



HAL
open science

The multi-dimensional aspect of uncertainty

Michel Dojat

► **To cite this version:**

Michel Dojat. The multi-dimensional aspect of uncertainty. Miccai, Oct 2024, Marrachech, Morocco. <inserm-04761437>

HAL Id: inserm-04761437

<https://inserm.hal.science/inserm-04761437v1>

Submitted on 31 Oct 2024

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-NC 4.0 - Attribution - Non-commercial use - International License

The multi-dimensional aspect of uncertainty



Michel DOJAT

Research Director Inserm
Deputy Scientific Director Inria

michel.dojat@inria.fr

I declare relationships with the company Pixyl
(pixyl.ai)

My main collaborators




Benjamin Lambert



Florence Forbes

RESEARCH TOPICS AND EXPERIMENTAL APPROACHES



/ Fundamental neurosciences

Cytoskeleton, Intracellular traffic, Synaptic plasticity, Mechanisms studied in normal and pathological conditions (neurobiological diseases, neurodegenerative diseases, myopathies)








/ Pre-clinical and clinical research

*Developing tools and concepts
Close links with networks such as GREEN, Neuropynov, NeuroCoG...*


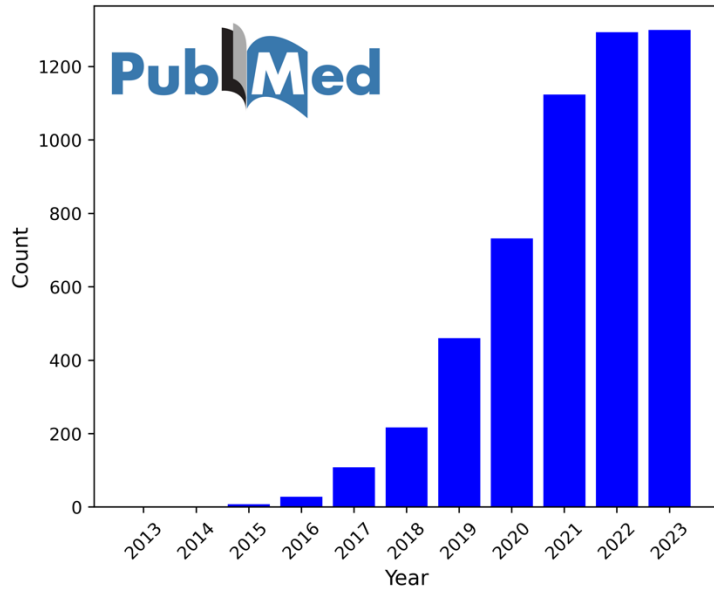


/ Innovative technologies and treatments

Multidisciplinary approaches including human social sciences, and methodological developments (optogenetics, reconstruction of neural networks, electrophysiology...)

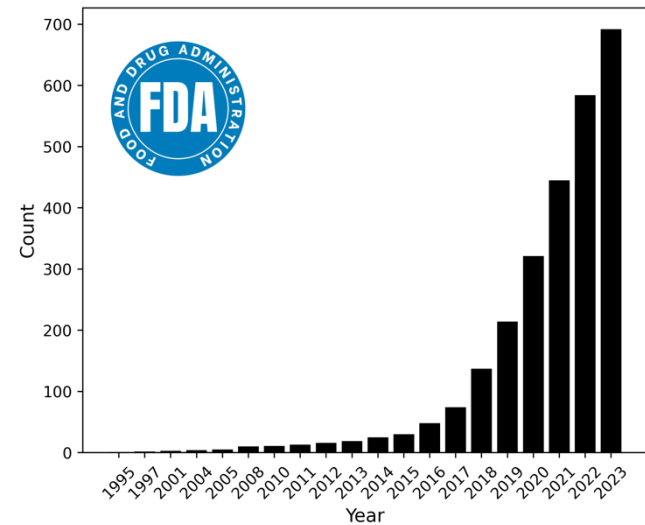

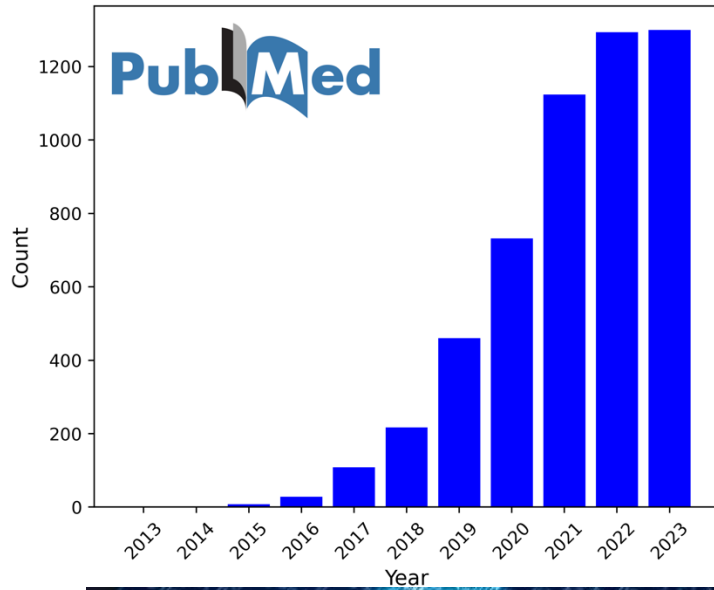


Automatic analysis of medical images




Kw: « Deep Learning » and « Medical Image Analysis »

Automatic analysis of medical images




The curse of the Black-Box

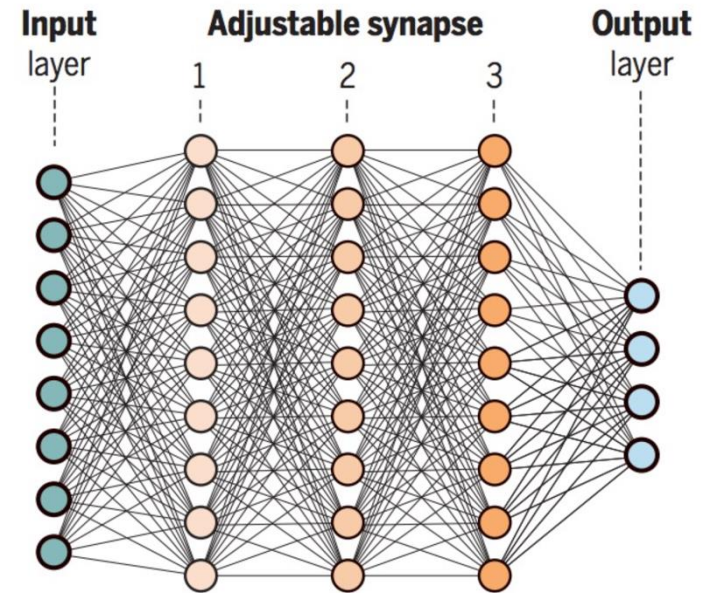
The curse of the Black-Box




[Shapson-Coe et al.
Science 384, 635 (2024)]



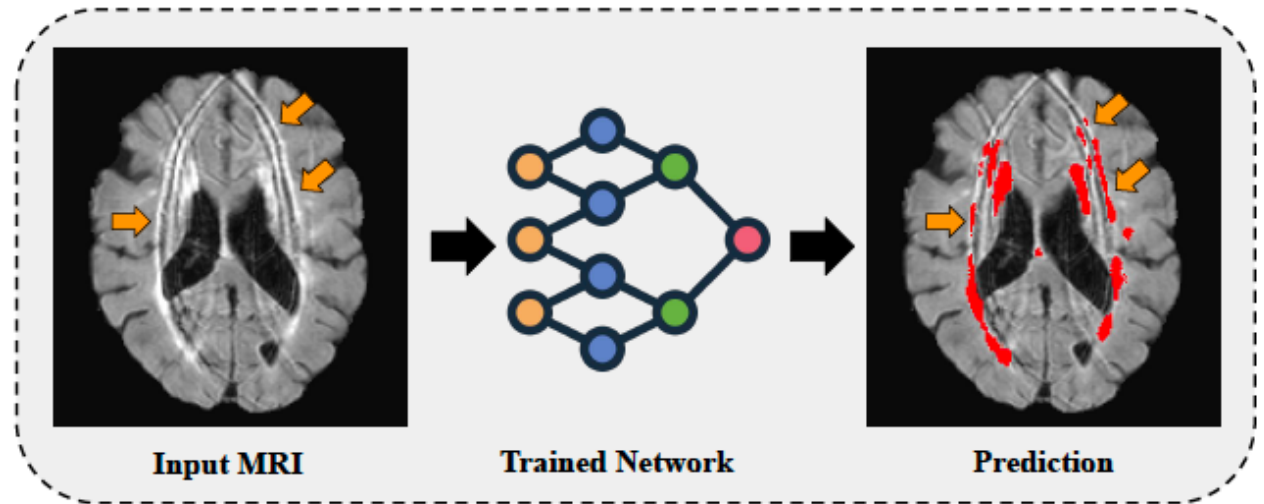
[Crumb, Says 1967 Num 1]



The curse of the Black-Box




Standard Deep Learning approach



End-user

The curse of the Black-Box




End-user


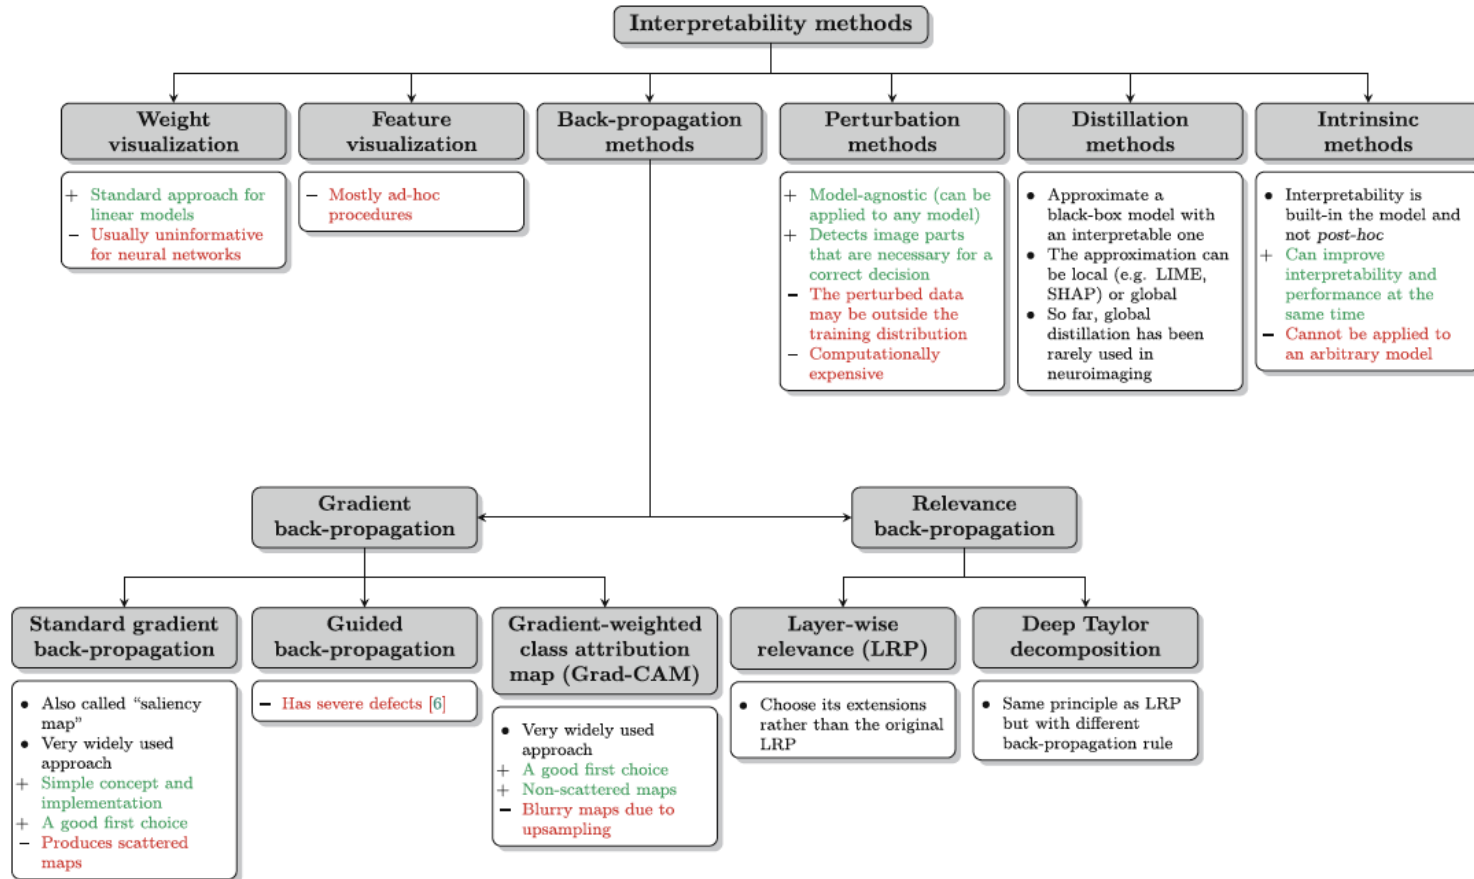
Explainability (XAI)
Interpretability
Understandability

[Erasmus et al 2021 Philosophy & Technology]

For who?
About what?
At which level?
...



Useful approaches : but mainly for developers ...



[Thibau-Sutre et al Neuromethods vol 197 2023]

AD vs Ctrl

The curse of the Black-Box




~~Explainability (XAI)
Interpretability
Understanding~~

*For who?
About what?
At which level?
...*

To Improve Confidence
To trust in AI



*Rigorous Validation
Usage conditions
Adverse effects
.....*




The curse of the Black-Box



~~Explainability
Interpretability~~

*For who?
About what?
What contents?
...*



Trusted AI

G. Hinton [In response to how do we trust systems?] ***You should regulate them based on how they perform.***



**Uncertainty Quantification
in NN decision**

Uncertainty: A keypoint for AI researchers-I

ERC survey (2023): the use of AI for data analysis and processing

- 1034 ERC (/14829)

- Kw: AI in Title, abstract and kw

« the current **underdevelopment of uncertainty quantification**, that is, the assessment of the reliability of models and simulations, and also concerns over the transparency of AI systems. »

Spread false information or inaccurate knowledge



Affect research integrity

https://erc.europa.eu/sites/default/files/2023-12/AI_in_science.pdf

Uncertainty: A keypoint for AI researchers-II

ERC survey (2023): the use of AI for data analysis and processing

- 1034 ERC (/14829)
- Kw: AI in Title, abstract and kw

« the current **underdevelopment of uncertainty quantification**, that is, the assessment of the reliability of models and simulations, and also concerns over the transparency of AI systems. »



<https://www.certain-trust.eu>

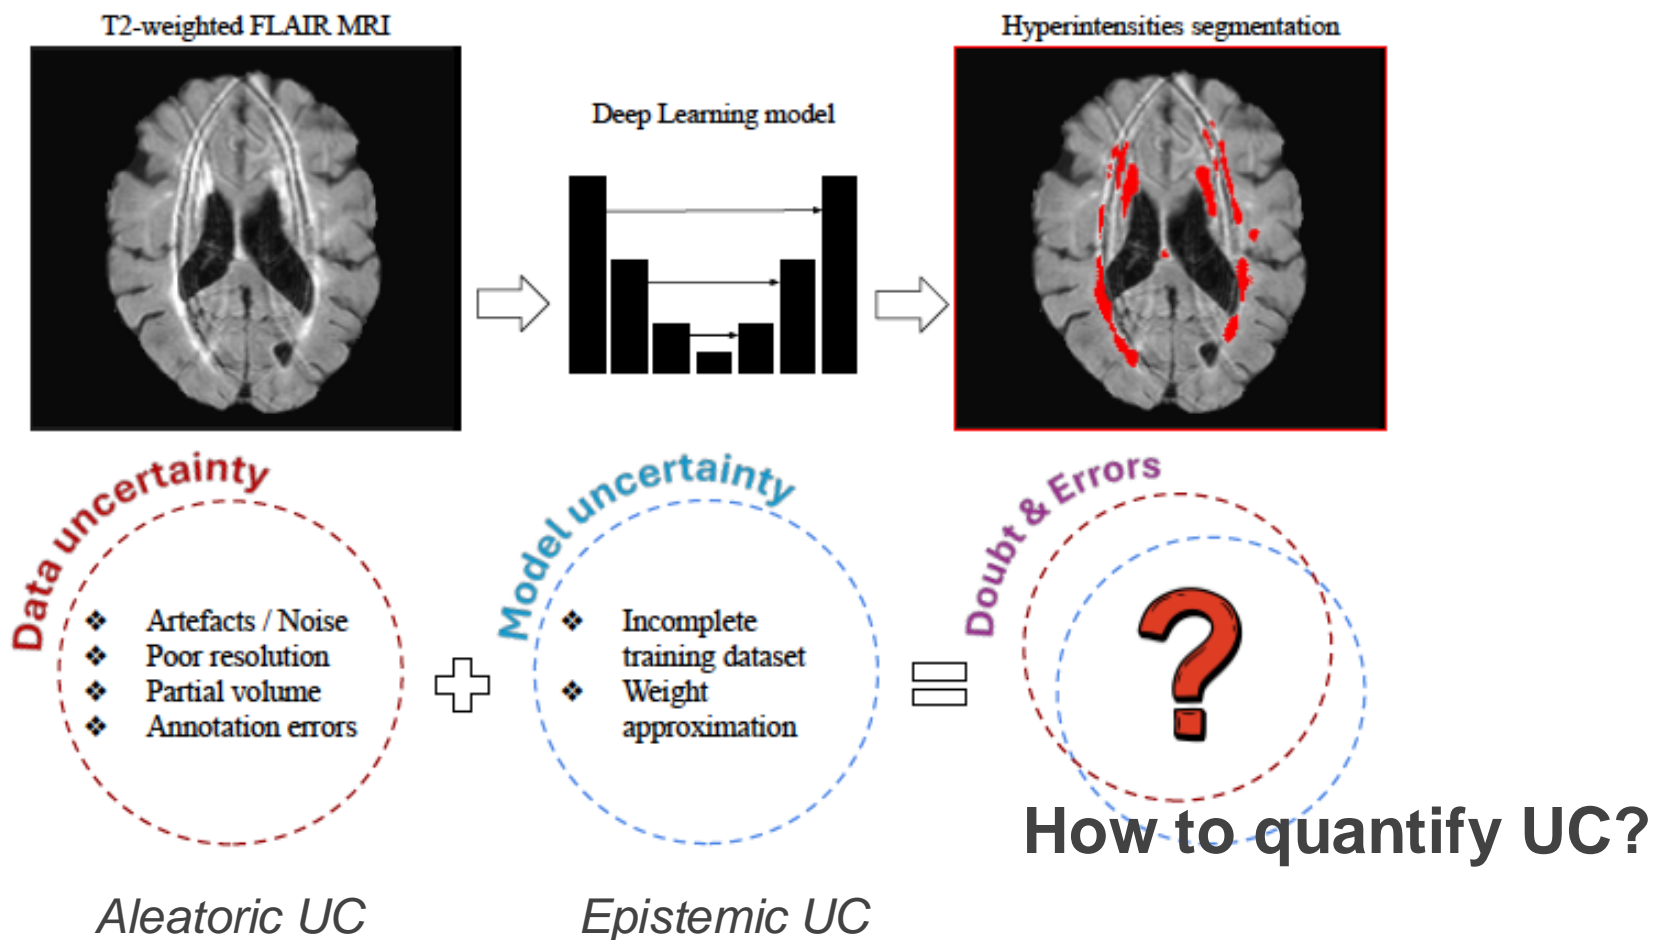
Spread false information or inaccurate knowledge



Affect research integrity

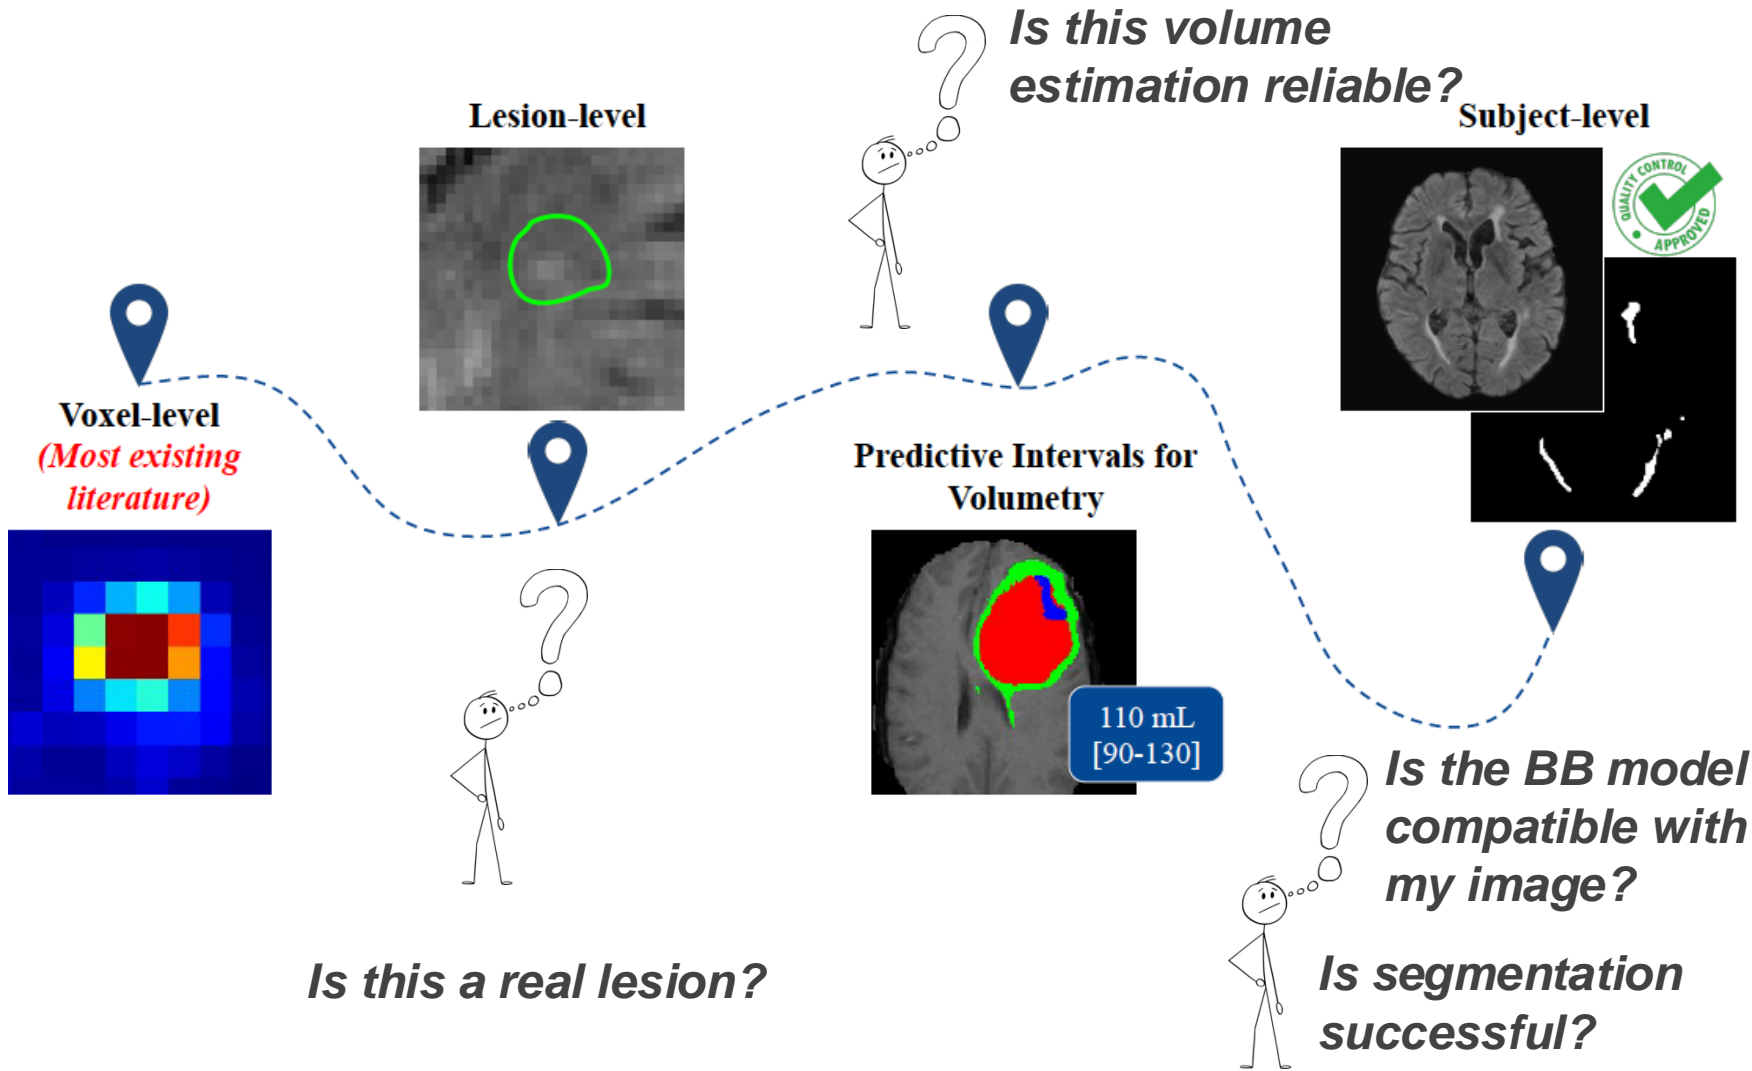
https://erc.europa.eu/sites/default/files/2023-12/AI_in_science.pdf

Where is uncertainty (UC) hidden?

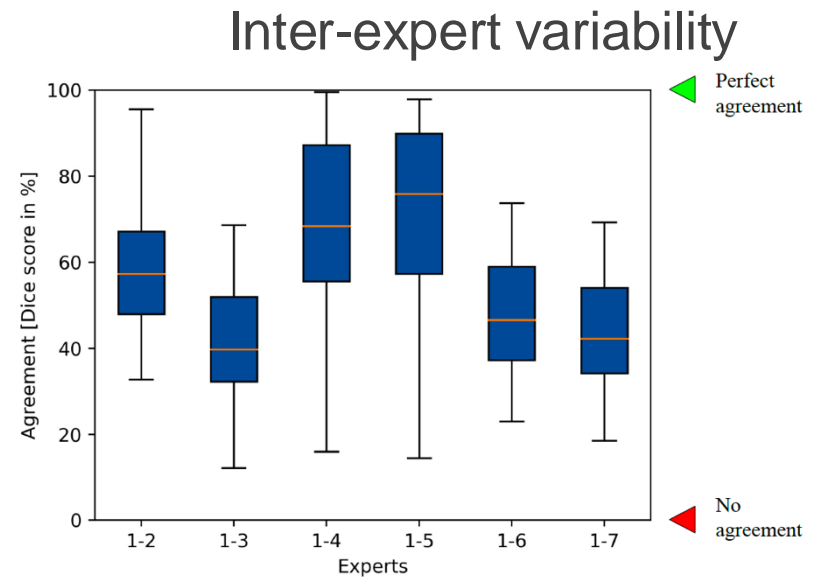
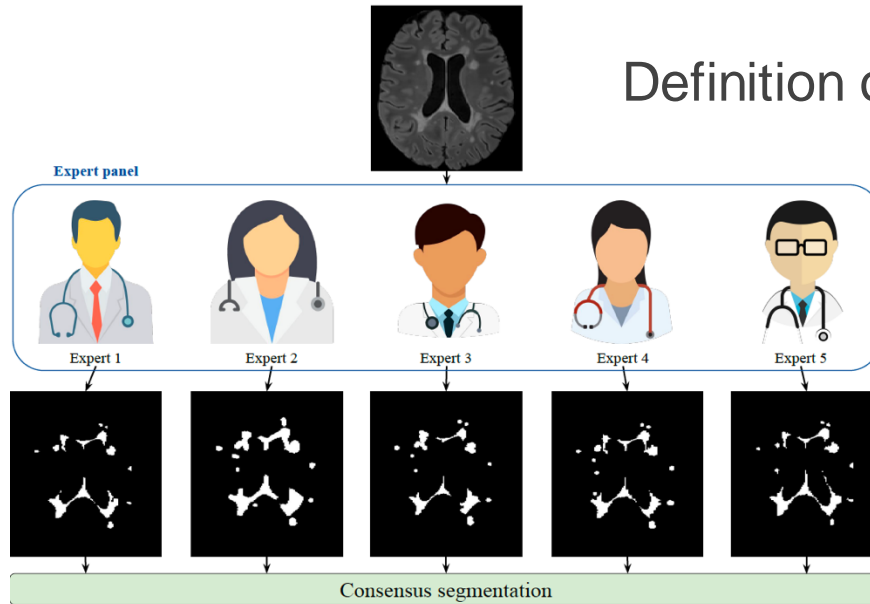


[Kendall & Gal 2017 Adv Neural Inf Process Syst]

The multi-dimensional aspect of uncertainty



Consensus meeting



[Commowick et al NeuroImage 2021]

Existing works

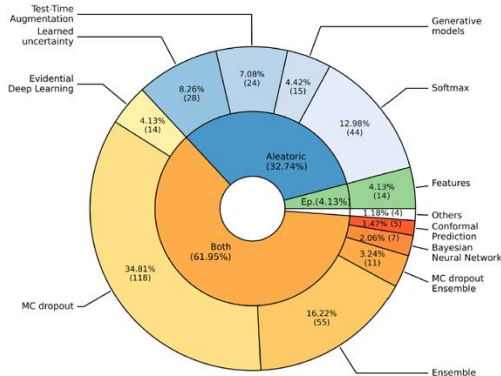
PubMed

218 papers reviewed
 KW: Deep Learning + Medical Image + Classification +
 Segmentation + Uncertainty


Google
 Scholar

11 frameworks identified


9/11 provide
 voxel-level uncertainty estimates



SEP lesions
 in a FLAIR image



Uncertainty
 Quantification



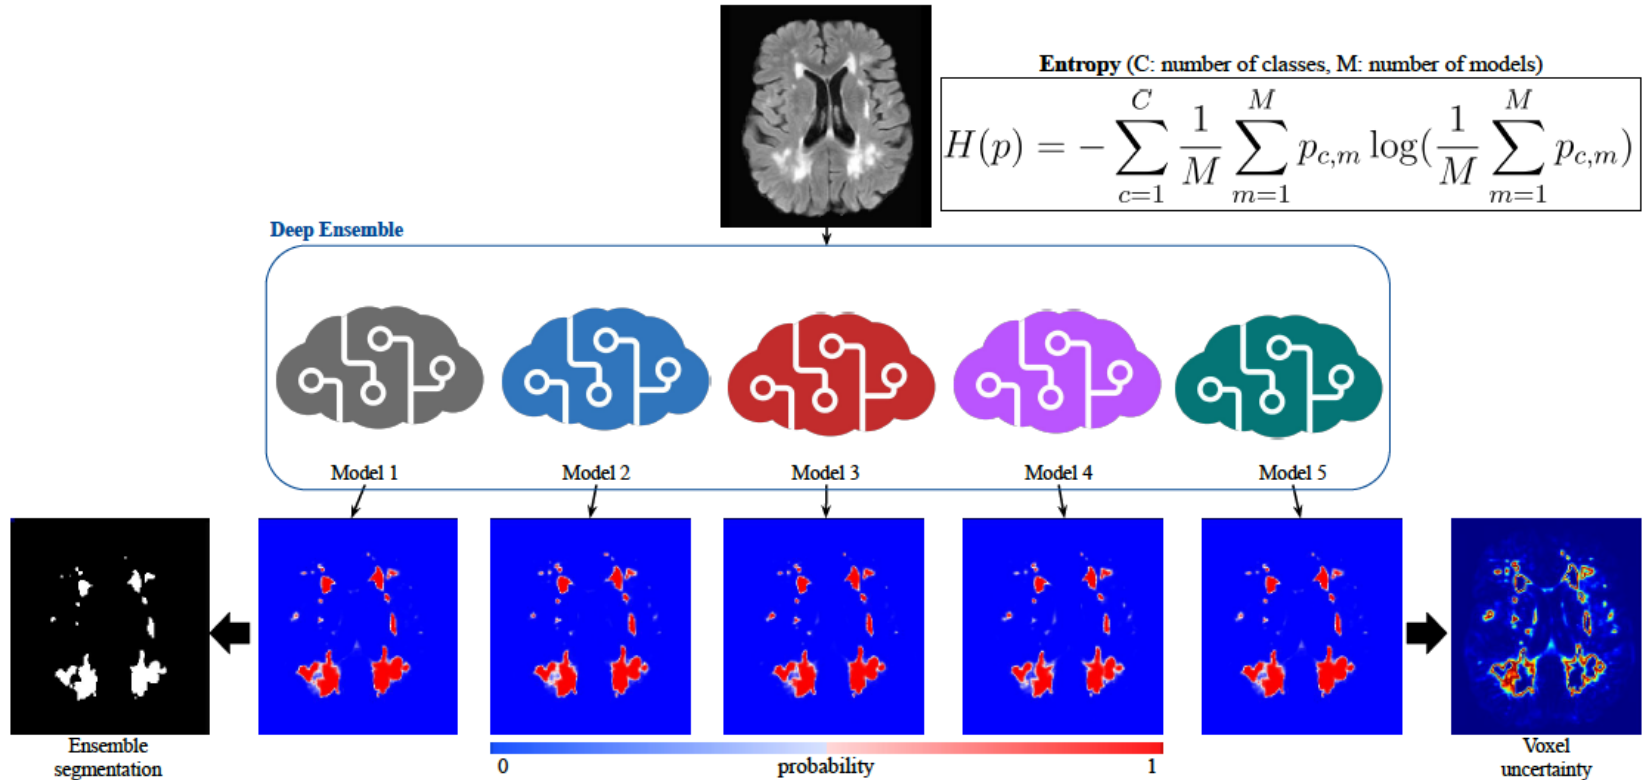
High
 uncertainty

Low
 uncertainty

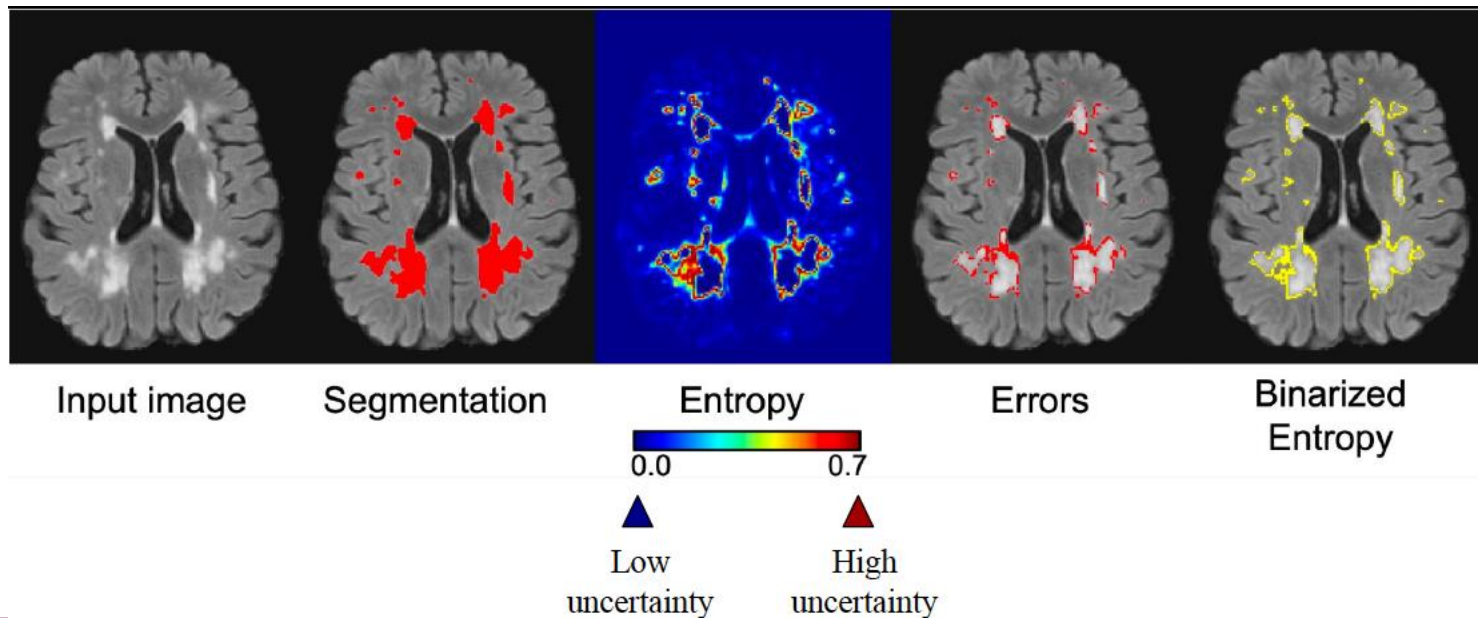
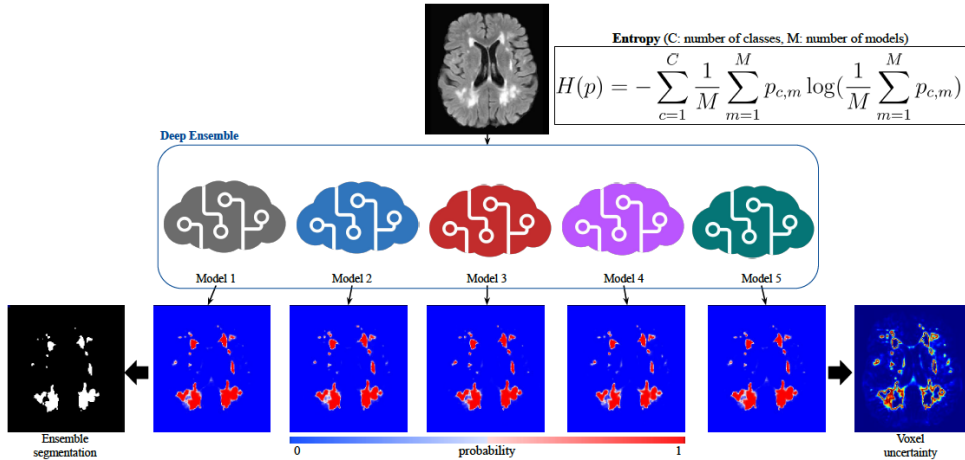
[Lambert et al. AIM 2024.

Trustworthy clinical AI solutions: A unified review of uncertainty quantification in Deep Learning models for medical image analysis]

Voxel level UQ: Ensembling




Voxel level UQ: Ensembling




Lesion level UQ

Goal: Identify FP


Estimate P_{FP} : proba that the lesion is FP using an auxillary classifier



Note: Lesion are highly variable in **shape**




and **position**




Lesions as Graphs

Lesion mask (2D view)




1 voxel = 1 node



Lesion graph (2D view)

Node features

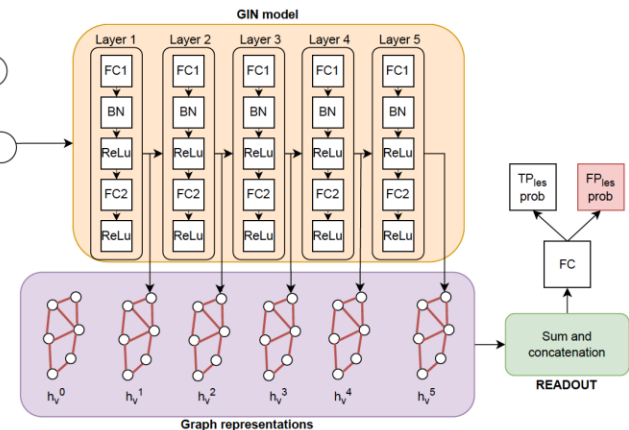
- Contour indicator
- Distance to contour
- Number of neighbors
- Voxel label
- Voxel Entropy
- Voxel Intensity



Mask Voxel uncertainty Intensity


Graph classification

Training to separate True and False Positive lesions





Lambert B et al. Beyond Voxel Prediction Uncertainty: Identifying brain lesions you can trust 2022, LNCS 13611 pp. 61-70.

The Graph Isomorphism Network




FC: Fully-connected layer. BN: Batch Normalization. ReLU: Rectified Linear Unit. Parameters: 26 700

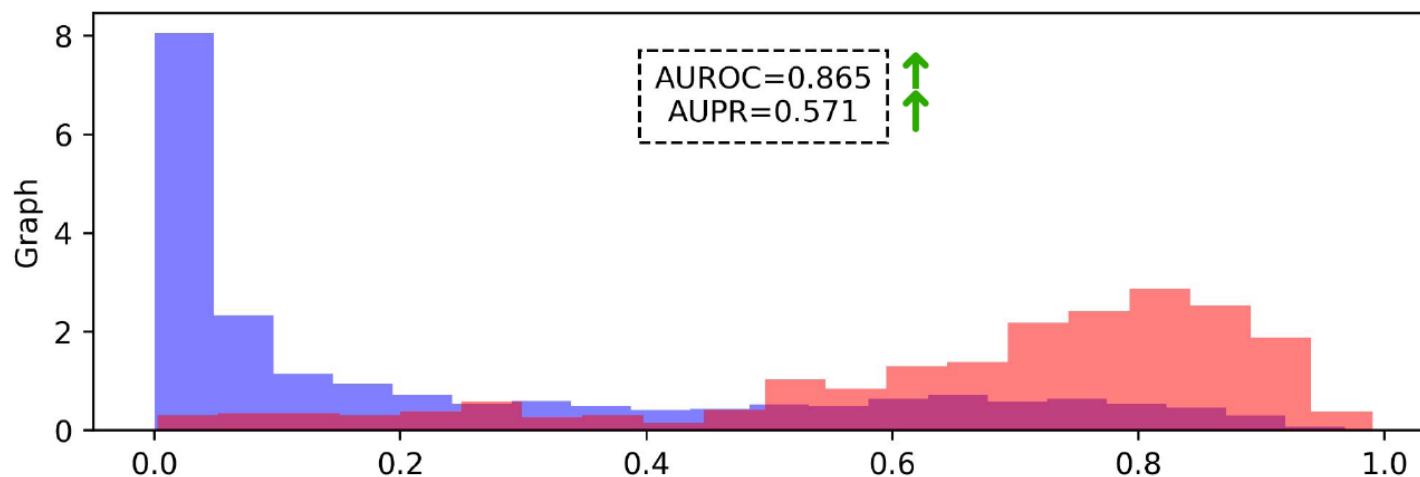
The Graph Isomorphism Network




The Graph Isomorphism Network



Score easily interpretable



The multi-dimensional aspect of uncertainty



Predictive intervals in volumetry

Definition

- $X \in \mathbb{R}^{N-1}$ are estimates of the true volumes $Y \in \mathbb{R}^{N-1}$, obtained from the segmentation.
- A predictive interval $\Gamma_\alpha(X)$ is a range of values intended to encompass Y with a specified degree of confidence $1 - \alpha$ (e.g. 90%, 95%), so that $P(Y \in \Gamma_\alpha(X)) \geq 1 - \alpha$

Sampling-based approaches

- Sample a set of estimated volumes X_1, \dots, X_K for the given image.
- Estimate the mean $\mu(X)$ and standard deviation $\sigma(X)$.
- Assuming $Y|X \sim \mathcal{N}(\mu(X), \sigma(X))$:
$$\Gamma_\alpha(X) = [\mu(X) - z\sigma(X), \mu(X) + z\sigma(X)] \quad (1)$$

Predictive intervals in volumetry

Definition

- $X \in \mathbb{R}^{N-1}$ are estimates of the true volumes $Y \in \mathbb{R}^{N-1}$, obtained from the segmentation.
- A predictive interval $\Gamma_\alpha(X)$ is a range of values intended to encompass Y with a specified degree of confidence $1 - \alpha$ (e.g. 90%, 95%), so that $P(Y \in \Gamma_\alpha(X)) \geq 1 - \alpha$

Sampling-based approaches

- Sample a set of estimated volumes X_1, \dots, X_K for the given image.
 - Estimate the mean $\mu(X)$ and standard deviation $\sigma(X)$.
 - Assuming $Y|X \sim \mathcal{N}(\mu(X), \sigma(X))$
- $$\Gamma_\alpha(X) = [\mu(X) - z\sigma(X), \mu(X) + z\sigma(X)]$$

Limitations

- **Inference time**, due to the sampling procedure.
- The **normality assumption**, which may not always hold.
- **Lack of flexibility**, as intervals are symmetrical by design.

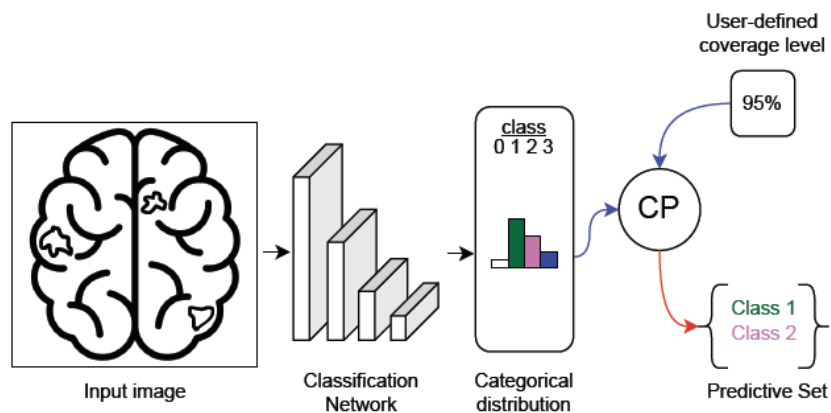
Predictive intervals in volumetry

Definition

- $X \in \mathbb{R}^{N-1}$ are estimates of the true volumes $Y \in \mathbb{R}^{N-1}$, obtained from the segmentation.
- A predictive interval $\Gamma_\alpha(X)$ is a range of values intended to encompass Y with a specified degree of confidence $1 - \alpha$ (e.g. 90%, 95%), so that $P(Y \in \Gamma_\alpha(X)) \geq 1 - \alpha$

Sampling-based approaches

- Sample a set of estimated volumes X_1, \dots, X_K for the given image.
- Estimate the mean $\mu(X)$ and standard deviation $\sigma(X)$.



Direct approaches

- Directly estimate the quantiles $\hat{t}_{\alpha/2}(X)$ and $\hat{t}_{1-\alpha/2}(X)$.
- The PI is computed as:

$$\Gamma_\alpha(X) = [\hat{t}_{\alpha/2}(X), \hat{t}_{1-\alpha/2}(X)] \quad (2)$$

Conformal prediction

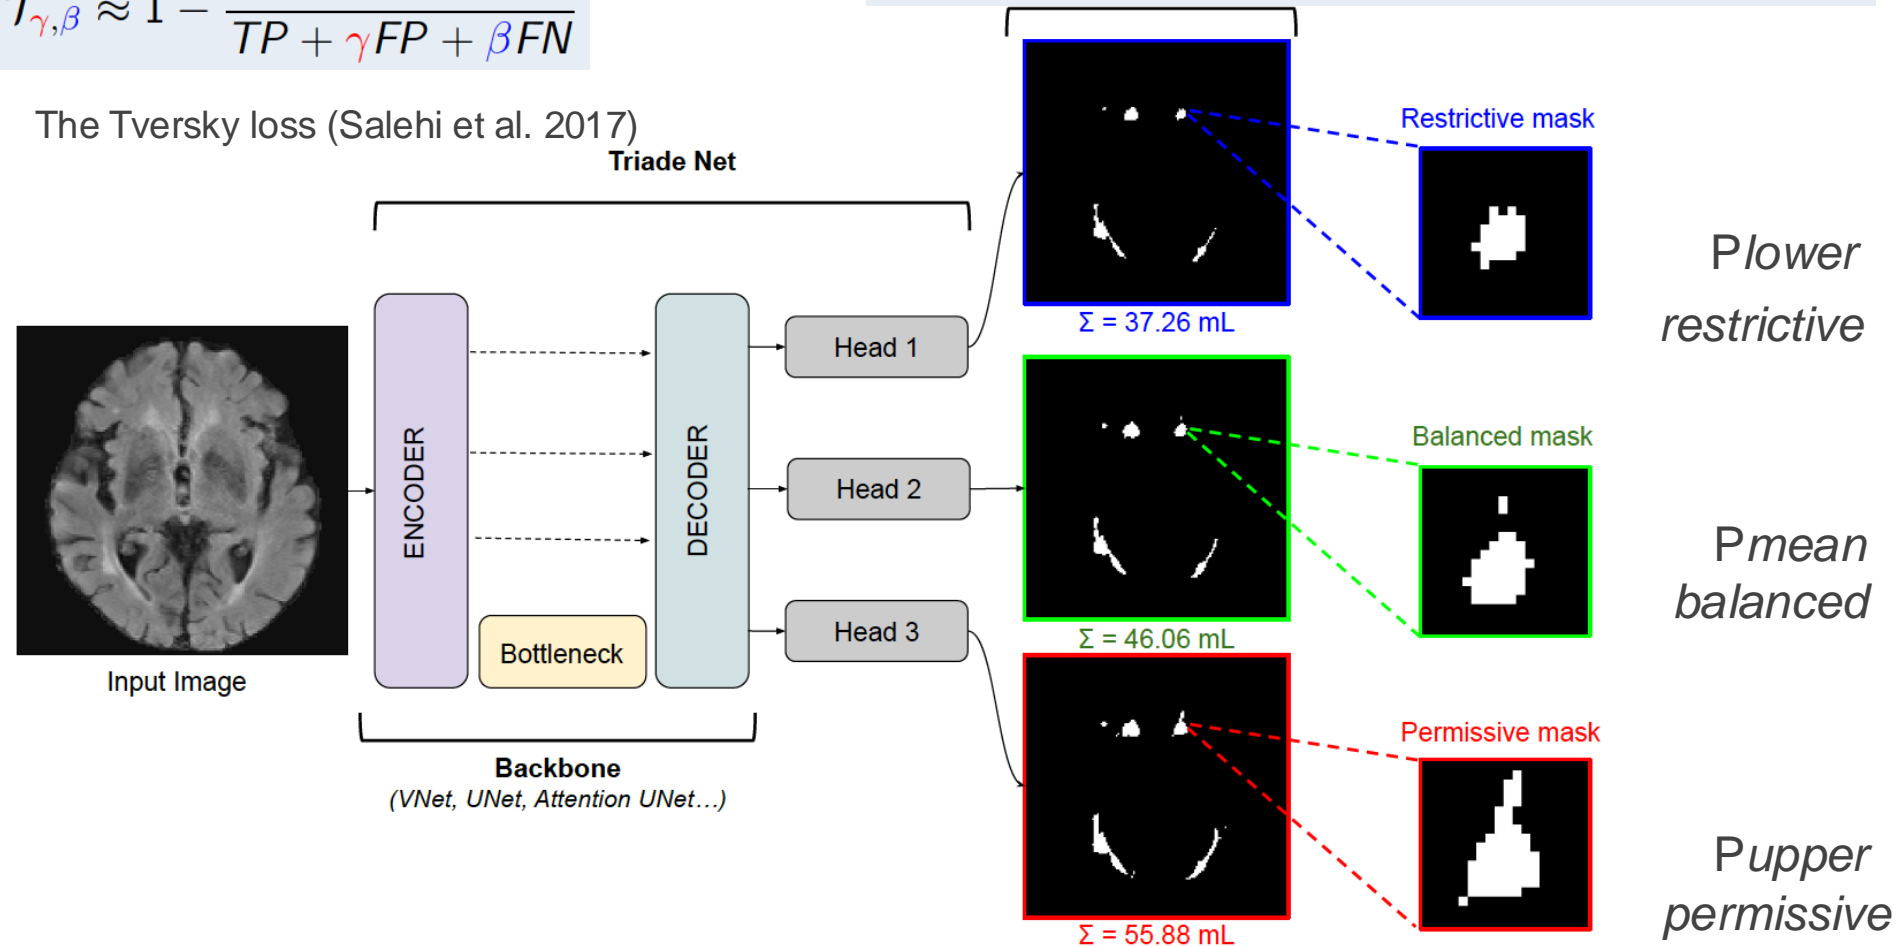
The TriadNet approach

$$T_{\gamma, \beta} \approx 1 - \frac{TP}{TP + \gamma FP + \beta FN}$$

The Tversky loss (Salehi et al. 2017)



$$\text{TriadLoss} = T_{0.8, 0.2}(p_{lower, y}) + T_{0.5, 0.5}(p_{mean, y}) + T_{0.2, 0.8}(p_{upper, y})$$

Triade Net



[B. Lambert et al. (2023). "TriadNet: Sampling-Free Predictive Intervals for Lesional Volume in 3D Brain MR Images". In: UNSURE 2023, LNCS 14291, pp. 32–41]


TriadNet predictions



TriadNet predictions

Need for calibration

- 120 subjects for training, 40 for calibration and 50 for in-distribution testing. (Multicentric - 3 Tesla: MSSEG 2016 / WMH 2017 / ISBI 2015)
- Intervals calibrated for a target coverage of 90%.
- Metrics (bootstrapping, $M = 15000$):
 - Mean Average Error: 3.08 ± 0.46 mL
 - Empirical Coverage: $92.06 \pm 5.34\%$




Tumor volume estimation

Glioblastoma

- 679 subjects for training, 227 for calibration, and 227 for testing (BraTS 2023 dataset)
- Intervals calibrated for a target coverage of 90%.


Necrosis:

- MAE: $3.10 \pm 0.46\text{mL}$
- Coverage: $90.78 \pm 2.71\%$



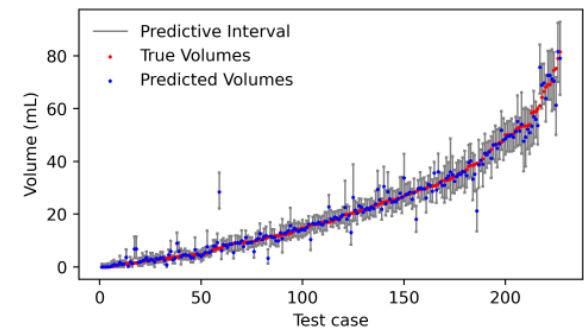
Edema:

- MAE: $8.22 \pm 0.57\text{mL}$
- Coverage: $90.76 \pm 2.70\%$

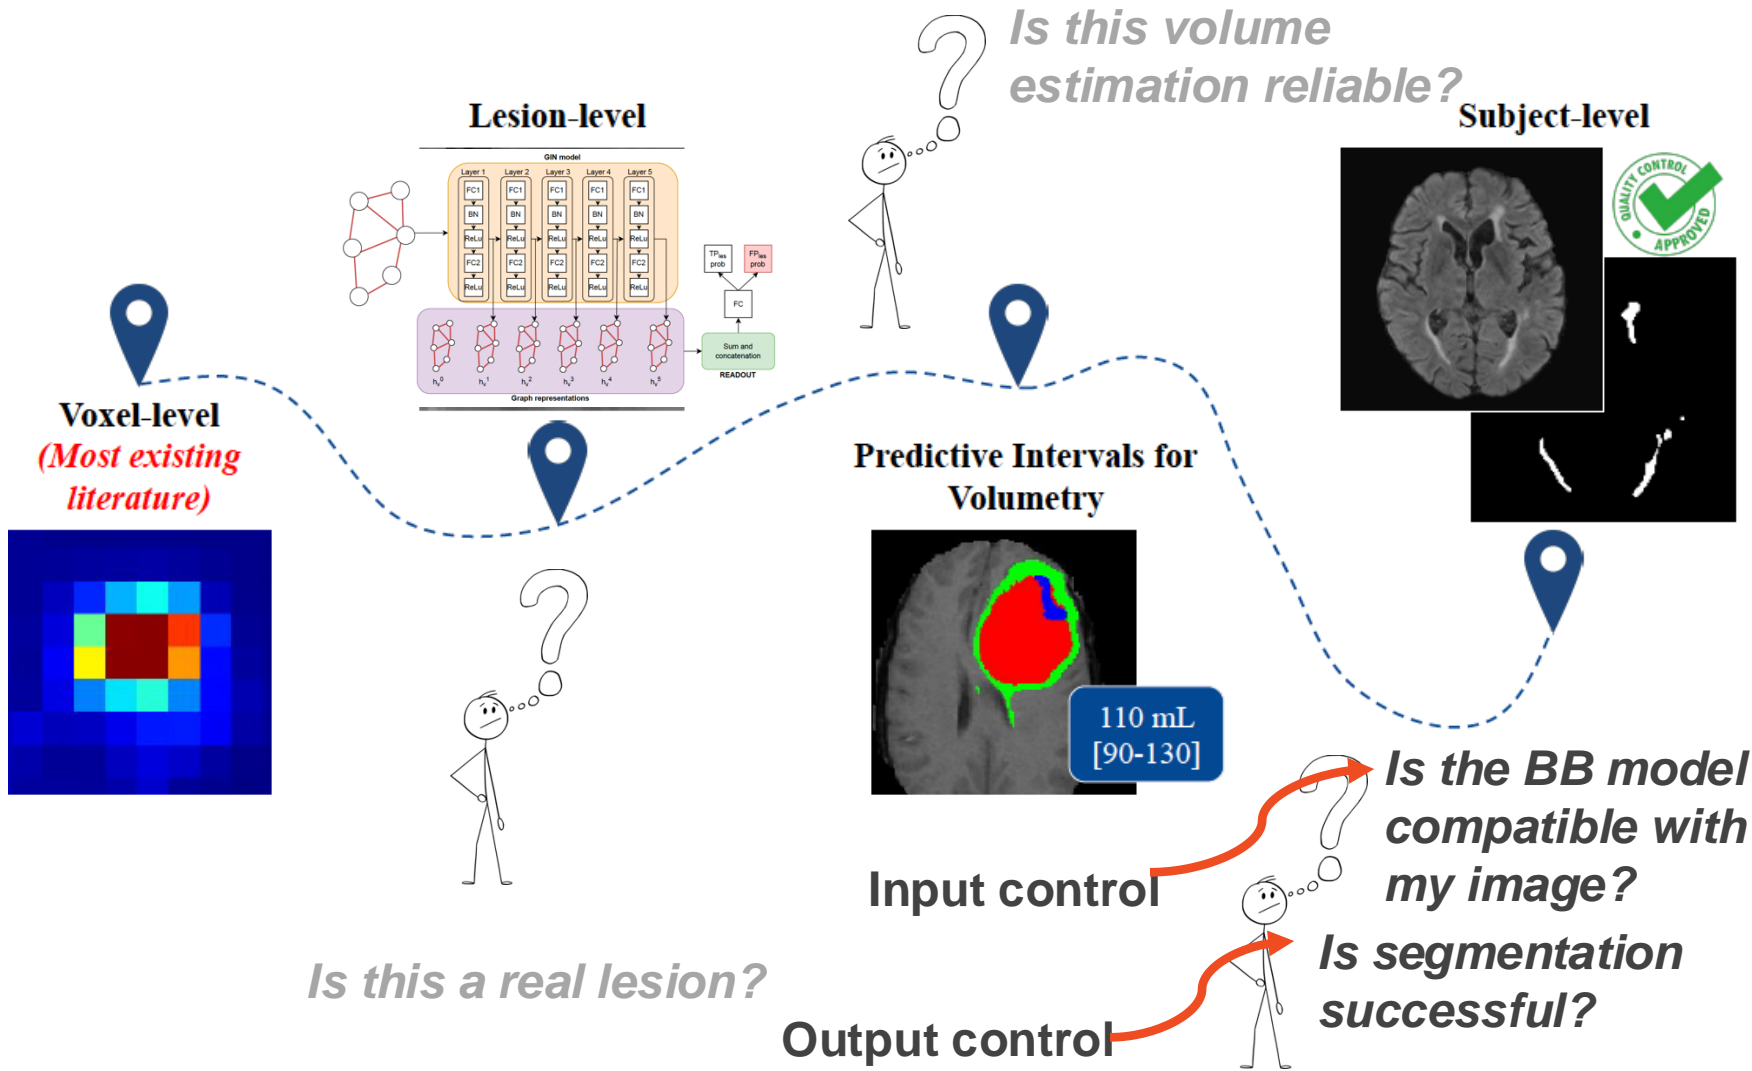


Enhancing tumor:


- MAE: $1.73 \pm 0.19\text{mL}$
- Coverage: $90.79 \pm 2.71\%$




The multi-dimensional aspect of uncertainty




DLL trained for Glioblastoma detection on T1w




Artefacted T1w



Healthy subject




FLAIR




Abdominal T1w

Know-it-all


DLL trained for Glioblastoma segmentation on T1w




Artefacted T1w







Healthy subject



FLAIR




Abdominal T1w




Input Quality Control

DLL trained for Glioblastoma segmentation on T1w



Artefacted T1w




Healthy subject




FLAIR




Abdominal T1w




Tumoral vol
251 ml



Tumoral vol
12 ml




Tumoral vol
239 ml



Tumoral vol
258 ml

Images far from the training image distribution (OOD) should be detected


Why an image is OOD?



- In-distribution \leftrightarrow training distribution (T1 MRI of Adult glioblastoma patients)
- Are out-of-distribution:
 - Images corrupted with artifacts.
 - Shifts in the imaged population.
 - Shifts in image modality.
 - Diseases not present in the training set.
 - Incorrect organs.



Input Quality Control

The latent-space distance



Input Quality Control

The latent-space distance



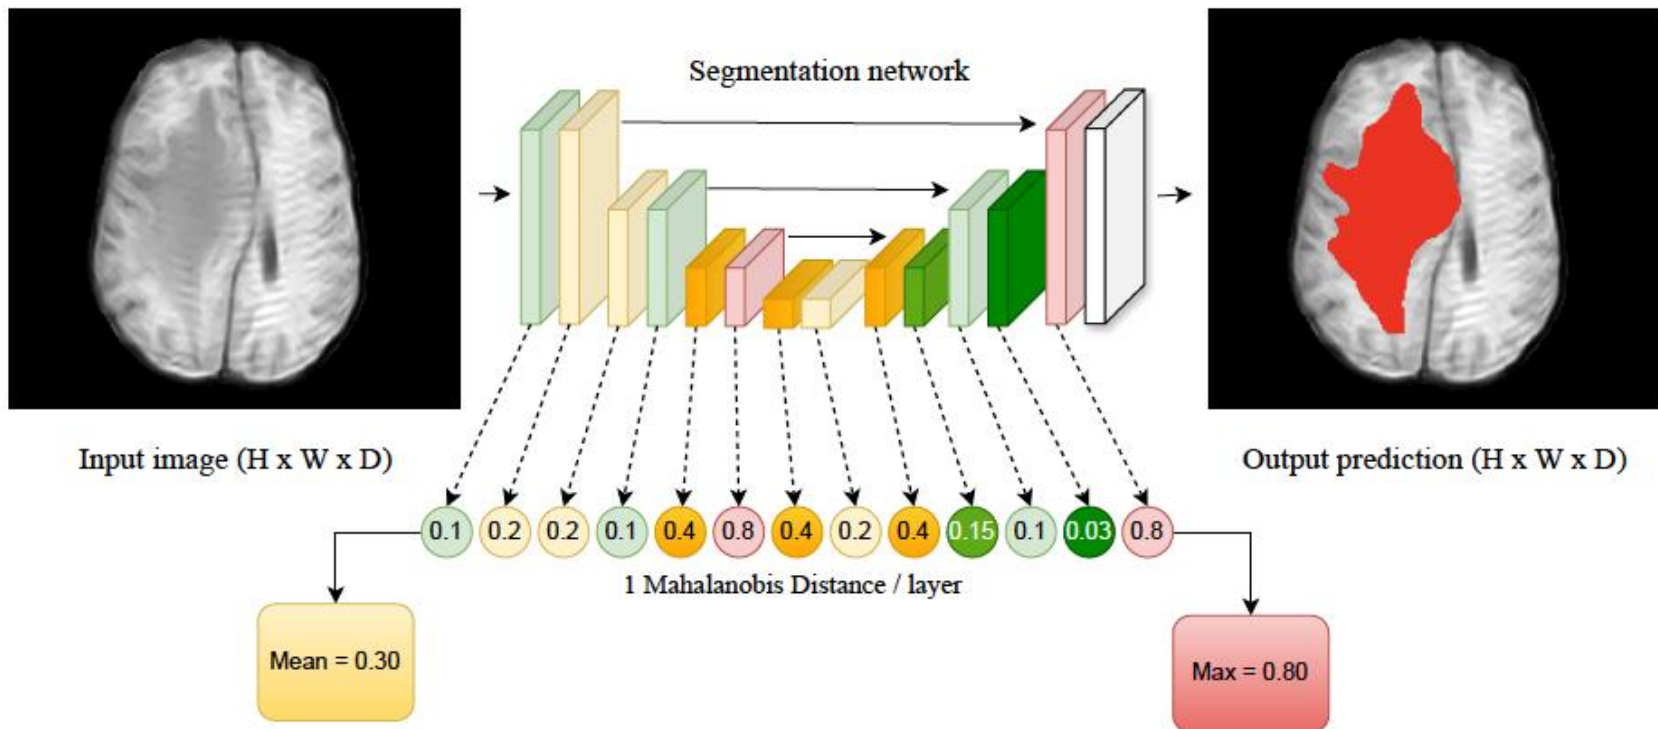
Computed with Mahalanobis distance (MD)

$$\mu = \frac{1}{N} \sum_{i=1}^N x_i$$

$$\Sigma = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)(x_i - \mu)^T$$

$$MD(x_{test}; \mu, \Sigma) = (x_{test} - \mu)^T \Sigma^{-1} (x_{test} - \mu)$$

A multi-layer aggregation of MD




[B. Lambert et al. (2023). “Multi-layer Aggregation as a key to feature-based OOD detection”. In: UNSURE 2023, Held in Conjunction with MICCAI 2023. LNCS 14291, pp. 104–114]

Input Quality Control

Brats: 876 subjects for training,
30 for validation,
227 for in-distribution testing




Dynamic U-Net



The optimal layer for OOD detection depends on the segmentation architecture.

Input Quality Control

The multi-layer scores (**Mean** and **Max**) provides high detection accuracy.




Latent-space distances efficient in detecting images far from the training distribution.

Output Quality Control

Goal: estimate the true segmentation accuracy.

How: Measure the segmentation variability among models.

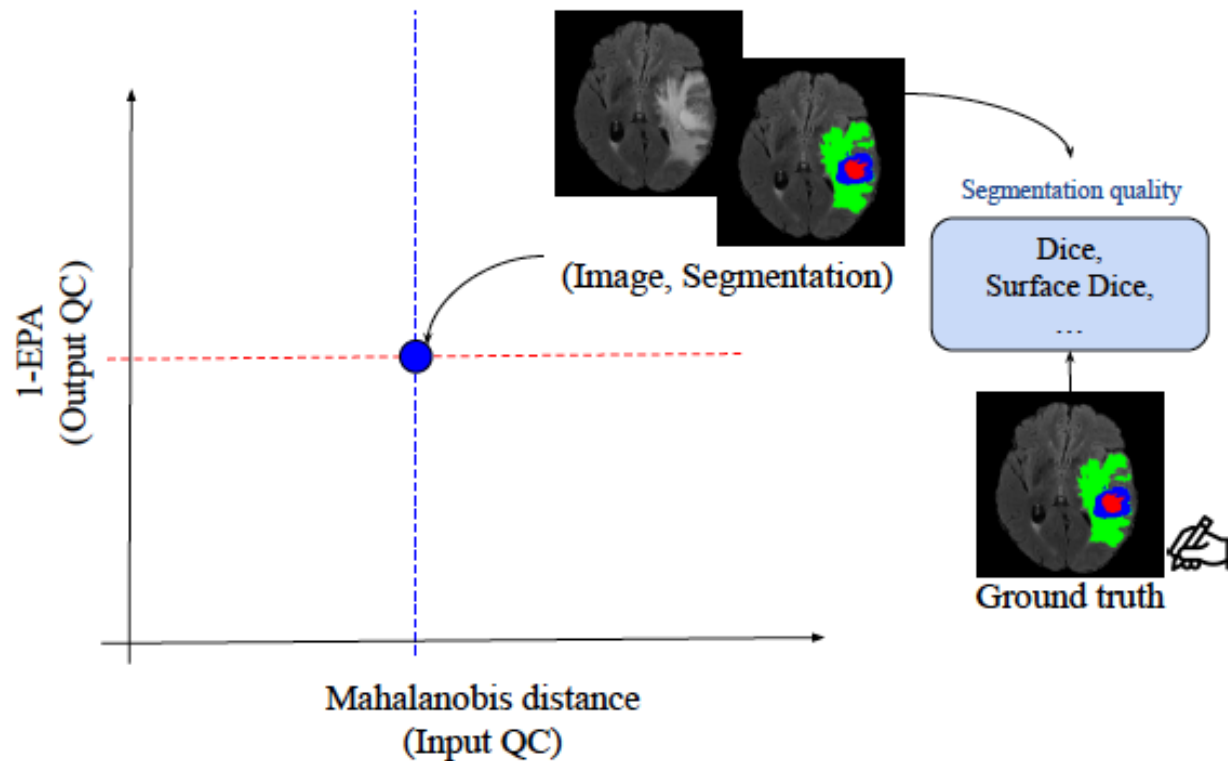


We note S_k the individual segmentations and MV the majority vote segmentation

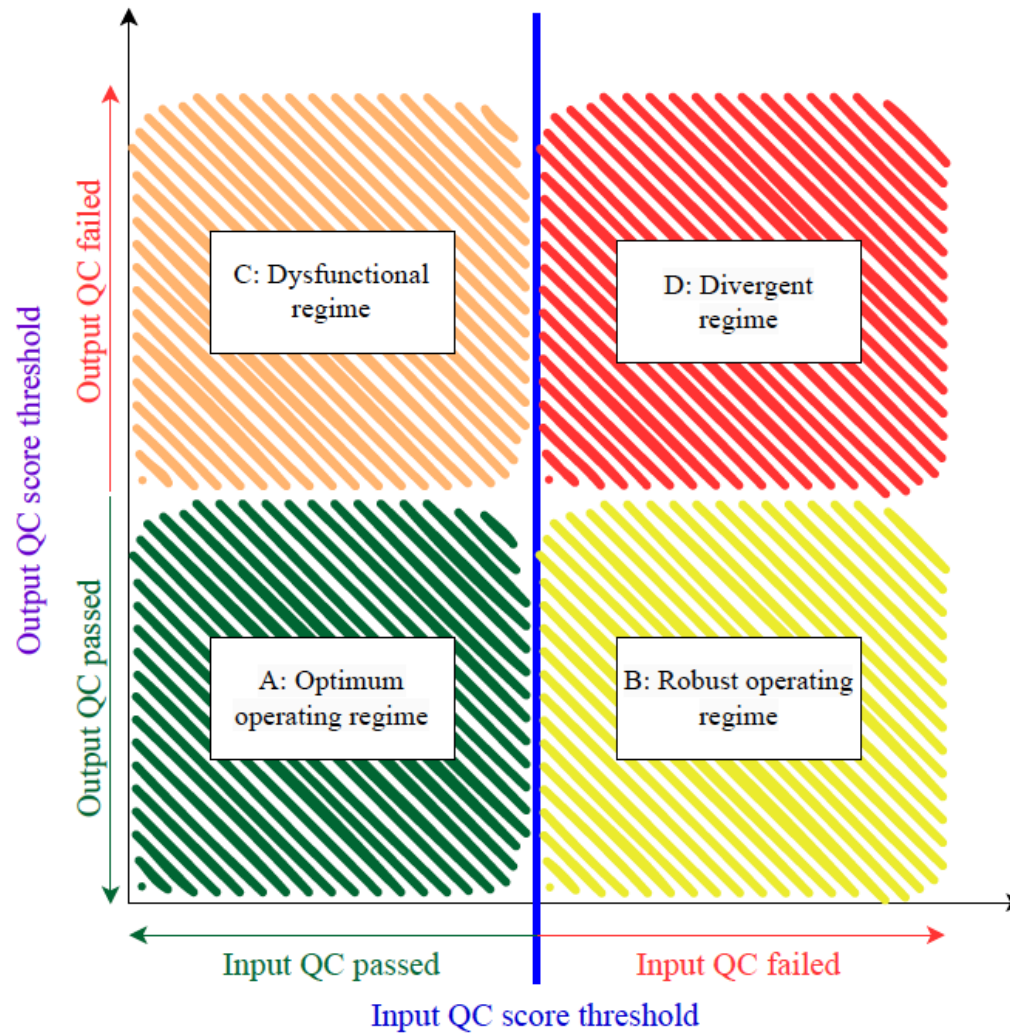
Ensemble Prediction Agreement (EPA): $EPA = \frac{1}{K} \sum_{i=1}^K \text{Dice}(S_k, MV)$

[B. Lambert et al. (2024) “From Out-of-distribution detection to Quality Control”.
In: Trustworthy AI in Medical Imaging, MICCAI book series]

Unified Input-Output control



Unified Input-Output control

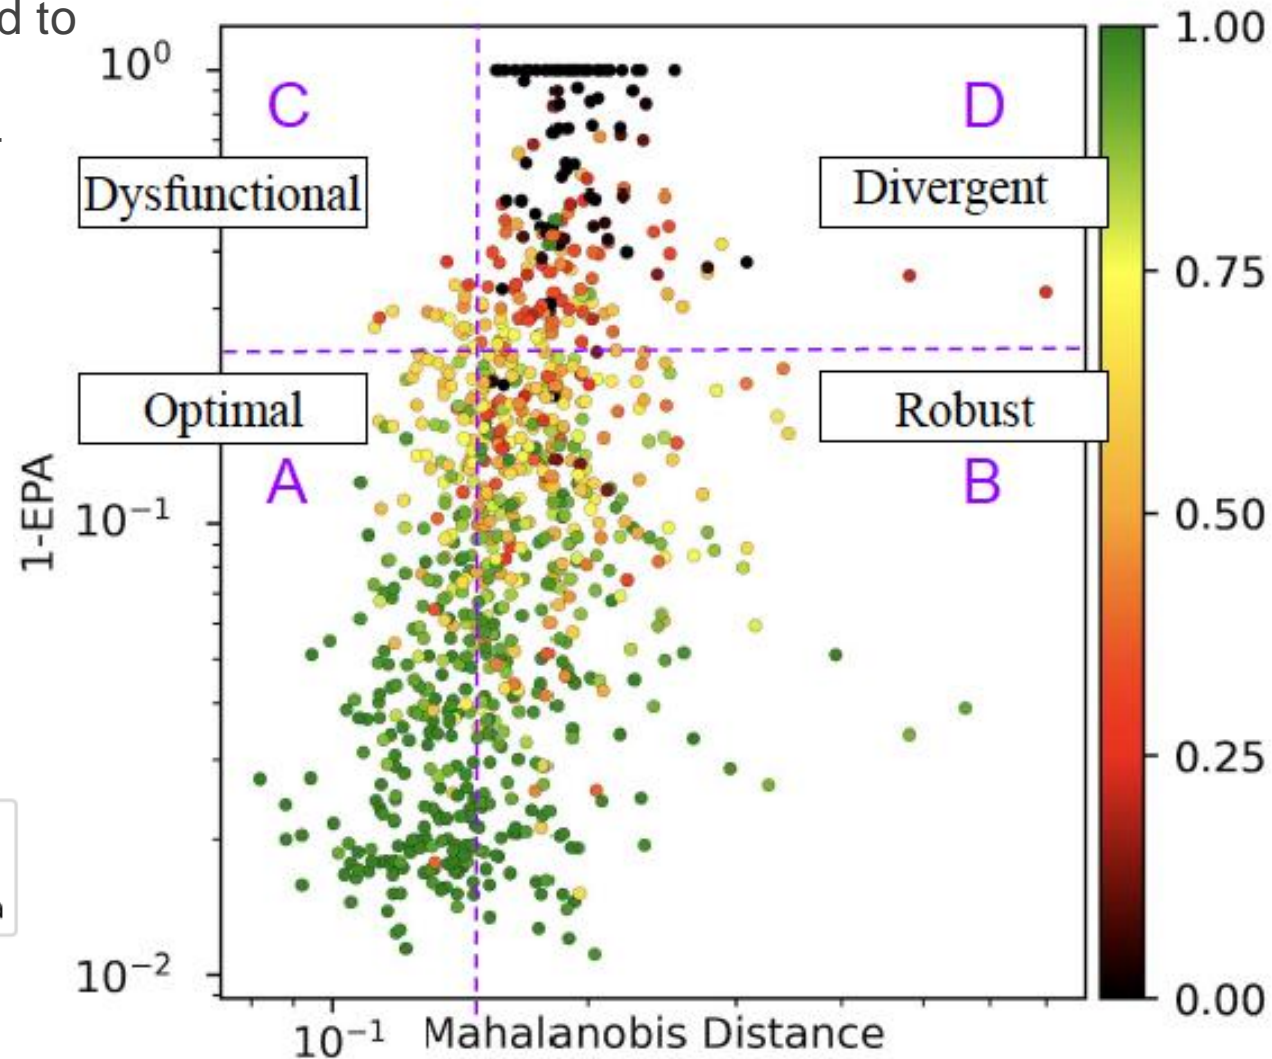


Unified Input-Output control

5 Dynamic U-Nets are trained to segment gliomas.

QC scores computed for 874 test subjects with variable difficulty.

thresholds fixed on a validation dataset (N=30).



- In Distribution
- Synthetic
- Africa
- Pediatric
- Metastases
- Meningioma

Take home messages

- ✓ Confidence is central for AI deployment
- ✓ Uncertainty quantification improves user's confidence
- ✓ Uncertainty is multidimensional

Lesion uncertainty scores
Predictive volume intervals
Unified input & output controls

} uncertainty quantification
for trusted AI

↓
for AI penetration in clinical
routine


NEXT STEP

What is the added-value in clinical routine applications?

TO BE EVALUATED



Trustworthy AI in Medical Imaging



Trustworthy AI in Medical Imaging

- 1st Edition - **December 1, 2024**
- Editors: Marco Lorenzi, Maria A Zuluaga
- eBook ISBN: 9780443237607

Section 1 – Robustness

Section 2 - Validation, Transparency and Reproducibility

Section 3 – Bias and Fairness

Section 4 - Explainability, Interpretability and Causality

Section 5 - Privacy-preserving ML

Section 6 - Collaborative Learning

Section 7 - Beyond the Technical Aspects

