



HAL
open science

Stroke Research Projects at Empenn

Youwan Mahé, Elise Bannier, Francesca Galassi

► **To cite this version:**

Youwan Mahé, Elise Bannier, Francesca Galassi. Stroke Research Projects at Empenn. Mind Seminar, Mar 2026, Saclay (CEA), France. 2026. <hal-05573442>

HAL Id: hal-05573442

<https://inria.hal.science/hal-05573442v1>

Submitted on 31 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License


Stroke Research Projects at Empenn

MIND Seminar
17th March 2026

Youwan MAHÉ

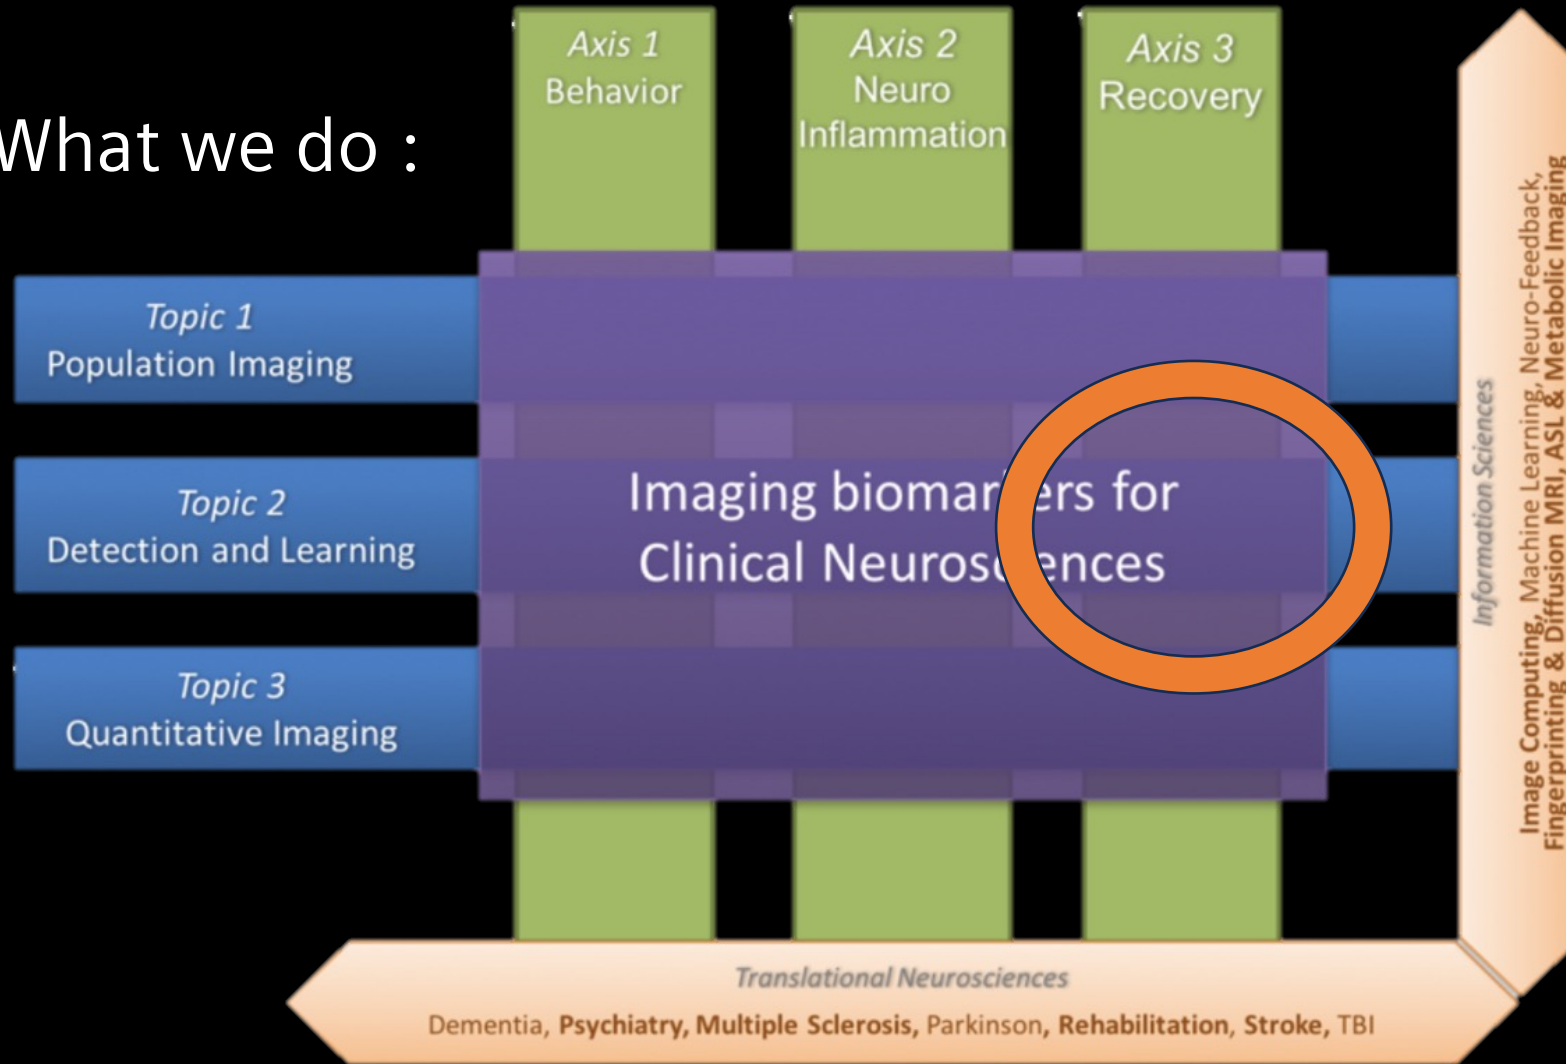
Univ Rennes, Inria, CNRS, Inserm, IRISA UMR 6074, Empenn

Siemens Healthineers



Empenn

What we do :



Today's scope

Context

Brain stroke affects:

1 in 5 people (mean age 74)

From which **1/3** suffer from **long-term disabilities**

Santé publique France, 4 mars 2025



Facial
paralysis



Hemiparesis



Speech
disorder

Stroke causes long-term disabilities such as: **hemiparesis, speech disorders**, etc.

- The accuracy of the **prognosis** at the individual level remains a challenge.
- Imaging research for recovery would benefit from automated tools to both **speed up** and **improve** the image processing.
- Then, in the long term, **personalize** the care by developing **predictive models**.

What we do at Empenn


1. Multimodal fusion for fully supervised stroke lesion segmentation
2. From the lab to the clinic, transferring models in a clinical research context
3. Unsupervised learning, detecting anomalies beyond lesions to better predict recovery pathways

What we do at Empenn

1. **Multimodal fusion for fully supervised stroke lesion segmentation**
2. From the lab to the clinic, transferring models in a clinical research context
3. Unsupervised learning, detecting anomalies beyond lesions to better predict recovery pathways

Method

Segmentation of chronic stroke lesions




- **Dataset :**
 - **ATLAS v2.0 dataset (T1-w, n=651)** was used for baseline training
 - An **internal dual-modality dataset (T1-w + FLAIR)** was used for fine-tuning (27 subjects) and testing (39 subjects, multi-centric).
- **Preprocessing:** skull-stripping (HD-BET), RAS reorientation, registration (T1-w to FLAIR), N4 bias correction.
- **Architecture:** 3D U-Net (nnU-Net, ResEnc U-Net size M) with deep supervision, dropout (0.2), and data augmentation
- **Postprocessing:** Binarisation threshold of 0.2 and lesion filtering (≥ 10 mm³).


Data-description



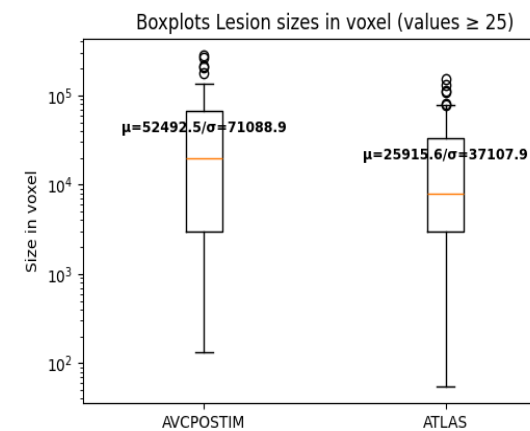
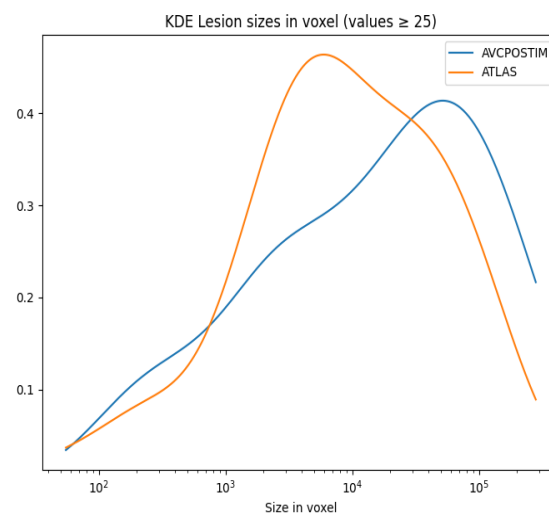
ATLAS



Internal

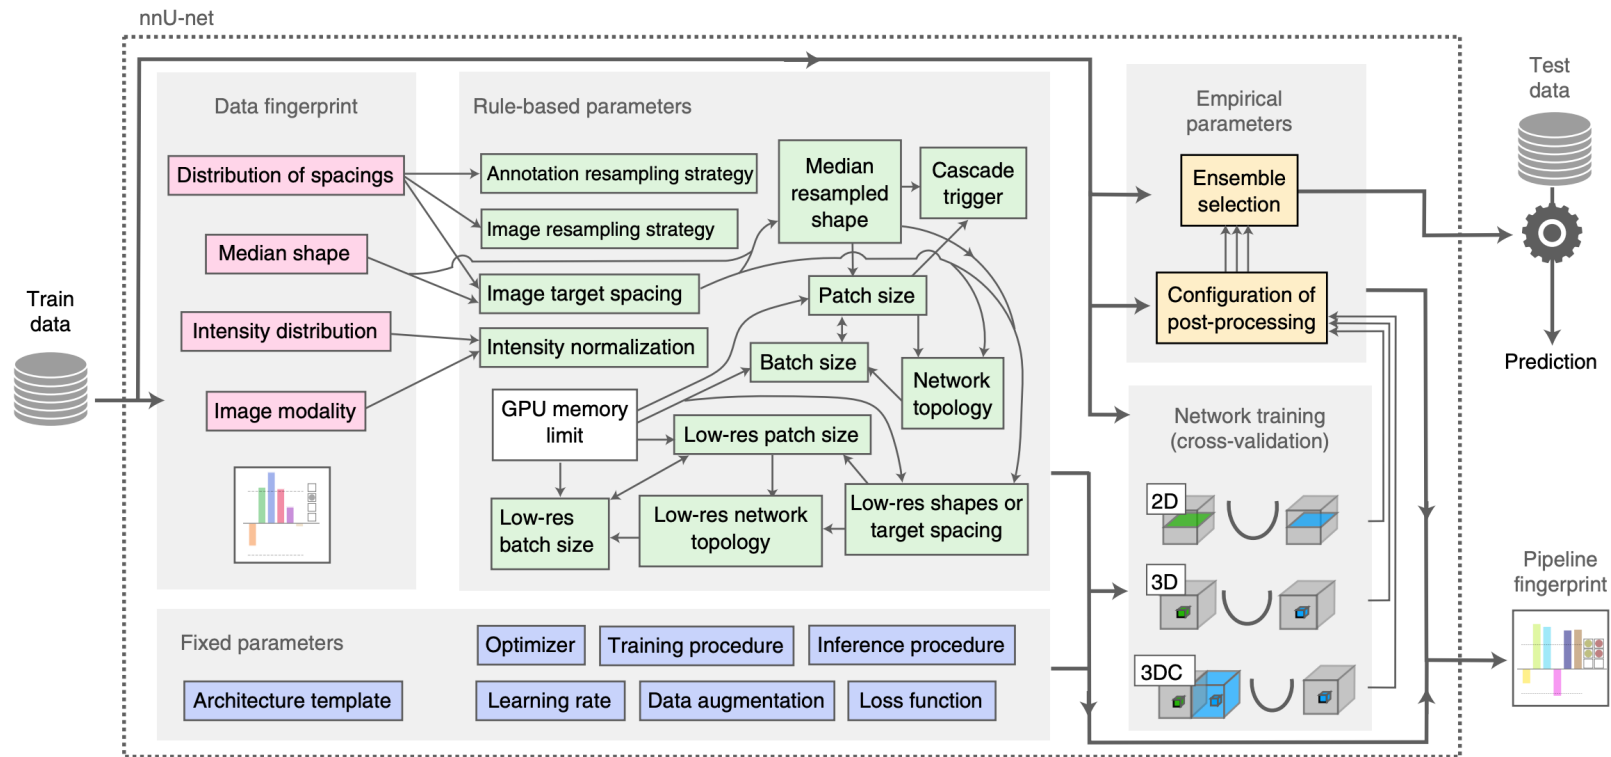


	ATLAS	Internal
Dataset size	655	52
Ratio Nb. Sub.	92.6%	7.3%
Ratio 3T / 1.5T	11.6	N/A



The training framework

The core Idea : « Self-configuration » of most hyperparameters



Isensee, F, *et al.* nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat Methods* 18, 203–211 (2021).

The training framework

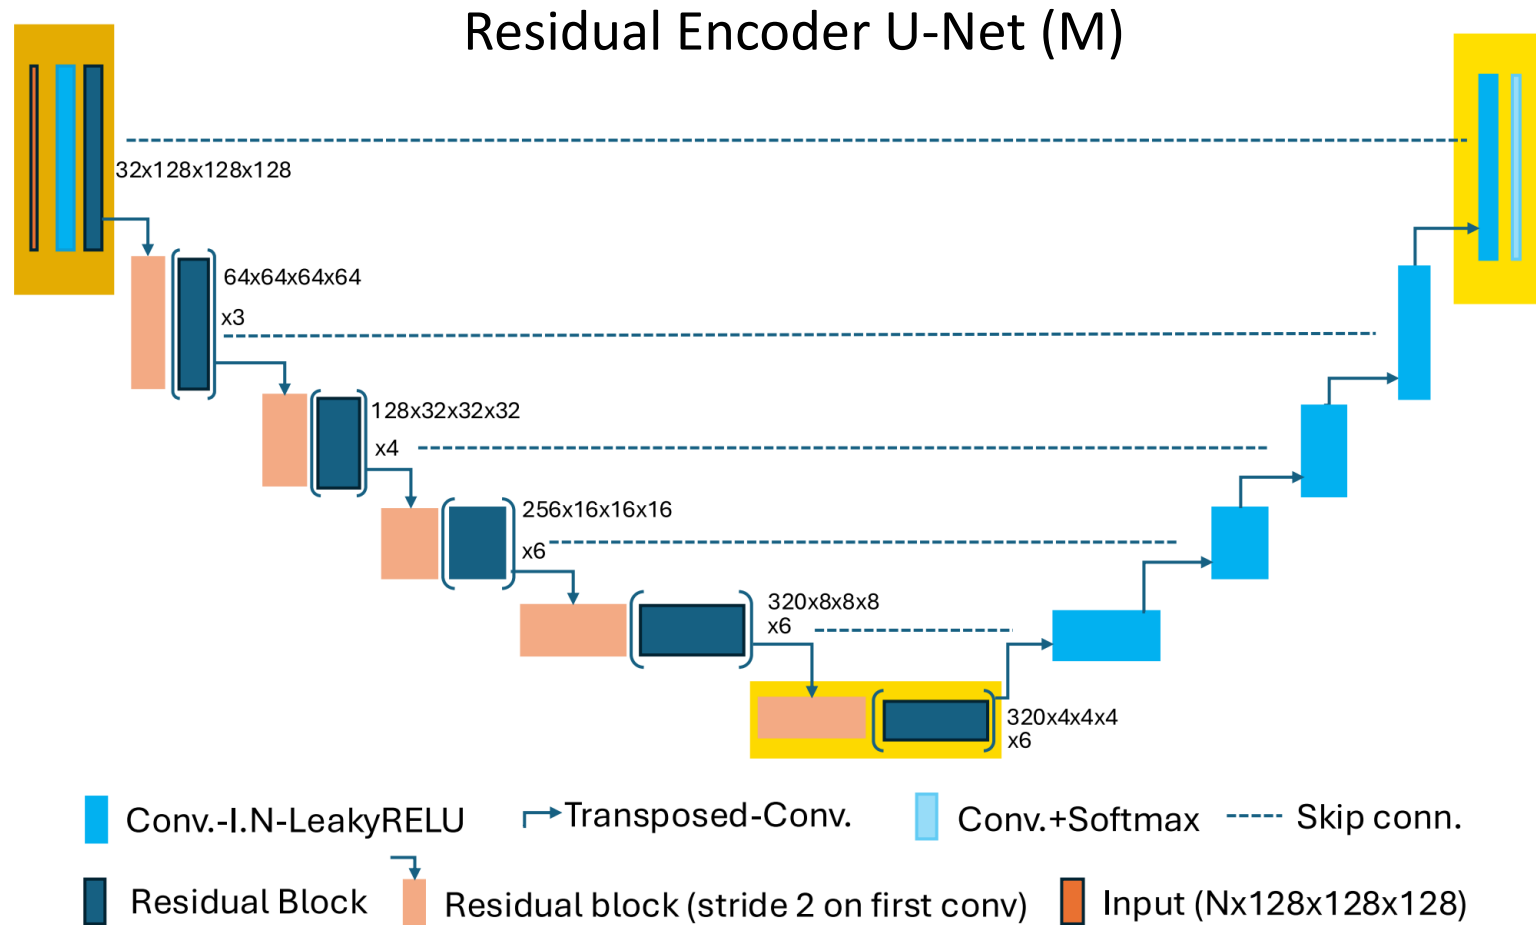
Comes in multiple flavours : 2D, 3D, (cascade 3D),

And multiple sizes : M-XL

	BTCV n=30	ACDC n=200	LiTS n=131	BraTS n=1251	KiTS n=489	AMOS n=360	VRAM [GB]	RT [h]
nnU-Net (org.) [21]	83.08	91.54	80.09	91.24	86.04	88.64	7.70	9
nnU-Net ResEnc M	83.31	91.99	80.75	91.26	86.79	88.77	9.10	12
nnU-Net ResEnc L	83.35	91.69	81.60	91.13	88.17	89.41	22.70	35
nnU-Net ResEnc XL	83.28	91.48	81.19	91.18	88.67	89.68	36.60	66

Runtime measured in GPU hours on A100 40GB PCIe

Base architecture



The notation $() \times n$ indicates repeated blocks.

Experiment Design

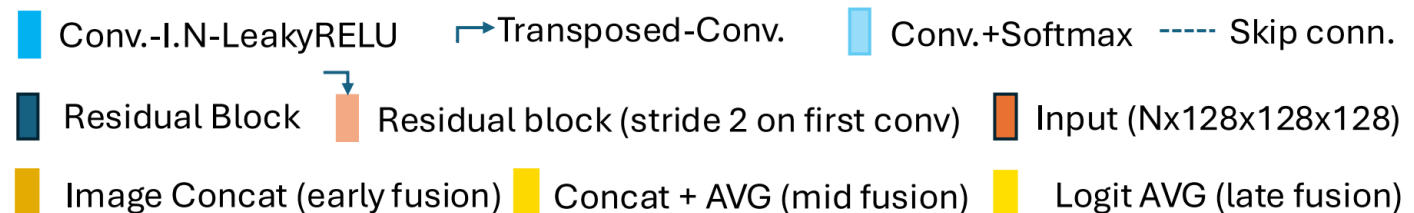
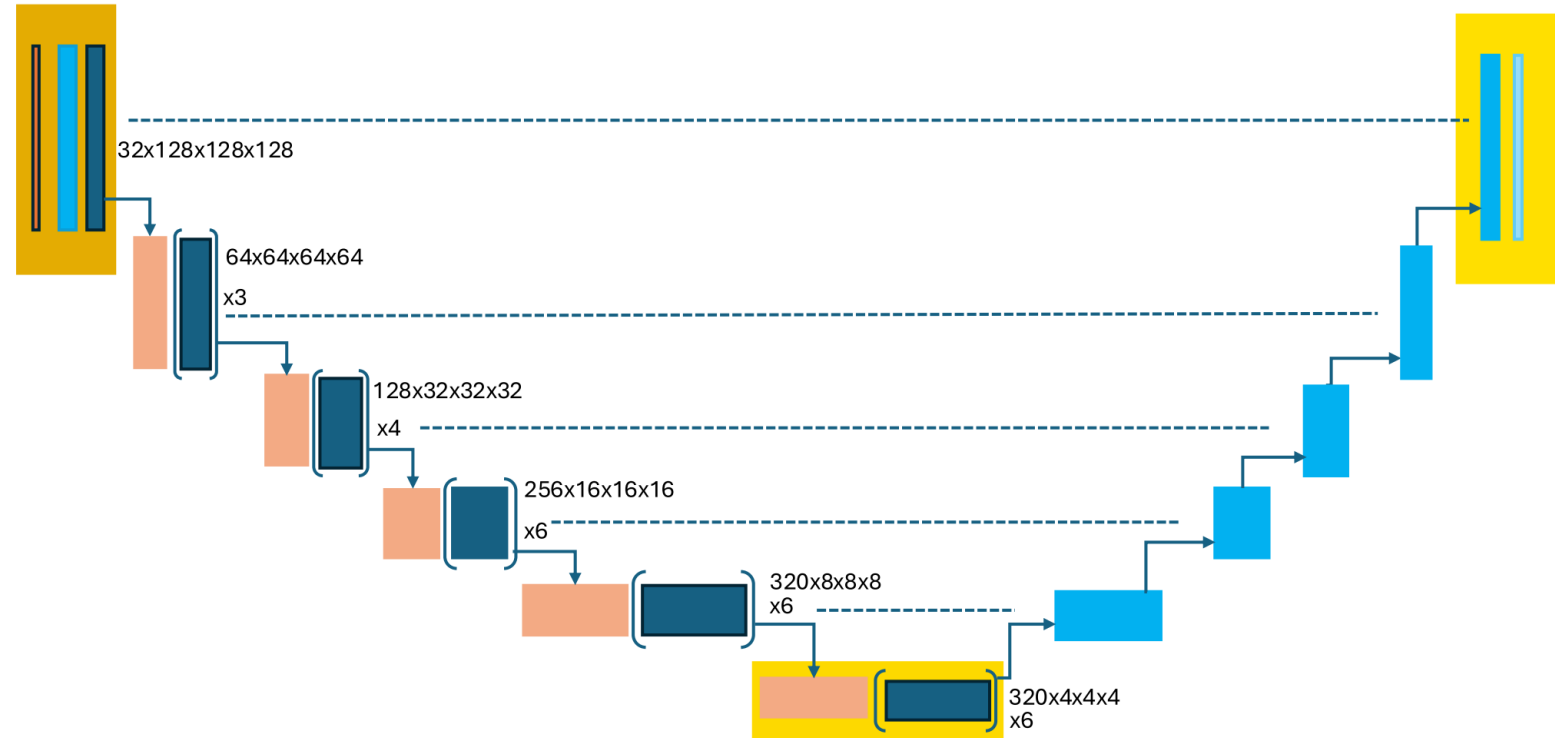
1. Preprocessing
2. Pre-training on ATLAS (n=651, T1-w)*
3. Fine-tuning on an internal bimodal dataset (T1-w, FLAIR; n=52)*.
 1. Fusion at **first layer**
 2. Fusion at **bottleneck**
 3. Fusion at **last layer**

* : 5-fold cross validation

Base architecture

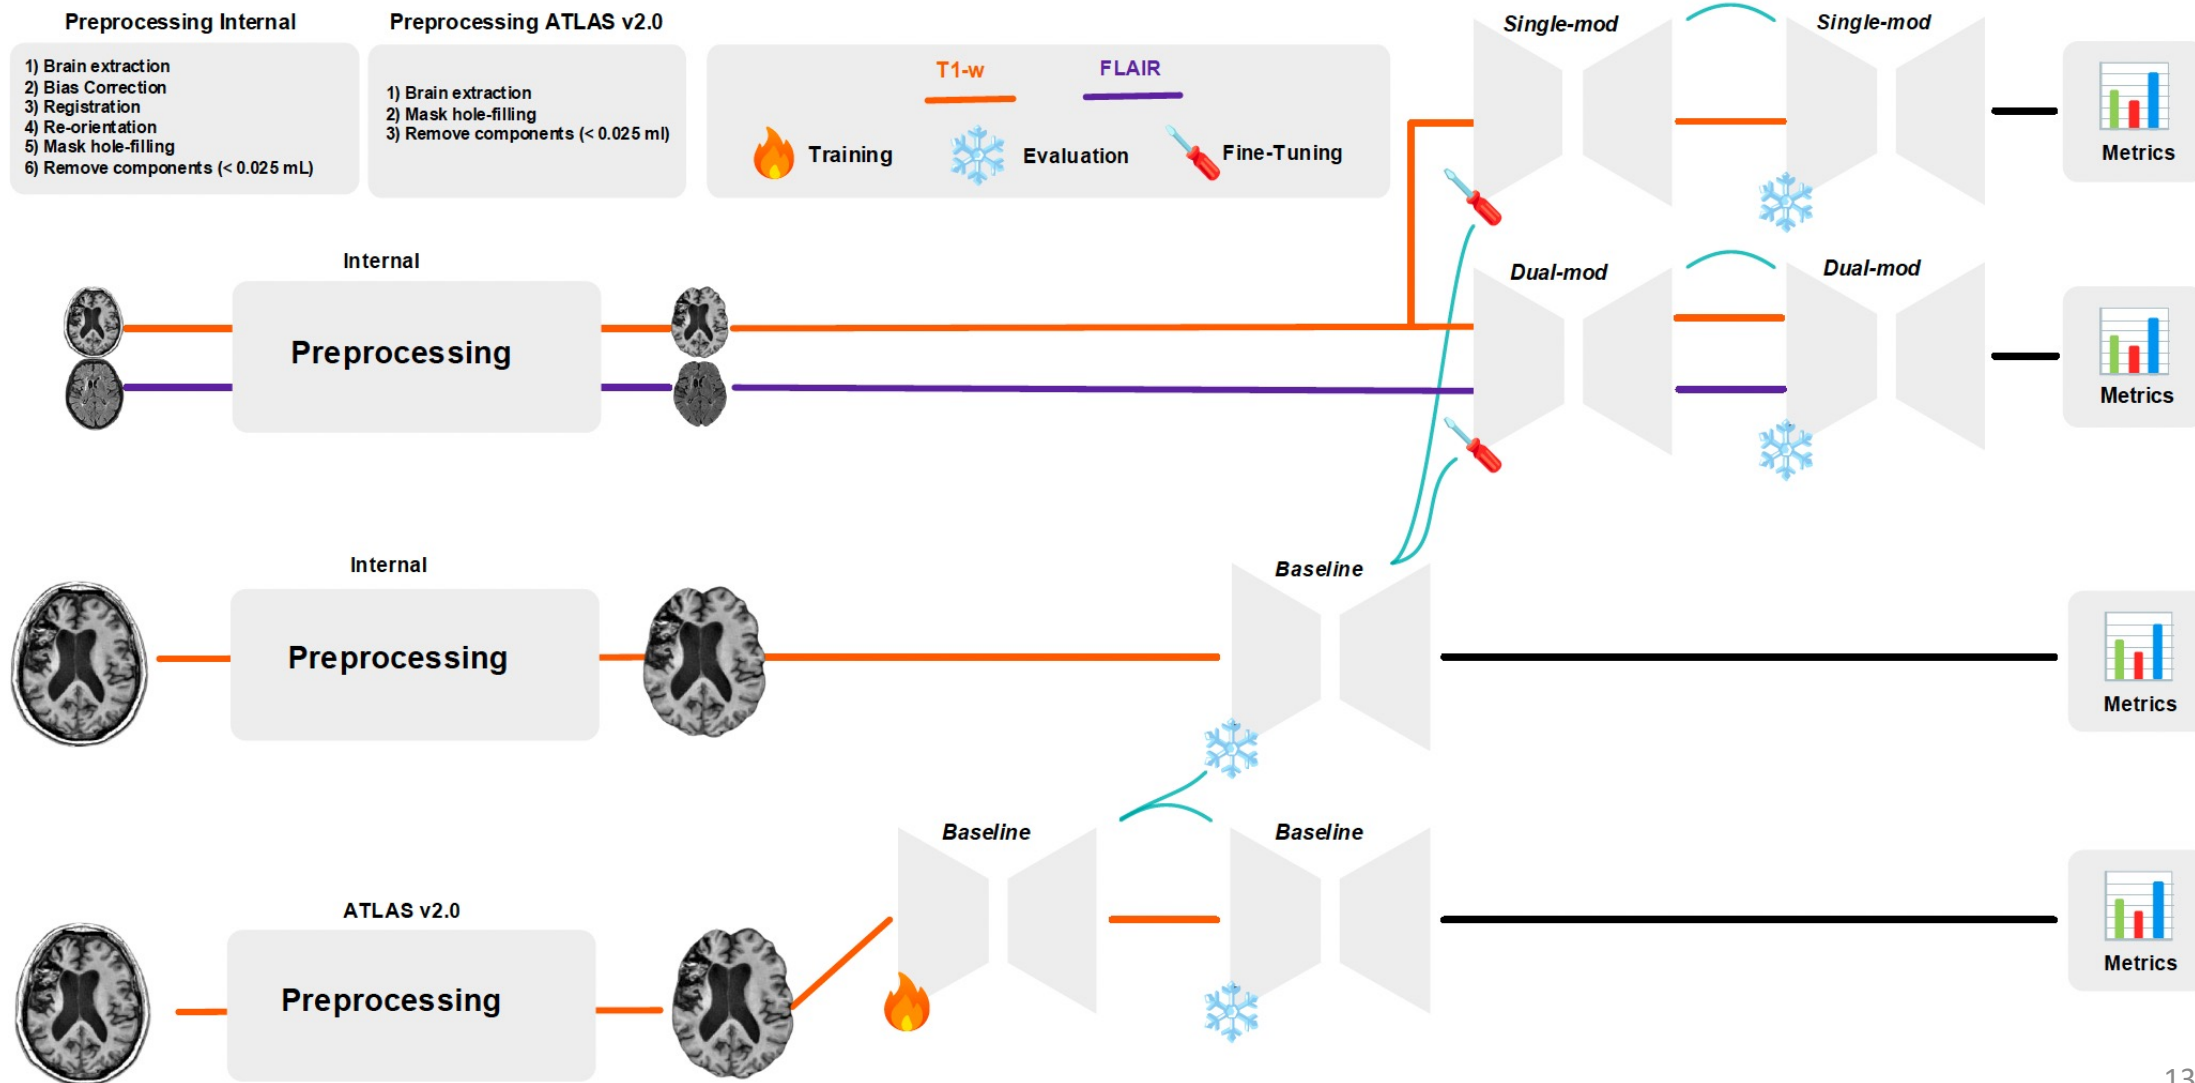
Residual Encoder U-Net

Colored blocks highlight the type of fusion strategy applied (early, mid, or late) with the T1w+FLAIR models



Di Matteo, A., Mahé, Y., Leplaideur, S., Bonan, I., Bannier, E., & Galassi, F. (2026). Deep learning and multi-modal MRI for the segmentation of sub-acute and chronic stroke lesions. Pattern Recognition Letters, 199, 225–231.

Experiment Design



Results

Baseline: model trained on ATLAS v2.0 (5-fold CV): Dice = 0.64, F1 = 0.67 (comparable to recent literature).

Dice = overlap; lesion F1 = lesion-level detection; ASD = average surface distance (mm).

Ensembling: voxel-wise mean of predictions from cross-validated models (ensemble).

Model	Dice \uparrow	Lesion F1 \uparrow	ASD (mm) \downarrow
Baseline model (T1)	0.68	0.73	6.95
Baseline + ensembling	0.70 ^{***}	0.79 ^{***}	3.47
Fine-tuning on clinical data (T1)	0.71 [*]	0.76 [*]	7.19
Fine-tuning (T1) + ensembling	0.72 ^{***}	0.78 ^{**}	3.37
T1+FLAIR early fusion + ensembling	0.74 ^{***}	0.79 [*]	3.14
T1+FLAIR mid-fusion + ensembling	0.69	0.77 [*]	7.68
T1+FLAIR late fusion + ensembling	0.75^{***}	0.80^{**}	2.94


Paired Wilcoxon test (two-sided), vs **baseline model**: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Results


Ground Truth

T1-w

T1-w + FLAIR



Di Matteo, A., Mahé, Y., Leplaideur, S., Bonan, I., Bannier, E., & Galassi, F. (2026). Deep learning and multi-modal MRI for the segmentation of sub-acute and chronic stroke lesions. *Pattern Recognition Letters*, 199, 225–231.



Take Home Message #1

- **FLAIR** + late fusion + ensembling yields the best average performance.
- The largest benefit of **FLAIR** is for **small lesions**, which are hardest to segment reliably.
- For **medium** and **large** lesions, gains are smaller: **most strategies converge**.

What we do at Empenn

1. Multimodal fusion for fully supervised stroke lesion segmentation
- 2. From the lab to the clinic, transferring models in a clinical research context**
3. Unsupervised learning, detecting anomalies beyond lesions to better predict recovery pathways


How to use ?

What is a terminal ?



1. Open a **terminal**
2. If you have a **GPU** : install CUDA

2 hours to install a driver ?!



1. Toolkit
2. Driver
3. Library

```
$ sudo dnf config-manager addrepo --from-repofile https://developer.nvidia.com/compute/cuda/repos/fedora42/x86_64/cuda-fedora42.repo  
$ sudo dnf clean all  
$ sudo dnf -y install cuda-toolkit-13-0
```

```
$ sudo dnf -y install nvidia-open
```

```
pip install nvidia-cudnn
```

3. Create a **virtual environment**

4. Install **Pytorch**

```
pip3 install torch torchvision --index-url https://download.pytorch.org/g/whl/cu130
```

5. Install **nnUnet**

```
pip install nnunetv2
```

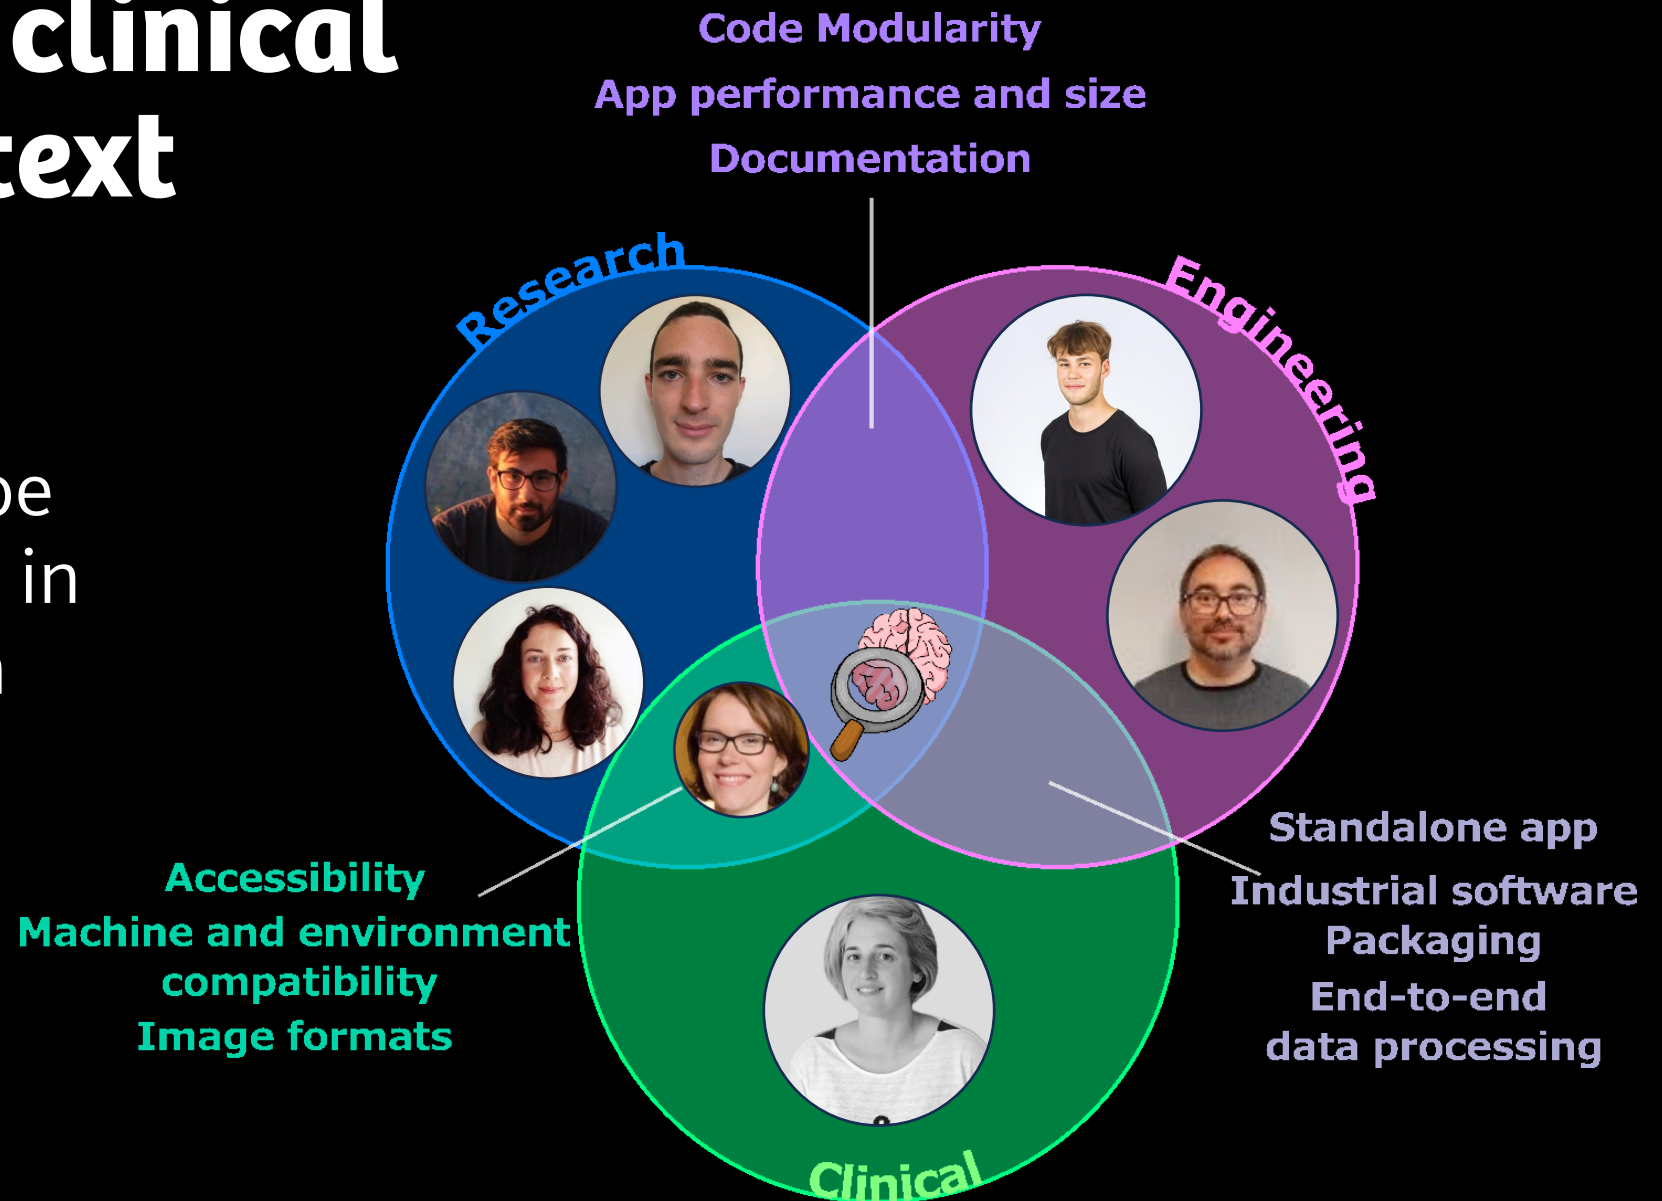
6. **Download** the model

7. **Preprocess** your data (Anima, FSL...)

8. Run the **inference**

Transfer to a clinical research context

How can this tool be made **easier to use** in a clinical research context?



Transfer to a clinical research context

Rationale



Clinical End User



Computer Vision End User



Software Engineering Good Practices

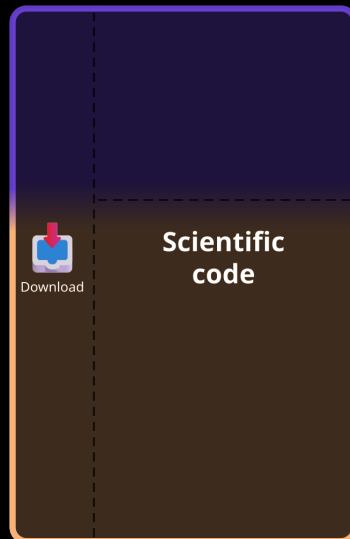
Dependencies



In-house



External



Transfer to a clinical research context

Rationale



Clinical End User



Computer Vision End User



Software Engineering Good Practices

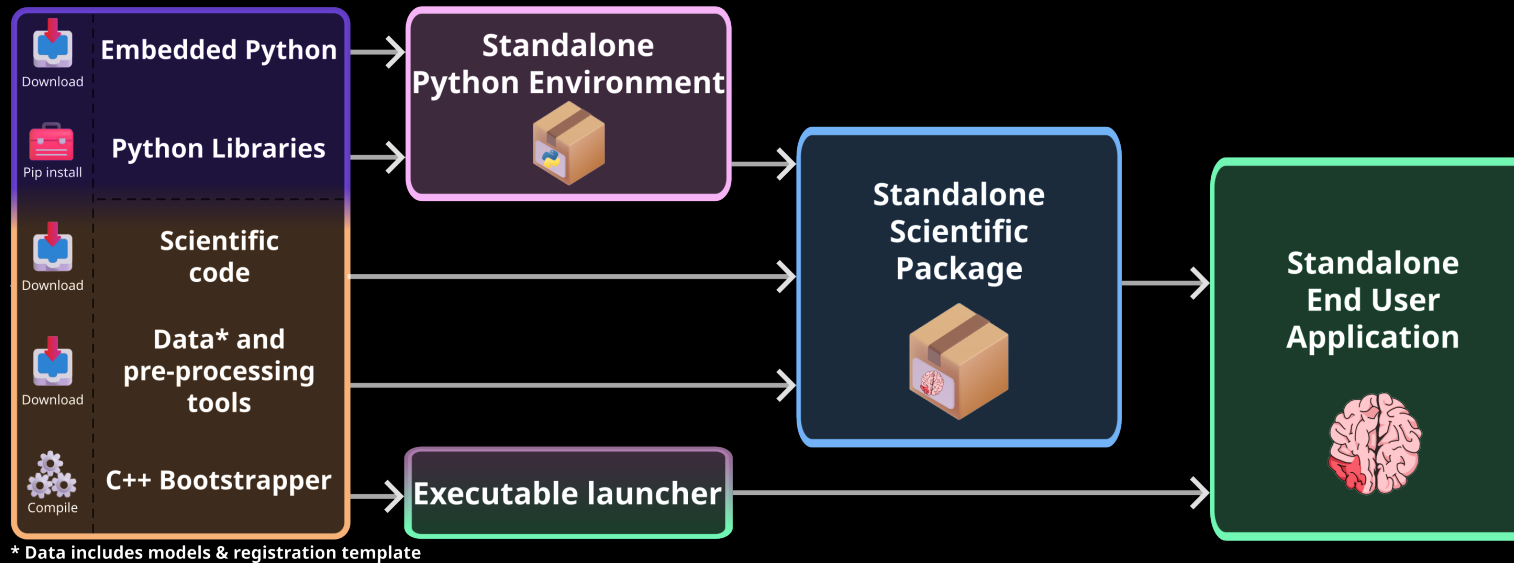
Dependencies



In-house



External



Transfer to a clinical research context

Rationale



Clinical End User



Computer Vision End User



Software Engineering Good Practices

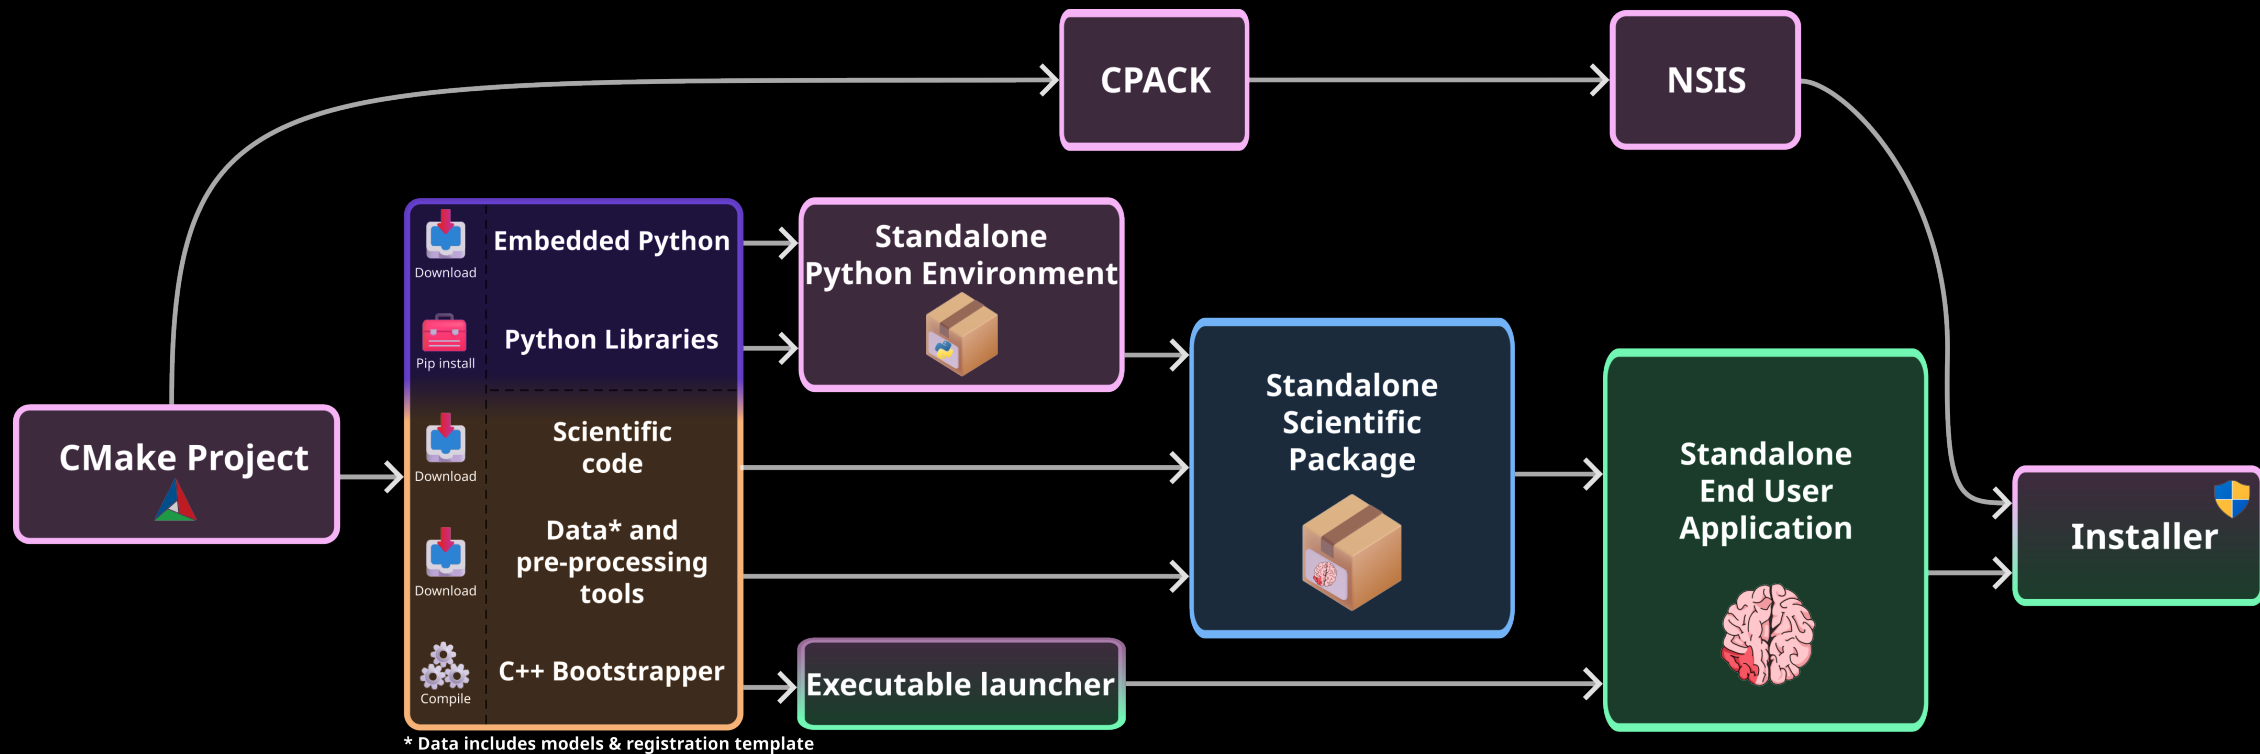
Dependencies



In-house



External



Transfer to a clinical research context

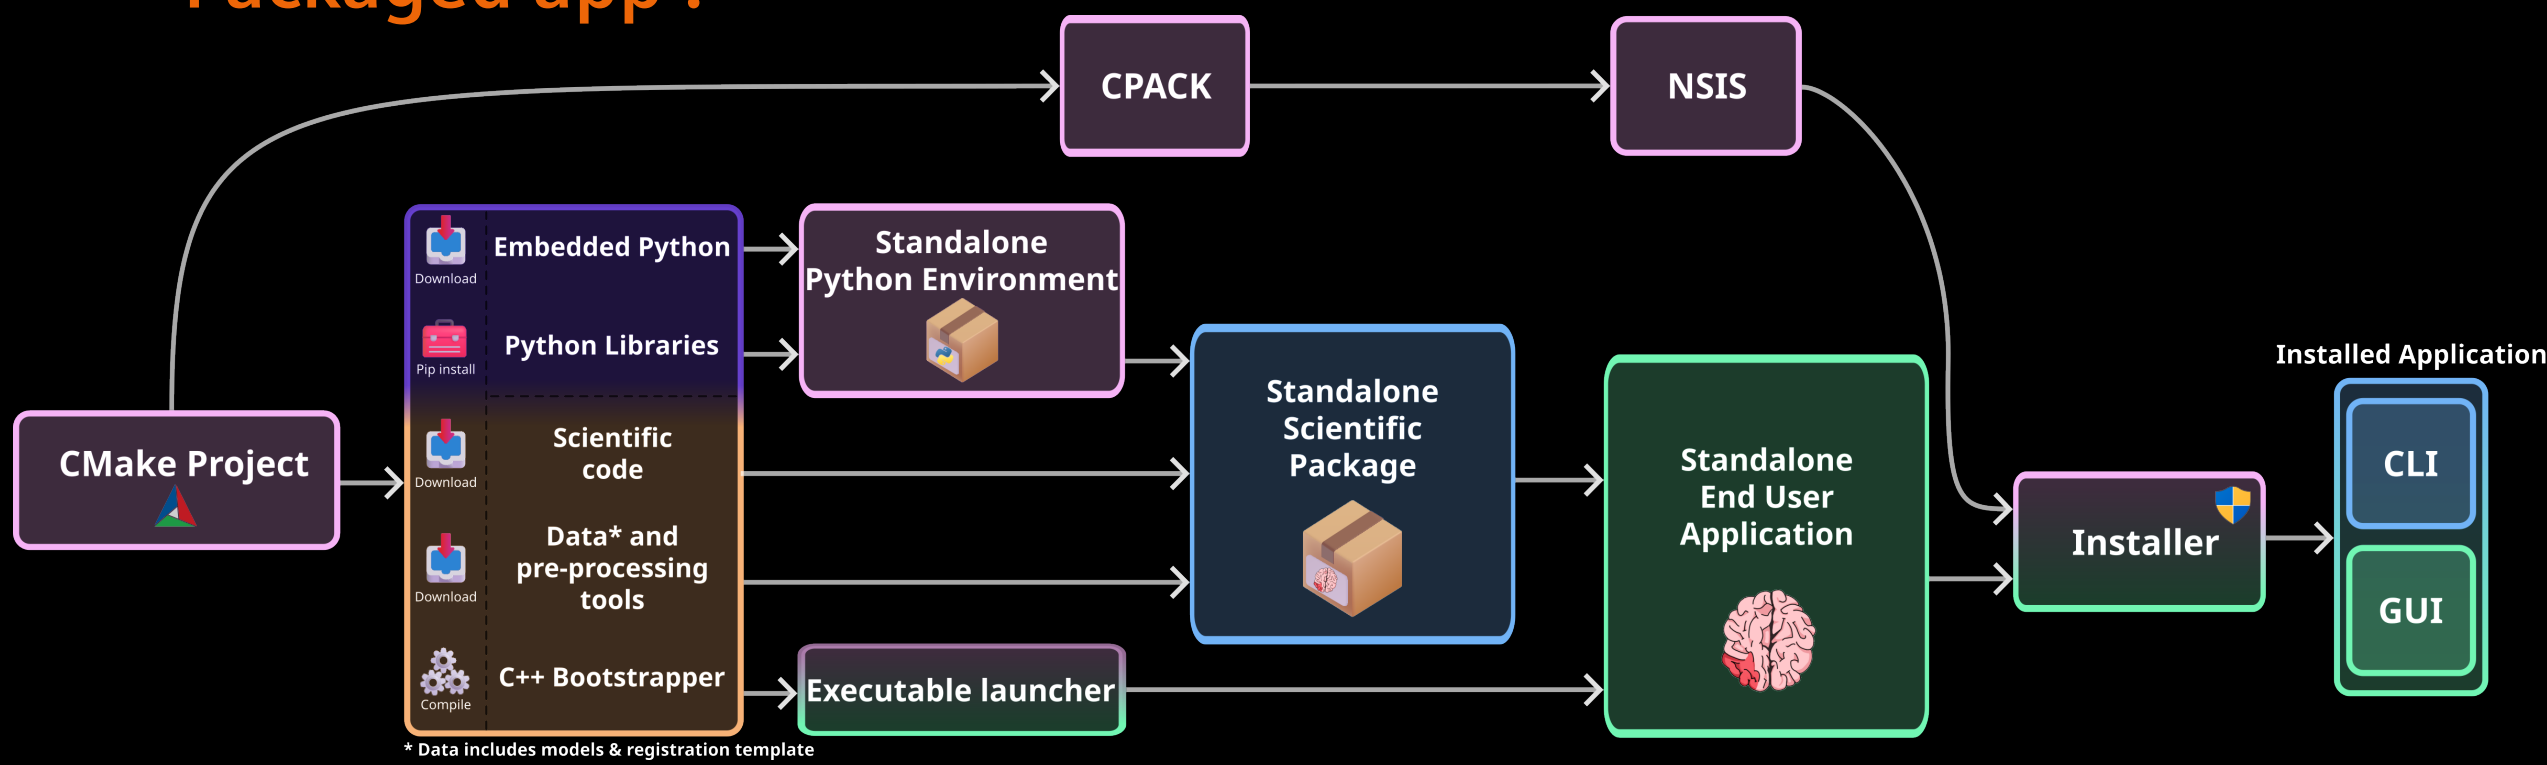
Rationale

- █ Clinical End User
- █ Computer Vision End User
- █ Software Engineering Good Practices

Dependencies

- █ In-house
- █ External

Packaged app :



We still face limitations

At this time, we already had replaced the overly heavy nnU-net framework with pyTorch.
With nnU-net, the global size was more than 10GB :

OS (Hardware)	Size
Fedora (GPU+CPU)	6.20 GB
Windows (GPU+CPU)	7.25 GB
Windows (CPU only)	2.10 GB

**Size
(GB)**

Task	Ultrabook laptop	Workstation laptop
Brain Extraction	198.5 s/volume	36.6 s/volume
Registration	15.1 s/volume	3.1 s/volume
Inference	149.2 s/volume	0.9 s/volume

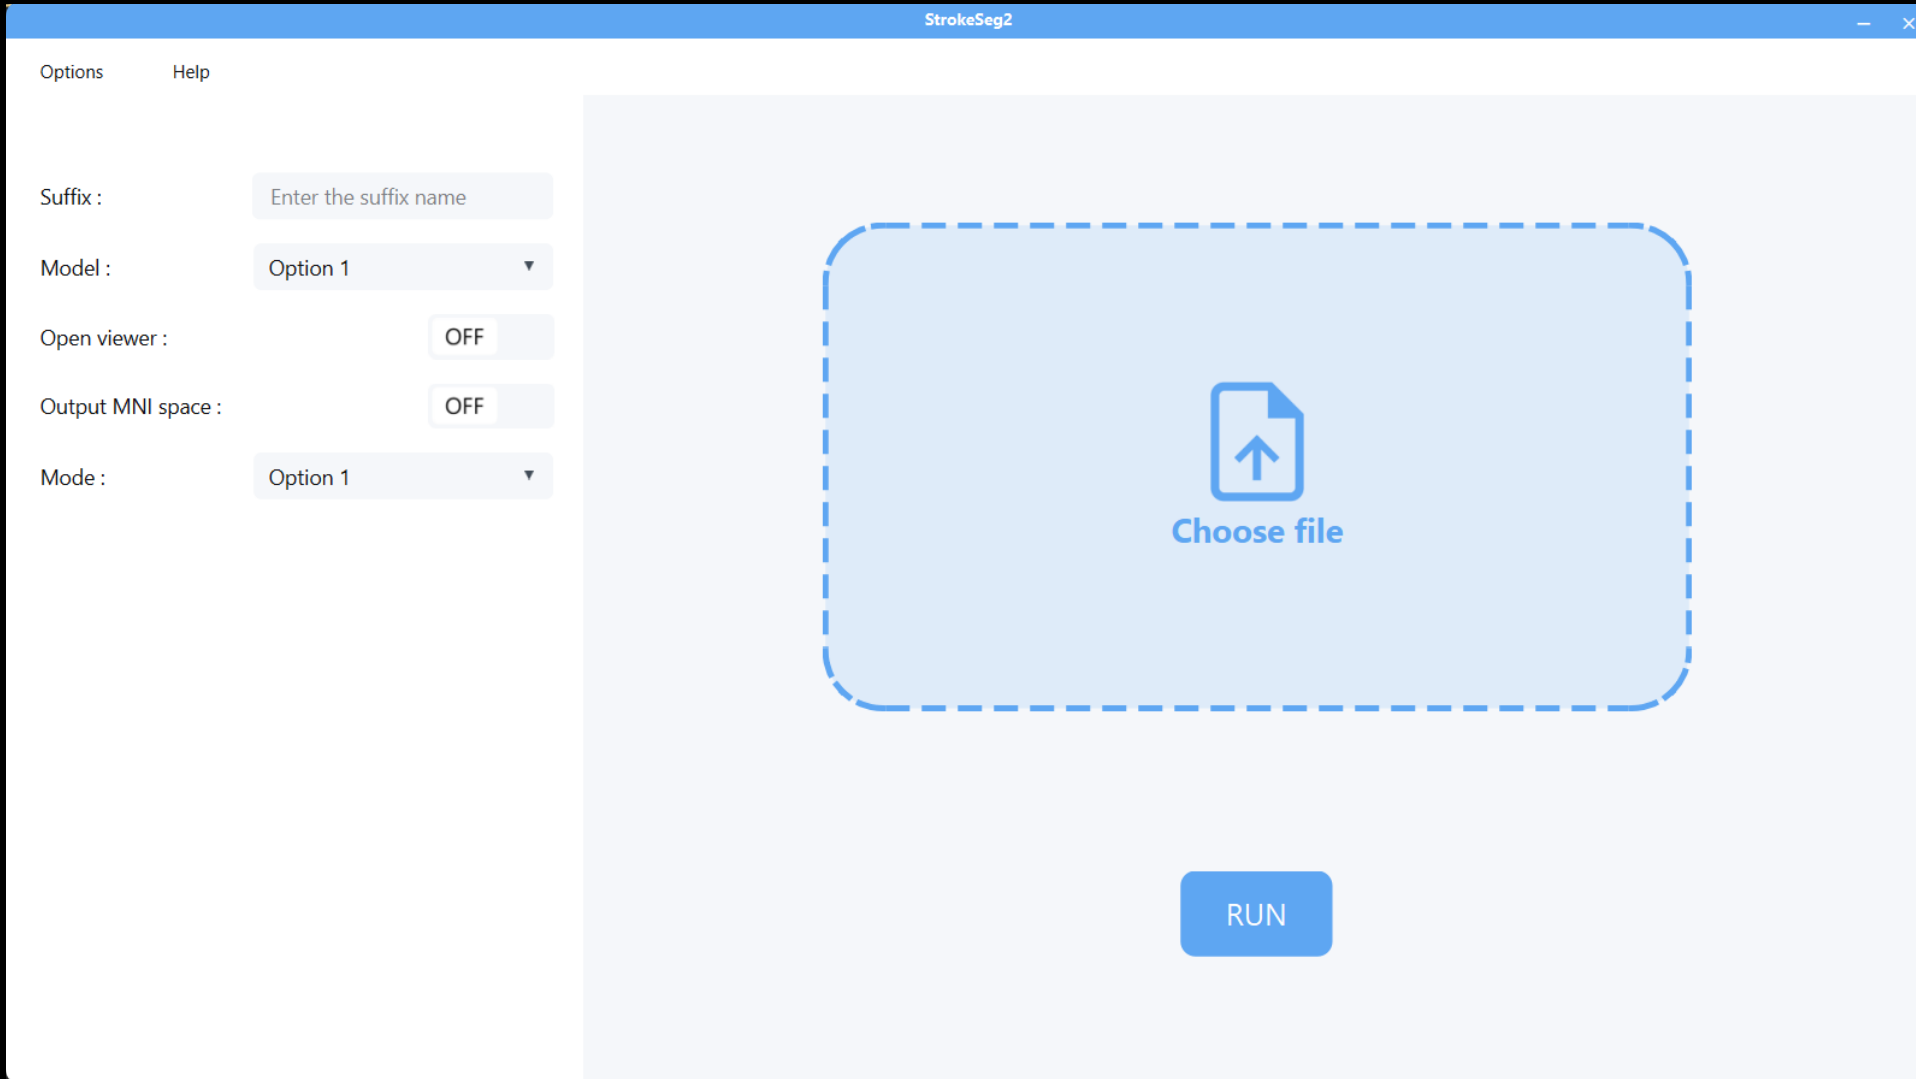
**Speed
(seconds/volume)**

We still face limitations

Changing the Deep-Learning backend **halves** app size

OS (Hardware)	PyTorch	ONNX	Reduction
Fedora (GPU+CPU)	6.20 GB	2.50 GB	60%
Windows (GPU+CPU)	7.25 GB	2.73 GB	62%
Windows (CPU only)	2.10 GB	0.80 GB	62%

Packaged application size



With **ESIR**
students :

Axel P.
Asma Z.
Damien V.
Eliot H.


What now ?

Engineering industrial project aiming to :

- Rationalize technologies used to raise the project
 - Replace the **python engine** with a **C/C++ engine**
 - Use agnostic hardware abstraction layer w/ Windows ML
- Add **missing features** (DICOM)
- Improve the **Graphical Interface** (Qt)
- Siemens Open Recon Integration

Internship :

Axel P.





What we do at Empenn

1. Multimodal fusion for fully supervised stroke lesion segmentation
2. From the lab to the clinic, transferring models in a clinical research context
- 3. Unsupervised learning, detecting anomalies beyond lesions to better predict recovery pathways**


My PhD project

- Explore methods for detecting and segmenting post-stroke brain abnormalities
- Anomalies can be: The **focal lesion** but also **cortical atrophy**, **enlargement of the ventricles**.
- First, exploration of **unsupervised methods** to compensate for the lack of annotated data.



The **unsupervised** segmentation paradigm

💡 This is not only about the lesion



State of the art: UAD in brain imaging

Condition	Family	Dice	AUPRC	Representative refs.
Tumors BraTS (1-4 seq.) ; in-house	AE	0.390– 0.773	0.300– 0.833	Lu'24 ; Kascenas'23
	VAE	0.301– 0.650	0.279– 0.819	Lambert'21 ; Lüth'23
	GAN	0.630– 0.770	–	Nguyen'21
	Diffusion / Flow	0.371– 0.851	0.417– 0.590	Beizae'26 ; Iqbal'23
SEP MSSEG ; MSLUB (1-4 seq.)	AE	0.173– 0.605	0.203	Baur'18 ; Ghorbel'23
	VAE	0.133– 0.378	0.272	Pinaya'22
	GAN	0.482	–	Dey'21
	Diffusion / Flow	0.055– 0.247	0.067– 0.106	Behrendt'23 ; Pinaya'22
WMH WMH 2017 (3 seq.)	VAE	0.133– 0.463	0.320	Lambert'21 ; Pinaya'22
	Diffusion / Flow	0.132– 0.569	–	Fontanella'24
AVC (lésion) ATLAS ; ISLES (1-2 seq.)	AE	–	0.705	Qu'26
	VAE	0.331	0.256– 0.549	Bengs'21 ; Lüth'23
	Diffusion / Flow	0.148– 0.416	0.145	Beizae'26 ; Bercea'23

✨ Mahé, Y., Bannier, E., Leplaideur, S., Fromont, E., & Galassi, F. (2025). Unsupervised deep generative models for ✨ anomaly detection in neuroimaging: A systematic scoping review. *arXiv preprint arXiv:2510.14462*.



Take Home Message #2

- UAD has been studied in **tumors, multiple sclerosis, white-matter hyperintensities**, and **stroke**.
- Best results are generally obtained for **large focal anomalies**.
- Performance drops for **subtle, small, or diffuse** abnormalities.
- **Stroke** remains relatively **under-explored**, especially beyond the focal lesion.
- There is no **standard evaluation method**.

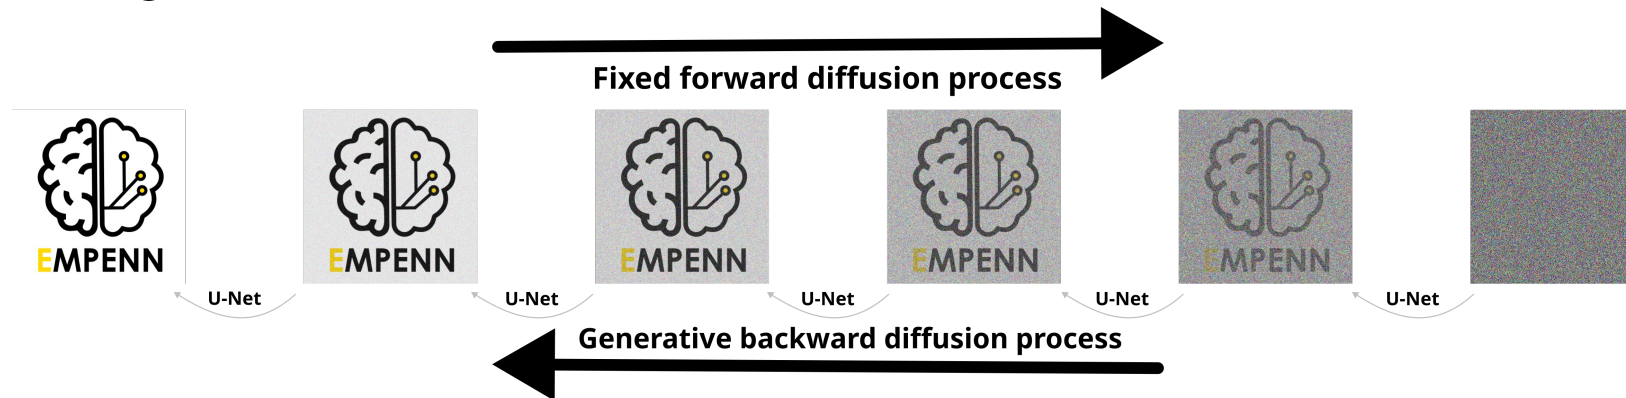
State of the art: UAD in brain imaging

Condition	Family	Dice	AUPRC	Representative refs.
Tumors BraTS (1-4 seq.) ; in-house	AE	0.390– 0.773	0.300– 0.833	Lu'24 ; Kascenas'23
	VAE	0.301– 0.650	0.279– 0.819	Lambert'21 ; Lüth'23
	GAN	0.630– 0.770	–	Nguyen'21
	Diffusion / Flow	0.371–0.851	0.417–0.590	Beizae'26 ; Iqbal'23
SEP MSSEG ; MSLUB (1-4 seq.)	AE	0.173– 0.605	0.203	Baur'18 ; Ghorbel'23
	VAE	0.133– 0.378	0.272	Pinaya'22
	GAN	0.482	–	Dey'21
	Diffusion / Flow	0.055– 0.247	0.067– 0.106	Behrendt'23 ; Pinaya'22
WMH WMH 2017 (3 seq.)	VAE	0.133– 0.463	0.320	Lambert'21 ; Pinaya'22
	Diffusion / Flow	0.132–0.569	–	Fontanella'24
AVC (lésion) ATLAS ; ISLES (1-2 seq.)	AE	–	0.705	Qu'26
	VAE	0.331	0.256– 0.549	Bengs'21 ; Lüth'23
	Diffusion / Flow	0.148–0.416	0.145	Beizae'26 ; Bercea'23

✨ Mahé, Y., Bannier, E., Leplaideur, S., Fromont, E., & Galassi, F. (2025). Unsupervised deep generative models for anomaly detection in neuroimaging: A systematic scoping review. *arXiv preprint arXiv:2510.14462*. ✨

Introduction to Flow-Matching

- Initially diffusion models where denoising machines
 - Denoising Diffusion Probabilistic Models (DDPM)



- Flow-Matching** is introduced by **Albergo et al.**, **Liu et al.** and **Lipman et al.** (ICLR 2023)
 - Backbone of **Stable Diffusion v3**, can be described as a generalisation of the older diffusion models (DDPM, 2020)

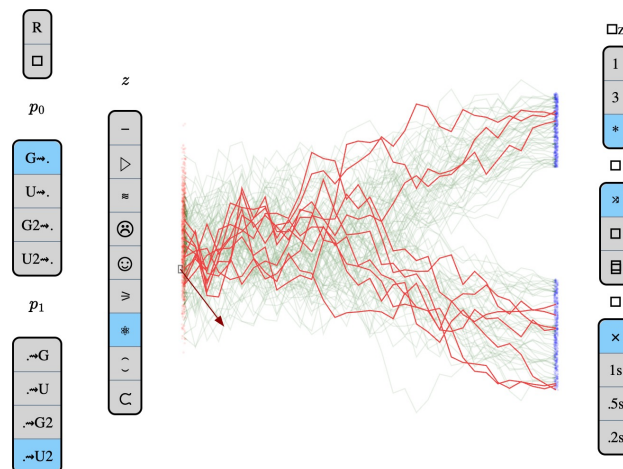
Introduction to Flow-Matching

	DDPM	Flow-Matching
Learning Objective	predict the noise	velocity field (between distribution)
Transformation Path	Stochastic Differential Equation (SDE)	Ordinary Differential Equation (ODE)
Efficiency	Requires typically 1000 U-Net calls	A few euler steps
Generalisability	Prone to overfitting	Prone to overfitting

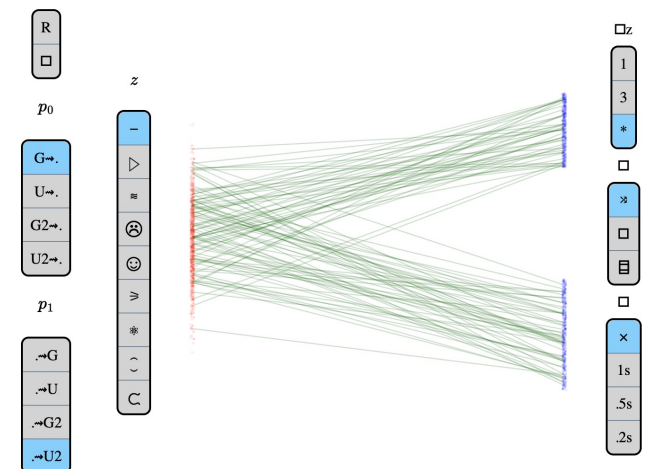
i Shout out to the team of R.Emonet (Lab. Hubert Curien) for this great blog post explaining Flow Matching in detail.

Gagneux, A., Martin, S., Emonet, R., Bertrand, Q., & Massias, M. (2025). A Visual Dive into Conditional Flow Matching. In *ICLR Blogposts 2025*.

CFM Playground




CFM Playground



Anomaly detection with flow-matching

UAD Method : Reflect

- Latent flow-matching for brain anomaly detection.



Beizae, J. (2025). Reflect: Rectified Flows for Efficient Brain Anomaly Correction Transport . In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*. Springer Nature Switzerland.

Anomaly detection with flow-matching


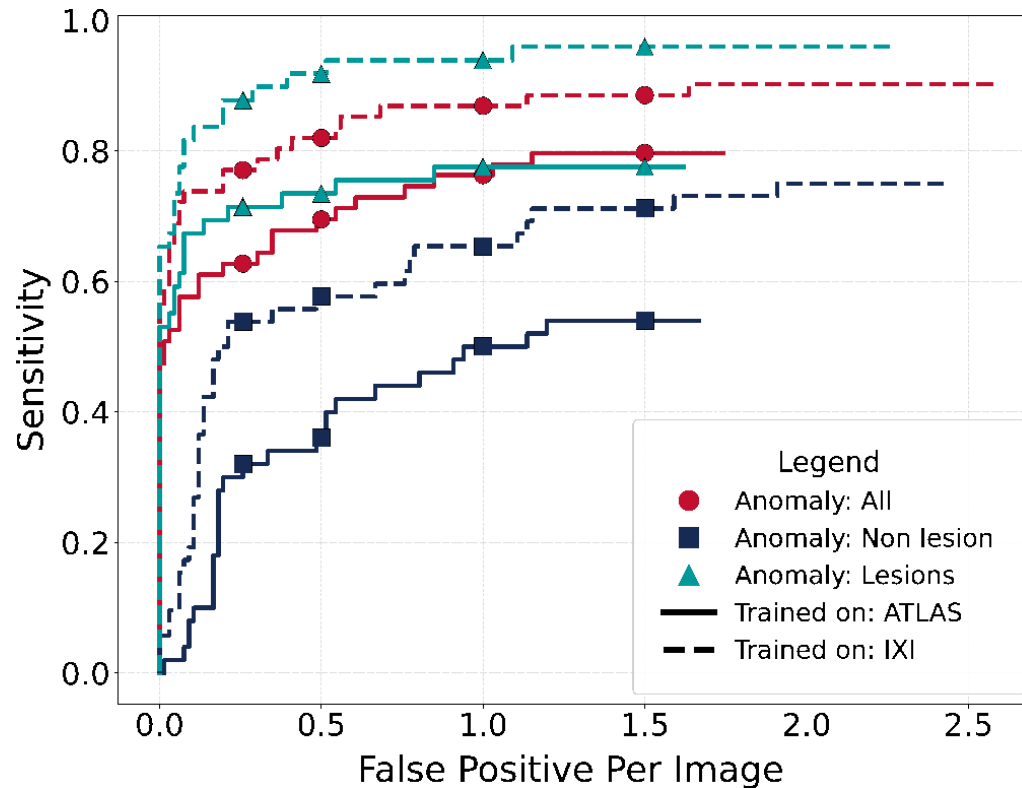
- Training with fully healthy data or non lesional slices of a pathological dataset
 - **Healthy:** The **IXI** dataset [1], consisting of 581 T1-w images of healthy participants.
 - **Pathological: ATLAS v2.0** [2] (T1-w, n=955) is a sub-acute and chronic stroke dataset with segmented **lesions** (n=655) and **microbleeds**. Lesion segmentation were traced by a single annotator and reviewed by two experts.
 - **Non-lesional abnormalities** independently annotated by two experts by point-click labels on a validation subset (n=66).

[1] IXI – Information eXtraction from Images (EPSRC GR/S21533/02), <https://brain-development.org/ixi-dataset/>

[2] Liew, S.L., Lo, B.P., Donnelly, M.R. *et al.* *Sci Data* **9**, 320 (2022), https://fcon_1000.projects.nitrc.org/indi/retro/atlas.html 42

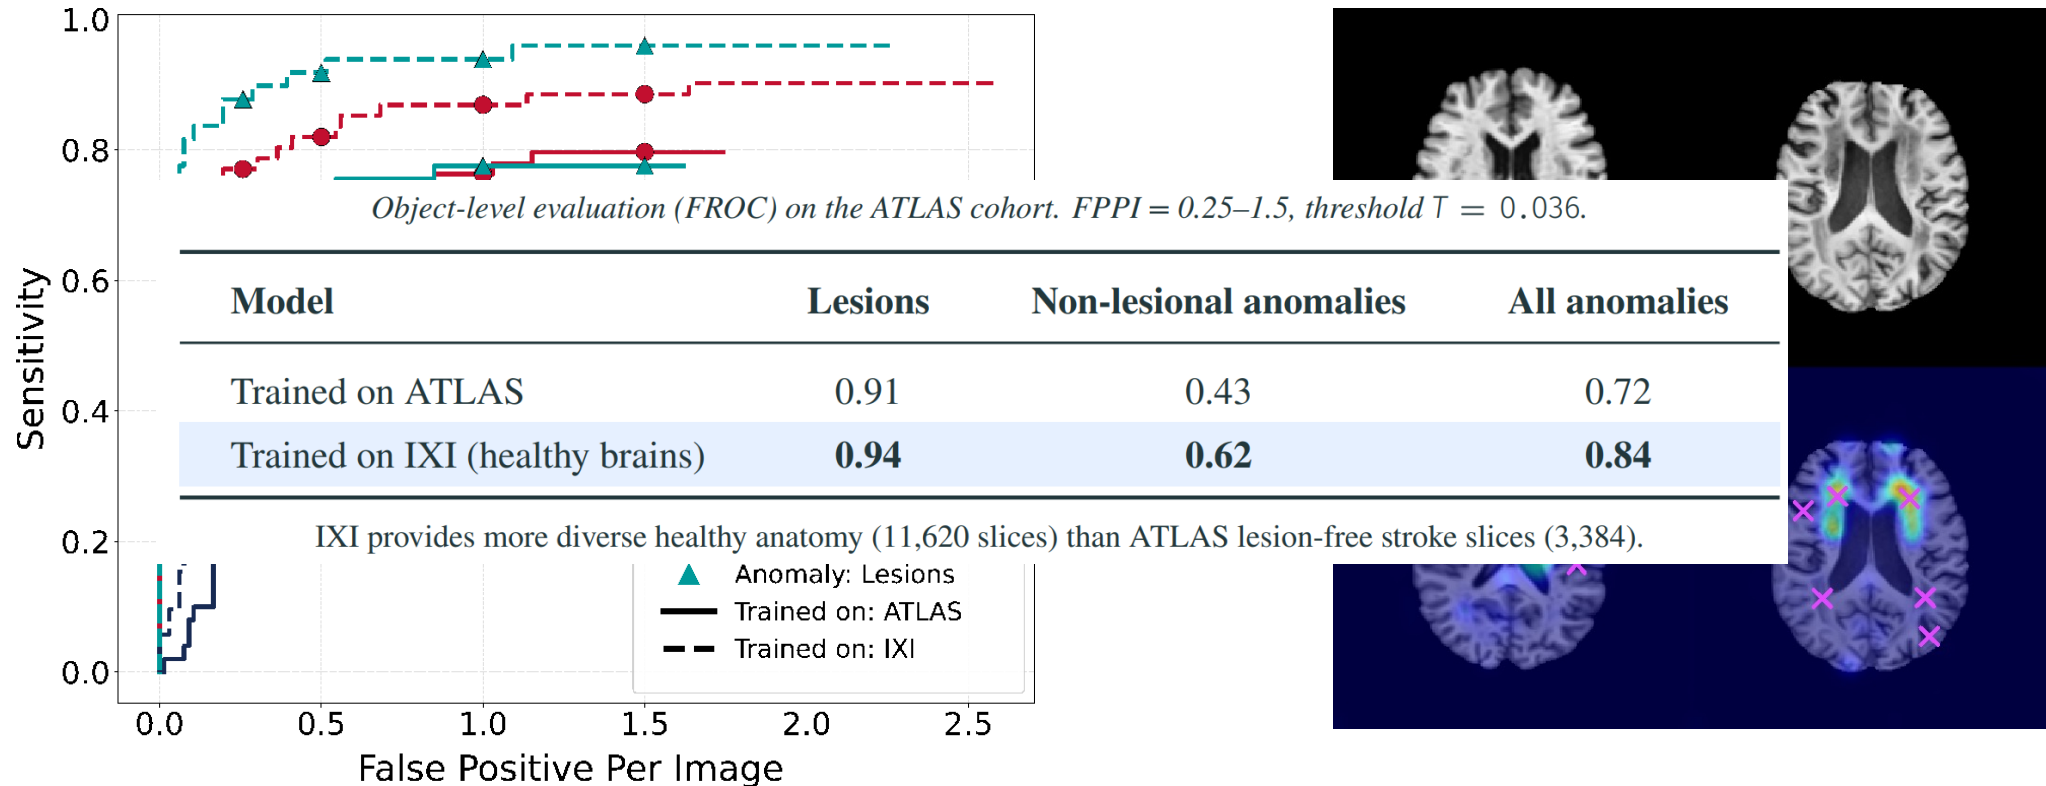
Anomaly detection with flow-matching

There is more than just the lesions




Anomaly detection with flow-matching

There is more than just the lesions



Anomaly detection with flow-matching

Unsupervised lesion delineation can benefit from a fully healthy training set



	Dice \uparrow	HD95 (mm) \downarrow	ASD (mm) \downarrow	F1 10% \uparrow
Trained on ATLAS				
All	0.270	22.4	4.18	0.326
Small	0.074	5.64	2.67	0.153
Medium	0.198	18.6	3.70	0.250
Large	0.600	28.1	4.799	0.640
Trained on IXI				
All	0.366*	31.3	10.9	0.444
Small	0.060	54.8	40.8	0.077
Medium	0.401*	25.2	6.50	0.558*
Large	0.609	32.6	7.4	0.603



Take Home Message #3

- **There's more than just the lesions**
- A methodology for quantifying the “**anomaly load**” is yet to be developed
- **Generative modelling** is a way to do it

What next ?

- **Foundation models** will be explored as pre-trained backbones to improve anomaly detection and segmentation with limited annotated stroke datasets
- Exploring foundation models :
 - **Pre-trained models** for anatomical representation learning
 - **Reduce dependence** on annotated datasets
 - **Improve anomaly detection** beyond lesions



PhD mobility project,
Boston (USA)


The stroke team



Stéphanie Leplaideur
Medical Doctor



Elise Bannier
MR Physicist



Simon Butet
Medical Doctor



Isabelle Bonan
Head of Rehabilitation
Medicine, CHU Rennes



Alessandro Di Matteo
PhD student



Youwan Mahé
PhD student



Francesca Galassi
Associate Prof.



Elisa Fromont
Full Prof.



Yann Kerverdo
Engineering student




Florent Leray
Engineer



Come see us at ISBI 2026

Anomaly Detection
4-page paper (Oral)



StrokeSeg Software
(poster)




IABM 2026, Lyon



Thank you
Do you have any questions?

Inria


IRISA


EMPENN

Malt

CHU
Rennes

SIEMENS
Healthineers 