Tópicos Especiais em IA

Inês Dutra
Vítor Santos Costa

Glory Days at CMU

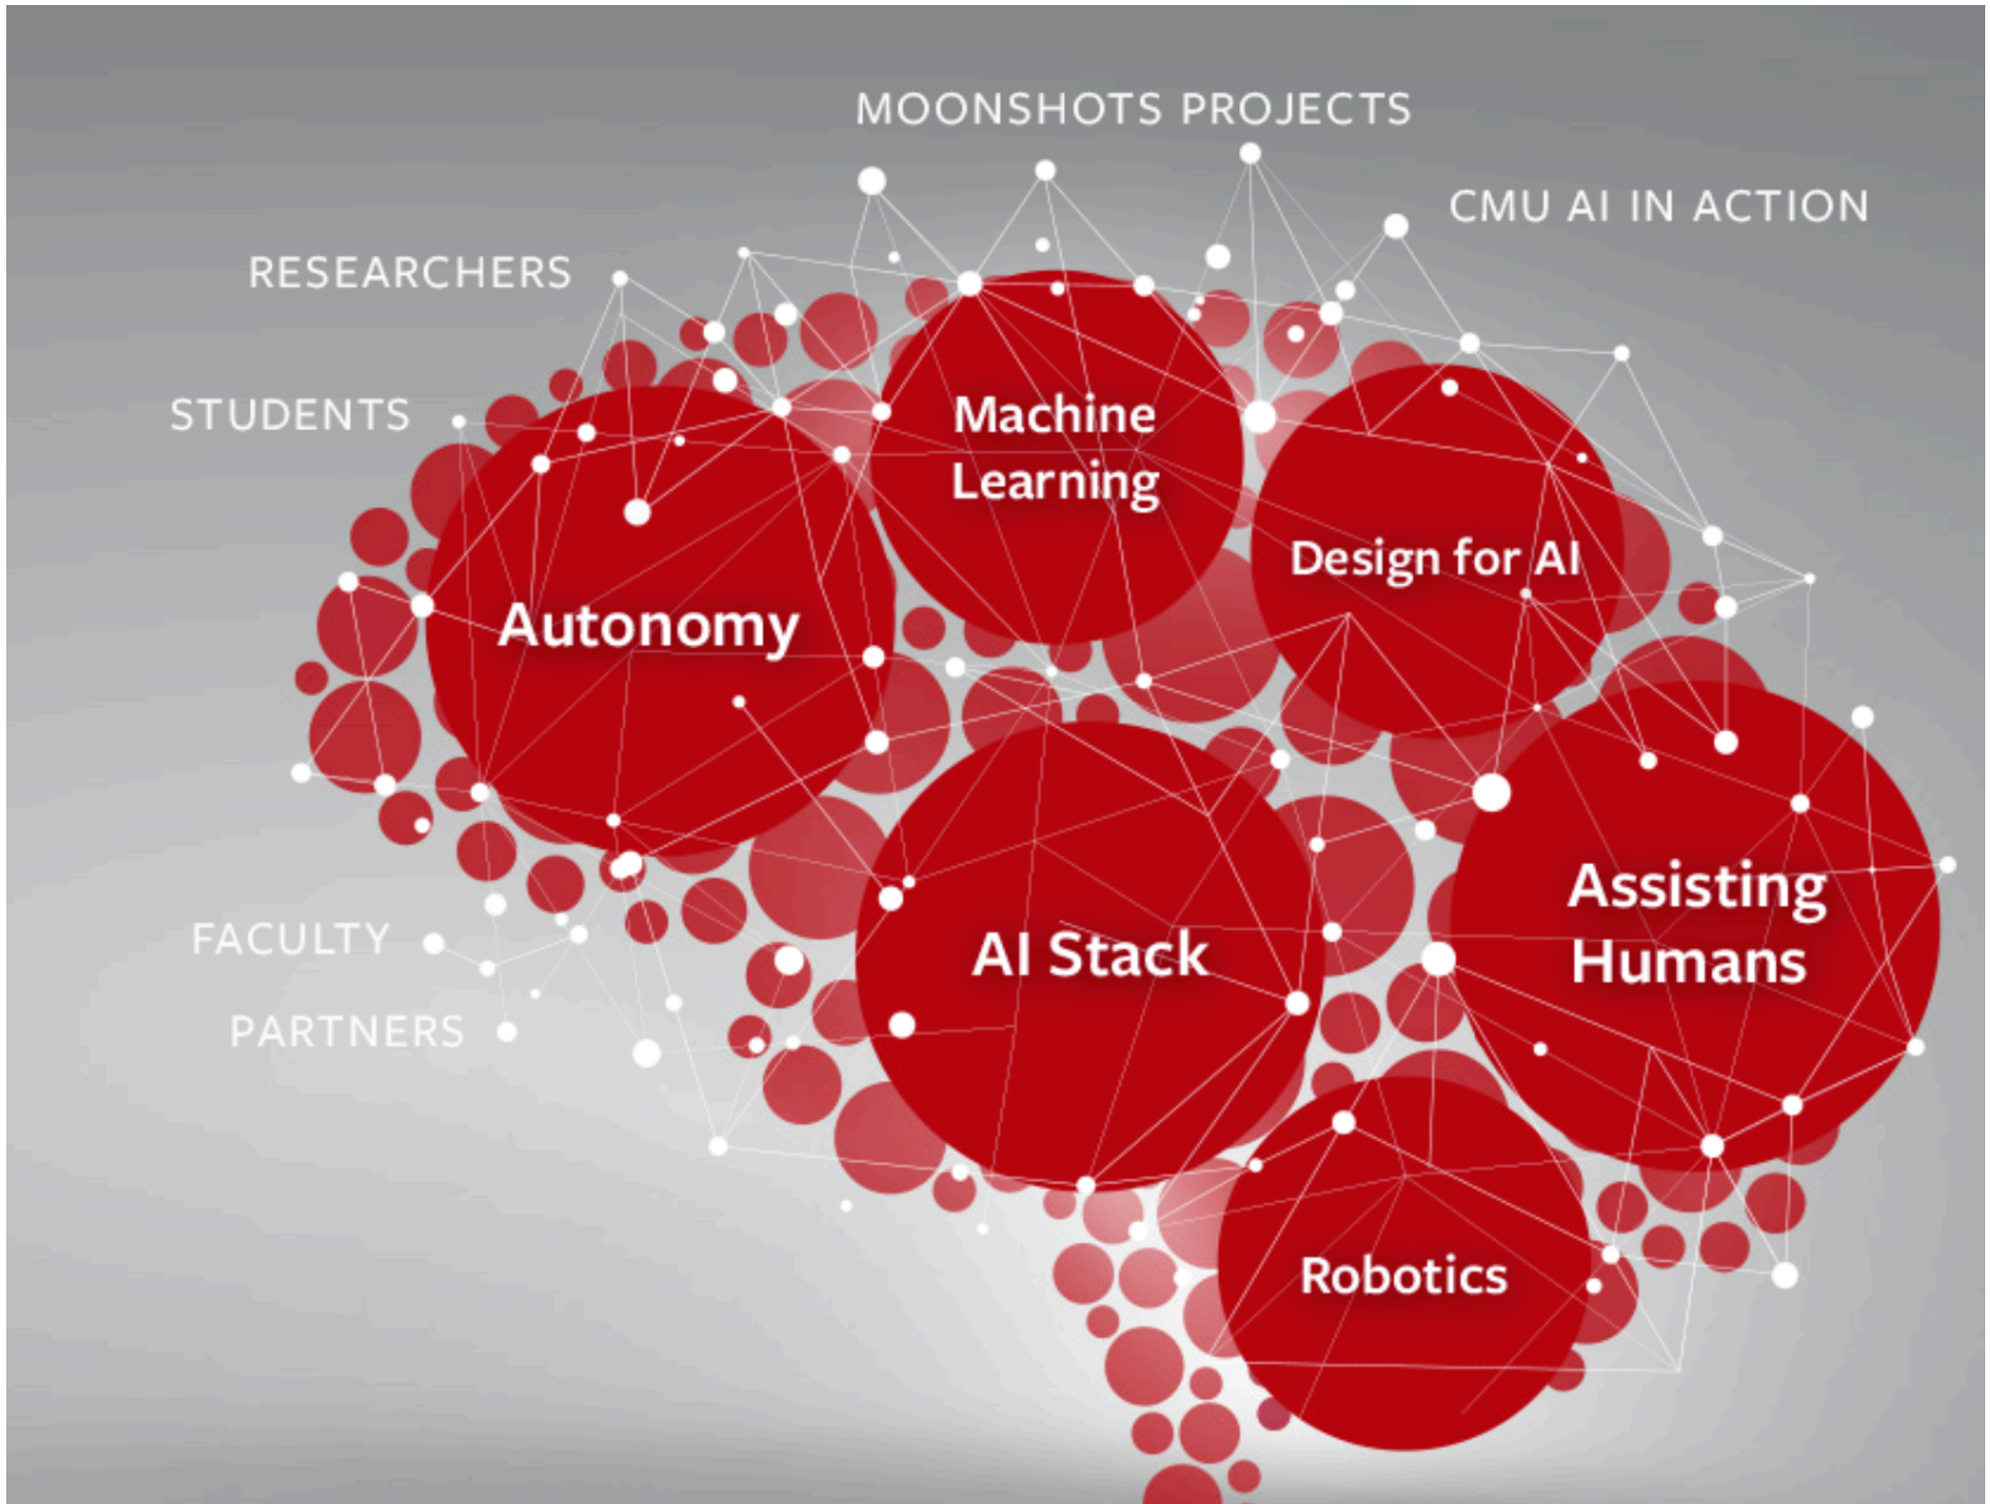

AAAI 2020

Vision: Tracking and Detection

ML: RL and Multiagent RL

NLP: Relational Learning

Application: Transportation

NLP: Generation

Application: Medical Imaging

Application: Web Search, Ranking and Recommendation

NLP: Machine Comprehension and Q&A

NLP: Knowledge Graph

NLP: Semantics and Summarization

Probabilistic Graphical Models and Probabilistic Methods

NLP: Semantics and Summarization

RL and Muliagent RL

Vision: Image Retrieval, Ranking, Recognition

ML: Neural Nets Theory, Models and Algorithm

: Constraint Satisfaction and Optimization

Planning and Scheduling

Knowledge Representation and Reasoning

Logic and KR: Inference and Reasoning

The ROBOTS are coming!

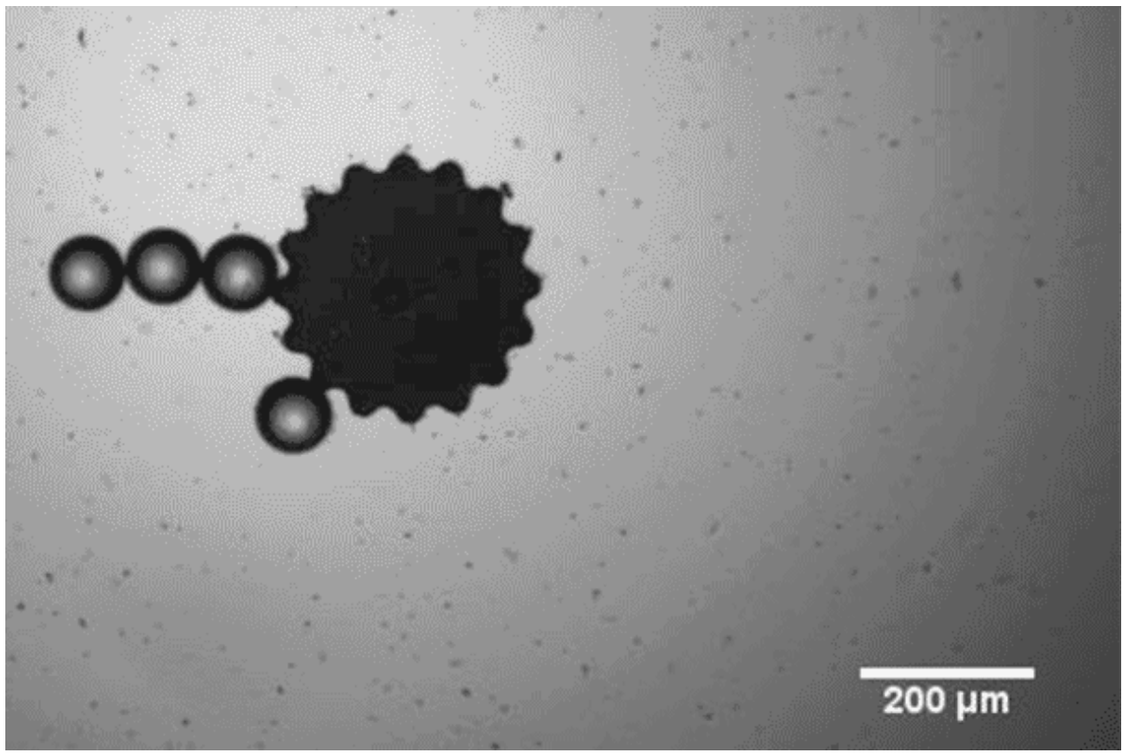

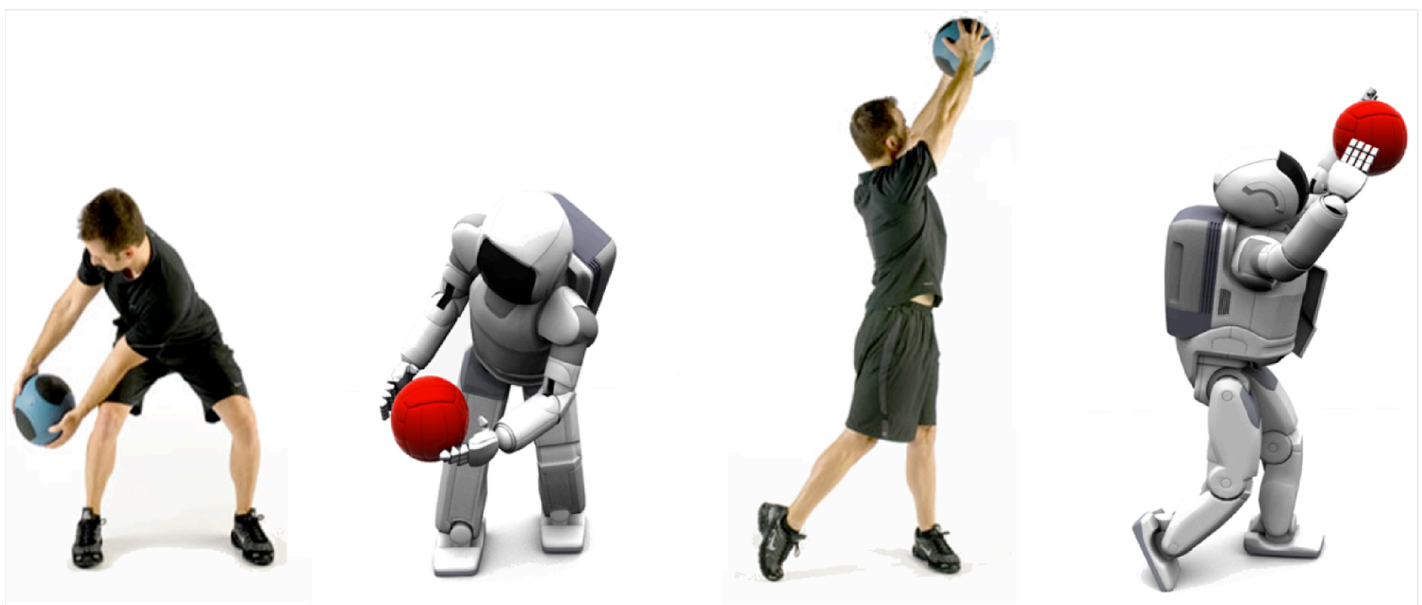

Tesla's NN Driver

Drug Discovery

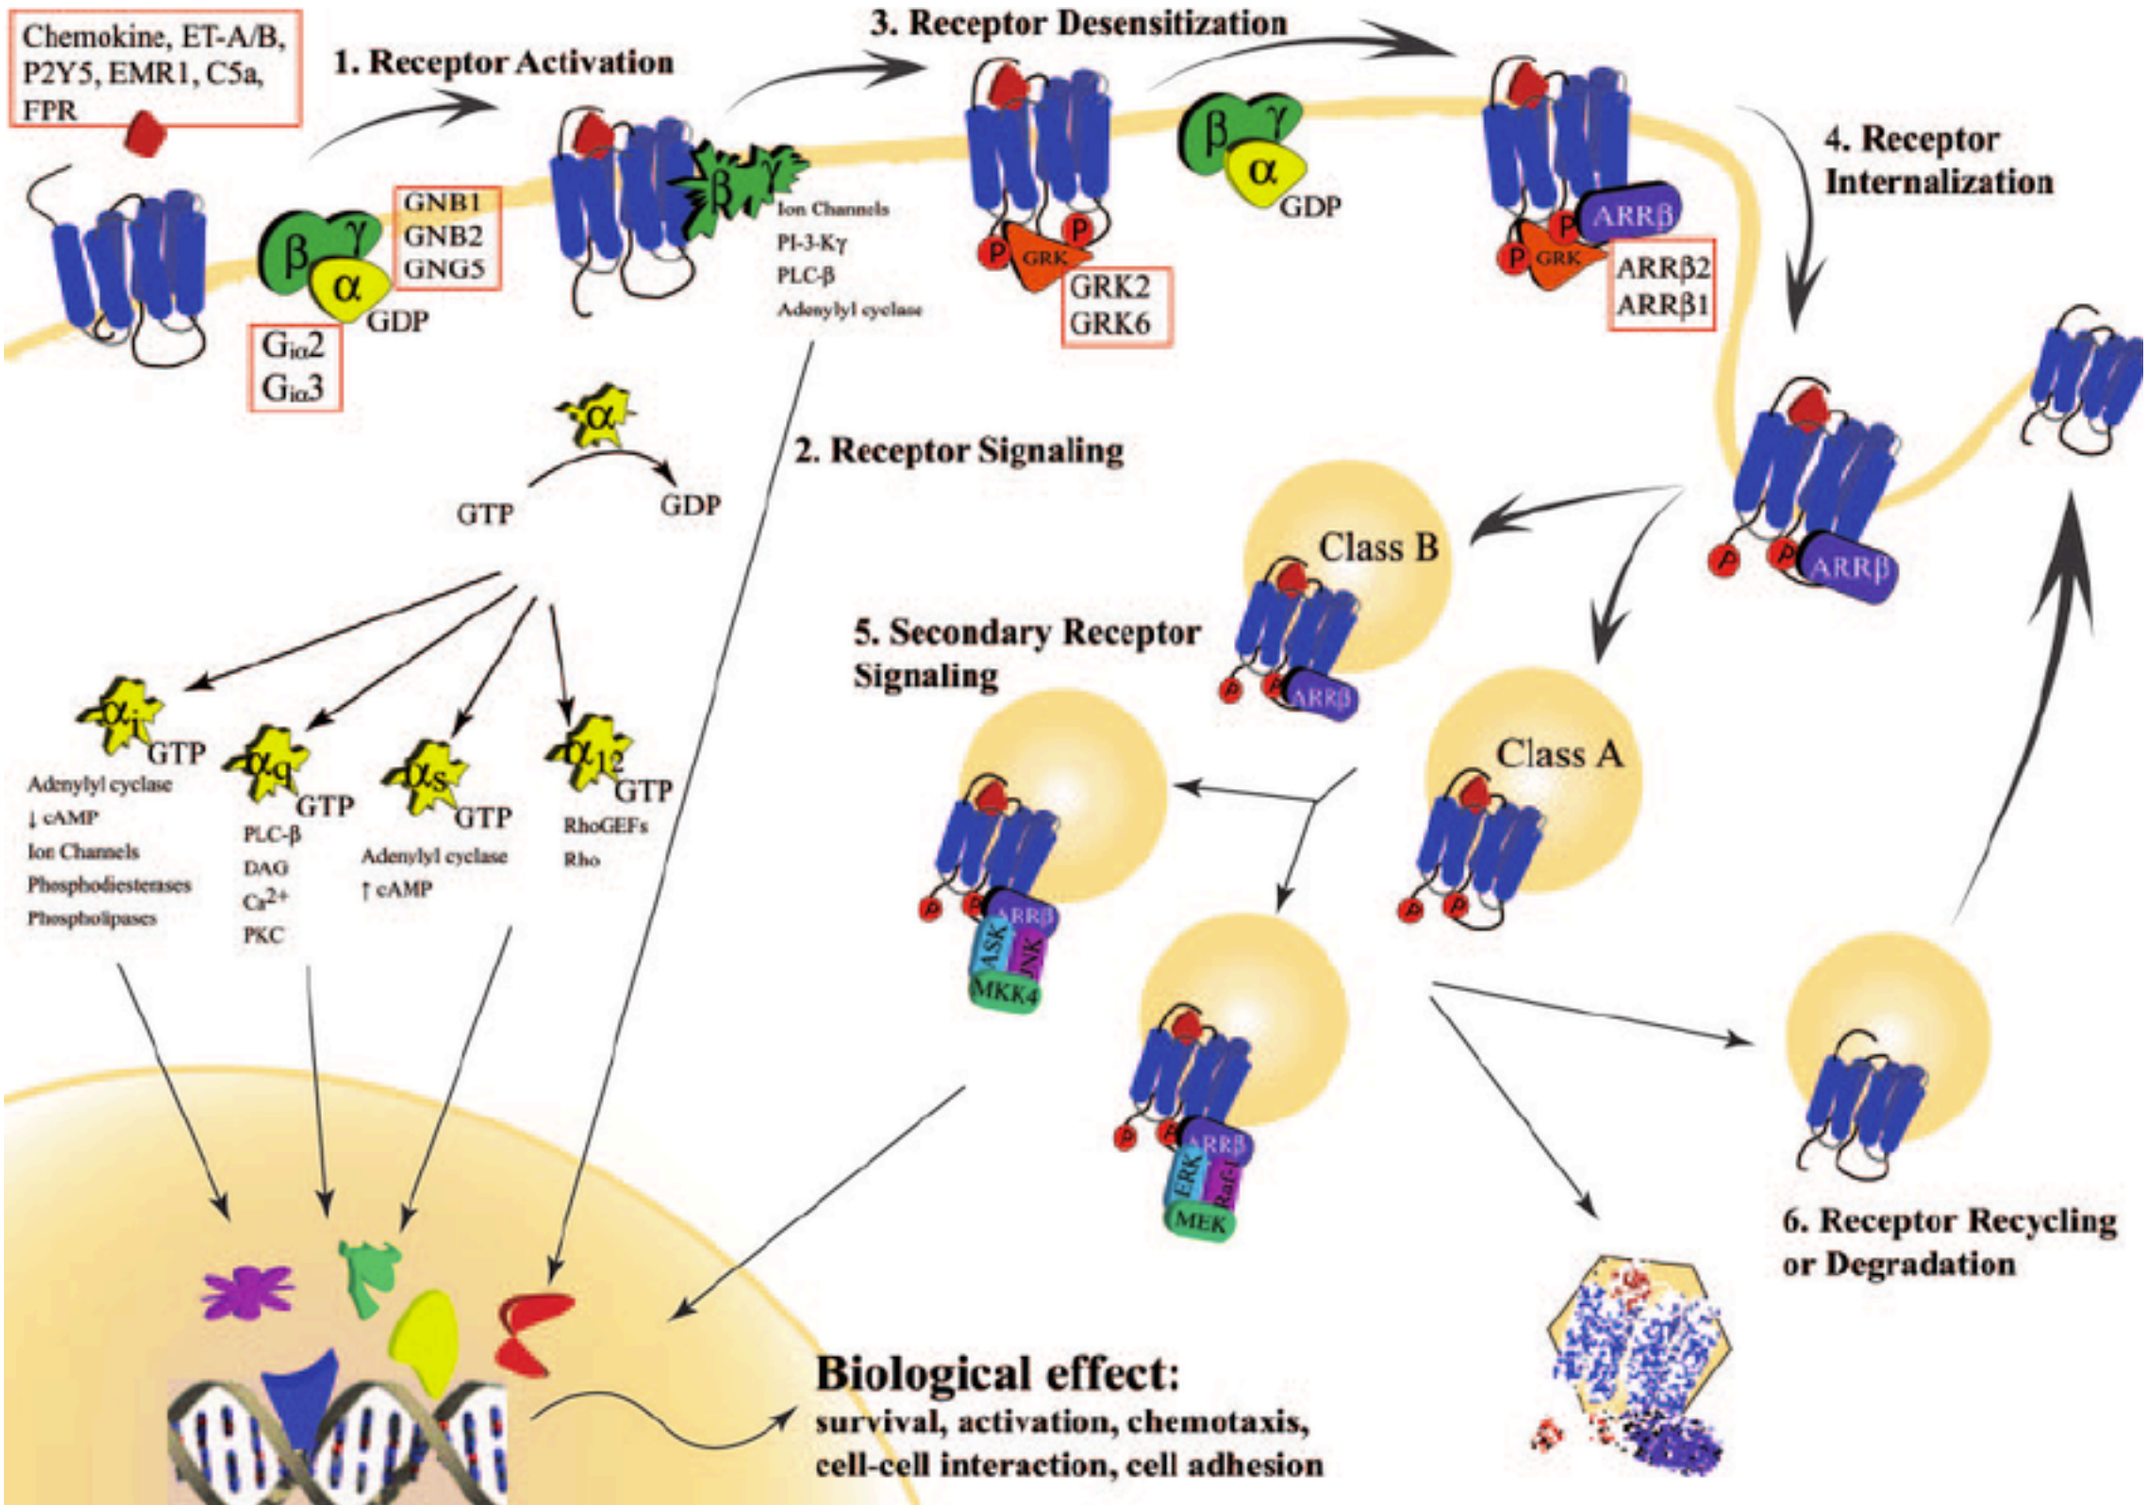

Single Cell Sequencing

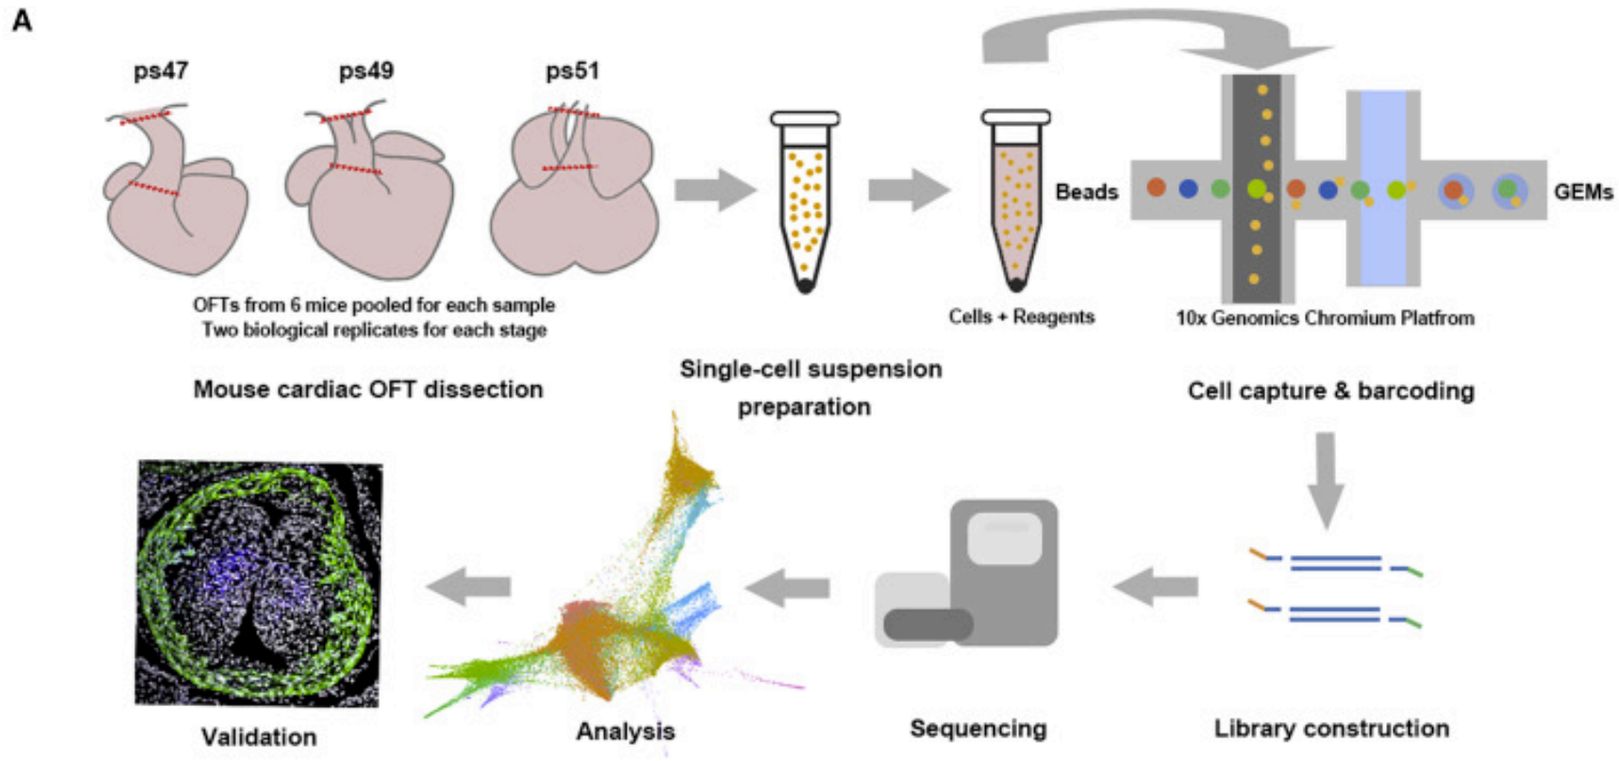

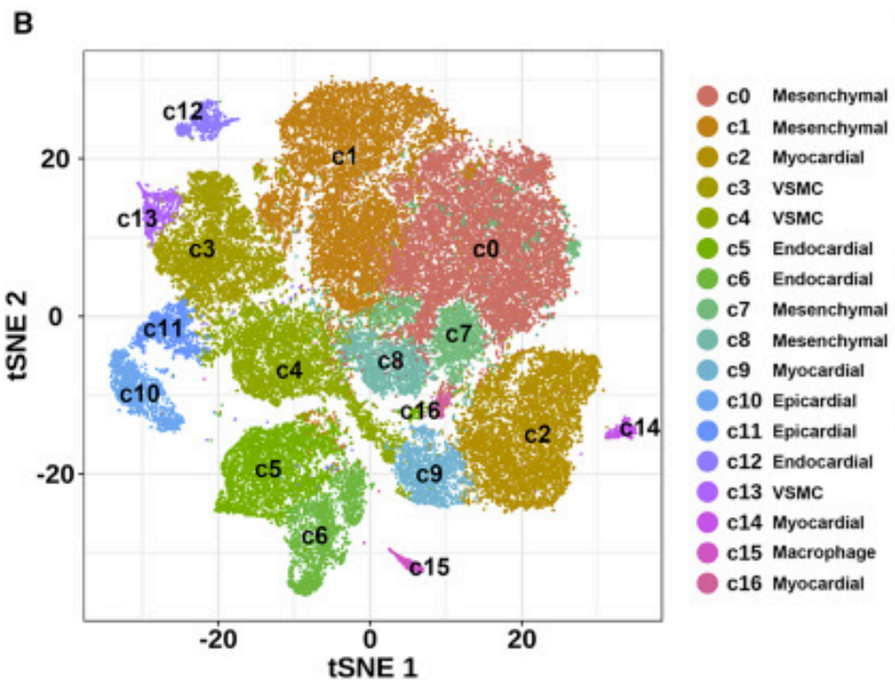

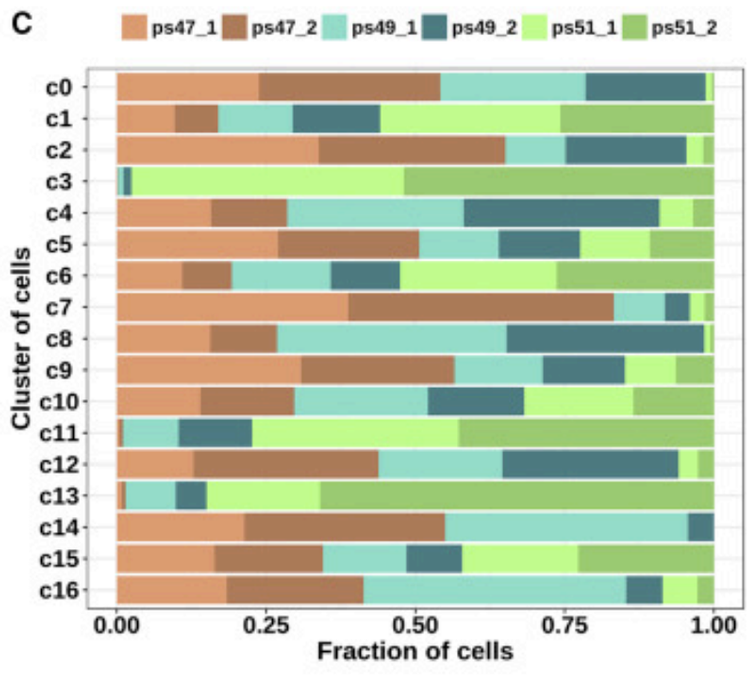

Vision

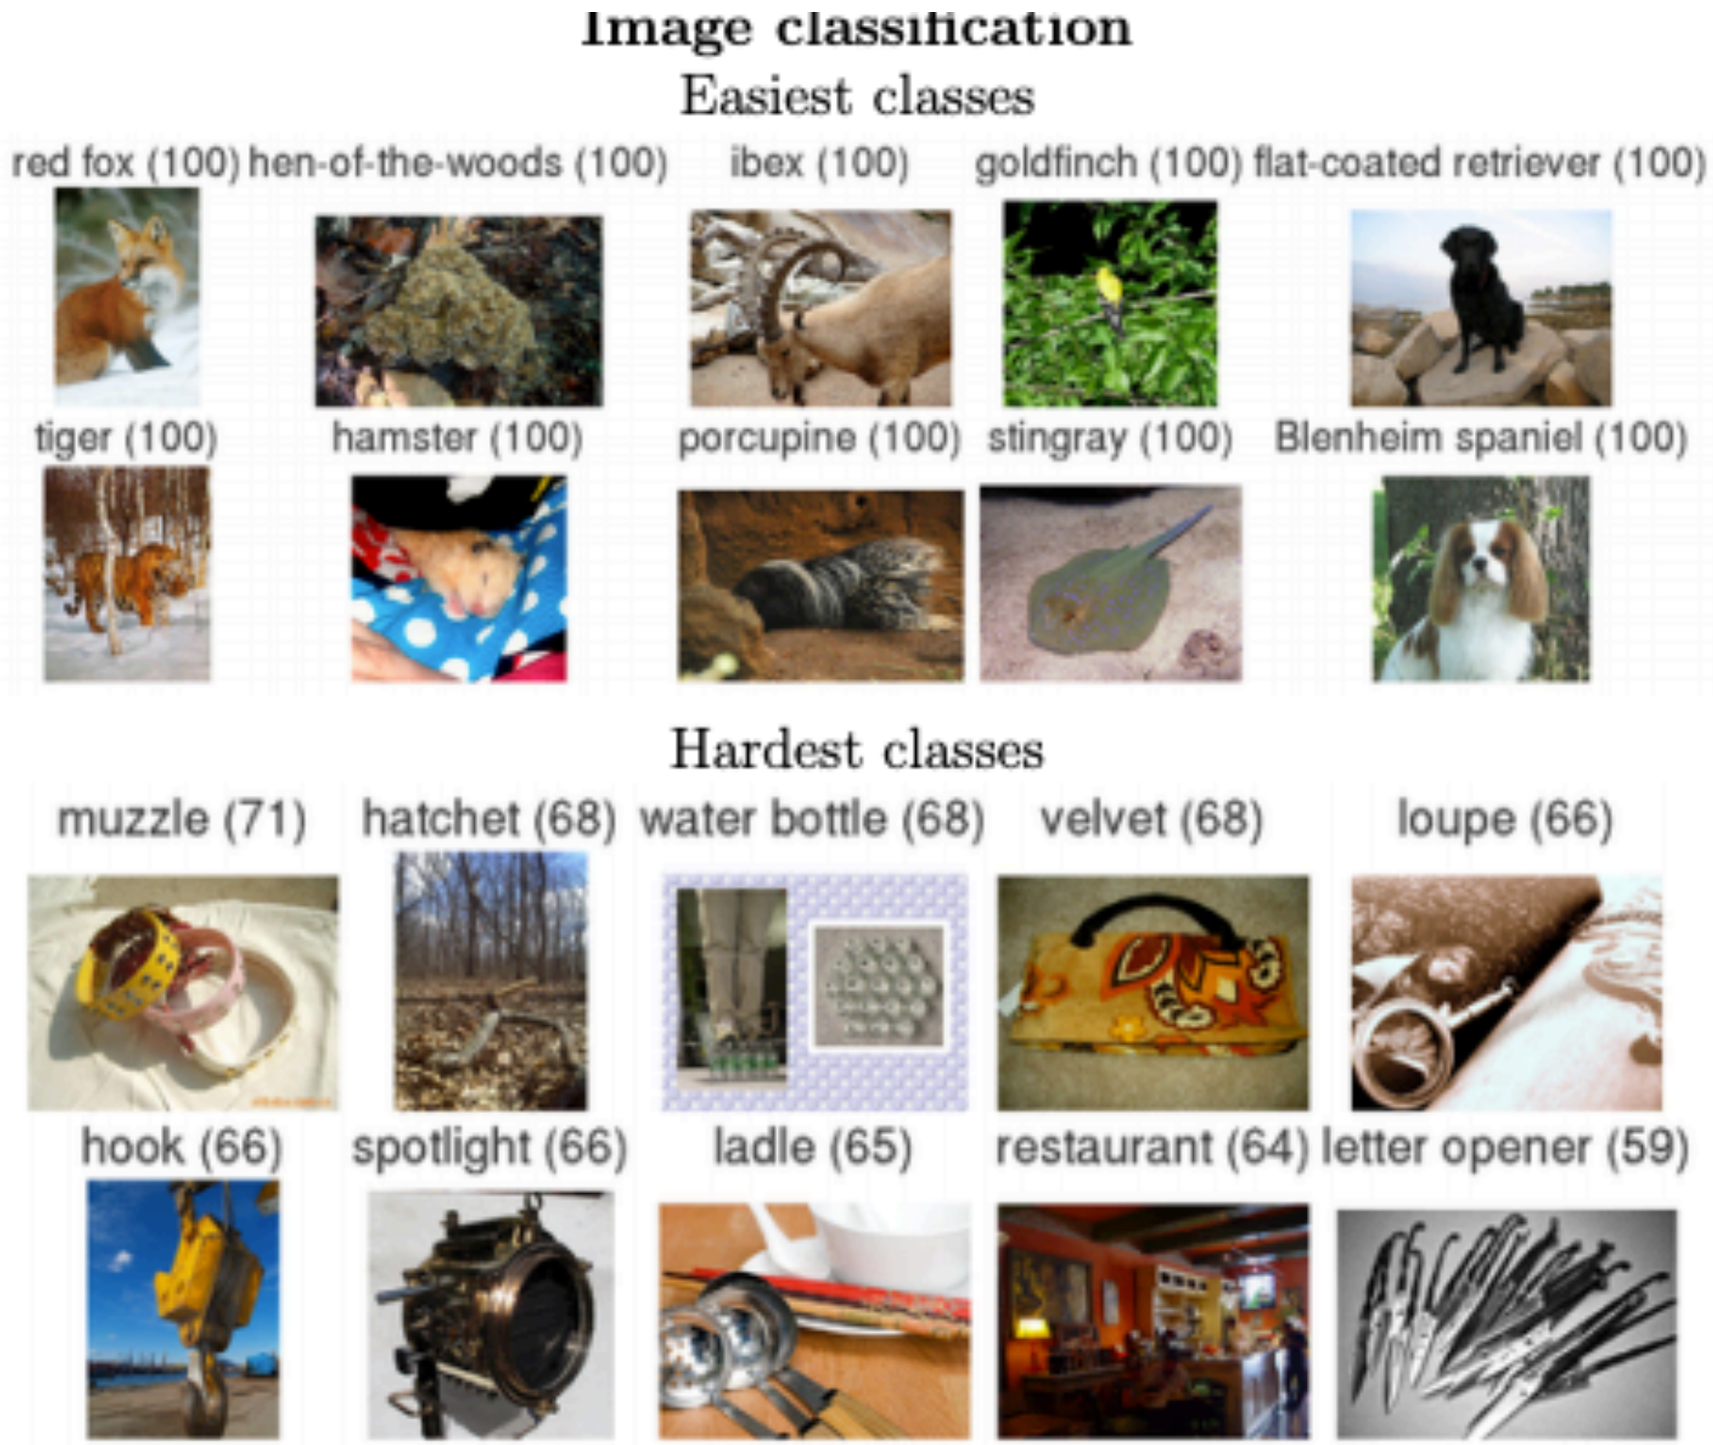
Image classification
Easiest classes
red fox (100)
hen-of-the-woods (100)
ibex (100)
goldfinch (100)
flat-coated retriever (100)
tiger (100)
hamster (100)
porcupine (100)
stingray (100)
Blenheim spaniel (100)
Hardest classes
muzzle (71)
hatchet (68)
water bottle (68)
velvet (68)
loupe (66)
hook (66)
spotlight (66)
ladle (65)
restaurant (64)
letter opener (59)

Avisa o Martins, que o jogo do Porto está na rádio

(a)

(b)

(c)

(d)

What drives this?

Machine Learning

What drives ML?

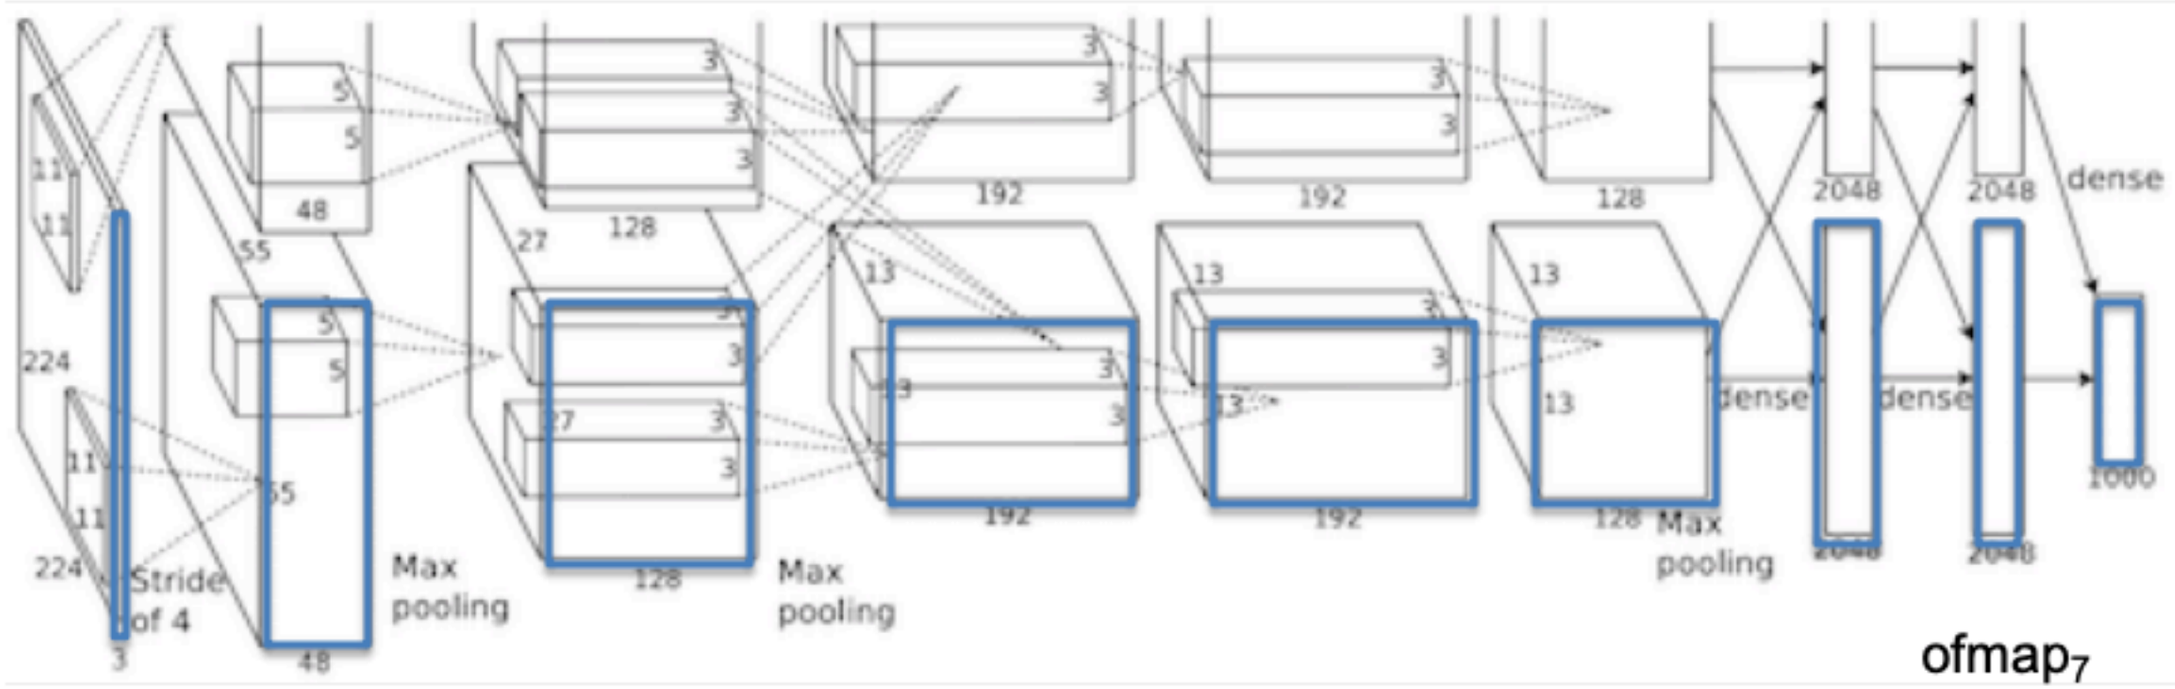

input ofmap$_1$ ofmap$_2$ ofmap$_3$ ofmap$_4$ ofmap$_5$ ofmap$_6$ ofmap$_8$

What drives DNN?

Convex Optimization

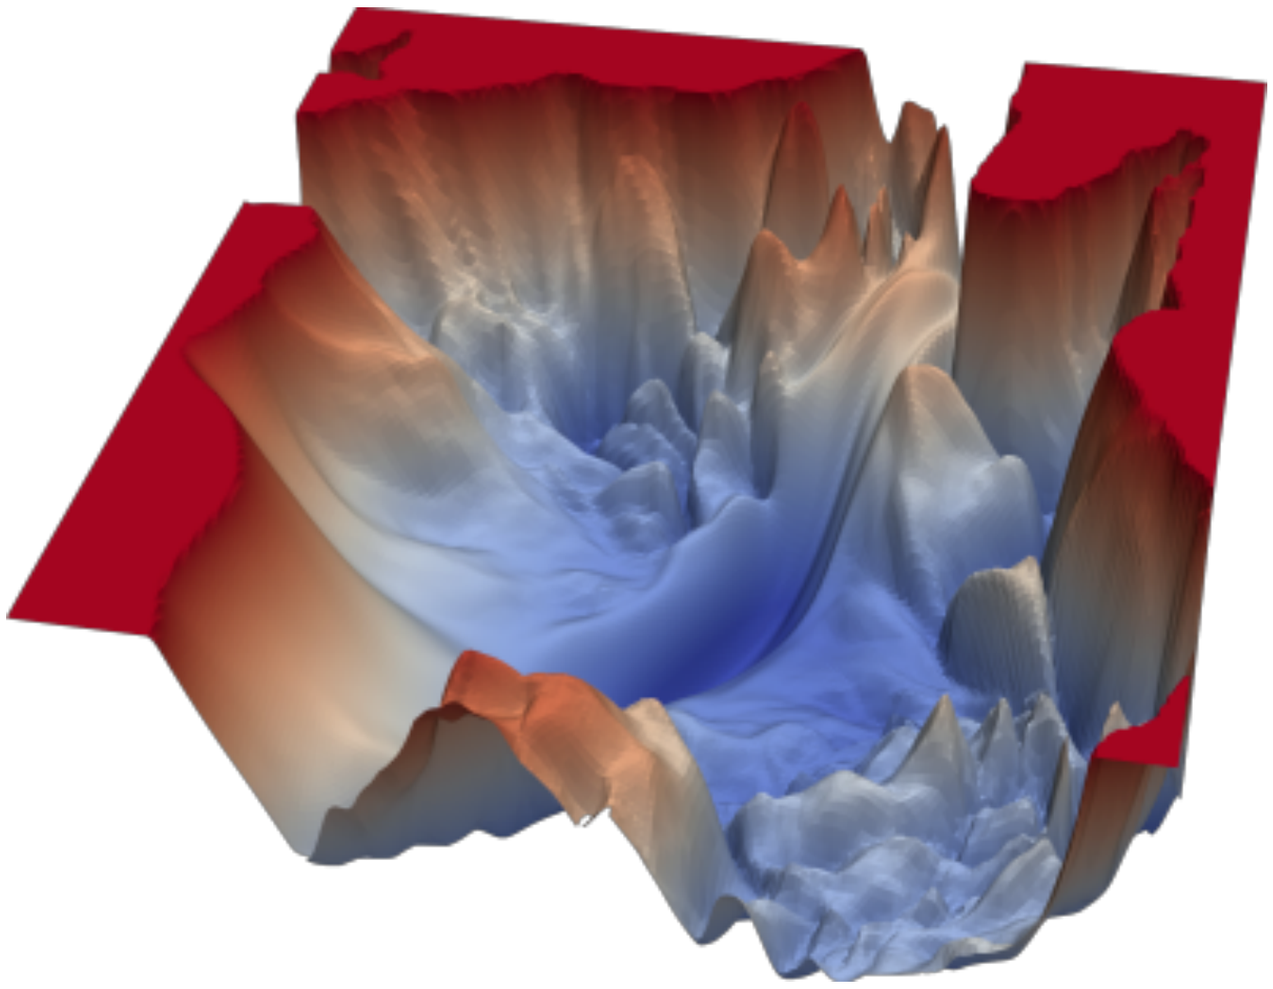

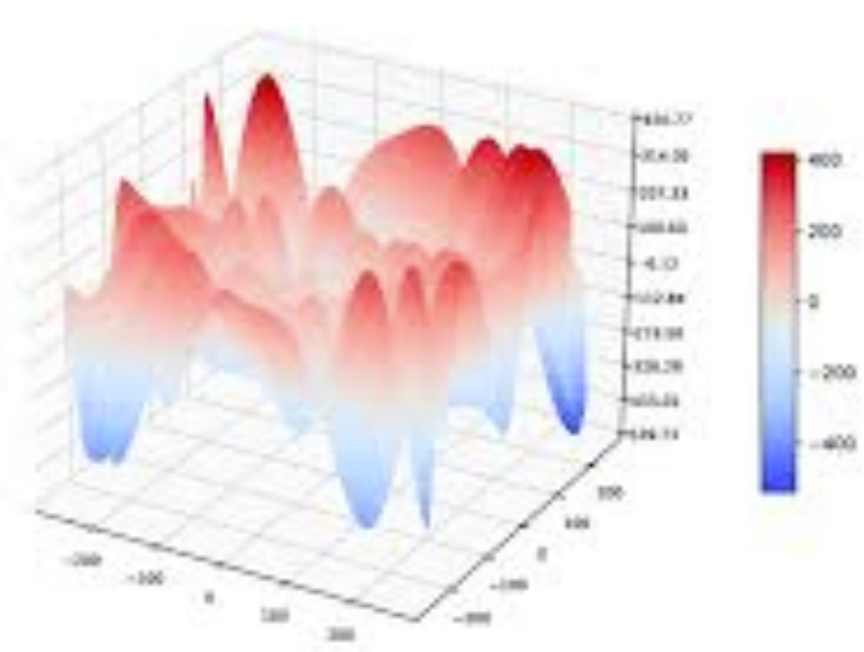

Overfitting Avoidance

- 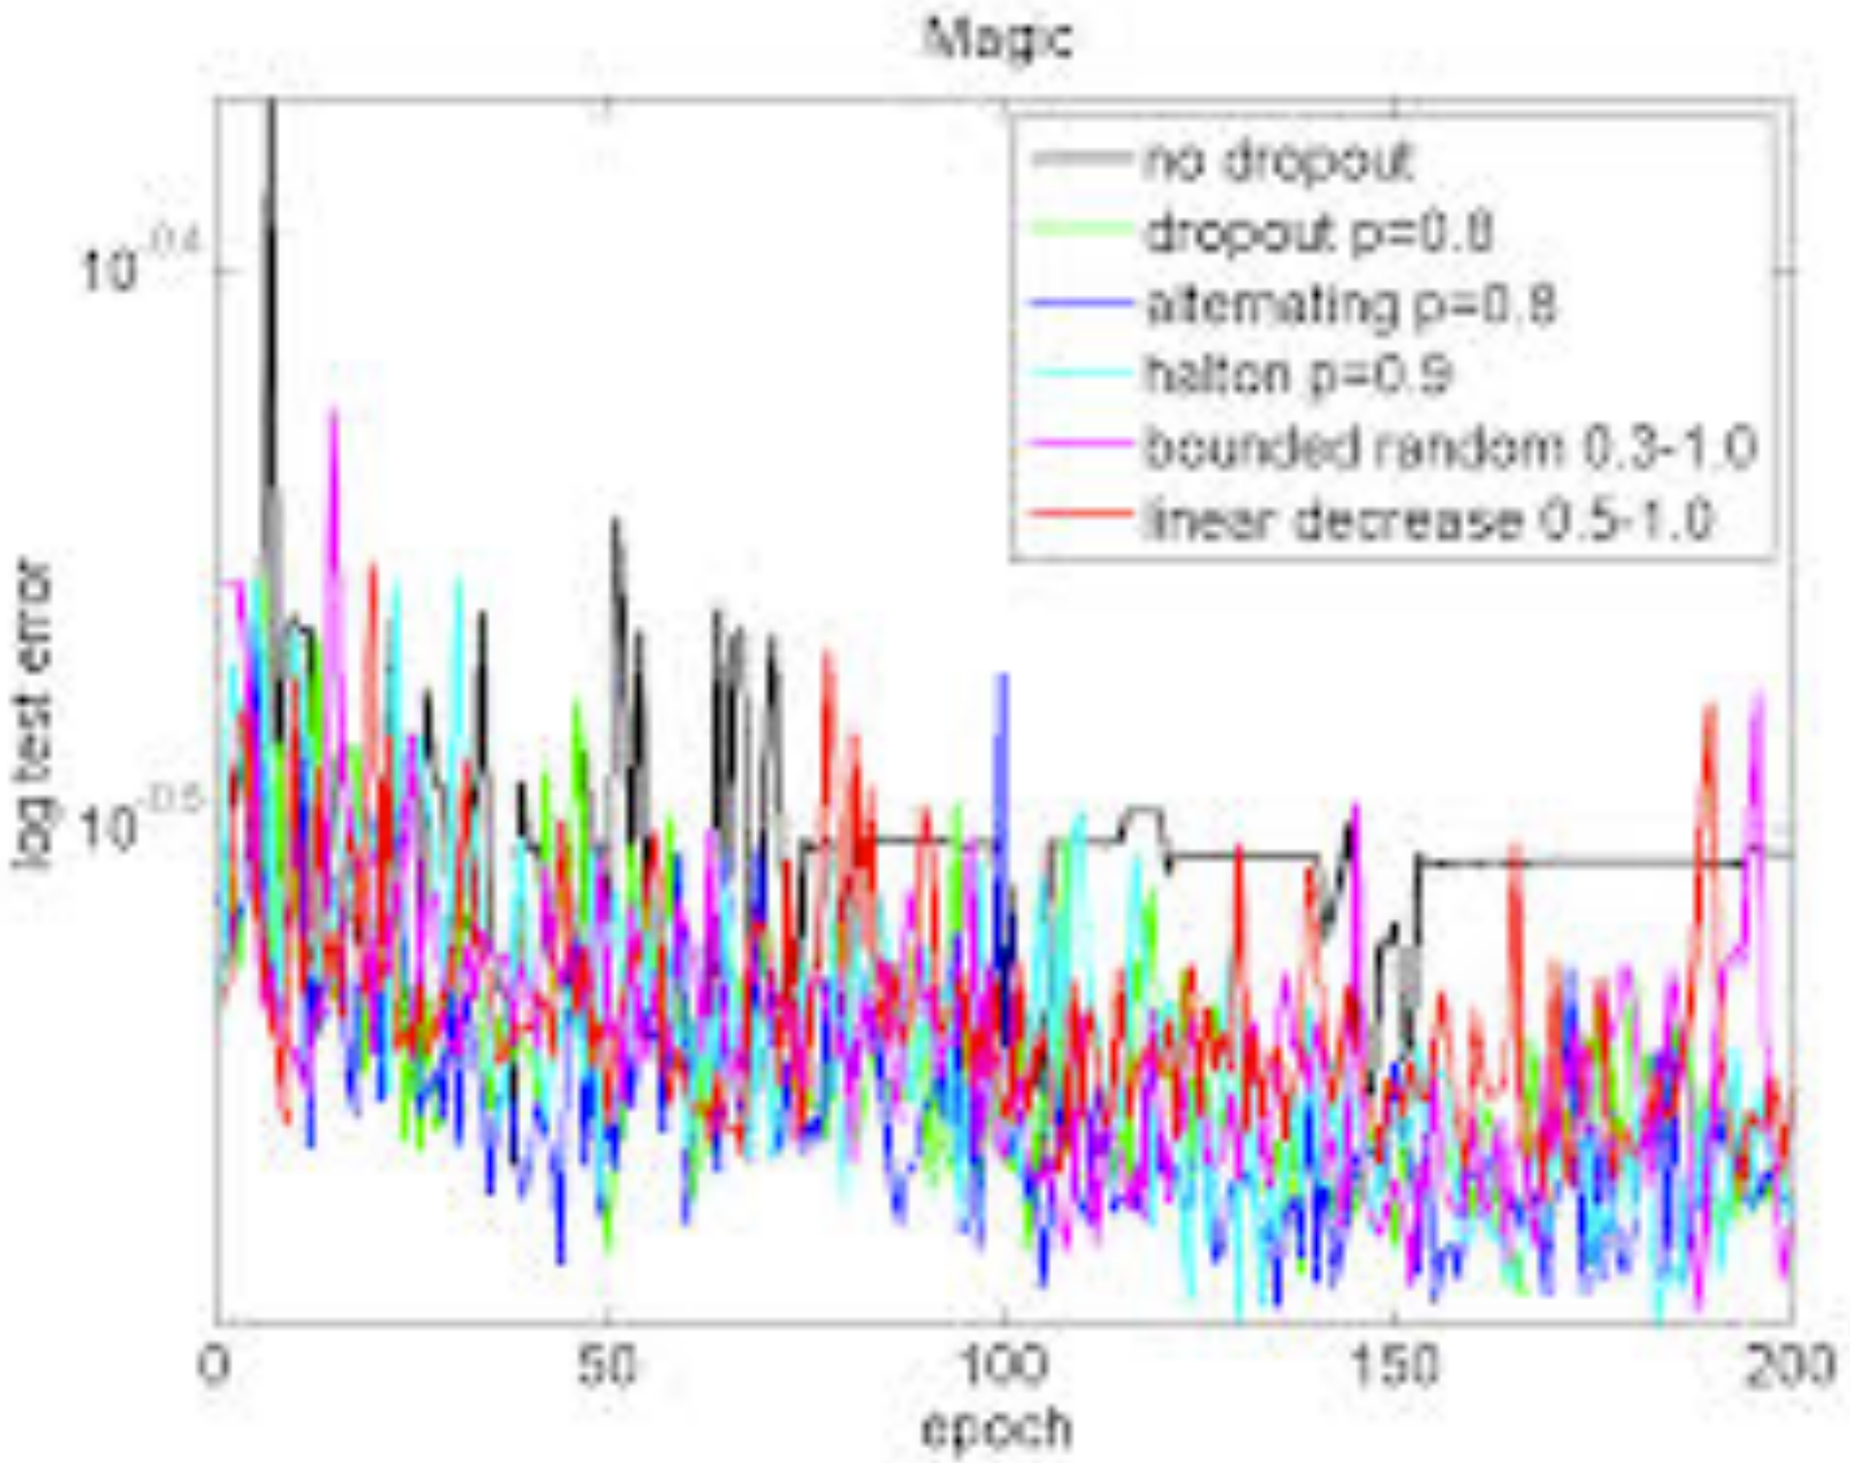

Structure

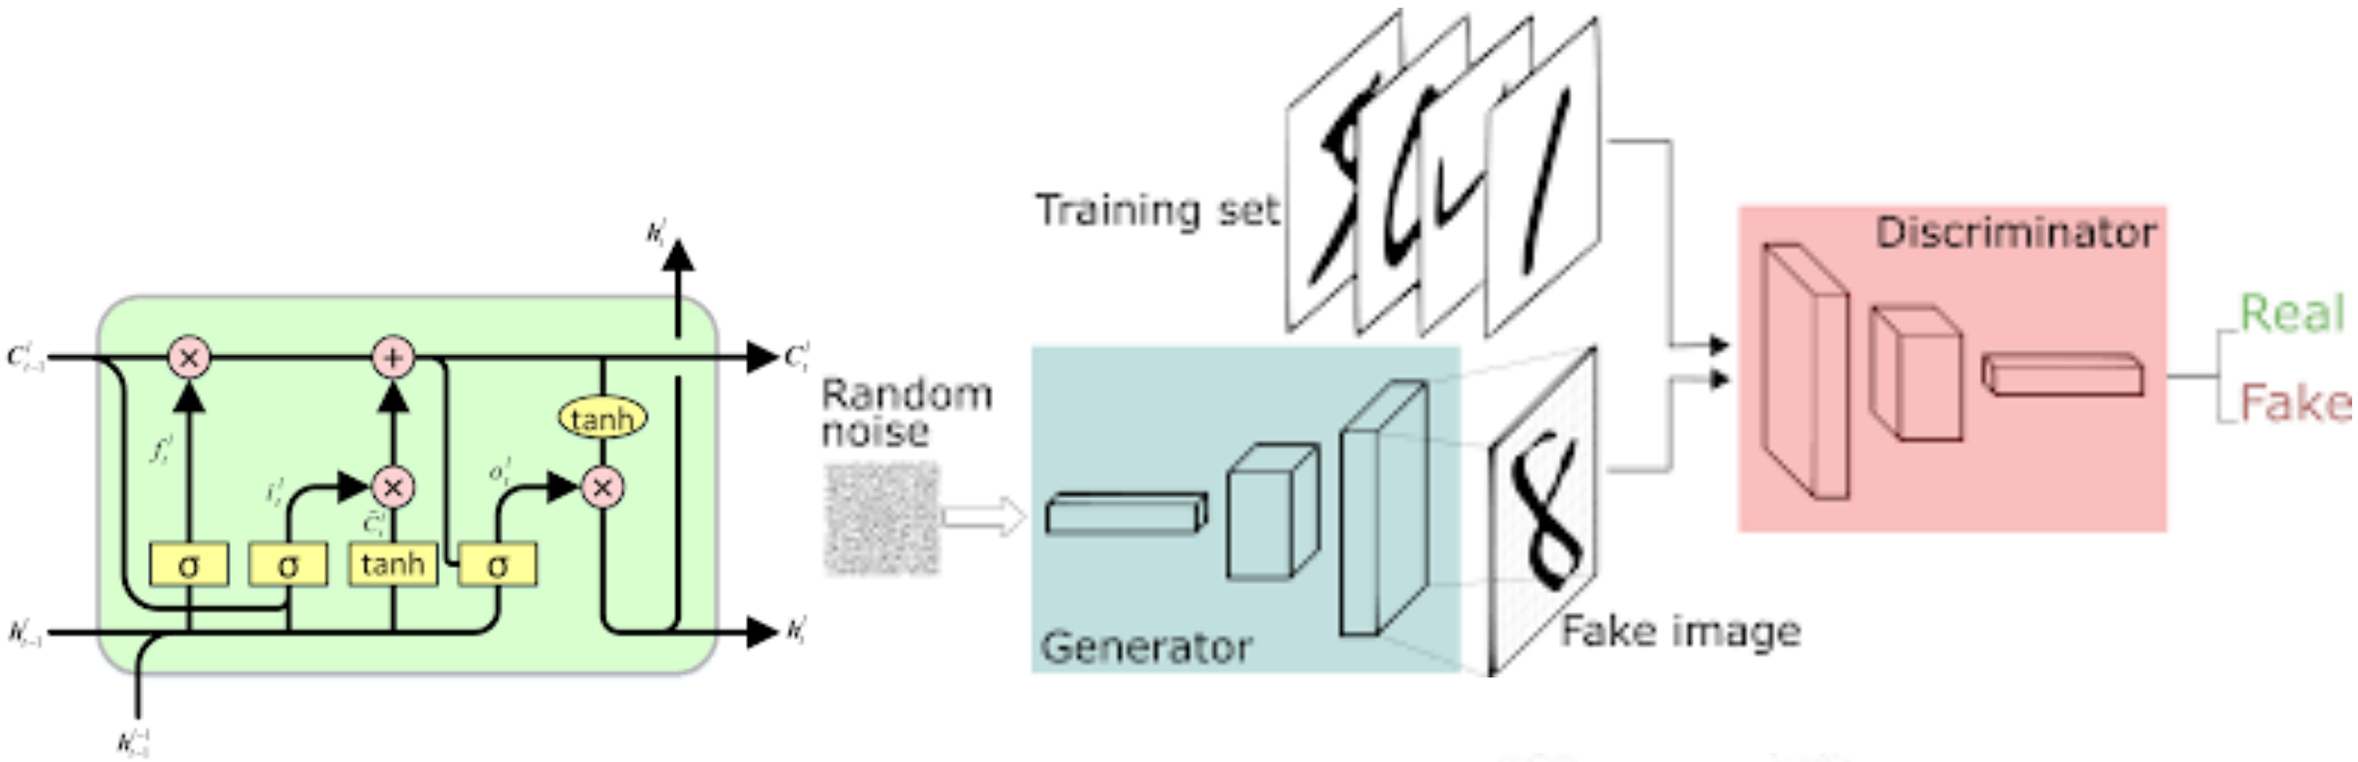
σ
σ
tanh
σ
tanh
Training set
Random noise
Generator
Fake image
Discriminator
Real
Fake

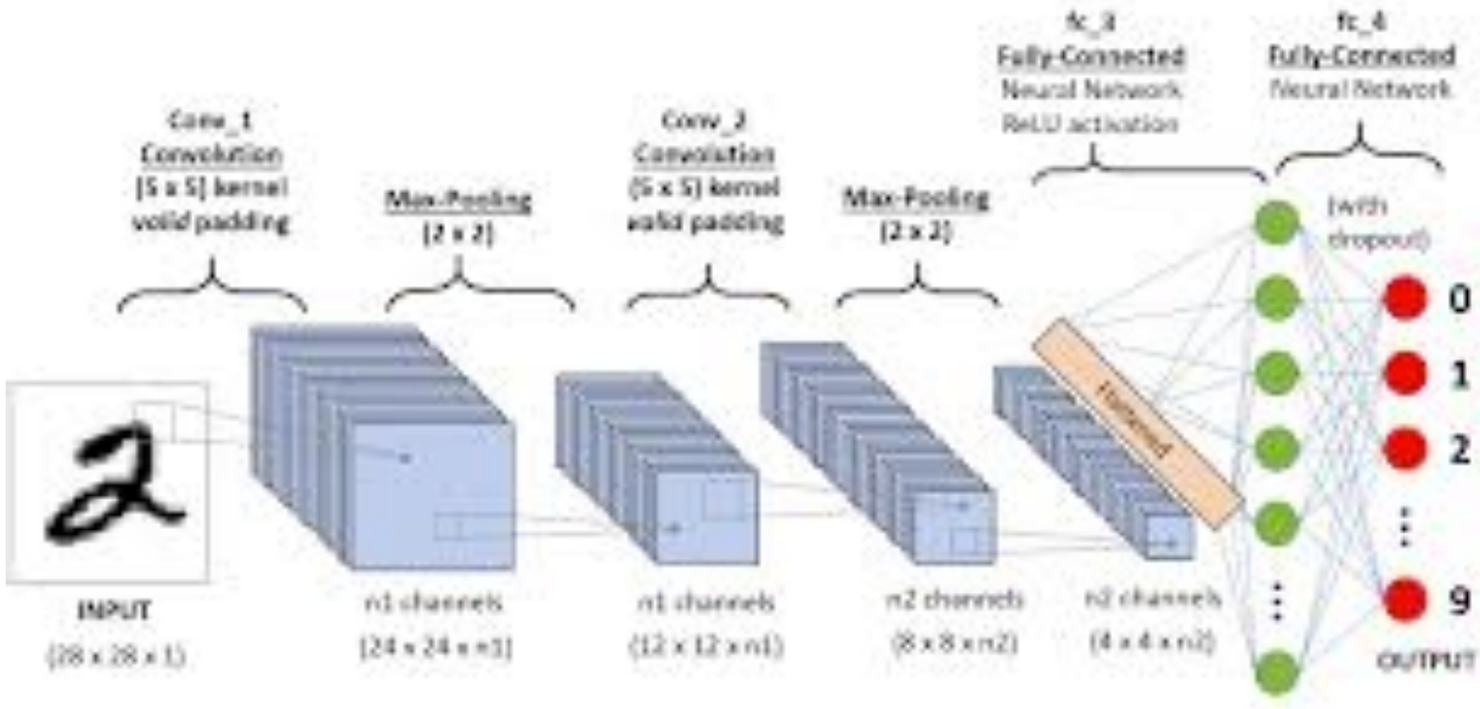
Conv_1
Convolution
[5 x 5] kernel
valid padding
Max-Pooling
(2 x 2)
Conv_2
Convolution
(5 x 5) kernel
valid padding
Max-Pooling
(2 x 2)
fc_3
Fully-Connected
Neural Network
ReLU activation
fc_4
Fully-Connected
Neural Network
(with dropout)
INPUT
(28 x 28 x 1)
n1 channels
(24 x 24 x n1)
n1 channels
(12 x 12 x n1)
n2 channels
(8 x 8 x n2)
n2 channels
(4 x 4 x n2)
Flattened
0
1
2
9
OUTPUT

But First

WE’LL START FROM THE BEGINNIG

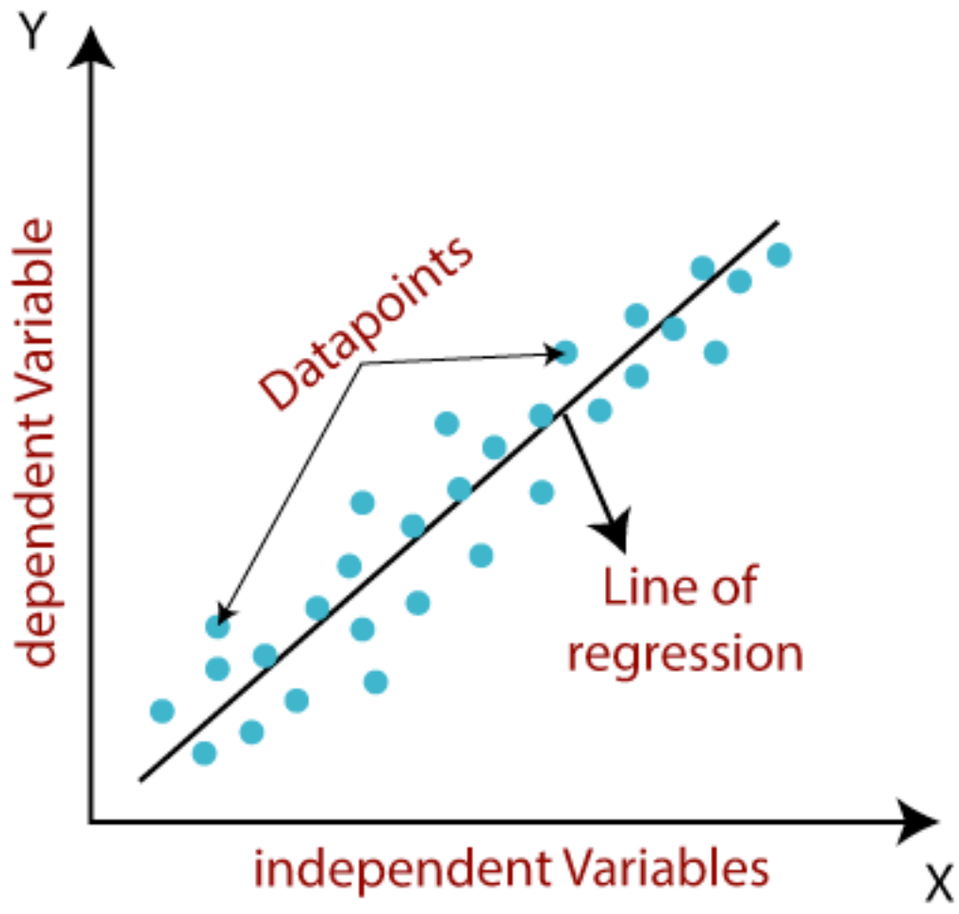

TAIA Future

You shall

- Revise basic ML
- Learn About Graphical Models
- Learn About NNs
- Learn Relational Learning

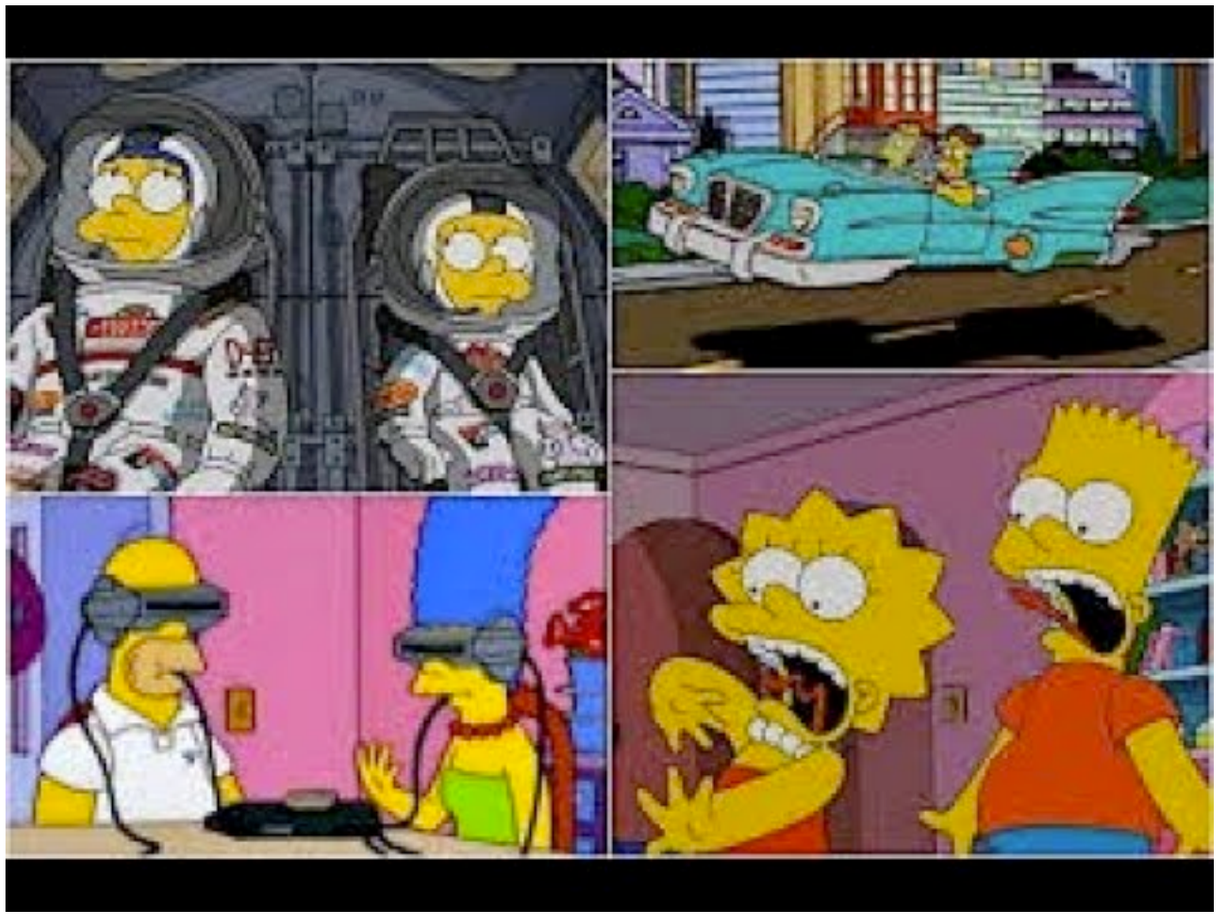

There will be two Halves

They shall be taught by two different people

Starting Point

- Problem Setting:
  - Instances or Examples X
  - Unknown Target Function
    $f : X \rightarrow Y$
  - Set of Possible Models
    $H = \{h \mid h : X \rightarrow Y\}$
- We are given: