Introduction to Deep Learning

The Team

Lecturers

Prof. Dr. Laura Leal-Taixé

Prof. Dr. Matthias Niessner

Tutors

Tim Meinhardt

Ji Hou

Andreas Rössler

What is Computer Vision?

- First defined in the 60s in artificial intelligence groups
- “Mimic the human visual system”
- Center block of robotic intelligence

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.

July 7, 1966

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

Some decades later…

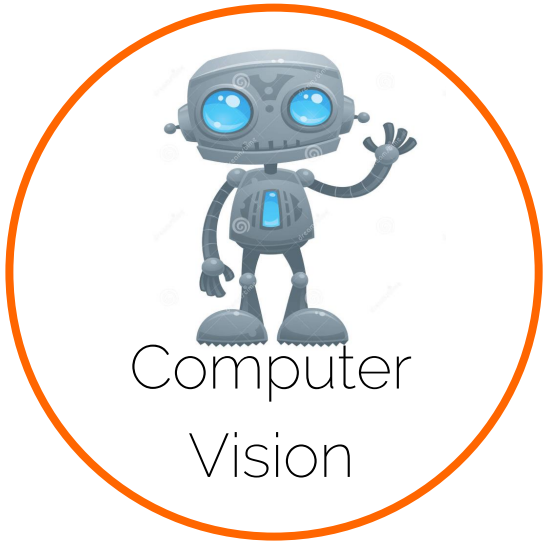

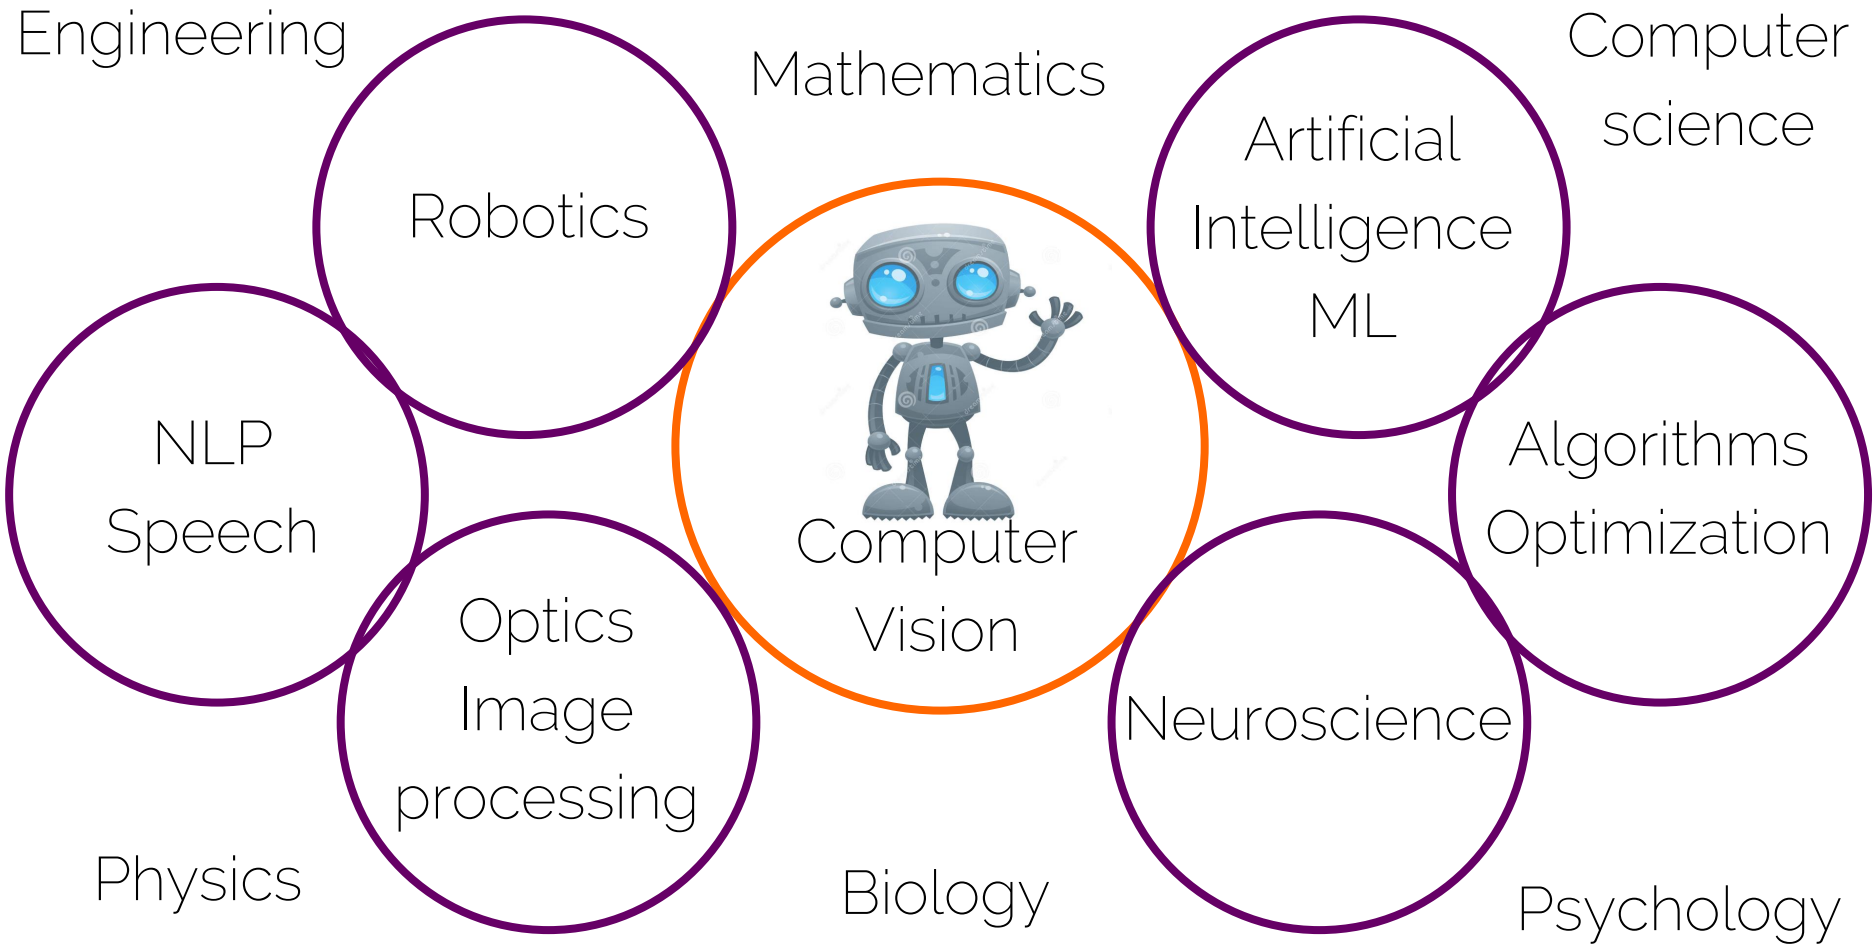
Engineering
Mathematics
Computer science
Robotics
Artificial Intelligence ML
NLP Speech
Computer Vision
Algorithms Optimization
Optics Image processing
Neuroscience
Physics
Biology
Psychology

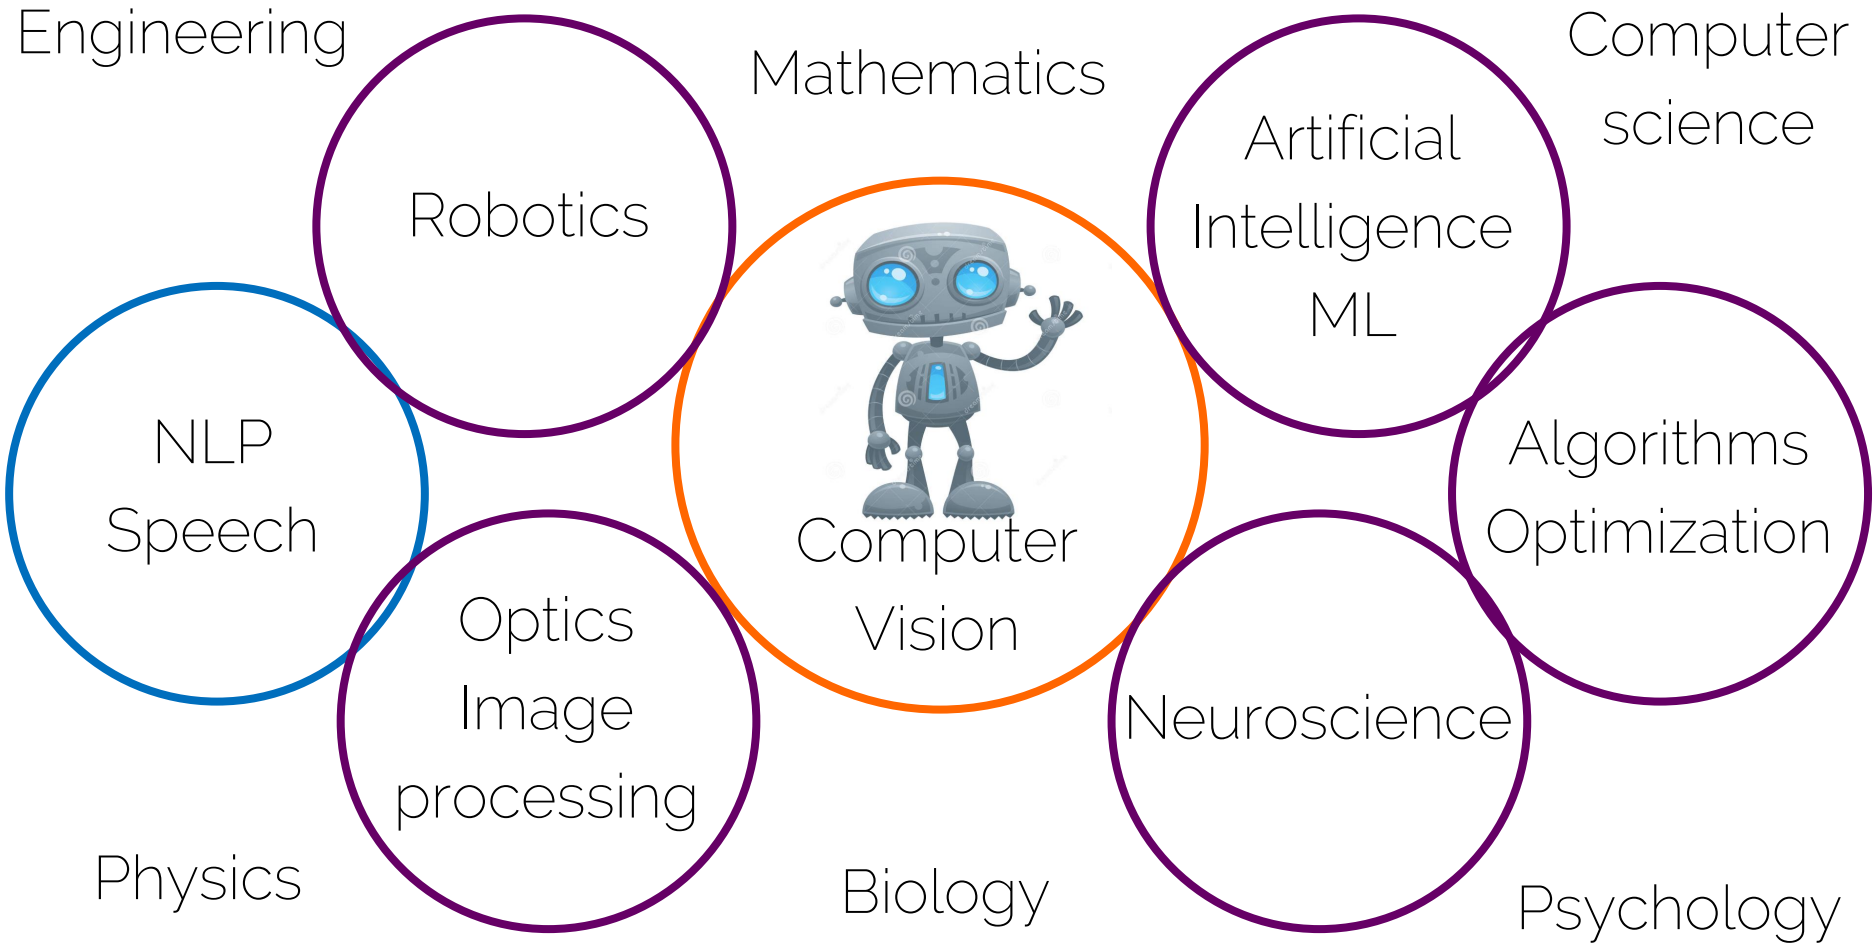
Engineering
Mathematics
Computer science
Robotics
Artificial Intelligence ML
NLP Speech
Computer Vision
Algorithms Optimization
Optics Image processing
Neuroscience
Physics
Biology
Psychology

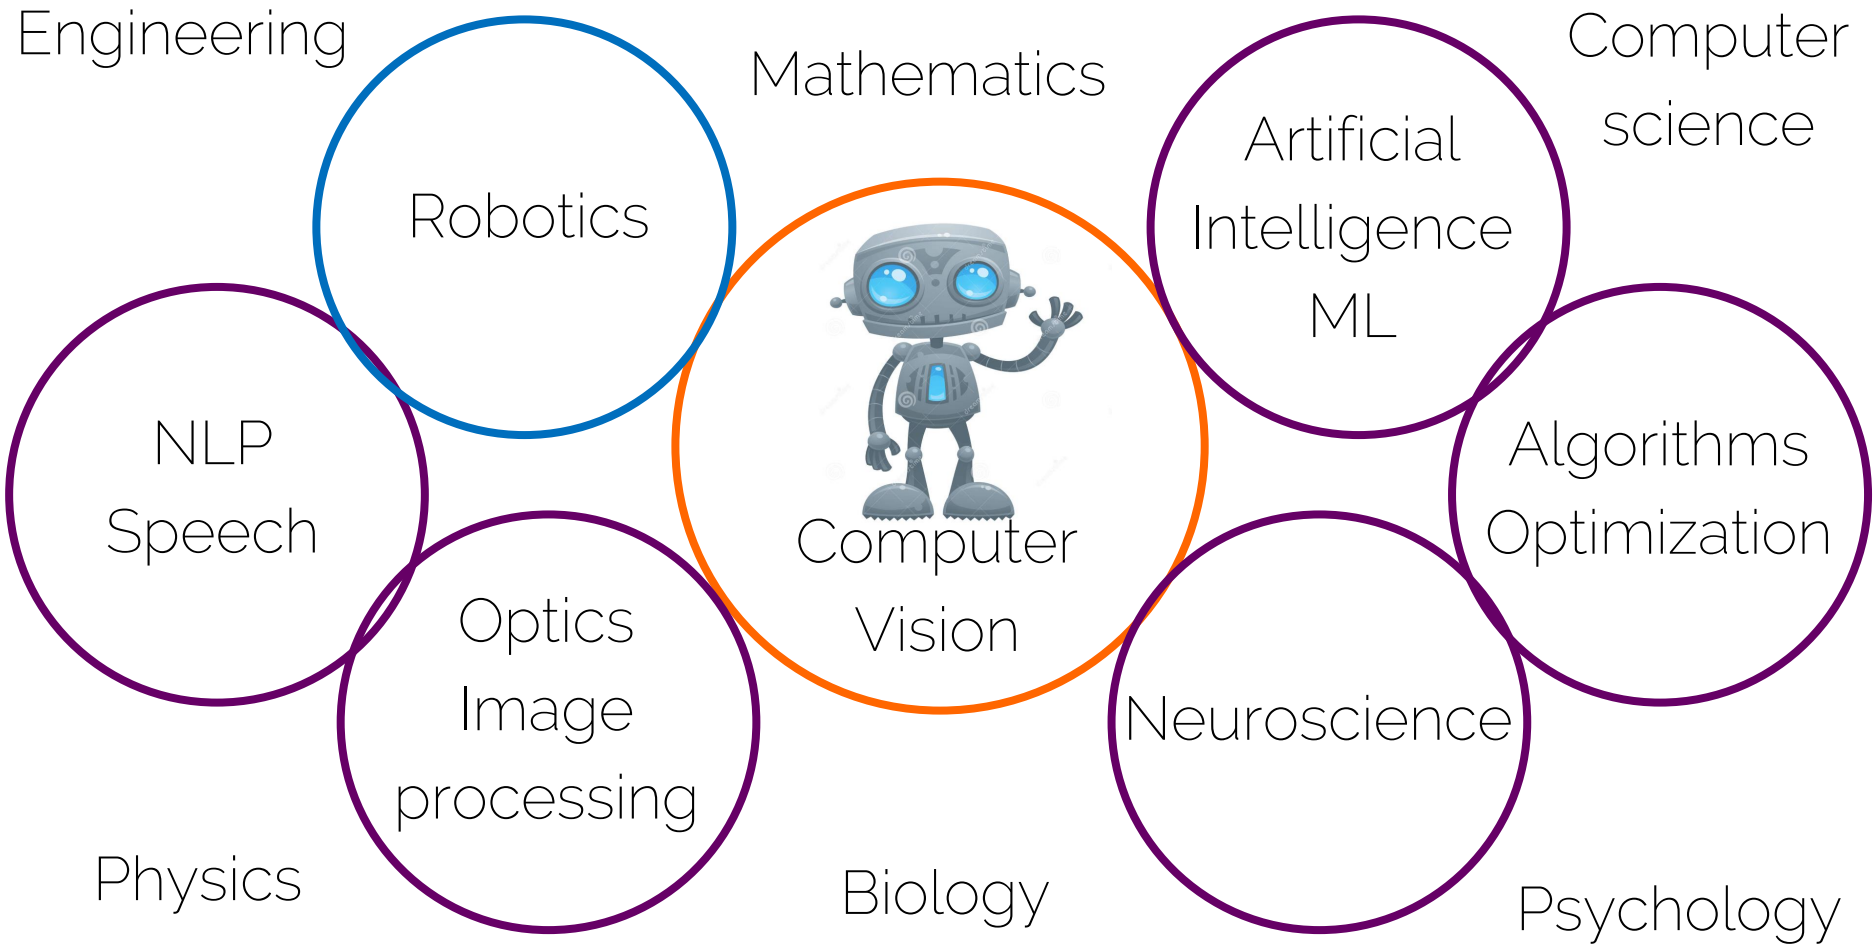
Engineering
Mathematics
Computer science
Robotics
Artificial Intelligence ML
NLP Speech
Computer Vision
Algorithms Optimization
Optics Image processing
Neuroscience
Physics
Biology
Psychology

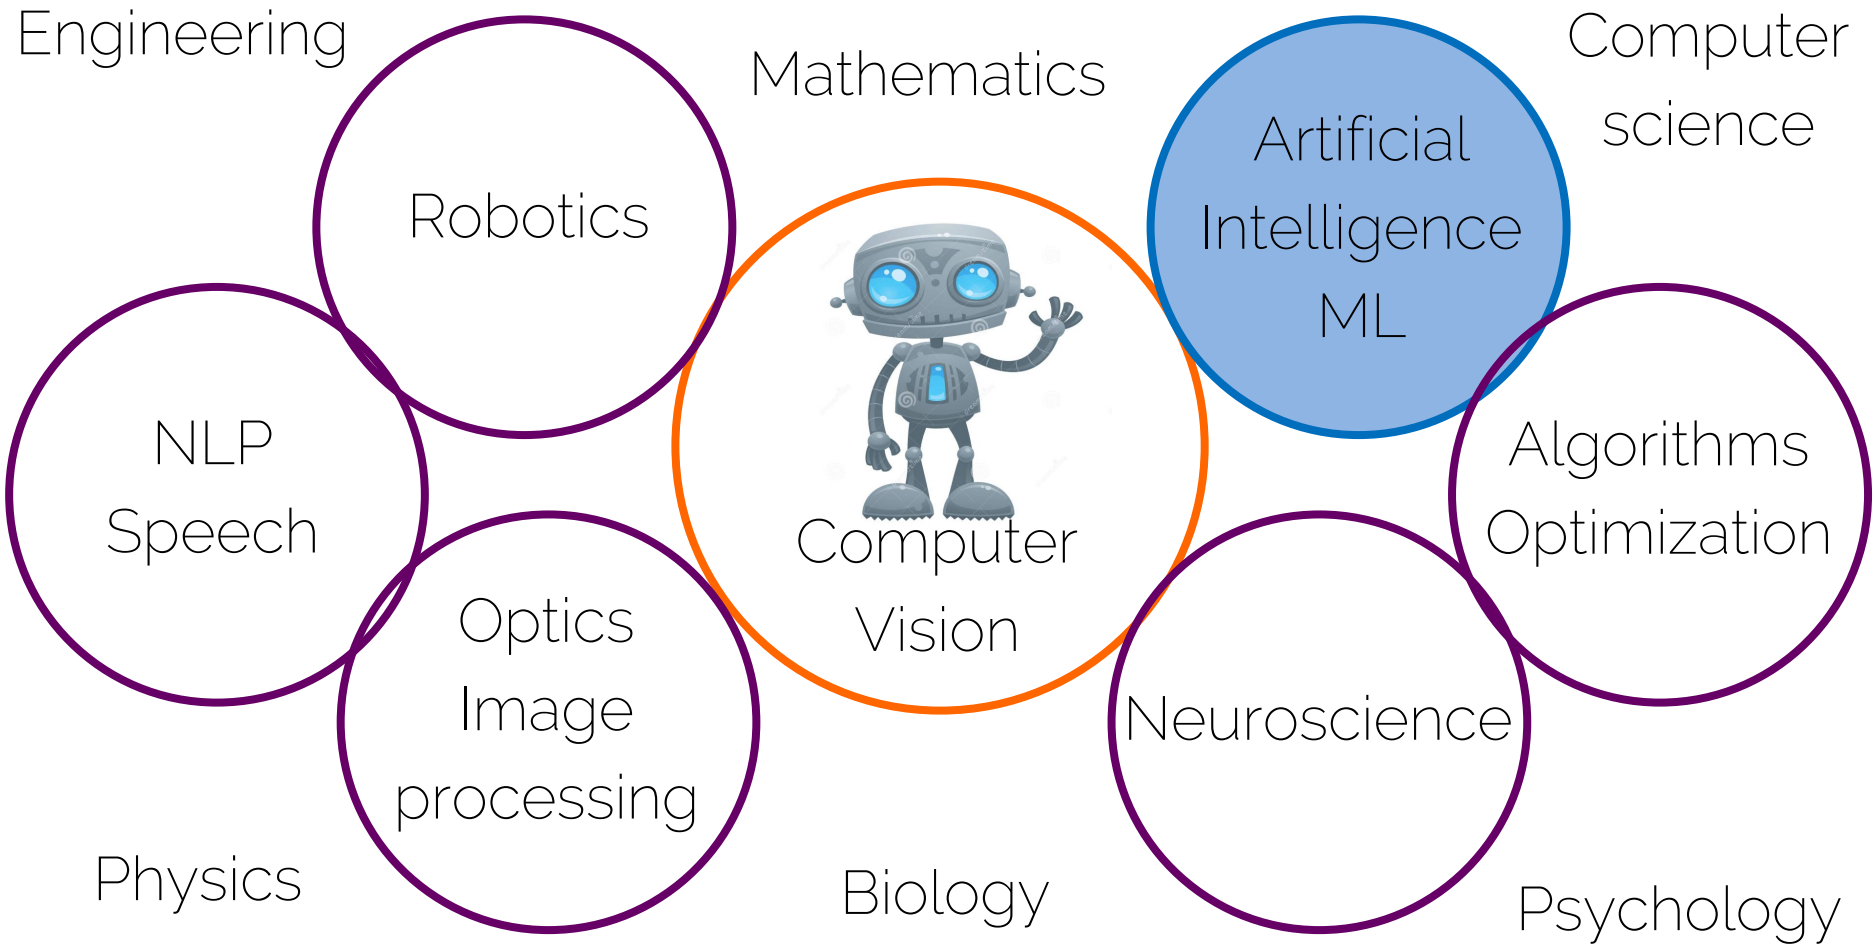
Engineering
Mathematics
Computer science
Robotics
Artificial Intelligence ML
NLP Speech
Computer Vision
Algorithms Optimization
Optics Image processing
Neuroscience
Physics
Biology
Psychology

Image classification

Input image

Preprocessing

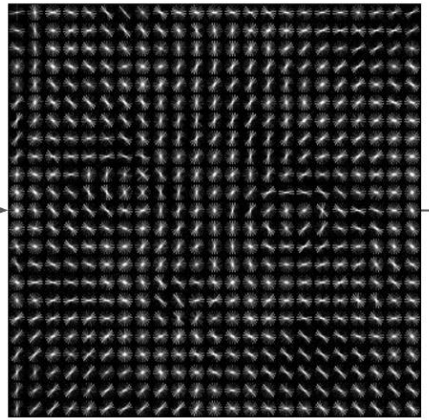

Features : HAAR, HOG, SIFT, SURF

Learning Algorithm

SVM, Random Forests, ANN

Label Assignment

Cat or Background

Image classification

Input image

Awesome magic box

Label Assignment

Cat or Background

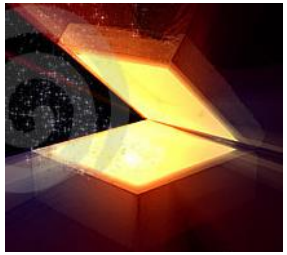

Open the box

Become magicians

Post 2012

Why Deep Learning?

Deep Learning History

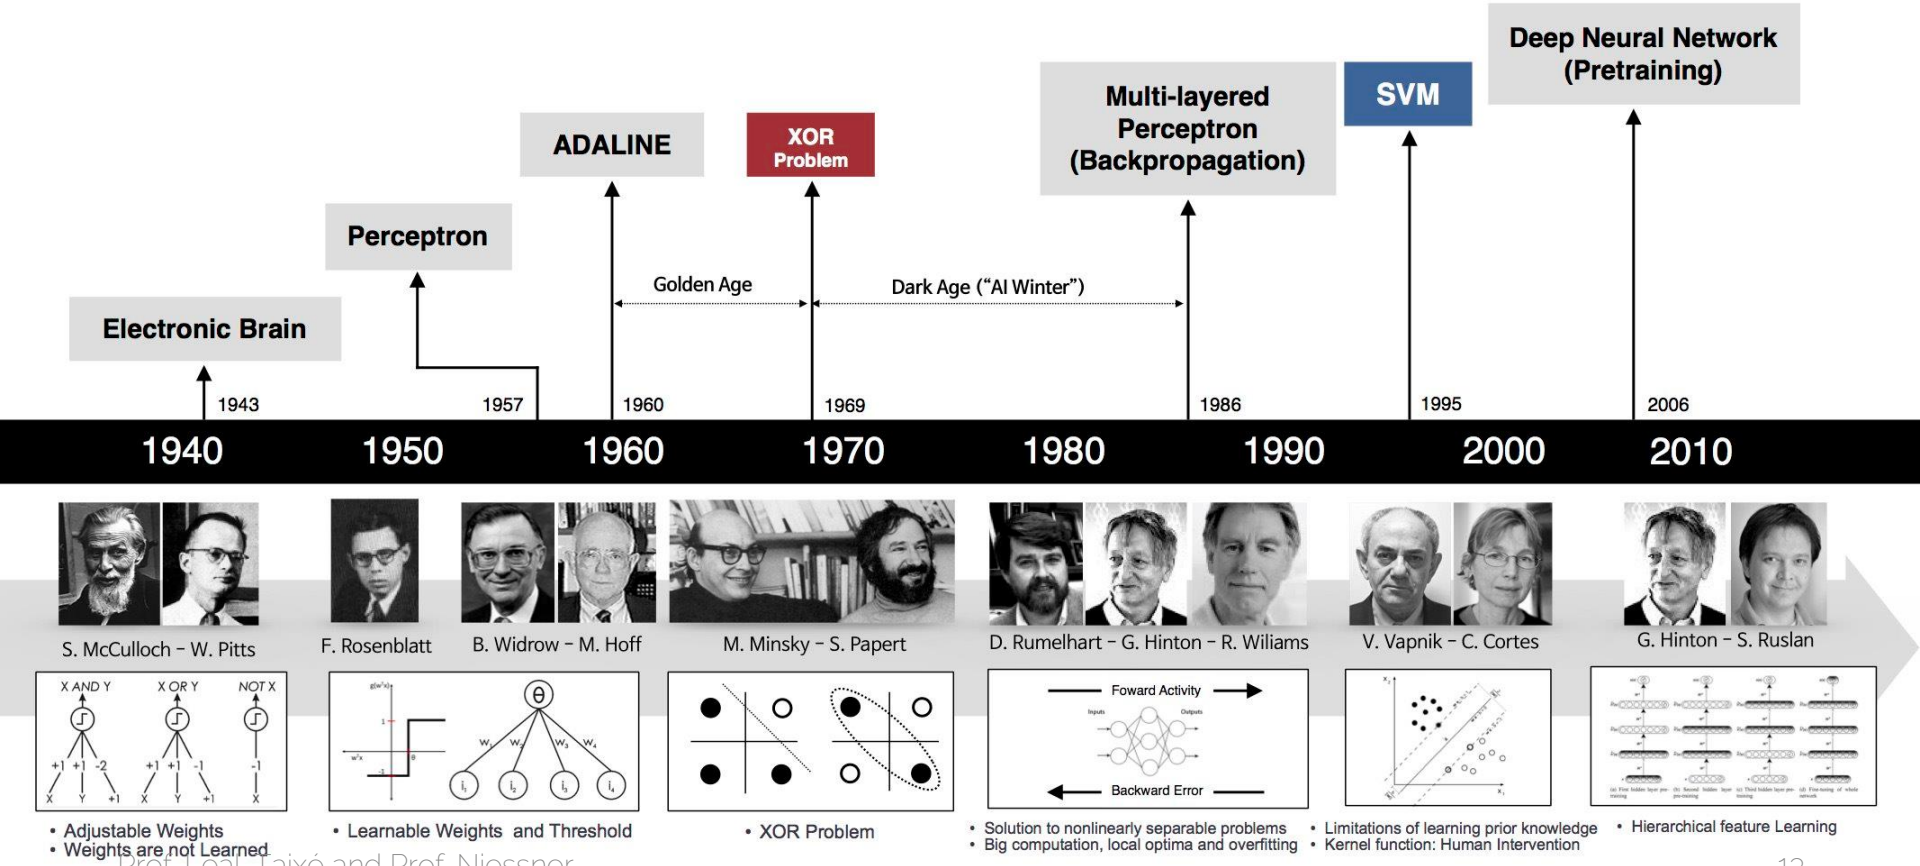

The empire strikes back

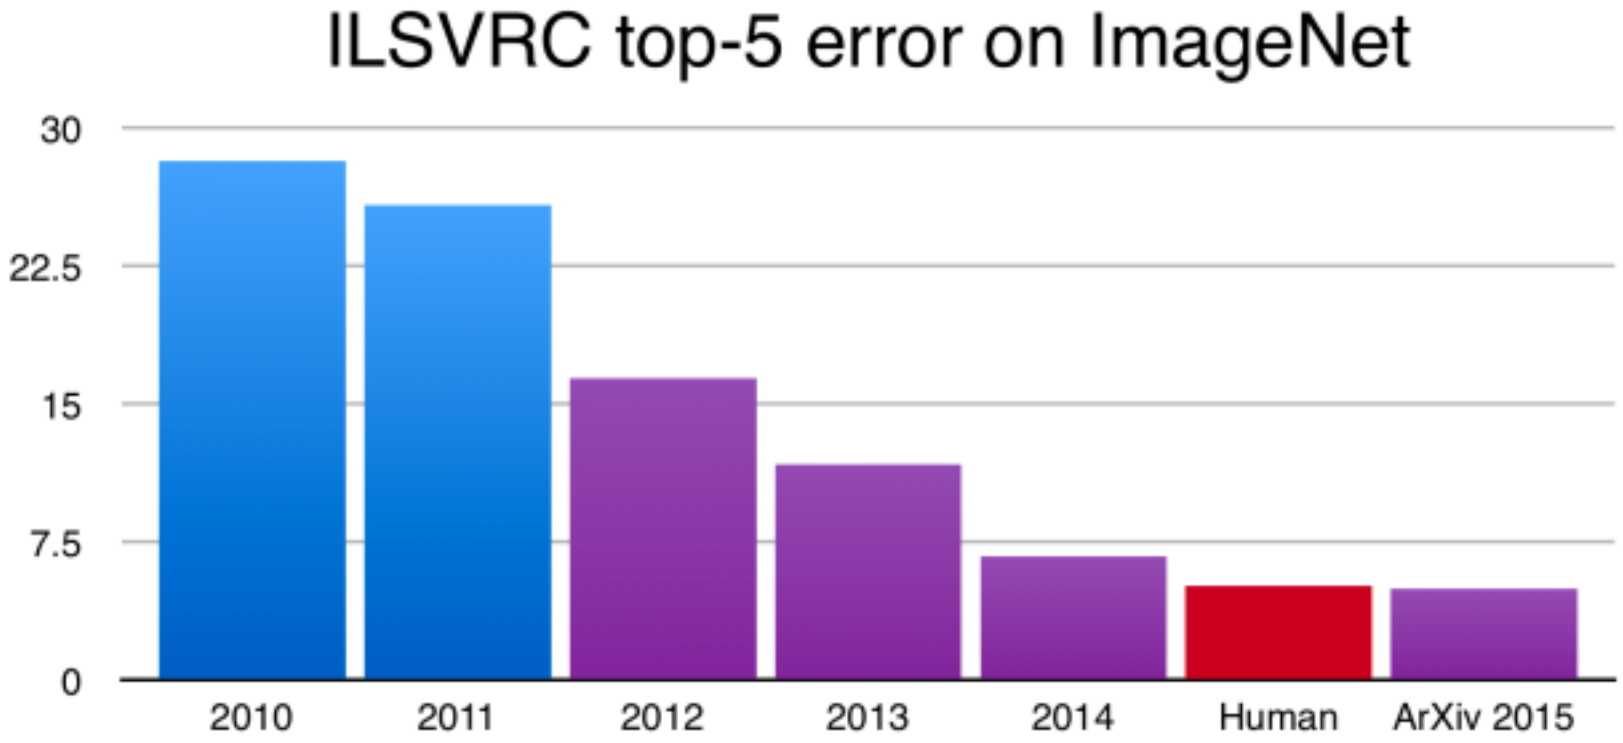

What has changed?

1988
LeCun
et al.

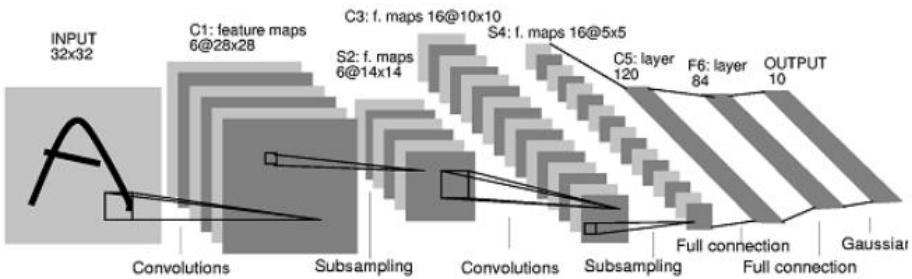

- MNIST digit recognition dataset
- 10^7 pixels used in training

2012
Krizhevsky
et al.

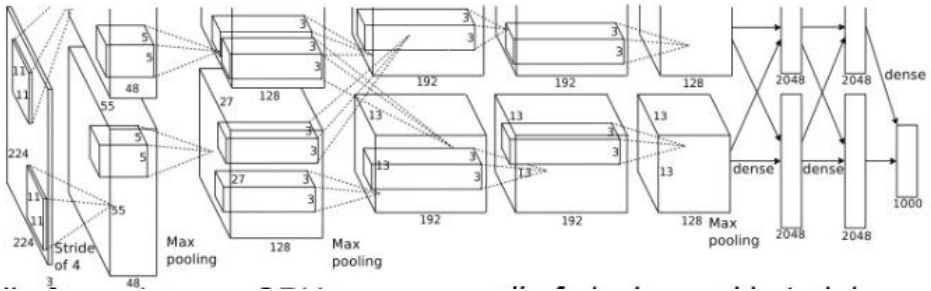

- ImageNet image recognition dataset
- 10^{14} pixels used in training

What made this possible?

Models know where to learn from

Models are trainable

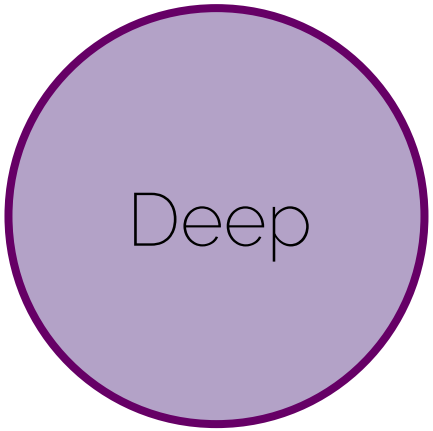

Models are complex

Deep Learning nowadays

AlphaGo

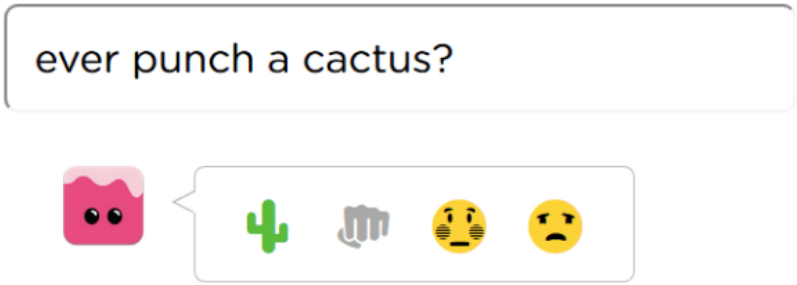

Emoticon suggestion

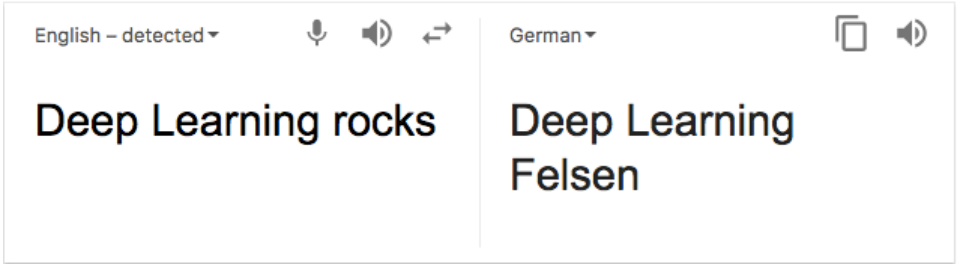

Machine translation

Deep Learning nowadays

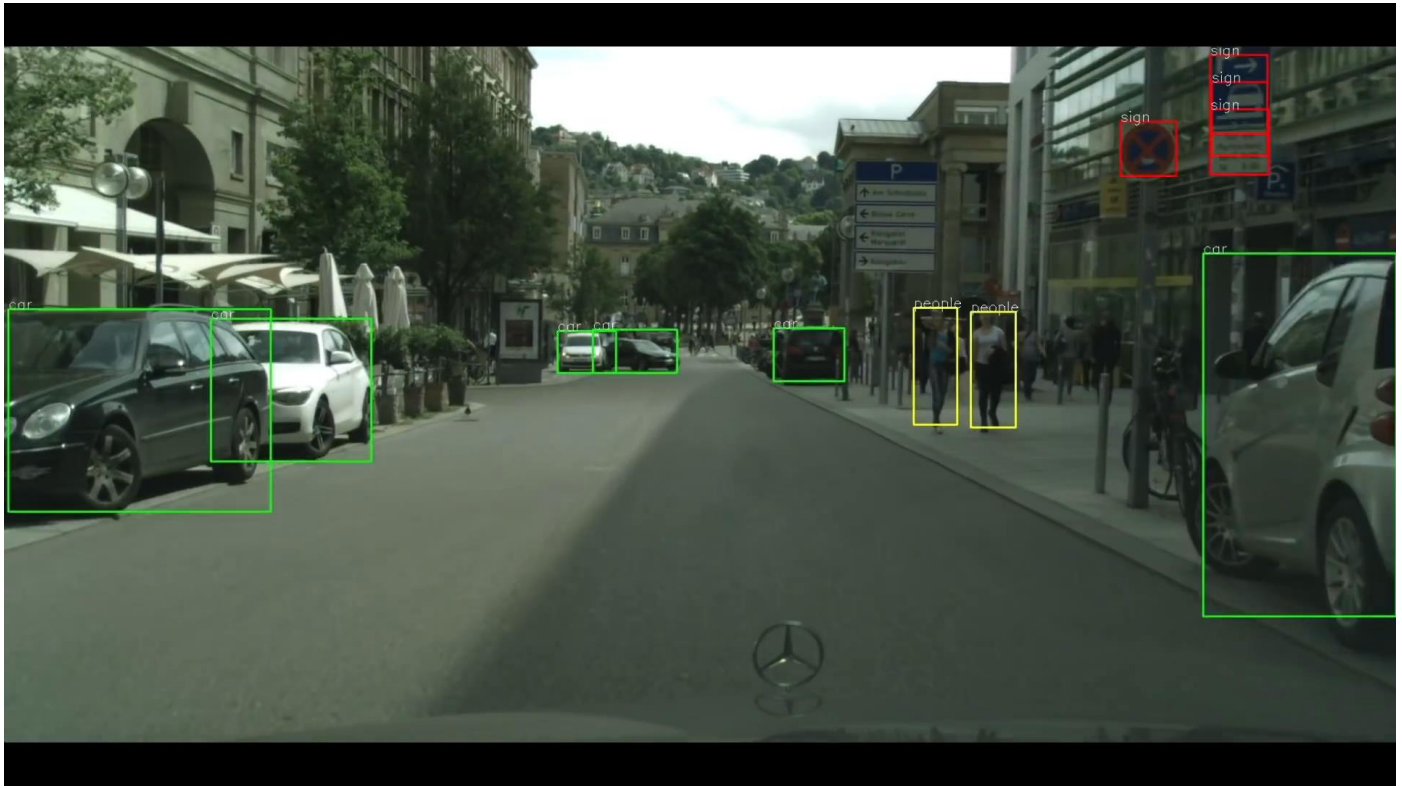

Self-driving cars

Deep Learning nowadays

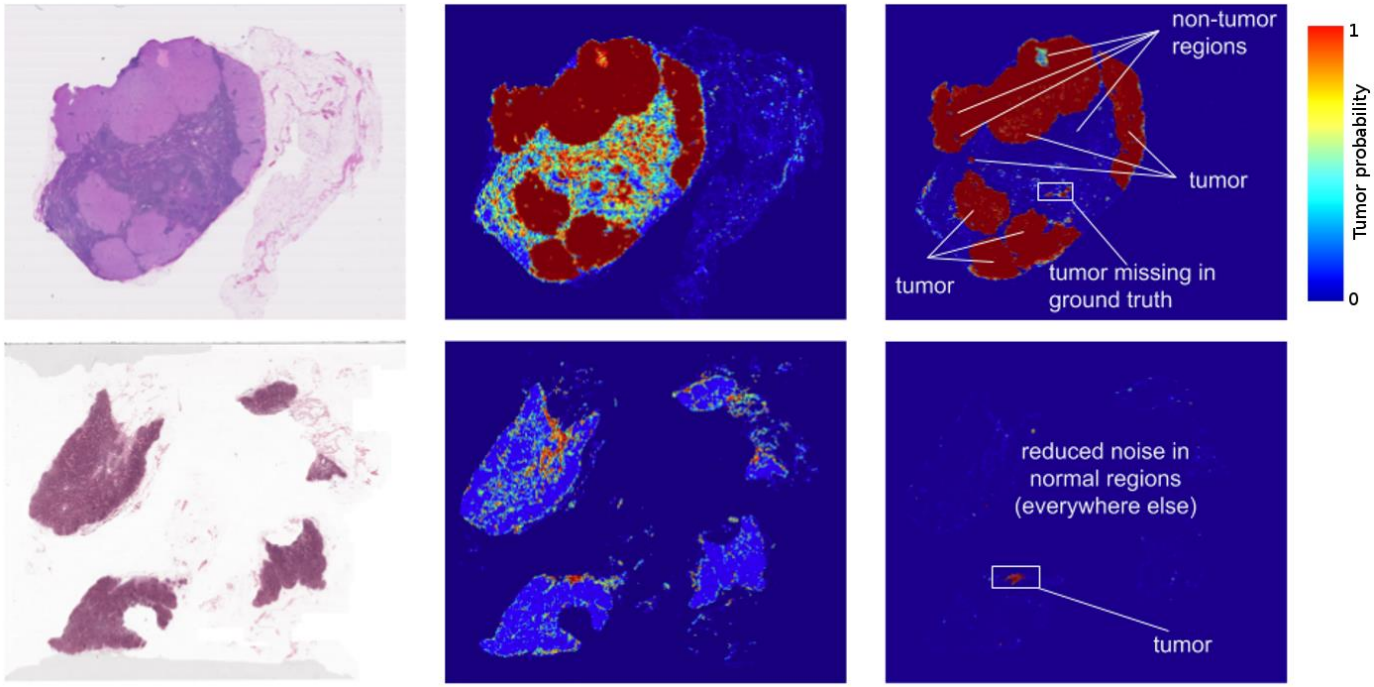

Healthcare, cancer detection

Deep Learning market

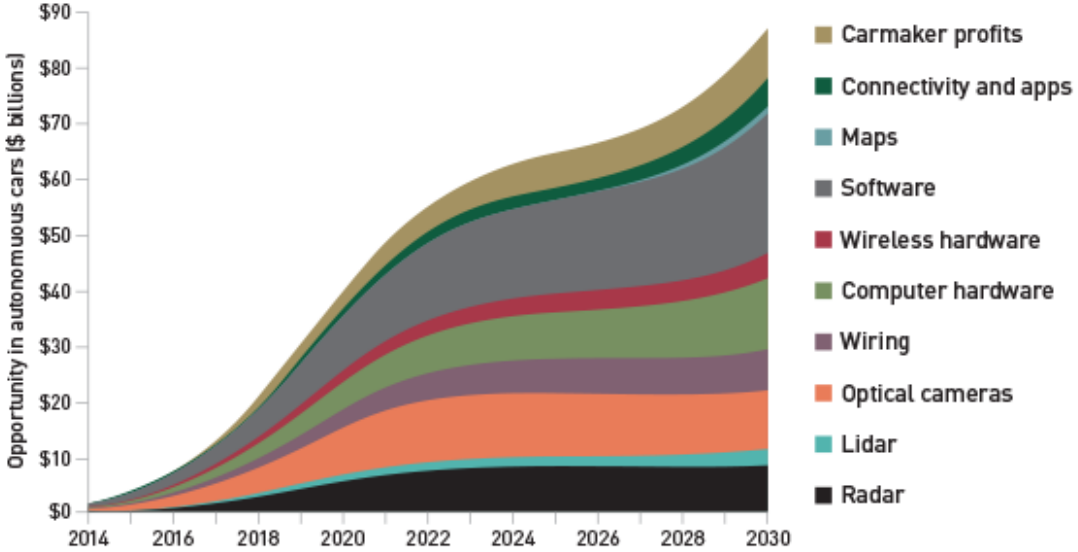

- […]market research report Deep Learning Market […] Global Forecasts to 2022", the deep learning market is expected to be worth **USD 1,722.9 Million by 2022.**

Deep Learning at TUM

S. Caelles, K.K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-Shot Video Object Segmentation, CVPR 2017.

Deep Learning at TUM

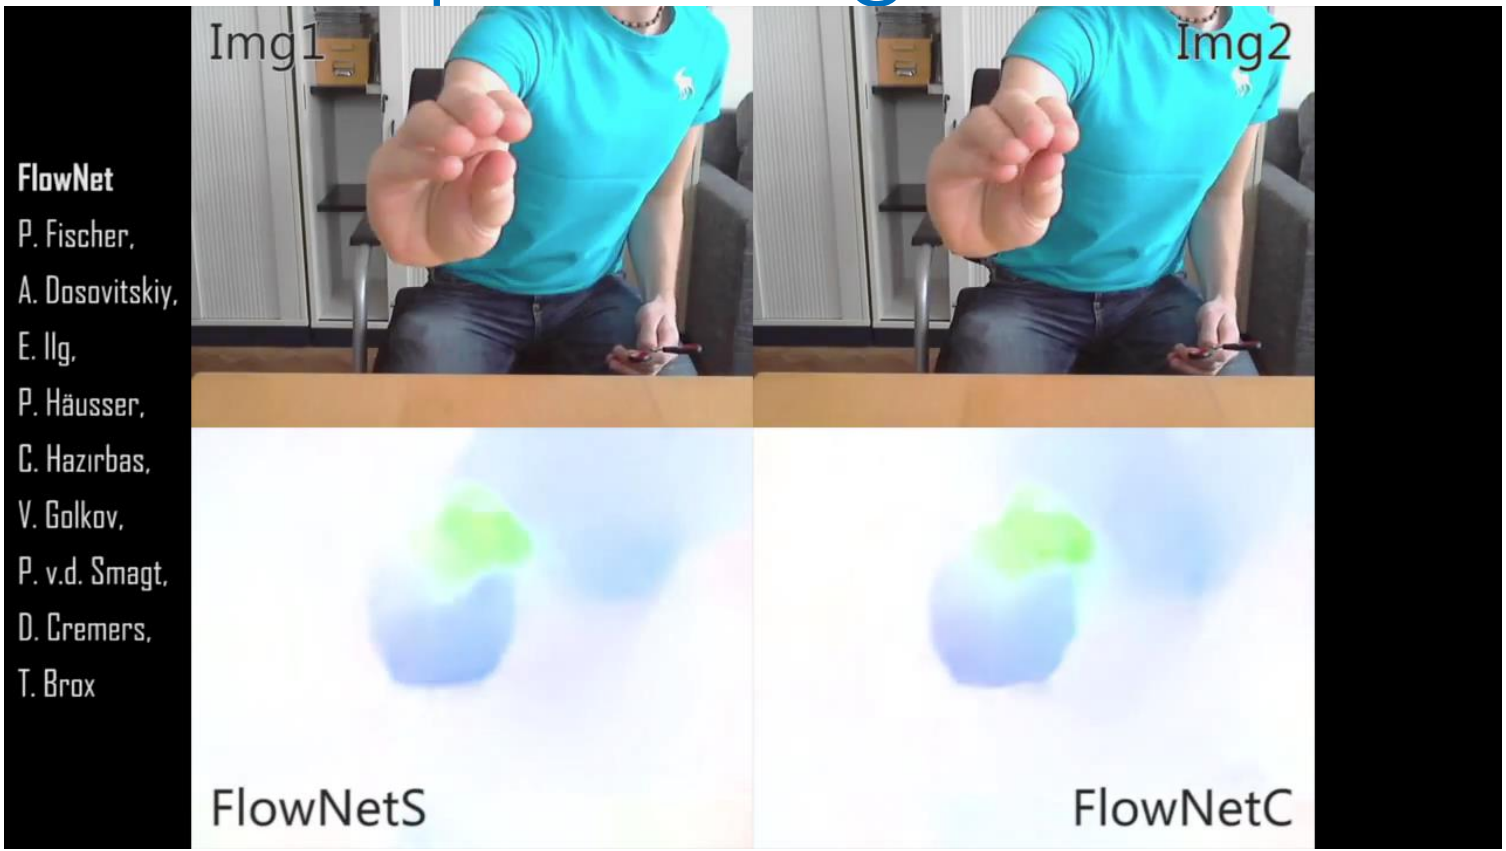

Deep Learning at TUM

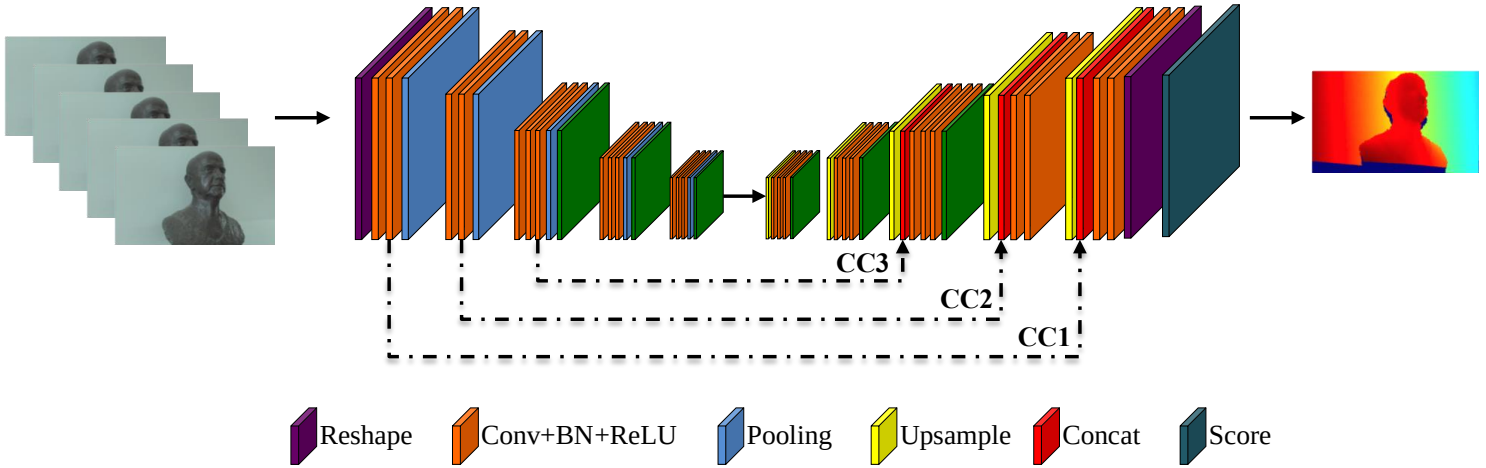

DDFF

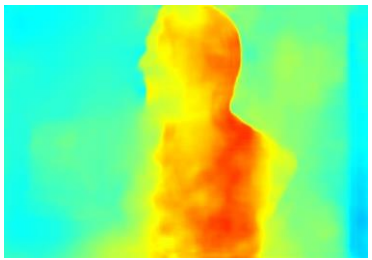

Computer Vision at TUM

ScanNet Stats:
-Kinect-style RGB-D sensors
-1513 scans of 3D environments
-2.5 Mio RGB-D frames
-Dense 3D, crowd-source MTurk labels
-Annotations projected to 2D frames

ScanNet: Dai, Chang, Savva, Halber, Funkhouser, Niessner., CVPR *2017.*

Deep Learning at TUM

Map

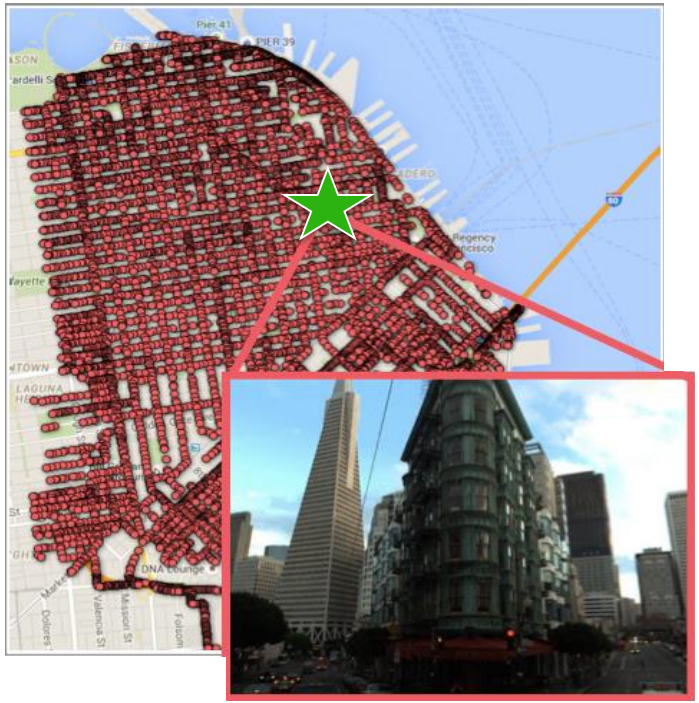

Photo

Introduction to Deep Learning

About the lecture

- Theory: 12 lectures
  - Every Monday 14-16h (MI HS 1)
- Practice: 3 exercises, practical sessions
  - Every Thursday 8-10h (Interim HS1)
- July 2nd: guest lecture by tba

https://dvl.in.tum.de/lectures/dl4cv-ss18.html

Grading system

- Exam: **July 16th**
- Review: allow until end of July for exam reviews
- Important: no retake exam

- Practice: 4 exercises (Thursdays)
  - Bonus 0.3 + questions in the final exam

https://dvl.in.tum.de/lectures/dl4cv-ss18.html

Exercise lectures

- Exercise 1: starting May 3rd

- Thursday lecture 1: DL math background
- Thursday lecture 2: DL math background
- Thursday lecture 3: Python introduction

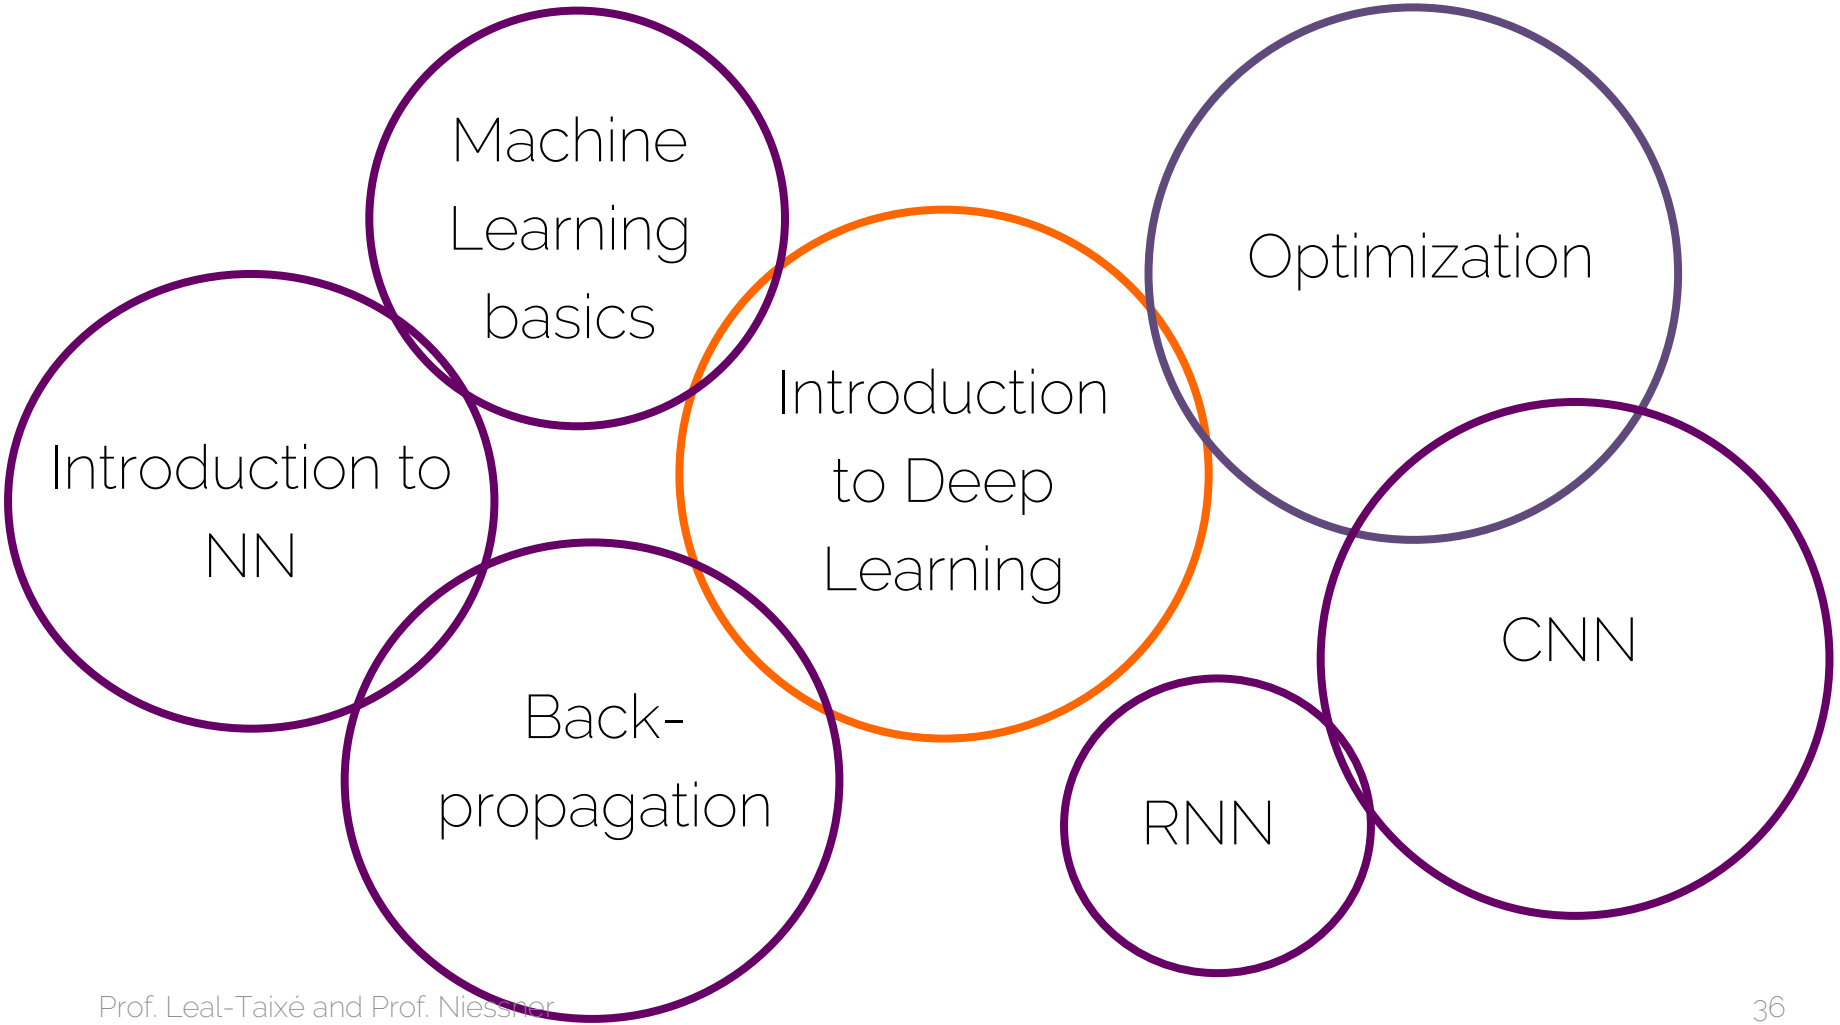
Machine Learning basics
Optimization
Introduction to NN
Introduction to Deep Learning
CNN
Back-propagation
RNN

Slides

- All material will be uploaded on Moodle
- Questions regarding the syllabus, exercises or contents of the lecture, use Moodle!
- Questions regarding organization of the course:

i2dl@dvl.in.tum.de

- Emails to our individual addresses will not be answered.

Deep Learning at TUM

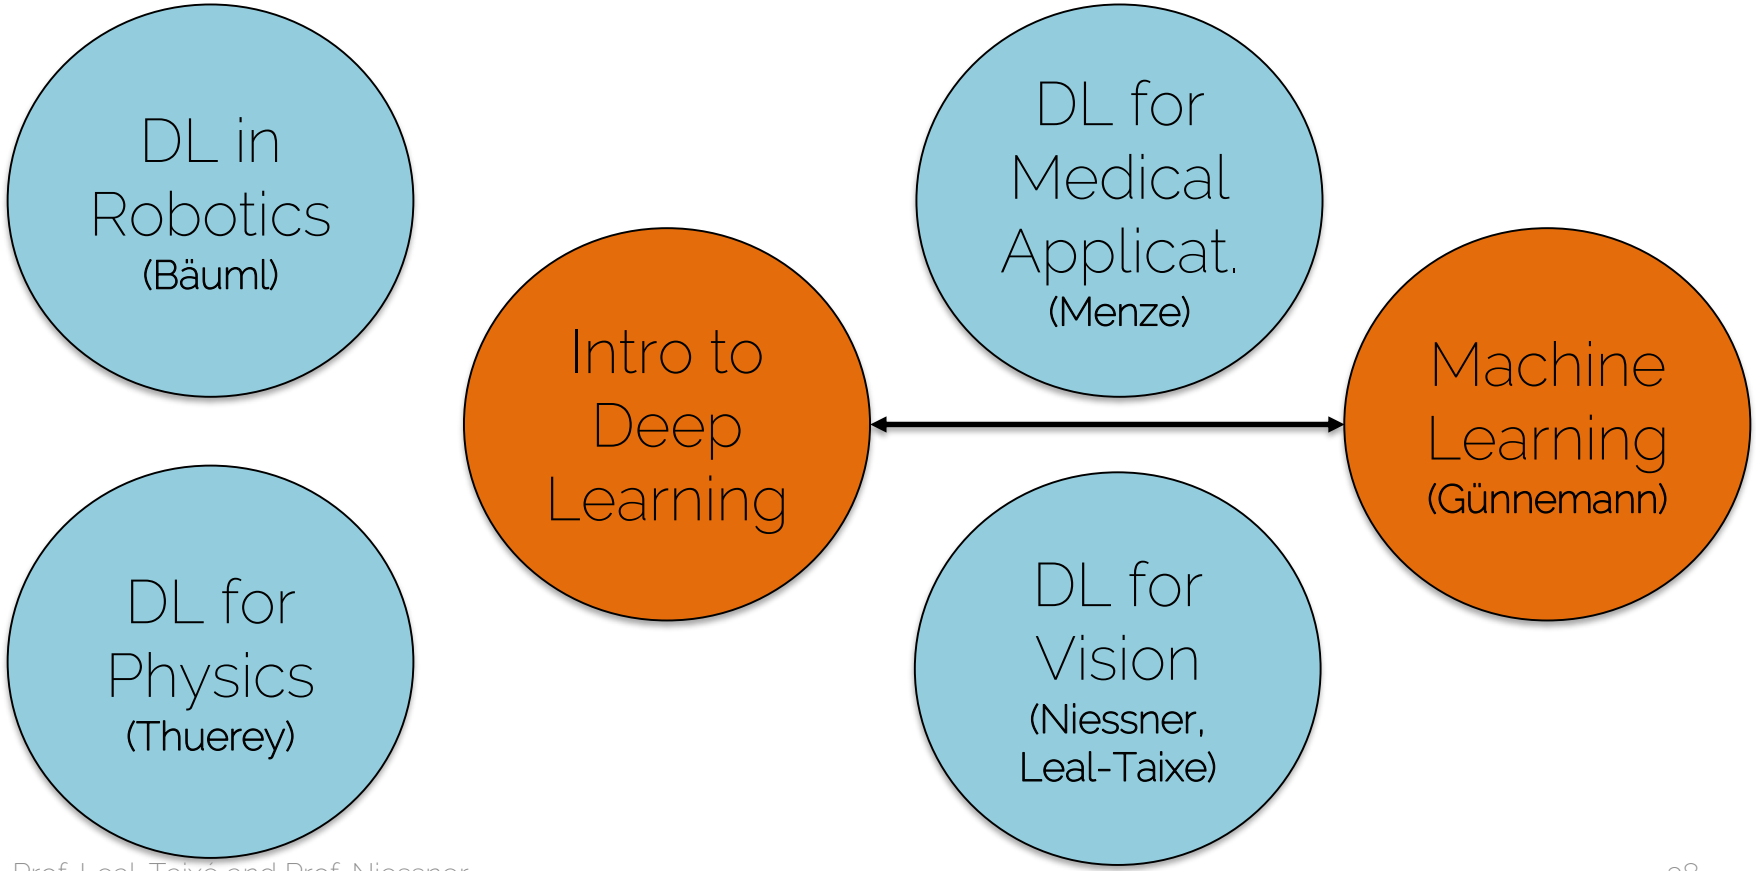

Machine Learning

Machine learning

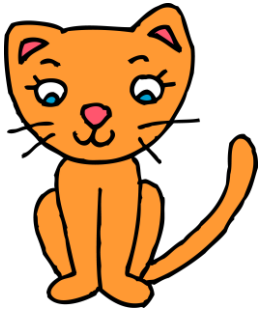

Task

Image classification

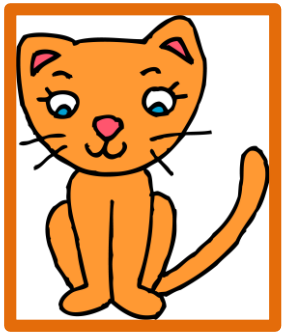





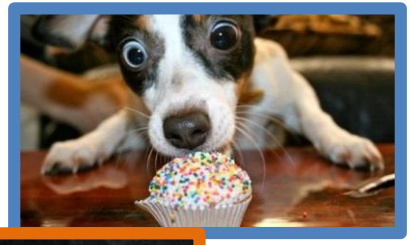

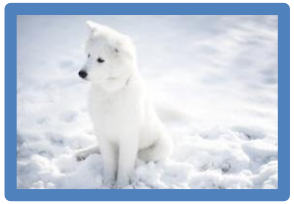
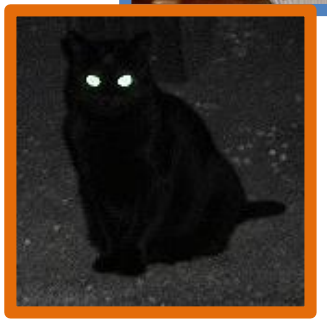

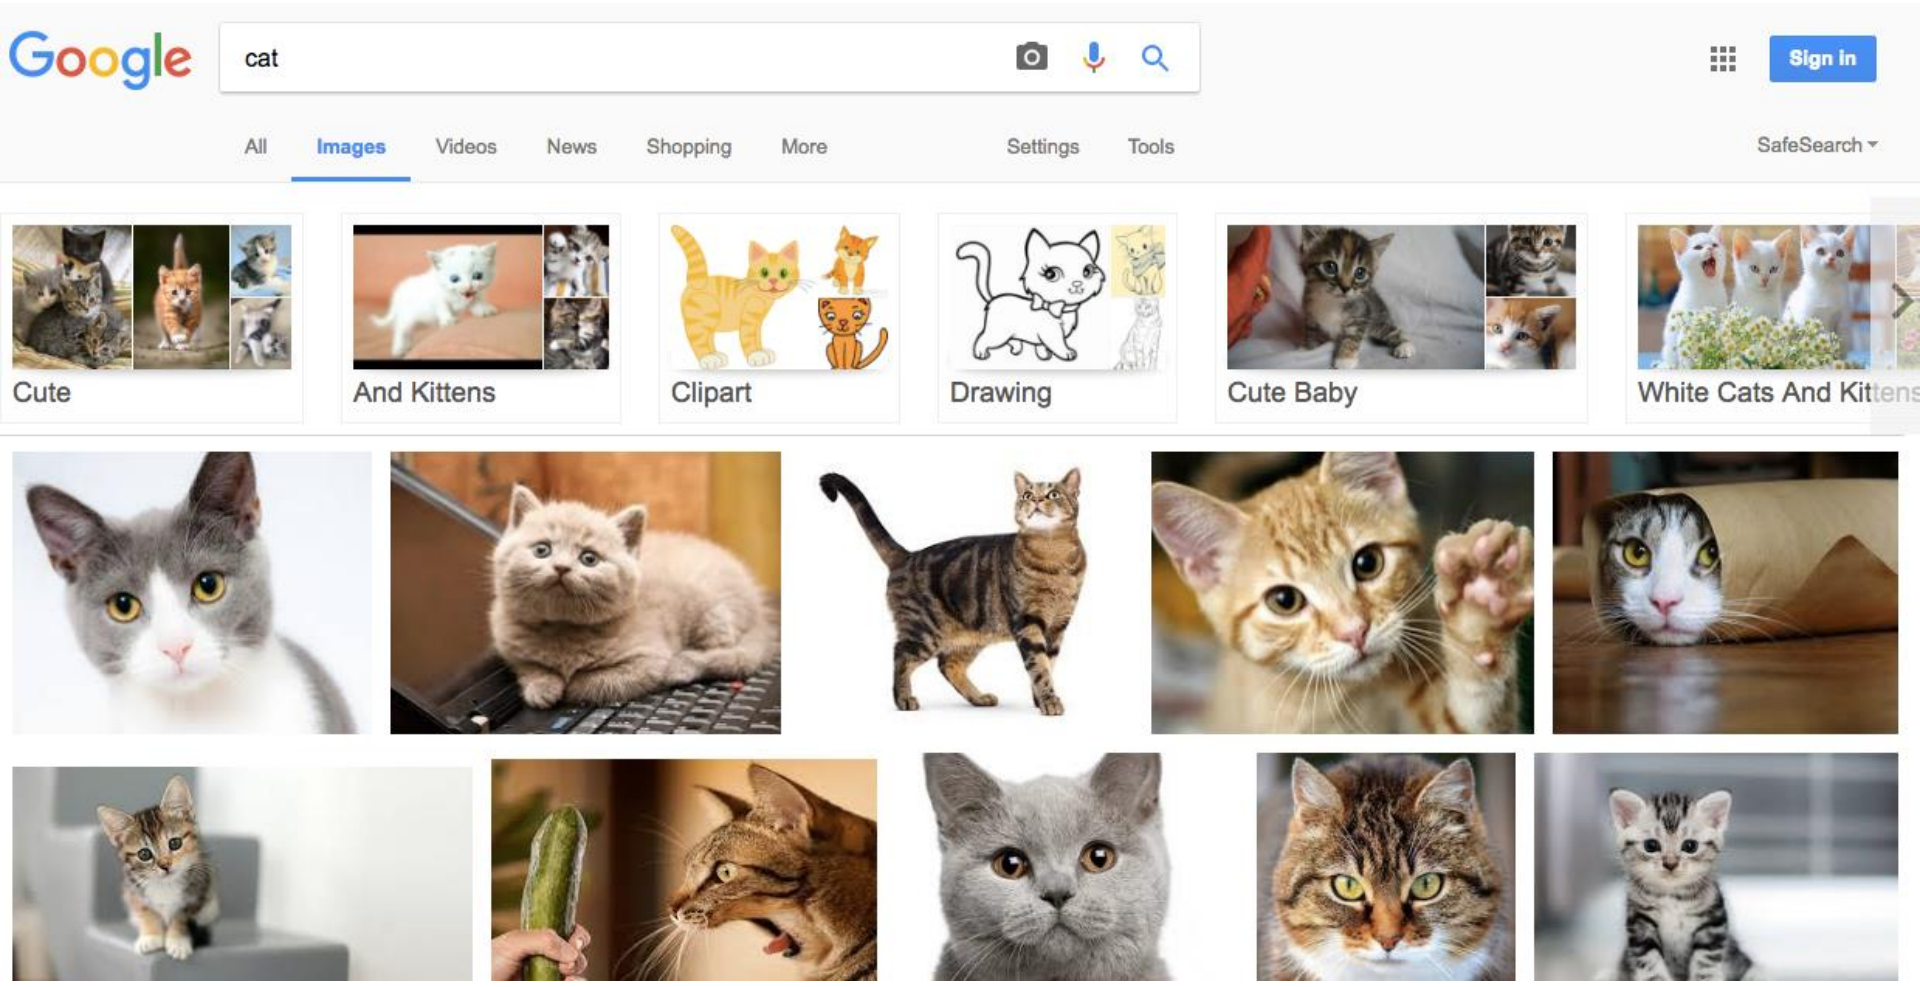

Pose Appearance Illumination

Image classification

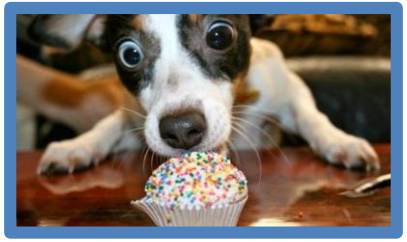

Occlusions

Image classification

Background clutter

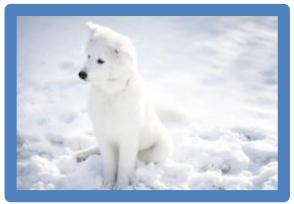

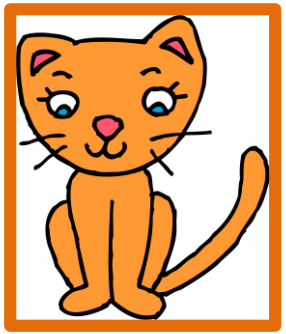

Image classification

Representation

Machine learning

- How can we learn to perform image classification?

Task

Image classification

Experience

Data

Machine learning

Unsupervised learning

- No label or target class
- Find out properties of the structure of the data
- Clustering (k-means, PCA)

Supervised learning

Machine learning

Unsupervised learning

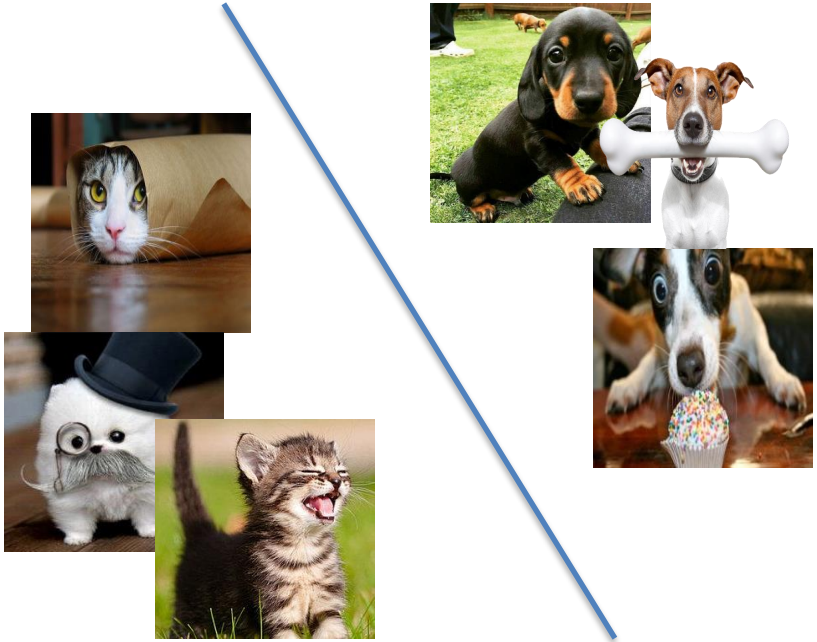

Supervised learning

Machine learning

Unsupervised learning

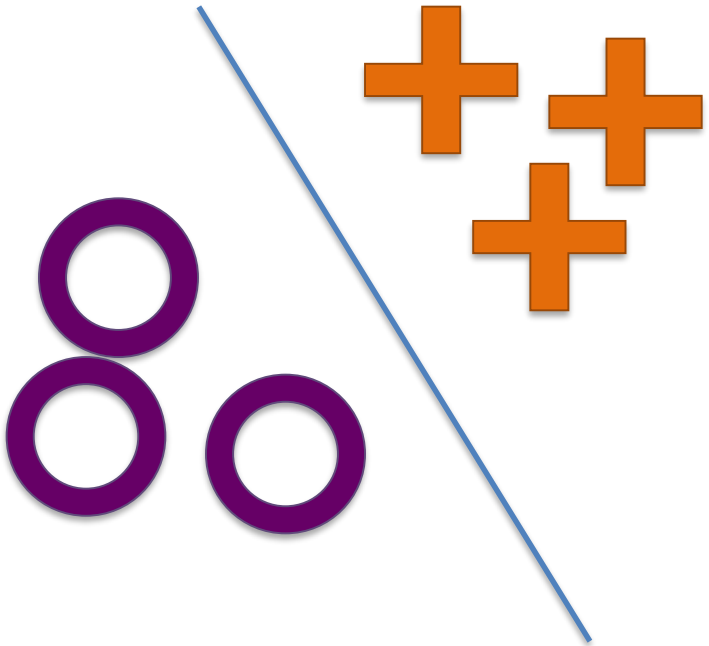

Supervised learning

- Labels or target classes

Machine learning

Unsupervised learning

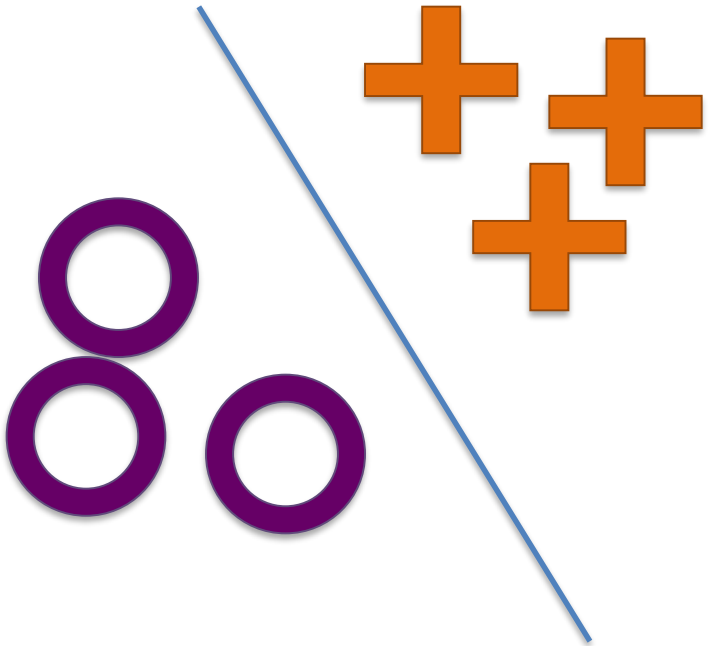

Supervised learning

CAT

DOG

DOG

CAT

CAT

DOG

Machine learning

- How can we learn to perform image classification?

Machine learning

Unsupervised learning

Supervised learning

Reinforcement learning

Machine learning

Unsupervised learning

Supervised learning

Reinforcement learning

Machine learning

- How can we learn to perform image classification?

Task

Image classification

Performance measure

Accuracy

Experience

Data

A simple classifier

Nearest Neighbor

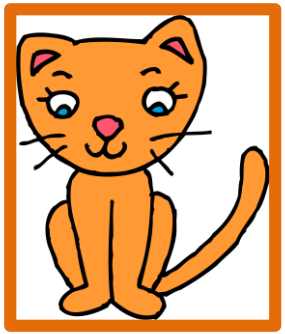






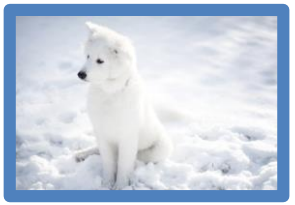


Nearest Neighbor

NN classifier = dog

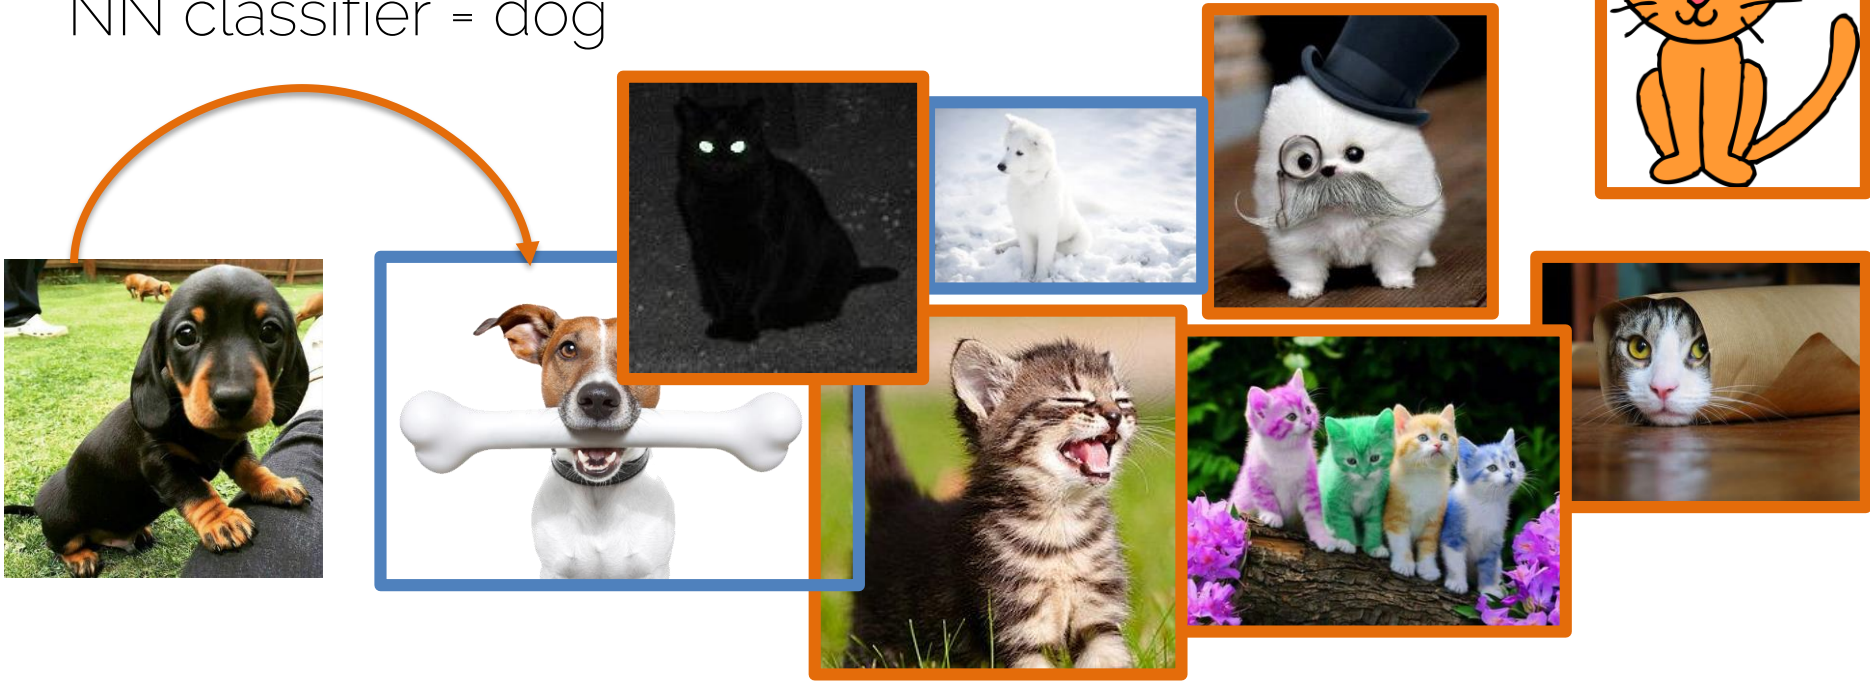

distance

Nearest Neighbor

k-NN classifier = cat

distance

Nearest Neighbor

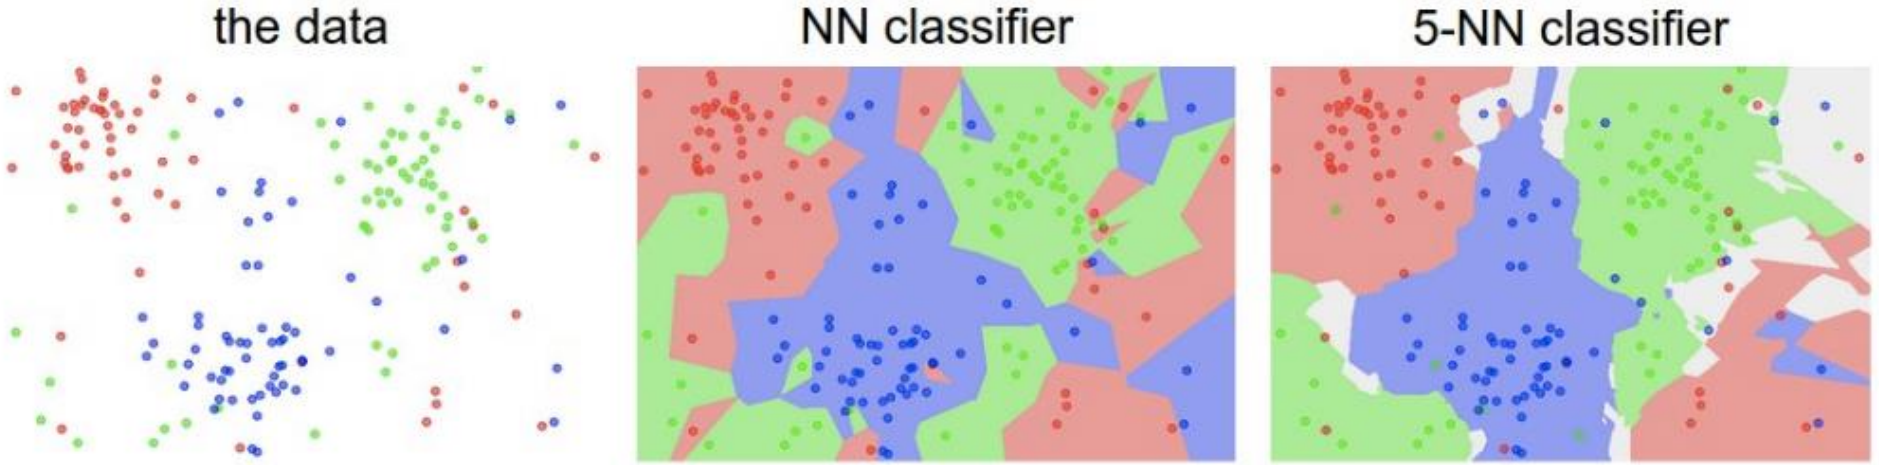

What is the performance on training data for NN classifier?

What classifier is more likely to perform best on test data?

Courtesy of Stanford course cs231n

Nearest Neighbor

- Hyperparameters
  - Distance (L1, L2)
  - k (number of neighbors)
- These parameters are problem dependent.
- How do we choose these hyperparameters?

Cross validation

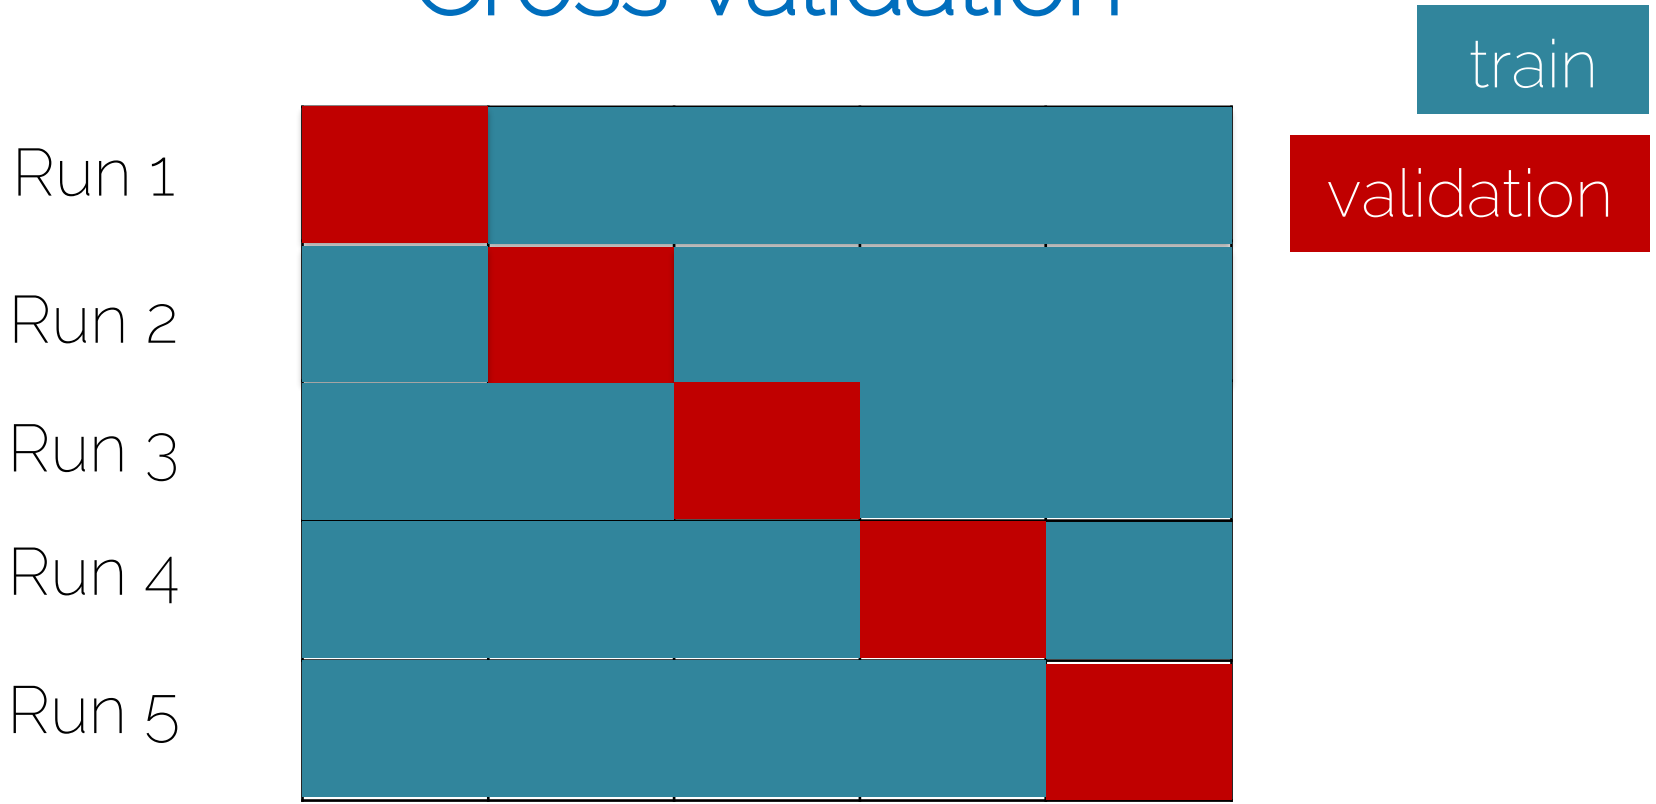

Split the **training data** into N folds

Cross validation

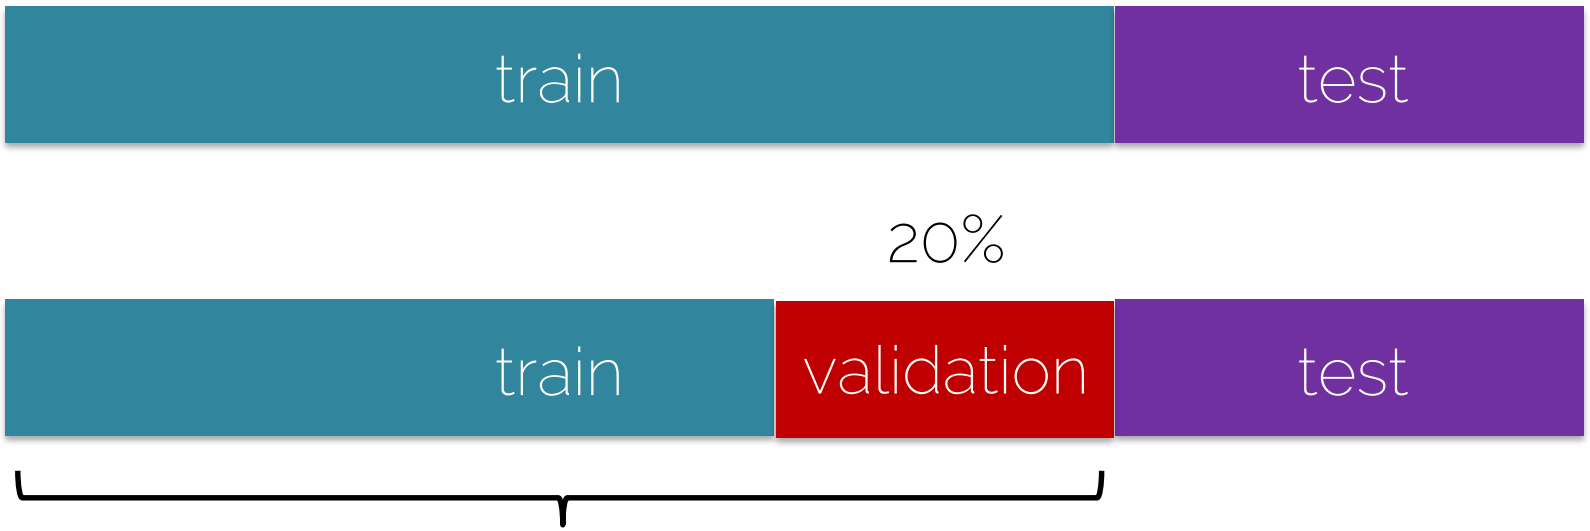

This lecture: improving our classifier

- Beyond linear classification
- How to train complex models → deep networks
- What is happening behind the scenes: optimization, CNN, regularization.

Upcoming lecture

- Next Monday: Lecture 2: Machine Learning basics

- Next Thursday: 1st practical lecture (DL math background)