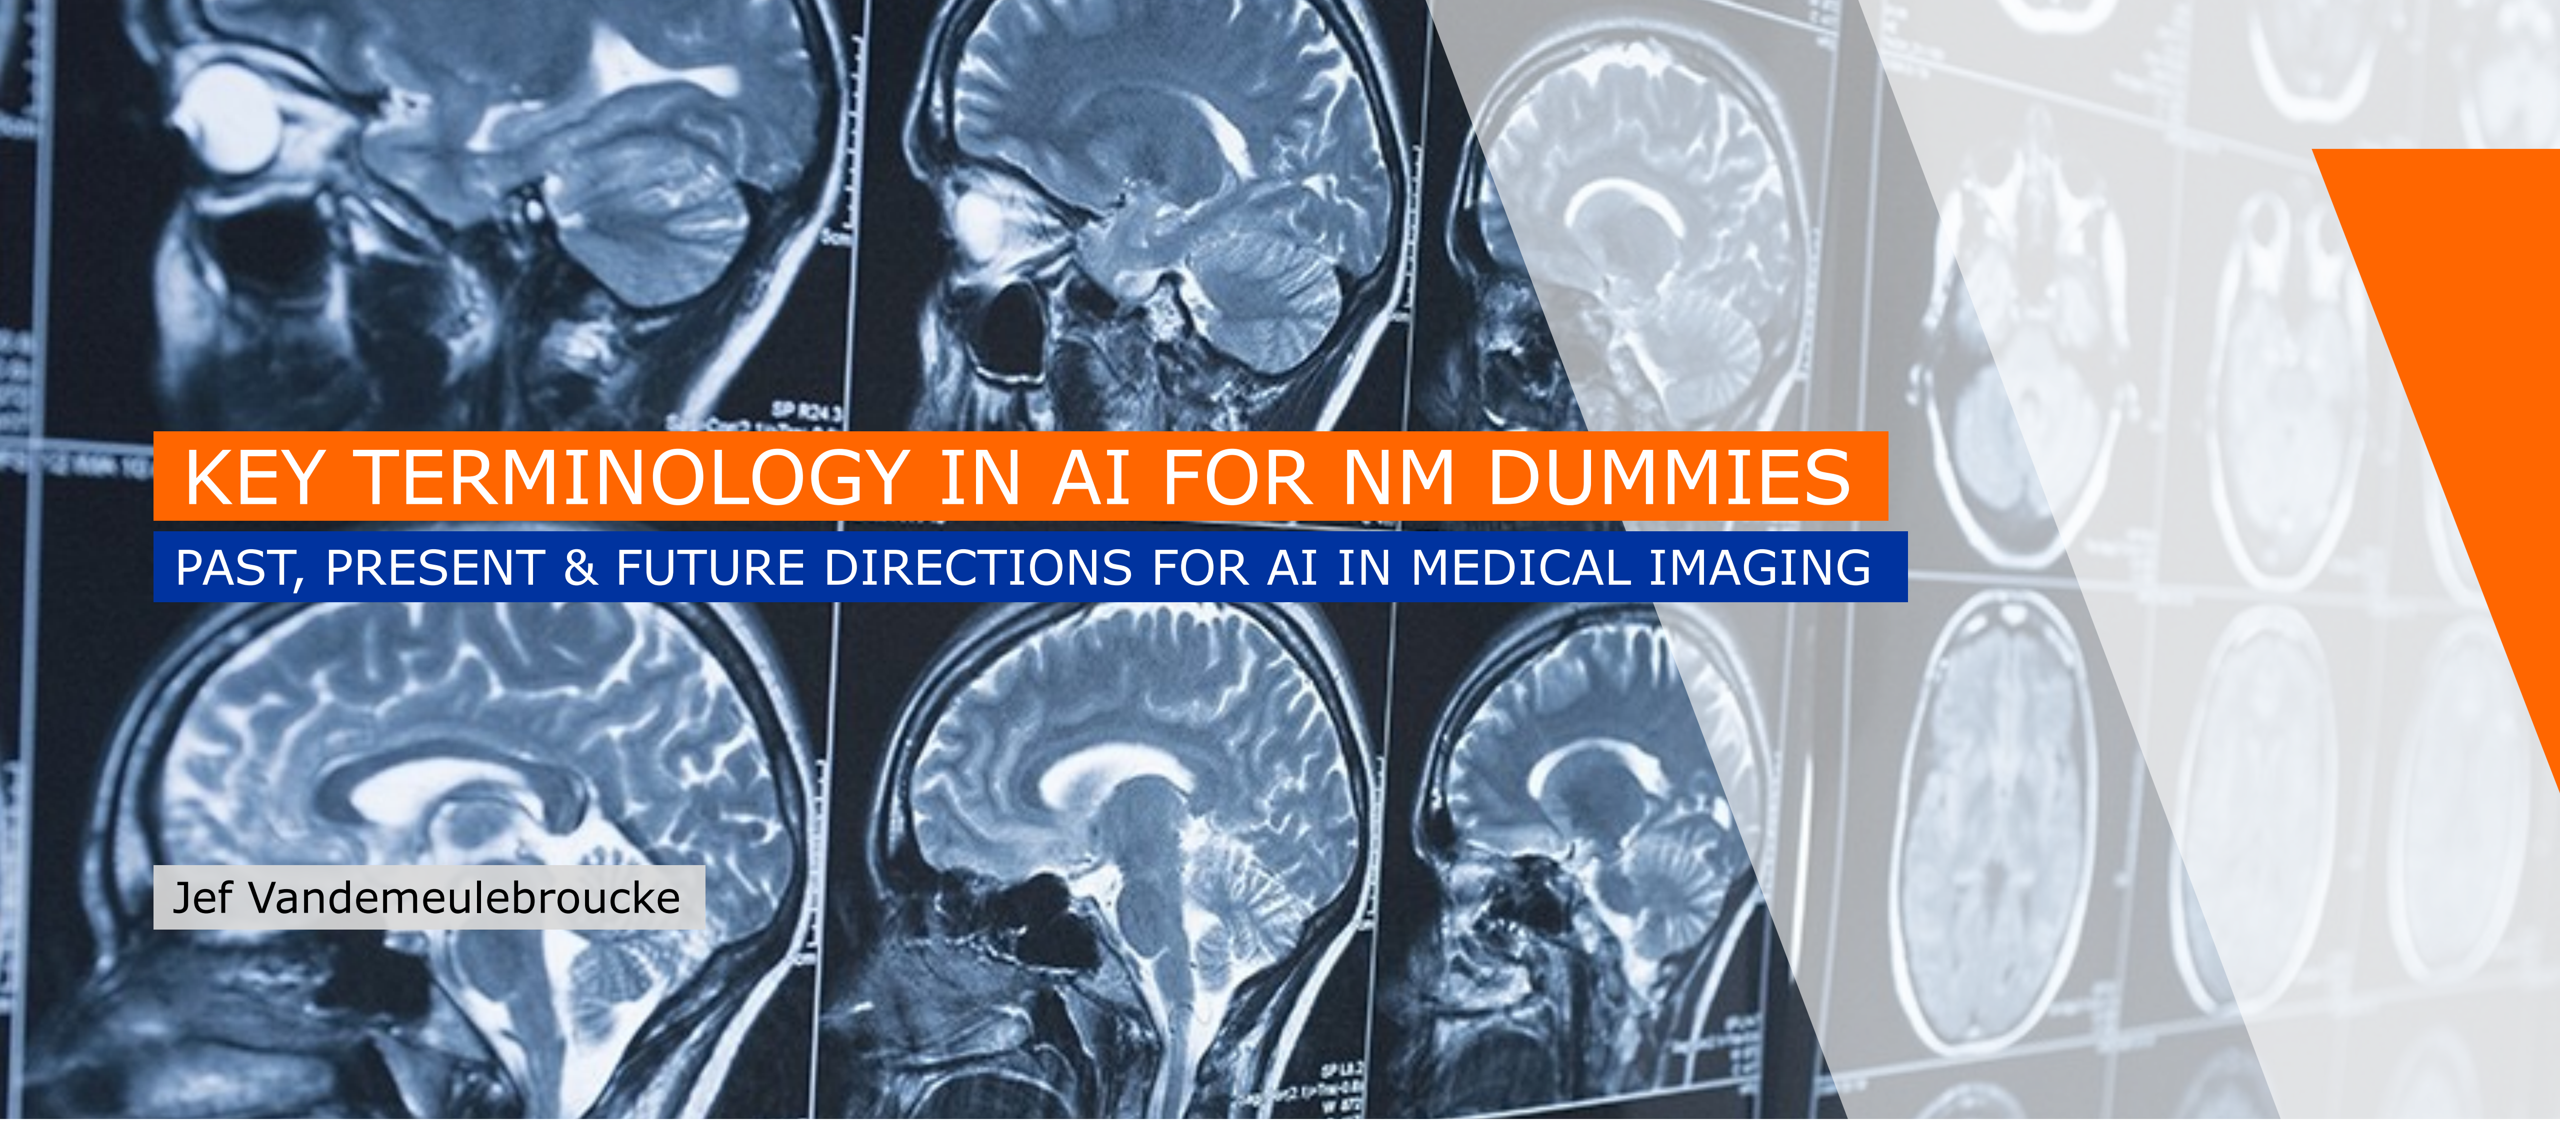

KEY TERMINOLOGY IN AI FOR NM DUMMIES

PAST, PRESENT & FUTURE DIRECTIONS FOR AI IN MEDICAL IMAGING

Jef Vandemeulebroucke

VUB VRIJE UNIVERSITEIT BRUSSEL
UZ Universitair Ziekenhuis Brussel
ETRO ELECTRONICS & INFORMATICS
imec
FARI Ai FOR THE COMMON GOOD INSTITUTE BRUSSELS

MEDICAL IMAGE ANALYSIS GROUP AT VUB

COMPUTER-AIDED DIAGNOSIS & THERAPY

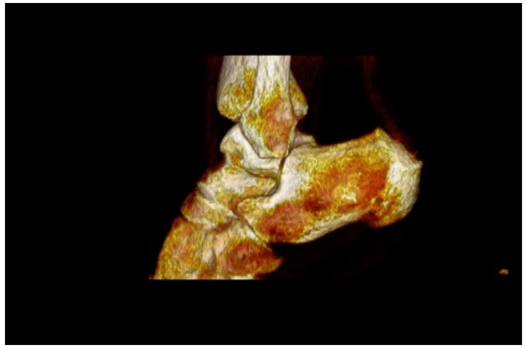

Motion analysis of joints

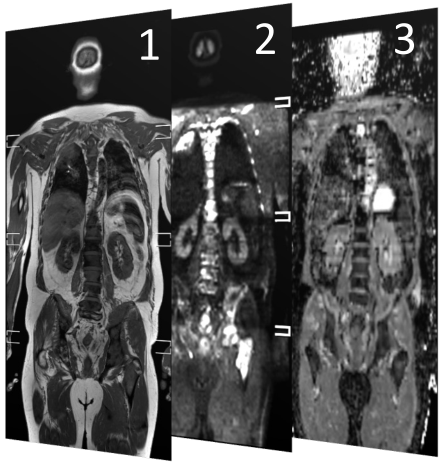

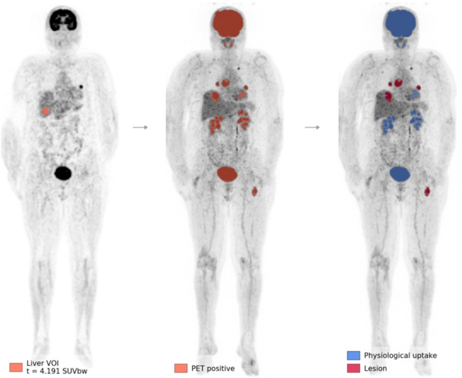

Tumour segmentation from whole-body imaging

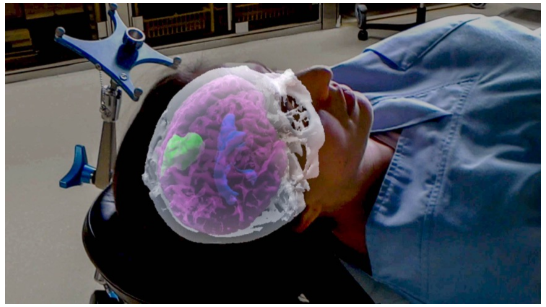

AR for surgical navigation

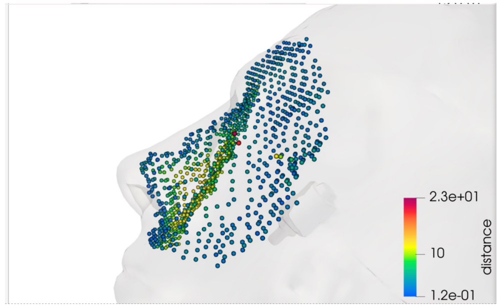

Intraoperative measurements

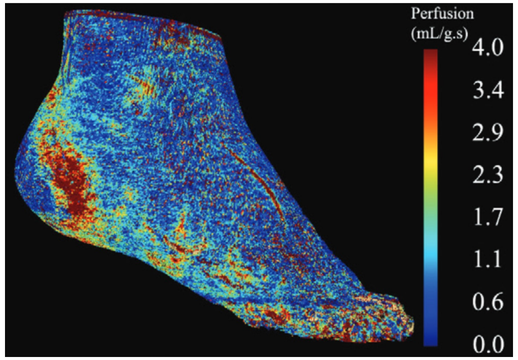

4D perfusion and angiography

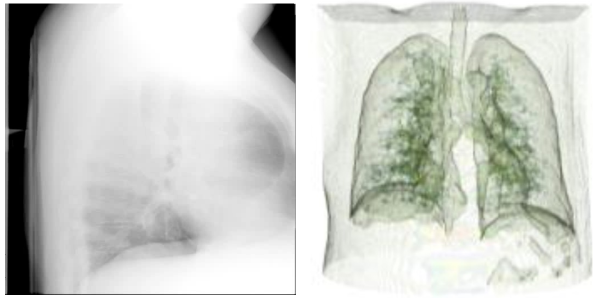

Sparse reconstruction for guidance

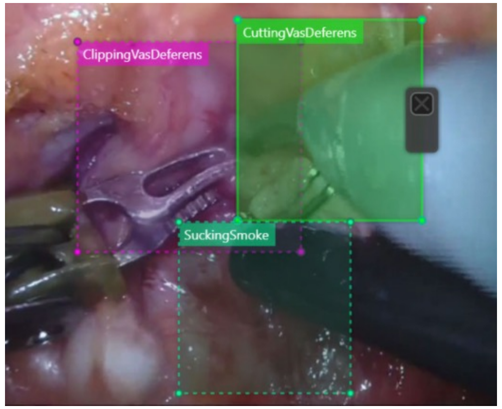

Surgical skill assessment

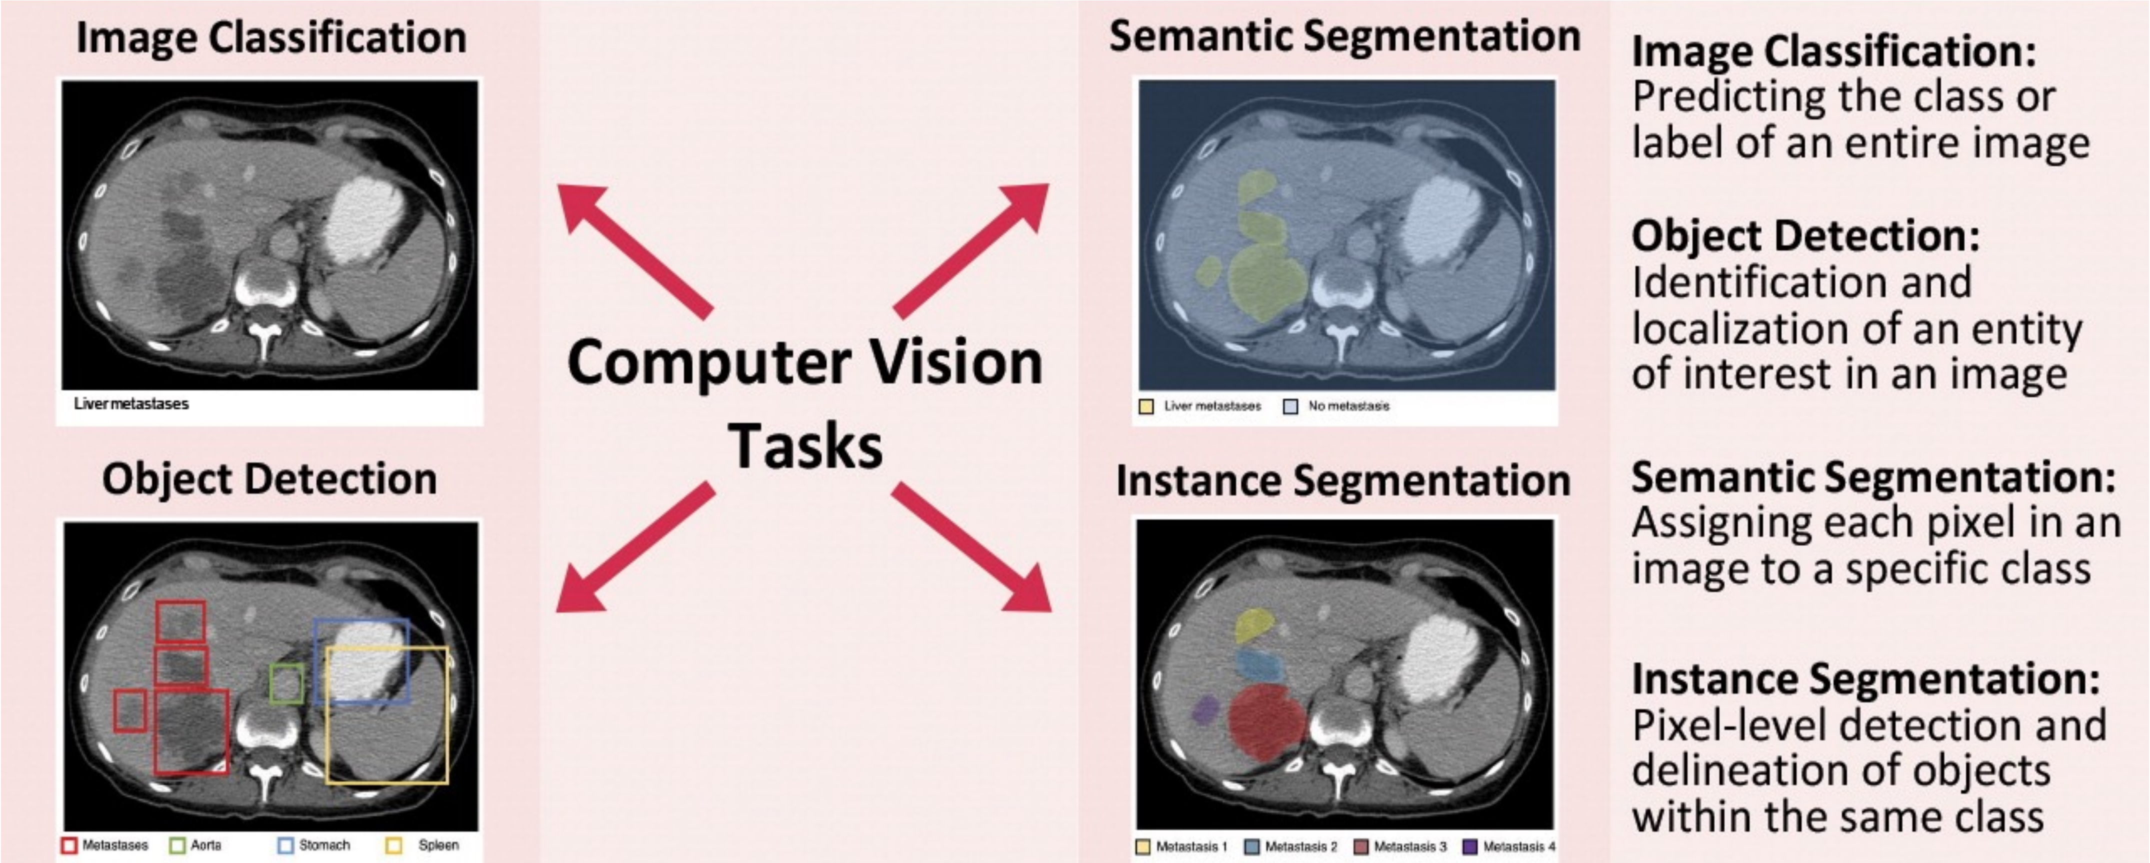

Deep Learning: An Update for Radiologists.
Cheng *et al.*, RadioGraphics 2021

TYPICAL TASKS

COMPUTER VISION

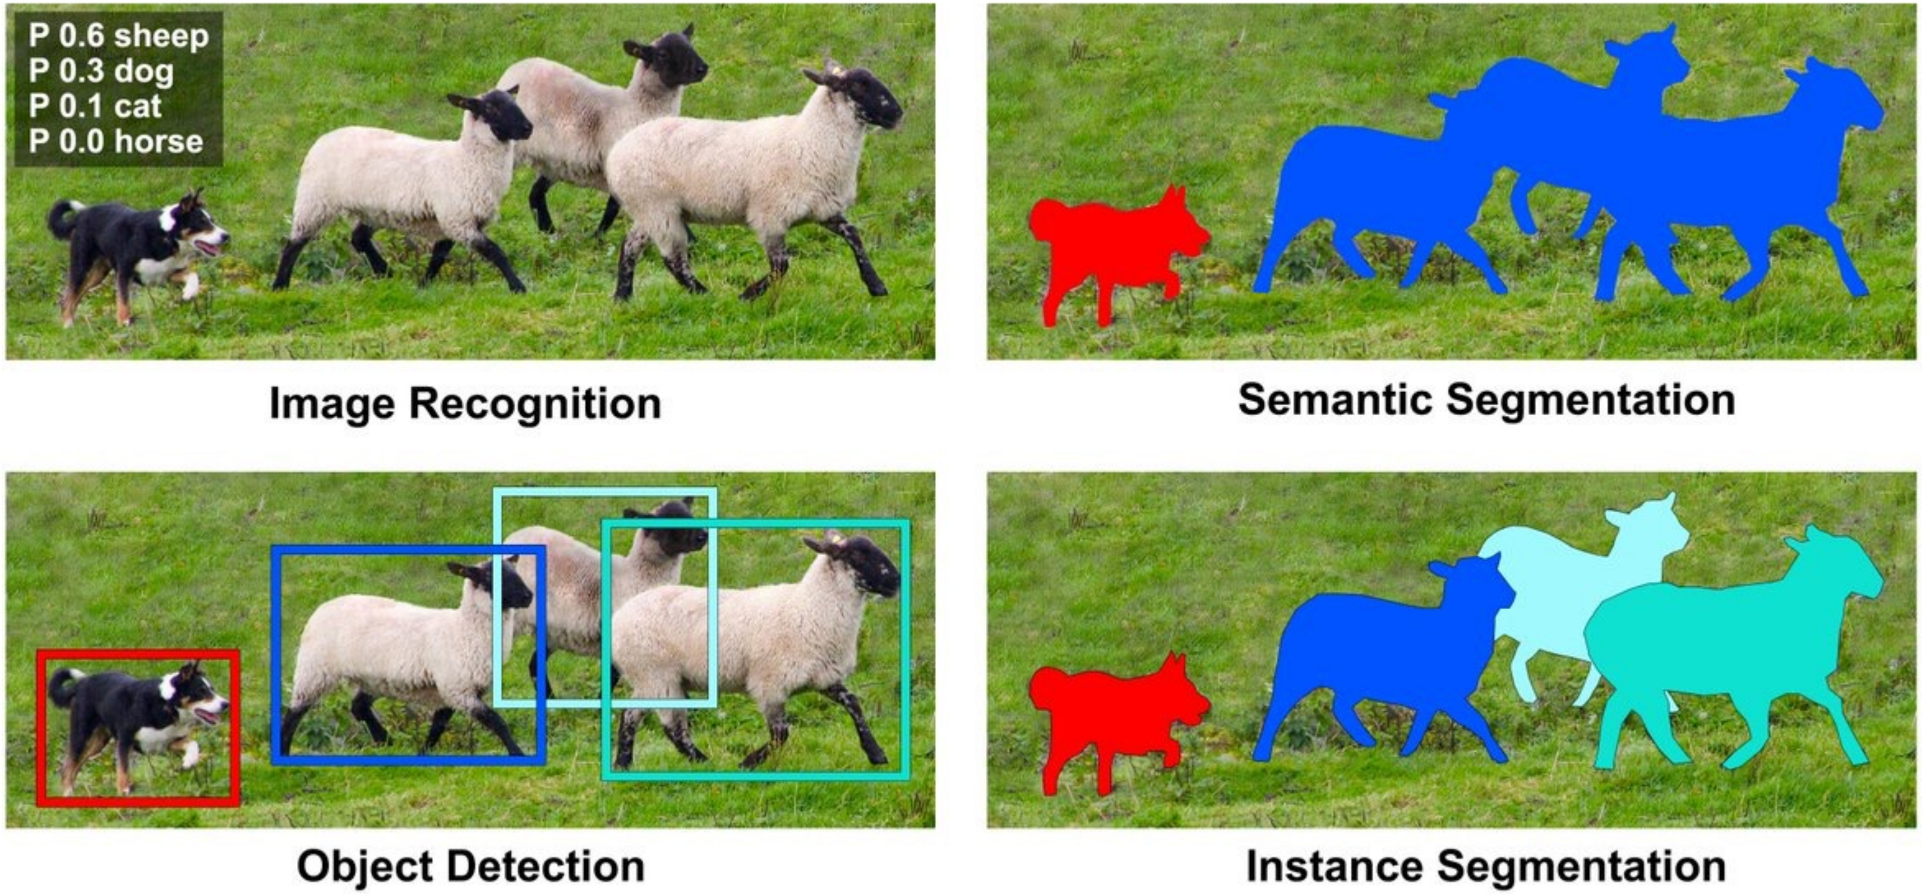

OTHER TASKS
RELEVANT FOR MEDICAL IMAGE ANALYSIS

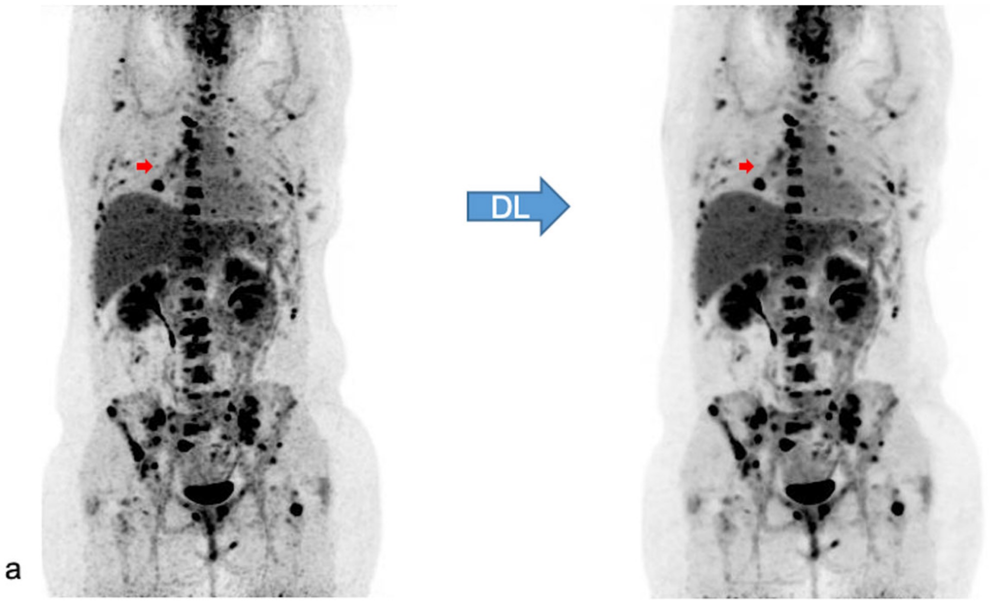

Denoising

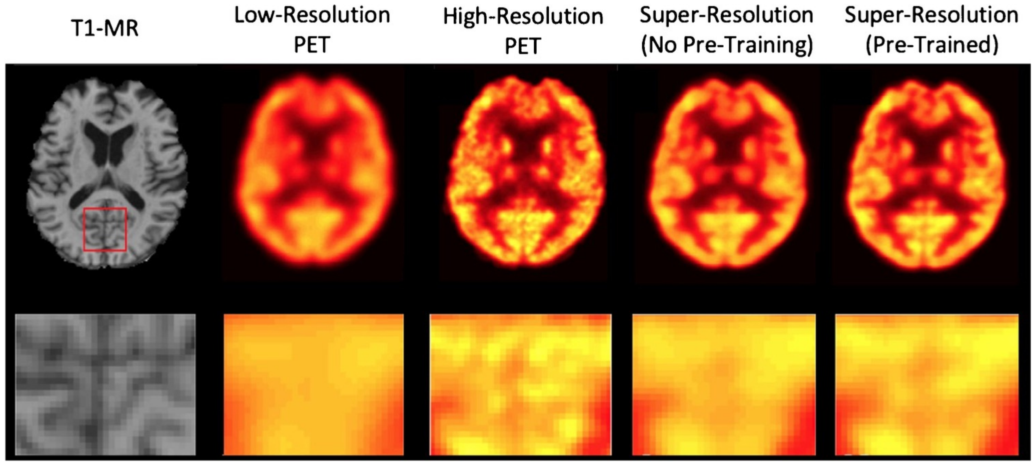

Super-Resolution

WHAT IS ARTIFICIAL INTELLIGENCE?

AI, ML & DL

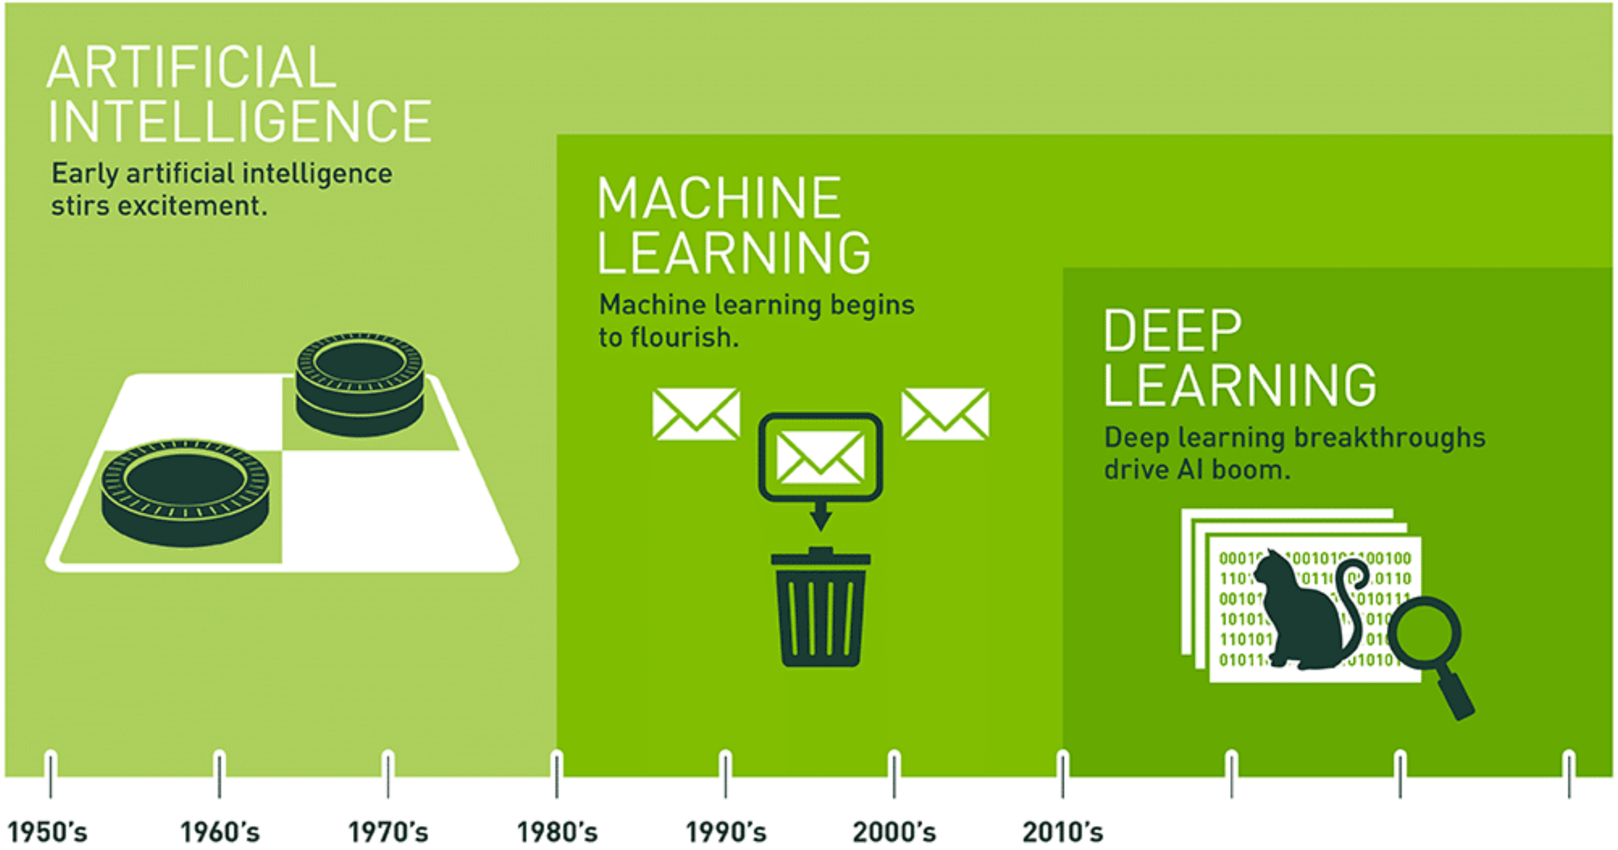

Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

MEDICAL IMAGE ANALYSIS HAS BEEN THROUGH A REVOLUTION

DEEP LEARNING IS THE NEW DOMINATING METHODOLOGY

Over the last 5-7 years

- Complete change of methods, workflow and challenges
- Huge increase in performance
- Huge increase in community

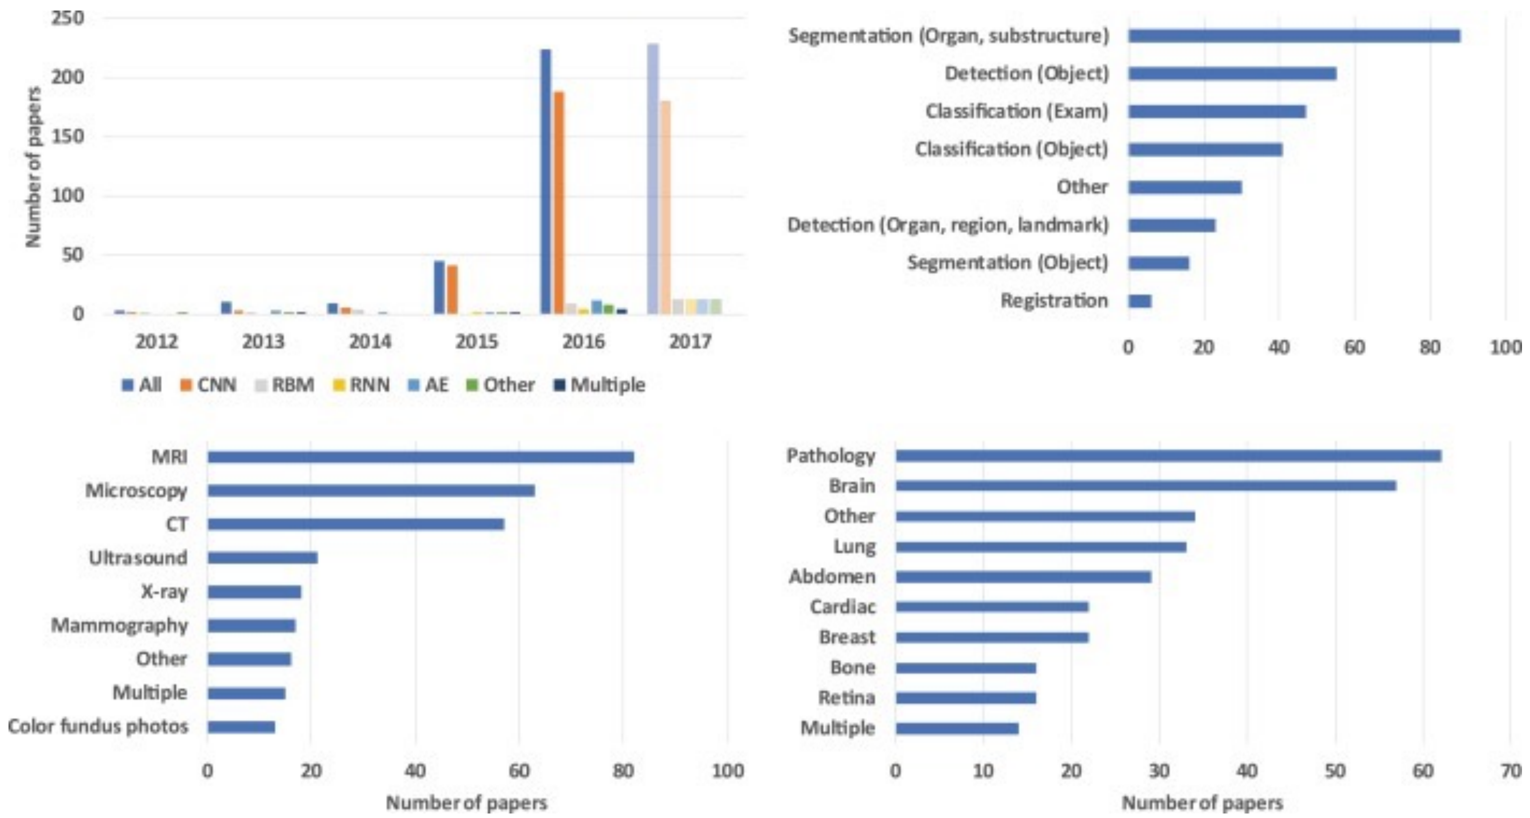

OVERVIEW

OUTLINE OF THIS TALK

A brief history of medical image analysis

- Typical tasks in medical image analysis
- From rule-based image processing, over ML to DL*

Where are we now?

- Current performance
- Challenges

Outlook

- The role of doctors in AI for medical imaging
- Novel innovations on the horizons

* Fifty years of computer analysis in chest imaging: rule-based, machine learning, deep learning.
van Ginneken, B., *Radiol Phys Technol,* 2017

BRAIN TUMOUR SEGMENTATION

RULE-BASED IMAGE PROCESSING ('80-'00)

A simplified example

- Segment the *skull*
- For all voxels inside the *skull*

	Features		Rules
If:	intensity on FLAIR	> X	AND
[If:	intensity on T2	> Y	OR
If:	intensity on T1c	> Z	]

- Then: add to *tumour*
- Else: add to *background*

- Fill holes in *tumour*

FLAIR T2 T1c

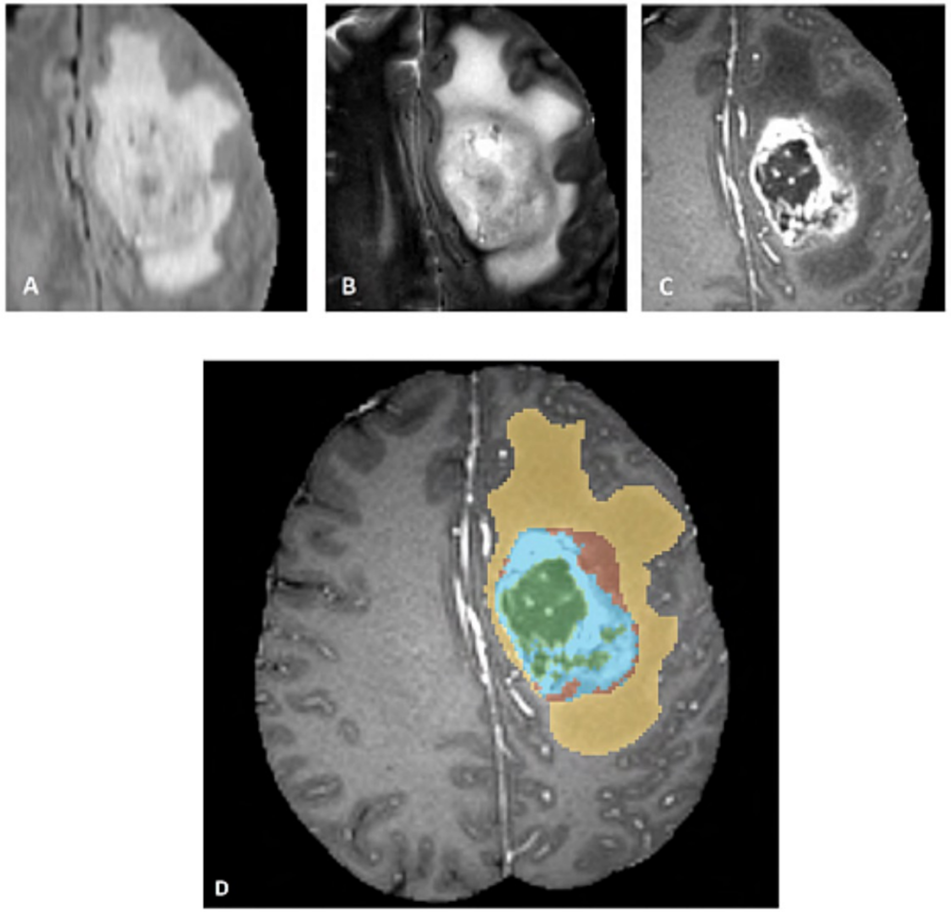

RULE-BASED IMAGE PROCESSING

PRO’S AND CON’S

Rule-based approaches leads to very intuitive, easy to understand processing

- **Typically achieved 60%-70% accuracy**
- Mimicking the MD’s reasoning

Complex and time-consuming to make

- Each problem requires another set of expert rules
- Difficult translate medical knowledge into mathematics & code
  - How to fix appropriate values for X, Y & Z?

	Features		Rules
If:	intensity on FLAIR	> X	AND
If:	intensity on T2	> Y	OR
If:	intensity on T1c	> Z	

BRAIN TUMOUR SEGMENTATION

MACHINE LEARNING USING HAND-CRAFTED FEATURES ('00-'10)

1) Extract large amount of features

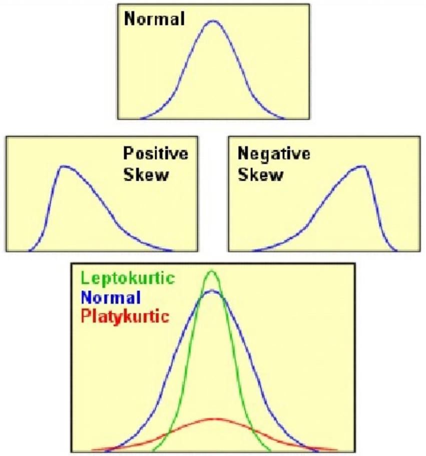

Intensity features

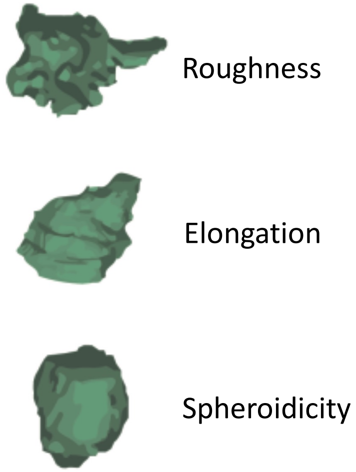

Shape features

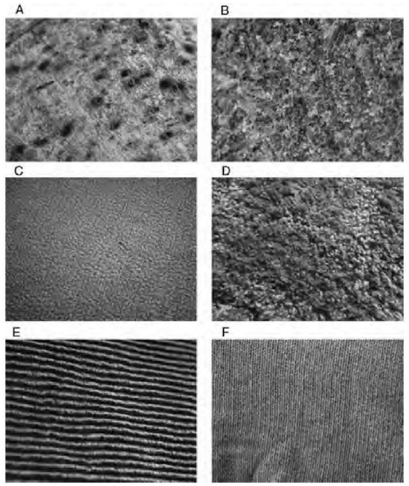

Texture features

2) Train a classifier on a set of annotated images

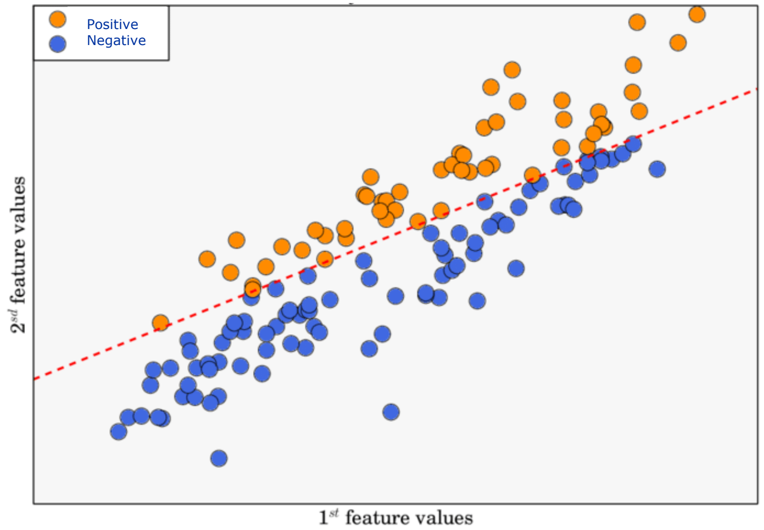

MACHINE LEARNING

TYPES OF LEARNING

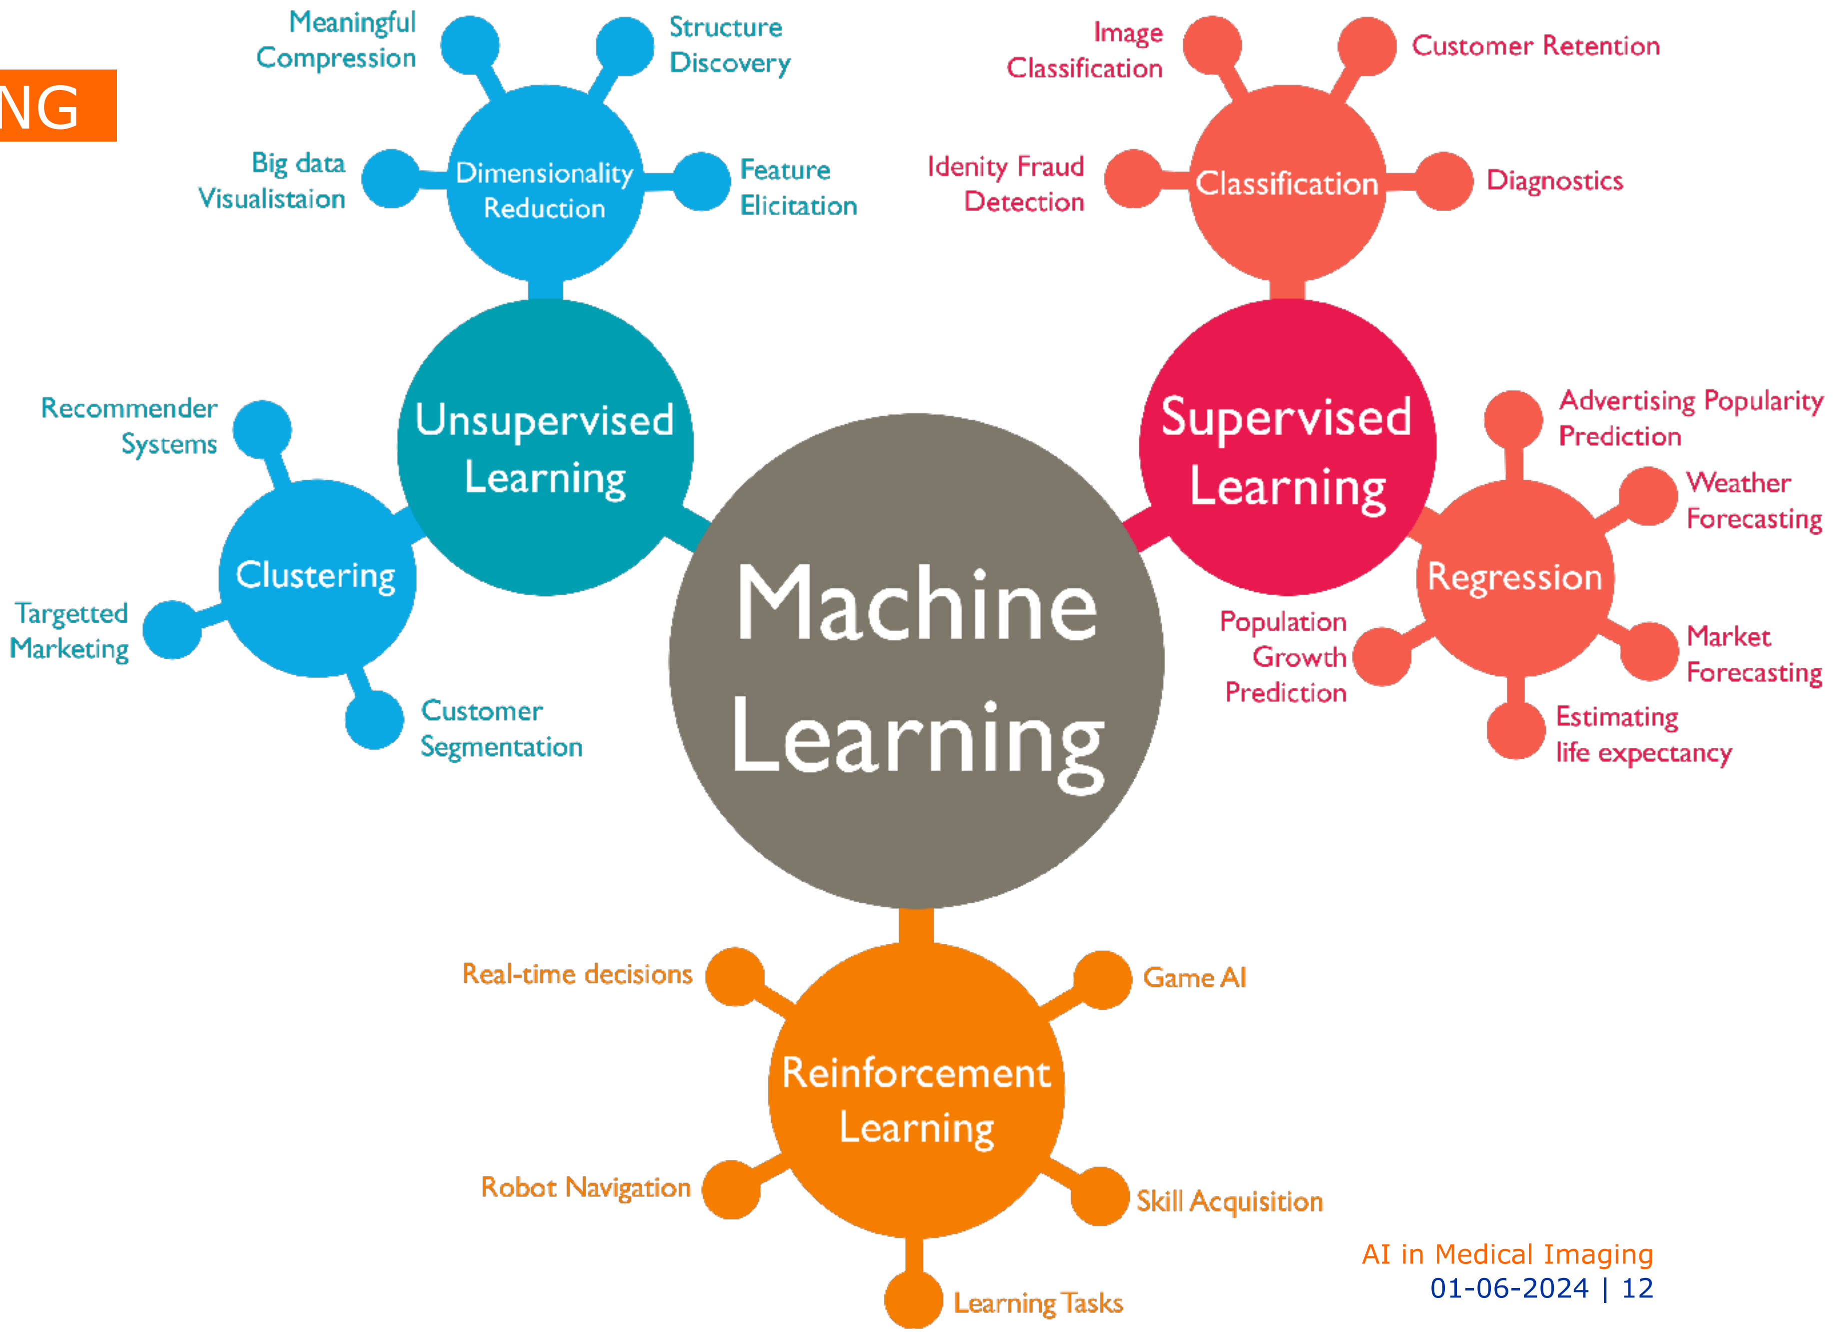

SUPERVISED LEARNING

LEARNING FROM EXAMPLES

Given

- Examples of a function
  - $(X, F(X))$

Predict

- $F(X)$ for new examples X
  - Discrete $F(X)$: Classification
  - Continuous $F(X)$: Regression
  - $F(X) = Probability(X)$: Probability estimation

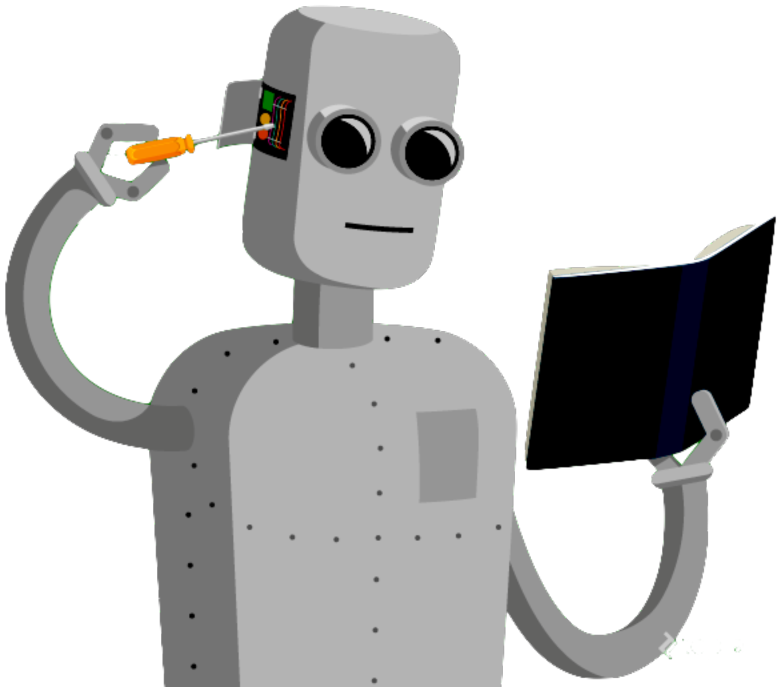

SUPERVISED LEARNING

BASIC PRINCIPLE

Build a mathematical model of sample data

- Known as "training data”

Use this model to make predictions or decisions on other, unseen data

- Testing data (or real-world data)

Predictions or decisions are made without being explicitly programmed to perform the task

- They are inferred from the data

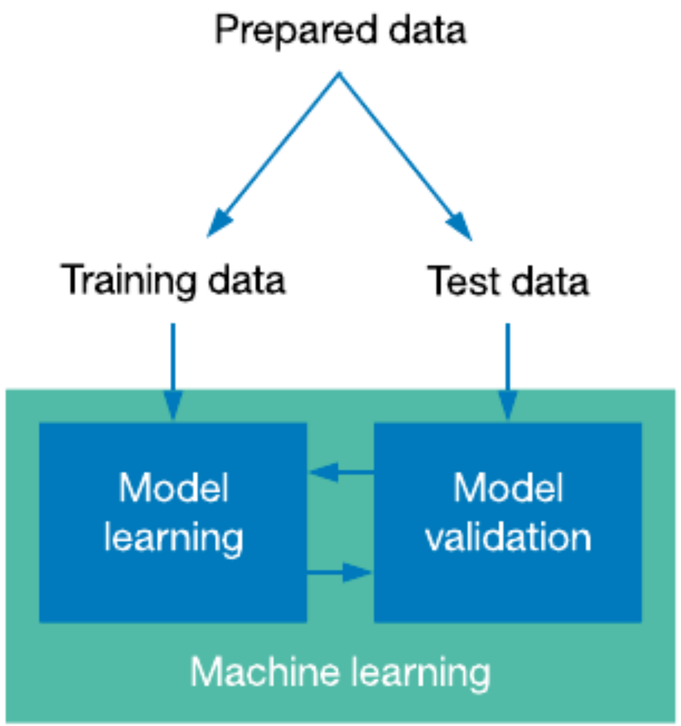

Validation set: tune the model hyperparameters

BRAIN TUMOUR SEGMENTATION

MACHINE LEARNING USING HAND-CRAFTED FEATURES (‘00-‘10)

Example: radiomics

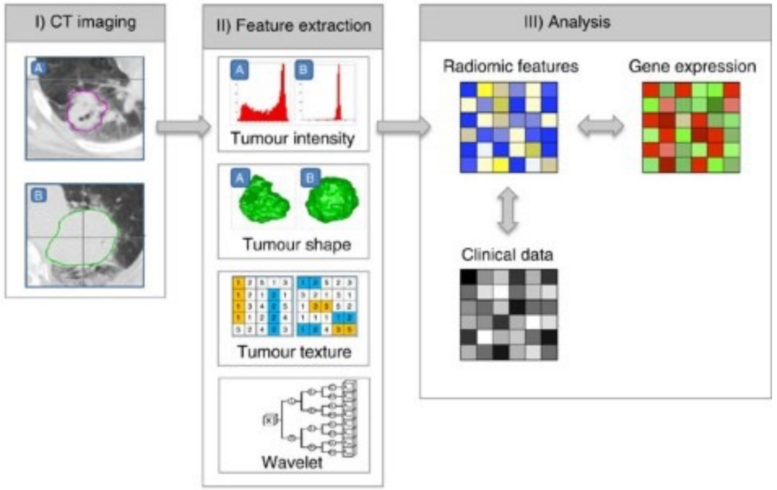

More powerful approach for medical image analysis

- **Achieves 75%-80% accuracy**
- ML allows to learn *how* features should be taken into account
  - You learn the rules from data

The “magic” is in the features

- Features are handcrafted for the considered task
- Designing your features is more of an art than a science

LARGE SCALE VISUAL RECOGNITION CHALLENGE

IMAGENET

Held for the first time in 2010

Data available for training:

- ~1M annotated images
- ~1k object classes

Task of the challenge

- Classify each image to one of the 1000 object classes

IMAGENET

PERFORMANCE OVER TIME

Rapid progression of performance

- First deep methods in 2012
- Models outperforming humans by 2015

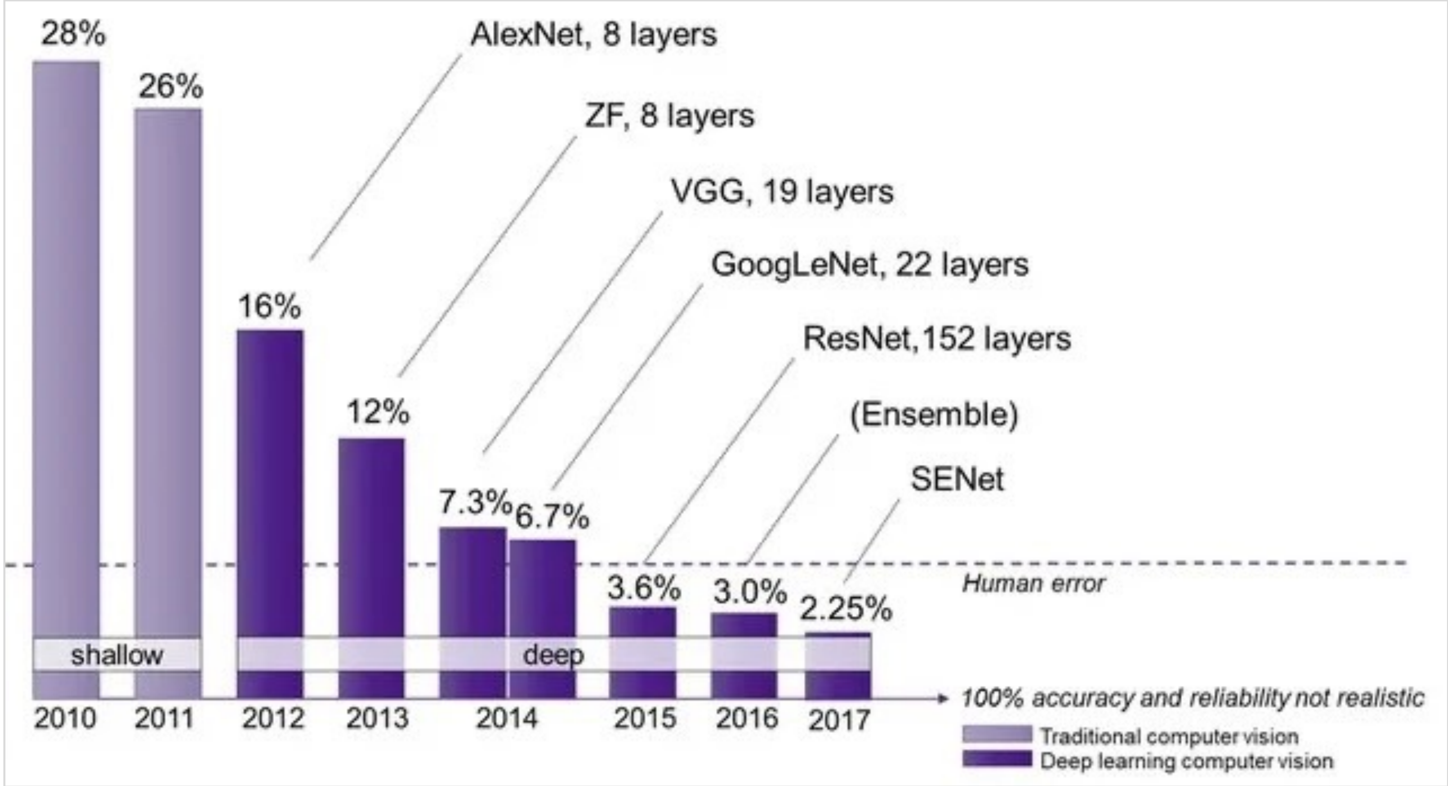

BRAIN TUMOUR SEGMENTATION
THE BRATS CHALLENGE

The medical domain soon followed

- Breakthrough results in 2013
- Whole tumour reaches over 90% Dice in current challenges

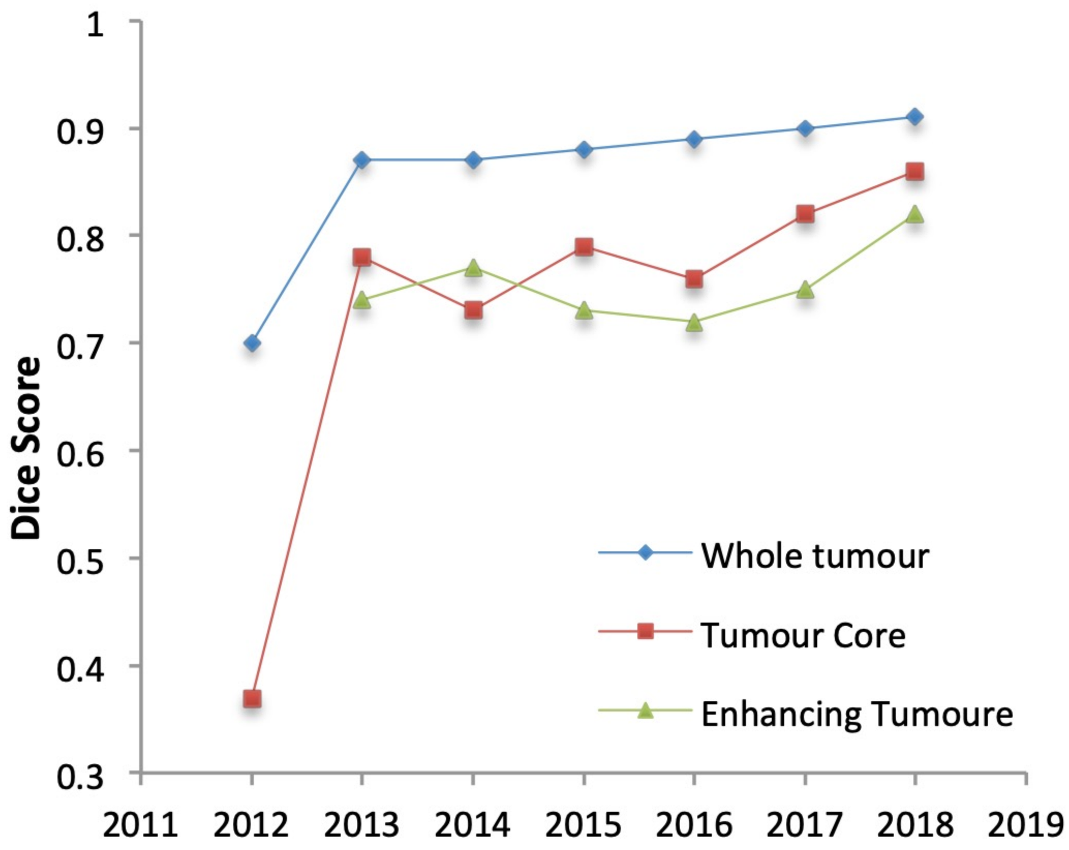

BRAIN TUMOR SEGMENTATION

DEEP LEARNING USING CNN (‘10-‘20)

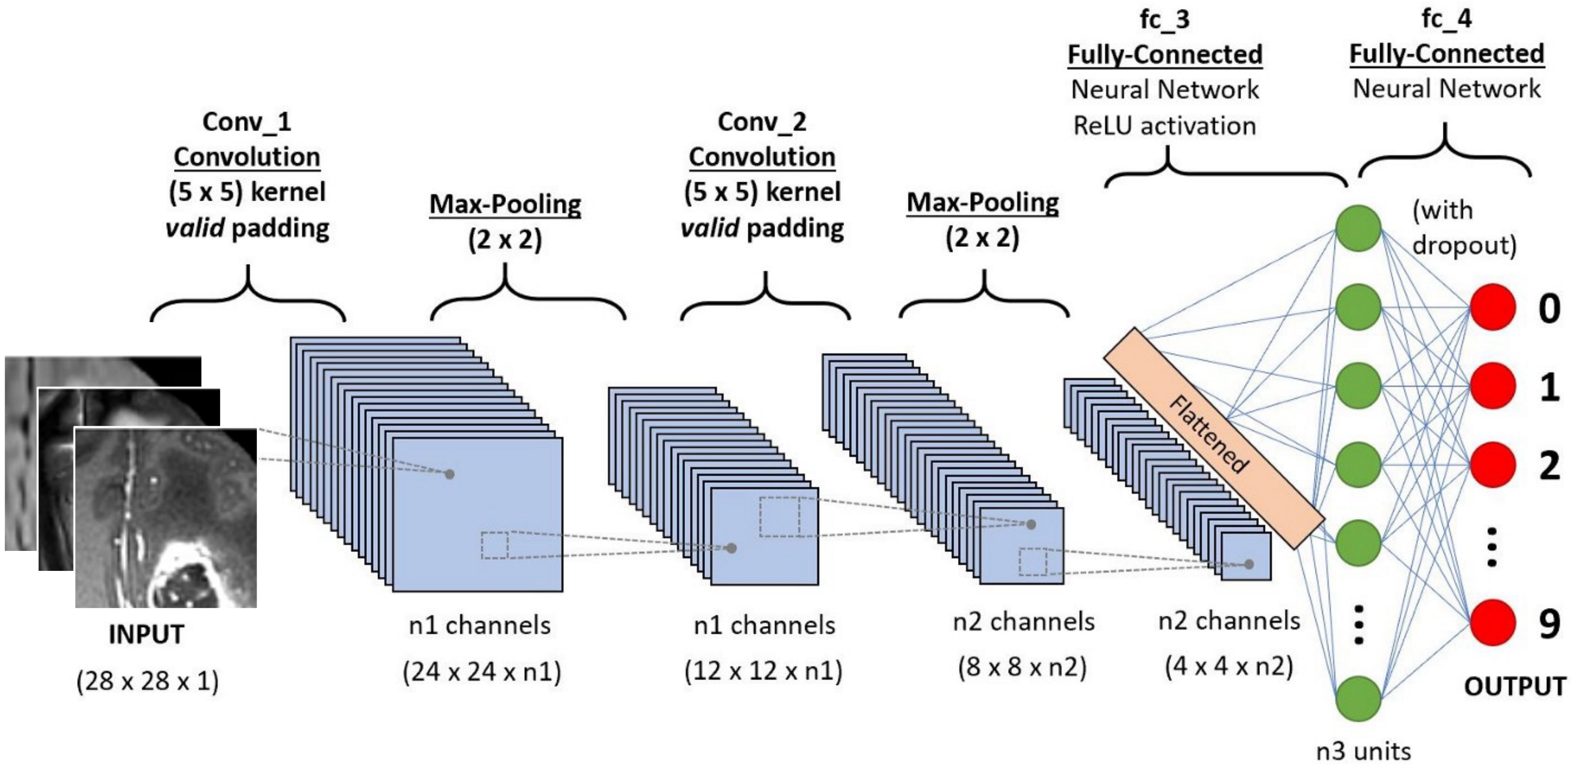

The input are (patches of) the images, often without preprocessing

Convolutions extract features

Neural network performs classification

During training, the network learns the features and the classifier, simultaneously

CONVOLUTION NEURAL NETWORKS

WHY CONVOLUTIONS?

Convolutional neural network (CNN)

- Neural network with some convolutional layers (and some other layers)
- A convolutional layer corresponds to applying a number of filters

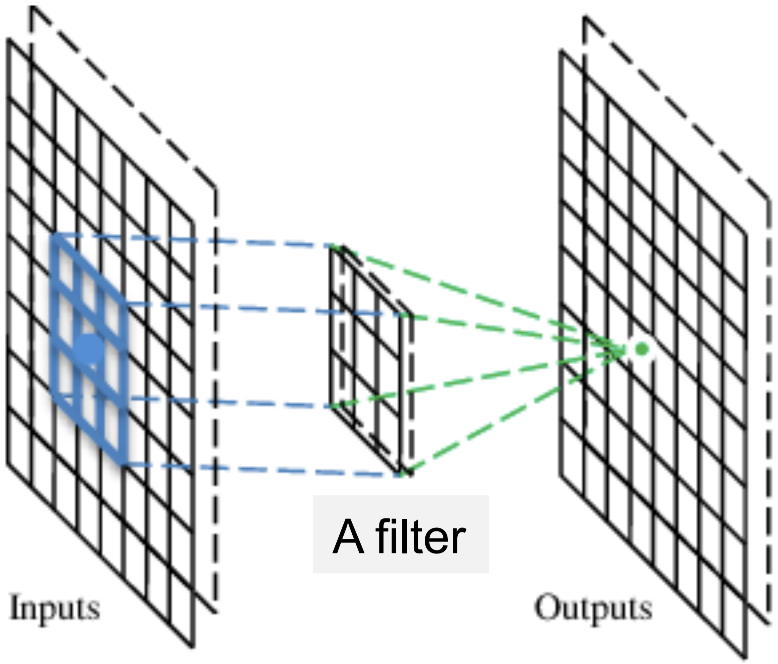

CONVOLUTION NEURAL NETWORKS

WHY (MAX) POOLING?

- Subsampling pixels will not change the overall appearance of the object
- Max pooling can be seen as a rough subsampling of the feature map
  - Dimensionality reduction

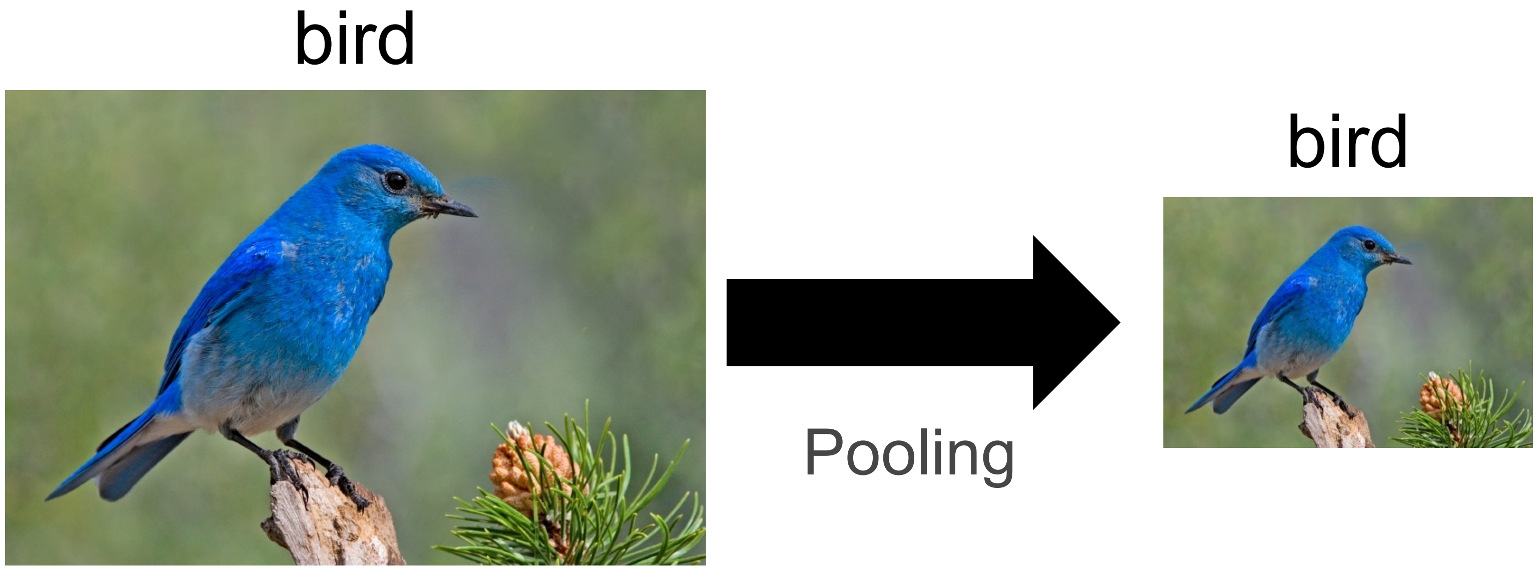

BRAIN TUMOR SEGMENTATION

DEEP LEARNING USING CNN (‘10-‘20)

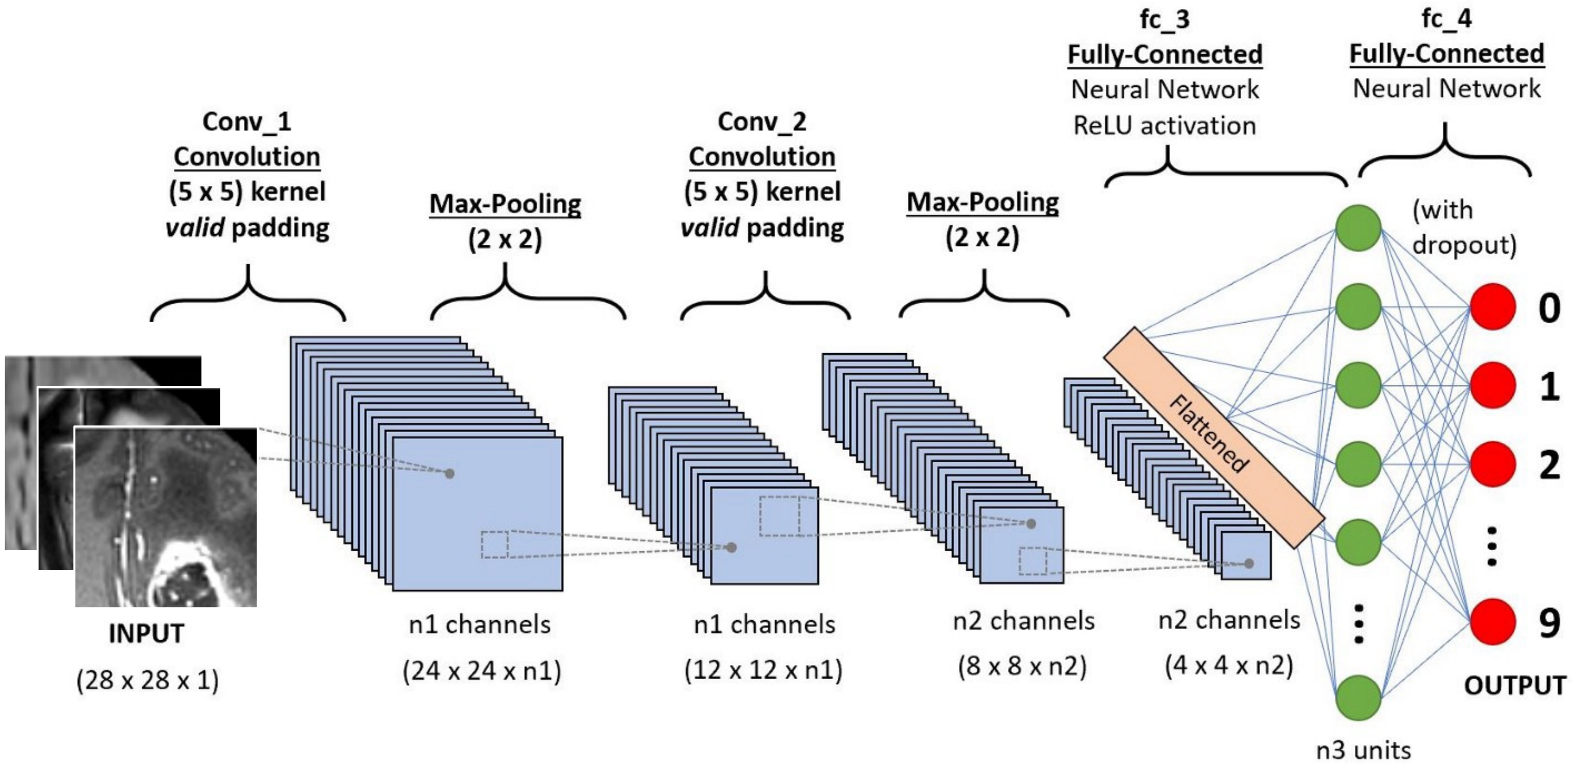

The input are (patches of) the images, often without preprocessing

Convolutions extract features

Neural network performs classification

During training, the network learns the features and the classifier, simultaneously

BRAIN TUMOUR SEGMENTATION

MACHINE LEARNING USING HAND-CRAFTED FEATURES (‘00-‘10)

Example: radiomics

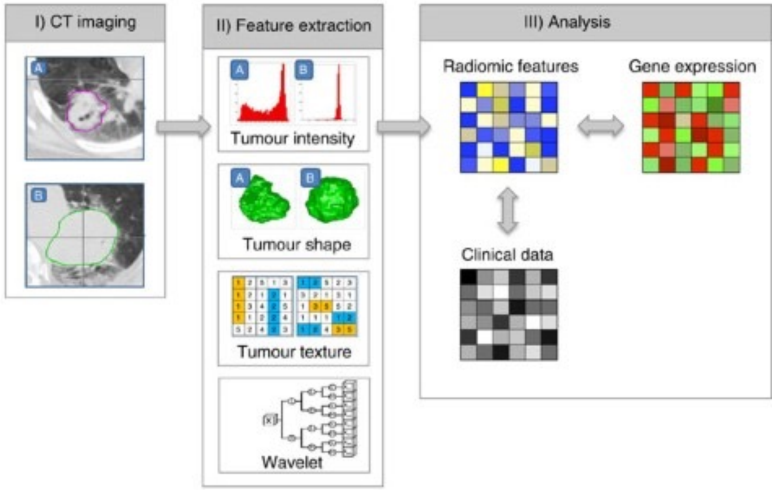

More powerful approach for medical image analysis

- **Achieves 75%-80% accuracy**
- ML allows to learn *how* features should be taken into account
  - You learn the rules from data

The “magic” is in the features

- Features are handcrafted for the considered task
- Designing your features is more of an art than a science

BRAIN TUMOUR SEGMENTATION

DEEP LEARNING USING CNN (‘10-‘20)

Breakthrough performance for medical image analysis

- **Achieves 90%-97% accuracy**
- Very generic method, same method applied to large variety of problems
    - Very little domain knowledge required for getting good results

Downsides

- As-is, offers little insight on “why is this tumour?”
- Training computationally very expensive, requiring specific hardware
- Requires large amounts of annotated data

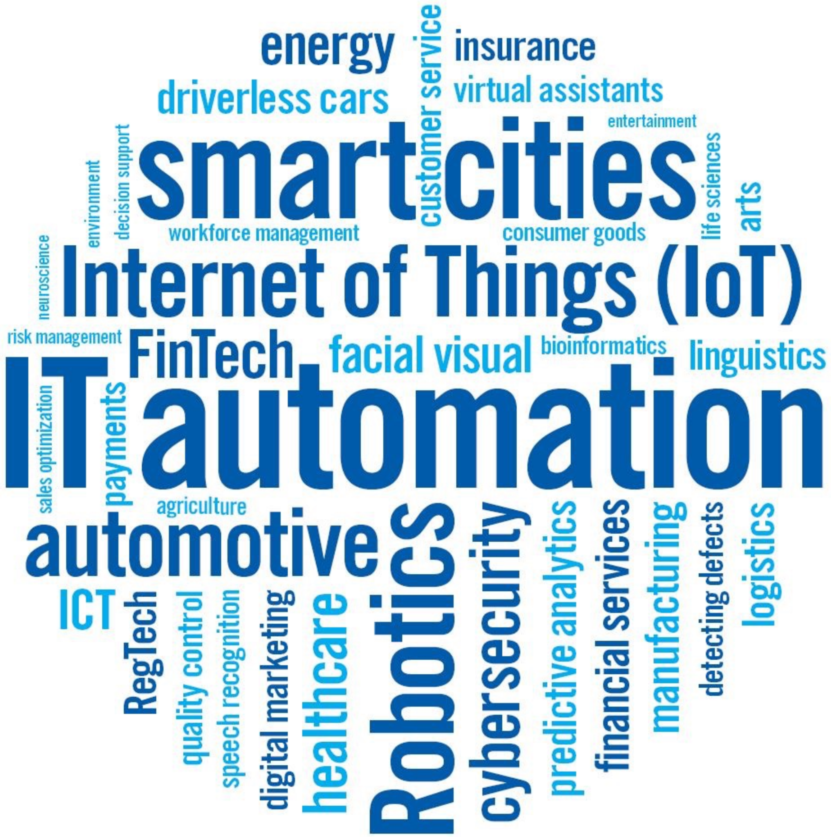

OVERVIEW

OUTLINE OF THIS TALK

A brief history of medical image analysis

- Typical tasks in medical image analysis
- From rule-based image processing, over ML to DL

Where are we now?

- **Current performance**
- **Challenges**

Outlook

- The role of doctors in AI for medical imaging
- Novel innovations on the horizons

THE FIRST DECADE OF DEEP LEARNING

AN INCOMPLETE OVERVIEW

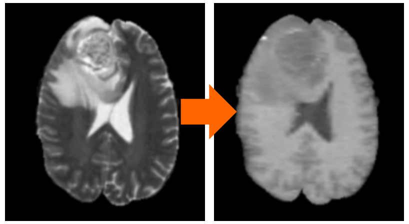

GANs competing networks improve each other

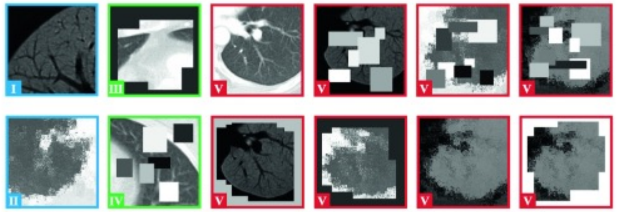

Self-supervised learning genesis, contrastive learning

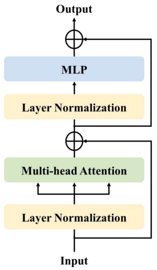

Visual Transformers Outperforming CNNs

2012 2014 2016 2018 2020 2022

IMAGENET

CNNs AlexNet, VGG16, GoogLeNet

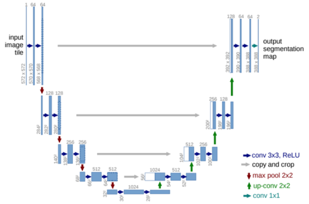

U-Net for image segmentation

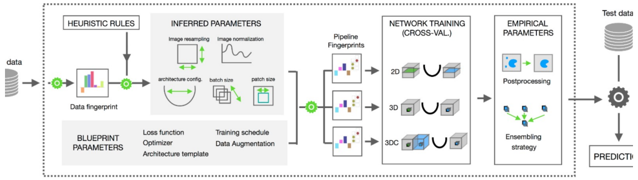

nnU-Net a self-configuring U-Net

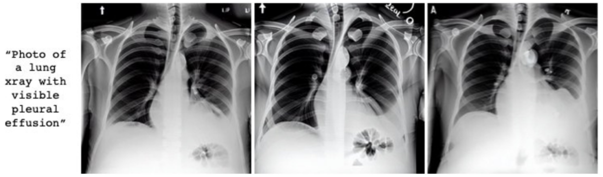

Generative models generate realistic images (even from text)

THE NEW WORKFLOW
IN THE DEEP LEARNING ERA

Workload has shifted to the collection of data and annotations

Performance is largely determined by the data
- Quantity
- Variety

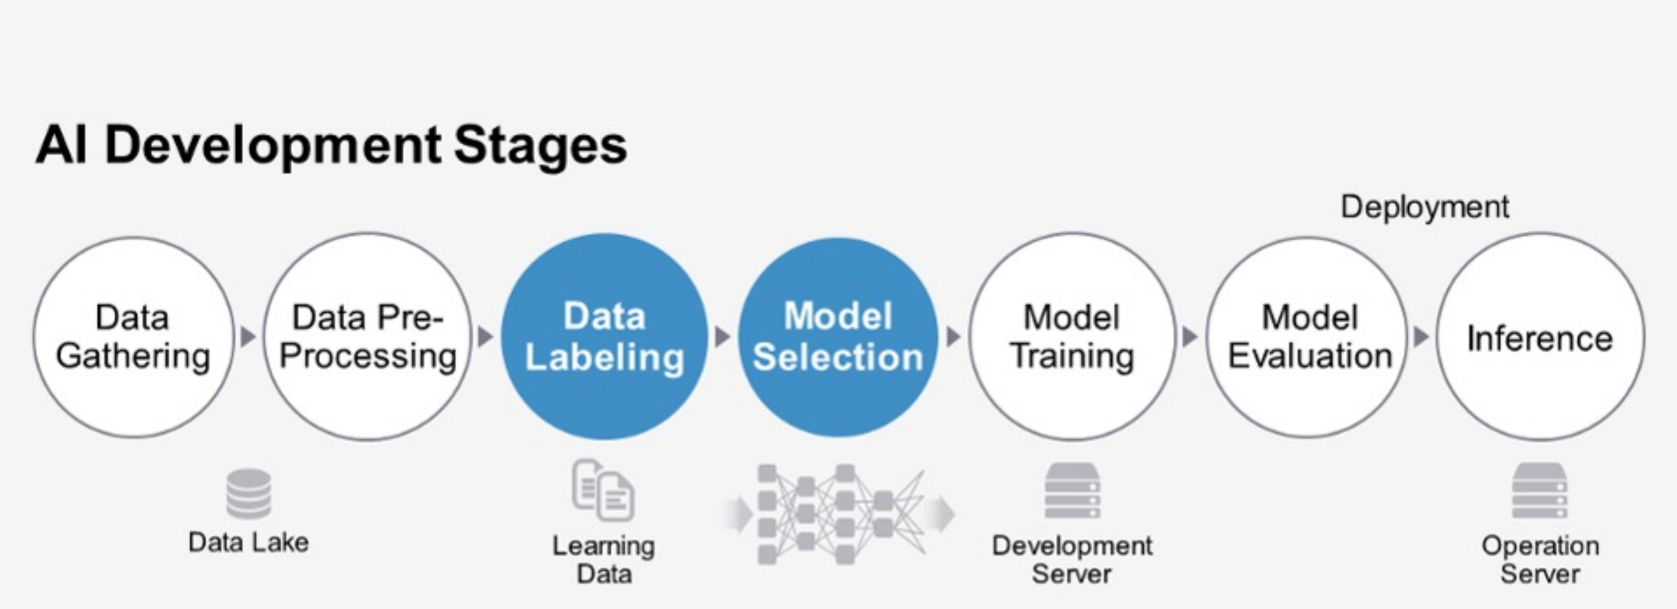

Cases needed for development / training

Before: 5-20 cases

Using deep learning: 100-1000 cases

Time needed for method to reach 80%

Before: 24-36 months

Using deep learning: 3-6 months

icovid

COVID-19

FAST DEVELOPMENT OF DIAGNOSTIC SUPPORT

Diagnostic support
Increase of CT imaging with AI

Image classification
Assign a COVID-19 likelihood to each CT image

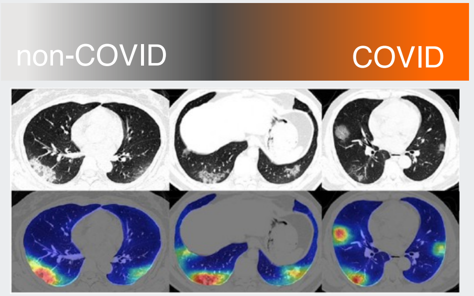

0.87 AUC

Lung involvement
To support the severity scoring

Image segmentation
Delineate COVID lesions within the CT image

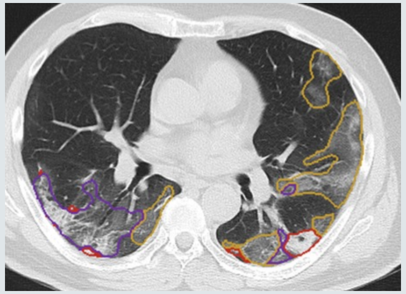

<50ml error

Risk assessment
Identify patients at risk

Predictive modelling
Combine multimodal features to predict patient outcome

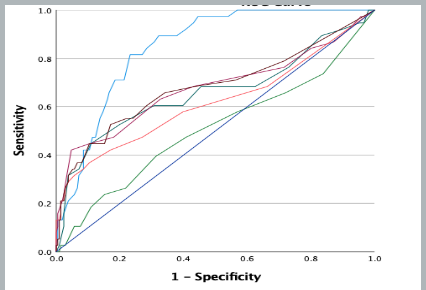

0.80 AUC

Working segmentation within 3 weeks, and (preliminary) certified in 6 weeks

MEASURING PERFORMANCE CAN BE TRICKY

HIDDEN STRATIFICATION

Performance largely depends on your data
Performance metrics depend on the data distribution
Bias is very difficult to fully remove

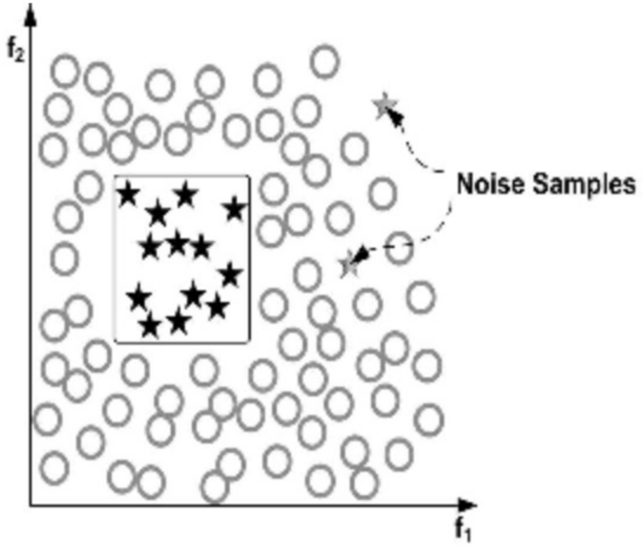

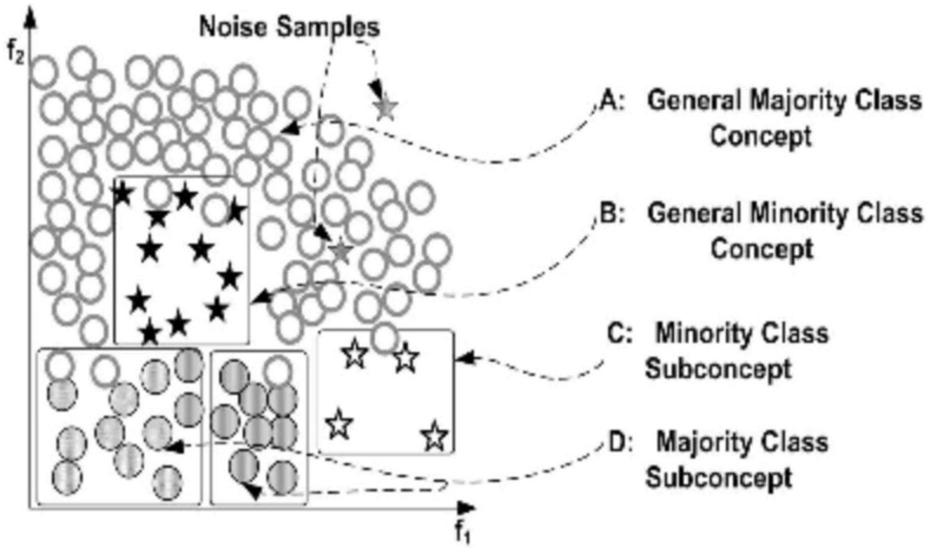

Overall performance: 90%

- 95% on the general concept class B
- 50% (random guesses) on the subconcept C

Learning from imbalanced data
He *et al.* IEEE Transactions on Knowledge and Data Engineering, 2009

Hidden Stratification Causes Clinically Meaningful Failures in Machine Learning for Medical Imaging
Oakden-Rayner *et al.* Conf Health Inference Learn, 2020

RISK OF BIAS

COVID-19

The domain needs to mature and adopt tools for bias assesment and reporting
- PROBAST
- TRIPOD

Prediction models for diagnosis and prognosis of covid-19: systematic review and critical appraisal. Wynants L., et al. BMJ (2020)

- "**Thirty three** diagnostic models were identified for detecting covid-19, in addition to **75** diagnostic models based on medical images, **10** diagnostic models for severity classification, and **107** prognostic models for predicting, among others, mortality risk, progression to severe disease"
- "Proposed models are poorly reported and at **high risk of bias**, raising concern that their predictions could be unreliable when applied in daily practice"
- "We cannot yet recommend **any** of the identified prediction models for widespread use in clinical practice"

OVERVIEW
OUTLINE OF THIS TALK

A brief history of medical image analysis

- Typical tasks in medical image analysis
- From rule-based image processing, over ML to DL

Where are we now?

- Current performance
- Challenges

Outlook

- **The role of doctors in AI for medical imaging**
- **Novel innovations on the horizons**

WILL WE GO THROUGH A NEW AI WINTER?

A PERIOD FOLLOWING DELILLUSIONEMENT

It seems unlikely, as AI seems to bring value, this time..

An example from United Imaging (Shanghai, China)

- Full stack: AI along the entire pipeline
- Full spectrum: all modalities

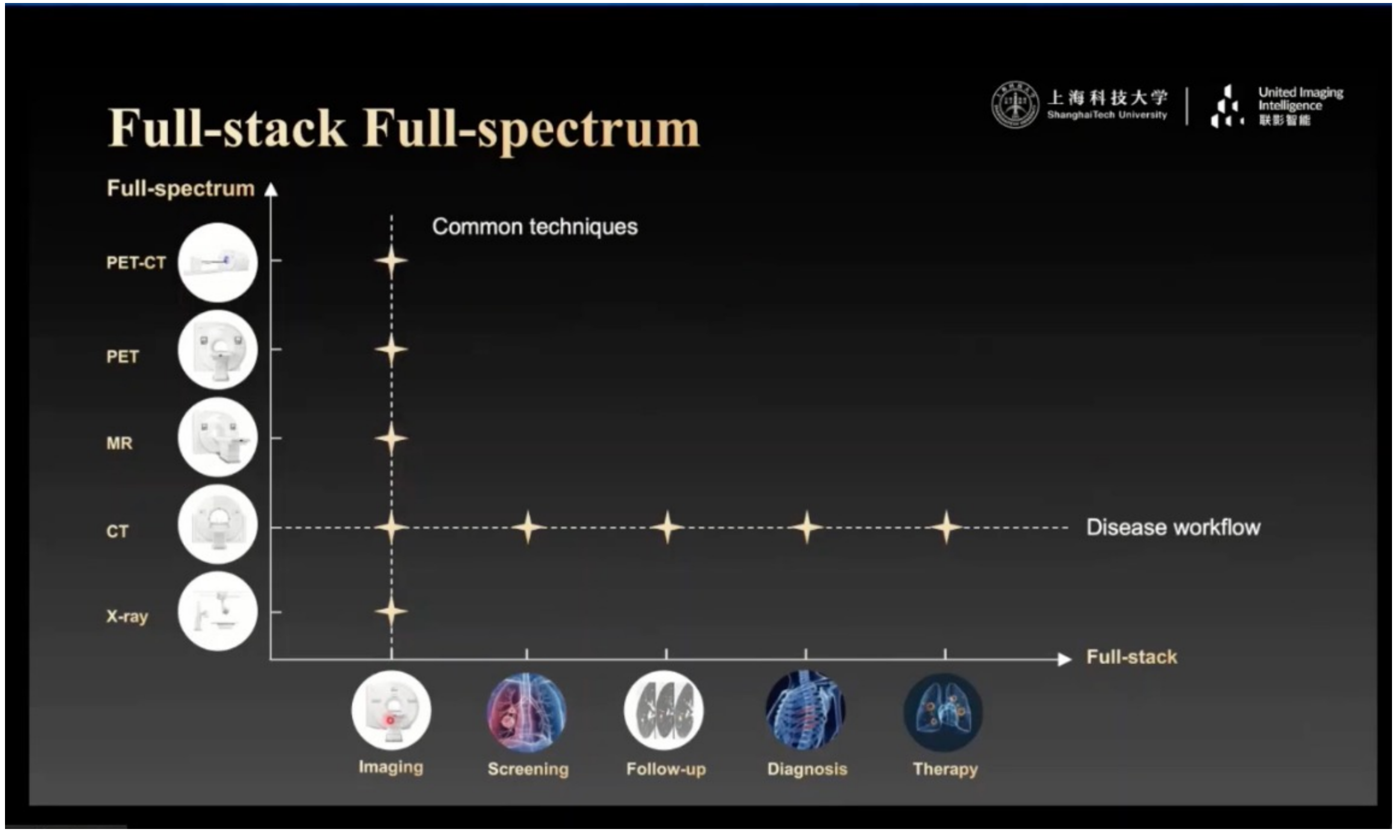

Full-Stack, Full-Spectrum AI in Medical Imaging
Dinggang Shen, Keynote at MICCAI 2022, Singapore

WILL AI REPLACE THE MD?

NO, BUT IT WILL IMPACT THE WAY OF WORKING

“Artificial intelligence will not replace the radiologist . Rather, radiologists who do not embrace AI will be replaced by those who do”
President of the American Radiological Associations

The real challenge is how to integrate AI in the workflow

- We mainly think about improving **sensitivity**
  - Double reading: AI+MD
- CAD systems are using MDs to optimize **specificity**
  - Alert fatigue
- Our health economic situation is pushing us to improve **efficiency**
  - Where do we trust AI?

AI FOR BUT ALSO BY MDS

CO-DEVELOPMENT

Medical practitioners will be actively involved in the entire process

- Data curation considering stratification
- AI-assisted annotation and labelling

Deep learning will be enhanced by medical domain knowledge

- Improve performance & stability
- Including uncertainty & explainability

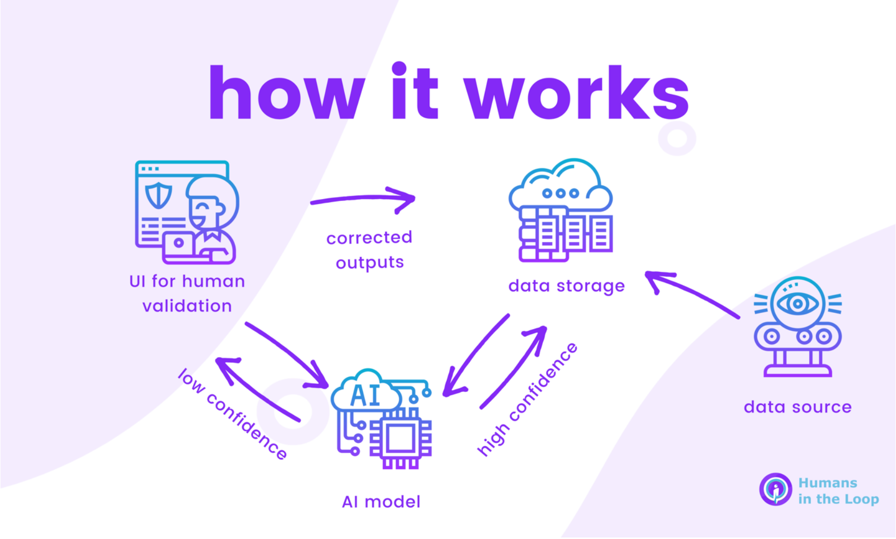

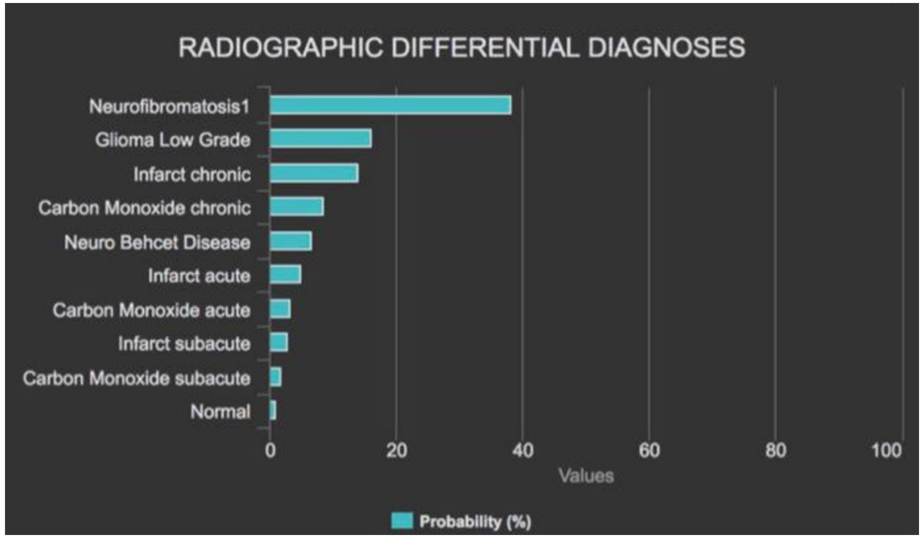

Brain MRI Deep Learning and Bayesian Inference System Augments Radiology Resident Performance
Rudie *et al.* J Digit Imaging. 2021

SELF-SUPERVISION FOR MEDICAL DATA

TOWARDS MULTI-MODAL CLINICAL FOUNDATION MODELS

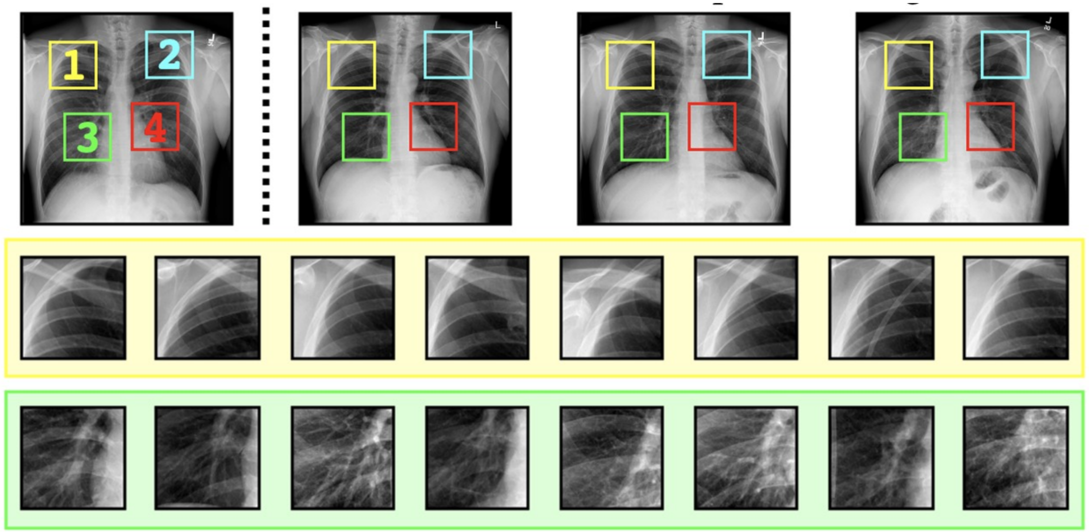

Anatomy and pathology-aware unsupervised pretraining

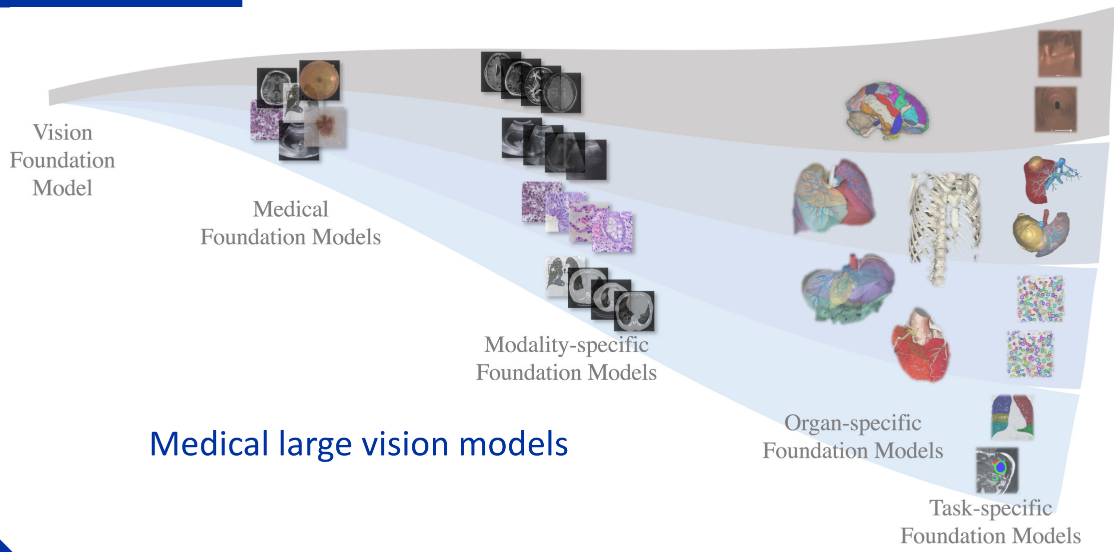

Medical large vision models

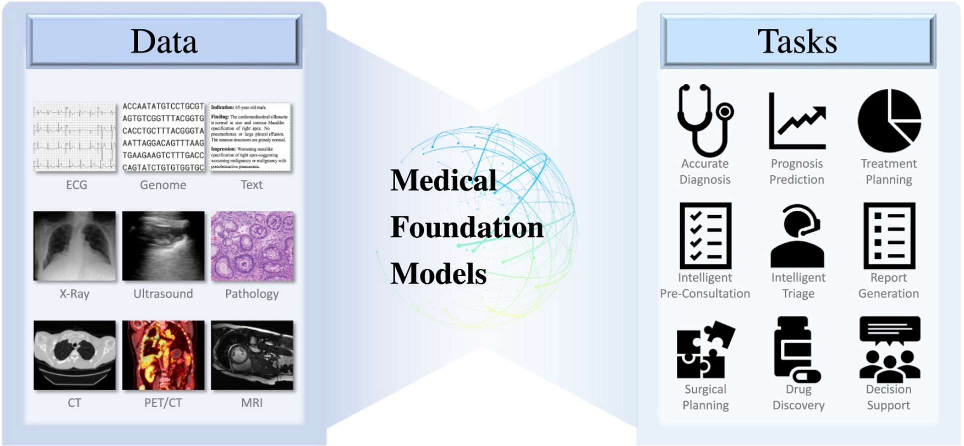

Multi-modal medical foundation models

Haghighi, et al. "Learning semantics-enriched representation via self-discovery, self-classification, and self-restoration." MICCAI 2020.

Shaoting Zhang & Dimitris Metaxas, On the challenges and perspectives of foundation models for medical image analysis, Medical Image Analysis, 2024.