MAKING AI MAKE SENSE: CONCEPT-BASED PATHOLOGY DIAGNOSIS AND UNCERTAINTY-AWARE MRI

Asst. Prof. Christian Baumgartner


Lucerne Medical AI Group

MEDICAL IMAGING IS FULL OF UNCERTAINTIES

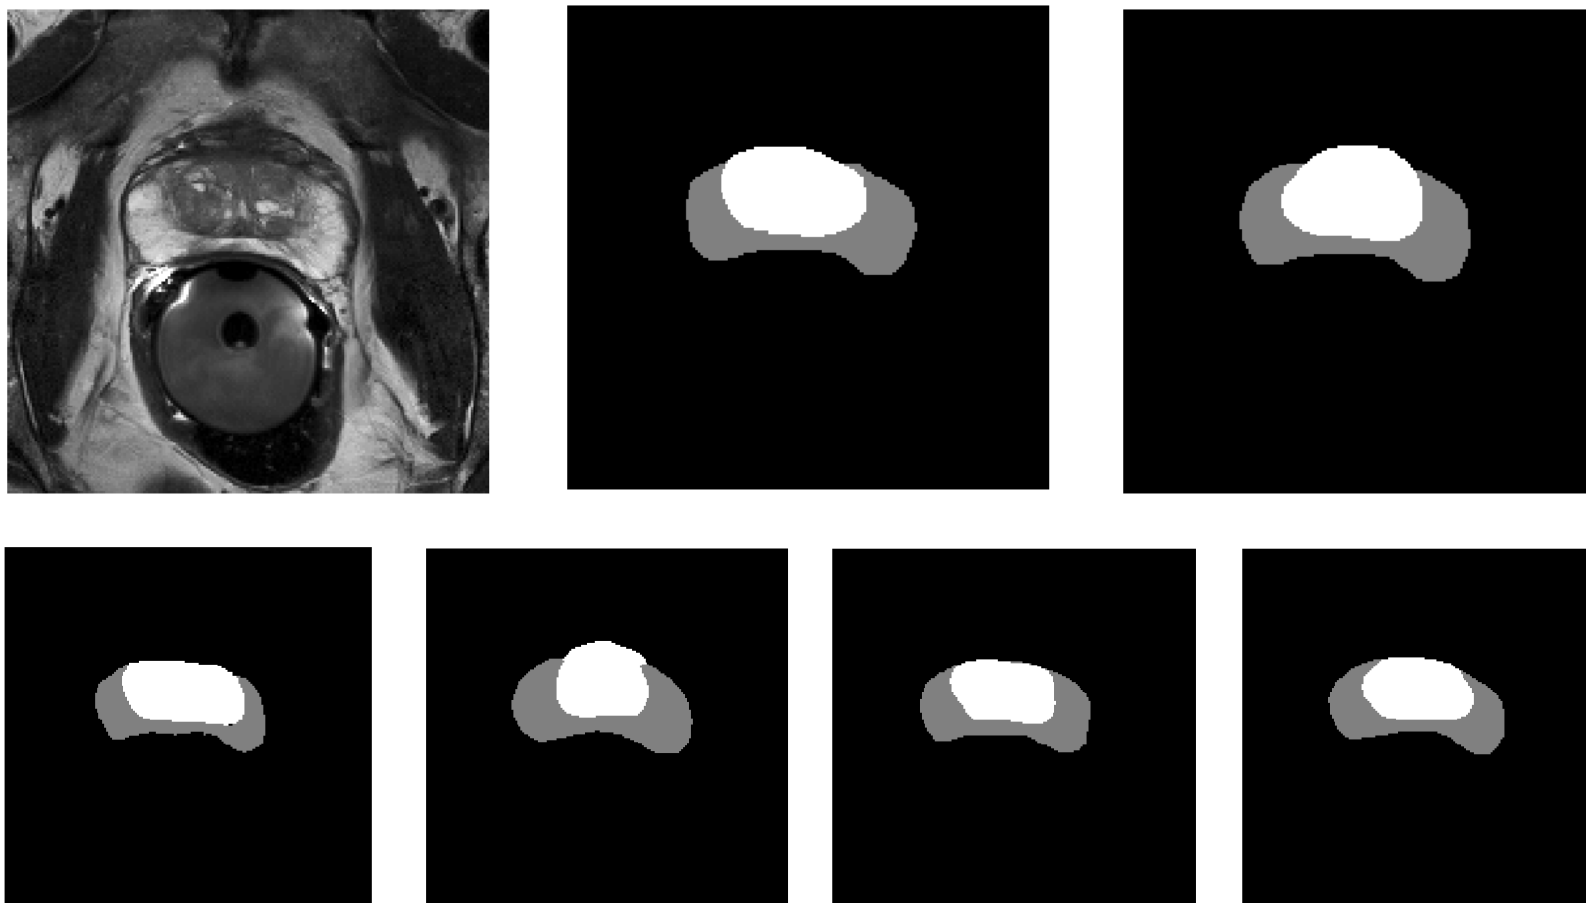

RESEARCH FOCUS: ESTIMATING UNCERTAINTIES

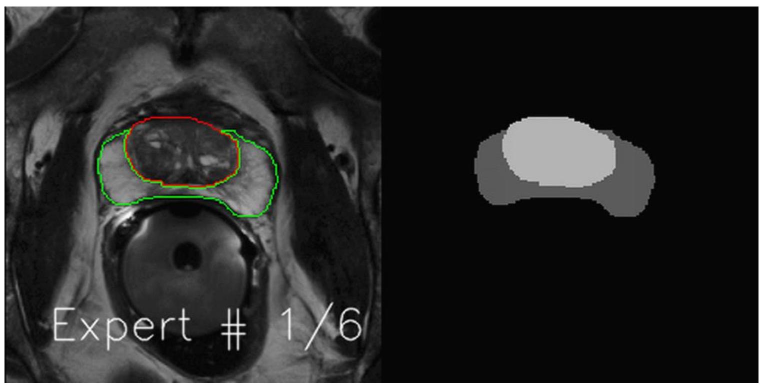

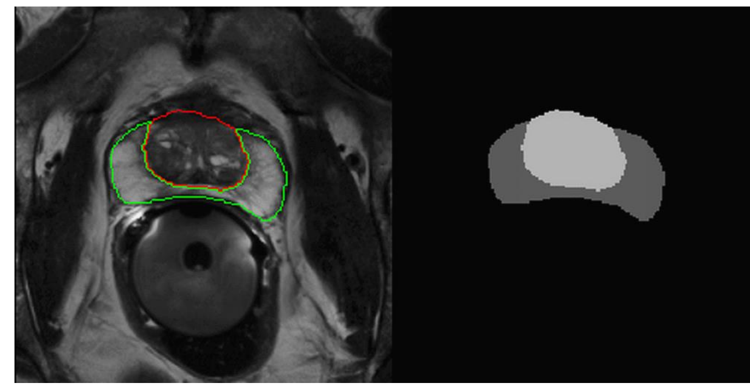

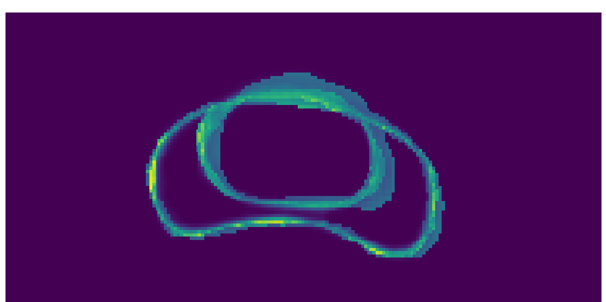

Annotator variance

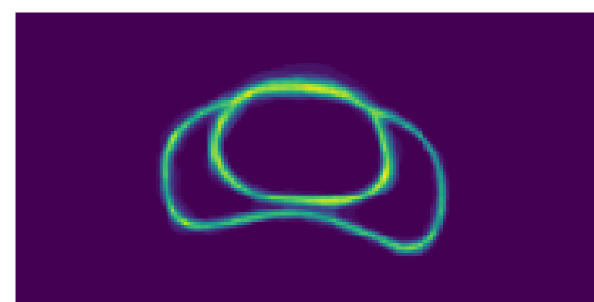

Predicted variance

Baumgartner, Christian F., et al. "PHiSeg: Capturing uncertainty in medical image segmentation." *MICCAI* 2019.

UNIVERSITÄT
LUZERN

RESEARCH GAPS

1) Usefulness of uncertainty of regarded as “self-evident”
   → *Uncertainty is only useful if influences some **downstream** clinical task*

2) Clinical decisions of rely on a cascade of clinical steps: how can we propagate uncertainty?

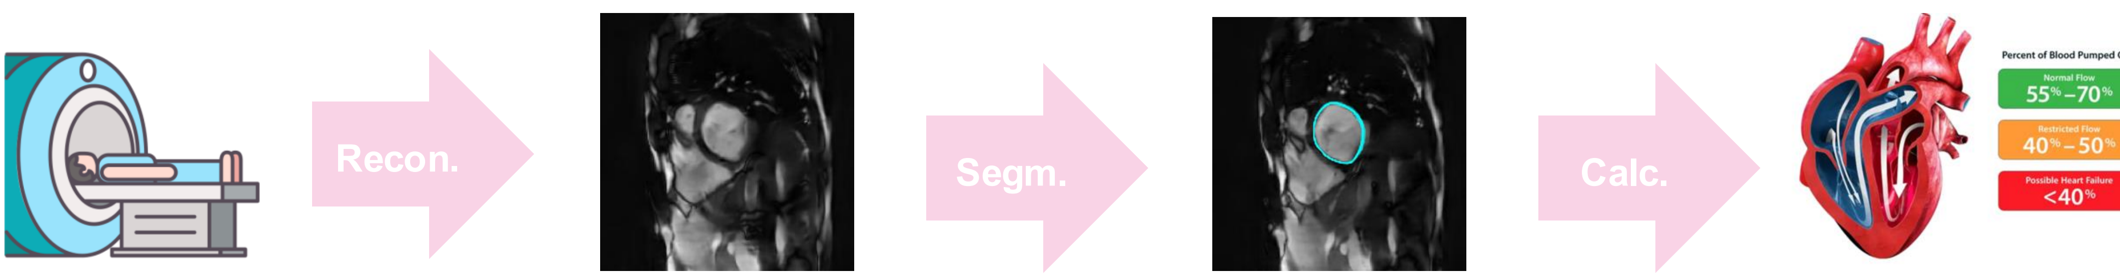

3) Could we use uncertainty to do something useful?
   → Personalise image acquisitions

FAST UNCERTAINTY-GUIDED MR ACQUISITION

Paul Fischer
PhD Student/
Post-doc Uni Basel

QUICK PRIMER ON MRI RECONSTRUCTION

Measurement data (k-space)

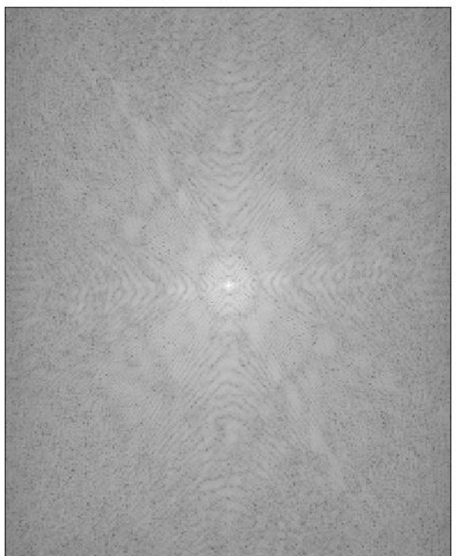

Image

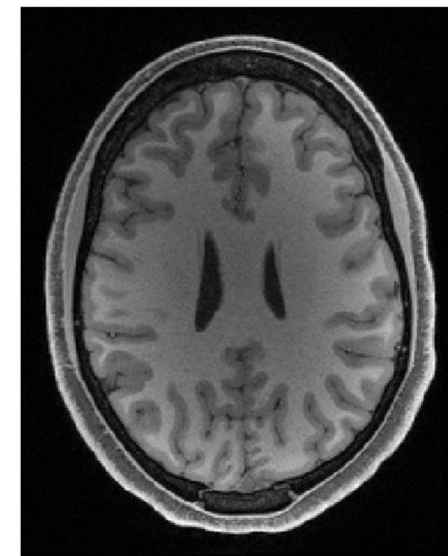

QUICK PRIMER ON MRI RECONSTRUCTION

Measurement data (k-space)

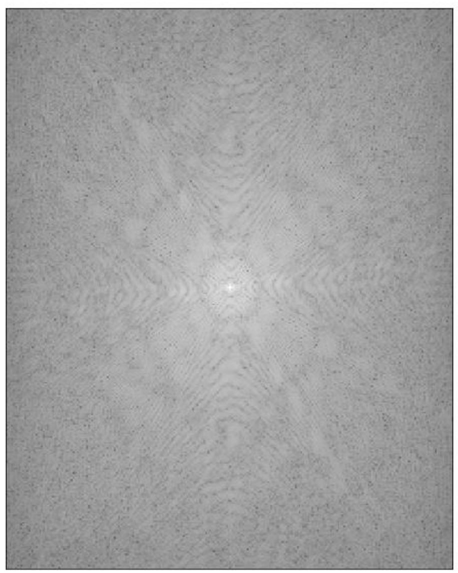

Fourier Transform

Image

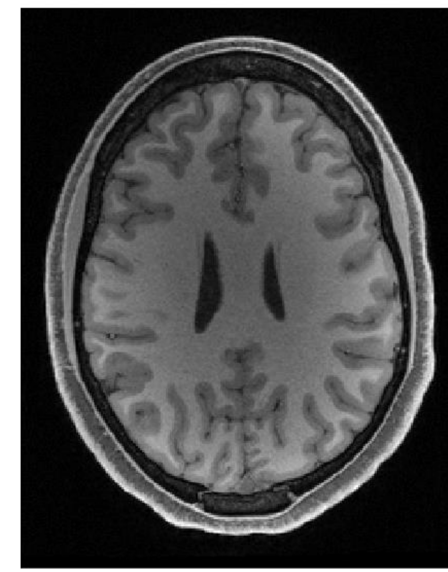

IMAGES CAN BE RECONSTRUCTED USING AI

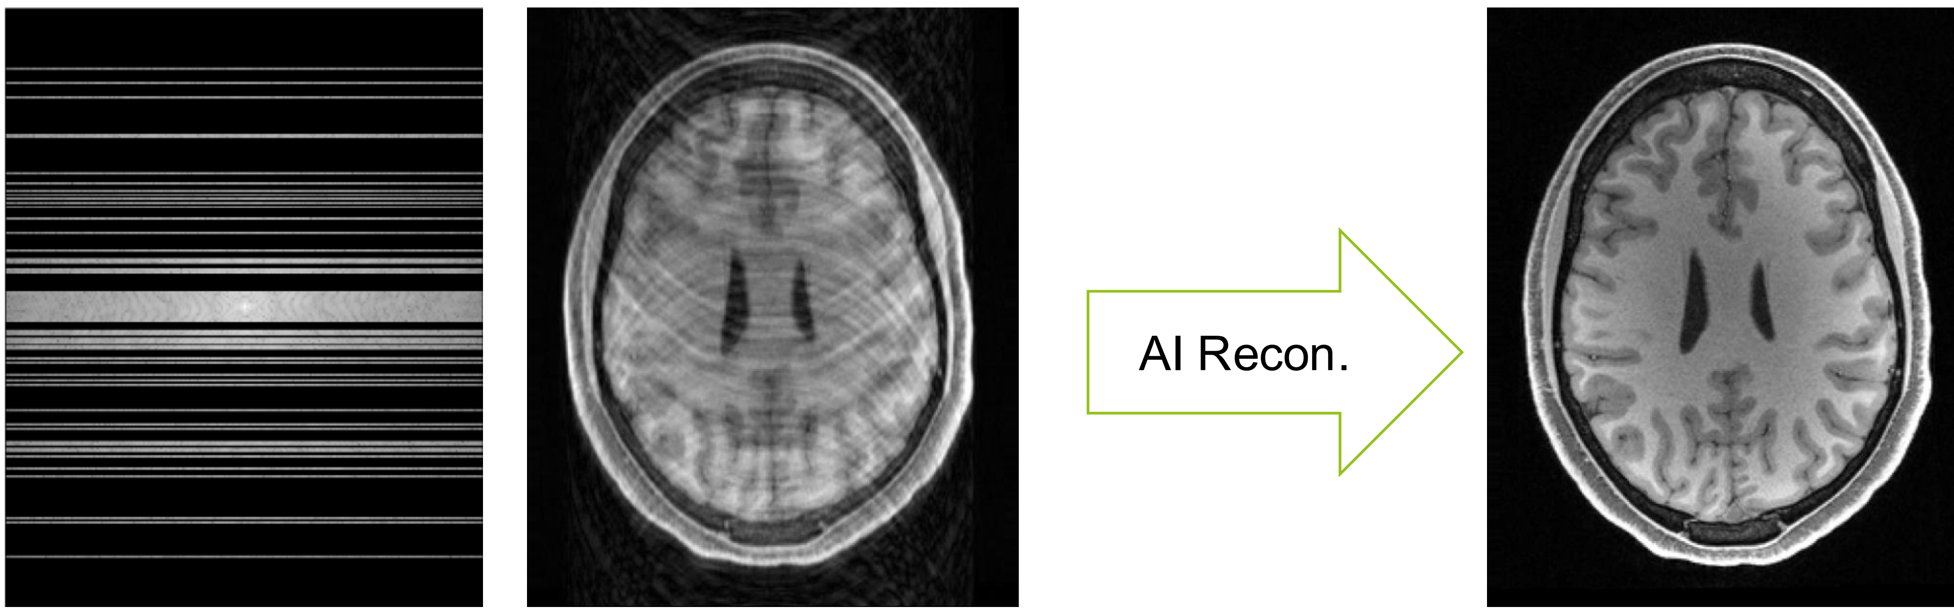

Infinitely many solutions! Some more likely than others.

FIRST CONTRIBUTION: PROBABILISTIC MR RECONSTRUCTION

State-of-the-art: MR Reconstruction based on Diffusion Models

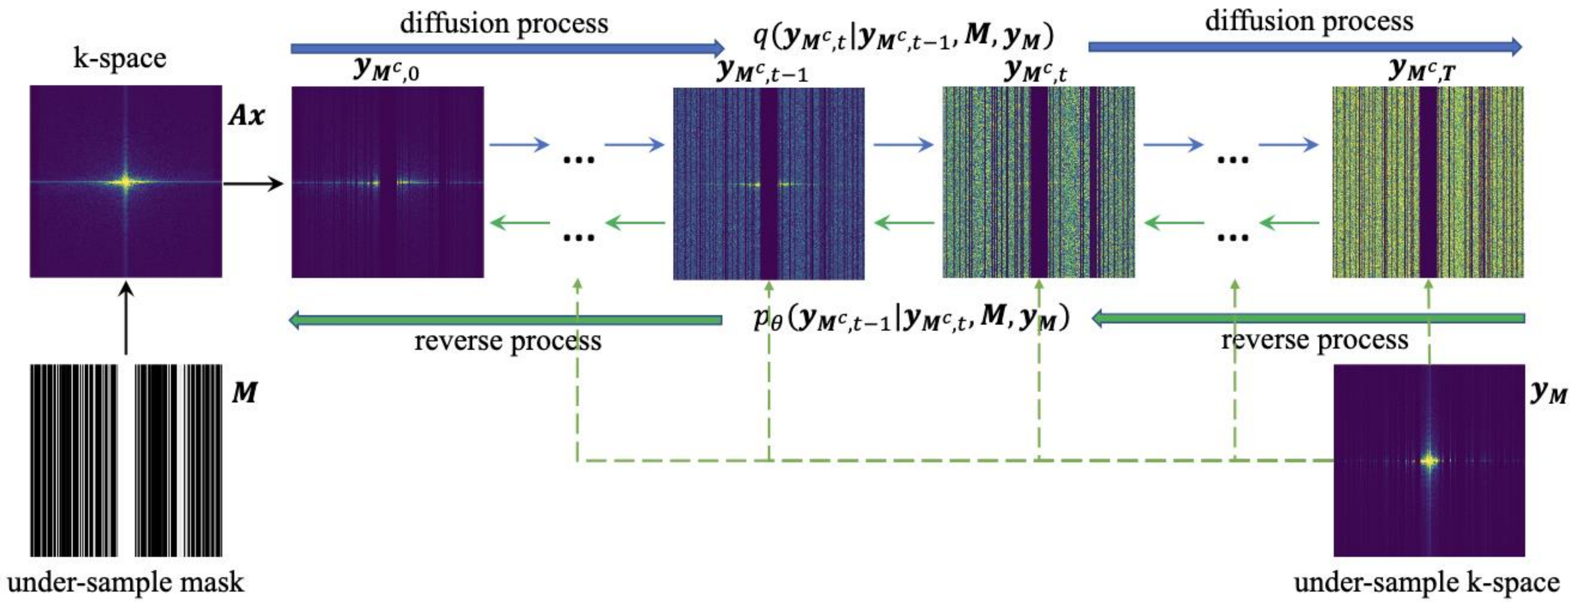

Xie, Yutong, and Quanzheng Li. "Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction." *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Cham: Springer Nature Switzerland, 2022.

Best reconstructions, but **very slow inference** (~10 seconds *per slice*)

OUR APPROACH IS BASED ON CONDITIONAL VAES

x_u

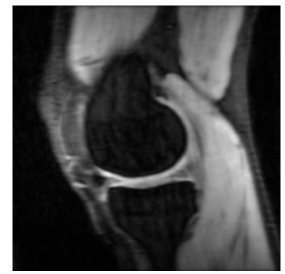

$p(x|x_u)$ x

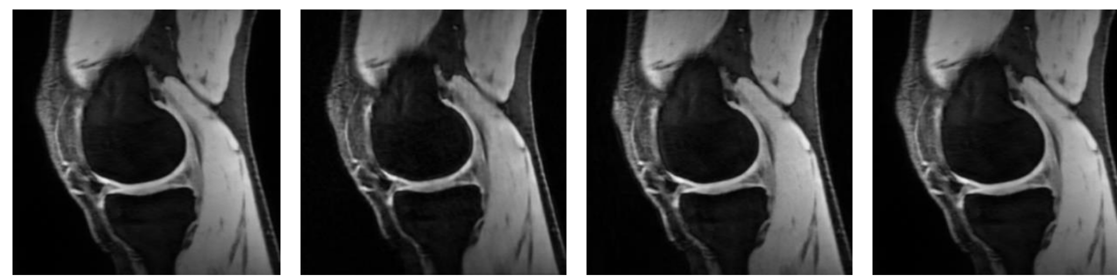

$\cdots$

Assumption of cVAE: low-dimensional *z* explains all the variation

$$p(x|x_u) = \int p(x|z)p(z|x_u)dz$$

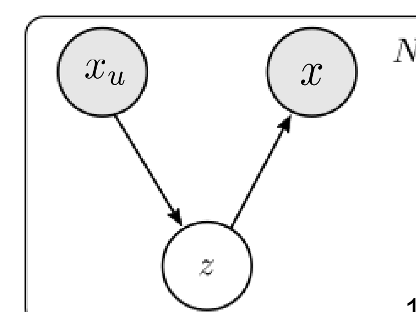

VARIATIONAL AUTOENCODERS: A BRIEF INTRO

$$\ln p(x) \geq \mathbb{E}_{q_\phi(z|x)}\left[\ln p_\theta(x|z)\right] - \mathrm{KL}\left[p(z)||q_\phi(z|x)\right]$$

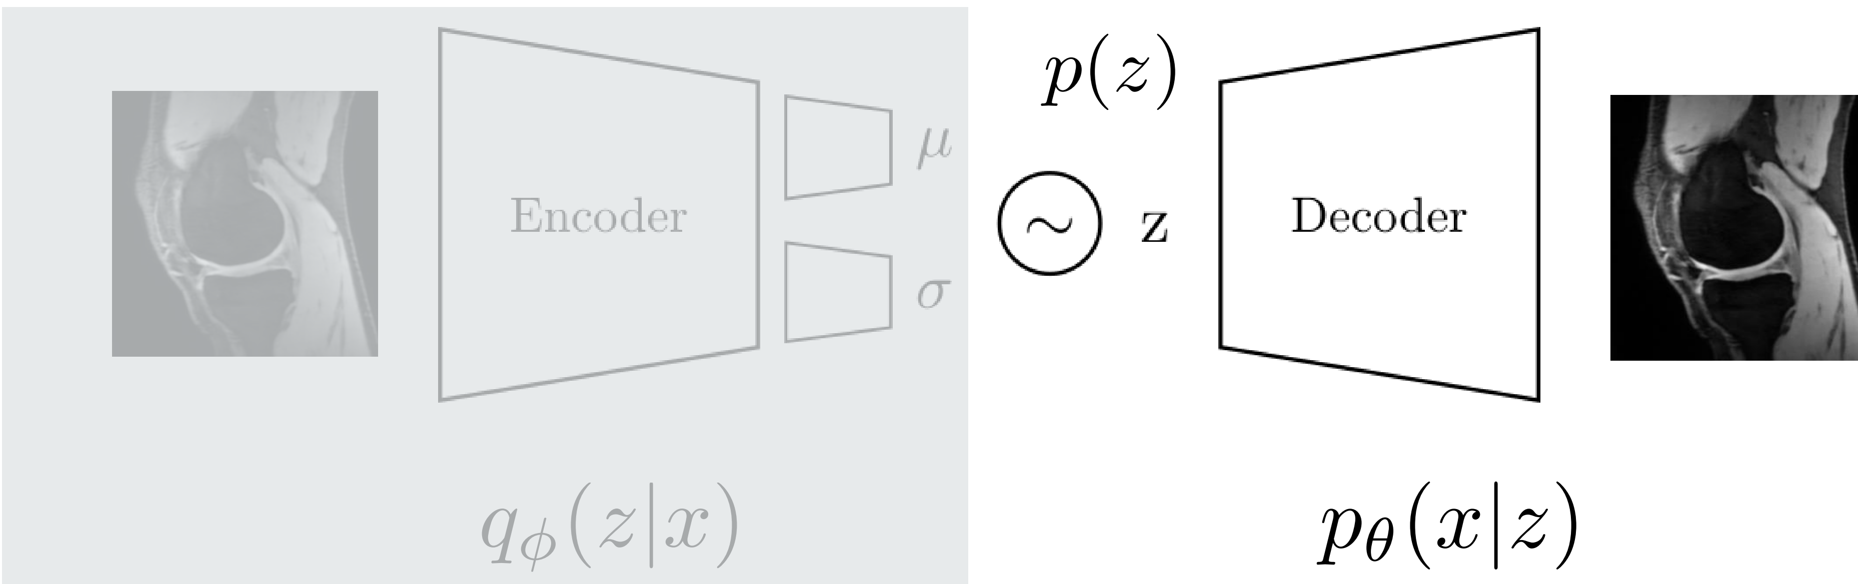

After training, discard posterior

CONDITIONAL VARIATIONAL AUTOENCODERS

VAE

$$p(s) = \int p(s|z)p(z)dz$$

$$\ln p(x) \geq \mathbb{E}_{q_\phi(z|x)}\left[\ln p_\theta(x|z)\right] - \mathrm{KL}\left[p(z)||q_\phi(z|x)\right]$$

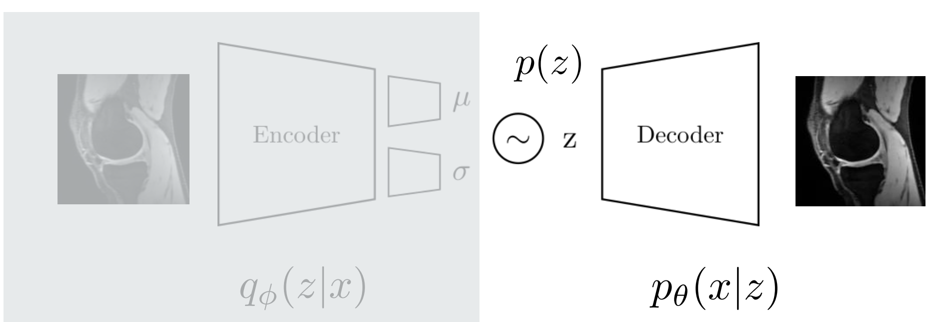

cVAE

$$p(x|x_u) = \int p(x|z)p(z|x_u)dz$$

$$\ln p(x|x_u) \geq \mathbb{E}_{q_\phi(z|x,x_u)}\left[\ln p_\theta(x|z)\right] - \mathrm{KL}\left[p(z|x_u)||q_\phi(z|x,x_u)\right]$$

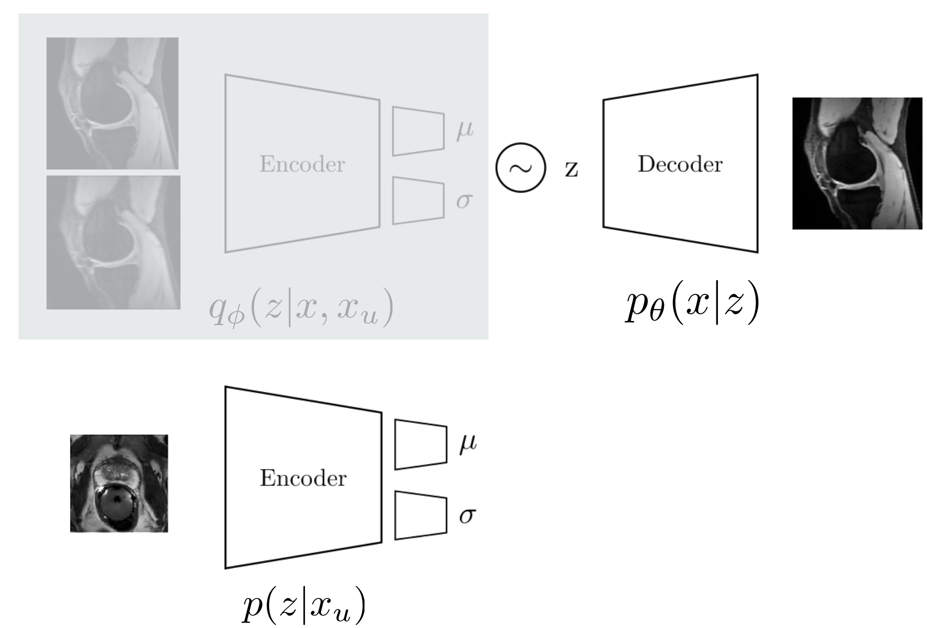

Sohn et al., Learning structured output representations using deep conditional generative models (NeurIPS 2015)

UNIVERSITÄT
LUZERN

FROM NORMAL CVAE TO HIERARCHICAL CVAE

One level cVAE

$$p(x|x_u) = \int p(x|z)p(z|x_u)dz$$

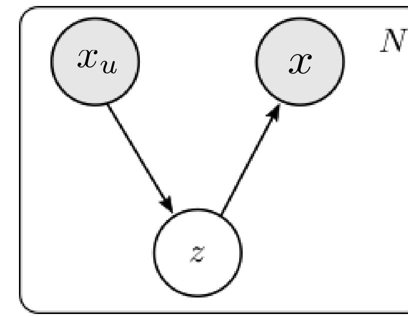

Hierarchical cVAE

$$p(x|x_u) = \int p(x|z_1, \ldots, z_L)p(z_1|z_2, x_u) \cdots p(z_{L-1}|z_L, x_u)p(z_L|x_u)dz_1 \cdots dz_L$$

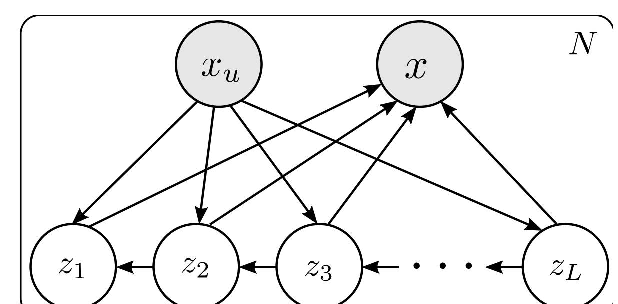

PROBABILISTIC HIERARCHICAL RECONSTRUCTION (PHIREC)

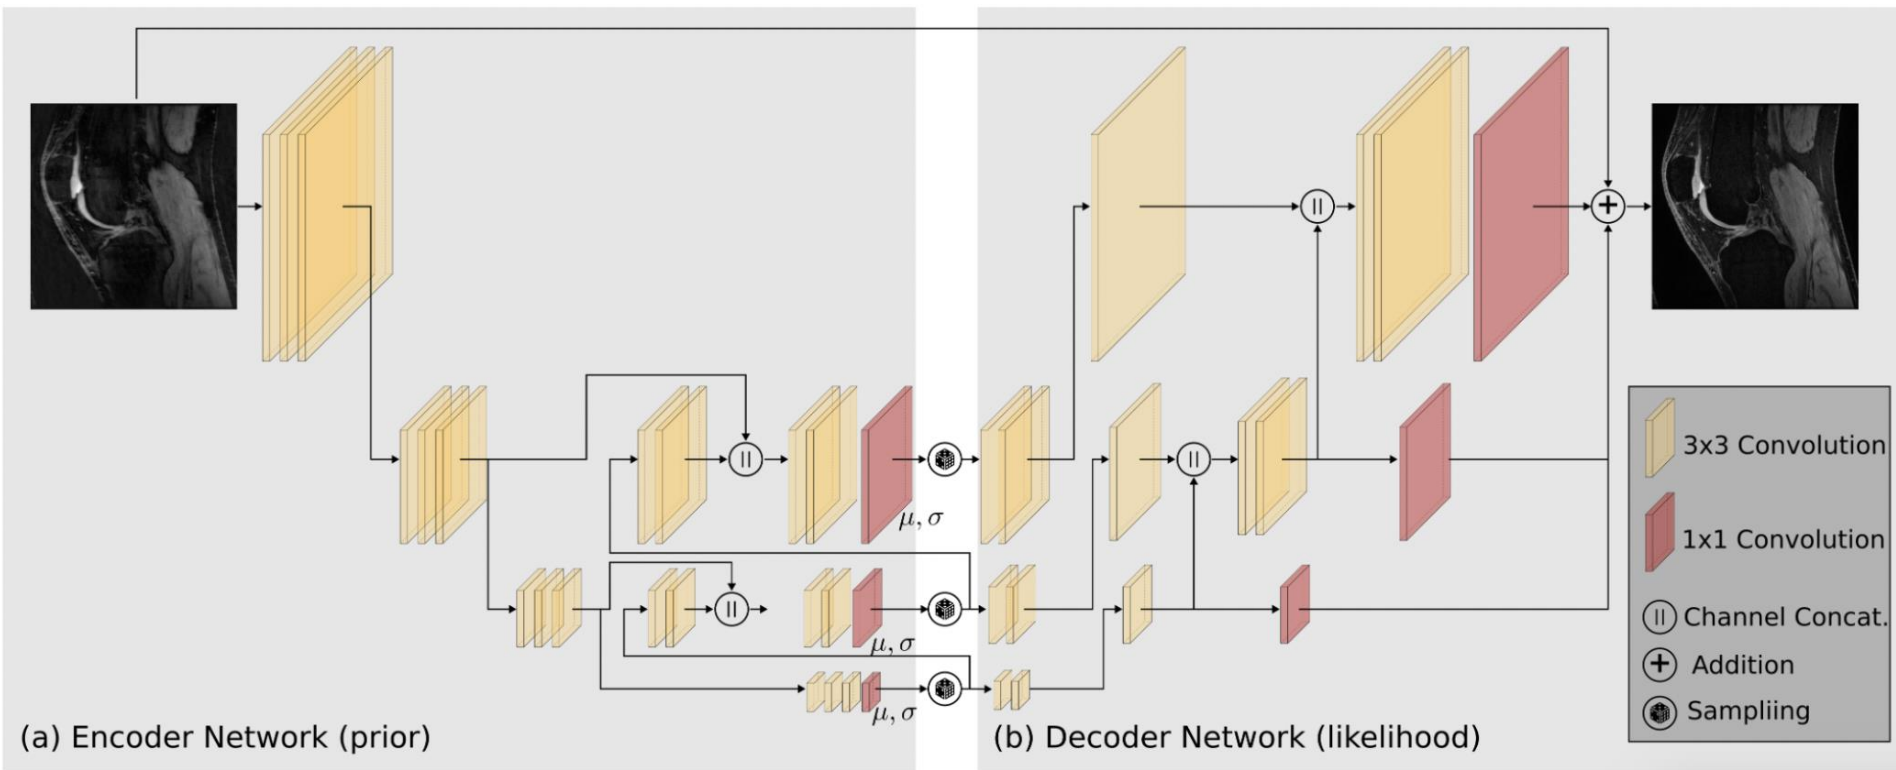

PHIREC FOR ACCELERATED MR IMAGING AND UNCERTAINTY PROPAGATION

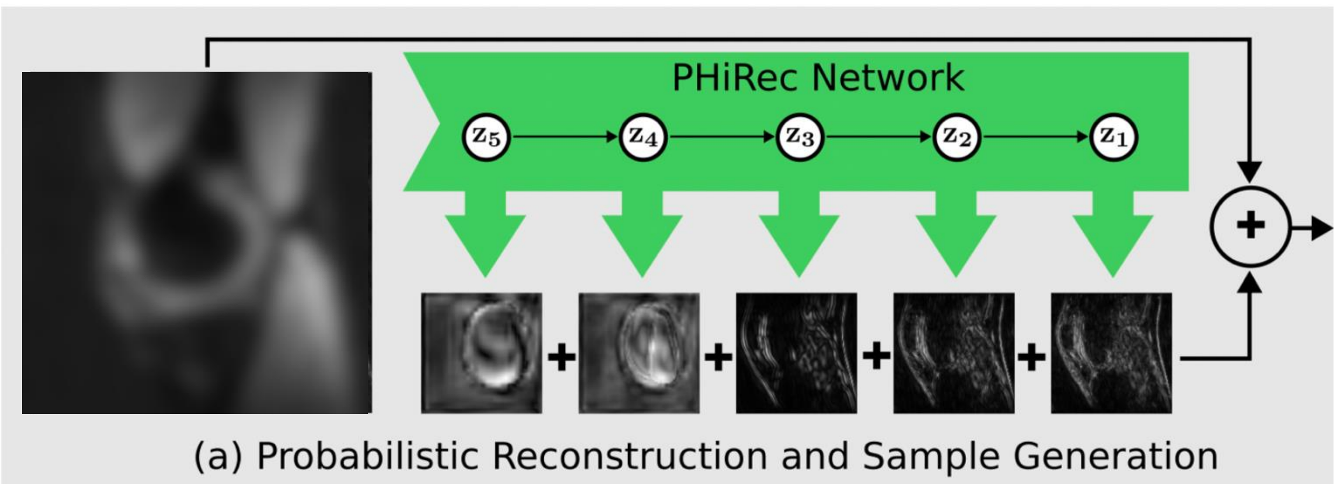

(a) Probabilistic Reconstruction and Sample Generation

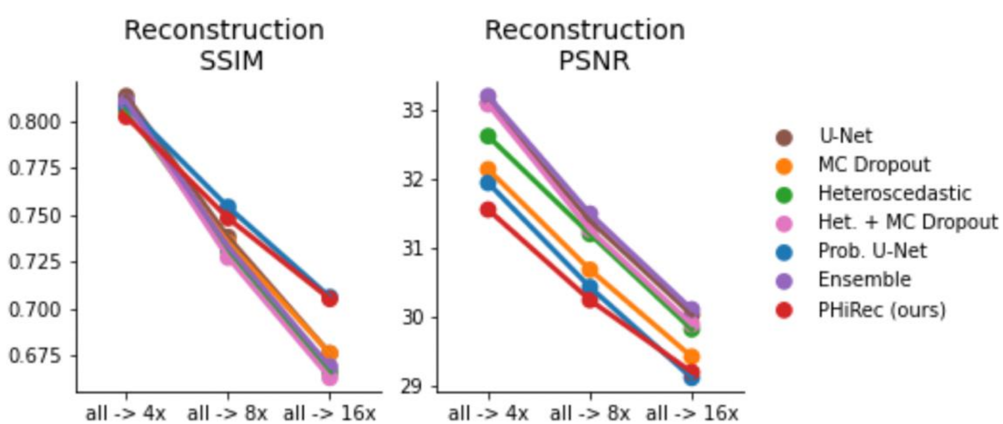

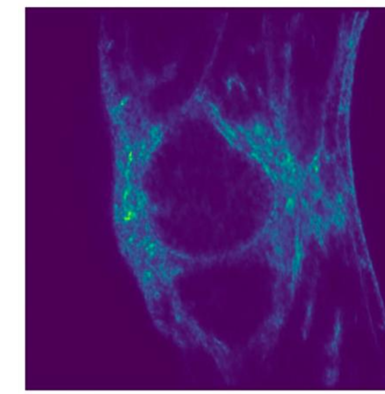

Infe e for 1 slice:

PHiRec: 1ms

Diffusion Model: ~10s

UNIVERSITÄT
LUZERN

UNCERTAINTY QUANTIFICATION PERFORMANCE

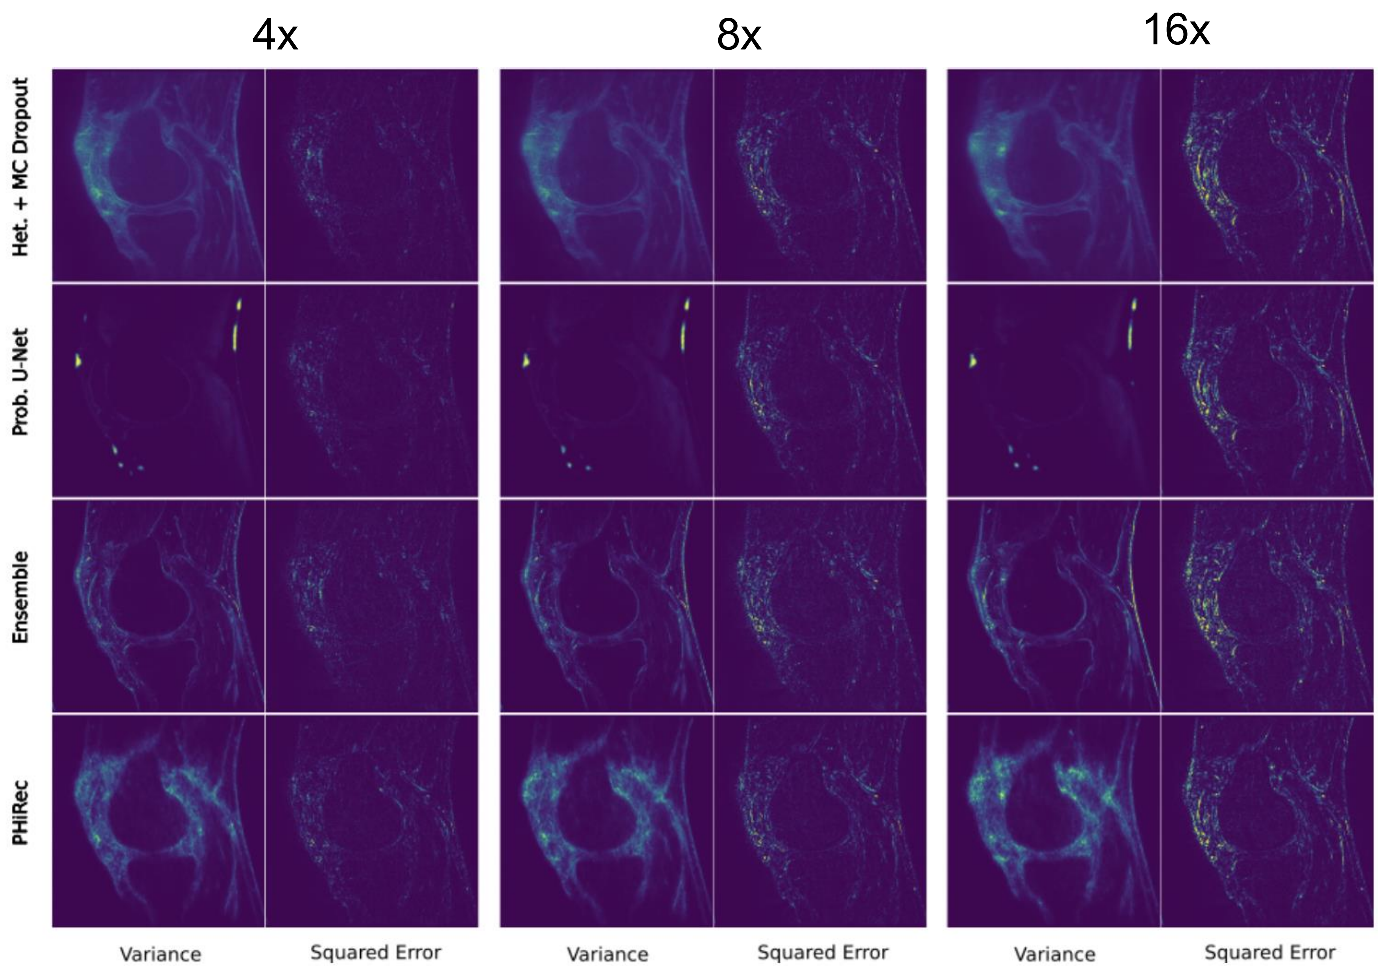

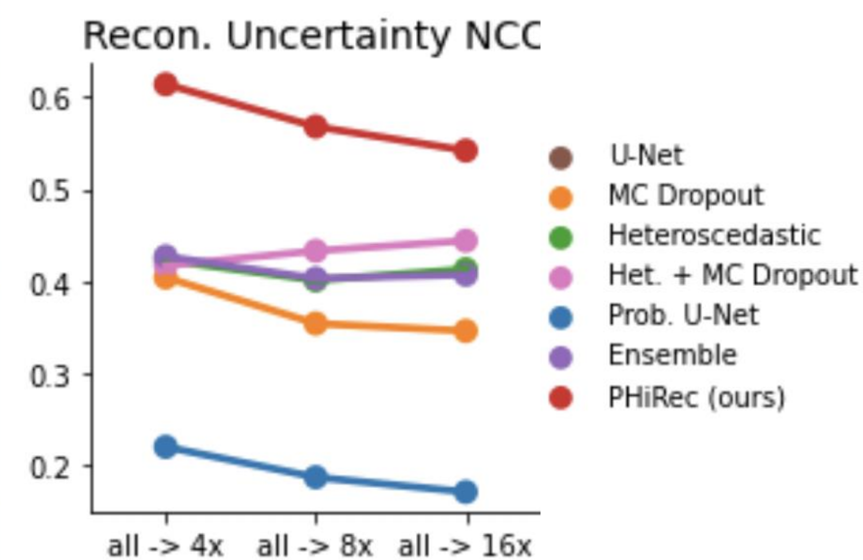

PART 2: UNCERTAINTY PROPAGATION

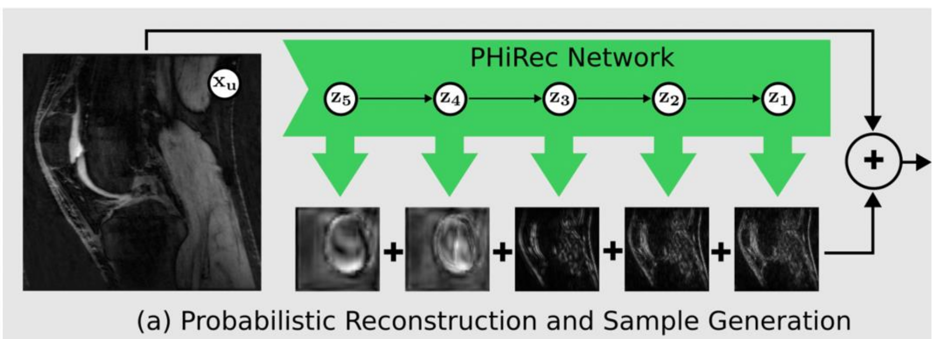

(a) Probabilistic Reconstruction and Sample Generation

$$p(x|x_u)$$

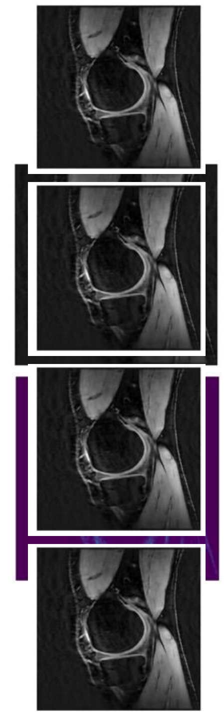

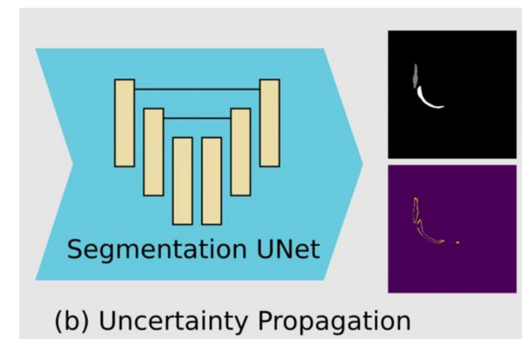

(b) Uncertainty Propagation

$$f : x \mapsto s$$

$$p(s|x_u) \approx \frac{1}{N} \sum_{i}^{N} \delta(s - f(X_i)), \quad X_i \sim p(x|x_u)$$

RESULTS: UNCERTAINTY PROPAGATION

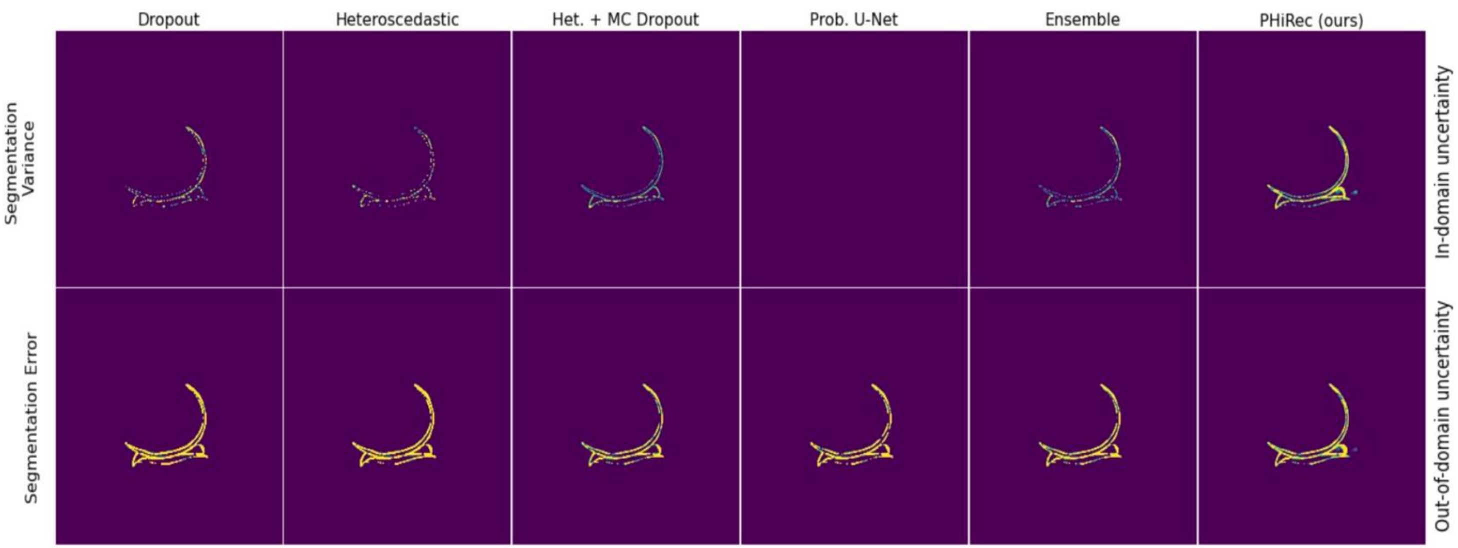

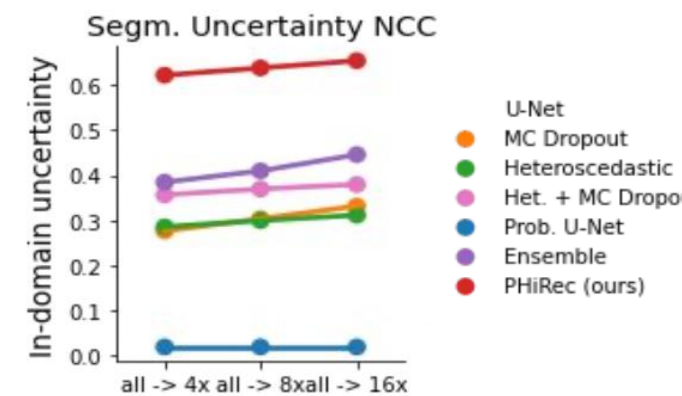

PERSONALISED ADAPTIVE MR AQUISITIONS

Goal: Use patient-specific uncertainty to stop the scan early if the certainty is high enough for a downstream decision

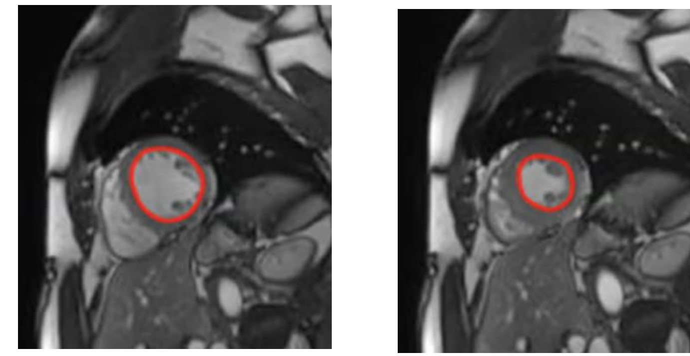

LVEF = [(End-Diastolic Volume - End-Systolic Volume) / End-Diastolic Volume]

Alice: Young healthy subject with normal cardiac anatomy. Faster acceleration possible.

Bob: Unusual Cardiac Anatomy, irregular breathing. Only low acceleration possible.

UNIVERSITÄT
LUZERN

UNCERTAINTY-GUIDED MR ACQUISITION

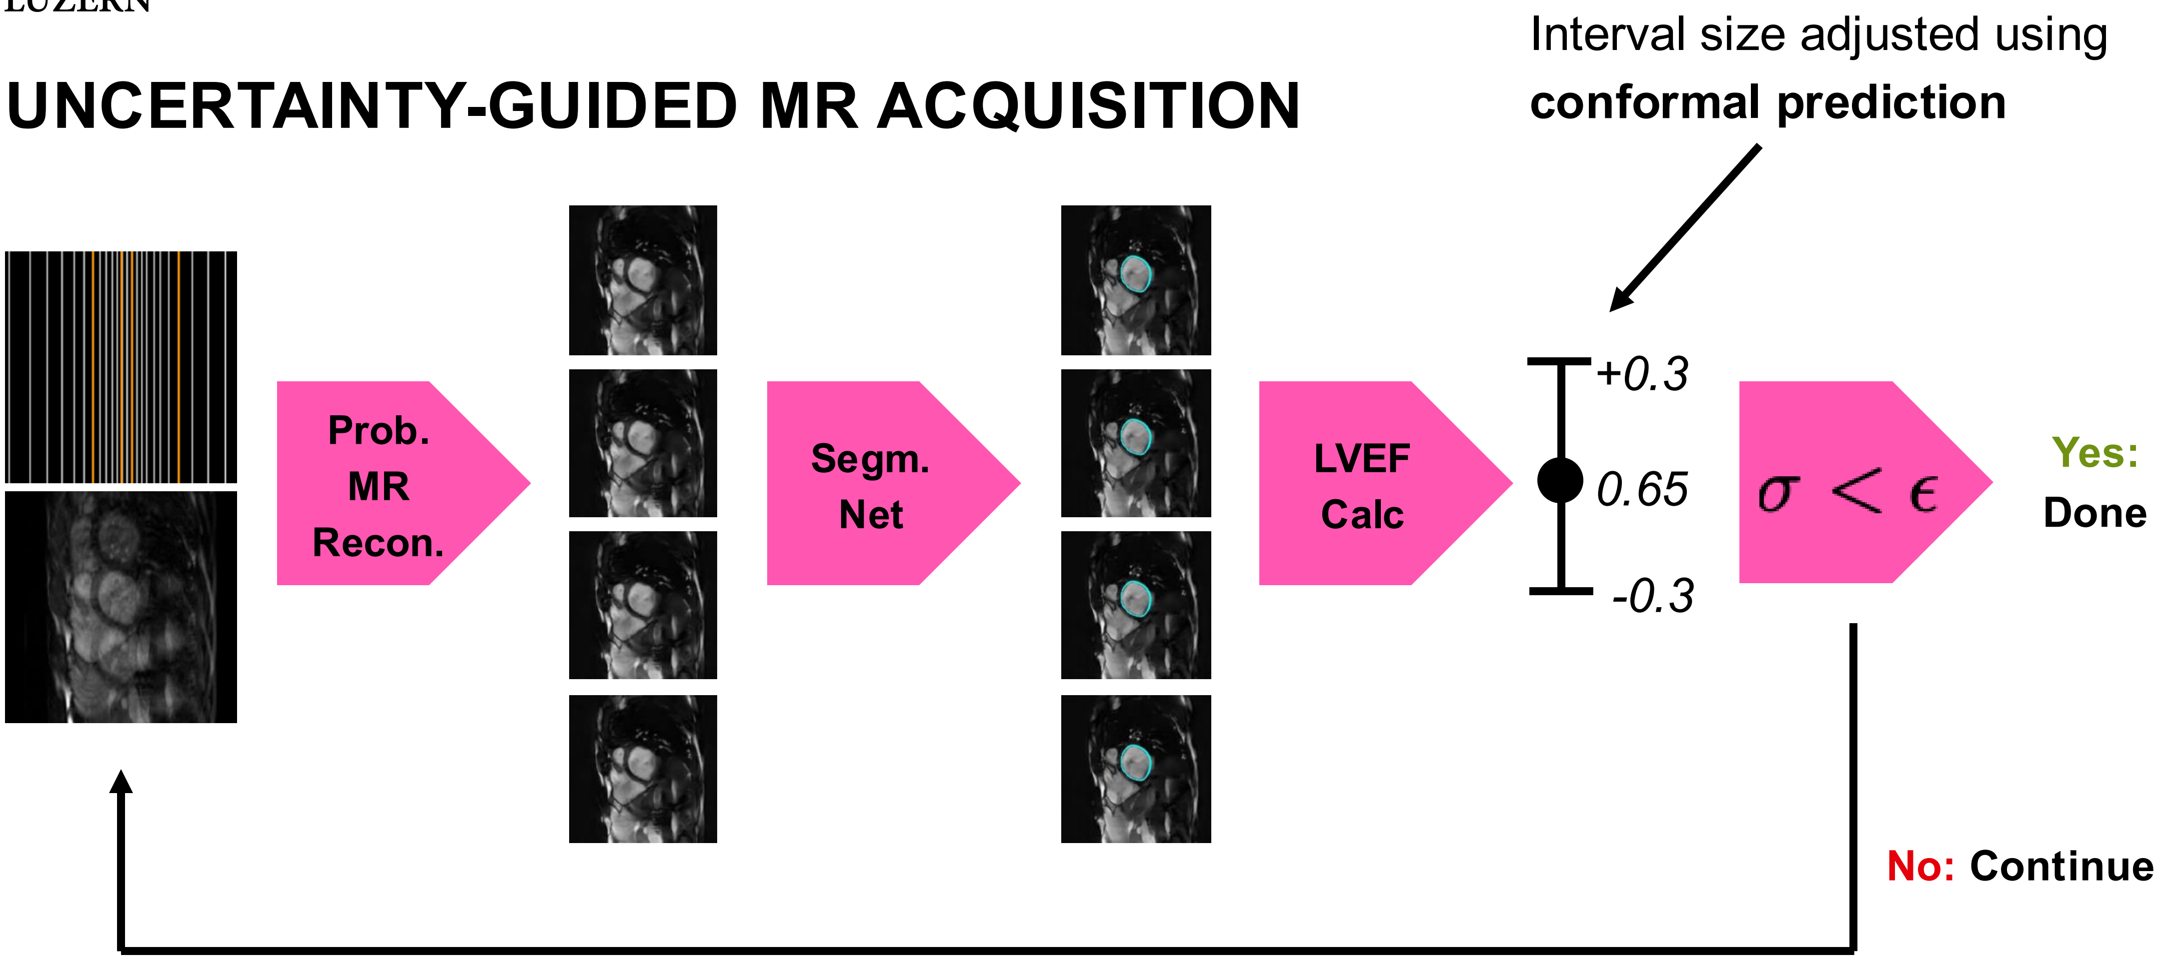

Paul Fischer, Jan Nikolas Morshuis, Thomas Küstner, Christian F Baumgartner, CUTE-MRI: Conformalized Uncertainty-based framework for Time-adaptivE MRI, Elsevier Medical Image Analysis (under review)

UNCERTAINTY-GUIDED MR ACQUISITION

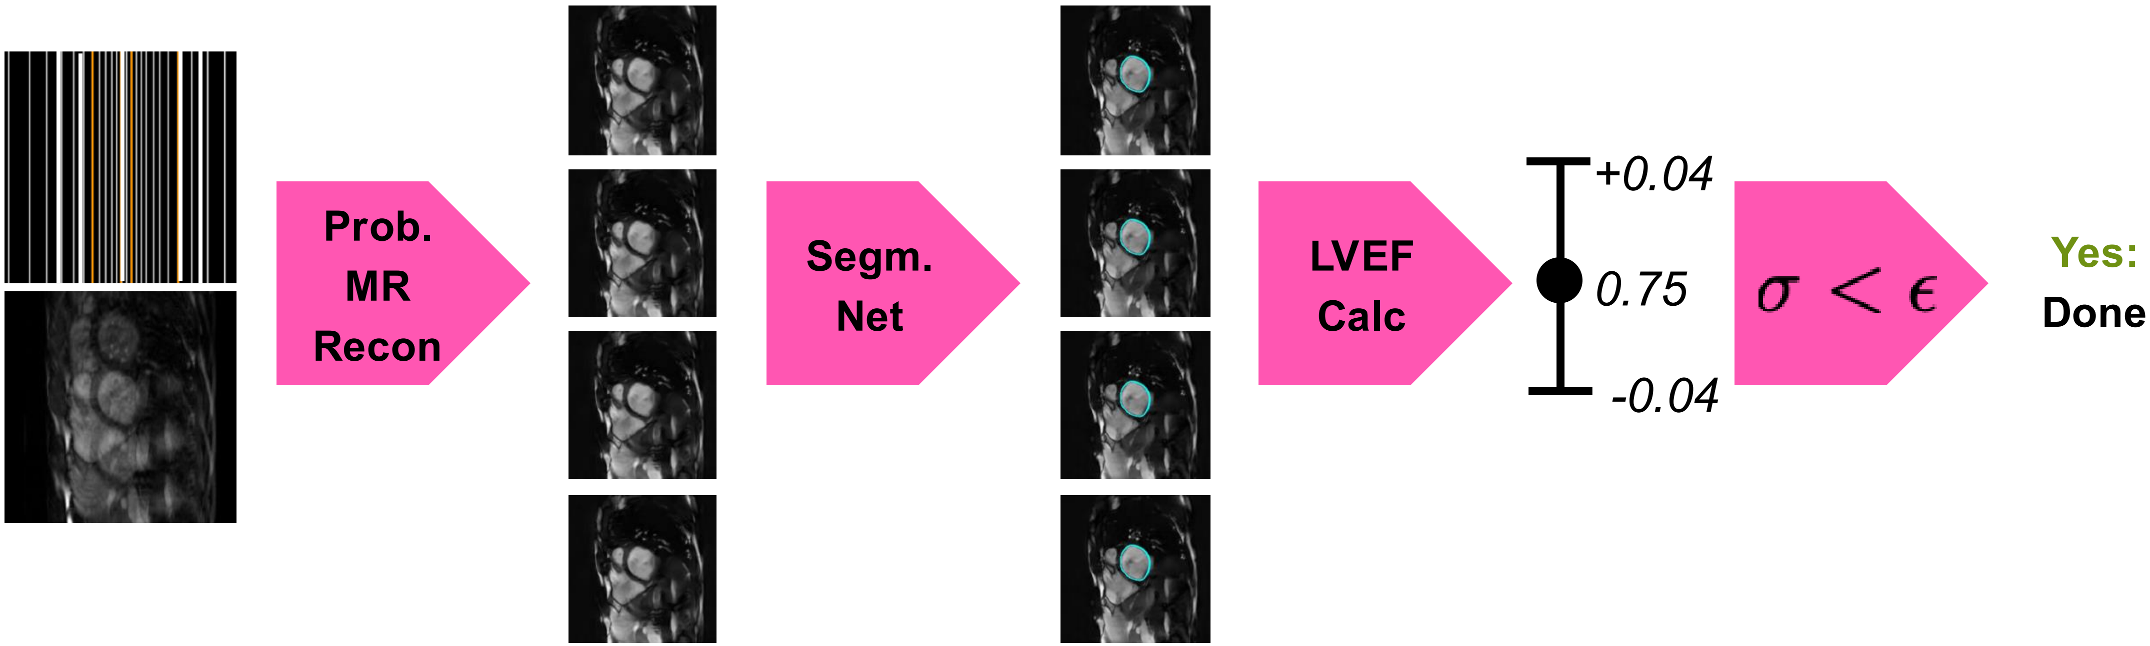

Paul Fischer, Jan Nikolas Morshuis, Thomas Küstner, Christian F Baumgartner, CUTE-MRI: Conformalized Uncertainty-based framework for Time-adaptivE MRI, Elsevier Medical Image Analysis (under review)

UNIVERSITÄT
LUZERN

CASE EXAMPLE: “EASY” CASE

CASE EXAMPLE: “DIFFICULT CASE”

EXPLAINABLE WHOLE SLIDE IMAGE DIAGNOSIS WITH AI

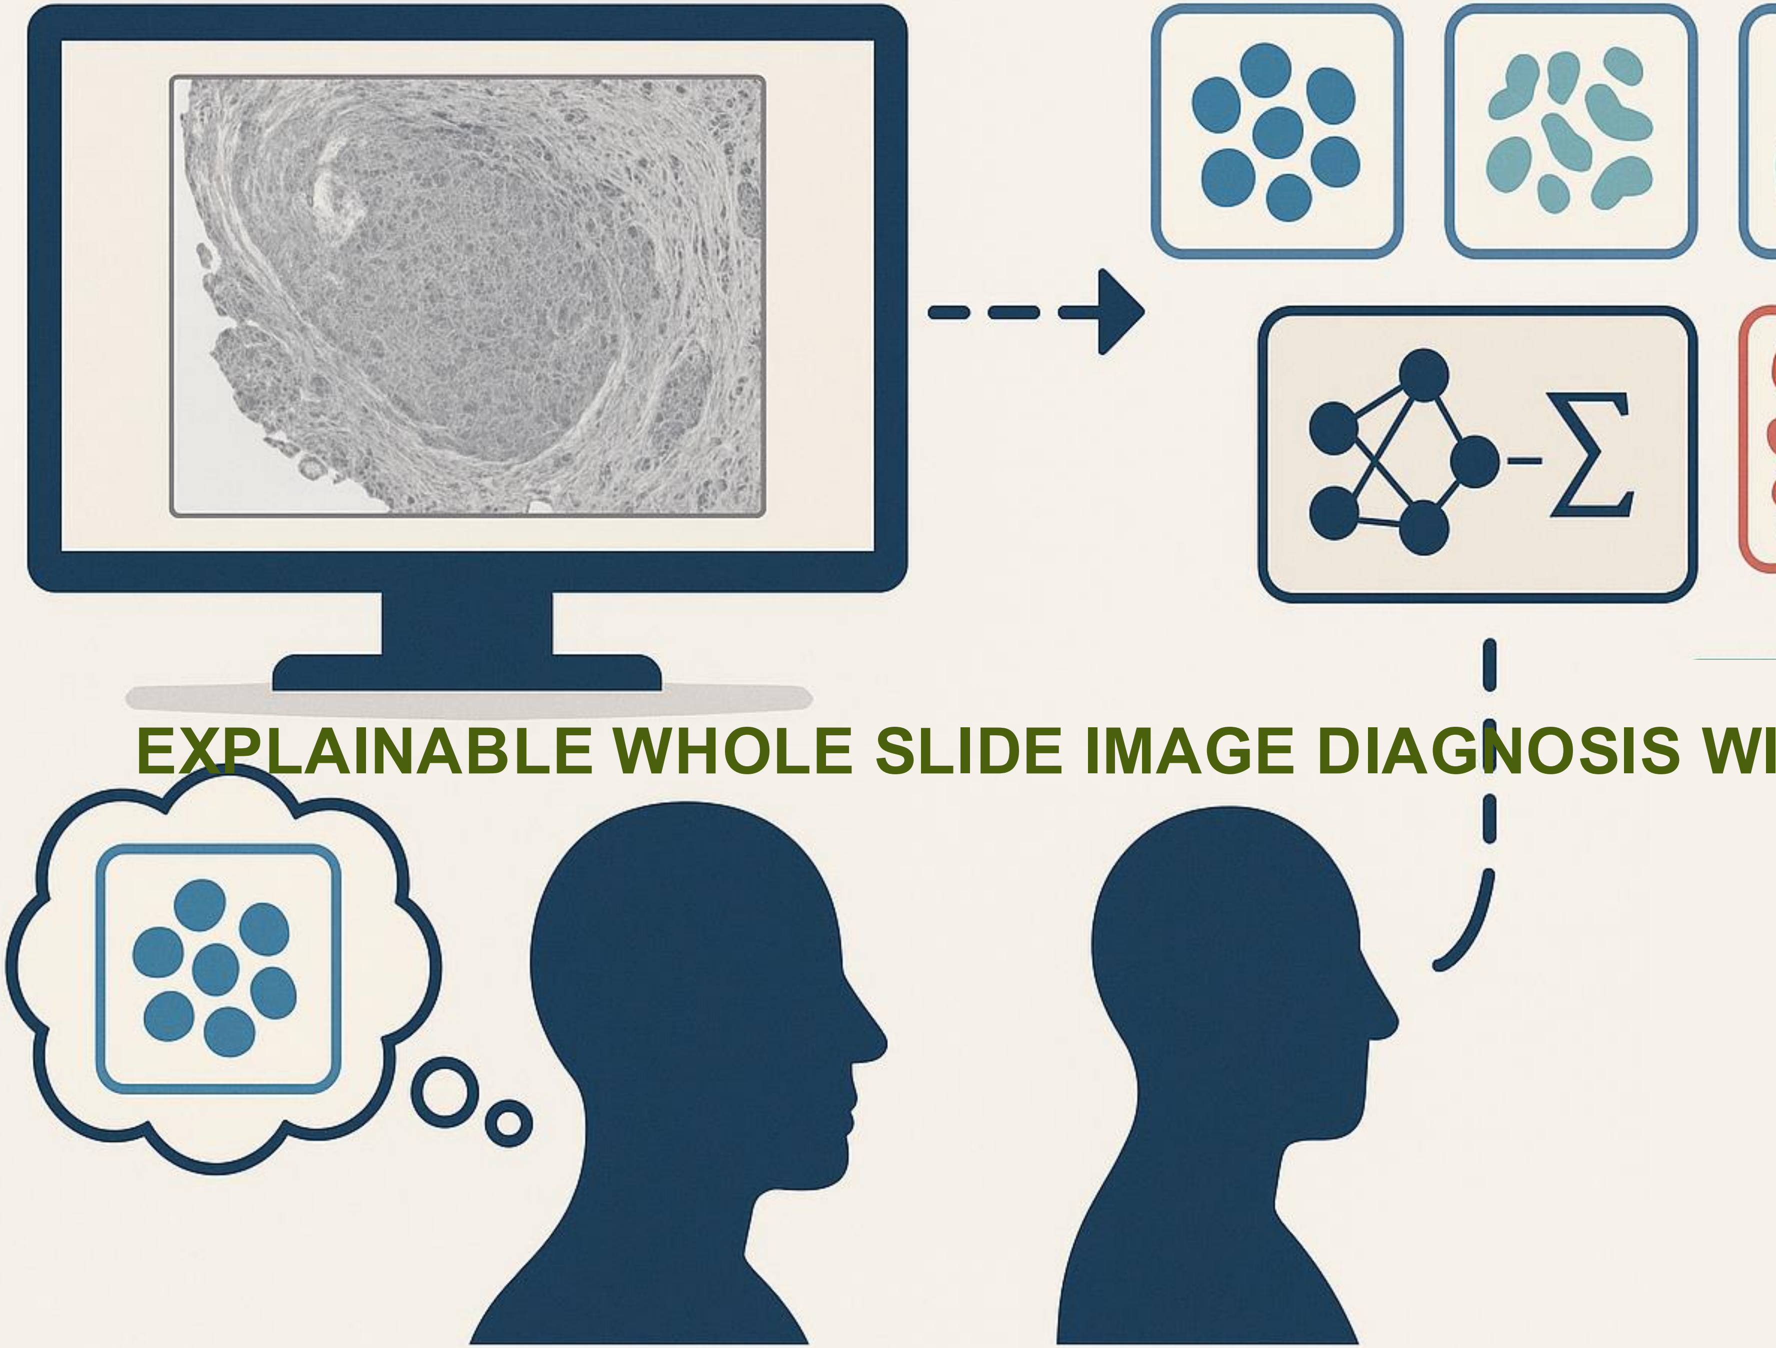

Susu Sun
PhD Student

WHOLE SLIDE IMAGES (WSI)

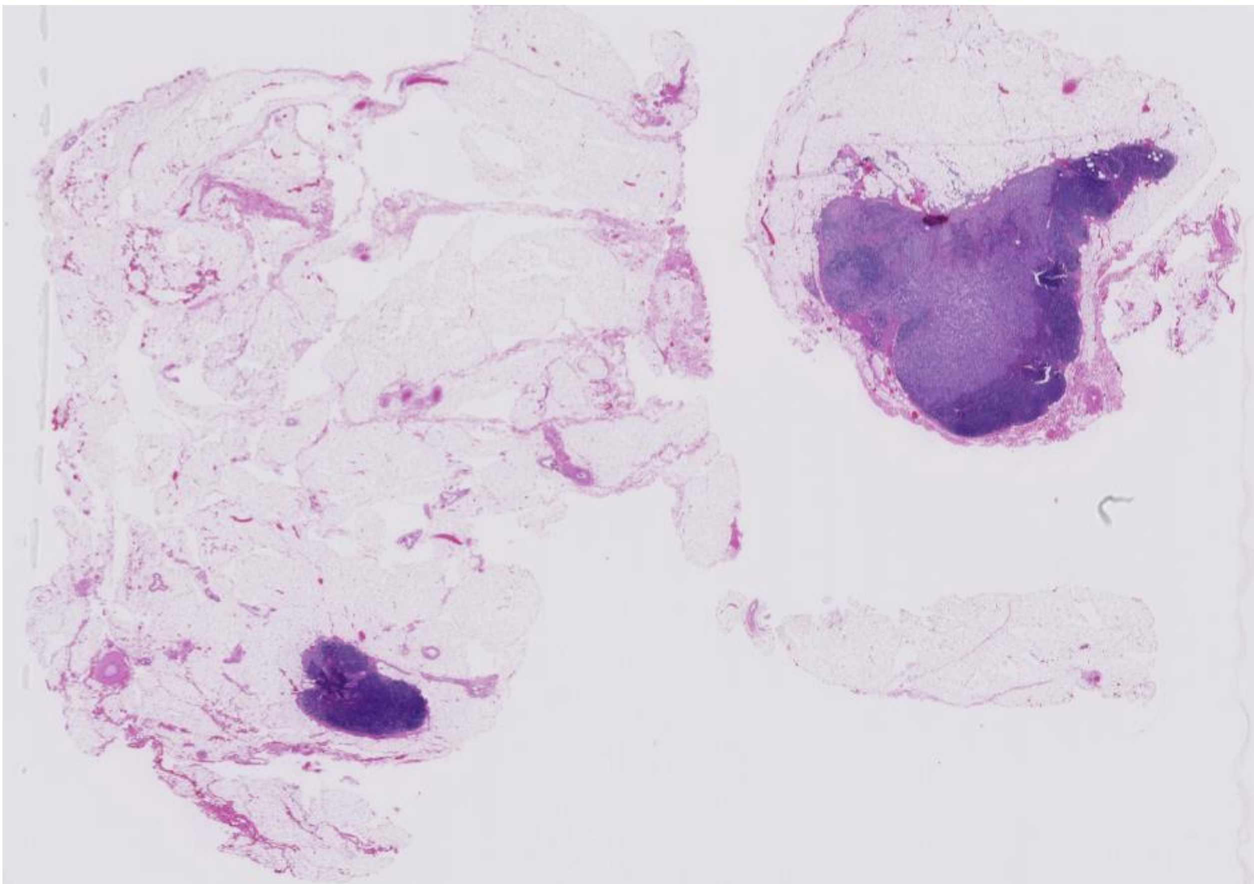

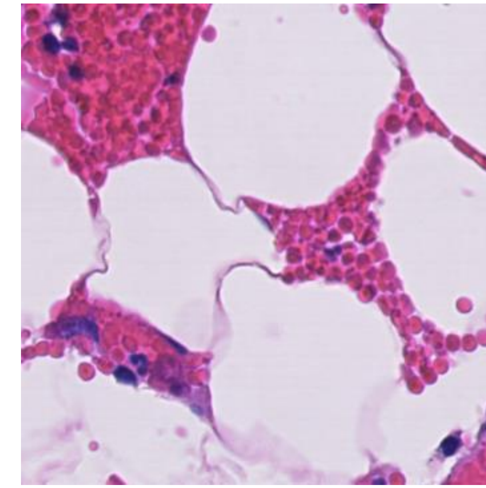

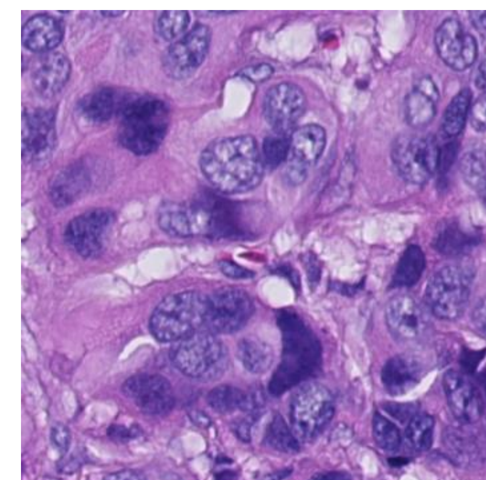

Challenge for AI: WSI are huge (on the order of 100’000 x 100’000 pixels)

UNIVERSITÄT
LUZERN

PRIOR STATE-OF-THE-ART

Multi-instance Learning (MIL)

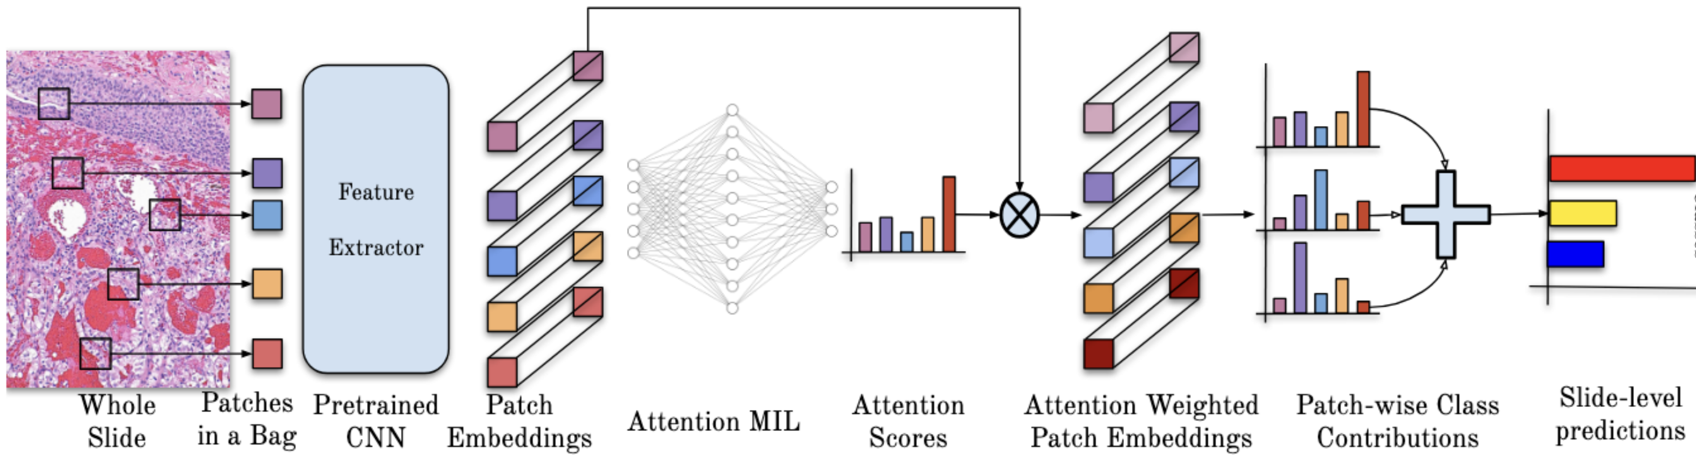

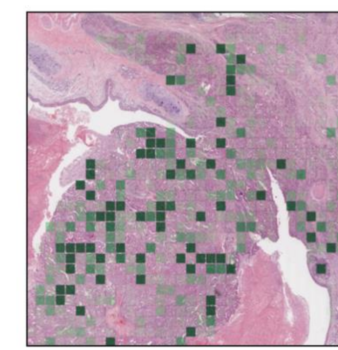

Attention Maps

Problems:

- Knowing *where* is not the same as knowing *why*
- Attention maps are known to not accurately reflect the model's decision

Javed, S. A., Juyal, D., Padigela, H., Taylor-Weiner, A., Yu, L., & Prakash, A. (2022). Additive mil: Intrinsically interpretable multiple instance learning for pathology. *Advances in Neural Information Processing Systems*, *35*, 20689-20702.

OUR IDEA: FIRST STEP - DISCOVER AND NAME CONCEPTS

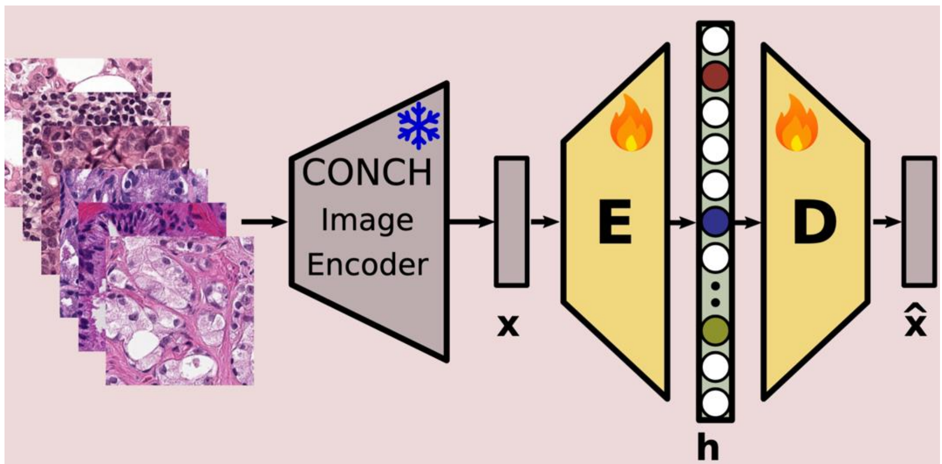

A **sparse autoencoder** is used to compress WSI patches to just a few informative neurons

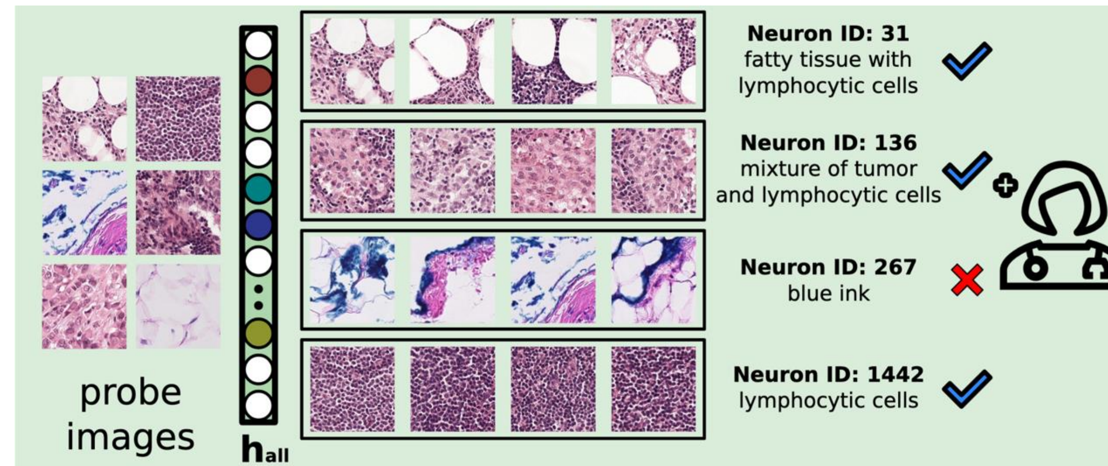

We check what type of patch activates each neuron the most and ask a pathologist to name them

Susu Sun, Dominique van Midden, Geert Litjens, and Christian F. Baumgartner. "Prototype-Based Multiple Instance Learning for Gigapixel Whole Slide Image Classification." *Proc. MICCAI* (2025).

SECOND STEP: ENCODE TRAINING INTO NEURONS AND TRAIN A MIL APPROACH ON THEM

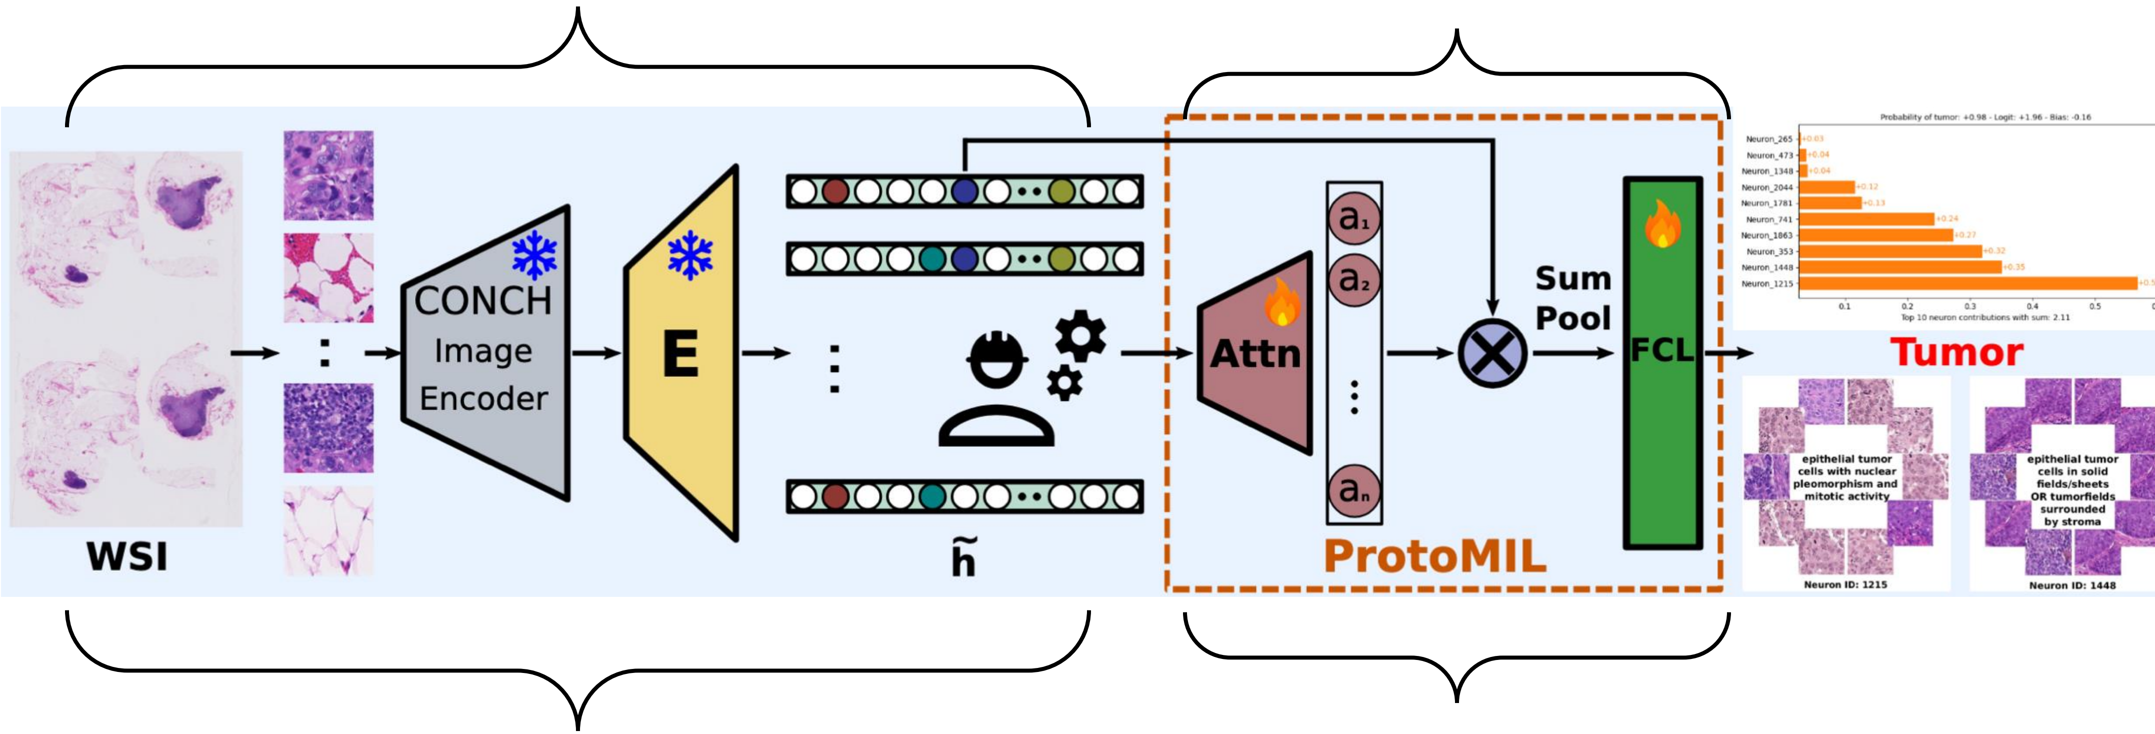

Applying the encoder part of the sparse autoencoder from before

Use a very simple MIL model for classification

OUTPUT OF OUR PROPOSED MODEL

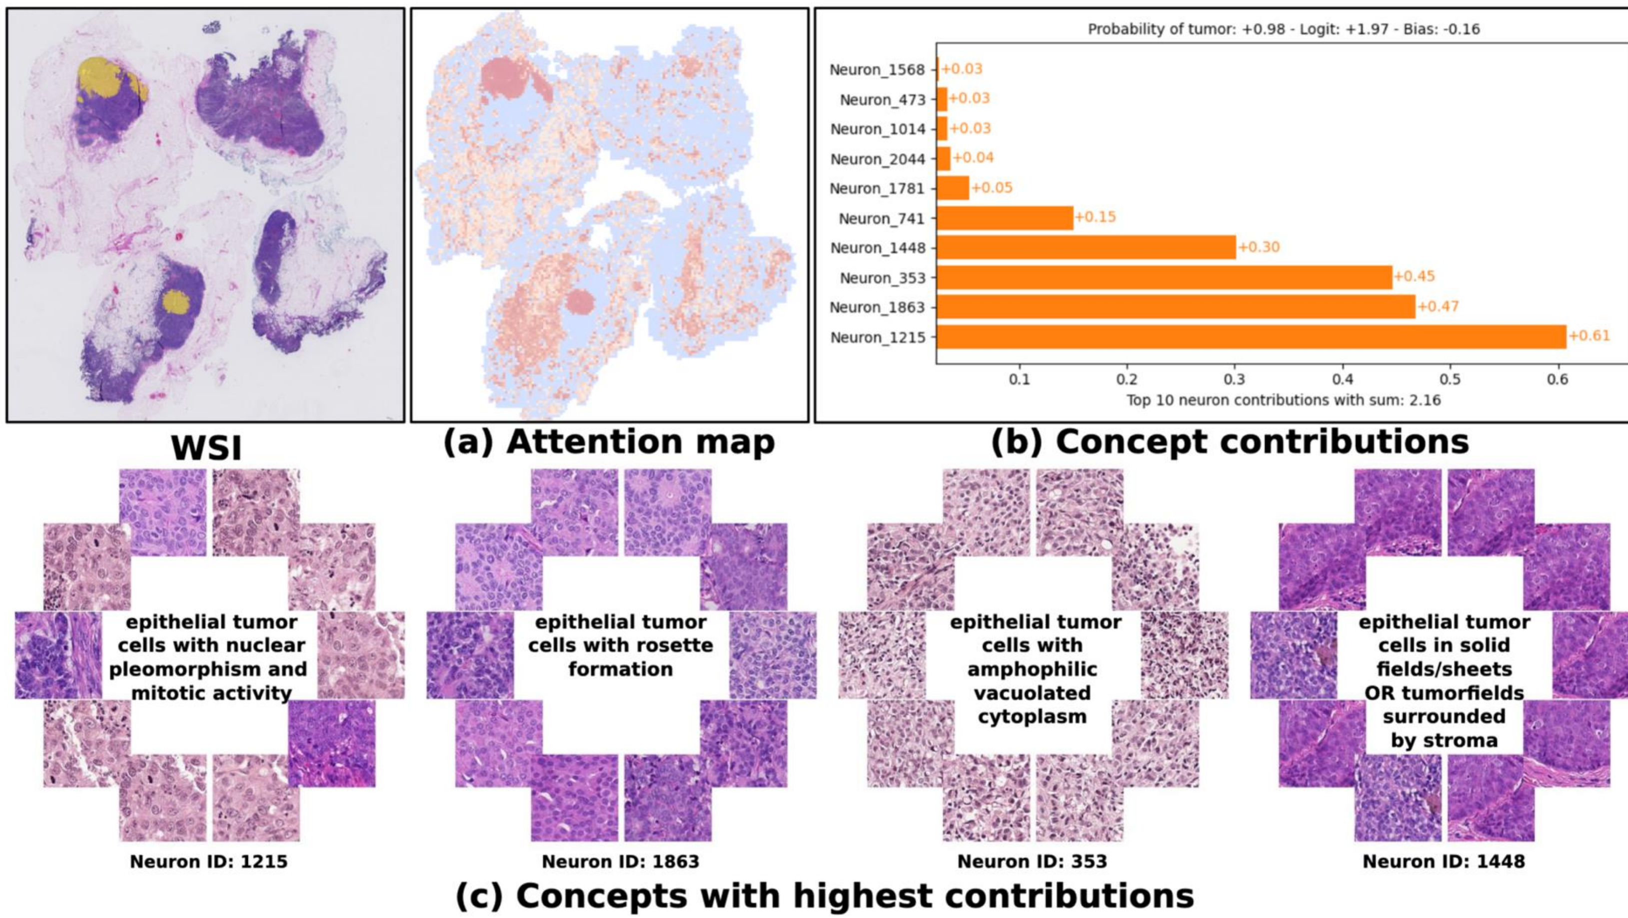

UNIVERSITÄT
LUZERN

REMOVING UNWANTED SHORTCUT LEARNING

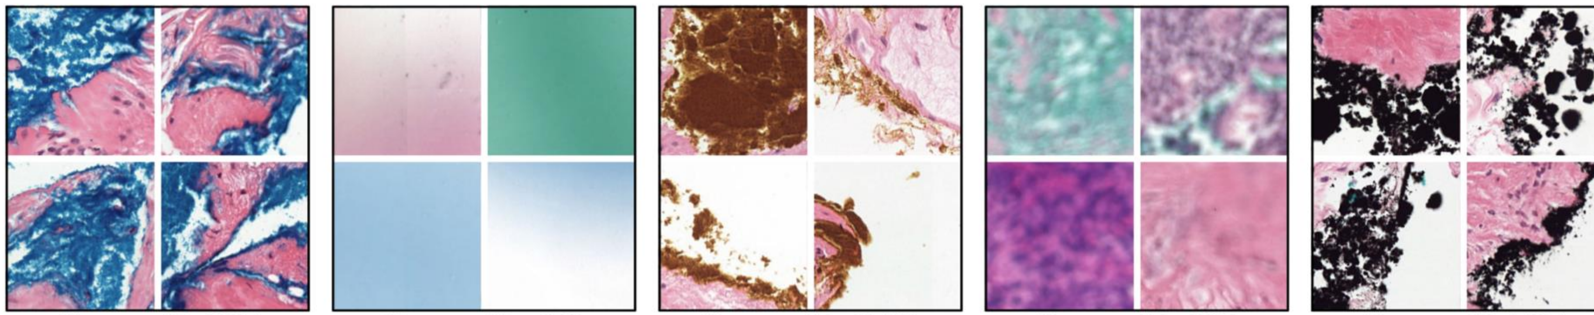

Neuron ID: 267 blue ink | Neuron ID: 439 pen marker | Neuron ID: 1509 yellow ink | Neuron ID: 908 out of focus | Neuron ID: 585 black pigments

Table 1. Classification performance measured by Accuracy and AUC.

Model	Camelyon16		PANDA	
	Acc.	AUC	Acc.	AUC
ABMIL (image)	0.922	0.908	0.892	0.953
CLAM (image)	0.915	0.966	0.884	0.979
TransMIL (image)	0.938	0.950	0.939	0.977
AdditiveMIL (image)	0.875	0.883	0.905	0.958
ProtoMIL (concept)	0.907	0.918	0.916	0.970
ProtoMIL (intervened concepts)	0.926	0.913	0.916	0.964

THANK YOU FOR YOUR ATTENTION

Paul Fischer
PhD Student/Post-doc Uni Basel

Susu Sun
PhD Student

Prof. Thomas Küstner
University Hospital Tübingen

Prof. Geert Litjens
Radboud University Medical Center

Get in touch: christian.baumgartner@unilu.ch