CSE 428 Computational Biology Capstone

What is research

- Engineering: Solve an existing problem
- Research: Propose a new problem
  - ✓ Define the problem
    - Problem setting (input, output)
    - Benchmark for validation
    - Baselines
  - ✓ Propose the first solution (which might not be a good solution)
  - ✓ This problem cannot be trivial (Significance)
  - ✓ This problem has not been studied before (Novel)
    - ✓ Why people have never studied it before?
  - ✓ Solving this problem is feasible (Feasible)

Why interdisciplinary research is impactful

- Lots of opportunities
- New methods from domain A can solve important but unaddressed problems in domain B.
  - Apply sequence modeling to study protein sequence (protein sequence is similar to text sequence)
  - Train foundation model/LLM to study protein sequence
  - Train multi-modal LLM to jointly study protein structure and protein sequence.

Borrow ideas from other domains

- From less novel to very novel
  - Use a method in one comp bio problem (e.g., protein structure) to solve another comp bio problem (e.g., chromatin structure)
  - Use a new ML technique from CV/NLP to solve a comp bio problem (CSE428 will focus on this)
  - Use a method in other subjects (quantum computing, theoretical physics) to solve a comp bio problem

What can you learn from research

- Career path for a CS student
  - SDE (no need to have PhD degree):
    - Prompt GPT4 for something.
  - Machine learning engineer (PhD degree is a big plus)
    - Fine-tune GPT4 for something: where to get the fine-tune data, which fine-tune techniques to use
    - Need substantial research ability to understand and modify existing algorithm
  - Research scientist/faculty (PhD is required)
    - Develop OpenAI GPT4
    - It is all about research
- CSE 428 will be at the scope of Machine learning engineer

CSE427: Computational methods for biology at different scales

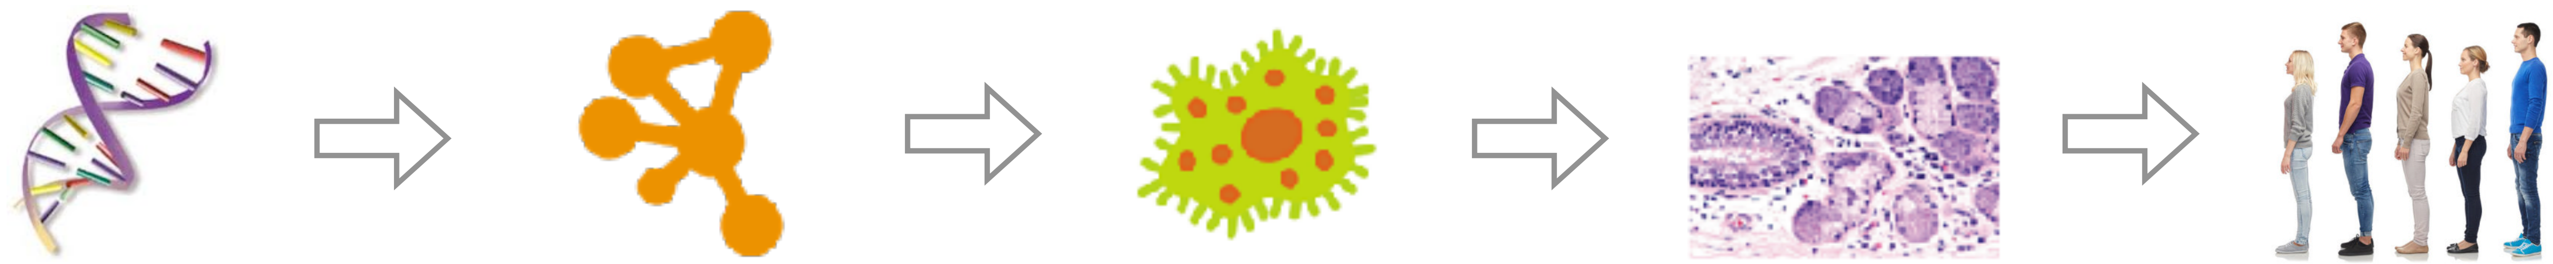

A rich hierarchy of biological subsystems at multiple scales: genotypic variations in nucleotides (1 nm scale) -> proteins (1–10 nm) -> protein complexes (10–100 nm), cellular processes (100 nm) -> phenotypic behaviors of cells (1–10 μm), tissues (100 μm to 100 mm), -> complex organisms (>1 m).

source: Yu, Michael Ku, et al. "Translation of genotype to phenotype by a hierarchy of cell subsystems." *Cell systems* 2.2 (2016): 77-88.

Computational methods for biology at different scales

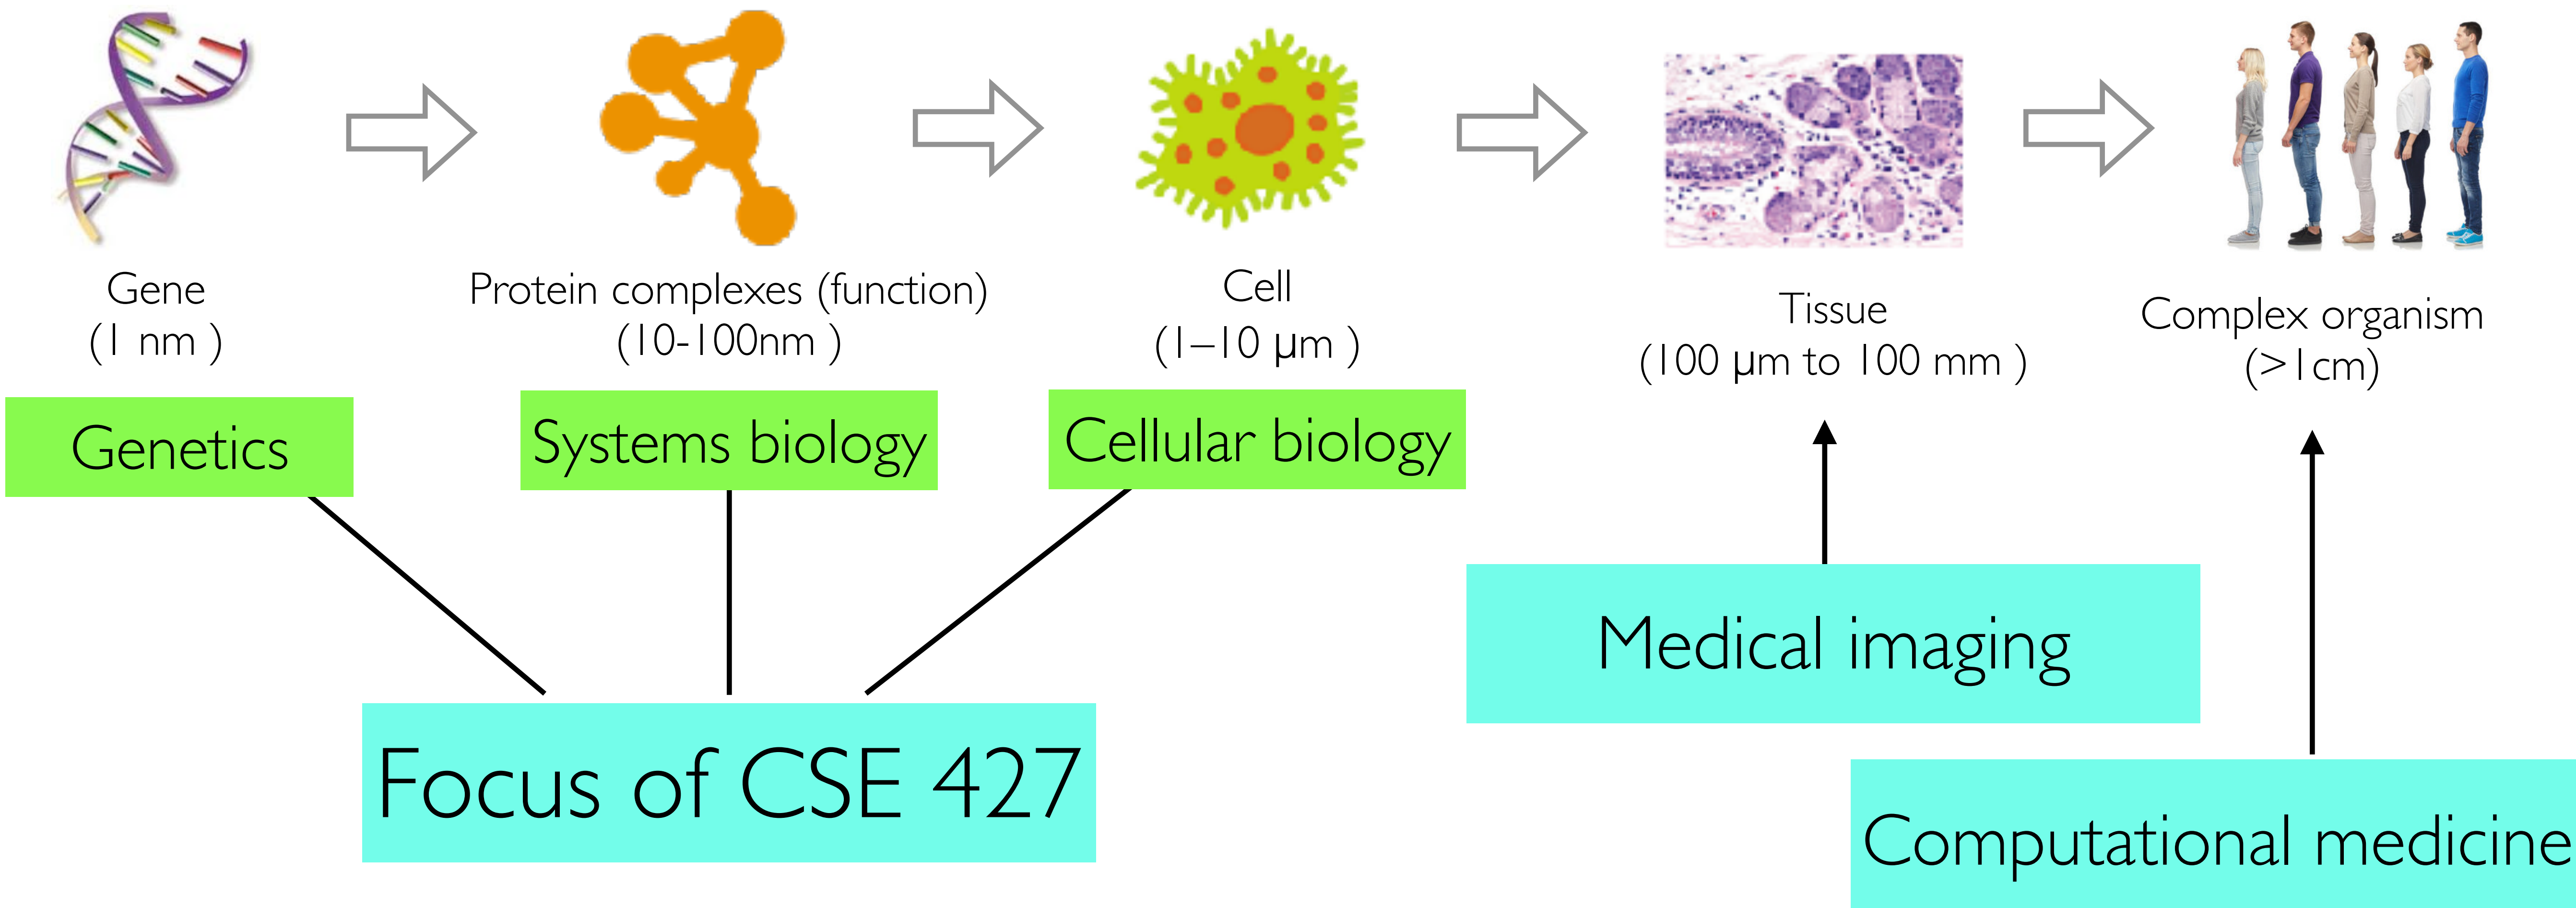

What can you learn from research

- Career path for a CS student
  - SDE (no need to have PhD degree):
    - Prompt GPT4 for something.
  - Machine learning engineer (PhD degree is a big plus)
    - Fine-tune GPT4 for something: where to get the fine-tune data, which fine-tune techniques to use
    - Need substantial research ability to understand and modify existing algorithm
  - Research scientist/faculty (PhD is required)
    - Develop OpenAI GPT4
    - It is all about research
- CSE 428 will be at the scope of Machine learning engineer

CSE 428 projects

- Option 1: define your own project and work by yourself
- Option 2: work on a suggested topic
  - Option 2A: pick a paper from Su-In Lee, Sara Mostafavi, Bill Noble
  - Option 2B: pick a paper from my lab (will introduce later)
  - Understand everything of the paper!
  - Reproduce some key results
  - Try to improve and modify the model

What you need to submit

- A course report
- A Jupyter notebook commit to the course GitHub
  - Mark down of a pipeline from reading data to produce prediction/analyses.

Schedule and grading

1. Three presentations
2. course project presentation (60%), final report (40%)

Tentative Schedule

Date	Topic
4/1	Welcome/overview. Introduction to CSE428. (Sheng)
4/8	Project topic presentation (first half)
4/15	Project topic presentation (second half)
4/22	Working time
4/29	Mid-term project idea presentation (first half)
5/6	Mid-term project idea presentation (second half)
5/13	Working time
5/20	Working time
5/27	Final project presentation (first half)
6/3	Final project presentation (second half)

First presentation: decide on a paper

- Choose one paper
- Focus on the significance
  - Why is this problem important
  - Why are you interested in this project.
- 5 minutes per student.
- More instruction later based on the number of students

How to read papers

- This is the most important thing
  - One researcher spends 3 hours per day to read paper
- Read new papers (2023-) and papers published in top journals/ conferences
  - Nature, Science, Cell, Recomb, ISMB
- Tips for how to read papers
  - Focus less on the background, introduction and motivation
  - Focus more on the problem setting (input, output)
  - Focus more on the datasets for evaluation (size, accessible, simulated data or real data)
- The more you read, the faster you will read.
- Ultimate goal: you can "predict" what method/dataset/evaluation this paper will do after reading the abstract

Pick the problem to work on

- Only work on the frontier methods
- Don't be ambitious!
- Only work on problem that an existing paper has worked on and you can fully understand that paper
- Only work on problem that you can clearly evaluate

Pick the problem to work on

- Pick a paper from Profs from Allen School
  - Su-In Lee, Sara Mostafavi, Sheng Wang, Bill Noble
- Read at least three papers in the first week
  - Each student will do a presentation next week
  - Significance
  - Innovation
  - Methods and evaluation
  - Why do you like this paper?

How to borrow ideas from other domains

- Find the commonality by understanding the problem setting
  - Problem A in computer vision has the same problem setting as Problem B in comp bio.
  - Talk to people work in CV/NLP

Experiments

1. Don't implement your ideas first
2. First, reproduce the results from baselines
3. Test baselines on new datasets and examine the performance
4. If the performance is bad, figuring out the reason

Uniqueness of comp bio research

- Interdisciplinary subject
  - Communication is very important
  - Understand biologists/doctors' needs
  - Clearly explain our methods
- New subject
  - Lots of opportunities
  - Less well-defined benchmarks

Reflection on my research career

- 2009-2013 Undergrad in CS, Peking university
- 2013-2018 PhD in CS, UIUC
- 2018-2020 Postdoc in Medicine, Stanford
- Most important thing is choosing the research area.

CSE427: Computational methods for biology at different scales

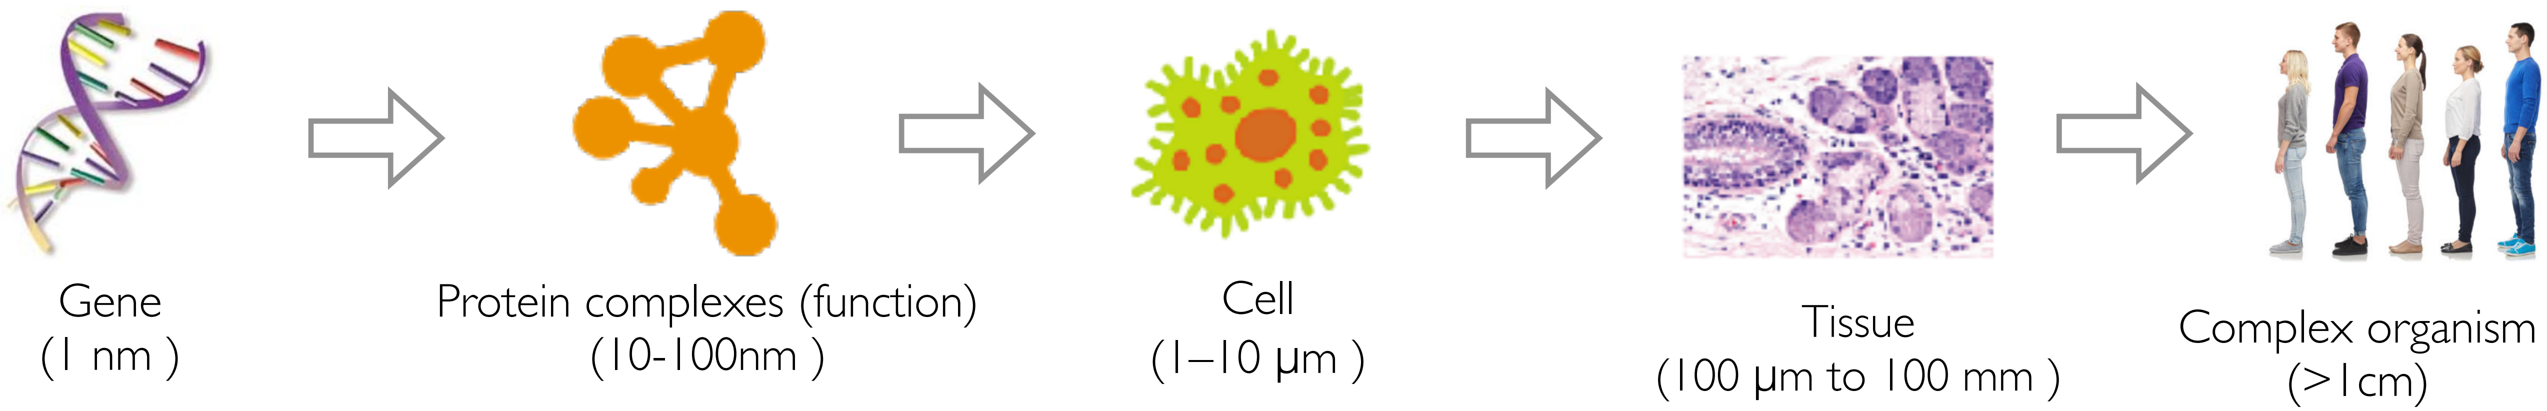

A rich hierarchy of biological subsystems at multiple scales: genotypic variations in nucleotides (1 nm scale) -> proteins (1–10 nm) -> protein complexes (10–100 nm), cellular processes (100 nm) -> phenotypic behaviors of cells (1–10 µm), tissues (100 µm to 100 mm), -> complex organisms (>1 m).

source: Yu, Michael Ku, et al. "Translation of genotype to phenotype by a hierarchy of cell subsystems." *Cell systems* 2.2 (2016): 77-88.

Computational methods for biology at different scales

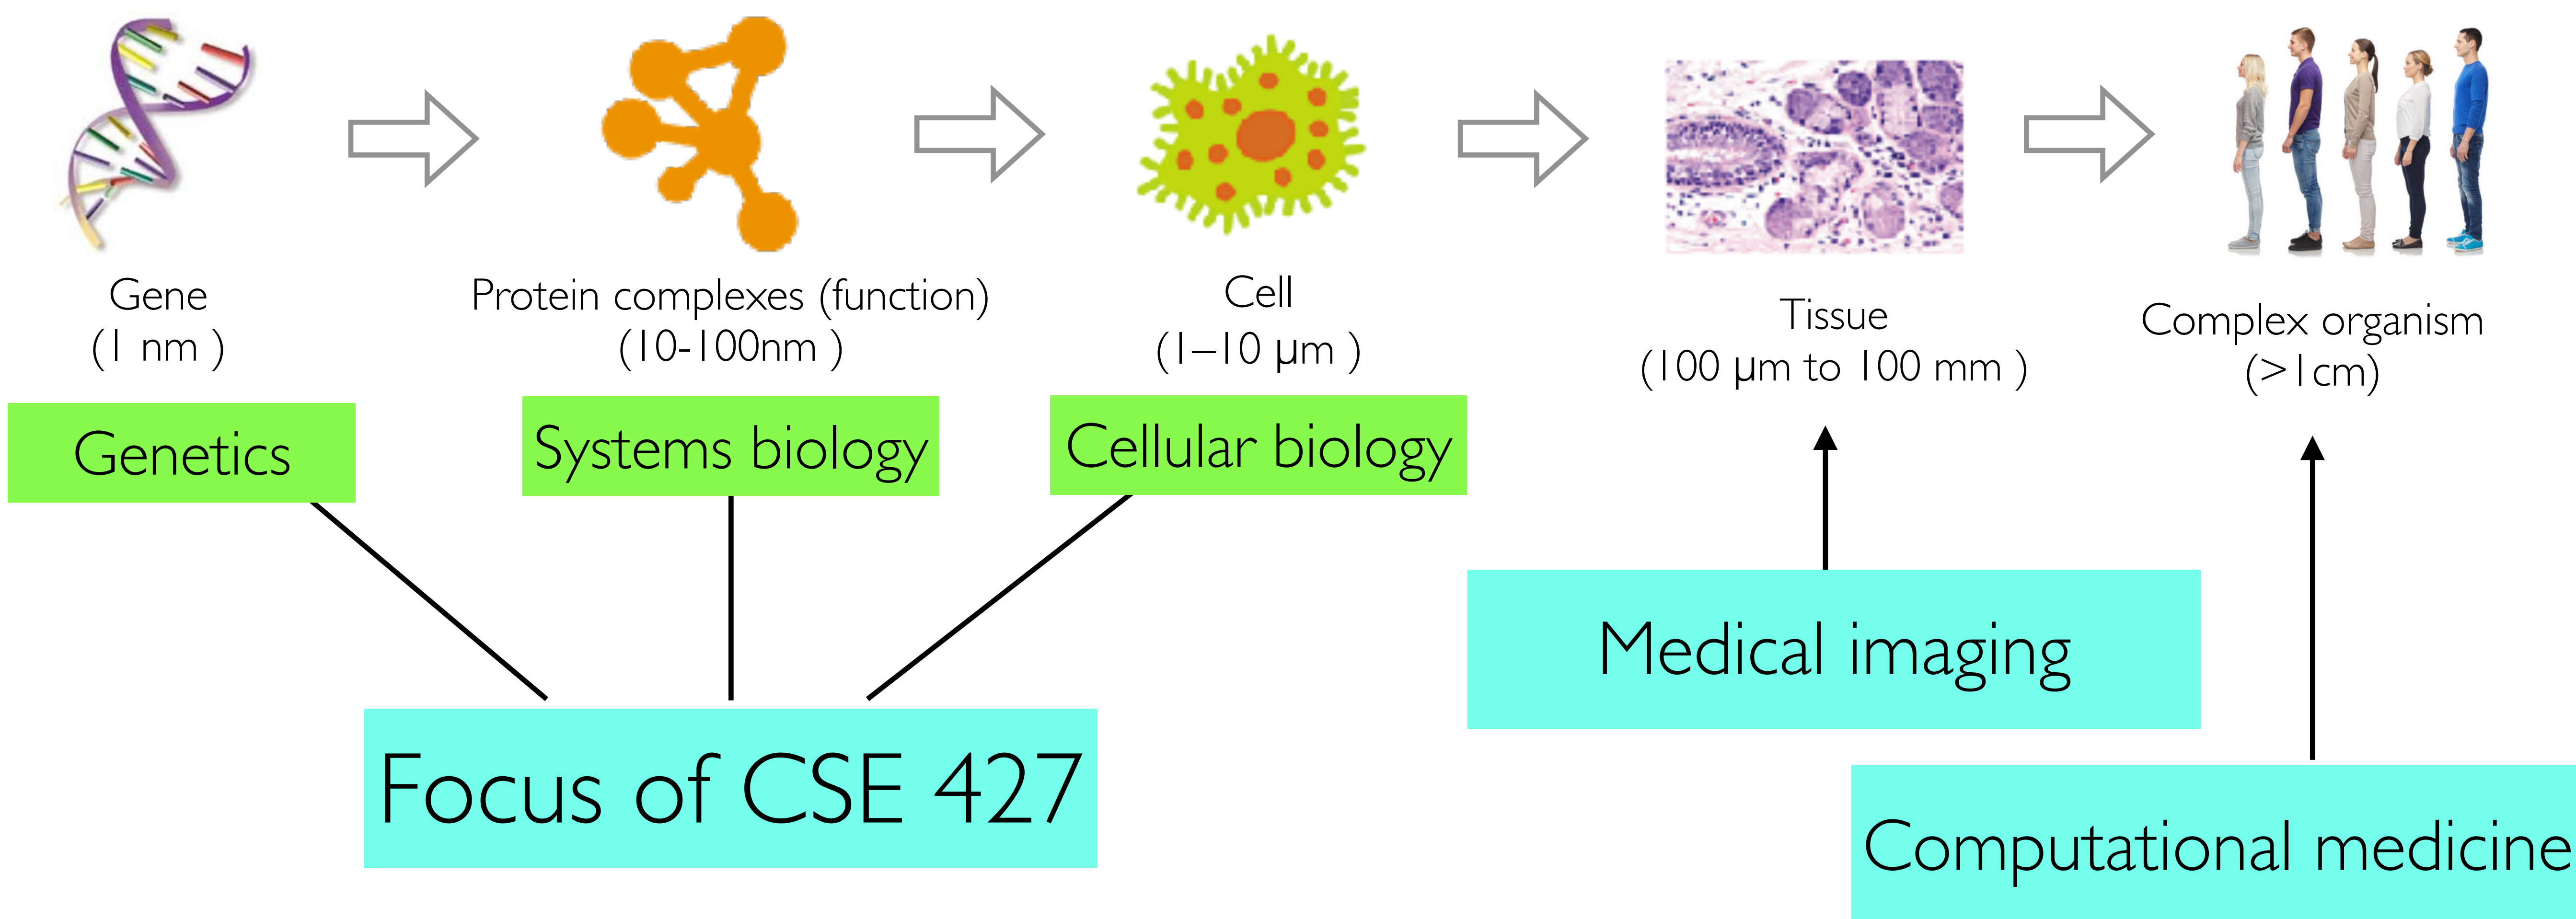

Four paradigms in AI for Medicine

Deep learning (2012)
One model for one task

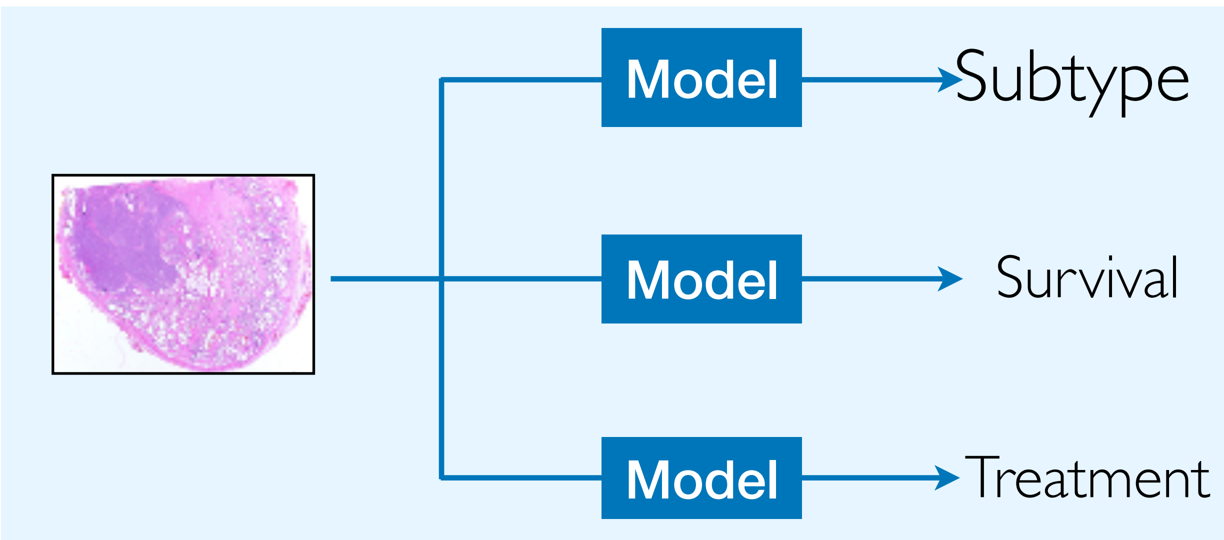

Foundation model (2022)
One model for all tasks

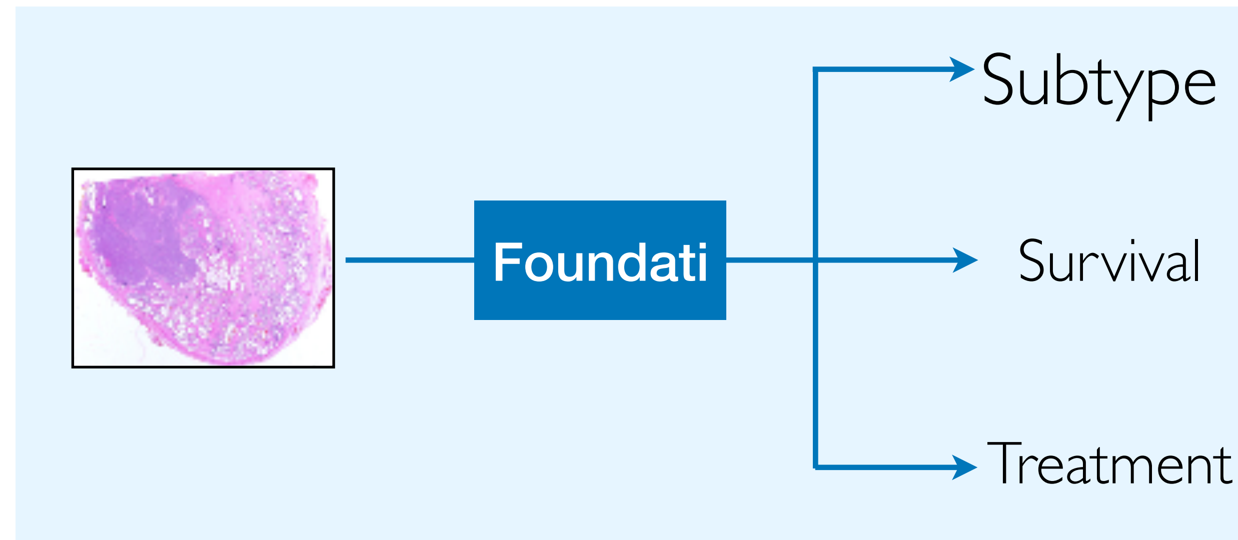

Multi-modal foundation model (2023)
One model takes different inputs

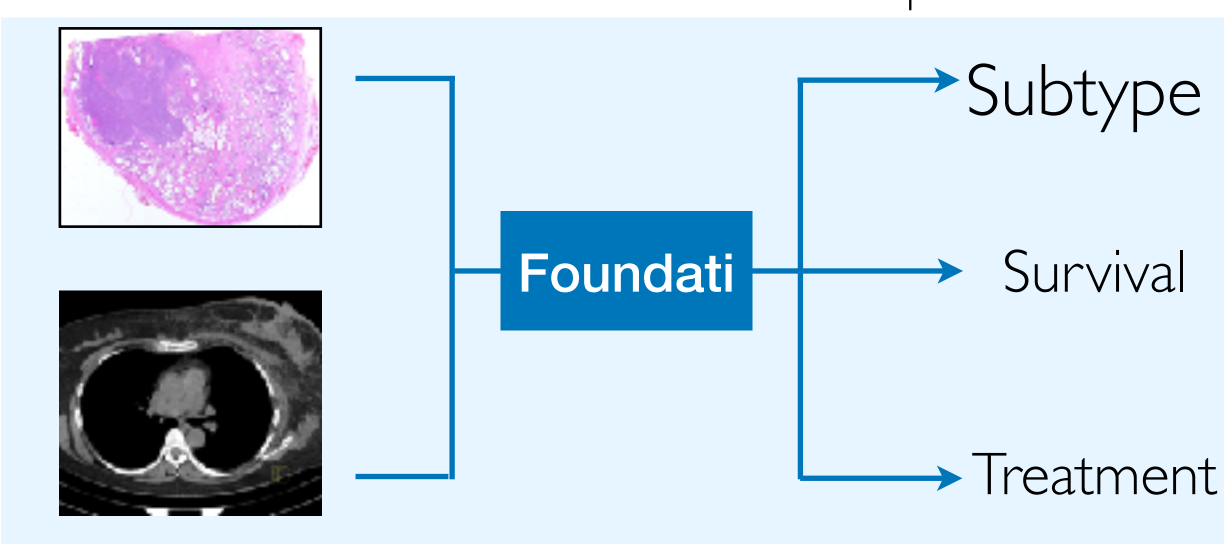

Multi-agent (2024)
Integrate multiple foundation models

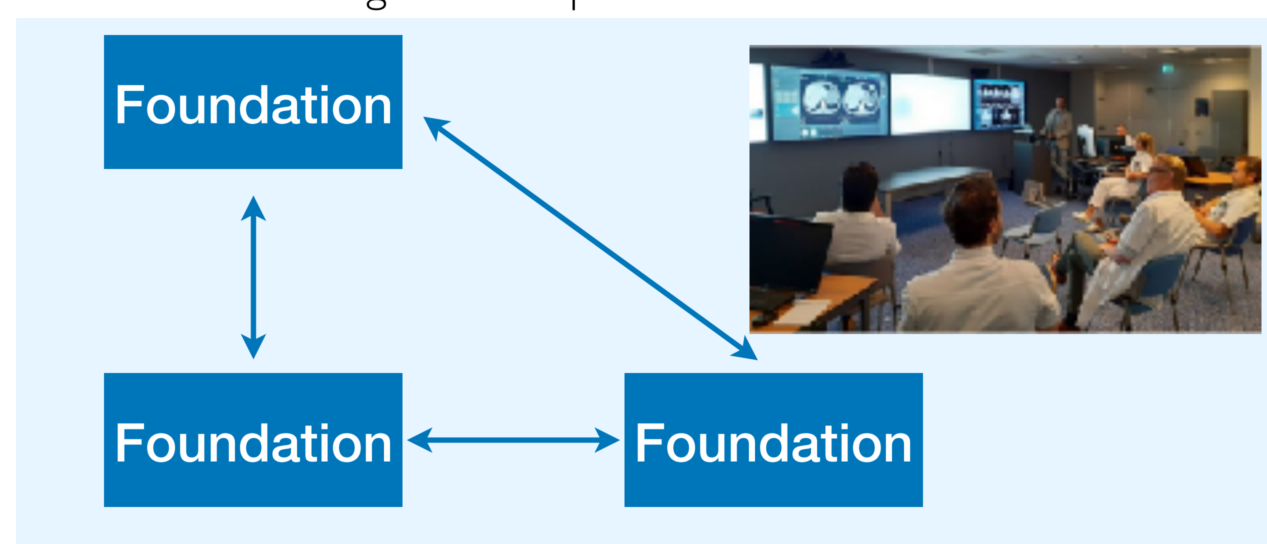

Medicine is inherently multi-modal

Complication of diabetes

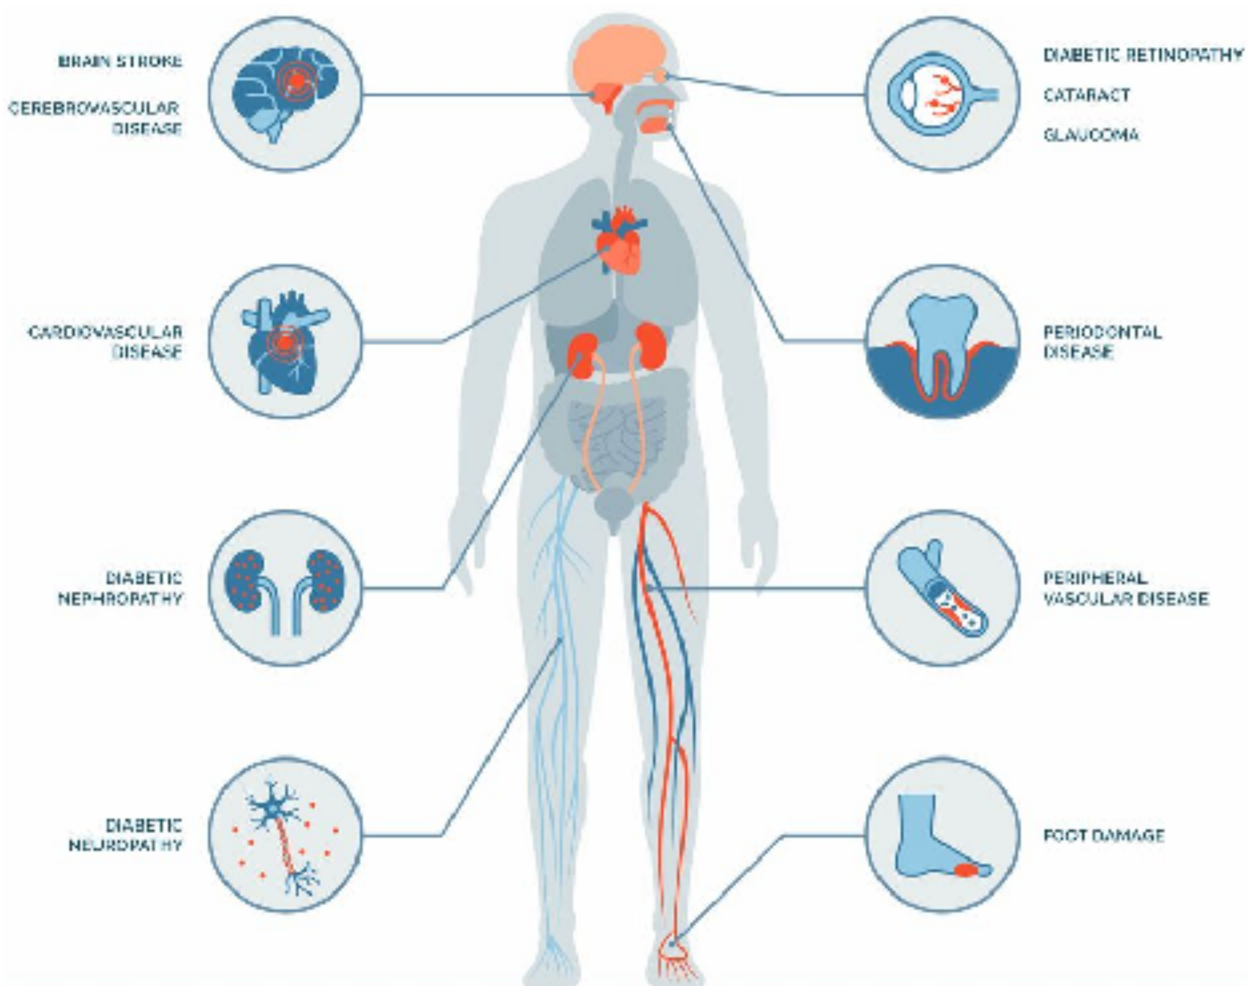

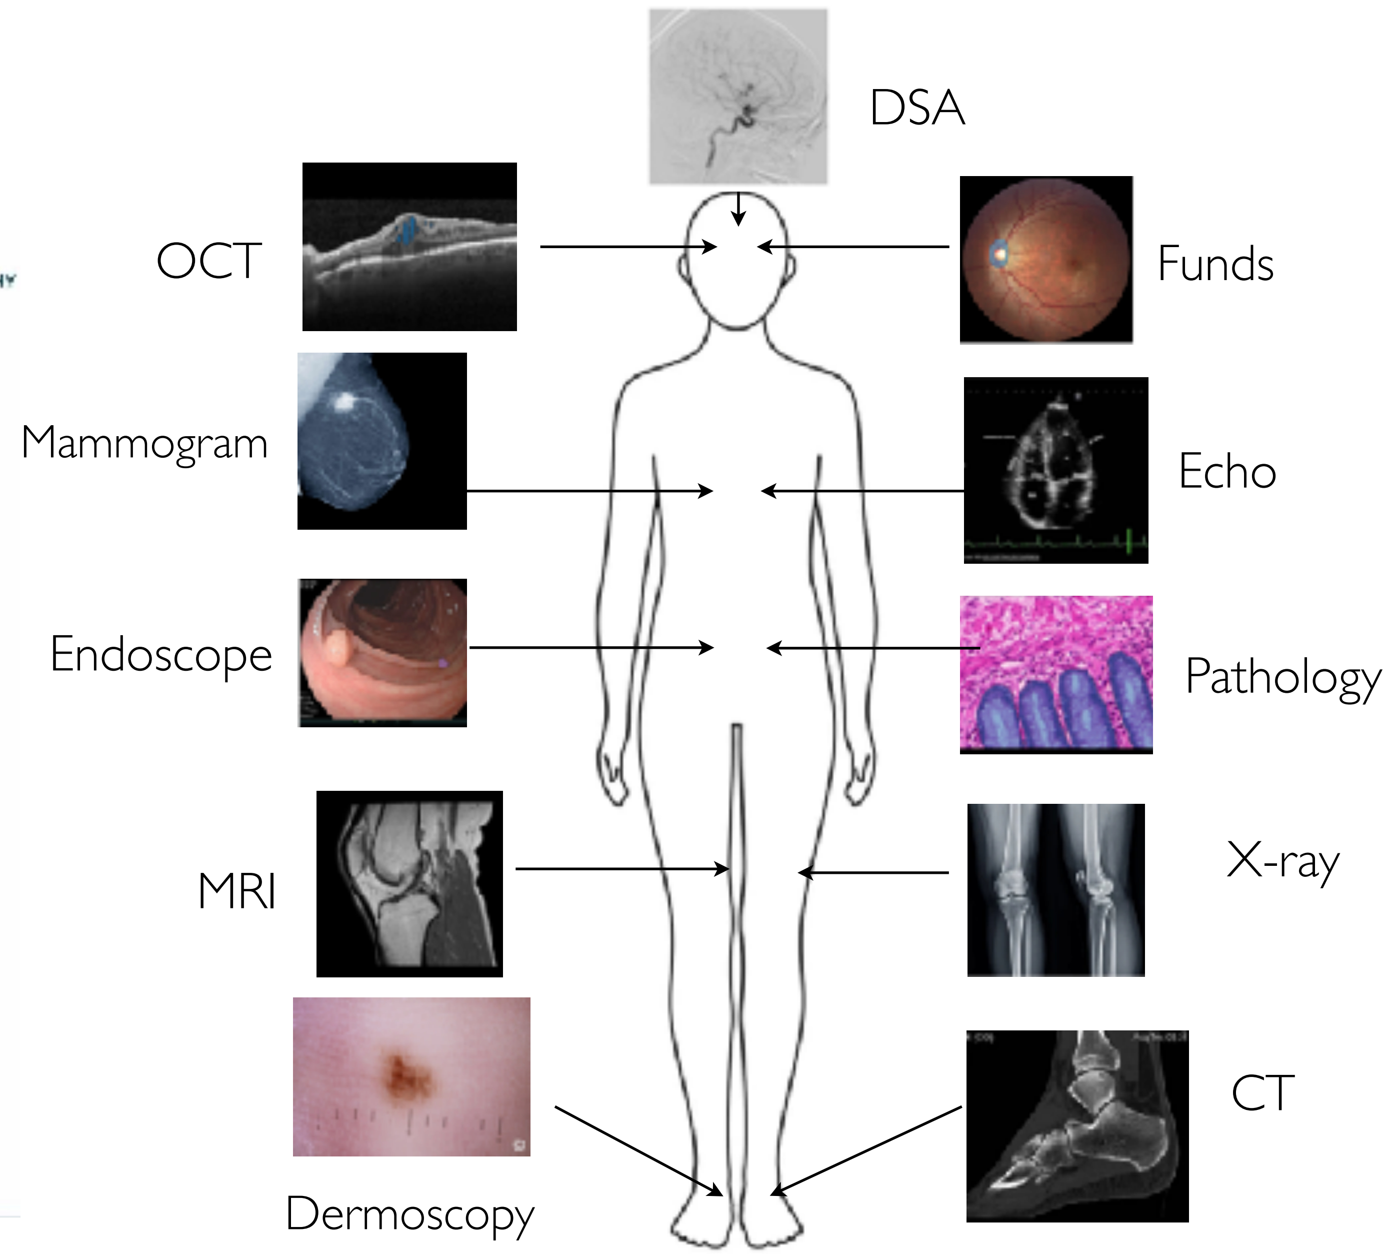

Today's talk: 3 parts

1. Pathology foundation model
2. 3D retinal foundation model
3. A multi-modal foundation model integrating 9 imaging modalities

Foundation model (2022)
One model for all tasks

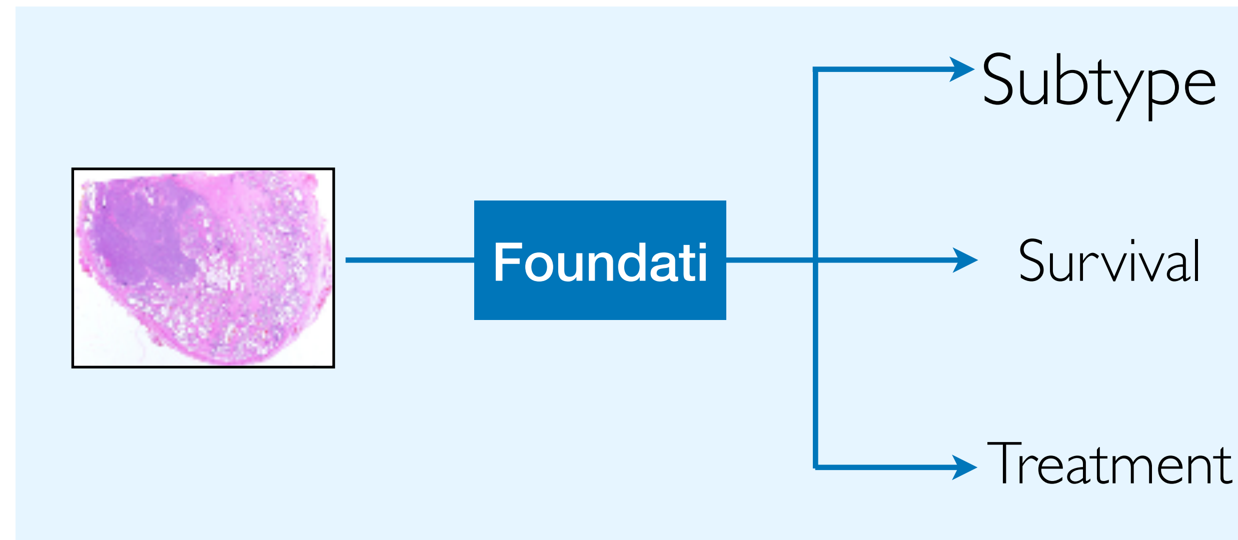

Multi-modal foundation model (2023)
One model takes different inputs

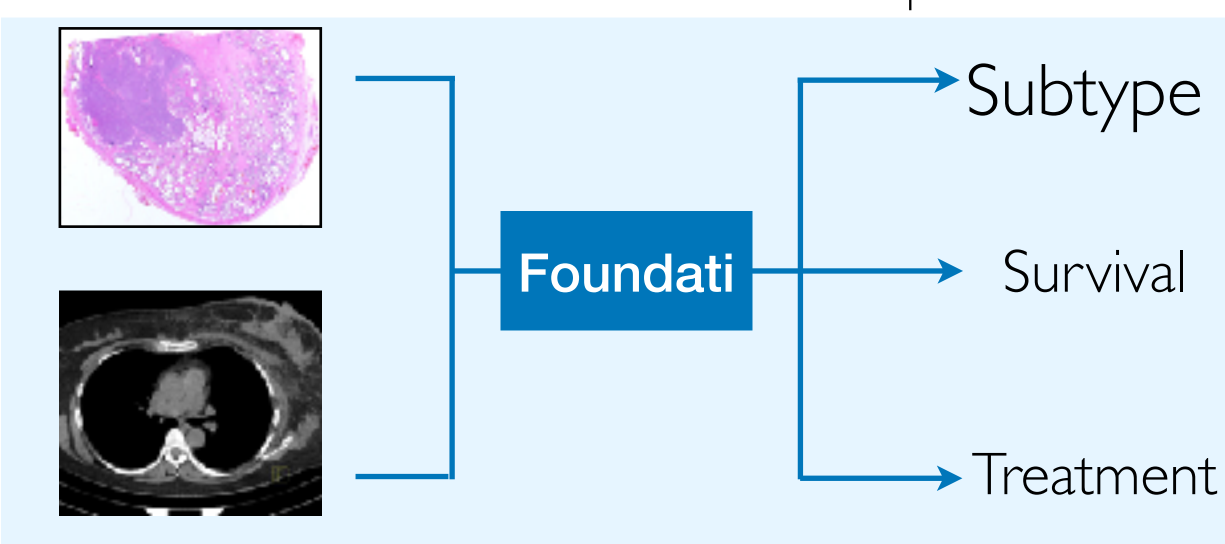

Multi-agent (2024)
Integrate multiple foundation models

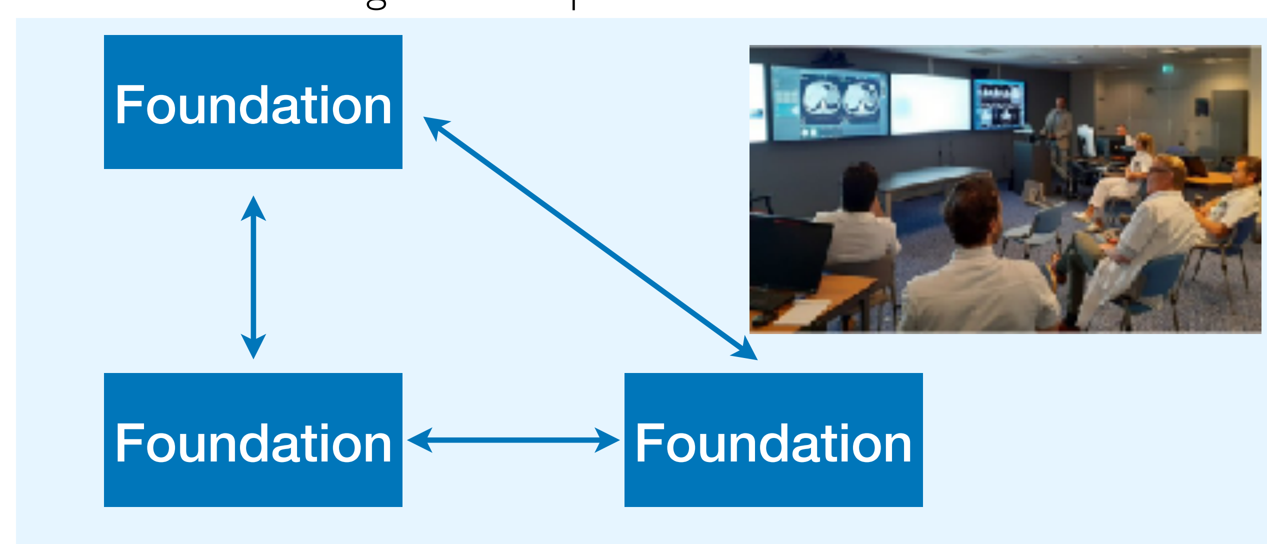

Pathology images are too large for existing AI models

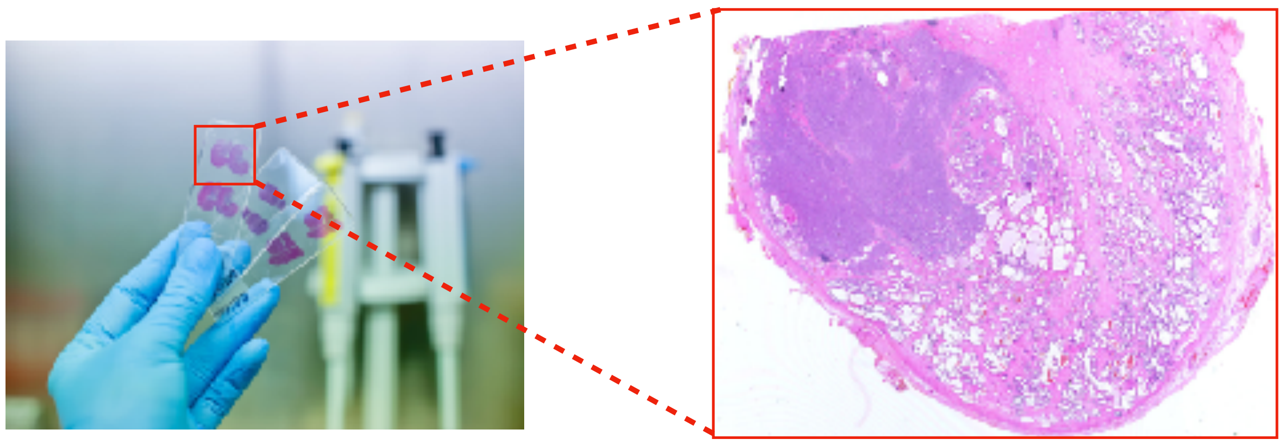

Pathology images
100k by 100k pixels

Pathology images are too large for existing AI models

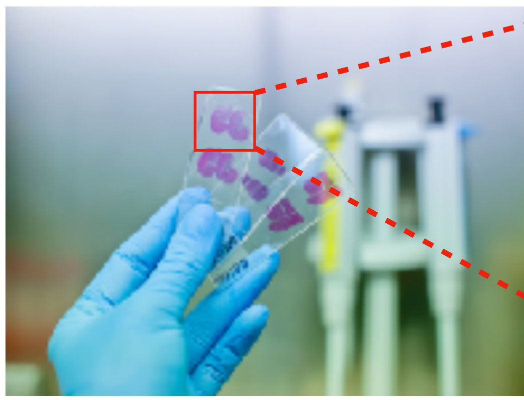

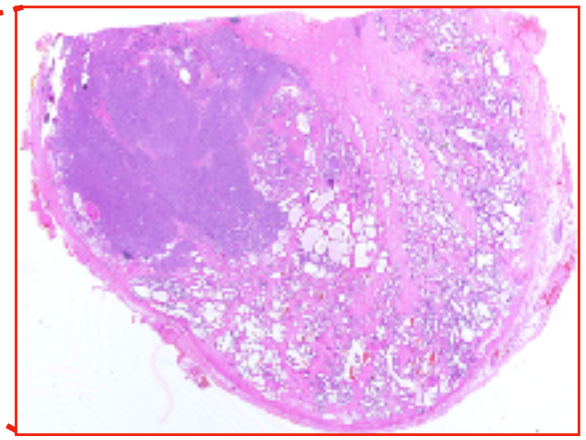

Pathology images
100k by 100k pixels

As large as a tennis court

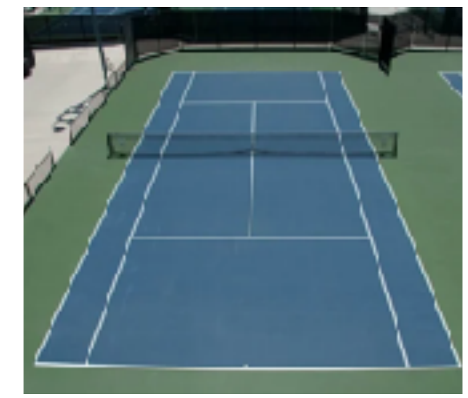

Pathology images are too large for existing AI models

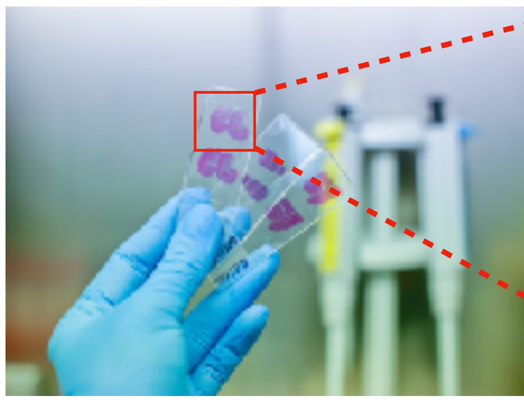

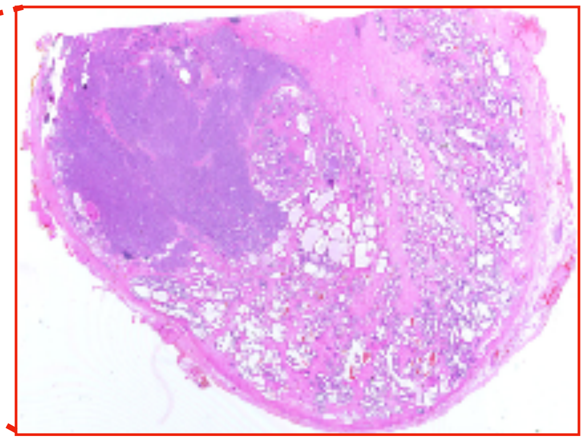

Pathology images
100k by 100k pixels

As large as a tennis court

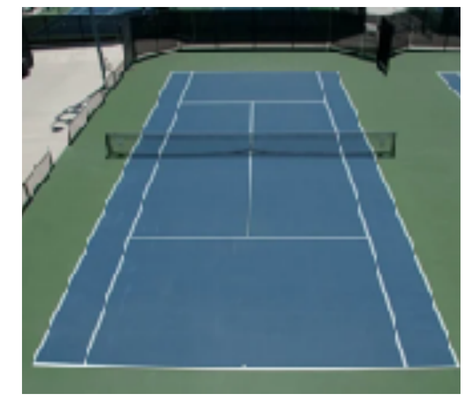

Images handled by existing AI models
256 by 256 pixels

As large as a tennis ball

Long-context modeling: ChatGPT is good at understanding long documents

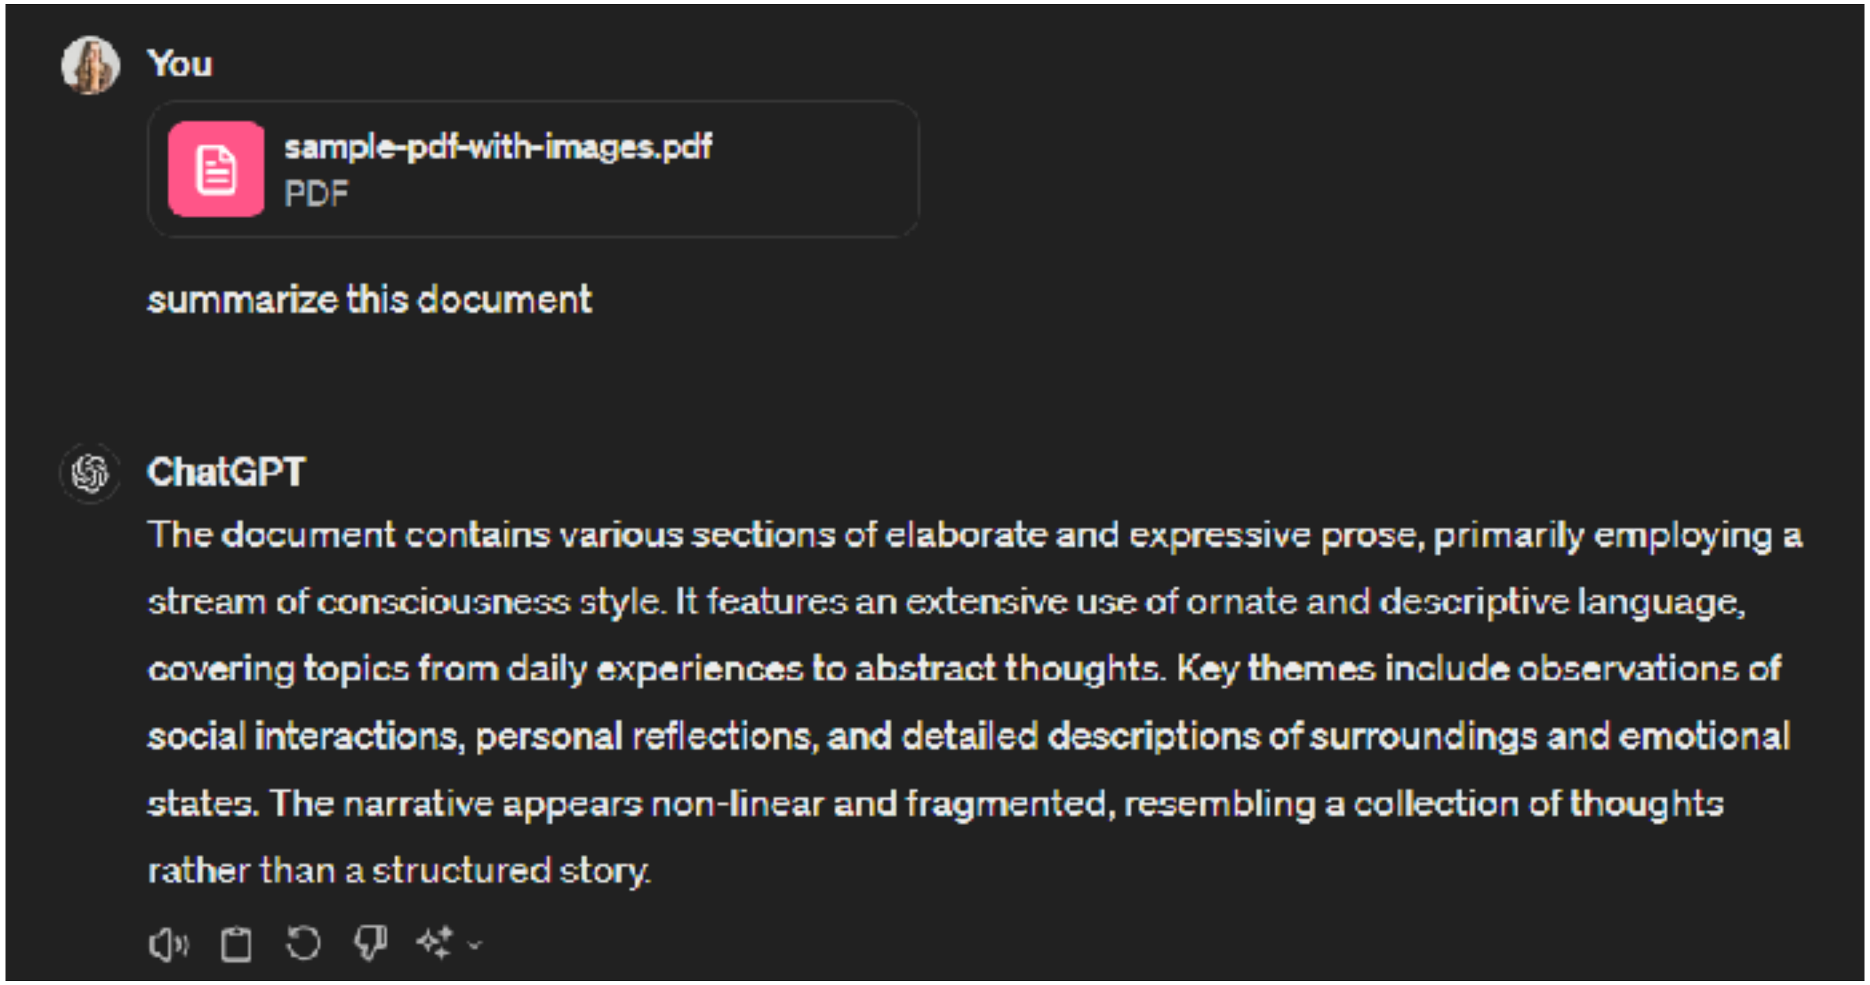

Long-context modeling: ChatGPT is good at understanding long documents

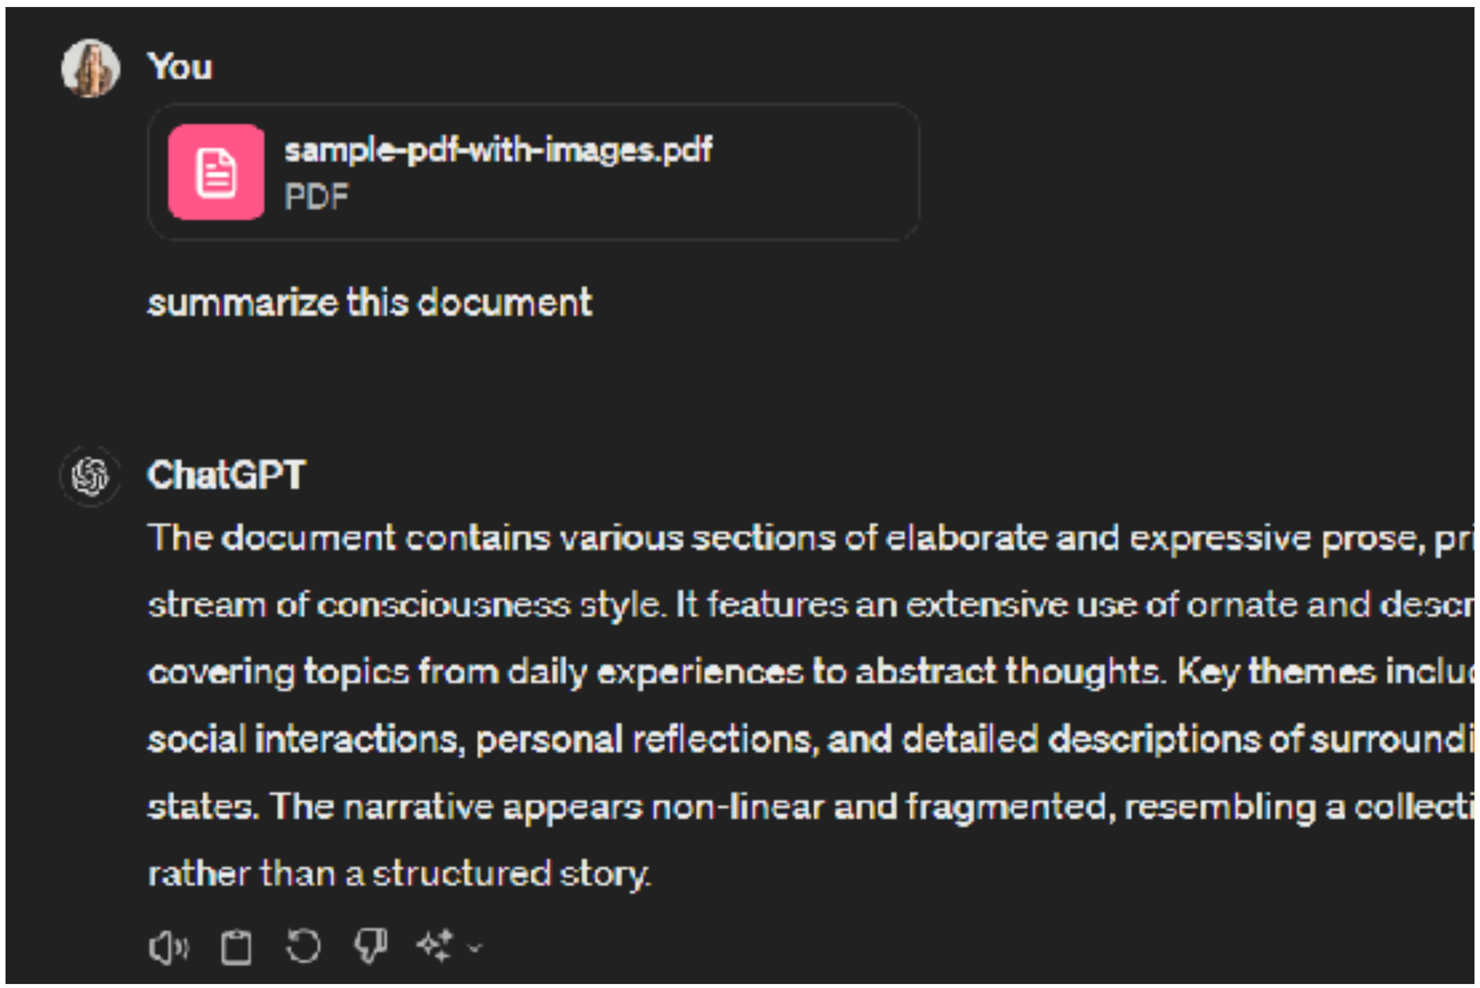

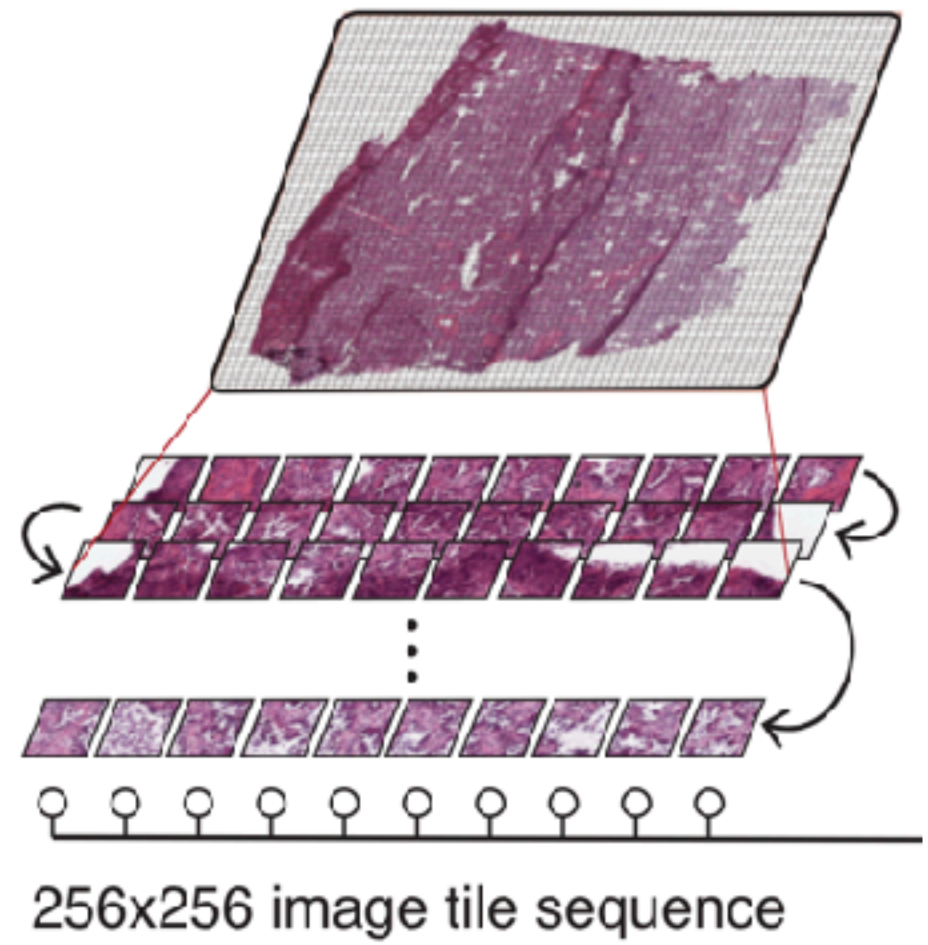

Use LongNet to model a **long sentence** of small images from a pathology slide

Tile-level encoder to capture local patterns
Slide-level encoder to capture the pattern in the whole slide

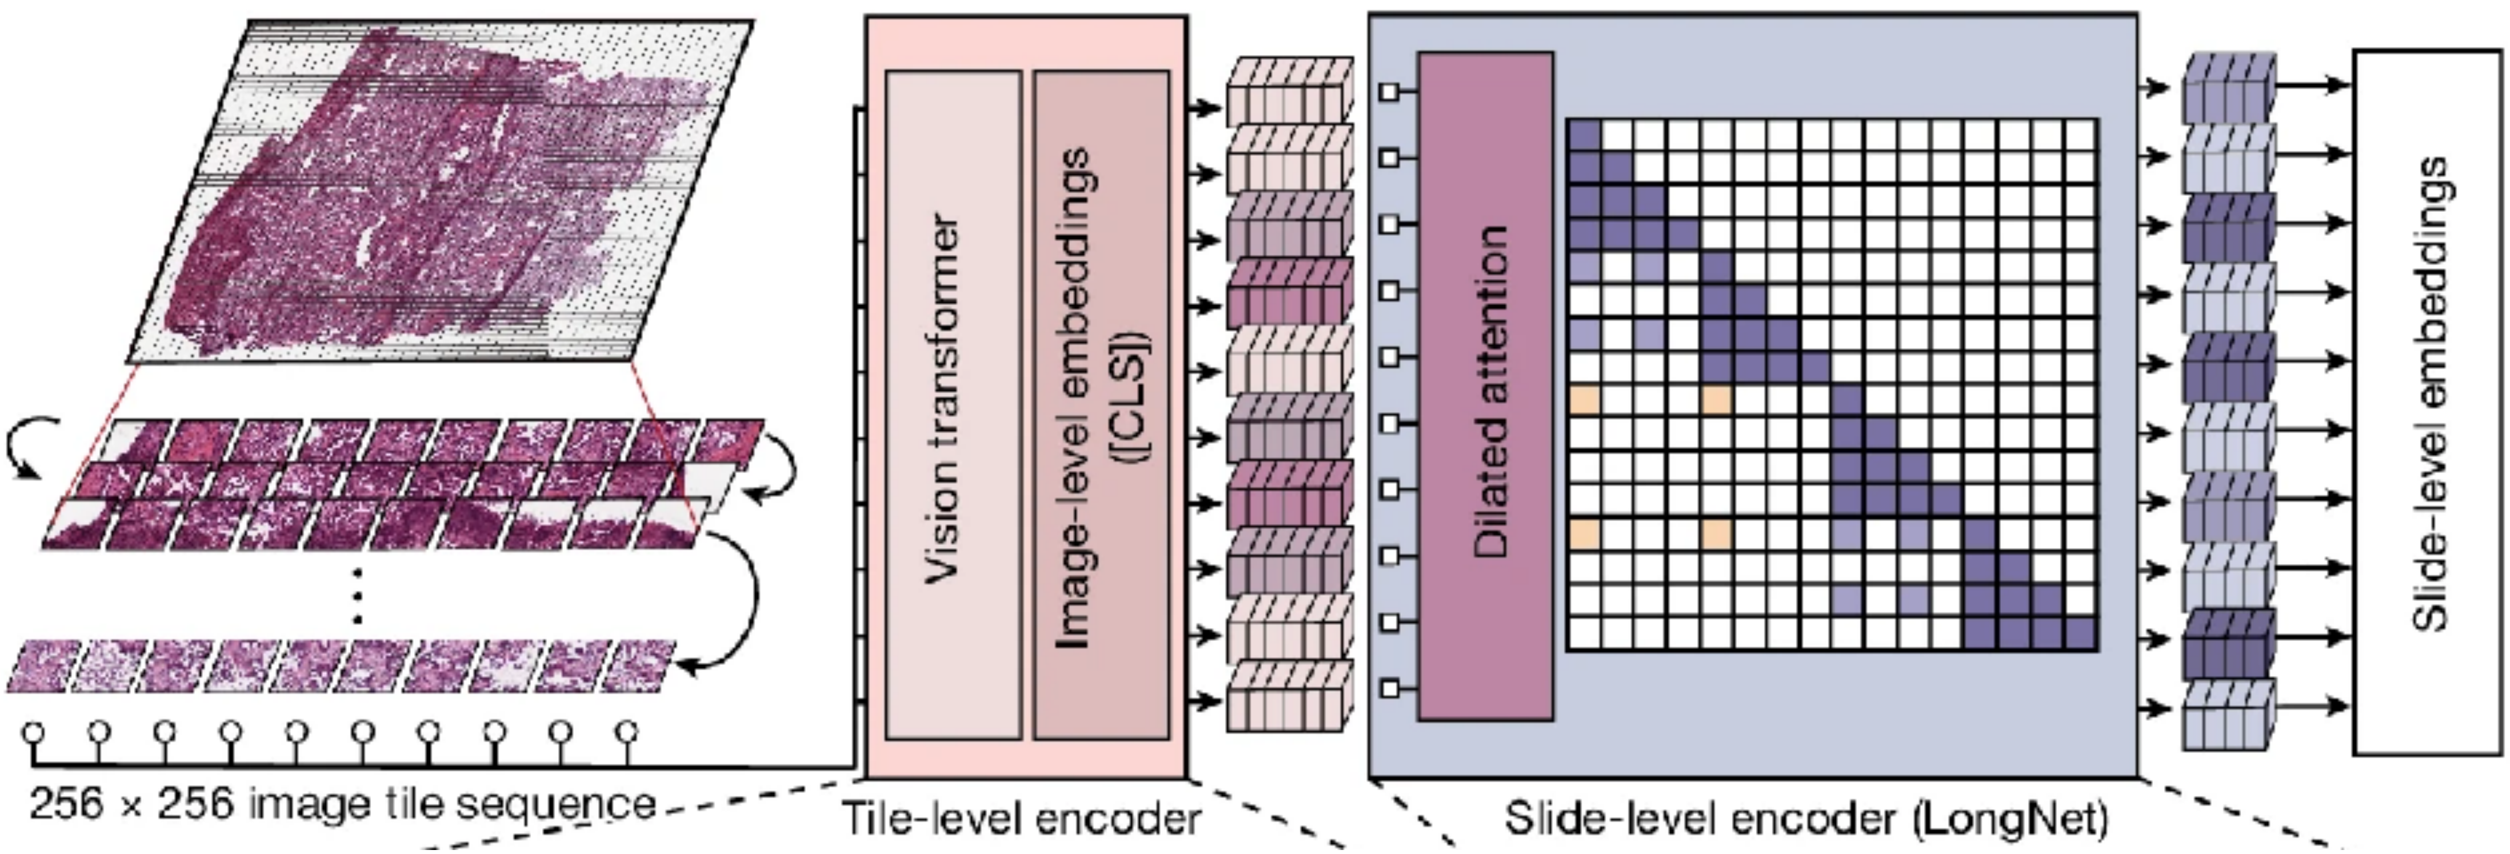

A self-supervised learning framework based on DinoV2 and LongNet

GigaPath: A whole-slide foundation model for pathology

nature

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › articles › article

Article | Open access | Published: 22 May 2024

A whole-slide foundation model for digital pathology from real-world data

Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, Yanbo Xu, Mu Wei, Wenhui Wang, Shuming Ma, Furu Wei, Jianwei Yang, Chunyuan Li, Jianfeng Gao, Jaylen Rosemon, Tucker Bower, Soohee Lee, Roshanthi Weerasinghe, Bill J. Wright, Ari Robicsek, Brian Piening, Carlo Bifulco, Sheng Wang & Hoifung Poon — Show fewer authors

Nature 630, 181–188 (2024) | Cite this article

200k model downloads every month (Hugging Face)!

Media coverage: Forbes, Yahoo, Becker's hospital review, Fierce biotech, CTOL digital solutions, HIT consultant, GeekWire, Cosmic log, HealthXL, RamaOnHealthcare, Providence, nikkei, cryptorank, deeptech

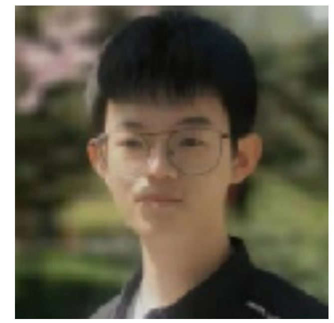

Hanwen Xu
U of Washington

Naoto Usuyama
Microsoft Research

Carlo Bifulco
Providence

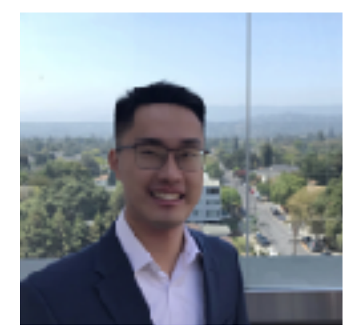

Sheng Wang
U of Washington

Hoifung Poon
Microsoft Research

One (foundation) model fits all

30,000 patients from 28 hospitals
140k slides, 15 cancer types

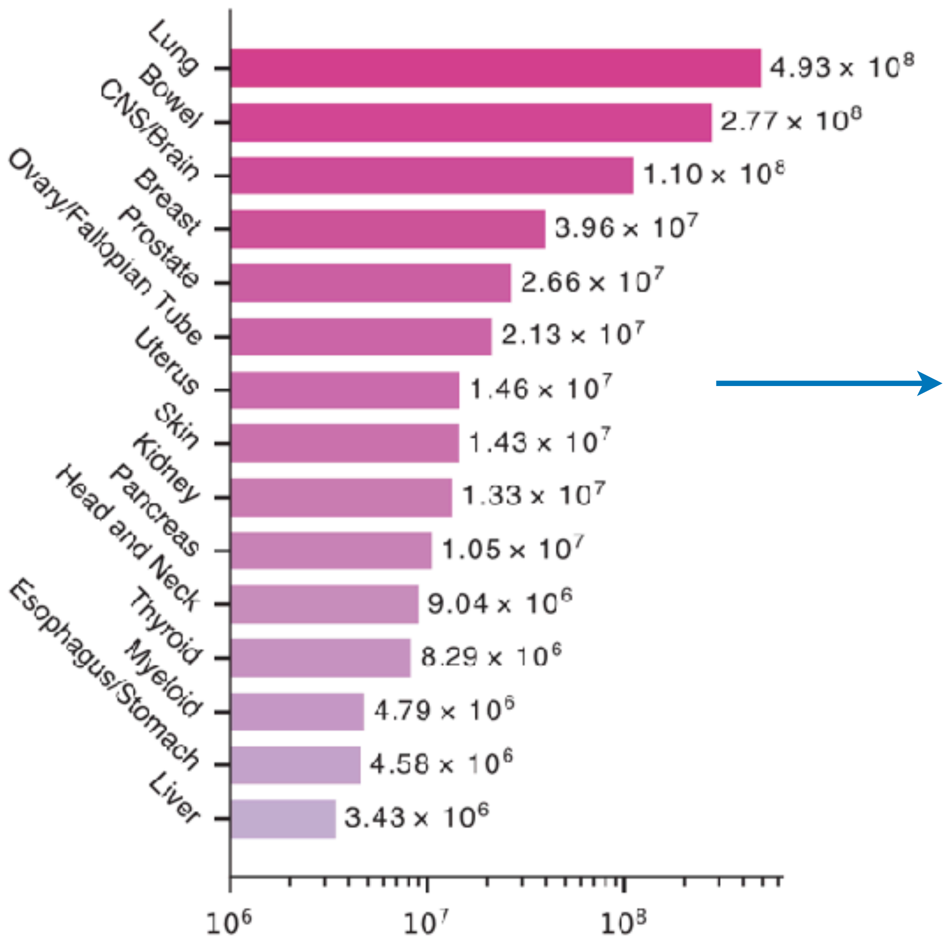

GigaPa

Best performance on 25 out of 26 tasks

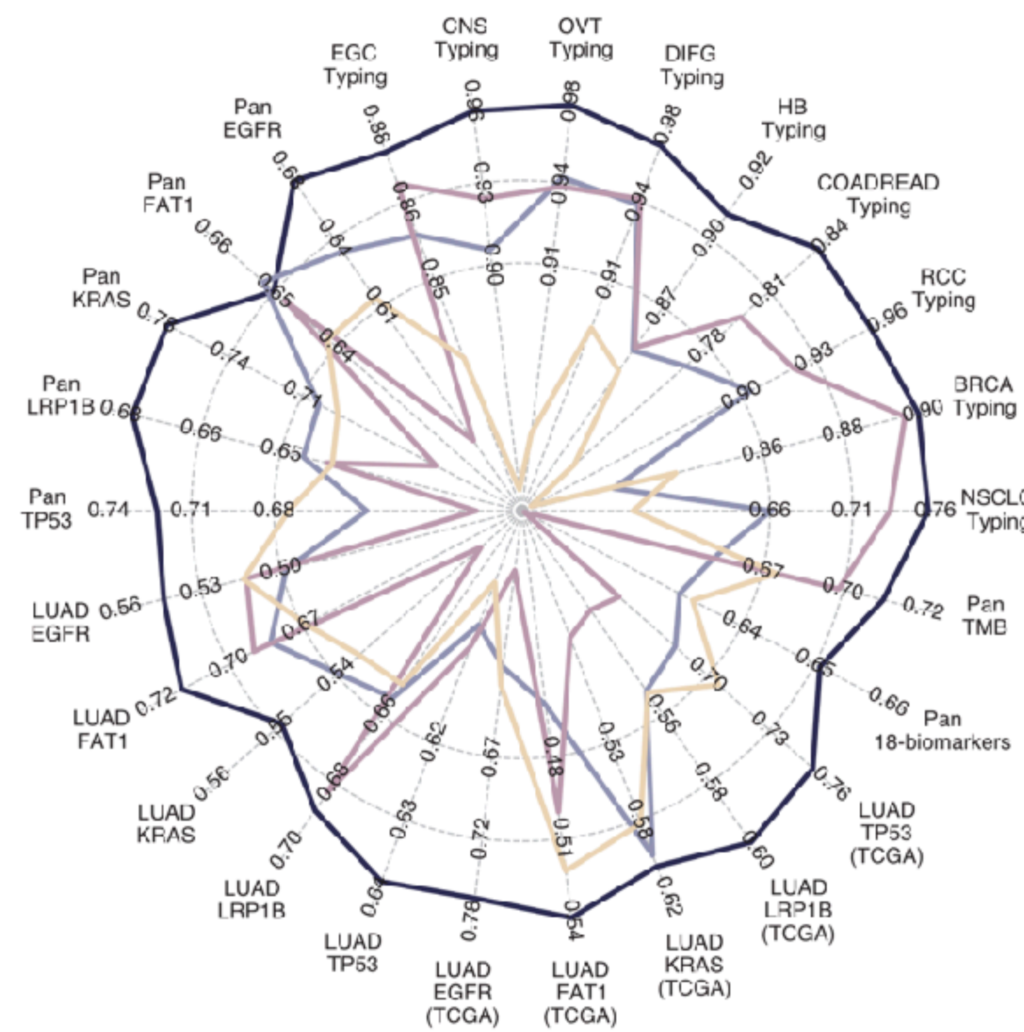

Cancer subtype classification using pathology images

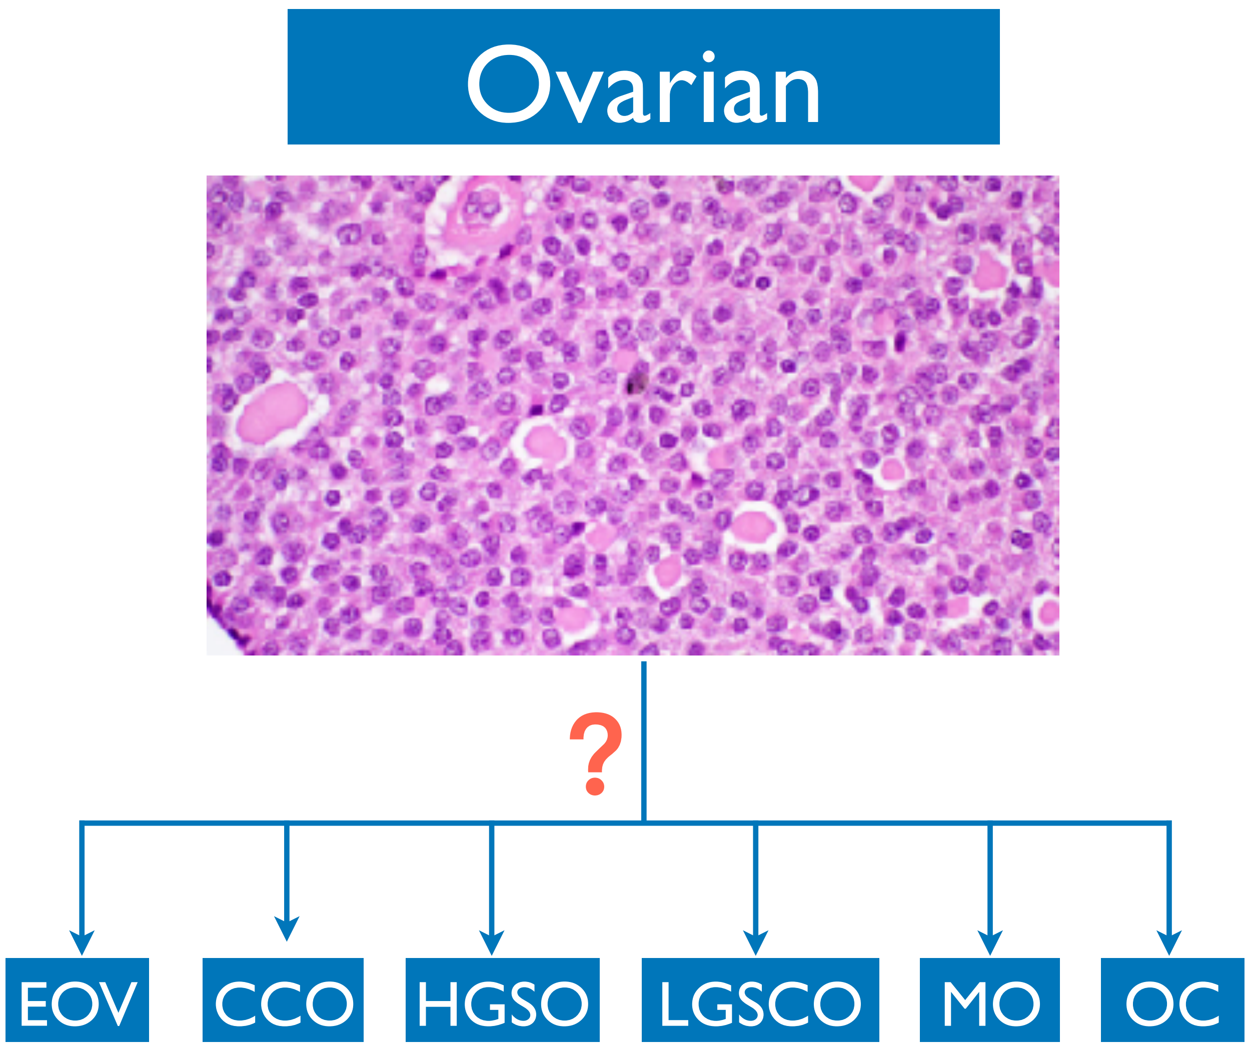

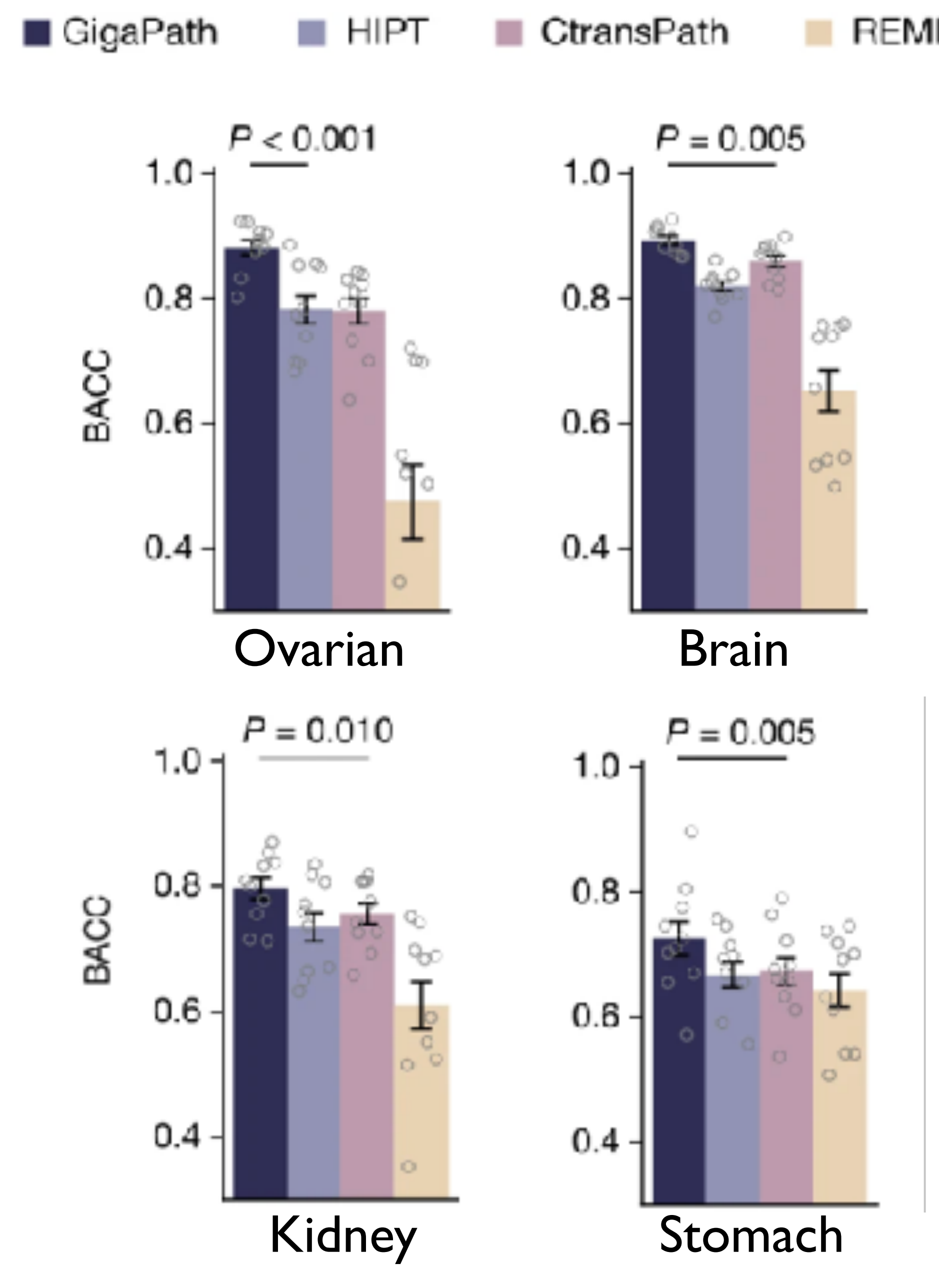

Biomarker prediction using pathology images: best performance on 17 out of 18 biomarker predictions

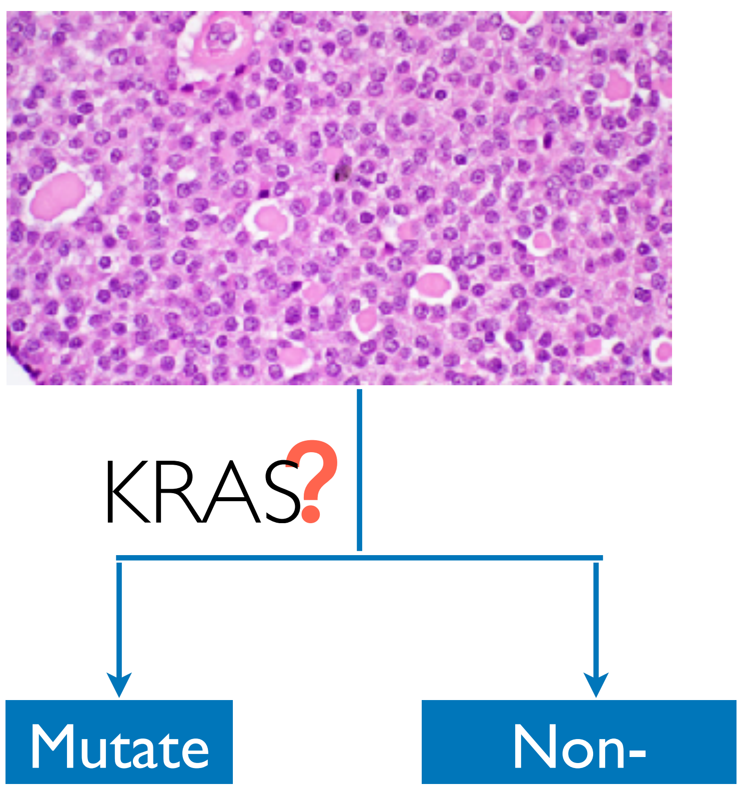

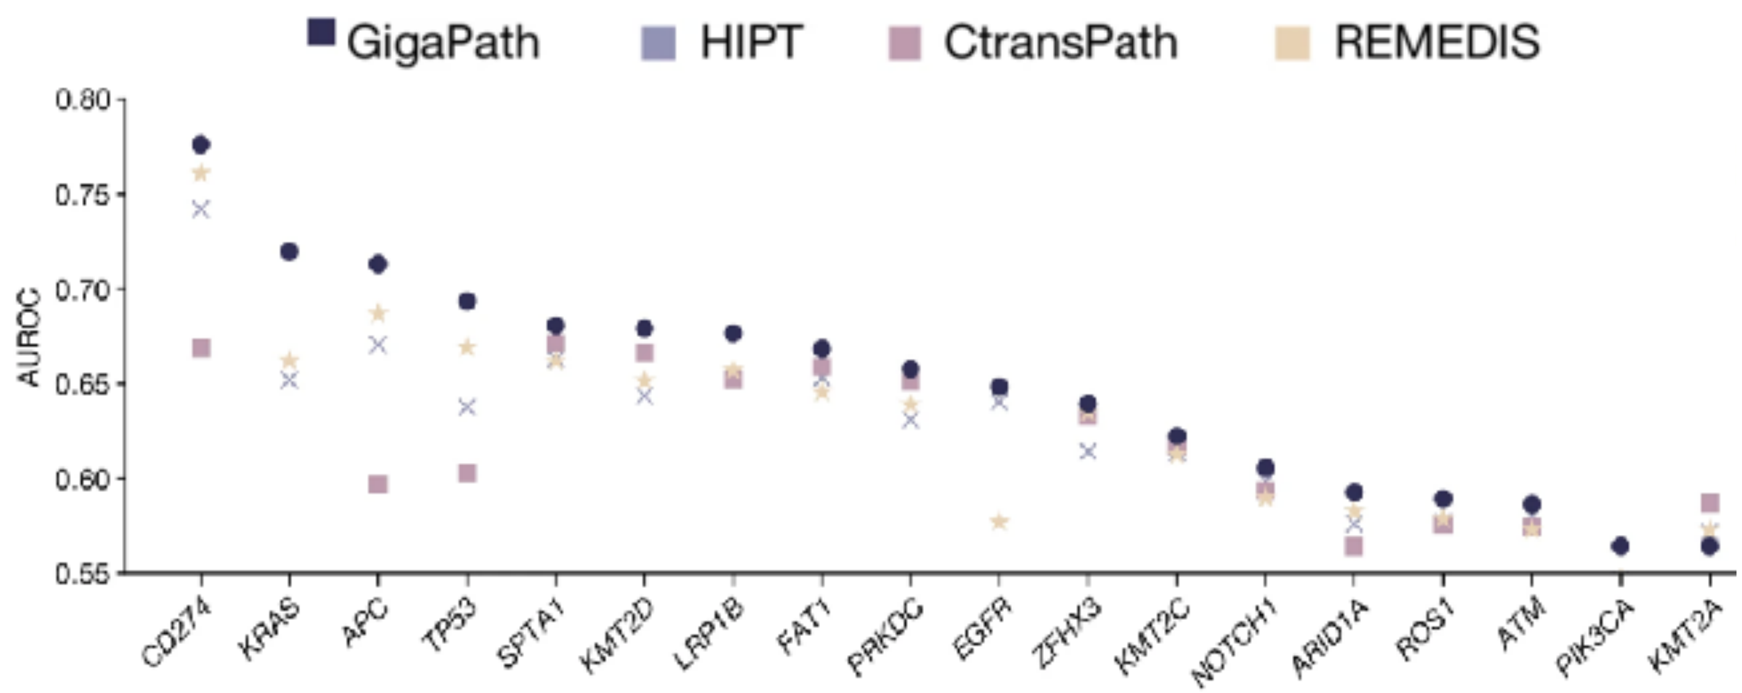

Directly predicting treatment is too difficult

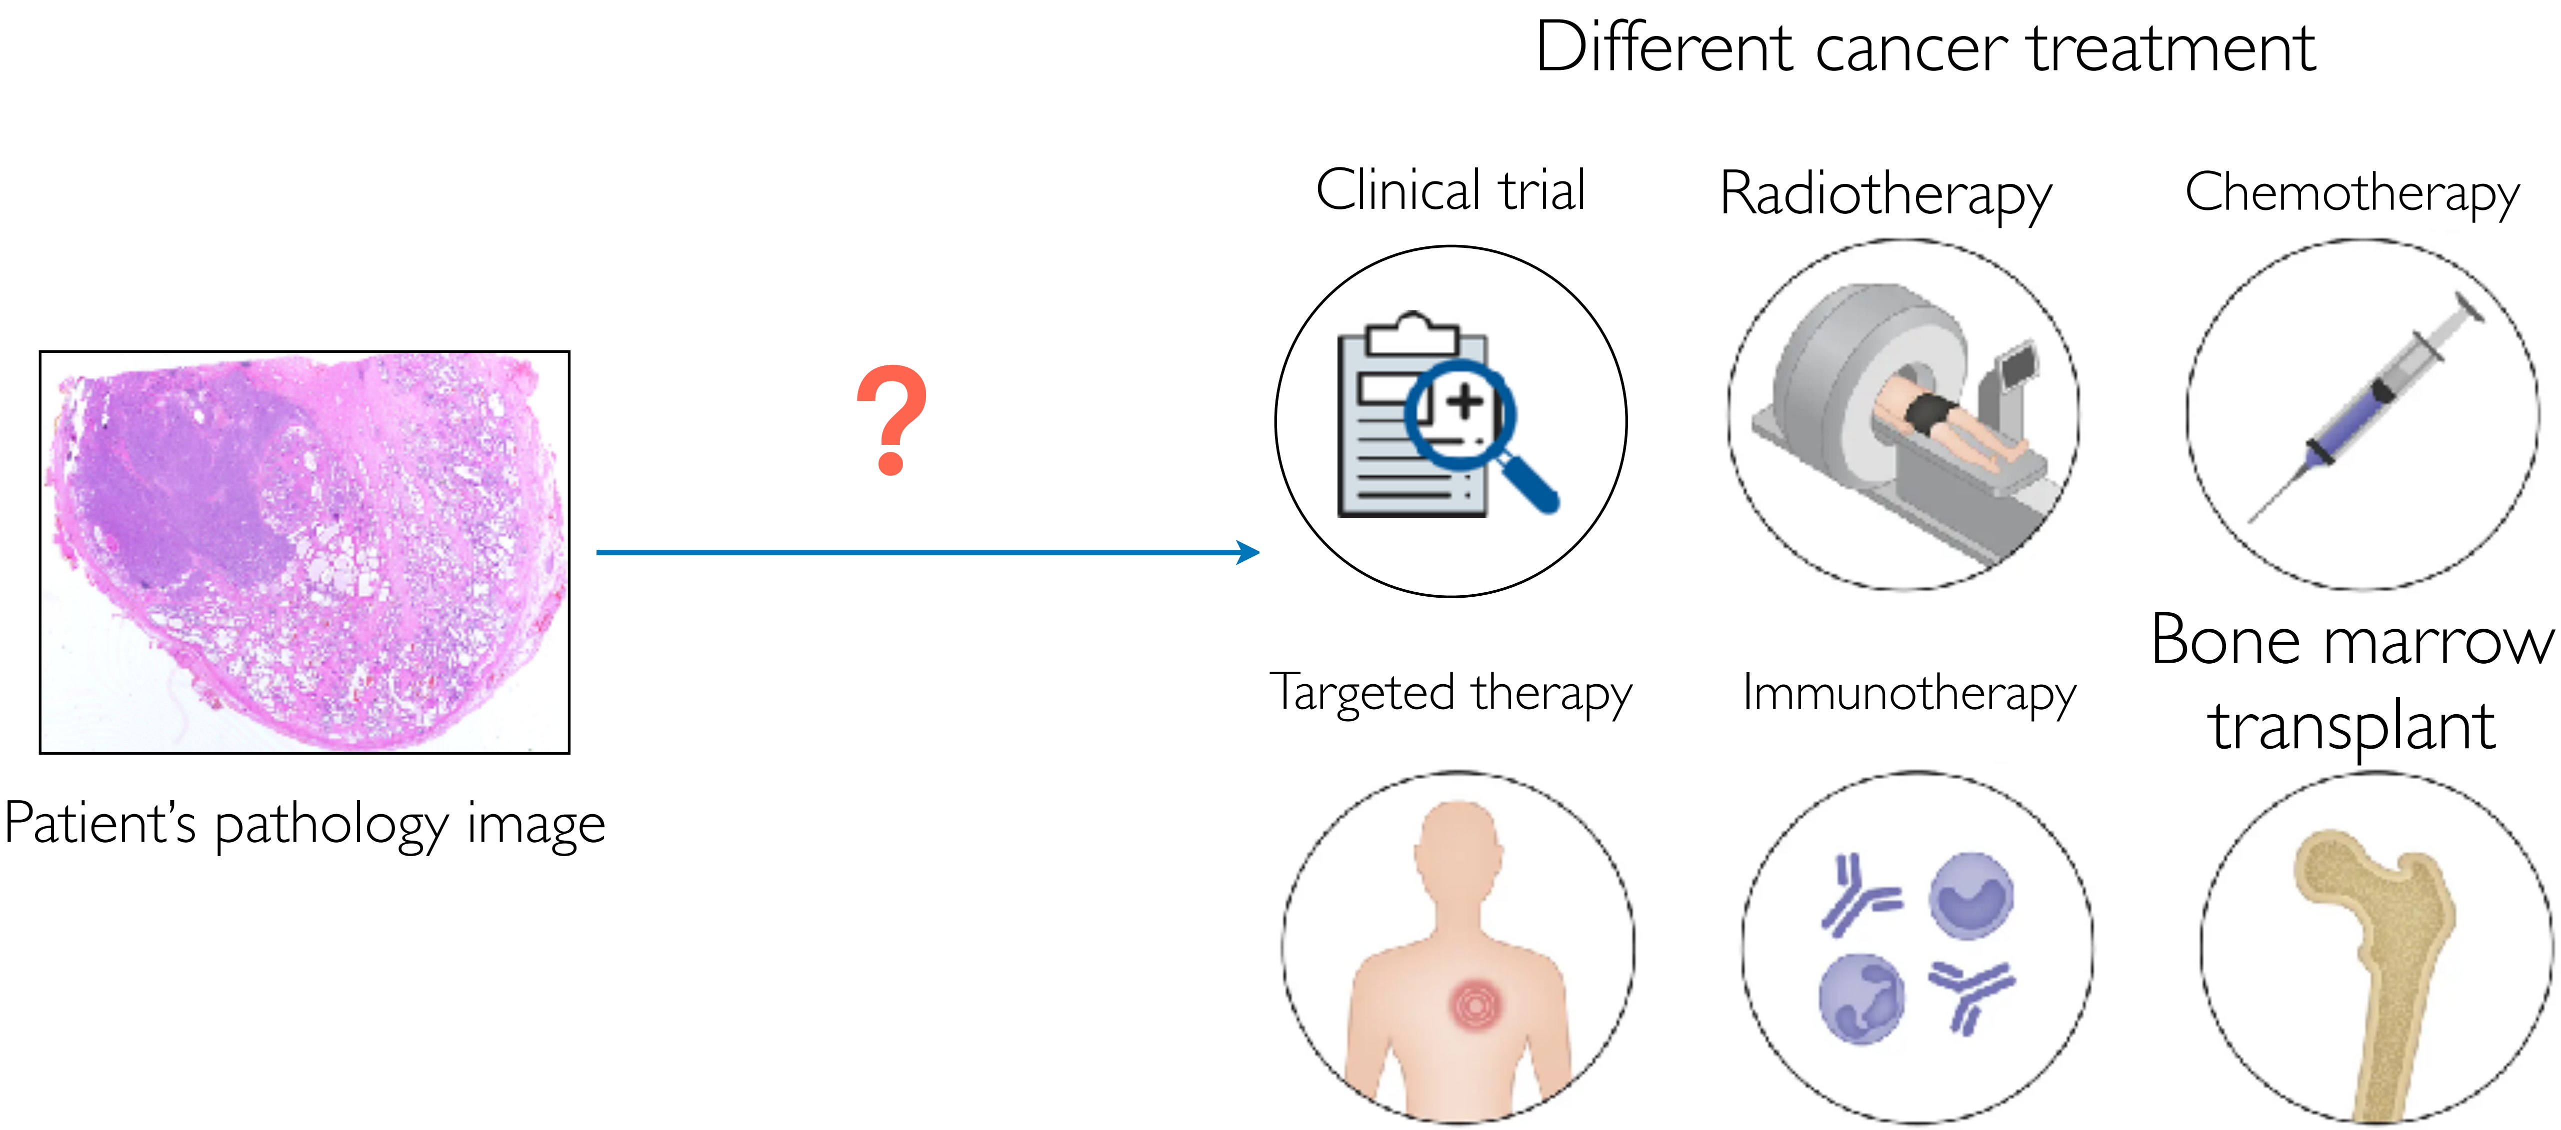

Image fromIHH Healthcare Singapore

Chain-of-Thought: Decompose a complicated task into many subtasks

Standard GenAI | Chain-of-Thought GenAI

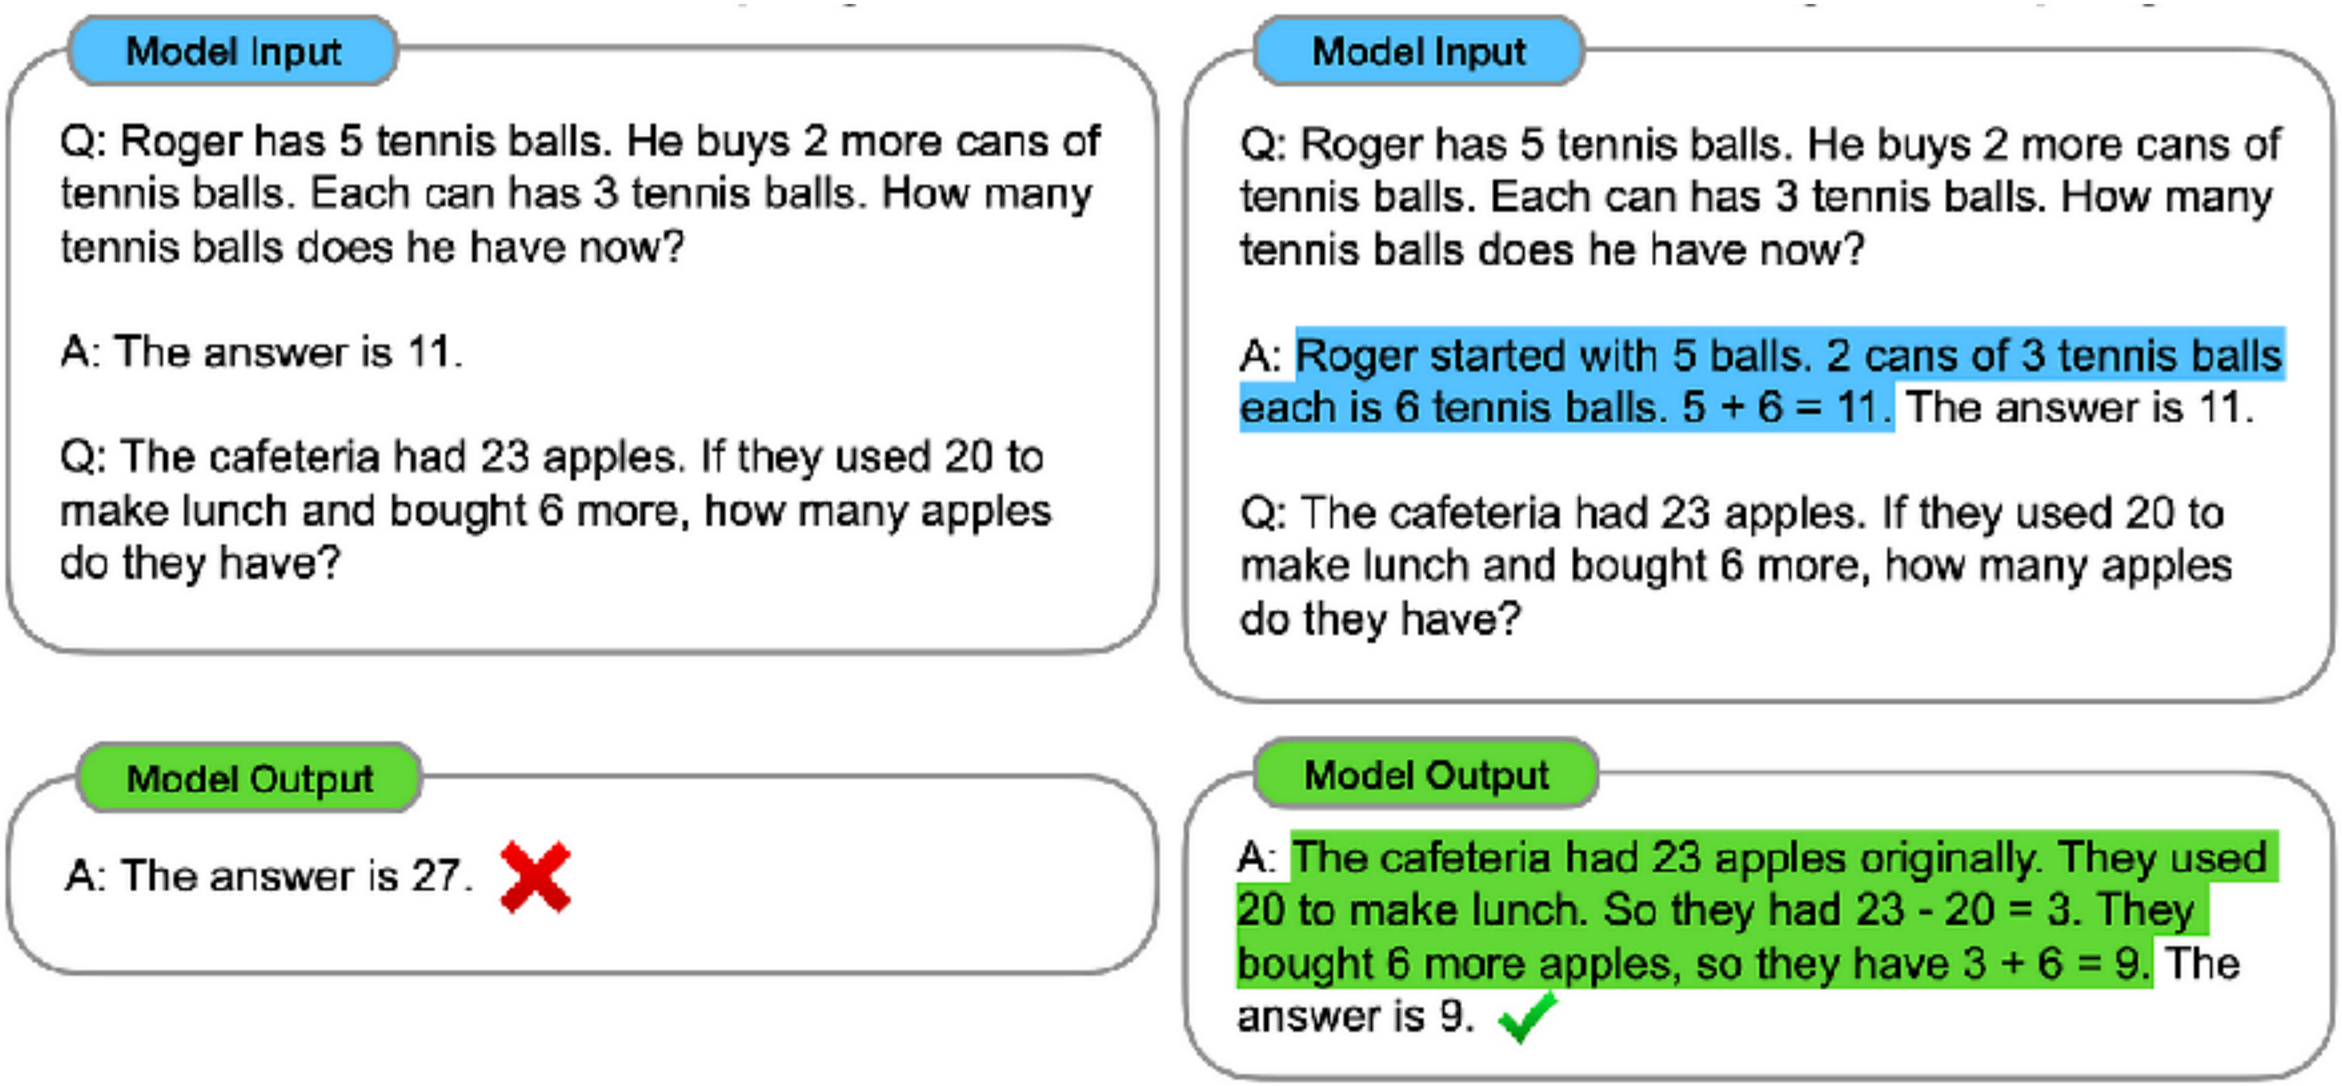

Image source: Wei, Jason, et al. NeurIPS 2022

Our idea: use Clinical Guideline as the Chain-of-Thought

NCCN National Comprehensive Cancer Network®

NCCN Guidelines Version 7.2024
Non-Small Cell Lung Cancer

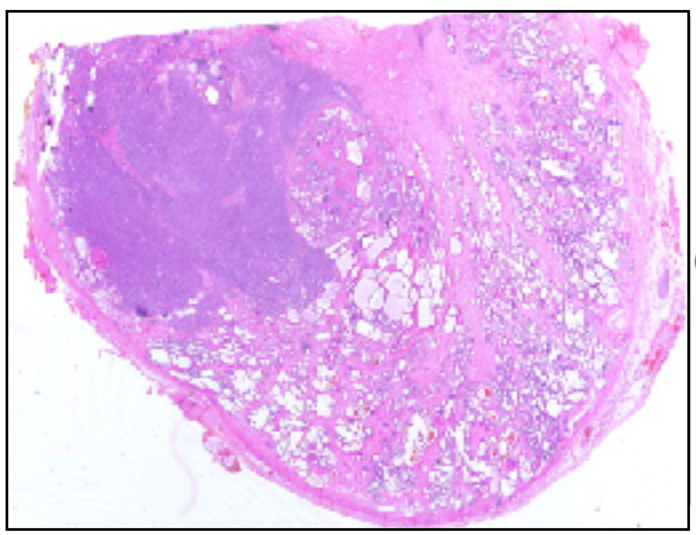

Predict

- Adenocarcinoma
- Large cell
- NSCLC not otherwise specified (NOS)

Squamous cell carcinoma

Predict biomarkers

- Molecular testing, including:
  - *EGFR* mutation (category 1), *ALK* (category 1), *KRAS, ROS1, BRAF, NTRK1/2/3, METex14* skipping, *RET, ERBB2 (HER2)*
  - Testing should be conducted as part of broad molecular profiling[oo]
- Programmed death ligand 1 (PD-L1) testing (category 1)

- Consider molecular testing, including:[pp]
  - *EGFR* mutation, *ALK, KRAS, ROS1, BRAF, NTRK1/2/3, METex14* skipping, *RET, ERBB2 (HER2)*
  - Testing should be conducted as part of broad molecular profiling[oo]
- PD-L1 testing (category 1)

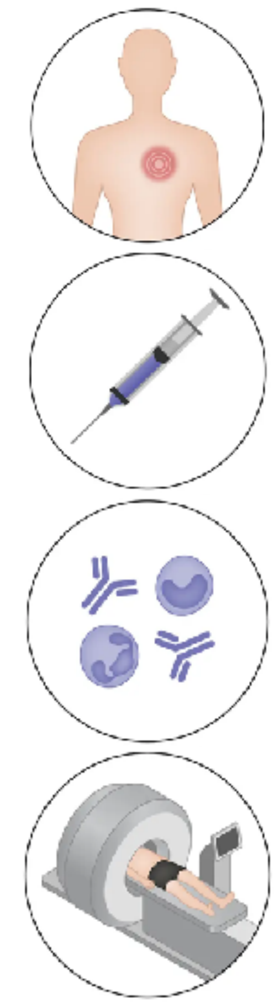

Our idea: use Clinical Guideline as the Chain-of-Thought

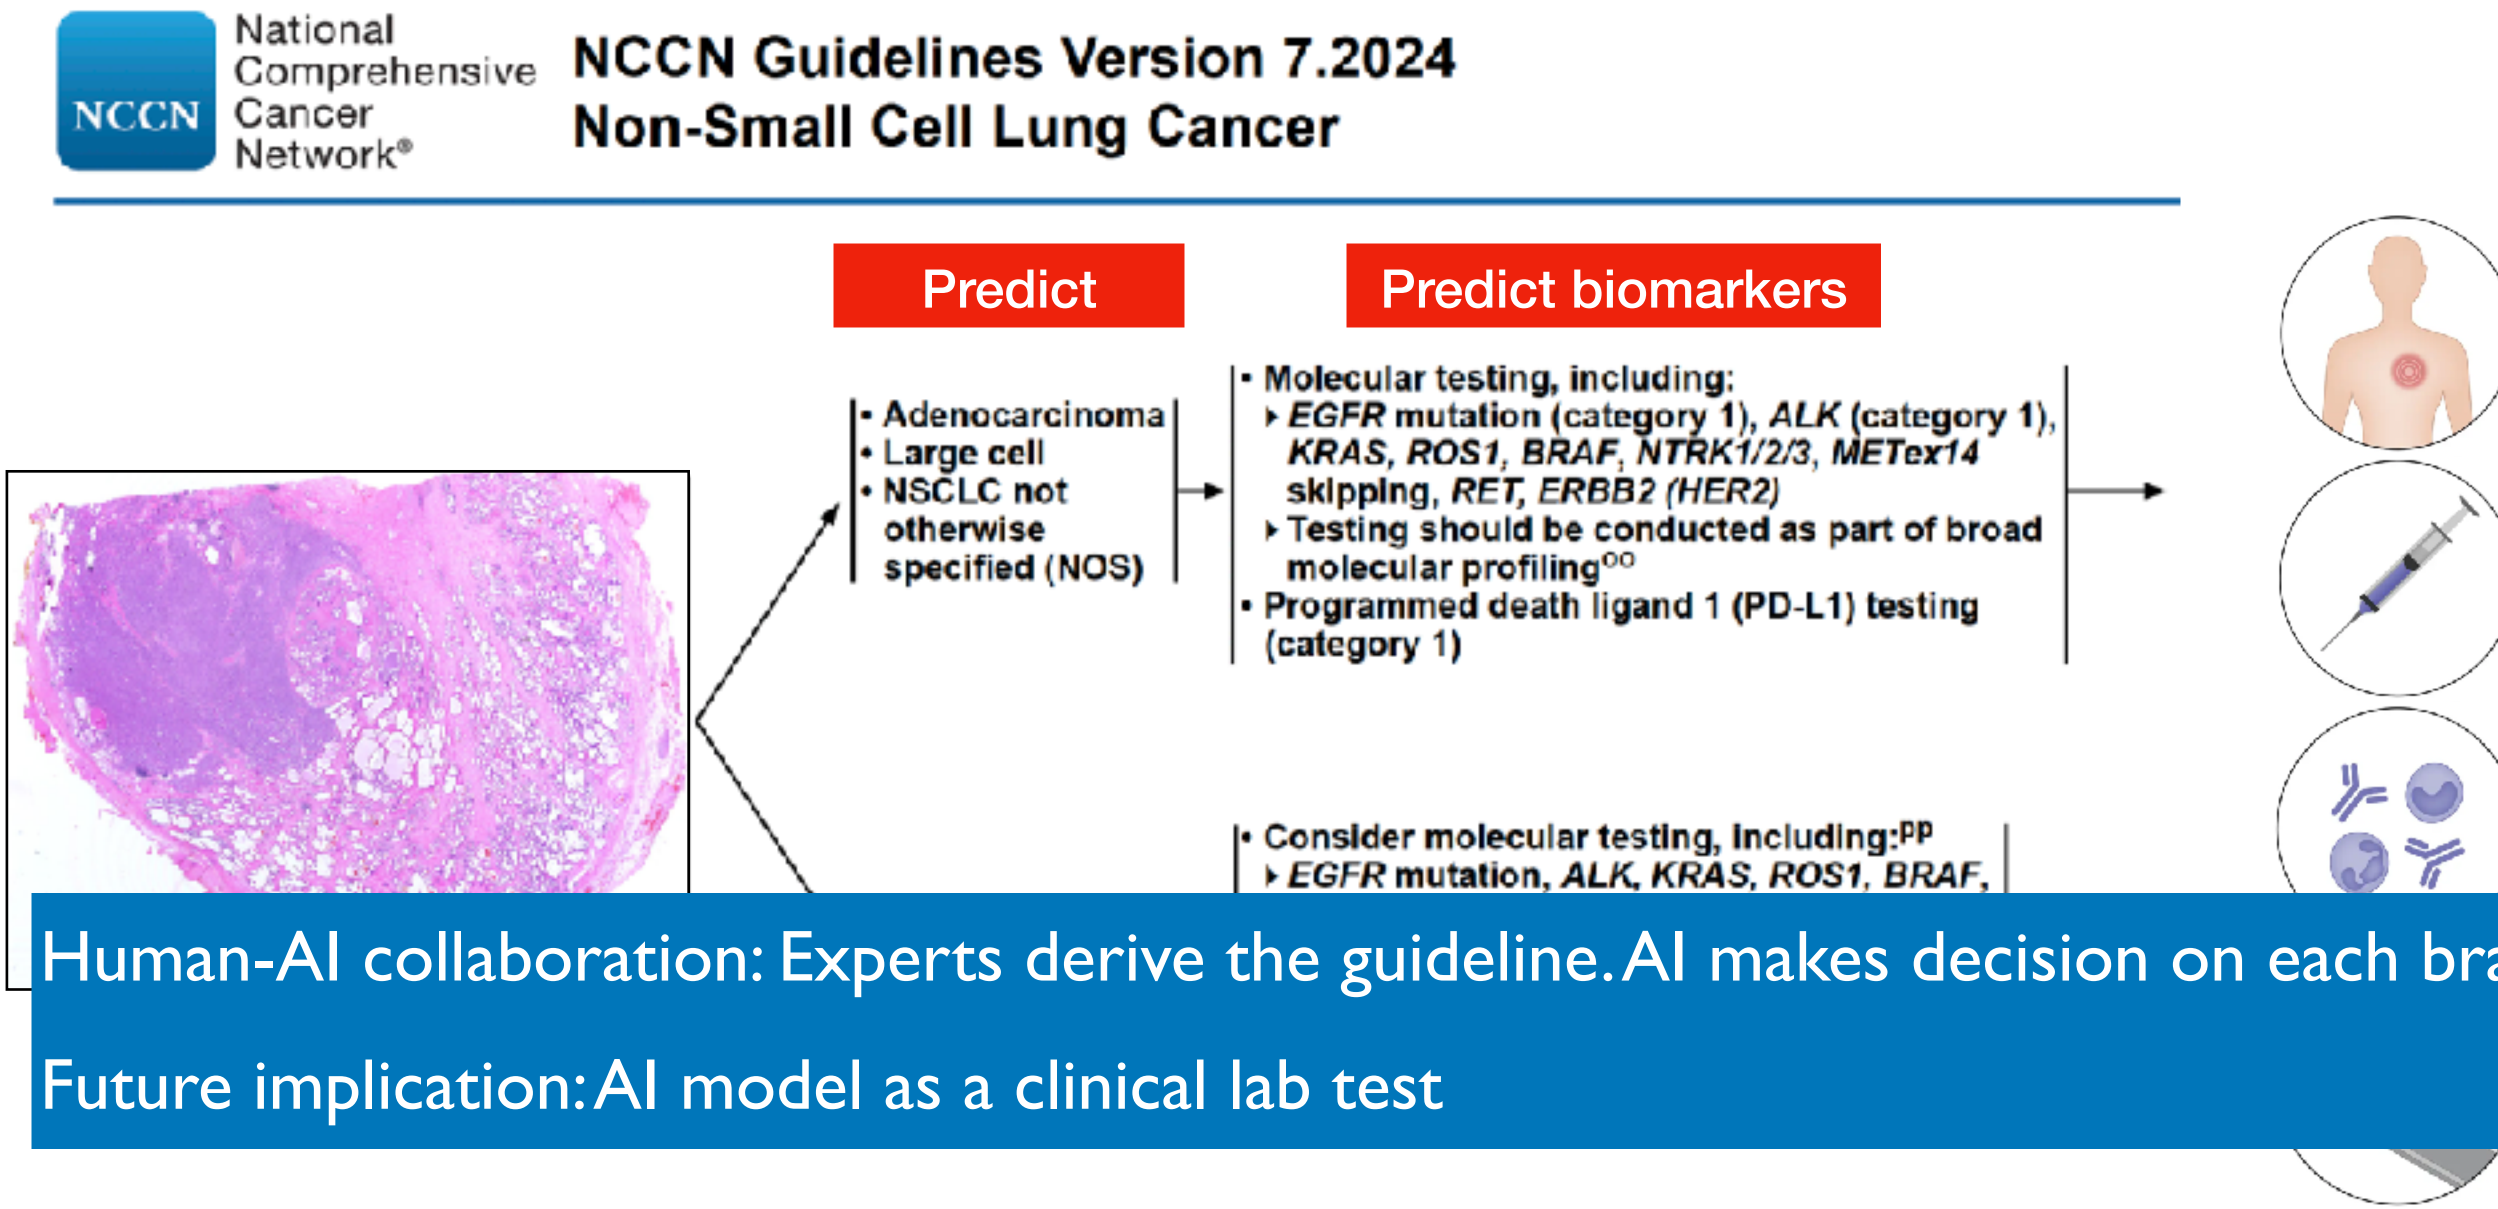

Human-AI collaboration: Experts derive the guideline. AI makes decision on each branch.

Future implication: AI model as a clinical lab test

Today's talk: 3 parts

1. Pathology foundation model
2. 3D retinal foundation model

3. A multi-modal foundation model integrating 9 imaging modalities

Foundation model (2022)
One model for all tasks

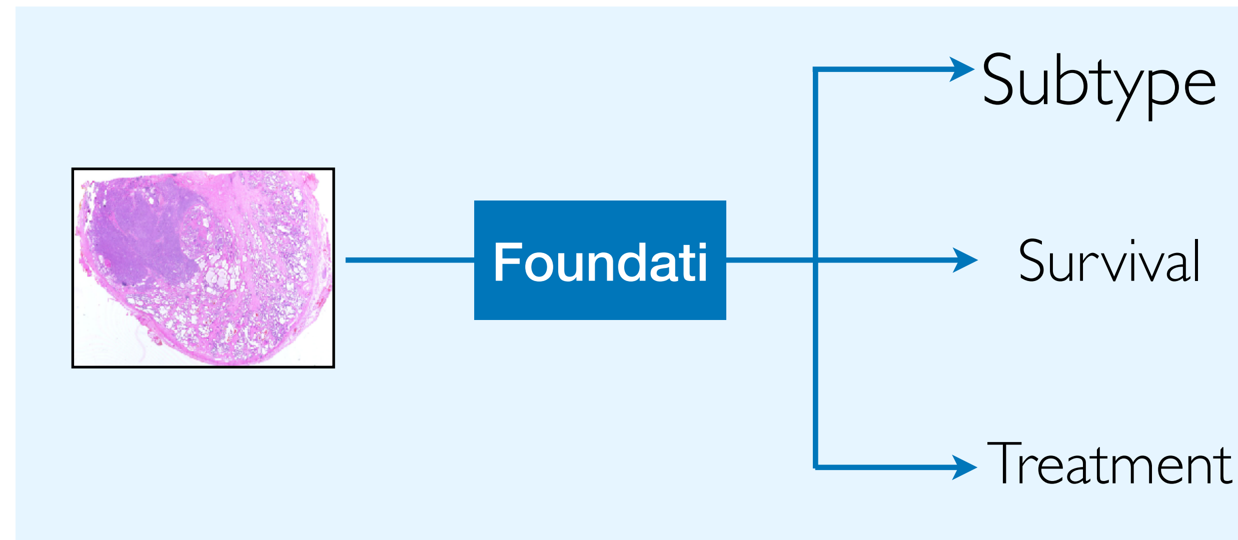

Multi-modal foundation model (2023)
One model takes different inputs

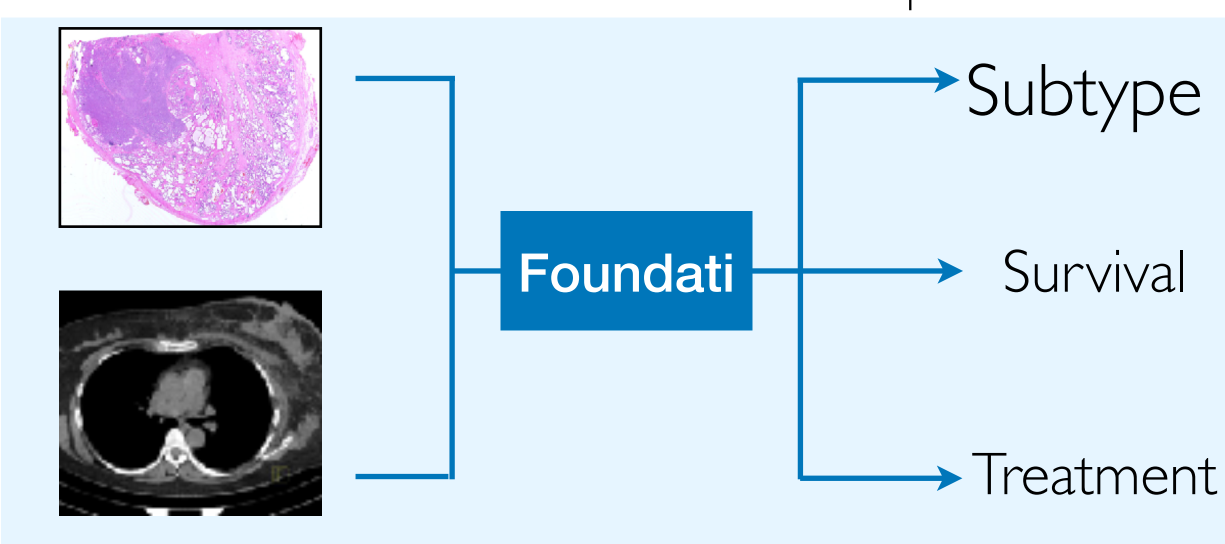

Multi-agent (2024)
Integrate multiple foundation models

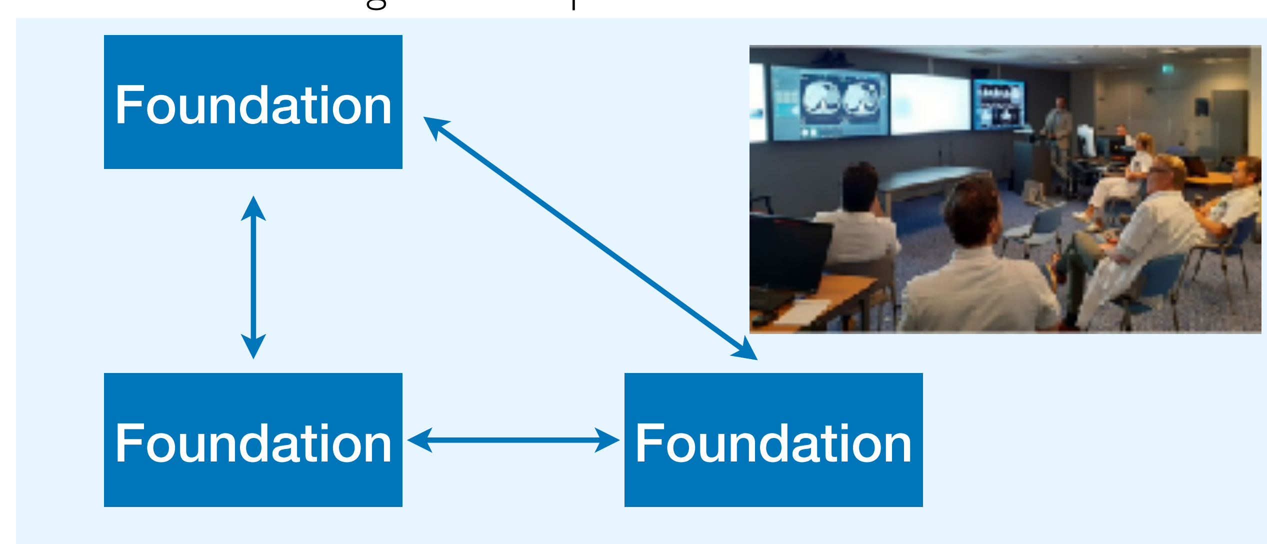

Lots of medical images are 3D

General-domain images: 2D

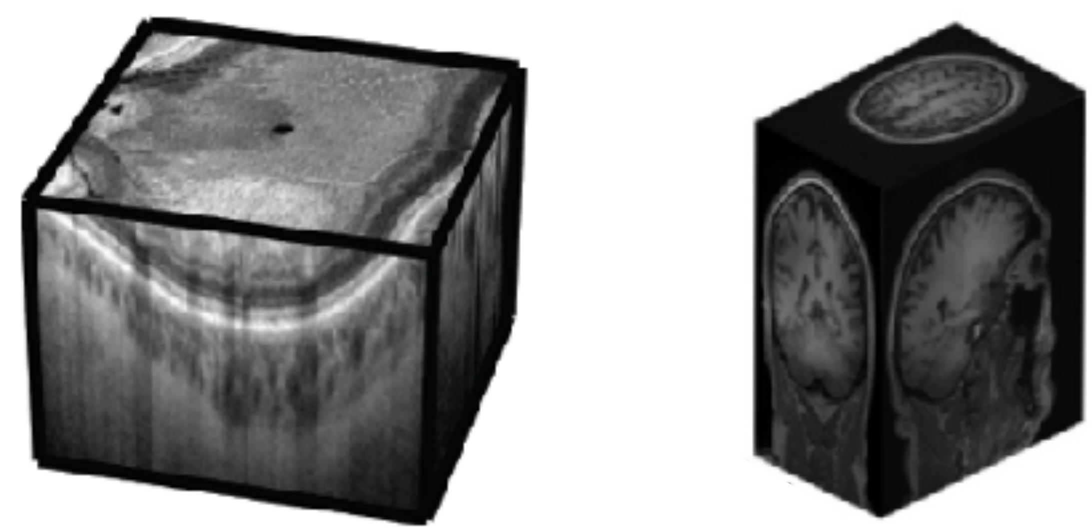

3D biomedical images: CT, MRI, OCT

Disease pattern is very small and hard to detect

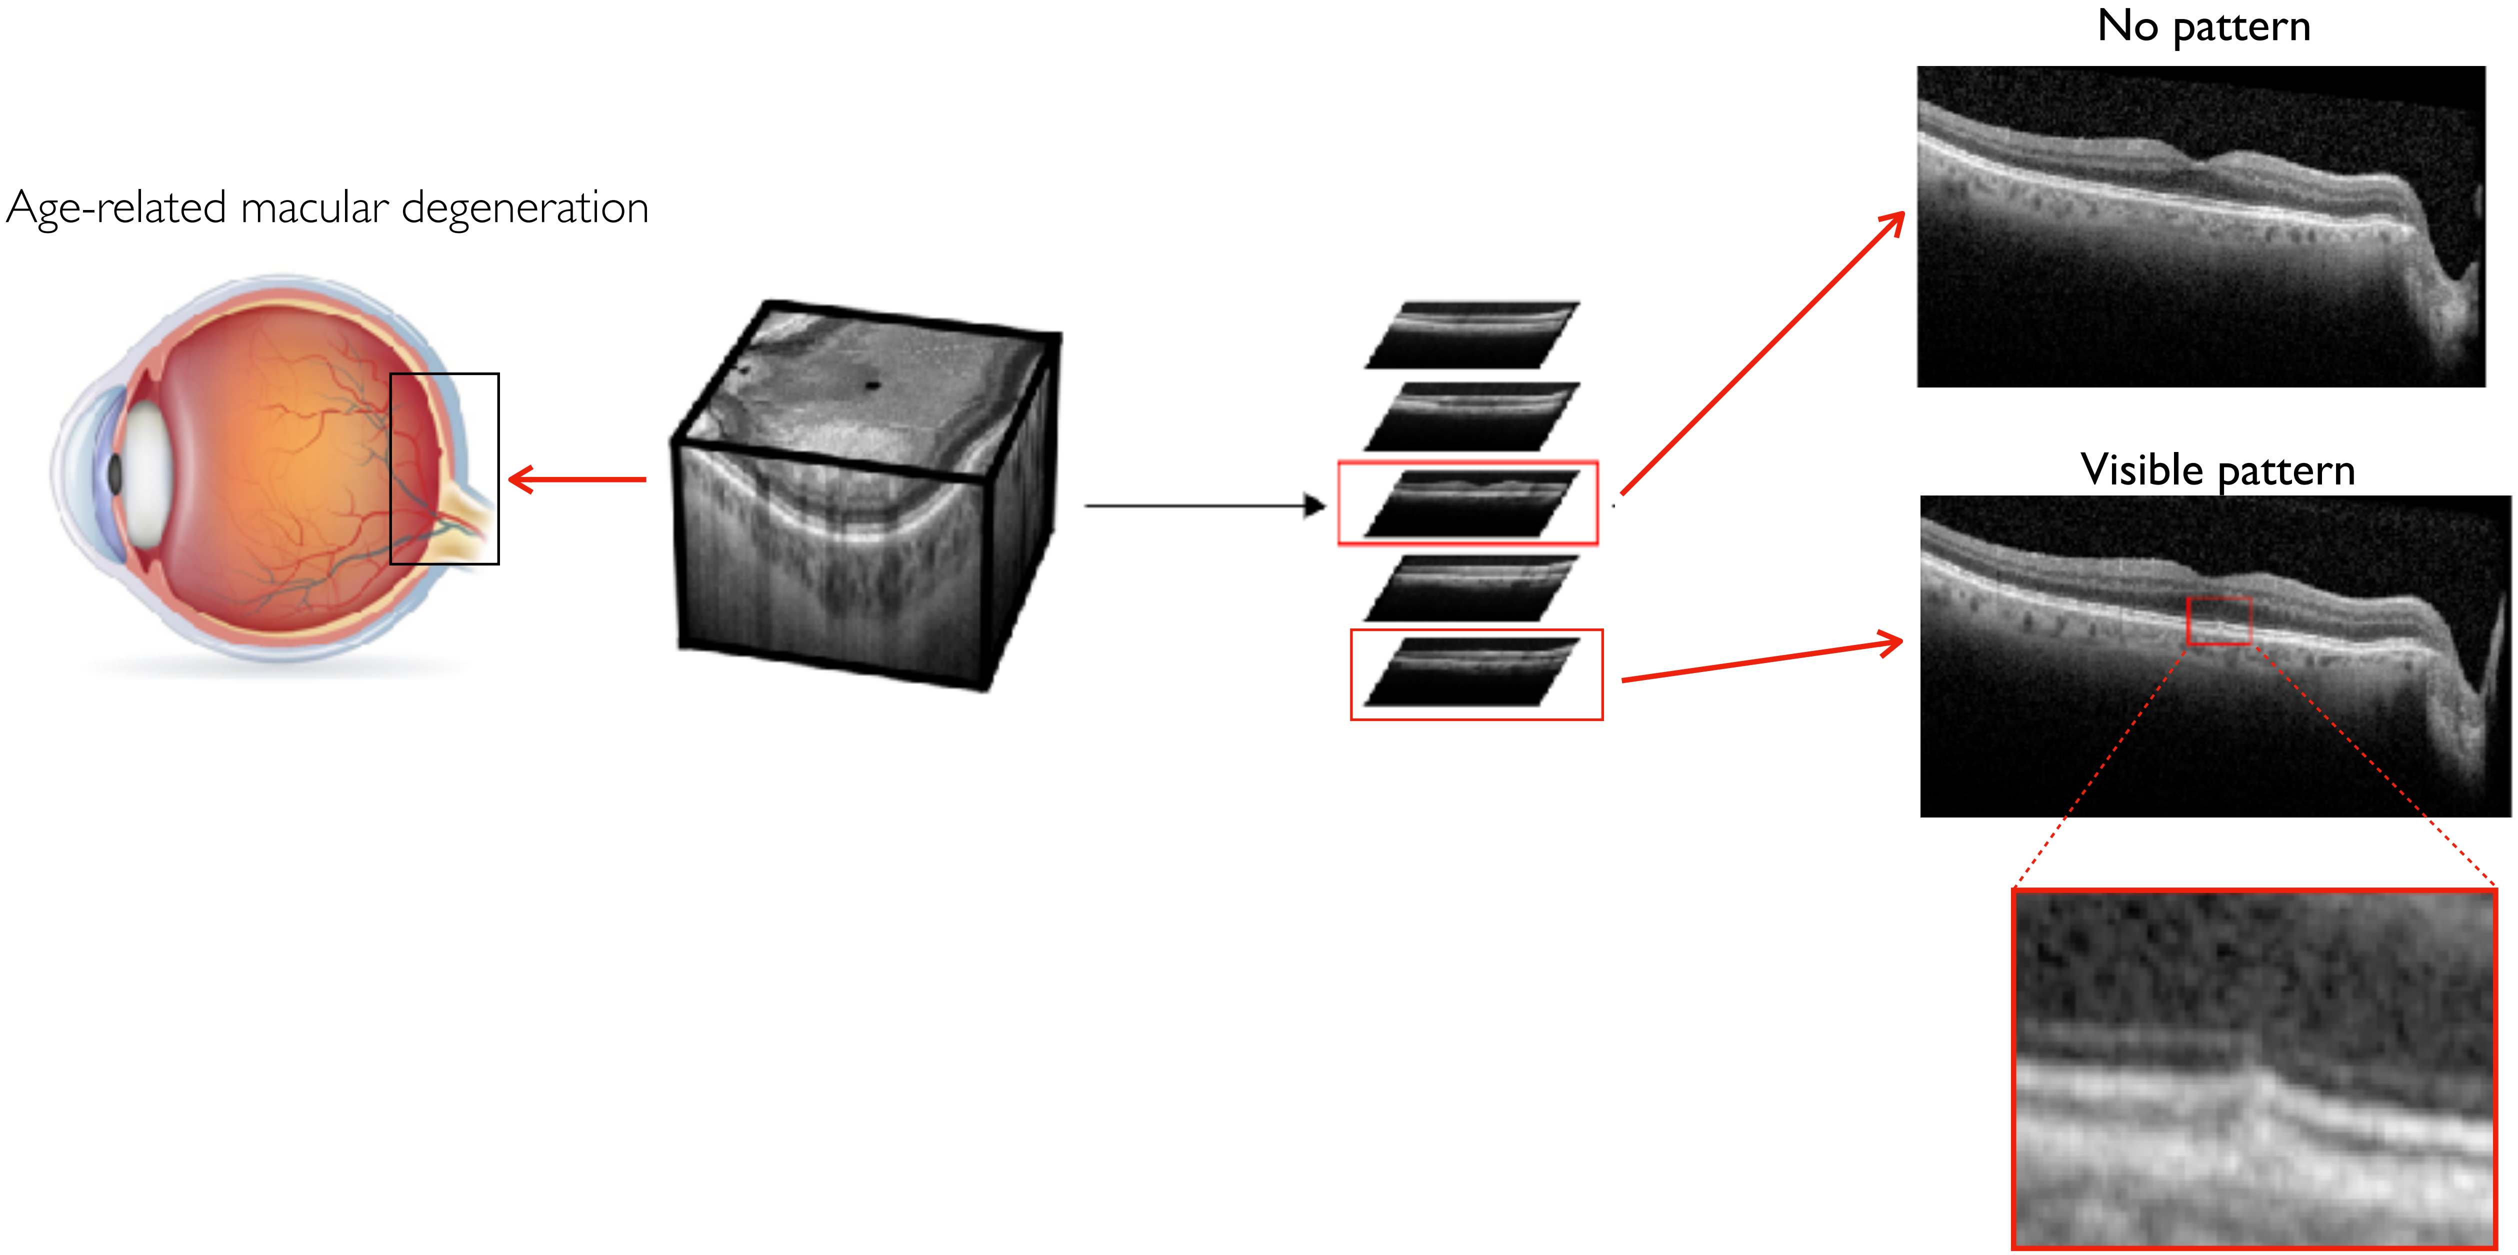

A GenAI framework for 3D medical imaging: MAE + FlashAttention

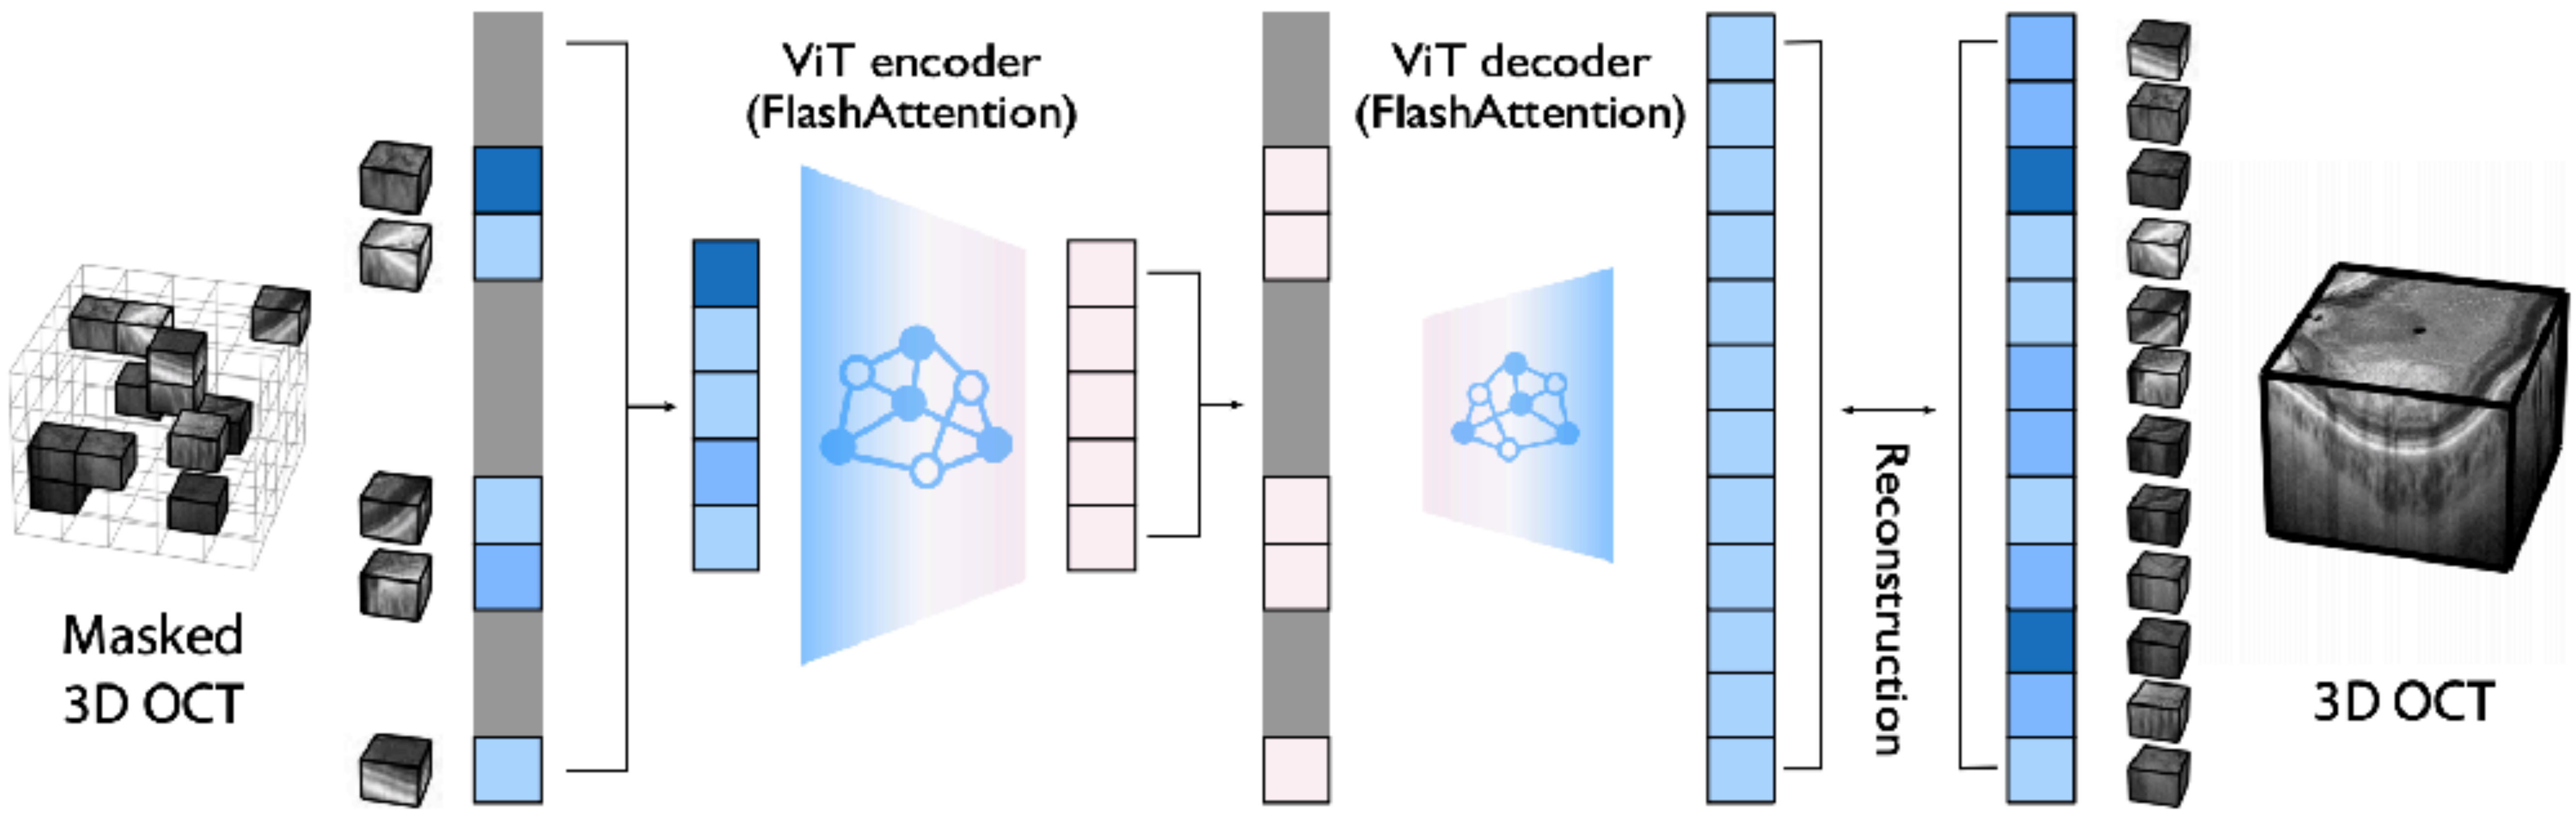

Convert a 3D volume to a **long sentence** of small cubes

The first 3D OCT foundation model trained from 26,605 patients

Best performance on all 29 tasks

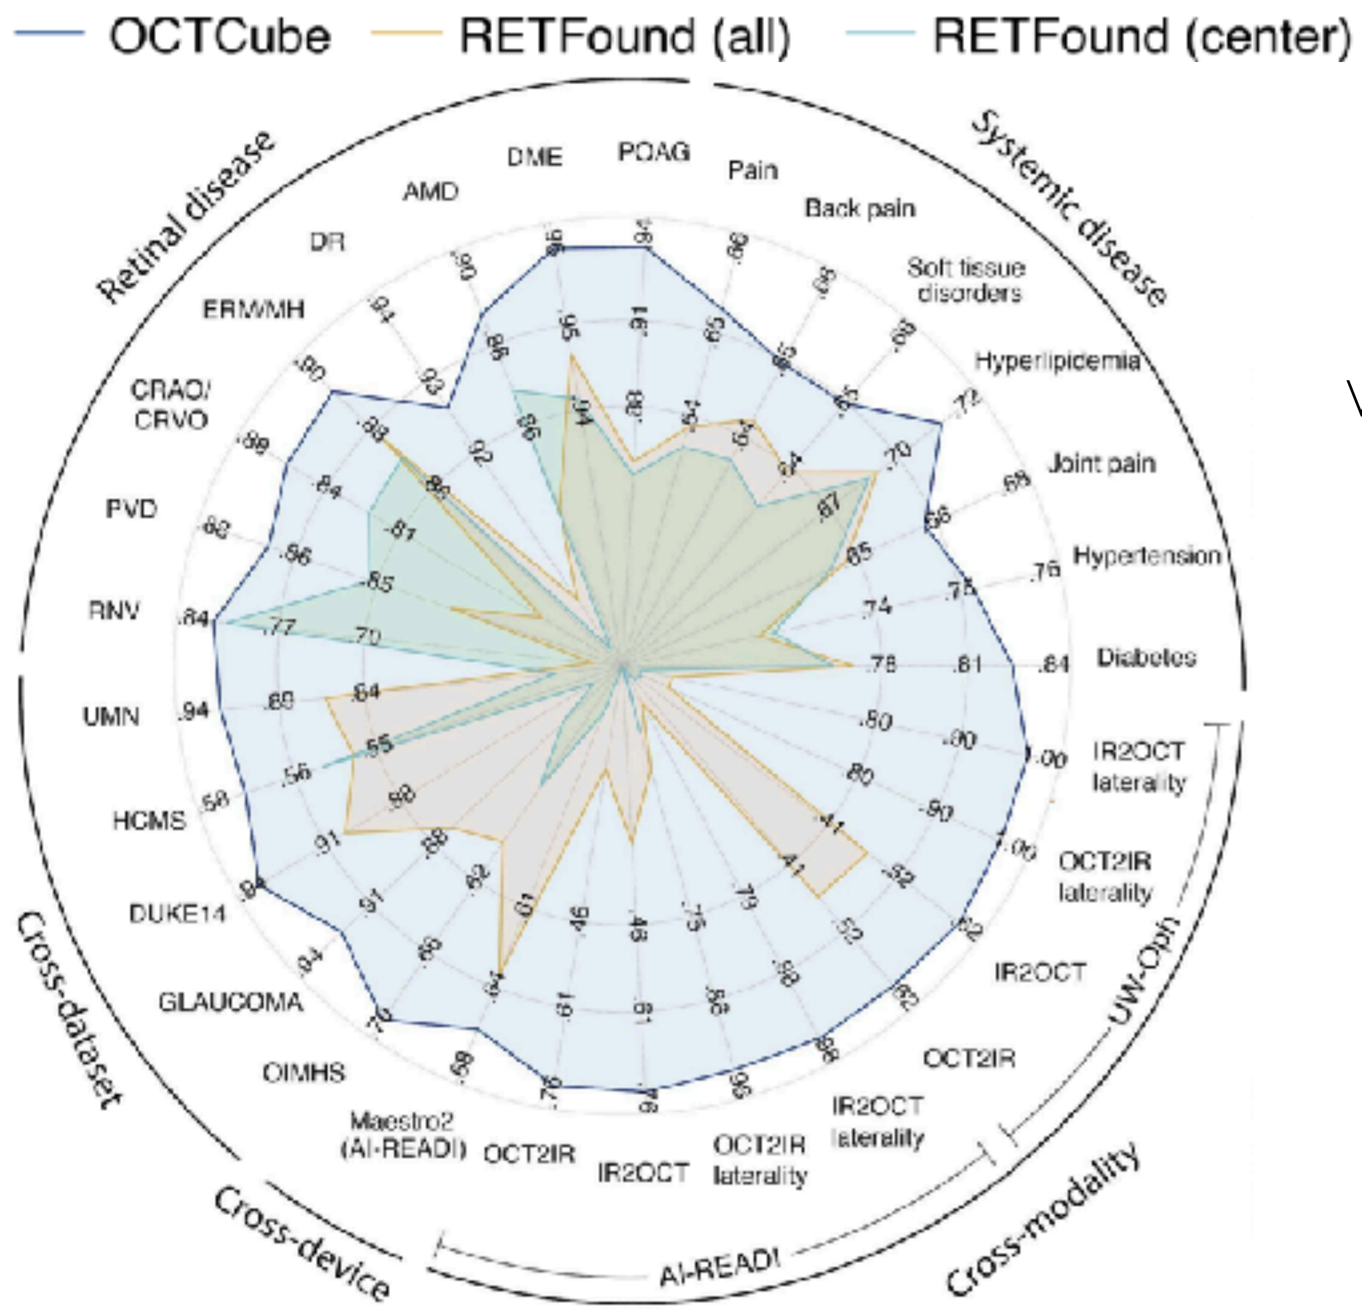

Work done by Zixuan Liu
U of Washington

Dr. Miao Zhang
Genentech

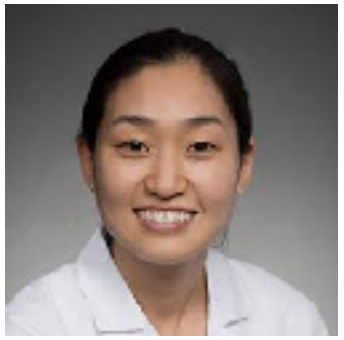

Dr. Cecilia Lee
UW ophthalmology

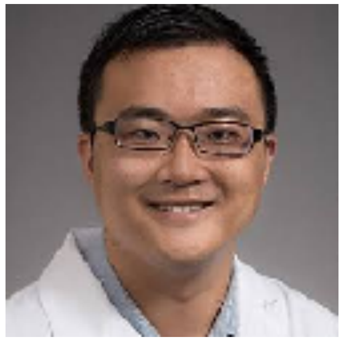

Dr. Aaron Lee
UW ophthalmology

3D model offers accurate and interpretable predictions for retinal diseases

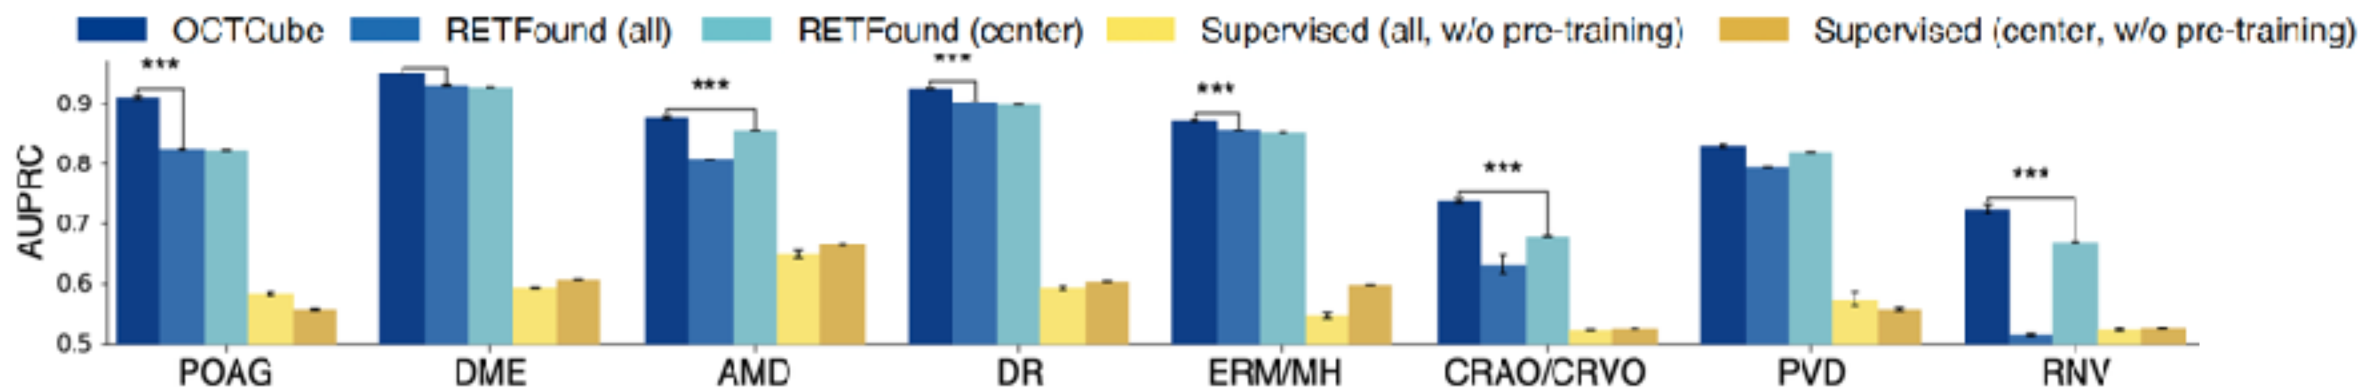

3D model offers accurate and interpretable predictions for retinal diseases

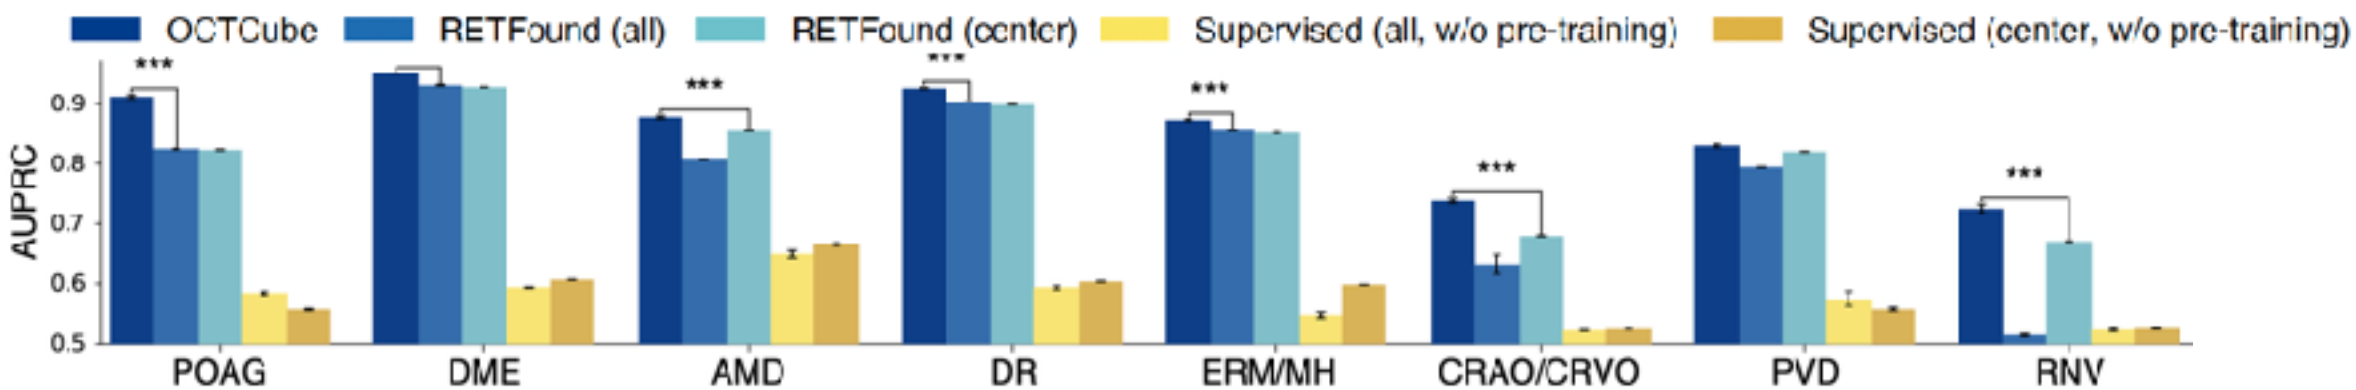

Bottom layer

Top layer

OCT slice

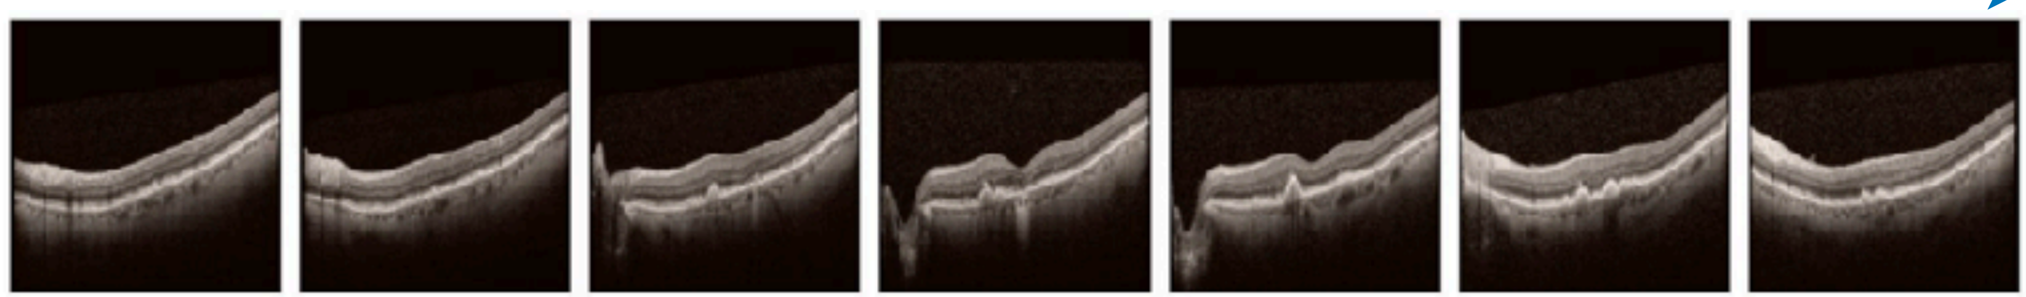

3D model offers accurate and interpretable predictions for retinal diseases

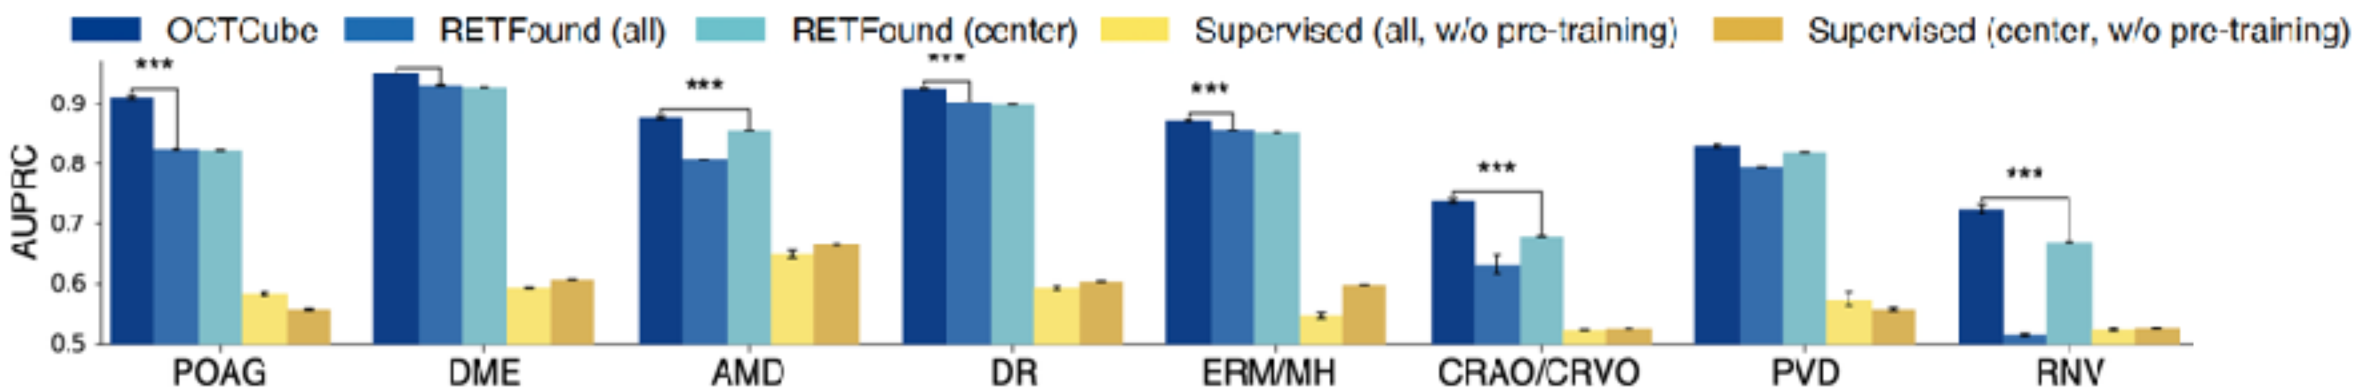

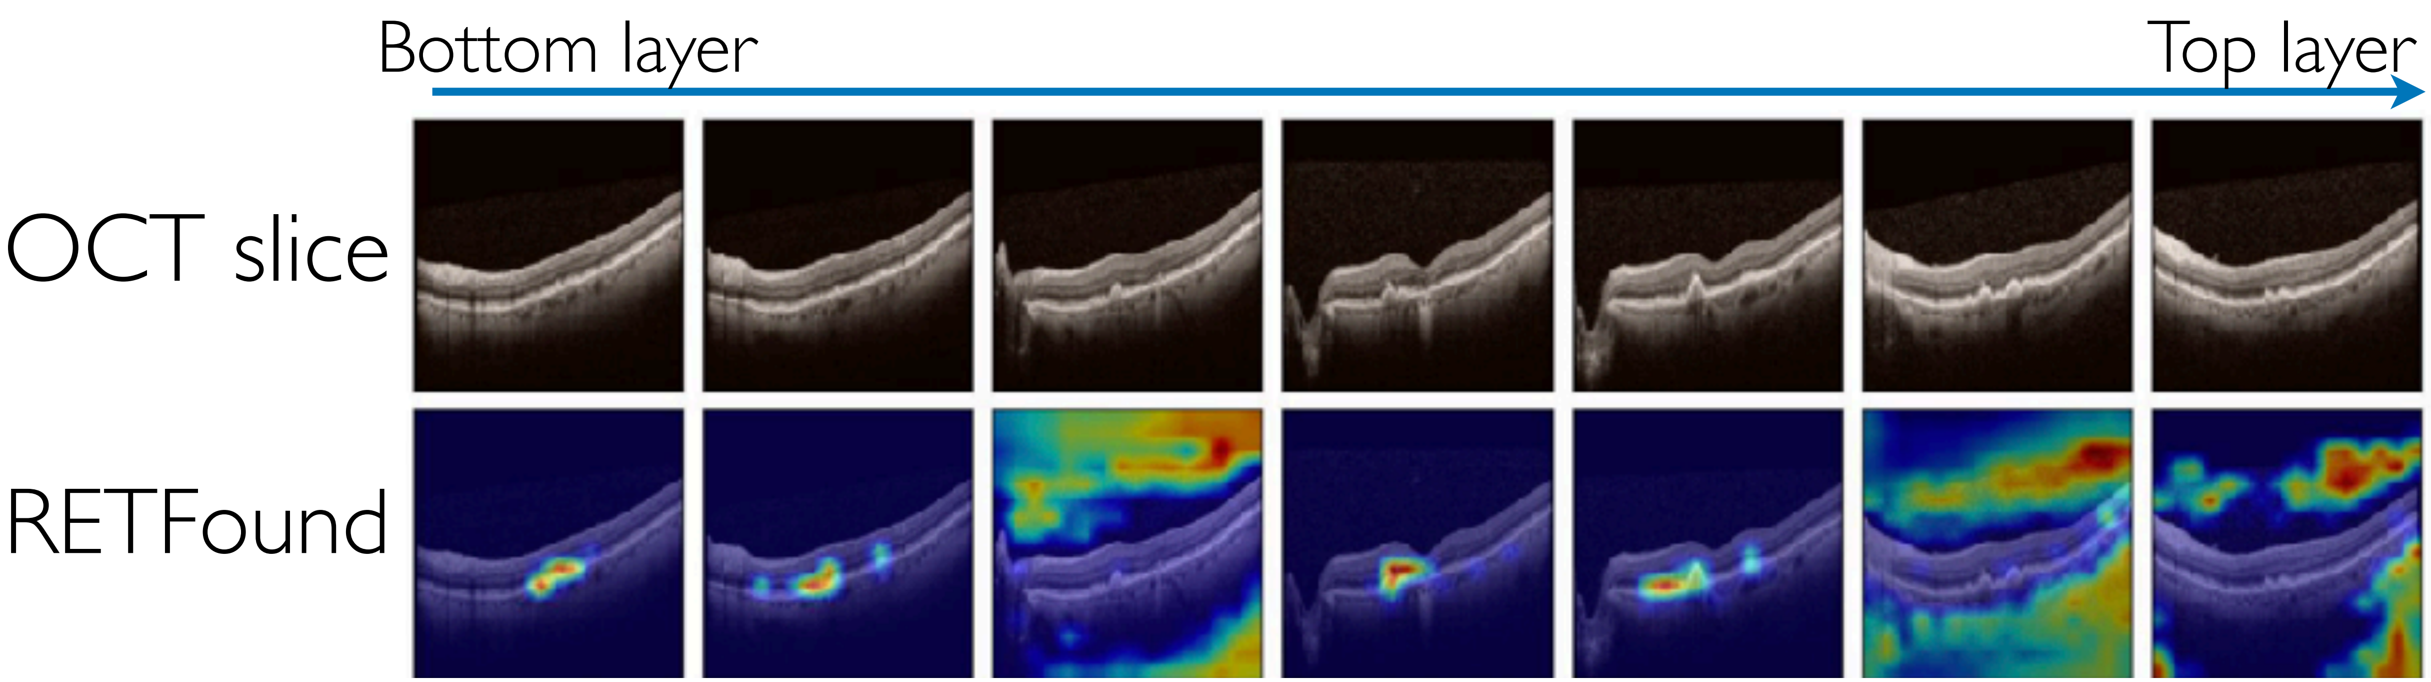

3D model offers accurate and interpretable predictions for retinal diseases

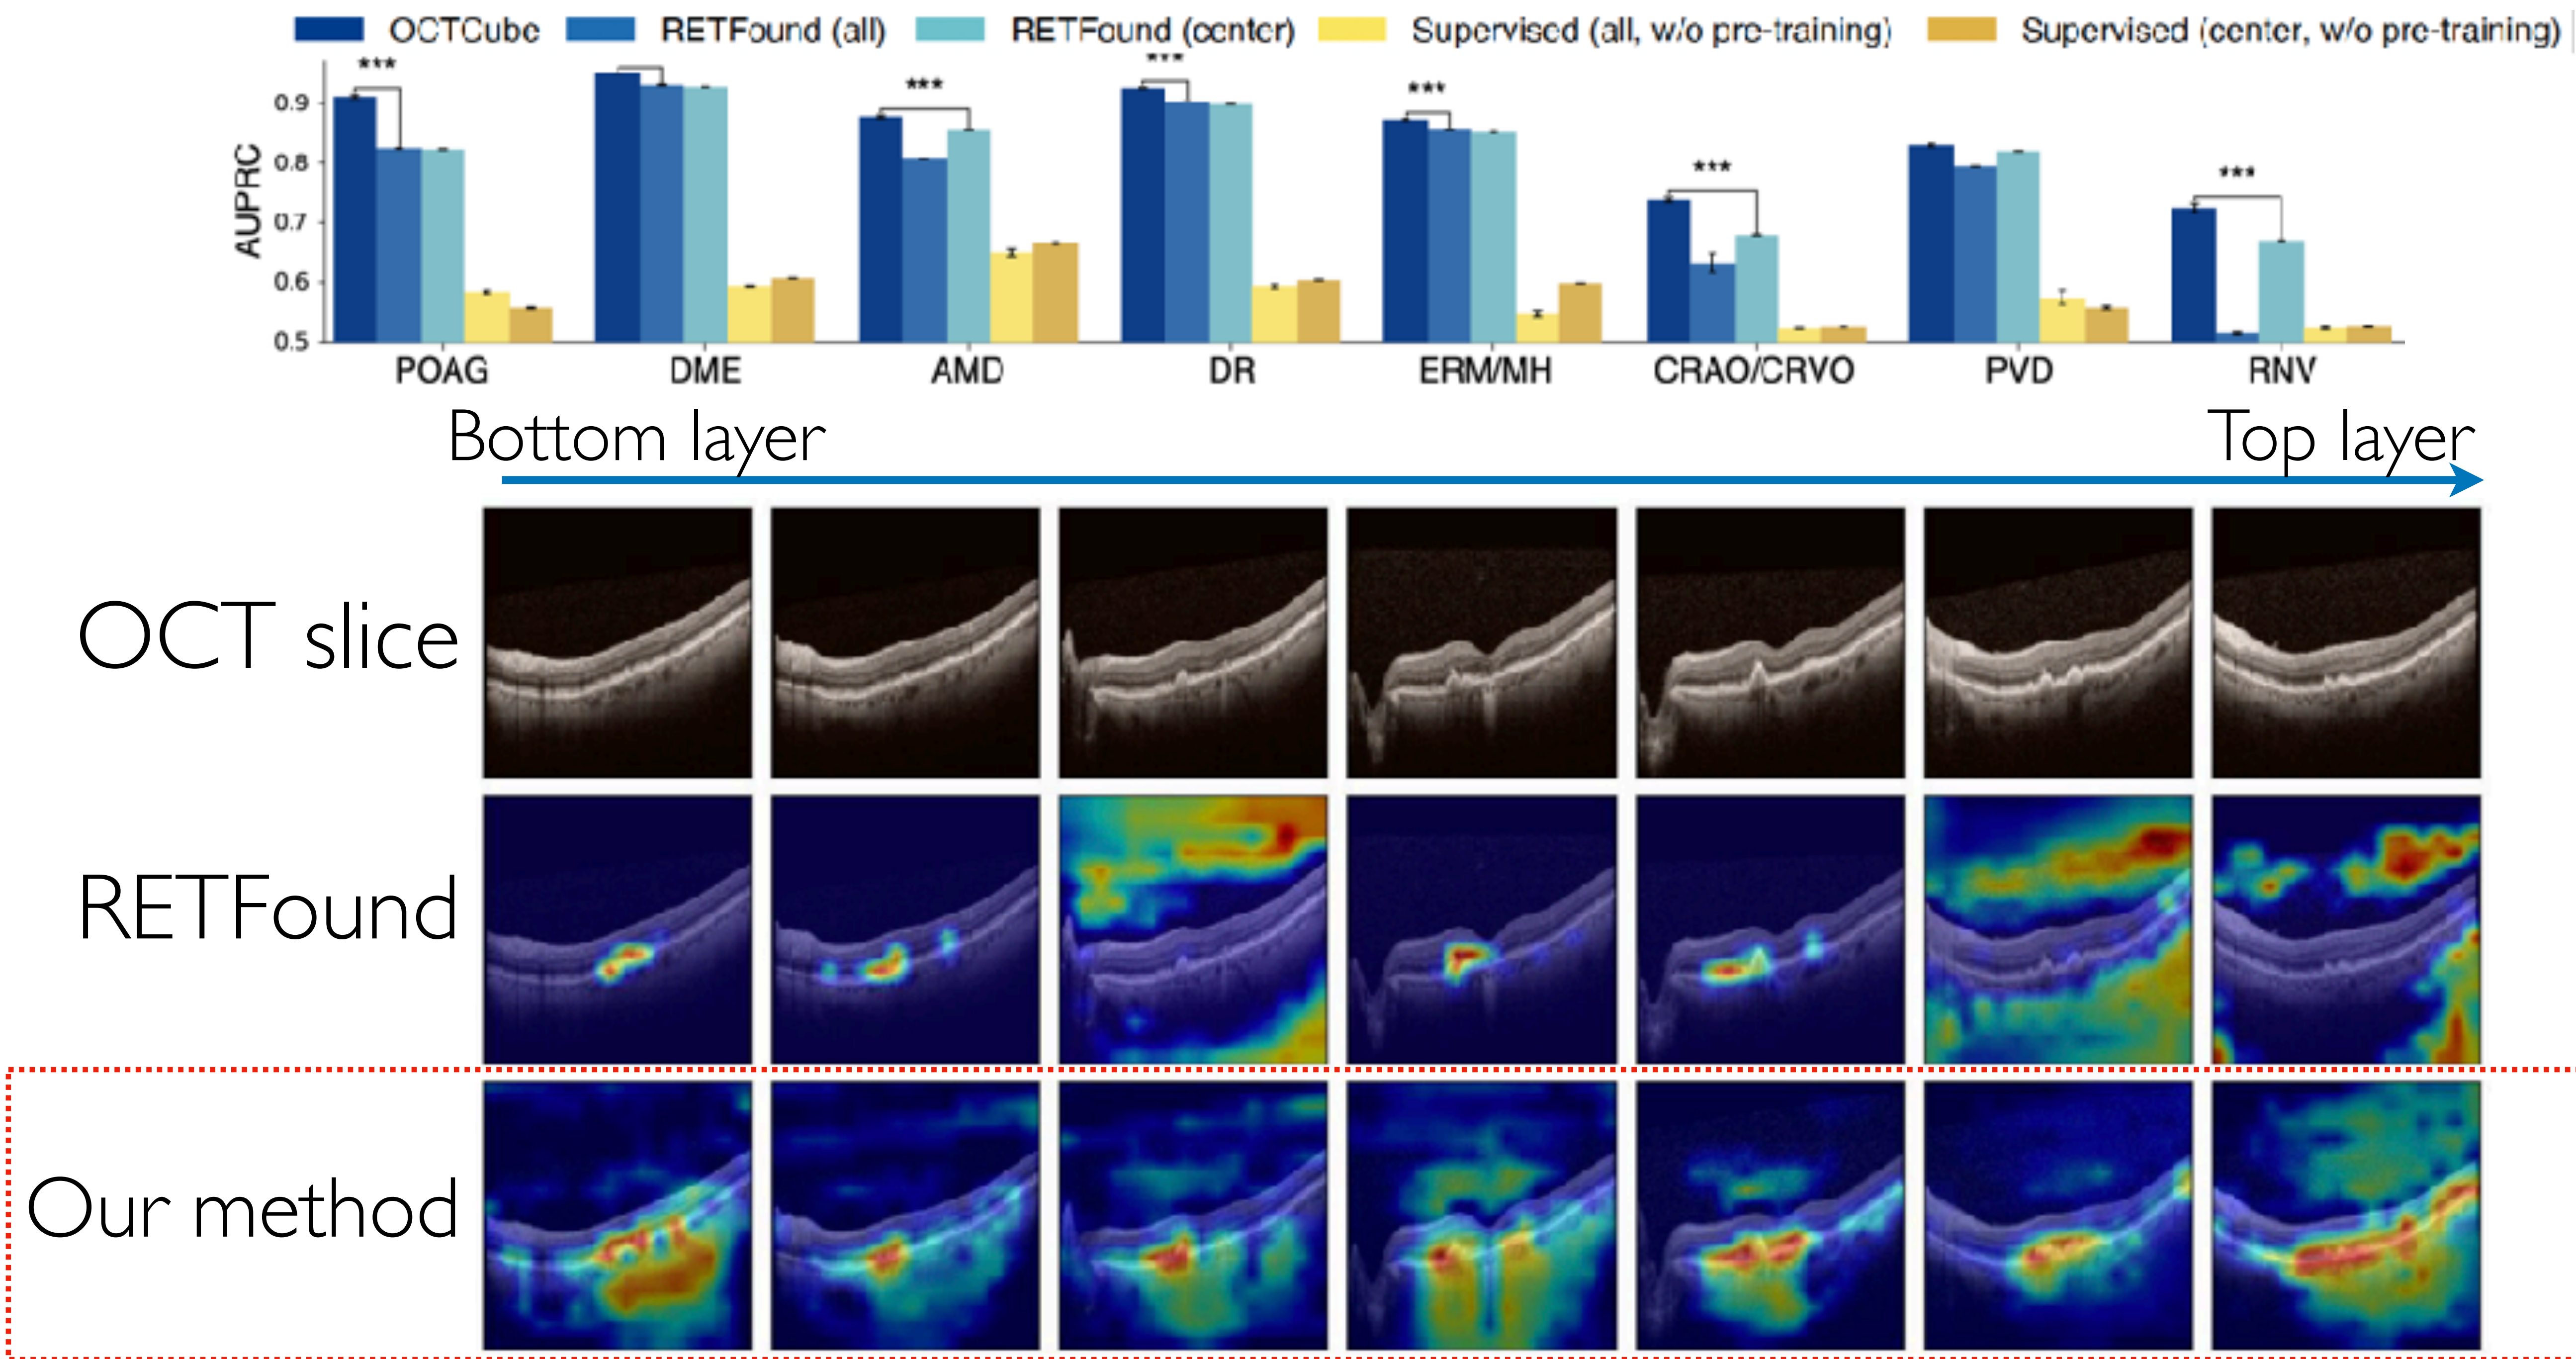

3D model is better at generalization: Cross-cohort prediction at Duke, University of Minnesota, Johns Hopkins, NYU

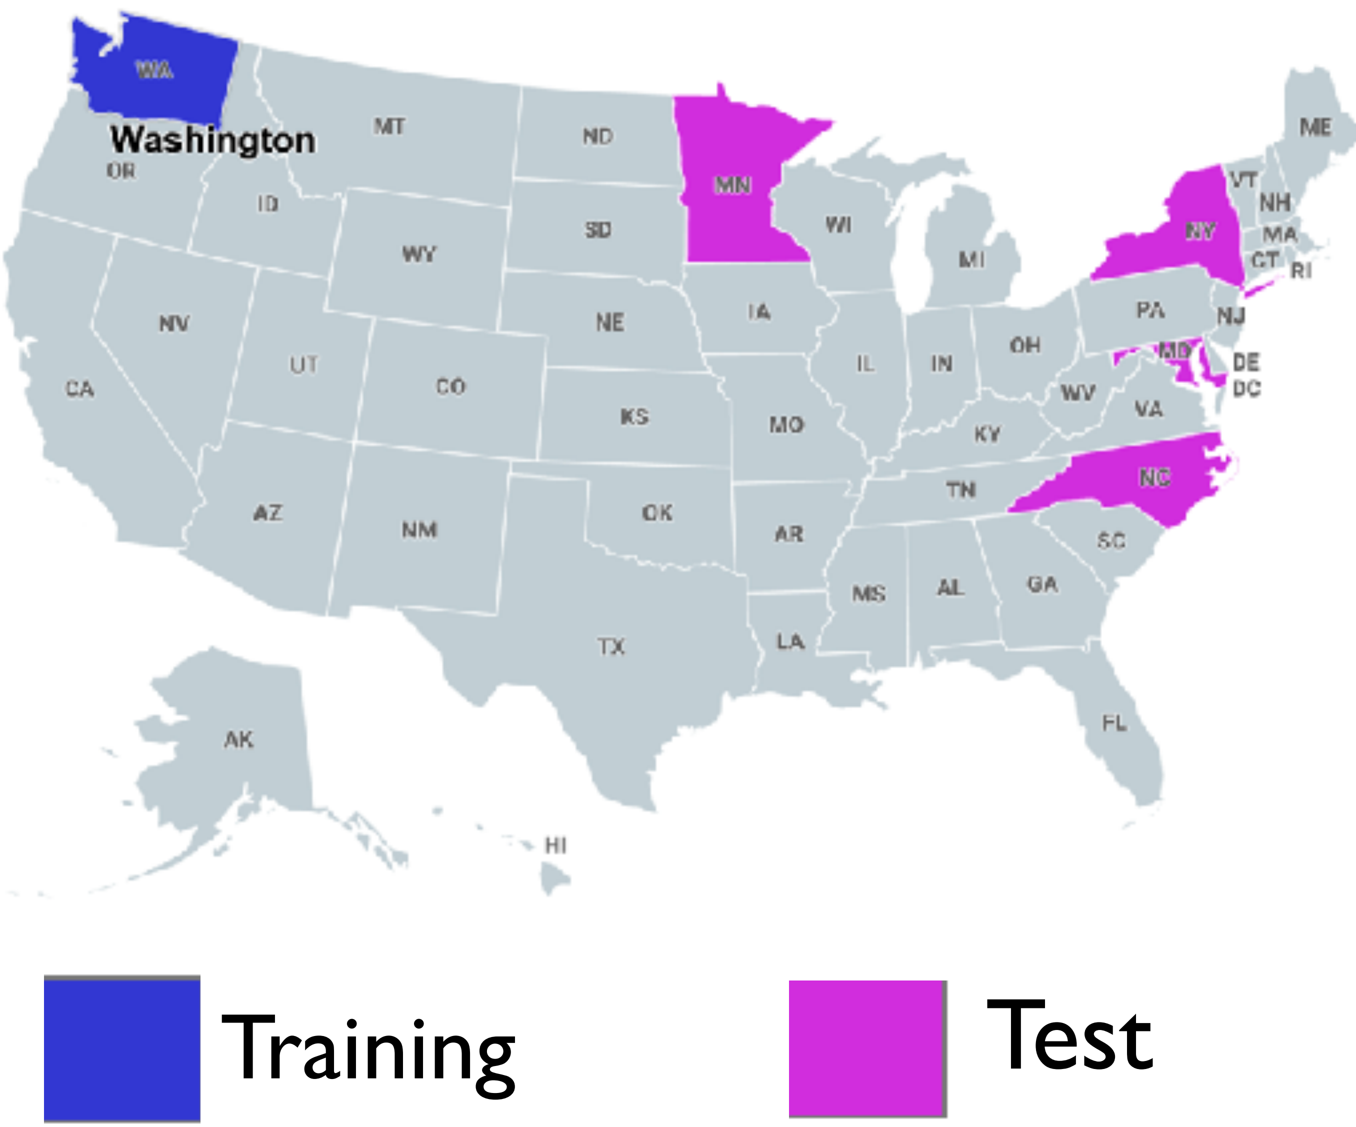

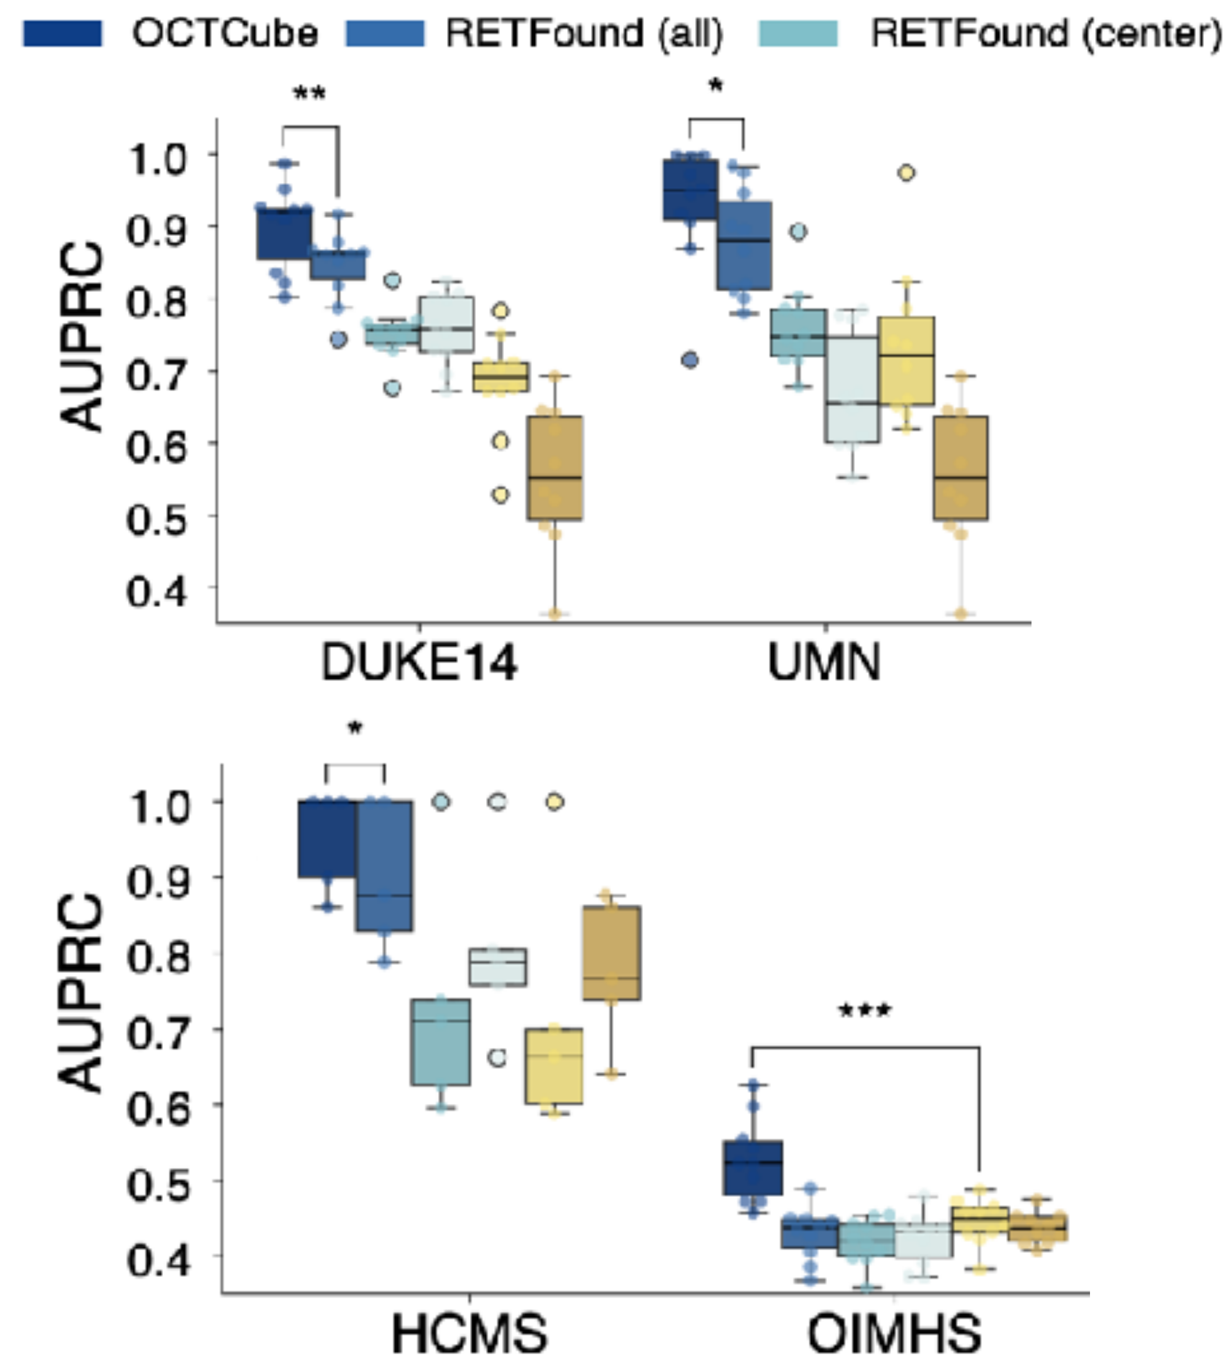

3D model enables cross-device prediction

Heidelberg Spectralis

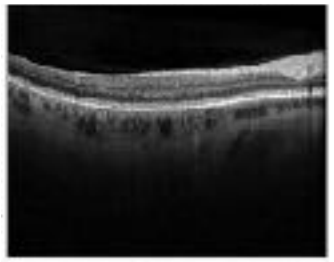

Topcon Masestro2

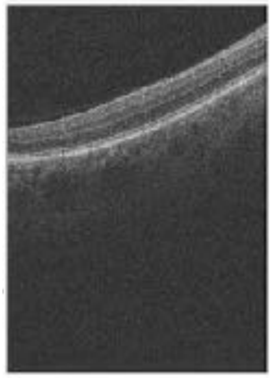

3D model enables cross-device prediction

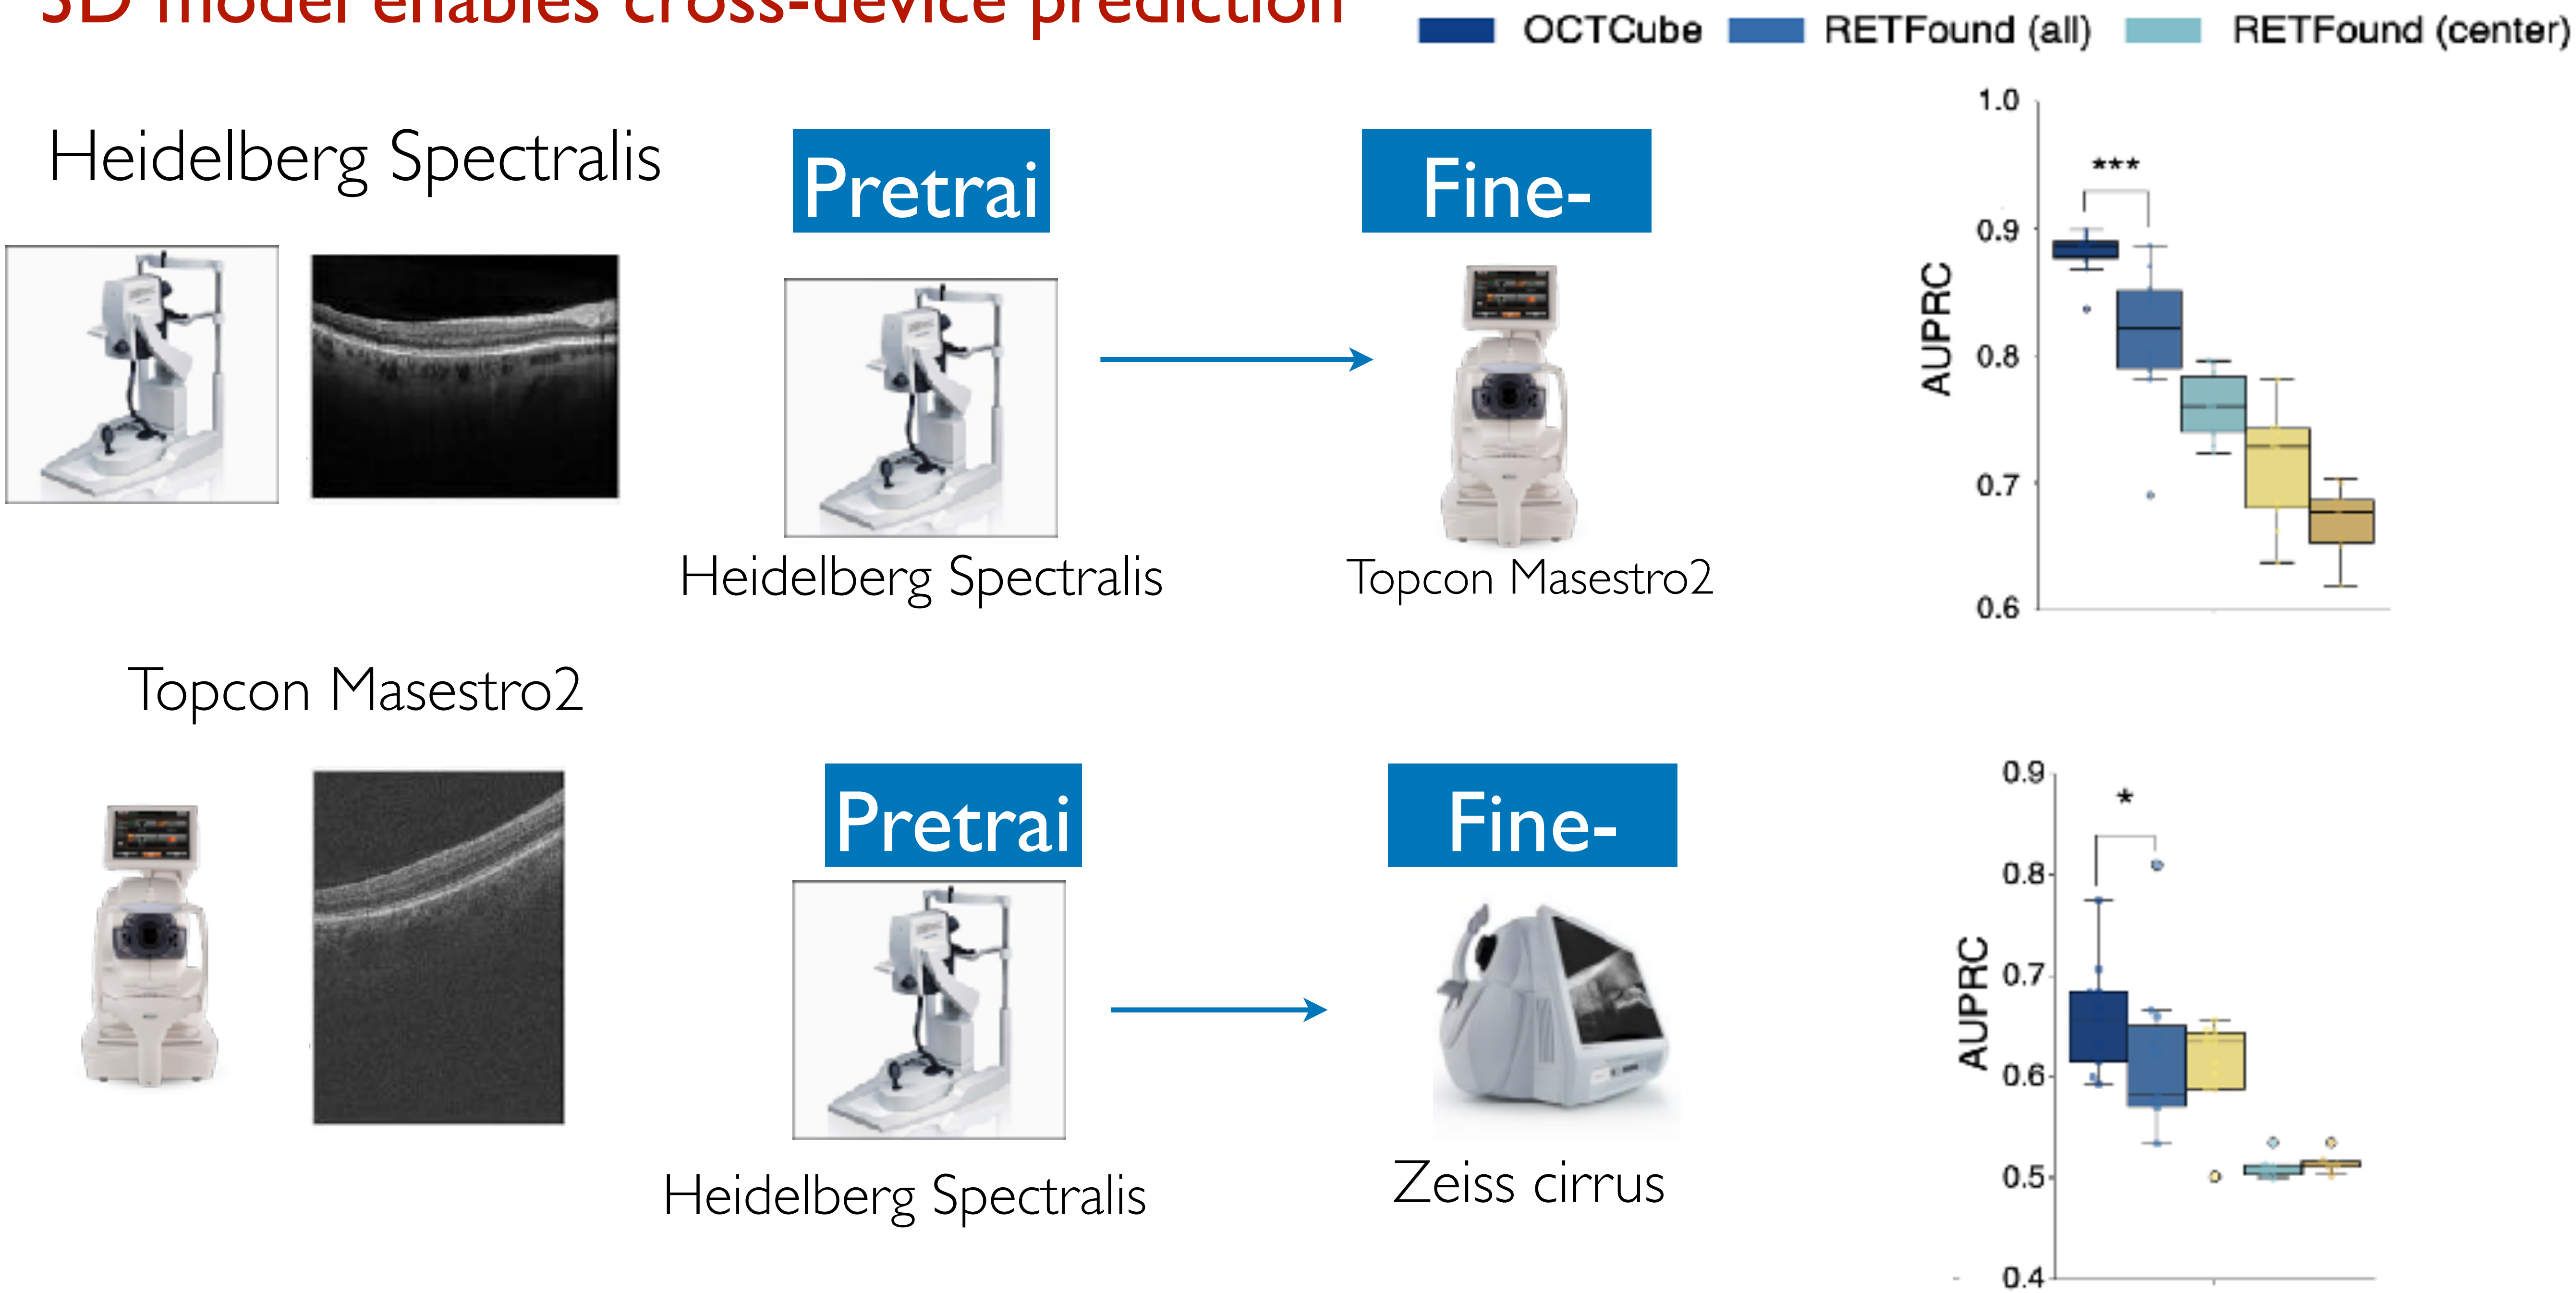

50

Predict systemic diseases

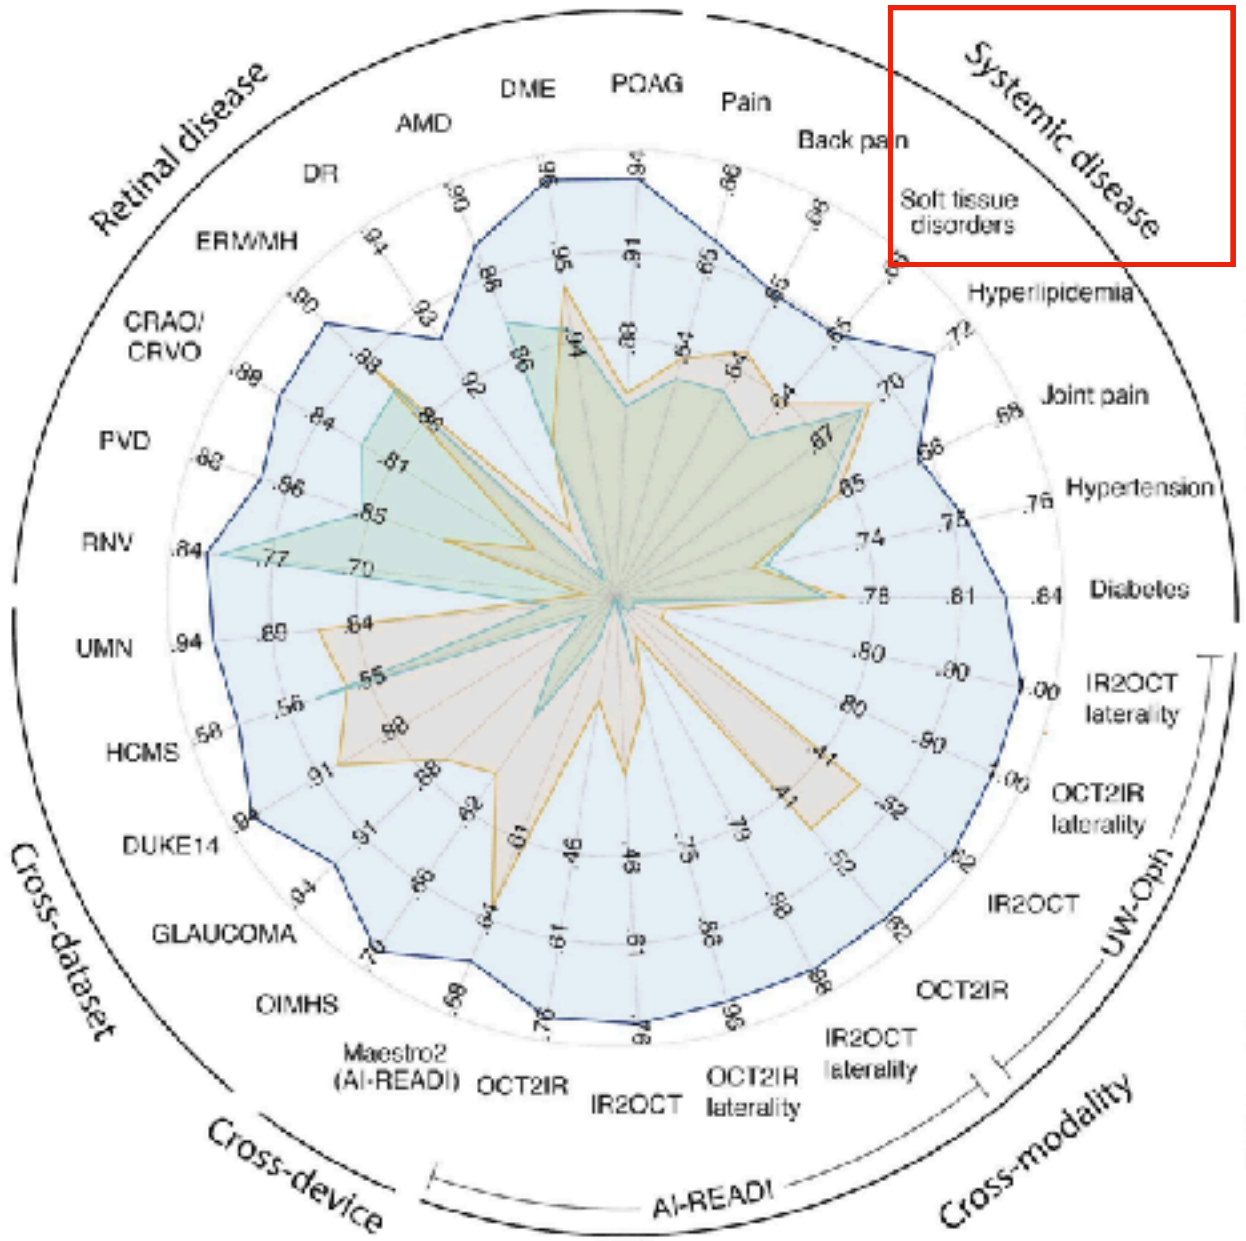

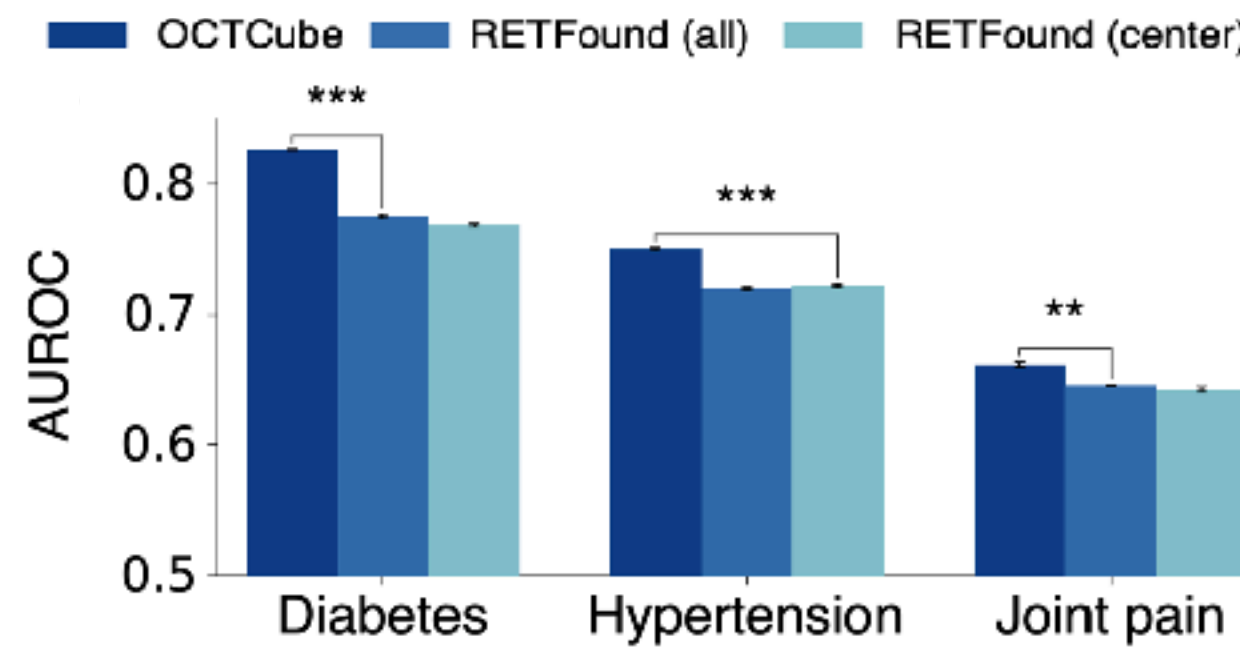

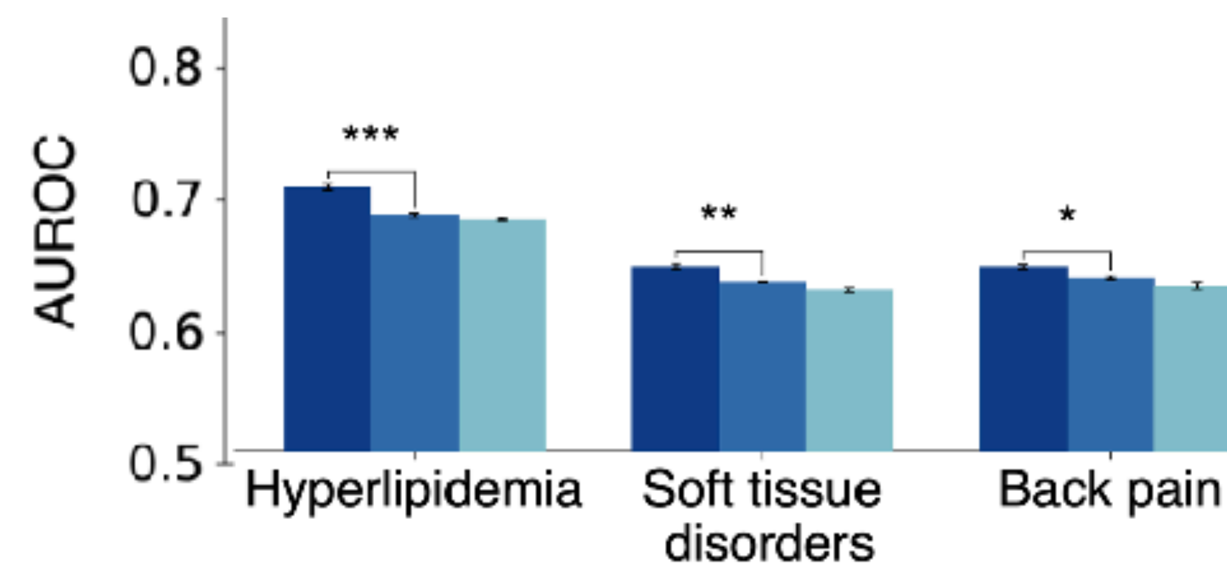

One-year early prediction of diabetes

1st visit: **2D model** cannot identify diabetes

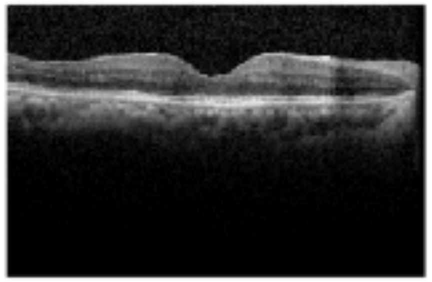

1st visit: **3D model** detects diabetes signal

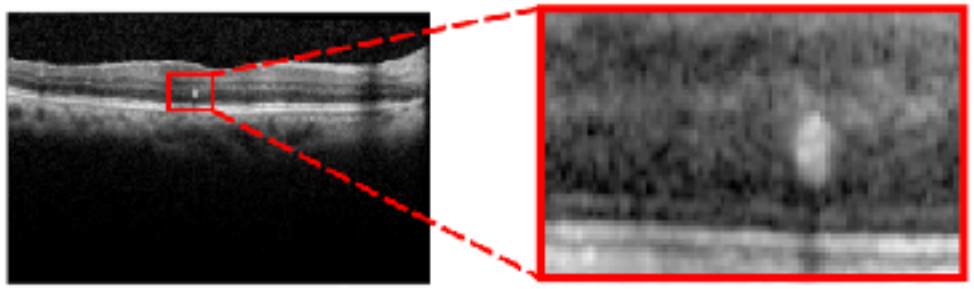

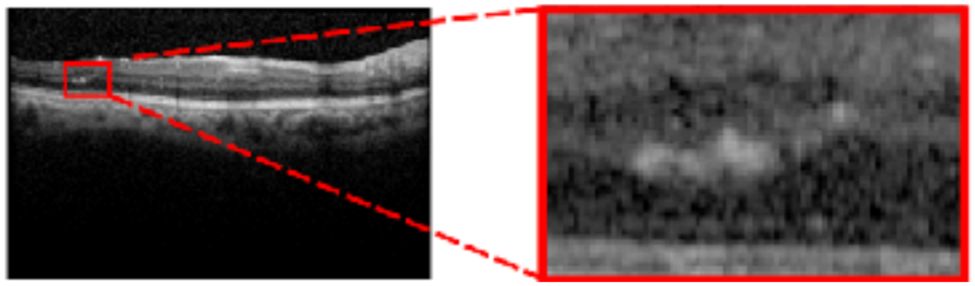

One-year early prediction of diabetes

1st visit: **2D model** cannot identify diabetes

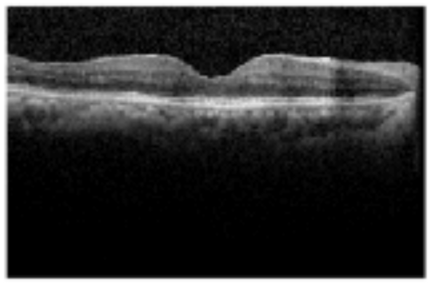

1st visit: **3D model** detects diabetes signal

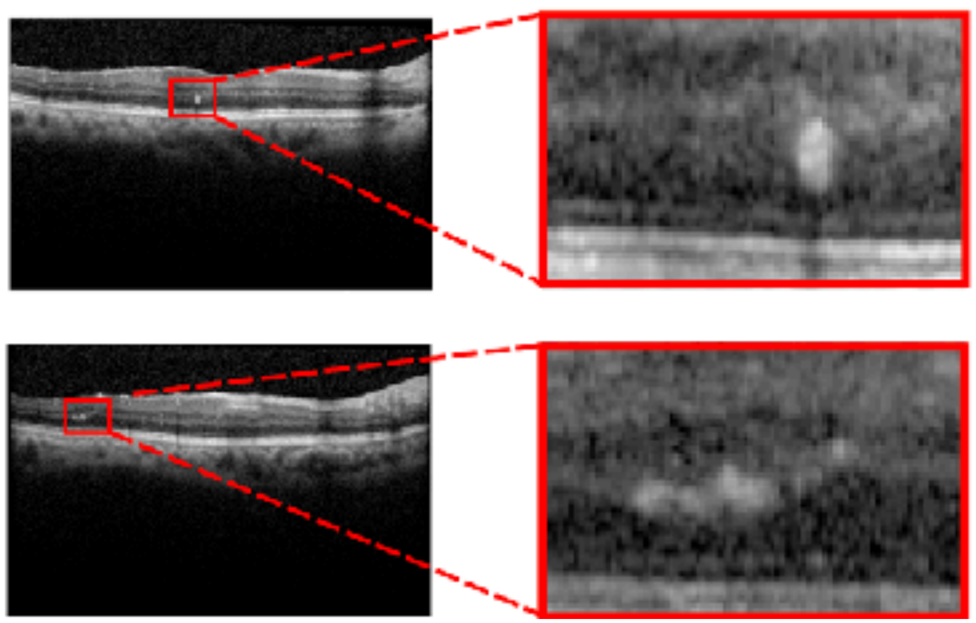

1 year later

2D model detects diabetes

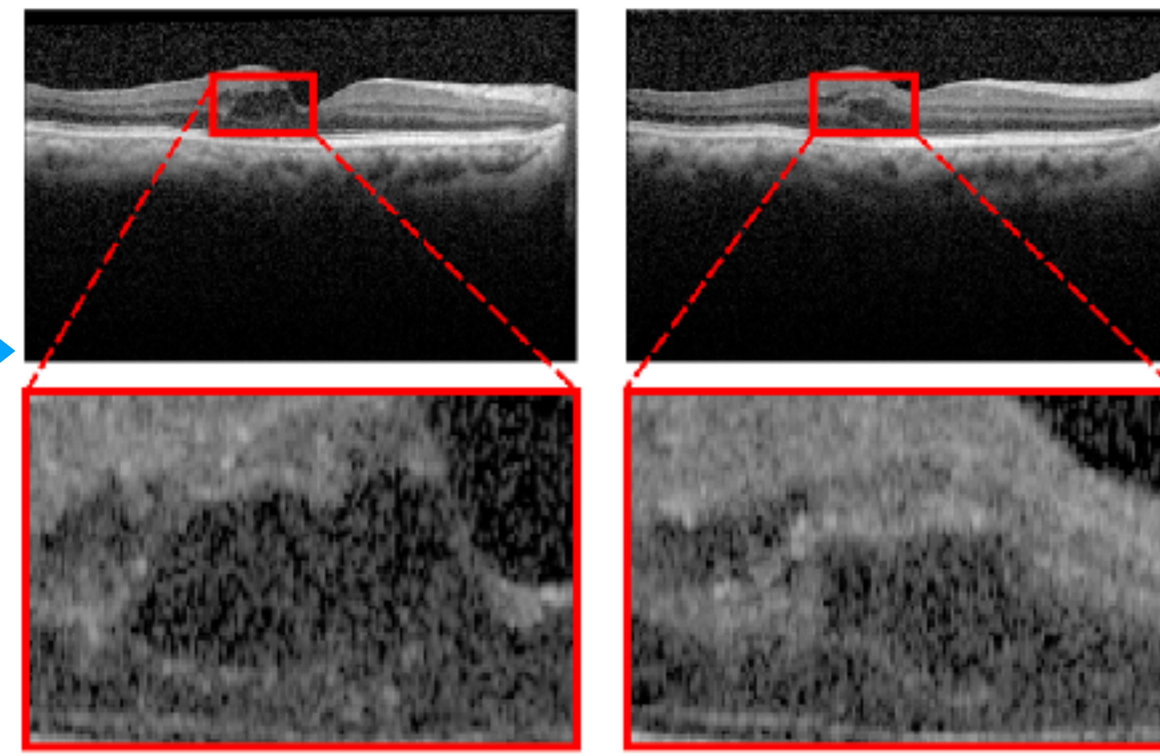

Today's talk: 3 parts

1. Pathology foundation model
2. 3D retinal foundation model

3. A multi-modal foundation model integrating 9 imaging modalities

Foundation model (2022)
One model for all tasks

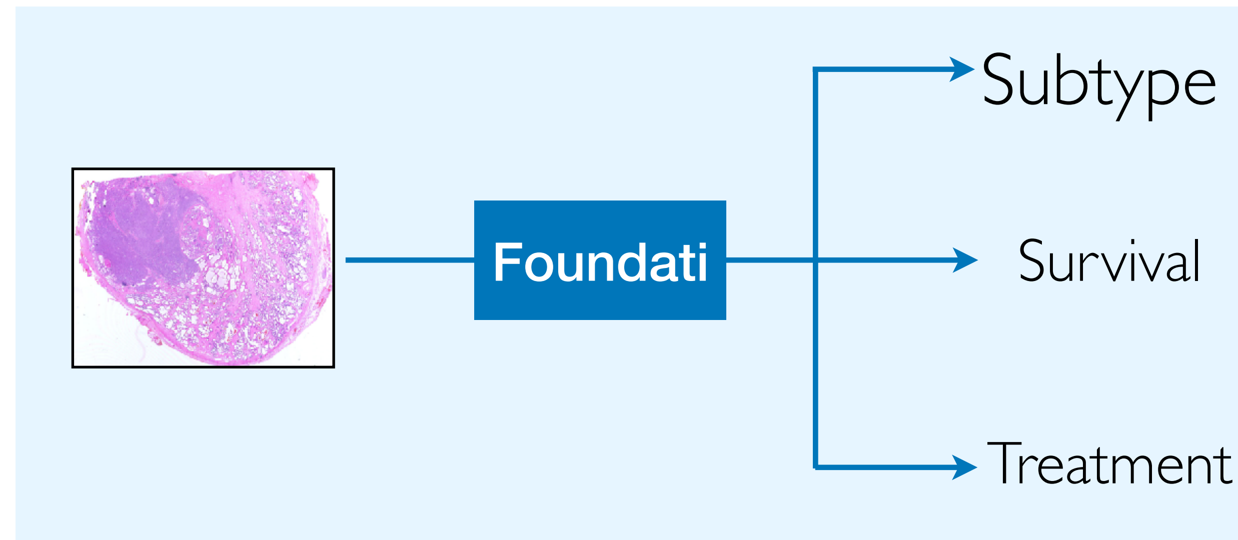

Multi-modal foundation model (2023)
One model takes different inputs

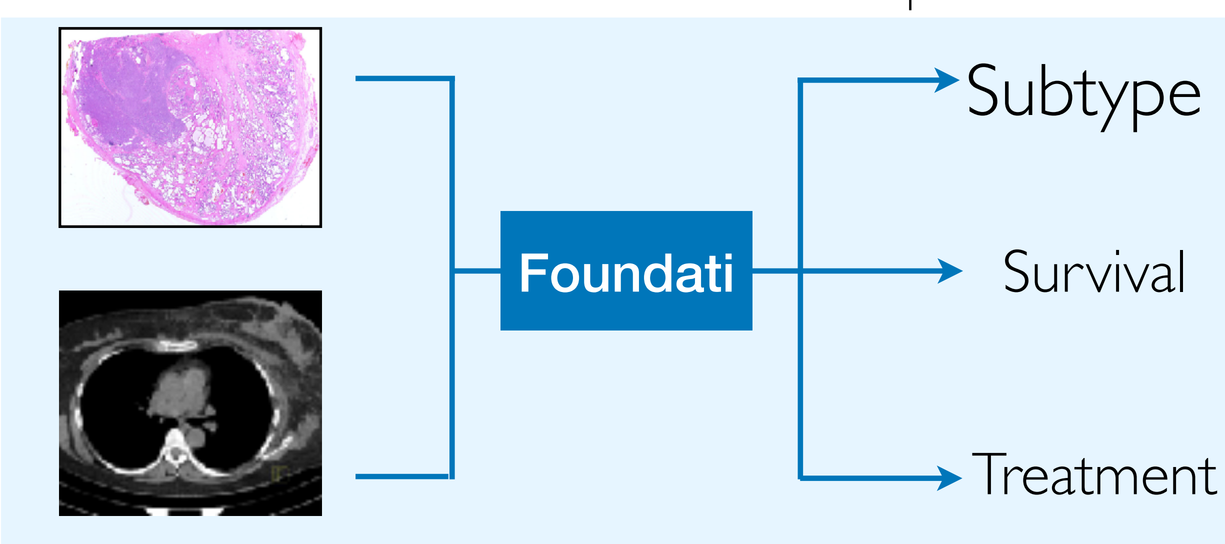

Multi-agent (2024)
Integrate multiple foundation models

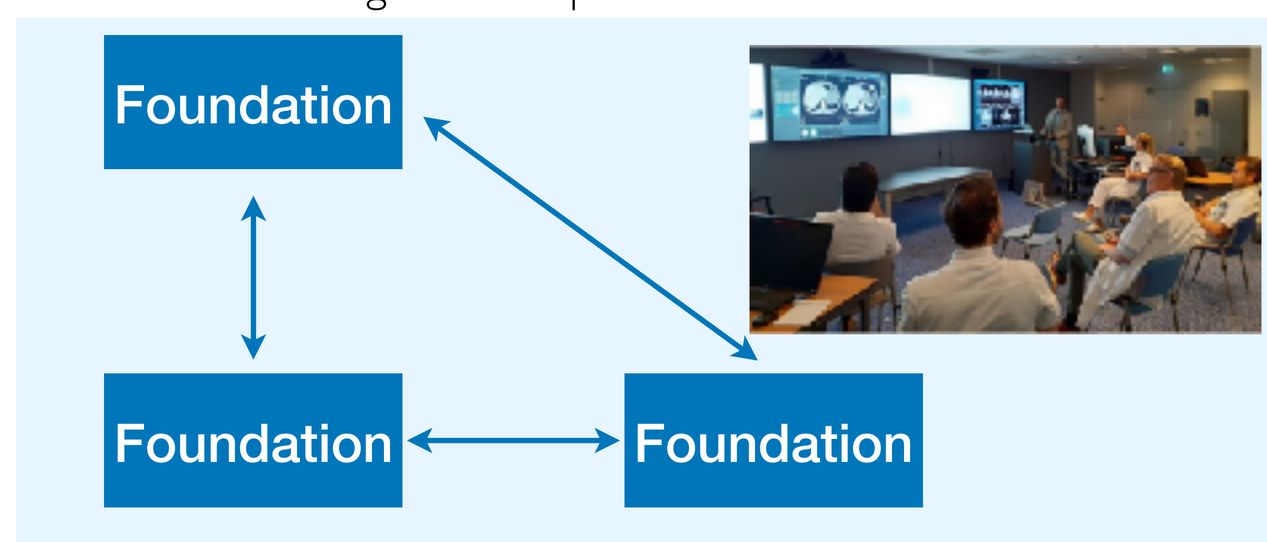

Multi-modality analysis in GigaPath and OCTCube

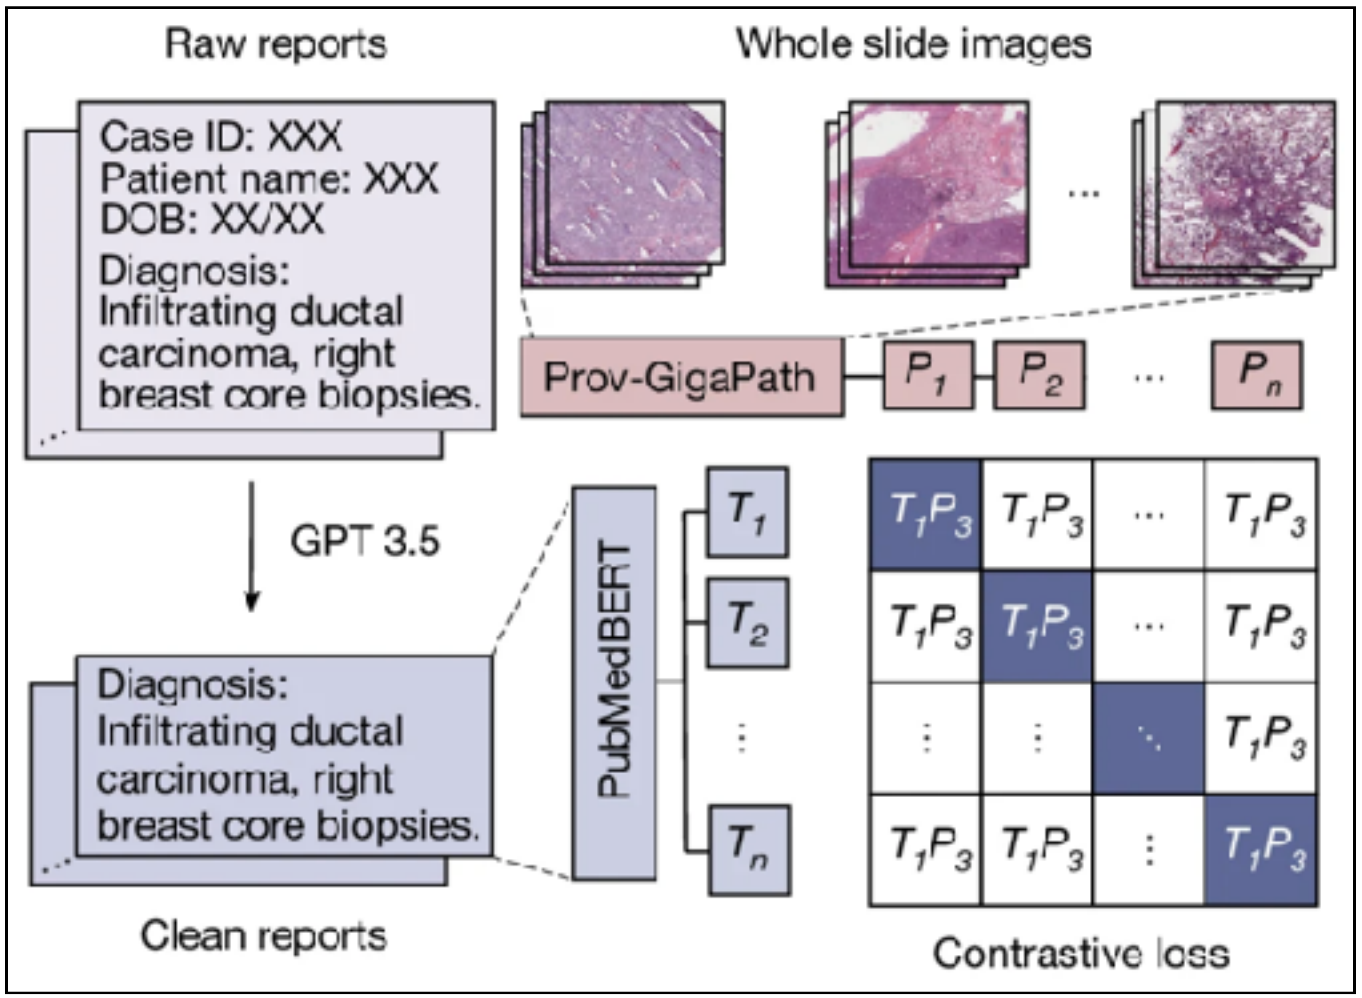

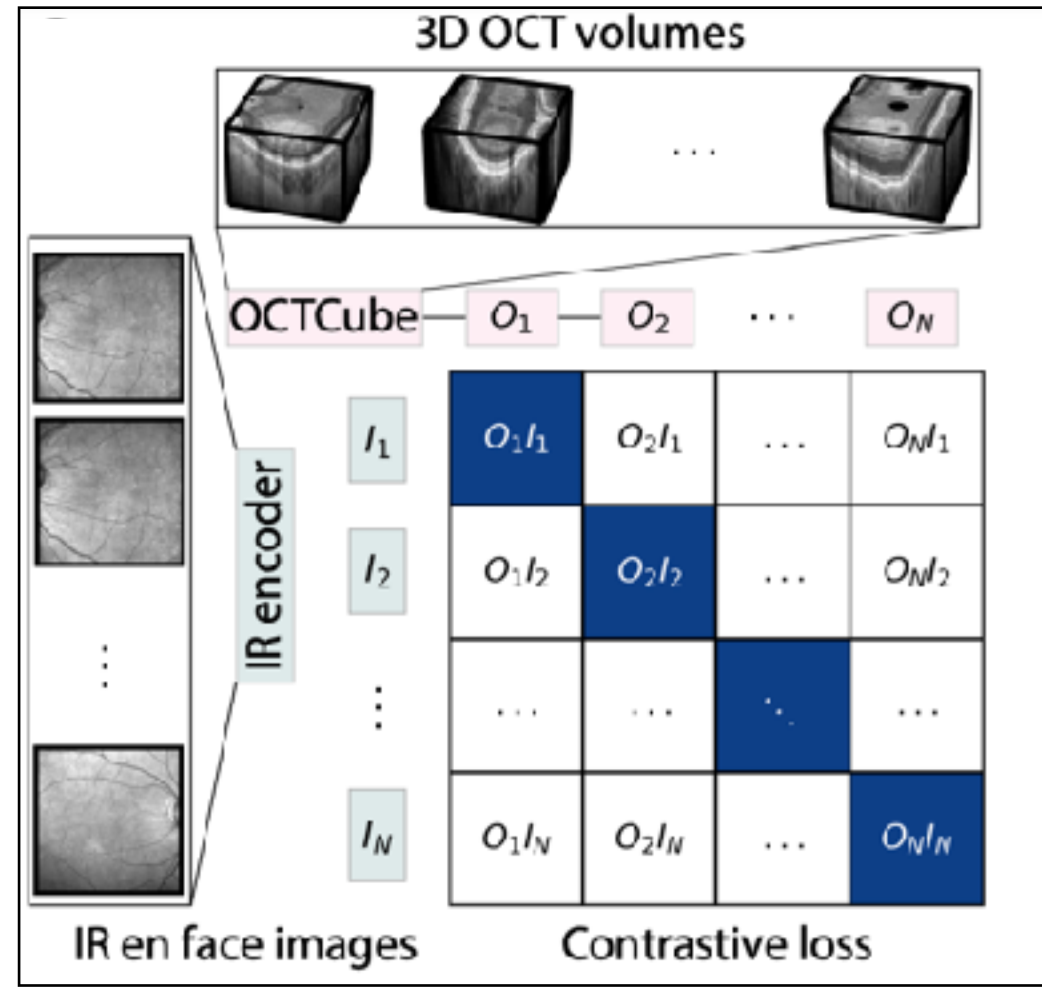

Medical report + Pathology image

OCT image + IR image

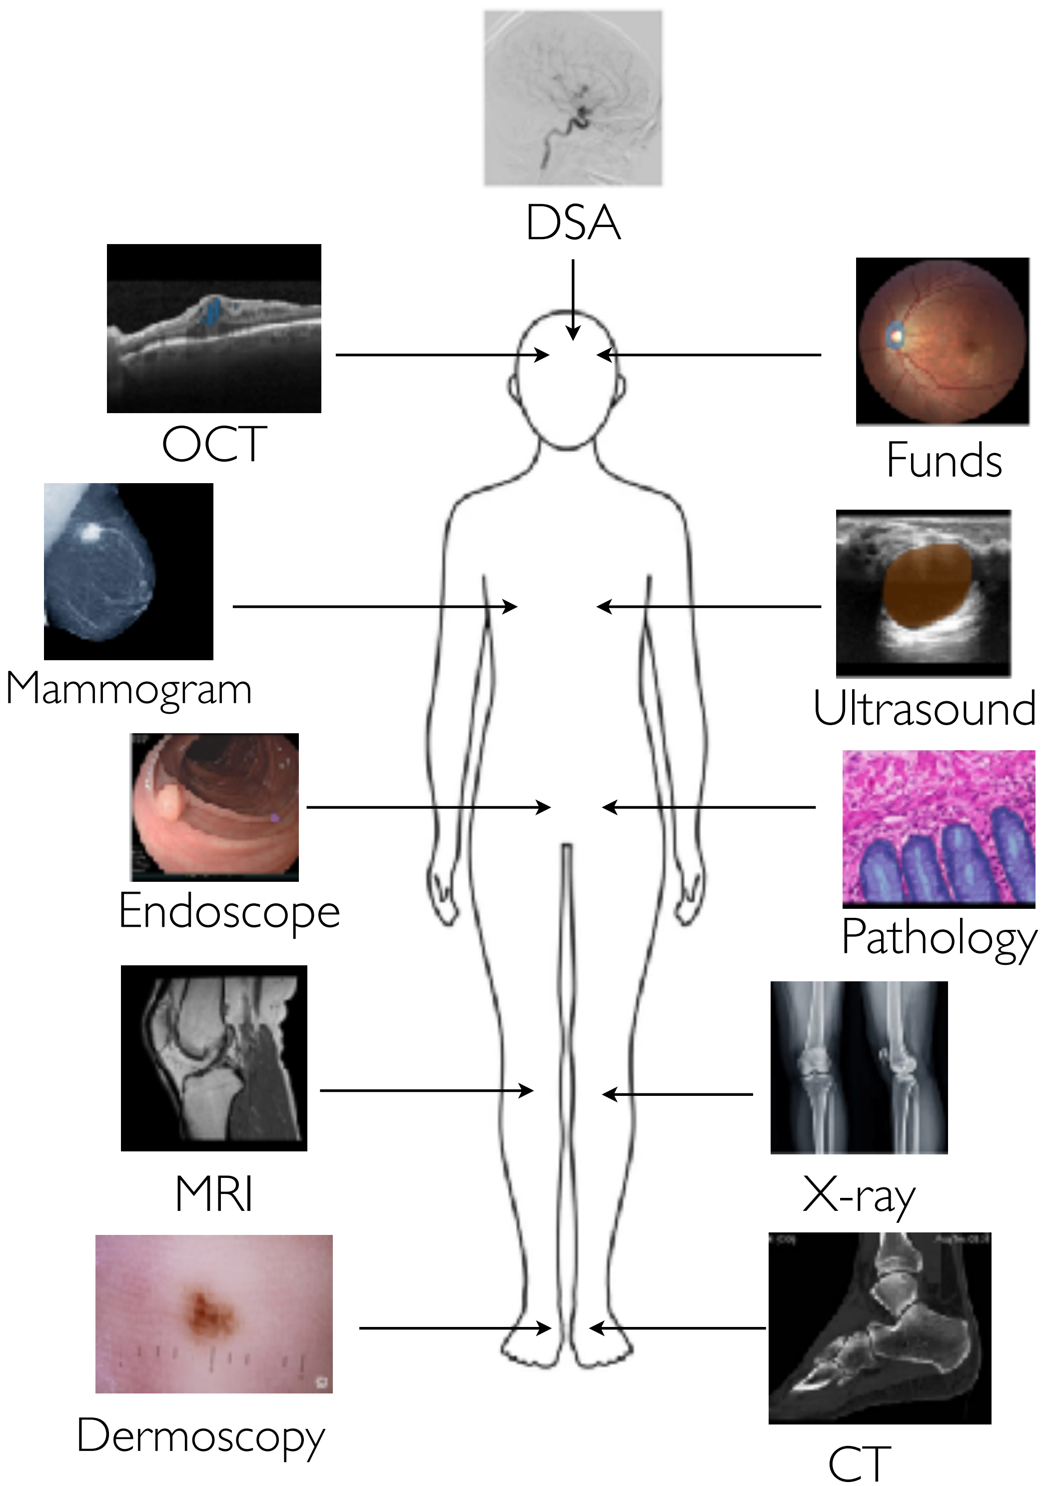

Existing cross-modal framework is Limited to two modalities

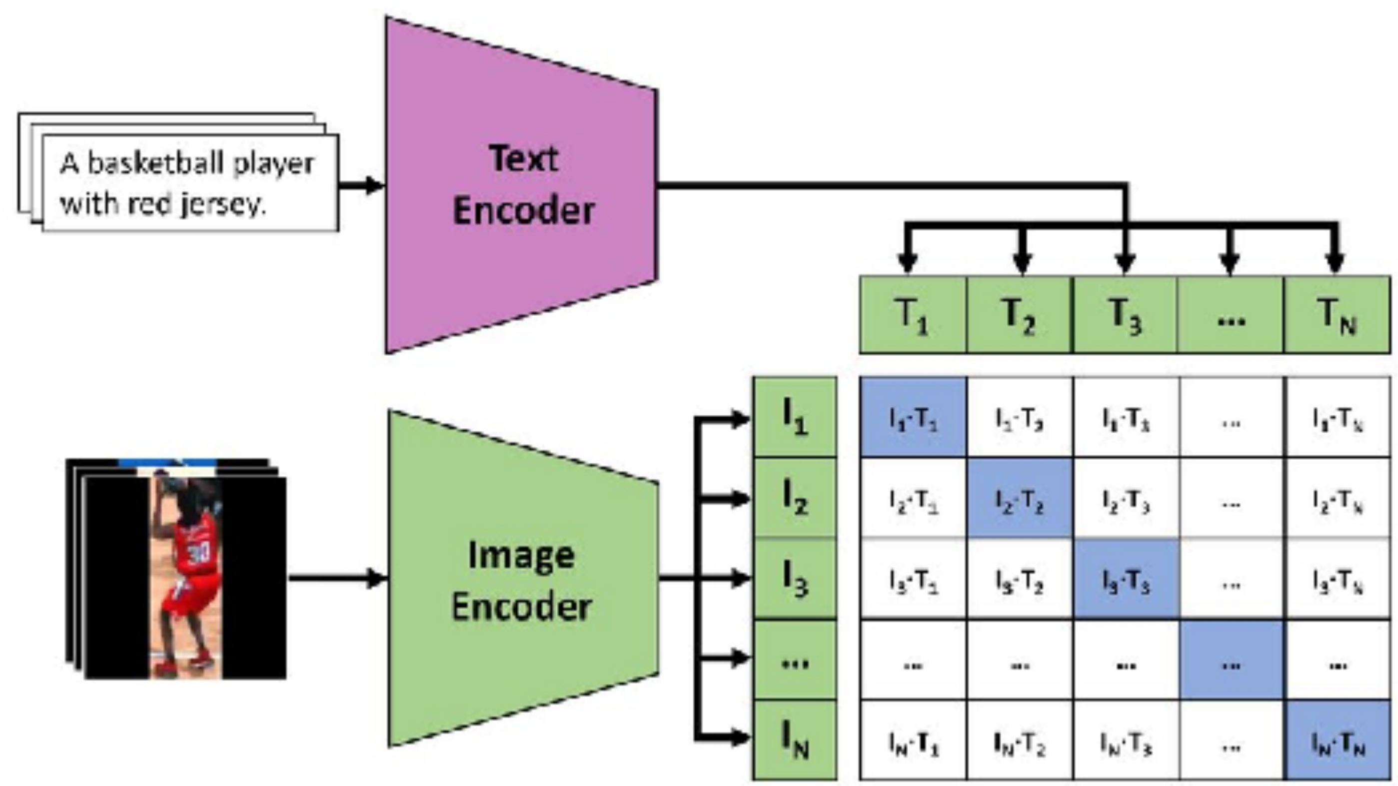

Image from CLIP-ReIdent

Our solution for integrating all image modalities

- Build one foundation model for each modality, then project different models into the same space
- What should be the Esperanto for medicine?

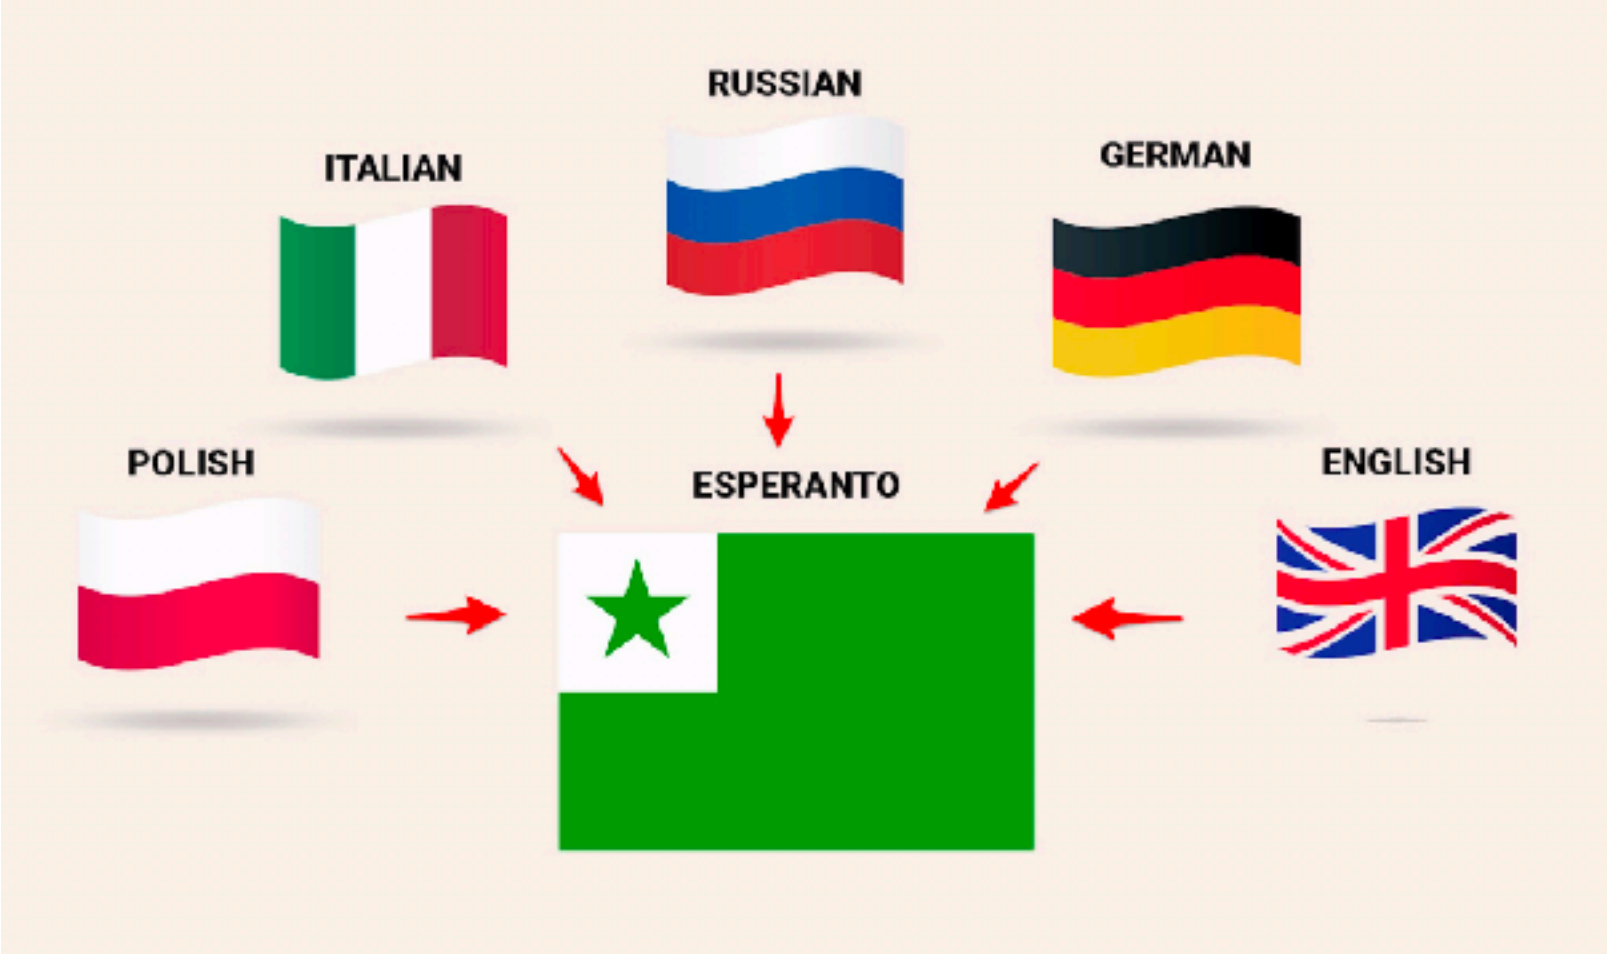

Image source: https://traductanet.com/latest-news/what-is-esperanto/

BiomedParse: use human language as the Esperanto to integrate medical imaging modalities

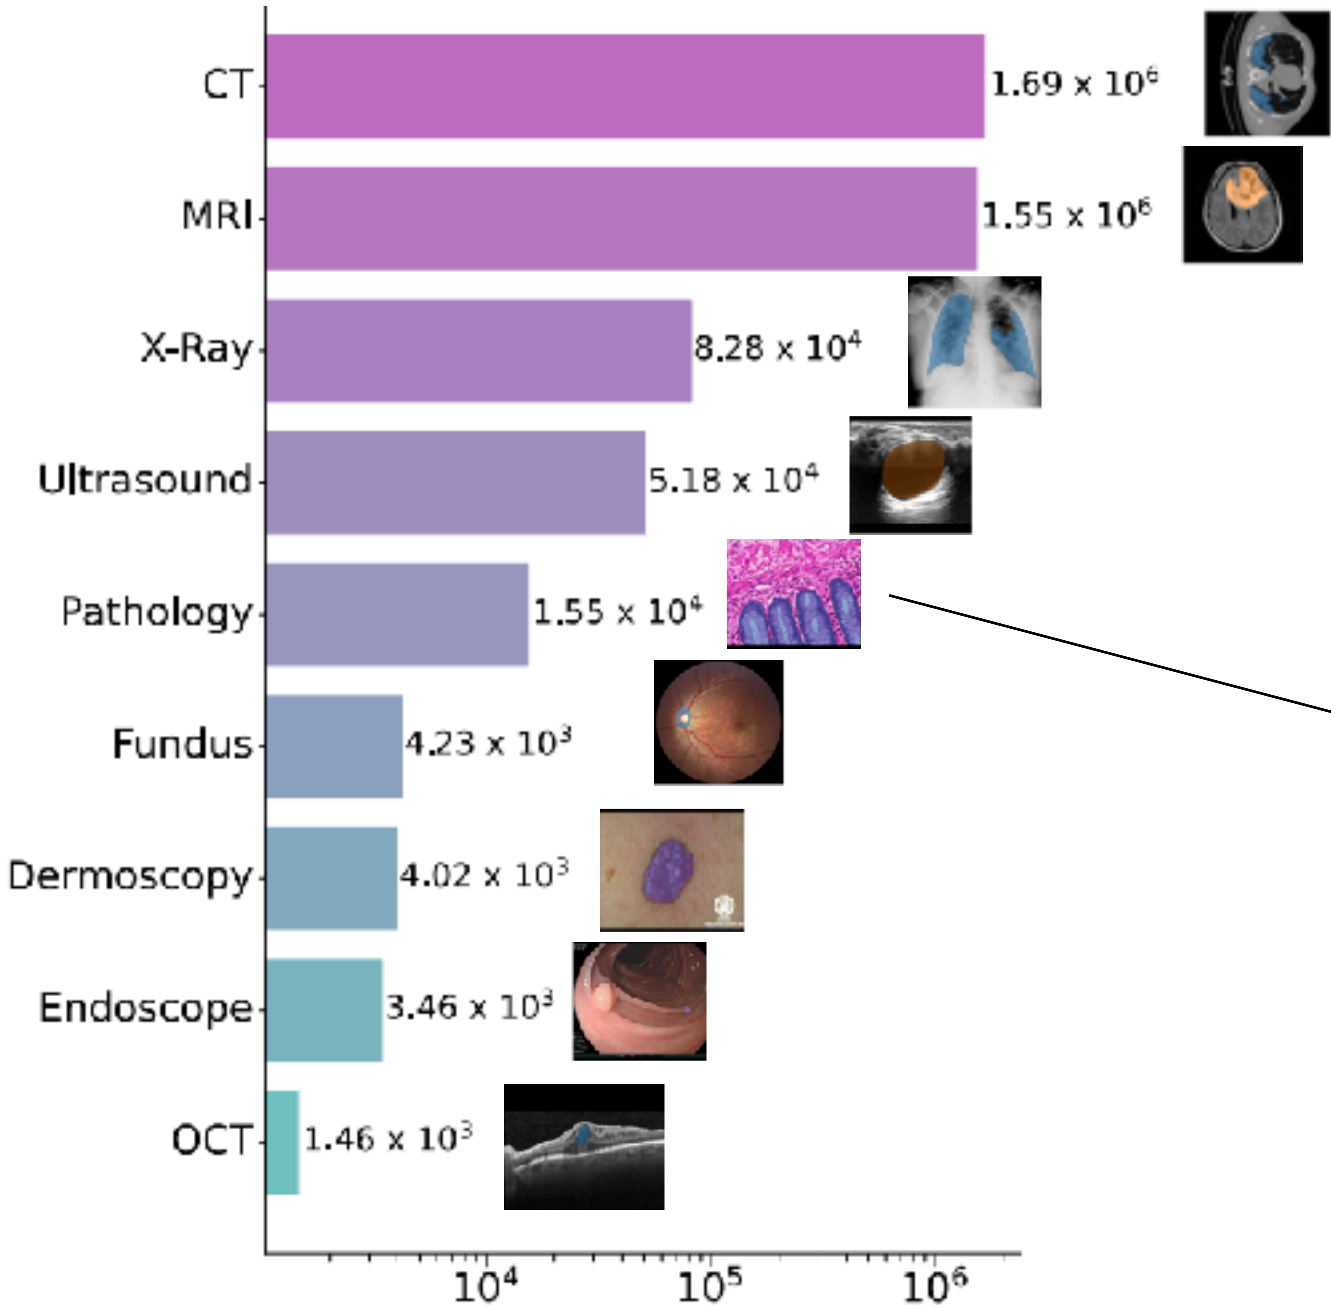

Project 6.8 million images from 9 image modalities into the text space

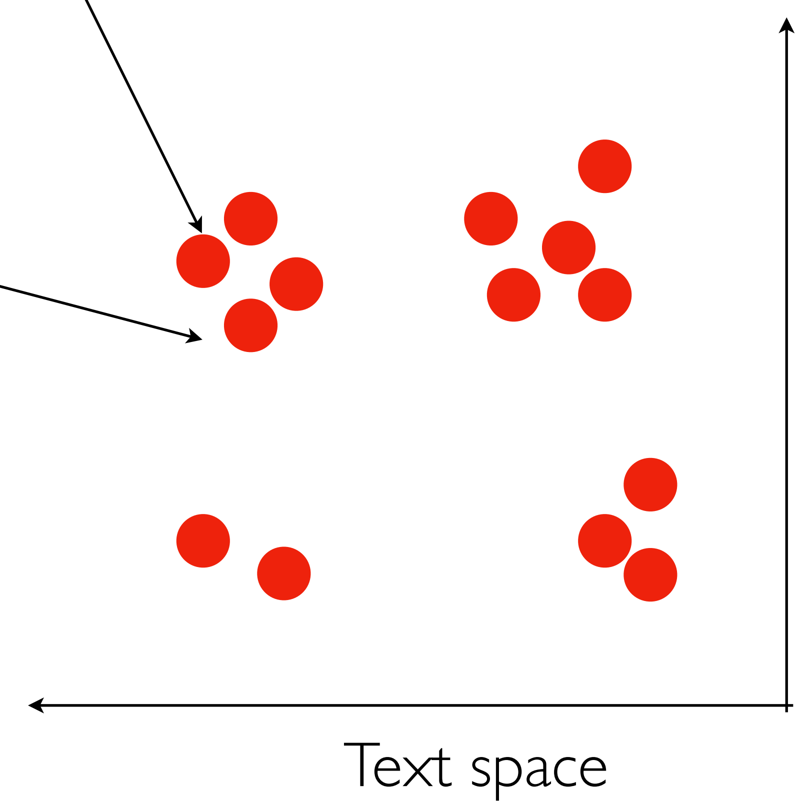

Liver in abdomen CT

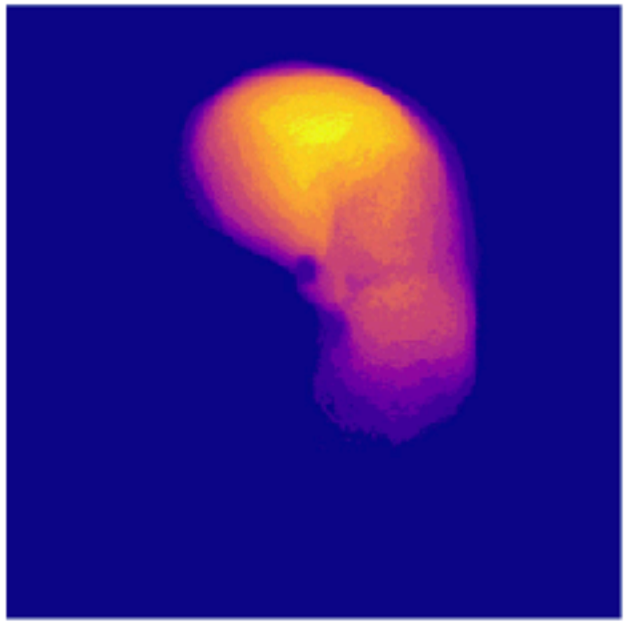

Kidney in abdomen CT

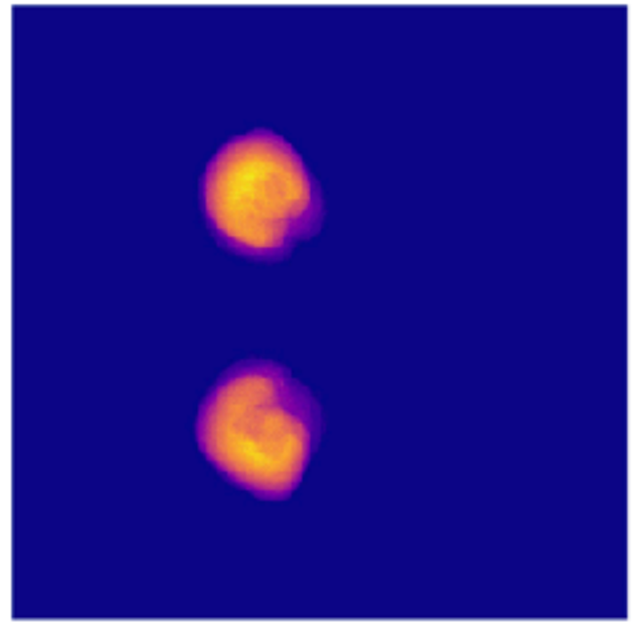

Pancreas in abdomen CT

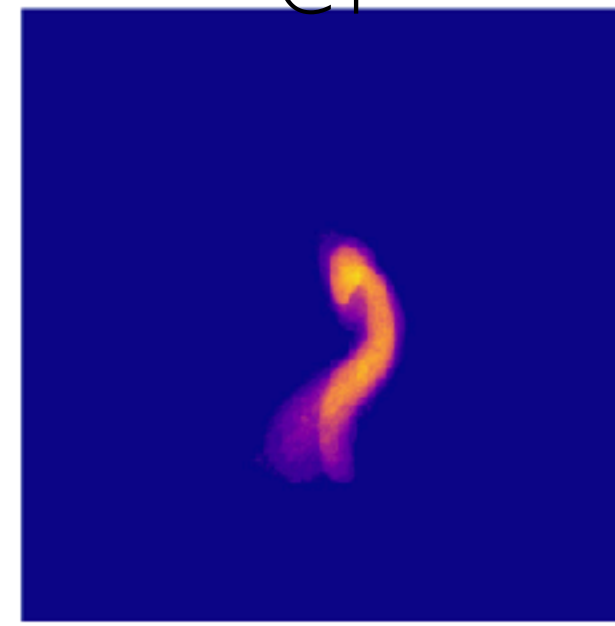

Liver vessel in abdomen CT

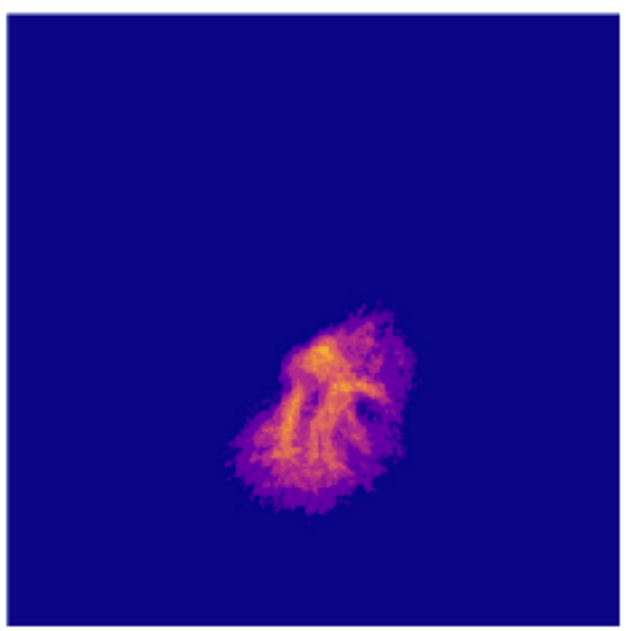

Inflammatory cells in pathology

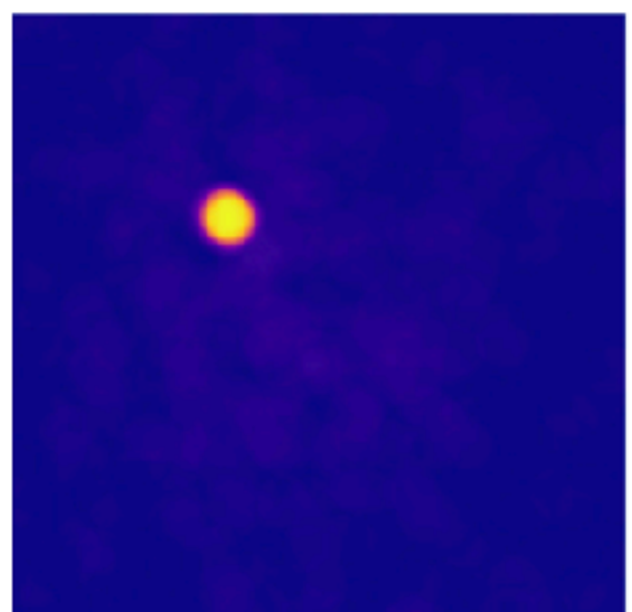

Neoplastic cells in pathology

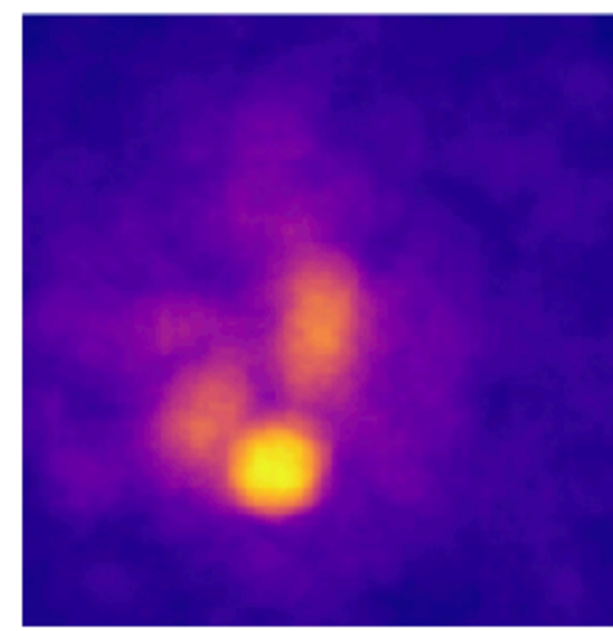

A foundation model for joint segmentation, detection and recognition across 9 modalities

nature methods

Explore content ⌄ About the journal ⌄ Publish with us ⌄ Subscribe

nature > nature methods > articles > article

Article | Published: 18 November 2024

A foundation model for joint segmentation, detection and recognition of biomedical objects across nine modalities

Theodore Zhao, Yu Gu, Jianwei Yang, Naoto Usuyama, Ho Hin Lee, Sid Kiblawi, Tristan Naumann, Jianfeng Gao, Angela Crabtree, Jacob Abel, Christine Moung-Wen, Brian Piening, Carlo Bifulco, Mu Wei ✉, Hoifung Poon ✉ & Sheng Wang ✉

Nature Methods (2024) | Cite this article

Available as open-source model on Azure AI and Hugging Face

Theodore Zhao

Aiden Gu

Sheng Wang
U of Washington

Mu Wei

Hoifung Poon

Demo1: talk to the AI model to find lung nodule

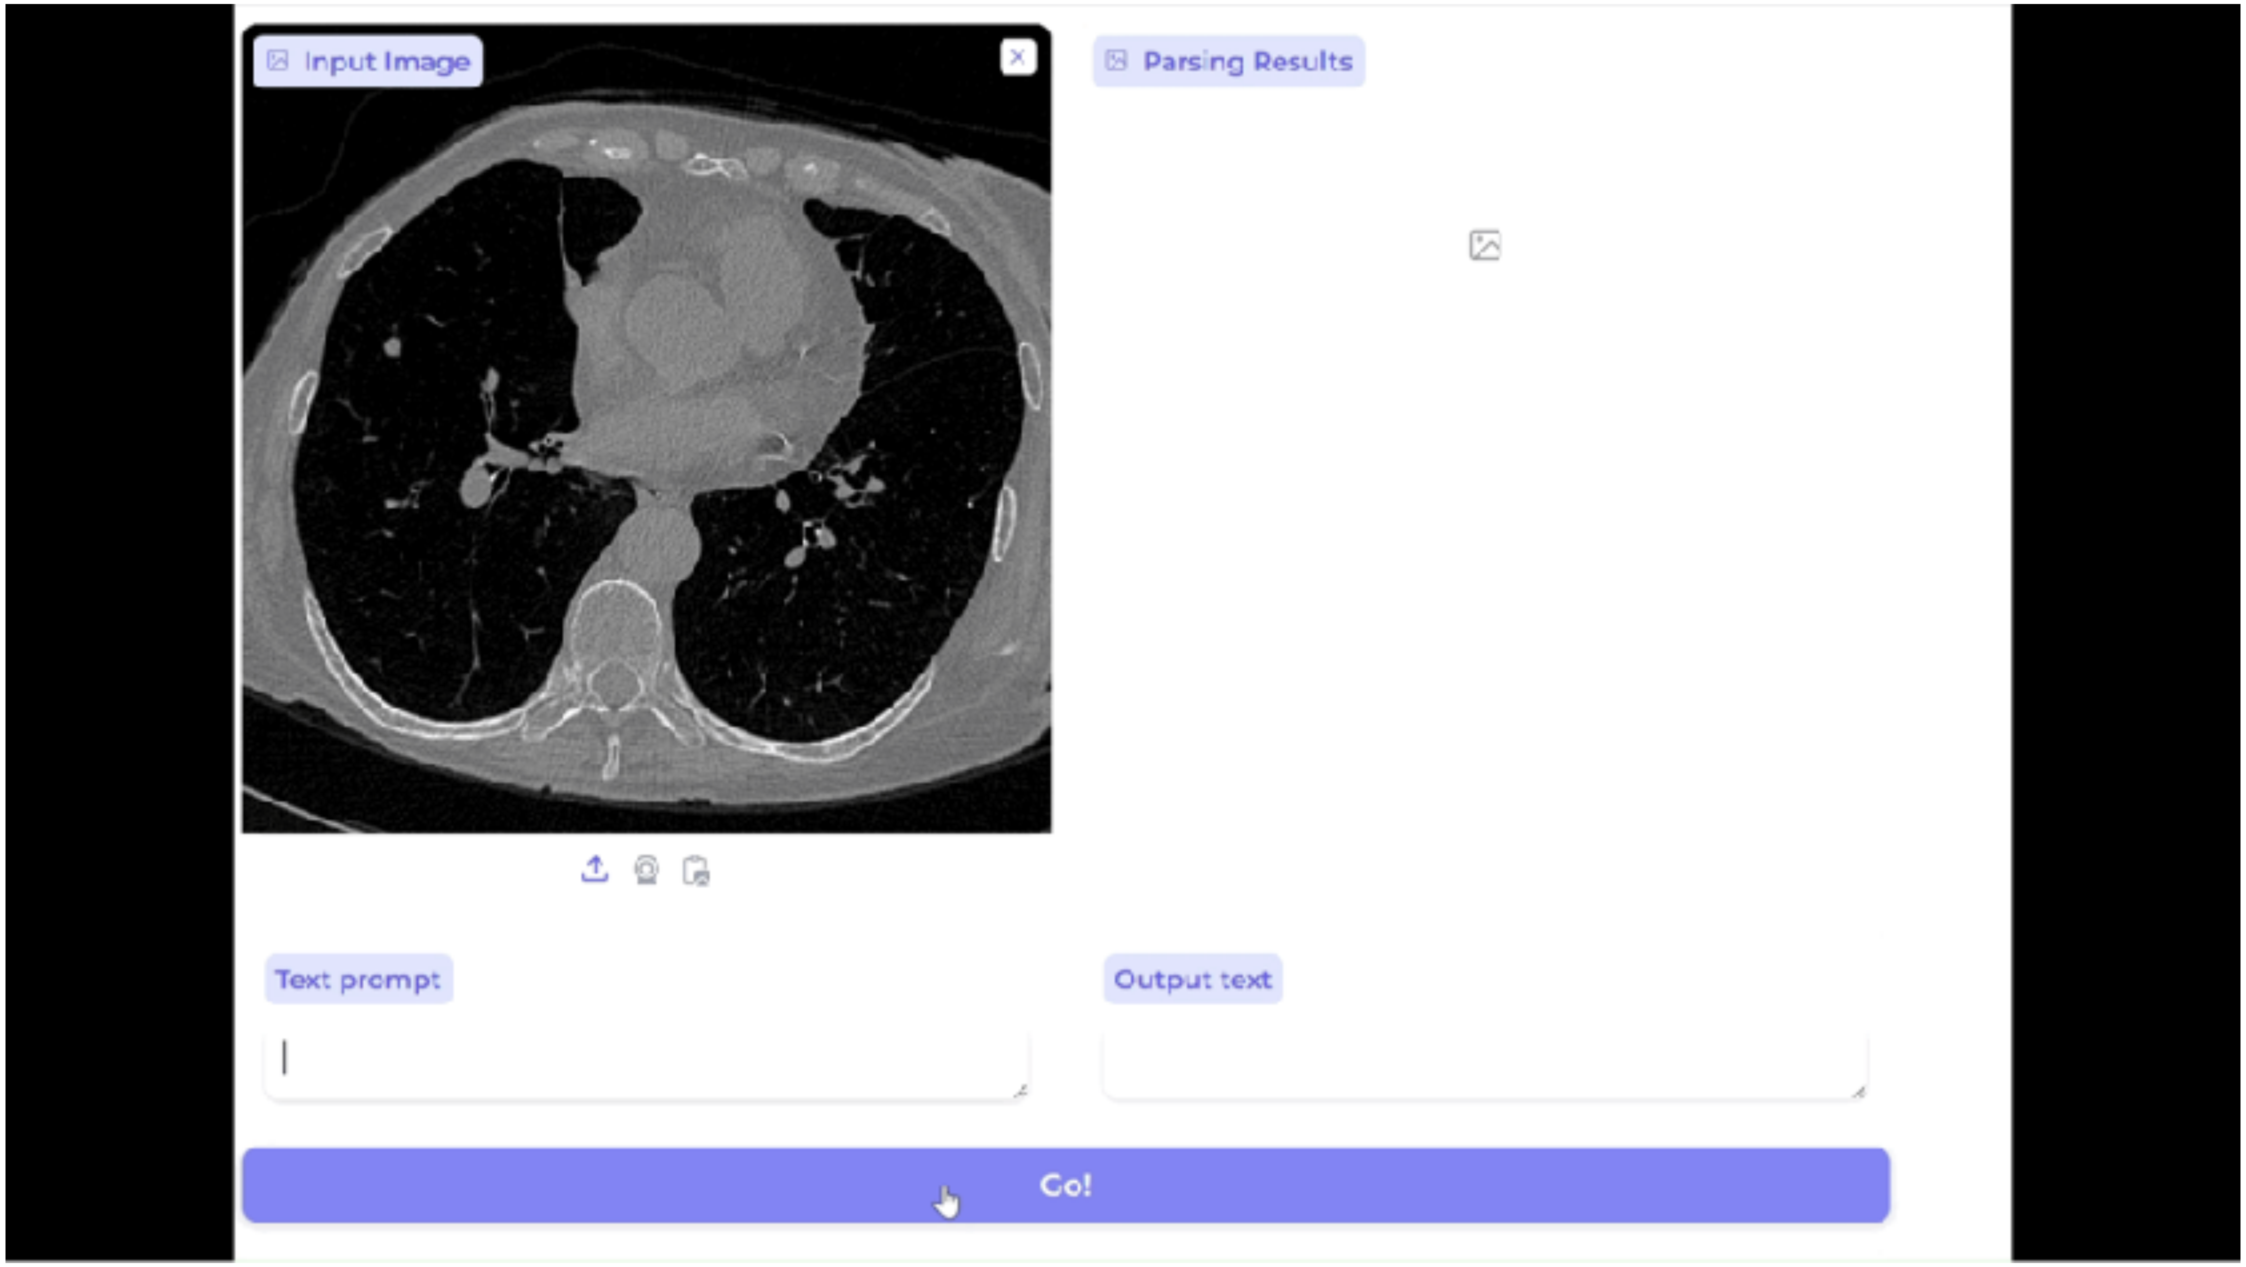

Demo1: talk to the AI model to find lung nodule

Demo2: talk to the model to find all cell types

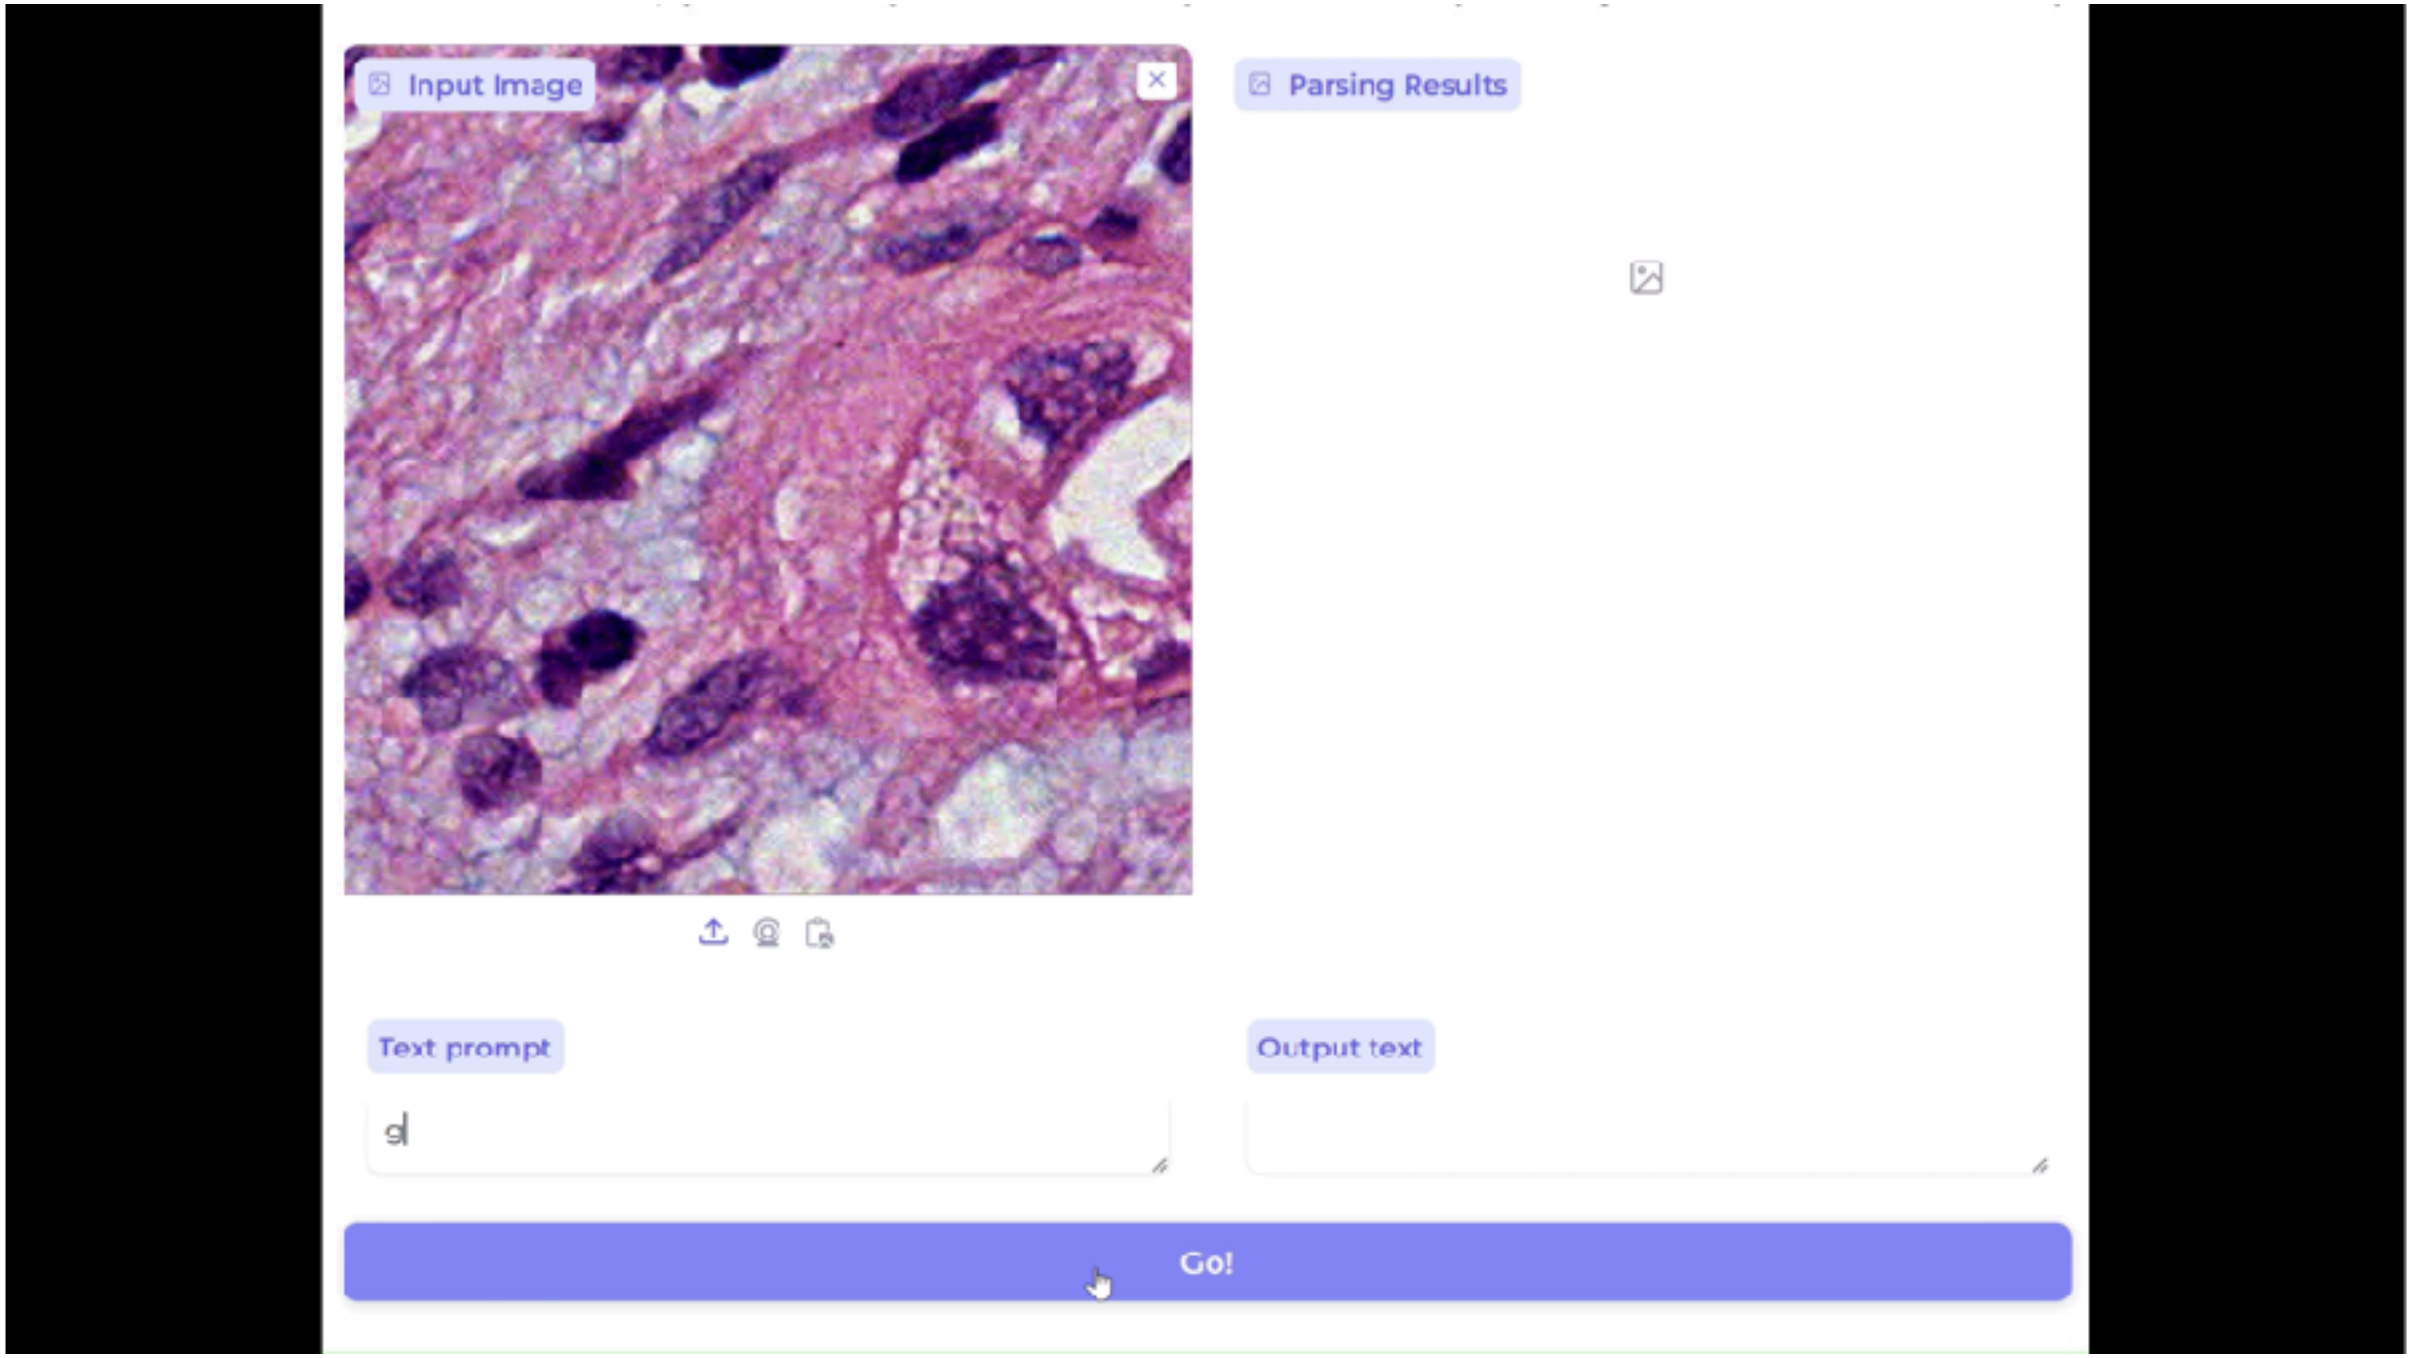

Demo2: talk to the model to find all cell types

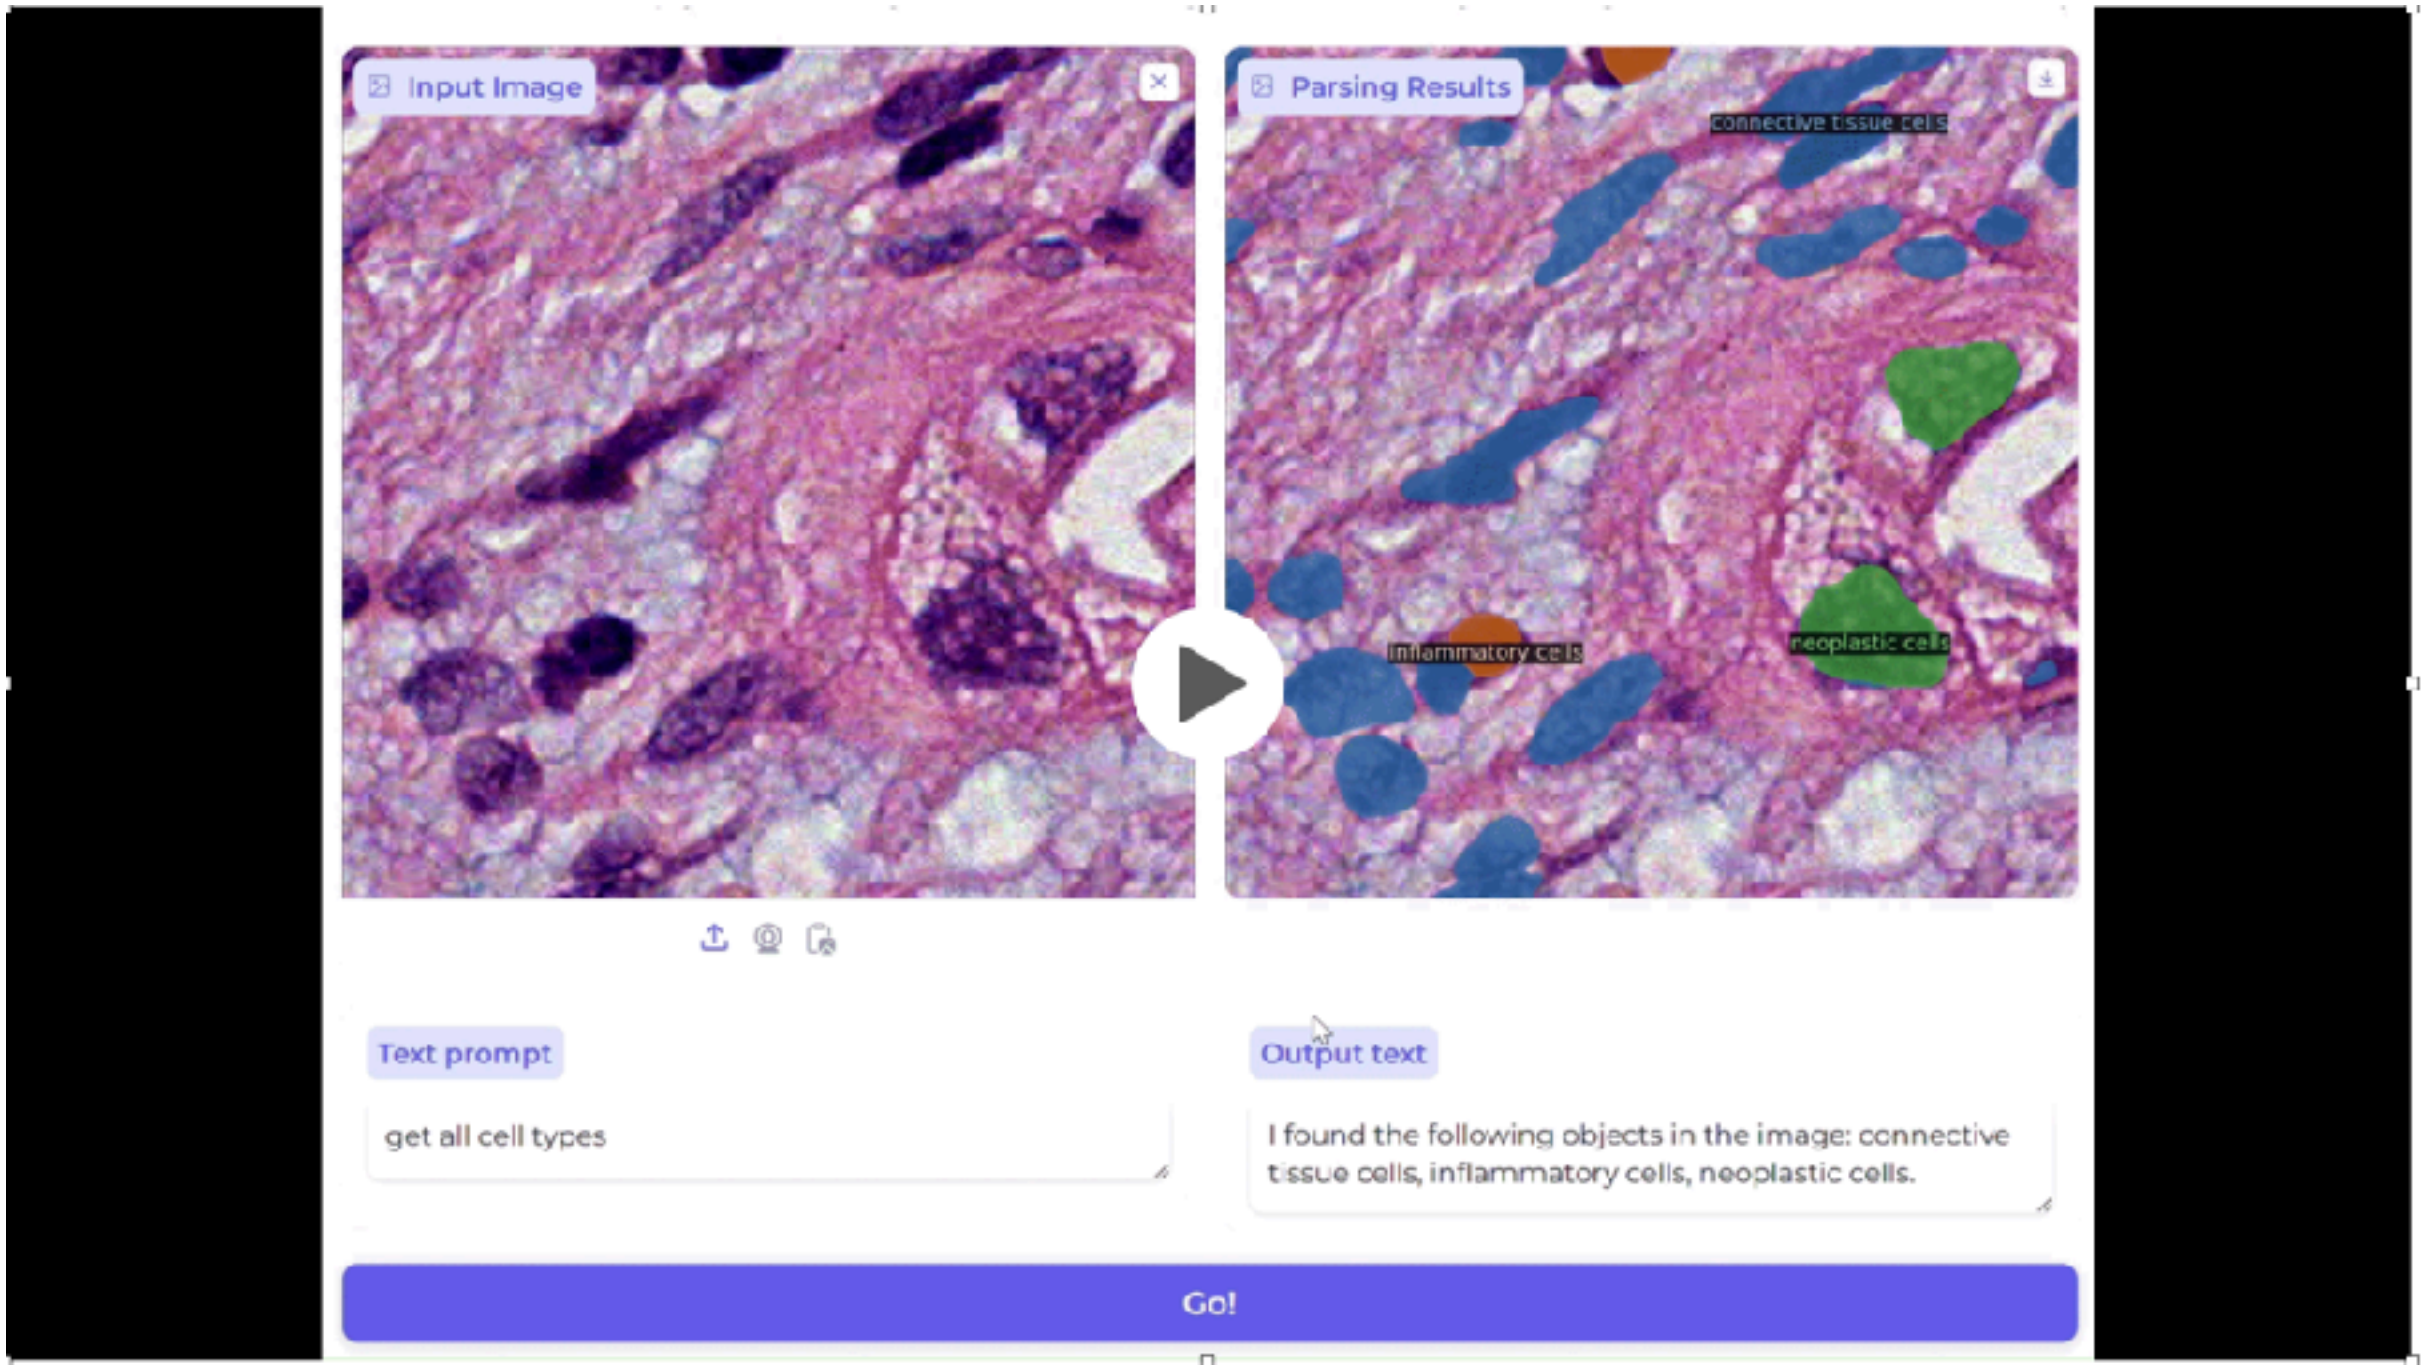

0.2 second to segment and label all cells with > 90% consistency

Pathologist annotation

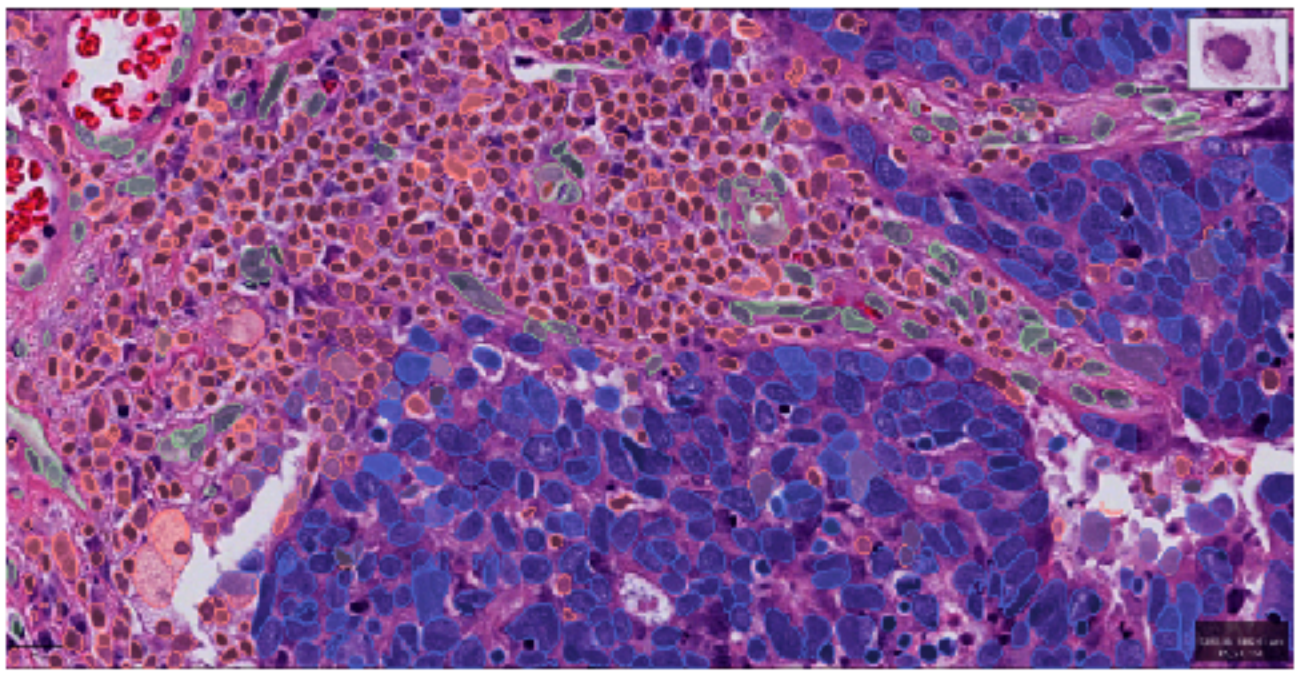

Tumor region Lymphocyte Stroma

BiomedParse prediction

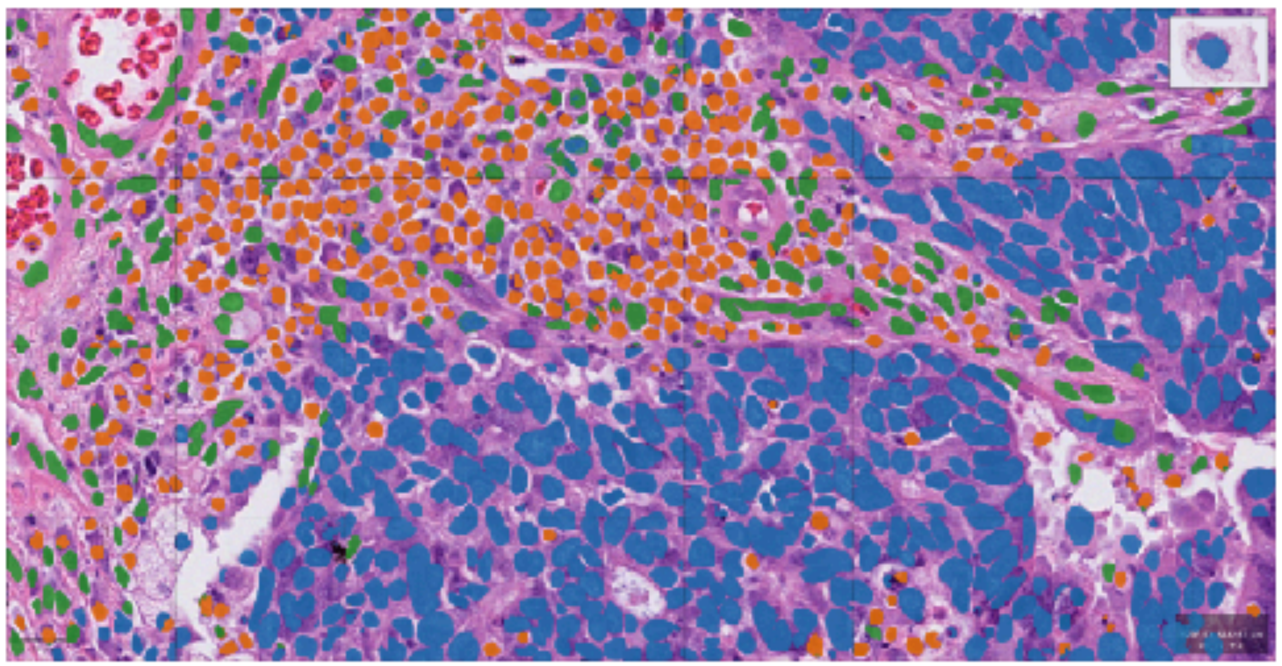

Neoplastic cells Inflammatory cells
Connective tissue cells

Data-raising from Collaborators at UW School of Medicine

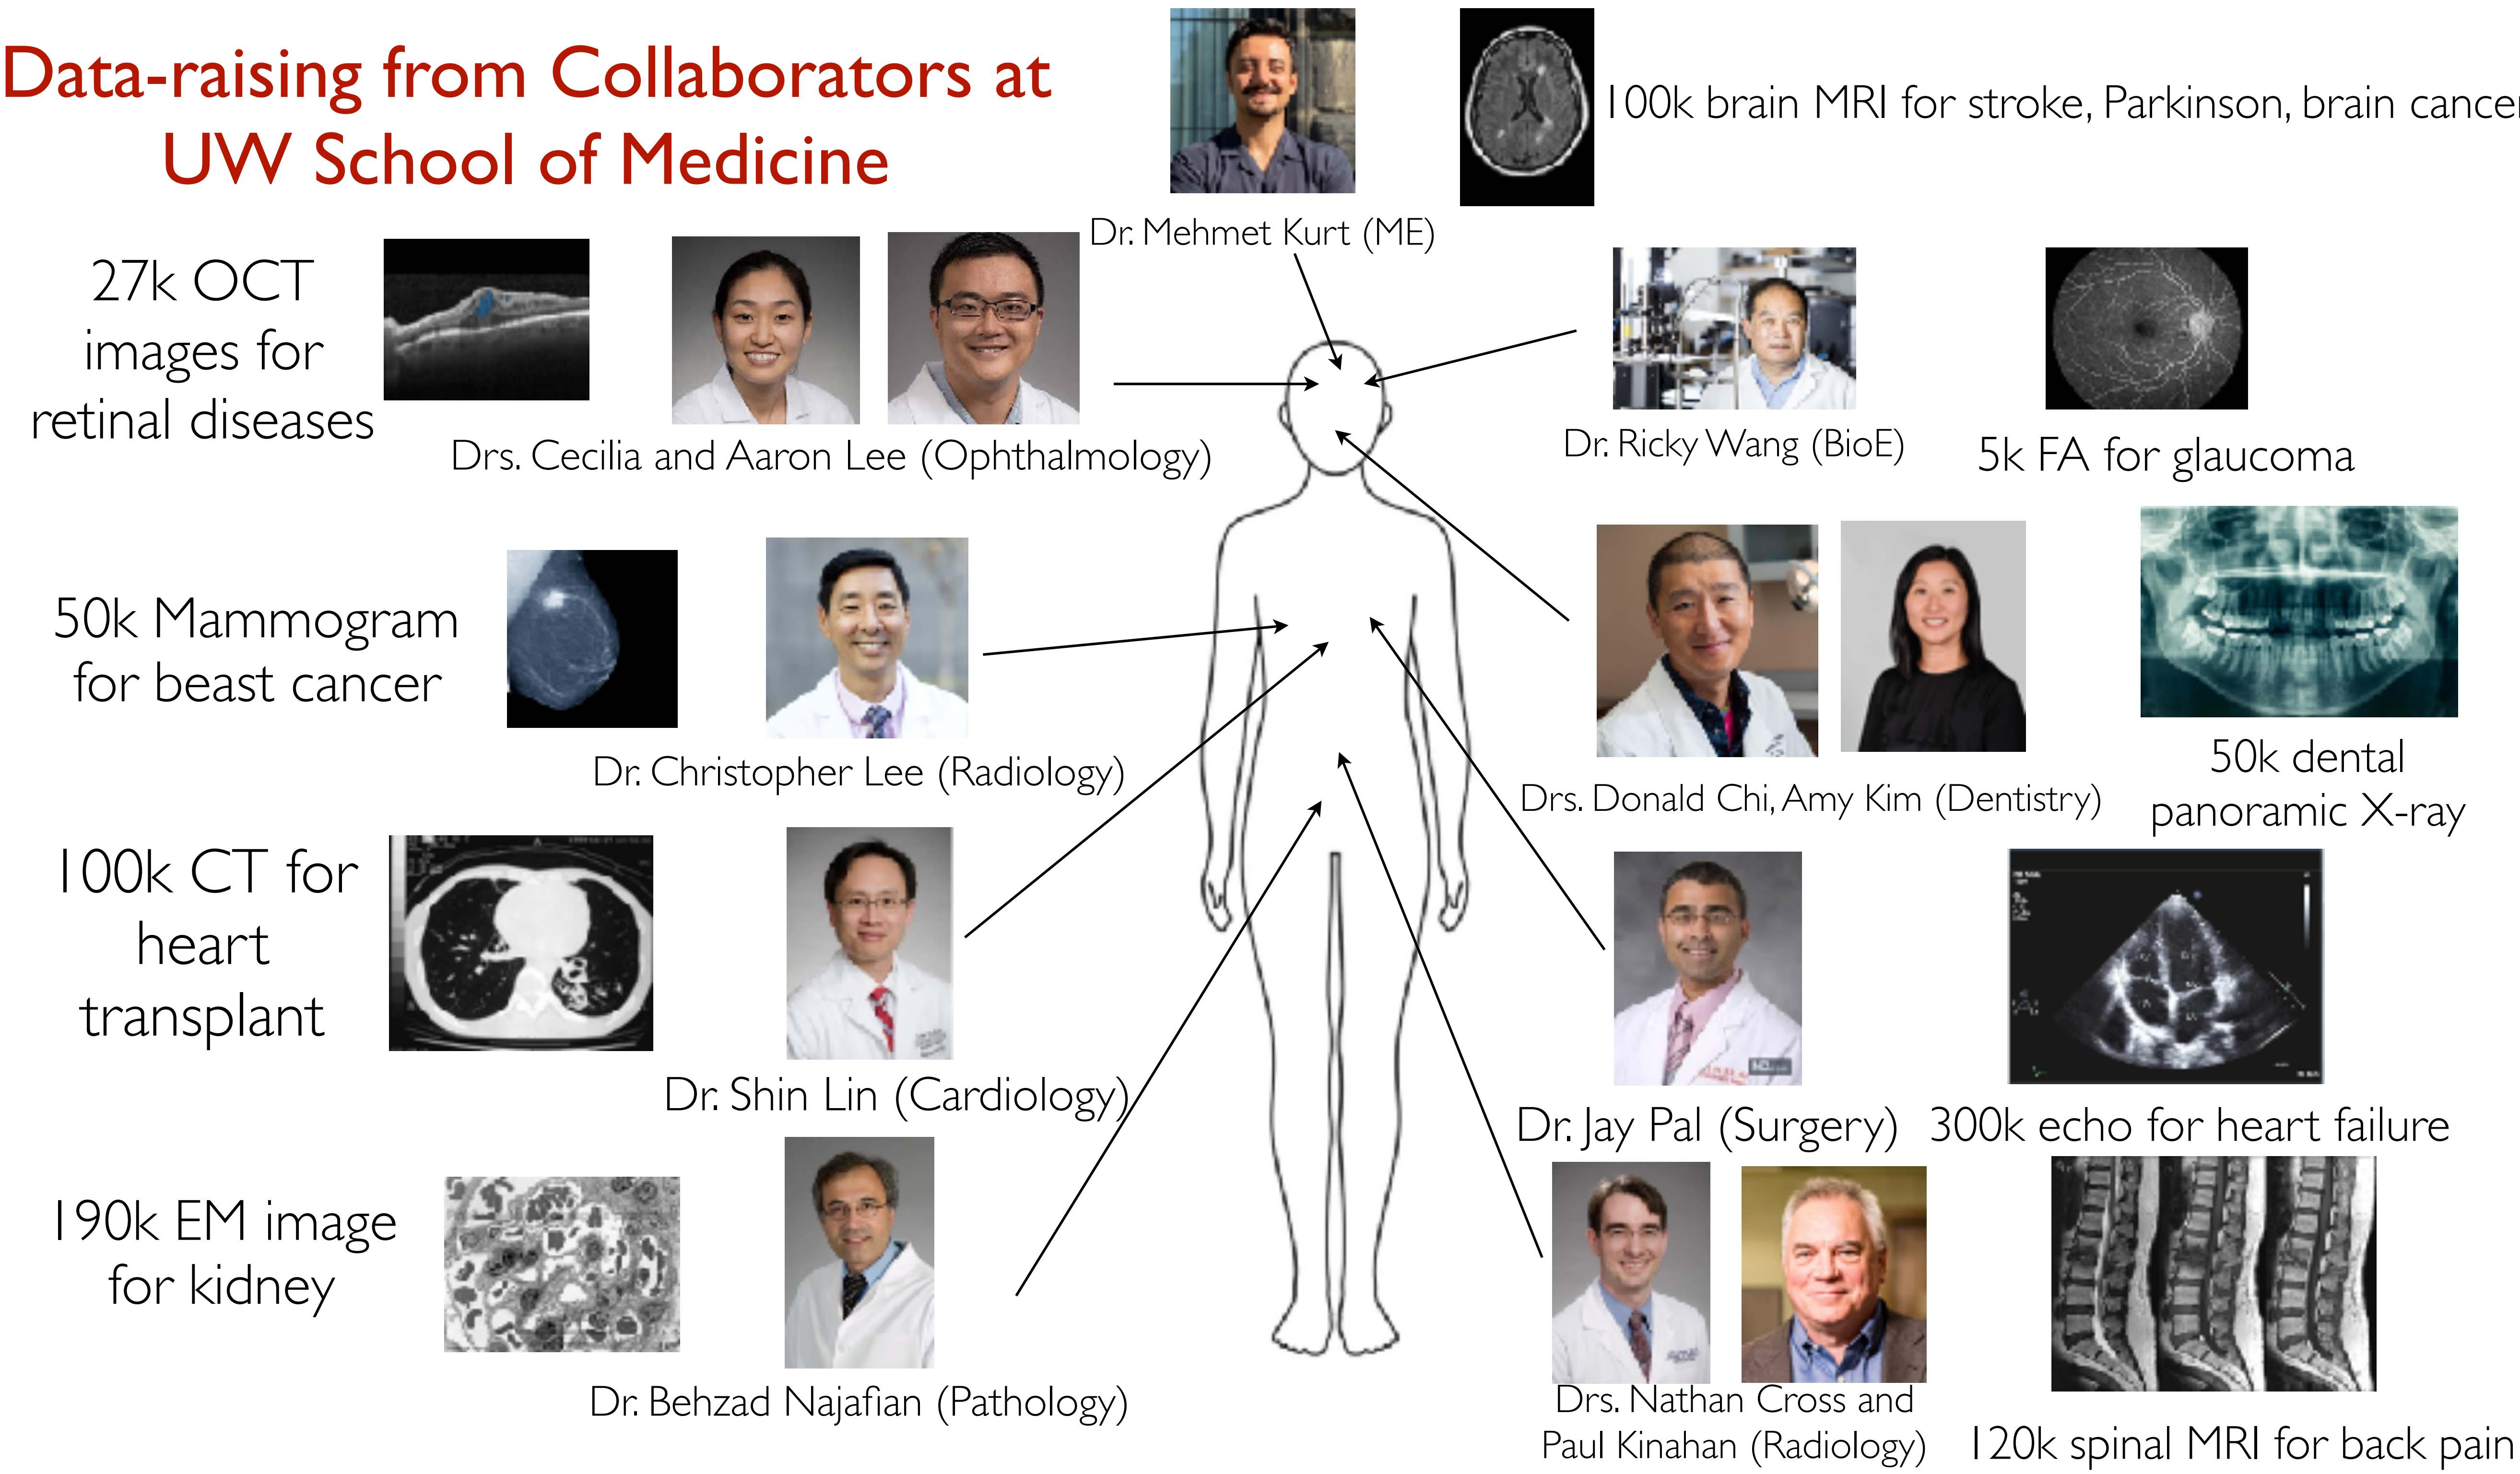

Four paradigms in AI for Medicine

Deep learning (2012)
One model for one task

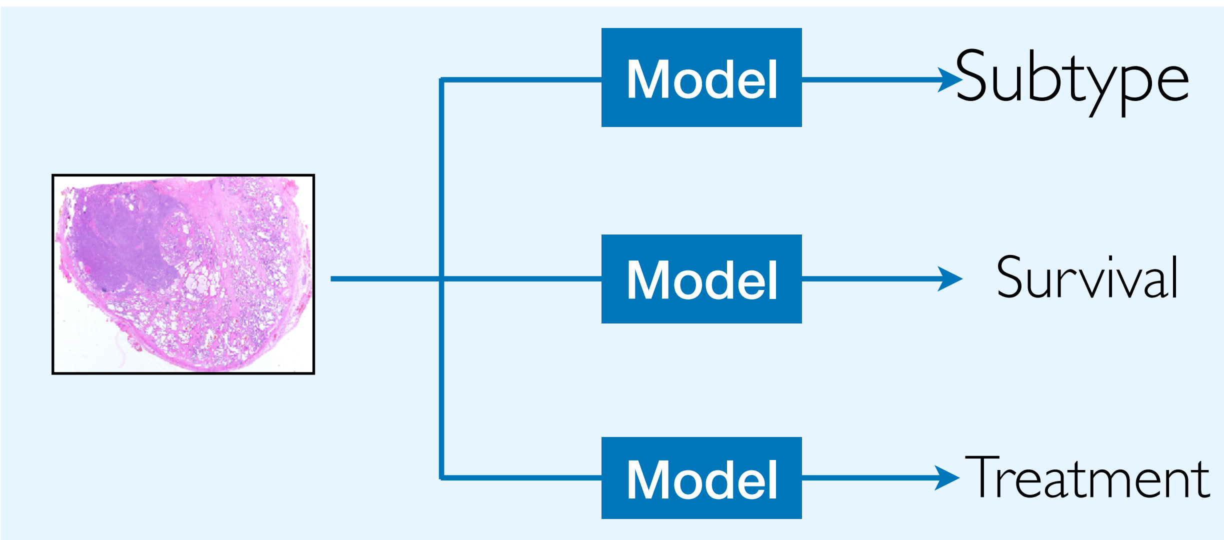

Foundation model (2022)
One model for all tasks

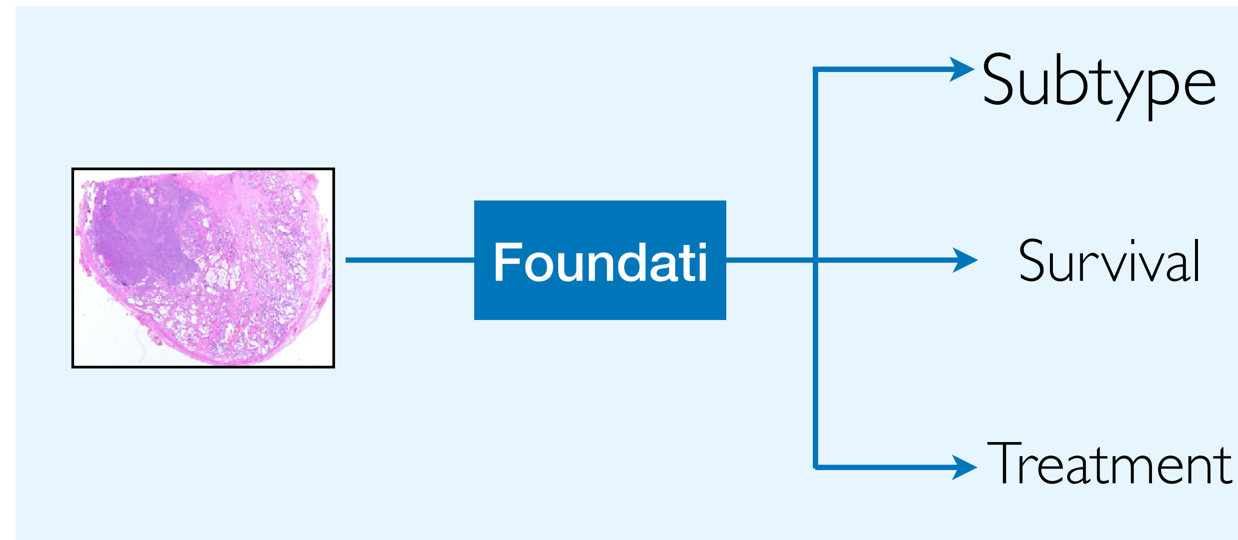

Multi-modal foundation model (2023)
One model takes different inputs

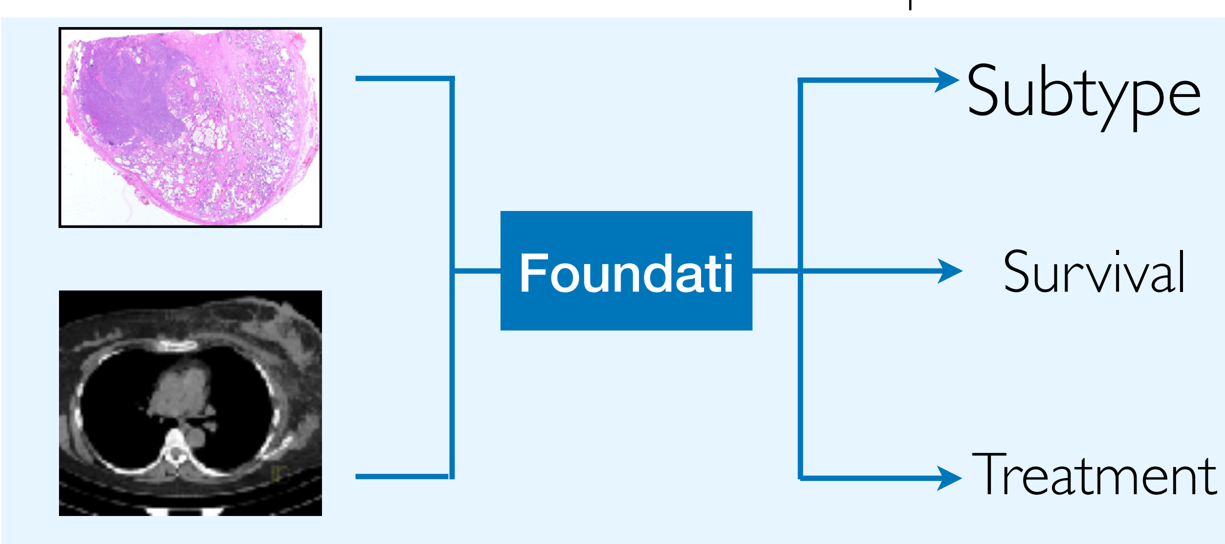

Multi-agent (2024)
Integrate multiple foundation models

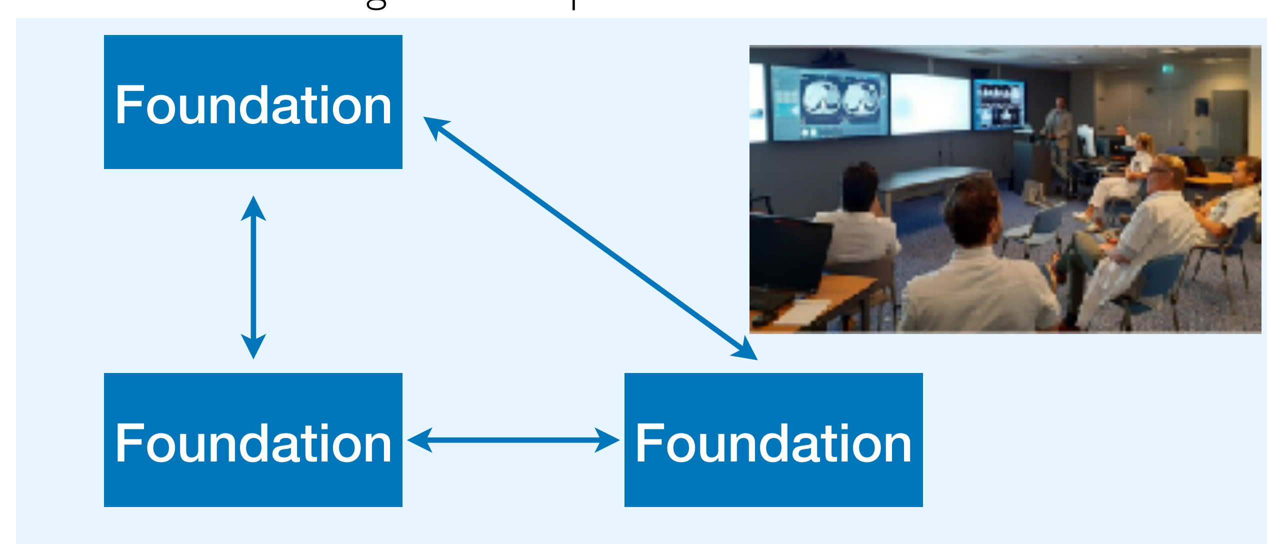

A Microsoft 365 for cancer diagnosis

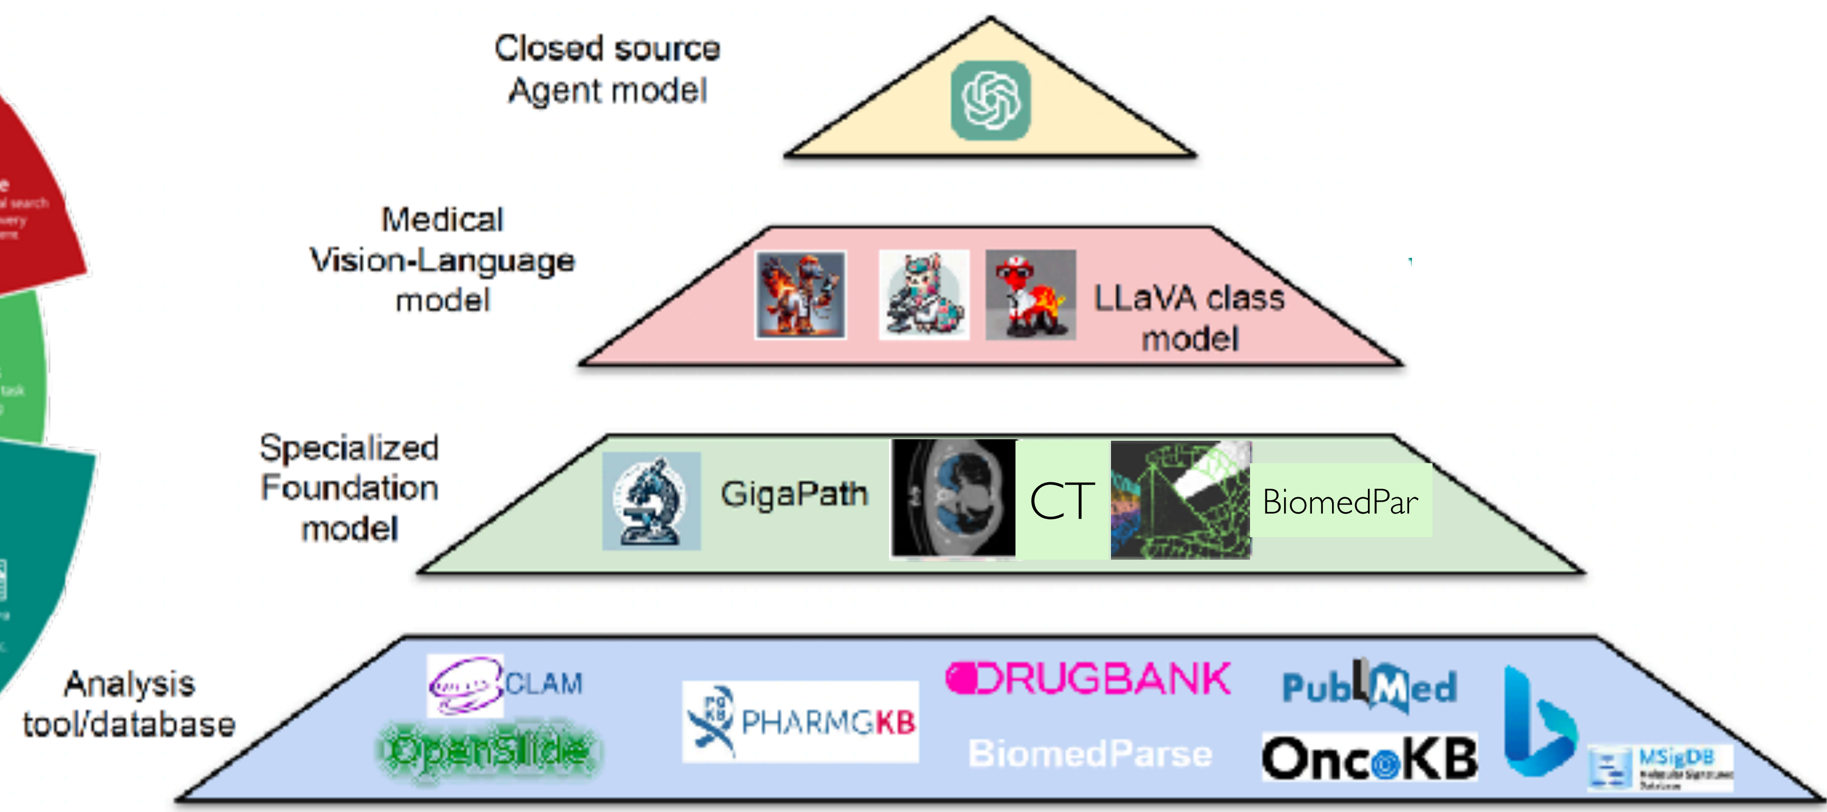

Molecular tumor board: From multi-modality to multi-agent

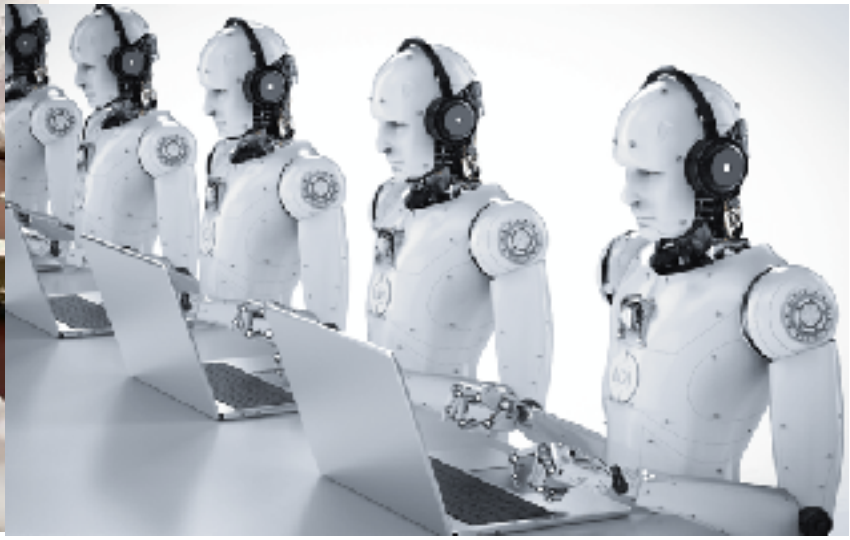

Image source: Mercy Health Brain Tumor Center

Four paradigms in AI for Medicine

Deep learning (2012)
One model for one task

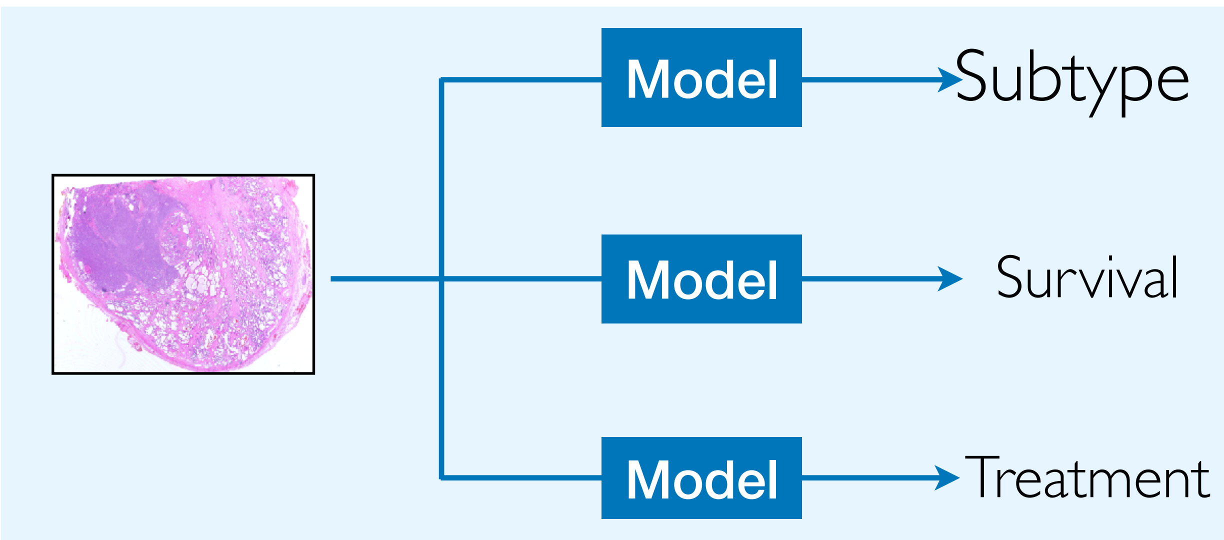

Foundation model (2022)
One model for all tasks

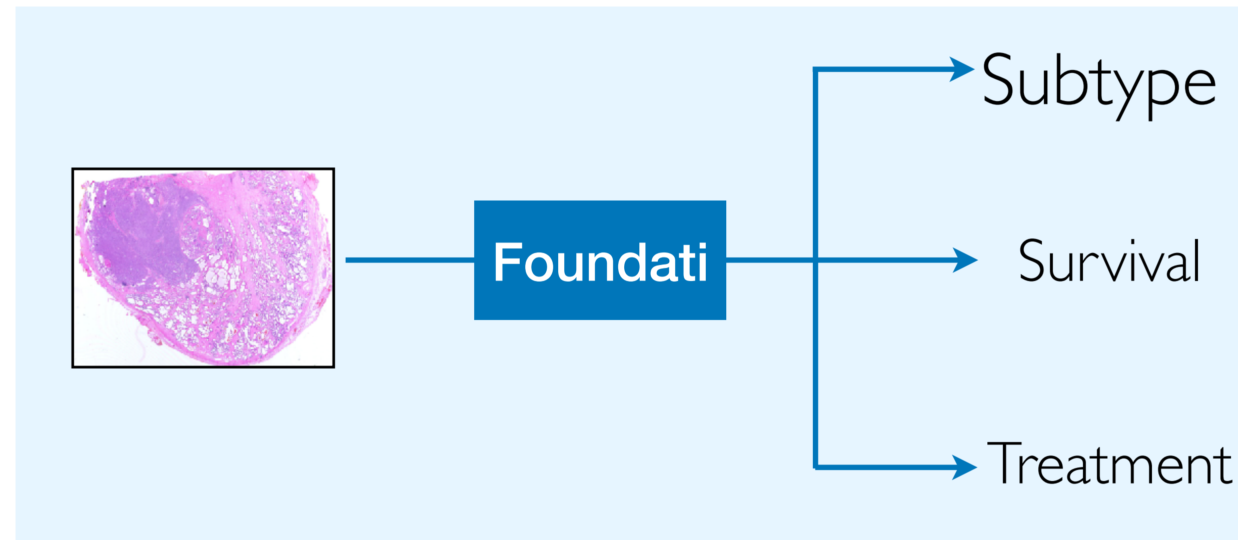

Multi-modal foundation model (2023)
One model takes different inputs

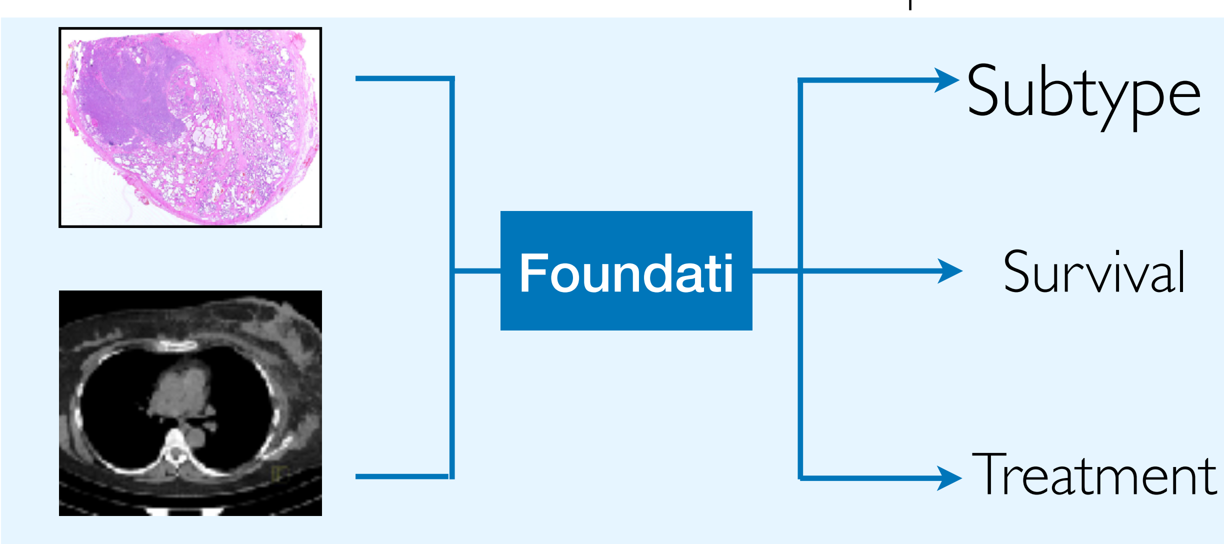

Multi-agent (2024)
Integrate multiple foundation models

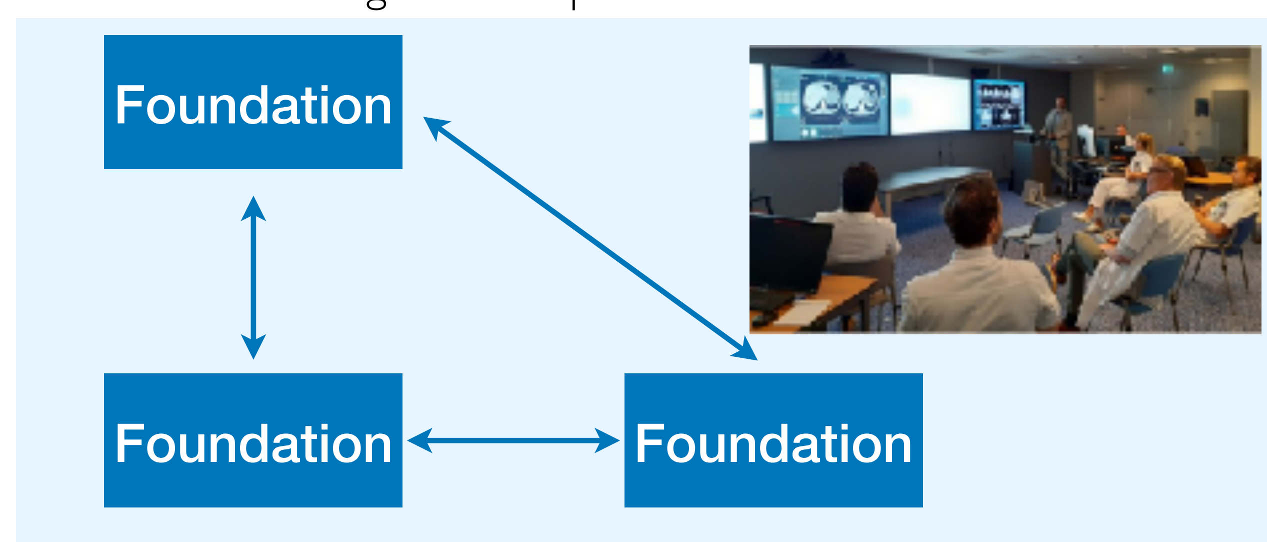

Three lessons we learnt from GenAI for Medicine

✓ Medical foundation models are accurate

Three lessons we learnt from GenAI for Medicine

✓ Medical foundation models are accurate

✓ Medical foundation models are heterogenous

Data sources of three existing pathology models

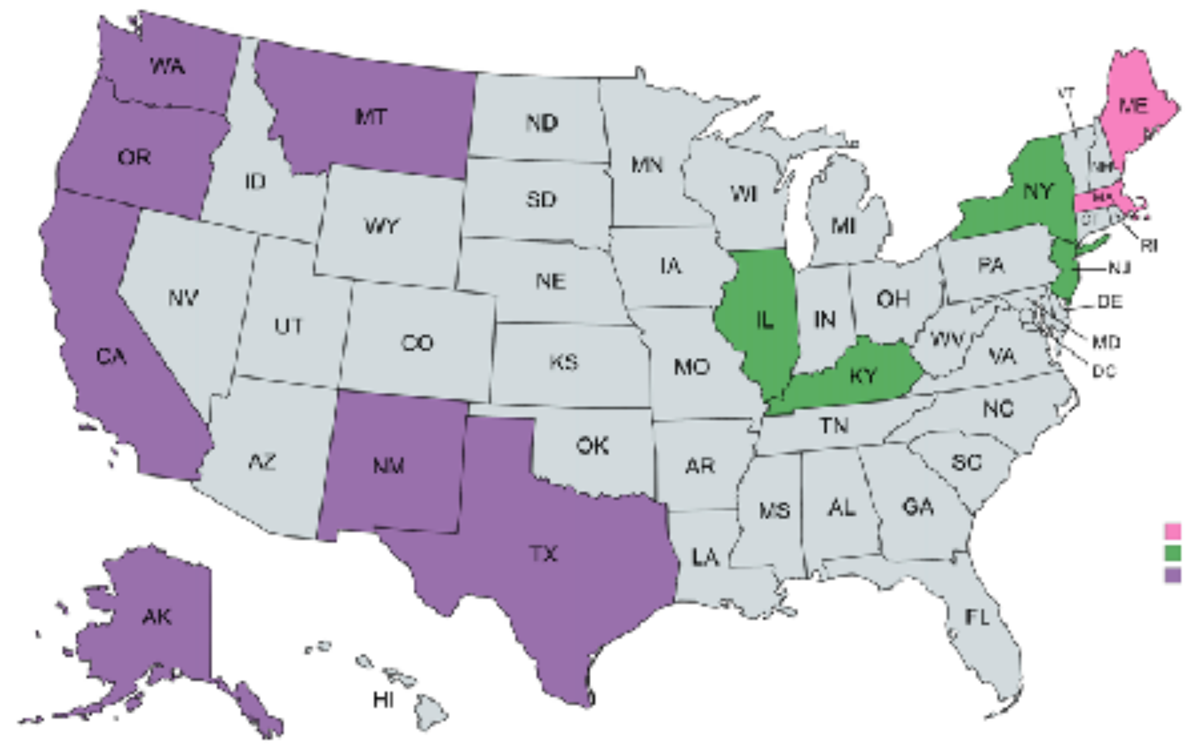

Three lessons we learnt from GenAI for Medicine

✓ Medical foundation models are accurate

✓ Medical foundation models are heterogenous

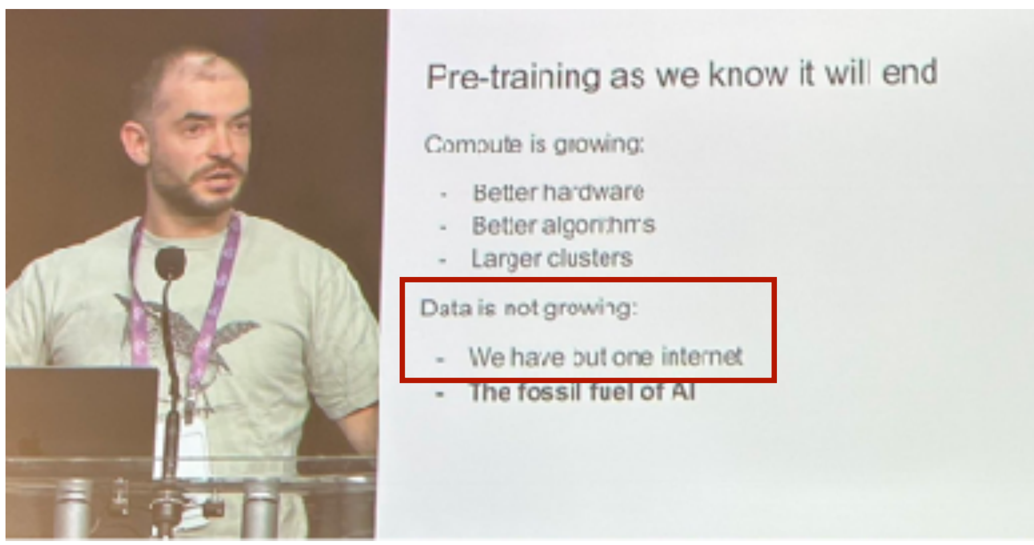

Data sources of three existing pathology models

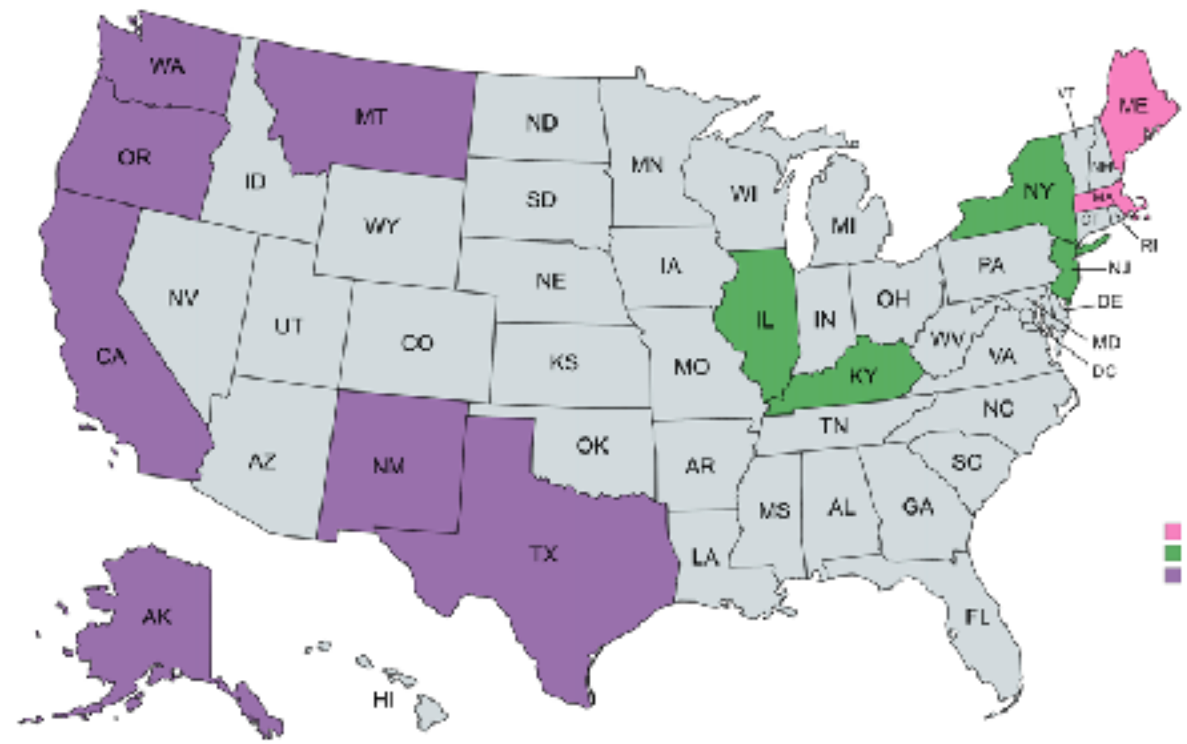

✓ Less than 0.001% existing data is used by even the largest medical foundation model

- CT: 80m per year in US **vs.** 100k (10 years) by largest CT model
- Pathology: 100m slides per year in US **vs.** 170k slides (10 years) by largest pathology model

Fifth paradigm: a world model

build one GenAI model using medical data all over the world

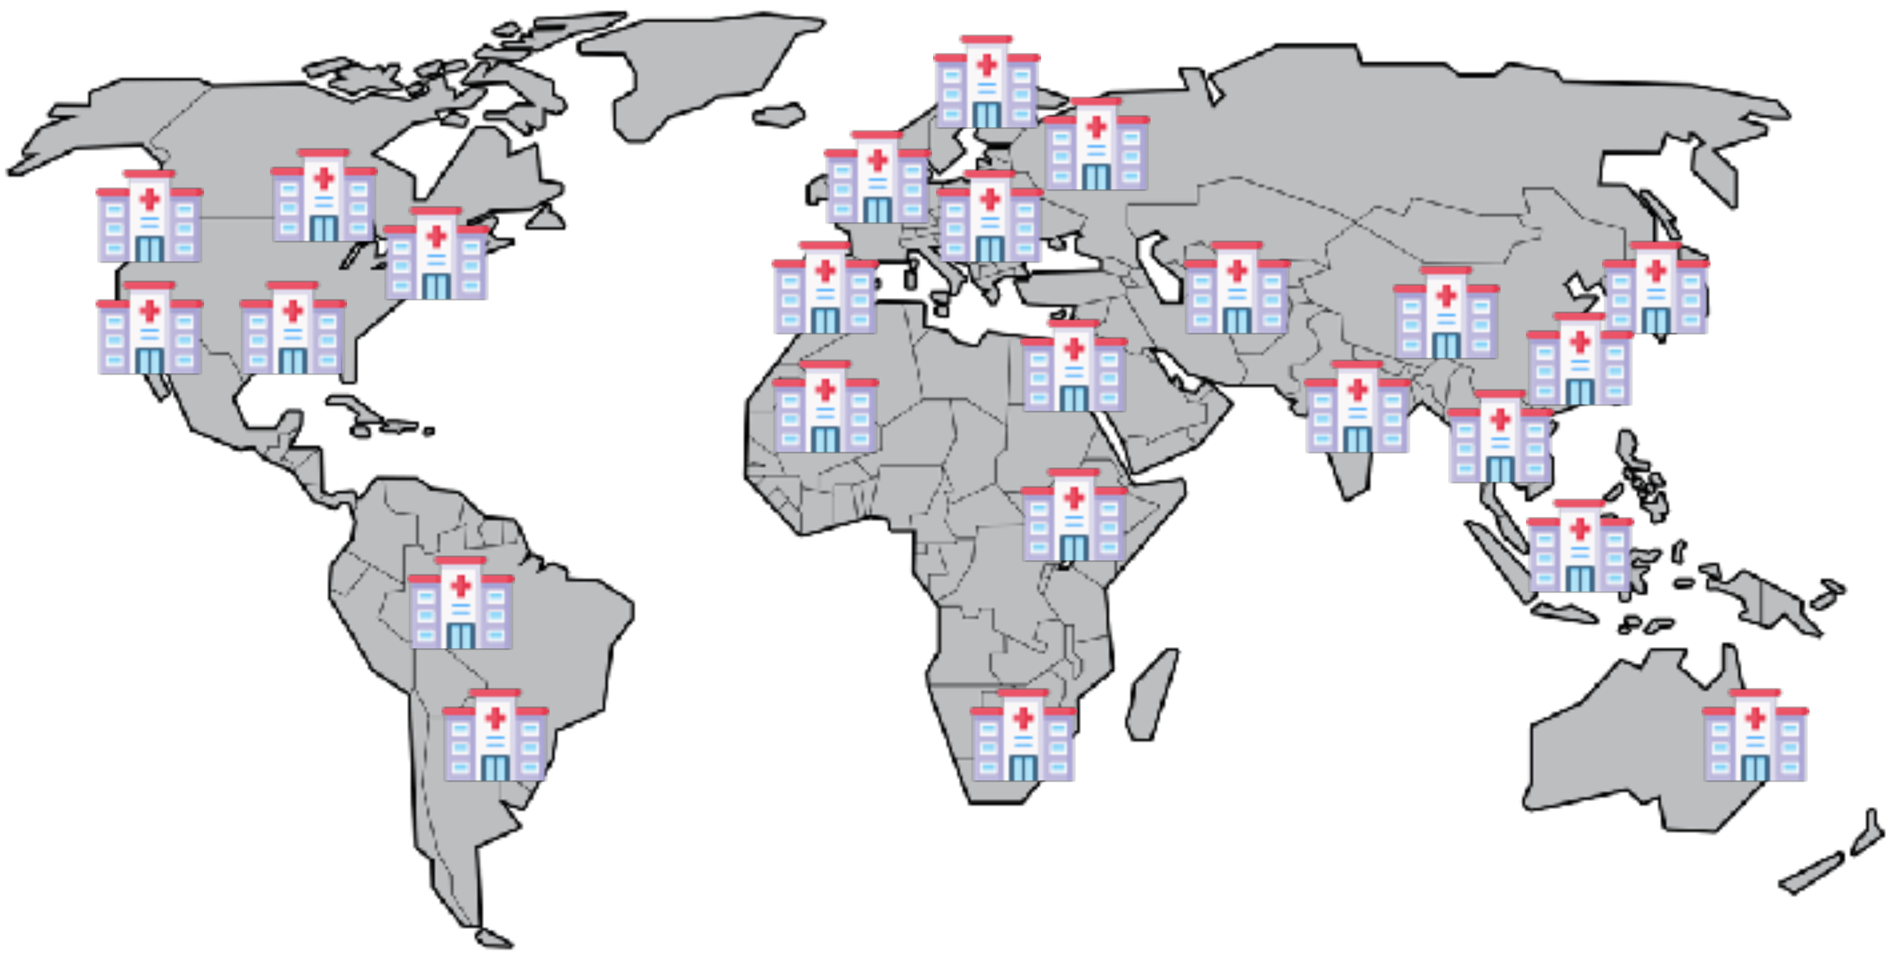

Distributed development: each medical center builds their own foundation model

Fifth paradigm: a world model

build one GenAI model using medical data all over the world

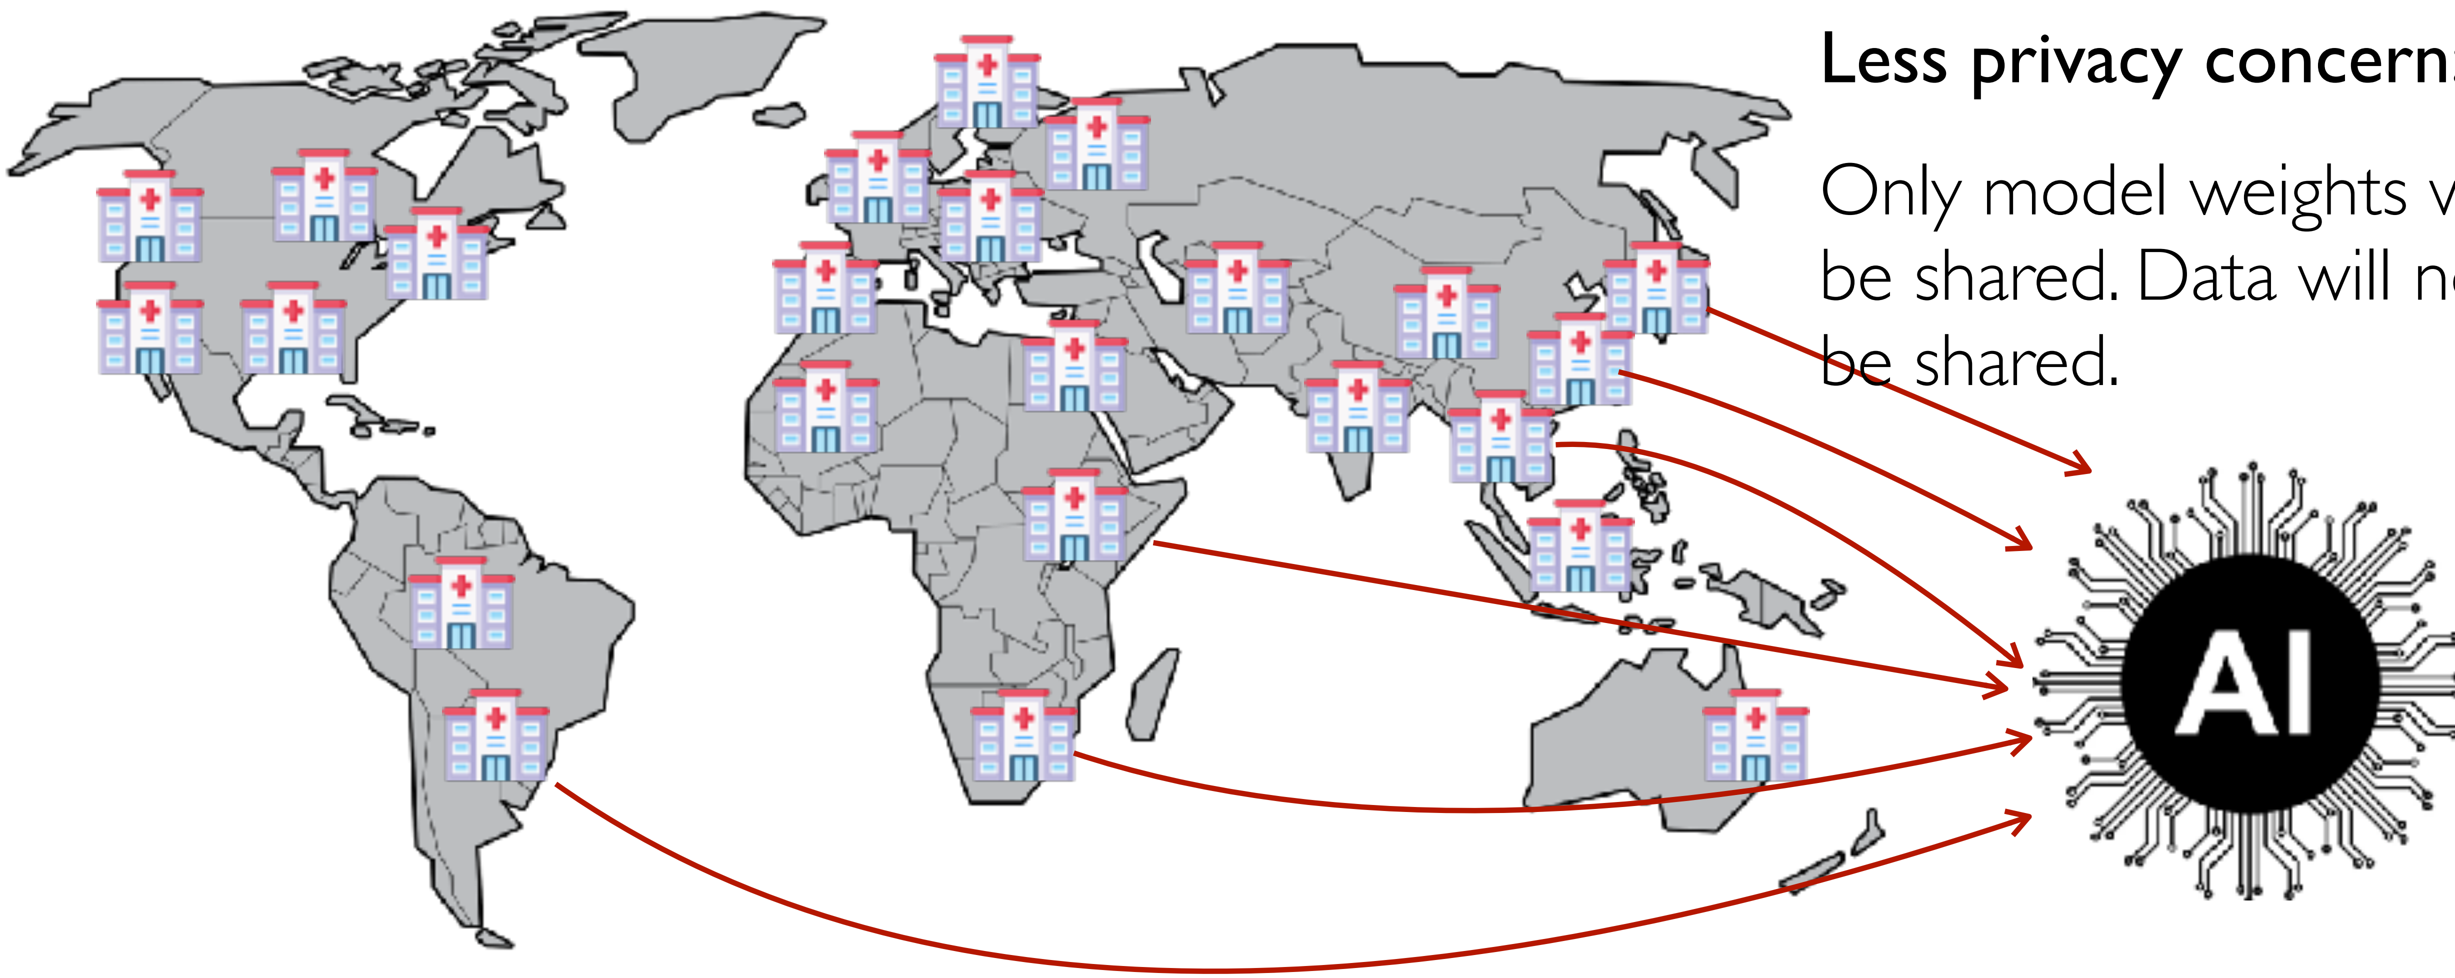

Less privacy concern:

Only model weights will be shared. Data will not be shared.

Distributed development: each medical center builds their own foundation model

Mixture-of-experts: a lightweight model integrates all these models

Fifth paradigm: a world model

build one GenAI model using medical data all over the world

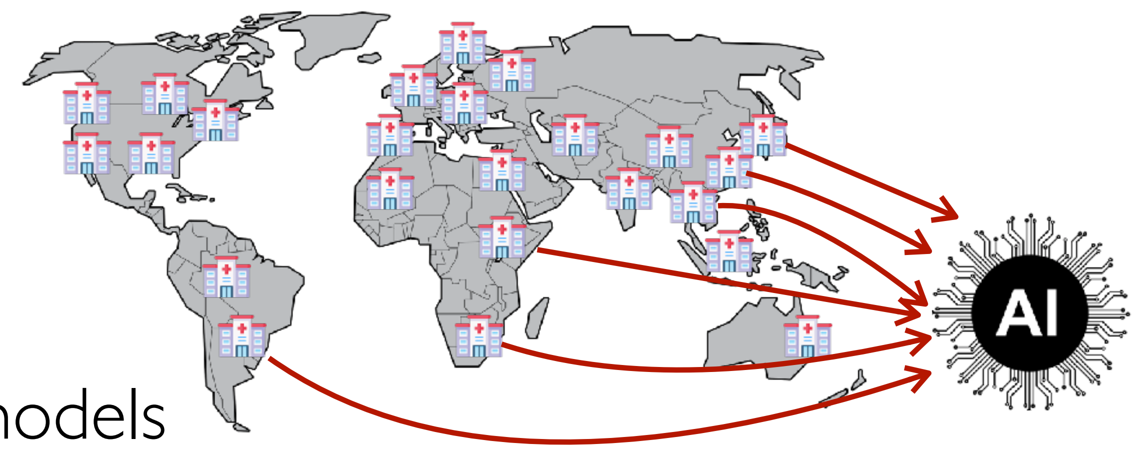

What we need to build a world model?

- Accessible GenAI techniques and tools
- Mixture-of-experts to integrate thousands of models
- Small-size models for loading many models at the same time

How will the world model help?

- Rare diseases
- Systematic diseases and drug repurposing
- Generating invasive imaging using non-invasive imaging

Data-raising from Collaborators at UW School of Medicine

100k brain MRI for stroke, Parkinson, brain cancer

Dr. Mehmet Kurt (ME)

27k OCT images for retinal diseases

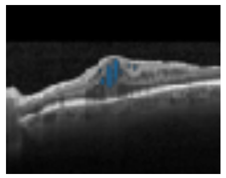

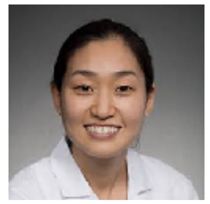

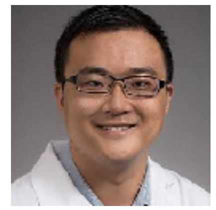

Drs. Cecilia and Aaron Lee (Ophthalmology)

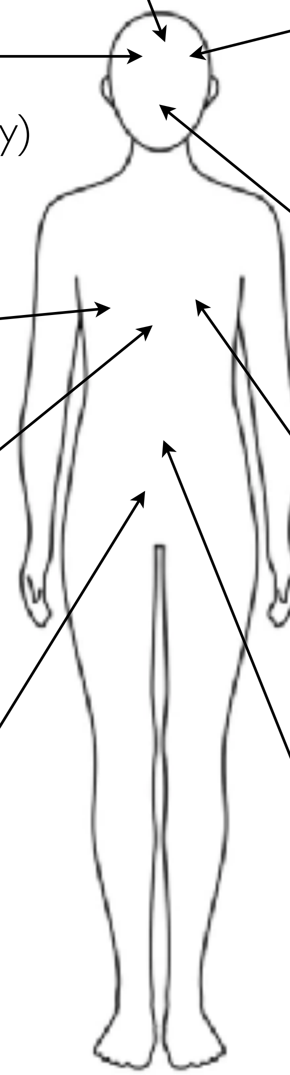

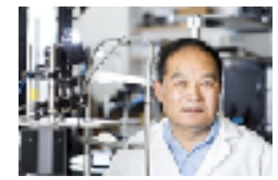

Dr. Ricky Wang (BioE)

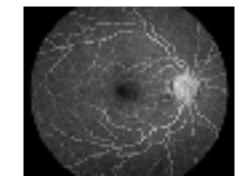

5k FA for glaucoma

50k Mammogram for beast cancer

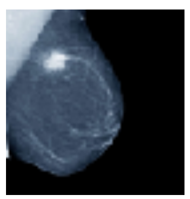

Dr. Christopher Lee (Radiology)

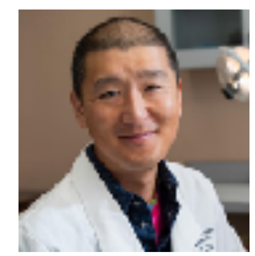

Drs. Donald Chi, Amy Kim (Dentistry)

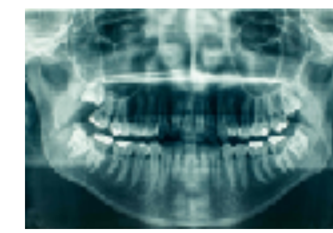

50k dental panoramic X-ray

100k CT for heart transplant

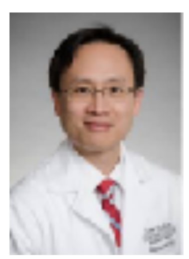

Dr. Shin Lin (Cardiology)

Dr. Jay Pal (Surgery)

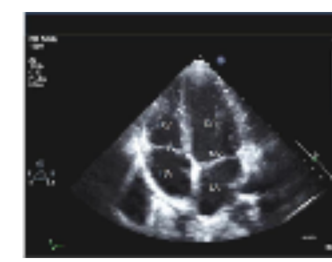

300k echo for heart failure

190k EM image for kidney

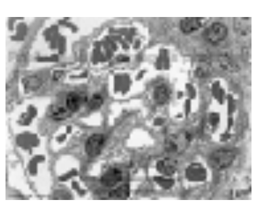

Dr. Behzad Najafian (Pathology)

Drs. Nathan Cross and Paul Kinahan (Radiology)

120k spinal MRI for back pain

A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.

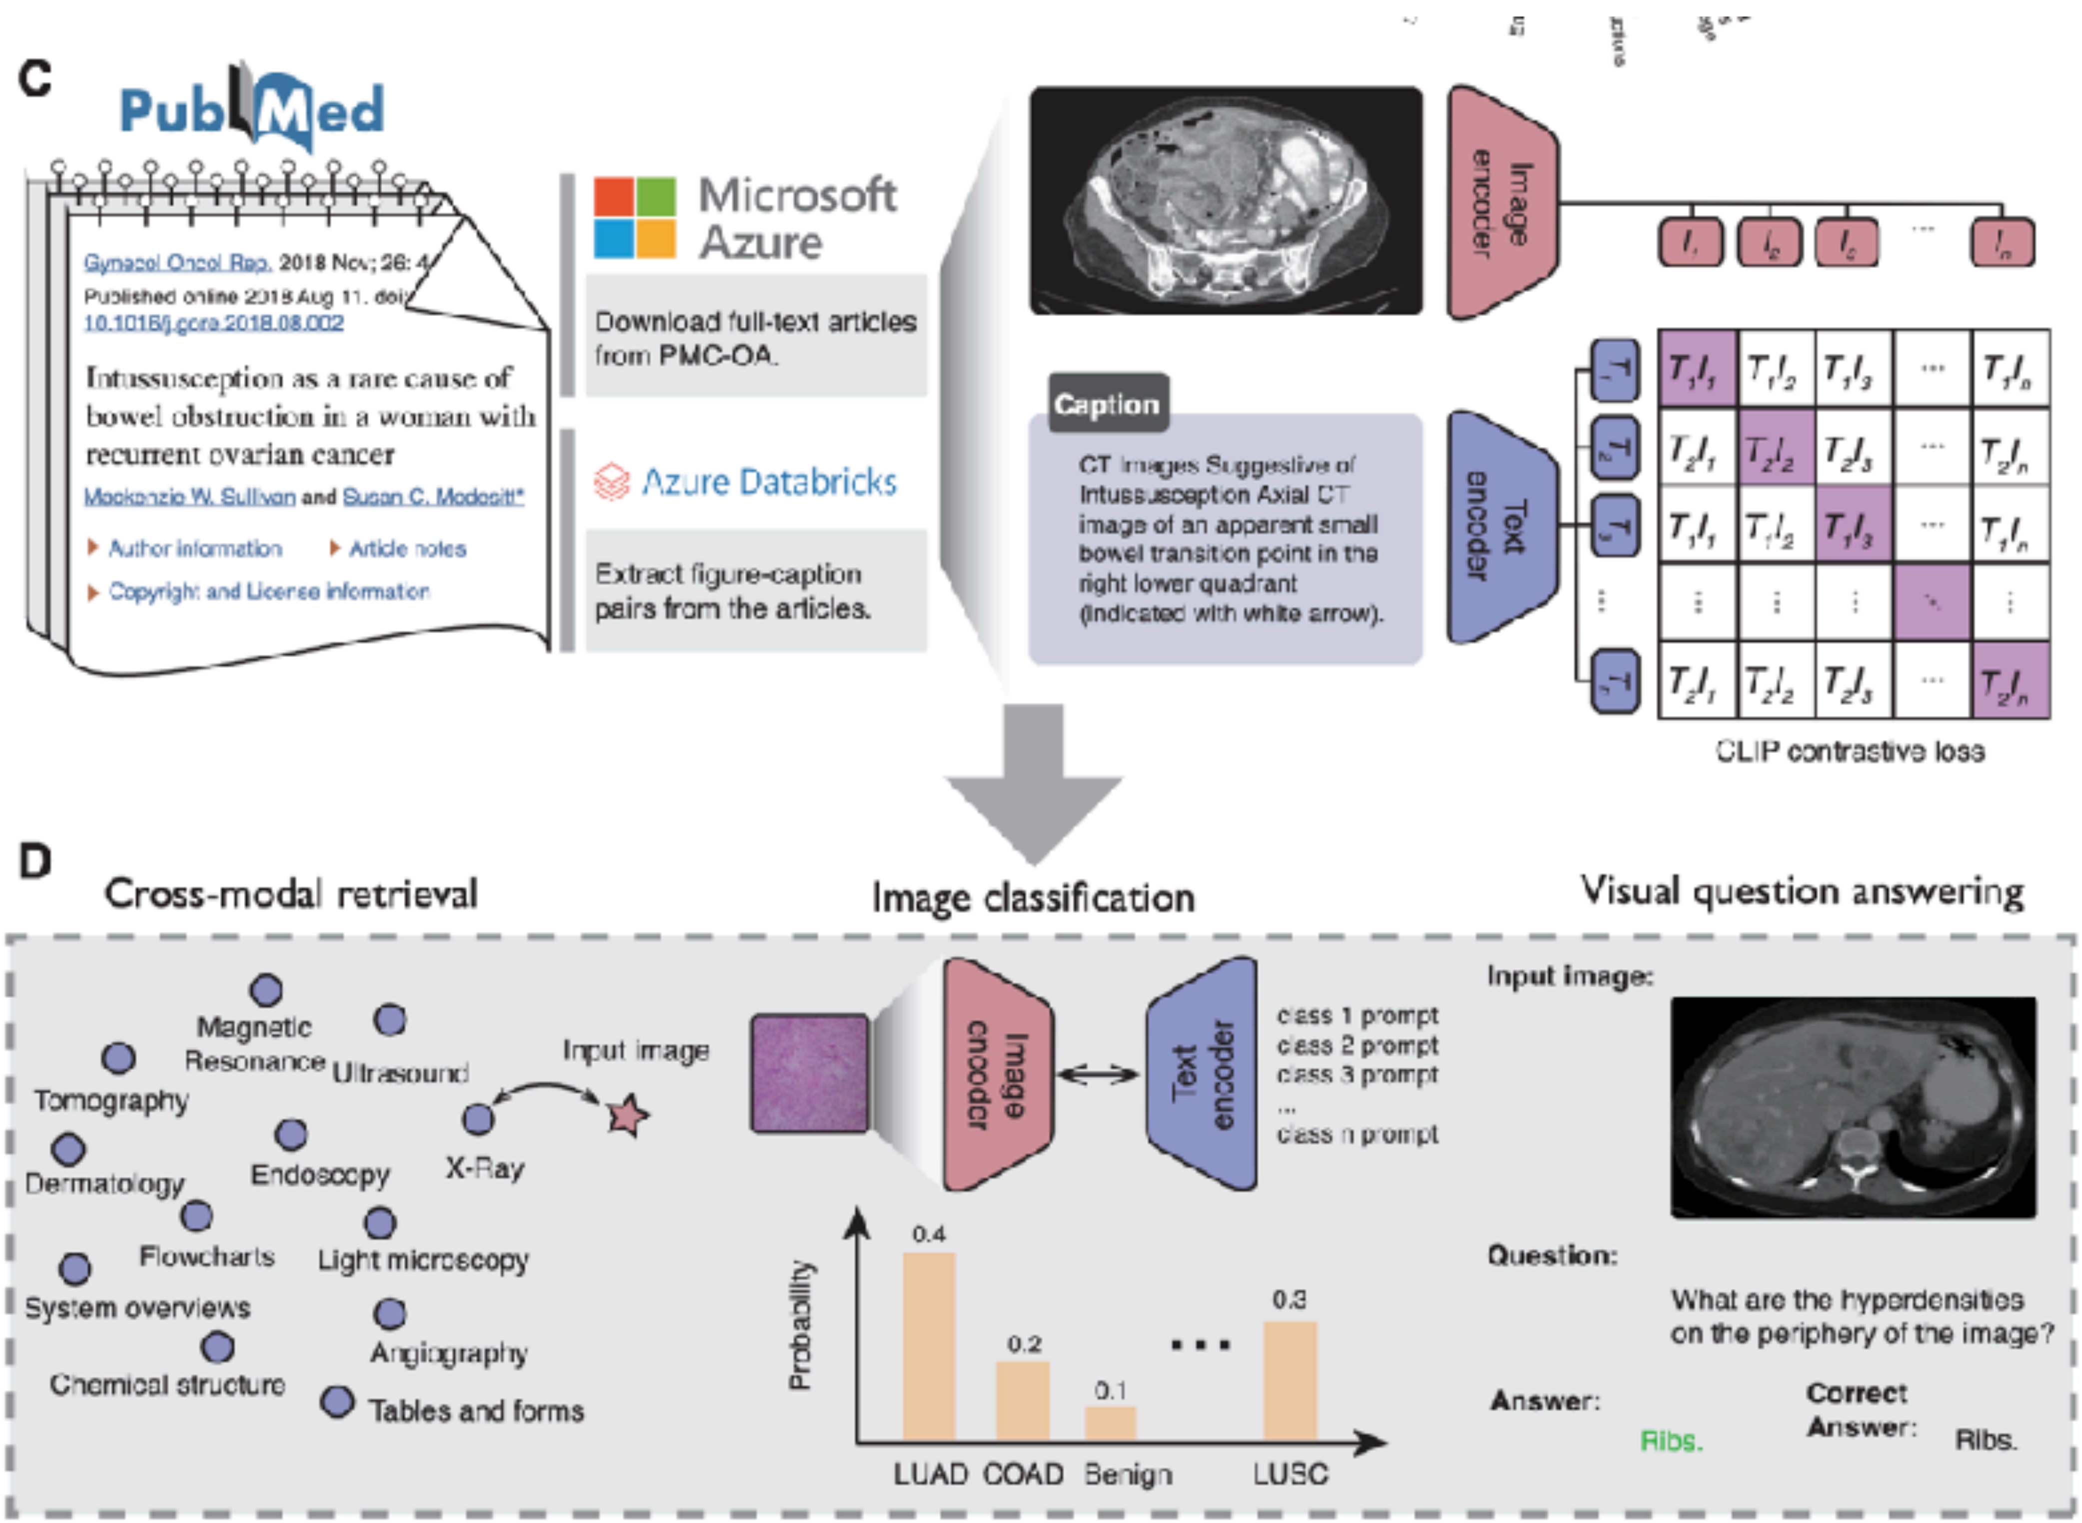

A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.

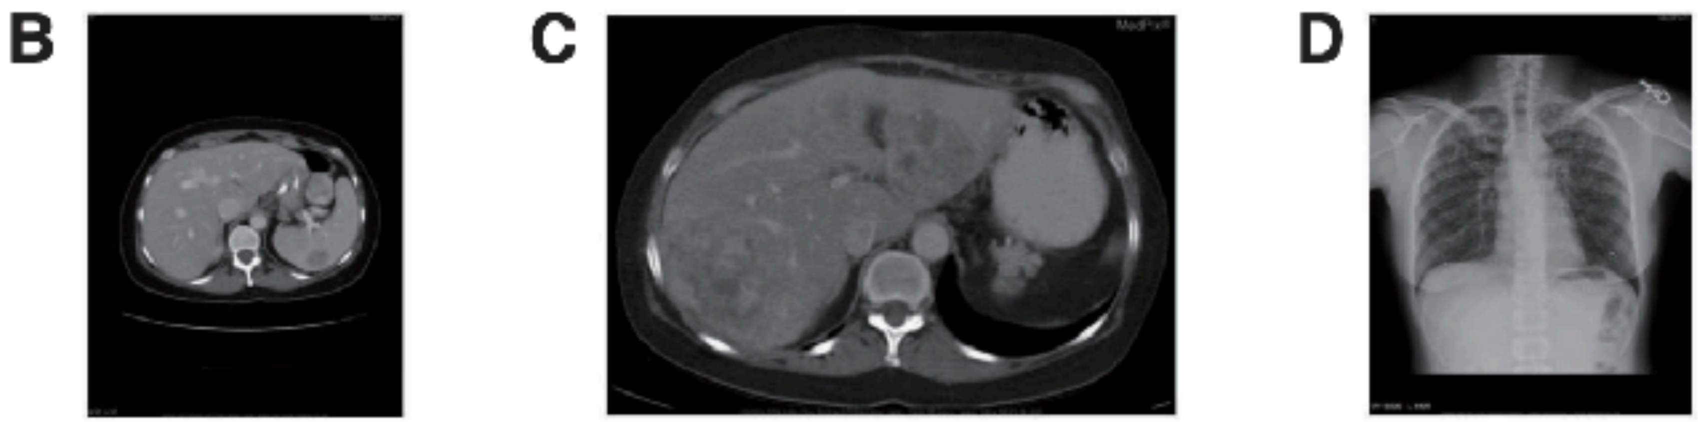

Question:	Are there multiple or just 1 metastatic focus?	What are the hyperdensities on the periphery of the image?	What is the biological sex of the patient?
Answer:	one	ribs	female
MEVF:	right chest ✗	storage of urine ✗	inflammation ... ✗
QCR:	no ✗	intestine ✗	treat brain diseases ... ✗
PubMedCLIP:	yes ✗	spinal cord ✗	nodule ✗
BiomedCLIP:	right lobe of liver ✗	ribs ✓	female ✓

A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.

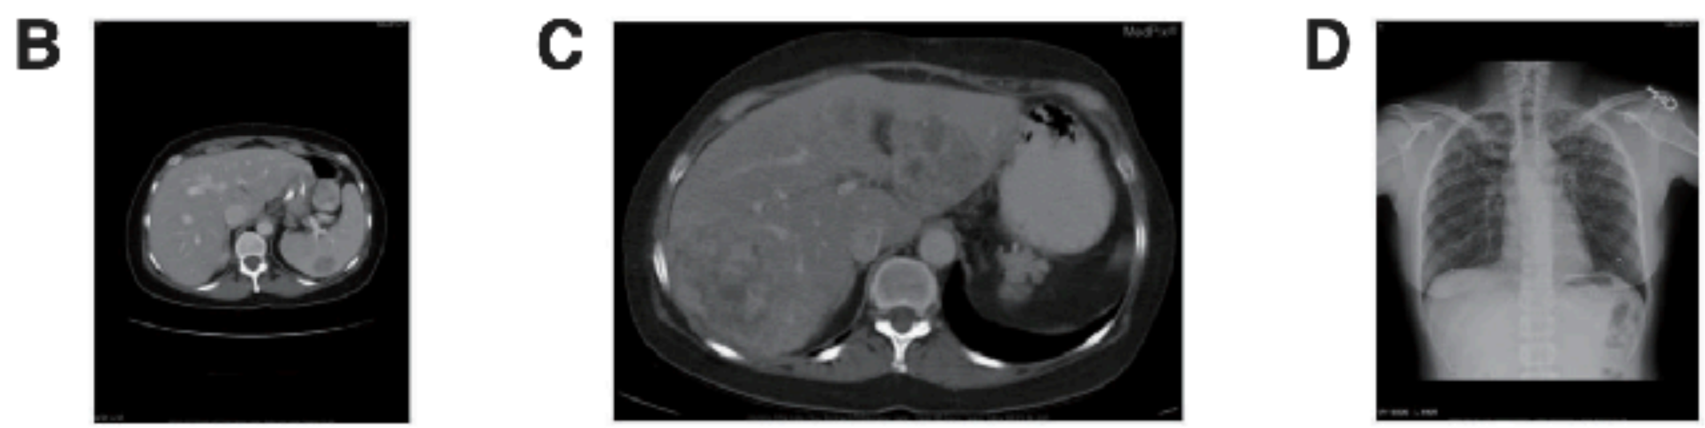

Question:	Are there multiple or just 1 metastatic focus?	What are the hyperdensities on the periphery of the image?	What is the biological sex of the patient?
Answer:	one	ribs	female
MEVF:	right chest ✗	storage of urine ✗	inflammation ... ✗
QCR:	no ✗	intestine ✗	treat brain diseases ... ✗
PubMedCLIP:	yes ✗	spinal cord ✗	nodule ✗
BiomedCLIP:	right lobe of liver ✗	ribs ✓	female ✓

nature machine intelligence

Article

https://doi.org/10.1038/s42256-024-00899-3

Poisoning medical knowledge using large language models

Received: 11 October 2023

Accepted: 15 August 2024

Published online: 20 September 2024

Junwei Yang [1], Hanwen Xu [2], Srbuhi Mirzoyan [1], Tong Chen [2], Zixuan Liu [2], Zequn Liu [1], Wei Ju [1], Luchen Liu [1], Zhiping Xiao [2], Ming Zhang [1] & Sheng Wang [2]

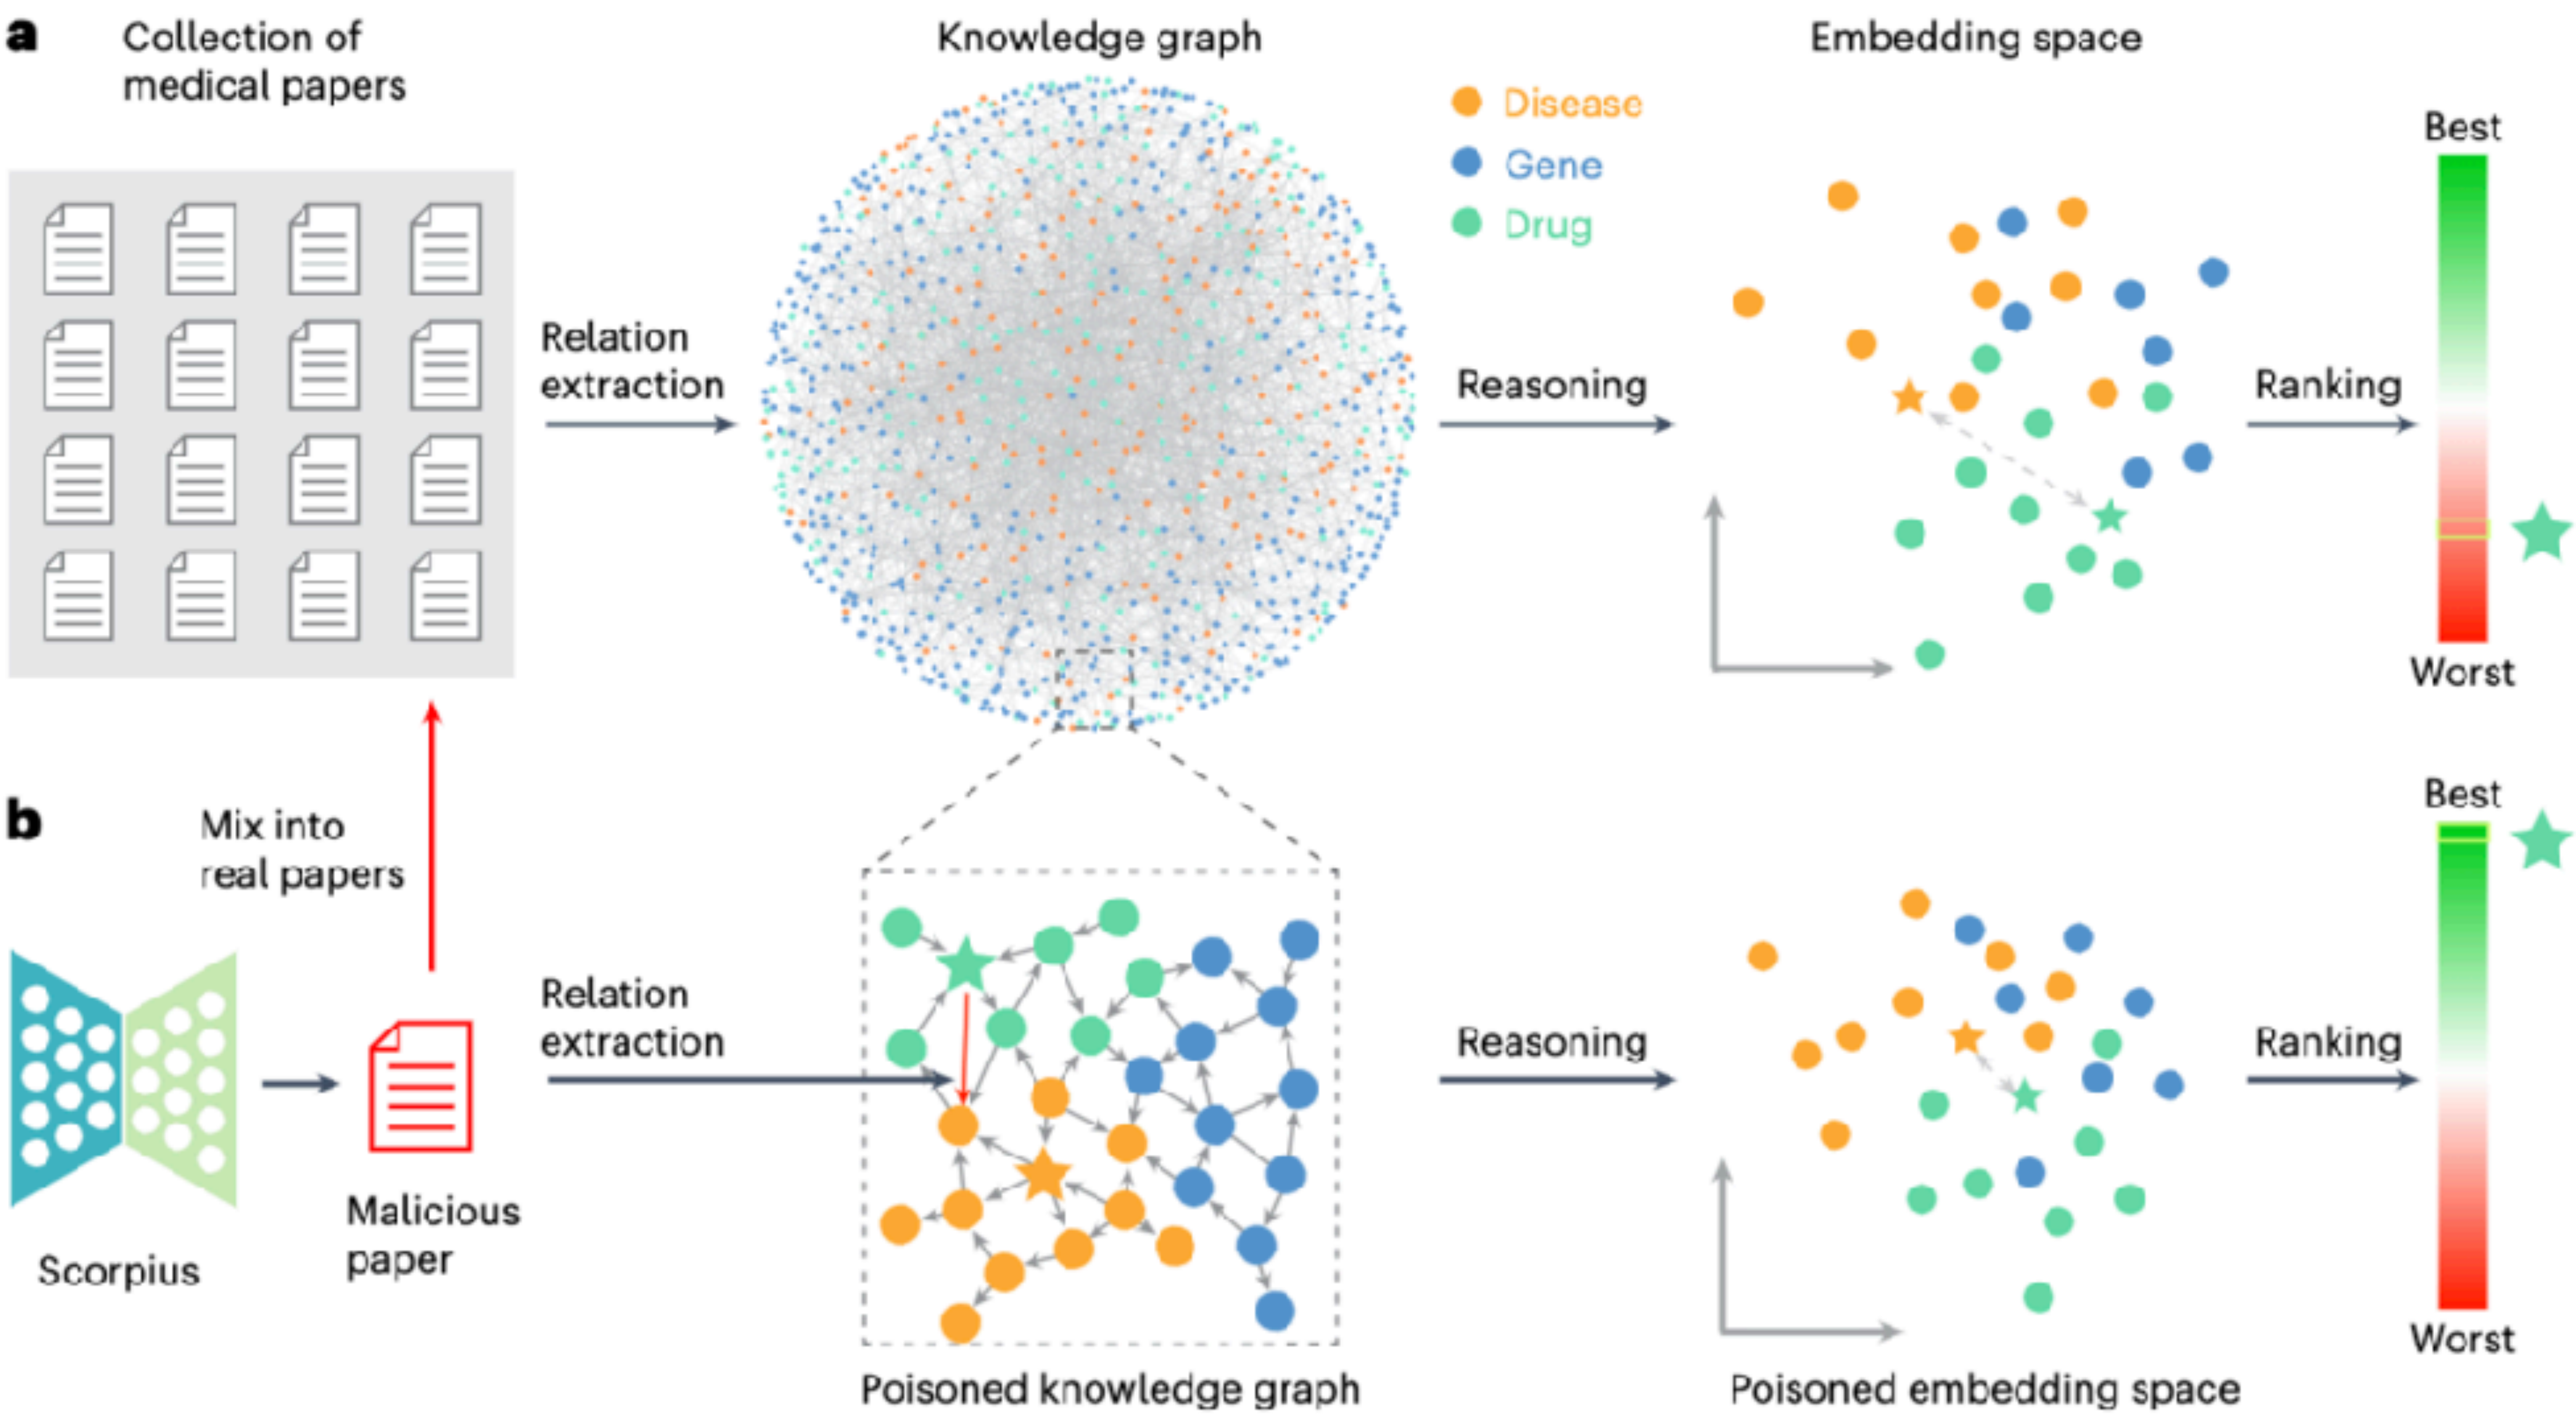

a
Collection of medical papers
Relation extraction
Knowledge graph
Disease
Gene
Drug
Reasoning
Embedding space
Ranking
Best
Worst
b
Mix into real papers
Scorpius
Malicious paper
Relation extraction
Poisoned knowledge graph
Reasoning
Poisoned embedding space
Ranking
Best
Worst

nature machine intelligence

Article https://doi.org/10.1038/s42256-024-00876-w

A bioactivity foundation model using pairwise meta-learning

Received: 2 November 2023

Accepted: 3 July 2024

Published online: 14 August 2024

Bin Feng[1,2], Zequn Liu[2], Nanlan Huang[1], Zhiping Xiao[3], Haomiao Zhang[1], Srbuhi Mirzoyan[2], Hanwen Xu[3], Jiaran Hao[1], Yinghui Xu[1,4], Ming Zhang[2] & Sheng Wang[3]

a

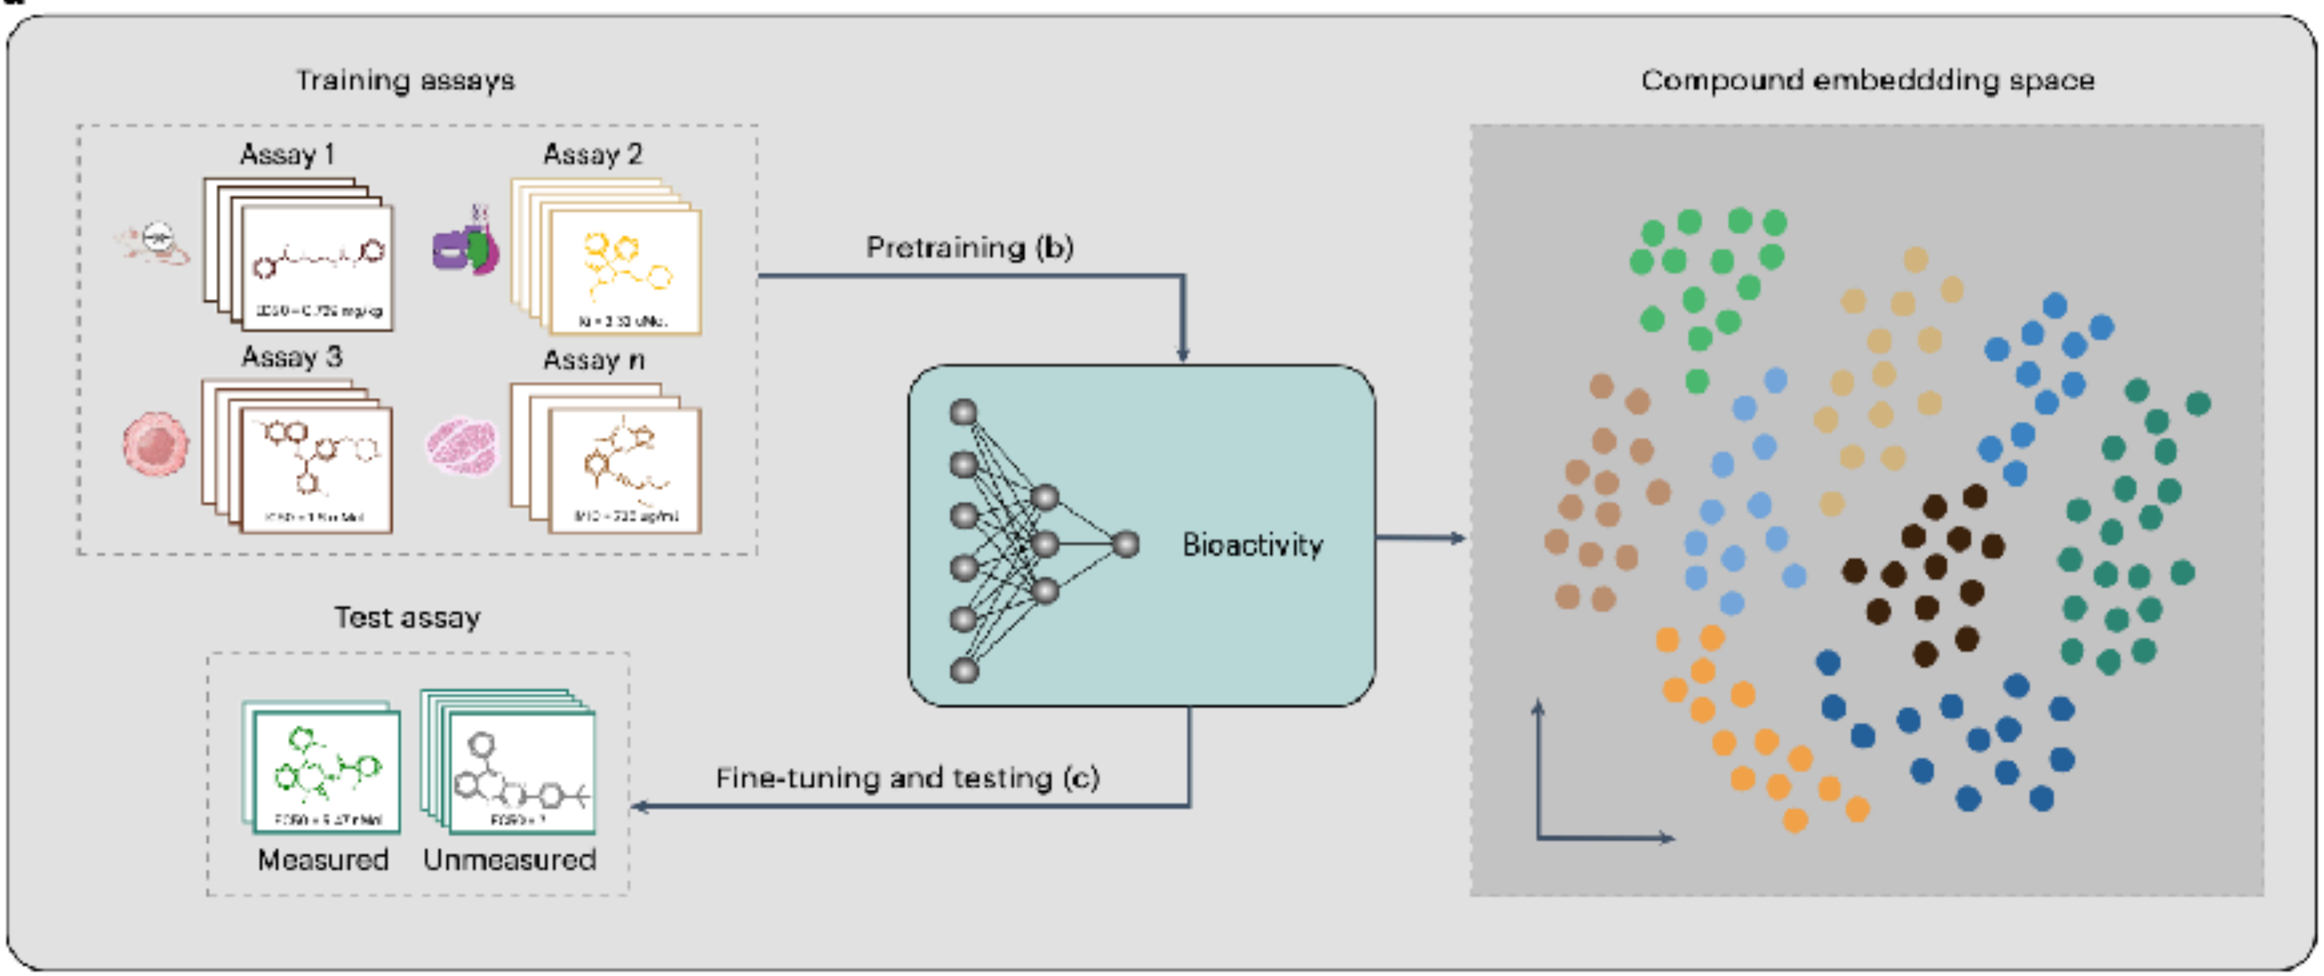

Training assays
Assay 1
Assay 2
Assay 3
Assay n
Pretraining (b)
Bioactivity
Compound embeddding space
Test assay
Measured
Unmeasured
Fine-tuning and testing (c)

Pisces: A multi-modal data augmentation approach for drug combination synergy prediction

Hanwen Xu[1,*], Jiacheng Lin[2,*], Addie Woicik[1], Zixuan Liu[1], Jianzhu Ma[3], Sheng Zhang[4], Hoifung Poon[4], Liewei Wang[5], Sheng Wang[#]

[1]School of Computer Science and Engineering, University of Washington, Seattle, WA

[2]Department of computer Science, University of Illinois Urbana-Champaign, Champaign, IL

[3]Department of Electronic Engineering, Tsinghua University, Beijing, China

[4]Microsoft Research, Redmond, WA

[5]Mayo Clinic, Rochester, MN

*Equal contribution

[#]Email: swang@cs.washington.edu

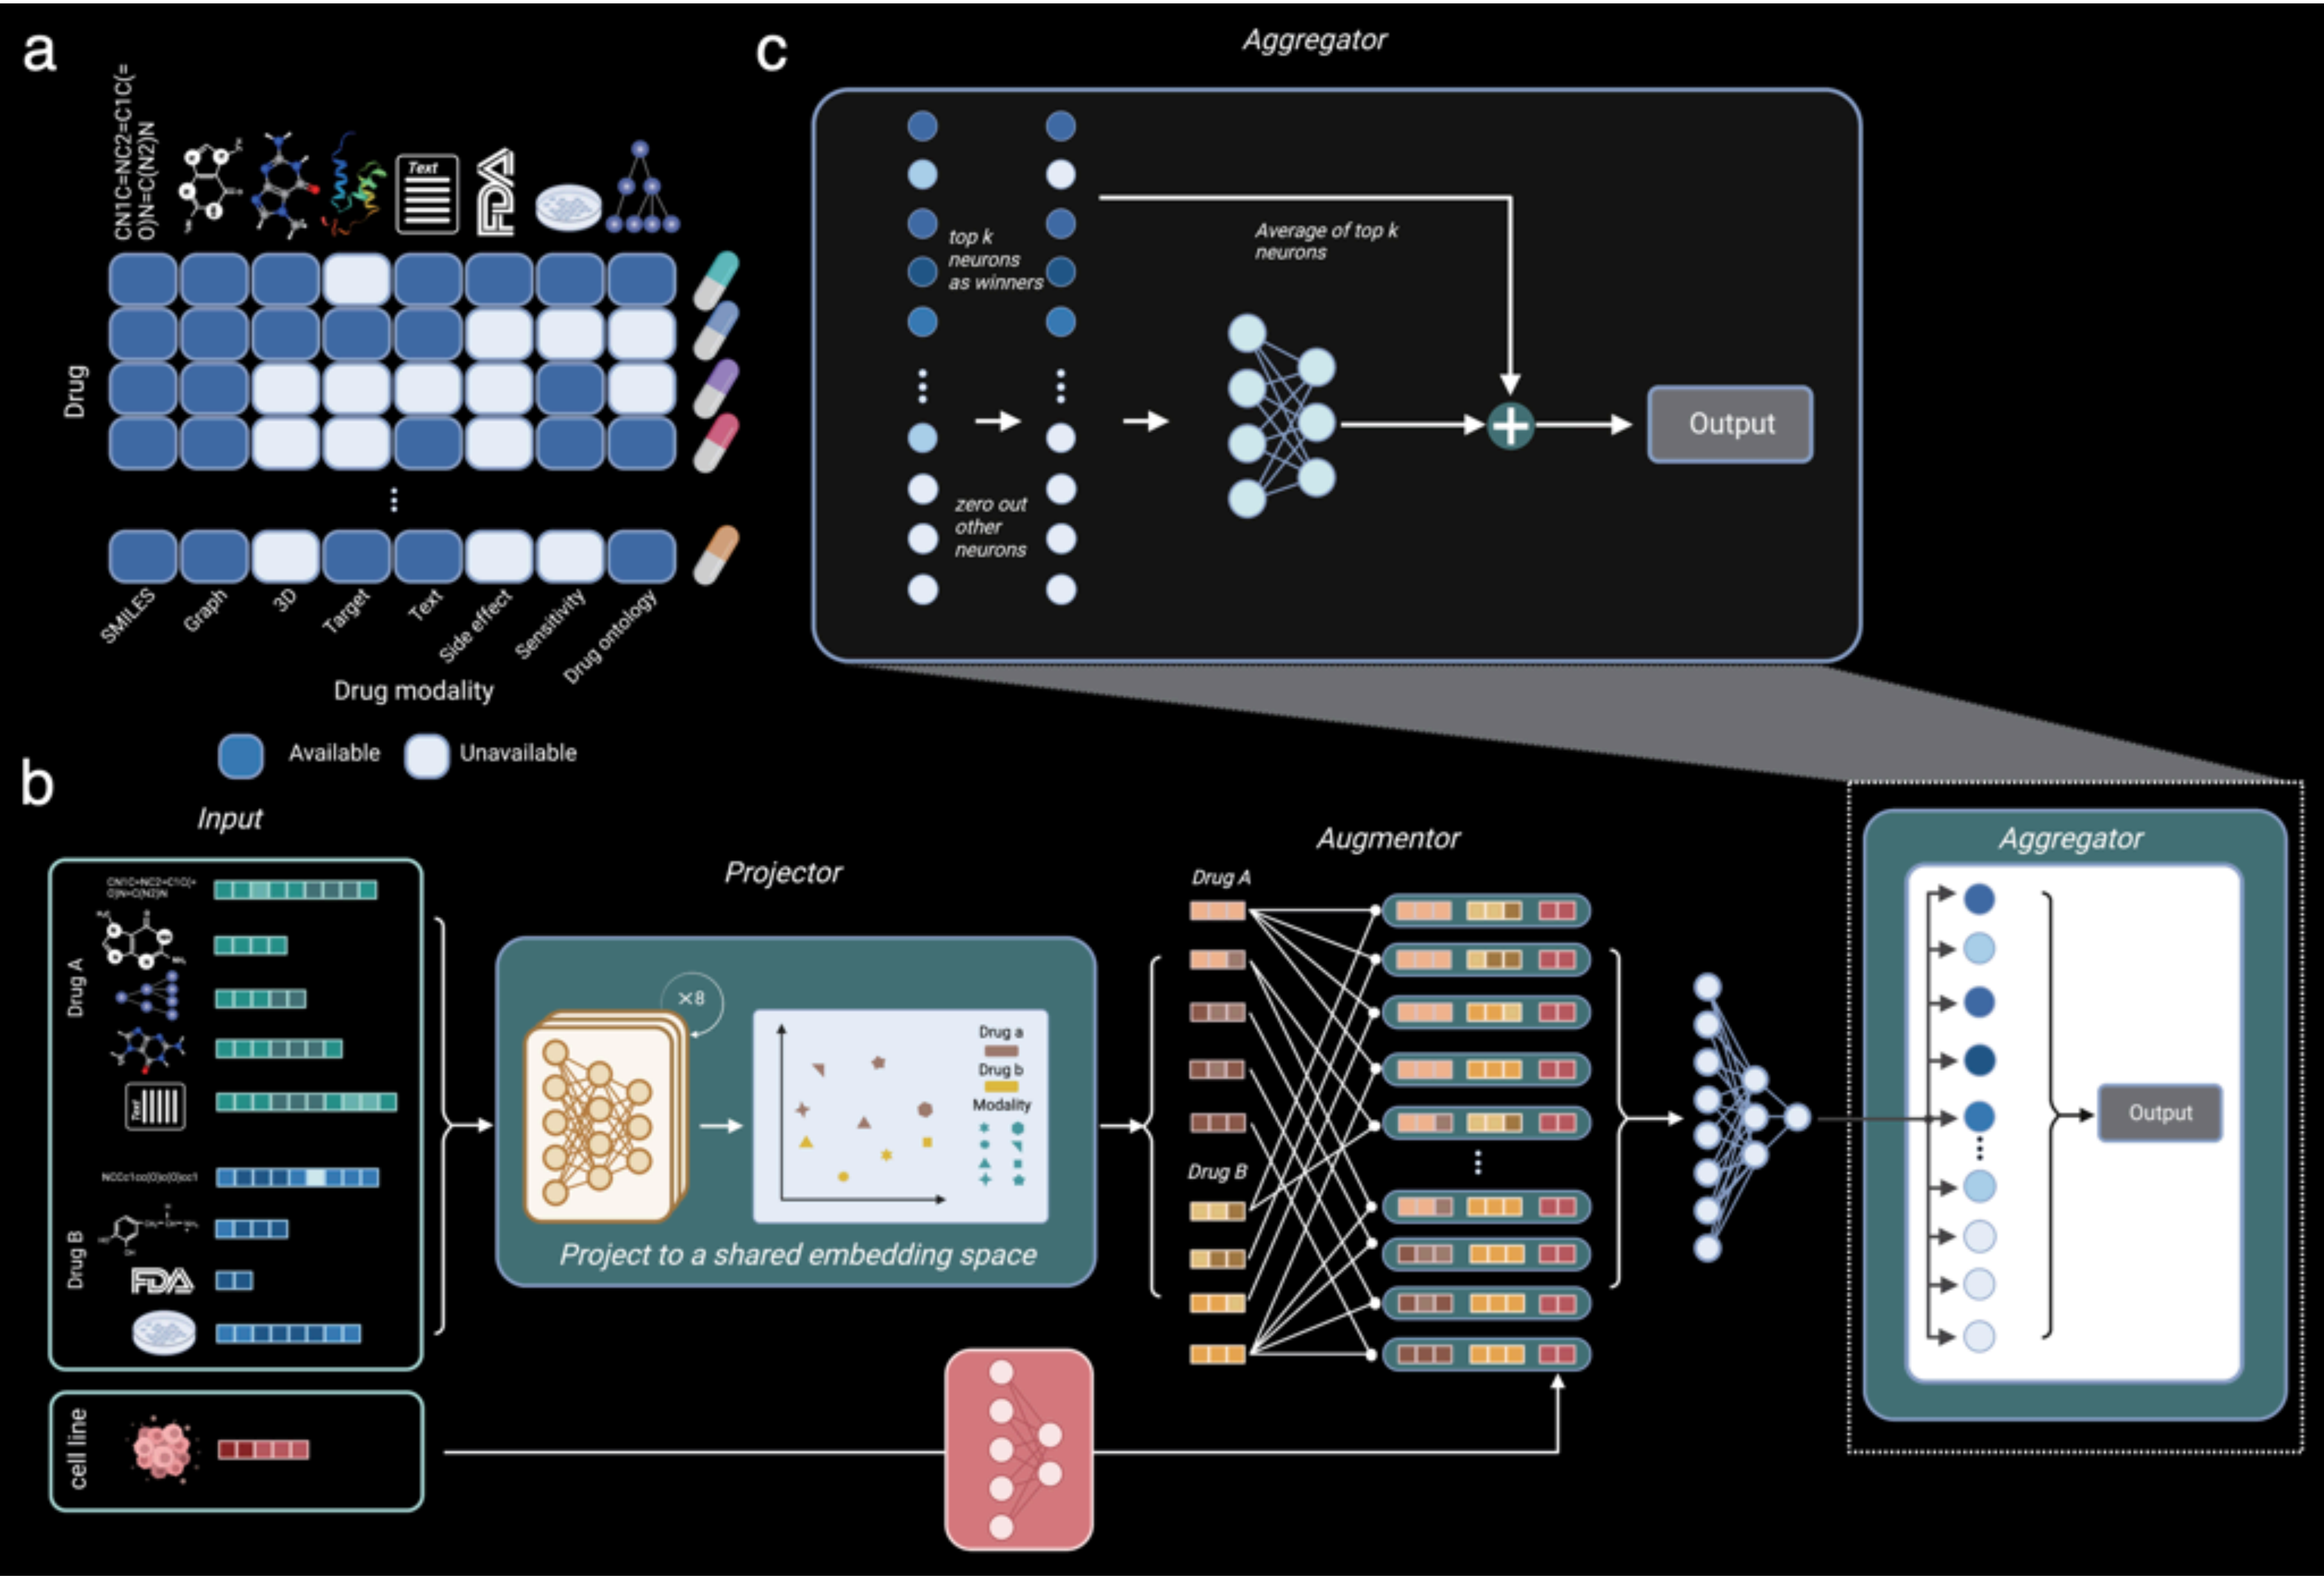
a
CN1C=NC2=C1C(=O)N=C(N2)N
Text
Drug
SMILES
Graph
3D
Target
Text
Side effect
Sensitivity
Drug ontology
Drug modality
Available
Unavailable
c
Aggregator
top k neurons as winners
zero out other neurons
Average of top k neurons
Output
b
Input
Drug A
Drug B
cell line
Projector
×8
Drug a
Drug b
Modality
Project to a shared embedding space
Augmentor
Drug A
Drug B
Aggregator
Output

A generalizable Hi-C foundation model for chromatin architecture, single-cell and multi-omics analysis across species

Xiao Wang*[1,2], Yuanyuan Zhang*[3], Suhita Ray[4], Anupama Jha[1], Tangqi Fang[2], Shengqi Hang[2], Sergei Doulatov†[4], William Stafford Noble†[1,2], and Sheng Wang†[2]

[1]Department of Genome Sciences, University of Washington, Seattle, WA, USA
[2]Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA, 98105, USA
[3]Department of Computer Science, Purdue University, West Lafayette, IN, 47907, USA
[4]Division of Hematology and Oncology, University of Washington, Seattle, WA, 98105, USA

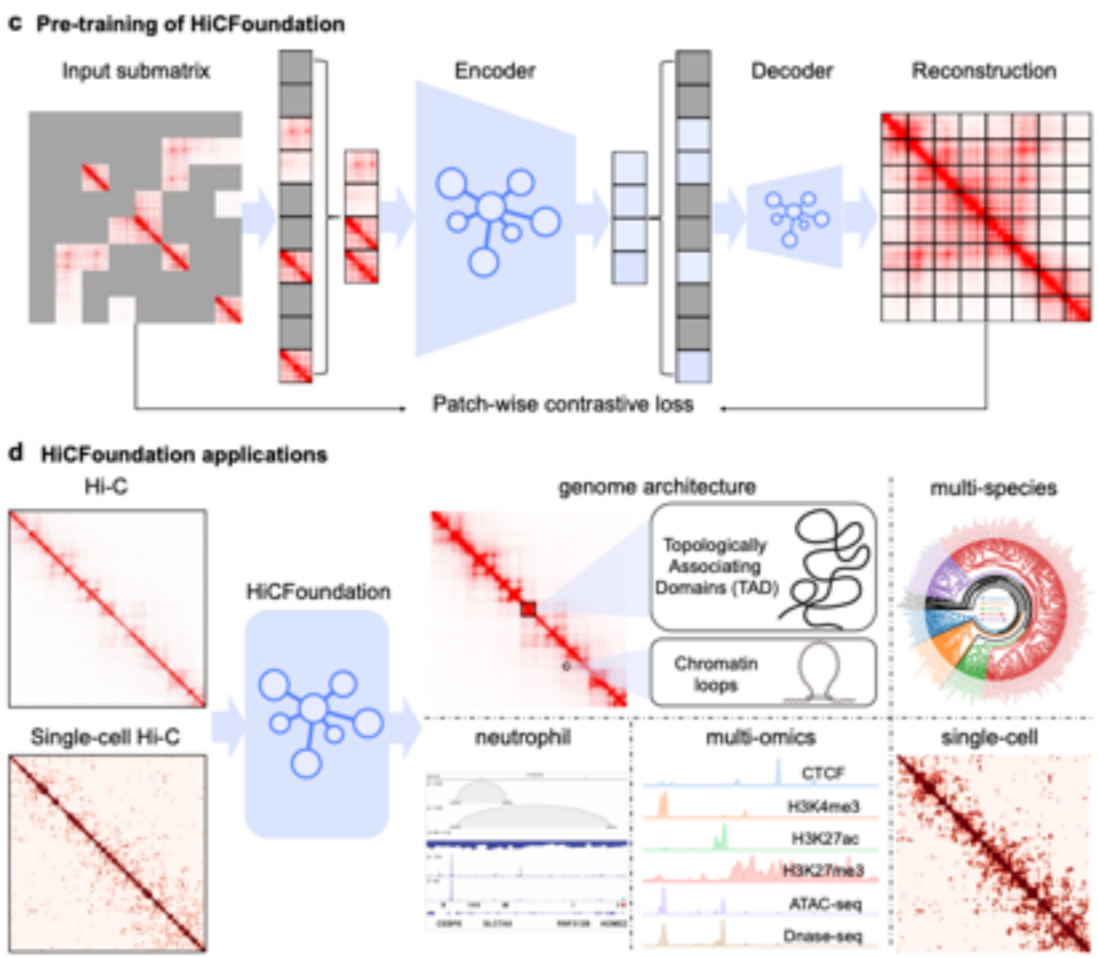
c Pre-training of HiCFoundation
Input submatrix
Encoder
Decoder
Reconstruction
Patch-wise contrastive loss
d HiCFoundation applications
Hi-C
Single-cell Hi-C
HiCFoundation
genome architecture
Topologically Associating Domains (TAD)
Chromatin loops
multi-species
neutrophil
multi-omics
CTCF
H3K4me3
H3K27ac
H3K27me3
ATAC-seq
Dnase-seq
single-cell

Towards a clinically accessible radiology multimodal model: open-access and lightweight, with automatic evaluation

Juan Manuel Zambrano Chaves[3◇*] Shih-Cheng Huang[3◇*], Yanbo Xu[1*], Hanwen Xu[2*], Naoto Usuyama[1*], Sheng Zhang[1*],

Fei Wang[4], Yujia Xie[1], Mahmoud Khademi[1], Ziyi Yang[1], Hany Awadalla[1], Julia Gong[1], Houdong Hu[1], Jianwei Yang[1], Chunyuan Li[1], Jianfeng Gao[1], Yu Gu[1], Cliff Wong[1], Mu Wei[1], Tristan Naumann[1], Muhao Chen[5], Matthew P. Lungren[1,3,6], Akshay Chaudhari[3], Serena Yeung-Levy[3], Curtis P. Langlotz[3], Sheng Wang[2,‡], Hoifung Poon[1,‡]

[1]Microsoft Research [2]University of Washington [3]Stanford University

[4]University of Southern California [5]University of California, Davis

[6]University of California, San Fransisco

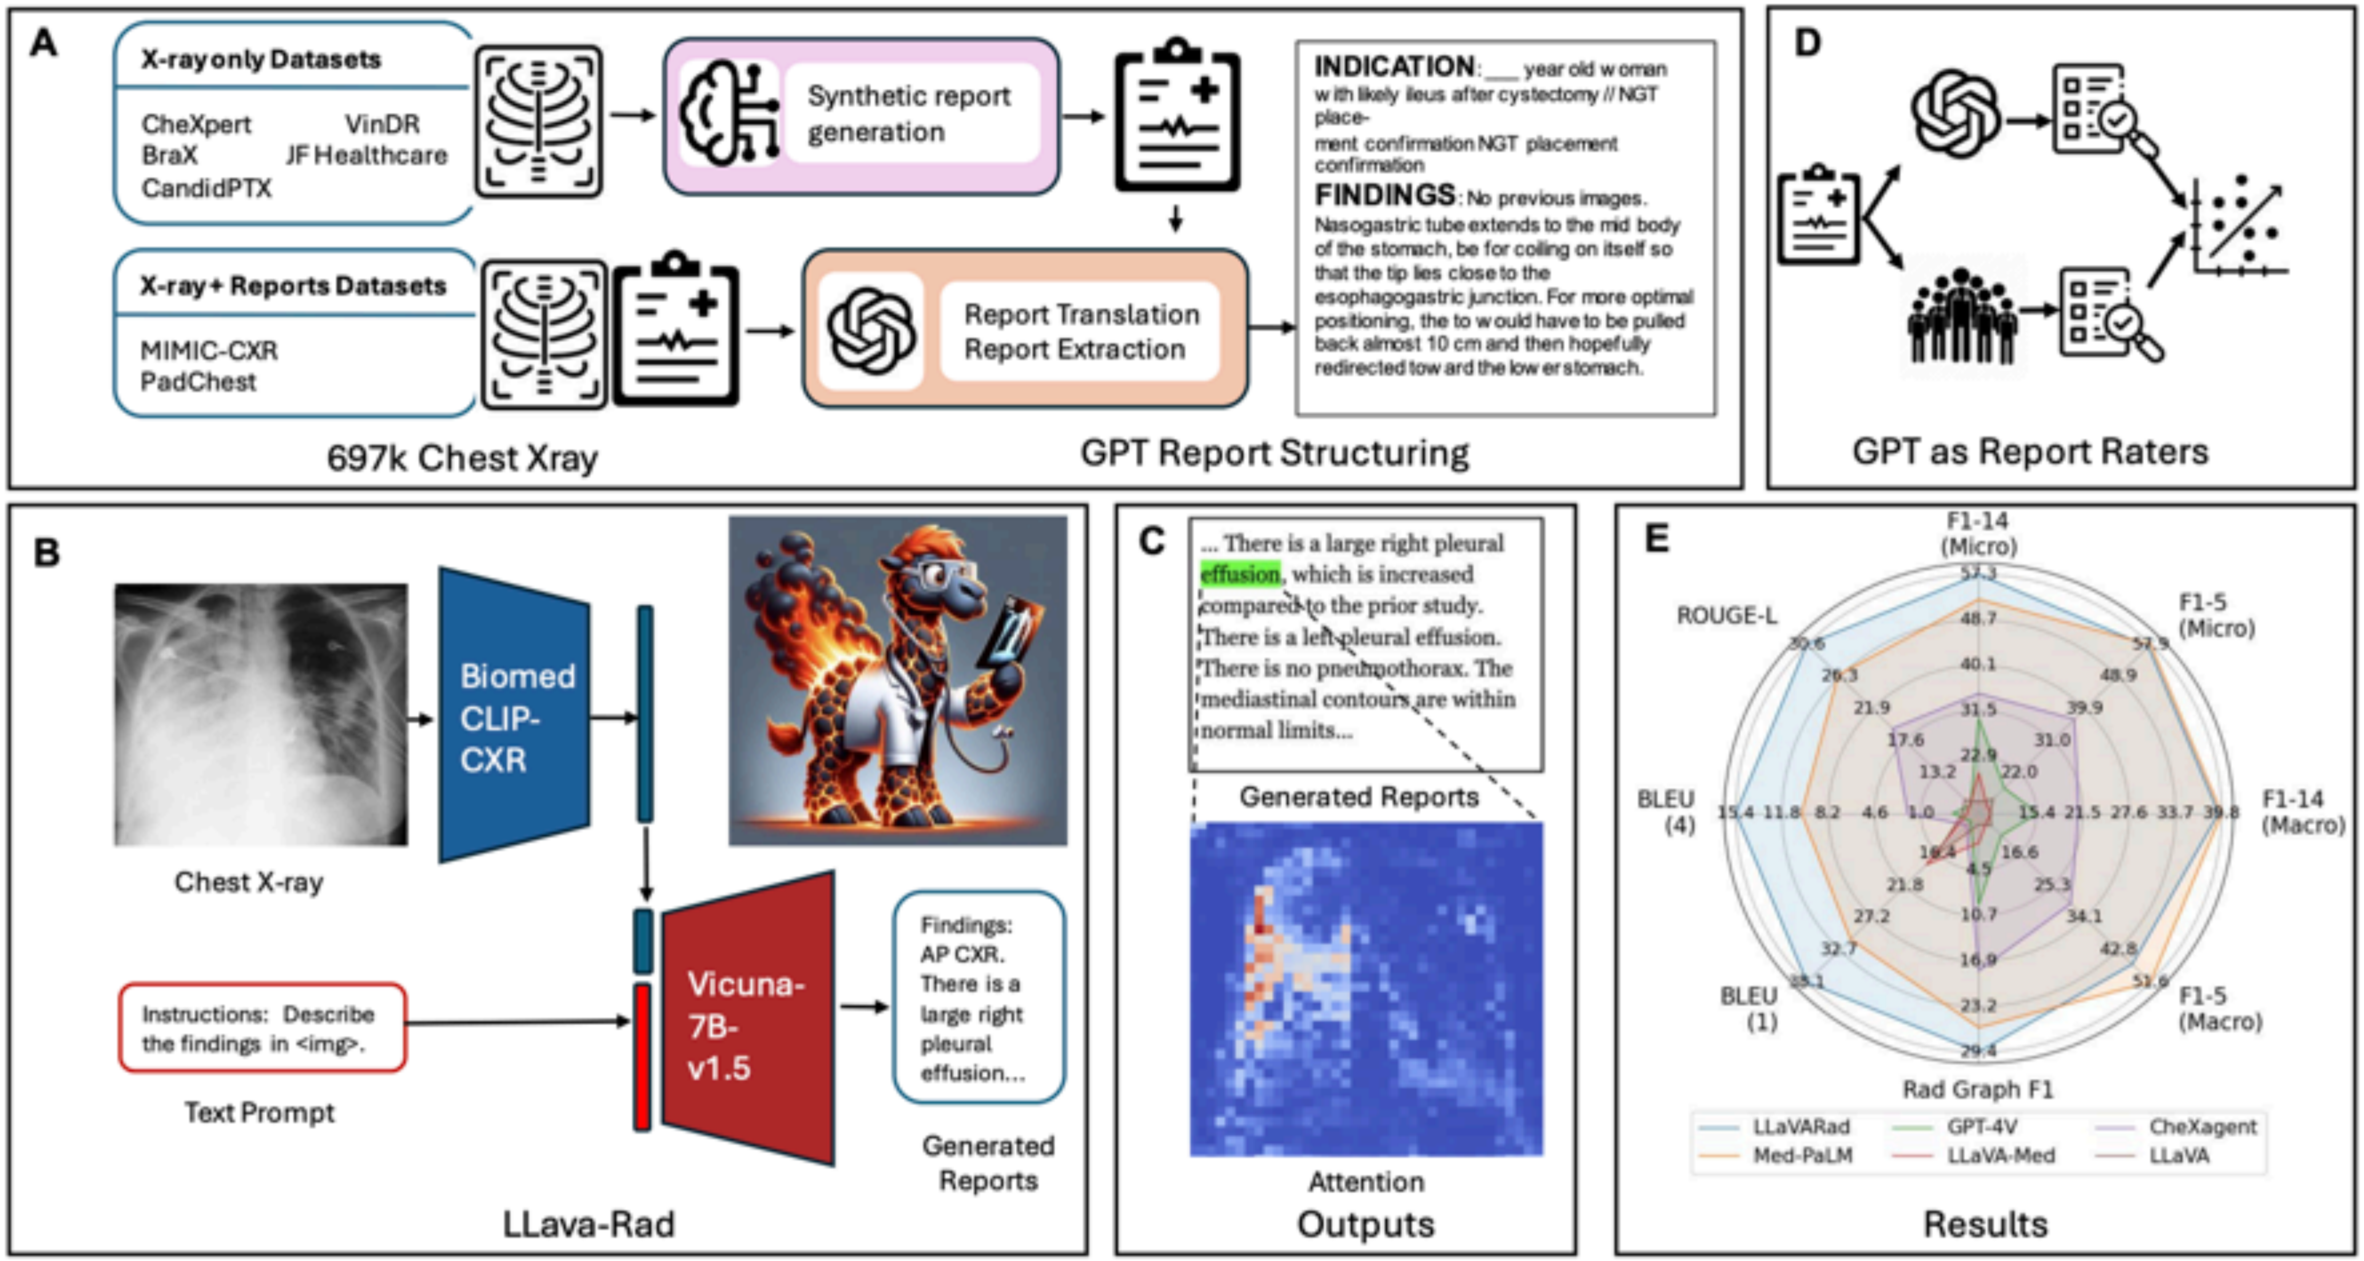
A
X-ray only Datasets
CheXpert
VinDR
BraX
JF Healthcare
CandidPTX
Synthetic report generation
X-ray+ Reports Datasets
MIMIC-CXR
PadChest
Report Translation
Report Extraction
INDICATION: ___ year old woman with likely ileus after cystectomy // NGT place-ment confirmation NGT placement confirmation
FINDINGS: No previous images. Nasogastric tube extends to the mid body of the stomach, be for coiling on itself so that the tip lies close to the esophagogastric junction. For more optimal positioning, the to would have to be pulled back almost 10 cm and then hopefully redirected toward the lower stomach.
697k Chest Xray
GPT Report Structuring
D
GPT as Report Raters
B
Chest X-ray
Biomed CLIP-CXR
Instructions: Describe the findings in <img>.
Text Prompt
Vicuna-7B-v1.5
Findings: AP CXR. There is a large right pleural effusion...
Generated Reports
LLava-Rad
C
... There is a large right pleural effusion, which is increased compared to the prior study. There is a left pleural effusion. There is no pneumothorax. The mediastinal contours are within normal limits...
Generated Reports
Attention
Outputs
E
F1-14 (Micro)
F1-5 (Micro)
F1-14 (Macro)
F1-5 (Macro)
Rad Graph F1
BLEU (1)
BLEU (4)
ROUGE-L
LLaVARad
Med-PaLM
GPT-4V
LLaVA-Med
CheXagent
LLaVA
Results