open talk

Machine Learning

...

opendot, Milano - 09.11.2016

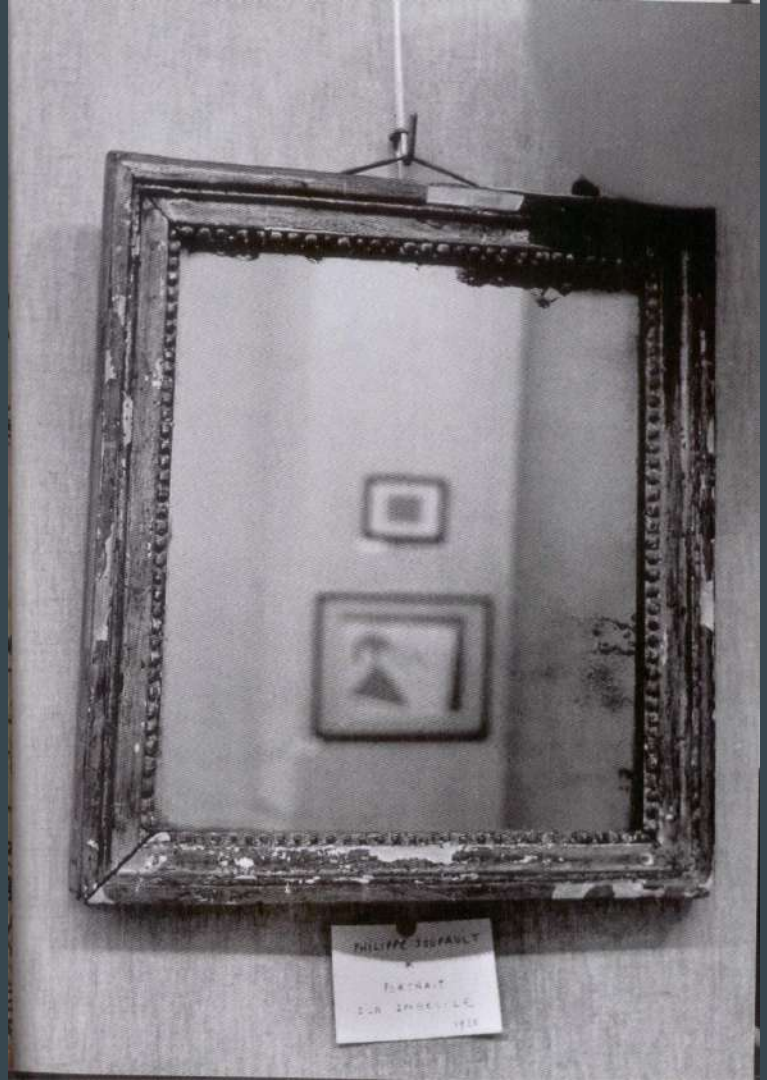

Philippe Soupault
Portrait of an Imbecile

framed mirror exhibited at the Salon Dada, June 1921.

“Man is a slow, sloppy, and brilliant thinker;
computers are fast, accurate, and stupid”

John Pfeiffer, *Fortune* (1961)

Intro

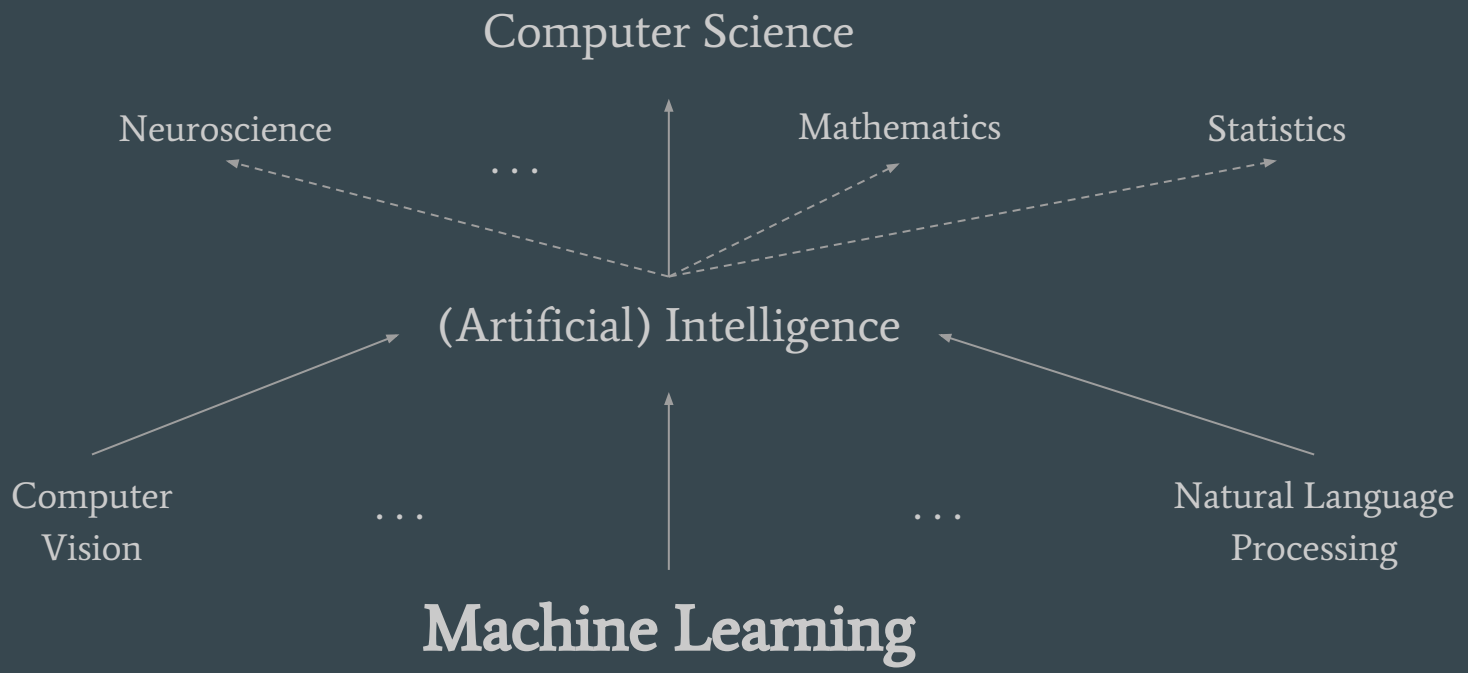
Computer Science
Neuroscience
. . .
Mathematics
Statistics
(Artificial) Intelligence
Computer Vision
. . .
. . .
Natural Language Processing
Machine Learning

Intro

Definitions

- “[The automation of] activities that we associate with **human** thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)
- “The art of creating machines that perform functions that require intelligence when performed by **people**.” (Kurzweil, 1990)
- “The study of how to make computers do things at which, at the moment, **people** are better.” (Rich and Knight, 1991)

Artificial Intelligence

Artificial Intelligence

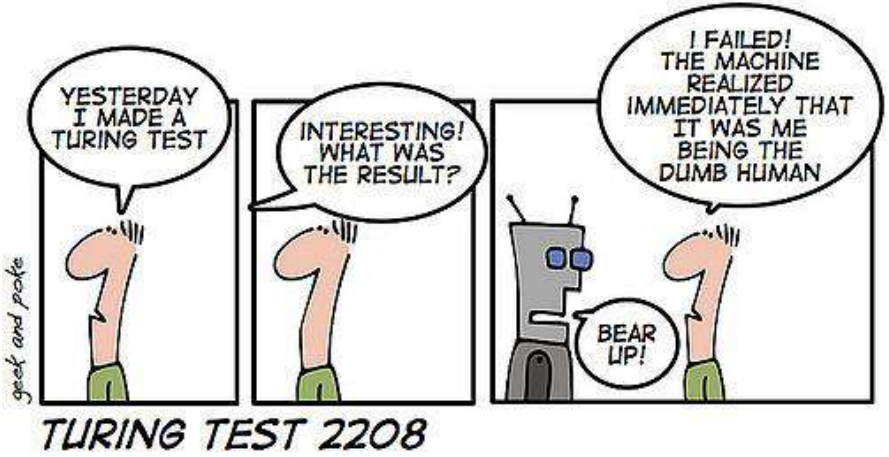

Machine Learning - What

- “Field of study that gives computers the ability to learn without being explicitly programmed” (Samuel, **1959**)

- “A computer program is said to learn from experience E with respect to some class of tasks T and a performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.” (Mitchell, 1997)

When the system needs to:

- **adapt** to the working environment
- **enhance** its performance regarding a specific task
- **discover** new information from empirical data
- **acquire** new computational abilities

Machine Learning - When

Why not using a classical algorithmic approach?

- impossible or, very hard, to **formalize** the problem
- presence of **noise** and/or uncertainty
- high **complexity** in formulating a solution
- lack of **knowledge** "filled in" with respect to the problem

Machine Learning - Why

Spam email detection

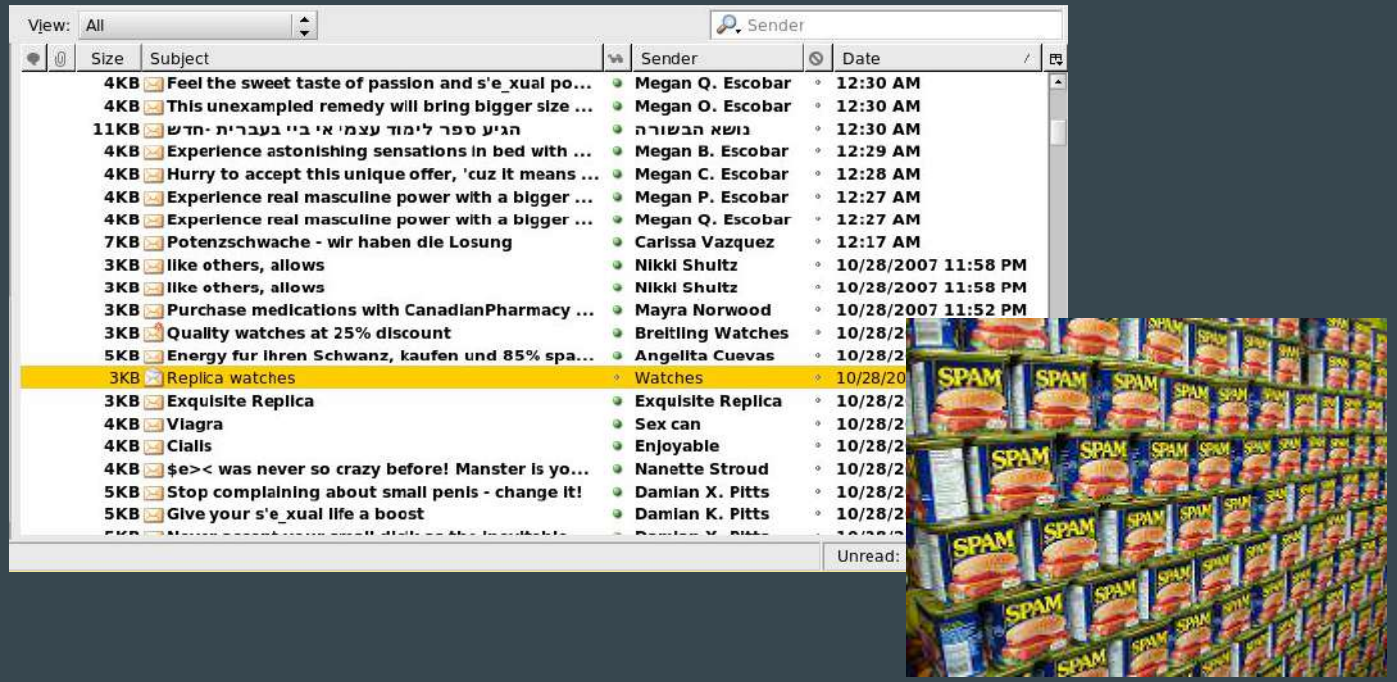

Machine Learning - Examples

Google Translate (Machine translation)

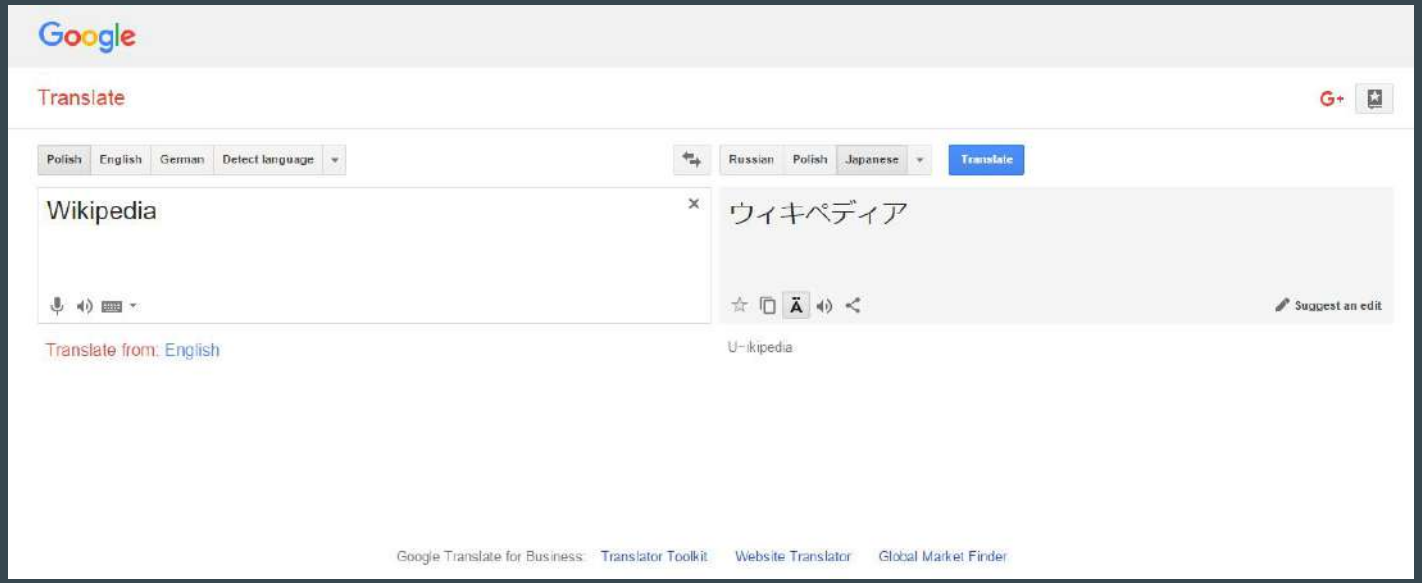

Machine Learning - Examples

Amazon Recommendations

Machine Learning - Examples

Shazam

Machine Learning - Examples

Facebook Tag Suggest

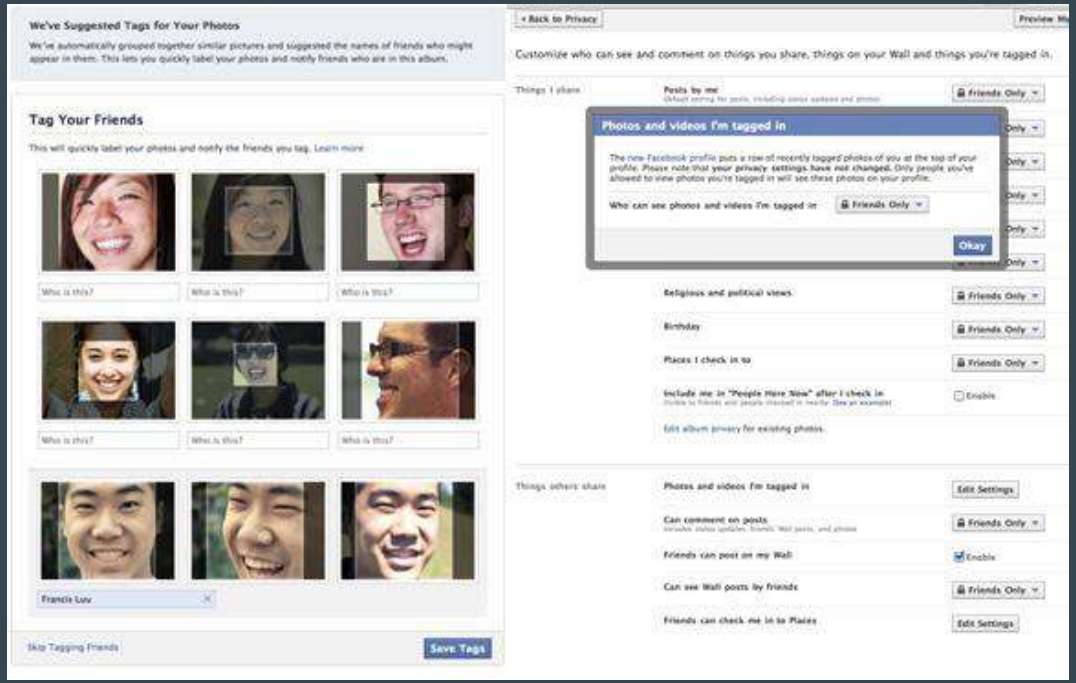

Machine Learning - Examples

Google Home / Now

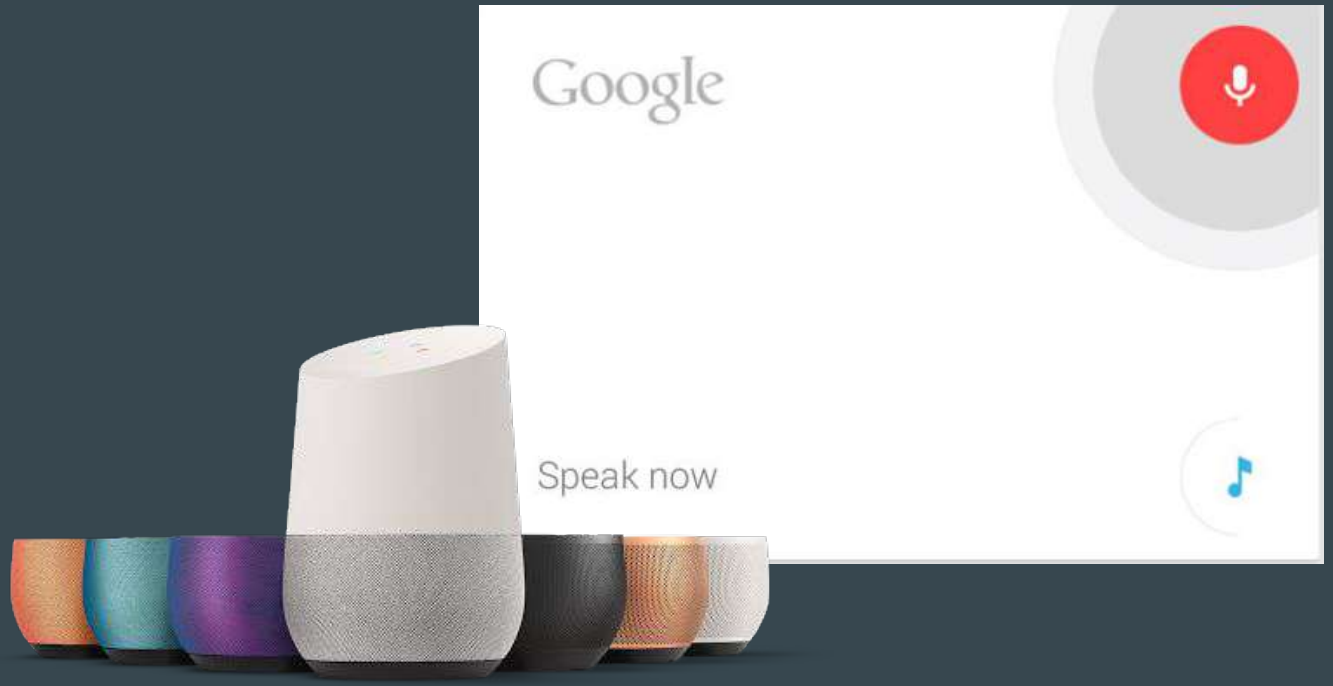

Machine Learning - Examples

Facebook Automatic Alt Text

Machine Learning - Examples

Facebook Automatic Alt Text

Machine Learning - Examples

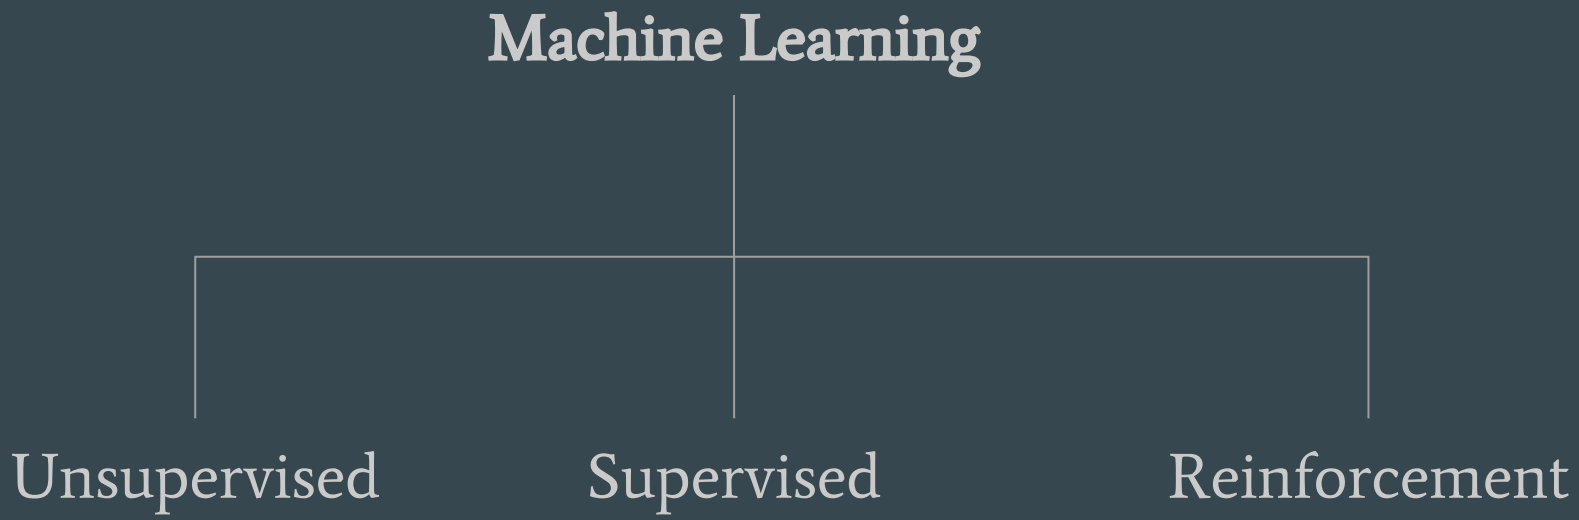

Machine Learning - Paradigms

Unsupervised learning is when you only have input data and **no corresponding labels**.

The goal for unsupervised learning is to model the **underlying structure** or distribution in the data in order to learn more about the data.

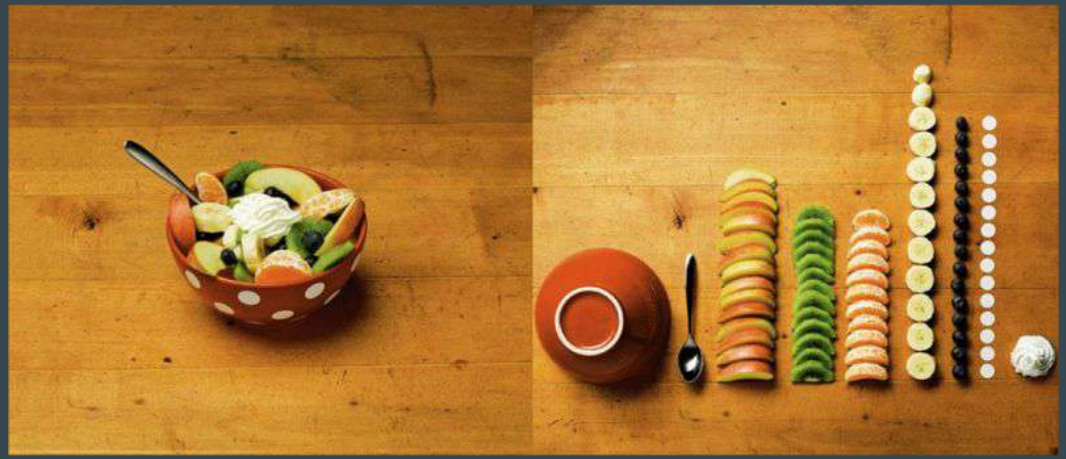

Machine Learning - Unsupervised

Clustering

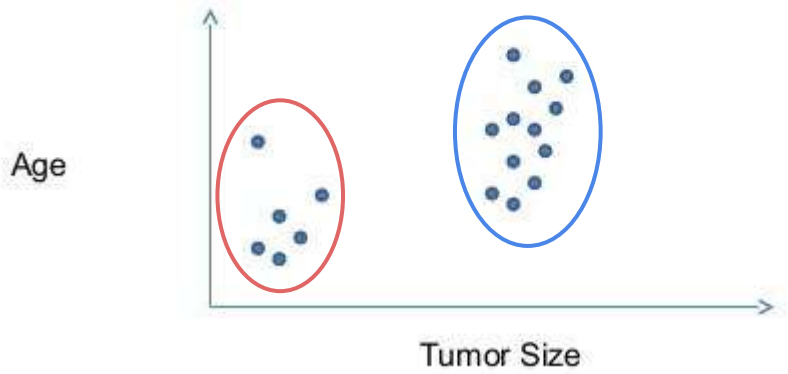

Machine Learning - Unsupervised

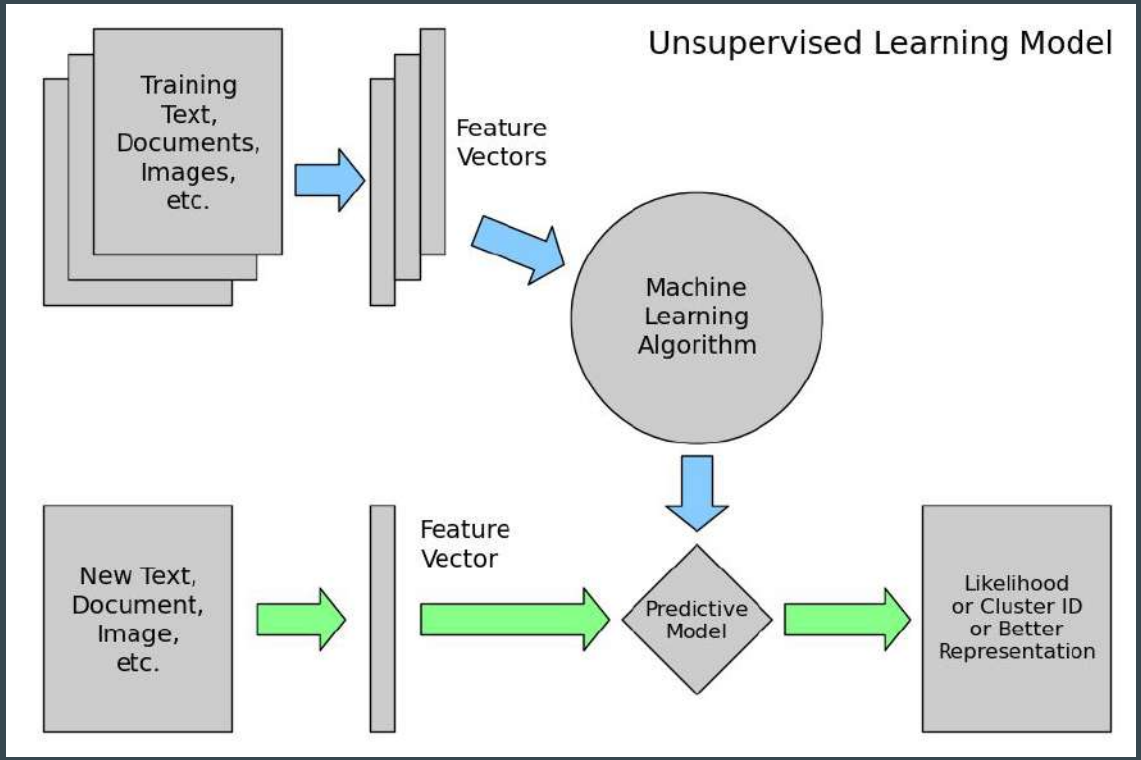

Machine Learning - Unsupervised

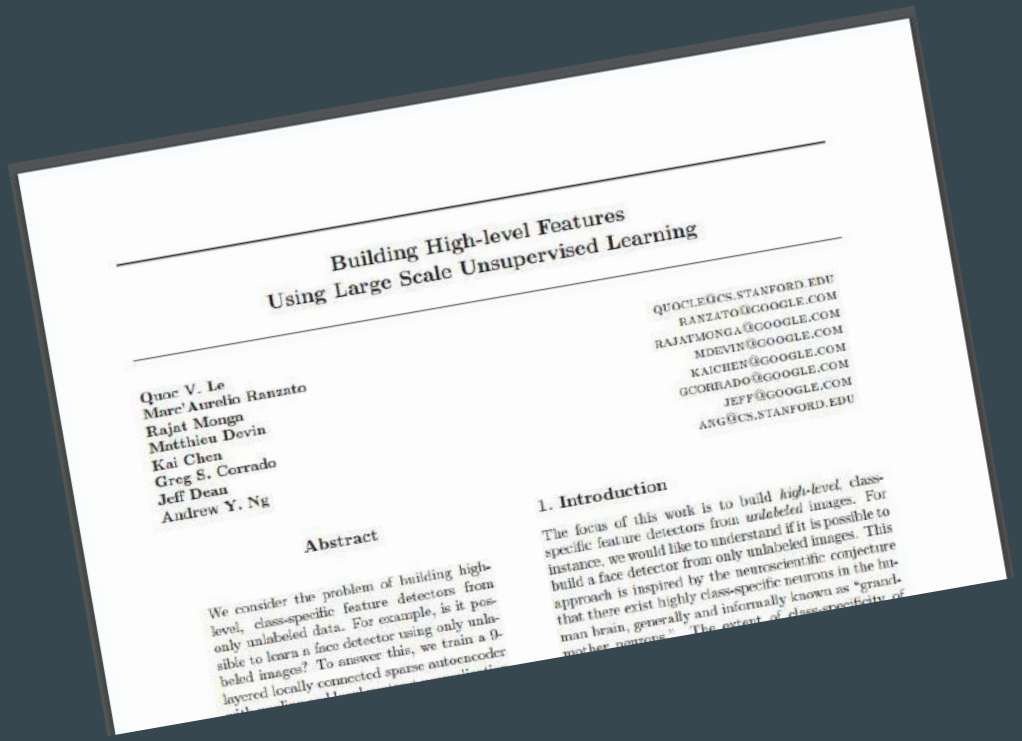

Building High-level Features
Using Large Scale Unsupervised Learning

Quoc V. Le
Marc'Aurelio Ranzato
Rajat Monga
Matthieu Devin
Kai Chen
Greg S. Corrado
Jeff Dean
Andrew Y. Ng

QUOCLE@CS.STANFORD.EDU
RANZATO@GOOGLE.COM
RAJATMONGA@GOOGLE.COM
MDEVIN@GOOGLE.COM
KAICHEN@GOOGLE.COM
GCORRADO@GOOGLE.COM
JEFF@GOOGLE.COM
ANG@CS.STANFORD.EDU

Abstract

We consider the problem of building high-level, class-specific feature detectors from only unlabeled data. For example, is it possible to learn a face detector using only unlabeled images? To answer this, we train a 9-layered locally connected sparse autoencoder

1. Introduction

The focus of this work is to build *high-level*, class-specific feature detectors from *unlabeled* images. For instance, we would like to understand if it is possible to build a face detector from only unlabeled images. This approach is inspired by the neuroscientific conjecture that there exist highly class-specific neurons in the human brain, generally and informally known as "grand-

Dataset of **10 million 200x200 unlabeled images** from YouTube/Web

Train on **1000 machines** (16000 cores) for **1 week**

1.15 billion parameters

Le, Quoc V. "Building high-level features using large scale unsupervised learning." 2013 *IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013.

Machine Learning - Unsupervised - Case history

Human faces

Cat faces

Machine Learning - Unsupervised - Case history

Supervised learning is when you have **input variables** (X) and an **output labels** (Y) and you use an algorithm to learn the **mapping function** from the input to the output.

$$Y = f(X)$$

The goal is to approximate the mapping function so well that when you have new input data (x) that you can **predict** the output variables (Y) for that data.

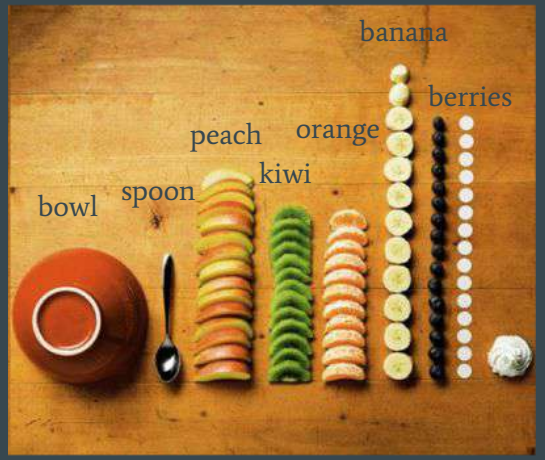

X =  ?

Y= banana!

Machine Learning - Supervised

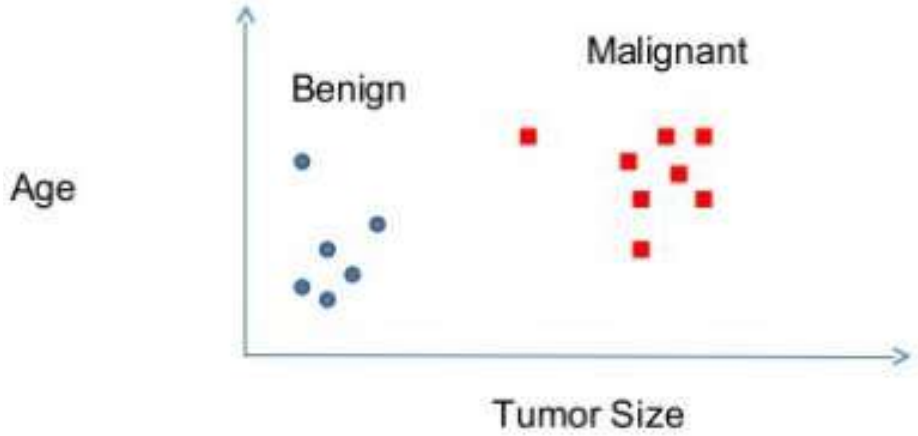

Machine Learning - Supervised

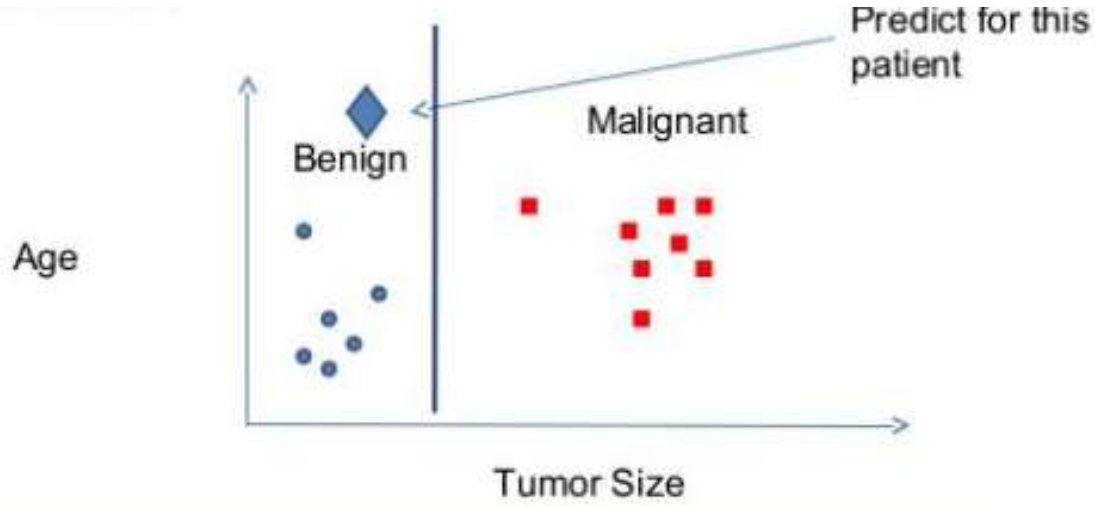

Machine Learning - Supervised

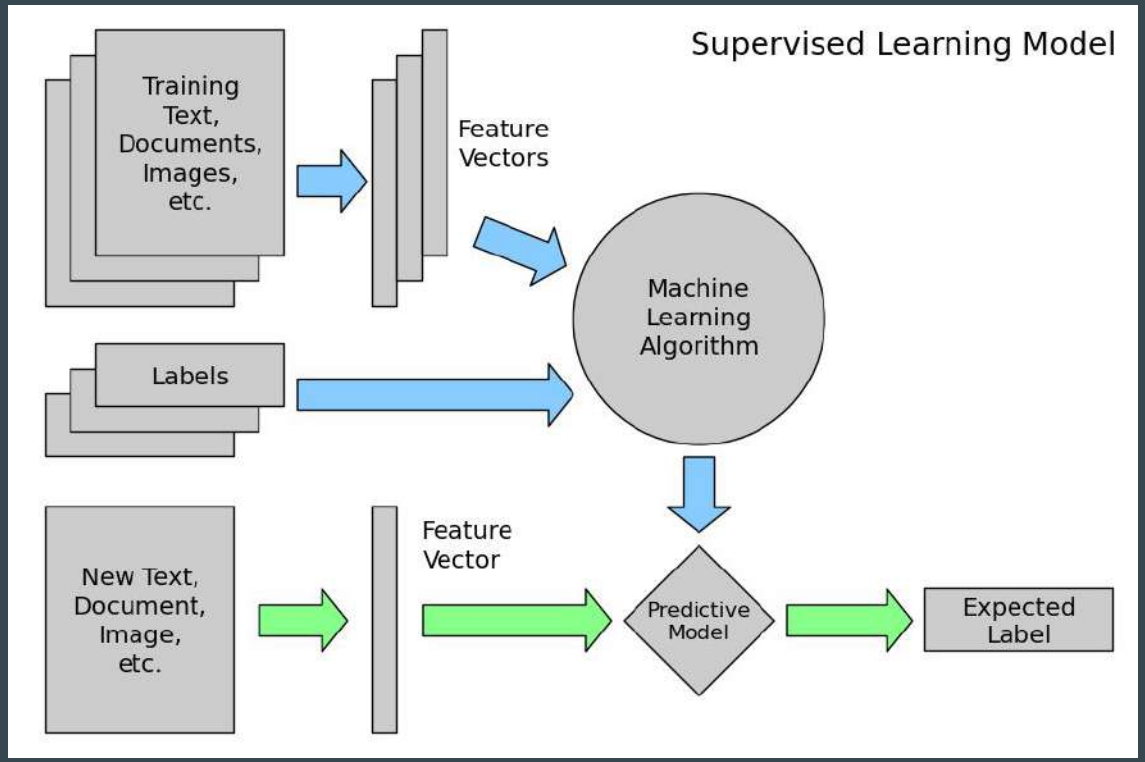

Machine Learning - Supervised

Affective Computing (AC) is an interdisciplinary field spanning computer science, psychology, and cognitive science.

“AC is computing that relates to, arises from, or deliberately influences emotion or other affective phenomena” (Picard, 1997)

The machine should interpret the emotional state of humans and adapt its behaviour to them, giving an appropriate response for those emotions.

Affective computing

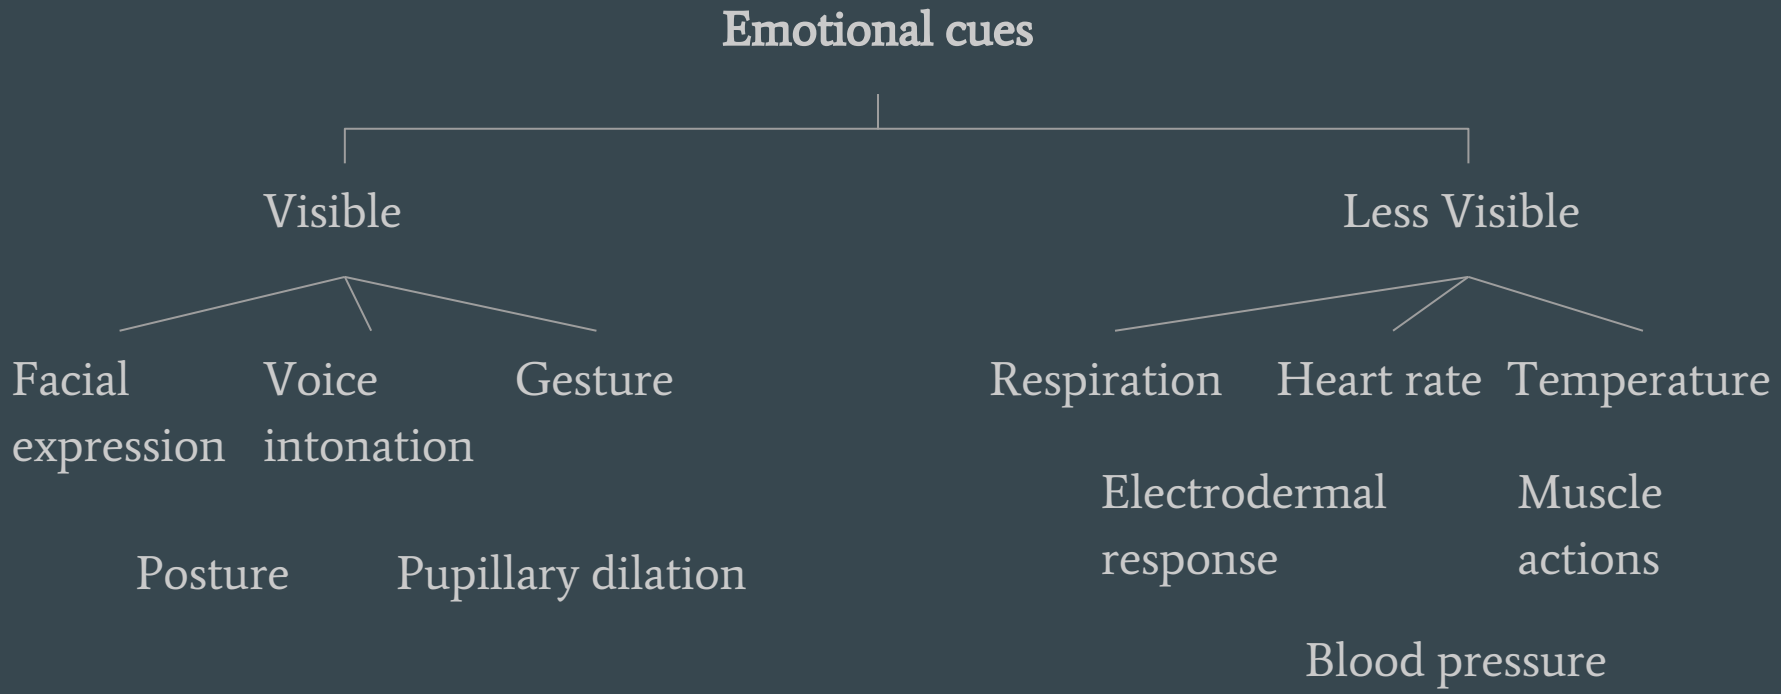

Affective computing

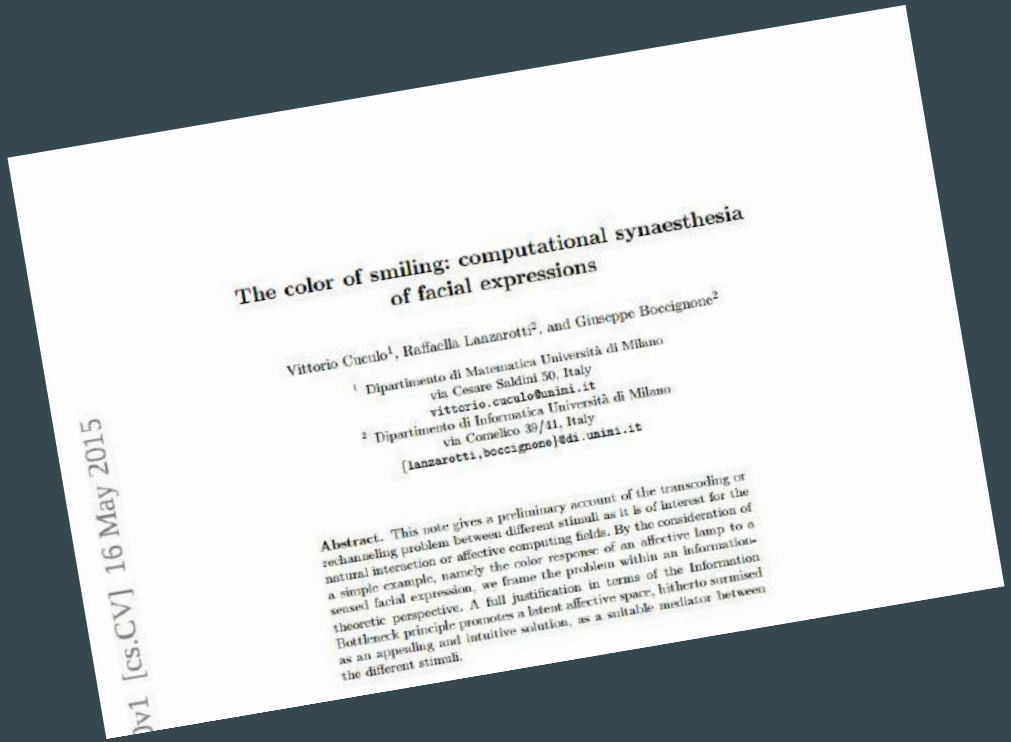

The color of smiling: computational synaesthesia of facial expressions

Vittorio Cuculo[1], Raffaella Lanzarotti[2], and Giuseppe Boccignone[2]

[1] Dipartimento di Matematica Università di Milano
via Cesare Saldini 50, Italy
vittorio.cuculo@unimi.it

[2] Dipartimento di Informatica Università di Milano
via Comelico 39/41, Italy
{lanzarotti,boccignone}@di.unimi.it

Abstract. This note gives a preliminary account of the transcoding or rechanneling problem between different stimuli as it is of interest for the natural interaction or affective computing fields. By the consideration of a simple example, namely the color response of an affective lamp to a sensed facial expression, we frame the problem within an information-theoretic perspective. A full justification in terms of the Information Bottleneck principle promotes a latent affective space, hitherto surmised as an appealing and intuitive solution, as a suitable mediator between the different stimuli.

Machine Learning - Supervised - Affective computing X Y

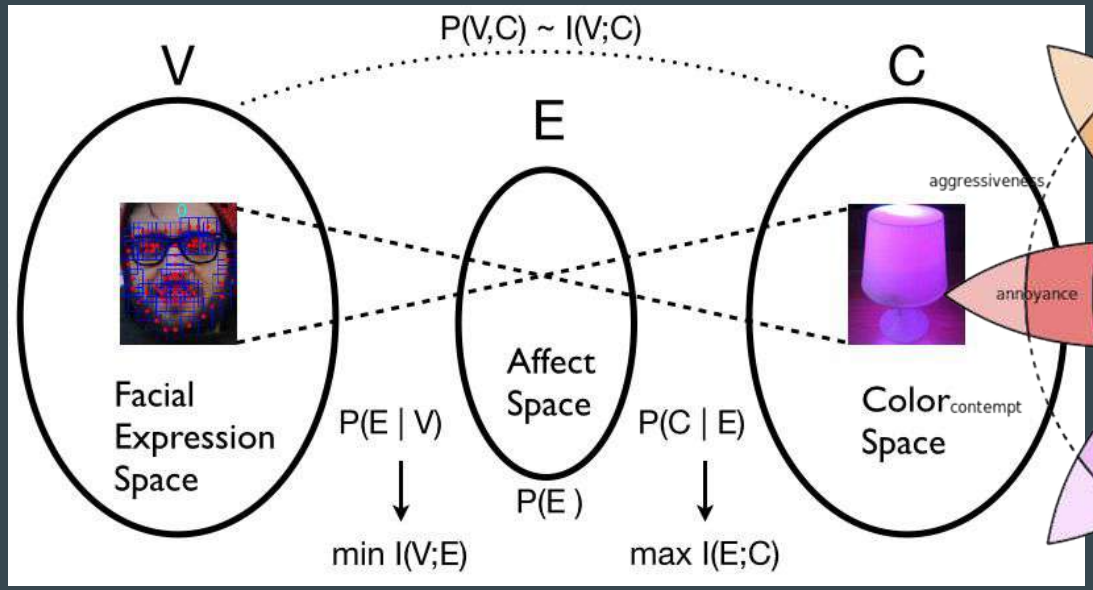

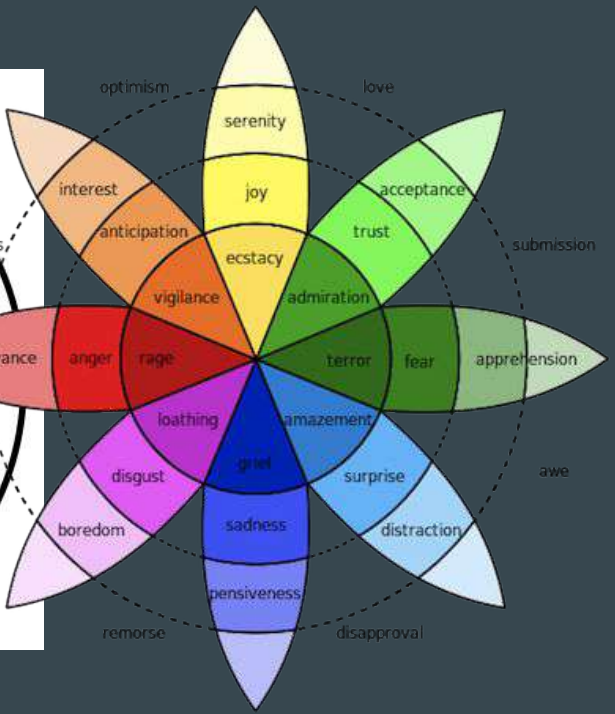

Machine Learning - Supervised - Affective computing

Reinforcement learning is learning by interacting with an environment.

An RL agent learns from the consequences of its actions, rather than from being explicitly taught and it selects its actions on basis of its past experiences (**exploitation**) and also by new choices (**exploration**), which is essentially trial and error learning.

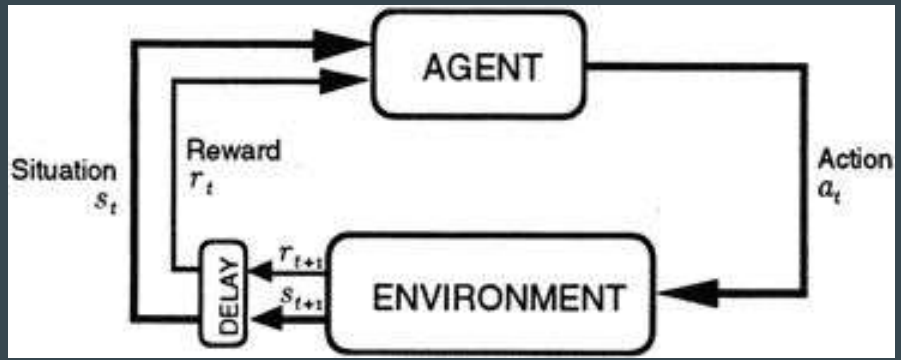

- a set of environment states **S**
- a set of actions **A**
- rules of transitioning between states
- rules that determine the scalar immediate reward of a transition
- rules that describe what the agent observes.

Machine Learning - Reinforcement

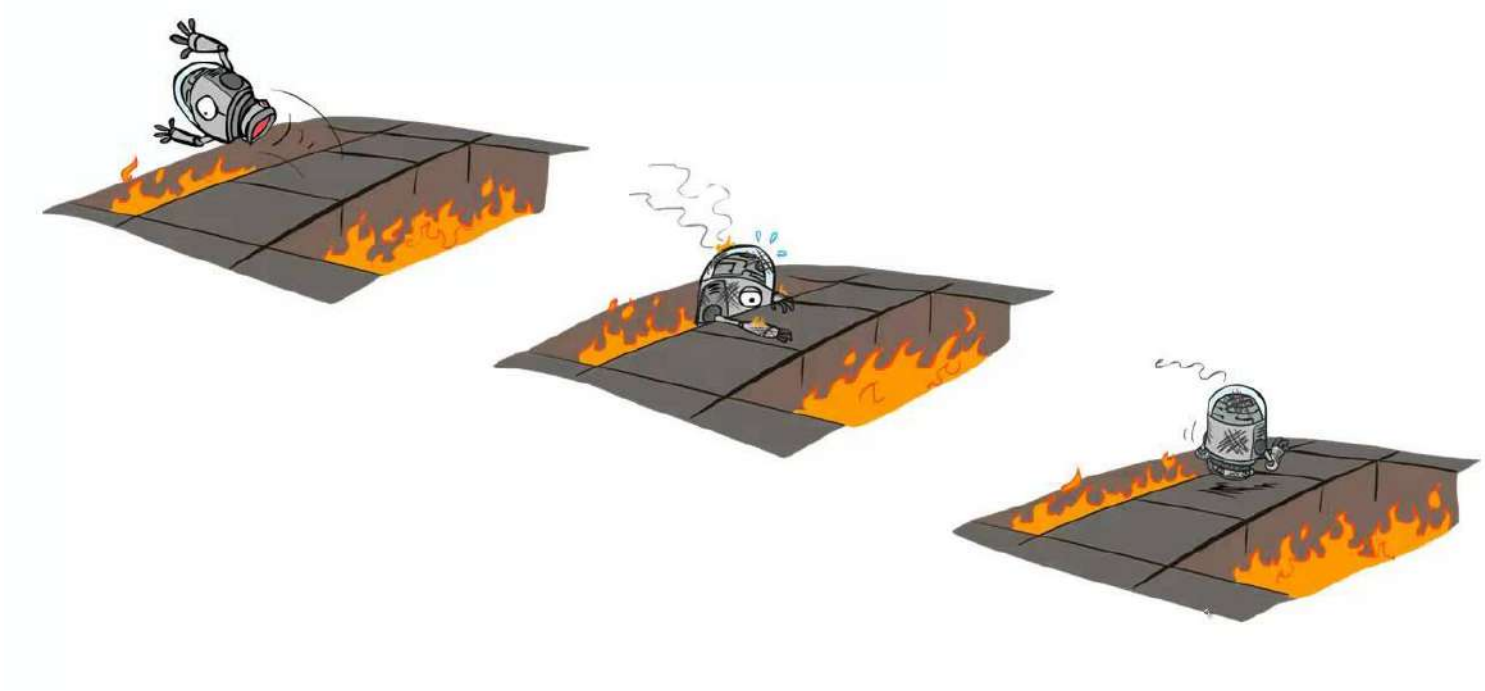

Machine Learning - Reinforcement

Machine Learning - Reinforced - Autonomous cars

Thanks!

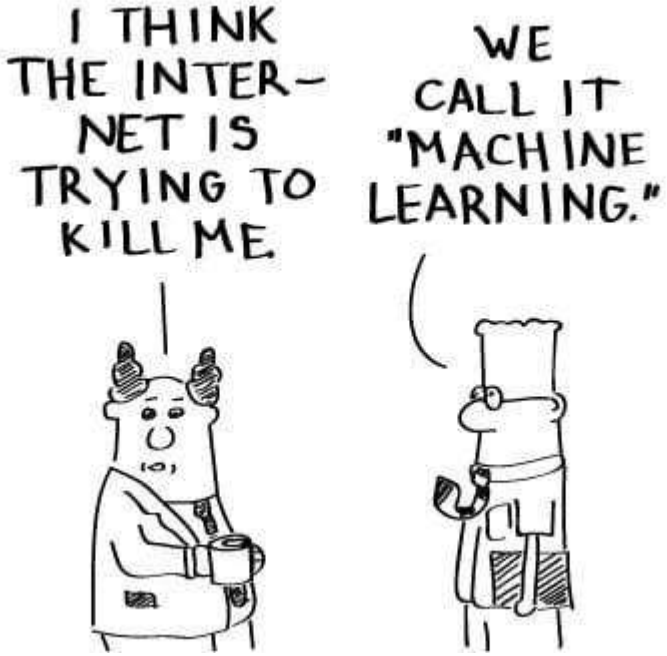

Vittorio Cuculo
tw: @vcuculo
mail: me@vcuculo.com