Flexible machine learning approaches for computer-assisted surgery

Tom Vercauteren, PhD

Professor of Interventional Image Computing
School of Biomedical Engineering & Imaging Sciences
Faculty of Life Sciences and Medicine
King's College London

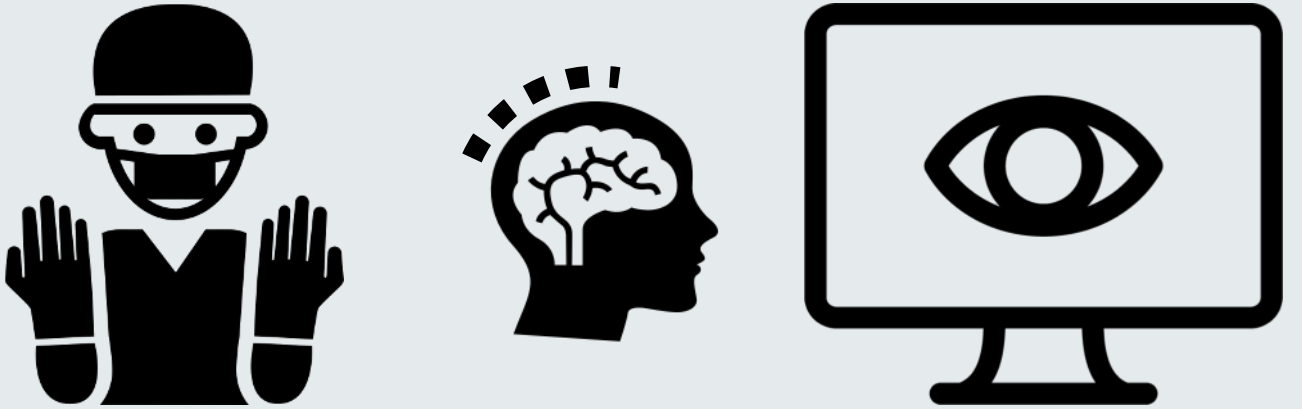

IPAM Workshop on Deep Learning and Medical Applications
January 2020

KING'S College LONDON

Forewords

- Most of the work presented in this talk stemmed from work from two of my PhD students

- Reuben Dorent, ongoing PhD, working on hetero-modal learning

- Guotai Wang, PhD between 2014-2018 on minimally interactive segmentation
  - Currently associate professor at the at University of Electronic Science and Technology of China

Personal impact of IPAM

http://www.ipam.ucla.edu/programs/summer-schools/summer-school-mathematics-in-brain-imaging/

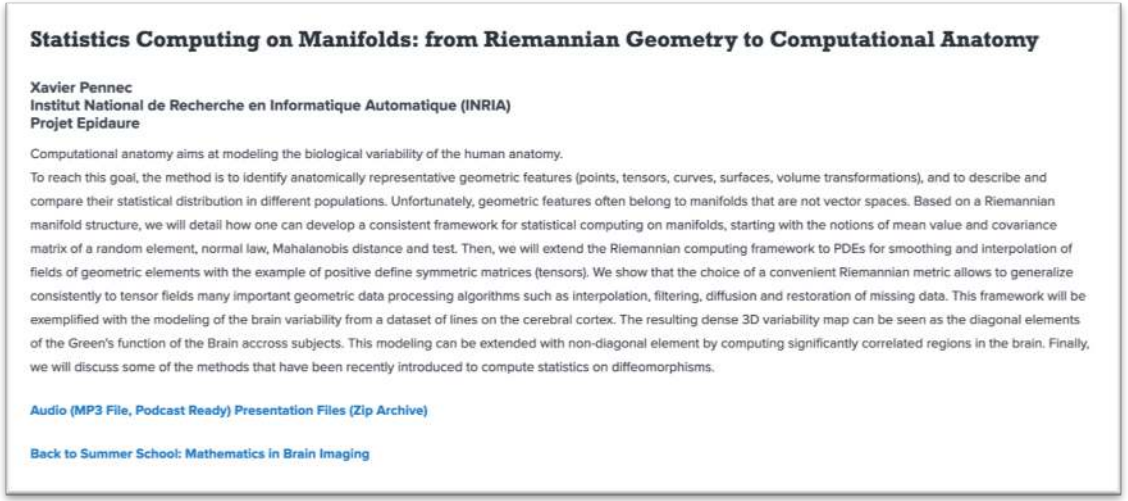

http://www.ipam.ucla.edu/abstract/?tid=7413&pcode=MBI2008

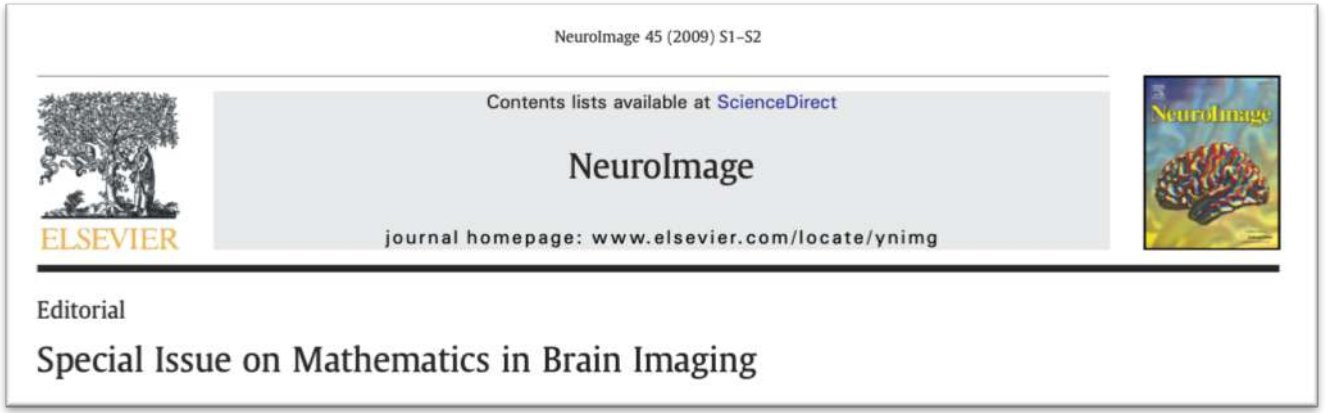

https://www.sciencedirect.com/science/journal/10538119/45/1/supp/S1

https://doi.org/10.1016/j.neuroimage.2008.10.040

Delivering flexible patient-specific decision support tools

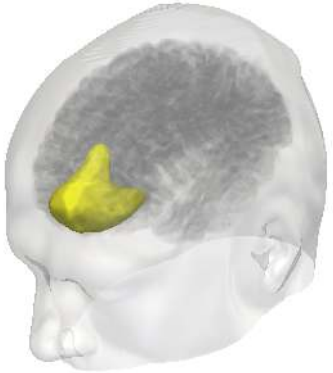

Image courtesy of Bilger et al.

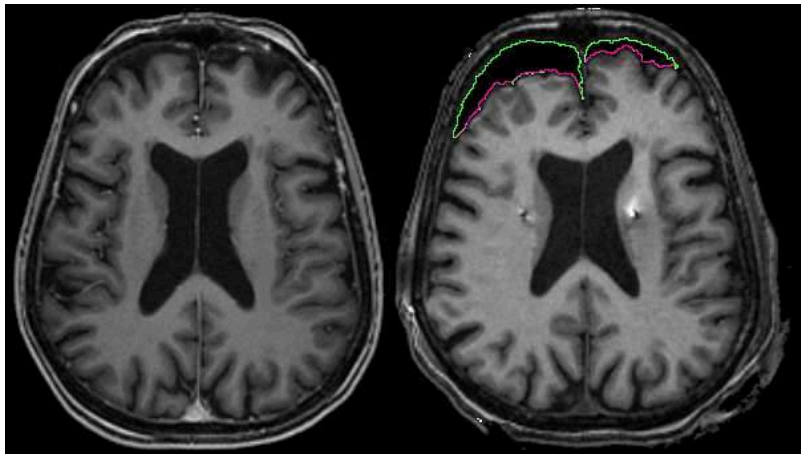

Image courtesy of Stechison

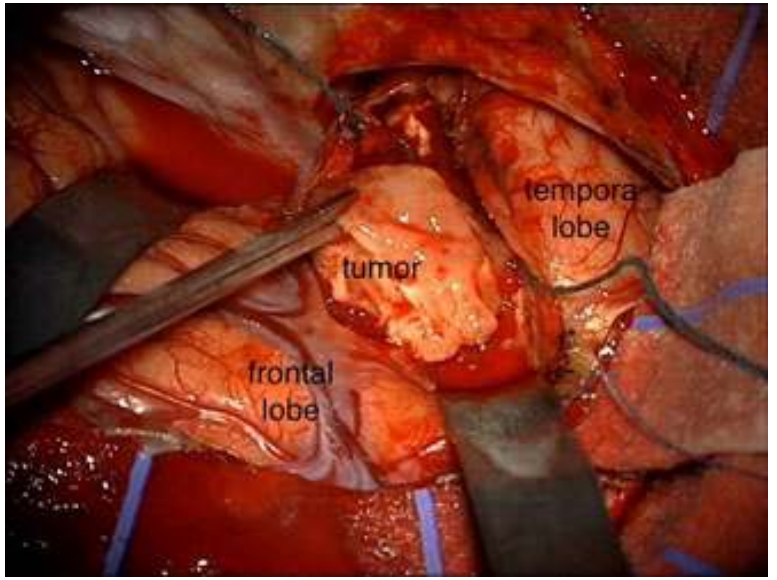

- Key role of medical imaging
  - Pre-operative planning
  - Intra-operative navigation / guidance
  - Post-operative assessment
- Challenges
  - Only partial standardisation of imaging
  - Intra-operative changes
  - Interpretation difficulties
  - Cognitive workload
- Requirements
  - Integrate in the workflow
  - Only partial annotations for learning
  - Leave the clinical team in control

Multimodal learning in medical imaging

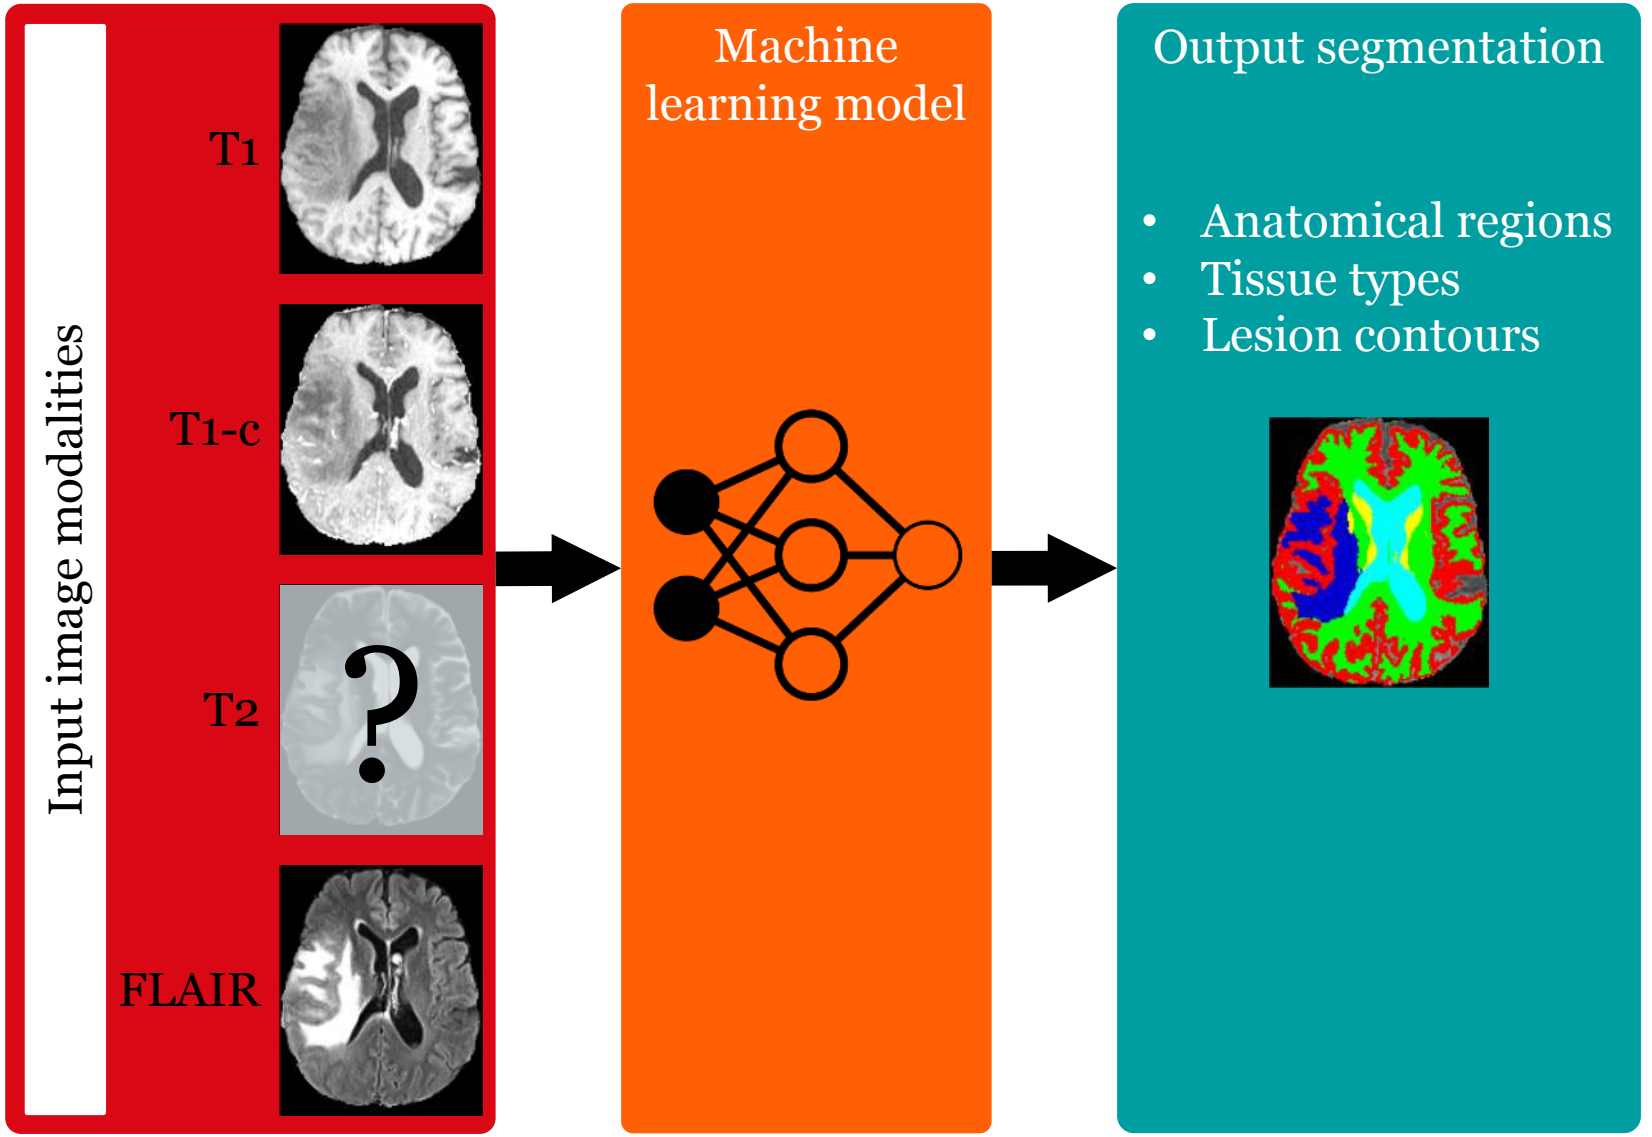

<u>Hetero-modality challenge</u>

Practical clinical workflows rarely provides all modalities for every patient
→ Need to handle hetero-modal scenarios

<u>Naïve solutions are suboptimal</u>

- Exploiting only common modality
- Zero filling
- Learning imputation as pre-processing

Training datasets in medical imaging

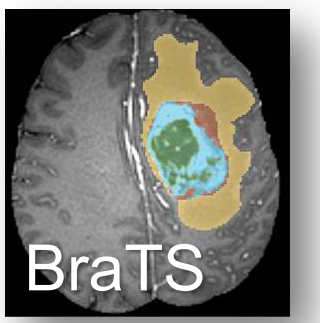

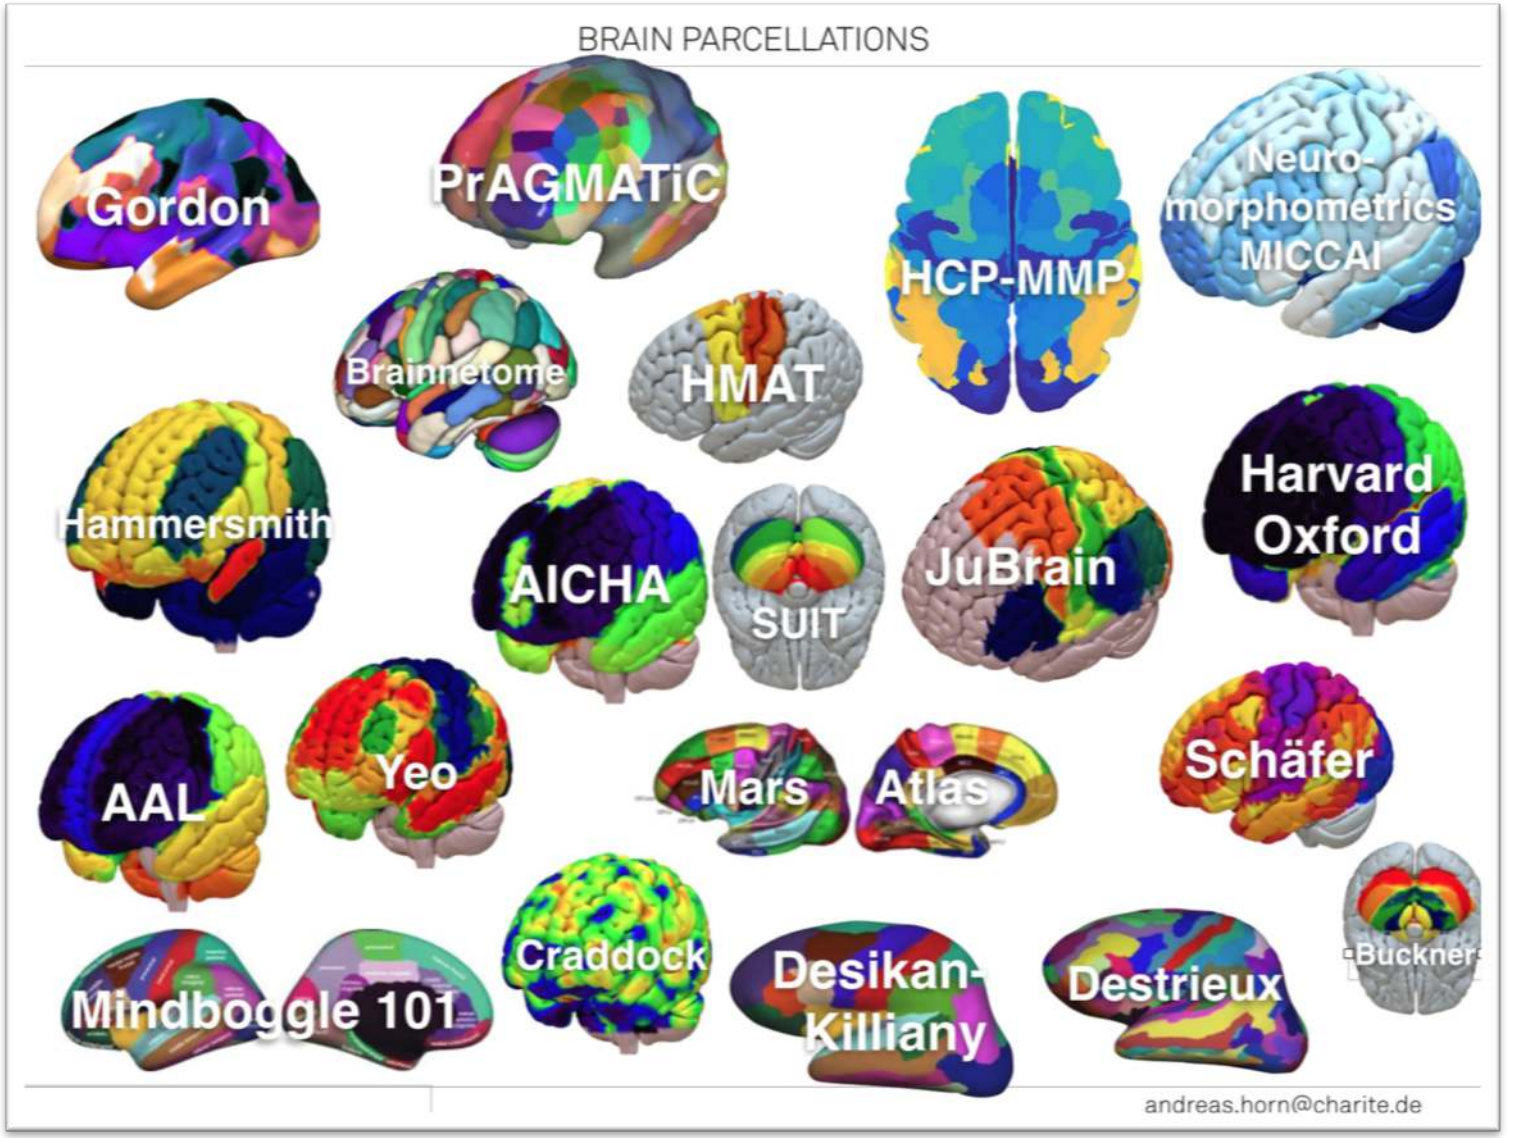

<u>Heterogeneous dataset challenge</u>

Training datasets are usually task specific
→ Need to handle multi-task learning innovatively

Quality control and human-in-the-loop for medical imaging

Robustness challenge

Despite some predictions, robustness will remain a key issue

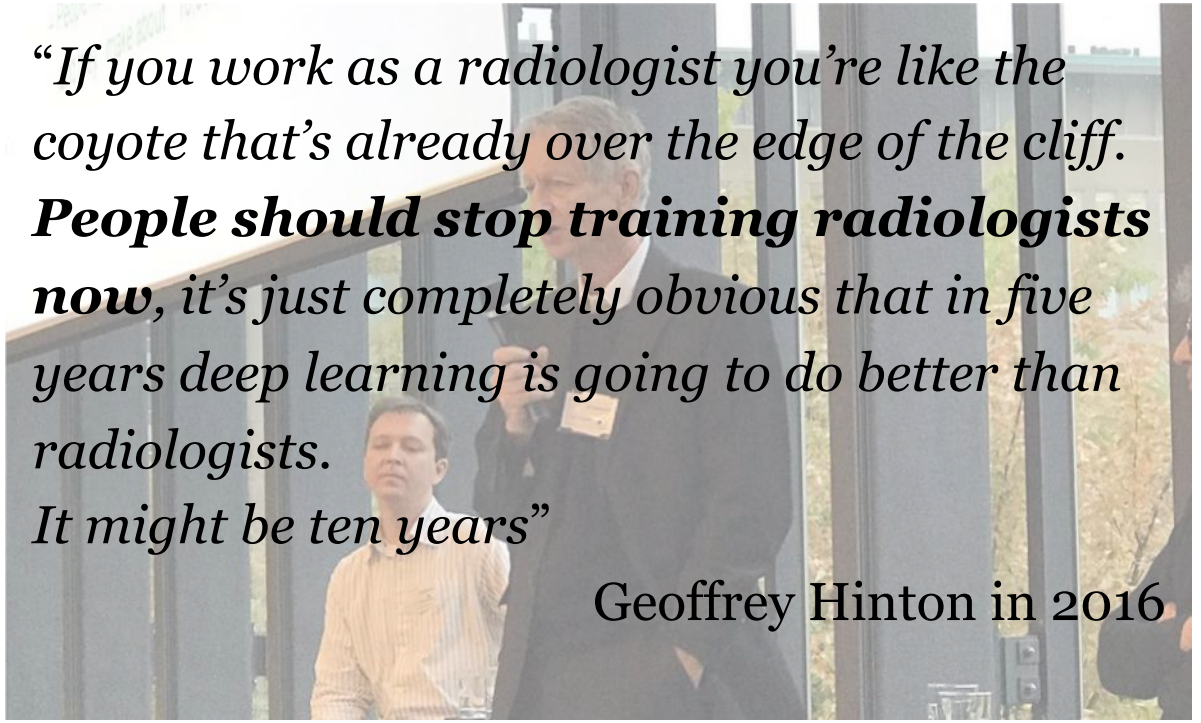

"If you work as a radiologist you're like the coyote that's already over the edge of the cliff. ***People should stop training radiologists now****, it's just completely obvious that in five years deep learning is going to do better than radiologists.*
It might be ten years"

Geoffrey Hinton in 2016

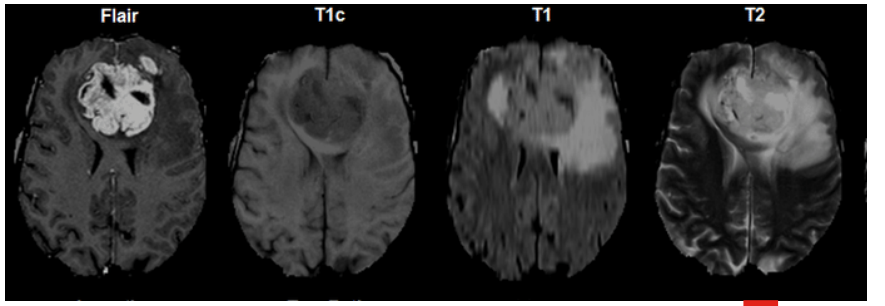

Automated

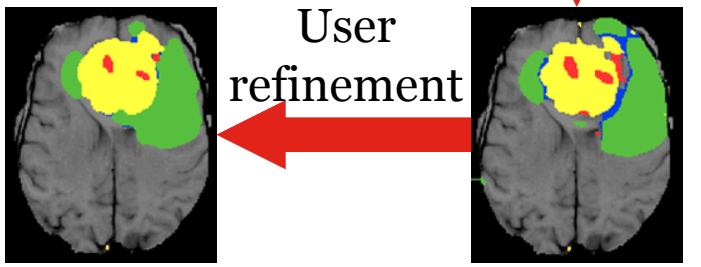

User refinement

Unknown ground truth

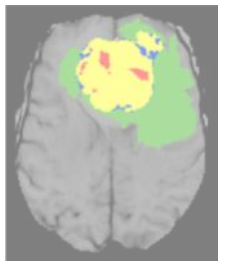

Dealing with hetero-modal inputs

Missing input modalities: the hetero-modality challenge

Goal: building a joint segmentation model for hetero-modal inputs

Possible solutions to deal with missing modalities

- Input filled with zeros

- Synthesising missing 3D modalities:
  → Generative models (ex: GAN, VAE)

- Creating a common feature space that encodes the shared information
  → Arithmetic operations (moments)
  → Hetero-modal variational encoder-decoder
  Treat the segmentation as a missing modality

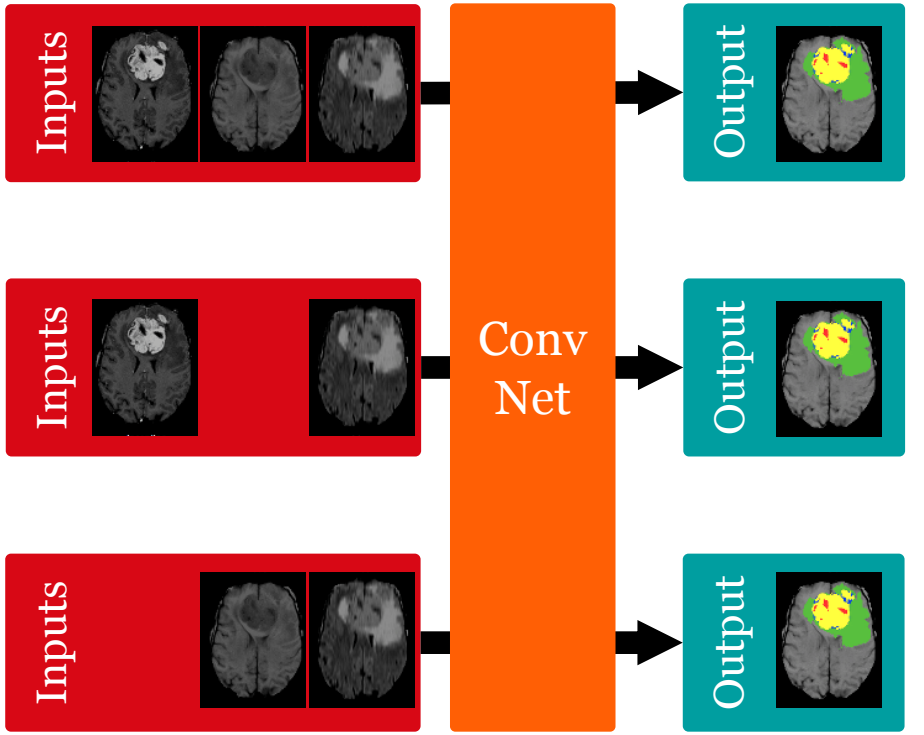

Hetero-modal architecture via arithmetic fusion

- Extract features from each modality independently
- Fuse *modality-agnostic* latent features using arithmetic operations (e.g. mean)

- Can we formalise the fusion in a probabilistic graphical model?

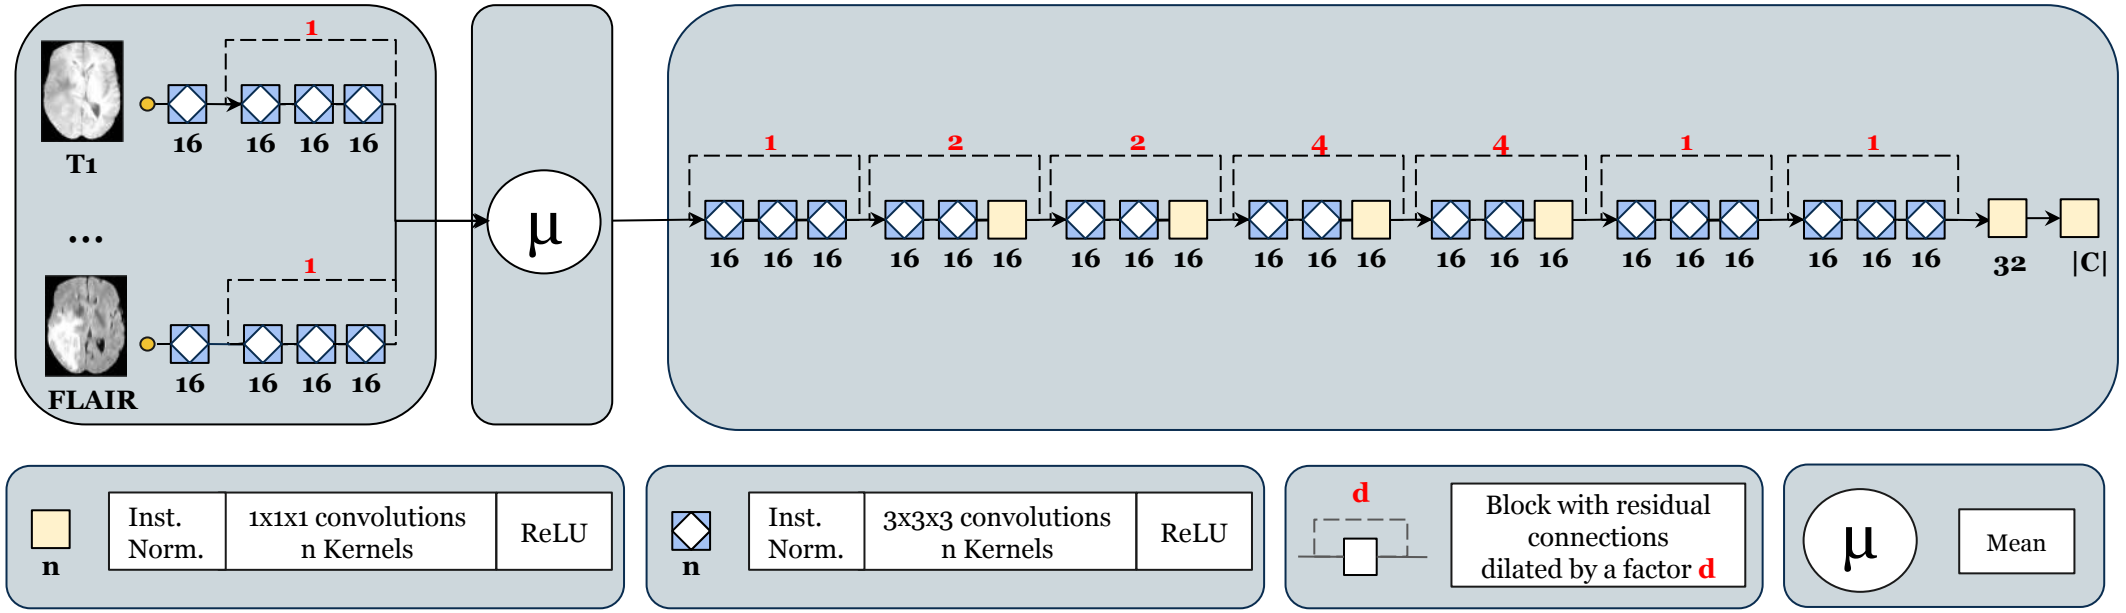

Havaei, *et al.* HeMIS: Hetero-modal image segmentation. MICCAI 2016.

Graphical model for multi-modal data auto-encoding

All modalities **and the segmentation** considered conditionally independent given latent features **z**

R Dorent et al. Hetero-modal variational encoder-decoder for joint modality completion and segmentation. Proc. MICCAI 2019.

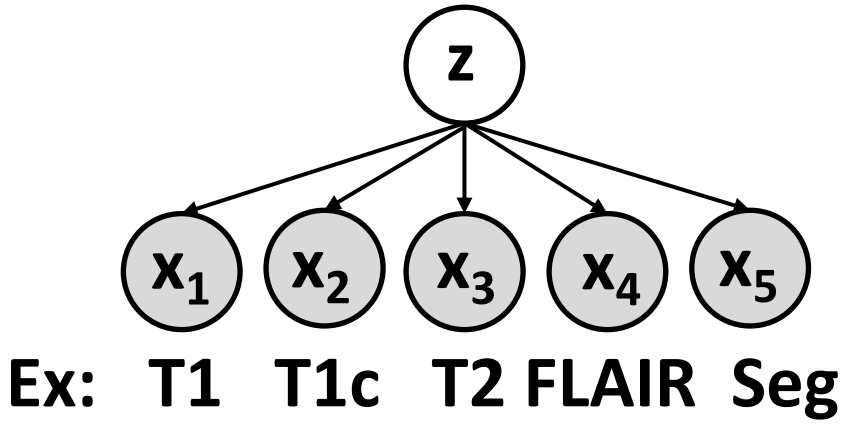

$$p_\theta(z, x_1, ..., x_n) = p(z) \prod_{i=1}^{n} p_\theta(x_i|z)$$

Prior (Typically modelled as Gaussian)

Parametrised with a neural network (**decoder** recons.)

Auto-encoder objective: maximize the marginal log-likelihood

$$\mathcal{L}(\mathbf{x};\theta) = \log(p_\theta(x_1, ..., x_n)) = \log\left(\int p_\theta(\mathbf{x}|z)p(z)\right)$$

intractable

Variational auto-encoder

Introduction of another parametric distribution (**encoder** provides parameters). Typically:

$$q_\phi(z|\mathbf{x}) = \mathcal{N}(z; \mu_\phi(\mathbf{x}), \Sigma_\phi(\mathbf{x}))$$

For any distribution $q_\phi(z|\mathbf{x})$ one can exploit the following lower-bound of the log-likelihood:

$$\mathcal{L}(\mathbf{x};\theta) \geq \text{ELBO}(\mathbf{x};\theta,\phi) \triangleq E_{q_\phi(z|\mathbf{x})_z}[\log(p_\theta(\mathbf{x}|z))] - \text{KL}[q_\phi(z|\mathbf{x})||p(z)]$$

Reconstruction loss
Estimated by sampling from Gaussian $q_\phi(z|\mathbf{x})$

Regularization term
Closed form between parameters (Gaussians)

ELBO maximal w.r.t. ϕ when $q_\phi(z|\mathbf{x}) = p_\theta(z|\mathbf{x})$

Kingma, D.P., Welling, M. Auto-encoding variational Bayes. ICLR 2014

Multi-modal variational encoder-decoder

Generalising multi-modal auto-encoder to missing modalities (including segmentation) by decomposing the latent posteriors as modality-specific latent posteriors

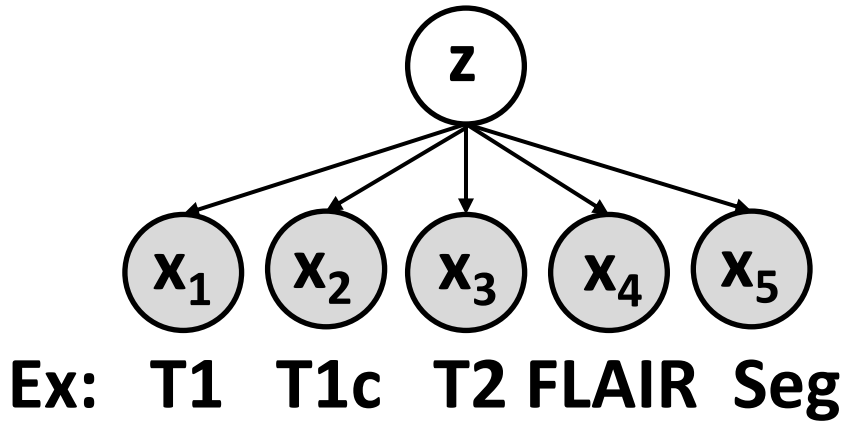

$$p_\theta(z|\mathbf{x}) \propto p(z) \prod_{i=1}^{n} \frac{p_\theta(z|x_i)}{p(z)}$$

$$q_\phi(z|\mathbf{x}) \propto p(z) \prod_{i=1}^{n} q_{\phi_i}(z|x_i)$$

Prior

Use modality-specific probabilistic encoder

Assume all modality-specific variational posteriors modeled with Gaussians:

$$q_i(z|x_i) = \mathcal{N}(z; \mu_{\phi_i}(x_i), \Sigma_{\phi_i}(x_i)) \qquad p(z) = \mathcal{N}(0, I)$$

$q_\phi(z|\mathbf{x})$ is Gaussian with mean and covariance defined by a closed-form formula:

$$\Sigma_\phi = (I + \sum_i \Sigma_{\phi_i}^{-1})^{-1} \text{ and } \mu_\phi = \Sigma_\phi^{-1}(\sum_i \Sigma_{\phi_i}^{-1} \mu_{\phi_i})$$ **Principled fusion operation**

Hetero-modal variational encoder-decoder

Let $\mathcal{P}$ denote the set of all possible non-empty combinations of the *n* modalities.

Goal: optimizing the log-likelihood when z encoded with a random subset $\pi \in \mathcal{P}$ drawn with a probability α_π

→ we model $q_\phi(z|\mathbf{x})$ as a mixture of Gaussian

$$q_\phi(z|\mathbf{x}) = \sum_{\pi \in \mathcal{P}} \alpha_\pi q_\phi^\pi(z|\mathbf{x}_\pi)$$

→ Given the convexity of the KL divergence and $\sum_{\pi \in \mathcal{P}} \alpha_\pi = 1$

$$\mathcal{L}(\mathbf{x};\theta,\phi) \geq \sum_{\pi \in \mathcal{P}} \alpha_\pi \underbrace{(E_{q_\phi^\pi(z|\mathbf{x}_\pi)}[\log(p_\theta(\mathbf{x}|z))] - \mathrm{KL}[q_\phi^\pi(z|\mathbf{x}_\pi)||p(z)])}_{\mathrm{ELBO}\pi(\mathbf{x})}$$

Multi-level hetero-modal variational encoder-decoder

① All **observed** (non-missing) modalities encoded **independently**
② Modality-specific features maps combined to produce a **multi-resolution** hidden variable
③ All modalities **and segmentation** independently decoded

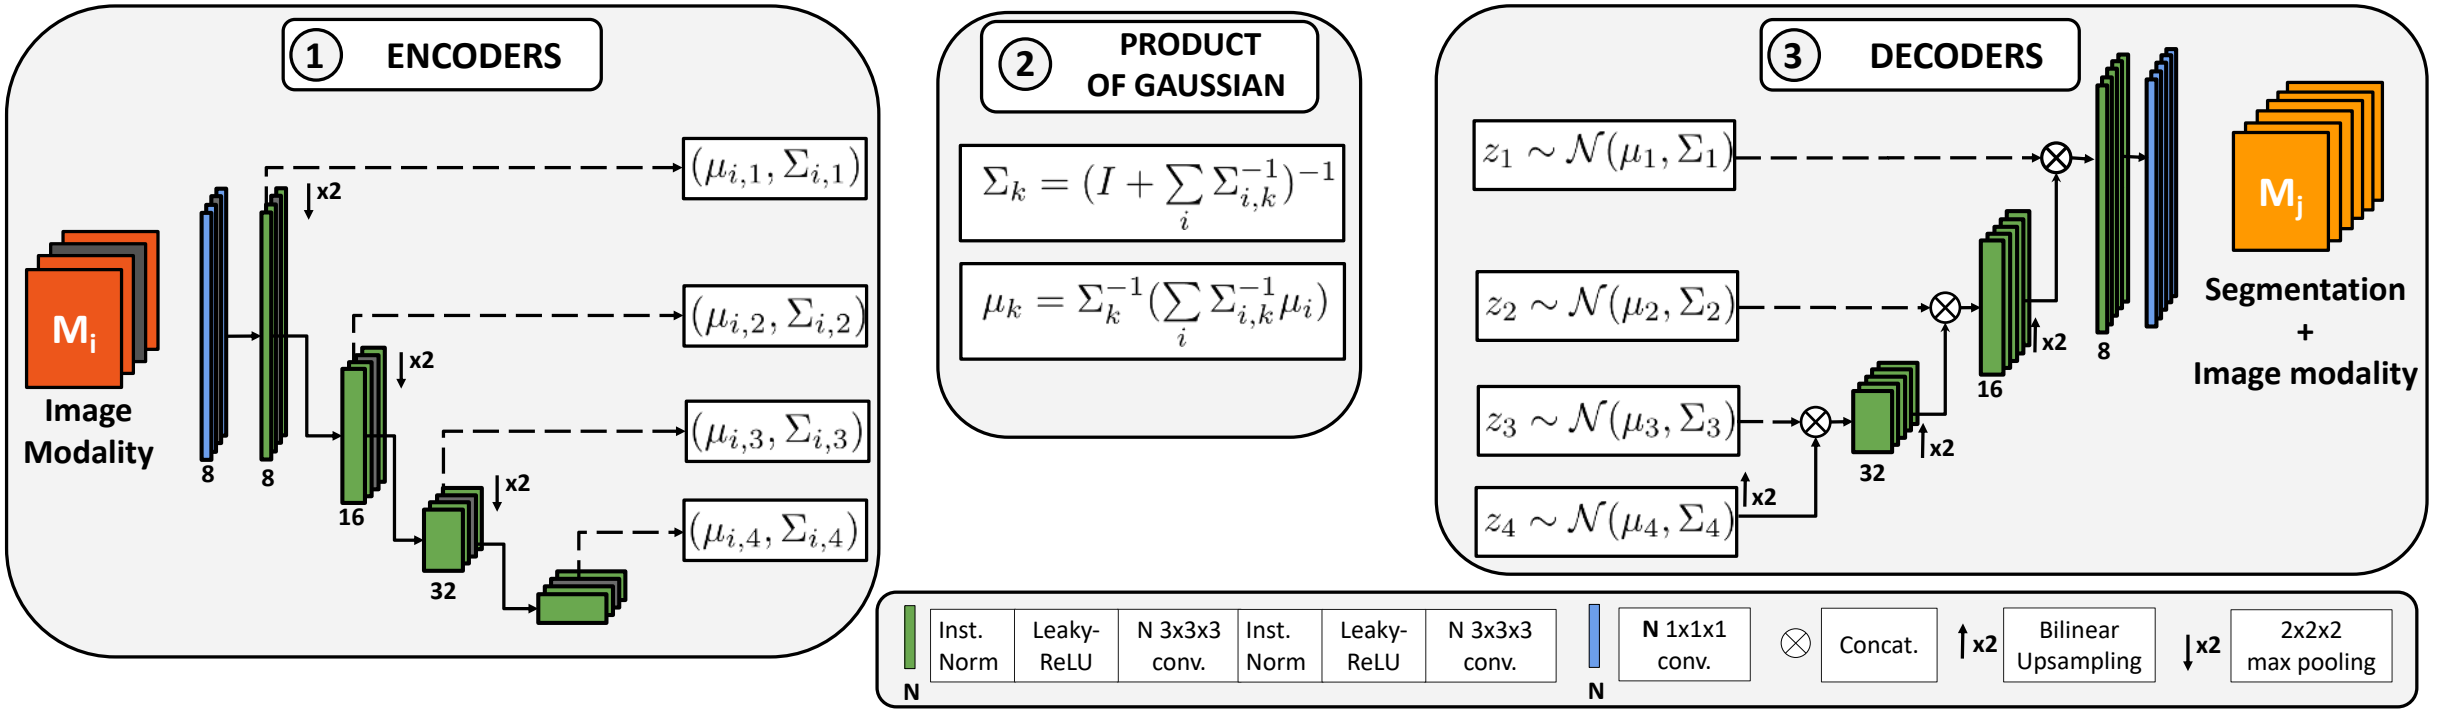

Example of outputs (segmentation and reconstructions)

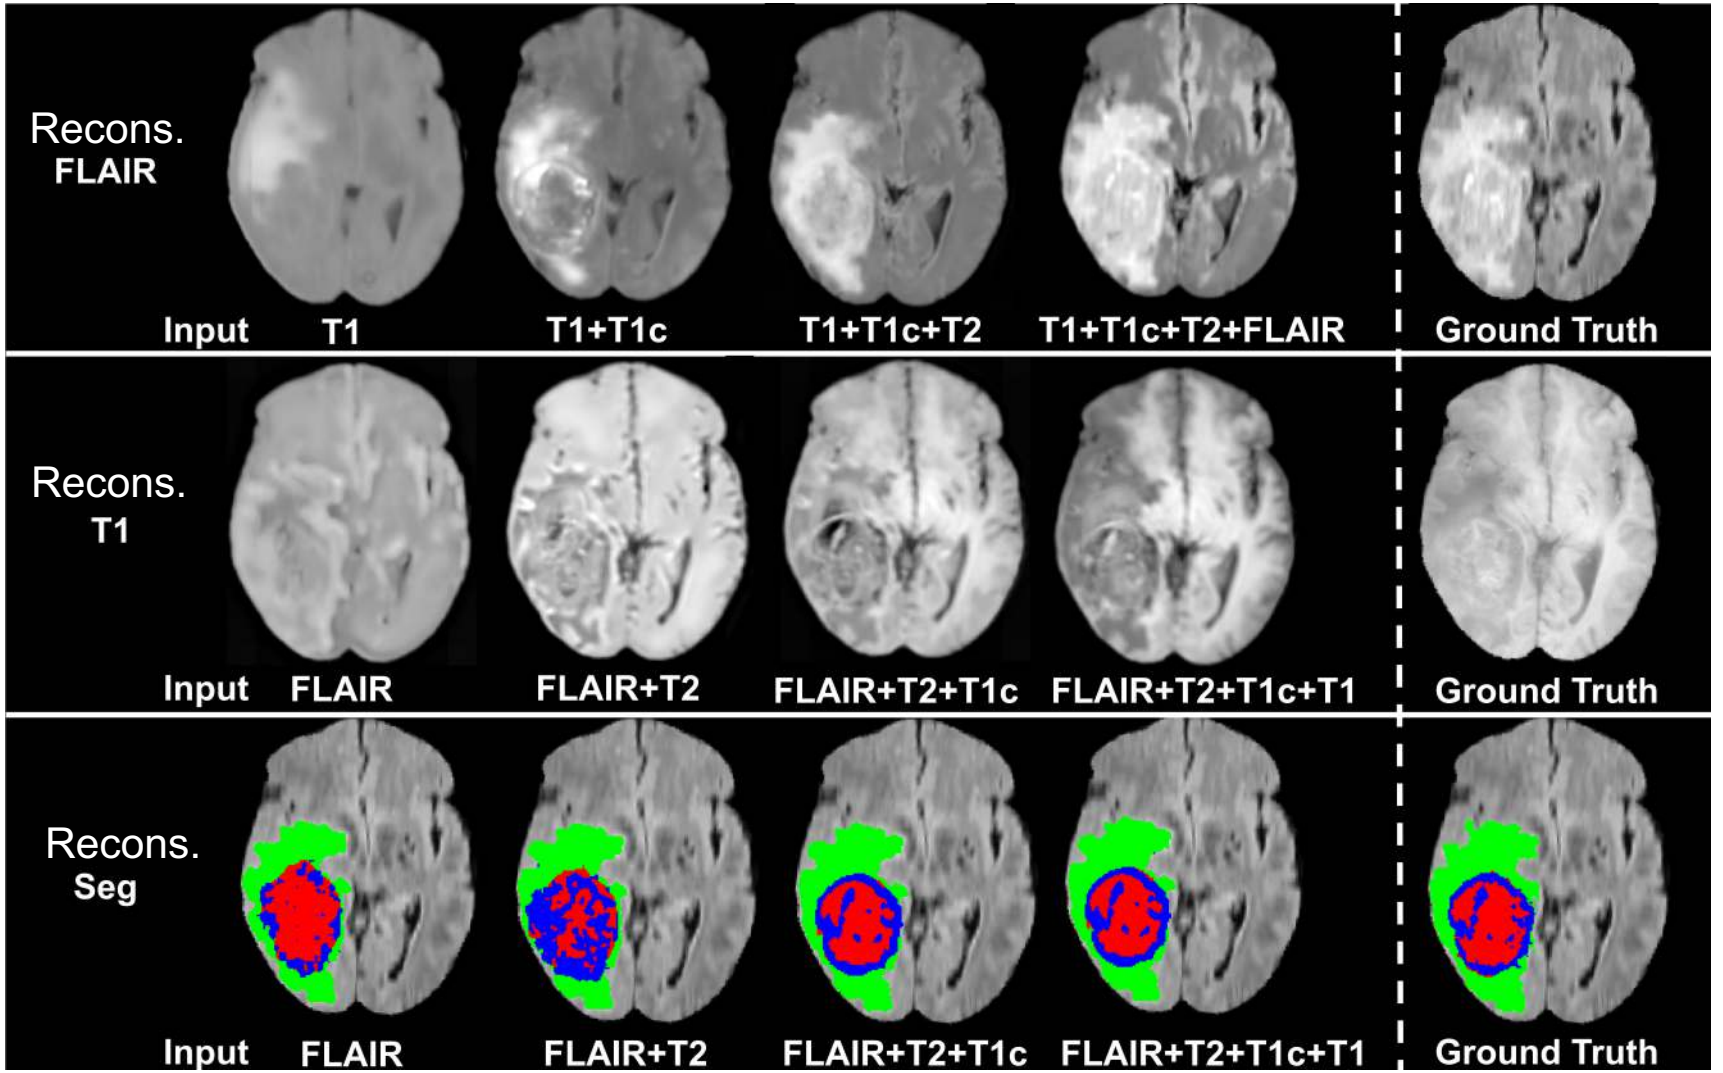

Comparing the hetero-modal architectures

- **Models comparison:**
  - **HeMIS**: architecture in [2] fusing modality-specific maps via arithmetic operations.
  - **U-HeMIS**: U-Net + modality specific features fused with first moment.
  - **U-HVED**: the hetero-modal multi-level variational auto-encoder
  - **Sing**: U-Net (1 encoder, 1 decoder) trained for each modality specific subset (i.e. 15 models)

Dice (%) for combinations of available (●) and missing (○) modalities. * significant improvement – Wilcoxon test p<0.05

Modalities				Complete				Core				Enhancing			
F	T_1	T_1c	T_2	HeMIS	U-HeMIS	U-HVED	Sing	HeMIS	U-HeMIS	U-HVED	Sing	HeMIS	U-HeMIS	U-HVED	Sing
○	○	○	●	38.6	79.2	**80.9***	82.6	19.5	50.0	**54.1***	54.9	0.0	23.3	**30.8***	34.2
○	○	●	○	2.6	58.5	**62.4***	70.4	6.5	58.5	**66.7***	71.5	11.1	60.8	**65.5***	70.4
○	●	○	○	0.0	**54.3***	52.4	72.7	0.0	**37.9**	37.2	59.2	0.0	12.4	**13.7***	32.2
●	○	○	○	55.2	79.9	**82.1***	81.5	16.2	49.8	**50.4**	55.5	6.6	**24.9**	24.8	26.3
○	○	●	●	48.2	81.0	**82.7***	83.2	45.8	69.1	**73.7***	73.3	55.8	68.6	**70.2***	70.1
○	●	●	○	15.4	63.8	**66.8***	70.6	30.4	64.0	**69.7***	73.9	42.6	65.3	**67.0***	71.9
●	●	○	○	71.1	83.9	**84.3**	83.3	11.9	**56.7***	55.3	54.3	1.2	**29.0***	24.2	30.7
○	●	○	●	47.3	80.8	**82.2***	83.1	17.2	53.4	**57.2***	59.7	0.6	28.3	**30.7***	33.4
●	○	○	●	74.8	86.0	**87.5***	86.3	17.7	58.7	**59.7**	57.7	0.8	28.0	**34.6***	31.0
●	○	●	○	68.4	83.3	**85.5***	85.3	41.4	67.6	**72.9***	72.0	53.8	68.0	**70.3***	69.9
●	●	●	○	70.2	85.1	**86.2***	85.1	48.8	70.7	**74.2***	74.9	60.9	69.9	**71.1**	70.1
●	●	○	●	75.2	87.0	**88.0***	85.7	18.7	61.0	**61.5**	57.9	1.0	33.4	**34.1**	34.1
●	○	●	●	75.6	87.0	**88.6***	85.8	54.9	72.2	**75.6***	75.2	60.5	69.7	**71.2***	72.2
○	●	●	●	44.2	82.1	**83.3***	81.5	46.6	70.7	**75.3***	74.7	55.1	69.7	**71.1***	71.1
●	●	●	●	73.8	87.6	**88.8***	87.5	55.3	73.4	**76.4***	78.4	61.1	70.8	**71.7***	72.7
Means				50.7	78.6	**80.1***	81.6	28.7	59.7	**64.0***	66.2	27.4	48.1	**50.0***	52.7

Dealing with heterogenous labelling

Learning a joint model for tissue and lesion segmentation

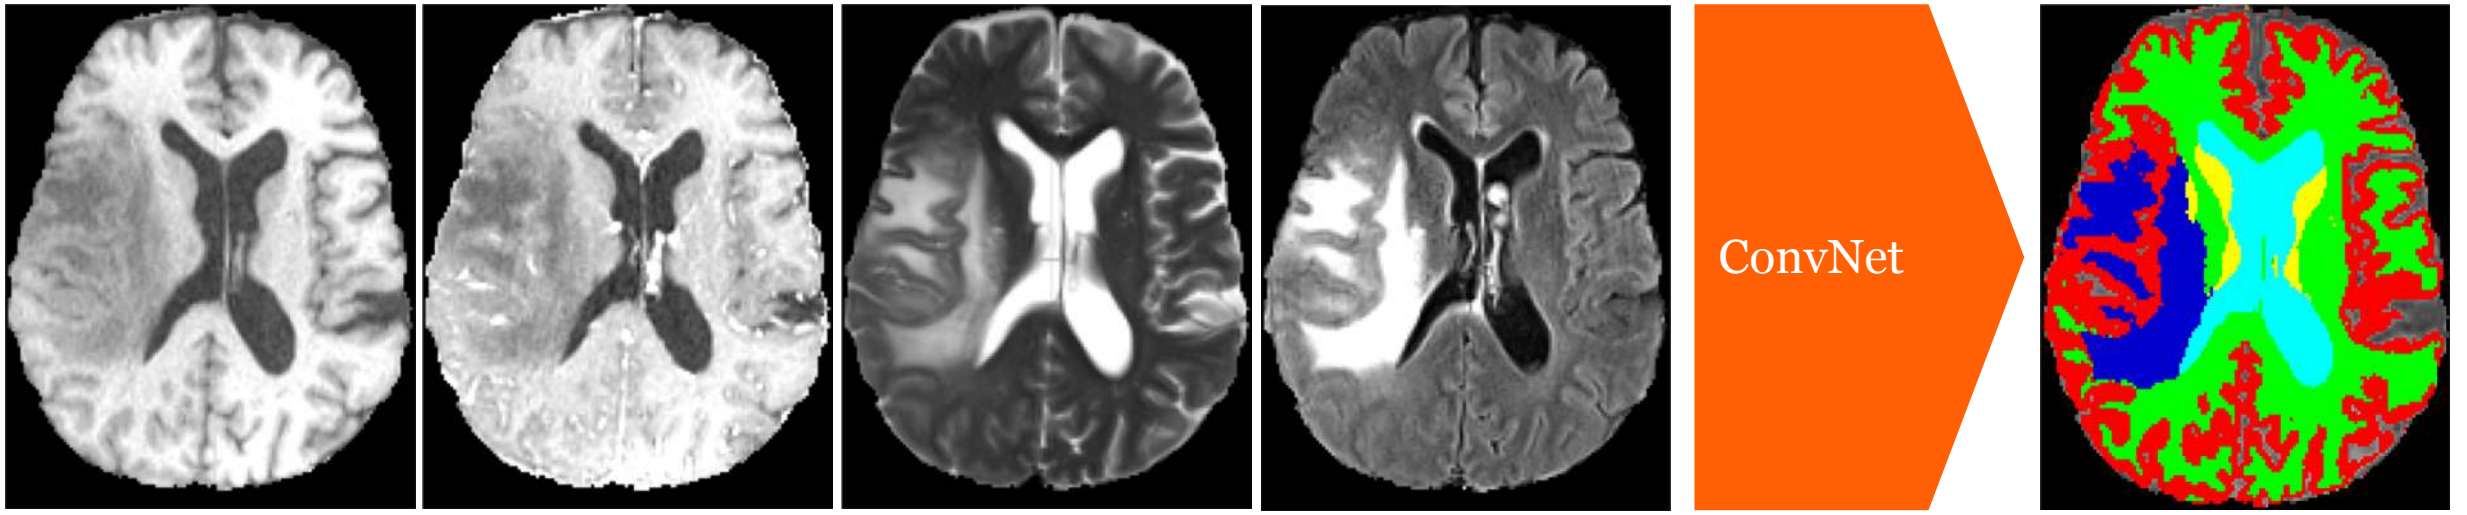

- Challenge: No large fully annotated dataset for joint tissue and lesion segmentation

- Can we leverage task-specific datasets?
  - Maybe, but these are also often hetero-modal
  - May even display domain shift in the common modalities (Not covered in this lecture)

Learning from task-specific hetero-modal and domain-shifted data

	Neuromorphometrics	**BraTS18**
	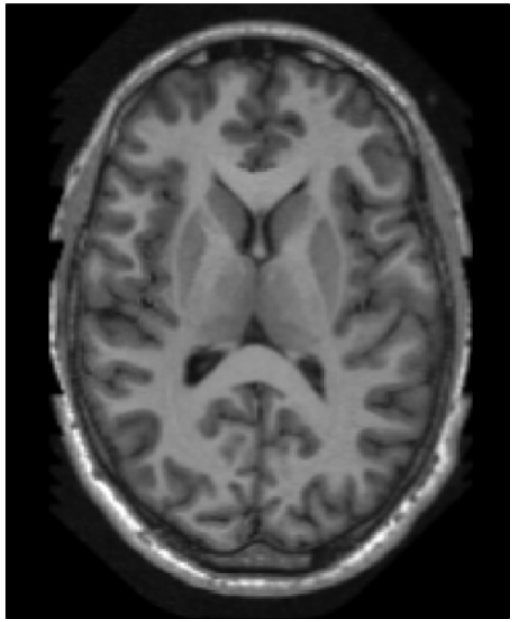	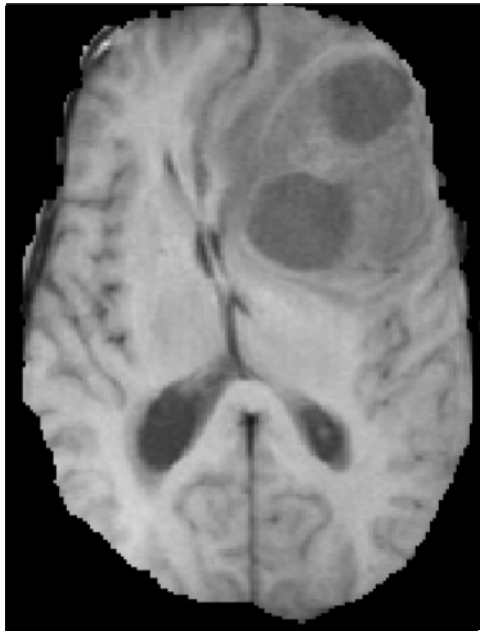
Available modalities:	T1	T1, T1c, T2, FLAIR
Illustrated modalities:	T1 Protocol 1	T1 Protocol 2
Hereby referred to as:	*Control*	*Lesion*

Loss decomposition: Tissue & lesion

R Dorent et al. Learning joint lesion and tissue segmentation from task-specific hetero-modal datasets. PMLR 102:164-174, 2019

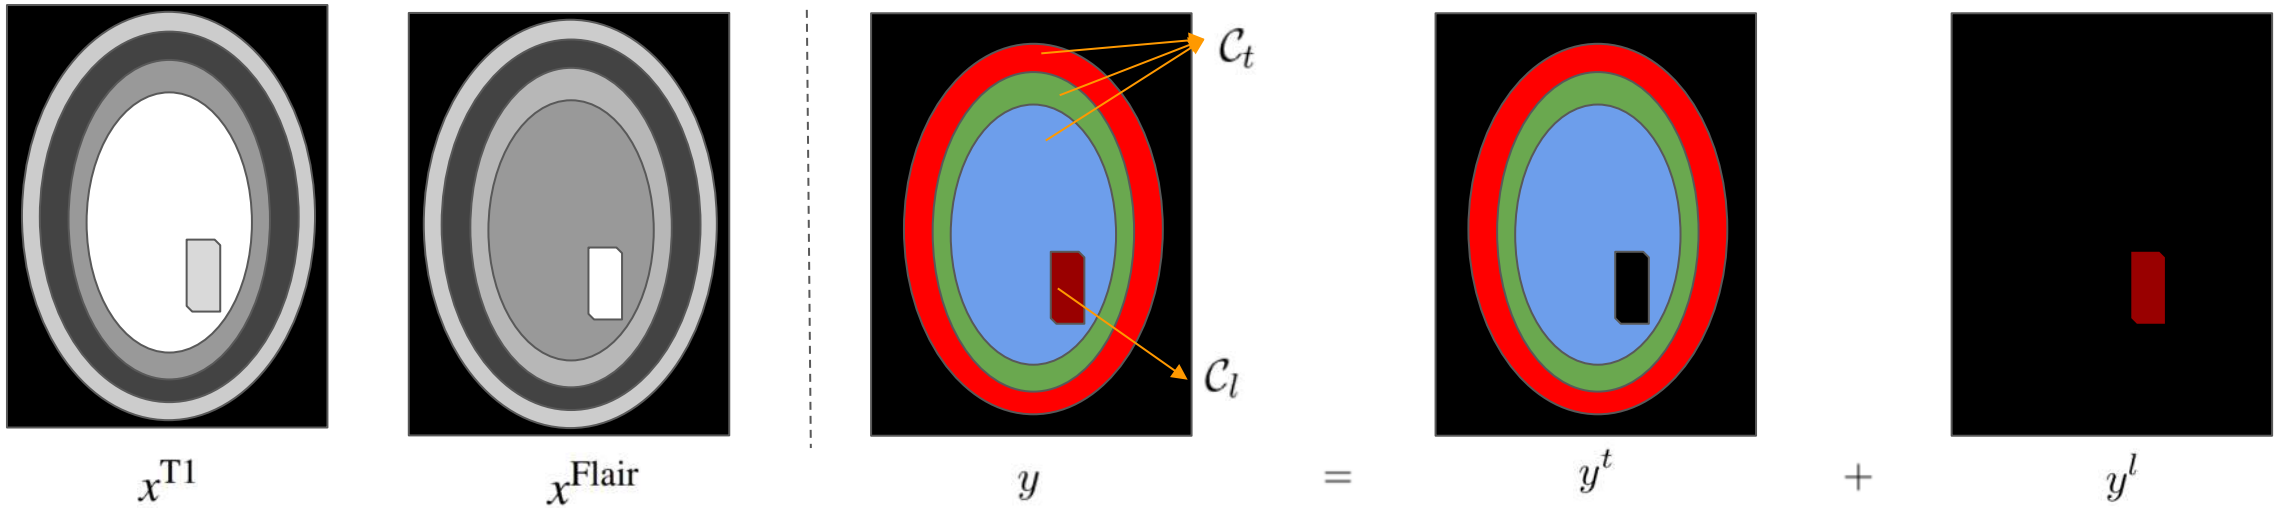

$$\mathcal{L}(h_\theta(x^{T1}, x^{Flair}), y) = \mathcal{L}^t(h_\theta(x^{T1}, x^{Flair}), y^t) + \mathcal{L}^l(h_\theta(x^{T1}, x^{Flair}), y^l)$$

CNN output $h_\theta(x^{T1}, x^{Flair})$ **Joint loss** **Tissue loss** **Lesion loss**

Decomposed loss and task-specific datasets

Focus on data distribution from the patient group: $(X,Y) \sim \mathcal{D}_{lesion}$

$$\theta^* = argmin_\theta \underbrace{E_{\mathcal{D}_{lesion}}[\mathcal{L}^t(h_\theta(X^{T1}, X^{Flair}), Y^t)]}_{\mathcal{R}^t(\theta)} + \underbrace{E_{\mathcal{D}_{lesion}}[\mathcal{L}^l(h_\theta(X^{T1}, X^{Flair}), Y^l)]}_{\mathcal{R}^l(\theta)}$$

Missing tissue annotations
Could we use the *Control* dataset?

Estimation using the *Lesion* dataset

Tissue loss upper-bound for missing modalities

$$\underbrace{E_{\mathcal{D}_{lesion}}[\mathcal{L}^t(h_\theta(X^{T1}, X^{Flair}), Y^t)]}_{\mathcal{R}^t(\theta)}$$

Use a loss satisfying the triangle inequality (e.g. Jaccard)

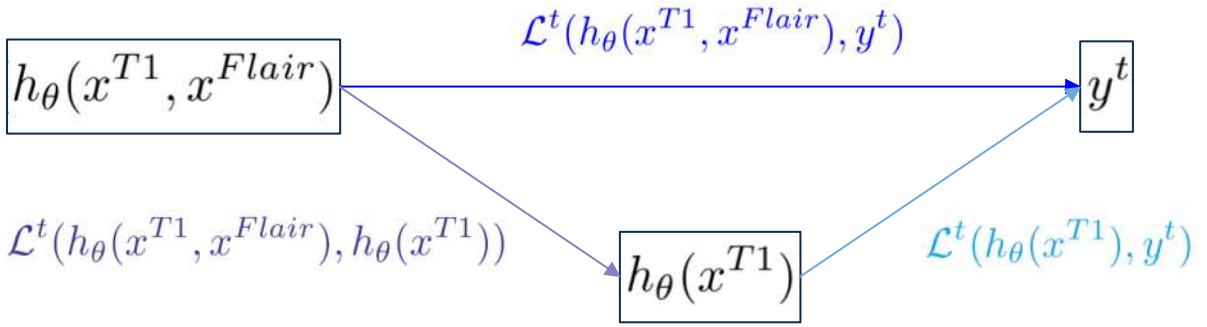

$$\mathcal{R}^t(\theta) \leq \underbrace{E_{(X,Y)\sim\mathcal{D}_{lesion}}[\mathcal{L}^t(h_\theta(X^{T1}, X^{Flair}), h_\theta(X^{T1}))]}_{\mathcal{R}_1^t(\theta)} + \underbrace{E_{(X,Y)\sim\mathcal{D}_{lesion}}[\mathcal{L}^t(h_\theta(X^{T1}), Y^t)]}_{\mathcal{R}_2^t(\theta)}$$

Hetero-modal consistency

Missing tissue annotations
Could we use the *Control* dataset?

Exploiting control data for tissue loss

Tissue loss only evaluated on healthy tissues

Scenario 1: T1 scans have been acquired with a similar setting

Intensity distribution similar on non-lesion parts of the brain

$$P_{\mathcal{D}_{lesion}}(x_i^{T_1}, y_i | y_i \in \mathcal{C}_T) = P_{\mathcal{D}_{control}}(x_i^{T_1}, y_i | y_i \in \mathcal{C}_T)$$

Assume the network outputs similar predictions on healthy tissues

$$\underbrace{E_{(X,Y)\sim\mathcal{D}_{lesion}}[\mathcal{L}^t(h_\theta(X^{T1}), Y^t)]}_{\mathcal{R}_2^t(\theta)} = \underbrace{E_{(X,Y)\sim\mathcal{D}_{control}}[\mathcal{L}^t(h_\theta(X^{T1}), Y^t)]}_{\mathcal{R}_2^t(\theta)}$$

Estimation using the *Control* dataset

Heterogeneous labelling learning summary

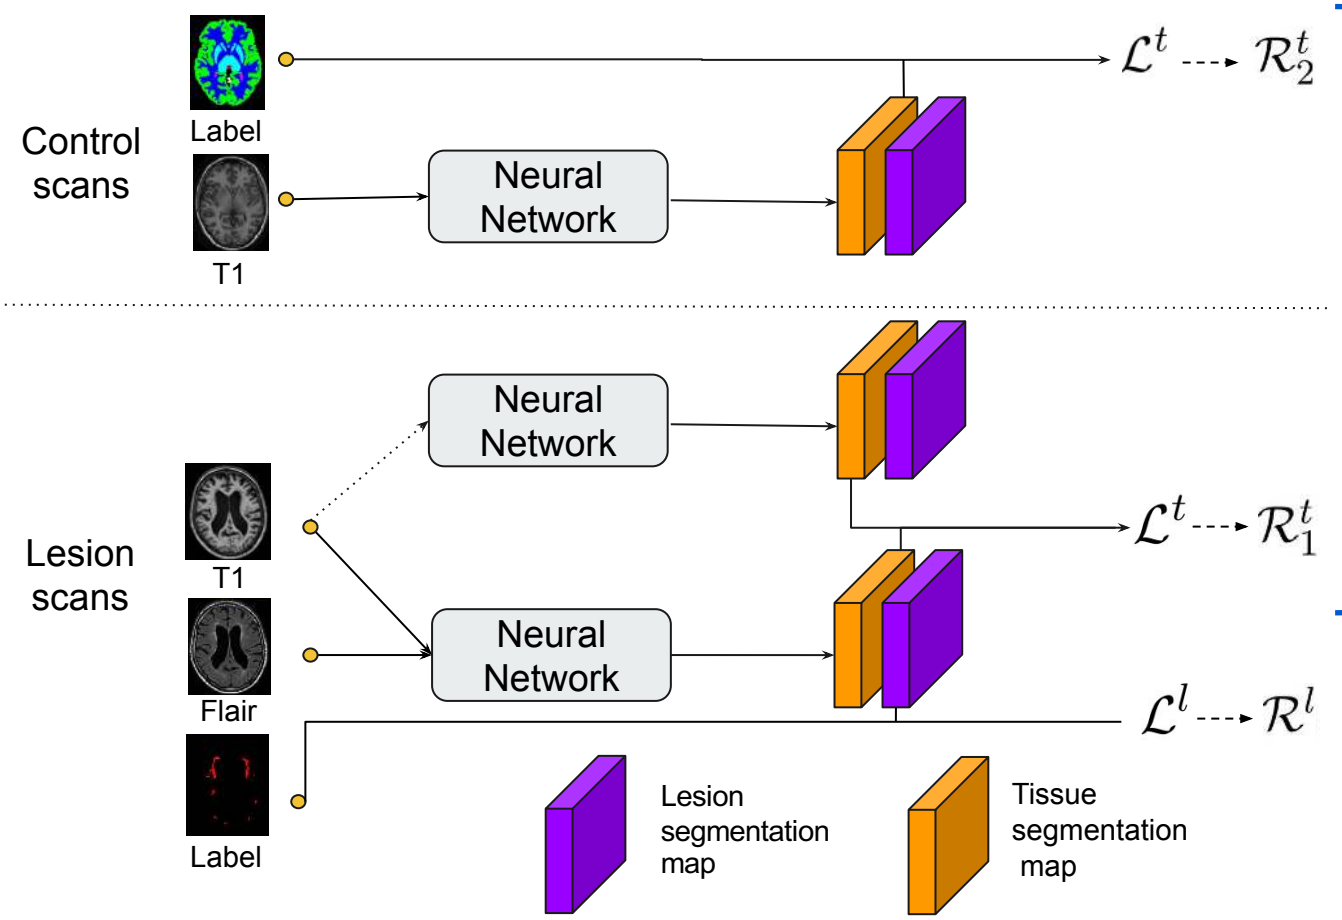

Tissue loss

***Control* tissue loss (supervised)**

Hetero-modal consistency (unsupervised)

Lesion loss (supervised)

White matter lesions: Datasets

Classes: White Matter Lesion + GM, WM, Cerebellum, Basal Ganglia, Brain Stem, Ventricles

Neuromorphometrics

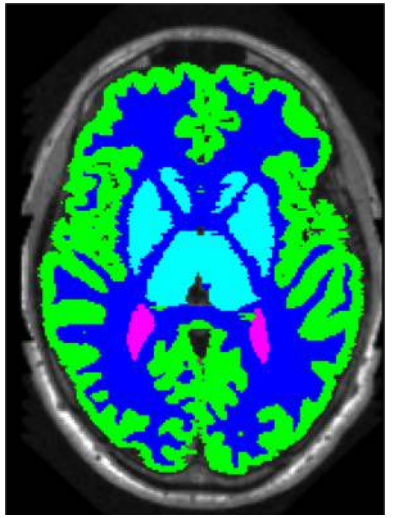

T1

*28+32 = 60**

WMH

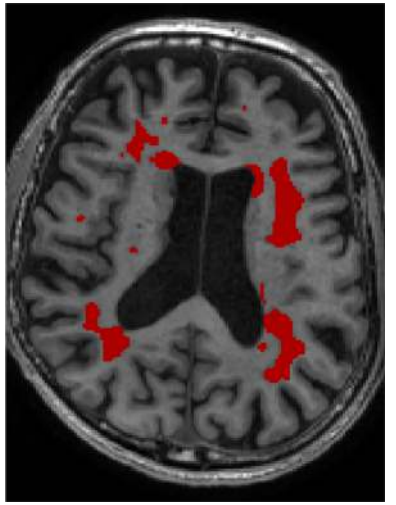

T1, Flair

60

MRBrainS18

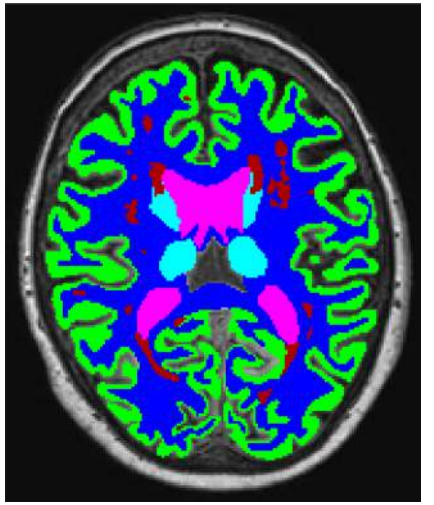

T1, Flair

7

[*] Cardoso, *et al.* Geodesic information flows: Spatially-variant graphs and their application to segmentation and fusion. 2015

Model comparison

Joint models:

- ***FS-MRB***: joint model trained on the 7 training scans of MRBrainS18 (T1+Flair)
- ***jSTABL***: our model trained on WMH and Neuromophometrics (T1+Flair)

Task-specific models:

- ***Tissue***: tissue segmentation model trained on Neuromorphometrics (T1)
- ***Lesion***: lesion segmentation model trained on WMH (T1+Flair)
- ***Pipeline = Tissue + Lesion***

Evaluation of our framework (jSTABL) on patients with White Matter Lesion (Bronze standard for WMH tissues)
Means and standard deviations for Dice scores

Classes	Neuromorphometrics			WMH			MRBrainS18		
	Tissue	FS-MRB	jSTABL	Pipeline	FS-MRB	jSTABL	SPM	FS-MRB	jSTABL
Grey matter	91.8 (1.6)	28.0 (4.4)	90.9 (1.5)	**84.6 (3.2)**	52.3 (21.8)	83.4 (1.9)	76.5	81.1 (2.0)	78.6 (3.1)
White mater	93.1 (1.6)	58.2 (6.3)	92.7 (1.3)	90.3 (1.9)	55.1 (21.5)	**90.4 (1.4)**	75.7	83.8 (2.1)	84.6 (2.4)
Brainstem	93.3 (0.9)	26.5 (5.4)	93.2 (0.8)	89.6 (3.9)	29.3 (34.8)	**91.4 (2.6)**	76.5	88.4 (1.8)	72.5 (1.7)
Basal ganglia	88.4 (2.5)	29.8 (6.1)	87.8 (2.1)	80.2 (7.2)	39.7 (20.9)	**81.9 (3.9)**	74.7	80.2 (2.9)	78.7 (2.7)
Ventricles	89.3 (4.3)	6.4 (5.0)	88.7 (4.6)	91.1 (6.0)	62.2 (25.6)	**94.0 (2.3)**	80.9	91.2 (3.9)	92.0 (2.5)
Cerebellum	95.7 (0.8)	18.6 (6.9)	95.4 (1.1)	93.3 (2.0)	30.7 (39.8)	**93.7 (1.9)**	89.4	89.8 (2.5)	89.9 (2.1)
White matter lesion				**76.4 (8.7)**	37.4 (26.9)	**76.4 (8.7)**	40.8	52.9 (22.5)	57.3 (23.5)

Brain tumour: Datasets

Classes: Whole tumour, Core tumour, Enhancing tumour + GM, WM, Cerebellum, Basal Ganglia, Brain Stem, Ventricles

Neuromorphometrics

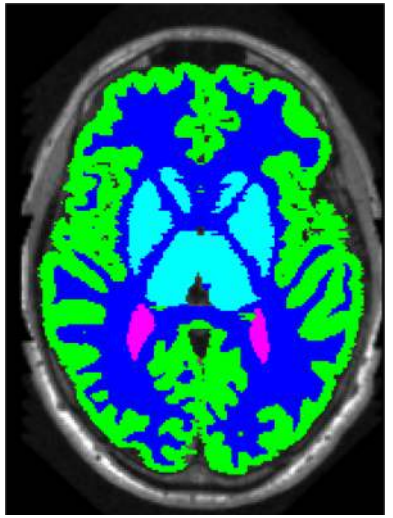

T1
*28+32 = 60**

BRaTS18

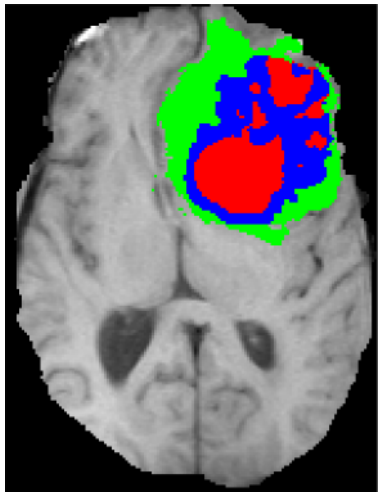

T1, T1-c, T2 Flair
285

BRaTS18 symmetrised

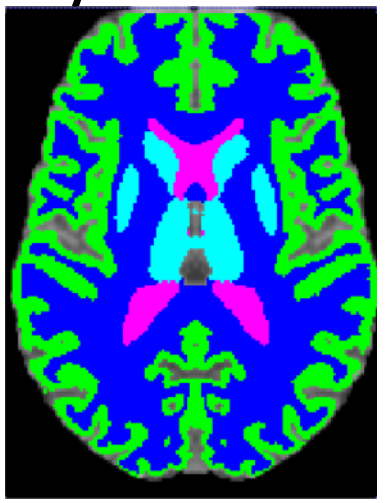

T1
129

[*] Cardoso, *et al.* Geodesic information flows: Spatially-variant graphs and their application to segmentation and fusion. 2015

[*] Cardoso, *et al.* Geodesic information flows: Spatially-variant graphs and their application to segmentation and fusion. 2015

Glioma: qualitative comparison

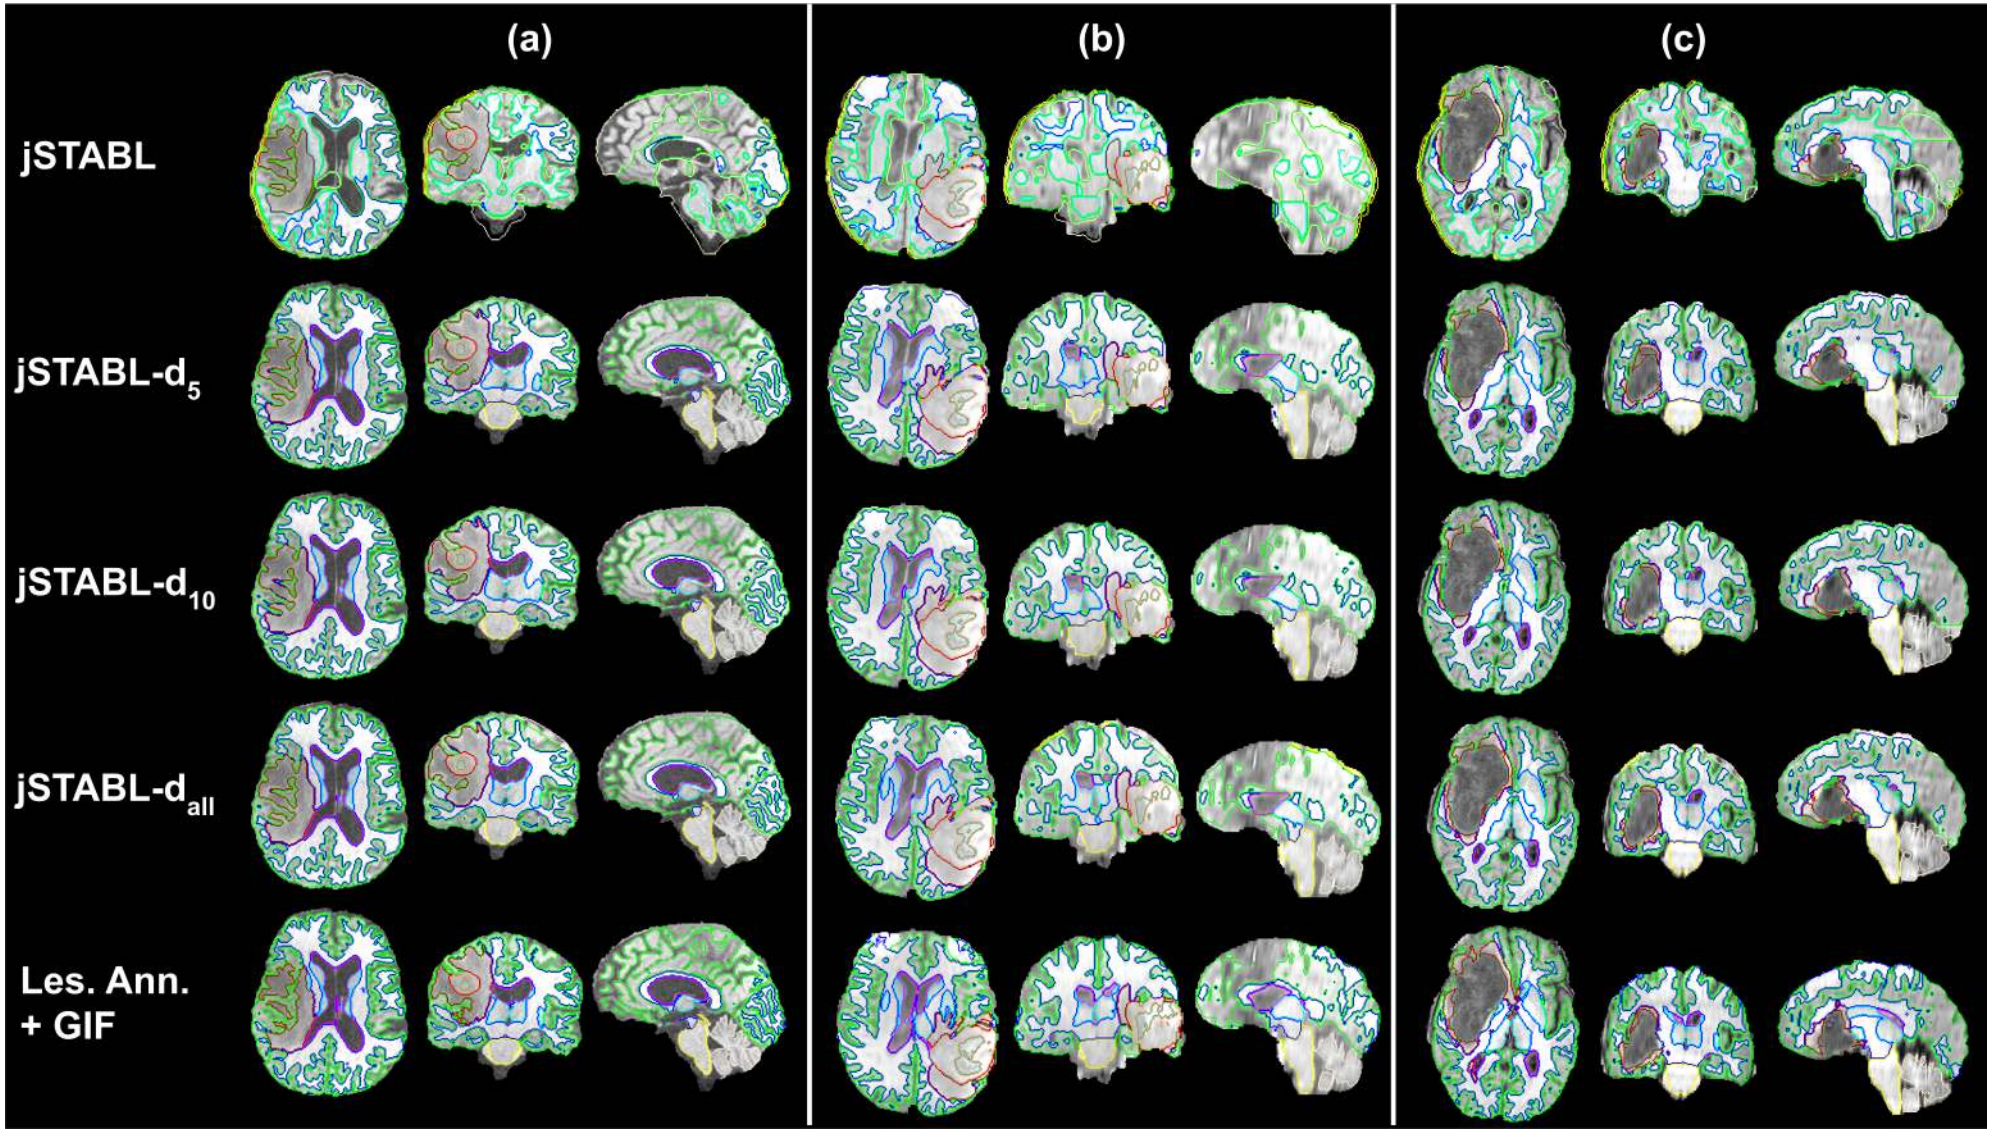

Model comparison

Joint models:

- ***jSTABL:*** joint model trained without Domain Adaptation (DA)
- ***jSTABL-d5:*** joint model trained with 5 symmetrised samples for DA
- ***jSTABL-d5:*** joint model trained with 10 symmetrised samples for DA
- ***jSTABL-dall:*** joint model trained with 129 symmetrised samples for DA

Evaluation of our framework (jSTABL) on patients with glioma in comparison.
Means and standard deviations for Dice scores.

Models	Neuromorphometrics				BRaTS18			
	jSTABL	jSTABL-d_5	jSTABL-d_{10}	jSTABL-d_{all}	jSTABL	jSTABL-d_5	jSTABL-d_{10}	jSTABL-d_{all}
Grey matter	93.1 (1.8)	92.3 (1.8)	92.2 (1.6)	92.6 (1.6)	67.0 (12.1)	87.1 (4.5)	86.9 (4.4)	87.8 (4.0)
Brainstem	92.9 (2.3)	91.6 (2.3)	92.4 (1.8)	89.8 (3.8)	0.0 (0.1)	80.5 (8.6)	86.2 (6.7)	81.9 (9.0)
Basal ganglia	86.6 (2.8)	83.7 (2.9)	85.4 (2.3)	85.8 (2.0)	2.1 (4.0)	73.6 (9.1)	76.8 (8.5)	80.7 (6.9)
Ventricules	89.5 (4.0)	88.5 (4.5)	88.8 (4.1)	89.0 (4.1)	0.3 (1.4)	91.4 (5.5)	91.5 (5.6)	92.0 (4.9)
Cerebellum	96.5 (1.6)	94.8 (2.1)	95.8 (1.3)	95.2 (1.7)	11.0 (8.4)	88.8 (6.2)	89.5 (5.0)	88.2 (5.3)
Whole tumour					86.3 (9.8)	86.1 (9.1)	86.3 (8.6)	86.4 (8.9)
Core tumour					78.3 (19.7)	75.2 (21.3)	76.4 (21.0)	76.8 (19.4)
Enhancing tumour					74.7 (20.2)	73.8 (21.2)	74.0 (21.0)	74.2 (20.4)

Dealing with user corrections

Interactive corrections pipeline

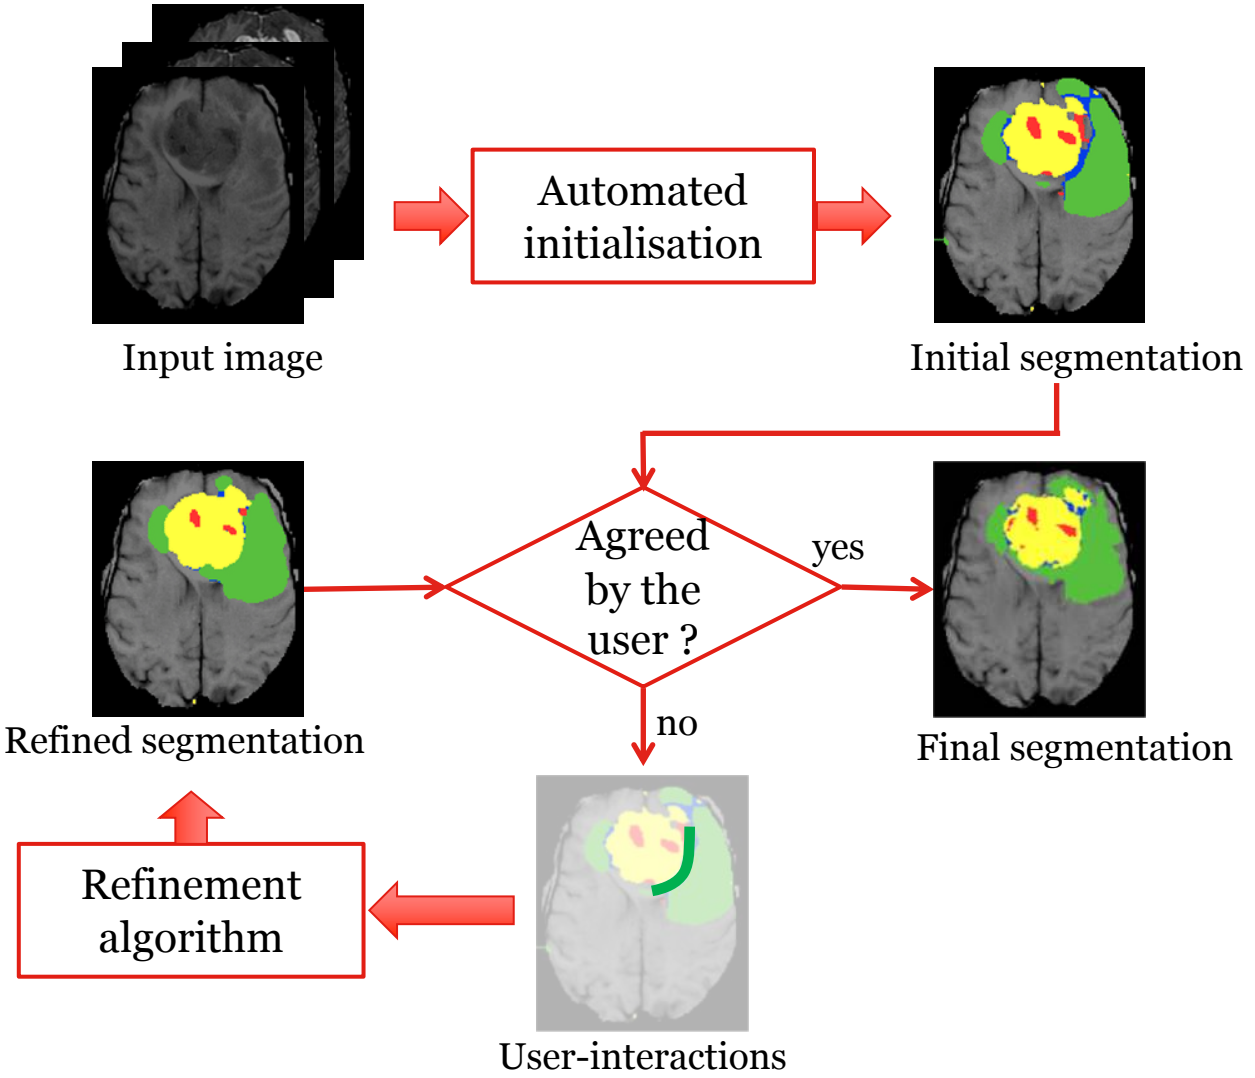

Classical geodesic interactive segmentation

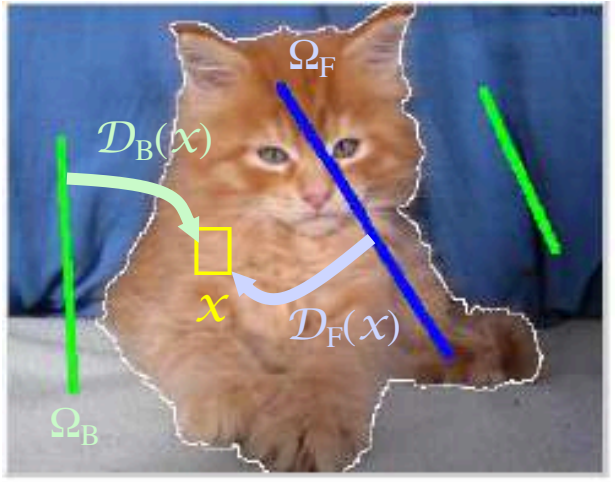

Input image with scribbles in green and blue
Segmentation in white

Scalar cost map (gradient magnitude, learned features, etc.)

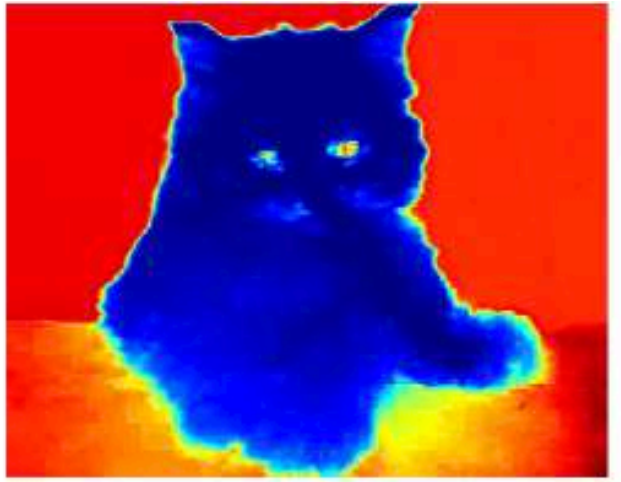

$\mathcal{D}_F(.)$ map

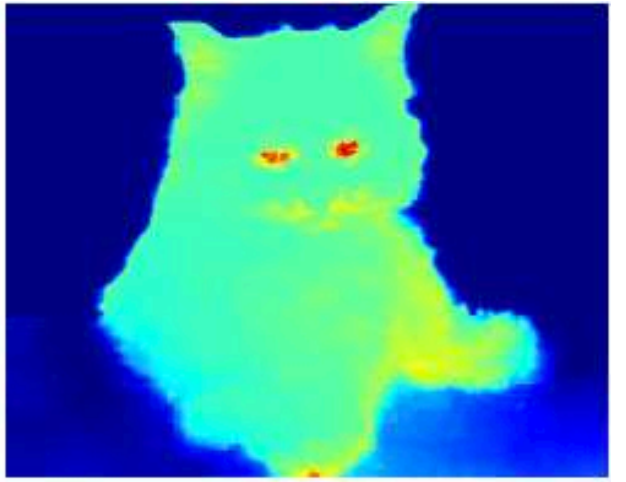

$\mathcal{D}_B(.)$ map

$$D_l(x) := \min_{s \in \Omega_l} d(s, x), \quad l \in \{\mathcal{F}, \mathcal{B}\},$$

where

$$d(s_1, s_2) := \min_{C_{s_1,s_2}} \int_{s_1}^{s_2} |W(x) \cdot \dot{C}_{s_1,s_2}(x)| dx,$$

Combining geodesic maps and deep learning

G Wang et al. DeepIGeoS: A Deep Interactive Geodesic Framework for Medical Image Segmentation. IEEE Trans Pattern Anal Mach Intell. 2019 Jul;41(7):1559-1572

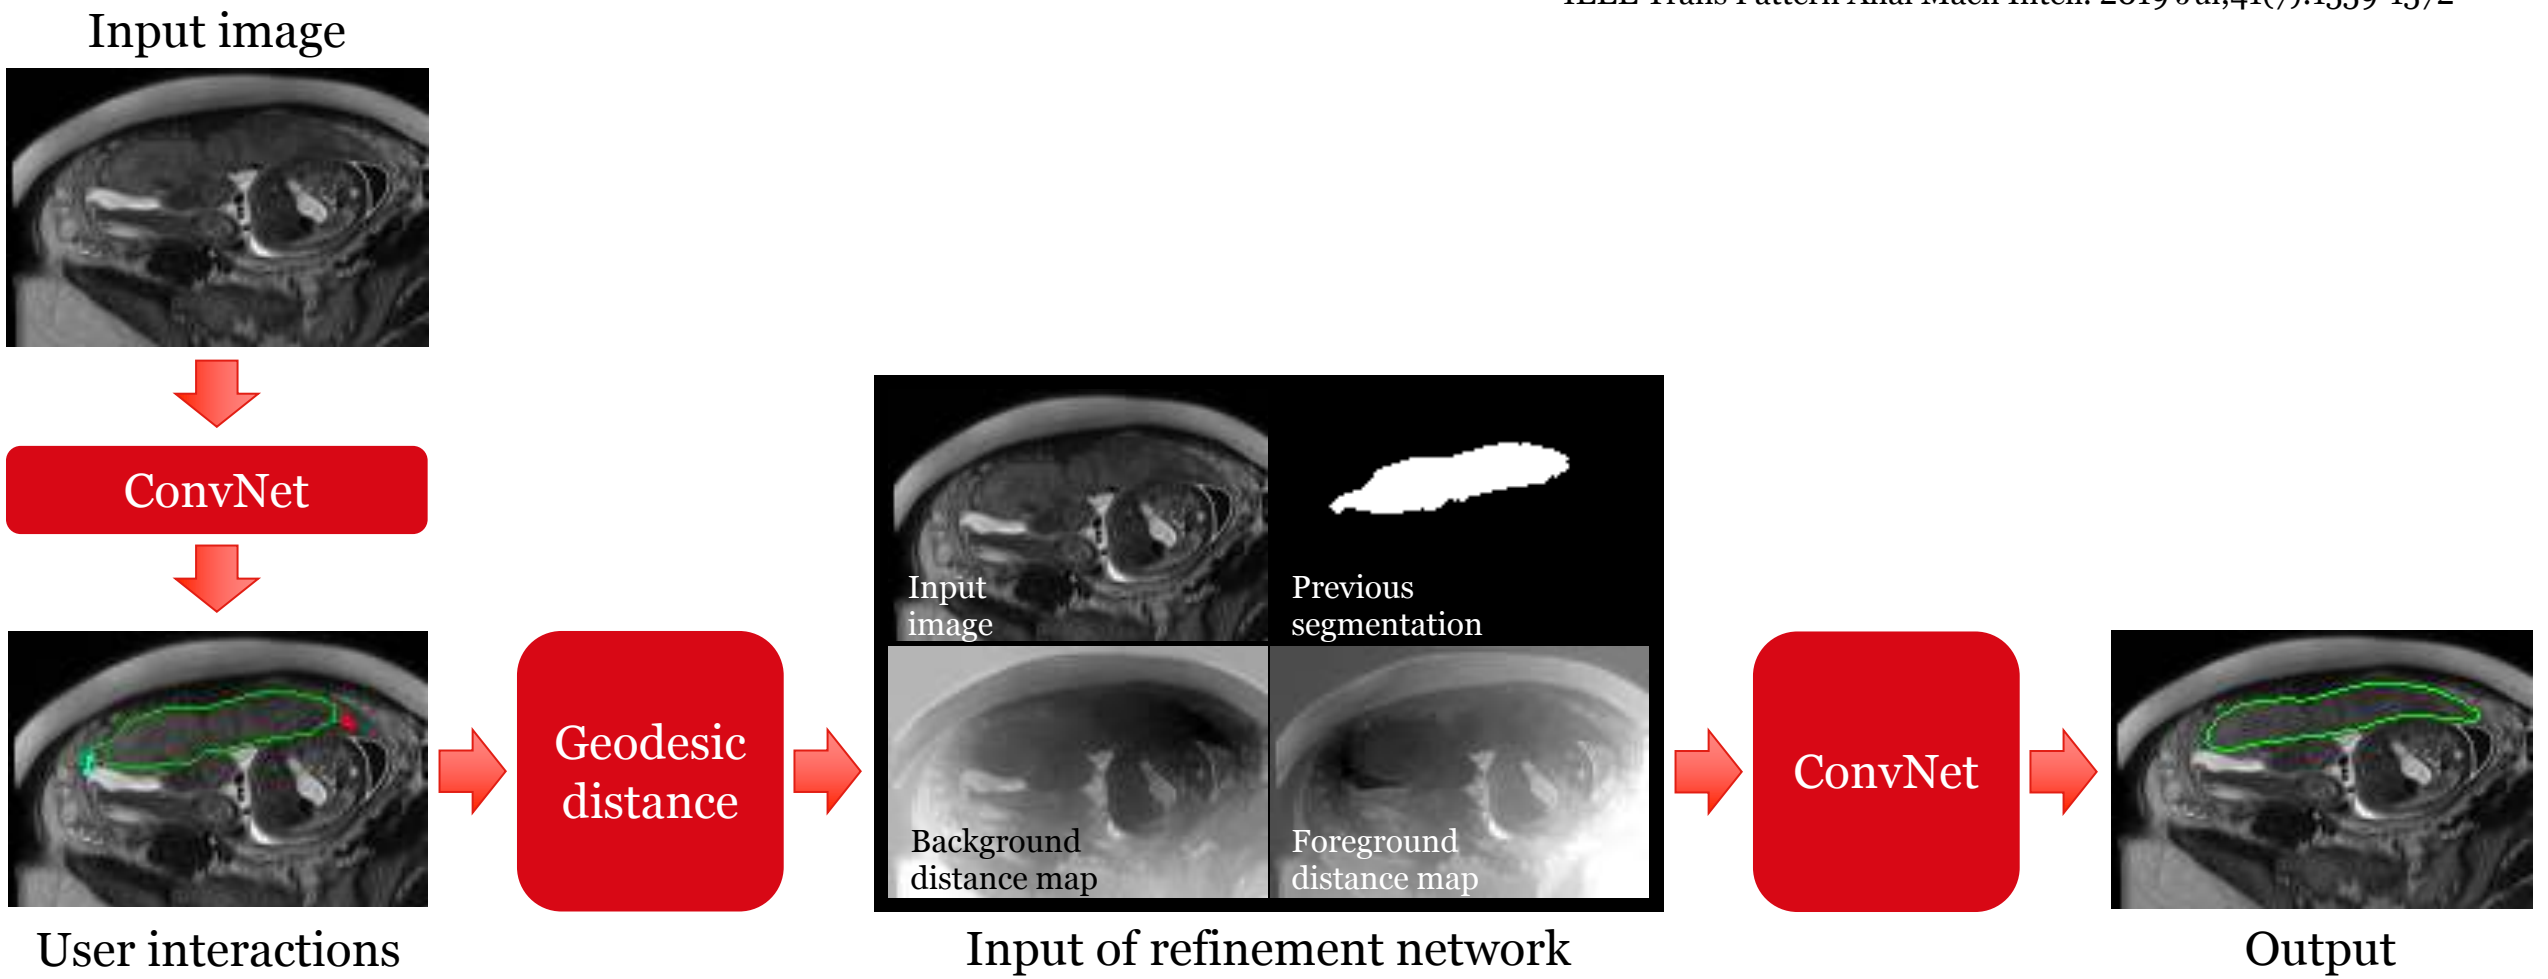

Scribbles as hard constraints

- Rely on conditional random field (CRF) for spatial regularisation and to impose hard constraints

$$E(\mathbf{x}) = \sum_i \psi_u(x_i) + \sum_{(i,j)\in\mathcal{N}} \psi_p(x_i, x_j)$$

Optimisation: CRF-RNN

Unary term: ConvNet output or Scribbles (∞)

Pairwise term: Fully trainable

- Trainable pairwise potentials

$$\psi_p(x_i, x_j) = \mu(l_i, l_j) f(\tilde{\mathbf{f}}_{ij}, d_{ij})$$

Label compat.

Free-form function

$\tilde{\mathbf{f}}_{ij} = \mathbf{f}_i - \mathbf{f}_j$

Features

d_{ij}

Euclid. Dist.

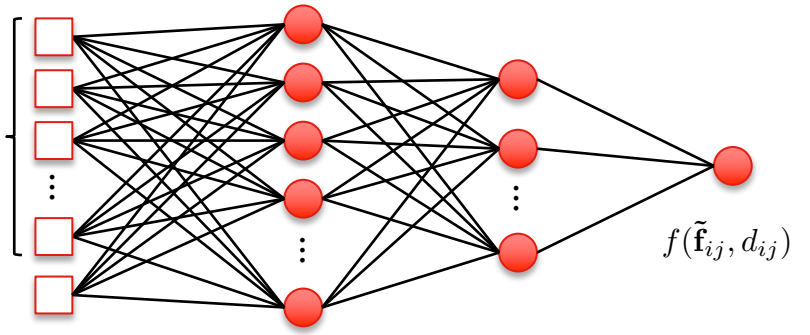

Classical image-specific fine-tuning for segmentation

- Motivation
  - Interactive segmentation
  - Learn from specific image context

- GrabCut
  - Rother et al. ACM Trans. Graph. 2004
  - Adapt a GMM to a specific image
  - Use a conditional random field (CRF) for spatial regularisation
  - Bounding box + refinement

- Deal with unseen objects
- Update the model on the fly

ConvNet: Pre-trained feature extraction + classifier

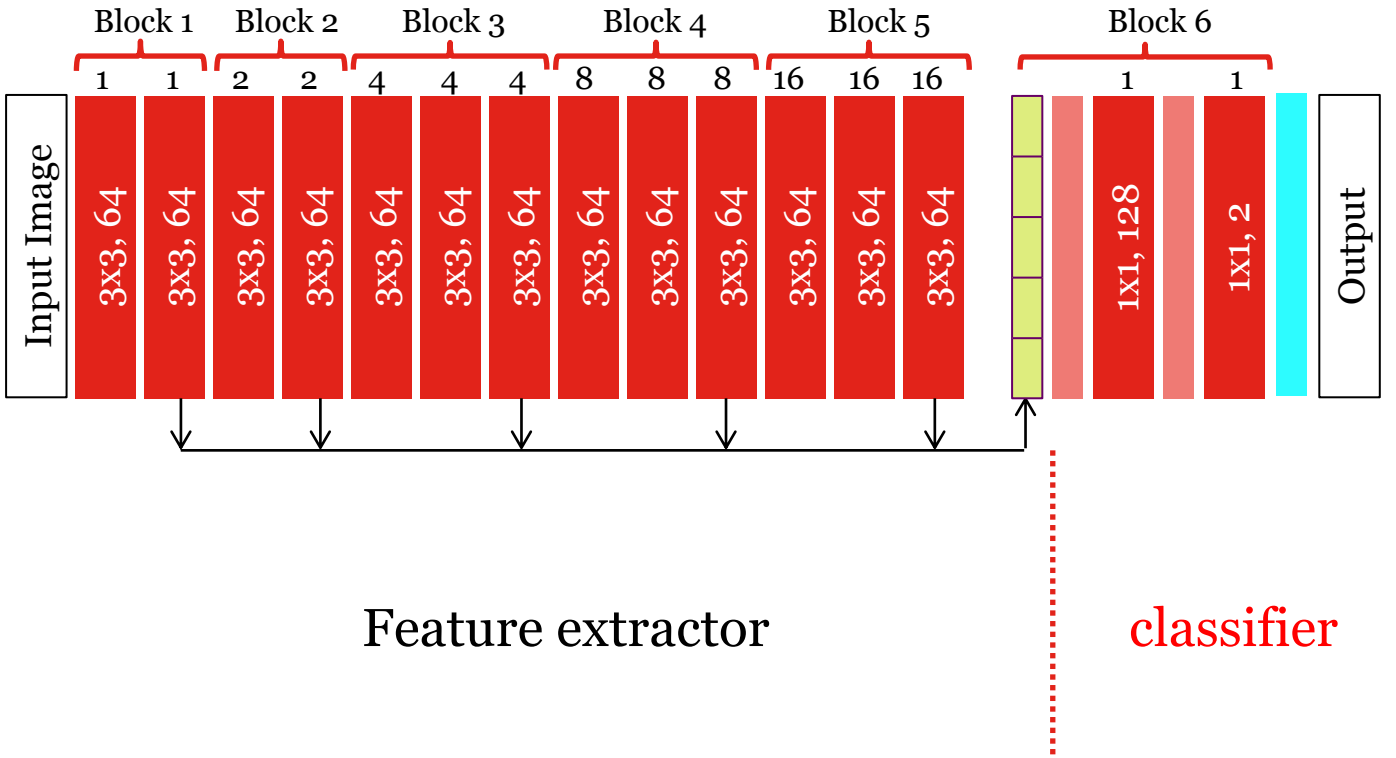

Image-specific fine-tuning with deep learning

- Method
  - Pre-train a ConvNet using cropped images of different organs
  - 1, Use ConvNet to get an initial segmentation inside a user-provided bounding box
  - 2, Fine-tune the ConvNet with optional scribbles
- Can deal with previously unseen objects

G Wang et al. Interactive Medical Image Segmentation Using Deep Learning With Image-Specific Fine Tuning. IEEE Trans Med Imaging. 2018 Jul;37(7):1562-1573.

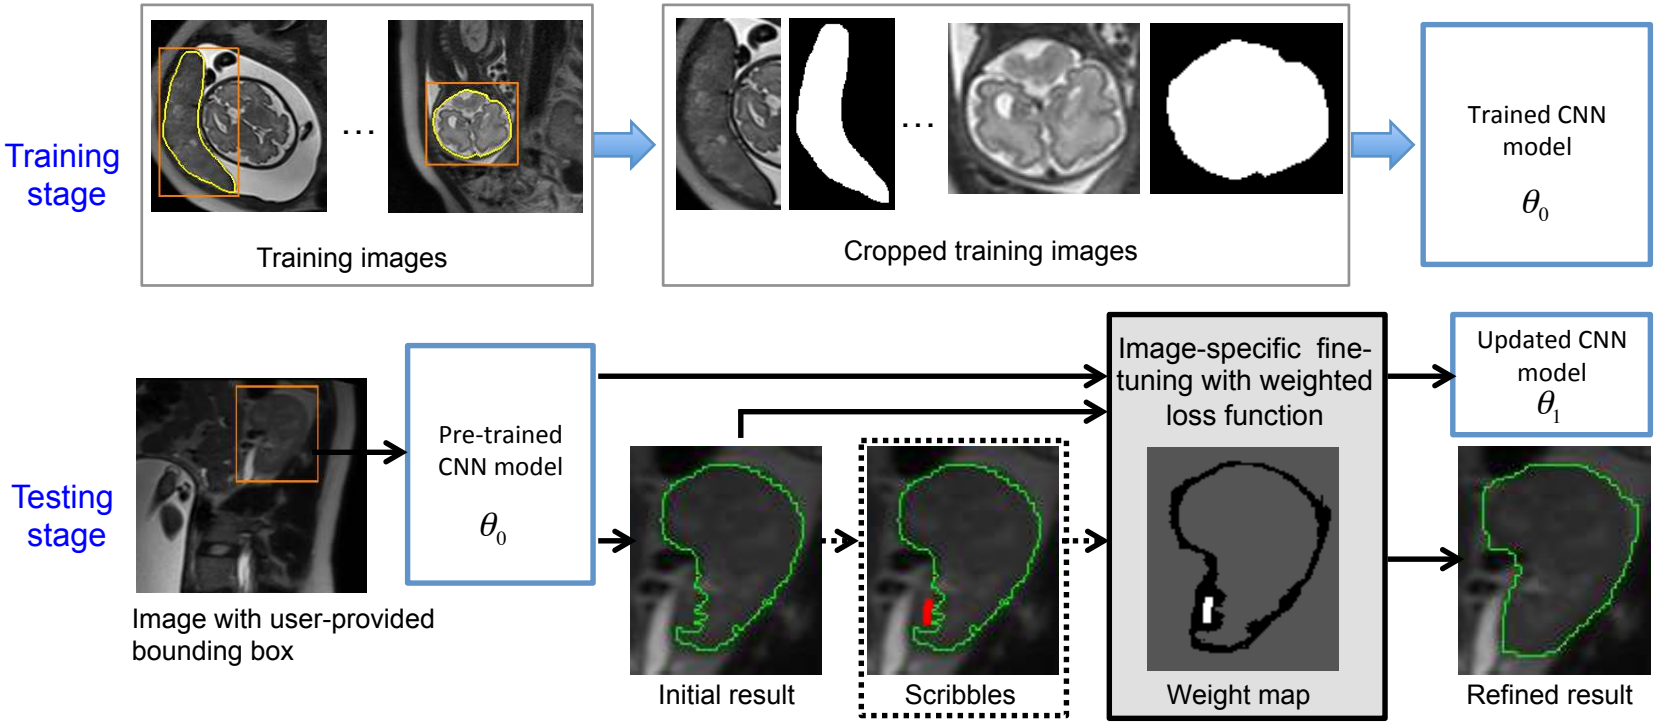

Alternate optimization for image-specific fine-tuning

- Joint optimization problem:

$$\arg\min_{\hat{Y},\theta} \left\{ E(\hat{Y},\theta) = \sum_i \phi(\hat{y}_i|\hat{X},\theta) + \lambda \sum_{i,j} \psi(\hat{y}_i,\hat{y}_j|\hat{X}) \right\}$$
$$\text{subject to : } \hat{y}_i = s_i \quad \text{if } i \in S$$

- When θ is fixed
  (Graph Cut type problem)

$$\arg\min_{\hat{Y}} \left\{ \sum_i \phi'(\hat{y}_i|\hat{X},\theta) + \lambda \sum_{i,j} \psi(\hat{y}_i,\hat{y}_j|\hat{X}) \right\} \quad (5)$$

$$\phi'(\hat{y}_i|\hat{X},\theta) = \begin{cases} +\infty & \text{if } i \in S \text{ and } \hat{y}_i = s_i \\ 0 & \text{if } i \in S \text{ and } \hat{y}_i \neq s_i \\ -\log P(\hat{y}_i|\hat{X},\theta) & \text{otherwise} \end{cases}$$

- When $\hat{Y}$ is fixed
  (Back propagation)

$$\arg\min_{\theta} \left\{ -\sum_i \left(\hat{y}_i \log p_i + (1-\hat{y}_i)\log(1-p_i) \right) \right\}$$

Image-specific fine-tuning for interactive corrections

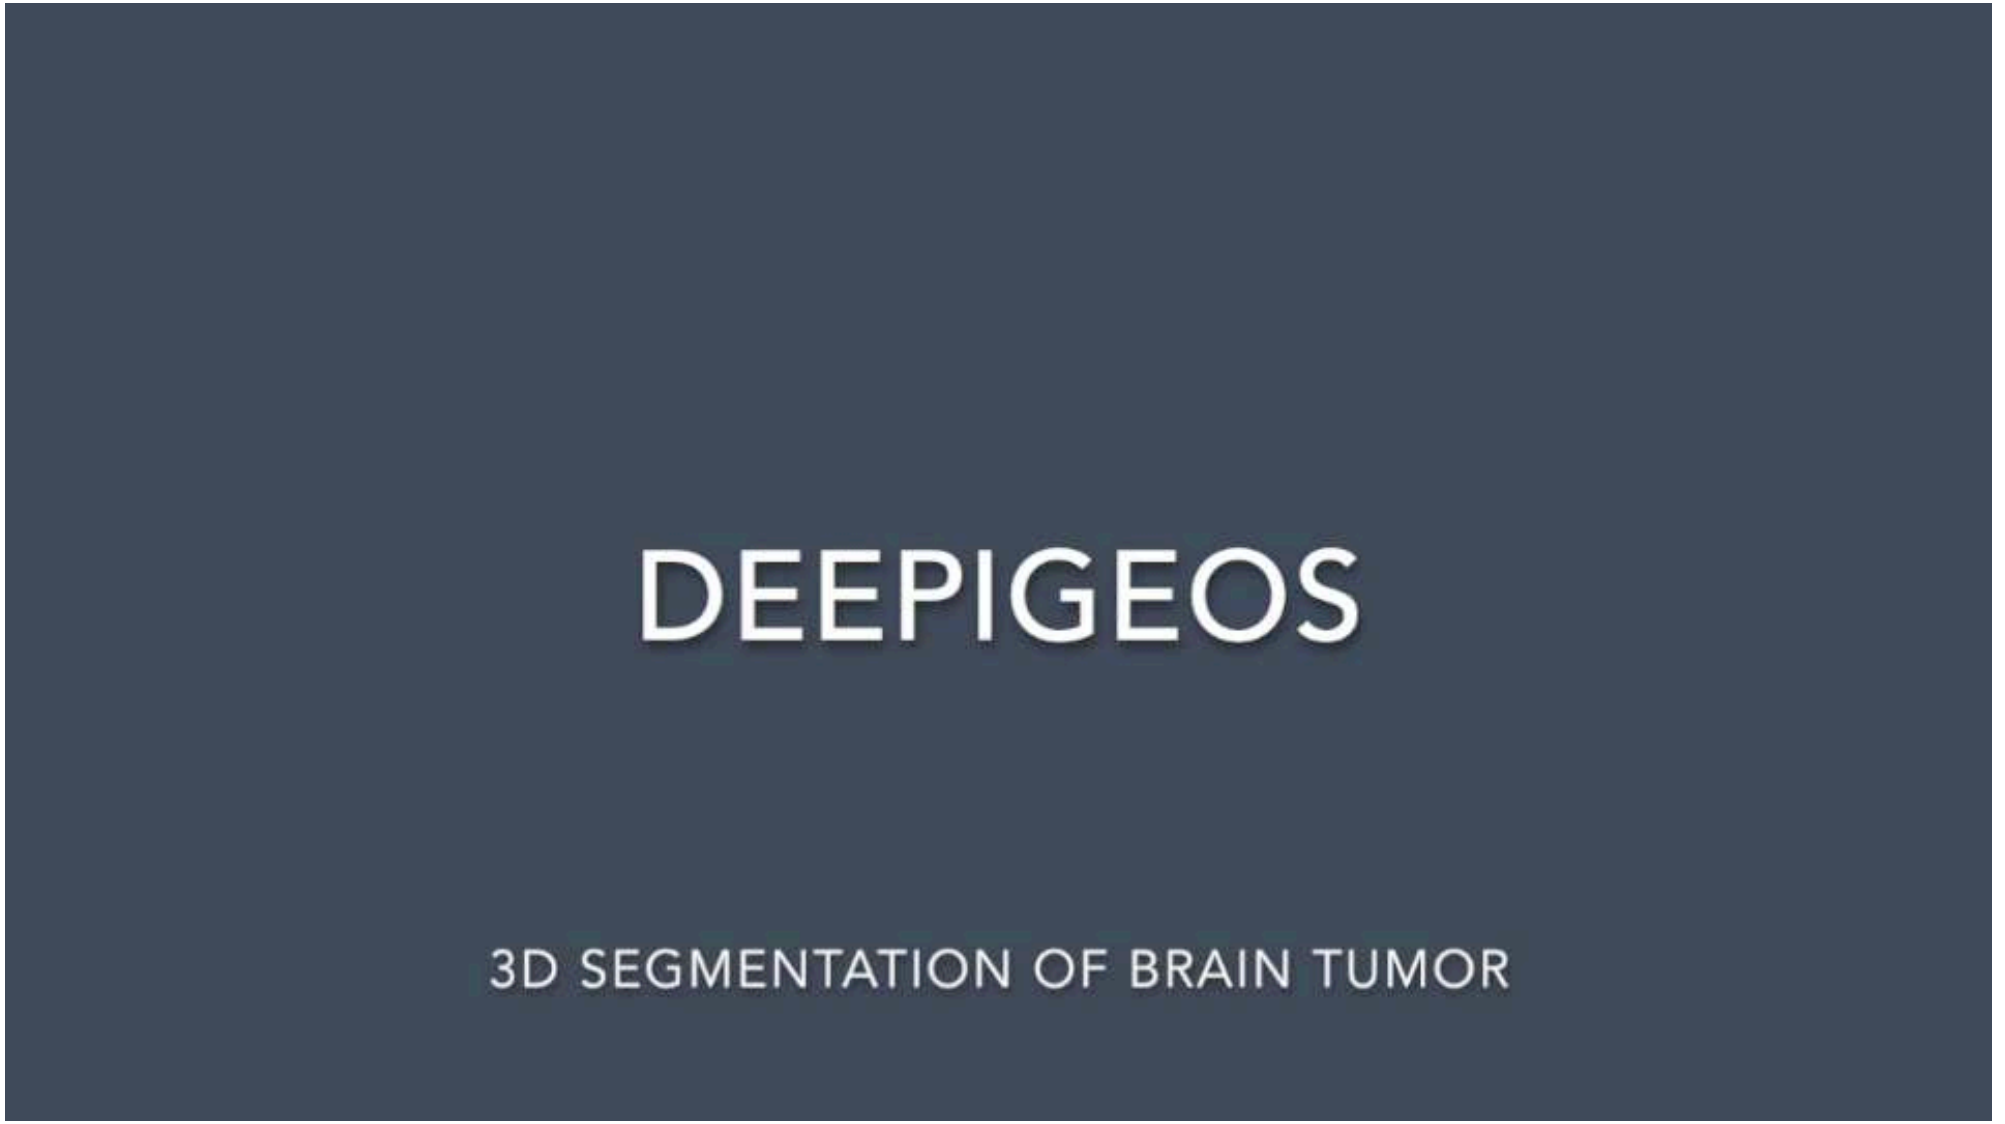

Wang et al.
IEEE TPAMI 2018

Wang et al.
IEEE TMI 2018

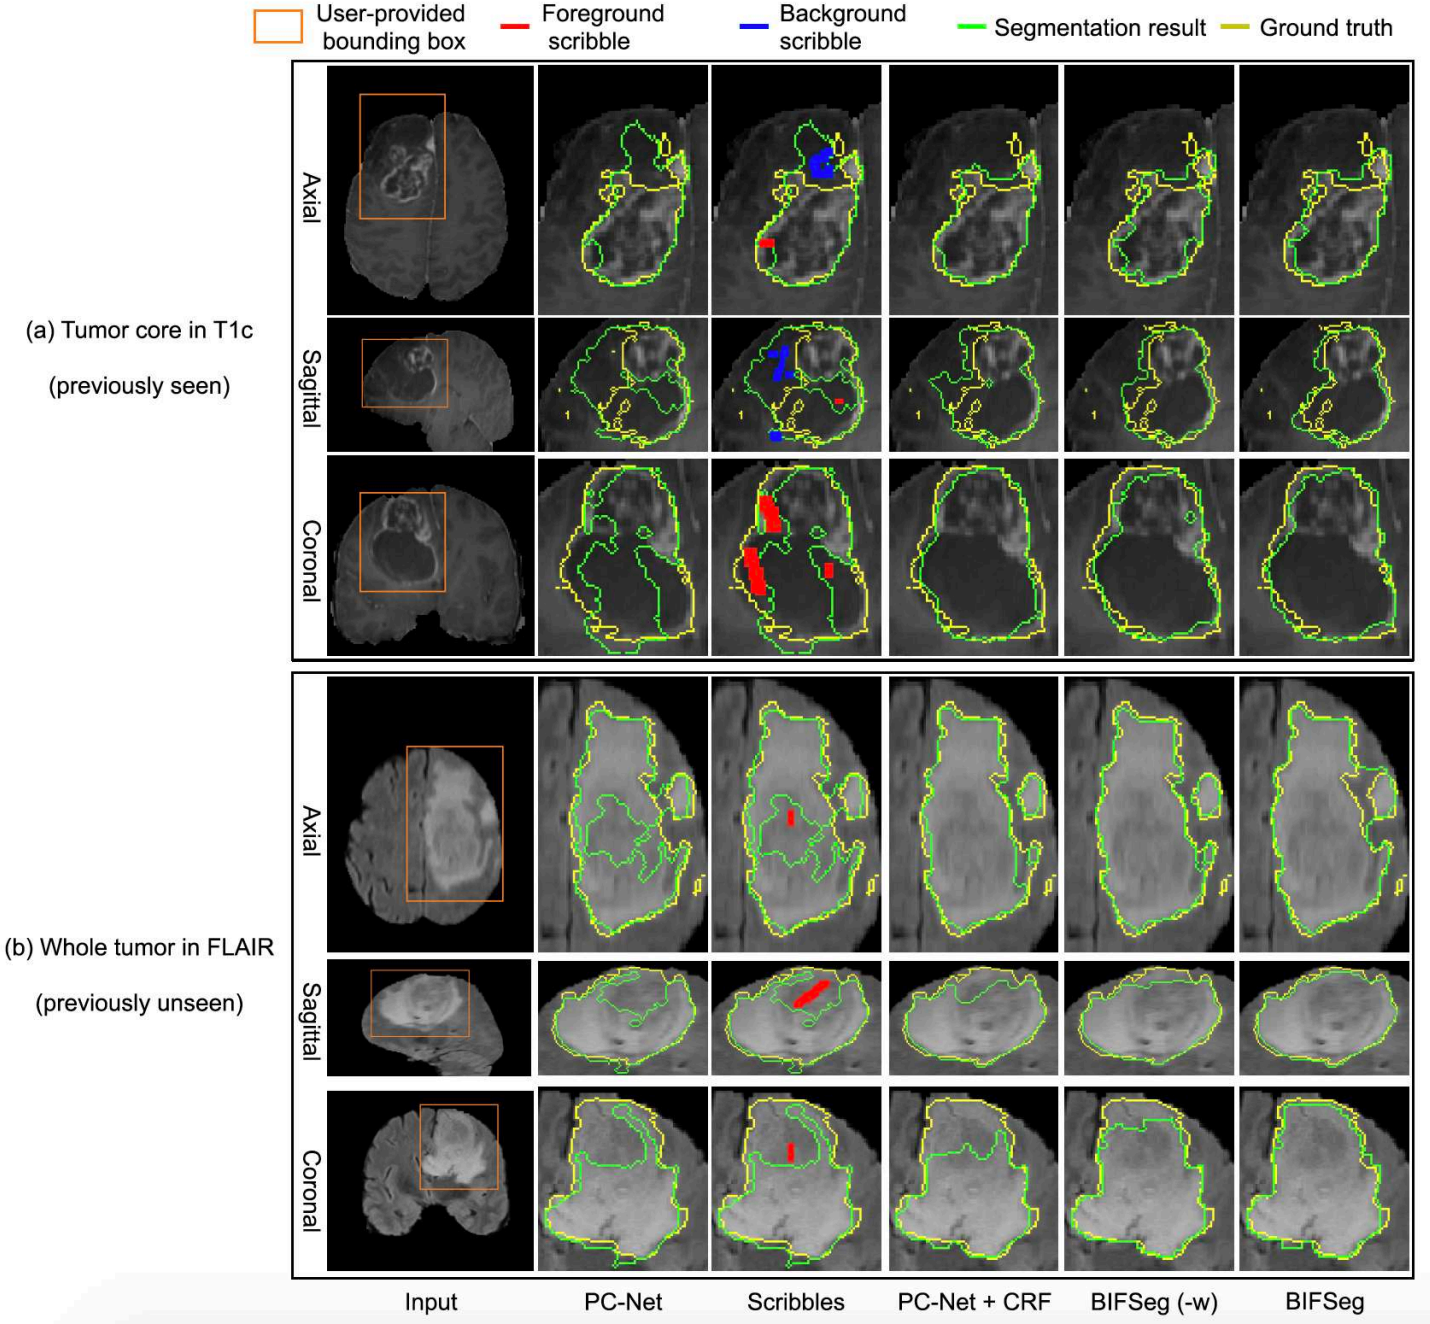
User-provided bounding box
Foreground scribble
Background scribble
Segmentation result
Ground truth
(a) Tumor core in T1c
(previously seen)
(b) Whole tumor in FLAIR
(previously unseen)
Axial
Sagittal
Coronal
Input
PC-Net
Scribbles
PC-Net + CRF
BIFSeg (-w)
BIFSeg

Conclusions

Take-home messages

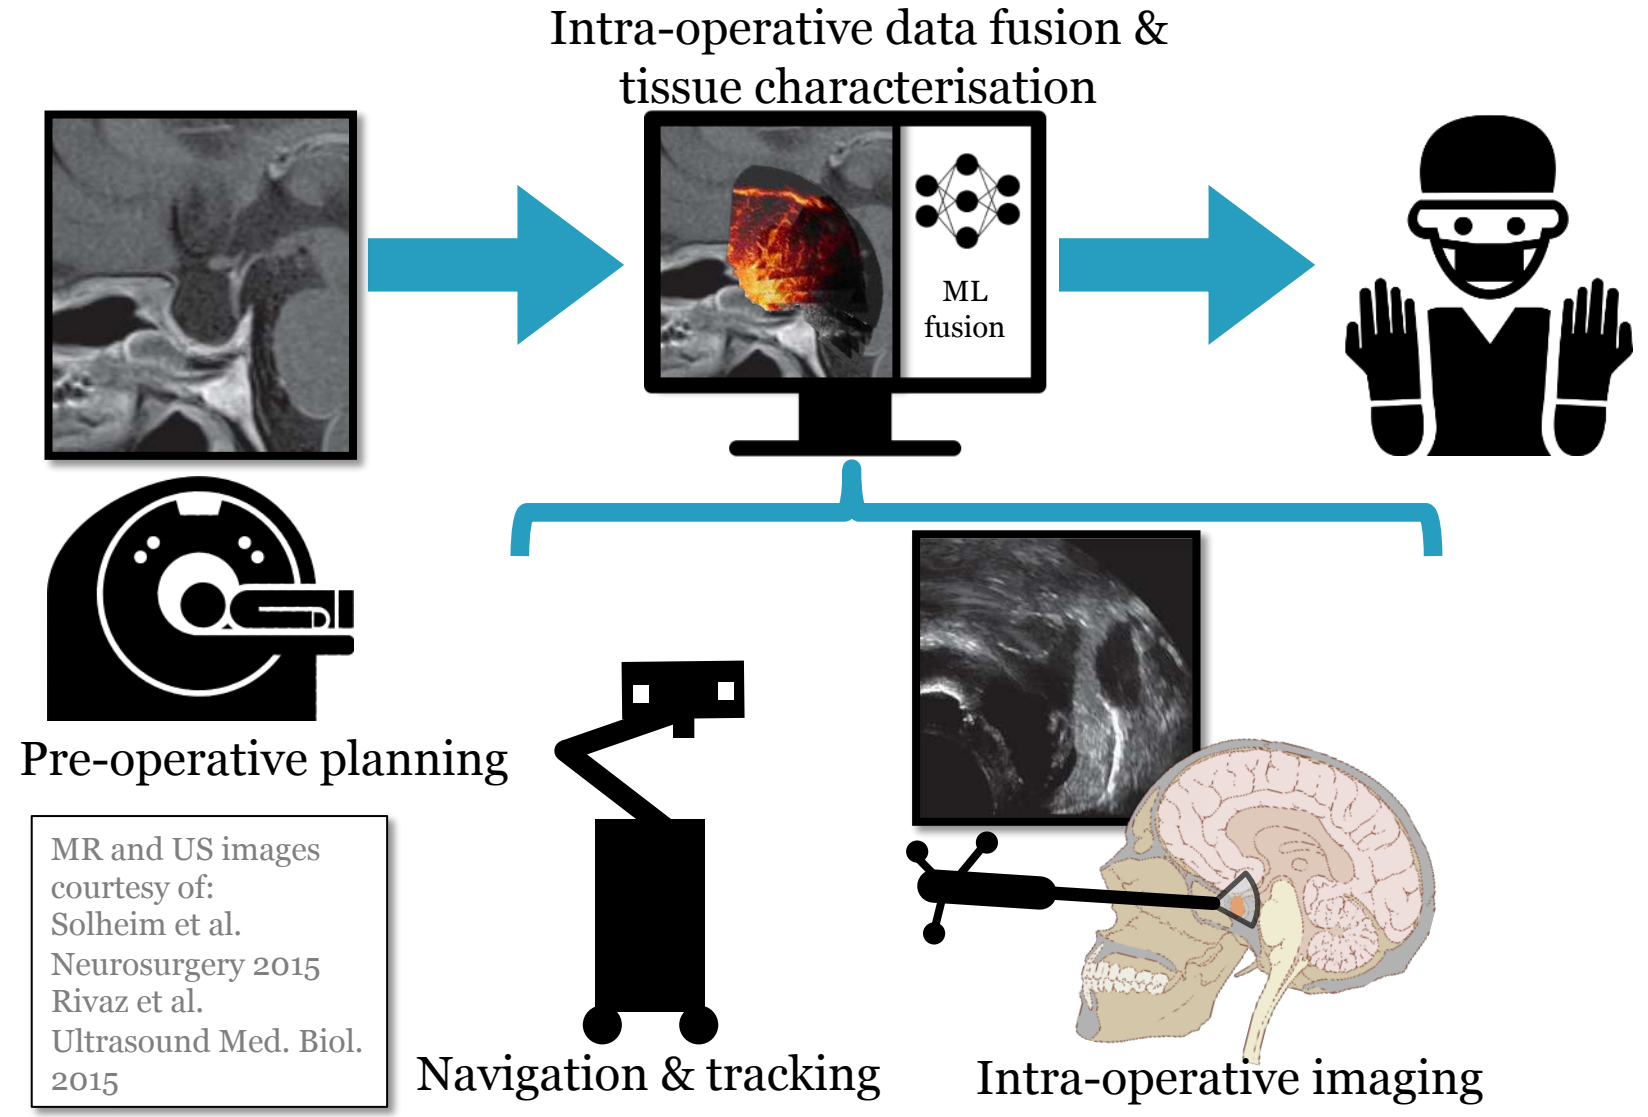

MR and US images courtesy of:
Solheim et al. Neurosurgery 2015
Rivaz et al. Ultrasound Med. Biol. 2015

- Flexible machine learning solutions are needed to address the complexity of clinical workflows
- Principled approaches can help us go beyond initial limitations of existing approaches
- Multi-disciplinarity is required to drive machine leaning towards patient impact

Thank you

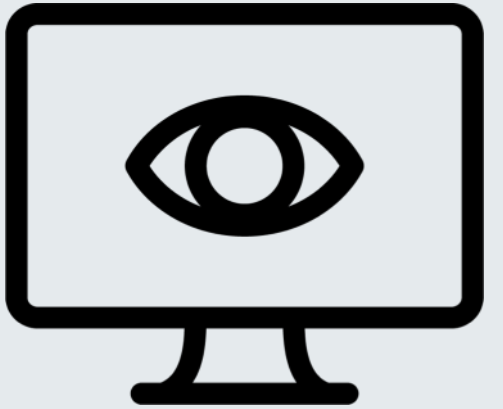

Contact details/for more information
tom.vercauteren@kcl.ac.uk