KONICA MINOLTA

Technische Universität München

IS *AI* GOING TO REPLACE YOUR DOCTOR ?

(Medical applications of Artificial Intelligence)

Dr. Marie Piraud

Konica Minolta Laboratory Europe
(Guest researcher TU München)

ON THE NEWS…

AI reaching human-level

theguardian

13 May 2015

Computers now better than humans at recognising and sorting images

Bathtub

Junco

Plane

ON THE NEWS…

AI reaching human-level

theguardian

13 May 2015

Computers now better than humans at recognising and sorting images

Bathtub | Junco | Plane

9 March 2016

KONICA MINOLTA

ON THE NEWS…

Also in healthcare!

Dermatologist-level classification of skin cancer with deep neural networks

Nature **542**, 115–118 (02 February 2017)

KONICA MINOLTA

ON THE NEWS…

Also in healthcare!

Dermatologist-level classification of skin cancer with deep neural networks

Nature **542**, 115–118 (02 February 2017)

Forget your GP, robots will 'soon be able to diagnose more ac almost any doctor'

Frank Chen
@withfries2

Folgen

Geoff Hinton: "we should stop training radiologists right now, in 5 years #deeplearning will have better performance" #mkt4intel

12:18 - 27. Okt. 2016 aus Toronto, Ontario

KONICA MINOLTA

ON THE NEWS…

Also in healthcare!

Dermatologist-level classification of skin cancer with deep neural networks

Nature **542**, 115–118 (02 February 2017)

- Exciting?
- Scary?

Forget your GP, robots will 'soon be able to diagnose more ac
almost any doctor'

Frank Chen
@withfries2

Folgen

Geoff Hinton: "we should stop training radiologists right now, in 5 years #deeplearning will have better performance" #mkt4intel

12:18 - 27. Okt. 2016 aus Toronto, Ontario

KONICA MINOLTA

OUTLINE

- **Introduction to *AI* and Deep learning in healthcare**
  - *Demystifying AI*
  - *A success story: skin cancer detection*
- Challenges and current limitations
  - *The data challenge*
  - *The curse of dimensionality*
  - *The ‘black-box’*
- Solutions: current research topics
  - *Reducing the need for annotations*
  - *Introducing prior knowledge*
  - *Dealing with volumetric data*
  - *Decision visualisation*

WHAT IS AI?

„Machines that learn, think and act like humans“

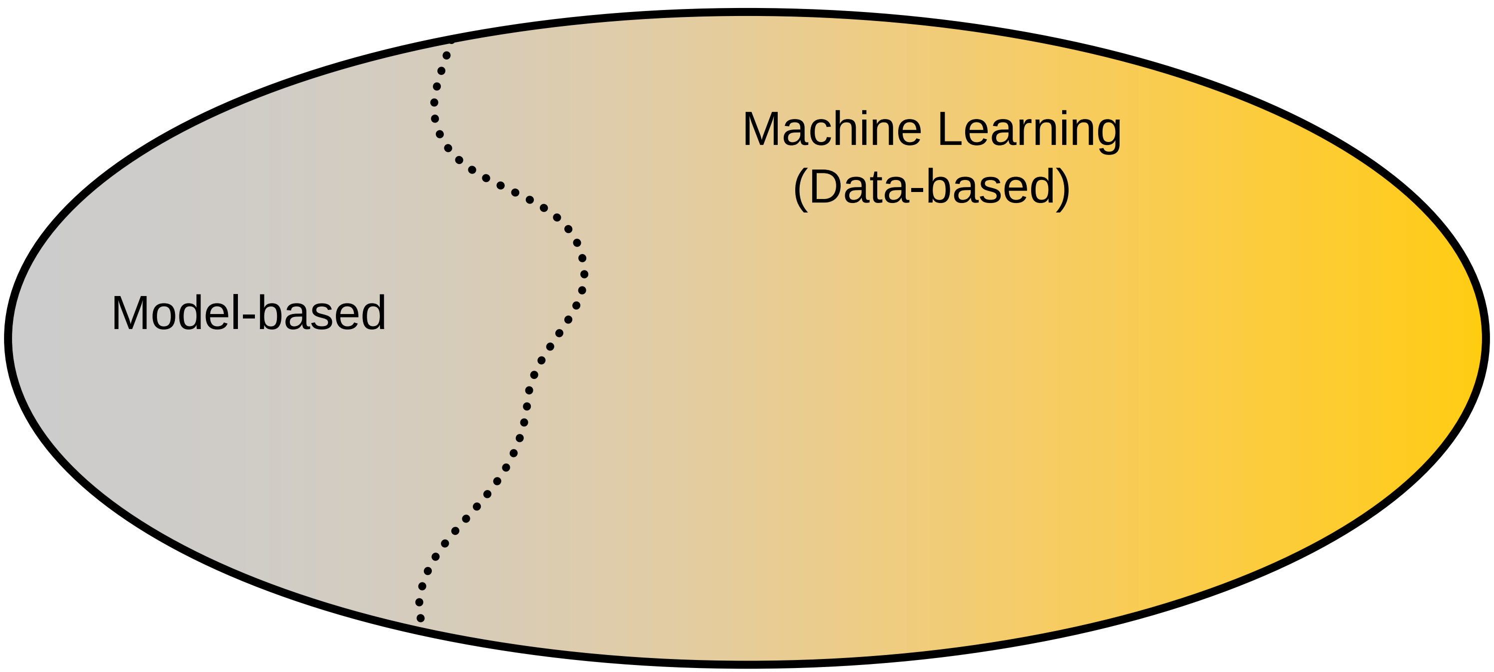

KONICA MINOLTA

WHAT IS AI?

„Machines that learn, think and act like humans“

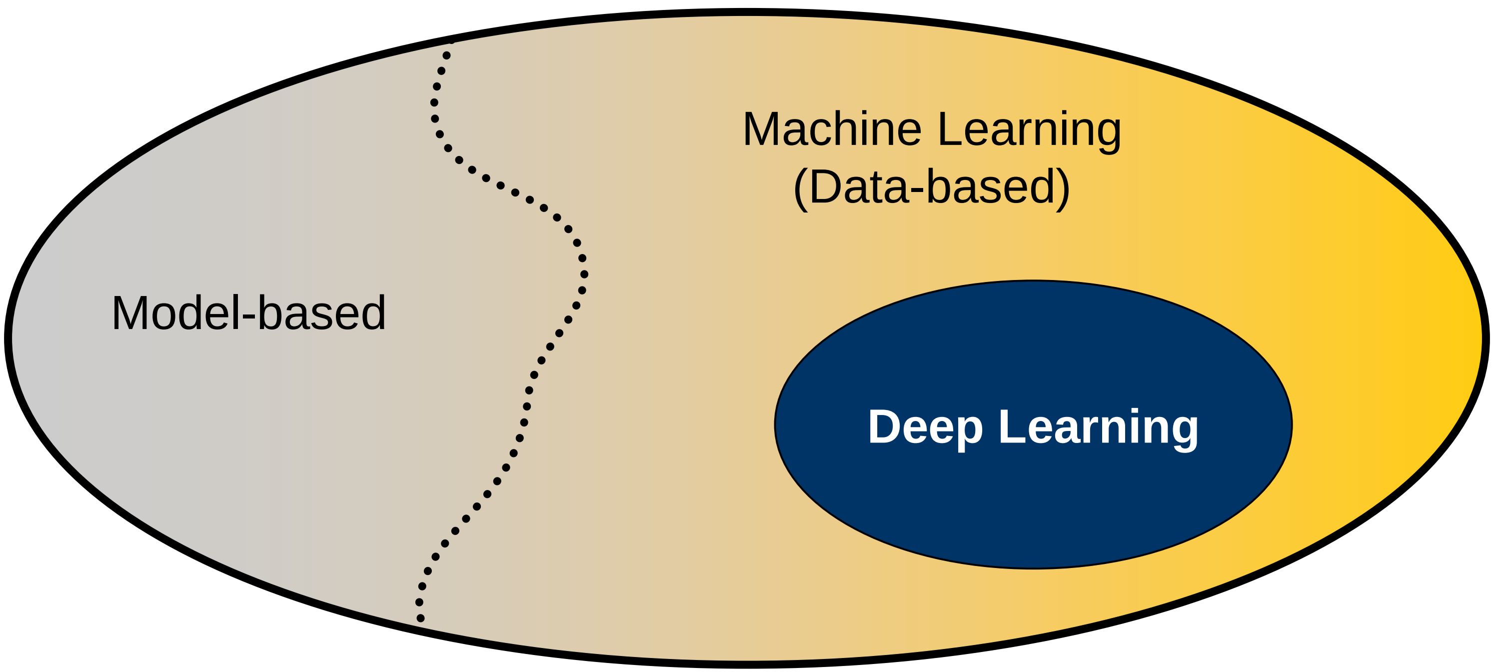

KONICA MINOLTA

LEARNING FROM DATA

LEARNING FROM DATA

Regression

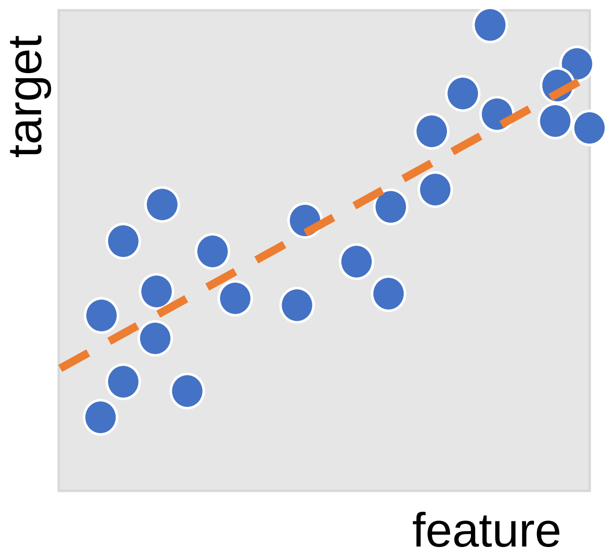

LEARNING FROM DATA

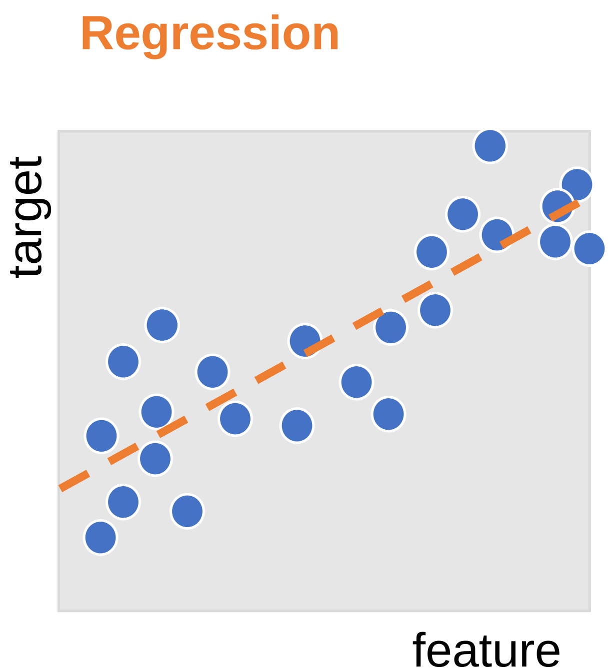

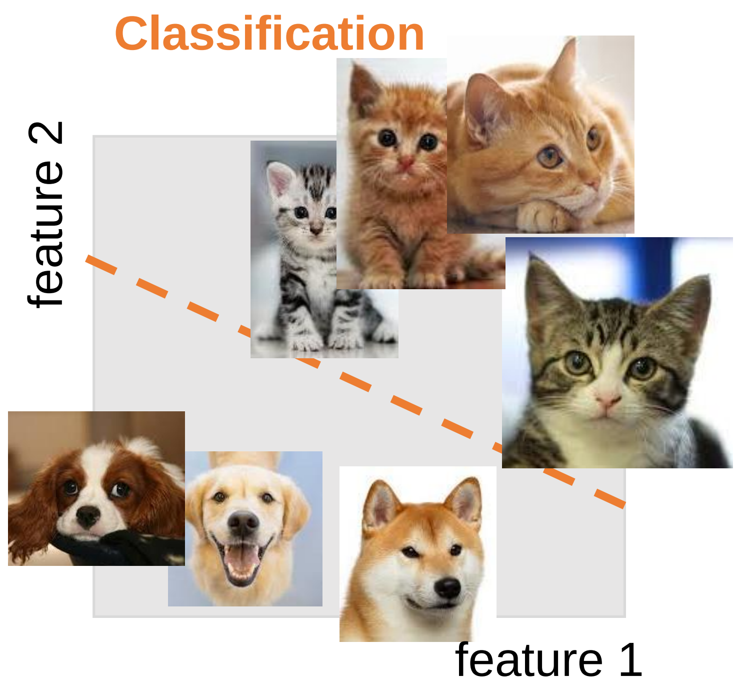

KONICA MINOLTA

LEARNING FROM DATA

Regression

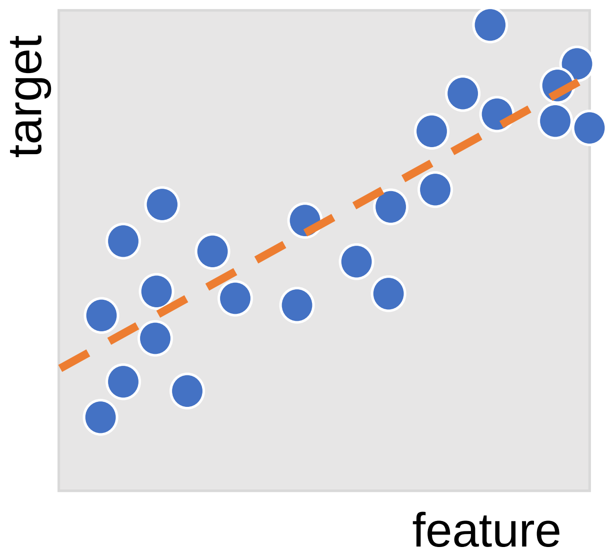

Classification

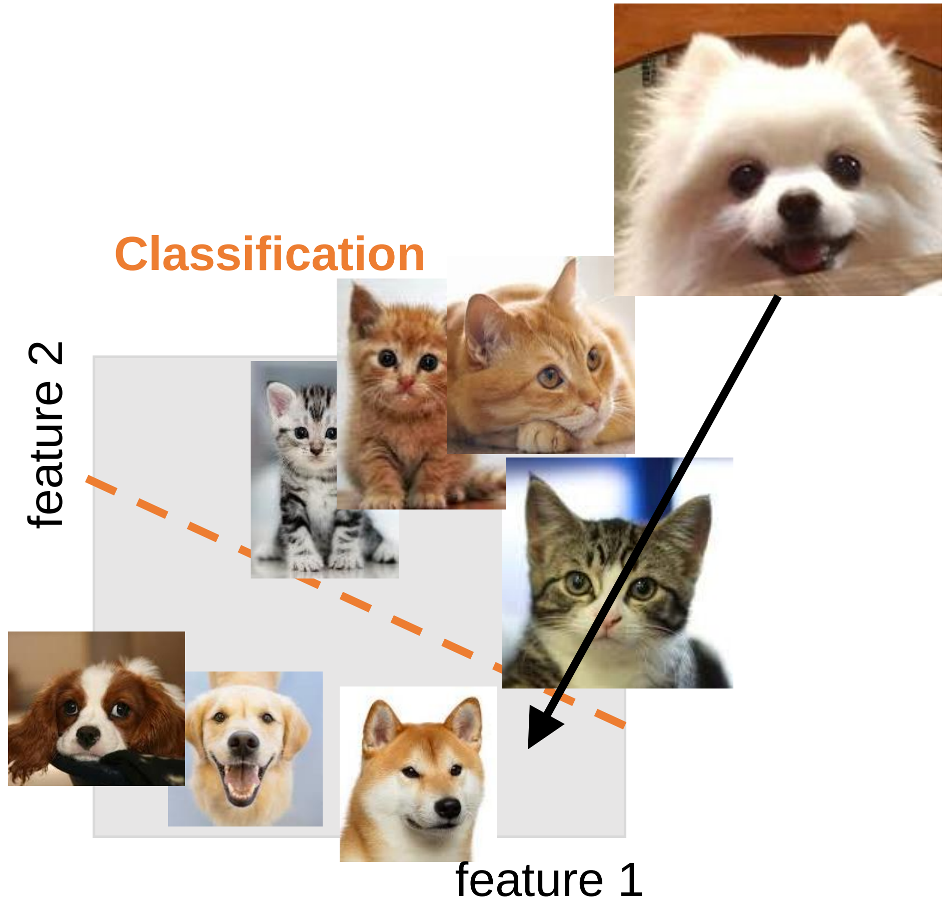

LEARNING FROM DATA

Regression

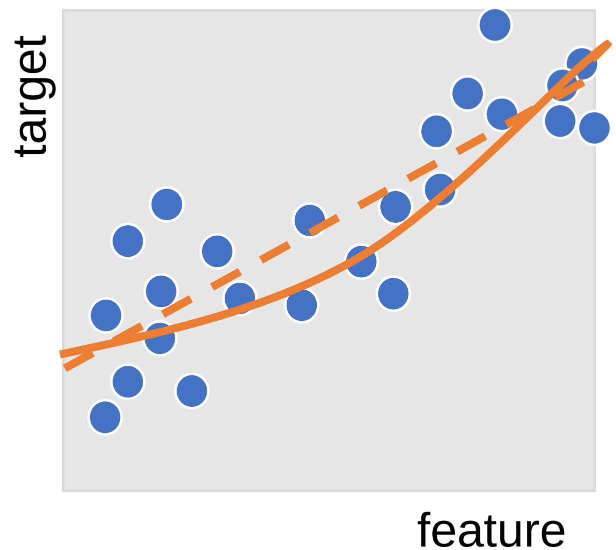

Classification

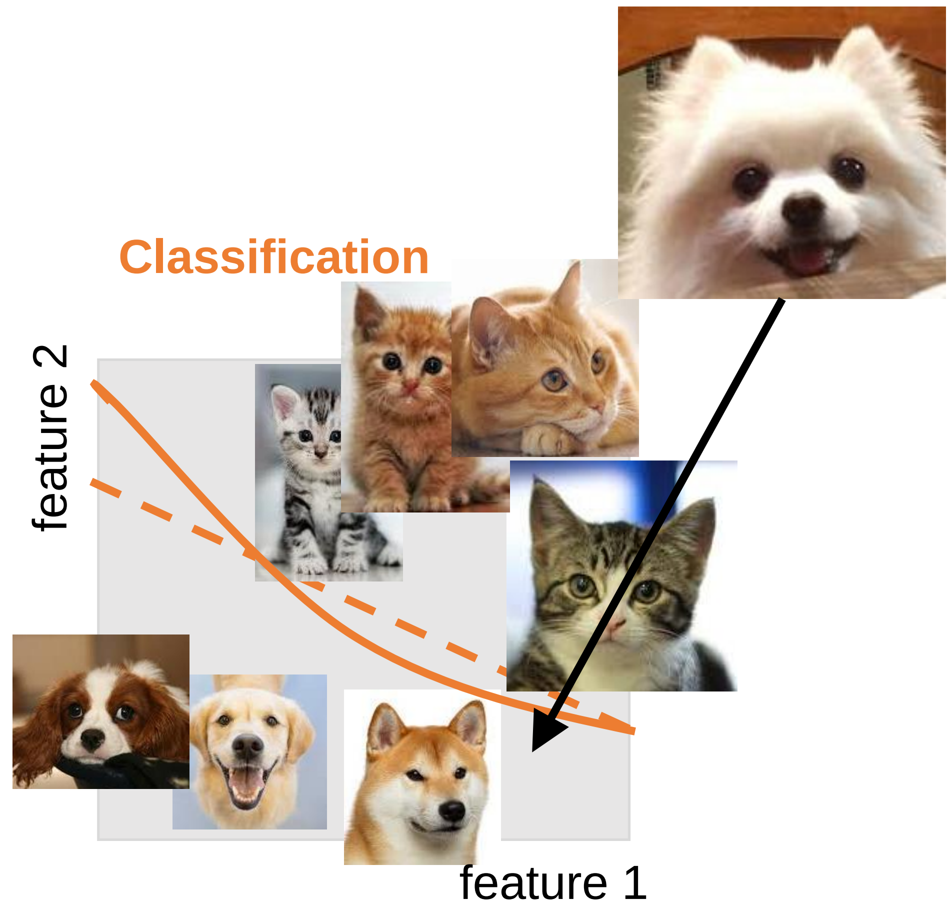

LEARNING FROM DATA

From regression to Deep Learning

Regression

- Linear: **z=ax+by**

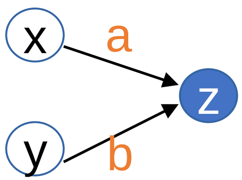

- Non-linear: **z=f(ax+by)**

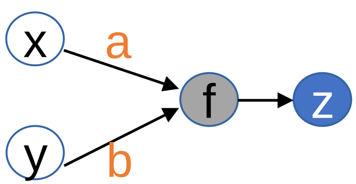

LEARNING FROM DATA

From regression to Deep Learning

Regression

- Linear: $z=ax+by$

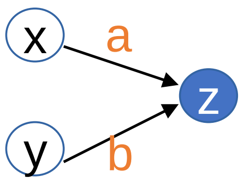

- Non-linear: $z=f(ax+by)$

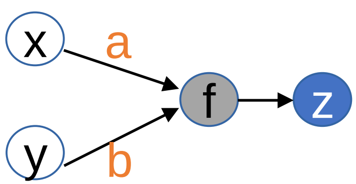

Deep Learning (~Highly non-linear)

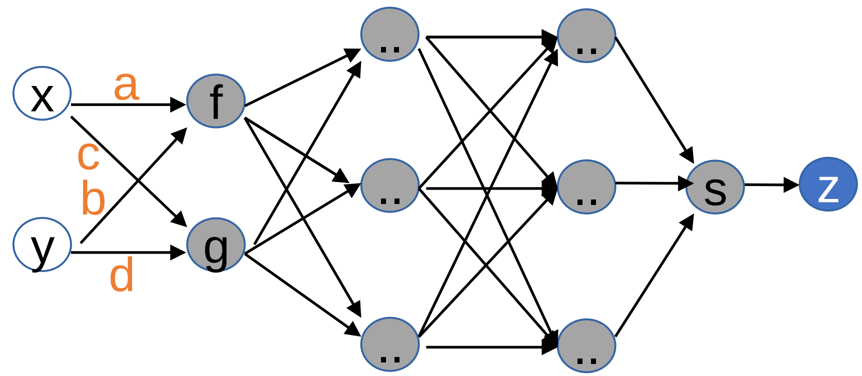

→ ‘trained’ on data

LEARNING FROM DATA

From regression to Deep Learning

‘Artificial Neural Networks’

Human brain: ~10^{14} neurons /

AlphaGo: ~10^{5} ‘neurons’ /

Regression

- Linear: **z=ax+by**

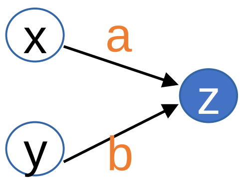

- Non-linear: **z=f(ax+by)**

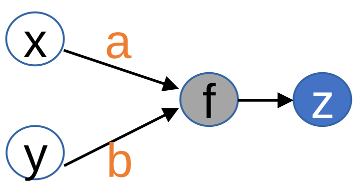

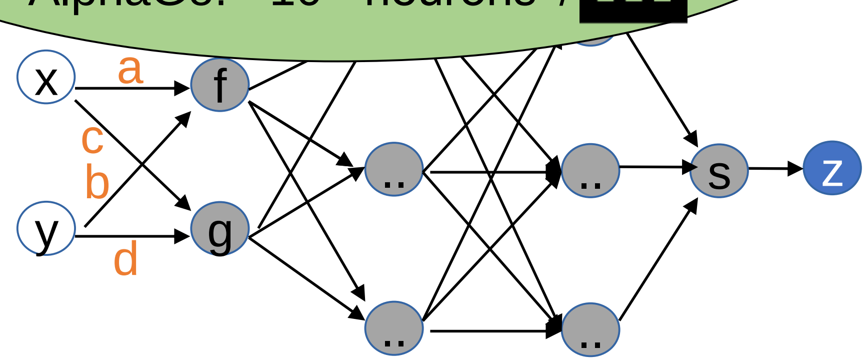

→ ‘trained’ on data

LEARNING FROM DATA

From regression to Deep Learning

Regression

- Linear: **z=ax+by**

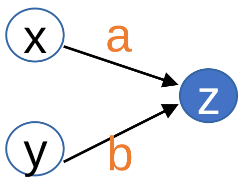

- Non-linear: **z=f(ax+by)**

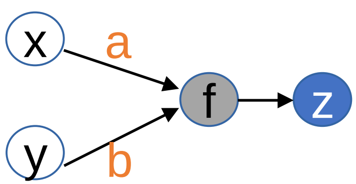

‘Artificial Neural Networks’

Human brain: ~10^{14} neurons /

AlphaGo: ~10^{5} ‘neurons’ /

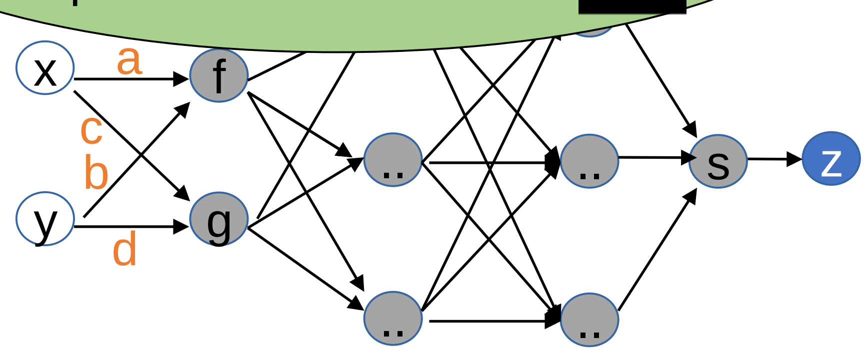

→ ‘trained’ on data

A lot of data !

ImageNet: 1.2 Mio

AlphaGo: 30 Mio

LEARNING FROM DATA

Convolutional Neural Networks for Computer Vision Applications

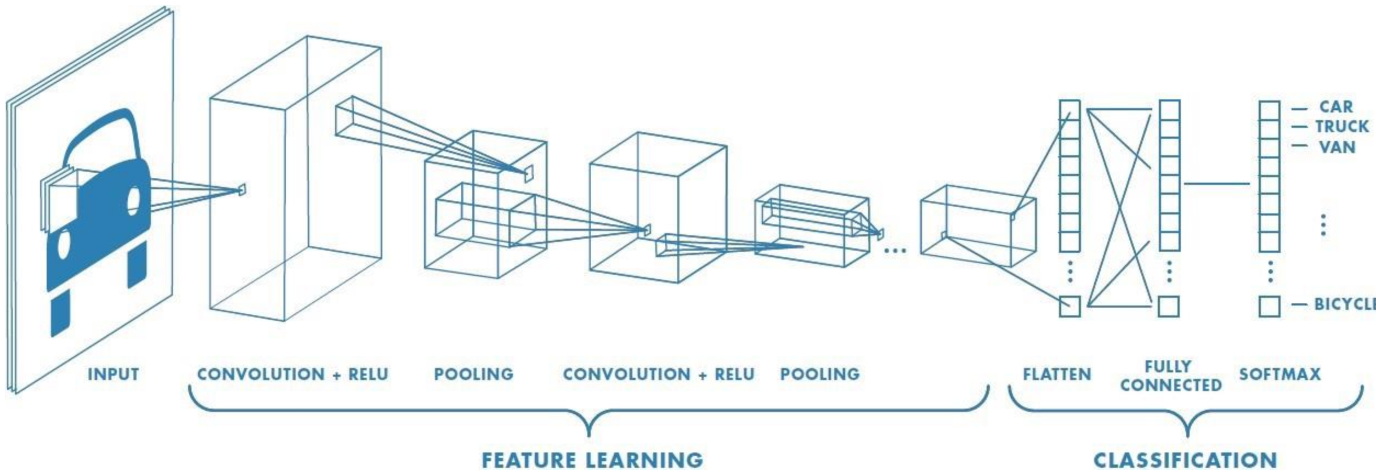

https://www.mathworks.com/videos/introduction-to-deep-learning-what-are-convolutional-neural-networks--1489512765771.html

Applications

- Classification of images → Skin-moles classification
- Detection of objects → Vertebrae detection
- Automated segmentation → Heart ventricle volume measurement
- Detection of abnormalities → Detection of eye fondus hemorrage

LEARNING FROM DATA

Summary

DL = fancy and highly non-linear regression

- Neural networks are trained by examples
- The features that are used for decision-making are ‚learned' by the network itself
- Lots of examples are needed to properly cover the variability of the data
- After training, networks can be used to predict information (e.g. classification of an image) on unseen data

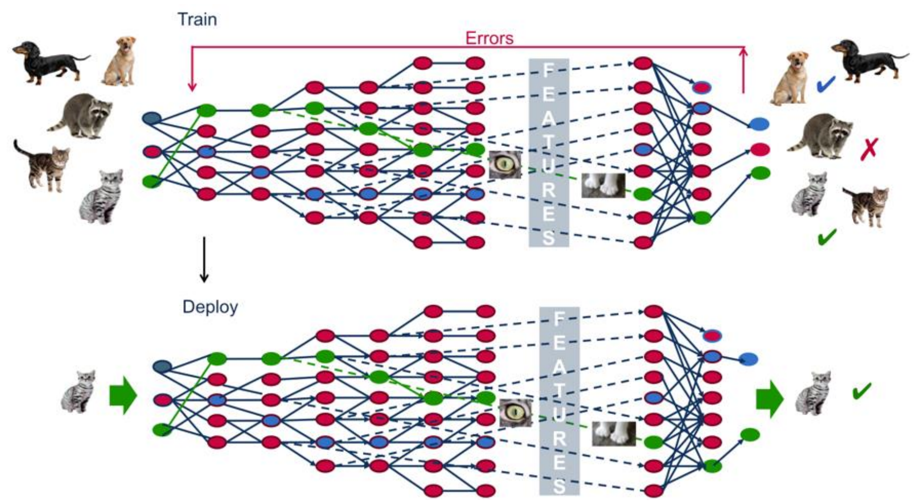

https://mapr.com/blog/demystifying-ai-ml-dl/

DEEP LEARNING IN HEALTHCARE

Scientific publications

Number of papers

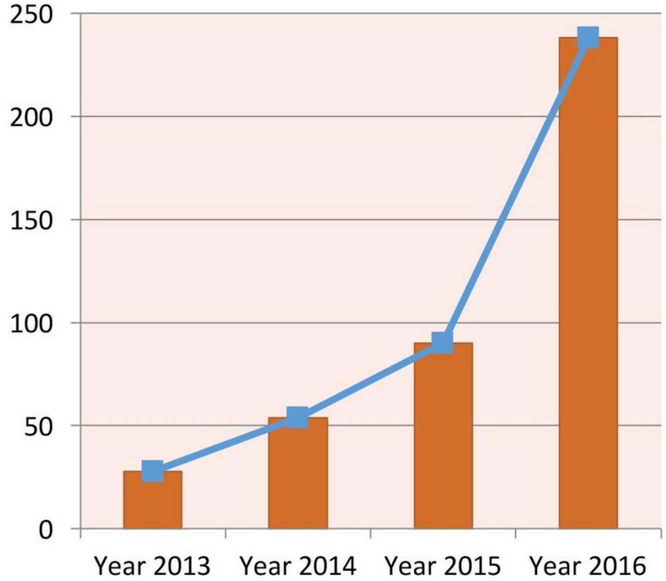

Data sources

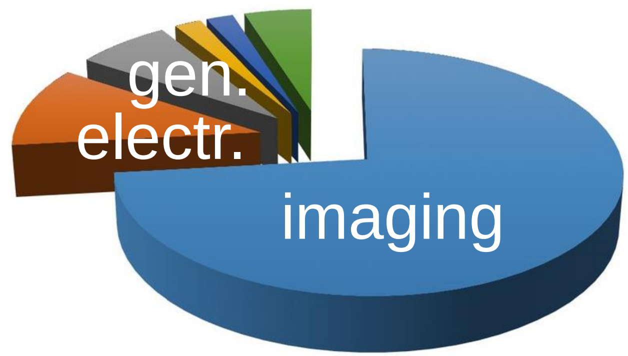

F. Jiang *etal,* Artificial intelligence in healthcare: past, present and future. Stroke and Vascular Neurology (2017)

KONICA MINOLTA

A SUCCESS STORY: SKIN CANCER DETECTION

Network architecture

Dermatologist-level classification of skin cancer with deep neural networks

A. Esteva *etal,* Nature **542**, 115-118 (2017)

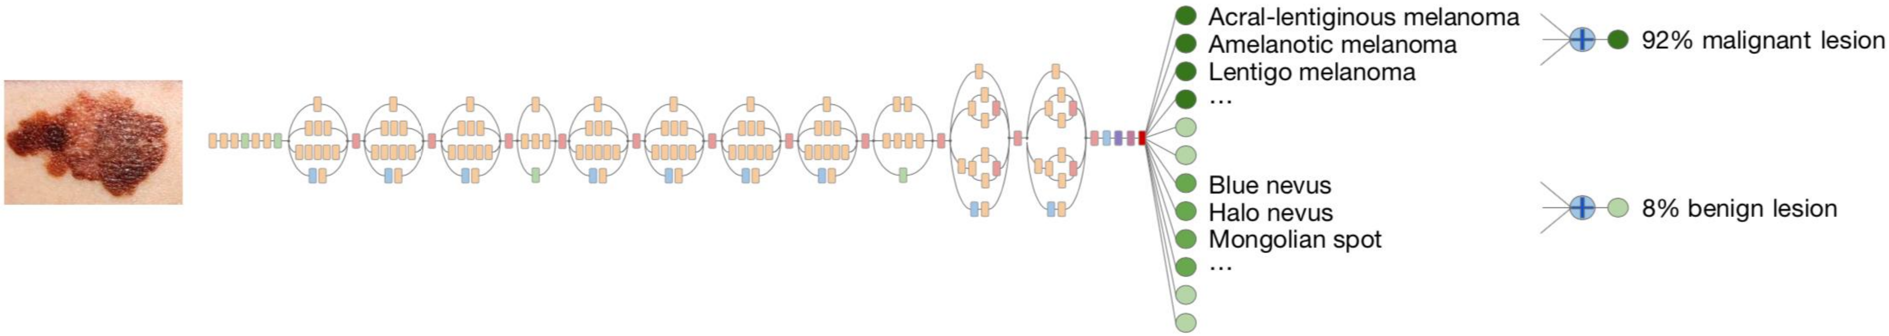

Inception v3: C. Szegedy *etal,* arXiv:1512.00567

KONICA MINOLTA

A SUCCESS STORY: SKIN CANCER DETECTION

Network architecture

Dermatologist-level classification of skin cancer with deep neural networks

A. Esteva *etal,* Nature **542**, 115-118 (2017)

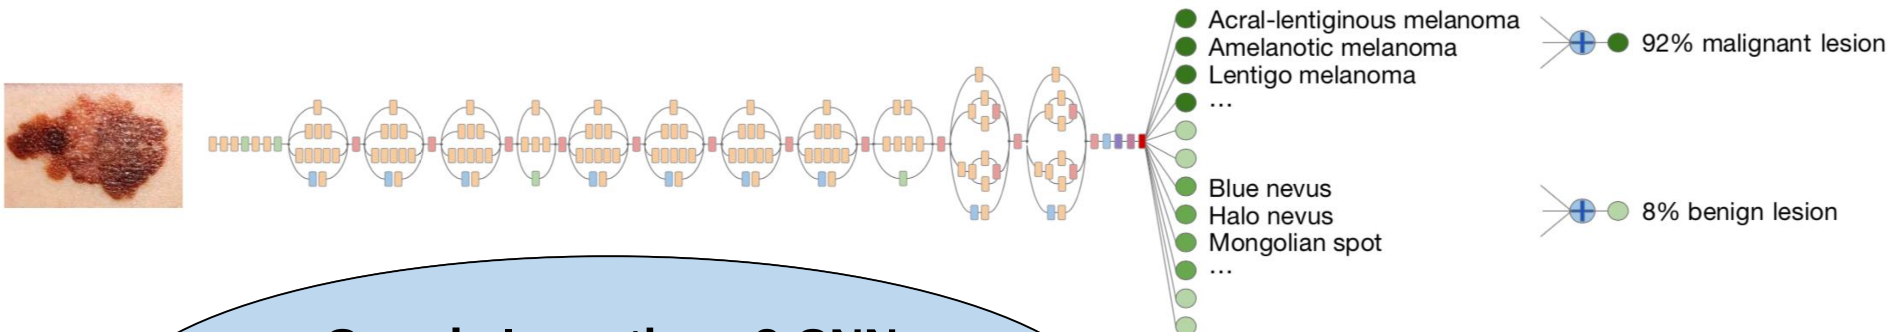

Google Inception v3 CNN

30 Mio parameters

130,000 / 757 classes

Inception v3: C. Szegedy *etal,* arXiv:1512.00567

A SUCCESS STORY: SKIN CANCER DETECTION

Dermatologist-level performance

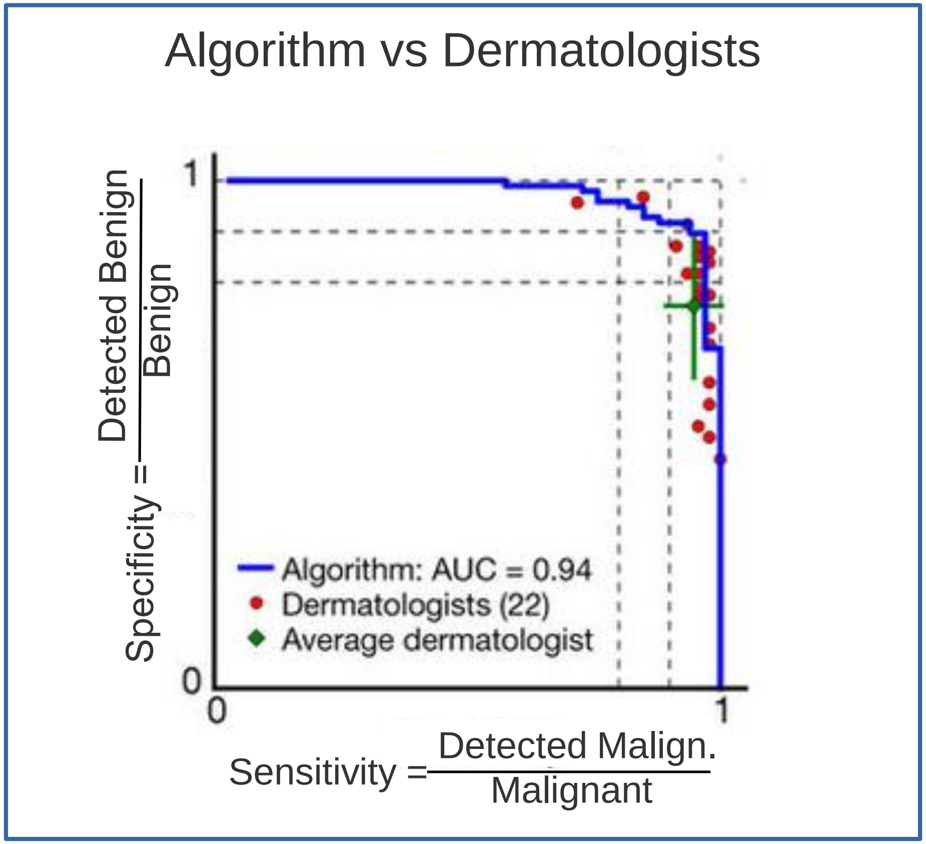

A. Esteva *etal,* Nature **542**, 115-118 (2017)

USE CASE 2: BLOOD CANCER

Technische Universität München

Enables better imaging biomarkers

Multiple Myeloma:

- plasma-cell disorder
- causes bone damage
- 3 stages:

MGUS/sMM/MM

KONICA MINOLTA

USE CASE 2: BLOOD CANCER

Enables better imaging biomarkers

Multiple Myeloma:
- plasma-cell disorder
- causes bone damage
- 3 stages:
MGUS/sMM/MM
→ **Risk stratification**

USE CASE 2: BLOOD CANCER

Enables better imaging biomarkers

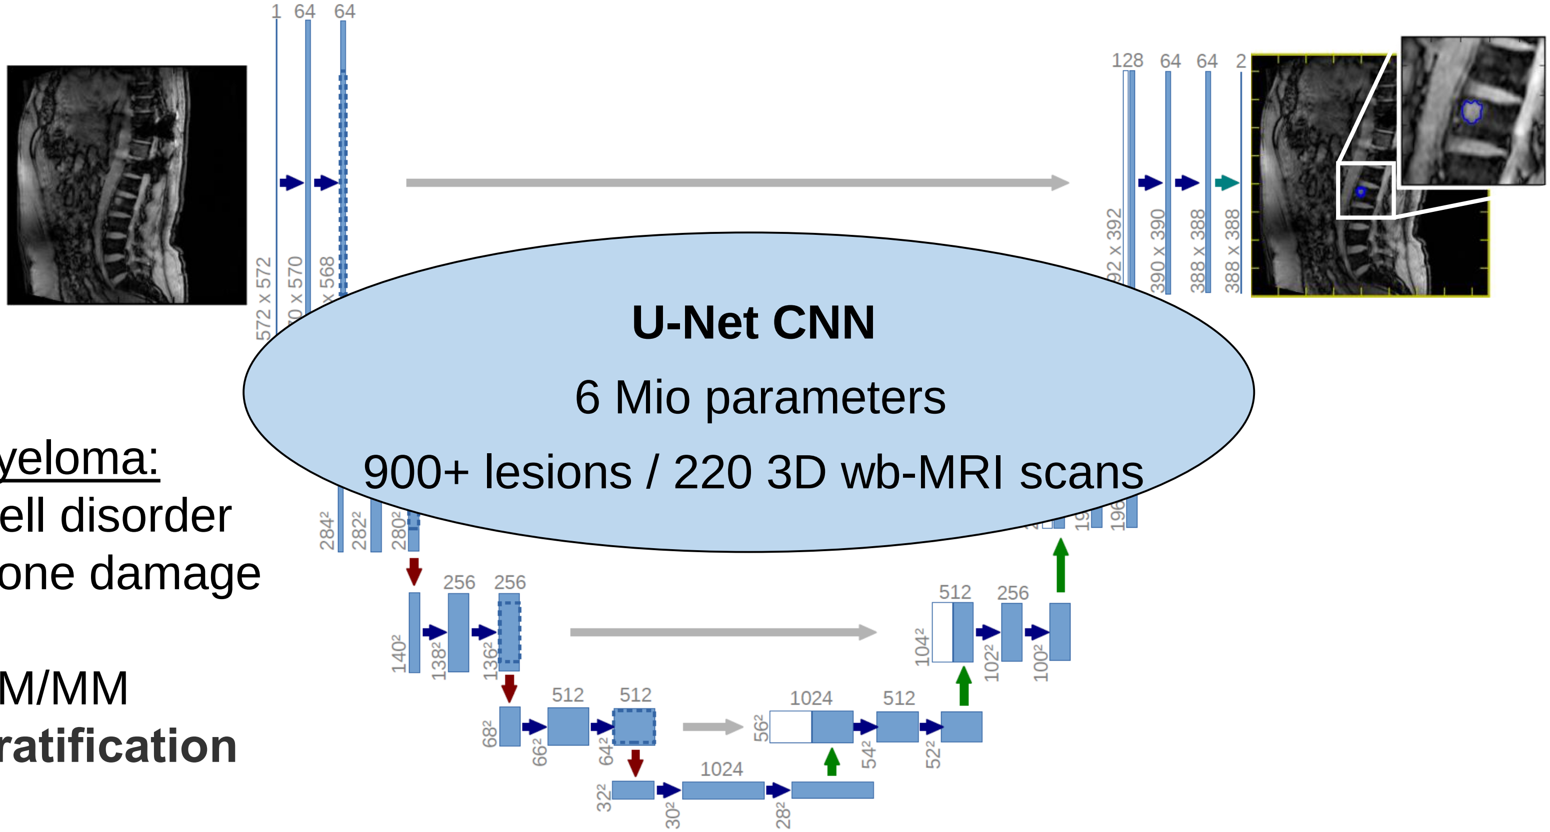

U-Net CNN

6 Mio parameters

900+ lesions / 220 3D wb-MRI scans

<u>Multiple Myeloma:</u>
- plasma-cell disorder
- causes bone damage
- 3 stages: MGUS/sMM/MM

→ **Risk stratification**

U-Net: O. Ronnerberger *etal,* MICCAI (2015)

USE CASE 2: BLOOD CANCER

Enables better imaging biomarkers

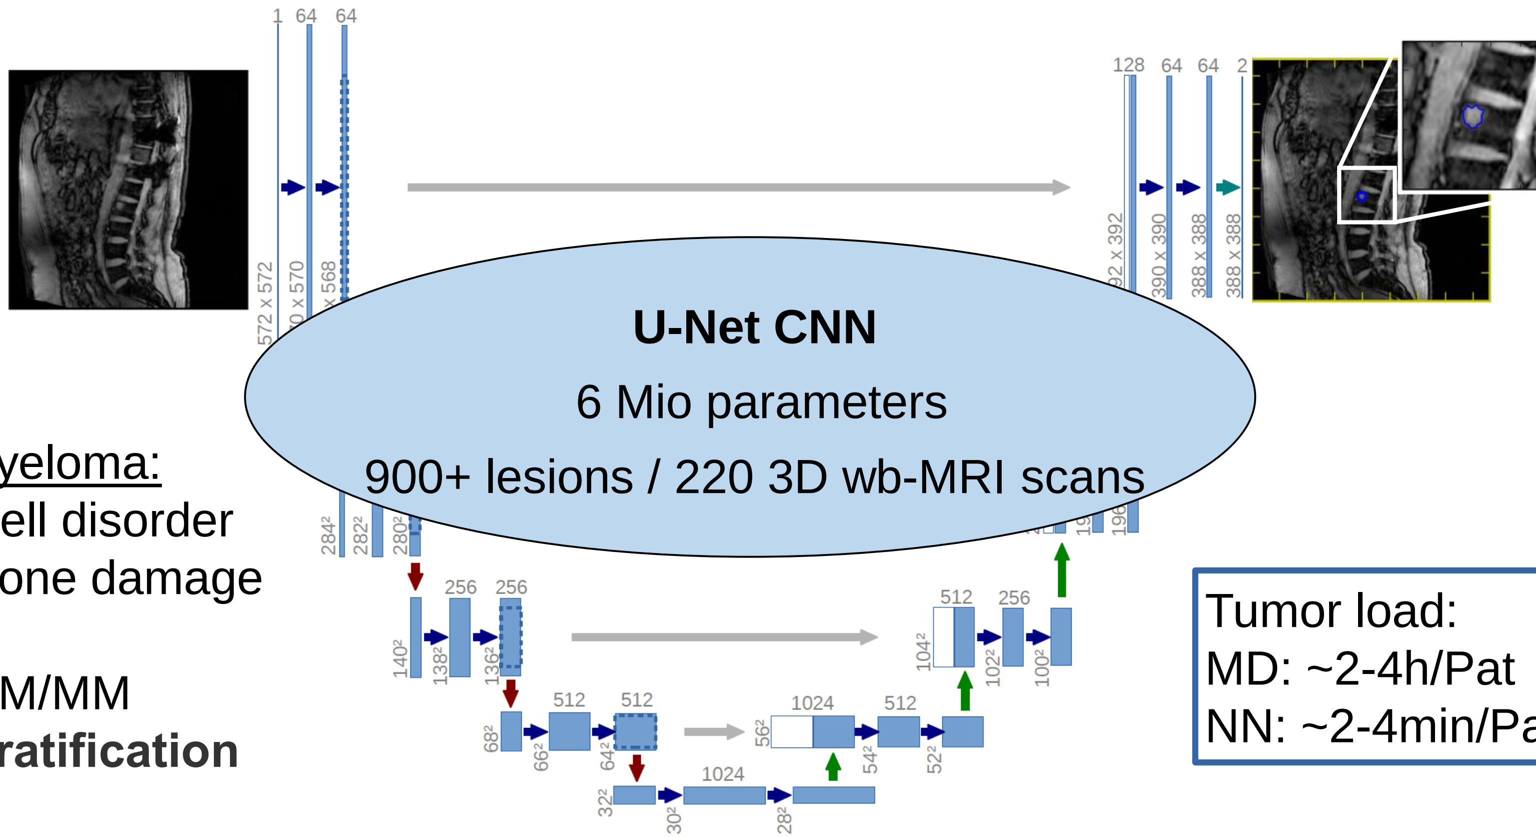

U-Net CNN

6 Mio parameters

900+ lesions / 220 3D wb-MRI scans

<u>Multiple Myeloma:</u>
- plasma-cell disorder
- causes bone damage
- 3 stages: MGUS/sMM/MM

→ **Risk stratification**

Tumor load:
MD: ~2-4h/Pat
NN: ~2-4min/Pat

M. Wennmann, …, <u>M. Piraud</u> *etal* Oncotarget (2018) / <u>M. Piraud</u> *etal,* bioRxiv: 613869 / M. Perkonigg, …, <u>M. Piraud</u> *etal,* Submitted (2019)

Technische Universität München

USE CASE 2: BLOOD CANCER

Enables better imaging biomarkers

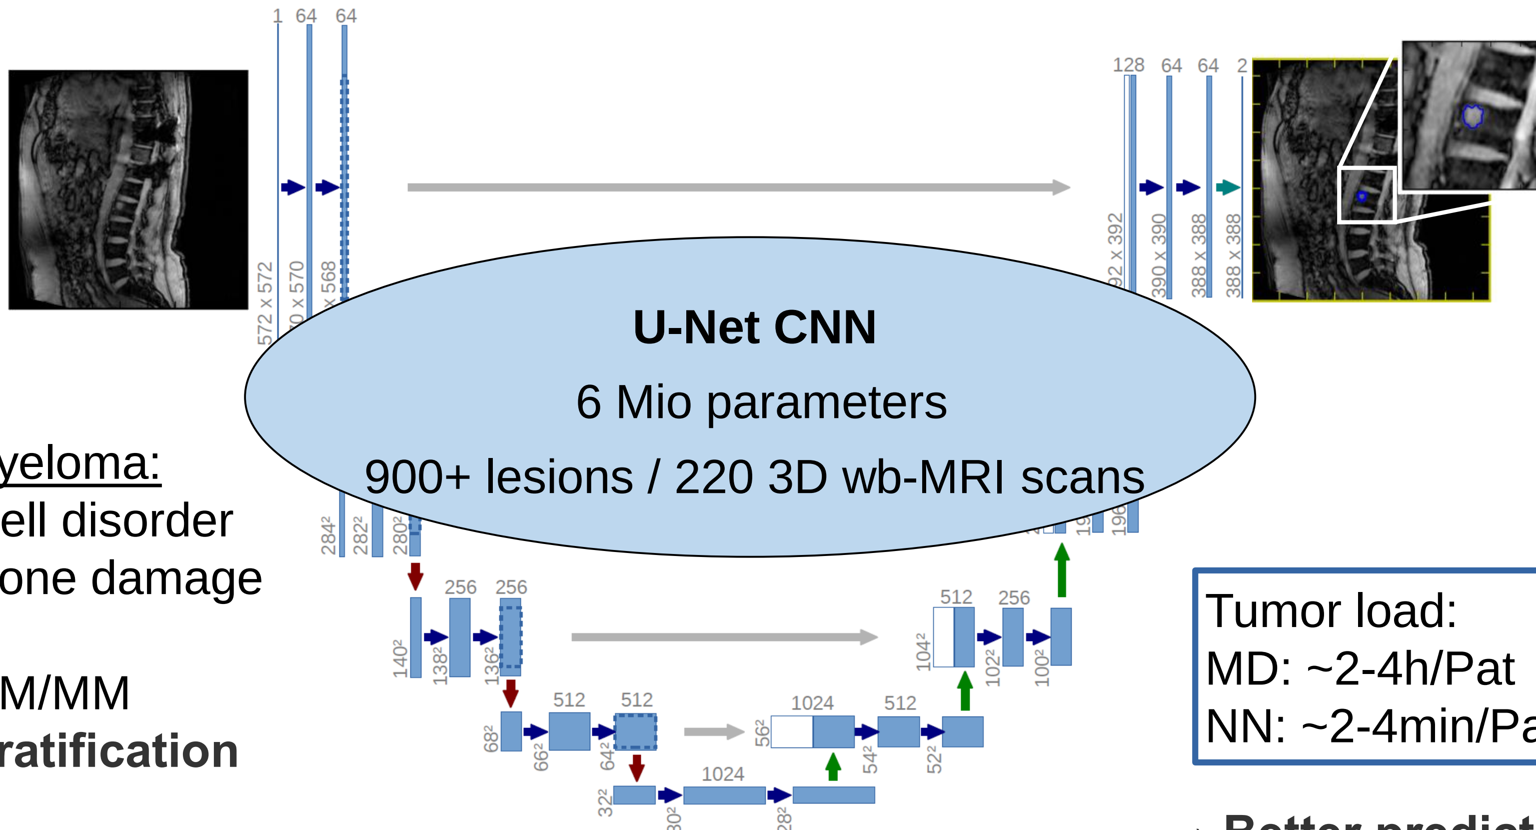

U-Net CNN

6 Mio parameters

900+ lesions / 220 3D wb-MRI scans

Multiple Myeloma:
- plasma-cell disorder
- causes bone damage
- 3 stages: MGUS/sMM/MM

→ **Risk stratification**

Tumor load:
MD: ~2-4h/Pat
NN: ~2-4min/Pat

→ **Better predictor!**

M. Wennmann, …, M. Piraud *etal* Oncotarget (2018) / M. Piraud *etal,* bioRxiv: 613869 /
M. Perkonigg, …, M. Piraud *etal,* Submitted (2019)

KONICA MINOLTA

OUTLINE

- Introduction to *AI* and Deep learning in healthcare
  - *Demystifying AI*
  - *A success story: skin cancer detection*
- **Challenges and current limitations**
  - *The data challenge*
  - *The curse of dimensionality*
  - *The 'black-box'*
- Solutions: current research topics
  - *Reducing the need for annotations*
  - *Introducing prior knowledge*
  - *Dealing with volumetric data*
  - *Decision visualisation*

CHALLENGES AND CURRENT LIMITATIONS

1) The data challenge

Variability: Medical data is ill-structured / highly-variable

Scarceness: expensive to collect and label

ImageNet: 1.2 Mio

Dermato article: 130,000

Med Study: 100s of cases

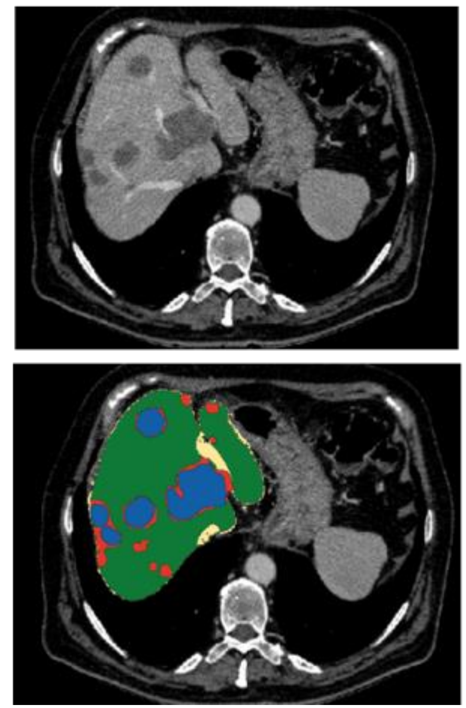

Christ *et al* MICCAI 2016

CHALLENGES AND CURRENT LIMITATIONS

2) The curse of dimensionality

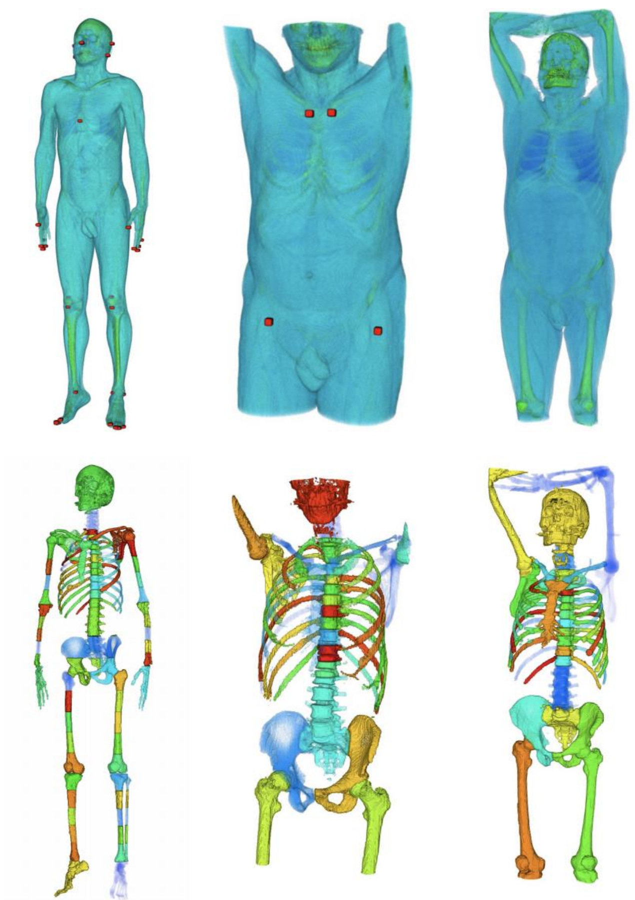

M. Bieth *etal* IEEE Trans. Med. Imag. 2017

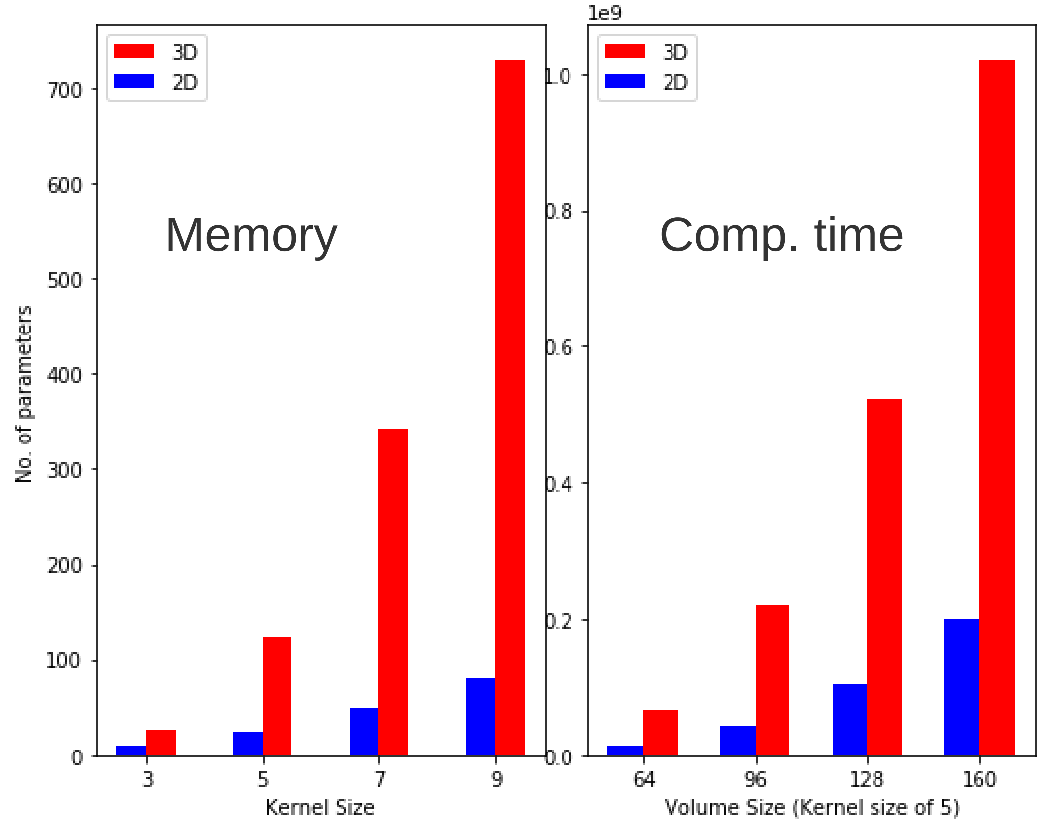

KONICA MINOLTA

CHALLENGES AND CURRENT LIMITATIONS

3) The ‚black-box‘ problem

No explicit programming! How are decisions reached?

Explanation:  to acceptance

GDPR: Automated individual decision-making should be **contestable**

CHALLENGES AND CURRENT LIMITATIONS

3) The ‚black-box' problem

No explicit programming! How are decisions reached?

Explanation:  to acceptance

GDPR: Automated individual decision-making should be **contestable**

→ **Testing/ Certification**

Jan 2017: 1st DL software to obtain FDA clearance...

ARTERYS www.arterys.com

CHALLENGES AND CURRENT LIMITATIONS

3) The ‚black-box' problem

No explicit programming! How are decisions reached?

Explanation: [key] to acceptance

GDPR: Automated individual decision-making should be **contestable**

→ Testing/ Certification

Jan 2017: 1st DL software to obtain FDA clearance...

www.arterys.com

April 2018: 1st DL diagnostic software to obtain FDA clearance

www.eyediagnosis.net

87% Sensitivity
90% Specificity

OUTLINE

- Introduction to *AI* and Deep learning in healthcare
  - *Demystifying AI*
  - *A success story: skin cancer detection*
- Challenges and current limitations
  - *The data challenge*
  - *The curse of dimensionality*
  - *The 'black-box'*
- **Solutions: current research topics**
  - *Reducing the need for annotations*
  - *Introducing prior knowledge*
  - *Dealing with volumetric data*
  - *Decision visualisation*

SOLUTIONS

Some of our research activities

Data scarceness and variability:

- Data augmentation
  & Semi-supervised learning
- Transfer learning
- Input prior knowledge

Volumetric Data:

- '2.5D' solutions
- Cross-hair filters

Hacking the black-box:

- Visualization
- Model-aware learning
- Link with dynamical systems

Reducing the need for annotations
in Breast Cancer detection

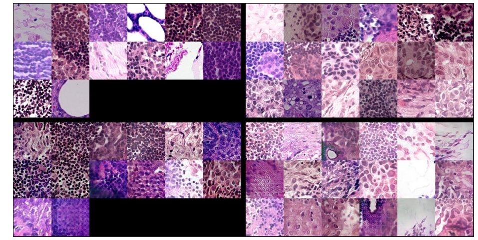

Multi-level activation for nuclei segmentation
BioComputing@ECCV 2018

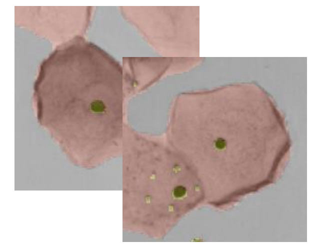

Cross-hair filters for Brain vessel segmentation
arXiv:1803.09340

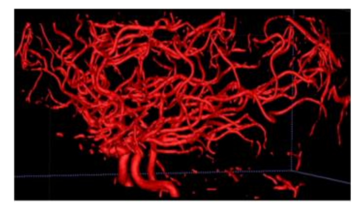

Class maximization for Breast Cancer classification

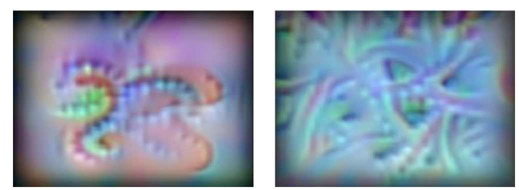

KONICA MINOLTA

REDUCING THE NEED FOR ANNOTATIONS

e.g. Whole-Slide Images for breast Cancer detection

CAMELYON 16 challenge
400 WSIs (100 000 px small side), >1TB total

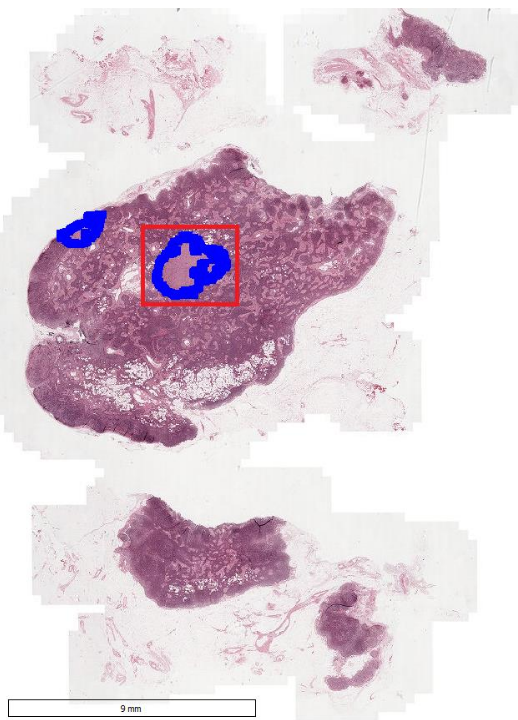

>

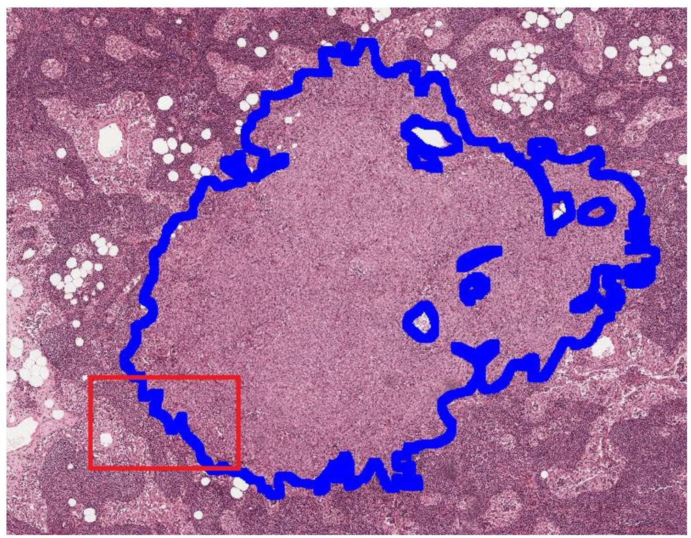

>

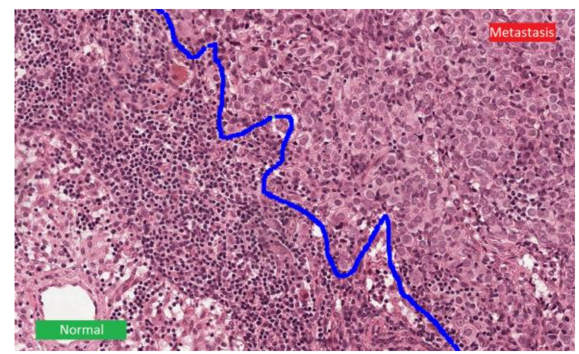

https://camelyon16.grand-challenge.org/

KONICA MINOLTA

REDUCING THE NEED FOR ANNOTATIONS

Data Augmentation and Generation

Patches generated by a FM-GAN

generated data

tumor patches

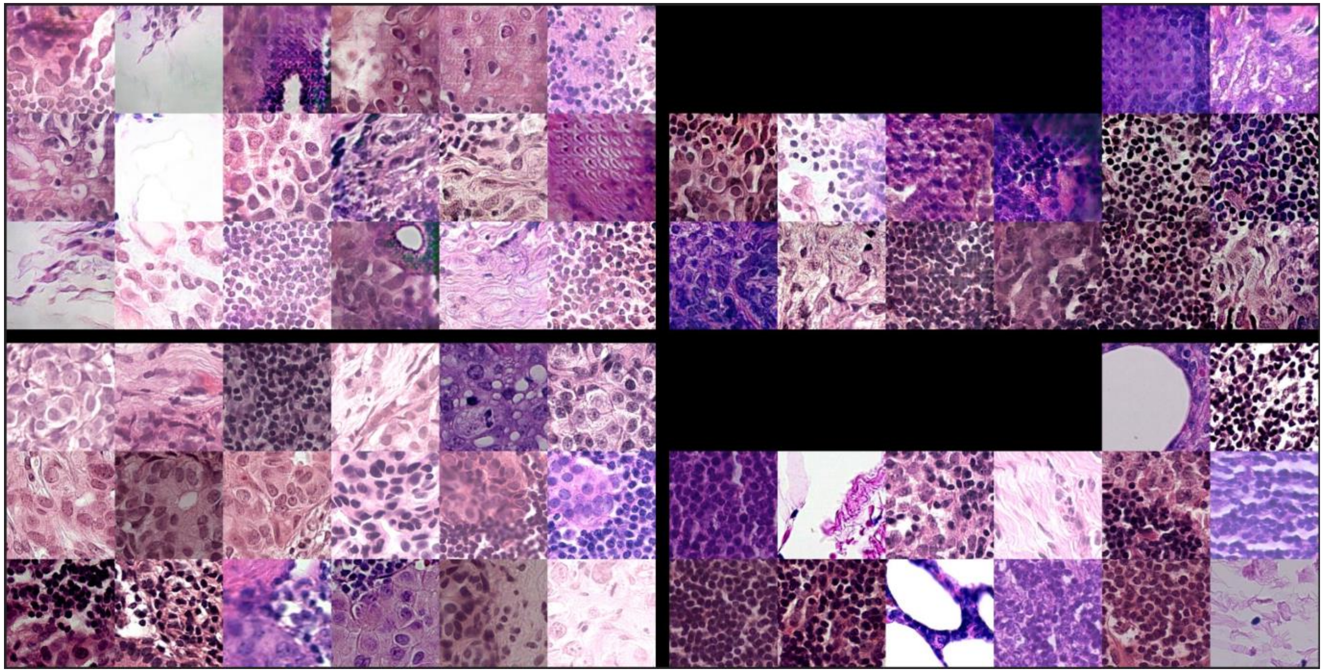

healthy patches

real data

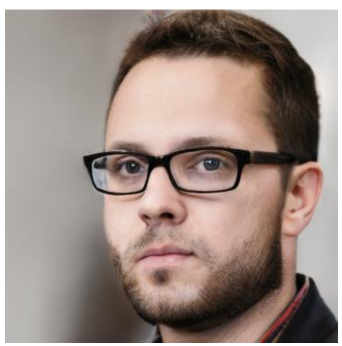

https://thispersondoesnotexist.com/

REDUCING THE NEED FOR ANNOTATIONS

Data Augmentation and Generation

Patches generated by a FM-GAN

generated data

tumor patches

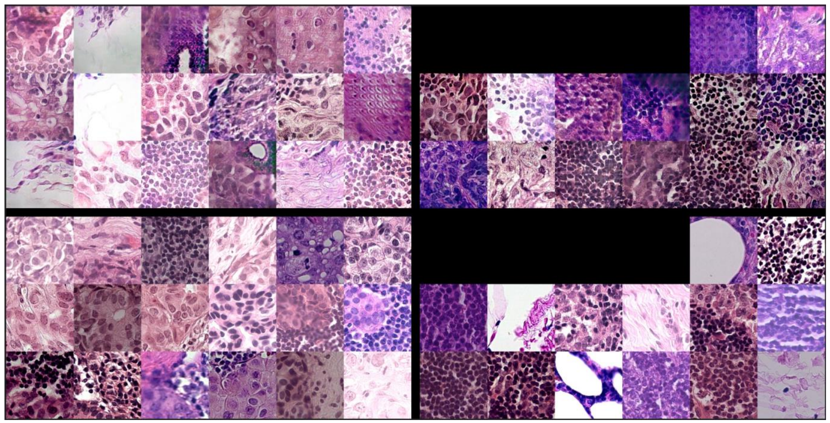

healthy patches

real data

CAMELYON dataset:
classify tumor patches

Method	FROC score (best: 100)
Supervised	61.1
Data augmentation (FM-GAN)	**65.1**

REDUCING THE NEED FOR ANNOTATIONS

Semi-supervised learning

Labeled set

Unlabeled set

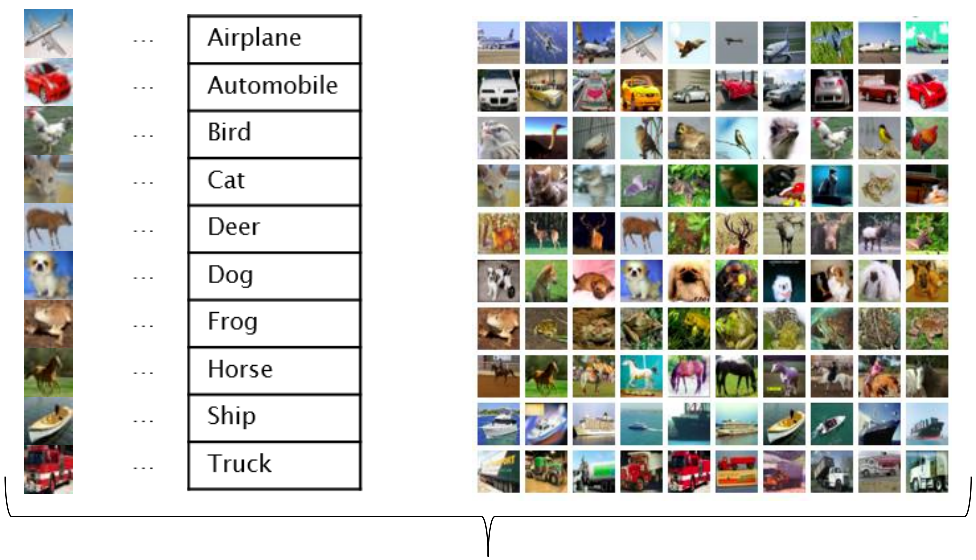

Full set

CAMELYON dataset: classify tumor patches

Method	FROC score (best: 100)
Supervised	61.1
Data augmentation (FM-GAN)	65.1
Semi-Supervised (Temporal ensembling)	**68.4**

10% labeled + 90% unlabeled examples

https://www.cs.toronto.edu/~kriz/cifar.html

INCLUDING PRIOR KNOWLEDGE

Technische Universität München

Multi-level activation for topological inclusion

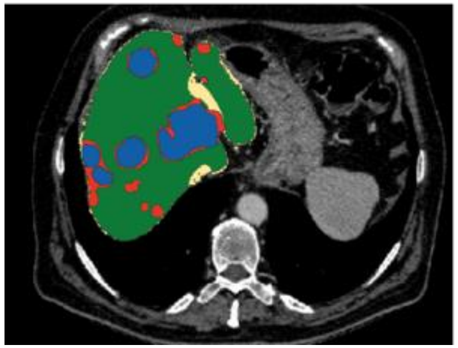

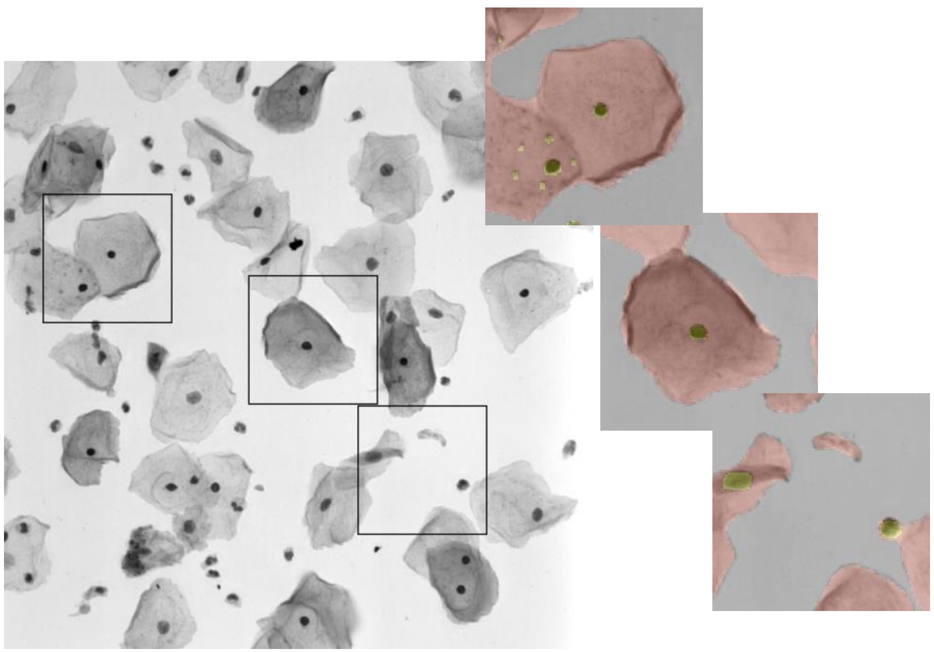

KONICA MINOLTA

INCLUDING PRIOR KNOWLEDGE

Technische Universität München

Multi-level activation for topological inclusion

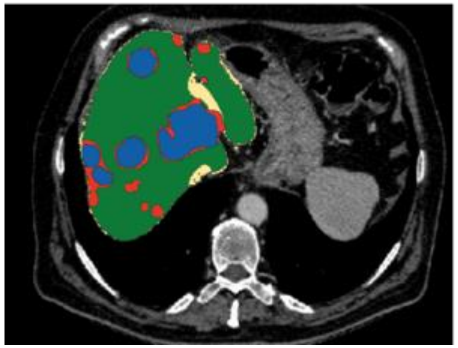

Standard multi-class: soft-max with cross-entropy loss

- exclusive classes
- no topological awareness

Logistic regression for hierarchically-nested classes

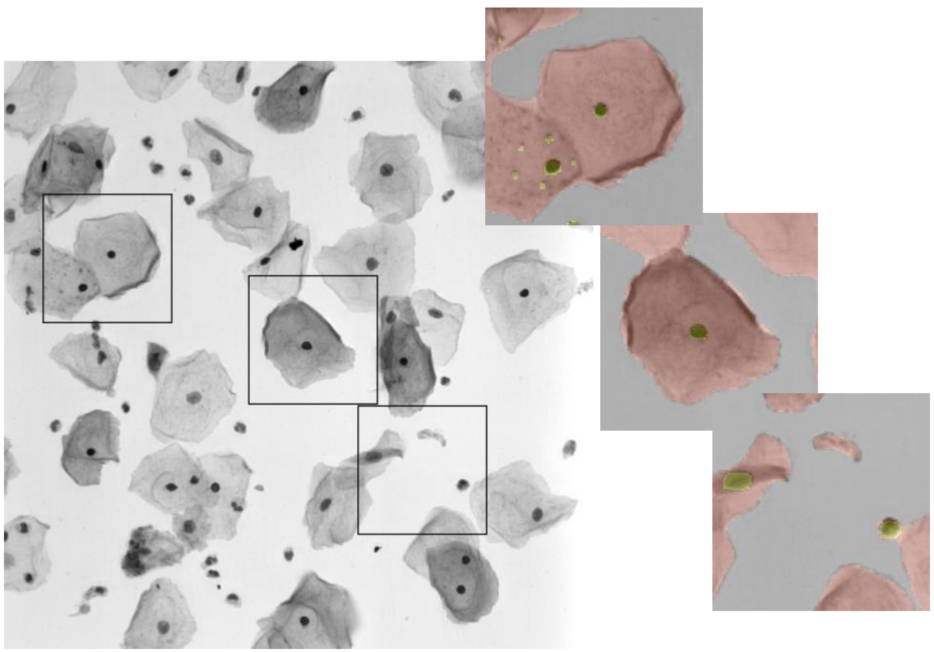

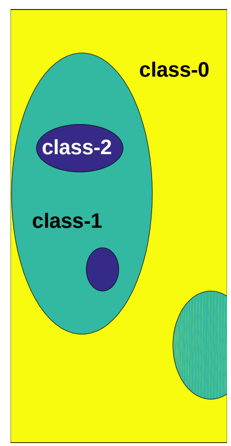

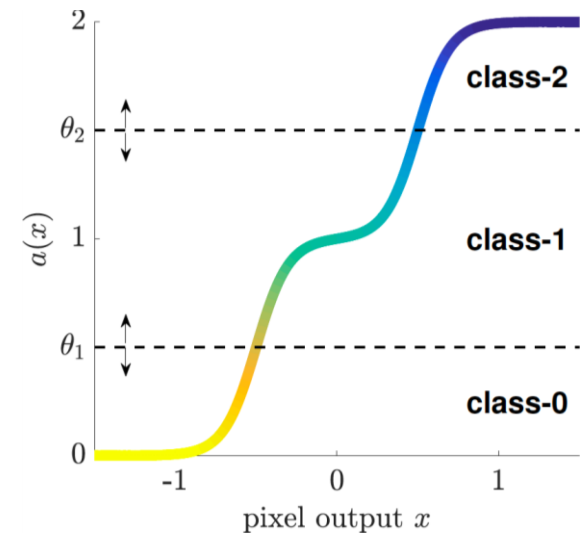

KONICA MINOLTA

INCLUDING PRIOR KNOWLEDGE

Technische Universität München

Achieving Faster and better segmentation

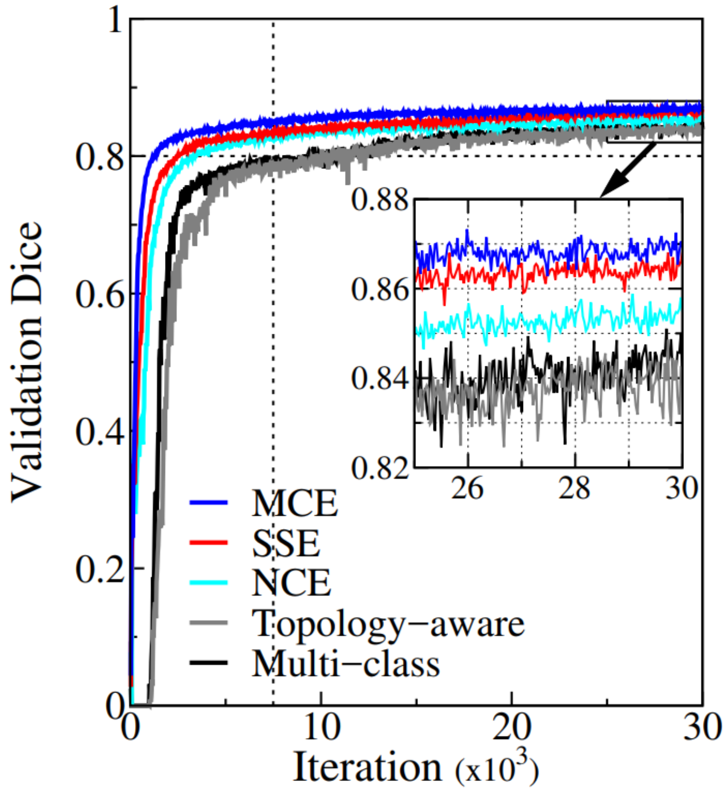

Kaggle challenge: only 16 images(!)
→ heavy online data augmentation
(flips, warping, rotations, translations, rescaling)

Method	Test Dice scores
Multi-class	0.832 (0.058)
NCE	0.844 (0.061)
SSE	**0.859 (0.052)**
MCE	**0.868 (0.082)**

M. Piraud, A. Sekuboyina and B. H. Menze, *Multi-level Activation for Segmentation of Hierarchically-nested Classes*, BioComputing@ECCV 2018

DEALING WITH VOLUMETRIC DATA

3D cross-hair filters

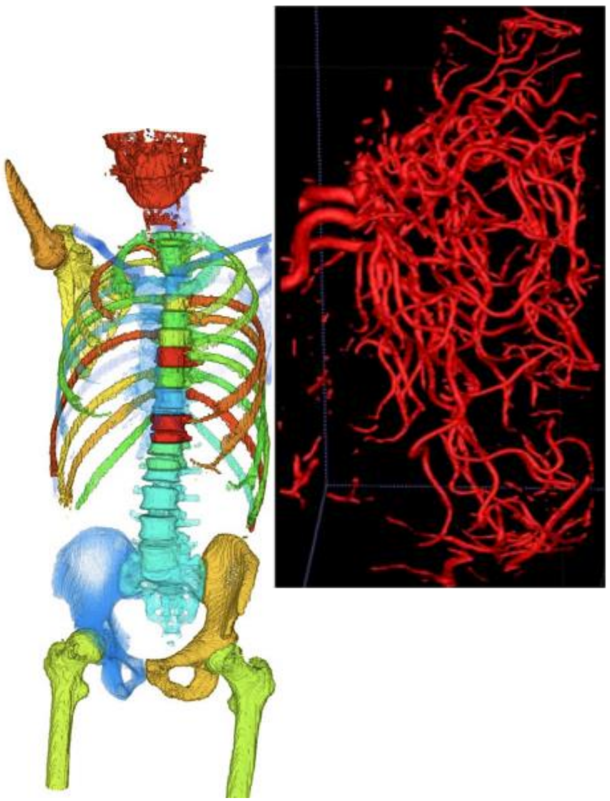

M. Bieth *etal* IEEE Trans. Med. Imag. 2017

DEALING WITH VOLUMETRIC DATA

Technische Universität München

3D cross-hair filters

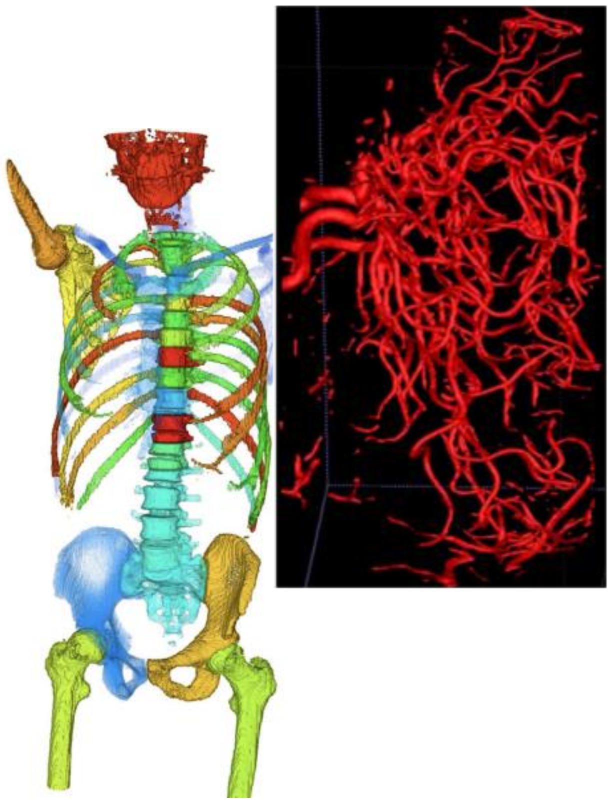

M. Bieth *etal* IEEE Trans. Med. Imag. 2017

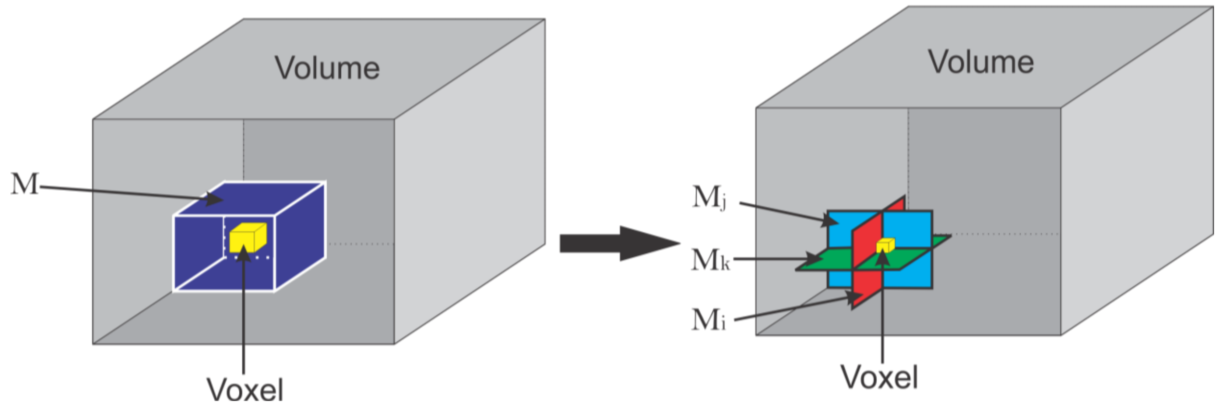

No. of param. → k^3

No. of param. → $3k^2$

$3k^2 \leq k^3$ if $k \geq 3$

DEALING WITH VOLUMETRIC DATA

Technische Universität München

Gain in memory and speed

V-Net: Milletari et al. 3DV 2016

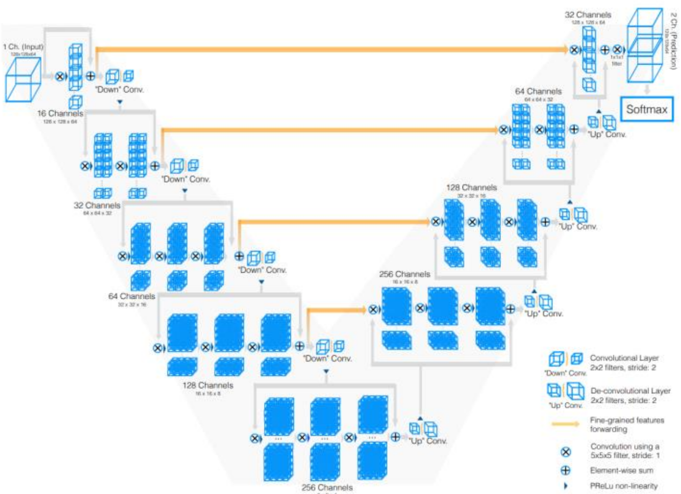

40% gain in memory

Tetteh et. al, M. Piraud, and B. H. Menze, *deepVesselNet: Vessel Segmentation, Centerline Prediction and Bifurcation Detection in Magnetic Resonance Angiography*, arXiv:1803.09340

KONICA MINOLTA

DEALING WITH VOLUMETRIC DATA

Technische Universität München

Gain in memory and speed

V-Net: Milletari et al. 3DV 2016

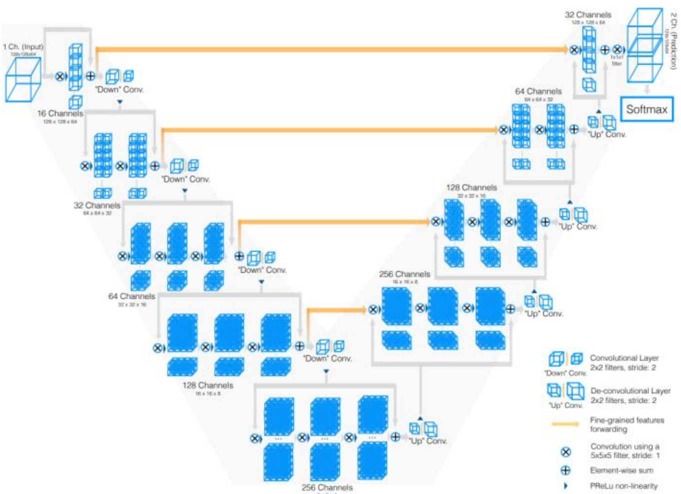

Vessel segmentation tasks (DeepVesselNet)

Dataset	Methods	F1 Score	Execution time
Synthetic (600x304x325)	CF filters	0.9956	**17s**
	3D filters	0.9949	23s
MRA (fine-tuning) (580x640x136)	CF filters	0.8425	**20s**
	3D filters	0.8497	26s
SRXTM (transfer learning) (254x254x254)	CF filters	0.9601	**7s**
	3D filters	0.9555	11s

40% gain in memory

23% gain in computation speed.

Tetteh et. al, M. Piraud, and B. H. Menze, *deepVesselNet: Vessel Segmentation, Centerline Prediction and Bifurcation Detection in Magnetic Resonance Angiography*, arXiv:1803.09340

KONICA MINOLTA

DECISION VISUALISATION

Feature inversion

Reconstruct image from features of a specific layer

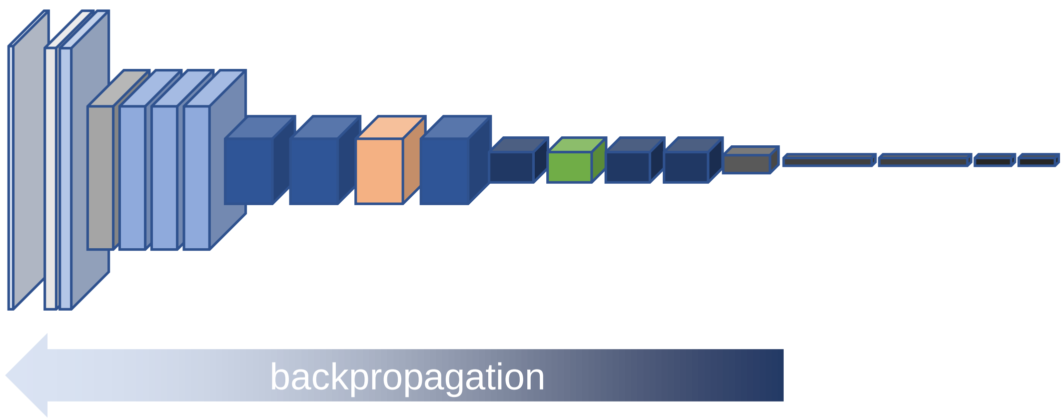

M. Baust *etal*, *Understanding regularization to visualize CNNs*, arXiv:1805.00071

DECISION VISUALISATION

Feature inversion

Reconstruct image from featu
specific layer

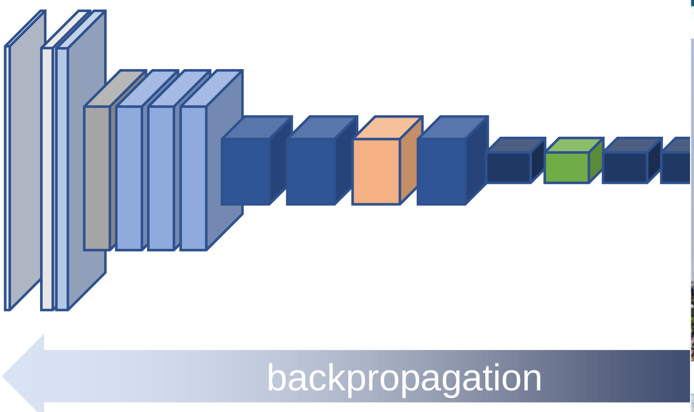

Network 1 Network 2

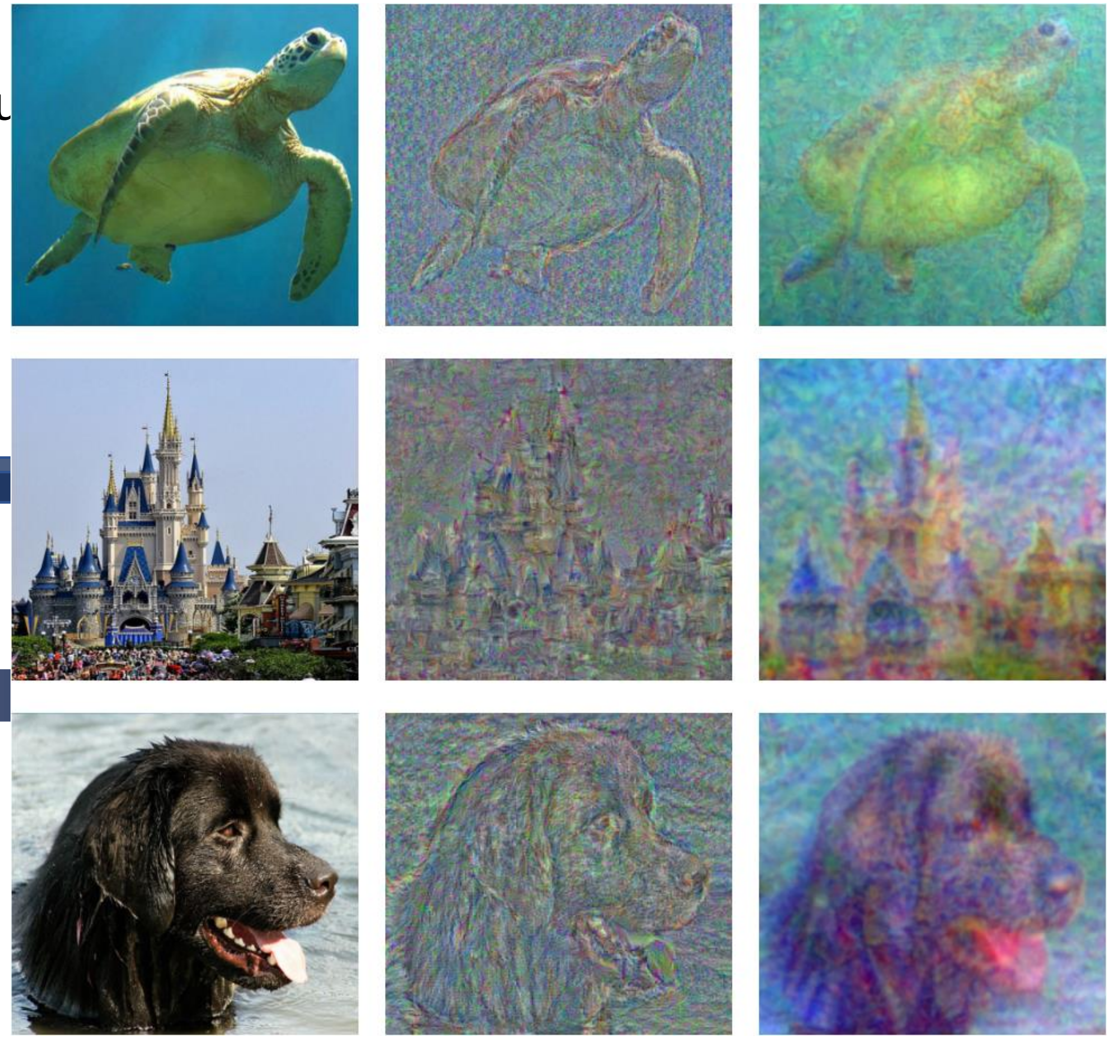

M. Baust *etal*, *Understanding regularization to visualize CNNs*, arXiv:1805.00071

DECISION VISUALISATION

Class maximization

Reconstruct ‘archetypical’ image which maximizes a specific class

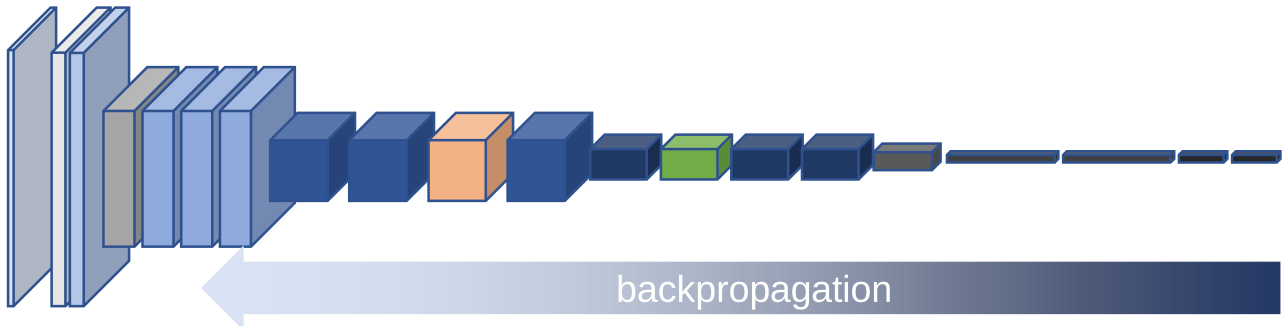

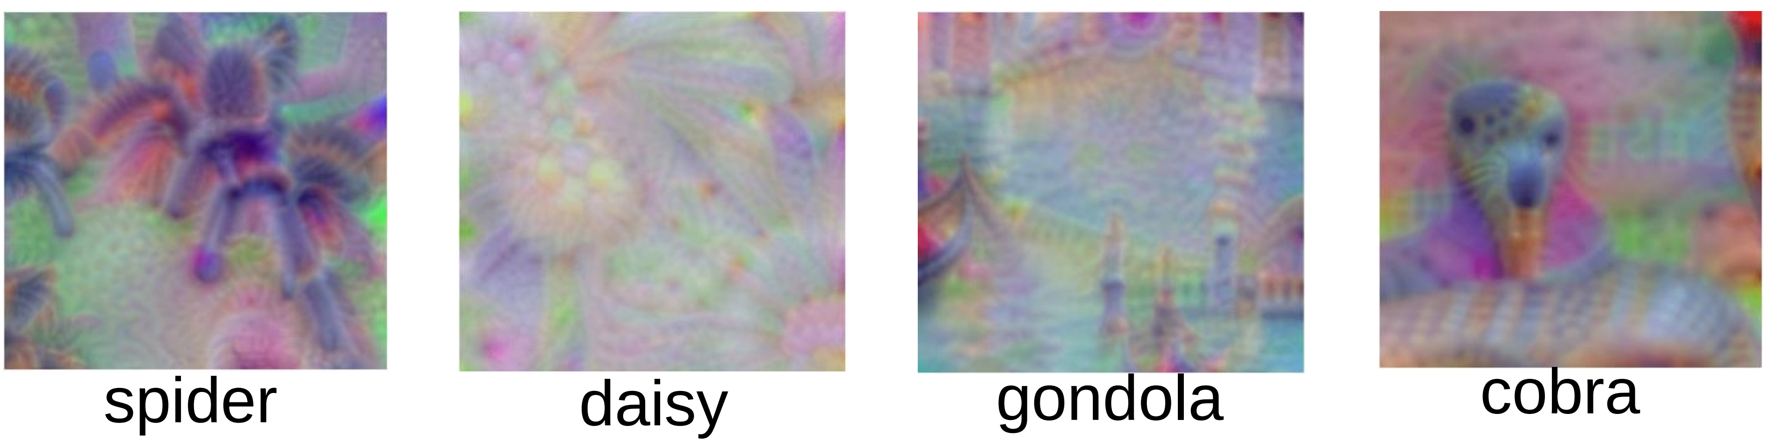

https://research.konicaminolta.eu/understanding-deep-convolutional-neural-networks-through-visualization/

KONICA MINOLTA

DECISION VISUALISATION

Class maximization

Reconstruct ‘archetypical’ image which maximizes a specific class

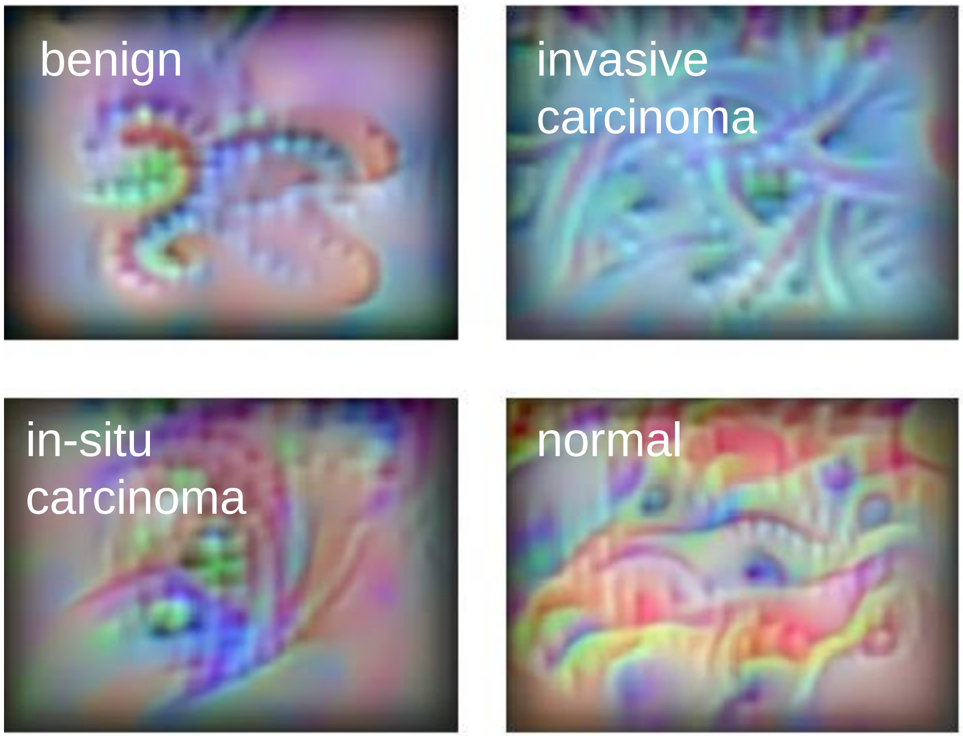

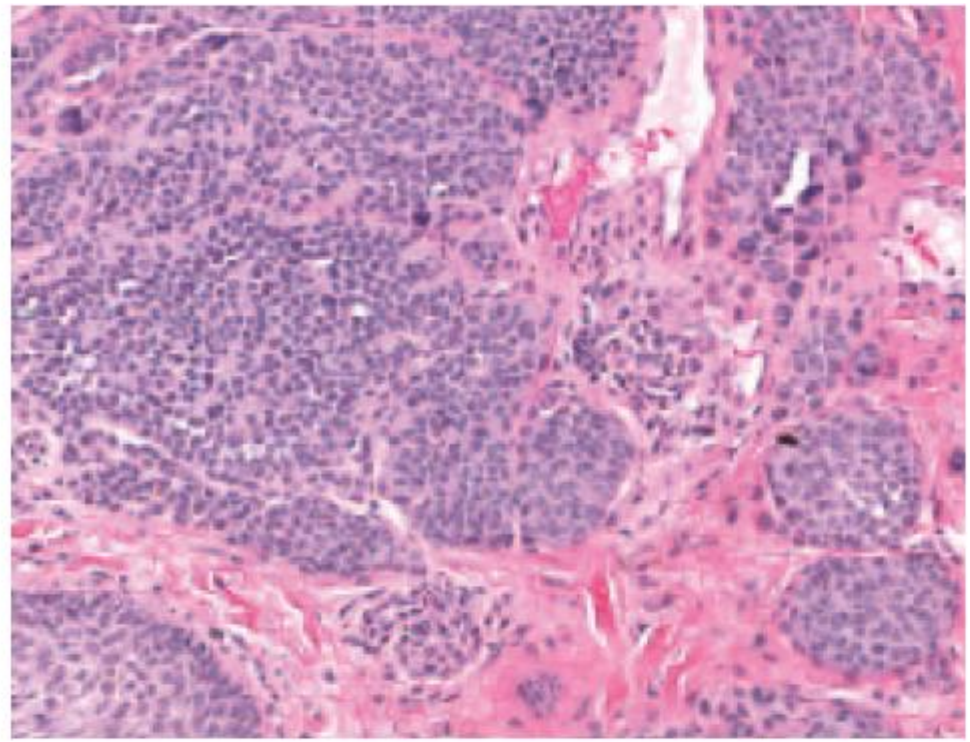

https://research.konicaminolta.eu/understanding-deep-convolutional-neural-networks-through-visualization/

CHALLENGES AND CURRENT LIMITATIONS

Summary

Data scarceness and variability

Semi-supervised learning/ Transfer Learning/ Use prior knowledge …

The 'black-box' problem

Some visualization methods… Very active research field!

Testing/ Certification for development in the clinic (15+ FDA-approved and 8+ CE softwares)

Currently: extensive 'heuristic' testing

ARTERYS

www.arterys.com

IDx

www.eyediagnosis.net

See e.g.: http://www.technologyreview.com/s/604271/deep-learning-is-a-black-box-but-health-care-wont-mind/

WHAT DEEP LEARNING IS BRINGING TO HEALTHCARE

− medical errors / + reproducibility

+ quantitative and personalized medicine

WHAT DEEP LEARNING IS BRINGING TO HEALTHCARE

− medical errors / + reproducibility

+ quantitative and personalized medicine

Make-up for shortage of medical doctors?

More time for ‘human’ practice?

WHAT DEEP LEARNING IS BRINGING TO HEALTHCARE

− medical errors / + reproducibility

+ quantitative and personalized medicine

Make-up for shortage of medical doctors?

More time for ‘human’ practice?

So…

- Exciting?
- Scary?

KONICA MINOLTA