UCLouvain
icteam Institute of Information and Communication Technologies, Electronics and Applied Mathematics

Coalitions intégratives pour une utilisation de confiance de l'intelligence artificielle en imagerie médicale

MIAM: Midi de l'Intelligence Artificielle pour la Médecine

Benoit Macq – benoit.macq@uclouvain.be

www.pilab.be

Artificial Intelligence

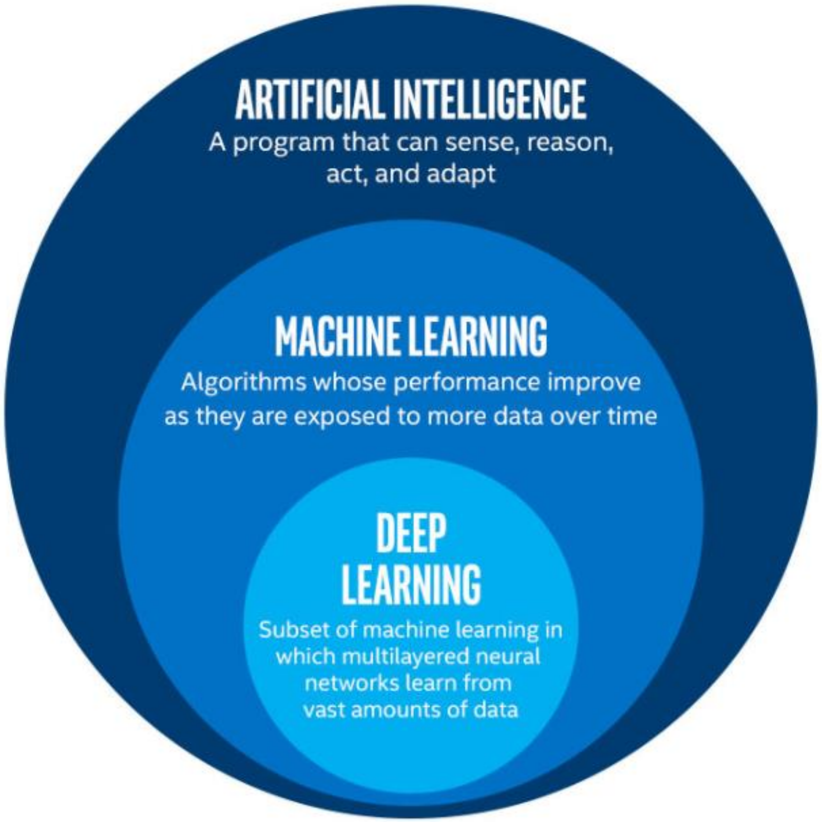

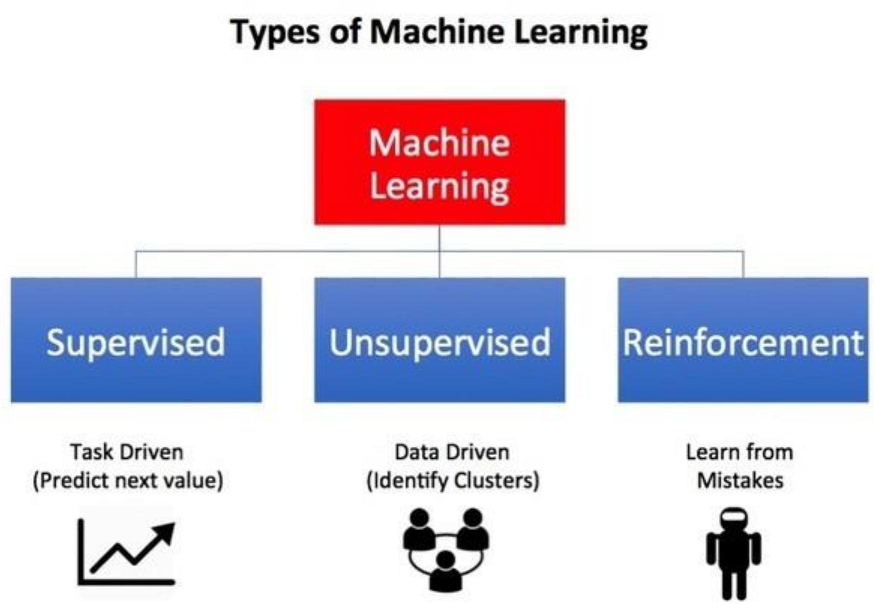

Deep learning (e.g. MILA-Montréal)

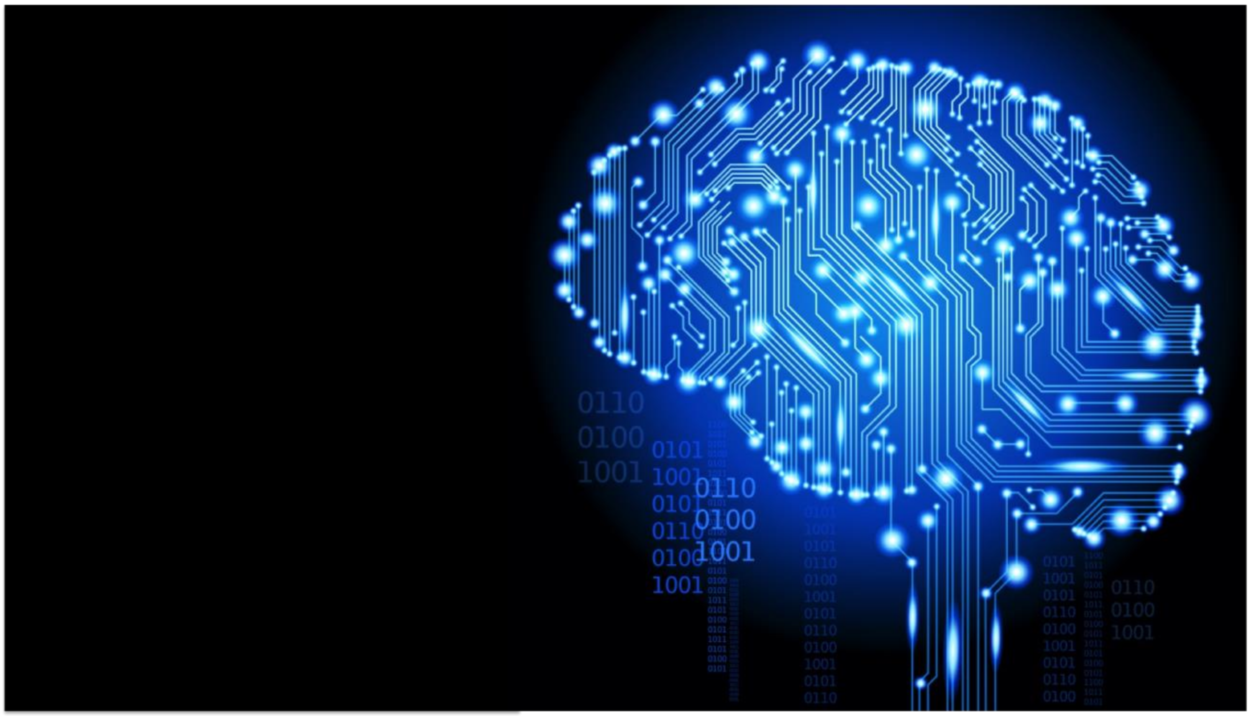

Deep learning success stories

Gaming

AlphaGo beat (4-1) world champion Lee Sedol *(March 2016)*

Self-driving car

Vehicules have driven 1.6m km
Fautive in 1 crash (June 2015)

Image recognition

ImageNet Challenge: classify 1.2m high-res. images
U. of Toronto team reaches 17% top-5 error rate (2012)

Why not medical imaging ?

enlitic

MIT technology 50 smartest companies
6000 lung cancer diagnoses
50% more accurate than human radiologists

IBM Watson

Watson Health medical imaging collaborative
15 health systems, medical centers and imaging comp.
Data from ~300m patients

Google DeepMind

Applying machine learning to RT planning for H&N cancer
Objective: segmentation process 4 hours → 1 hour

ImagX, BidMed, … Telemis, Intuitim, DNAlytics, Oncoradiomics

Image-based decision in the previous millenium

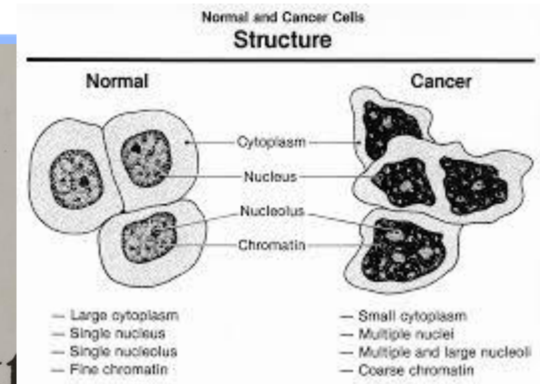

IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, VOL. 43, NO. 10, OCTOBER 1996

Morphological Feature Ext... for the Classification of Digital Images of Cancerous Tissues

Jean-Philippe Thiran,* *Student Member, IEEE*, and Benoît Macq, *Member, IEEE*

Abstract—This paper presents a new method for automatic recognition of cancerous tissues from an image of a microscopic section. Based on the shape and the size analysis of the observed cells, this method provides the physician with nonsubjective numerical values for four criteria of malignancy. This automatic approach is based on mathematical morphology, and more specifically on the use of Geodesy. This technique is used first to remove the background noise from the image and then to operate a segmentation of the nuclei of the cells and an analysis of their shape, their size and their texture. From the values of the extracted criteria, an automatic classification of the image (cancerous or not) is finally operated.

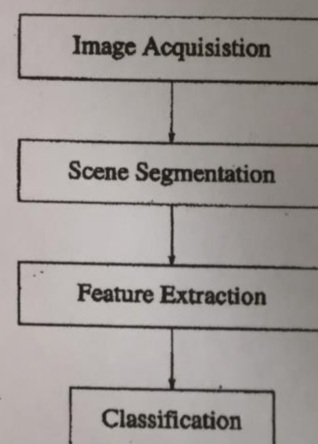

Fig. 1. General structure of the method

I. Introduction

OBJECTIVE analysis of microscopic images of cells and

Deep Learning high predicitive power (e.g. facial recognition)

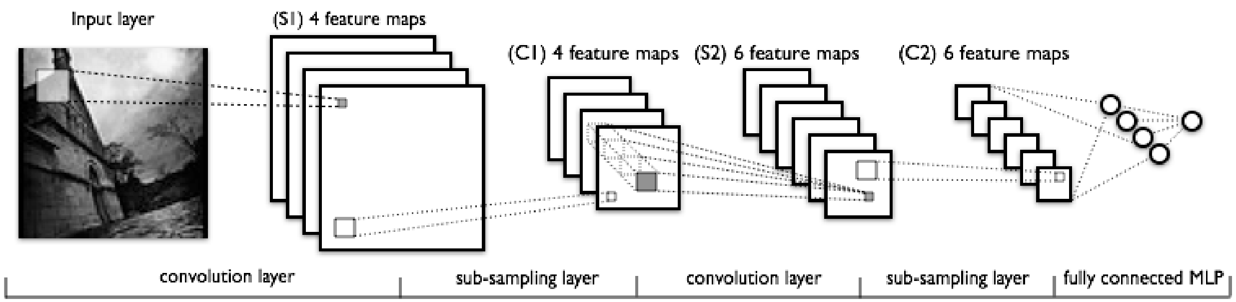

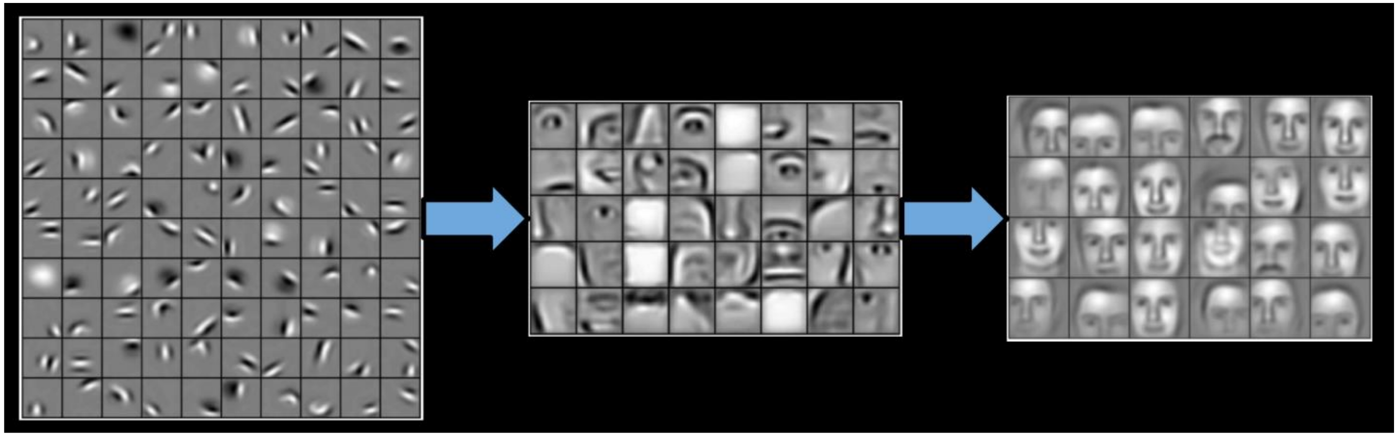

The move based on the high predictivity power of deep learning

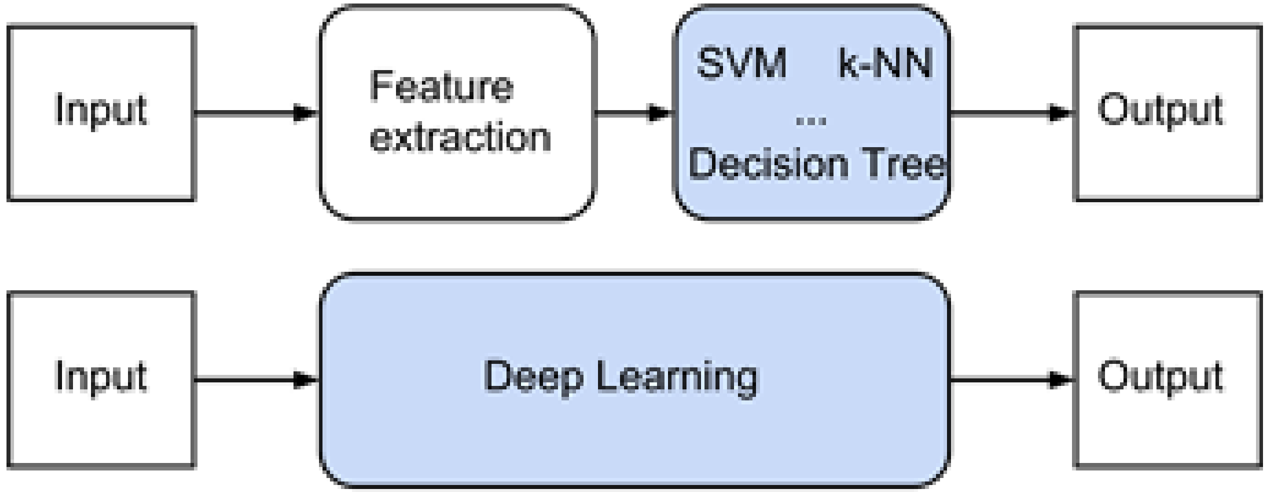

Convolutional Neural Networks (CNN)

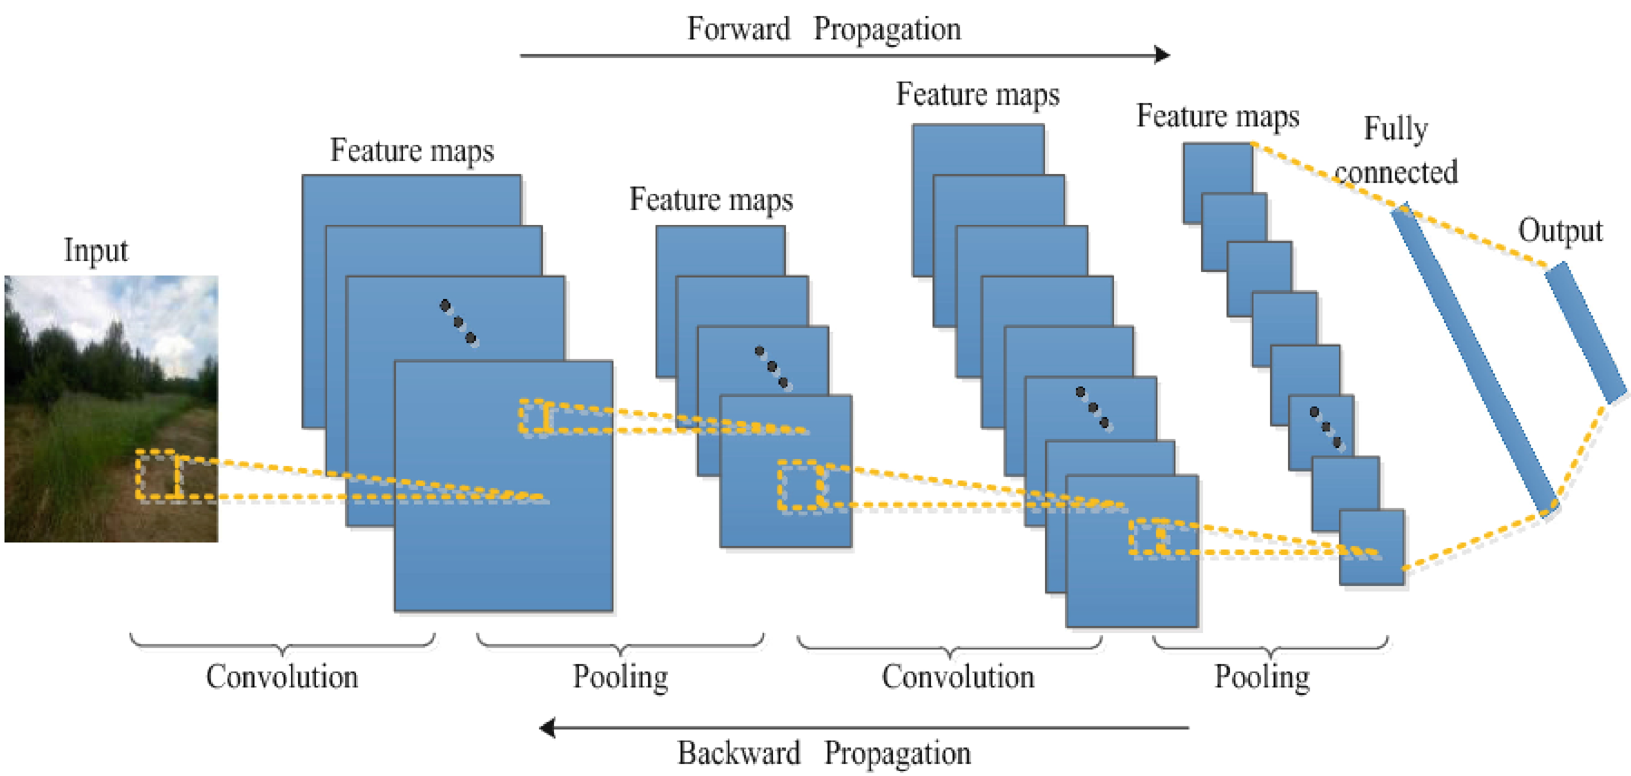

Fig. 1. The CNN structure

Pro and cons of Deep Learning

- Unique structure (CNN) for many problems
- « Generalisable » with regards the training set
- There exist a lot of opensource tools
- But: need for large annotated training sets
- But: lack of explainability of the deep features and their co-action
- But: lack of actionability

And ….

- Intriguing properties of neural networks (C. Szegedy et al.)

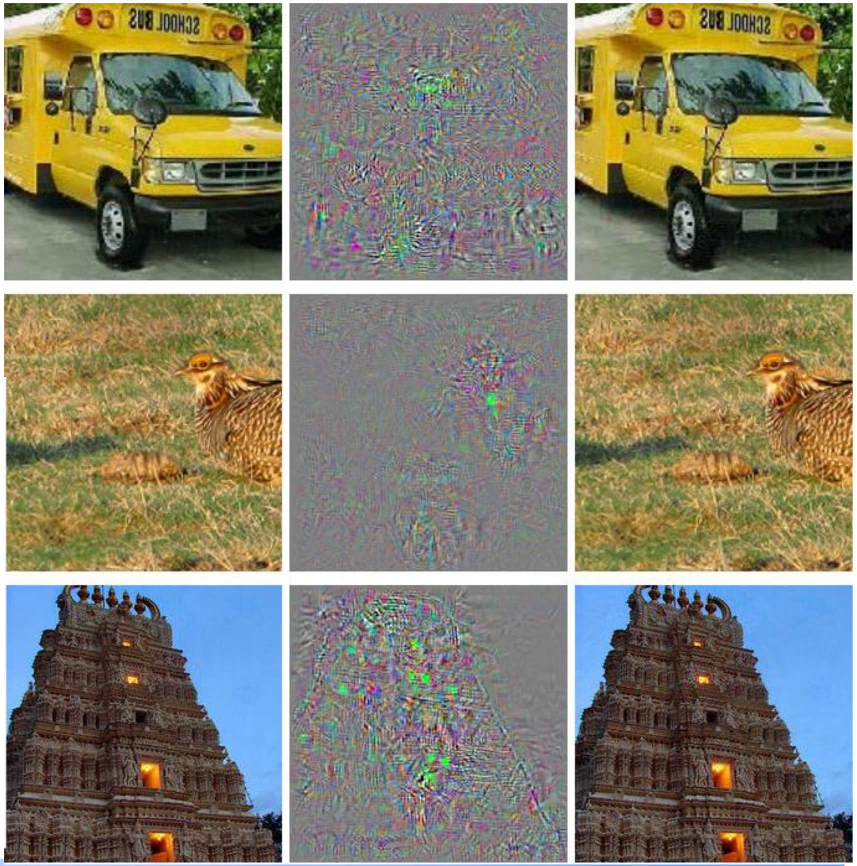

https://www.pluribus-one.it/research/sec-ml/wild-patterns

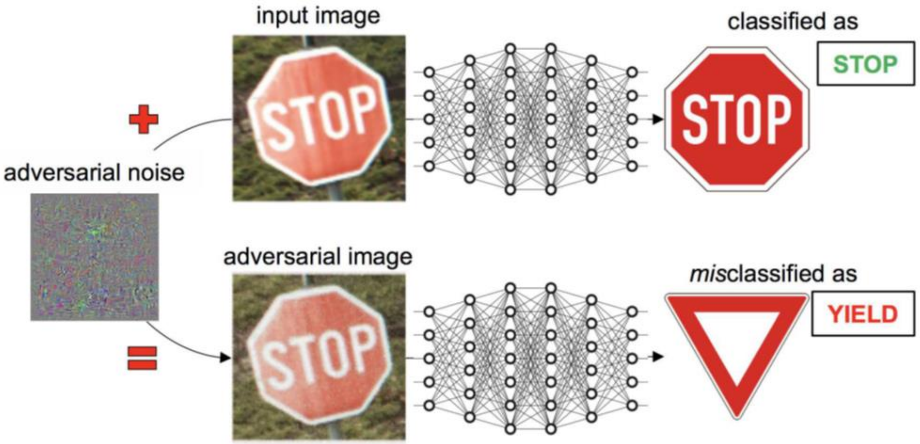

Deep Learning has an outstanding accuracy in difficult problems but hard to explain outliers

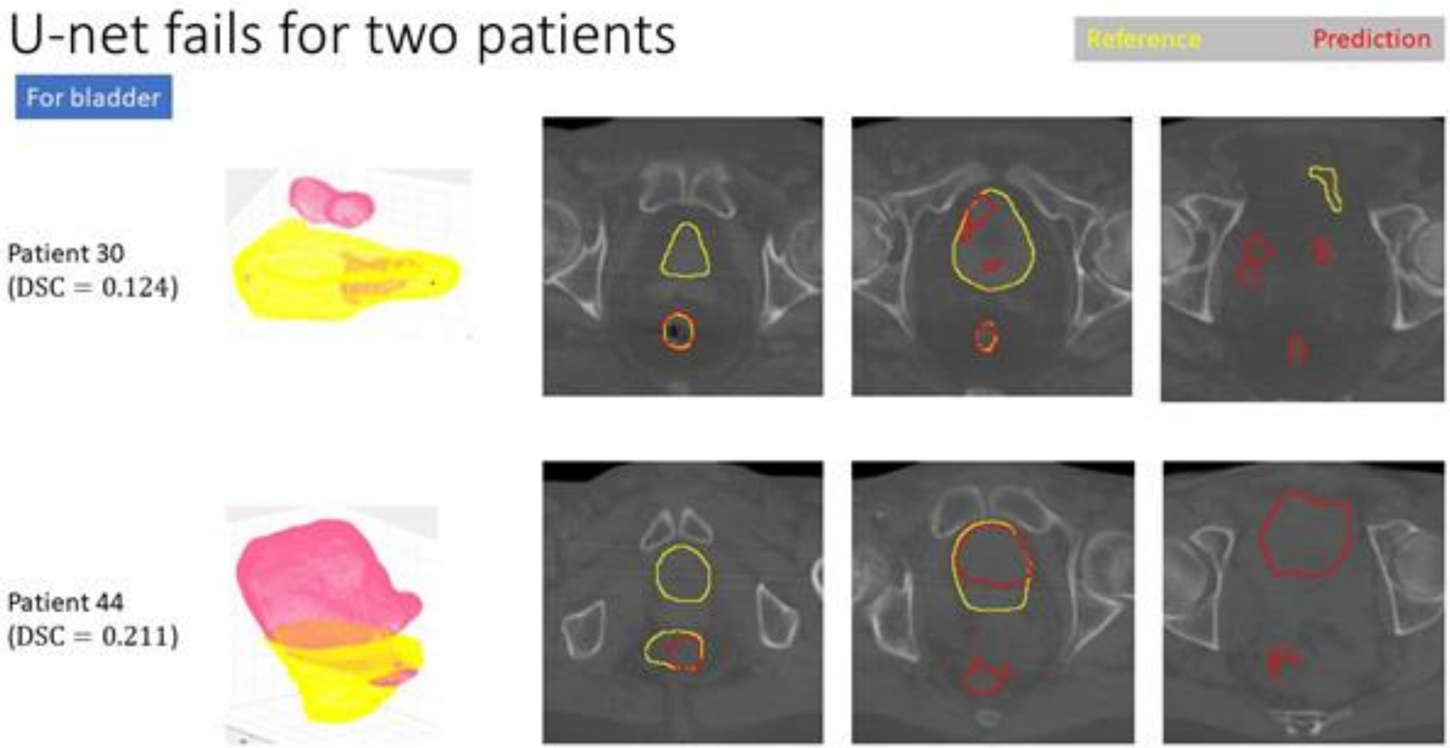

L'Intelligence Artificielle et les radiologues

« Les radiologues qui utiliseront l'IA remplaceront ceux qui ne l'utilisent pas »

1. « Unpredictable » outliers (reliability ?)
2. Explainability of the decision ?
3. Actionability and commitment
4. Data privacy (blockchained distributed learning)
5. Evolution of expertise

« Prédire n'est pas comprendre »

Explainability of algorithms (GDPR)

Deep learning in medical imaging

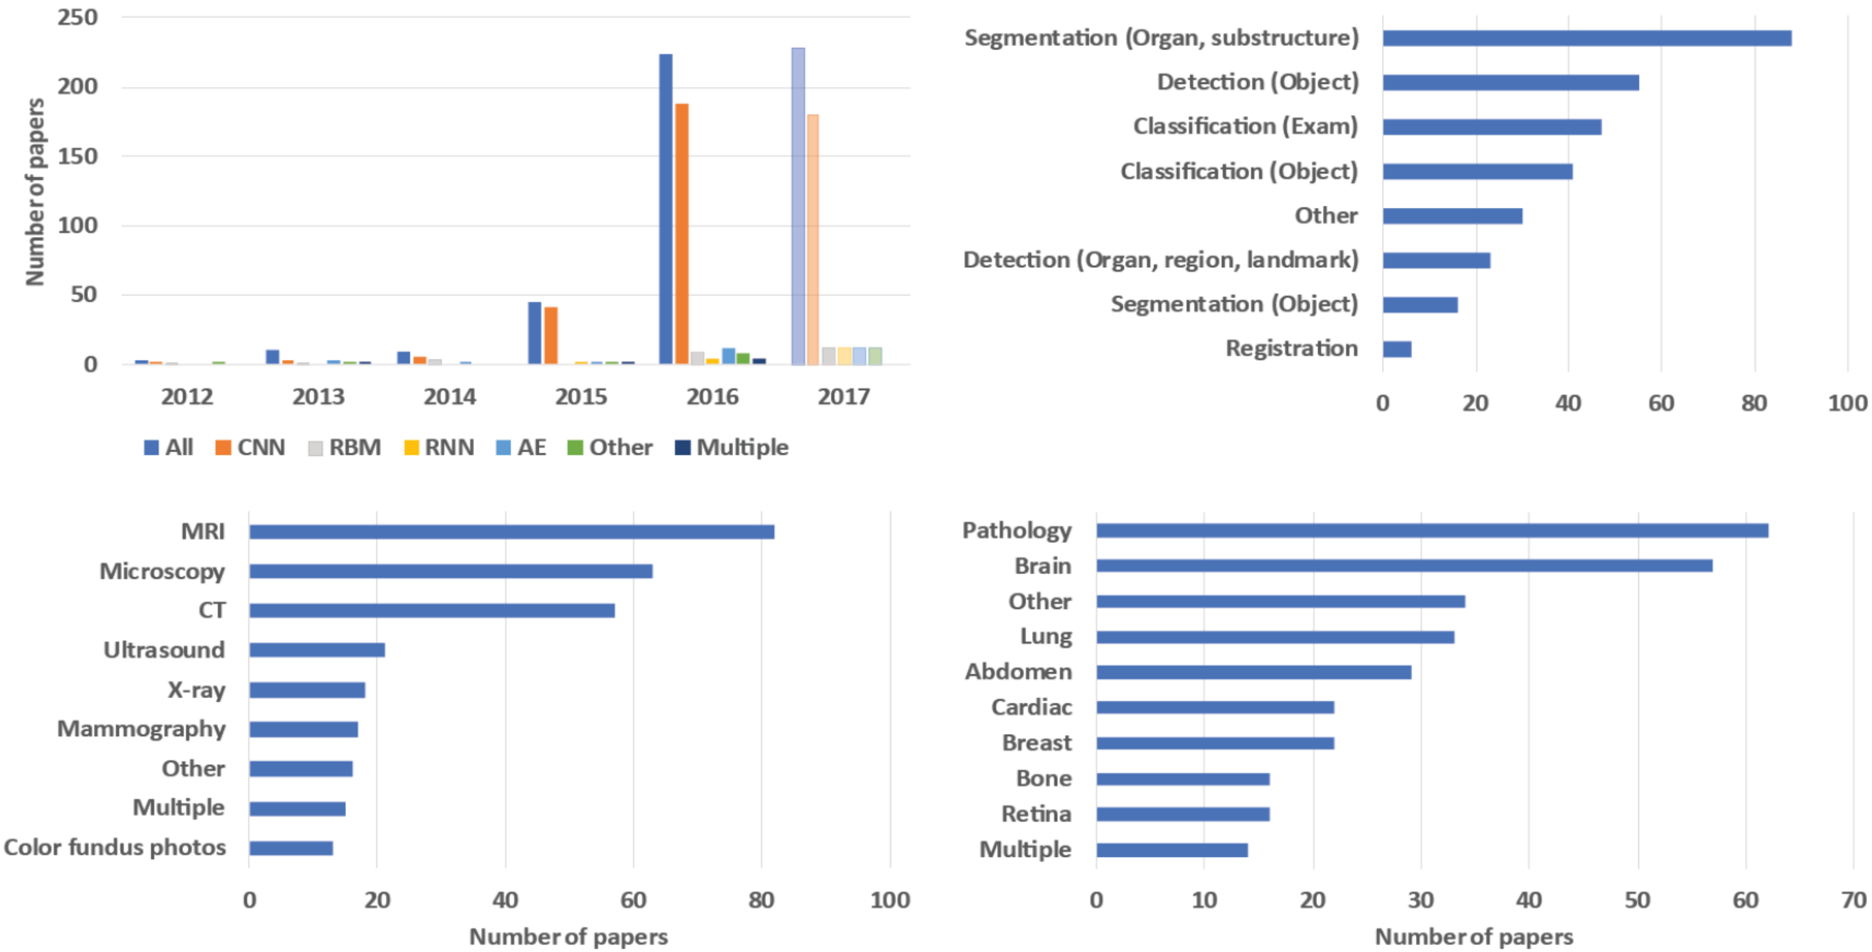

Litjens, *et al*. **a survey on deep learning in medical image analysis**. *Med Image Anal.* 2017

Advanced deep learning

- U-Net segmentation
- Generative Adversarial Networks (GANs)
- Deep Reinforcement learning

U-Net Segmentation

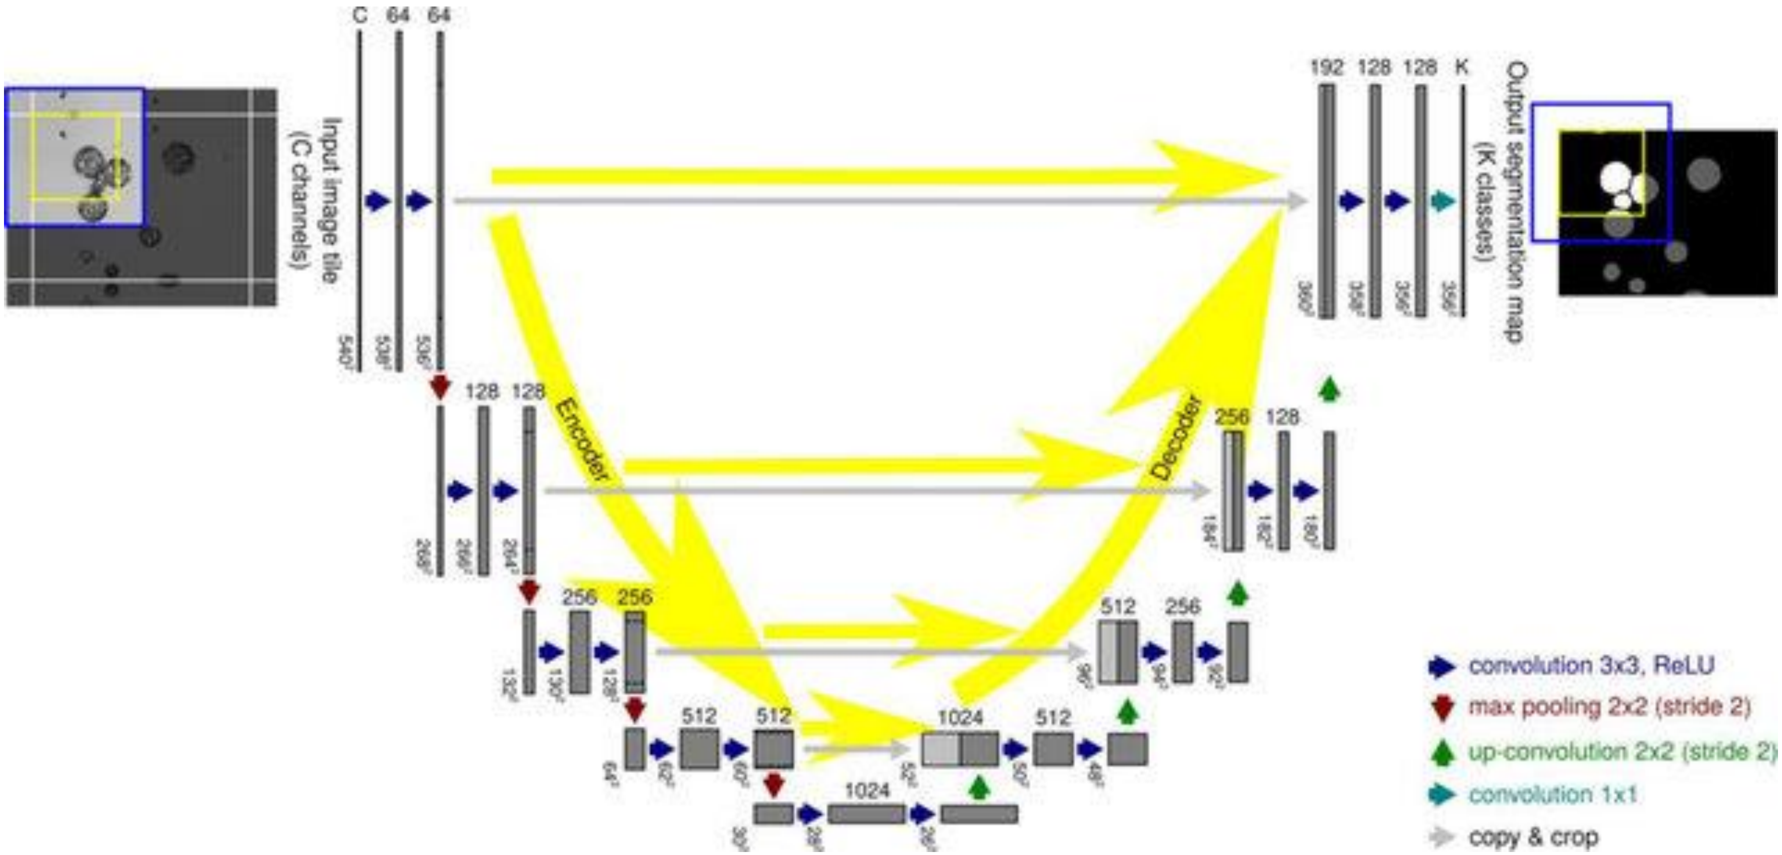

GANs

Generative Adversarial Network

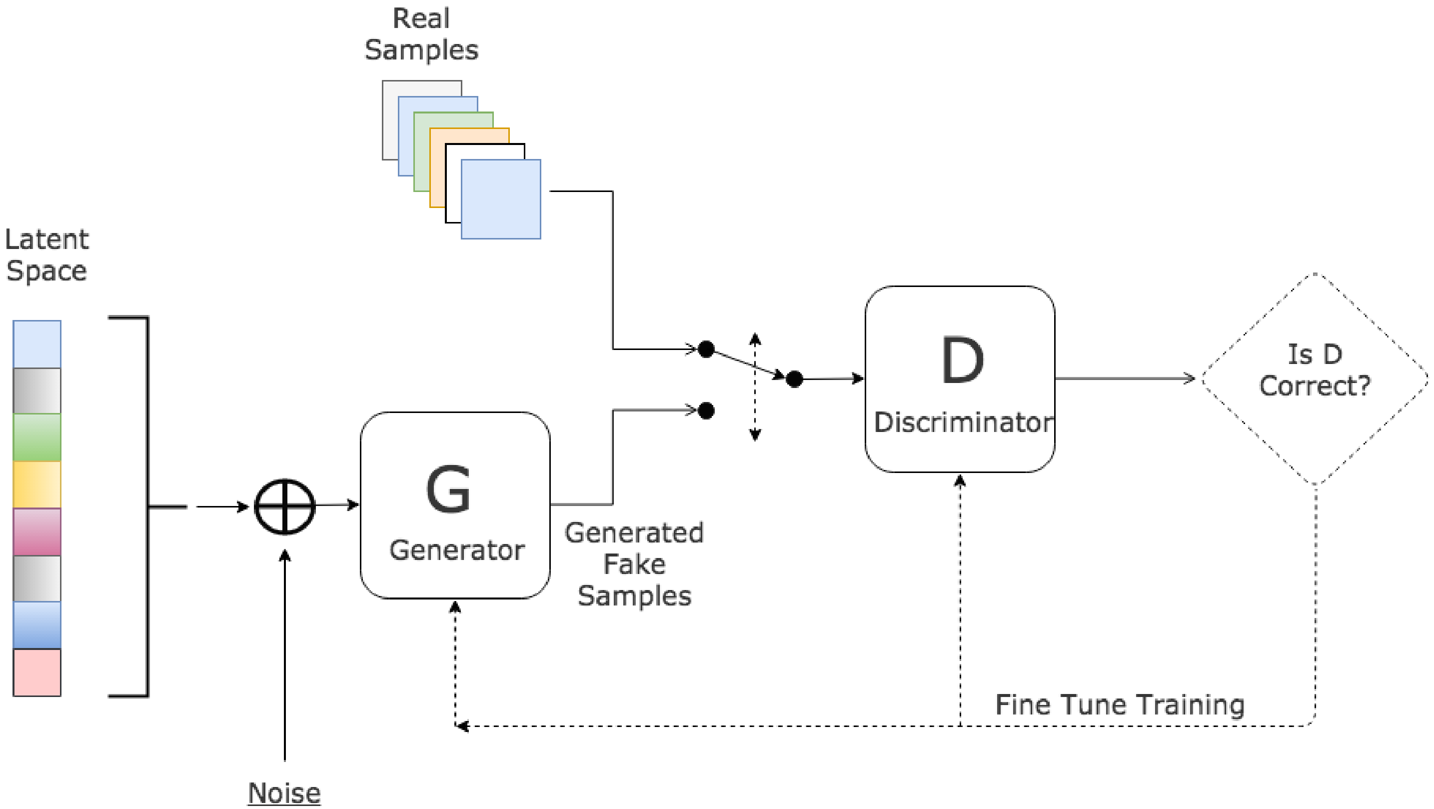

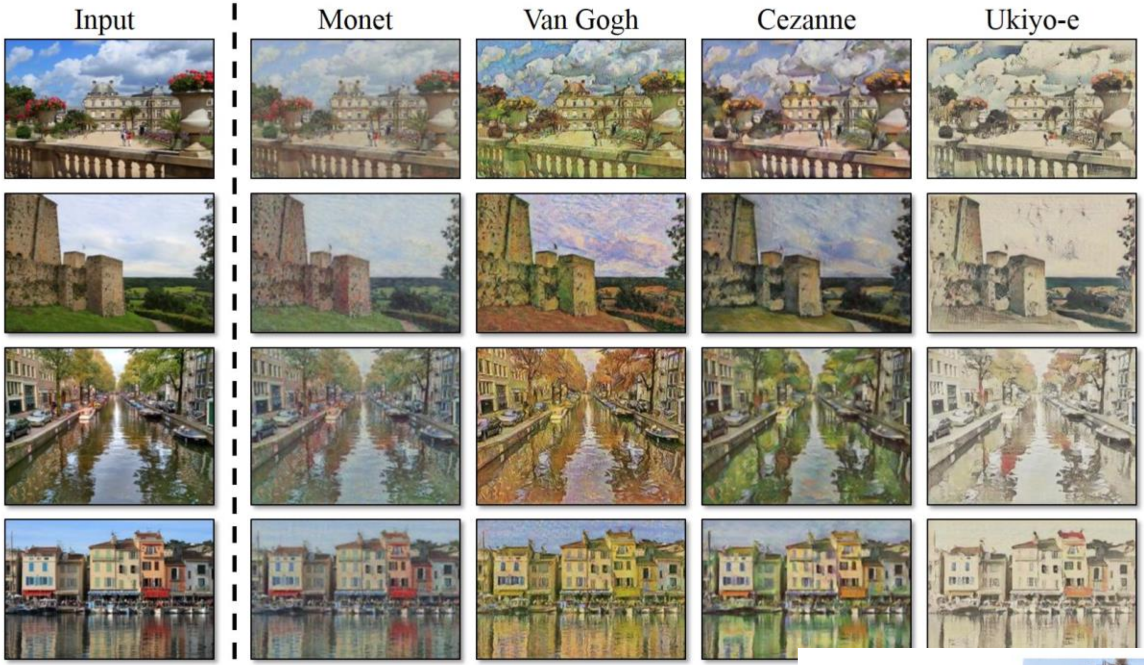
Input
Monet
Van Gogh
Cezanne
Ukiyo-e

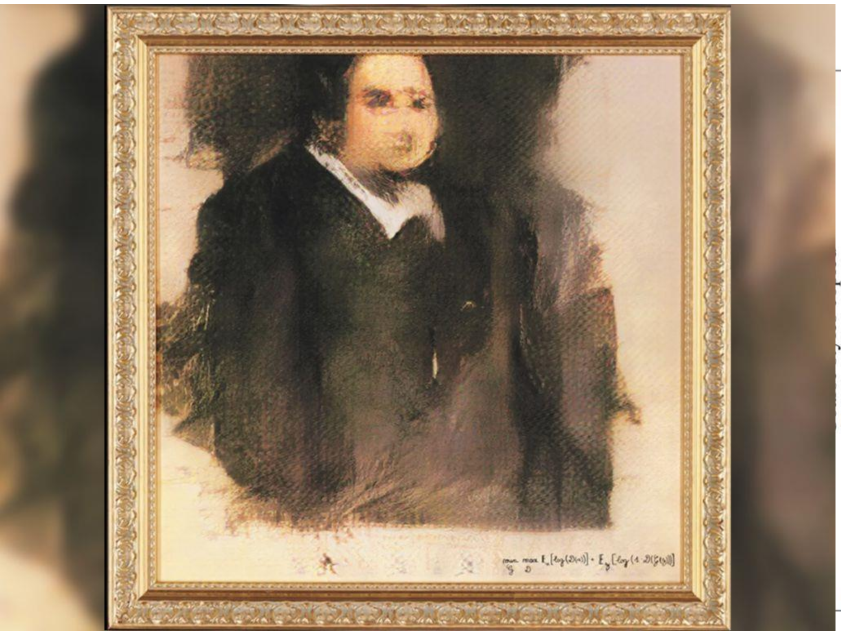


source
destination

Use of GAN in radiology

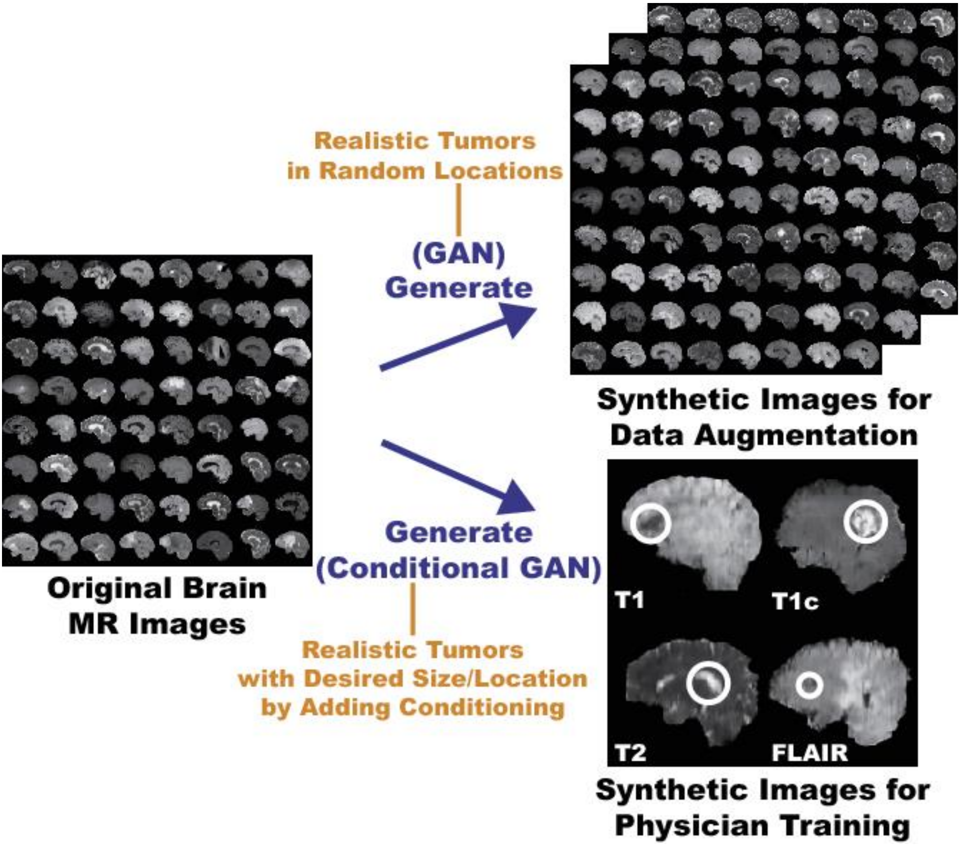

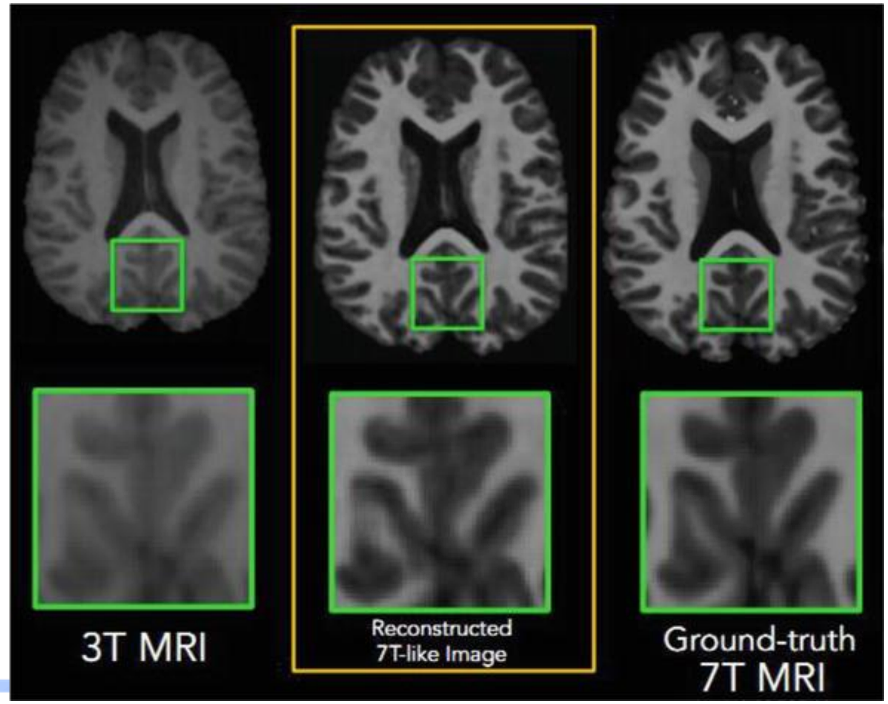

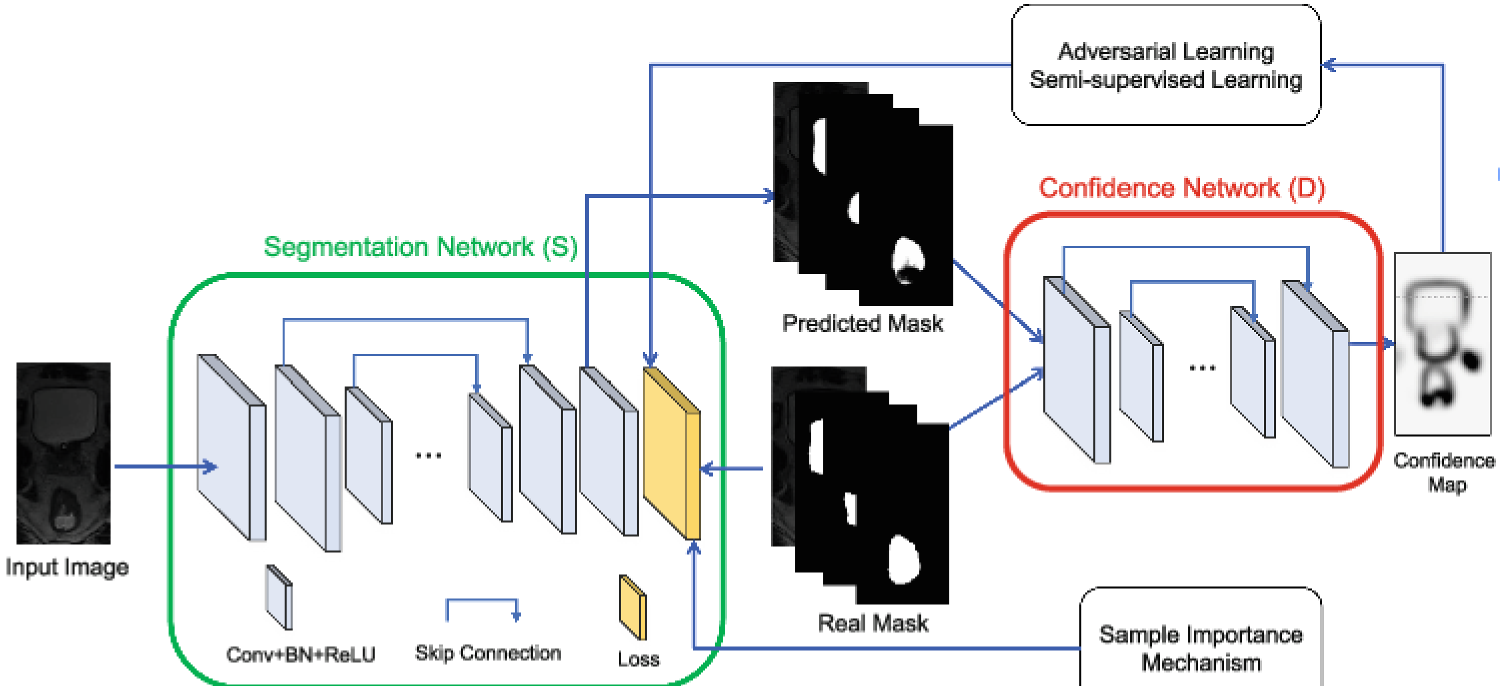

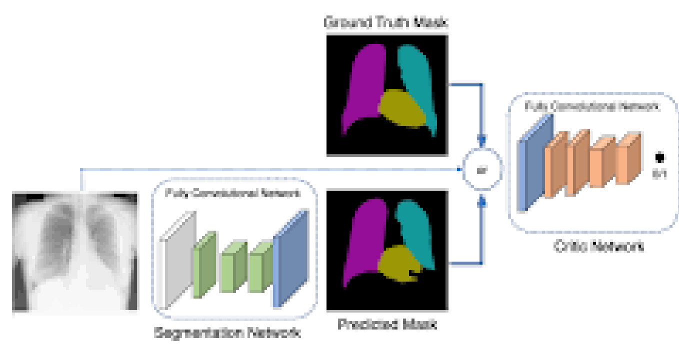

Figure 3: Overview of the proposed SCAN framework that jointly trains a segmentation network and a critic network with an adversarial mechanism. The segmentation network produces per-pixel class prediction. The critic takes either the ground truth label or the prediction by the segmentation network, optionally with the CXR image, and output the probability estimate of whether the input is the ground truth (with training target 1) or the segmentation network prediction (with training target 0).

Deep Reinforcement learning

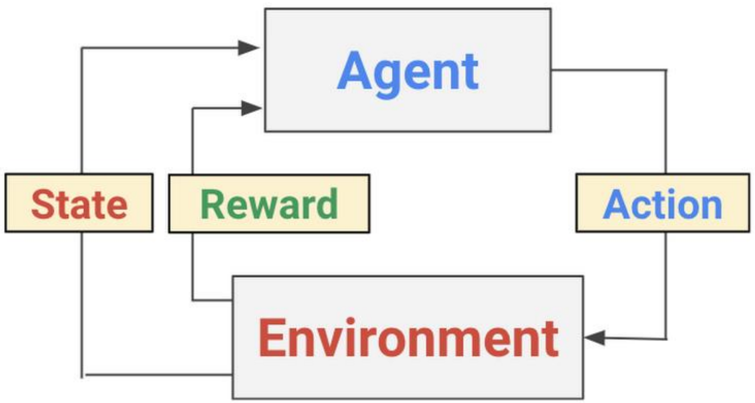

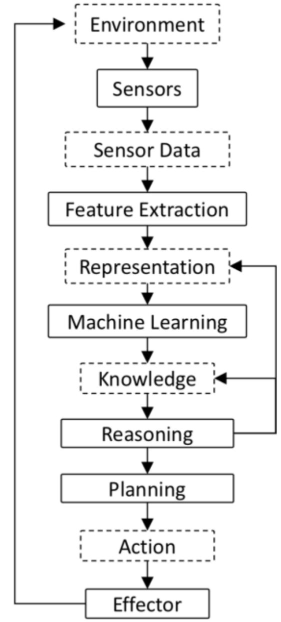

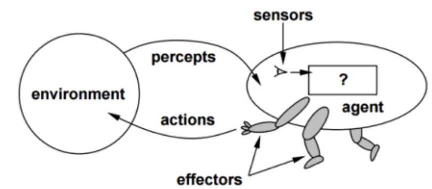

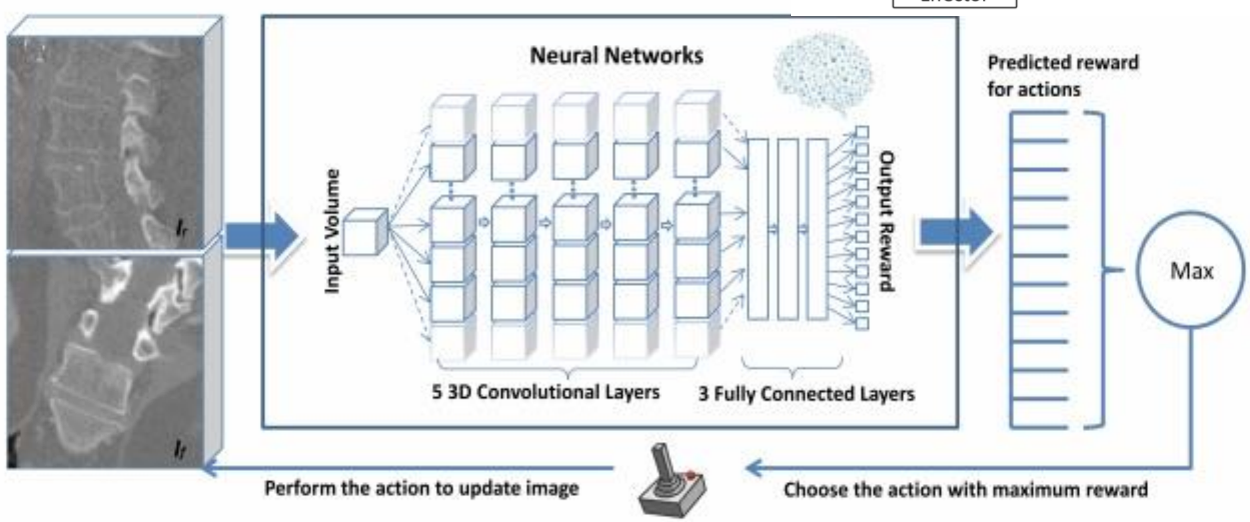

Optimizing by Deep (R)L

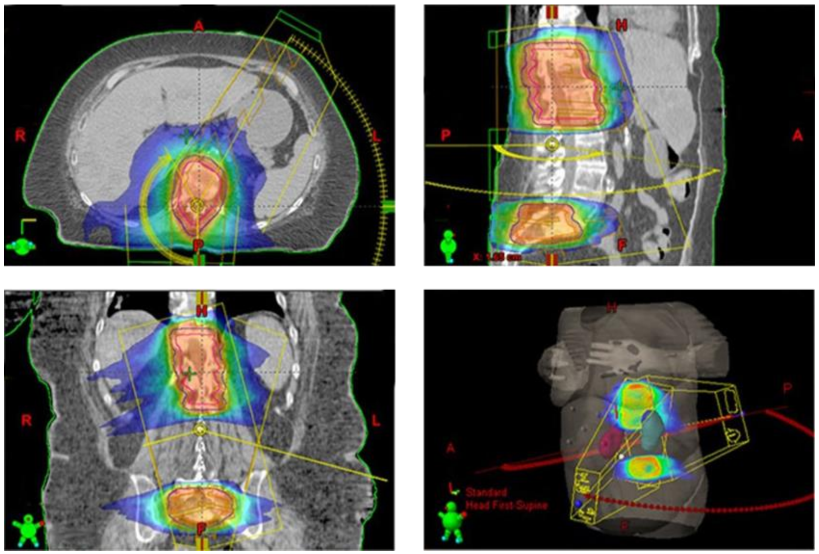

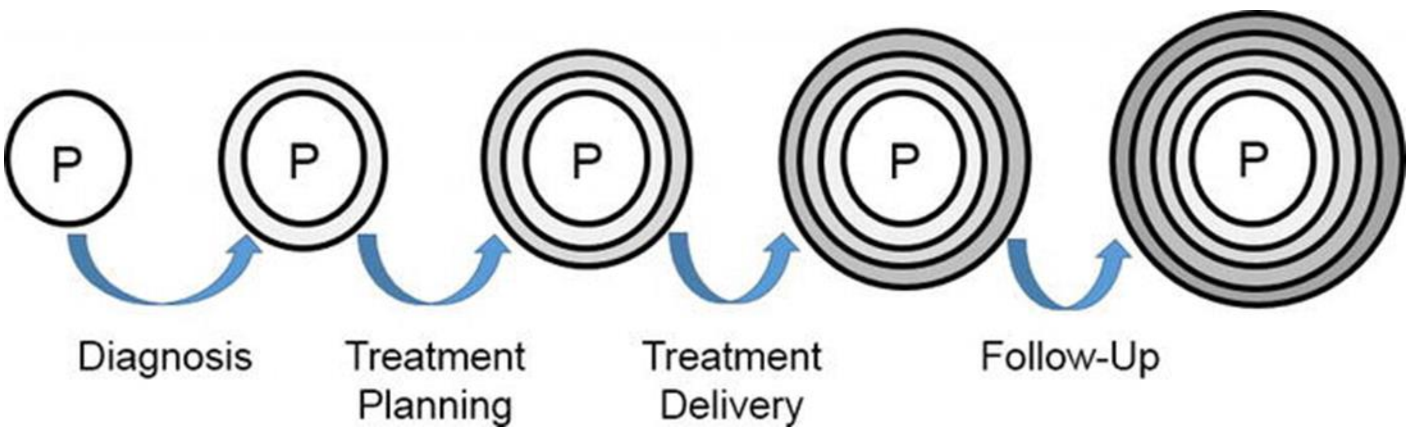

The Radiomics challenge

- Predictive and personnalized medicine

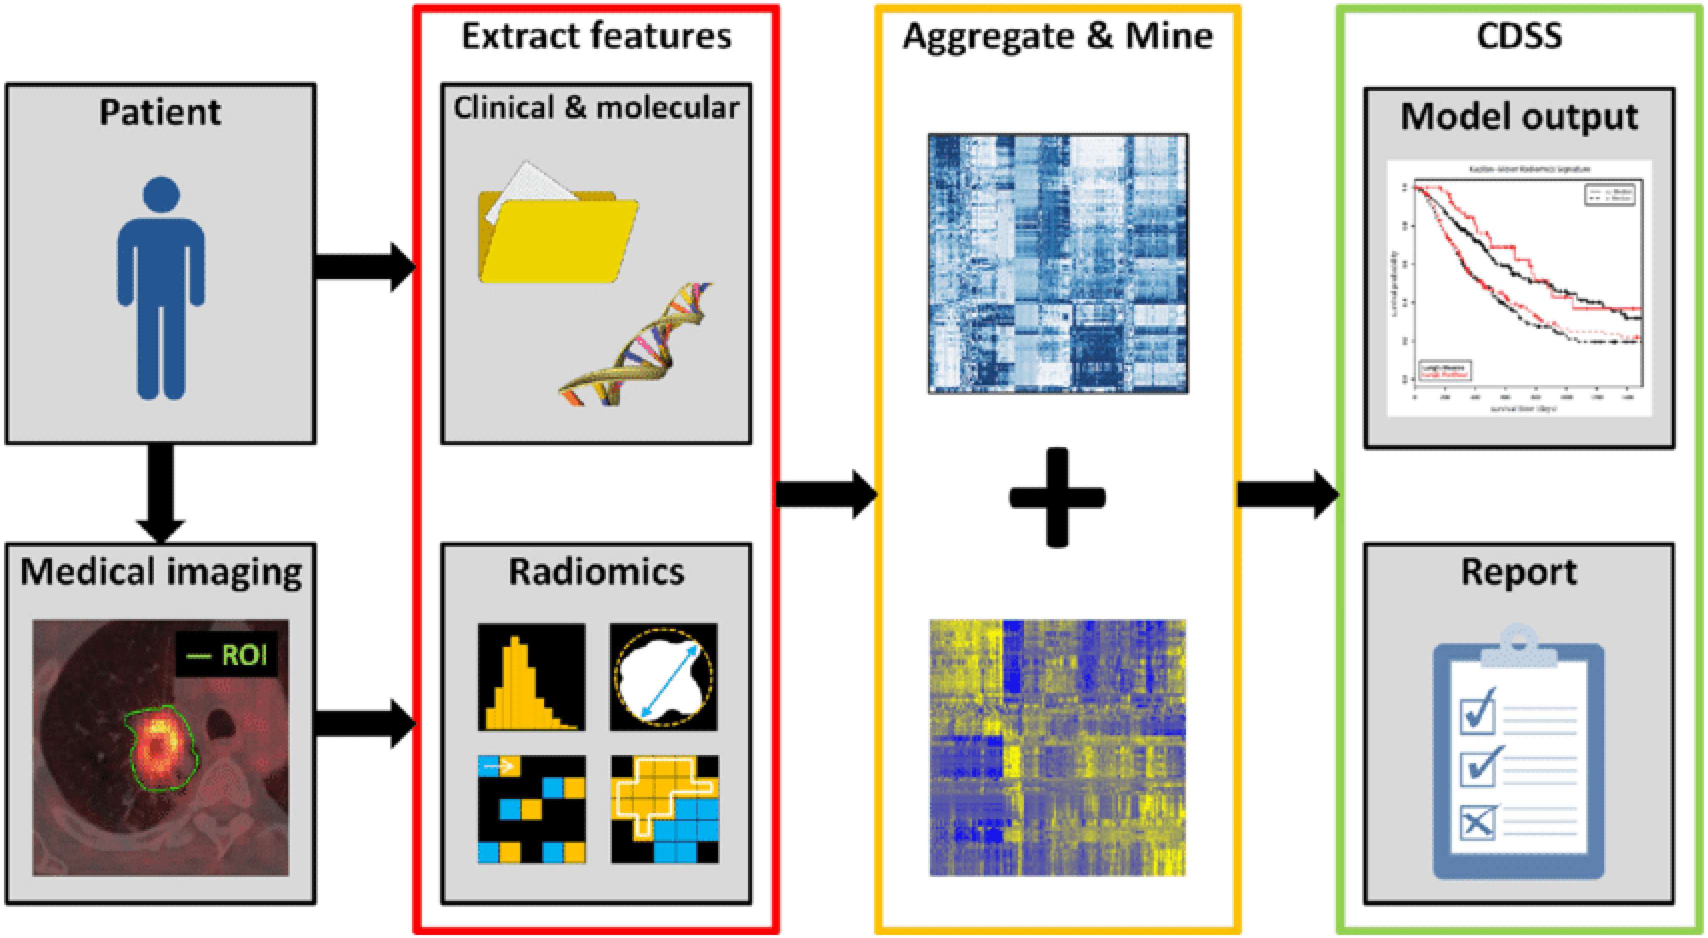

Research assumption 1: three kinds of latent spaces by multi-agents

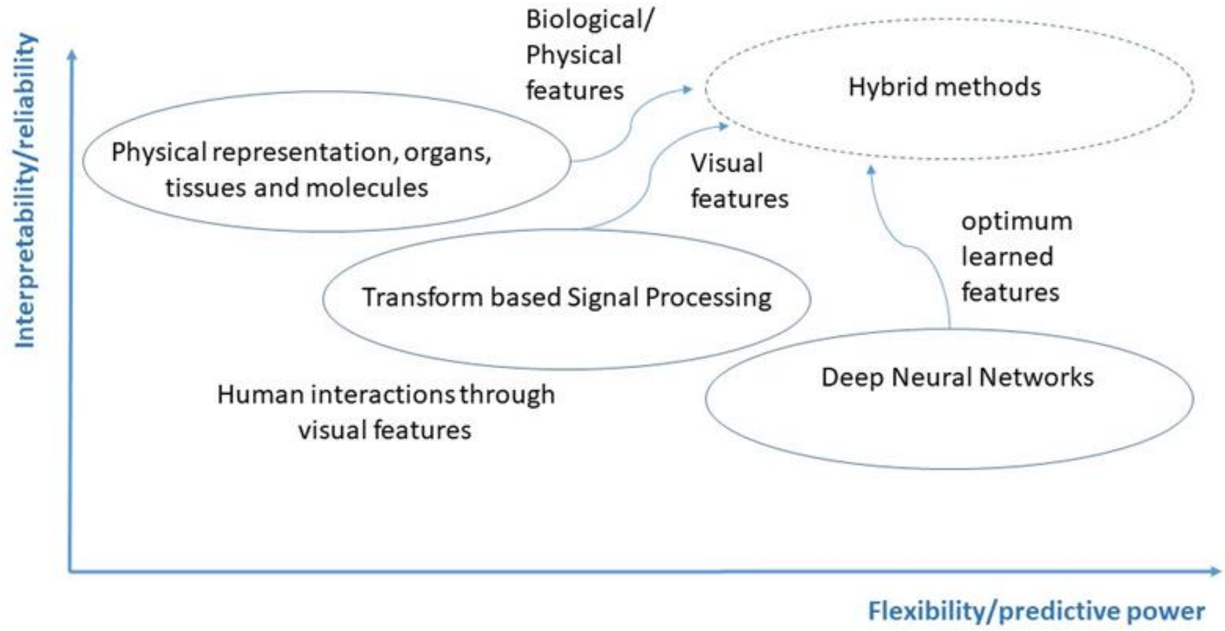

Fingerprints (physical) latent space

• Rensonnet, G., Scherrer, B., Girard, G., Jankovski, A., Warfield, S. K., Macq, B., ... & Taquet, M. (2019). Towards microstructure fingerprinting: Estimation of tissue properties from a dictionary of Monte Carlo diffusion MRI simulations. NeuroImage, 184, 964-980.

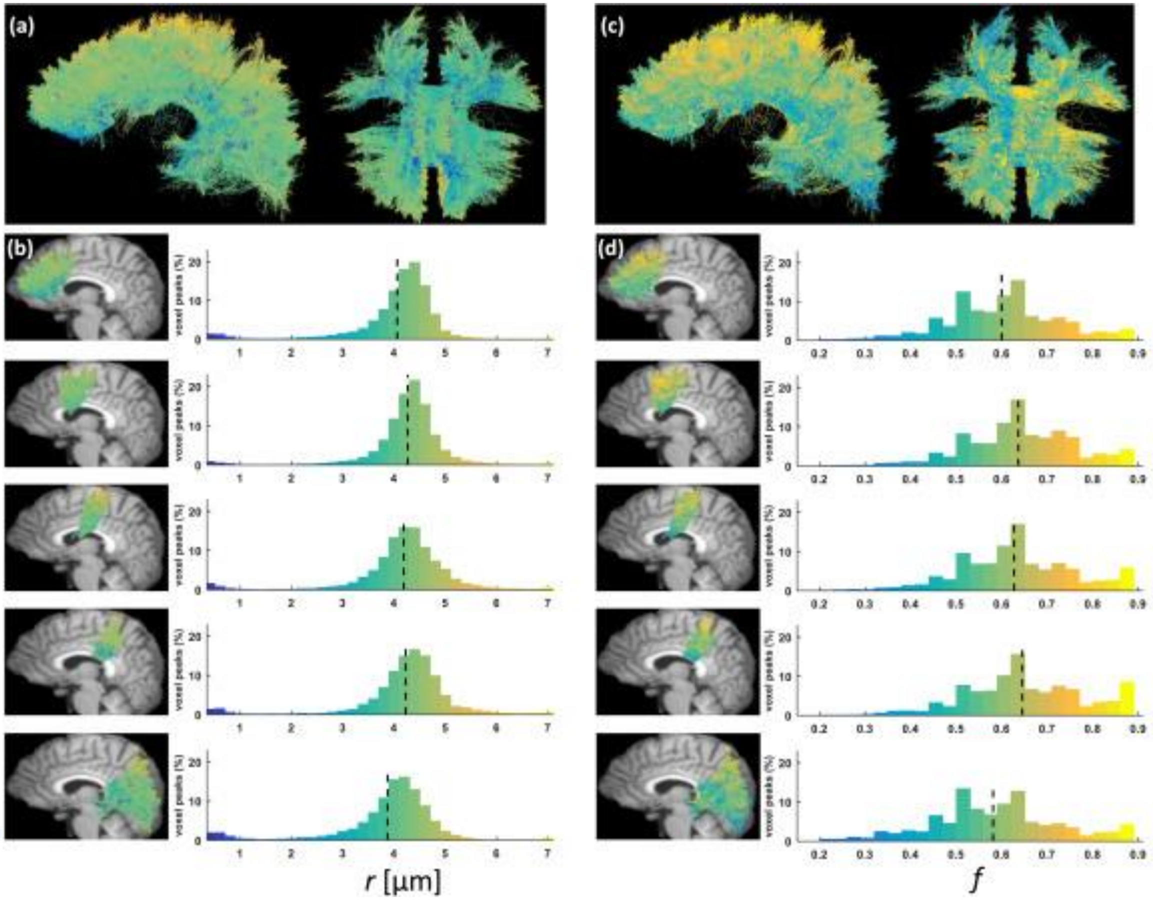

Visual feature (actionable) latent space

Motion measurements 4DCT and IRM (coronal) after co-registration (2D on 3D)

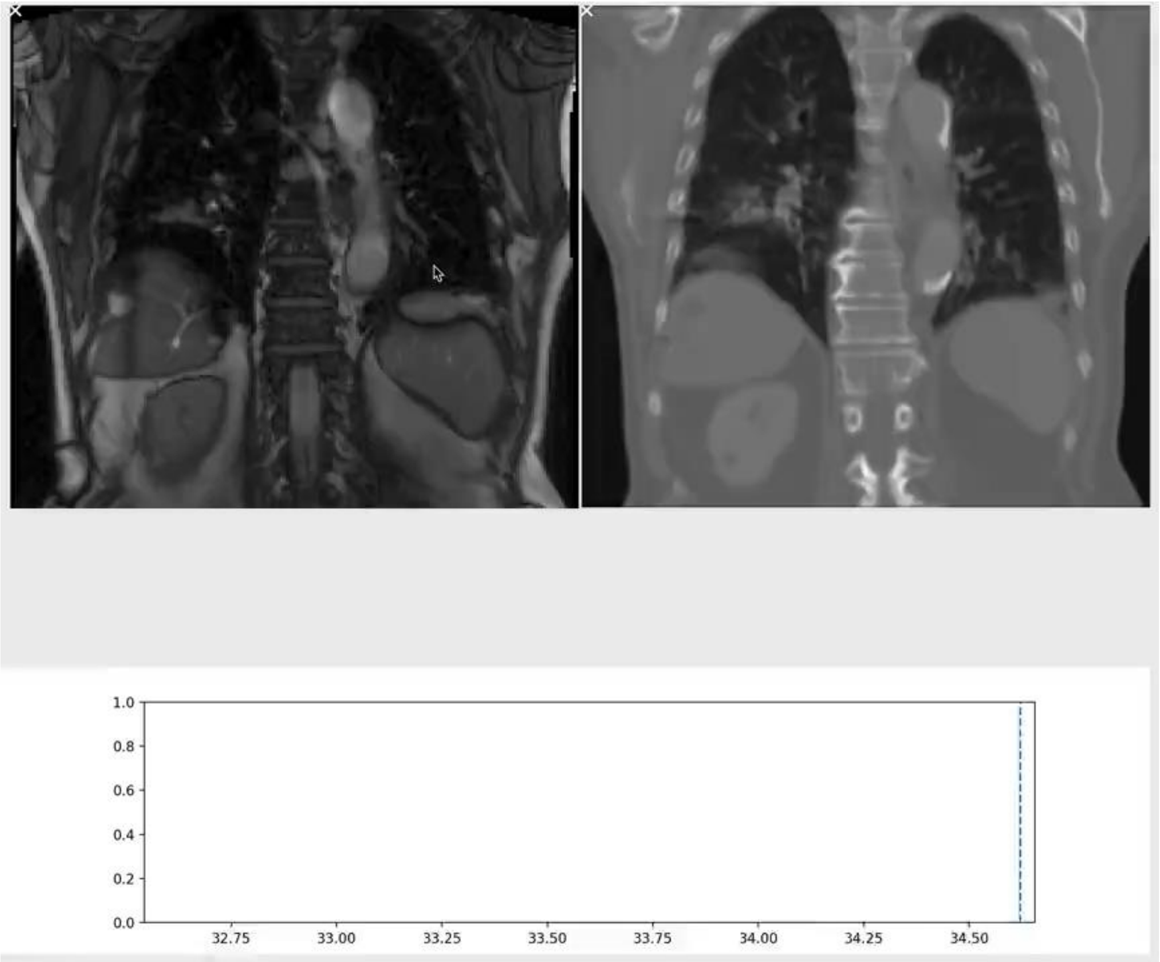

Validation

a. Comparison with motion at the same position
b. Comparison with motions at other positions

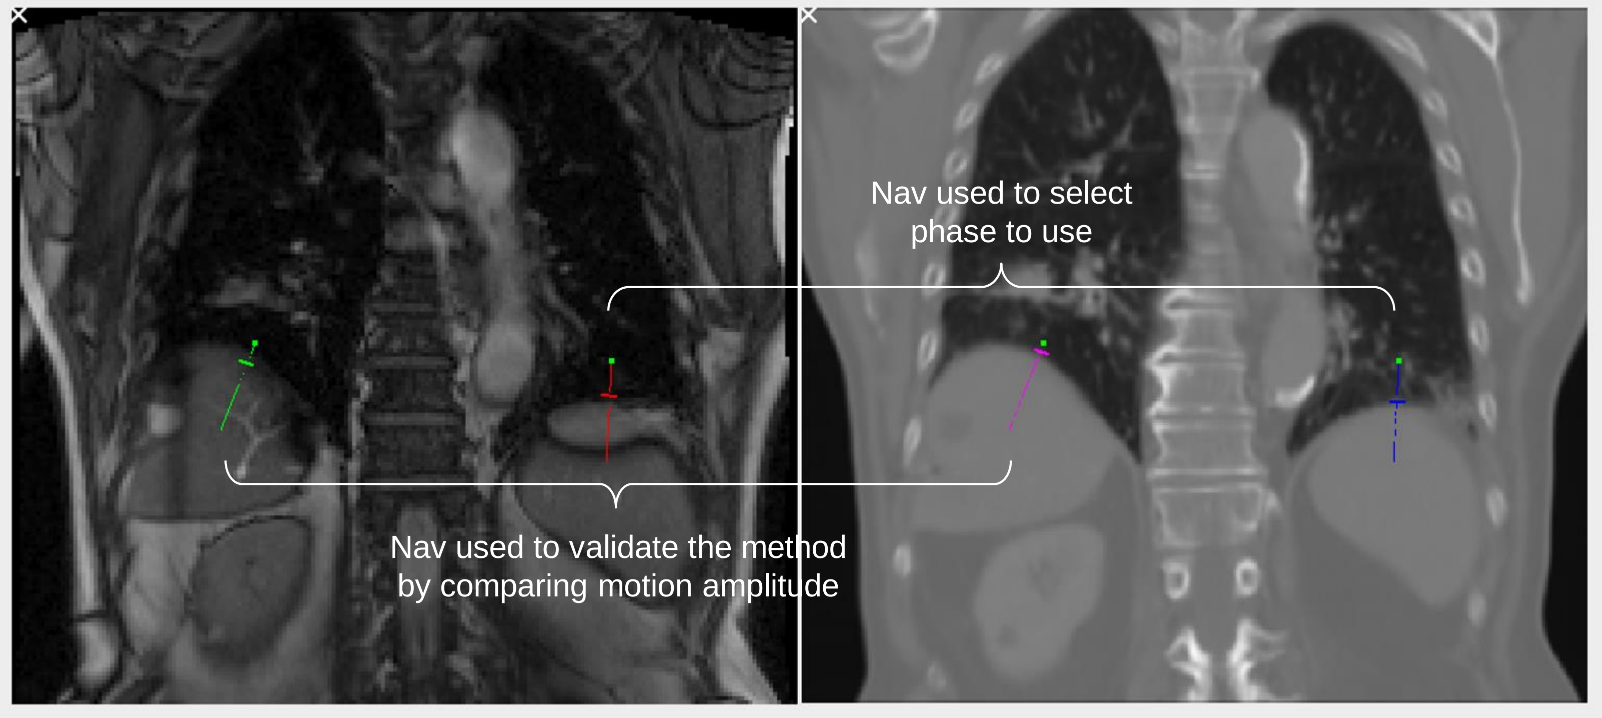

Complete example with dose delivery observation

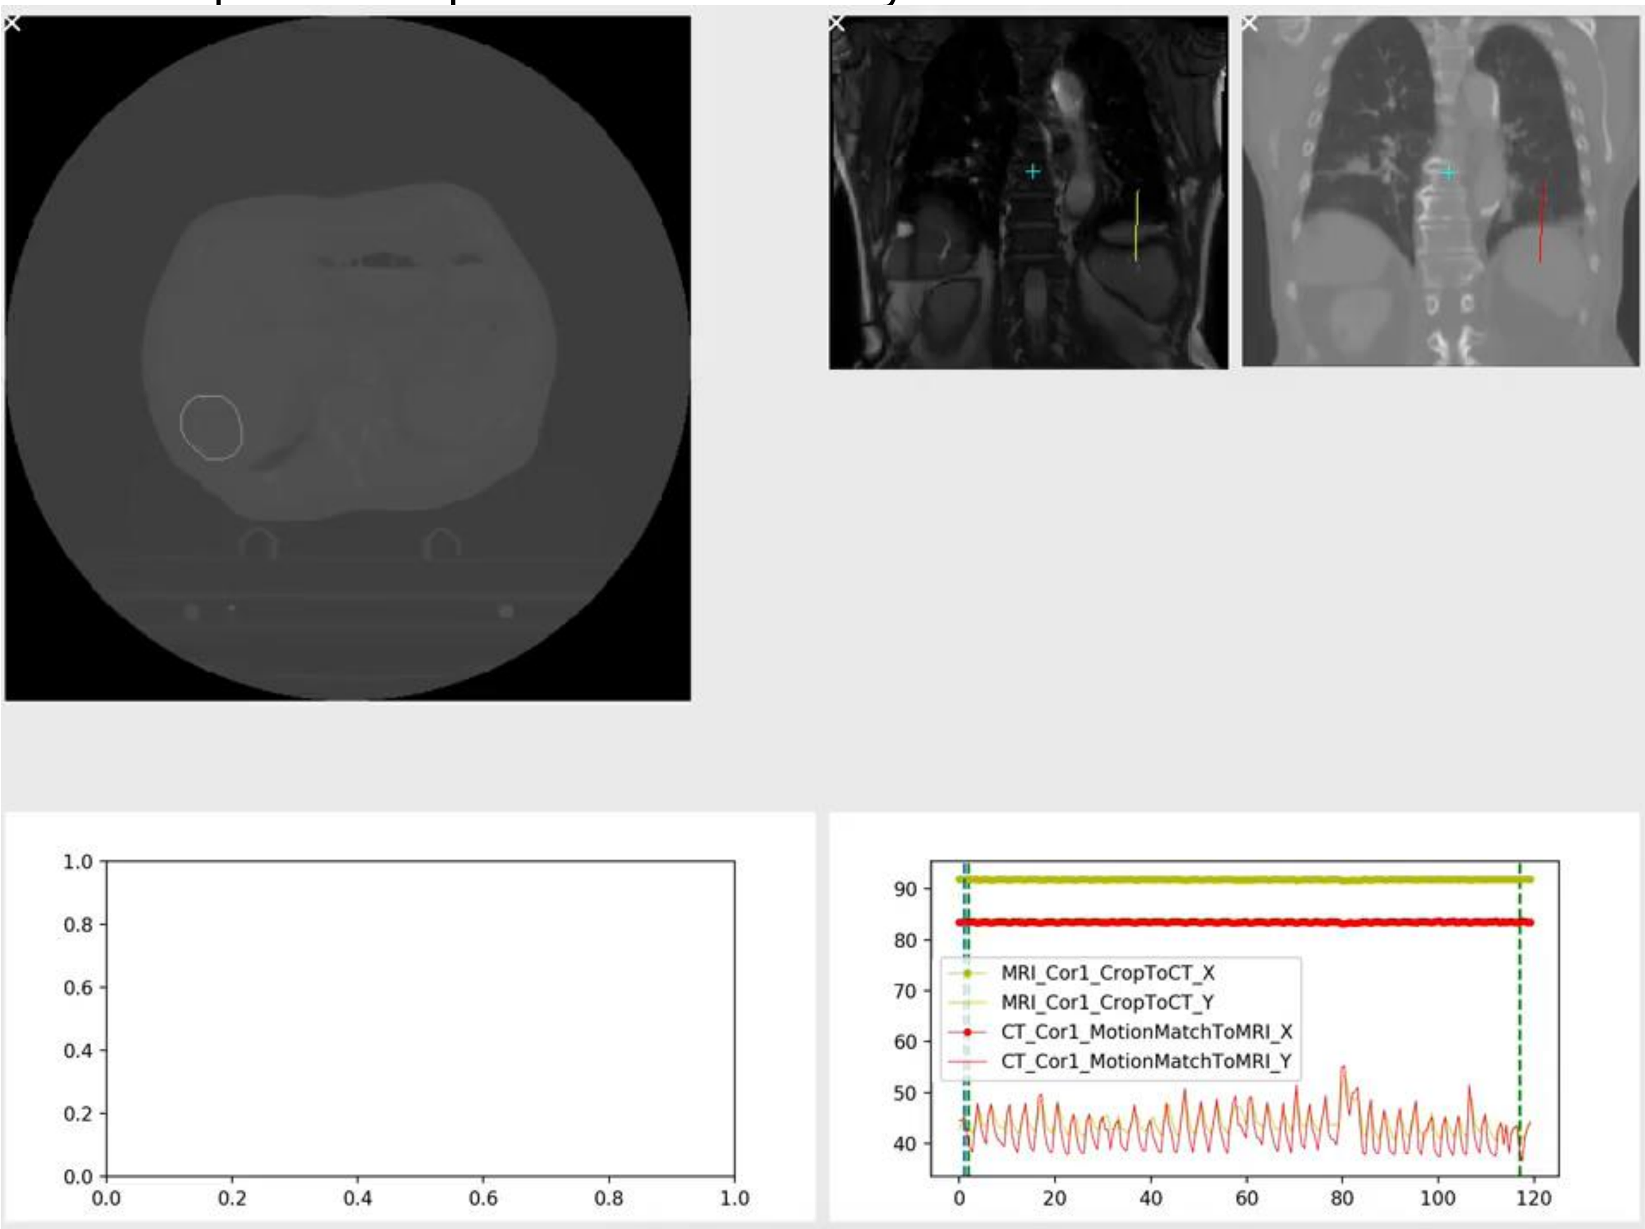

Deep Learning latent space

Measuring anatomical variations between treatment sessions would improve dose conformity

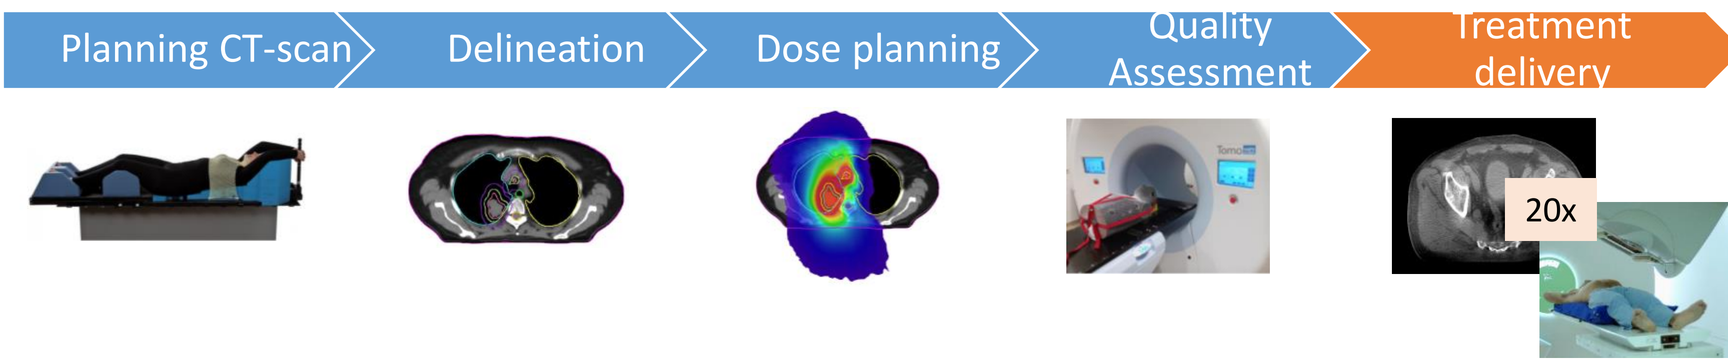

Measuring anatomical variations between treatment sessions would improve dose conformity

Planning CT-scan → Delineation → Dose planning → Quality Assessment → Treatment delivery

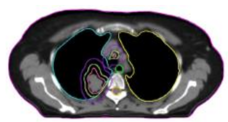

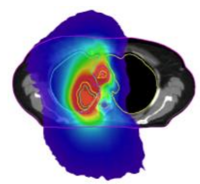

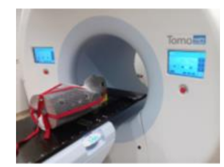

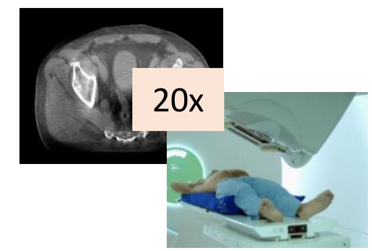

Problem	Question
Scarcity of annotated CBCTs to train a deep neural network	Add (abundant) annotated CTs in training set?

Methods: The network architecture is u-net

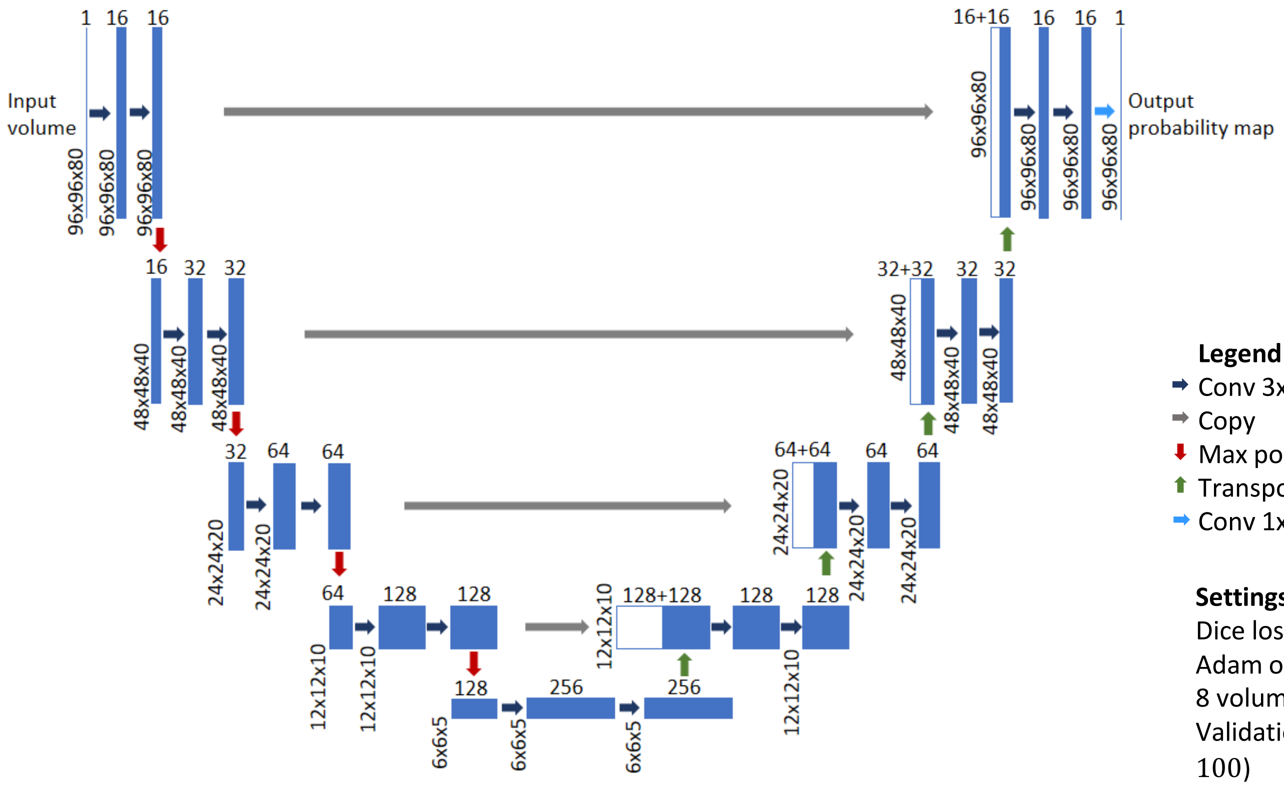

Legend
- Conv 3x3x3, ReLu, « same » padding
- Copy
- Max pooling 2x2x2
- Transpose conv 2x2x2, « same » padding
- Conv 1x1x1, Sigmoid

Settings
Dice loss
Adam optimizer ($LR = 10^{-4}$)
8 volumes per batch
Validation set to earlystop training ($n_{\text{max epochs}} = 100$)
Data augmentation (shift, rotation, shear)

Adapted from Ronneberger *et al.*, 2015

Performance assessment

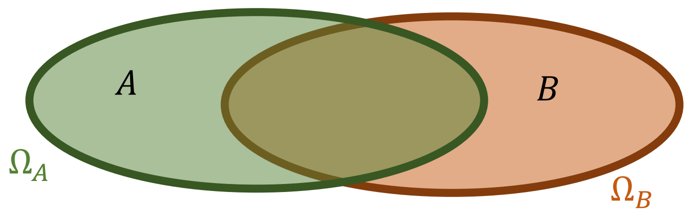

Dice similarity coefficient

$$\mathrm{DSC} = \frac{2|A \cap B|}{|A| + |B|}$$

Jaccard index

$$\mathrm{JI} = \frac{|A \cap B|}{|A \cup B|}$$

Symmetric mean boundary distance

$$\mathrm{SMBD} = \frac{\overline{D}(A,B) + \overline{D}(B,A)}{2}$$

where $D(A,B) = \left\{ \min_{x \in \Omega_B} \|x - y\|, y \in \Omega_A \right\}$

Comparison baselines

Deformable image registration

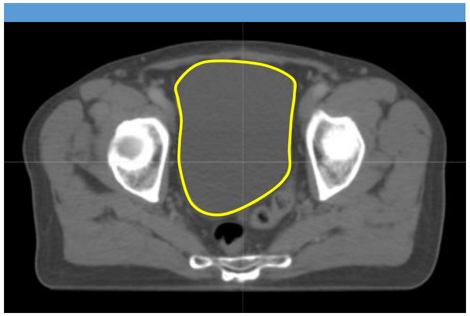

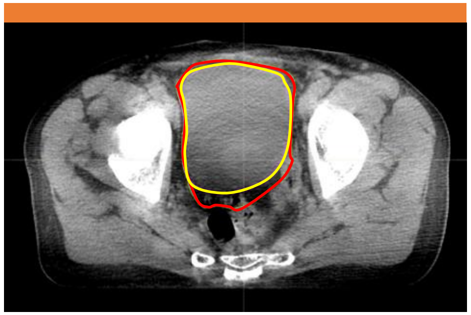

Results: Our approach outperforms a state-of-the-art DIR-based software on a representative patient

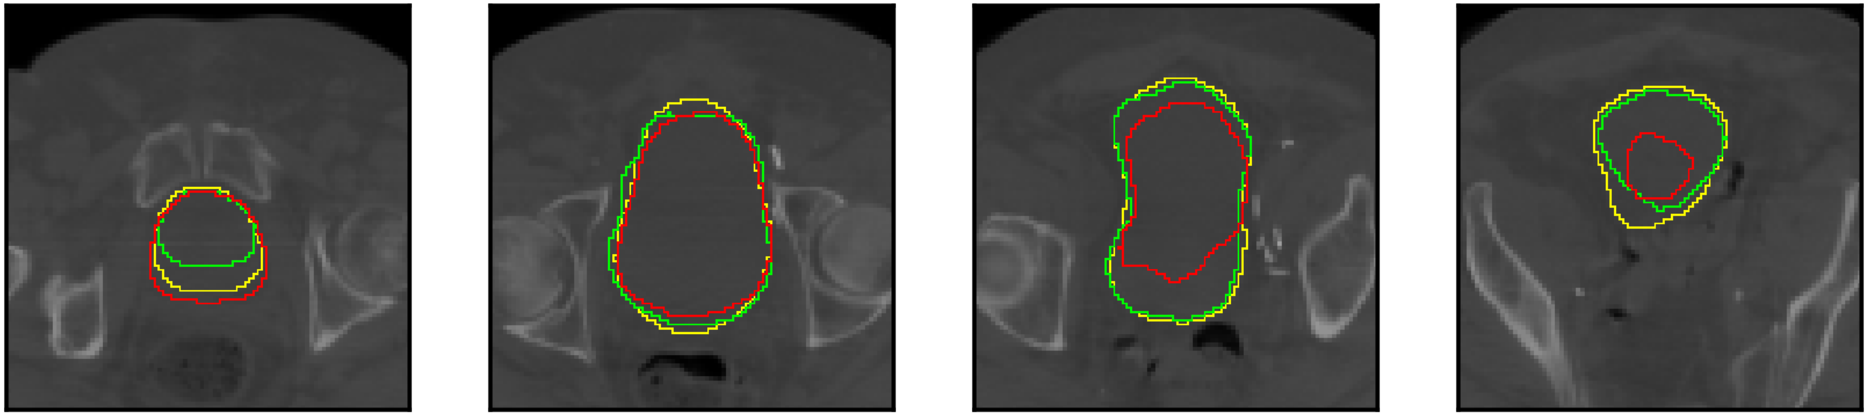

Ground truth segmentation
Deformable image registration, Raystation (DSC = 0.788)
U-net (DSC = 0.892, setting $n_{CBCT} = 32$, $n_{CT} = 64$)

Results

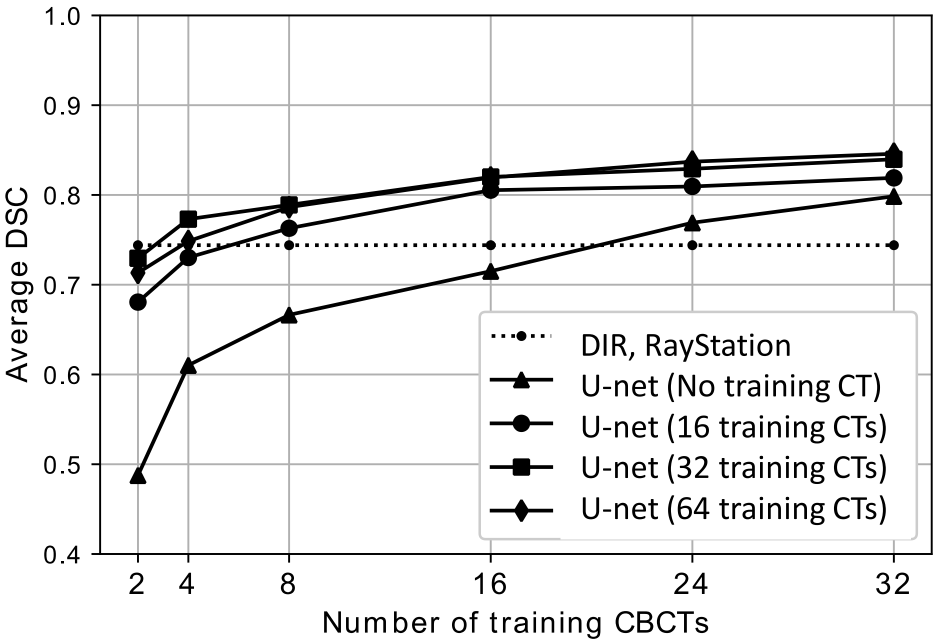

1.0
0.9
0.8
0.7
0.6
0.5
0.4
Average DSC
2
4
8
16
24
32
Number of training CBCTs
DIR, RayStation
U-net (No training CT)
U-net (16 training CTs)
U-net (32 training CTs)
U-net (64 training CTs)

3 latent spaces cooperating in a multi-agent approach (incl HITL)

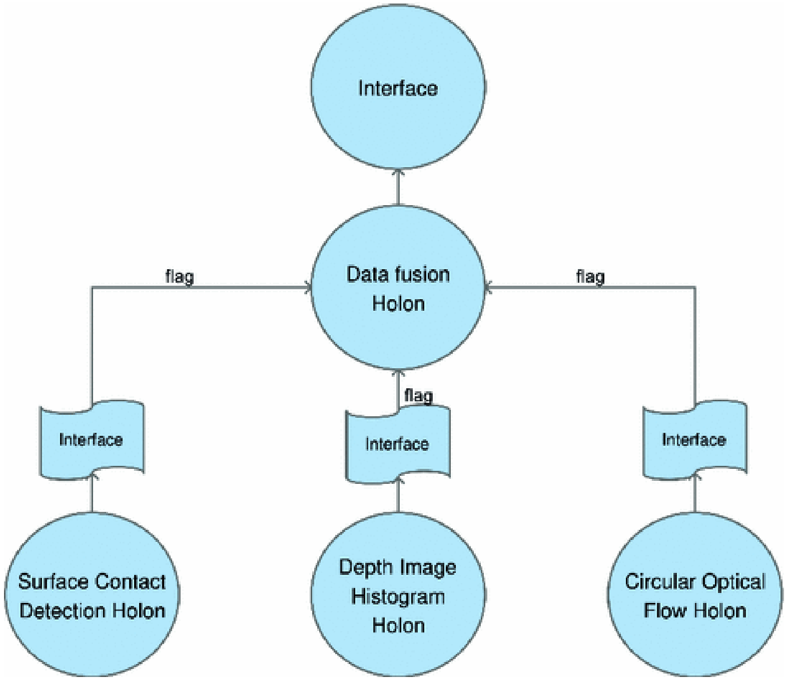

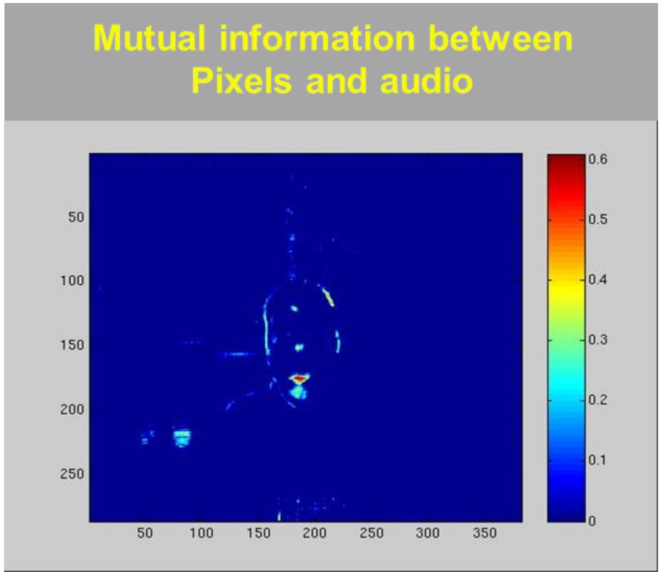

Research assumption 2: Byzantine learning for sharing data and expertise

- Need of integrative coalitions
  - -To share data (privacy and relevance)
  - -To explore complementarity, redundancy and equivalence of the algorithms
  - -To asssess co-evolution of algorithms and human expertise

  - -By the use of consensus mechanisms (Federated Byzantine Agreements- blockchain)

The needs to better use Deep L

-Coalitions for Image Processing

-Distributed machine learning for larger data sets -Trusted Image Processing through Integrative Coalitions

- Security (blockchains)

-Reliability (mutimodality-multiagents)

-Human in-the-loop (regular update-how to poll)

The Oncoradiomics Model (Ph. Lambin)

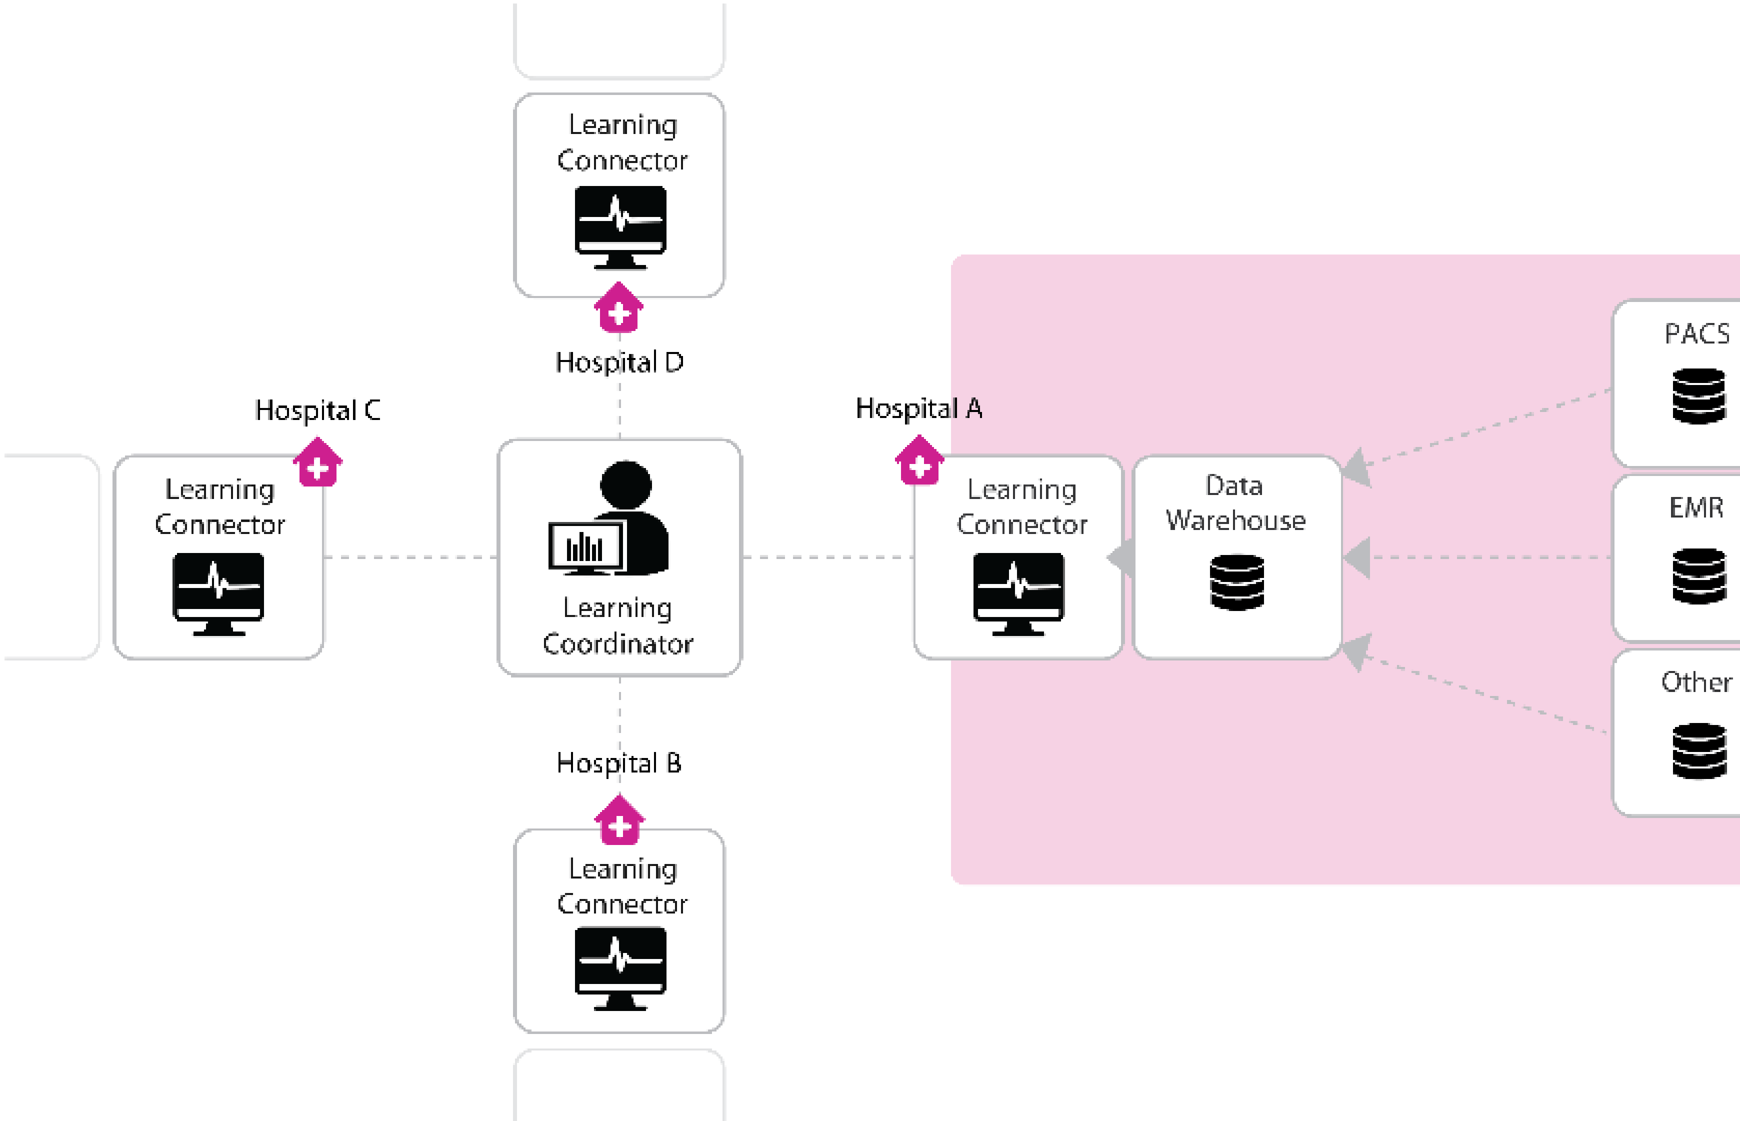

Distributed learning: an abundant litterature

- Distributed SVM: convergence equivalent to central learning can be proven
  - -**Boyd**, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." Foundations and Trends® in Machine learning 3.1 (2010): 1-122
  - -Forero, P. A., Cano, A., & Giannakis, G. B. (2010). Consensus-based distributed support vector machines. Journal of Machine Learning Research, 11(May), 1663-1707.
- Distributed DNN – Federated learning convergence similar to central learning can be shown
  - -McMahan, B., & Ramage, D. (2017). Federated learning: Collaborative machine learning without centralized training data. Google Research Blog, 3.

Security requirements

Challenge 1

Data privacy of the datasets used for the training (leakage effect of the gradients) : working by batches- differential privacy is the "crypto" model

Challenge 2

Protection of the model against degradation by training on inadequate data: steps validation by the coalition and blockchained public ledger with hash of the iterative versions of the model

Challenge 3

Confidentiality of the model and the gradients: homomorphic operations and/or access control of the model vault

Challenge 4

Traceability of the model: DNN watermarking

WHERE MIGHT A DISTRIBUTED LEDGER USE CRYPTOGRAPHY?

Initiation and Broadcasting of Transaction	• *Digital Signatures* • *Private/Public Keys*
Validation of Transaction	• *Proof of Work and certain alternatives*
Chaining Blocks	• *Hash Function*

The hash function: SHA (one-way!!)

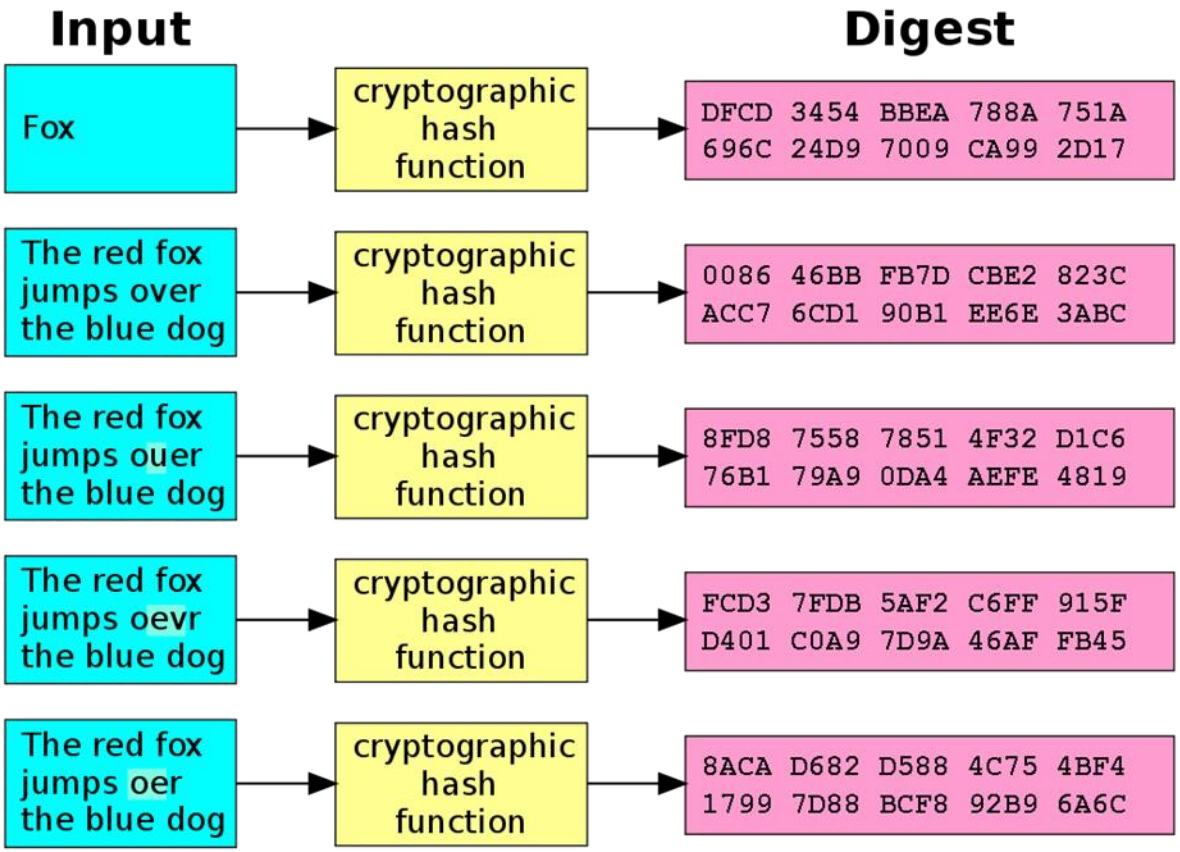

Homomorphic encryption

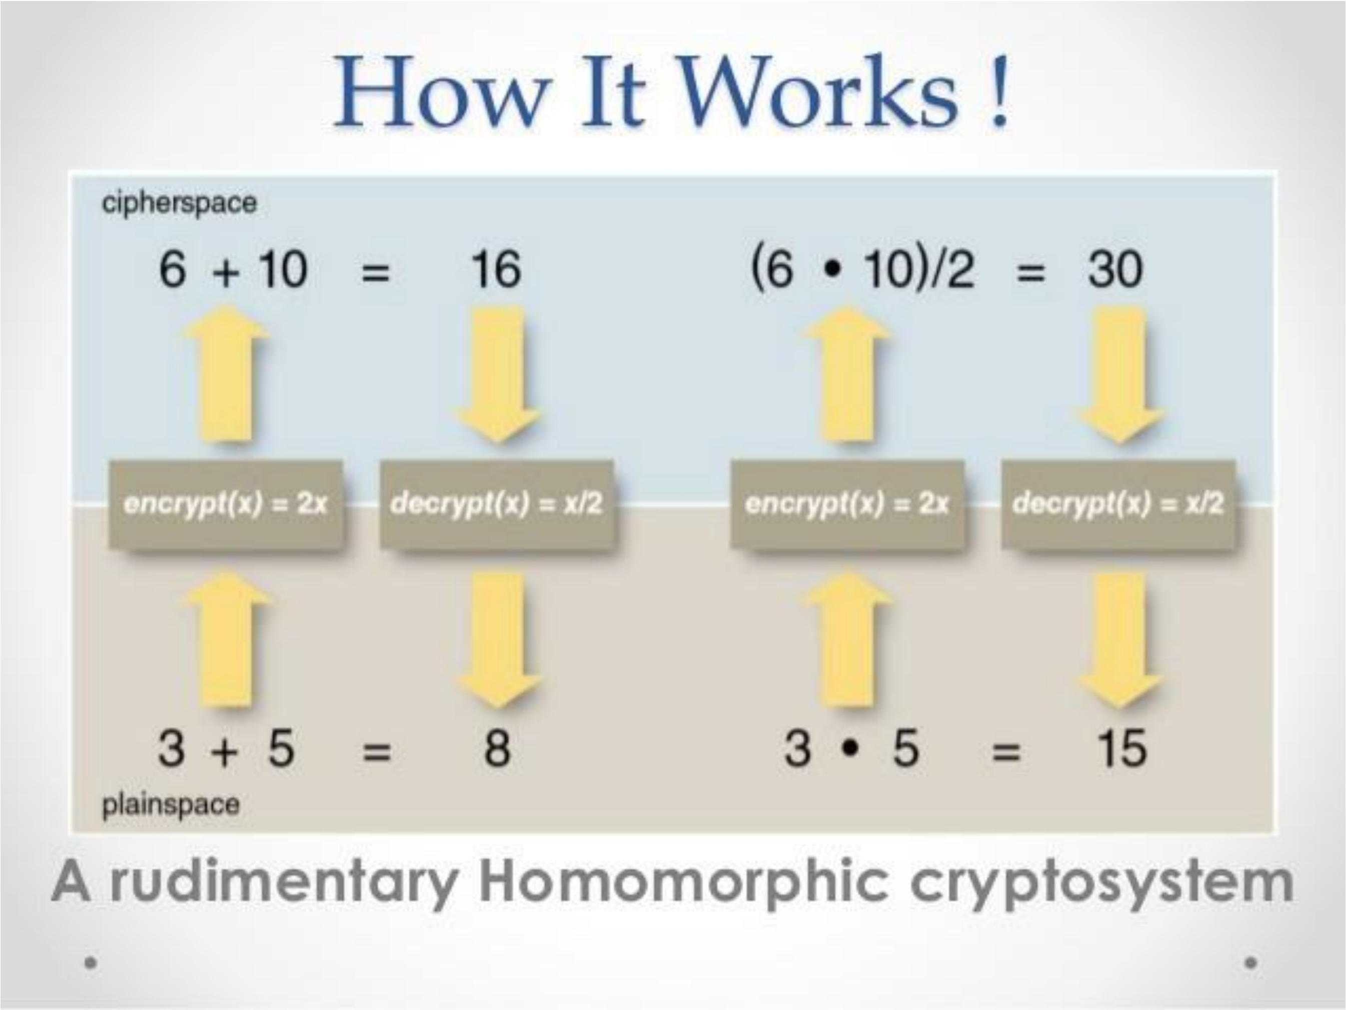

Watermarking

Secret marks in audivisual contents:

-Authentication

-Copyright

-Fingerprinting

Watermarks can be embedded into DNN:

-Uchida, Y., Nagai, Y., Sakazawa, S., & Satoh, S. I. (2017, June). Embedding watermarks into deep neural networks. In Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval (pp. 269-277). ACM.

Scalable security architectures for trusted coalitions

TCLearn-A

Learned model is *public*
Each member is accountable for the privacy protection of its own data

Solution to security challenge1
(Data privacy of the datasets used for the training):

Local training of the model by each member with their own datasets

Generated gradients are uploaded and merged with the previous model

Batches of a minimum size to mitigate the long term memory effect

Solution to security challenge 2
(Protection of the model against degradation by training on inadequate data):

Blockchain storing cryptographic hashes of every training step

Federated Byzantine Agreement (FBA) to prevent corrupted increments

Federated Byzantine Agreement

- Two types of test databases:
  global test database (G),
  local test database (L)
- A "general" is randomly selected among the validators
- The "general" creates a new candidate block referencing the new model
- Every validator validates the viability (model) and integrity of this new candidate block
- Each validator broadcasts its opinion (positive or negative)
- The FBA process ends when 2/3 of the validators agree

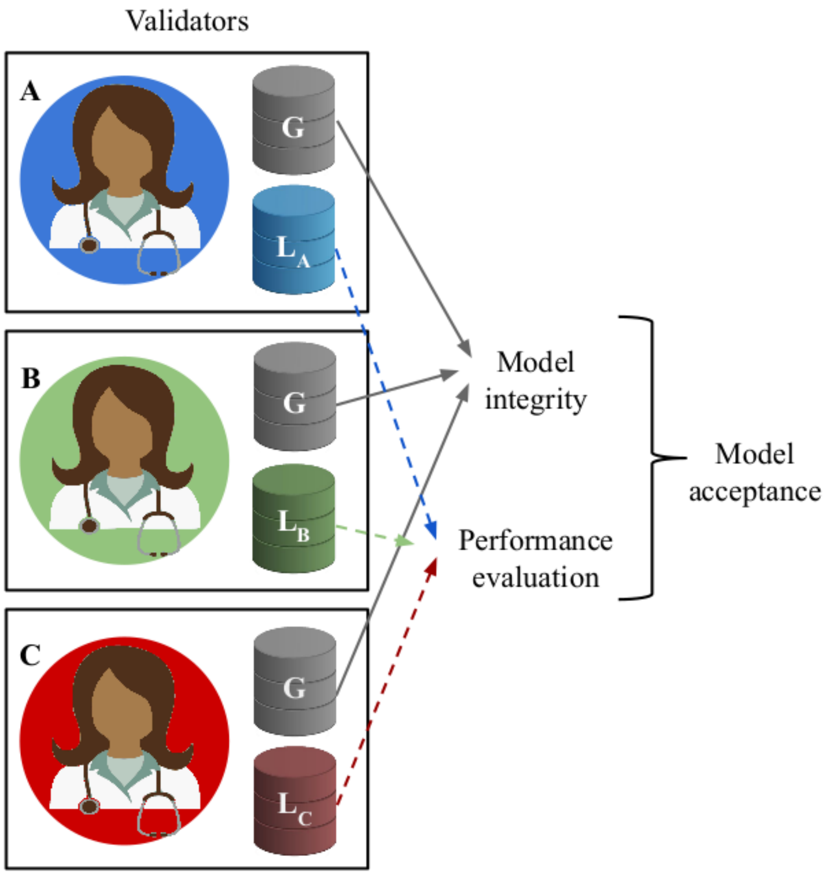

Scalable security architectures for trusted coalitions

TCLearn-B

Learned model is *private*, the members of the coalition trust each other.

Solution to security challenges 1 & 2:

Same as for TCLearn-A

Solution to security challenge 3:
(Confidentiality of the model and the gradients):

Storage of all iterations of the model in an off-chain storage

Iterations only referenced by links in the blockchain

Secure, encrypted transport of the model (using e.g. TLS or S/MIME)

Solution to security challenge 4:
(Traceability of the model):

Access control and audit mechanisms to protect the models and parameters

Scalable security architectures for trusted coalitions

TCLearn-C

The members of the coalition do no trust each other.

Solution to security challenges 1 & 2:

Same as for TCLearn-A

Solution to security challenges 3 & 4:

Storage of all iterations of the model in an off-chain storage

Each member is provided with a homomorphically encrypted model and the corresponding public key, used to encrypt their datasets, by a supervisor

Prediction could be performed locally on encrypted data, but the result must be decrypted by the supervisor

Full traceability since the encrypted model cannot be used without the associated public key, itself associated with the partner which received it

Summary of our blockchained D DNN

- New architecture for distributed learning based on a blockchain using a federated Byzantine agreement
- Performance of the model ensured through shared evaluation of individual contributions (leading to acceptance or rejection)
- Trusted coalitions, actions for updating the model stored on a public ledger implemented as a blockchain
- Three kinds of coalitions with increasing security levels depending on the requirements for the distribution of the model
- Solutions based on effective cryptographic tools and homographic encryption
- Data privacy protection through encryption and off-chain storage

- https://arxiv.org/abs/1906.07690 (Lugan .... Macq)