Deep Learning

Introduction

Danijel Skočaj
University of Ljubljana
Faculty of Computer and Information Science

Academic year: 2022/23

Media hype

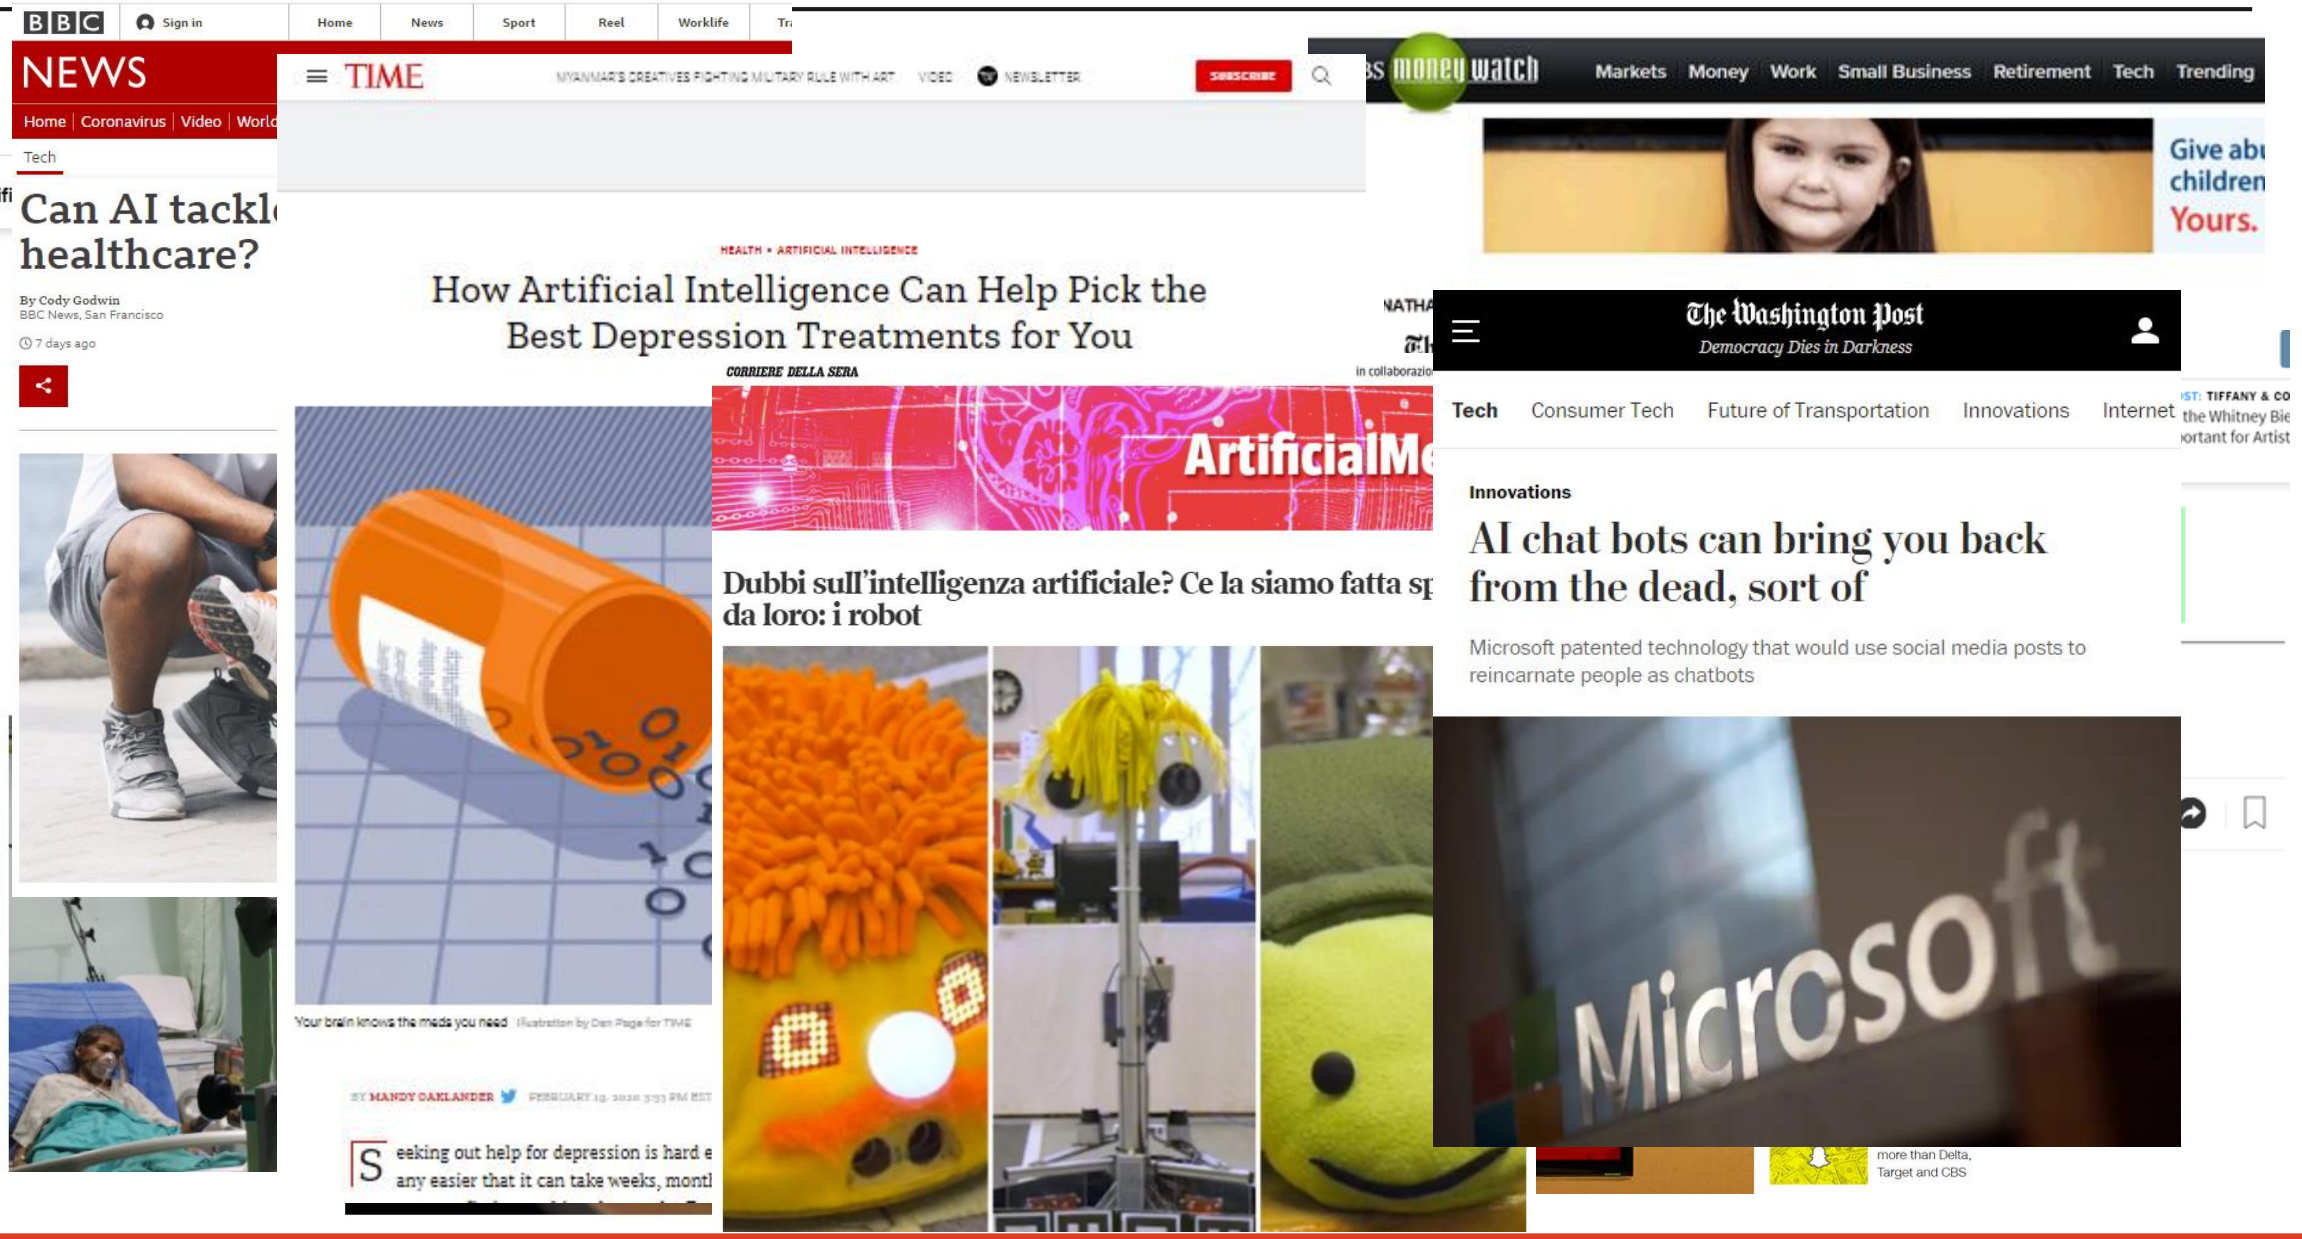
BBC
Sign in
Home
News
Sport
Reel
Worklife
NEWS
Home | Coronavirus | Video
Tech
Can AI tackl
healthcare?
By Cody Godwin
BBC News, San Francisco
7 days ago
TIME
MONEYWATCH
Markets
Money
Work
Small Business
Retirement
Tech
Trending
Give abu
children
Yours.
HEALTH • ARTIFICIAL INTELLIGENCE
How Artificial Intelligence Can Help Pick the Best Depression Treatments for You
CORRIERE DELLA SERA
Dubbi sull'intelligenza artificiale? Ce la siamo fatta sp
da loro: i robot
The Washington Post
Democracy Dies in Darkness
Tech
Consumer Tech
Future of Transportation
Innovations
Internet
Innovations
AI chat bots can bring you back from the dead, sort of
Microsoft patented technology that would use social media posts to reincarnate people as chatbots
Microsoft
Your brain knows the meds you need
BY MANDY OAKLANDER
Seeking out help for depression is hard e
any easier that it can take weeks, montl
more than Delta,
Target and CBS

Superior perfomance

IMAGENET

1k categories

1,3M images

Top5 classification

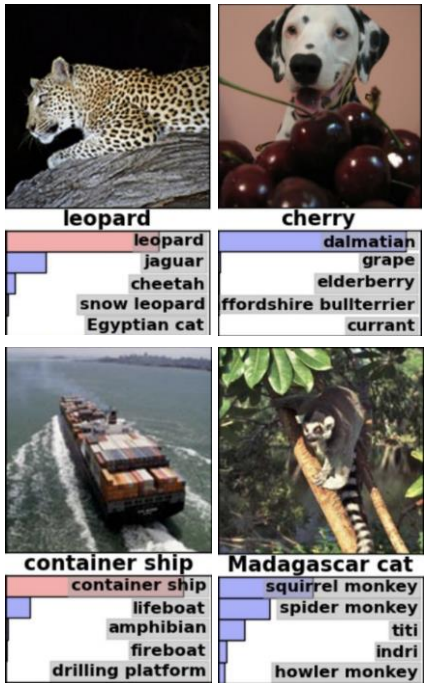

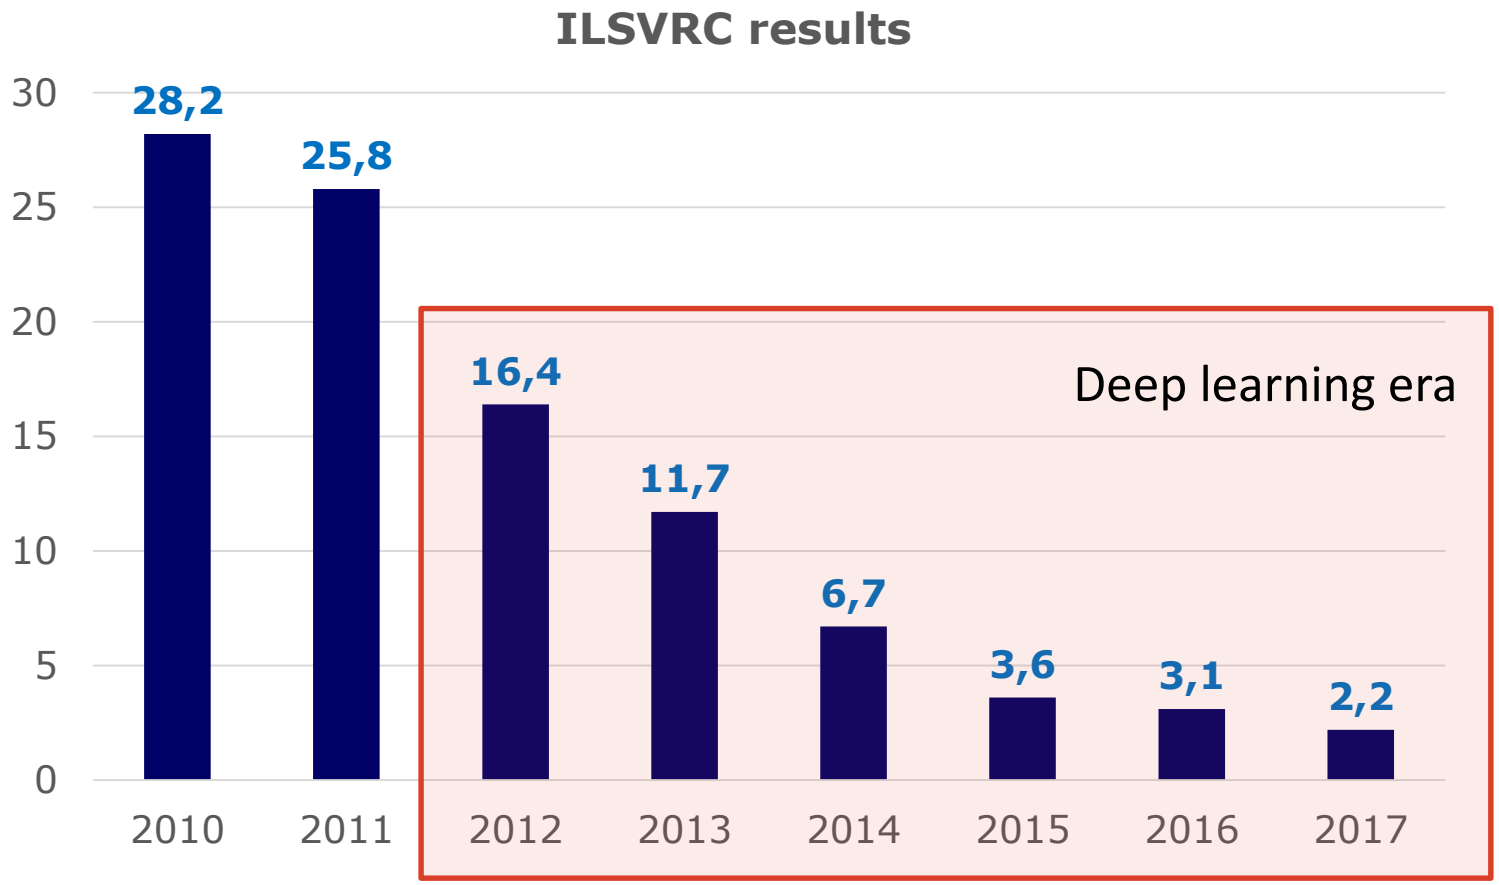

New deep learning era

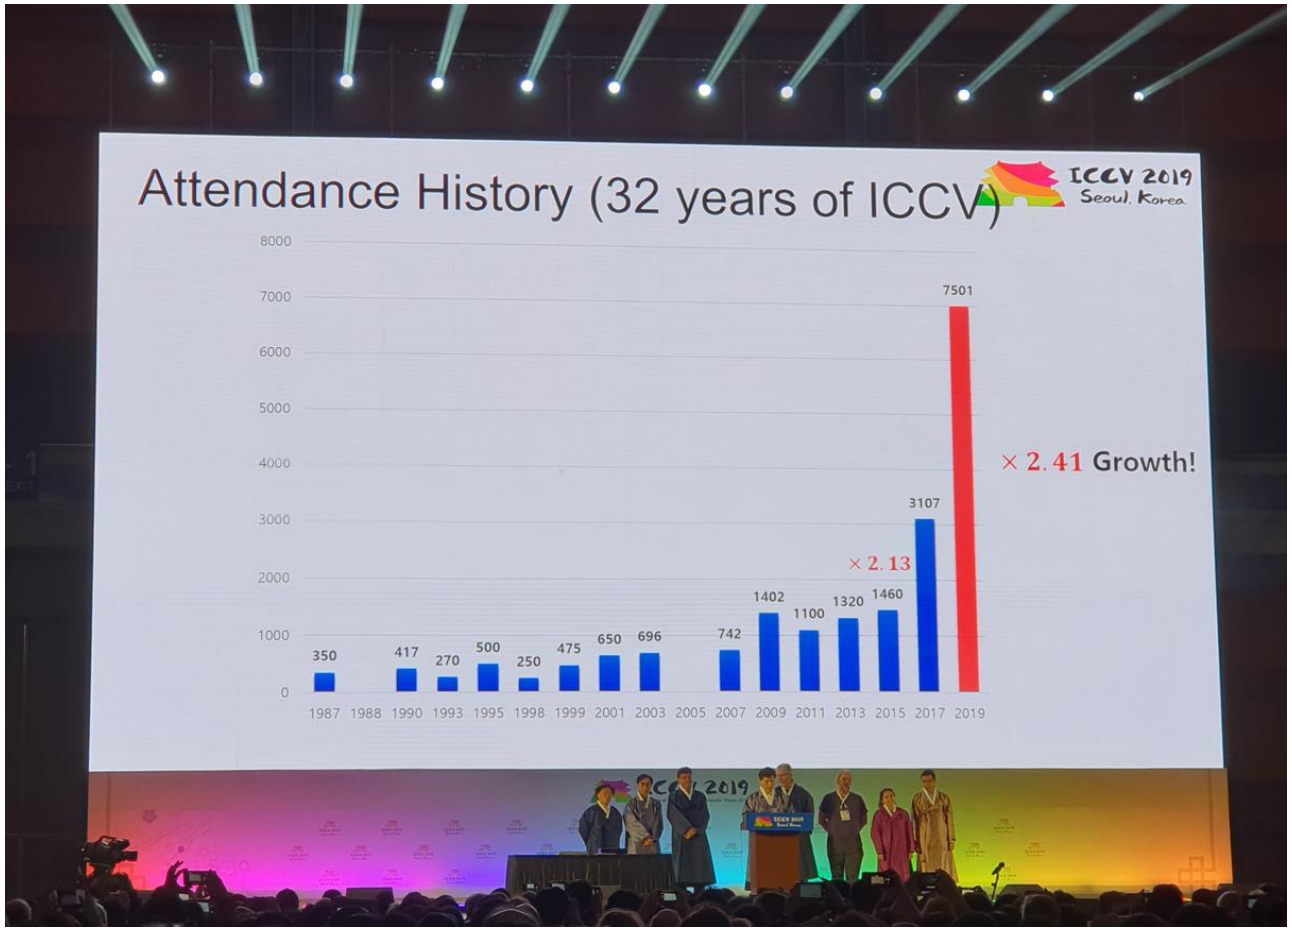

ICCV 2019, Seoul, Korea, 27. 10. - 2. 11. 2019

Turing award 2018

- „Nobel prize“ in computer science

FATHERS OF THE DEEP LEARNING REVOLUTION RECEIVE ACM A.M. TURING AWARD

Bengio, Hinton, and LeCun Ushered in Major Breakthroughs in Artificial Intelligence

Yoshua Bengio

Geoffrey E Hinton

Yann LeCun

Jürgen Schmidhuber

History - Perceptron

Rosenblatt, 1957

- Mark I Perceptron machine
- Frank Rosenblatt, 1957
- Binary classifier
- 20x20 images
- Character recognition

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0, \\ 0 & \text{otherwise} \end{cases}$$

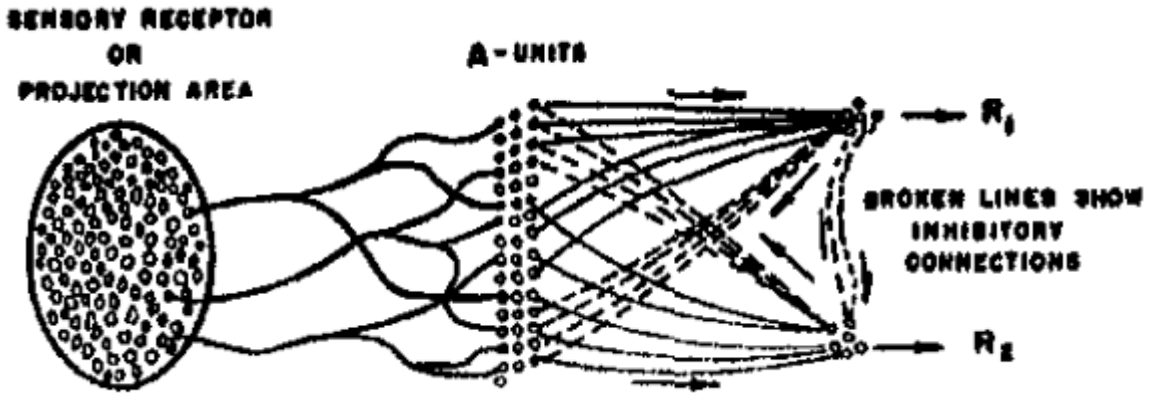

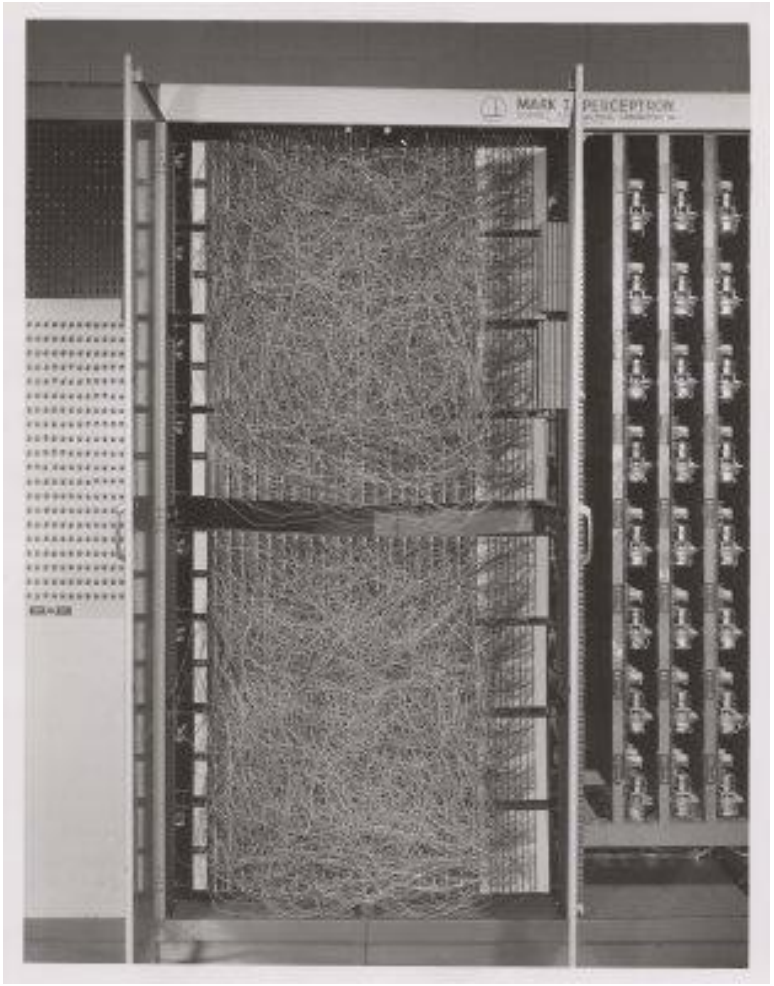

[wikipedia]

- NYT: „the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence."

History – Hubel and Wiesel

Hubel & Wiesel, 1961

- The neural basis of visual perception
- 1950's and 60's
- Neurons fire on a particular orientation of a line
- Receptive fields
- Topographical map in the visual cortex
- Simple and complex cells
- Nobel price in medicine in 1981

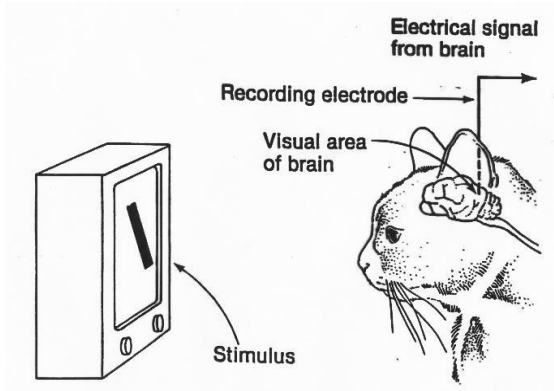

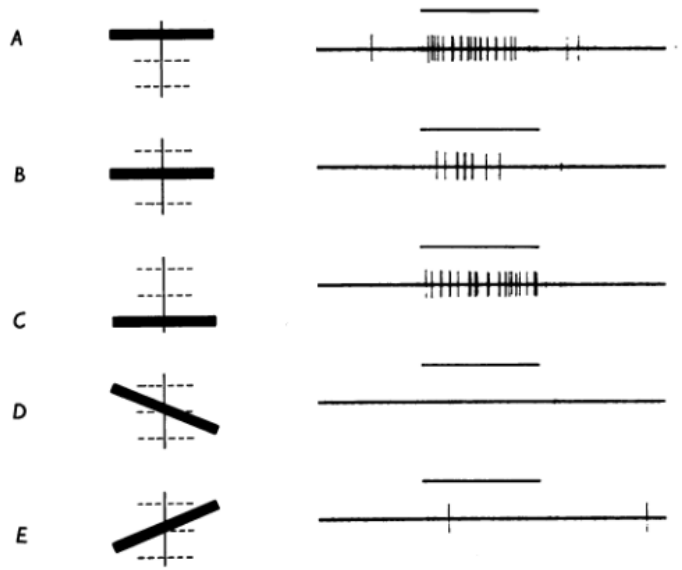

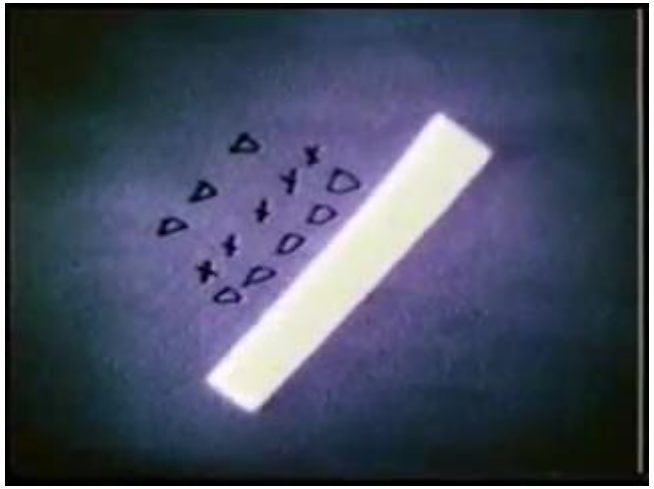

[https://www.youtube.com/watch?v=Cw5PKV9Rj3o]

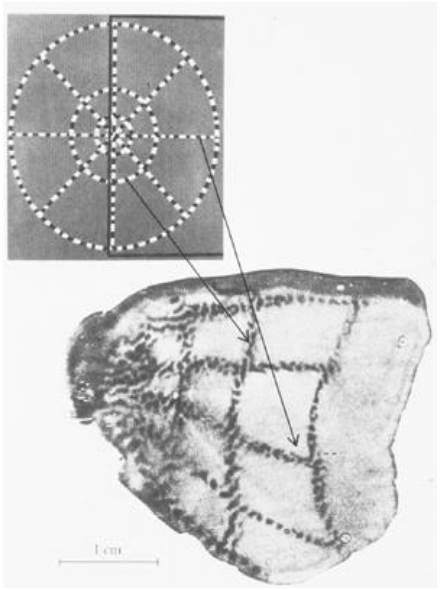

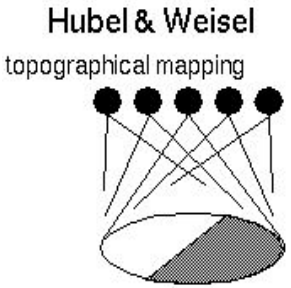

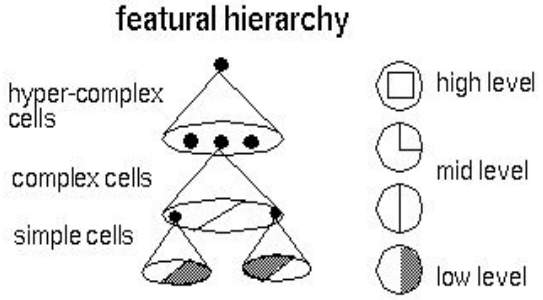

History - Neocognitron

Fukushima, 1980

- Fukushima, 1980
- Pattern recognition
- Hierarchical multilayered neural network
- Simple an complex cells

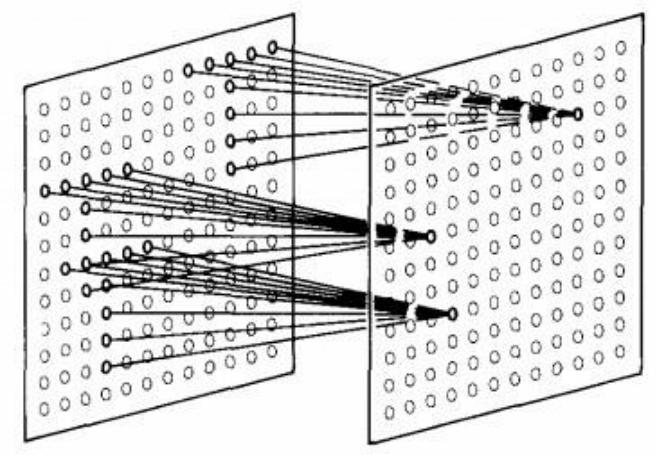

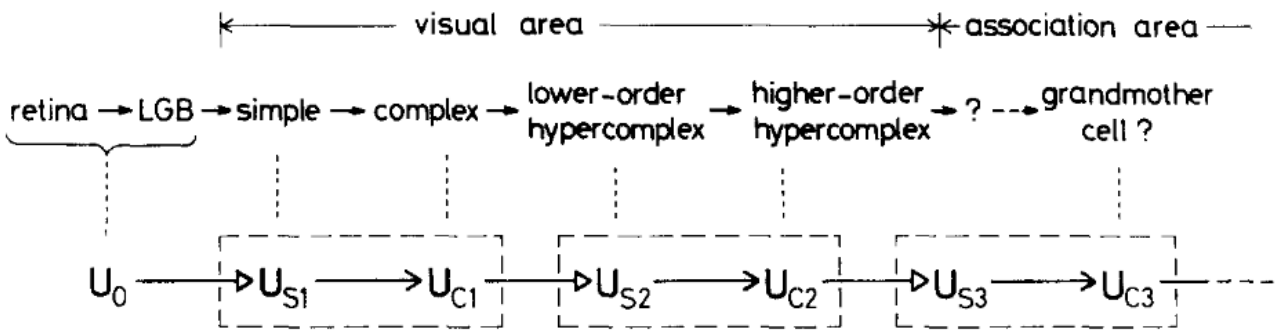

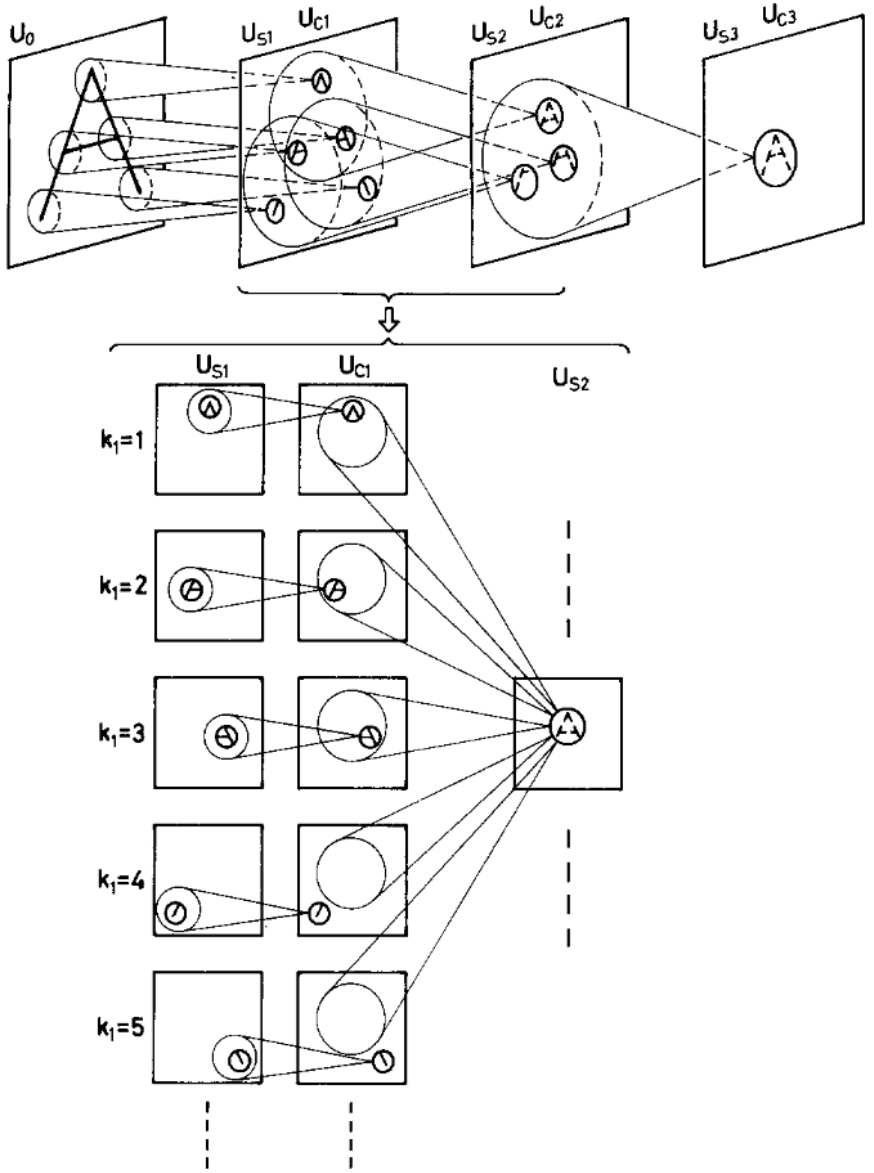

History - Backpropagation

Rumelhart et al., 1986

- Learning representations by backpropagating errors
- Rumelhalt, Hinton, Williams, 1986
- Backpropagation for yielding representations in hidden layers

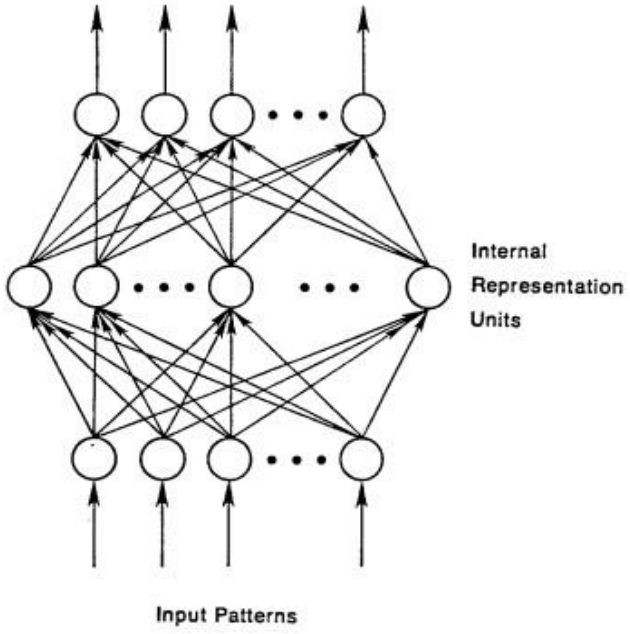

The backward pass starts by computing $\partial E/\partial y$ for each of the output units. Differentiating equation (3) for a particular case, c, and suppressing the index c gives

$$\partial E/\partial y_j = y_j - d_j \tag{4}$$

We can then apply the chain rule to compute $\partial E/\partial x_j$

$$\partial E/\partial x_j = \partial E/\partial y_j \cdot \mathrm{d}y_j/\mathrm{d}x_j$$

Differentiating equation (2) to get the value of $\mathrm{d}y_j/\mathrm{d}x_j$ and substituting gives

$$\partial E/\partial x_j = \partial E/\partial y_j \cdot y_j(1-y_j) \tag{5}$$

This means that we know how a change in the total input x to an output unit will affect the error. But this total input is just a linear function of the states of the lower level units and it is also a linear function of the weights on the connections, so it is easy to compute how the error will be affected by changing these states and weights. For a weight w_{ji}, from i to j the derivative is

$$\begin{aligned}\partial E/\partial w_{ji} &= \partial E/\partial x_j \cdot \partial x_j/\partial w_{ji} \\ &= \partial E/\partial x_j \cdot y_i \end{aligned} \tag{6}$$

and for the output of the i^{th} unit the contribution to $\partial E/\partial y_i$

History - LSTM

- Hochreiter & Schmidhuber,*1997*
- Improved recurrent neural networks
- Overcoming vanishing gradient problem
- Long short-term memory

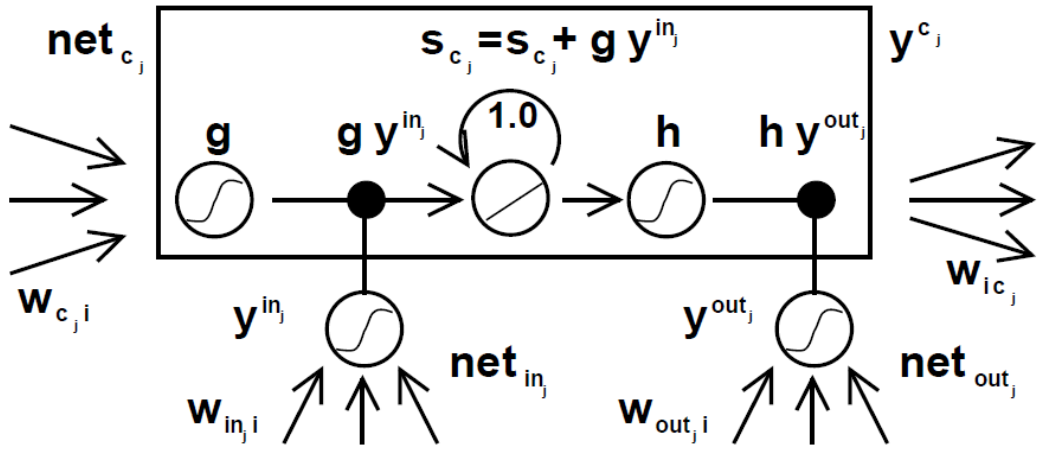

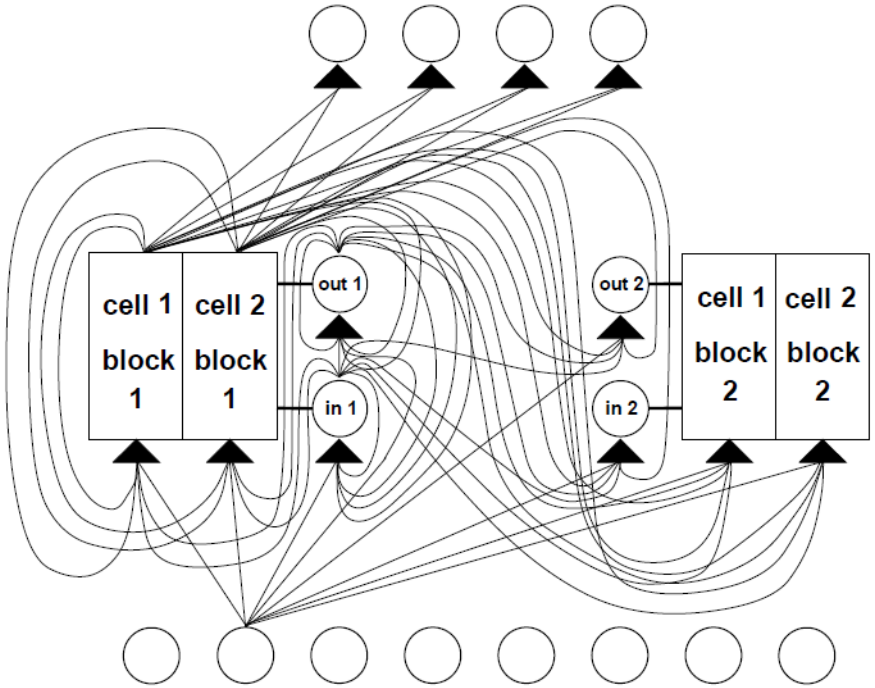

Hochreiter & Schmidhuber, 1997

History – LeNet-5

- Gradient-based learning
  Applied to Document Recognition
- LeCun, Bottou, Bengio, Haffner, 1998
- Convolutional Neural Networks
- MNIST dataset

- Architecture LeNet-5

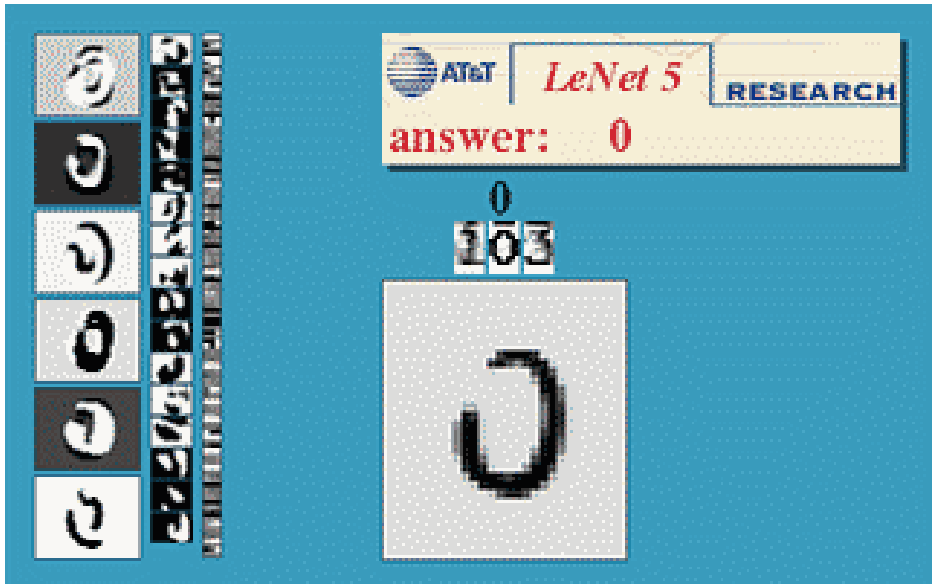

[http://yann.lecun.com]

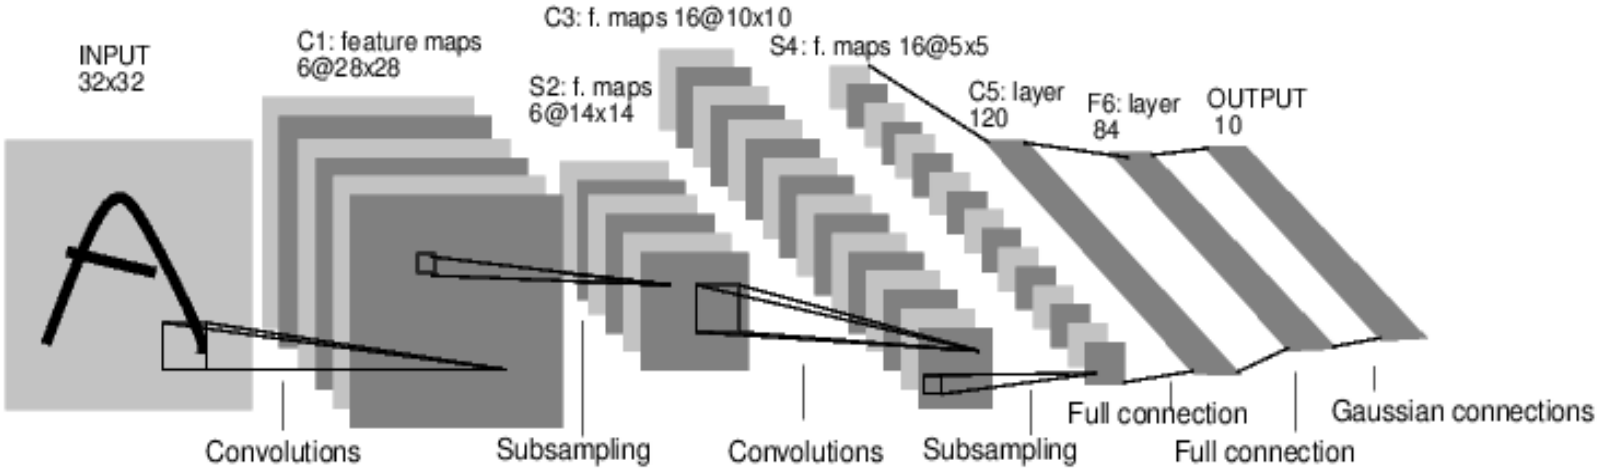

LeCun et al., 1998

History - pretraining

- Reducing the Dimensionality of Data with Neural Networks
- Hinton & Salakhutdinov, 2006
- Autoencoder
- Effective way of initializing the weights
- Pretraining, fine-tuning

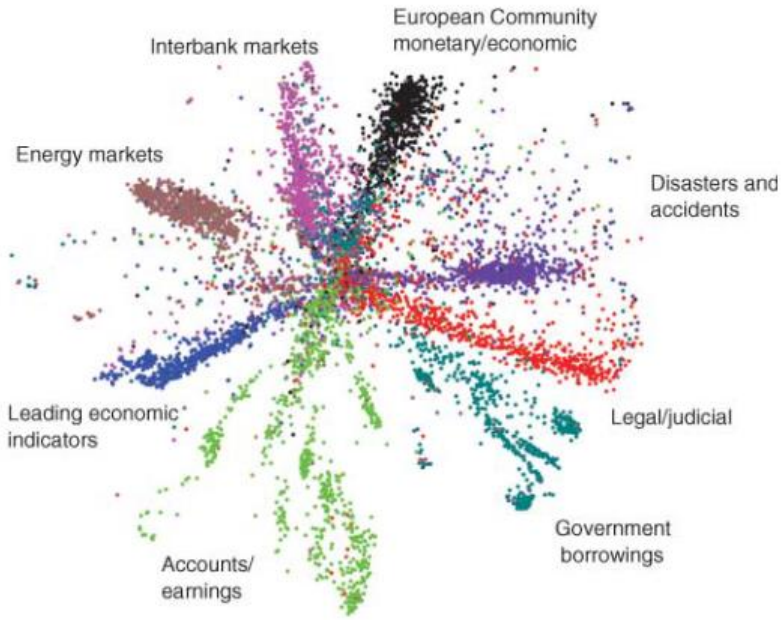

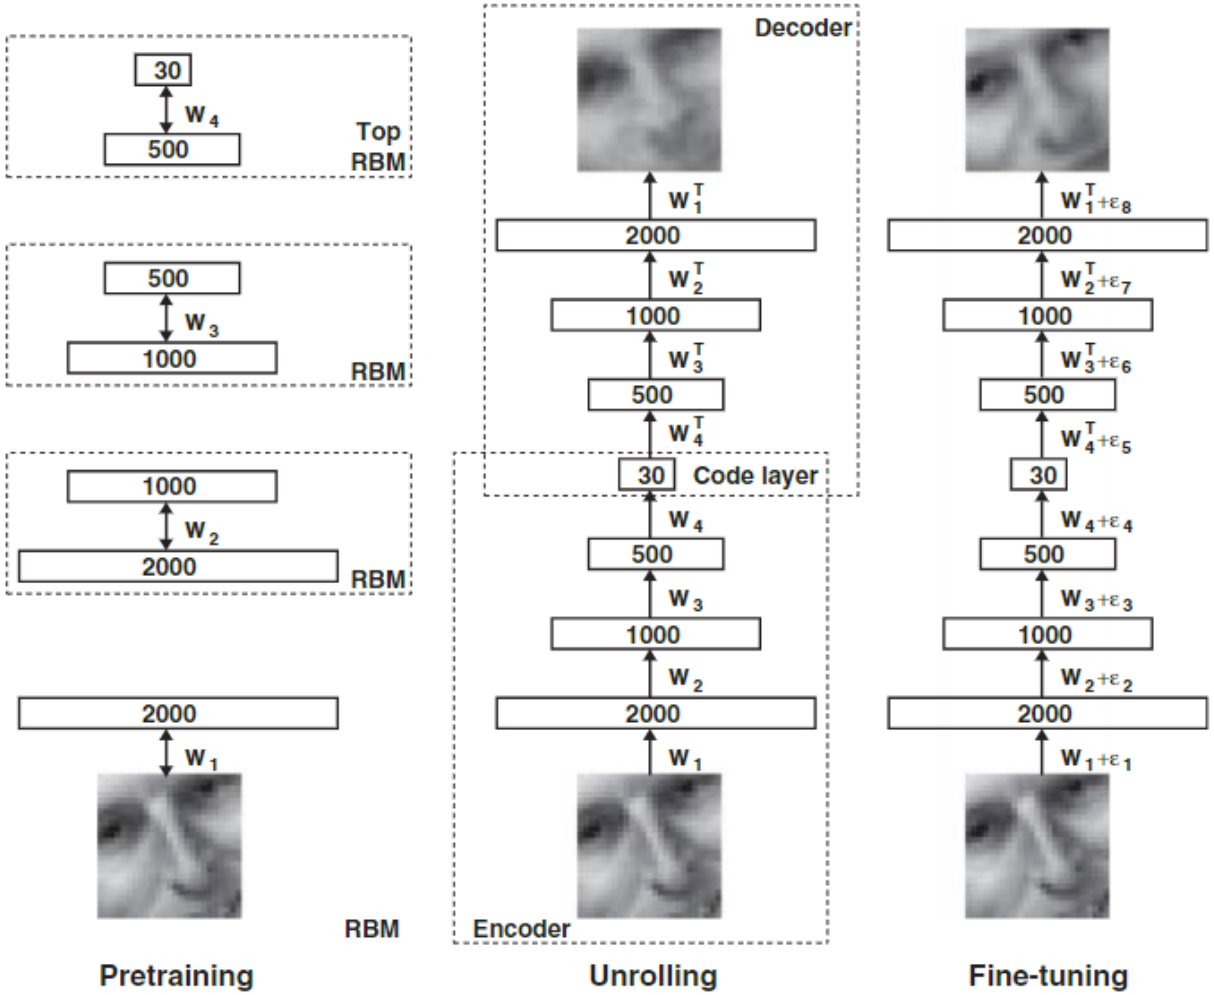

Hinton & Salakhutdinov, 2006

History – AlexNet: beginning of new era

- ImageNet classification with deep convolutional neural networks
- Alex Krizhevsky, Ilya Sutskever, Geoffrey Hinton, 2012
- ILSVRC 2012 Winner
- Significantly improved the results in several computer vision tasks
- AlexNet architecture

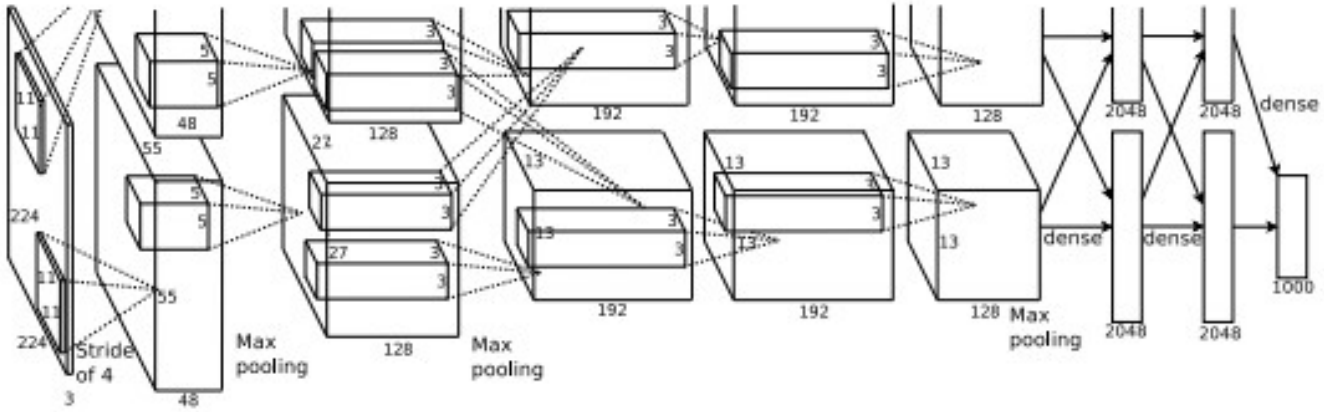

Krizhevsky, 2012

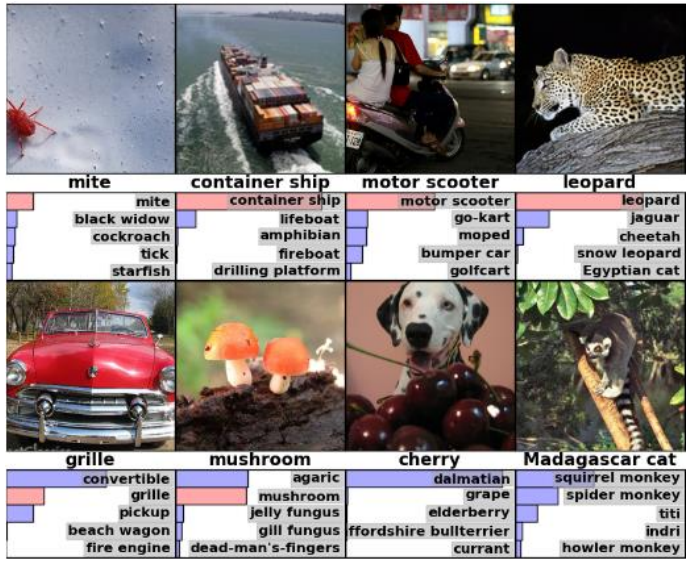

New deep learning era

- More data!

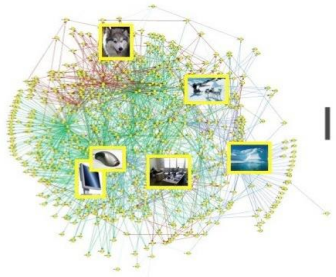

IMAGENET

- More computing power - GPUs!

- Better learning algorithms!

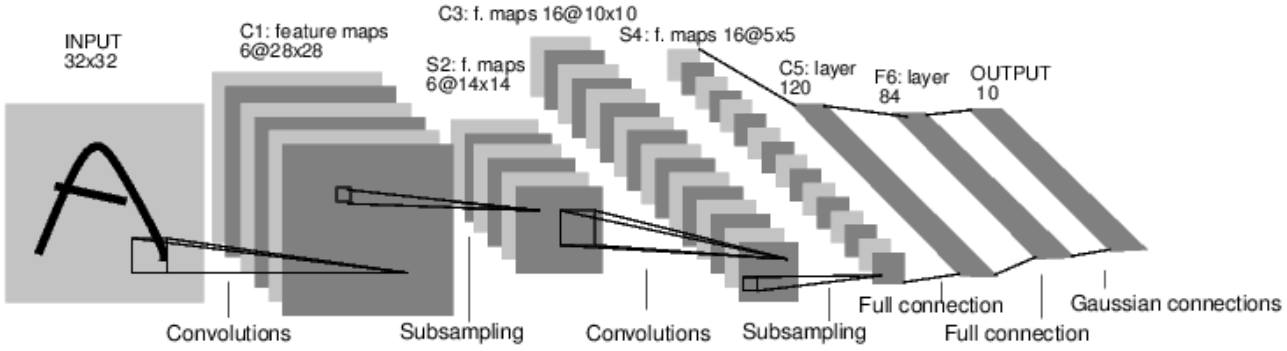

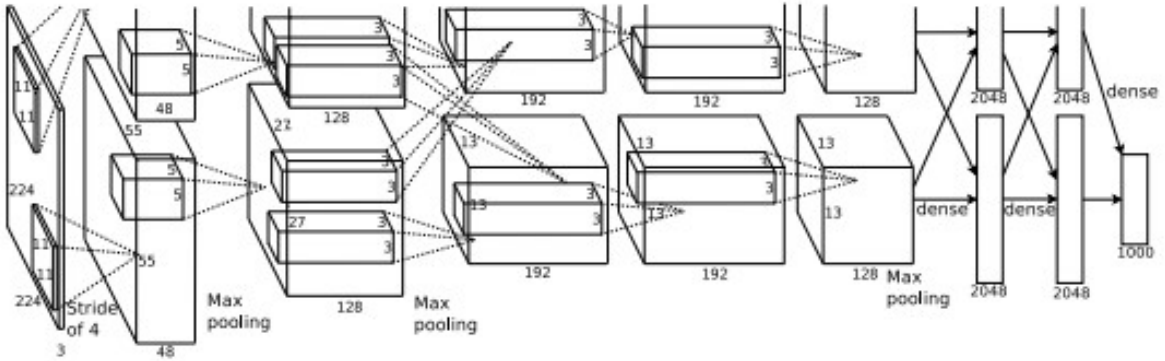

Dataset size

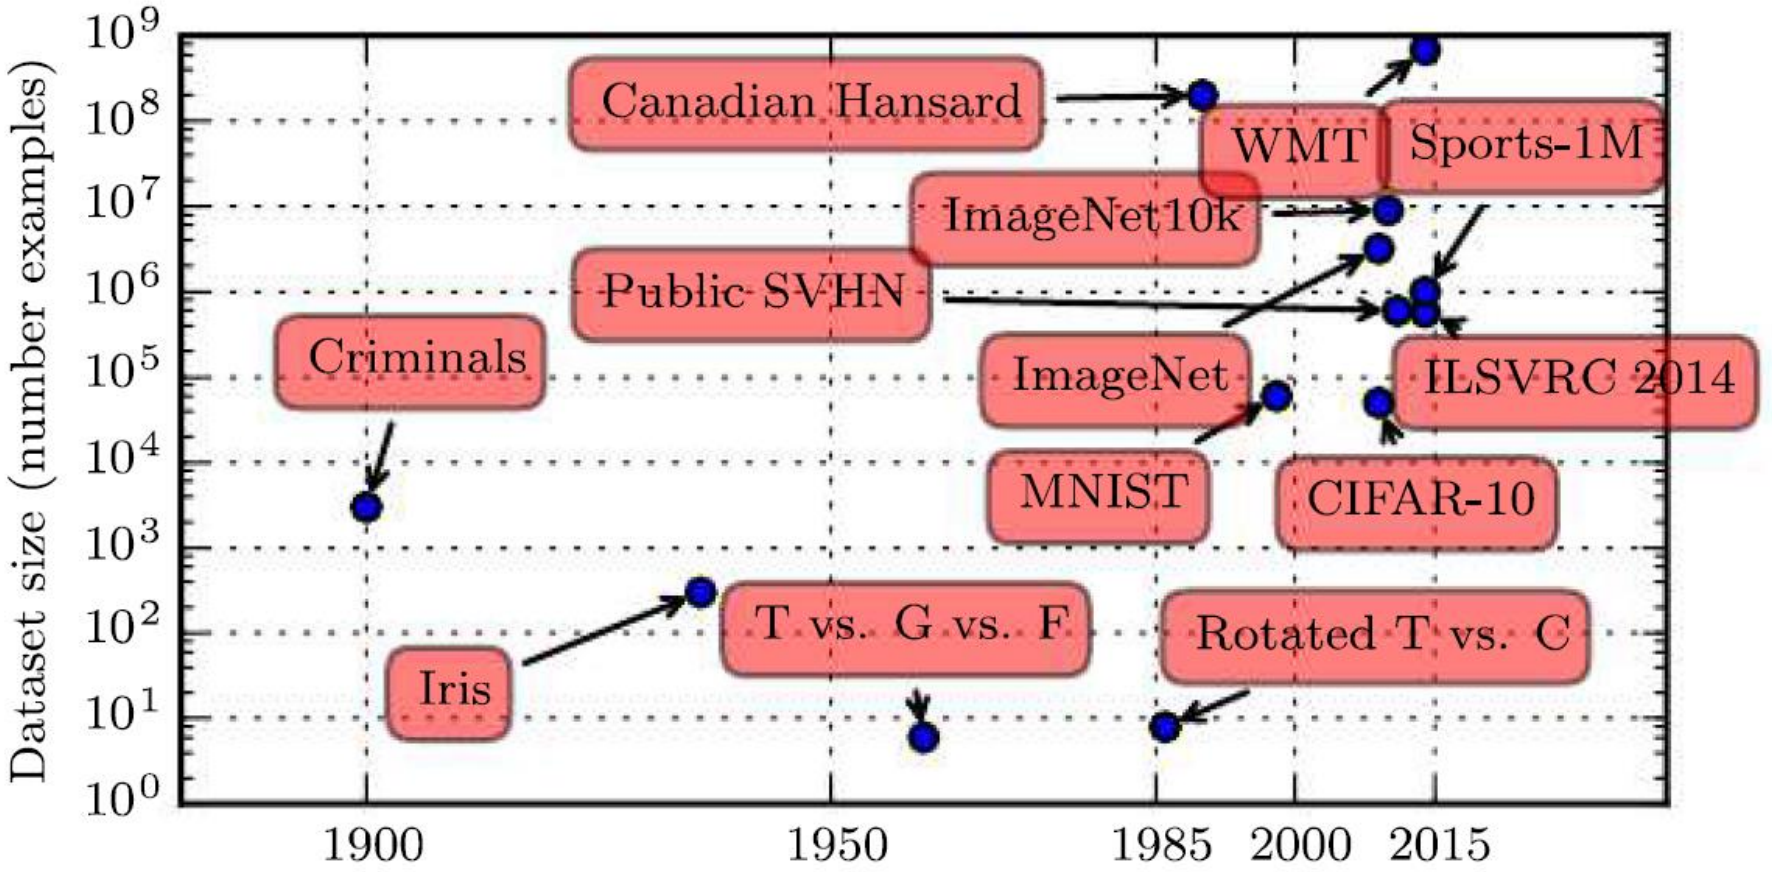

Goodfellow et al., *2016*

Neural network size

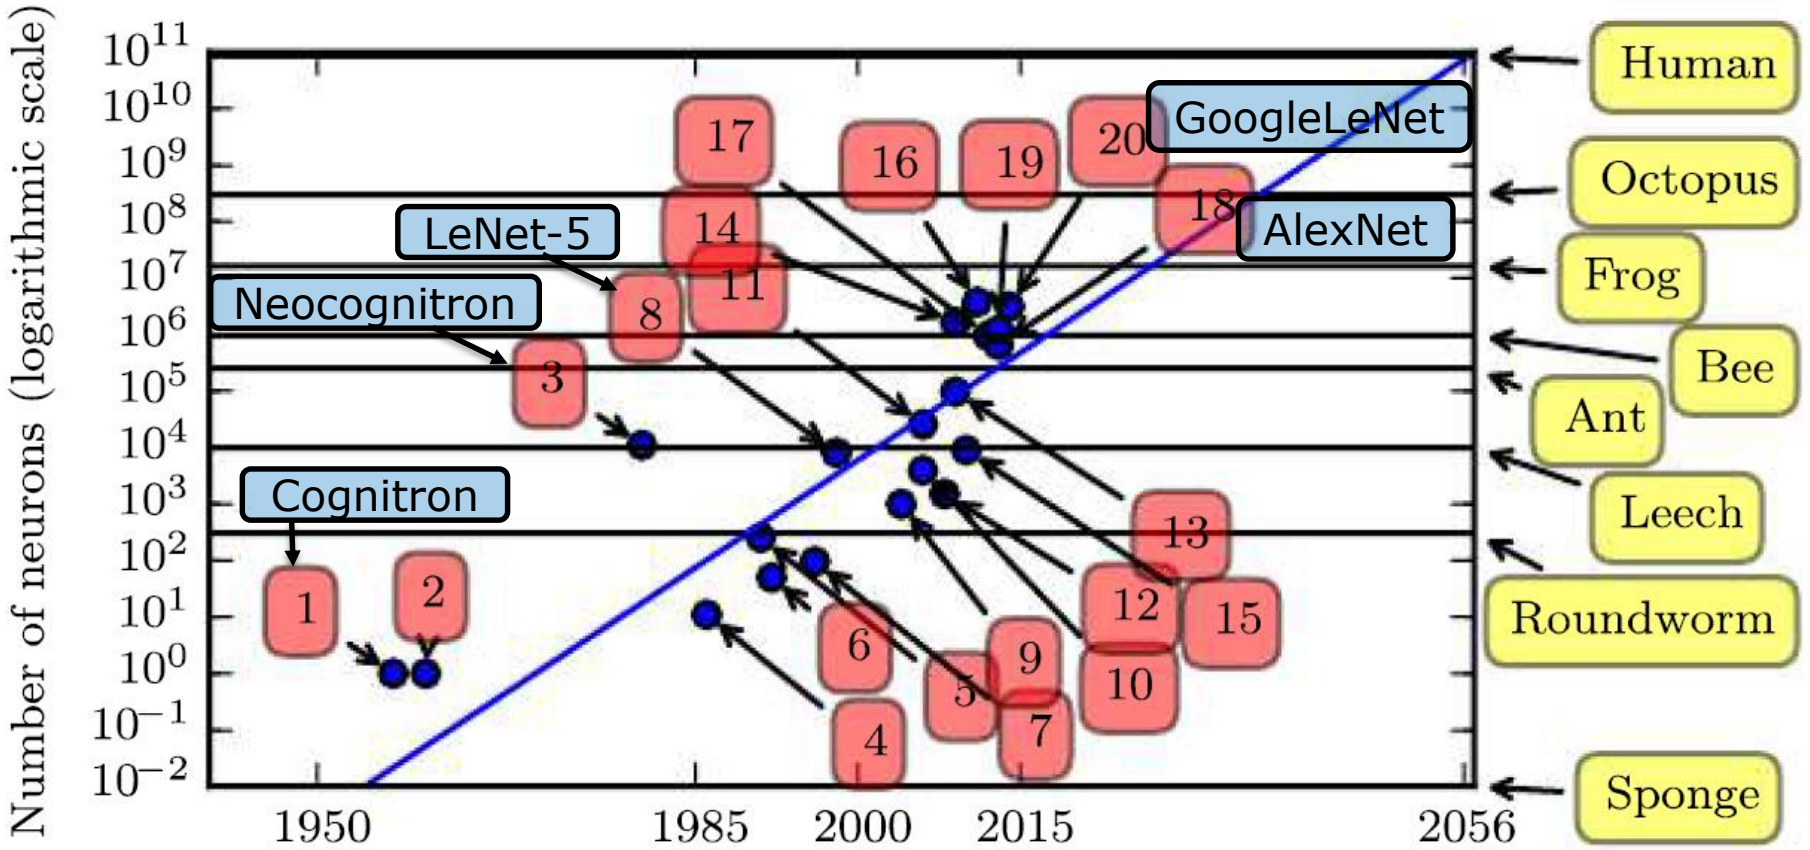

Deep learning – the main concept

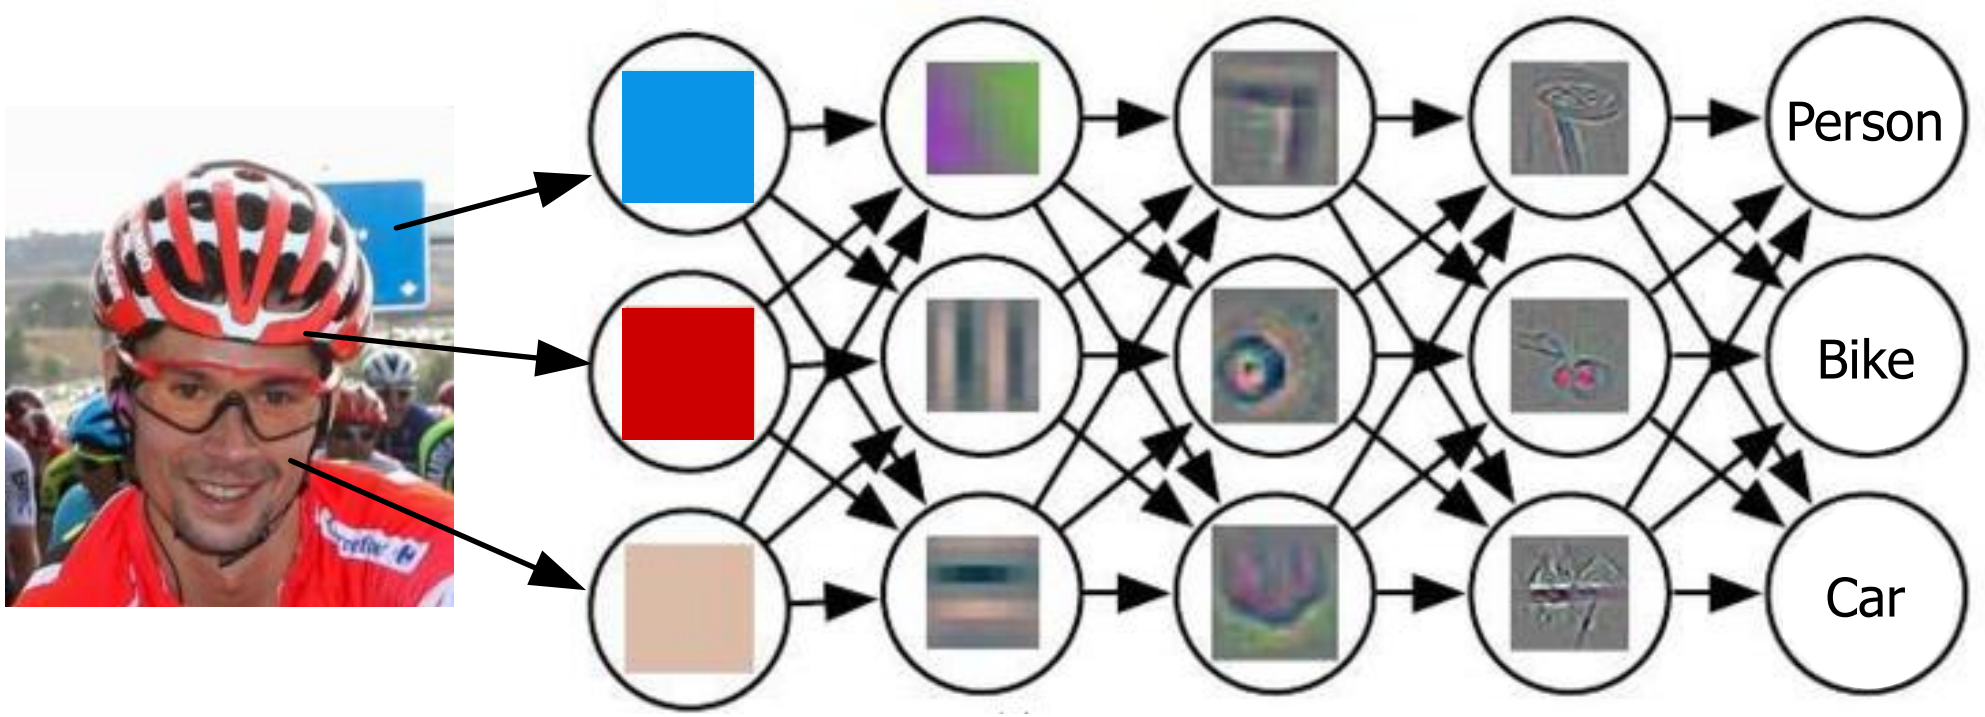

Representation learning

- Conventional two-stage computer vision/machine learning approach

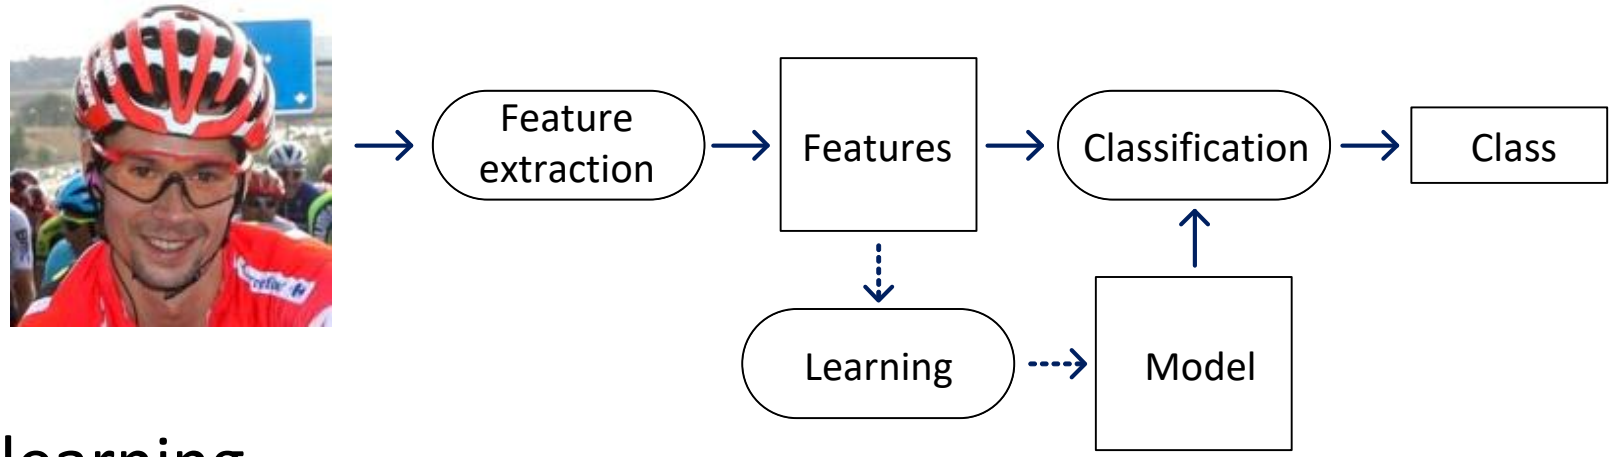

- Deep learning

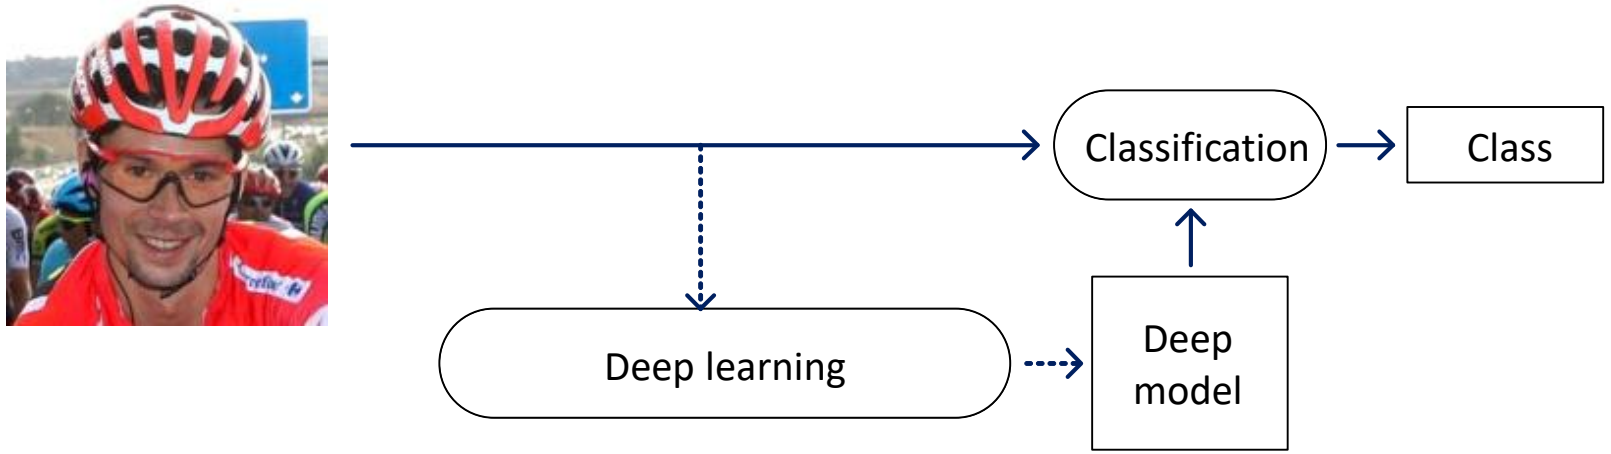

Deep learning and AI

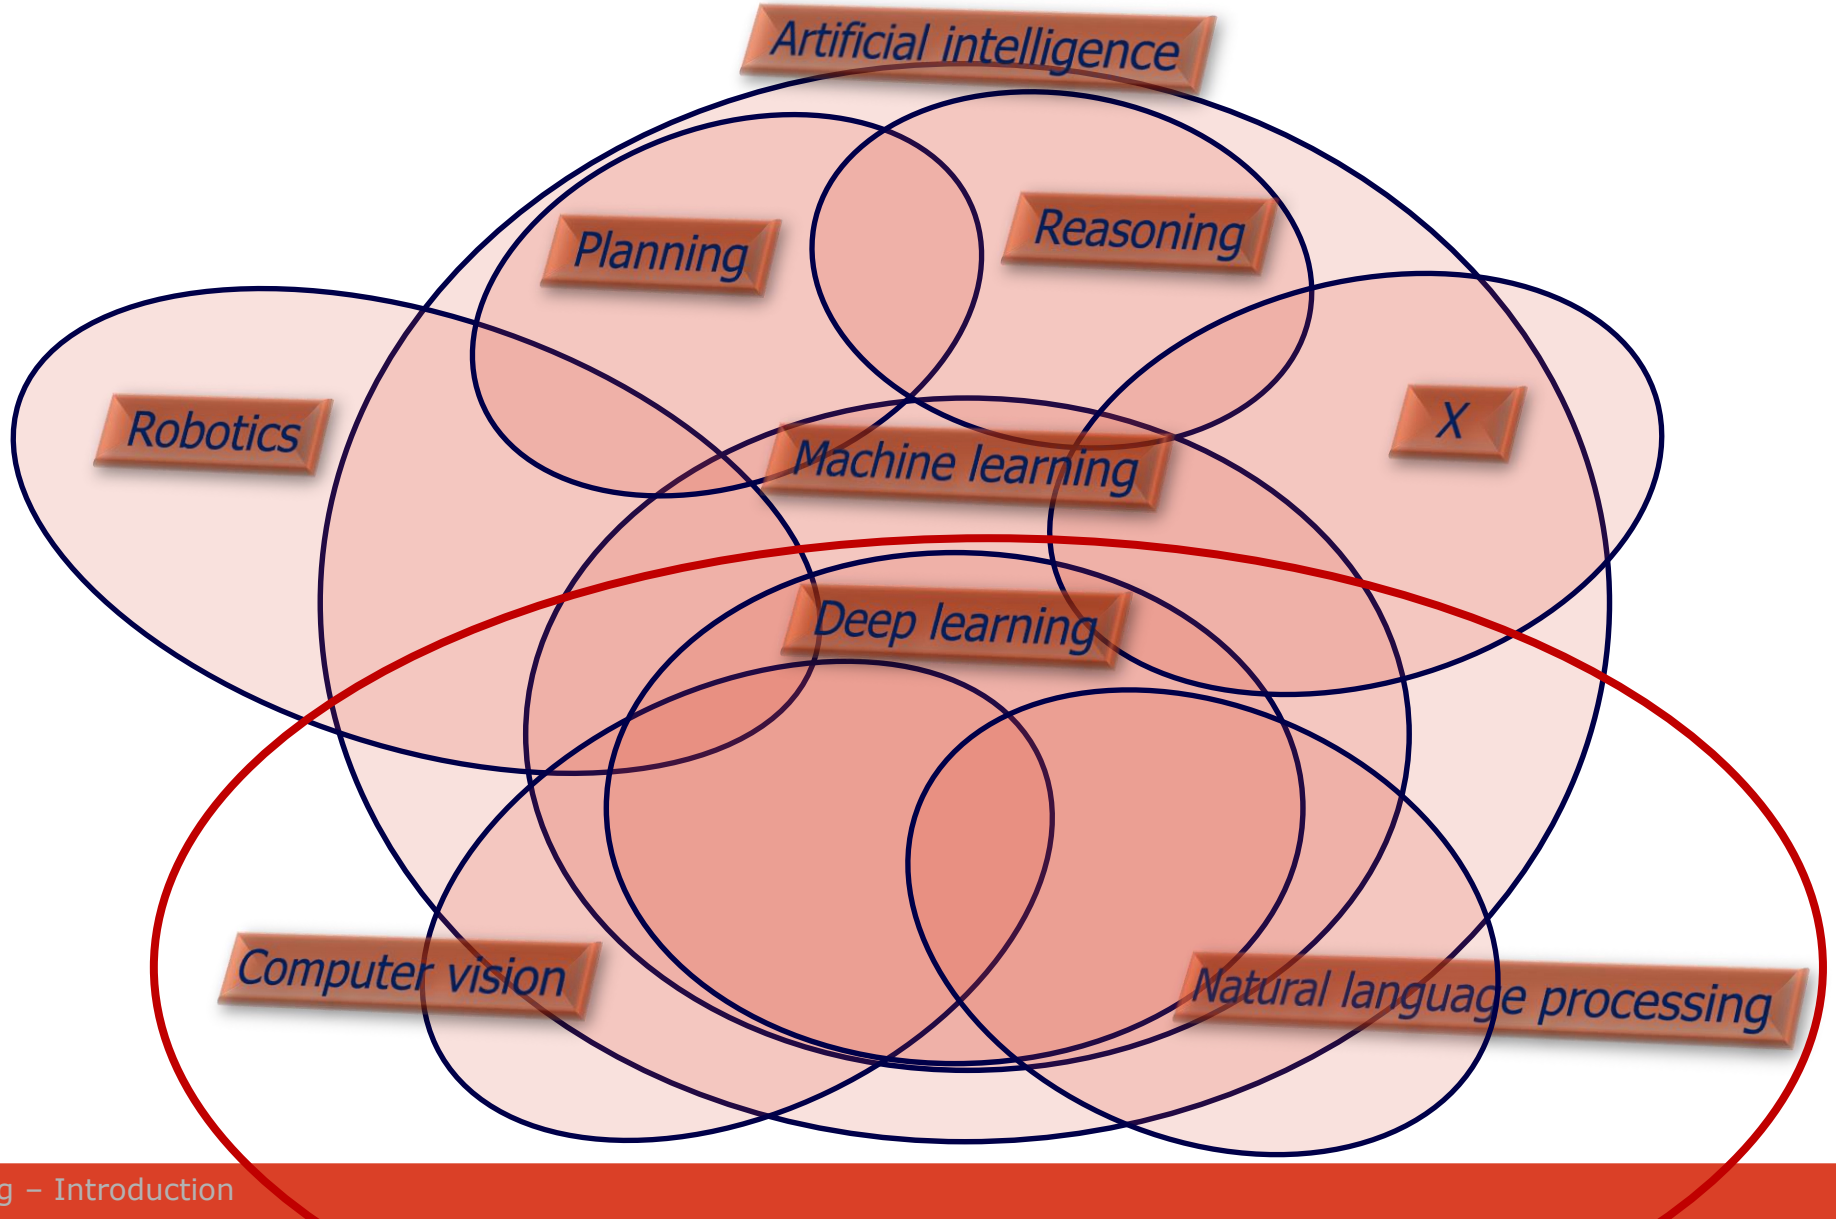

Deep learning and AI

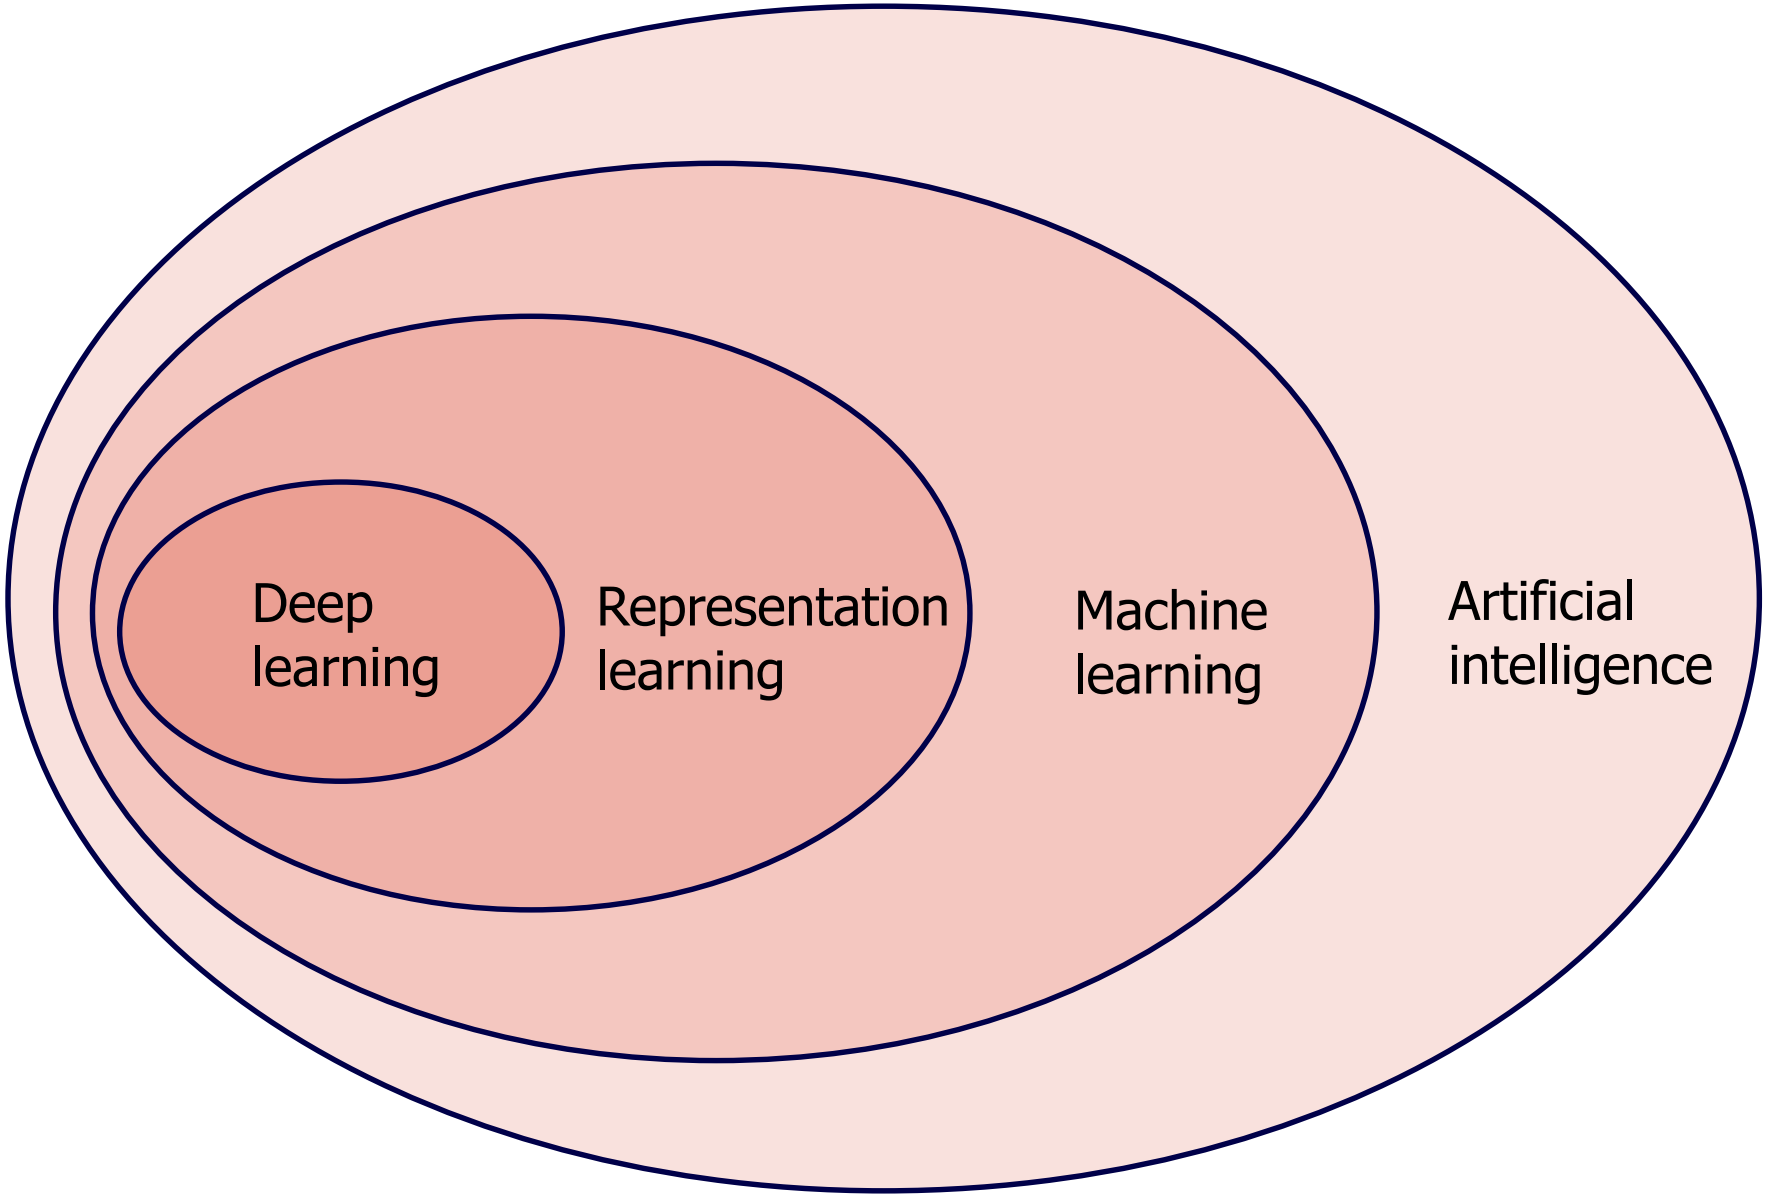

End to end learning

- Representations as well as classifier are being learned

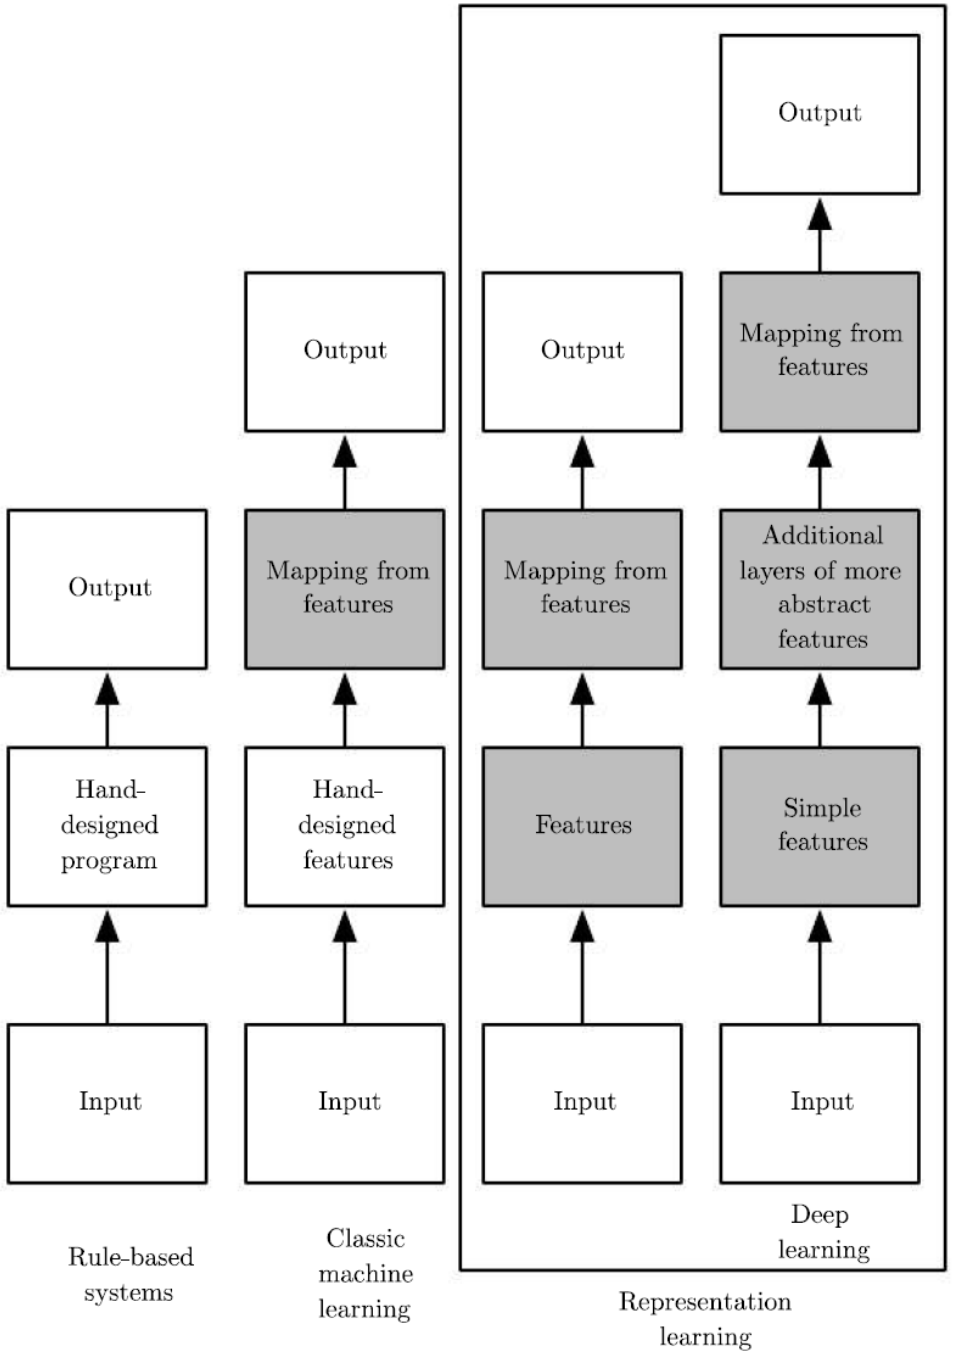

Typical solution

Korak 1: Zajem podatkov

Arhitecture

- Training the model
- Inference

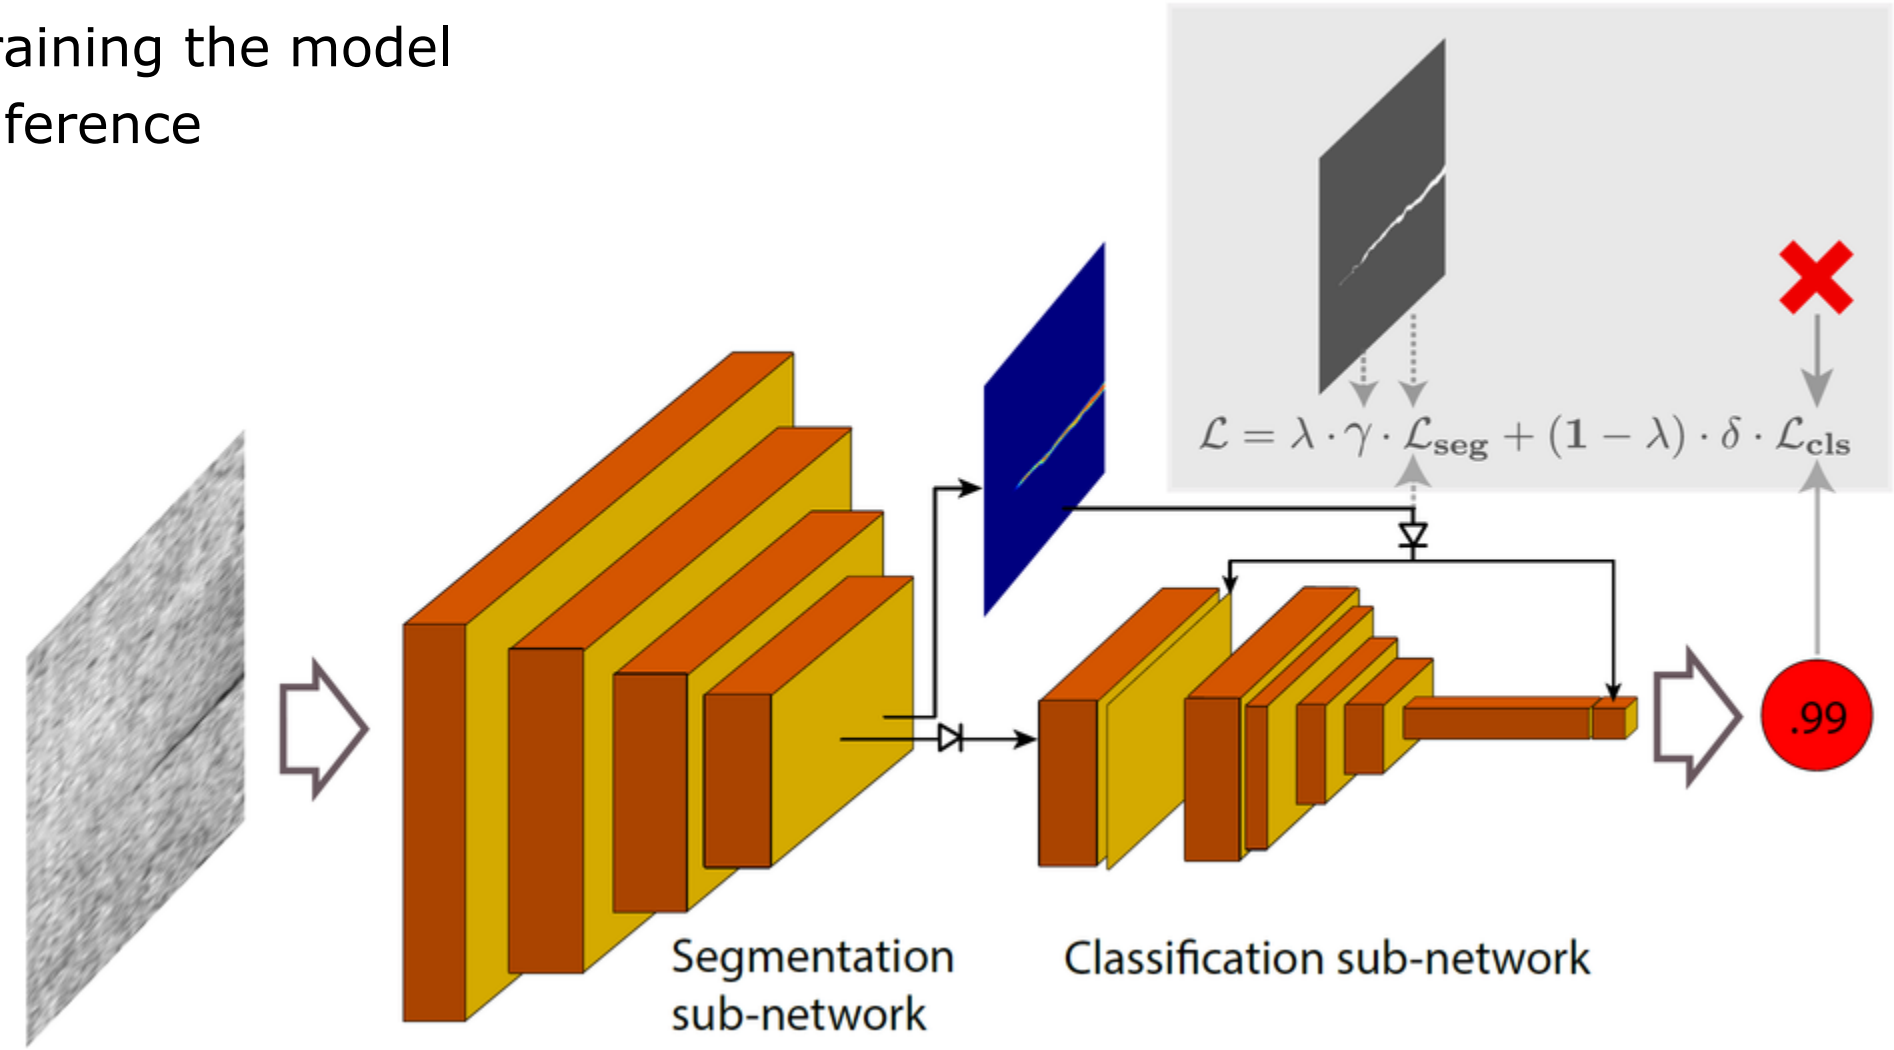

Detection of surface defects

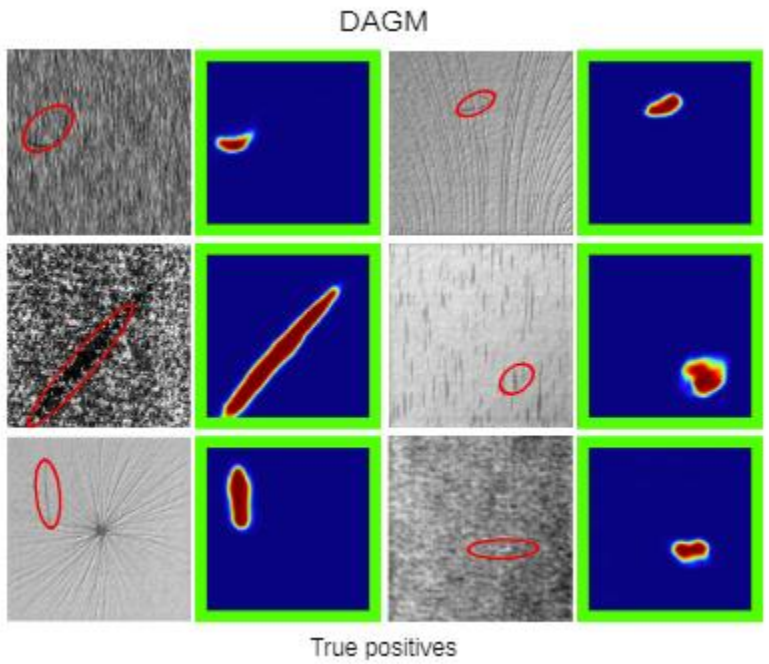

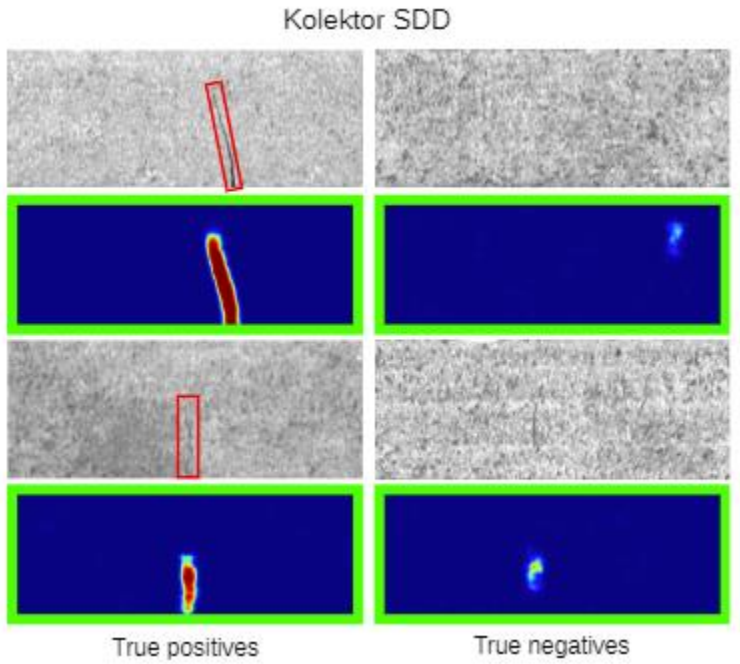

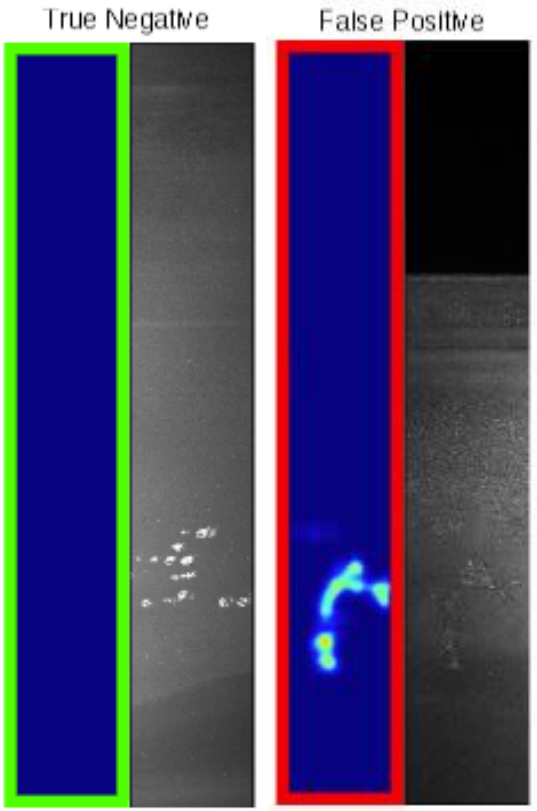

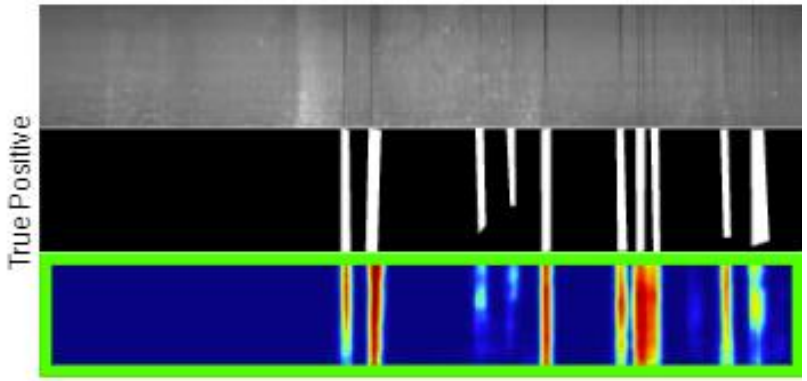

Object (traffic sign) detection

Obstacle detection on autonomous boat

USV equipped with different sensors:

- stereo camera
- IMU
- GPS
- compass

Segmentation based on RGB + IMU

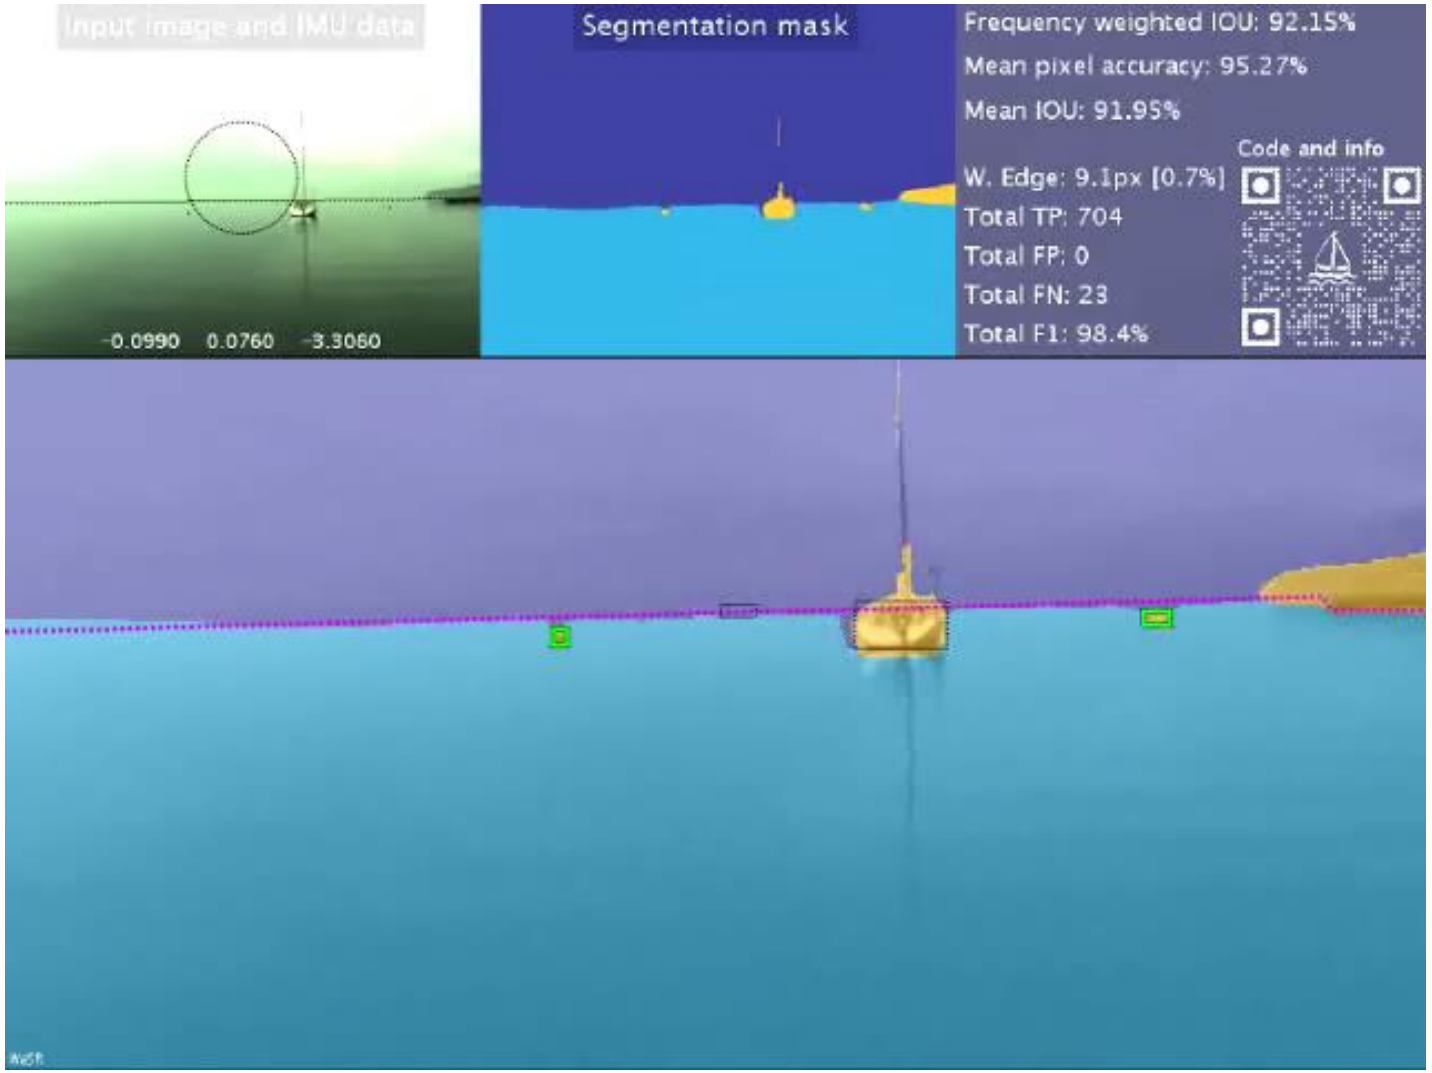

Ski jump style scoring

Judge 1	Judge 2	Judge 3	Judge 4	Judge 5	Virt.jud.
0,22	0,10	0,17	0,22	0,16	**0,21**

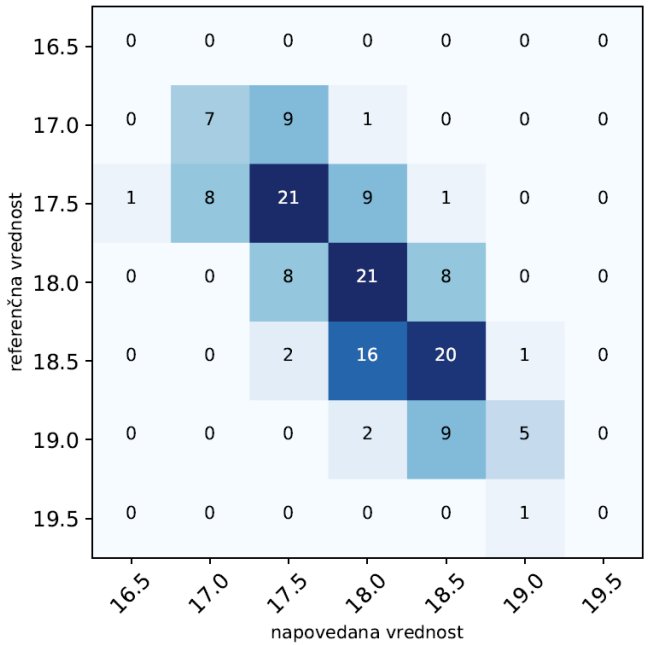

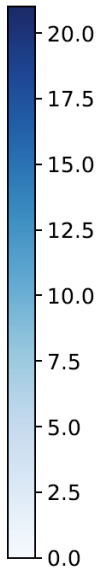

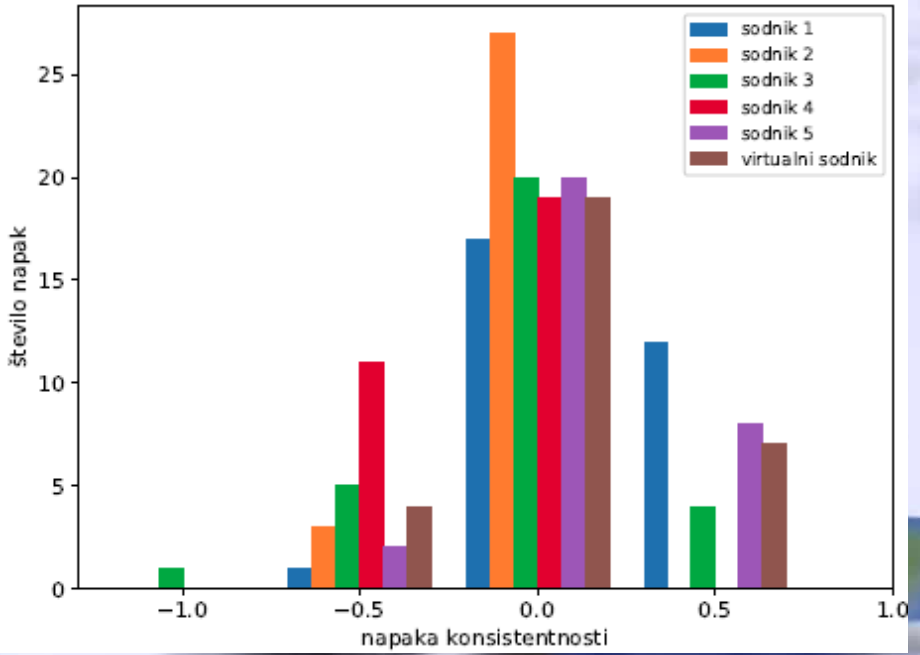

Sentiment analysis

Ni najboljša ampak za ceno je solidna, usb priključek je malo čudn ker zazna samo usb-2.0 in nima najbolj čvrstega stojala. ★★

Kamero uporabljam za šolske potrebe in sem popolnoma zadovoljna. Poceni, enostavna uporaba, dela popolnoma vredu. ★★★★★

Slika ni preveč dobra, je pa ok glede na ceno. ★★★

Odlična kamera za ta denar. postaviš, vtakneš v usb režo prižgeš računalnik in vse dela kot mora. Slika odlična zvok tudi. Nekaj težav ko sem jo priklopil med delovanjem računalnika. Reboot je vse rešil ★★★★

poceni web kamera, win10 ti sam namesti gonilnike. slika je obupna, vendar zadostljiva ce rabis zacasno kamero hitro. ★

Za podobno ceno so tudi externe webkamere s HD tehniko ★★

V slabi svetlobi bolj slaba slika, drugač pa za silo v redu kamera. ★★★★★

Na kameri mikrofon prekinja. zato jo bom reklamiral. ★

V specifikaciji piše, da dela tudi na USB 2.0, a se izkaže, da je nestabilna, ker odvzame preveč elektrike in je potrebno za dobro delovanje kupiti vmes usb hub z dodatnim napajanjem. ★★★

Slika ok glede na denar. Mikrofon neuporaben. Drzalo neuporabno. ★★

Priklopljena na stacionarni računalnik, deluje super + vgrajen mikrofon. Idealna kombinacija za nadgradnjo računalnika. ★★★★★

Slaba kvaliteta barvni spekter kamere na nuli kk pride malo vec svetlobe zravn vglavnem skoda 15ih eurov ★

Image generation

Deep reinforcement learning

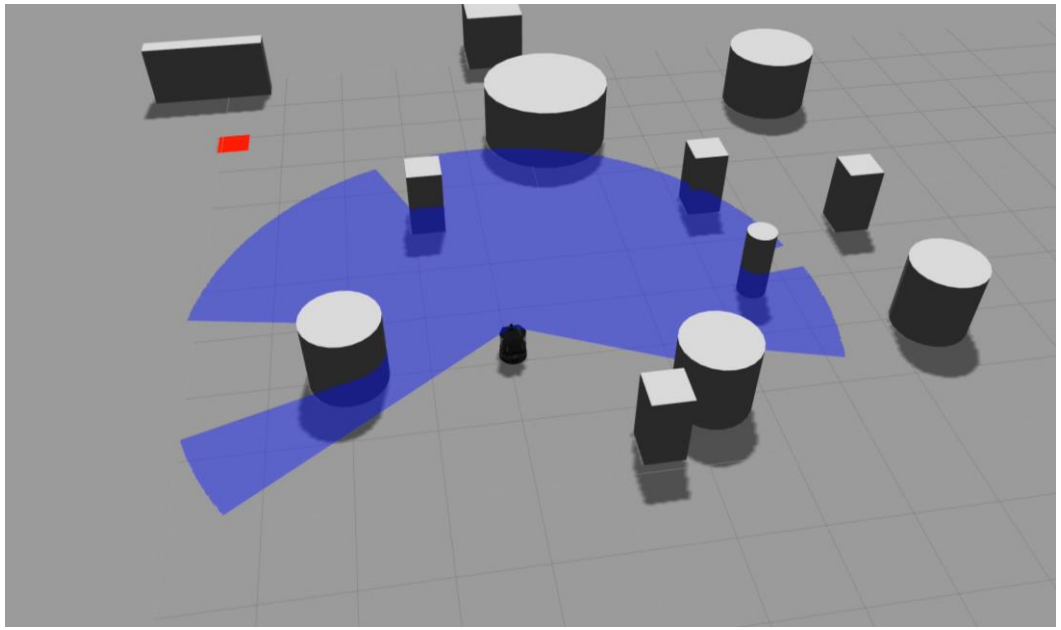

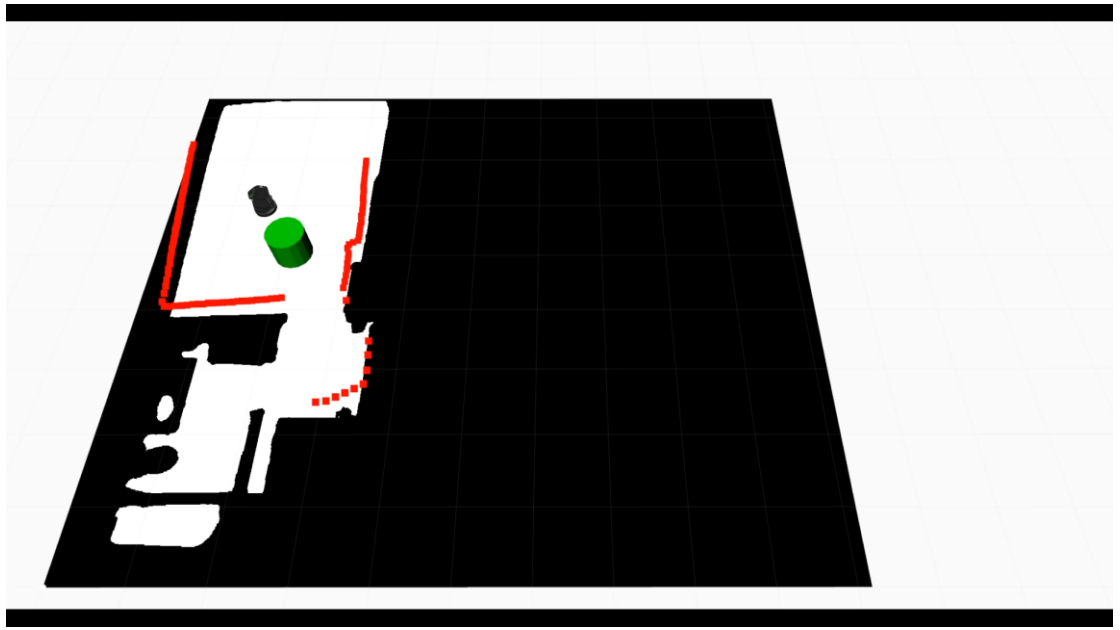

Problem solving

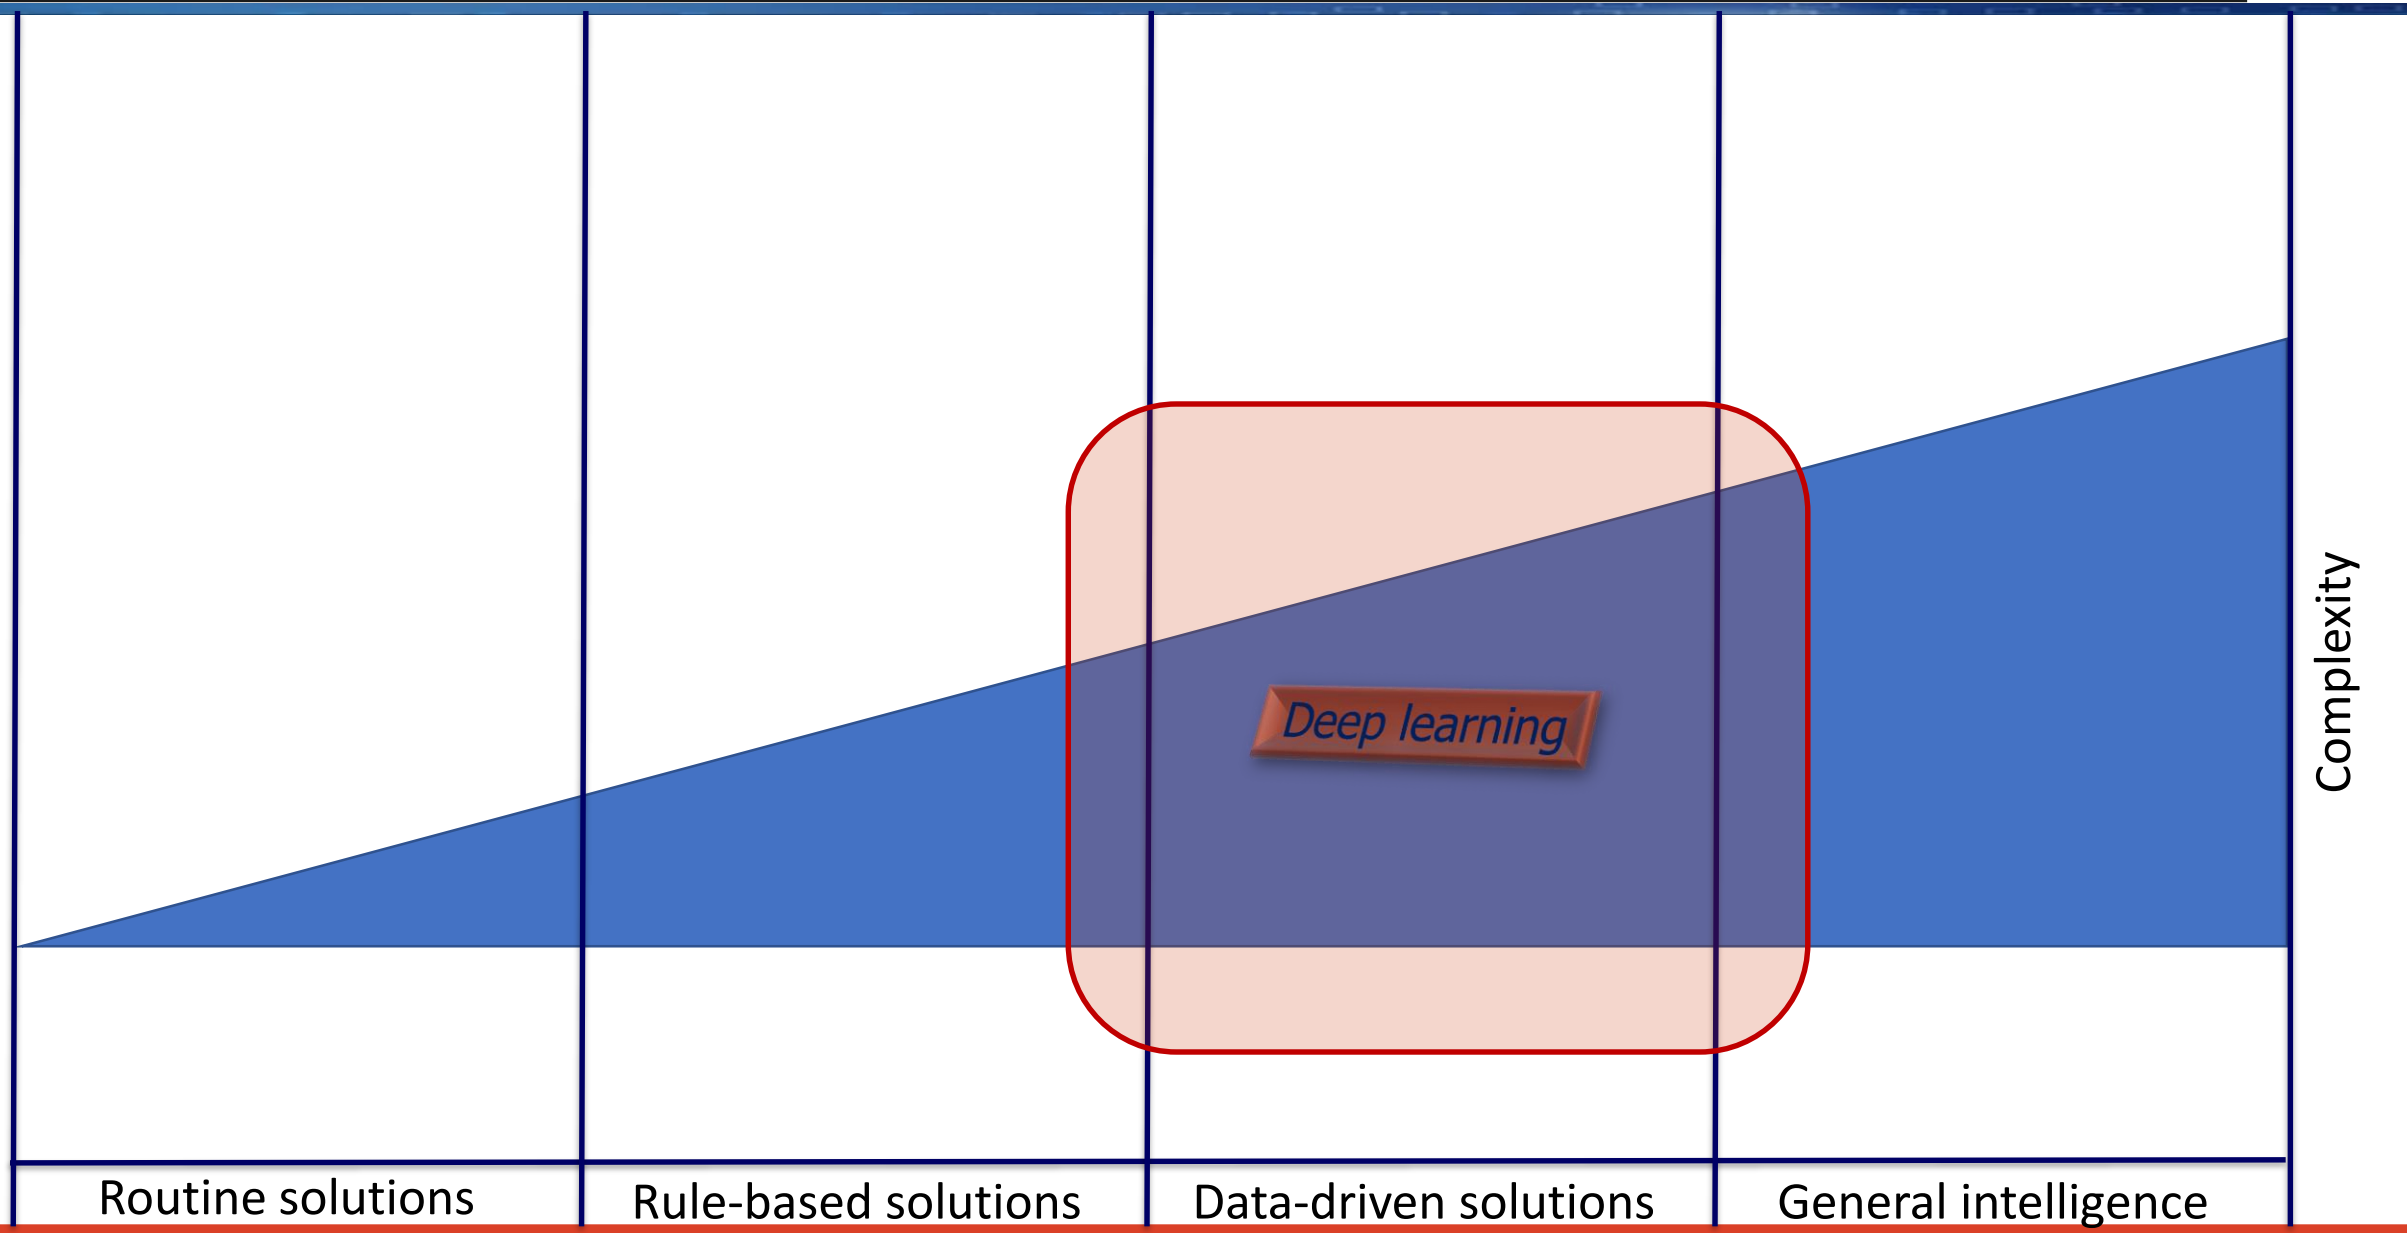

Development of deep learning solutions

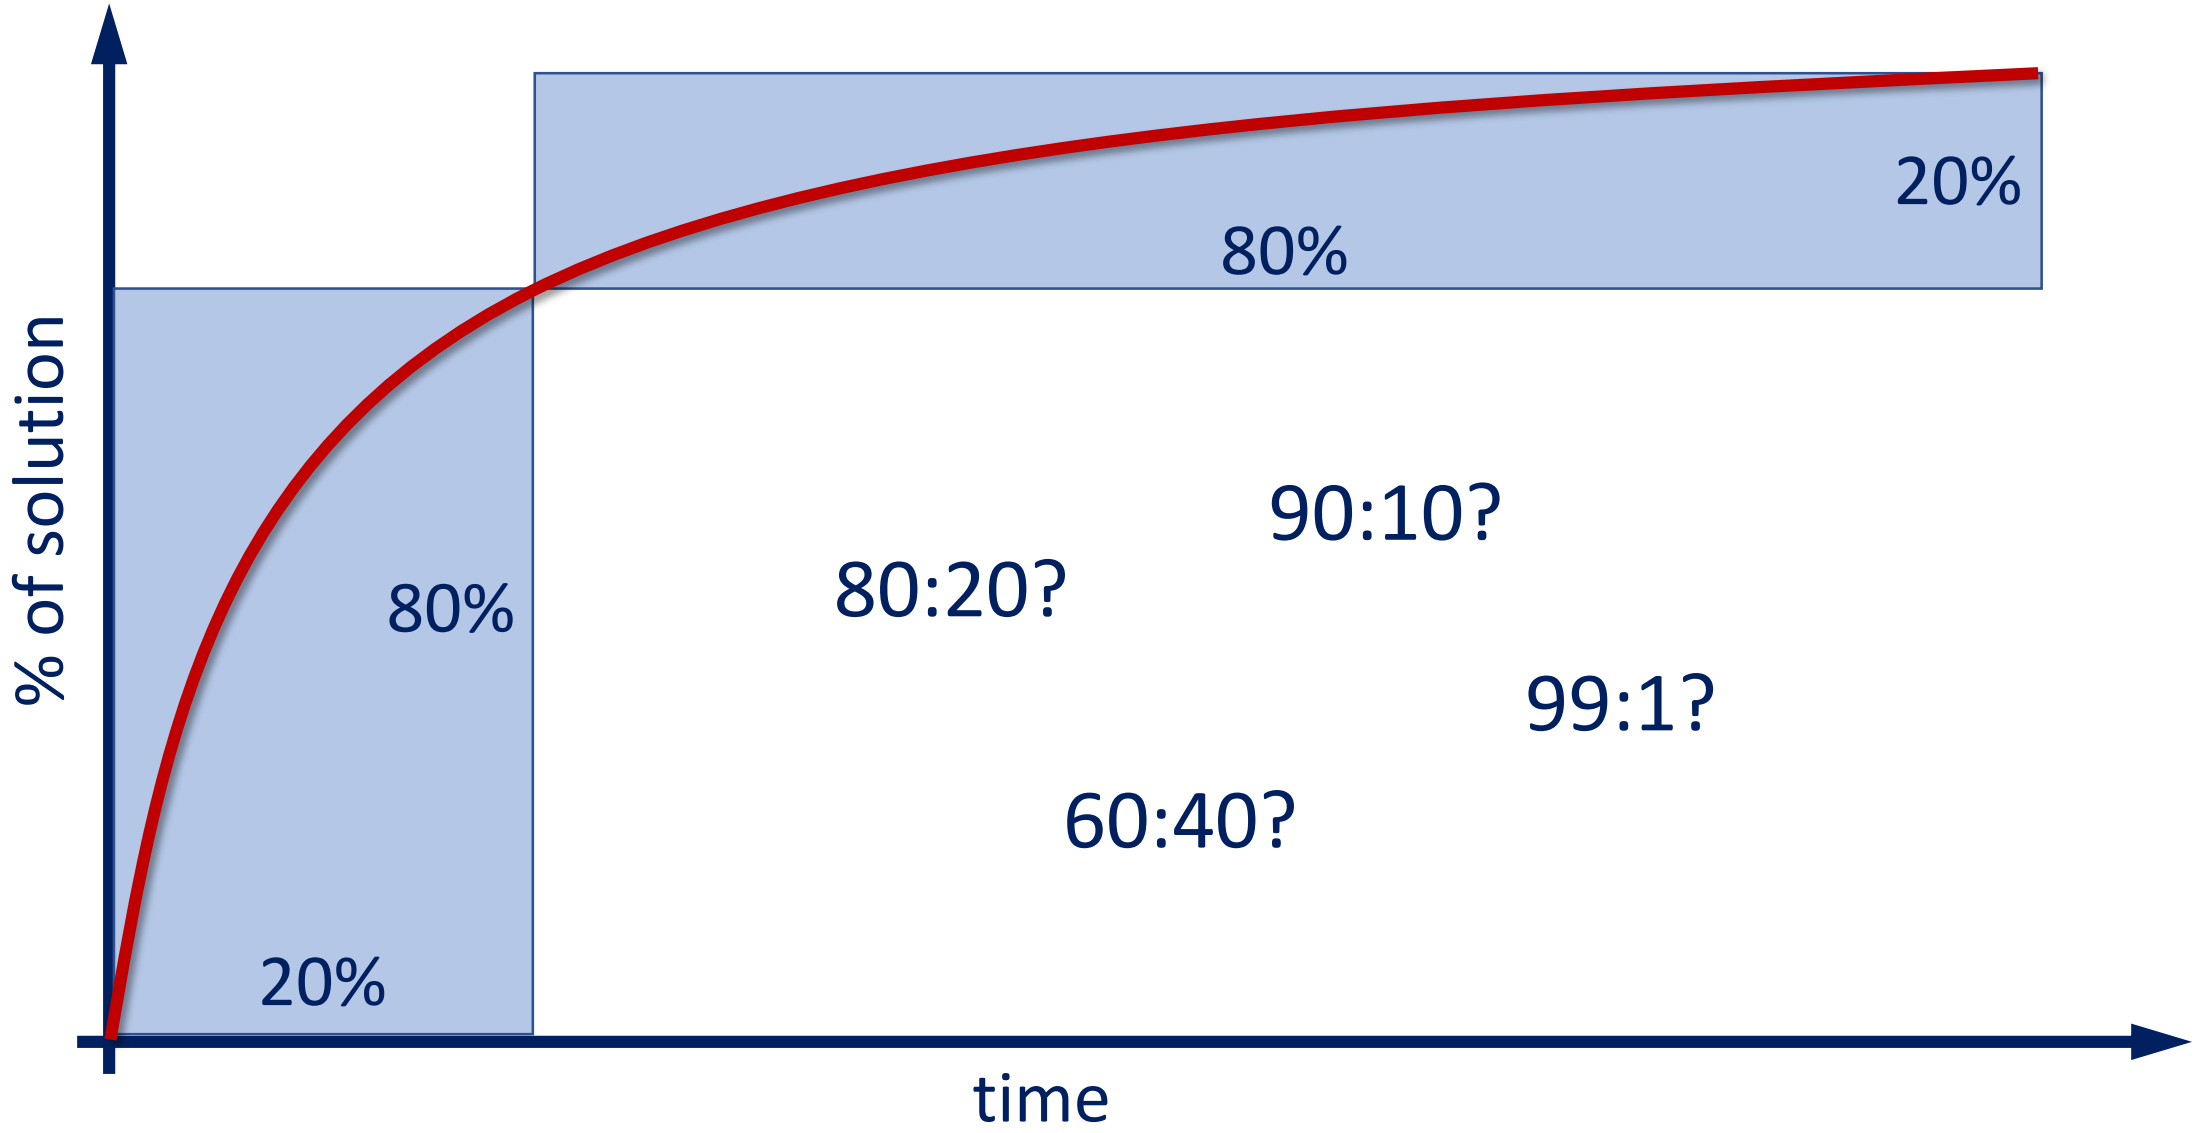

Knowledge and experience count

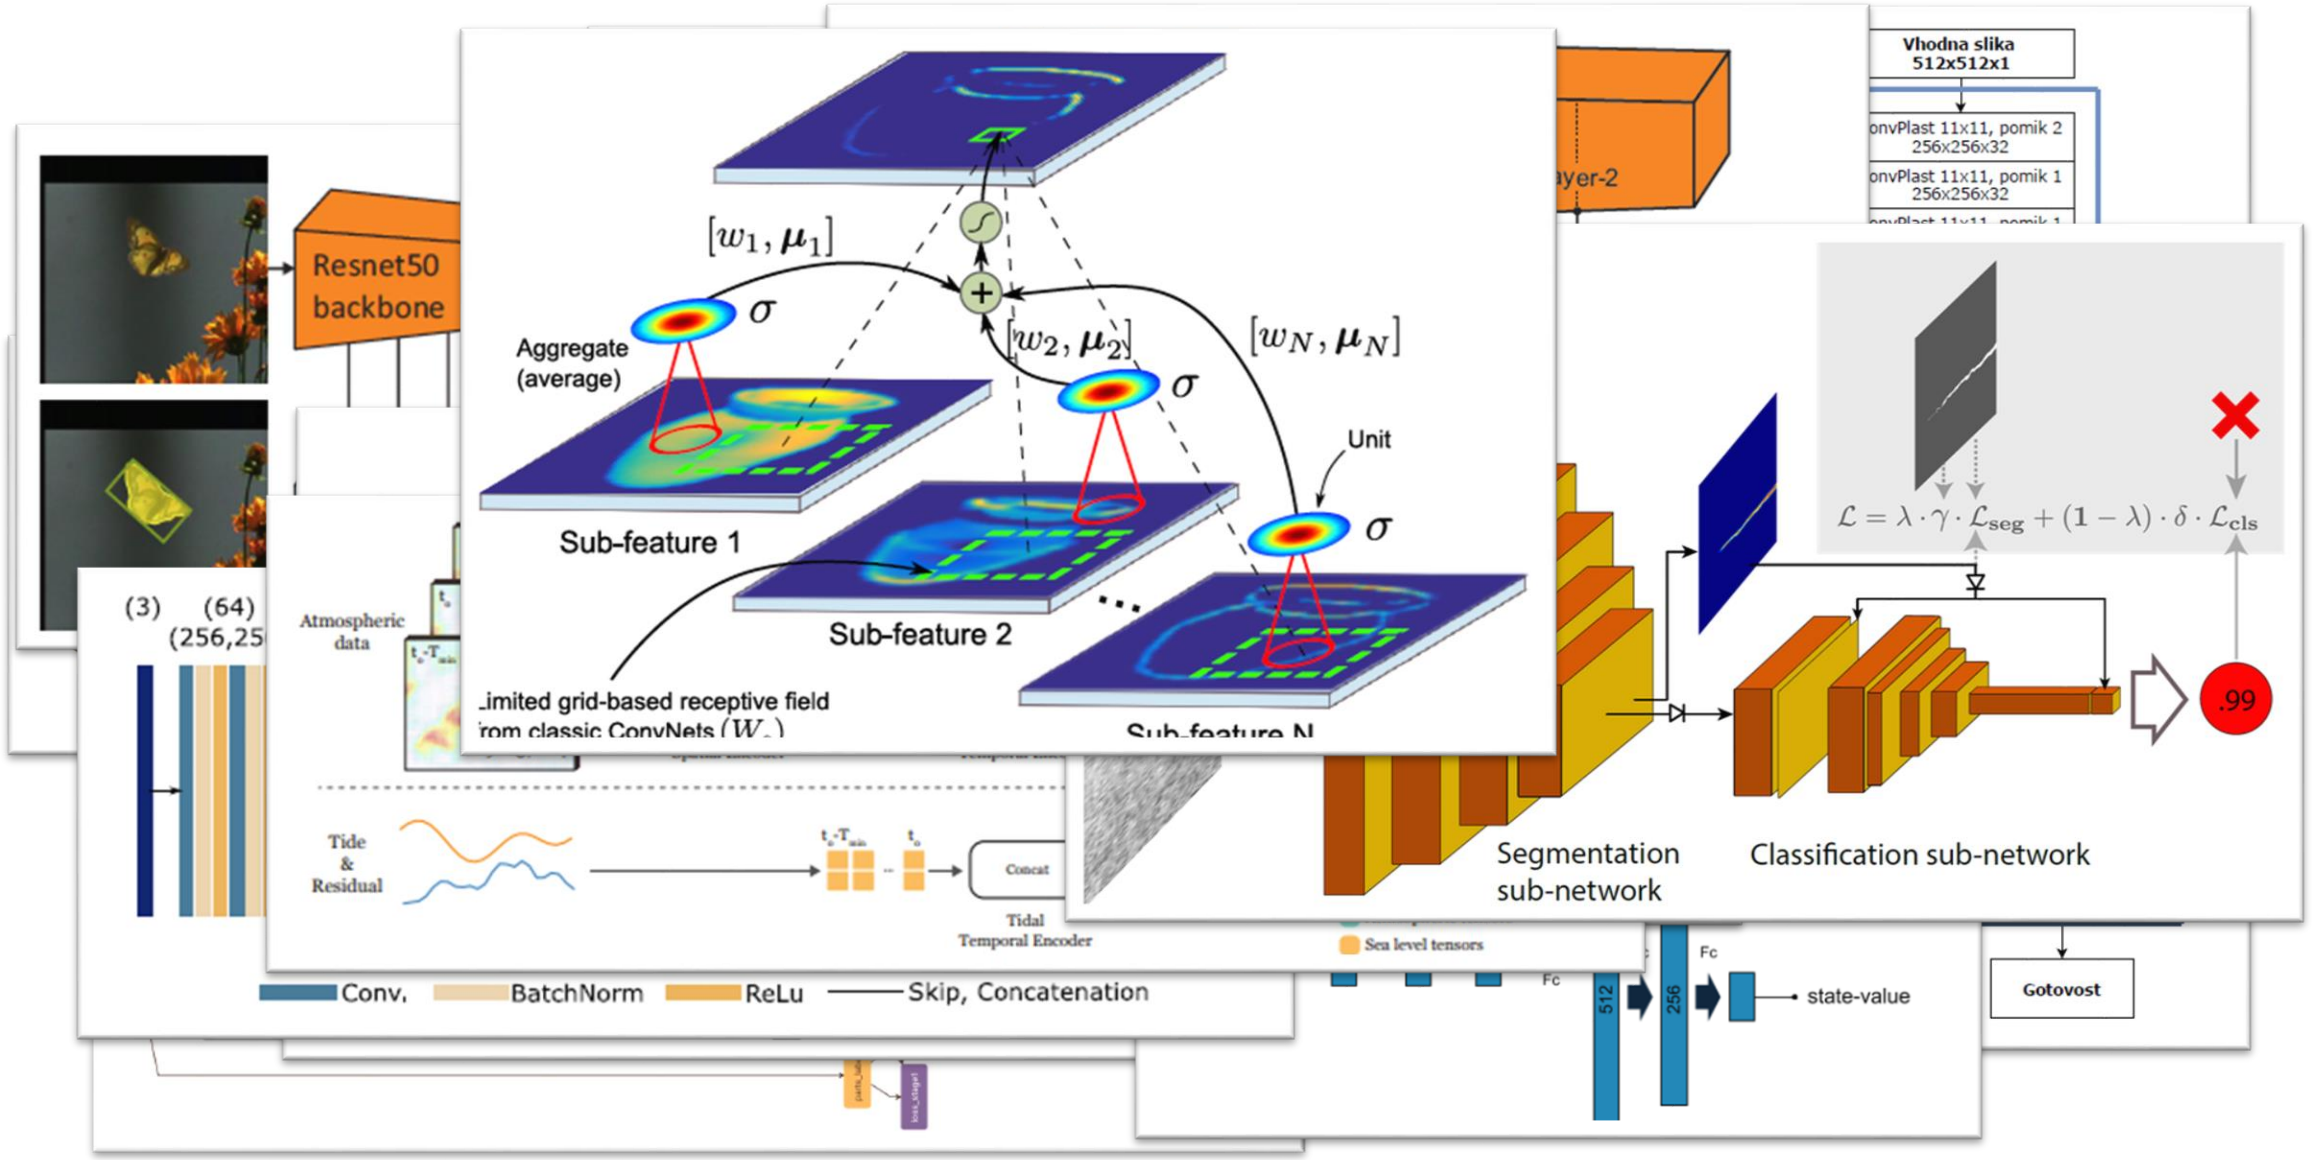

Acknowledgements

Literature:

- Ian Goodfellow and Yoshua Bengio and Aaron Courville,
  Deep Learning,
  MIT Press, 2016
  http://www.deeplearningbook.org/

- UL FRI LUVSS