Machine Learning (AIMS) - MT 2018
0. (My) Introduction to ML

Varun Kanade

University of Oxford
November 5, 2018

Machine Learning in Action

(Using https://www.betafaceapi.com/demo.html)

Machine Learning in Action

(Using https://www.betafaceapi.com/demo.html)

Machine Learning in Action

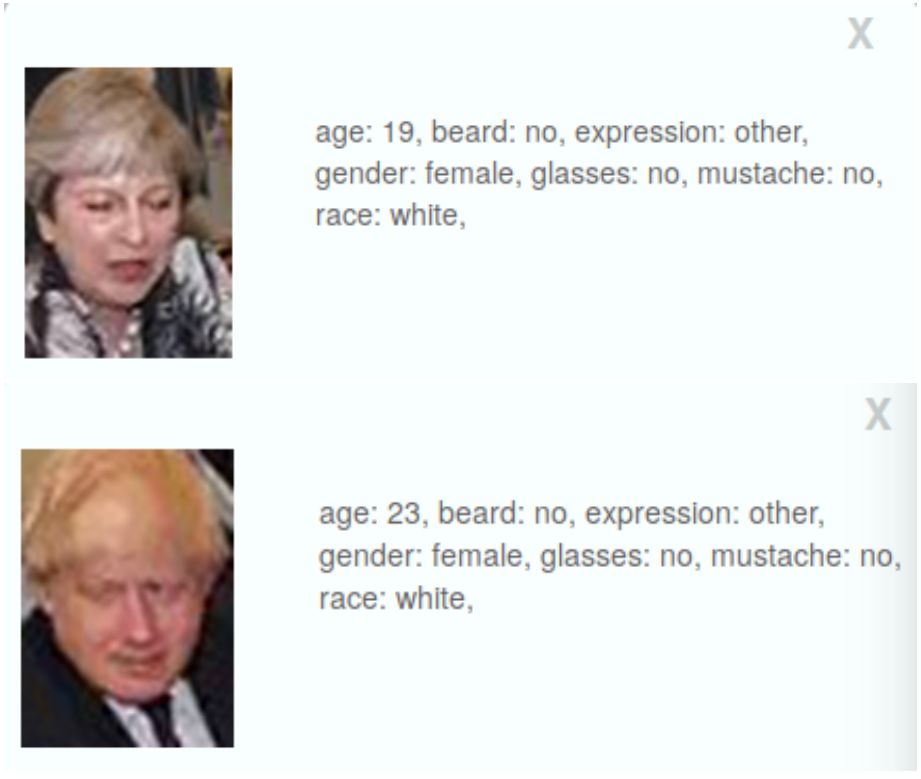

(Using https://www.betafaceapi.com/demo.html)

What is machine learning?

What is machine learning?

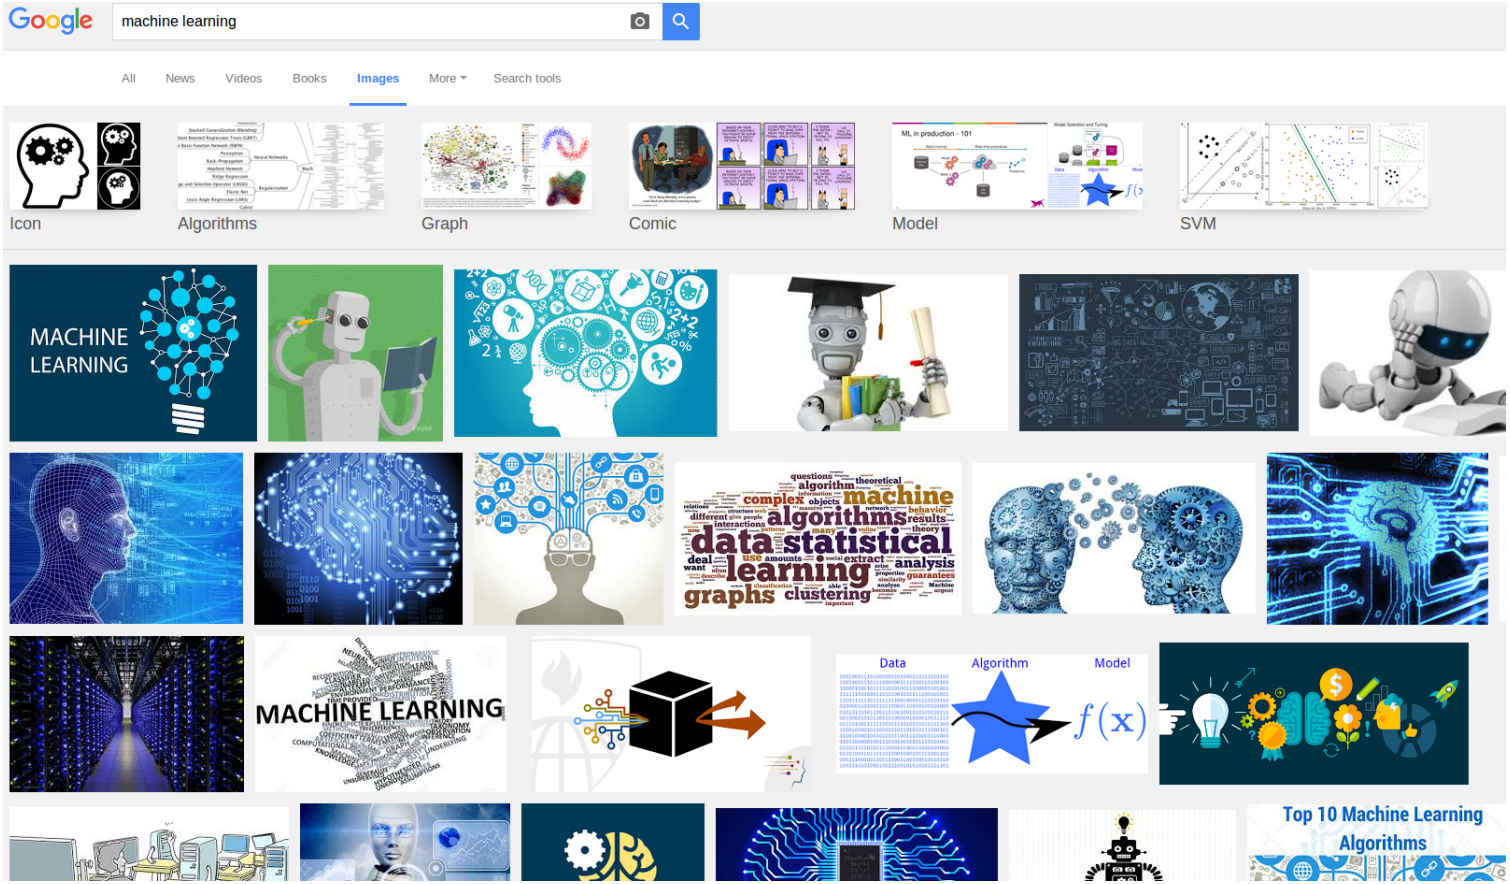

What is machine learning?

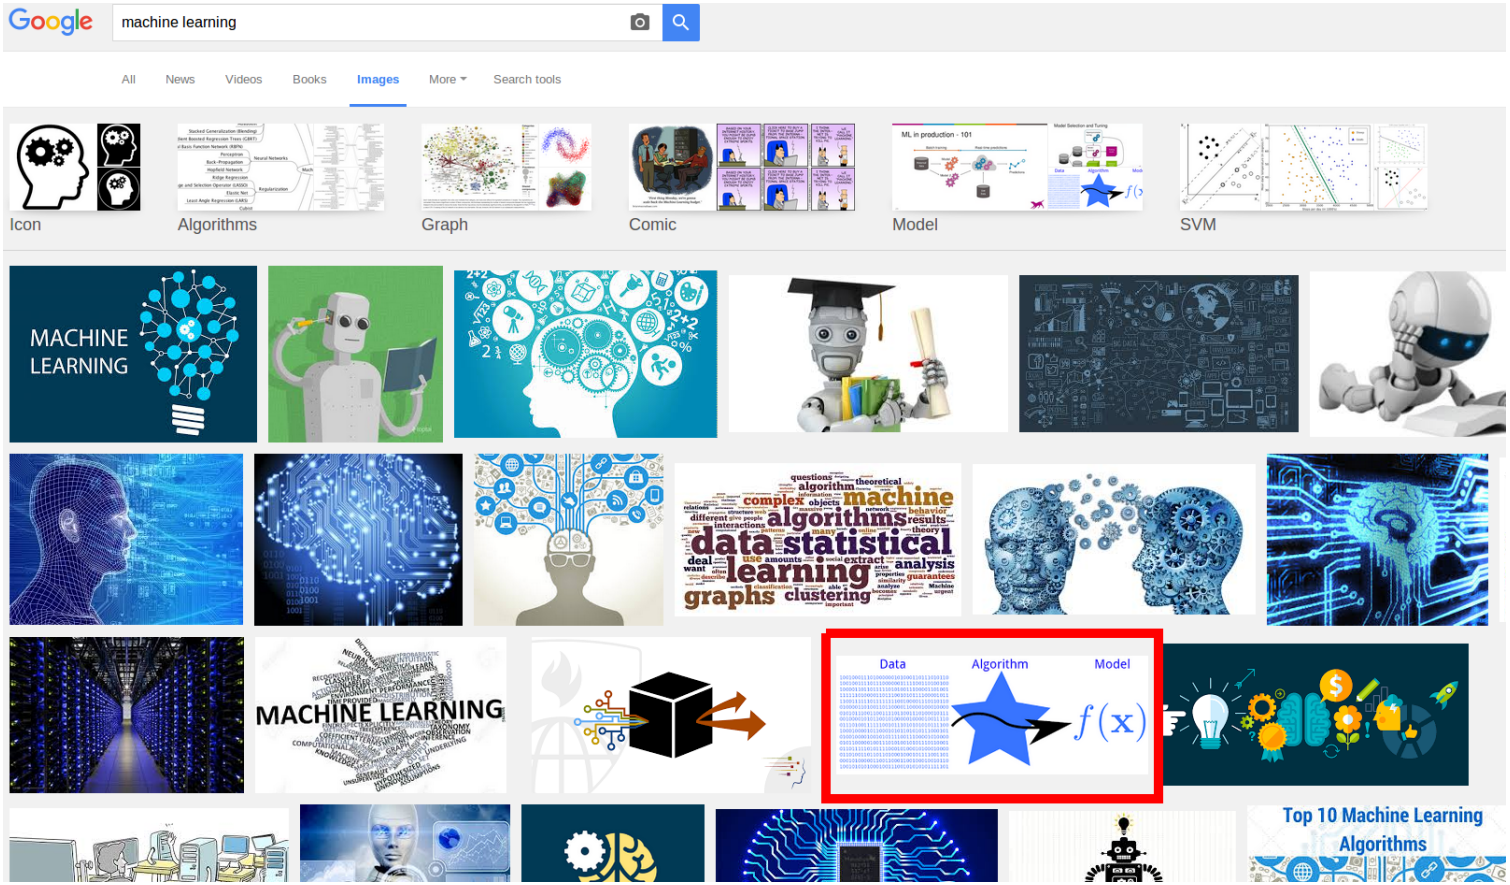

circa October 2016

What is machine learning?

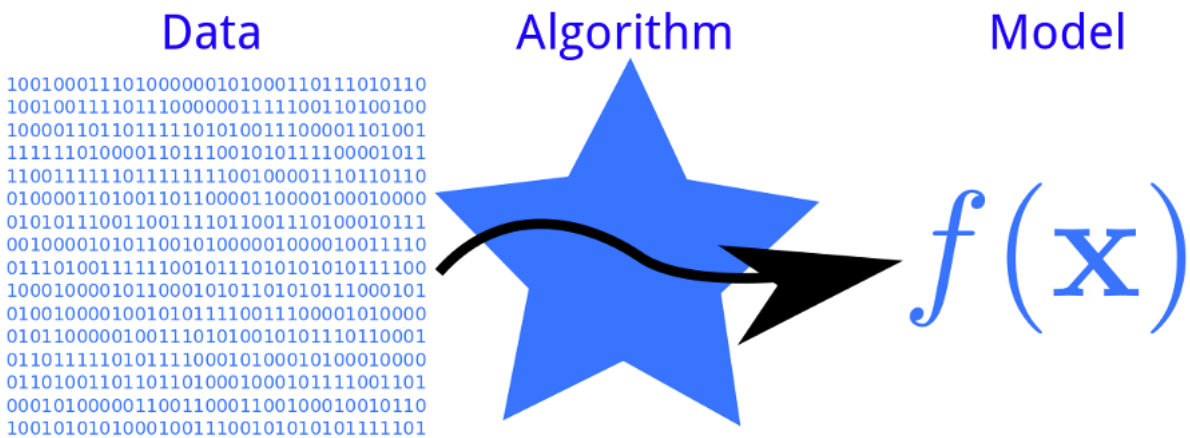

circa October 2016

What is machine learning?

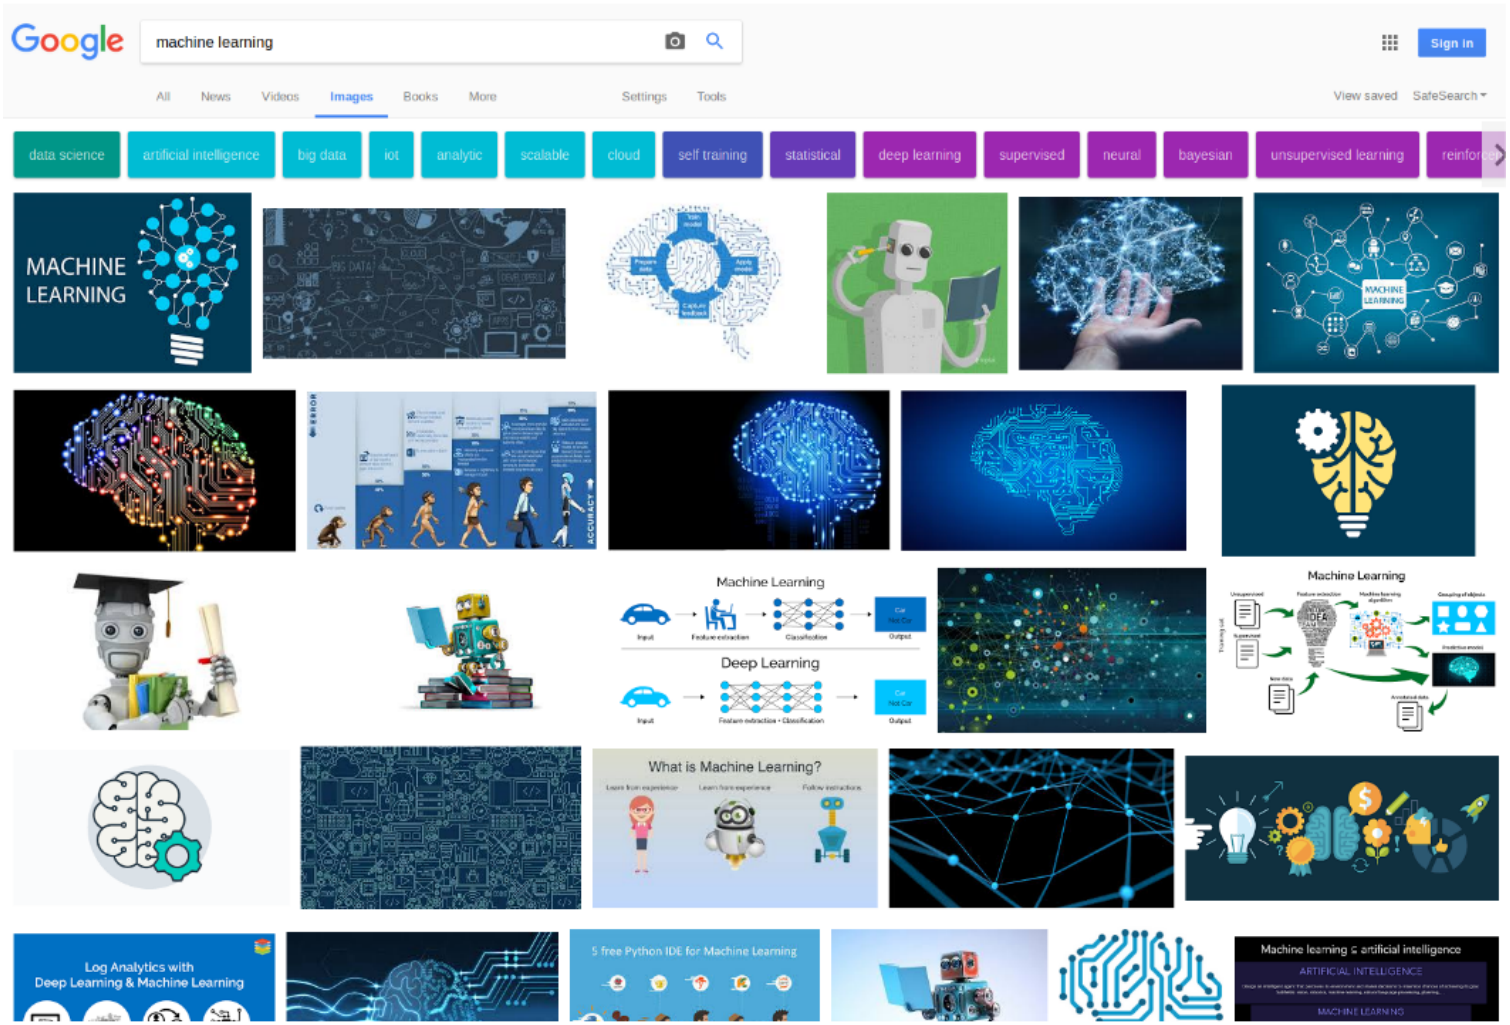

circa October 2017

What is machine learning?

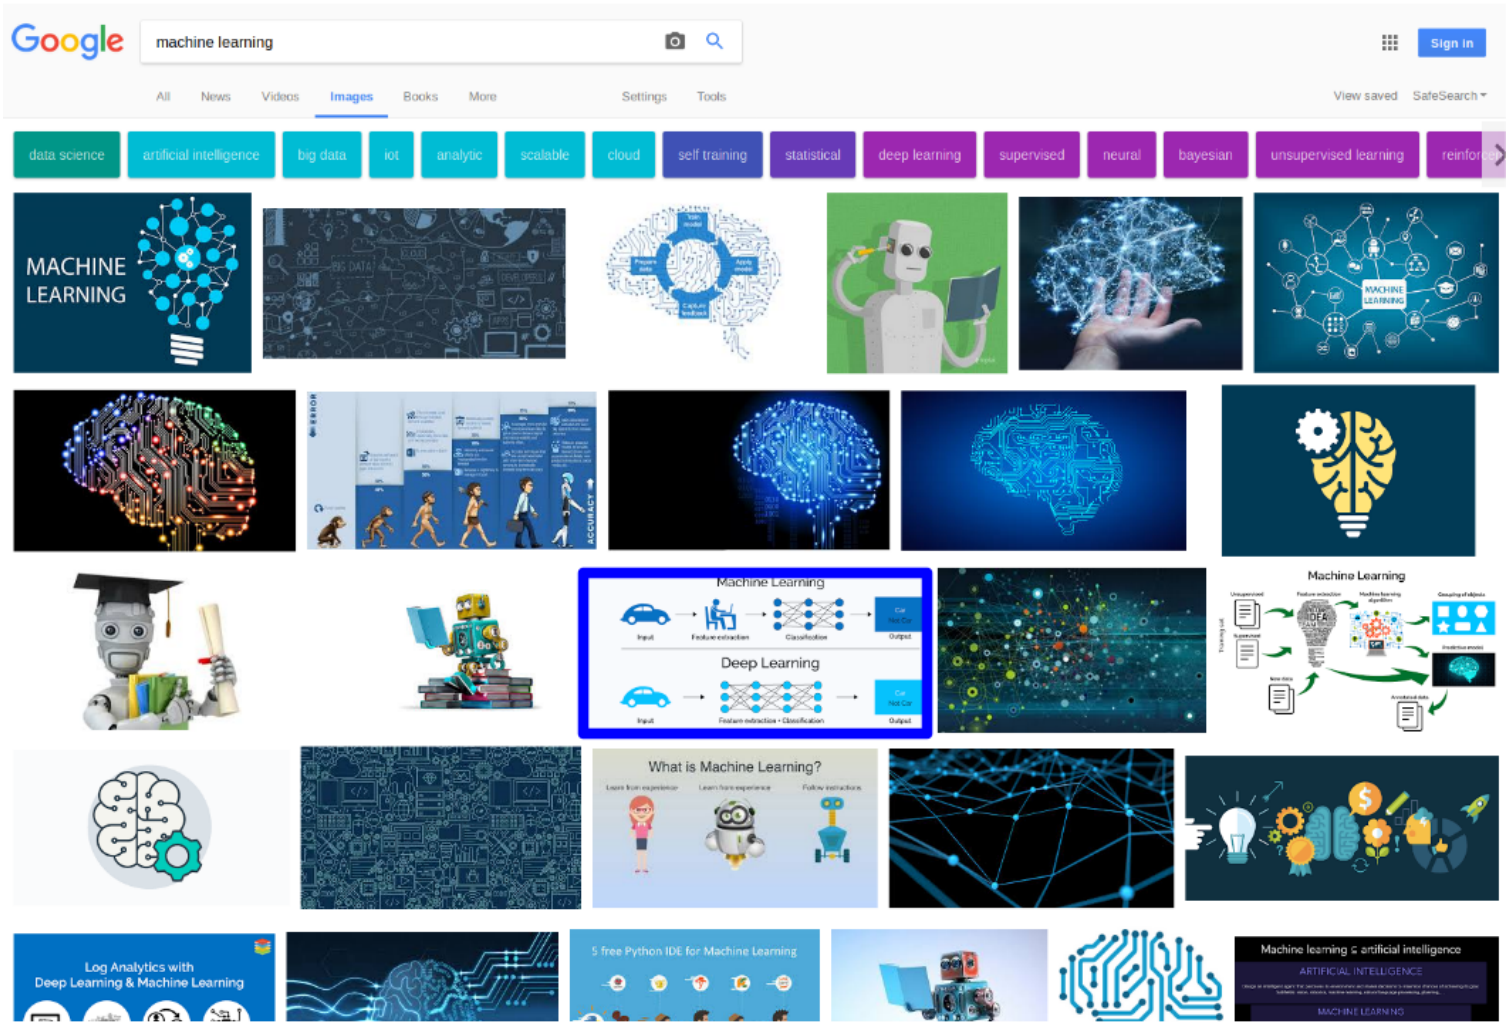

circa October 2017

What is machine learning?

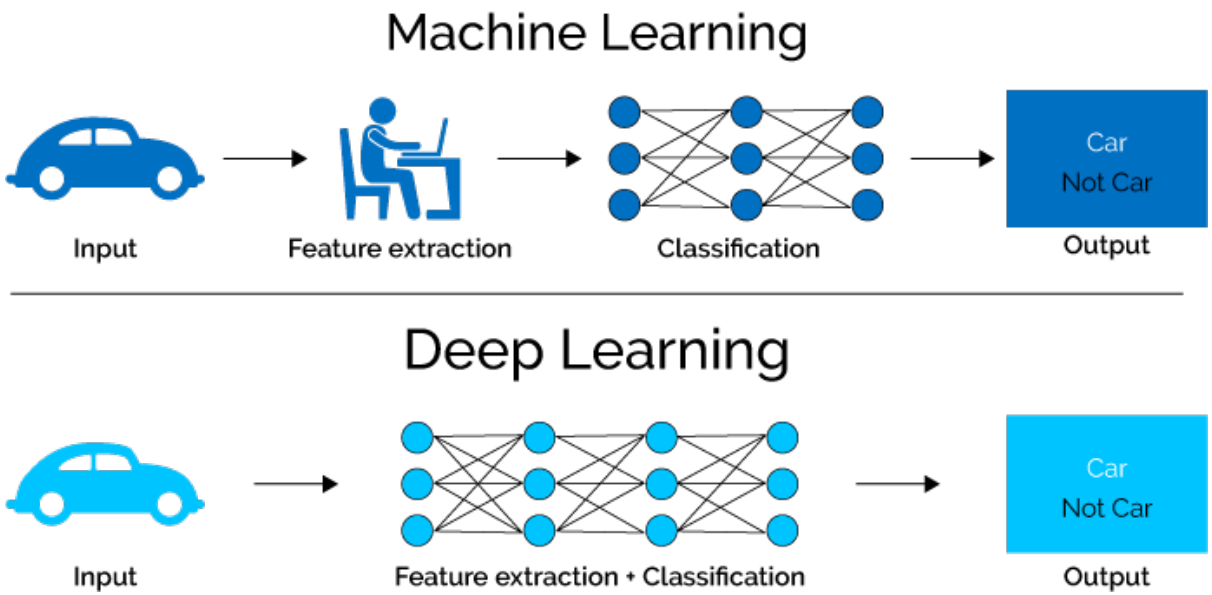

circa October 2017

What is machine learning?

What is artificial intelligence?

What is machine learning?

What is artificial intelligence?

“Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child’s? If this were then subjected to an appropriate course of education one would obtain the adult brain.”

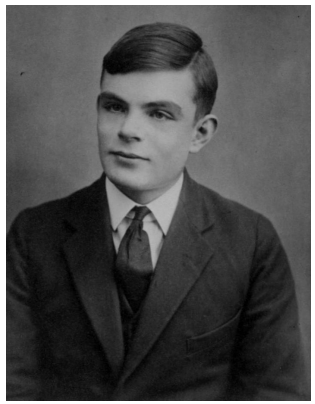

Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59, 433-460.

What is machine learning?

Definition by Tom Mitchell

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

What is machine learning?

Definition by Tom Mitchell

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

Face Detection

What is machine learning?

Definition by Tom Mitchell

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

Face Detection

- E : images (with bounding boxes) around faces

What is machine learning?

Definition by Tom Mitchell

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

Face Detection

- E : images (with bounding boxes) around faces
- T : given an image without boxes, put boxes around faces

What is machine learning?

Definition by Tom Mitchell

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

Face Detection

- E : images (with bounding boxes) around faces
- T : given an image without boxes, put boxes around faces
- P : number of faces correctly identified

Outline

Course Logistics

Some Machine Learning Applications

Course Information

Website

`www.cs.ox.ac.uk/people/varun.kanade/teaching/ML-AIMS-MT2018/`

Lectures

Mon-Thu: 10h30 - 12h30; Fri: 9h-11h (Lecture Room 7 all days)

Practicals

Mon-Thu: 15h-17h (?)

Demonstrators: Philip Lazos, David Martínez

Textbooks

Kevin Murphy - Machine Learning: A Probabilistic Perspective

- Online access through Bodleian library

Other posted online material

A few last remarks about this course

As ML developed through various disciplines - CS, Stats, Neuroscience, Engineering, *etc.*, there is no consistent usage of notation or even names among the textbooks. At times you may find inconsistencies even within a single textbook.

You will be required to read, both before and after the lectures. I will post suggested reading on the website.

Resources:

- Wikipedia has many great articles about ML and background material
- Online videos: Andrew Ng on coursera, Nando de Freitas on youtube, *etc.*
- Many interesting blogs, podcasts, *etc.*

NIPS Papers!

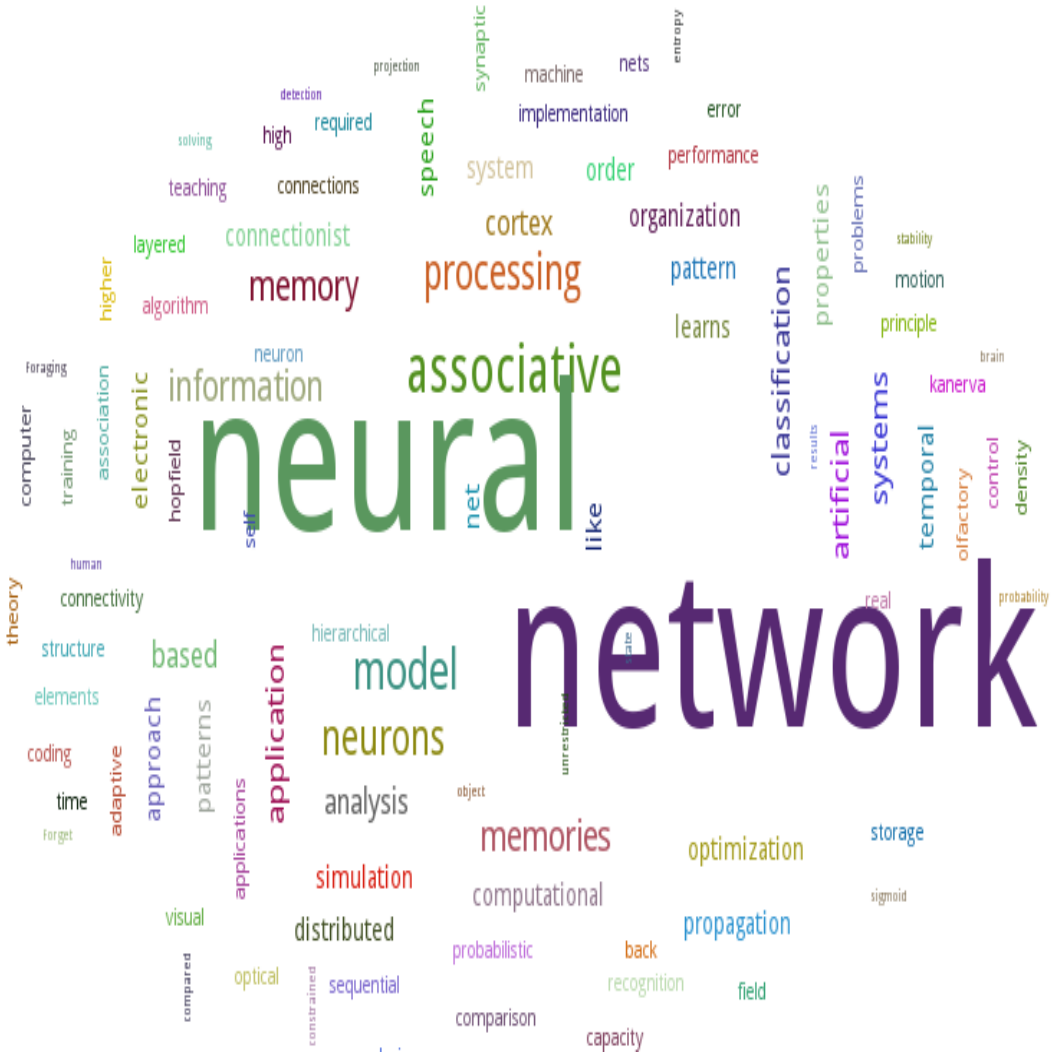

Advances in Neural Information Processing Systems 1988

NIPS Papers!

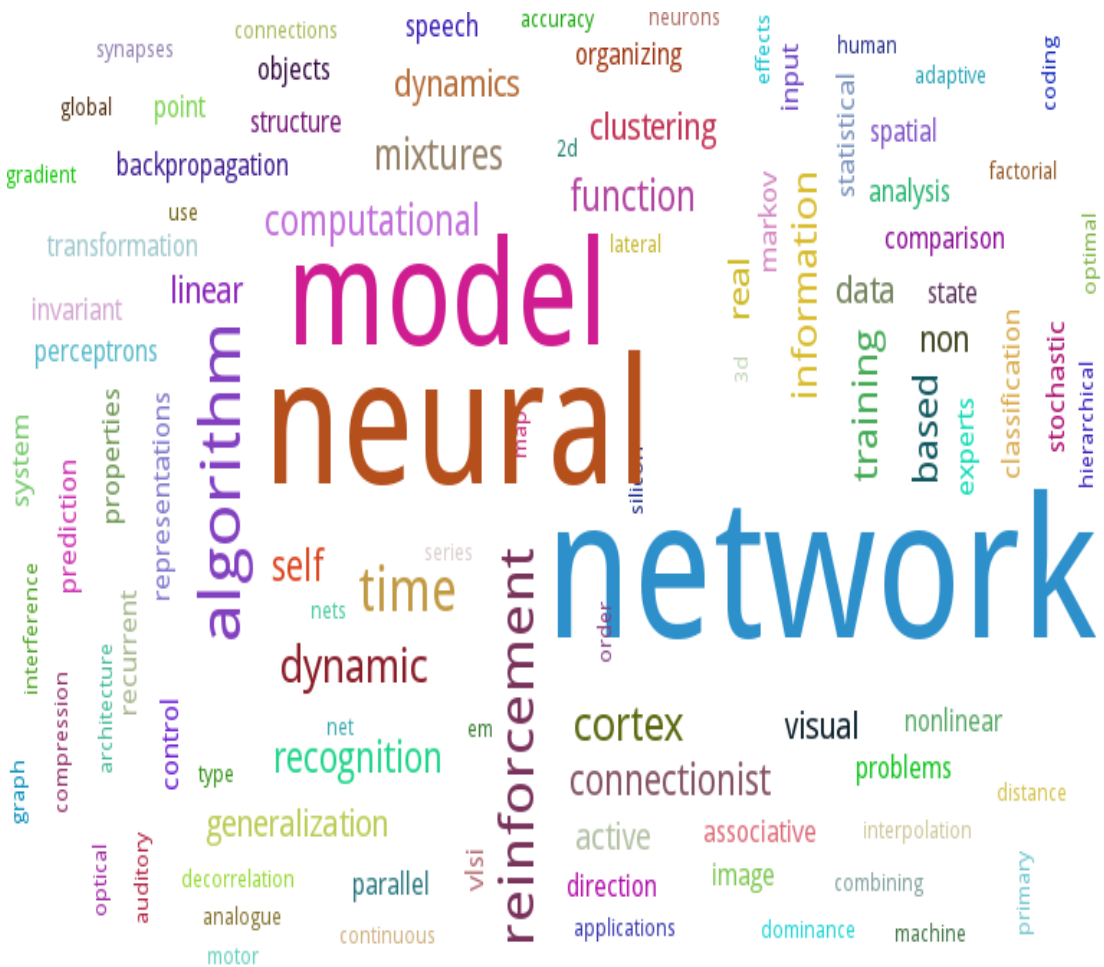

Advances in Neural Information Processing Systems 1995

NIPS Papers!

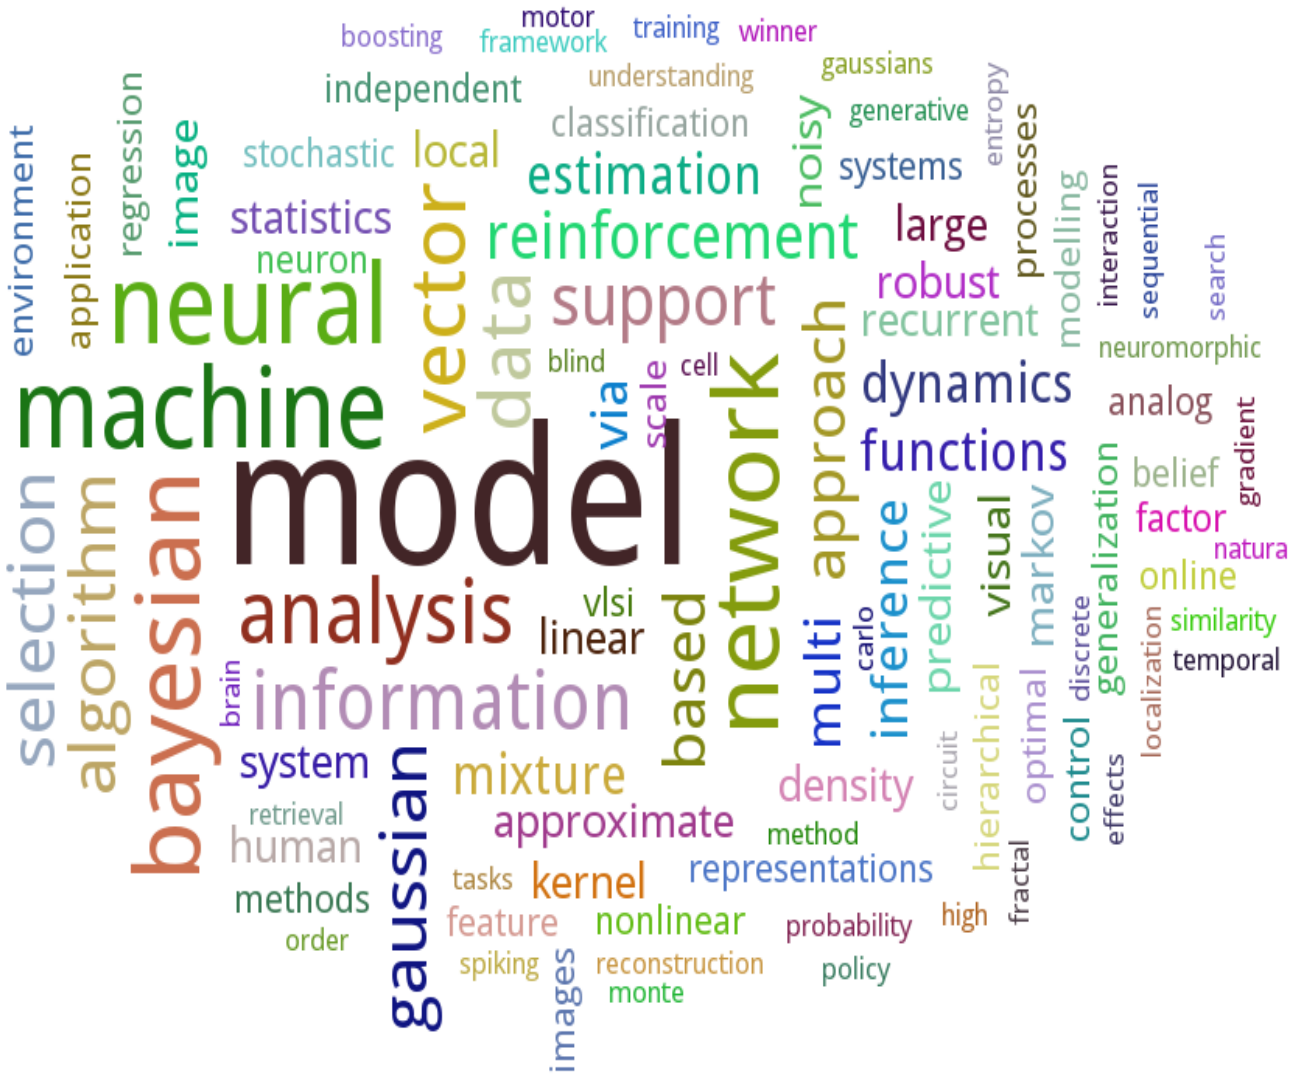

Advances in Neural Information Processing Systems 2000

NIPS Papers!

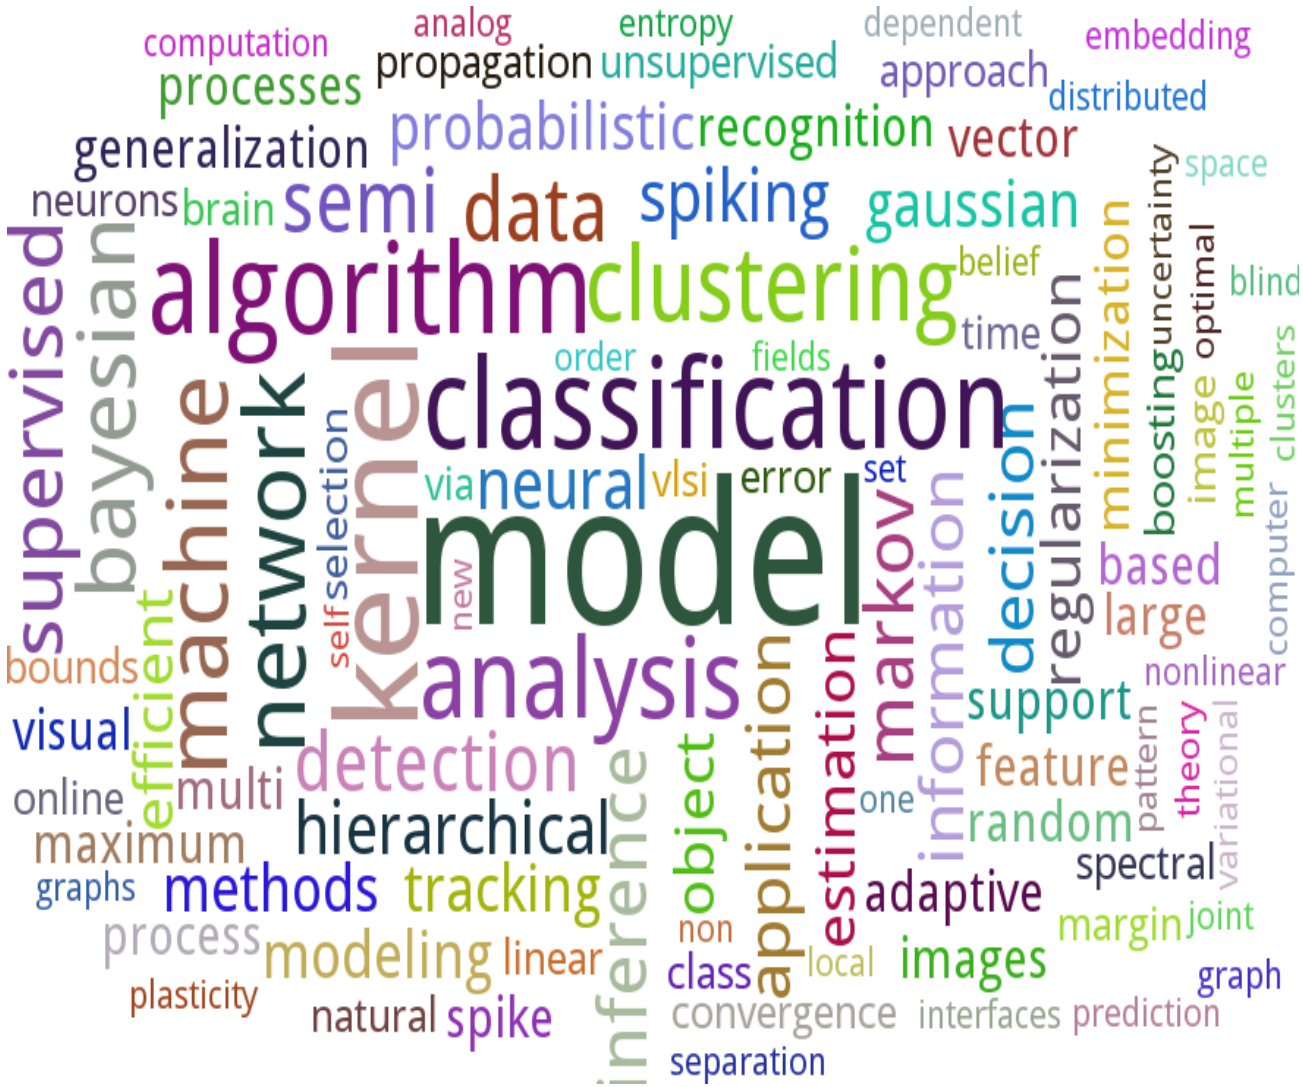

Advances in Neural Information Processing Systems 2005

NIPS Papers!

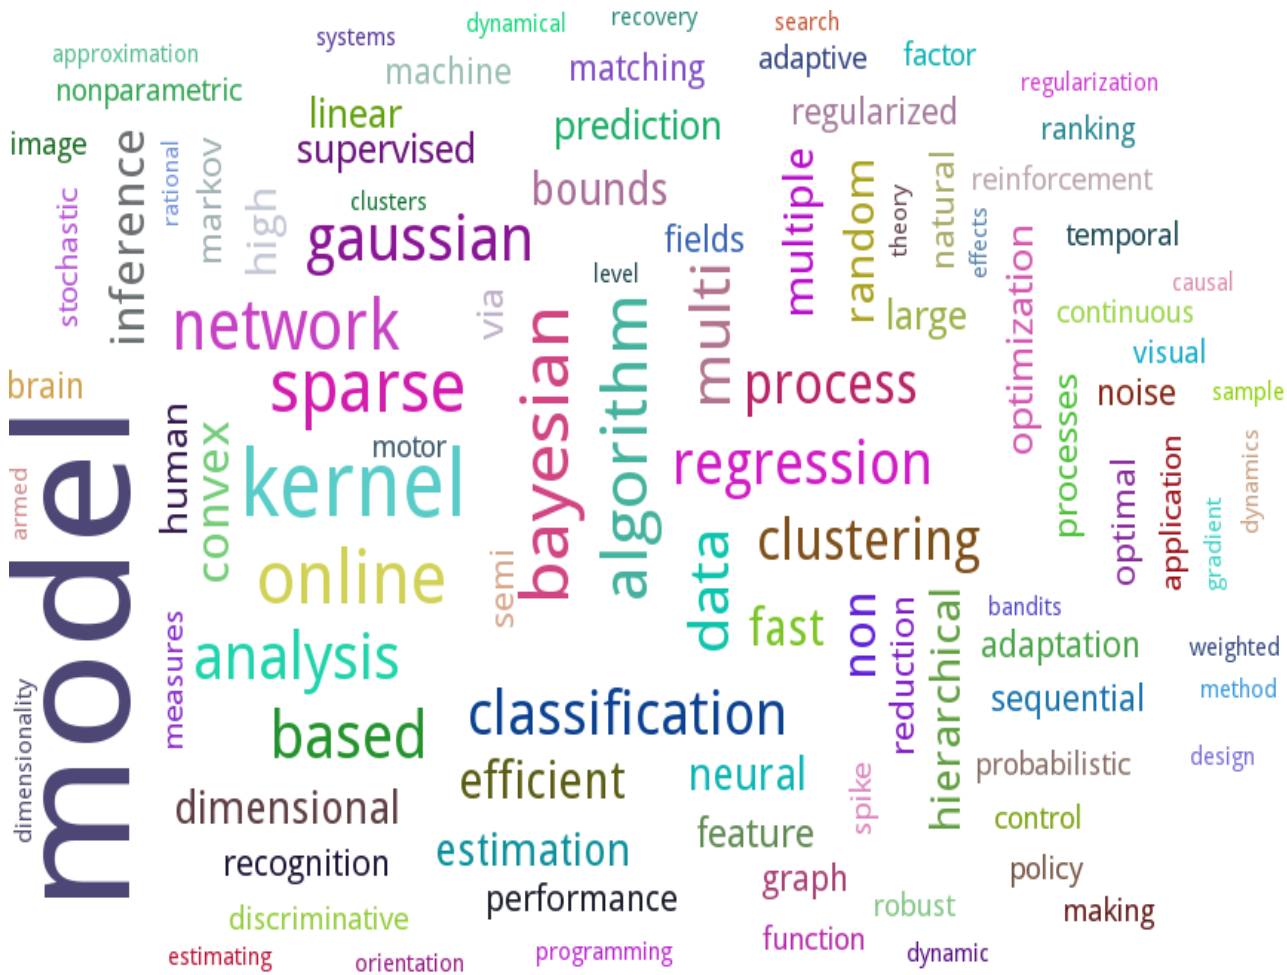

Advances in Neural Information Processing Systems 2009

NIPS Papers!

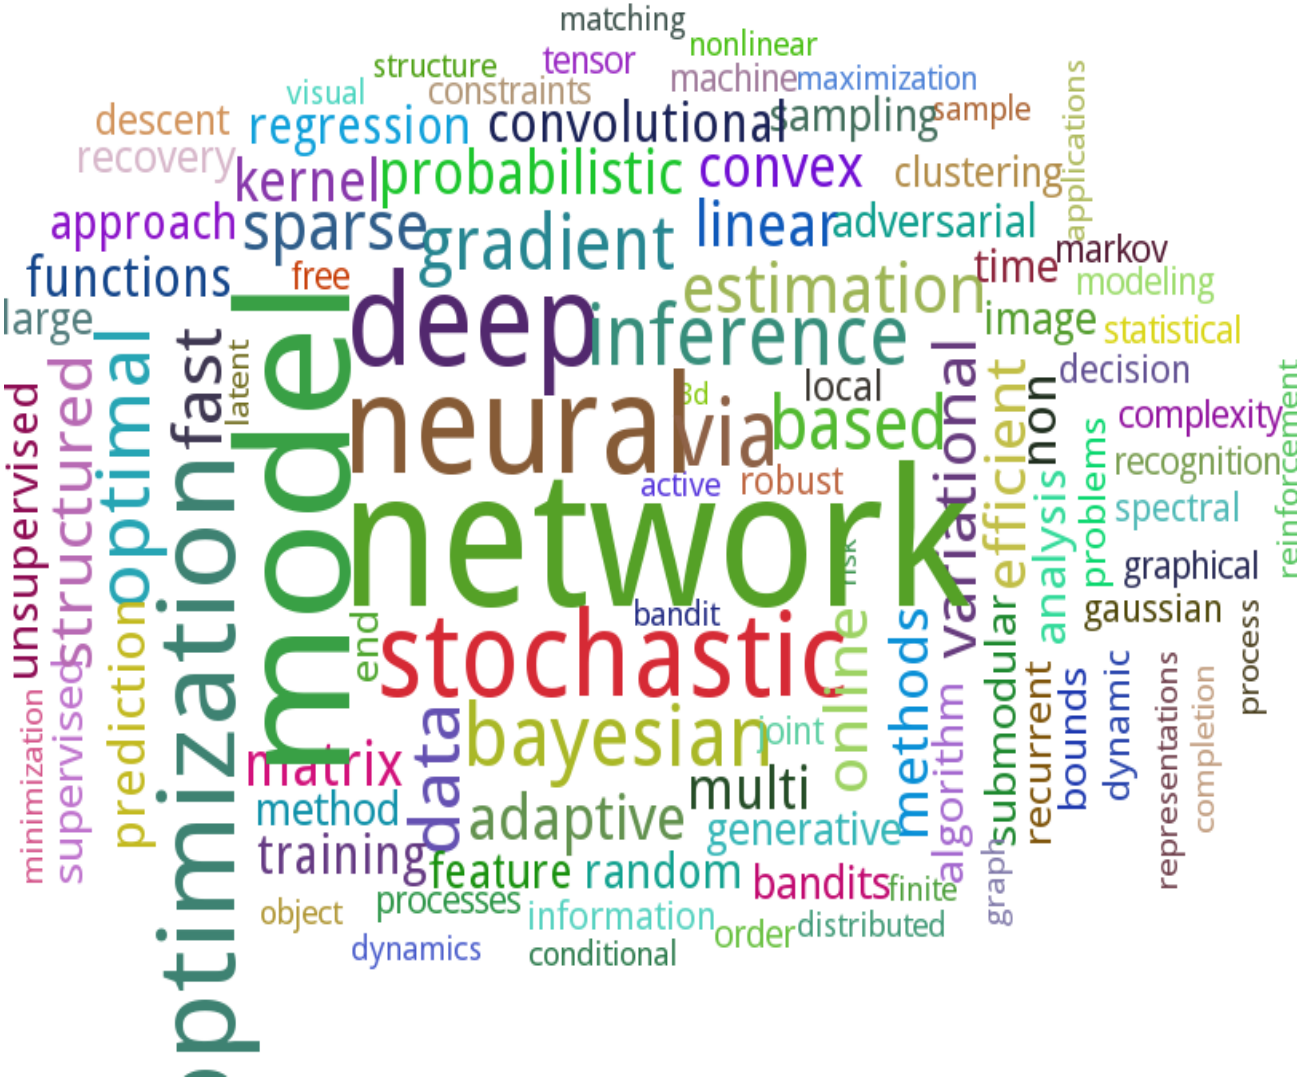

Advances in Neural Information Processing Systems 2016

NIPS Papers!

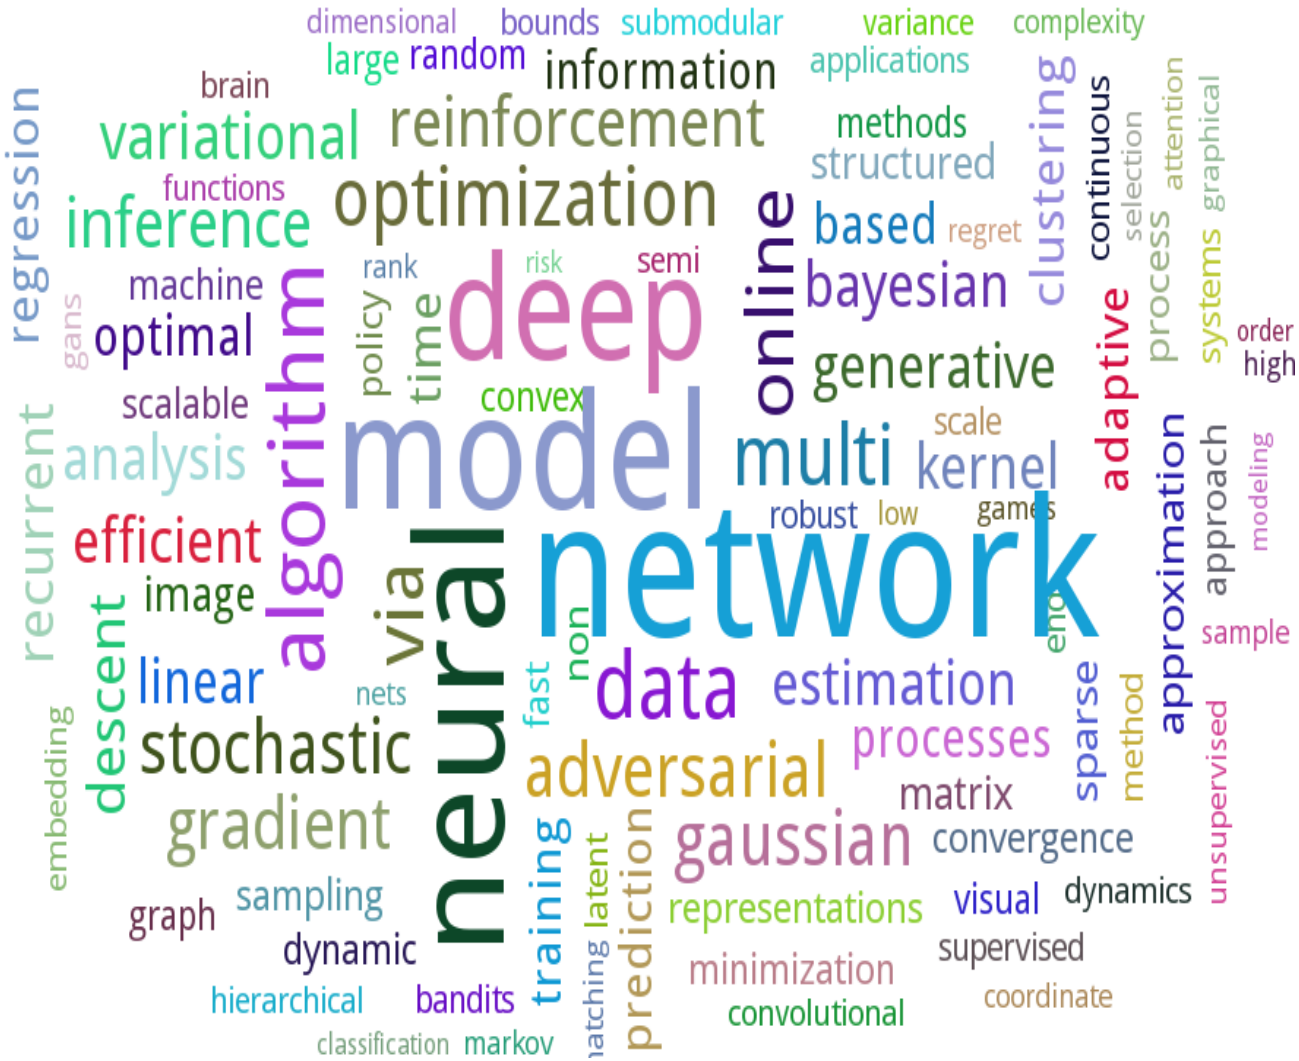

Advances in Neural Information Processing Systems 2017 [video]

Outline

Course Logistics

Some Machine Learning Applications

Application: Boston Housing Dataset

Numerical attributes

- Crime rate per capita
- Non-retail business fraction
- Nitric Oxide concentration
- Age of house
- Floor area
- Distance to city centre
- Number of rooms

Categorical attributes

- On the Charles river?
- Index of highway access (1-5)

Source: UCI repository

Predict house cost

Application: Object Detection and Localisation

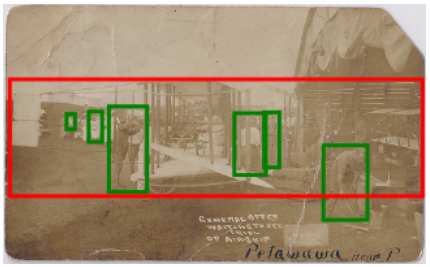




- 200-basic level categories
- Here: Six pictures containing airplanes and people
- Dataset contains over 400,000 images
- Imagenet competition (2010-)
- All recent successes through very deep neural networks!

Supervised Learning

Training data has inputs $\mathbf{x}$ (numerical, categorical) as well as outputs y (target)

Regression: When the output is real-valued, *e.g.,* housing price

Classification: Output is a category

- Binary classification: only two classes *e.g.,* spam
- Multi-class classification: several classes *e.g.,* object detection

Unsupervised Learning : Genetic Data of European Populations

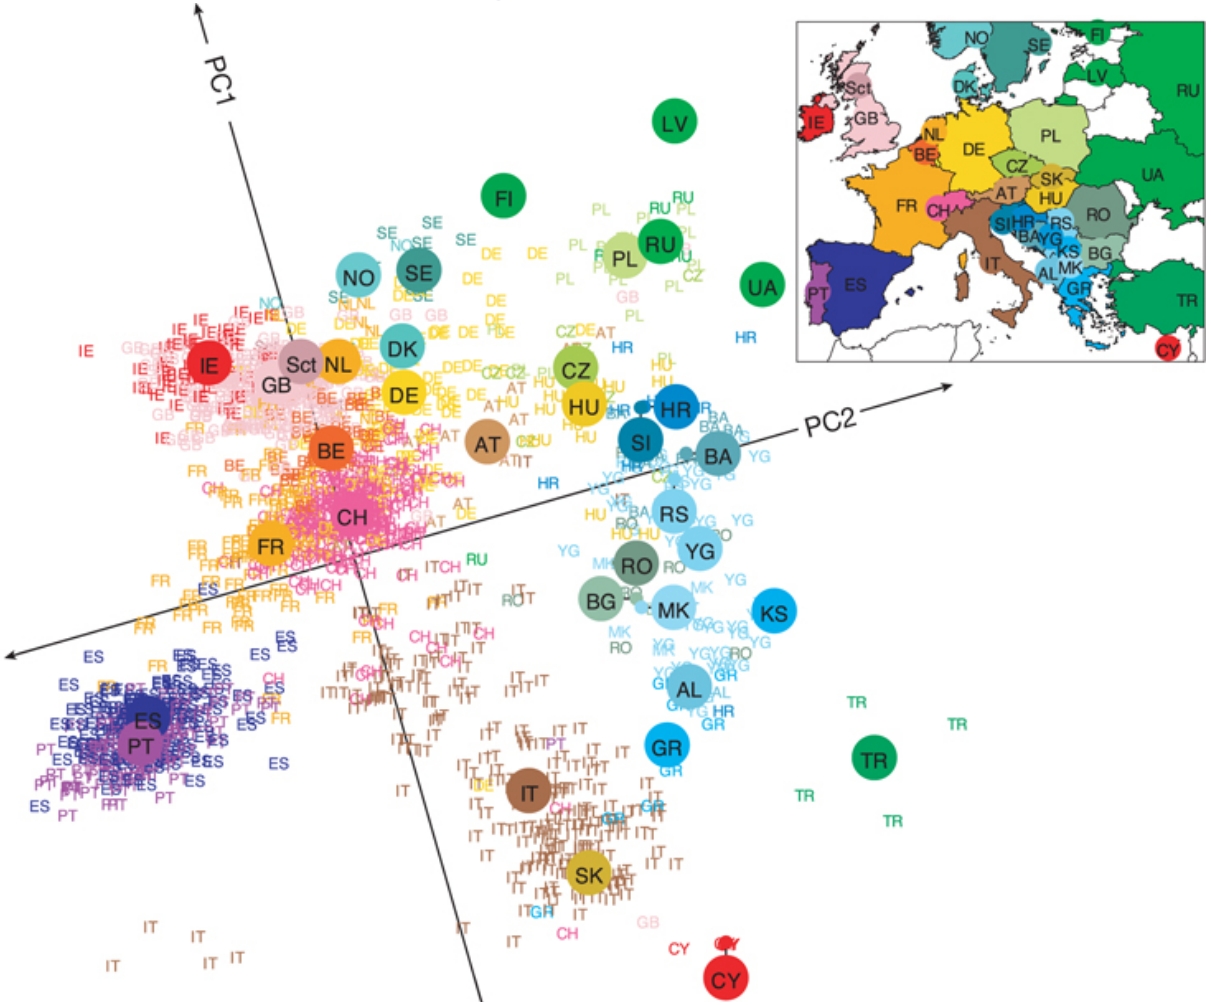

Source: Novembre *et al.*, Nature (2008)

Unsupervised Learning : Genetic Data of European Populations

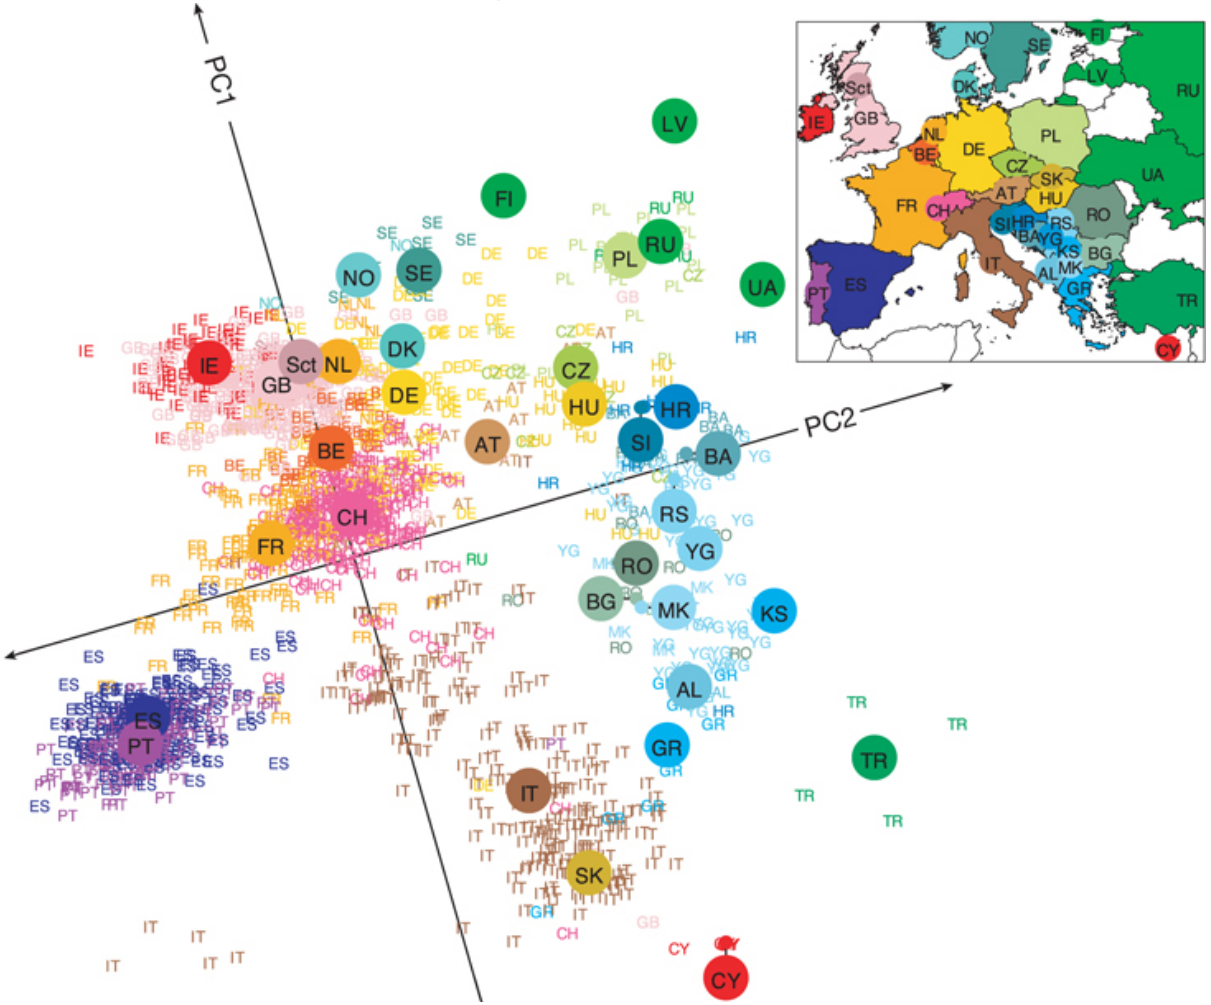

Experience (E)

Task (T)

Performance (P)

Source: Novembre *et al.*, Nature (2008)

Unsupervised Learning : Genetic Data of European Populations

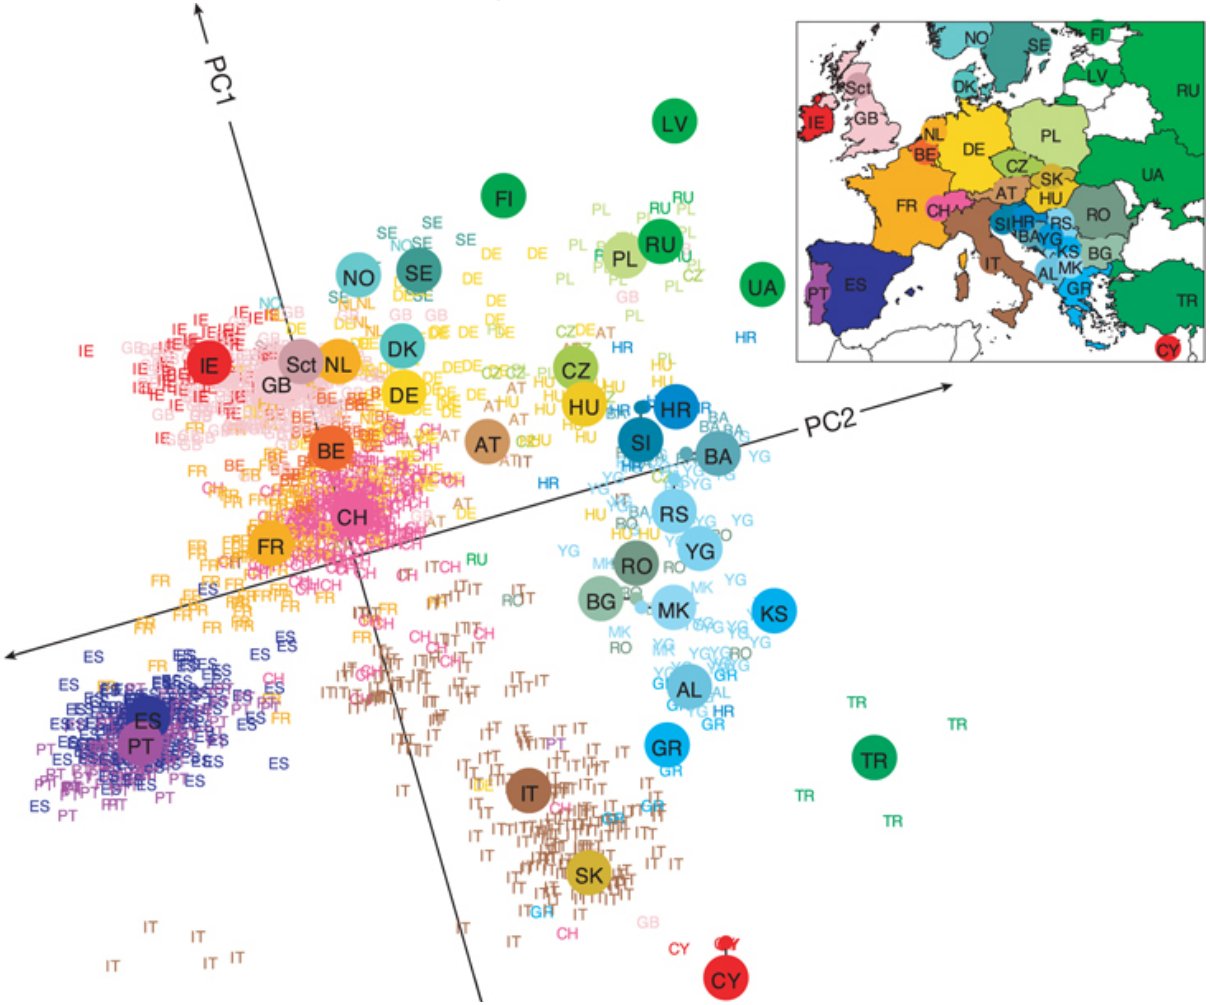

Experience (E)

Task (T)

Performance (P)

Source: Novembre *et al.*, Nature (2008)

Dimensionality reduction - Map high-dimensional data to low dimensions

Clustering - group together individuals with similar genomes

Unsupervised Learning : Group Similar News Articles

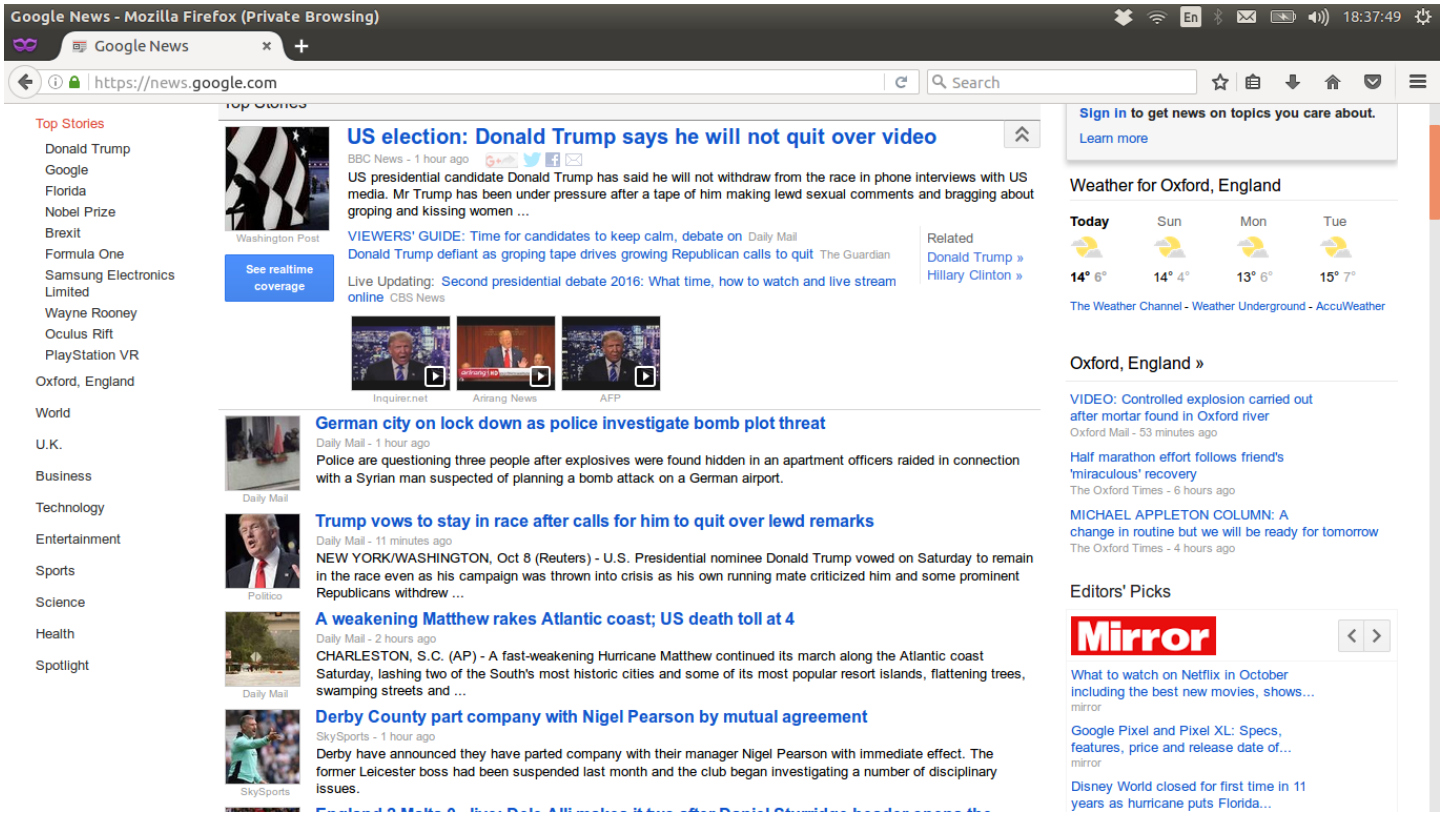

Group similar articles into categories such as politics, music, sport, *etc.*

In the dataset, there are no labels for the articles

Active and Semi-Supervised Learning

Active Learning

- Initially all data is unlabelled
- Learning algorithm can ask a human to label some data

Semi-supervised Learning

- Limited labelled data, lots of unlabelled data
- How to use the two together to improve learning?

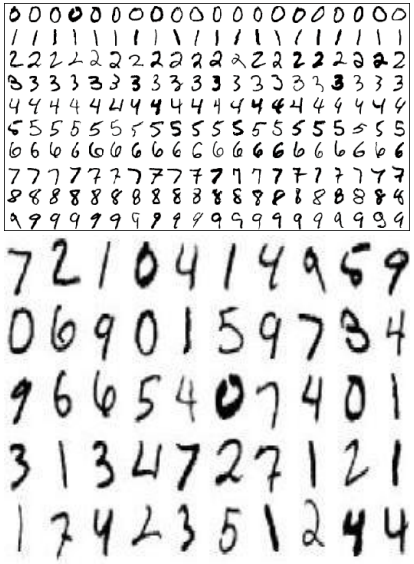

Collaborative Filtering : Recommender Systems

Movie / User	Alice	Bob	Charlie	Dean	Eve
The Shawshank Redemption	7	9	9	5	2
The Godfather	3	?	10	4	3
The Dark Knight	5	9	?	6	?
Pulp Fiction	?	5	?	?	10
Schindler's List	?	6	?	9	?

Netflix competition to predict user-ratings (2008-09)

Any individual user will not have used most products

Most products will have been use by some individual

Reinforcement Learning

- Automatic flying helicopter; self-driving cars
- Cannot conceivably program by hand
- Uncertain (stochastic) environment
- Must take sequential decisions
- Can define reward functions
- Fun: Playing Atari breakout! [video]