looking at the external world?
- How much prior structure is necessary?

Neuroscience perspective

- Does the cortex «run» a single, general learning algorithm?

Artificial Intelligence Perspective

- What is the fundamental
- How do we build a
- What is the architectu

Deep learning addresses the problem of learning hierarchical representations with a single algorithm.

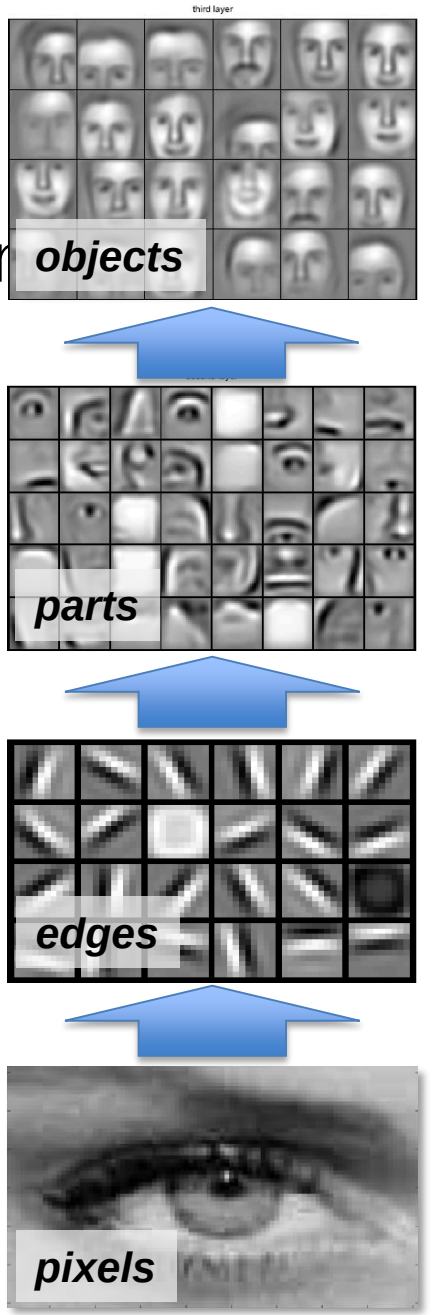

Trainable Features Hierarchy

Deep learning assumes it is possible to «learn» a hierarchy of descriptors with increasing abstraction, i.e., layers are trainable feature transforms

In image recognition

- Pixel → edge → texton → motif → part → object

In text analysis

- Character → word → word group → clause → sentence → story

In speech recognition

- Sample → spectral band → sound → phone → phoneme → word

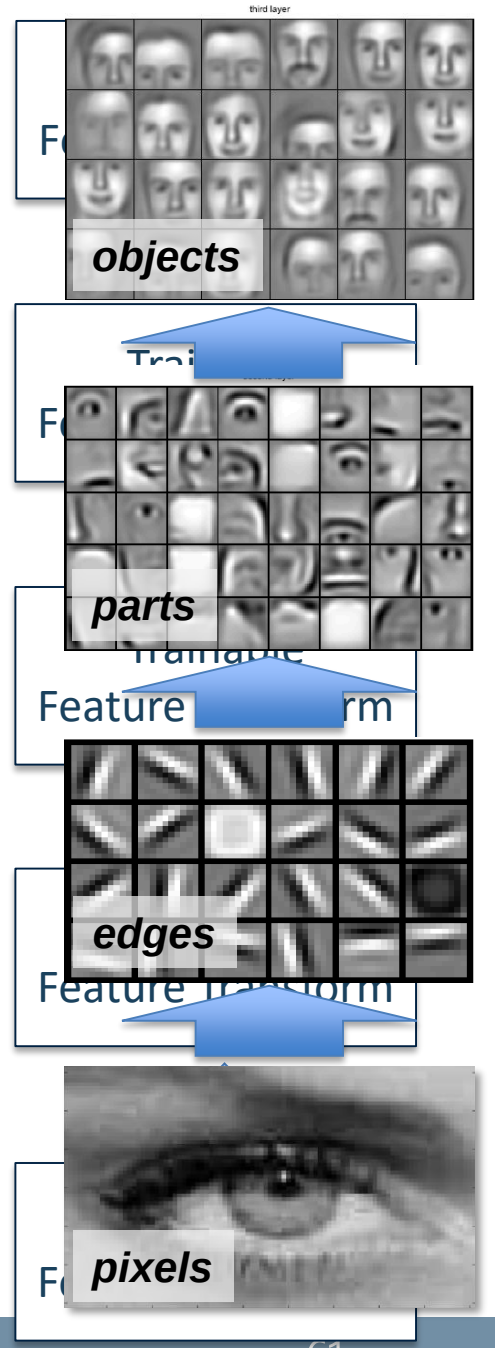

Architectures and Algoriths

Depending on the direction of the information flow we can have different architectures for the hierarchy of features

- Feed forward (e.g., Multilayer Neural Nets, Convolutional Nets)
- Feed back (e.g., Stacked Sparse Coding, Deconvolutional Nets)
- Bi-directional (e.g, Deep Boltzmann Machines, Autoencoders)

We can have also different kind of learning protocols

- Purely supervised
- Unsupervised (layerwise) + supervised on top
- Unsupervised pre-training through regularized auto-encoders + …
- …

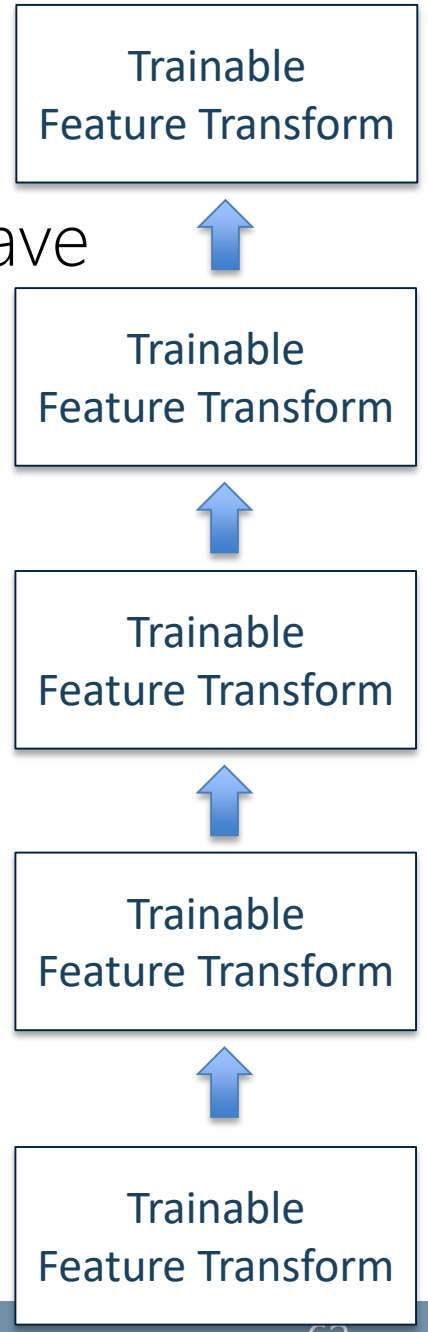

Question Time!