Image Processing with Artificial Intelligence

Eleftherios Trivizakis, PhD

Computational BioMedicine Laboratory (CBML)
Institute of Computer Science (ICS)
Foundation for Research and Technology – Hellas (FORTH)

Contact: trivizakis@hmu.gr

Course Outline

- Segmentation & Localization in Medical Imaging
- Segmentation Uncertainty Methods
- Interactive Segmentation
- Label-free Segmentation
- Denoising

Segmentation & Localization in Medical Imaging

While classification tell us ***if a pathology exist***, object detection tell us ***where***

- Localization is critical for:
  - Surgical planning
  - Longitudinal tracking of tumor growth
  - Feature Extraction
  - Lesion Measuring

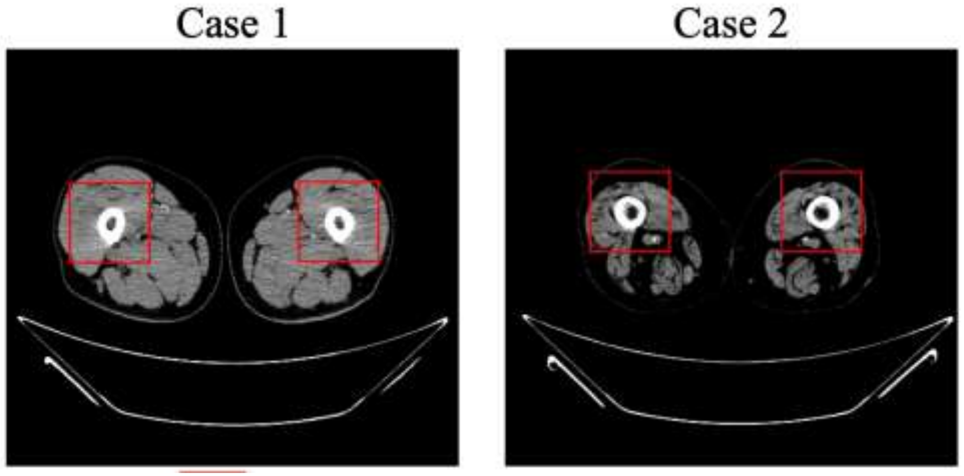

Localization of the left and right femur through Deep Learning

Koutoulakis, E., Trivizakis, E., Koutoulidis, V., Moulopoulos, L. A., Terpos, E., Ntanasis-Stathopoulos, I., Malandrakis, P., Grigoropoulos, P., Papadopoulos, P., Nikiforaki, K., Papanikolaou, N., Fotiadis, D. I., & Marias, K. (2023). Fully Automated Detection and Segmentation Pipeline for the Bone Marrow of the Lytic Bone of Multiple Myeloma Patients. *2023 IEEE EMBS Special Topic Conference on Data Science and Engineering in Healthcare, Medicine and Biology*, 39–40. https://doi.org/10.1109/IEEECONF58974.2023.10405013

Object Detection in HealthCare

Addressing the **"Search-and-Find"** bottleneck in high-volume clinical workflows

- Manual annotation and hand-crafted feature extraction
- Computer-aided Detection (CADe)
- Modern Deep Learning: Detection + Localization

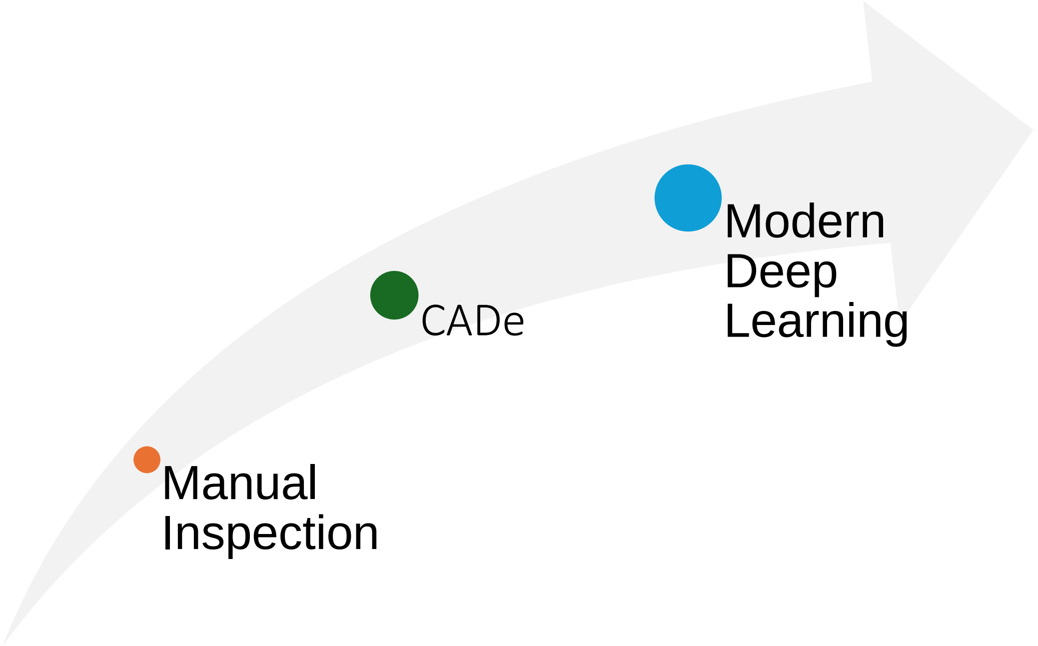

Use cases

- Localization for Attention-Guided Classification
  1. Find the lungs
  2. Find the right lung
  3. Find the left lung

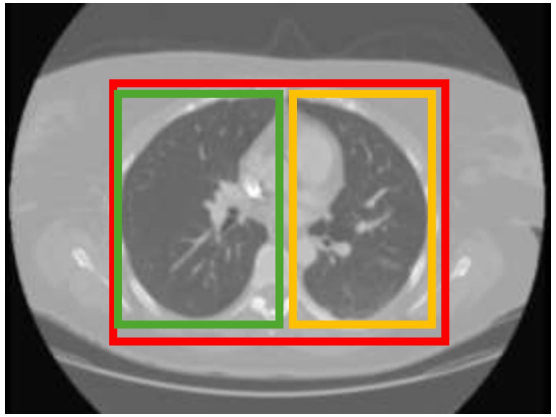

- Localization of organs at risk for Image-Guided Radiation Therapy (IGRT)
  1. Femurs left - right
  2. Bladder
  3. Rectum

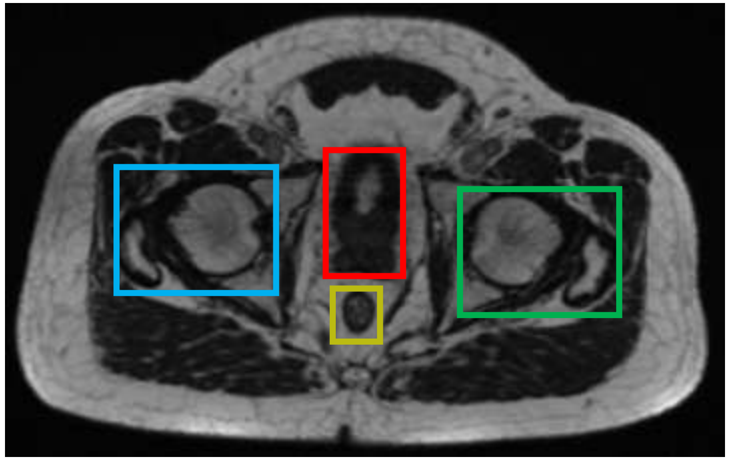

NLST. (n.d.). *The Cancer Imaging Archive (TCIA)*. Retrieved April 3, 2026, from https://www.cancerimagingarchive.net/collection/nlst/

Koutoulakis, E., Marage, L., Markodimitrakis, E., Aubignac, L., Jenny, C., Bessieres, I., & Lalande, A. (2023). Automatic Multiorgan Segmentation in Pelvic Region with Convolutional Neural Networks on 0.35 T MR-Linac Images. *Algorithms*, *16*(11), 521. https://doi.org/10.3390/a16110521

Architectural Taxonomy

Paradigm	Mechanism	Key advantage in healthcare
Two-Stage (Faster R-CNN)	RPN → RoI Pooling	Superior for small lesions (e.g microcalcifications)
One-Stage (YOLOv11)	Dense Prediction	Real-time intraoperative guidance (endoscopy)
Set-Prediction (RT-DETR)	Bipartite Matching	Eliminates heuristic anchor-box design

- The two stage detectors prioritize precision through region proposals newtowks
- One-stage detectors optimize for throughput

Ren, S., He, K., Girshick, R., & Sun, J. (2016). *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks* (arXiv:1506.01497). arXiv. https://doi.org/10.48550/arXiv.1506.01497
Khanam, R., & Hussain, M. (2024). *YOLOv11: An Overview of the Key Architectural Enhancements* (arXiv:2410.17725). arXiv. https://doi.org/10.48550/arXiv.2410.17725
Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., & Chen, J. (2023, April 17). *DETRs Beat YOLOs on Real-time Object Detection*. arXiv.Org. https://arxiv.org/abs/2304.08069v3

Quantitative Evaluation

- **Intersection over Union**: Measures the overlap between the predicted bounding box (B_p) and ground truth (B_{gt})

$$\mathrm{IoU} = \frac{\mathrm{Area}(B_p \cap B_{gt})}{\mathrm{Area}(B_p \cup B_{gt})}$$

- **Free-response Receiver Operating Characteristic (FROC)**: Plots the sensitivity (probability of detection) against the average number of False Positives per image
- **Means Average Precision (mAP)**: Evaluates the model’s performance across all possible confidence thresholds and all classes

$$\mathbf{Sensitivity} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}$$

- **Sensitivity**: quantifies the model's ability to correctly localize pathological regions of interest (ROIs) out of all ground-truth instances

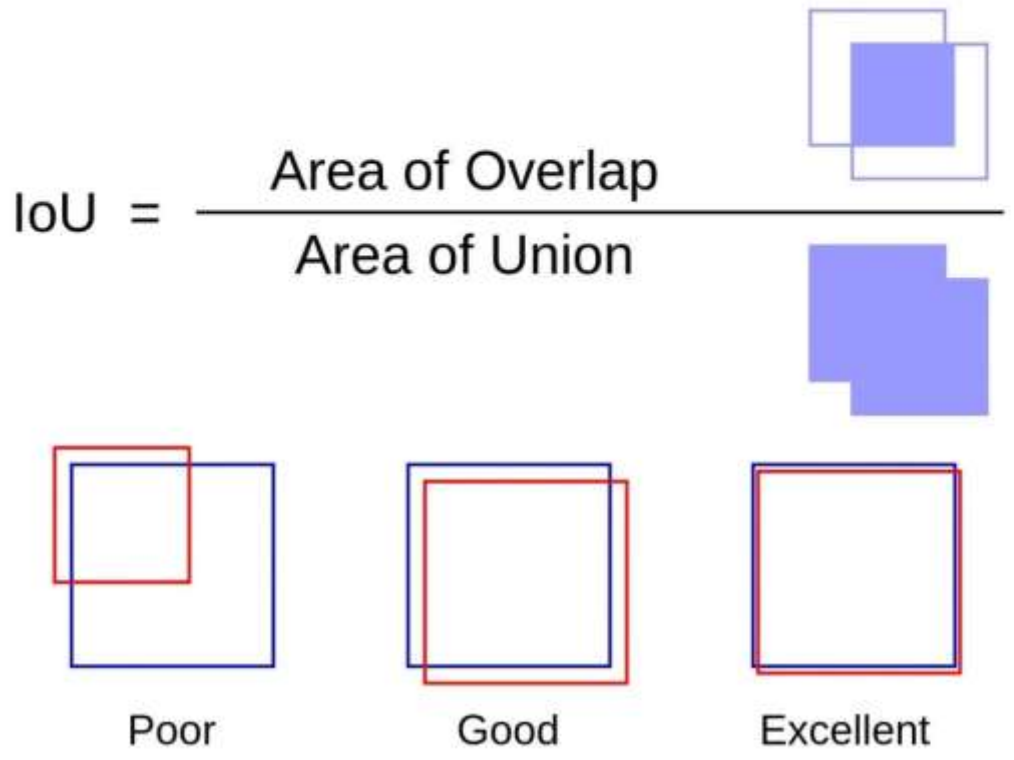

Object Detection with CNNs: R-CNN

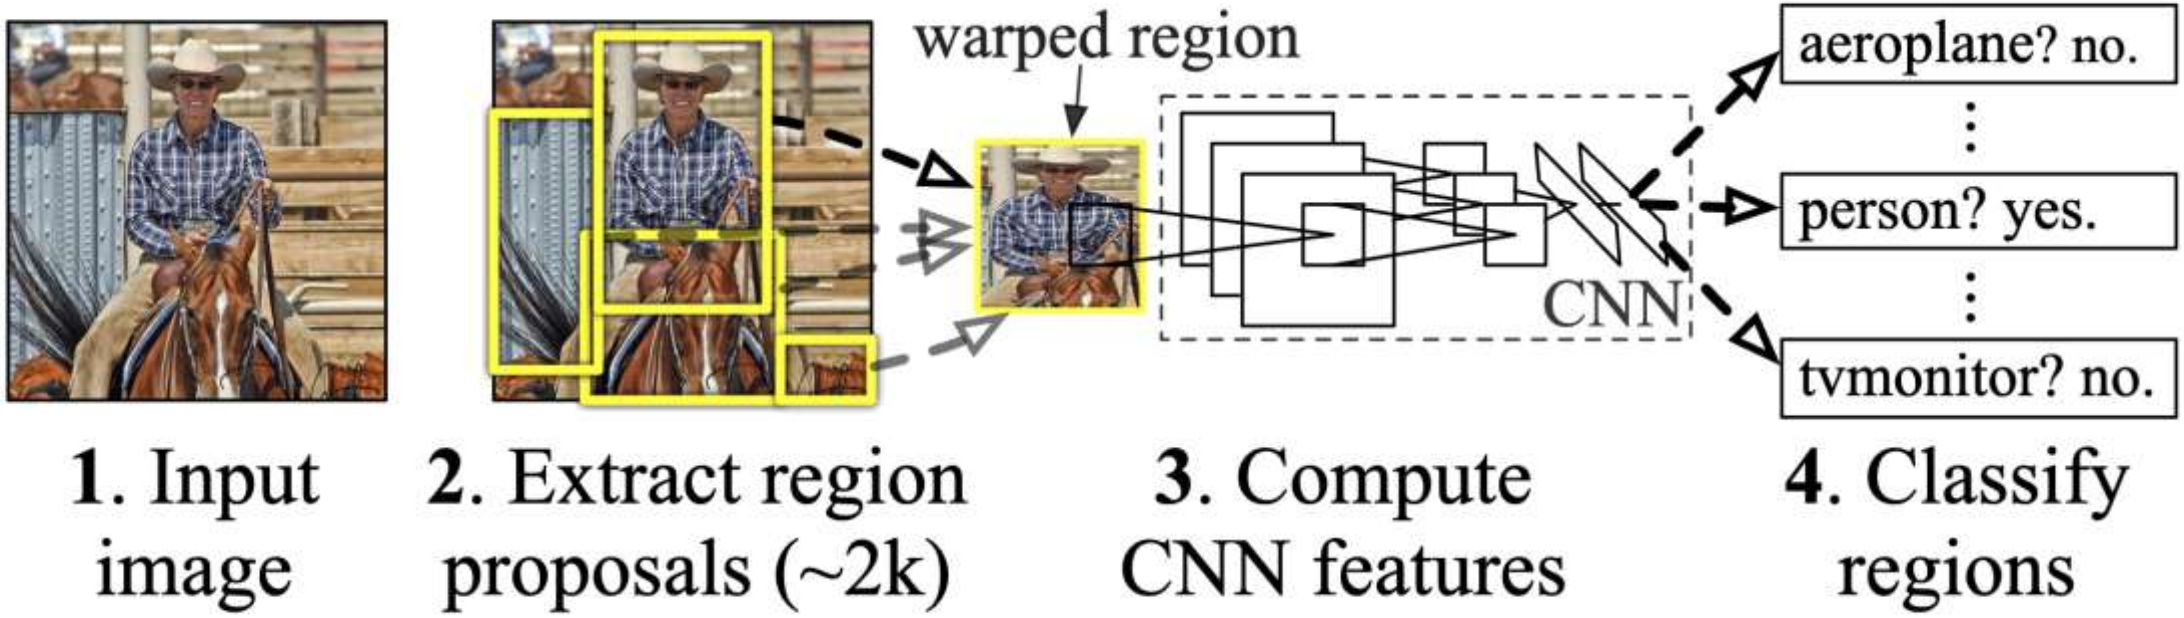

Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). *Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation*. 580–587. https://www.cv-foundation.org/openaccess/content_cvpr_2014/html/Girshick_Rich_Feature_Hierarchies_2014_CVPR_paper.html

Yolo Architecture – Detection Model

- One of the **most popular** detection models with many variations
- **Backbone Blocks**: Spatial feature extraction
- **Attention Module**: Allows the model to prioritize critical spatial regions
- **Head:** A decoupled architecture for classification and regression tasks ensuring more precise **localization**

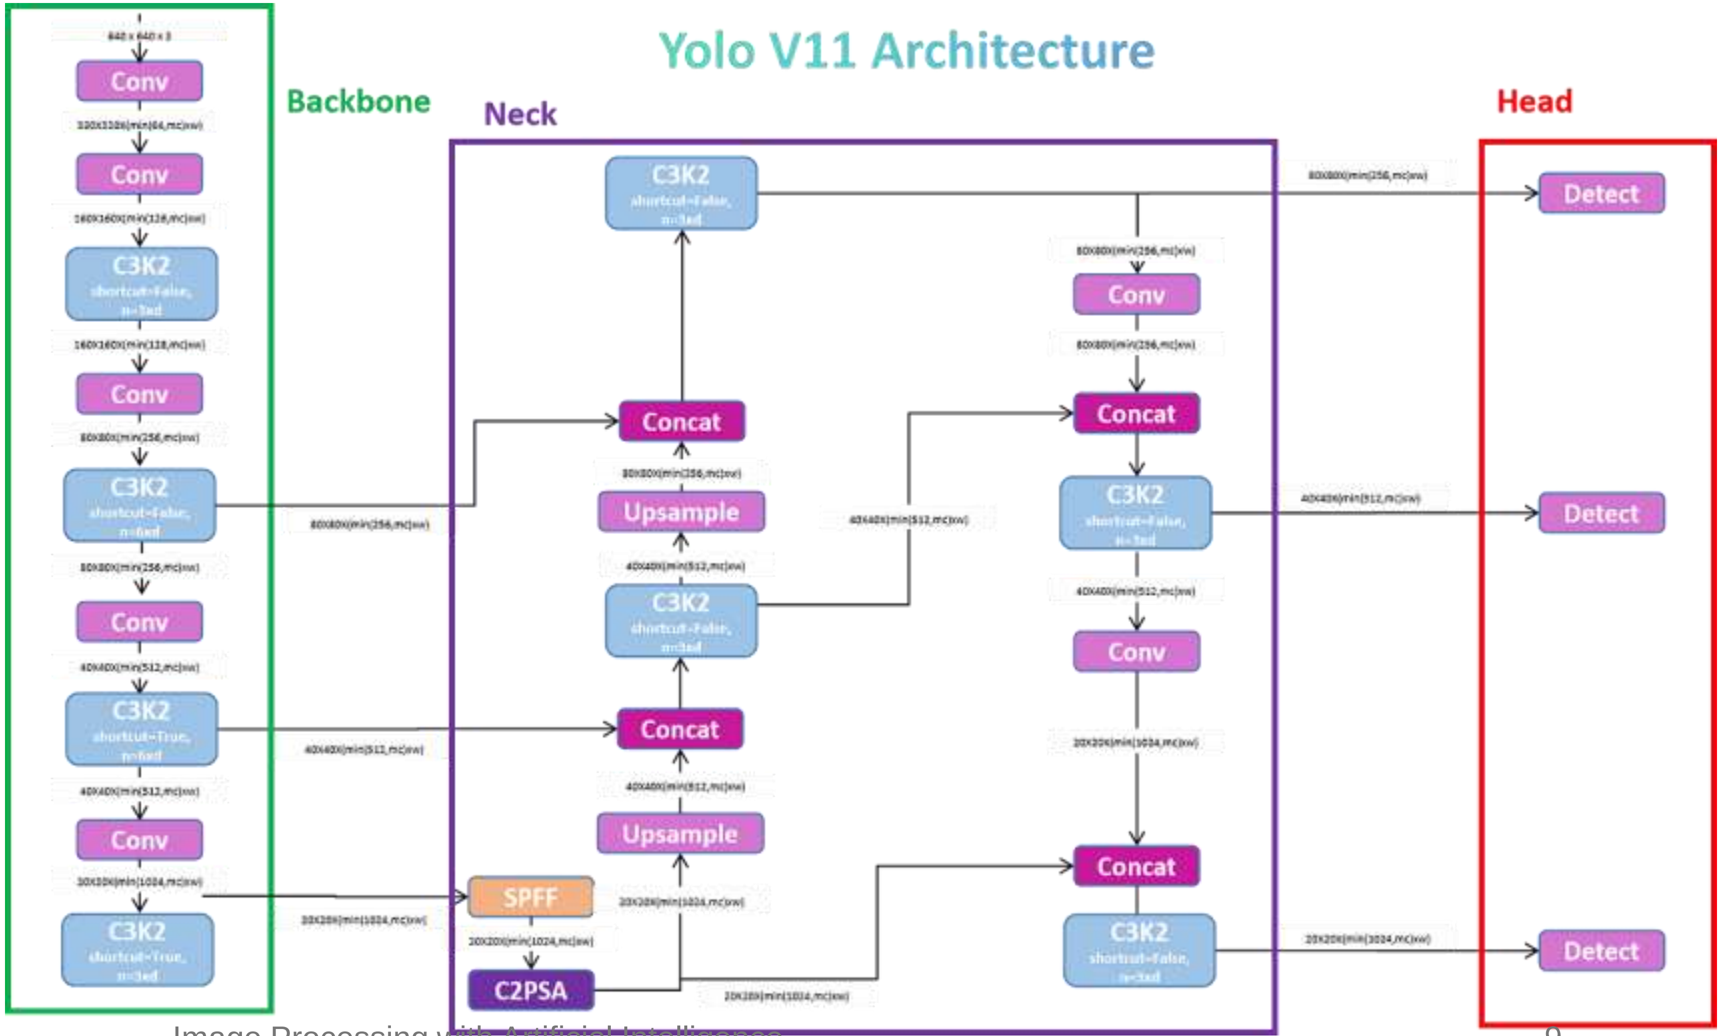

Aorta Segmentation

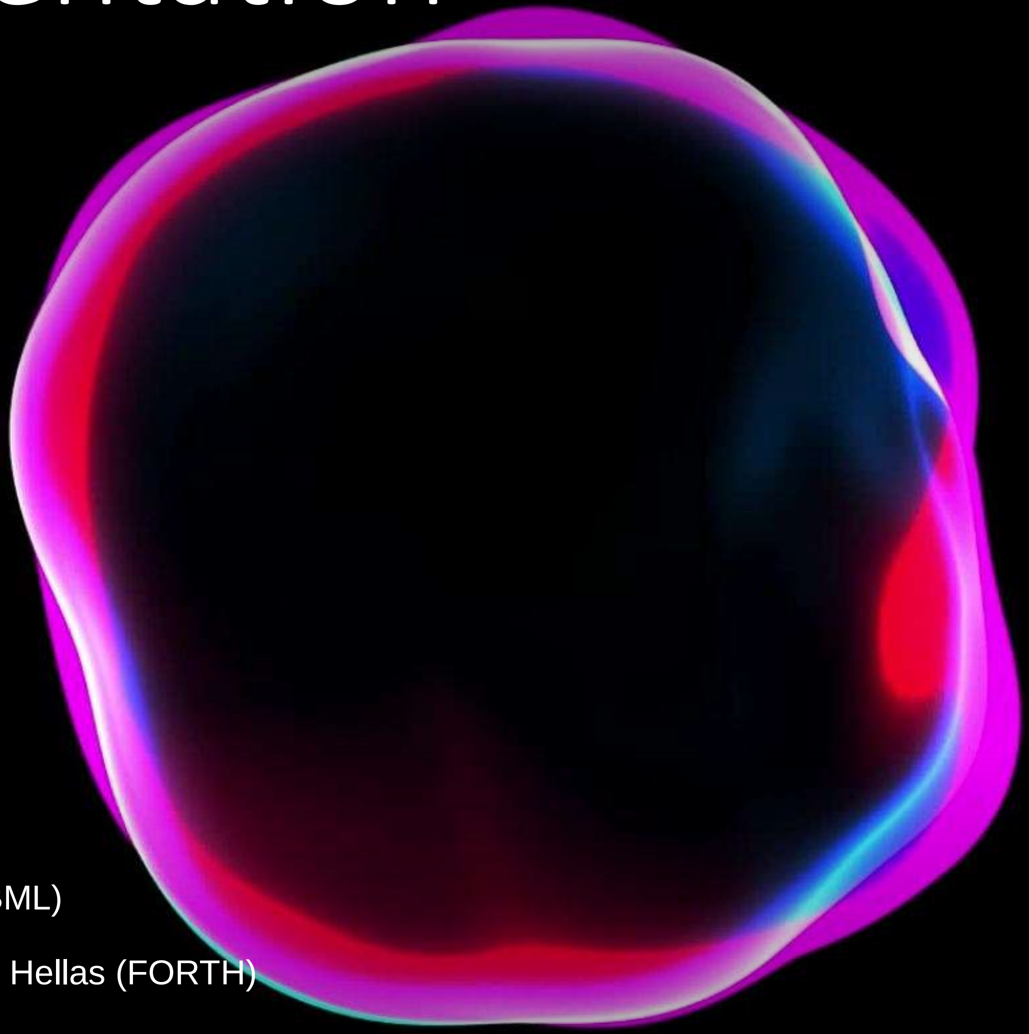

Manos Koutoulakis, PhDc

Computational BioMedicine Laboratory (CBML)
Institute of Computer Science (ICS)
Foundation for Research and Technology – Hellas (FORTH)

Aorta Segmentation

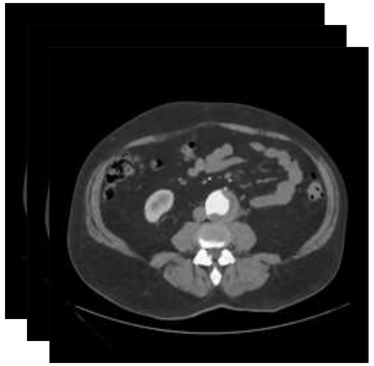

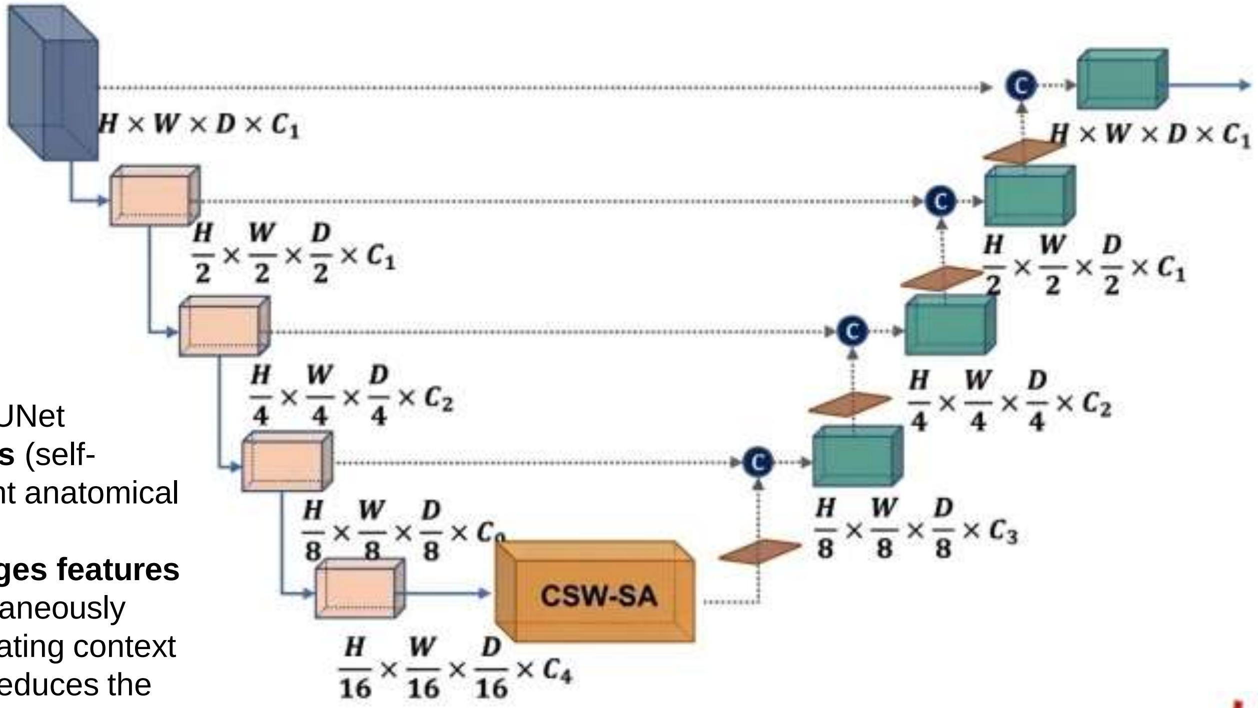

(a) The architecture of CIS-UNet

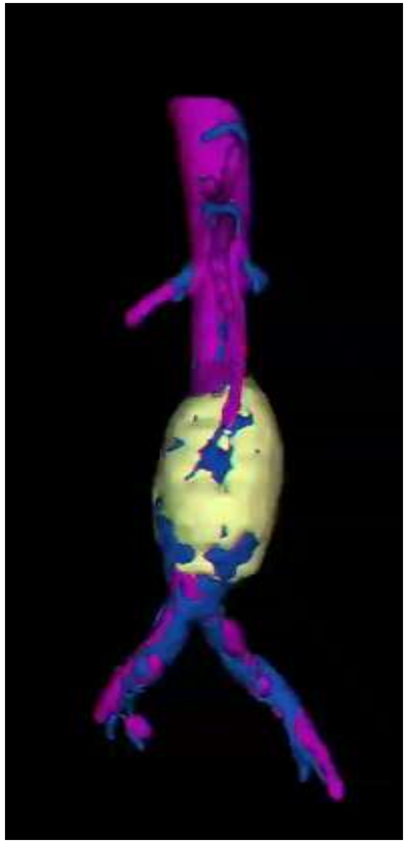

- **Context-Integrated Design:** CIS-UNet integrates **global context modules** (self-attention) to understand how distant anatomical structures relate to one another
- **Multi-Scale Feature Fusion: Merges features** from multiple encoder levels simultaneously
- **Semantic Consistency:** By integrating context earlier in the encoding process, it reduces the "semantic gap" between the encoder and decoder, leading to **more precise boundaries** in segmented tumors or organs.

sAVe AortAI

Results – Multi-Class Aorta Segmentaiton

Anatomical Regions	Dice*	Dice (σ)	IoU*	IoU (σ)
Lumen	0.9060	±0.035	0.8299	±0.056
Thrombus	0.8092	±0.155	0.7009	±0.173
Calcifications	0.3386	±0.134	0.2114	±0.099

* ↑ = Better

Expert Segmentation

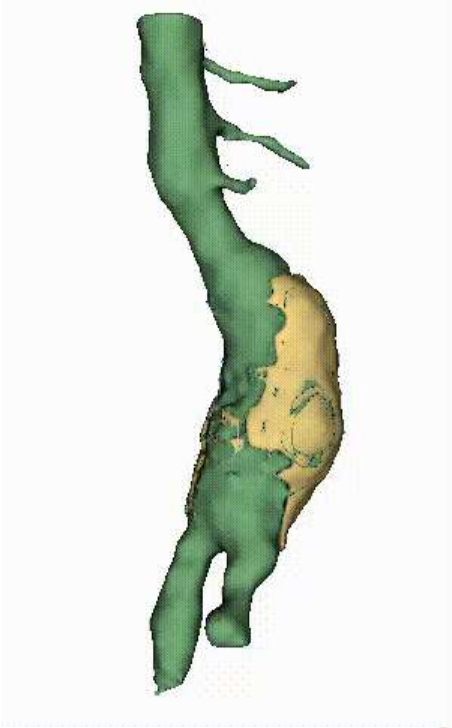

AI Segmentation

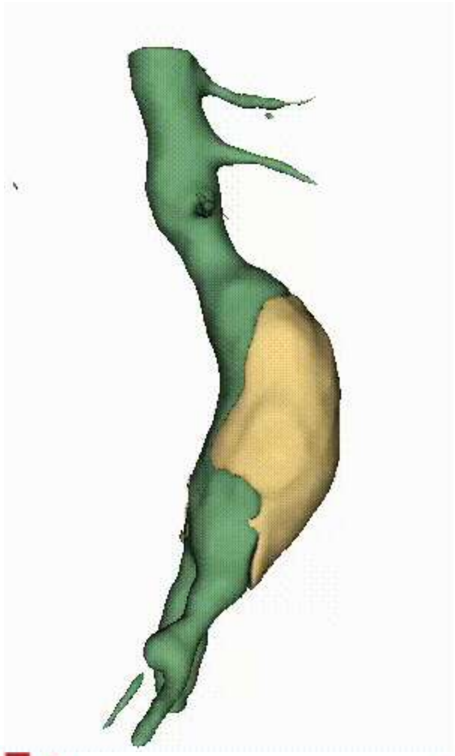

Image Segmentation – Uncertainty

Medical image segmentation is not just about producing a "mask," but understanding the **confidence behind every pixel prediction**. Uncertainty quantification identifies regions of low clinical confidence, highlighting areas that **require expert human review**.

- **Sources of Uncertainty:**
  - **Data:** Noise in the scan (MRI/CT), subtle contrasts, blurry tumor boundaries, or poor image quality.
  - **Model:** Lack of sufficient training data (low variability in patient cohort) or the model's architectural limitations in handling complex anatomy.
- **Confidence Maps**
  - Moving from binary segmentation (Malignant/Benign) to probabilistic segmentation. The model outputs a continuous "Confidence Score" for each pixel, from 0.0 to 1.0.
- **Visualization Methods**
  - **Uncertainty Heatmap,** overlaid on the segmentation, highlights the areas where the model’s predictions are the most volatile and likely to be incorrect.

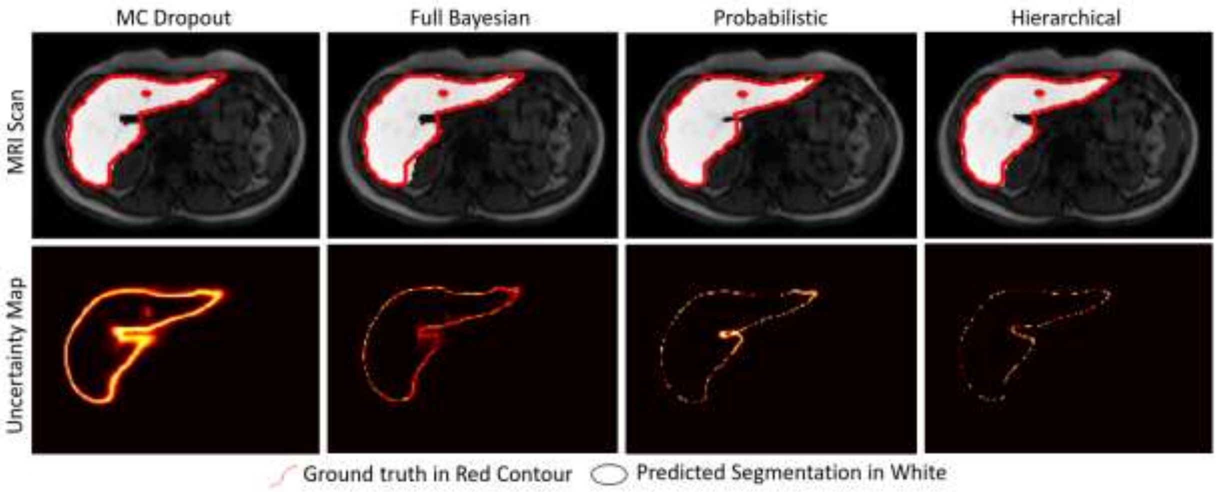

Senapati, J., Roy, A. G., Pölsterl, S., Gutmann, D., Gatidis, S., Schlett, C., ... & Wachinger, C. (2020, September). Bayesian neural networks for uncertainty estimation of imaging biomarkers. In *International Workshop on Machine Learning in Medical Imaging* (pp. 270-280). Cham: Springer International Publishing.

Interactive segmentation

Stathis Kyriazis, PhDc

Computational BioMedicine Laboratory (CBML)
Institute of Computer Science (ICS)
Foundation for Research and Technology – Hellas (FORTH)

Interactive vs. Fully Automatic Segmentation

Fully Automatic

Standard deep learning approach where the network operates independently

- **Domain-Specific:** Predefined classes seen during training
- **Passive output:** Errors require tedious, manual correction
- **Rigidity:** Struggles with out-of-distribution or rare clinical cases
- **Throughput:** Optimized for high-volume, where speed is important

Interactive

User guides the network in real-time to get the final segmentation

- **Class agnostic:** Focuses on boundaries and regions not just static labels.
- **Human-in-the-Loop:** Interactions guide the attention of the network
- **Generalizable:** Capable of segmenting targets not present in the training set
- **Clinical Accuracy:** User guides the network to a reliable result

nnInteractive: 3D interactive segmentation framework

- **nnU-Net Core:** Native 3D fully convolutional architecture
- **Massive Pre-training:** Trained on 120+ public datasets
- **Universal Modality:** Validated across CT, MRI, Ultrasound, PET, and 3D Microscopy
- **Multi-Modal Prompting:** Supports Points, Bounding Boxes, Lassos, and Scribbles
- **True 3D Context:** Native 3D input/output. Interaction placed on any plane (Axial, Sagittal, Coronal) of the entire volume.
- **Clinical-Grade Accuracy:** Outperforms existing methods in "Accuracy-per-Click"

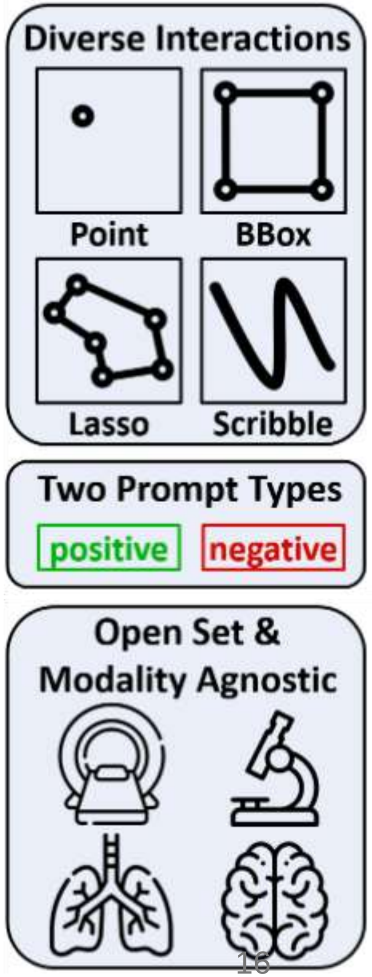

nnInteractive: Segmentation loop

- **Early prompting strategy:** User interactions are encoded as additional guidance channels alongside the raw image.
- **Volumetric Input:** The network processes the 3D volume and interaction maps simultaneously.
- **3D Output:** Generates a consistent 3D binary mask in a single forward pass.
- **Iterative Refinement:** Current output mask is fed back into the input
- **Next Interaction:** User provides corrective a interaction to recursively "fine-tune" the result.

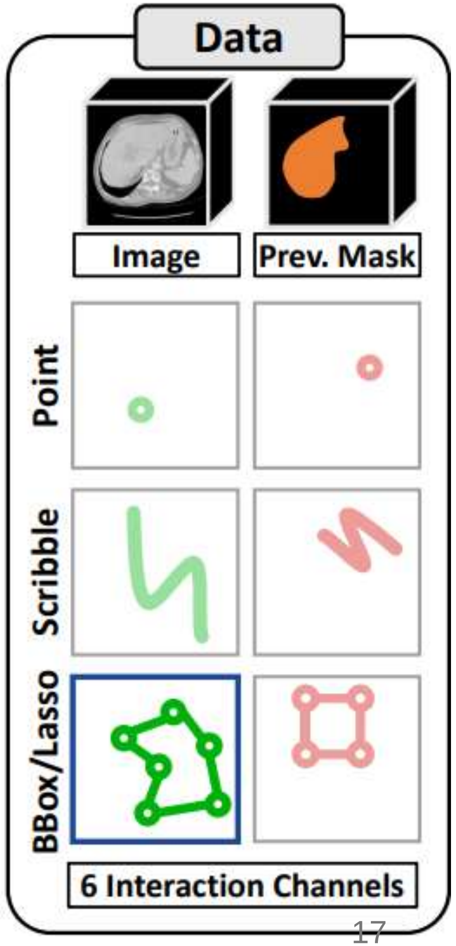

nnInteractive: Segmentation loop

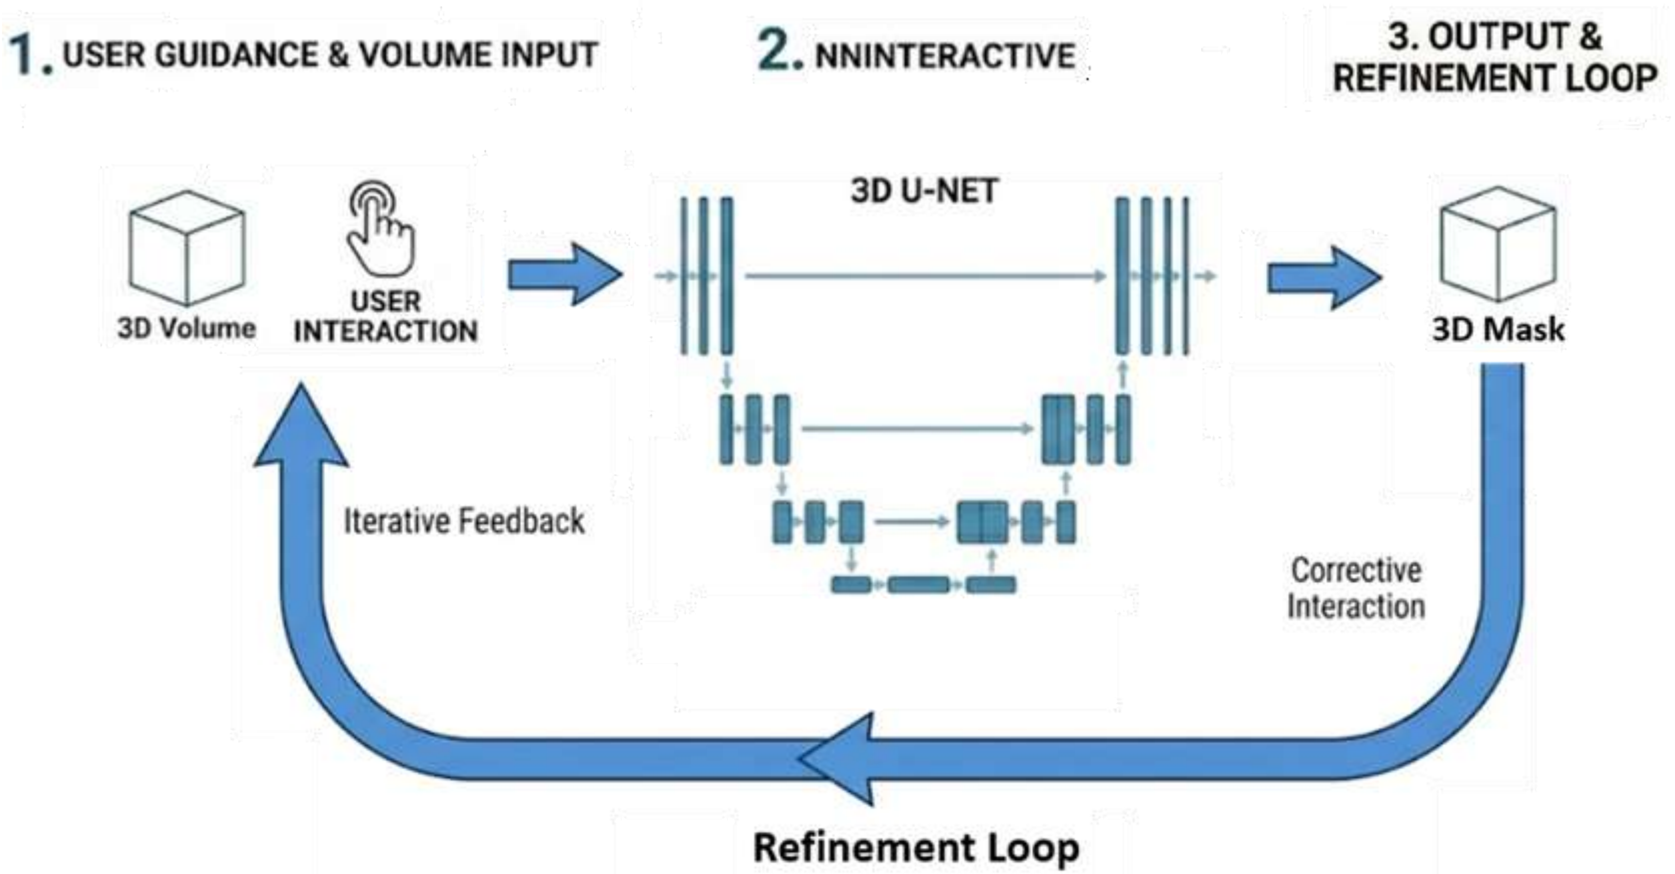

nnInteractive: Demo

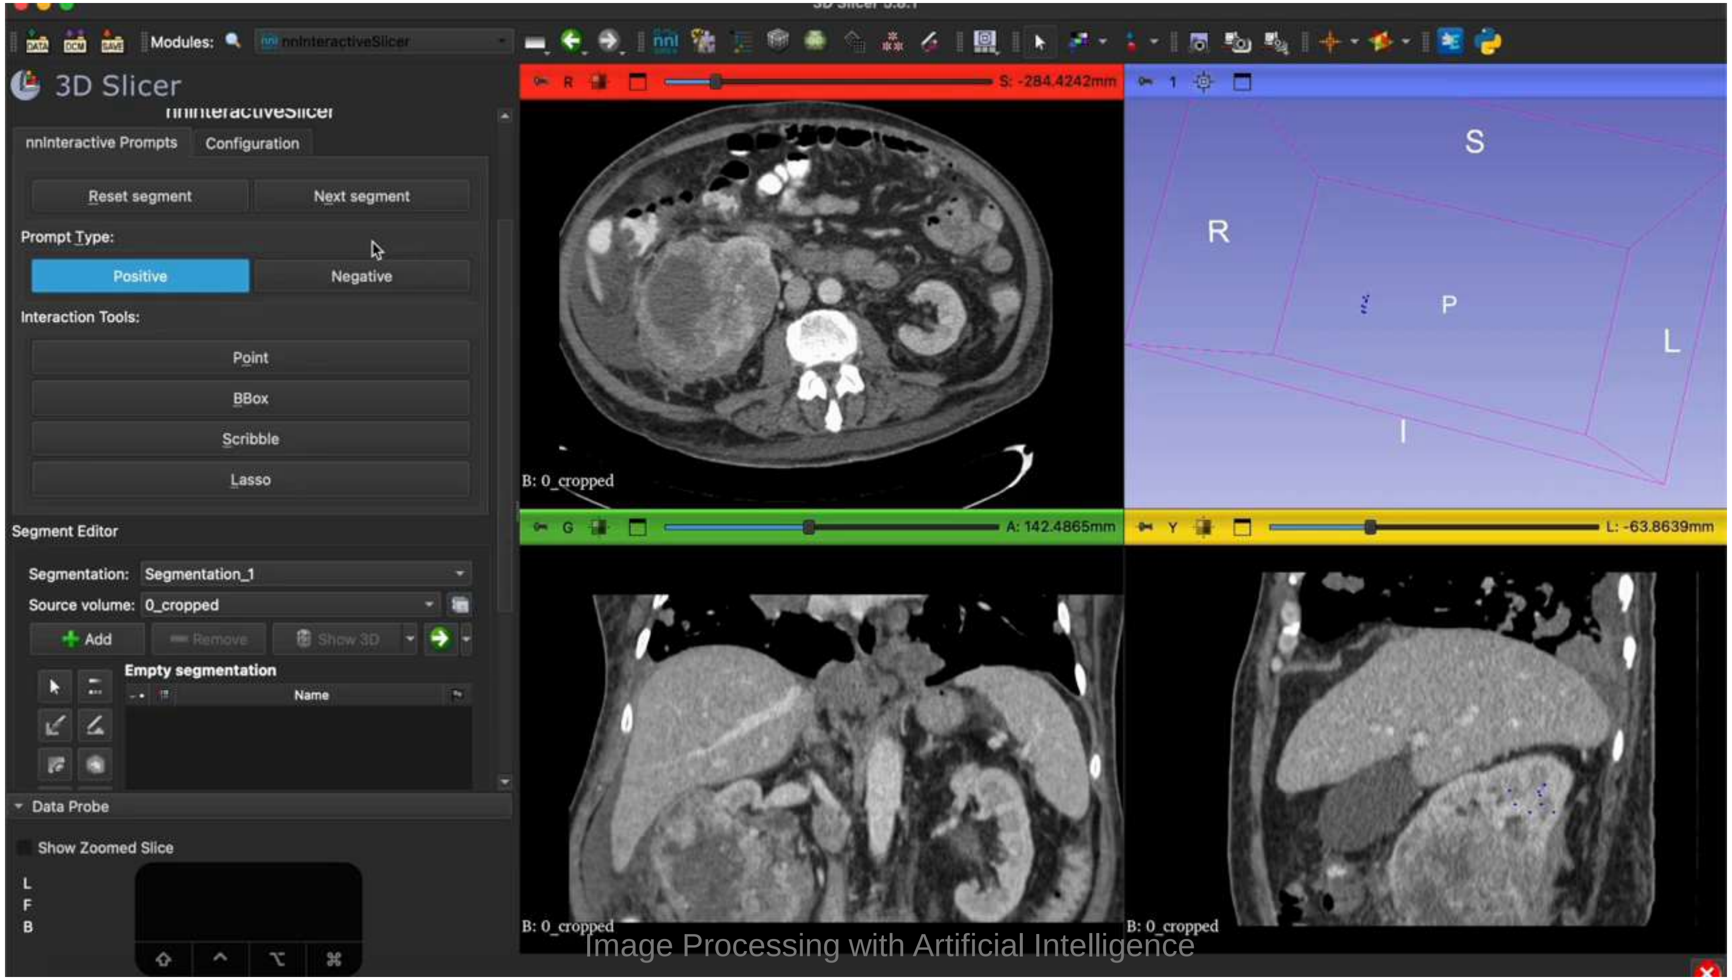

SAM-based interactive segmentation models

- **Natural Image Origin:** Adapted from the Segment Anything Model (SAM) designed for general-purpose natural image interactive segmentation
- **Transformer-Based:** Built on a Vision Transformer (ViT) backbone
- **Medical Adaptations:** Specialized frameworks fine-tuned to bridge the domain gap between natural and medical images. (e.g., MedSAM, SAM-Med2D)
- **Zero-Shot Potential:** High "out-of-the-box" performance on 2D slices for diverse anatomical structures.
- **The 2D vs. 3D Trade-off:** Require additional modules to support 3D volumetric segmentation

Label-free Segmentation

When labels and annotations are scarce!

DL-based Bone Segmentation

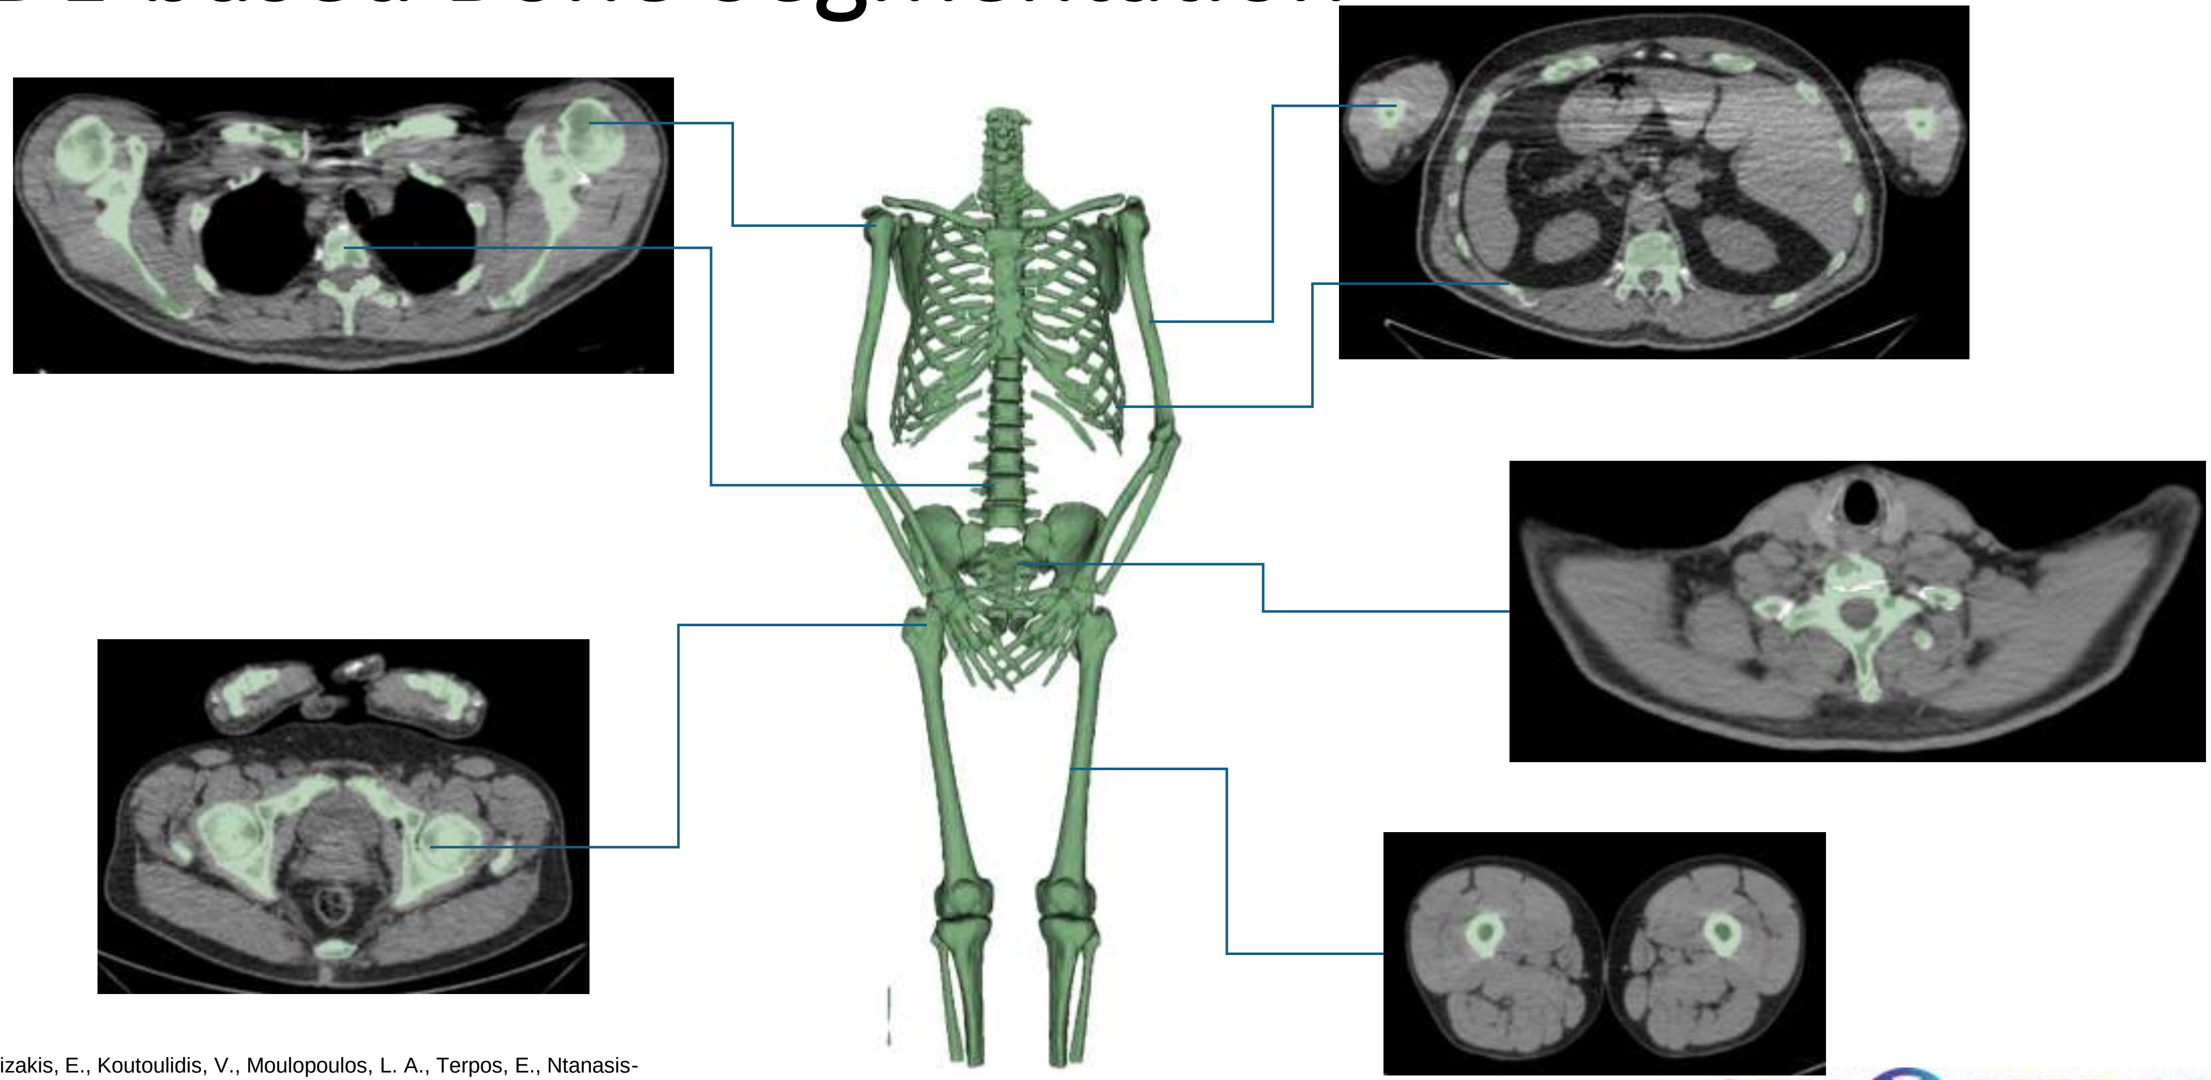

Koutoulakis, E., Trivizakis, E., Koutoulidis, V., Moulopoulos, L. A., Terpos, E., Ntanasis-Stathopoulos, I., ... & Marias, K. (2025, May). Label-Free Machine Learning-Based Segmentation of Whole-Body Bone Marrow Imaging in Multiple Myeloma. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 1-6). IEEE.

Unsupervised clustering

- Identify tissue based on the **natural distribution** of pixels
- **No ground truth** is required
- **Unsupervised** Machine learning analysis
  - o Segmented bones

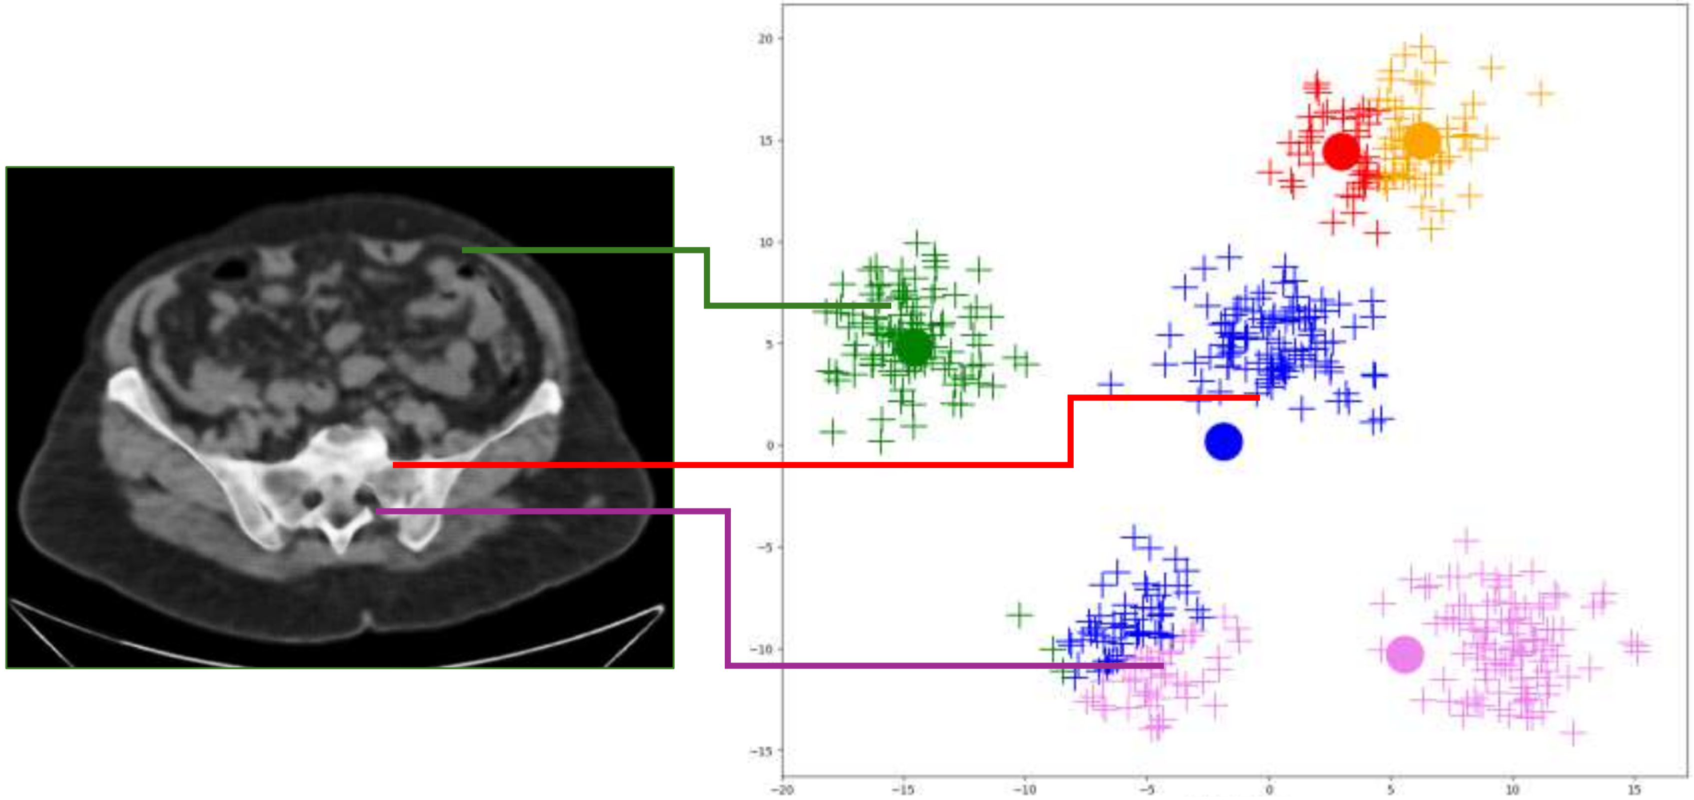

Koutoulakis, E., Trivizakis, E., Koutoulidis, V., Moulopoulos, L. A., Terpos, E., Ntanasis-Stathopoulos, I., ... & Marias, K. (2025, May). Label-Free Machine Learning-Based Segmentation of Whole-Body Bone Marrow Imaging in Multiple Myeloma. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 1-6). IEEE.

Label-free Segmentation – Bone Marrow

Results
- Good resilience to false positive pixels
- DSC 79% (available expert annotations, femur bone only)

Limitations
- Quantitative evaluation
- Tomographic imaging
- Thin bones – difficult to segment
- Bone marrow → fewer pixels than bones

Impact
- New opportunities in quantitative analysis for hematology oncology
- Allows multiple myeloma patient monitoring
- Minimizes segmentation time compared to expert clinicians (couple of minutes VS 5+ hours)

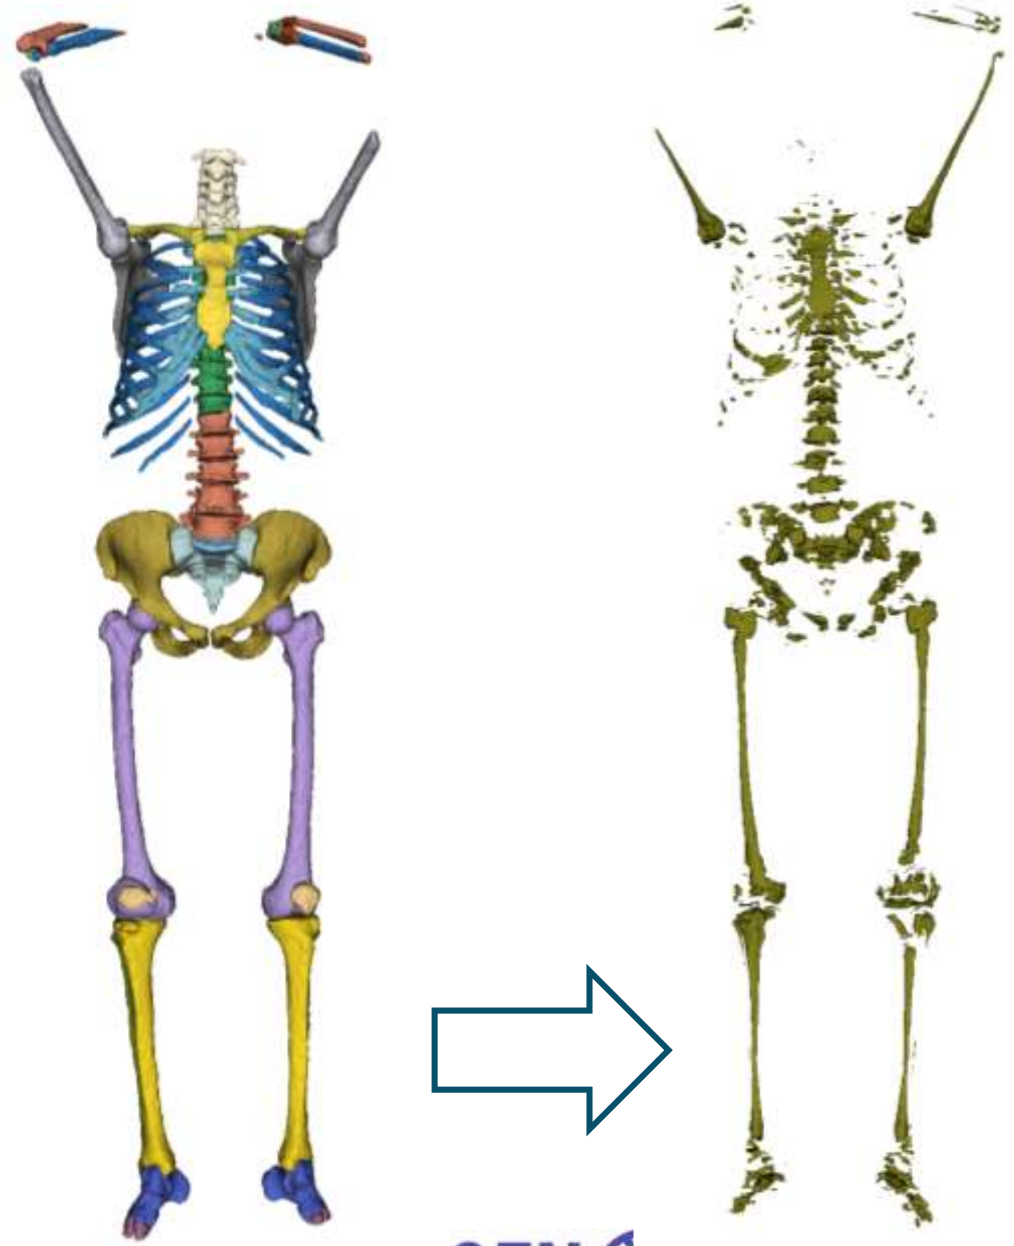

Koutoulakis, E., Trivizakis, E., Koutoulidis, V., Moulopoulos, L. A., Terpos, E., Ntanasis-Stathopoulos, I., ... & Marias, K. (2025, May). Label-Free Machine Learning-Based Segmentation of Whole-Body Bone Marrow Imaging in Multiple Myeloma. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 1-6). IEEE.

Image Denoising

AI against the medical imaging physics.

Data Generation – Synthetic Noise

- Prostate T2WI (a)
- Gaussian noise
- Random blending
- Noise levels: 4 to 12% (b-f)
- The noisy image is the **input**, the original image minus the noise is the output of the DL model
- The model attempts to assess the level of noise in the image

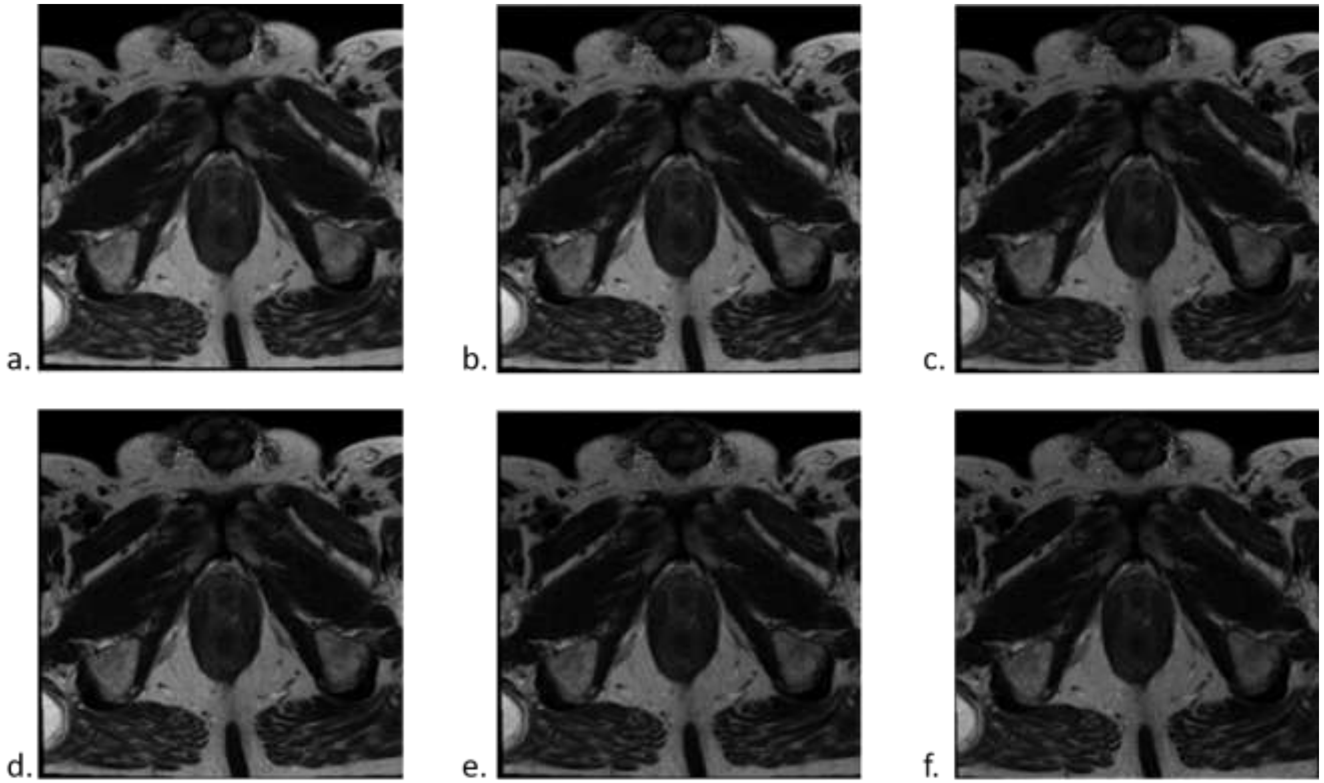

Trivizakis, E., Mylona, E., Nikiforaki, K., Zaridis, D. I., Tsiknakis, M., Fotiadis, D. I., ... & Marias, K. (2025, May). Texture Preserving Deep Learning-Based Noise Reduction for Anatomical Magnetic Resonance Images and its Impact on Imaging Features. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 000223-000230). IEEE.

EUCAIM CANCER IMAGE EUROPE ProCAncer-I

Proposed DLNR

- Adopted the architecture for medical imaging
- Increase the number of learnt filters
- Structural Differences Index
  - Loss function
  - Better texture reconstruction
  - Better than MSE
  - **The DL model learns to minimize the <u>textural differences</u> of an image**

$$l(x,y) = \frac{2\mu_x\mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1}$$

$$c(x,y) = \frac{2\sigma_x\sigma_y + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}$$

$$s(x,y) = \frac{\sigma_{xy} + c_3}{\sigma_x\sigma_y + c_3}$$

$$SSIM(x,y) = l(x,y)^{\alpha} \cdot c(x,y)^{\beta} \cdot s(x,y)^{\gamma}$$

$$SDI = 1 - SSIM(x,y)$$

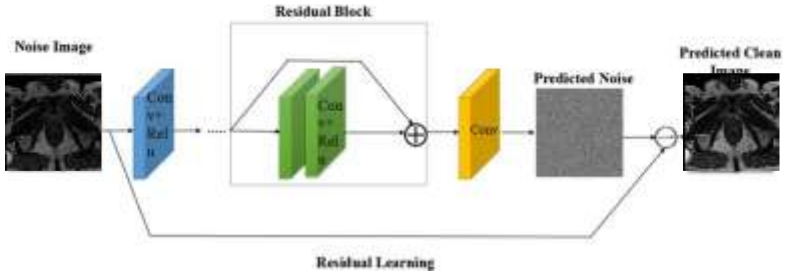

Trivizakis, E., Mylona, E., Nikiforaki, K., Zaridis, D. I., Tsiknakis, M., Fotiadis, D. I., ... & Marias, K. (2025, May). Texture Preserving Deep Learning-Based Noise Reduction for Anatomical Magnetic Resonance Images and its Impact on Imaging Features. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 000223-000230). IEEE.

Impact of Denoising on Radiomics – Horizontally applied

- Region of Interest: Prostate
- Shape, First order, GLSZM
  - Not affected
- GLCM, GLDM, GLRLM
  - Affected

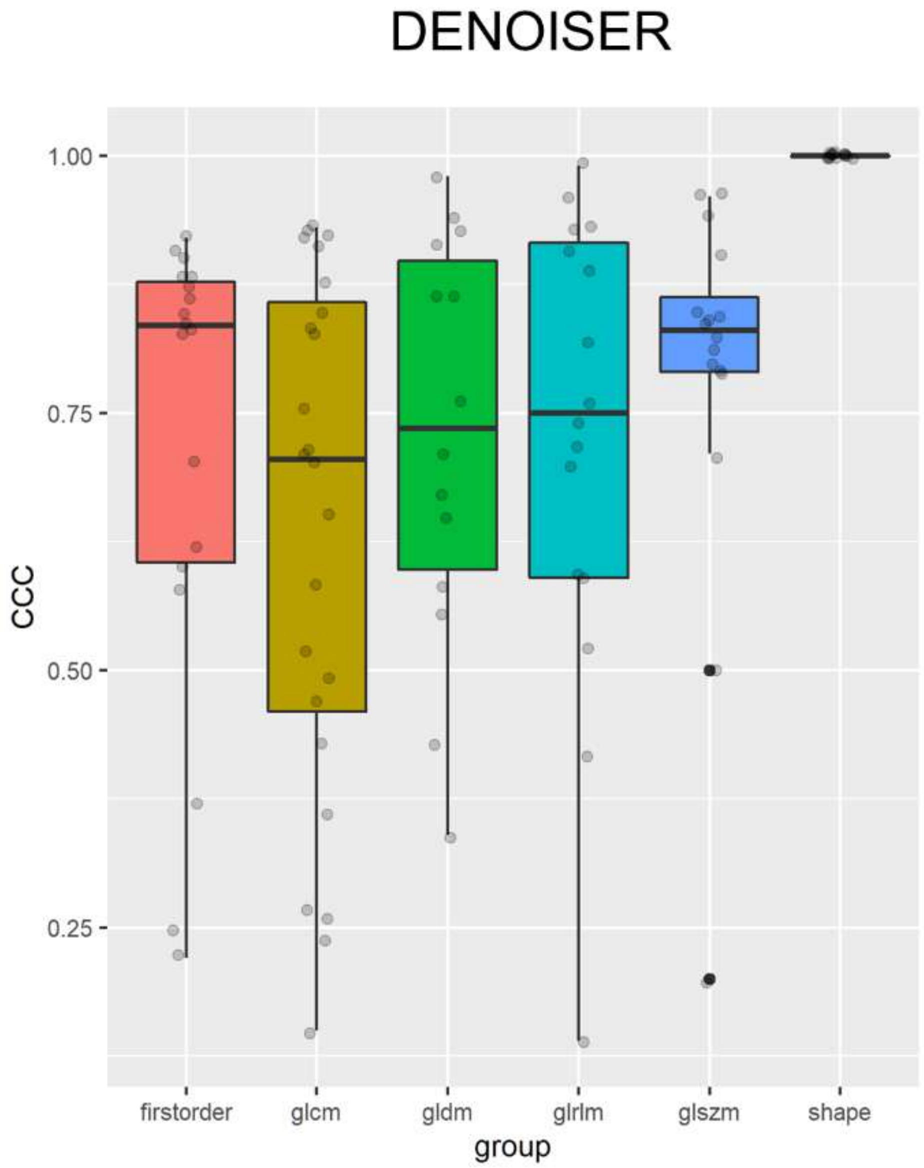

Trivizakis, E., Mylona, E., Nikiforaki, K., Zaridis, D. I., Tsiknakis, M., Fotiadis, D. I., ... & Marias, K. (2025, May). Texture Preserving Deep Learning-Based Noise Reduction for Anatomical Magnetic Resonance Images and its Impact on Imaging Features. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 000223-000230). IEEE.

DLNR – Results

- DLNR
  - Reduces noise
  - Reconstructs the texture of T2W
- Texture reconstruction
  - Noise 4-8%: Good
  - Noise >10%: Challenging
- Some radiomics are affected
  - DLNR → <u>only on noisy images</u>

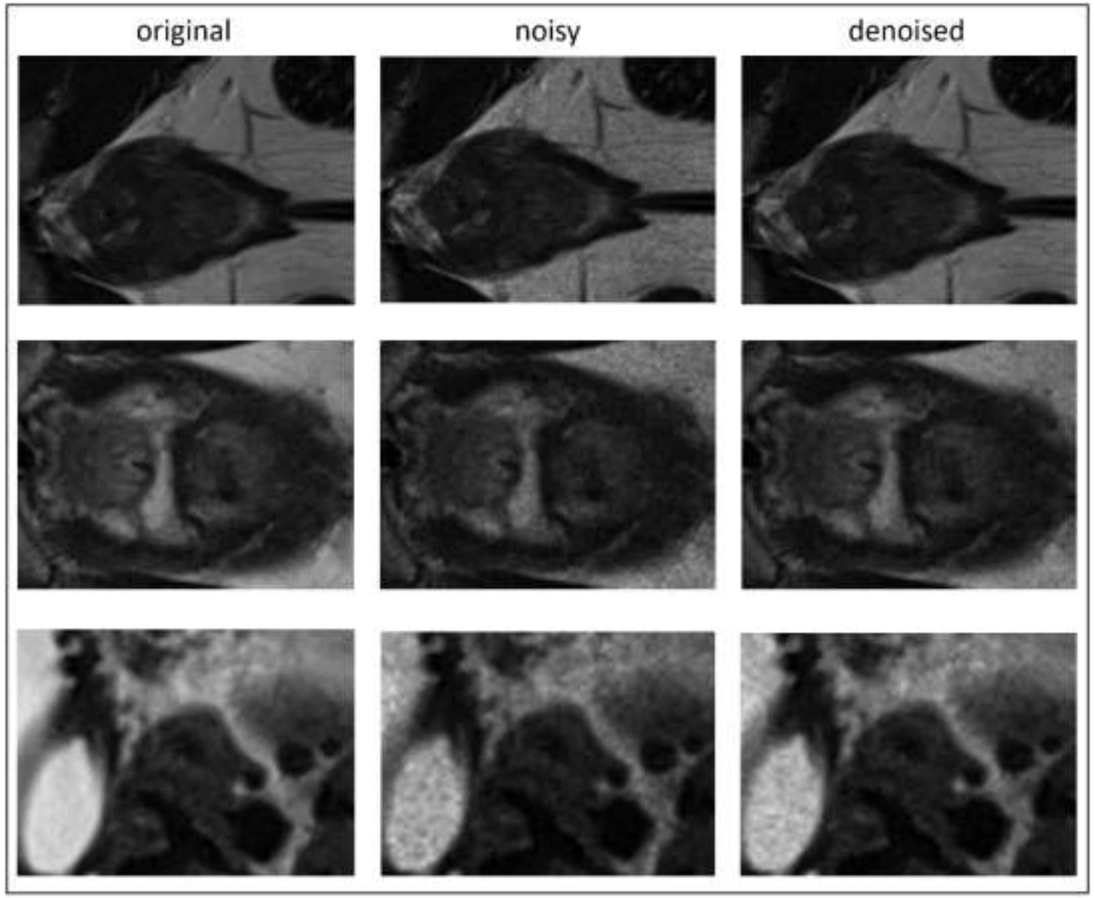

Trivizakis, E., Mylona, E., Nikiforaki, K., Zaridis, D. I., Tsiknakis, M., Fotiadis, D. I., ... & Marias, K. (2025, May). Texture Preserving Deep Learning-Based Noise Reduction for Anatomical Magnetic Resonance Images and its Impact on Imaging Features. In *2025 IEEE 19th International Symposium on Applied Computational Intelligence and Informatics (SACI)* (pp. 000223-000230). IEEE.

Thanks!

Any question?