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

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 Matteucci, 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."

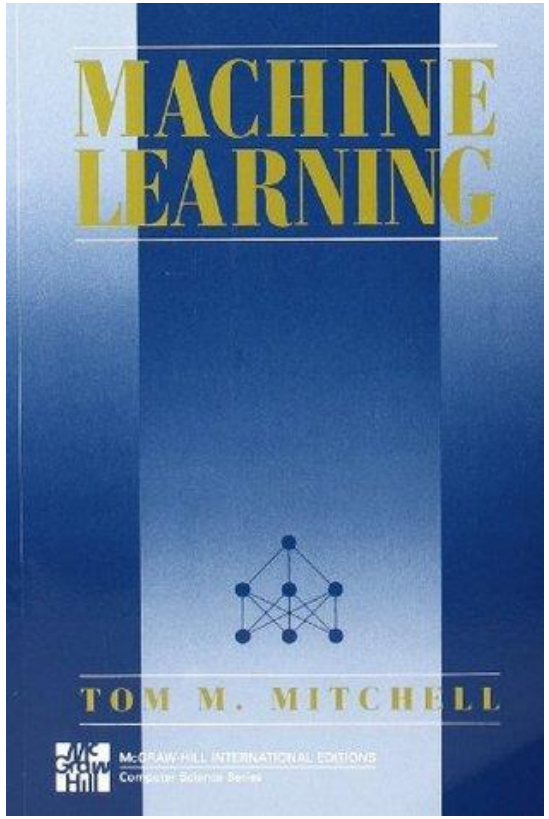

Machine Learning Paradigms

Imagine you have a certain experience D, 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

Supervised learning: Regression


12000 $


15000 $

6000 $


2000 $


8000 $


22000 $


4000 $


28000 $


6000 $


35000 $

Hand-crafted Features

Learned Regressor

3800 $

Machine Learning Paradigms

Imagine you have a certain experience D, 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



















Machine Learning Paradigms

Imagine you have a certain experience D, 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 the environment, and receiving rewards r_1, r_2 to maximize rewards in the long term

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

What about Deep Learning?

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.

What is Deep Learning after all?

… let's say it with flowers!

The Iris Dataset

Collected by Ronald Fisher in 1936

Iris Setosa

Iris Virginica

Iris Versicolor

What is Deep Learning after all?

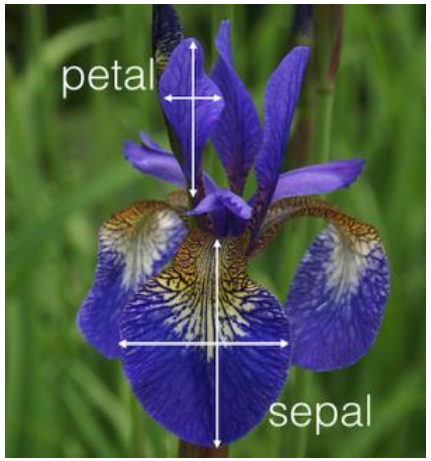

Hand-crafted Features

Learned Classifier

Machine learns how to take the Iris apart

setosa

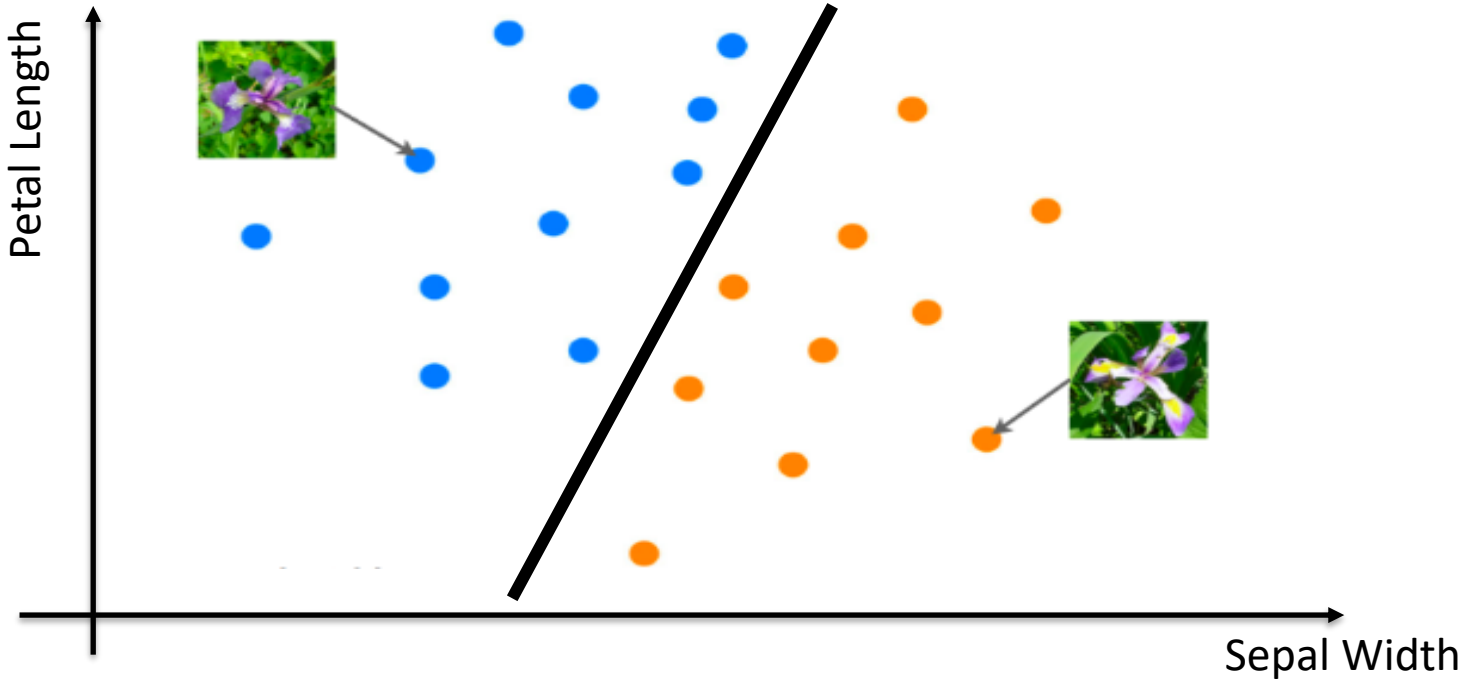

What is Deep Learning after all?

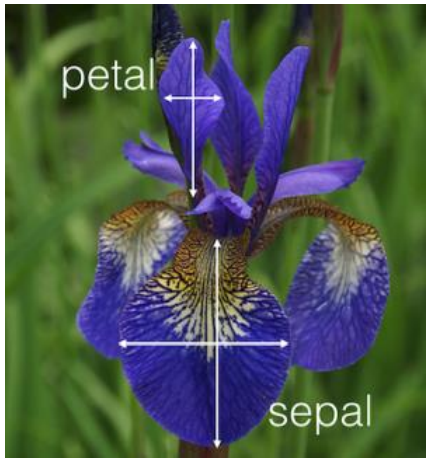

Hand-crafted Features

Learned Classifier

setosa

Machine learns how to take the Iris apart

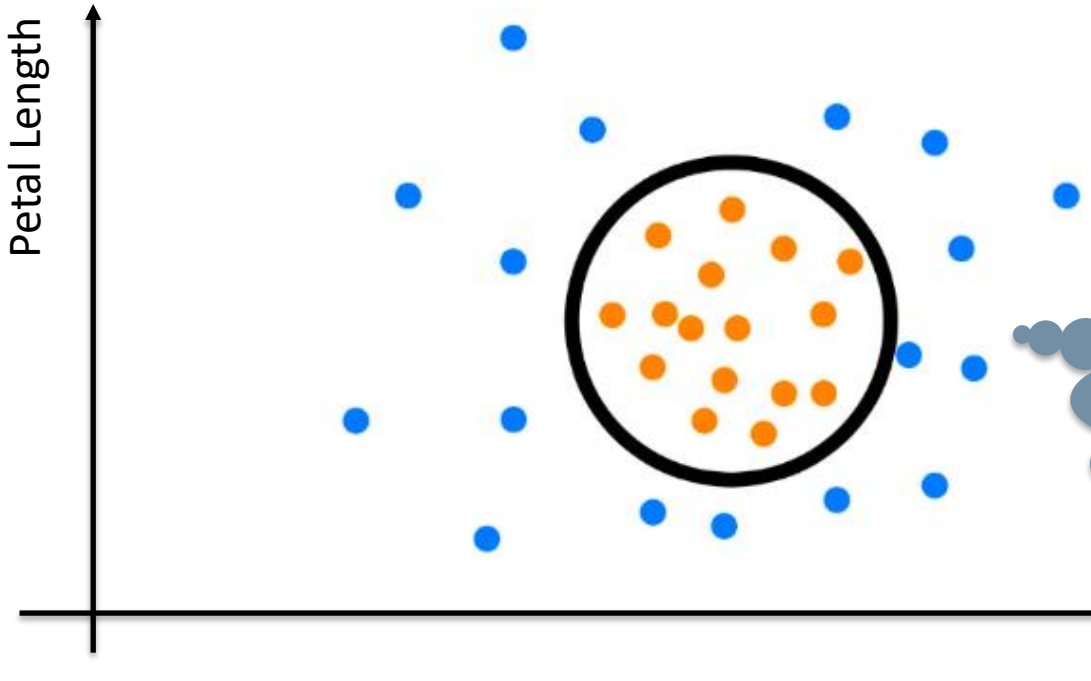

Sometimes the decision might be more complex

What is Deep Learning after all?

Machine learns how to take the Iris apart

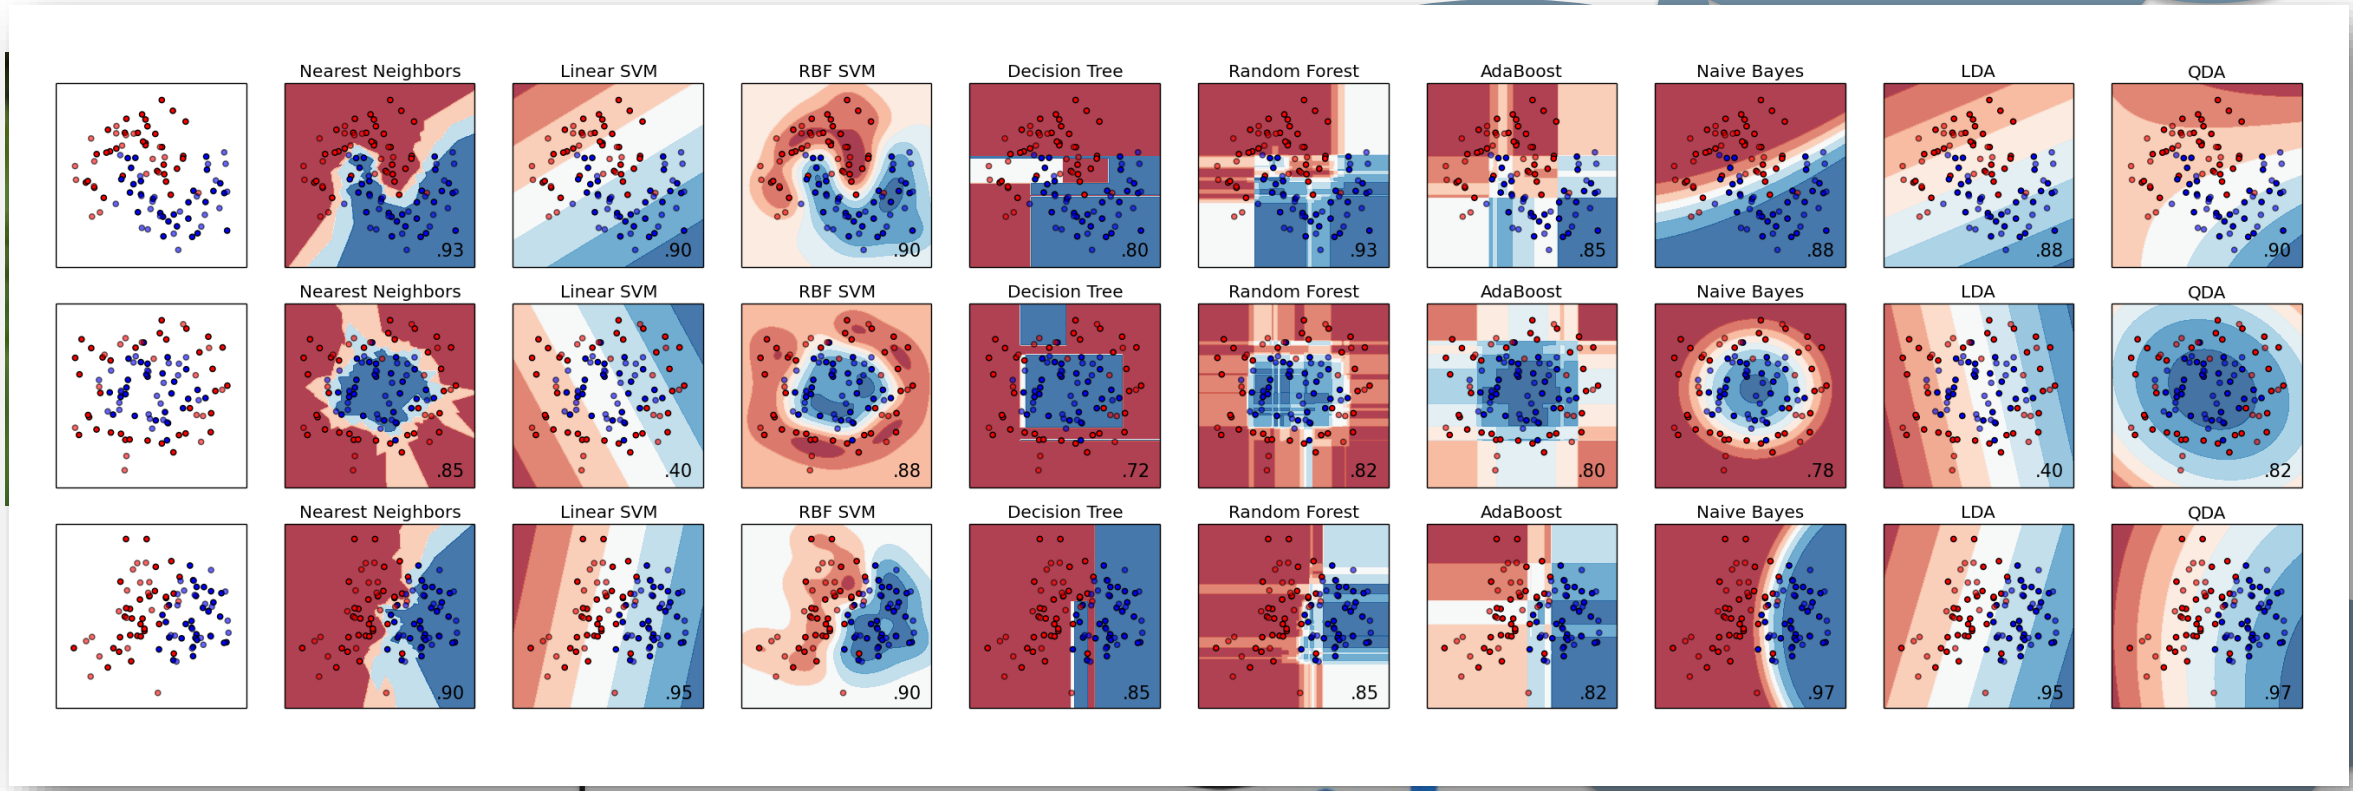

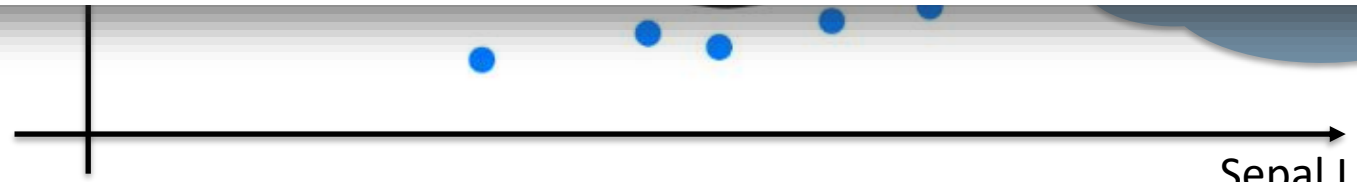

What is Deep Learning after all?

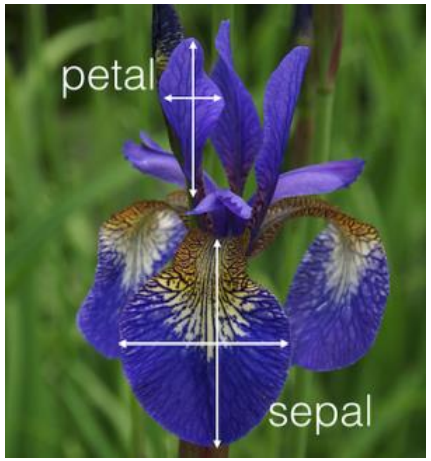

Hand-crafted Features

Learned Classifier

setosa

Machine learns how to take the Iris apart

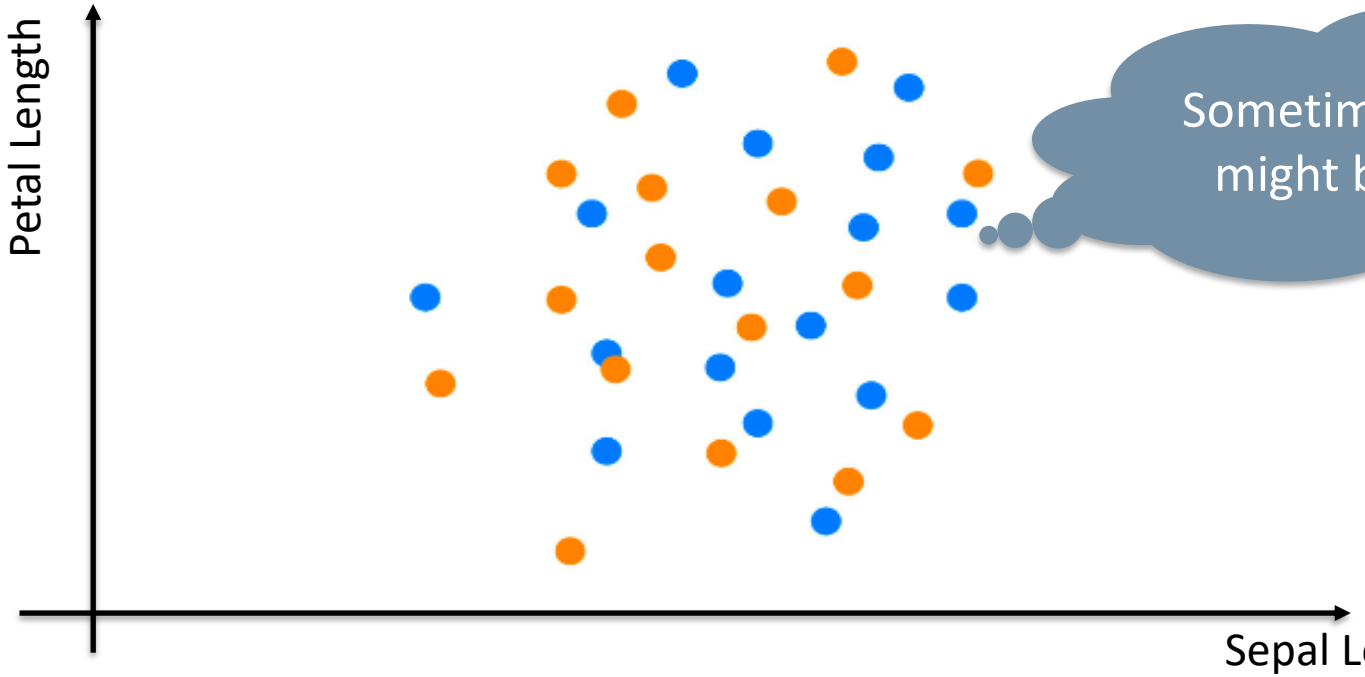

Sometimes the decision might be Impossible!

What is Deep Learning after all?

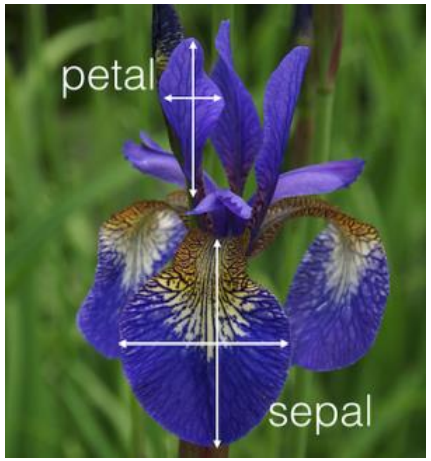

This happens if you do not know which features to extract!!!

Hand-crafted Features

Learned Classifier

setosa

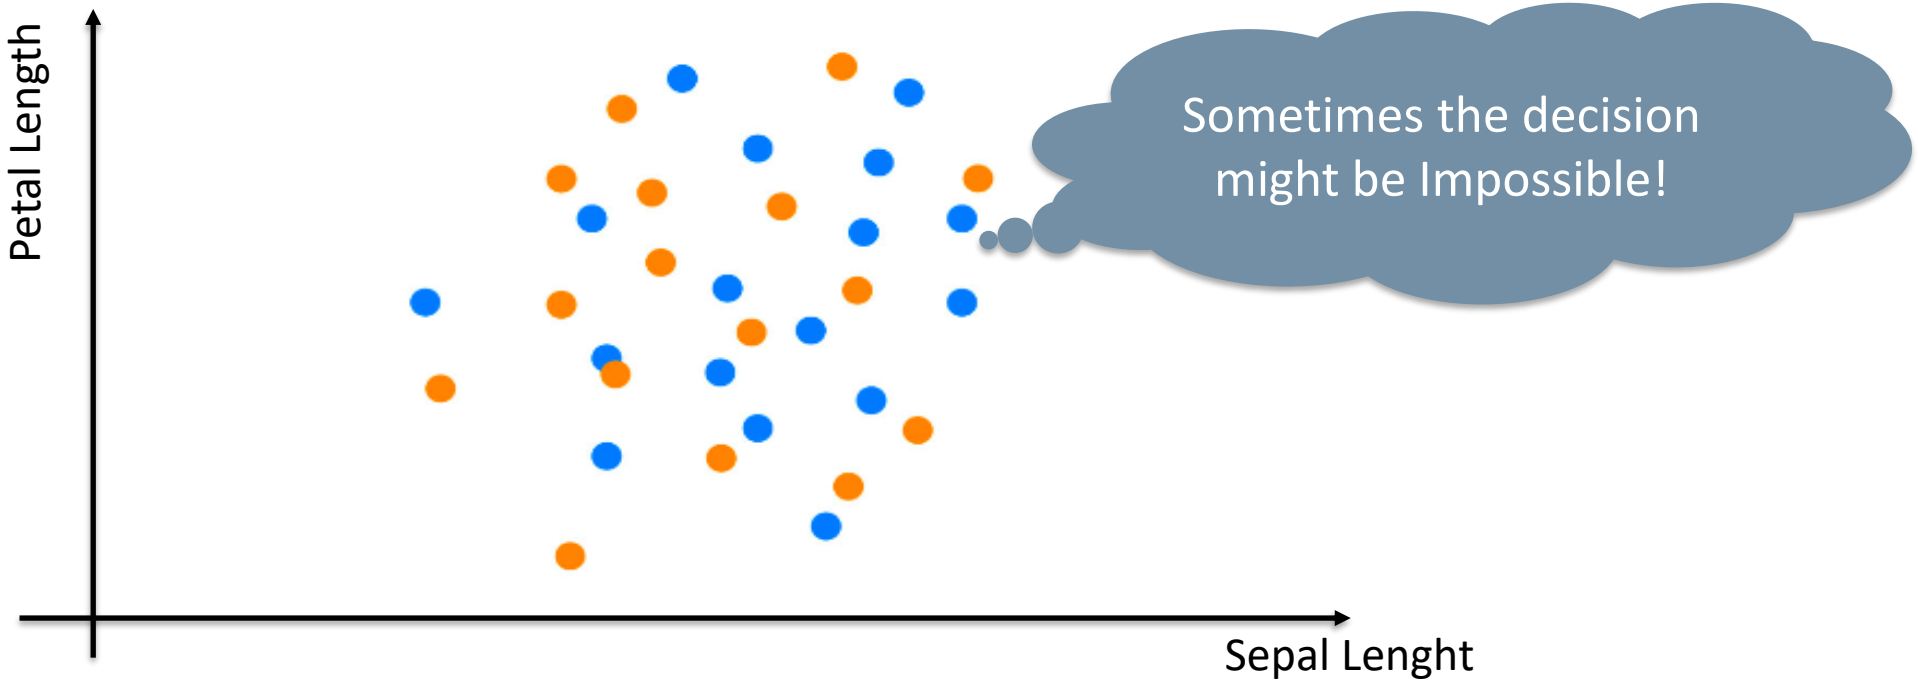

Sometimes the decision might be Impossible!

Modern Pattern Recogniton

Speech recognition (early 90's – 201

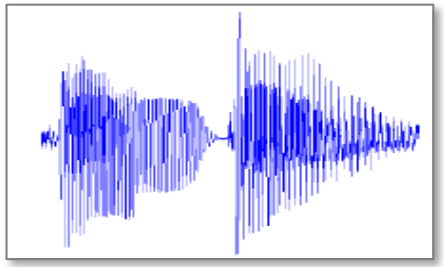

Low level features

MFCC

Fixed Features

Is this the best thing to do?

Mixture of Gaussian

Unsupervised Learning

Learned Classifier

Supervised Learning

Horn

Object recognition (2006 – 201

Low level features

SIFT/HoG

Fixed Features

And what about these features?

K-means Sparse Coding

Unsupervised Learning

Pooling

Learned Classifier

Supervised Learning

Car

What is Deep Learning after all?

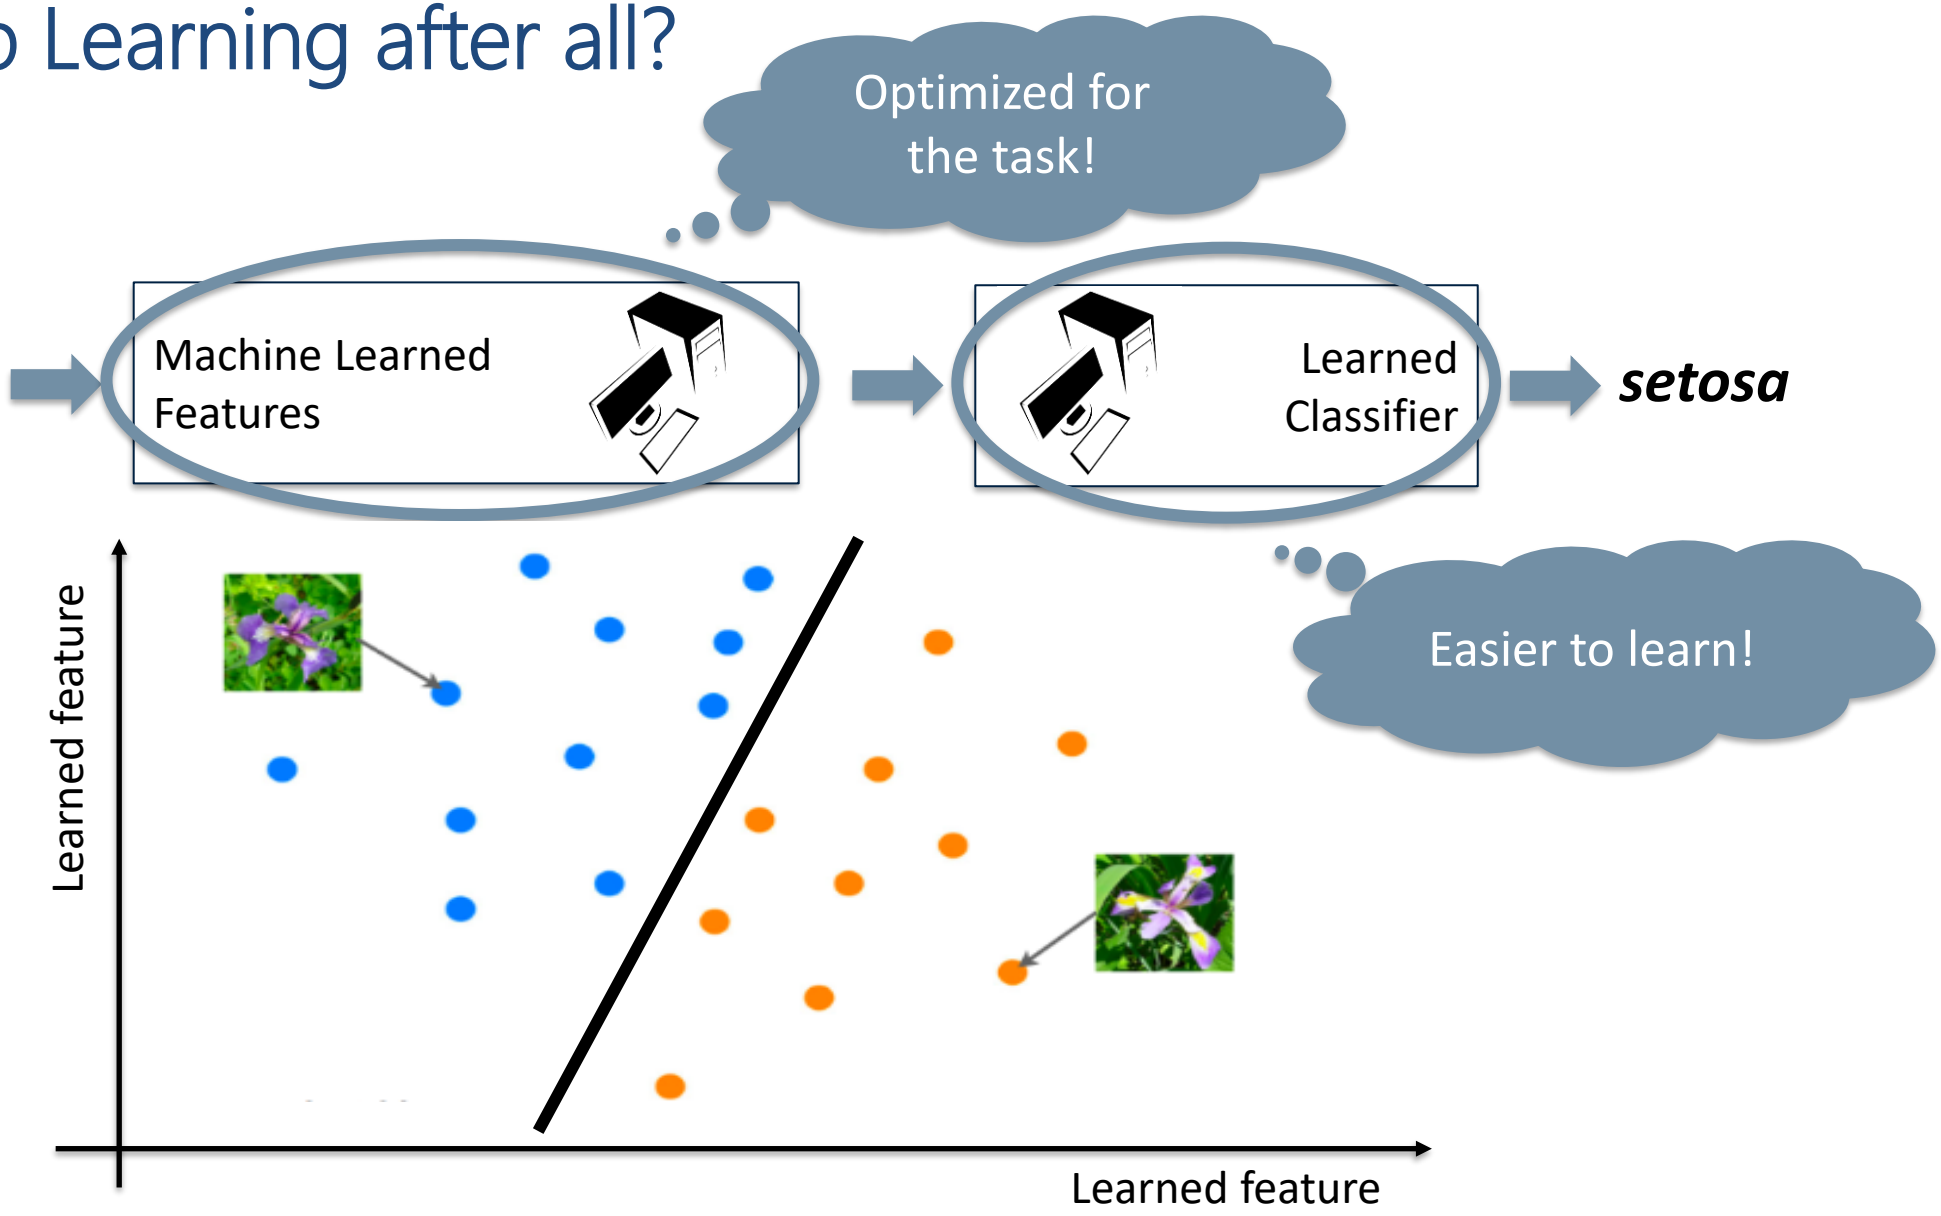

What is Deep Learning after all?

Learn from data!

Hierarchical representation optimized for the task!

Learned features

Learned features

Learned features

Learned Classifier

setosa

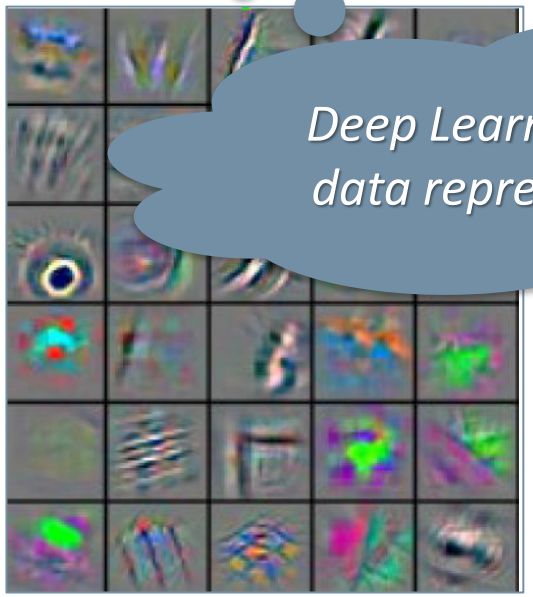

Deep Learning is about learning data representation from data!

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

What's behind Deep Learning?

MIT Technology Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Introduction | The 10 Technologies | Past Years

According to MIT, it was all about massive computational power

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 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.

What's behind Deep Learning?

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

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

IMAGENET

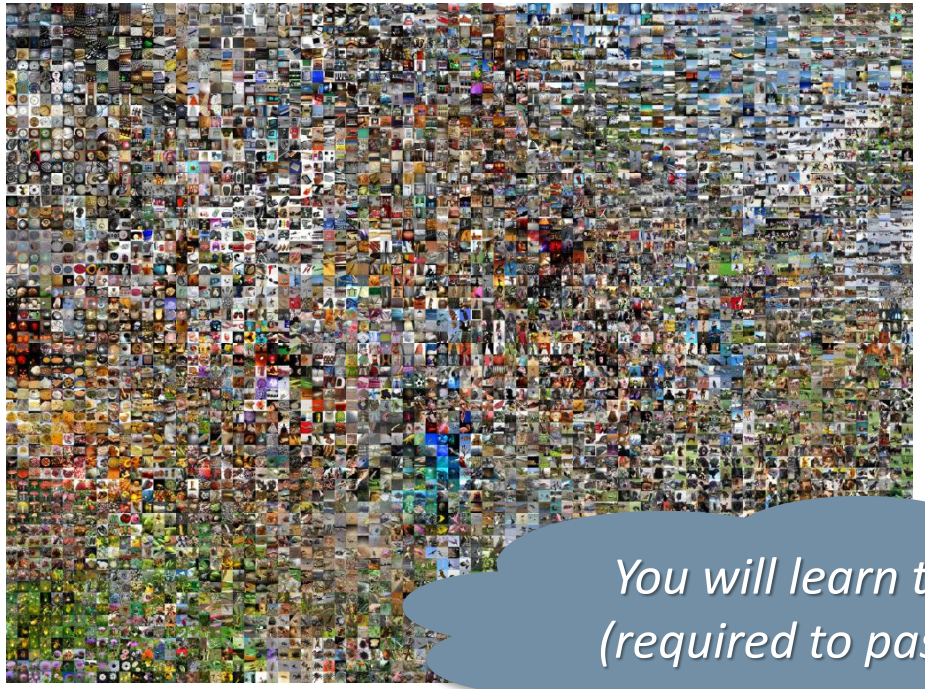

You will learn to read this!
(required to pass the exam)

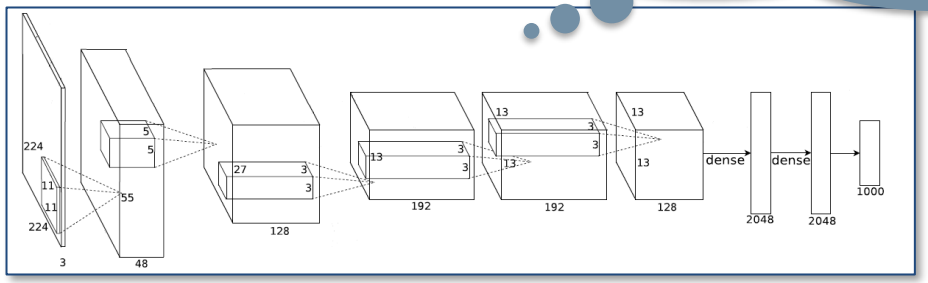

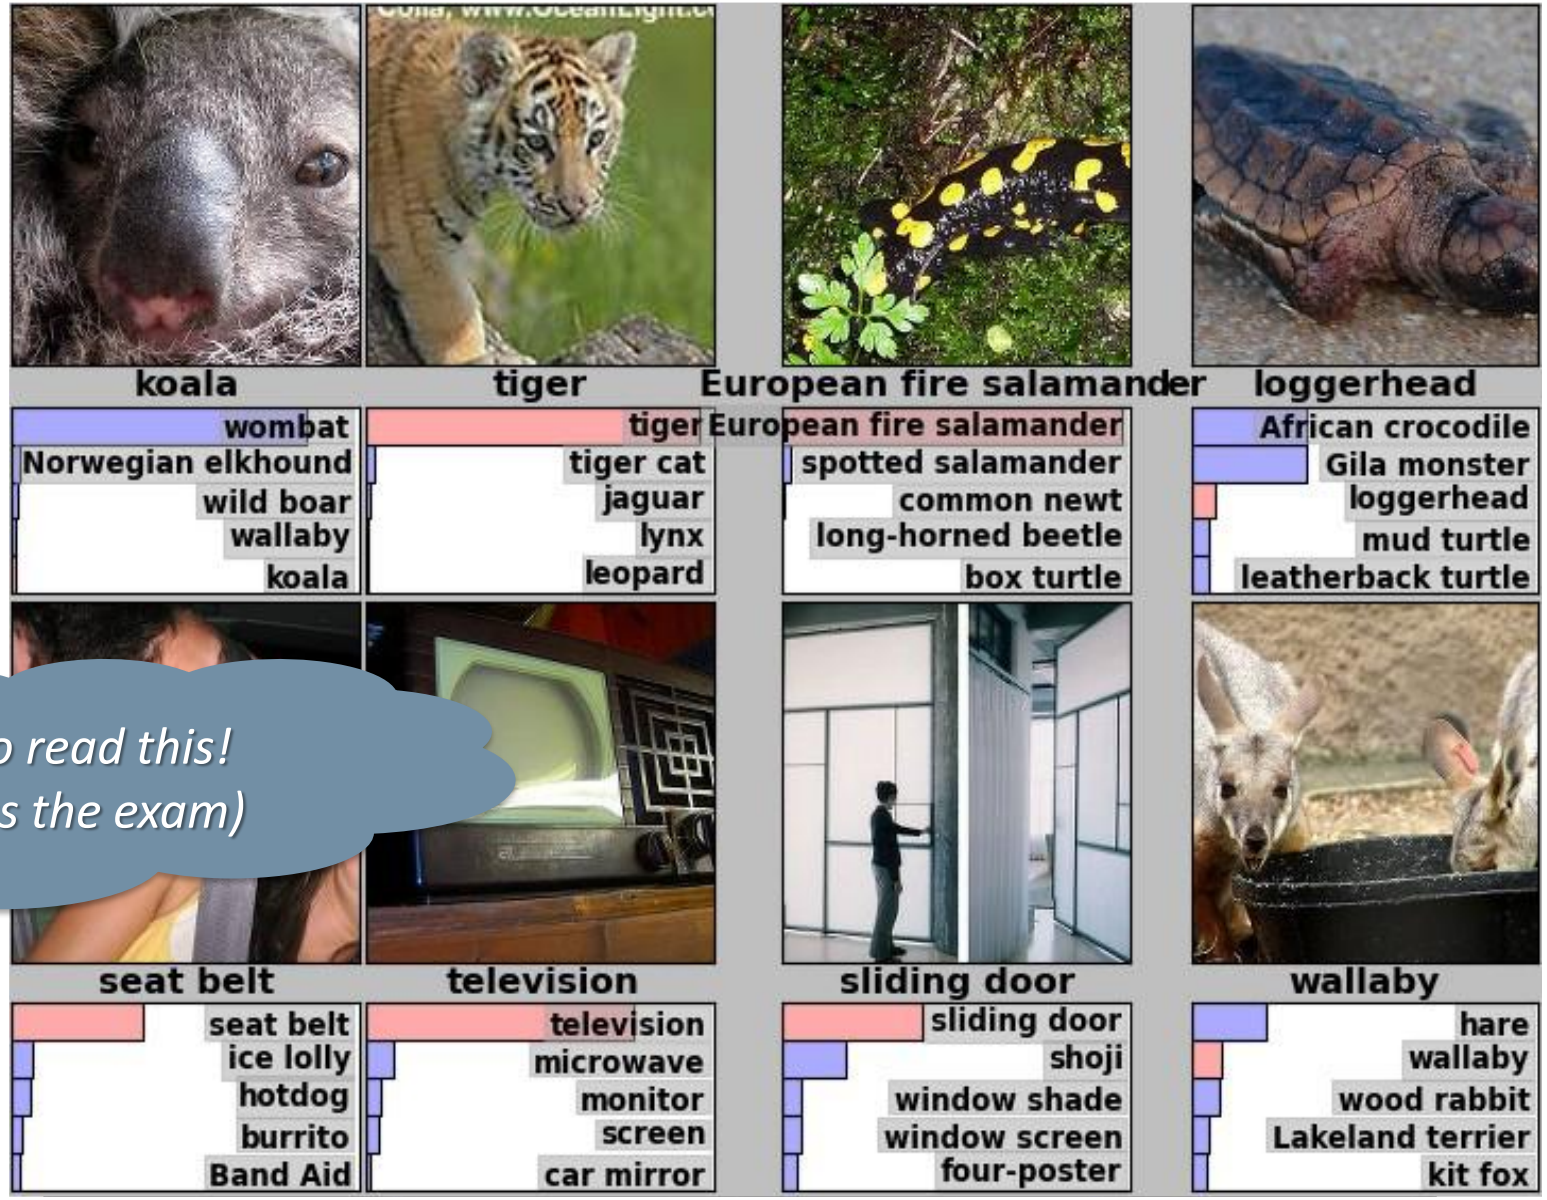

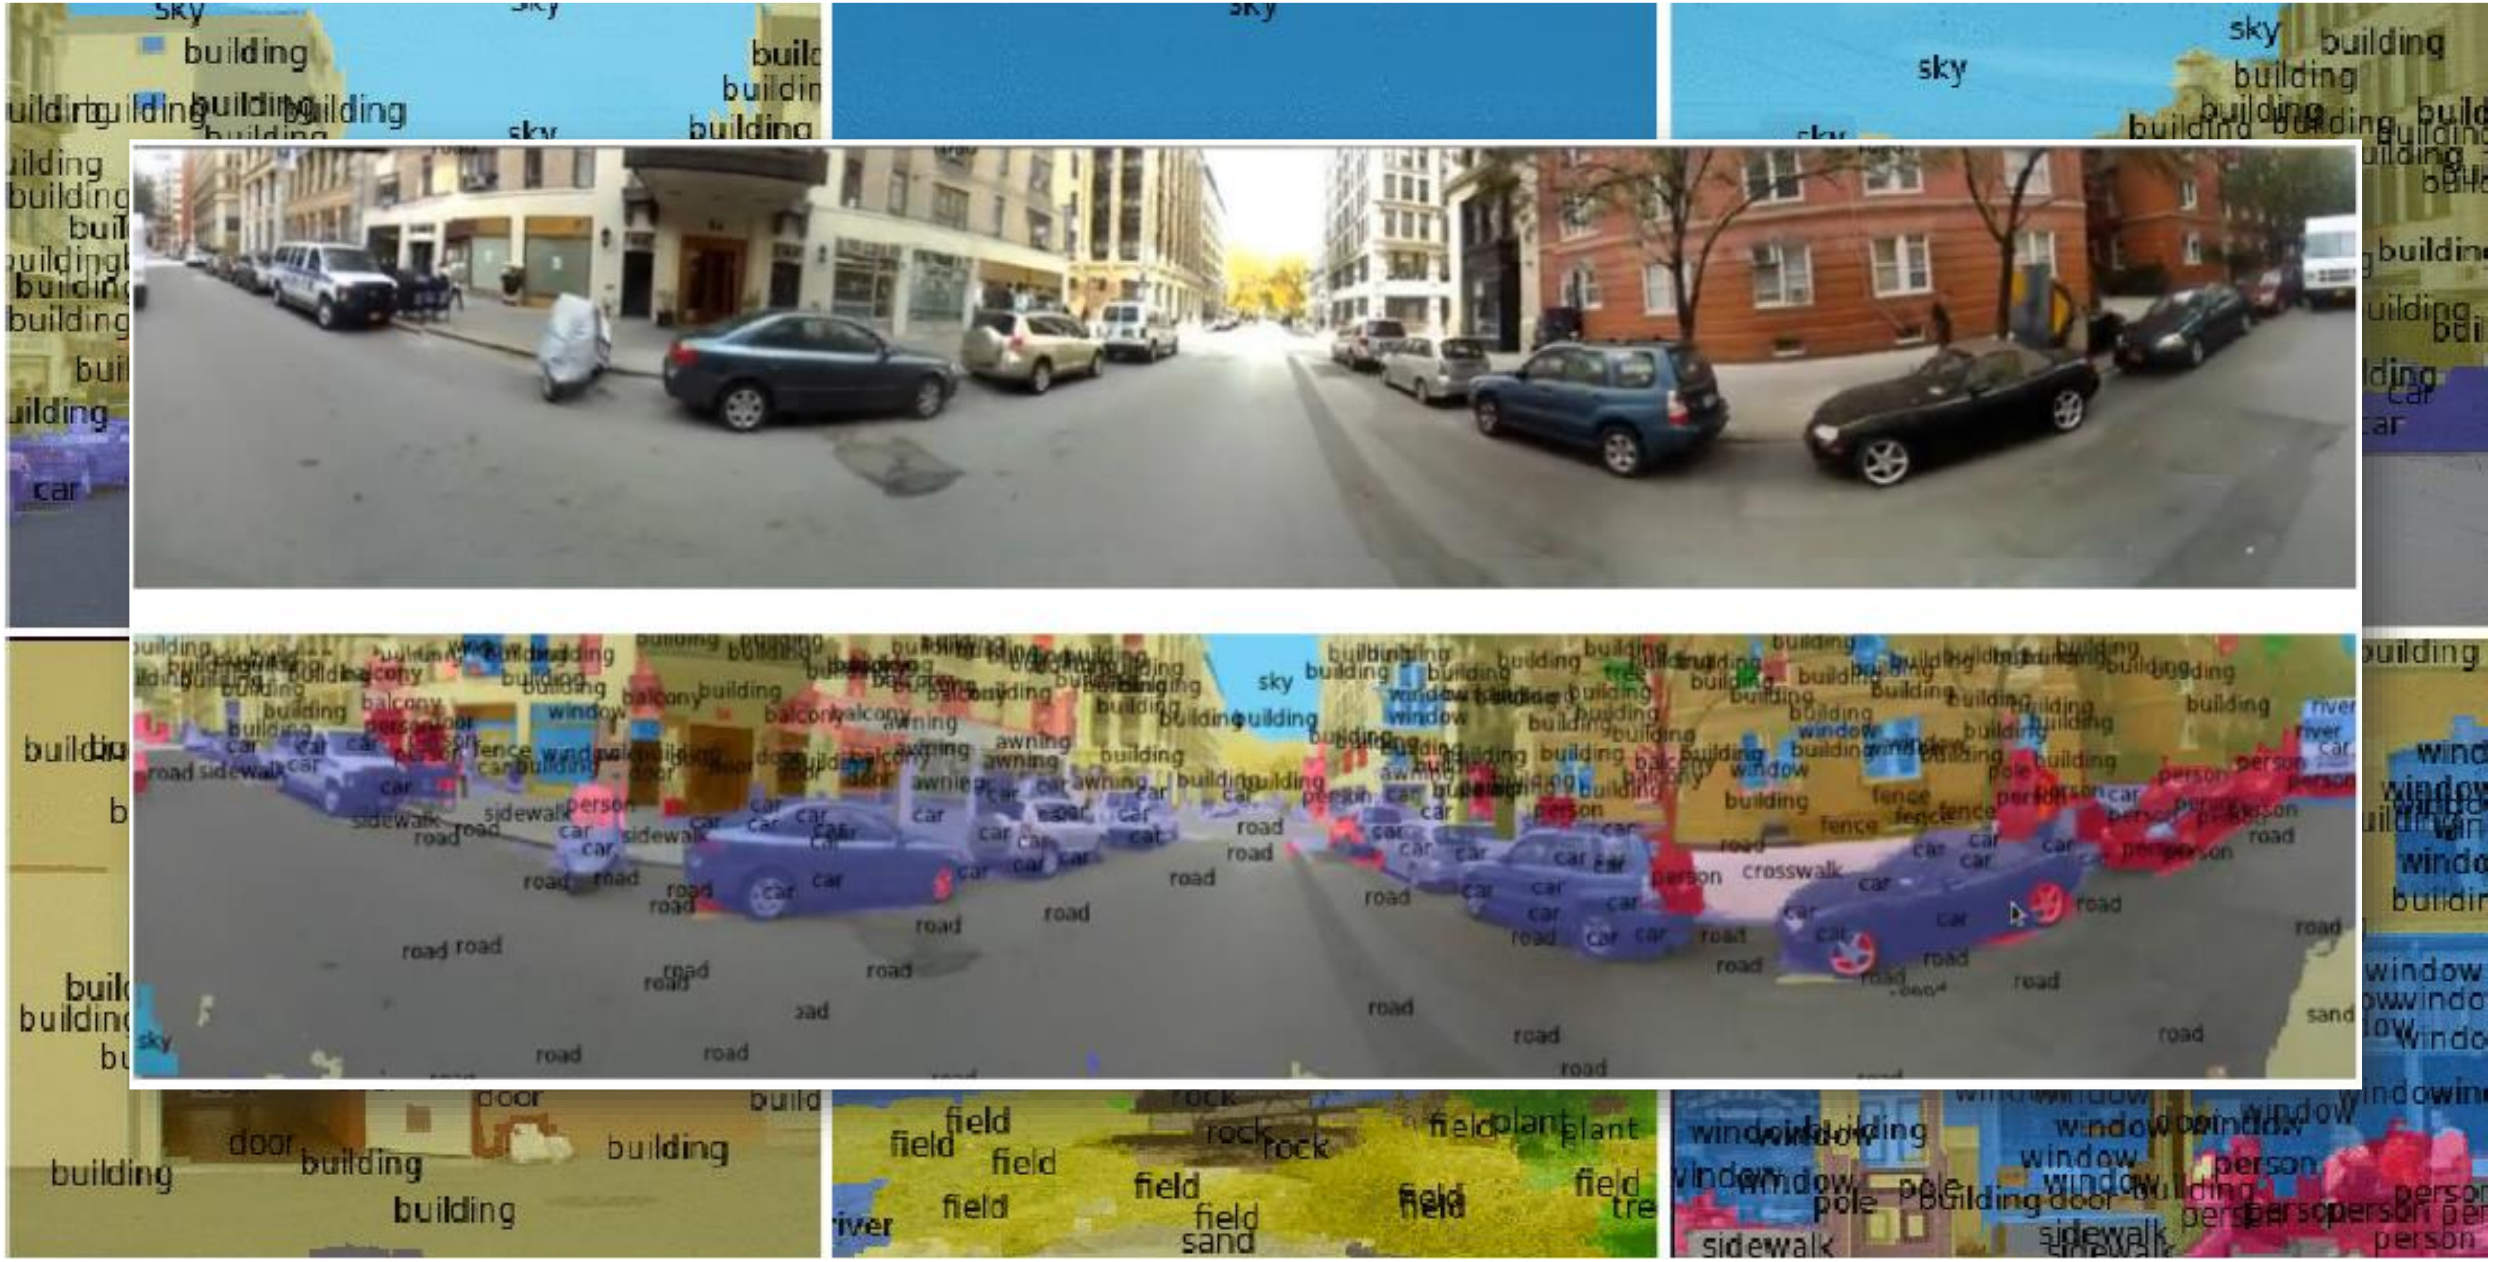
sky
building
car
road
sidewalk
person
crosswalk
door
field
sand
rock
plant
window
pole
fence
balcony
awning


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

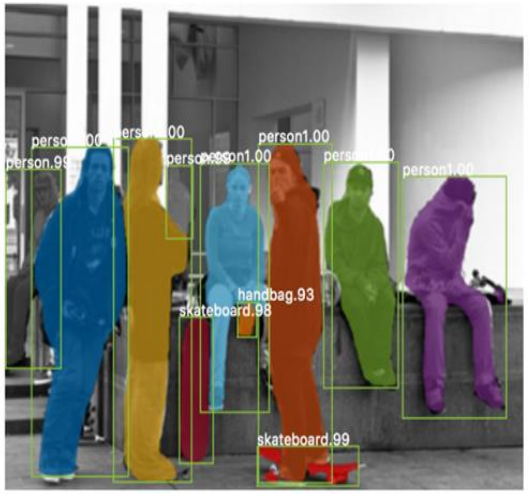
handbag.93
skateboard.98
skateboard.99

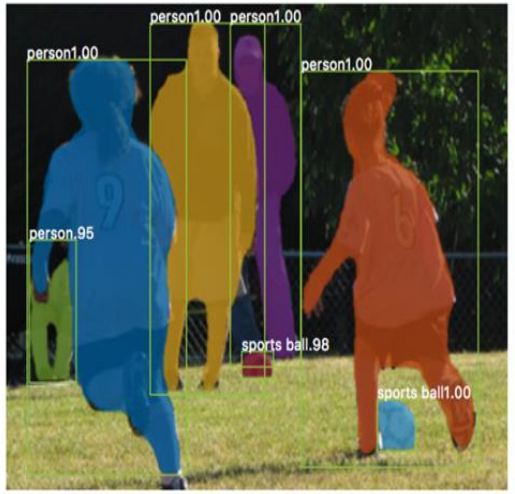
person1.00
person1.00
person1.00
person1.00
person.95
sports ball.98
sports ball1.00

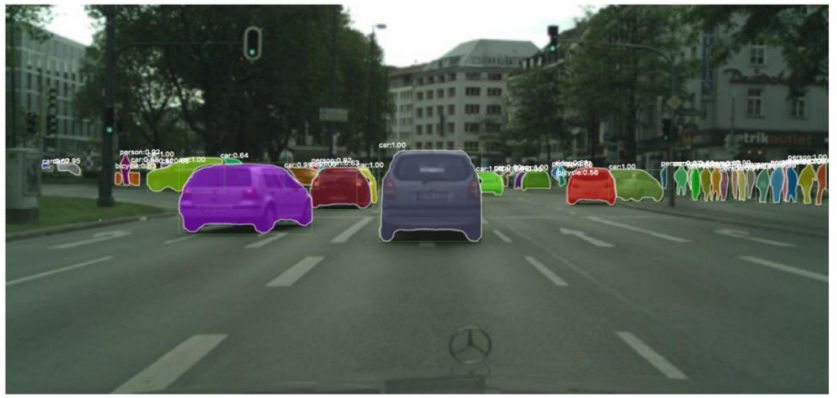

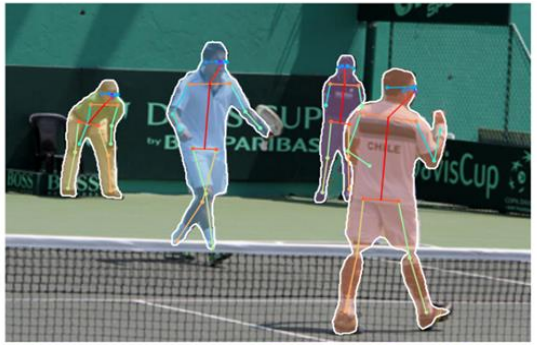

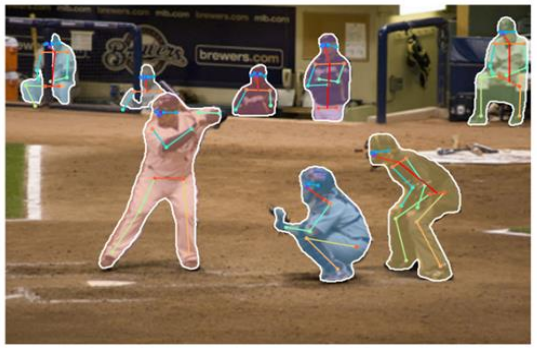

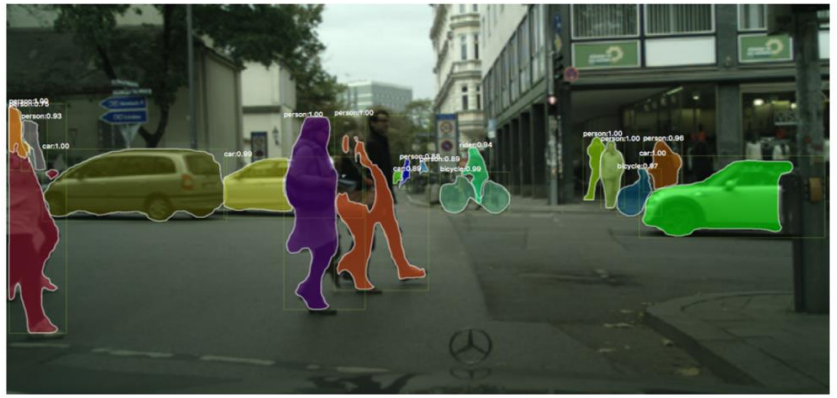

https://github.com/luanfujun/deep-photo-styletransfer

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

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

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

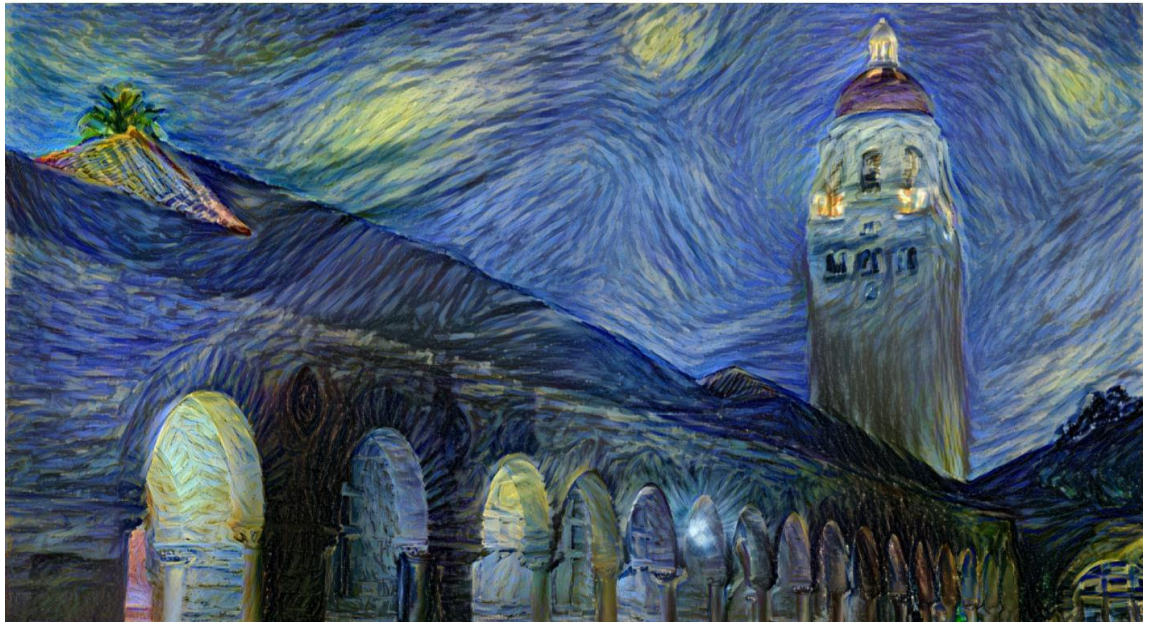

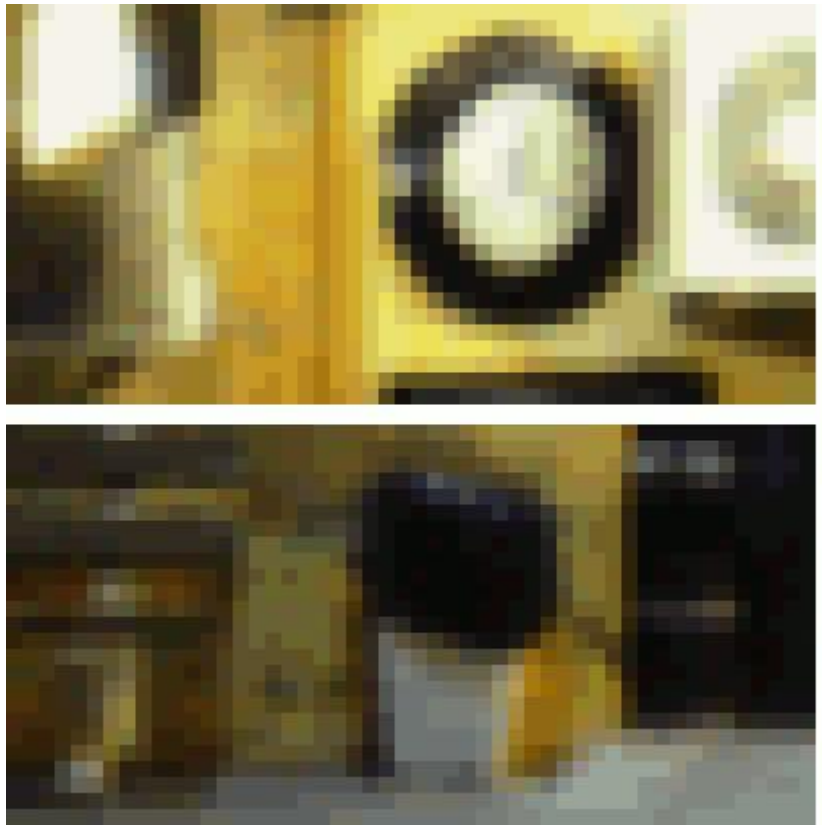





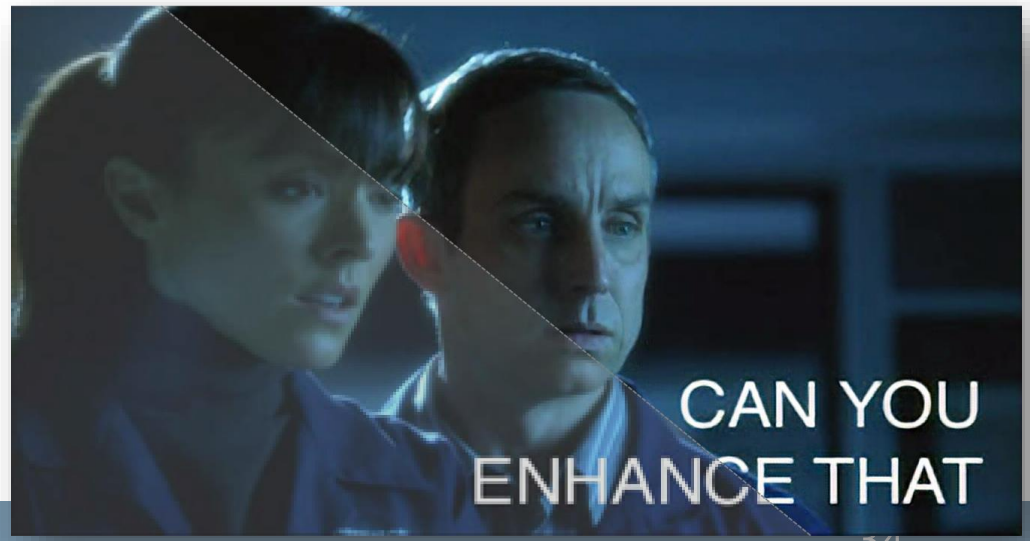

https://github.com/alexjc/neural-enhance

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							

It's all about features …

Learn from data!

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

Deep Learning (2012 – …)

That' what you'll learn in this course!

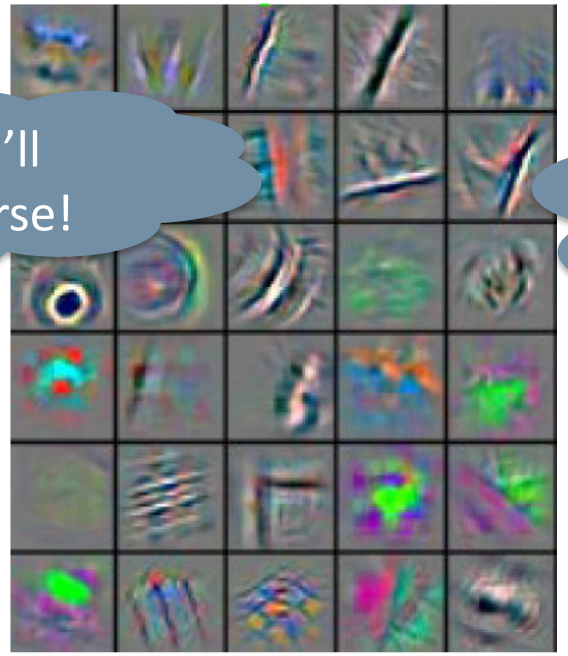

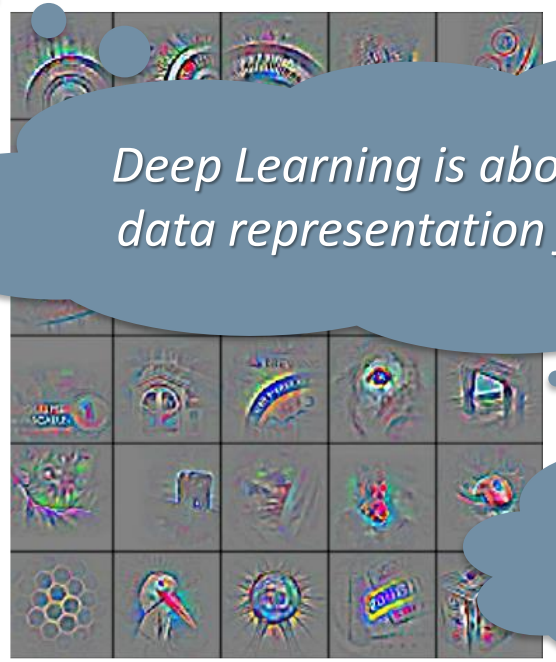

Deep Learning is about learning data representation from data!

But which data?