Early Cancer Detection by Computed Tomography and Artificial Intelligence

Zongwei Zhou, PhD
Dept. Computer Science, Johns Hopkins University
zzhou82@jh.edu | (480)738-2575

JOHNS HOPKINS
UNIVERSITY

2018 — **FELIX**
Lustgarten

2021 — Joined JHU

2023 — **FELIX-Civitas**
Lustgarten &
McGovern

2025 — **FELIX-Civitas**
Lustgarten &
NIH

This talk summarizes a lot of research over the last four years

ZONGWEI ZHOU AWARDED $2.8 MILLION NIH GRANT

Early Detection of Cancer (#2 Killer)

- Early detection of cancer enables surgery and will save many lives.
- For pancreatic cancer, the five-year survival rate increases from about **7-10%** to **40-45%** if detected at an early stage.
- **80,000,000** Computer Tomography (CT) scans taken each year in the US, enabling to screen many people.
- Radiologists can detect pancreatic cancer from CT scans, but the sensitivity of early pancreatic cancer is only **33-44%**.
- This motivates the development of AI algorithms for detecting and localizing early cancer from CT scans, less than 2 cm, and even before tumors are visible.

A Successful Story

- We formulate this problem as *Semantic Segmentation*.
- We developed an AI algorithm and train it to classify voxels as Healthy Pancreas, Tumor, or Background.
- For the pancreas, our AI has achieved very high performance, and can detect tumors 8 months earlier than radiologists.

	Sensitivity early tumors ≤2 cm	Sensitivity all-size tumors	Specificity
Radiologists	33–44%	76–92%	82–96%
Our AI	94%	97%	99%

A Successful Story

8 months earlier …

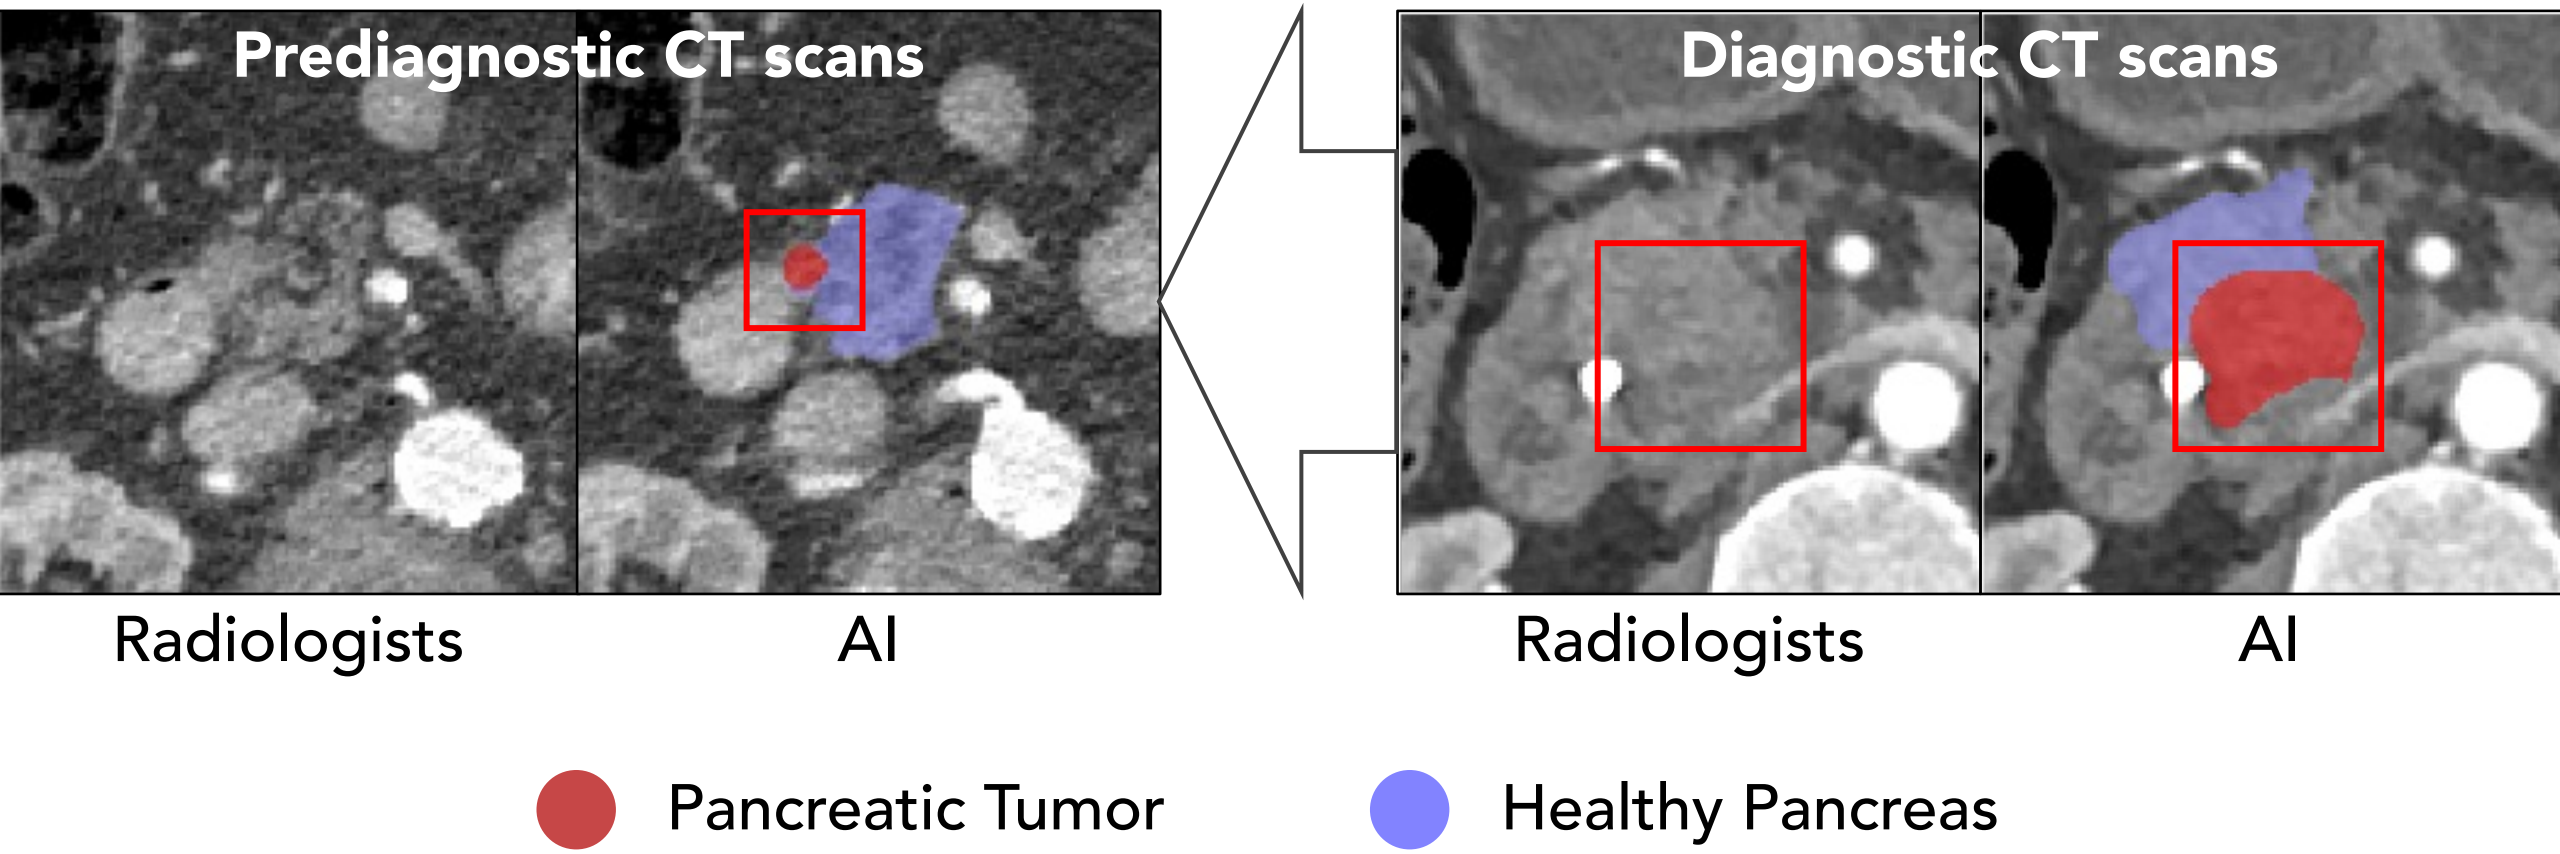

Tumors (0.0001%) in 3D medical images vs. objects (5-50%) in 2D natural images

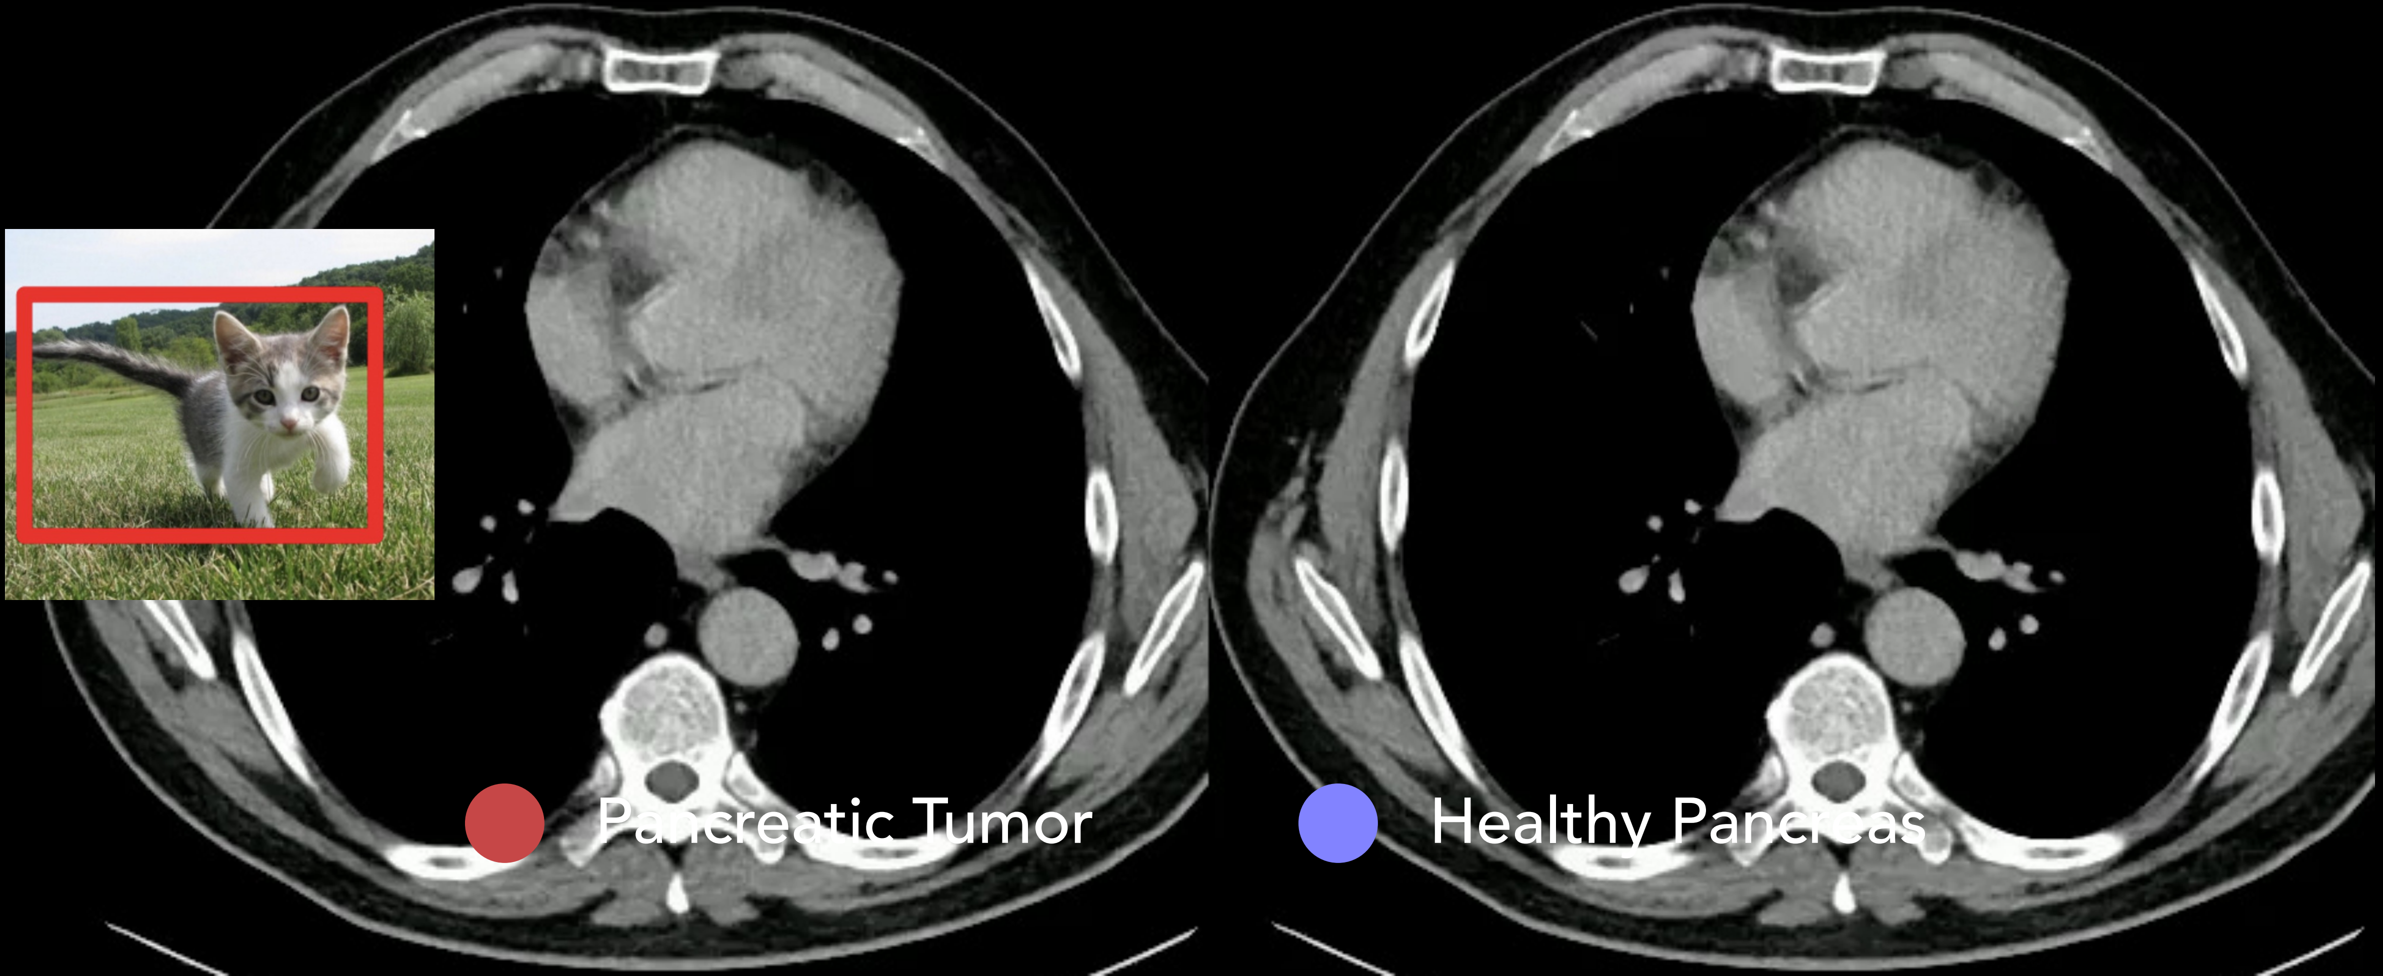

Challenges: Detecting More Cancer Types

- There's a huge data gap in medical AI right now, particularly when you have rare diseases, uncommon conditions.
- We don't have enough "AI-ready" data to train these models.
- AI-ready? Voxel-wise annotation is very time consuming and requires experts (e.g., FELIX@JHU required **25 person years**).
  - In FELIX@JHU, joint project between CS and Radiology, the radiologist team has collected and annotated more than 5,000 CT scans (over 2.5 million images). This is largest dataset in the world to our knowledge dedicated to pancreatic cancer.
- But scaling this effort to many cancer types is not feasible.

Key Questions

- Where to get the data?
- How to annotate the data?
- How to collect data from a variety of hospitals?
- Which algorithms to use?
- How to integrate AI into radiology practice?

Our AI Strategy (Four Chapters)

I. Segmentation algorithms, using voxel-wise annotations
II. Active learning, quickly creating voxel-wise annotations
III. Novel strategies, reducing the need of voxel-wise annotations
    A. Foundation models transfer from organ to tumor tasks
    B. Radiology reports as weak supervision
    C. Synthetic tumors as additional training data
IV. AI helps radiologists, clinical integration
    A. Enable earlier tumor detection with longitudinal, prediagnostic data
    B. Assist in writing radiology reports using tumor/organ segmentation
    C. Use LLMs to automatically retrieve CT scans (15 min → 5 sec per scan)

Chapter I. Segmentation Architectures

- AI is an extremely dynamic research field. Novel AI algorithms are continually being created and improved.

Medical Image Segmentation

1.	U-Net	(O. Ronneberger et al., 2015)	120,000 cites
2.	UNet++	(Z. Zhou et al., 2019)	15,000 cites
3.	TransU-Net	(J. Chen et al., 2021, JHU)	7,800 cites
4.	nnU-Net	(F. Isensee et al., 2020)	7,000 cites
5.	…		

Developing Segmentation Architectures

- UNet++ is suitable for segmenting tumors of a wide range of sizes (Z. Zhou et al., TMI 2019). *Evolution from U-Nets to UNet++*

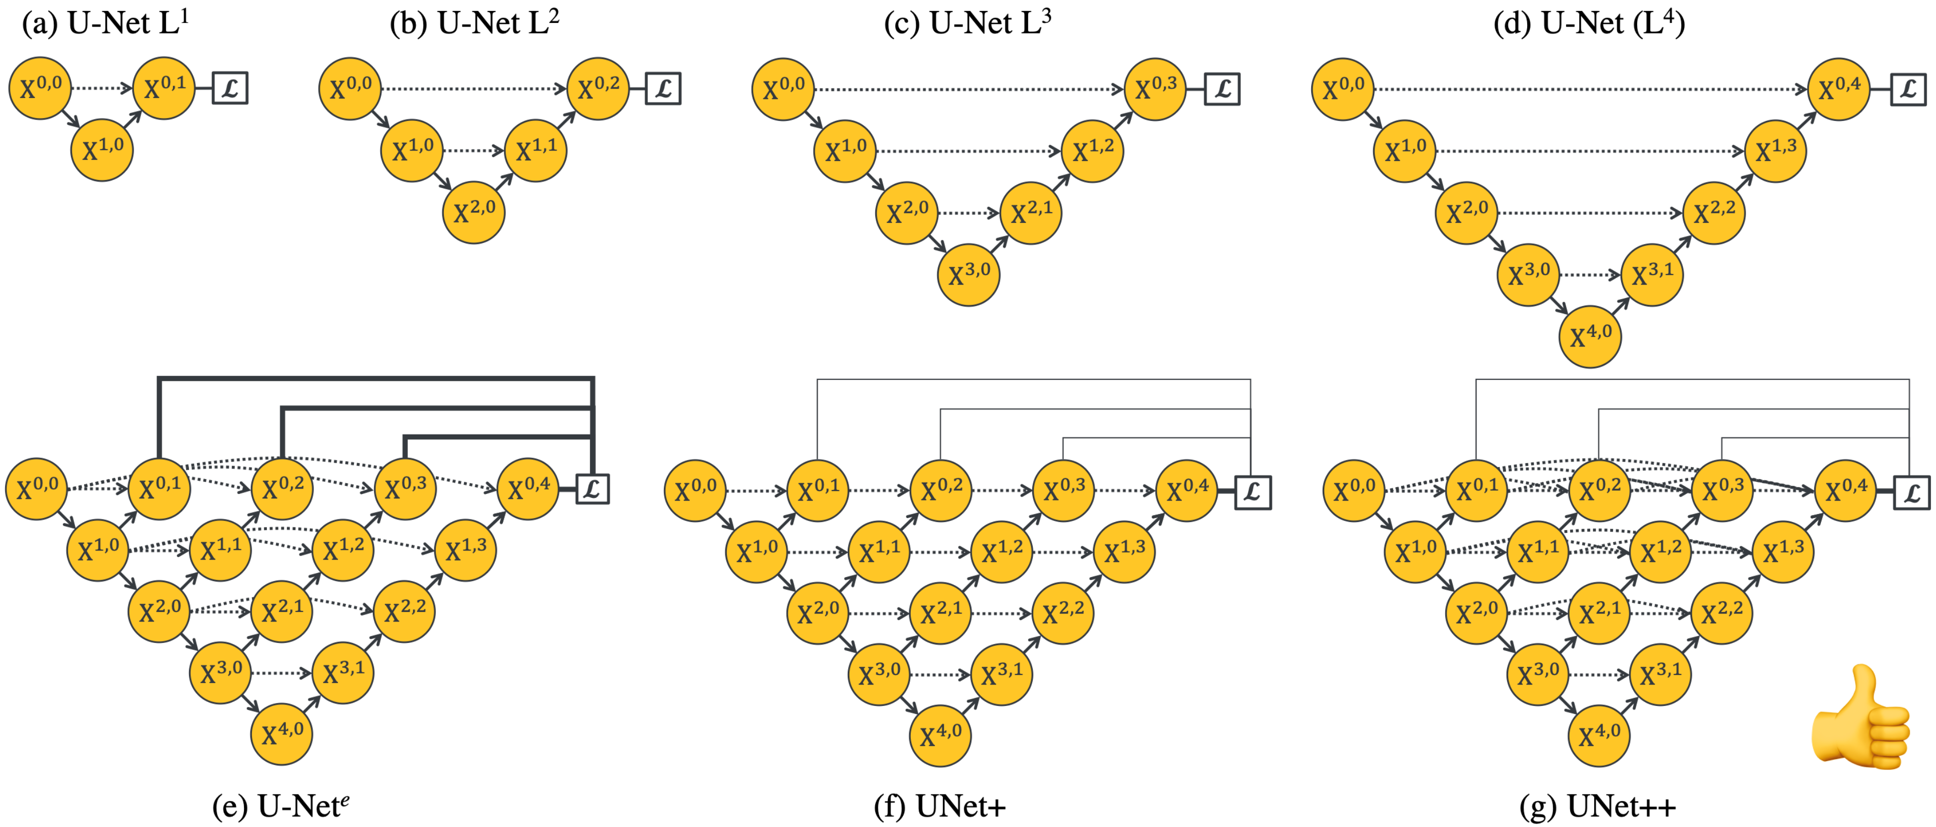

Developing Segmentation Architectures

- UNet++ is faster and more effective.
- UNet++ has inspired many stronger segmentation architectures
- UNet++ has also demonstrated broad clinical and scientific adoption across modalities and *beyond medicine*.

43.9% → 58.1% (U-Net → UNet++) Covid-19 segmentation (CT) [Fan et al., IEEE TMI]	**78.6% → 82.9% (U-Net →UNet++)** Fiber tracing (corneal confocal microscopy) [Mou et al., MICCAI]	**86.5% → 89.5% (U-Net → UNet++)** Spleen segmentation (MRI) [Li et al., Computers & Graphics]
51.2% → 58.6% (U-Net → UNet++) Esophagus segmentation (CT) [Huang et al., IEEE Access]	**63.7% → 66.3% (U-Net → UNet++)** Liver tumor segmentation (CT) [Bajpai et al., Master Thesis]	**90.7% → 91.6% (U-Net → UNet++)** Heart segmentation (MRI) [Ji et al., MICCAI]

Evaluating Segmentation Architectures

- The AI algorithms were trained and tested on CT datasets developed at Johns Hopkins Hospital (JHH).
- The algorithms are also tested on CT scans with pathology-proven results from other institutions in the US, Germany, Poland, and China.
- This project is ongoing with more data from new hospitals.

- *Note: The test data only requires weak annotations, such as radiology/pathology reports (readily available in hospitals), and does not necessarily require voxel-wise annotations.*

Evaluating Segmentation Architectures

- It is critical to test CT scans from other hospitals, as they may use different scanners and imaging protocols, and patient demographics (e.g., race, gender, age) can vary even within the same hospital.
- This is called the *Domain Transfer (DT)* problem (A. Lubonja et al., MICCAIW 2025).

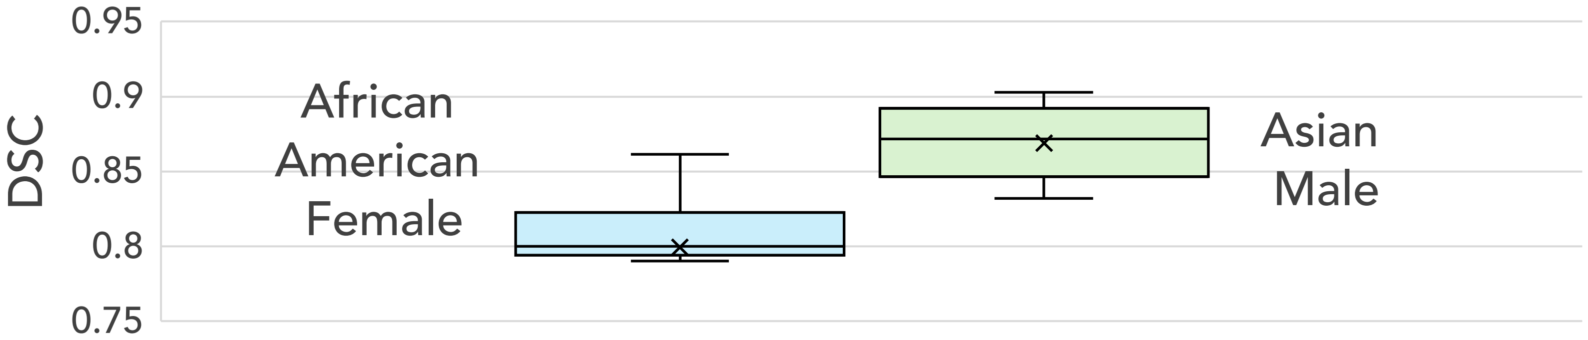

GitHub.com/ariellubonja/RankInsight

Evaluating Segmentation Architectures

- Our preliminary results of DT are promising, but need to be fine-tuned and applied to more scanning protocols.
- We are designing algorithms to perform DT to ensure that our algorithms will transfer to other scanning protocols. See later.

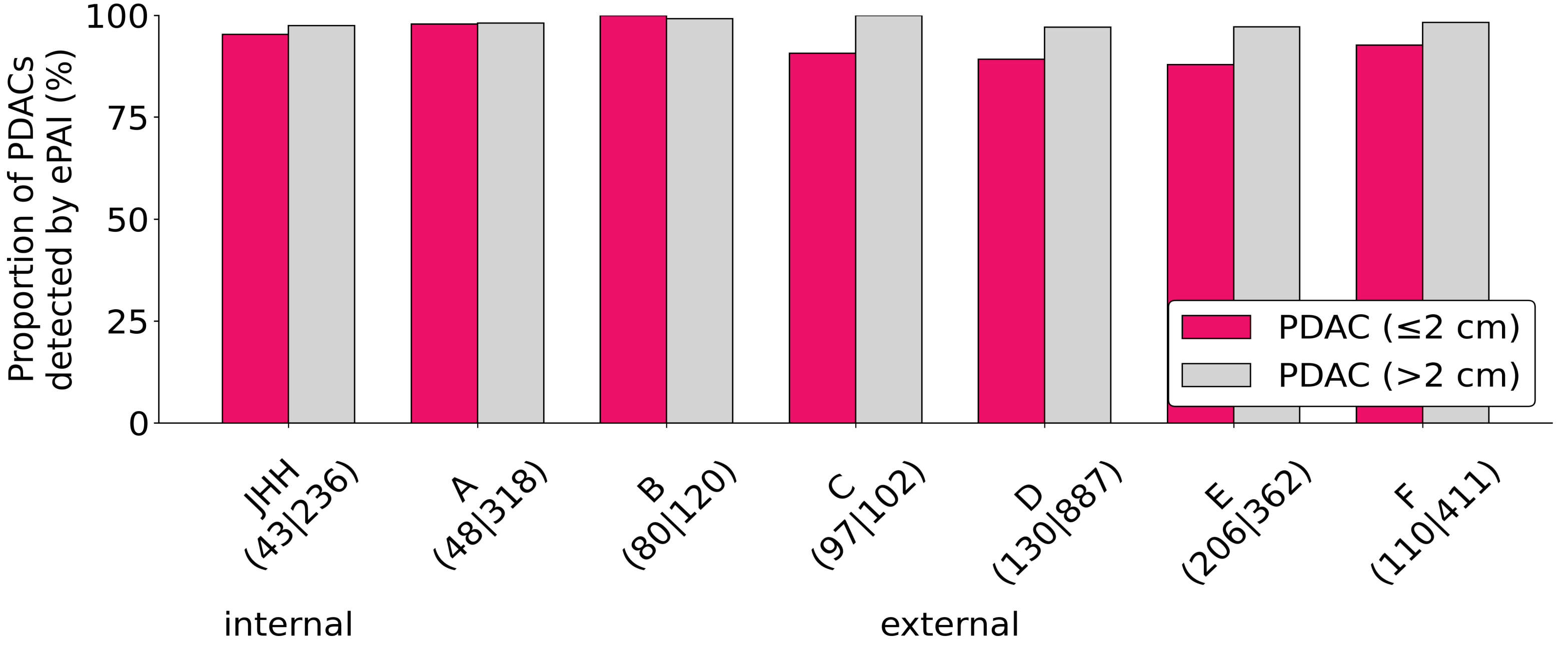

Evaluating Segmentation Architectures

- We initiated a new standard for evaluating medical AI algorithms to improve AI **generalizability** across demographics and hospitals (Bassi et al., NeurIPS 2024).
  - Large, out-of-distribution test set (n = 5,903).
  - Large, multicenter training set (n = 5,196, from 76 hospitals).
  - Inventor-involved training, third-party evaluation.
  - Long-term investigation (Transformers, Mamba, newer architectures).
- New A benchmark 2.0 for tumor segmentation is in preparation.


GitHub.com/MrGiovanni/Touchstone

Chapter II. Annotations with Active Learning

- We created the largest, annotated, public dataset of CT scans with
  - voxel-wise annotations of **6 types of cancer**
  - voxel-wise annotations of **25 organs**
  - patient-wise paired **radiology reports**
- It provided **9,262** patients' CT scans of human subjects with and without cancer collected from **138** hospitals worldwide.
- We created voxel-wise annotations in this dataset by active learning.
  - Speed up organ annotations by **533x**; tumor annotations by **80x**




https://www.zongweiz.com/dataset

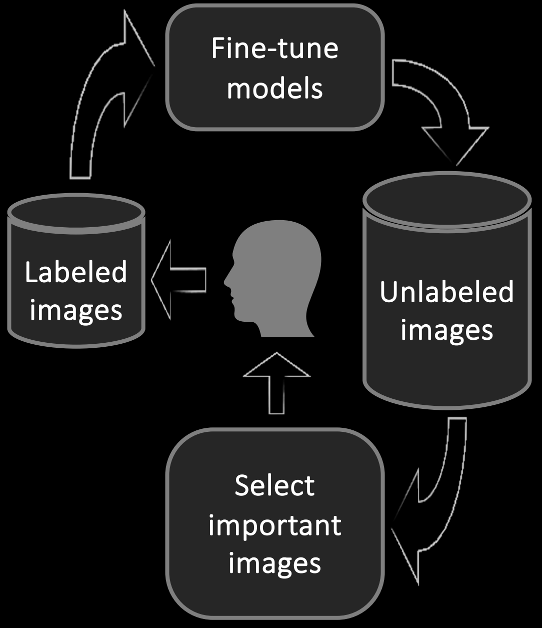

Active Learning

Annotated
25
organs

Annotated
6
cancers

Integrated
15
public datasets

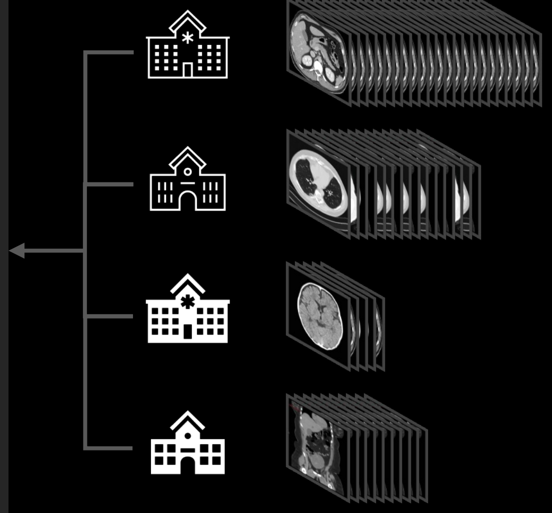

Collected from
138 hospitals
worldwide

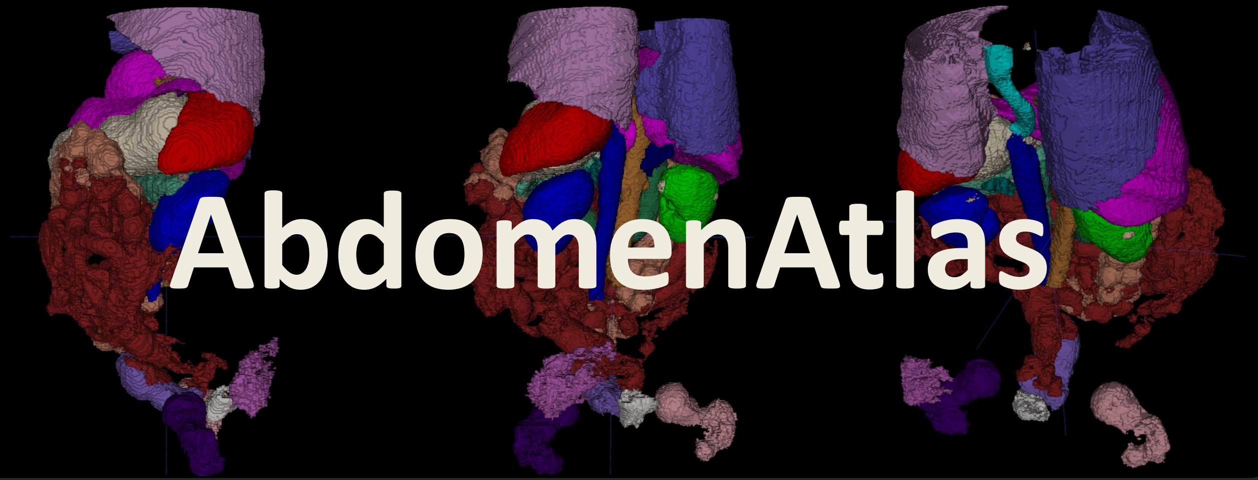

Up to
533x faster
than previous strategies

MONAI

Annotated
3.2M
images

Annotated
9,262
CT volumes

Created in
3 Weeks
by 1 annotator

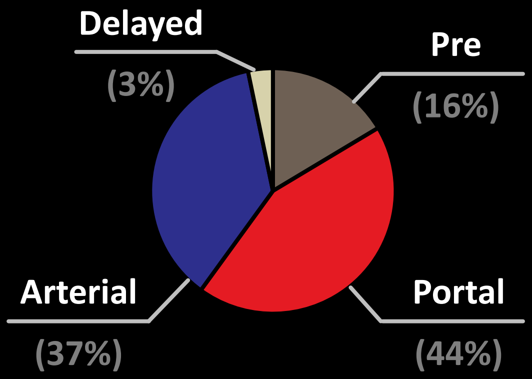

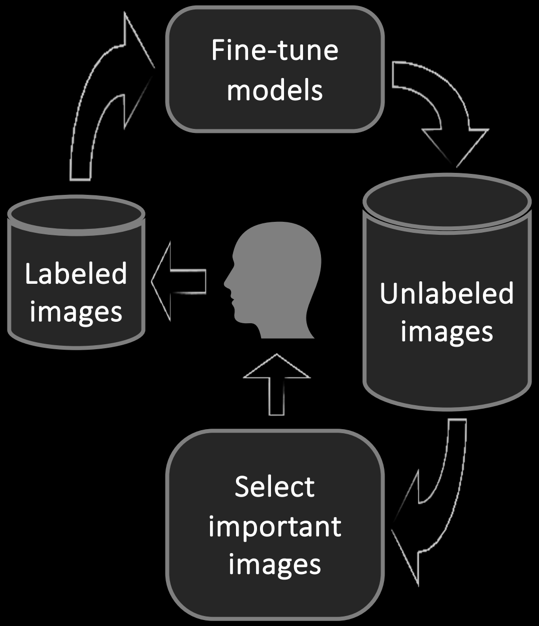

Active Learning

Annotated
25
organs

Annotated
6
cancers

Integrated
15
public datasets

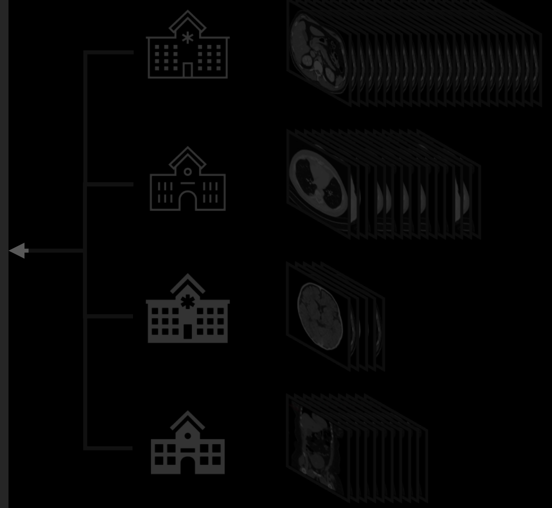

Collected from
138 hospitals
worldwide

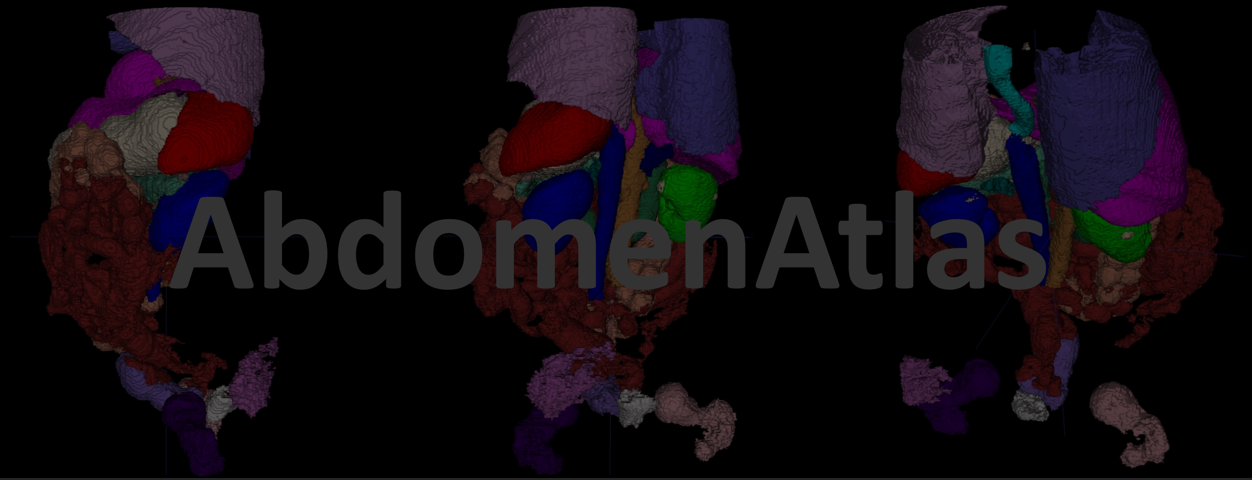

Up to
533x faster
than previous strategies

MONAI

Annotated
3.2M
images

Annotated
9,262
CT volumes

Created in
3 Weeks
by 1 annotator

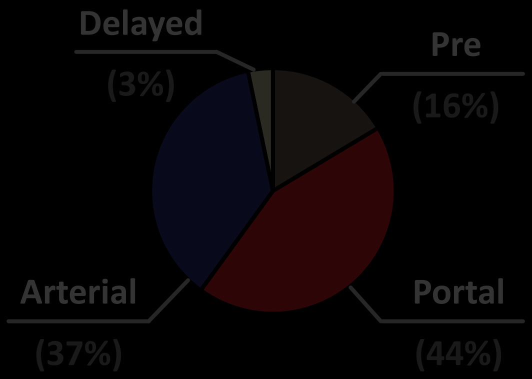

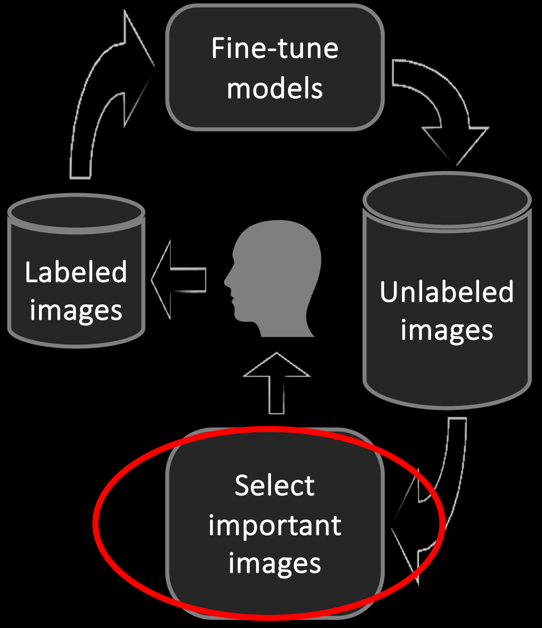

Active Learning

Up to

533x faster

than previous strategies

CT Scan | Inconsistency | Uncertainty | Overlap | Attention Map

We summarized a **taxonomy** of common errors made by AIs and humans (Qiao et al., RSNA 2023 Oral)

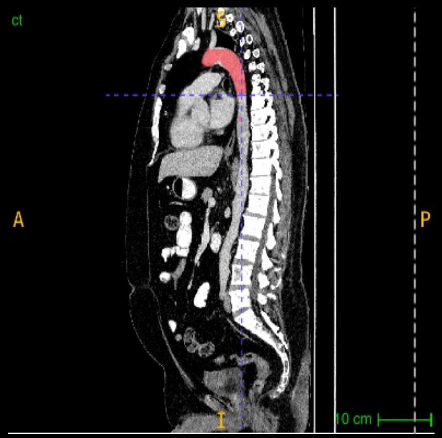

Aorta
Inconsistent labeling protocols

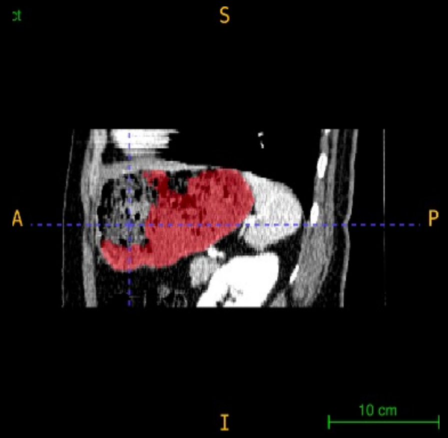

Stomach
uncertainty in empty areas

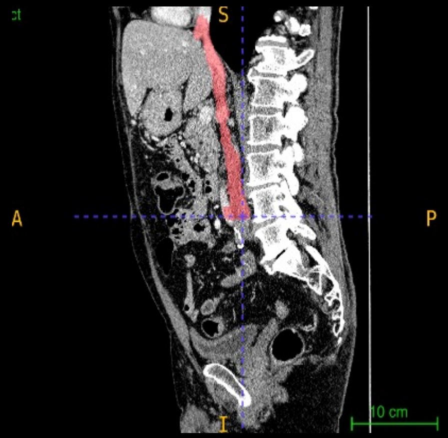

Postcava
ambiguous & blurry boundaries

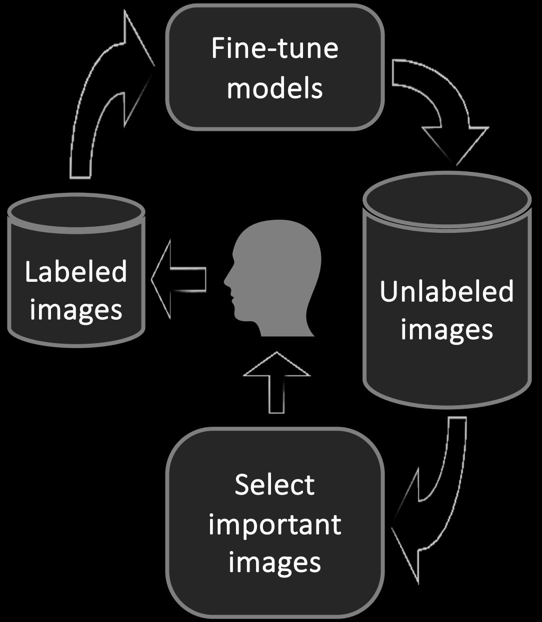

Active Learning

Up to

533x faster

than previous strategies

MONAI

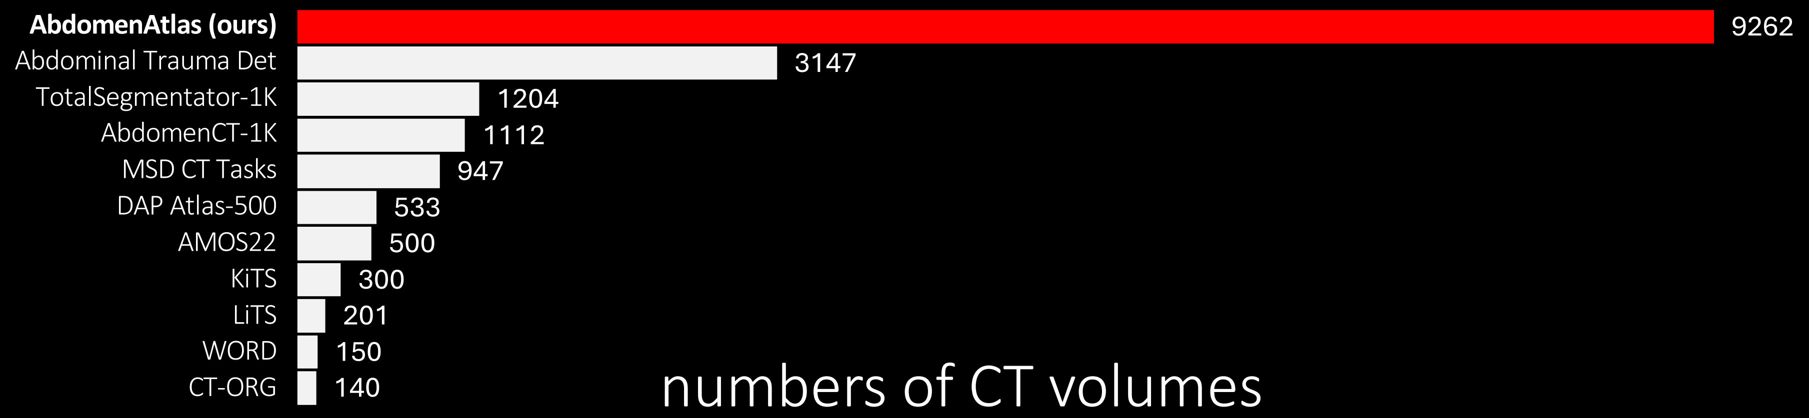

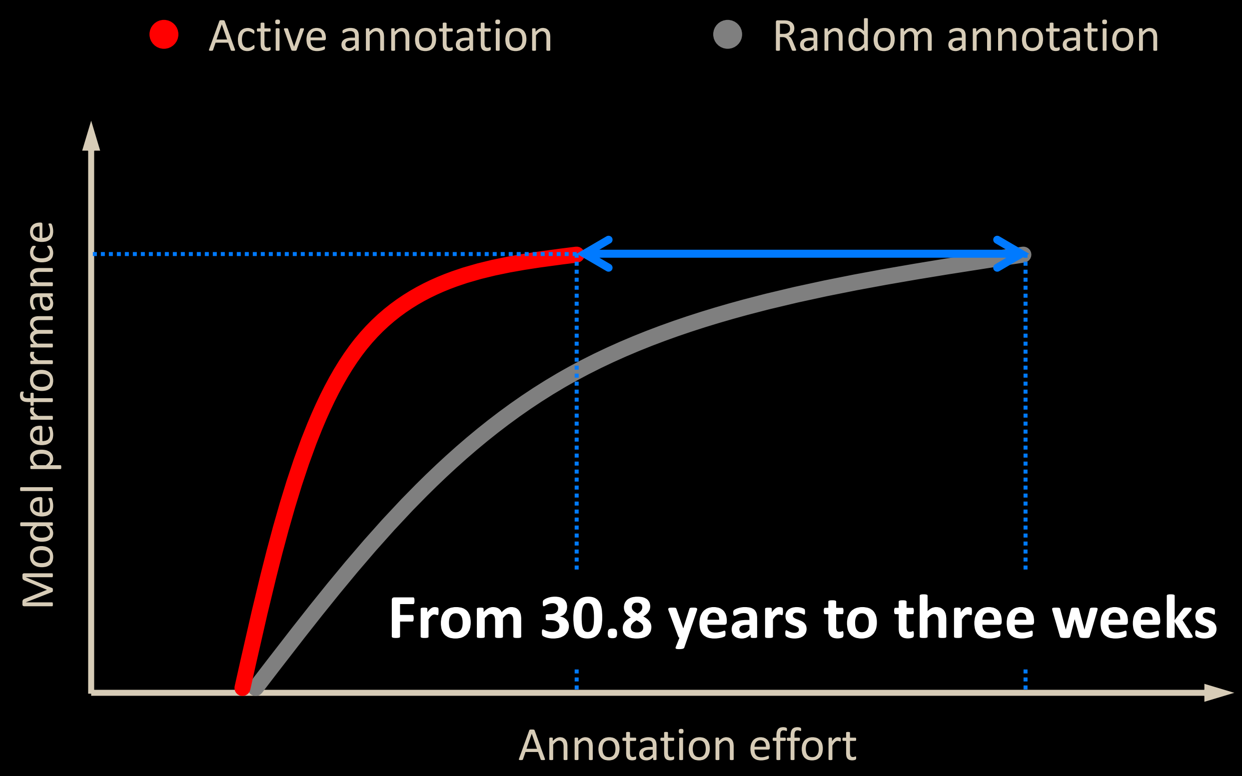

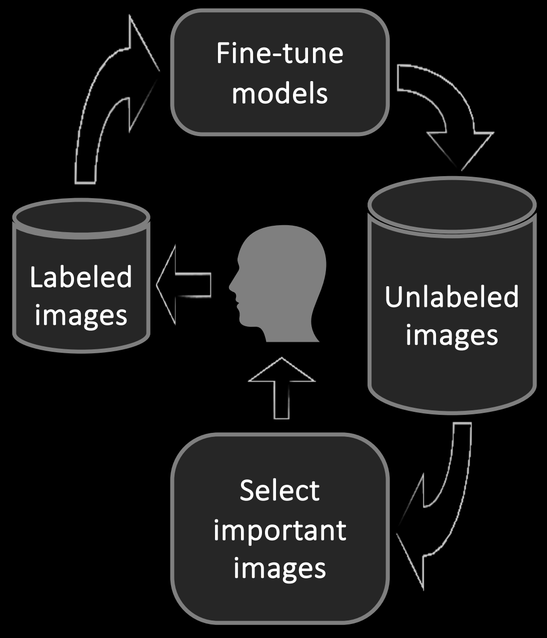

Active Learning

Up to

533x faster

than previous strategies

MONAI

We have released **AbdomenAtlas** of **9,626** CT volumes and **41K** organ masks

(Qu et al., NeurIPS 2023; Li et al., ICLR 2024 Oral)

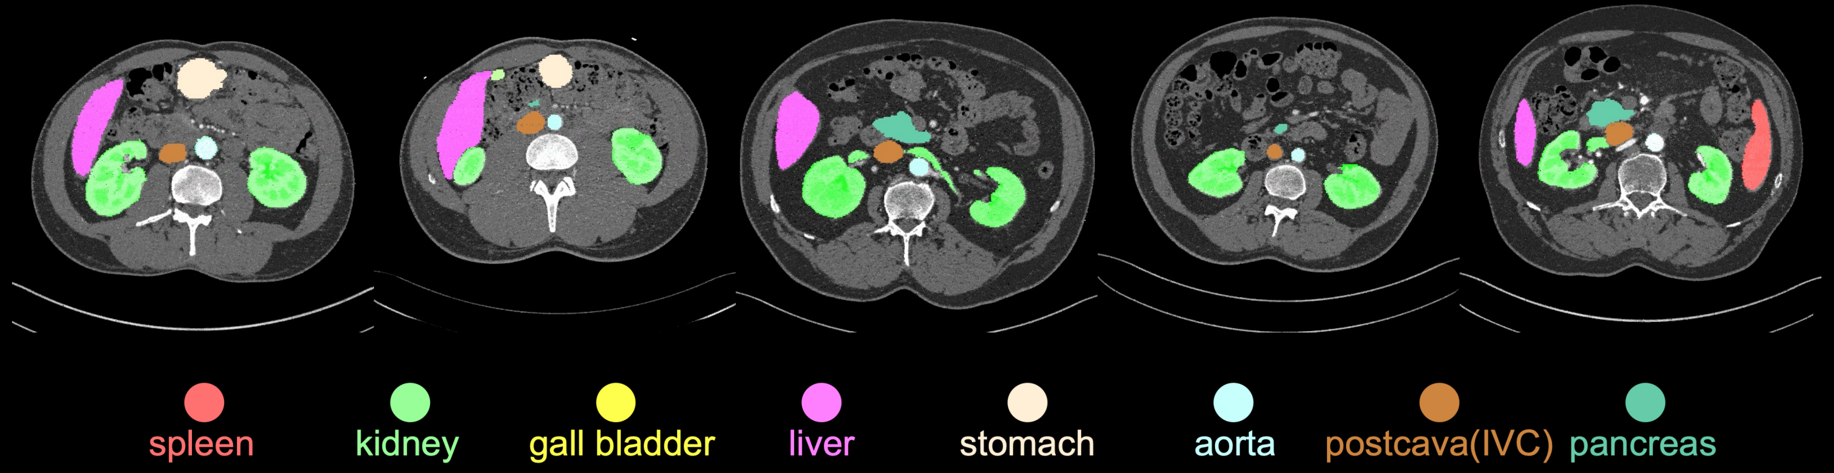

Efficient Tumor Annotations

- Make the AI highly sensitive, offering a strong starting point for radiologist review and edit at least **80×** faster (Zhou et al., ISBI 2024).

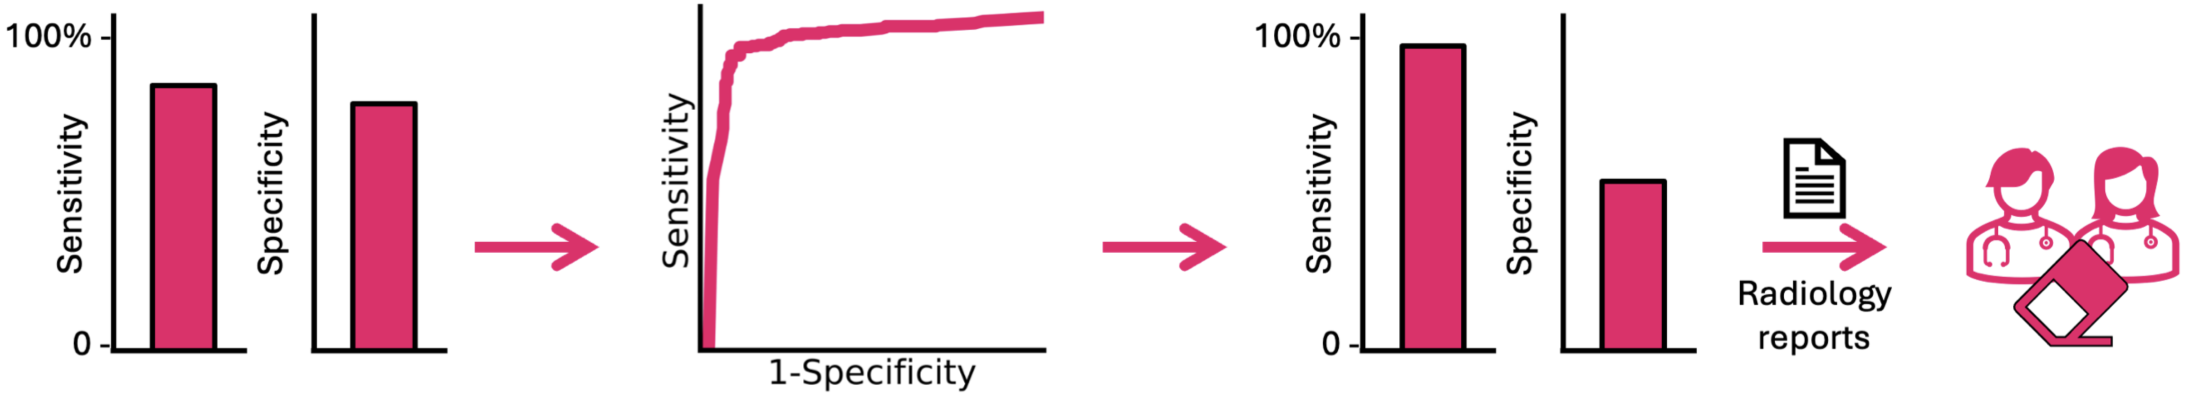

Adjust AI prediction threshold to make the AI overly sensitive,
i.e., false negatives -- false positives ++

Efficient Tumor Annotations

- Make the AI highly sensitive, offering a strong starting point for radiologist review and edit at least **80×** faster (Zhou et al., ISBI 2024).
- (I) Editing an AI-generated tumor takes **~1 minute**. (rarely needed)
- (II) Removing a false positive takes **<5 seconds**.
- In contrast, manual annotation from scratch takes **4–5 minutes**.

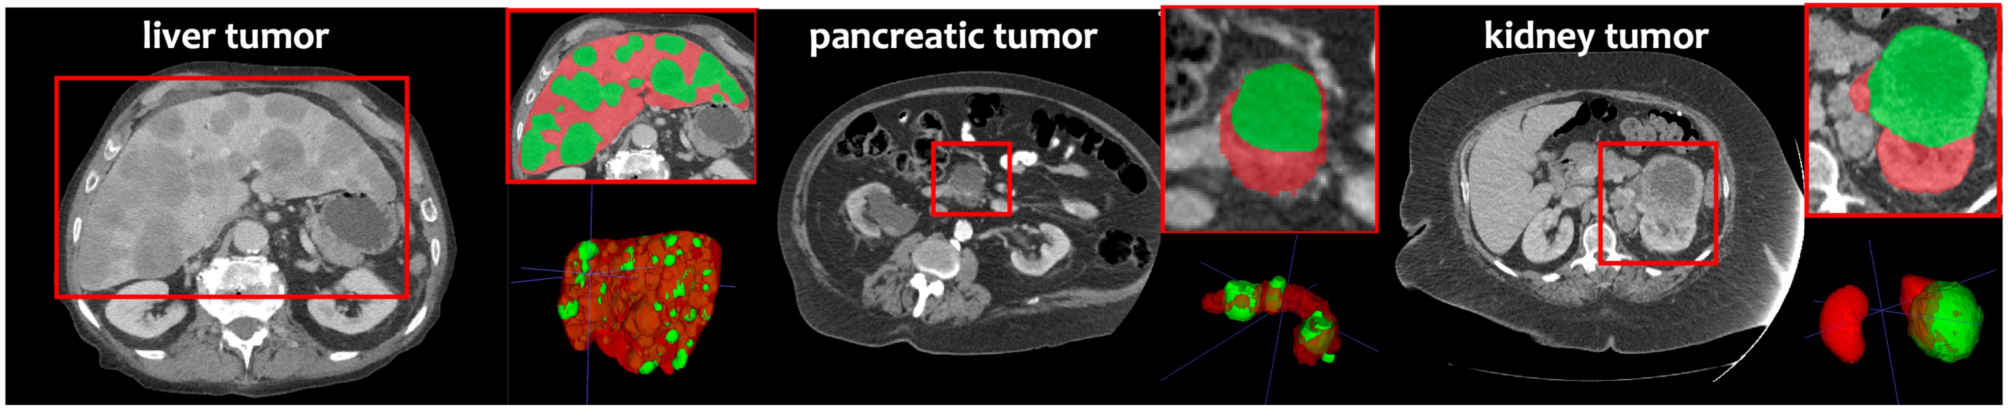

PanTS

Pancreatic Tumor Segmentation

  

GitHub.com/MrGiovanni/PanTS

```
git clone https://github.com/MrGiovanni/PanTS.git; cd PanTS
bash download_PanTS_data.sh
bash download_PanTS_label.sh
http://www.cs.jhu.edu/~zongwei/dataset/PanTSMini_Label.tar.gz
```

PanTS is a large-scale, multi-institutional dataset, containing **36,390** three-dimensional CT volumes from **145** medical centers, with expert-validated, voxel-wise annotations of over **993,000** anatomical structures, including ***pancreatic tumors, pancreas head, body, and tail, and 24 surrounding anatomical structures such as vascular/skeletal structures and abdominal/thoracic organs.***

(Li et al., NeurIPS 2025)



GitHub.com/MrGiovanni/PanTS

A Huge AI-Ready Internal Dataset

BodyMaps

Internal use only; open for collaboration

Funded by

NIH R01 (PI: Zongwei Zhou, Yang Yang, Kang Wang),
Lustgarten Foundation (PI: Alan Yuille), and
McGovern Foundation (PI: Alan Yuille)

81.7 million	241,336	300	145
2D CT images	3D CT volumes	anatomical structures	hospitals

300 anatomical structures: adrenal gland left · adrenal gland right · airway · anterior scalene left · anterior scalene right · aorta · artery brachiocephalic · artery common carotid left · artery common carotid right · artery internal carotid left · artery internal carotid right · artery subclavian left · artery subclavian right · atrial appendage left · atrium left · atrium right · auditory canal left · auditory canal right · autochthon left · autochthon right · bladder · body · body extremities · body trunc · bone · brachiocephalic trunk · brachiocephalic vein left · brachiocephalic vein right · brain · brain ventricle · brainstem · breast left · breast right · bronchus · carpal · caudate nucleus · celiac trunk · central sulcus · cerebellum · cheek left · cheek right · clavicula left · clavicula right · colon · common bile duct · common carotid artery left · common carotid artery right · coronary artery · costal cartilages · cricoid cartilage · digastric left · digastric right · duodenum · esophagus · eye lens left · eye lens right · eyeball left · eyeball right · fat · femur left · femur right · fibula · frontal lobe · gall bladder · gluteus maximus left · gluteus maximus right · gluteus medius left · gluteus medius right · gluteus minimus left · gluteus minimus right · gonads · hard palate · heart · heart atrium left · heart atrium right · heart myocardium · heart tissue · heart ventricle left · heart ventricle right · hepatic vessel · hip left · hip right · humerus left · humerus right · hyoid · hypopharynx · iliac artery left · iliac artery right · iliac vena left · iliac vena right · iliopsoas left · iliopsoas right · inferior oblique muscle left · inferior oblique muscle right · inf ... ernal capsule · internal ... tery

A team of 23 board-certified radiologists

A Huge AI-Ready Internal Dataset

BodyMaps

Internal use only; open for collaboration

Funded by

NIH R01 (PI: Zongwei Zhou, Yang Yang, Kang Wang),

Lustgarten Foundation (PI: Alan Yuille), and

McGovern Foundation (PI: Alan Yuille)

81.7 million	**241,336**	**300**	**145**
2D CT images	3D CT volumes	anatomical structures	hospitals

Voxel-wise annotated 16 tumor types: adrenal · bladder · bone · breast · colon · duodenum · esophagus · gallbladder · kidney · liver · lung · pancreas · prostate · spleen · stomach · uterus

 

A team of 23 board-certified radiologists

A Series of AI-Ready Datasets

"Annotating **240,000 CT scans** with **72 million anatomical shapes** would require an expert radiologist to have started working around 420 BCE—the era of Hippocrates—to complete the task by 2025. *We did it in two years*." says lead author Zongwei Zhou

Chapter III. Strategies to Further Reduce Need of Voxel-Wise Annotations

A. Vision foundation models: transfer from organ to tumor tasks
B. Radiology reports as weak supervision for multi-cancer detection
C. Synthetic tumors as additional training data for small tumors

Chapter III.A. Foundation Models

- Two major strategies
- Self-supervised pre-training (Z. Zhou et al., MICCAI 2019 Young Scientist Award; MIA 2020 Best Paper Award)
  - Mask image modeling, no need for voxel-wise annotations
- Supervised pre-training (Li et al., ICLR 2024 Oral)
  - Organ segmentation, requiring voxel-wise annotations
  - The models are pre-trained on **9,262** voxel-wise annotated CT scans
  - The dataset & annotation used for training are public (Li et al., MEDIA 2024).

GitHub.com/MrGiovanni/SuPreM

Scaling Laws in Foundation Models

- Supervised pre-training helps the model to learn image features that are relevant to downstream tumor tasks (e.g., organ segmentation).
- The need for voxel-wise annotated tumor scans was reduced by **52.6%**.

GitHub.com/MrGiovanni/SuPreM

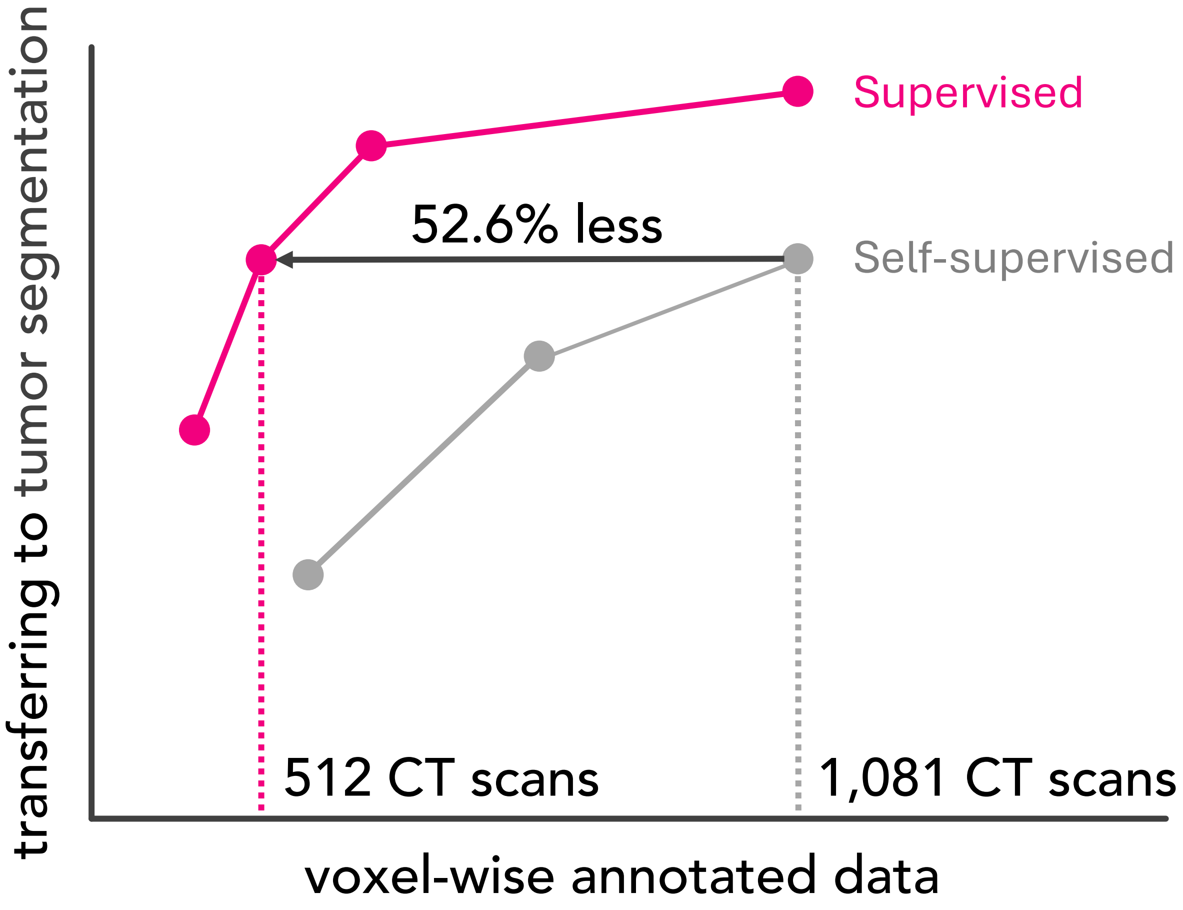

Chapter III.B. Reports as Weak Supervision

- Public datasets have few tumor Image-Mask pairs, only **10s to 100s**.
- By contrast, reports are written every day by radiologists—public datasets have more than **500,000** Image-Report pairs.
- We enable AI to learn tumor segmentation directly from these reports (P. Bassi et al., MICCAI 2025 Best Paper Award, Runner-up).
- This is a collaboration with UCSF and other institutions.

GitHub.com/MrGiovanni/R-Super

Chapter III.B. Reports as Weak Supervision

- R-Super, a novel AI training method that enforce the consistency between AI segmented tumors and report descriptions such as tumor **number, size, and location**.

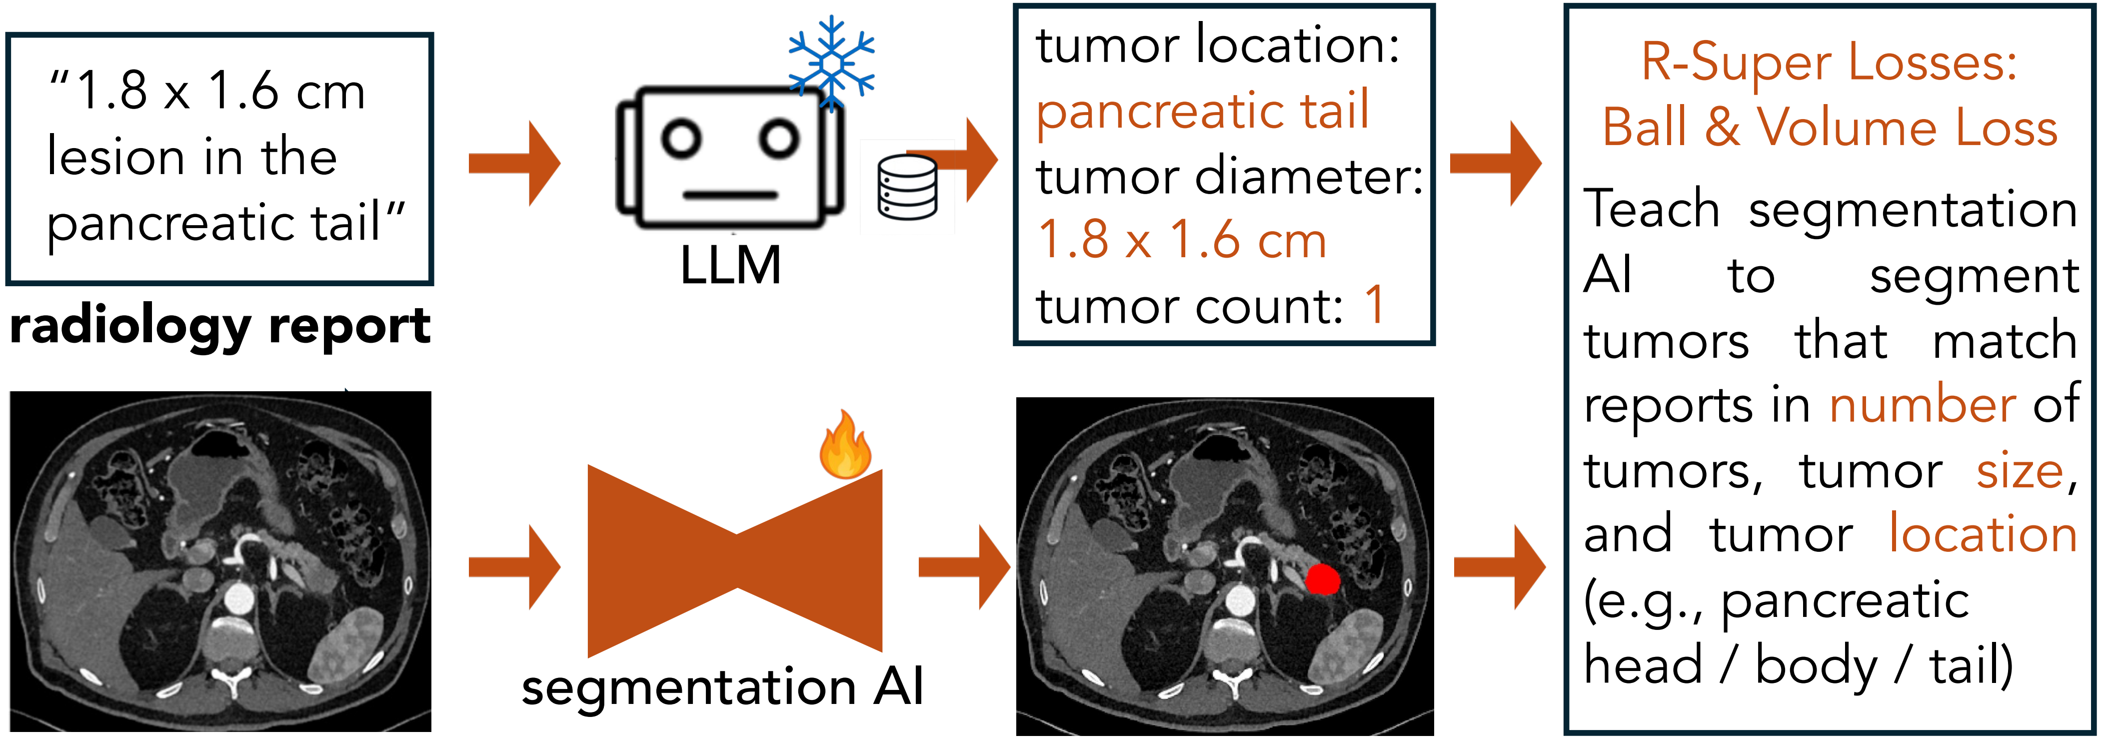

Scaling Laws in Reports Supervision

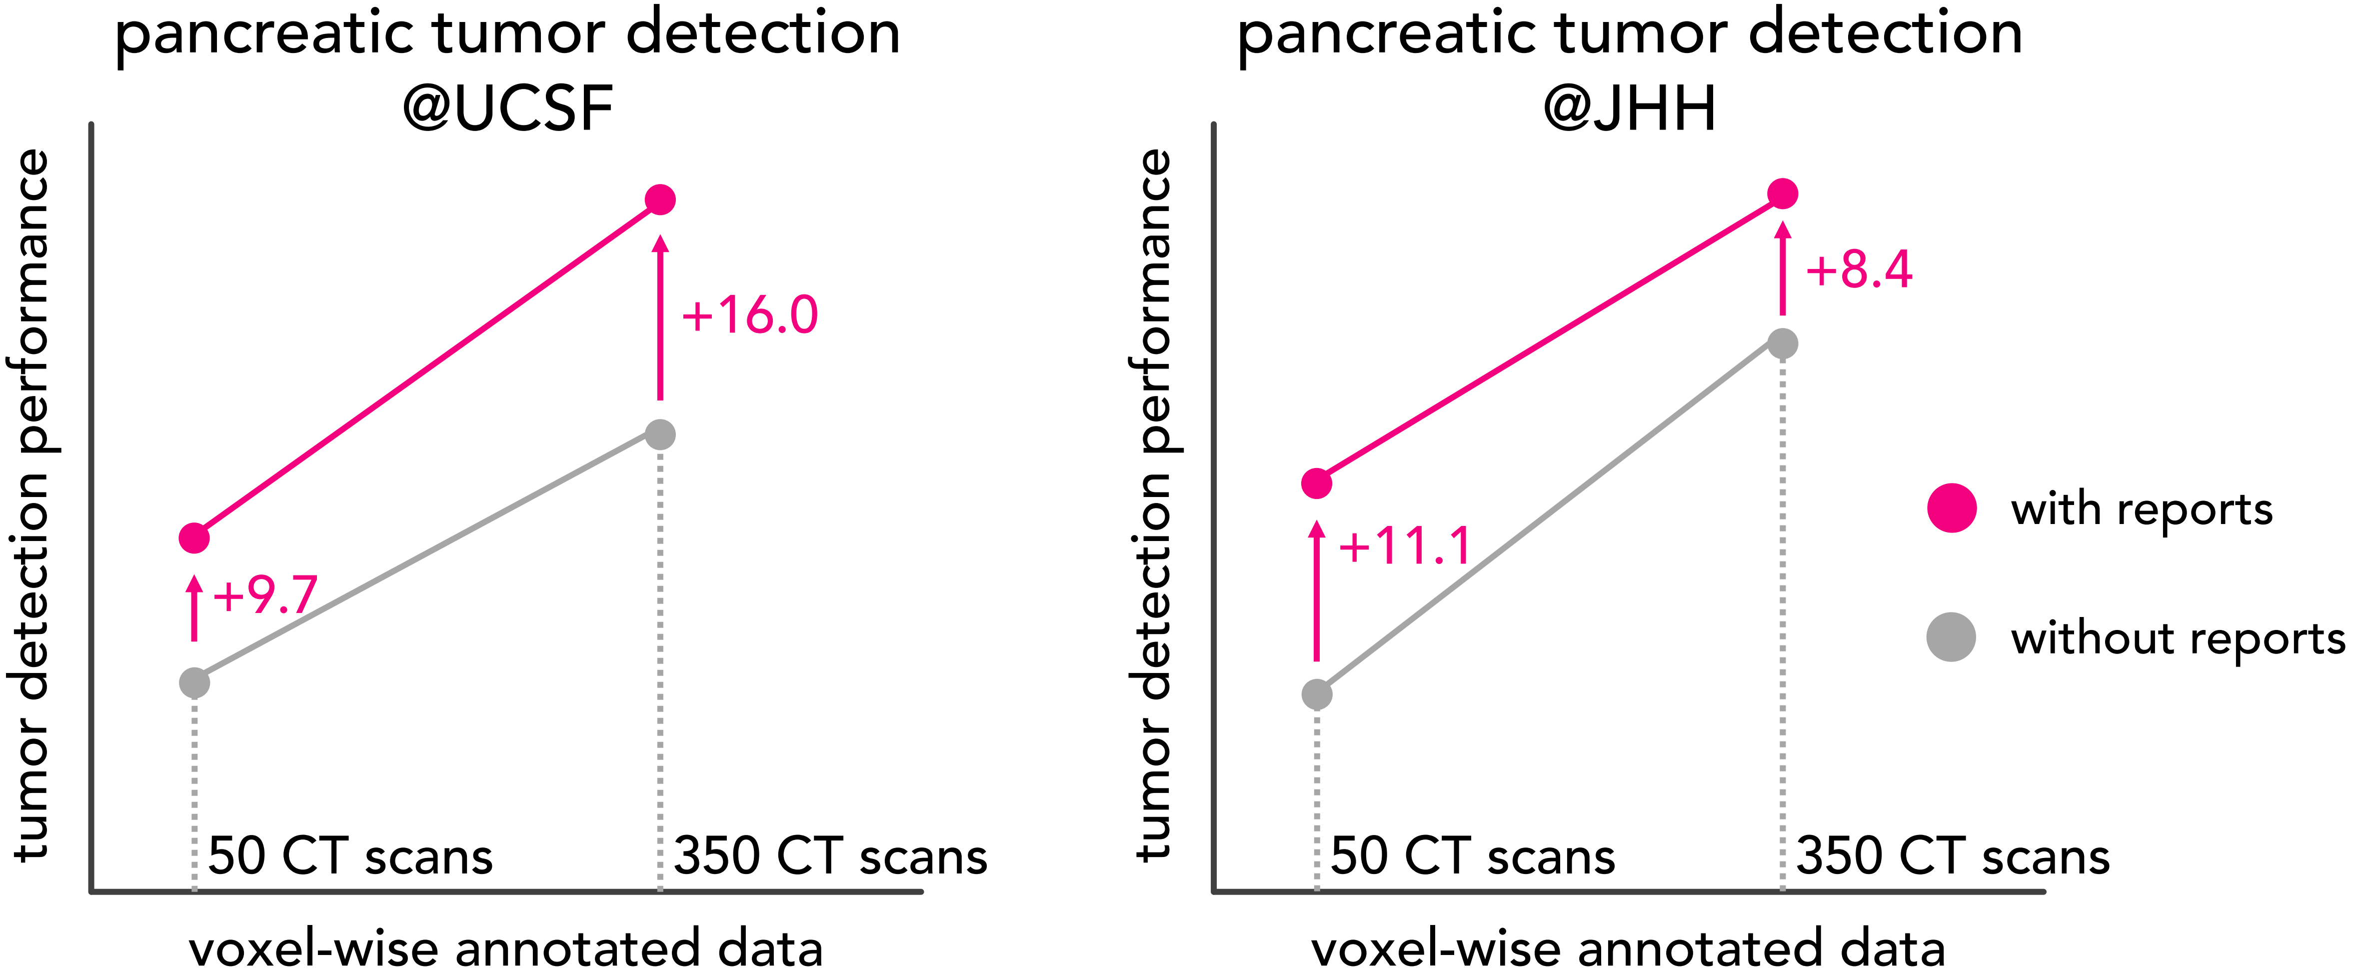

Reports Supervision for Multi-Cancer

- We have curated a dataset of **117,000** Image–Report and **270** CT-Mask pairs for tumors in the adrenal, bladder, esophagus, gallbladder, prostate, spleen, and uterus.
- No publicly available Image–Mask pairs exist for these tumor types.
- We will release the first AI model that can segment these tumor types.

GitHub.com/MrGiovanni/R-Super

Reports Supervision for Multi-Cancer

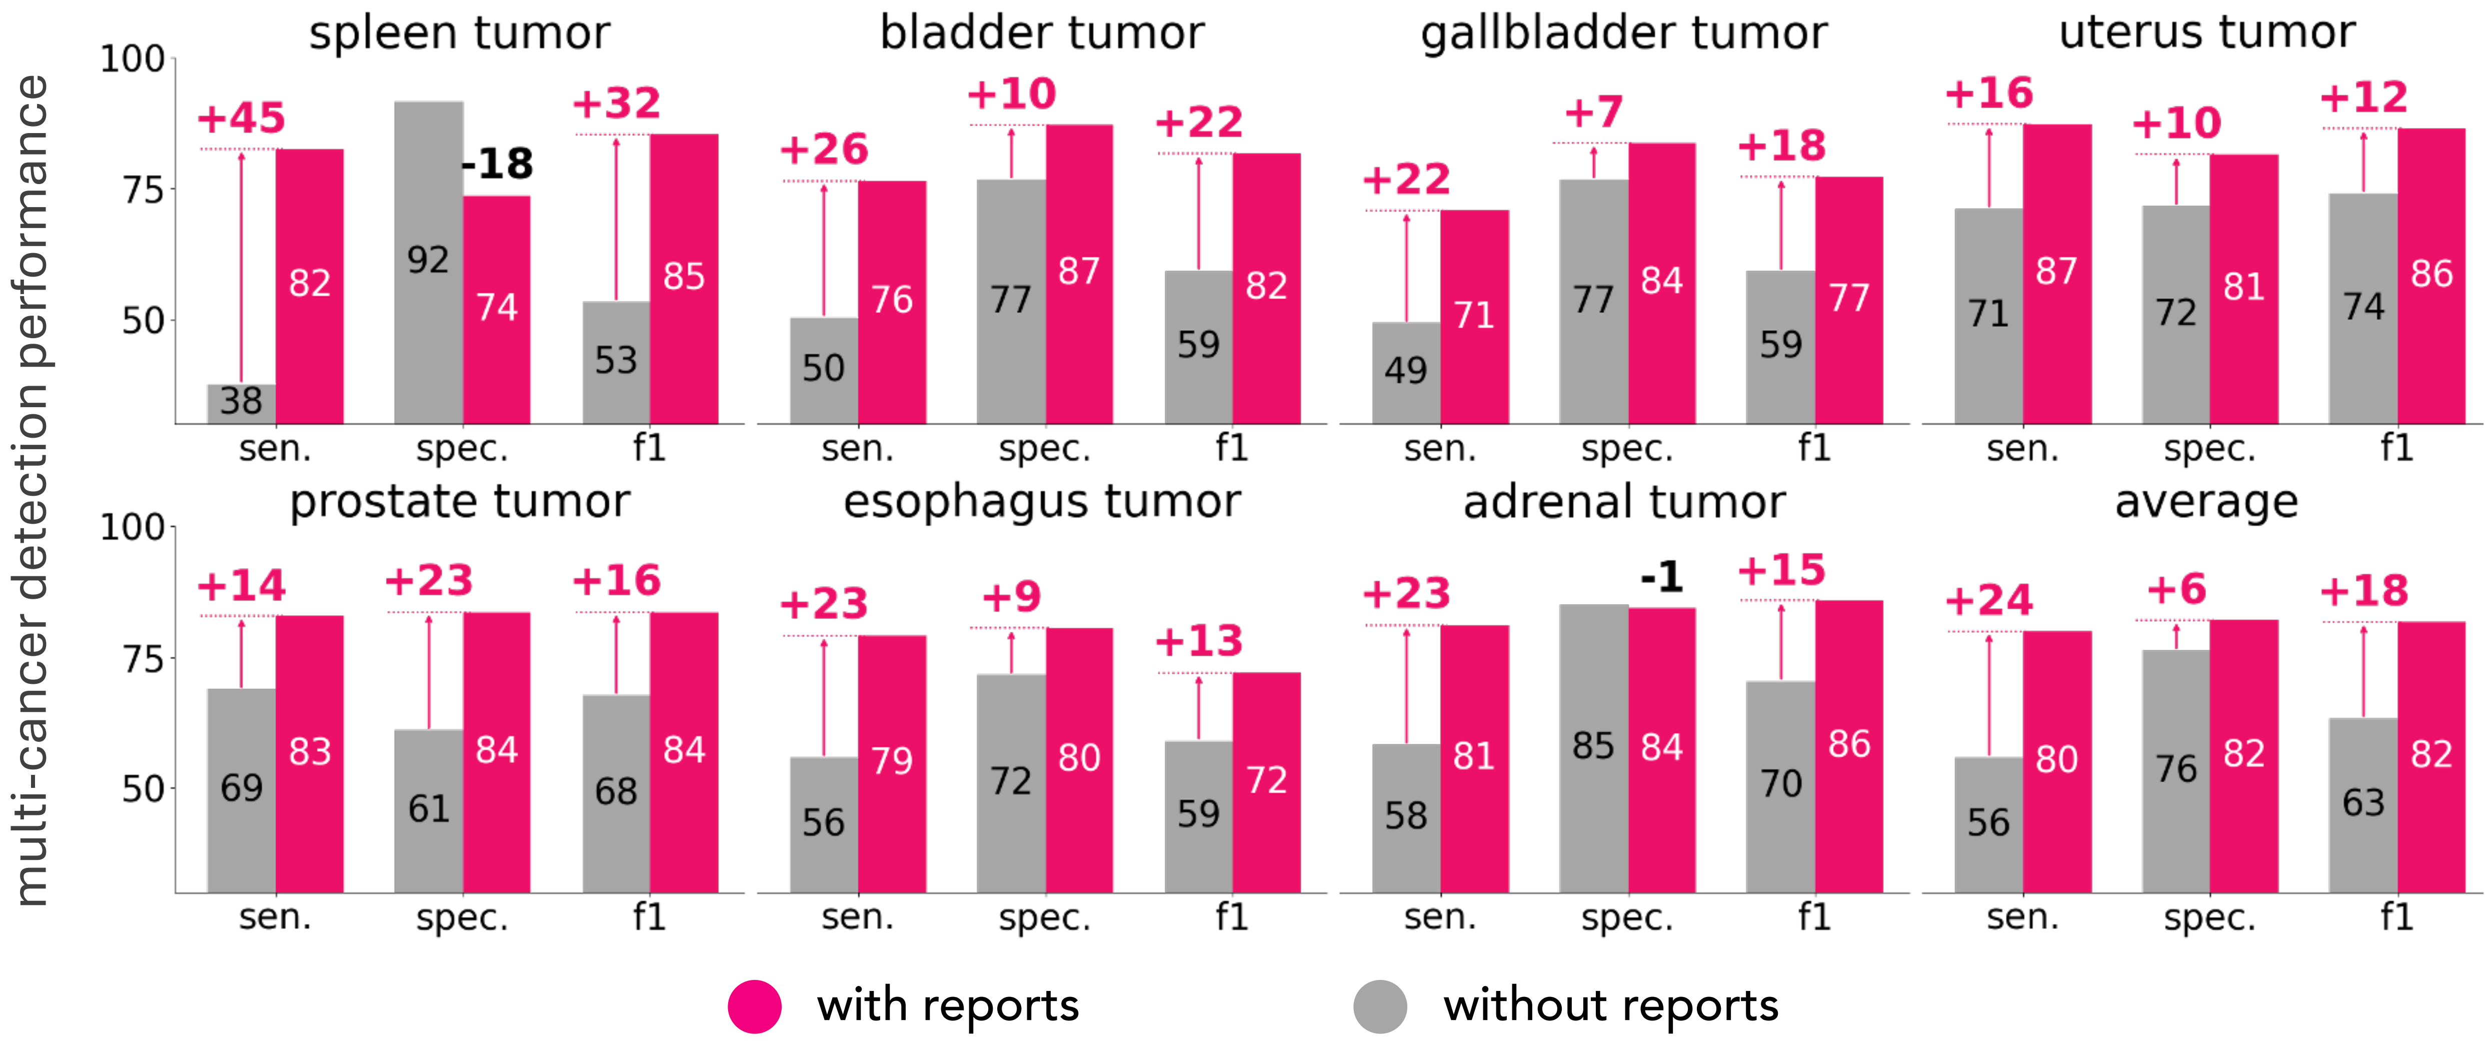

Chapter III.C. Synthetic Data

- There's a huge data gap in medical AI right now, particularly when you have rare diseases, uncommon conditions.
- Early-stage tumor scans are 10–20 times less common than late-stage scans in clinical datasets.
- We don't have enough early-stage tumor scans to train these models; unfortunate these are the tumors we must detect to improve survival.
- Synthetic data can be a big piece of that puzzle (Lai et al., MICCA 2024).
- *Note: synthetic tumors are used for training AI only – not for testing.*

GitHub.com/MrGiovanni/Pixel2Cancer

Synthetic Tumors as Time Machine

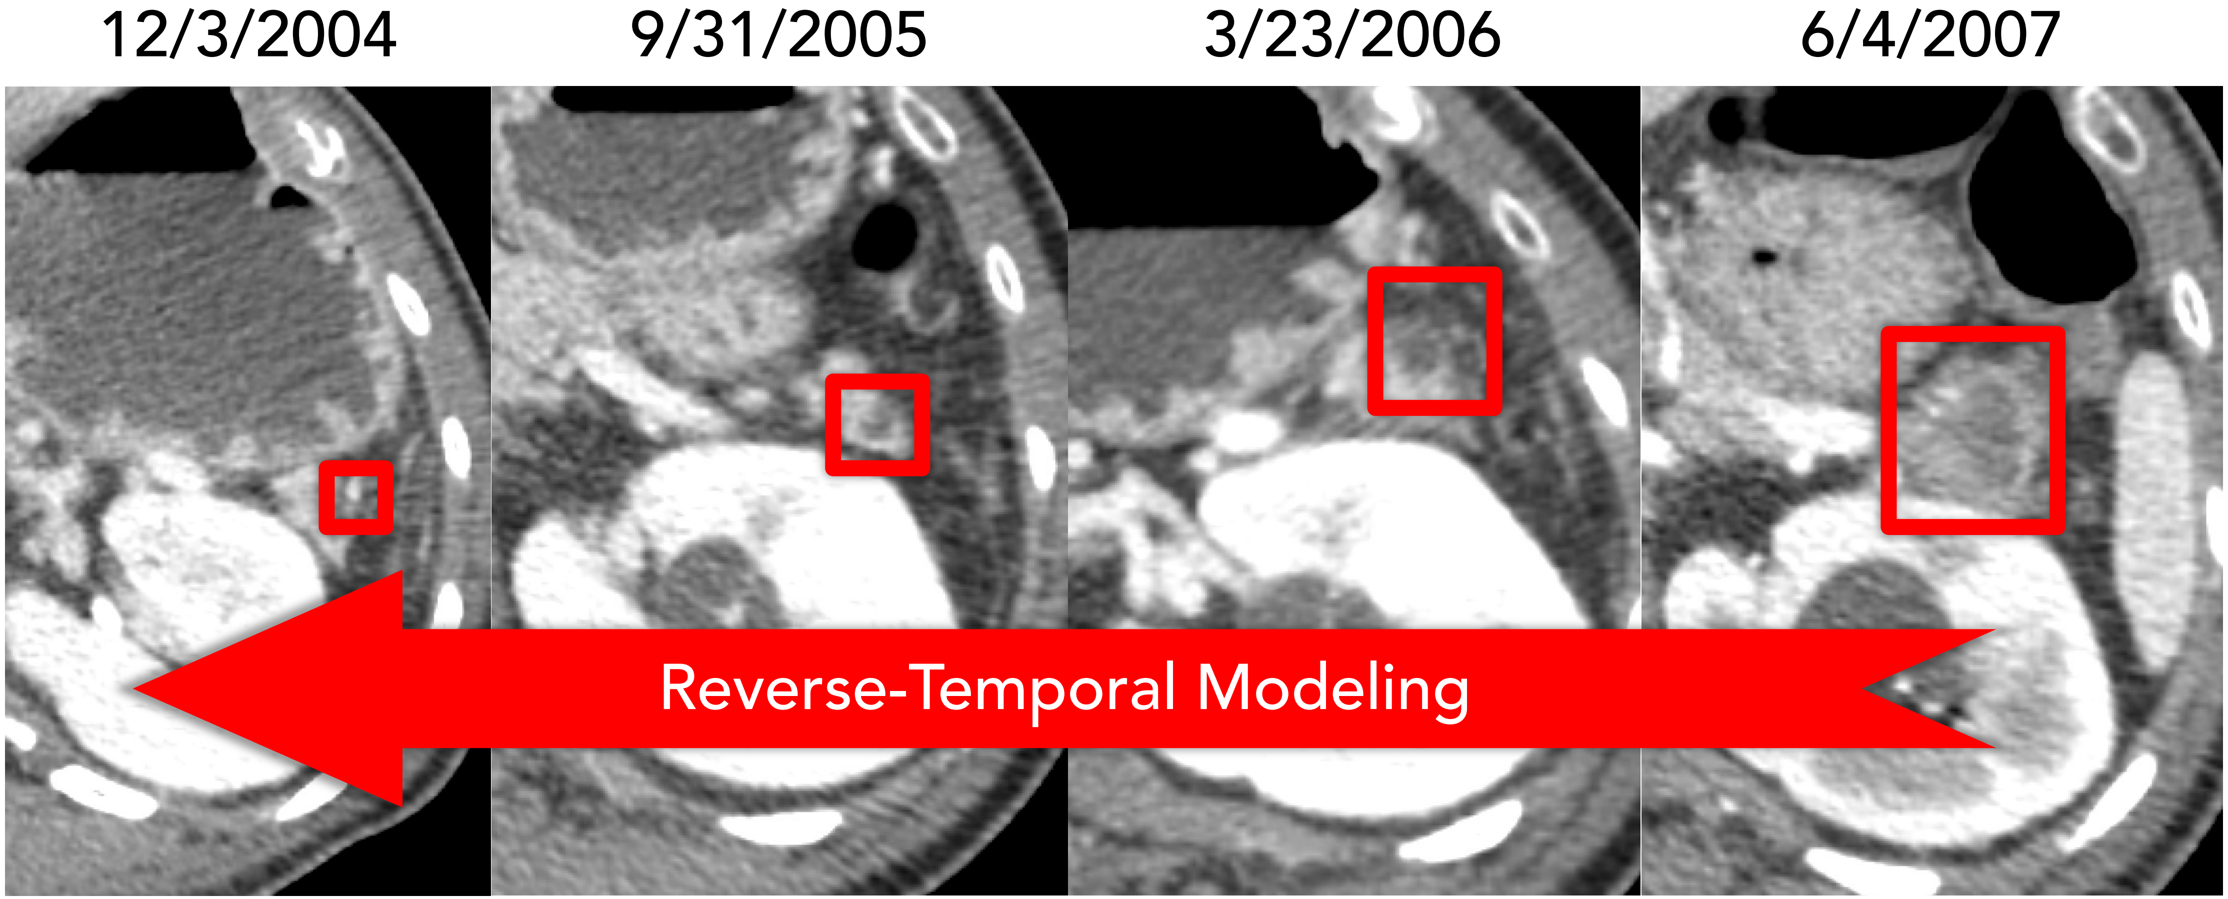

Synthetic Tumors as Time Machine

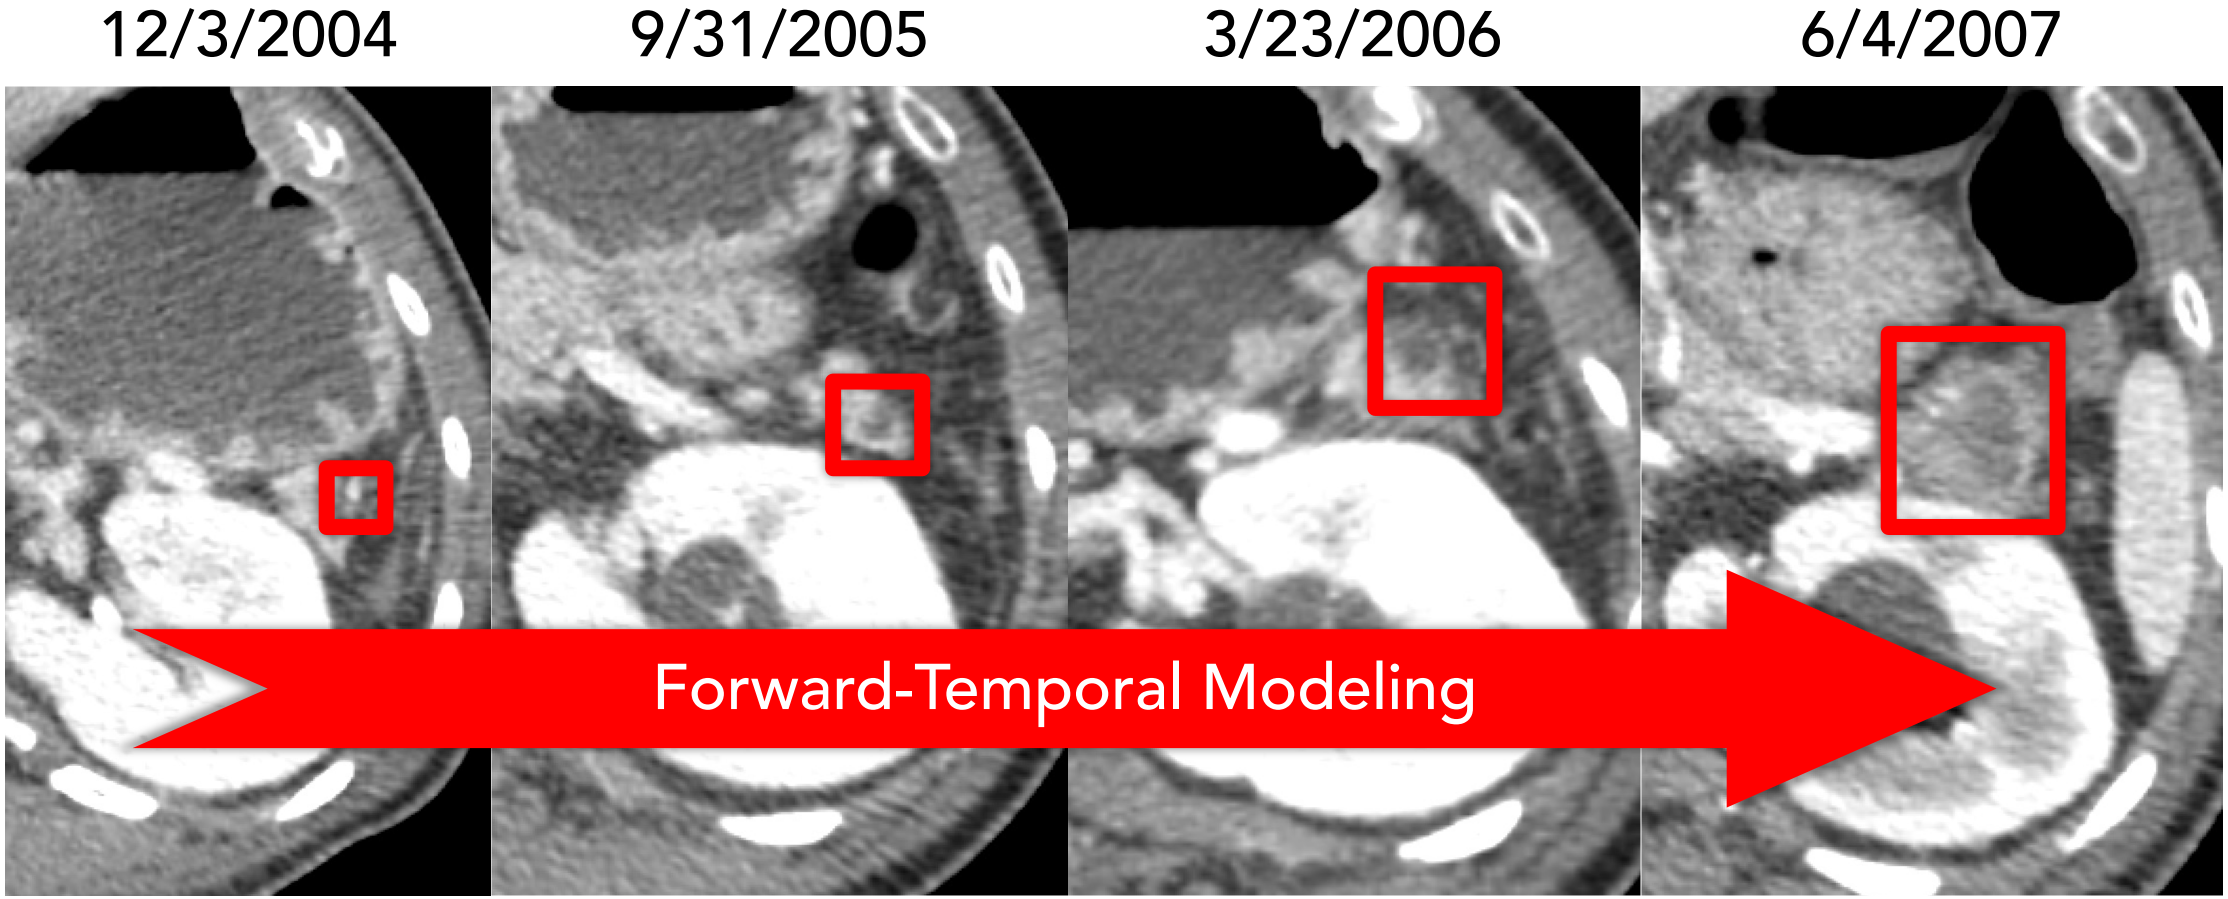

Visual Turing Test for Radiologists

(a) real or fake test

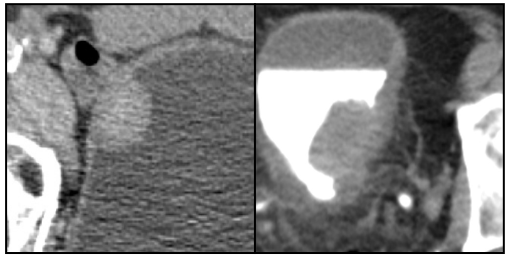

(b) bladder tumor

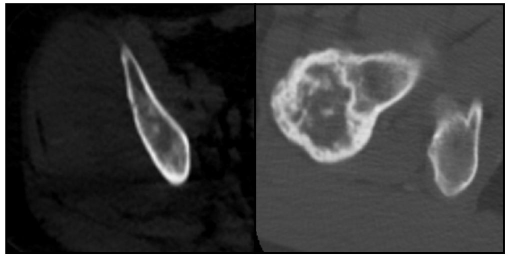

(c) bone tumor

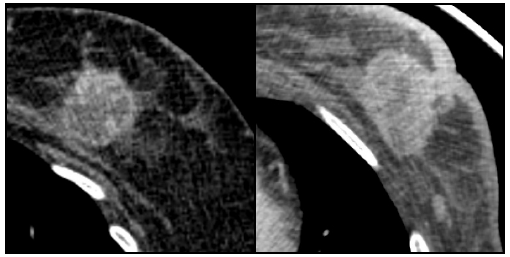

(d) breast tumor

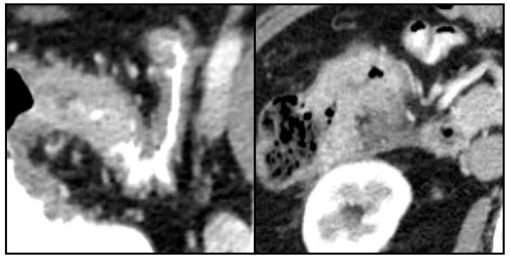

(e) colon tumor

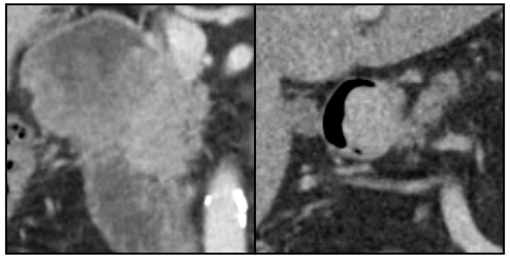

(f) duodenum tumor

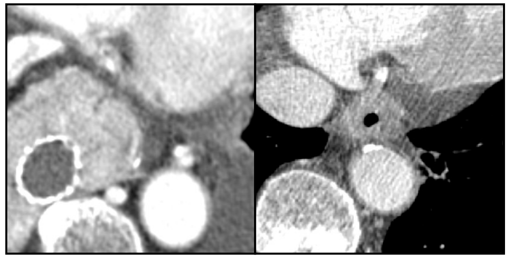

(g) esophagus tumor

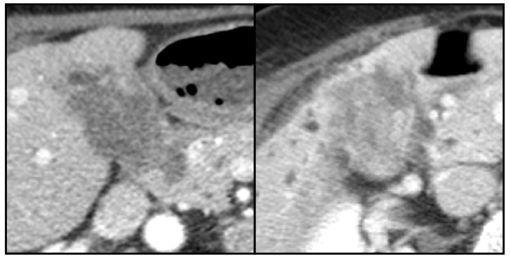

(h) gallbladder tumor

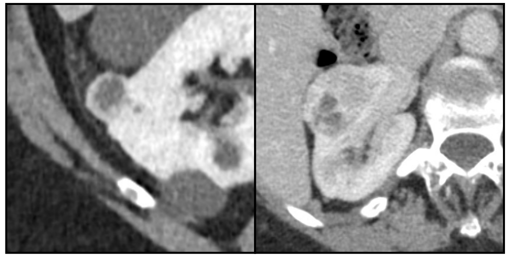

(i) kidney tumor

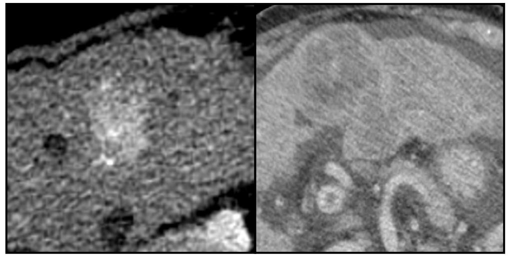

(j) liver tumor

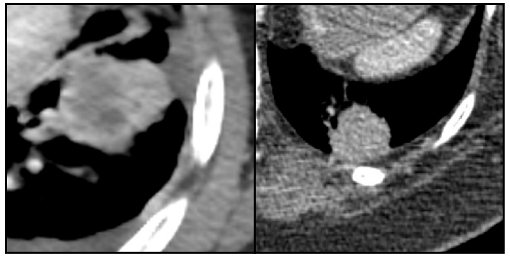

(k) lung tumor

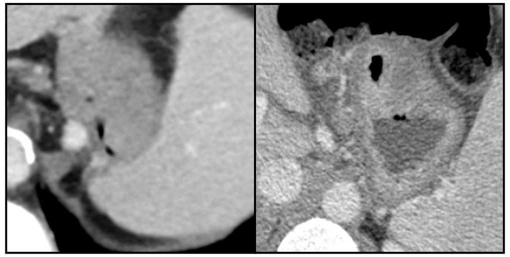

(l) stomach tumor

Tumor/Vessel/Duct/Organ Synthesis

- We developed "game of life" to simulate tumor development (Lai et al., MICCAI 2024) and applied diffusion models to create synthetic tumors.

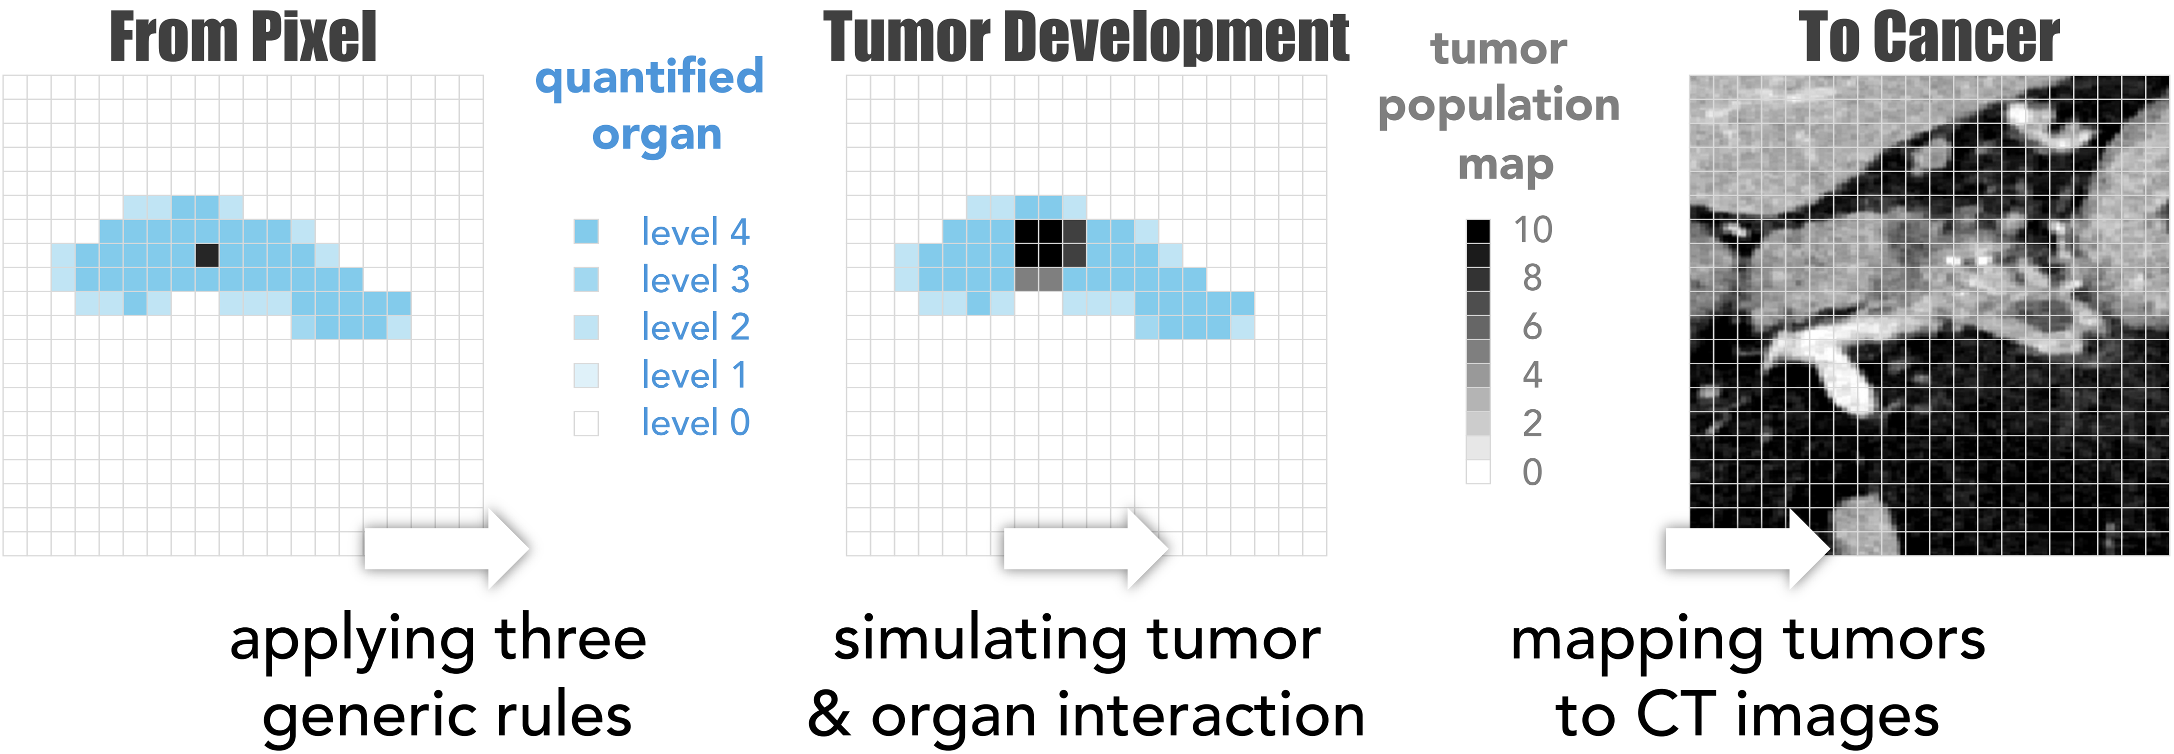

Tumor/Vessel/Duct/Organ Synthesis

Cellular Automata

a mathematical model that uses simple rules to simulate complex systems

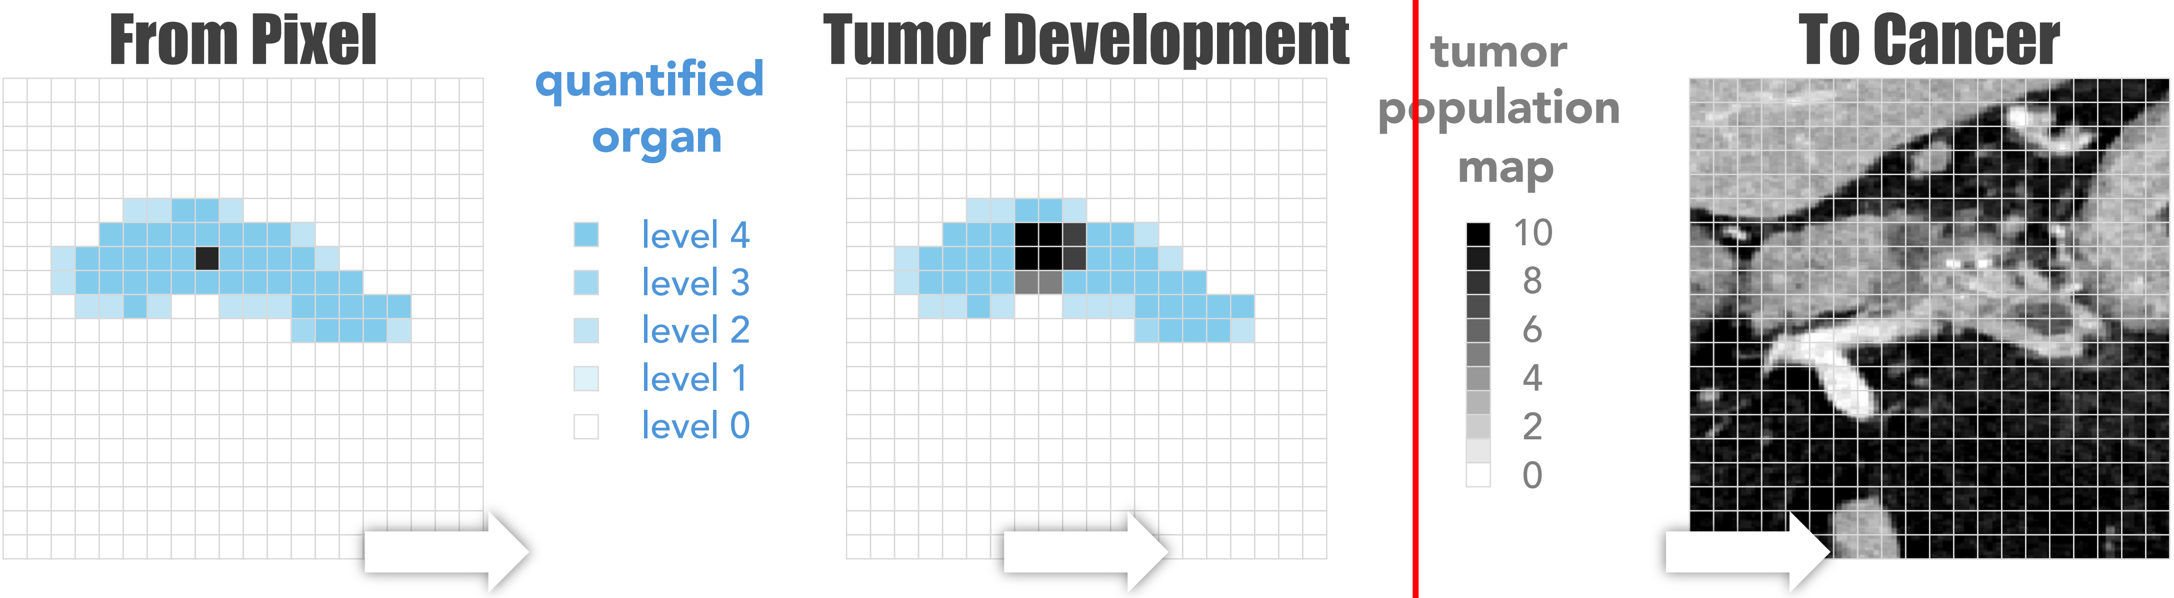

applying three generic rules

simulating tumor & organ interaction

mapping tumors to CT images

Tumor/Vessel/Duct/Organ Synthesis

From Pixel

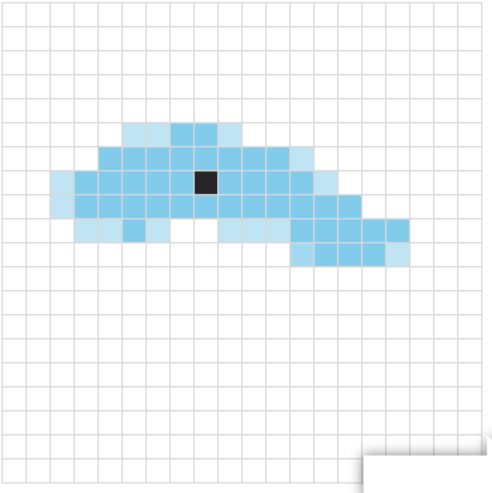

quantified organ

- level 4
- level 3
- level 2
- level 1
- level 0

applying three generic rules

Diffusion Models

conditioned on tumor/vessel/duct/organ shapes simulated by cellular automata

Tumor Development

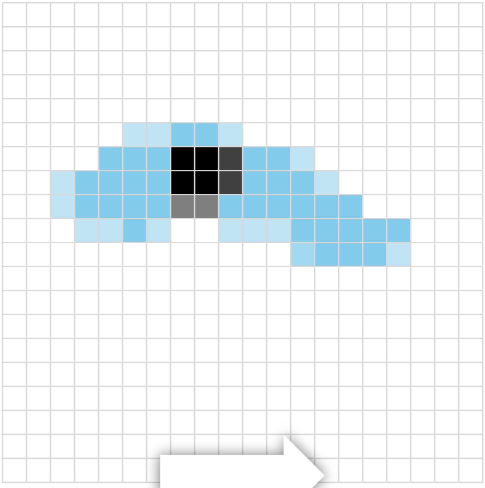

tumor population map

simulating tumor & organ interaction

To Cancer

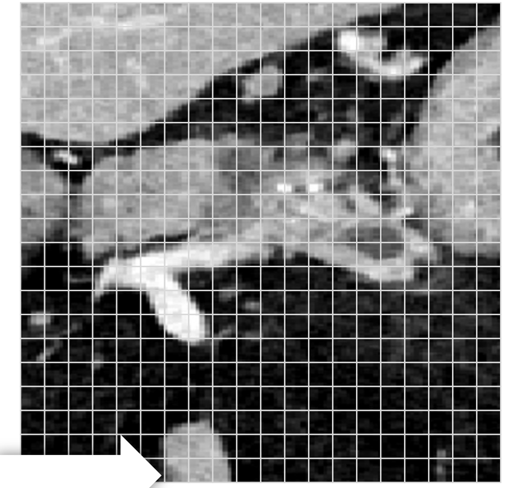

mapping tumors to CT images

Scaling Laws in Synthetic Tumors

- Synthetic data reduces the need for voxel-wise annotated real data from 1,500 down to 500. (Chen et al., ICCV 2025).

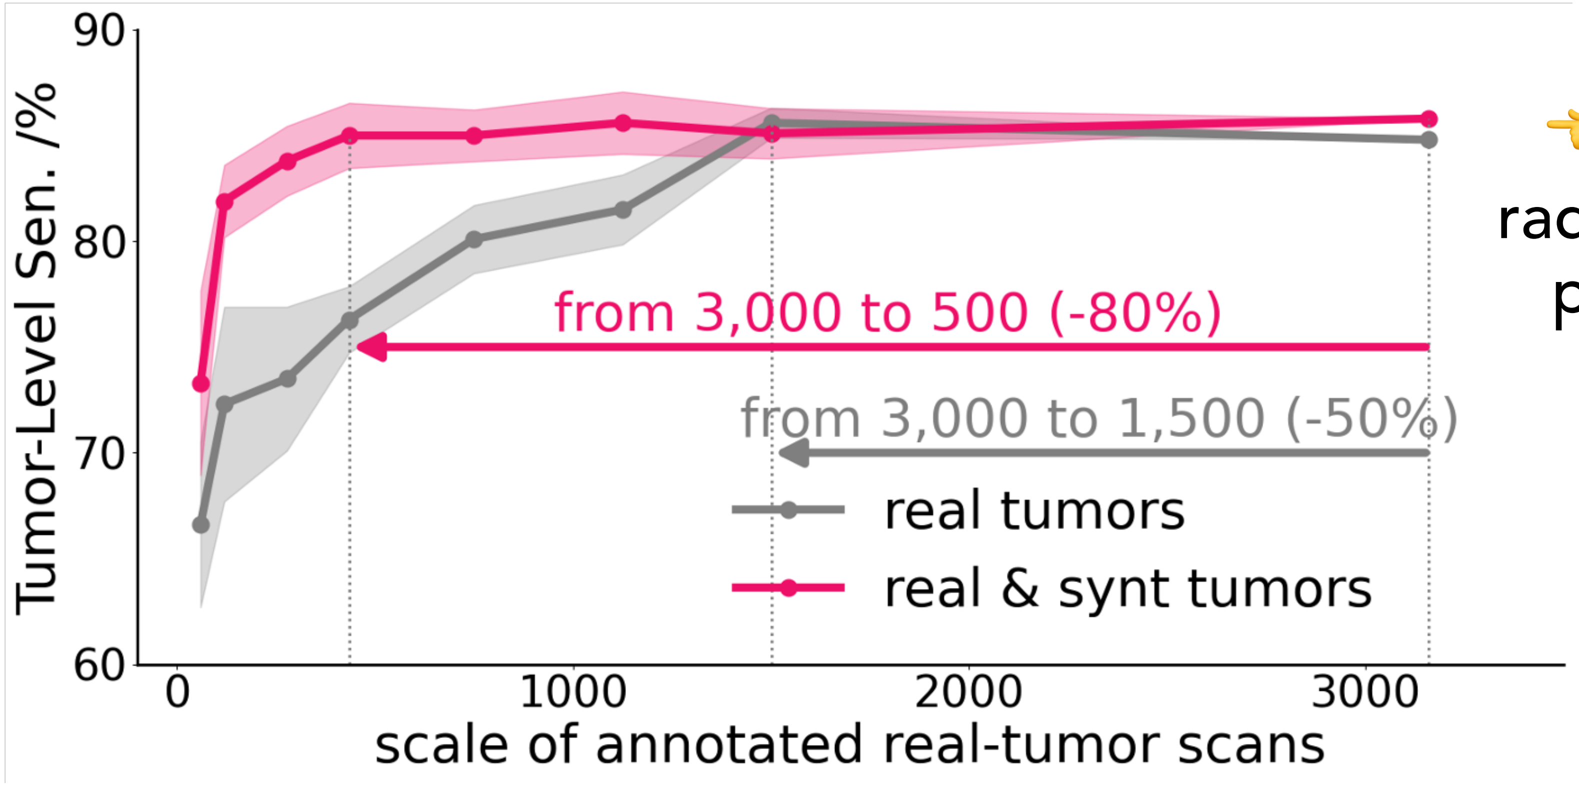

👈 exceeding radiologist-level performance

Synthetic Data Helps Small Tumor Detection

- Synthetic data improves sensitivity of detecting small tumors (≤2 cm) by **5% (89% → 94%)** (Chen et al., CVPR 2024; Hu et al., CVPR 2023)
- The smallest lesion we detected was 2 mm.

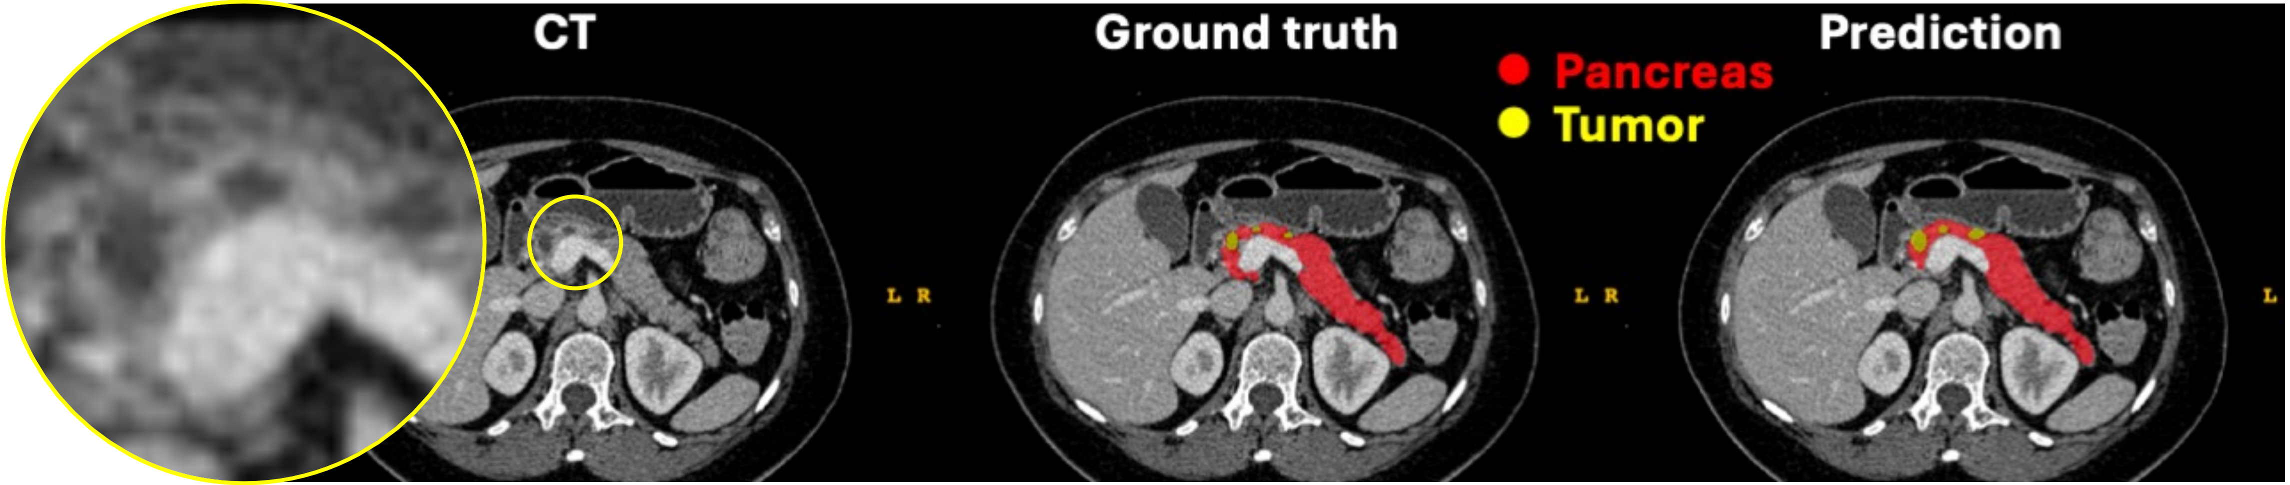

GitHub.com/MrGiovanni/SyntheticTumors

Chapter IV. AI Helps Radiologists

A. Enable early detection (< 2 cm) and even earlier tumor detection with longitudinal, prediagnostic data
B. Assist in writing radiology reports using tumor/organ segmentation
C. Use Large Language Models (LLMs) to automatically retrieve CT scans (15 min → 5 sec per study)

Chapter IV.A. Prediagnostic Detection

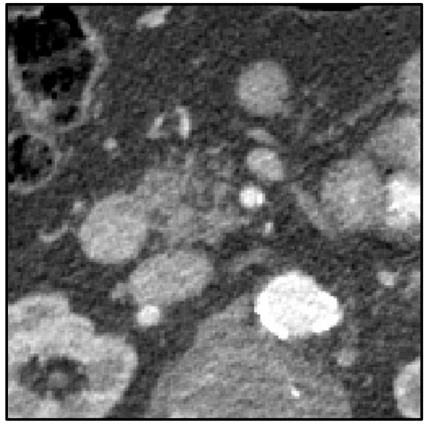

Prediagnostic
CT scans

Diagnostic
CT scans

2021

3-36 months
before diagnosis

2024

8 Month Earlier Than Radiologists

- Our AI algorithms successfully detected pancreatic tumors in **36 of 58** patients (sensitivity = **62%**) that had been overlooked by radiologists (Li et al., arXiv 2025)
- AI detect early tumors with a median lead time of **244** days before clinical diagnosis.

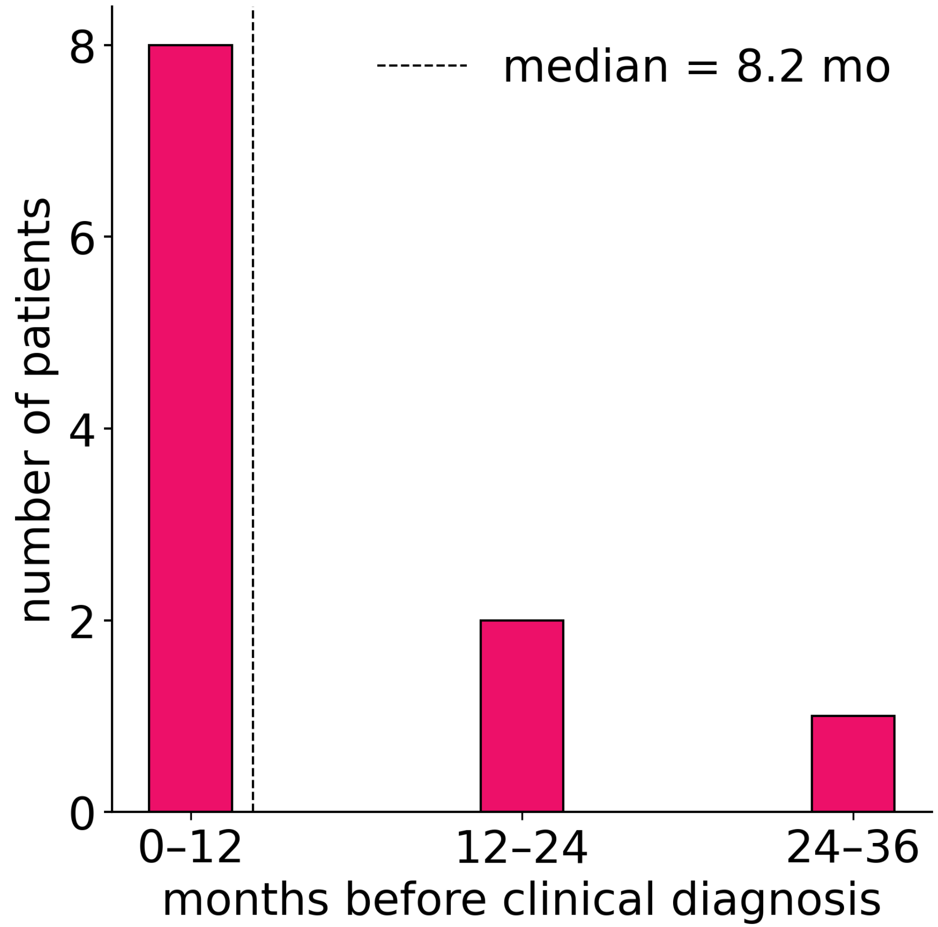

GitHub.com/MrGiovanni/ScaleMAI

Clinical Integration & Multi-reader Study

- A user-friendly, desktop AI system for algorithm running, inference, and visualization (co-developed with **Nvidia**).

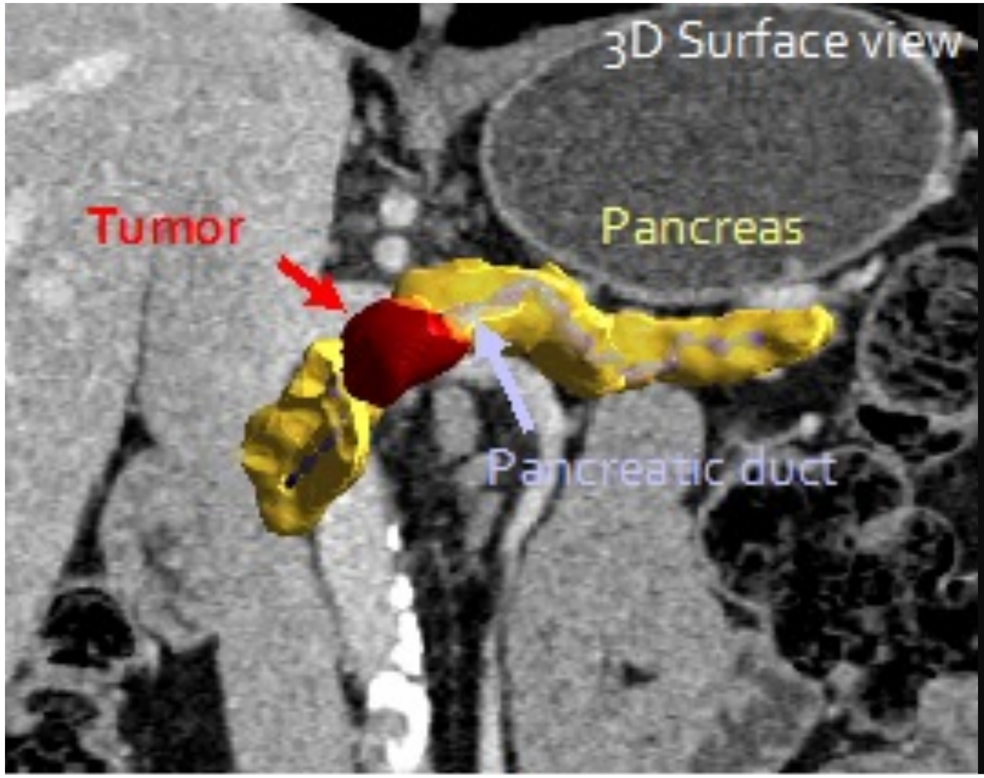

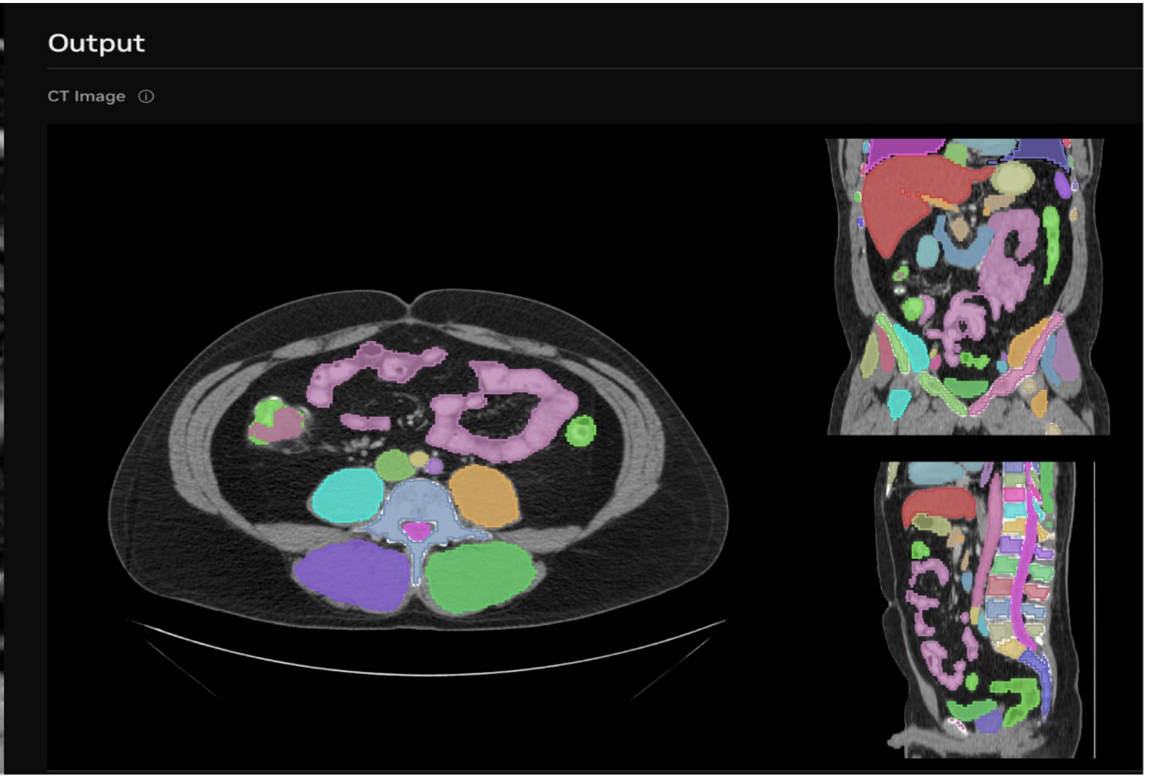

Clinical Integration & Multi-Reader Study

- Comparing with radiologists.
- Few studies report *radiologist performance* on these tasks. Existing studies show low sensitivity (**30–40%**) for detecting small tumors (≤2 cm).
- In our ongoing study, a team of **30** radiologists achieved **34%** sensitivity and **94%** specificity.
- The AI obtained **2x** sensitivity than radiologists with similar specificity.

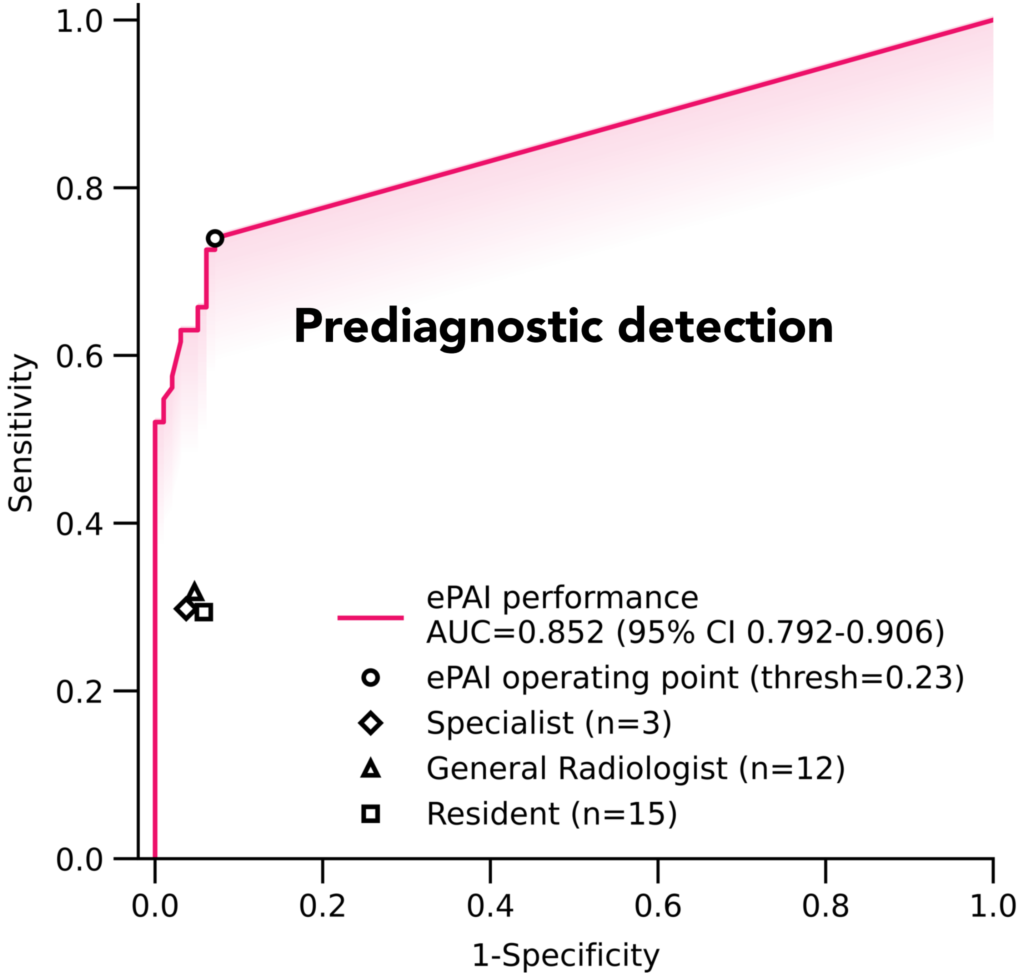

Chapter IV.B Report Generation

- Two major strategies
- (I) End-to-End Vision-Language Modeling
- (II) Segmenting-then-Reporting (Bassi et al., ICCV 2025)

GitHub.com/MrGiovanni/RadGPT

End-to-End Vision-Language Modeling

Black Box
Vision-Language Models

✓ End-to-end
✗ Hallucination

Examples
CLIP · CT-CHAT · CT2Rep · M3D · RadFM · Merlin

Pancreas:
Pancreas is enlarged (volume: 84.6 cm^3).
Mean HU value: 105.7 +/- 33.1.

Pancreas lesions:
Pancreas tumor 1:
Location: pancreas head/body.
Size: 2.9 x 2.2 cm (image 298).
Volume: 8.2 cm^3.
Tumor Stage (T stage): T2.
Enhancement relative to pancreas:
Hypoattenuating (HU value is 52.6 +/- 26.8).

Segmenting-then-Reporting

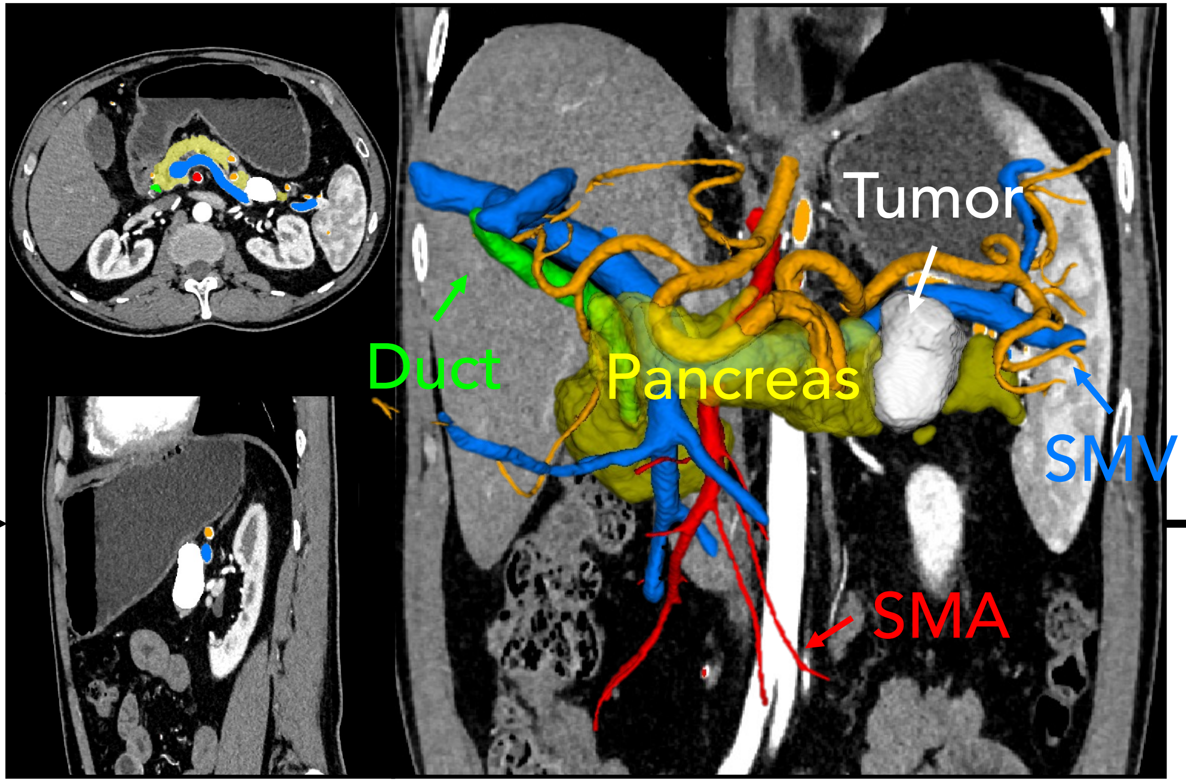

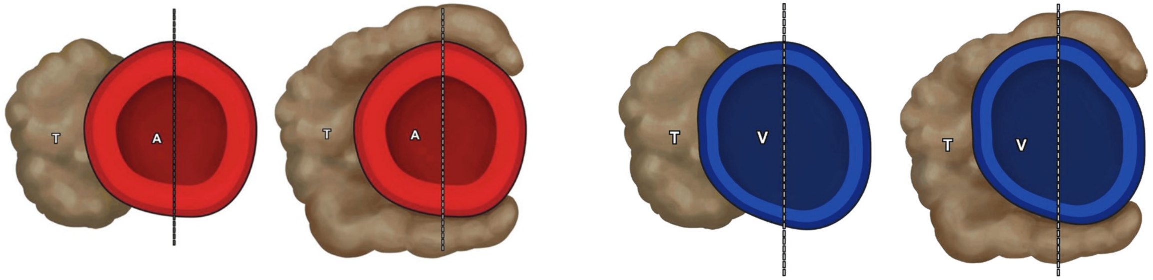

Pancreas:
Pancreas is enlarged (volume: 84.6 cm³).
Mean HU value: 105.7 +/- 33.1.

Pancreas lesions:
Pancreas tumor 1:
Location: pancreas head/body.
Size: 2.9 x 2.2 cm (image 298).
Volume: 8.2 cm³.
Tumor Stage (T stage): T2.
Enhancement relative to pancreas:
Hypoattenuating (HU value is 52.6 +/- 26.8).

Evaluating Tumor Detection in Reports

“End-to-end vision-language” approaches

Model	Pancreatic Tumor (%) Sen. (≤2 cm)	Pancreatic Tumor (%) Sen. (>2 cm)	Pancreatic Tumor (%) Spec.	Liver Tumor (%) Sen. (≤2 cm)	Liver Tumor (%) Sen. (>2 cm)	Liver Tumor (%) Spec.
CT-CHAT	66.7	51.9	61.2	5.7	3.2	94.7
CT2Rep	0.0	0.0	92.5	35.8	49.2	70.4
M3D	0.0	7.4	97.2	9.4	12.7	86.0
RadFM	0.0	0.0	99.9	3.3	5.7	93.9
Merlin	33.3	51.9	71.8	30.2	41.3	95.9
RadGPT (ours)	66.7	81.5	93.2	39.6	96.8	64.4

“Segmenting-then-reporting” approach

GitHub.com/MrGiovanni/RadGPT

Automated Report Generation

Pancreas:
Pancreas is enlarged (volume: 84.6 cm^3).
Mean HU value: 105.7 +/- 33.1.
Pancreas lesions:
Pancreas tumor 1:
Location: pancreas head/body.
Size: 2.9 x 2.2 cm (image 298).
Volume: 8.2 cm^3.
Tumor Stage (T stage): T2.
Enhancement relative to pancreas:
Hypoattenuating (HU value is 52.6 +/- 26.8).

MEDICAL REPORT

PATIENT INFORMATION
BDMAP ID: BDMAP_00000037
Sex: N/A
Age: N/A

IMAGING DETAIL
spacing: [0.9 0.9 5.0]
scanner: N/A
shape: (512, 512, 44)
contrast: N/A

AI MEASUREMENTS

	organ Volume (cc)	total lesion #	total lesion volume (cc)
liver	2107.6	1	193.4
pancreas	109.3	1	0.1
kidney	282.9	1	70.7

NARRATIVE REPORT
The patient has a liver mass located in hepatic segment 5, measuring 7.7 x 6.8 cm (image 13), with a volume of 193.4 cc. The lesion is hyperattenuating, with a mean HU value of 98.3 +/- 23.4.
The pancreas is enlarged, with a volume of 109.3 cc and a mean HU value of 71.8 +/- 34.4. There are multiple pancreatic lesions, including a hypoattenuating mass in the pancreas head, measuring 0.6 x 0.2 cm (image 24), with a volume of 0.1 cc. This lesion is consistent with a biopsy-proven pancreatic ductal adenocarcinoma (PDAC). The tumor does not contact the SMA, aorta, portal vein, SMV, IVC, CA, CHA, or SA, and is considered resectable. Additionally, there are two hyperattenuating masses in the pancreas, consistent with pancreatic neuroendocrine tumors (PNETs). The larger mass measures 5.3 x 3.9 cm (image 33), with a volume of 32.7 cc, and the smaller mass measures 0.8 x 0.7 cm (image 31), with a volume of 0.4 cc. Both tumors do not contact the SMA, aorta, portal vein, SMV, IVC, CA, CHA, or SA, and are considered resectable.
There is also a hypoattenuating cystic lesion in the pancreas, measuring 0.5 x 0.4 cm (image 24), with a volume of 0.1 cc.
The spleen is normal in size, with a volume of 234.0 cc and a mean HU value of 124.6 +/- 32.9.
The kidneys are enlarged, with a total volume of 534.0 cc. There is a large hypoattenuating mass in the left kidney, measuring 5.6 x 5.0 cm (image 9), with a volume of 70.7 cc.
IMPRESSION:
The patient has a large hyperattenuating liver mass, an enlarged pancreas with multiple lesions including a biopsy-proven PDAC and two PNETs, and enlarged kidneys with a large hypoattenuating mass in the left kidney.

STRUCTURED REPORT
CT Venous Phase
FINDINGS:

KEY IMAGES

LIVER TUMORS

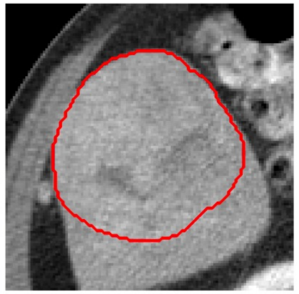

PANCREAS TUMORS

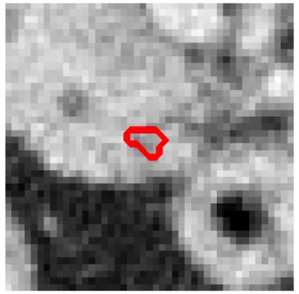

KIDNEY TUMORS

Chapter IV.C LLM-Enabled CT Retrieval

- Large language models (LLMs) analyze radiology reports to identify tumor cases and retrieve pre-diagnostic scans from 3–36 months prior.
- Radiologists need **15 minutes** to identify a suitable tumor scan.
- Our developed LLMs make it less than **5 seconds (180x)**, and even faster with more powerful computers.
- Scalable Medical Artificial Intelligence (ScaleMAI) accelerates the development of trusted datasets and AI models (Li et al., arXiv 2025)


GitHub.com/MrGiovanni/ScaleMAI

Chapter IV.C LLM-Enabled CT Retrieval

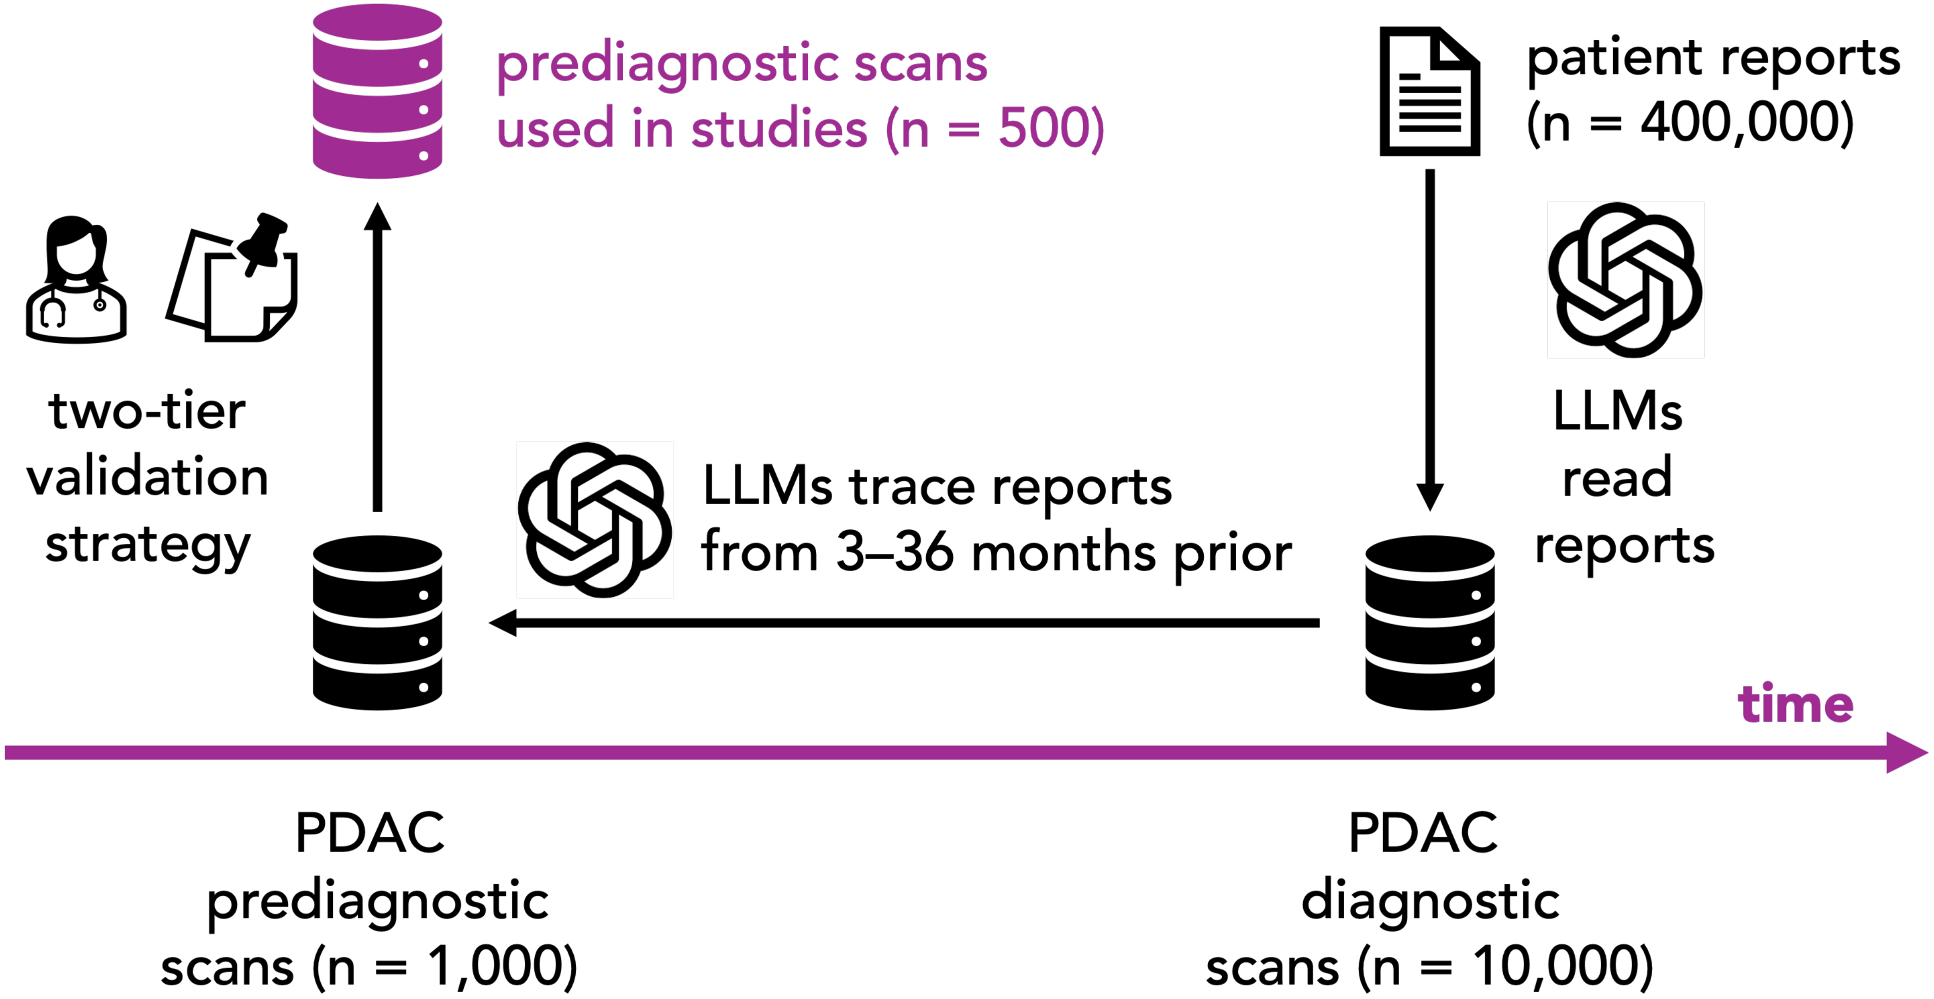

Other Projects & Future Work

- Even bigger datasets to ensure algorithms perform well at most institutions (1-10 Million).
- Multi-model – CT, MRI, Histopathology and risk factors (AI studies of how best to combine them).
- Understanding cancer development prior to diagnosis through causal models and world models that incorporate causality and treatment. (Yang et al., ICCV 2025)
- Pre-diagnostic detection – years before conventional detection in all abdominal organs.

Multimodal + Longitudinal Analysis

2018 2021 2025 2030

Reverse- and Forward Temporal Modeling

Retrieving paired data

Images: CT/MRI/PET/WSI
Labels: Voxel-Wise Annotations
Reports: Rad/Path Reports
EHR: Patient Information
Therapy: Treatment Record

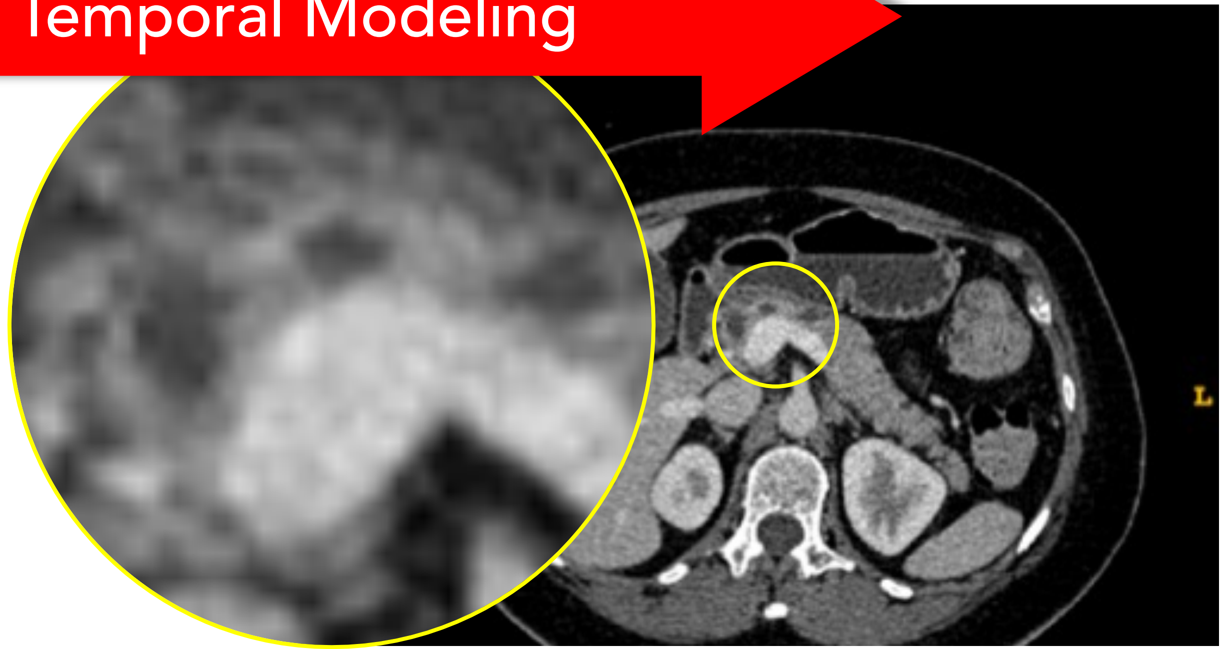

Multimodal + Longitudinal Analysis

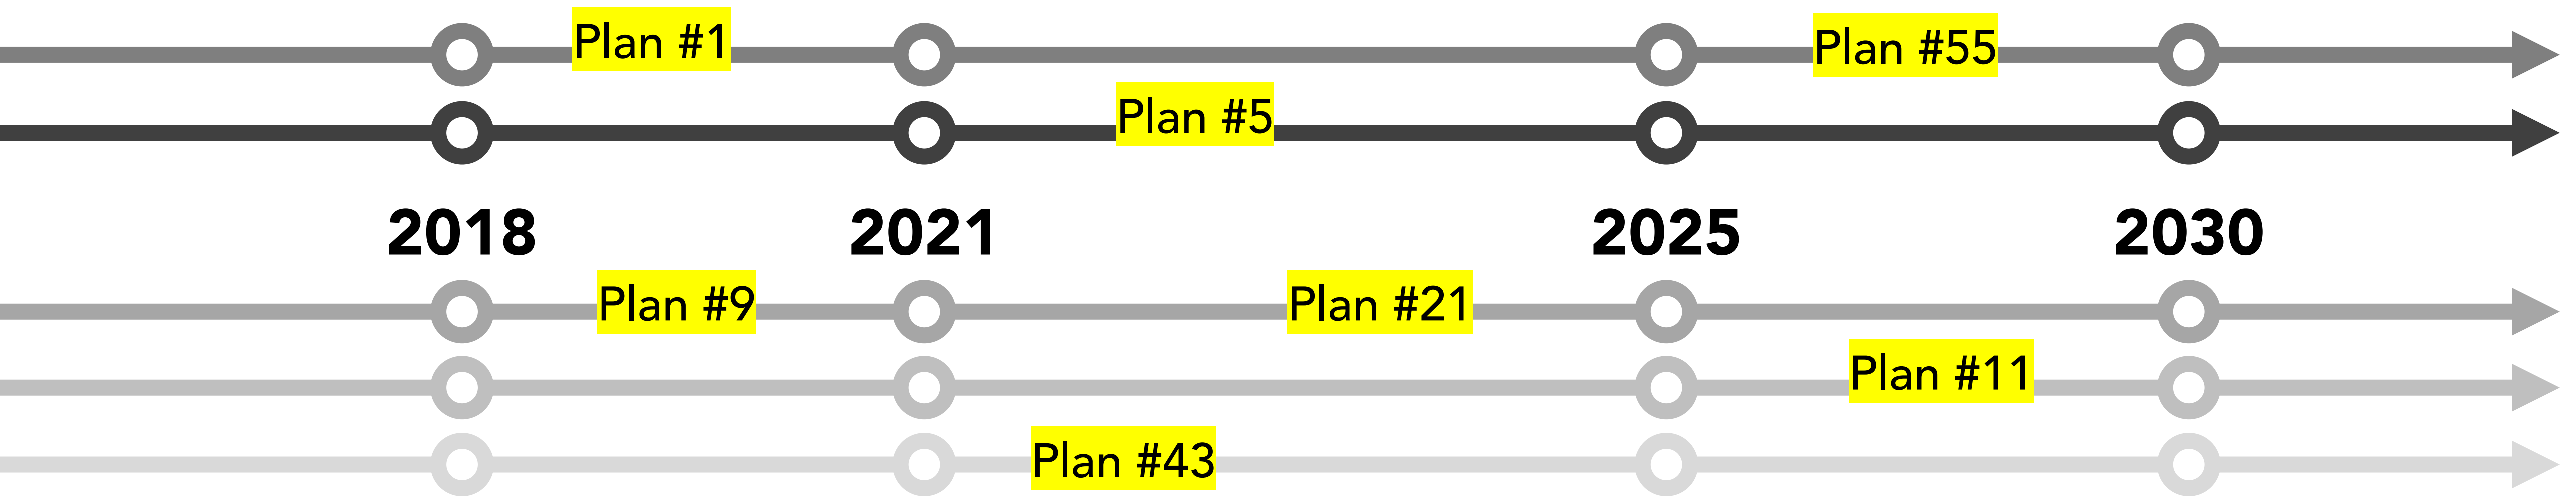

Retrieving paired multimodal, longitudinal data for training
Images: CT/MRI/PET/WSI
Labels: Voxel-Wise Annotations
Reports: Rad/Path Reports
EHR: Patient Information
Therapy: Treatment Record

Multimodal + Longitudinal Analysis

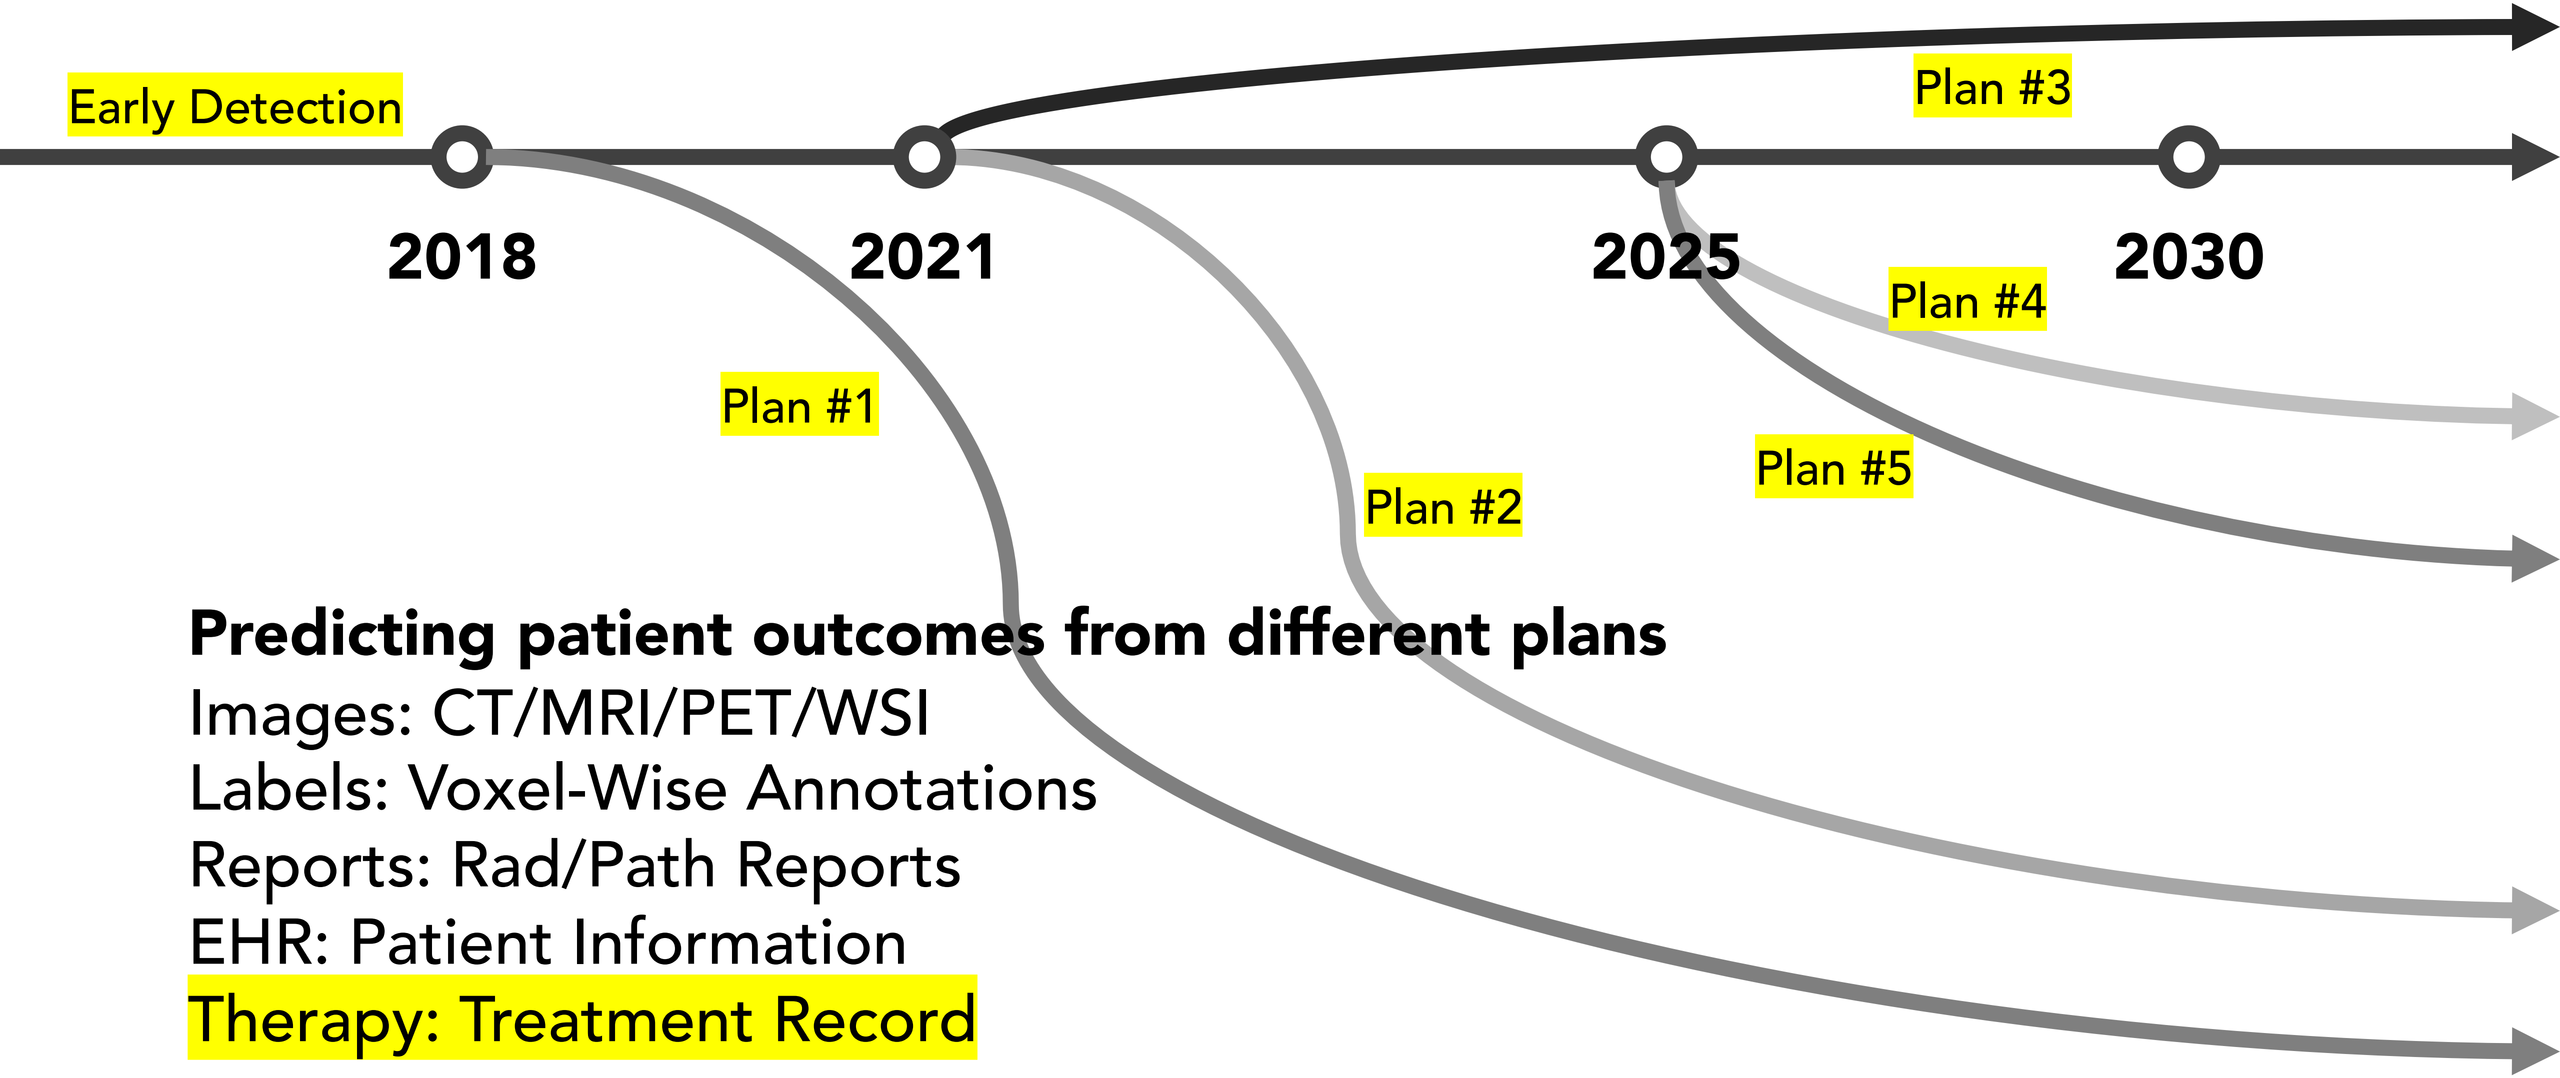

Predicting patient outcomes from different plans

Images: CT/MRI/PET/WSI

Labels: Voxel-Wise Annotations

Reports: Rad/Path Reports

EHR: Patient Information

Therapy: Treatment Record

Multimodal + Longitudinal Analysis

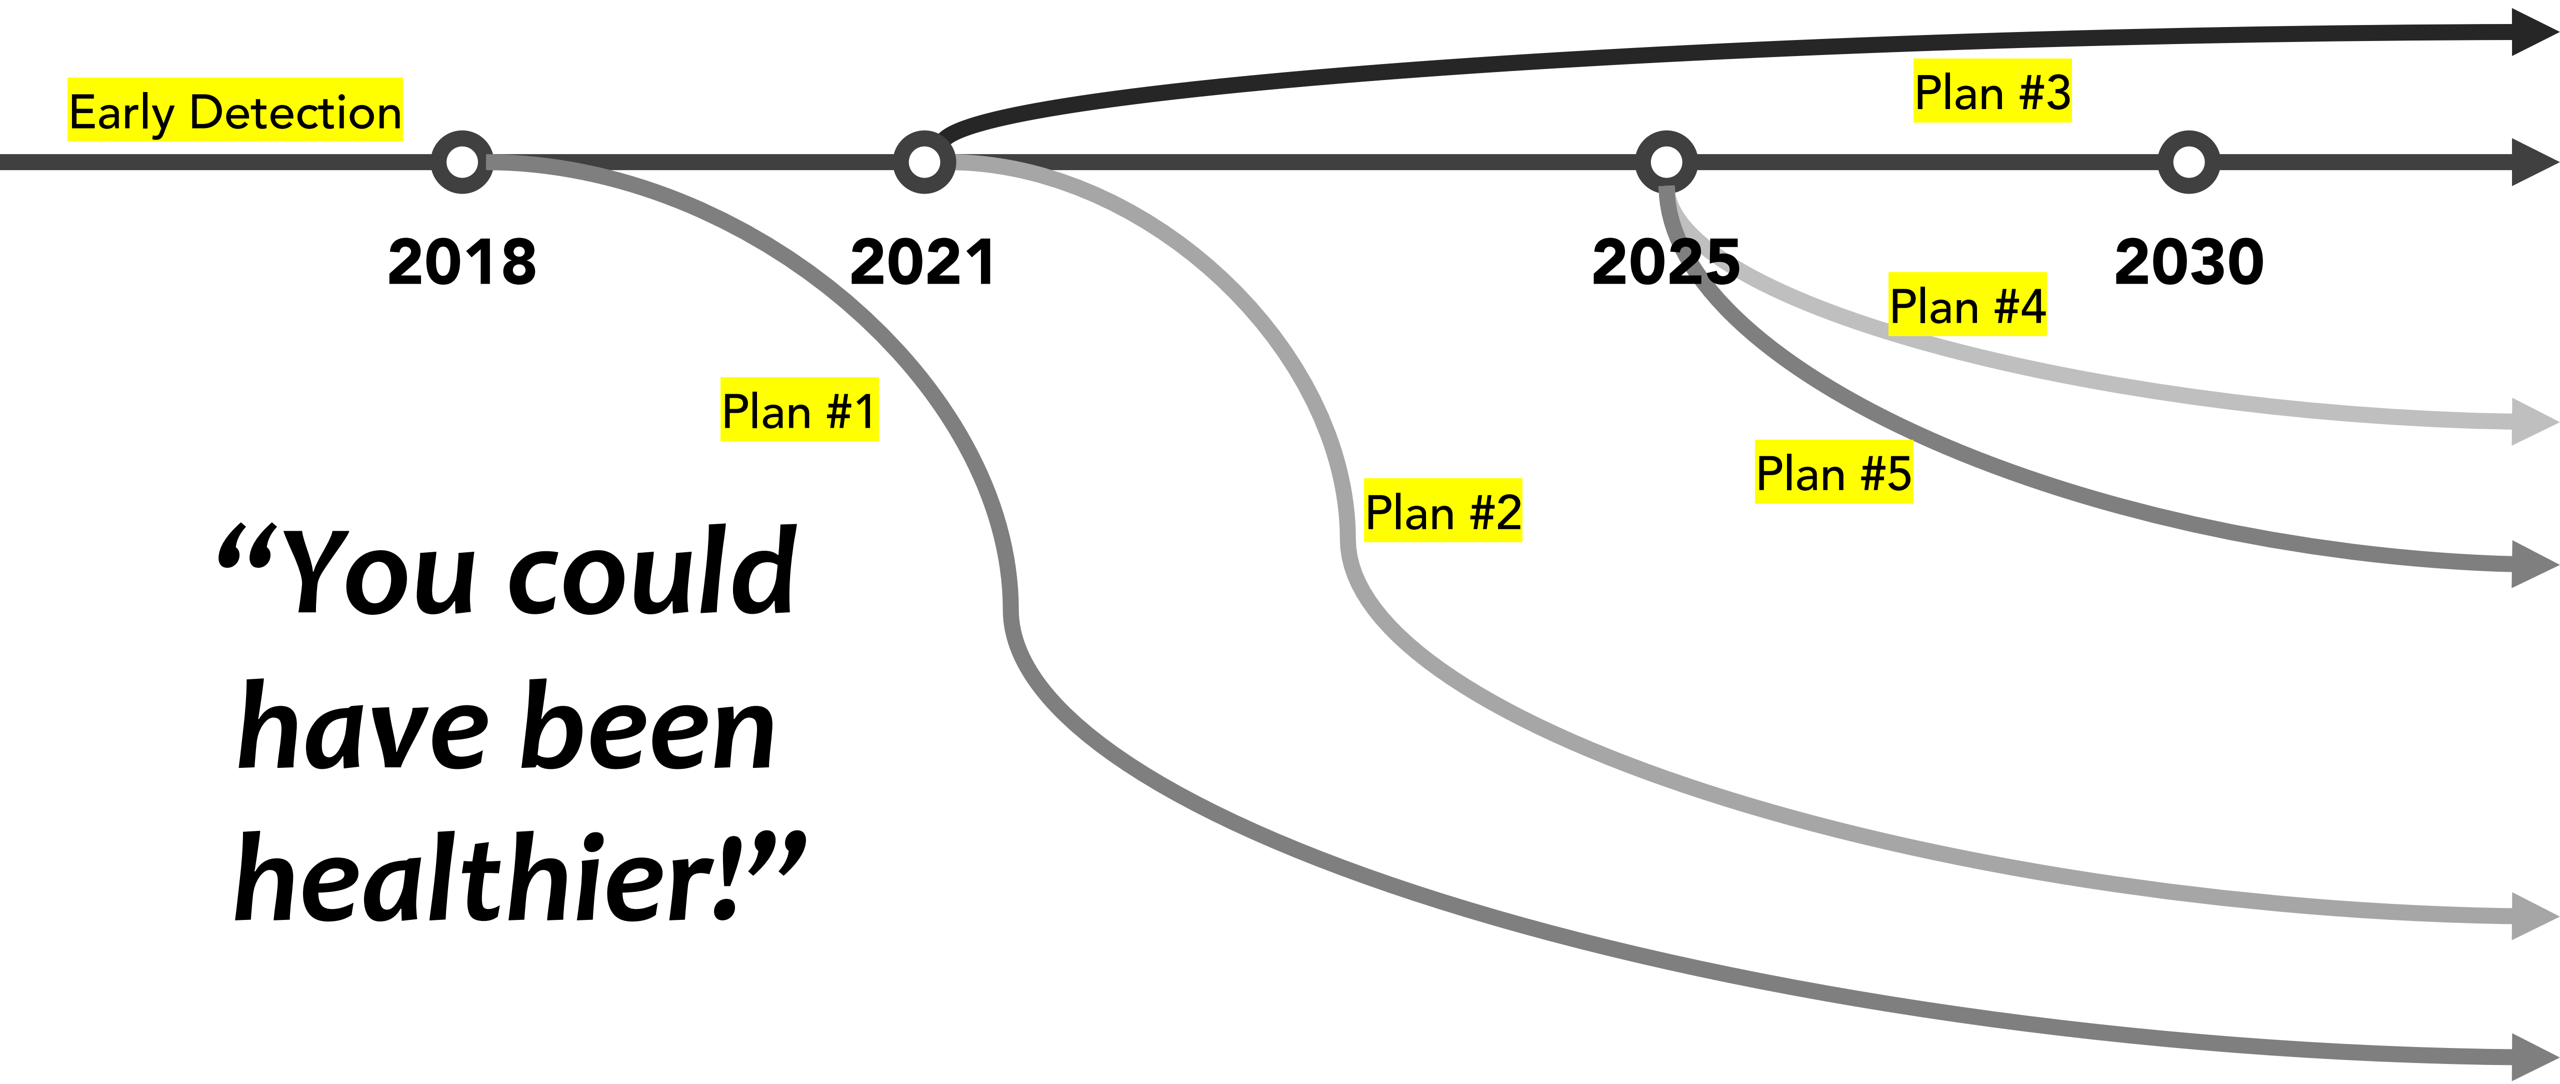

"You could have been healthier!"


The One Possibility
Over 14 Million Futures

Report Generation

Early Cancer Detection

Organ Segmentation

Patient Retrieval

Applications

Algorithms

Research Topics	Featured Achievements	Peer-Reviewed Publications
I. Segmentation Architectures	UNet++, 15,000 Citations	TMI, MIA, ICCV, NeurIPS, RSNA
II. Annotations with Active Learning	AbdomenAtlas, N = 240,000	MIAx2, CVPRx2, MICCAI, NeurIPS, RSNAx12
III.A. Medical Foundation Models	Models Genesis, MICCAI Best Paper Award & MIA Best Paper Award	TPAMI, TMI, MIA, CVPRx2, ICLRx2, ICCV, MICCAIx4, RSNAx12
III.B. Vision & Language	Finalist, MICCAI Best Paper Award	ICCV, MICCAI, ISBI, RSNAx2
III.C. Tumor Synthesis & Generation	Segmentation of 16 Cancer Types	CVPRx2, ICCVx3, MICCAI, RSNAx10

Key References

- Bassi, Pedro RAS, Wenxuan Li, ..., Alan Yuille, and **Zongwei Zhou**. "Touchstone benchmark: Are we on the right way for evaluating ai algorithms for medical segmentation?" *NeurIPS, 2024.*
- Bassi, Pedro RAS, …, **Zongwei Zhou**. "Learning Segmentation from Radiology Reports." *MICCAI, 2025* (Runner-up, Best Paper Award).
- Bassi, Pedro RAS, …, **Zongwei Zhou**. "RadGPT: Constructing 3D Image-Text Tumor Datasets." *ICCV, 2025.*
- Chen, Qi, …, and **Zongwei Zhou**. "Towards generalizable tumor synthesis." *CVPR, 2024.*
- Lai, Yuxiang, …, **Zongwei Zhou**. "From pixel to cancer: Cellular automata in computed tomography." *MICCAI, 2024.*
- Li, Wenxuan, …, **Zongwei Zhou**. "How well do supervised 3d models transfer to medical imaging tasks?" *ICLR, 2025* (Oral).
- Li, Wenxuan, …, **Zongwei Zhou**. "PanTS: The Pancreatic Tumor Segmentation Dataset." NeurIPS, 2025.
- Li, Wenxuan, …, **Zongwei Zhou**. "Abdomenatlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking." *MEDIA, 2024.*
- Lubonja, Ariel, …, **Zongwei Zhou**. "Auditing Significance, Metric Choice, and Demographic Fairness in Medical AI Challenges." *MLMI, 2025.*
- Xia, Yingda, Qihang Yu, …, **Zongwei Zhou**, …, Alan Yuille. "The felix project: Deep networks to detect pancreatic neoplasms." *medRxiv, 2022.*