BodyMaps

Towards 3D Atlas of Human Body

Zongwei Zhou

Johns Hopkins University

zzhou82@jh.edu

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation** (5 min) – We annotated 25 anatomical structures in 9,262 CT volumes
   - NeurIPS-2023 – Semi-automated annotation
   - Three versions: AbdomenAtlas 1.0, 1.1, Pro
2. **Dataset** (10min) – We developed data synthesis strategies to enrich tumor examples
   - CVPR-2023 – Liver tumor synthesis
   - CVPR-2024 – Liver, pancreatic, & kidney tumor synthesis
   - Next version: AbdomenAtlas X
3. **Algorithm** (10min) – We released a set of supervised pre-trained vision-language models
   - ICCV-2023 & MICCAI-2023 – Vision-language models
   - ICLR-2024 – Supervised pre-training
4. **Ending** (2 min) – A large, open algorithmic benchmark

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation** (5 min) – We annotated 25 anatomical structures in 9,262 CT volumes
    - NeurIPS-2023 – Semi-automated annotation
    - Three versions: AbdomenAtlas 1.0, 1.1, Pro

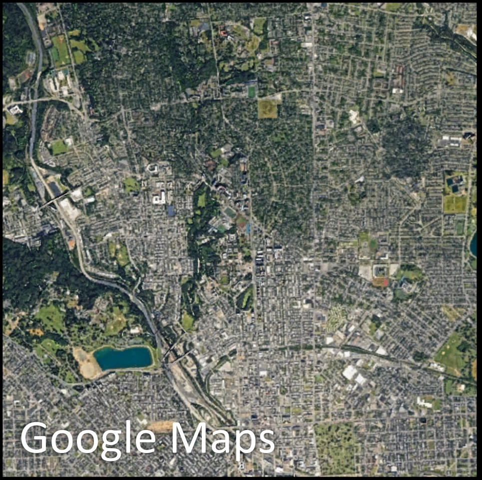

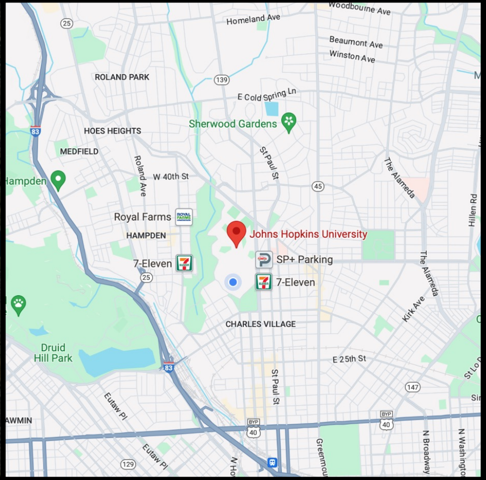

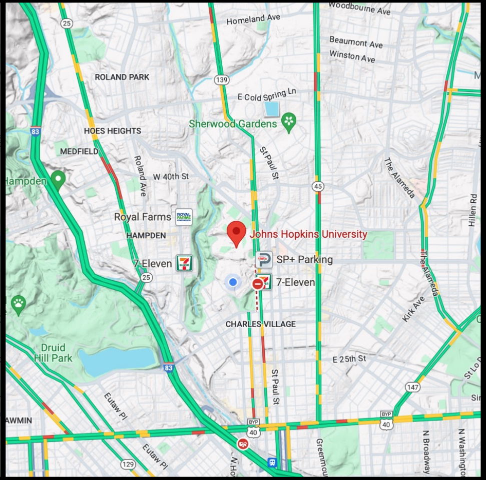

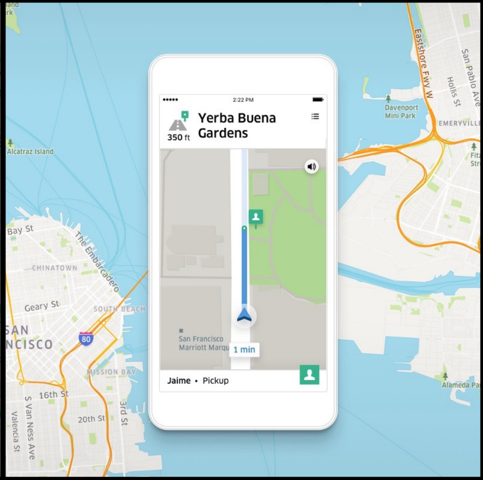

Data → Information → Knowledge → Application

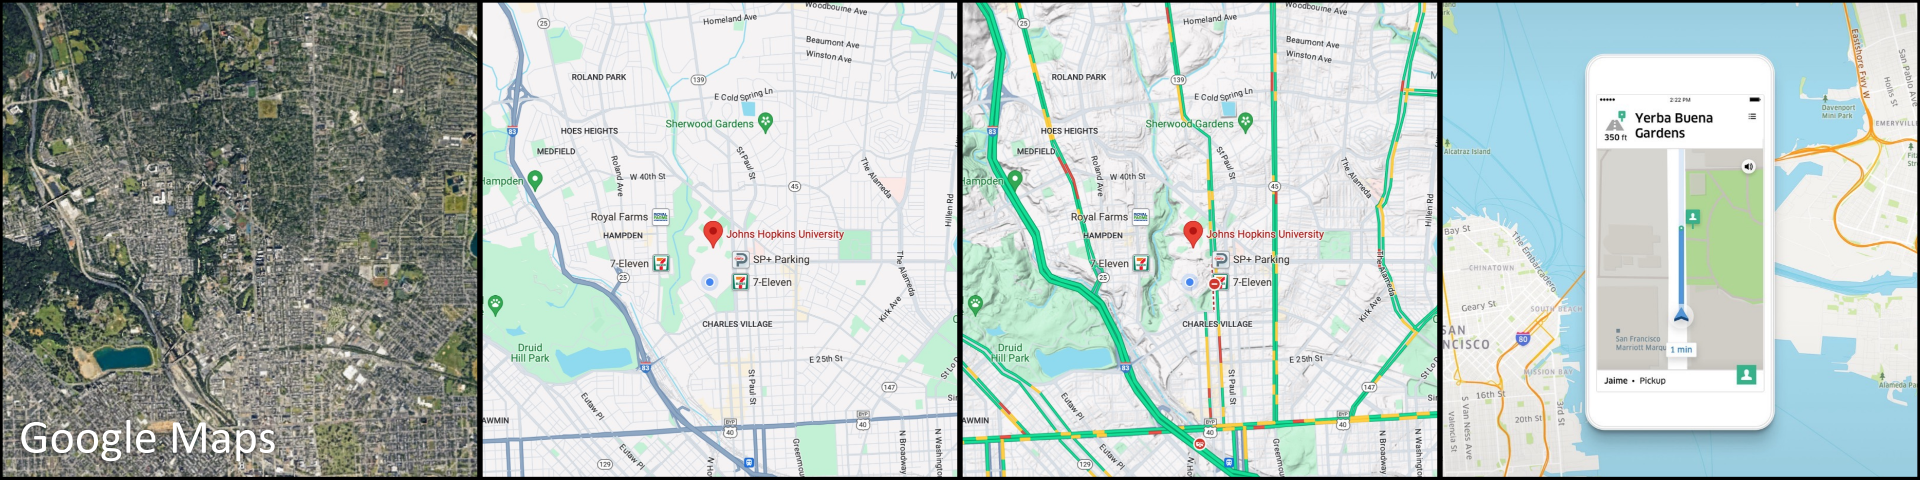

Data → Information → Knowledge → Application

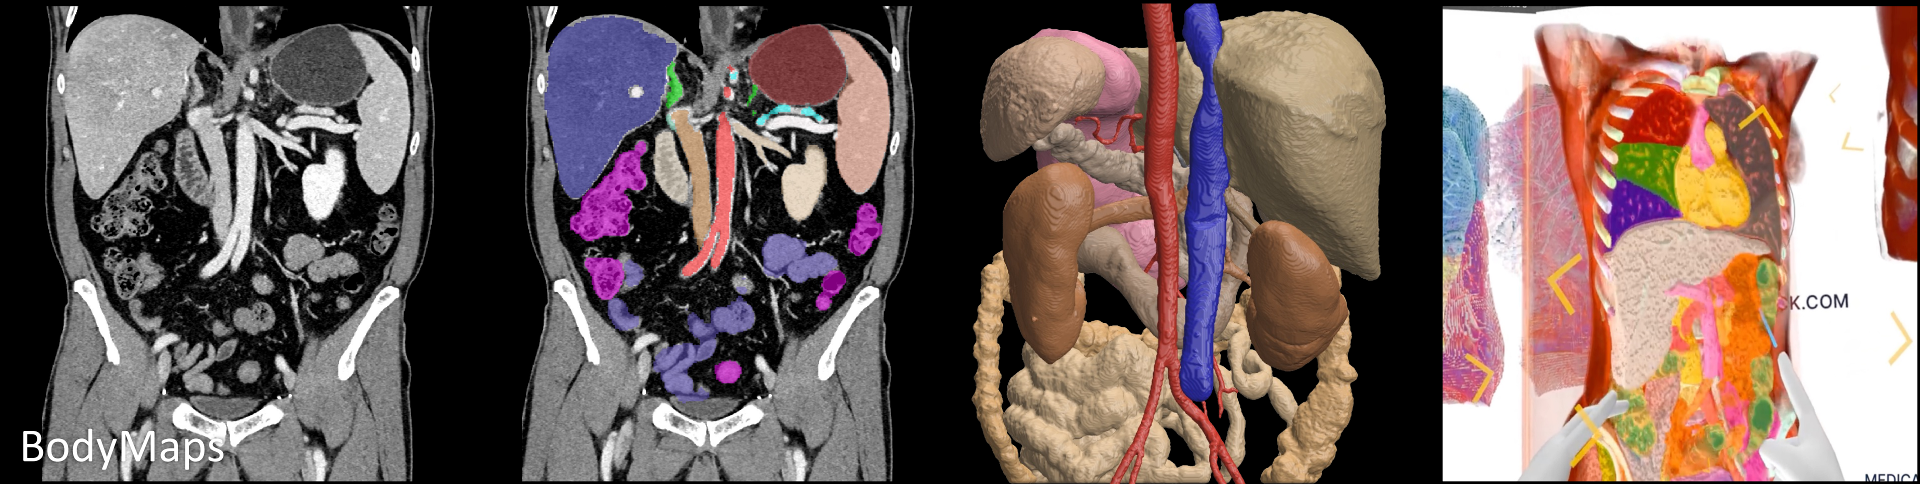

We created AbdomenAtlas of 22,682 CT volumes and 3.2M annotated masks

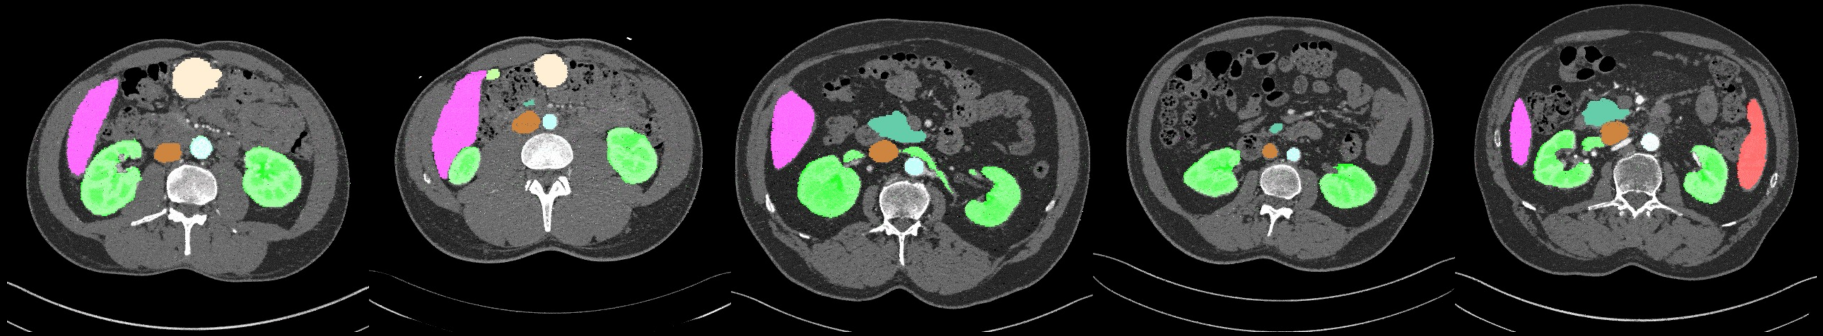

spleen kidney gall bladder liver stomach aorta postcava(IVC) pancreas

AbdomenAtlas 1.0

publicly available!

5,195 CT volumes and 46K annotated masks

https://www.zongweiz.com/dataset

AbdomenAtlas 1.1

coming in MICCAI-2024 challenge

9,262 CT volumes and 231K annotated masks

https://www.zongweiz.com/dataset

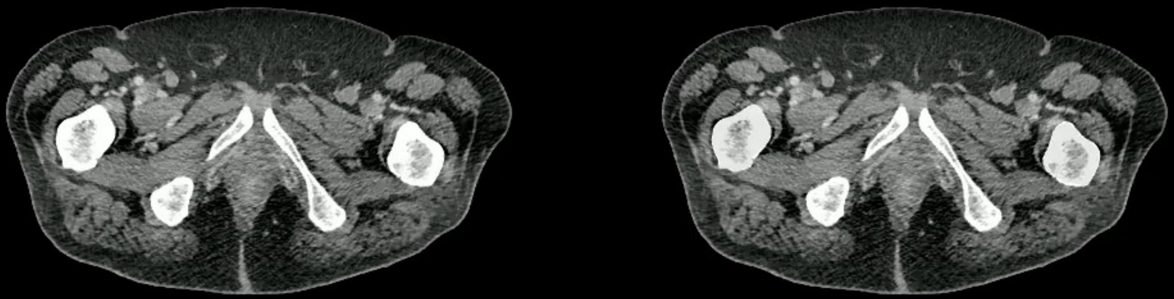

Highlight 1. Multicenter CT scans (3M)

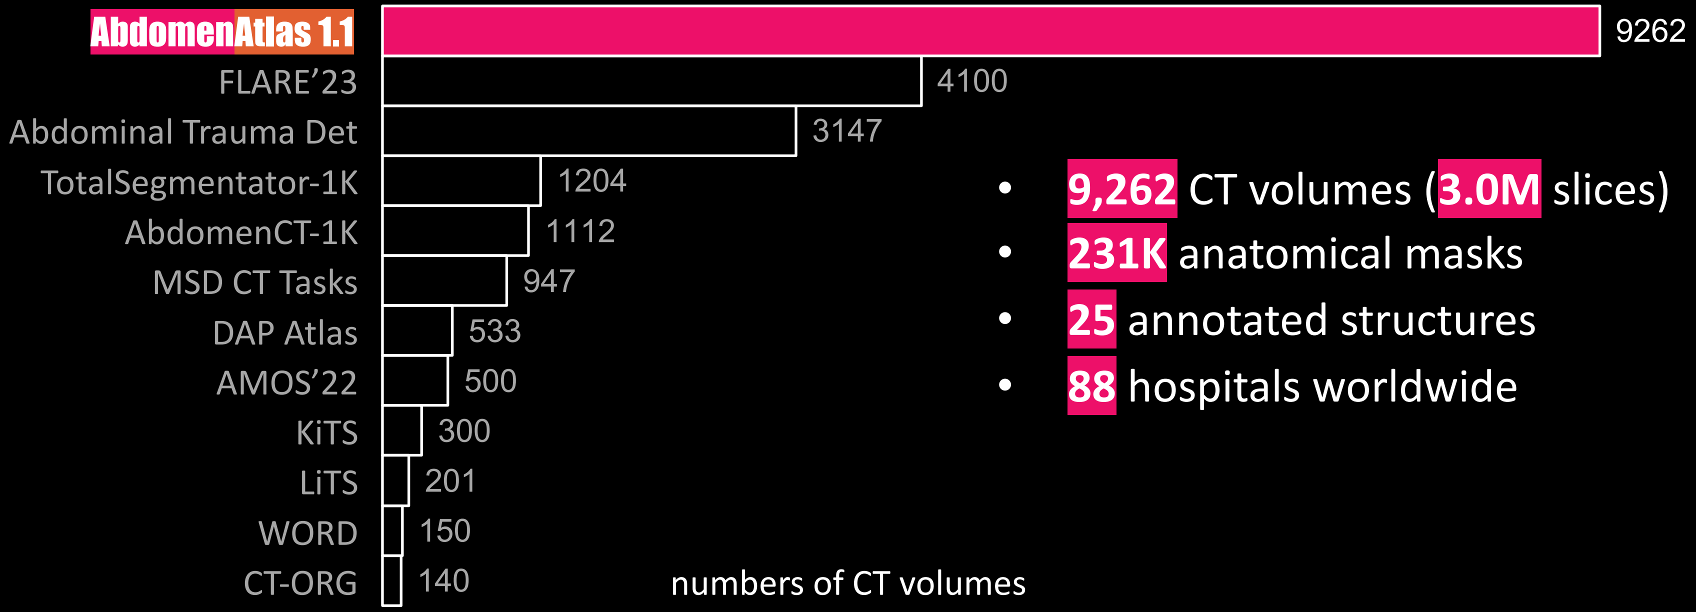

- **9,262** CT volumes (**3.0M** slices)
- **231K** anatomical masks
- **25** annotated structures
- **88** hospitals worldwide

Highlight 2. Fully annotated 25 structures

public dataset ↓ / structure name →	spleen	right kidney	left kidney	gall bladder	esophagus	liver	stomach	aorta	postcava	portal & splenic vein	pancreas	right adrenal gland	left adrenal gland	duodenum	hepatic vessel	right lung	left lung	colon	intestine	rectum	bladder	prostate	left head of femur	right head of femur	celiac trunk
BTCV	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓											
Pancreas-CT											✓														
CHAOS						✓																			
LiTS						✓																			
KiTS		✓	✓																						
WORD	✓	✓	✓	✓	✓	✓	✓				✓	✓	✓	✓				✓	✓	✓	✓		✓	✓	
AbdomenCT-1K	✓	✓	✓			✓					✓														
AMOS	✓	✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓							✓	✓			
MSD-Colon																		✓							
MSD-Lung																									
MSD-Pancreas											✓														
MSD-HepaticVessel															✓										
MSD-Spleen	✓																								
CT-ORG		✓	✓			✓										✓	✓				✓				
AbdomenAtlas 1.1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

unlabeled

labeled

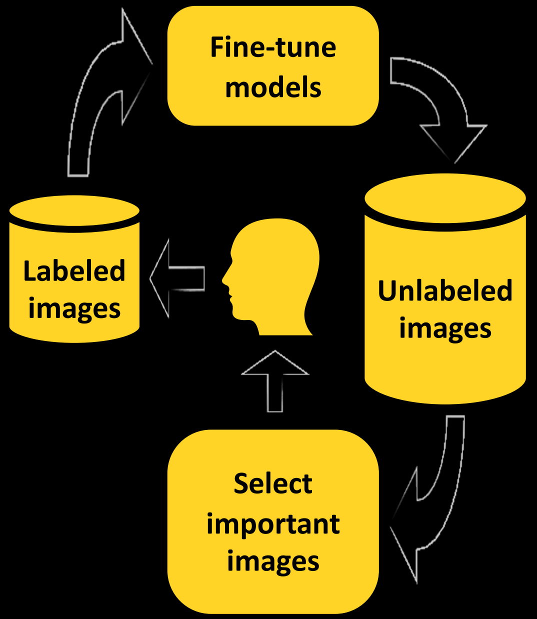

Active annotation

Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally*

Zongwei Zhou[1], Jae Shin[1], Lei Zhang[1], Suryakanth Gurudu[2], Michael Gotway[2], and Jianming Liang[1]

[1]Arizona State University
{zongweiz,sejong,lei.zhang.10,jianming.liang}@asu.edu
[2]Mayo Clinic
{gurudu.suryakanth,gotway.michael}@mayo.edu

Abstract

Intense interest in applying convolutional neural networks (CNNs) in biomedical image analysis is wide spread, but its success is impeded by the lack of large annotated datasets in biomedical imaging. Annotating biomedical images is not only tedious and time consuming, but also demanding of costly, specialty-oriented knowledge and skills, which are not easily accessible. To dramatically reduce annotation cost, this paper presents a novel method called AIFT (active, incremental fine-tuning) to naturally integrate active learning and transfer learning into a single framework. AIFT starts directly with a pre-trained CNN to seek "worthy" samples from the unannotated for annotation, and the (fine-tuned) CNN is further fine-tuned continuously by incorporating newly annotated samples in each iteration to enhance the CNN's performance incrementally. We have evaluated our method in three different biomedical imaging applications, demonstrating that the cost of annotation can be cut by at least half. This performance is attributed to the several advantages derived from the advanced active and incremental capability of our AIFT method.

but its success is impeded by the lack of such large annotated datasets in biomedical imaging. Annotating biomedical images is not only tedious and time consuming, but also demanding of costly, specialty-oriented knowledge and skills, which are not easily accessible. Therefore, we seek to answer this critical question: *How to dramatically reduce the cost of annotation when applying CNNs in biomedical imaging.* In doing so, we present a novel method called AIFT (active, incremental fine-tuning) to naturally integrate active learning and transfer learning into a single framework. Our AIFT method starts directly with a pre-trained CNN to seek "salient" samples from the unannotated for annotation, and the (fine-tuned) CNN is continuously fine-tuned by incrementally enlarging the training dataset with newly annotated samples. We have evaluated our method in three different applications including colonoscopy frame classification, polyp detection, and pulmonary embolism (PE) detection, demonstrating that the cost of annotation can be cut by at least half.

This outstanding performance is attributed to a simple yet powerful observation: To boost the performance of CNNs in biomedical imaging, multiple patches are usual-

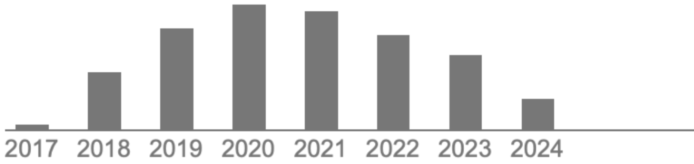

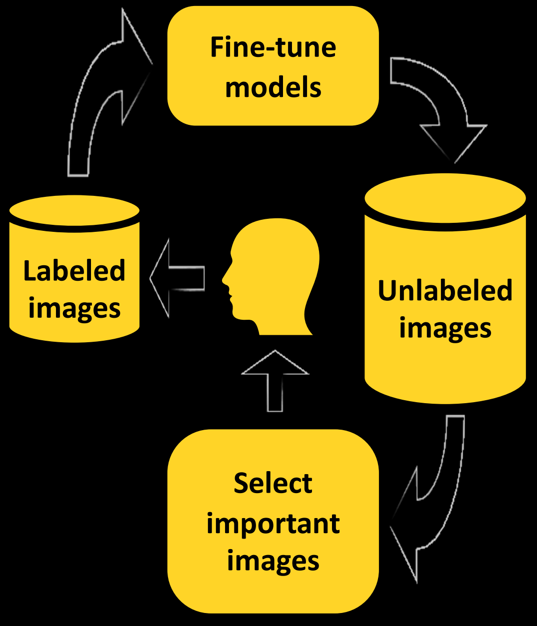

Active annotation

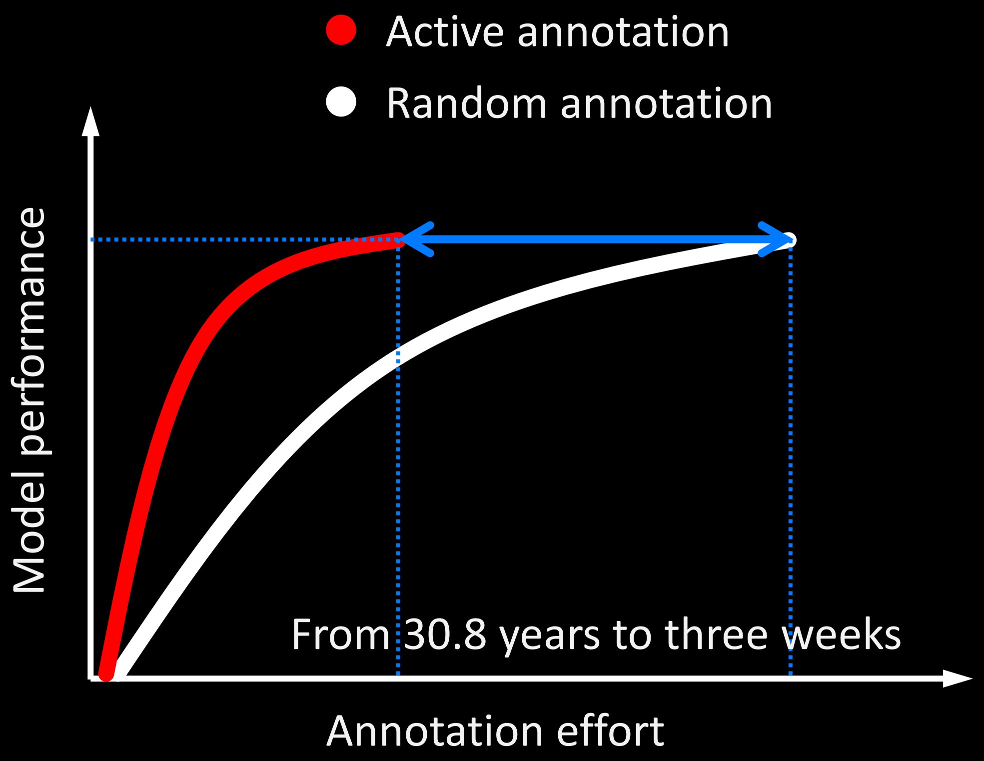

Entropy

CT Scan　Inconsistency　Uncertainty　Overlap　Attention Map

Diversity

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

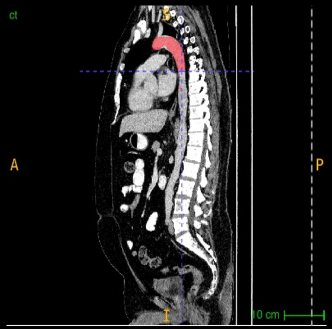

Aorta
Inconsistent labeling protocols

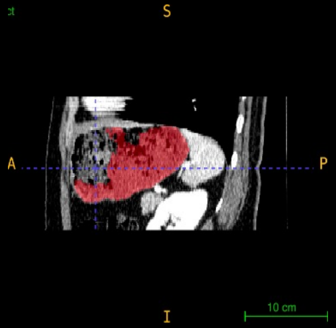

Stomach
uncertainty in empty areas

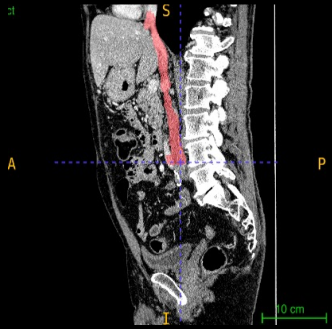

Postcava
ambiguous & blurry boundaries

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation - Summary**

- We created AbdomenAtlas 1.1 for 25 anatomical structures
- But, scaling annotations for tumors remains challenging
  - Pathology reports
  - Manual annotations
  - Collaborations (academia, industry, & hospital)

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation** (5 min) – We annotated 25 anatomical structures in 9,262 CT volumes
   - NeurIPS-2023 – Semi-automated annotation
   - Three versions: AbdomenAtlas 1.0, 1.1, Pro
2. **Dataset** (10min) – We developed data synthesis strategies to enrich tumor examples
   - CVPR-2023 & MICCAI-2024 – Liver tumor synthesis
   - CVPR-2024 – Liver, pancreatic, & kidney tumor synthesis
   - Next version: AbdomenAtlas X

Observation: early-stage tumors (< 2cm) tend to have *similar imaging characteristics* in computed tomography (CT), whether they originate in the liver, pancreas, or kidneys.

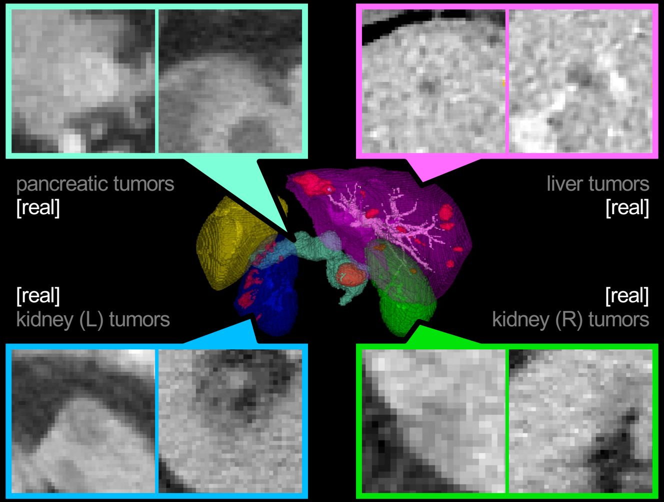

Observation: early-stage tumors (< 2cm) tend to have *similar imaging characteristics* in computed tomography (CT), whether they originate in the liver, pancreas, or kidneys.

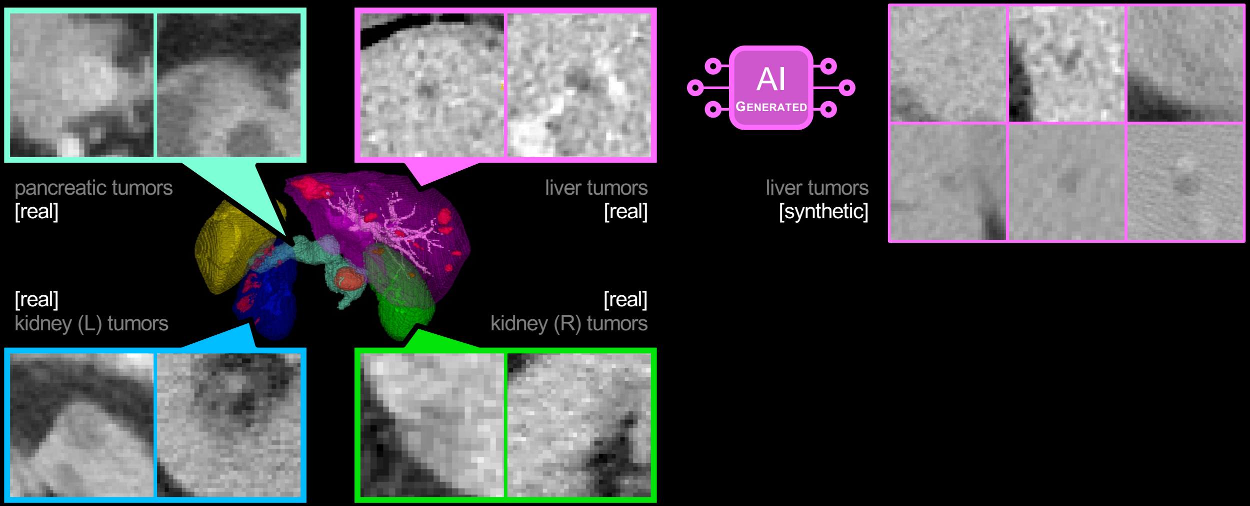

Generative AI models (e.g., diffusion models) trained on **liver tumor** examples

Observation: early-stage tumors (< 2cm) tend to have *similar imaging characteristics* in computed tomography (CT), whether they originate in the liver, pancreas, or kidneys.

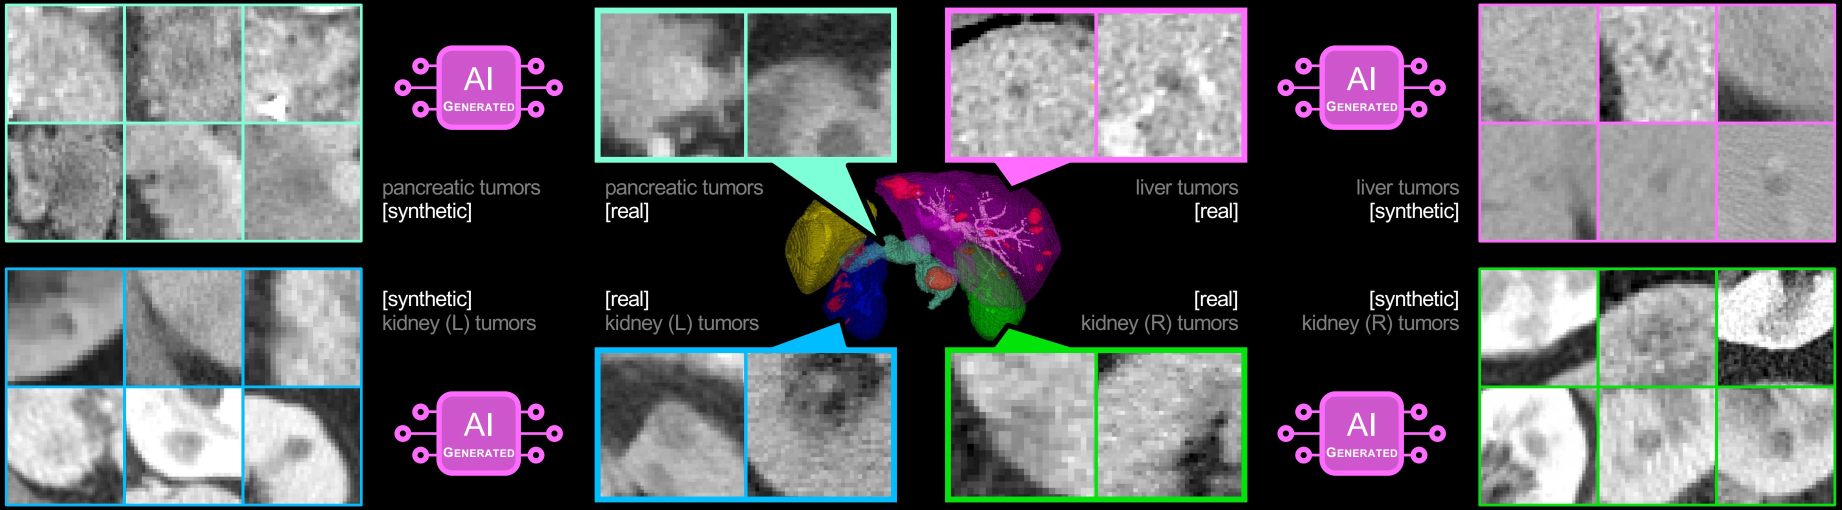

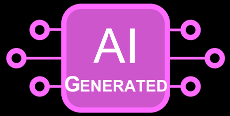

Generative AI models (e.g., diffusion models) trained on **liver tumor** examples

1. We used Diffusion Model to overfit tumor appearance
2. We address out-of-distribution by using diverse CT volumes

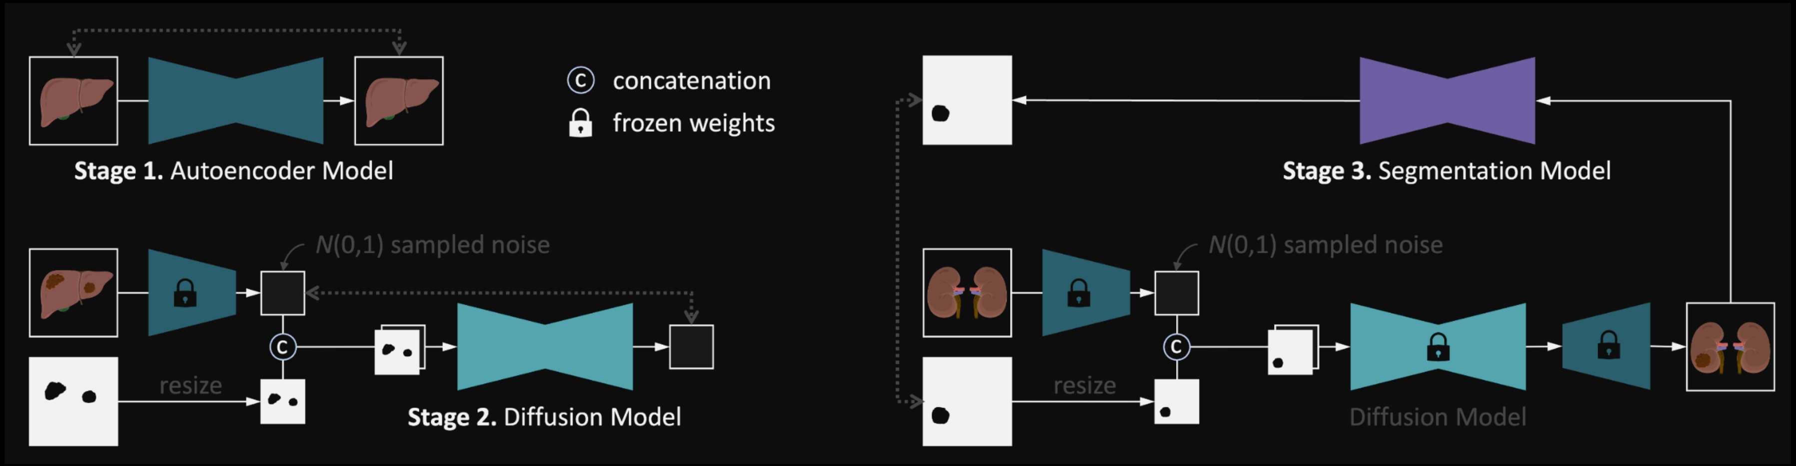

Generative AI models (e.g., diffusion models) trained on **liver tumor** examples

Medical professionals cannot tell which are real and which are synthetic tumors

Can you?

Training AI on synthetic tumors performs as well as training it on real tumors

CT

AI prediction
trained on real tumors
with per-voxel annotation
DSC = 58% [52% - 63%]

AI prediction
trained on synthetic tumors
with no annotation
DSC = 60% [55% - 65%]

Liver
Liver tumor

Diffusion Model can generate enormous small tumors for training AI models

Observation: Compared with real tumors,

AI trained on synthetic tumors improves Sensitivity

from 52% to 62% for detecting small tumors (0-5mm).

We evaluated on real tumors!

Diffusion Model can generate enormous small tumors for training AI models

Observation: Compared with real tumors,
AI trained on synthetic tumors improves Sensitivity
from 52% to 62% for detecting small tumors (0-5mm).
We evaluated on real tumors!

- Needed for early detection
    - Early signs of cancer can be subtle
    - 1/2 of liver cancer are missed by radiologists

Diffusion Model can generate enormous small tumors for training AI models

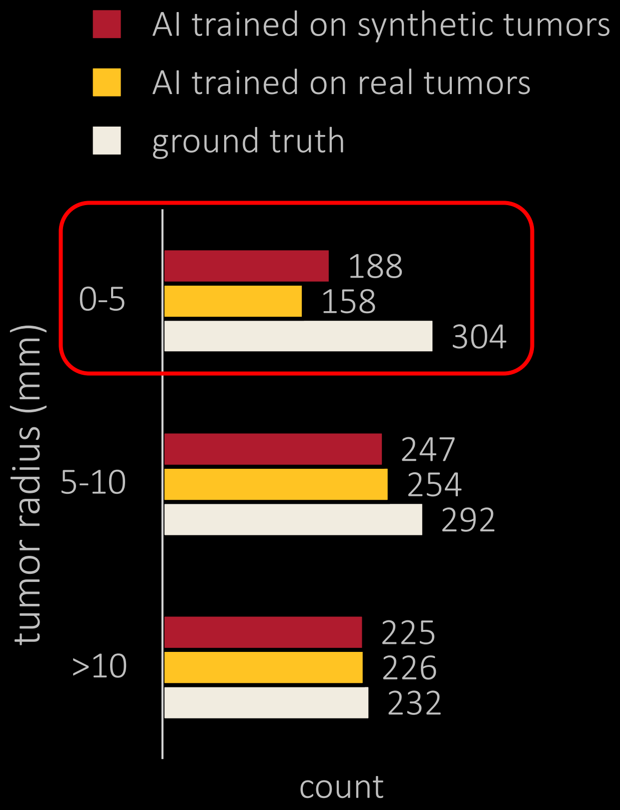

Observation: Compared with real tumors,

AI trained on synthetic tumors improves Sensitivity

from 52% to 62% for detecting small tumors (0-5mm).

We evaluated on real tumors!

- Needed for early detection
  - Early signs of cancer can be subtle
  - 1/2 of liver cancer are missed by radiologists
- Needed for AI development
  - CT scans with early cancer are limited
  - Annotations for early cancer are hard

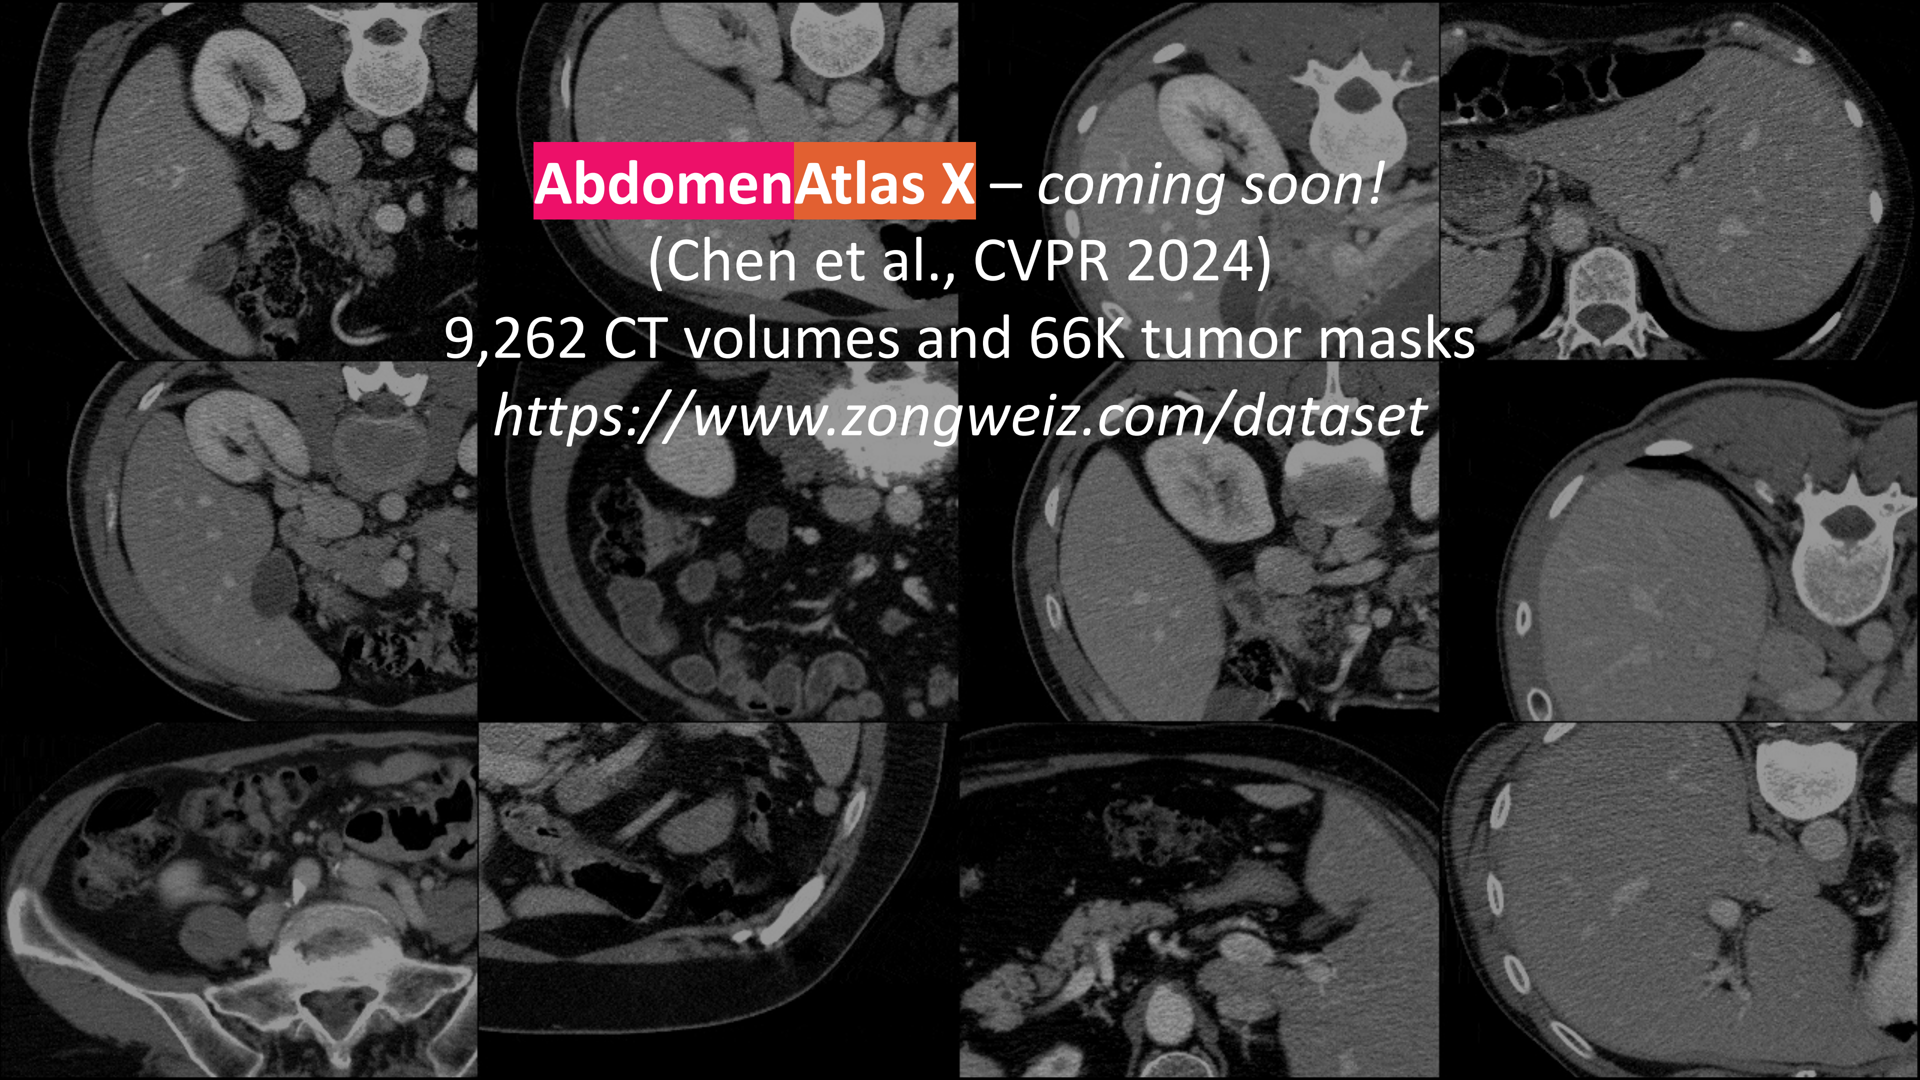

AbdomenAtlas X – *coming soon!*

(Chen et al., CVPR 2024)

9,262 CT volumes and 66K tumor masks

https://www.zongweiz.com/dataset

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

2. Dataset - Summary

- We generated synthetic tumor data in the liver, pancreas, & kidneys
- But, generalizable tumor synthesis to other organs is challenging
    - The assumption—early tumors are similar—is problematic
    - Medical knowledge is needed – CV plus Radiology

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation** (5 min) – We annotated 25 anatomical structures in 9,262 CT volumes
   - NeurIPS-2023 – Semi-automated annotation
   - Three versions: AbdomenAtlas 1.0, 1.1, Pro
2. **Dataset** (10min) – We developed data synthesis strategies to enrich tumor examples
   - CVPR-2023 – Liver tumor synthesis
   - CVPR-2024 – Liver, pancreatic, & kidney tumor synthesis
   - Next version: AbdomenAtlas X
3. **Algorithm** (10min) – We released a set of supervised pre-trained vision-language models
   - ICCV-2023 & MICCAI-2023 – Vision-language models
   - ICLR-2024 – Supervised pre-training

Universal Model

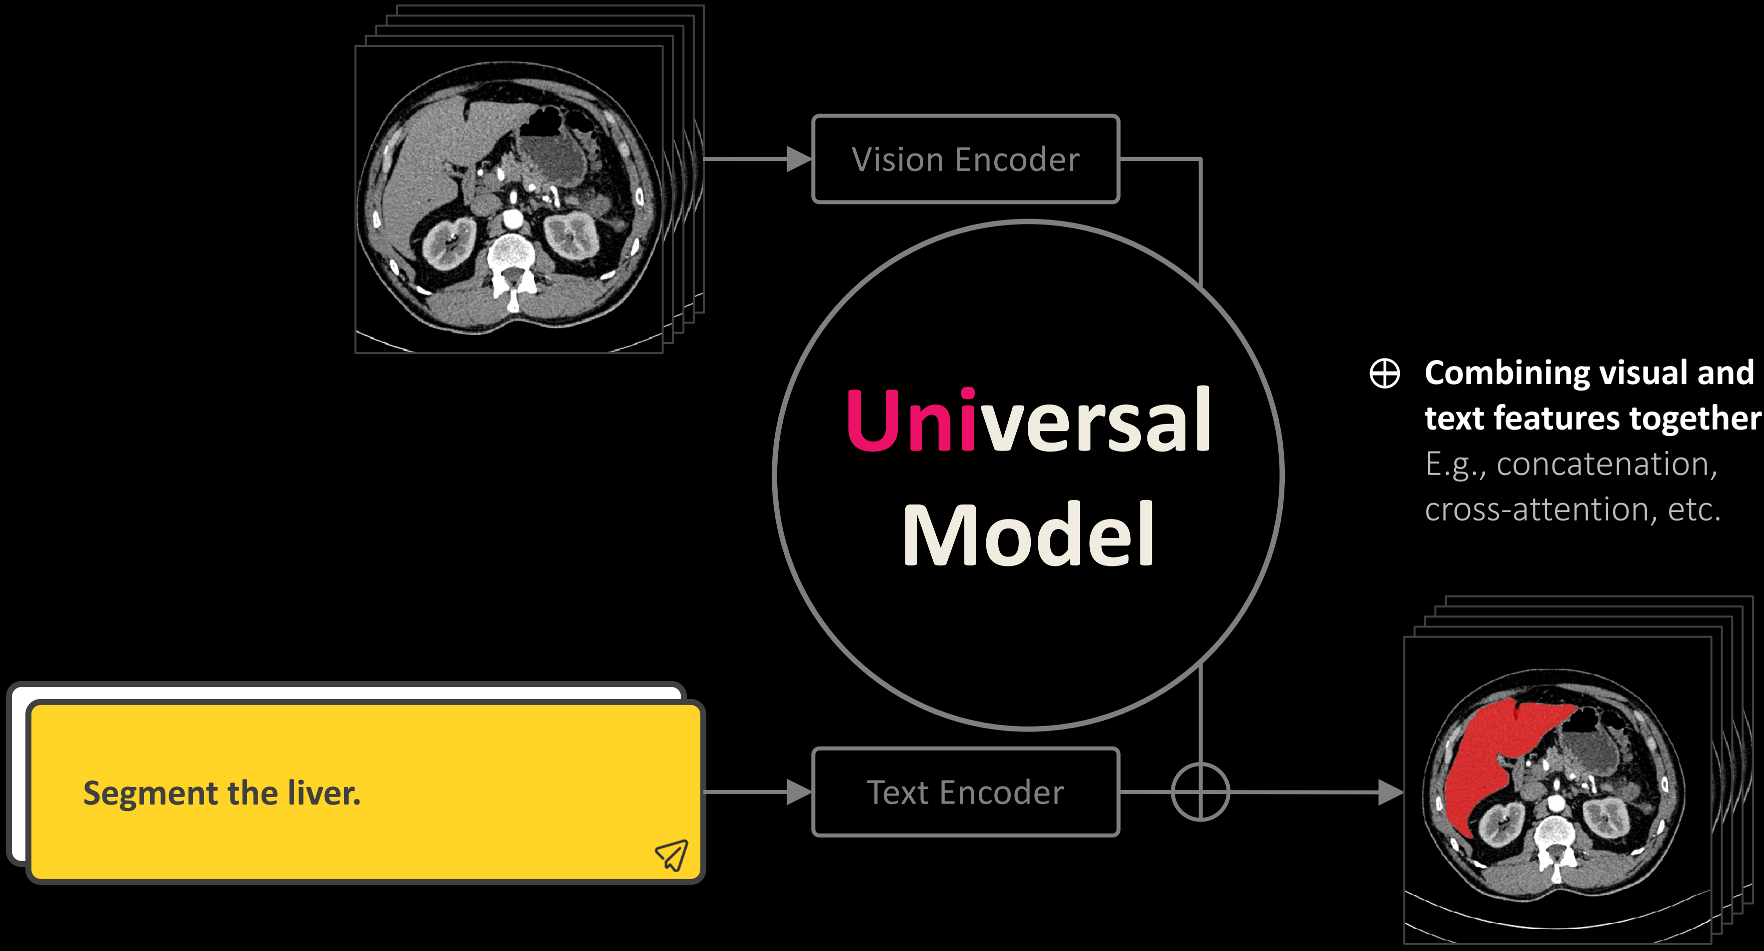
Vision Encoder
Universal Model
Combining visual and text features together
E.g., concatenation, cross-attention, etc.
Segment the liver.
Text Encoder

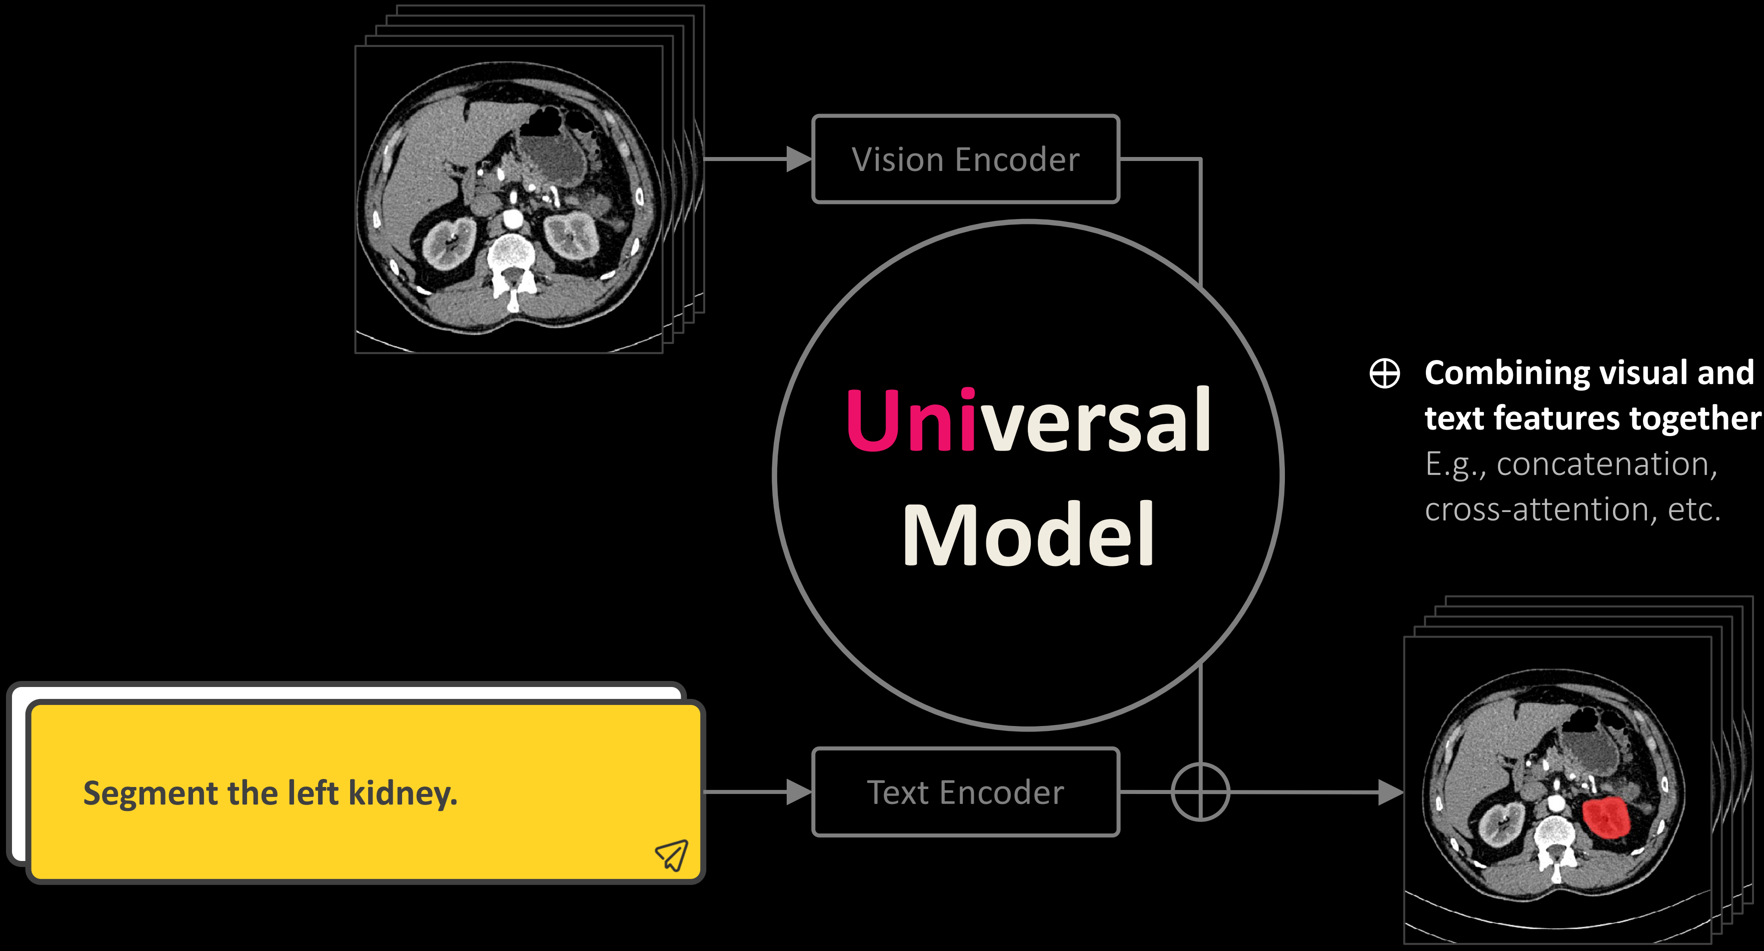

Vision Encoder
Universal
Model
Combining visual and text features together
E.g., concatenation, cross-attention, etc.
Segment the left kidney.
Text Encoder

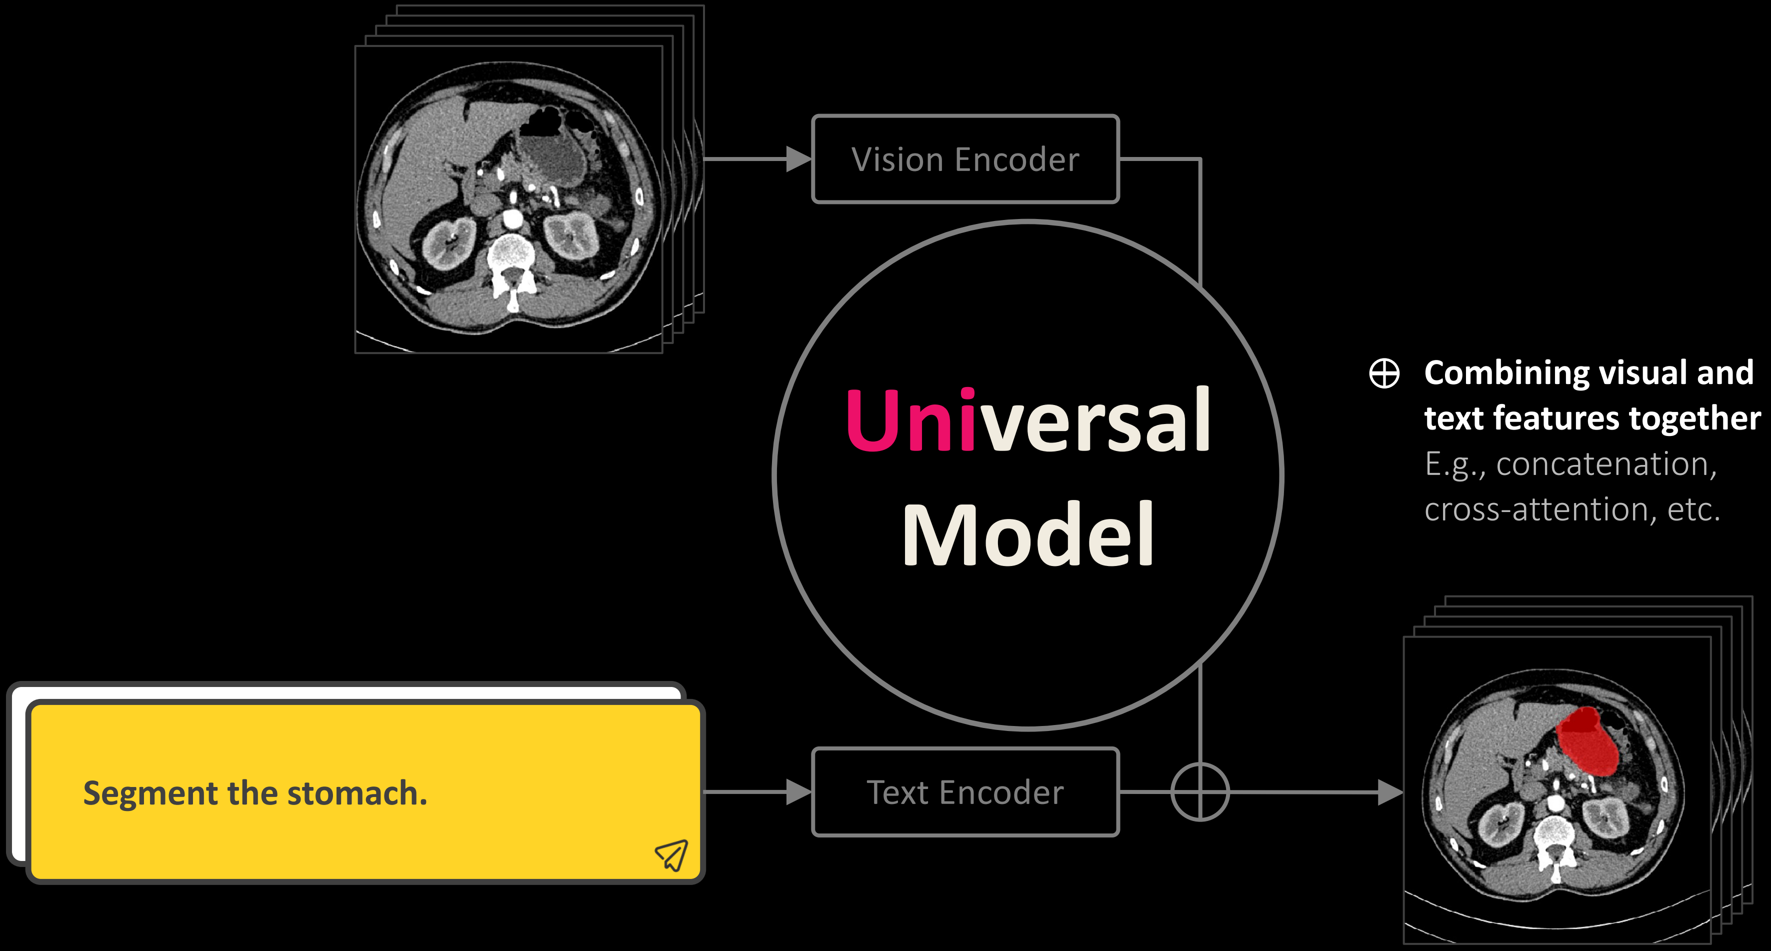

Vision Encoder
Universal
Model
Combining visual and text features together
E.g., concatenation, cross-attention, etc.
Segment the stomach.
Text Encoder

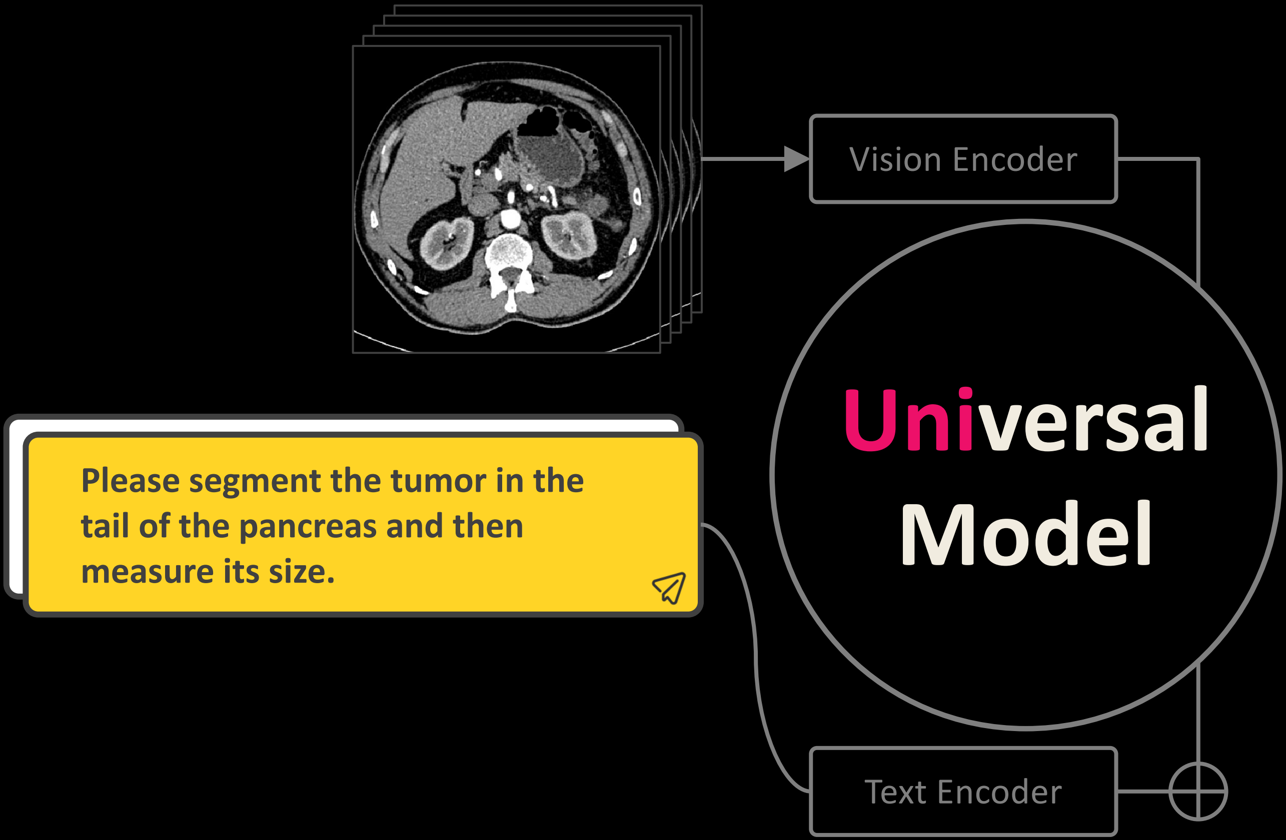
Vision Encoder
Universal Model
Please segment the tumor in the tail of the pancreas and then measure its size.
Text Encoder

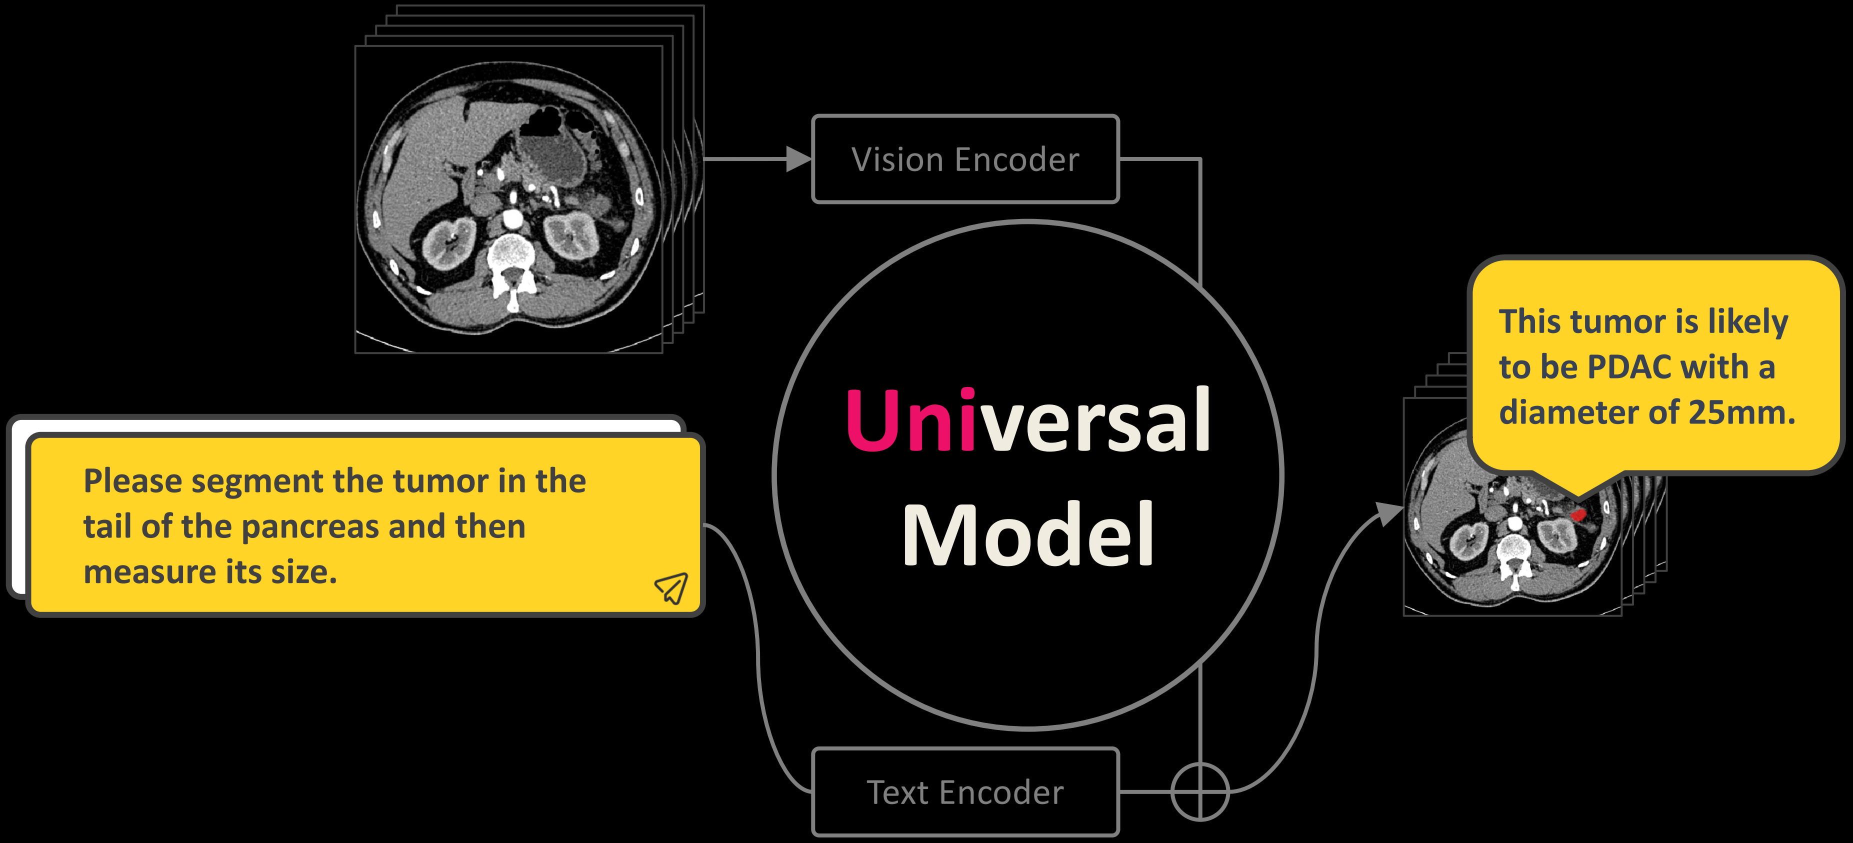

Vision Encoder
Universal
Model
Please segment the tumor in the tail of the pancreas and then measure its size.
Text Encoder
This tumor is likely to be PDAC with a diameter of 25mm.

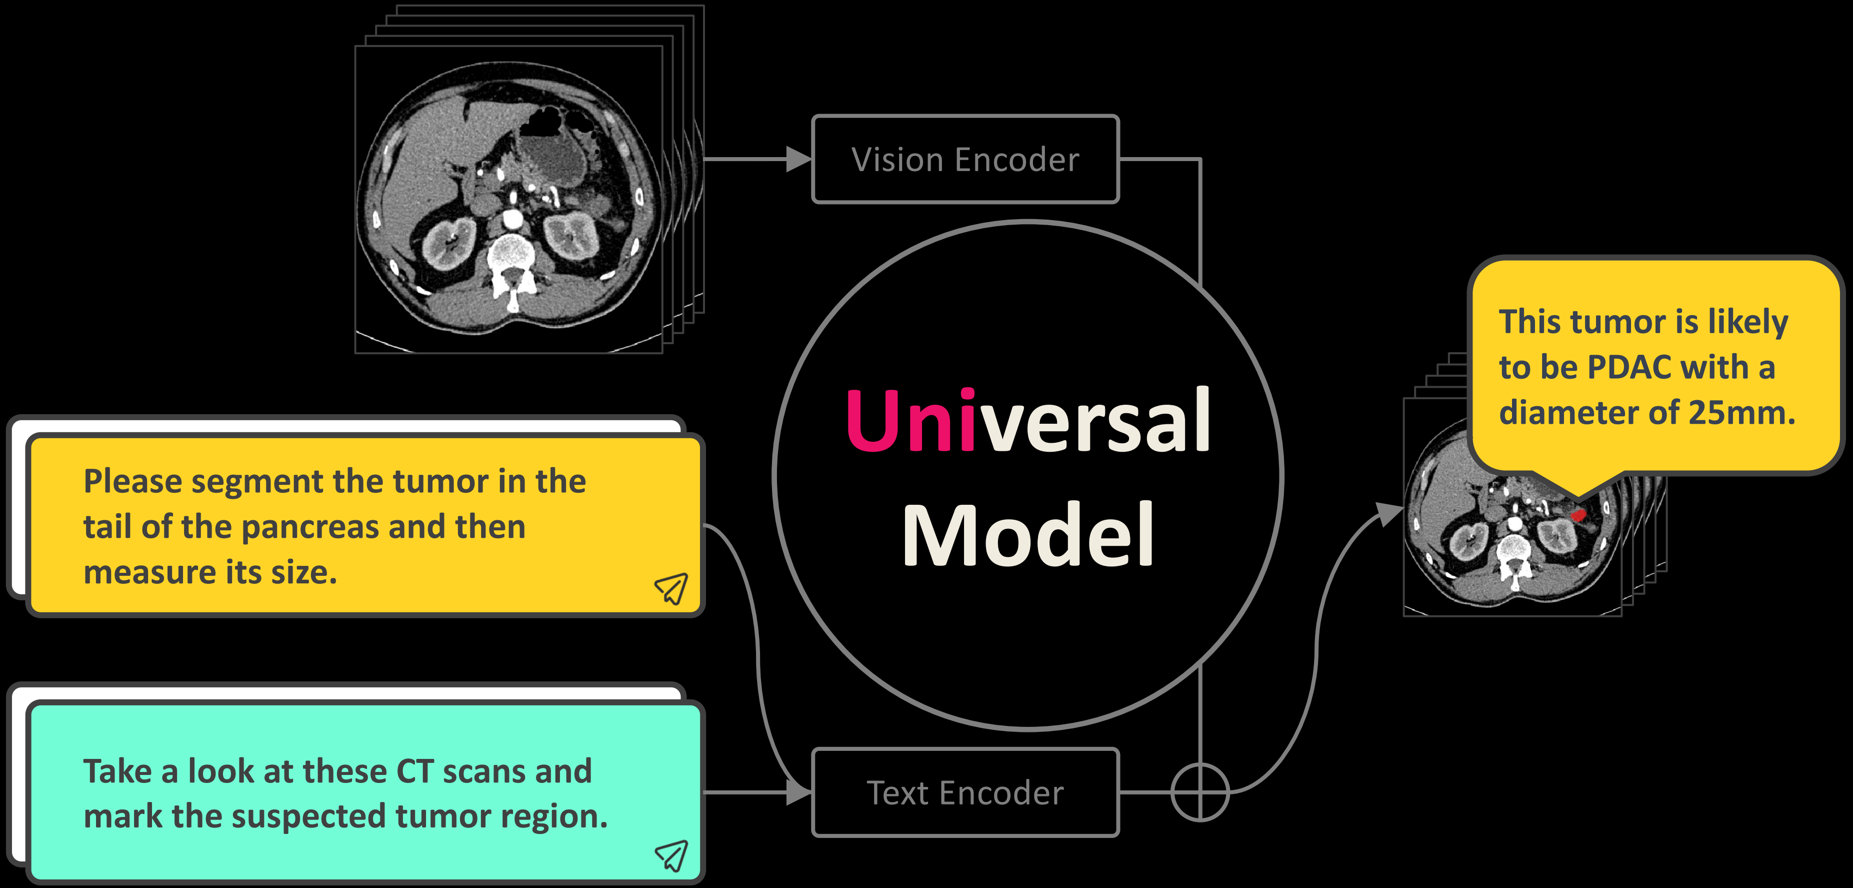
Vision Encoder
Universal Model
This tumor is likely to be PDAC with a diameter of 25mm.
Please segment the tumor in the tail of the pancreas and then measure its size.
Take a look at these CT scans and mark the suspected tumor region.
Text Encoder

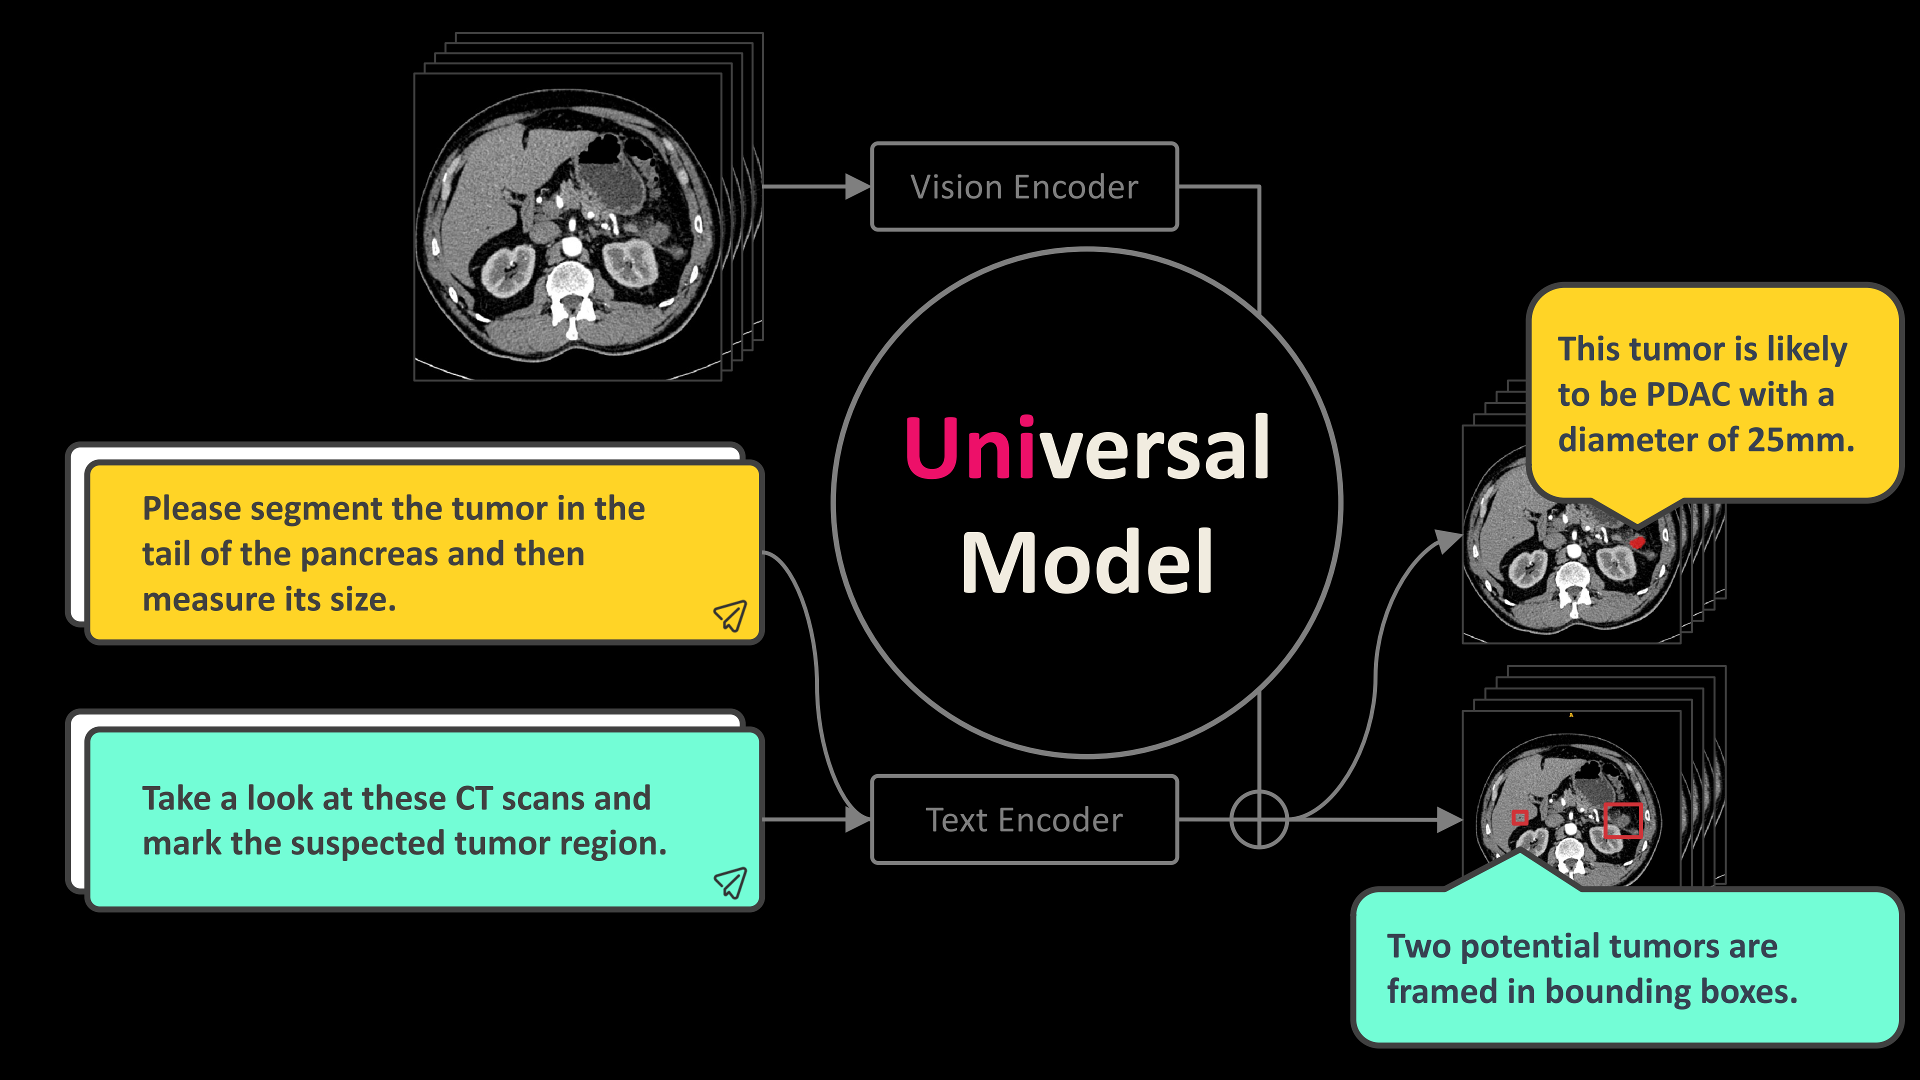
Vision Encoder
Universal
Model
Please segment the tumor in the tail of the pancreas and then measure its size.
Take a look at these CT scans and mark the suspected tumor region.
Text Encoder
This tumor is likely to be PDAC with a diameter of 25mm.
Two potential tumors are framed in bounding boxes.

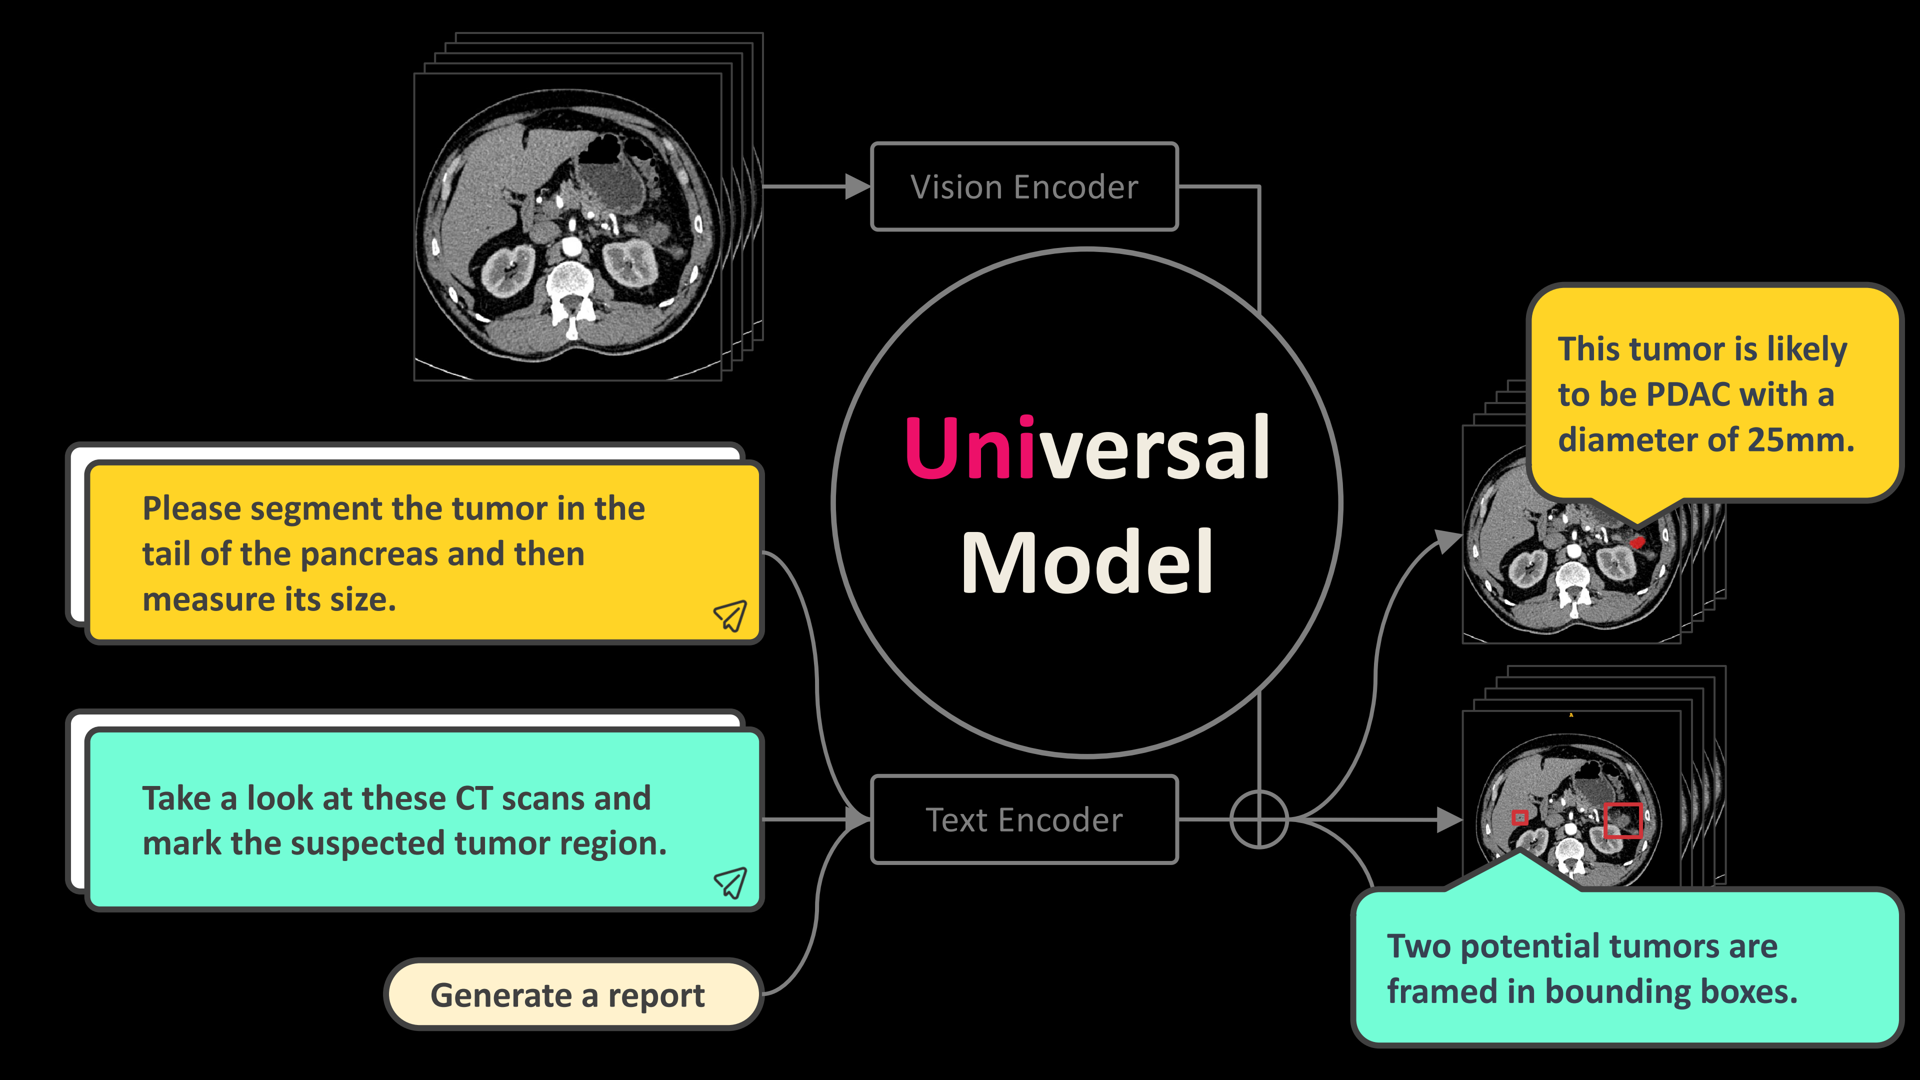
Vision Encoder
Universal
Model
Please segment the tumor in the tail of the pancreas and then measure its size.
Take a look at these CT scans and mark the suspected tumor region.
Generate a report
Text Encoder
This tumor is likely to be PDAC with a diameter of 25mm.
Two potential tumors are framed in bounding boxes.

#	User (Team)	Created	Mean Position
1st	zongwei.zhou (universal_model)	13 Feb. 2023	5.6
2nd	Swin_UNETR	12 Nov. 2021	10.1
3rd	ahatamiz2	12 Nov. 2021	10.1
4th	Isensee	6 Dec. 2019	12.3
5th	AndyL	24 Nov. 2022	12.5
6th	heyufan1995	30 Oct. 2020	12.8
7th	qsyeung	5 Jan. 2023	12.9
8th	vishwesh.nath	11 Nov. 2021	13.5

Acknowledgement: Universal Model adopted the checkpoint released by Tang et al., CVPR 2022 (NVIDIA). This checkpoint was ***self-supervised*** pre-trained from 5,000 CT volumes.

But wait, we have AbdomenAtlas 1.1 now!

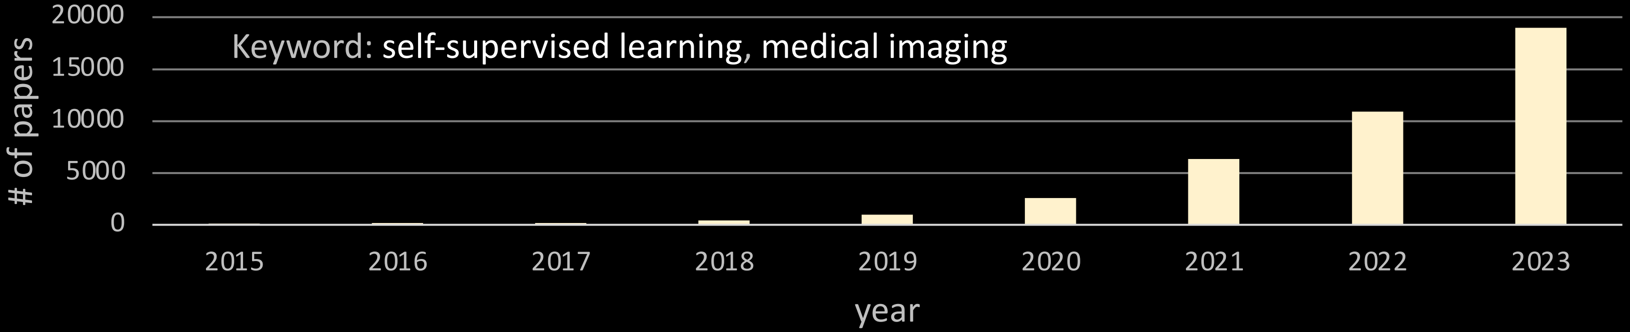

Option 1. Self-supervised pre-training

De-noising [Vincent et al., 2010]
Positioning **[Doersch et al., 2015]**
In-painting [Pathak et al., 2016]
Jigsaw [Noroozi and Favaro, 2016]
DeepCluster [Caron et al., 2018]
Restoration [Chen et al., 2019]

Models Genesis
[Zhou et al., MICCAI 2019]
MICCAI Young Scientist Award
MedIA Best Paper Award

Rubiks Cube
[Zhuang et al., MICCAI 2019]

Swin UNETR
[Tang et al., CVPR 2022]

UniMiSS
[Xie et al., ECCV 2022]

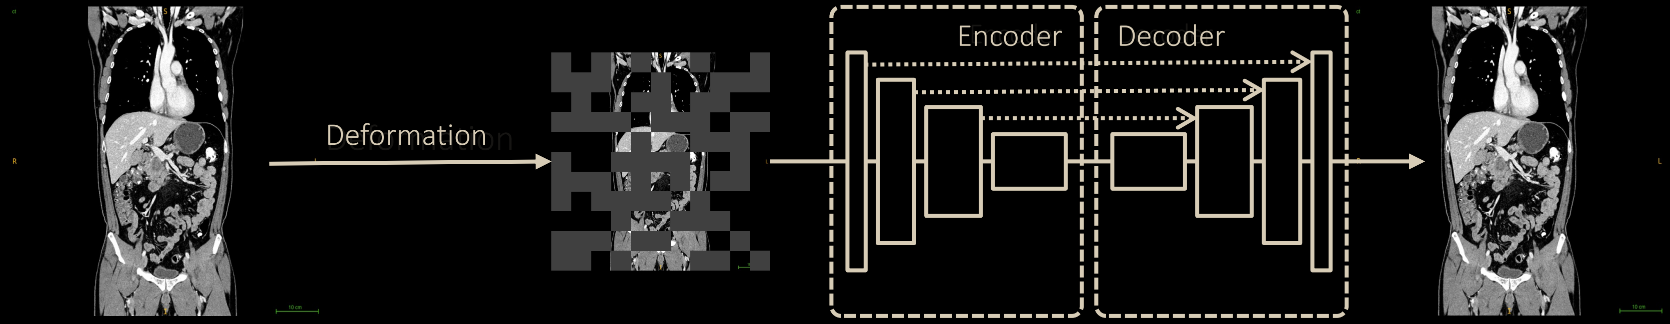

Option 1. Self-supervised pre-training

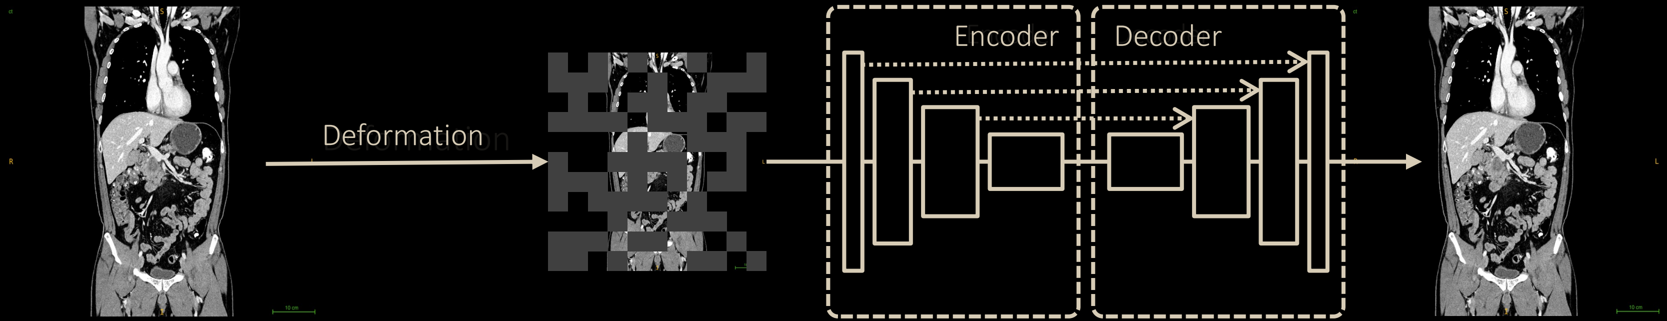

Option 1. Self-supervised pre-training

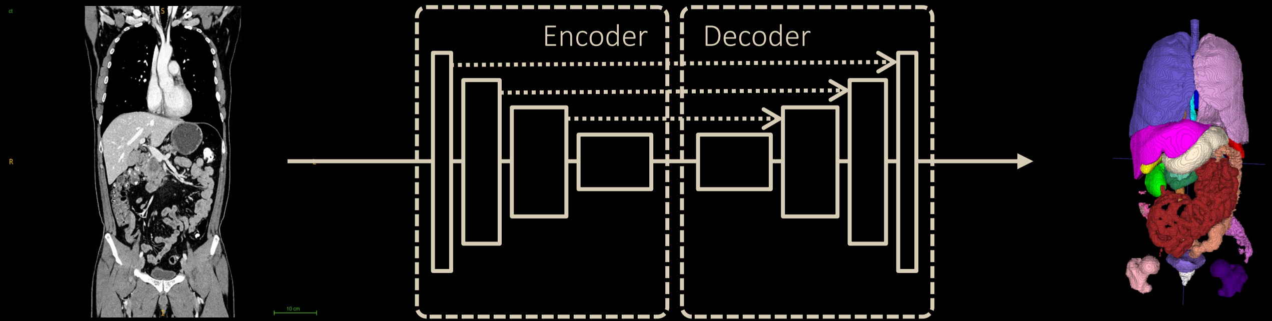

Option 2. Supervised pre-training

AbdomenAtlas 1.1

Supervised > **Self-supervised**
data & computation efficiency

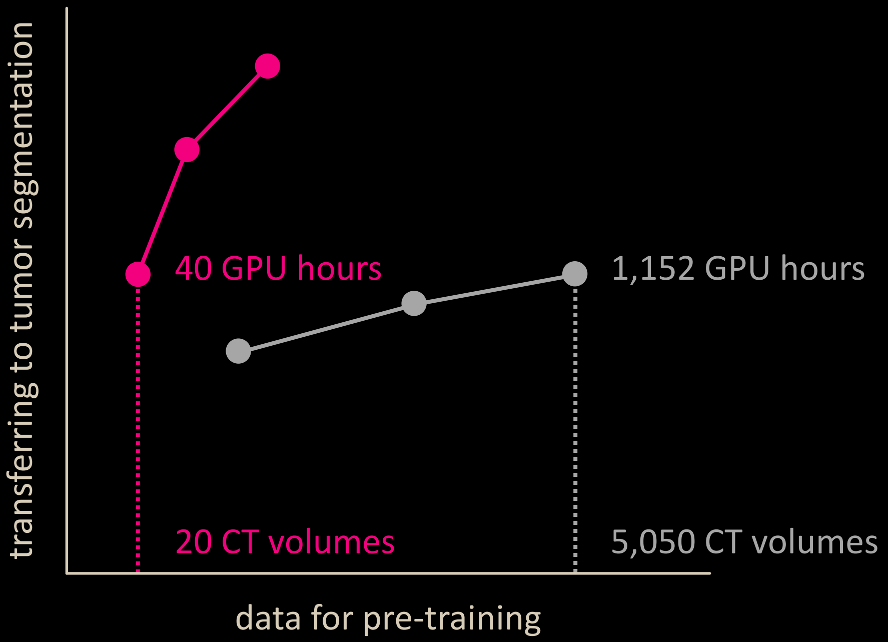

Supervised > **Self-supervised**
annotation & learning efficiency

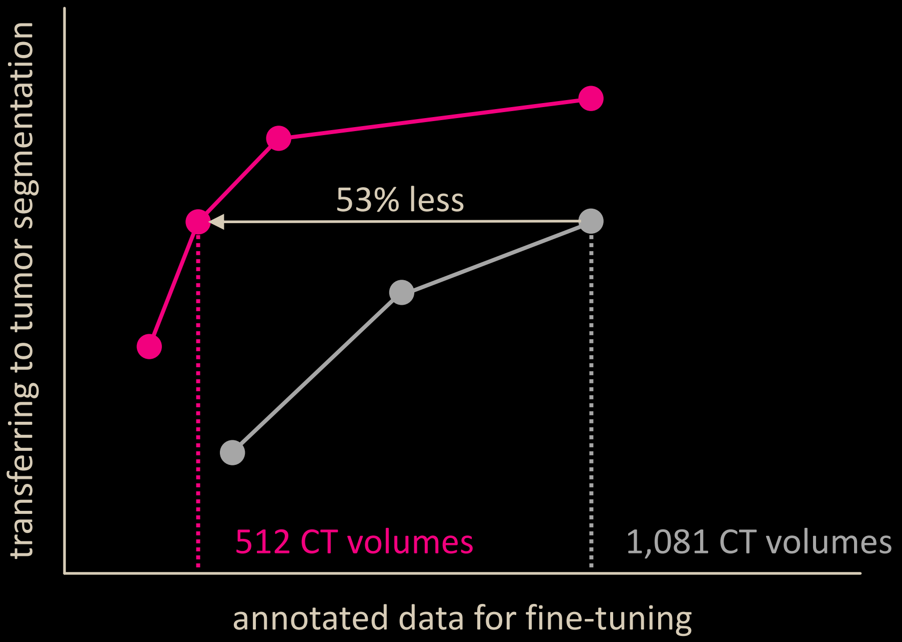

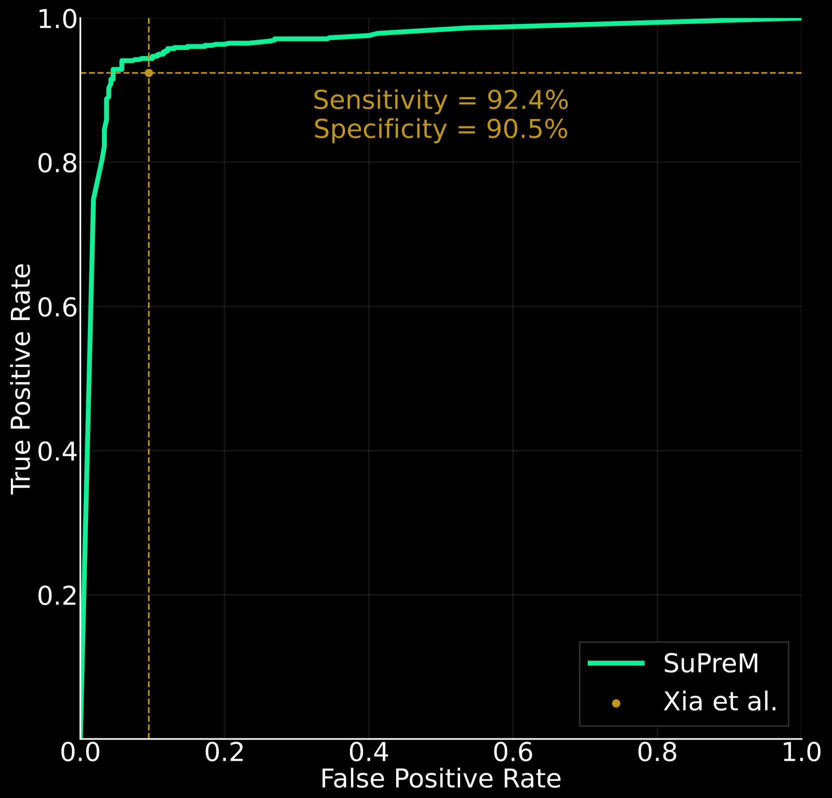

The FELIX Project: Deep Networks To Detect Pancreatic Neoplasms

Yingda Xia[1,†], Qihang Yu[1,†], Linda Chu[2,†], Satomi Kawamoto[2,†], Seyoun Park[2], Fengze Liu[1], Jieneng Chen[1], Zhuotun Zhu[1], Bowen Li[1], Zongwei Zhou[1], Yongyi Lu[1], Yan Wang[1], Wei Shen[1], Lingxi Xie[1], Yuyin Zhou[1], Christopher Wolfgang[3], Ammar Javed[3], Daniel Fadaei Fouladi[2], Shahab Shayesteh[2], Jefferson Graves[2], Alejandra Blanco[2], Eva S. Zinreich[2], Benedict Kinny-Köster[3], Kenneth Kinzler[4,6,7,8], Ralph H. Hruban[5], Bert Vogelstein[4,6,7,8,9], Alan L. Yuille[1,*], Elliot K. Fishman[2,*]

[1]Department of Computer Science, Johns Hopkins University; [2]Department of Radiology and Radiological Science, Johns Hopkins Medicine; [3]Department of Surgery, New York University; [4]Department of Oncology, Johns Hopkins Medicine; [5]Department of Pathology, Johns Hopkins Medicine; [6]Ludwig Center, Johns Hopkins University School of Medicine; [7]Sidney Kimmel Comprehensive Cancer Center, Johns Hopkins University School of Medicine; [8]Sol Goldman Pancreatic Cancer Research Center, Johns Hopkins University School of Medicine; [9]Howard Hughes Medical Institute, Johns Hopkins Medical Institutions

Abstract Tens of millions of abdominal images are performed with computed tomography (CT) in the U.S. each year but pancreatic cancers are sometimes not initially detected in these images. We here describe a suite of algorithms (named FELIX) that can recognize pancreatic lesions from CT images without human input. Using FELIX, >90% of patients with pancreatic ductal adenocarcinomas were detected at a specificity of >90% in patients without pancreatic disease. FELIX may be able to assist radiologists in identifying pancreatic cancers earlier, when surgery and other treatments offer more hope for long-term survival.

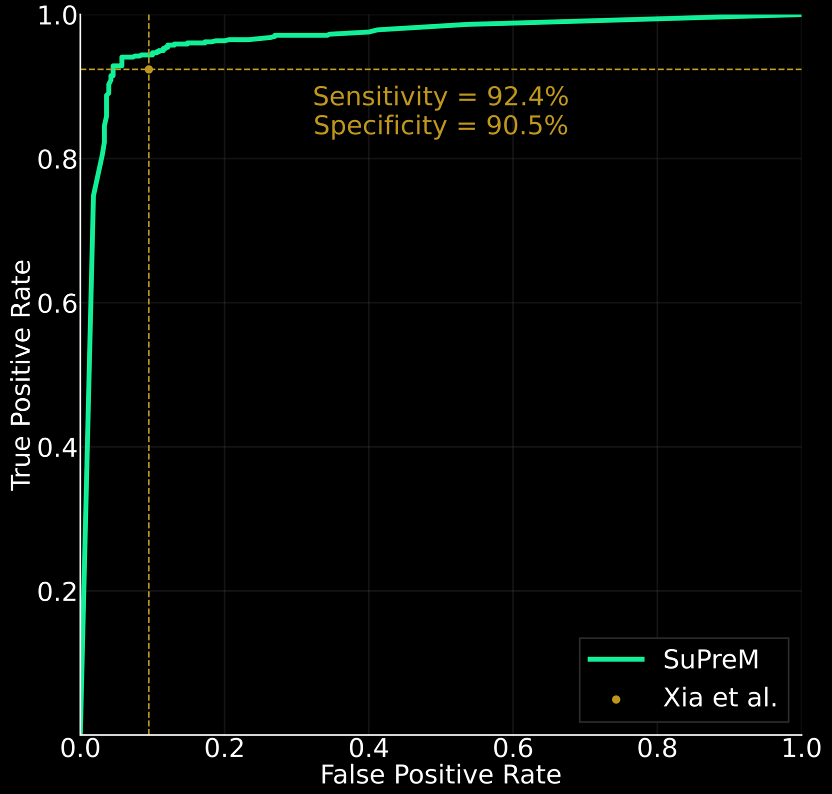

Extending Supervised Pretraining on Large Organ Dataset to Universal Lesion Segmentation

Yannick Kirchhoff, Maximilian R. Rokuss, Saikat Roy, Klaus H. Maier-Hein
German Cancer Research Center (DKFZ), Heidelberg, Germany

Method Description:
Our method uses a lightweight nnUNet [1] pretrained on a large organ segmentation dataset. We used supervised pretraining on the Abdomen Atlas 1.0 [2,3] dataset with 8 annotated organs over 5000 CT samples. Our work demonstrates that we can efficiently train the smallest possible models to learn robust representations of CT images from a large organ database which does not have any label overlaps with the ULS23 labelset. Our priority is not solely about reaching the very best Dice score. Instead, we aim for a very lightweight design while maintaining comparable accuracy.

References:

[1] Isensee, Fabian, et al. "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation." Nature methods 18.2 (2021): 203-211.

[2] Qu, Chongyu, et al. "Abdomenatlas-8k: Annotating 8,000 CT volumes for multi-organ segmentation in three weeks." Advances in Neural Information Processing Systems 36 (2024).

[3] Li, Wenxuan, Alan Yuille, and Zongwei Zhou. "How well do supervised models transfer to 3d image segmentation?." The Twelfth International Conference on Learning Representations. 2023.

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

3. Algorithm - Summary

- Language is the key to unify multiple vision tasks
  - High-performance & generalizable
  - Accommodating new classes/tasks
- Let’s start the debate between self-supervised vs. supervised pre-training

The BodyMaps Project

Goal: Detect, segment, and classify anatomical structures and abdominal tumors

1. **Annotation** (5 min) – We annotated 25 anatomical structures in 9,262 CT volumes
    - NeurIPS-2023 – Semi-automated annotation
    - Three versions: AbdomenAtlas 1.0, 1.1, Pro
2. **Dataset** (10min) – We developed data synthesis strategies to enrich tumor examples
    - CVPR-2023 – Liver tumor synthesis
    - CVPR-2024 – Liver, pancreatic, & kidney tumor synthesis
    - Next version: AbdomenAtlas X
3. **Algorithm** (10min) – We released a set of supervised pre-trained vision-language models
    - ICCV-2023 – Vision-language models
    - ICLR-2024 – Supervised pre-training
4. **Ending** (2 min) – A large, open algorithmic benchmark

AbdomenAtlas 1.0

@ISBI 2024 Challenge

IEEE INTERNATIONAL SYMPOSIUM ON BIOMEDICAL IMAGING
ISBI 2024
27-30 MAY, 2024 – ATHENS, GREECE
MEGARON ATHENS INTERNATIONAL CONFERENCE CENTRE (MAICC)

Backbone	Author	Institute	Publication	Backbone	Author	Institute	Publication
U-Net	O. Ronneberger	Uni Freiburg	MICCAI	nnU-Net	Fabian Isensee	DKFZ	Nat. Methods
nnFormer	Hong-Yu Zhou	HKU	TIP	SAT	Ziheng Zhao	SJTU	arXiv
CoTr	Yutong Xie	NPU	MICCAI	Swin UNETR	Ali Hatamizadeh	NVIDIA	MICCAIW
UniverSeg	Victor Ion Butoi	MIT	ICCV	UniSeg	Yiwen Ye	NPU	MICCAI
UNet++	Zongwei Zhou	ASU	TMI	MagicNet	Duowen Chen	ECNU	CVPR
TransUNet	Jieneng Chen	JHU	ICMLW	MedSegDiff	Junde Wu	NUS	AAAI
Swin-Unet	Hu Cao	Huawei	ECCVW	3D UNeXt	Jeya Maria Jose	JHU	MICCAI
DiNTS	Yufan He	JHU	CVPR	……			

So far, **53 groups** have confirmed the contribution—we will invite more authors of **famous backbones** for medical segmentation.

Zongwei Zhou
zzhou82@jh.edu
Code, Dataset, & Model:
https://github.com/MrGiovanni/AbdomenAtlas

Thank you!

Acknowledgement: Chongyu Qu · Xiaoxi Chen · Pedro R. A. S. Bassi · Tiezheng Zhang · Hualin Qiao · Jessica Han · Jie Liu · Yucheng Tang · Pedro Ricardo Ariel Salvador Bassi · Yijia Shi · Lihua Gao · Wenlan Zhou · Guiqin Zhang · Damin Li · Yixiao Zhang · Jieneng Chen · Linda Chu · Satomi · Kawamoto · Seyoun Park · Christopher Wolfgang · Ammar Javed · Daniel Fadaei · Shahab Shayesteh · Jefferson Graves · Alejandra Blanco · Eva S. Zinreich · Benedict Kinny-Köster · Kenneth Kinzler · Ralph H. Hruban · Bert Vogelstein · Elliot K. Fishman · Jaimie Patterson · Junfei Xiao · Yongyi Lu · Xinyi Li · Huimiao Chen · Yaoyao Liu · Qi Chen · Jianning Li · Yu-Cheng Chou · Angtian Wang · Yixiong Chen · Yuxiang Lai · Jincheng Wang · Huimin Xue · Yining Cao · Haoqi Han · Xiaorui Lin · Yutong Tang · Meihua Li · Yujiu Ma · Jinghui Xu · Jiawei Liu · Zheyuan Zhang

Reference

Scaling Annotations

C. Qu, T. Zhang, H. Qiao, J. Liu, Y. Tang, A. Yuille, and Z. Zhou*. "Annotating 8,000 Abdominal CT Volumes for Multi-Organ Segmentation in Three Weeks." **NeurIPS 2023**. https://github.com/MrGiovanni/AbdomenAtlas

W. Li, A. Yuille, Z. Zhou*. "How Well Do Supervised 3D Models Transfer to Medical Imaging Tasks?" **ICLR 2024**. https://github.com/MrGiovanni/SuPreM

Scaling Datasets

Q. Hu, Y. Chen, J. Xiao, S. Sun, J. Chen, A. Yuille, Z. Zhou*. "Label-Free Liver Tumor Segmentation." **CVPR 2023**. https://github.com/MrGiovanni/SyntheticTumors

Q. Chen, X. Chen, H. Song, Z. Xiong, A. Yuille, C. Wei, Z. Zhou*. "Towards Generalizable Tumor Synthesis." **CVPR 2024**. https://github.com/MrGiovanni/DiffTumor

Y. Lai, X. Chen, A. Wang, A. Yuille, Z. Zhou*. "From Pixel to Cancer: Cellular Automata in Computed Tomography." **MICCAI 2024**. https://github.com/MrGiovanni/Pixel2Cancer

Scaling Algorithms

J. Liu, Y. Zhang, J. Chen, Y. Lu, Y. Yuan, A. Yuille, Y. Tang*, Z. Zhou*. "CLIP-Driven Universal Model for Organ Segmentation and Tumor Detection." **ICCV 2023**. https://github.com/ljwztc/CLIP-Driven-Universal-Model

Y. Zhang, X. Li, H. Chen, A. Yuille, Y. Liu*, Z. Zhou*. "Learning without Forgetting for Continual Abdominal Multi-Organ and Tumor Segmentation." **MICCAI 2023**. https://github.com/MrGiovanni/ContinualLearning