1

Fast detection of slender bodies in high density microscopy data

Albert Alonso

Task

Solution

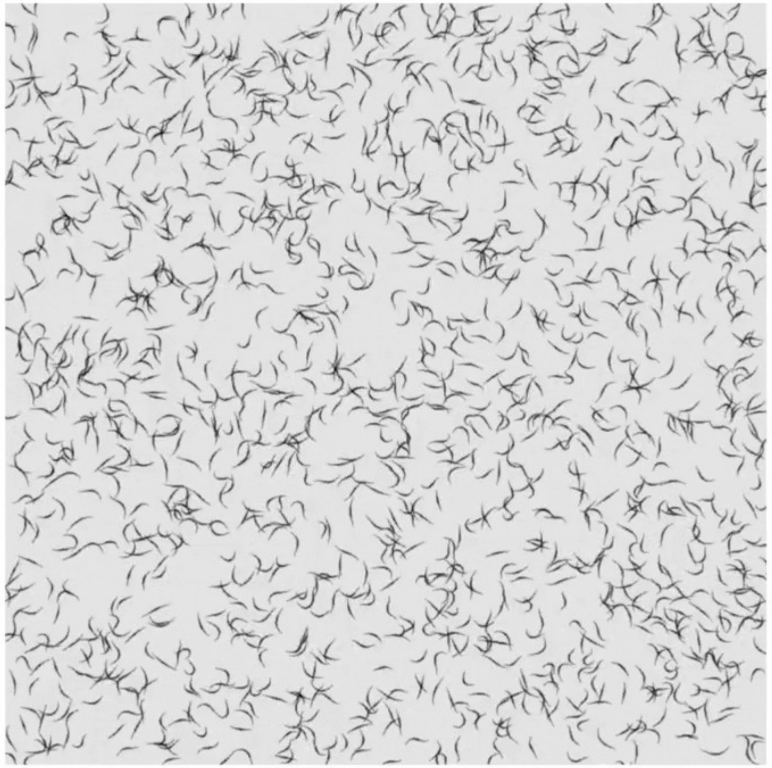

de(ep)tangle

Fast detection of slender bodies in high density microscopy data

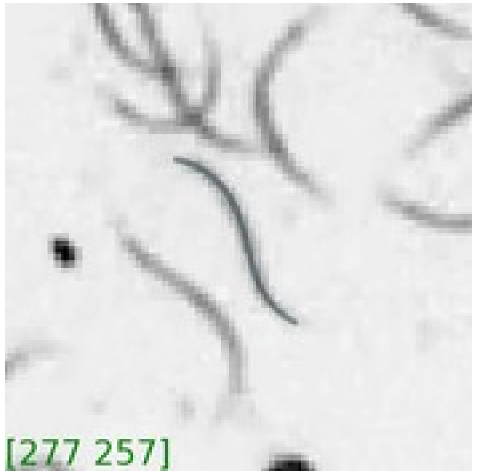

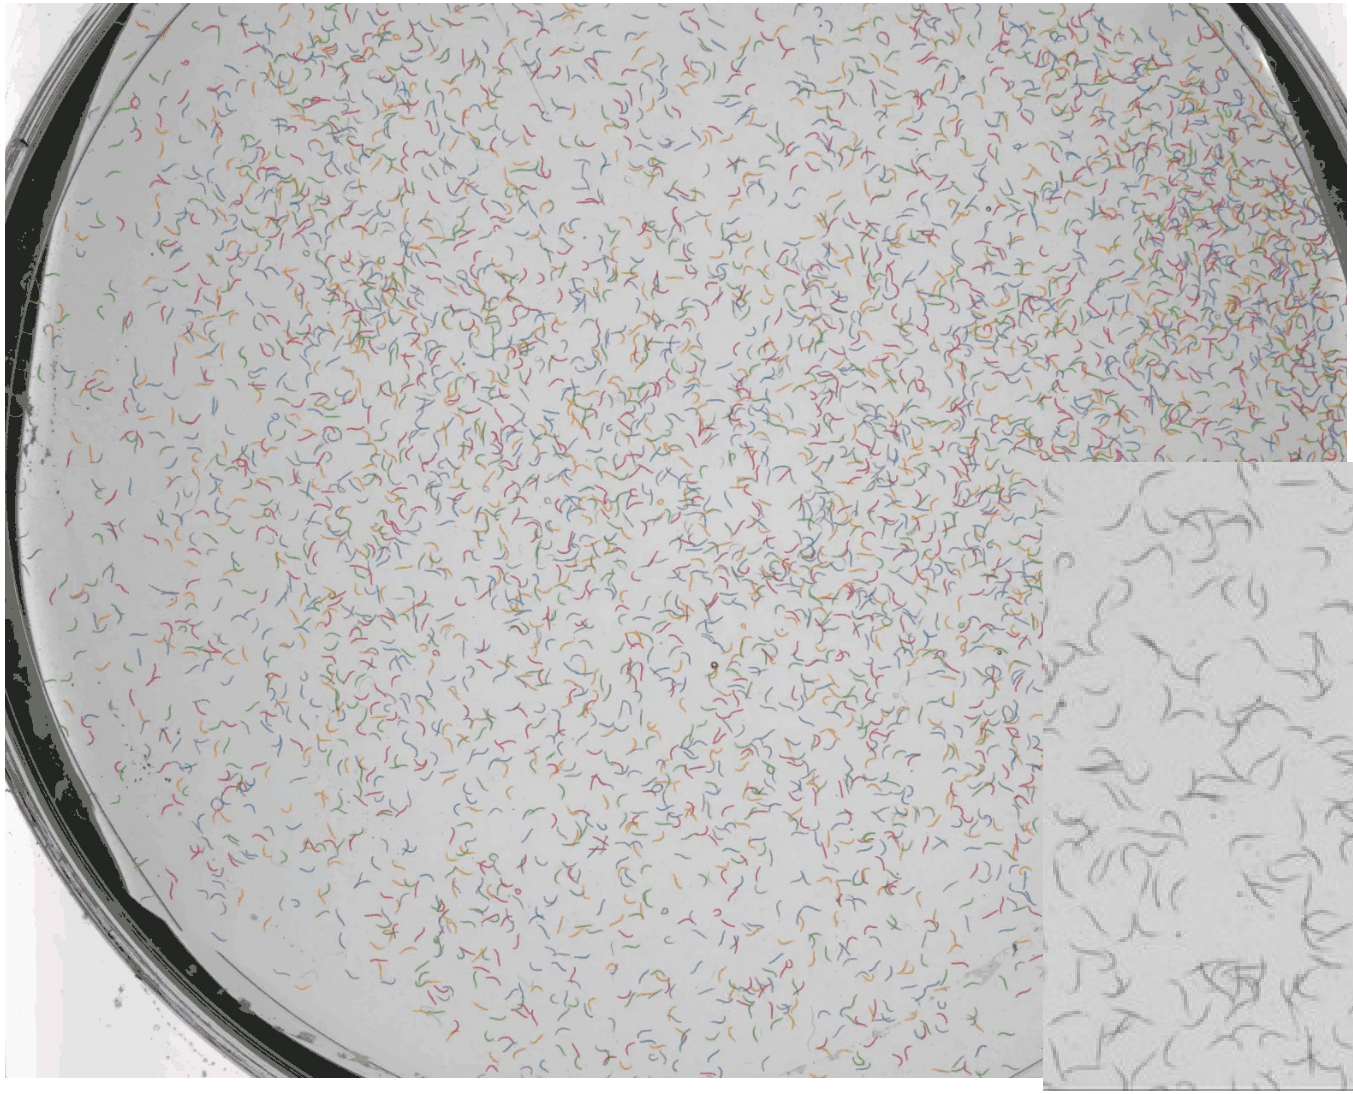

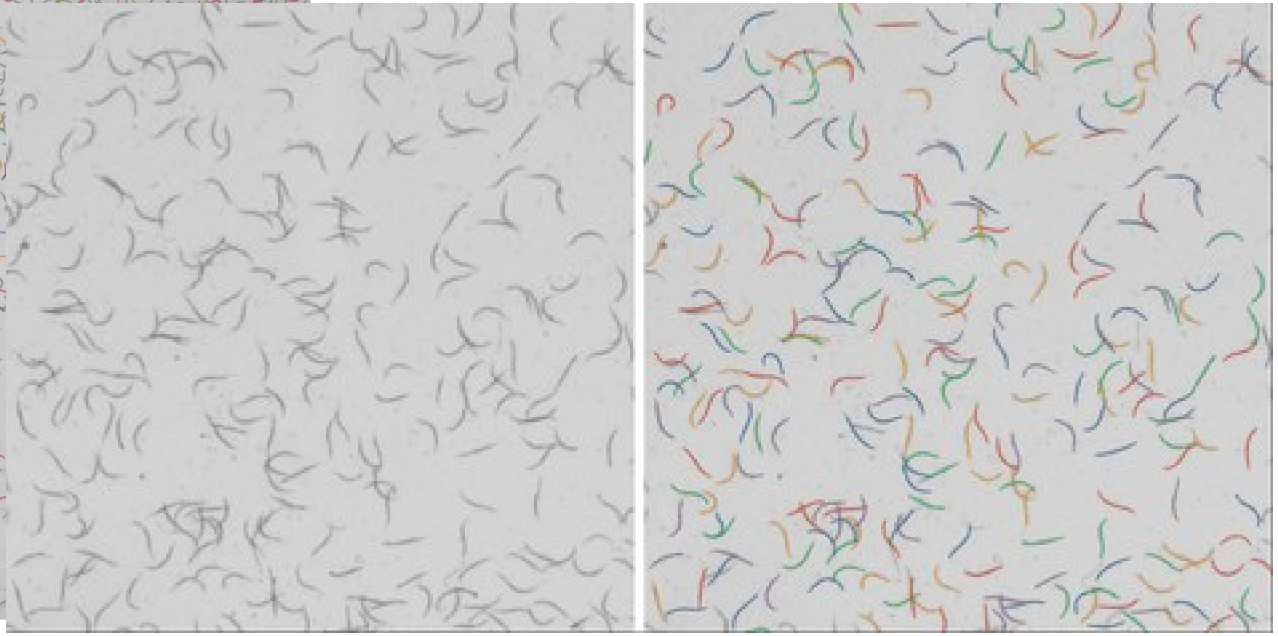

de(ep)tangle

2

Longitudinal Self-Supervised Deep Learning for Radiotherapy

Alejandro Cortina Uribe

Motivation

- Radiation therapy for cancer treatment.

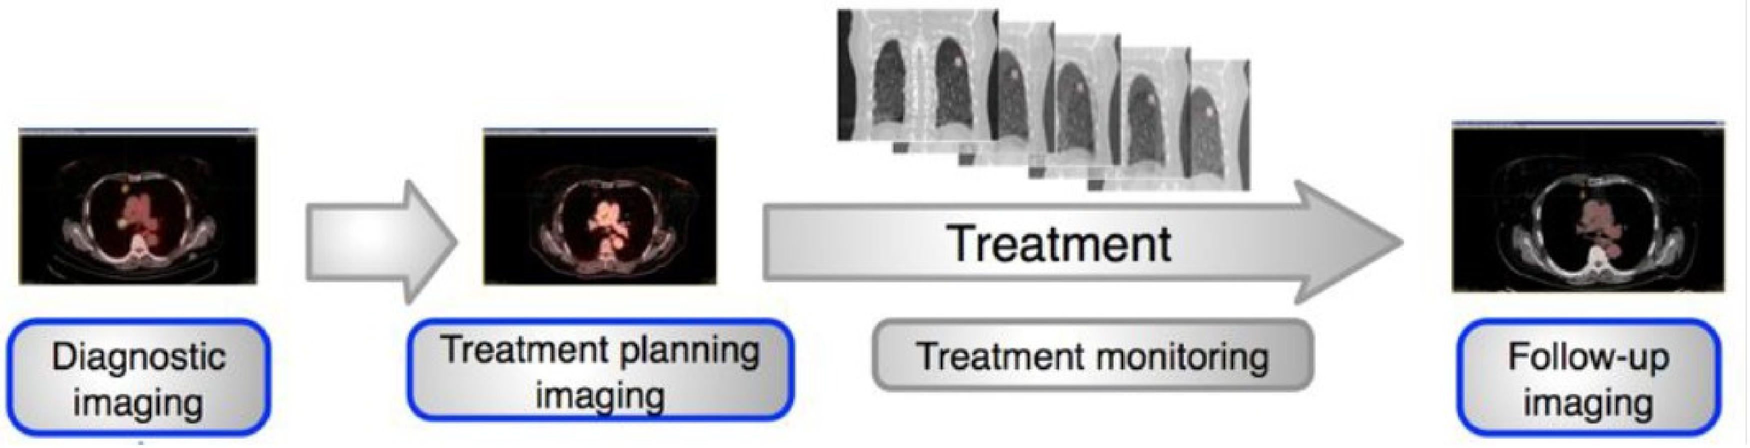

- Sequential imaging data (treatment CBCTs) we could potentially extract more information from, about the **development** of the treatment.

Longitudinal Self-Supervised Deep Learning for Radiotherapy

Alejandro Cortina Uribe

- ❑ Unlabeled data → SSL

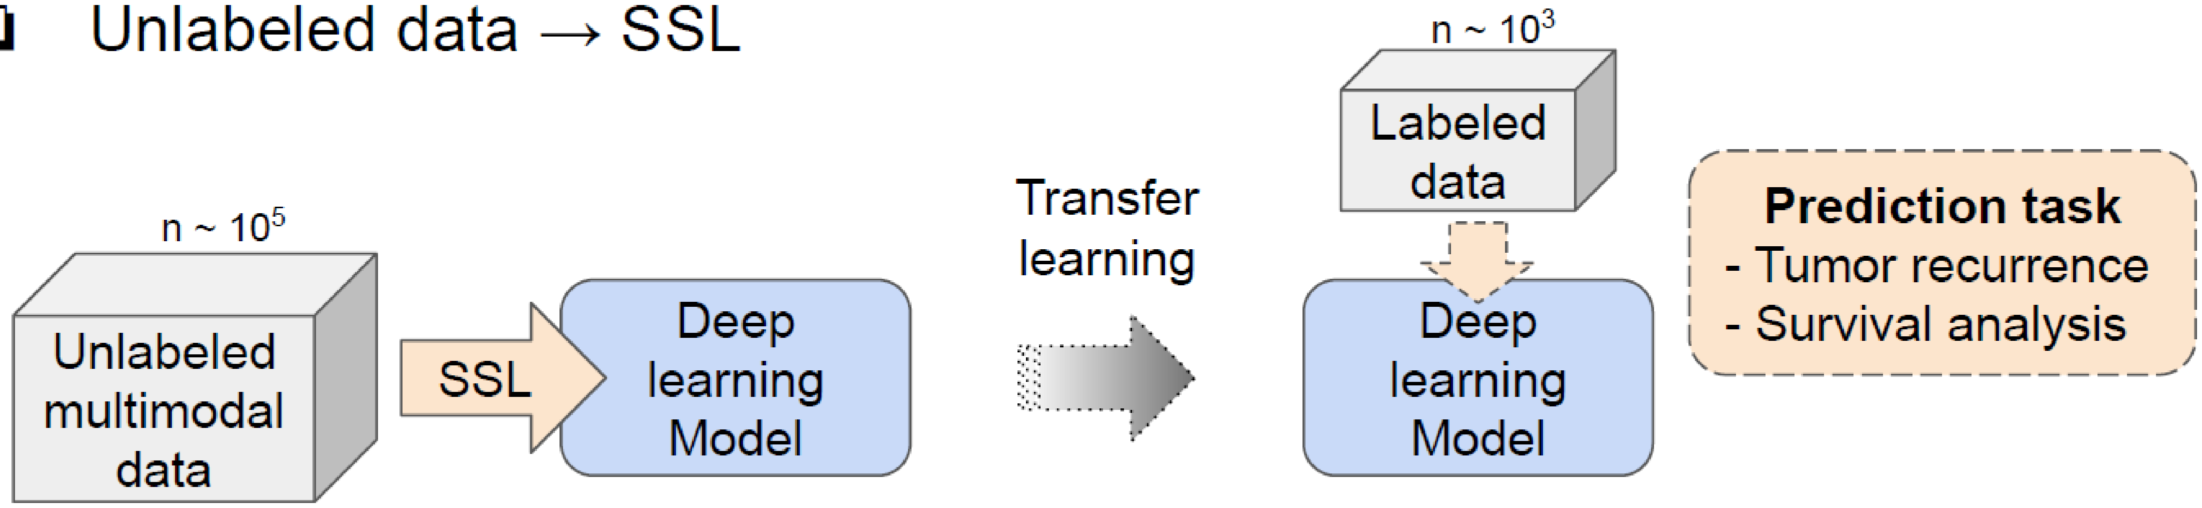

Hypotheses

- ❑ SSL improves performance for downstream tasks.
- ❑ By modeling time we can identify **image-based responses** to treatment.

3

PHENOTYPING SEED SHATTERING OF PERENNIAL RYEGRASS

INTRODUCTION

Perennial ryegrass (*Lolium perenne* L.) is one of the most extensively produced forage crops in temperate areas used for feed and as turf.

It is an important activity to increase seeds for commercial market by achieving high seed yield however, seed shattering is a major constraint.

Information on seed shattering of perennial ryegrass, the stage at which it occurs and an efficient method of evaluating shattering is limited.

OBJECTIVE

The study seeks to develop a high throughput system to phenotype seed shattering in perennial ryegrass

METHODODLGY

Four (4) methods of measuring seed shattering will be employed in this study

- Bagged inflorescence (Kiesbauer et al., 2023)
- Rolling on inflorescence (Tubbs & Chastain, 2023).
- Tray under plant
- Imaging (Ortiz et al., 2023)

ANALYSIS

Abscision zone of the spike will be assessed under microscope at varying growth stages of the plant.

Seed shattering methods will be analysed statistically using linear mixed effects models accounting for the block design, genotype and weather conditions in the field.

EXPECTED OUTCOME

- A high throughput protocol for phenotyping seed shattering in perennial ryegrass will be established
- Identify genotypes with low shattering traits under optimum environmental conditions.
- Development of high throughput protocol to guide breeding of high seed retention of perennial ryegrass.

4

Automated estimation of cardiac stroke volumes from computed tomography

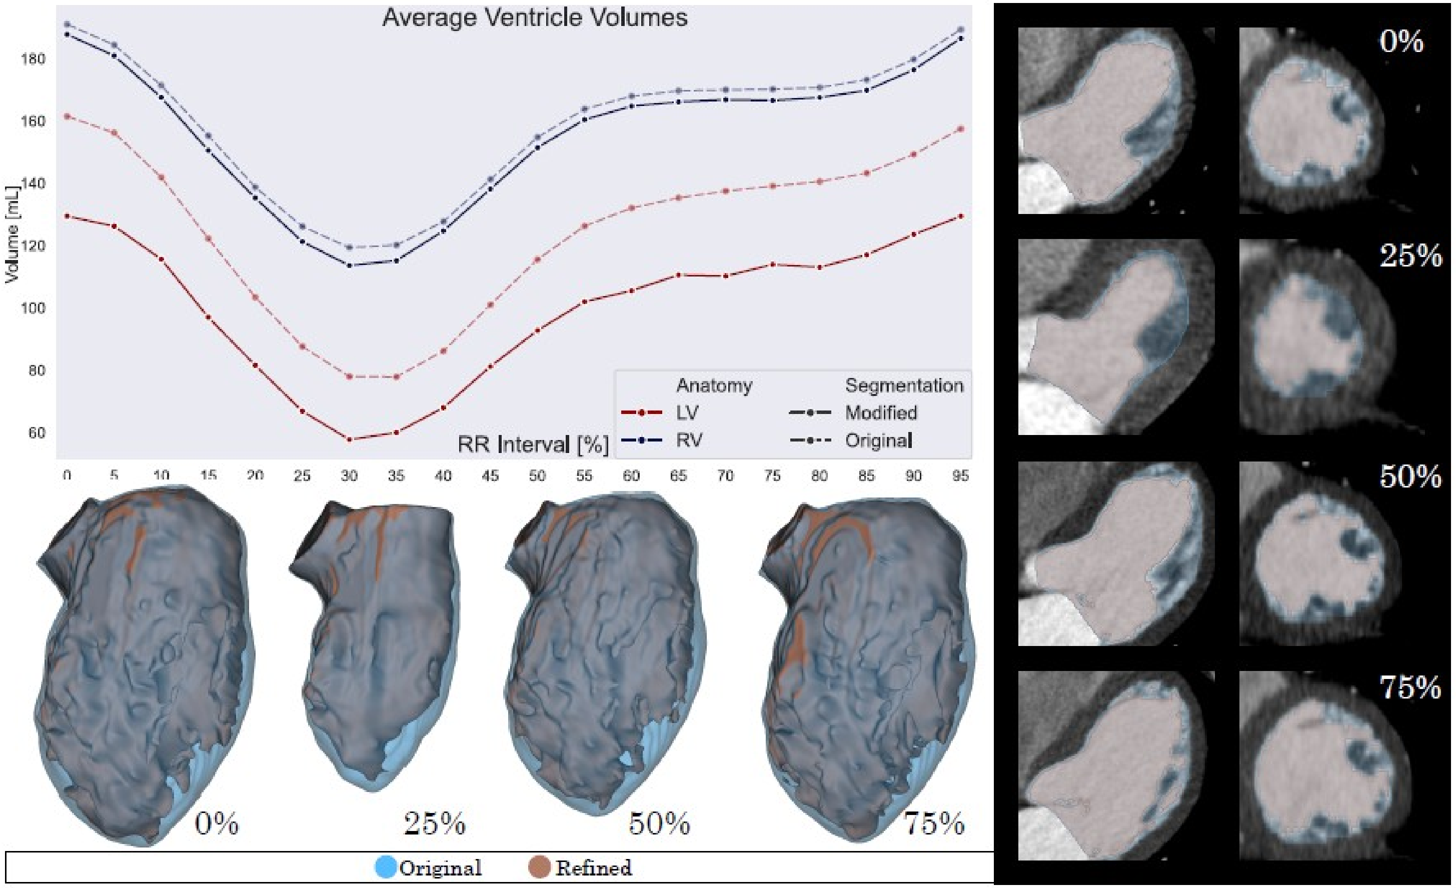

5

Using AI for Vertebral Segmentation and Fracture Detection

Andreas With Aspe

Segmentation framework

- Three-step model

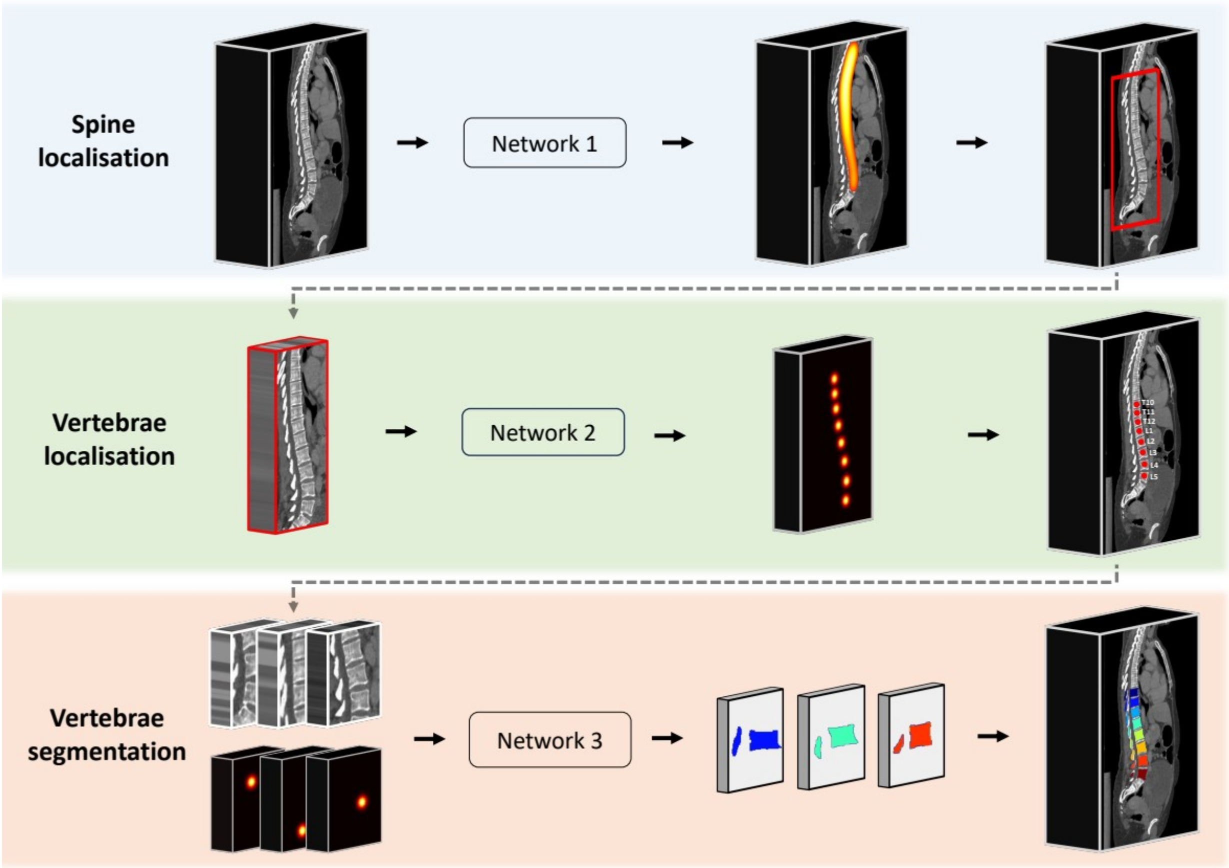

Fracture detection

- Fracture = deformation

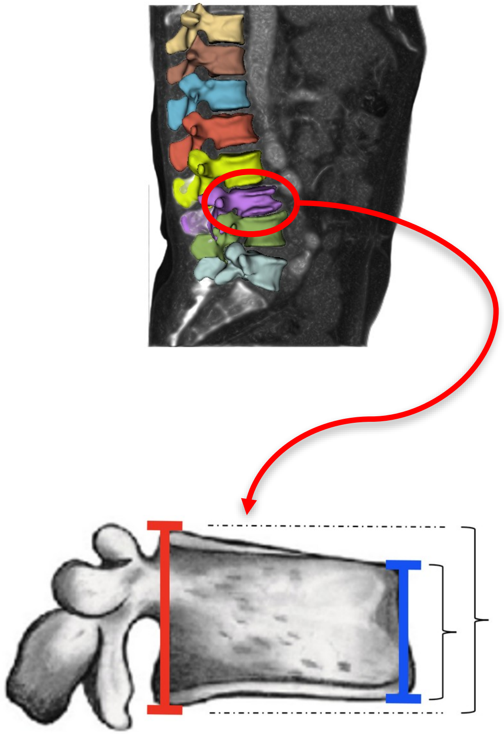

6

Eye-tracking for assessing X-rays image interpretation

Anna Anikina

UNIVERSITY OF COPENHAGEN

Purpose: Explore the potential of eye tracking data in predicting errors made by radiologists during the analysis of chest X-ray images

Input data: Features, describing the movement of the doctor's gaze while reading x-ray chest images

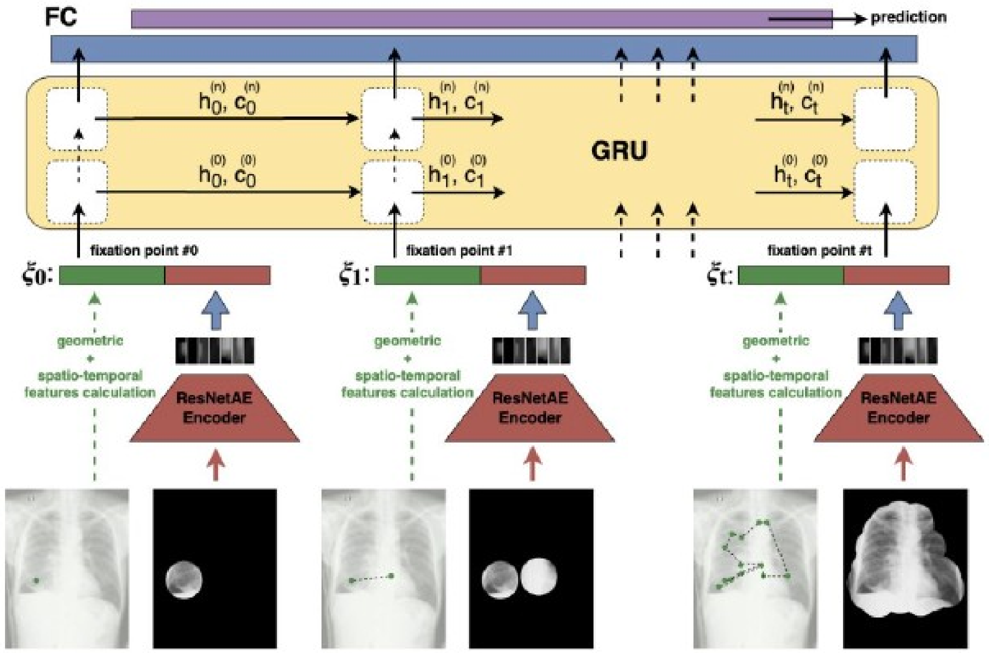

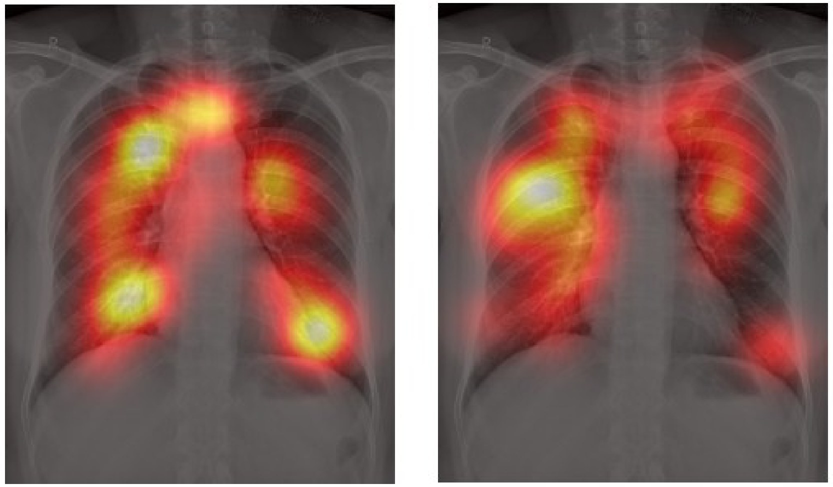

Model	Accuracy
GRU Tabular + image features	0.83
GRU Tabular features	0.82

7

AMAES: Augmented Masked Autoencoder Pretraining on Public Brain MRI Data for 3D-Native Segmentation

Asbjørn Munk, DIKU

- Brain MRI
- Masked Autoencoder pretraining for 3D segmentation
- 🧠**BRAINS-45K:** The largest public pretraining dataset

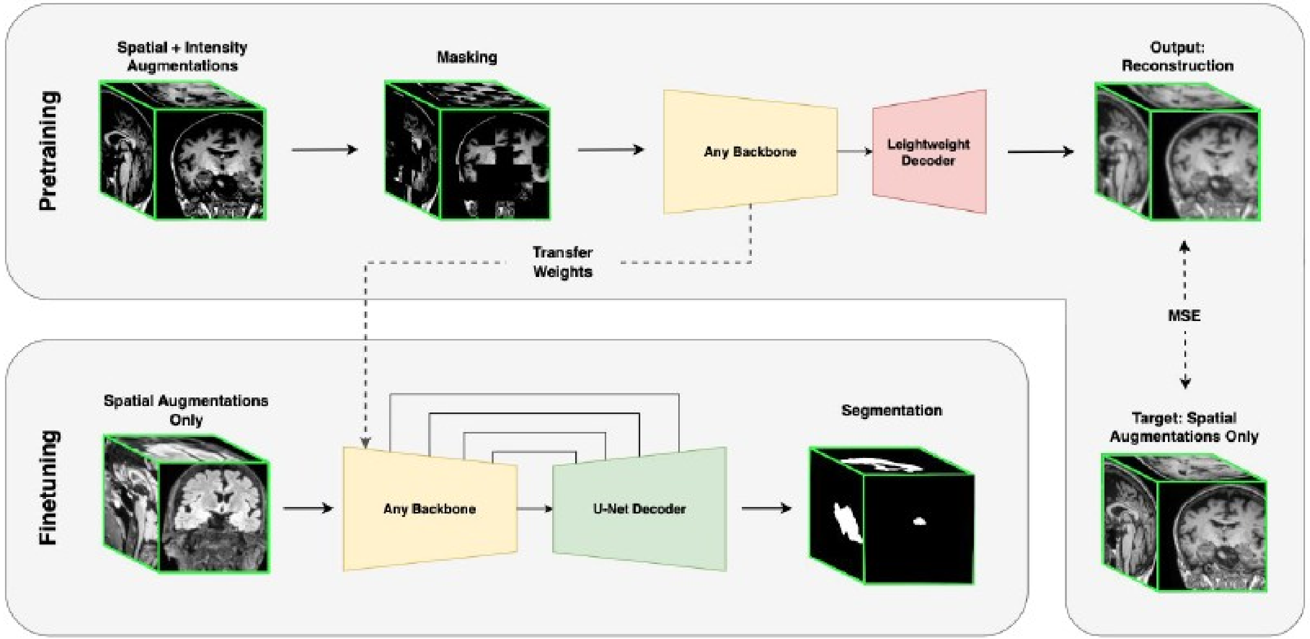

8

Unify tasks in Video Segmentation

Semi-supervised Video Object Segmentation (VOS)

Referring VOS

"A man in a red sweatshirt performing breakdance"

Interactive VOS

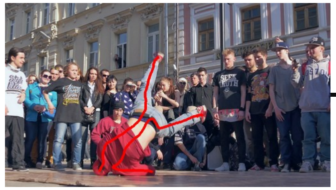

Unsupervised VOS

-

Video Instance Segmentation

Set of labels = {human, dog, cat, ...}

Propagation Model

[Caelles CVPR17, Caelles arXiv18, Khoreva ACCV18]

9

Hatzakis Lab

Cell segmentation and cell-specific intercellular PSF detections

+ automations for cell detection

How do we detect 3D PSFs in cells and quantify their characteristics?

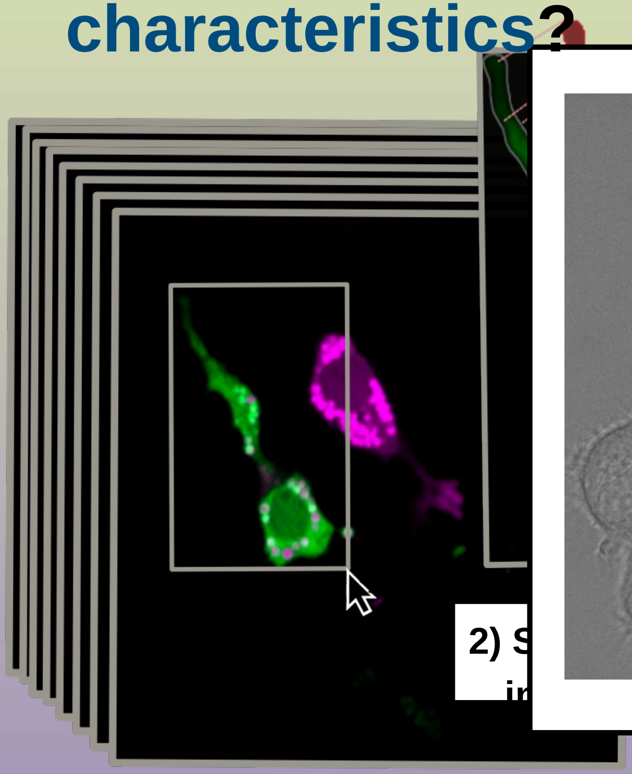

1) Select a cell with dual expression

2) S
in

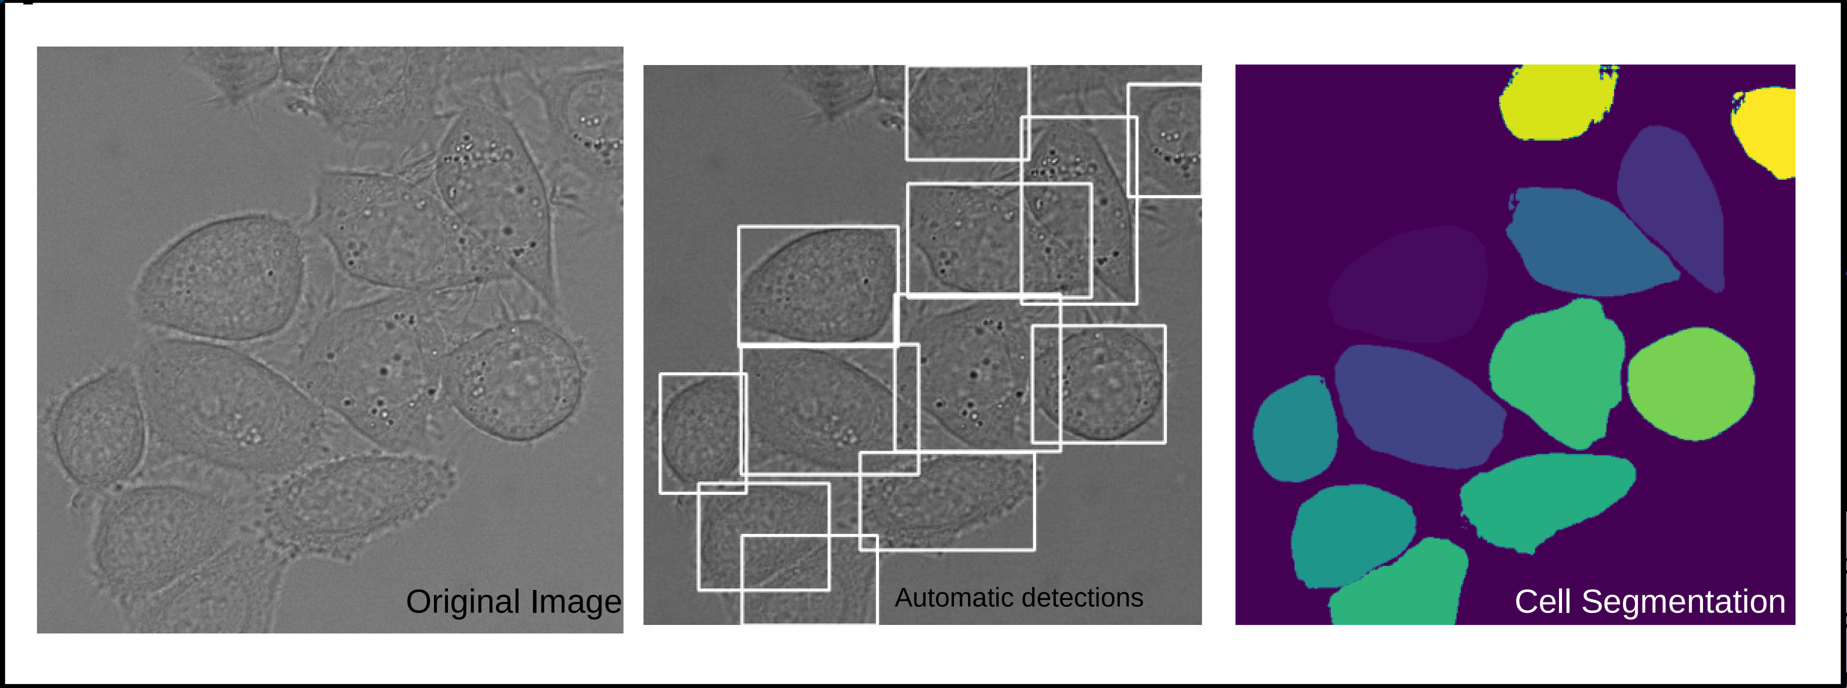

on 3D through tracking

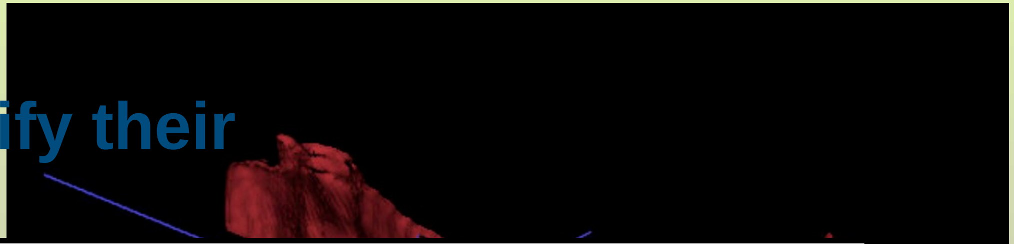

10

Volumetric Super-Resolution via Multi-Scale Transformers

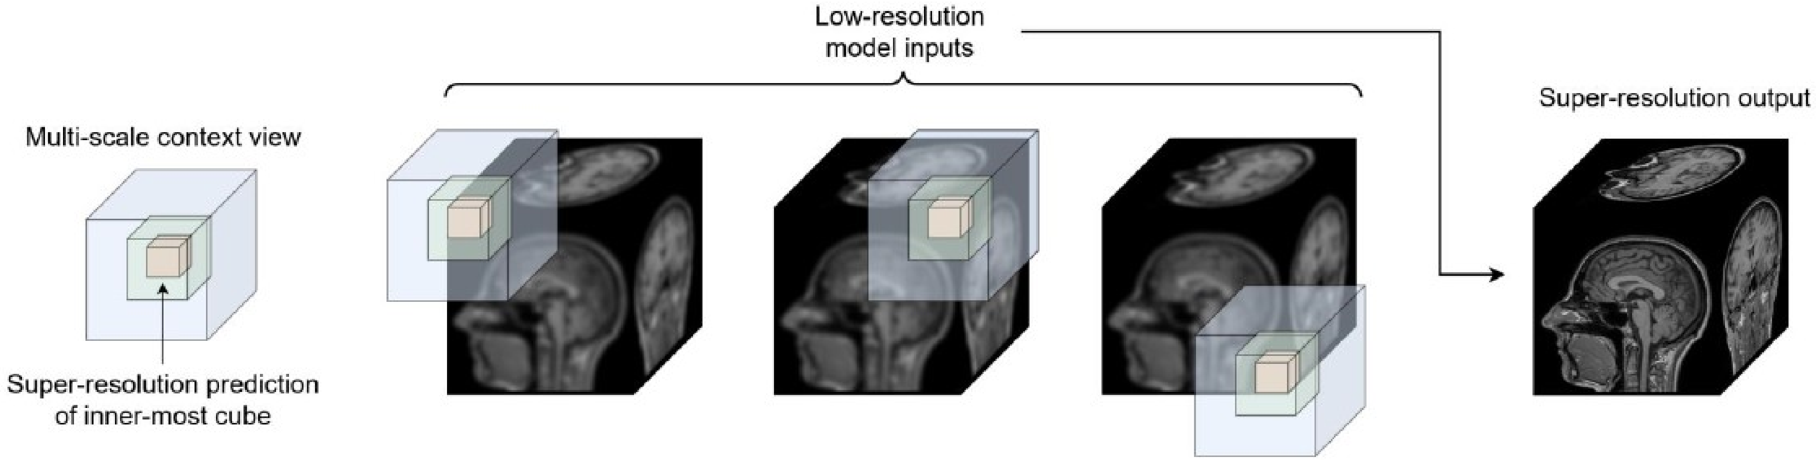

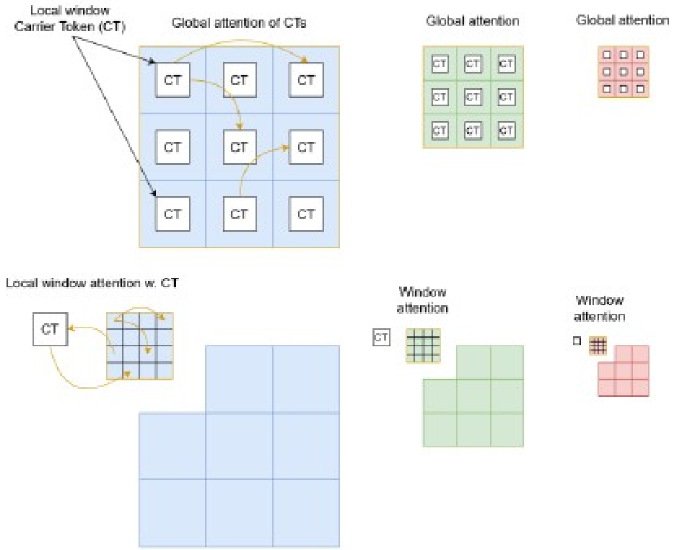

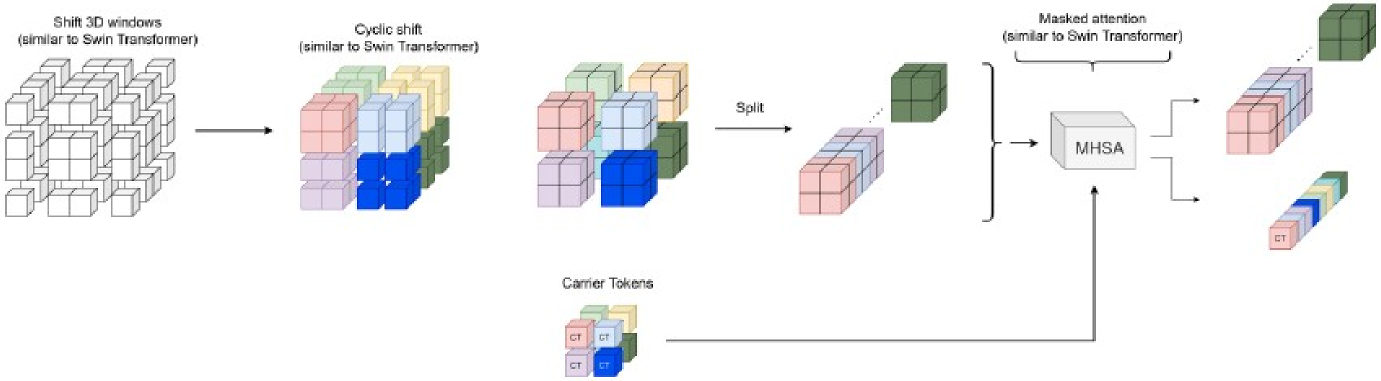

11

Synthesis of Medical Images

- Problem:
  - Medical data is sparse and data augmentation can be problematic
  - Medical images are usually very big
- Solution
  - Synthesize medical images with context taken into consideration
  - Utilize LDM to reduce memory footprint

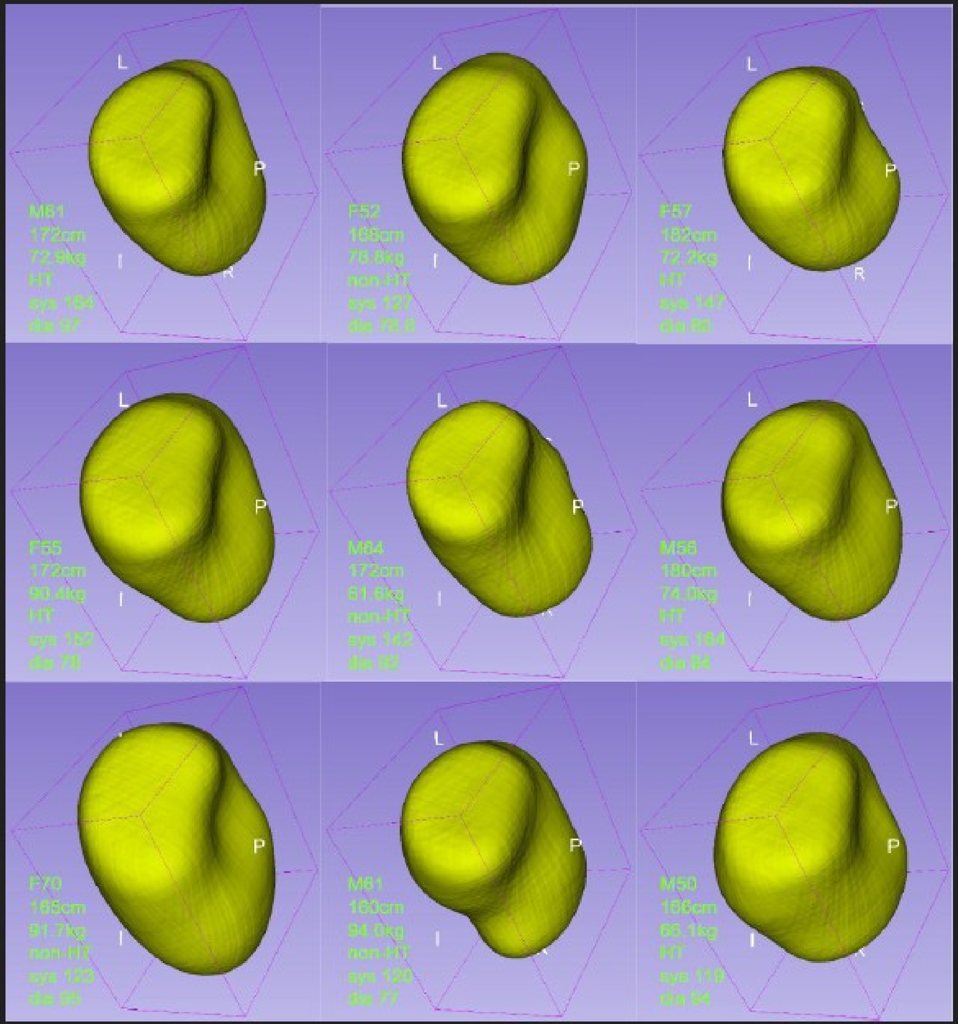

12

DO YOU WANNA KNOW WHAT YOUR UNSUPERVISED MODEL IS LOOKING AT?

DO YOU WANNA KNOW WHAT YOUR UNSUPERVISED MODEL IS LOOKING AT?

THEN YOU NEED **NEM-U**

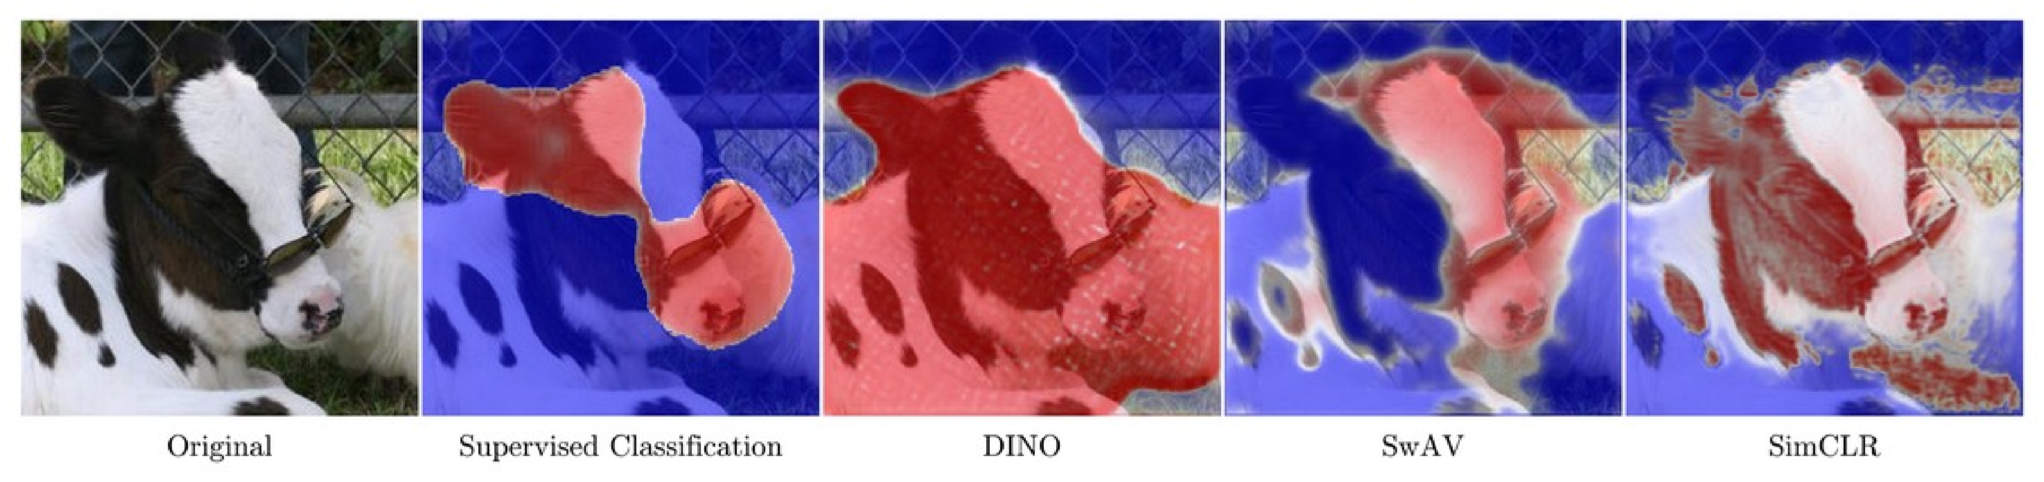

ICML
International Conference
On Machine Learning

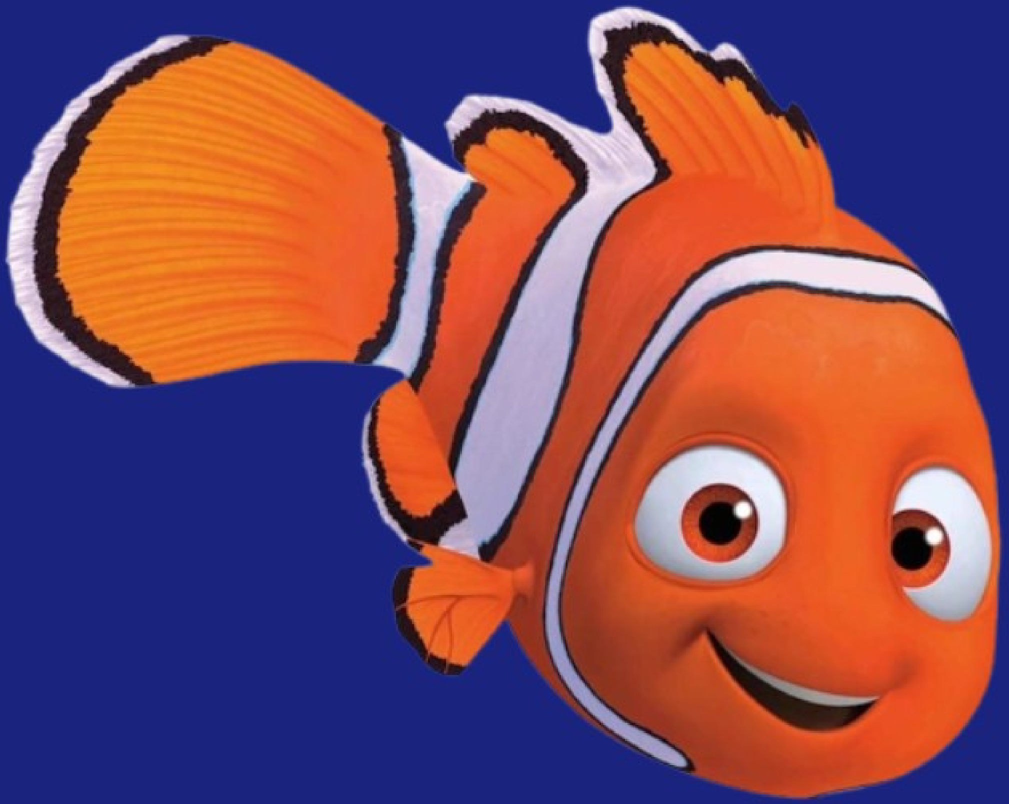

Fast, **accurate** and **compact**
occlusion-based **explanations** for
unsupervised representation learning

ICML
International Conference
On Machine Learning

Finding NEM-U

COME AND CHAT TO DISCOVER IF NEM-U WORKS FOR YOU

Fast, **accurate** and **compact**
occlusion-based **explanations** for
unsupervised representation learning

13

Channel Attention Separable Convolution Network for Skin Lesion Segmentation

Changlu Guo

Technical University of Denmark, Kgs. Lyngby, Denmark

chagu@dtu.dk

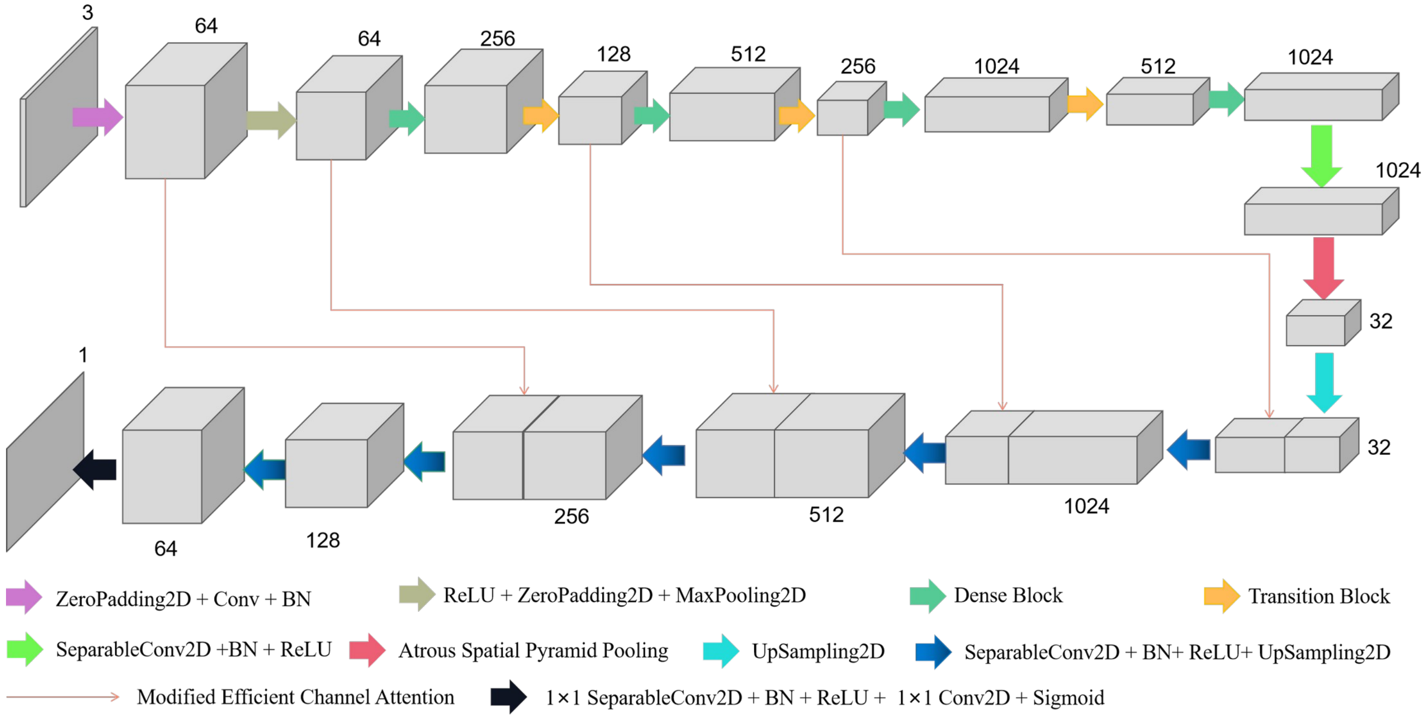

Inspired by advanced mechanisms such as U-Net, DenseNet, Separable Convolution, Channel Attention, and Atrous Spatial Pyramid Pooling (ASPP), we propose a novel network called Channel Attention Separable Convolution Network (CASCN) for skin lesions segmentation.

Channel Attention Separable Convolution Network for Skin Lesion Segmentation

Changlu Guo

Technical University of Denmark, Kgs. Lyngby, Denmark

chagu@dtu.dk

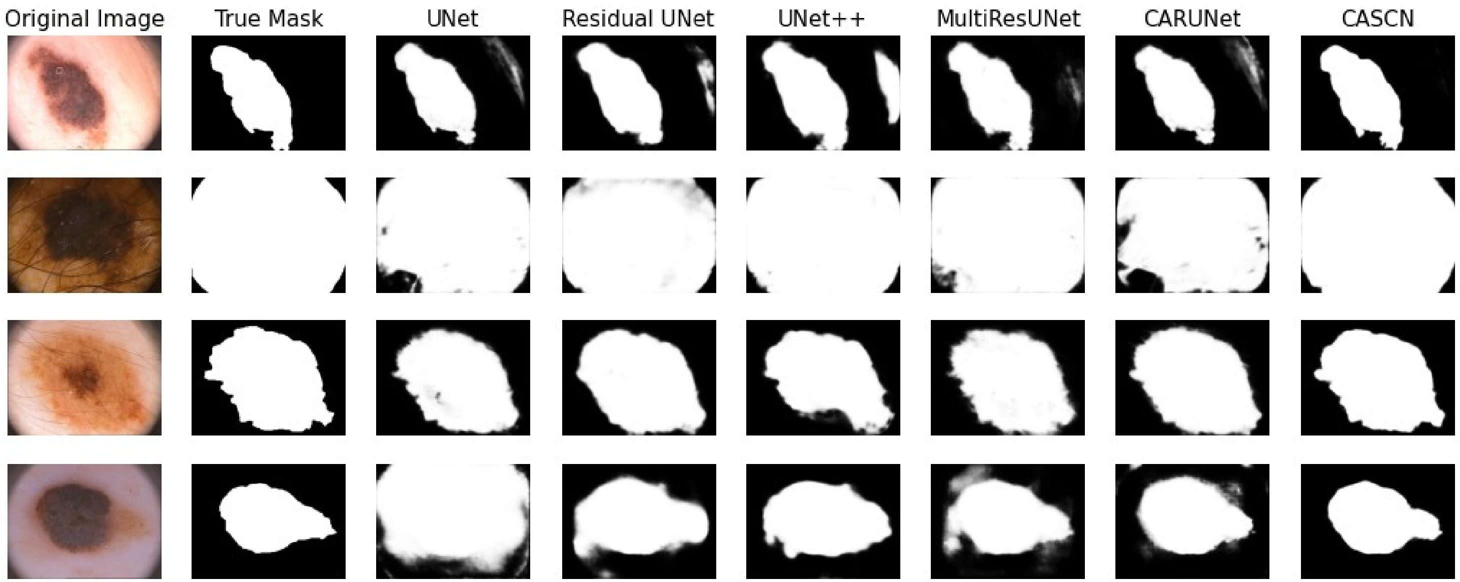

Table 1. Experiments on PH2 Dataset

Models	$SE(\%)$	$SP(\%)$	$AC(\%)$	$DI(\%)$	$JA(\%)$
U-Net [3]	94.67	93.61	94.89	92.56	86.84
Residual U-Net [4]	94.14	94.50	95.07	92.53	86.84
U-Net++ [5]	94.84	94.01	95.35	92.81	87.11
MultiResUNet [6]	94.88	94.87	95.92	93.56	88.48
CAR-UNet [7]	93.77	94.95	95.24	92.85	87.30
CASCN (ours)	**95.79**	**96.21**	**96.45**	**94.61**	**90.18**

ICONIP 2023

November 20-23, 2023, Changsha, China

The experimental results demonstrate that CASCN **achieves state-of-the-art performance on the PH2 dataset**.

14

Deploying Deep Learning Model in Real World Clinical Setting: a case study in obstetric ultrasound

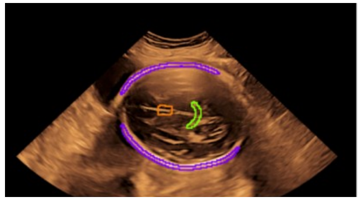

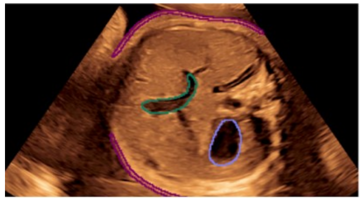

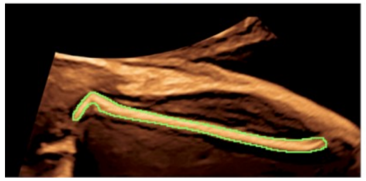

we trained a model

we want to use it in the clinic

how?

15

Source Matters: Source Dataset Impact on Model Robustness in Medical Imaging

ImageNet *or* RadImageNet?

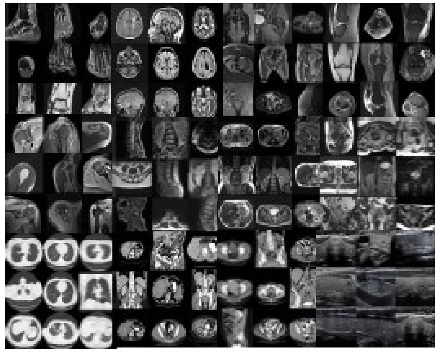

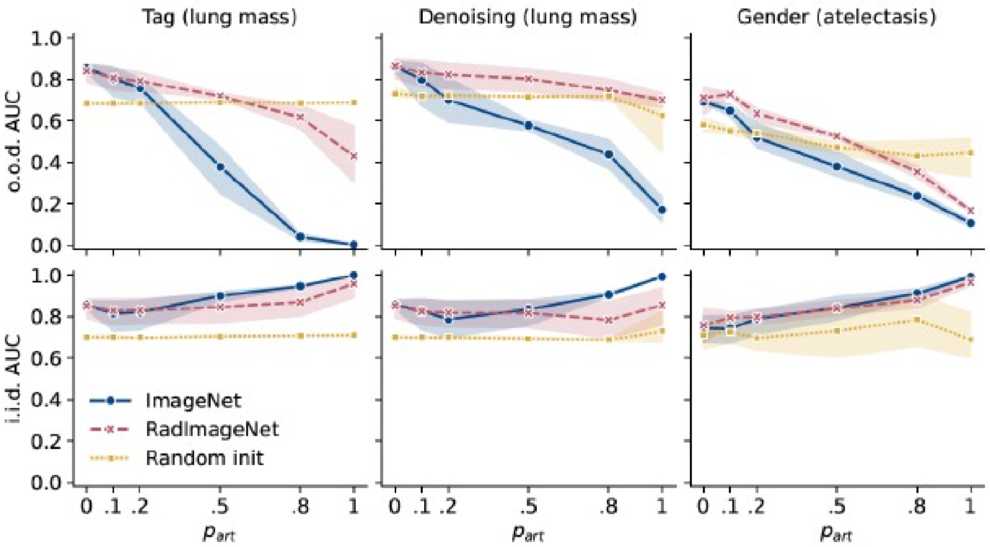

Train ($p_{art} \in [0,1]$)

No findings (M) Lung mass (F)

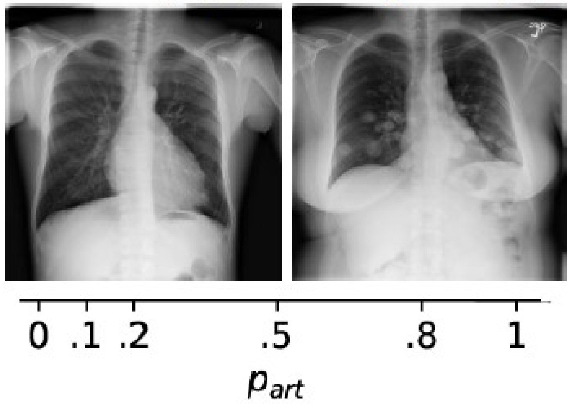

Test ($p_{art} = 1$)

No findings (F) Lung mass (M)

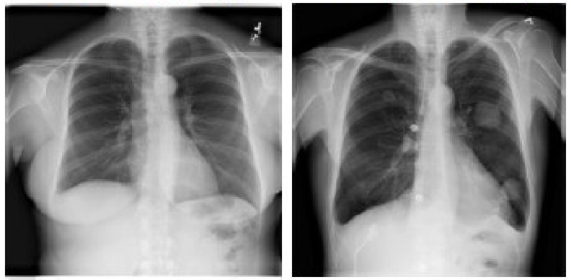

IT UNIVERSITY OF COPENHAGEN

Universidad Nacional de San Martín

UNL CONICET

sinc(i)

16

Direct Observation and Kinetic Quantification of Stochastic Protein-Protein Interactions

Emily Winther Sørensen and Nikos Hatzakis

Experimental setup

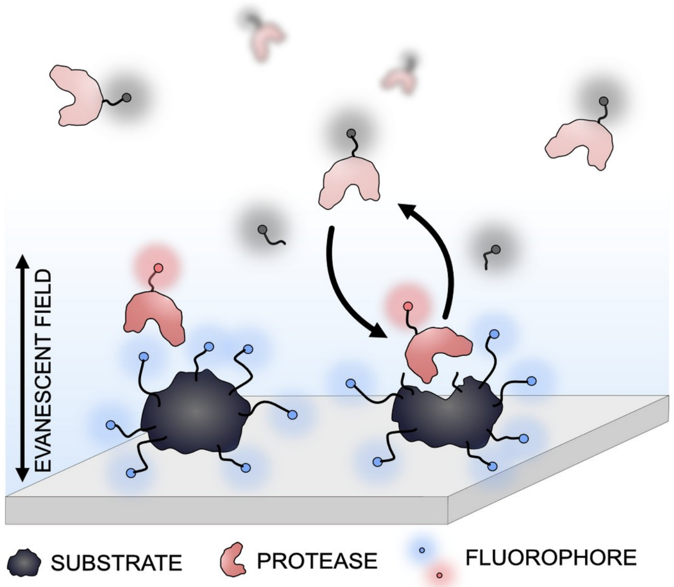

Direct observation

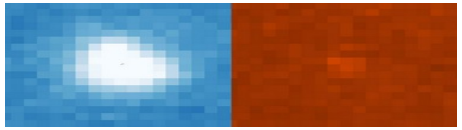

Quantification

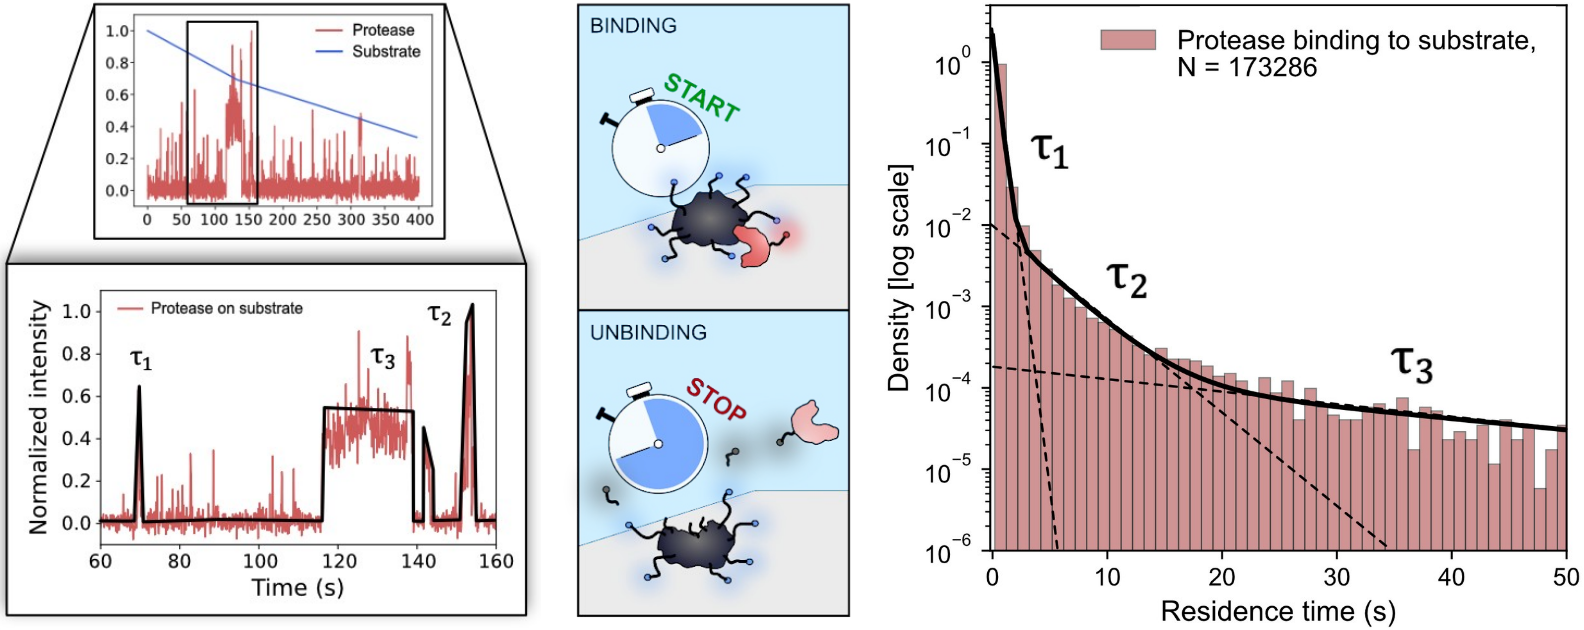

Applications

Toolbox to **directly observe and quantify** distinct protease binding processes and substrate degradation **synchronously**

Streamlined quantification of nanoparticle interactions of liposomes, biomolecular motors, drug-target interactions, enzymatic reactions etc.

17

Deep Generative Models for Characterising Atomic Structures of Nanomaterials

The Data

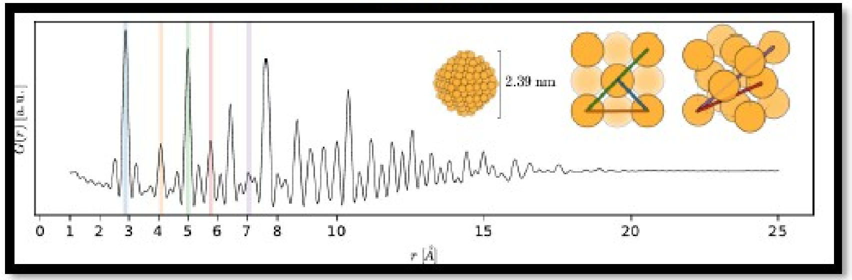

Accelerated Data Generation [3]

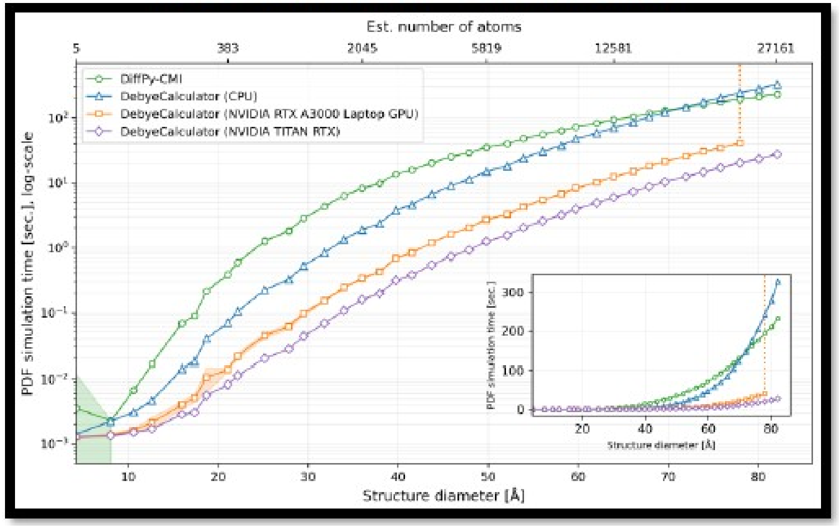

Data Curation for Chemically-Informed Graph Datasets [2]

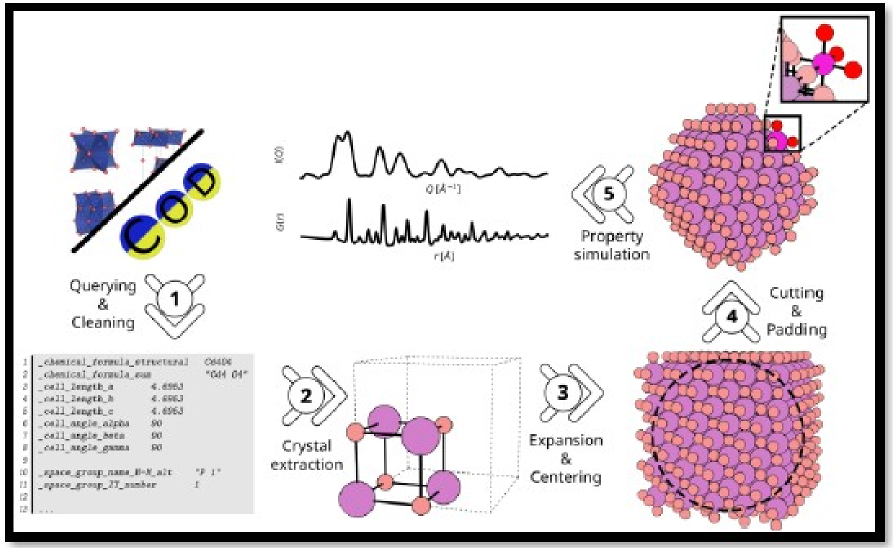

Autoregressive Models [1]

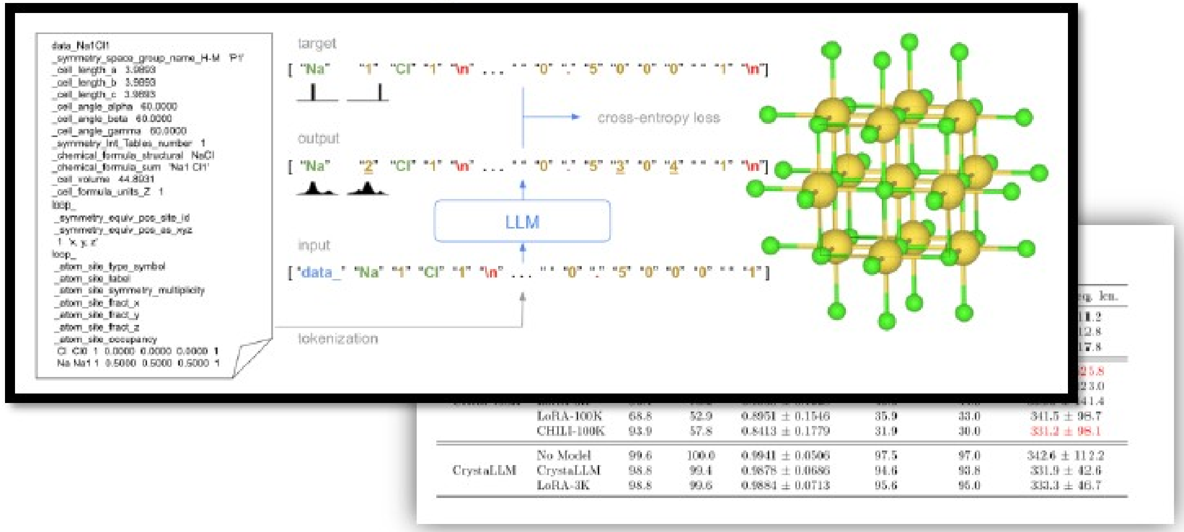

[1] L. M. Antunes, K. T. Butler, and R. Grau-Crespo. Crystal structure generation with autoregressive large language modeling, 2024.

[2] U. Friis-Jensen, F. L. Johansen, A. S. Anker, E. B. Dam, K. M. Jensen, and R. Selvan. Chili: Chemically-informed large-scale inorganic nanomaterials dataset for advancing graph machine learning, 2024.

[3] F. L. Johansen, A. S. Anker, U. Friis-Jensen, E. B. Dam, K. M. Jensen, and R. Selvan. A gpu-accelerated open-source python package for calculating powder diffraction, small-angle-, and total scattering with the debye scattering equation. *Journal of Open Source Software*, 9(94):6024, 2024.

18

Optimized KiU-Net: Lightweight Convolutional Neural Network for Retinal Vessel Segmentation in Medical Images

Hazrat Bilal
(CRT-AI, School Of Computer Science, University of Galway)

Supervised By: **Dr. Malika Bendechache**

Introduction

- Blood vessels and circulatory system
- Diseases diagnose
  - Diabetic retinopathy, Macular edema, Arteriosclerosis etc.
- Detection of Eye Diseases, main factors, and structural points
  - vessel diameter, branch angle, branch length
- Manual diagnosis is difficult, subjective, and time-consuming
- Access to medical specialists and infrastructure
- Automatic detection of retinal diseases can address the gap

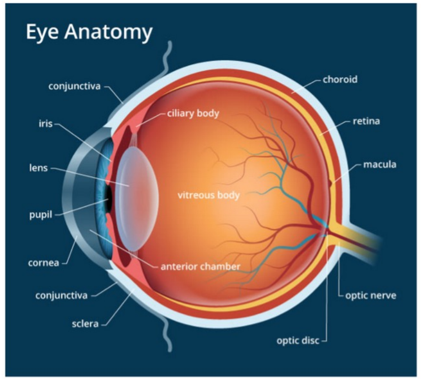

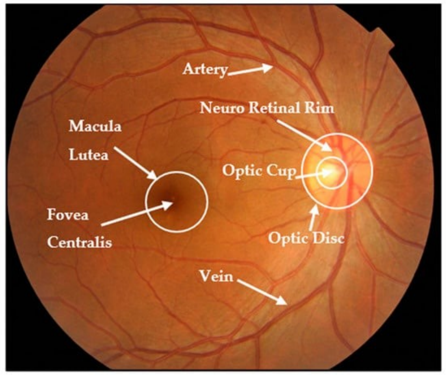

Problem Statement

- Low pixel proportion
- Low contrast
- Semantic information on small-scale vessels

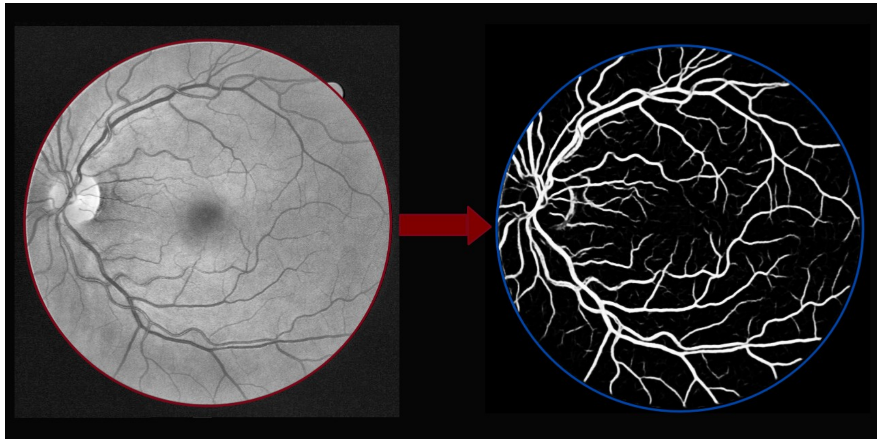

Methodology and Results

- **Optimized KiU-Net**

  Conv Channels of KiU-Net: [32,64,128]

  Conv Channels of Optimized KiU-Net: [16,32,48,64]

Table: Result of the proposed model on GlaS Dataset

Method	F1 Score	IoU
U-Net	77.78	65.34
Res-UNet	78.83	65.95
Deeplabv3+	76.01	67.04
MedT	81.02	69.61
HistoSeg	**98.07**	**76.73**
Optimized KiU-Net (ours)	82.21	71.03

Table: Result of the proposed model on RITE Dataset

Method	F1 Score	IoU	Parameters
Seg-Net	52.23	39.14	12.5M
U-Net	55.24	31.11	3.1M
KiU-Net	75.17	60.37	0.29M
Optimized KiU-Net (ours)	**79.80**	**66.30**	**0.18M**

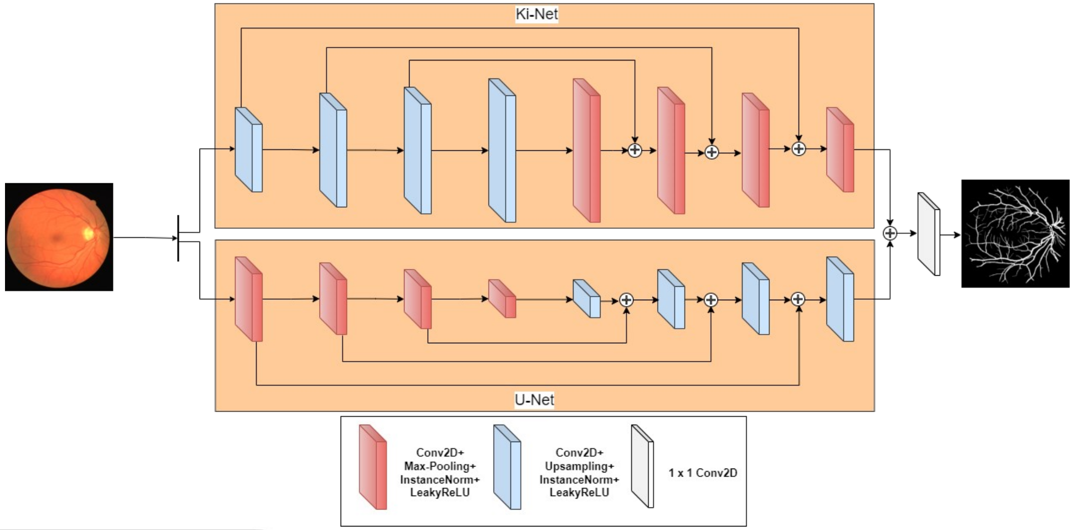

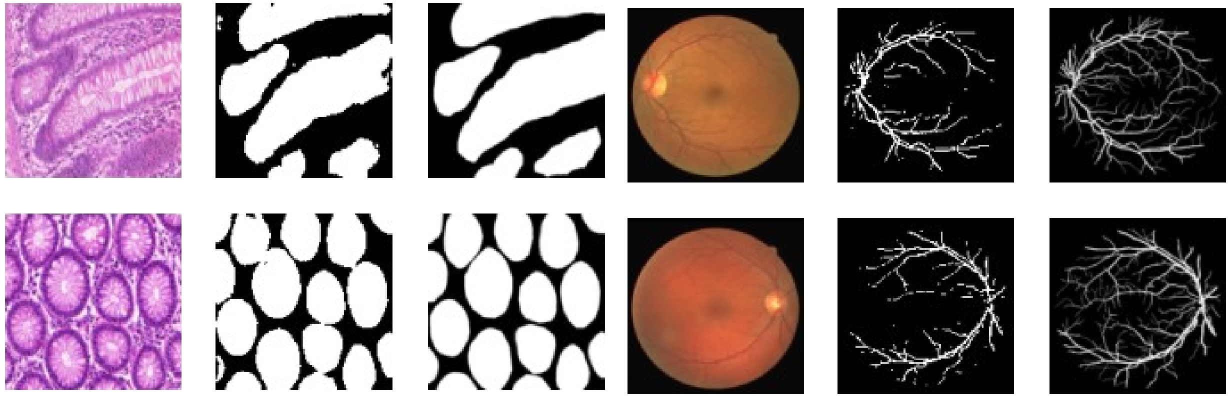

Acknowledgements

This work was conducted with the financial support of the Science Foundation Ireland Centre for Research Training in Artificial Intelligence under Grant No. 18/CRT/6223.

Thank You!

19

Predicting urban tree cover from incomplete point labels and limited background information

Hui Zhang, Ankit Kariryaa, Venkanna Babu Guthula, Christian Igel, Stefan Oehmcke

University of Copenhagen

UNIVERSITY OF COPENHAGEN

Learning from incomplete & sparse labels - mask & objectness

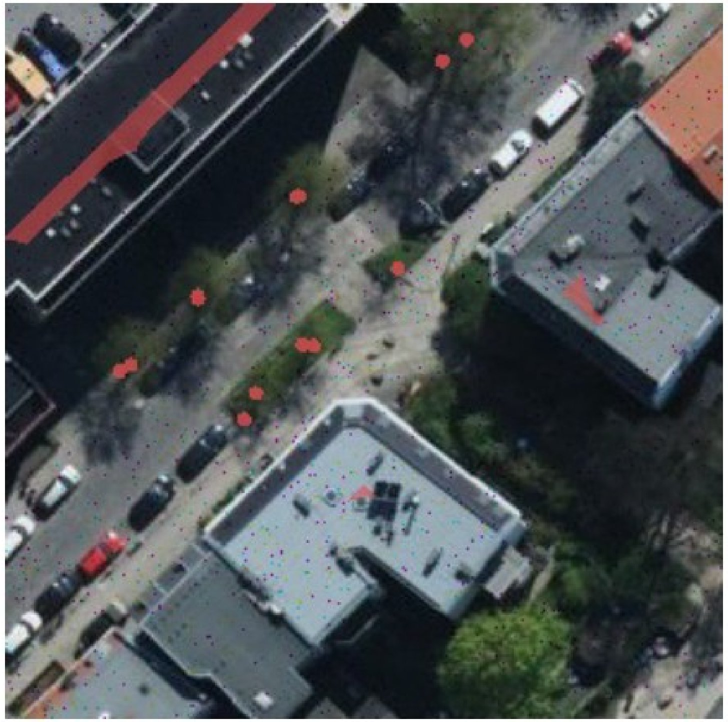

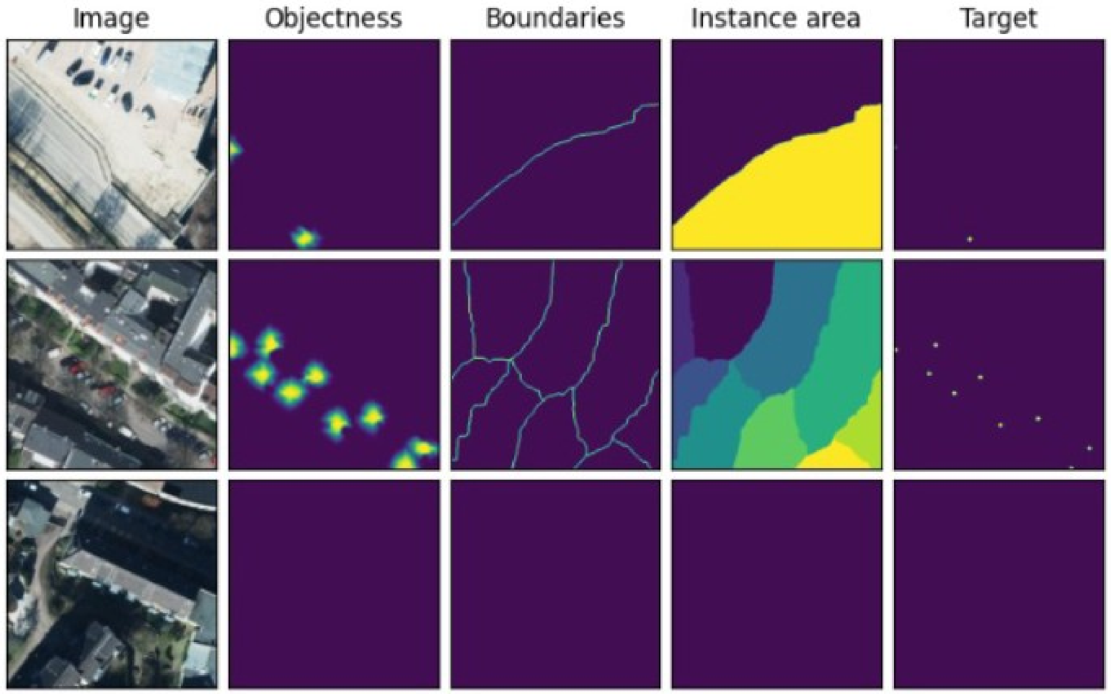

$$\mathcal{L} = \mathcal{L}^{\text{sup}}(f(x)\odot \boldsymbol{m}, \boldsymbol{y}\odot \boldsymbol{m})$$

$$\mathcal{L}^{\text{obj}}(f(x)\odot \boldsymbol{r}, \boldsymbol{o}\odot \boldsymbol{r}, \hat{\boldsymbol{r}}\odot \boldsymbol{r}) = -\frac{1}{|o|}\sum_{i=1}^{|o|} \mathrm{BCE}(f(x)\odot \hat{\boldsymbol{r}}_i, \boldsymbol{o}\odot \hat{\boldsymbol{r}}_i)$$

15. Shijie Li, Neel Dey, Katharina Bermond, Leon von der Emde, Christine A. Curcio, Thomas Ach, and Guido Gerig. 2021. Point-Supervised Segmentation Of Microscopy Images And Volumes Via Objectness Regularization. In International Symposium on Biomedical Imaging (ISBI). IEEE, 1558–1562. https://doi.org/10.1109/ISBI48211.2021.9433963 2, 4, 5, 6, 7

Results

Name	Description	Setting	Delineated			Sparse	
			$\mathrm{IoU_d}$	$\mathrm{F1_d}$	$\mathrm{BA_d}$	$\mathrm{Recall_s}$	$\mathrm{BA_s}$
baseline	Neither masking nor objectness and enlarging positive labels to a circle with 1.5 m radius.	$\beta = 0$ $\boldsymbol{m} = \mathbf{1}^{w \times h}$ $\boldsymbol{y} = \boldsymbol{m}^{\mathrm{disk}}$	–	–	–	0.0005	0.5002
Obj	Reimplementation of [15].	$\beta = 1$ $\boldsymbol{m} = \boldsymbol{y} \cup \boldsymbol{b}$ $\hat{\boldsymbol{r}} = \mathbf{1}^{w \times h}$	0.1191	0.5479	0.5575	0.5908	0.7844
Mask **(ours)**	Only supervised loss and masking out unknown areas.	$\beta = 0$ $\boldsymbol{m} = \boldsymbol{y} \cup (\boldsymbol{m}^{\mathrm{OSM}} \setminus \boldsymbol{m}^{\mathrm{disk}})$	**0.4839**	0.7551	**0.8119**	**0.8994**	0.8205
MaskObj **(ours)**	As [15] but restricting objectness loss to 1.5 m radius around points.	$\beta = 1$ $\hat{\boldsymbol{r}} = \boldsymbol{m}^{\mathrm{disk}}$ $\boldsymbol{m} = \boldsymbol{y} \cup (\boldsymbol{m}^{\mathrm{OSM}} \setminus \boldsymbol{m}^{\mathrm{disk}})$	0.3364	0.7289	0.6978	0.7771	**0.8399**
MaskObjThresh **(ours)**	As MaskObj but removing objectness smaller than the specified threshold ($t = 0.2$).	$\beta = 1$ $\hat{\boldsymbol{r}} = \boldsymbol{m}^{\mathrm{disk}} \cap (\boldsymbol{o} \geq t)$ $\boldsymbol{m} = \boldsymbol{y} \cup (\boldsymbol{m}^{\mathrm{OSM}} \setminus \boldsymbol{m}^{\mathrm{disk}})$	0.4805	**0.7660**	0.7870	0.8345	0.8135

20

Unsupervised Detection of Fetal Brain Anomalies using Denoising Diffusion Models

Markus Ditlev Sjøgren Olsen[1], Jakob Ambsdorf[2,4], Manxi Lin[1], Caroline Taksøe-Vester[3], Morten Bo Søndergaard Svendsen[1], Anders Nymark Christensen[1], Mads Nielsen[2,4], Martin Grønnebæk Tolsgaard[3], Aasa Feragen[1,4], Paraskevas Pegios[1,4]
[1]Technical University of Denmark, [2]University of Copenhagen, [3]CAMES Rigshospitalet, [4]Pioneer Centre for AI

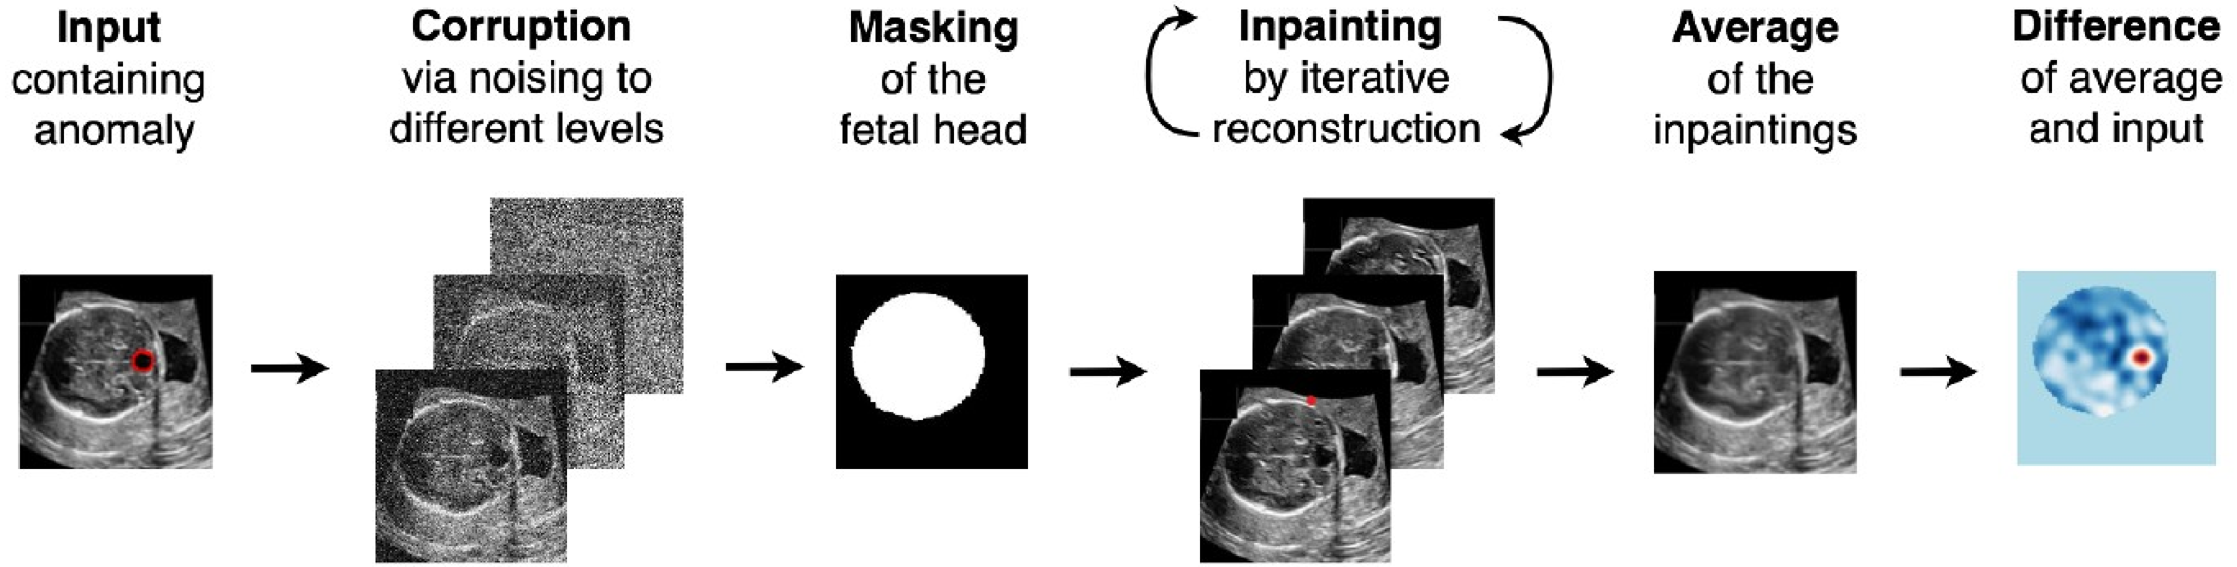

21

Universal Image Segmentation with Diffusion Models

- Segment any image, either completely agnostically, or conditioned on
  - Number of classes, Label names
  - Bounding boxes
  - Foreground/background points
  - Instance/semantic
  - Related images (e.g. few shot)

- Talk with me about
  - Diffusion, Generative models
  - Anything really
  - Programming

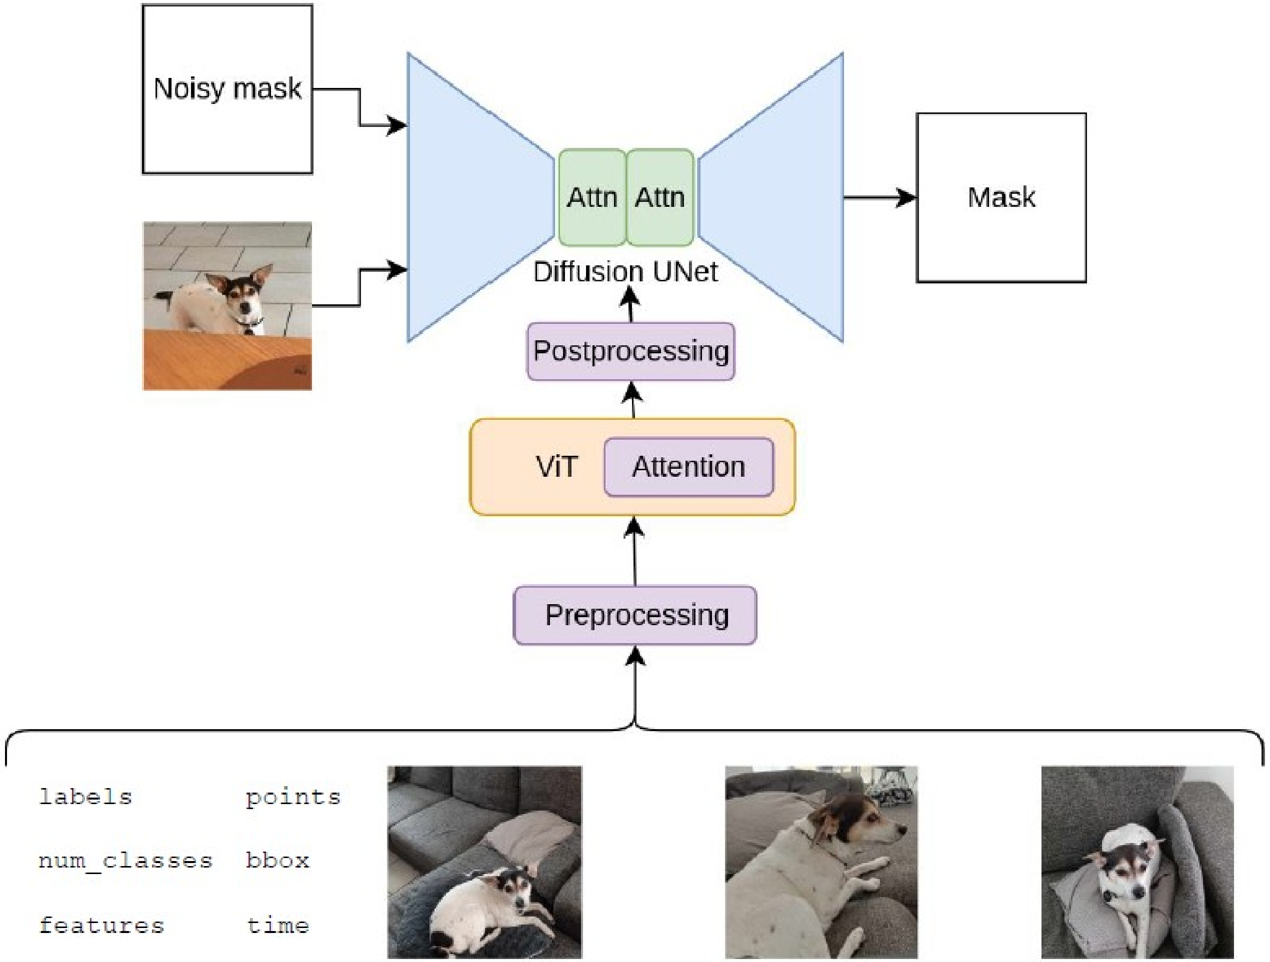

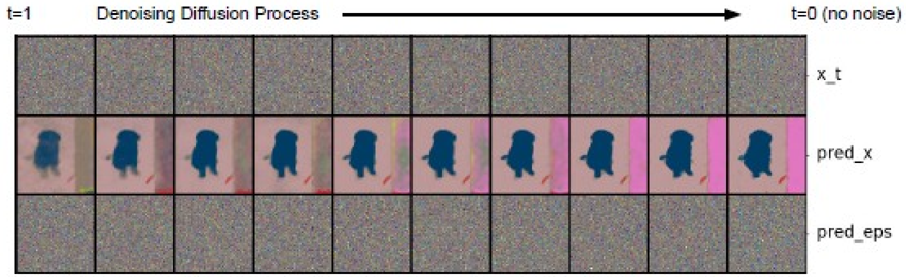

22

Deep Learning Based Localization and Characterization of White Matter Lesions

Julia Machnio[1], Mads Nielsen[1], Mostafa Mehdipour Ghazi[1]
[1] Pioneer Centre for AI, Department of Computer Science, University of Copenhagen, Denmark

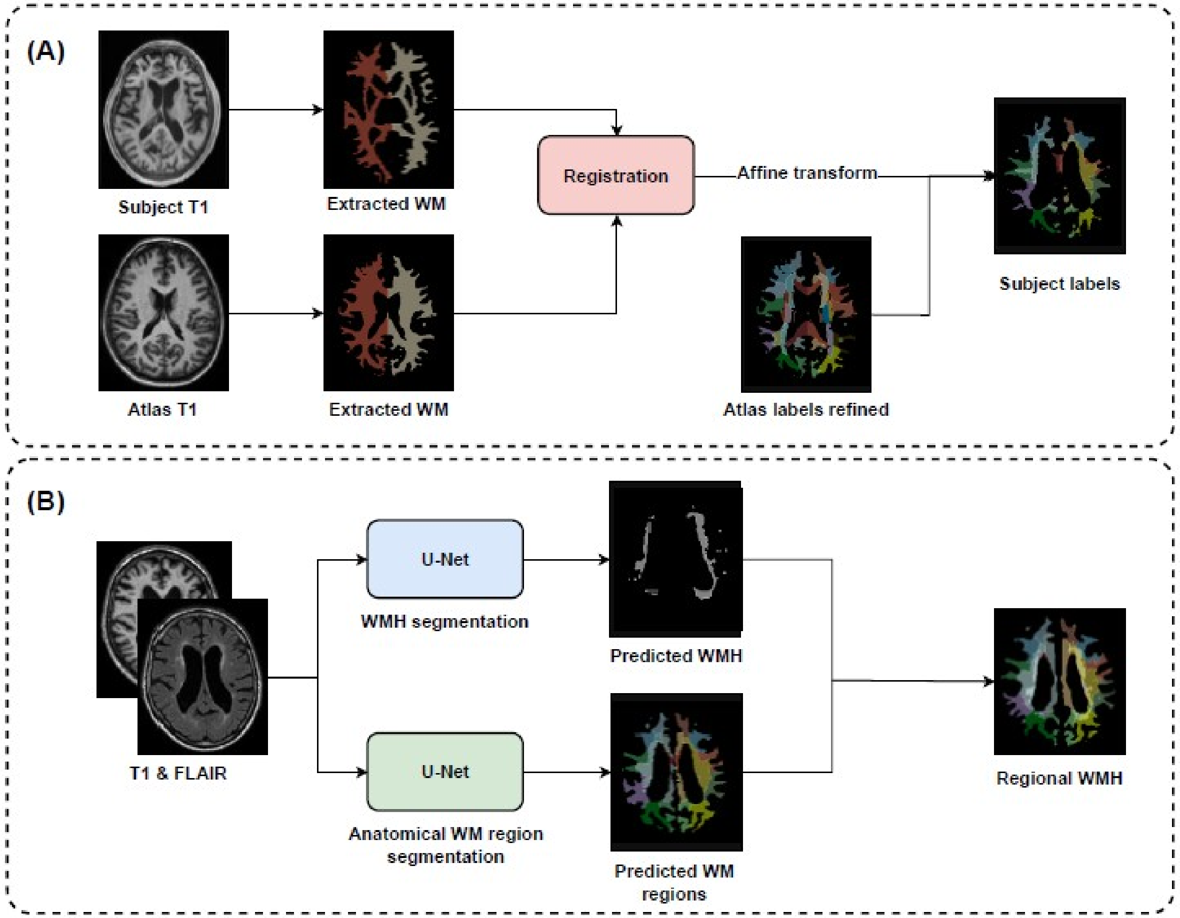

23

AI-Driven Outlier Detection of Human Vertebrae from CT Scans

Detecting vertebral fractures using AI

- Testing different neural networks with autoencoder architecture.
- CT scans of Lumbar 1 vertebra.
- Optimal convolutional U-Net model.
- Trained on healthy vertebra images.
- Investigate performance with unseen data.

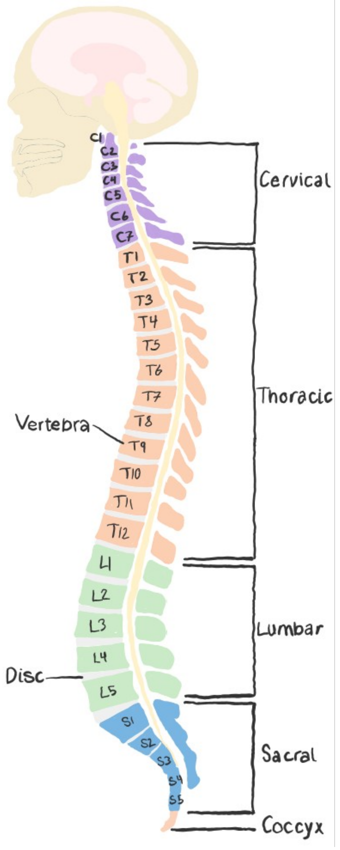

Figure 1: Illustration of the spine and the vertebrae.

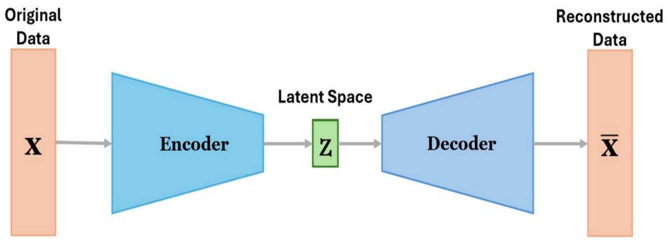

Figure 2: An illustration of the components in the autoencoder architecture.

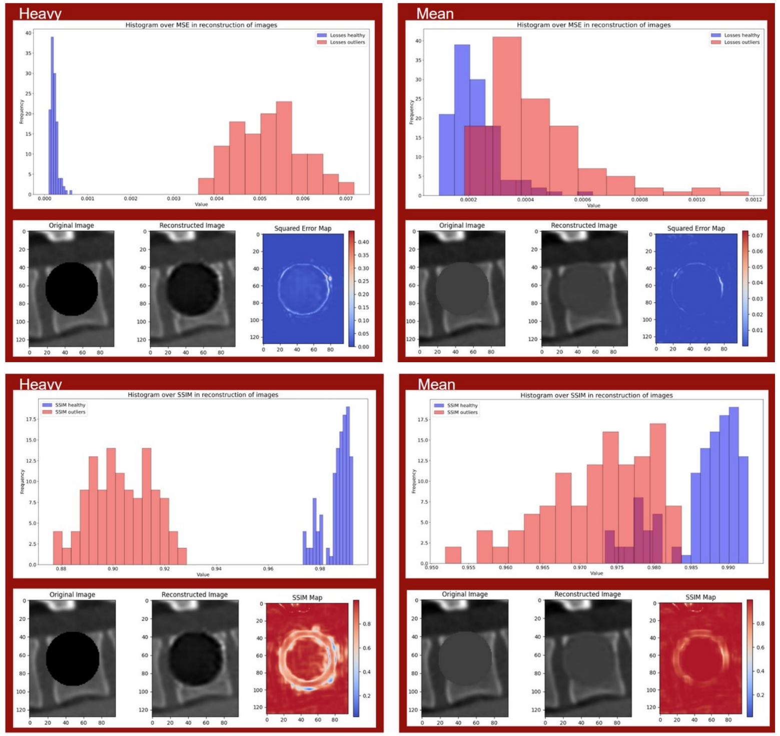

Figure 3: Overview of some of the evaluations performed in the project.

Julie Boel og Katja Norsker

24

Diffusion-weighted MRS

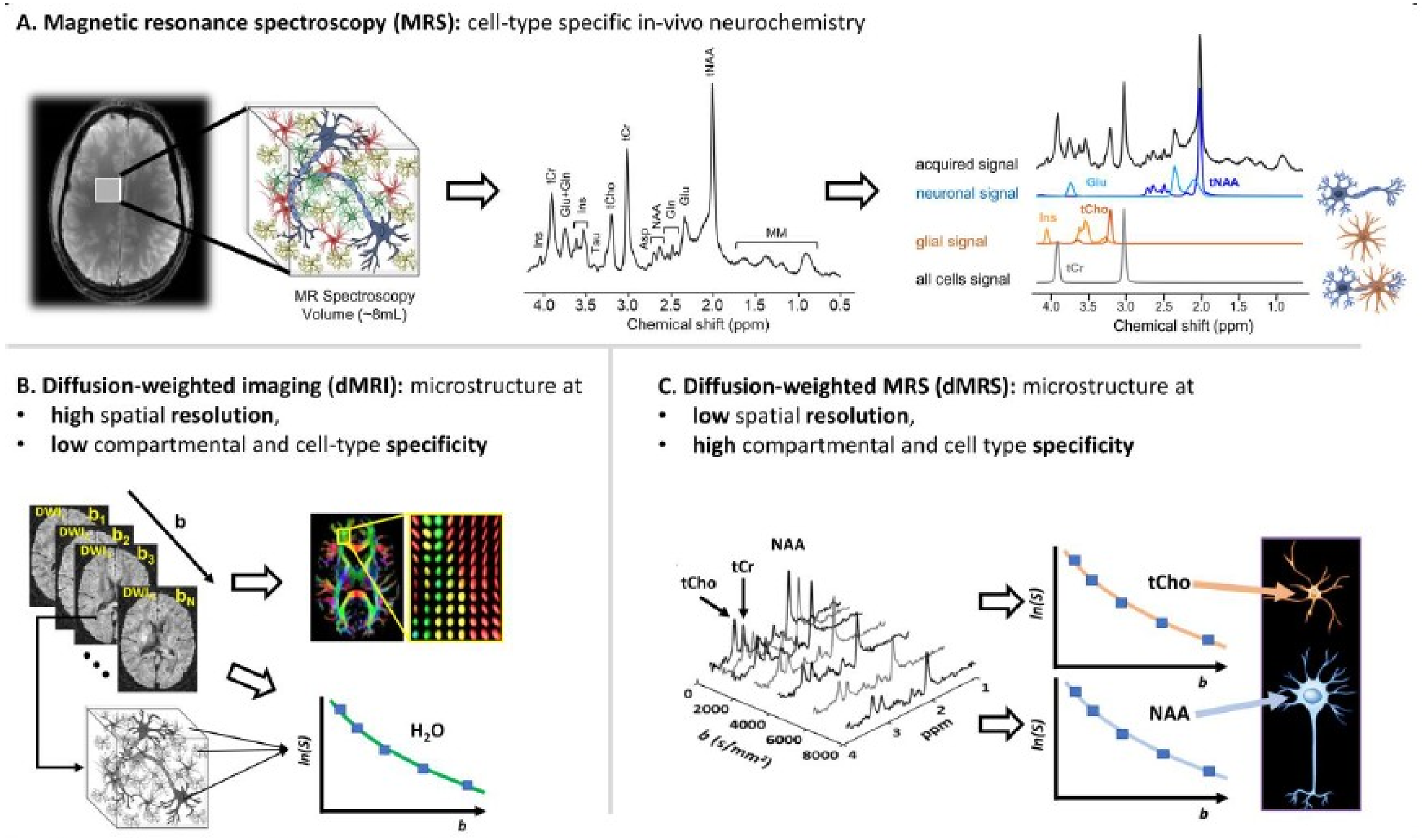

Ligneul, Clémence, et al. "Diffusion-weighted MR spectroscopy: Consensus, recommendations, and resources from acquisition to modeling." *Magnetic resonance in medicine* 91.3 (2024): 860-885.

Eddy Currents

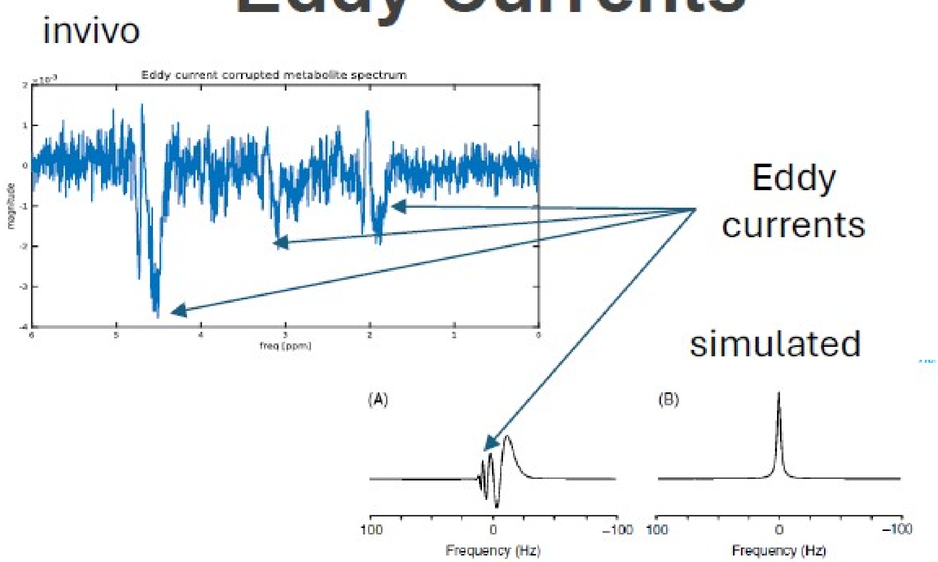

Sensitivity to motion

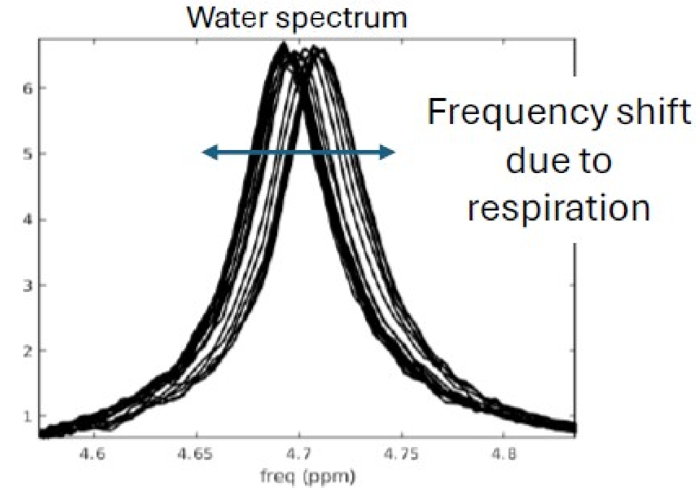

25

Perfusion Defect in Cardiovascular

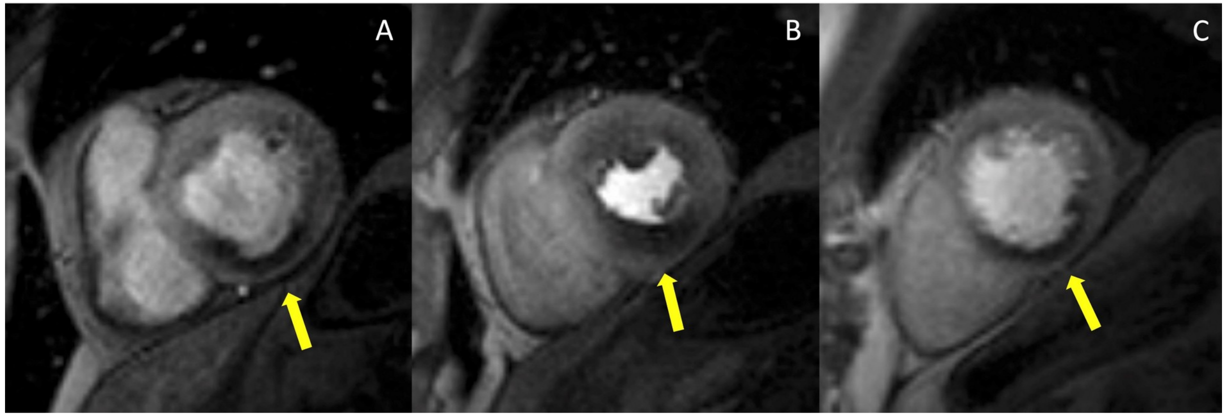

Perfusion imaging- short axis CMRI

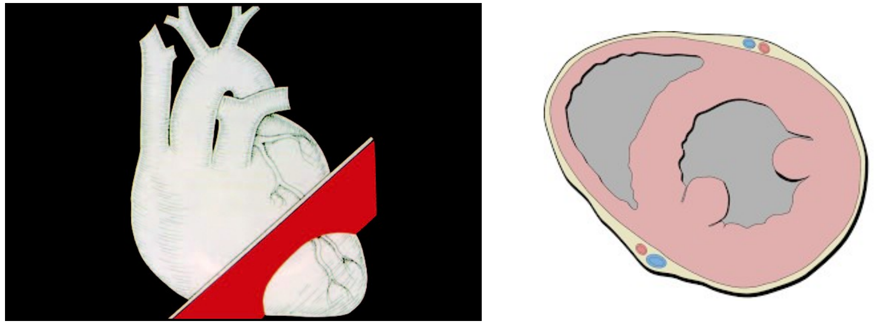

Short axis images

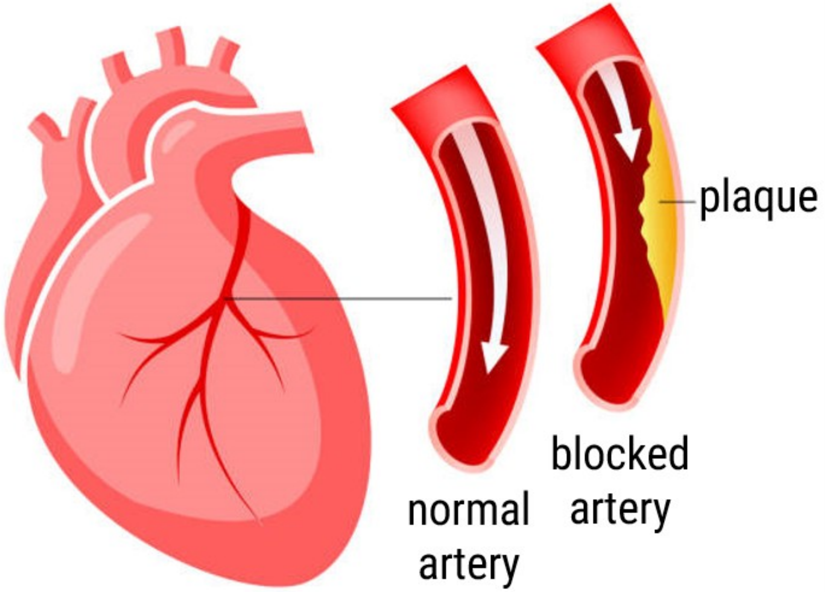

Perfusion defect

UNIVERSITY OF LEICESTER

NIHR | Leicester Biomedical Research Centre

Problem Definition

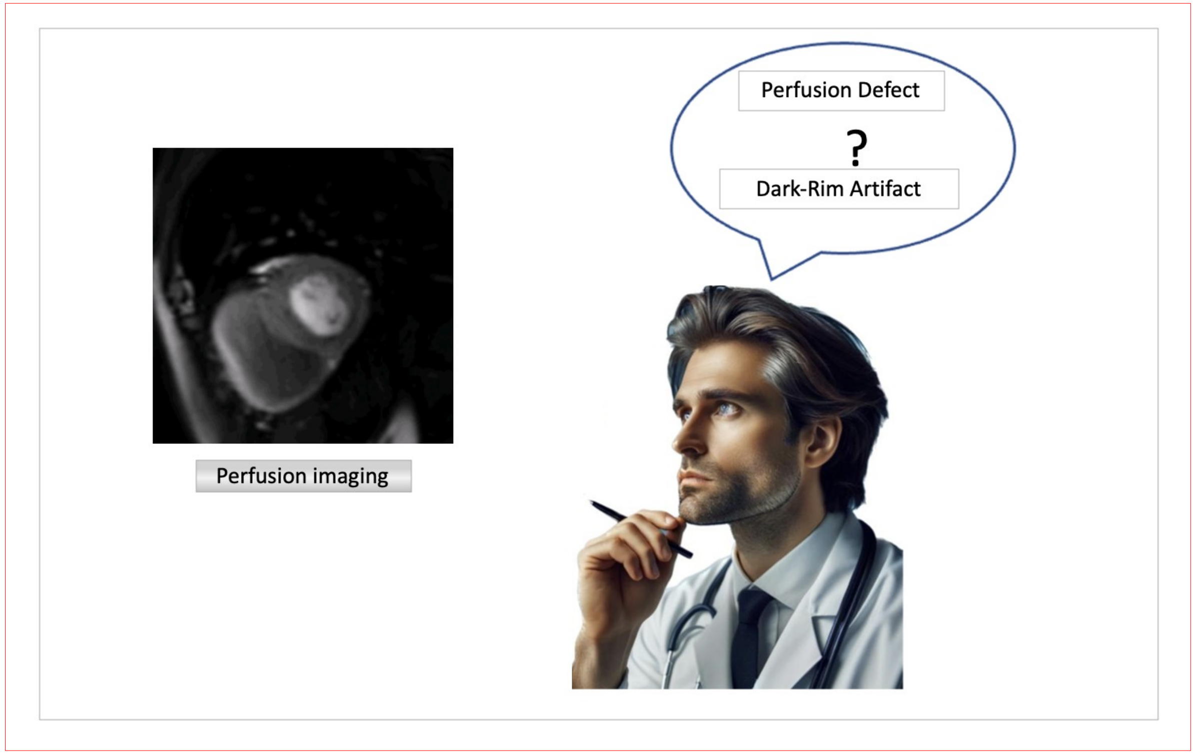

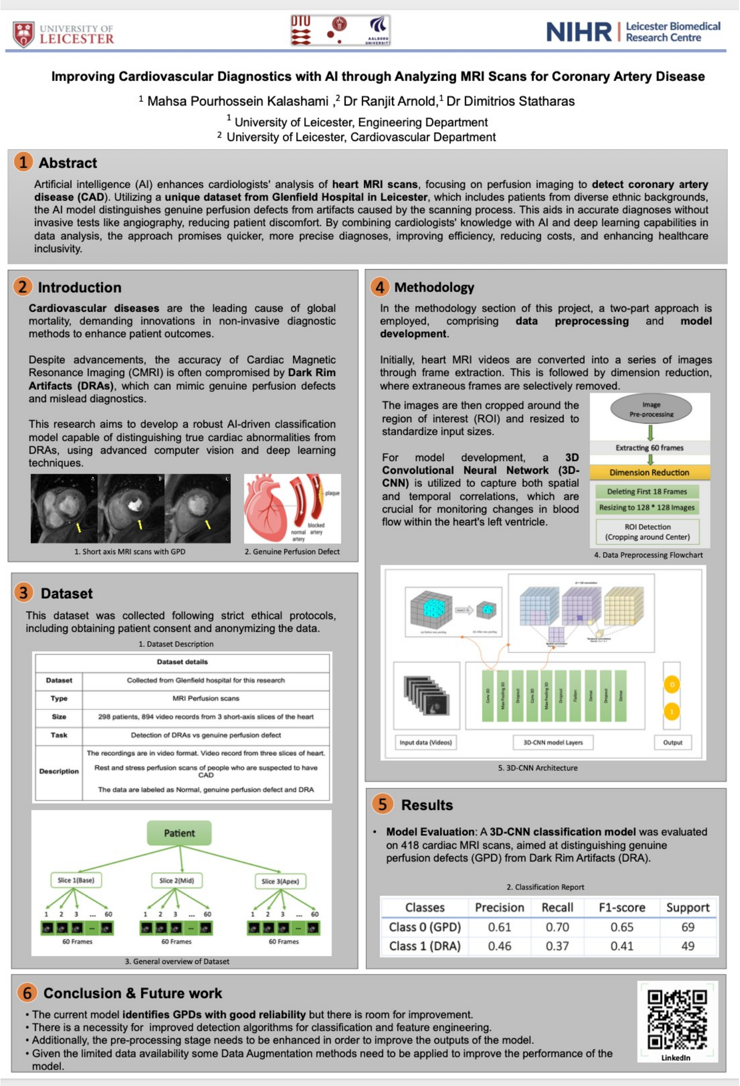
UNIVERSITY OF LEICESTER

NIHR | Leicester Biomedical Research Centre

Improving Cardiovascular Diagnostics with AI through Analyzing MRI Scans for Coronary Artery Disease

[1] Mahsa Pourhossein Kalashami ,[2] Dr Ranjit Arnold,[1] Dr Dimitrios Statharas

[1] University of Leicester, Engineering Department
[2] University of Leicester, Cardiovascular Department

1 Abstract

Artificial intelligence (AI) enhances cardiologists' analysis of **heart MRI scans**, focusing on perfusion imaging to **detect coronary artery disease (CAD)**. Utilizing a **unique dataset from Glenfield Hospital in Leicester**, which includes patients from diverse ethnic backgrounds, the AI model distinguishes genuine perfusion defects from artifacts caused by the scanning process. This aids in accurate diagnoses without invasive tests like angiography, reducing patient discomfort. By combining cardiologists' knowledge with AI and deep learning capabilities in data analysis, the approach promises quicker, more precise diagnoses, improving efficiency, reducing costs, and enhancing healthcare inclusivity.

2 Introduction

Cardiovascular diseases are the leading cause of global mortality, demanding innovations in non-invasive diagnostic methods to enhance patient outcomes.

Despite advancements, the accuracy of Cardiac Magnetic Resonance Imaging (CMRI) is often compromised by **Dark Rim Artifacts (DRAs)**, which can mimic genuine perfusion defects and mislead diagnostics.

This research aims to develop a robust AI-driven classification model capable of distinguishing true cardiac abnormalities from DRAs, using advanced computer vision and deep learning techniques.

1. Short axis MRI scans with GPD

2. Genuine Perfusion Defect

3 Dataset

This dataset was collected following strict ethical protocols, including obtaining patient consent and anonymizing the data.

1. Dataset Description

Dataset details	
Dataset	Collected from Glenfield hospital for this research
Type	MRI Perfusion scans
Size	298 patients, 894 video records from 3 short-axis slices of the heart
Task	Detection of DRAs vs genuine perfusion defect
Description	The recordings are in video format. Video record from three slices of heart. Rest and stress perfusion scans of people who are suspected to have CAD The data are labeled as Normal, genuine perfusion defect and DRA

3. General overview of Dataset

4 Methodology

In the methodology section of this project, a two-part approach is employed, comprising **data preprocessing** and **model development**.

Initially, heart MRI videos are converted into a series of images through frame extraction. This is followed by dimension reduction, where extraneous frames are selectively removed.

The images are then cropped around the region of interest (ROI) and resized to standardize input sizes.

For model development, a **3D Convolutional Neural Network (3D-CNN)** is utilized to capture both spatial and temporal correlations, which are crucial for monitoring changes in blood flow within the heart's left ventricle.

4. Data Preprocessing Flowchart

5. 3D-CNN Architecture

5 Results

- **Model Evaluation**: A **3D-CNN classification model** was evaluated on 418 cardiac MRI scans, aimed at distinguishing genuine perfusion defects (GPD) from Dark Rim Artifacts (DRA).

2. Classification Report

Classes	Precision	Recall	F1-score	Support
Class 0 (GPD)	0.61	0.70	0.65	69
Class 1 (DRA)	0.46	0.37	0.41	49

6 Conclusion & Future work

- The current model **identifies GPDs with good reliability** but there is room for improvement.
- There is a necessity for improved detection algorithms for classification and feature engineering.
- Additionally, the pre-processing stage needs to be enhanced in order to improve the outputs of the model.
- Given the limited data availability some Data Augmentation methods need to be applied to improve the performance of the model.

LinkedIn

26

Mitochondrial Motility and Subcellular Localisation in Pancreatic Alpha Cells

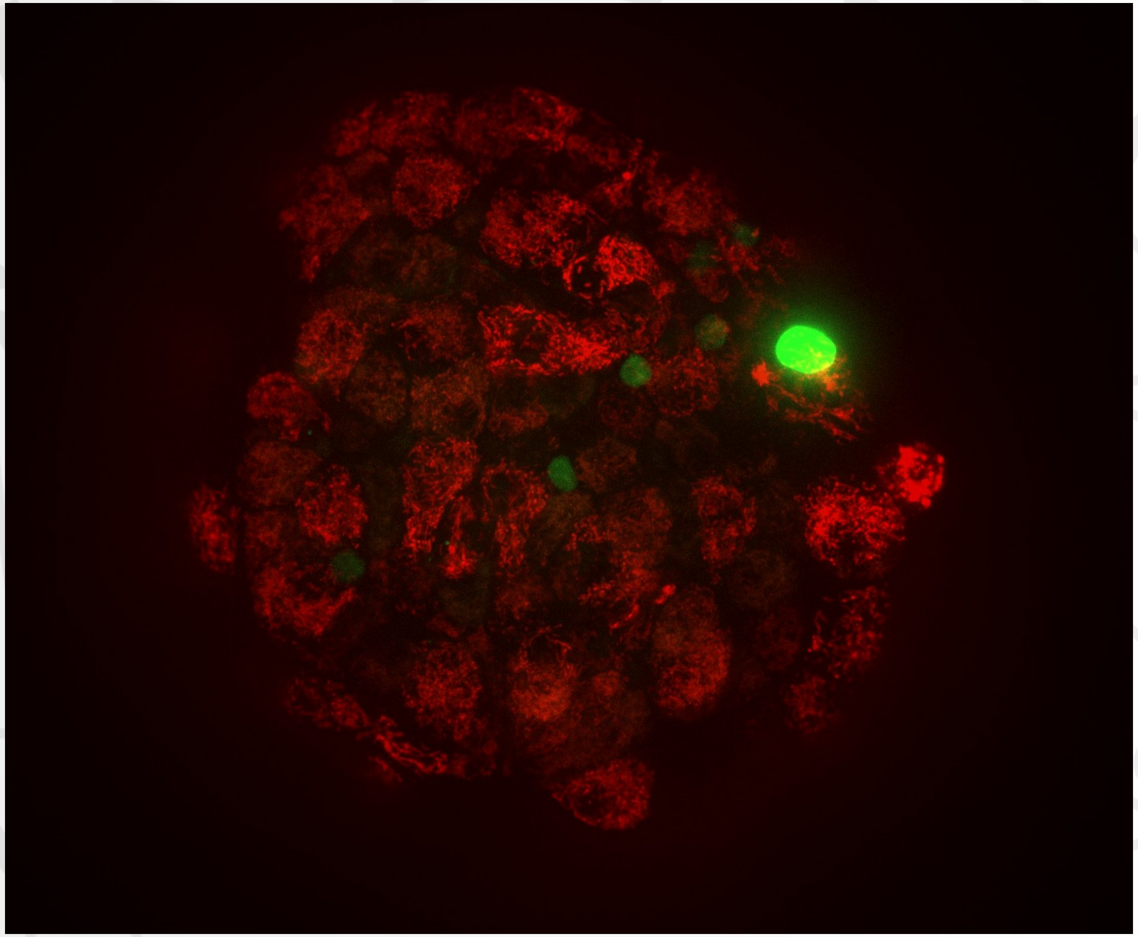

Poster Pitch

Maia H. Ekstrand
maia.ekstrand@bio.ku.dk
Ph.D. fellow in Knudsen Lab
Section for Cell Biology and Physiology

Mitochondrial Motility and Subcellular Localisation in Pancreatic Alpha Cells

Aim of Study:

How are mitochondrial motility and localisation regulated in alpha cells?
How do mitochondrial motility and localisation impact glucagon secretion?

- Glucagon helps to increase blood glucose when needed
- Type 2 Diabetic patients suffer from dysregulated glucagon secretion

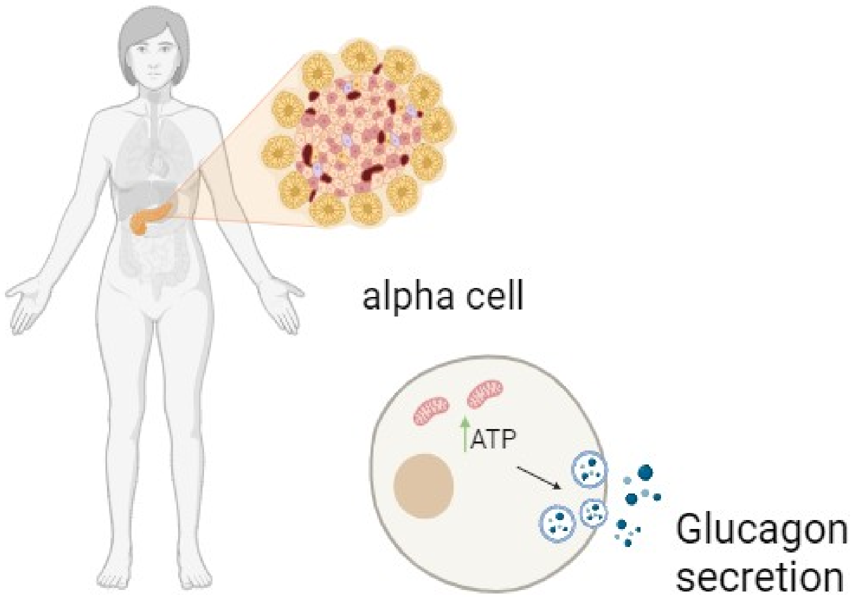

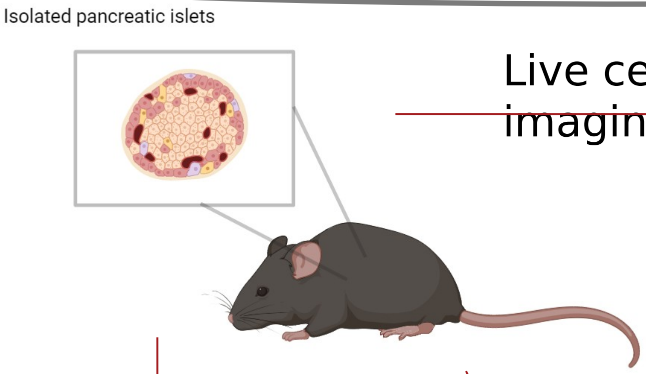

Live cell confocal imaging

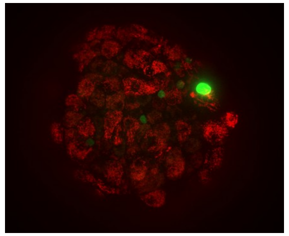

Glucose stimulations and TEM

Glucagon secretion measurements

Analysis of mitochondrial motility

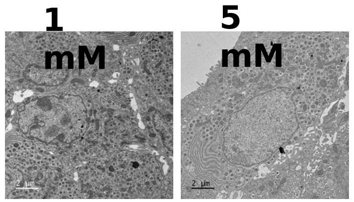

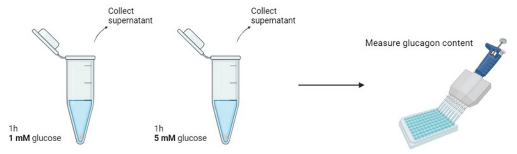

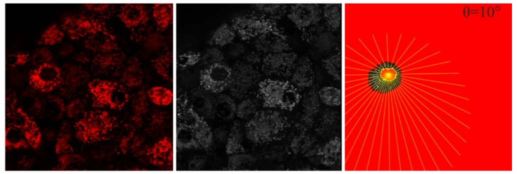

Graphical abstract

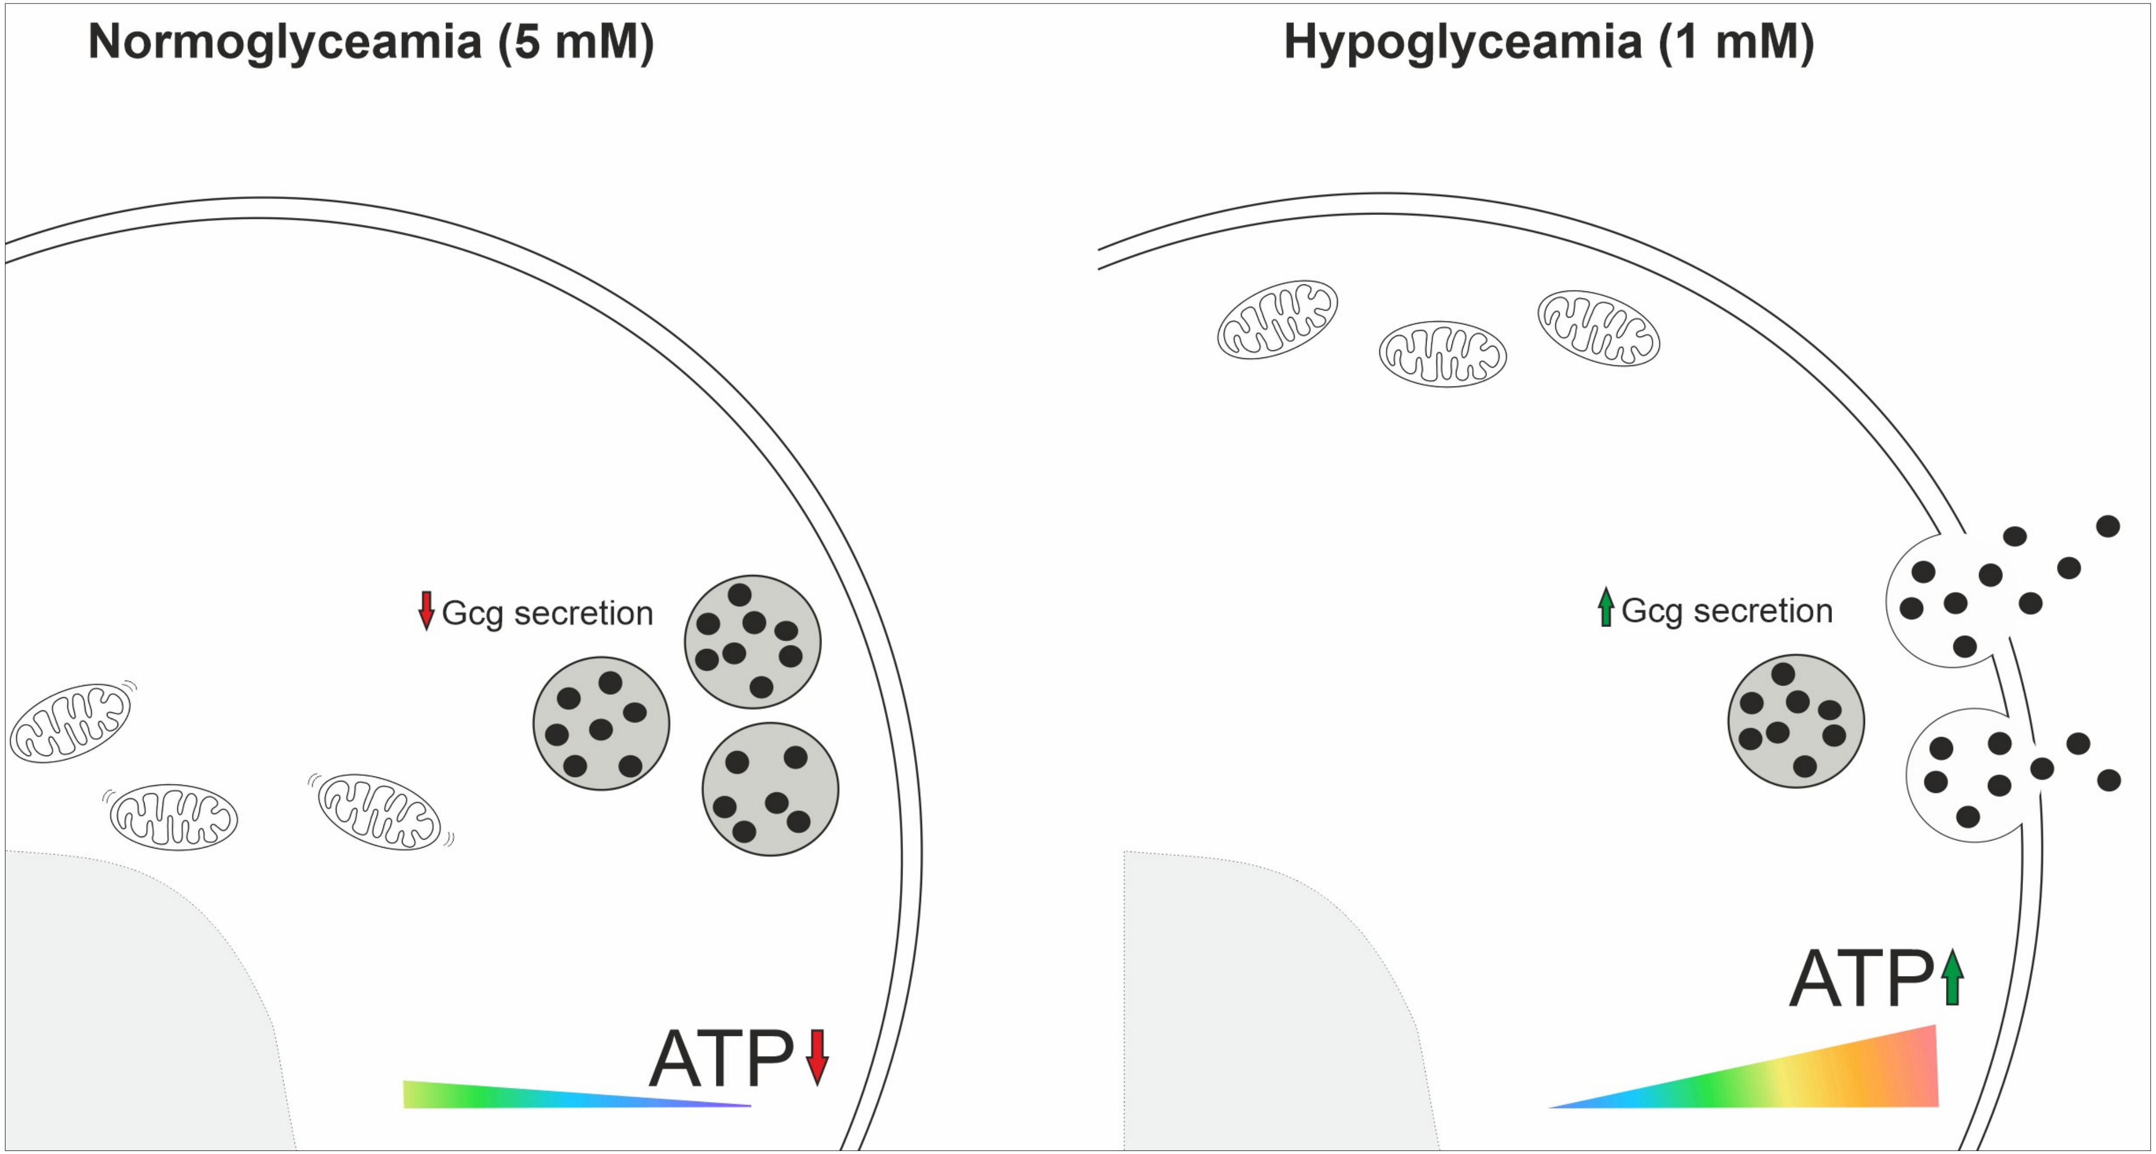

27

Incorporating Clinical Guidelines through Adapting Multi-modal Large Language Model for Prostate Cancer PI-RADS Scoring

Tiantian Zhang*, Manxi Lin*, Hongda Guo, Xiaofan Zhang, Ka Fung Peter Chiu, Aasa Feragen, Qi Dou

Motivation

- **Improving Diagnostic Accuracy**: Current deep learning methods for PI-RADS scoring often neglect the PI-RADS Clinical Guideline (PICG), which radiologists use, potentially compromising the accuracy of prostate cancer diagnosis through MRI.
- **Incorporating Clinical Expertise**: There is a need to seamlessly integrate clinical guidelines like PICG into AI models to ensure that the nuanced information used by radiologists is effectively captured, thus enhancing model trustworthiness and reducing subjective biases.
- **Minimizing Model Modifications**: Developing a method that incorporates PICG into existing scoring networks without requiring extensive architectural changes or additional annotations, to streamline the integration process and improve usability in clinical settings.

Experimental Results

- Performance of different methods on the private test set. Results are reported with the average and standard deviation over three independent runs.

Model	Accuracy % ↑	MSE ↓	MAE ↓	Precision%↑	Recall%↑	F1%↑
VGG [23]	31.6±1.4	1.38±0.2	0.92±0.5	17.4±1.5	22.6±0.8	17.4±1.8
VGG [23]+PICG (ours)	38.6±2.1(+7.0)	1.09±0.1	0.77±0.0	21.1±0.3	22.9±1.2	21.0±0.6
Kang et al., [11]	30.0±4.0	1.43±0.3	0.93±0.1	13.2±2.6	20.1±2.5	14.5±2.0
Kang et al., [11]+PICG (ours)	36.4±1.0 (+6.4)	1.25±0.0	0.83±0.0	16.1±3.6	20.9±1.8	15.9±2.3
Sanford et al., [21]	30.4±0.4	1.61±0.2	0.97±0.1	17.7±1.3	22.1±1.7	16.1±1.1
Sanford et al., [21]+PICG (ours)	35.7±0.9(+5.3)	1.38±0.2	0.87±0.1	18.4±1.2	22.0±0.5	17.4±1.5
Yu et al., [27]	33.8±0.6	1.22±0.2	0.83±0.1	18.2±6.1	21.8±2.3	13.4±3.1
Yu et al., [27]+PICG (ours)	38.6±0.4(+4.8)	1.17±0.1	0.79±0.0	20.4±0.6	23.8±1.4	20.6±0.8

- Model performance with different α and the effect of two-stage fine-tuning on model performance.

Loss weight $\alpha = 0.2$			Loss weight $\alpha = 0.4$			Loss weight $\alpha = 0.6$		
Acc.% ↑	MSE↓	MAE↓	Acc.% ↑	MSE↓	MAE↓	Acc.% ↑	MSE↓	MAE↓
36.2	1.24	0.83	38.6	1.26	0.82	32.4	1.70	0.99

Model	w/o PICG	with PICG	Baseline MLLM
Kang et al. [11]	33.8%	37.5%	33.1%
Sanford et al. [21]	30.7%	36.5%	29.7%
Yu et al. [27]	34.5%	38.9%	35.8%

Methodology-Overview

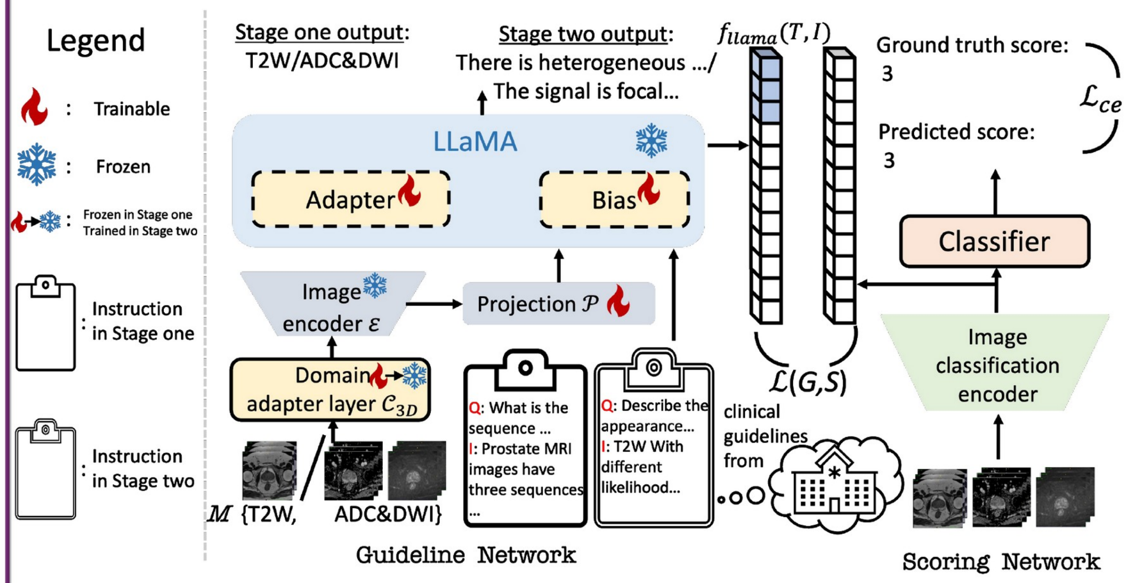

- **Two-Stage Fine-Tuning Strategy**:
  - ✓ Stage One: Developed a domain adapter layer to handle 3D MRI inputs and instructed the MLLM to differentiate MRI sequences.
  - ✓ Stage Two: Translated PICG into guiding instructions to generate PICG-guided image features and aligned these with the scoring network's features.
- **Feature Distillation**: Incorporated PICG information into the scoring network through feature distillation, minimizing the need for architectural changes.

28

Prediction of Breast Cancer Risk in Women Aged 40-50 Using BERT-Based Model

Maria Elkjær Montgomery[1], Mads Nielsen[1]
[1]Pioneer Center for AI, Department of Computer Science, University of Copenhagen, Denmark

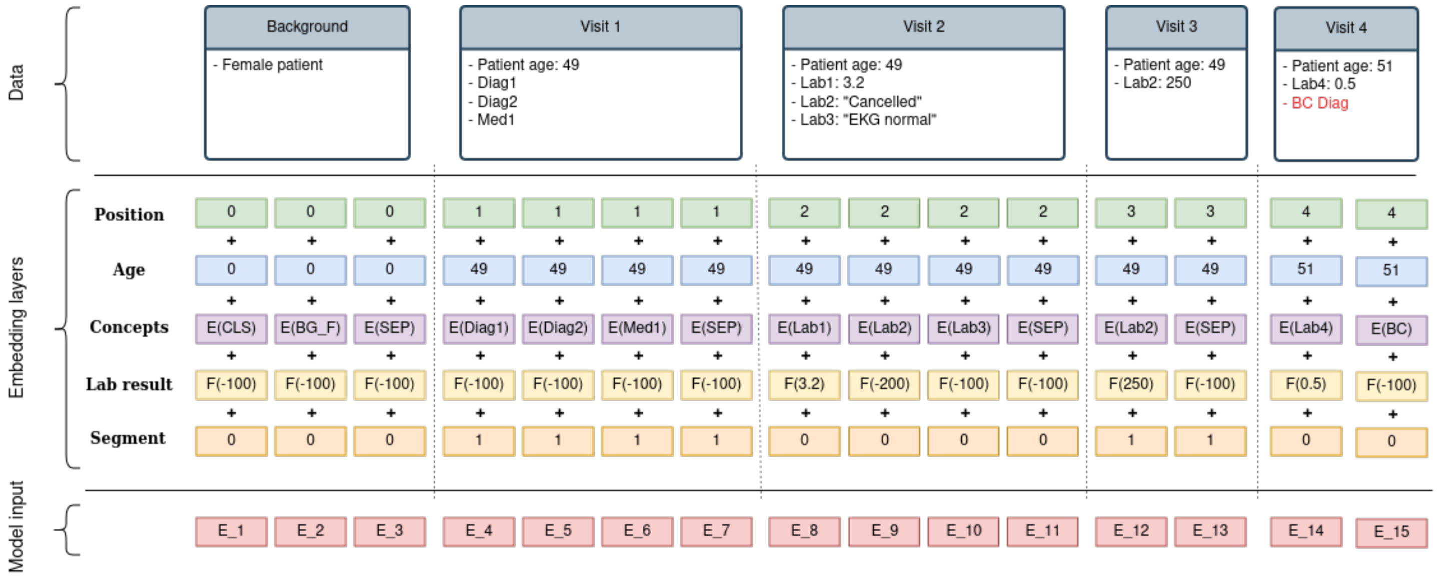

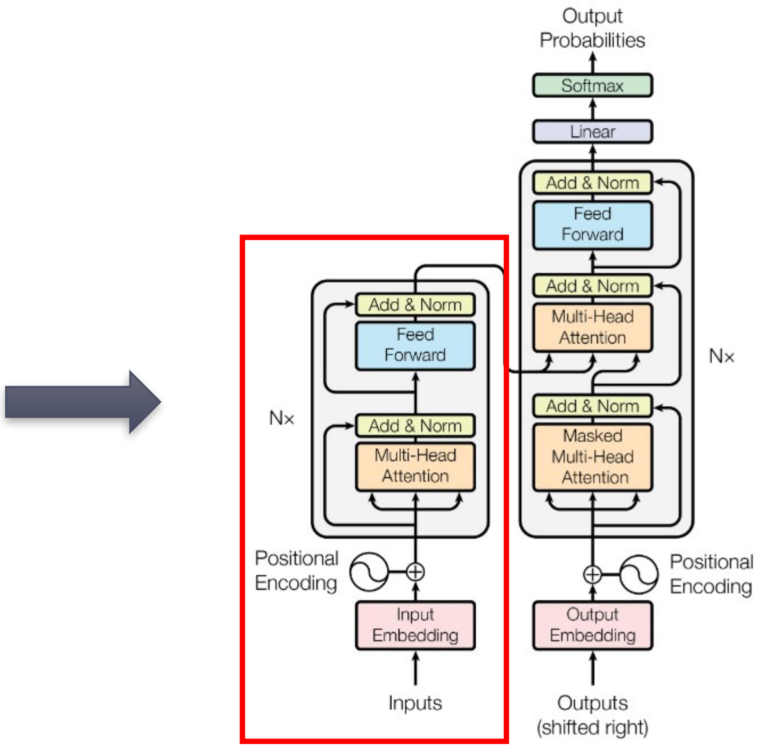

29

Implicit Neural Representations for Registration of Left Ventricle Myocardium During a Cardiac Cycle

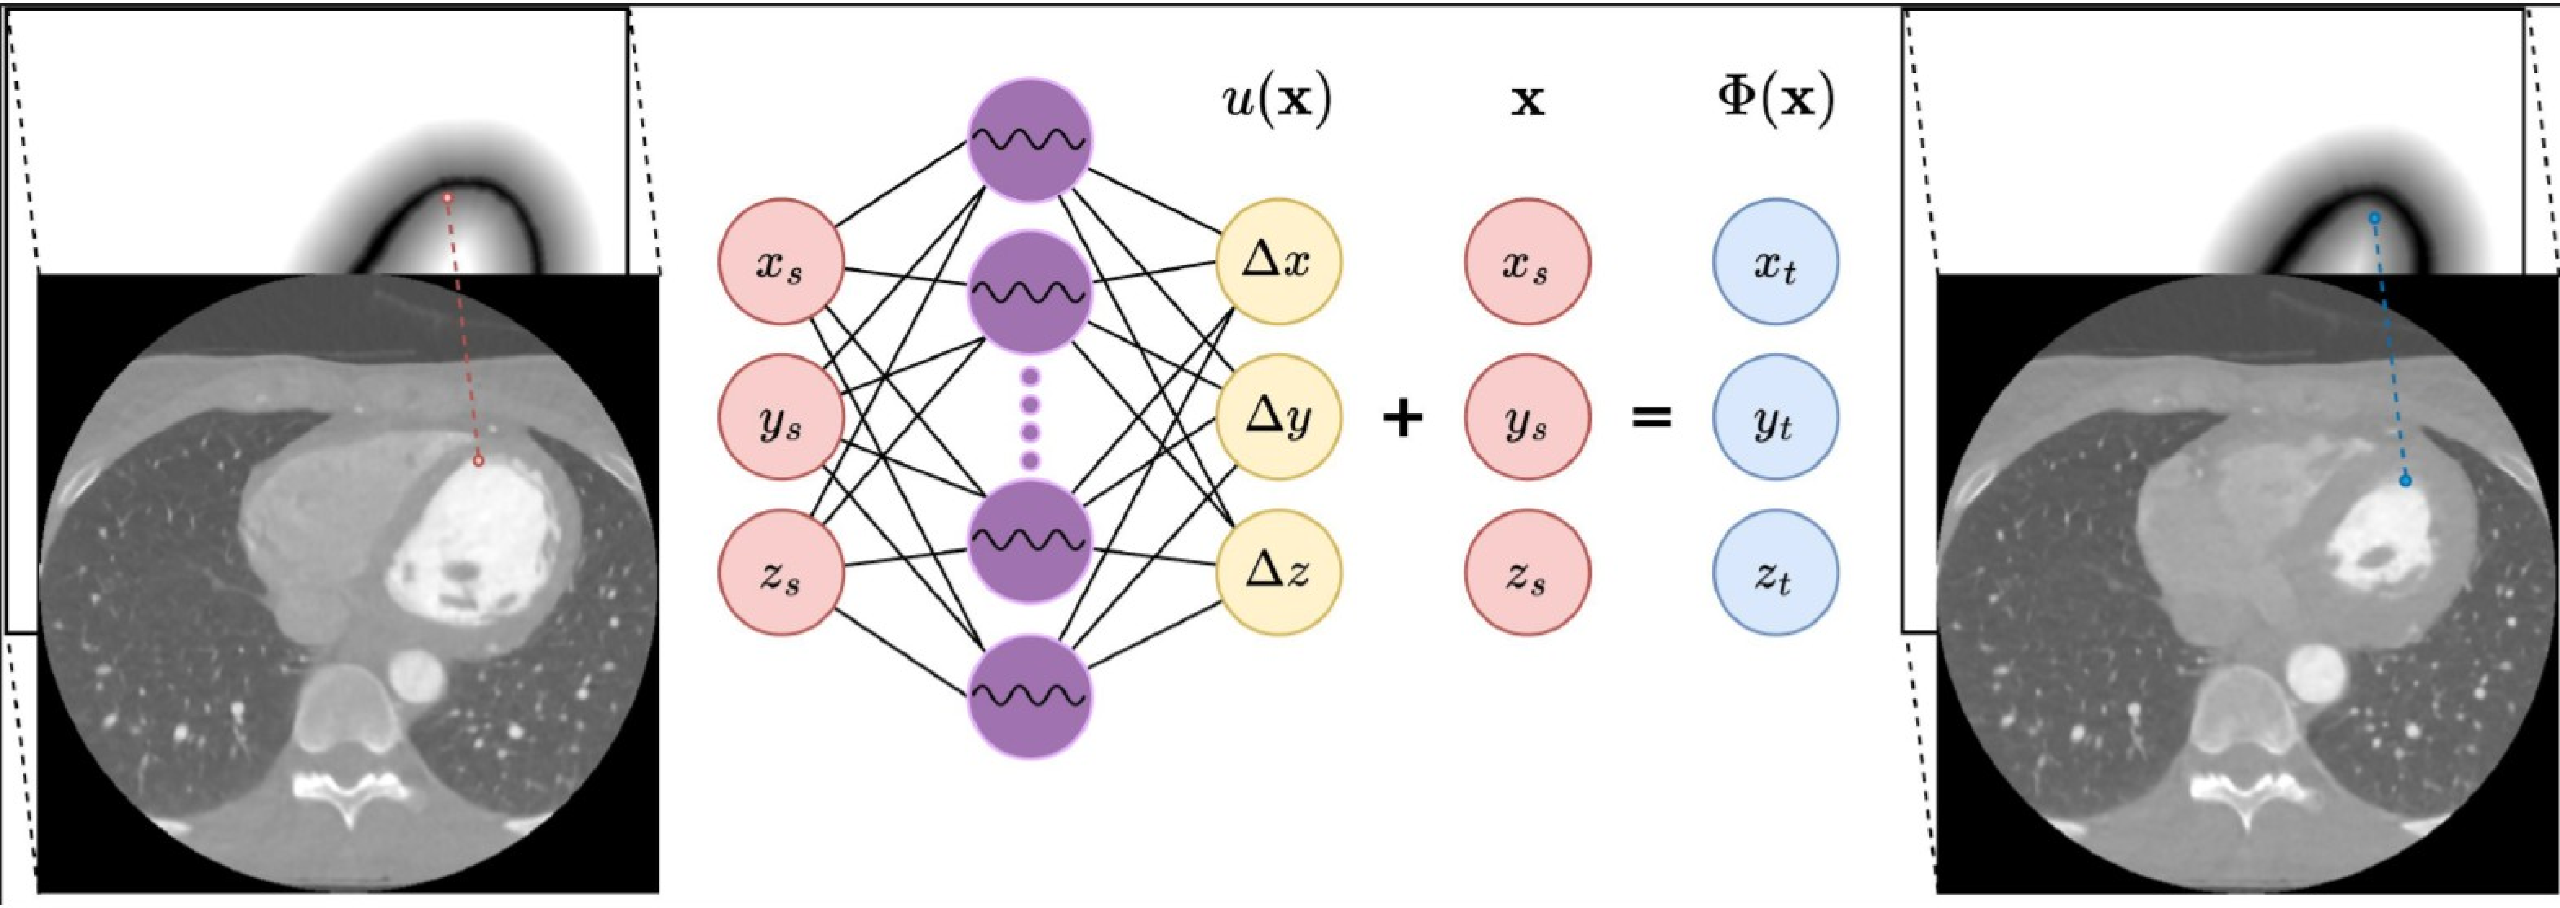

30

milestone

Detecting Anomalies in Videos

Mia Siemon

Industrial PhD

Can we simplify it?
What is the minimum?

Could it be the analysis of mere *bounding boxes*?

Can we simplify it?
What is the minimum?

How about object *silhouettes*?

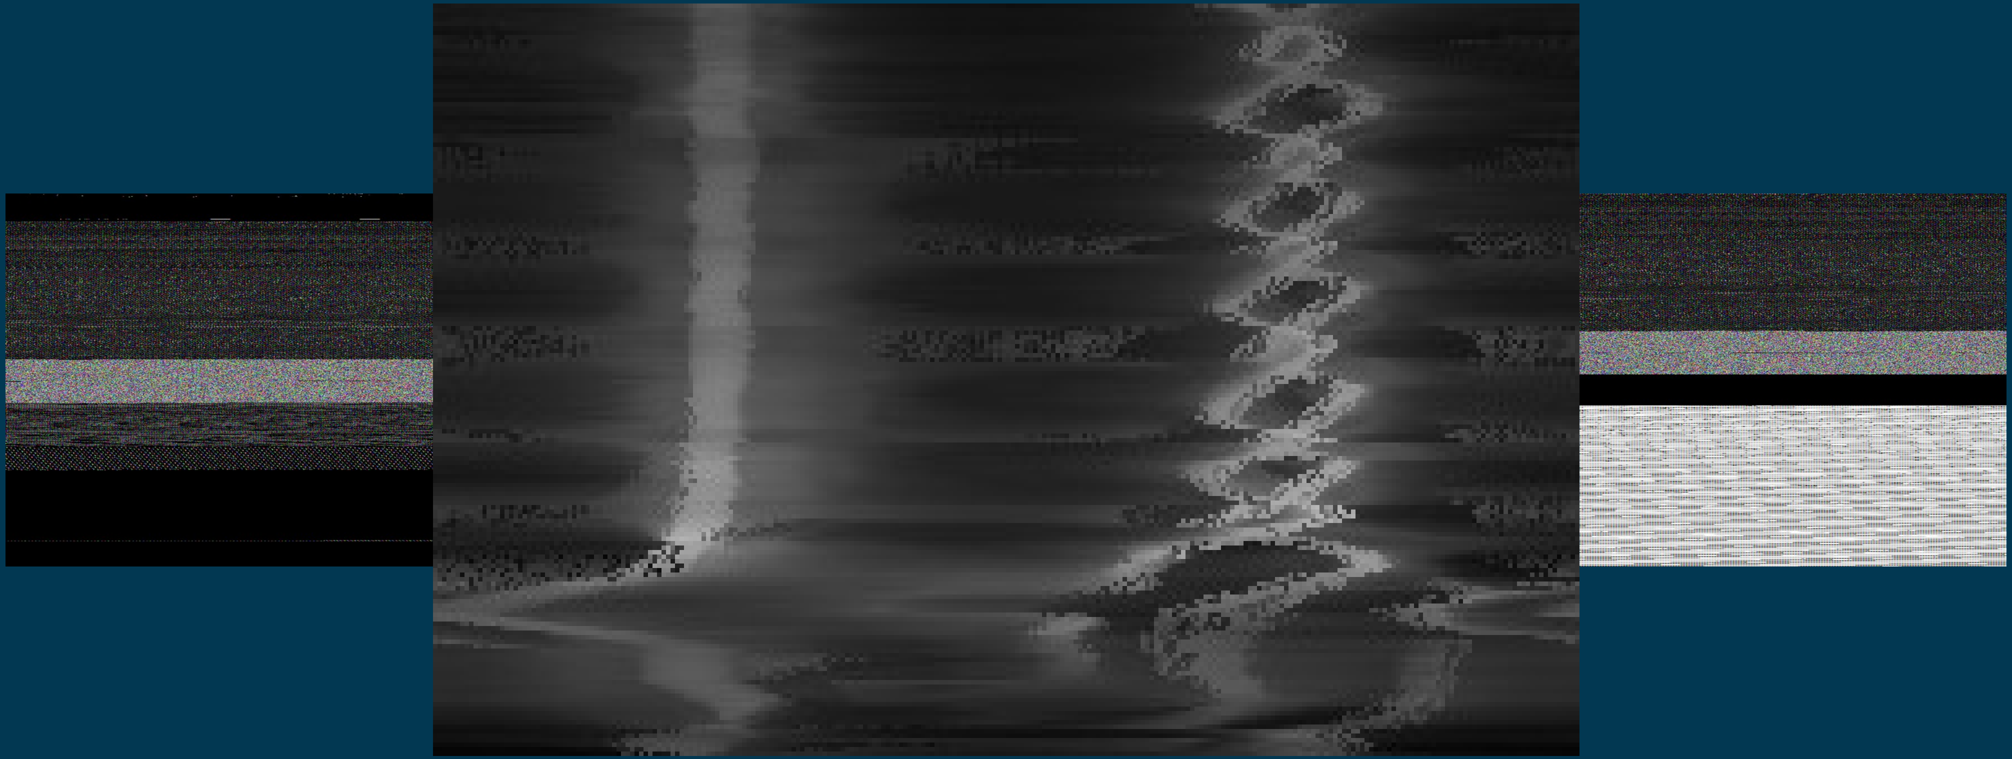

Can we simplify it?
What is the minimum?

+ 4%

Detection Accuracy
Compared to Prior Art

Can we simplify it?
What is the minimum?

+ 4%

Detection Accuracy
Compared to Prior Art

< 10s

Model Training Time

Can we simplify it?
What is the minimum?

Yes, we can.

Bounding Boxes and Silhouettes.

Can we simplify it?
What is the minimum?

31

Policy-Space Diffusion
for Physics based Character Animation

Michele Rocca, PhD student
Supervised by:
Kenny Erleben, Sune Darkner, and Sheldon Andrews

DIKU-DTU-AAU Summer School, 2024

UNIVERSITY OF COPENHAGEN

ÉTS Le génie pour l'industrie
ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
Université du Québec

Dance Motion Policy

On trained character

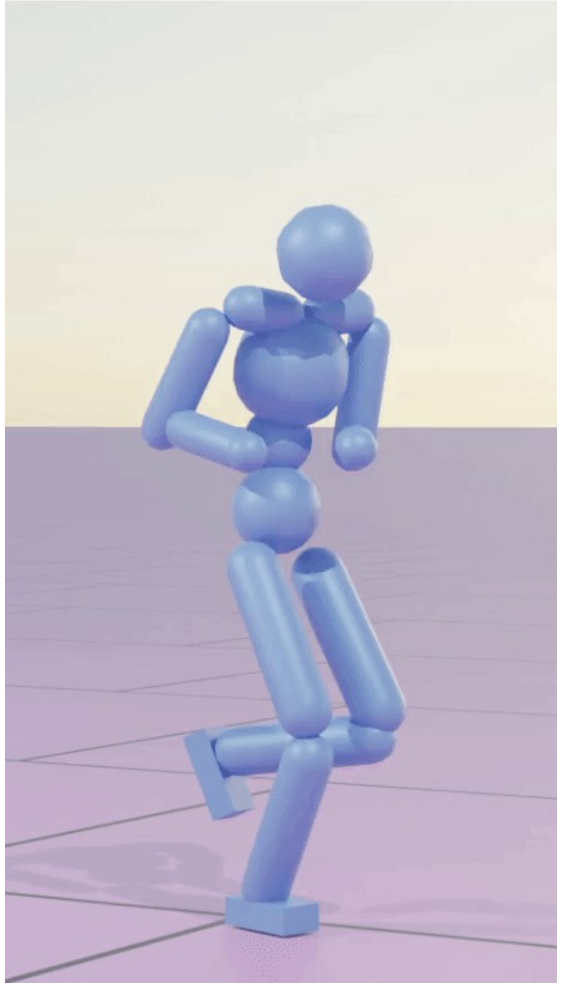

On unseen characters

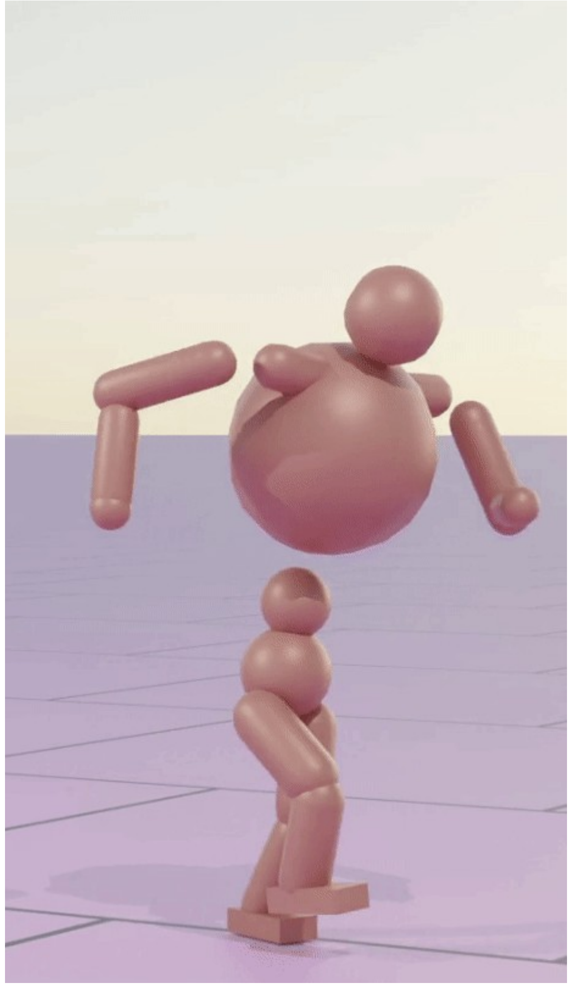

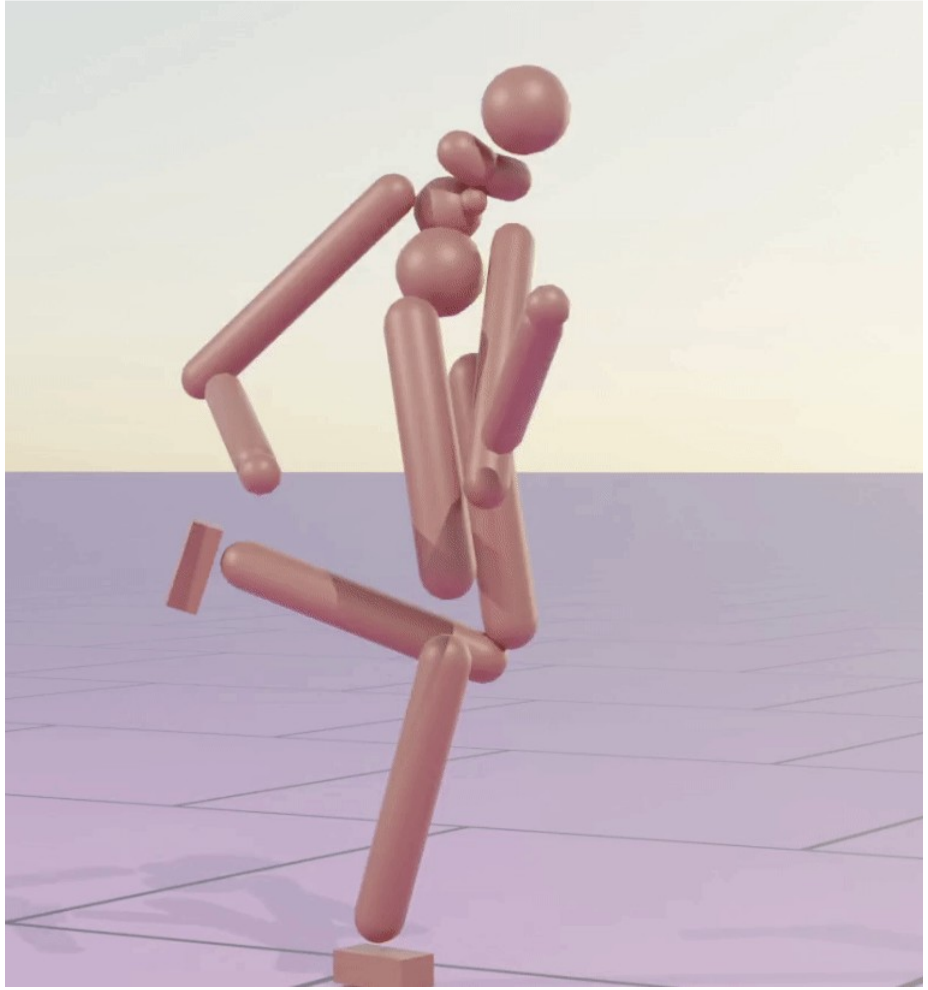

Regularized training:

similar motions->similar policies

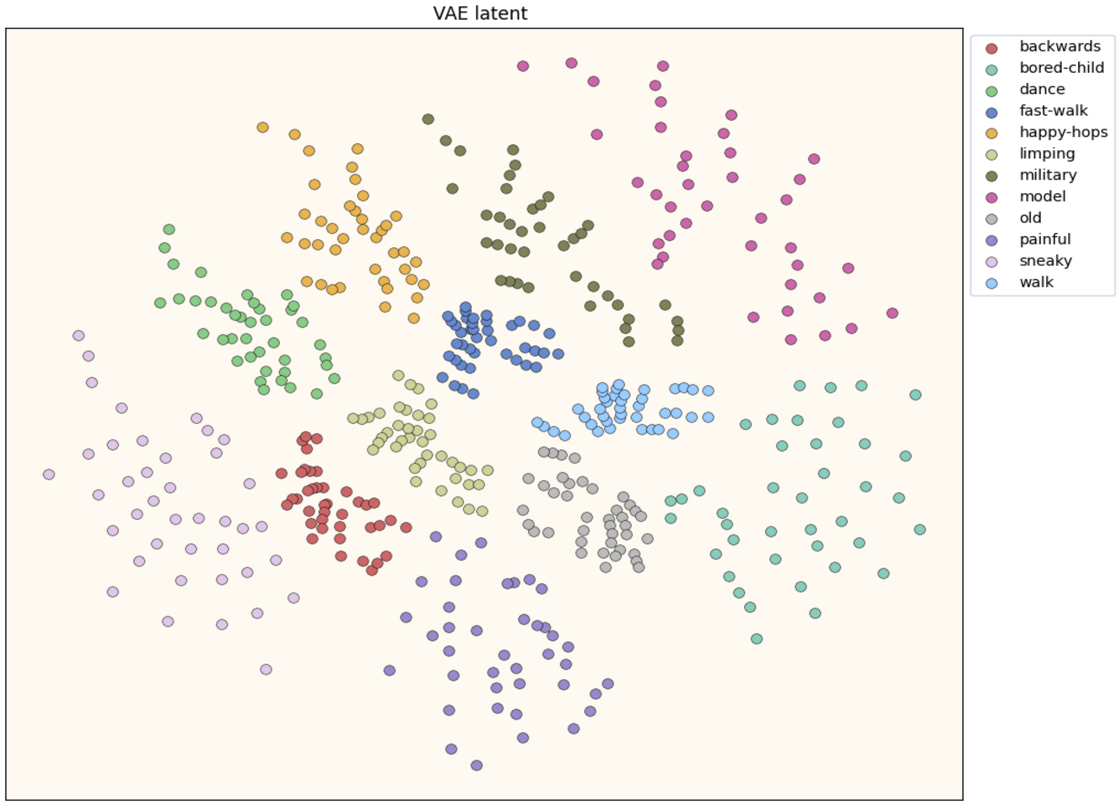

Policy Diffusion Model:

Discrete set of trained policies -> Continuous sampling

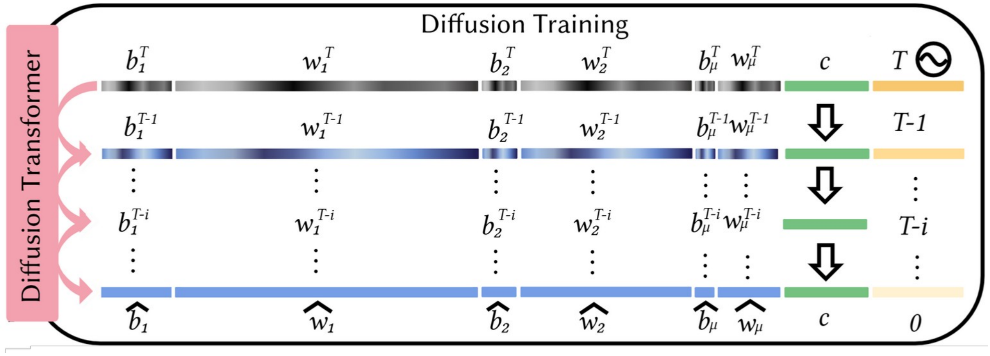

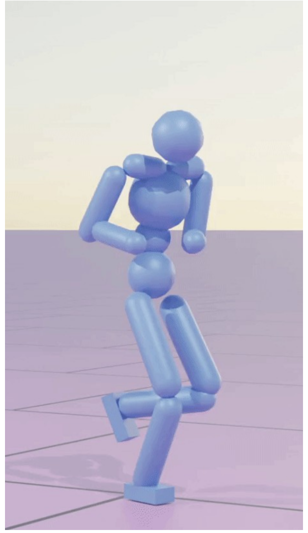

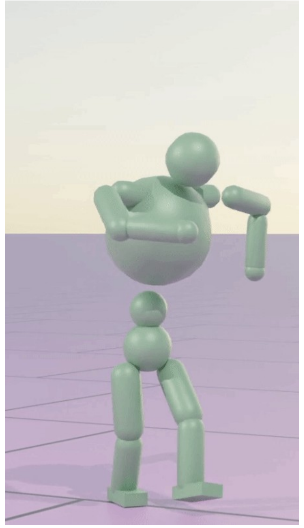

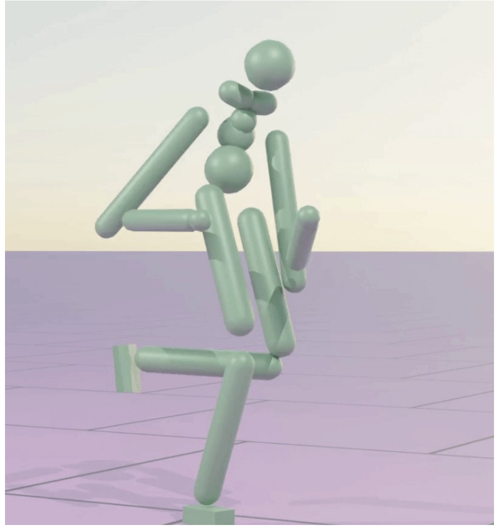

32

Counterfactuals is all about ***“What if…”***
Fx. What if Harry Potter went to Slytherin? 😈

Shortcut learning happens when there’s ***spurious correlation***
Fx. Scans from patients are more likely with support devices 🤔

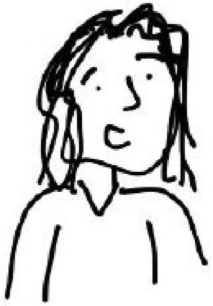

So what if the support devices are removed?

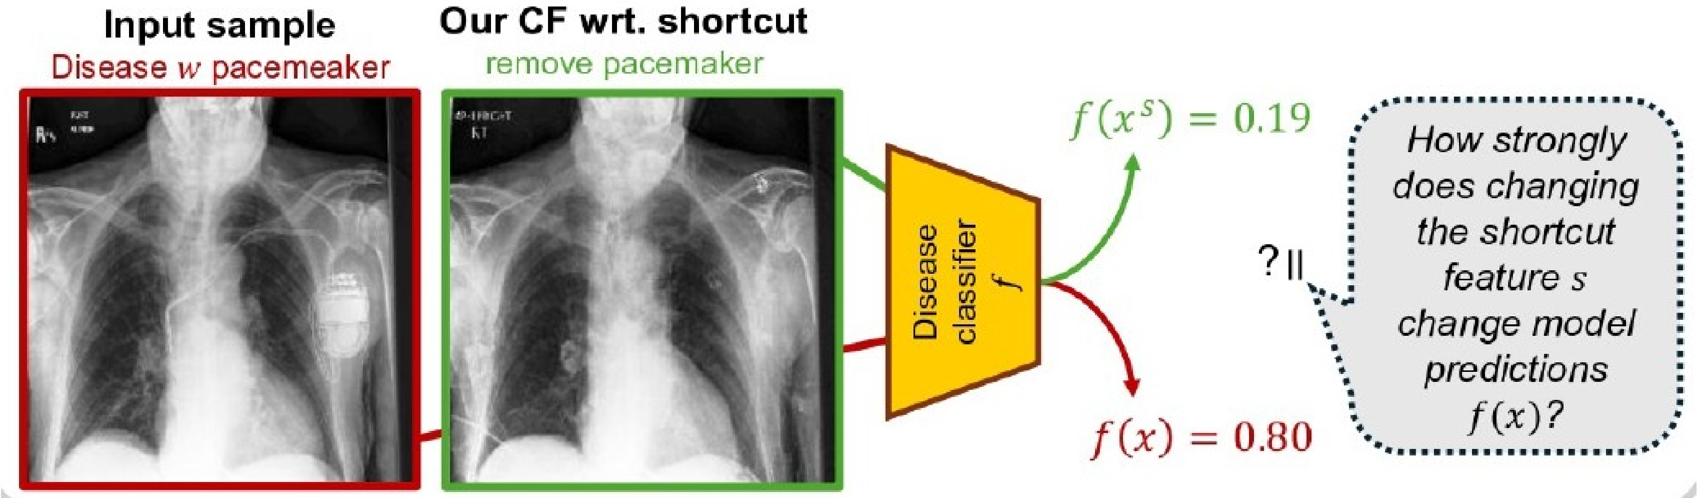

Sounds interesting? Check our poster!

Fast Diffusion-Based Counterfactuals for Shortcut Removal and Generation

Nina Weng*, Paraskevas Pegios*, Eike Petersen, Aasa Feragen, Siavash Bigdeli
Technical University of Denmark, Denmark   * Equal contribution

33

Structural analysis of mozzarella cheese

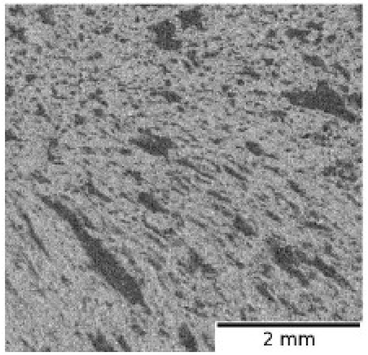

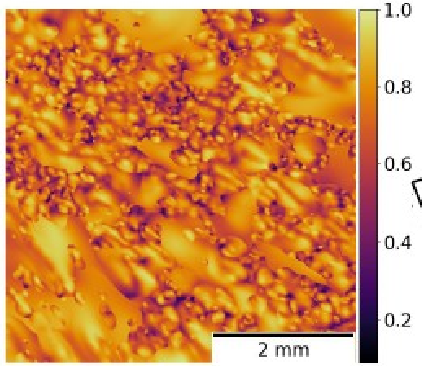

???

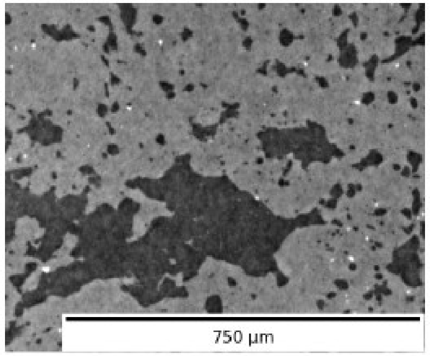

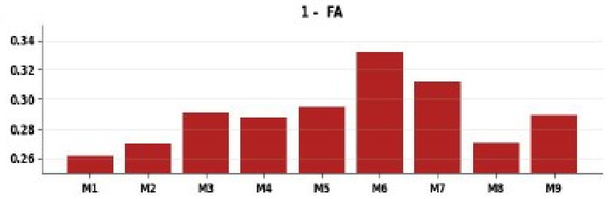

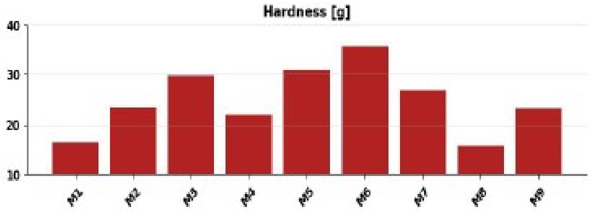

34

Stochastic subset selection in partially Bayesian transformers

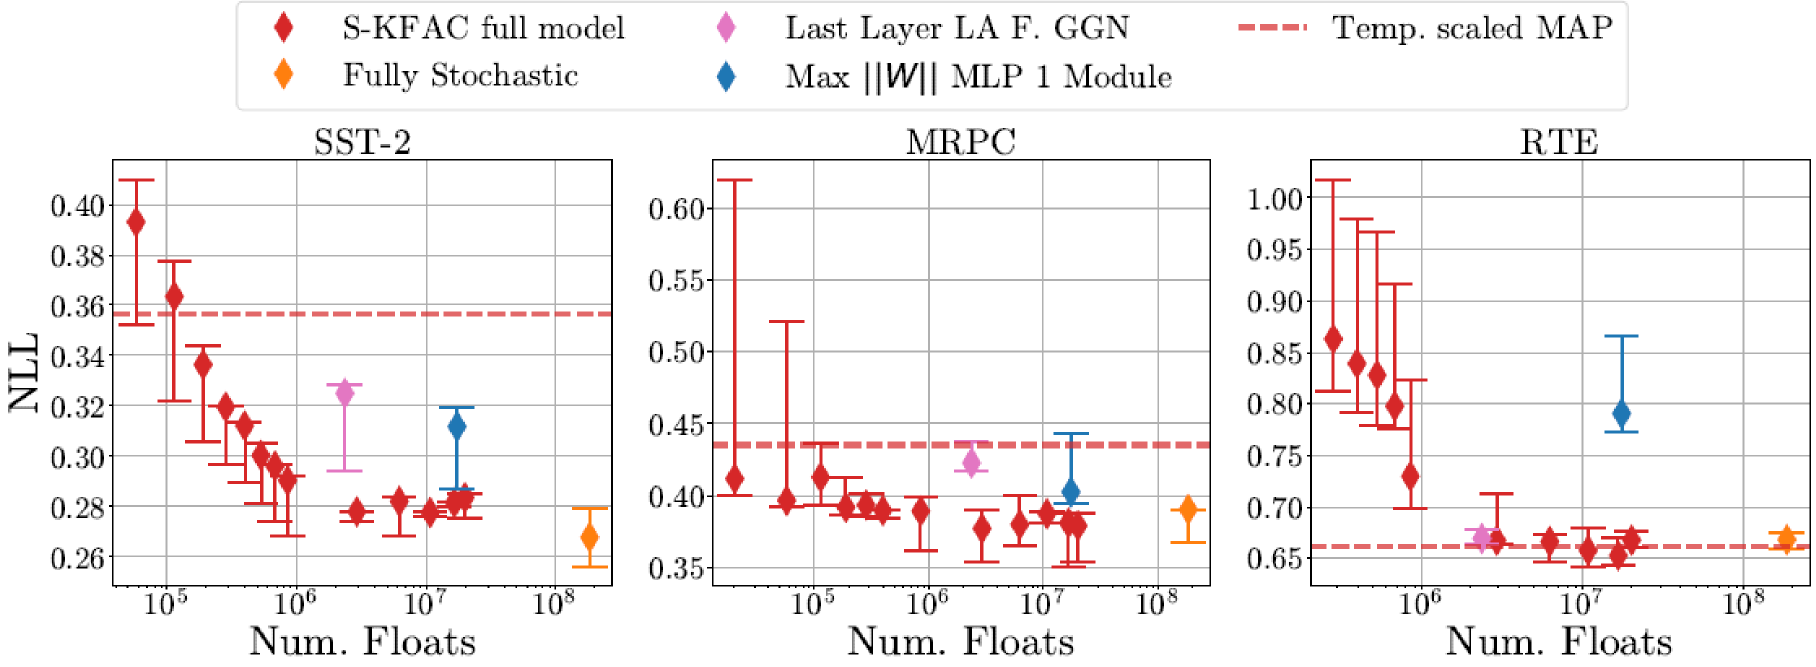

What can we do with partially stochastic models?

- We can achieve fully (approximate) Bayesian behaviour from only partially stochastic models if...
  - We select the weights (somewhat) intelligently
  - We include the prior variance as a hyperparameter for optimization

35

Search in free-text radiology report database using large language model cooperation

Reza Karimzadeh, Bulat Ibragimov

Search in free-text radiology report database using large language models

- Itemization: Llama 2, Orca 2, Yi chat models
- Find the closest embedding of items for the query

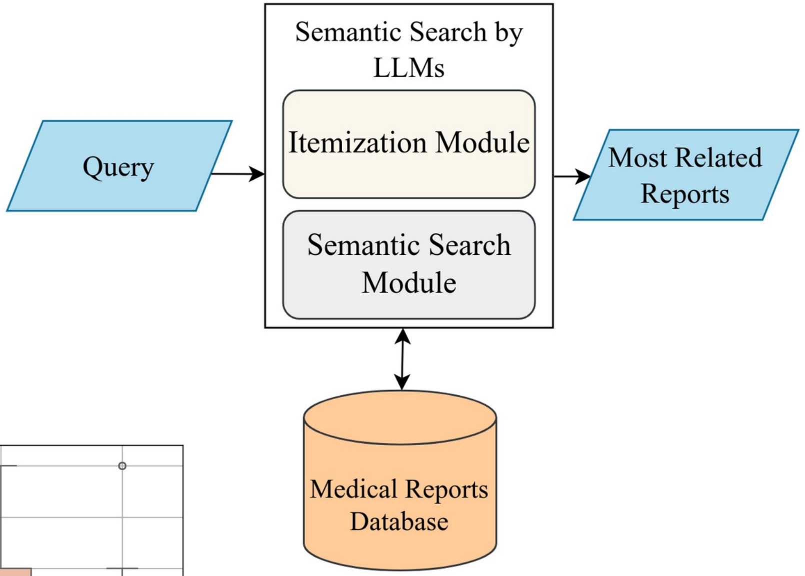

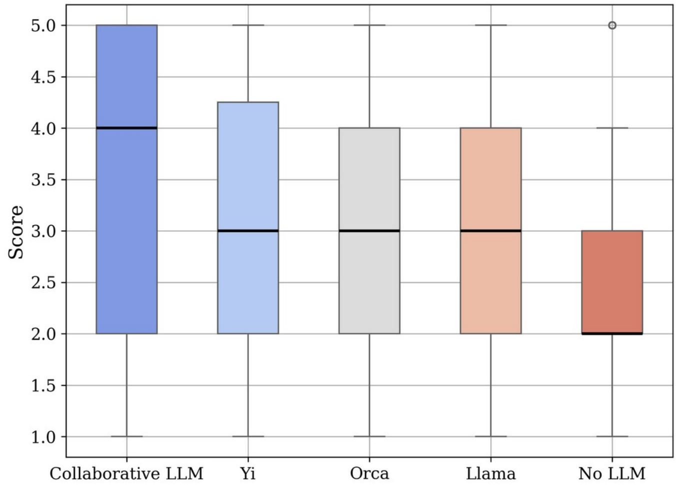

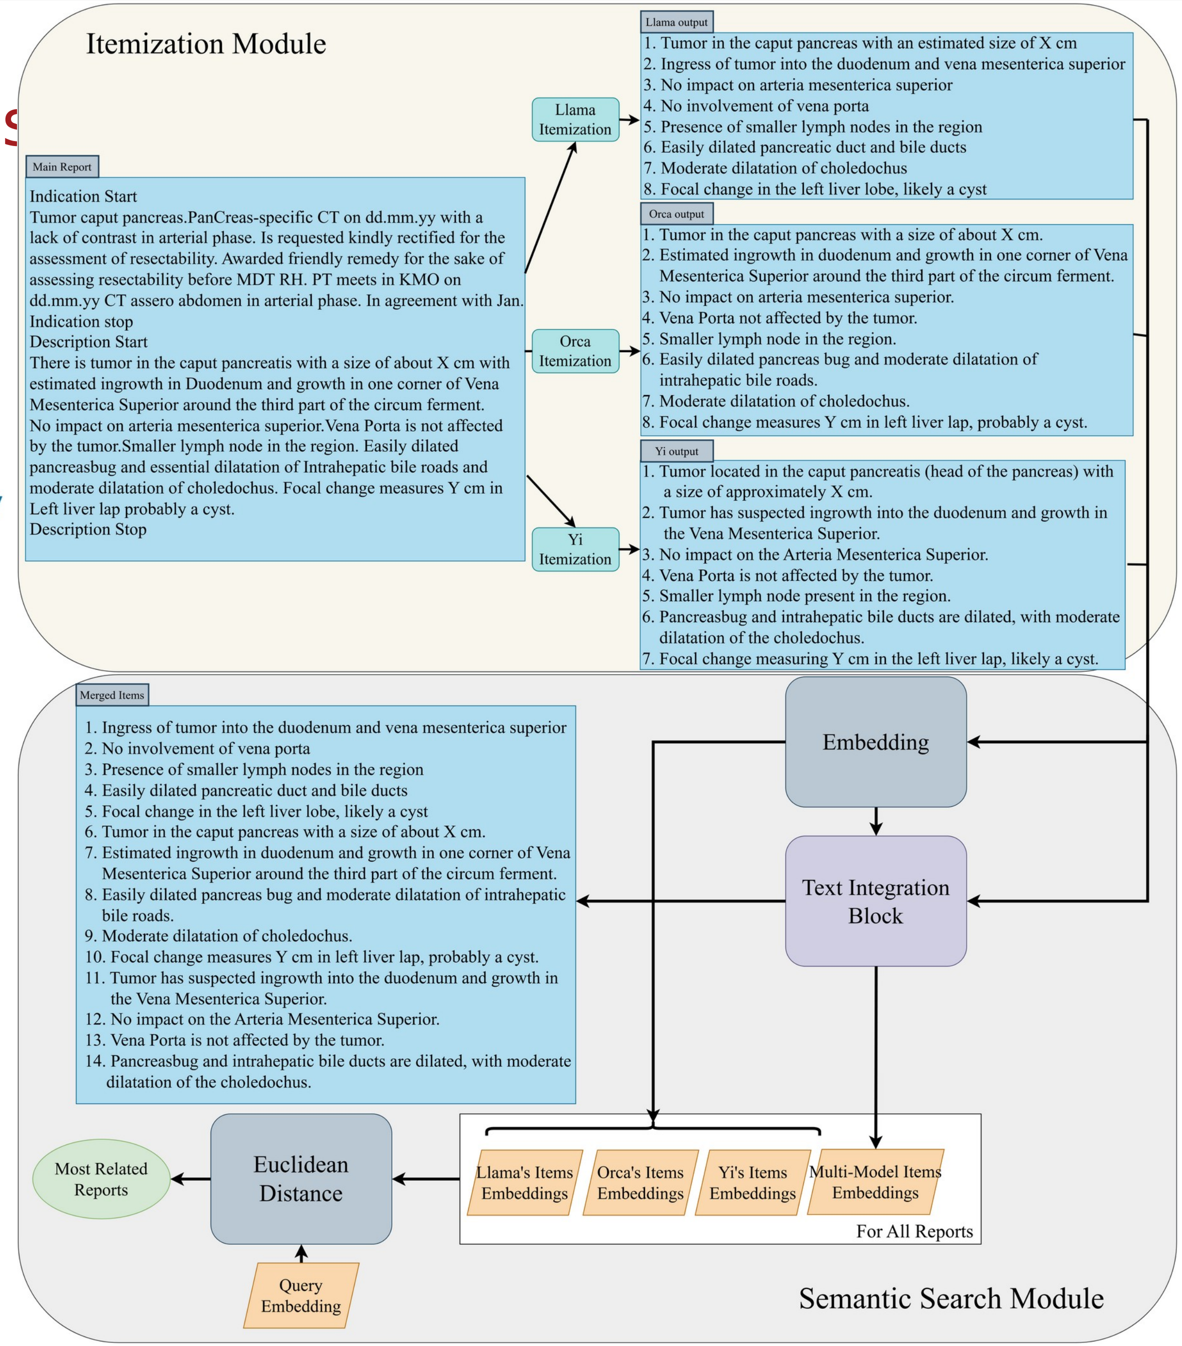

36

Synthesis of Geometric Models for Axons

Ruiqi Cui, Sidsel Winther, Tim B. Dyrby, Andreas Bærentzen

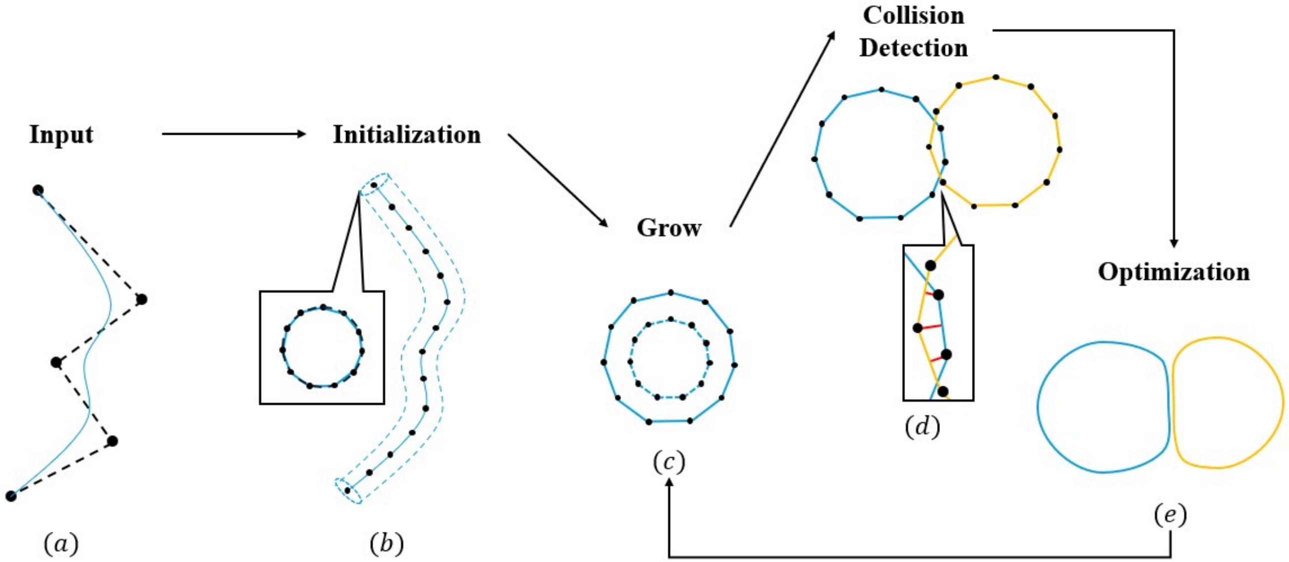

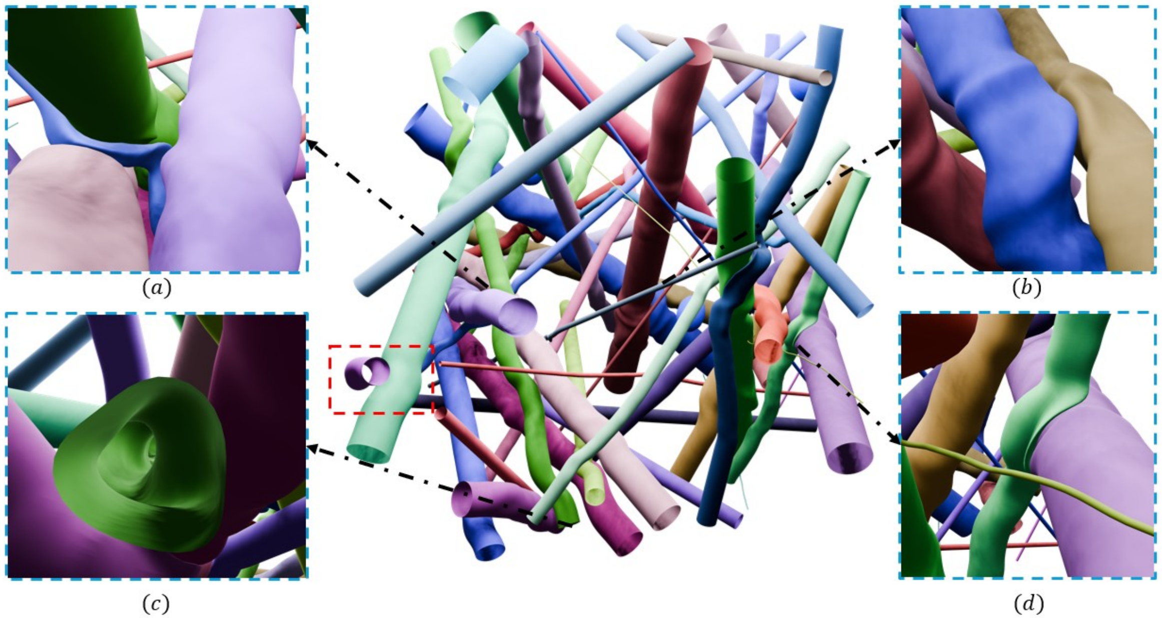

37

Pre-training LLMs require big GPU clusters

Pre-training a **LLaMA 13B** model with Adam and a *batch size of one* requires over **100 GB** of memory.

Pre-training LLMs require big GPU clusters

Pre-training a **LLaMA 13B** model with Adam and a *batch size of one* requires over **100 GB** of memory.

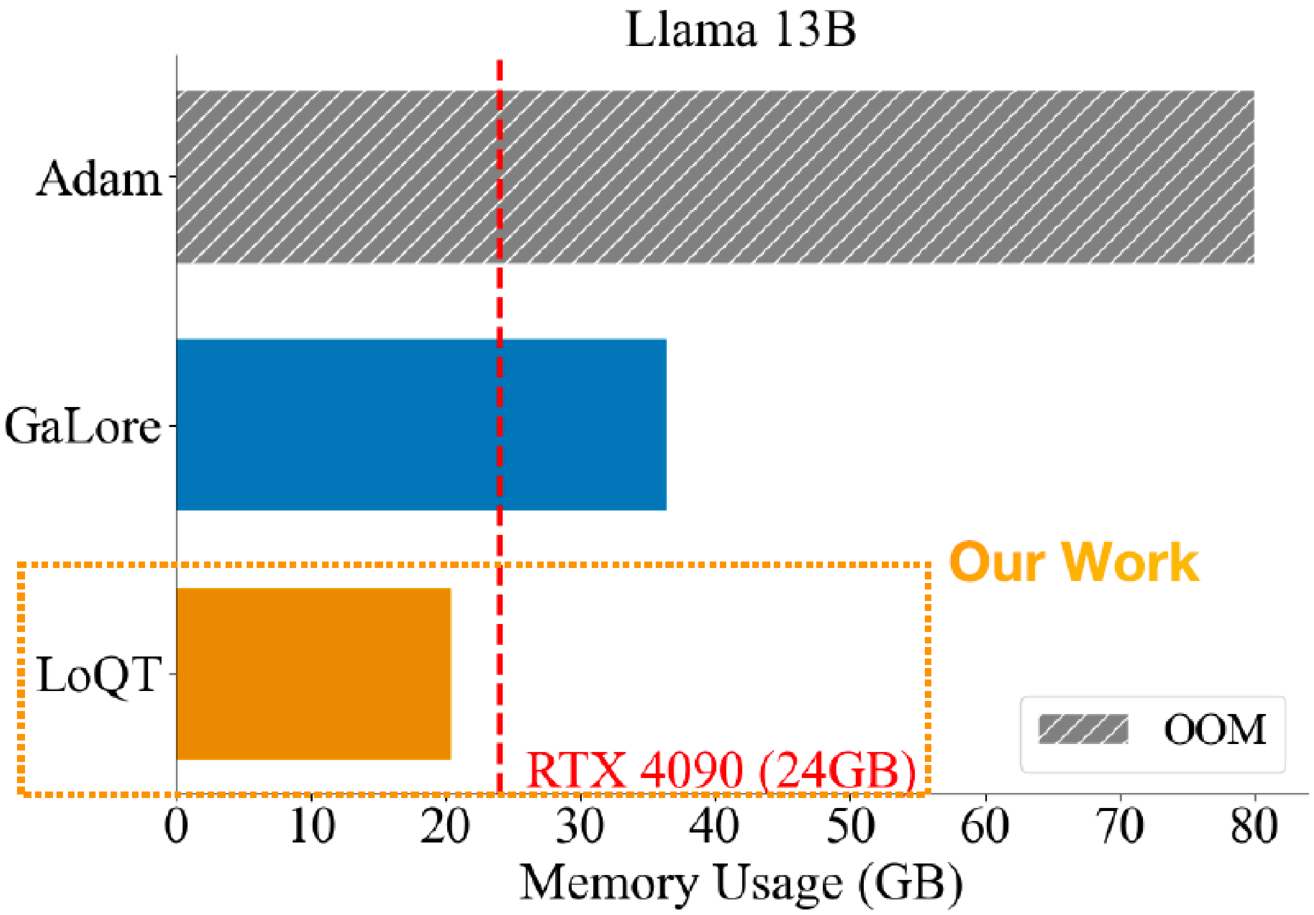

LoQT: Low Rank Adapters for Quantized Pre-Training

- **Pre-training from scratch** while only *optimizing a **single low-rank factor** per layer.*

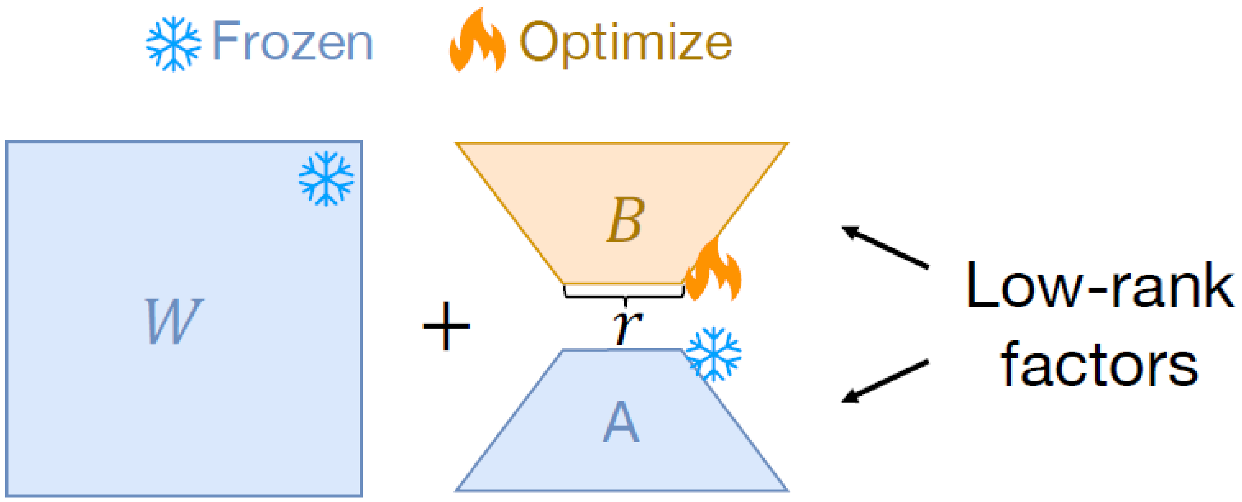

LoQT: <u>Lo</u>w Rank Adapters for <u>Q</u>uantized Pre-<u>T</u>raining

- **Pre-training from scratch** while only *optimizing a **single low-rank factor*** *per layer.*
- Keep most *weights in **4-bit precision.***
- <u>*Match regular FP16 training*</u>

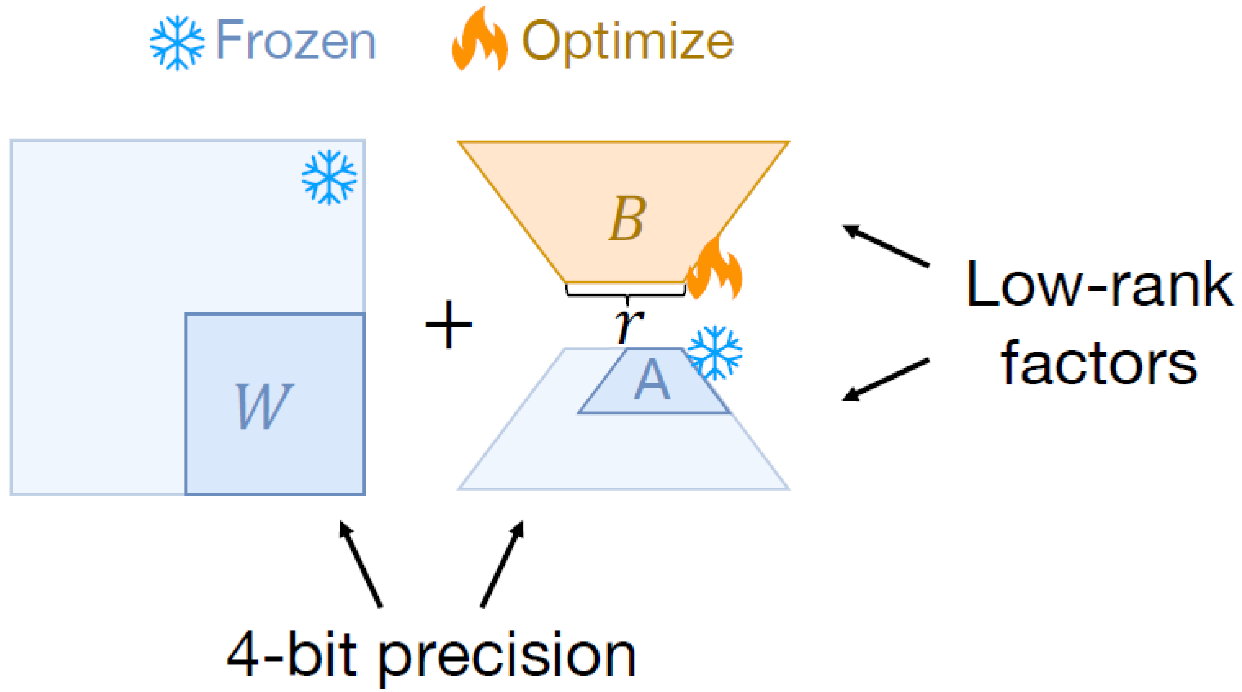

LoQT: Low Rank Adapters for Quantized Pre-Training

Sebastian Loeschcke*, Mads Toftrup*, Michael J. Kastoryano, Serge Belongie, and Vésteinn Snæbjarnarson

UNIVERSITY OF COPENHAGEN

AARHUS UNIVERSITY

IT UNIVERSITY OF CPH

PIONEER CENTRE FOR ARTIFICIAL INTELLIGENCE

*Equal Contribution

38

SUPSI

Protoporphyrin fluorescence quantification in glioblastoma tumor phantoms

Marinelli S.[1], Oberli C.[1], Mazevet M.[1], Kaelin A.[2], Marchi F.[2], Cardia A.[2], Reinert M.[3], Allegri D.[1], Gardenghi R.[1]

[1]SUPSI-Department of Innovative Technologies, Institute of Systems and Applied Electronics, [2]Neurocenter-Instituto di Neuroscienze cliniche della svizzera italiana, [3]USI-Università della svizzera italiana

Sebastiano Marinelli, Eng.

31 luglio 2024

Content

- **Context**: Protoporphyrin (PpIX), 5-ALA metabolite, is used in fluorescence-guided surgery for tumours resection.

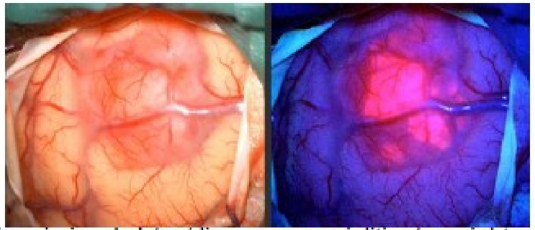

https://neurochirurgie.insel.ch/en/diseases-specialities/special-techniques/5-ala-fluorescence

- **Aim:** Need for quantification to reduce inter-observer variability.
- **How**:
  - Development of two models.

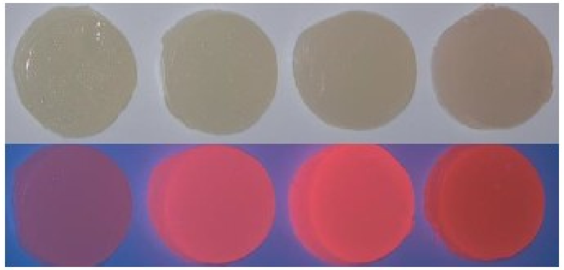

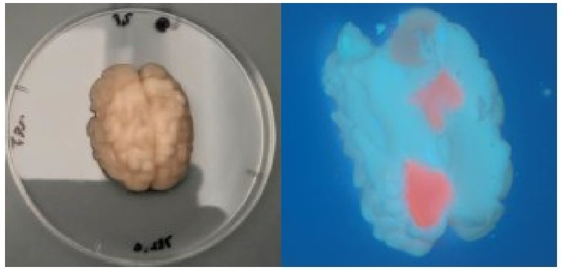

  - Development of a quantification technique with calibration curve
  - Comparison in two systems (Orbeye from Olympus and FLUO custom made setup from SUPSI).
- **Results**:
  - Accurate quantification in gel using calibration curve
  - Relative quantification in small brain model.

39

3D Whole-heart fibrosis: Can we quantify it?

Collagen in a ReninAAV UNx db/db mouse heart

3D Whole-heart fibrosis: Can we quantify it?

db/db UNx-ReninAAV mouse model

reninAAV - Hypertension
Unx - Kidney disease
db/db Mice - Diabetes

Whole-heart collagen staining with 'Fast Green' dye

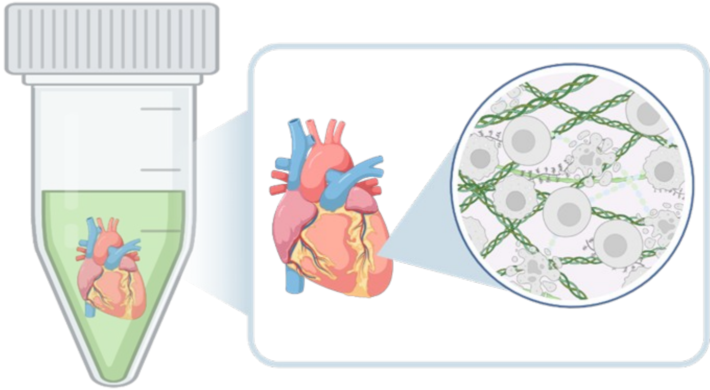

Clearing and 3D imaging

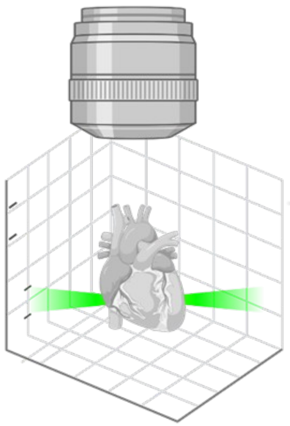

AI-assisted analysis

iDISCO:
+iDISCO standard protocol
+Fast green dye

Image analysis
+3D Whole-hea
quantifica collagen
volume

3D Light sheet imaging data

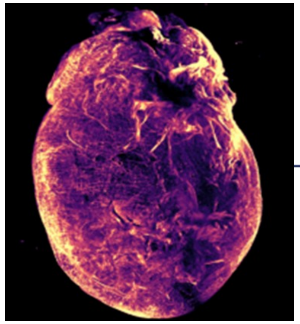

Preprocessing

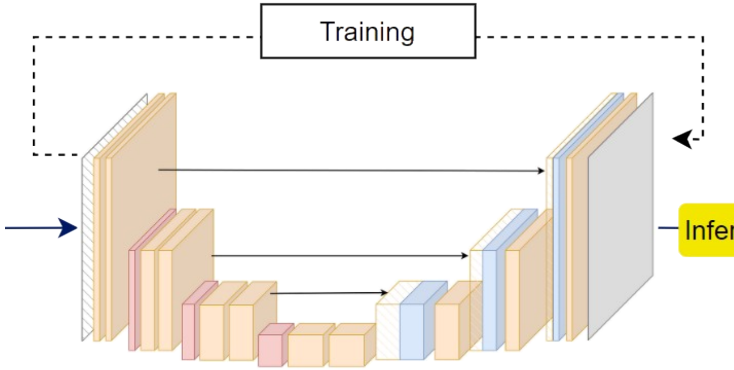

Inference

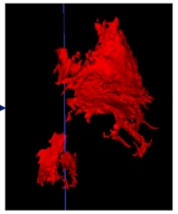

3D Fibrosis segmentation

High throughput cardiac analysis platform that opens the door to testing new therapeutics.

40

If we want models that can **enhance the resolution** of medical CT images, we need **real data** and better **performance metrics**.

Paper title: Superresolution of real-world multiscale bone CT verified with clinical bone measures

FACTS dataset

(**F**emur **A**rchaeological **CT** **S**uperresolution)

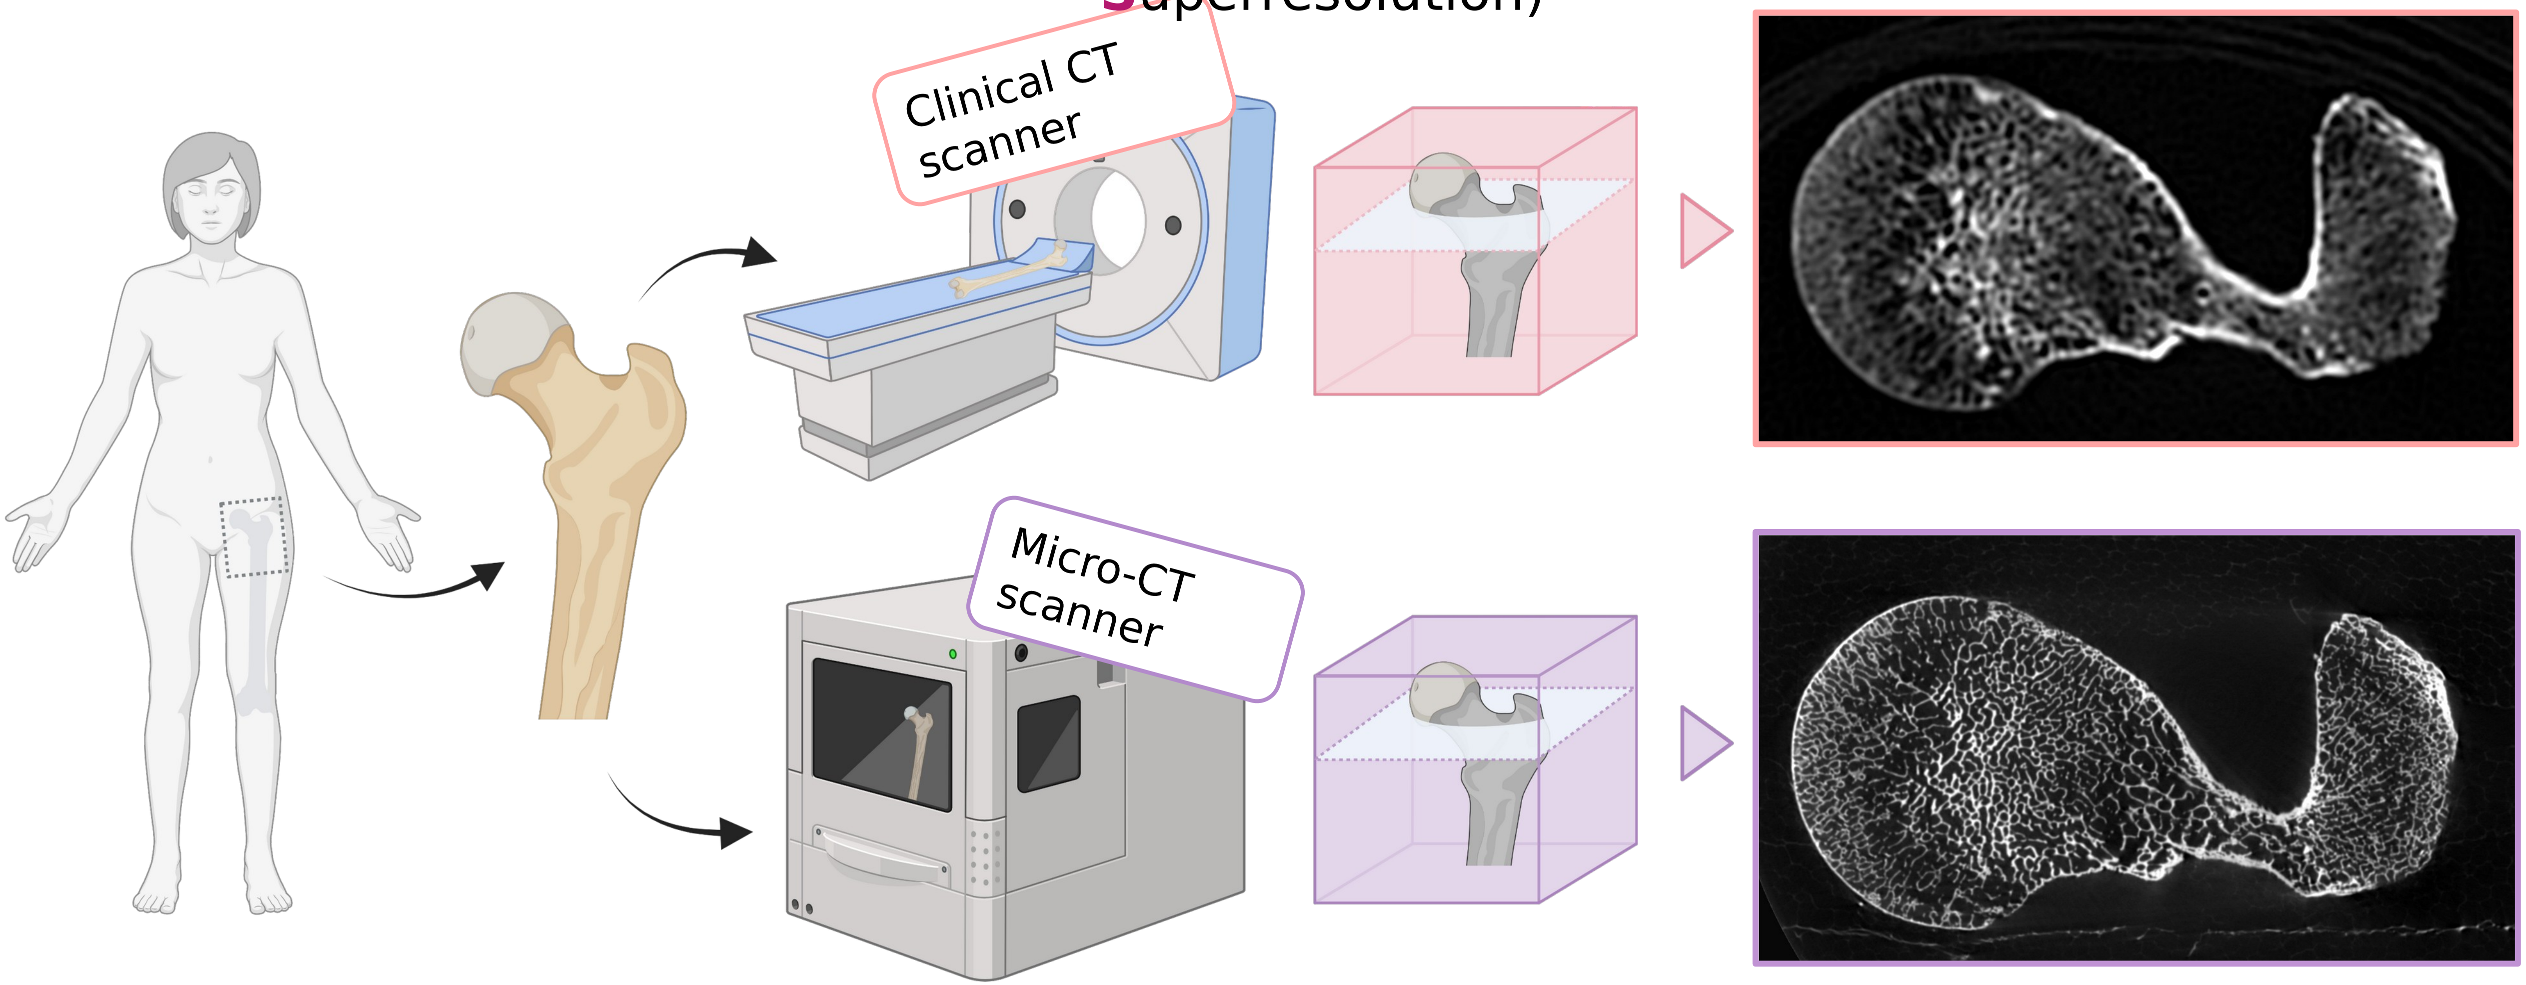

- 13 proximal femurs
- 2 CT resolutions
- Publicly available

41

Validating YOLO v8 and SAM Foundation Models for Robust Point-of-Care Ultrasound Segmentation

Sumit Pandey
Phd Fellow

Supervisor: Prof. Erik B Dam

Methodology

- Medical image segmentation performed using YOLOv 8 with SAM and HQ SAM models and comparing with variations of UNets

- All six models were trained on 510 images and corresponding masks from 175 patients and tested on an independent cohort 375 patients

- Analysis revealed that the YOLOv8 model outperformed both Unets models and YOLOv8 + SAM.

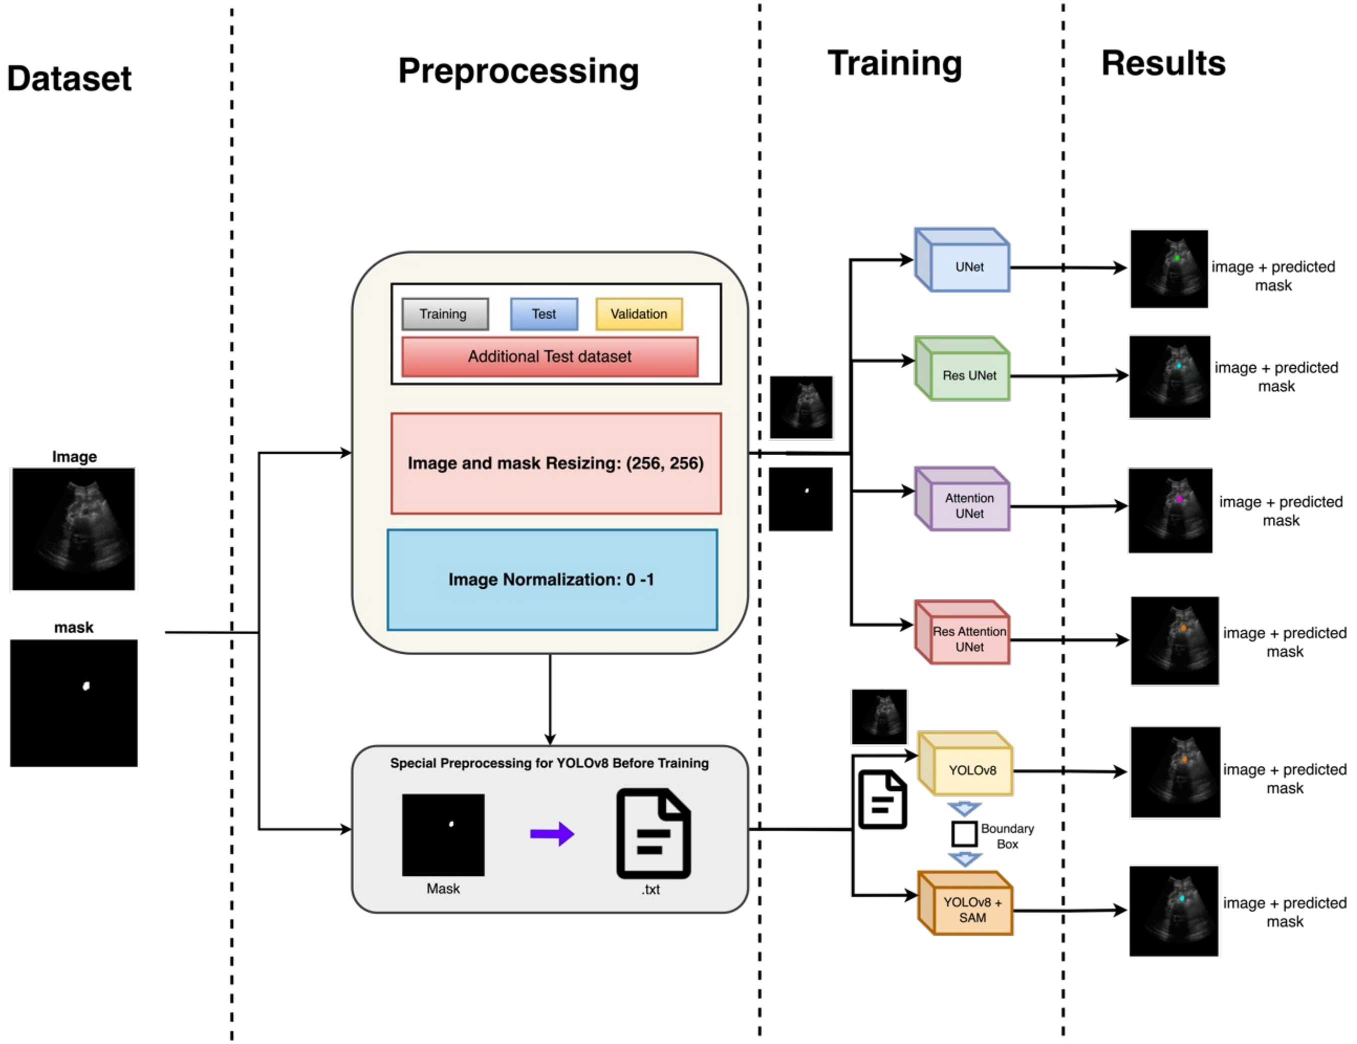

Result

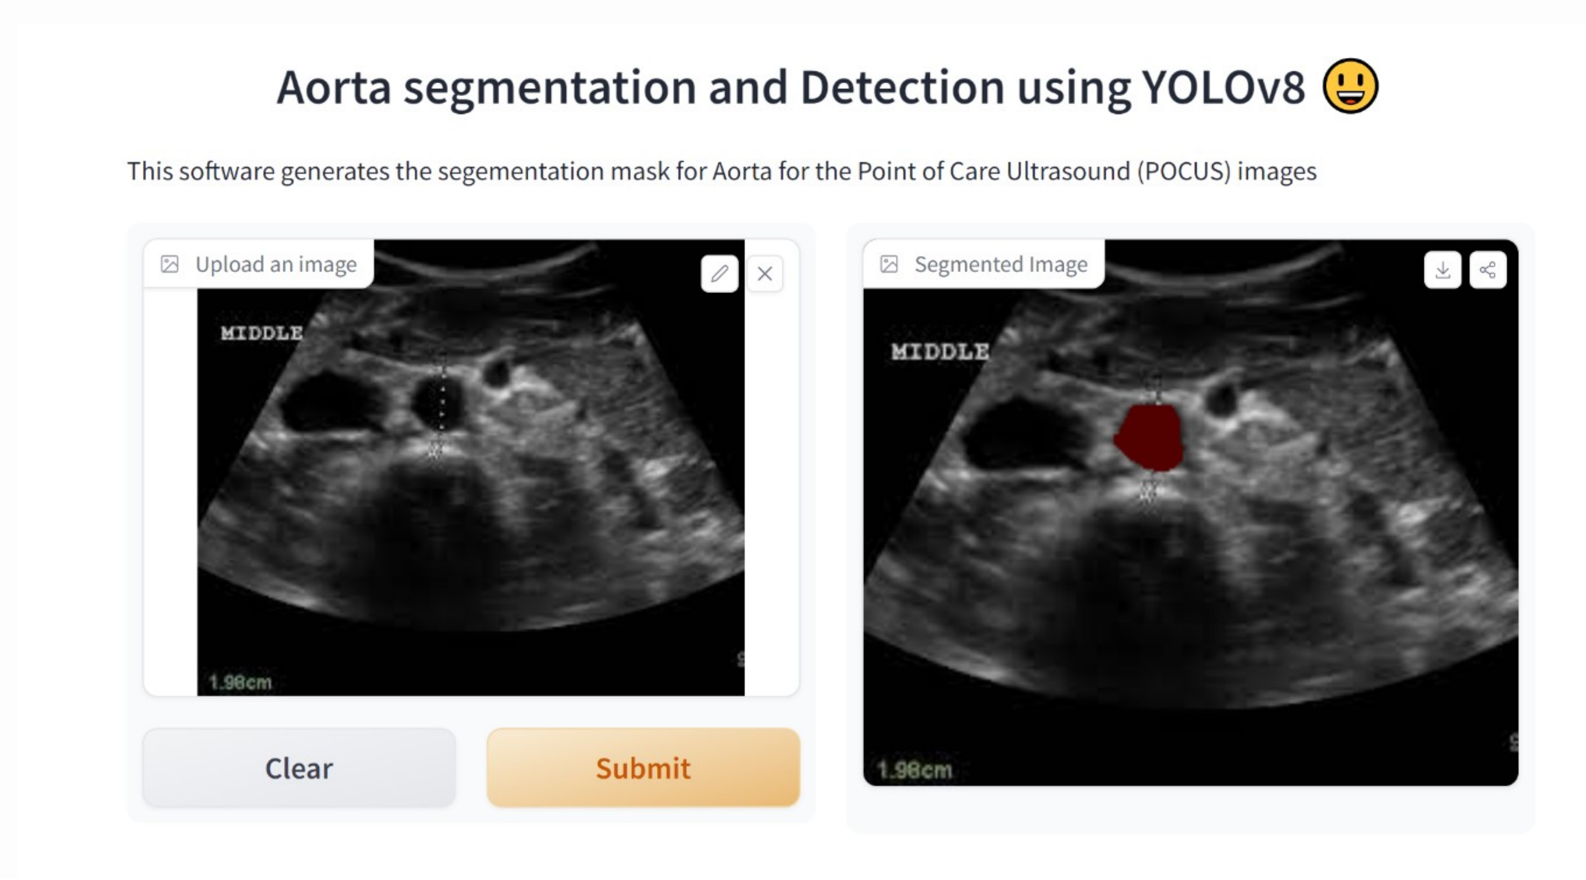

Please click here to check the web-app

42

FreqRISE: Explaining time series using frequency masking

- Time series data is **ubiquitous** in several **critical domains**
  - Health, finance, climate monitoring
- **Explainability** -> salient information in **latent domains**
- Can we provide **explanations** in **dual domains** where salient information is sparse?

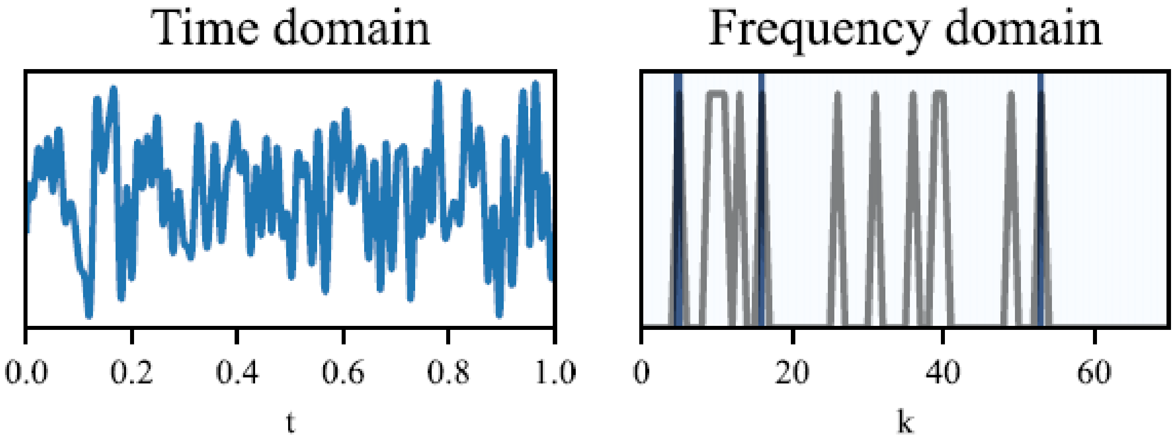

FreqRISE: Explaining time series using frequency masking

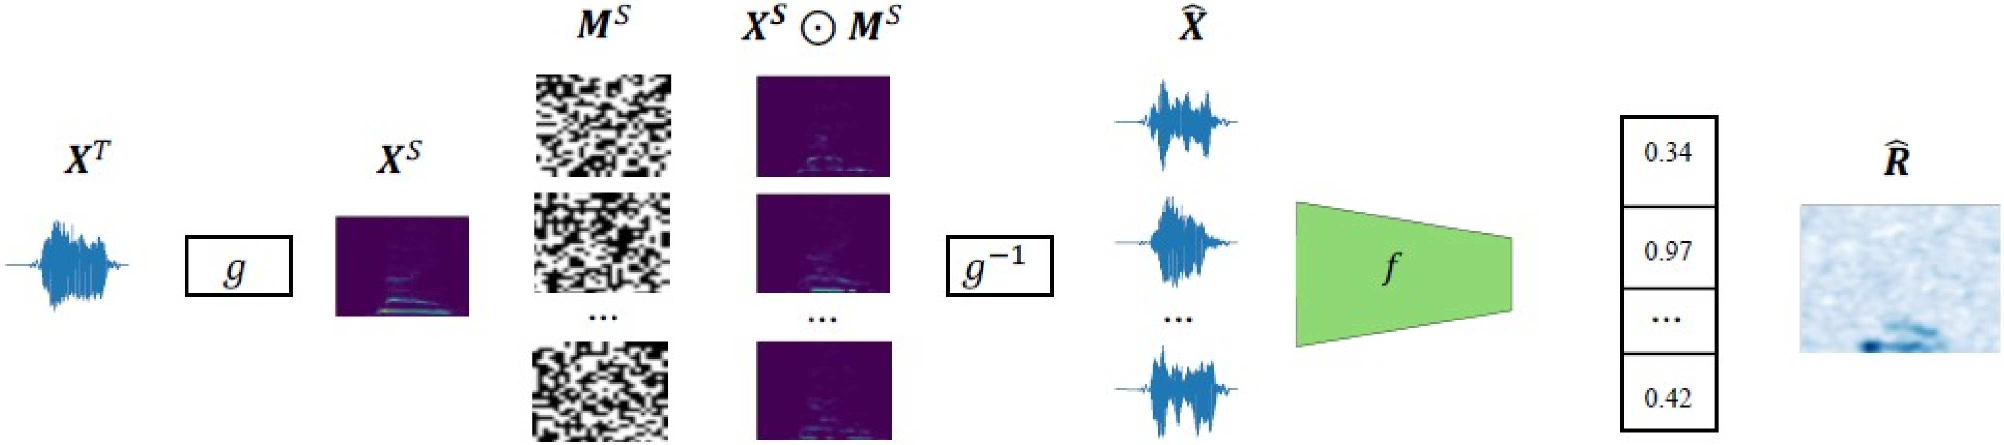

43

Quad mesh generation using RL

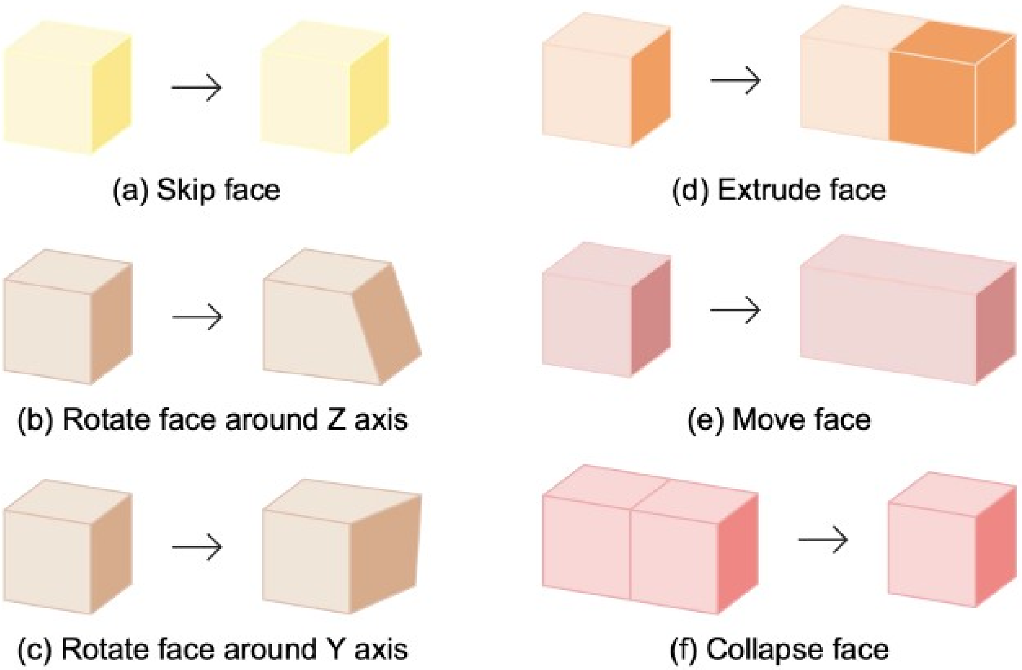

Figure by Alba Reinders Sánchez, Master thesis Project at DTU, 2024

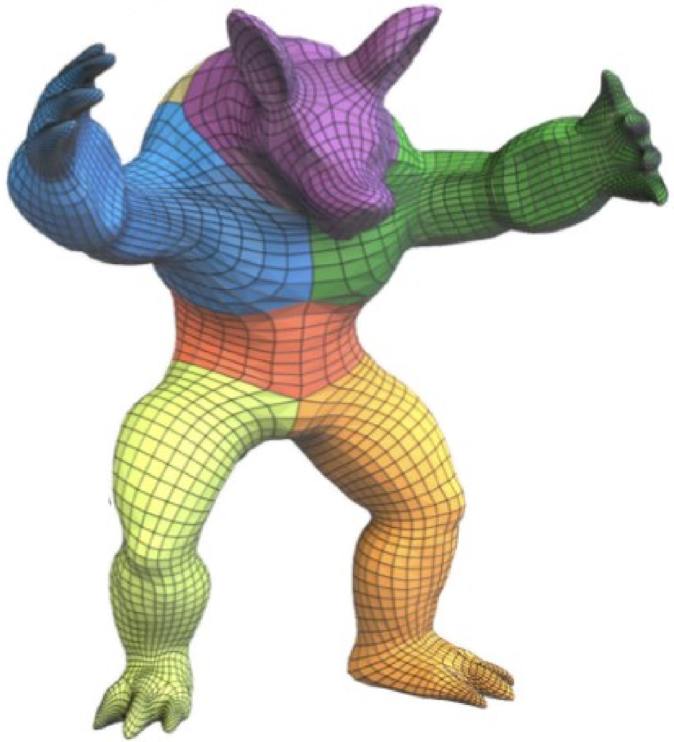

Figure by Pandey et. al., “Face Extrusion Quad Meshes”, 2022

Quad mesh generation using RL

Initial cube

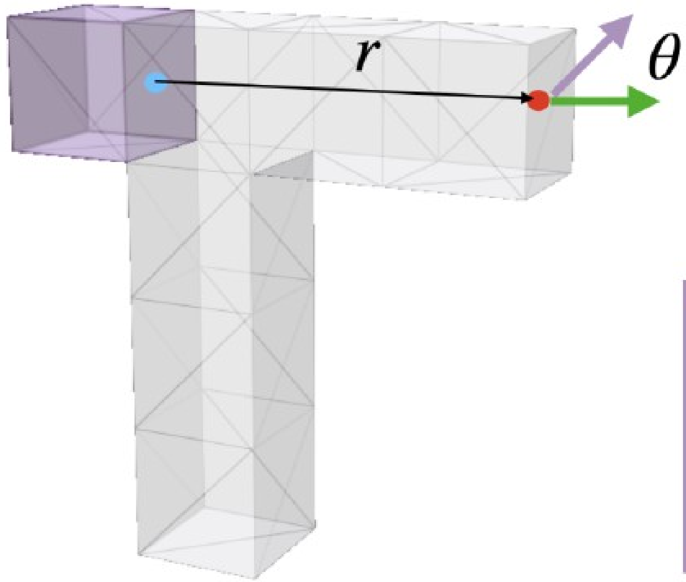

Target mesh

Feature vector: $[\text{Area}(f), \alpha, r, \theta]$

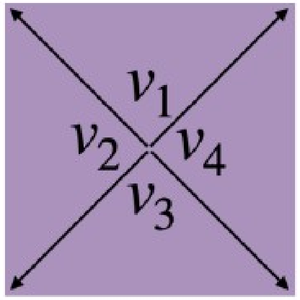

$$\alpha = \sum_{i=1}^{4} v_i$$

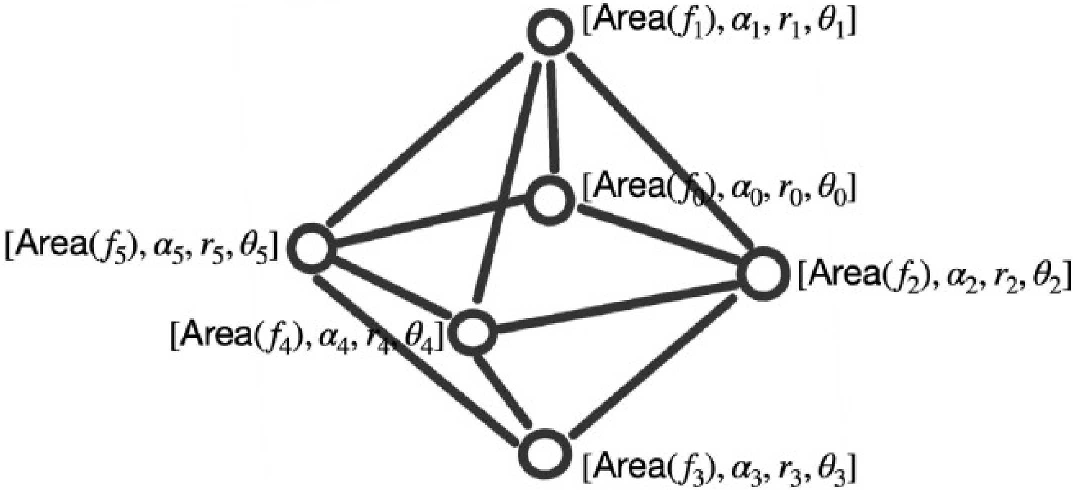

44

Using Deep Generative Models for Atomic Structure Prediction of Metal Oxide Nanoparticles from X-ray Sc

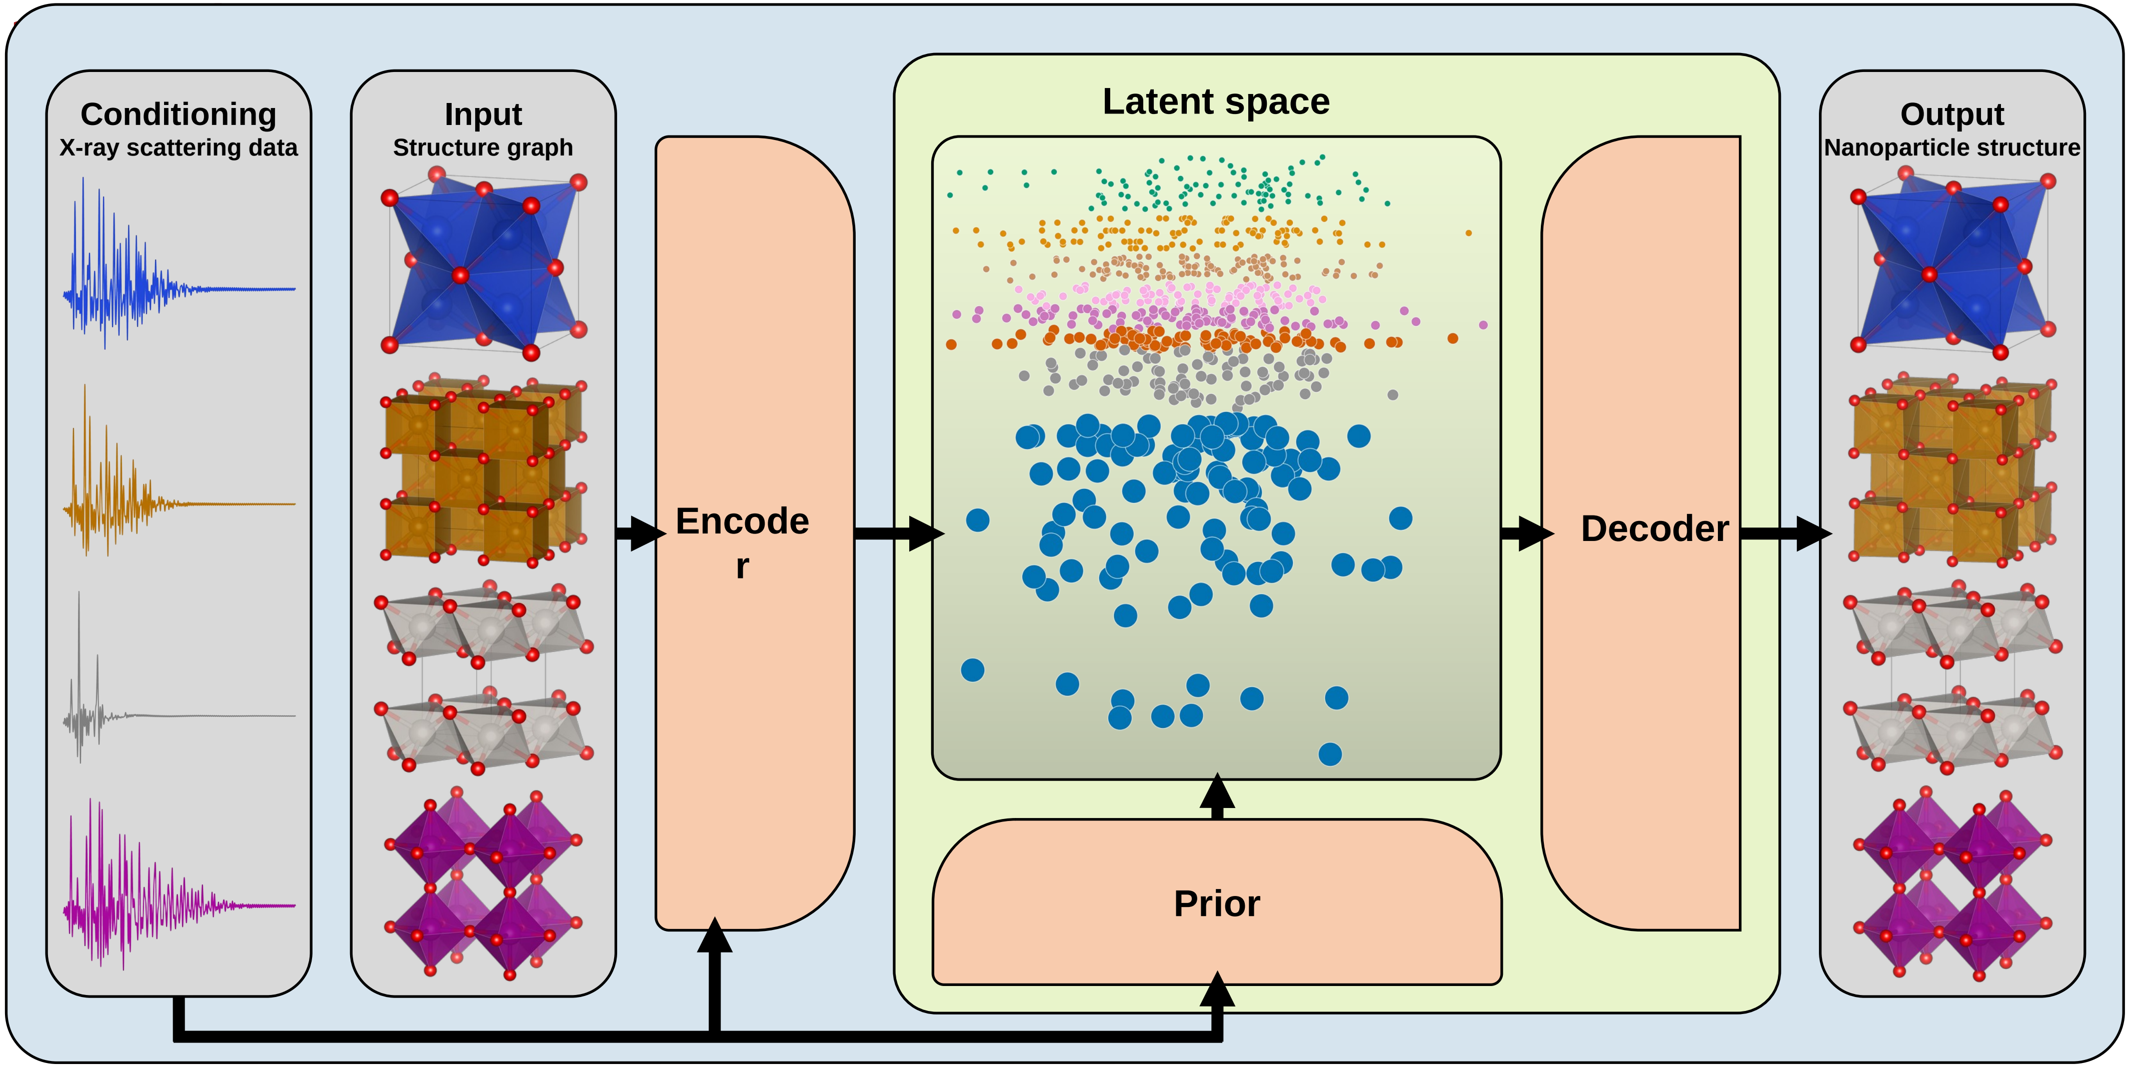

45

Nacala-Roof-Material: Drone Imagery for Roof Detection, Classification, and Segmentation to Support Mosquito-borne Disease Risk Assessment

Venkanna Babu Guthula, Stefan Oehmcke, Remigio Chilaule, Hui Zhang, Nico Lang, Ankit Kariryaa, Johan Mottelson, Christian Igel

UNIVERSITY OF COPENHAGEN

> Machine Learning

UNIVERSITY OF COPENHAGEN

novo nordisk foundation

Data

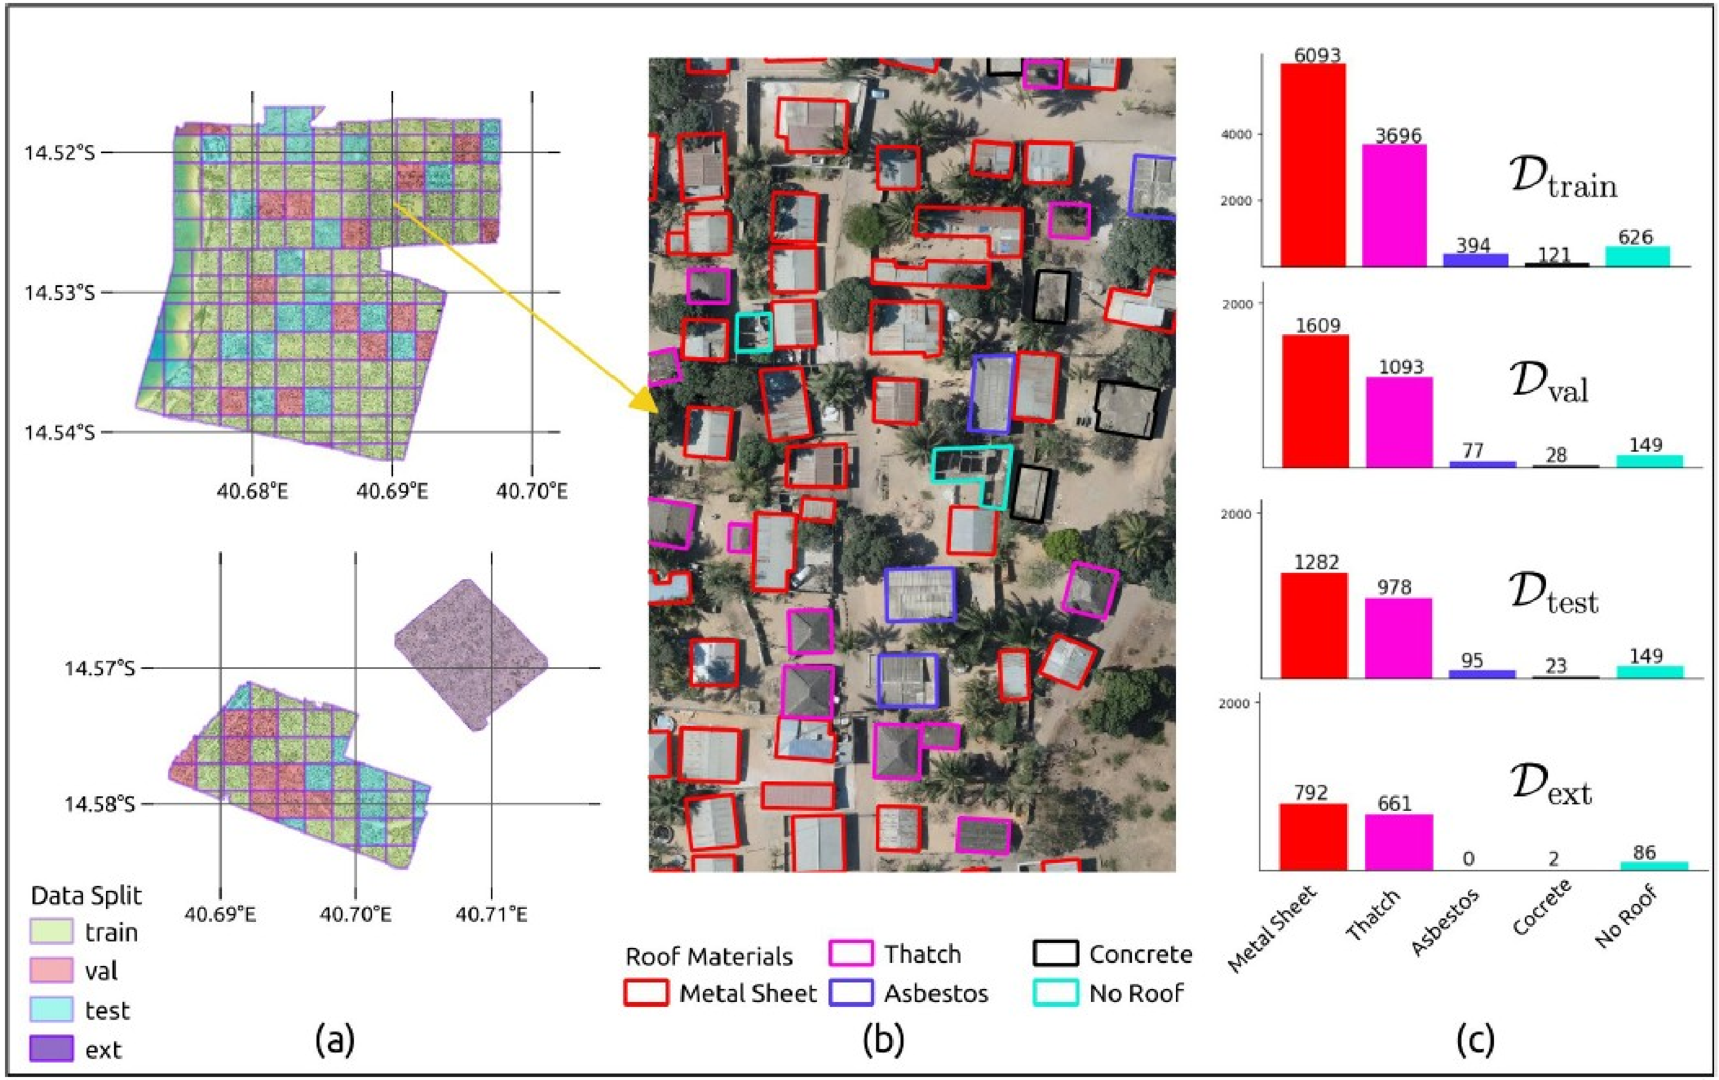

Figure-1: Nacala-Roof-Material data

Results

Model Name	$\mathcal{D}_{test}$						$\mathcal{D}_{ext}$			
	pixel level			object level			pixel level		object level	
	IoU	$mIoU^3$	$mIoU^5$	AP_{50}	mAP^3_{50}	mAP^5_{50}	IoU	$mIoU^3$	AP_{50}	mAP^3_{50}
YOLOv8	0.866 ± 0.012	0.713 ± 0.019	0.568 ± 0.015	**0.941** ± 0.003	0.815 ± 0.011	0.698 ± 0.018	0.896 ± 0.002	0.761 ± 0.006	**0.963** ± 0.005	0.846 ± 0.008
DINOv2	0.833 ± 0.002	0.755 ± 0.004	0.562 ± 0.003	0.882 ± 0.004	0.789 ± 0.006	0.683 ± 0.008	0.905 ± 0.000	0.747 ± 0.011	0.919 ± 0.005	0.806 ± 0.008
$DINOv2_{DOW}$	0.884 ± 0.001	0.763 ± 0.002	0.565 ± 0.004	0.930 ± 0.005	0.836 ± 0.002	0.725 ± 0.004	0.905 ± 0.001	0.852 ± 0.007	0.956 ± 0.001	0.852 ± 0.007
U-Net	**0.895** ± 0.003	0.757 ± 0.024	0.570 ± 0.016	0.910 ± 0.005	0.810 ± 0.008	0.688 ± 0.014	0.909 ± 0.001	0.748 ± 0.007	0.929 ± 0.000	0.787 ± 0.011
$U\text{-}Net_{DOW}$	**0.895** ± 0.002	0.775 ± 0.013	0.577 ± 0.009	0.935 ± 0.001	0.836 ± 0.005	**0.730** ± 0.011	**0.911** ± 0.002	0.764 ± 0.006	0.947 ± 0.004	0.812 ± 0.008
$YOLOv8_{Multi}$	0.824 ± 0.023	0.708 ± 0.010	0.550 ± 0.017	0.910 ± 0.005	0.816 ± 0.009	0.597 ± 0.007	0.885 ± 0.002	0.785 ± 0.006	0.948 ± 0.003	0.849 ± 0.015
$DINOv2_{Multi}$	0.880 ± 0.002	0.774 ± 0.004	0.699 ± 0.012	0.899 ± 0.003	0.820 ± 0.010	0.689 ± 0.025	0.899 ± 0.002	0.818 ± 0.005	0.946 ± 0.001	**0.880** ± 0.011
$DINOv2_{DOW\text{-}Multi}$	0.885 ± 0.001	**0.786** ± 0.006	**0.734** ± 0.006	0.918 ± 0.003	0.824 ± 0.011	0.702 ± 0.013	0.902 ± 0.001	**0.819** ± 0.006	0.950 ± 0.005	0.875 ± 0.010
$U\text{-}Net_{Multi}$	0.879 ± 0.012	0.783 ± 0.010	0.634 ± 0.024	0.924 ± 0.004	**0.850** ± 0.011	0.716 ± 0.018	0.903 ± 0.002	0.805 ± 0.020	0.943 ± 0.010	0.844 ± 0.039
$U\text{-}Net_{DOW\text{-}Multi}$	0.892 ± 0.001	0.777 ± 0.012	0.672 ± 0.042	0.928 ± 0.002	0.829 ± 0.011	0.671 ± 0.022	0.904 ± 0.002	0.794 ± 0.014	0.942 ± 0.005	0.812 ± 0.021

Table-1: Results on Nacala Roof Material data set

46

Masked Autoencoders for Hyperspectral Imaging

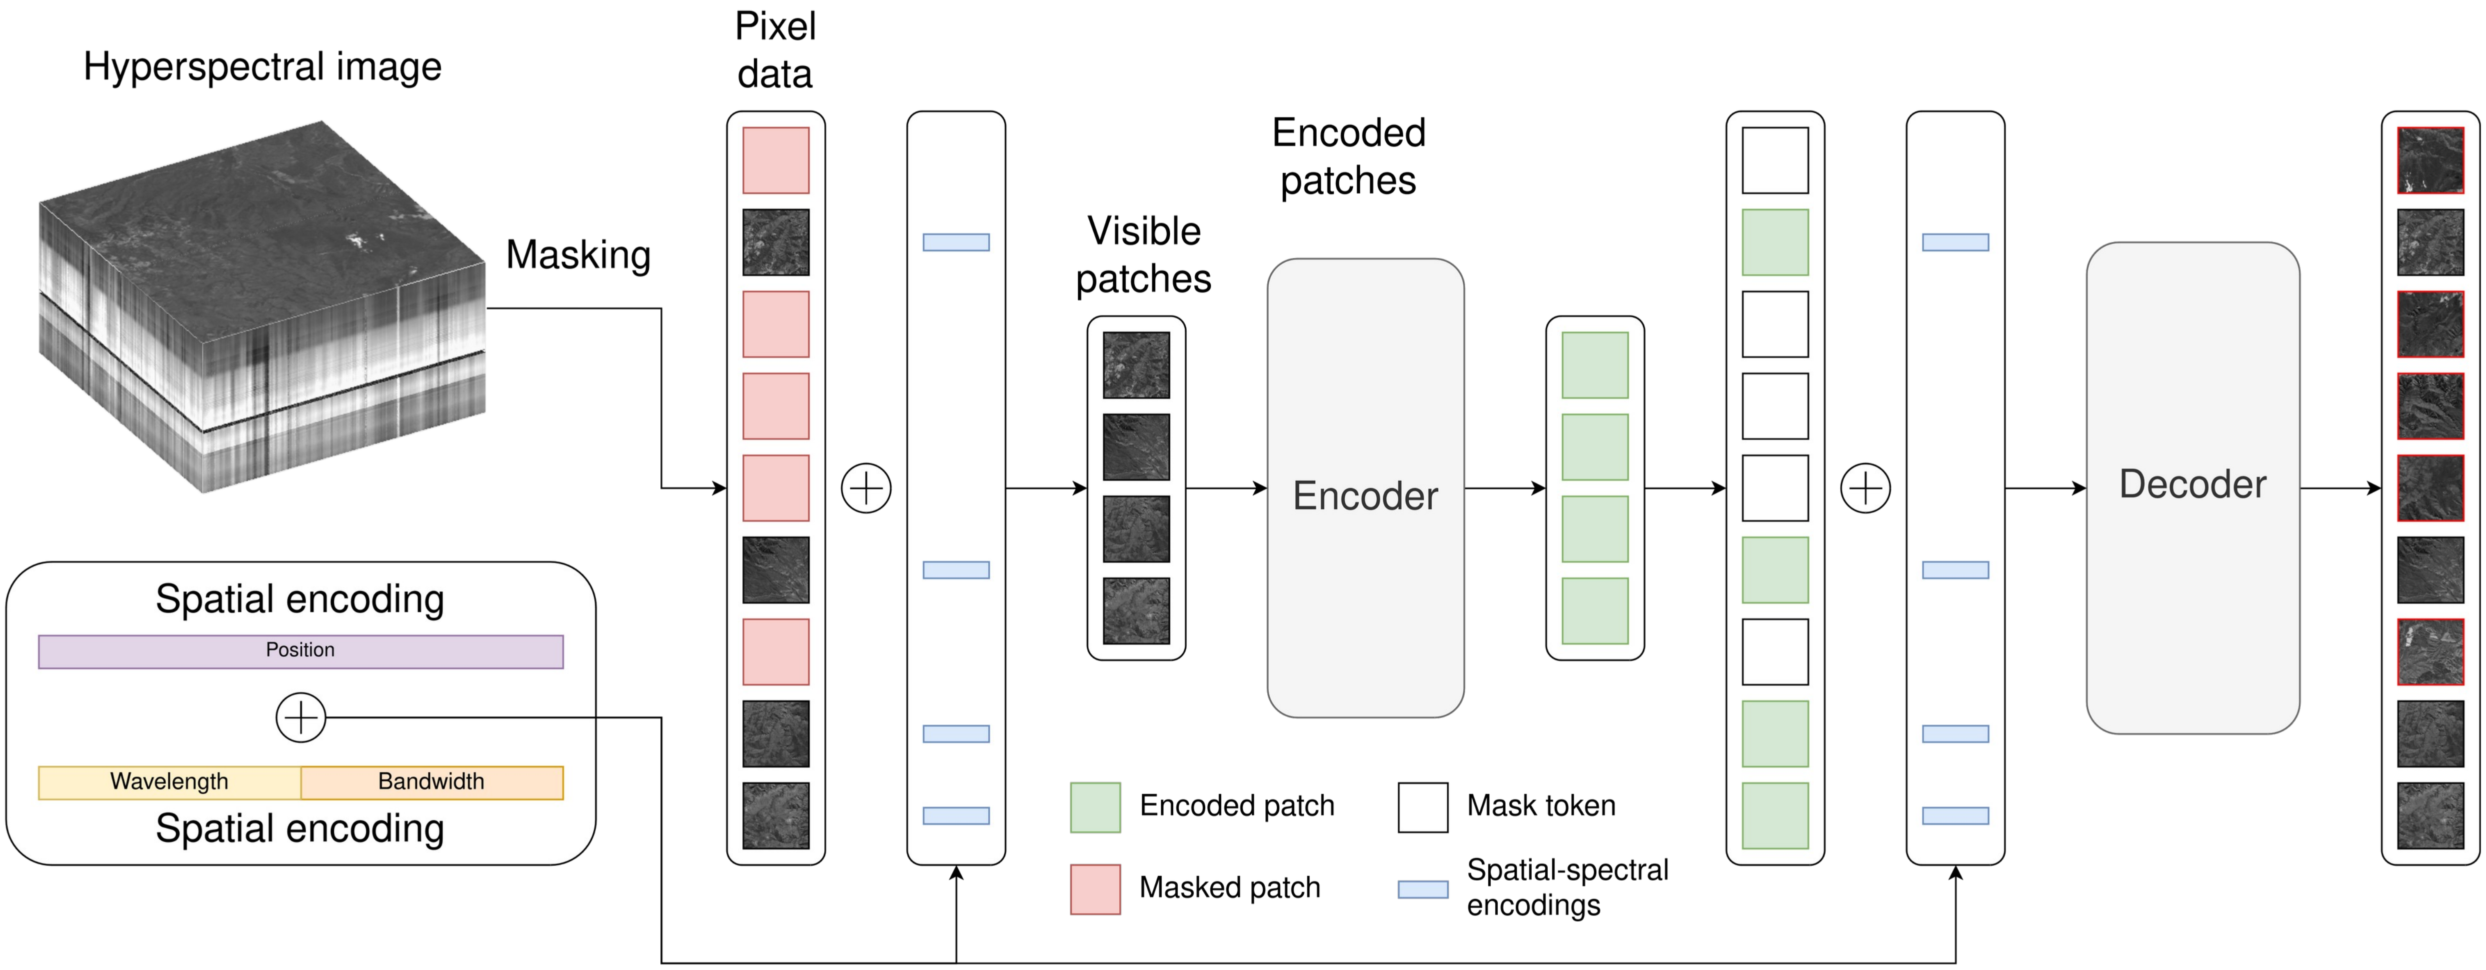

47

Combining Physics and Deep Learning: A New Framework for Image Denoising

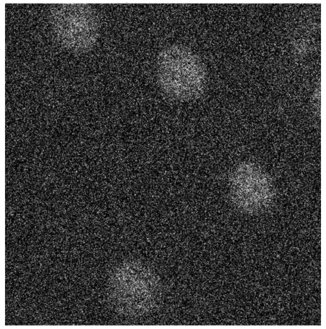

Raw image

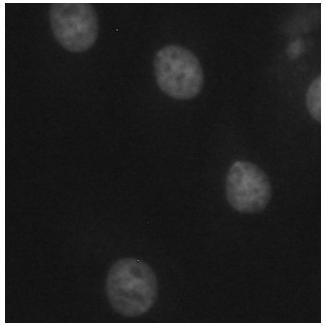

GT image

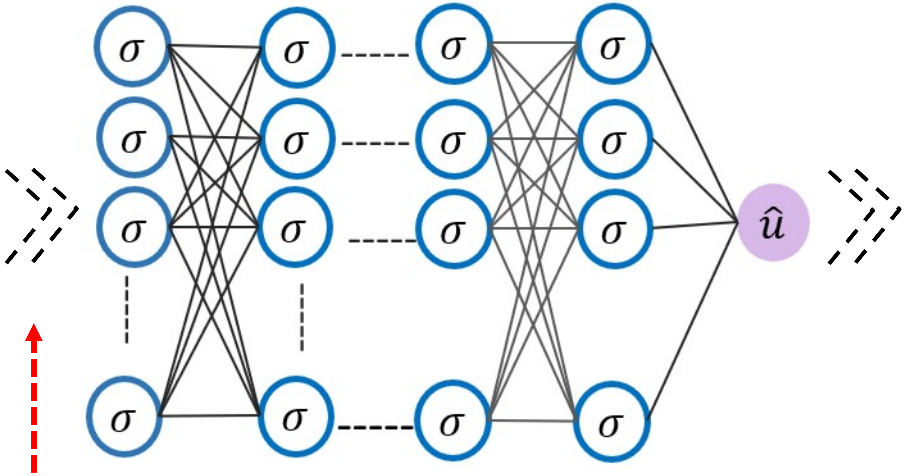

$$\text{MSE} = \overset{\text{Mean}}{\boxed{\frac{1}{n}\sum_{i=1}^{n}}}\overset{\text{Error}}{\boxed{(Y_i - \hat{Y}_i)}}^{\overset{\text{Squared}}{\boxed{2}}}$$

Combining Physics and Deep Learning: A New Framework for Image Denoising

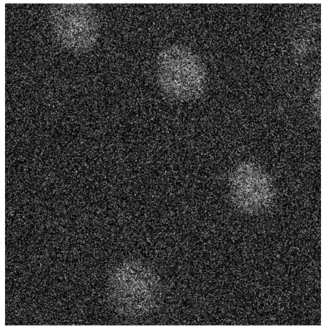

Raw image

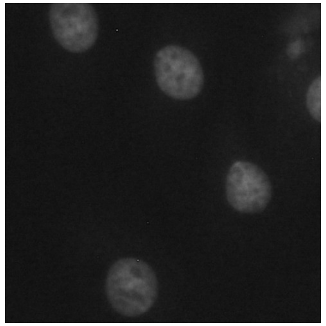

GT image

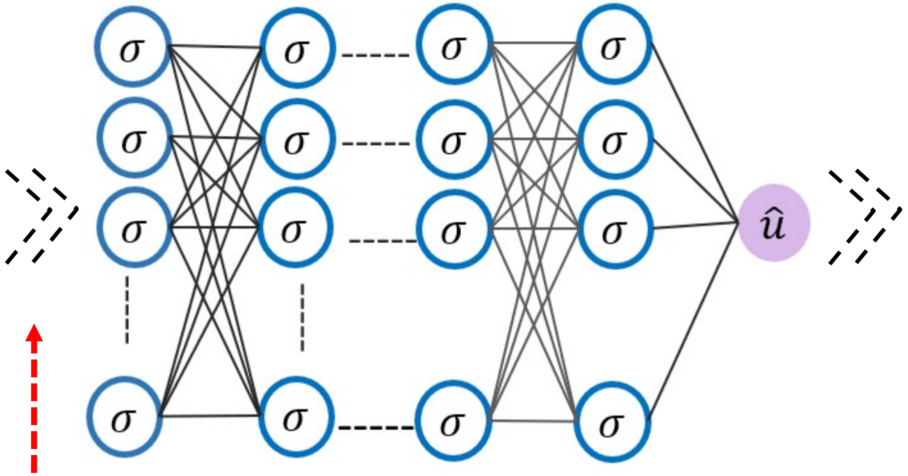

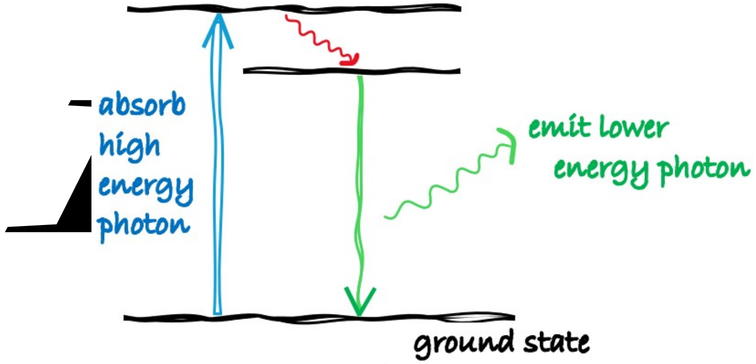

Physics of Fluorescence Microscopy

48

Fetal Ultrasound Image Segmentation for Measuring Biometric Parameters Using Multi-Task Deep Learning

Zahra Sobhaninia[1], Shima Rafiei[2], Ali Emami[3], Nader Karimi[3], Kayvan Najarian[4], Shadrokh Samavi[4], S.M.Reza Soroushmehr[4]

1. Pioneer Center for AI, Department of Computer Science, University of Copenhagen, Denmark

2. Department of Computer Science, McMaster university, Canada

3. Department of Computer Engineering, Isfahan University of Technology, Iran

4. Electrical and Computer Engineering, Seattle University, USA