Applied Statistics for Neuroscientists
Part IIa: Neural networks and deep learning

Dr. Seyed-Ahmad Ahmadi

16.11.2017

Outline – Neural networks and Deep Learning

- Artificial neural networks
  - Structure of a MLP
  - Inference via forward-propagation
  - Learning via back-propagation
  - Cost functions: classification vs. regression
- Machine learning vs. representation learning
- Deep learning
  - Successes
- Supervised DL with CNNs
  - Image filters
  - Stacking of hierarchical filters → CNN
- Auto-encoders
- Outlook

Recap: supervised learning pipeline

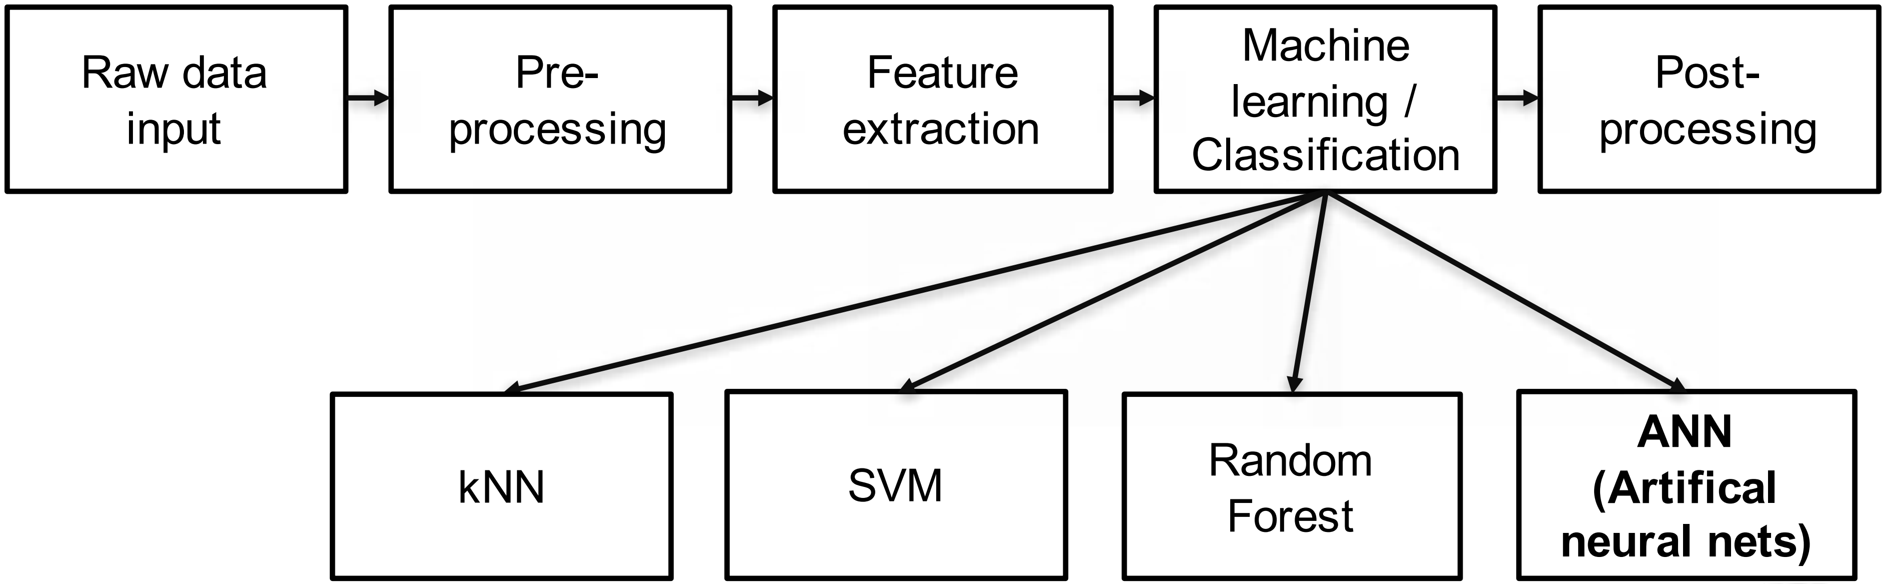

Artifical neural networks (ANN): inspiration

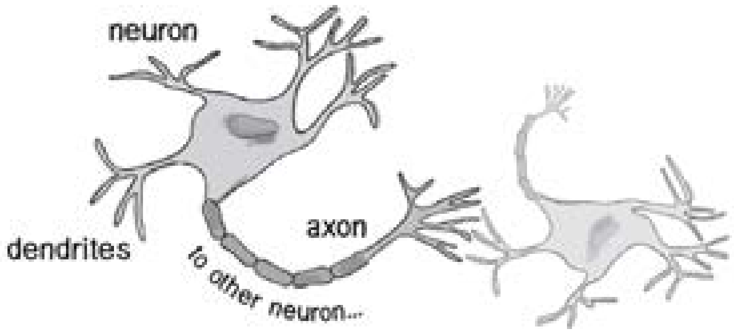

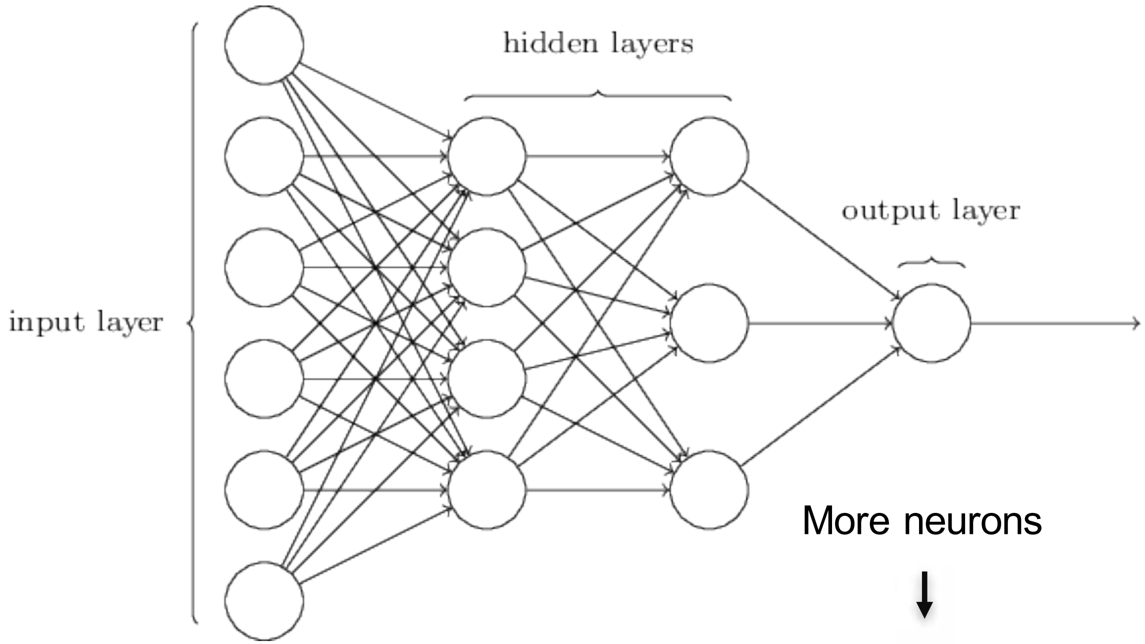

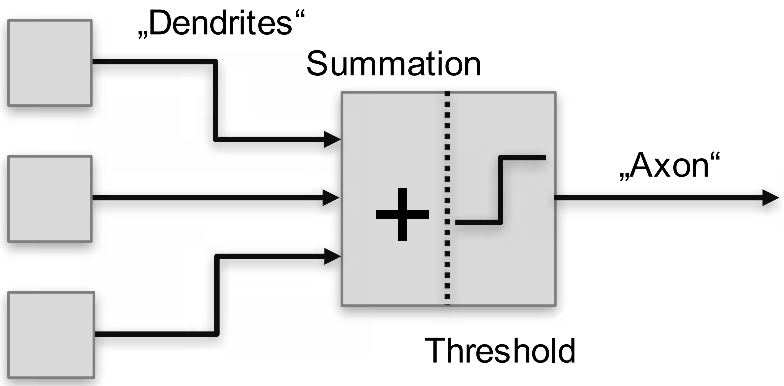

More neurons

↓

more modeling capacity

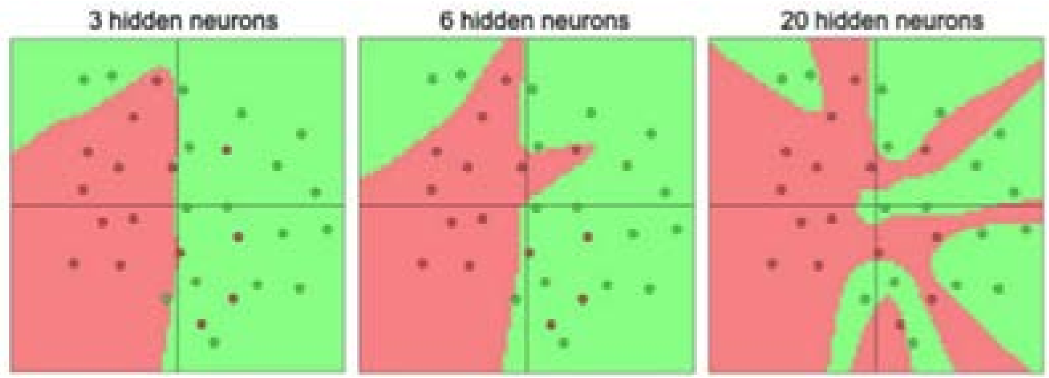

Mathematical modeling of a neuron

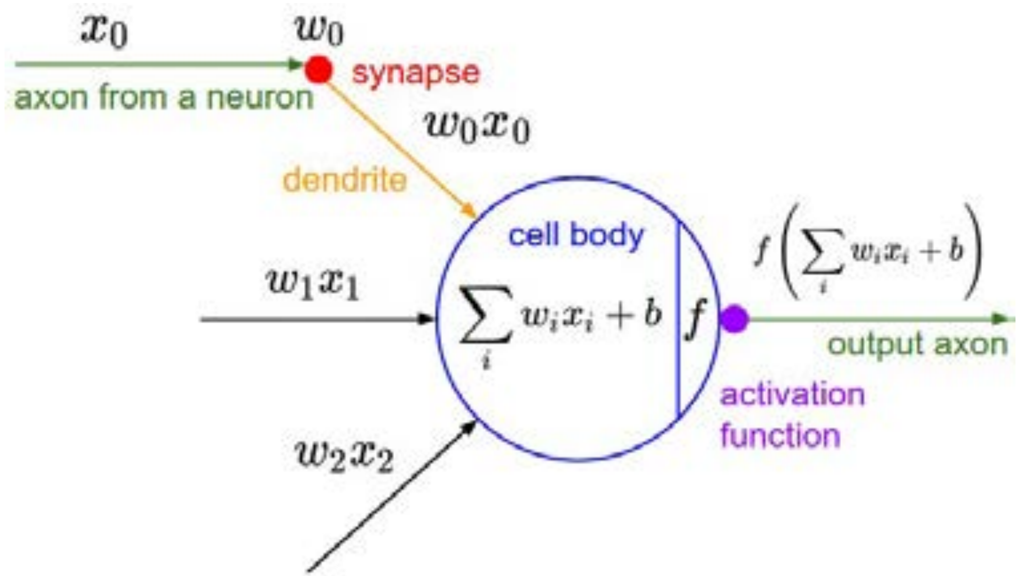

Non-linear activation function

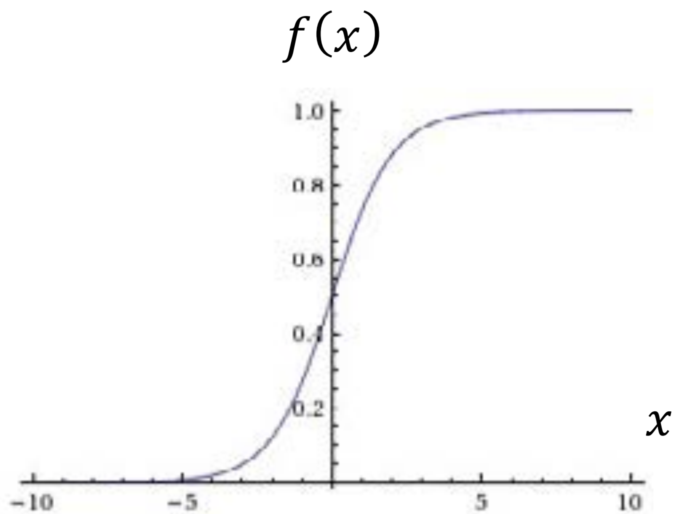

$$f(x) = \frac{1}{1 + e^{-x}} \longrightarrow f\left(\sum_i w_i x_i + b\right) = \frac{1}{1 + e^{-(\sum_i w_i x_i + b)}}$$

Images from CS231n: http://cs231n.github.io/

Neural Networks: Classification of hand-written digits using a fully connected ANN

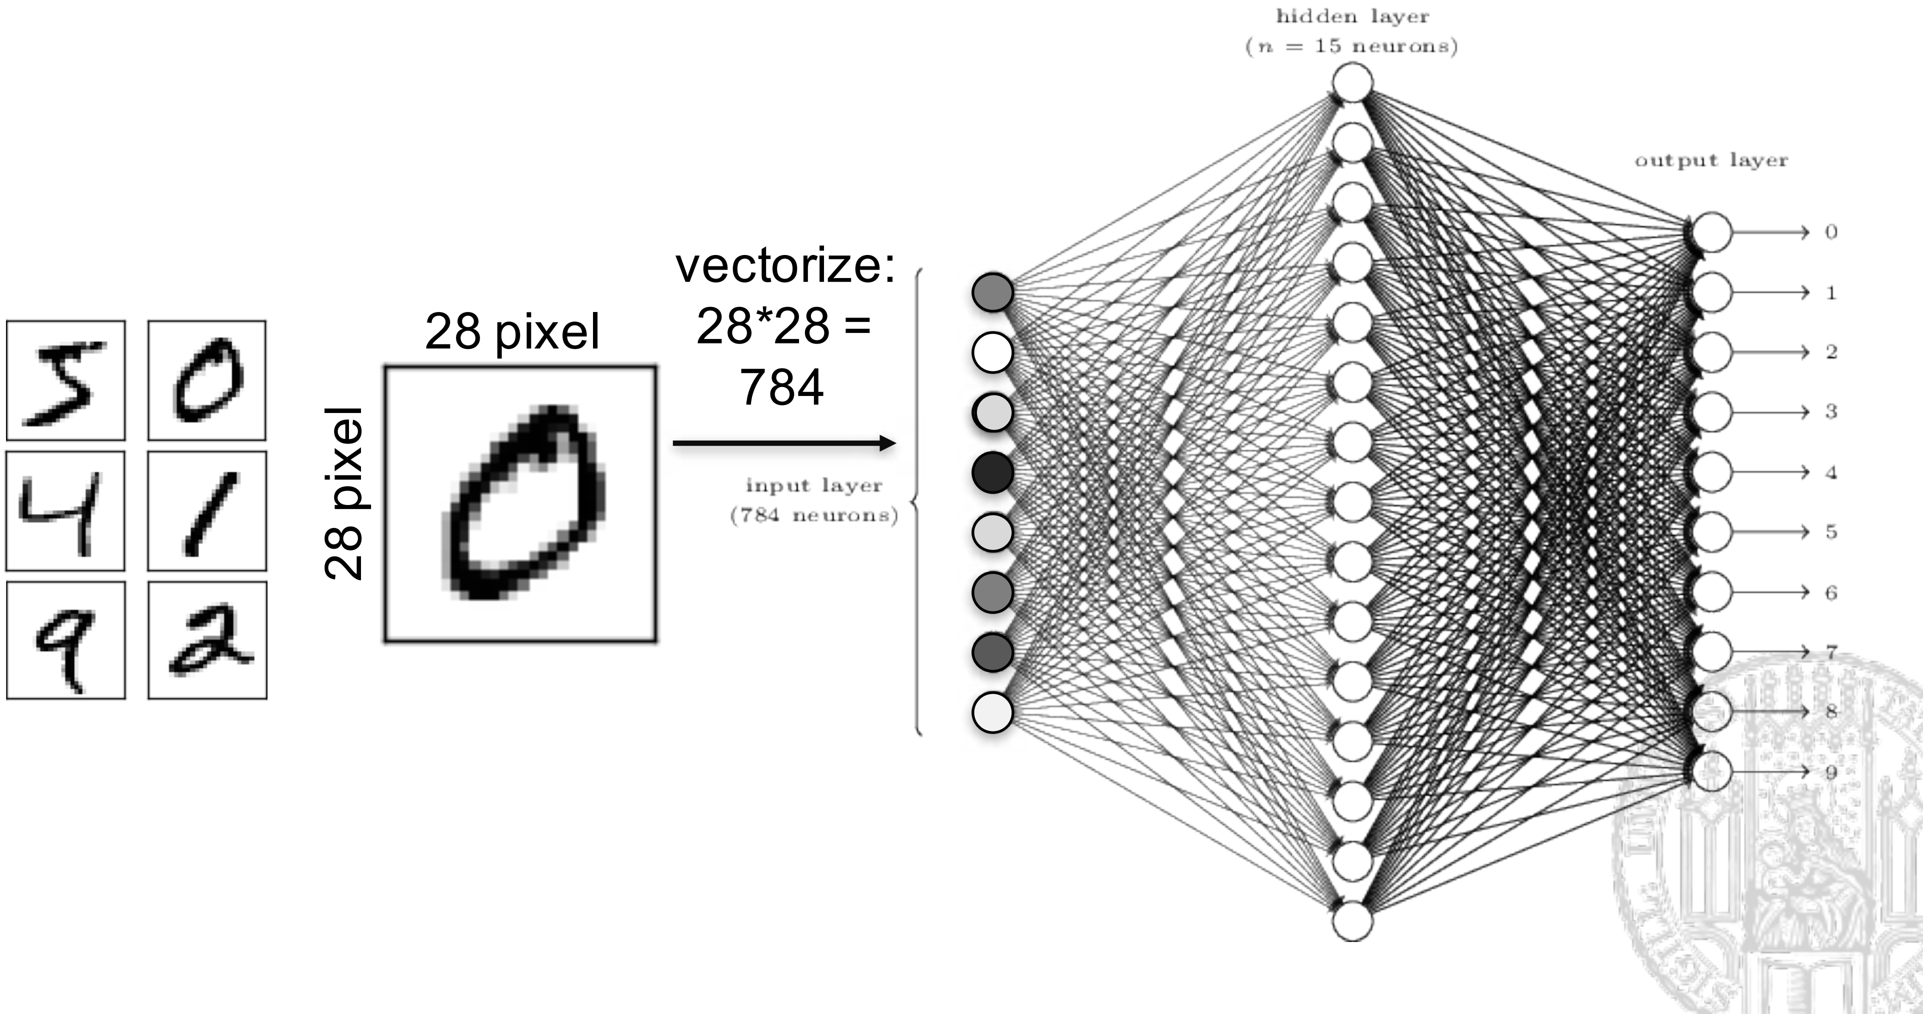

Neural Networks: classification by feed-forward, training with backpropagation

Classification:
Feed-forward

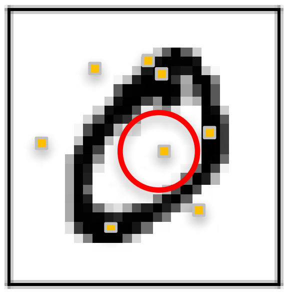

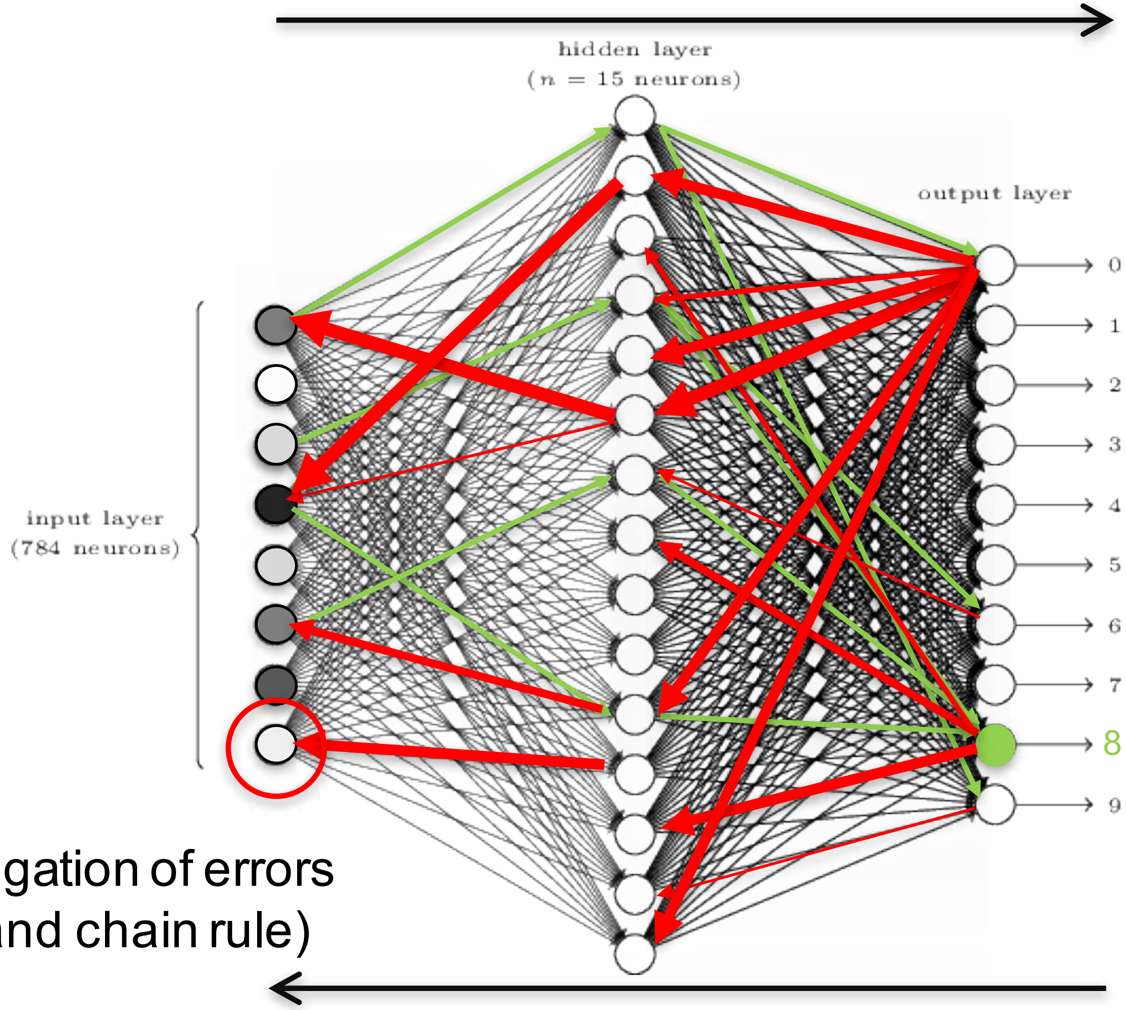

Out
0.12
0
0
0
0
0
0.2
0
0.85
0.1

Training:
Back-propagation of errors
(gradients and chain rule)

ANN example: posturography platform

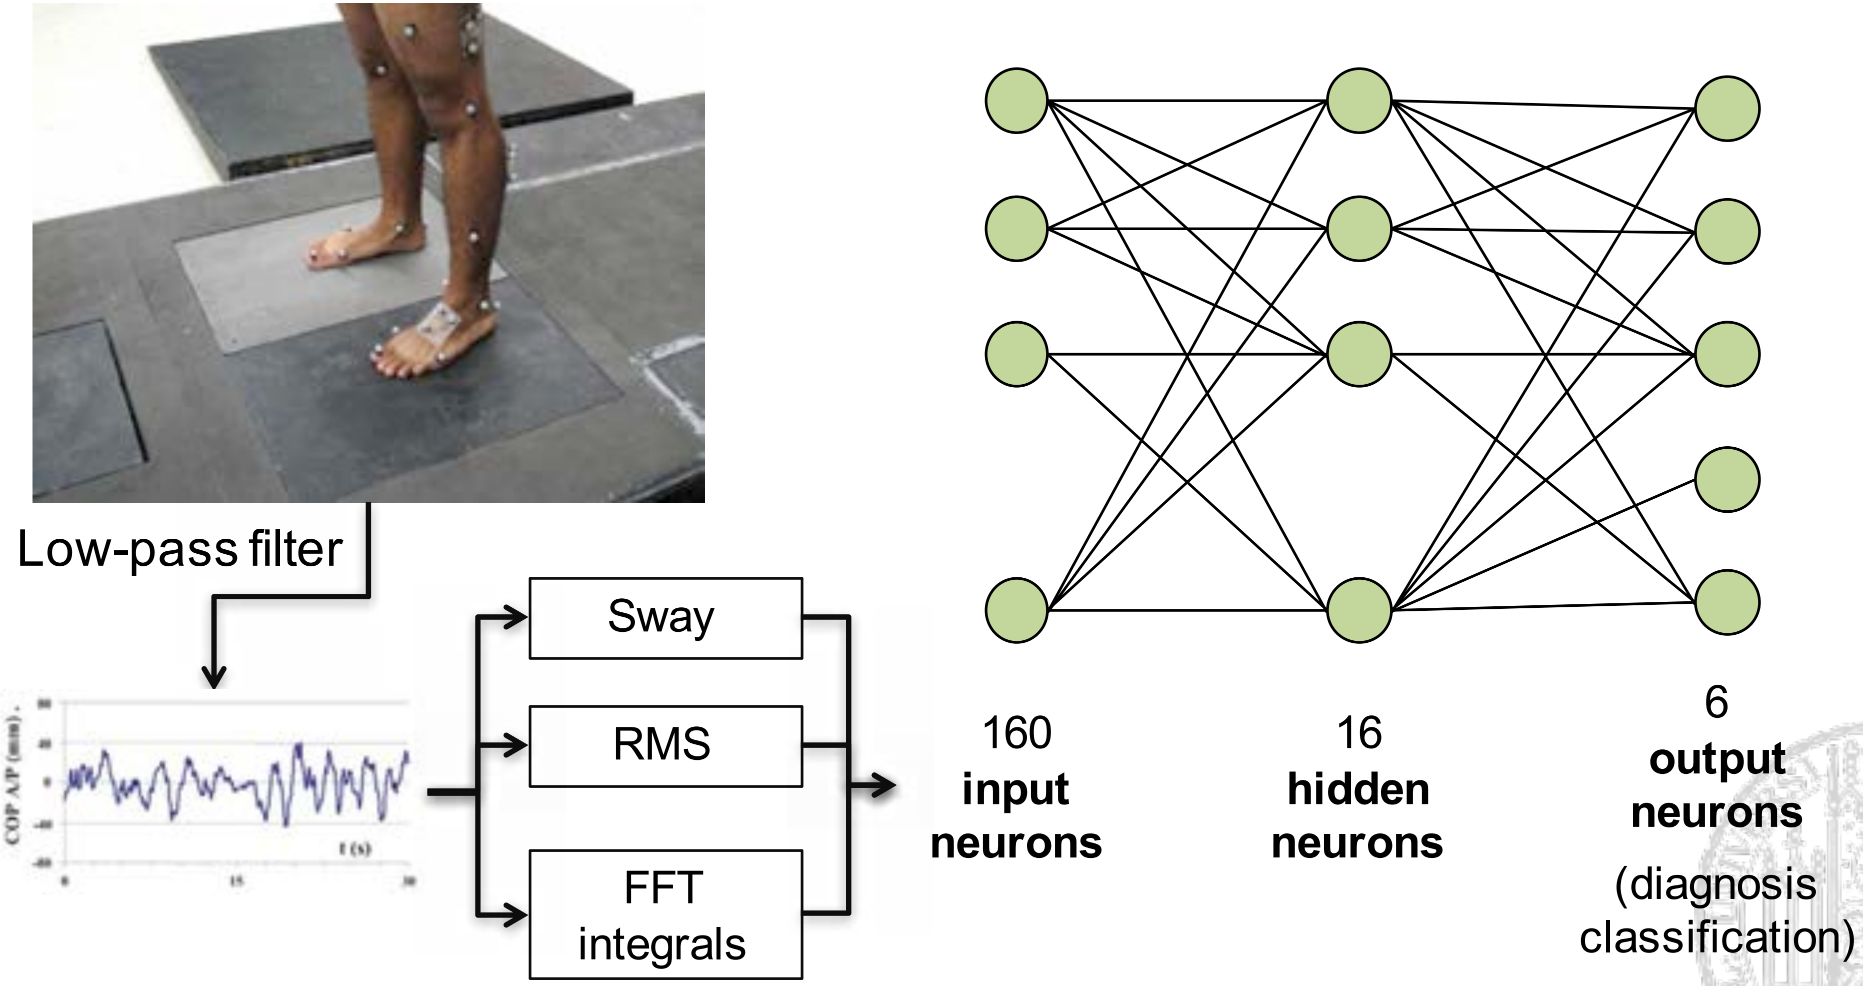

Krafczyk S, Tietze S, Swoboda W, Valkovic P, Brandt T., Artificial neural network: a new diagnostic posturographic tool for disorders of stance., Clin Neurophysiol, 117(8):1692-8, 2006

ANN example: posturography platform

2006 model

- 50 train / 50 test, no cross-validation
- One hidden layer MLP (16 hidden neurons)

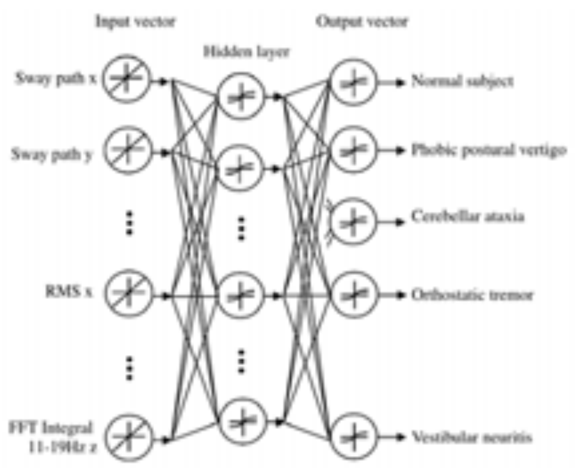

2017 models, revisited

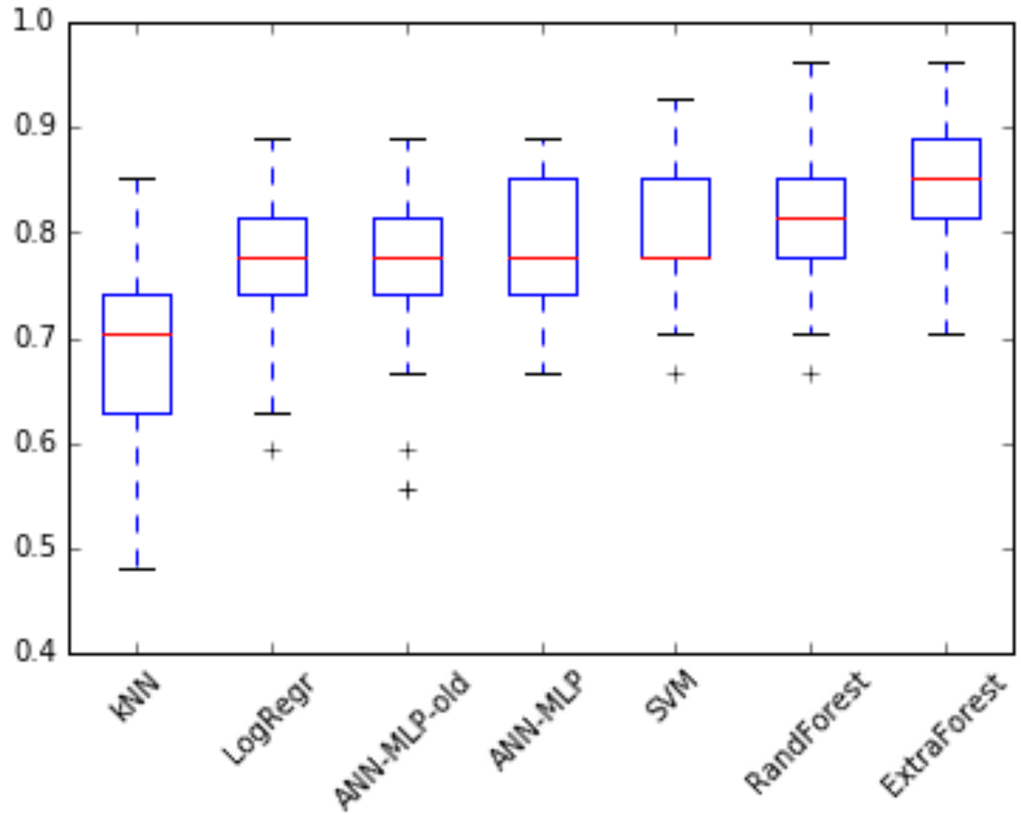

Posturography feature distribution in 2D

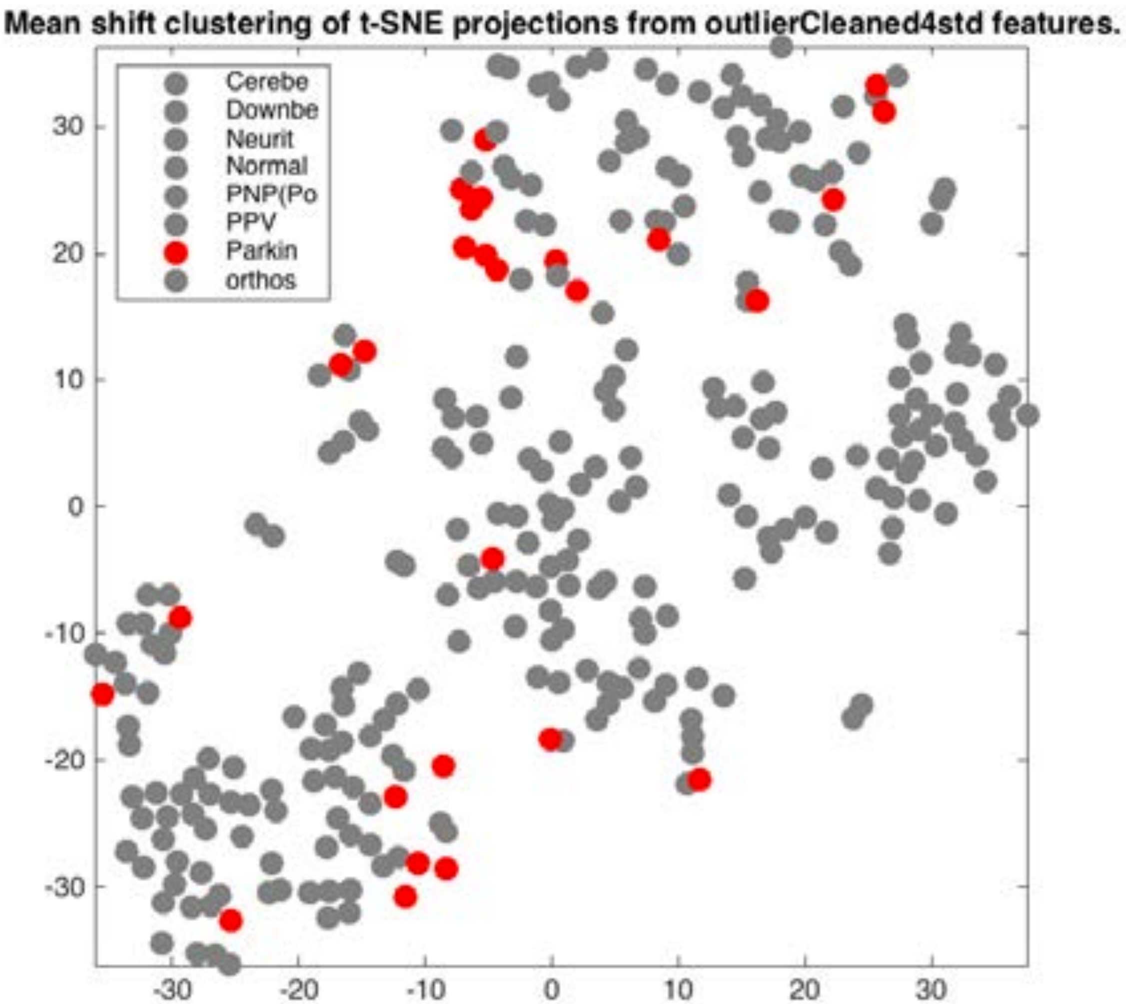

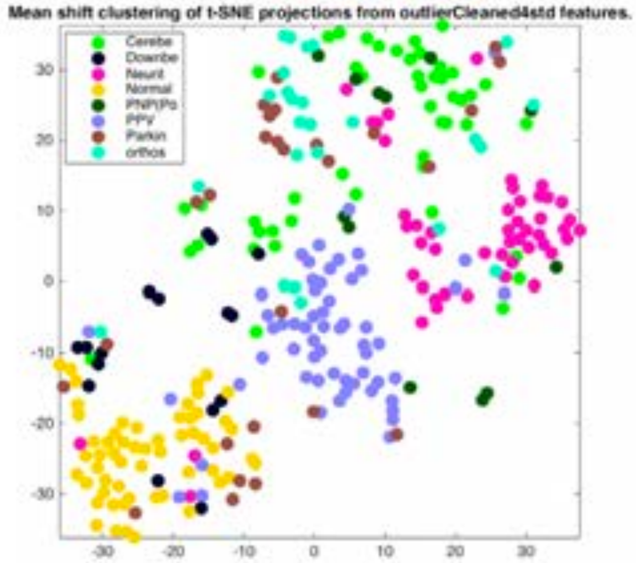

ANN example: posturography platform

Are these really the best features to extract from posturography data?

Are there maybe features which we are missing in order to distinguish Parkinson patients better from the other classes?

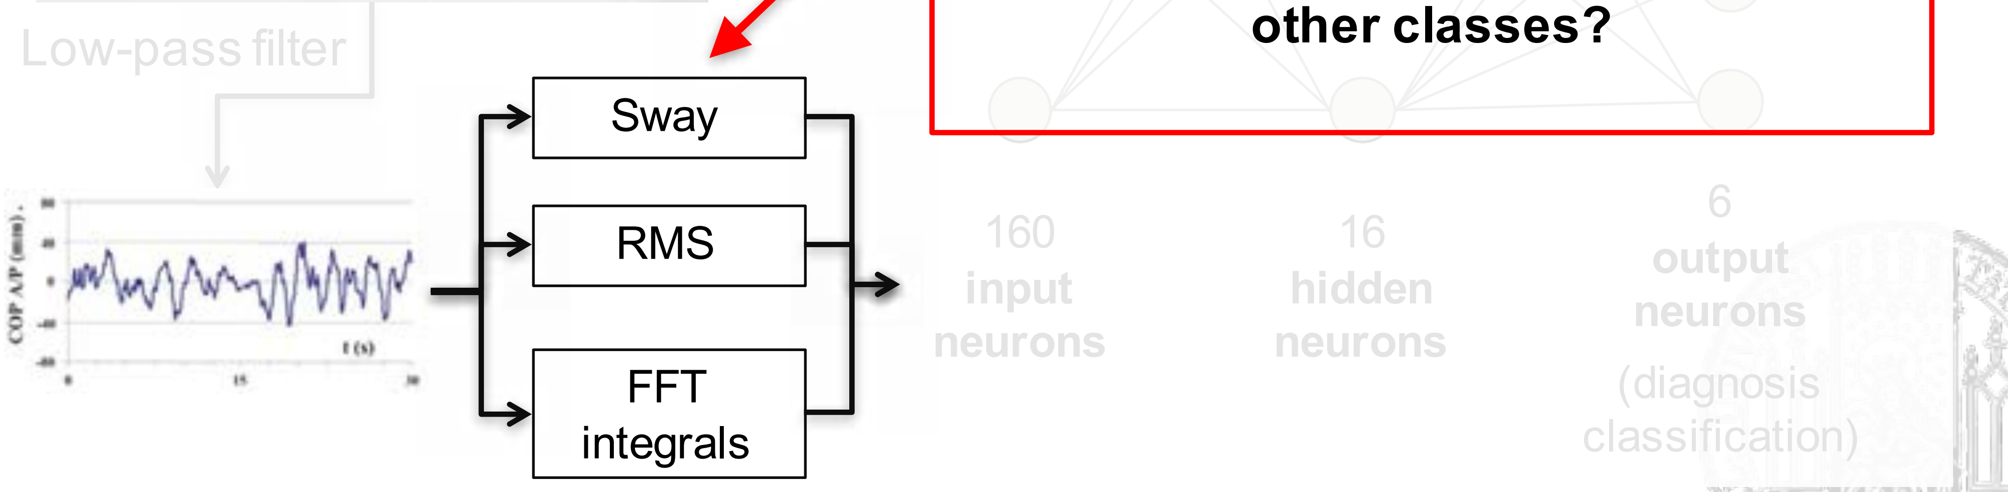

Krafczyk S, Tietze S, Swoboda W, Valkovic P, Brandt T., Artificial neural network: a new diagnostic posturographic tool for disorders of stance., Clin Neurophysiol, 117(8):1692-8, 2006

Vision: „representation learning“ on posturography data

Traditional Machine Learning:
„Learn the optimal classifier output.“

Deep Learning:
„Learn the optimal input AND classifier output.“

Deep Learning = Representation learning!

Raw spatio-temporal data

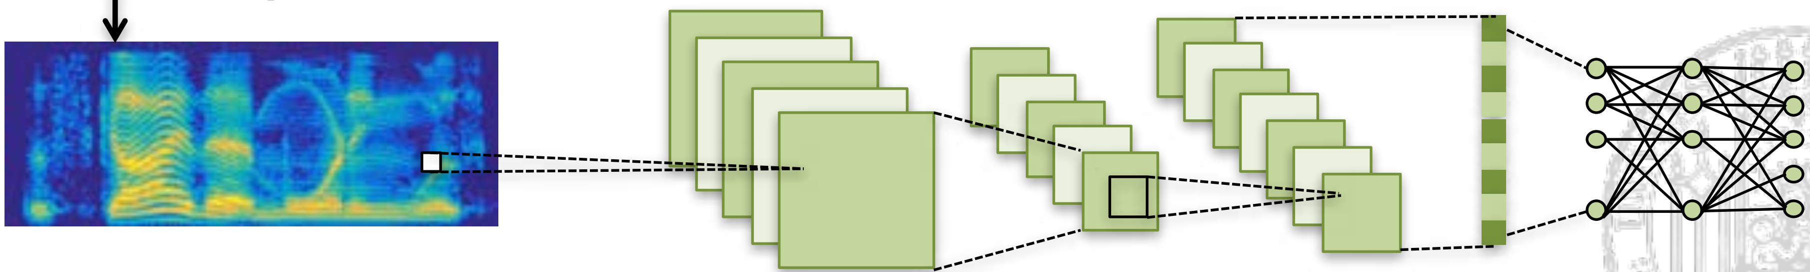

Machine learning based image segmentation pipeline

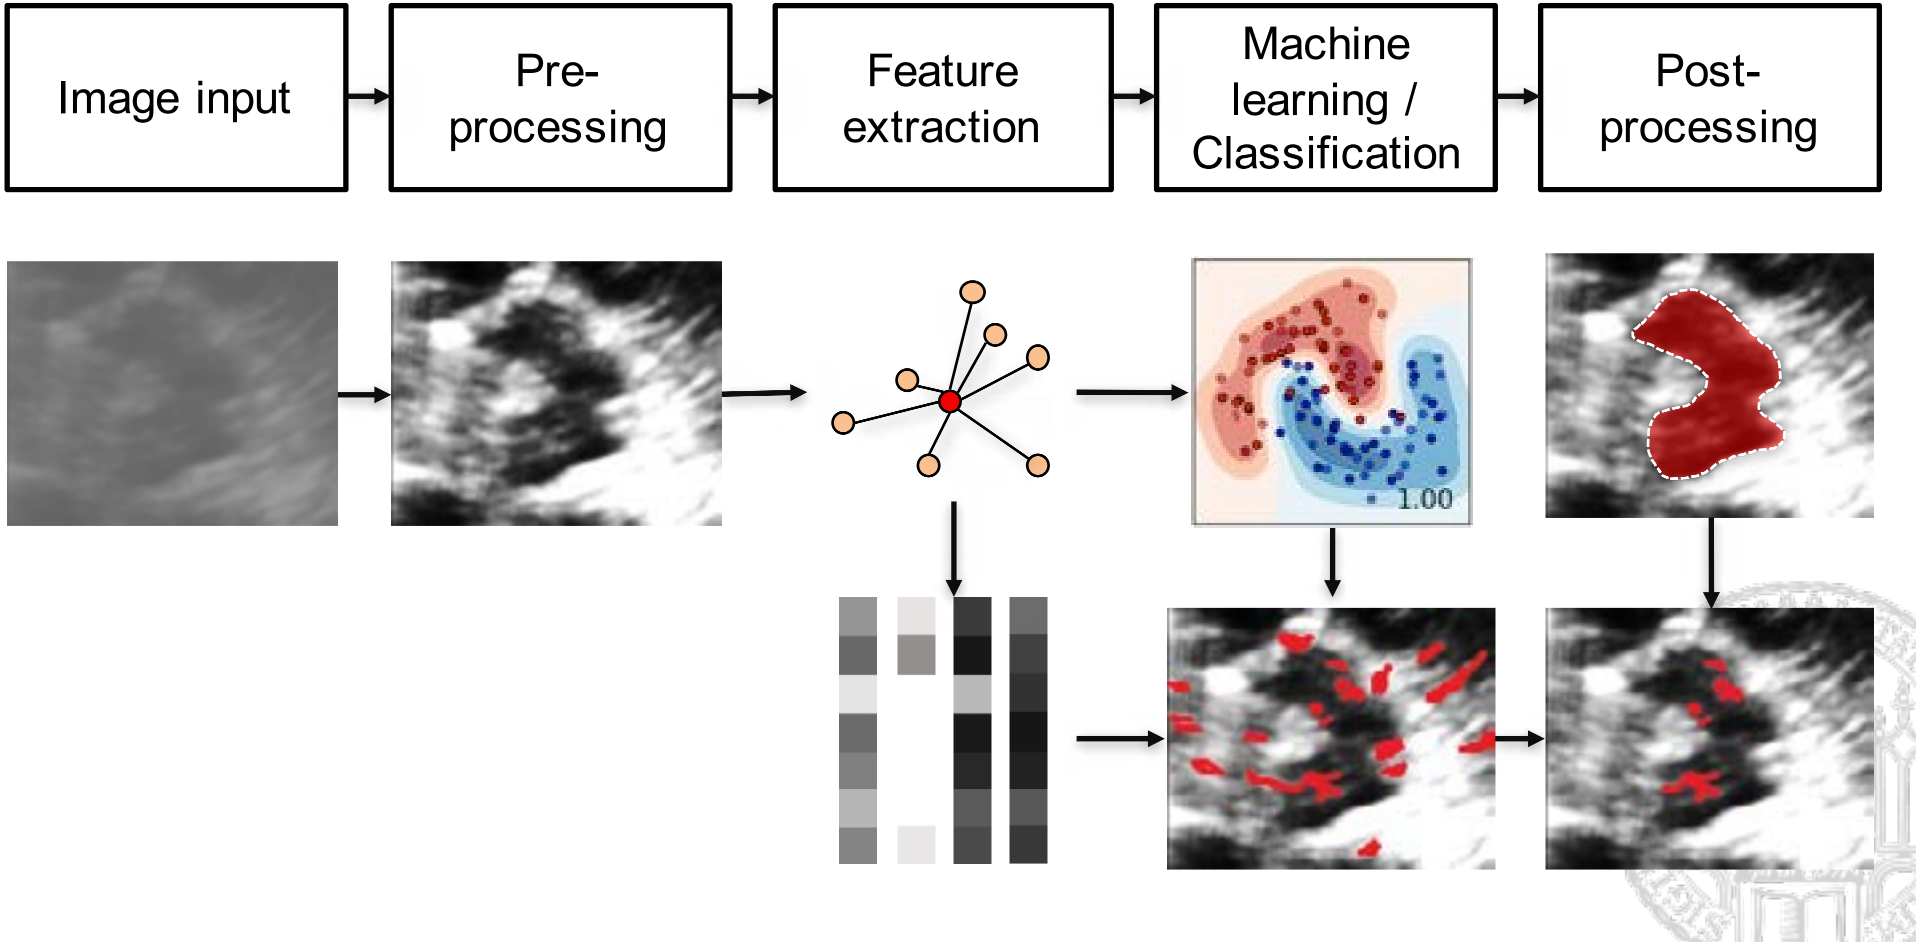

Most critical point in segmentation pipeline

Image input → Pre-processing → **Feature extraction** → Machine learning / Classification → Post-processing

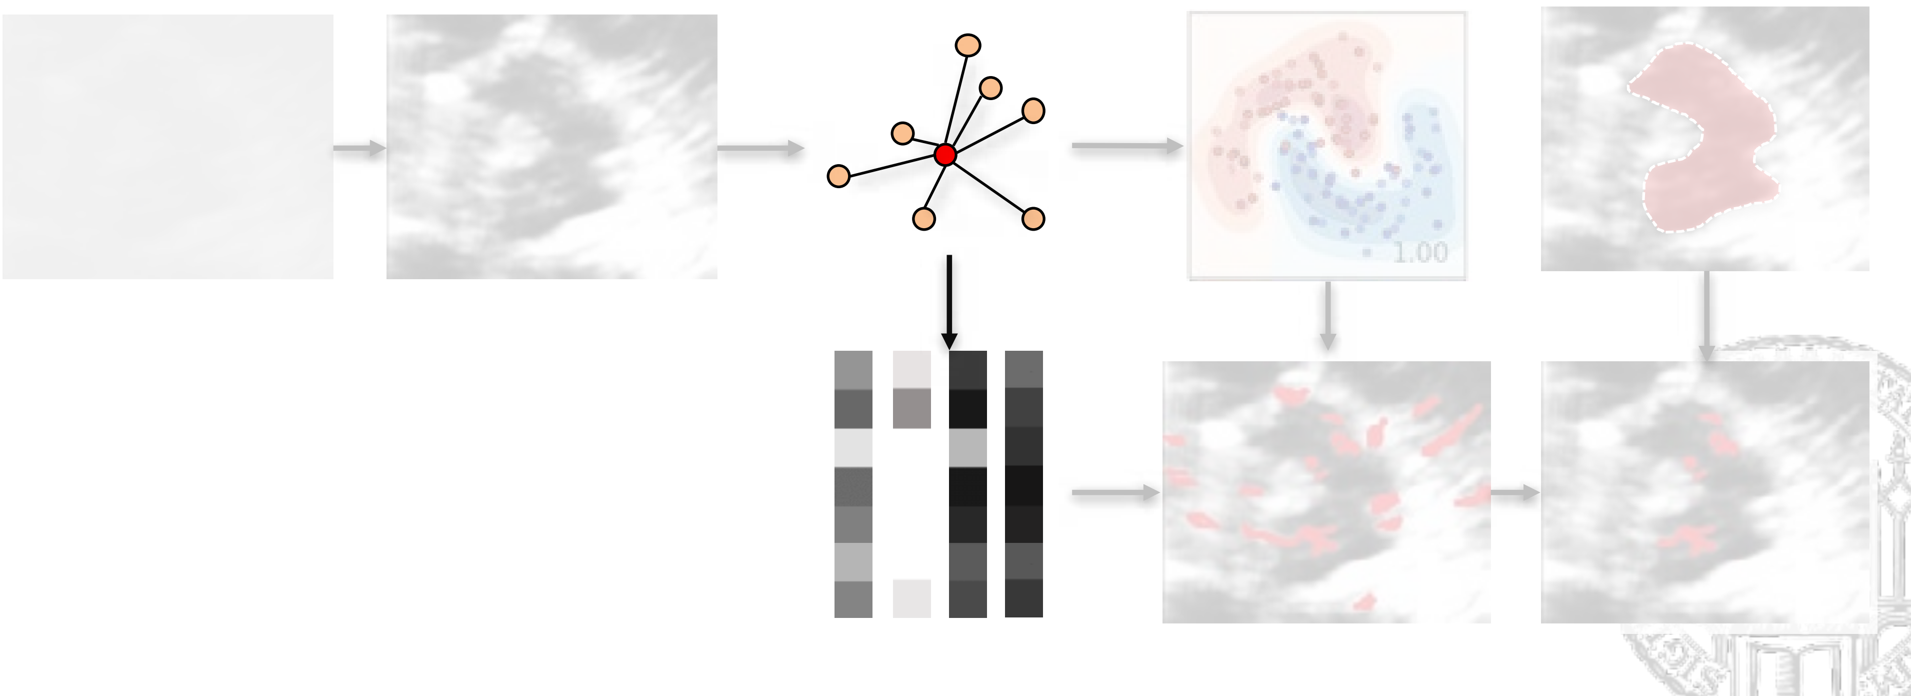

Short sidestep into computer vision:

Filter convolution for feature extraction

Computer vision: image filters and convolution

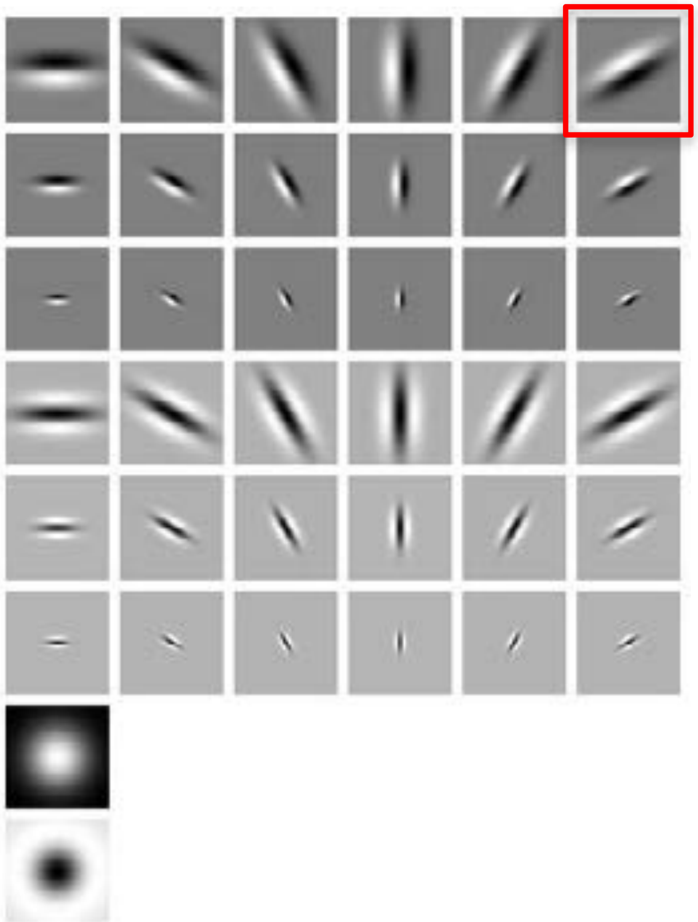

Maximum Response (MR) Filter Bank

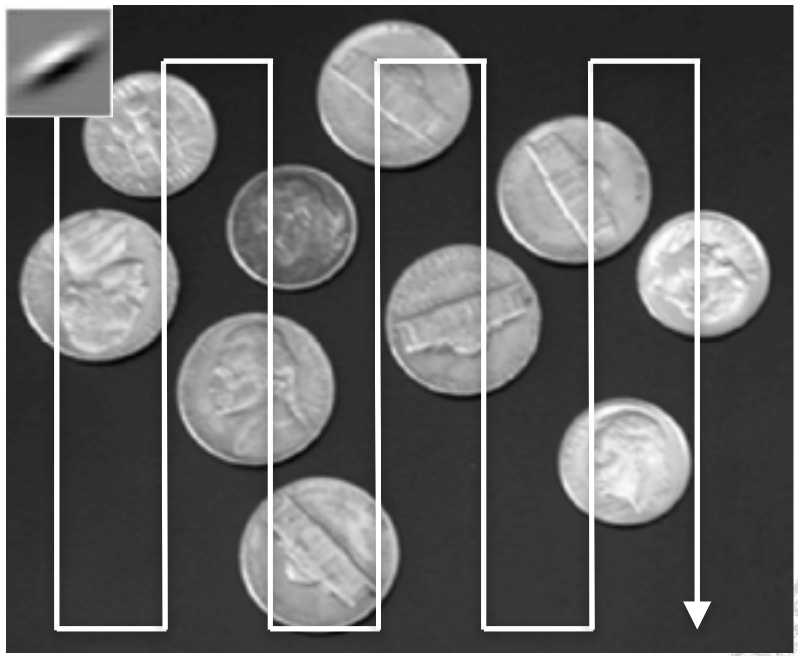

Computer vision: image filters

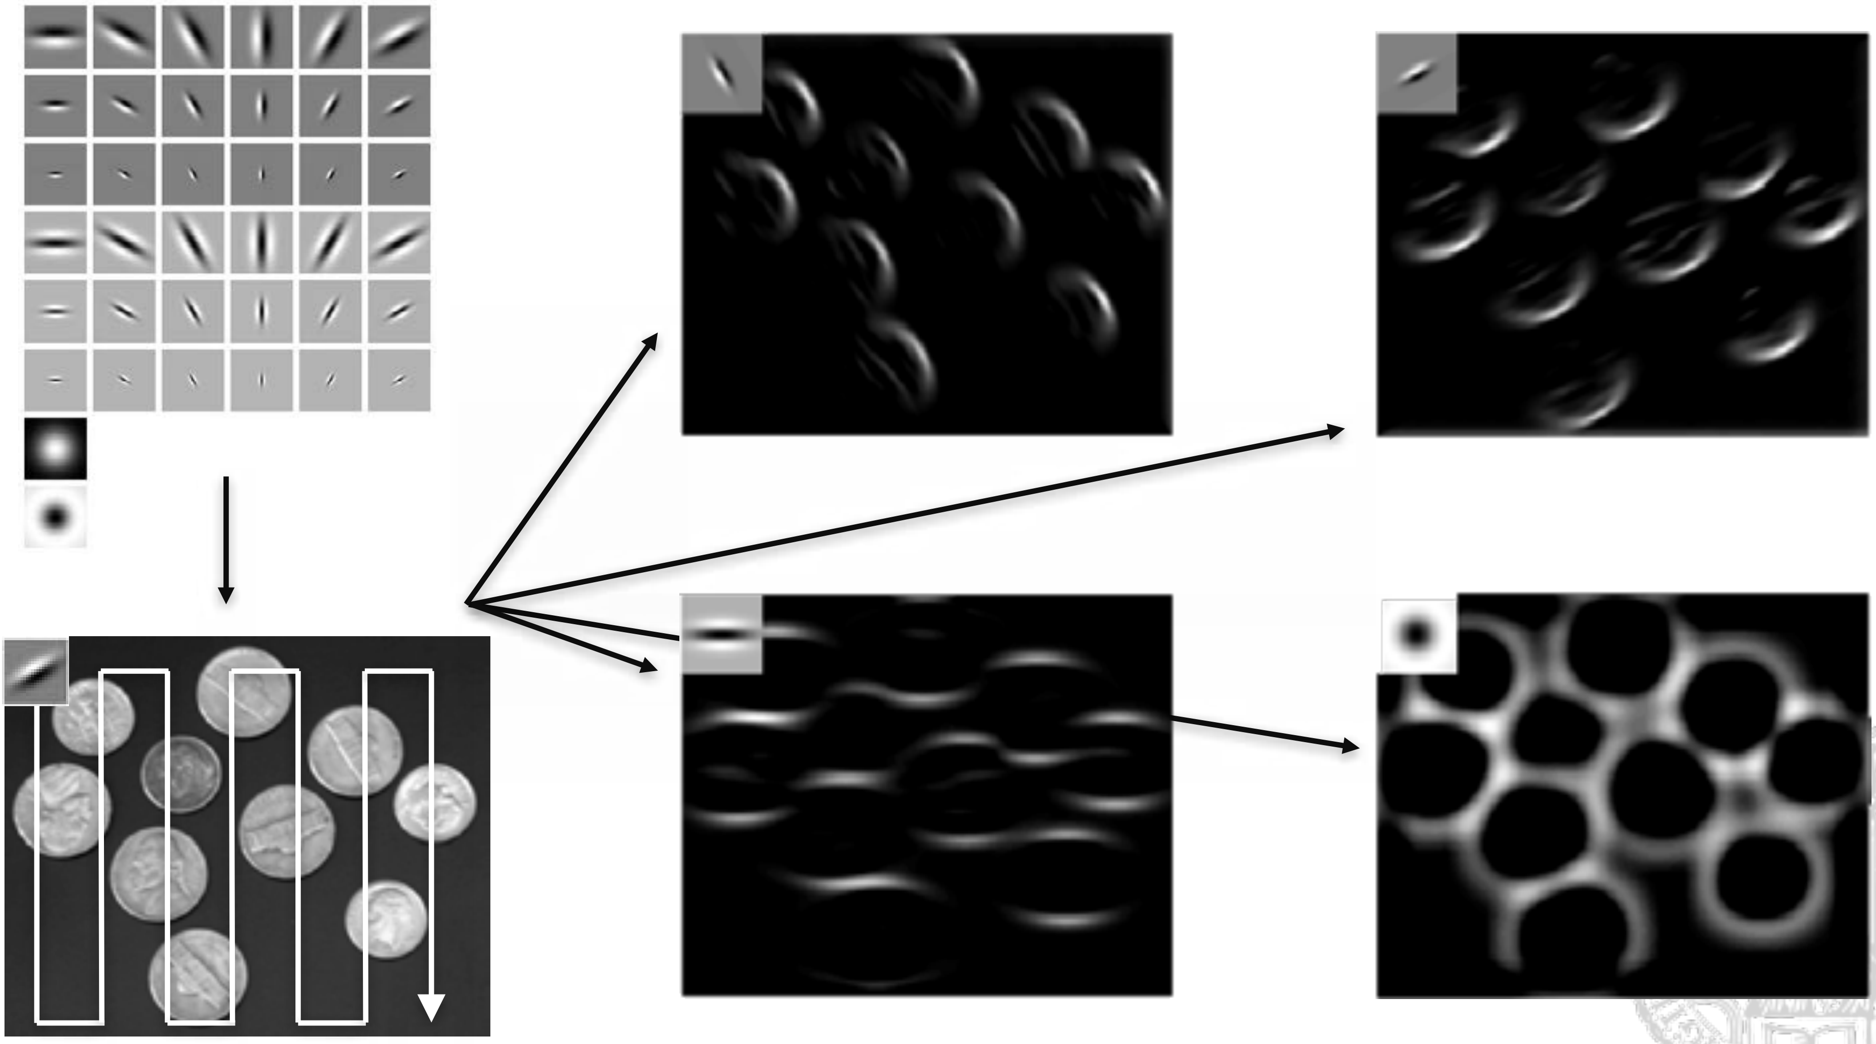

Image classification / segmentation with filter banks

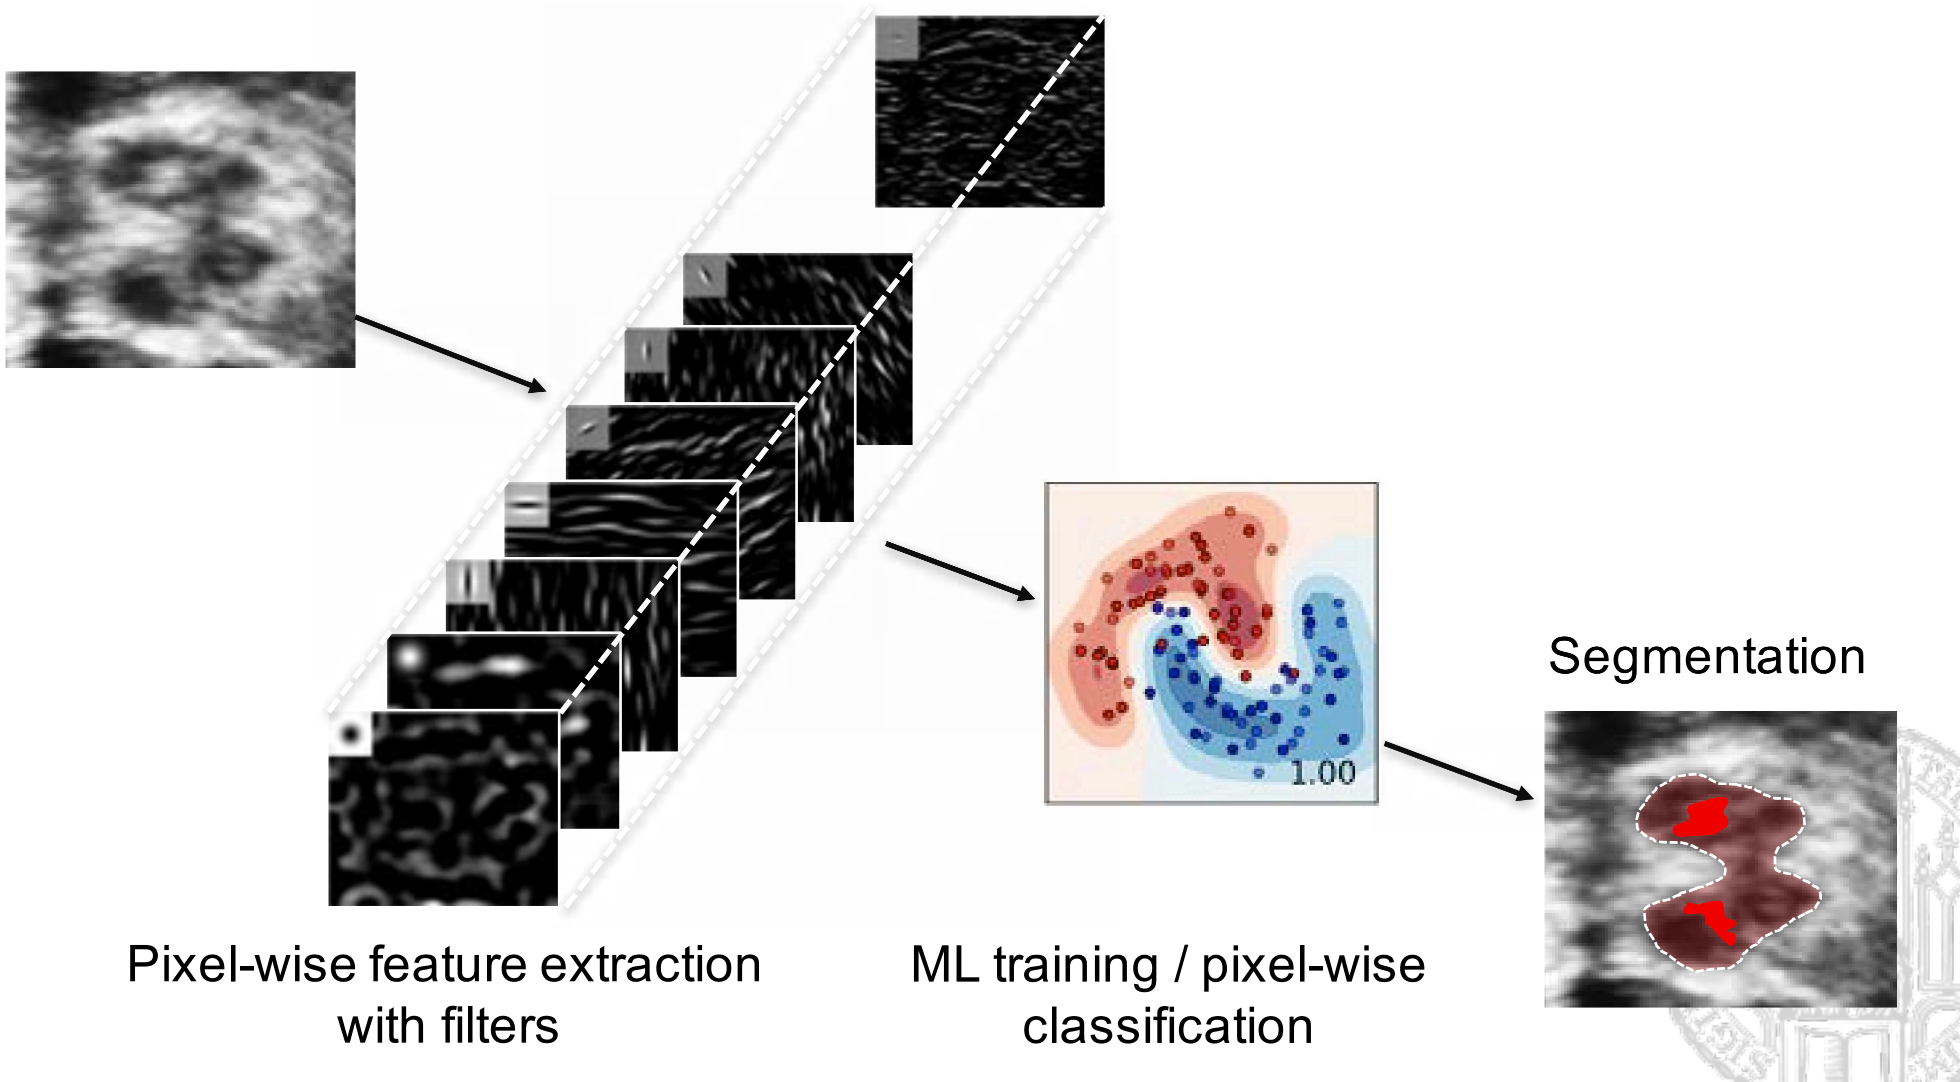

Back to neural networks and deep learning…

LMU | KLINIKUM DER UNIVERSITÄT MÜNCHEN | DSGZ

Convolutional Neural Networks: combining image filters with neural nets

- Convolutional filter: extracting local salient features
- Provide the ANN with a hierarchical architecture of filter banks
- Instead of applying pre-defined filters, initialize the filter with random “noise” and let the filters be learned by backpropagation!

LeNet (1989)

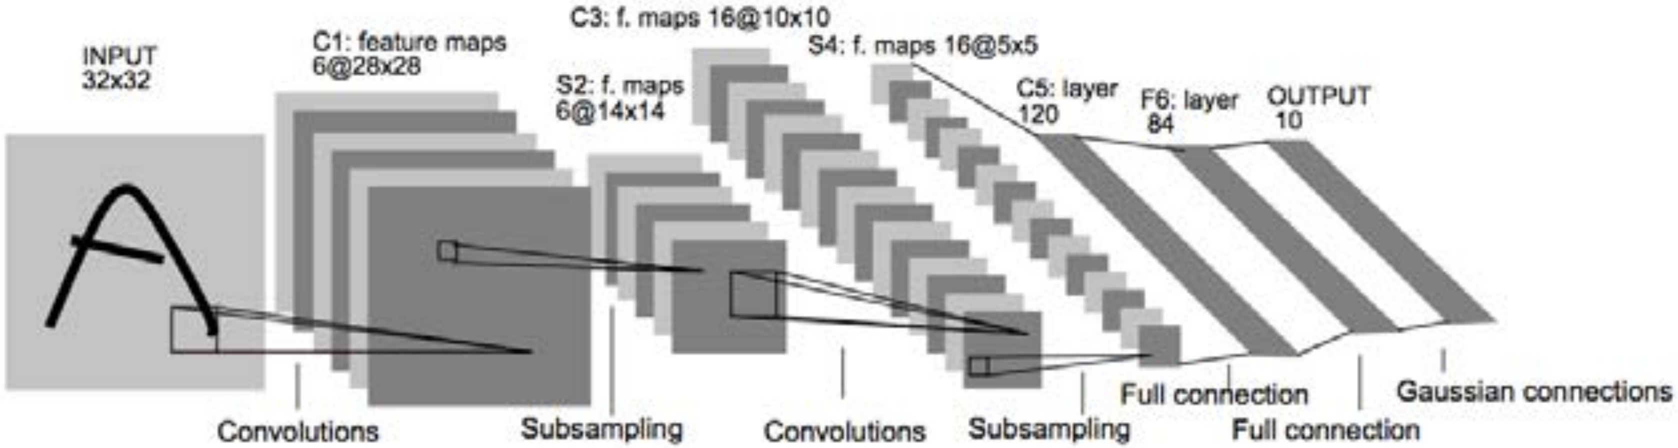

LeCun et al., Backpropagation applied to handwritten zip code recognition, Neural computation 1 (4), 541-551, 1989

CNN components: convolutional and pooling layers

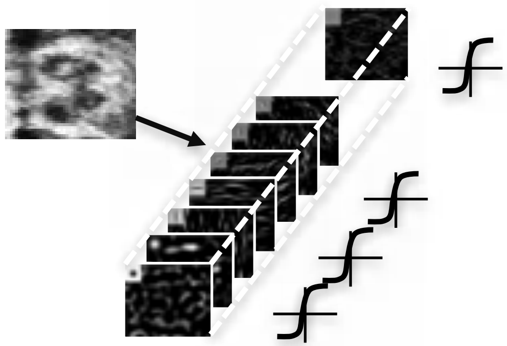

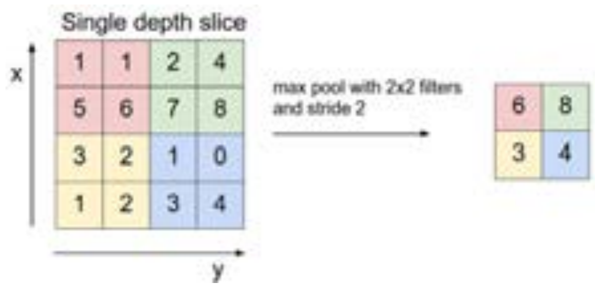

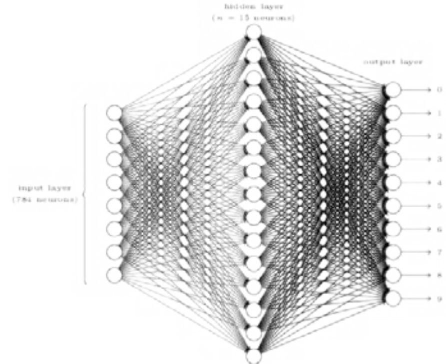

Convolutional layer (+act.fct.) Pooling layer Multi-layer perceptron

LeNet (1989)

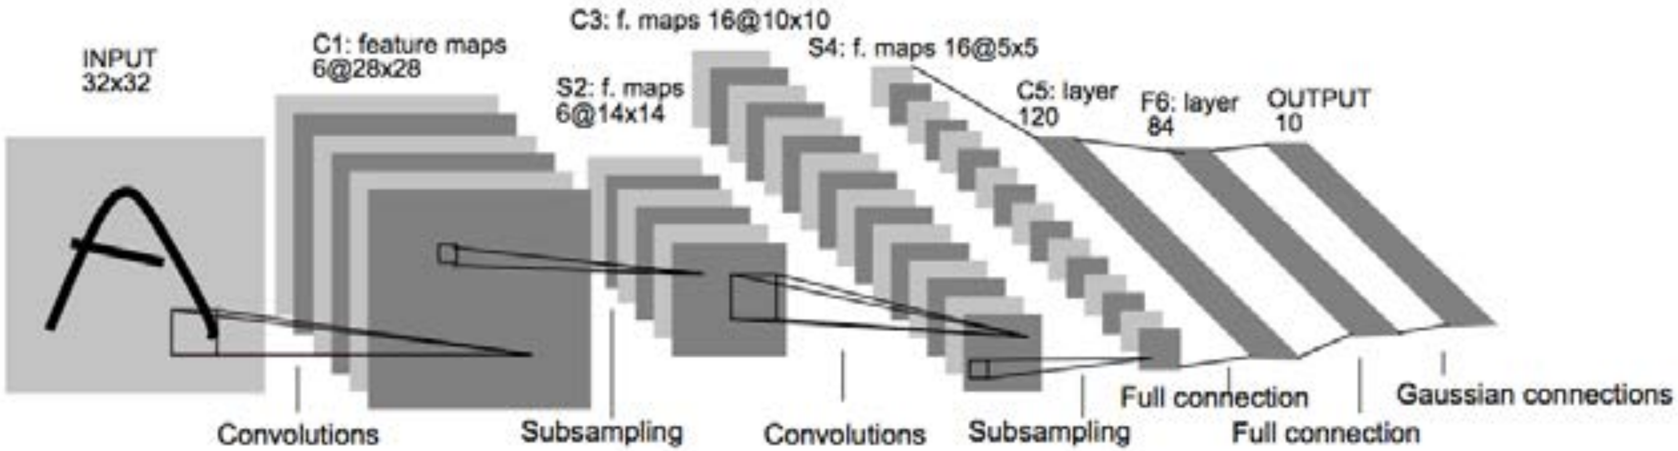

LeCun et al., Backpropagation applied to handwritten zip code recognition, Neural computation 1 (4), 541-551, 1989

LMU KLINIKUM DER UNIVERSITÄT MÜNCHEN DSGZ

CNN Demo: LeNet architecture from 1989

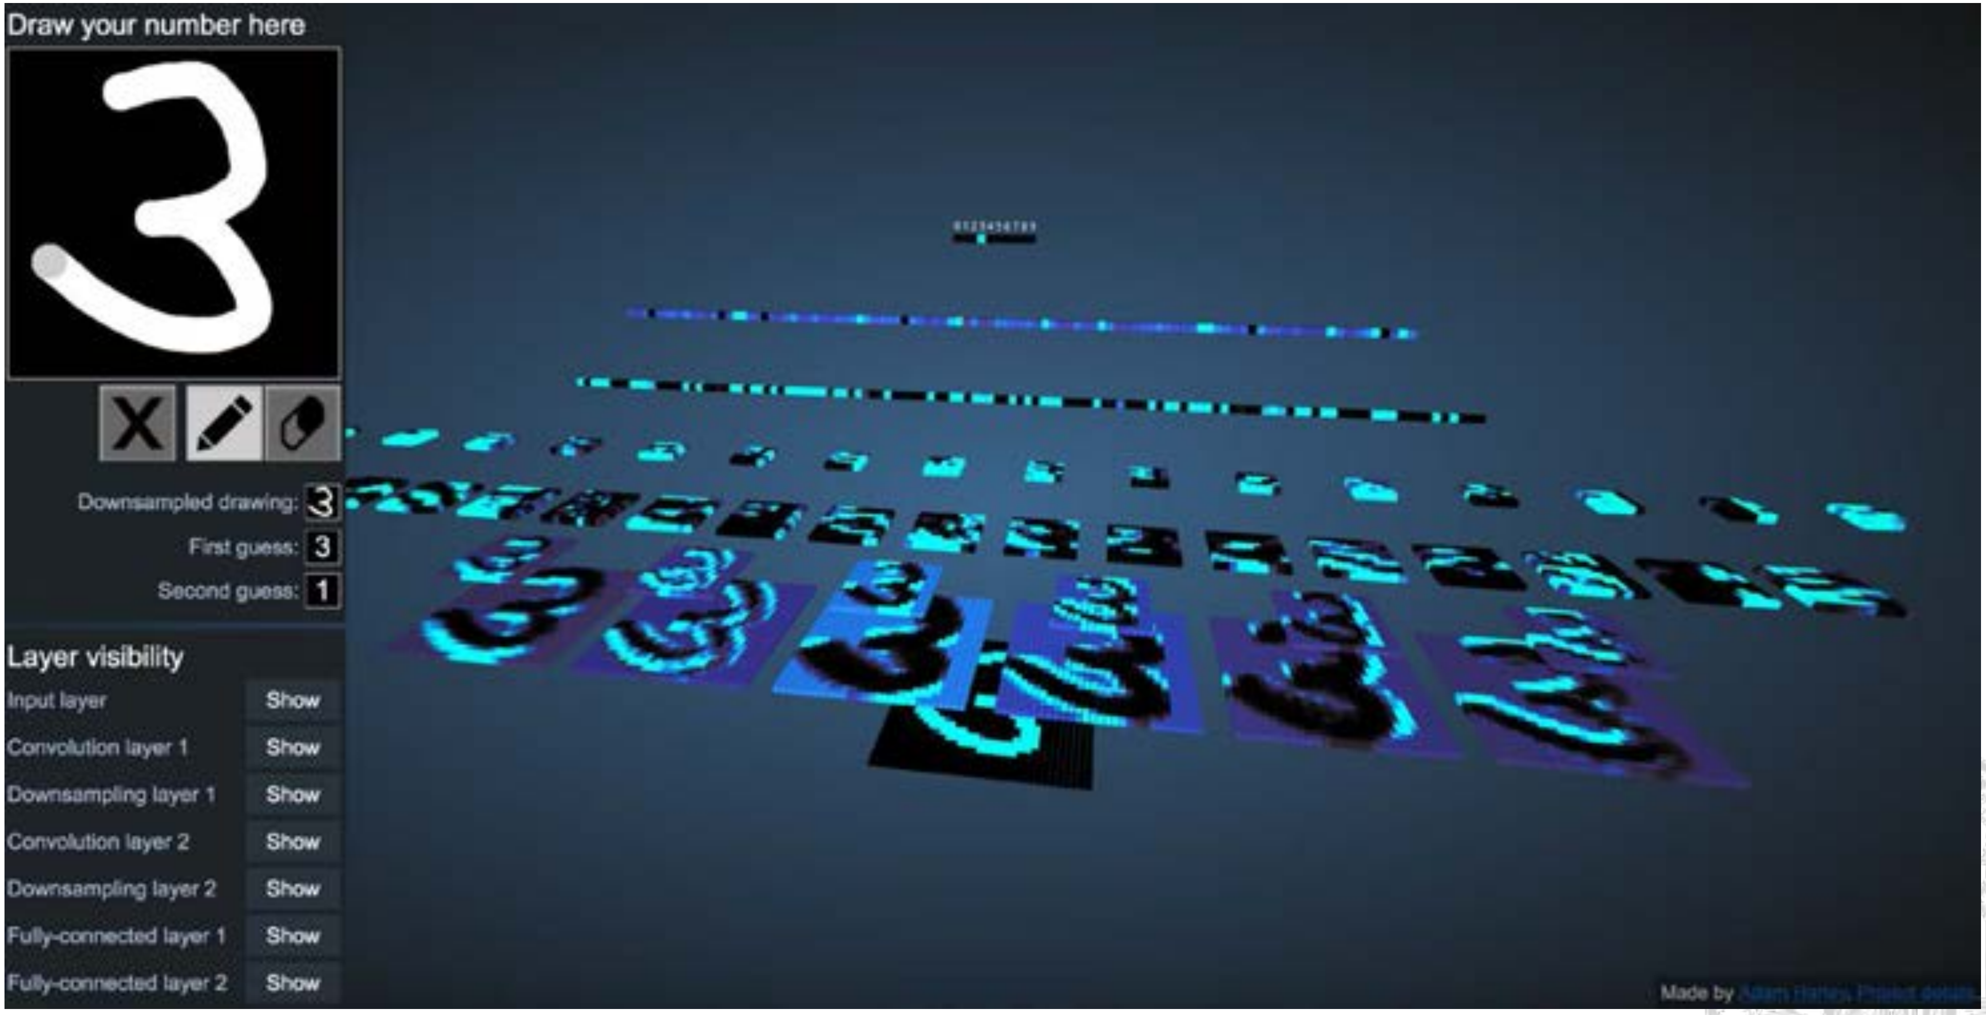

http://scs.ryerson.ca/~aharley/vis/conv/

Modern computer vision datasets: ImageNet

- Published in 2010
- 10 million labeled images with 10.000+ classes

Breakthrough performances in 2012 and 2015

- ImageNet 2012: 1.3 million images, 1000 classes
- 2012: AlexNet architecture: 5 conv layers, 500.000 neurons, 60 million parameters
- GPU acceleration of training
- 2015: Better than human performance (He et al., ICCV, 2015)

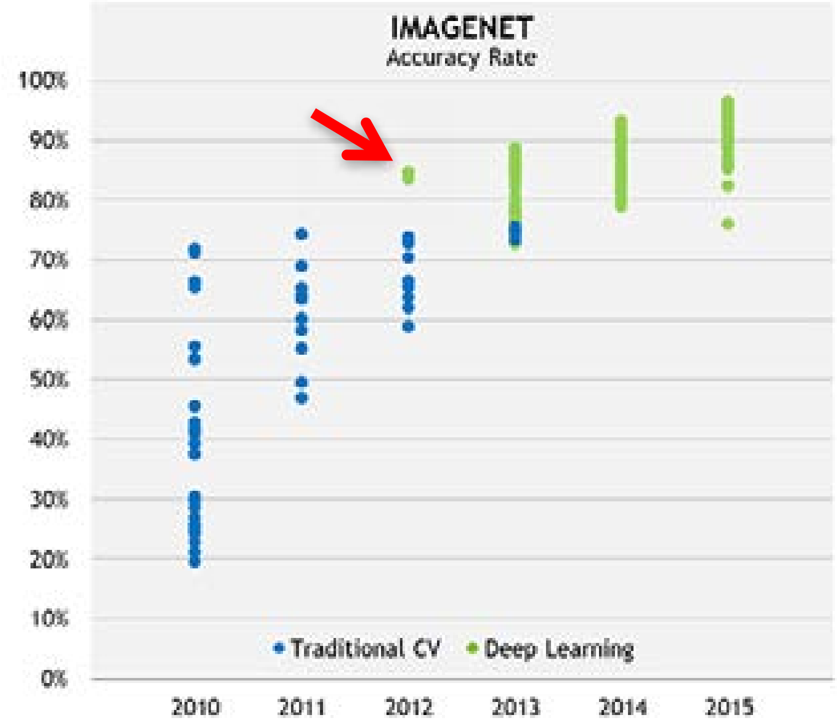

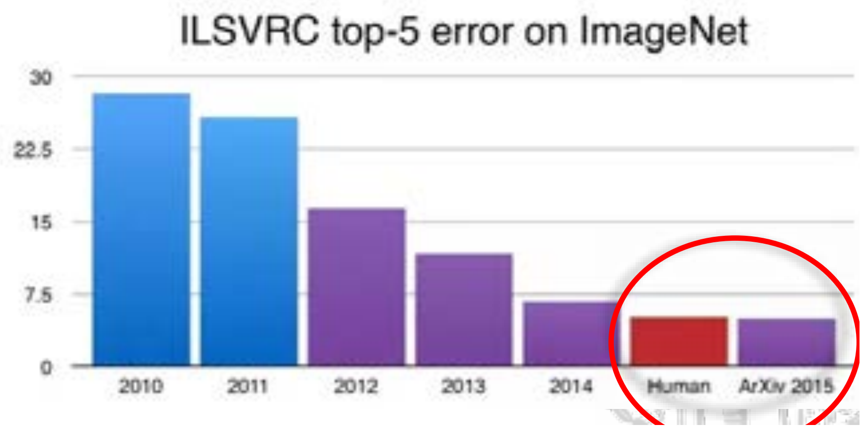

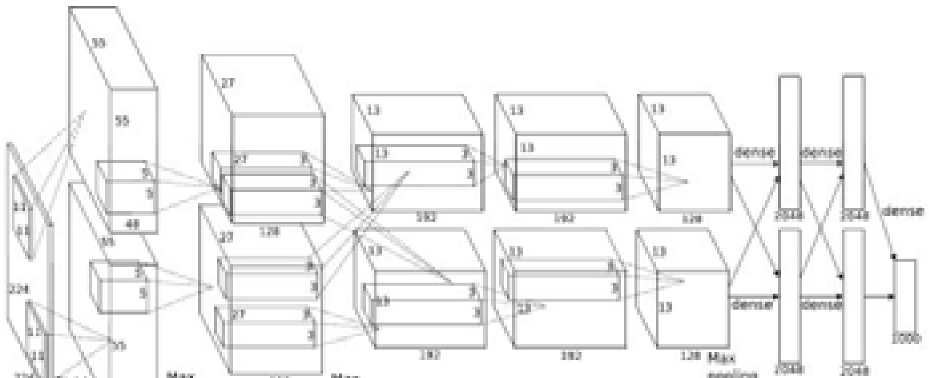

AlexNet

Visualizing and Understanding CNNs

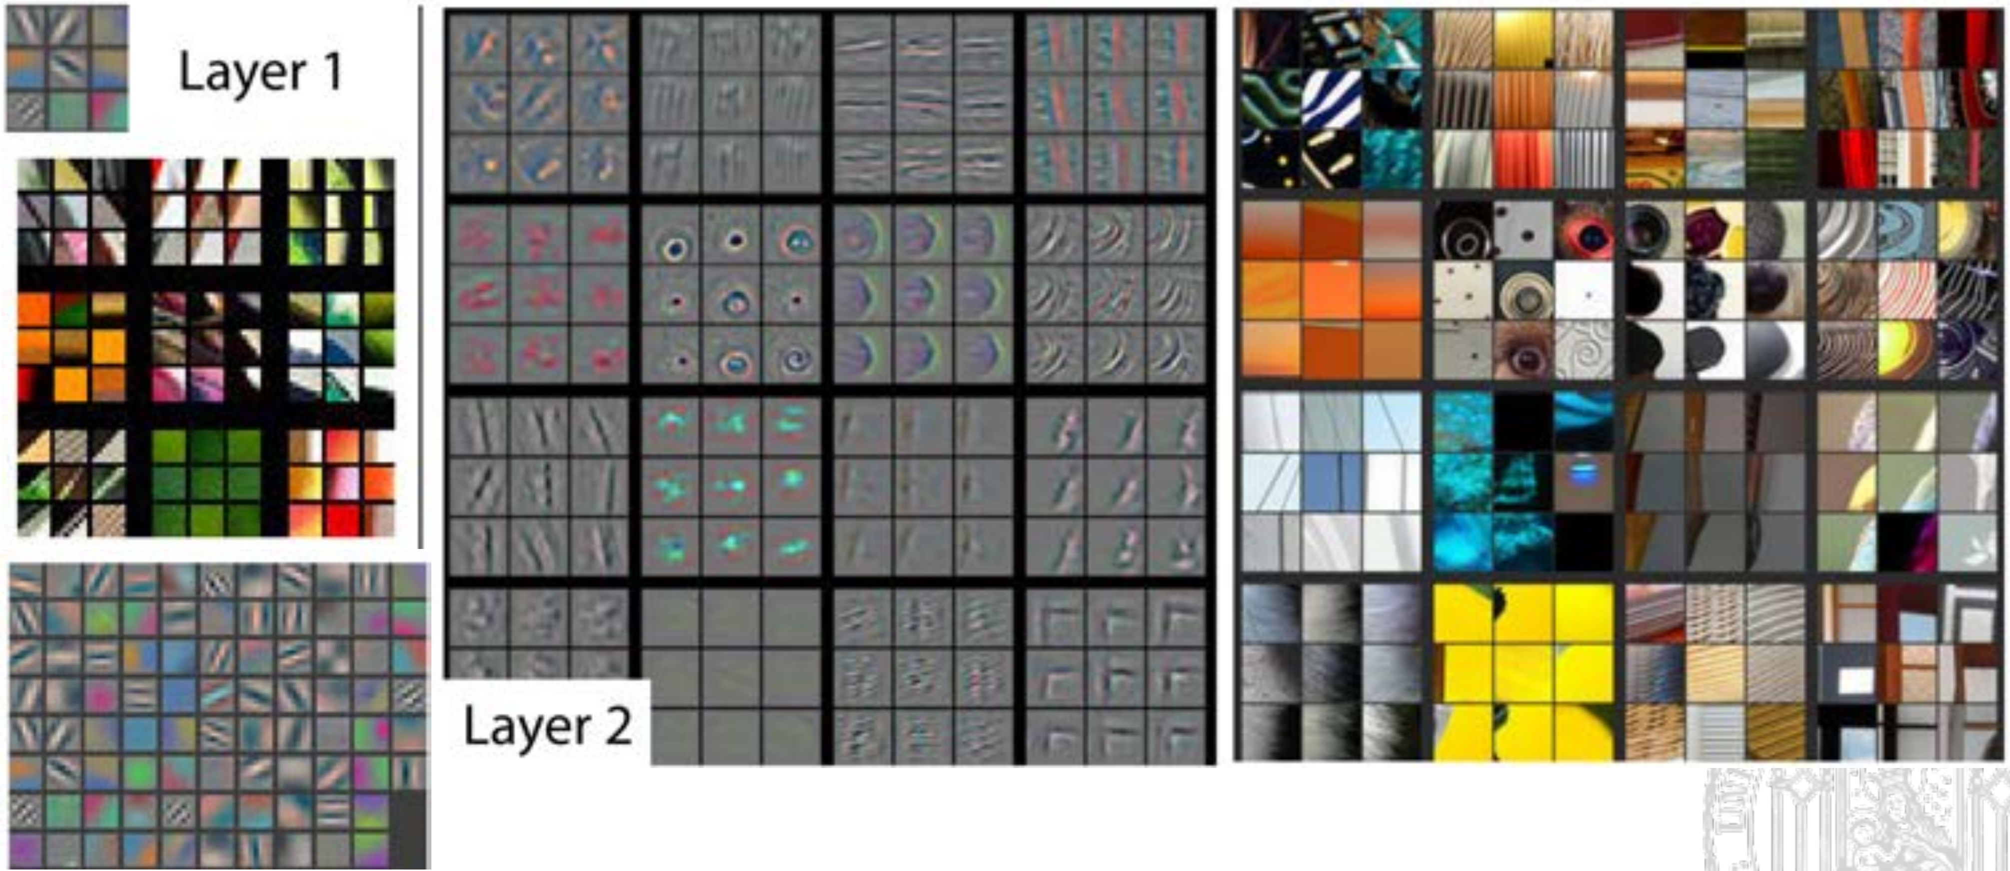

Zeiler and Fergus, Visualizing and Understanding Convolutional Networks, ECCV, 2014

Visualizing and Understanding CNNs

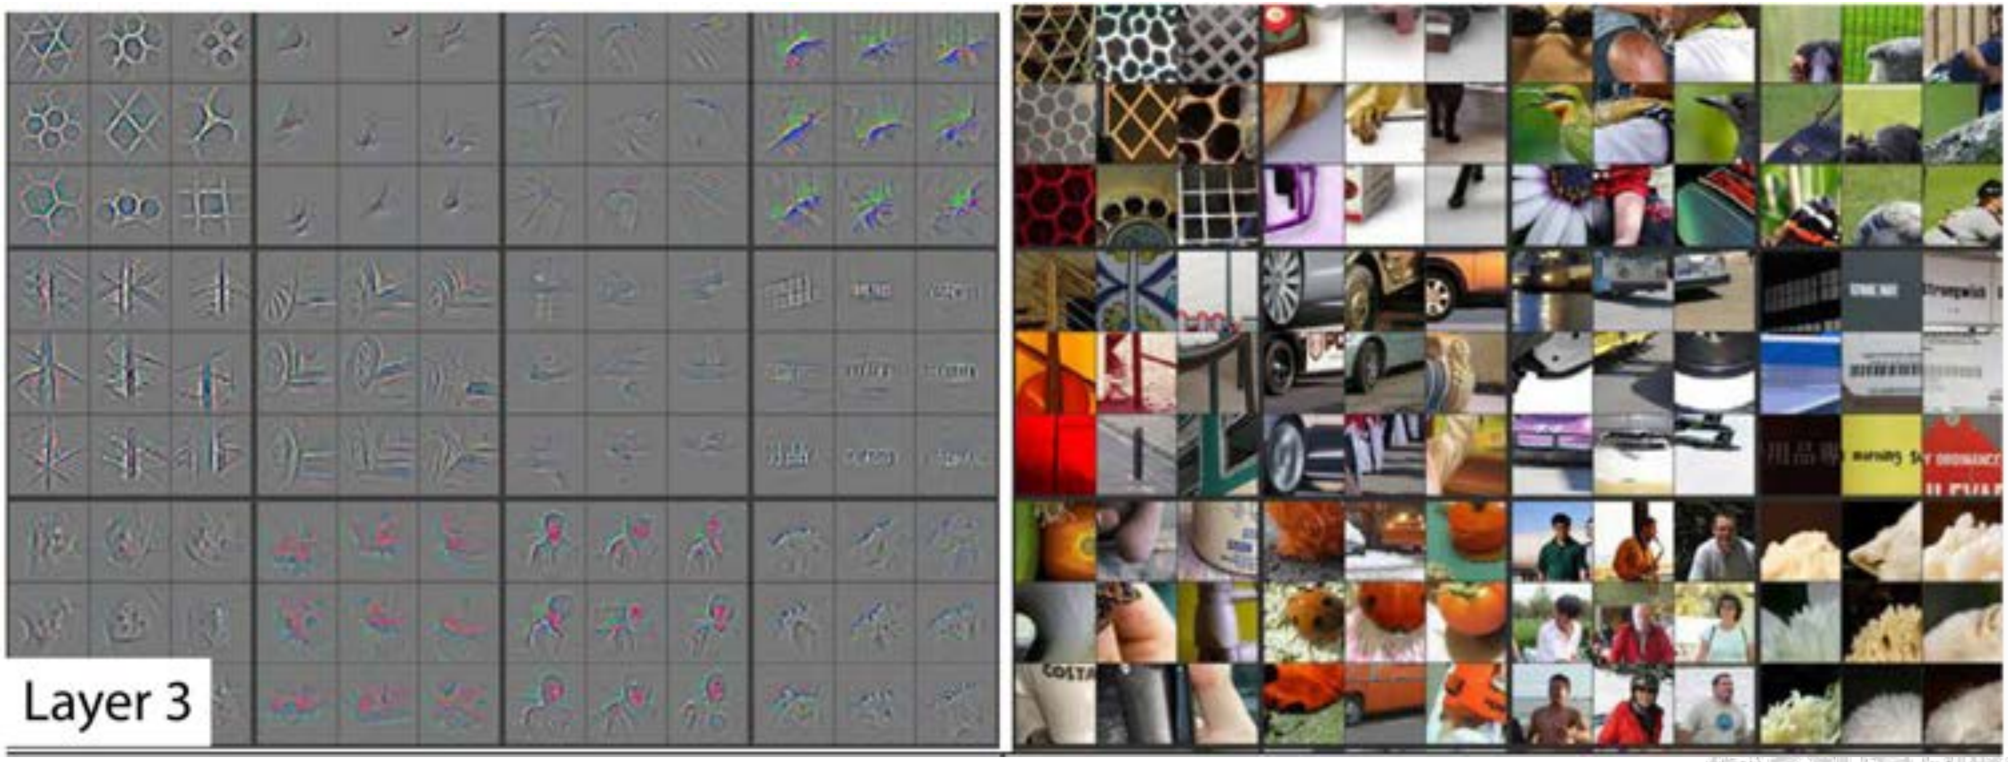

Zeiler and Fergus, Visualizing and Understanding Convolutional Networks, ECCV, 2014

Visualizing and Understanding CNNs

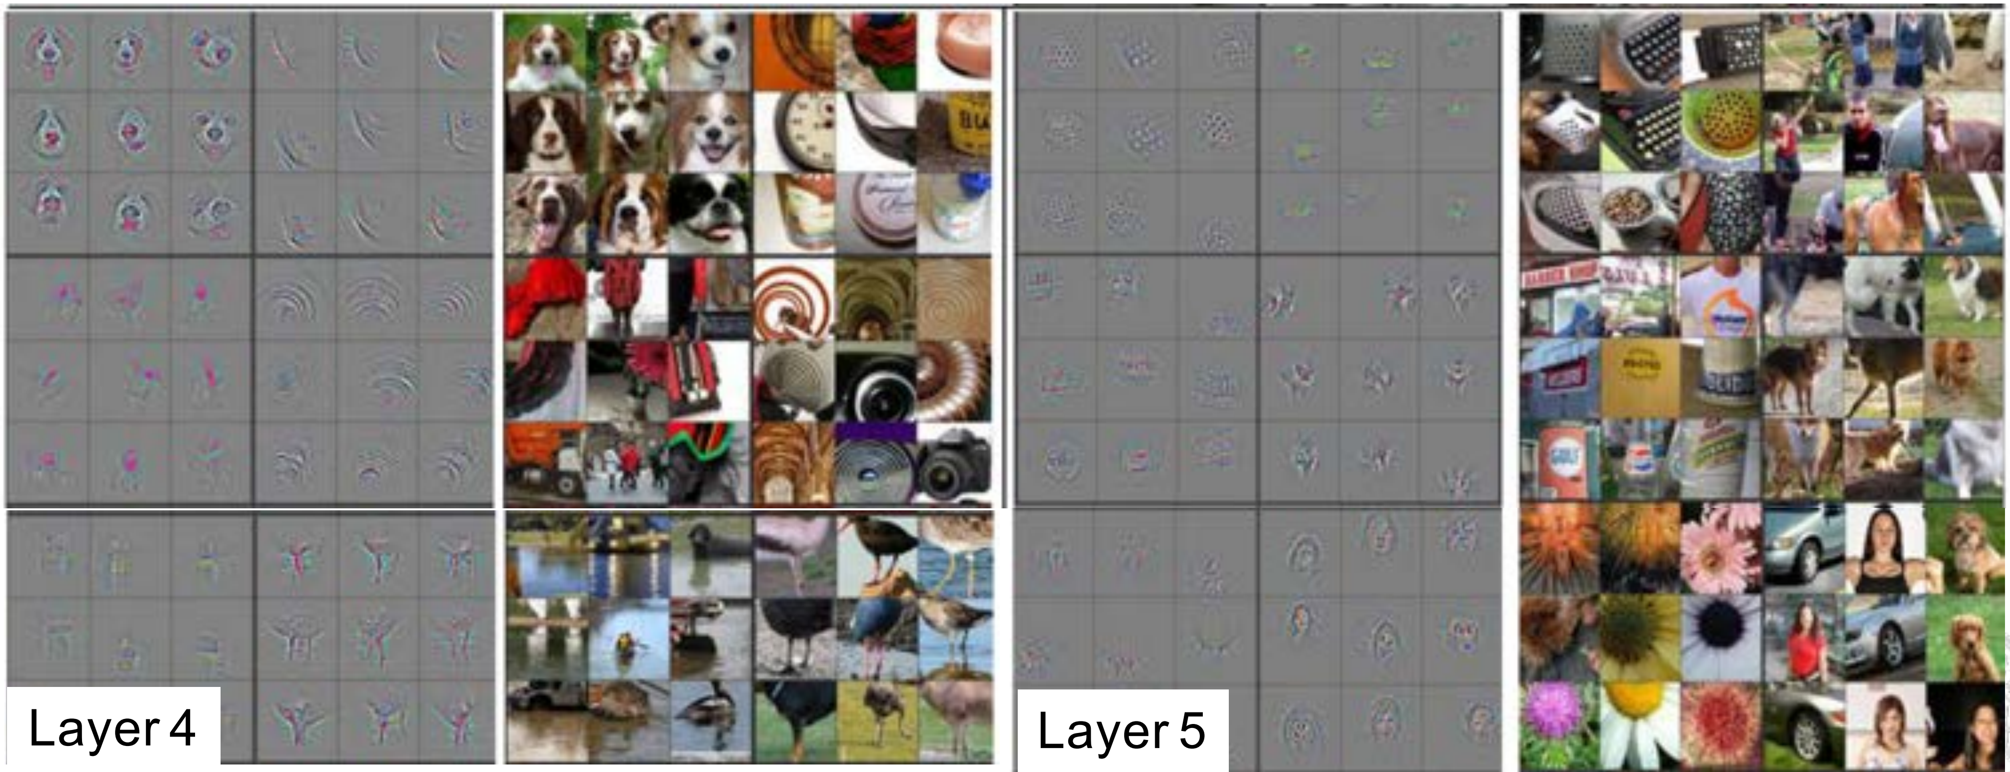

Zeiler and Fergus, Visualizing and Understanding Convolutional Networks, ECCV, 2014

Deep Learning success

Google

Citation indices	All	Since 2010
Citations	117128	47516
h-index	113	86
i10-index	273	200

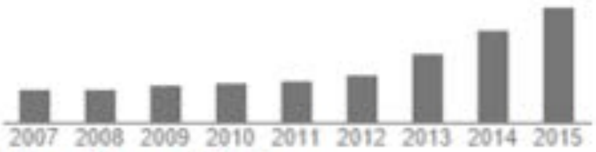

Geoffrey Hinton

Citation indices	All	Since 2010
Citations	29582	17815
h-index	77	59
i10-index	179	141

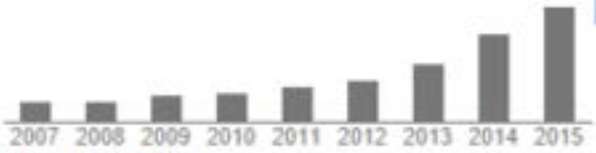

Yann LeCun

facebook

McGill UNIVERSITY

ELEMENT AI

Citation indices	All	Since 2010
Citations	32736	25285
h-index	73	65
i10-index	245	200

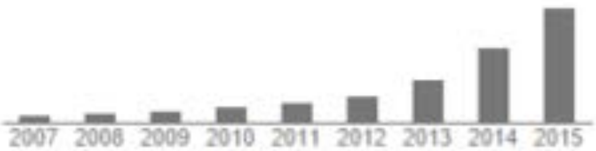

Yoshua Bengio

Citation indices	All	Since 2010
Citations	15412	10292
h-index	64	48
i10-index	242	178

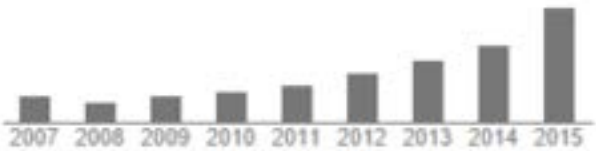

Juergen Schmidhuber

USI/SUPSI IDSIA

nnaisense

LMU

KLINIKUM
DER UNIVERSITÄT MÜNCHEN

DSGZ

Deep Learning success

RACE FOR AI: MAJOR ACQUIRERS IN ARTIFICIAL INTELLIGENCE
2011 - 2016 YTD (12/1/16)

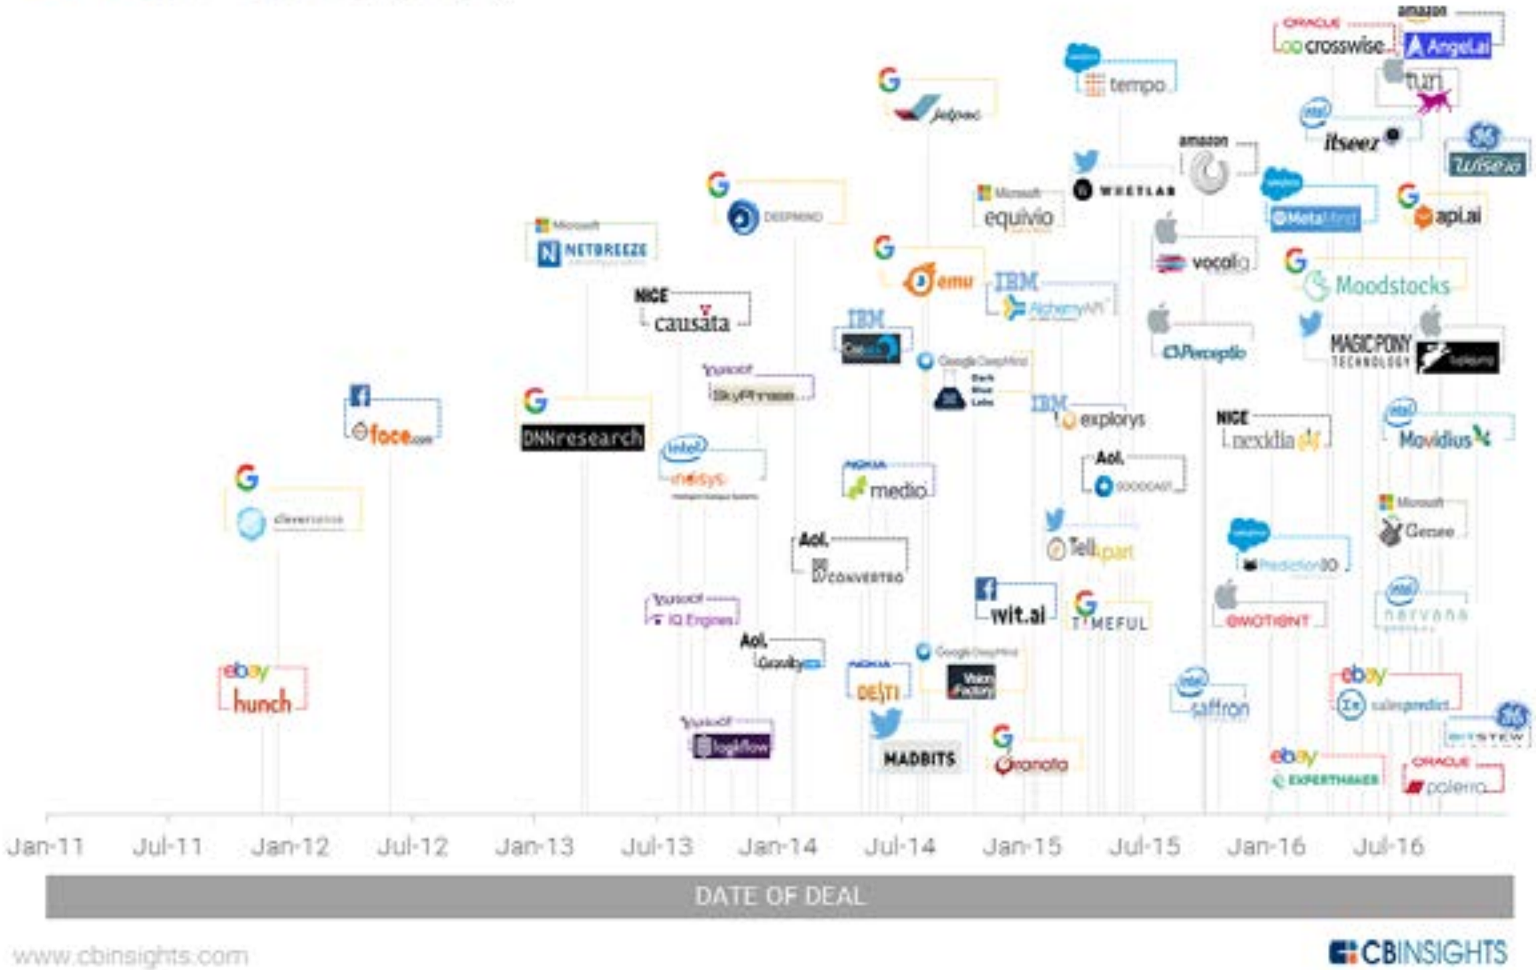

Better training algorithms, faster hardware, deeper models, higher accuracy

AlexNet

VGG

GoogLeNet

ResNet

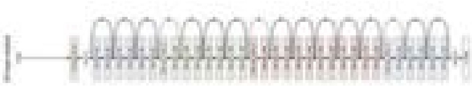

Error on ImageNet
(top 5 classification)

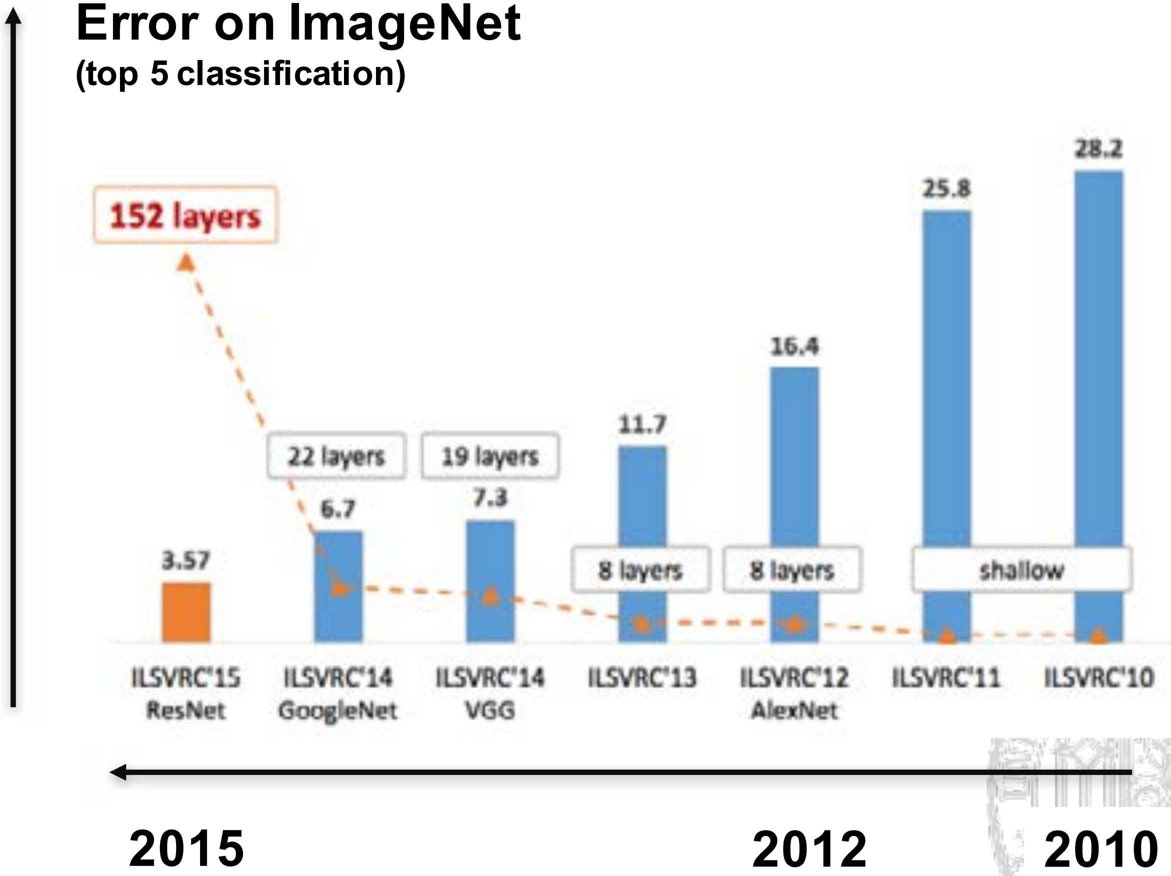

Major Deep Learning achievements

- 2012: AlexNet beats ImageNet challenge by wide margin
- 2013: DeepMind‘s deep learner teaches itself to play Atari games and beats humans at 31 different games (Google buys DeepMind, starting the „Race for AI“)
- 2014: Facebook‘s DeepFace achieves human-level performance on face detection of 4000 subjects (from 4 million images)
- 2015: Microsoft surpasses human performance on ImageNet using ResNet architecture with residual connections
- 2016: AlphaGo beats world champion Lee Sedol in Go using deep CNNs and reinforcement learning
- 2017: Dermatologist-level skin cancer identification

Dematologist-level skin cancer classification

- Using InceptionV3 CNN architecture
- Pre-trained on natural images!
- 130.000 images, 760 classes

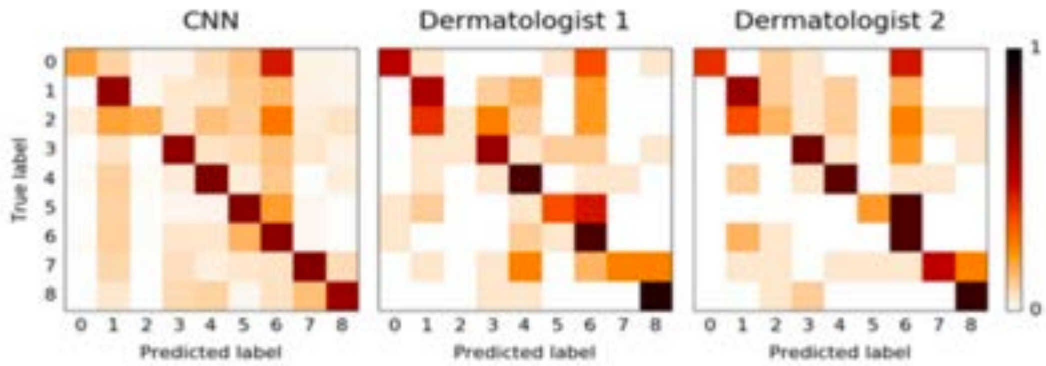

Esteva et al., Dermatologist-level classification of skin cancer with deep neural networks, Nature 542, 115–118, Feb. 2017, doi:10.1038/nature21056

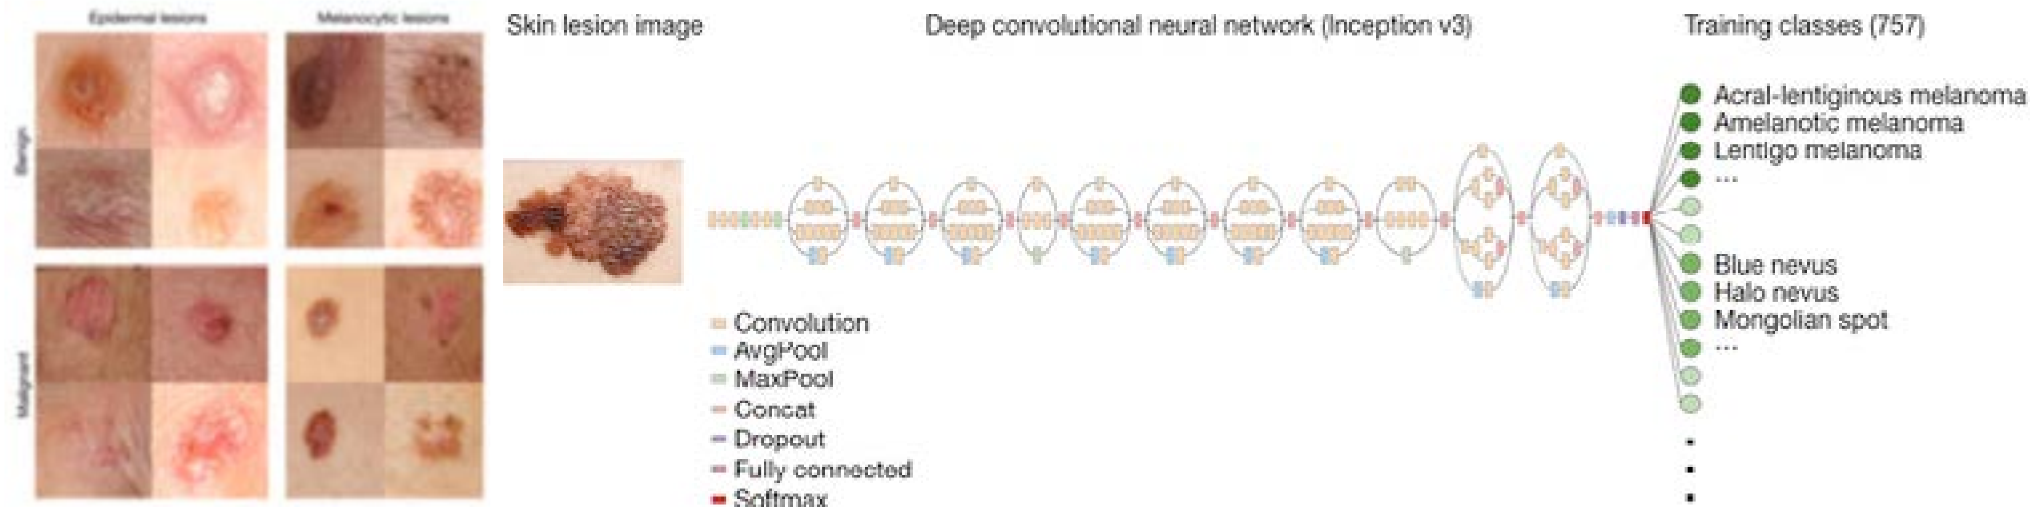

Hough-CNN: anatomy localization & segmentation

- Perform voxel-wise classification
- Each voxel classified as foreground votes for centroid of classified region

  → a peak in the vote map emerges
- All votes that voted within a radius R of the peak are considered "correct"
- These patches can now re-project the segmentation patches of kNN patches from the dictionary
- → segmentation volume
- → thresholding yields final contour

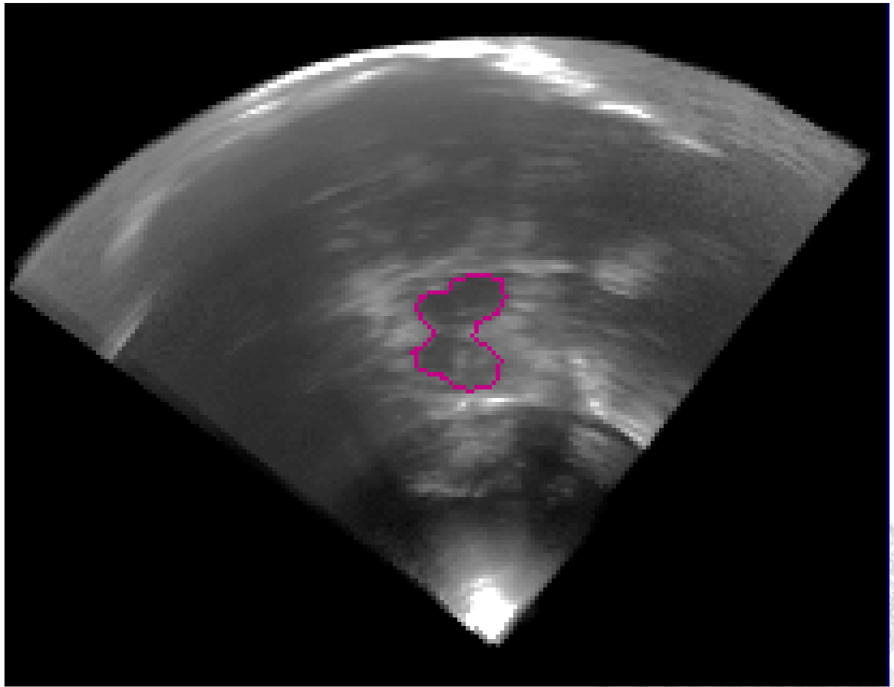

Milletari, Ahmadi, et al., Hough-CNN, https://arxiv.org/abs/1601.07014

Hough-CNN Results: Ultrasound

- Training via (5,10,15,…,45) volumes, testing on 114 volumes
- Successful segmentation close to human inter-rater performance in 112 volumes

Segmentation via voxel-wise classification of volumetric patches:

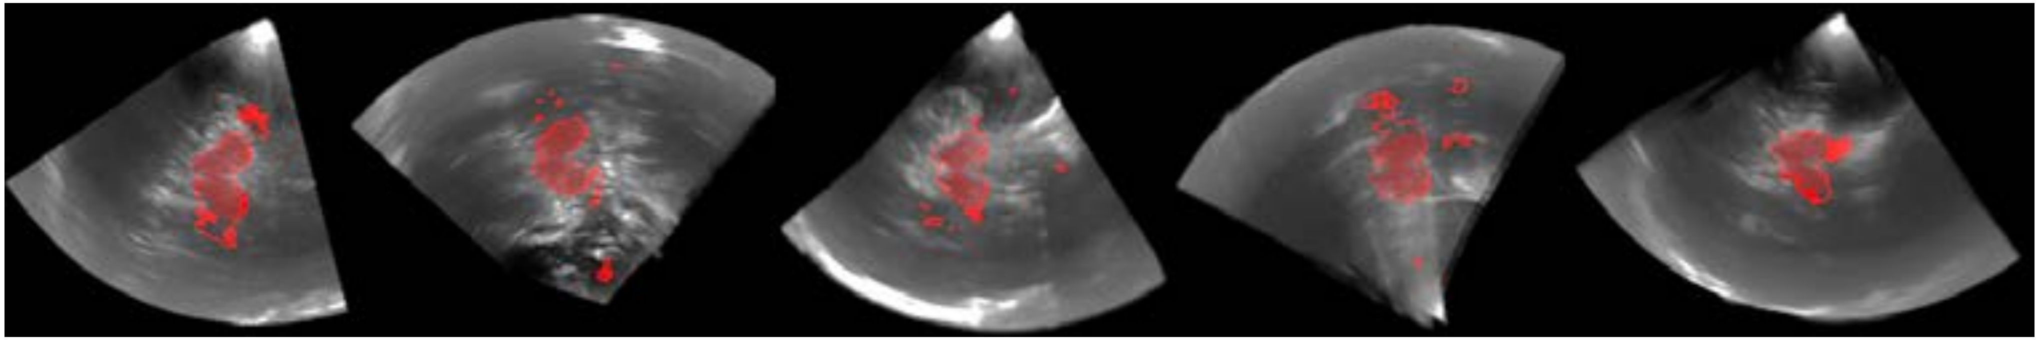

Segmentation via Hough-CNN

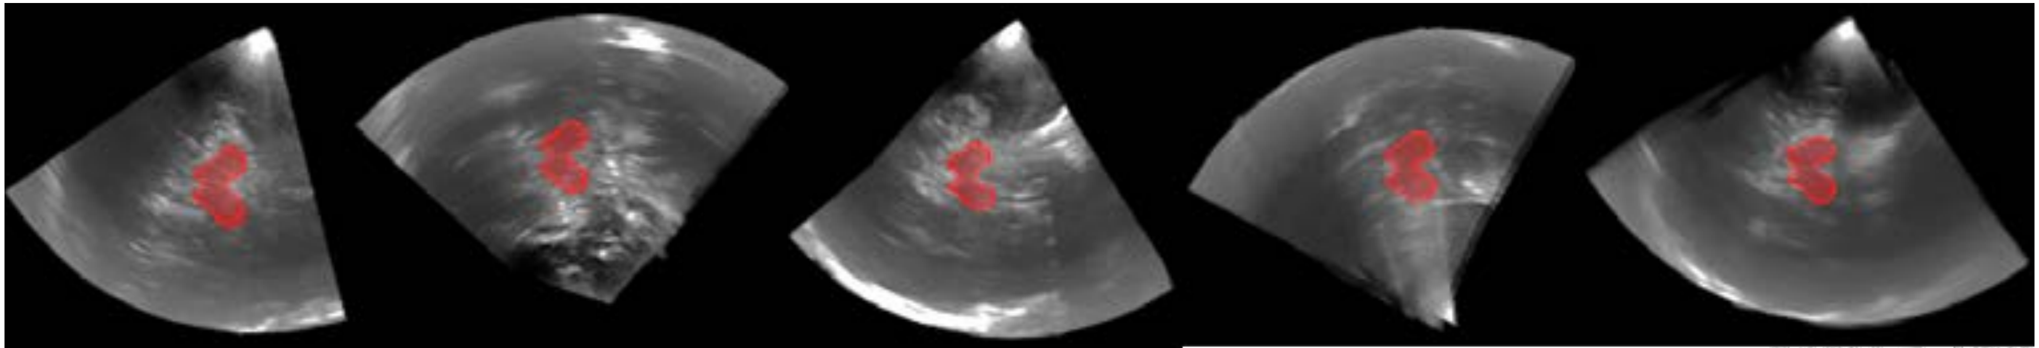

Hough-CNN Results: Structural MRI

- Voxel-wise classification with patch-wise CNN

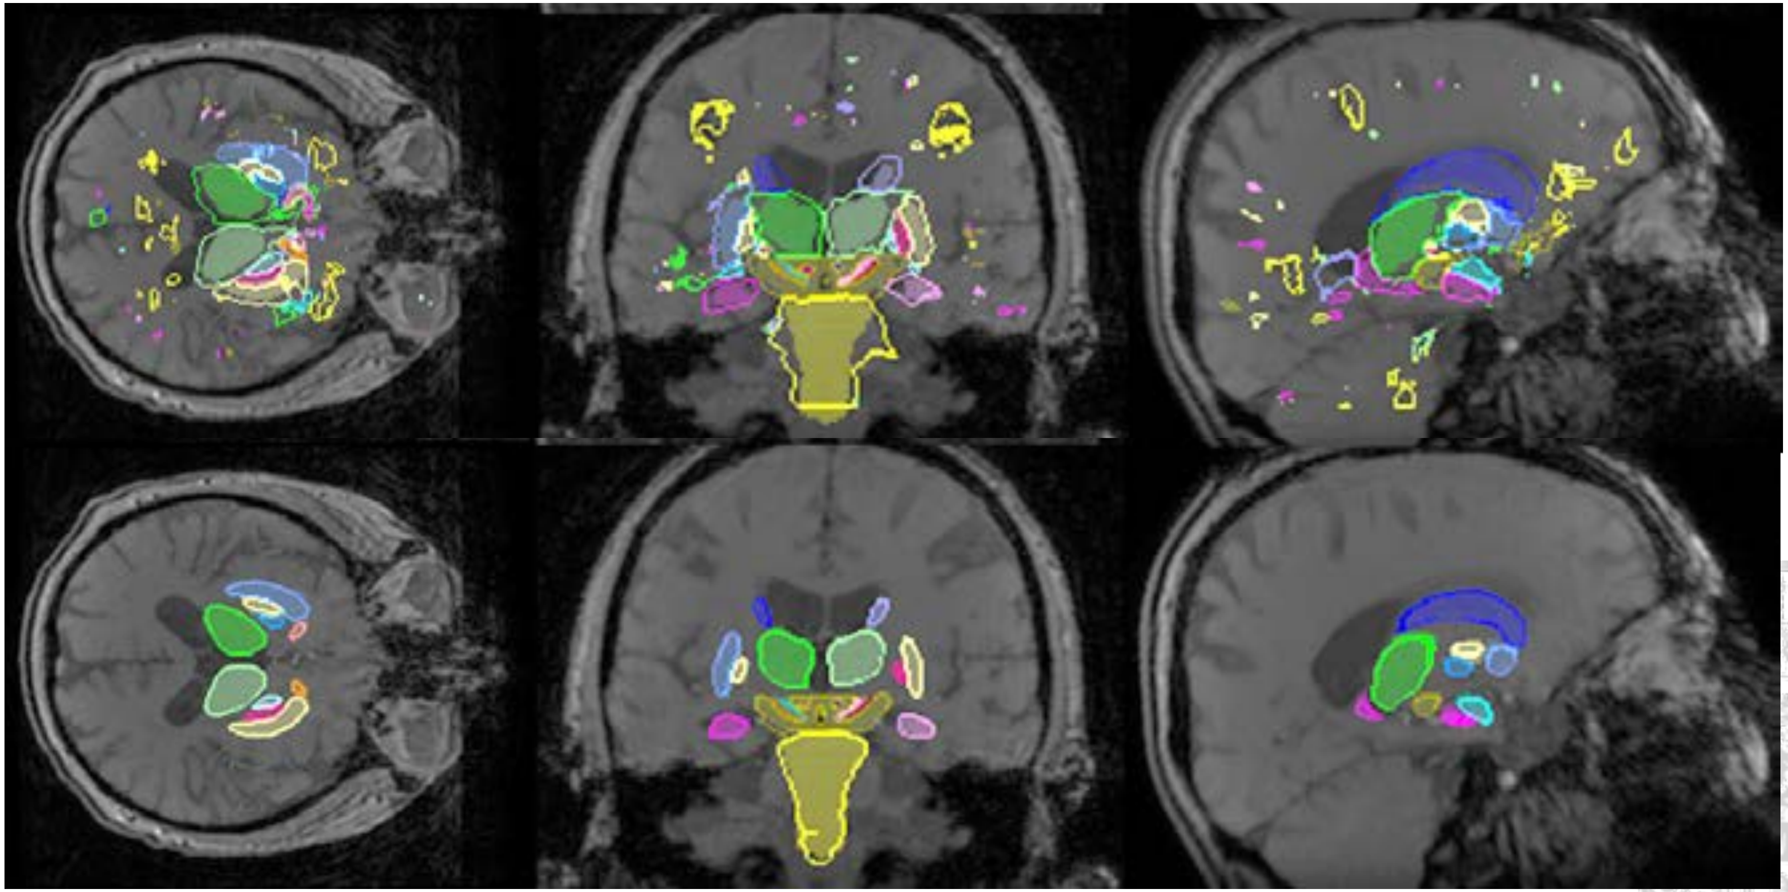

Hough-CNN Results: Structural MRI

- Voxel-wise classification with Hough-CNN

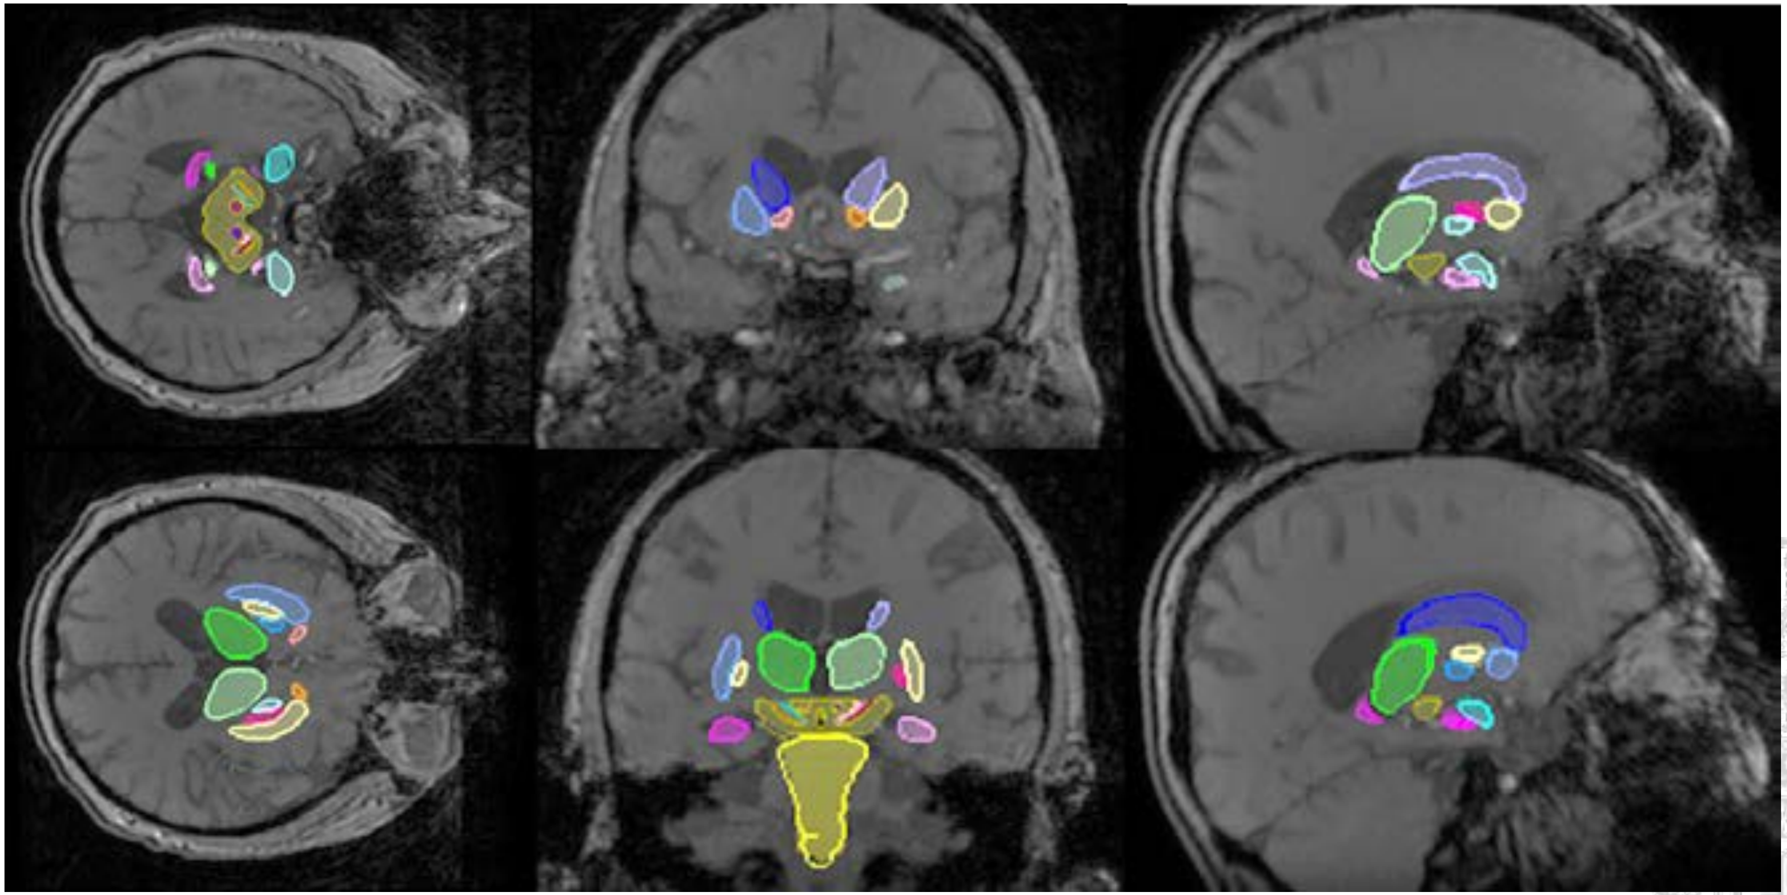

V-Net Method

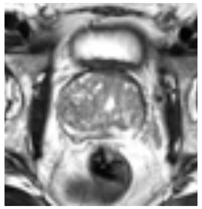

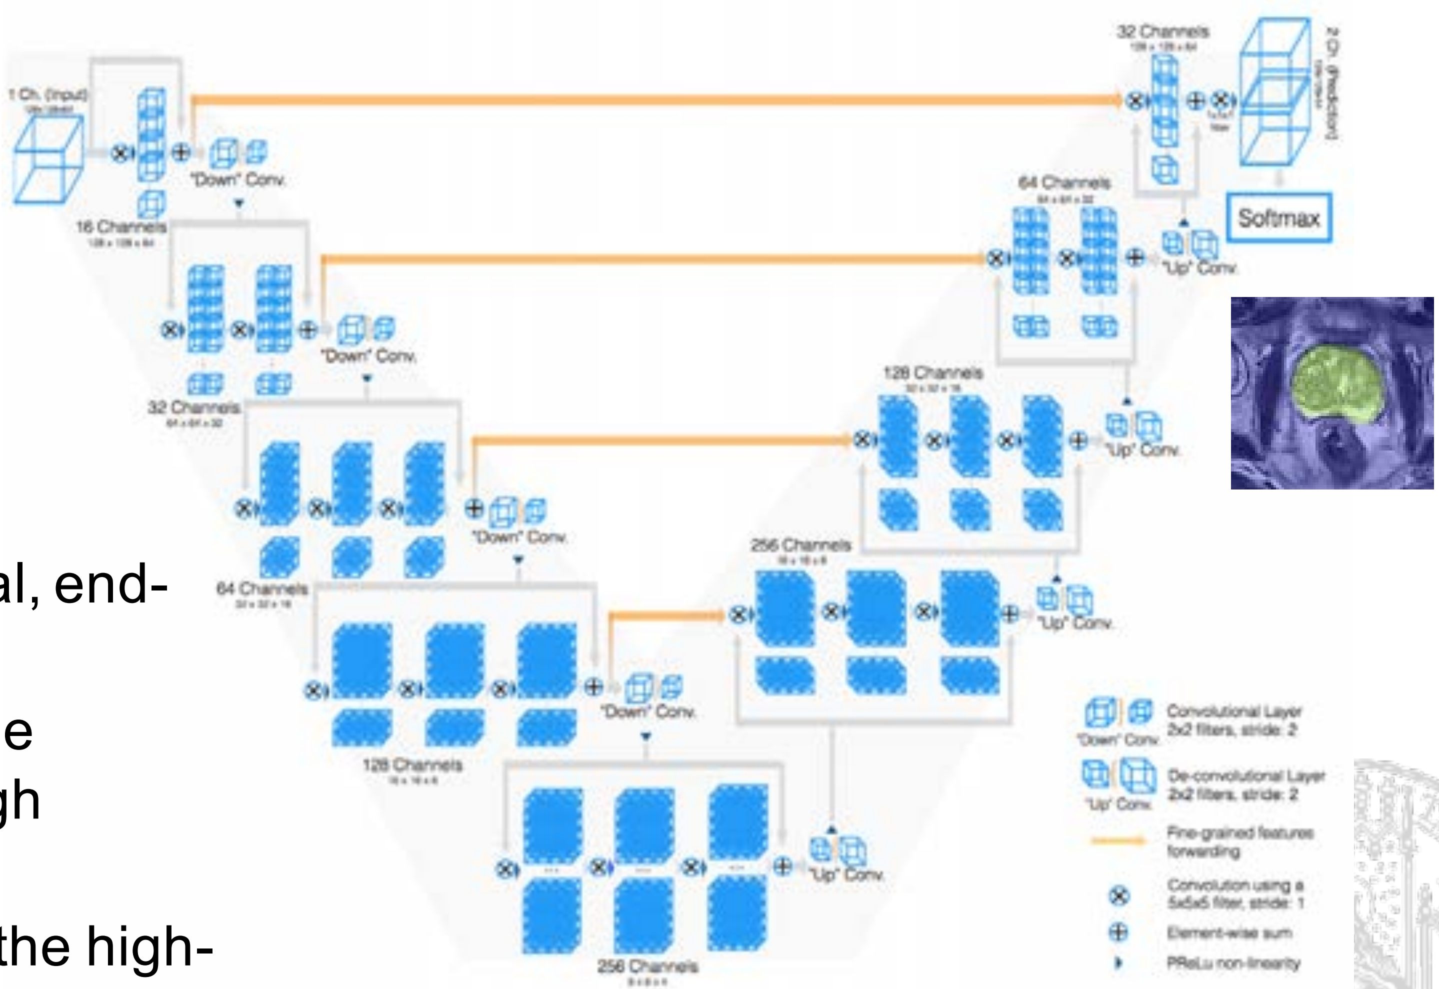

- Fully convolutional, end-to-end training
- Funneling of image information through
- Low-dimensional representation of the high-dimensional input

Milletari, Ahmadi, et al., V-Net
https://arxiv.org/abs/1606.04797

V-Net results on prostate MRI

Top 3 on PROMISE12 challenge dataset

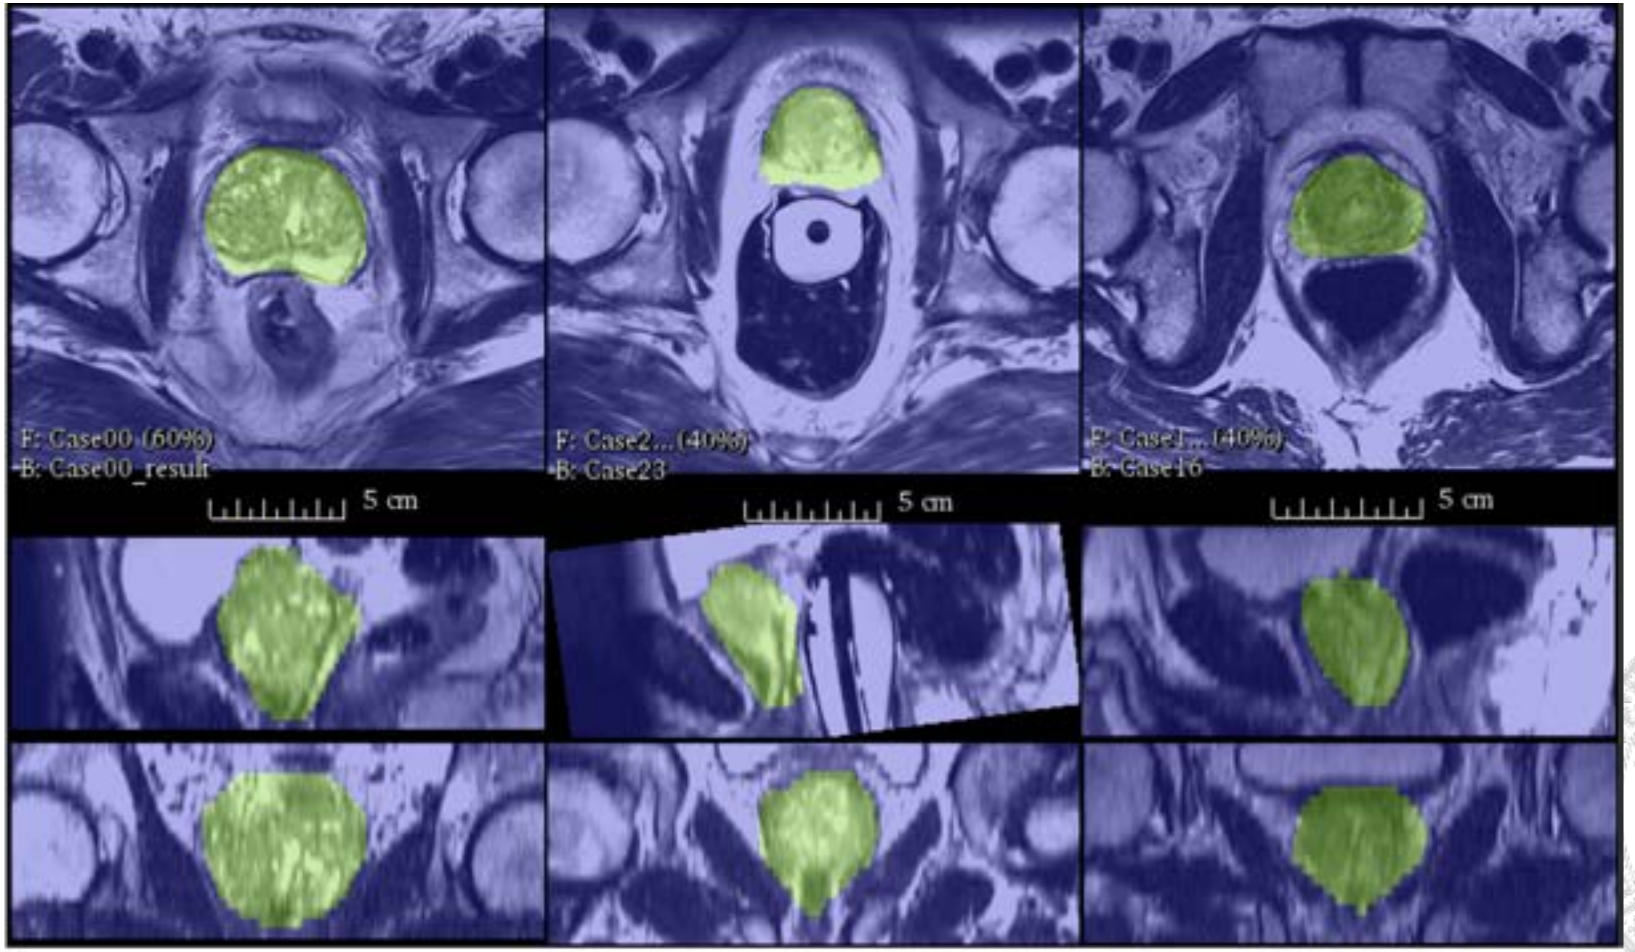

Auto-encoders: when labeled data is scarce…

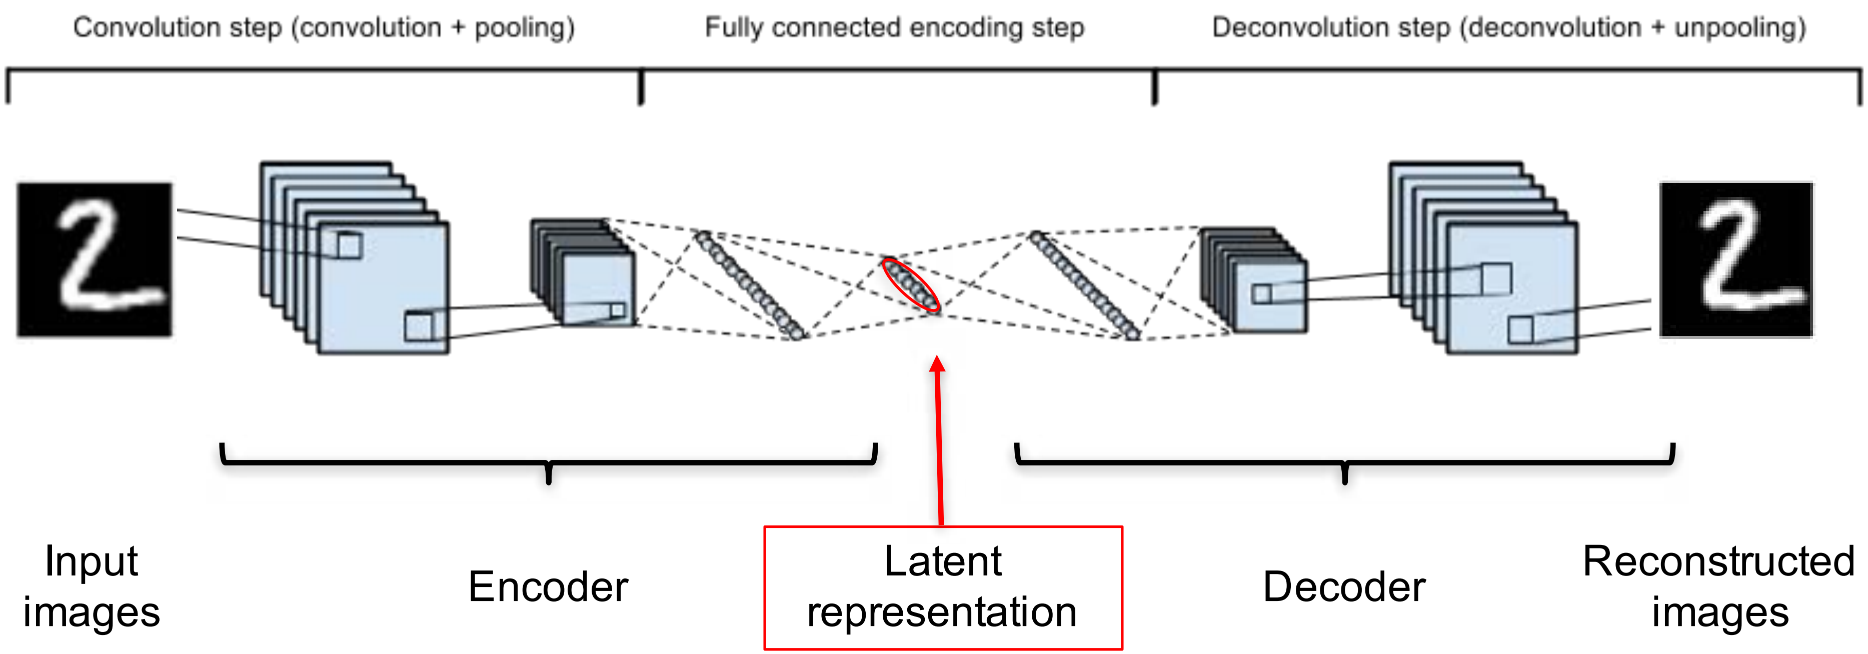

Input images:

Reconstructed images:

Denoising auto-encoders for more robust features

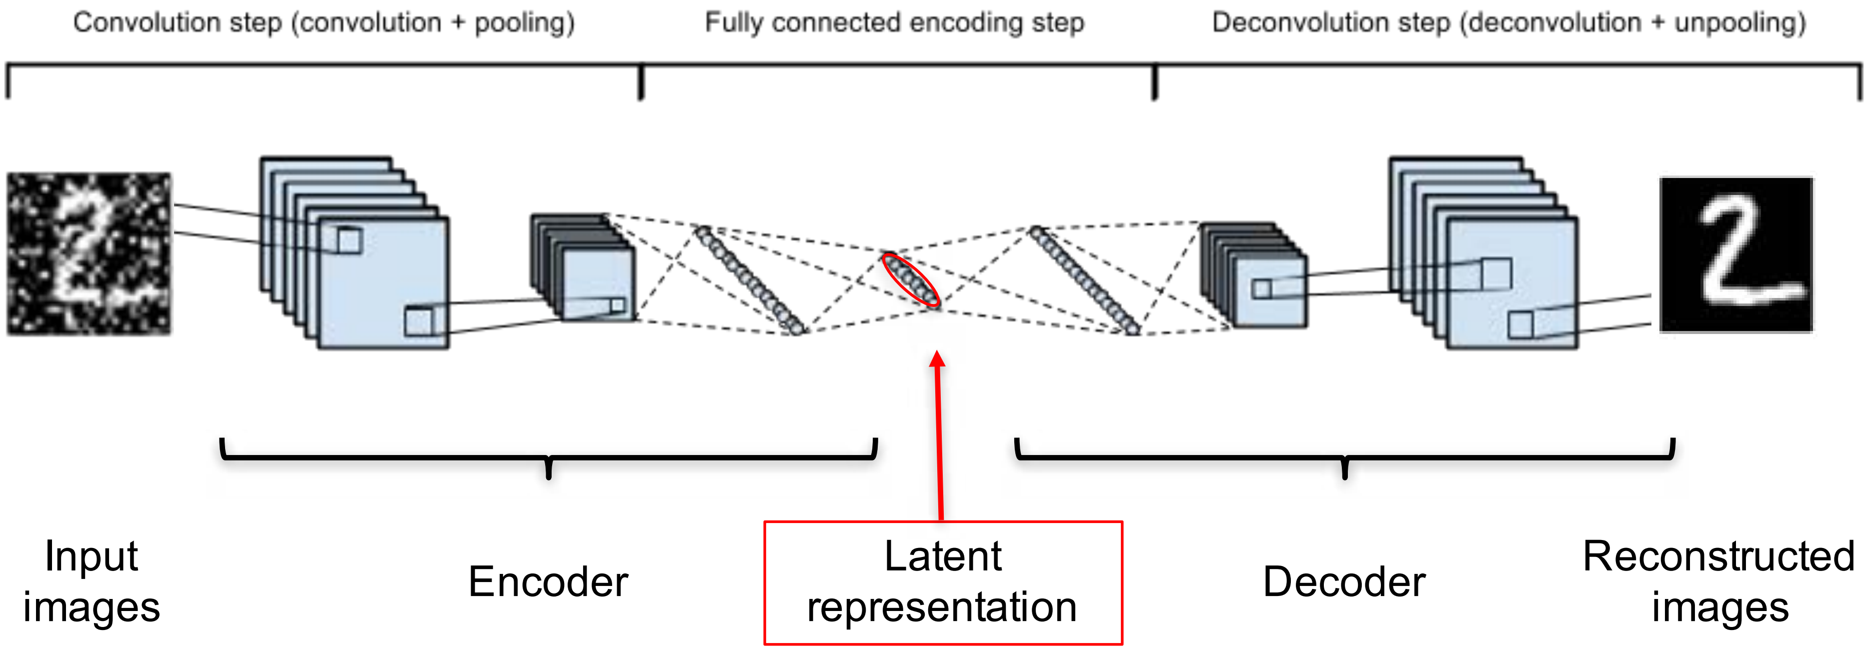

Input images:

Reconstructed images:

Conclusion

Recent (dramatic) advances in computer vision using recurrent and convolutional artificial neural networks and deep learning

Application to various medical image and sensor data:
Cutting-edge data analysis for the multi-modal vestibular system

- Structural image segmentation of vestibular system
  - Multi-modal, multi-region using shape-models, atlasses and CNNs
- Functional, computer-aided VOG
  - Novel acquisition workflows and features from dynamic analysis with RNNs
- Mining in multi-modal big data
  - Combined structural/functional assessment of the vestibular system

German Center for Vertigo and Balance Disorders

Questions?

Applied Statistics for Neuroscientists
Part IIa: Neural networks and deep learning

Dr. Seyed-Ahmad Ahmadi