Machine Learning for Medical Images Processing

Michel Dojat

Part II

michel.dojat@inserm.fr

Comuter vision: Image classification

ImageNet 14 M images
20000 different object categories
2022 91% accuracy

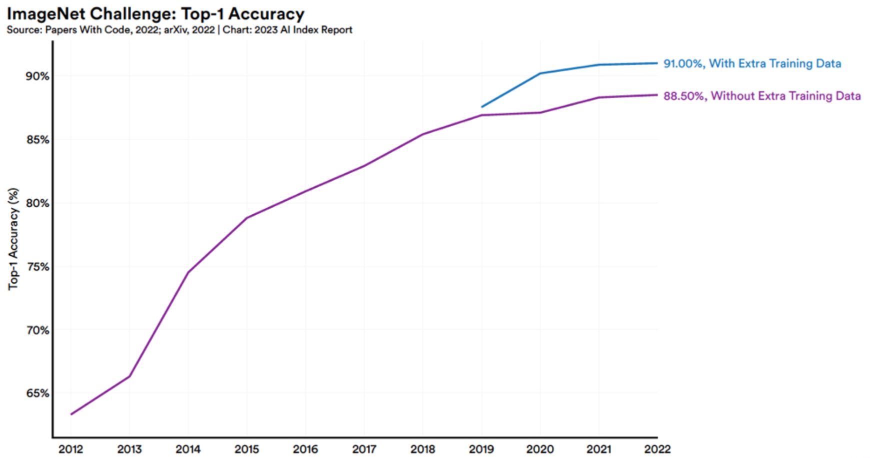

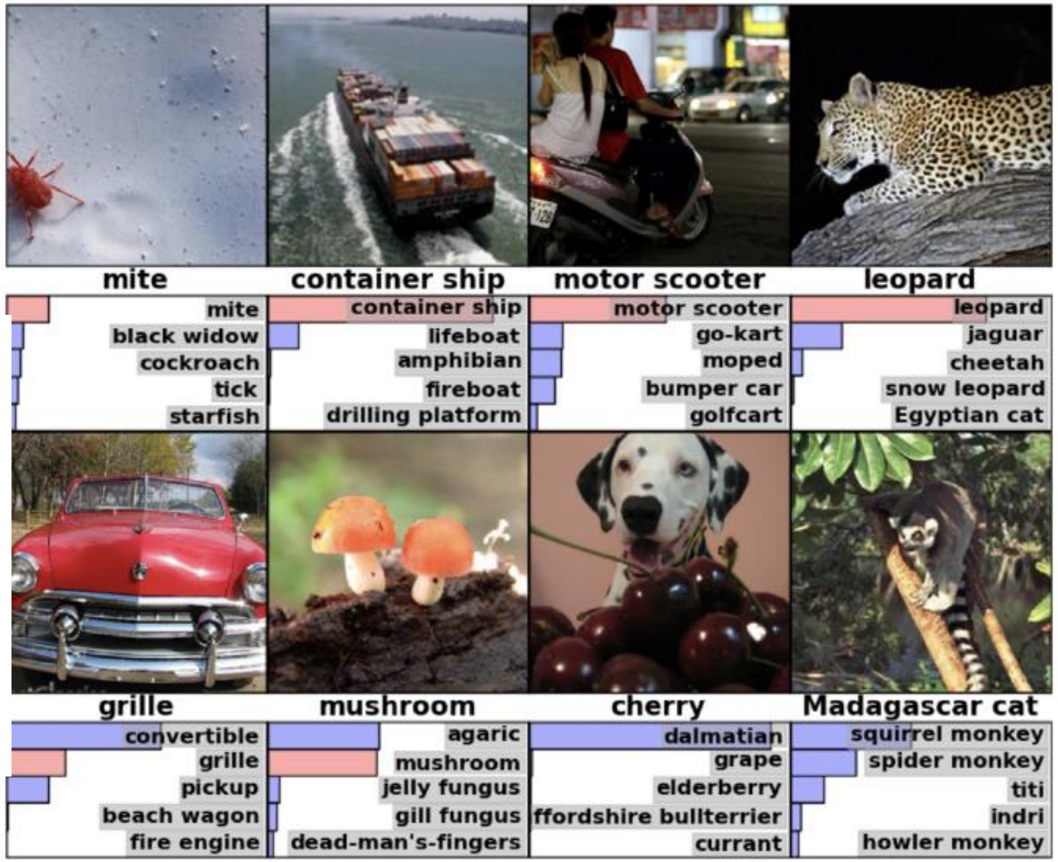

[Krizhevsky et al 2012]

Explosion …

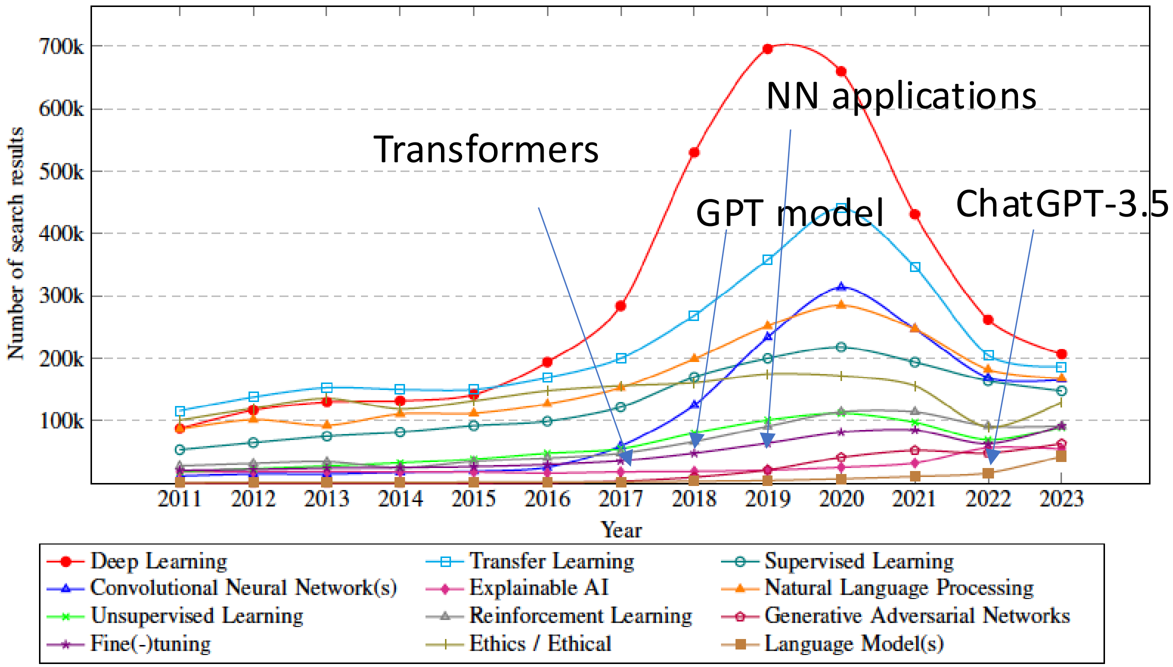

Google scolar "(AI OR artificial OR (machine learning) OR (neural network) OR computer OR software) AND ([specific keyword])

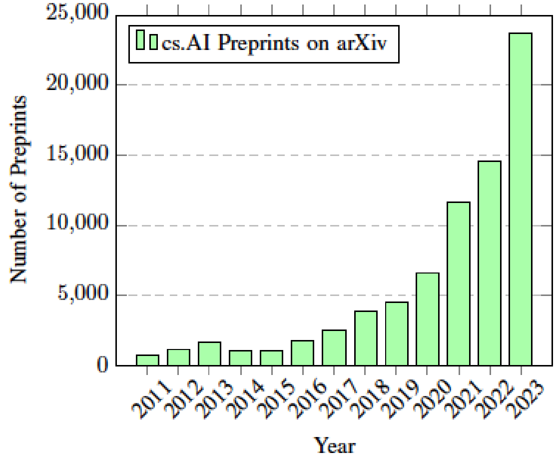

Cs.AI category on arXiv.org

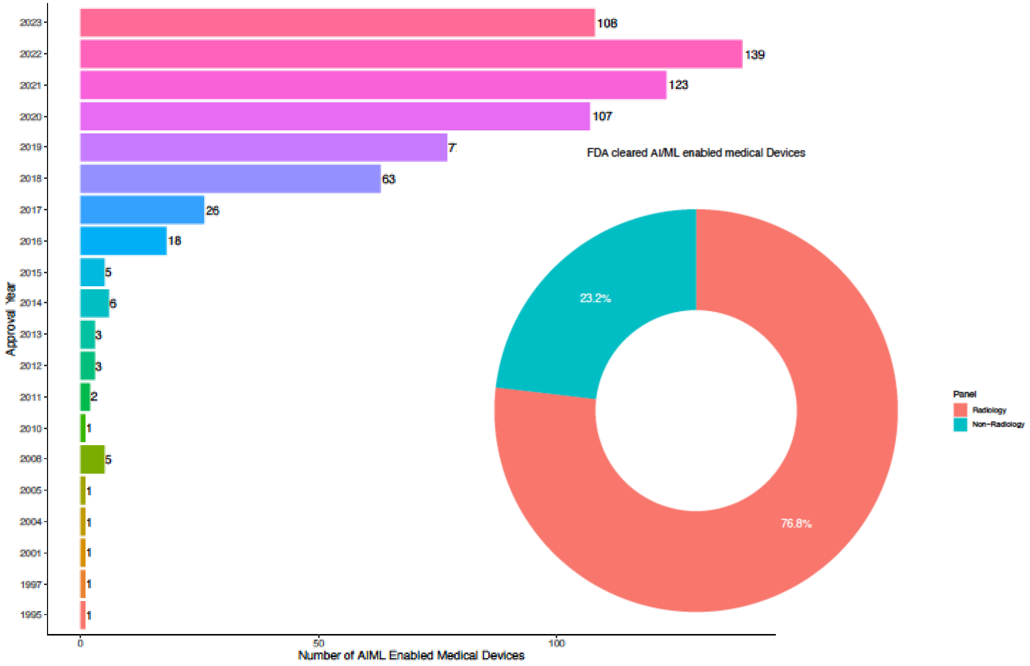

Joshi et al. Electronics 2024

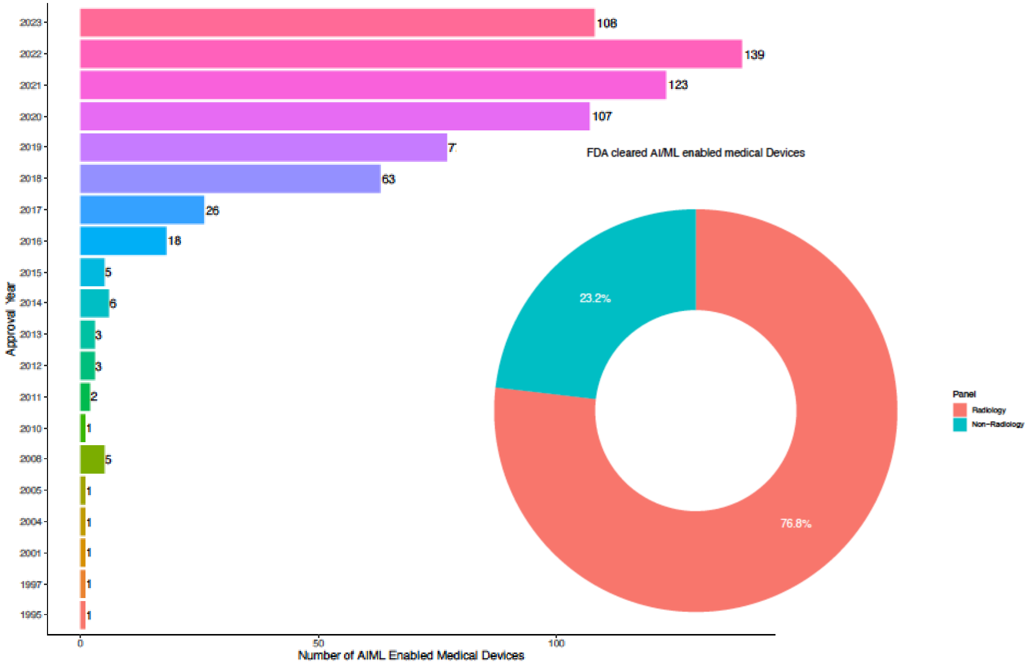

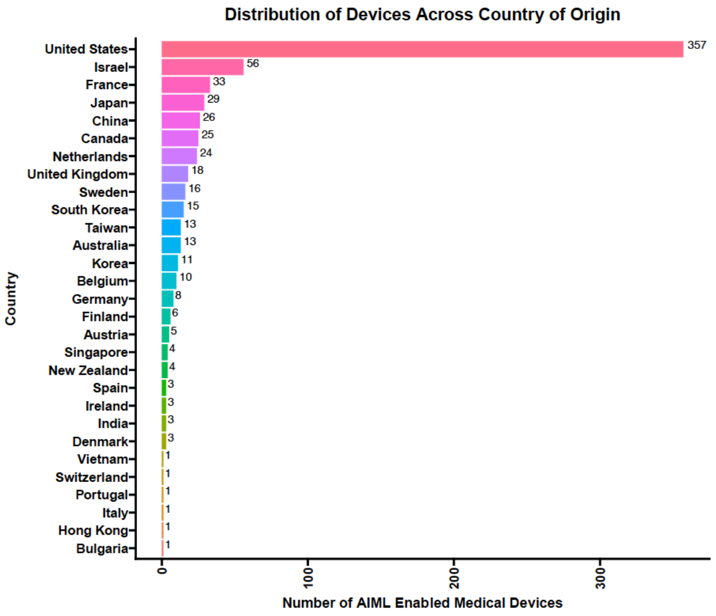

Joshi et al. Electronics 2024

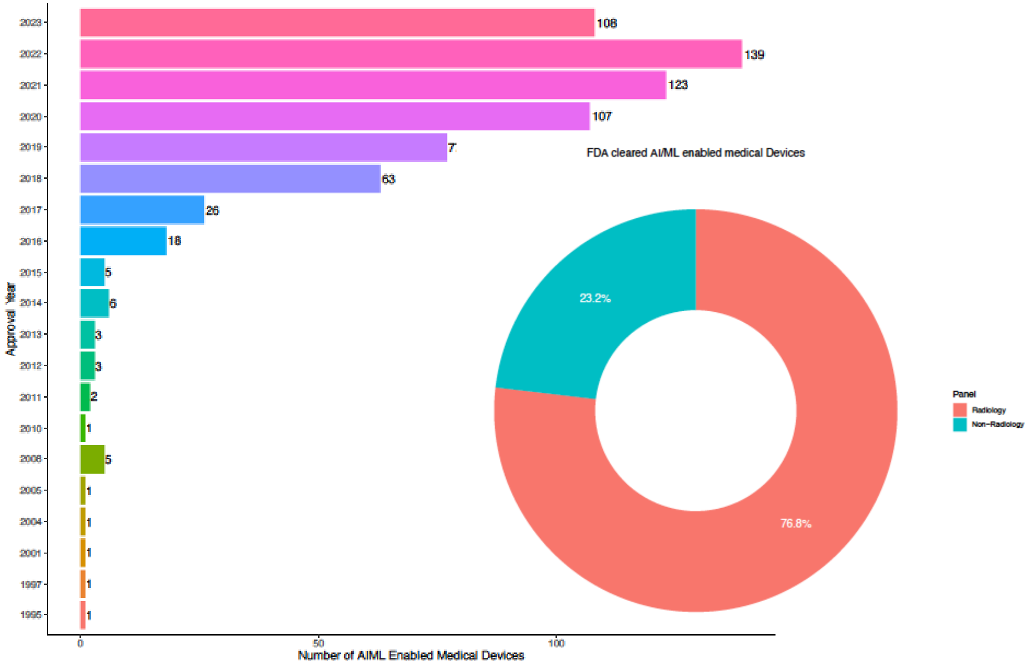

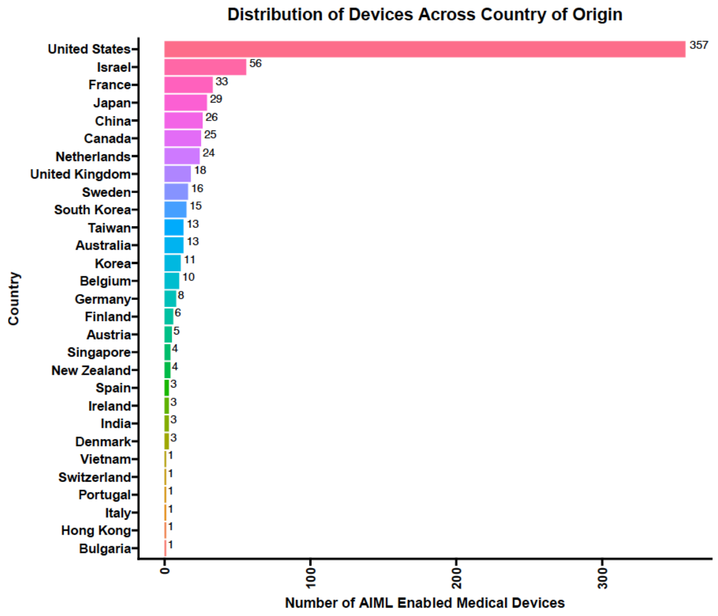

Joshi et al. Electronics 2024

The majority (97%) of FDA approvals was approved through the 510(k)-clearance pathway, relying on the demonstration of substantial equivalence that circumvents the necessity for exhaustive clinical trials (3%).

Automatic segmentation

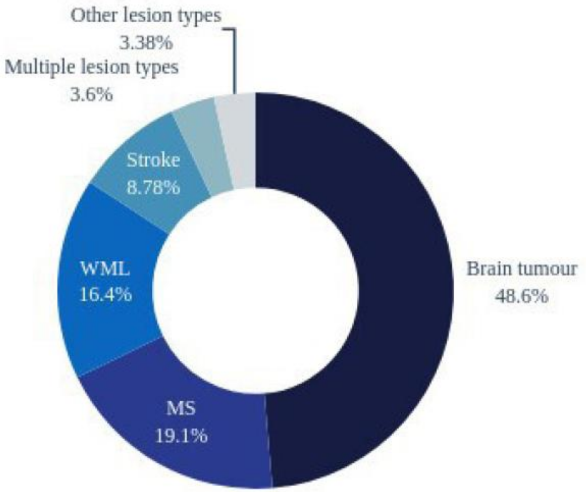

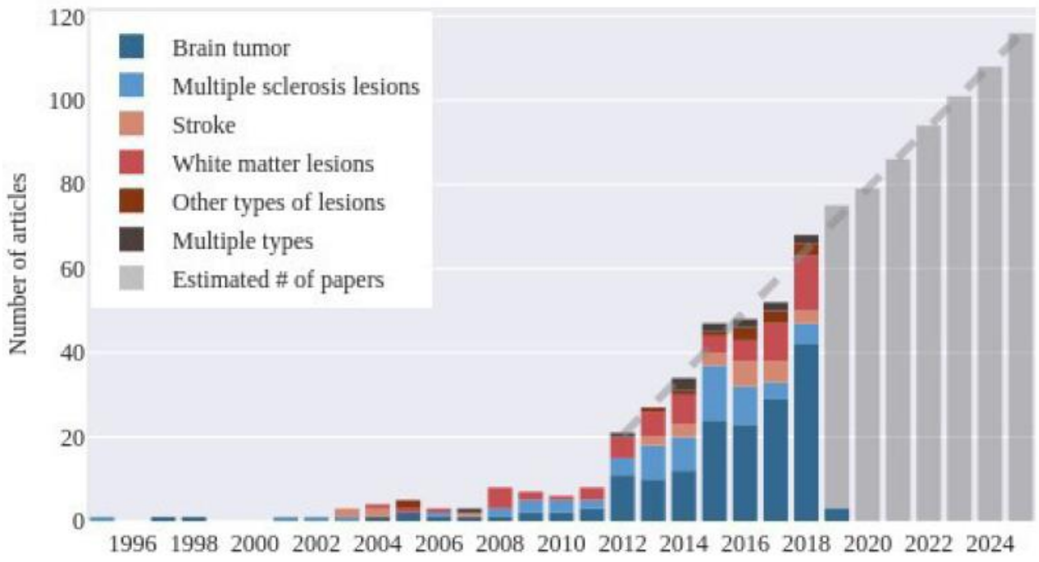

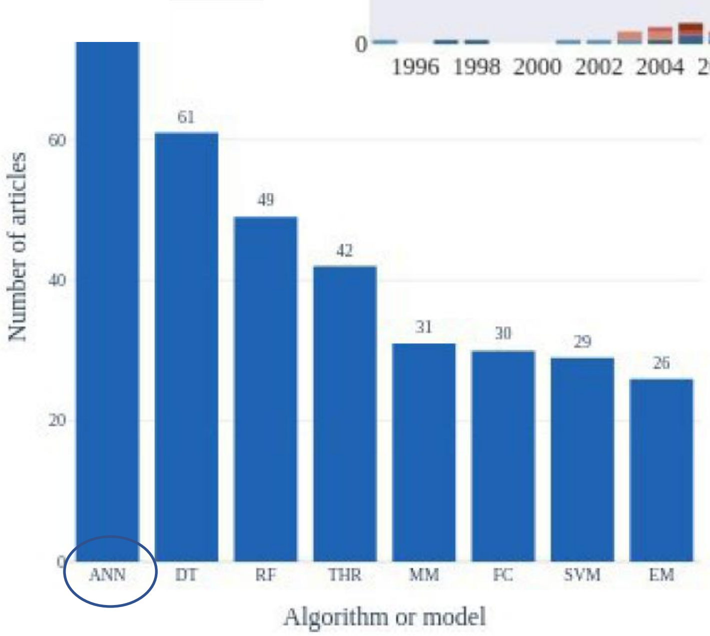

[Gryska et al BMJ *Open* 2

CNN in Radiology

- PubMed 2013-18 ("DL" or "CNN") and ("image" or "imaging" or "radiology")
- 744 art. 180 relevant

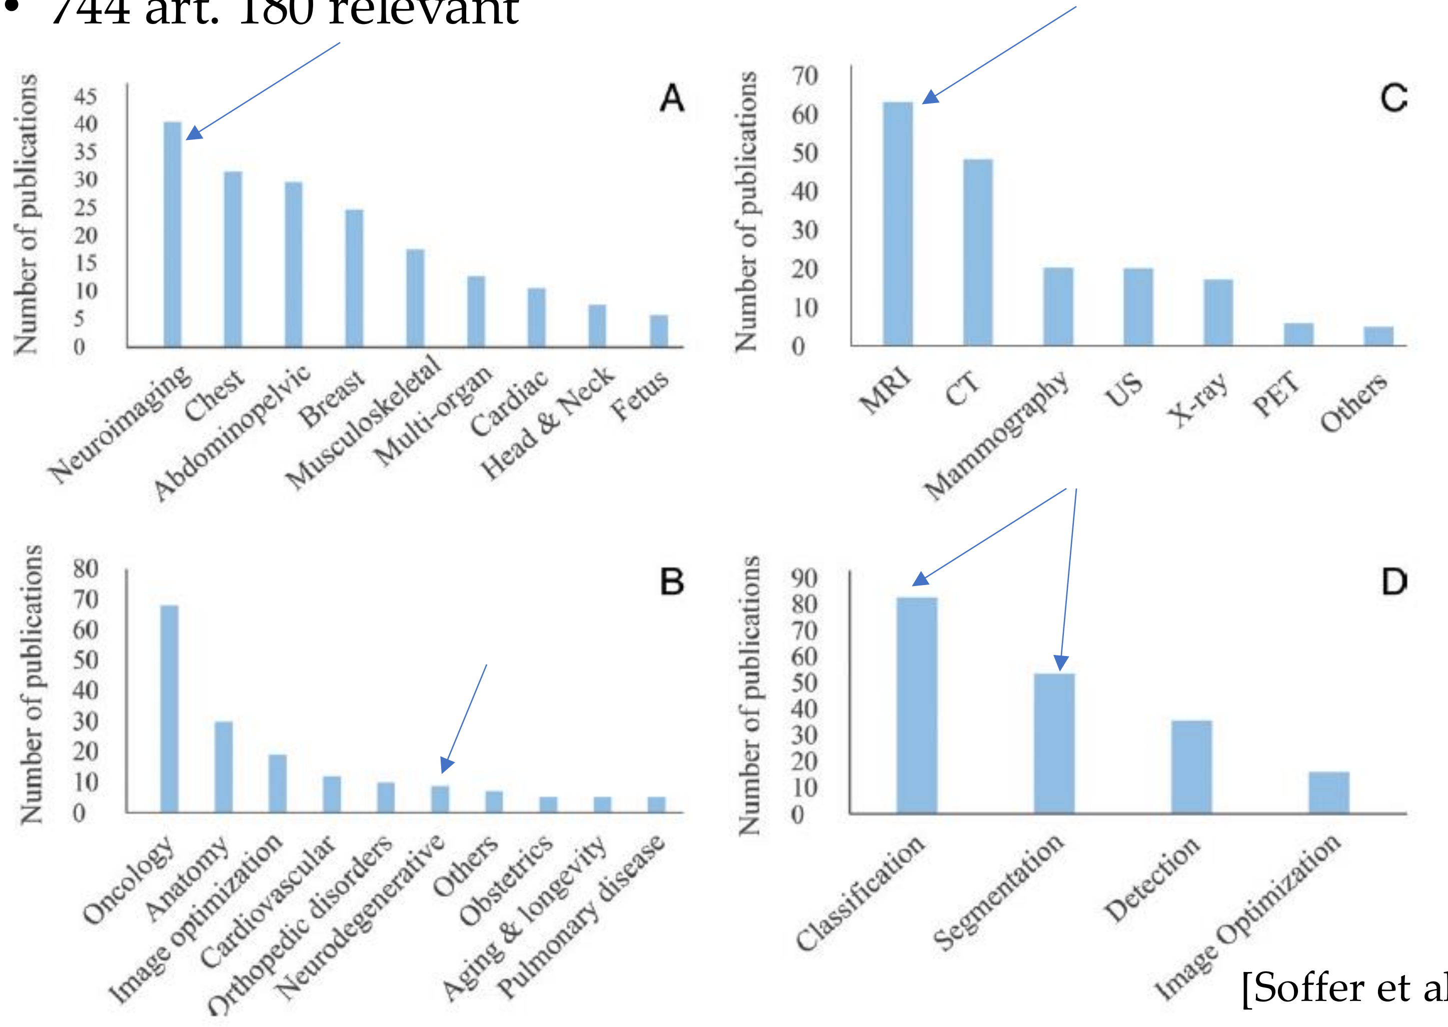

[Soffer et al 2019 Radiology]

Clinical Deployment

Table 4 | Clinical integration of medical AI at different developmental stages

	Areas where AI performance is more reliable than that of a human expert	Areas where AI performance is at the expert level	Areas where AI performance is reasonable	Areas where AI performance is not yet good enough	Areas where the nature of the clinician–patient interaction is fundamentally different from that of the AI–patient interaction
Examples	Serum analyser[144,145]; alert systems (such as drug–drug interaction checkers[146,147])	Assessment of certain radiology images (for example, annotation of cardiovascular MRI images[57,58] or evaluation of X-ray images for distal radius fracture[148]); dermoscopic melanoma diagnosis[149]; fundus photograph evaluation for DR[5,7]	ECG reading[11]	Surgery; full interaction with patients	Emotional support and rapport
Potential clinical integrations	Delegate to AI	AI does the majority of the task, clinicians confirm the diagnosis	AI does a portion of the task (such as screening), clinicians confirm the diagnosis	Clinicians lead the clinical evaluations and intervention, AI assists in routine sub-tasks	Clinicians continue to provide the service

Yu et al Nat Bio Eng 2018

Challenges

Rajpurkar et al 2022 Nat Med

Specific tools & infrastructures

Clinical use

Legal / Societal

Applications

Ethical

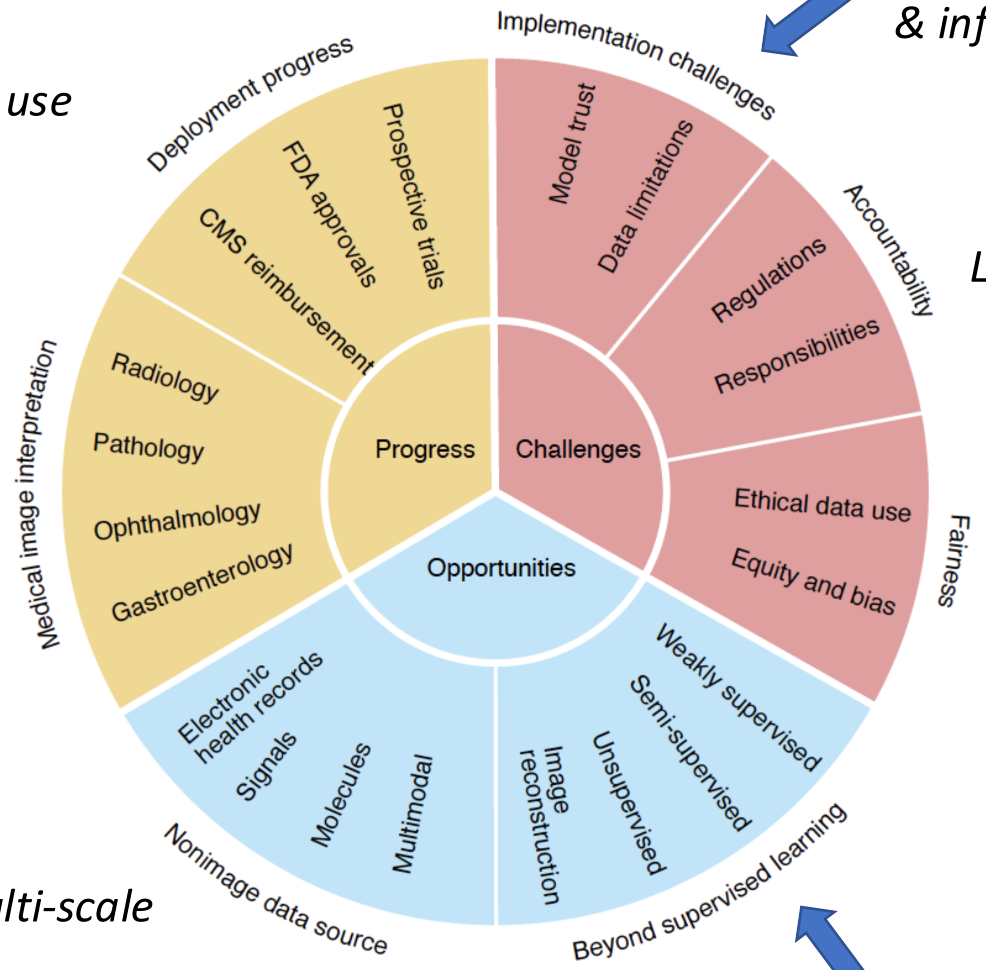

Multi-modal/Multi-scale

Original ML techniques

ML & CNN

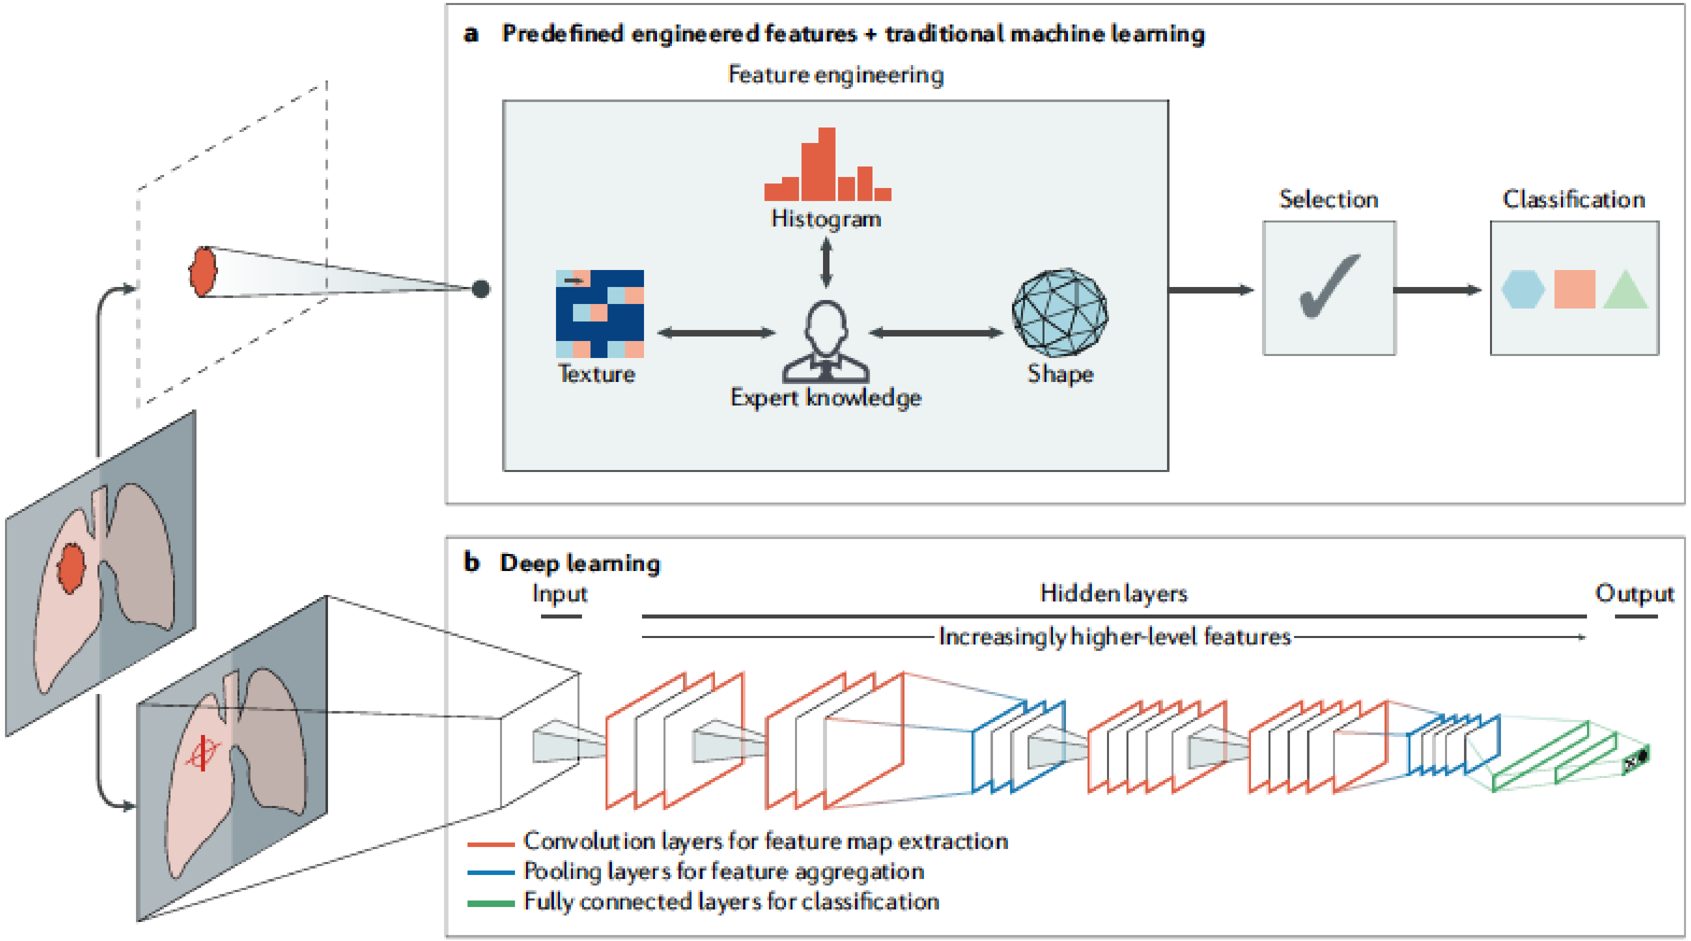

[Hosny et al 2018 Nature]

Life Science applications

- Classification/segmentation/detection

Moeskops et al TMI 2016; Rajchl et al MIDL 2018 (*brain tissue*); Dolz et al. Neuroimage 2018 (*brain structures*); Kamnitsas et al MediA 2016 (*brain lesions*)
Kleesiek et al Neuroimage 2016 (*brain extraction*); Havaei et al MediA 2016 (*brain tumors*)
Suk et al NeuroImage 2014 (*AD/MCI*)
Zhao et al MediA 2017; Suk et al NeuroImage 2016; Kin et al Neuroimage 2016 (*functional brain networks*)
Ciompi et al Scient Rep 2017 (*lung nodules*); **Esteva et al Nature 2017** (*skin cancer*)

- Synthetic image generation

Nie et al Miccai 2017 (*MR-CT*); Liu et al Radiology 2017 (*MR-CT*); Zhao et al Media 2018 (*retinal images*)

- Predictive models

Polpin et al Nat Bio Eng 2018 (*cardiovascular risk*); Miotto et al. Scient Rep 2016 (*deepPatient*)

- Processing

Rajchl et al MIDL 2018 (*automatic process for segmentation*) denoising, registration (see Litjens Media 2017)

- Retrieval

Anavi et al Spie 2016

- Brain encoding/decoding models

Classification example

- Training set: train the model
- Validation set: Refine hyperparameters, avoid overfitting
- Test set: new set of data to verify the performances

- 3 hidden layers with 100 nodes
- ReLu
- Softmax output

Architecture example

$$y_i(\mathbf{x}) = g(\mathbf{W_i x} + \mathbf{b_i})$$

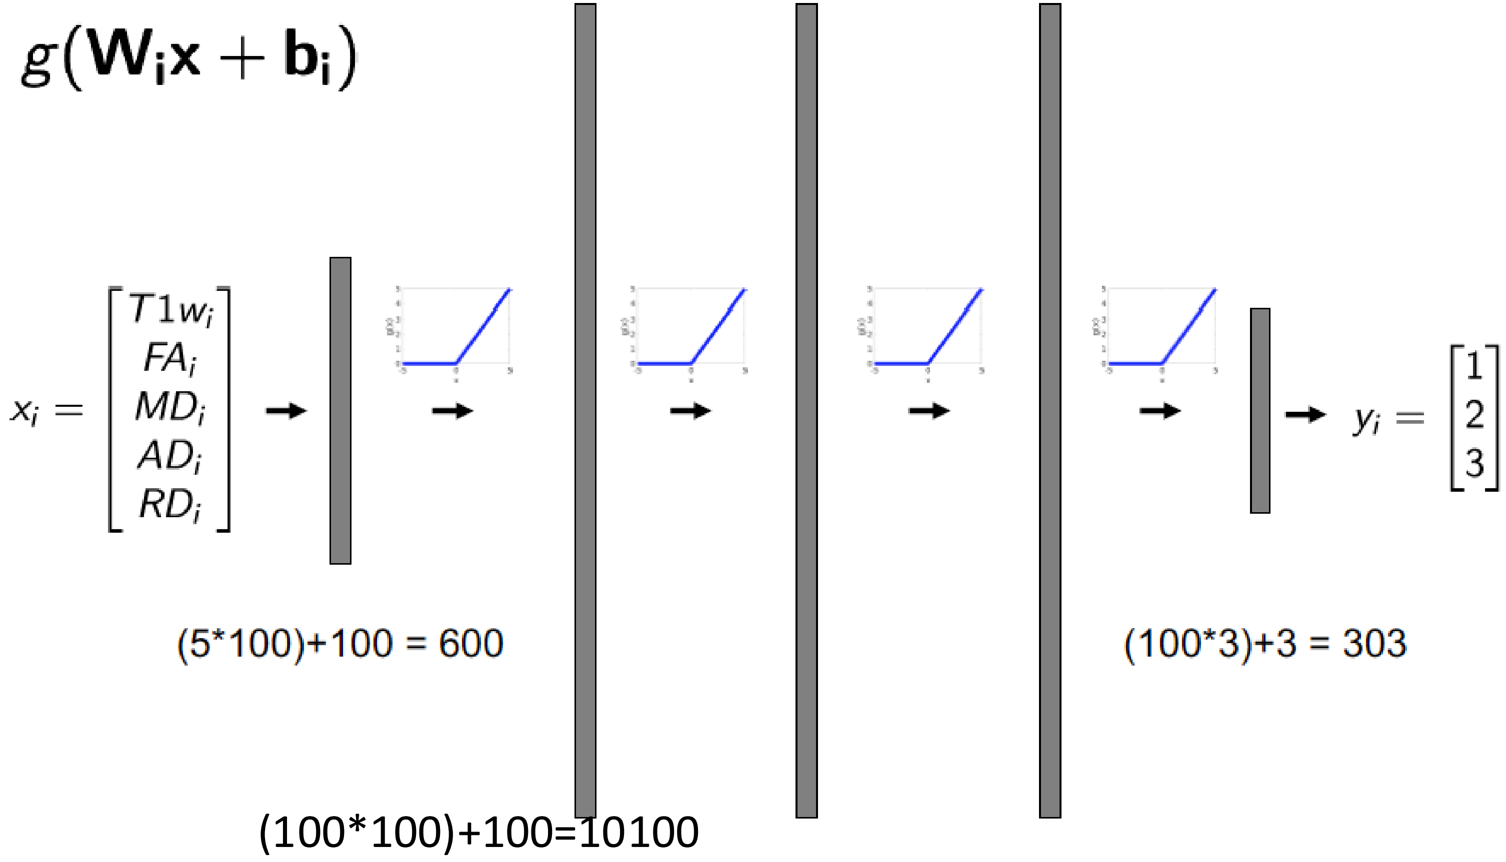

10100x3+600+303=31203 parameters to be estimated

From Knoll Ismrm 2018

Implementation

Matlab

```
%% NN with 3 HIDDEN LAYER NEURONS
nElements = 100;
nLayers = 3;
inputLayer=imageInputLayer([nFeatures,1,1]);
f1=fullyConnectedLayer(nElements);
f2=fullyConnectedLayer(nElements);
f3=fullyConnectedLayer(nElements);
f4=fullyConnectedLayer(nClasses);
s1=softmaxLayer();
outputLayer=classificationLayer();

architecture = [inputLayer; f1; f2; f3; f4; s1; outputLayer];
disp(architecture);
```

```
epochs = 250;
miniBatchSize = 1024;
InitialLearnRate = 0.001;

% Training options: Note that we set the validation patience stopping
% criterion to the number of epochs. This is a stupid thing to do, but we
% want force the training to go to the defined number of epochs so that it
% is consistent with Tensorflow and Pytorch
options = trainingOptions('adam','MaxEpochs',epochs,'InitialLearnRate',InitialLearnRate,...
    'MiniBatchSize',miniBatchSize,'ExecutionEnvironment','cpu','Plots','training-progress',...
    'ValidationData',{x_val,y_val},'ValidationPatience',epochs);
```

```
%% Train
tic
[net,op] = trainNetwork(x_train,y_train,architecture,options);
toc
```

From Knoll Ismrm 2018

Results

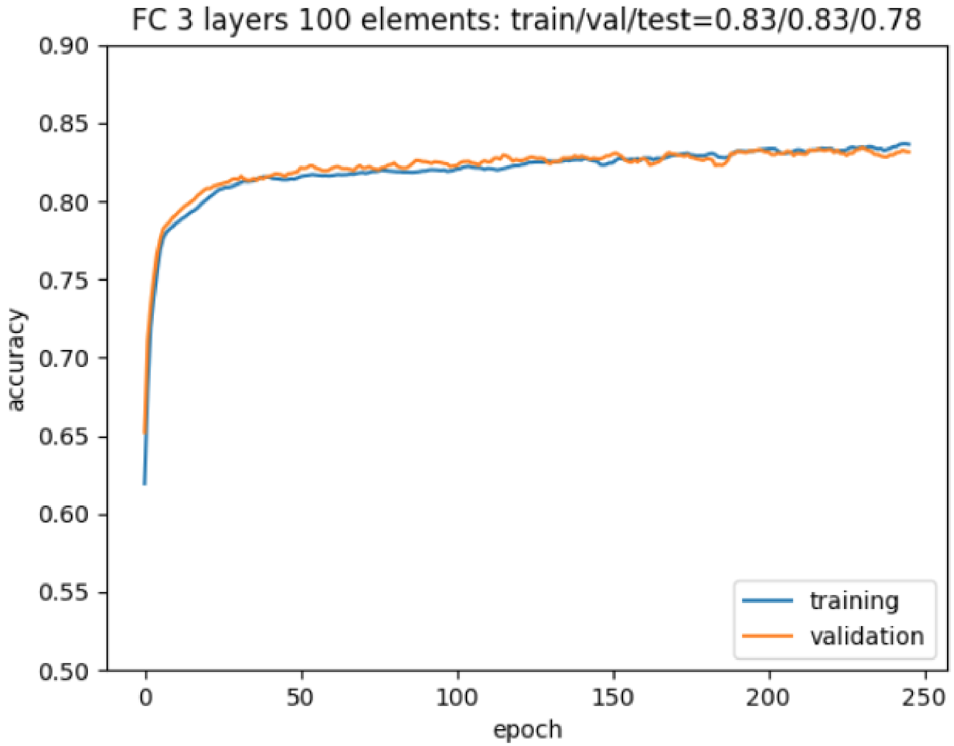

FC 3 layers 100 elements: train/val/test=0.83/0.83/0.78
accuracy
0.90
0.85
0.80
0.75
0.70
0.65
0.60
0.55
0.50
0
50
100
150
200
250
epoch
training
validation

Skin cancer classification

Gather many data

- Dermatologists vs CNN
  - 127.463 biopsy images for training
  - 1942 for validation
  - Inception v3

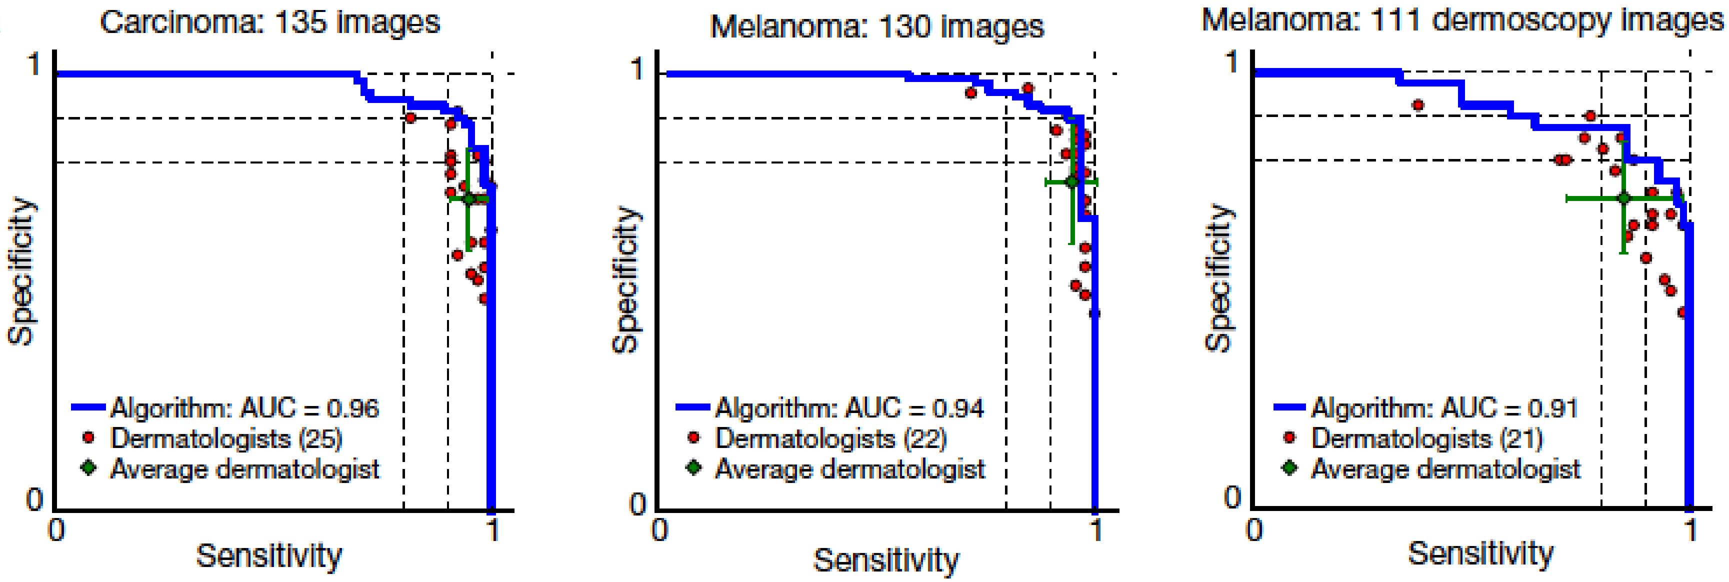

Esteva et al Nature 2017

Skin cancer classification

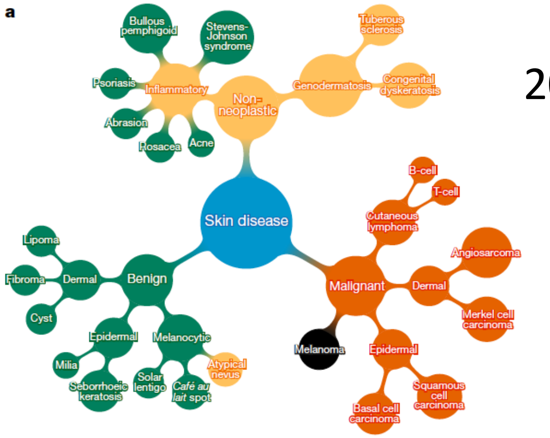

2032 diseases

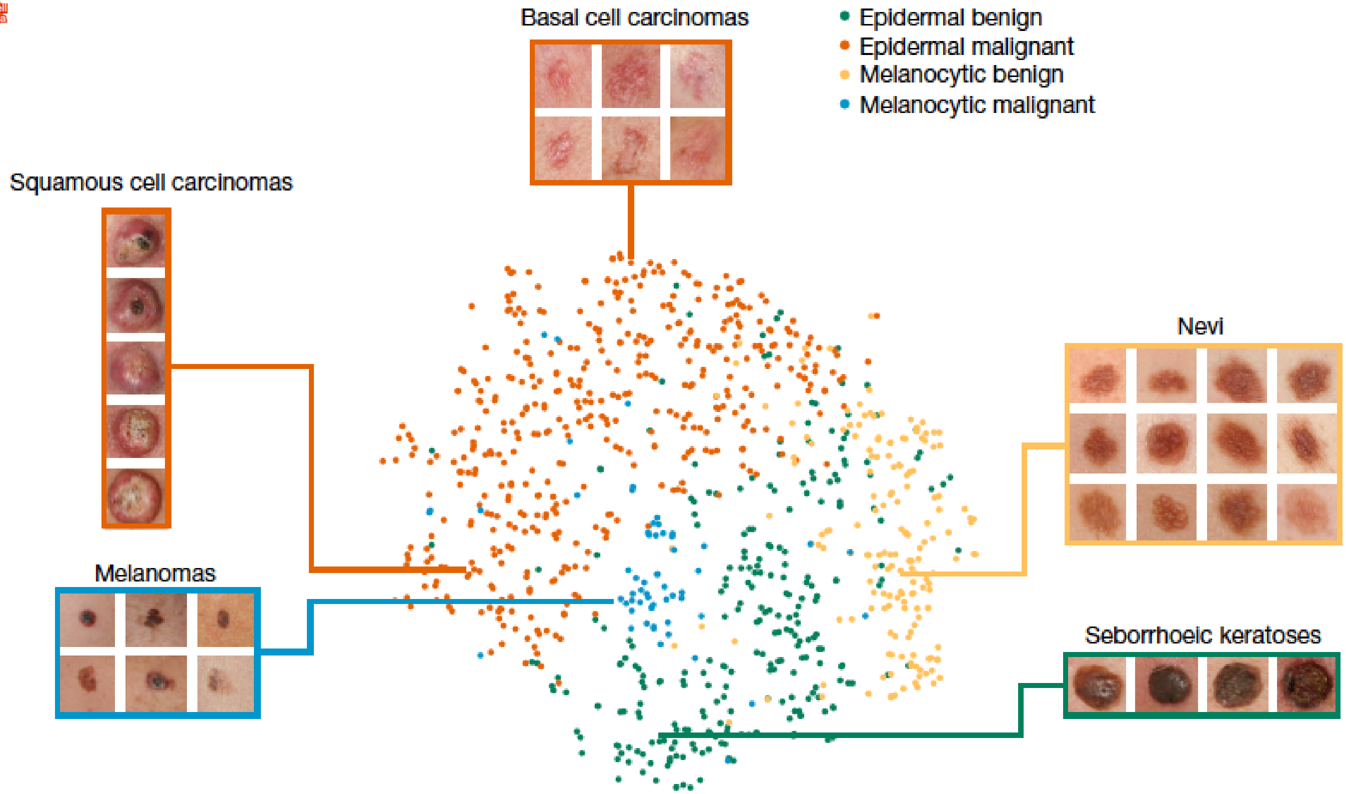

Esteva et al
Nature 2017

Glioblastome classification

Liechty et al 2022 Scient Report

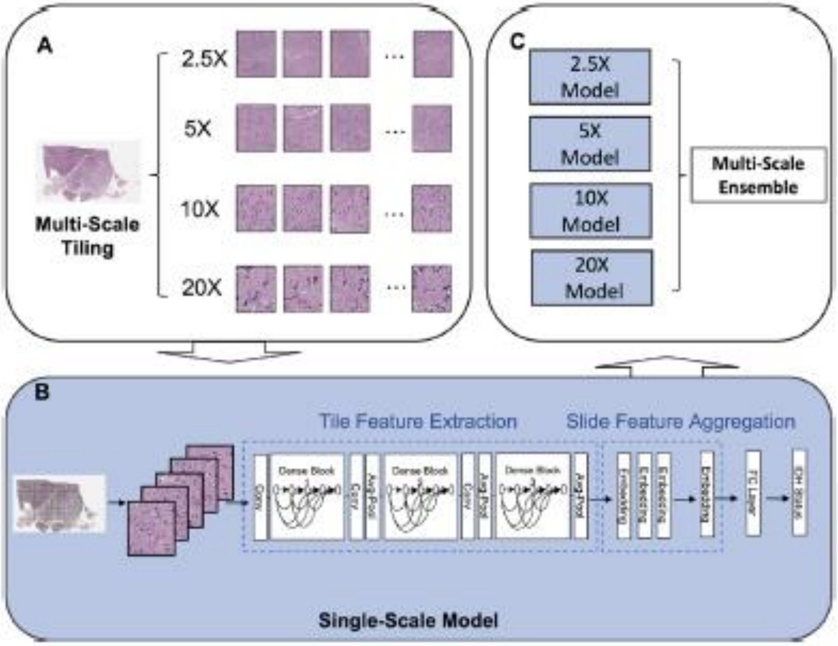

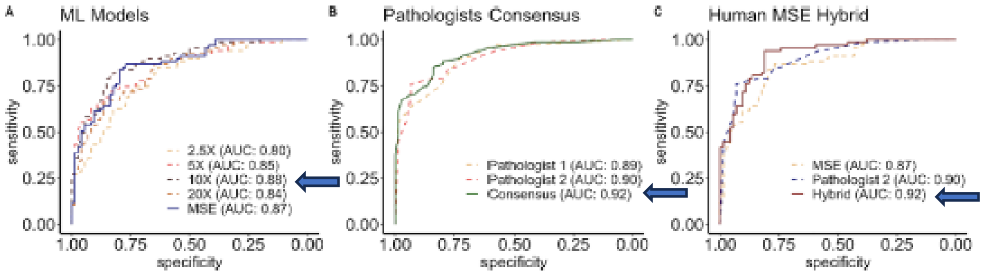

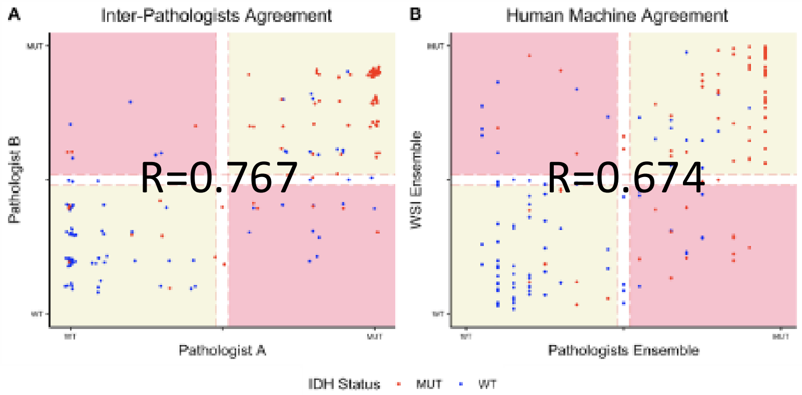

Computational pathology

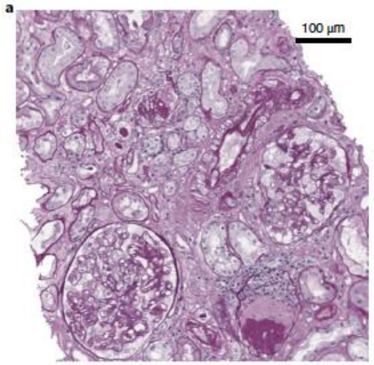

Original image

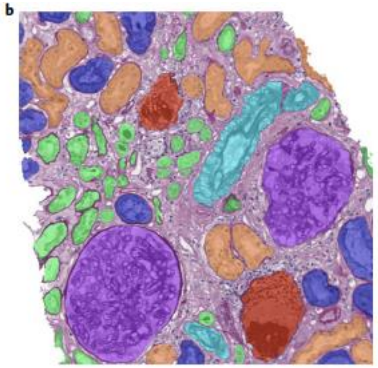

manually annotated by an expert

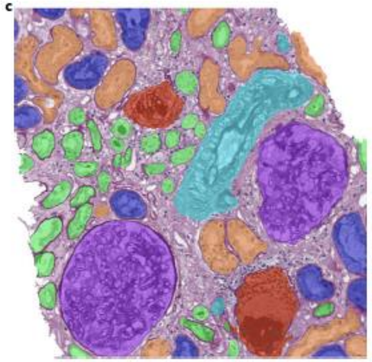

using a deep learning algorithm

Kidney tissue

Purple: glomerulus; red: sclerotic glomerulus;
Dark blue: proximal tubule; orange: distal tubule;
Green: atrophic tubule; turquoise: artery or arteriole

[van des Laak et all Nat Med 2021]

For brain tissue and structures classification

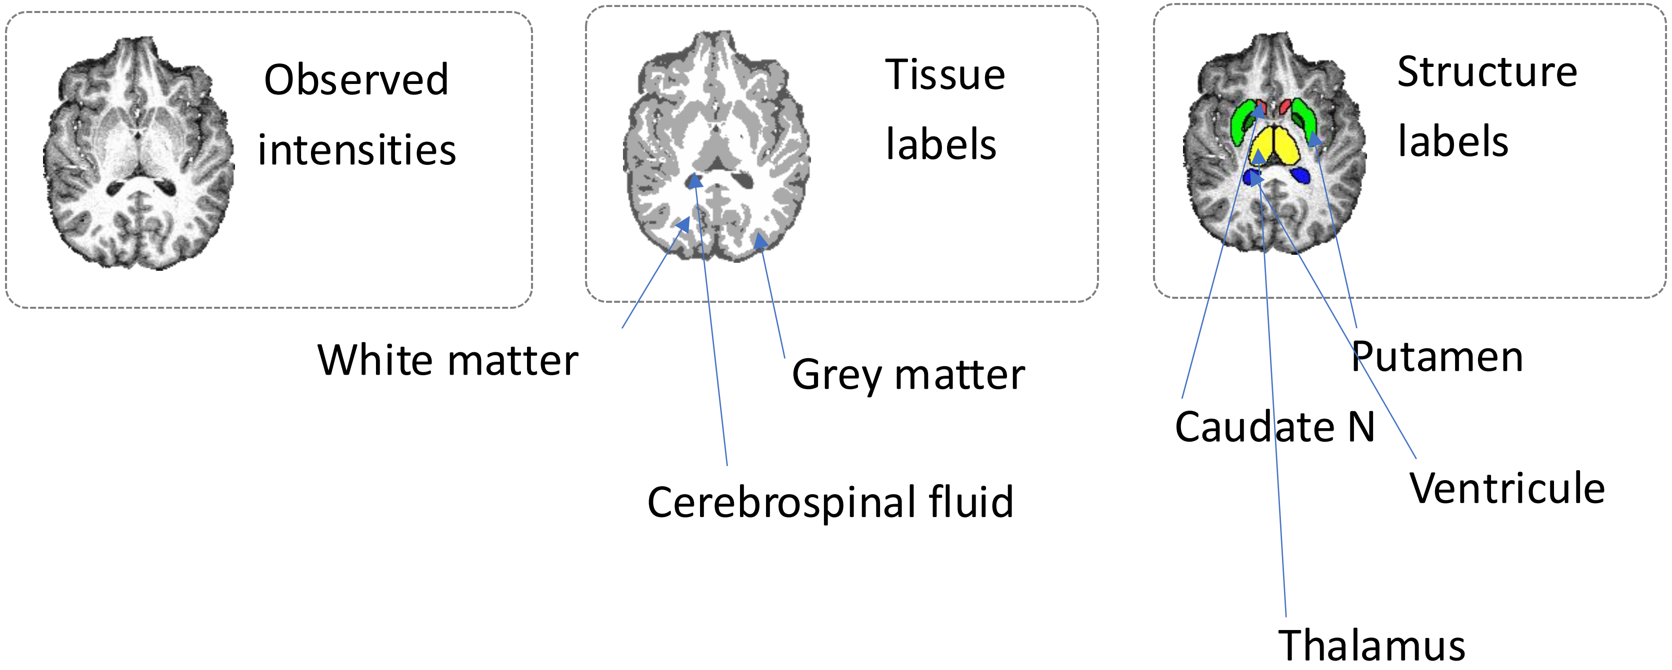

For brain tissue and structures classification

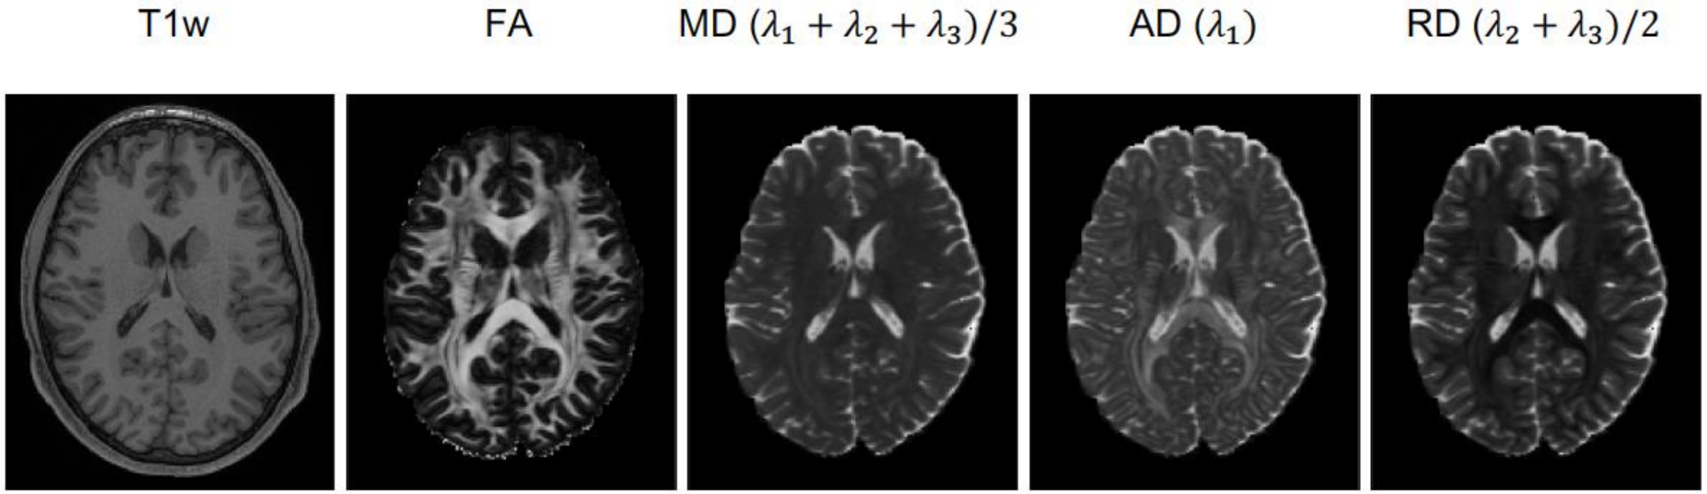

X_i=[$T1w_i$, FA_i, MD_i, AD*i*, RD*i*]

Y_i=[1,2,3, 4, 5, …]

Specific considerations in medical domain

- Few training sets
- Few ground truth
- Need of localization not just classification
- Touching objects of the same class to be separated

Brain Lesion Segmentation: CNN example

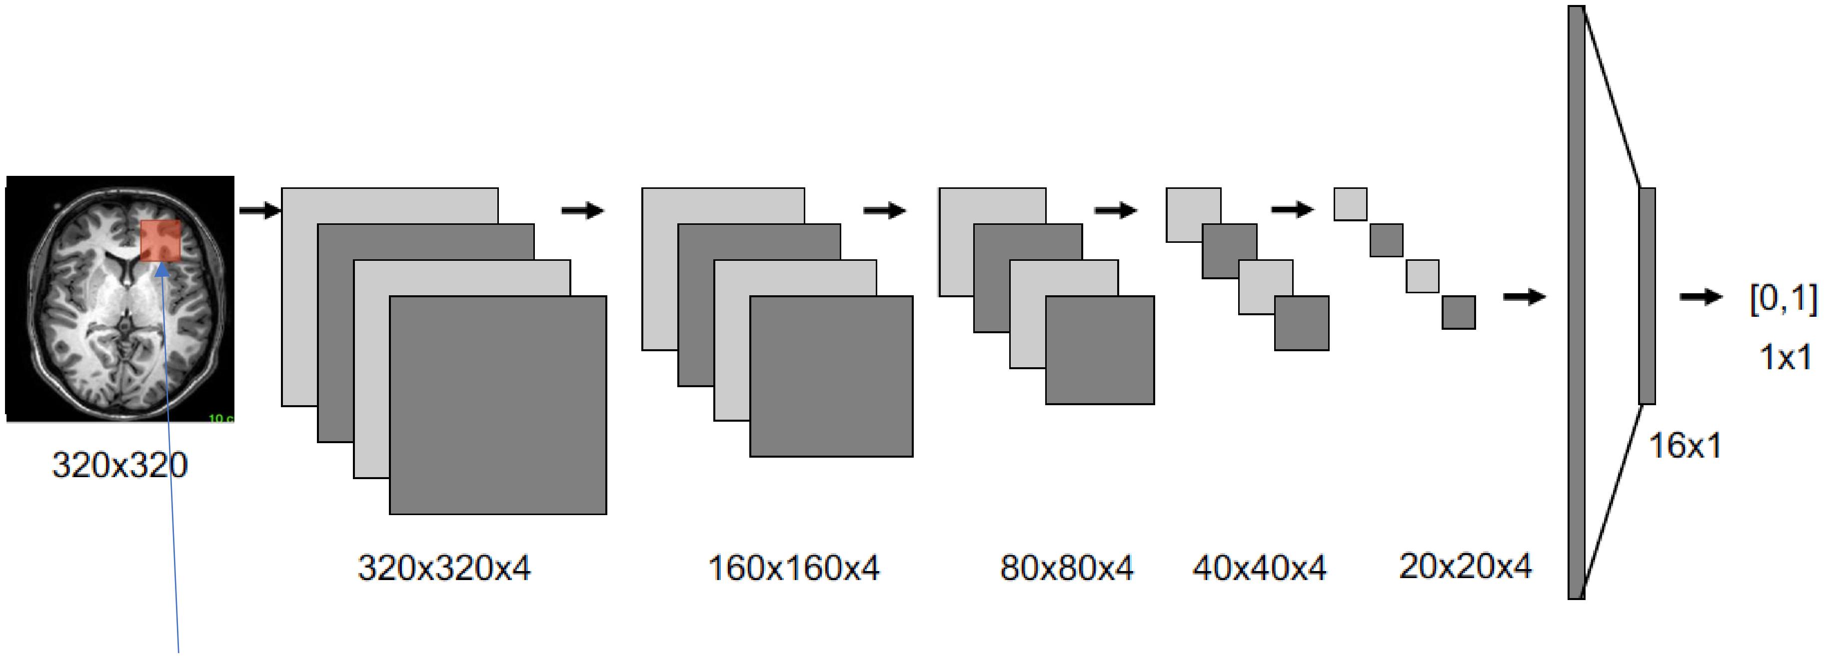

4 3x3x4 convolutional layers, Relu activation
2D MaxPooling
1 fully connected layer with 16 elements
Sigmoid Output

26117 parameters to be estimated

Adapted from Knoll Ismrm 2018

Brain Lesion Segmentation

Kamnitsas et al MedIA 2017

$$\mathbf{y}_l^m = f\left(\sum_{n=1}^{C_{l-1}} \mathbf{k}_l^{m,n} * \mathbf{y}_{l-1}^n + b_l^m\right)$$

learned bias
learned bias

(3D) kernel matrix

$\mathbf{k}_l^{m,n}$

Previous layer

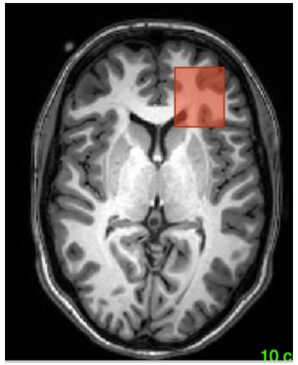

l: layer (1:L)

n: neuron

m: feature (1:Cl)

Cl: max feature map or channel

CL= segmentation class

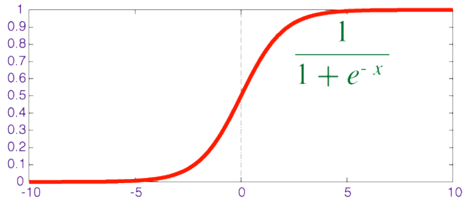

Softmax function

$$p_c(\mathbf{x}) = \exp(\mathbf{y}_L^c(\mathbf{x})) / \sum_{c=1}^{C_L} \exp(\mathbf{y}_L^c(\mathbf{x}))$$

Registration

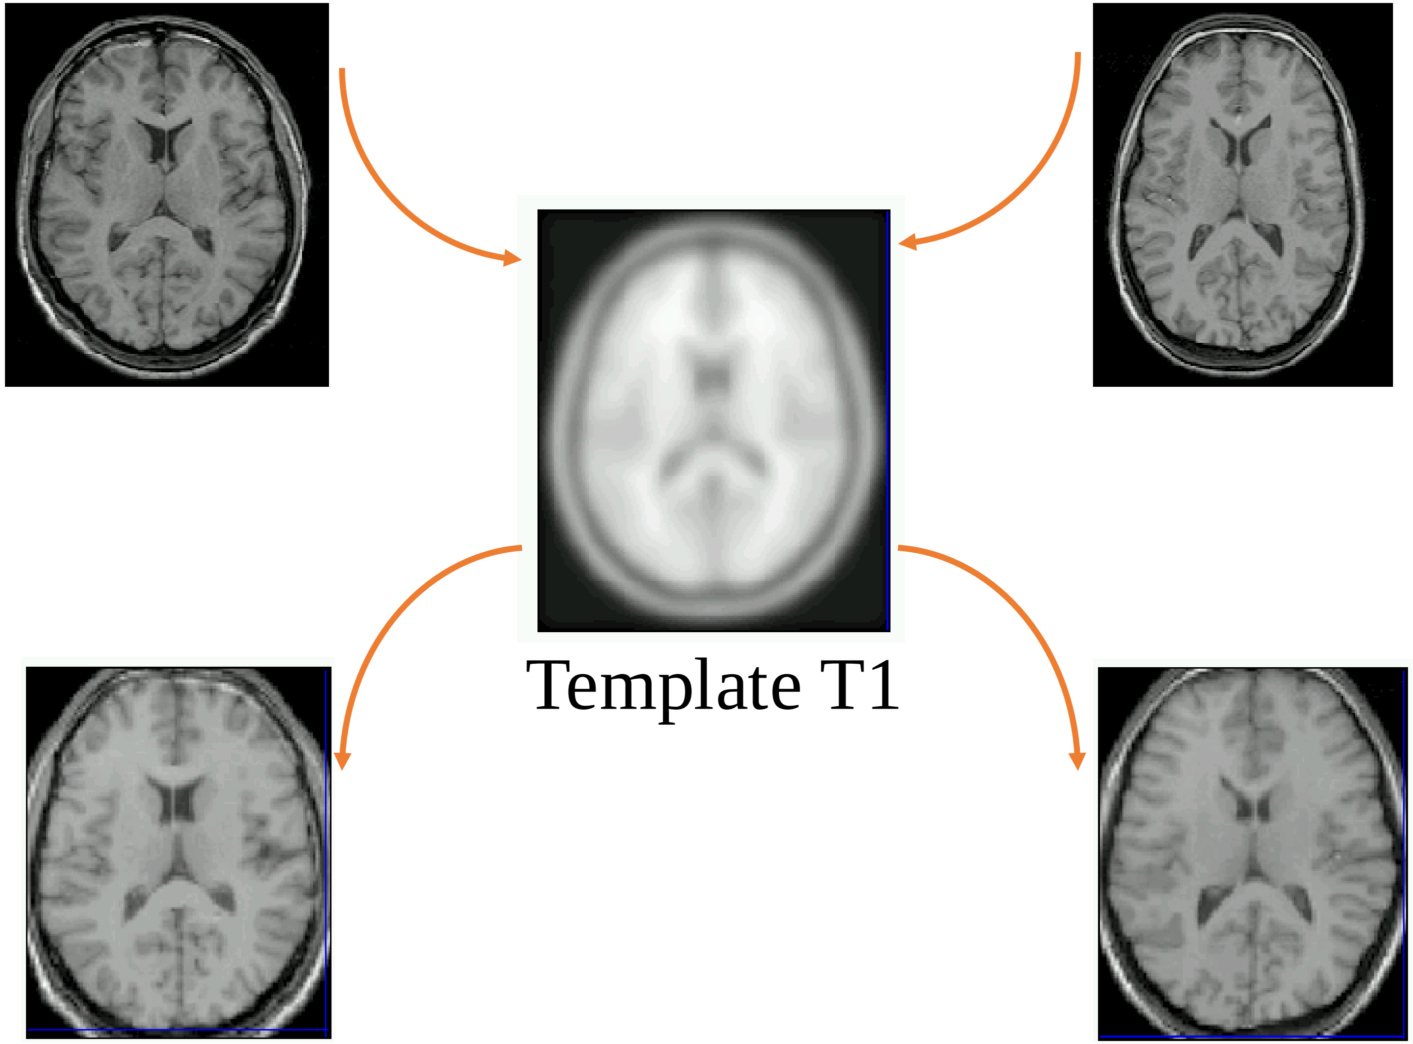

Transformation Affine

Translation

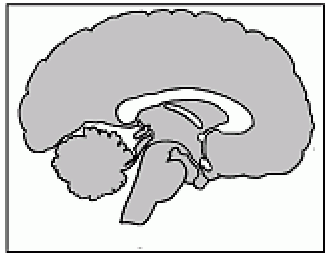

Rotation

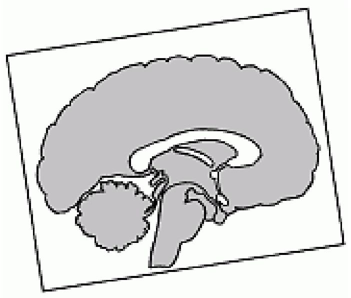

Zoom

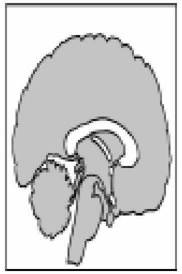

Shear

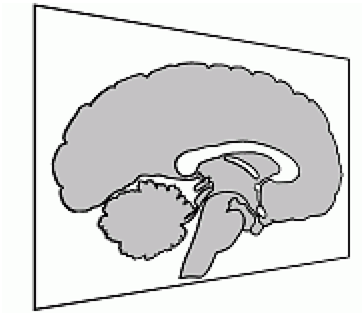

Rotation & translation

zoom

shear

$$\begin{bmatrix} 1 & 0 & 0 & q1 \\ 0 & 1 & 0 & q2 \\ 0 & 0 & 1 & q3 \\ 0 & 0 & 0 & 1 \end{bmatrix} x \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(q4) & \sin(q4) & 0 \\ 0 & -\sin(q4) & \cos(q4) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x \begin{bmatrix} \cos(q5) & 0 & \sin(q5) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(q5) & 0 & \cos(q5) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x \begin{bmatrix} \cos(q6) & \sin(q6) & 0 & 0 \\ -\sin(q6) & \cos(q6) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x \begin{bmatrix} q7 & 0 & 0 & 0 \\ 0 & q8 & 0 & 0 \\ 0 & 0 & q9 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x \begin{bmatrix} 1 & q10 & q11 & 0 \\ 0 & 1 & q12 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

12 paramètres

Insuffisante

Affine = rigide + isométrie (12 paramètres)

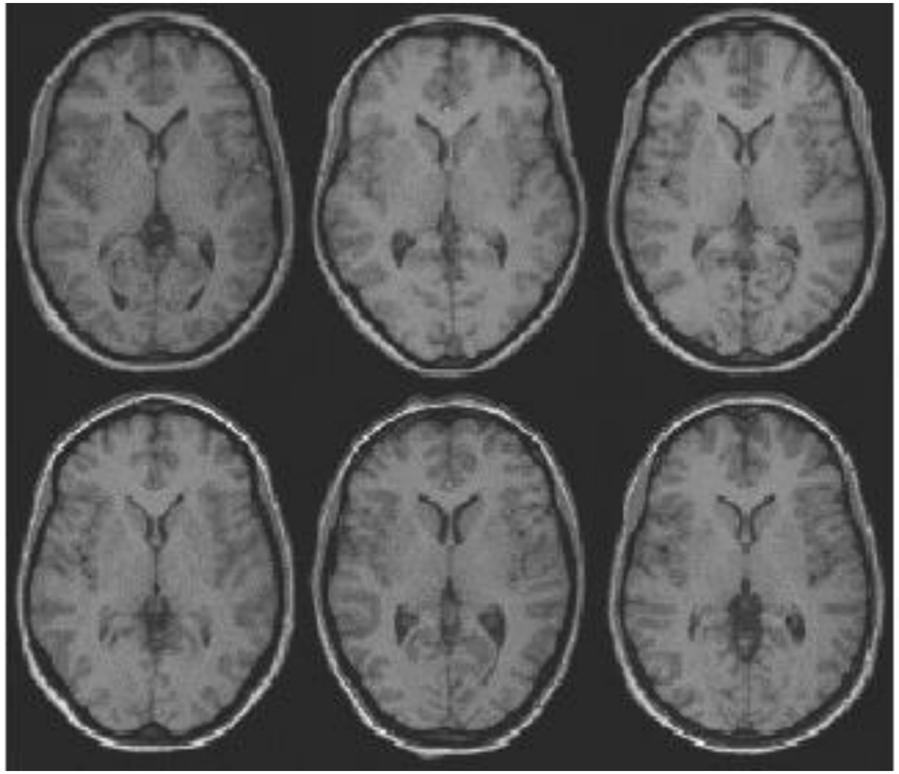

[from Cachia workshop Inserm/CNRS anatomie cérébrale 2006]

Nécessité de recalage non-linéaire

$$T = \underset{\forall T}{\arg\min}\{Mis(O_c, O_s, T) + \gamma Reg(T)\}$$

Registration toward a common template

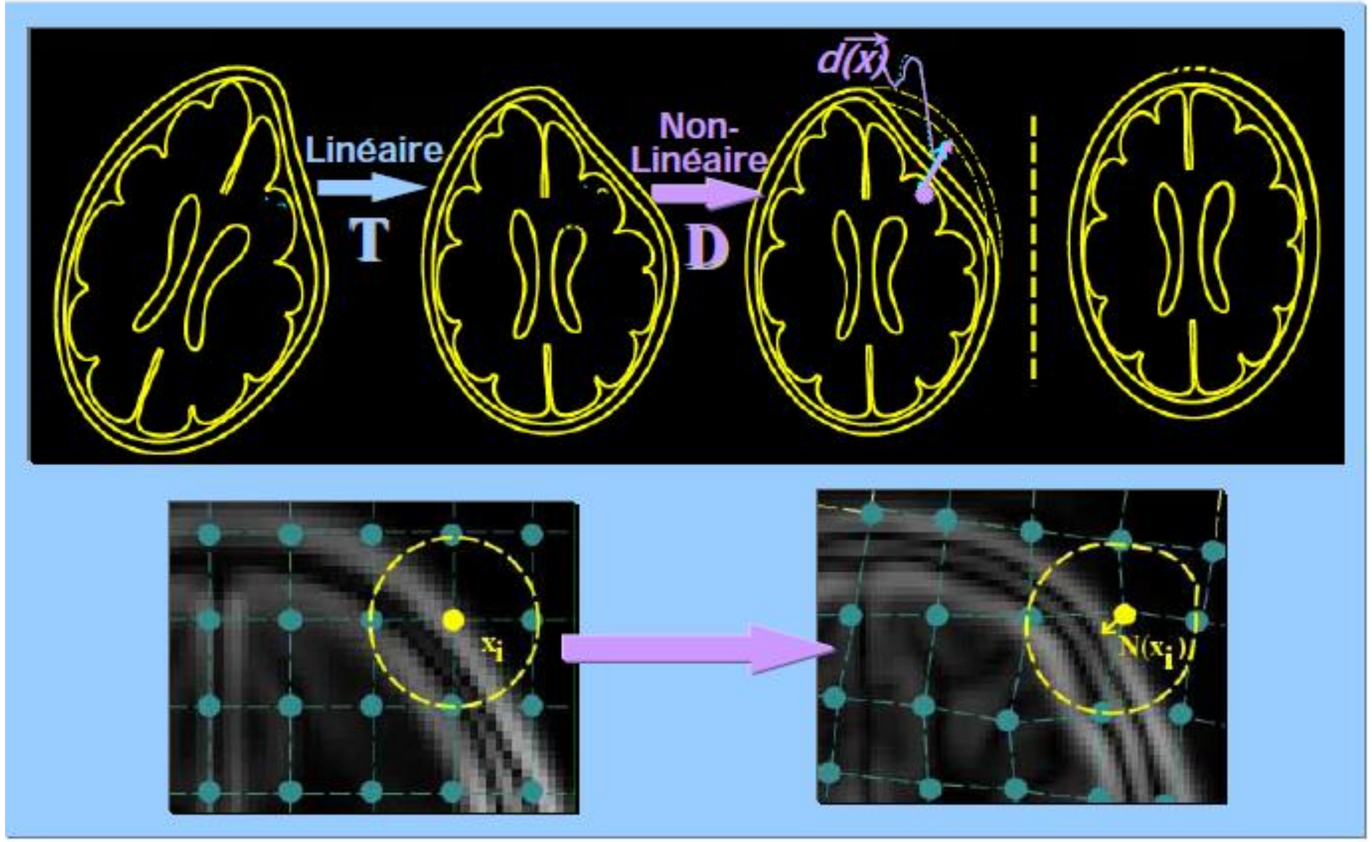

[Collins D. et al., *HBM*, 1995.]

Some templates

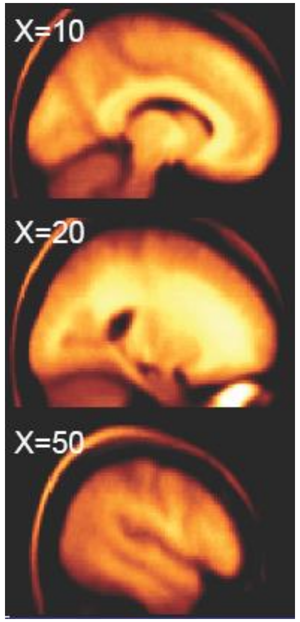

MNI250
250 subjects
Manual
registration
No cerebellum
[Evans et al. 92]

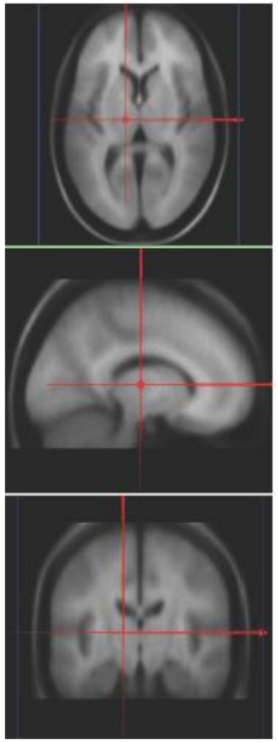

MNI305
350 subjects
Automatic
registration
No cerebellum
[Collins et al. 94]

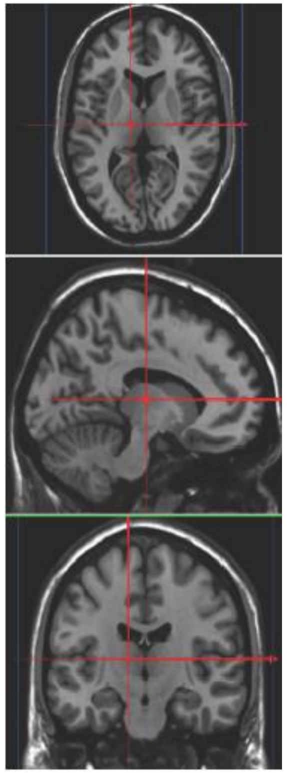

Colin27
1 subject, 27 scans
Automatic
registration
[Holmes et al. 98]

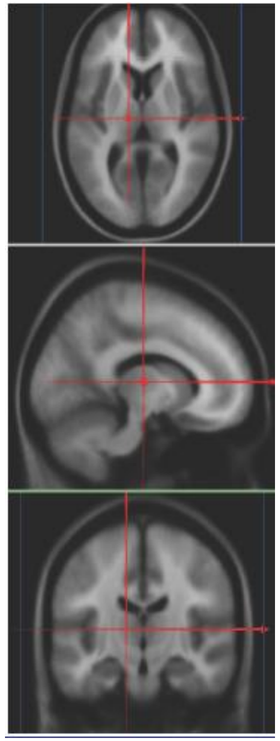

ICBM152
152 subjects
Automtic registration
T1,T2,PWD,
No cortical info
[Maziotta et al. 01]

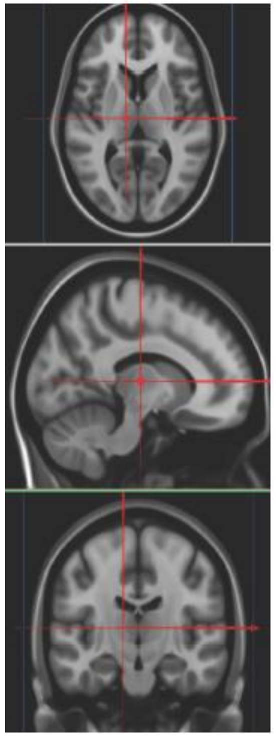

ICBM152 Non linear
152 sujets

[Maziotta et al. 01]

CNN for tissue segmentation

24 filters

Max pooling 2x2

32 filters
Max pooling 2x2

48 filters

Max pooling 2x2

9 classes

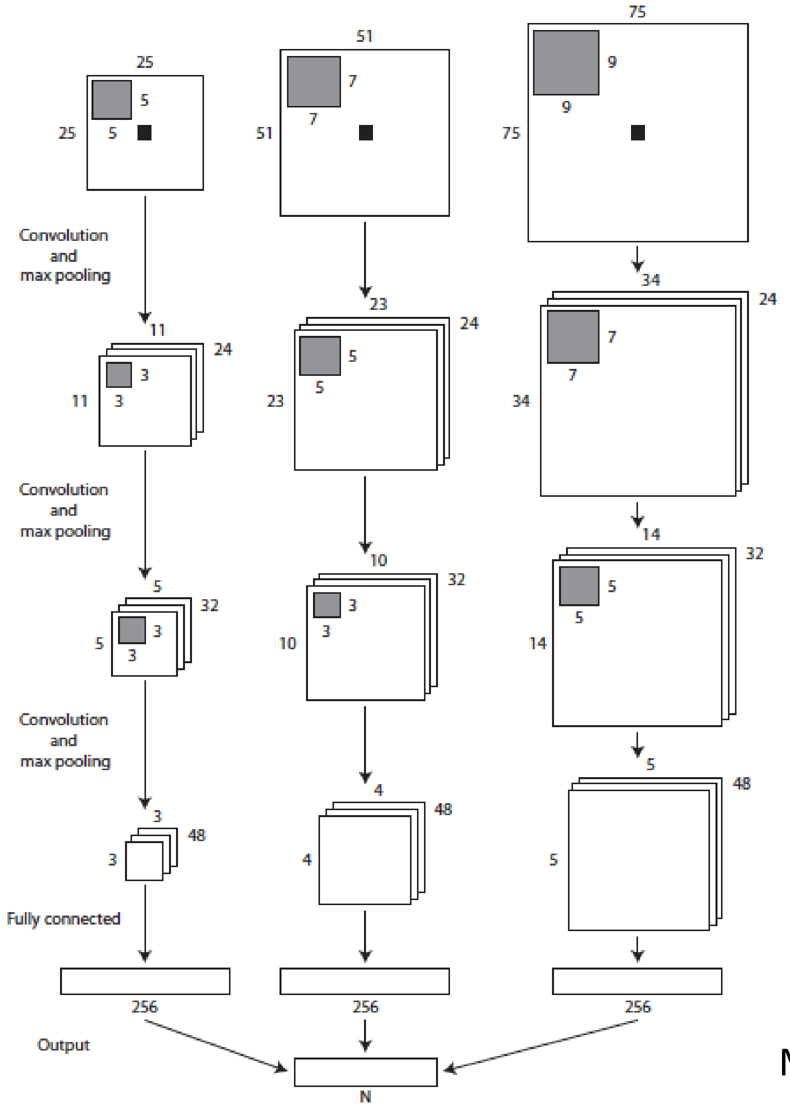

22 neonatal images 3T
20 adult images 3T
15 young adult 1.5T

Moeskops et al TMI 2016

Lesion detection in MR brain scans

- 3D CNN for MRI
  - Computationally expensive so limitation of the number of layers
  - Internal covariate shift > convergence pb

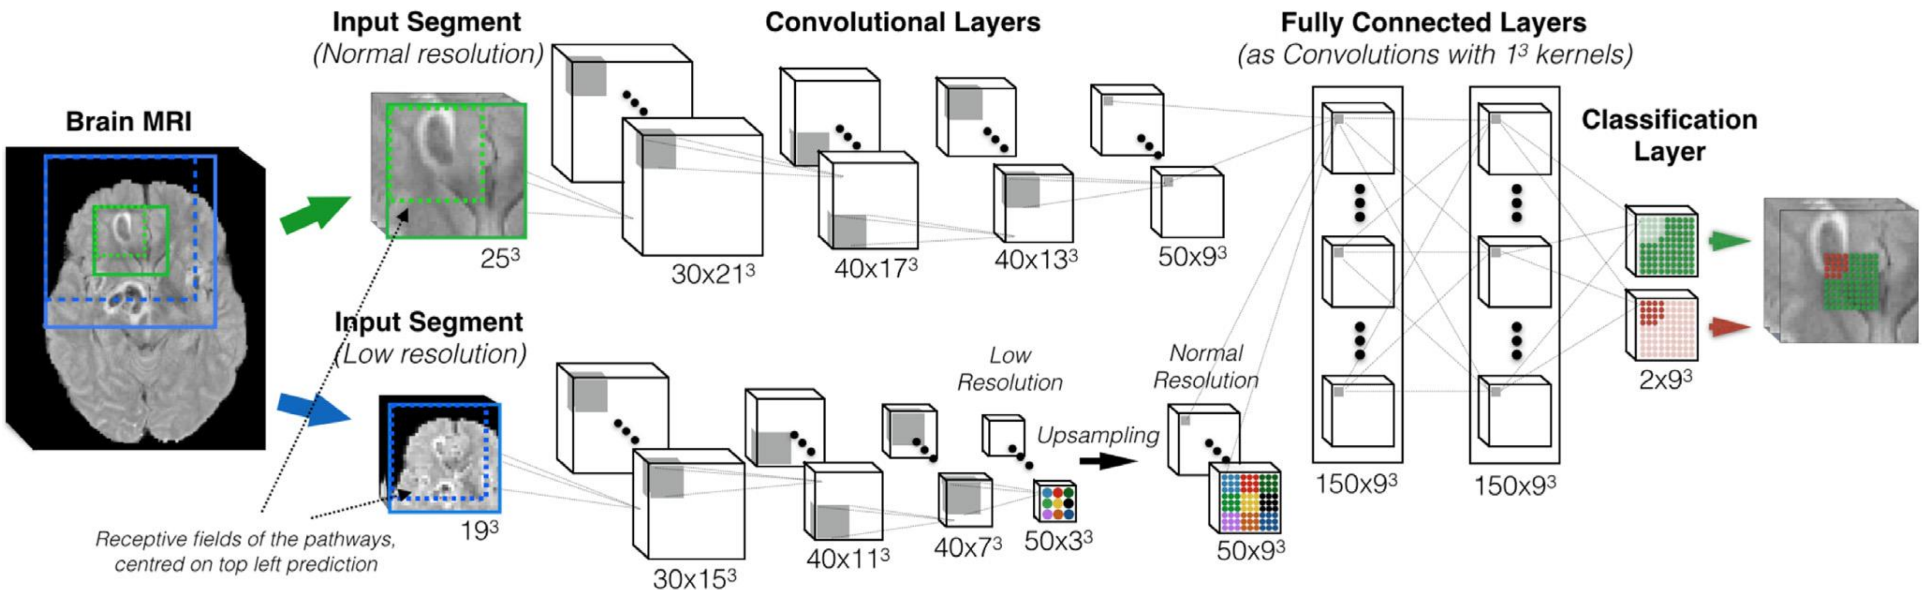

Kamnitsas et al Media 2017

Lesion detection in MR brain scans

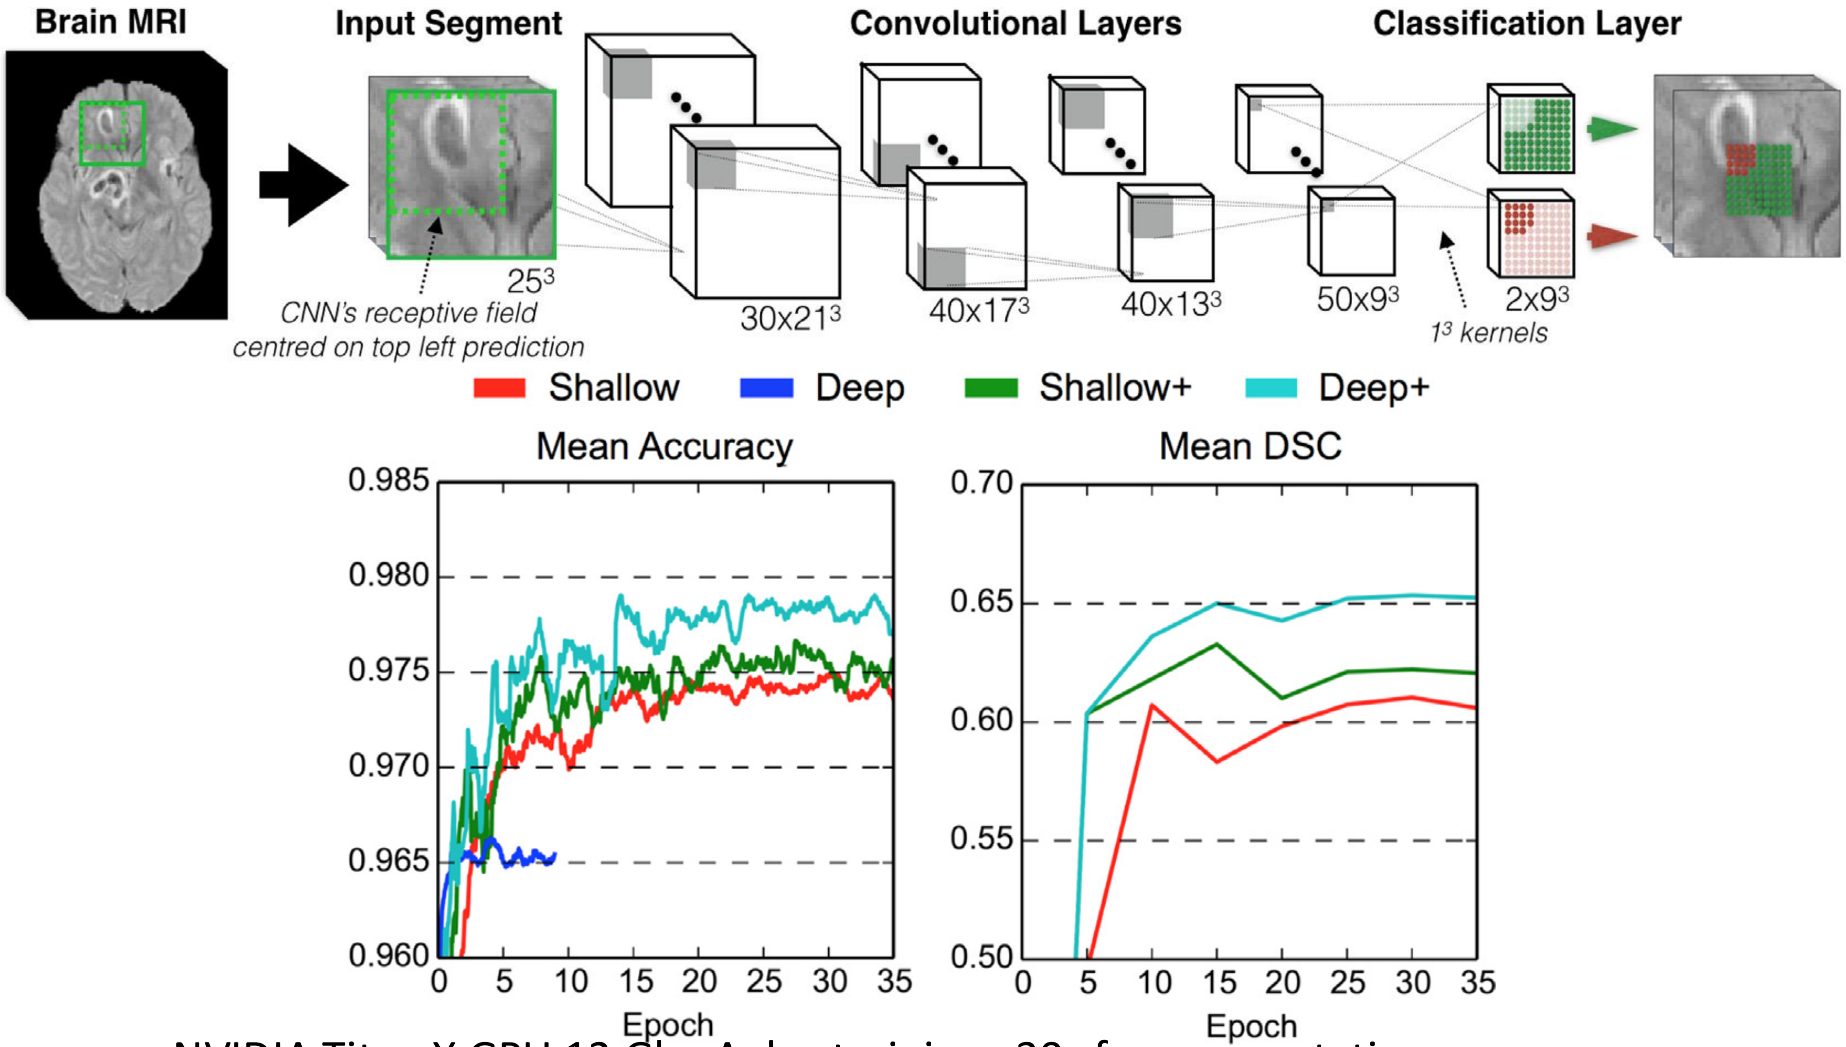

NVIDIA Titan X GPU 12 Gb. A day training. 30s for segmentation.

Traumatic Brain Injury

9 real cases, 9 centers

Normal values computed on 3 healthy cases for each center

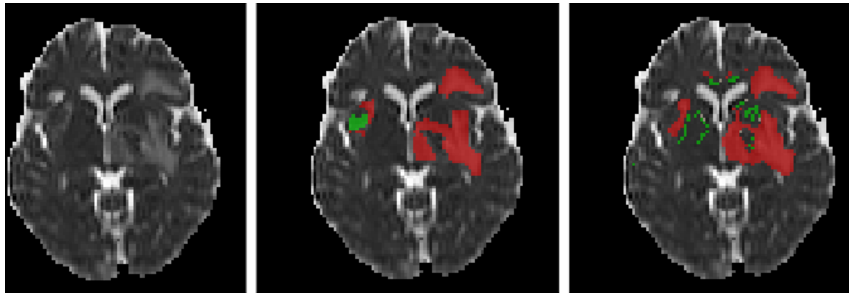

	DICE	Precision	Sensitivity	ICC
Inter-Raters	0.60± 0.03	0.68± 0.04	0.86± 0.04	0.70
AUTOMATIC	0.59± 0.06	0.59± 0.07	0.66± 0.06	0.70
Stroke (Maier et al 2015)	0.73	0.84	0.69	

Kamitsas et al Media 2017

	DSC	Precision	Sensitivity	ASSD	Haussdorf
R. Forest	51.1(20.0)	50.1(24.4)	60.1(15.8)	8.29(6.76)	64.17(15.98)
R. Forest+CRF	**54.8(18.5)****	58.6(23.1)	56.9(17.4)	6.71(5.01)	59.45(15.52)
DeepMedic	62.3(16.4)	65.3(18.8)	64.4(16.3)	4.24(2.64)	56.50(15.88)
DeepMedic+CRF	**63.0(16.3)****	67.7(18.2)	63.2(16.7)	4.02(2.54)	55.68(15.93)
Ensemble	64.2(16.2)	67.7(18.3)	65.3(16.3)	3.88(2.33)	54.38(15.45)
Ensemble+CRF	**64.5(16.3)***	69.8(17.8)	63.9(16.7)	3.72(2.29)	52.38(16.03)

Subcortical segmentation

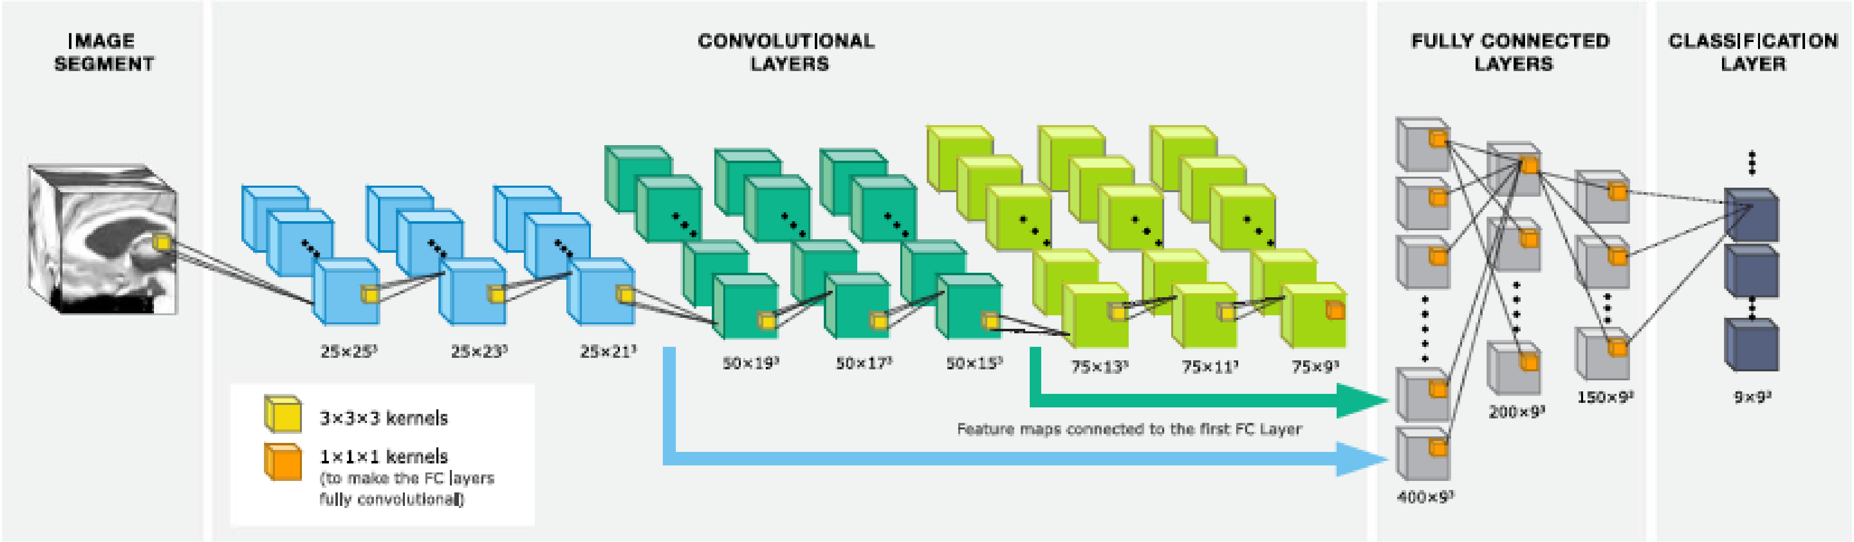

9 structures

Dolz et al Neuroimage 2018

U-Net approach

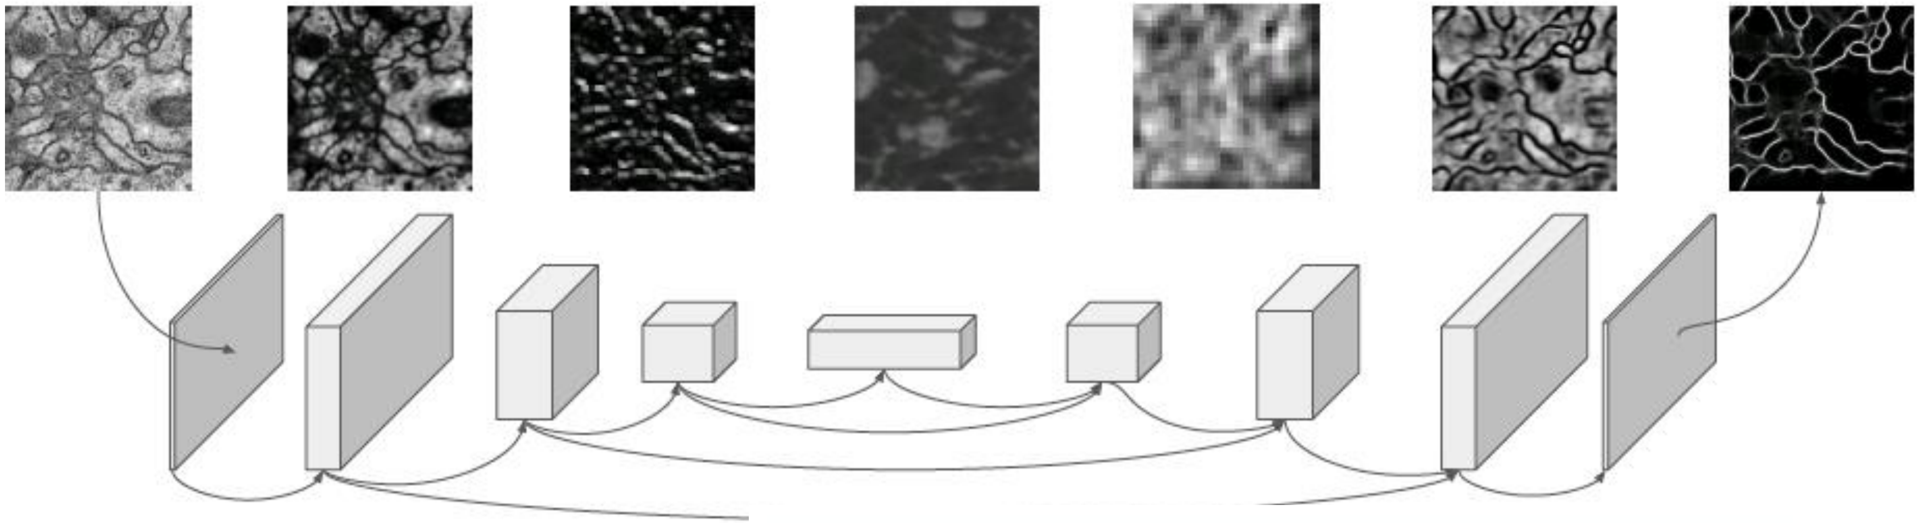

23 convolutional networks

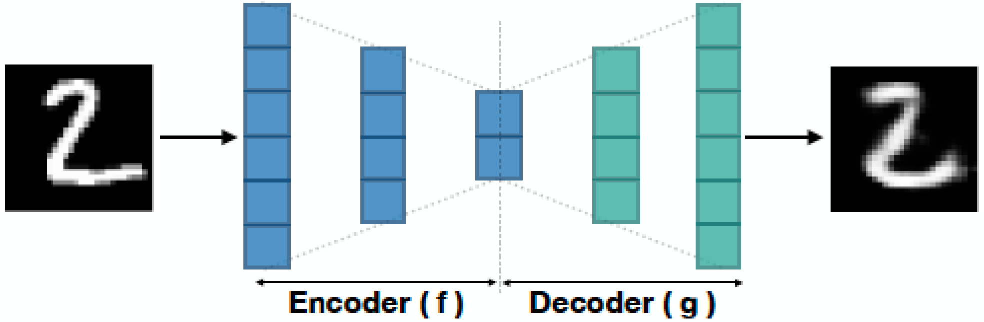

$$z = f(x), \qquad y = g(z)$$

$$\sum_{x \in N} L(x, g(f(x)))$$

[Ronneberger et al 2015]

U-Net approach

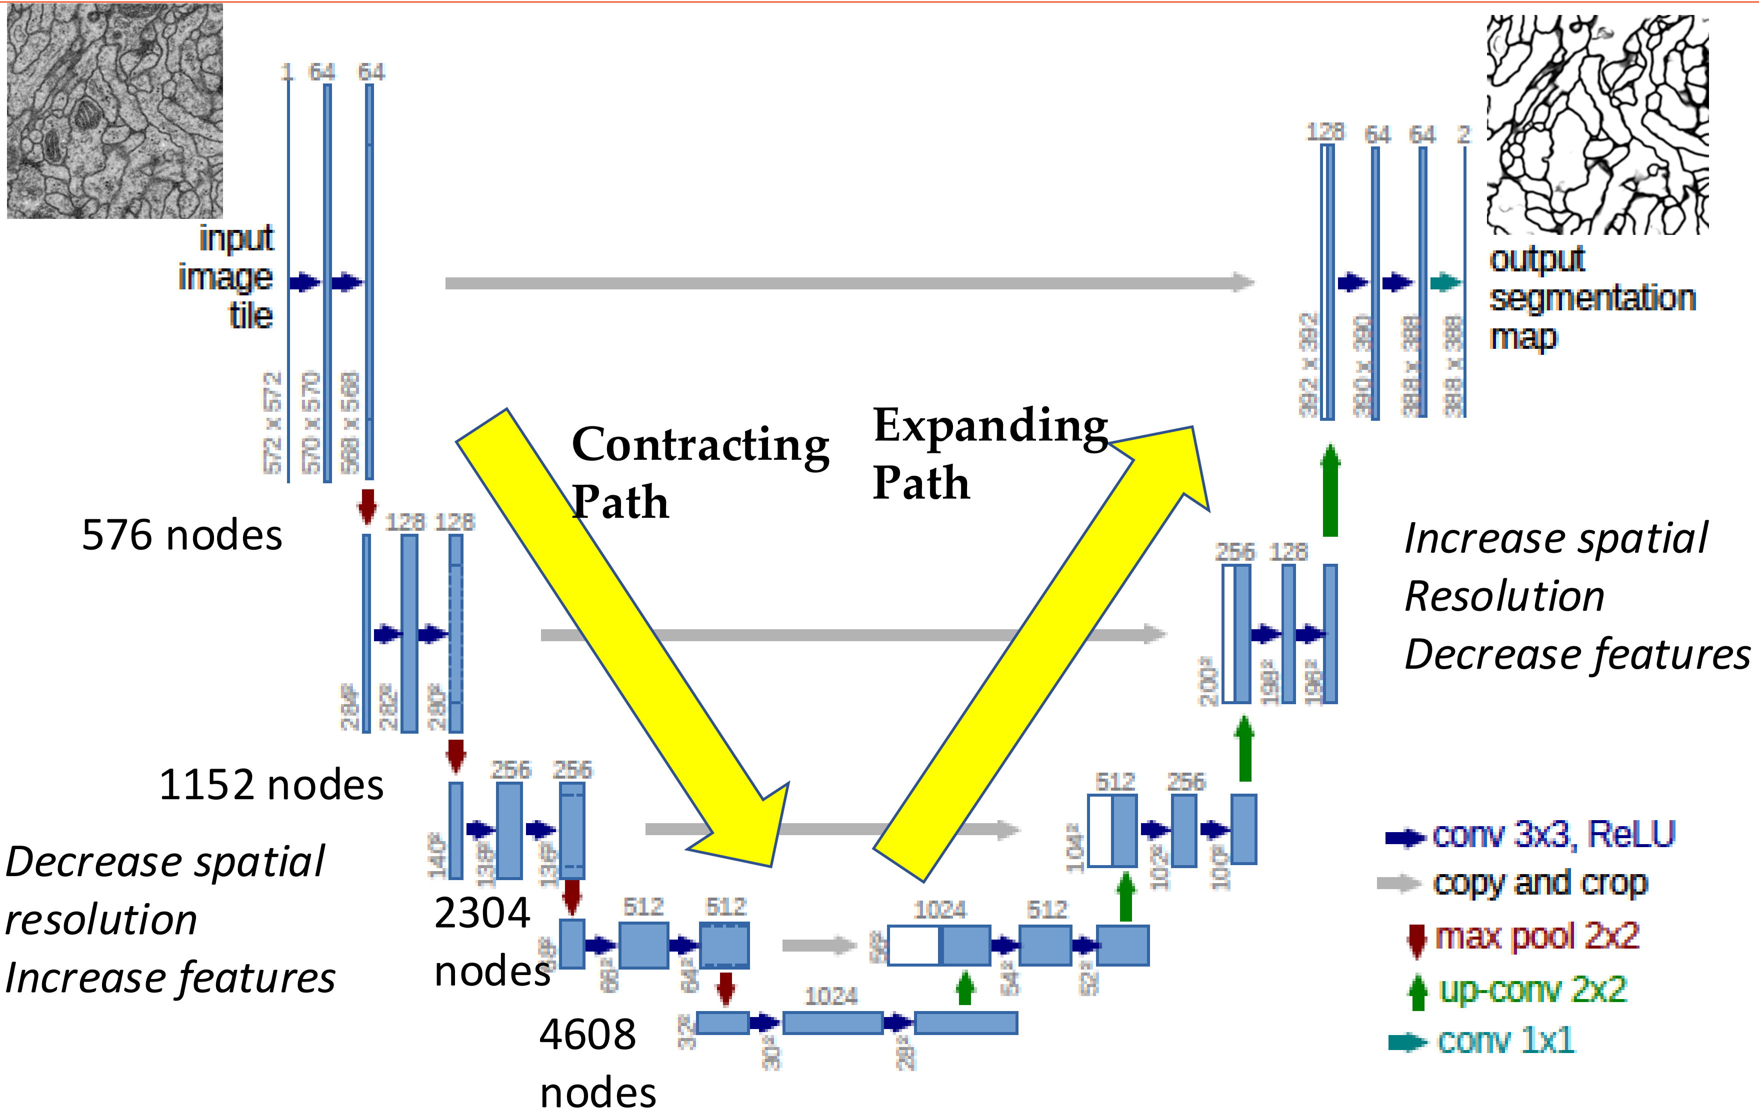

3x3 Convolution + ReLU

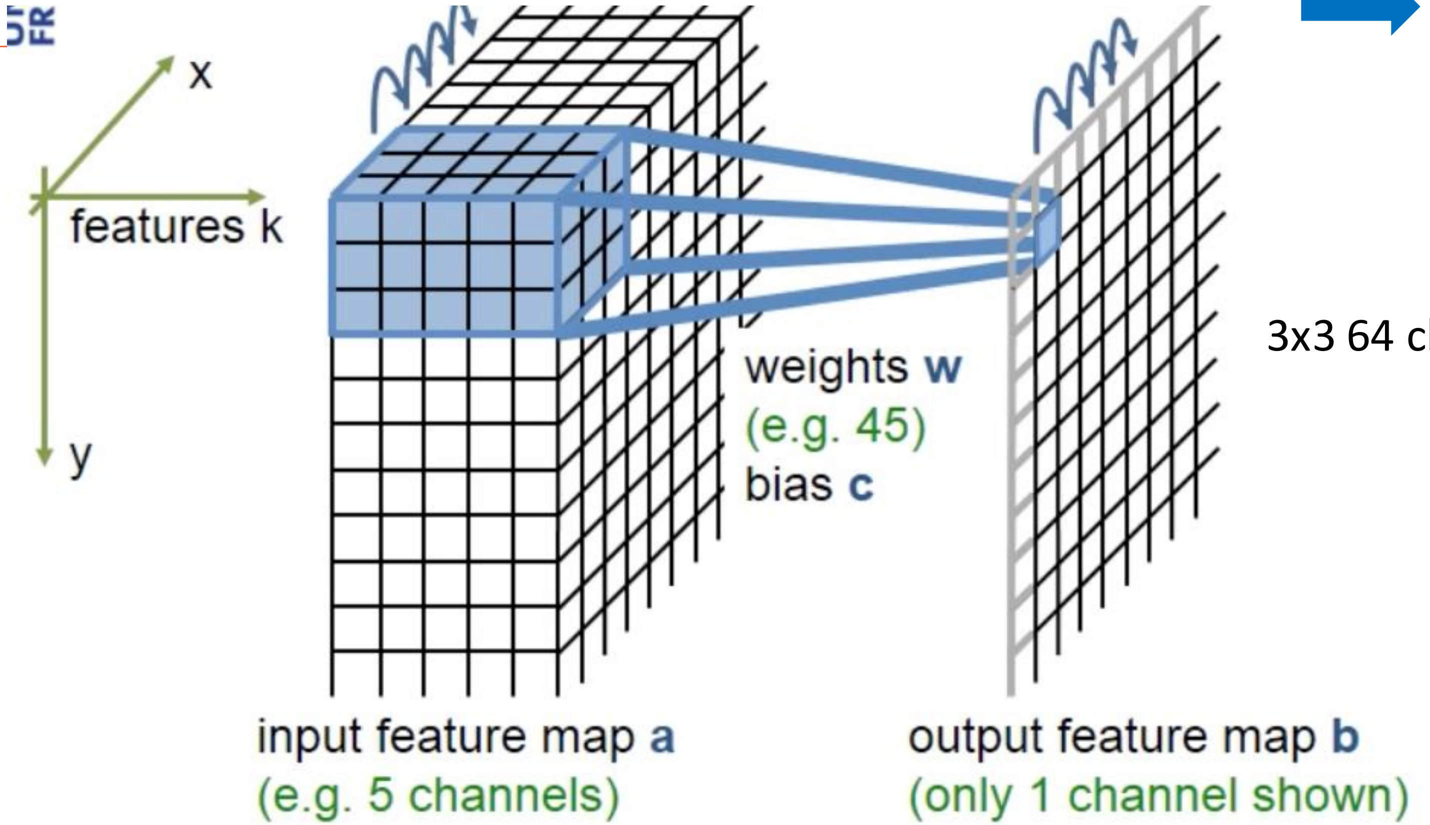

3x3 64 channels=576 nodes

$$b_{x,y,l} = \mathrm{ReLU}\left(\sum_{\substack{i\in\{-1,0,1\}\\ j\in\{-1,0,1\}\\ k\in\{1,\dots,K\}}} w_{i,j,k,l} \cdot a_{x+i,y+j,k} + c_l \right)$$

https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

2x2 Max Pooling

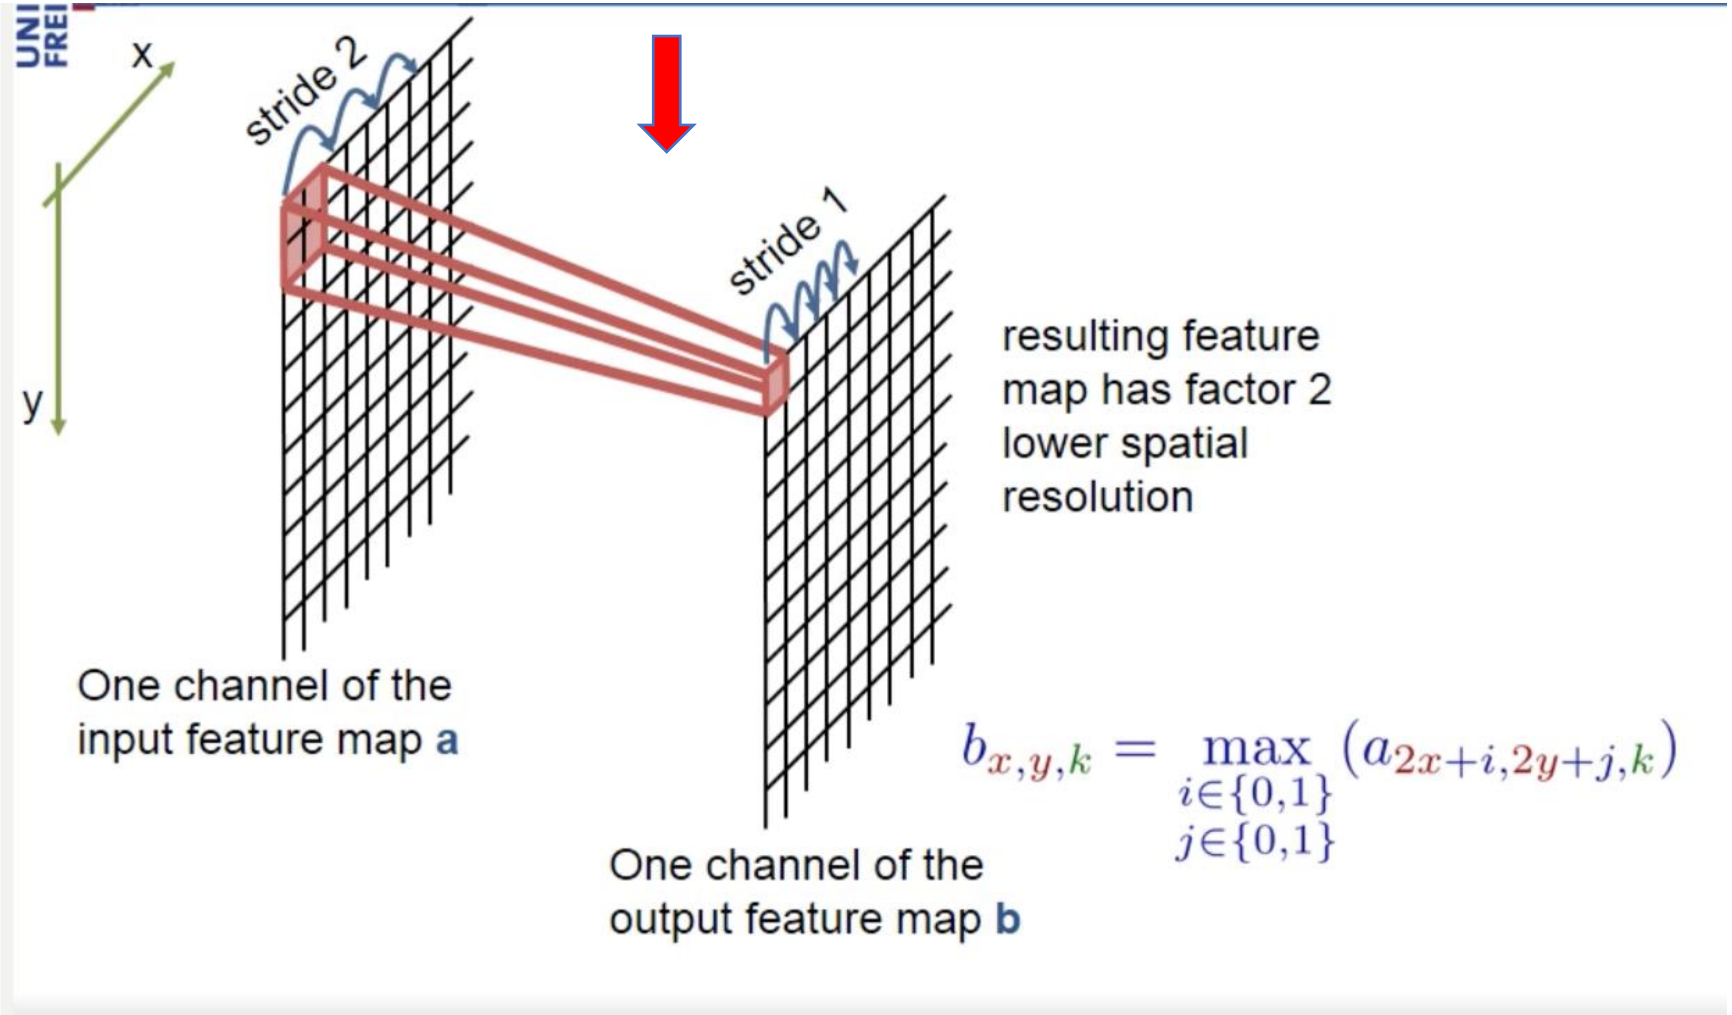

2x2 convolution

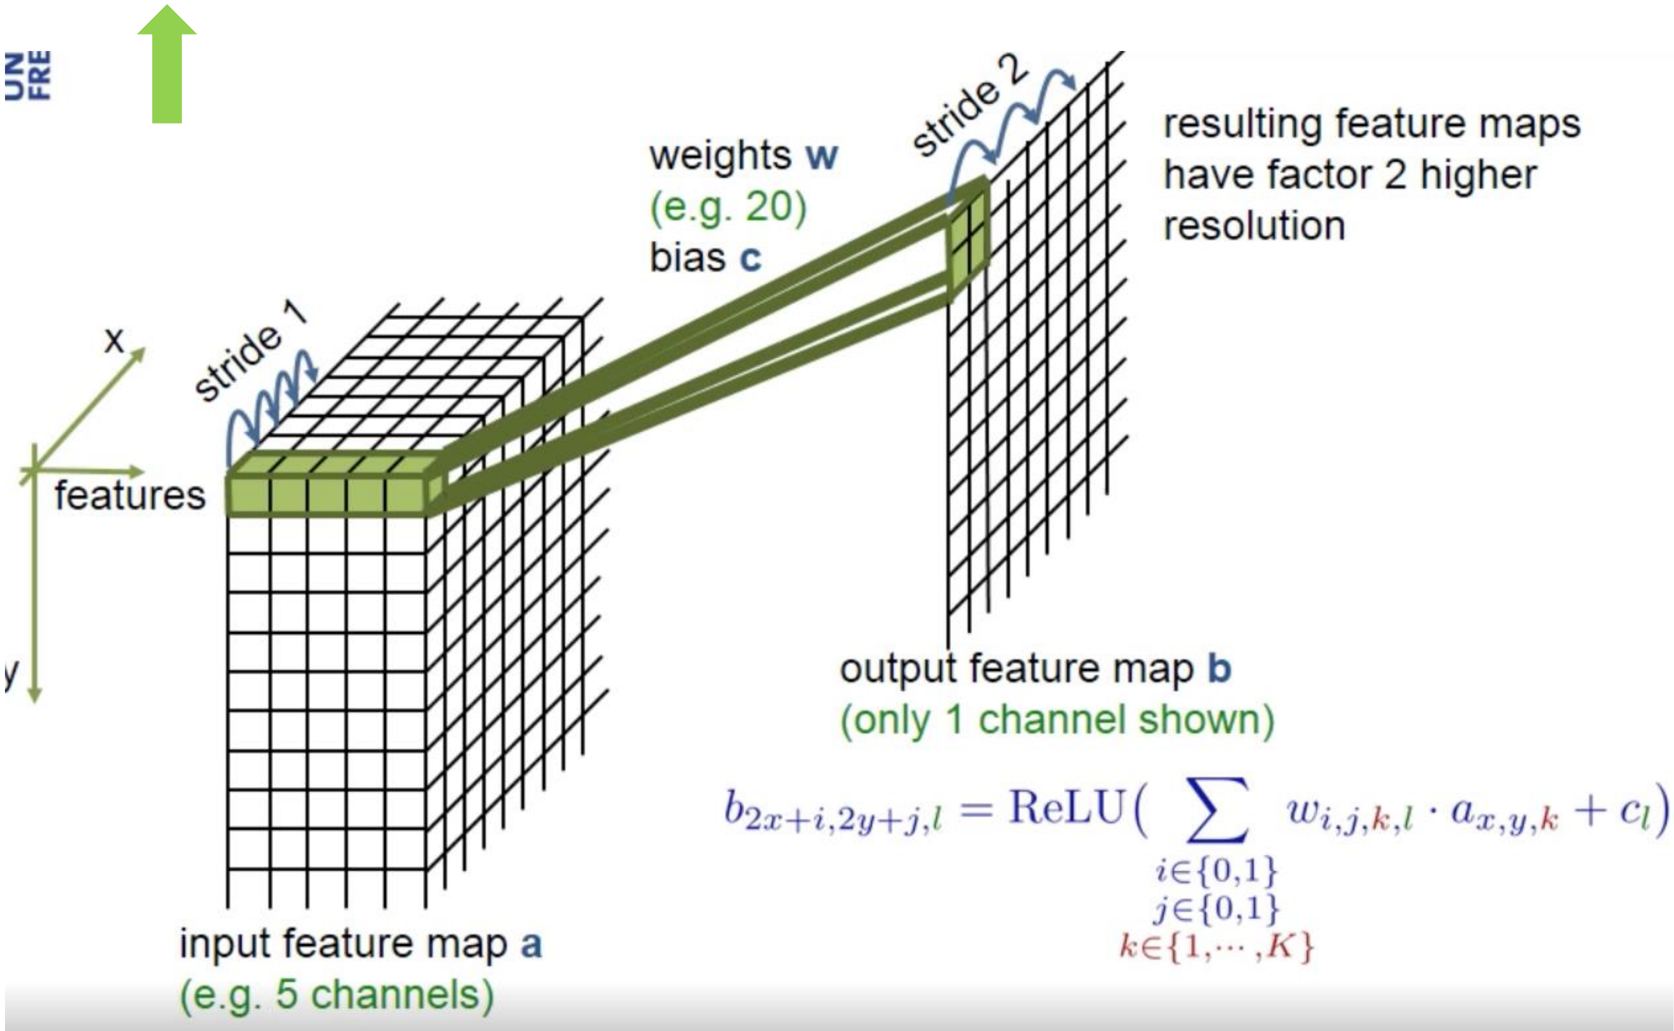

NeuroNet

Rajchl et al MIDL 2018

- 3D Unet structure (DLTK+Tensorflow)
  - 16, 32, 64, 128 features
  - 1,2,2,2 stride
  - Upsampling: 64, 32, 16
  - Training 5000 T1w, segmented FSL, SPM &MALP-EM
  - Validation 713

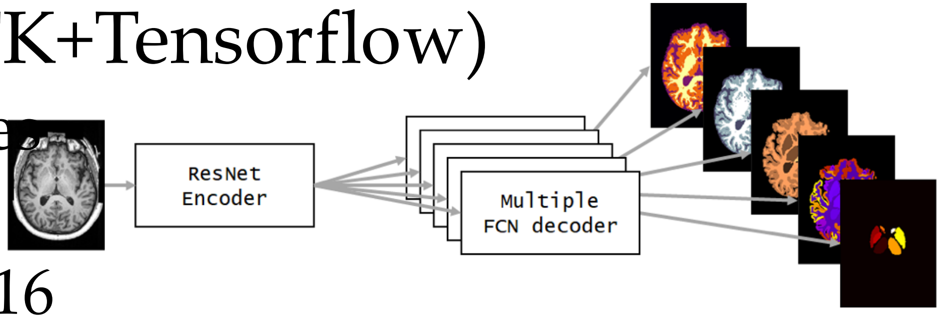

- UK-biobank

protocol	fsl_fast	fsl_first	malp_em	malp_em_tissue	spm_tissue
mean	93.1	88.8	85.8	93.2	93.4
std	2.4	4.9	3.1	1.9	4.3
min	73.7	28.8	57.3	75.8	6.2
max	96.5	92.7	89.7	95.9	96.1
time	20 min		1 h. 8 CPU		>30 min

vs 90s

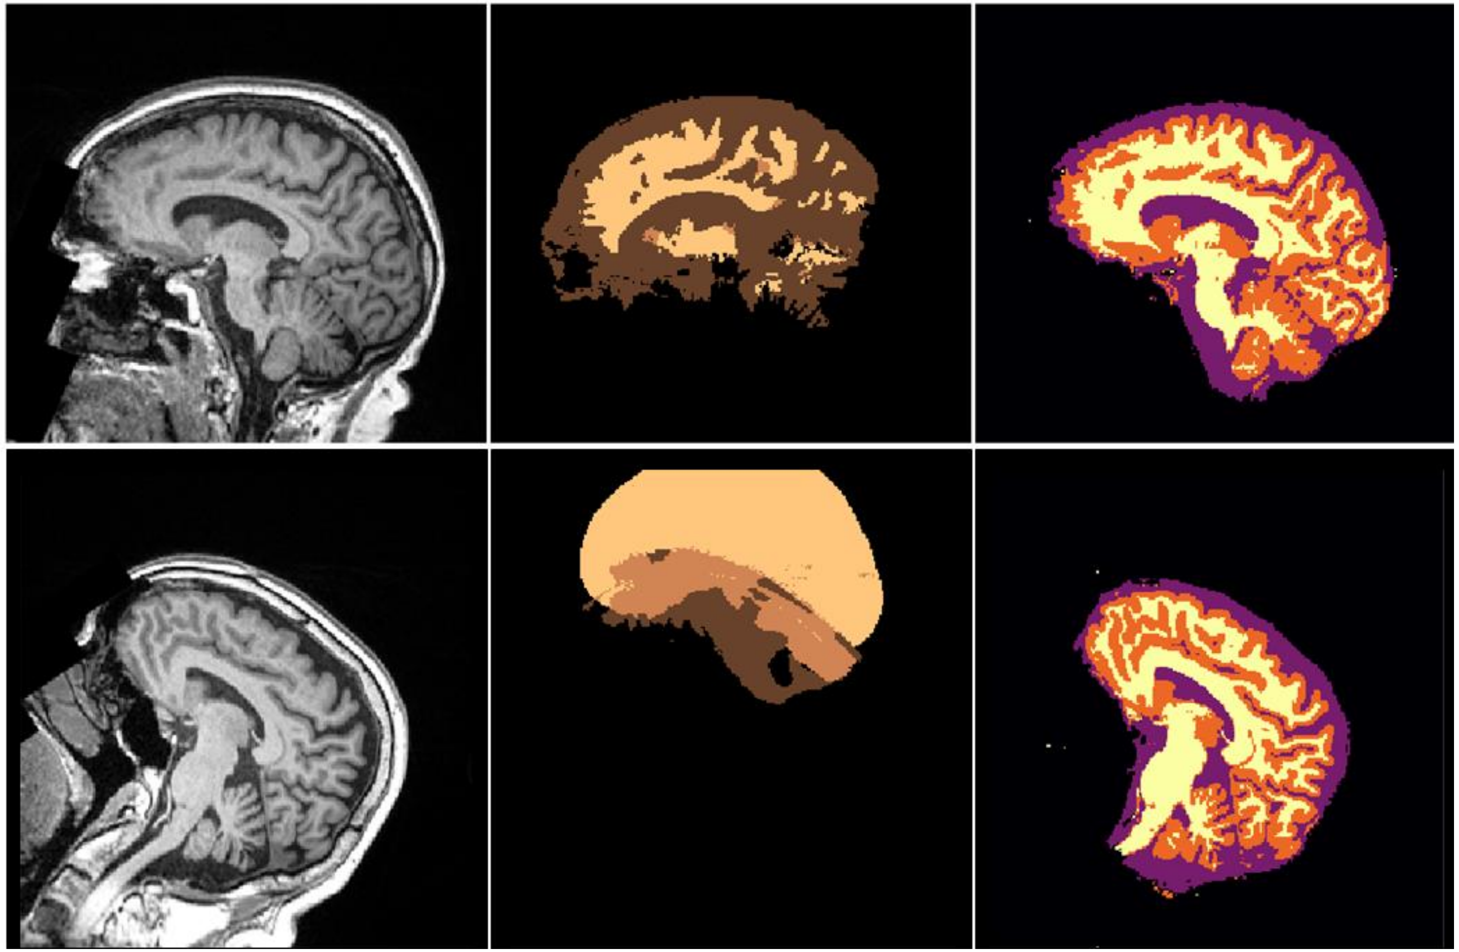

SPM12 NeuroNet

Data hungry systems

- Gather many data
- Artificially increase the number of examples
- Transfert Learning
- Generate synthetic data
- Pooling from many centers
- Anomaly Detection

UK BioBank

Gather many data

How to scan 100.000 people Brain Imaging in UK Biobank

4 scanners, 54 pat/day, 7/7, 5 years
2022 Jv 50000 subjects

Cost: about 10k€ for data import

Miller et al Nat Neurosc
Multimodal population brain imaging in the UK Biobank prospective epidemiological study.2016

Modality	Duration	Voxel, Matrix	Key Parameters
T1	4:54	1.0x1.0x1.0 mm 208x256x256	3D MPRAGE, sagittal, R=2, TI/TR=880/2000 ms
T2 FLAIR	5:52	1.05x1.0x1.0 mm 192x256x256	FLAIR, 3D SPACE, sagittal, R=2, PF 7/8, fat sat, TI/TR=1800/5000 ms, elliptical
swMRI	2:34	0.8x0.8x3.0 mm 256x288x48	3D GRE, axial, R=2, PF 7/8 TE1/TE2/TR=9.4/20/27 ms,
dMRI	7:08	2.0x2.0x2.0 mm 104x104x72	MB=3, R=1, fat sat, b=0(5x + 3x phase-encoding-reversed), 1000(50x), 2000(50x)
rfMRI	6:10	2.4x2.4x2.4 mm 88x88x64	TE/TR=39/735 ms, MB=8, R=1, flip angle 52°, fat sat
tfMRI	4:13	2.4x2.4x2.4 mm 88x88x64	Acquisition same as rfMRI. Task is faces/shapes "emotion" task.

R = in-plane acceleration factor, MB = multiband factor, PF=partial Fourier. All non-EPI scans are pre-scan normalized (on-scanner bias-field corrected). Gradient distortion correction is turned off on the scanner and applied in post-processing.

Genome-wide association studies of brain imaging phenotypes in UK Biobank T Eliott et al Nat Neurosc 2018

The UK Biobank imaging enhancement of 100,000 participants: rationale, data collection, management and future directions

Phenotypic and genetic associations of quantitative magnetic susceptibility in UK Biobank brain imaging Wang et al Nat Neurosc 2022 LittleJohns TJ et al Nature Comm 2020

SARS-CoV-2 is associated with changes in brain structure in UK Biobank Douaud et al Nat Neurosc 2022

Ivy Glioblastoma Atlas

Gather many data

Alignment of histologic features, genomic alterations and gene expression patterns

41 patients,
42 tumors,
440 tissue blocks,
270 transcriptomes,
11500 annotated H&E,
23000 ISH images,
400 MRI scans

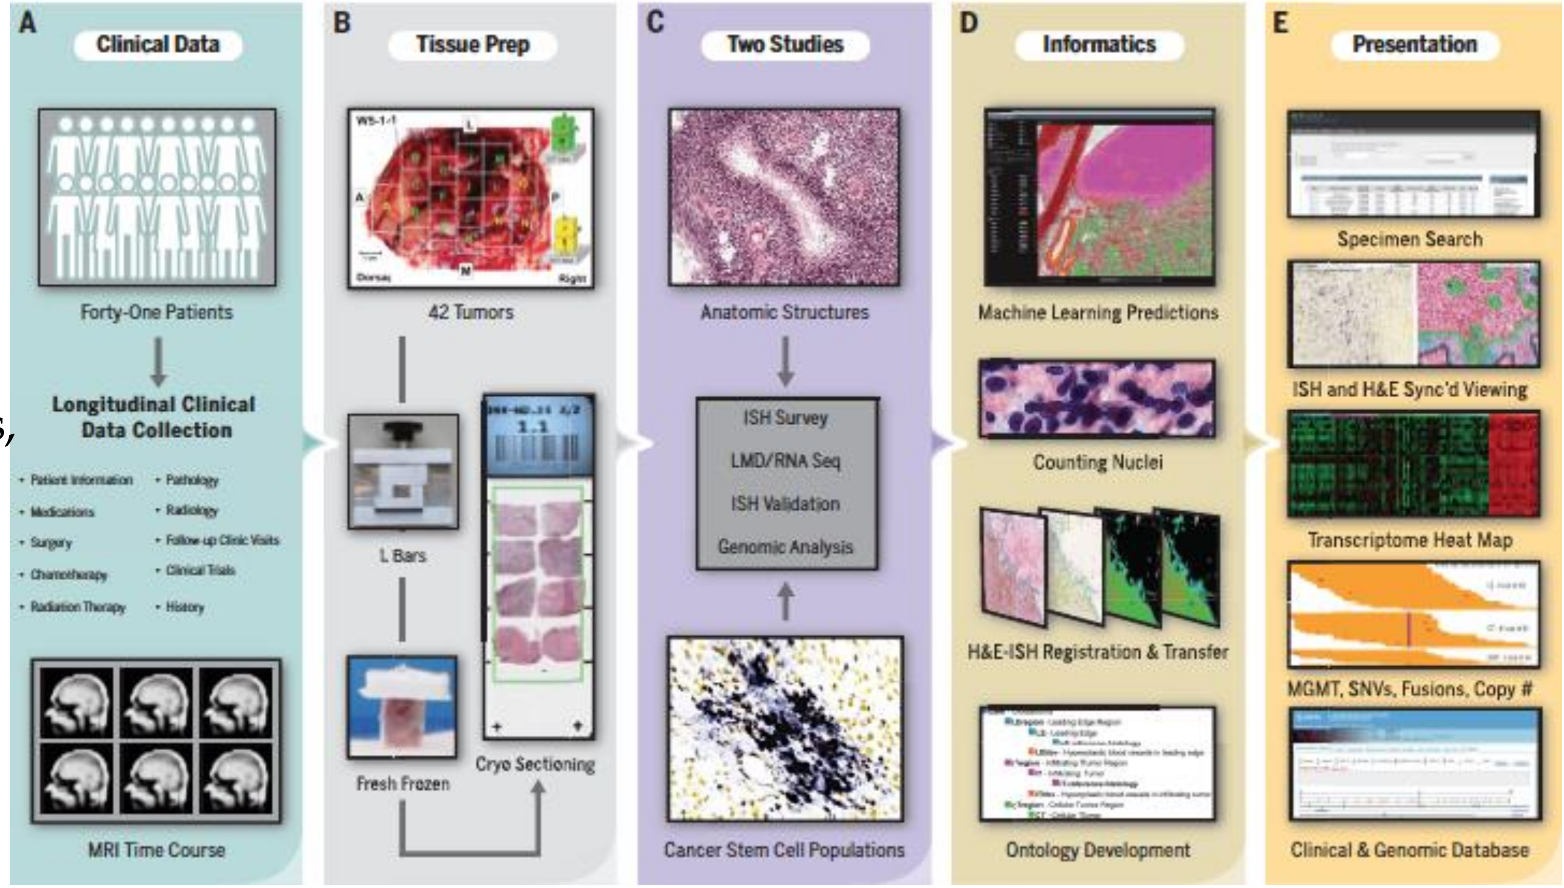

Puchalski et al Science 2018

ivygap.org

Segment anything in medical images

A foundation model for medical image segmentation

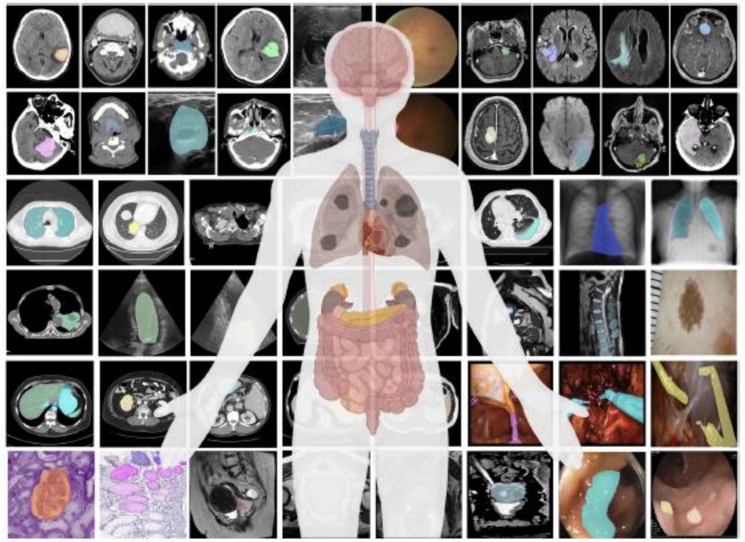

MedSam

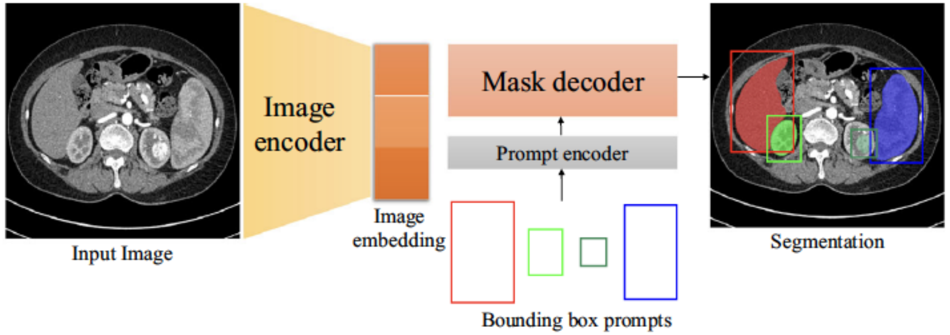

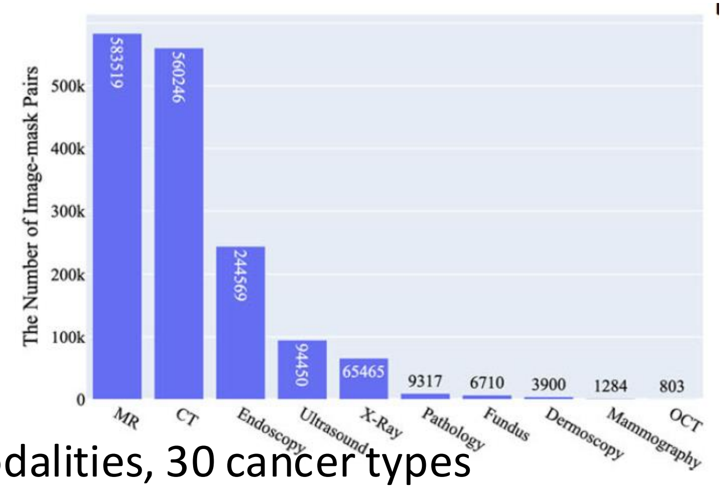

1.570.263 images, 10 imaging modalities, 30 cancer types

Ma et al Nat Comm 2024

Paradigm change

Gather many data

Ivy Gliobastoma atlas [Puchalski et al Science 2018]

41 patients, 42 tumors, 440 tissue blocks, 270 transcriptomes
11500 annotated H&E images
23000 ISH images (**400 Gb/image**), 400 MRI scans

Data for connectomics (neural networks): [Lichtman et al; Nat Neuro 2014]

1mm^3 rat cortex =>2M Gb =2 x10^{15}=2 Pb=2x10^3 TB
total cortex 500mm3=>10^{3}PB (1exabyte, 10^6 TB)
Man = 1000xlarger =10^3 exabyte (10^9 TB)
(source lichtman et la; Nat Neuro 2014)

Data rate

1mm^3 =>800h (33j) 2.5 Tb/h => 45y on one machine

[Miller et al; Nat Neuro 2016]

UK BioBank

100000 subjects (2016-22)
6 MR imaging modalities: tT1w, T2w, swMRI, dMRI, tfMRI, rfMRI
2501 individual measures of brain structure and function (2Gb p.sub)
1100 other non imaging variables
about **0.2 PB (200 TB)**

Dermatologists vs CNN [Esteva et al Nature 2017]

127463 biopsies for training
1942 for validation
Inception v3 (GoogleNet, using several replicas on Nvidia Titan X Gpu)

Effect of training dataset size

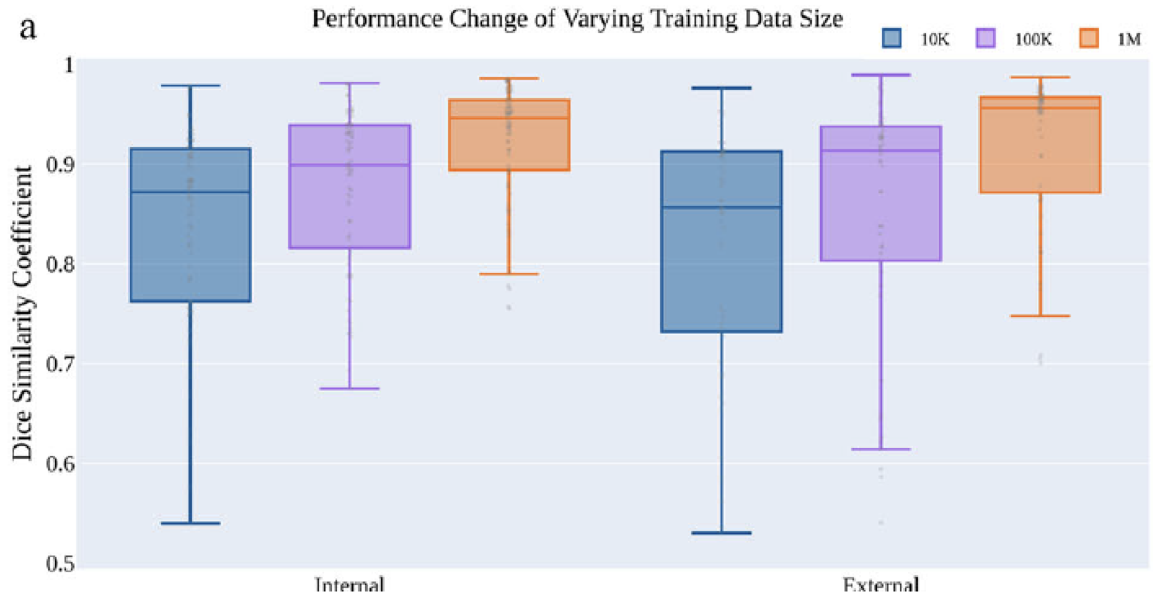

Scaling up the training image size to one million can significantly improve the model performance on both internal and external validation sets.

Ma et al Nat Comm 2024

Generation of artificial examples

Introduce morphing operations, symetry, rotation ….

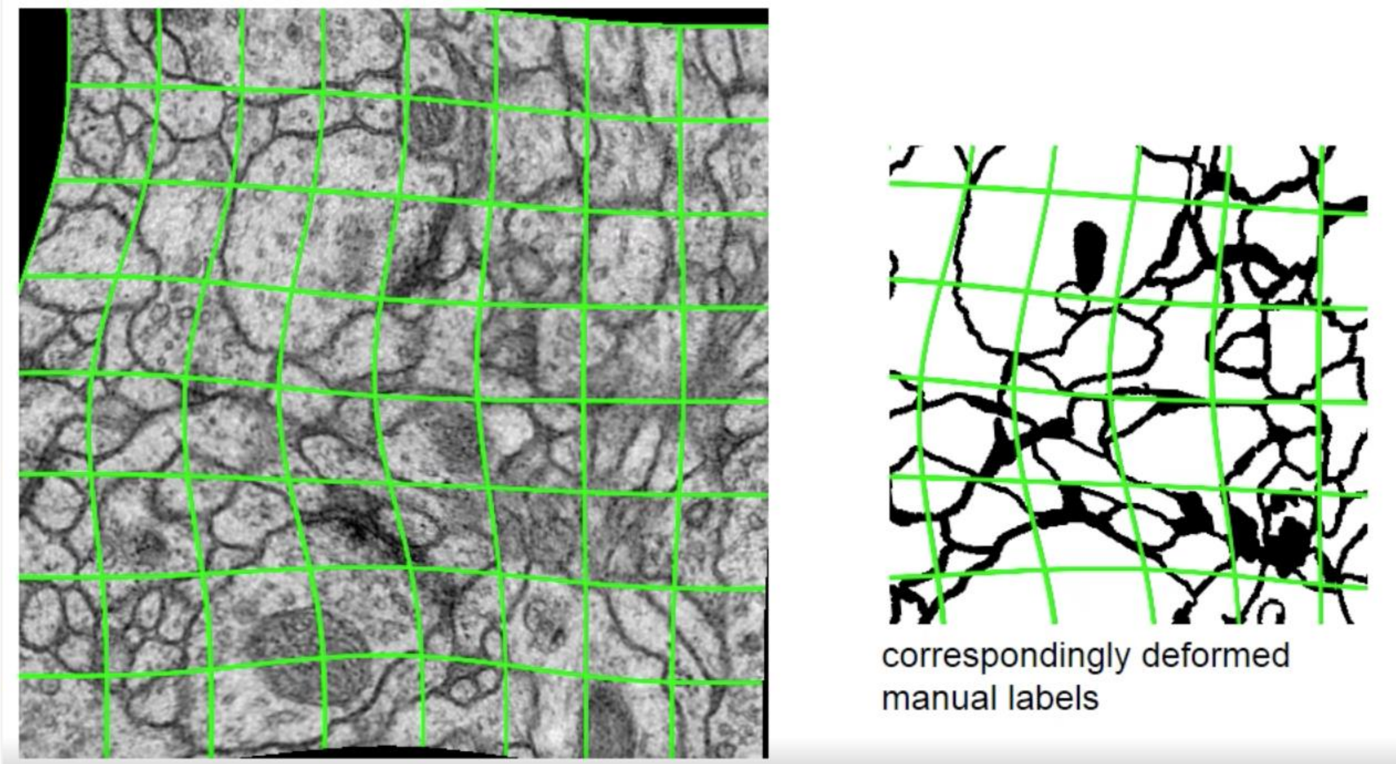

Train the NN to learn invariance to such elastic deformations

Automatic pulmonary module management

Create examples

Ciompi et al Scient Rep 2017

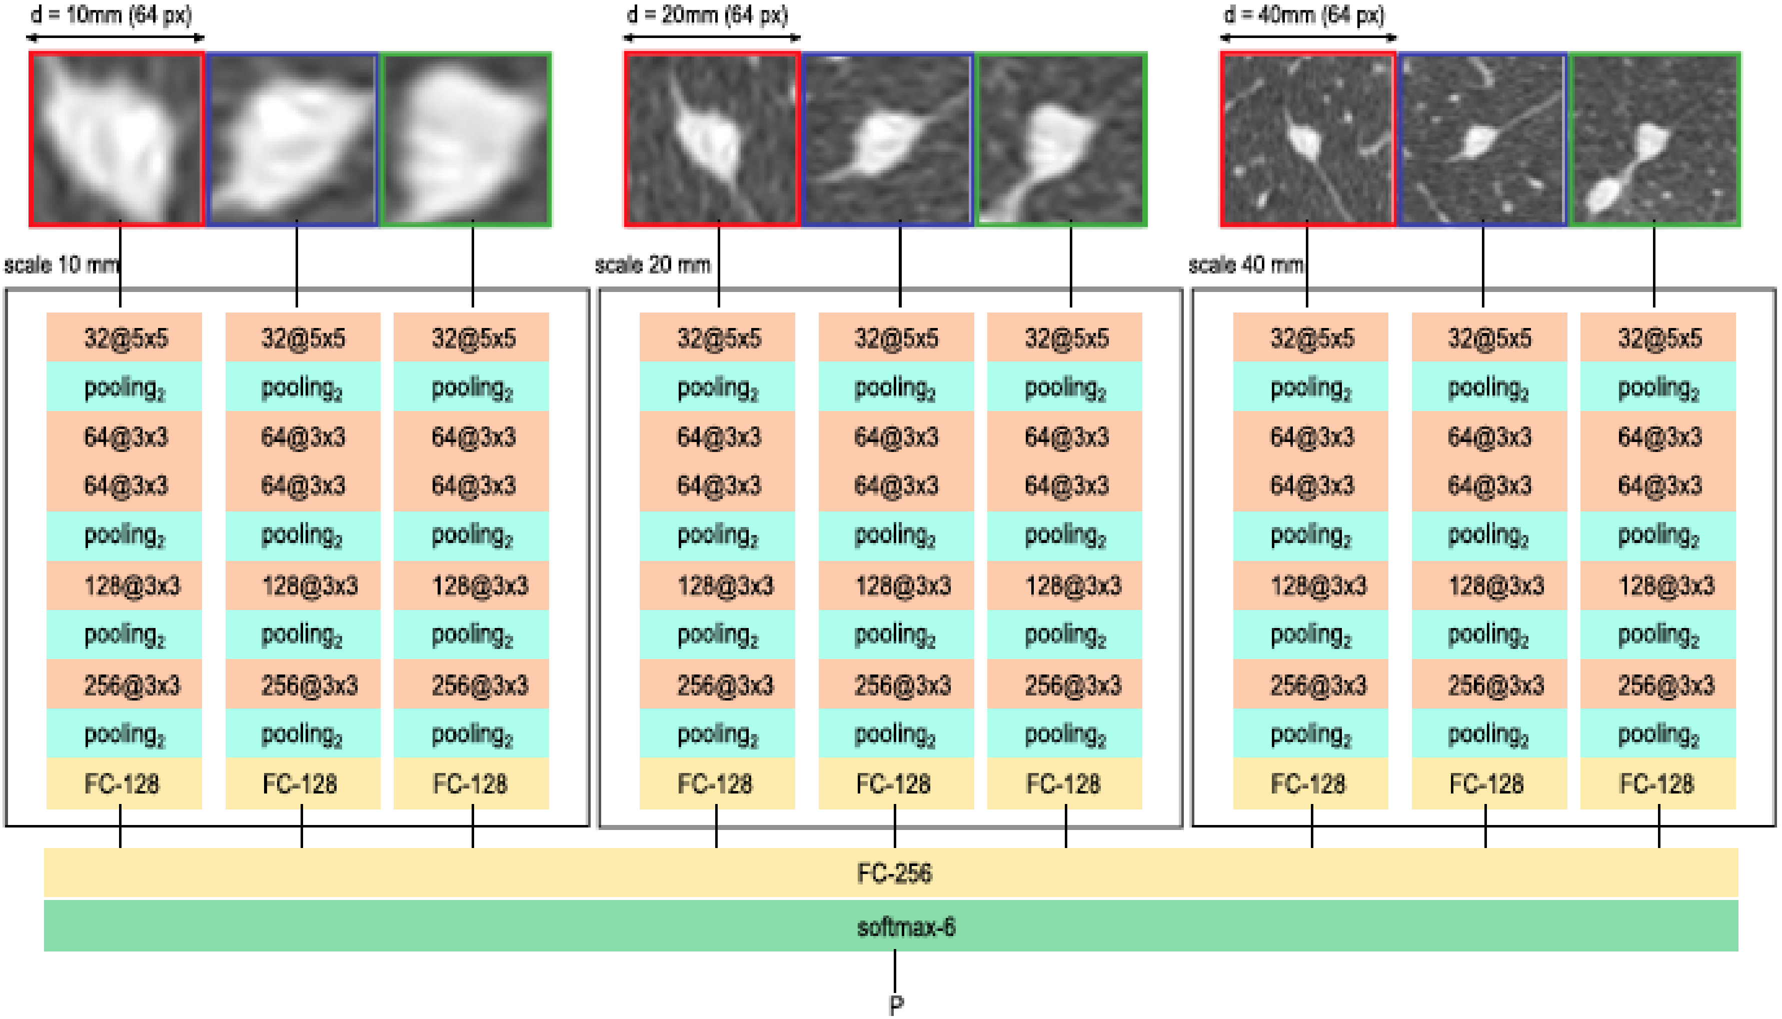

Automatic pulmonary module management

Ciompi et al Scient Rep 2017

	MILD (943 patients)				DLCST (468 patients)
	Training nodules	N	Training samples	Validation nodules	Test nodules $test_{ALL}/test_{OBS}$
Solid	694	8	88,832	232	382/27
Calcified	233	22	82,016	78	58/27
Part-solid	63	80	80,640	21	37/27
Non-solid	152	33	80,256	50	87/27
Perifissural	181	28	81,088	62	48/27
Spiculated	29	167	77,488	10	27/27
Total	1,352	—	490,320	453	639/162

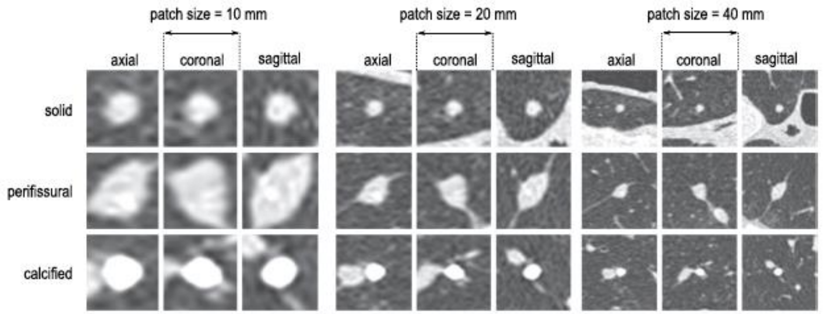

Automatic pulmonary module management

Ciompi et al Scient Rep 2017

	Accuracy	F_{solid}	$F_{Calcified}$	$F_{Part\text{-}solid}$	$F_{non\text{-}solid}$	$F_{Perifissural}$	$F_{Spiculated}$	$F_{Not\text{-}a\text{-}nodule}$
O_1 vs. Computer (3 scales)	71.5%	60.8%	88.4%	66.7%	86.3%	62.2%	71.4%	—
O_2 vs. Computer (3 scales)	66.2%	62.6%	82.4%	47.8%	72.7%	80.0%	56.4%	—
O_3 vs. Computer (3 scales)	67.7%	56.8%	85.1%	59.1%	78.3%	75.6%	60.9%	—
O_4 vs. Computer (3 scales)	72.8%	64.2%	88.9%	71.7%	80.0%	77.3%	62.7%	—
Average	69.6%	61.1%	86.2%	61.3%	79.3%	73.8%	62.9%	—
O_1 vs. O_2	66.0%	52.7%	84.0%	51.3%	79.2%	63.6%	83.3%	50.0%
O_1 vs. O_3	71.0%	55.0%	87.0%	66.7%	80.0%	81.5%	74.4%	40.0%
O_1 vs. O_4	72.8%	64.8%	90.9%	66.7%	71.7%	75.5%	89.4%	0.0%
O_2 vs. O_3	76.5%	74.7%	88.9%	61.5%	81.0%	77.3%	75.7%	66.7%
O_2 vs. O_4	72.2%	64.4%	88.5%	70.8%	71.1%	79.1%	73.2%	0.0%
O_3 vs. O_4	79.0%	68.4%	95.8%	71.1%	80.9%	90.6%	79.2%	0.0%
Average	72.9%	63.3%	89.2%	64.7%	77.3%	77.9%	79.2%	26.1%

Ilastik

Create examples

Interactive training
Shallow learning can be trained interactively with immediate feedback
Adapting from similar training data

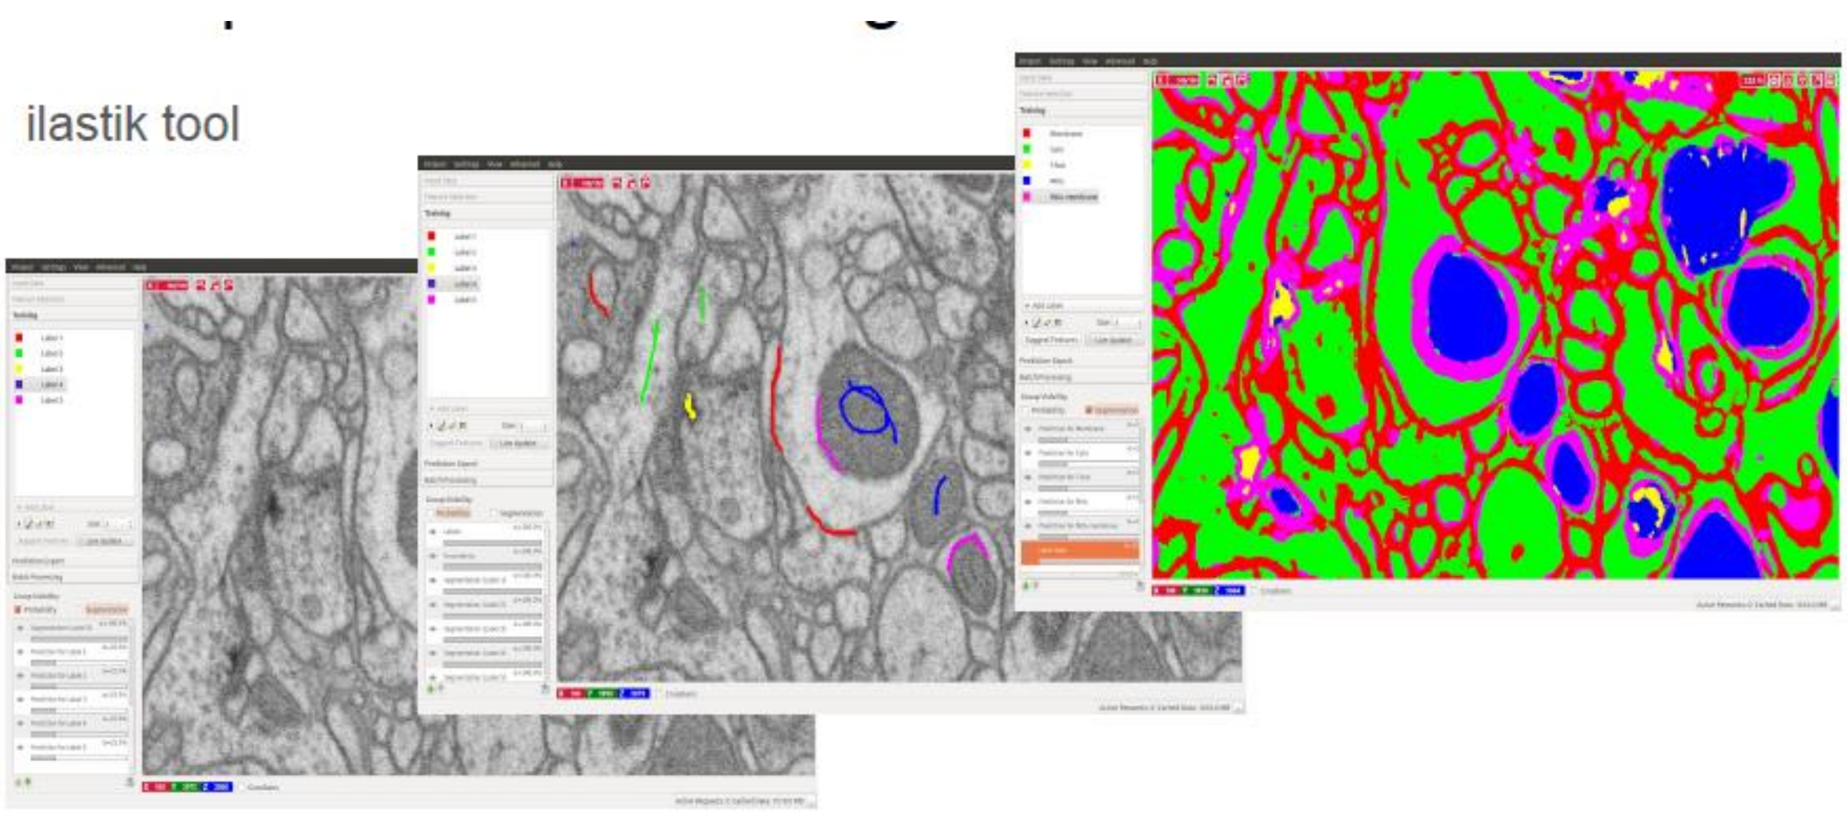

Transfert learning

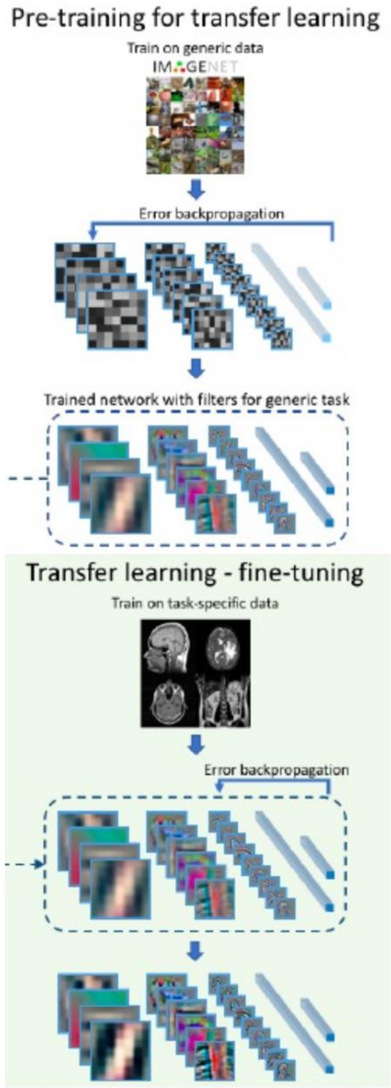

Mazurowski et al. JMRI 2018

Tranfert learning

T. Coudert Master 2021

- Few examples : 9 for training / 11 for testing

Dataset	WMH DSC	Sensitivity	Precision	95HD (mm)
9 EpiBrainRad LR=5e-3	0.557 ± 0.203	0.507 ± 0.229	0.703 ± 0.228	21.089 ± 15.469
9 EpiBrainRad LR=1e-5	0.705 ± 0.088	**0.869 ± 0.140**	0.624 ± 0.144	**6.838 ± 4.468**
9 EpiBrainRad dropout=0.3	**0.740 ± 0.127**	0.719 ± 0.182	**0.799 ± 0.112**	13.105 ± 17.237
9 EpiBrainRad max epochs set to 150	0.578 ± 0.257	0.578 ± 0.275	0.623 ± 0.258	19.848 ± 24.362

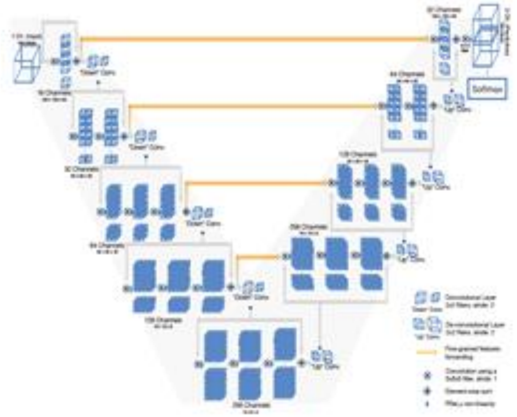

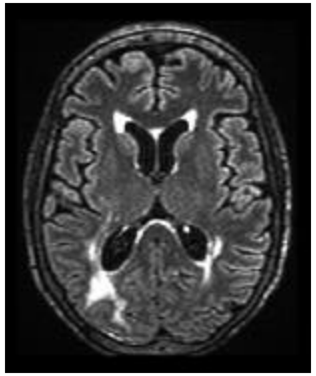

FLAIR

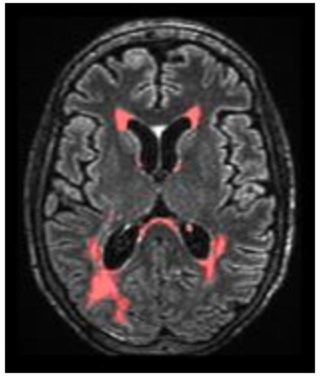

Ground truth

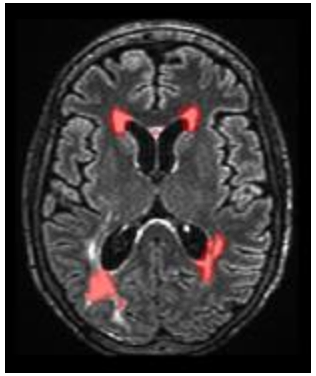

Prediction

Use a pretrained network

Transfert Learning

T. Coudert Master 2021

Gliomas data from **BraTS Challenge**[1] dataset
2016 and 2017

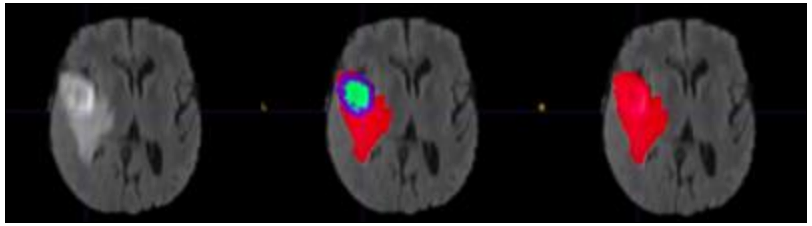

	Dataset distribution	WMH DSC	Sensitivity	Precision	95HD (mm)
without	200 BraTS + 9 EpiBrainRad into one training	0.836 ± 0.075	0.879 ± 0.106	0.803 ± 0.073	10.246 ± 12.439
with	200 BraTS pre-trained model, transferred on 9 EpiBrainRad	$\mathbf{0.858 \pm 0.064}$	$\mathbf{0.890 \pm 0.104}$	$\mathbf{0.835 \pm 0.057}$	$\mathbf{2.806 \pm 3.523}$

test set of 11 images with **13% of variability with a neurologist segmentation** on this 11 images

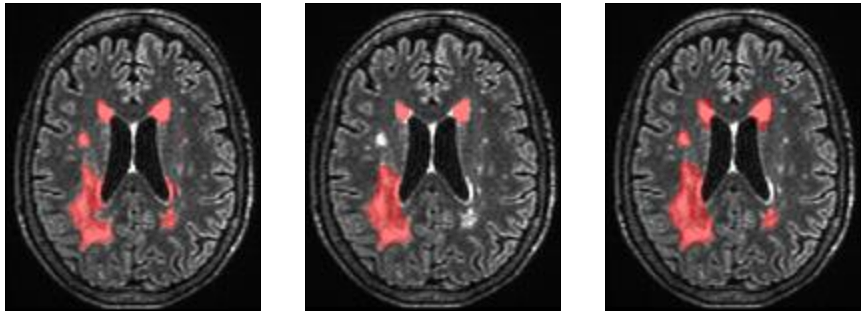

Ground Truth **Without TL** **With TL**

Clear benefit of transfer learning, **best results with exactly the same amount of input images**

[1] B. H. Menze et al. **The multimodal brain tumor image segmentation benchmark (brats)**. IEEE Transactions on Medical Imaging, 34(10):1993–2024, 2015.

Augmented training

PJ Lartaud PhD 2022

Synthetic data

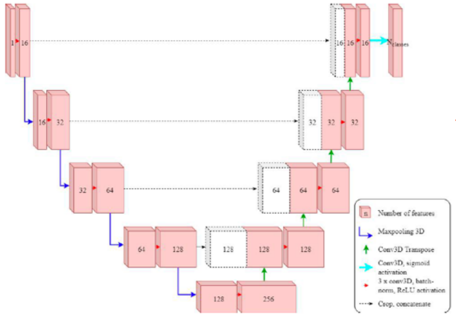

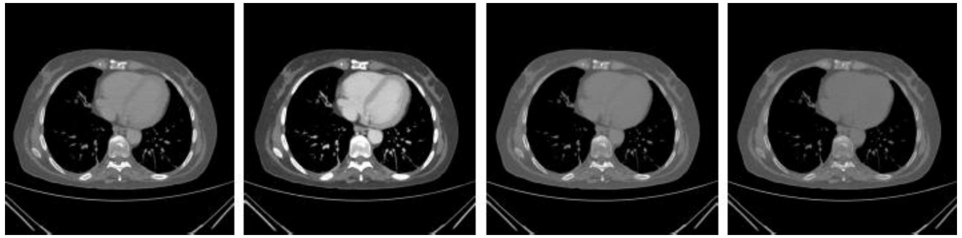

classic spectral

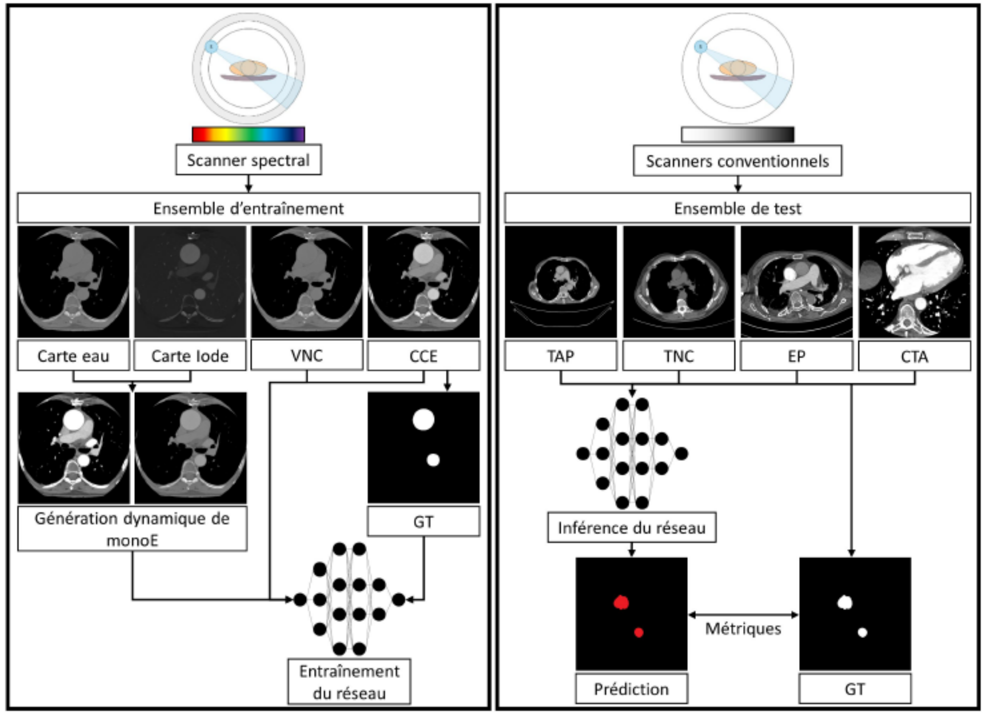

classic spectralAug.

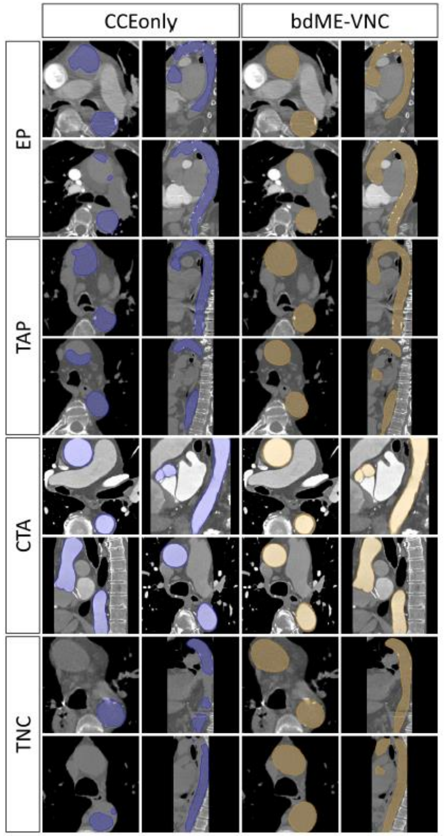

Image Generation

PJ Lartaud PhD 2022

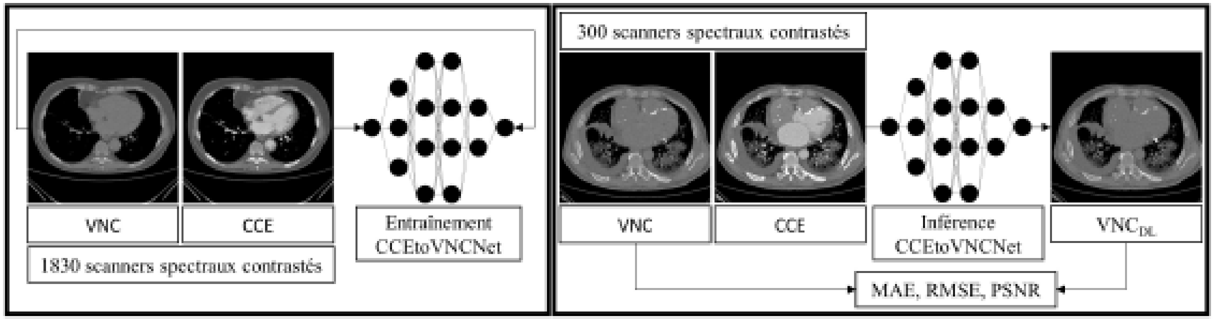

Generation
Spectral data

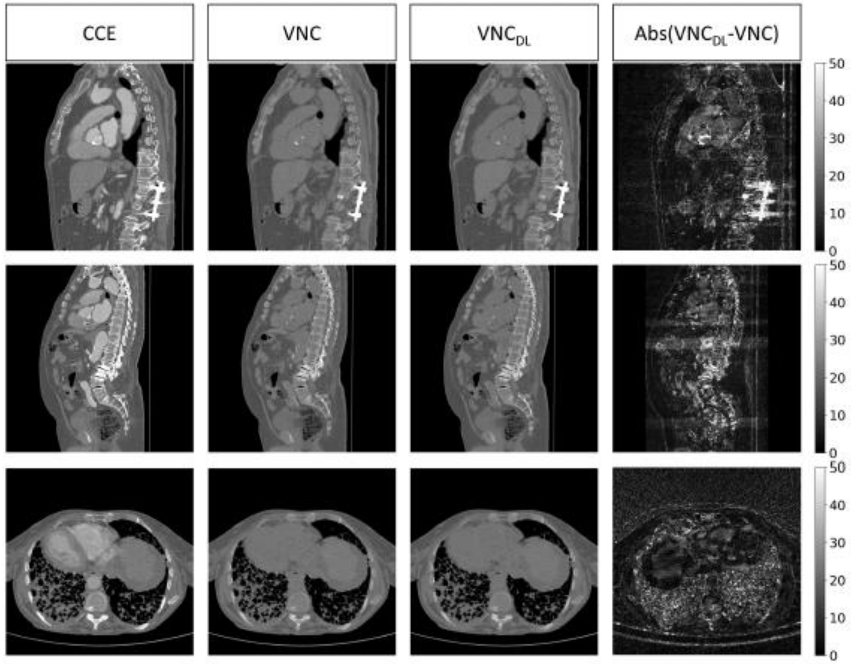

Results

Image Generation
PJ Lartaud PhD 2022

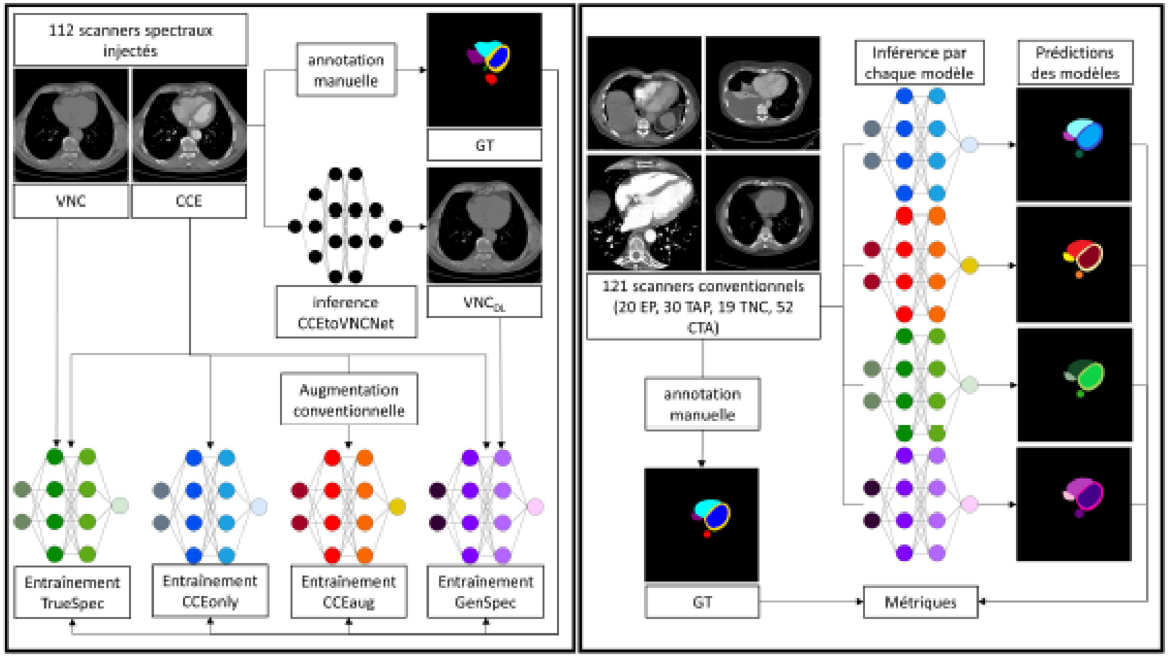

Training with artificial data …

Results

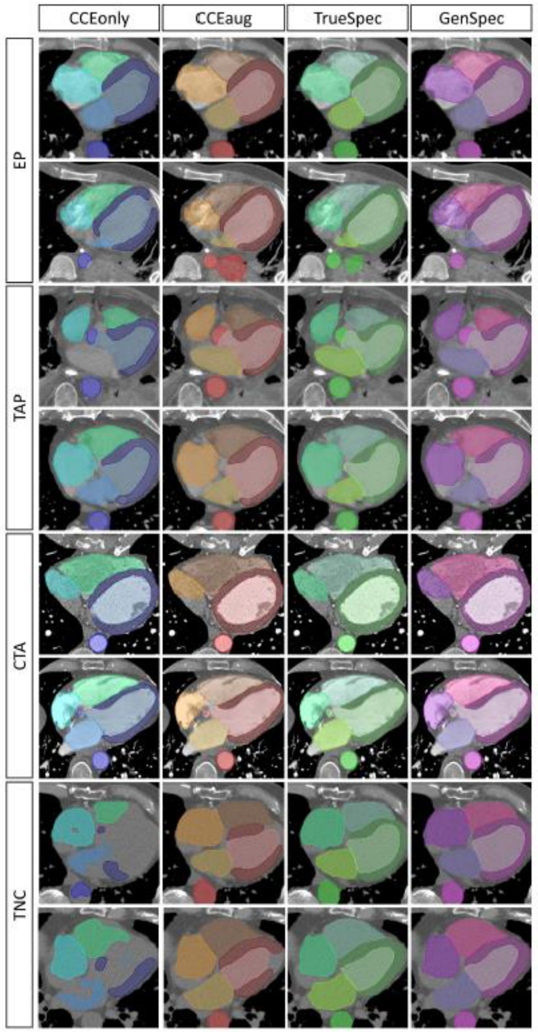

Producing new data

Synthetic data

- Generative Adversarial Networks (GAN): produces output undistinguable from “real” images.

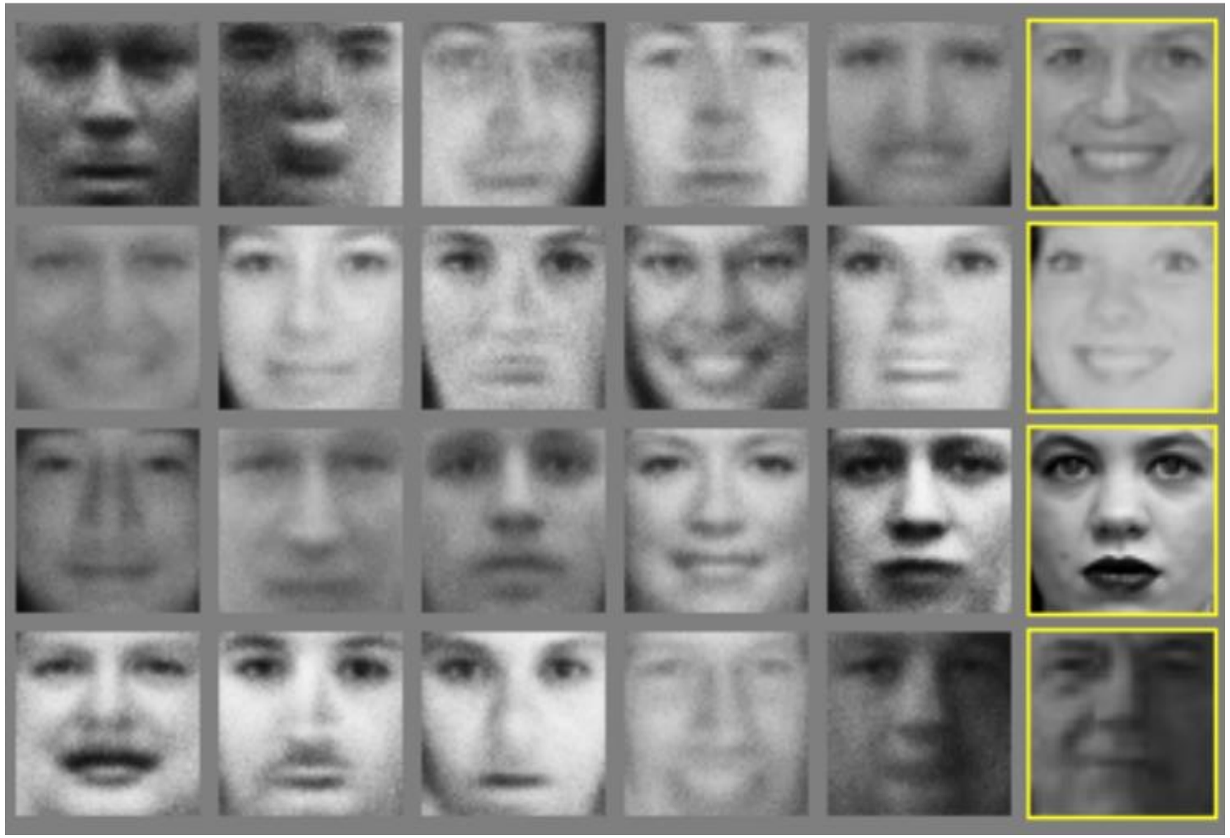

Goodfellow et al
2014
arXiv:1406.2661

GAN: Generative Adversarial Network

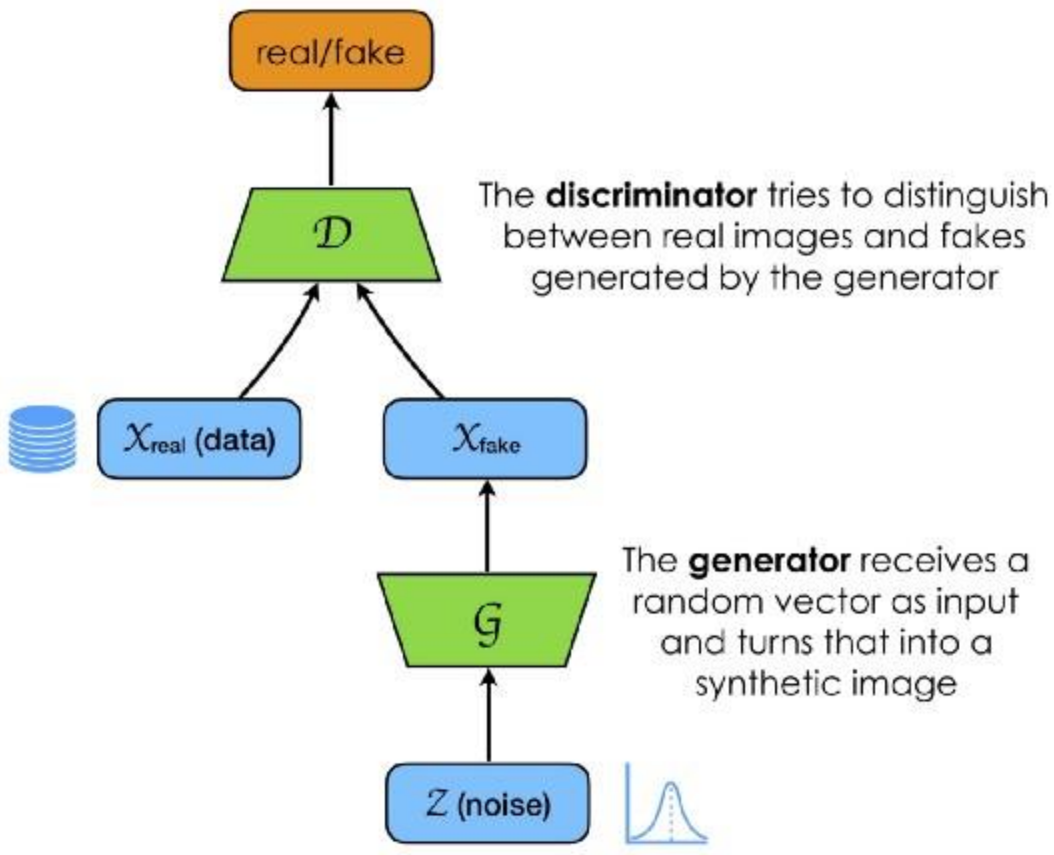

Role of GAN in HealthCare

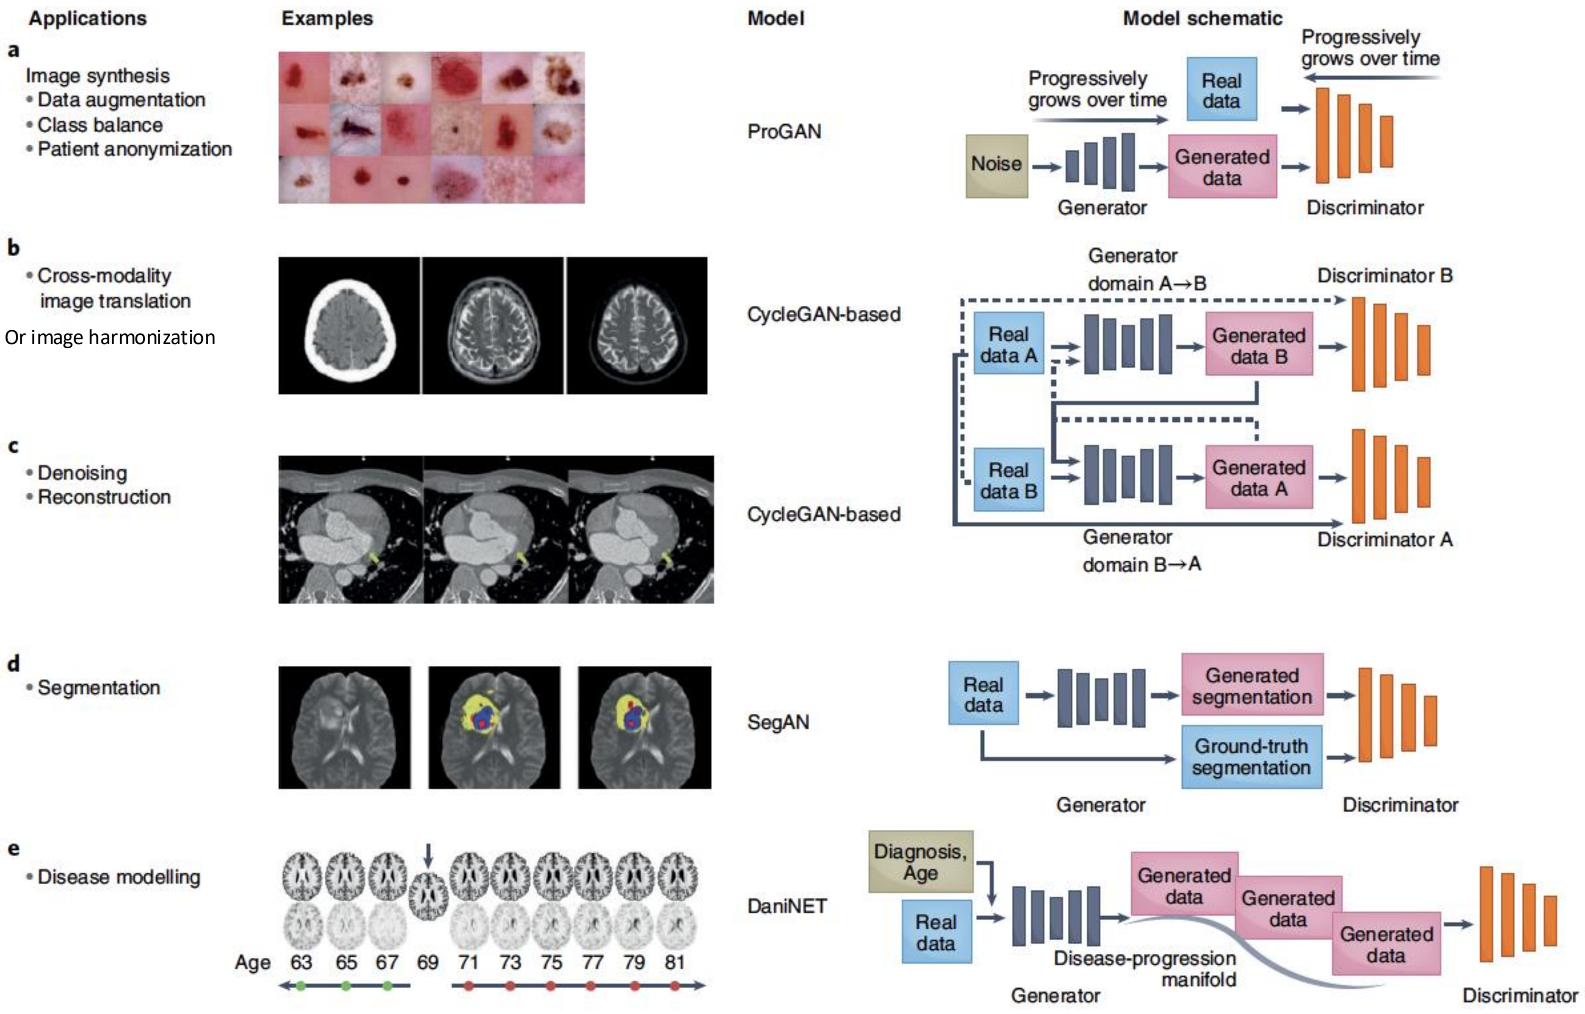

Zhang et al Nat Med Eng 2022

Generative Adversarial Network

Liu et al Radiology 2017

To generate realistic synthetic images

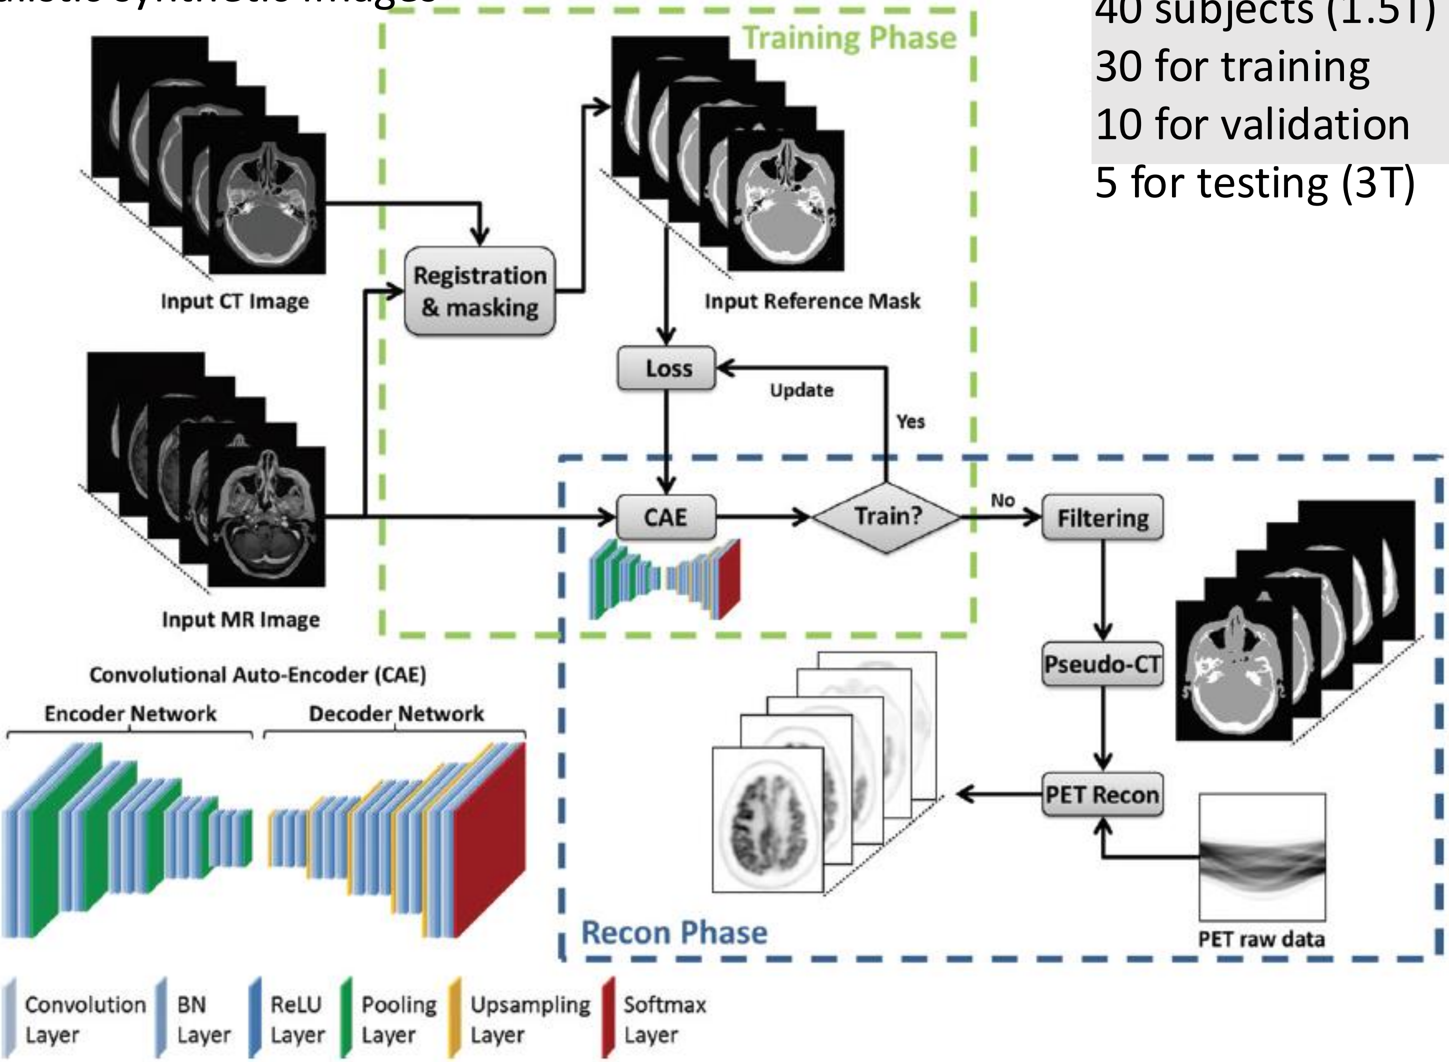

40 subjects (1.5T)
30 for training
10 for validation
5 for testing (3T)

13 layers
VGG16
Filters 64 to 512
Stride: 2x2

Generative Adversarial Network

To generate realistic synthetic images

A cascade of GANs

32 32 32 64 64 64 32 32 32

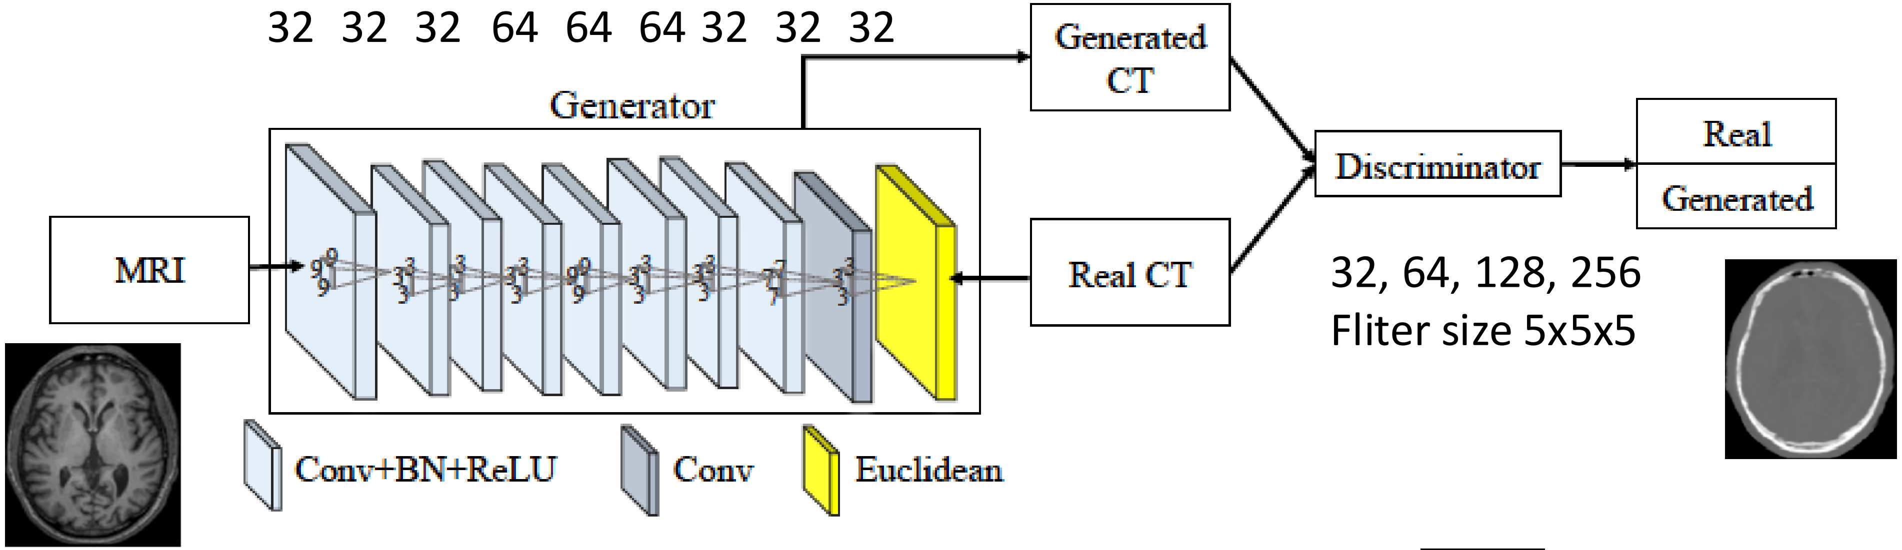

32, 64, 128, 256
Fliter size 5x5x5

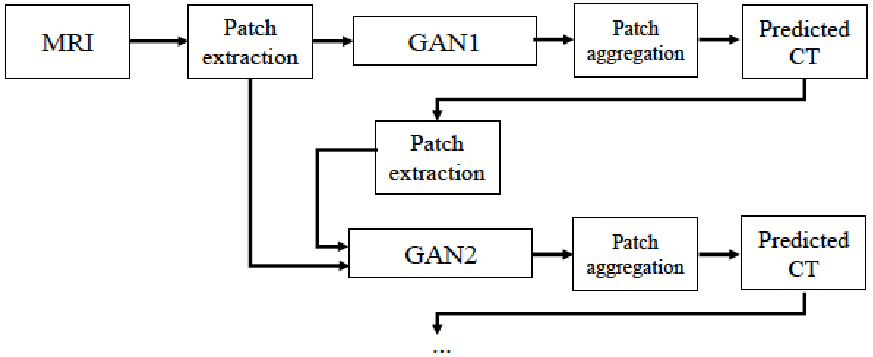

Nii et al Miccai 2017

Image Harmonization - I

Pooling data

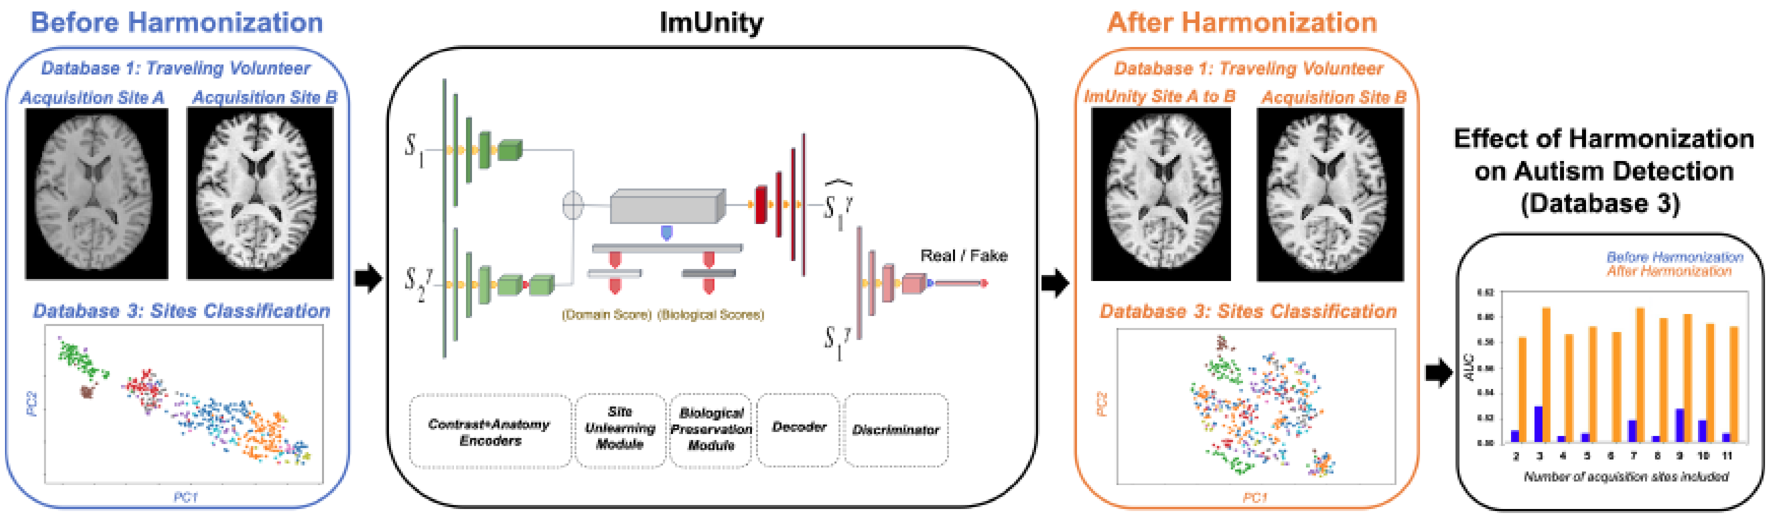

Cakowsky et al 2022 AIM

Image Harmonization - II

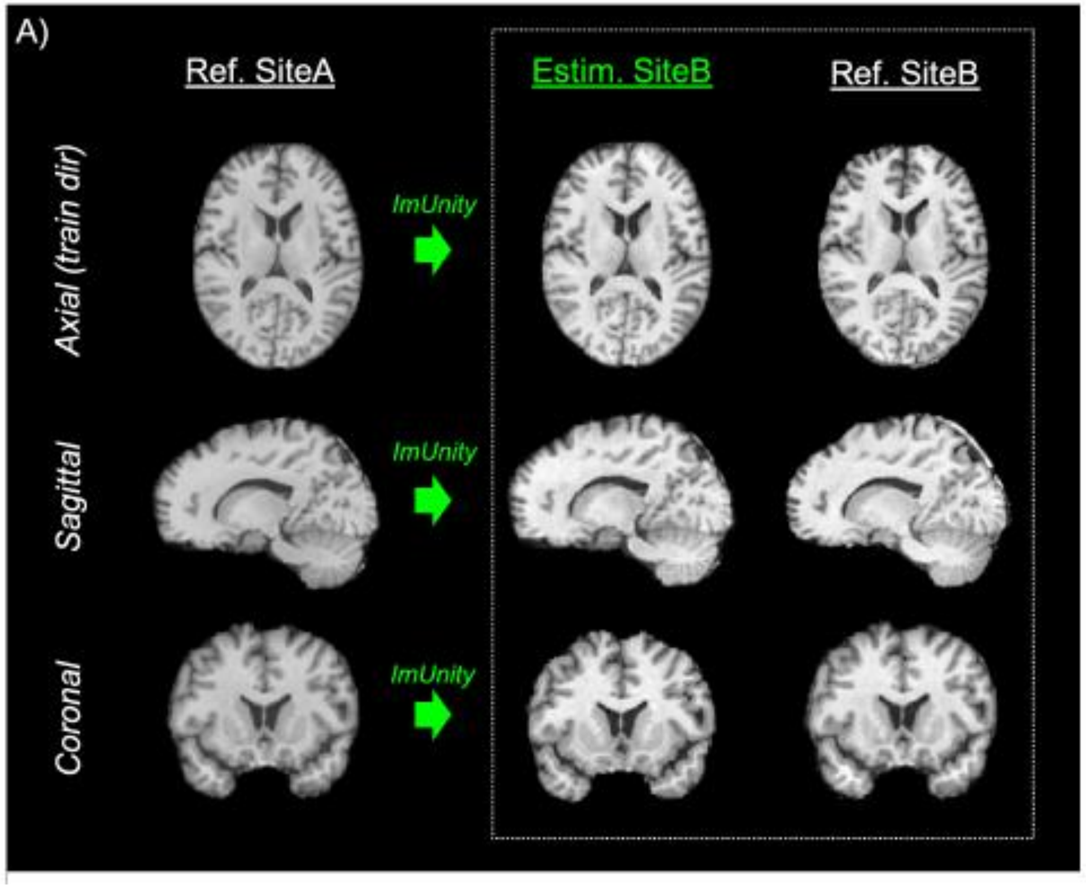

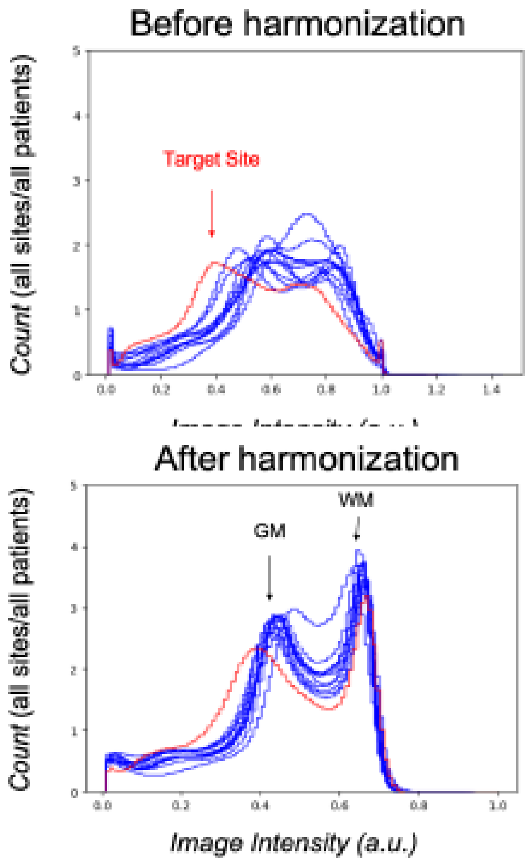

Image Harmonization - III

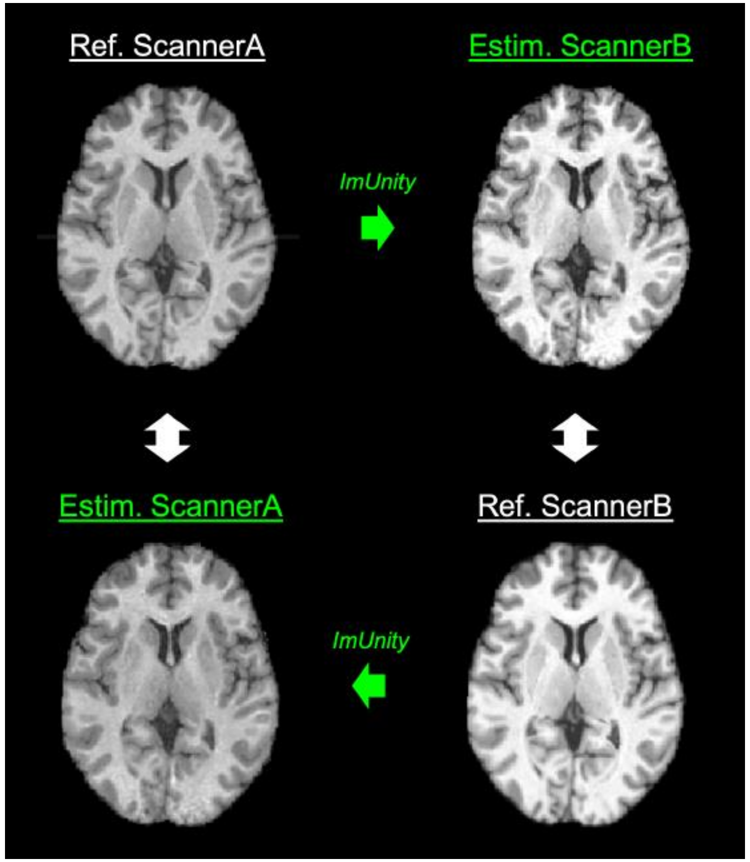

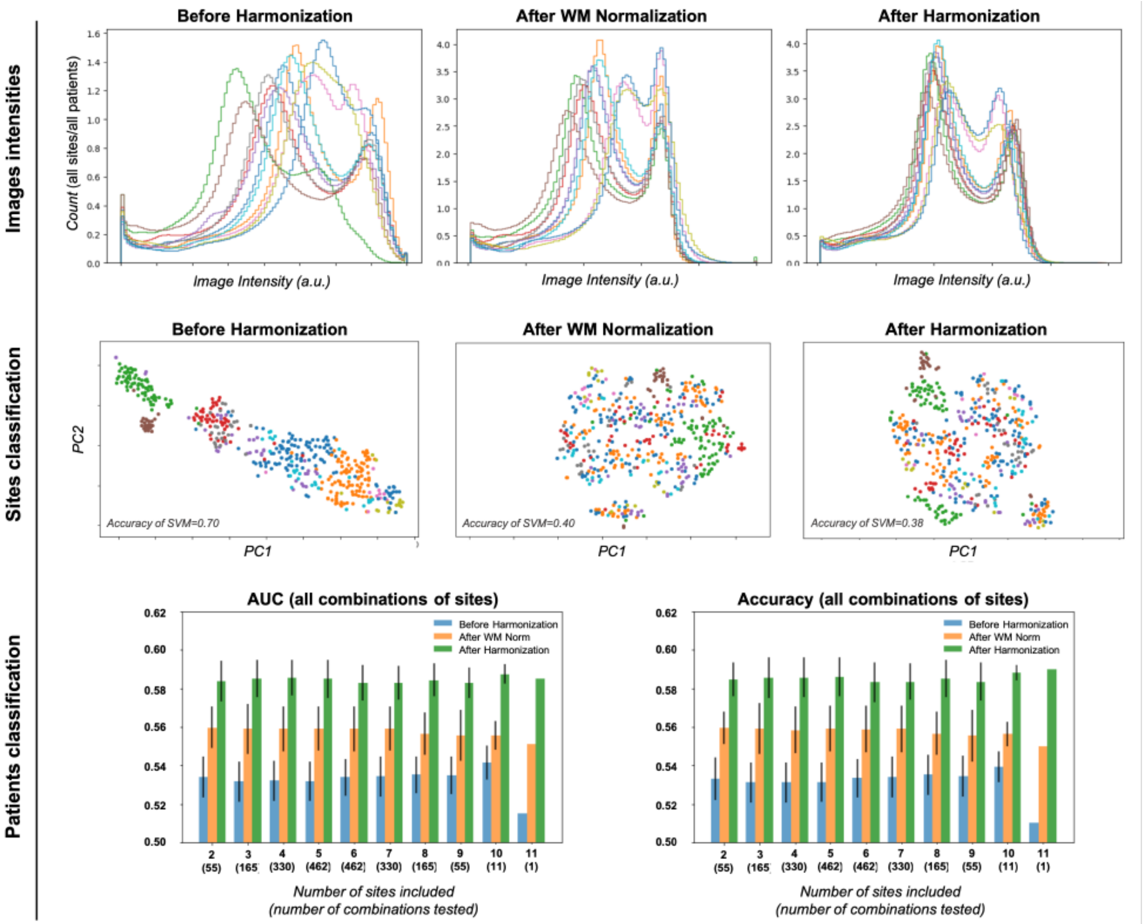

Variational Auto-encoder: Anomalies detection

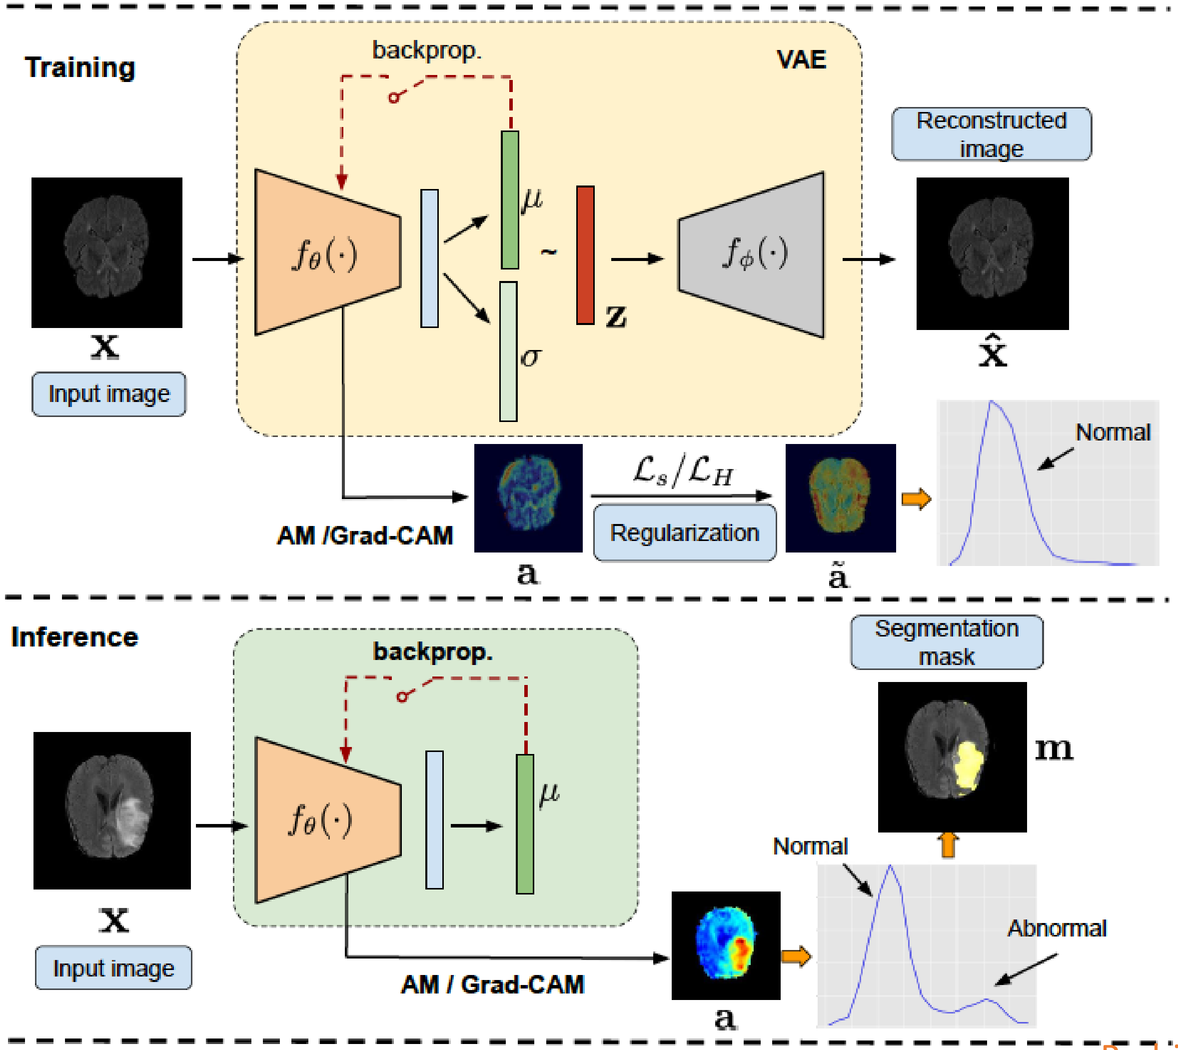

Rodriguez et al Media 2022

Parkinson

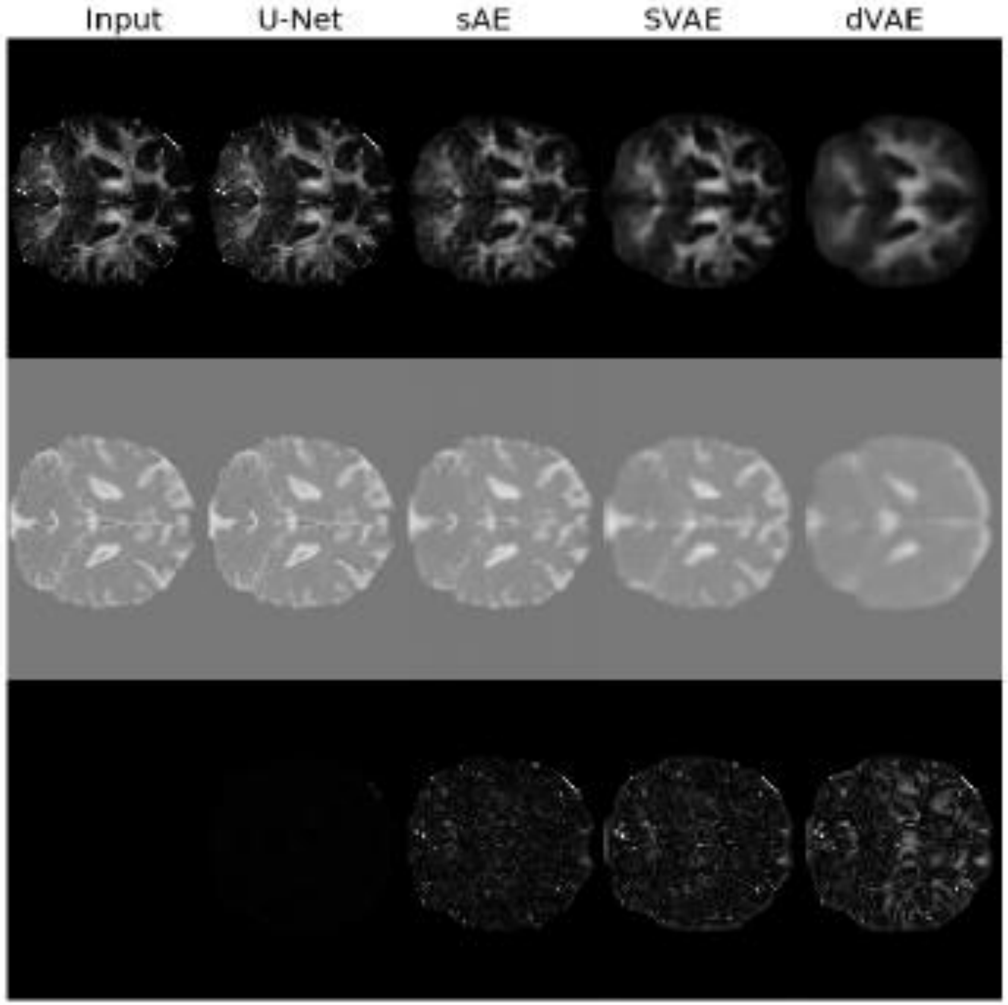

	U-Net	sAE	sVAE	dVAE
Subcortical structures/Brainstem	**0.70**	**0.76**	**0.77**	**0.70**
Frontal lobe	0.56	0.58	0.56	0.59
Parietal lobe	0.45	0.62	0.62	0.62
Temporal lobe	0.59	0.67	**0.72**	0.69
Occipital lobe	0.35	0.43	0.41	0.43
White matter	0.62	**0.80**	**0.78**	**0.76**
Insular/Cingulate cortex	**0.70**	0.68	**0.76**	0.68
Ventricles	0.68	0.53	0.50	0.52
Red Nucleus	0.51	0.69	0.64	0.58
Substantia Nigra	0.61	**0.73**	**0.72**	0.66
Subthalamic Nucleus	0.67	0.69	0.66	0.50
Caudate	0.60	0.57	0.53	0.62
Putamen	0.46	0.62	0.61	0.60
Globus Pallidus external	0.54	0.56	0.58	0.60
Globus Pallidus internal	0.54	0.64	0.67	0.64
Thalamus	0.67	0.65	**0.73**	0.68
Superior Colliculus	0.66	0.45	0.54	0.52
Inferior Colliculus	0.55	0.40	0.42	0.43

Table 4. ROC AUC contemplating the eight regions (top) and the ten sub-cortical structures (bottom)

Kmetzsch et al 2019

Beyond CNN

- Limitations
  - No model of long spatial dependencies (limit by kernel)
  - Fixed weights at the inference time
- Tranformers
  - To manage
  - VisioTranf

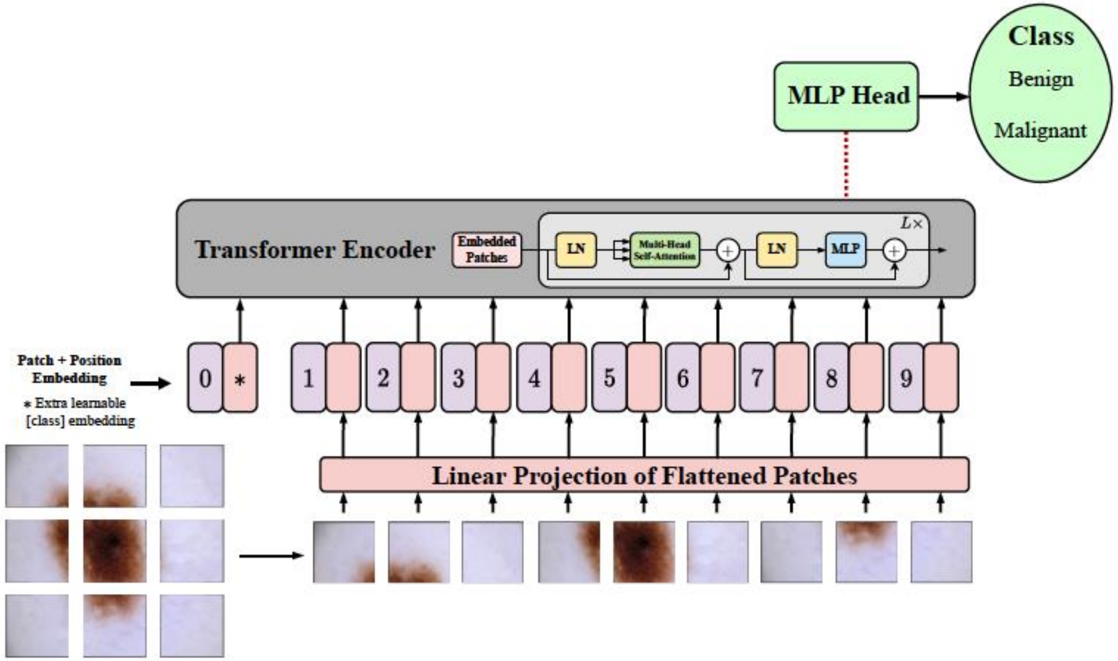

(Azad et al 2023 arXiv)

An example

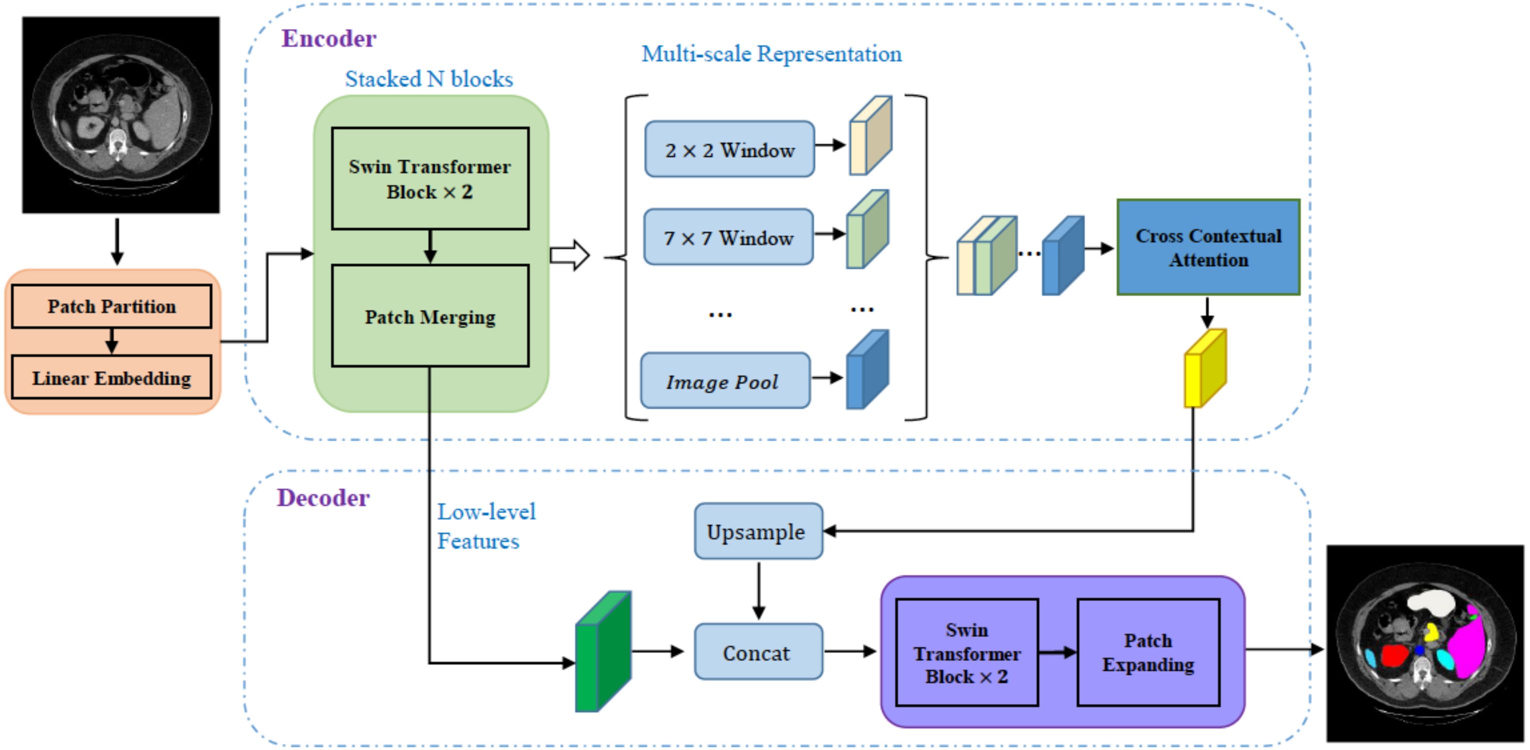

DeepLabv3++

(Azad et al 2023 arXiv)

Unsupervised Learning

On Mimic-cxr

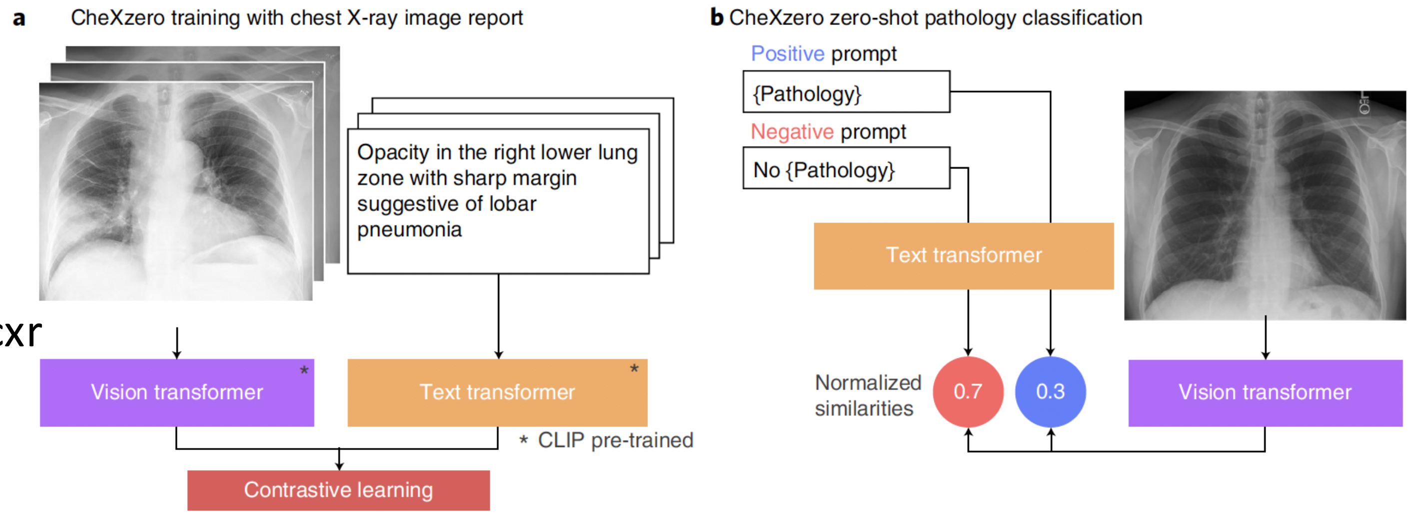

On cheXpert test

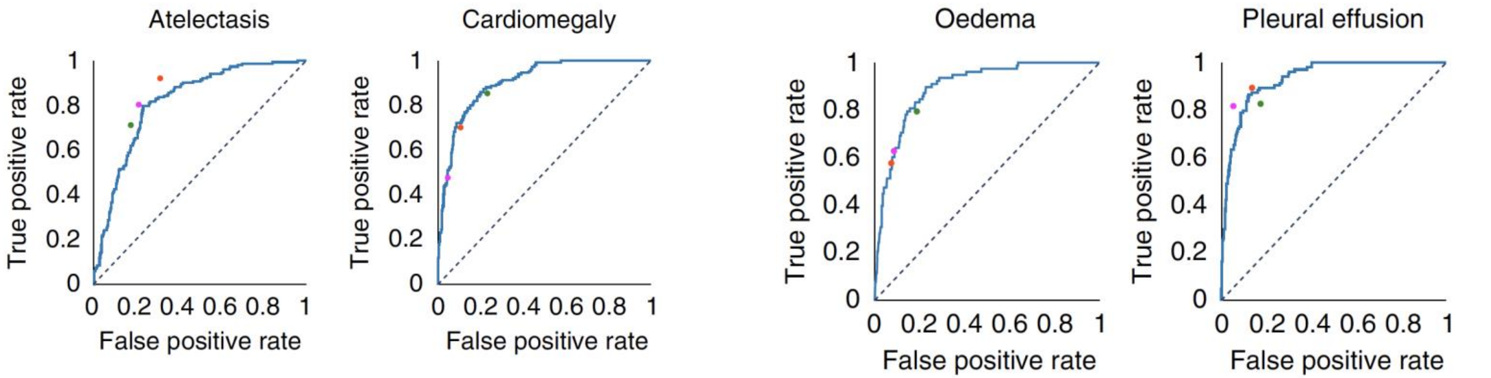

Tiu et al Nature Bio Eng 2022

• Radiologist 1 • Radiologist 2 • Radiologist 3

Data pipeline for ML

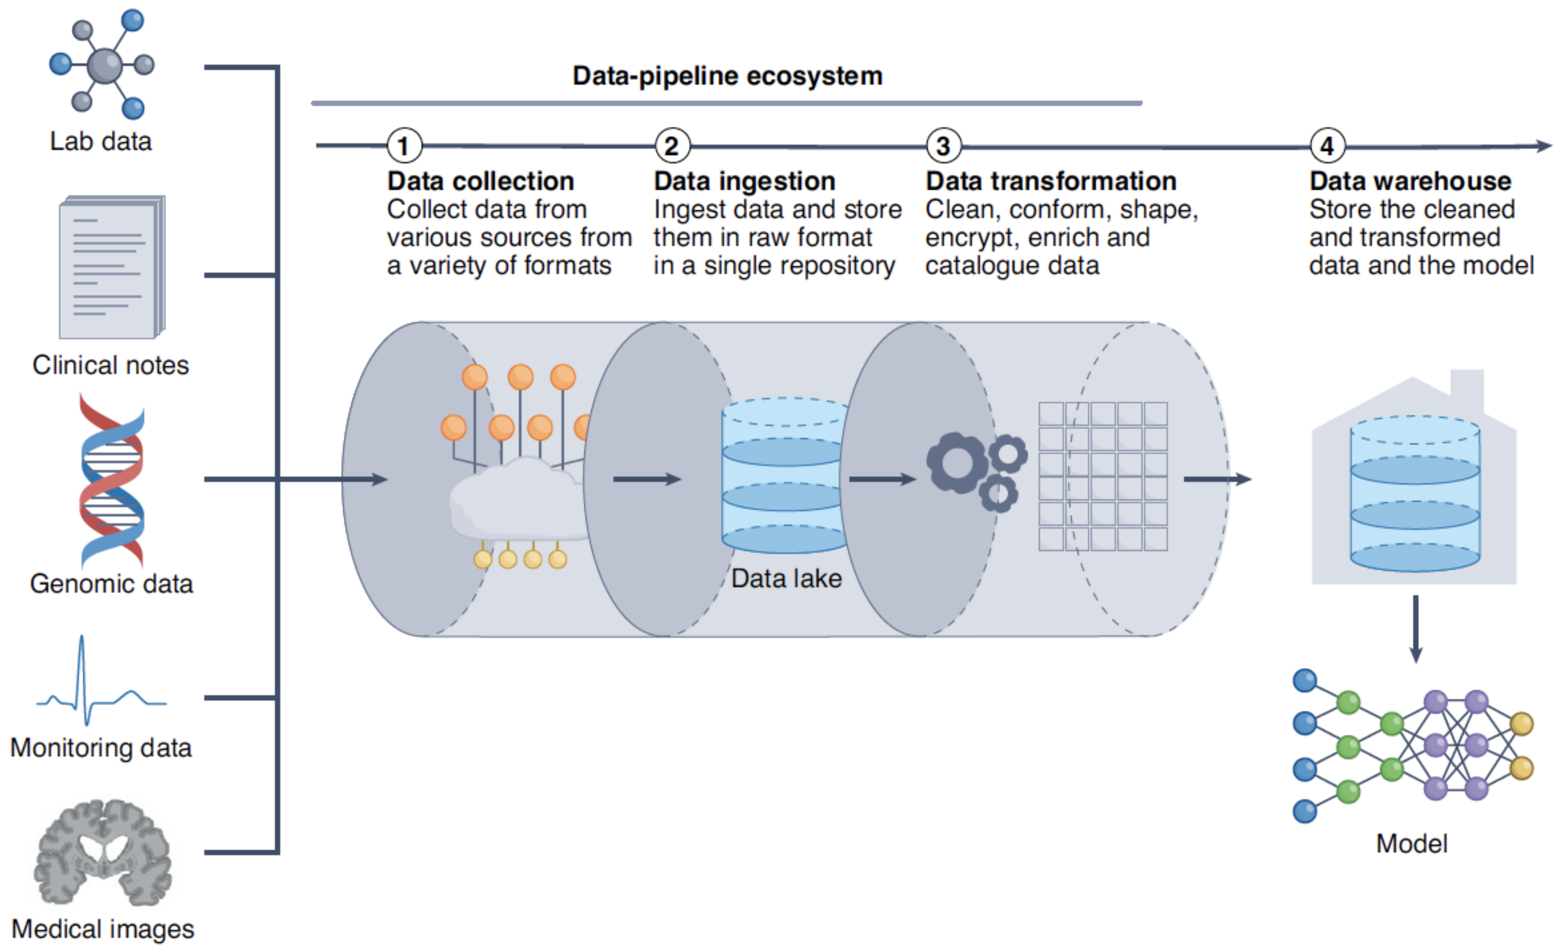

Zhang et al Nat Med Eng 2022

Specific ethical considerations - I

- Data collection and respect private for life
- GPRD, securisation

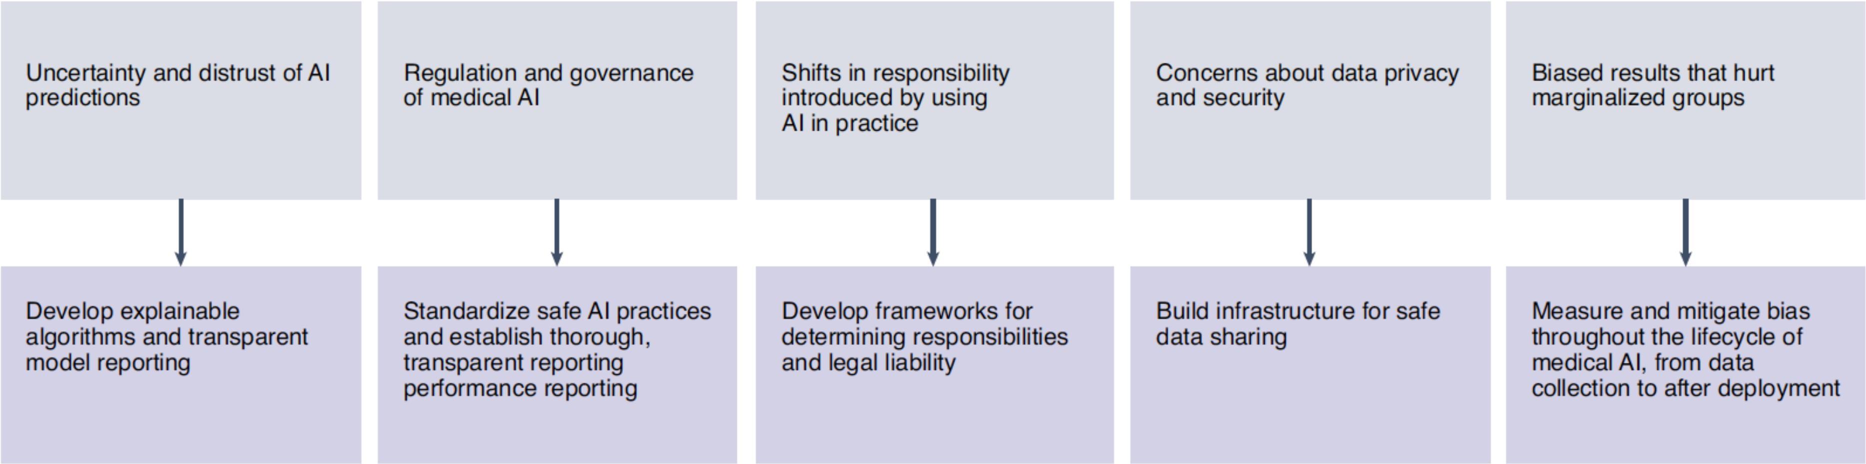

Rajpurkar et al 2022 Nat Med

Specific ethical considerations - II

- Data collection and respect private for life
- GPRD, securisation

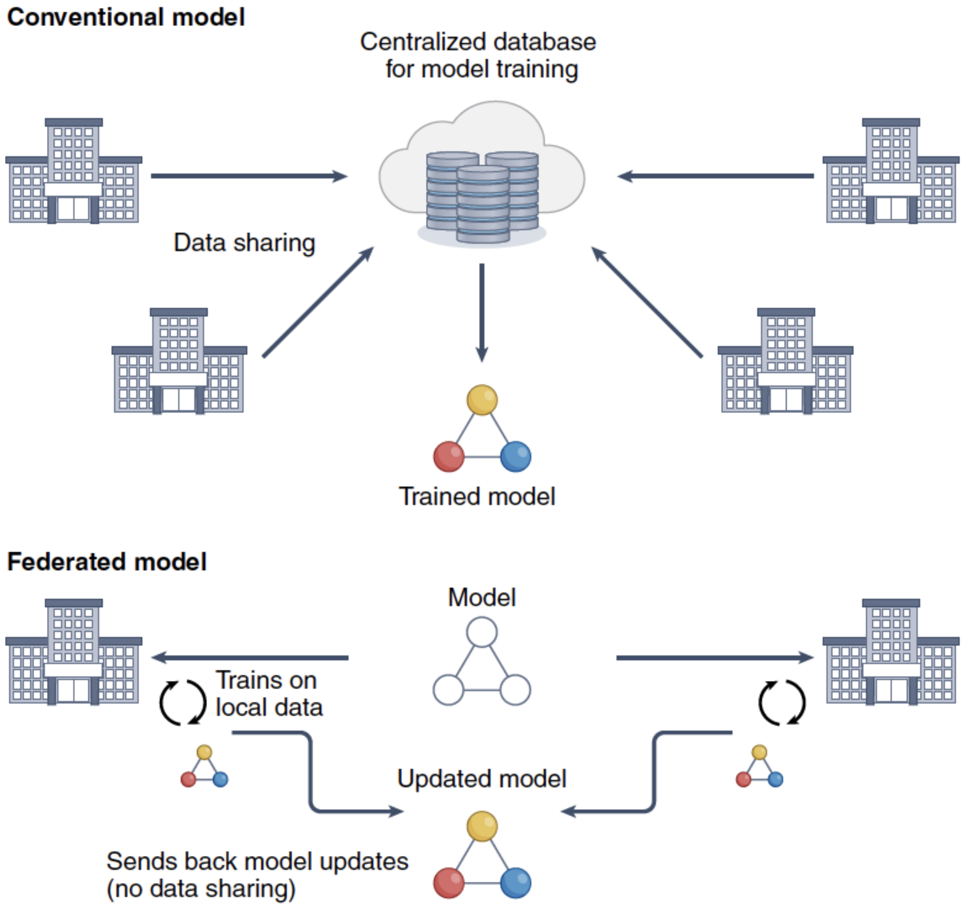

Rajpurkar et al 2022 Nat Med

Cross-silo federated learning for healthcare

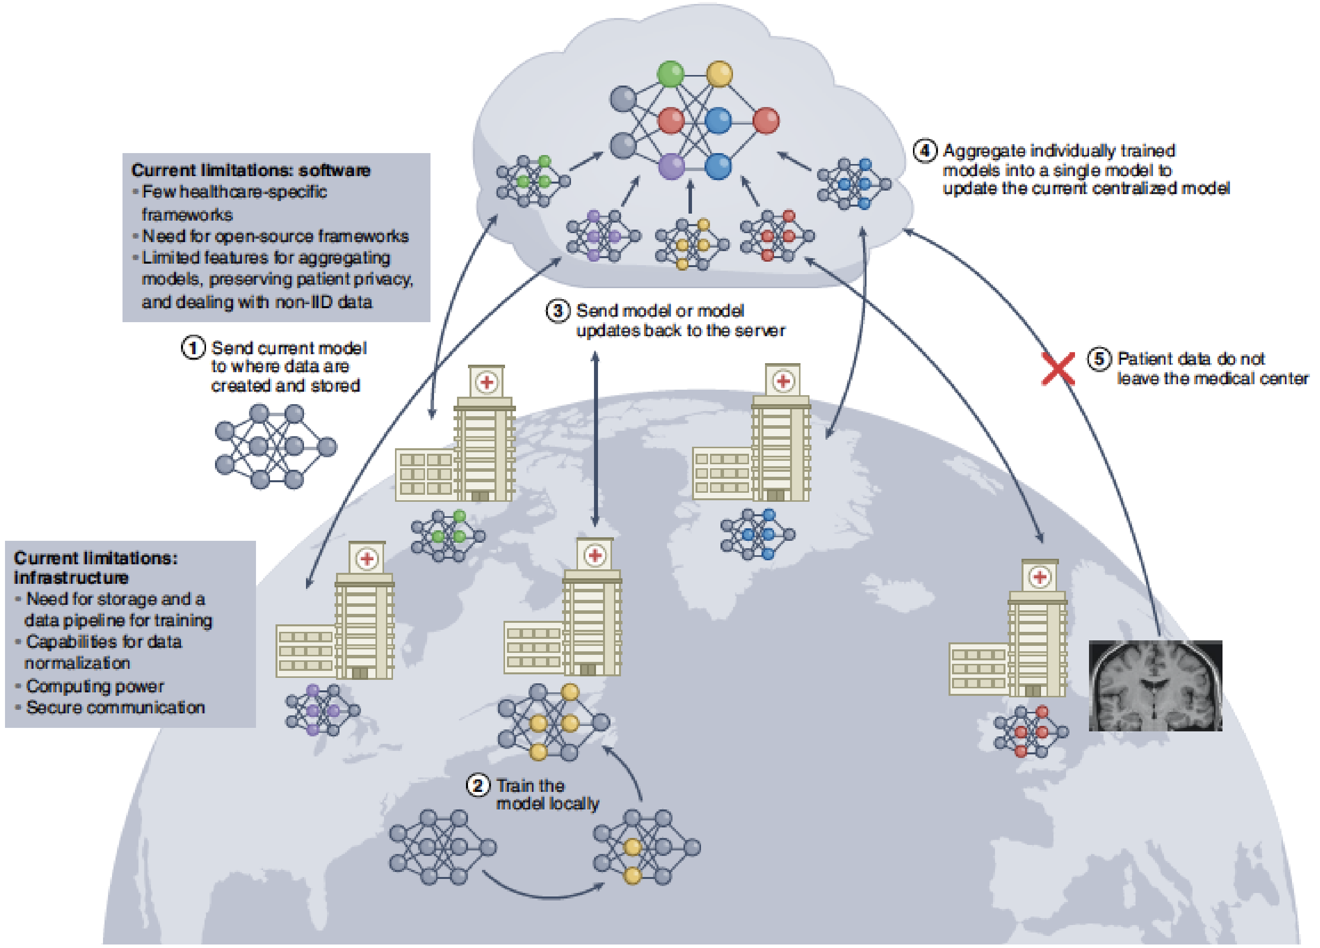

Zhang et al Nat Bio Eng 2022

Specific architectures

For Storage and ML algorithms execution / comparison

- Lesion load estimation is crucial for therapy adaptation
- Manual deliniation too variable and time-consumming
- How to obtain a fair comparison of automatic segmentation?

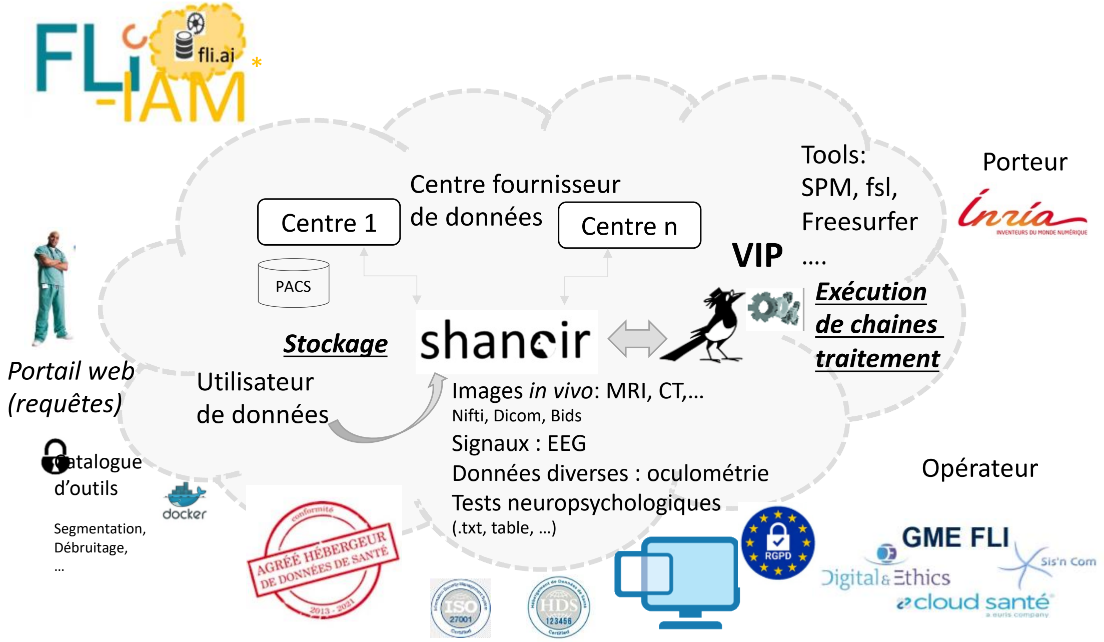

MS-SEG

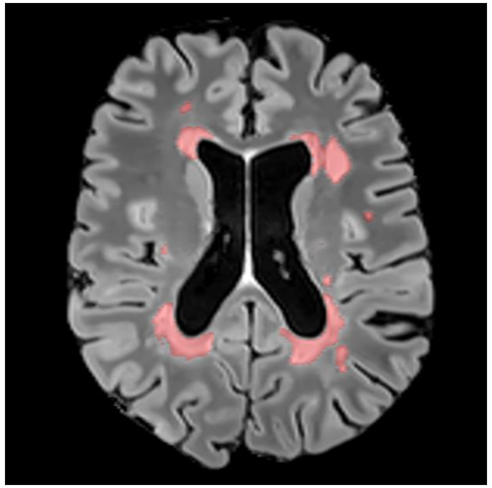

FLAIR from
center 01

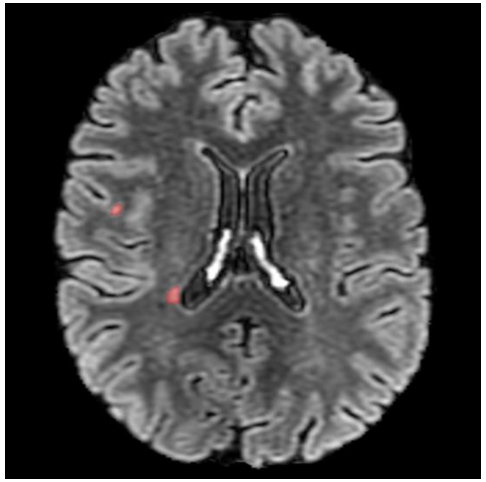

FLAIR from
center 03

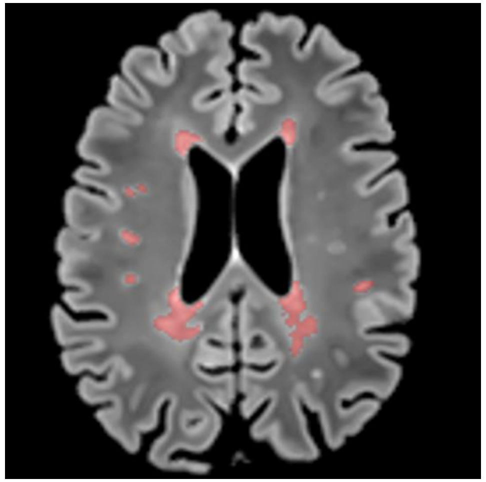

FLAIR from
center 07

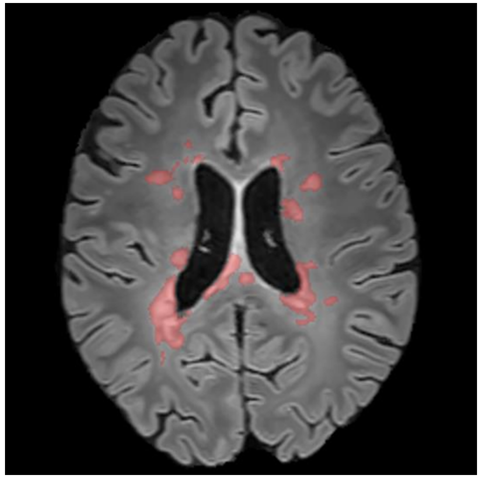

FLAIR from
center 08

Challenge Miccai October, 21, 2016

- 53 Ground truth
- MRI from three different centers: Bordeaux, Lyon, Rennes
- Acquisitions from 4 different scanners
  Siemens Verio 3T (01) – Rennes – 10 cases
  GE Discovery 3T (03) – Bordeaux – 8 cases
  Siemens Aera 1.5T (07) – Lyon – 10 cases
  Philips Ingenia 3T (08) – Lyon – 10 cases

  MRI following the OFSEP protocol
  - 3D T1, FLAIR, axial DP and T2
  - 13 challengers (4 CNN)
  - 7 expert manual delineation per case => Ground Truth

Process …

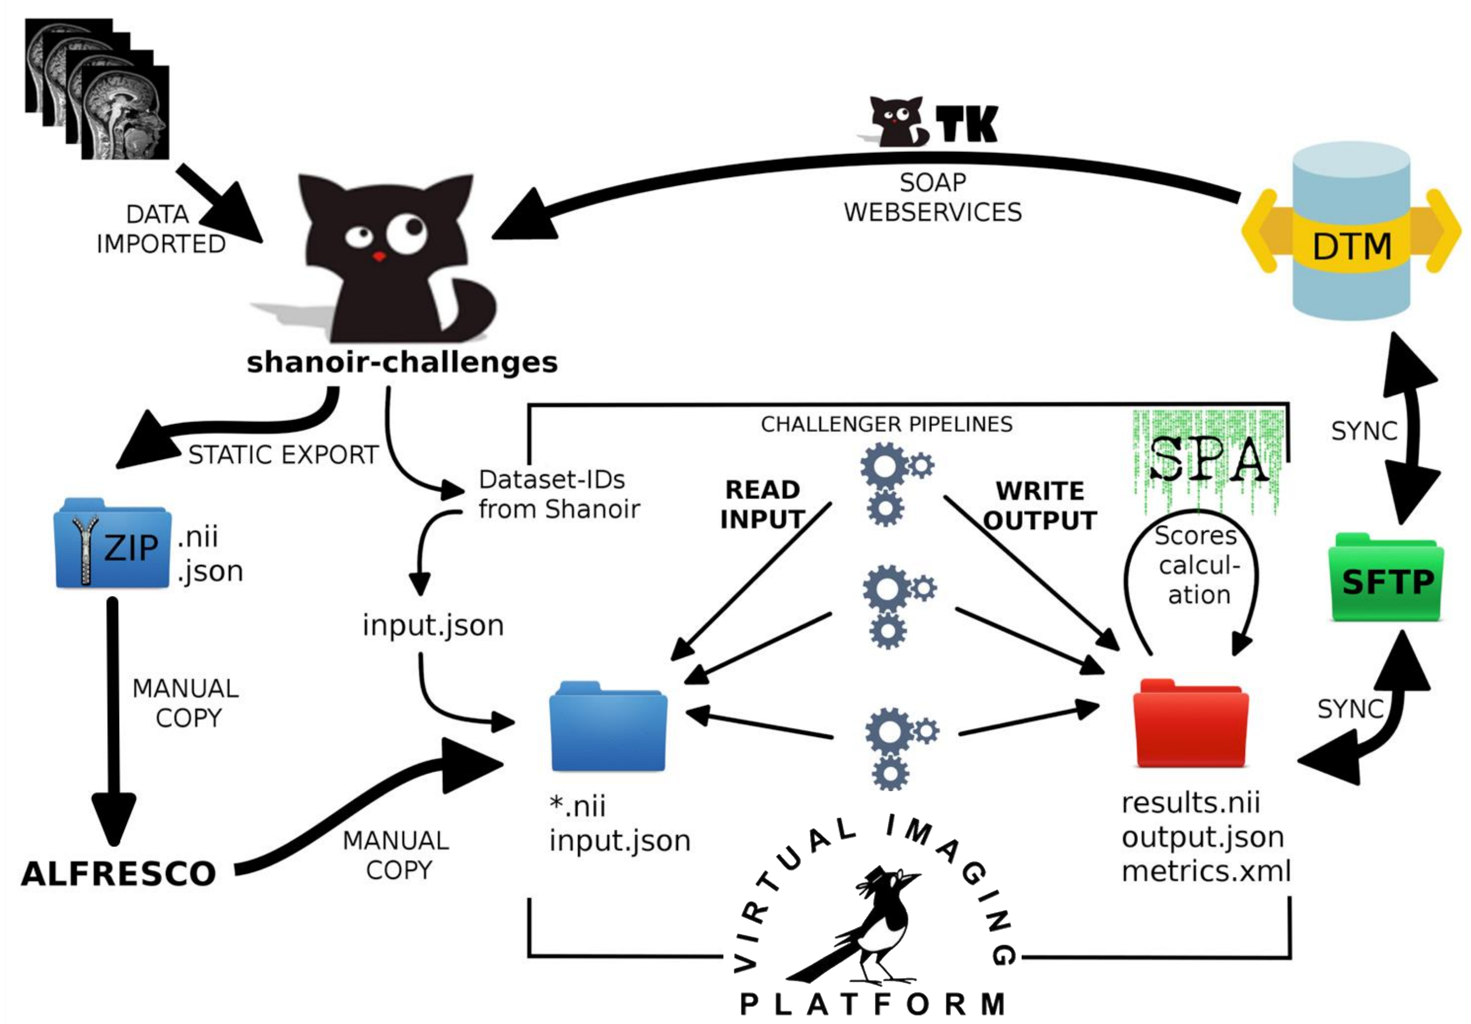
TK
SOAP
WEBSERVICES
DATA
IMPORTED
DTM
shanoir-challenges
CHALLENGER PIPELINES
SYNC
STATIC EXPORT
SPA
Dataset-IDs
from Shanoir
READ
INPUT
WRITE
OUTPUT
ZIP
.nii
.json
Scores
calcul-
ation
SFTP
input.json
MANUAL
COPY
SYNC
*.nii
input.json
results.nii
output.json
metrics.xml
ALFRESCO
MANUAL
COPY
VIRTUAL IMAGING
PLATFORM

Results

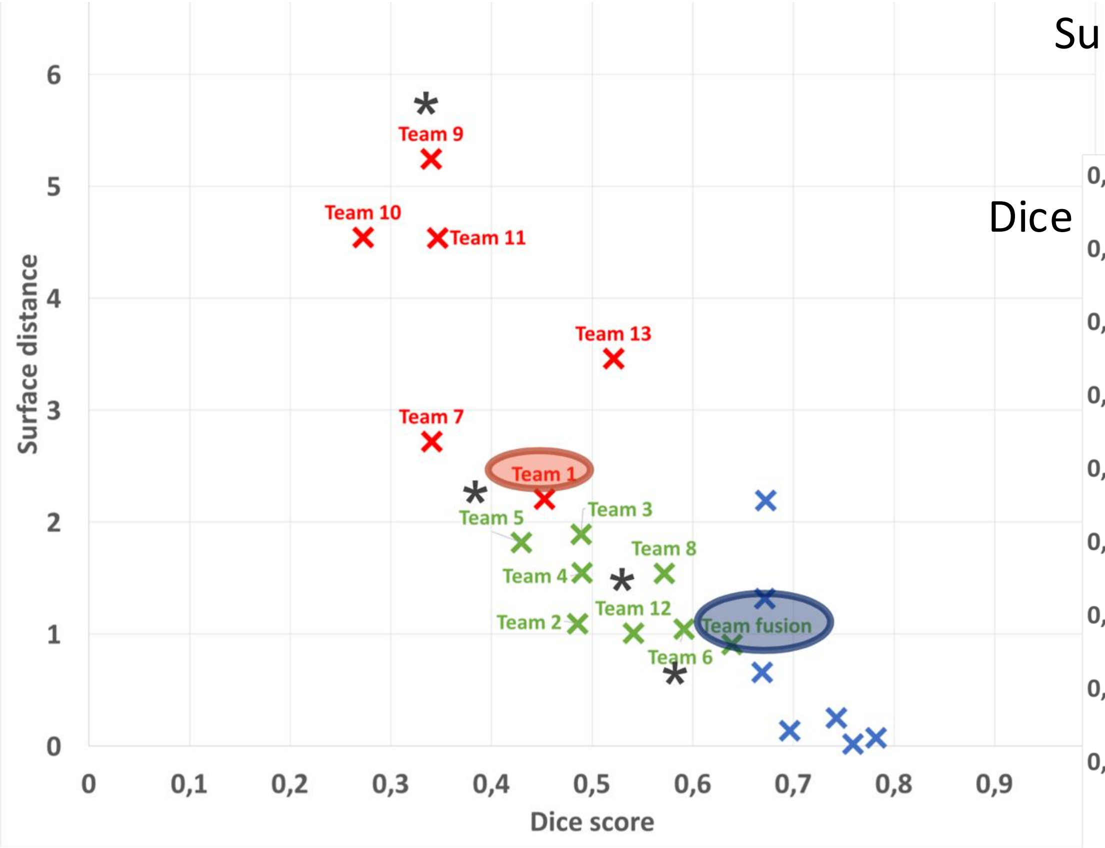

Surface

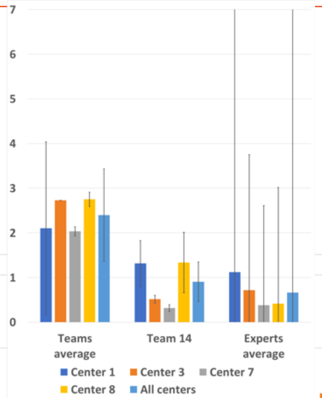

Dice

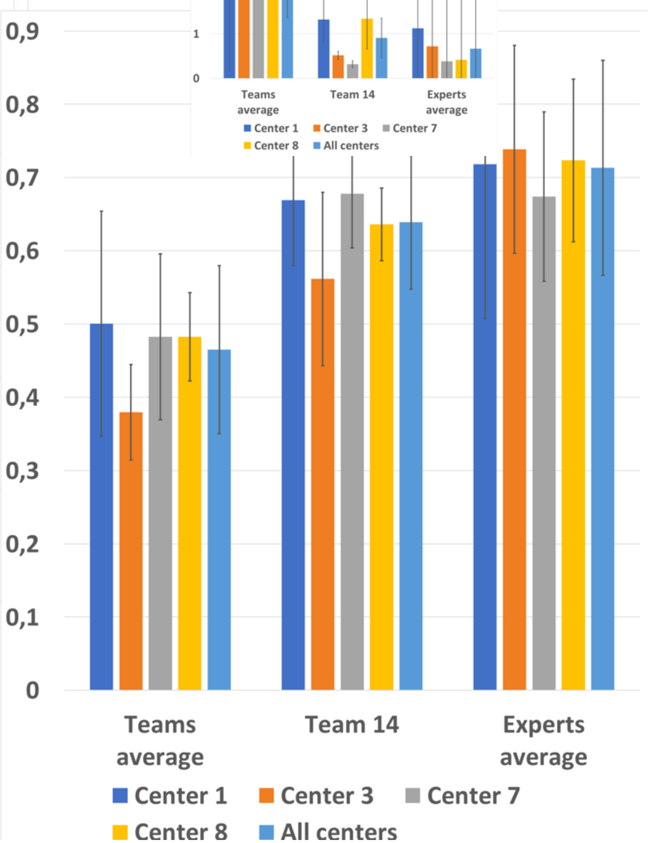

Results

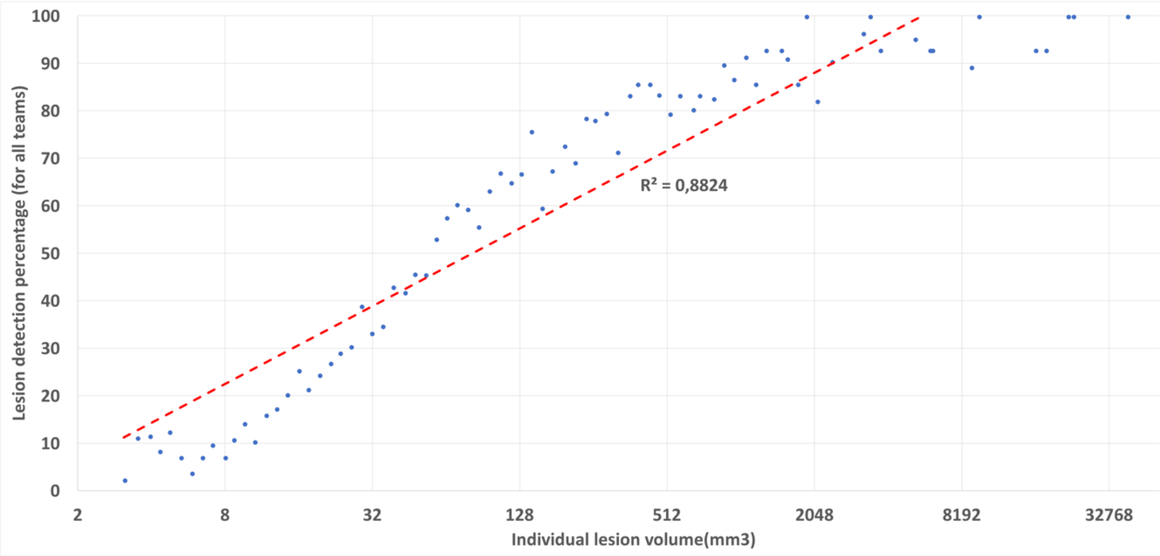

	Lesions volume (cm^3)	Number of lesions
Expert 1	0.052	2
Expert 2	0.090	2
Expert 3	10.887	8
Expert 4	**0**	**0**
Expert 5	0.017	1
Expert 6	**0**	**0**
Expert 7	0.029	2
Team 1	8.252	18
Team 2	**0**	**0**
Team 3	**0**	**0**
Team 4	NA	NA
Team 5	28.436	522
Team 6	0.473	7
Team 7	5.990	168
Team 8	**0**	**0**
Team 9	2.545	33
Team 10	11.085	31
Team 11	3.436	42
Team 12	0.056	1
Team 13	0.074	4

Conclusion

- Results still far from experts
    - Mainly on images with low lesion load
    - All methods sensitive to image quality (center 03)

[Commonwick et al Scient Reports 2018]