BMIF203/BMI702
Week 6: Medical Imaging II

Kun-Hsing "Kun" Yu, MD, PhD

Department of Biomedical Informatics

Harvard Medical School

March 4th, 2025

Outline

- Human-interpretable features + machine learning

- Deep learning methods
  - Multi-modal foundation models
  - Model interpretation

- Clinical applications in cancer pathology diagnoses

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Method 1: Image Analyses by Human-Interpretable Features

- Define human-interpretable features
  - e.g., size and shape of an object
- Extract these features computationally
- Connect these features with outcomes of interest

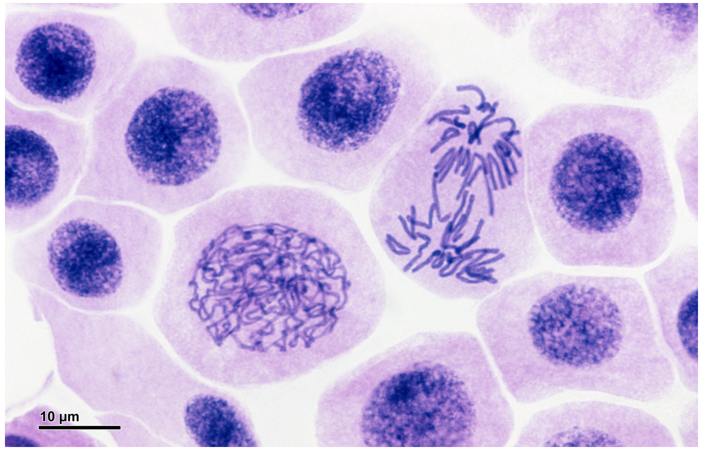

Image credit: Wikipedia

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Basic Features

- Size and shape (two-dimensional)
  - Area
    - Number of pixels in the region of interest
  - Perimeter
    - The total number of pixels around the boundary of each region

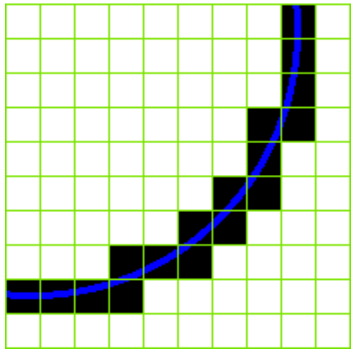

- Size and shape (three-dimensional)
  - Volume
    - Number of voxels in the region of interest
  - Surface area
    - The total number of voxels around the boundary of each region in the image

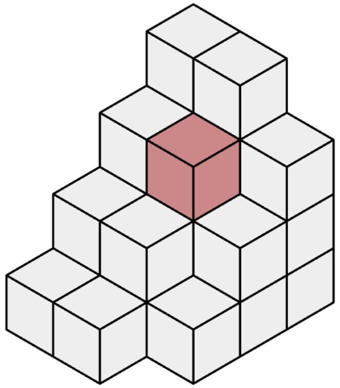

Image credit: Wikipedia

Basic Features

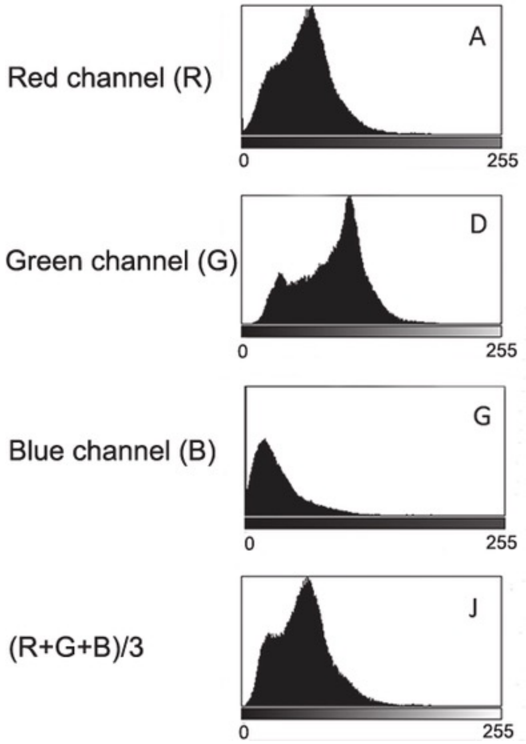

- Form factor
  - $4 * \pi * Area/Perimeter^2$
  - Form factor of a perfect circle = 1
- Intensity metrics
  - Mean intensity
    - Mean of pixel intensity values in the region/image
  - Median intensity, standard deviation of intensity values, median absolute deviation of pixel intensity values
- Saturation metrics
  - Percent of pixels at the maximum/minimum intensity value of the image

Image credit: Francisco Guilhien Gomes Jr et al.

Texture Features

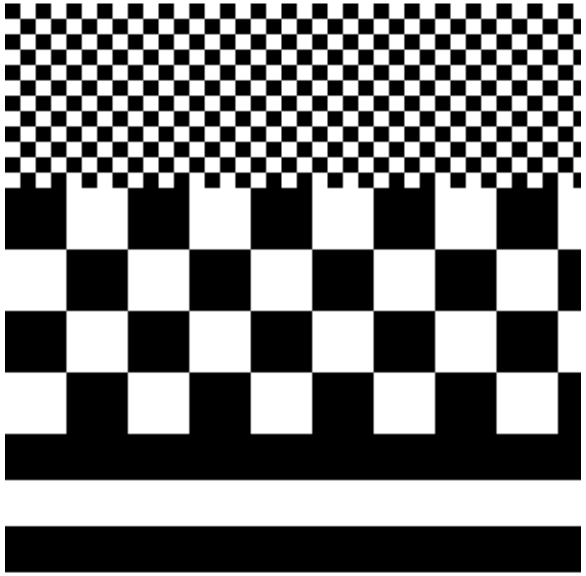

- Quantify the spatial arrangement of pixel intensities
- Example: Haralick texture features
  - Goal: to distinguish between rough and smooth patterns
  - Method: compute summary statistics of the gray-level co-occurrence matrices (GLCM; next slide)

Image credit: Wikipedia

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Gray-Level Co-occurrence Matrix (GLCM)

- Example:

1	3	2
2	3	1
3	3	2

Left-right pairs of adjacent pixels:

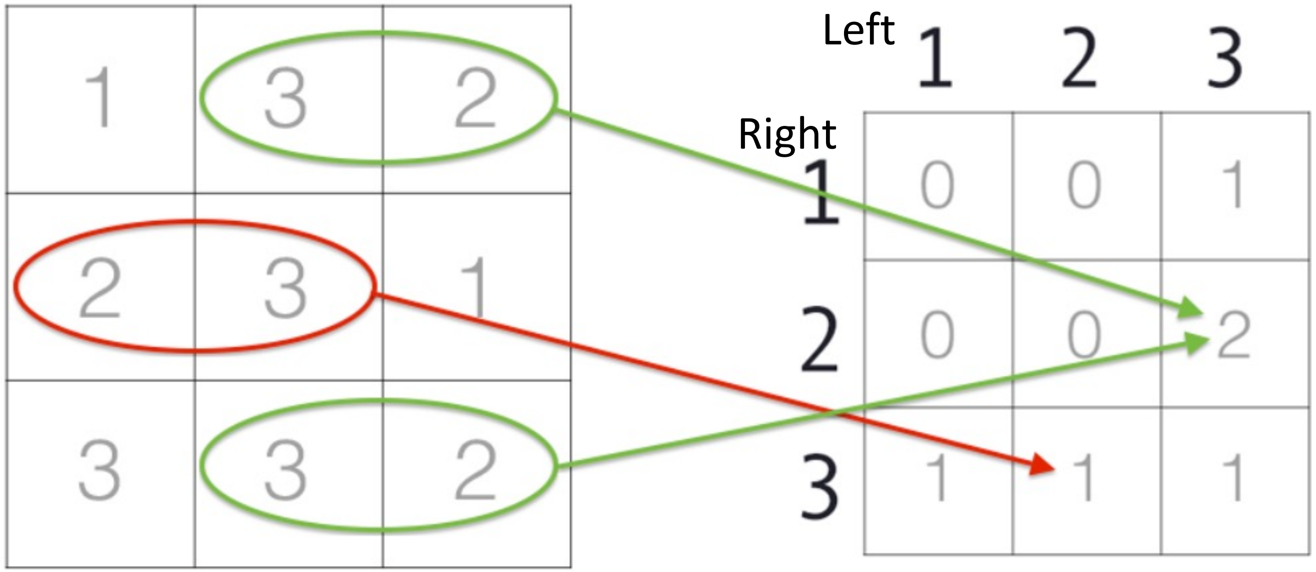

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Image credit: CVExplained

Gray-Level Co-occurrence Matrix (GLCM)

- Four directions of adjacency:

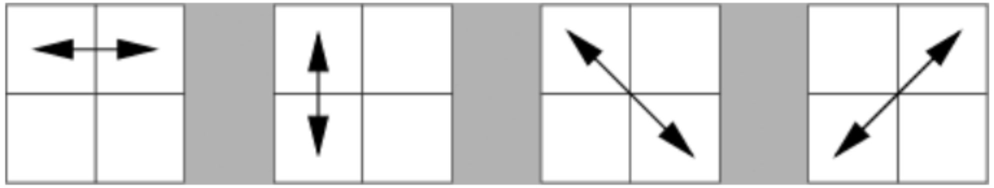

- Construct GLCMs for each direction

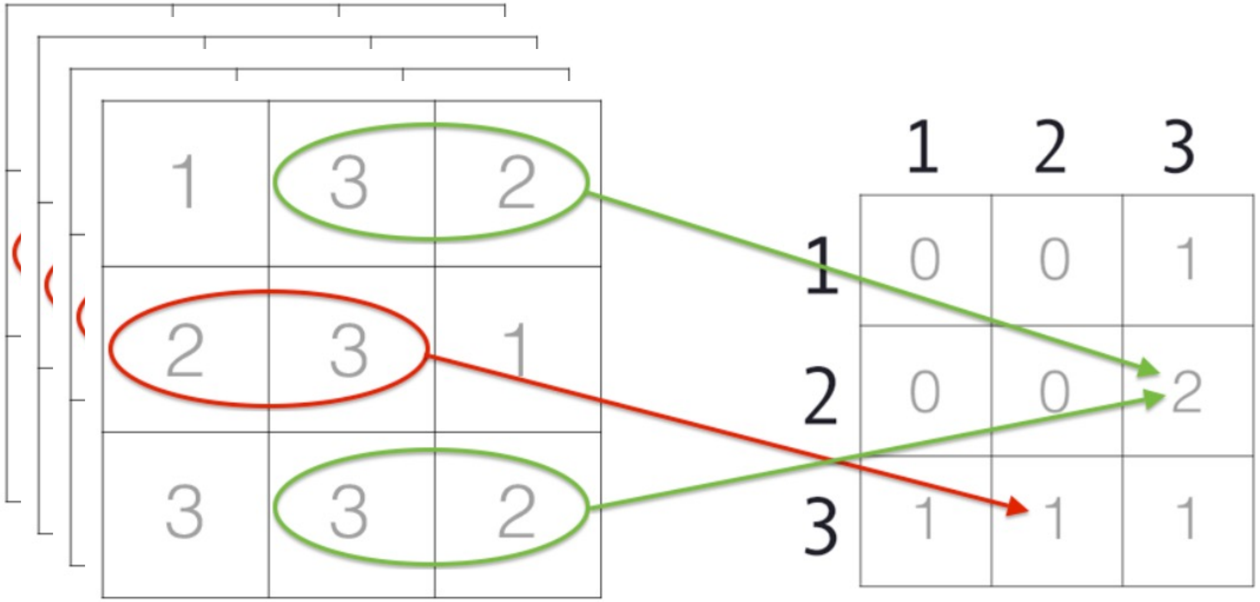

Image credit: CVExplained

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Haralick Texture Features

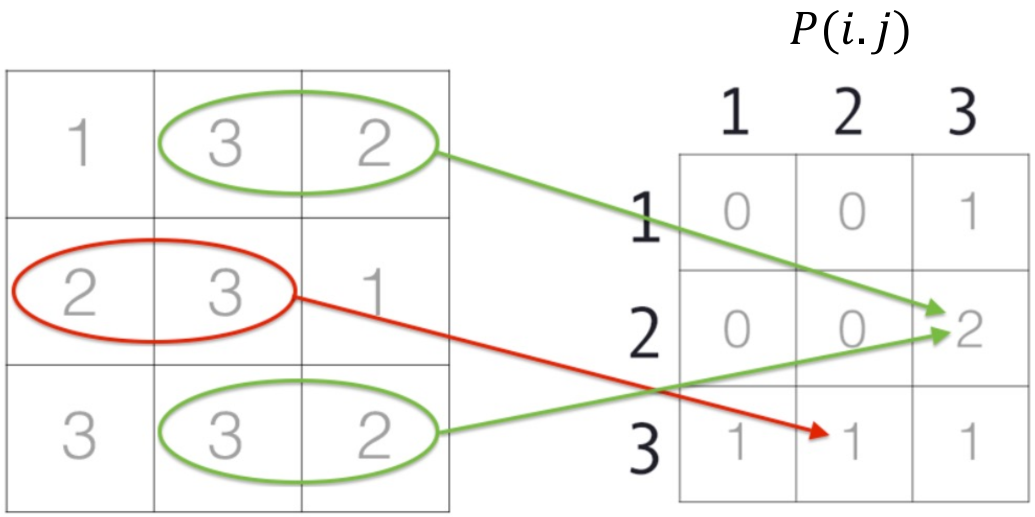

- Notation
  - R: The sum of all entries in a GLCM
  - $p(i.j)$: $(i.j)^{\text{th}}$ entry in a normalized GLCM
    - $p(i.j) = P(i.j)/R$
- Texture features for each GLCM
  - Angular second moment
    - $\sum_i \sum_j (p(i,j))^2$
  - Sum of squares
    - $\sum_i \sum_j (i-\mu)^2 p(i,j)$
  - Inverse difference moment
    - $\sum_i \sum_j \frac{1}{1+(i-j)^2} p(i,j)$
  - ... and 10 other features

- Finally, concatenate the features from each of the 4 GLCMs
  - Or simply take their average

Haralick RM et al. *IEEE Transactions on systems, man, and cybernetics*. 1973 Nov(6):610-21.

Zernike Shape Features

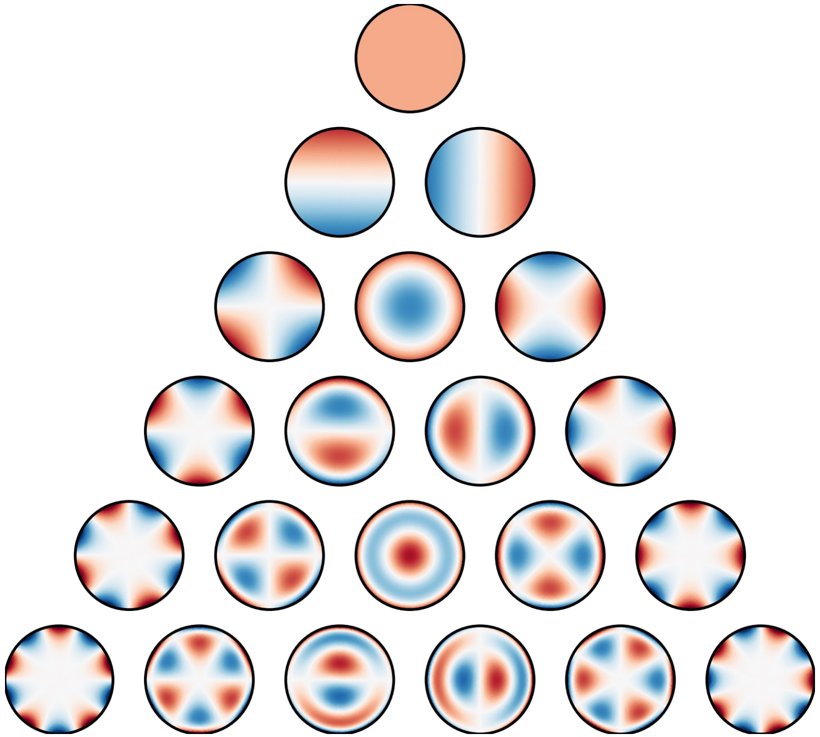

- Characterize the distribution of intensity across the object
  - Zernike polynomials: a sequence of polynomials that are orthogonal on the unit disk (a set of point whose distance from a given point P is less than 1)
  - We can decompose a region of interest into a weighted sum of a sequence of Zernike polynomials
- An example
  - Zernike (1,1): a prototype with a low intensity on one side and high on the other
- Note: these features are rotationally invariant

Zernike F. *Monthly Notices of the Royal Astronomical Society* 94 (1934): 377-384.

Clinical Applications

- Dermatology
  - Melanoma screening
- Ophthalmology
  - Diabetic retinopathy assessment using fundus photographs
- Radiology
  - Automated region of interest identification
- Pathology
  - Cancer diagnosis and subtyping

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Example: Human Interpretable Pathology Features Diagnose Cancers

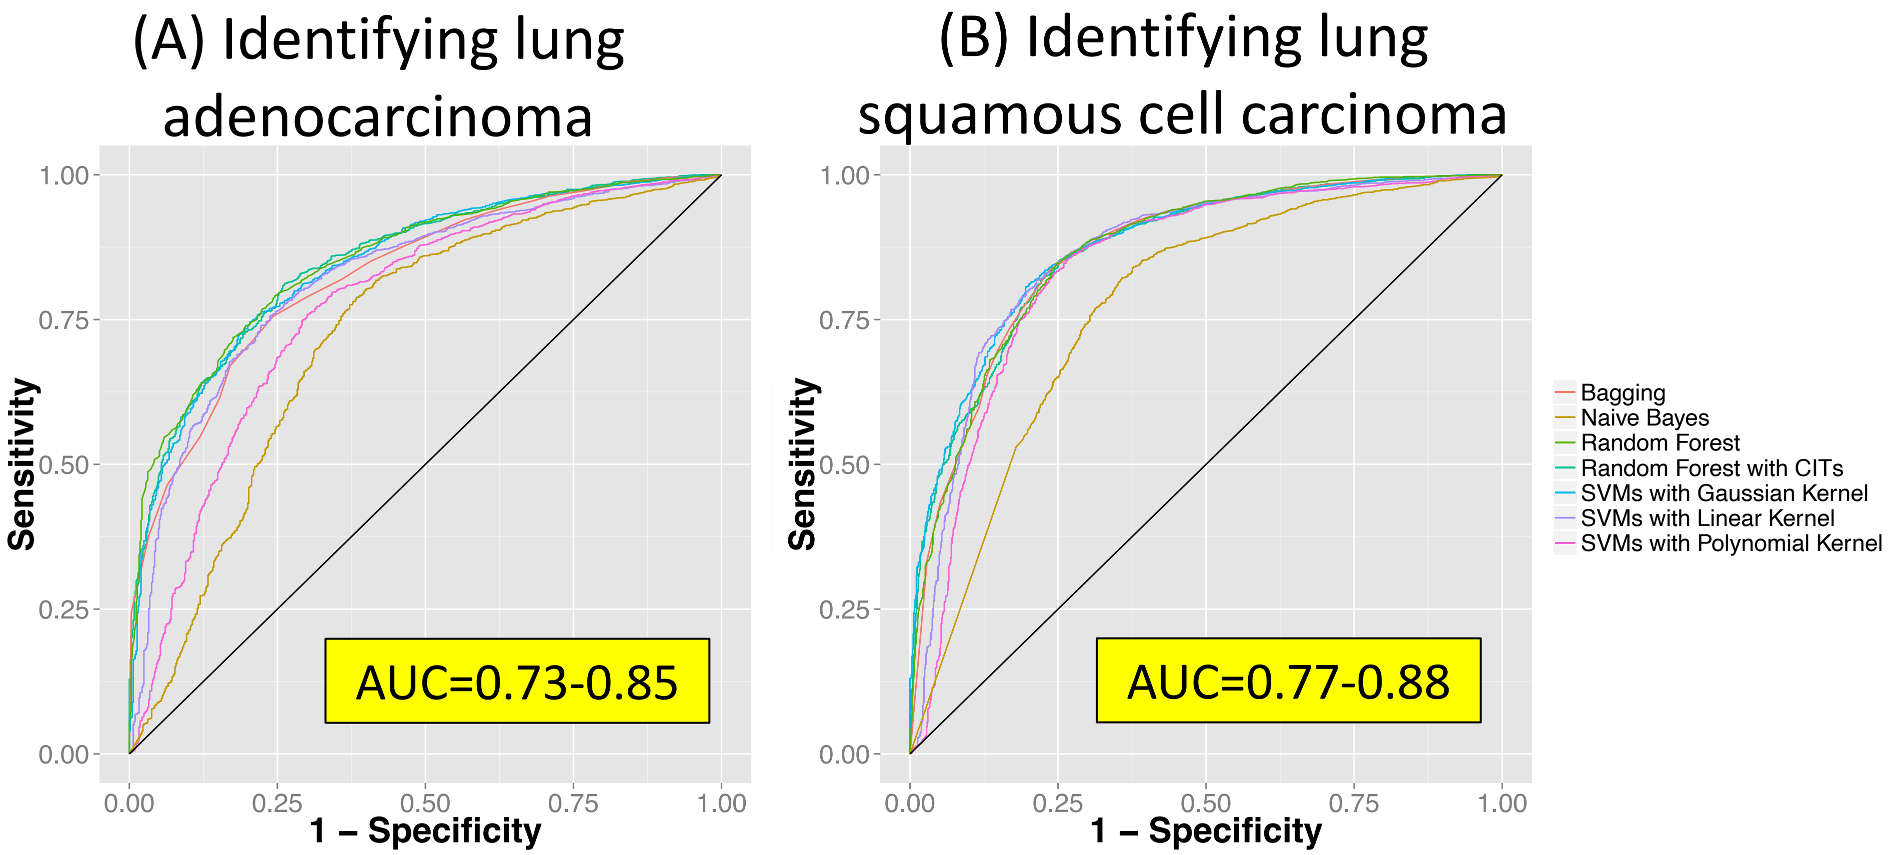

Top features: Textures (pixel correlations, intensity variance) of the nuclei

Yu KH et al. *Nature Communications*. 2016 Aug 16;7:12474.

Clinical Baseline: Stage and Grade are Often Insufficient to Predict Patient Survival

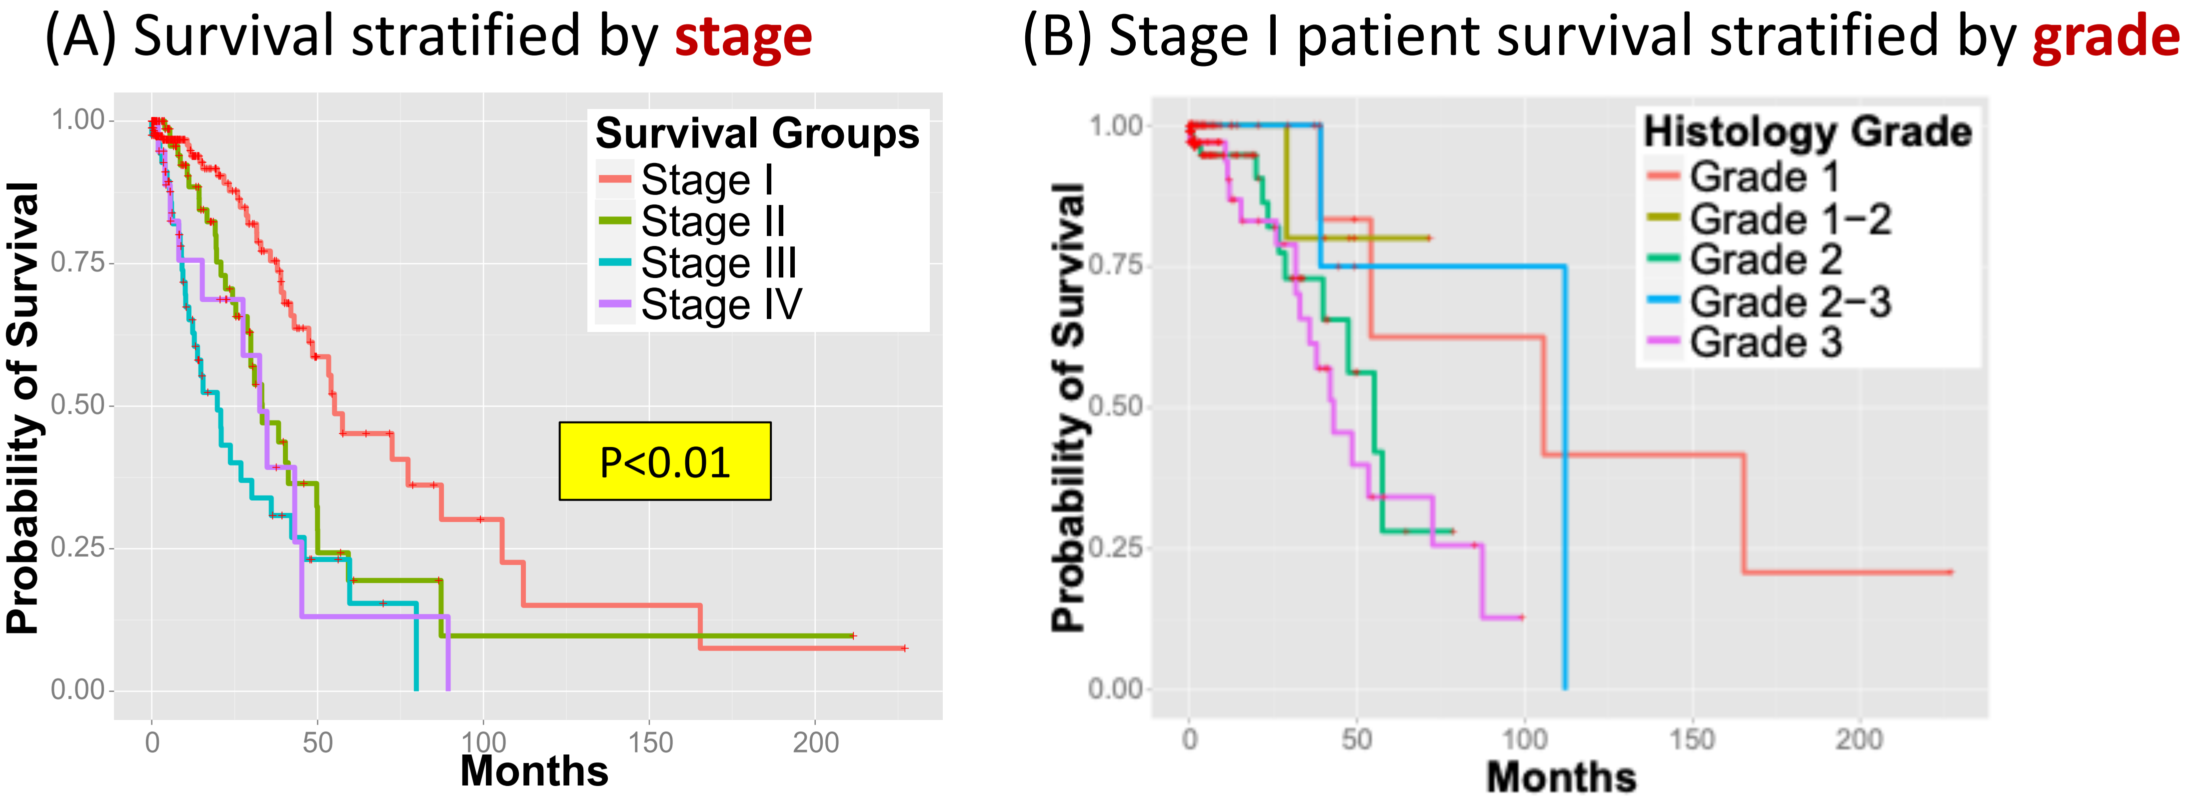

Dissimilar survival outcomes among stage I patients
Pathology grade did NOT significantly correlate with survival

Image Features Predicted Prognosis in Stage I Lung Adenocarcinoma Patients

(A) Image features predicted the prognosis of stage I patients

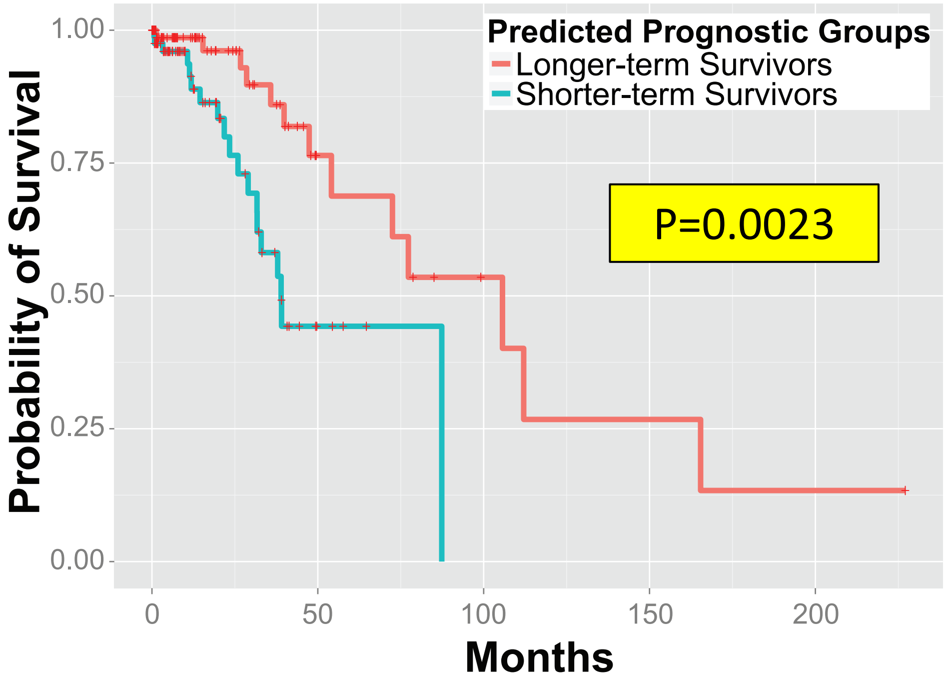

(B) Validated in an independent validation set

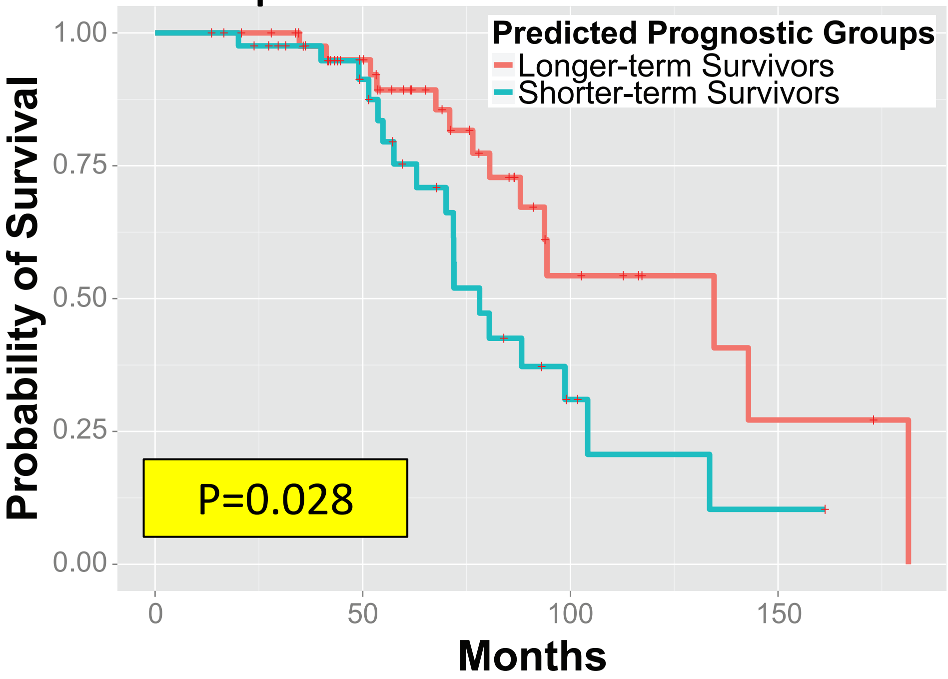

Quantitative features predicted survival, validated in TMA

Top features: Zernike shape features of the nuclei

Yu KH et al. *Nature Communications*. 2016 Aug 16;7:12474.

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Method 2: Deep Learning Approaches for Medical Image Analyses

- Reusing the neural network architectures for nature image analyses + fine-tuning
  - Convolutional neural networks
    - AlexNet, VGGNet, ResNet, DenseNet, EfficientNet
  - Vision transformers

- Designing specific models for the tasks/image types of interest
  - Automated hyperparameter search for model optimization
  - Pathology/radiology foundation models

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Foundation Models

- Machine learning models trained on vast datasets and can be applied across a wide range of use cases

- Examples of foundation models:
  - GPT series
  - BERT
  - DALL-E (image generation)

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Early Foundation Models for Pathology Imaging

- Vision-focused: CTransPath
  - A transformer-based feature extractor for pathology images

- Vision-language: Pathology language-image pretraining (PLIP)
  - A multimodal model trained with pathology images and natural language descriptions

- Other tile-level pathology foundation models
  - Lunit, Phikon, UNI, Virchow, etc.

Review article: Chanda D et al. arXiv preprint arXiv:2408.14496. 2024 Aug 23.

A Frequently-Employed Module in Image Foundation Models: Contrastive Learning

- Goal: Enhances the model performance by
  - Maximizing the differences between samples of different categories
  - Minimizing the differences between samples of the same category

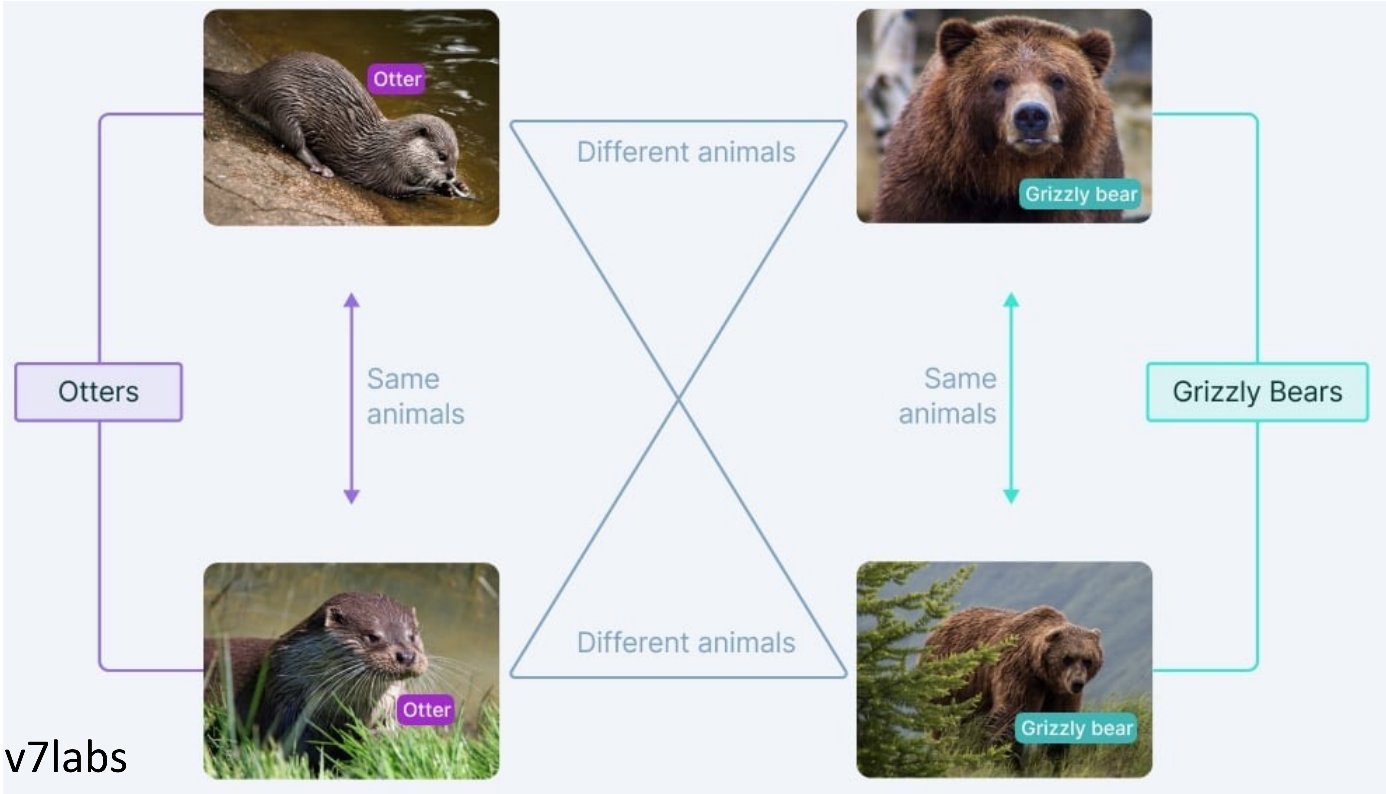

Image credit: v7labs

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Self-Supervised Contrastive Learning

- What can we do if we don't have a lot of labeled data?
  - Use augmented data as the positive instances!

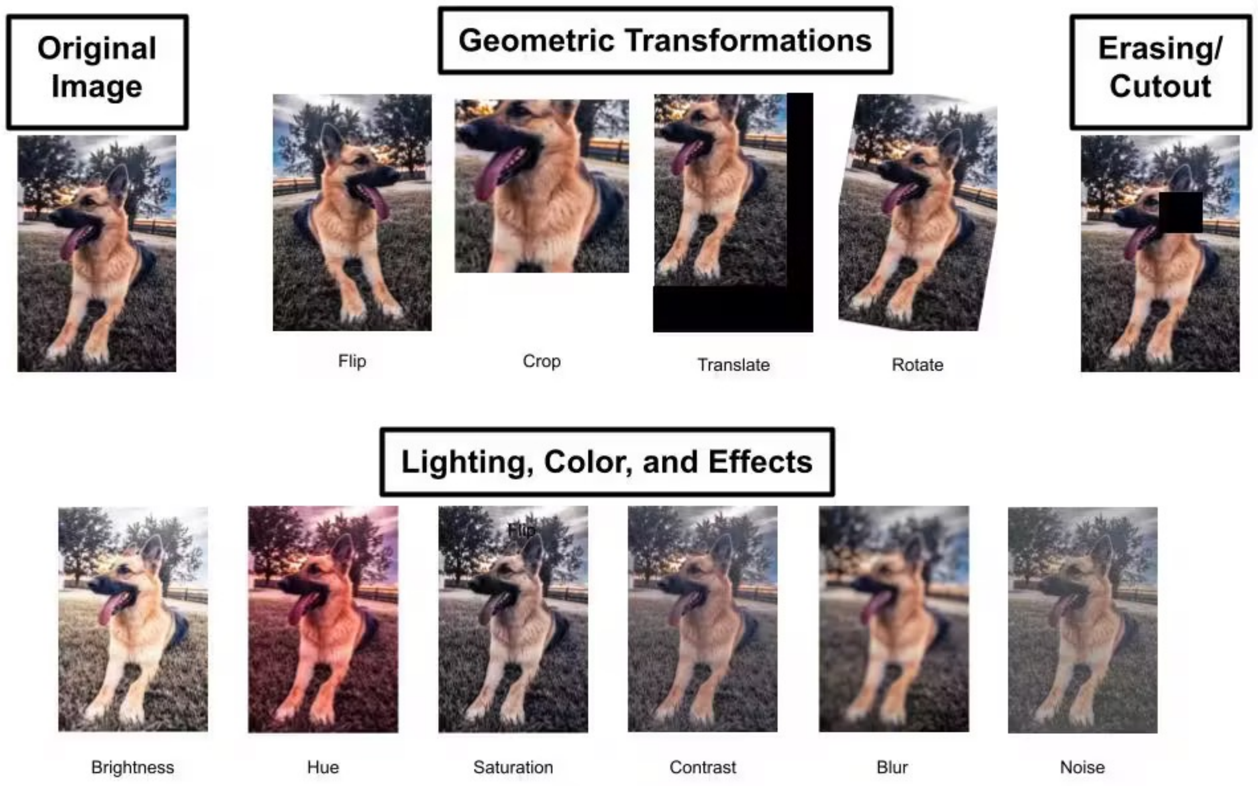

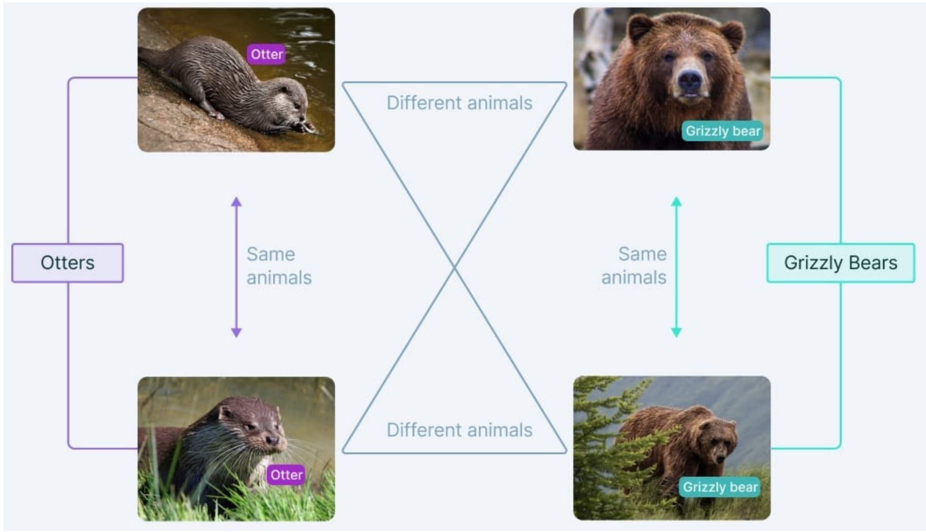

Image credit: v7labs

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Loss Functions for Contrastive Learning

- Contrastive loss
  - Maximize the agreement between positive pairs (instances from the same category) in the embedding space
  - Minimize the agreement between negative pairs (instances from different categories) in the embedding space
- Triplet loss
  - Triplets of instances: an anchor instance, a positive sample (similar to the anchor), and a negative sample (dissimilar to the anchor)
  - Goal: Distance (anchor, positive sample) < Distance (anchor, negative sample) + ε
- …

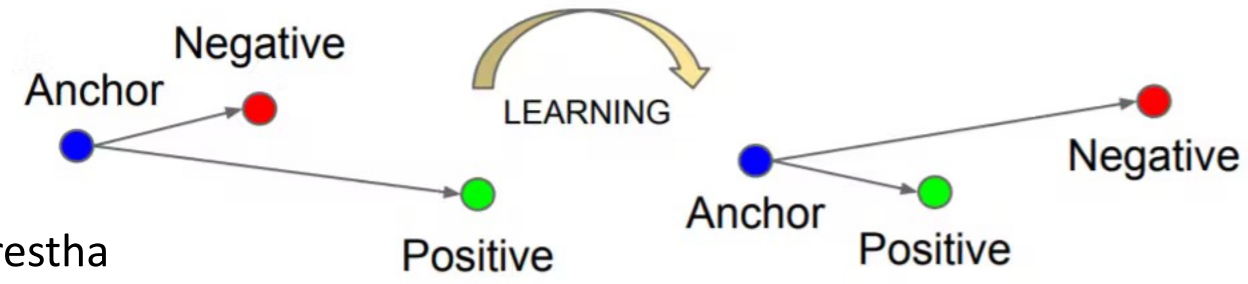

Image credit: Enosh Shrestha

Example: Clinical Histopathology Imaging Evaluation Foundation (CHIEF) Model

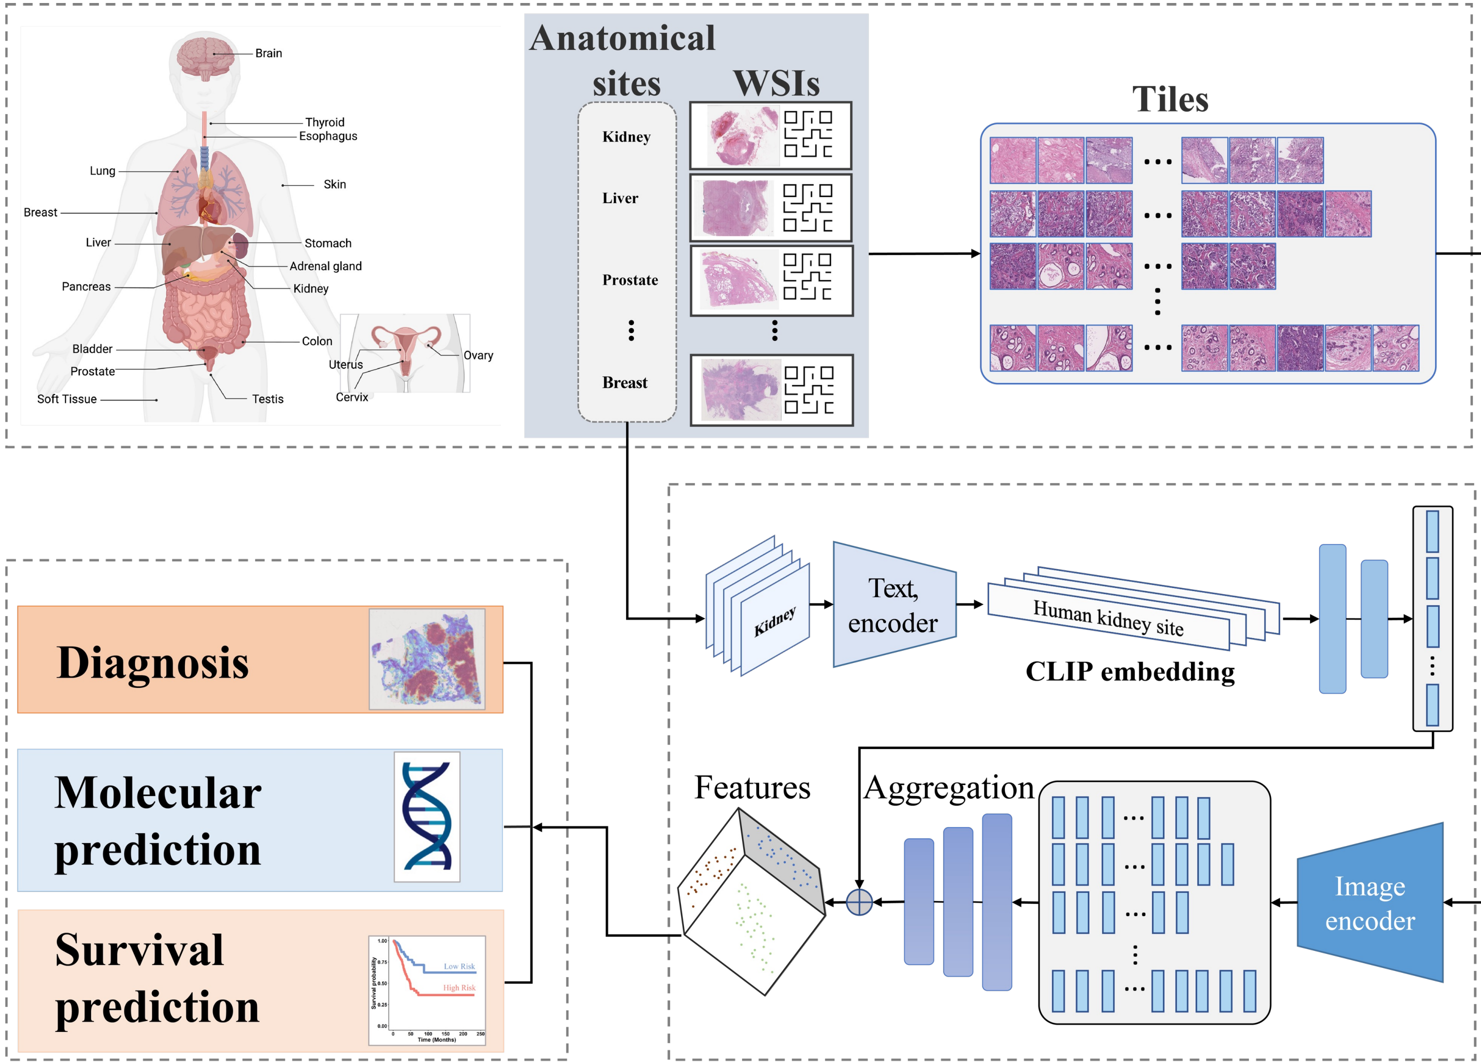

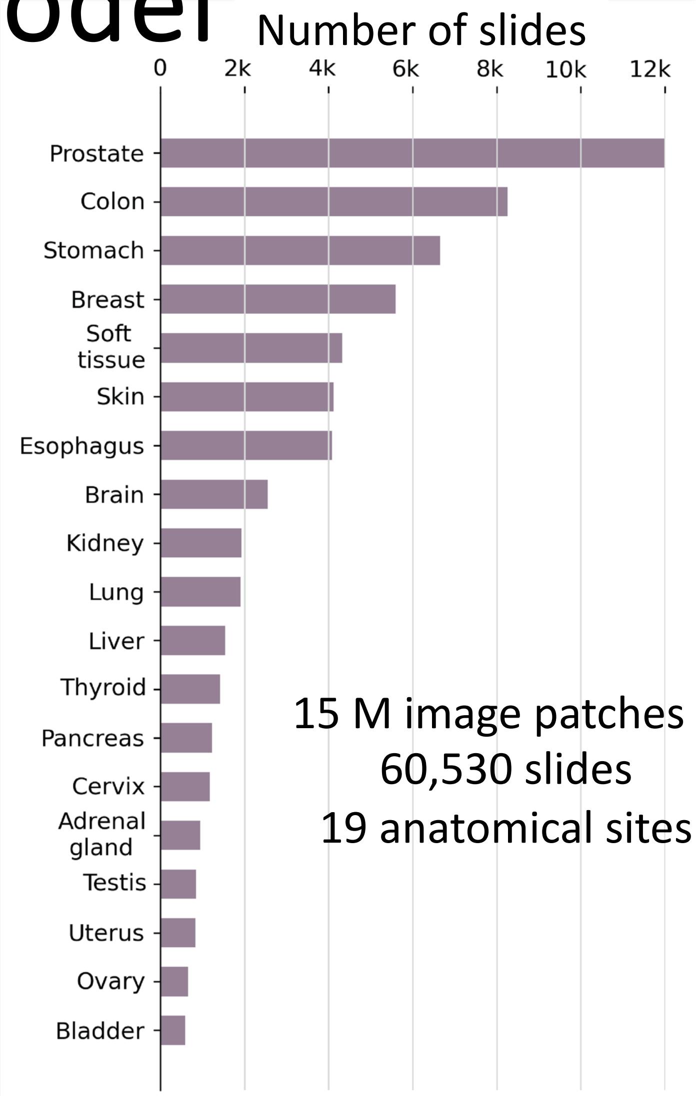

15 M image patches +
60,530 slides
19 anatomical sites

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

CHIEF's Image Feature Aggregation Framework

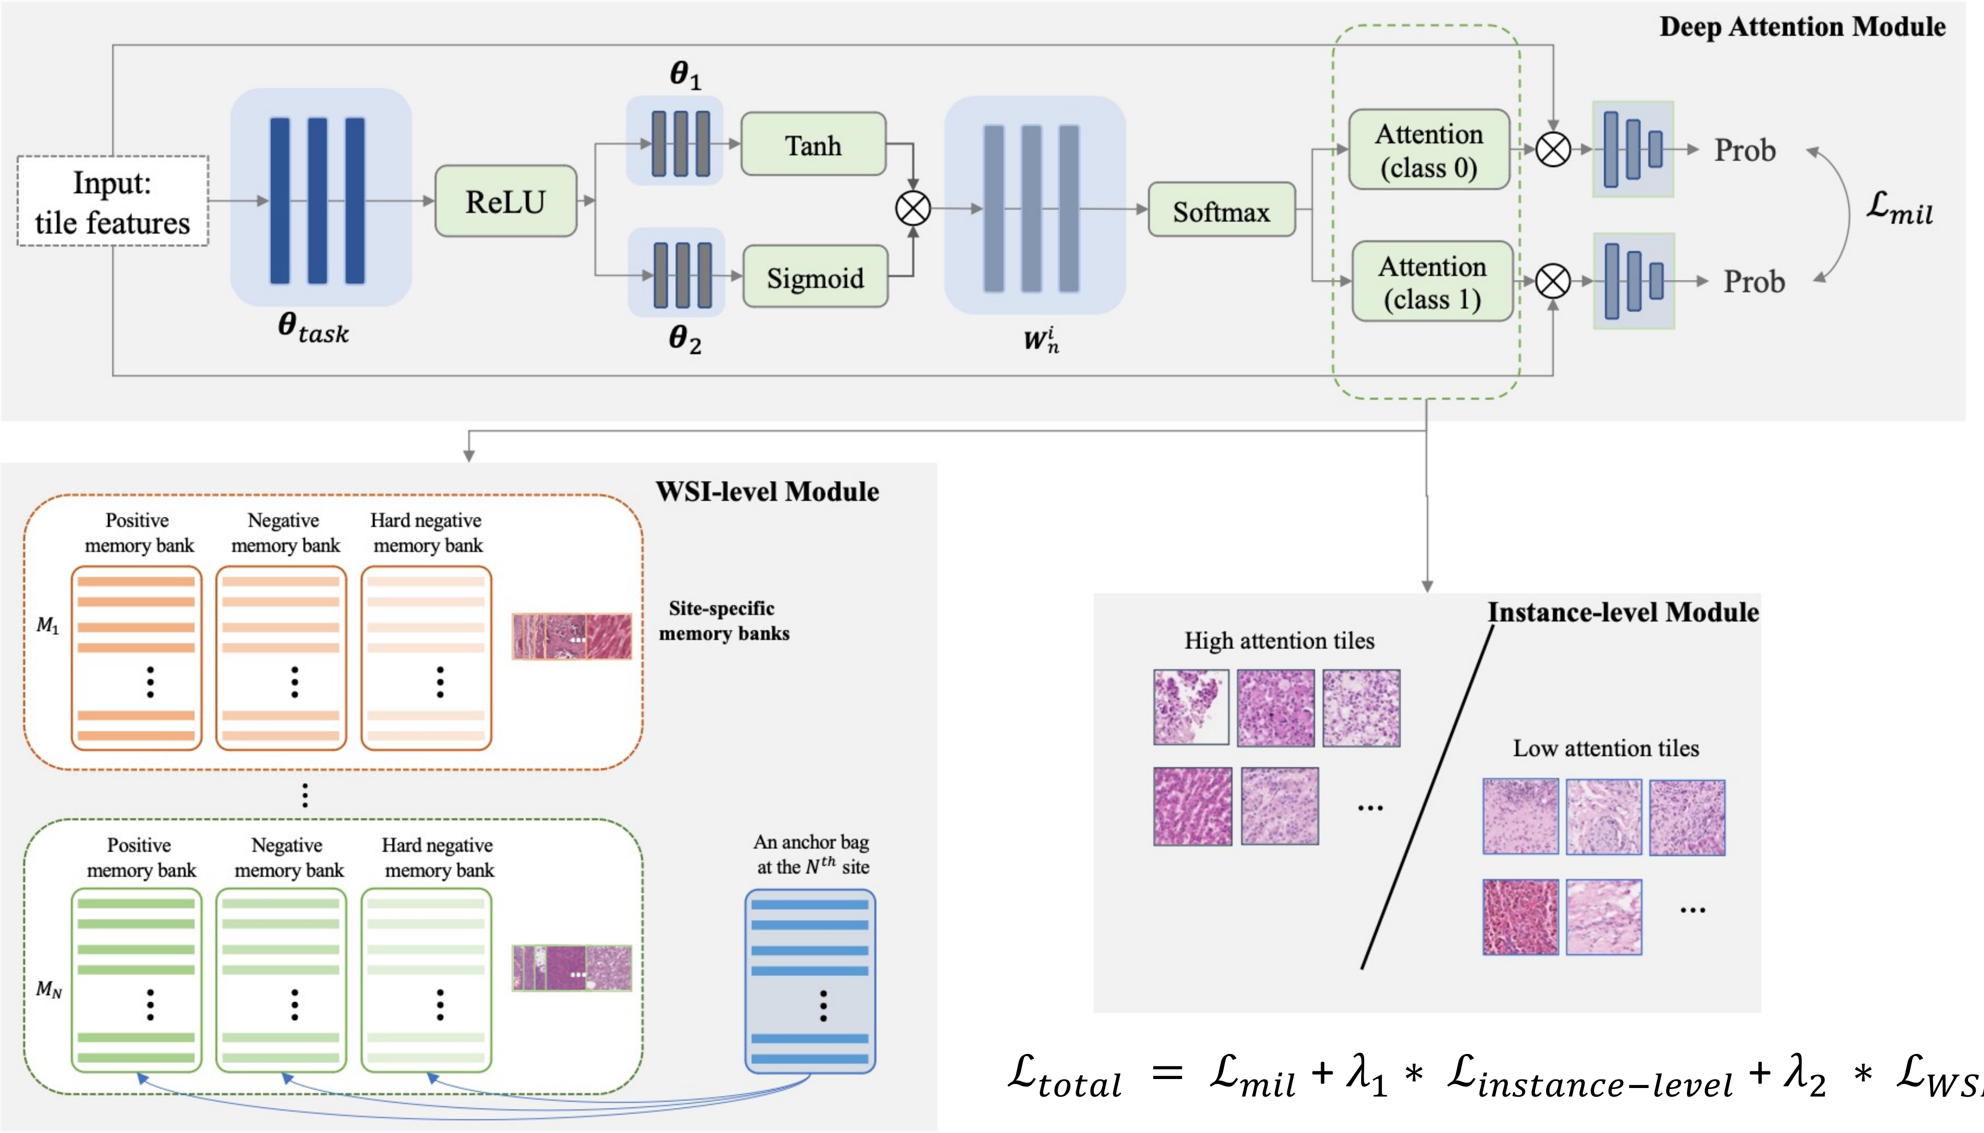

$$\mathcal{L}_{total} = \mathcal{L}_{mil} + \lambda_1 * \mathcal{L}_{instance-level} + \lambda_2 * \mathcal{L}_{WSI-level}$$

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

Cancer Cell Detection

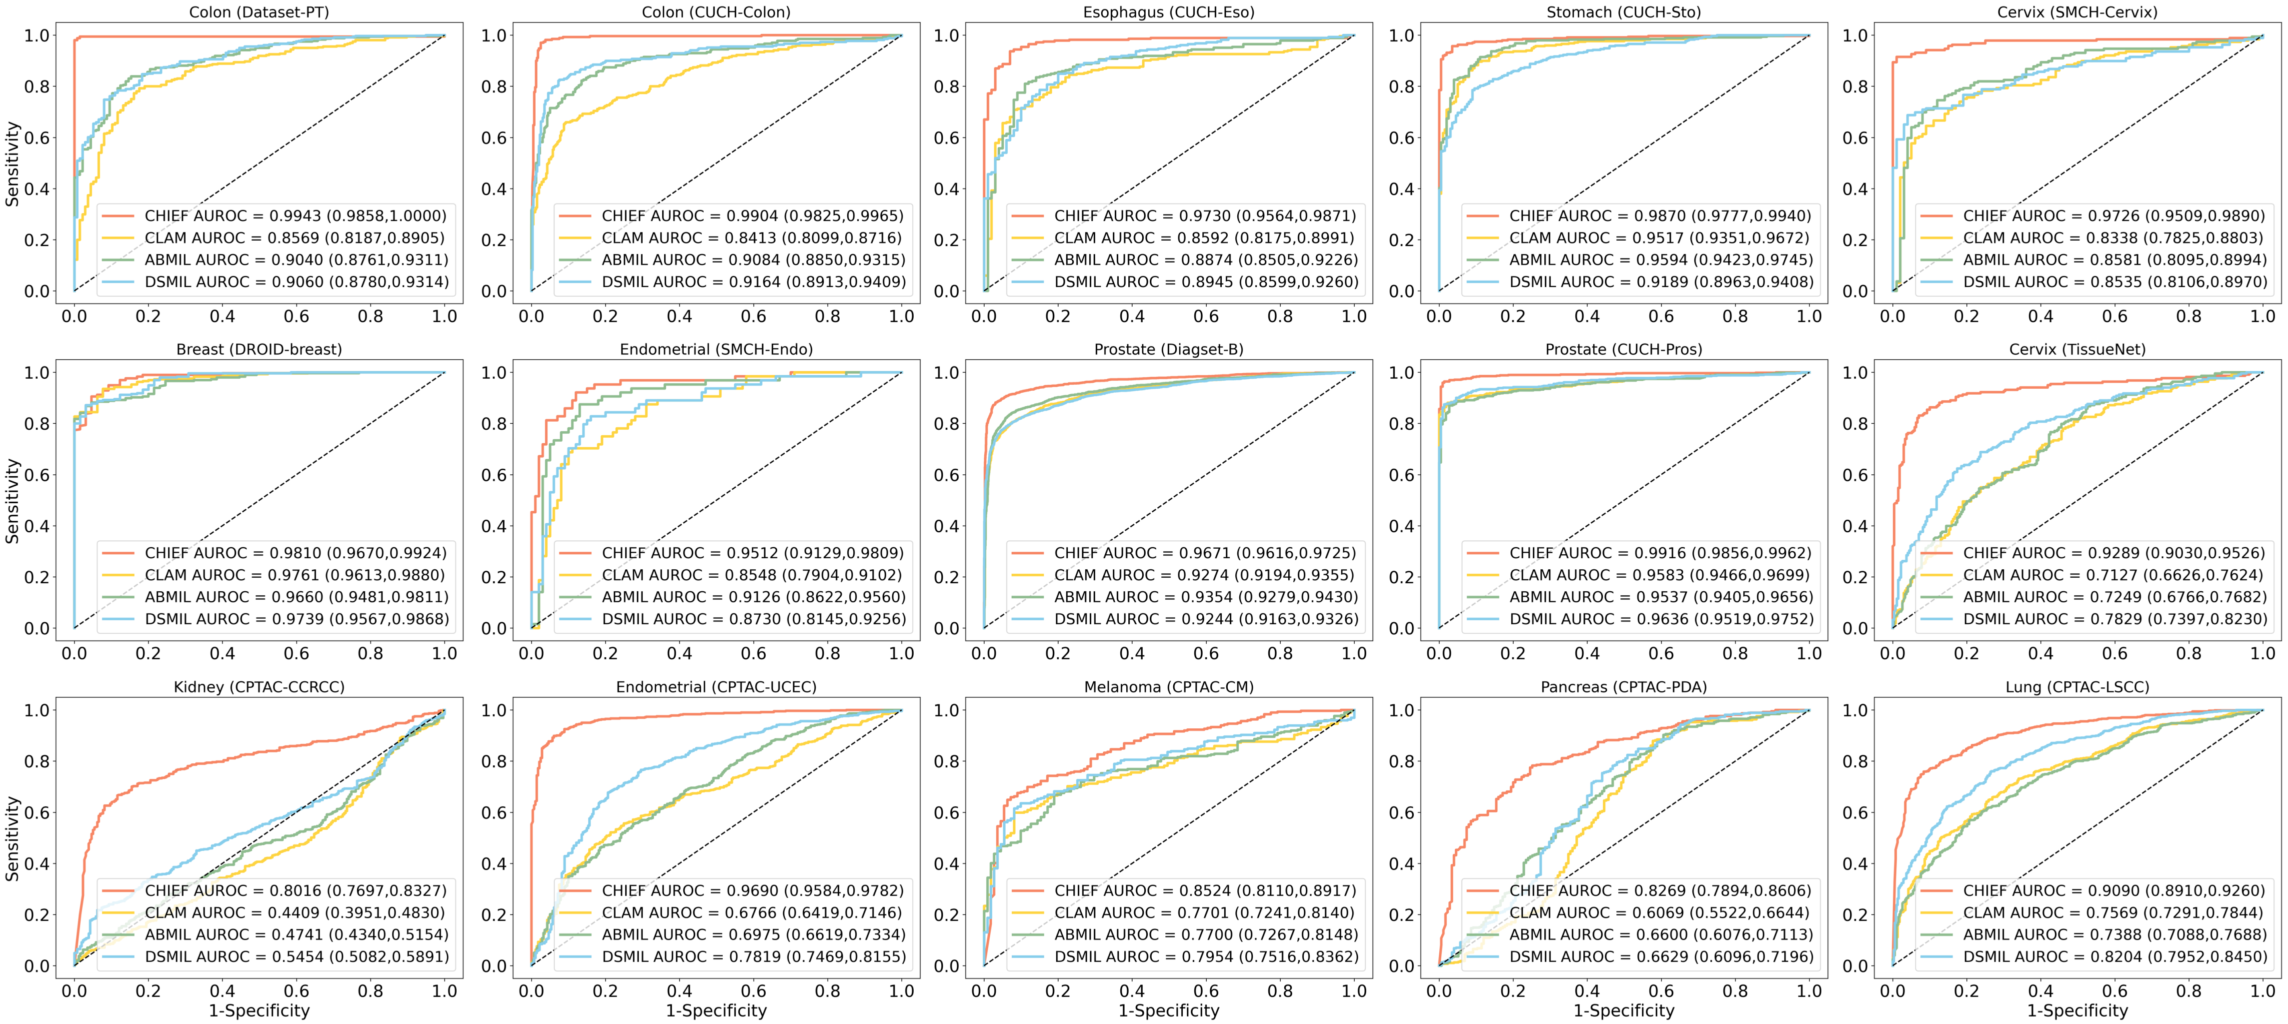

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

Weakly-Supervised Foundation Models Identified Cancer Cells Without Pixel-Level Segmentation

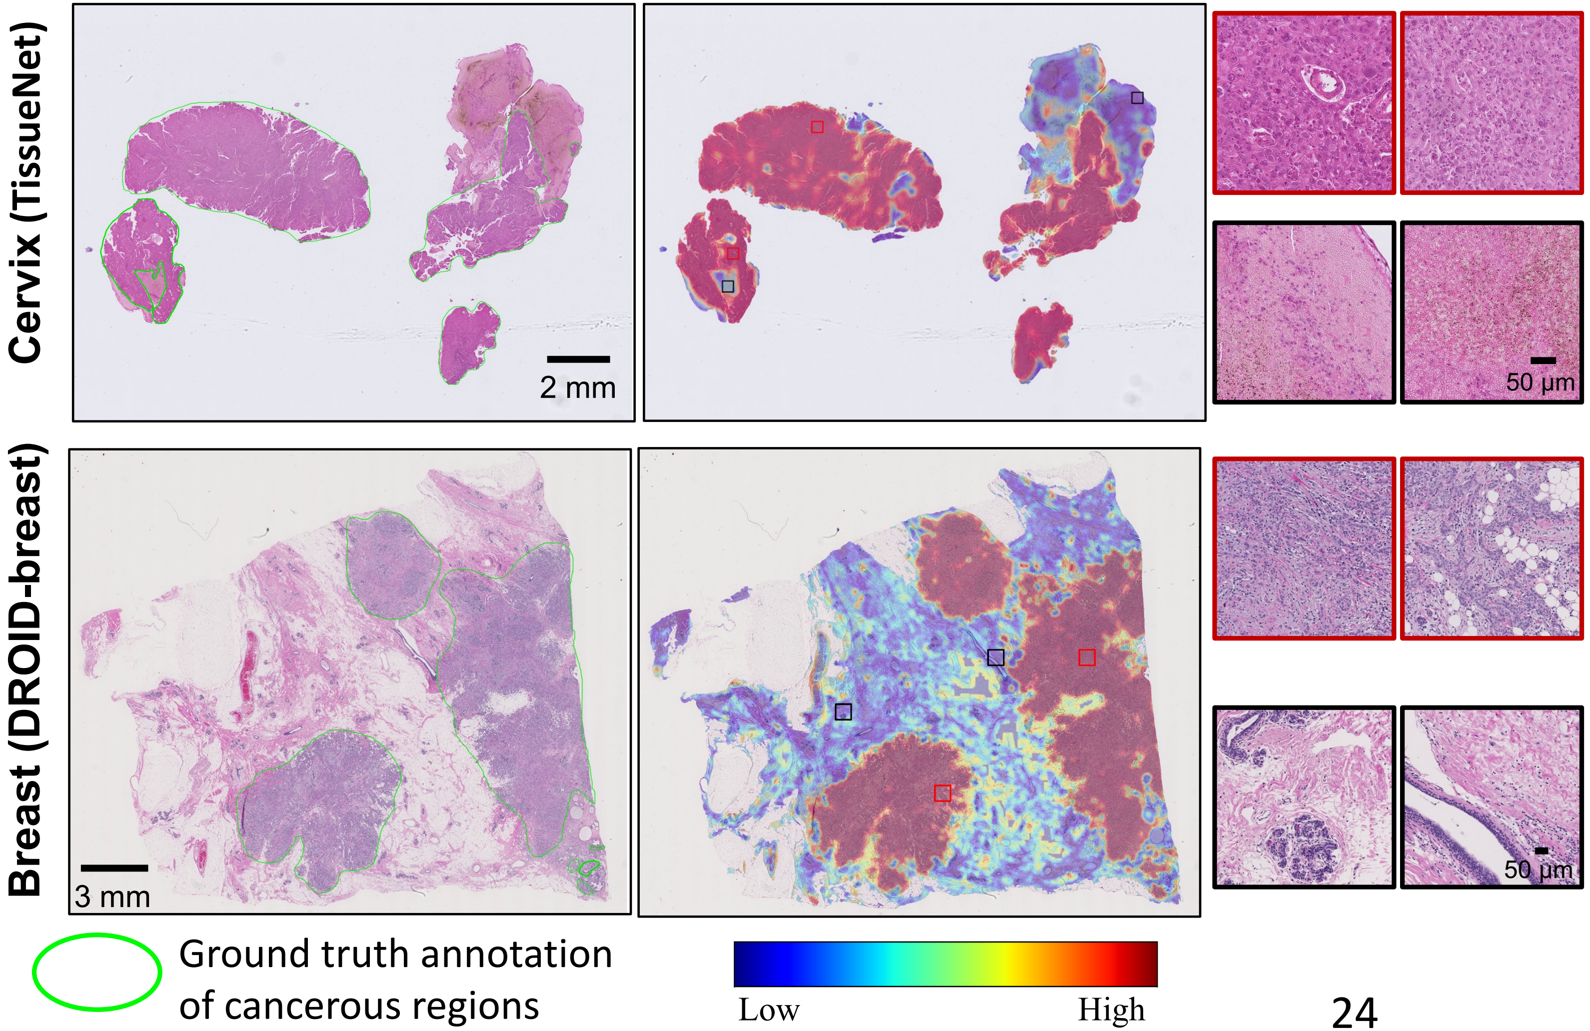

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

Cancer Origin Identification

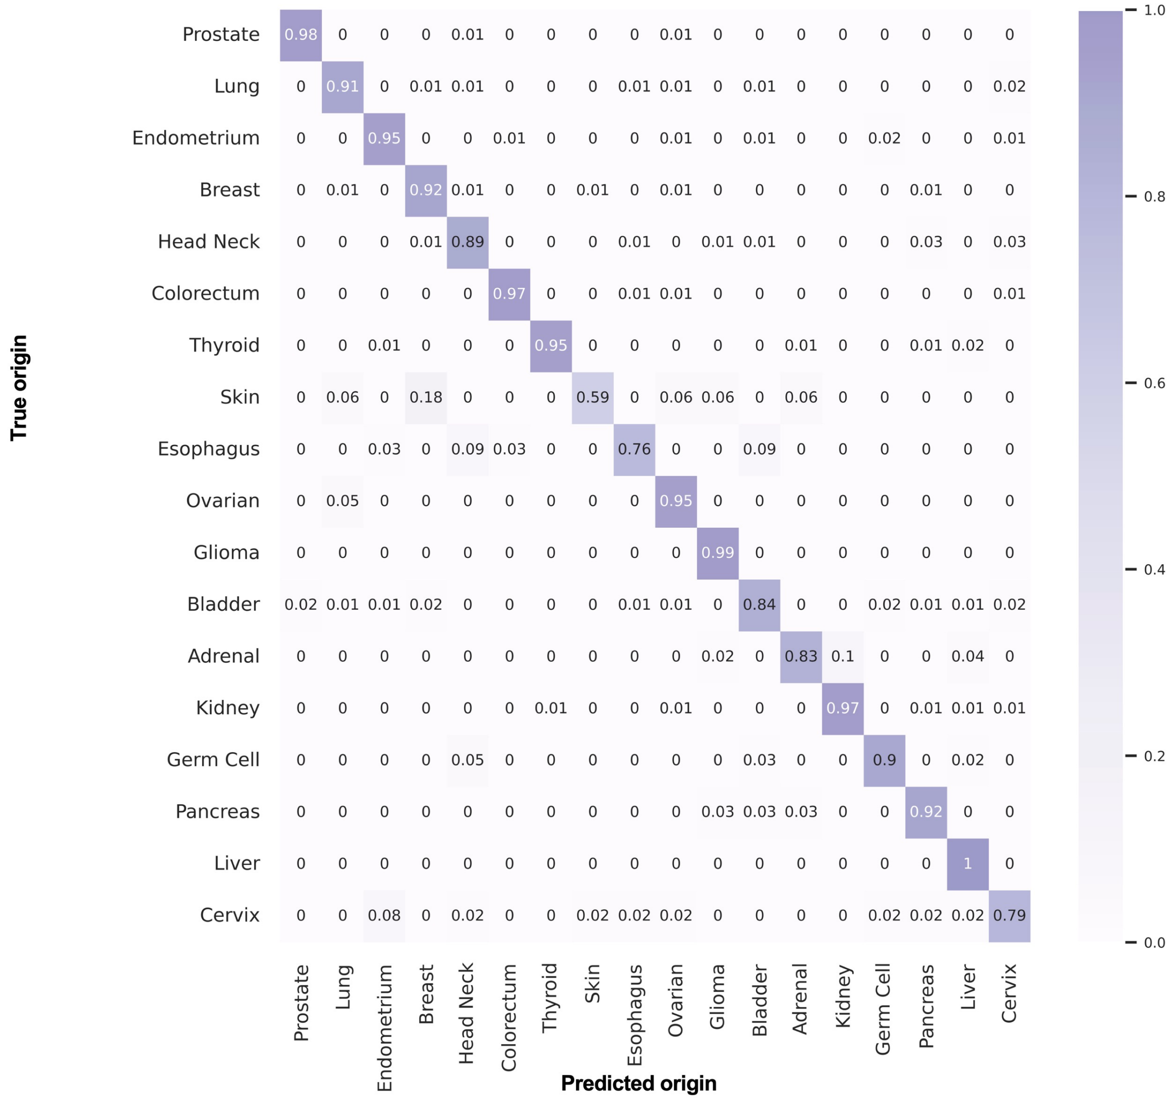

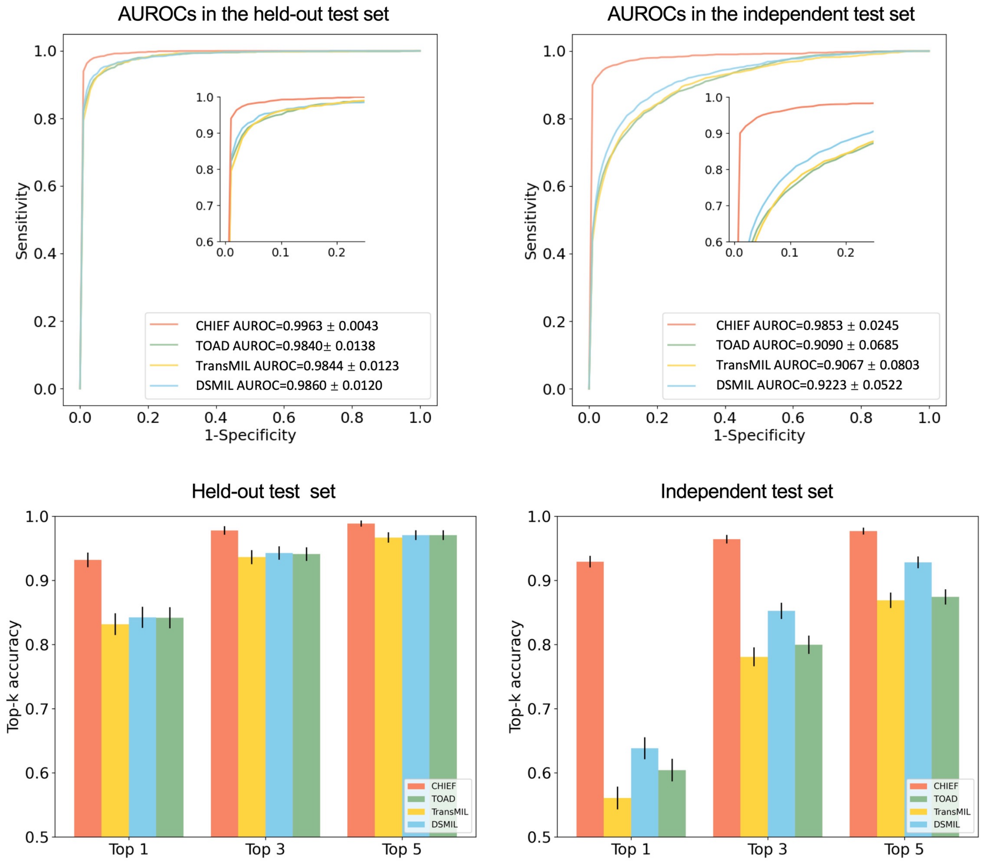

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

- Is there ***hidden*** information in histopathology images?
  - e.g., **genomic variations?**

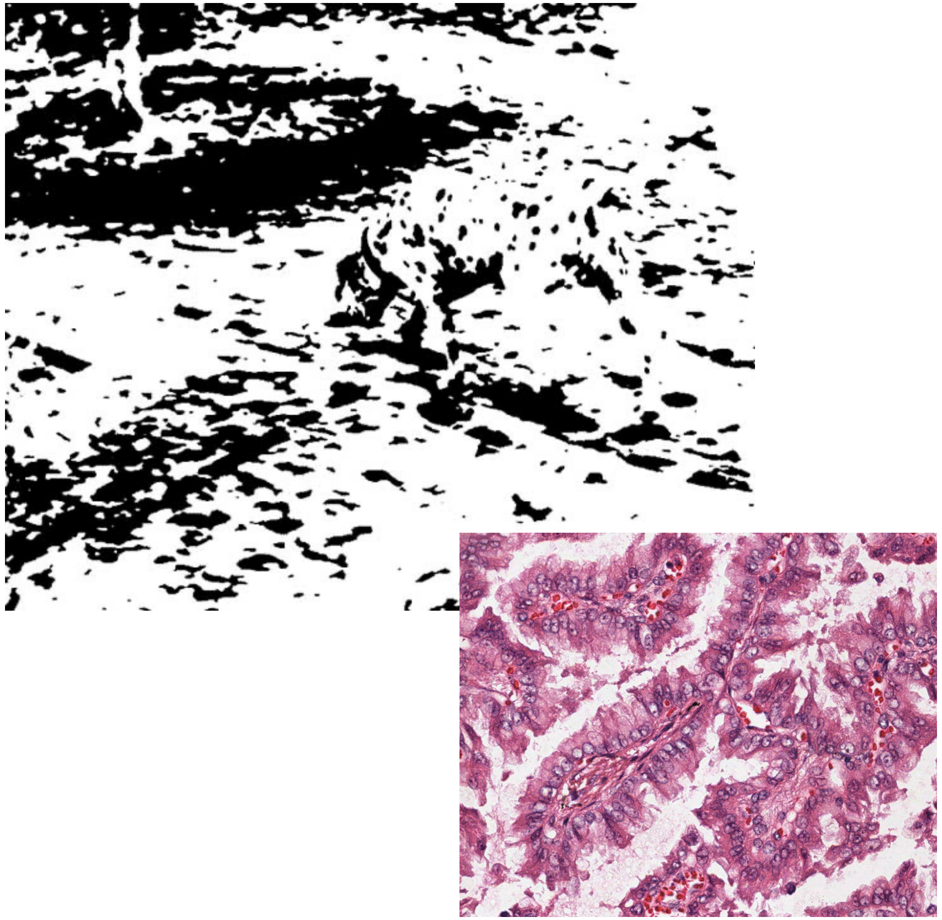

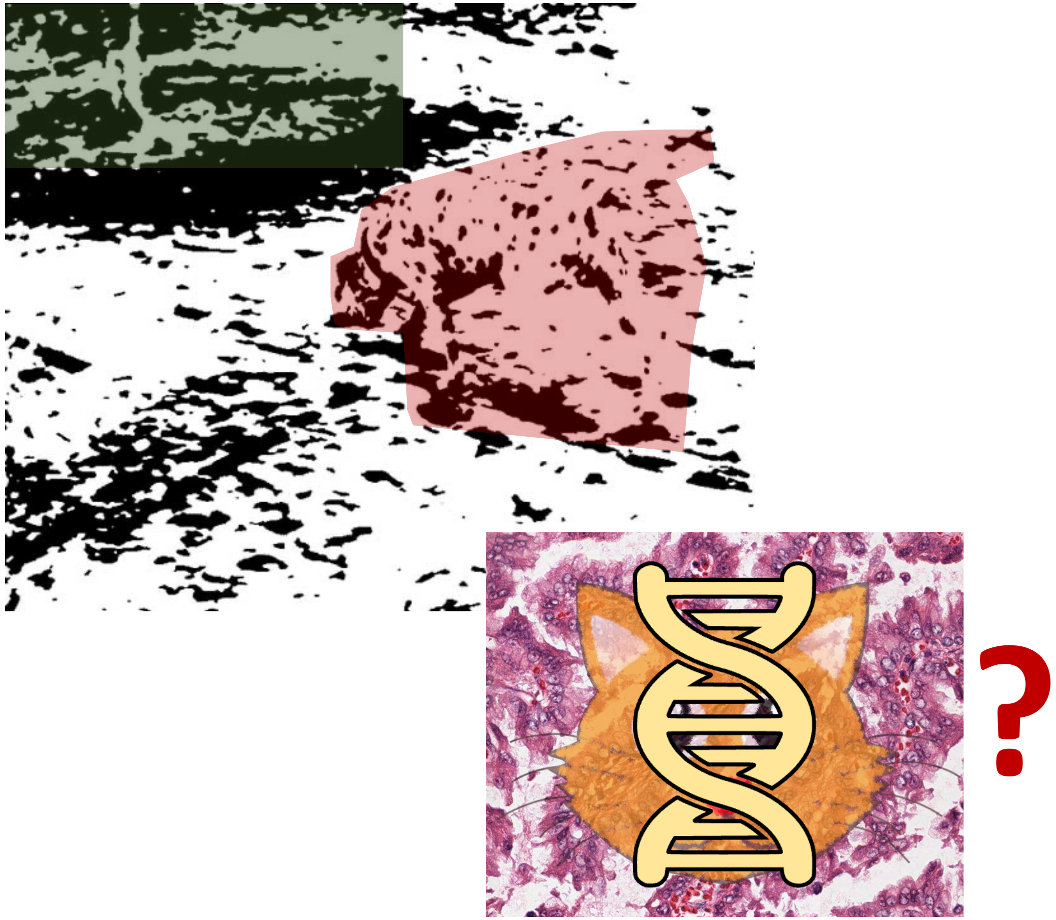

Image from: Gregory RL. *Phil. Trans. R. Soc. B* 2005; 360,1231–1251.

Predicting the Mutation Statuses of Clinically Important Genes

Prevalent mutations

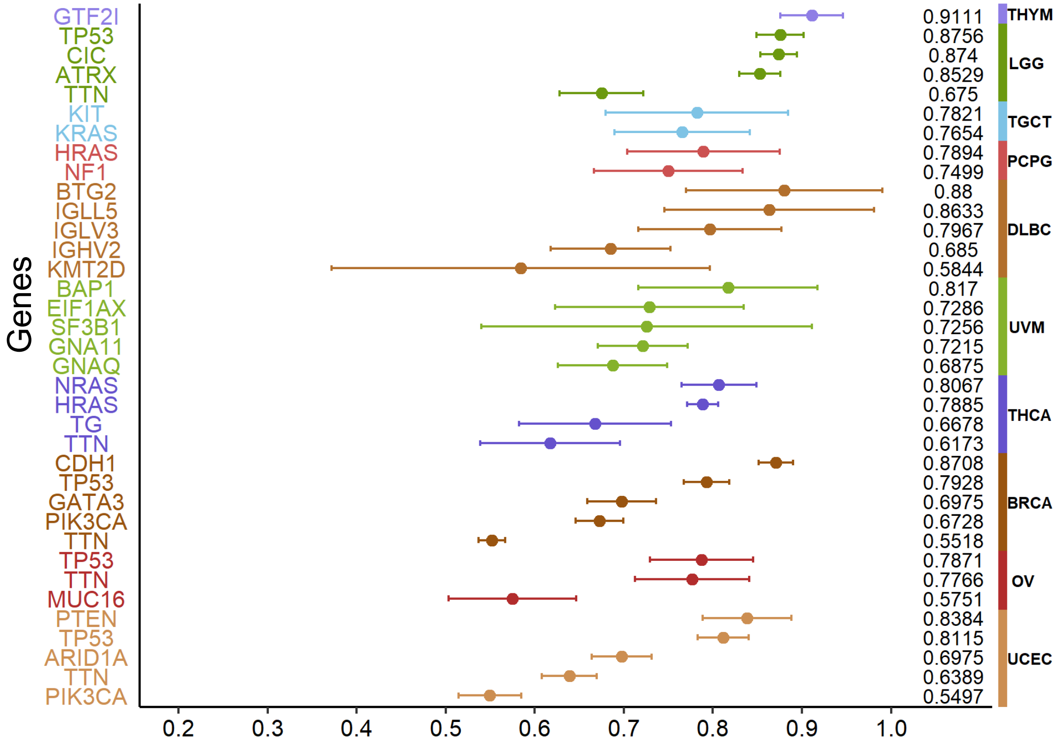

Mutations related to targeted therapies

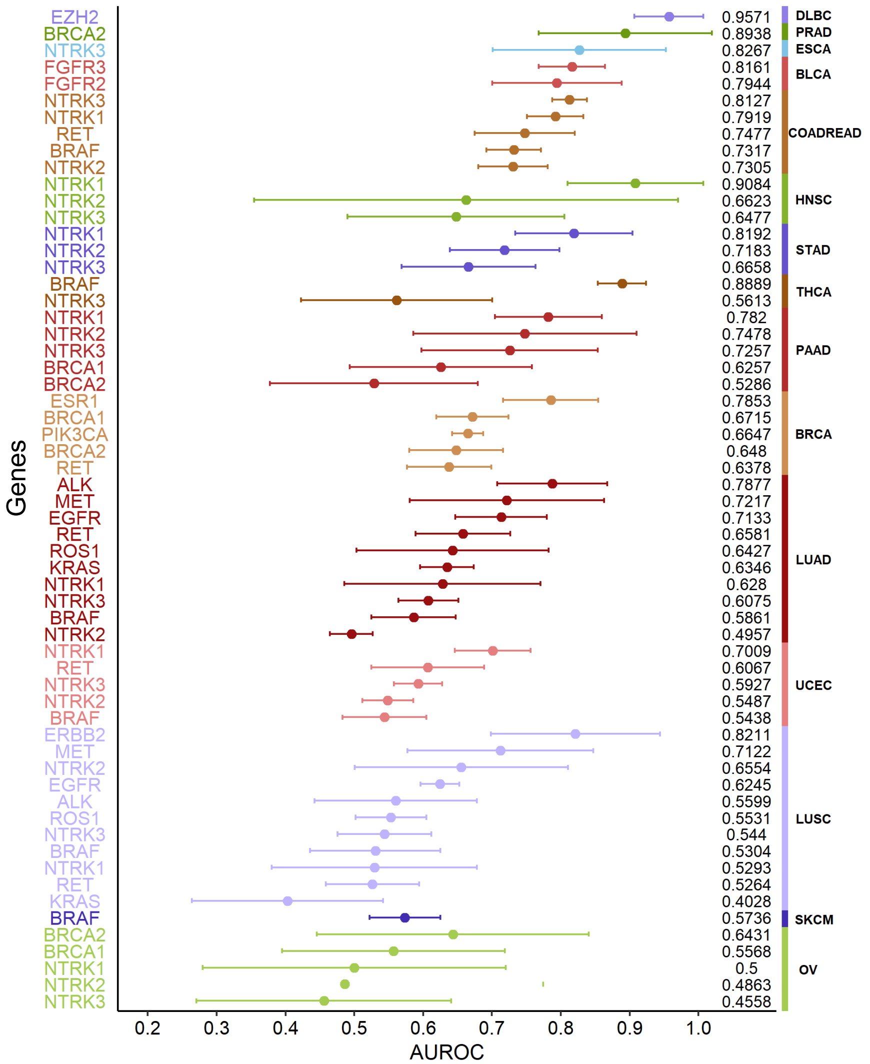

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

Identifying Molecular Profiles Related to Responses to Immune Checkpoint Blockade

Microsatellite instability (MSI) prediction for colorectal cancers:

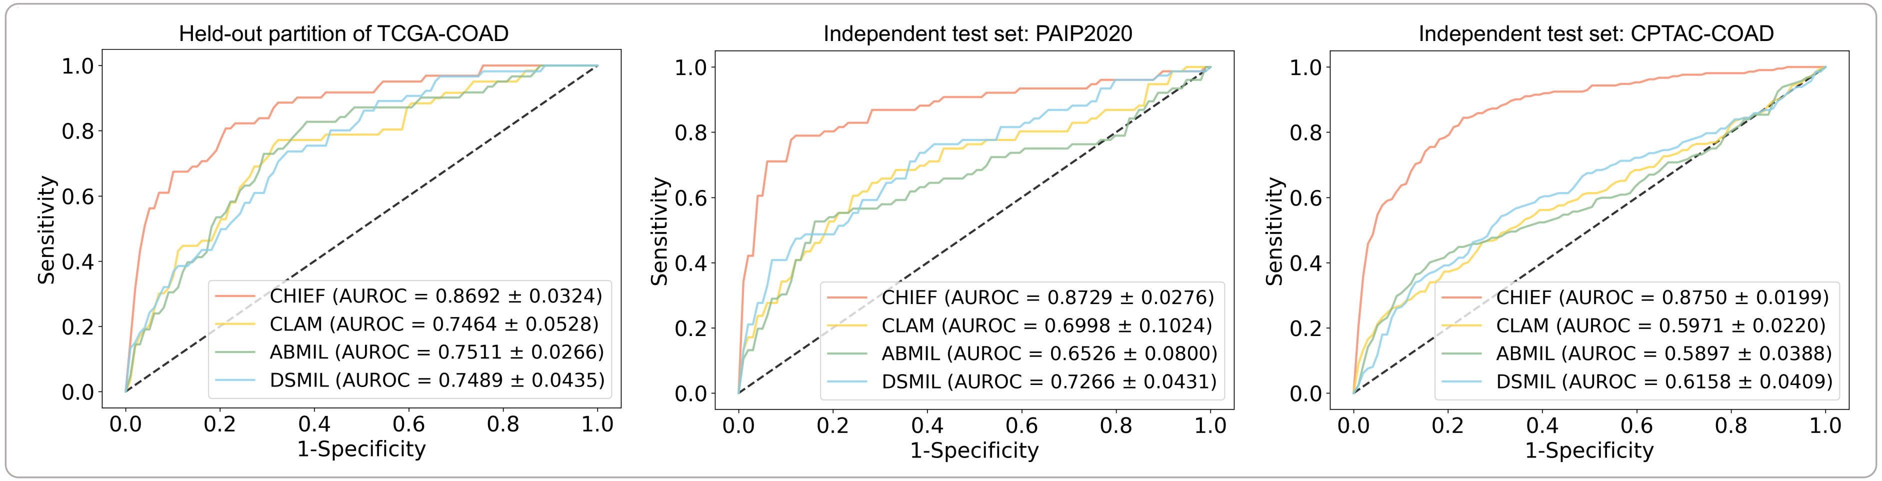

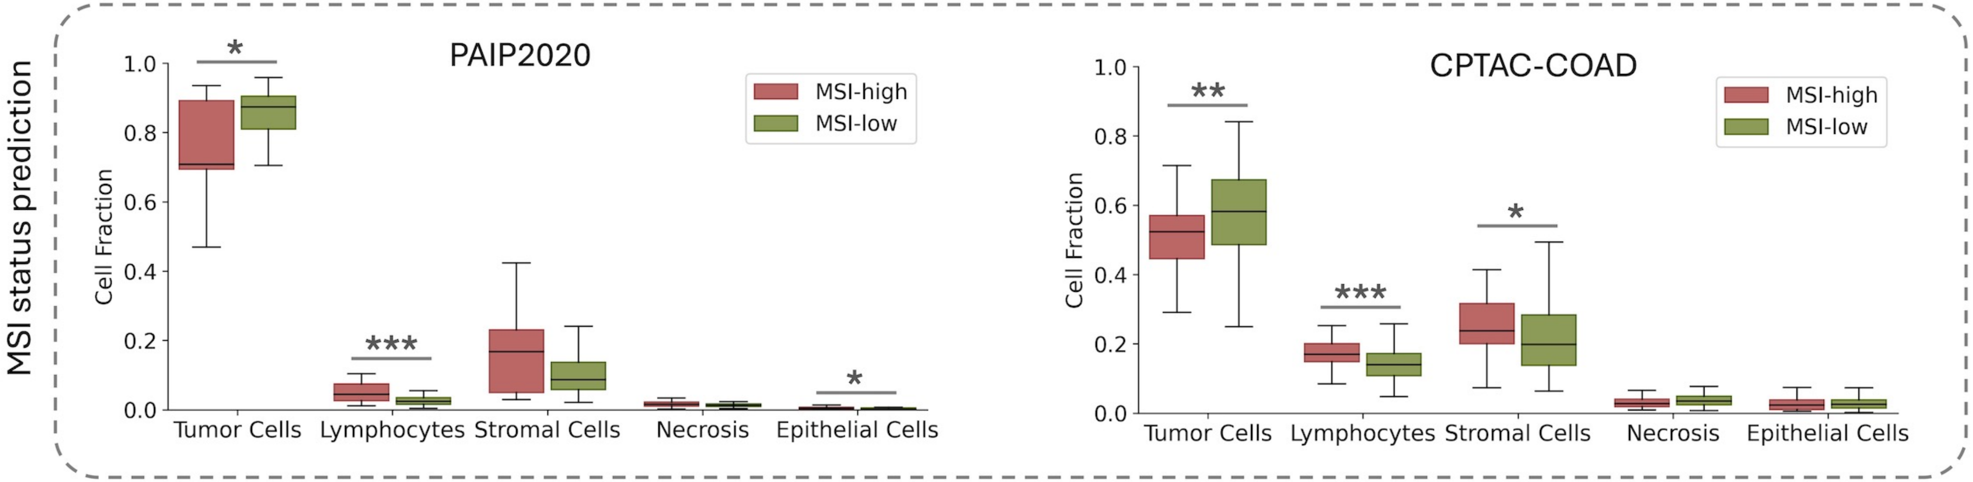

* P-value < 0.05
** P-value < 0.01
*** P-value < 0.001

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

CHIEF Predicts Cancer Patients' Survival Outcomes

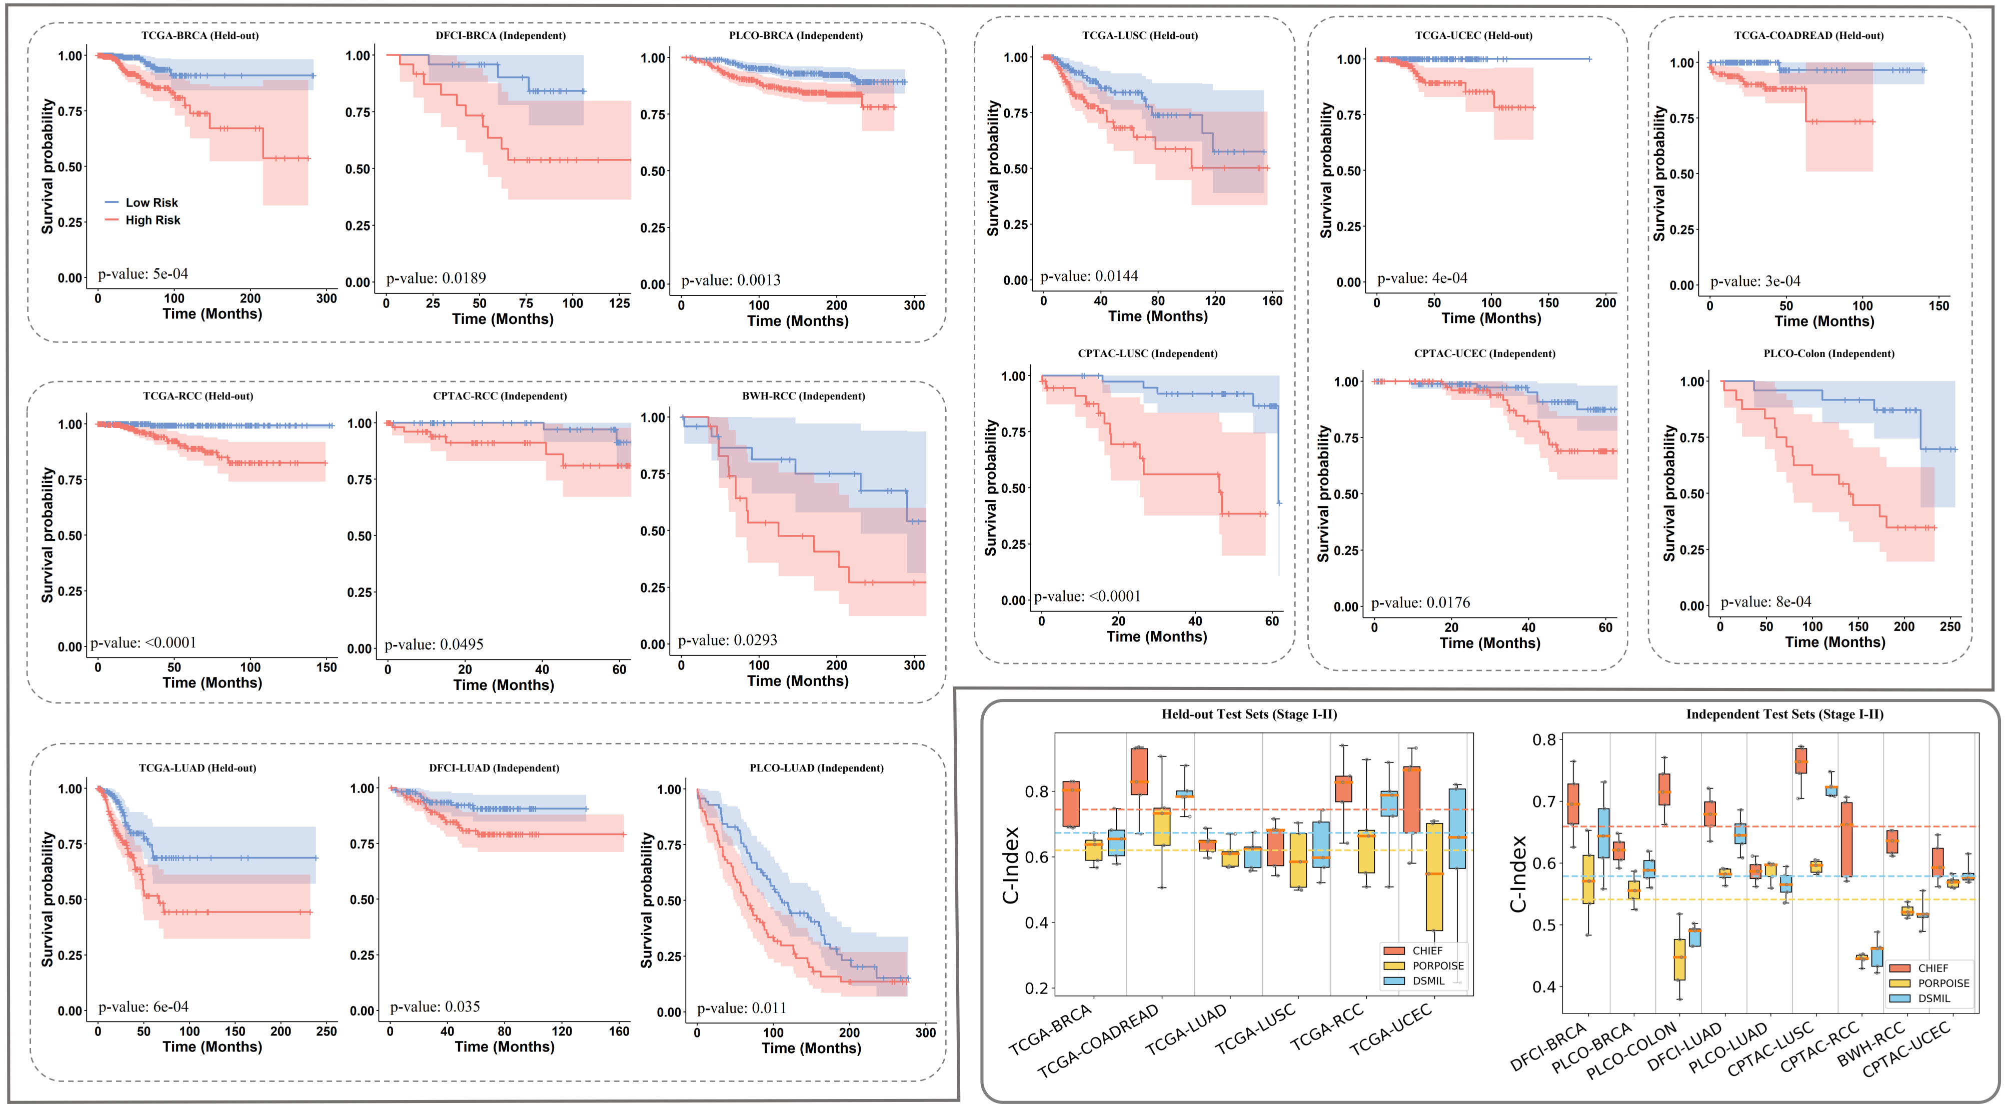

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

TMENT OF
nedical Informatics

Tumor Microenvironment Patterns Associated with Survival Outcomes

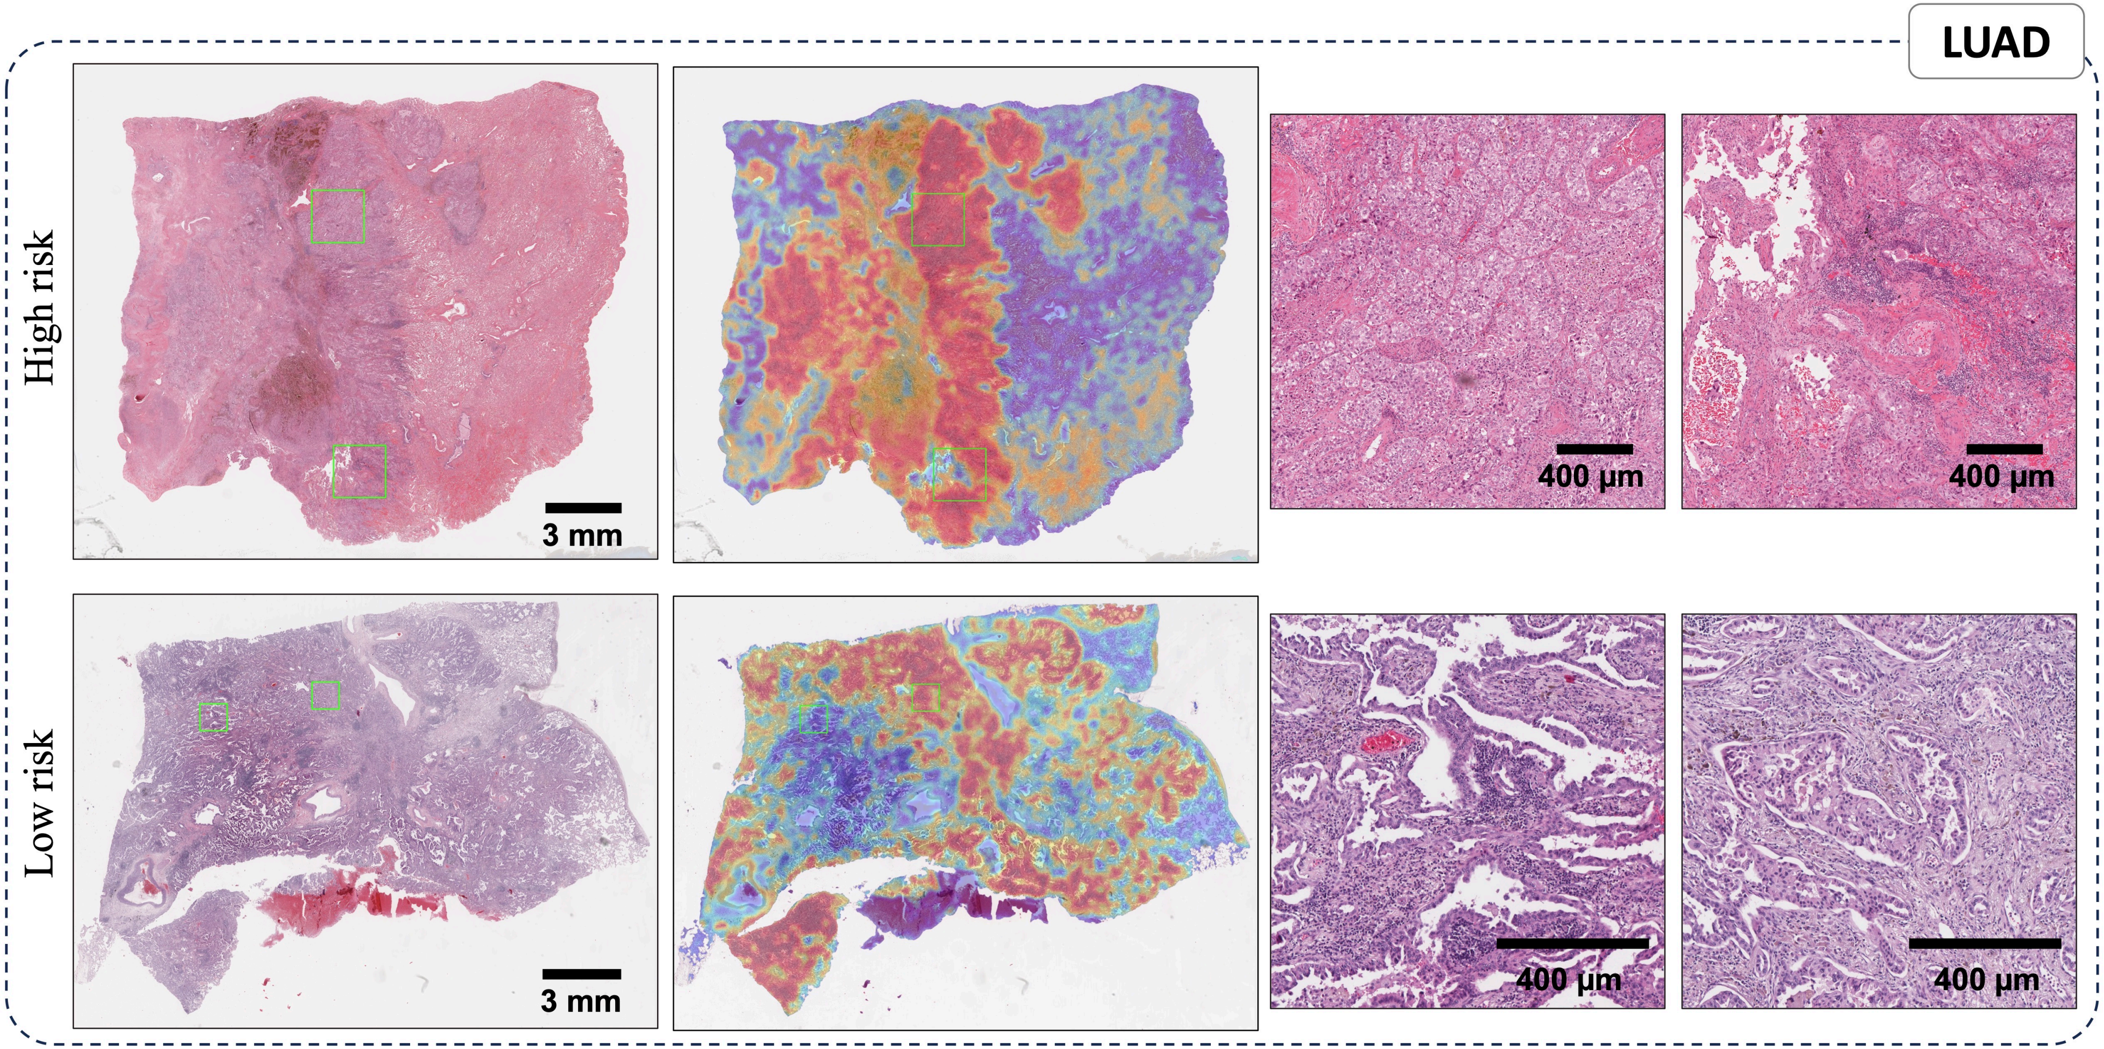

Wang X et al. *Nature*. 2024 Oct;634(8035):970-978.

Multimodal transformer with Unified maSKed modeling (MUSK)

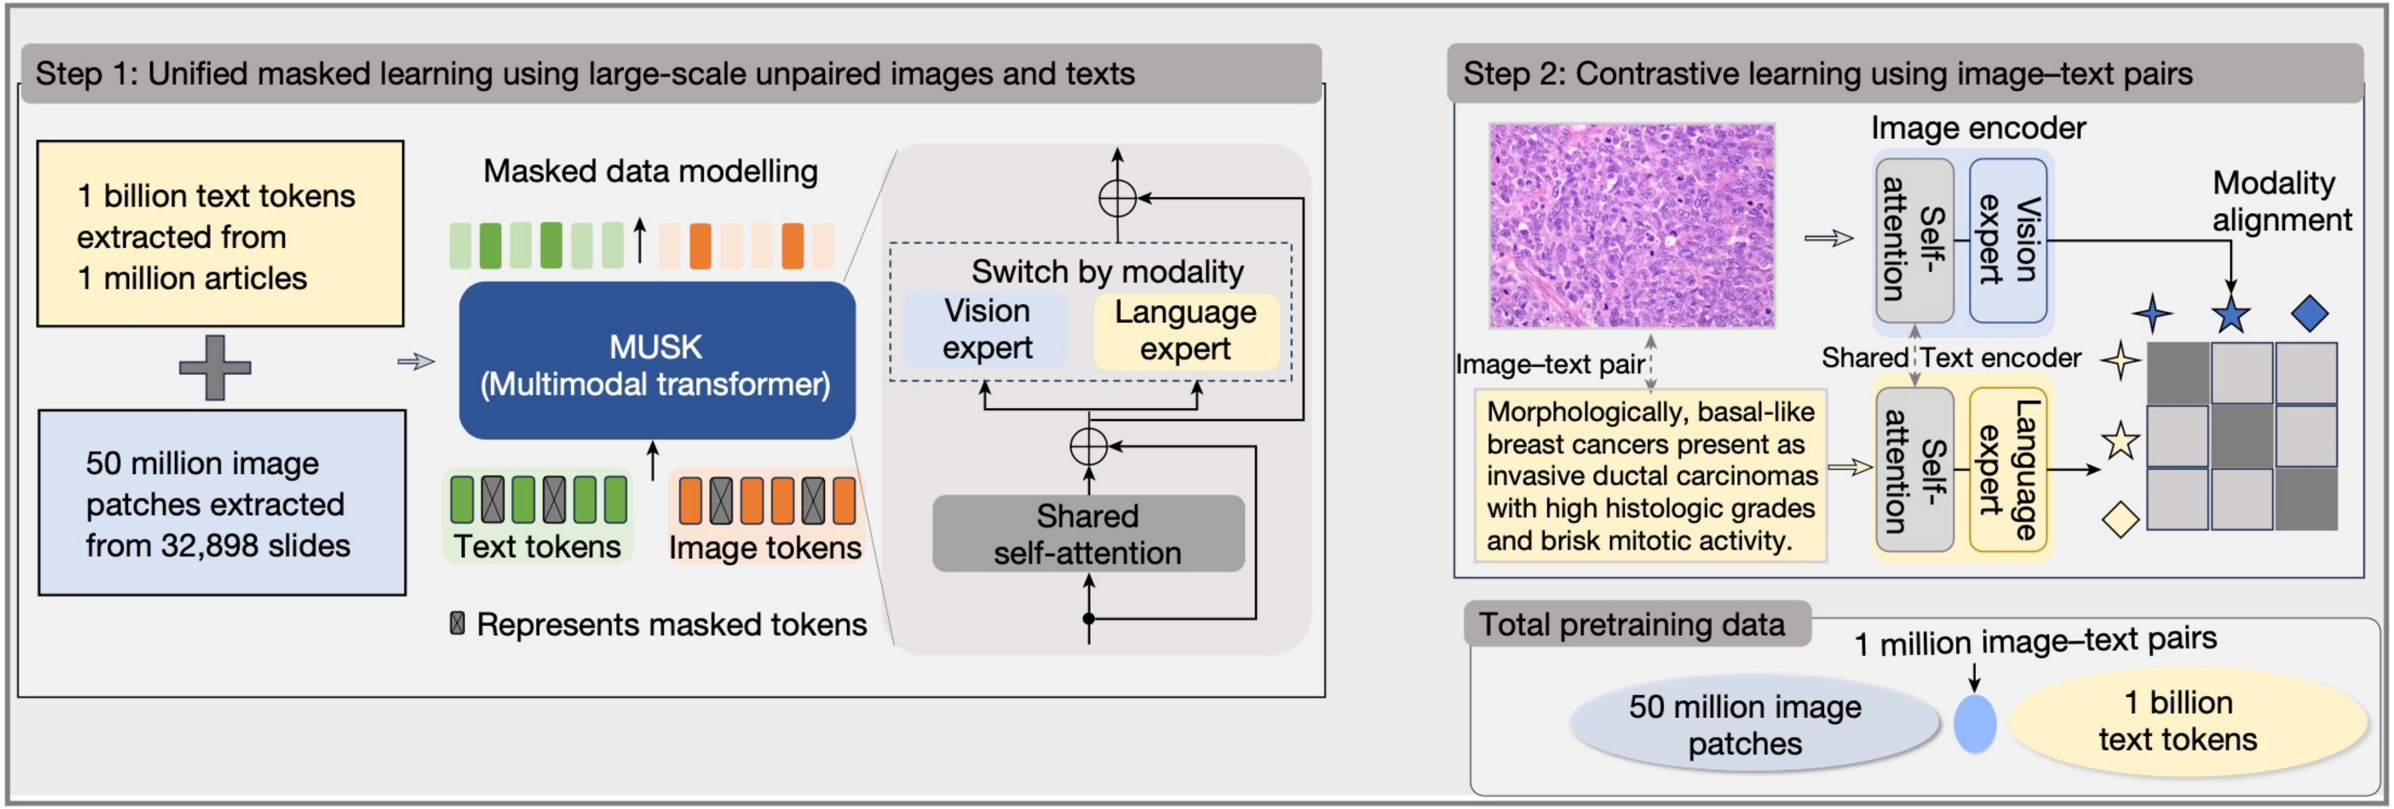

Xiang J et al. *Nature*. 2025 Feb;638(8051):769-778.

Prov-GigaPath

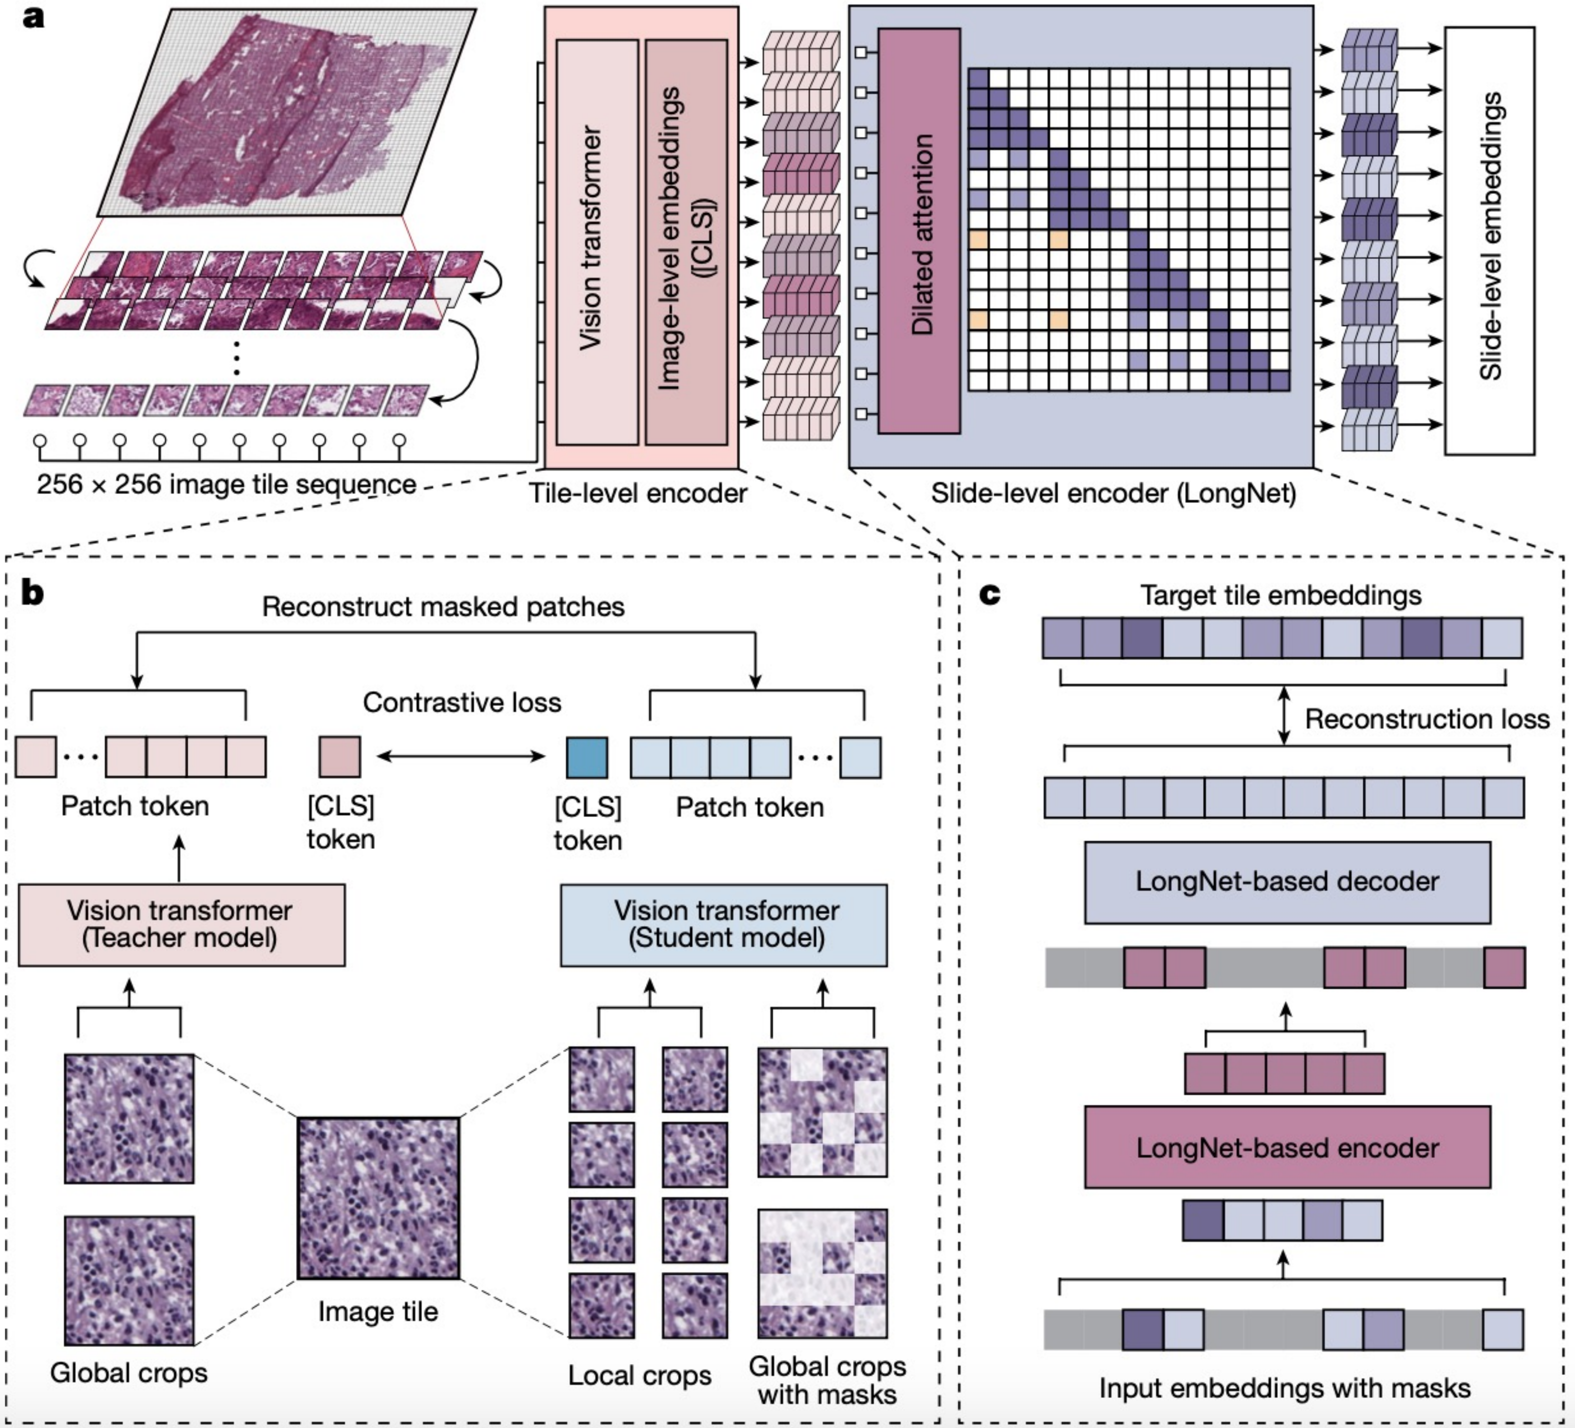

Xu H et al. *Nature*. 2024 Jun 6;630(8015):181-8.

Interpreting Deep Learning Models

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Approaches for Model Interpretation

- Visualize image patches that maximally activated neurons
- Visualize the “feature” space
- Visualize the convolution filters
- Occlusion
- Attention maps
- “Deconv”
- Optimize to image
- …

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Method 1: Visualize Image Patches that Activate the Selected Neurons

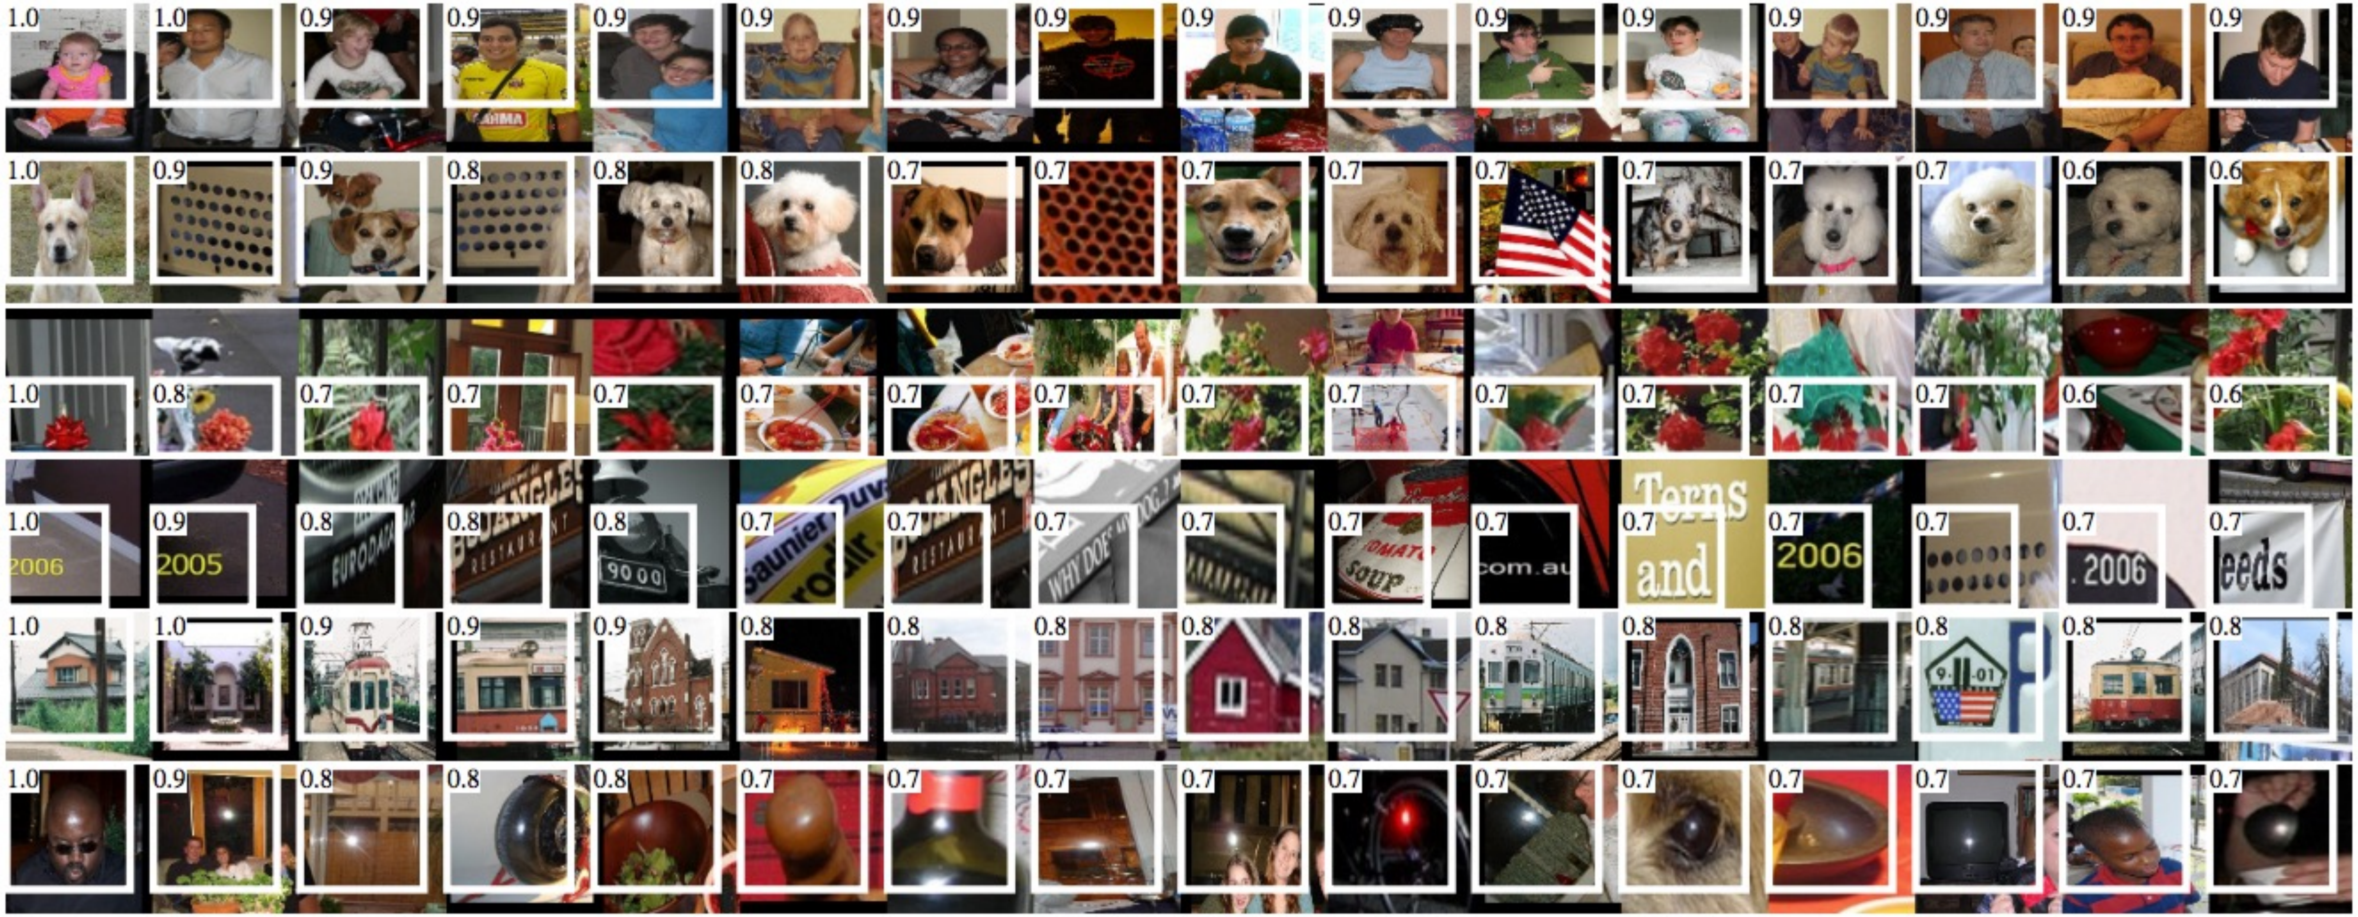

Girshick R et al. *CVPR*. 2014:580-587.

Method 2: Visualize the “Feature” Space

- Treat the values in the fully-connected layer as “features”

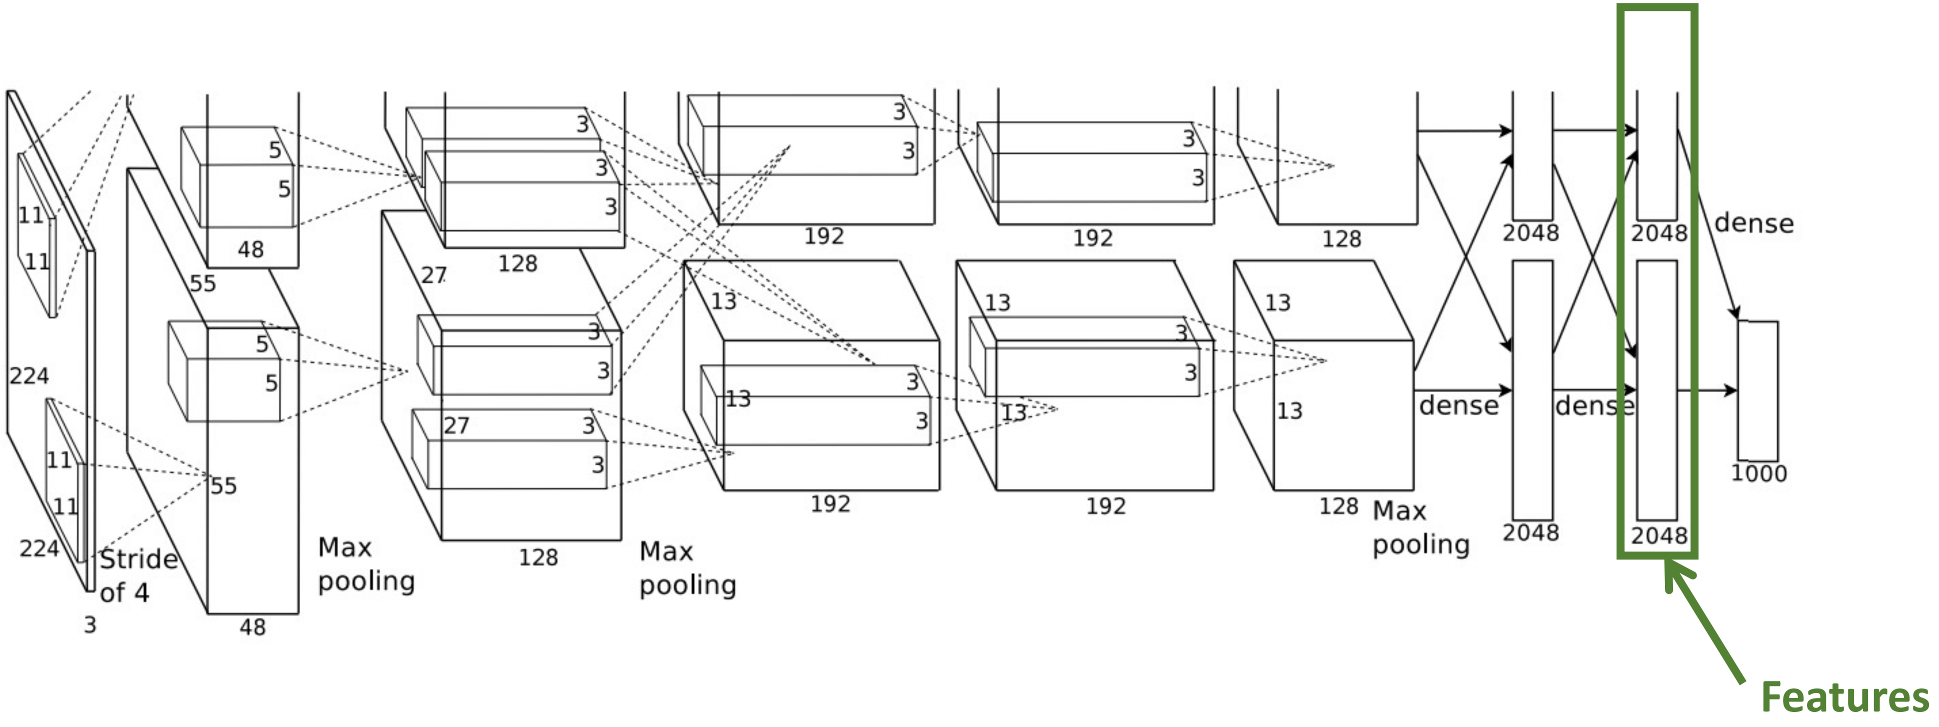

Krizhevsky A et al. Advances in Neural Information Processing Systems 2012:1097-1105.

Method 2: Visualize the "Feature" Space

- Principal component analysis
  - Convert a set of features into a set of linearly uncorrelated variables
- t-SNE (t-distributed stochastic neighbor embedding)
  - Similar objects have a high probability of being picked as neighbors

t-SNE

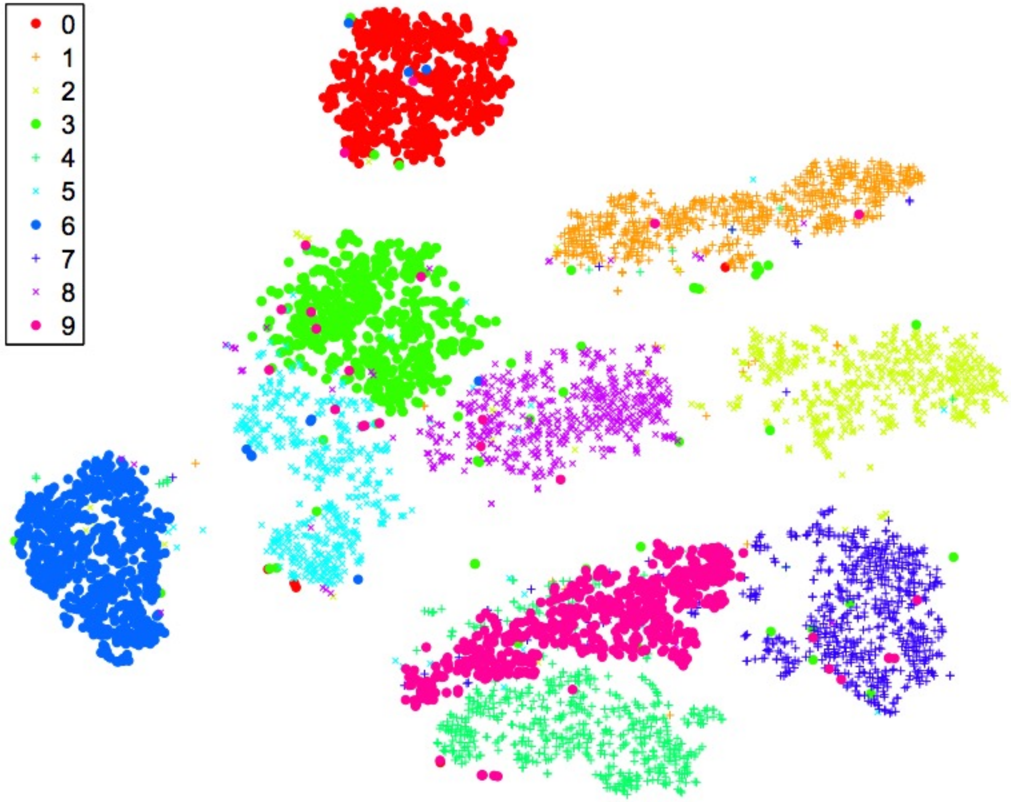

Van Der Maaten L. *JMLR*. 2014 Jan 1;15(1):3221-45.

Method 2: Visualize the "Feature" Space

- U-Map (Uniform Manifold Approximation and Projection)
  - Assumptions
    - The data is uniformly distributed on a Riemannian manifold
    - The Riemannian metric can be approximated as constant locally
    - The manifold is locally connected
  - Preserves more global structure
  - Faster than t-SNE

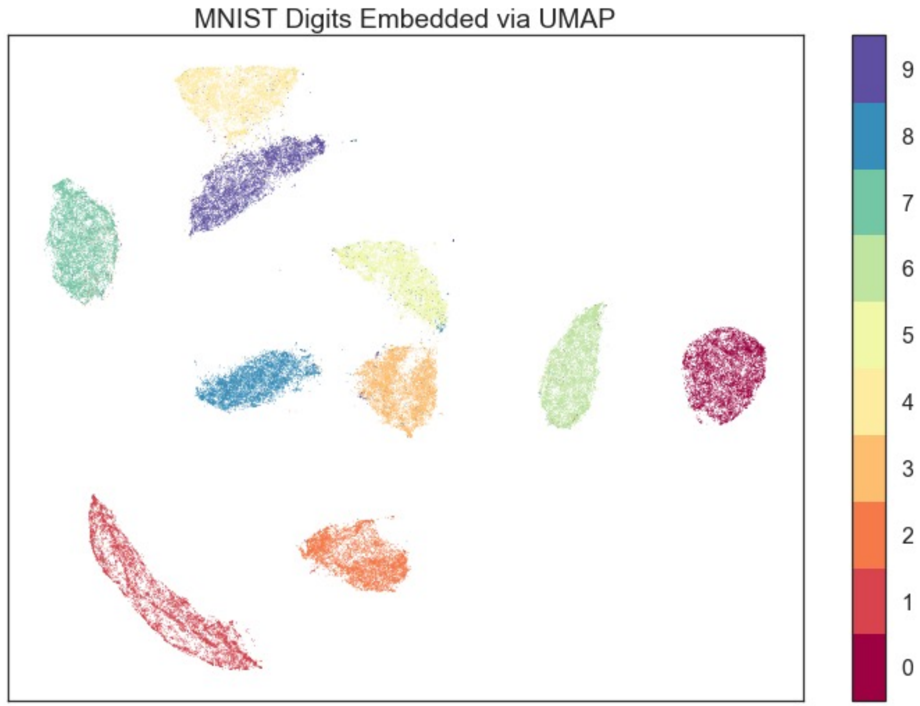

McInnes L et al. arXiv preprint arXiv:1802.03426. 2018 Feb 9.
Figure credit: Leland McInnes

Method 3: Visualize the Convolutional Filters

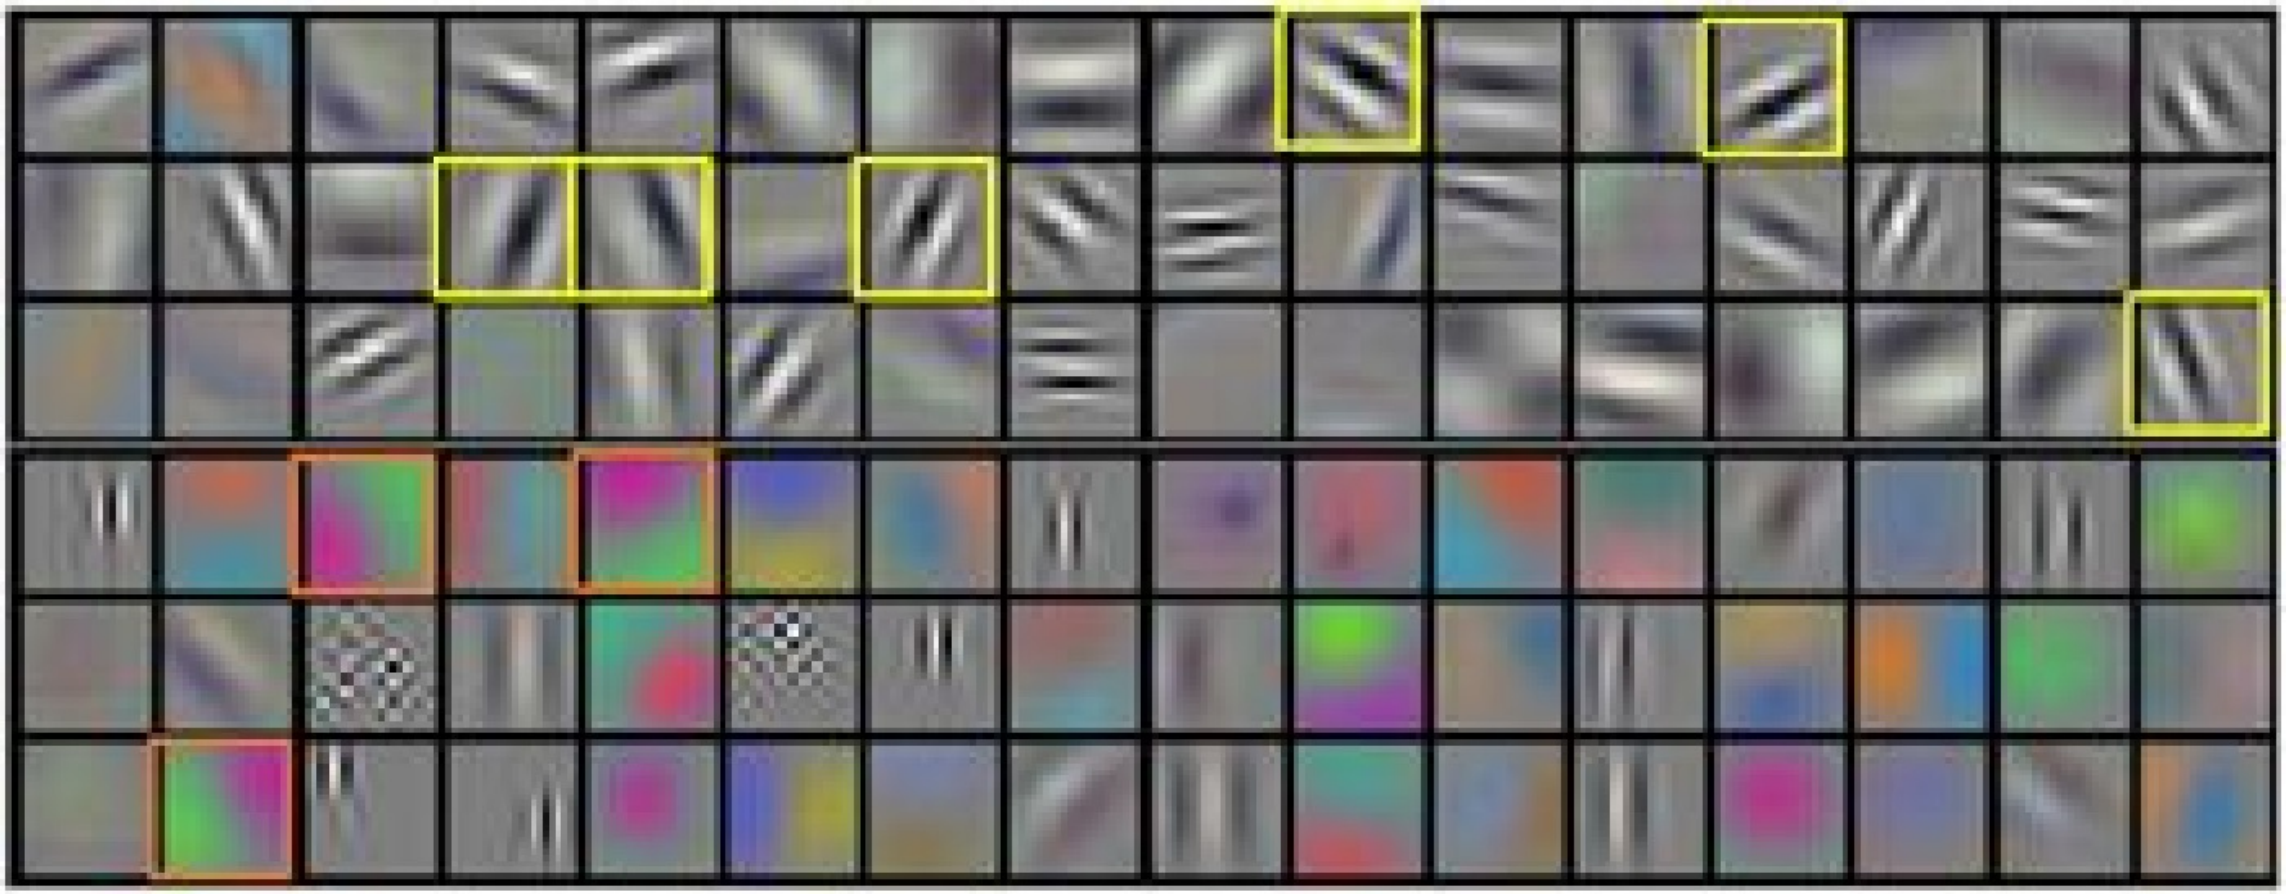

- Show the raw weights of the filters

Luan S et al. *IEEE Transactions on Image Processing*. 2018 Sep;27(9):4357-66.

Method 4: Occlusion

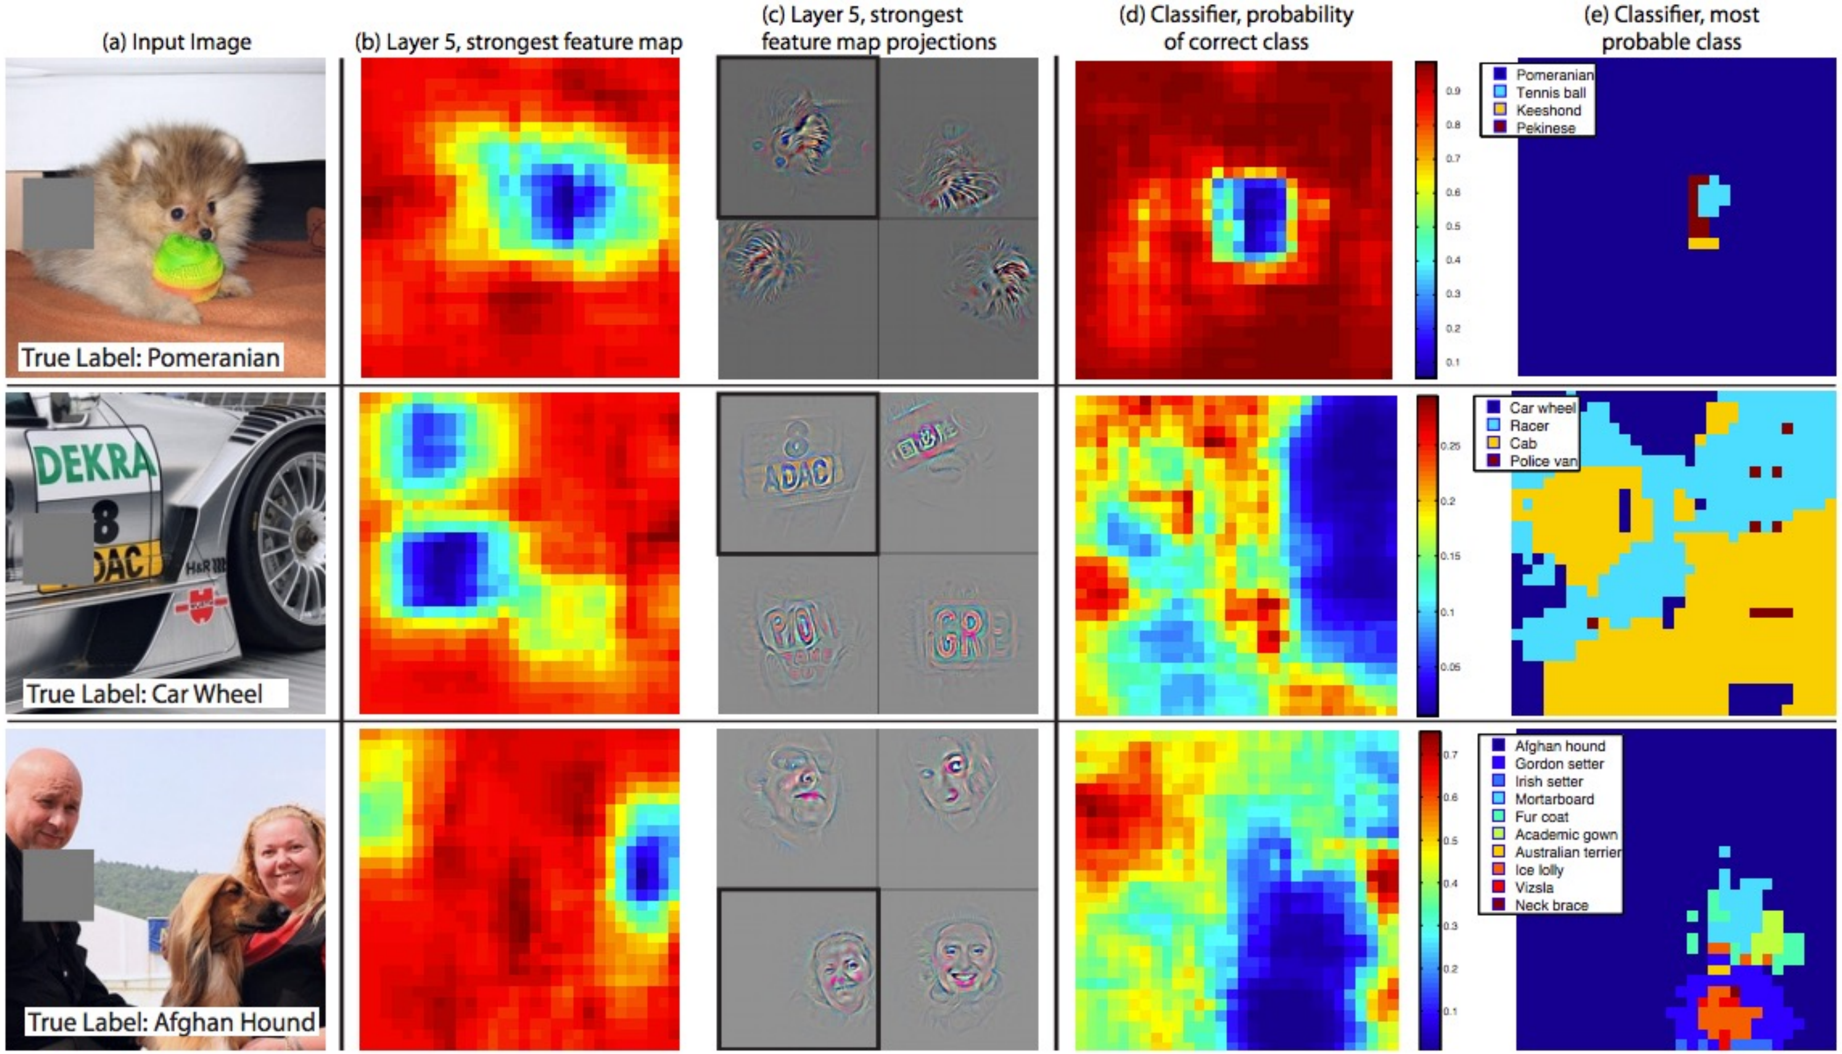

Zeiler MD and Fergus R. *ECCV*. 2014 Sep 6;818-833.

Method 5: Attention Maps

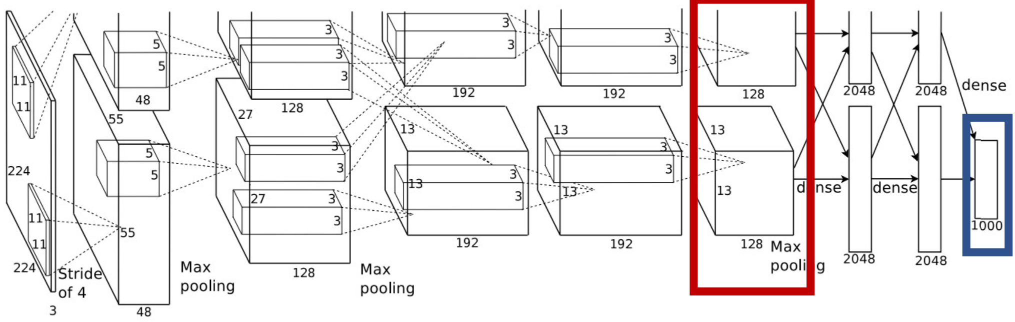

- Saliency map
  - Visualize how the output category would change if we tweak the input image pixels (i.e., visualize the gradient with respect to the **output category**)
- Class activation map
  - Visualize the gradient with respect to **the layer right before the fully connected layer**

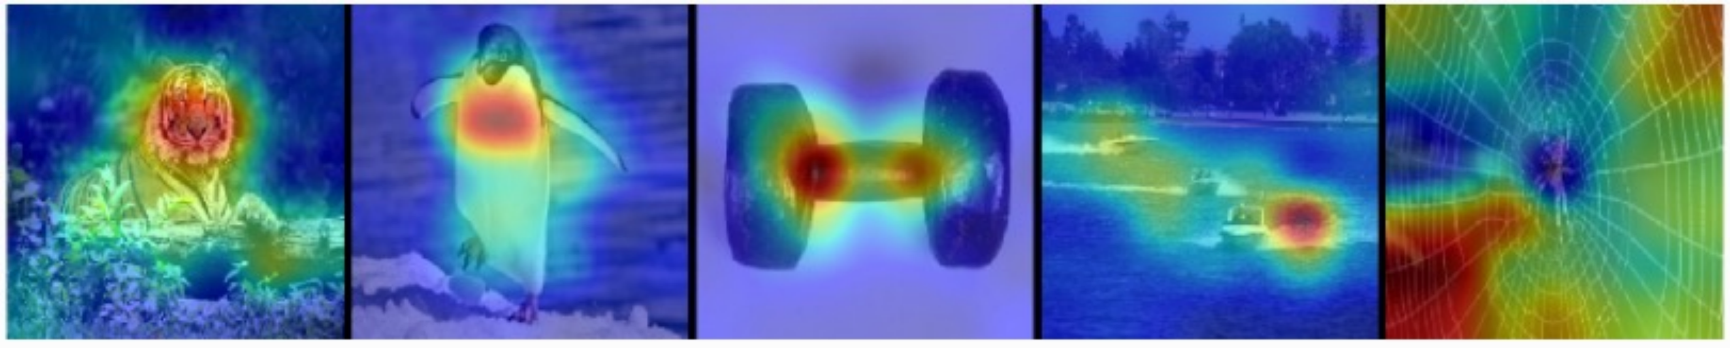

Simonyan K et al. arXiv preprint arXiv:1312.6034. 2013 Dec 20.
Image credit: Raghavendra Kotikalapudi

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Method 6: “Deconv”

- Feed image into a CNN
- Pick a layer, set the gradient of that layer to be [0, 0, …, 1, …, 0]
- Backprop to image

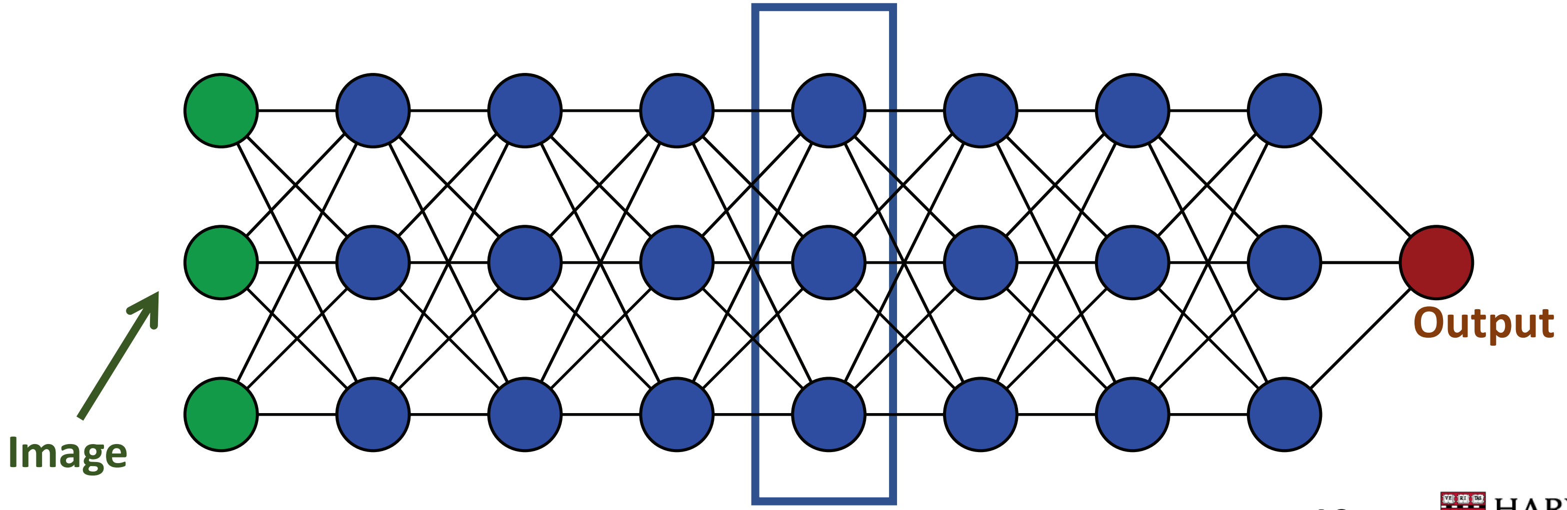

Zeiler MD and Fergus R. *ECCV*. 2014 Sep 6;818-833.

Method 6: “Deconv”

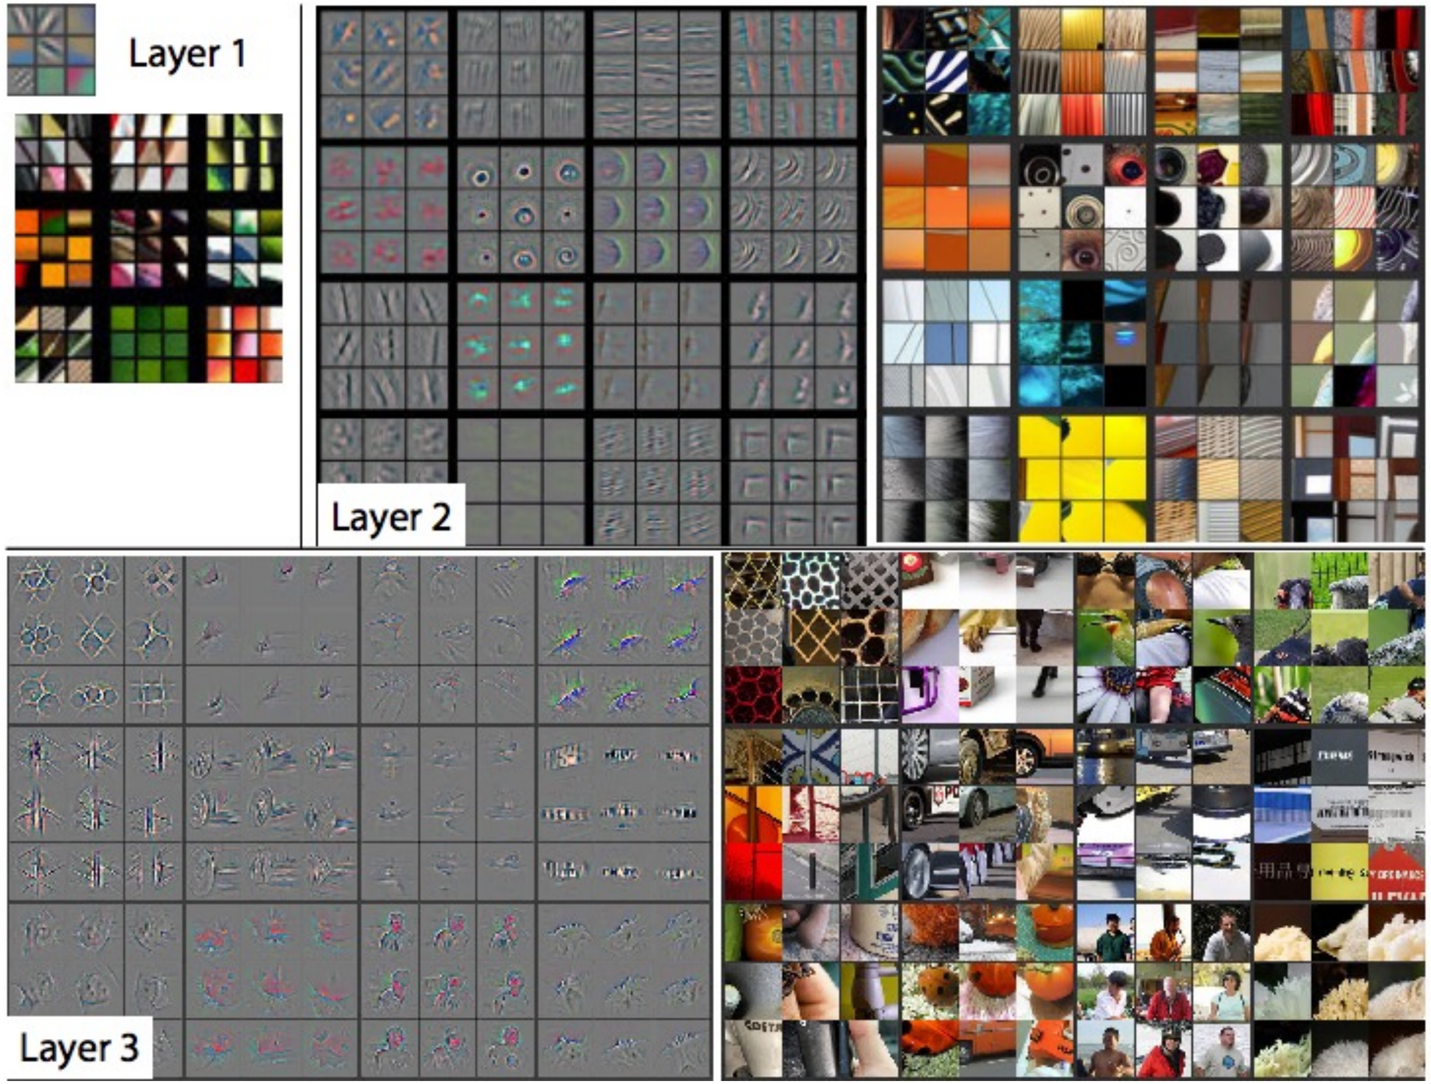

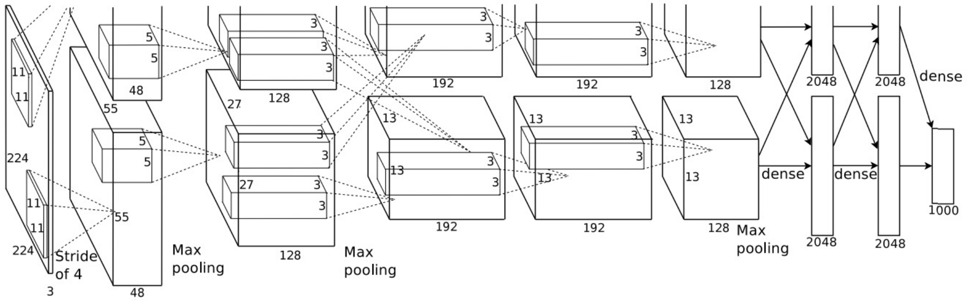

Image

Zeiler MD and Fergus R. *ECCV*. 2014 Sep 6;818-833.

Method 6: “Deconv”

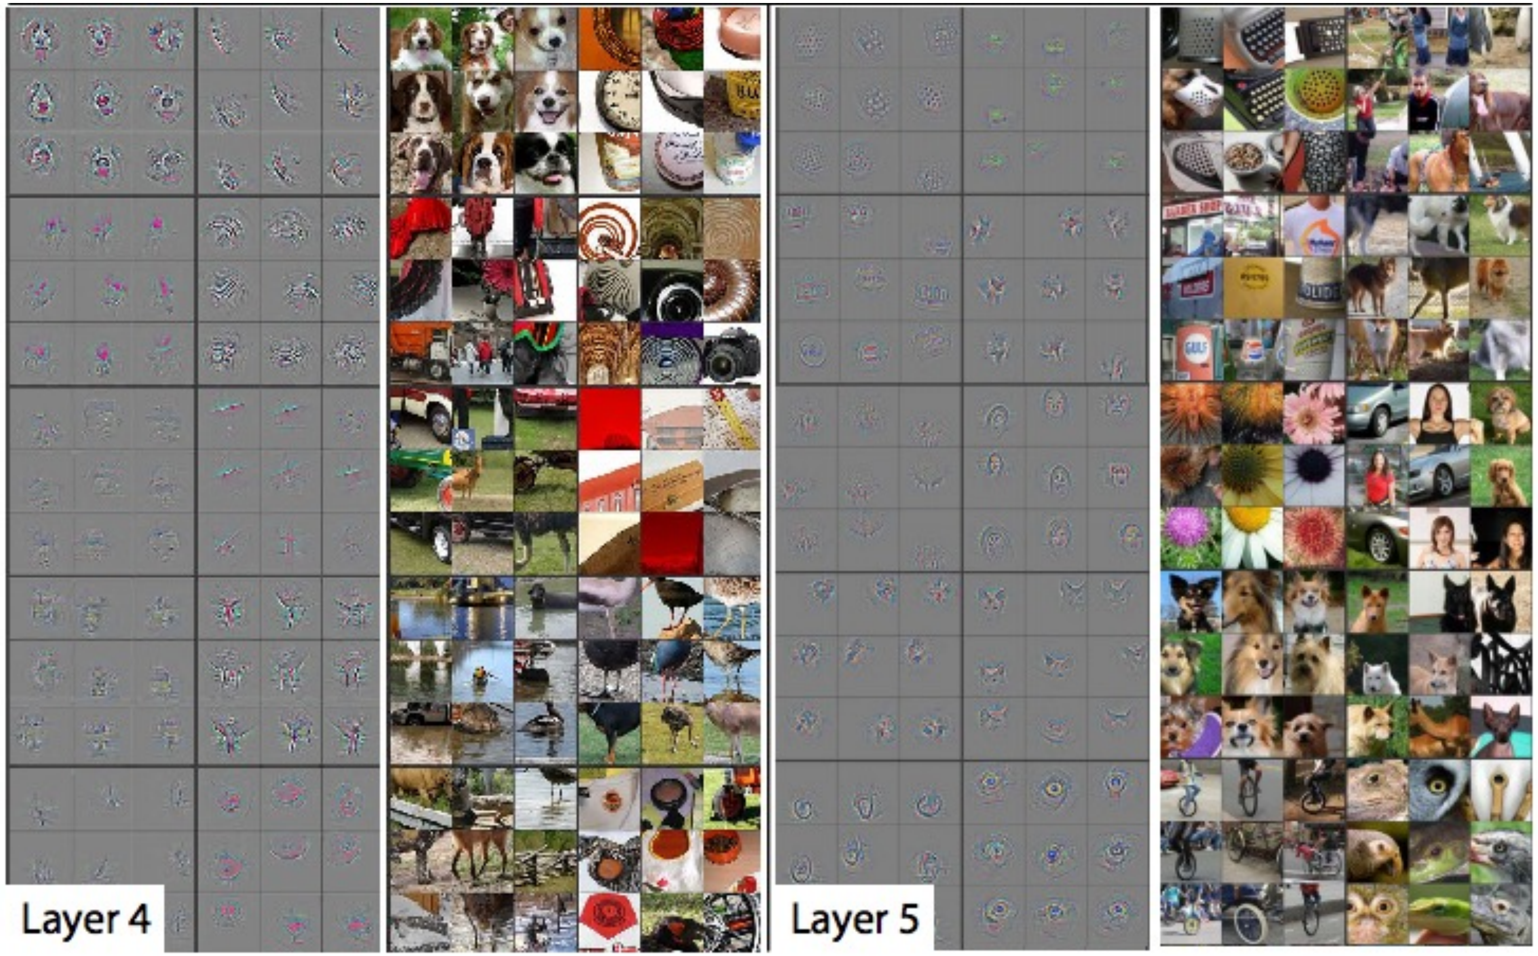

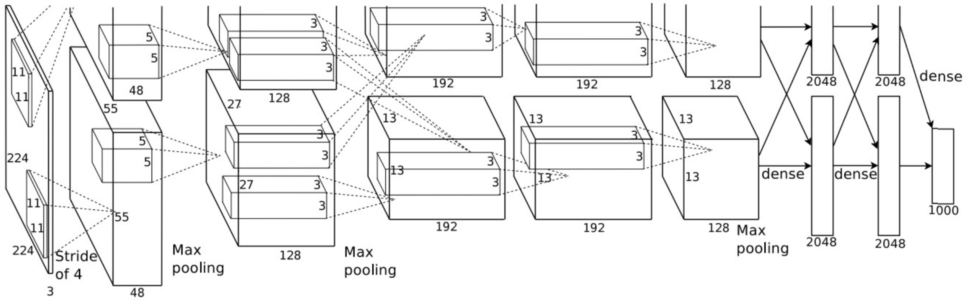

Image

Note: “Deconv” consists of a single backward pass

Zeiler MD and Fergus R. *ECCV*. 2014 Sep 6;818-833.

Method 7: "Optimize" the Image

- Start with a neutral image
- Do
    - Set the gradient of the score vector to be [0, 0, …, 1, …, 0]
    - Backprop to the image
    - Forward the image through the network

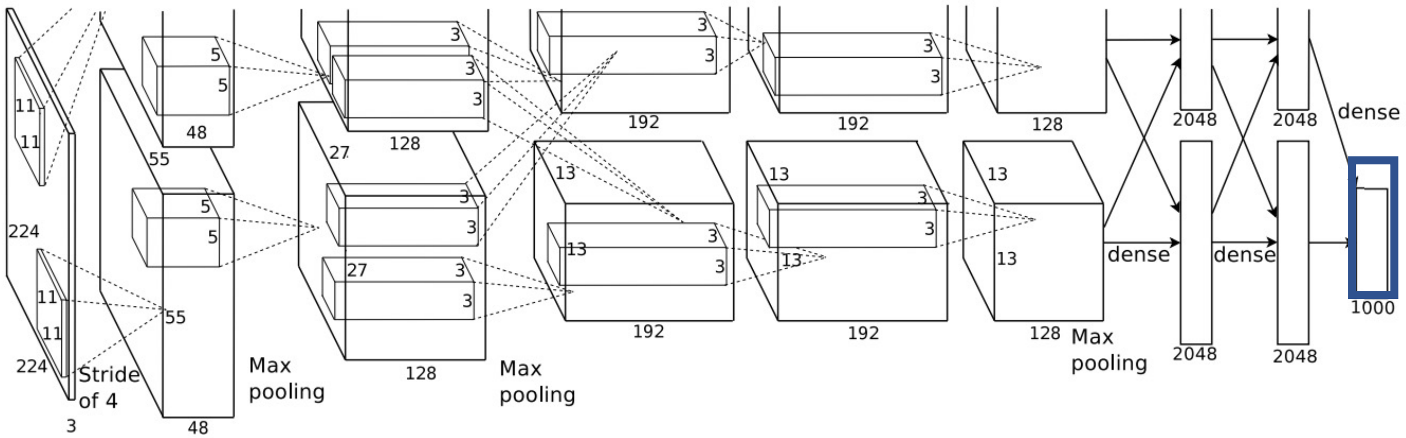

Image

Simonyan K et al. arXiv preprint arXiv:1312.6034. 2013 Dec 20.

Method 7: "Optimize" the Image

- Start with a neutral image
- Do
  - Set the gradient of the score vector to be [0, 0, …, 1, …, 0]
  - Backprop to the image
  - Forward the image through the network

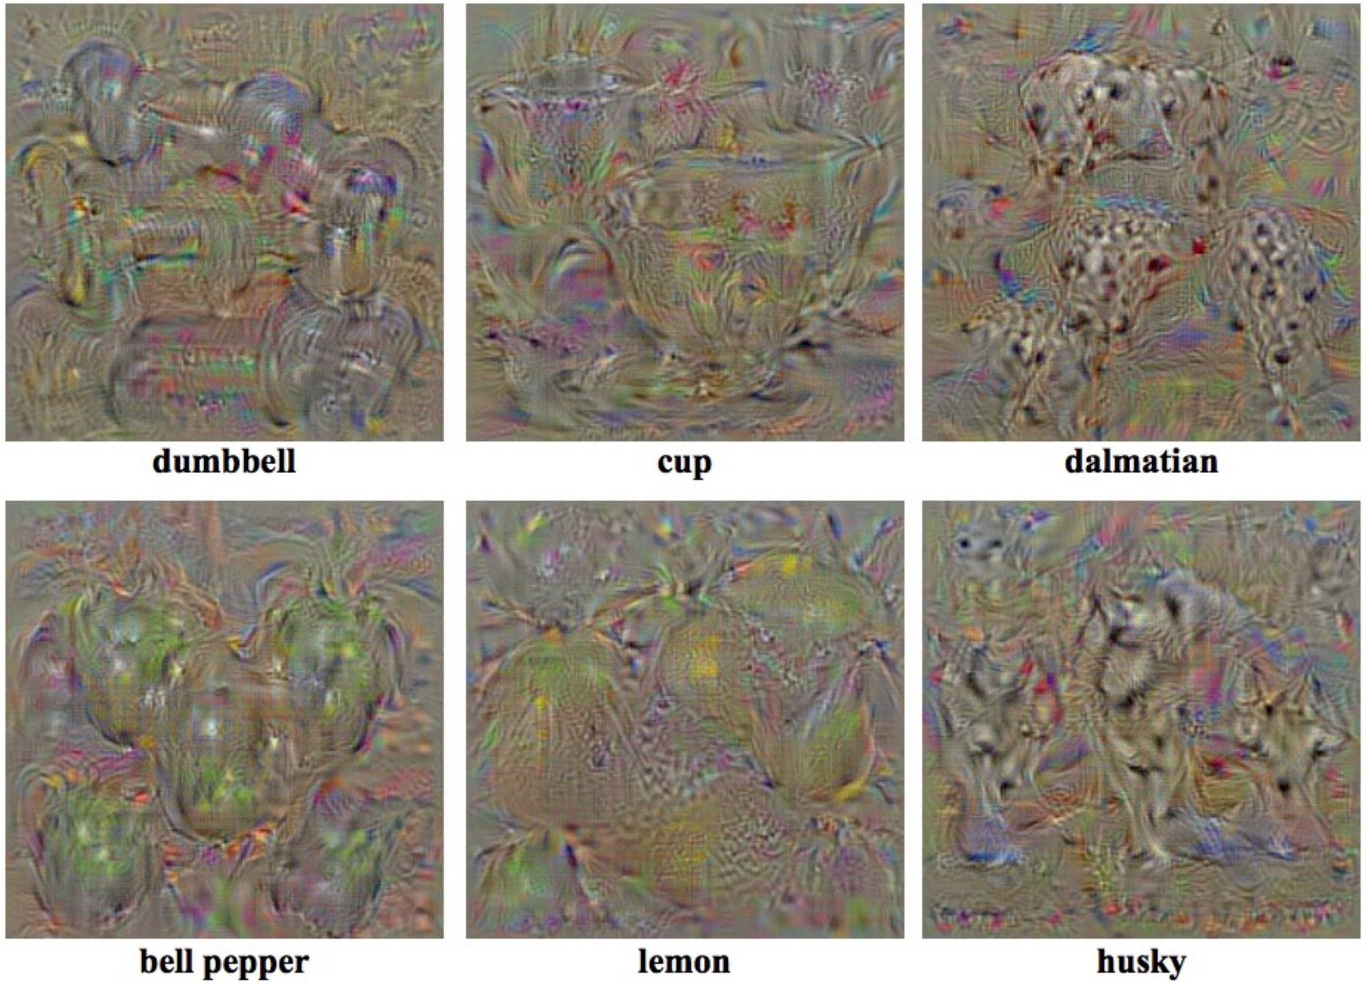

Simonyan K et al. arXiv preprint arXiv:1312.6034. 2013 Dec 20.

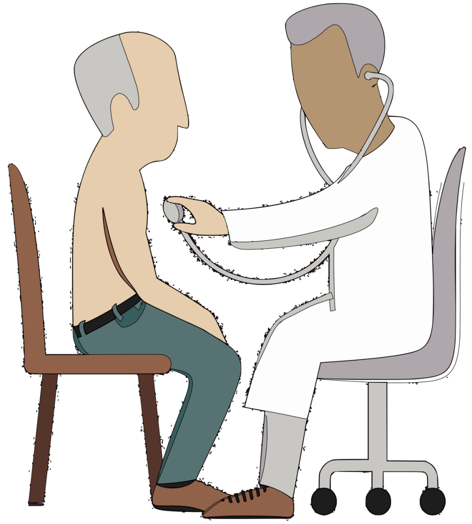

How Can AI Assist in Real-World Clinical Settings?

1. Real-time brain cancer evaluation during surgery
2. Multi-omics prediction for personalized colorectal cancer treatments
3. Multi-task AI for genomic profile identification

Example 1: The Challenge of Intra-Operative Neuropathology Diagnosis

Brain cancer: 204,000 death/year

Tumor sample

Diagnostic results guide surgery

Costly

Human Error

Physician Burnout

"Frozen Section" Diagnosis

Siegel RL et al. *CA Cancer J Clin.* 2023 Jan 1;73(1):17-48.

AI-Based Cryosection Histopathology Assessment and Review Machine (CHARM)

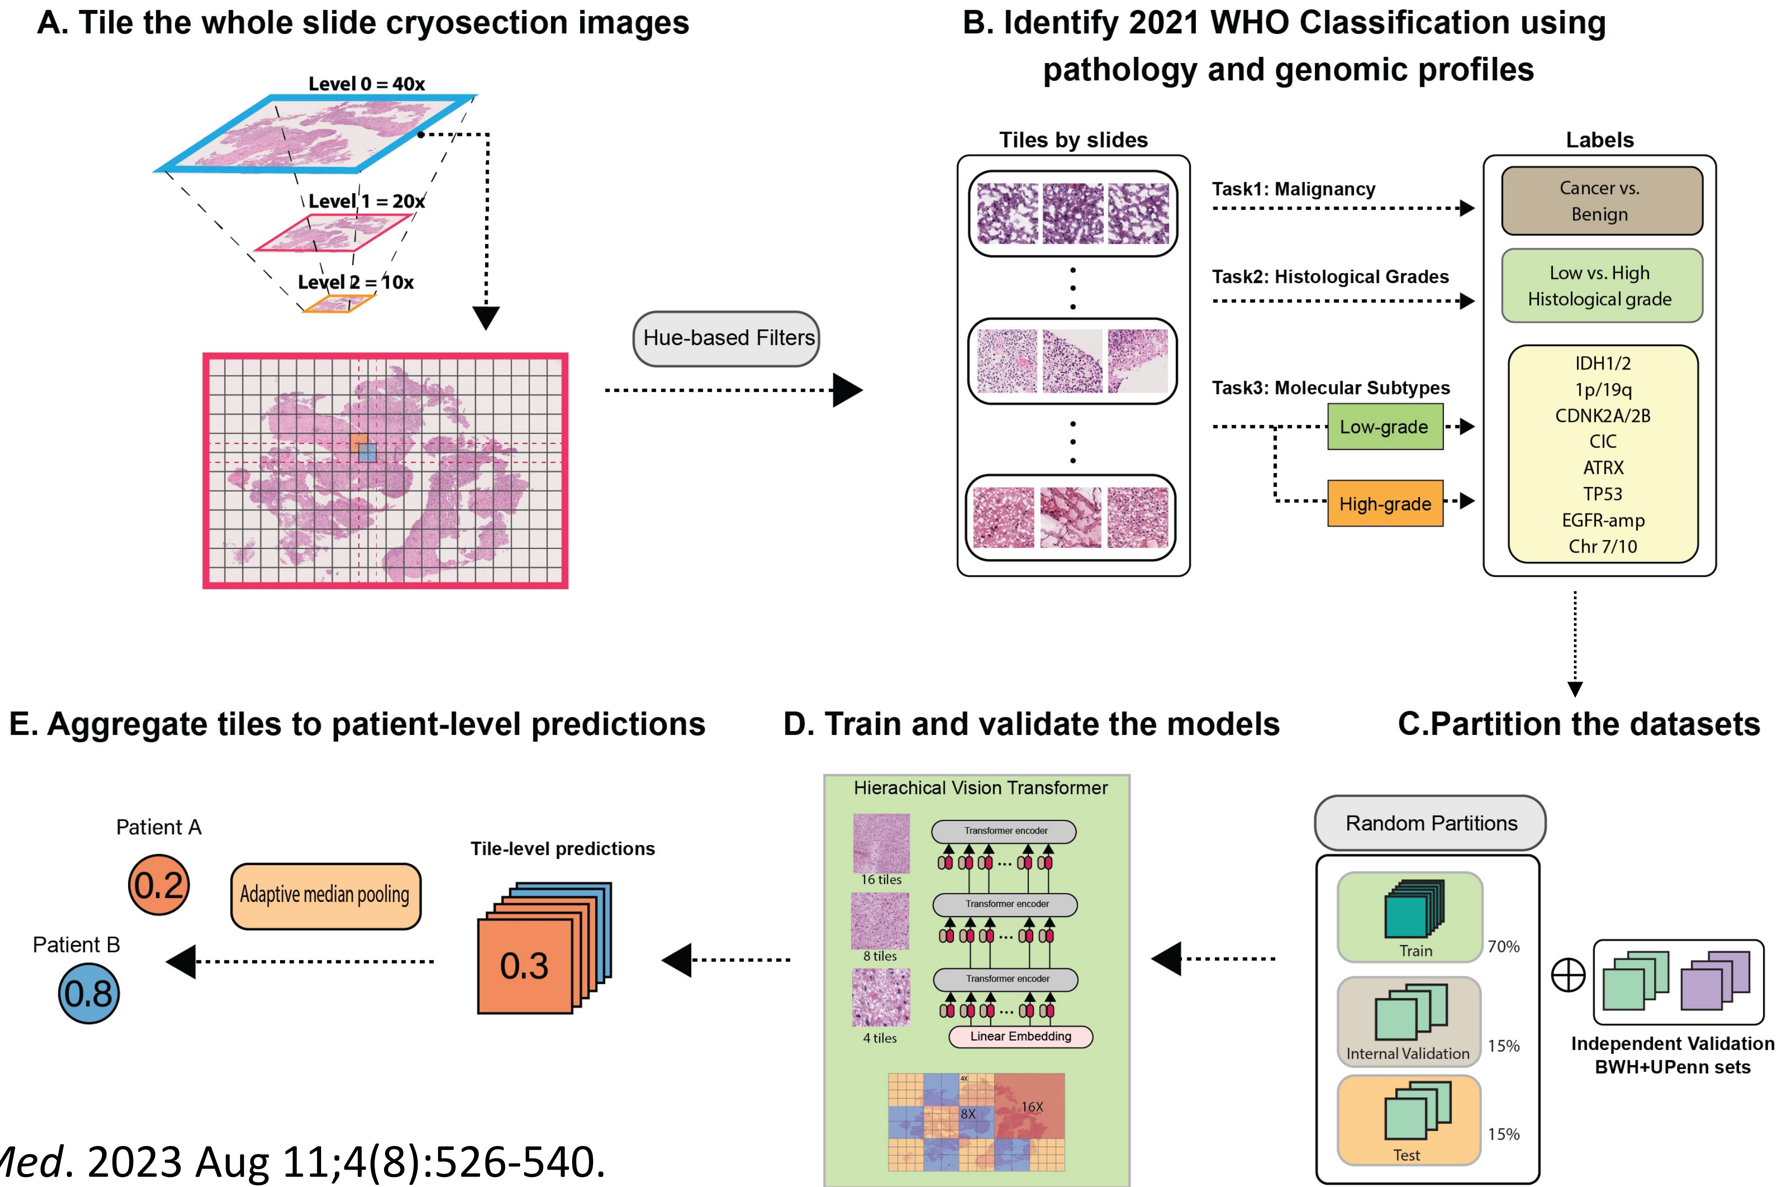

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

DEPARTMENT OF
Biomedical Informatics

Identifying Malignant Cells in Cryosection Samples

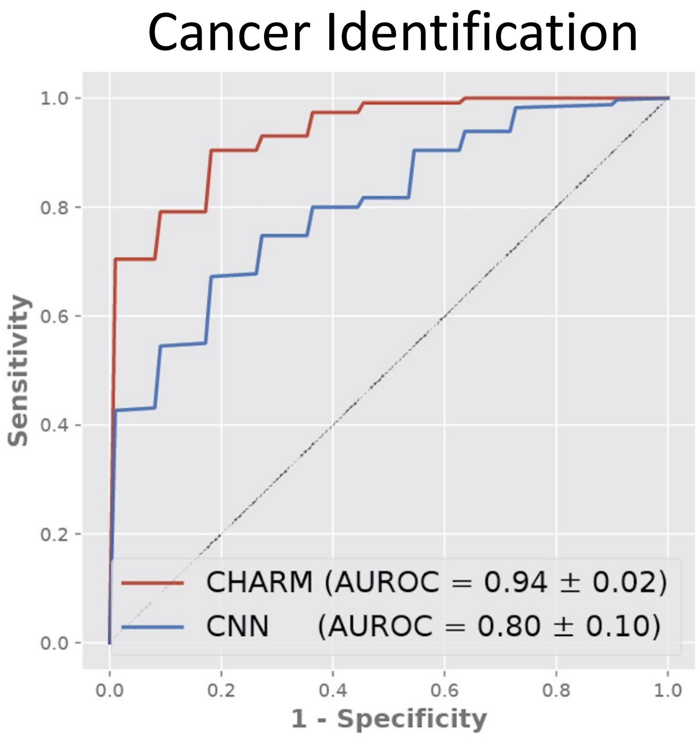

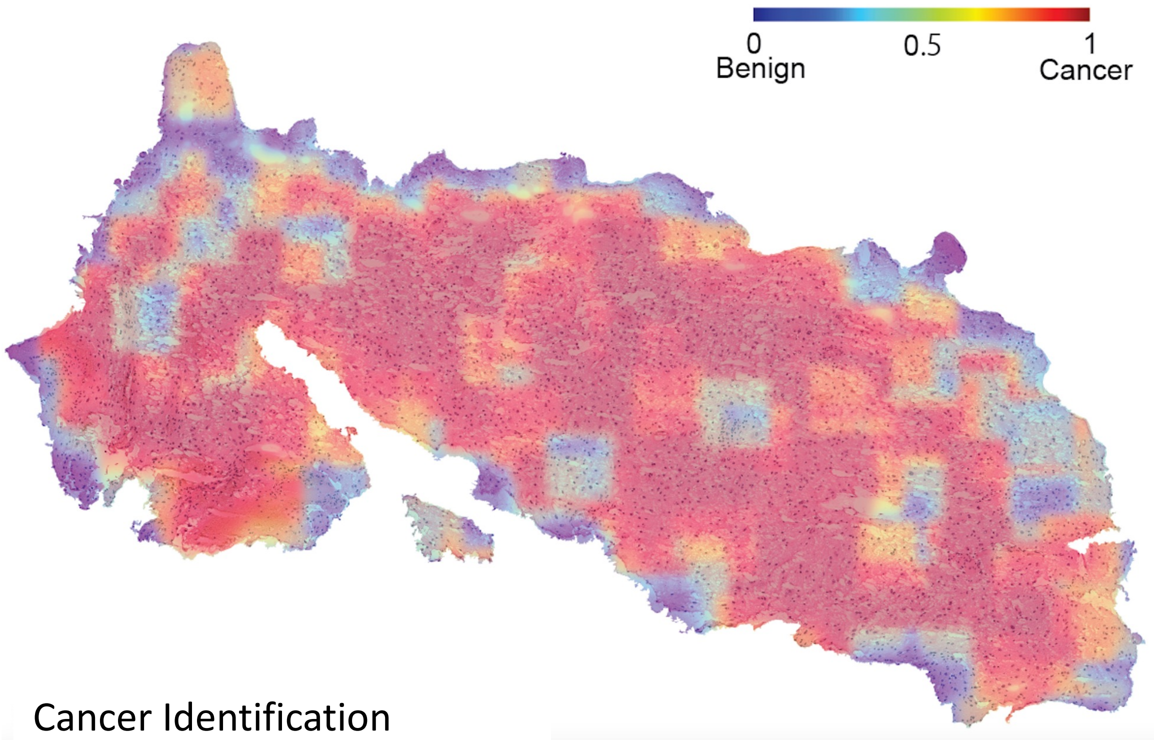

Cancer Identification

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

CHARM Successfully Differentiated Histological Grades

Histological Grade Classification

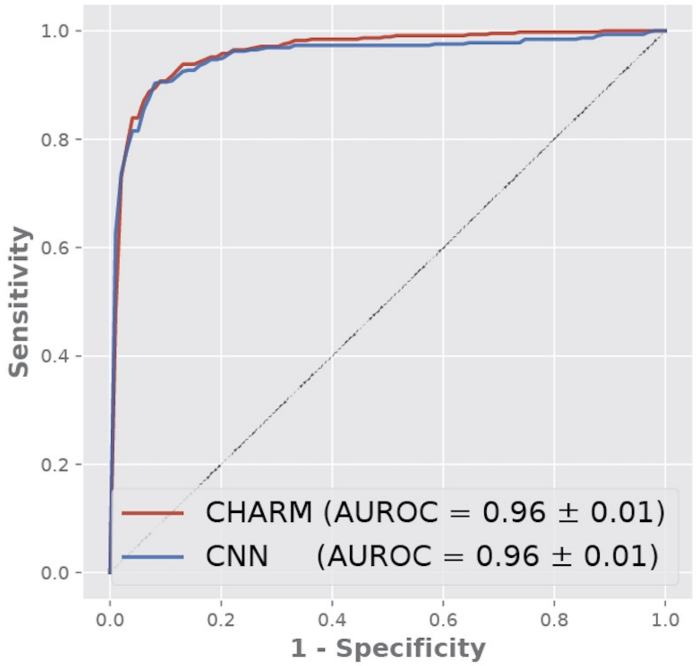

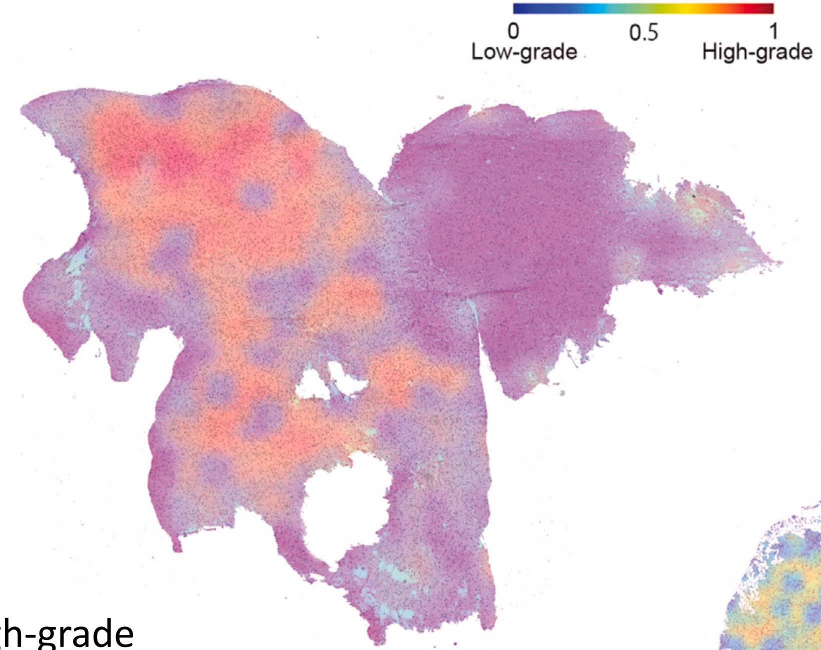

High-grade

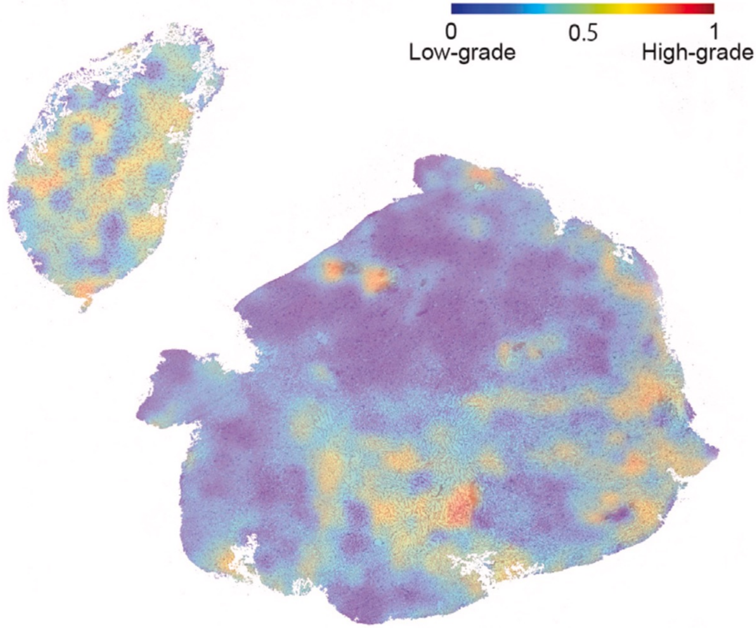

High histological grade: dense glioma cells and necrosis

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

Low-grade

However, the updated WHO Classification Included Molecular Profiles in the Definition of Glioma Grades

Neuro-Oncology

XX(XX), 1–21, 2021 | doi:10.1093/neuonc/noab106 | Advance Access date 29 June 2021

The 2021 WHO Classification of Tumors of the Central Nervous System: a summary

David N. Louis, Arie Perry, Pieter Wesseling, Daniel J.
Dominique Figarella-Branger, Cynthia Hawkins, H. K. I
Riccardo Soffietti, Andreas von Deimling, and David W

Can AI infer IDH mutation status from pathology images?

Table 1 2021 WHO Classification of Tumors of the Central Nervous System. Provisional Entities are in Italics

World Health Organization Classification of Tumors of the Central Nervous System, fifth edition	
Gliomas, glioneuronal tumors, and neuronal tumors	
Adult-type diffuse gliomas	
Astrocytoma, IDH-mutant	Low grade
Oligodendroglioma, IDH-mutant, and 1p/19q-codeleted	Low grade
Glioblastoma, IDH-wildtype	High grade
Pediatric-type diffuse low-grade gliomas	
Diffuse astrocytoma, *MYB*- or *MYBL1*-altered	
Angiocentric glioma	
Polymorphous low-grade neuroepithelial tumor of the young	
Diffuse low-grade glioma, MAPK pathway-altered	

Predicting IDH Mutation Status

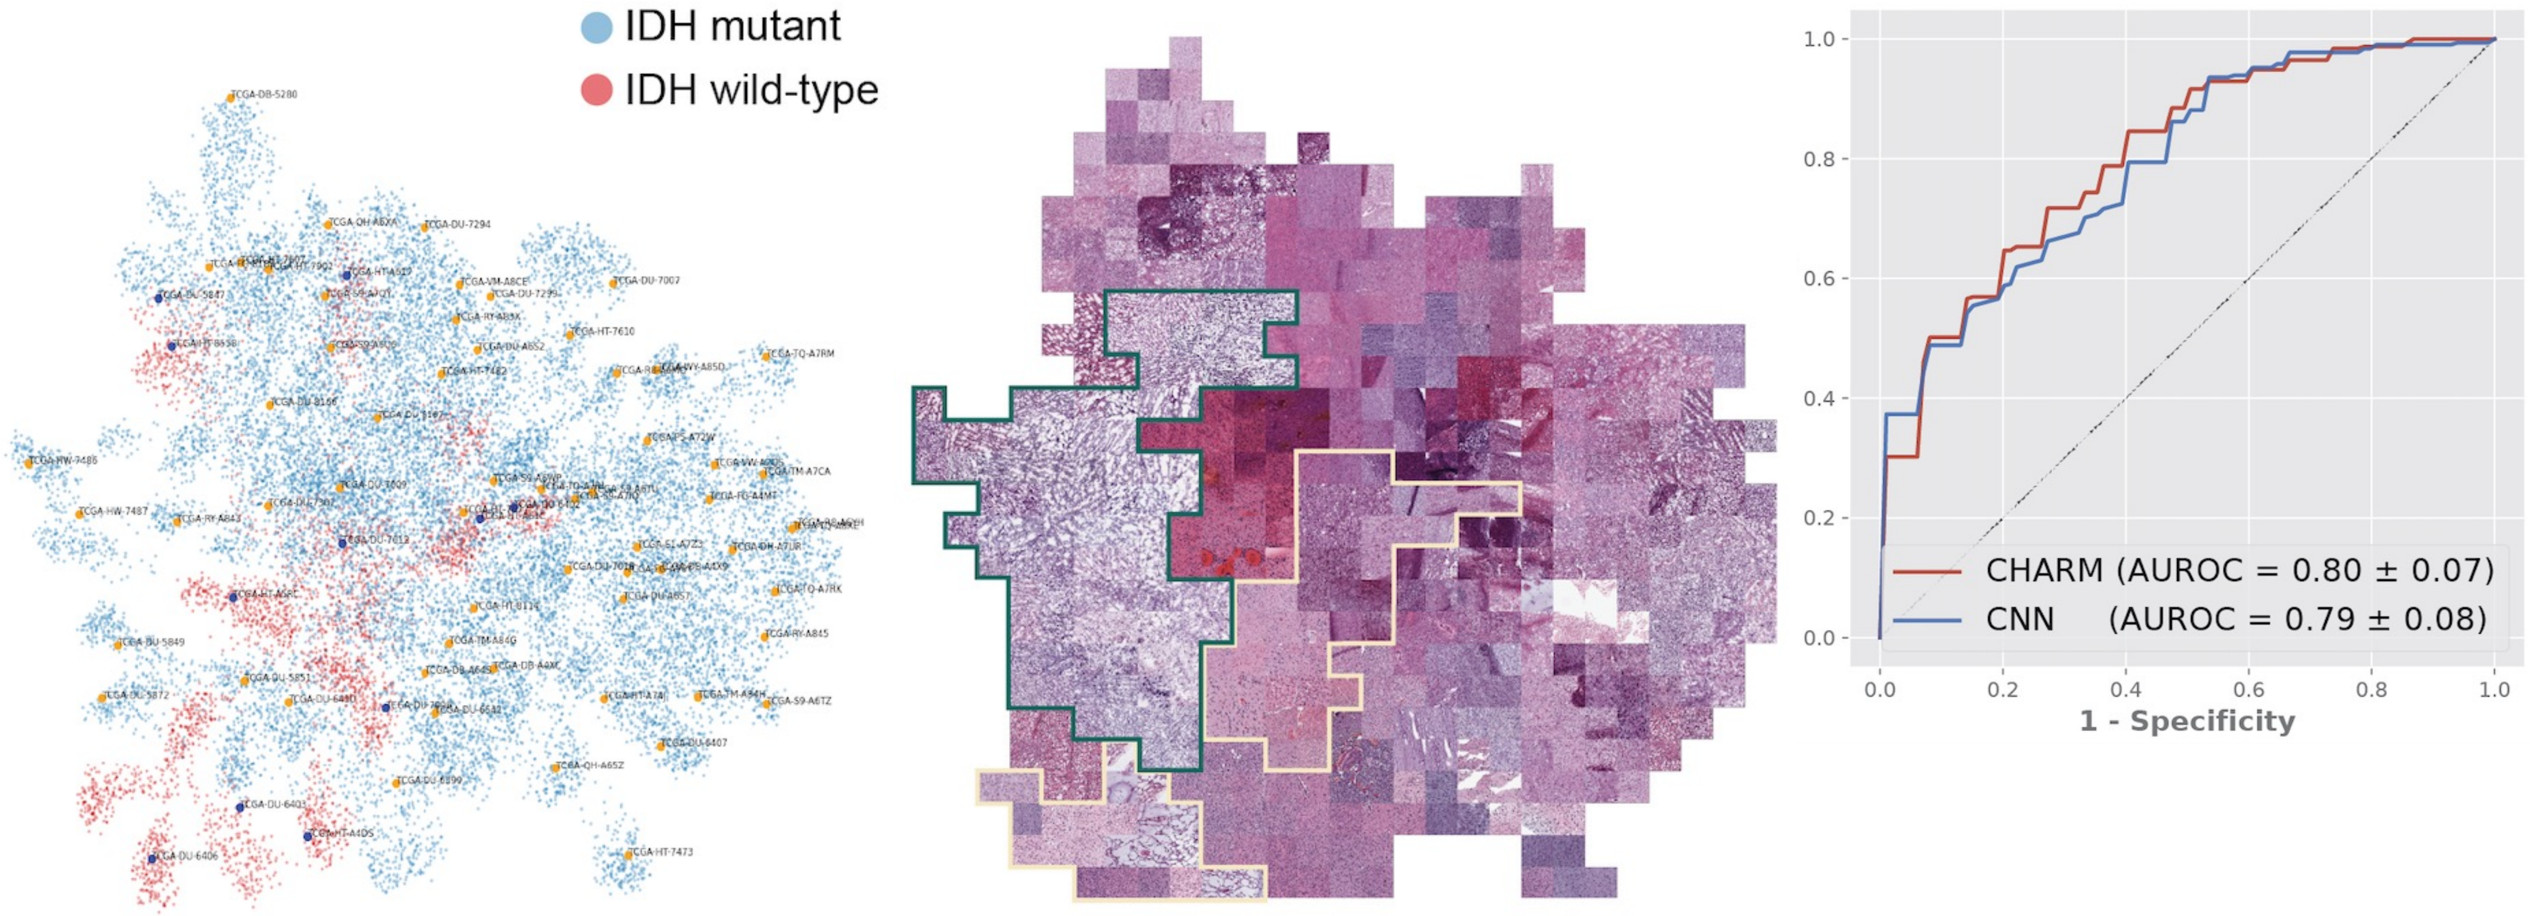

IDH mutant: highly edematous specimens with lower cellularity
IDH wild-type: greater cellularity and atypia

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

Predicting WHO CNS5 Classification

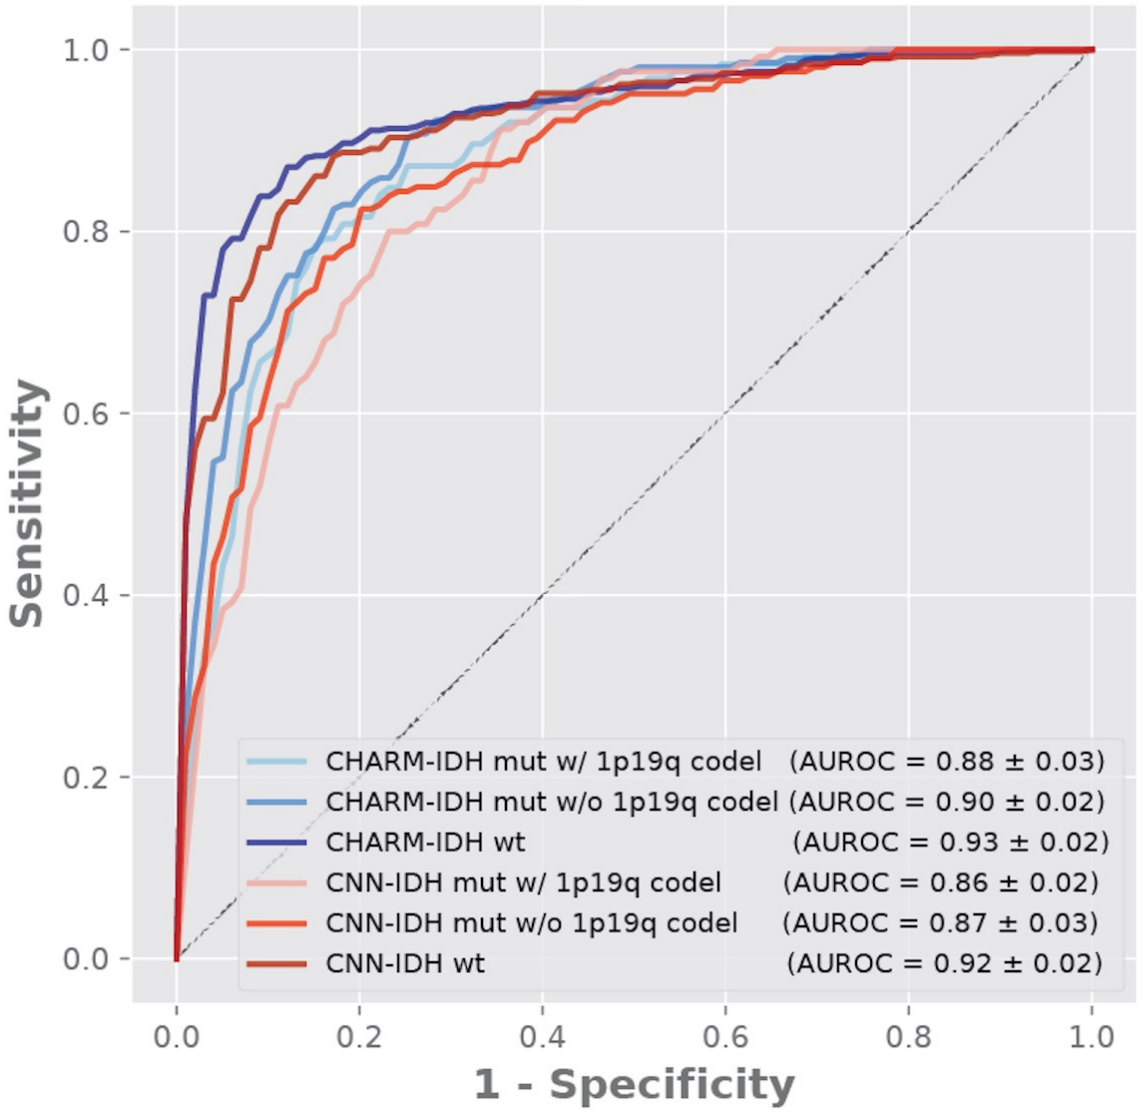

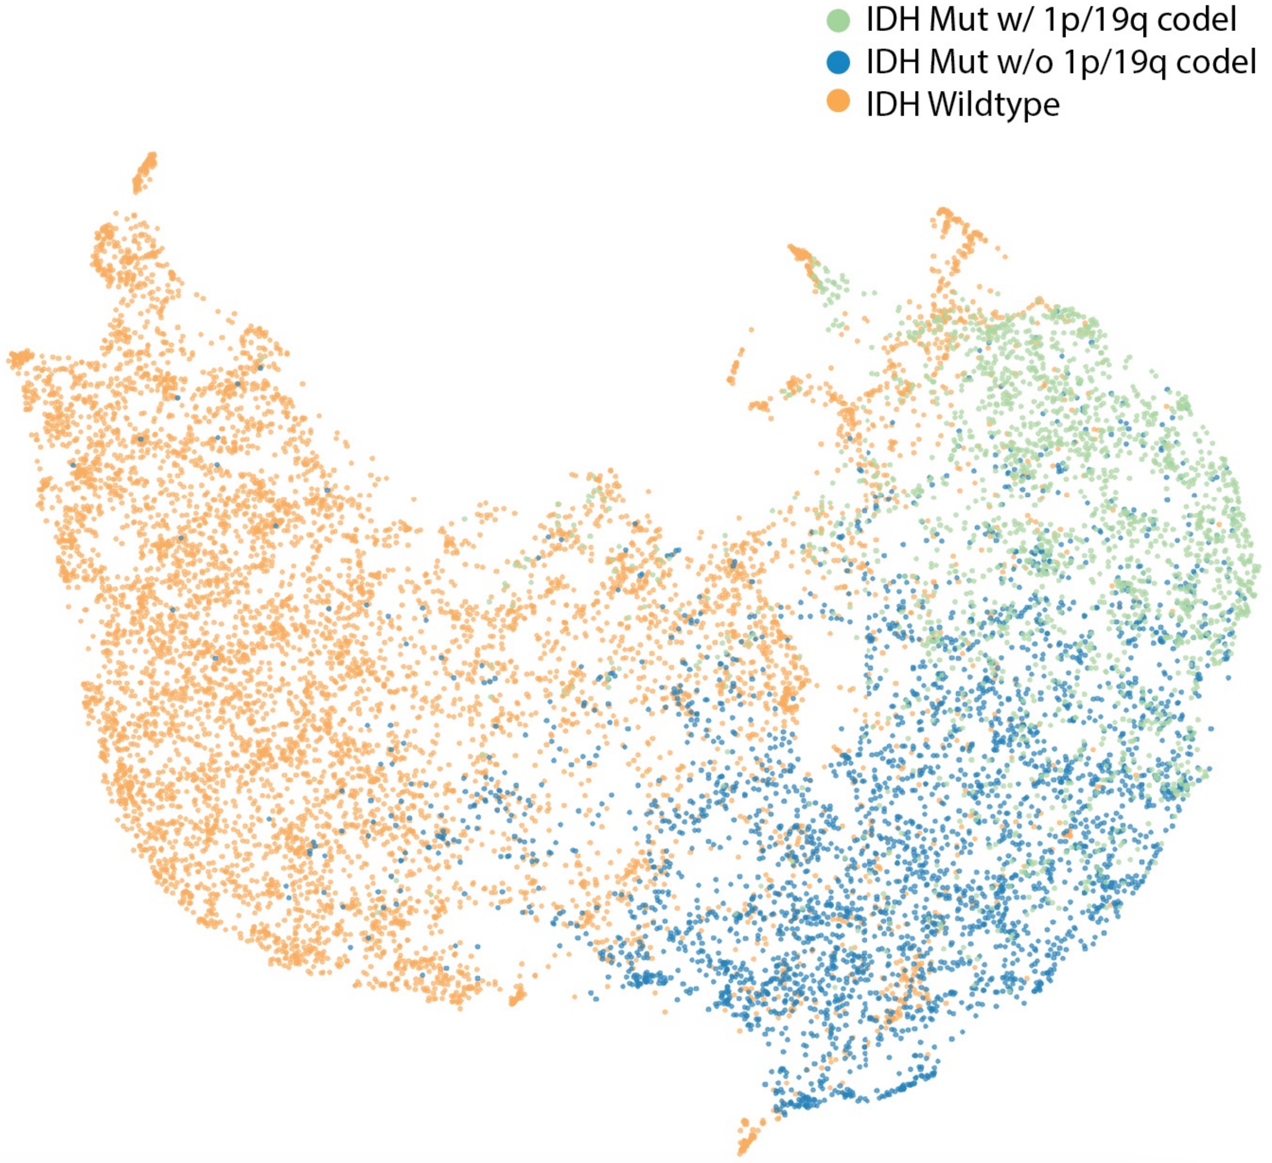

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

Predicting Key Genomic Profiles Related to Prognosis

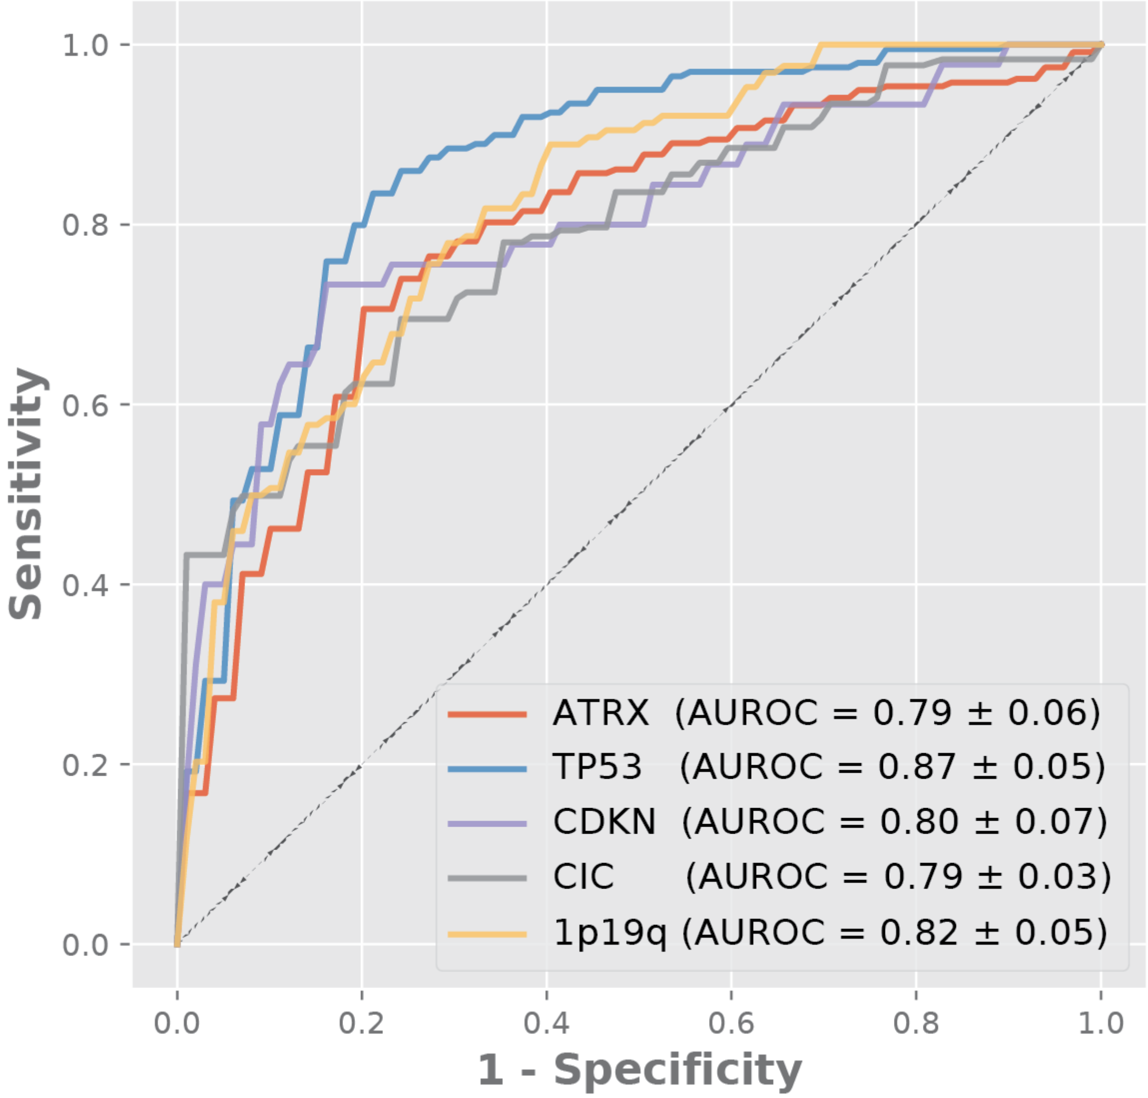

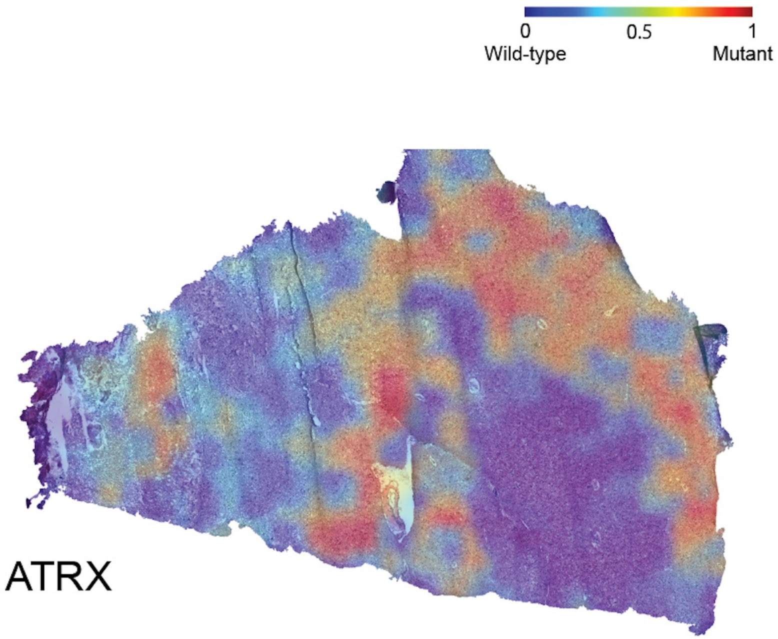

ATRX

ATRX: cortical infiltration by the tumor

CDKN2A/2B homozygous deletion: greater atypia and cellularity

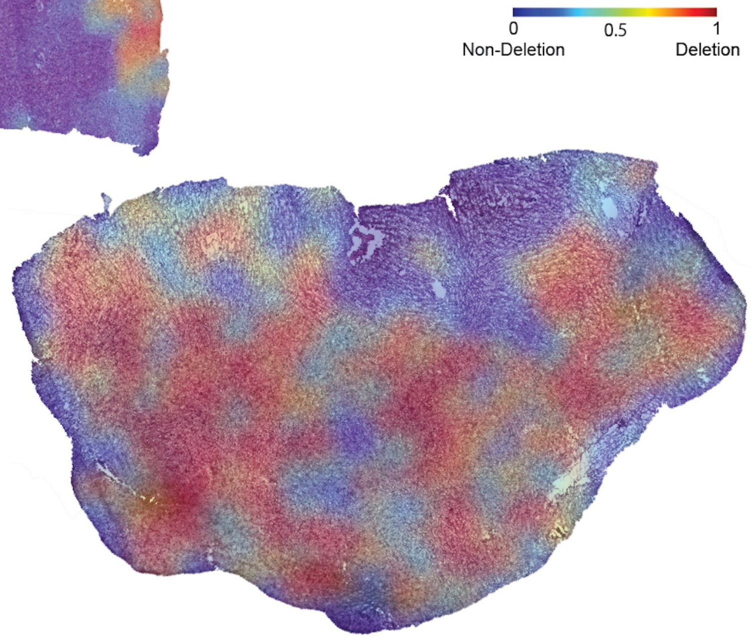

Homozygous deletion of CDKN2A/2B

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

Associating Histology Findings with Tumor Mutation Burden (TMB)

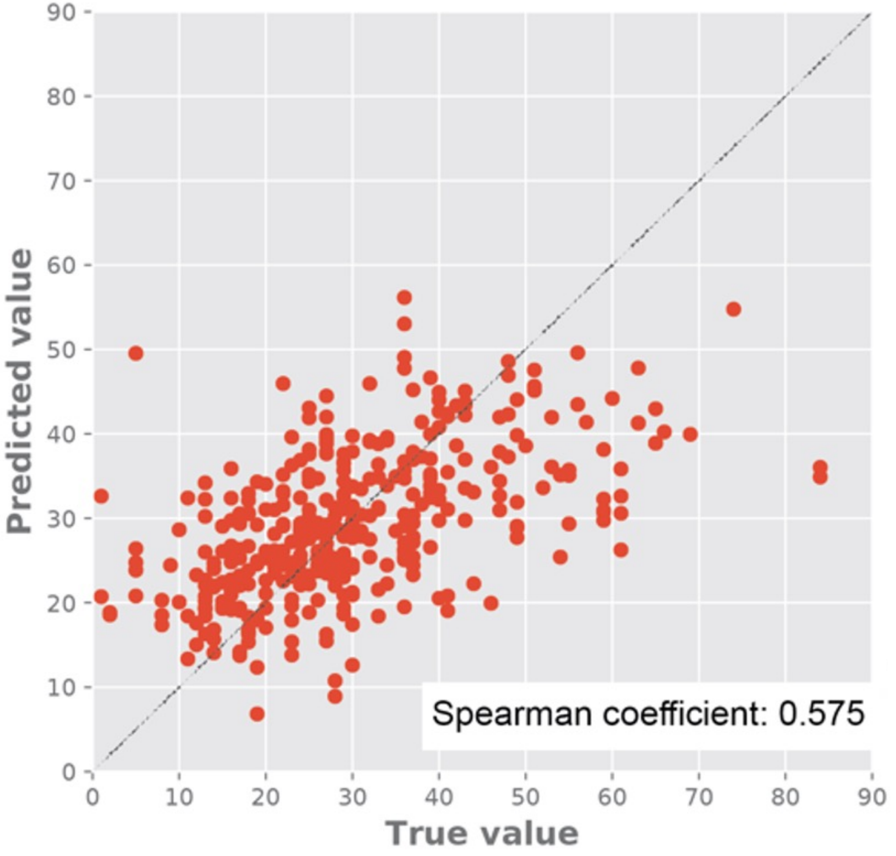

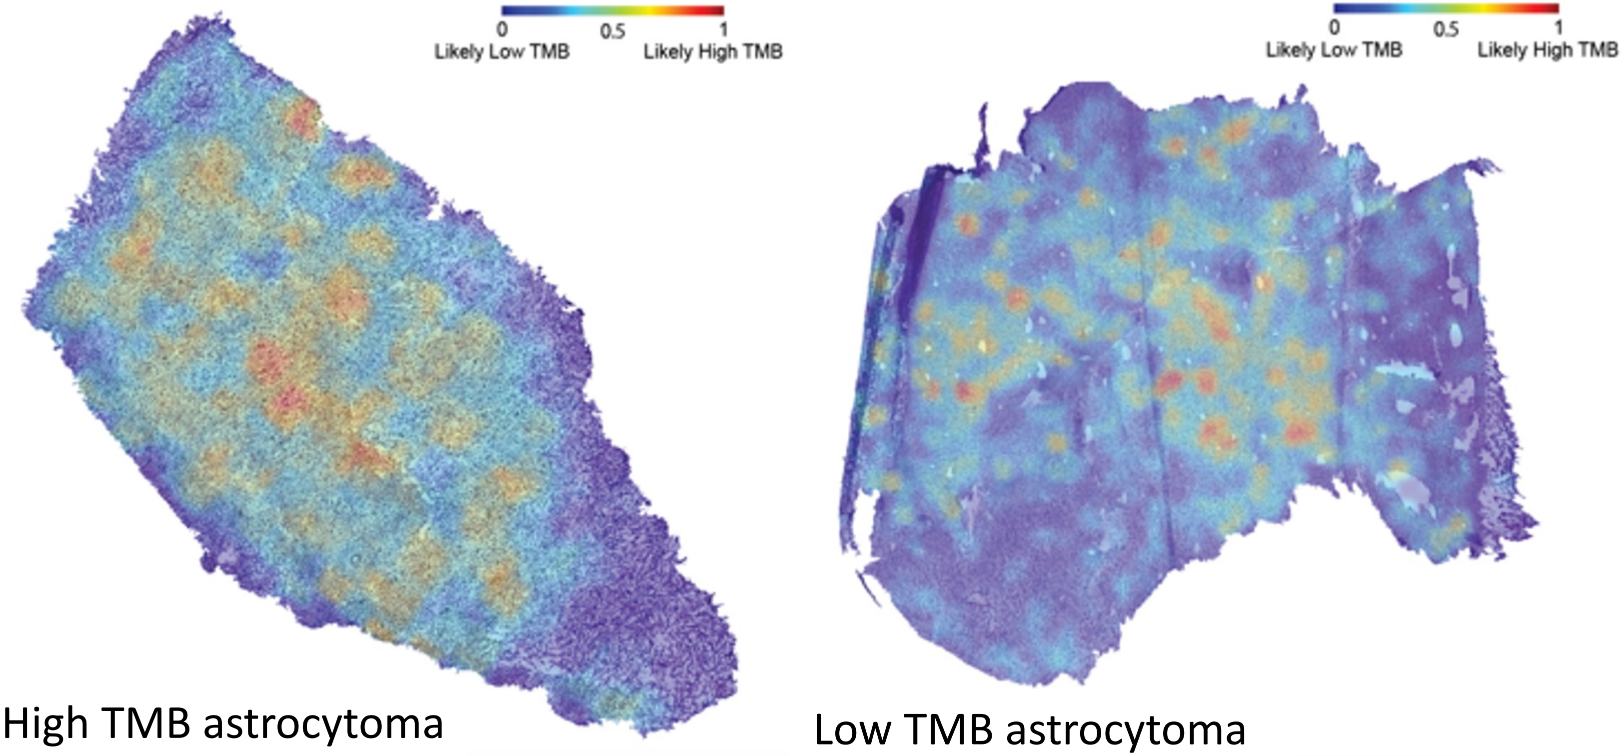

High TMB astrocytoma

Low TMB astrocytoma

High TMB: greater cellularity and atypia, with less edematous regions

HARVARD MEDICAL SCHOOL

DEPARTMENT OF Biomedical Informatics

Nasrallah M et al. *Med*. 2023 Aug 11;4(8):526-540.

Example 2: Multi-Omics Multi-cohort Assessment (MOMA) Platform for Molecular and Prognostic Prediction

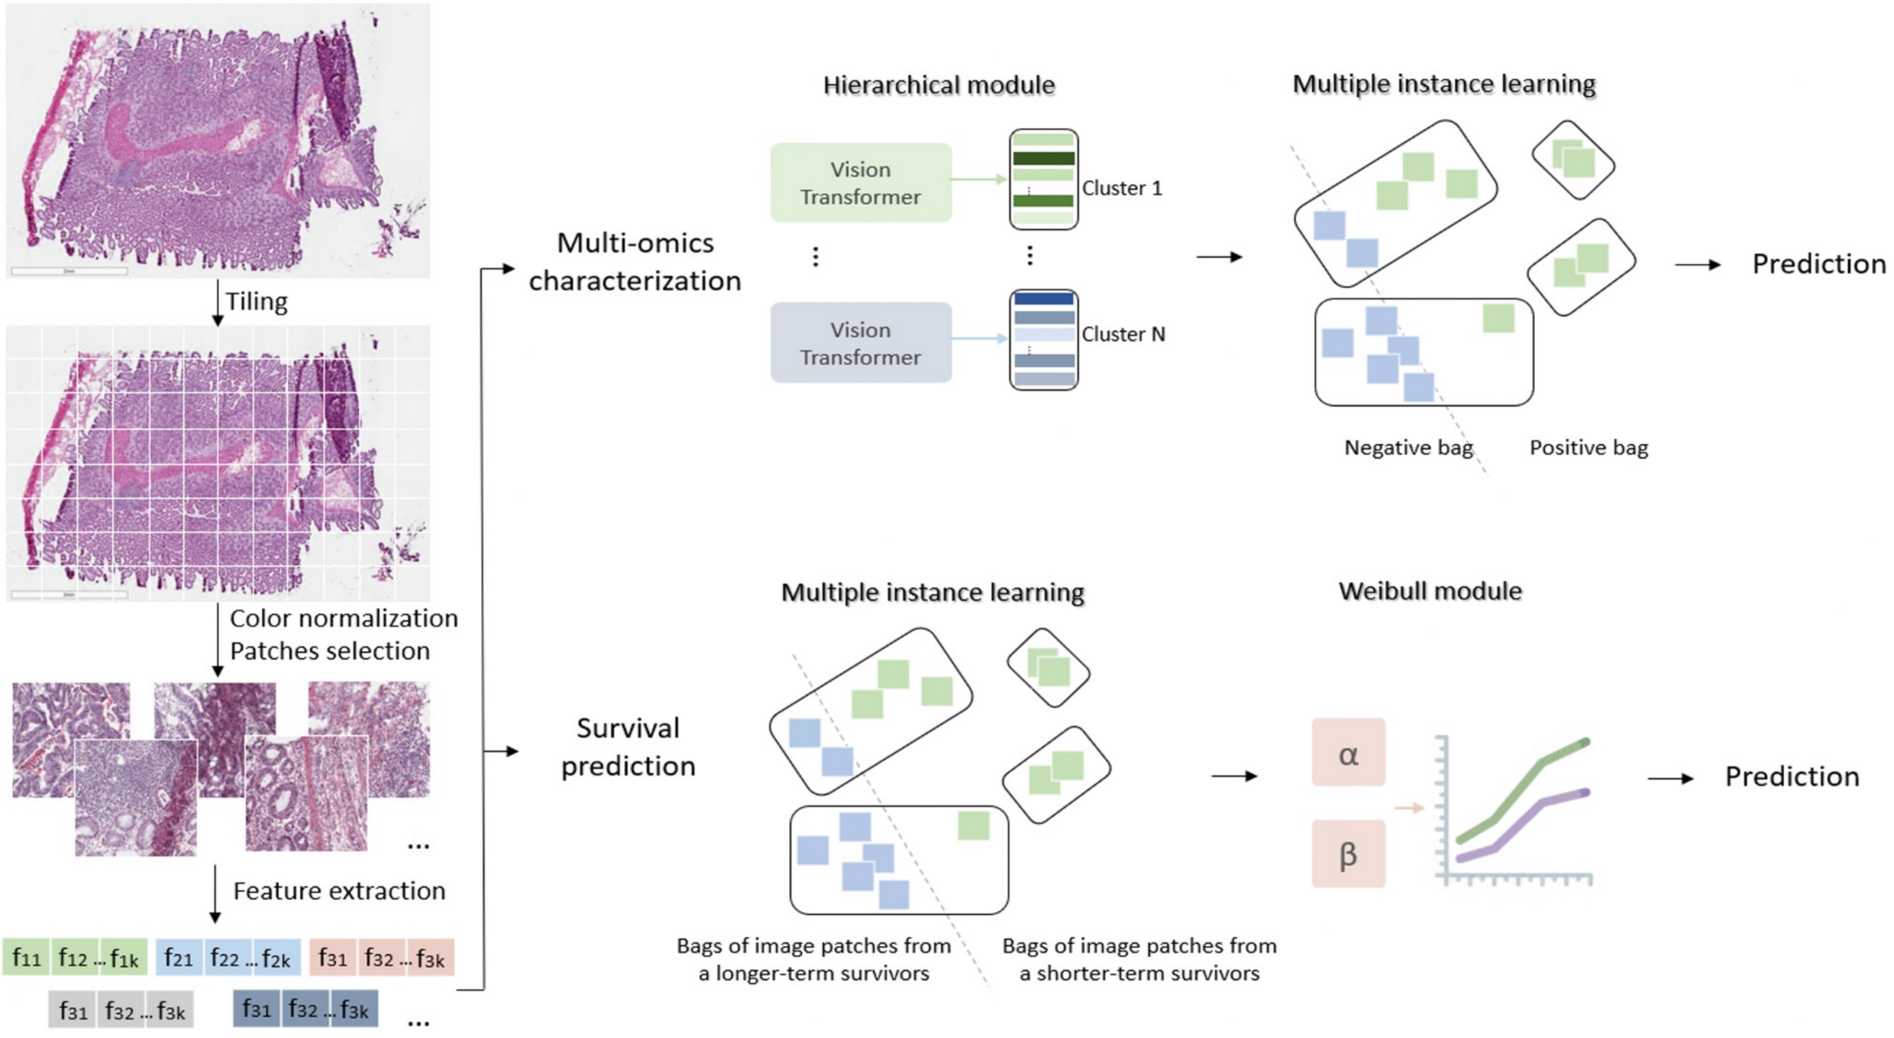

- The Cancer Genome Atlas (TCGA)
- PLCO Cancer Screening Trial
- Nurses' Health Study (NHS)
- Health Professional Follow-Up Study (HPFS)

Whole-slide images

Tissue microarray images

Tsai PC et al. *Nature Communications*. 2023 Apr 13;14(1):2102.

AI Predicts Overall Survival and Disease-Free Survival of Colorectal Cancer Patients in Multiple Cohorts

The Cancer Genome Atlas

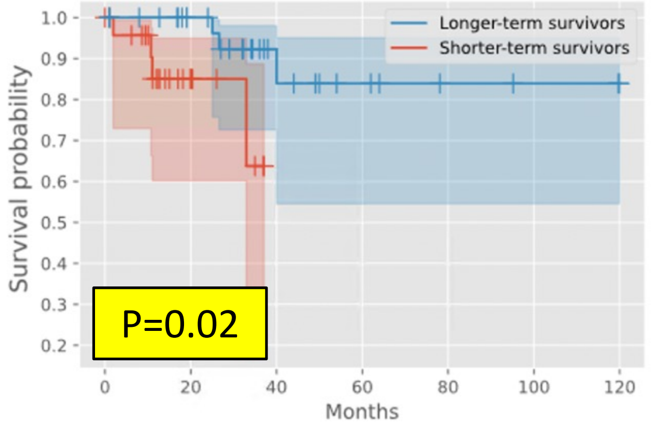

NHS and HPFS

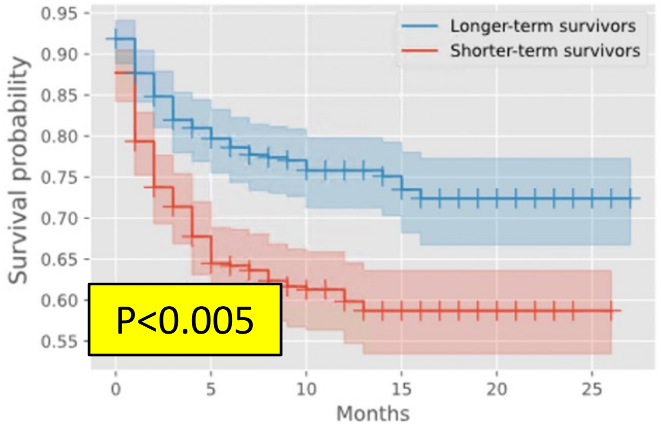

Long-term survivor (124.27 months)

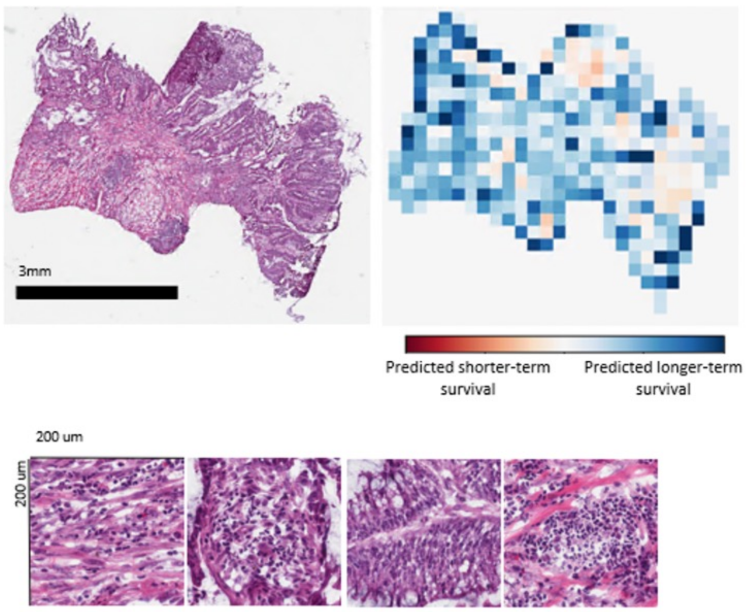

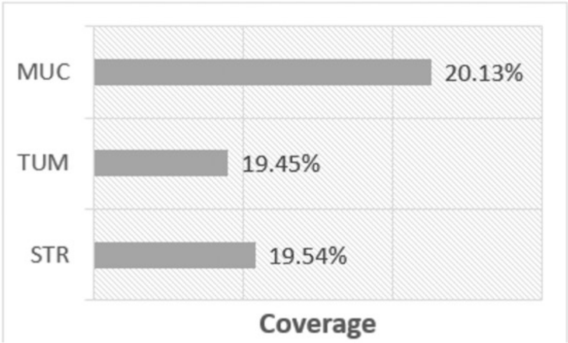

Short-term survivor (2.99 months)

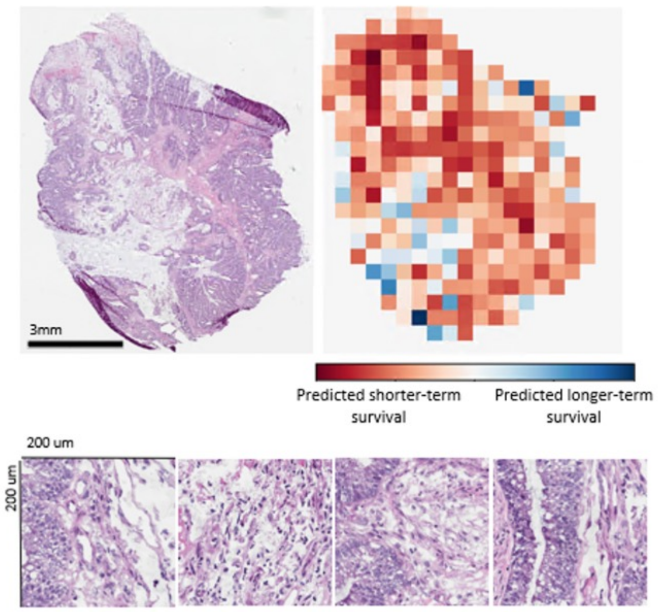

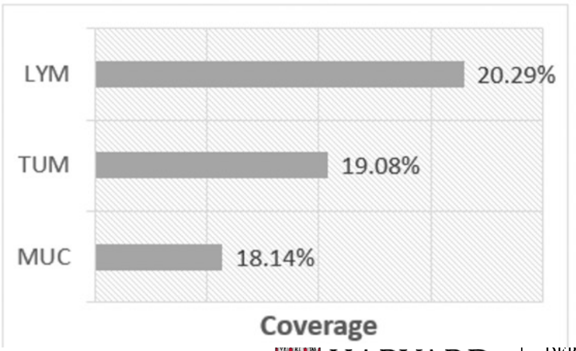

Tsai PC et al. *Nature Communications*. 2023 Apr 13;14(1):2102.

AI Predicts Multi-Omics Profiles from Pathology Images

- MSI prediction

The Cancer Genome Atlas

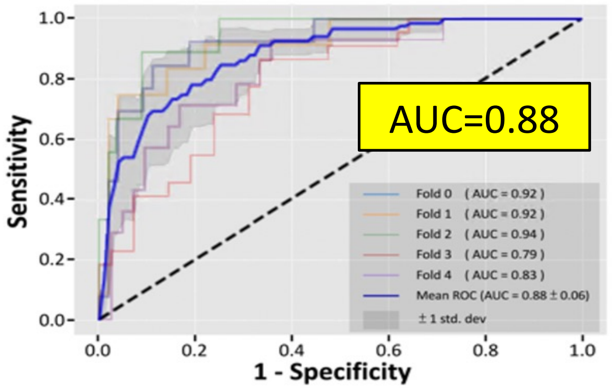

NHS and HPFS

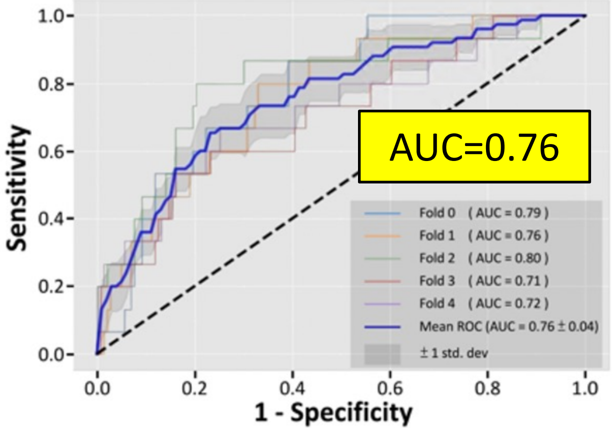

- Copy number alteration

Deletion

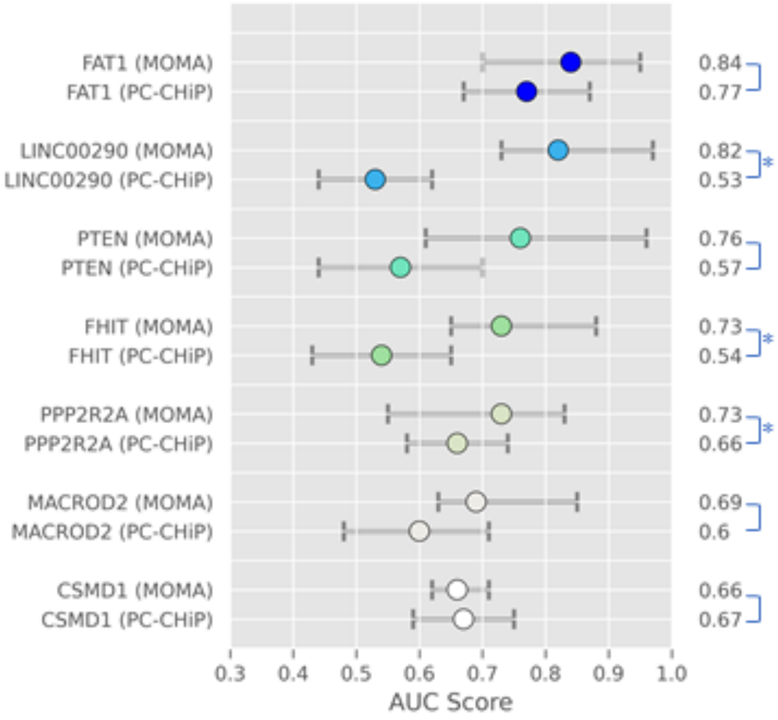

Amplification

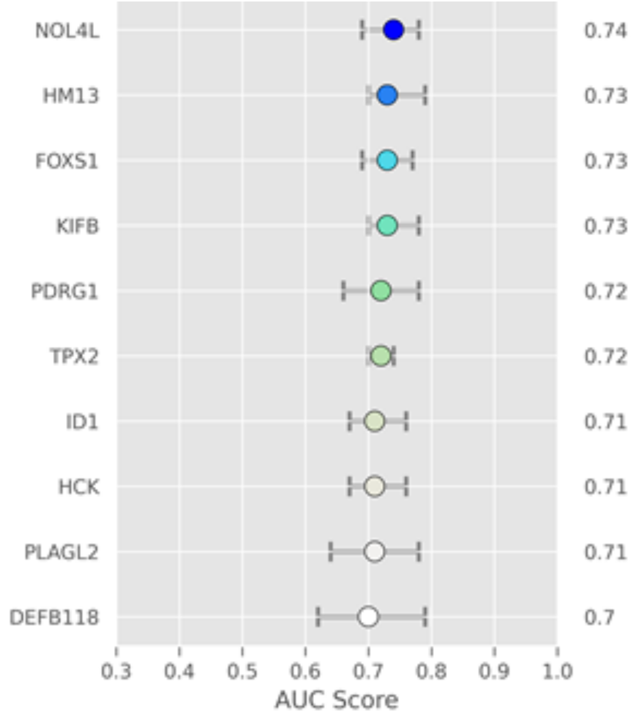

Tsai PC et al. *Nature Communications*. 2023 Apr 13;14(1):2102.

Explainable AI Describes Novel Imaging Patterns Using Pathology Concepts

- Connecting pathology knowledge with AI-derived features using the weight map

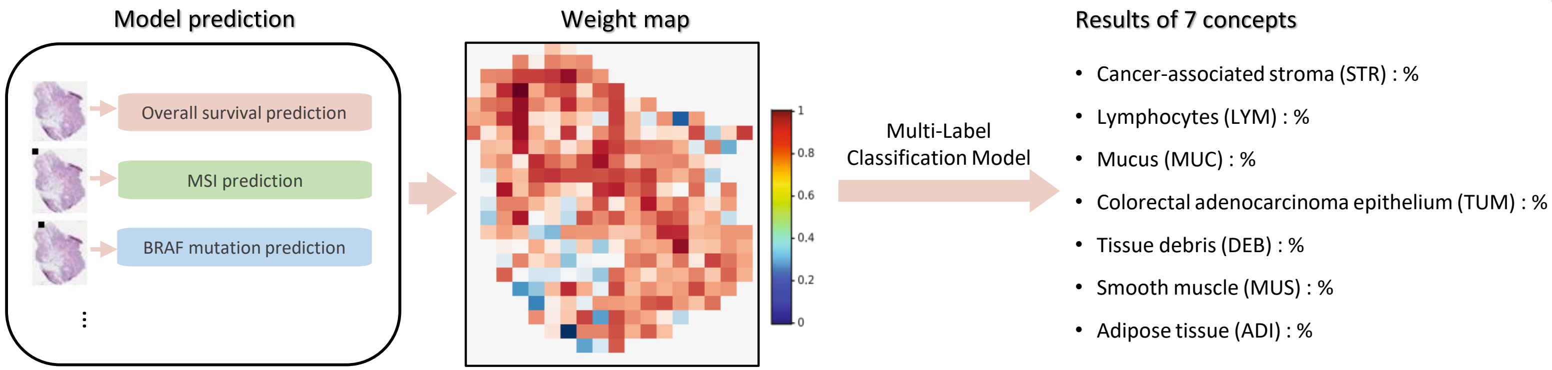

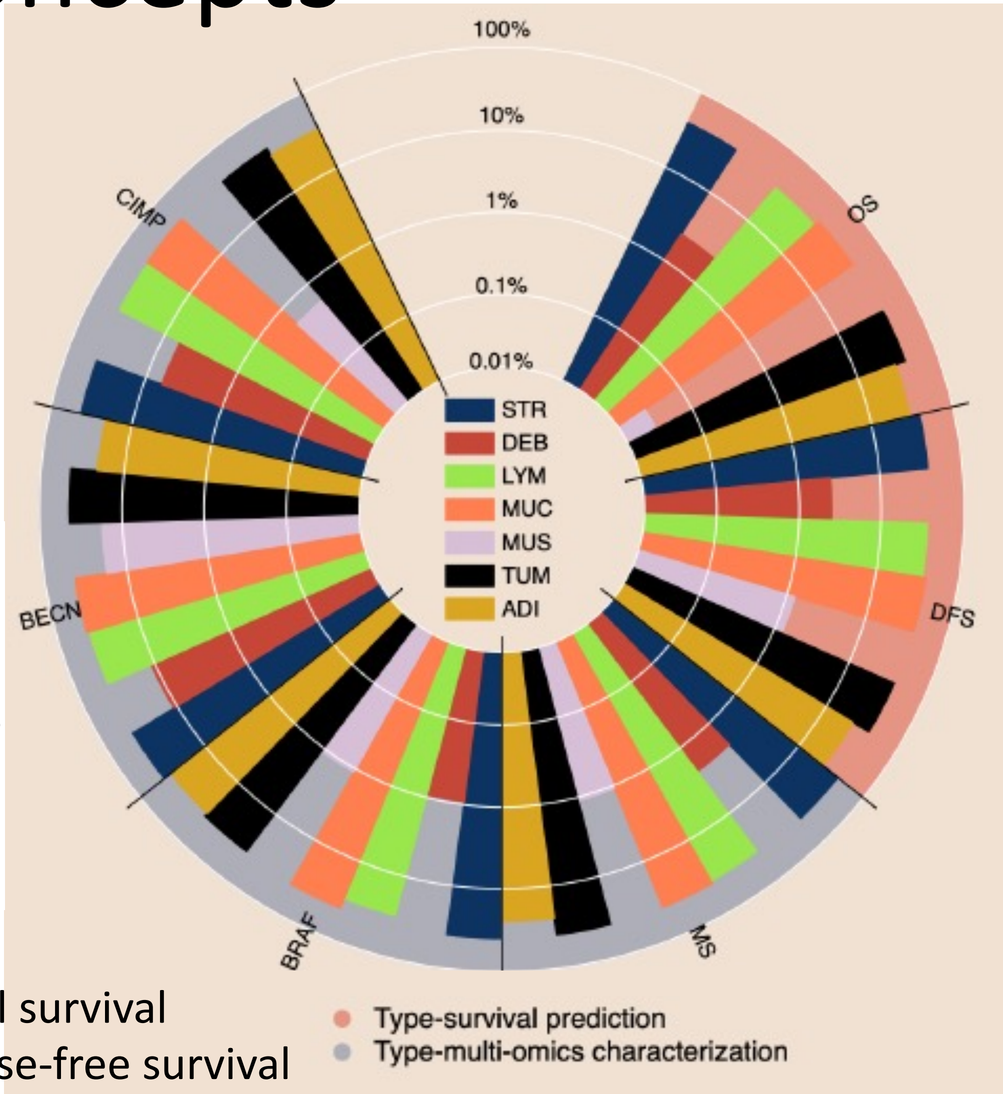

OS: Overall survival
DFS: Disease-free survival
MS: Microsatellite instability

Tsai PC et al. *Nature Communications*. 2023 Apr 13;14(1):2102.

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics

Summary

- Human-interpretable features + machine learning

- Deep learning methods
  - Multi-modal foundation models
  - Model interpretation

- Clinical applications in cancer pathology diagnoses

HARVARD MEDICAL SCHOOL
DEPARTMENT OF Biomedical Informatics