

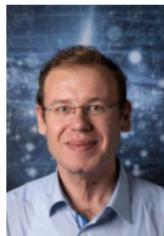
Deep Learning for Beginners - Introduction 1

A. Maier, K. Breininger, V. Christlein, S. Seitz, A. Sindel, T. Würfl

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg



Who are we? - Lab Members



Andreas
Maier



Vincent
Christlein



Aline
Sindel



Katharina
Breininger



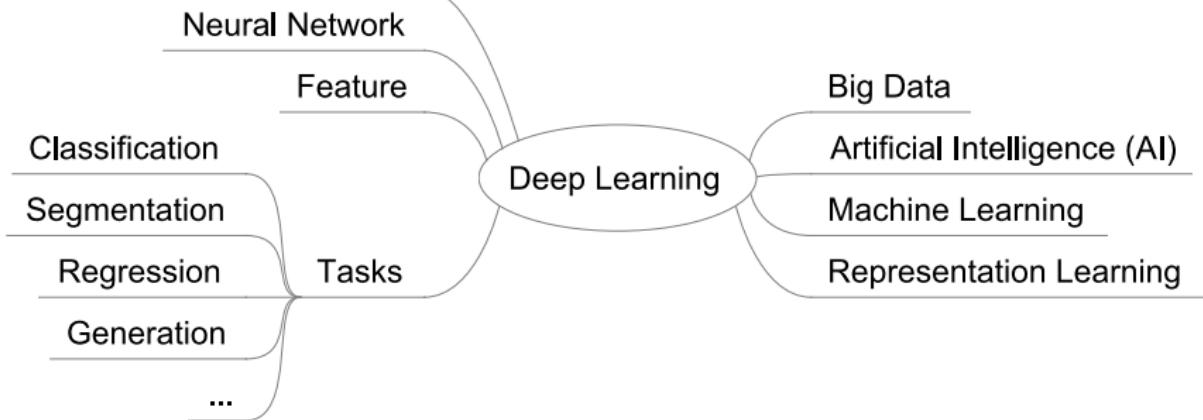
Stephan
Seitz



Tobias
Würfl

Deep Learning – Buzzwords

Supervised vs. unsupervised



Outline

Motivation

Machine Learning and Pattern Recognition

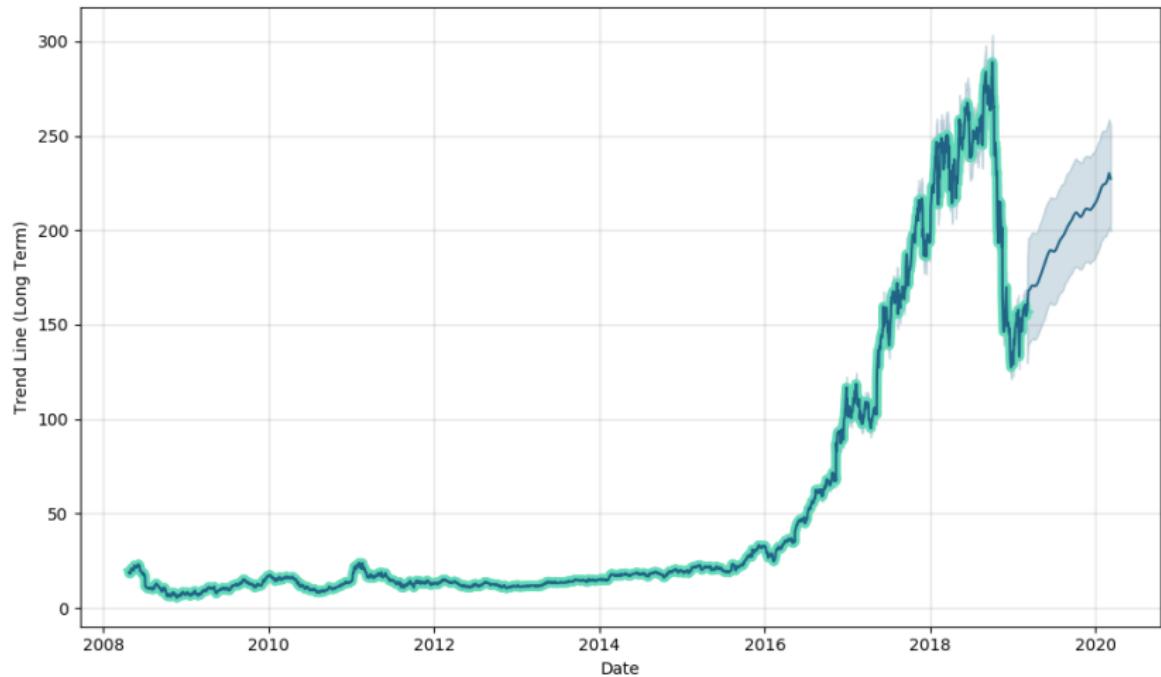
Perceptron

Organizational Matters

Motivation



NVIDIA Stock Market



Source: <https://walletinvestor.com/stock-forecast/nvda-stock-prediction/chart>

The Big Bang of Deep Learning

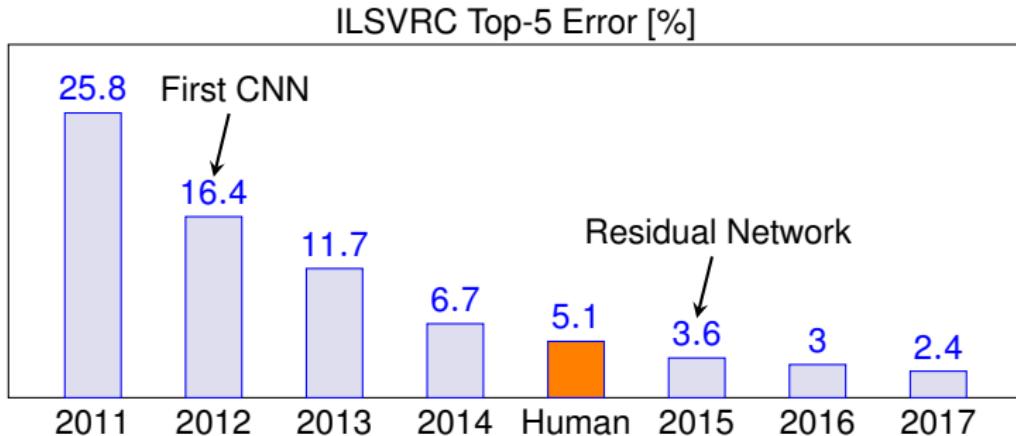


ImageNet [1] Dataset

- \approx 14 mio. images, labeled into \approx 20.000 **synonym sets**
- ImageNet Large Scale Visual Recognition Challenge using \approx 1000 classes
- Images downloaded from the Internet, **single** label per image
- **2012: Breakthrough** by Krizhevsky et al. [2]

Source: Krizhevsky et al. 2012

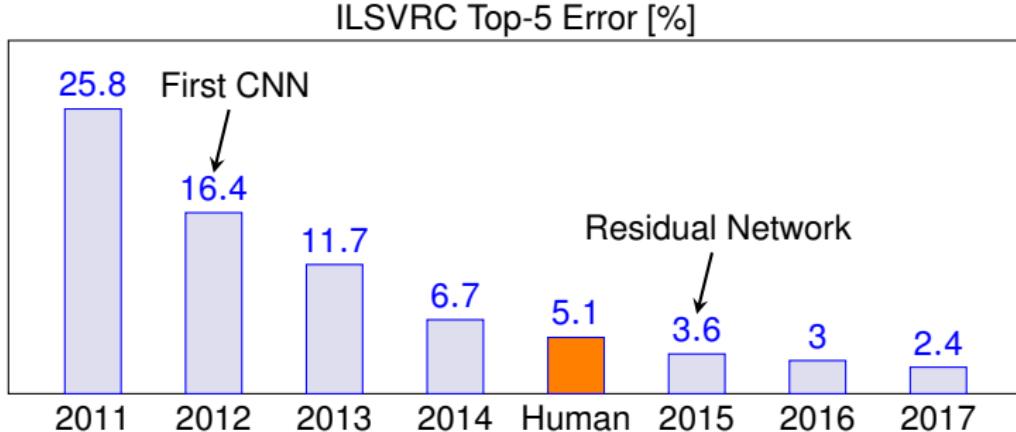
ImageNet Large Scale Visual Recognition Challenge



- First CNN approach now famous as **AlexNet** [2]

Source: image-net.org, Russakovsky et al. 2015

ImageNet Large Scale Visual Recognition Challenge

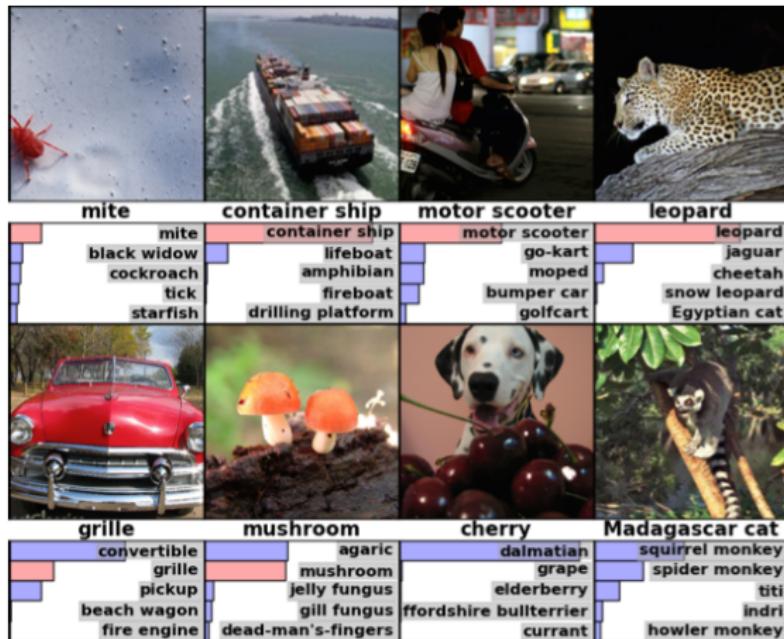


- First CNN approach now famous as **AlexNet** [2]
- “Superhuman” should be Super-Karpathy-an performance



Source: image-net.org, Russakovsky et al. 2015

ImageNet Large Scale Visual Recognition Challenge



Source: Krizhevsky et al. 2012

Deep Learning Users

NETFLIX

DAIMLER

IBM

xerox



Microsoft



 **Lunit**



SIEMENS

Google

 **DeepMind**



SAMSUNG

Playing Go

- 1997: Deep Blue beats Garry Kasparov
- Go as a next challenge
- Large branching factor



Source: <https://commons.wikimedia.org/wiki/File:FloorGoban.jpg>

Playing Go

- 1997: Deep Blue beats Garry Kasparov
- Go as a next challenge
- Large branching factor
- 2016: AlphaGo [3] beats a professional



Source: <https://commons.wikimedia.org/wiki/File:FloorGoban.jpg>

Playing Go

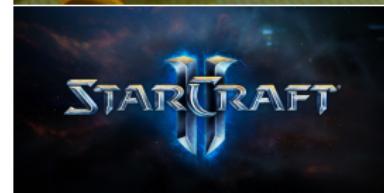
- **1997:** Deep Blue beats Garry Kasparov
- **Go** as a next challenge
- Large branching factor
- **2016:** AlphaGo [3] beats a professional
- **2017:** AlphaGoZero [4] surpasses every human in Go by self-play
- **2017:** AlphaZero [5] generalizes to a number of other board games



Source: <https://commons.wikimedia.org/wiki/File:FloorGoban.jpg>

Playing Go

- **1997:** Deep Blue beats Garry Kasparov
- **Go** as a next challenge
- Large branching factor
- **2016:** AlphaGo [3] beats a professional
- **2017:** AlphaGoZero [4] surpasses every human in Go by self-play
- **2017:** AlphaZero [5] generalizes to a number of other board games
- **2019:** AlphaStar beats professional StarCraft players



Source: <https://commons.wikimedia.org/wiki/File:FloorGoban.jpg>

Google DeepDream

Attempt to understand the inner workings of the network: What it "dreams" about when presented with images

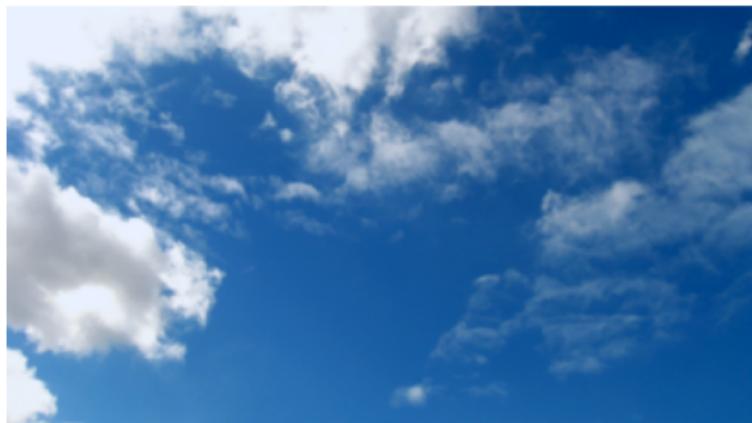
Idea:

- Arbitrary image or noise as input
- Instead of adjusting network parameters, tweak image towards high activations
- Different layers enhance different features (low or high level)



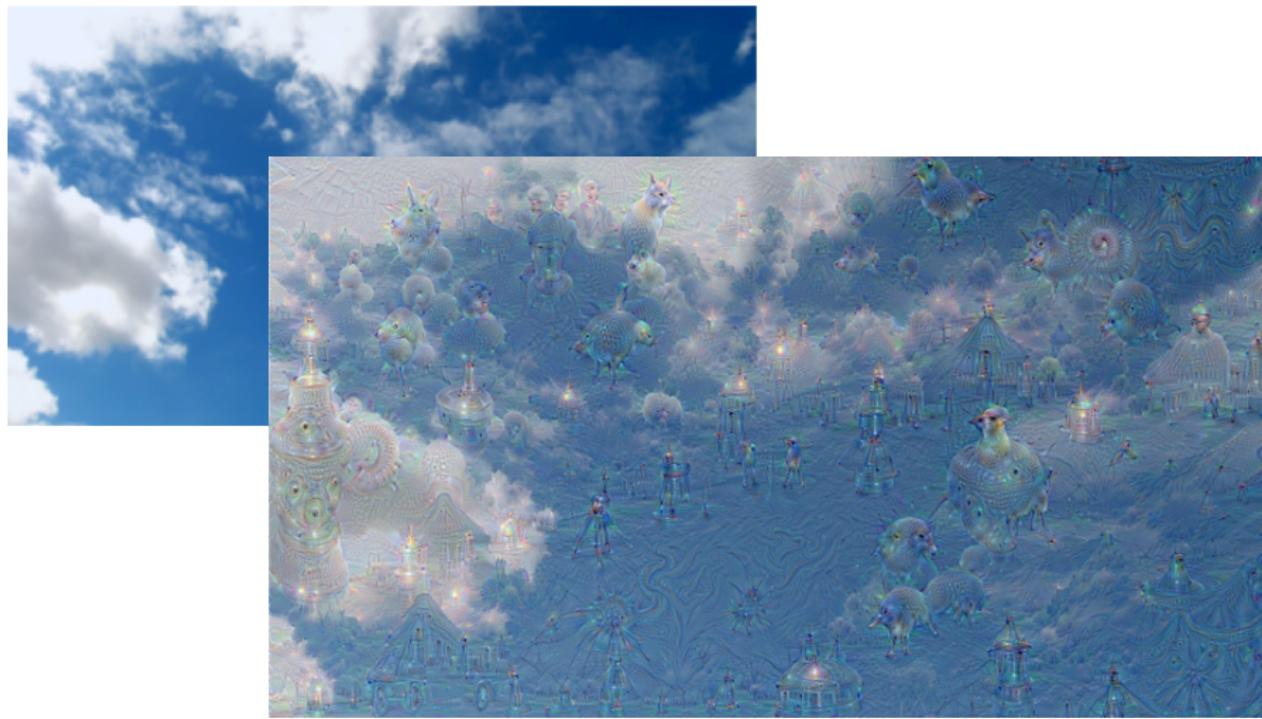
Source: <https://research.googleblog.com>

Google DeepDream



Source: <https://research.googleblog.com>

Google DeepDream



Source: <https://research.googleblog.com>

Google DeepDream

Looking for new animals in the clouds



"Admiral Dog!"



"The Pig-Snail"



"The Camel-Bird"



"The Dog-Fish"

Source: <https://research.googleblog.com>

Real-Time Object Detection: YOLO, YOLO9000, YOLOv3 [6]–[8]



Click for video

- YOLO: You only look once
- Prior systems → Use classifiers at multiple locations and scales
- YOLO → Simultaneous regression of bounding box and label
- FAST: 40-90 frames/second on a NVIDIA Titan X

Source: [www.youtube.com](https://www.youtube.com/watch?v=9qDzJLWVQHk), Redmon and Farhadi 2016

Every Day Use



Siri

Siri: Speech Interpretation and Recognition Interface



"Hey Siri, call Mom"

You can activate Siri and make your request all at once
— without using the Home button.*

Source: www.apple.com/ios/siri/

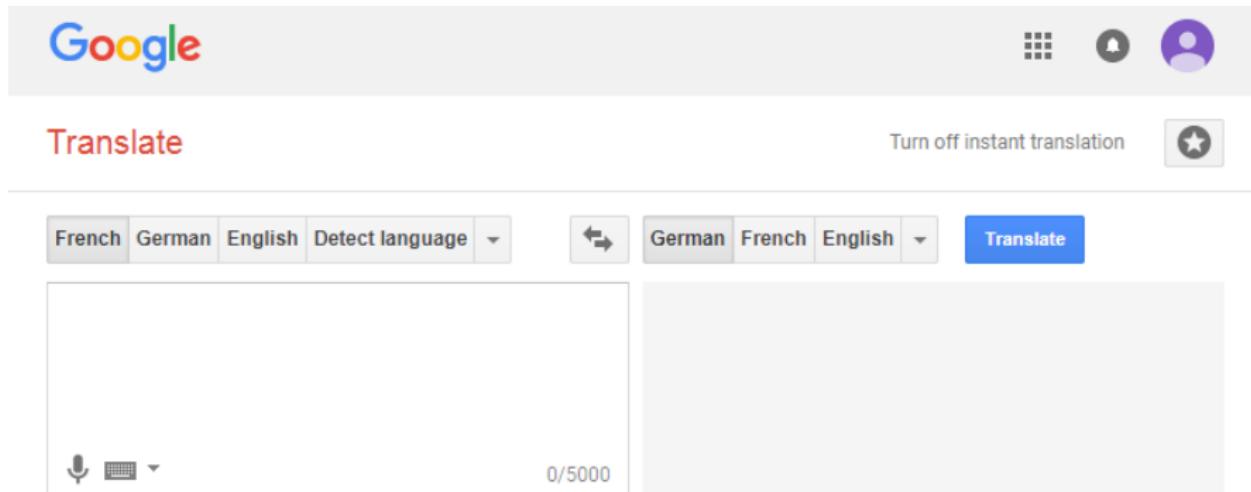
Google Echo & Amazon Alexa Voice Service

W H A T I S
ECHO DOT?



Source: www.amazon.com

Google Translate



The screenshot shows the Google Translate homepage. At the top, there's a navigation bar with the Google logo, a grid icon, a bell icon, and a user profile icon. Below the bar, the word "Translate" is displayed in red, along with a "Turn off instant translation" link and a star icon.

The main interface features two language selection dropdowns. The left one has options: French, German, English, Detect language, and a dropdown arrow. The right one has options: German, French, English, and a dropdown arrow. Between them is a double-headed arrow icon. To the right of the second dropdown is a blue "Translate" button.

Below the dropdowns are two large input fields. The left field is empty and contains icons for microphone and keyboard, with a character count of 0/5000. The right field is also empty.

Type text or a website address or translate a document.

Source: translate.google.de

**NEXT TIME
ON DEEP LEARNING**

Deep Learning for Beginners - Introduction 2

A. Maier, K. Breininger, V. Christlein, S. Seitz, A. Sindel, T. Würfl

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg



Research at the Pattern Recognition Lab

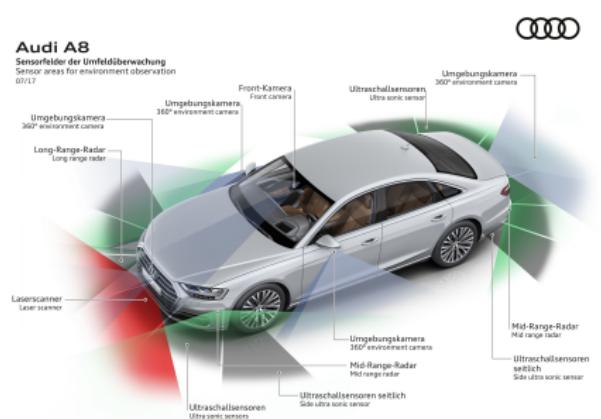


Assisted and Automated Driving

Goal

Find new ways to train and update deep learning mechanisms in environments with high safety requirements

- Assisted and automatic driving relies on sensor data
- Cameras to detect dynamic objects, driving lanes and free space
- Detection and segmentation tasks → deep learning



Source: Audi AG

Assisted and Automated Driving

- Currently: neural networks trained and thoroughly tested before deployment
- Requires huge amounts of manually labeled data
- Regular test drives cannot verify system reliability in all traffic scenarios



Click for video

Assisted and Automated Driving

- Currently: neural networks trained and thoroughly tested before deployment
- Requires huge amounts of manually labeled data
- Regular test drives cannot verify system reliability in all traffic scenarios
- **Challenge:** New ways to test algorithms in simulated environments and utilize data collected in production cars equipped with appropriate hardware



Click for video

Smart Devices

Problem statement

Renewable energy power \neq energy demand

- Underproduction \rightarrow backup power plants
- Overproduction \rightarrow energy lost
- \rightarrow Real-Time-Pricing to match energy demand and supply
- Needs *smart devices* to shift workload automatically



Smart Devices

Goal

Establish energy equilibrium by predicting energy consumption

- Example: Interrupt fridge cooling cycle when price is high, start washing machine when price is low
- Dependencies between tasks, user information and action necessary (e.g., washer/dryer)
- Task: Identify time-shiftable loads and assess appropriate time frame
- Approach: Train **recurrent neural networks** to identify usage patterns and dependencies between devices

Cloud Detection for Power Forecast [9]

Goal

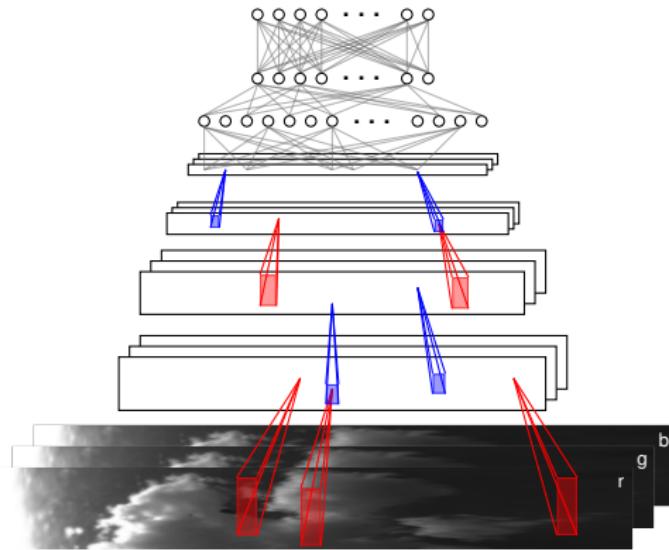
Power forecast for solar power plants with a high temporal and spatial resolution

Approach

1. Monitor the sky
2. Detect clouds
3. Estimate the cloud motion
4. Establish power forecasts



Cloud Detection for Power Forecast [9]



Input: Sky moving towards the sun

Output: Clear Sky Index = values betw. 0 (overcast sky) to 1 (clear sky)

Writer Recognition

Goal

Writer identification with limited training data (few pages per writer)

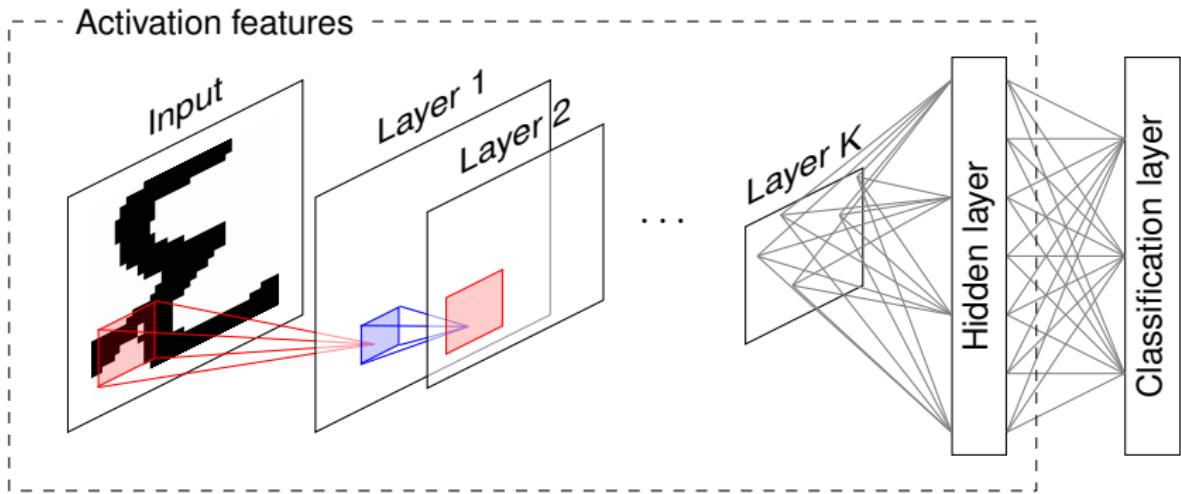
If we desire to
desire to secure
rising prosperity
for war.



Source: ICDAR'13 dataset, QUWI'15 dataset, freepik.com

Writer Recognition using CNN Activation Features [10]

Use Neuronal Network for feature extraction



Medical Applications



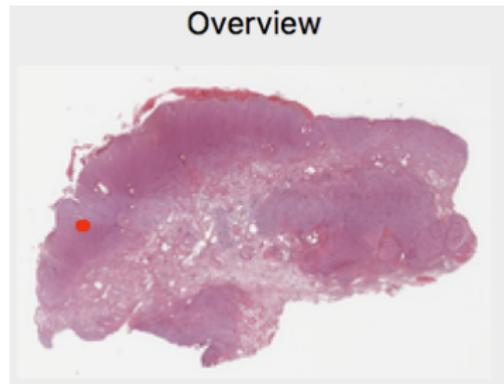
Cell Classification for Tumor Diagnostics [11]

Goal

Identify cells undergoing mitosis to asses tumor proliferation and aggressiveness in histological images

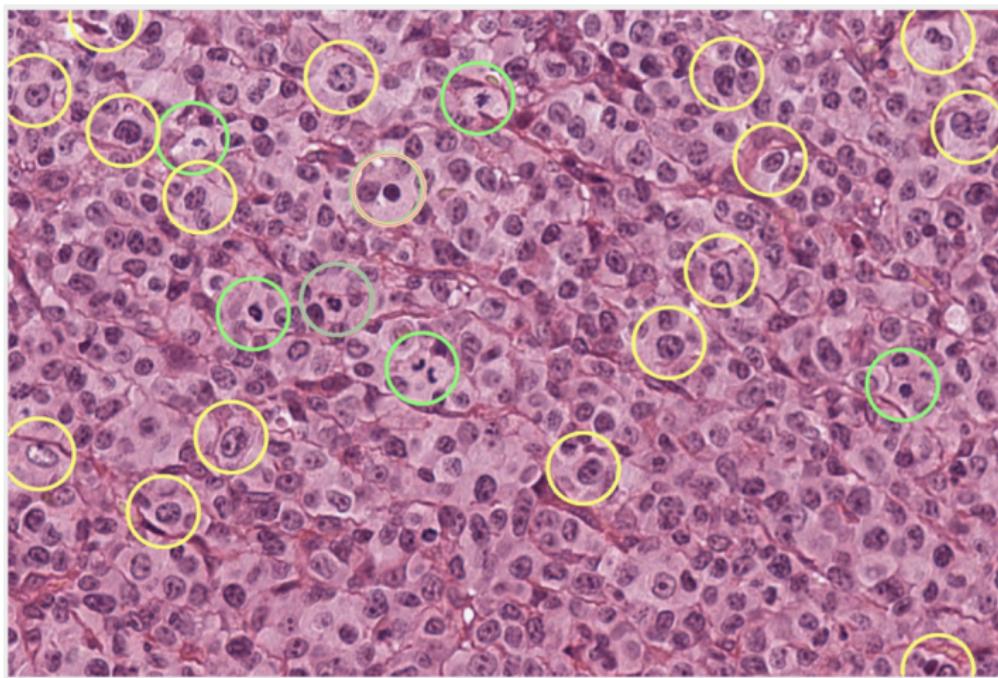
Challenge

- Histological images: large number of cells
- Full annotations not feasible
- Sparse annotations
- Cells vary significantly in size/shape/etc



Source: Aubreville et al. 2017

Cell Classification for Tumor Diagnostics [11]

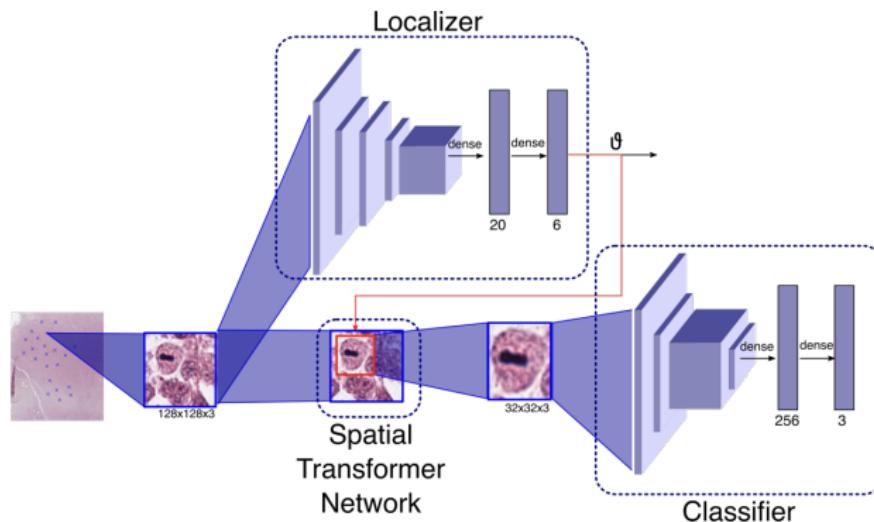


Source: Aubreville et al. 2017

Cell Classification for Tumor Diagnostics [11]

Approach

Use *spatial transformer networks* (STNs) to learn affine transformation **and** classification



Source: Aubreville et al. 2017

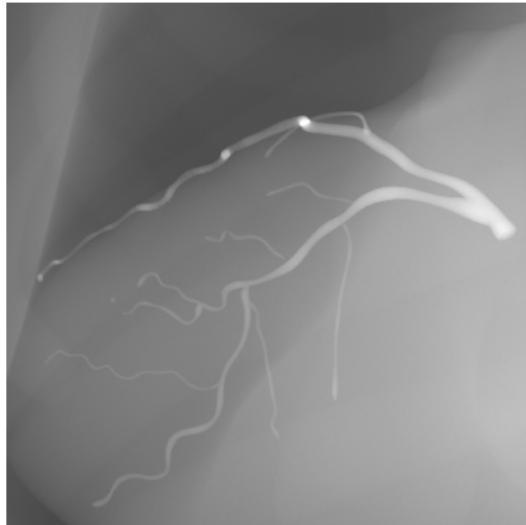
Defect Pixel Interpolation

Goal

- Reconstruction of coronaries based on truncated X-ray images
- Create “virtual” digital subtraction angiography

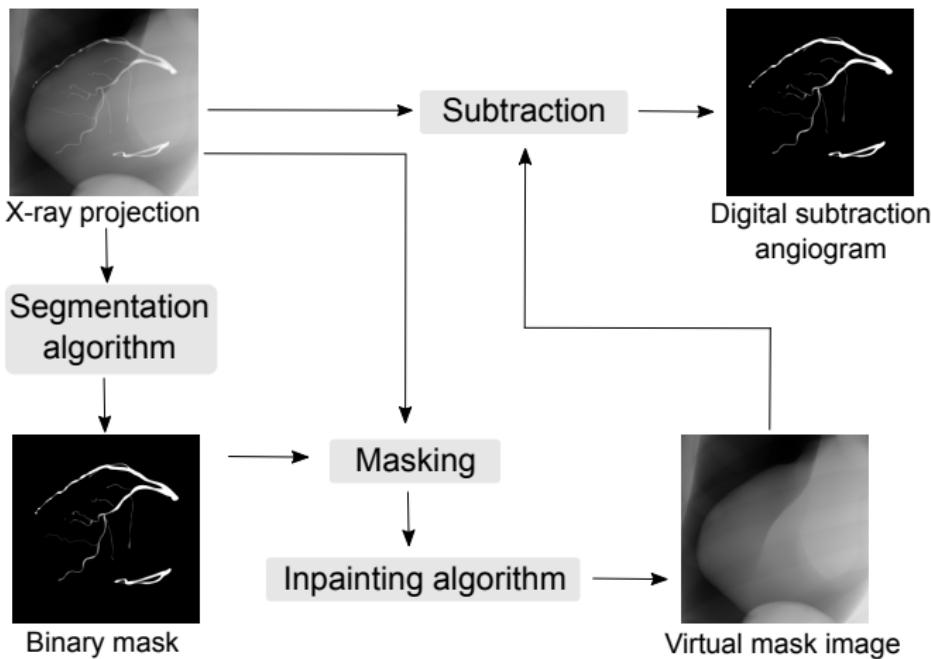
Approach

1. Segment coronary vessels
2. Mask fluoroscopic image
3. Inpaint using U-net
4. Subtract inpainted image to get untruncated data



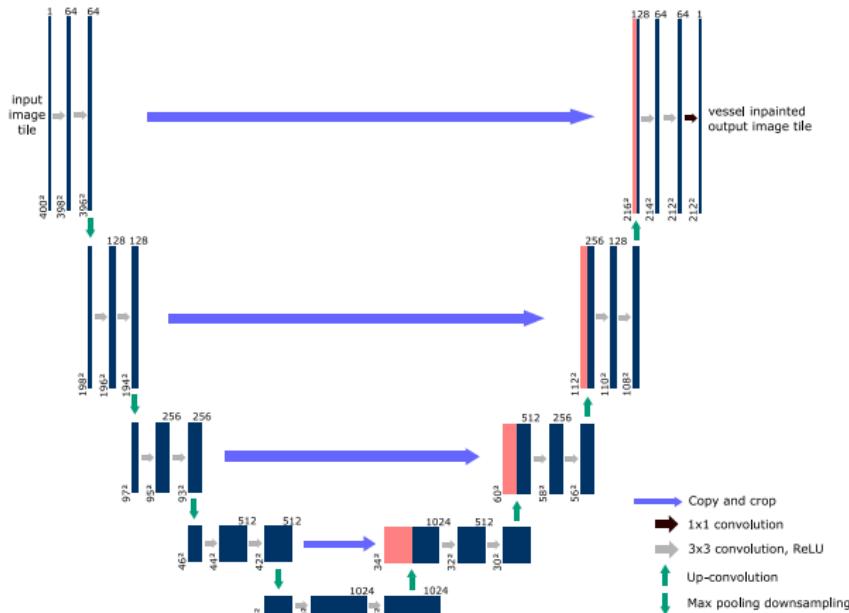
Defect Pixel Interpolation

Processing pipeline



Defect Pixel Interpolation

Deep learning for inpainting



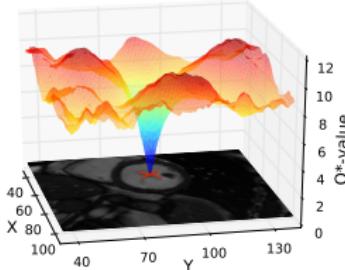
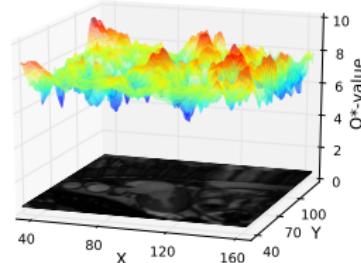
Organ Search [12]

Goal

Locate anatomic structures automatically

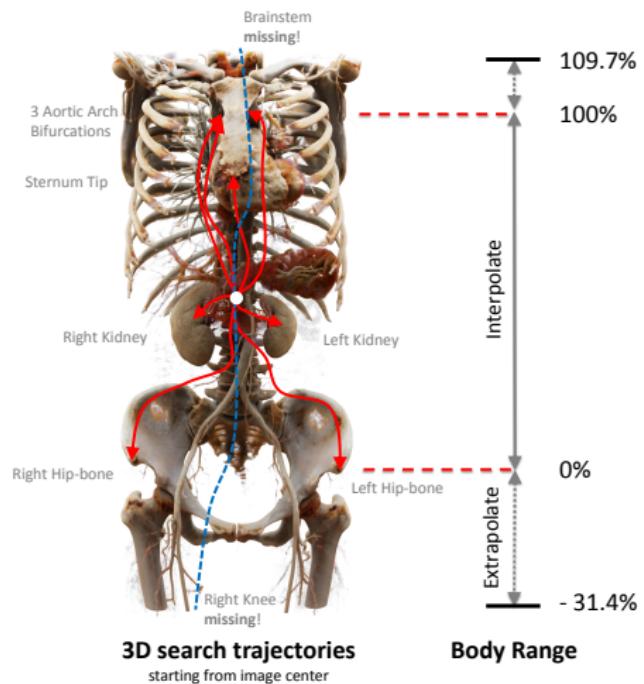
Approach

- Deep reinforcement learning
- Learn strategies how to search objects
 - Learn optimal shortest search through image volume to different landmarks
- Hierarchical approach to improve speed and robustness



Source: Ghesu et al. 2016, Ghesu et al. 2017

Organ Search [12]



Organ Search [12]



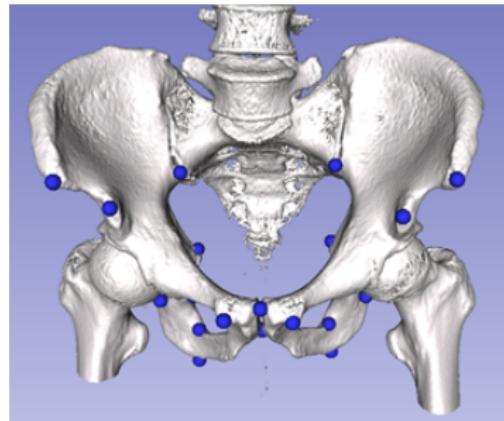
X-ray-transform Invariant Anatomical Landmark Detection

Goal

- Detect landmarks in X-ray images
- Knowing correspondences enables symbolic reconstruction
- Classic computervision reconstruction

Challenge

- Transmission imaging
- Overlap/superposition of structures
- High variance due to projection
- Artifacts e.g. interventional devices



Source: Bier et al. 2018

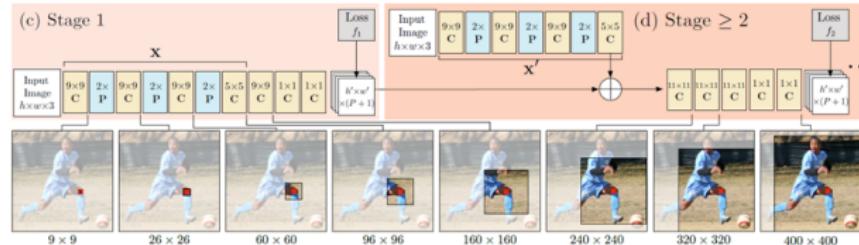
X-ray-transform Invariant Anatomical Landmark Detection

Approach: Convolutional Pose Machine (CPM) [13]

- Sequential prediction framework to detect landmarks
 - Yields 2D belief maps

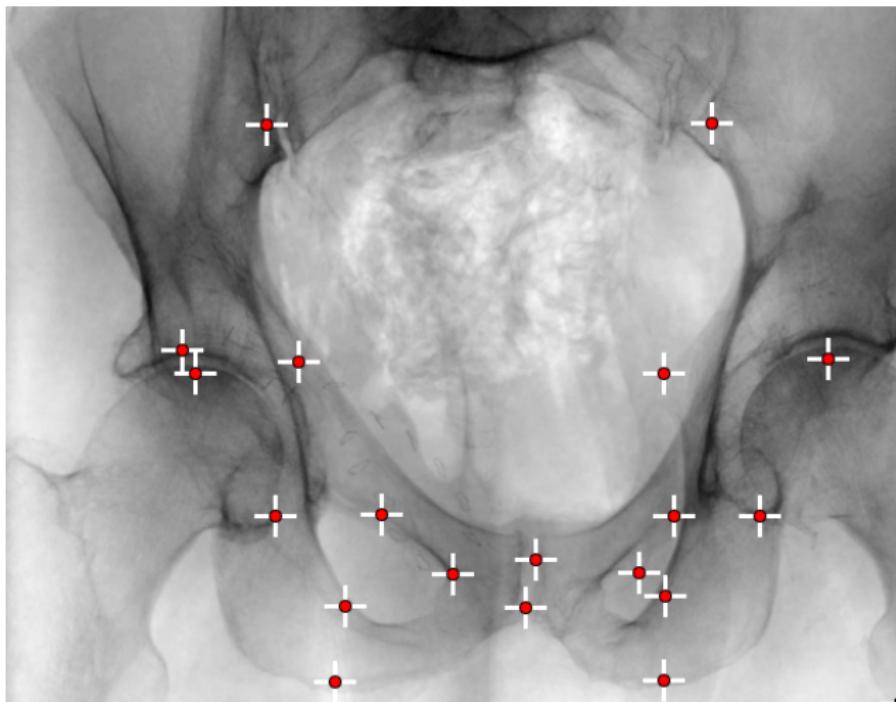
Properties

- Large receptive fields enable learning of configurations
- Estimation is refined over stages



Source: Wei et al. 2016

X-ray-transform Invariant Anatomical Landmark Detection

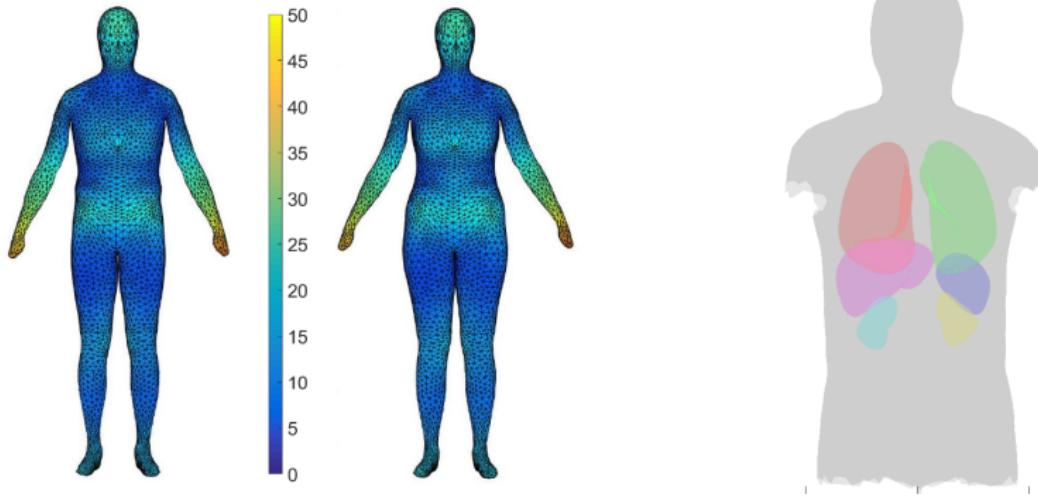


Source: Bier et al. 2018

Organ Prediction

Goal

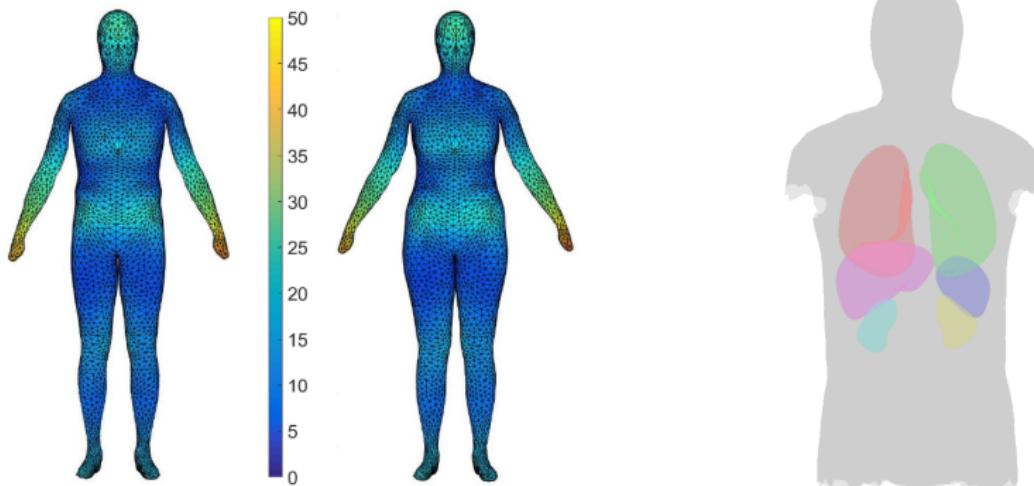
Estimation of body and organ shapes based on patient's height and weight for X-ray exposure estimation.



Organ Prediction

Goal

Estimation of body and organ shapes based on patient's height and weight for X-ray exposure estimation.



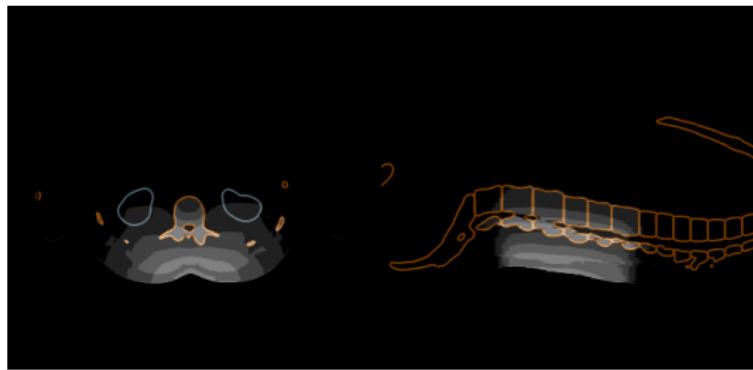
Could we achieve more if we had old CT data of a patient?

Action Learning for 3D Point Cloud Based Organ Segmentation

Goal: Versatile organ segmentation for:

- Use it in computer aided diagnosis
- Treatment planning
- Dose management

Dose estimation in interventions with overlays

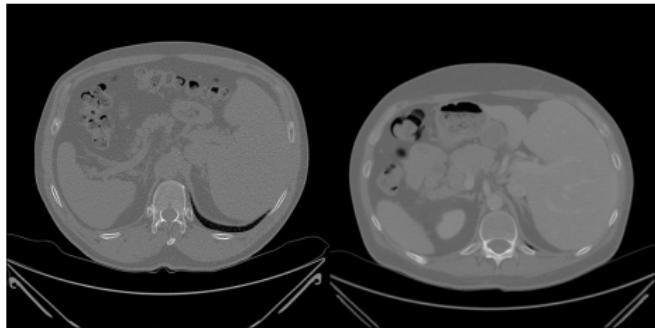


Action Learning for 3D Point Cloud Based Organ Segmentation

Challenges for clinical applications

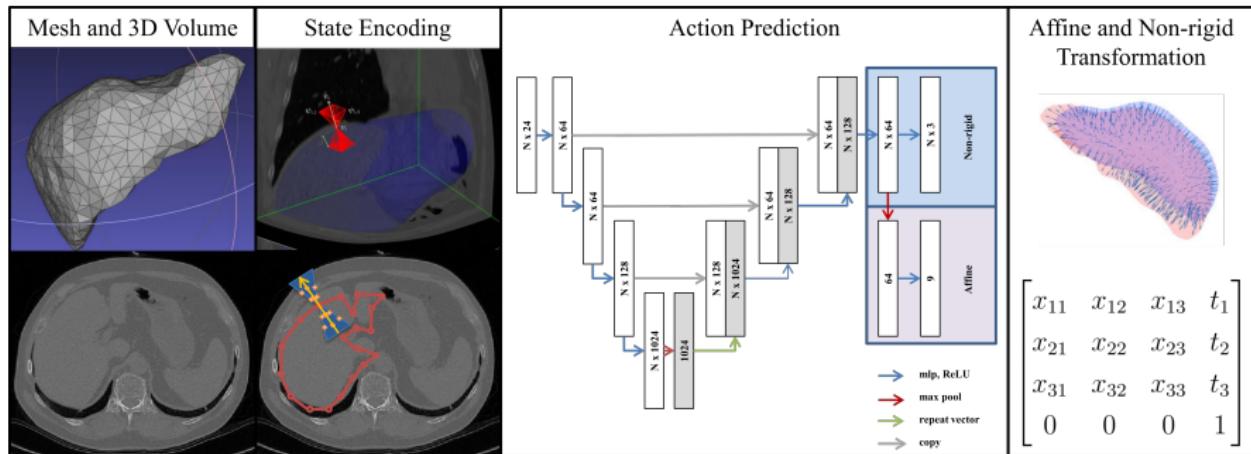
- Robustness w.r.t.
 1. Individual anatomy
 2. Scan protocols
- Time constraints

Pre-operative CT (left) and contrast enhanced CT (right)



Action Learning for 3D Point Cloud Based Organ Segmentation

- Reinforcement learning
- Predict the transformation at given state

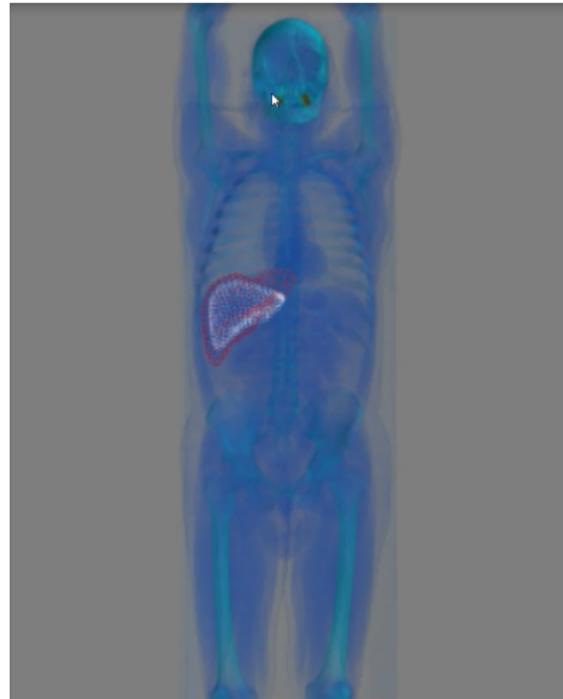


Action prediction pipeline for 3D point cloud based organ segmentation

Source: Zhong et al. 2018

Action Learning for 3D Point Cloud Based Organ Segmentation

- Runtime:
 1. **0.3 - 2.6s per volume**
 2. **50 - 100 speedup** from U-net [14]
- Very accurate
- Robust to:
 1. scan protocol
 2. contrast agent
 3. organ initialization



Source: Zhong et al. 2018

**NEXT TIME
ON DEEP LEARNING**

Deep Learning for Beginners - Introduction 3

A. Maier, K. Breininger, V. Christlein, S. Seitz, A. Sindel, T. Würfl

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg



Limitations

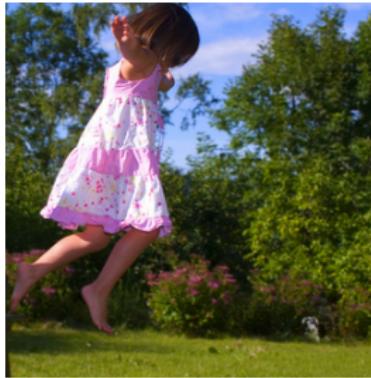


Image Captioning

Image captioning (e.g., Karpathy et al. 2014 [15]) often yields impressive results:



"baseball player is throwing ball in game."



"girl in pink dress is jumping in air."



"man in black shirt is playing guitar."

Source: <http://cs.stanford.edu/people/karpathy/deepimagesent>

Image Captioning

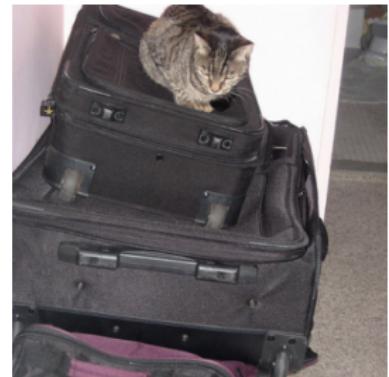
“Straightforward” errors:



"a young boy is holding a baseball bat."



"a cat is sitting on a couch with a remote control."



"black cat is sitting on top of suitcase."

Source: <http://cs.stanford.edu/people/karpathy/deepimagesent>

Image Captioning

Plainly wrong:



"a horse is standing in the middle of a road."



"a woman holding a teddy bear in front of a mirror."

Source: <http://cs.stanford.edu/people/karpathy/deepimagesent>

Challenges with Training Data

- Deep learning applications often rely on **huge**, manually-annotated data sets
- Hard to obtain, time-consuming, expensive, ambiguous
- To err is human: Mislabeled ground-truth annotation
 - May cause a significant drop in performance

Challenges with Training Data

- Deep learning applications often rely on **huge**, manually-annotated data sets
- Hard to obtain, time-consuming, expensive, ambiguous
- To err is human: Mislabeled ground-truth annotation
 - May cause a significant drop in performance
- Question: How far can we get with simulations?

Challenges with Trust and Reliability

- Verification is mandatory for high risk applications
- End-to-end learning prohibits verification of parts
- Largely unsolved

Challenges with Trust and Reliability

- Verification is mandatory for high risk applications
- End-to-end learning prohibits verification of parts
- Largely unsolved
- Possible solution: Reformulate classical algorithms

Future Directions



Learning of Algorithms

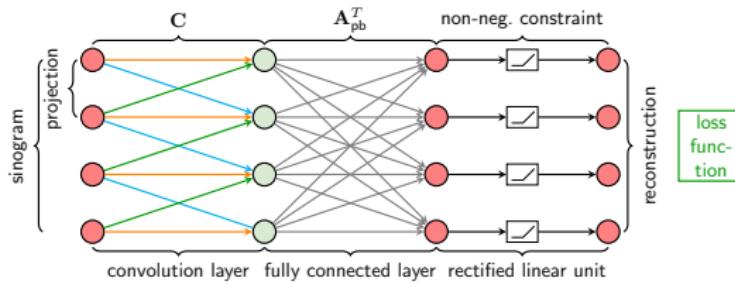
- Computed Tomography
- Efficient solution via filtered back-projection:

$$f(x, y) = \int_0^{\pi} p(s, \theta) * h(s)|_{s=x \cos \theta + y \sin \theta} d\theta$$

- Three steps:
 - Convolution along s
 - Back-projection along θ
 - Suppress negative values

Reconstruction Networks

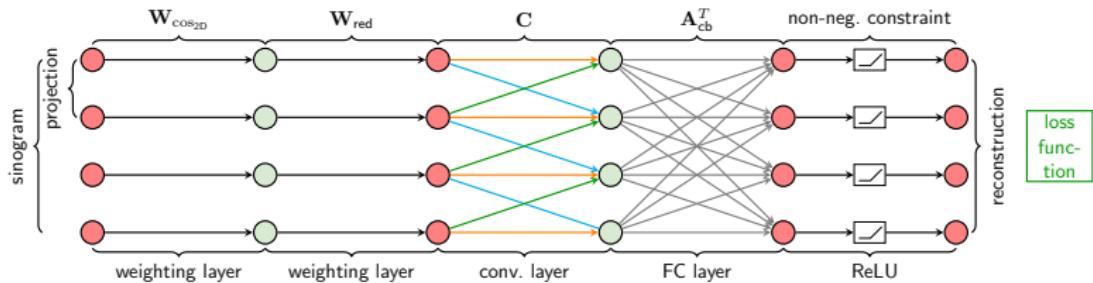
- All three steps can be modeled as a neural network:



- All weights are known from FBP

Reconstruction Networks

- Reconstruction Networks can be expanded



- Embedding of "heuristics" for artifact reduction possible

Application to Incomplete Scans [16]



Reconstruction with 360°

Application to Incomplete Scans [16]



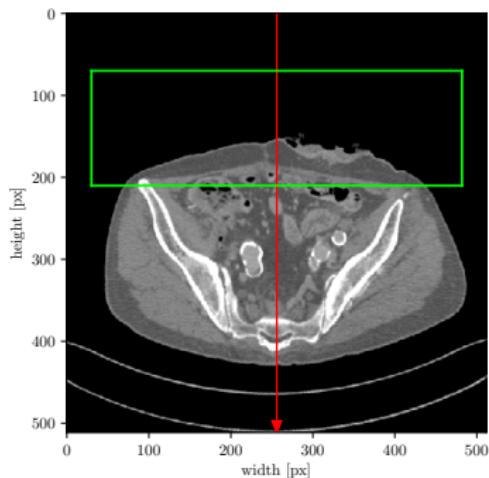
Reconstruction with 180° (FBP)

Application to Incomplete Scans [16]

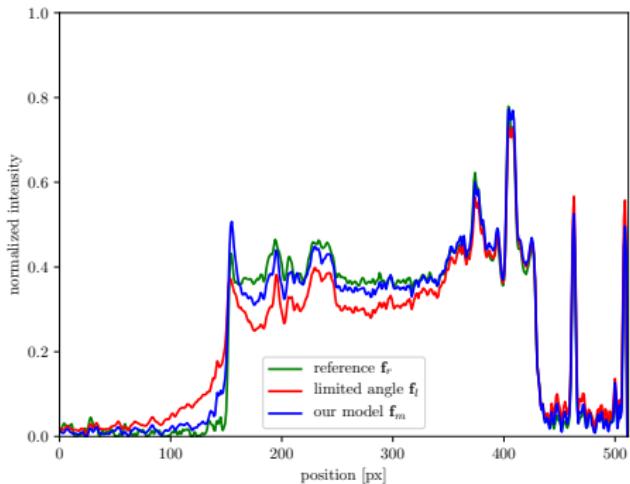


Reconstruction with 180° (NN)

Application to Incomplete Scans [16]

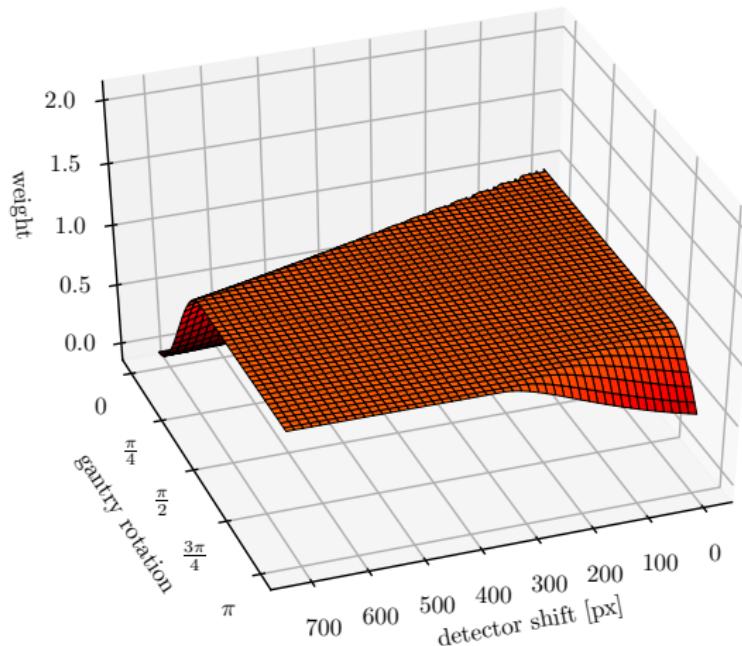


Location of the lineplot



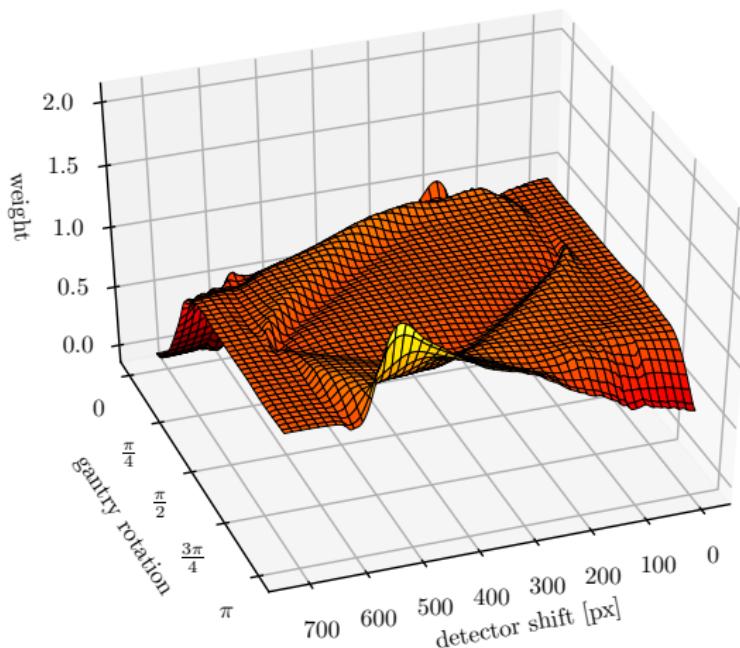
Lineplot

Parker Weights



Parker weights before learning

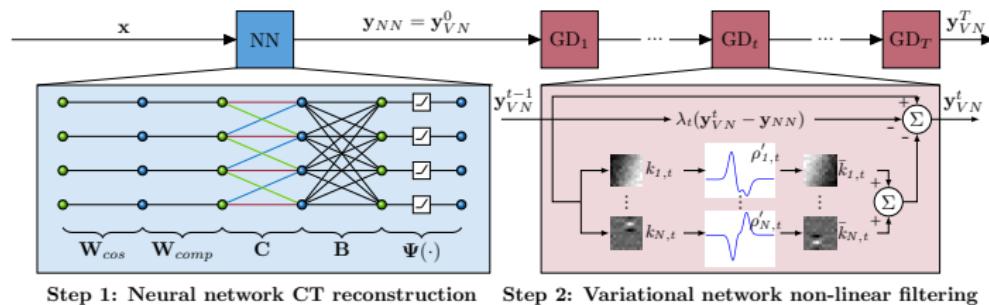
Parker Weights



Parker weights after learning

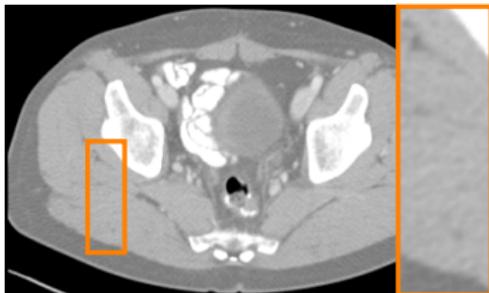
Further Extensions

- Add non-linear de-streaking and de-noising step:



Further Extensions

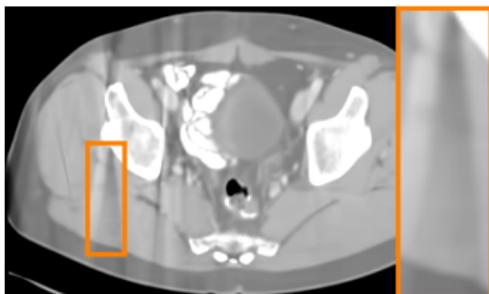
Full Scan Reference



Neural Network Input



BM3D



Variational Network ($k = 13$)



**NEXT TIME
ON DEEP LEARNING**

Deep Learning for Beginners - Introduction 4

A. Maier, K. Breininger, V. Christlein, S. Seitz, A. Sindel, T. Würfl

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg



Machine Learning and Pattern Recognition



Terminology and Notation

Throughout these slides, we will use the following notation:

- Matrices: bold, uppercase, e.g., \mathbf{M} , \mathbf{A}
- Vectors: bold, lowercase, e.g., \mathbf{v} , \mathbf{x}
- Scalars: italic, lowercase, e.g., y , w , α
- Gradient of a function: ∇ , partial derivative: ∂

Terminology and Notation

Throughout these slides, we will use the following notation:

- Matrices: bold, uppercase, e.g., \mathbf{M} , \mathbf{A}
- Vectors: bold, lowercase, e.g., \mathbf{v} , \mathbf{x}
- Scalars: italic, lowercase, e.g., y , w , α
- Gradient of a function: ∇ , partial derivative: ∂

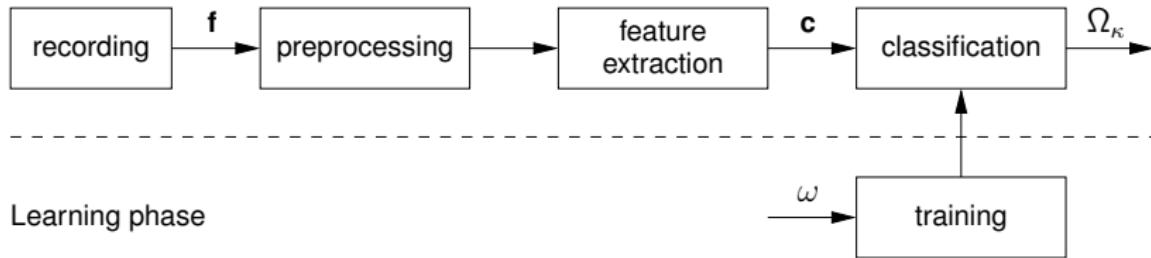
Notation regarding deep learning:

- Trainable parameters (“weights”): w
- Features/input: \mathbf{x}
- Ground truth label/target: y
- Estimated output: \hat{y}
- Index denoting iteration will be in superscript, e.g., $\mathbf{x}^{(i)}$

The notation and the terminology will be further developed throughout the lecture.

“Classical” Image Processing Pipeline

Classification phase

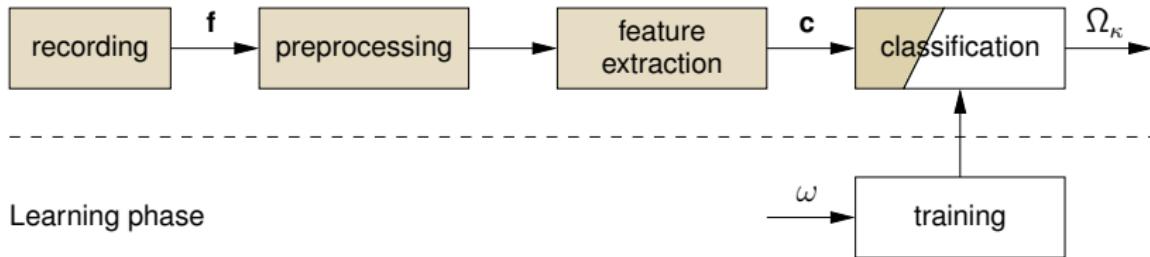


Learning phase

“Classical” Image Processing Pipeline

Lecture Introduction to Pattern Recognition

Classification phase

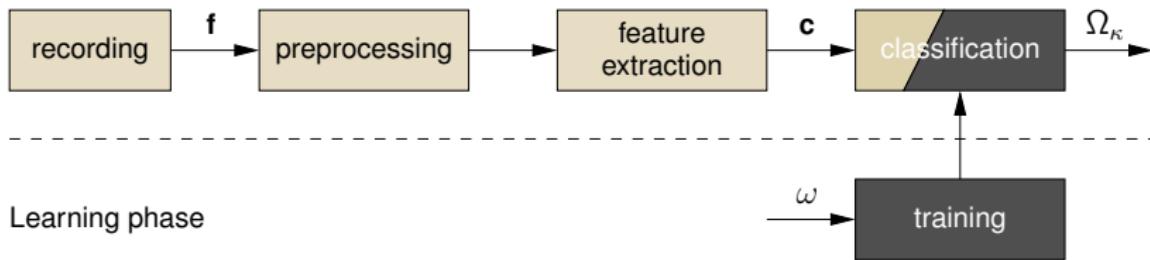


Learning phase

“Classical” Image Processing Pipeline

Lecture Introduction to Pattern Recognition

Classification phase



Learning phase

Lecture Pattern Recognition

“Classical” Image Processing Pipeline: Apple vs. Pears



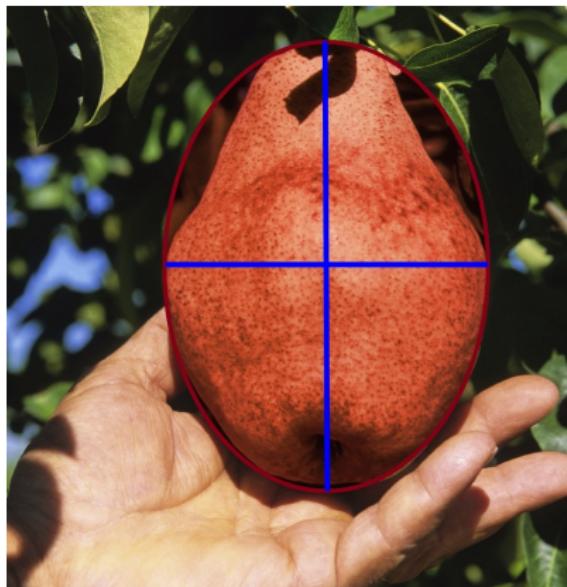
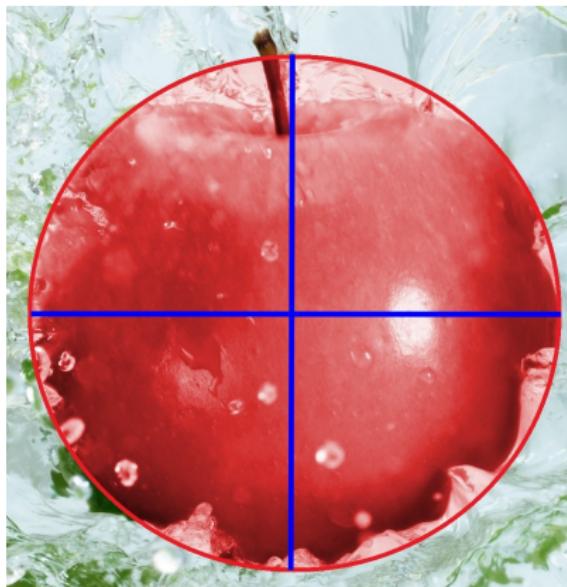
Source: <https://commons.wikimedia.org>

“Classical” Image Processing Pipeline: Apple vs. Pears



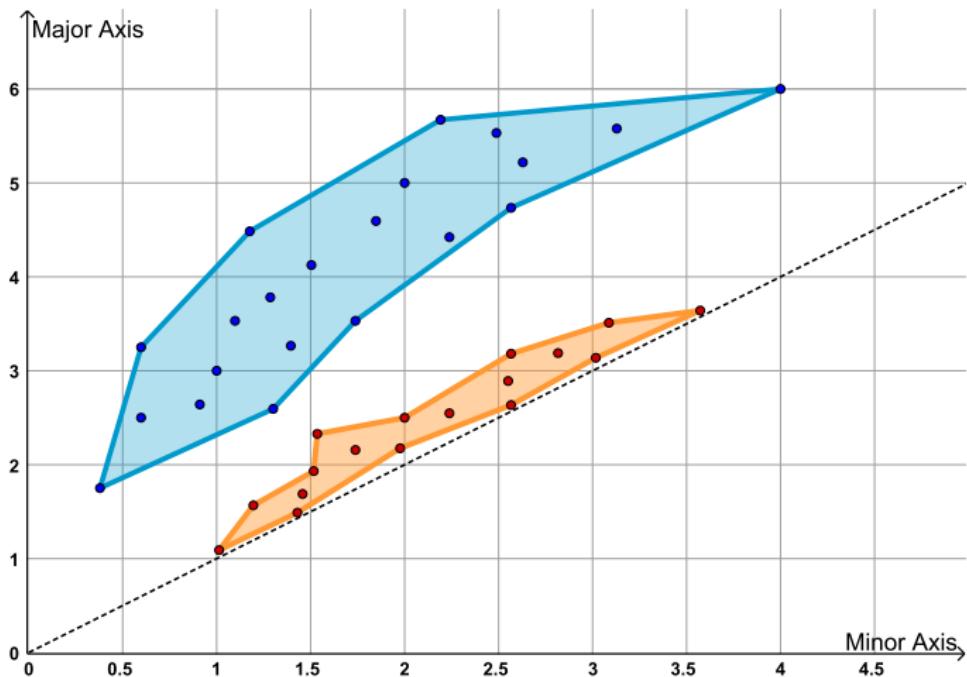
Source: <https://commons.wikimedia.org>

“Classical” Image Processing Pipeline: Apple vs. Pears

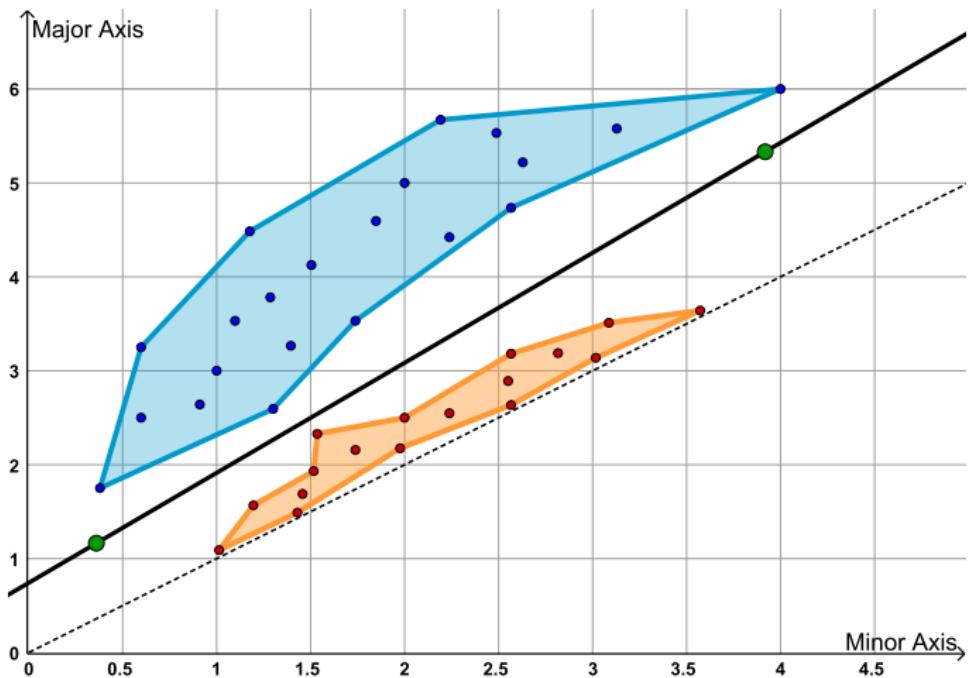


Source: <https://commons.wikimedia.org>

“Classical” Image Processing Pipeline: Apple vs. Pears



“Classical” Image Processing Pipeline: Apple vs. Pears



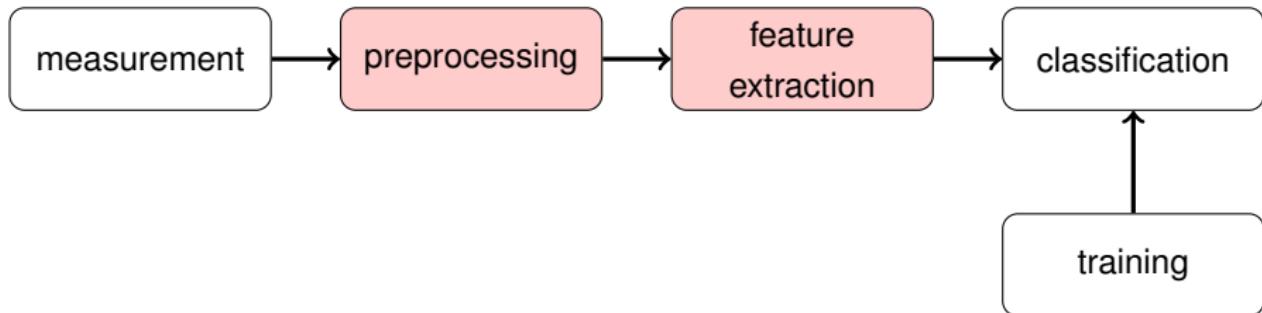
Pipeline in Deep Learning



Source: <https://xkcd.com/1838/>

Pipeline in Deep Learning

Reminder



Pipeline in Deep Learning

Now



Postulates for Pattern Recognition

6 Postulates:

1. Availability of a **representative sample** ω of **patterns** ${}^i\mathbf{f}(\mathbf{x})$ for the given field of problems Ω

$$\omega = \{{}^1\mathbf{f}(\mathbf{x}), \dots, {}^N\mathbf{f}(\mathbf{x})\} \subseteq \Omega.$$

Postulates for Pattern Recognition

6 Postulates:

1. Availability of a **representative sample** ω of **patterns** ${}^i\mathbf{f}(\mathbf{x})$ for the given field of problems Ω

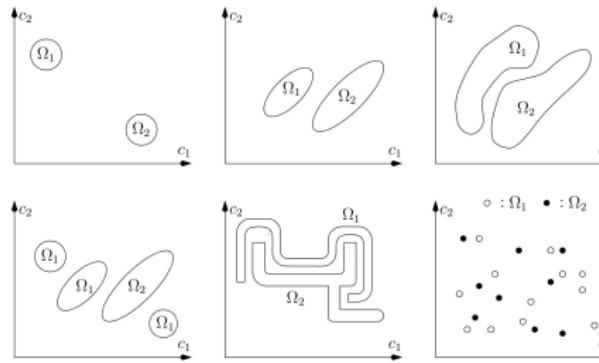
$$\omega = \{{}^1\mathbf{f}(\mathbf{x}), \dots, {}^N\mathbf{f}(\mathbf{x})\} \subseteq \Omega.$$

2. A (simple) pattern has **features**, which characterize its membership in a certain class Ω_κ .

Postulates for Pattern Recognition (cont.)

3. Compact domain of features of the same class; domains of different classes are (reasonably) separable.
 - small **intra-class distance**
 - high **inter-class distance**

Example of an increasingly less compact domain in the feature space:



Postulates for Pattern Recognition (cont.)

4. A (complex) pattern consists of **simpler constituents**, which have certain relations to each other. A pattern may be decomposed into these constituents.

Postulates for Pattern Recognition (cont.)

4. A (complex) pattern consists of **simpler constituents**, which have certain relations to each other. A pattern may be decomposed into these constituents.
5. A (complex) pattern $f(x) \in \Omega$ has a certain **structure**. Not any arrangement of simple constituents is a valid pattern. Many patterns may be represented with relatively few constituents.

Postulates for Pattern Recognition (cont.)

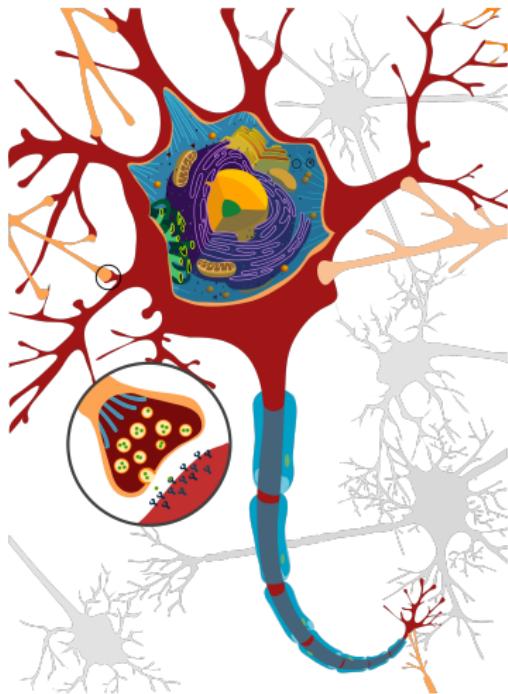
4. A (complex) pattern consists of **simpler constituents**, which have certain relations to each other. A pattern may be decomposed into these constituents.
5. A (complex) pattern $f(x) \in \Omega$ has a certain **structure**. Not any arrangement of simple constituents is a valid pattern. Many patterns may be represented with relatively few constituents.
6. Two patterns are **similar** if their features or simpler constituents differ only slightly.

Perceptron



Perceptron Biology - Neural Excitation (simplified)

- Neurons are **connected** by synapses / dendrites
- If the **sum** of incoming (excitatory and inhibitory) **activations** is large enough, an action potential is created
- The action potential activates synapses to other neurons, “transmitting” information
- All-or-none response: A **higher** stimulus does **not** cause a **higher** response
→ “binary classifier”



Source: <https://commons.wikimedia.org>

Rosenblatt's Perceptron

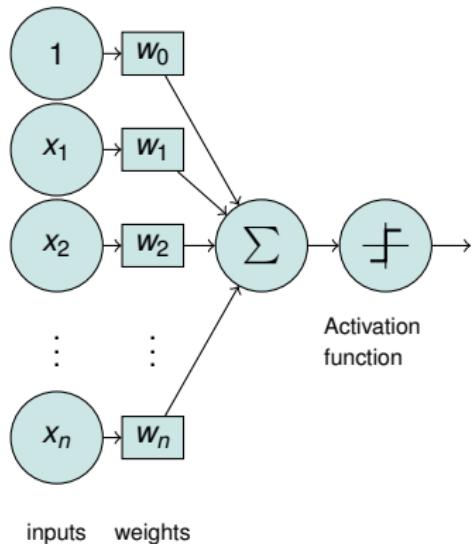
- In 1957, Frank Rosenblatt [17] invented the Perceptron
- Binary classification $y \in \{-1, 1\}$.
- It computes the function

$$\hat{y} = \text{sign}(\mathbf{w}^T \mathbf{x}),$$

where

$\mathbf{w} = (w_0, \dots, w_n)$: set of weights
(w_0 =bias)

$\mathbf{x} = (1, x_1, \dots, x_n)$: input feature vector



Perceptron Objective Function

Task: Find weights that minimize the distance of misclassified samples to the decision boundary

Assumptions

- Let $S = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_m, y_m)\}$ be a training data set
- Let \mathcal{M} be the set of misclassified feature vectors $y_i \neq \hat{y}_i = \text{sign}(\mathbf{w}^\top \mathbf{x}_i)$ according to a given set of weights \mathbf{w}
- Optimization problem:

$$\underset{\mathbf{w}}{\operatorname{argmin}} \quad \left\{ D(\mathbf{w}) = - \sum_{\mathbf{x}_i \in \mathcal{M}} y_i \cdot (\mathbf{w}^\top \mathbf{x}_i) \right\}$$

Perceptron Objective Function – Observations

- Objective function depends on misclassified feature vectors $\mathcal{M} \rightarrow$ iterative optimization
- In each iteration, the cardinality and composition of \mathcal{M} may change
- The gradient of the objective function is:

$$\nabla D(\mathbf{w}) = - \sum_{x_i \in \mathcal{M}} y_i \cdot \mathbf{x}_i$$

Perceptron Training

- Strategy 1: Process all samples, then perform weight update
- Strategy 2: Take an update step right after each misclassified sample
- Update rule in iteration $(k + 1)$ for the misclassified sample \mathbf{x}_i simplifies to:

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} + y_i \cdot \mathbf{x}_i$$

- Optimization until convergence or for a predefined number of iterations

**NEXT TIME
ON DEEP LEARNING**

Deep Learning for Beginners - Introduction 5

A. Maier, K. Breininger, V. Christlein, S. Seitz, A. Sindel, T. Würfl

Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg





Pattern
Recognition
Lab



virtuelle
hochschule
bayern

FAU

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Organizational Matters



Grading

- Module consists of lecture **and** exercises (together 2.5 ECTS)
- Written exam in the semester break, determines grade
- Exercises are **highly recommended**
- Unit tests available for exercises, no exercise submission

Exercise Content

- Fundamentals in signal and image processing
 - Fast Fourier transform
 - Convolution with filter kernels
- Developing a basic neural network framework from scratch
 - Feed Forward Neural Networks
 - Convolutional Neural Networks
 - Regularization
- Using Python jupyter notebooks
 - Implementation in pure Python (using Numpy, Scipy etc.)
 - No large hardware requirements (code runs on CPU)
 - No deep learning frameworks such as PyTorch or TensorFlow

Exercises

- Exercise requirements
 - Basic knowledge of Python and Numpy
 - Linear algebra, -
 - Image processing, -
 - Pattern recognition fundamentals
- How it works
 - Six exercises throughout the semester
 - Unit tests for all exercises

Summary

- Deep learning more and more present in day to day life
- Huge support and interest from industry
- **Very** active area of research!
- Perceptron as binary classifier motivated by biological neurons

**NEXT TIME
ON DEEP LEARNING**

Next Lecture Block

- Extending the Perceptron to obtain a universal function approximator
- Gradient based training algorithm for these models
- Efficient automatic computation of gradients

Comprehensive Questions

- What are the six postulates of pattern recognition?
- What is the Perceptron objective function?
- Can you name three applications successfully tackled by deep learning?

Further Reading

- [Link](#) - Deep learning book
- [Link](#) - Research and publications at the Pattern Recognition Lab
- [Link](#) - Google AI Blog with posts on e.g. [Deep dream](#) or [Alpha Go](#)

Questions?

References



References I

- [1] Jia Deng, Wei Dong, Richard Socher, et al. "Imagenet: A large-scale hierarchical image database". In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. IEEE. 2009, pp. 248–255.
- [2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: Advances in Neural Information Processing Systems 25. Curran Associates, Inc., 2012, pp. 1097–1105.
- [3] David Silver, Aja Huang, Chris J. Maddison, et al. "Mastering the game of Go with deep neural networks and tree search". In: Nature 529.7587 (Jan. 2016), pp. 484–489.
- [4] David Silver, Julian Schrittwieser, Karen Simonyan, et al. "Mastering the game of go without human knowledge". In: Nature 550.7676 (2017), p. 354.

References II

- [5] David Silver, Thomas Hubert, Julian Schrittwieser, et al. "Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm". In: [arXiv preprint arXiv:1712.01815](#) (2017).
- [6] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, et al. "You Only Look Once: Unified, Real-Time Object Detection". In: [CoRR abs/1506.02640](#) (2015).
- [7] J. Redmon and A. Farhadi. "YOLO9000: Better, Faster, Stronger". In: [ArXiv e-prints](#) (Dec. 2016). arXiv: 1612.08242 [cs.CV].
- [8] Joseph Redmon and Ali Farhadi. "YOLOv3: An Incremental Improvement". In: [arXiv](#) (2018).
- [9] David Bernecker, Christian Riess, Elli Angelopoulou, et al. "Continuous short-term irradiance forecasts using sky images". In: [Solar Energy](#) 110 (2014), pp. 303–315.

References III

- [10] Vincent Christlein, David Bernecker, Florian Höning, et al. "Writer Identification Using GMM Supervectors and Exemplar-SVMs". In: Pattern Recognition 63 (2017), pp. 258–267.
- [11] M. Aubreville, M. Krappmann, C. Bertram, et al. "A Guided Spatial Transformer Network for Histology Cell Differentiation". In: ArXiv e-prints (July 2017). arXiv: 1707.08525 [cs.CV].
- [12] Florin Cristian Ghesu, Bogdan Georgescu, Tommaso Mansi, et al. "An Artificial Agent for Anatomical Landmark Detection in Medical Images". In: Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016 Athens, 2016, pp. 229–237.
- [13] S. E. Wei, V. Ramakrishna, T. Kanade, et al. "Convolutional Pose Machines". In: CVPR. 2016, pp. 4724–4732.

References IV

- [14] Patrick Ferdinand Christ, Mohamed Ezzeldin A Elshaer, Florian Ettlinger, et al. "Automatic liver and lesion segmentation in CT using cascaded fully convolutional neural networks and 3D conditional random fields". In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer. 2016, pp. 415–423.
- [15] A. Karpathy and L. Fei-Fei. "Deep Visual-Semantic Alignments for Generating Image Descriptions". In: ArXiv e-prints (Dec. 2014). arXiv: 1412.2306 [cs.CV].
- [16] Tobias Würfl, Florin C Ghesu, Vincent Christlein, et al. "Deep learning computed tomography". In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer International Publishing. 2016, pp. 432–440.

References V

- [17] Frank Rosenblatt. The Perceptron—a perceiving and recognizing automaton. 85-460-1. Cornell Aeronautical Laboratory, 1957.
- [18] Olga Russakovsky, Jia Deng, Hao Su, et al. “ImageNet Large Scale Visual Recognition Challenge”. In: International Journal of Computer Vision 115.3 (2015), pp. 211–252.