A review of unsupervised anomaly detection models for neuroimaging applications

Carole Lartizien
carole.lartizien@creatis.insa-lyon.fr

- AI for medical image analysis
- AI for neuroimaging analysis
- Basics of machine learning for image analysis
- Supervised semantic segmentation : use case #1
- Unsupervised anomaly detection : use cases #2 and #3.
- Conclusion

AI for medical image analysis

Data

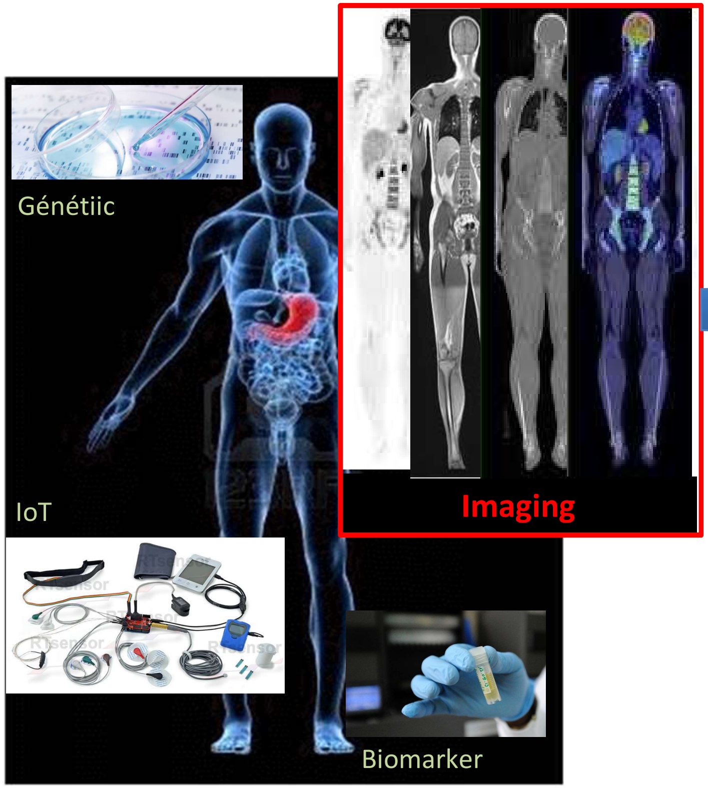

AI – assisted physicians

Precision medicine

- Refine Diagnosis
- Improve therapy planning
- Predict therapy outcome
- Develop preventive medicine

Figure 2: Diagnostic exams and MRI exams per 1,000 population

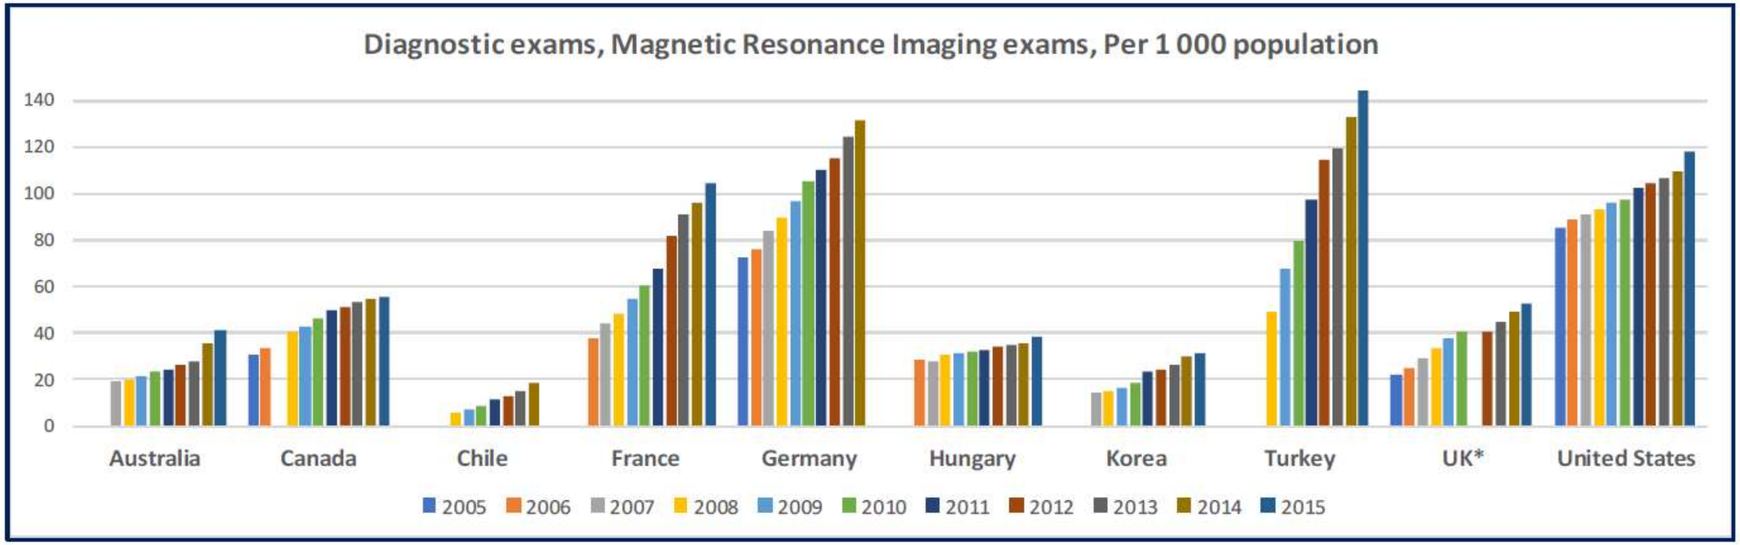

https://emea.gehealthcarepartners.com/images/pdfs/Rapid-Review--Radiology-Workforce-Review-FINAL.pdf

AI in clinical practice

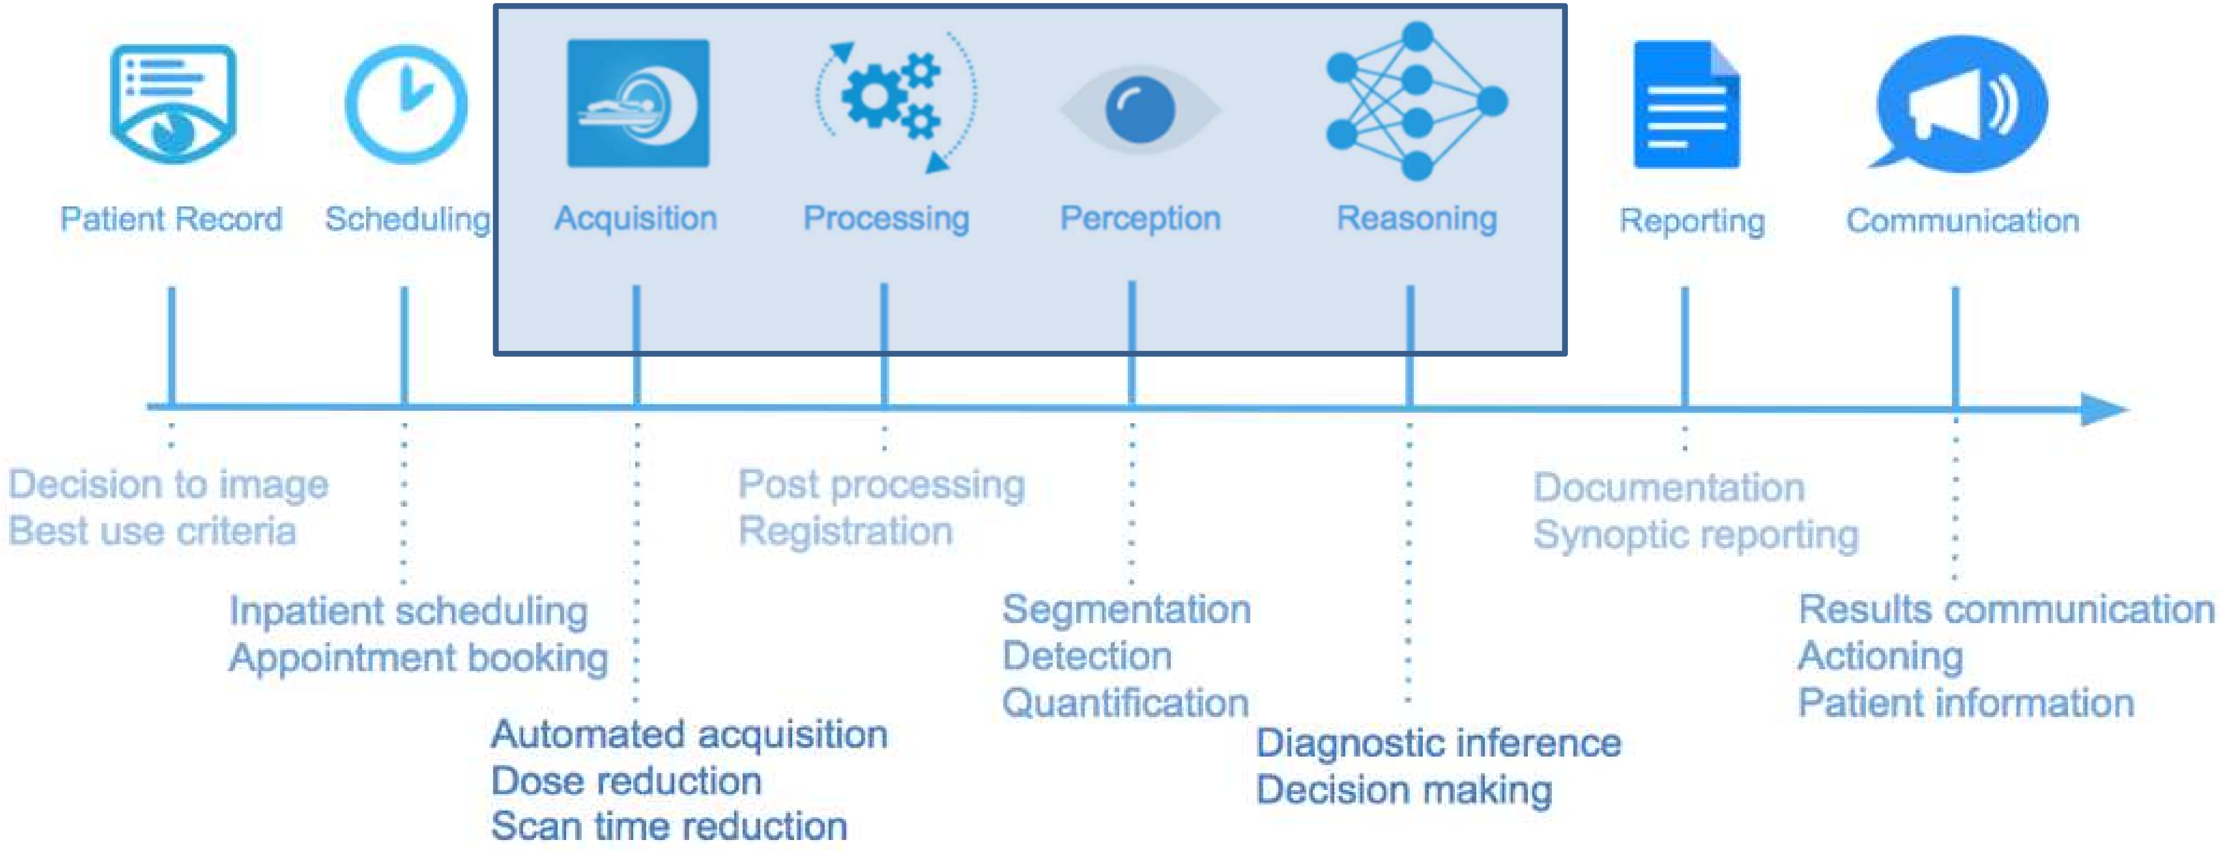

https://towardsdatascience.com/why-ai-will-not-replace-radiologists-c7736f2c7d80

From imaging data to wisdom

DECISION
diagnosis, prognosis

WISDOM — Diagnostic and Prognostic models

KNOWLEDGE — Segmentation, detection, classification

INFORMATION — feature extraction

DATA — Multi-modality imaging

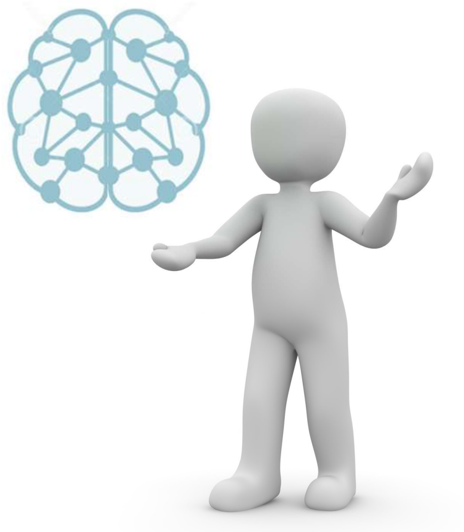

AI ⟷ human expert

AI for neuroimaging analysis

AI for neuroimaging data analysis

Is this patient affected by preliminary symptoms of the Alzheimer disease (AD) ?

Are there **imaging biomarkers** of the pathology ?

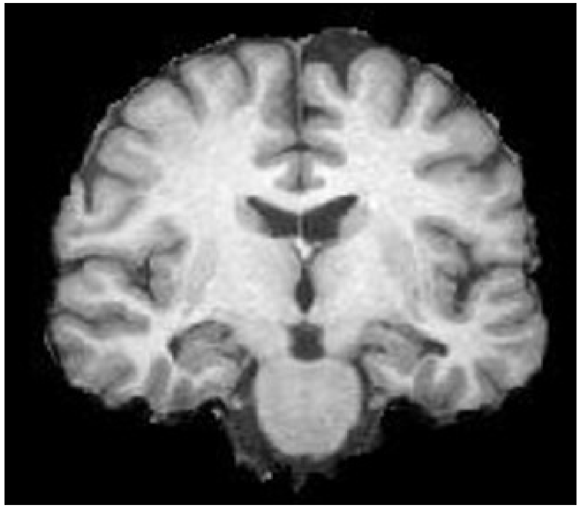

Hippocampus atrophy is one imaging biomarkers of AD

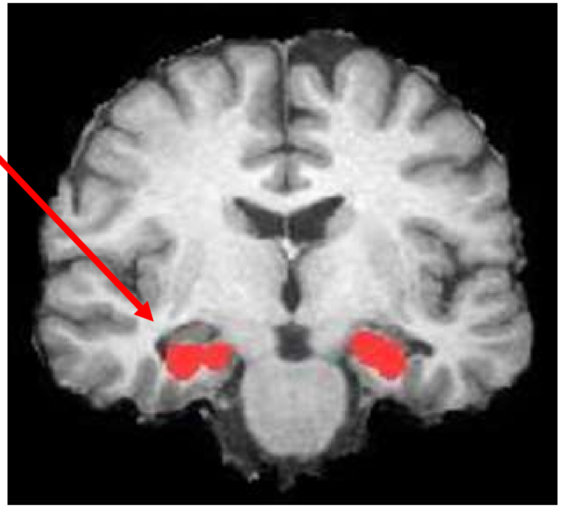

Pathological classes : Normal, Mild cognitive impairement, AD

Prediction of a patient-level malignancy score

Segmentation of anatomical structures of interest

AI for neuroimaging data analysis

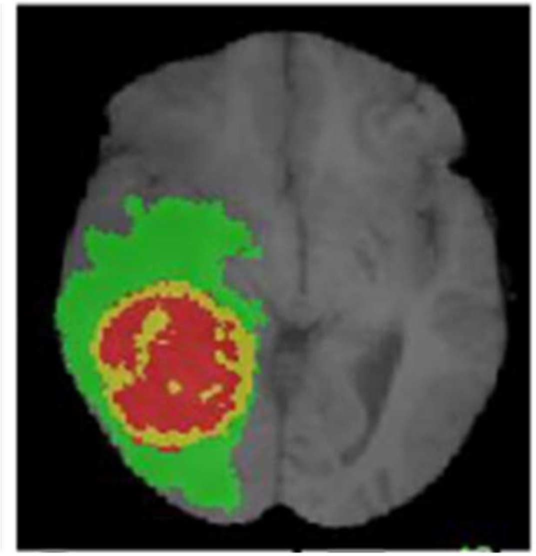

Segmentation of the different components of a brain tumor

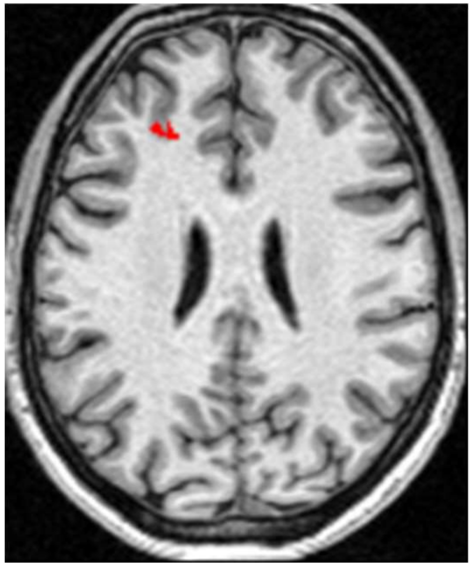

Detection of subtle abnormalities in cortical gyration leading to epileptogenic seizures

Localisation of suspicious area

AI for neuroimaging data analysis

Will this patient who has just had a stroke benefit from a surgical thrombectomy ?

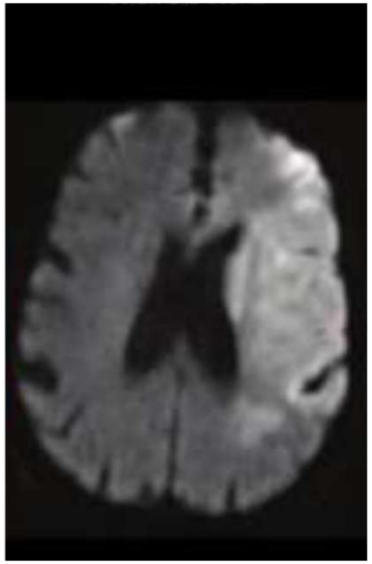

MRI image at ICU

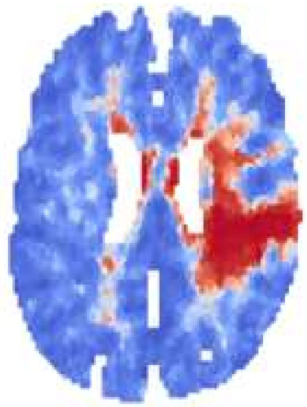

Predicted lesion 6 months after thrombectomy

Prediction of outcome from surgery or therapy

AI for neuroimaging data analysis

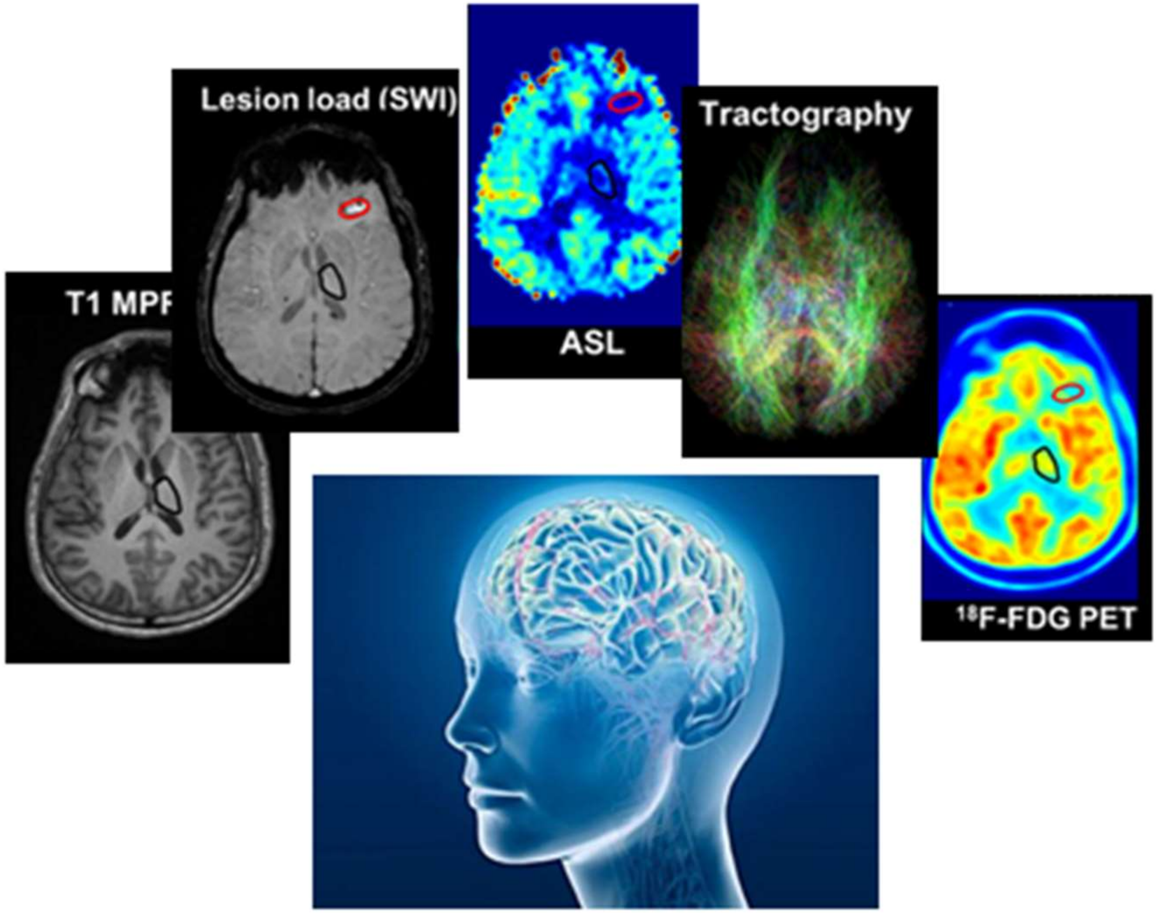

Multimodal heterogenous data analysis….

AI for neuroimaging data analysis

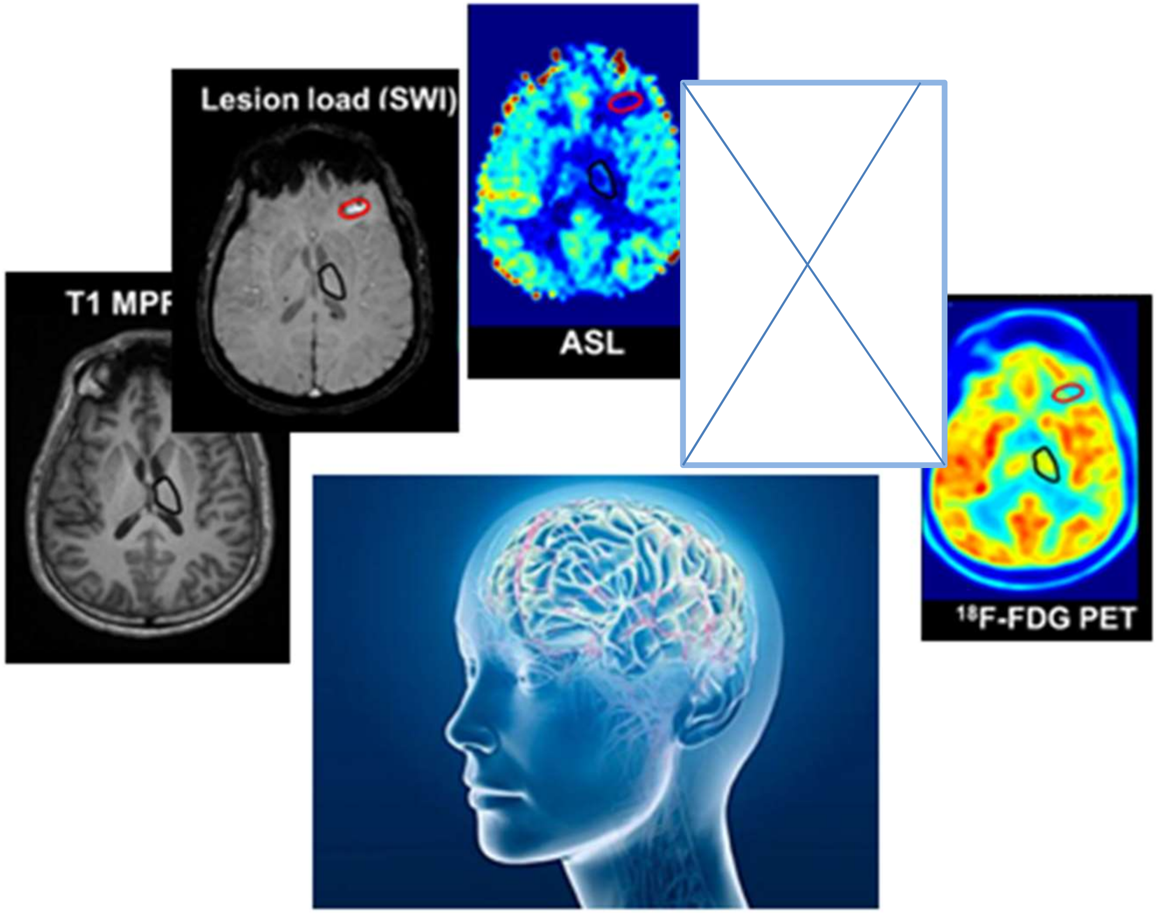

with missing data….

Basics of machine learning for image analysis

(in a few seconds)

Basics of machine learning

1. Define a **task**
2. Formulate this task as a **decision model**
3. Learn the hyperparameters of the decision model based on samples **data** and a **performance metric**
4. Infer decision from this model on new samples

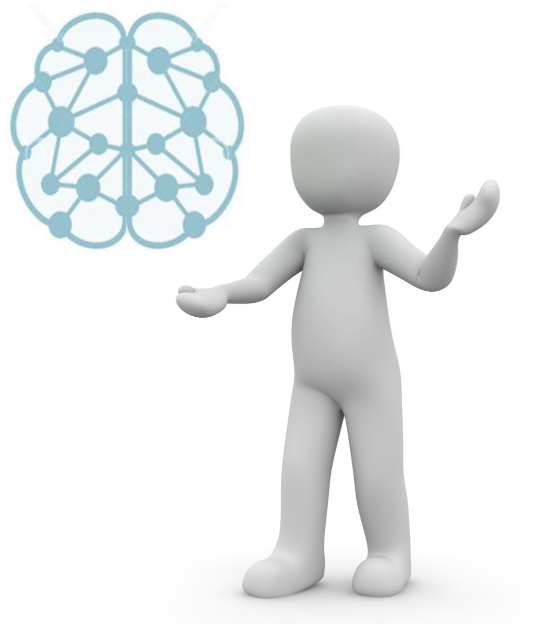

Design your AI algorithm

1. Task definition

Detect lesions on brain T1 MRI

2. Problem formulation as a decision task

Decide whether each voxel is a ‘lesion’ or ‘normal tissue’

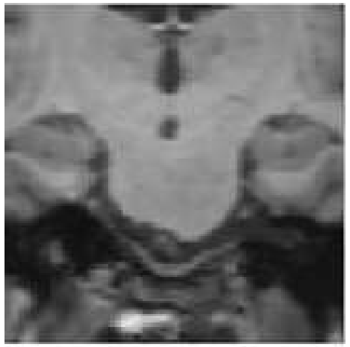

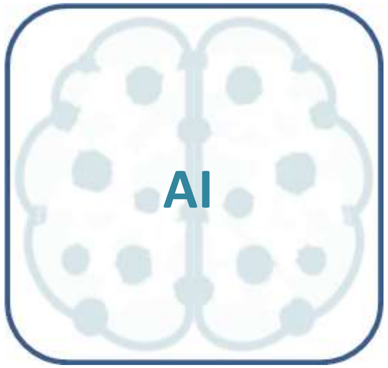

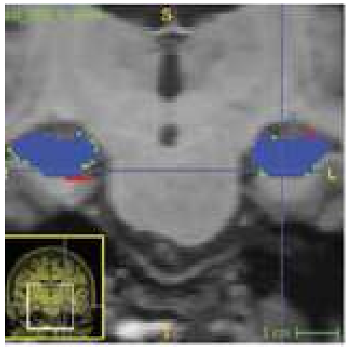

- **Binary classification problem**
- **At the voxel level**

Design your AI algorithm

3. Characteristic of the database

How many samples?

Are they annotated?

…

4. Expert knowledge

Insert manually engineered features? Priors on the expected output etc…

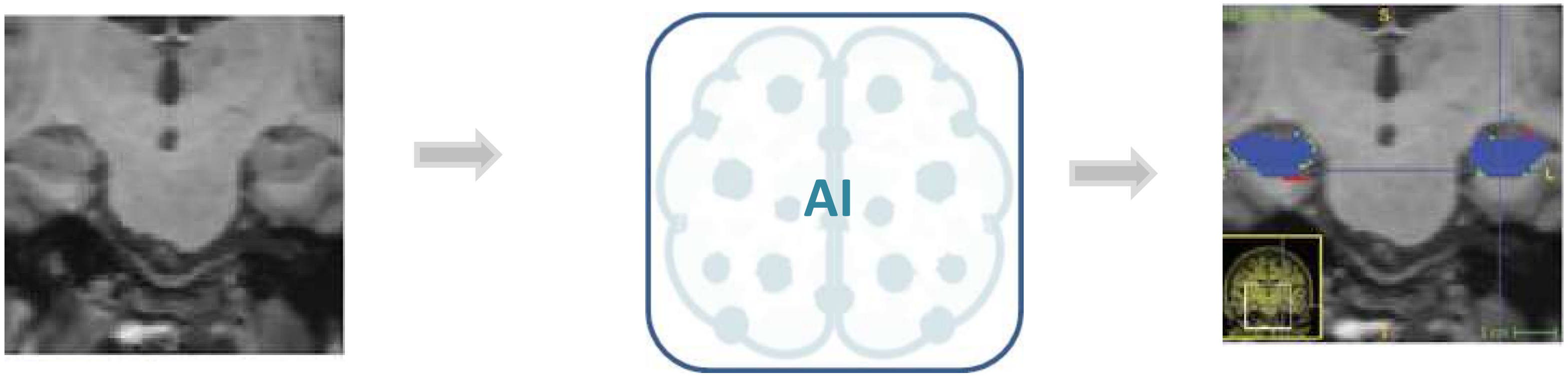

- Binary classification problem
- At the voxel level
- Supervised learning
- Clinically driven feature maps

Design your AI algorithm

3. Learn the decision mode based on training samples and performance metric

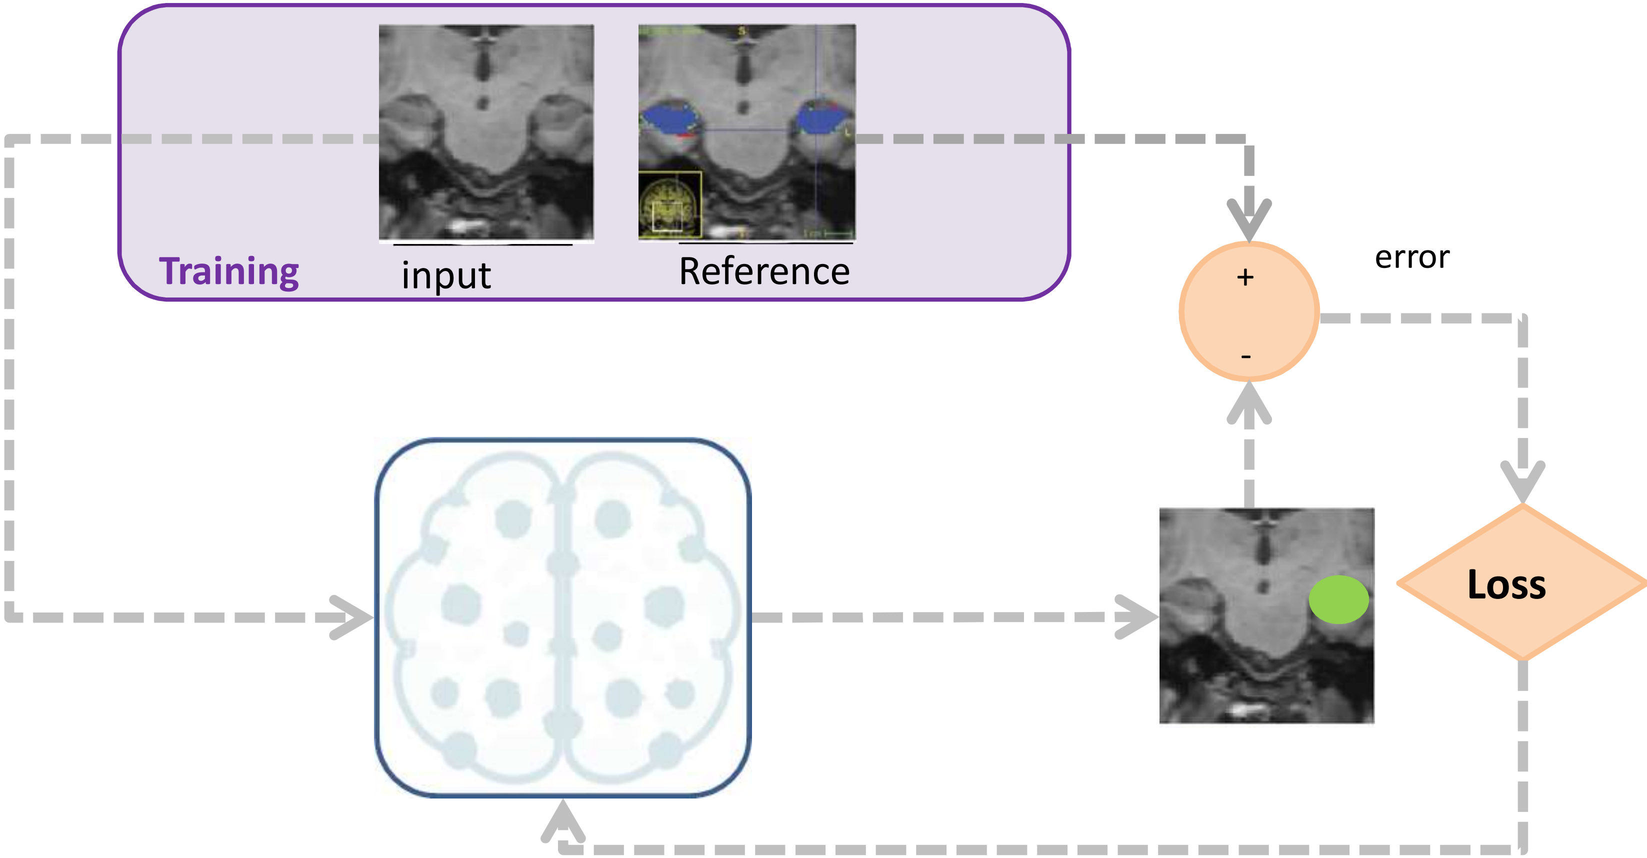

Design your AI algorithm

4. Infer decision on new samples

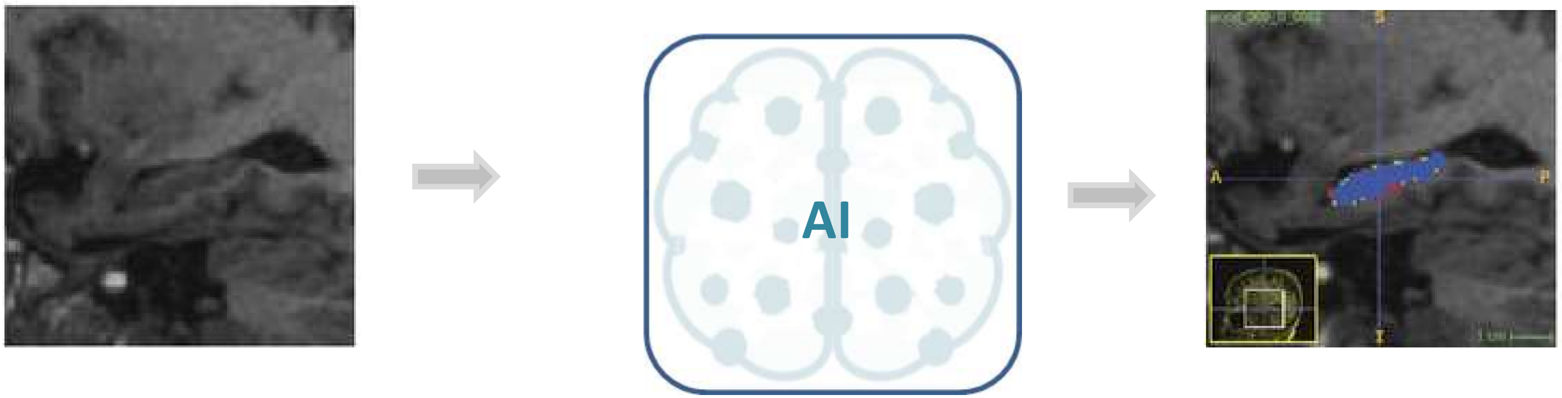

From standard machine learning…

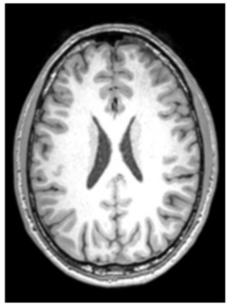

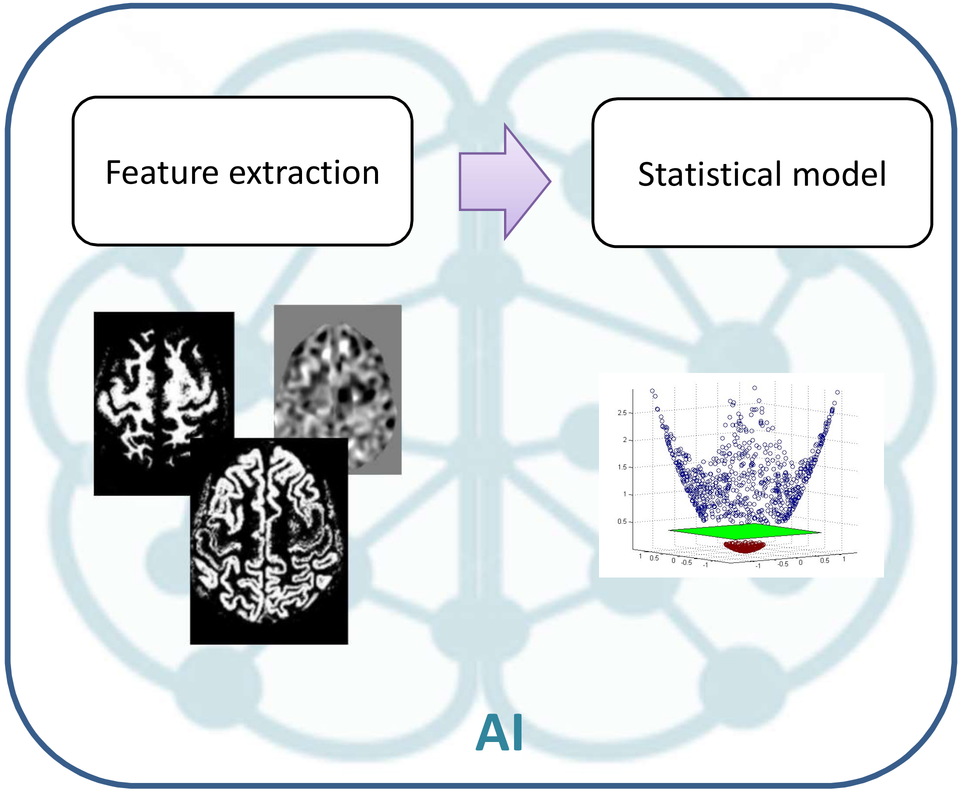

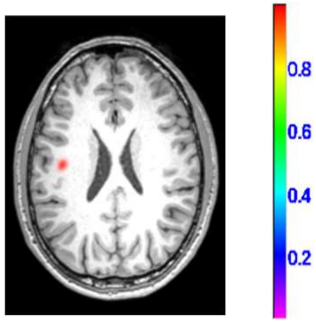

…to deep learning

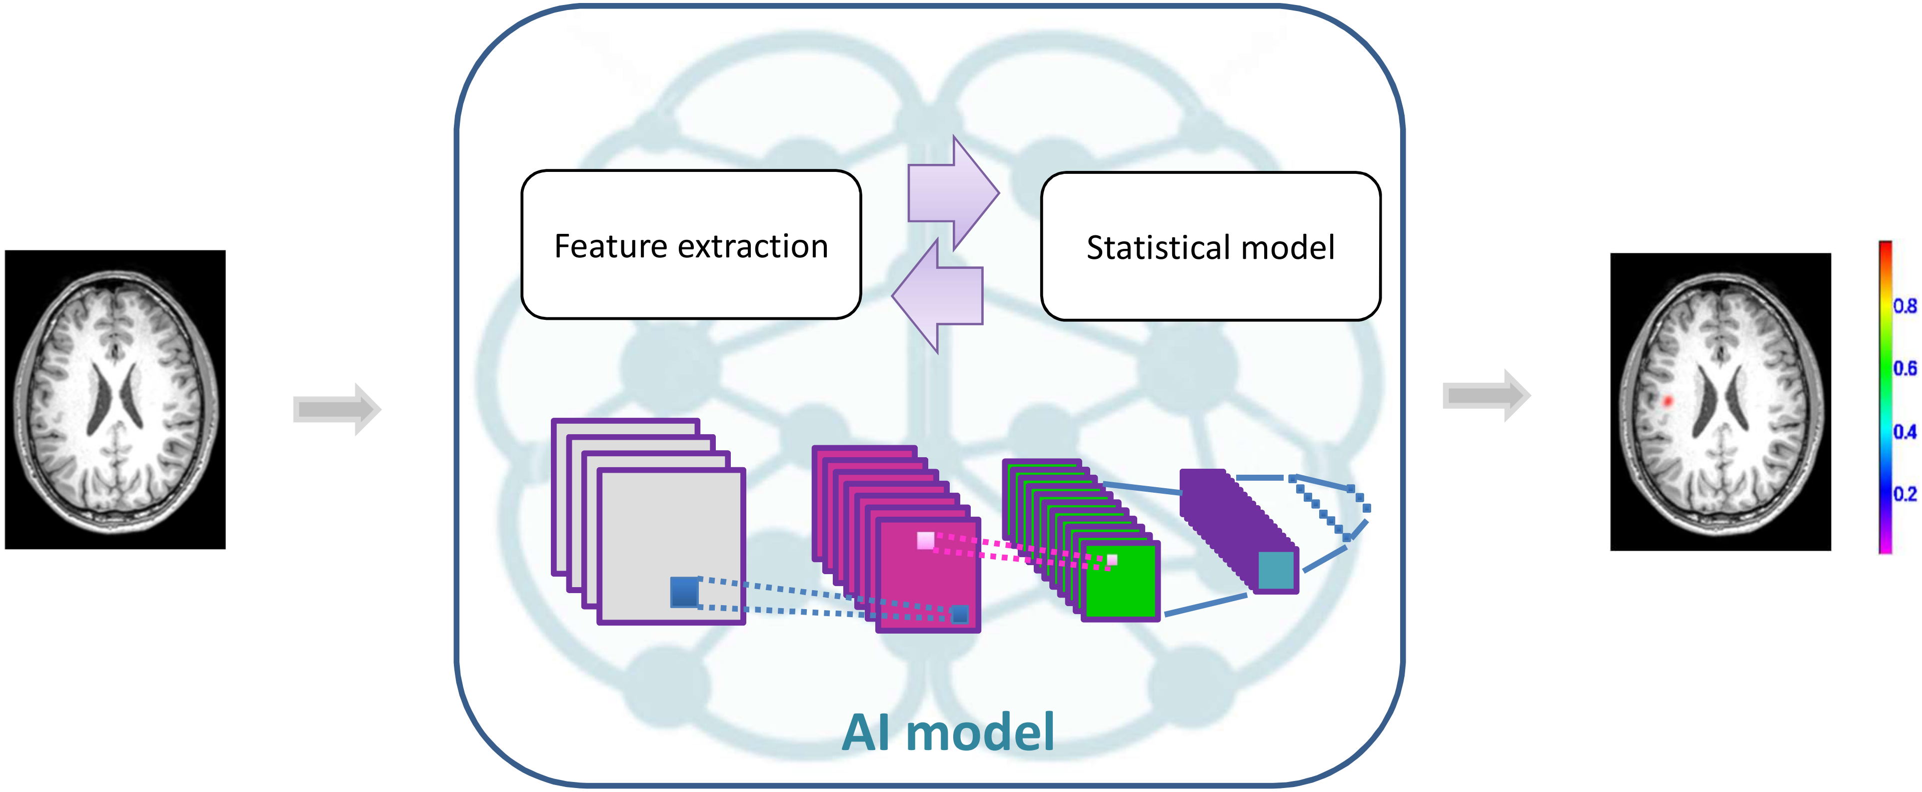

USE CASE # 1

Automatic multi-class segmentation of brain tumors

Supervised learning

Automatic muti-class segmentation of brain tumors

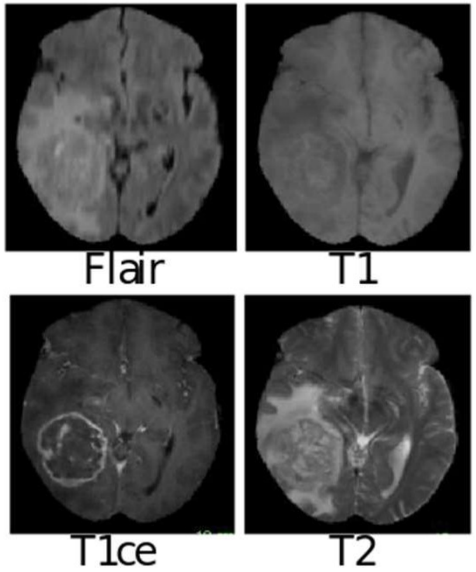

Multiparametric Magnetic Resonance Imaging (MRI)

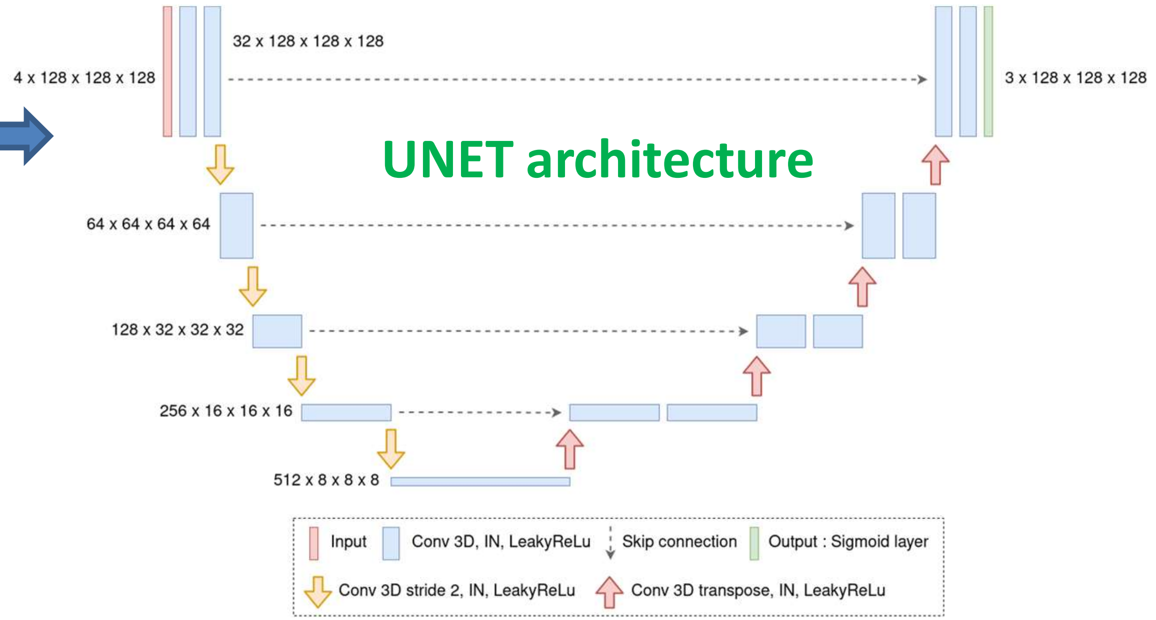

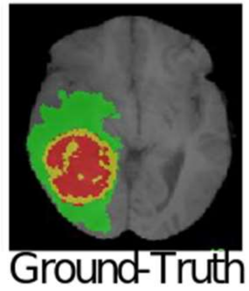

Whole Tumor
Tumor core
Enhancing Tumor

→ Segmentation of the brain tumor is valuable at each step of the patient care : from diagnosis to prognosis, treatment planning and follow-up up to outcome prediction

Automatic muti-class segmentation of brain tumors

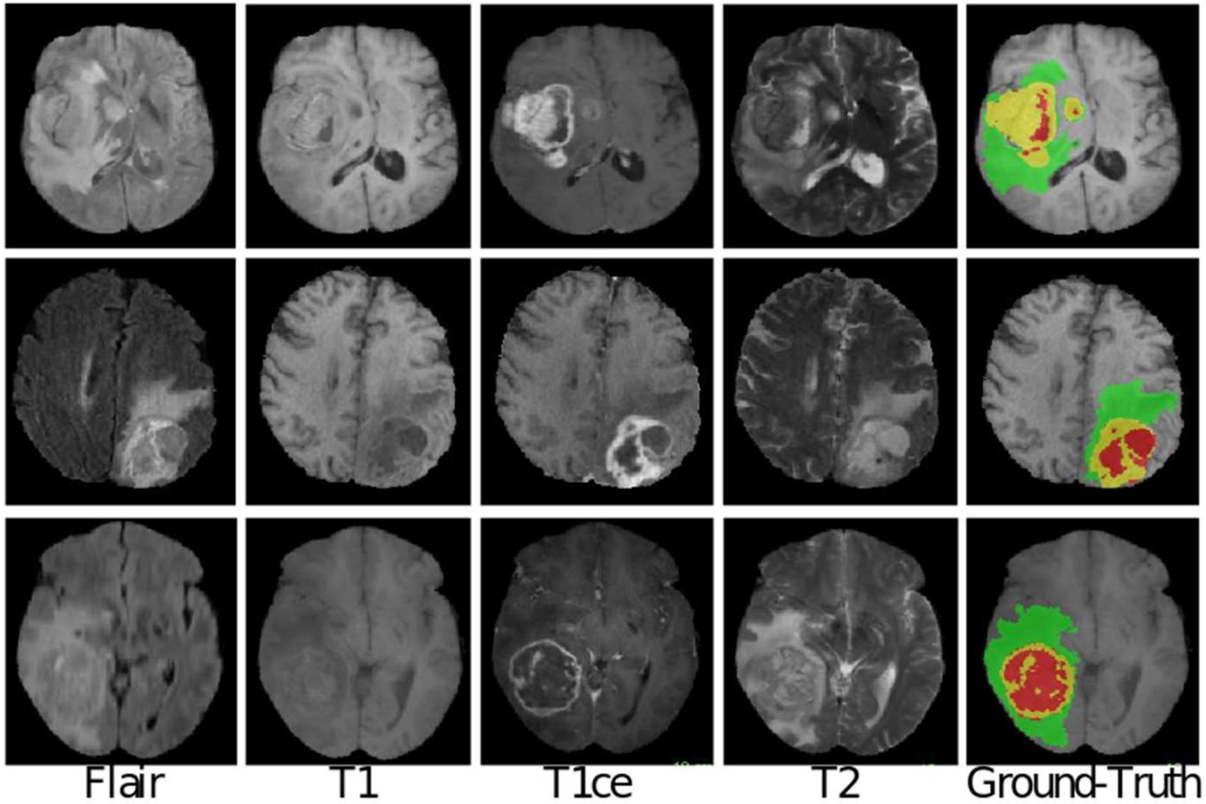

Manual annotation by expert clinicians on 3D image of dimension : ~240 x 240 x 155 with $1mm^3$ isotropic voxels

Manual annotations are time consuming but « feasible »

→ International initiatives to gather large datasets

Automatic muti-class segmentation of brain tumors

Federated Brain Tumor Segmentation (FeTS2022) Challenge Dataset

- **Dataset size** : 1251 patients,
- **MRI scan**
  - **voxel size** : 240 x 240 x 155 of $1mm^3$
  - 4 modalities : Flair, T1, T1ce, T2,
- **Partitioning :** original**,** institution-wise
  - 23 institutions, heterogeneous distributions.
- **Manual annotations :** 3D multi labels
  - Whole Tumor, Tumor core, Enhancing Tumor

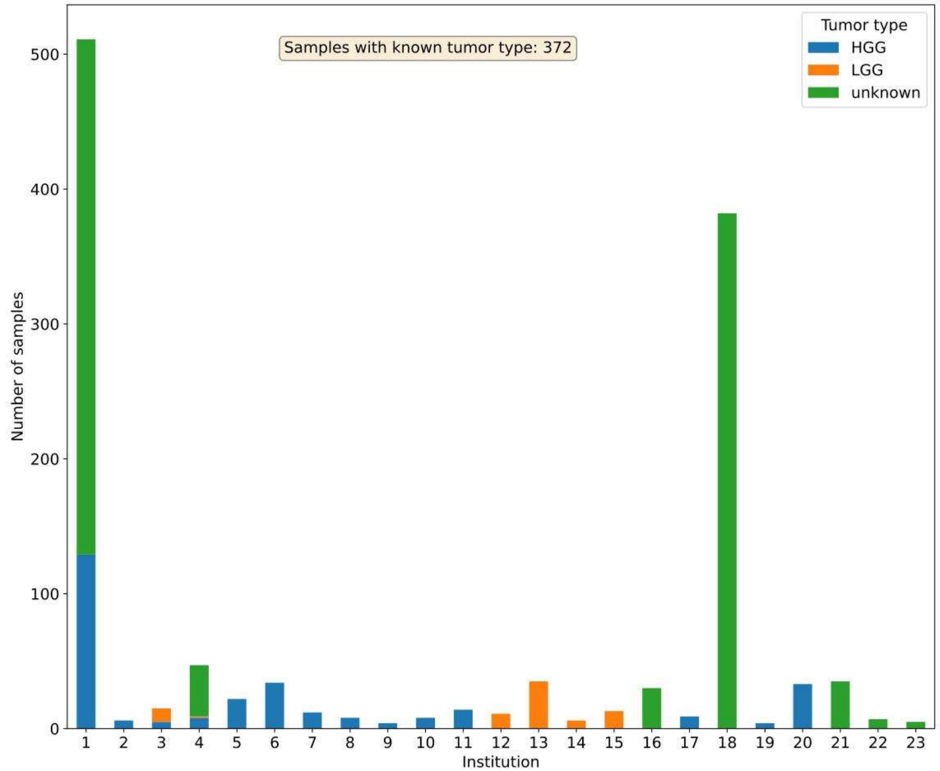

M. Islam et al. « Brain Tumor Segmentation and Survival Prediction using 3D Attention UNet » Apr. 2021

Performance generalization and domain shift

- The **generalization** of a model measures its ability to make good predictions on new unseen data
- These new data come potentially from different "populations" (domain) than those to train the model → **domain shift**

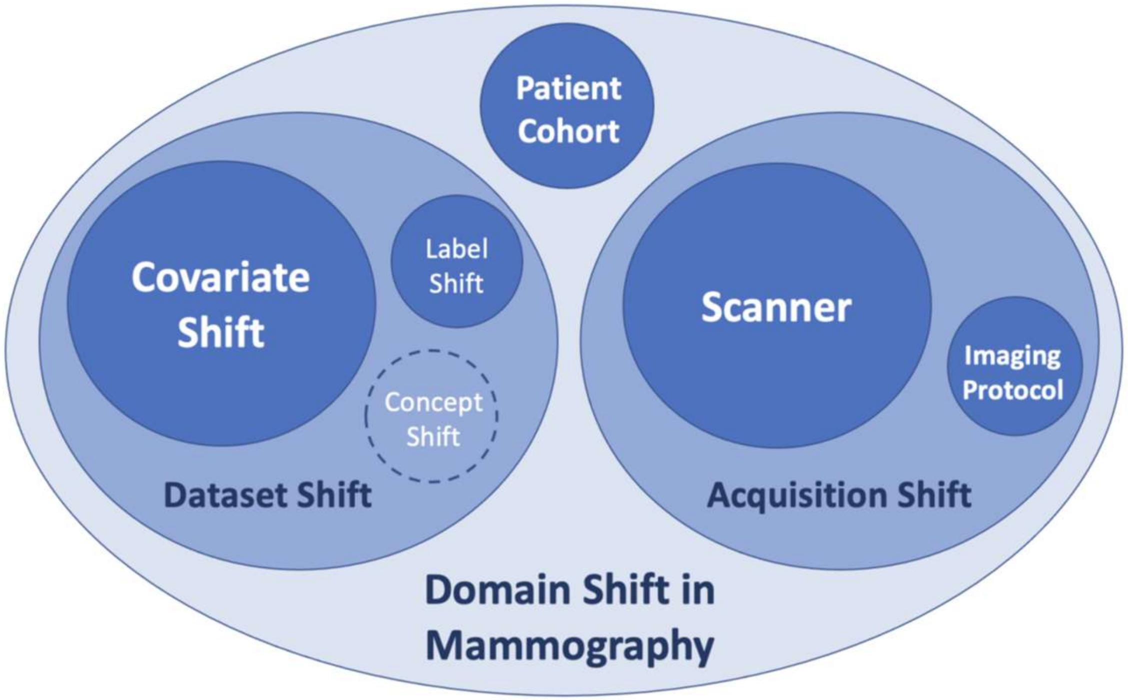

[Garrucho et al, arxiv 2022]

Performance generalization and domain shift

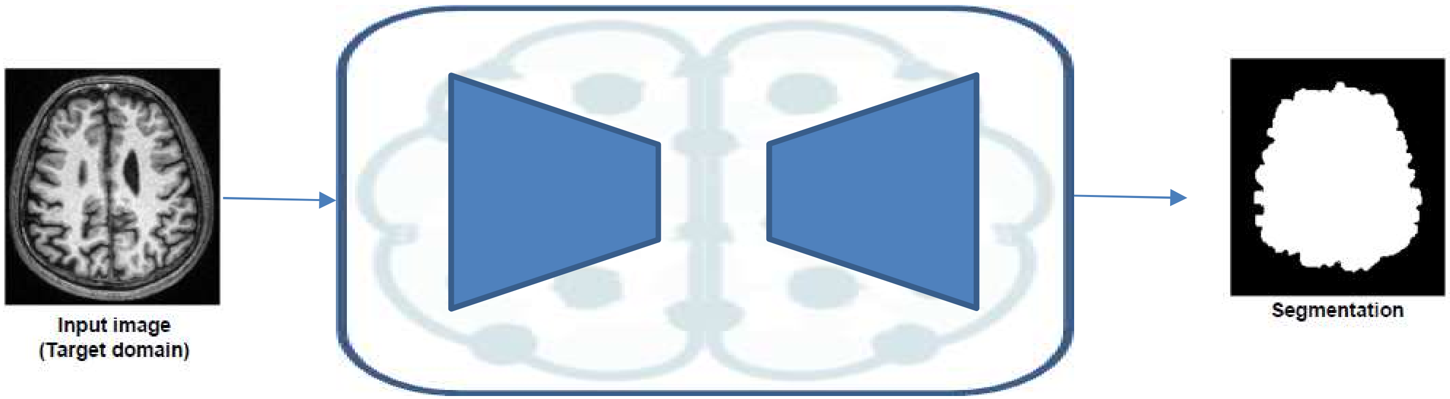

Task : brain Segmentation (skull stripping)

Database : CC359* : 359 T1 MRI of healthy subjects acquired on 6 different scanners

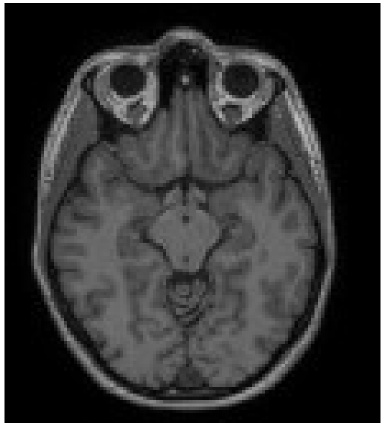

Philips 1.5 T

GE 1.5 T

Siemens 3T

*https://www.ccdataset.com/download

[Zakazov, MICCAI21]

Performance generalization and domain shift

	Sm, 1.5T	Sm, 3T	GE, 1.5T	GE, 3T	Ph, 1.5T	Ph, 3T
Sm, 1.5T	**.86 (.04)**	.61 (.14)	.70 (.08)	.66 (.13)	.69 (.10)	.78 (.06)
Sm, 3T	.70 (.07)	**.87 (.04)**	.63 (.08)	.73 (.07)	.65 (.08)	.72 (.06)
GE, 1.5T	*.36 (.12)*	*.11 (.10)*	**.84 (.06)**	*.36 (.14)*	.62 (.16)	.55 (.11)
GE, 3T	.74 (.07)	.66 (.15)	.58 (.11)	**.89 (.03)**	.59 (.10)	.74 (.05)
Ph, 1.5T	.58 (.07)	.40 (.11)	**.81 (.05)**	.52 (.07)	**.88 (.04)**	.74 (.09)
Ph, 3T	.55 (.13)	*.37 (.15)*	.63 (.11)	.36 (.11)	.53 (.12)	**.86 (.05)**

Table 1: columns - source domains for training, rows - target domains for testing. Metric: surface dice, tolerance=1.

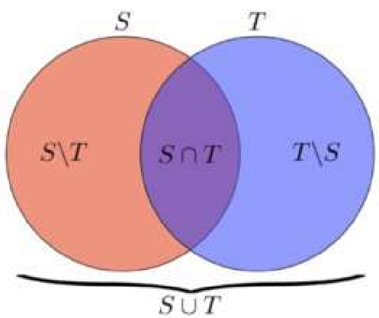

$$Dice(S,T) = \frac{2|S \cap T|}{|S| + |T|}$$

[Zakazov, MICCAI21]

Performance generalization and domain shift

Data harmonization

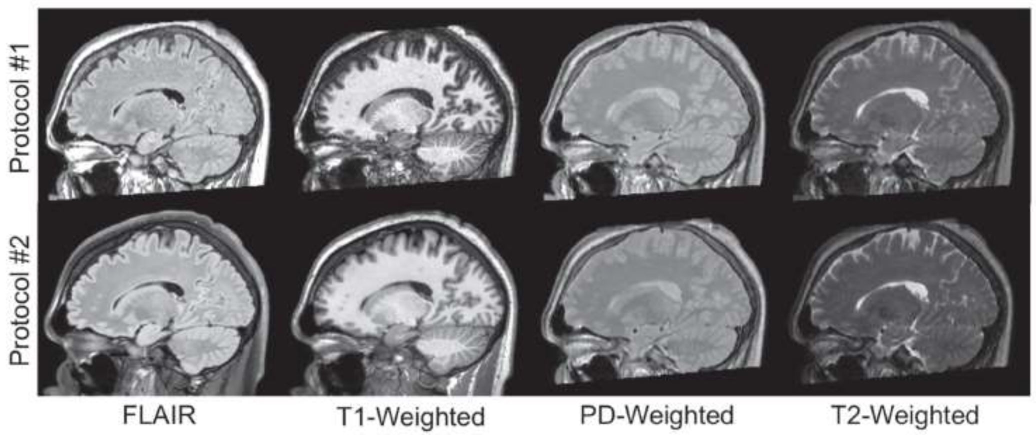

Fig. 1. Preprocessed images from one subject from the overlap cohort depicting the four input (Protocol #1) and four target (Protocol #2) contrasts.

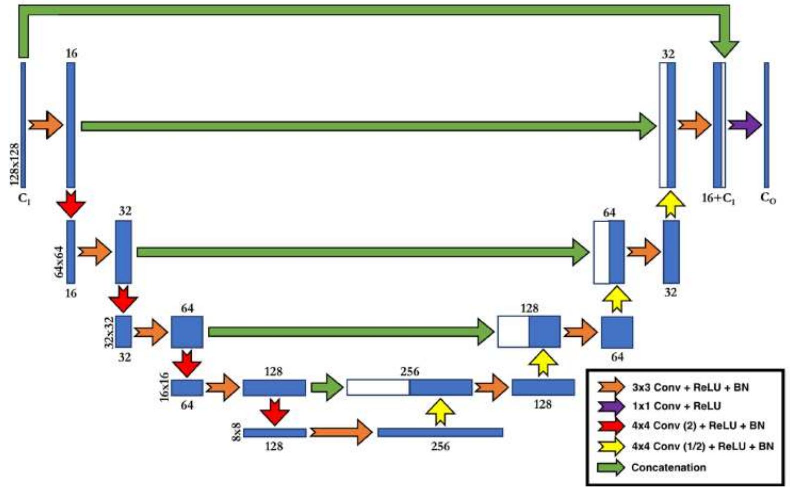

DeepHarmony: A deep learning approach to contrast harmonization across scanner changes

[Dewey, Mag Res Imag 19]

- UNET
- 1 UNET per view (coronal, axial, sagittal)
- 2 versions :
  - O2O one-to-one : 1 input modality → 1 output modality
  - MO2 many-to-one : 4 input modalities → 1 output modalities

Performance generalization and domain shift

Data harmonization

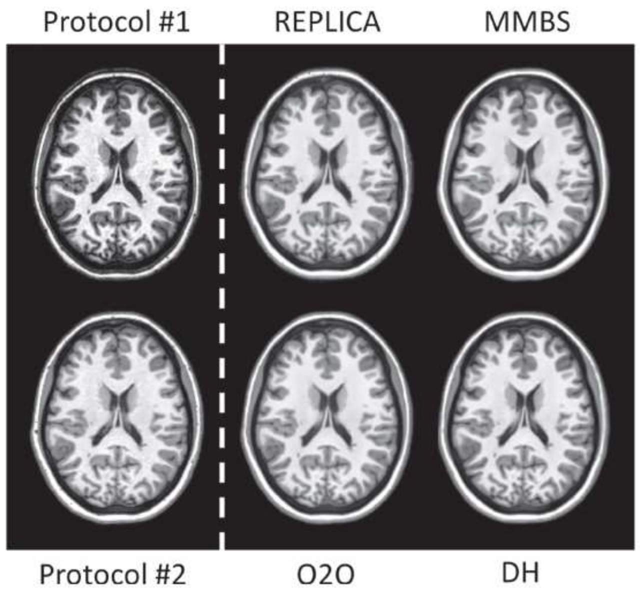

Fig. 4. Harmonized Protocol #1 T1-weighted images using REPLICA, MMBS, O2O, and DeepHarmony (DH). For comparison, the input contrast (Protocol #1) and the target contrast (Protocol #2) are displayed on the left side of the white, dashed line.

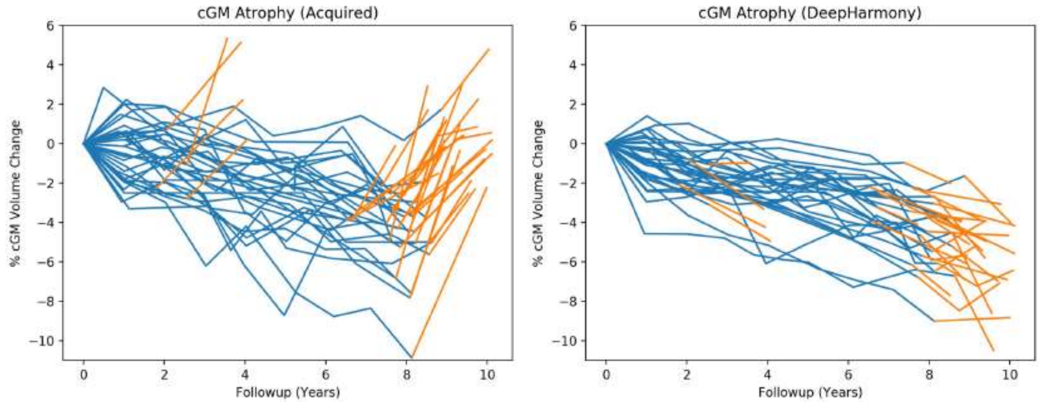

Fig. 8. Longitudinal trajectories for cortical grey matter (in % from baseline). Protocol #1 is shown in blue and Protocol #2 is shown in orange. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

[Dewey, Mag Res Imag 19]

Federated learning

A privacy-preserving decentralized training paradigm

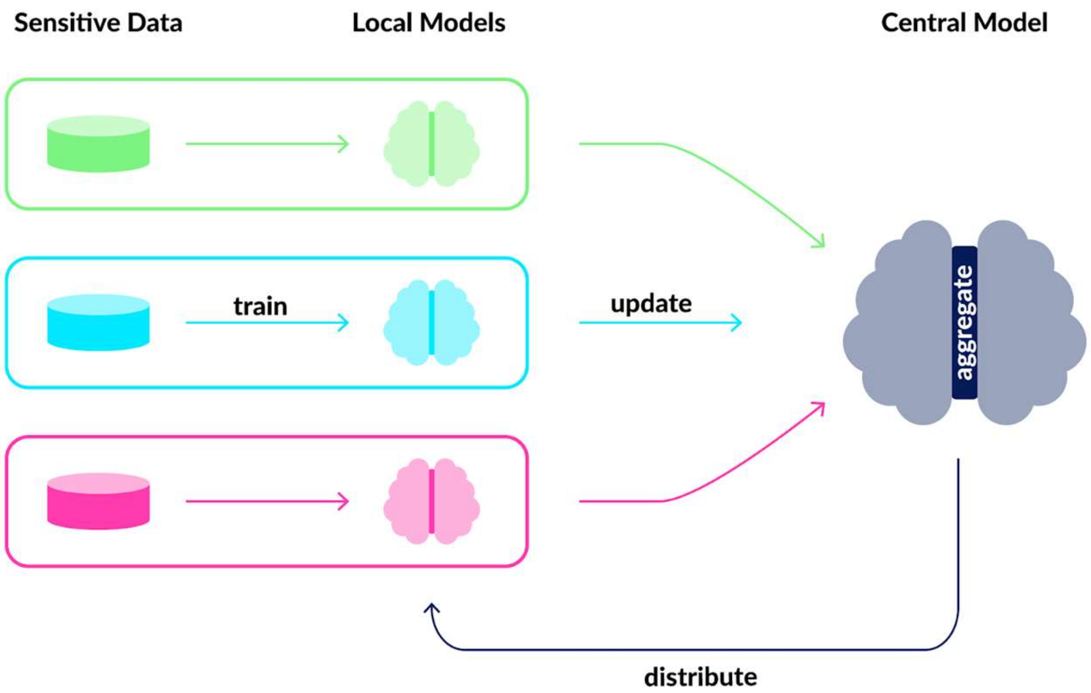

- **First algorithm** : Federated Averaging (FedAvg, **2016**).
- A growing interest in the medical community
- Many open questions regarding FL in the context of heterogeneous data : fairness, personalization ,....

Federated learning

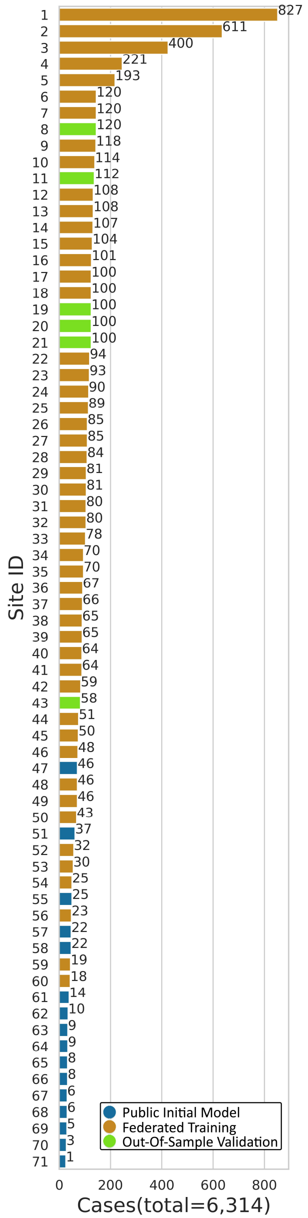

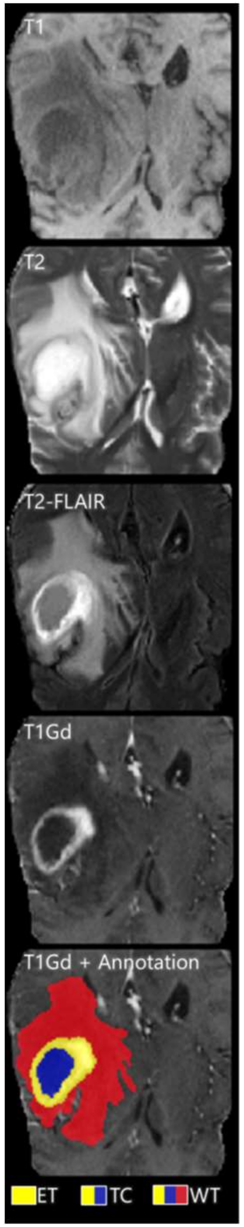

Federated Brain Tumor Segmentation (FeTS)

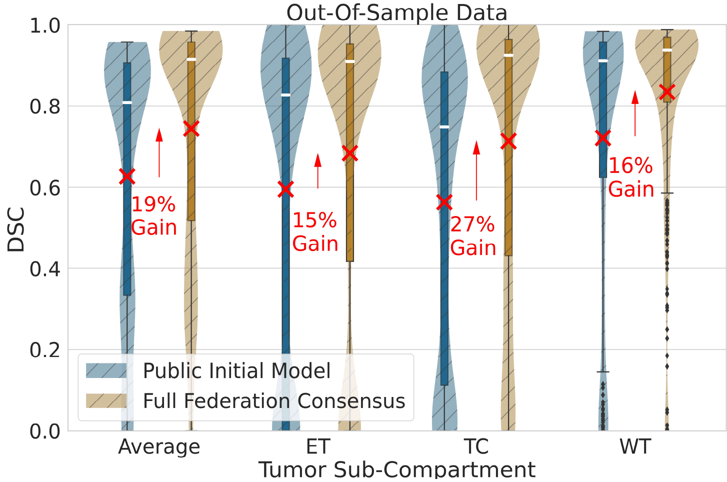

- **public initial model** : trained on 231 cases from 16 sites
- **Out-of-sample site** : did not participate in model training

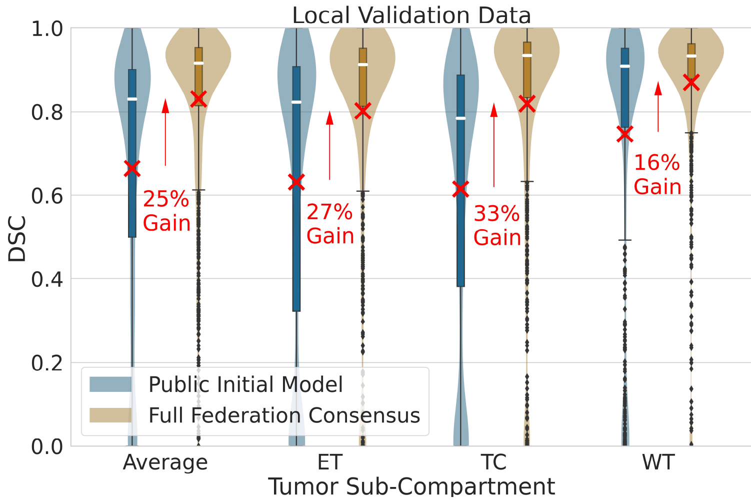

[Pati et al arxiv 22]

USE CASE # 2

Detection and localisation of subtle brain anomalies

Unsupervised anomaly detection

Learning with no annotation

Annotations are time consuming and sometimes hard/impossible to generate

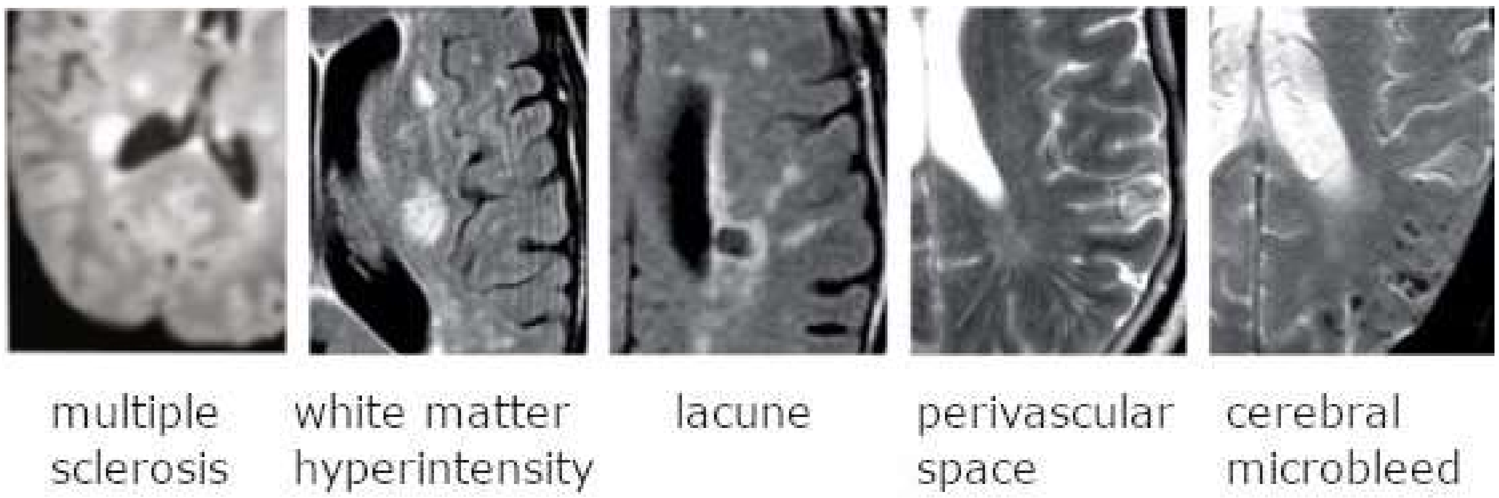

Exemple neuropathologies with subtle lesions

Unsupervised anomaly detection (UAD)

Illustration from Wardlaw et al. , 2012

Problem formulation: objectives & challenges

- Multi-modality data (MRI, PET)
- High dimensional data
- No pathological examples
- Noisy images

- Labelled cluster map
- Probabilistic outputs

Problem formulation

- Formulate the problem as an anomaly detection problem
- Model the distribution of the normal class
- Detect outliers from the normative distribution

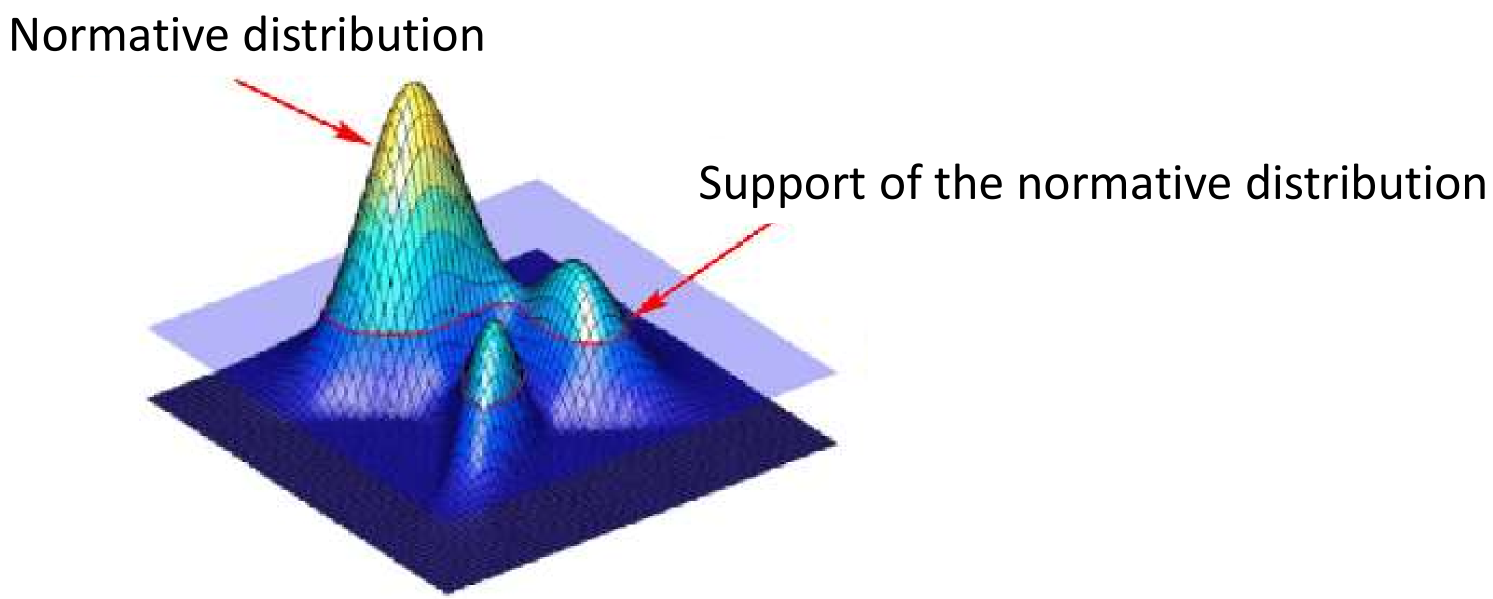

A Unifying Review of Deep and Shallow Anomaly Detection

Lukas Ruff, Jacob R. Kauffmann, Robert A. Vandermeulen, Grégoire Montavon, Wojciech Samek, *Member, IEEE*,
Marius Kloft*, *Senior Member, IEEE*, Thomas G. Dietterich*, *Member, IEEE*,
Klaus-Robert Müller*, *Member, IEEE*.

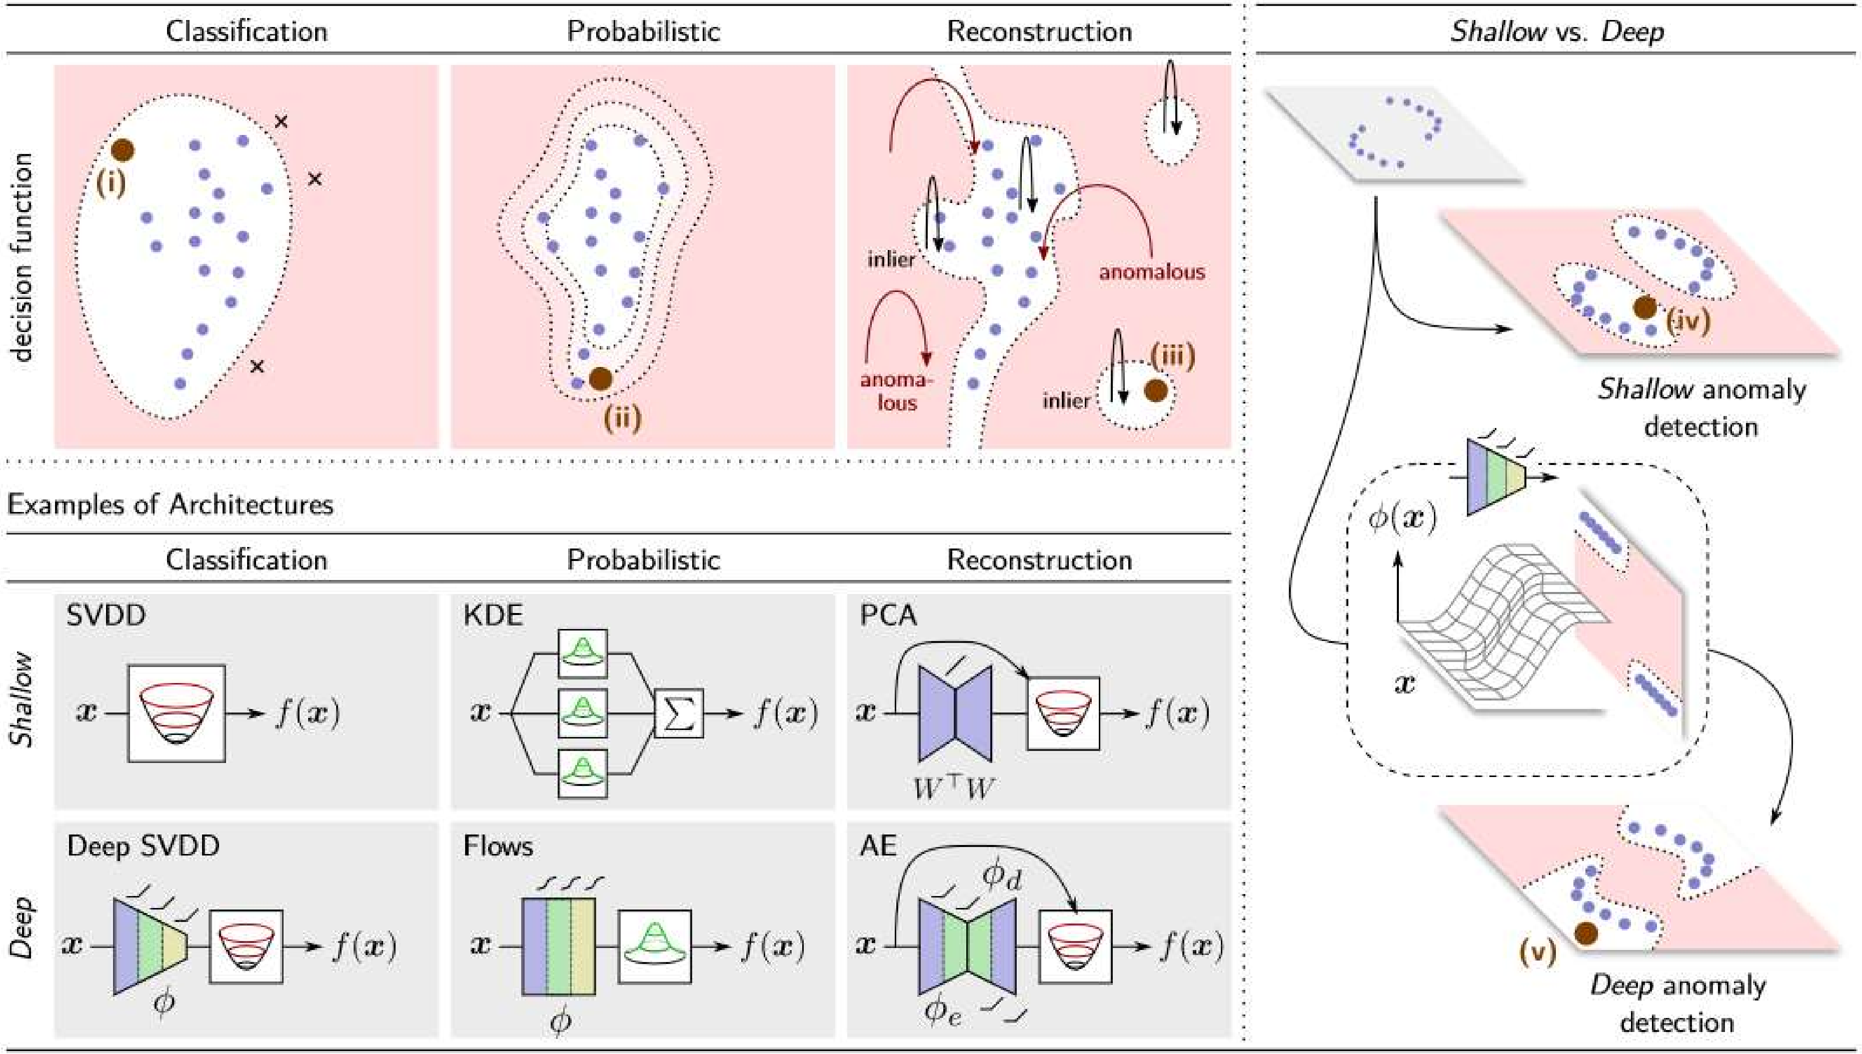

Proceedings of the IEEE (2021) 1-40

Anomaly detection based on reconstruction error

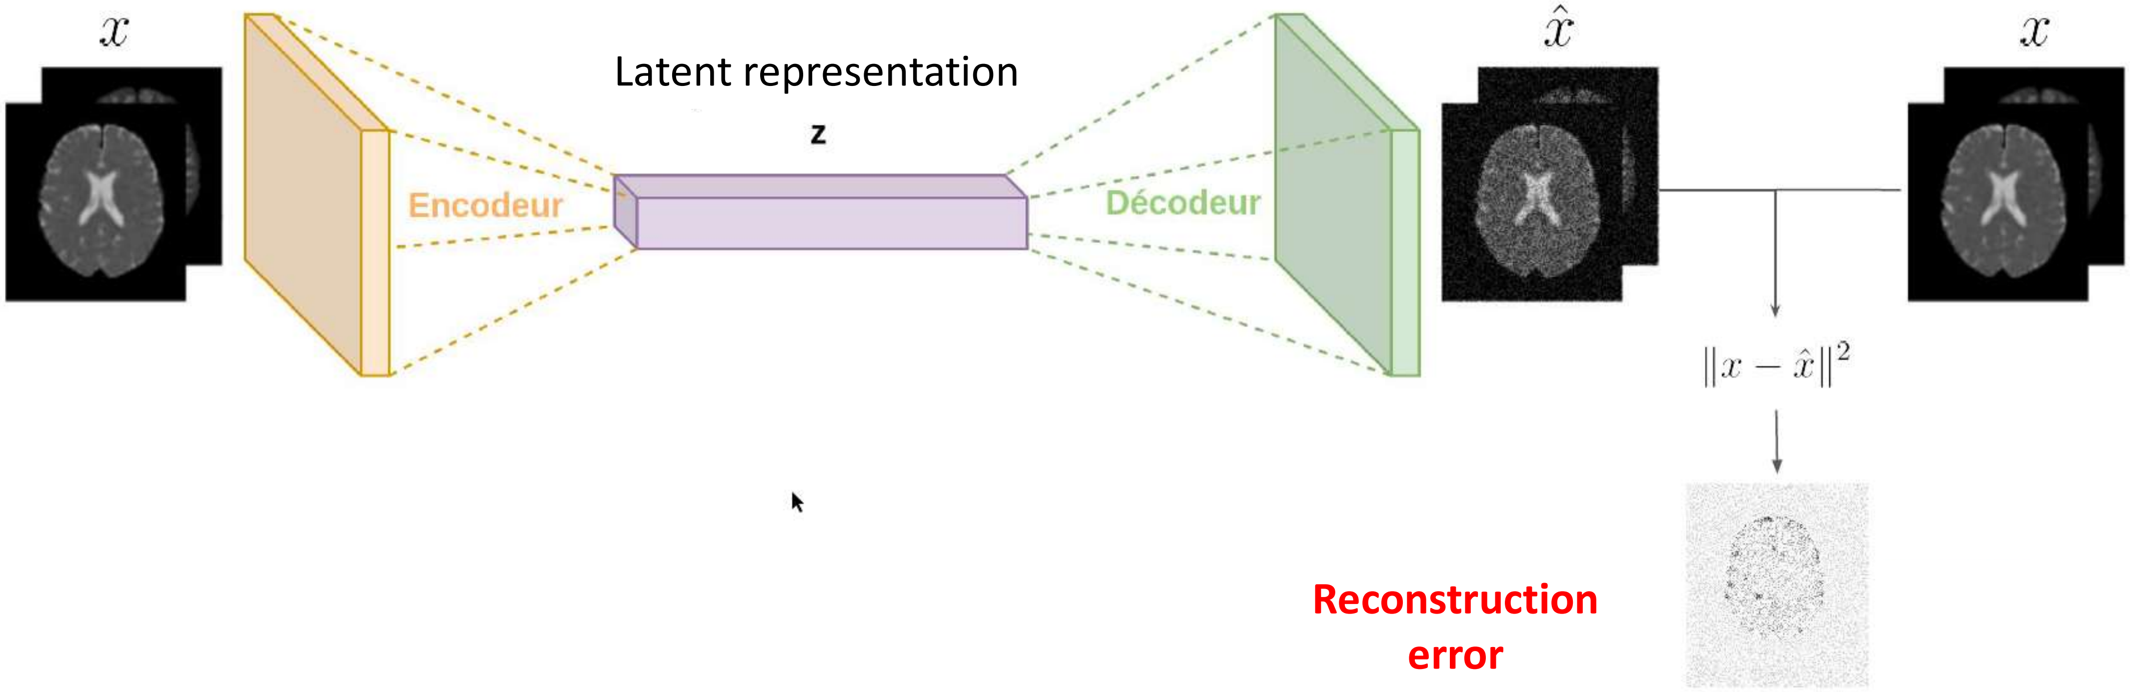

Anomaly detection based on reconstruction error after restauration of the latent representation

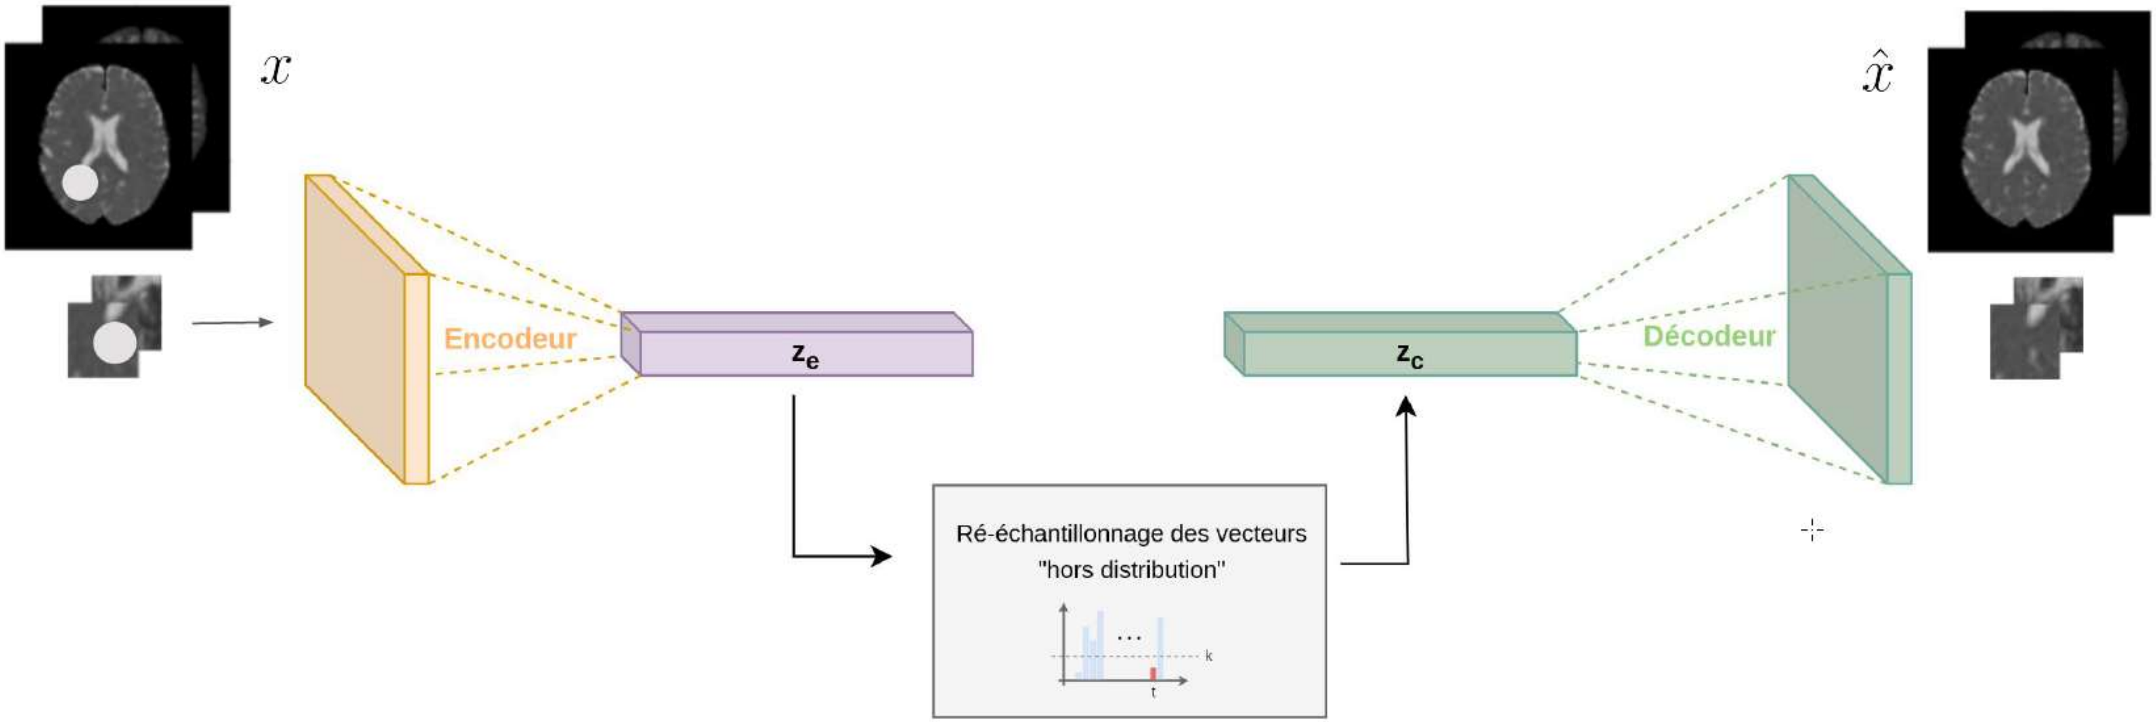

Vector quantized –variational auto-encoder (VQ-VAE)

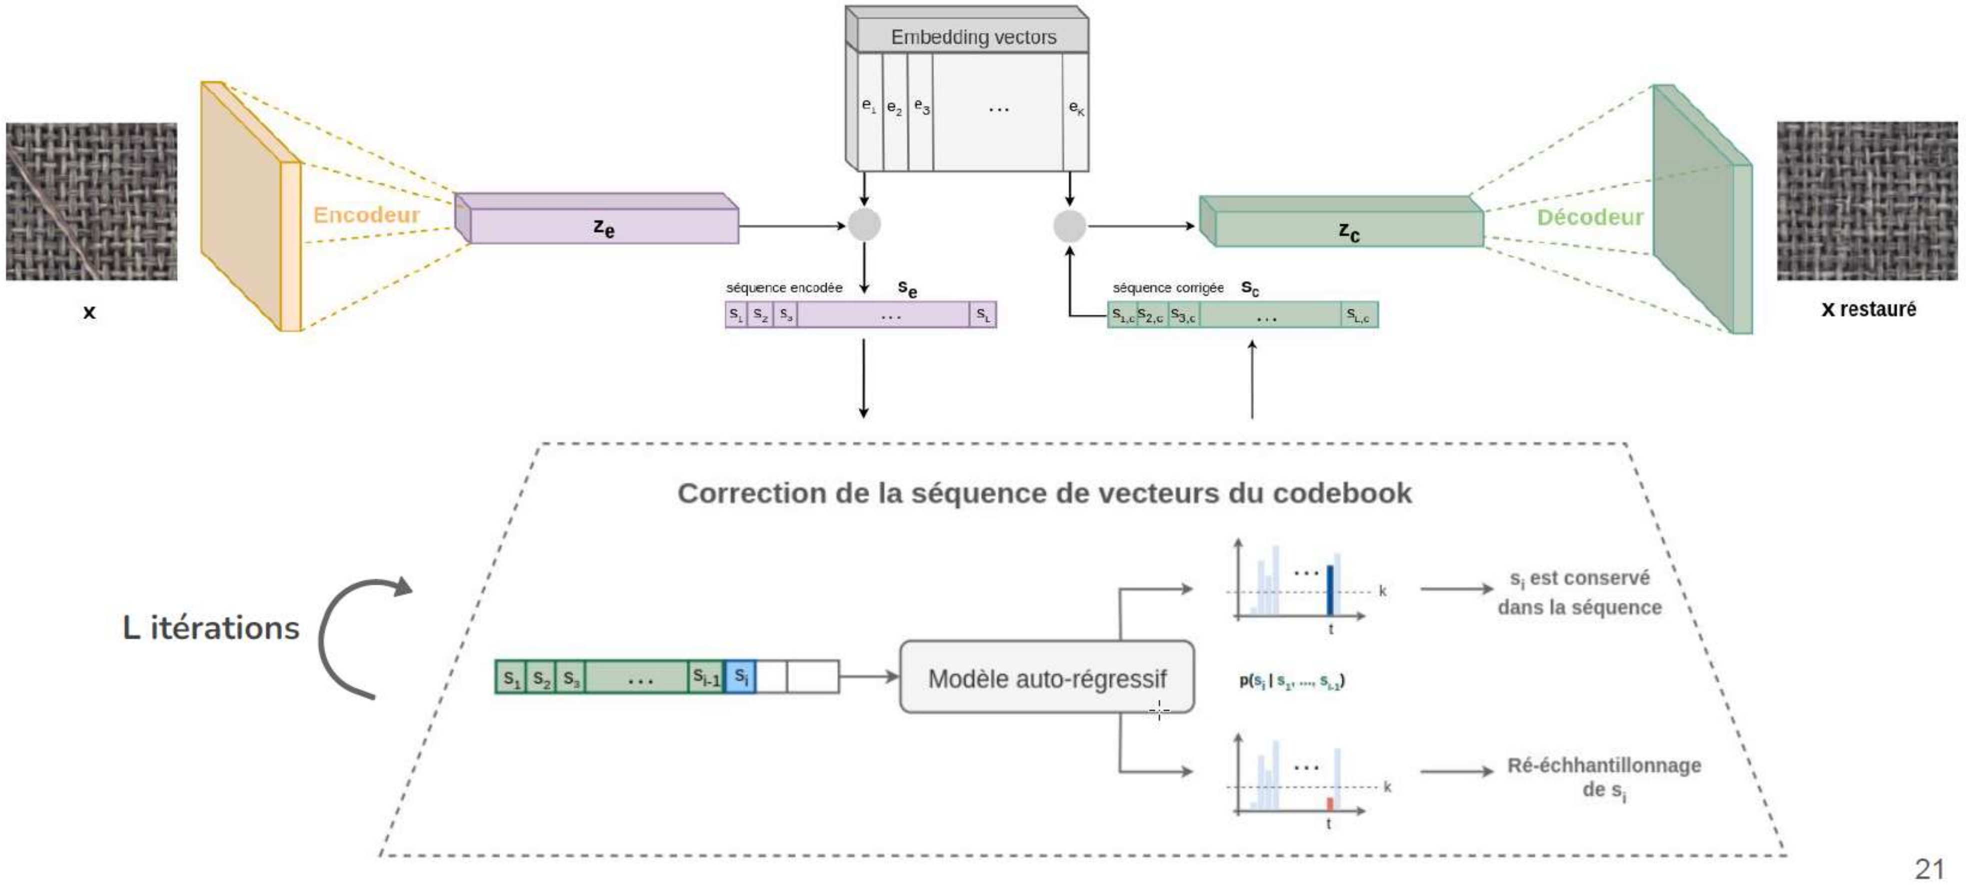

[Oord et al, 2017]

Anomaly detection based on estimation of the support of the normative distribution in the latent space

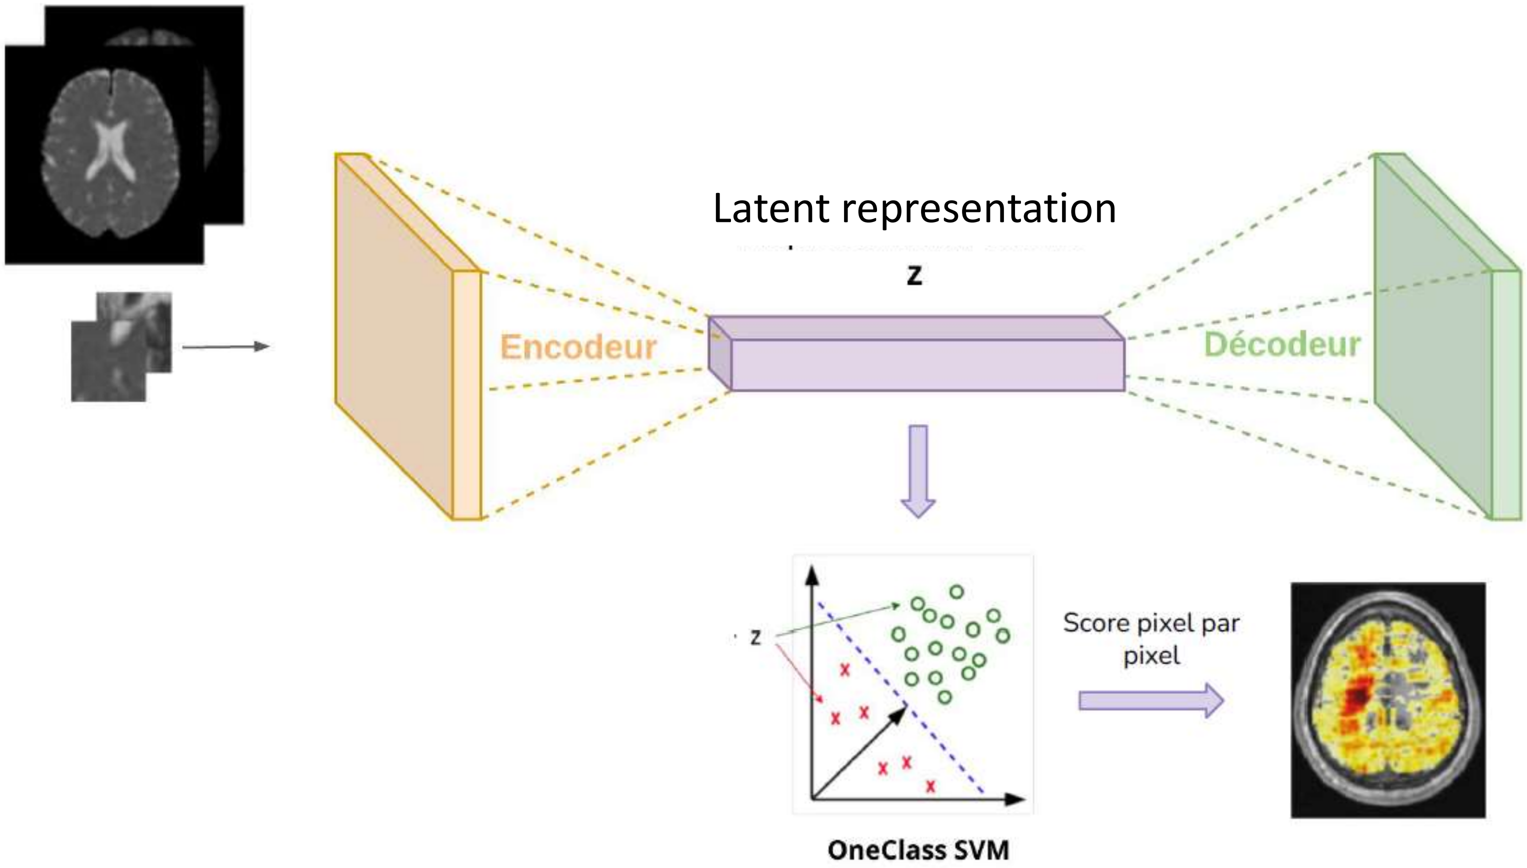

Anomaly score map

[Alaverdyan MEDIA 2020]

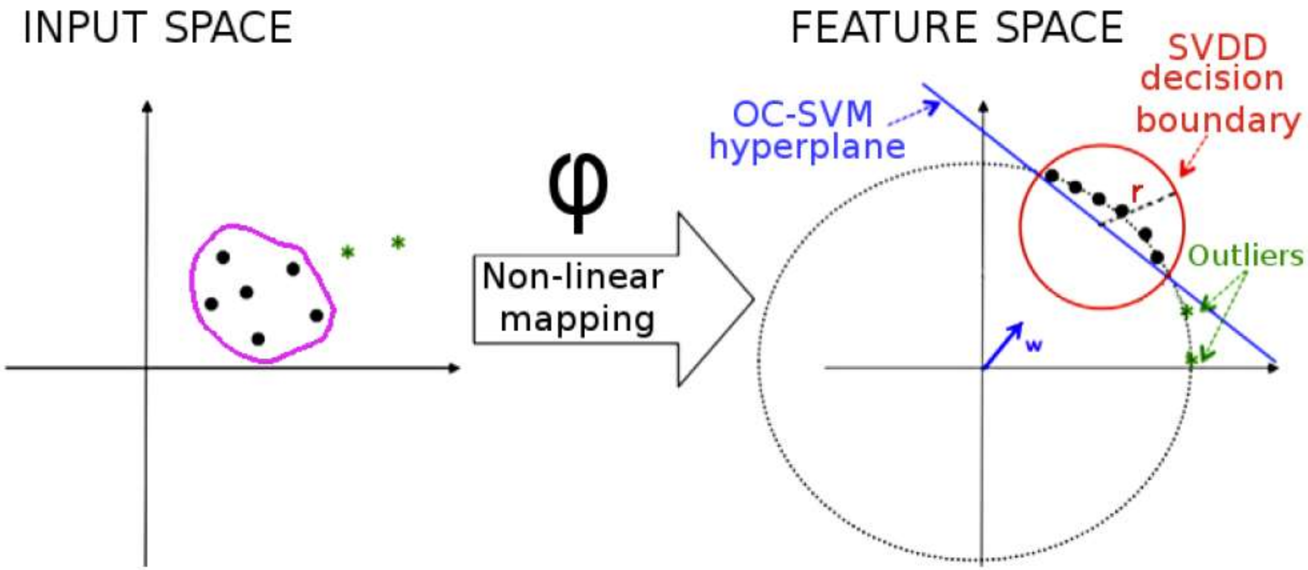

One class-SVM

$$\begin{cases} \min\limits_{\mathbf{w},\rho} & \frac{1}{2}\|\mathbf{w}\| - \rho \\ \textit{with} & \mathbf{w}^\top \mathbf{x}_i \geq \rho + \frac{1}{2}\|\mathbf{x}_i\|^2 \end{cases}$$

Support Vector Data Description

$$\begin{cases} \min\limits_{R \in \mathbb{R}, \mathbf{c} \in \mathbb{R}^d} & R^2 \\ \textit{with} & \|\mathbf{x}_i - \mathbf{c}\|^2 \leq R^2, \quad i = 1, \ldots, n \end{cases}$$

with $\rho = \frac{1}{2}(\|\mathbf{c}\|^2 - R)$ *and* $\mathbf{w} = \mathbf{c}$.

Benchmarking UAD models on the challenge WMH dataset

The WMH challenge dataset

White Matter Hyperintensities (WMH)

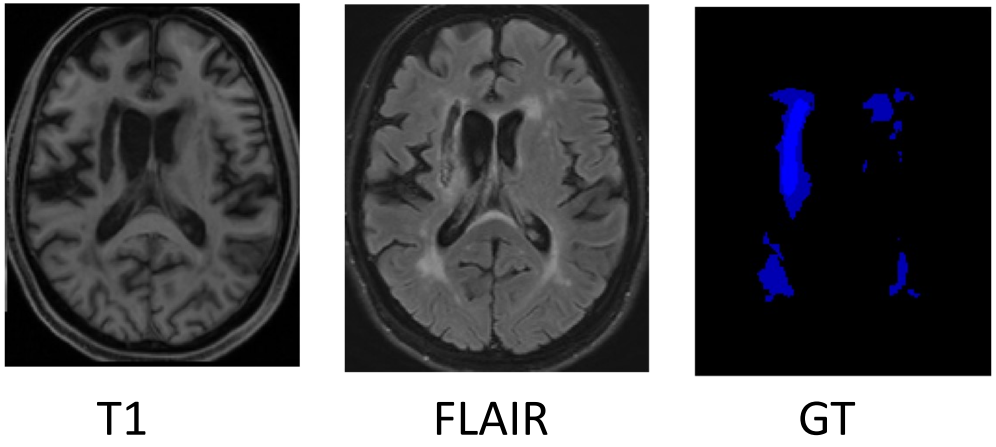

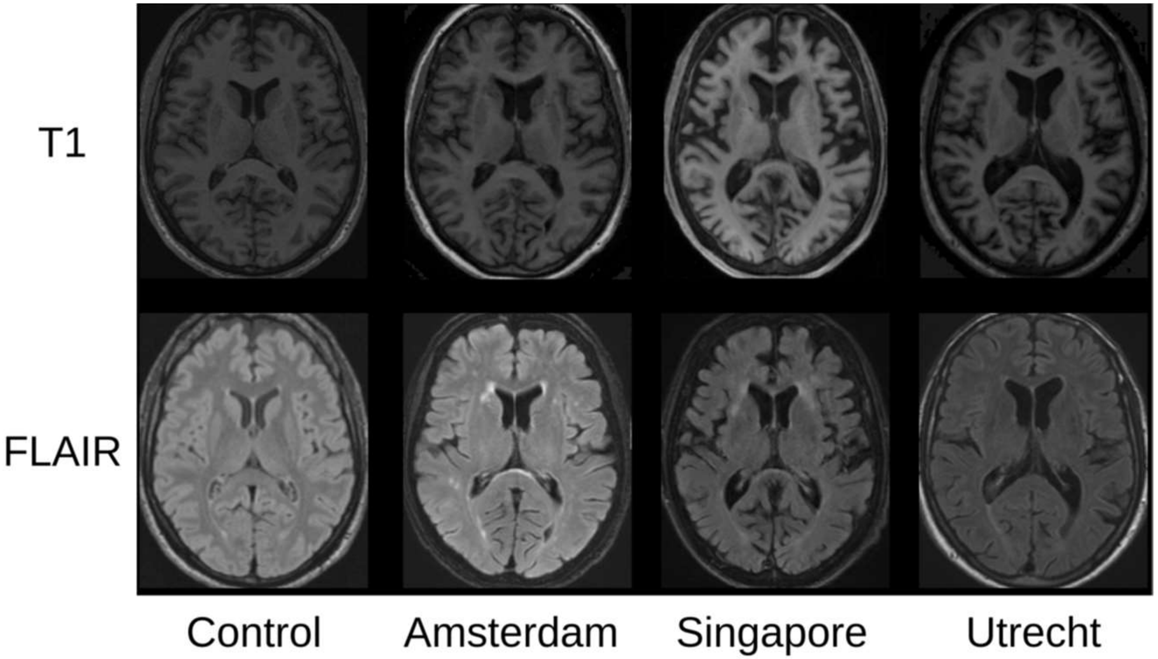

- The **control private dataset** :
  - 75 paired T1w and FLAIR MRI scans of healthy subjects
  - Acquired on a 1.5T Siemens Sonata scanner.
- The **WMH Challenge dataset:**
  - 60 T1w and FLAIR images
  - acquired on 3 different hospitals with 3 scanners of different manufacturers,
  - each image as its associated 3D lesion mask.

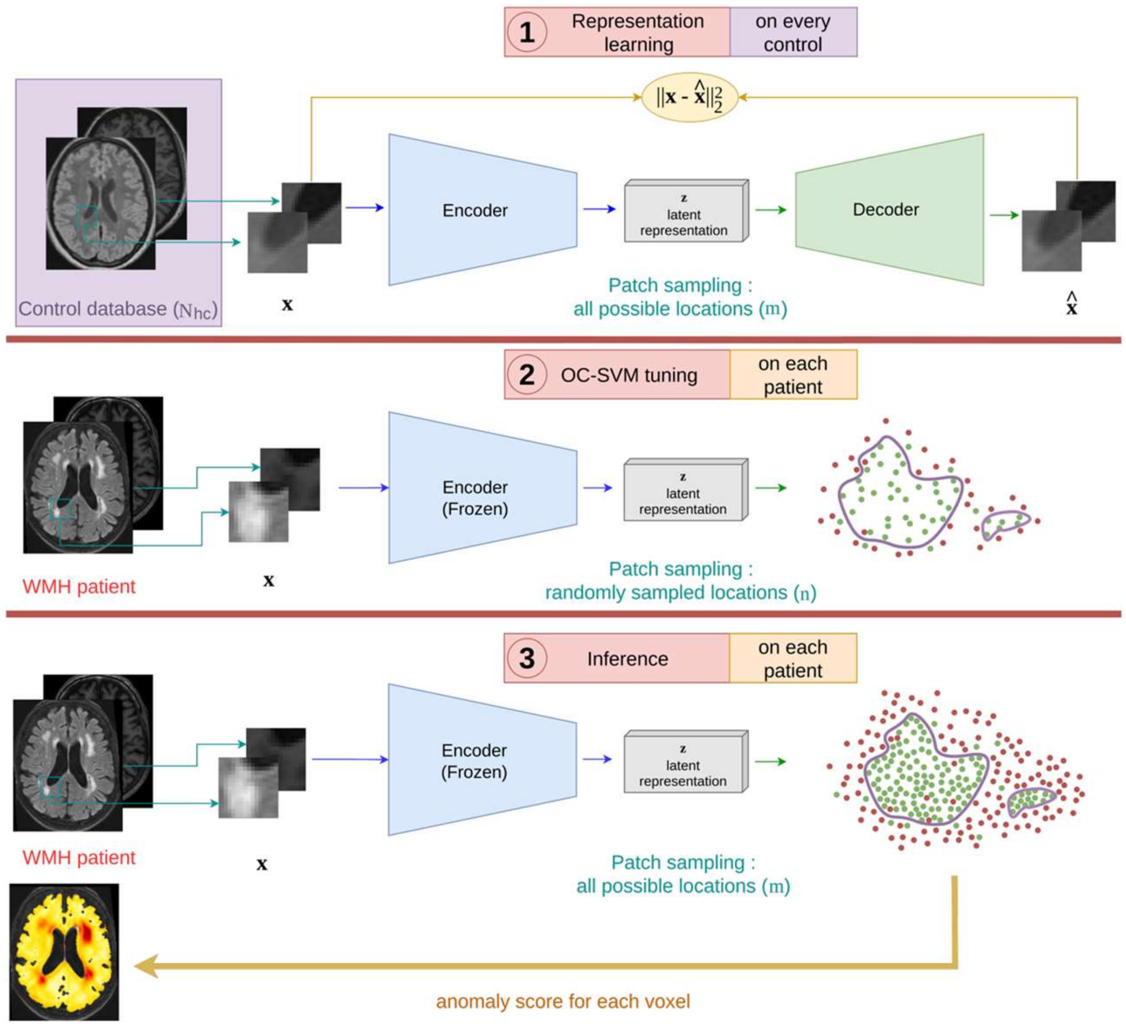

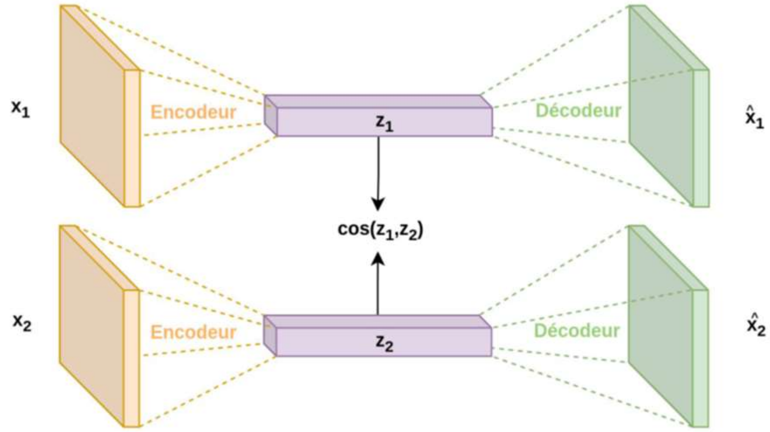

Siamese autoencoder

$$L_{SAE}(\mathbf{x_1},\mathbf{x_2}) = \sum_{t=1}^{2} ||\mathbf{x_t} - \hat{\mathbf{x}}_\mathbf{t}||_2^2 - \alpha \cdot cos(\mathbf{z_1},\mathbf{z_2})$$

[Pinon MIDL 2023]

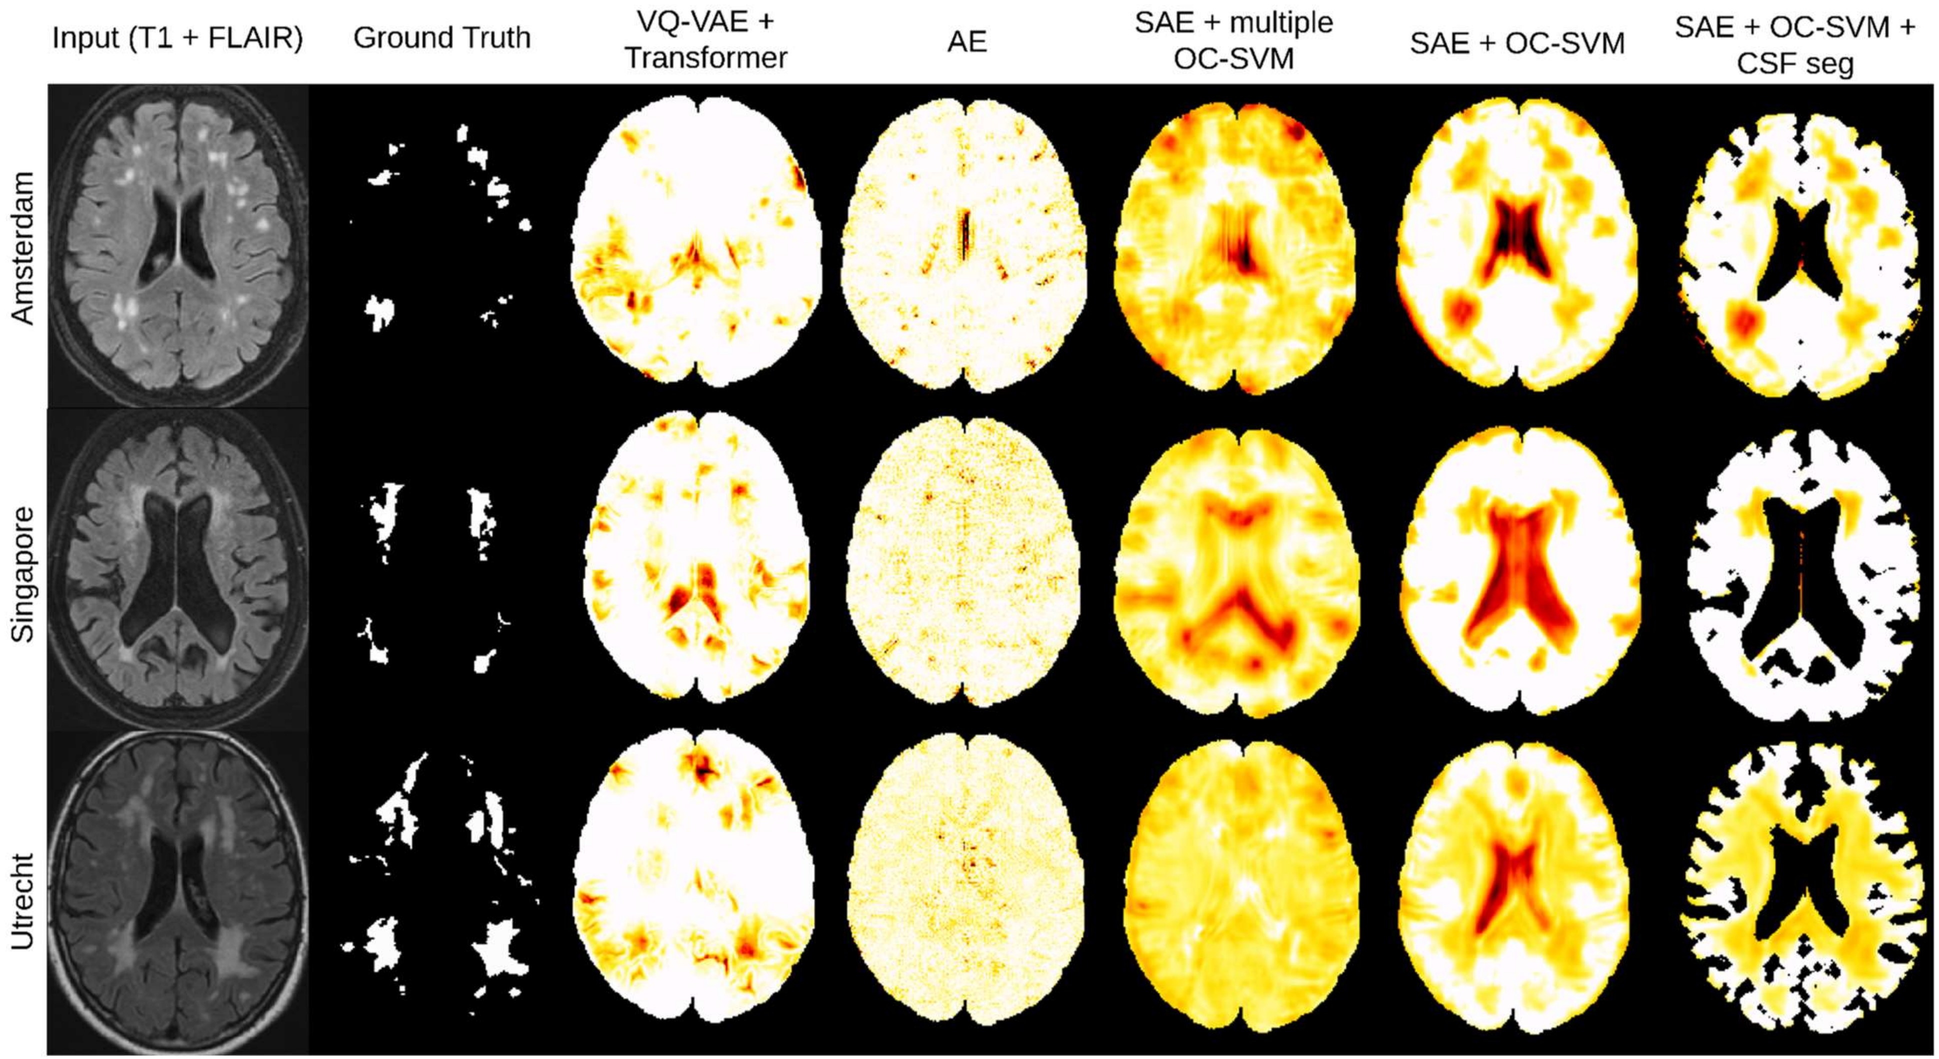

[Pinon MIDL 2023]

3 hospitals	VQ-VAE + Transformer (Pinaya)	AE (Baur)	SAE + multiple OC-SVM (Alaverdyan)	SAE + OC-SVM (Ours)	SAE +OC-SVM +CSF seg (Ours)
AU ROC	0.69 ± 0.13	0.53 ± 0.09	0.52 ± 0.19	$\mathbf{0.80} \pm 0.09$	$\mathbf{0.81} \pm 0.10$
AU ROC 30	0.40 ± 0.20	0.20 ± 0.12	0.19 ± 0.16	$\mathbf{0.48} \pm 0.20$	$\mathbf{0.59} \pm 0.17$
AU PRC	0.065 ± 0.079	0.028 ± 0.030	0.023 ± 0.031	$\mathbf{0.084} \pm 0.099$	$\mathbf{0.165} \pm 0.168$
AU PRO	0.55 ± 0.10	0.50 ± 0.08	0.43 ± 0.17	0.71 ± 0.11	$\mathbf{0.80} \pm 0.07$
AU PRO 30	0.19 ± 0.13	0.15 ± 0.07	0.09 ± 0.13	0.33 ± 0.18	$\mathbf{0.48} \pm 0.13$
$\lceil$ Dice $\rceil$	0.11 ± 0.10	0.06 ± 0.05	0.05 ± 0.05	$\mathbf{0.14} \pm 0.13$	$\mathbf{0.22} \pm 0.17$

[Pinon MIDL 2023]

- The proposed SAE+ OC-SVM model performs on par with the state of the art UAD model on the WMH dataset
- Limits of the WMH dataset as a reference dataset for UAD benchmarking :
  - Anomaly are 'easy' to detect on FLAIR images
  - Anomaly are not visible on T1 images
  - Median age of the population is high → physiological normal anomalies due to brain aging process (cortex shrinkage..)

Unsupervised anomaly detection of the epileptogenic zone (EZ)

Management of the patient with medically refractory epilepsy

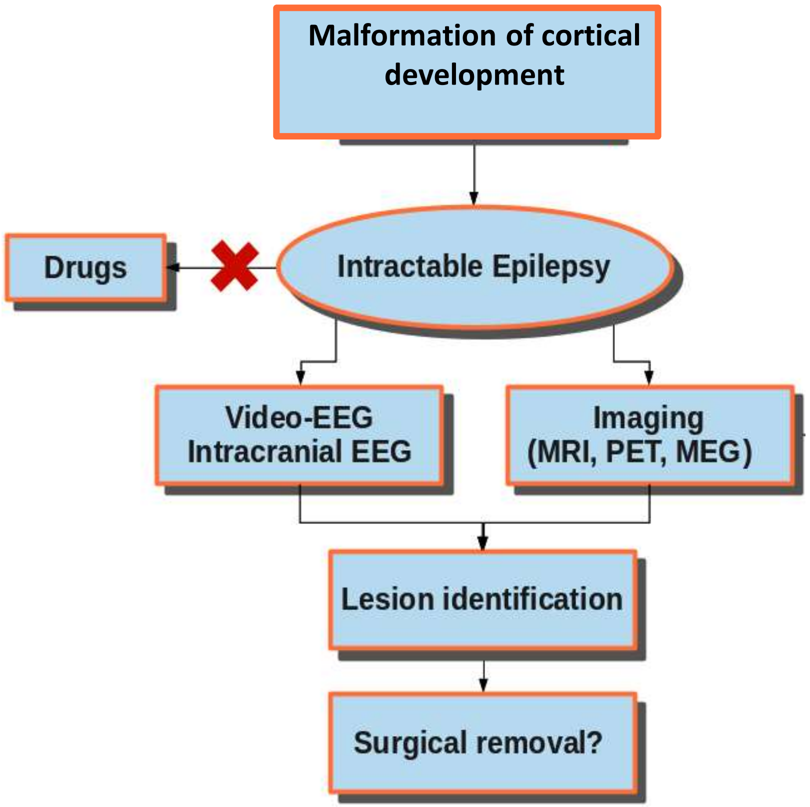

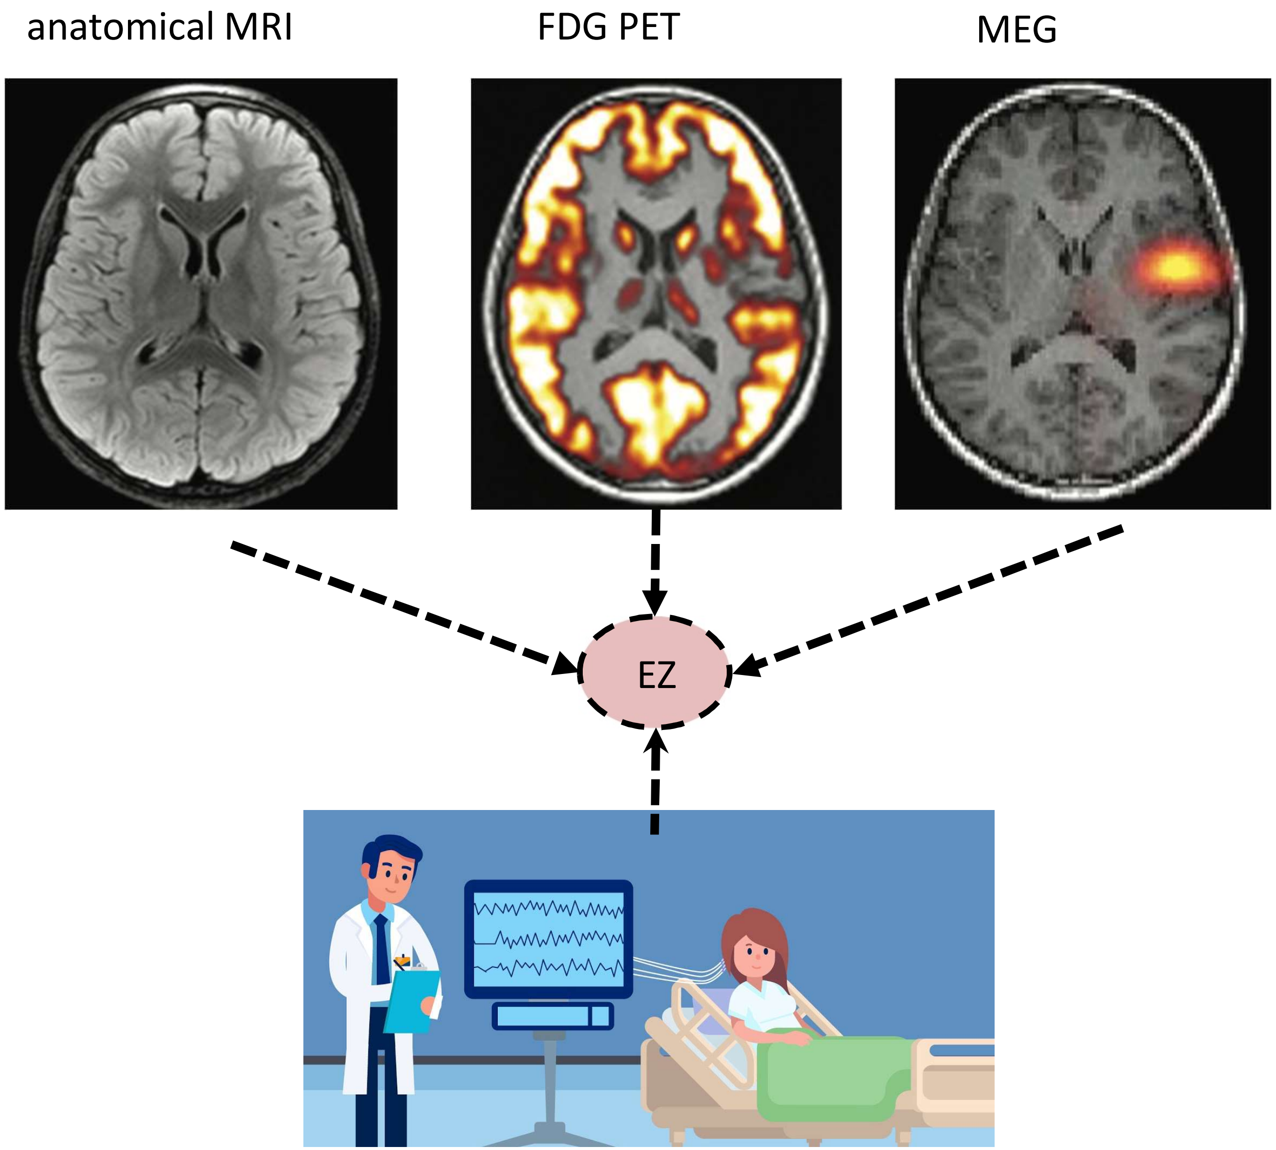

Long-Term Video EEG

EZ = Epileptogenic Zone

- Imaging plays a crucial role in the surgical planning

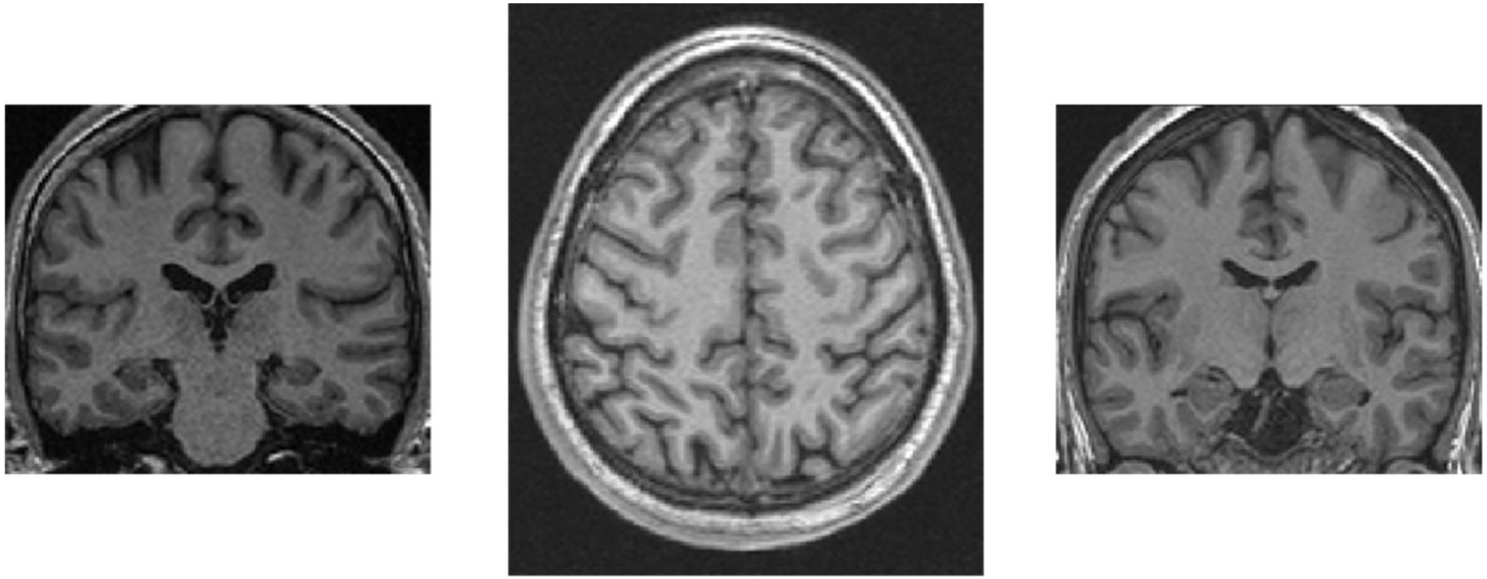

- But the detection task is hard

- 20-30% of patient MRI are « negative » meaning the physician considered the exam as « non pathological »

- Detecting and localizing the epileptogenic zone → significant improvement of good surgery outcome

[Nagae et al. 2016]

Medical Image Analysis 60 (2020) 101618

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Regularized siamese neural network for unsupervised outlier detection on brain multiparametric magnetic resonance imaging: Application to epilepsy lesion screening

Zaruhi Alaverdyan[a], Julien Jung[b], Romain Bouet[b], Carole Lartizien[a,*]

[a] Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1206, F69621, Lyon, France
[b] Lyon Neuroscience Research Center, CRNL, INSERM U1028, CNRS UMR5292, University Lyon 1, Lyon, France

References	Subjects	Imaging modality	Imaging features	Classifiers	Main outcomes
Alaverdyan et al. (2020)	21 FE, 75HC	T1, FLAIR	Signals	SVM, RSN	Sens. = 0.62 to detect anomaly lesion

Dataset

- The **private control dataset** :
  - 75 paired T1w and FLAIR MRI scans of healthy subjects
  - Acquired on a 1.5T Siemens Sonata scanner.
- The **private epilepsy dataset:**
  - 21 T1w and FLAIR images
  - Acquired on a 1.5T Siemens Sonata scanner
  - Référence of EZ localization is based on
    - Post-surgery Engel score
    - Manual annotation of the lesion on the MRI based on clinical report and sEEG analysis

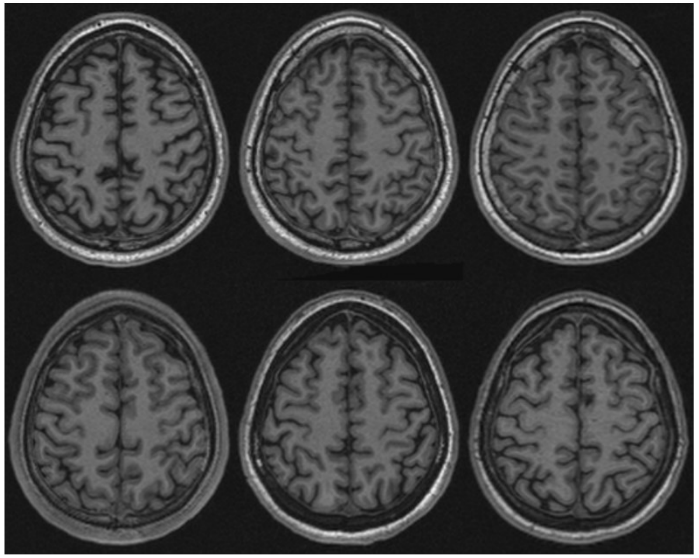

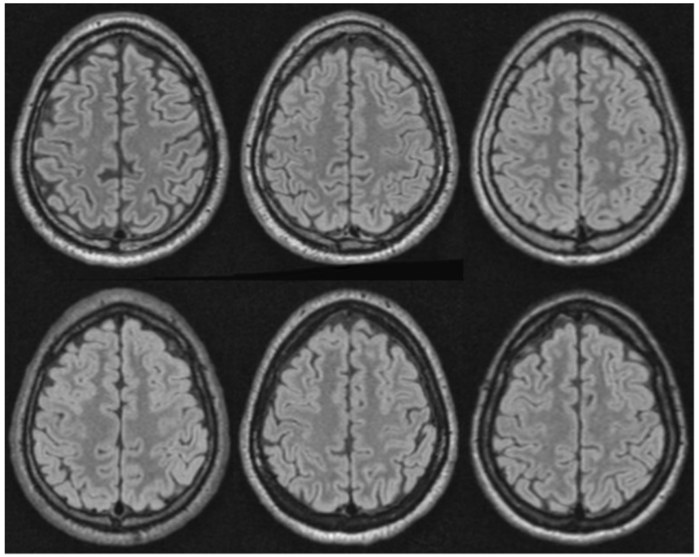

Example T1 and FLAIR images of patients

Post-surgical scans
thermocoagulation reports

Positive outcome :
seizure-free patients

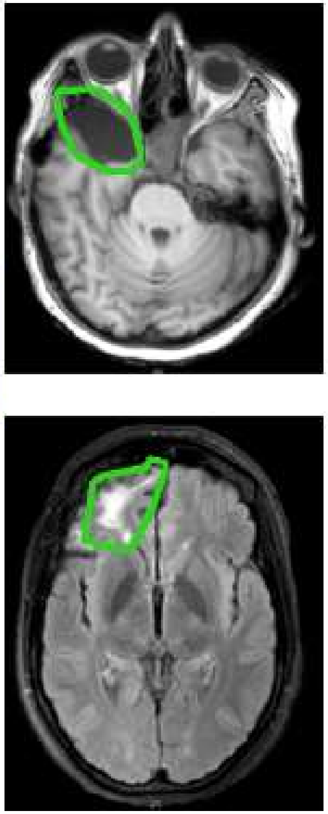

Ground truth : 18 *MRI-negative* patients

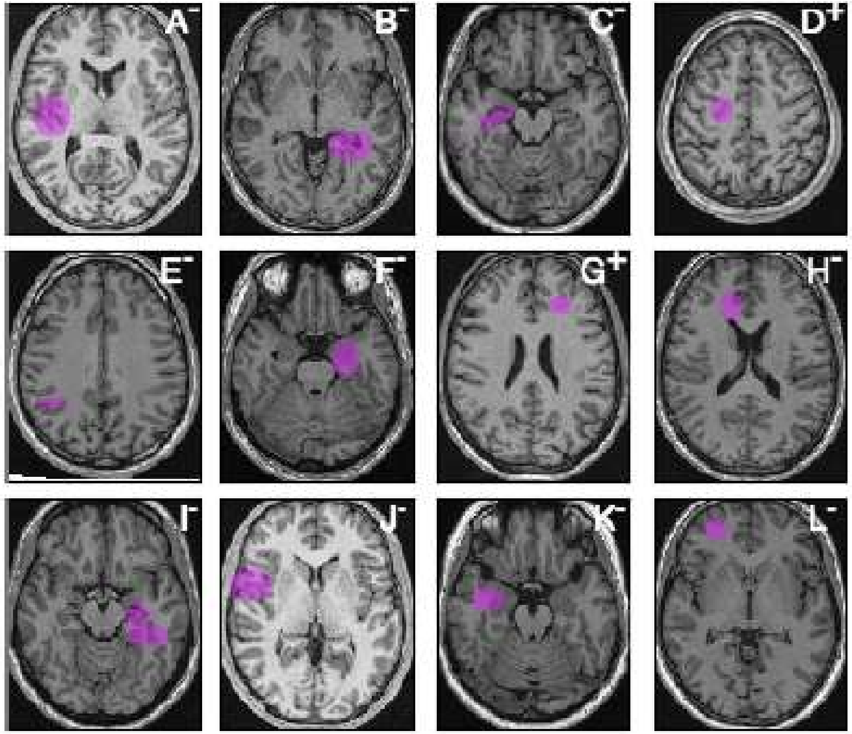

Lesion annotations obtained in collaboration with Dr J. Jung

Unsupervised anomaly detection

Step 1 : train a self-supervised model on healthy control data

Healthy control database

Deep siamese autoencoder for representation learning

Outlier detection with oc-SVM in the latent space

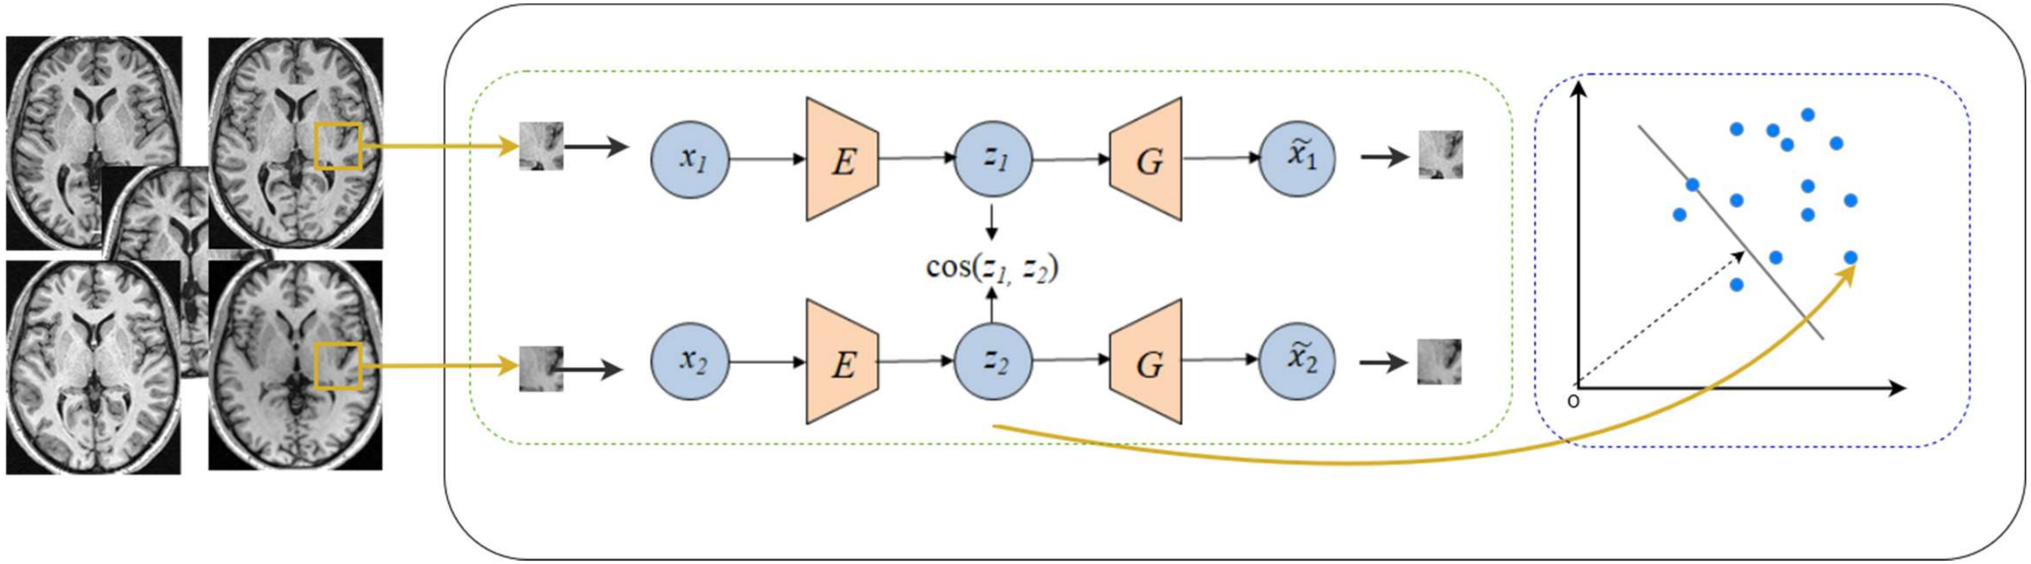

[Alaverdyan et al MEDIA 2020]

Unsupervised anomaly detection

Step 2 : detect anomalous pattern in epilepsy patients

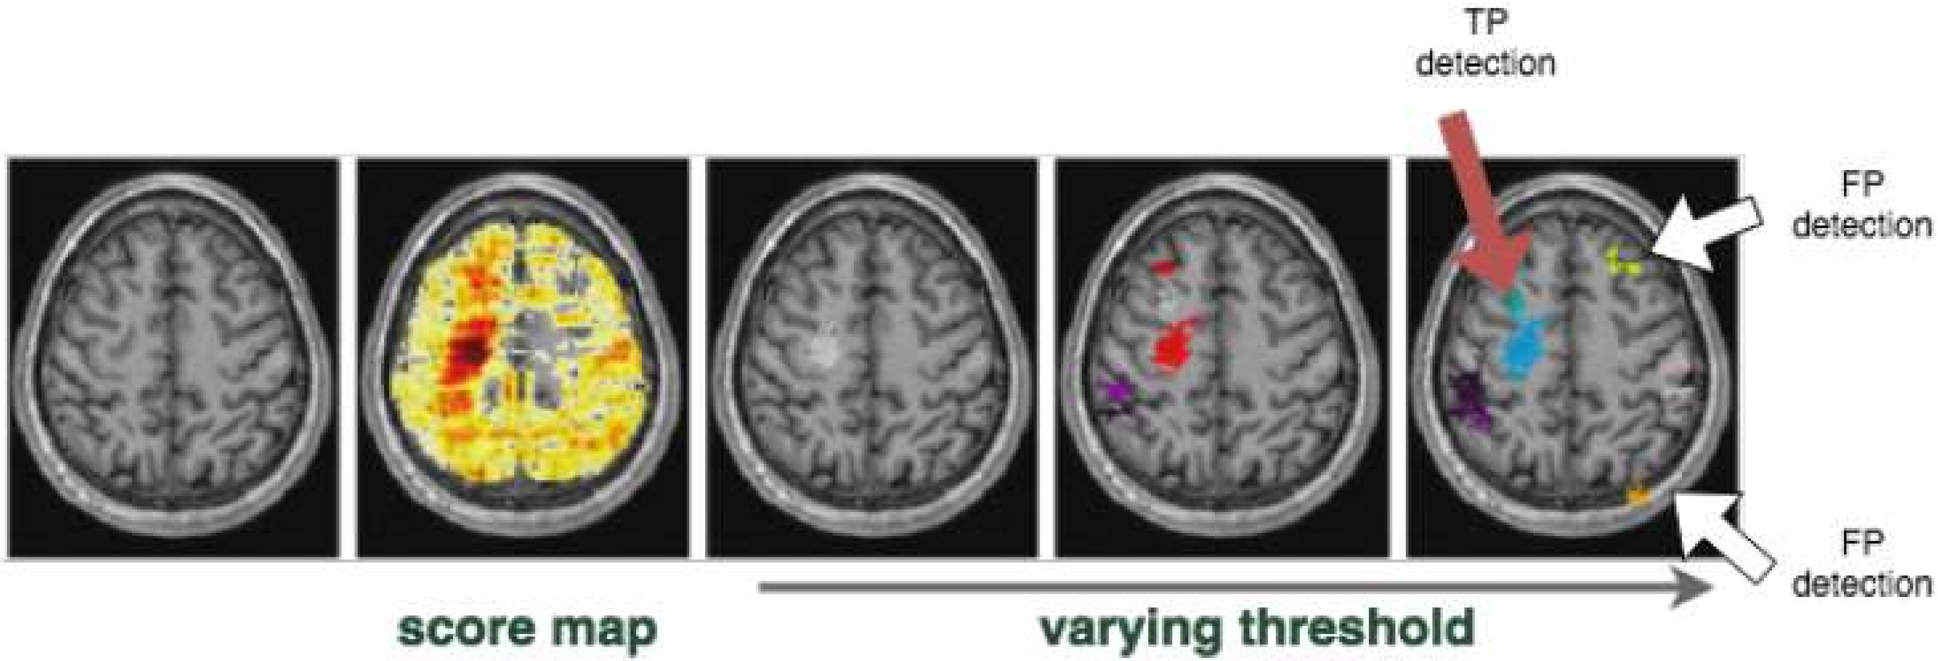

Predicted lesion map (MIP).
The brighter colour, the most suspicious the clusters

[Alaverdyan et al MEDIA 2020]

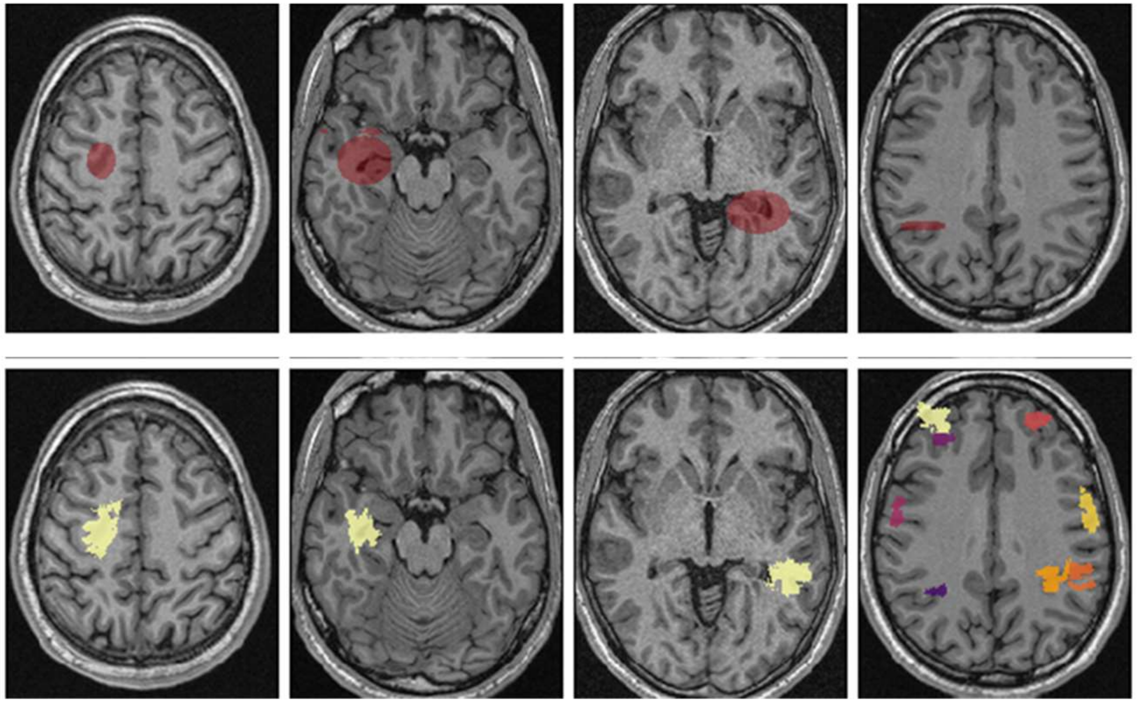

Predicted lesion map (MIP).
The brighter colour, the most suspicious the clusters

- **Sensitivity** :
  - 62% on 21 negative MRI negative exams of epilepsy patients
- **Specificity**
  - Mean rank of the detected clusters : 3.5

[Alaverdyan et al MEDIA 2020]

Dataset

- How to improve ?
- Include PET modality → complementary information

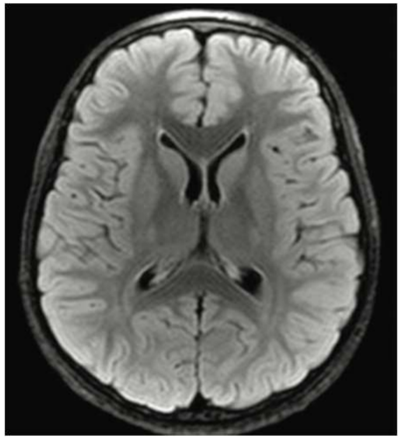

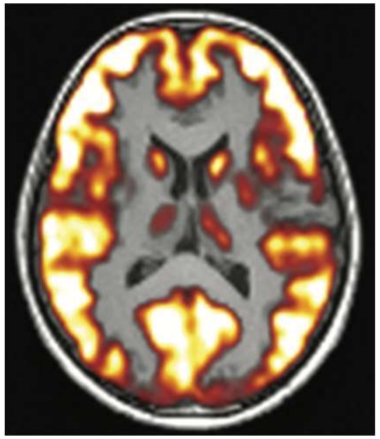

Dataset

- The **private control dataset** :
  - **35 paired** T1w MRI and FDG scans of healthy subjects
  - Acquired on a 1.5T Siemens Sonata scanner.
  - 40 T1w MRI of healthy subject with NO paired FDG PET examen

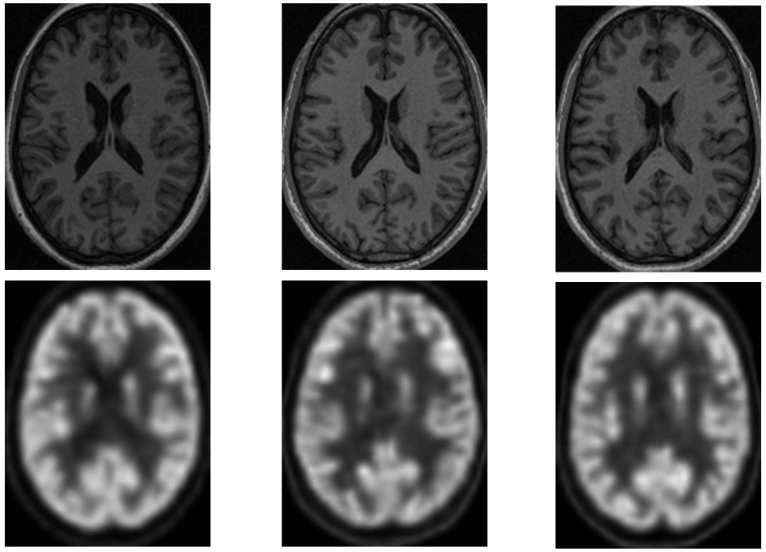

- The **private epilepsy dataset:**
  - 21 T1w and PET images
  - Acquired on a 1.5T Siemens Sonata scanner
  - Référence of EZ localization is based on
    - Post-surgery Engel score
    - Manual annotation of the lesion on the MRI based on clinical report and sEEG analysis

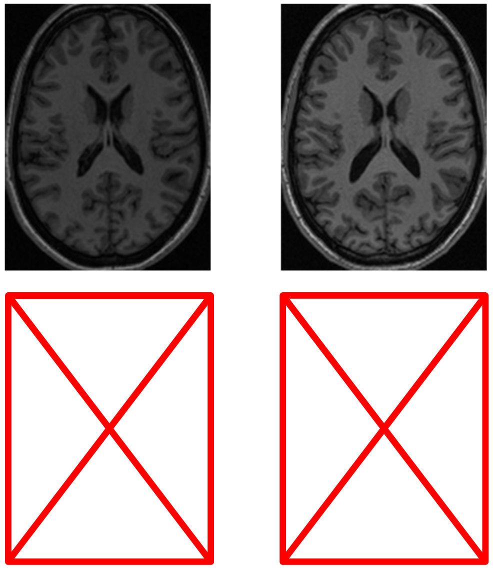

Challenge : Missing data in the control population

GAN-based synthetic FDG PET images from T1 brain MRI can serve to improve performance of deep unsupervised anomaly detection models

Daria Zotova[1], Julien Jung[2], and Carole Lartizien[1][0000-0001-7594-4231]

[1] Univ Lyon, CNRS, Inserm, INSA Lyon, UCBL, CREATIS, UMR5220, U1206, F-69621, Villeurbanne, France
{Daria.Zotova,Carole.Lartizien}@creatis.insa-lyon.fr
[2] Lyon Neuroscience Research Center, CRNL, INSERM U1028, CNRS UMR5292, University Lyon 1, Lyon, France

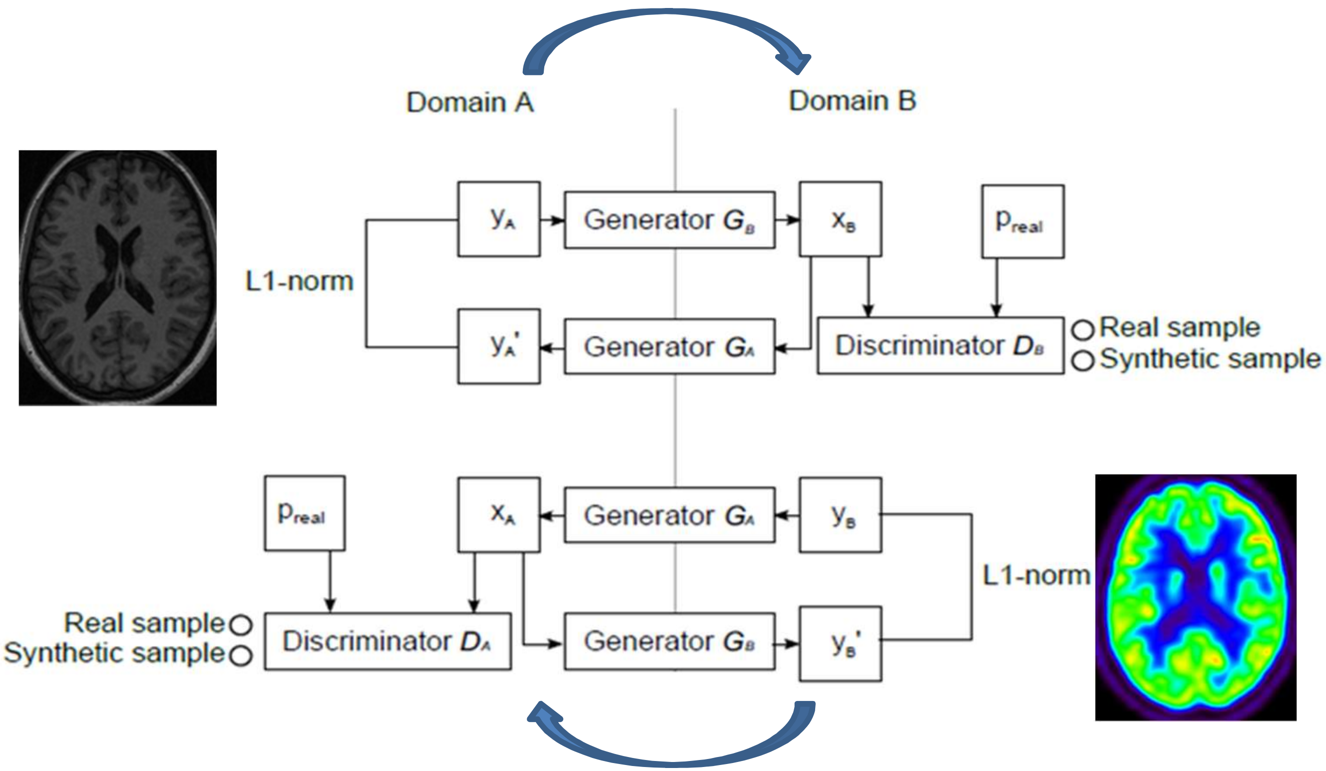

Structure of a CycleGAN

Generative modeling of noramtive PET data with GAN (generative adversarial networks)

[Zotova et al MICCAI SASHIMI 2021]

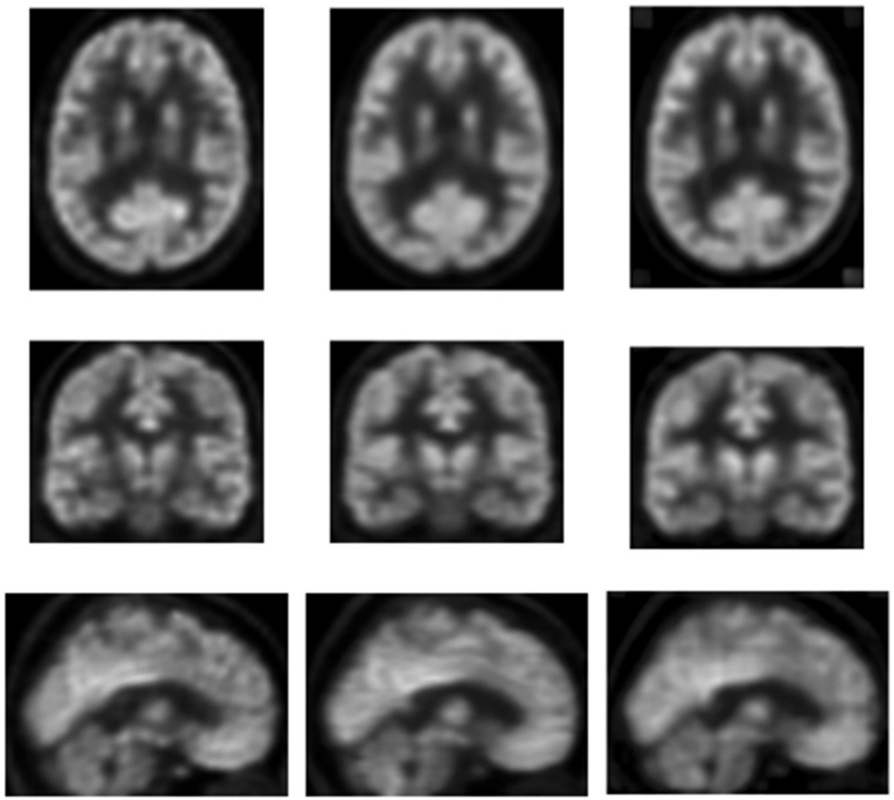

	35 T1+PET real	35 T1+fake PET semi3d	35 T1+fake PET 3dpatch	Color code
Patient A	-	-	-	lobe temporal L
Patient B	x (5)	x (4)	x (1)	lobe temporal R
Patient C	x (3)	x (4)	x (2)	pole temporal R
Patient D	x (8)	-	x (1)	insula
Patient E	-	-	-	other
Patient F	-	x (4)	x (5)	
Patient G	x (1)	x (1)	x (3)	
Patient H	-	x (3)	x (5)	
Patient I	x (7)	x (1)	x (1)	
Patient J	x (4)	x (3)	-	
Patient K	-	-	-	
Patient L	-	-	x (2)	
# of detected	6	7	8	
mean rank	4,7	2,9	1,9	

Qualitative analysis of the synthetic normative PET data

Performance of the UAD model trained with real or synthetic normative PET data

[Zotova et al MICCAI SASHIMI 2021]

Unsupervised anomaly detection

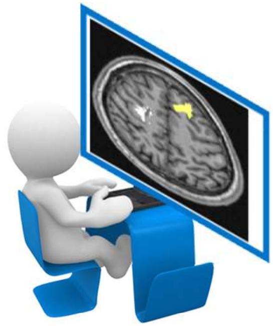

From Research..

Some methodological challenges to address ..

..to clinic

- **Improve performance**
  - Inject some priors → toward weak supervision
  - Efficiently fuse multi-modality imaging and non imaging data, accounting for missing modalities -> unsupervised representation modeling (couplage LLM and CV modeling..)
- **Generalize** well
  - Same level of performance regardless of data origin and quality
- Provide some **confidence level** on predictions
- Be **respectful of privacy**

Thank you for your attention!

Acknowledgments

Most of the illustrations and the reported results of this talk were produced by PhD students of my group, Meriem El Azami, Zaruhi Alaverdyan, Daria Zotova, Nicolas Pinon, Matthis Manthe, Robin Trombetta.

More on our work

https://scholar.google.fr/citations?user=na6f2dQAAAAJ&hl=fr