QUANTITATIVE PATHOLOGIC ASSESSMENT USING AI-BASED WHOLE-SLIDE IMAGE ANALYSIS

AI4HEALTH – PHAROS AI FACTORY TRAINING SERIES

Georgios Manikis

Assistant Professor, Hellenic Mediterranean University

Collaborating Researcher at the Computational BioMedicine Laboratory (CBML), FORTH and the Karolinska Institutet, Department of Oncology-Pathology, Sweden

Marie Sklodowska-Curie Postdoctoral Fellow at the University of Cyprus

manikis@hmu.gr

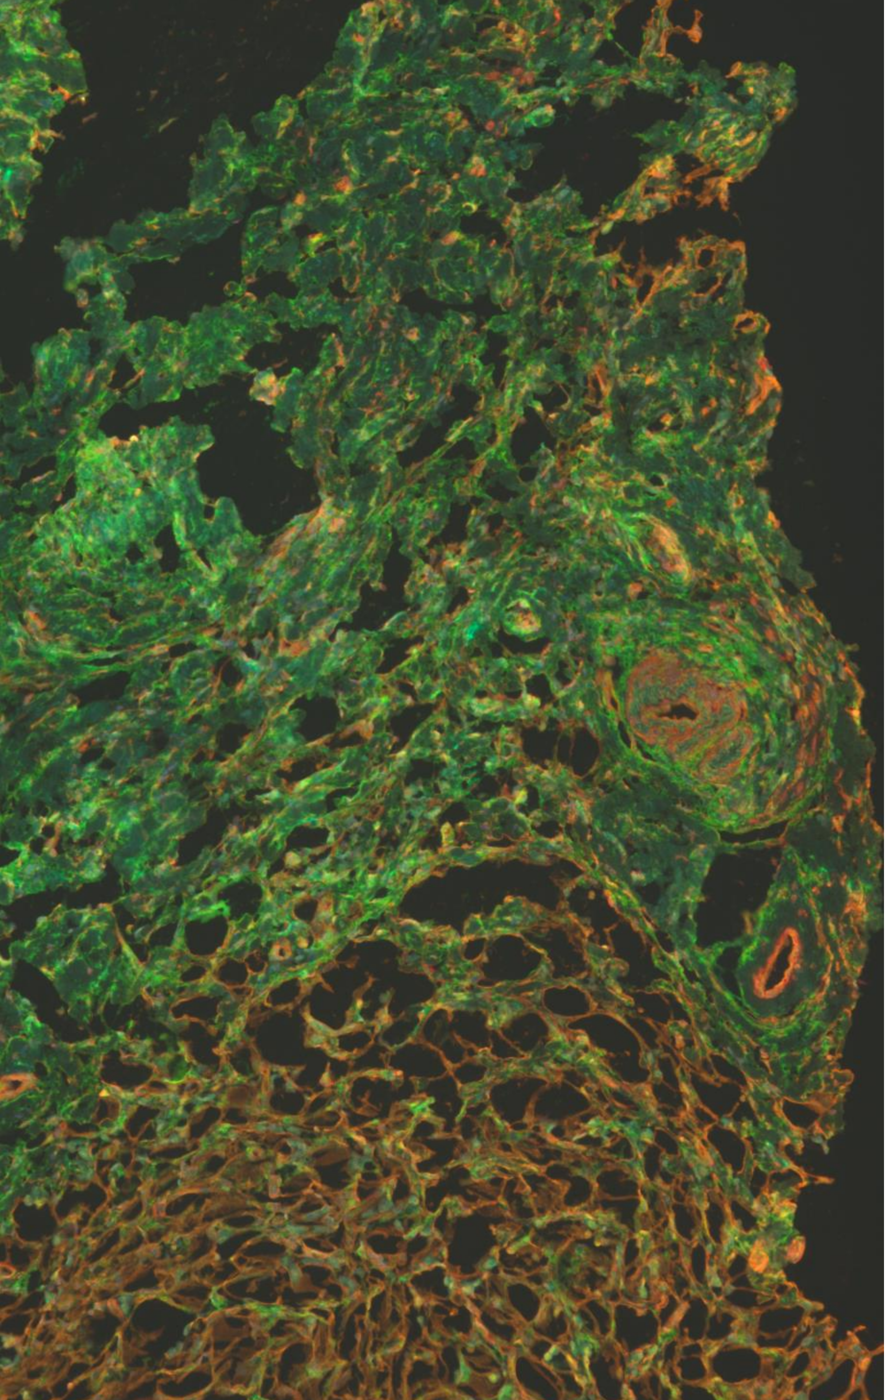
Karolinska Institutet, Department of Oncology-Pathology, Sweden

THE TRANSFORMATION OF PATHOLOGY

- Digital pathology has fundamentally reshaped how tissue is analyzed
- Instead of visual inspection under a microscope, we now work with gigapixel whole-slide images (WSIs) that enable computational interrogation of millions of cells simultaneously
- This shift allows us to move from subjective visual grading to reproducible, quantitative, and scalable analysis pipelines powered by AI

THE DATA SCALE PROBLEM IN DIGITAL PATHOLOGY

- Digital pathology fundamentally differs from standard medical imaging (like MRI or CT) due to its sheer scale
- A Whole-Slide Image (WSI) is not a single image, but a multi-resolution pyramidal structure where the highest magnification (40x) contains billions of pixels
- **The Memory Footprint**: A single, uncompressed 40x slide requires up to 30 GB of active storage and memory, pushing the boundaries of standard image processing libraries
- Processing this data at a cohort level transitions the problem from basic algorithm design to advanced High-Performance Computing (HPC)

A WSI at 40x magnification can be of size 100,000 x 100,000 pixels

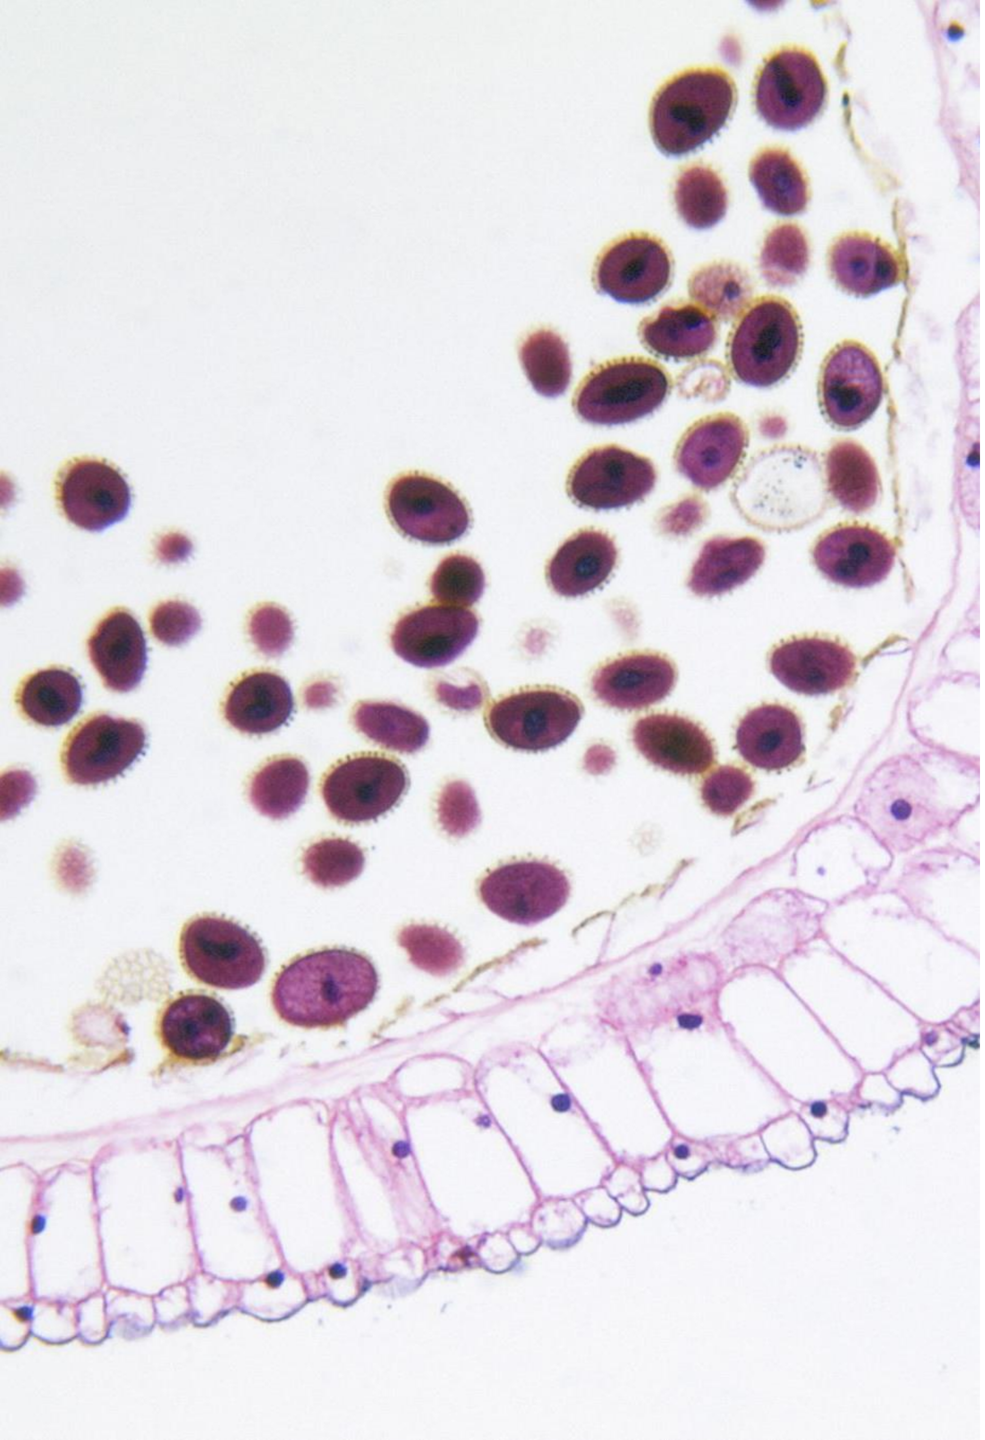

HEMATOXYLIN AND EOSIN (H&E) THE GOLD STANDARD OF STRUCTURAL MORPHOLOGY

- **The Biological Reality**: H&E has been the global diagnostic standard for over a century
- Hematoxylin stains cell nuclei a **dark blue/purple** (binding to DNA/RNA)
- Eosin stains the cytoplasm and extracellular matrix **pink** (binding to proteins)
- It provides excellent structural context, allowing pathologists to see tumor shapes, cellular atypia, and tissue invasion
- **The Limitation**: H&E only shows structure; it cannot identify specific molecular targets (like PD-L1 expression)
- **The Compute Reality**: Digitally, an H&E slide is a standard 3-channel (Red, Green, Blue) image tensor

BEYOND H&E: MULTIPLEX IMMUNOHISTOCHEMISTRY

DEEP PHENOTYPING WITH CODEX AND OPAL

- **The Biological Reality**: We must move beyond tissue structure to deep molecular phenotyping
- **Opal Multiplexing**: Uses multispectral fluorescence to visualize 6 to 9 distinct protein markers simultaneously on a single slide
- **CODEX Technology**: Uses DNA-barcoded antibodies to visualize 40 to 100+ markers on the exact same tissue section
- This allows us to definitively identify complex immune states, such as exhausted PD-1+ CD8+ T-cells interacting with tumor cells
- **The Compute Reality**: We move from a 3-channel RGB image to a massive 40+ channel biological tensor
- **HPC Execution**: This massive increase in tensor depth instantly exhausts standard GPU memory bandwidth. We must utilize large-memory A100 (80GB) nodes just to load these hyper-plexed image blocks into VRAM

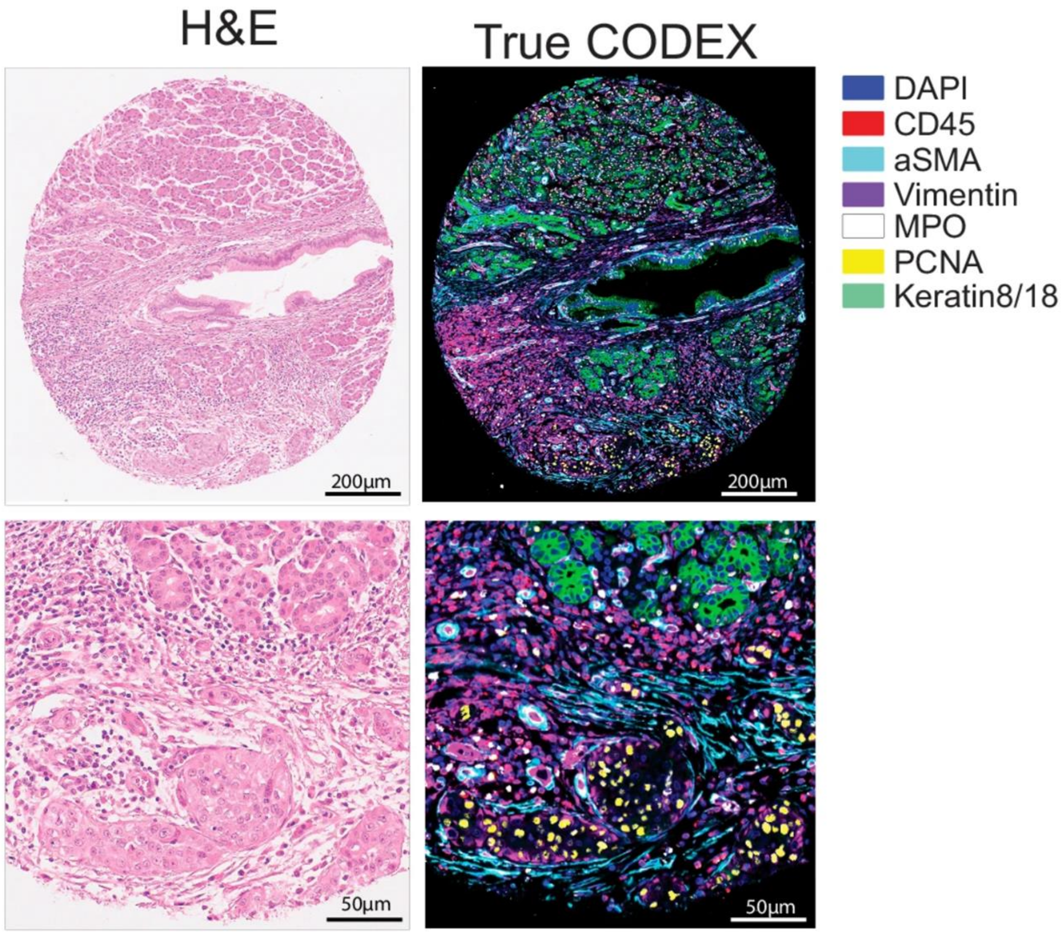

Lin, JR., Chen, YA., Campton, D. et al. High-plex immunofluorescence imaging and traditional histology of the same tissue section for discovering image-based biomarkers. Nat Cancer 4, 1036–1052 (2023)

THE DATA SCALE PROBLEM IN DIGITAL PATHOLOGY

- While a localized academic study might analyze 100 to 500 slides, true clinical validation requires massive scale
- A modern clinical trial or national biobank cohort can analyze thousands of patients (processing a 5,000-patient cohort generates over 150 Terabytes of raw, high-dimensional imaging data)
- Standard hospital IT infrastructure and basic GPU workstations completely fail at this scale due to severe network bottlenecks, limited RAM, and insufficient storage IOPS (Input/Output Operations Per Second)
- ***Designing a novel neural network architecture is no longer the primary hurdle in medical AI.***
- ***The true bottlenecks are Data I/O speeds, GPU VRAM capacity, and multi-node cluster networking***

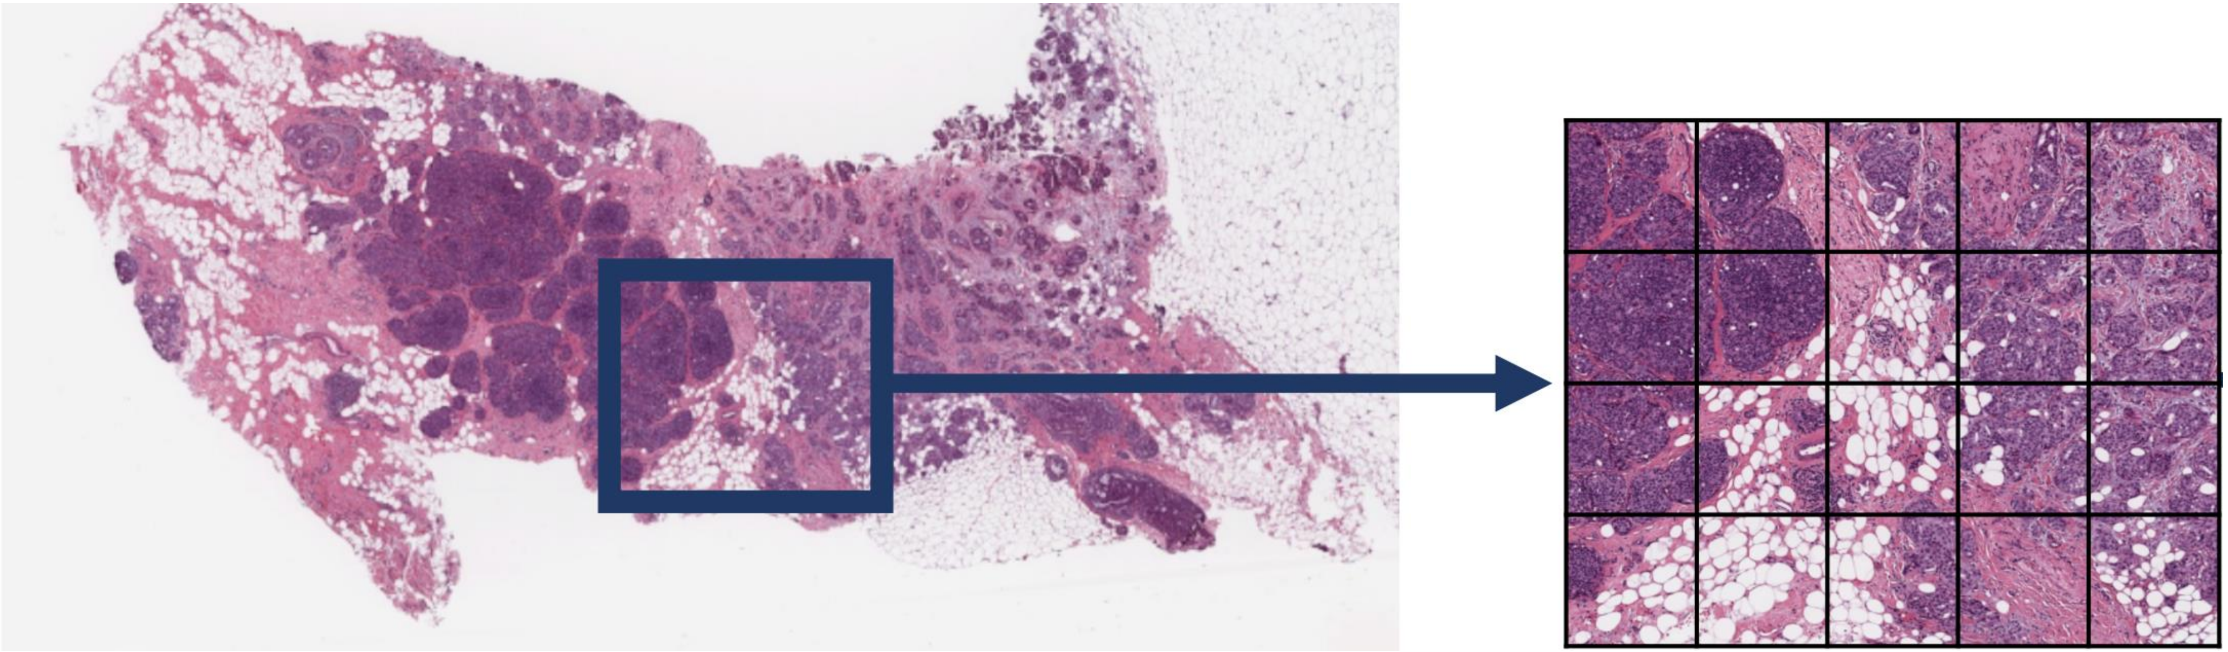

THE INFRASTRUCTURE BOTTLENECK

- In large-scale computational pathology, the main limitation is often not GPU power but data movement (The problem is not the GPUs, is how the data is stored and loaded)
- Extracting patches from hundreds of thousands of WSIs results in millions to hundreds of millions of small files, which significantly strain storage systems
- During distributed training on Alvis we observed that GPU utilization can drop dramatically when data loading pipelines are not optimized

WHY MORE GPUS DON'T HELP!

- Whole-slide images are broken into millions of small image patches.

- If each patch is saved as a separate small file (e.g., JPEG):
  - The storage system has to constantly open and close millions of tiny files.
  - This overwhelms the storage servers.
  - GPUs end up waiting for data instead of training.

- <u>What happened on Alvis</u>:
  - GPUs dropped to very low usage
  - Training slowed down instead of speeding up
  - The bottleneck was the storage system, not the AI model

We bypass file-system limits by pre-processing raw WSI data into contiguous, sequential formats like HDF5 chunks

If we don't build our system to do this "prep-work" automatically, our new, multi-million-euro supercomputer will crash the moment we try to process millions of histopathology slides.

THE ALVIS SUPERCOMPUTER (C3SE.CHALMERS.SE/ABOUT/ALVIS/)

- **The Core Mission**: Alvis is Sweden's premier national compute cluster strictly dedicated to Machine Learning (C3SE / NAISS)
- **Massive GPU Density**: The cluster is powered by hundreds of NVIDIA A100 Tensor Core GPUs (40GB and high-memory 80GB variants)
- **Node Topology**: These GPUs are packed into dense, 4-way and 8-way multi-GPU compute nodes
- **The Network Fabric**: Compute nodes are wired together using ultra-low latency HDR InfiniBand interconnects
- **Why InfiniBand Matters**: Standard ethernet cannot handle AI scale
- InfiniBand allows GPUs across different physical servers to share memory and model gradients with near-zero latency

THE ALVIS SUPERCOMPUTER (C3SE.CHALMERS.SE/ABOUT/ALVIS/)

- **The CPU Misconception**: In AI, people often ignore CPUs, but spatial pathology heavily relies on them
- **High-Core Processors**: Alvis features massive AMD EPYC and Intel Xeon CPU partitions (up to 128 cores per node)
- **The CPU Payload**: We utilize these cores for the tasks GPUs fail at: calculating H&E color normalizations and computing complex KD-Trees for our Spatial Graphs
- **The Storage Tier**: Persistent cohort data (petabytes) lives on the central "Cephyr" parallel file system
- **The Compute Tier (NVMe)**: Active training data is staged to ultra-fast, node-local NVMe scratch drives ($SNIC_TMP) to prevent network starvation

---

- ***The Greek Connection: This balanced ecosystem of GPU vision and CPU geometry serves as the exact engineering blueprint for the new PHAROS AI Factory***

Writing the Bash File

```bash
#!/bin/bash
#SBATCH -A ...........          # Alvis allocation
#SBATCH -p alvis                # Partition
#SBATCH -t 24:00:00             # Time limit (24 hours)
#SBATCH --gpus-per-node=A100:4  # 4x A100 GPUs
#SBATCH -N 1                    # 1 node
#SBATCH -J vessel_seg           # Job name
#SBATCH -o vessel_seg_%j.out    # Output file
#SBATCH -e vessel_seg_%j.err    # Error file

# ============================================================
# SLURM SUBMISSION SCRIPT FOR ALVIS CLUSTER
# Deep Vascular Segmentation V4.3 - 4x A100 Training
# ============================================================
# ============================================================

# Configuration
DATA_DIR="./training_data"
OUTPUT_DIR="./models_a100"
CONFIG="a100_4gpu"

# Environment setup
module purge
module load PyTorch/2.1.0-foss-2023a-CUDA-12.1.1

# Activate the virtual environment
# source /path/to/............/bin/activate

# Print job info
echo "=================================================="
echo "DEEP VASCULAR SEGMENTATION V4.3 - A100 TRAINING"
echo "=================================================="
echo "Job ID:        $SLURM_JOB_ID"
echo "Node:          $SLURMD_NODENAME"
echo "GPUs:          4x A100"
echo "Config:        $CONFIG"
echo "Data dir:      $DATA_DIR"
echo "Output dir:    $OUTPUT_DIR"
echo "Start time:    $(date)"
echo "=================================================="

# NCCL settings for multi-GPU
export NCCL_DEBUG=INFO
export NCCL_IB_DISABLE=0
export NCCL_NET_GDR_LEVEL=2

# PyTorch optimizations for A100
export TORCH_CUDA_ARCH_LIST="8.0"  # A100 architecture
export PYTORCH_CUDA_ALLOC_CONF=max_split_size_mb:512

# Get master address and port for distributed training
export MASTER_ADDR=$(hostname)
export MASTER_PORT=$(python -c "import socket; s=socket.socket(); s.bind(('',0)); print(s.getsockname()[1]); s.close()")

echo ""
echo "Distributed settings:"
echo "  MASTER_ADDR: $MASTER_ADDR"
echo "  MASTER_PORT: $MASTER_PORT"
echo ""

# Check GPU availability
nvidia-smi

# Run distributed training
echo ""
echo "Starting distributed training..."
echo ""

torchrun --nproc_per_node=4 \
    --master_addr=$MASTER_ADDR \
    --master_port=$MASTER_PORT \
    deep_vascular_seg_v4_3_integrated.py train \
    --data-dir $DATA_DIR \
    --output-dir $OUTPUT_DIR \
    --config $CONFIG

# Check exit status
if [ $? -eq 0 ]; then
    echo ""
    echo "=================================================="
    echo "TRAINING COMPLETED SUCCESSFULLY"
    echo "=================================================="
    echo "End time:    $(date)"
    echo "Output dir:  $OUTPUT_DIR"
    echo "=================================================="
else
    echo ""
    echo "=================================================="
    echo "TRAINING FAILED"
    echo "=================================================="
    echo "Check error log: vessel_seg_${SLURM_JOB_ID}.err"
    exit 1
fi
```

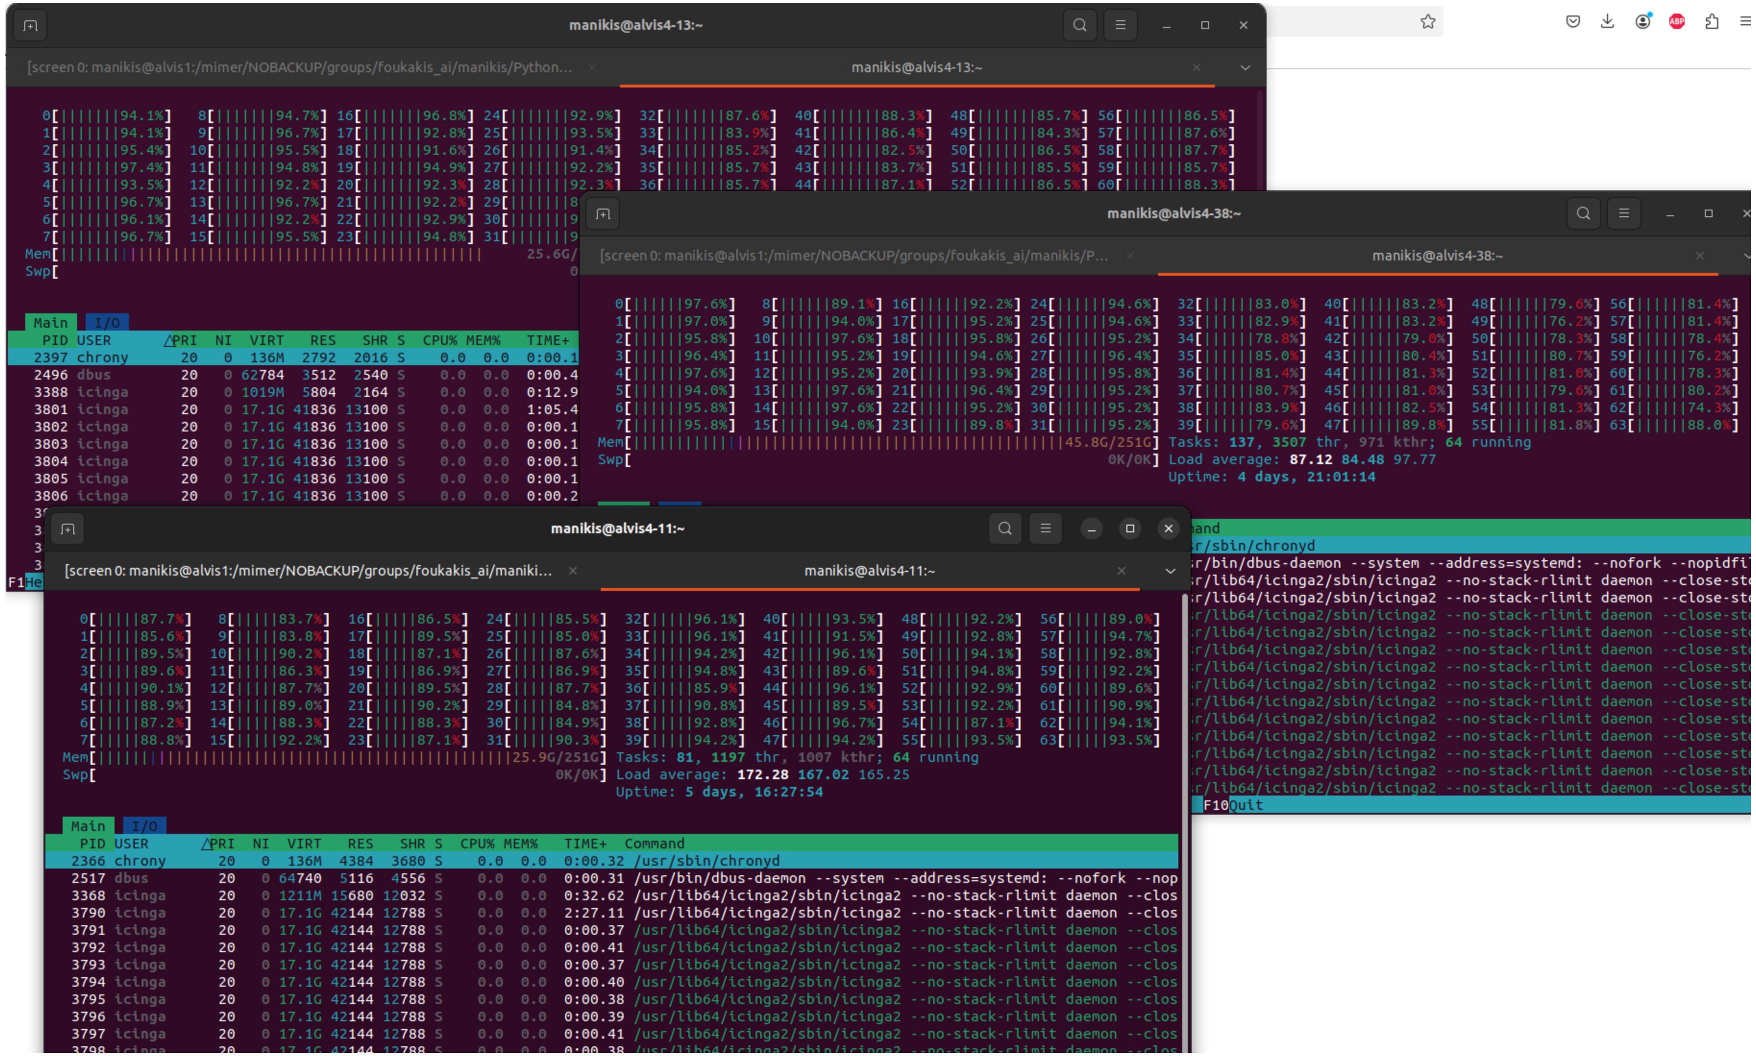

RUNNING CPUS

THE EVOLUTION OF PATHOLOGY AI

THE TRANSITION

Trying to solve this with attention-based mechanisms

Handcrafted Era

Find and segment cells

Code the exact shape, size, and texture

Deploy machine learning models

Deep Learning Era

Give the model raw pixels, the final diagnosis and say: 'figure out the pattern'

Because the images are so massive, we used Multiple Instance Learning, chopping the slide into thousands of tiny puzzle pieces

The model looks at every piece and makes a decision

However, looking at tiny isolated patches means the AI loses the 'big picture'—the overall spatial architecture of the tumor

That exact limitation is what drove the invention of the Gigapixel Foundation Models and Spatial Graphs that we use today

Karolinska Institutet, Department of Oncology-Pathology, Sweden

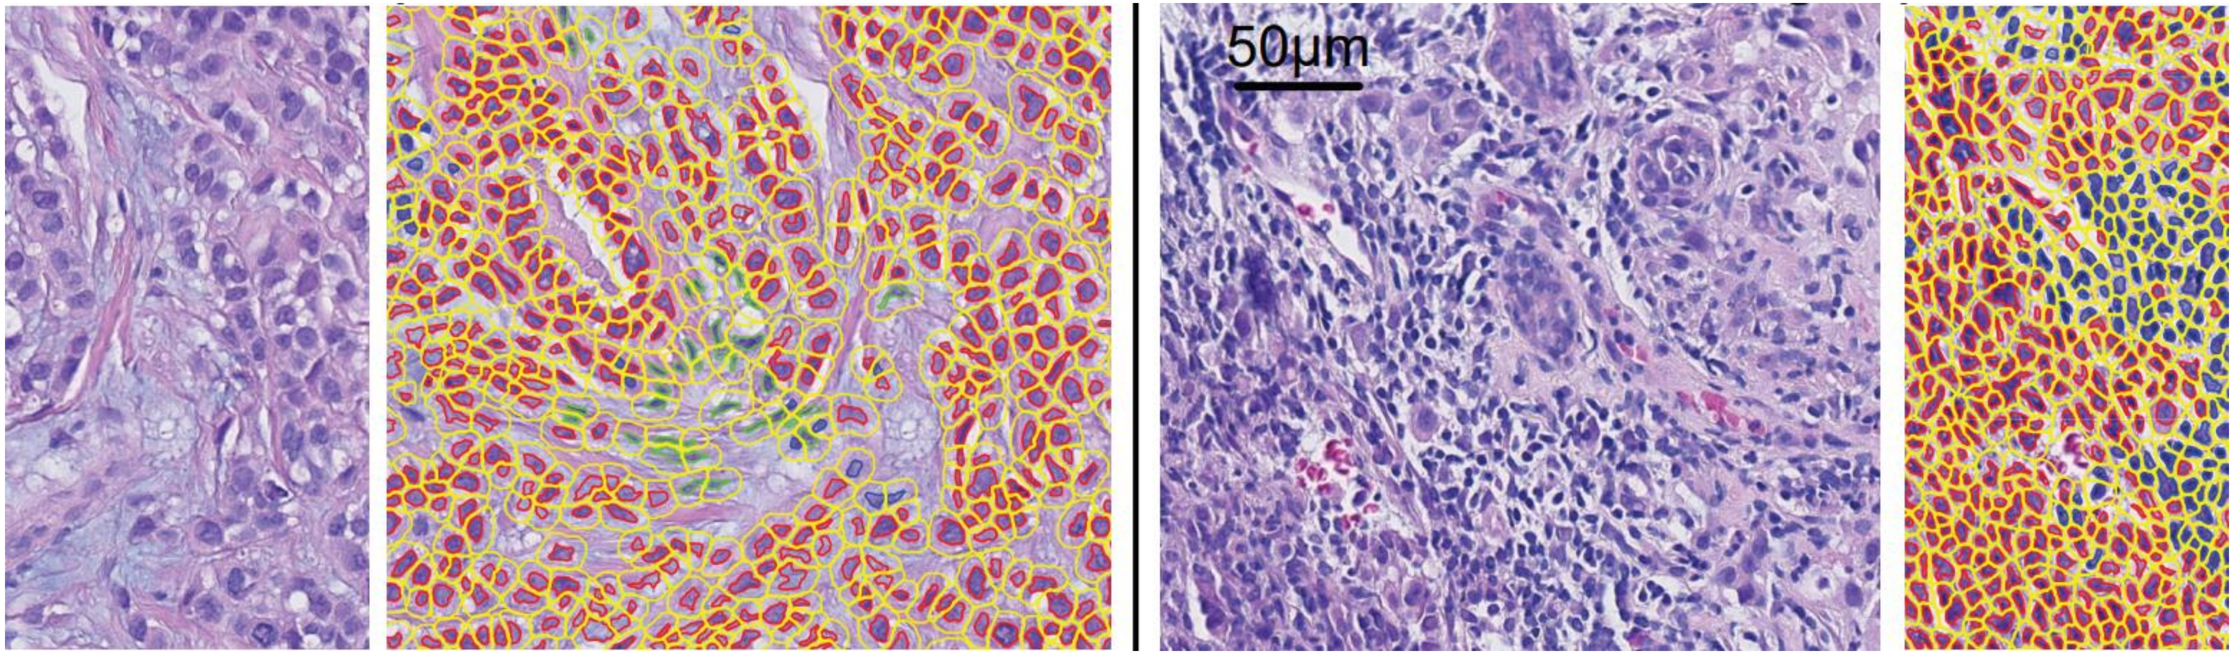

THE HANDCRAFTED MACHINE LEARNING PARADIGM

- The traditional computational pathology pipeline relies on explicit feature engineering
- After segmenting nuclei or tissue, researchers extract handcrafted features such as nuclear size, shape irregularity, texture descriptors (e.g., Haralick features), and spatial statistics
- These features were then fed into classical machine learning models such as Support Vector Machines, Random Forests, or Cox proportional hazards models

THE HANDCRAFTED MACHINE LEARNING PARADIGM

- This approach provides interpretability and biological insight but requires substantial domain expertise and often struggled to generalize across institutions

- ***This era is strictly limited to standard CPU processing!***

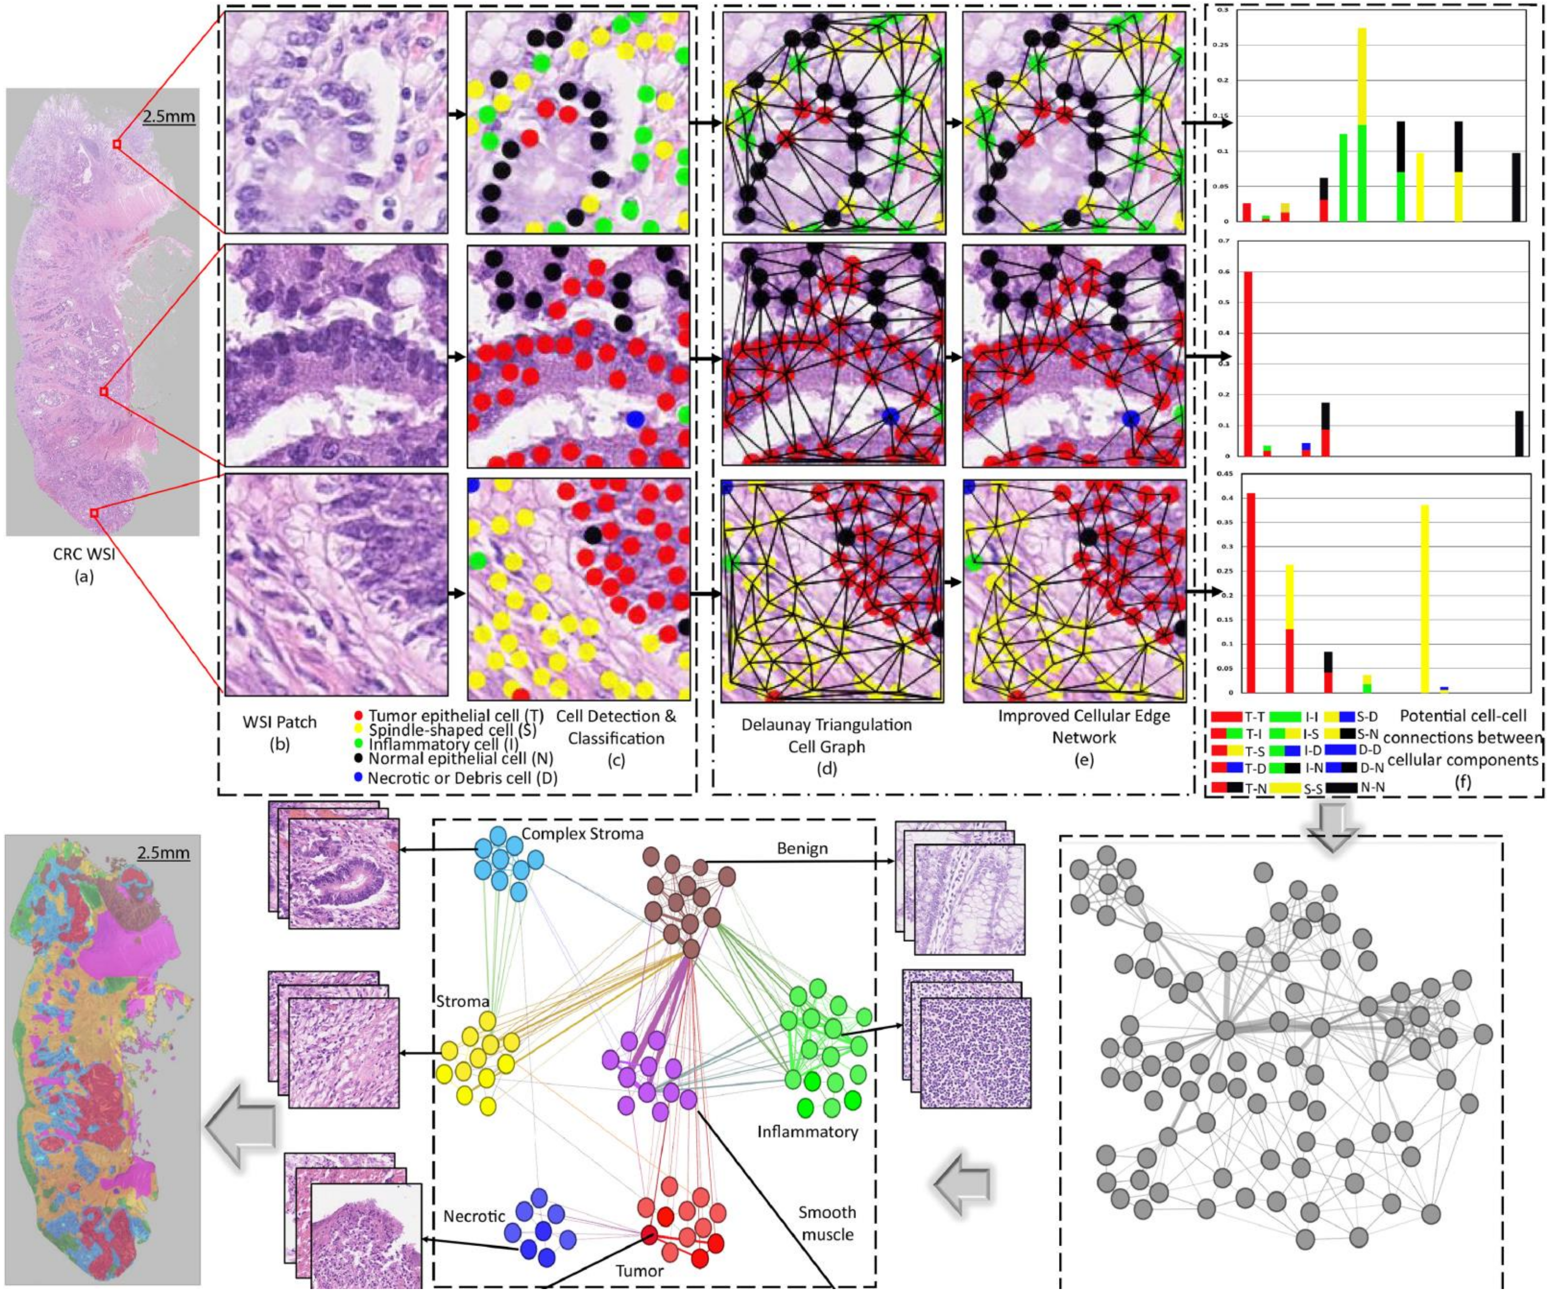

Javed S, Mahmood A, Fraz MM, Koohbanani NA, Benes K, Tsang YW, Hewitt K, Epstein D, Snead D, Rajpoot N. Cellular community detection for tissue phenotyping in colorectal cancer histology images. Med Image Anal. 2020 Jul;63:101696.

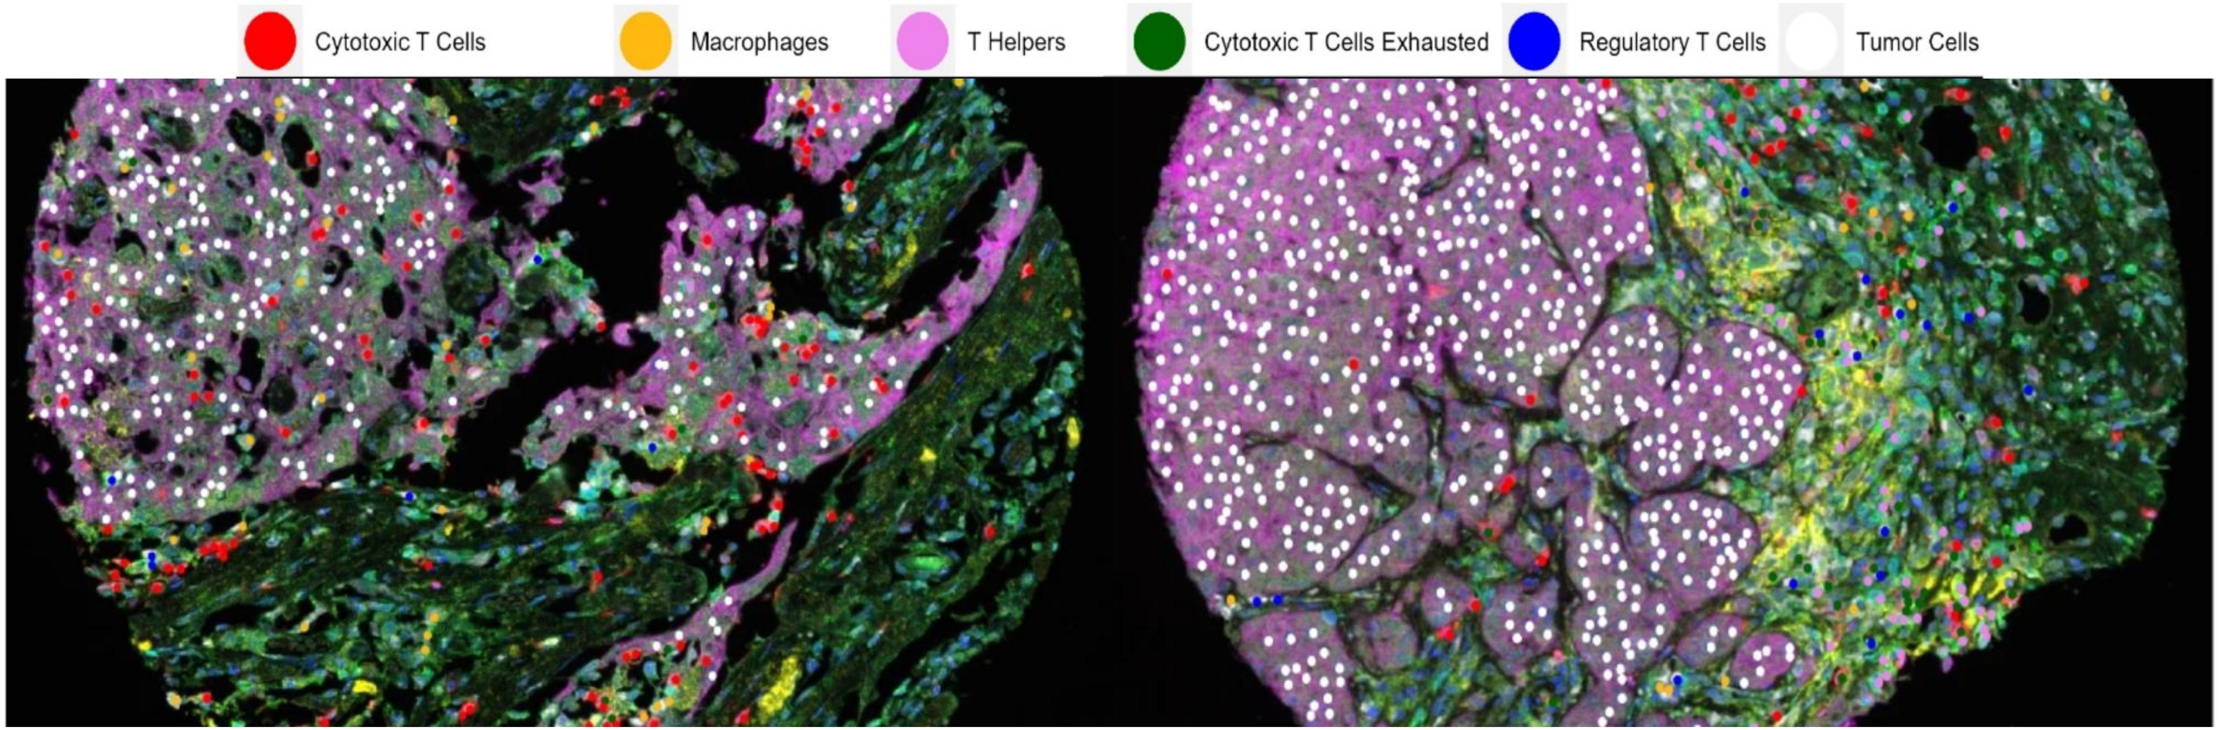

SHOWCASING THE HANDCRAFTED ERA

- High-Throughput Cell Spatial Analytics Using Parallel CPU Processing
- Millions of detected cells per cohort
- Spatial graph construction at whole-slide scale (and multiple TME)
- ***Parallel CPU-based processing (multi-core execution)***
- Scalable to large multi-center datasets

Zerdes, I., Matikas, A., Mezheyeuski, A., Manikis, G., et al. Machine learning-based spatial characterization of tumor-immune microenvironment in the EORTC 10994/BIG 1-00 early breast cancer trial. npj Breast Cancer 11, 23 (2025)

- High-Throughput Cell Spatial Analytics Using Parallel CPU Processing
- Two panels (immune + stromal)
- Multiple WSIs, millions of cells
- ***HPC analysis using multiple CPUs***

SHOWCASING THE HANDCRAFTED ERA

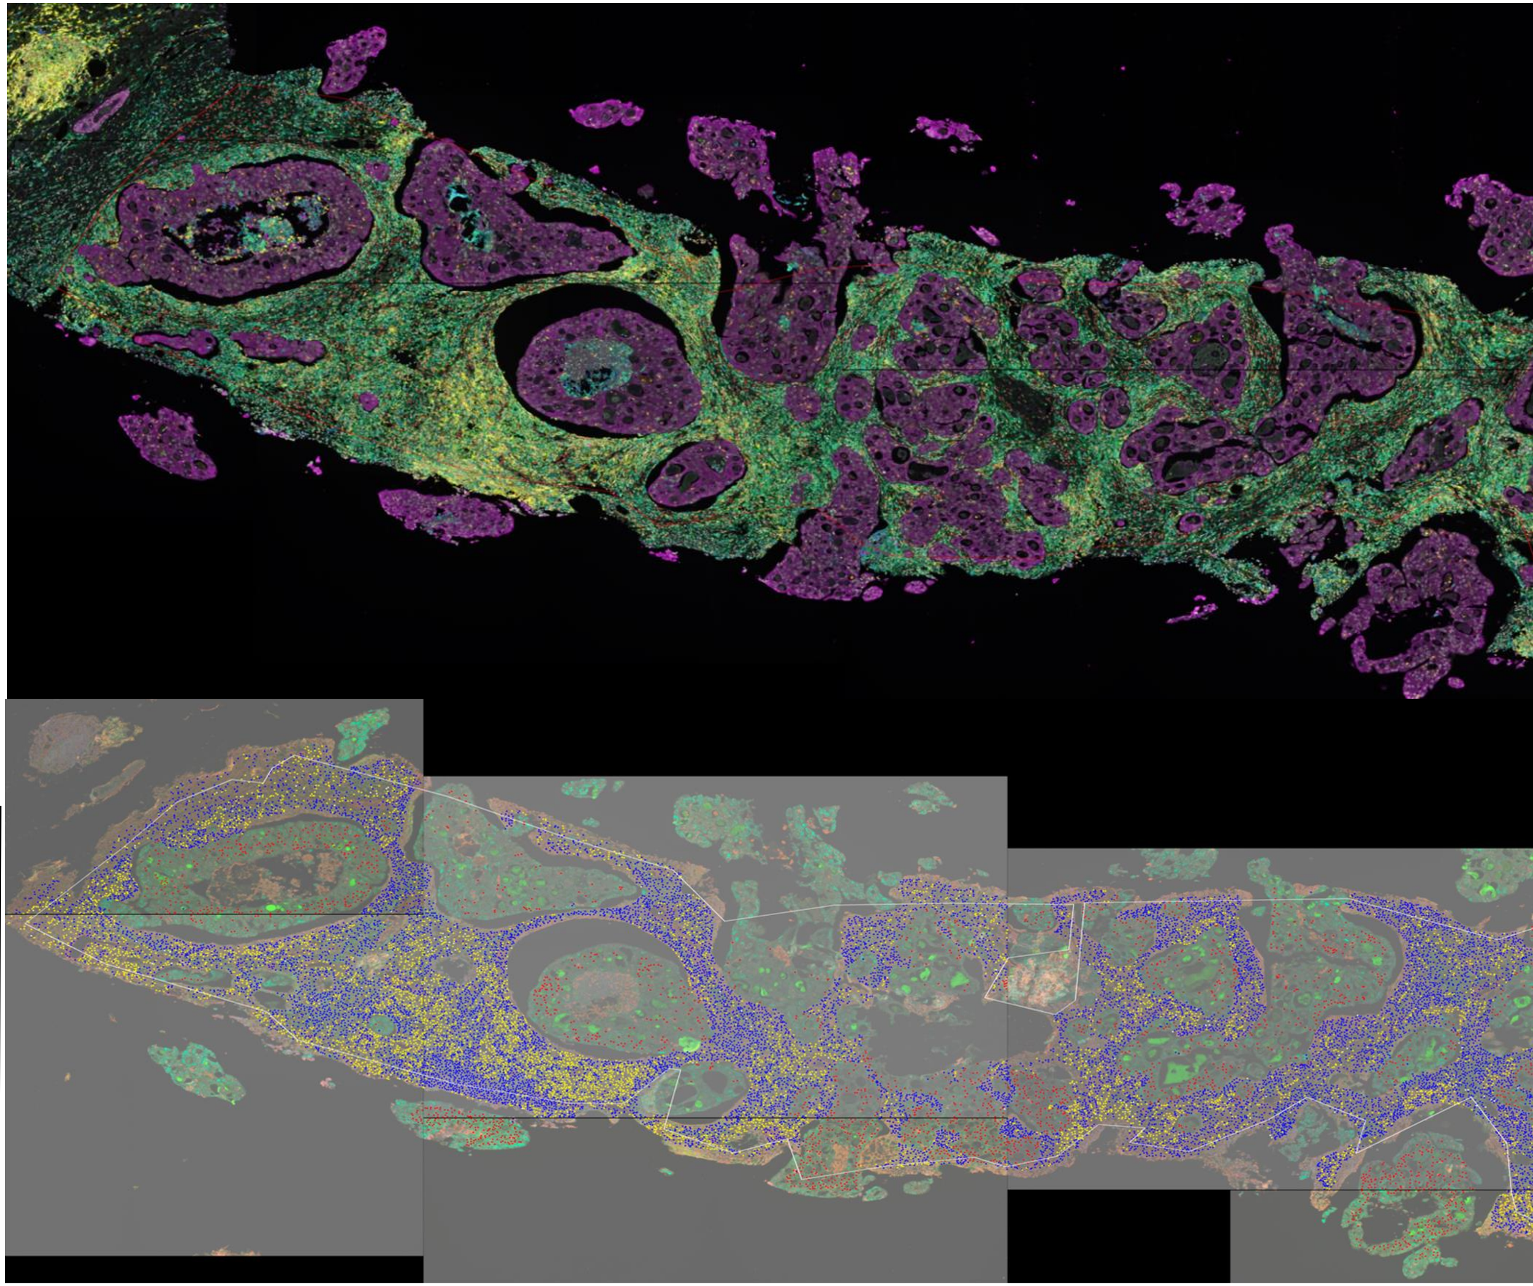

Manikis, Georgios, et al. "PDGFRa plus/FAP-fibroblasts at multiplex in situ resolution are associated with resistance to HER2 blockade in neoadjuvant treated HER2+ breast cancer: results from the Swedish PREDIX HER2 trial." CLINICAL CANCER RESEARCH. Vol. 31. No. 12. 615 CHESTNUT ST, 17TH FLOOR, PHILADELPHIA, PA 19106-4404 USA: AMER ASSOC CANCER RESEARCH, 2025.

SHOWCASING THE HANDCRAFTED ERA

- High-Throughput Cell Spatial Analytics Using Parallel CPU Processing
- H&E WSIs
- > 10 spatial analysis techniques
- Machine learning analysis for pathological complete response (pCR)

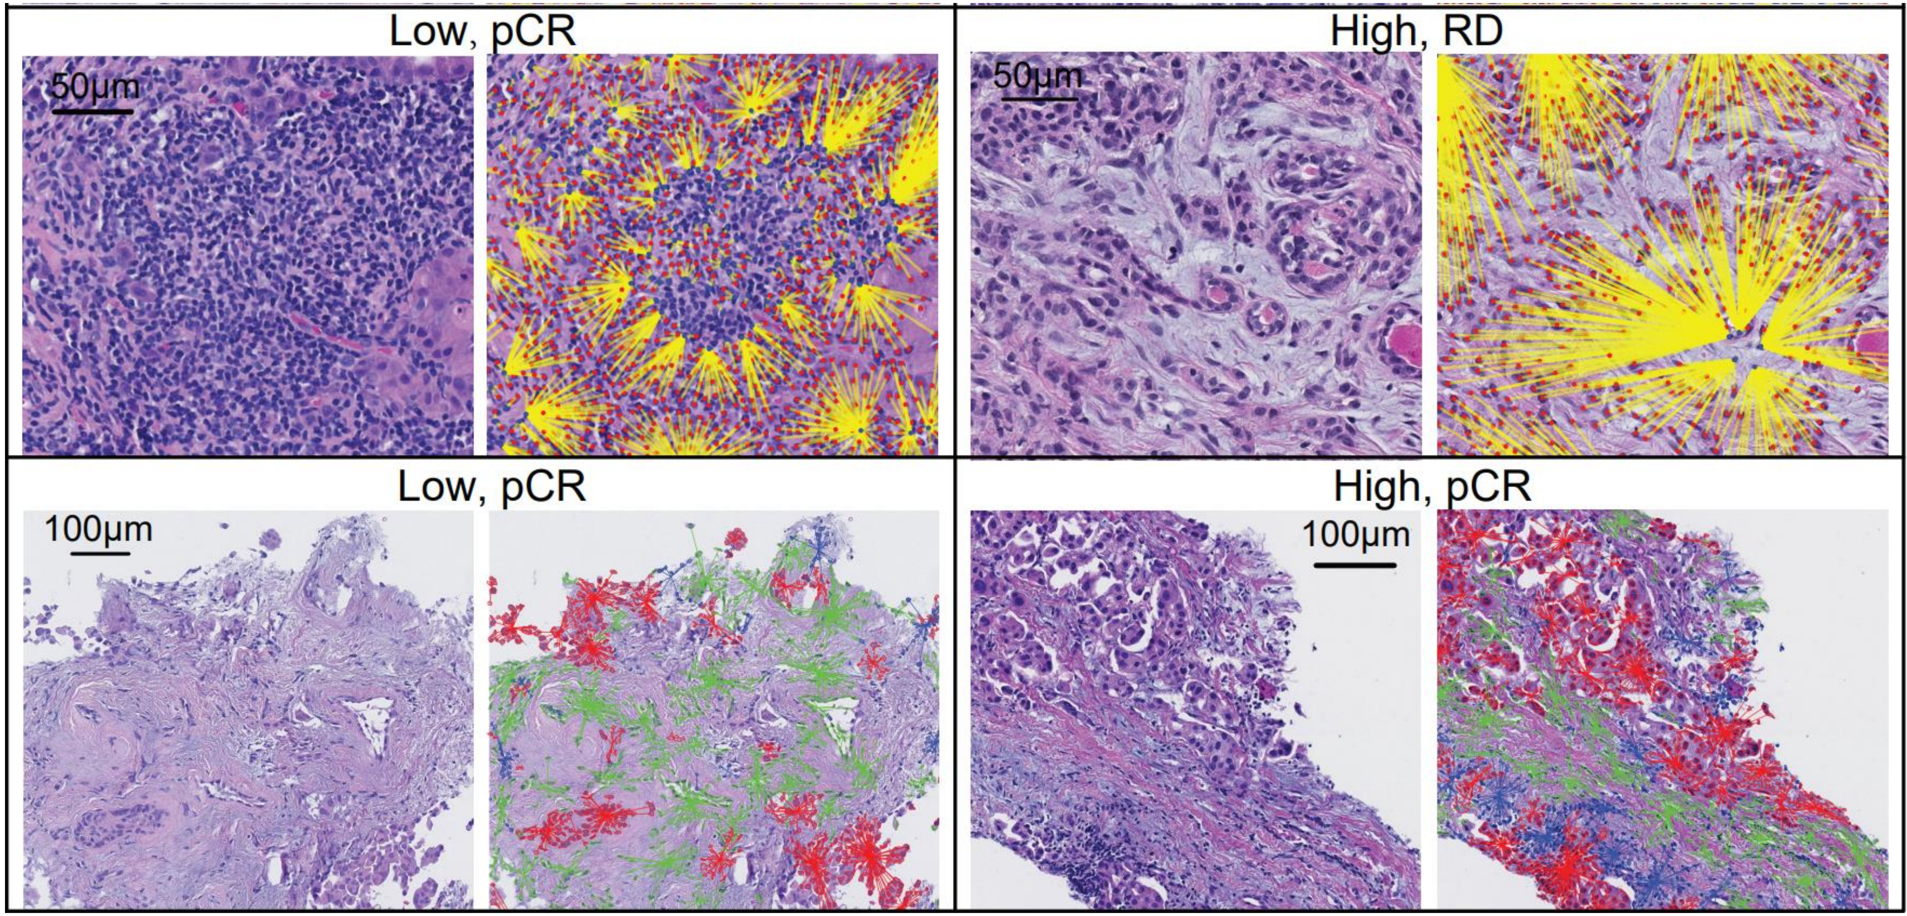

THE END-TO-END DEEP LEARNING PARADIGM

- End-to-end deep learning shifted the paradigm by allowing neural networks to learn representations directly from raw image patches
- Instead of manually defining features, deep learning models automatically discover hierarchical visual patterns

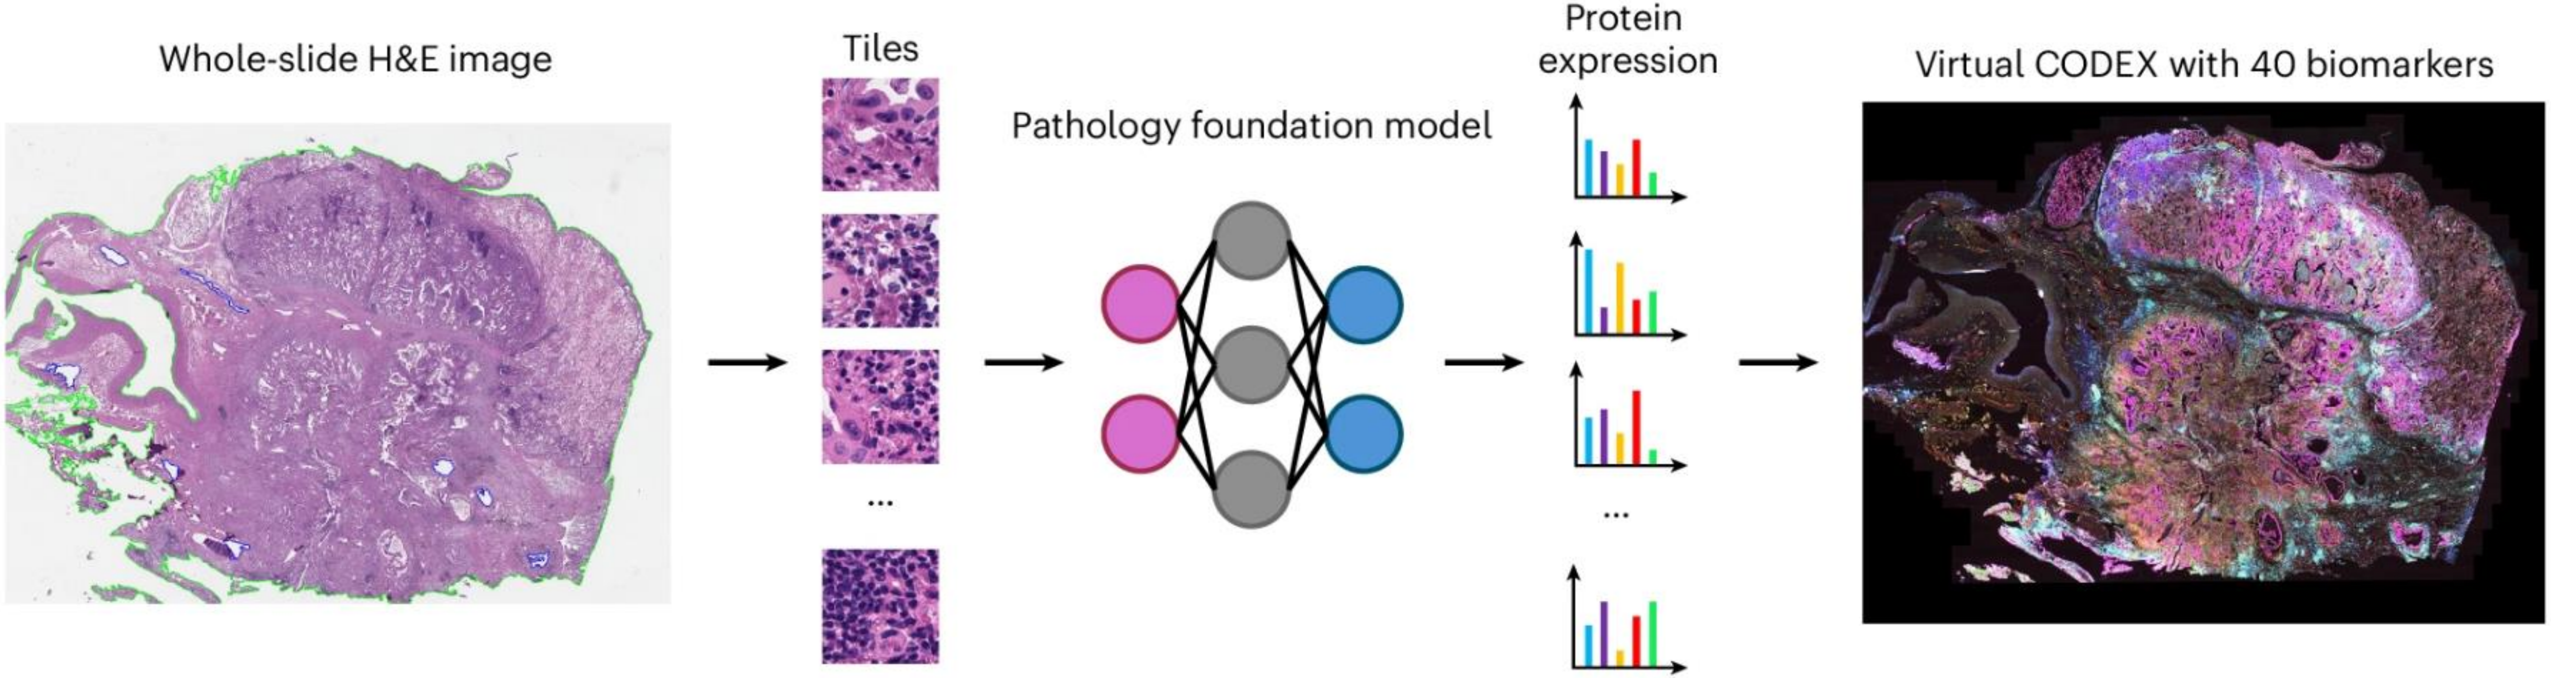

Li, Z., Li, Y., Xiang, J. et al. AI-enabled virtual spatial proteomics from histopathology for interpretable biomarker discovery in lung cancer. Nat Med 32, 231–244 (2026)

The price of automation is scale….. and scale requires infrastructure

THE END-TO-END DEEP LEARNING PARADIGM

- The computational workload shifted entirely from the CPU to the parallel processing cores of the GPU

- **This created an immediate engineering crisis**: a gigapixel WSI requires 30GB of memory, which far exceeds the VRAM limits of standard GPUs when calculating deep learning activation maps

- Because a full WSI exceeds VRAM limits, the standard solution over the last decade was to chop the gigapixel image into thousands of tiny 256x256 pixel tiles (instances).

- The GPU analyzes each tile independently. If a single tile is classified as malignant by the model, the entire slide (the "bag" of instances) is flagged as positive for cancer

New architectures use dense neural layers with attention mechanisms to predict outcome

UNVEILING THE POWER OF MODEL-AGNOSTIC MULTISCALE ANALYSIS FOR ENHANCING ARTIFICIAL INTELLIGENCE MODELS IN BREAST CANCER HISTOPATHOLOGY IMAGES

- A model-agnostic multiresolution feature aggregation framework tailored for the analysis of H&E slides in the context of breast cancer (TP53 mutation status prediction and survival prediction), on a multicohort dataset of 2038 patient samples

- Image patches from a high (20x) and low (10x) magnification

- Cross-scale analysis benefits the explainability aspects of attention-based architectures, since one can extract attention maps at the tissue- and cell-levels, improving the interpretation of the model's decision

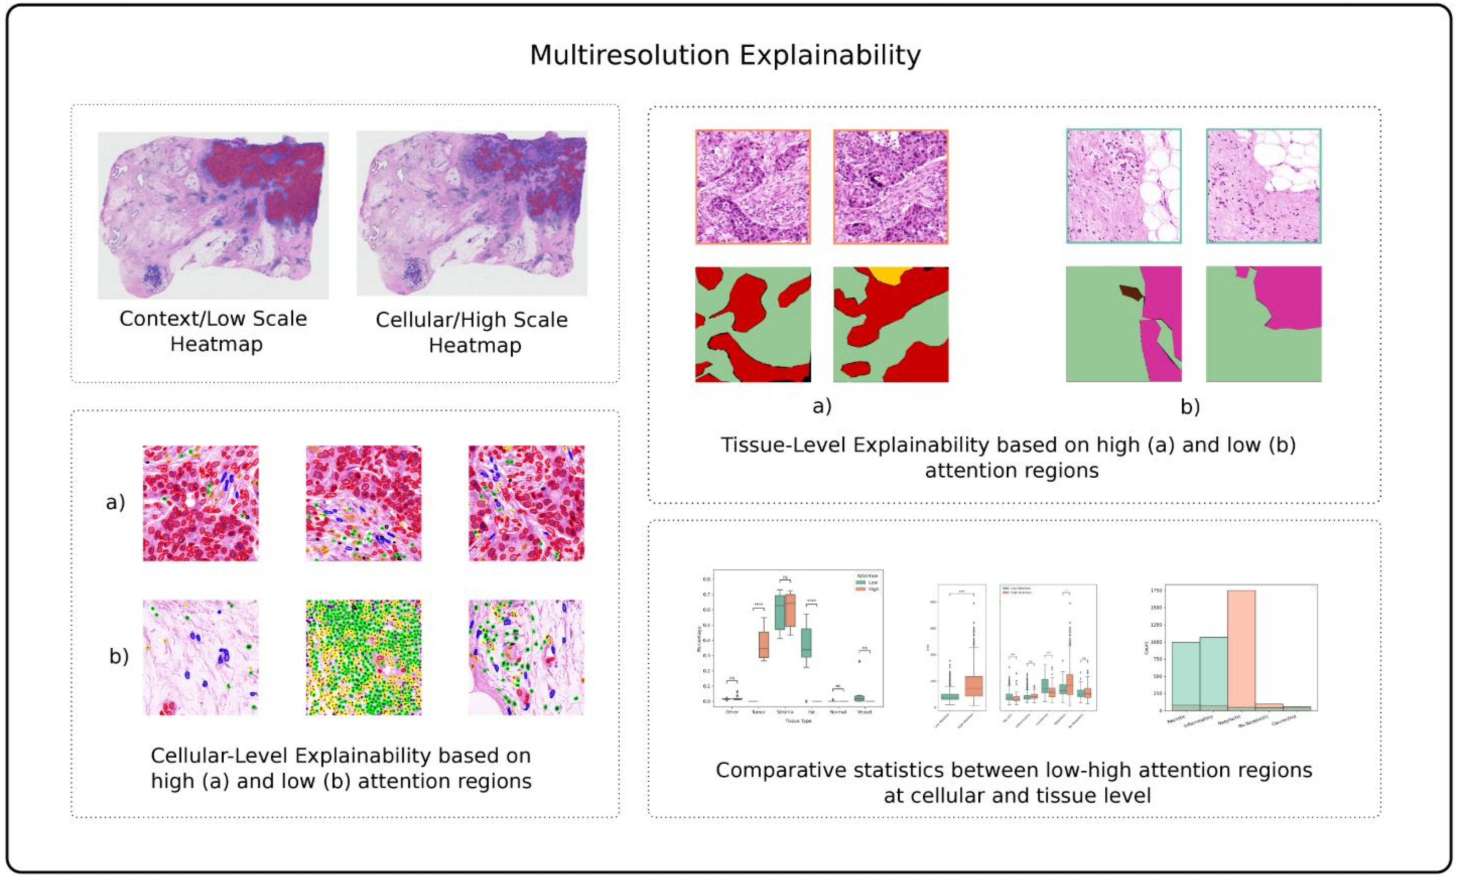

Analysis of high- and low-attention regions to identify characteristics of the cellular and tissue-architectural compositions

N. Tsiknakis, G. Manikis, et al., "Unveiling the Power of Model-Agnostic Multiscale Analysis for Enhancing Artificial Intelligence Models in Breast Cancer Histopathology Images," in IEEE Journal of Biomedical and Health Informatics, vol. 28, no. 9, pp. 5312-5322, Sept. 2024

AI FOUNDATION MODELS - A NEW CO-PILOT FOR CLINICAL DECISION MAKING IN HISTOPATHOLOGY

- **The Challenge**: We face an unprecedented explosion of data for every patient—genomics, imaging, pathology, clinical notes, and rapidly evolving literature
- **The Problem**: Synthesizing this complex information at the point of care is becoming humanly impossible
- **The Opportunity**: What if we had an intelligent tool to help us connect the dots, personalized for each patient?

Traditional AI models:
Train on a
single, specific task

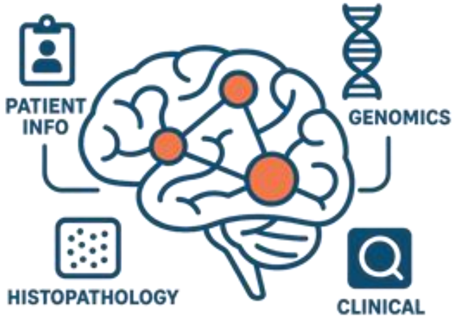

- Pre-trained on a massive, diverse set of general data (millions of oncology cases, research papers, clinical trials, and imaging data)
- It learns the fundamental patterns, context, and "language" of the data
- One model, many applications (vs. building separate AI for each task)

BUILDING AN ONCOLOGY-SPECIFIC FOUNDATION MODEL

- **How do we make it an expert in cancer care? We fine-tune it with our data**
- An oncology foundation model is trained on a comprehensive range of cancer-specific data sources to understand the unique language and complexities of our field:
- Electronic Medical Records (EMR): Patient histories, lab values, treatment cycles, and physician notes
- Genomic and Molecular Data: NGS reports, tumor markers, and gene expression profiles
- Imaging & Pathology: Radiology reports, digital pathology slide interpretations, and molecular pathology results
- Medical Literature: Clinical trial publications, treatment guidelines (e.g., NCCN), and research papers

---

- **The Result**: A model that understands oncology not as isolated data points, but as an integrated clinical picture

THE COMPUTATIONAL PIPELINE: FROM RAW TISSUE TO CLINICAL OUTPUT

RAW DATA INGESTION

- **The Task**: Reading a 30GB multi-resolution ndpi or svs scanner file into memory
- **Handcrafted Approach**: Loading single, small Regions of Interest (ROIs) manually into local RAM for simple observation
- **AI Approach**: Parsing the entire 40x magnification layer of 5,000 slides simultaneously to build a training dataset
- **Alvis HPC Execution**: We utilize the high-core-count CPU nodes on Alvis. We launch a Slurm job using OpenSlide to asynchronously read hundreds of gigapixel chunks in parallel from the Cephyr central storage, bypassing single-thread limits.

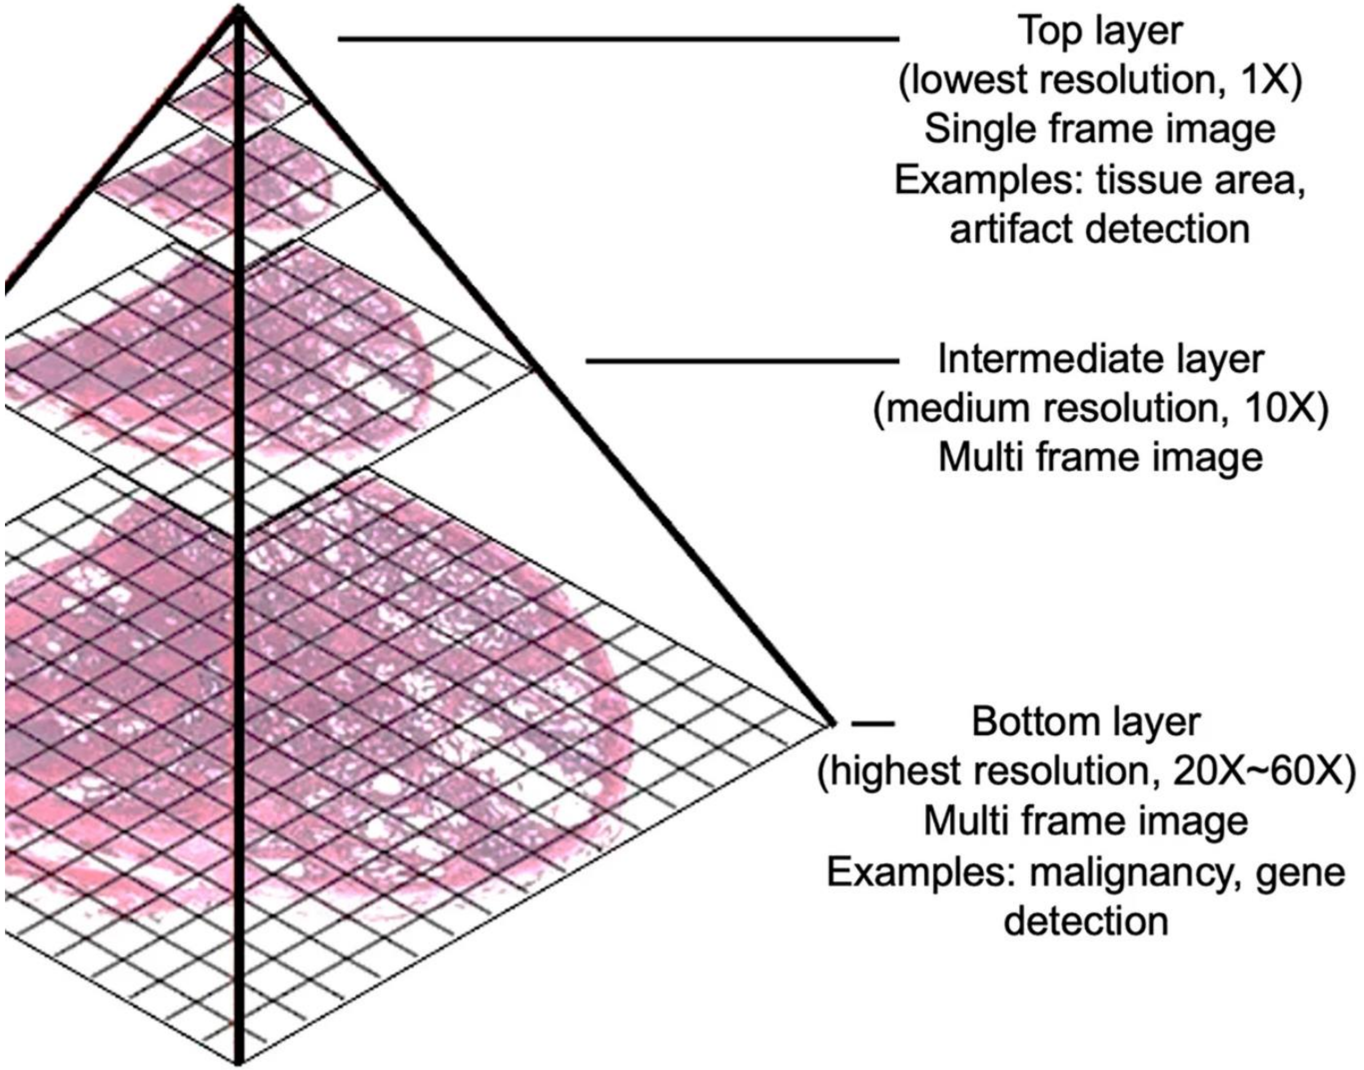

Hossain, M.S., Shahriar, G.M., Syeed, M.M.M. et al. Region of interest (ROI) selection using vision transformer for automatic analysis using whole slide images. Sci Rep 13, 11314 (2023)

TISSUE MASKING & TILING

- **The Task**: ~70% of a WSI is empty white glass or wax. We must filter this out and cut the tissue into thousands of analyzable tiles

- **Handcrafted Approach**: Simple Otsu thresholding to find tissue, saving individual tiles as millions of .png files (e.g. Size of 256x256)

- **AI Approach**: AI-driven semantic segmentation to identify active tumor zones, saving coordinates rather than discrete images

- **Alvis HPC Execution**: Saving millions of .png files instantly crashes the Alvis Lustre file system. We strictly save tiled data into continuous HDF5 or WebDataset (.tar) formats to protect the C3SE metadata servers

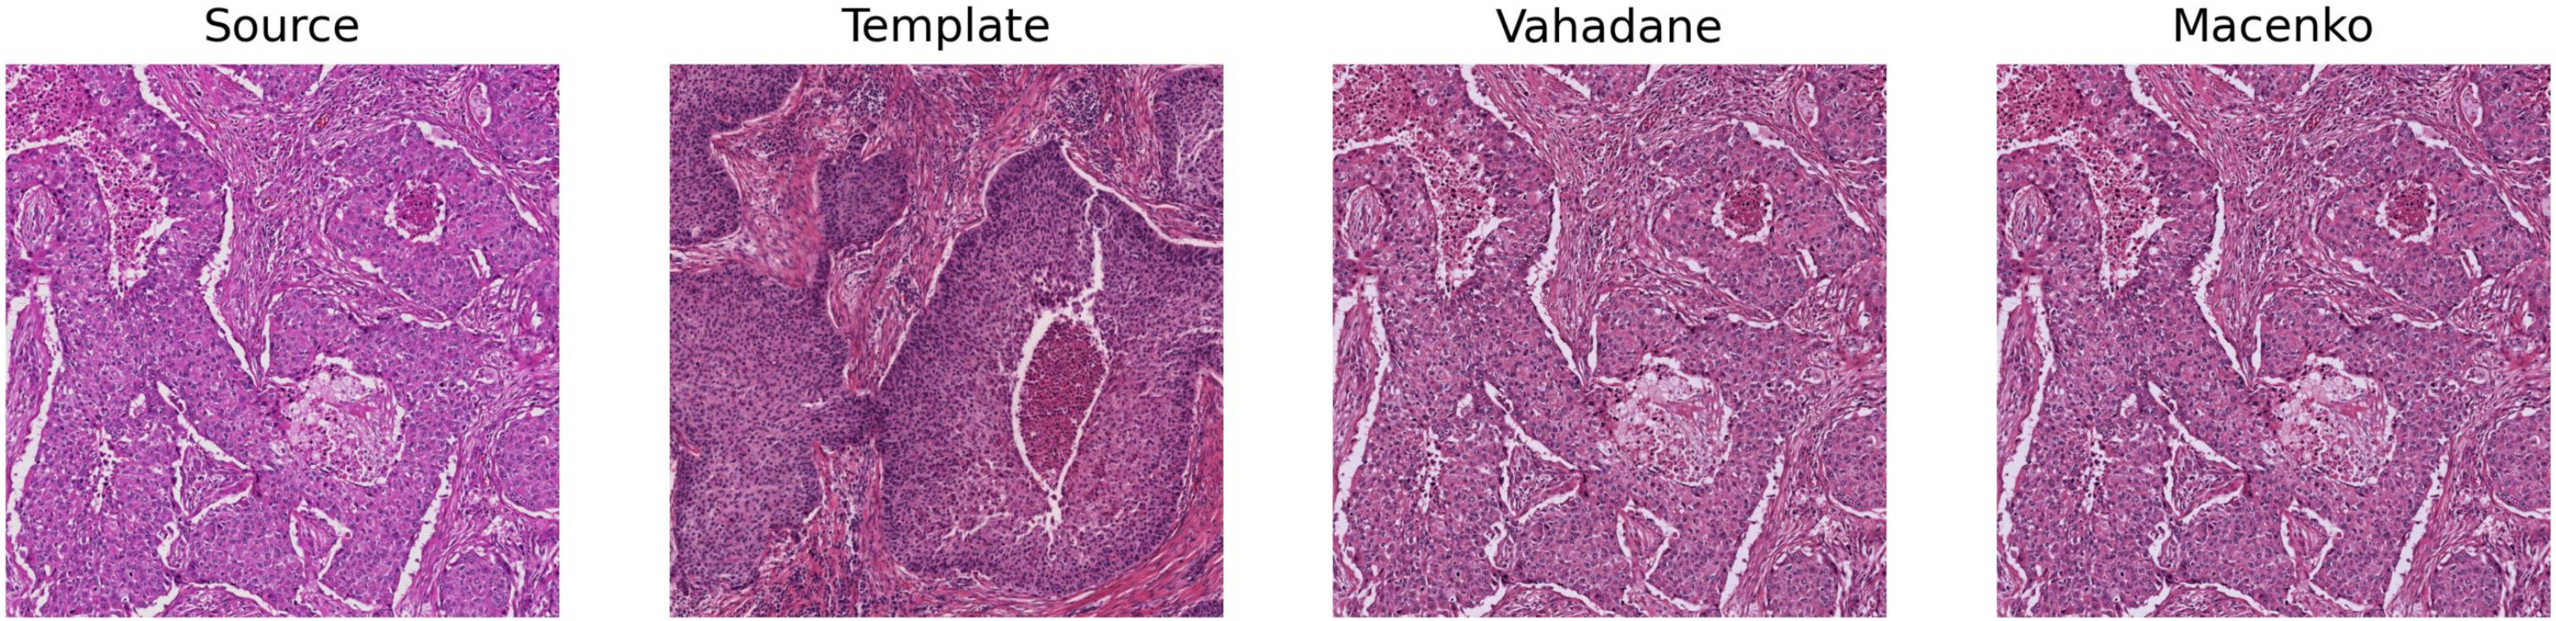

github.com/CielAl/torch-staintools

COLOR NORMALIZATION STANDARDIZING THE LABORATORY STAINS

- **The Task**: Correcting the massive color variations caused by different hospital staining chemicals and digital scanners
- **Handcrafted Approach**: Ignoring color shift or manually adjusting RGB histograms on a slide-by-slide basis
- **AI Approach**: Applying complex mathematical algorithms (Macenko, VAHADANE) to map every pixel to a universal reference color space
- **HPC Execution**: We execute a Slurm array job across tens of **Alvis CPU cores**. This distributed processing normalizes hundreds/thousands slides offline, saving the clean tensors back to Cephyr storage before we ever request an expensive GPU node

SEMANTIC TISSUE SEGMENTATION

- **The Task**: We must accurately map the macroscopic landscape of the slide before looking at individual cells
- **The Goal**: Separating malignant tumor nests from healthy stroma, necrosis, and immune aggregates
- **AI Approach**: We train AI-based Segmentation models using pathologist-annotated regional polygons. The AI learns to classify every individual pixel into a specific biological tissue category.
- **Alvis HPC Execution: Training these models requires massive batch sizes to learn spatial context. On Alvis, we request multi-GPU nodes (#SBATCH --gpus-per-node=A100_80GB:4). We use PyTorch Distributed Data Parallel (DDP) to synchronize the model gradients across all 4 GPUs, cutting training time from weeks to hours.**

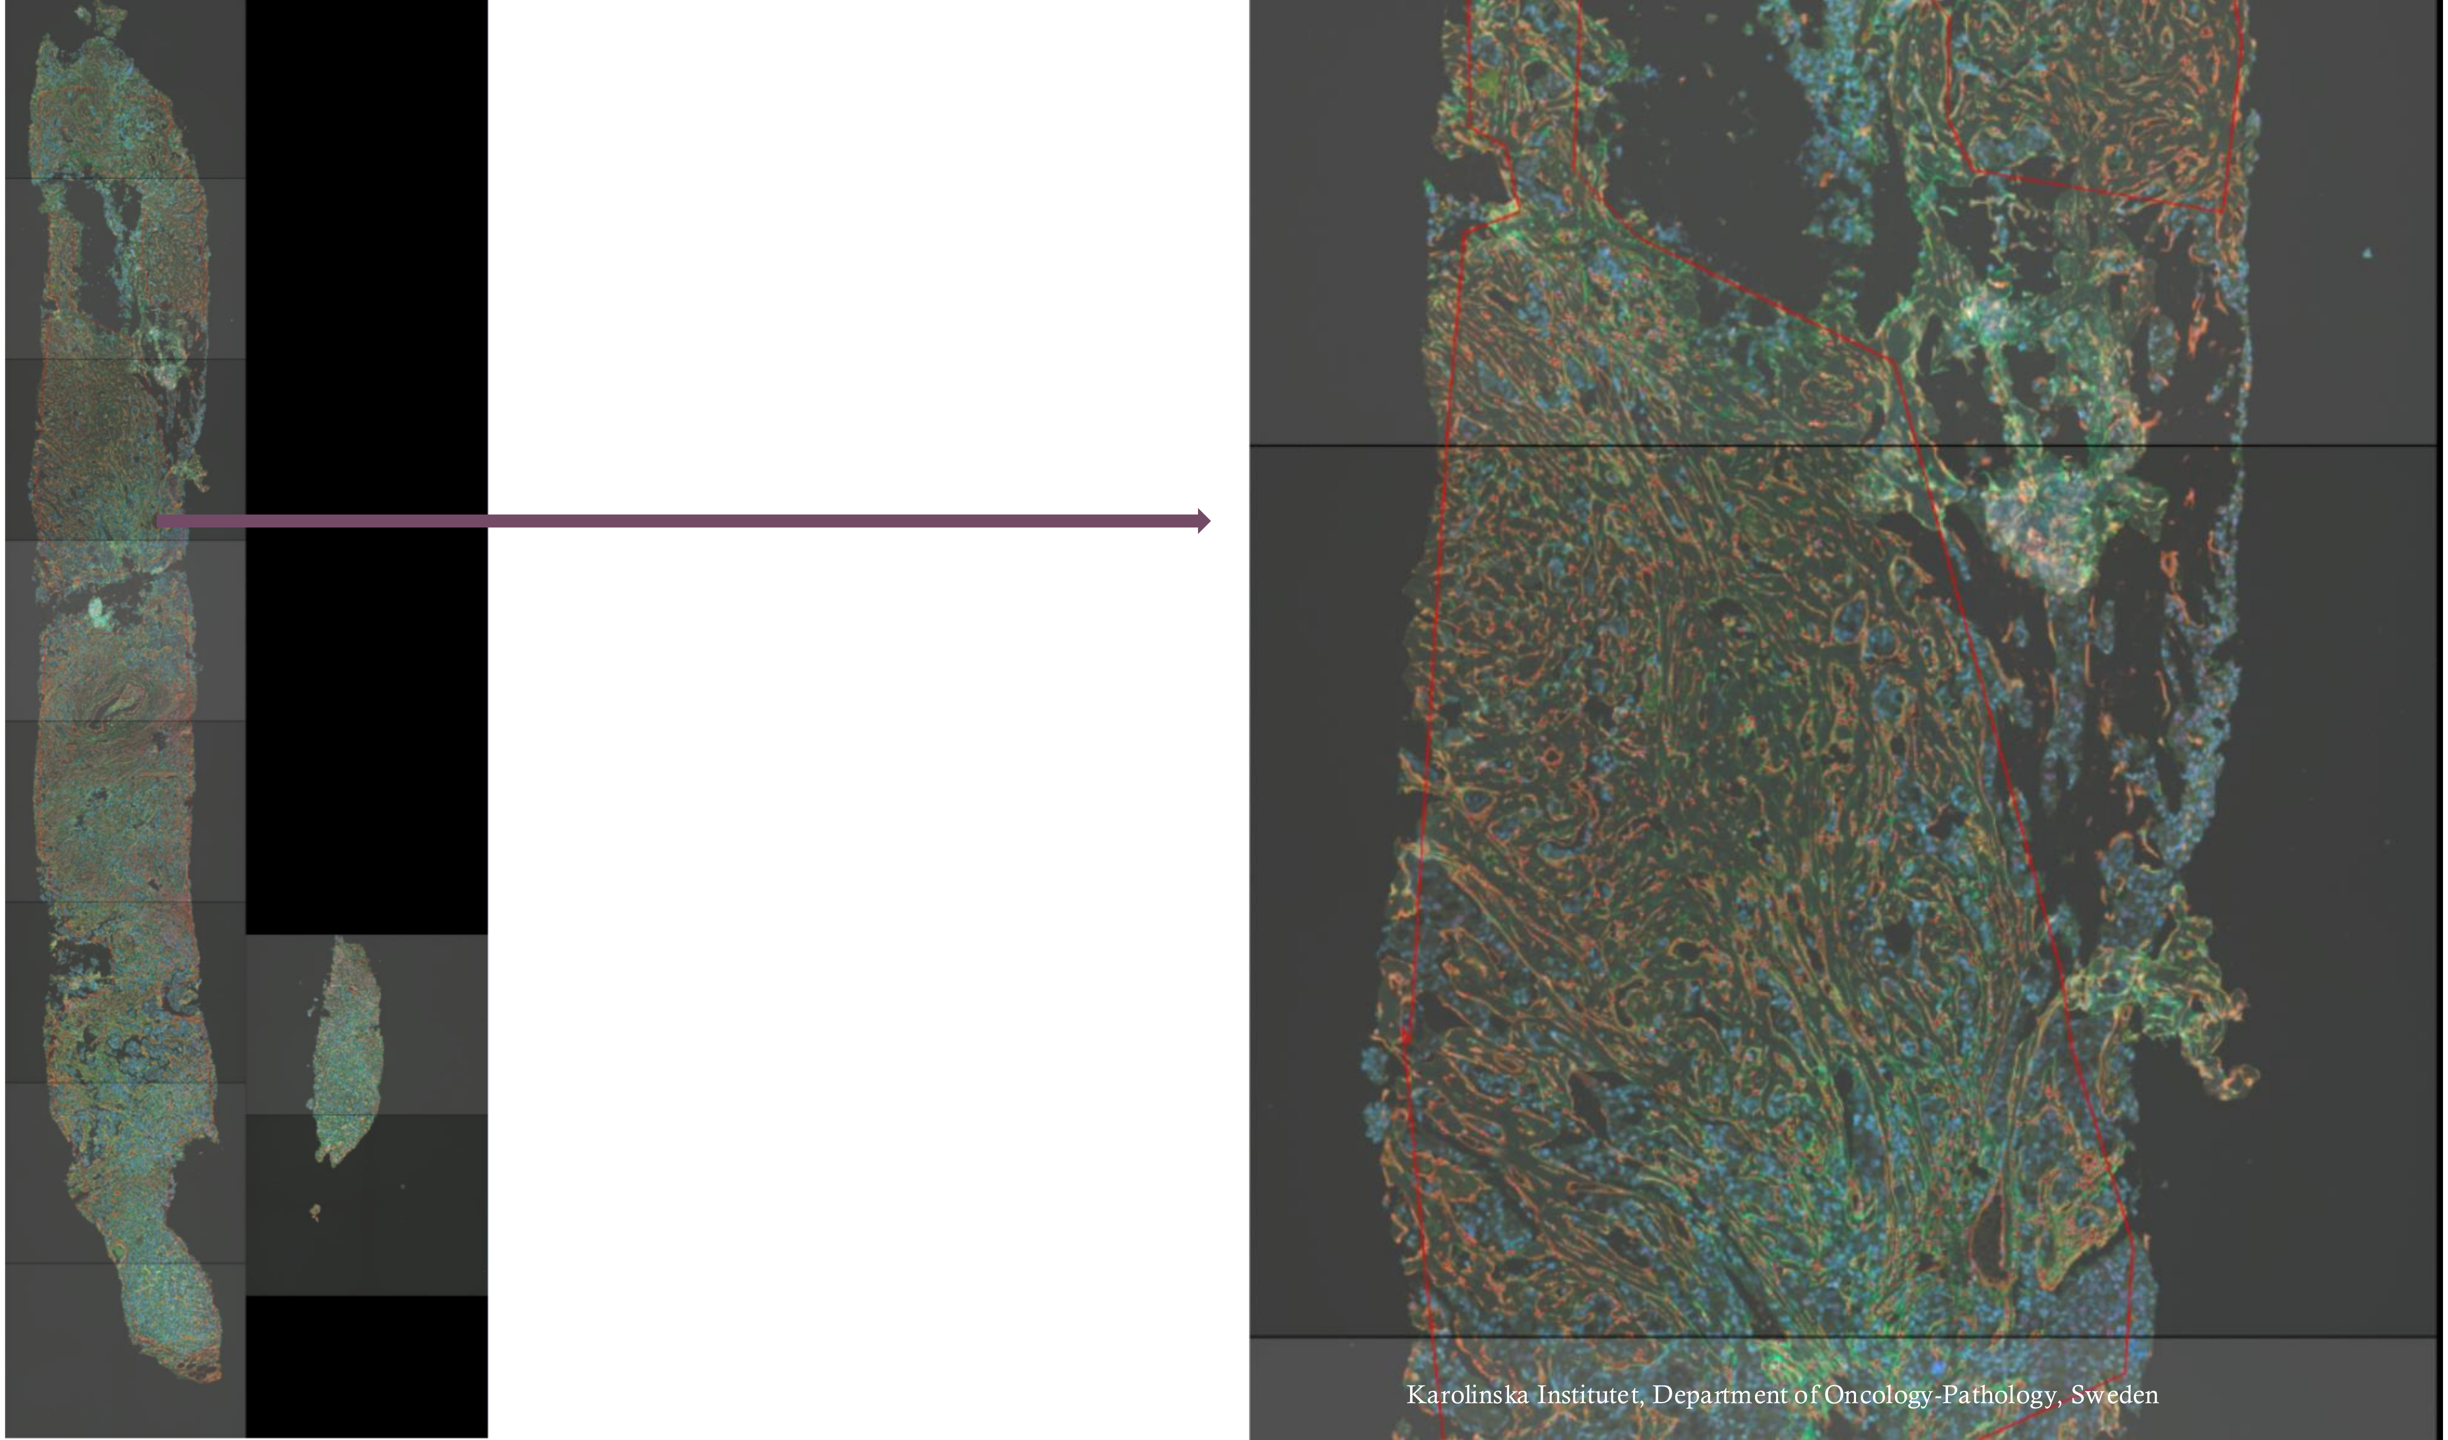
Karolinska Institutet, Department of Oncology-Pathology, Sweden

CELL INSTANCE SEGMENTATION

- **The Task**: We zoom in to 40x magnification to precisely locate and map over 2 million individual cell nuclei per slide
- **The Challenge**: Classical algorithms failed when cells physically overlapped; Deep Learning solves this by predicting the geometric center of each cell
- **AI Approach**: We train Instance Segmentation models (like HoVer-Net or StarDist) on heavily annotated nuclear datasets. The model outputs horizontal and vertical distance maps to mathematically force apart touching cells
- **Alvis HPC Execution**: Predicting dense instance maps for 2 million cells requires trillions of FLOPS. On Alvis, we compile our PyTorch models using torch.compile or NVIDIA TensorRT. This performs "Kernel Fusion" on the A100 Tensor Cores, massively accelerating inference and allowing us to segment a full WSI in minutes instead of hours.

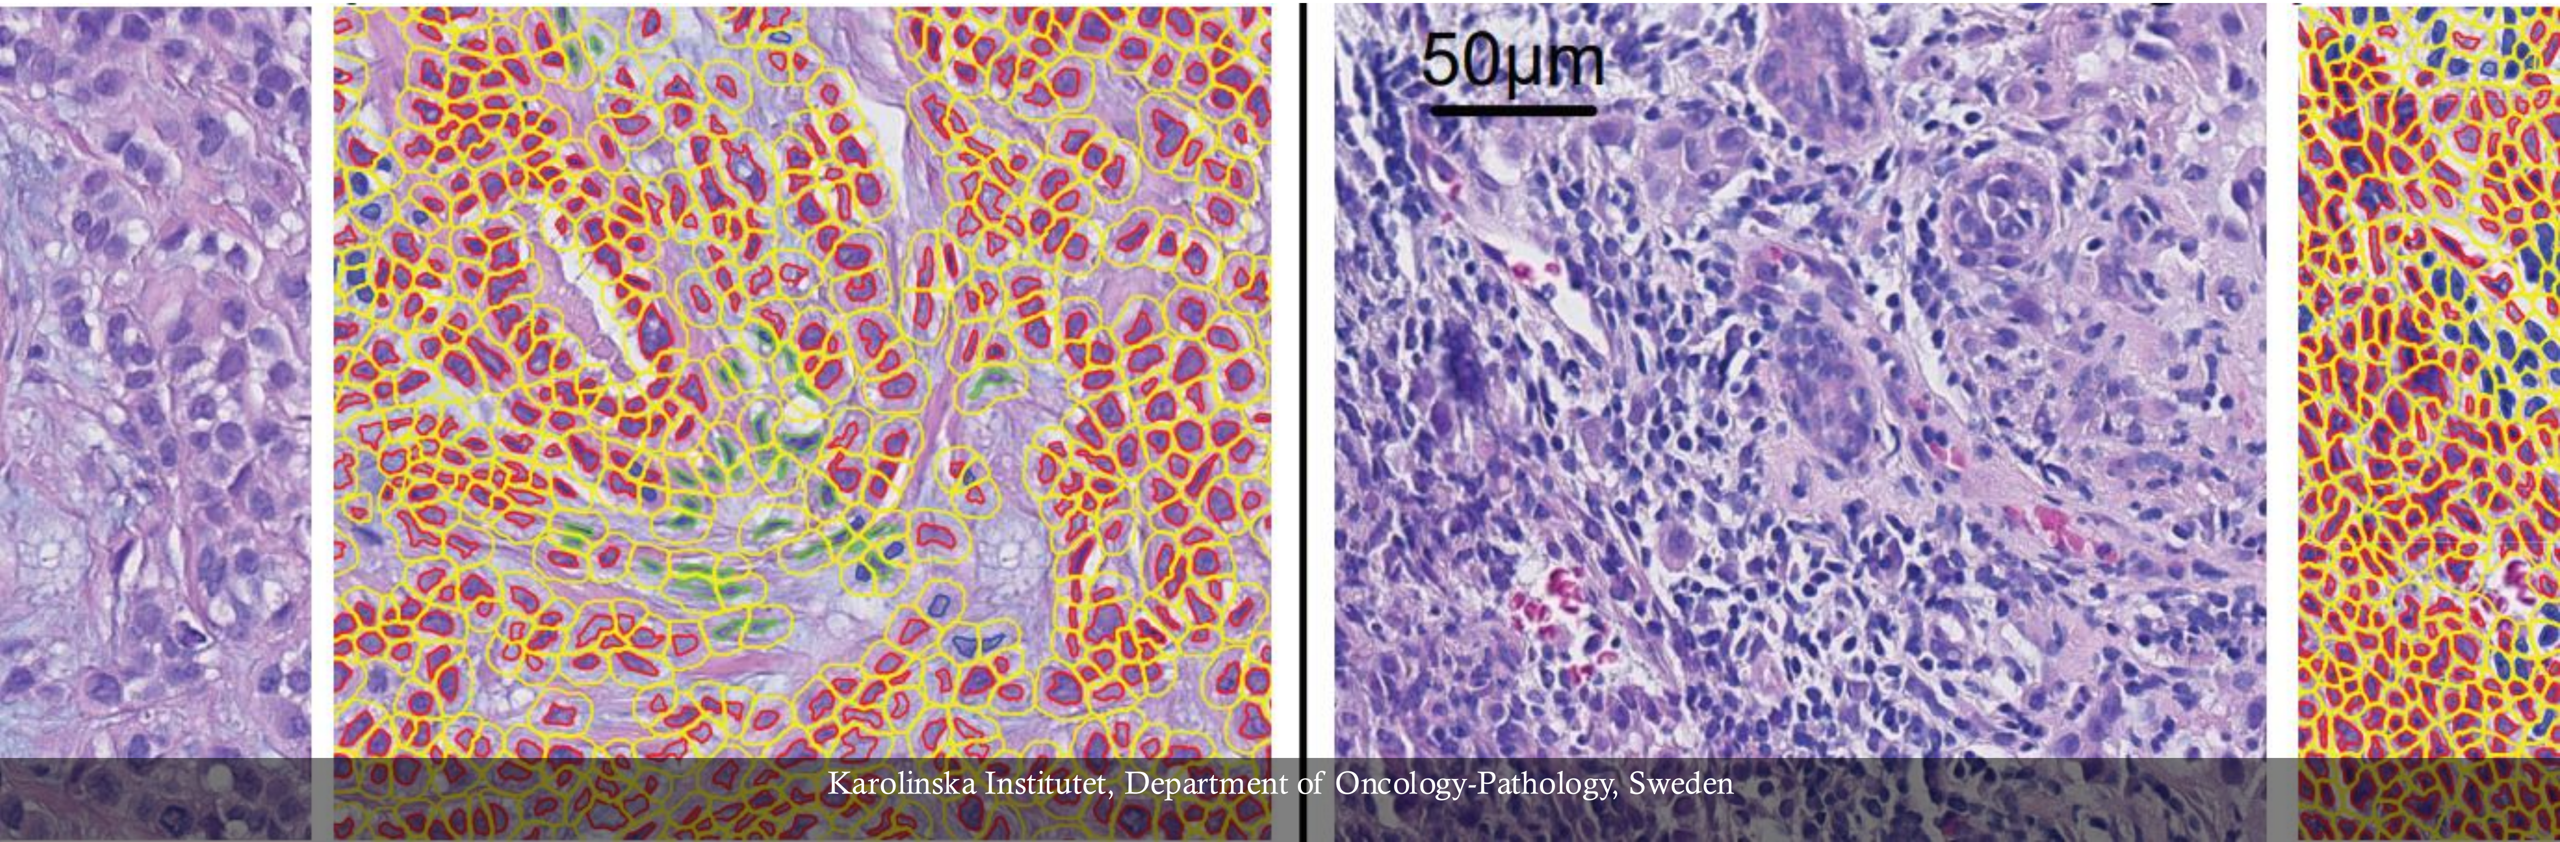

Karolinska Institutet, Department of Oncology-Pathology, Sweden

HIGH-THROUGHPUT CELLULAR PHENOTYPING

- **The Biological Need**: To understand the immune system's response to the cancer, we must accurately detect and identify millions of individual cells per slide.

- **The Deep Learning Pipeline**: Specialized architectures (like HoVer-Net) categorize every detected nucleus into distinct phenotypes: neoplastic epithelial, non-neoplastic epithelial, lymphocytes, stroma, necrotic.

- **The Output**: This process fundamentally transforms an unstructured visual image into a highly structured, queryable database of cellular identities and physical locations

HIGH-THROUGHPUT CELLULAR PHENOTYPING

Cell Colors

- CAFs: blue
- Negative Cells: white
- Cancer Cells: red
- TILs: yellow

Tissue contours

- Peritumoral: white
- Vessel: red
- Invasive cancer area: yellow

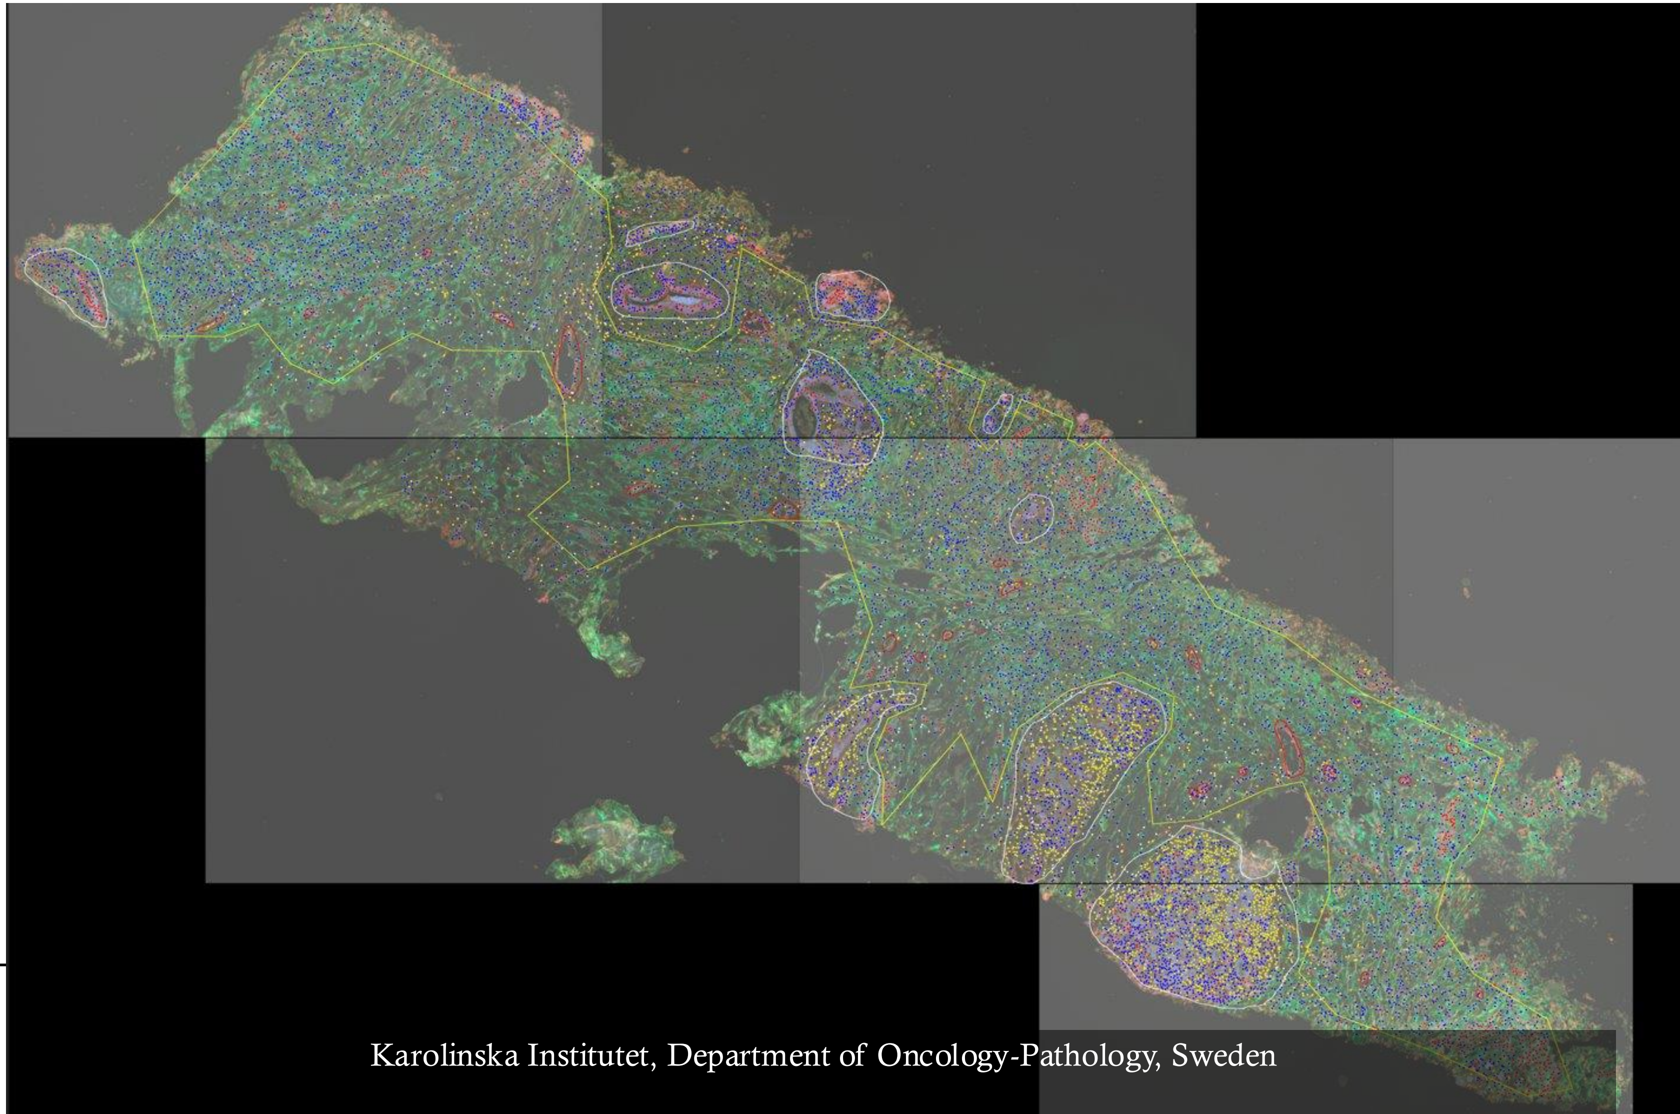

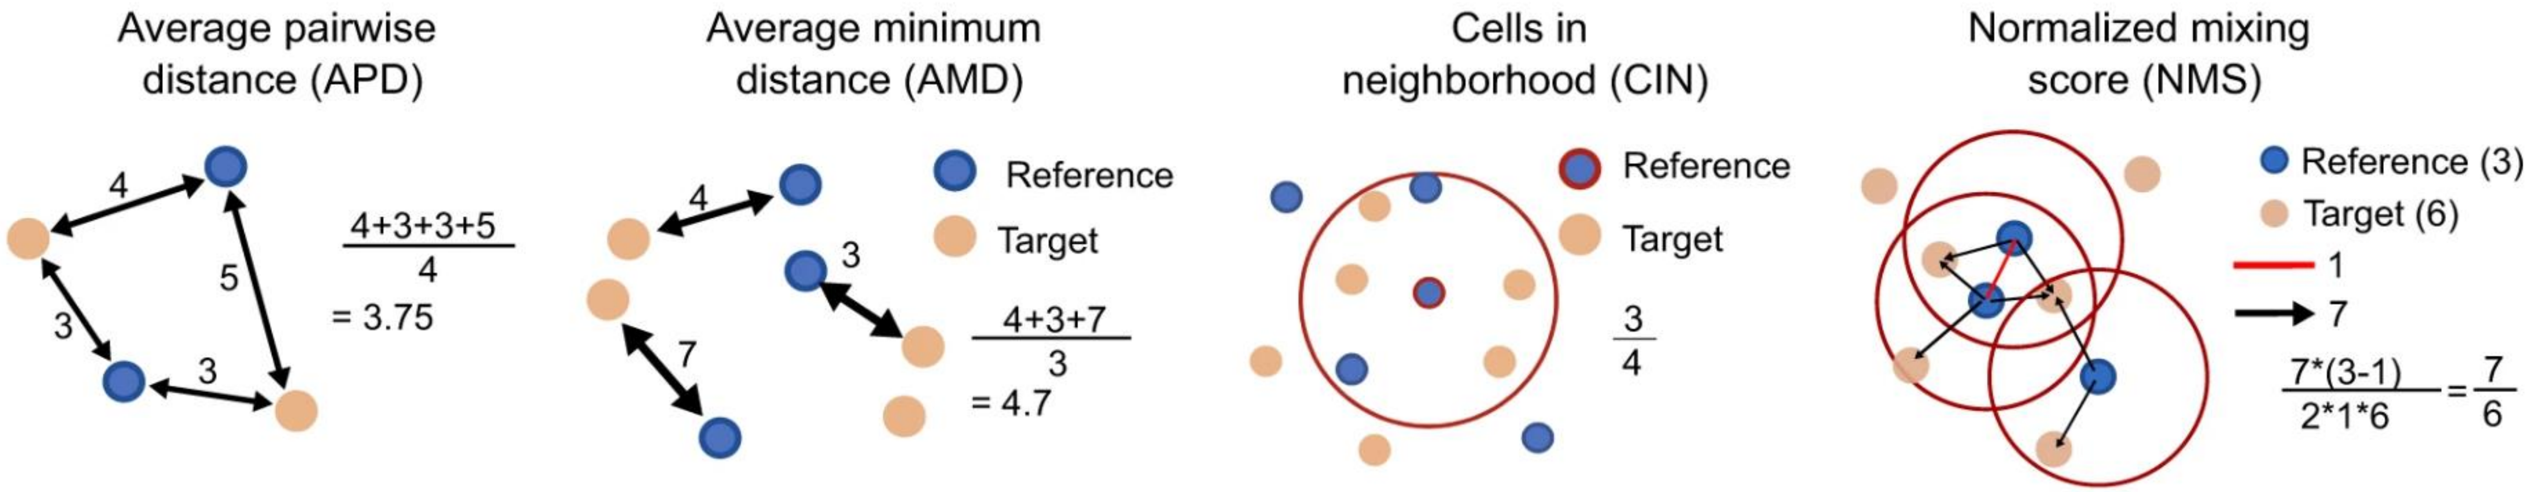

Feng, Y., Yang, T., Zhu, J. et al. Spatial analysis with SPIAT and spaSim to characterize and simulate tissue microenvironments. Nat Commun 14, 2697 (2023)

MORPHOLOGICAL, TEXTURAL AND SPATIAL FEATURE EXTRACTION

- **The Task**: Once cells were found, we had to define their shape mathematically
- **Handcrafted Approach**: Engineers wrote explicit formulas to calculate area, perimeter, and nuclear compactness
- **Limitation**: Human math cannot account for irregular, chaotic tumor morphologies, textural nucleus patterns and distances between cell types
- **Alvis HPC Execution**: Calculating thousands of features for 2 million cells takes days. We parallelized this task on Alvis to calculate features for WSIs or/and TMAs

SPATIAL GRAPH CONSTRUCTION

- **The Task**: Converting the 2 million identified cells into a mathematical network (nodes = cells, edges = physical proximity)
- AI Approach: Defining the Tumor Microenvironment (TME) (e.g. calculating exactly which immune cells are physically touching tumor cells)
- **Alvis HPC Execution**: GPUs are terrible at branching geometry logic. We drop our Slurm allocation to Alvis CPU nodes. We utilize chunked parallel processing to calculate KD-Trees, Delaunay triangulations, etc. at lightning speed across 64 cores

MORPHOLOGICAL, TEXTURAL AND SPATIAL FEATURE EXTRACTION

- Structural differences between tumor and immune local graphs

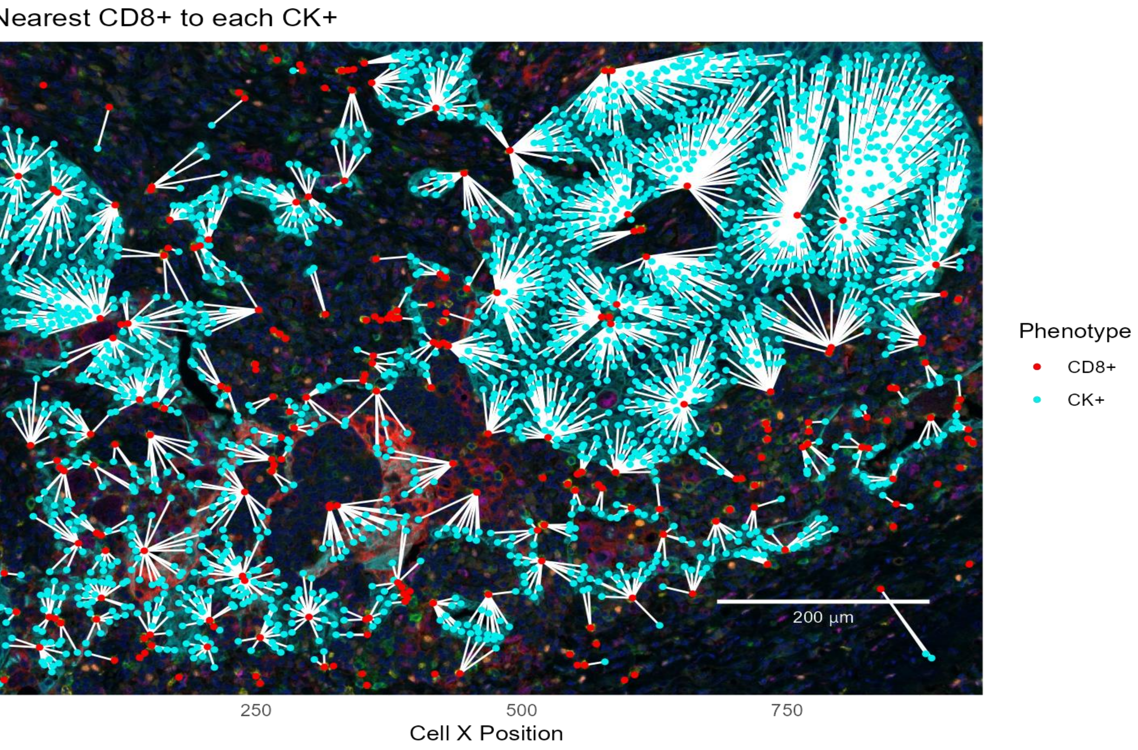

Karolinska Institutet, Department of Oncology-Pathology, Sweden

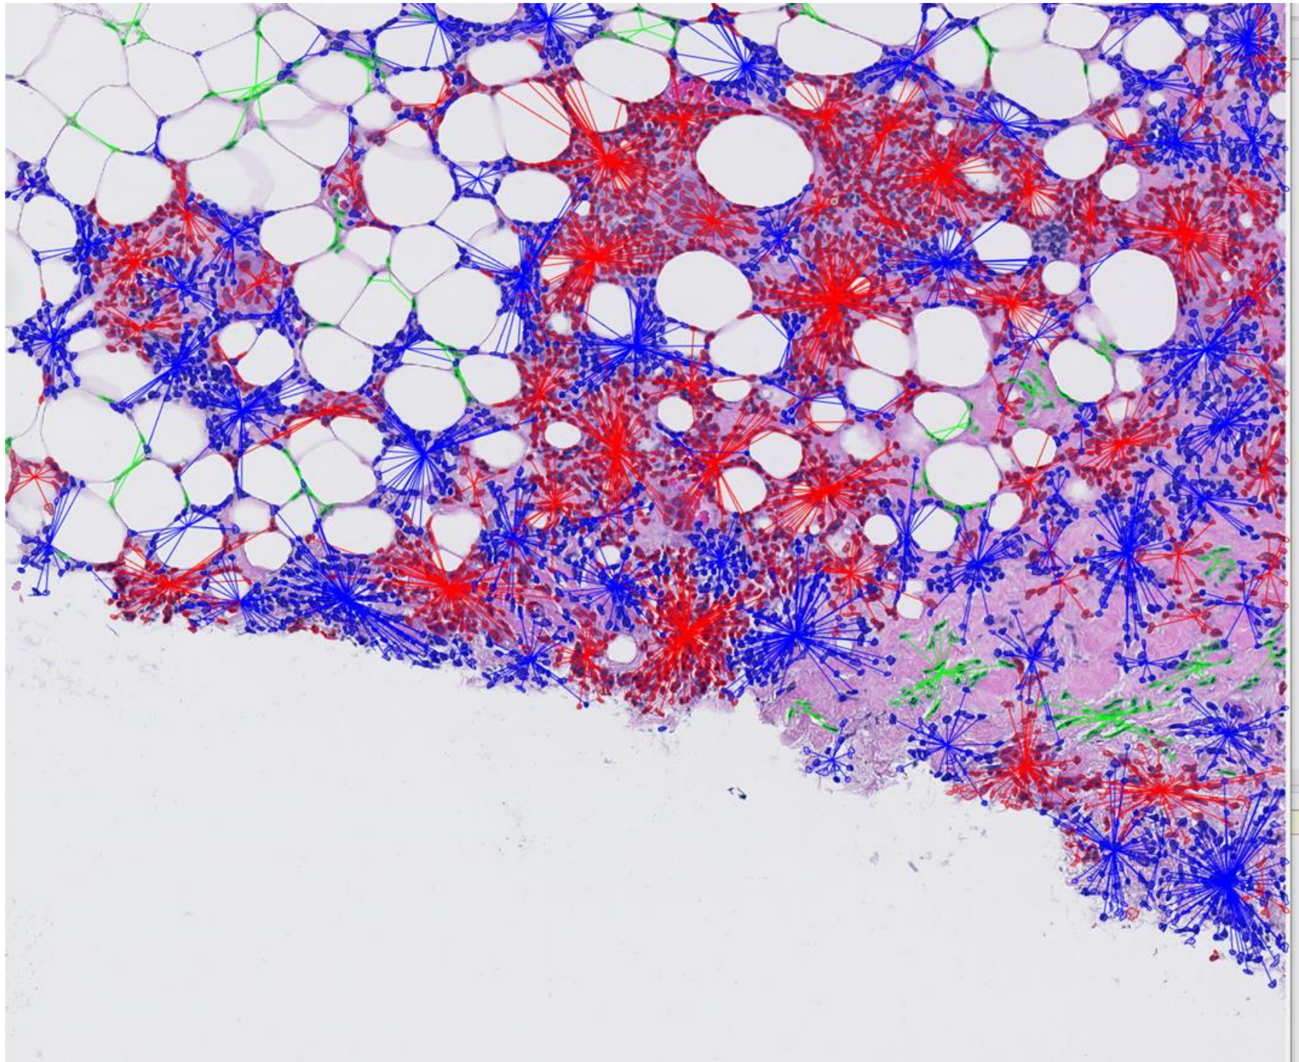

REPRESENTATION LEARNING
SHIFTING FROM CPUS TO HPC GPUS

- **The Task**: Bypassing human feature engineering

- **AI Approach**: Feeding raw pixels directly into a deep neural network. The AI learns its own optimal features

- **Advantage**: Discovers prognostic biological patterns entirely unknown to human pathologists

- **Alvis HPC Execution**: The workload shifts entirely to the alvis GPU partitions. We request A100_80GB nodes. The GPU Tensor Cores perform thousands of matrix multiplications per microsecond, rendering classical CPU extraction obsolete.

MULTIPLE INSTANCE LEARNING (MIL)

- **The Task**: A gigapixel slide cannot fit into an 80GB Alvis A100 GPU
- **Deep Learning Approach**: We break the WSI into thousands of "instances" (patches) for independent CNN processing
- **Limitation**: We lose the macro-spatial architecture of the tumor ecosystem
- **Alvis HPC Execution**: To maximize patch batch size per GPU, we program the PyTorch environment within our Apptainer using Automatic Mixed Precision (AMP). This utilizes FP16 to halve the VRAM requirement without losing diagnostic accuracy.

MULTIPLE INSTANCE LEARNING (MIL) USING ATTENTION-BASED MECHANISMS

- WSI of healthy buccal oral mucosa that has been segmented into tiles of 2000 X 2000 pixels
- Nuclei centroids were annotated and labelled as either epithelial (red), fibroblast or endothelial (blue), inflammatory (green) or lymphocytic (yellow)

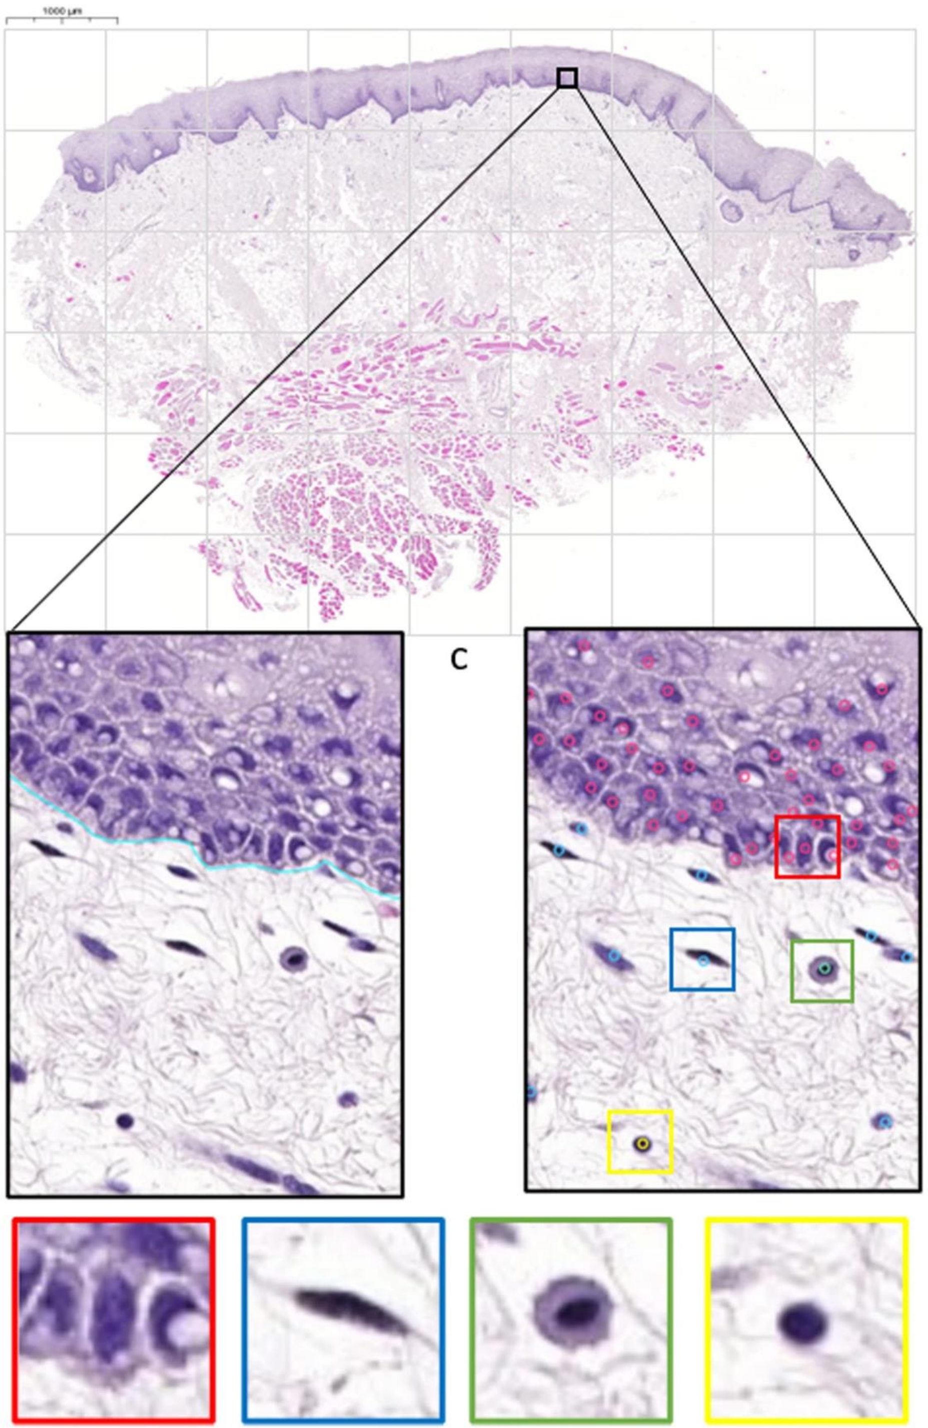

Nair, A., Arvidsson, H., Gatica V., J.E. et al. A graph neural network framework for mapping histological topology in oral mucosal tissue. BMC Bioinformatics 23, 506 (2022)

GRAPH NEURAL NETWORKS (GNNS) PREDICTING FROM THE TOPOLOGY

- **The Task**: Passing the spatial graph through an AI to predict patient survival or drug response
- **Modern Approach**: GNNs allow cells to mathematically "communicate" their states to neighboring cells, discovering complex spatial motifs.
- **Alvis HPC Execution**: Biological graphs are mostly empty space. We encode the data in Compressed Sparse Row (CSR) formats. We use NVIDIA cuSPARSE libraries on the Alvis GPUs to ensure we only execute math on actual physical cell connections.

GRAPH NEURAL NETWORKS (GNNS) PREDICTING FROM THE TOPOLOGY

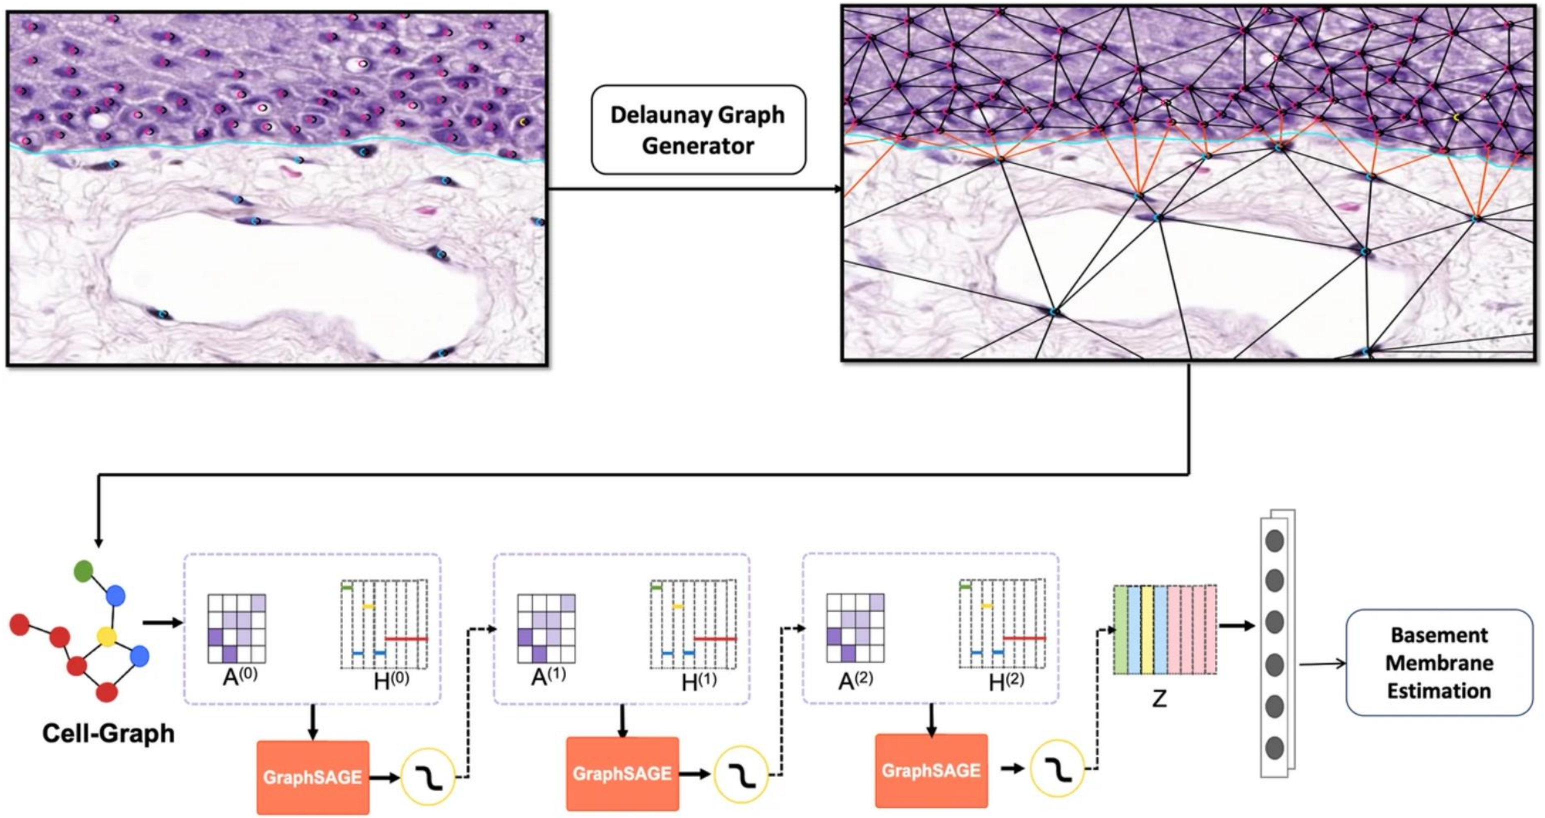

GRAPH NEURAL NETWORKS (GNNS)
PREDICTING FROM THE TOPOLOGY

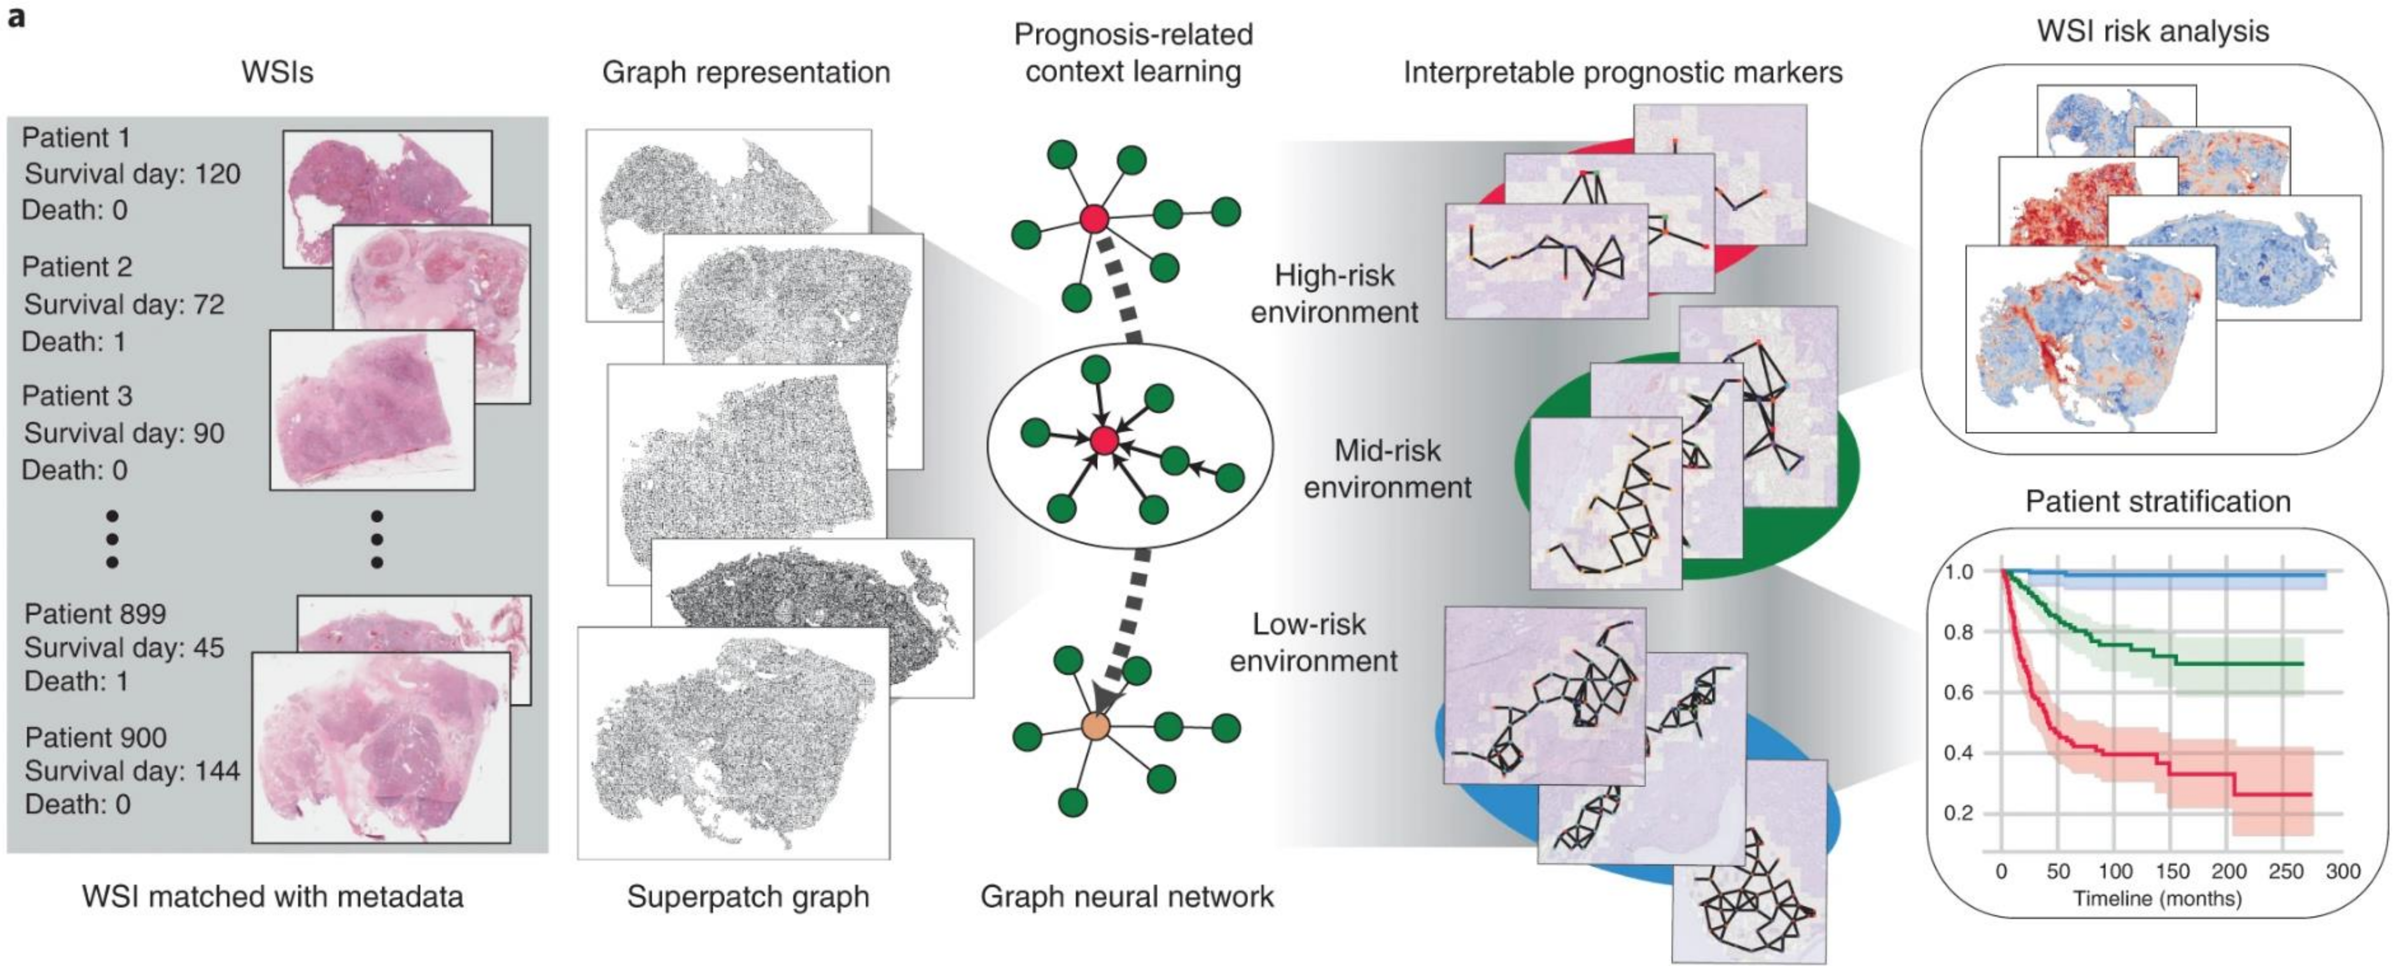

Lee, Y., Park, J.H., Oh, S. et al. Derivation of prognostic contextual histopathological features from whole-slide images of tumours via graph deep learning. Nat. Biomed. Eng 6, 1452–1466 (2022)

INTEGRATIVE AI & MULTIMODAL FUSION FUSING H&E INTELLIGENCE FOR CLINICAL ENDPOINTS

- **The Clinical Goal**: We have successfully extracted millions of morphological and spatial features from the slides.
- We now fuse this massive feature set into a single Integrative AI architecture.
- **Target 1 (pCR Prediction)**: Predicting if a tumor will completely disappear after pre-surgical chemotherapy.
- **Target 2 (Treatment Response)**: Identifying exactly which patients will respond to targeted immunotherapies.
- **Target 3 (Prognostication)**: Calculating highly accurate, long-term patient survival risk scores.

- **Alvis HPC Execution**: Training these integrative models requires optimizing complex, custom loss functions (like Deep Cox Proportional Hazards). We synchronize these calculations across multiple GPU nodes to maintain mathematical accuracy across the entire patient batch.

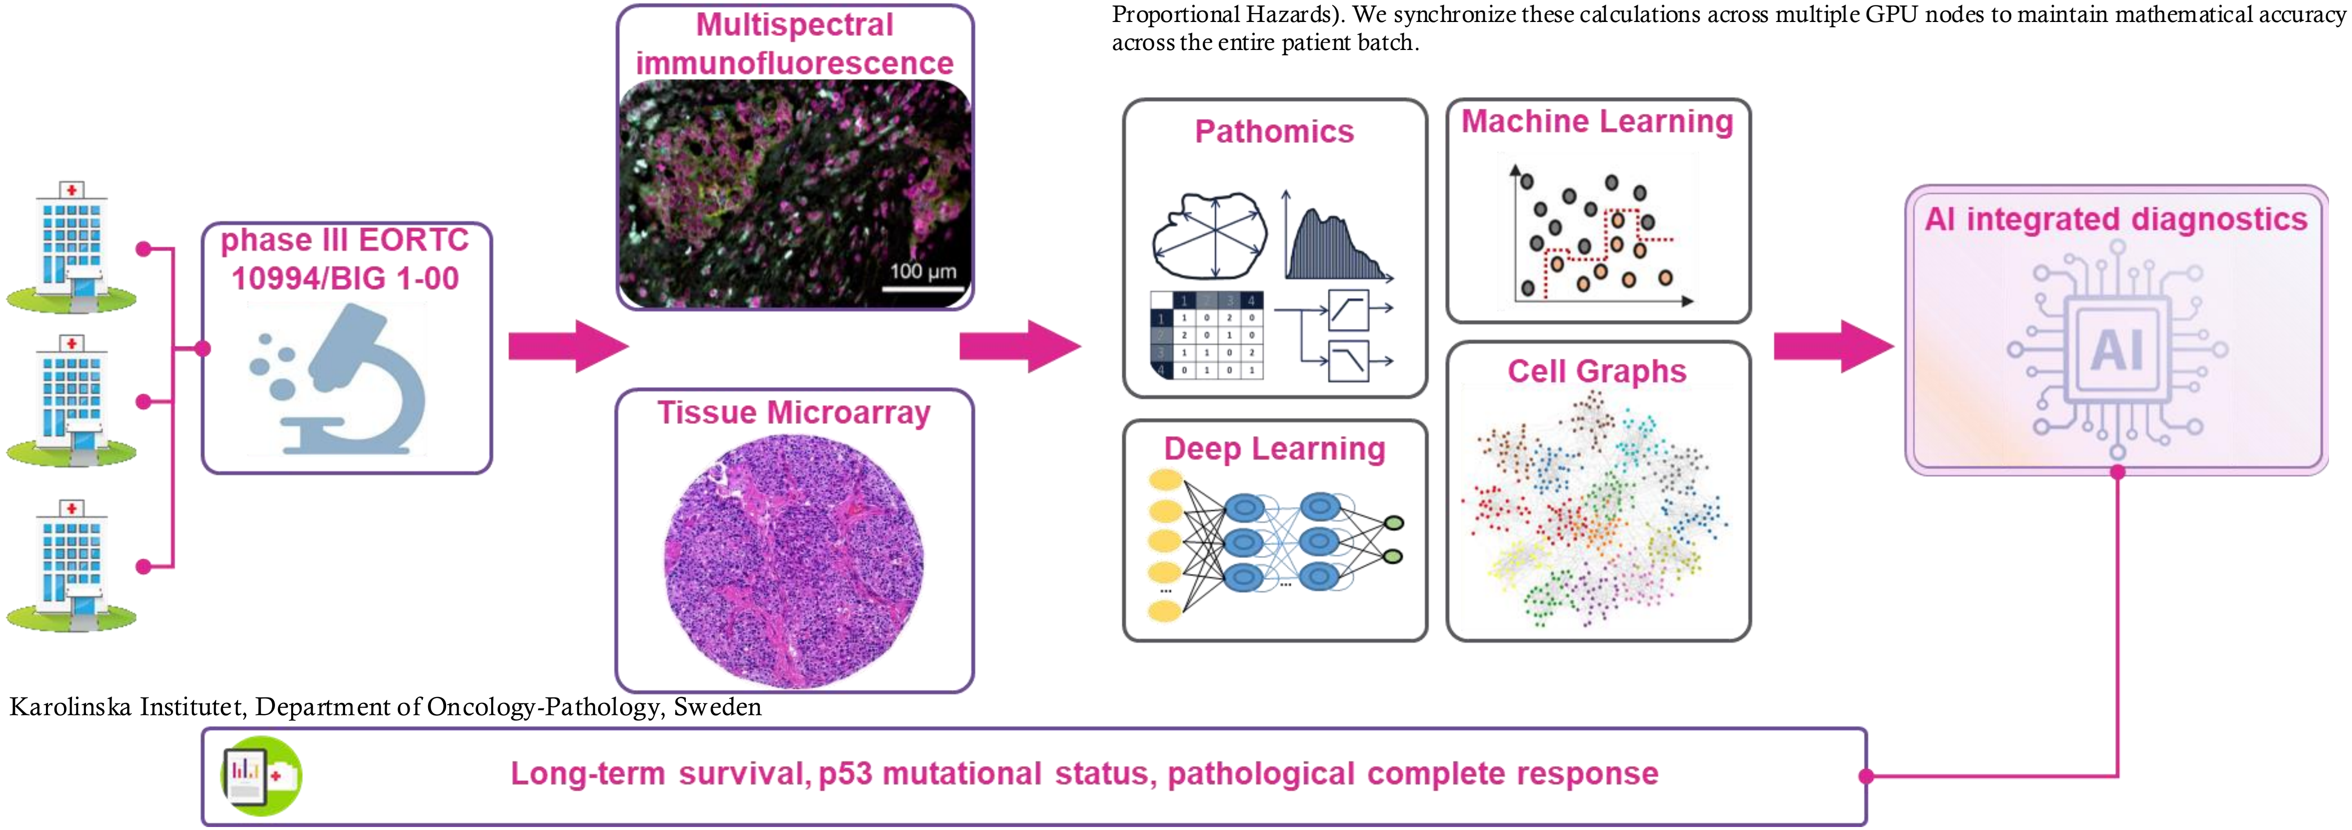

Karolinska Institutet, Department of Oncology-Pathology, Sweden

A NEW ERA: FOUNDATION AI

STATE-OF-THE-ART FOUNDATION MODELS

A foundation model that looks across the entire slide at multiple magnifications to make slide-level assessments

A whole-slide pathology foundation model pre-trained on 171,189 whole slides from Providence, a large US health network comprising 28 cancer centers

Predicts **microsatellite instability (MSI)** status (high number of mutations), **tumor mutation burden** (TMB) (number of mutations found in the DNA of cancer cells), and **homologous recombination deficiency** (HRD) (phenotype that is characterized by the inability of a cell to effectively repair DNA double-strand breaks) without additional testing

Validated using a prospective study across 12 hospitals showing 91% sensitivity

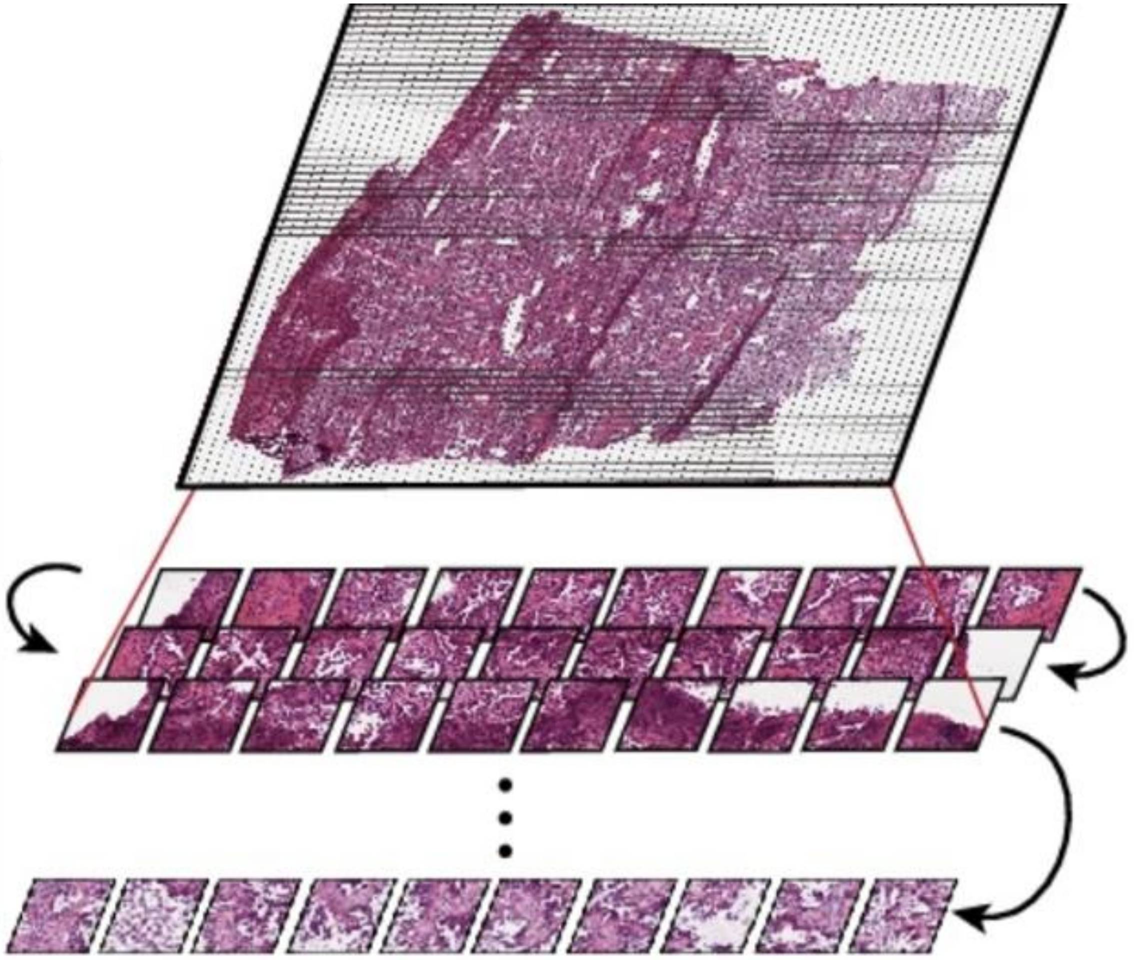

Xu, H., Usuyama, N., Bagga, J. et al. A whole-slide foundation model for digital pathology from real-world data. Nature 630, 181–188 (2024) *Xu, H., Usuyama, N., Bagga, J. et al. A whole-slide foundation model for digital pathology from real-world data. Nature 630, 181–188 (2024)*

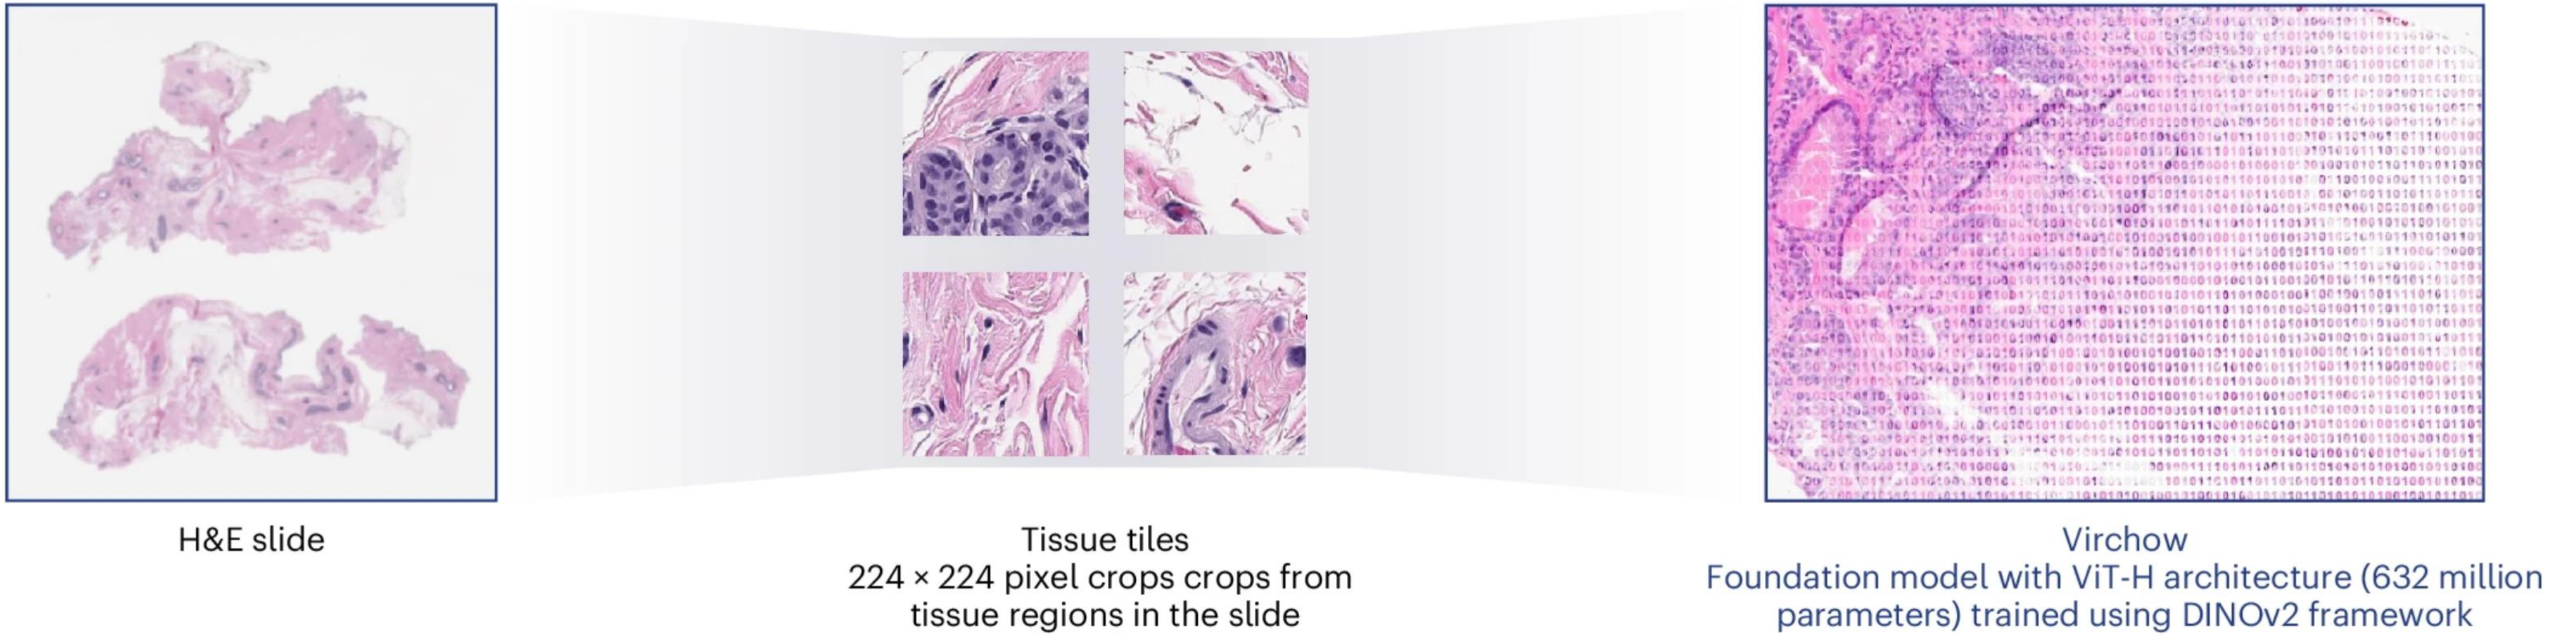

STATE-OF-THE-ART FOUNDATION MODELS - VIRCHOW

- Trained on 1.5M whole slide images (WSIs) from 100K+ patients
- Pan-cancer detection, achieving 95% AUC across nine common and seven rare cancers
- Achieves state-of-the-art performance across diverse pathology tasks including cancer classification, tumor subtyping, biomarker prediction (MSI, TMB, gene mutations), and survival outcome prediction

Vorontsov, E., Bozkurt, A., Casson, A. et al. A foundation model for clinical-grade computational pathology and rare cancers detection. Nat Med 30, 2924–2935 (2024)

STATE-OF-THE-ART FOUNDATION MODELS - CTRANSPATH

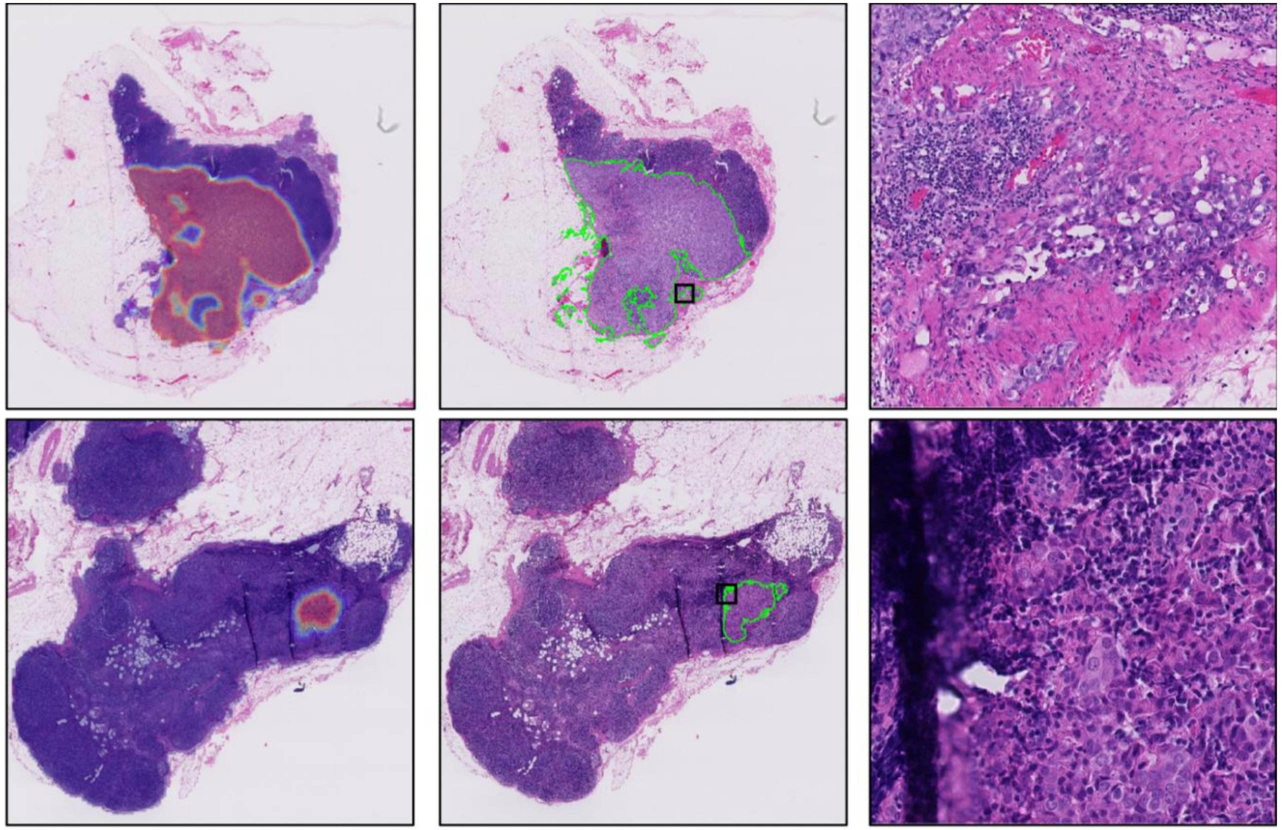

Left: automatically computed attention heatmaps. Warmer colors of the attention map indicate higher estimated probabilities of being tumorous tissue. **Middle-Right**: Ground truth with green lines representing the ground truth of the cancer metastasis, while dark rectangles indicate local ROIs highlighting the boundaries between metastatic and normal tissues

- Trained on 15+ million tissue image patches from >30,000 patients, focusing on 9 most common cancers (breast, lung, colorectal, prostate)
- Cancer detection & grading: Identifies malignant regions, Gleason scoring, nuclear grading
- Subtype classification: Distinguishes histological subtypes (adenocarcinoma vs. squamous, etc.)
- Biomarker prediction: Estimates MSI status, HER2 expression, PD-L1 scoring from H&E alone
- Prognostic assessment: Risk stratification based on morphological patterns
- Treatment response: Predicts likely response to immunotherapy or targeted therapy

STATE-OF-THE-ART FOUNDATION MODELS - UNI

- A single unified model that trained to understand both the image and the radiology reports at the same time
- Works across all organs and cancer types without organ-specific training
- A pre-trained model for pathology, using more than 100 million images from over 100,000 diagnostic H&E-stained WSIs across 20 major tissue types
- The Breakthrough: It doesn't just see patterns; it understands the clinical meaning and context behind them, creating a bridge between visual evidence and diagnostic language
- Pan-cancer diagnosis, Rare tumor recognition, Molecular features from morphology, Estimates immunotherapy response markers from routine H&E, Identifies micro-metastases and isolated tumor cells

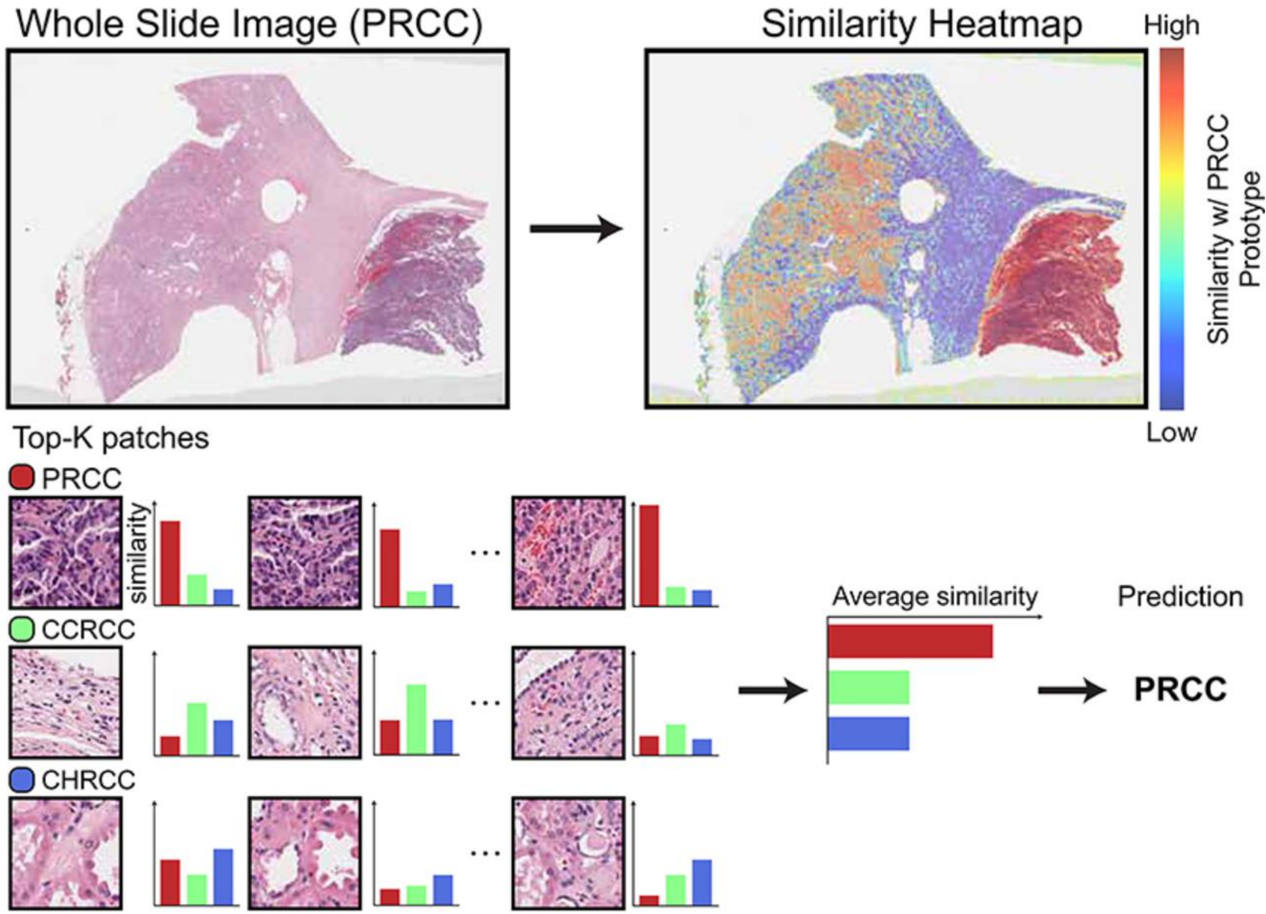

Chen RJ, Ding T, Lu MY, Williamson DFK, Jaume G, Song AH, Chen B, Zhang A, Shao D, Shaban M, Williams M, Oldenburg L, Weishaupt LL, Wang JJ, Vaidya A, Le LP, Gerber G, Sahai S, Williams W, Mahmood F. Towards a general-purpose foundation model for computational pathology. Nat Med. 2024

STATE-OF-THE-ART FOUNDATION MODELS - TITAN

- A multimodal whole-slide foundation model pretrained using 335,645 whole-slide images with corresponding pathology reports across 20 organ types

- Encodes entire WSIs into slide-level general-purpose feature representations

- TITAN produces general-purpose slide representations that can readily be applied to slide-level tasks such as cancer subtyping, biomarker prediction, outcome prognosis and slide retrieval tasks

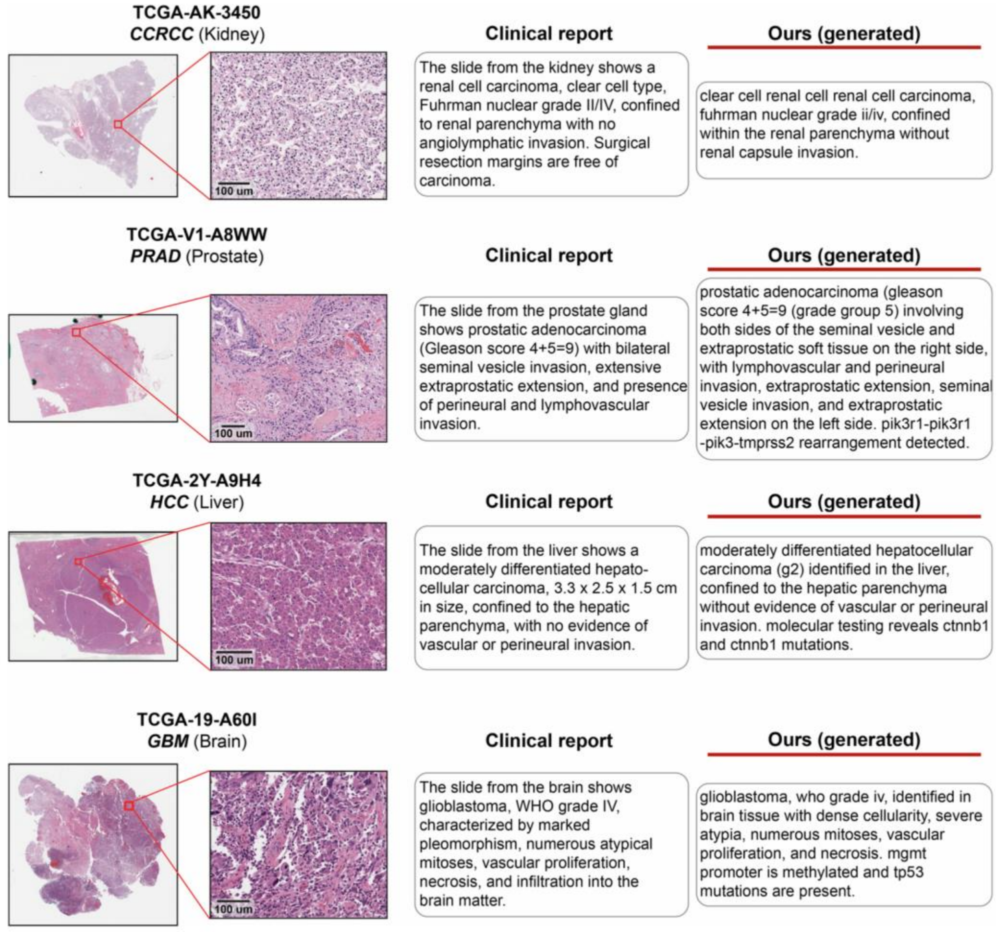

Ding, T., Wagner, S.J., Song, A.H. et al. A multimodal whole-slide foundation model for pathology. Nat Med 31, 3749–3761 (2025)

BUILDING THE HELLENIC FOUNDATION MODEL A NATIONAL STRATEGIC INITIATIVE

- We propose utilizing PHAROS to build a national WSI Foundation Model
- We will ingest federated diagnostic data from all major Greek research hospitals
- This ensures the AI accurately represents the Mediterranean genetic profile
- Global models trained strictly on North American data often underperform here
- This creates highly valuable, sovereign intellectual property for Greece

INDUSTRIAL SYNERGIES (PHARMA ROI) COMPUTE-AS-A-SERVICE FOR DRUG DISCOVERY

- The PHAROS cluster has massive commercial value for the biotech sector
- We can offer High-Throughput Computational Biomarker discovery
- AI can create synthetic patient cohorts for in-silico clinical trial simulation
- Spatial GNNs can predict patient response to expensive immunotherapies
- This infrastructure will drastically improve clinical trial success rates in Greece

CONCLUSION & THE FUTURE
COMPUTE AS THE MODERN MICROSCOPE

- Gigapixel AI is fundamentally an I/O, Memory, and Parallel Compute challenge
- The evolution of AI was entirely driven by underlying hardware capabilities
- Foundation Models and Spatial Graphs now see patterns beyond human vision
- Mastering HPC systems engineering is the only way to deploy these tools clinically
- ***PHAROS provides the sovereign compute backbone necessary to revolutionize Greek healthcare***

THANK YOU FOR YOUR ATTENTION