Towards multimodal and context-aware representations of the tumor microenvironment

Marianna Rapsomaniki
Assistant Professor
Biomedical Data Science Center

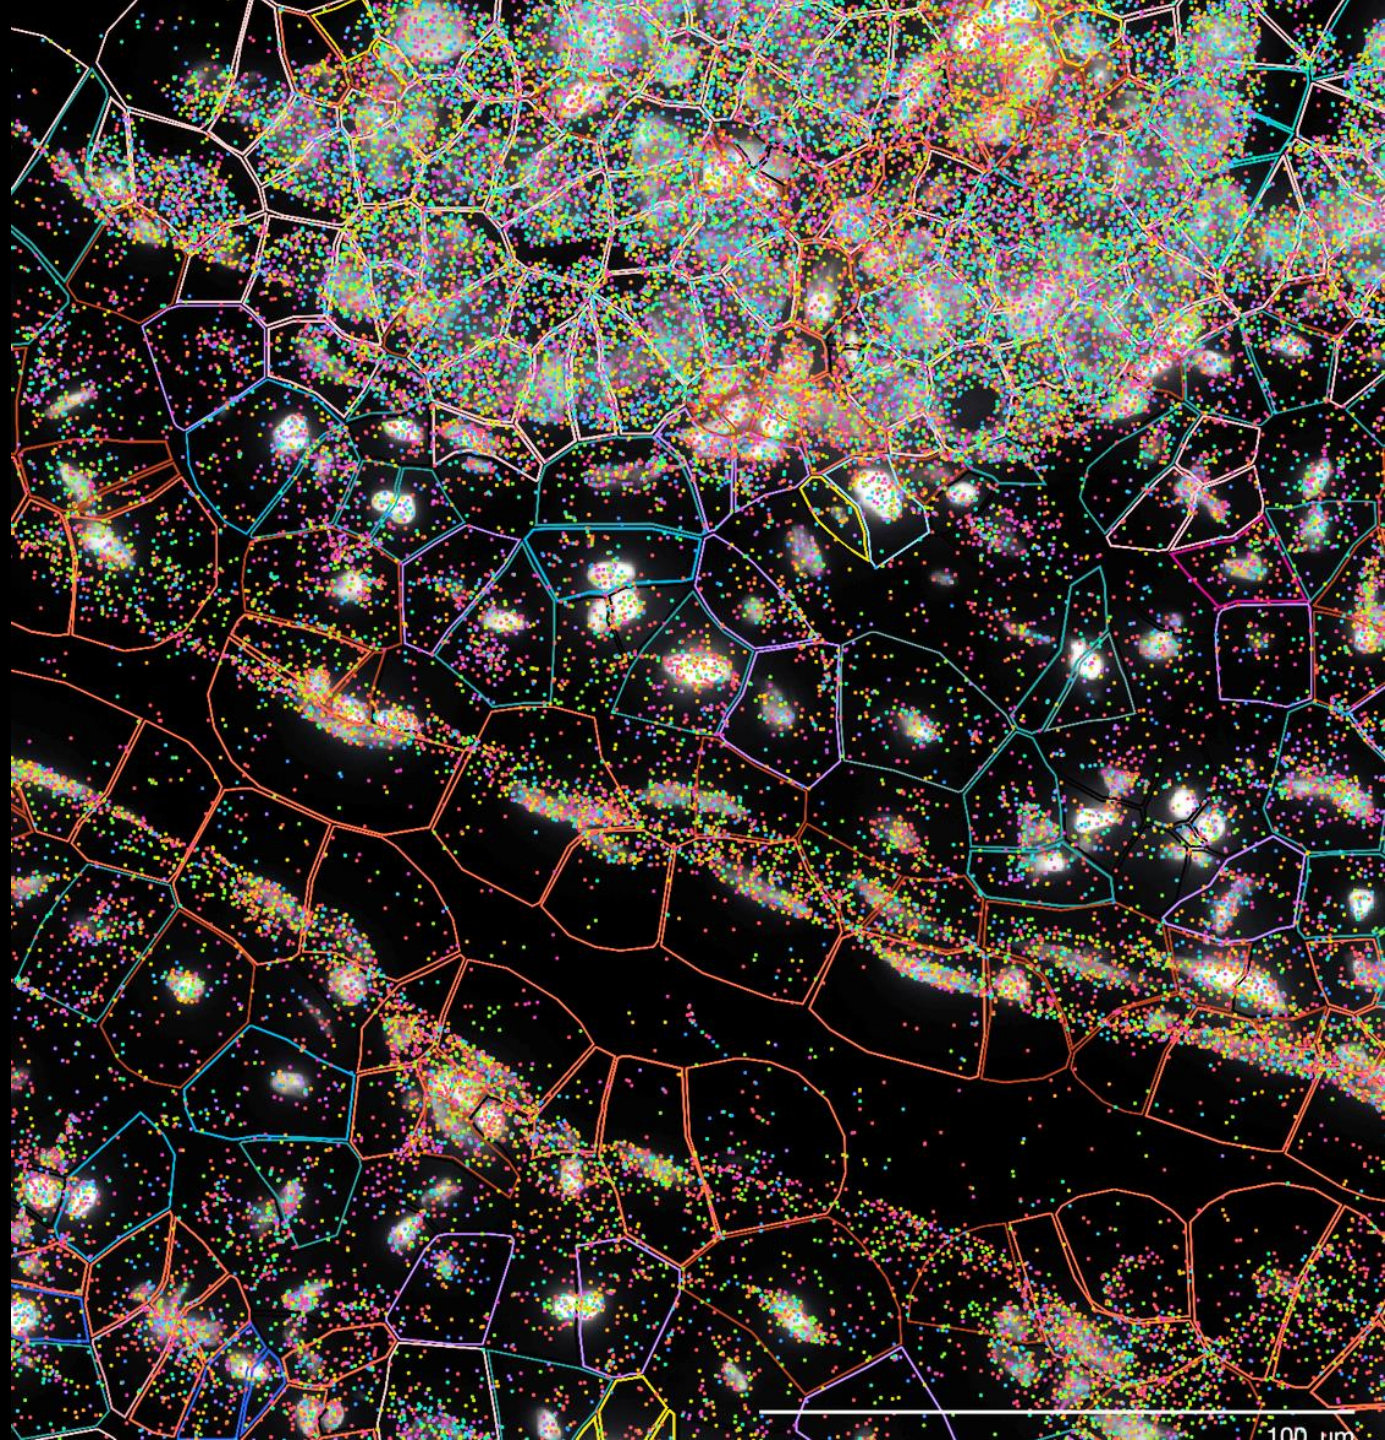

Lyon AI/ML sc-spatial omics workshop
Oct. 16th, 2025

From biological stochasticity to tumor heterogeneity

Intra-cell

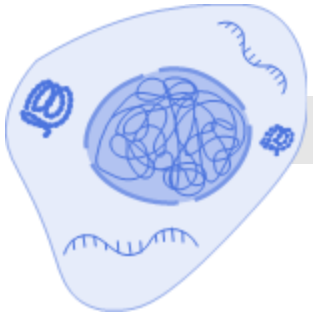

Intra-tumor

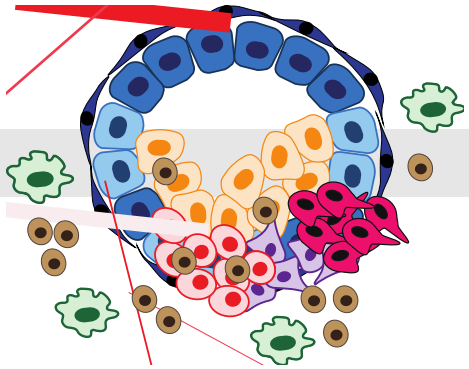

Intra-patient

Inter-patient

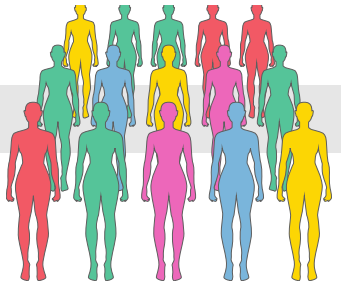

From single-cell biology…

Single-cell omics
Nature Method of the year 2013

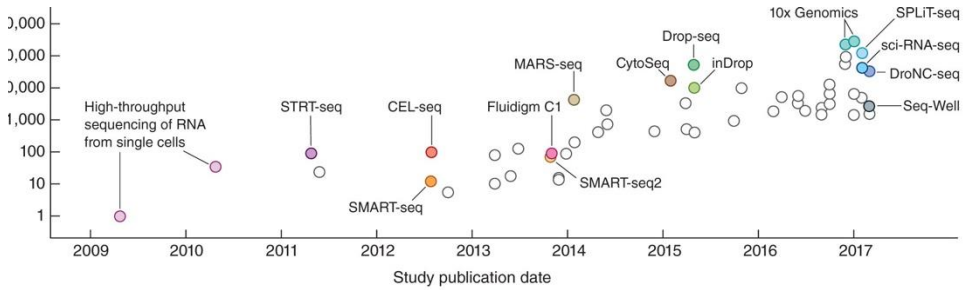

Svensson et al., Nat Protocols, (2018)

… to multi-modal single-cell biology

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

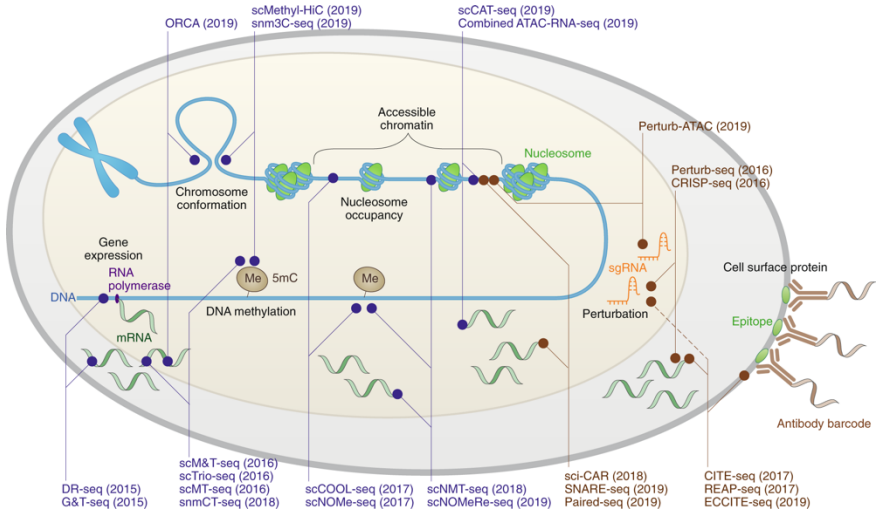

Zhu et al., *Nat Methods*, (2020)

… to spatial single-cell biology

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

10X Xenium - Breast Cancer

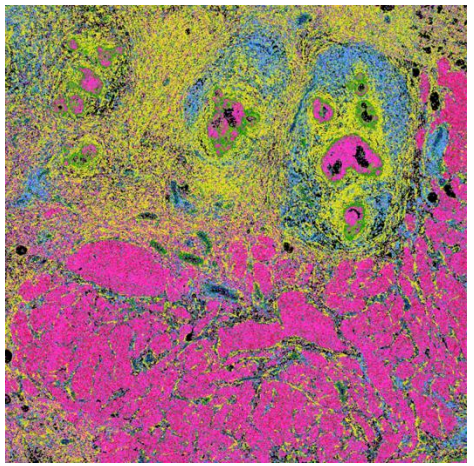

Nanostring CosMx – Human Lung

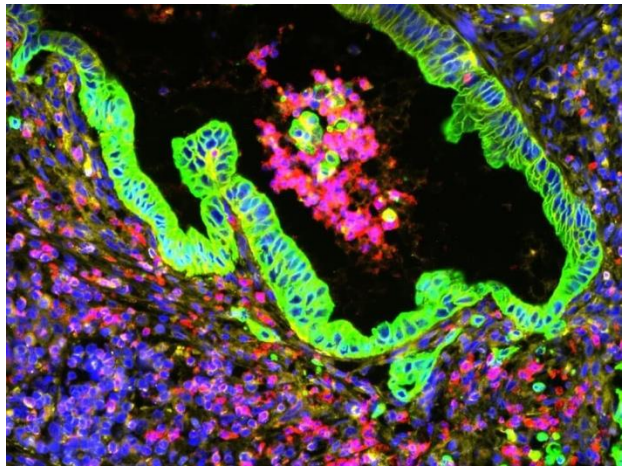

Imaging Mass Cytometry - Human Placenta

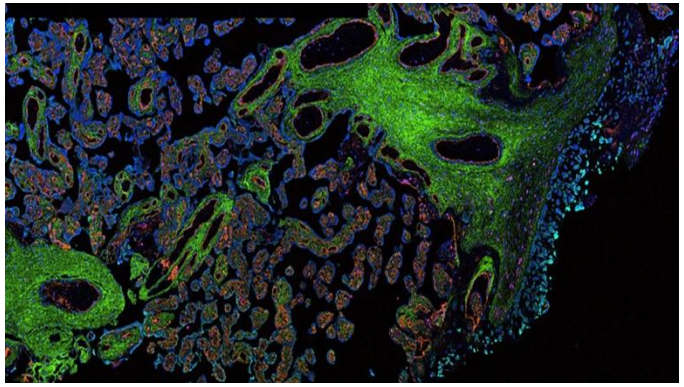

CODEX – Lymph node

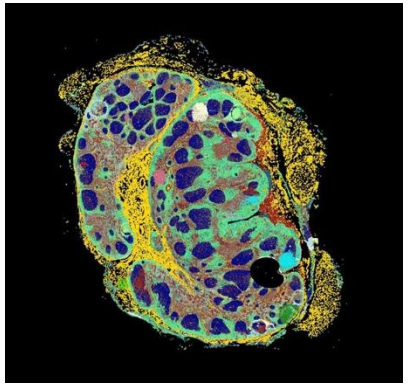

Images from https://hubmapconsortium.org, 10Xgenomics.com and Nanostring.com

… and perturbation biology

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

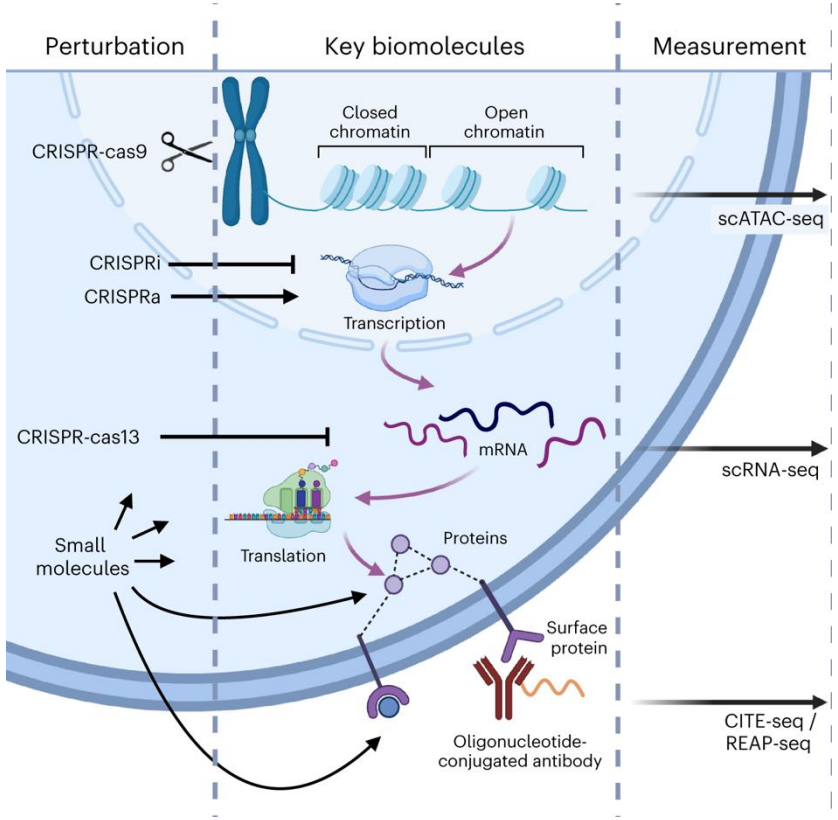

Peidl et al., Nature Methods, 2024

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Whole tumors in health and disease, over time or under treatment
down to the last single-cell, for millions of cells

… to new paradigms in AI

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Variational Autoencoders & Generative Adversarial Networks
2013 - 2014

Whole tumors in health and disease, over time or under treatment down to the last single-cell, for millions of cells

… to new paradigms in AI

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Variational Autoencoders & Generative Adversarial Networks
2013 - 2014

Transformers and Foundation models
2017 - 2019

Whole tumors in health and disease, over time or under treatment down to the last single-cell, for millions of cells

… to new paradigms in AI

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Variational Autoencoders & Generative Adversarial Networks
2013 - 2014

Transformers and Foundation models
2017 - 2019

Multimodal models (DALL-E, CLIP)
2021

Whole tumors in health and disease, over time or under treatment down to the last single-cell, for millions of cells

… to new paradigms in AI

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Variational Autoencoders & Generative Adversarial Networks
2013 - 2014

Transformers and Foundation models
2017 - 2019

Multimodal models (DALL-E, CLIP)
2021

Agentic AI & Reasoning
Today

Whole tumors in health and disease, over time or under treatment down to the last single-cell, for millions of cells

… to new paradigms in AI

Single-cell omics
Nature Method of the year 2013

Single-cell multiomics
Nature Method of the year 2019

Spatial (single-cell) (multi) omics
Nature Method of the year 2020 & 2024

Drug or Genetic Perturbations

Variational Autoencoders & Generative Adversarial Networks
2013 - 2014

Transformers and Foundation models
2017 - 2019

Multimodal models (DALL-E, CLIP)
2021

Agentic AI & Reasoning
Today

Whole tumors in health and disease, over time or under treatment down to the last single-cell, for millions of cells

+

New modeling paradigms able to capture complex multimodal, spatiotemporal data, interact and reason

From molecules to whole tissues

Prediction of protein structure
(Jumper *et al.*, *Nature,* 2021)

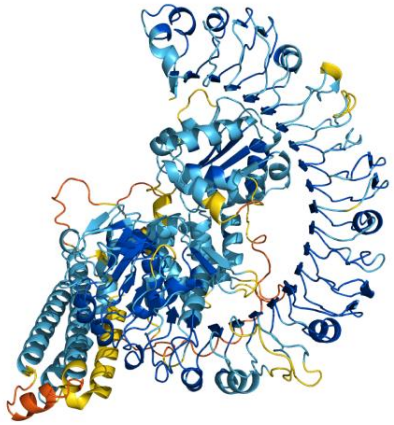

Structure of protein complexes
(Humphreys *et al.*, *Science,* 2021)

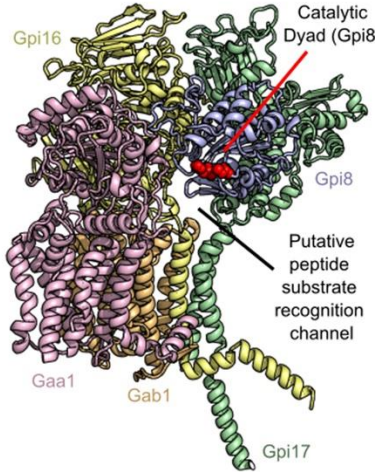

De novo protein or enzyme design
(Watson *et al.*, *Nature,* 2023)

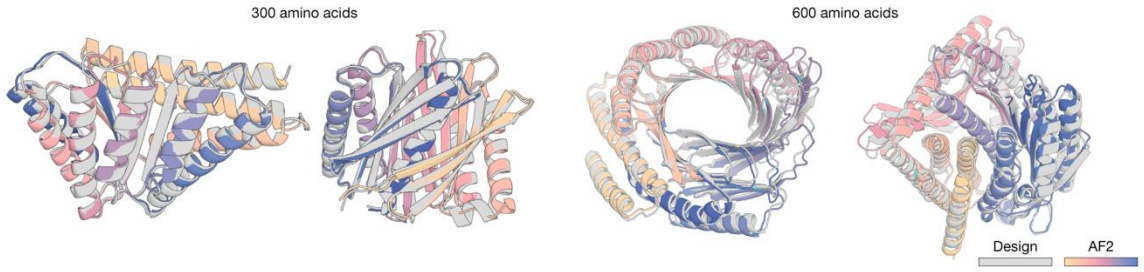

A Multimodal Generative AI Copilot for Human Pathology
(Lu *et al.*, *Nature,* 2024)

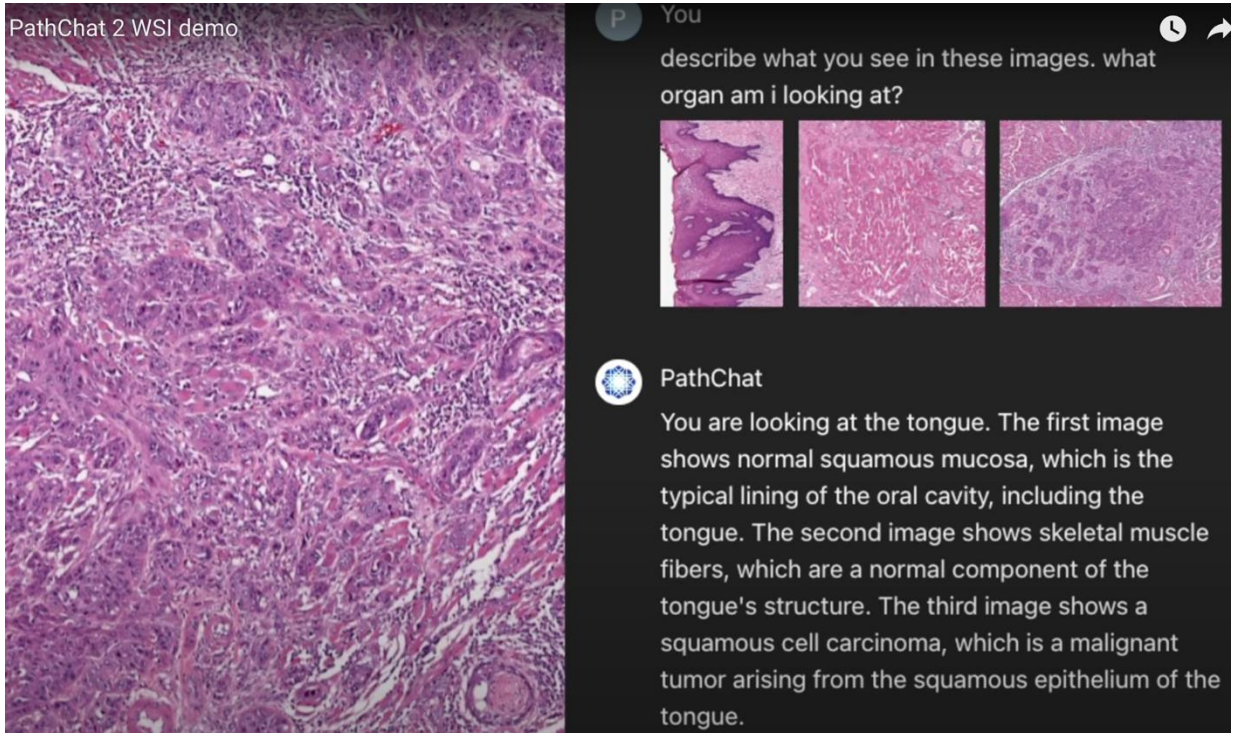

Towards a general-purpose foundation model for computational pathology
(Chen *et al.*, *Nature Medicine,* 2024)

A visual-language foundation model for computational pathology (Lu *et al.*, *Nature Medicine,* 2024)

A reality check

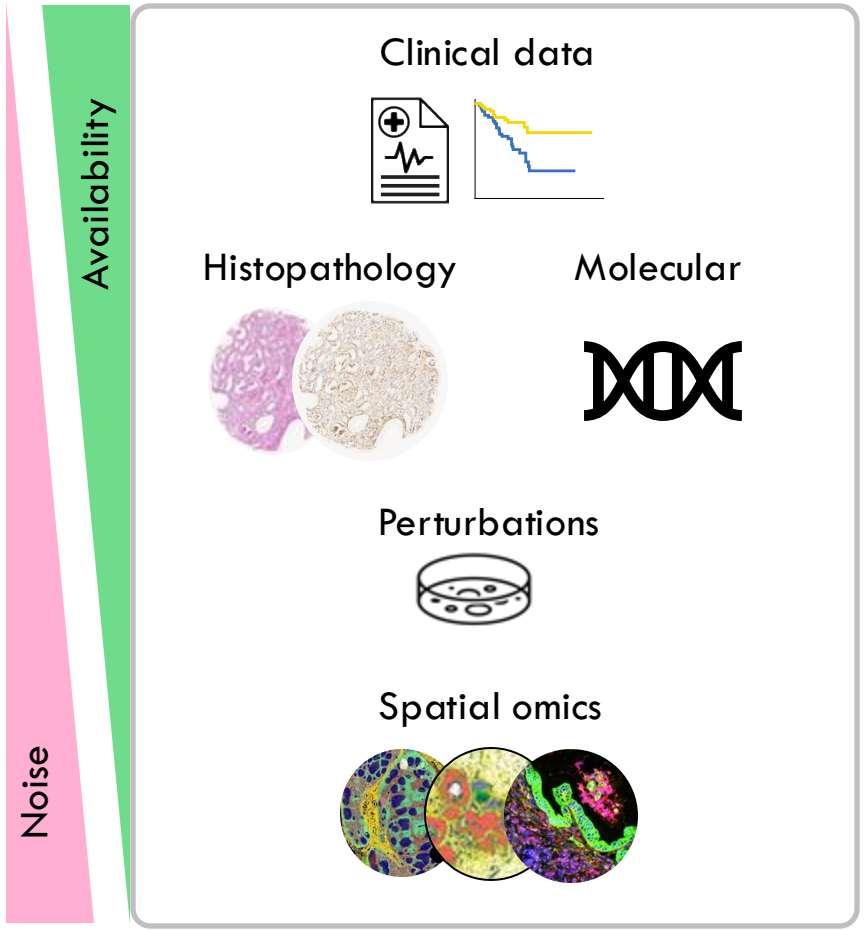

Modeling limitations

Models are not directly transferrable

Model selection/tuning

Validation and assessment

Generalizability

Reliable and Safe

Explainability

Avoiding Bias

Privacy and Security

Represent

Histopathology

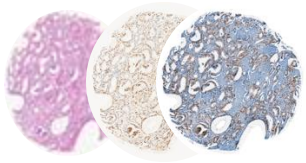

Spatial omics

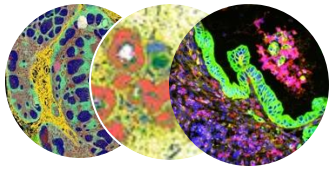

Clinical data

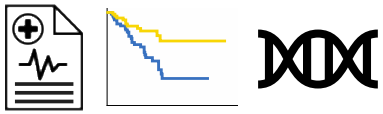

Interpret

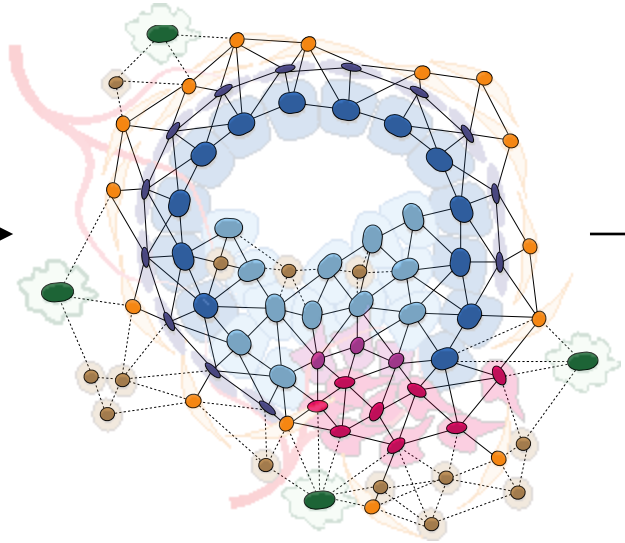

Perturb

Perturbations

Biological discovery

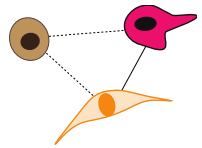

Treatment selection

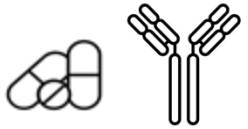

Represent

Histopathology

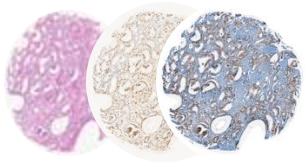

Spatial omics

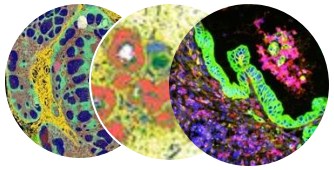

Clinical data

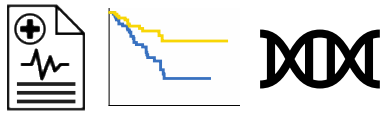

Interpret

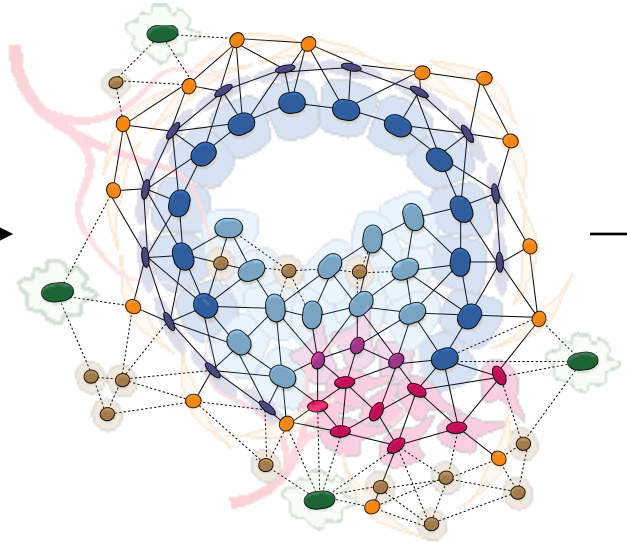

Perturb

Perturbations

Biological discovery

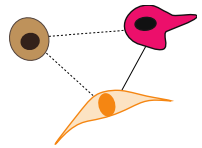

Treatment selection

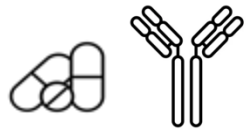

How can we learn multimodal **representations** from spatial omics?

How can we **interpret** these representations of the TME?

How can we predict the effects of drug **perturbations**?

Represent

Histopathology

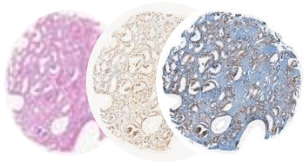

Spatial omics

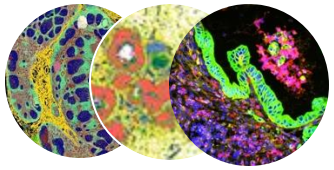

Clinical data

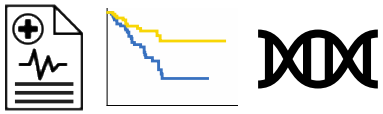

Interpret

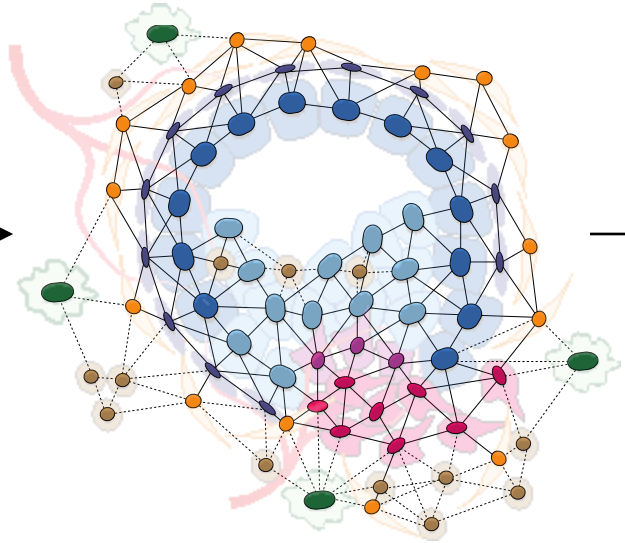

Perturb

Perturbations

Biological discovery

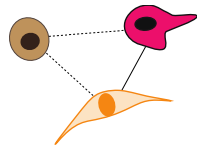

Treatment selection

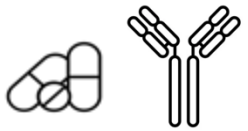

How can we learn multimodal **representations** from spatial omics?

How can we **interpret** these representations of the TME?

How can we predict the effects of drug **perturbations**?

Tumors are complex ecosystems!

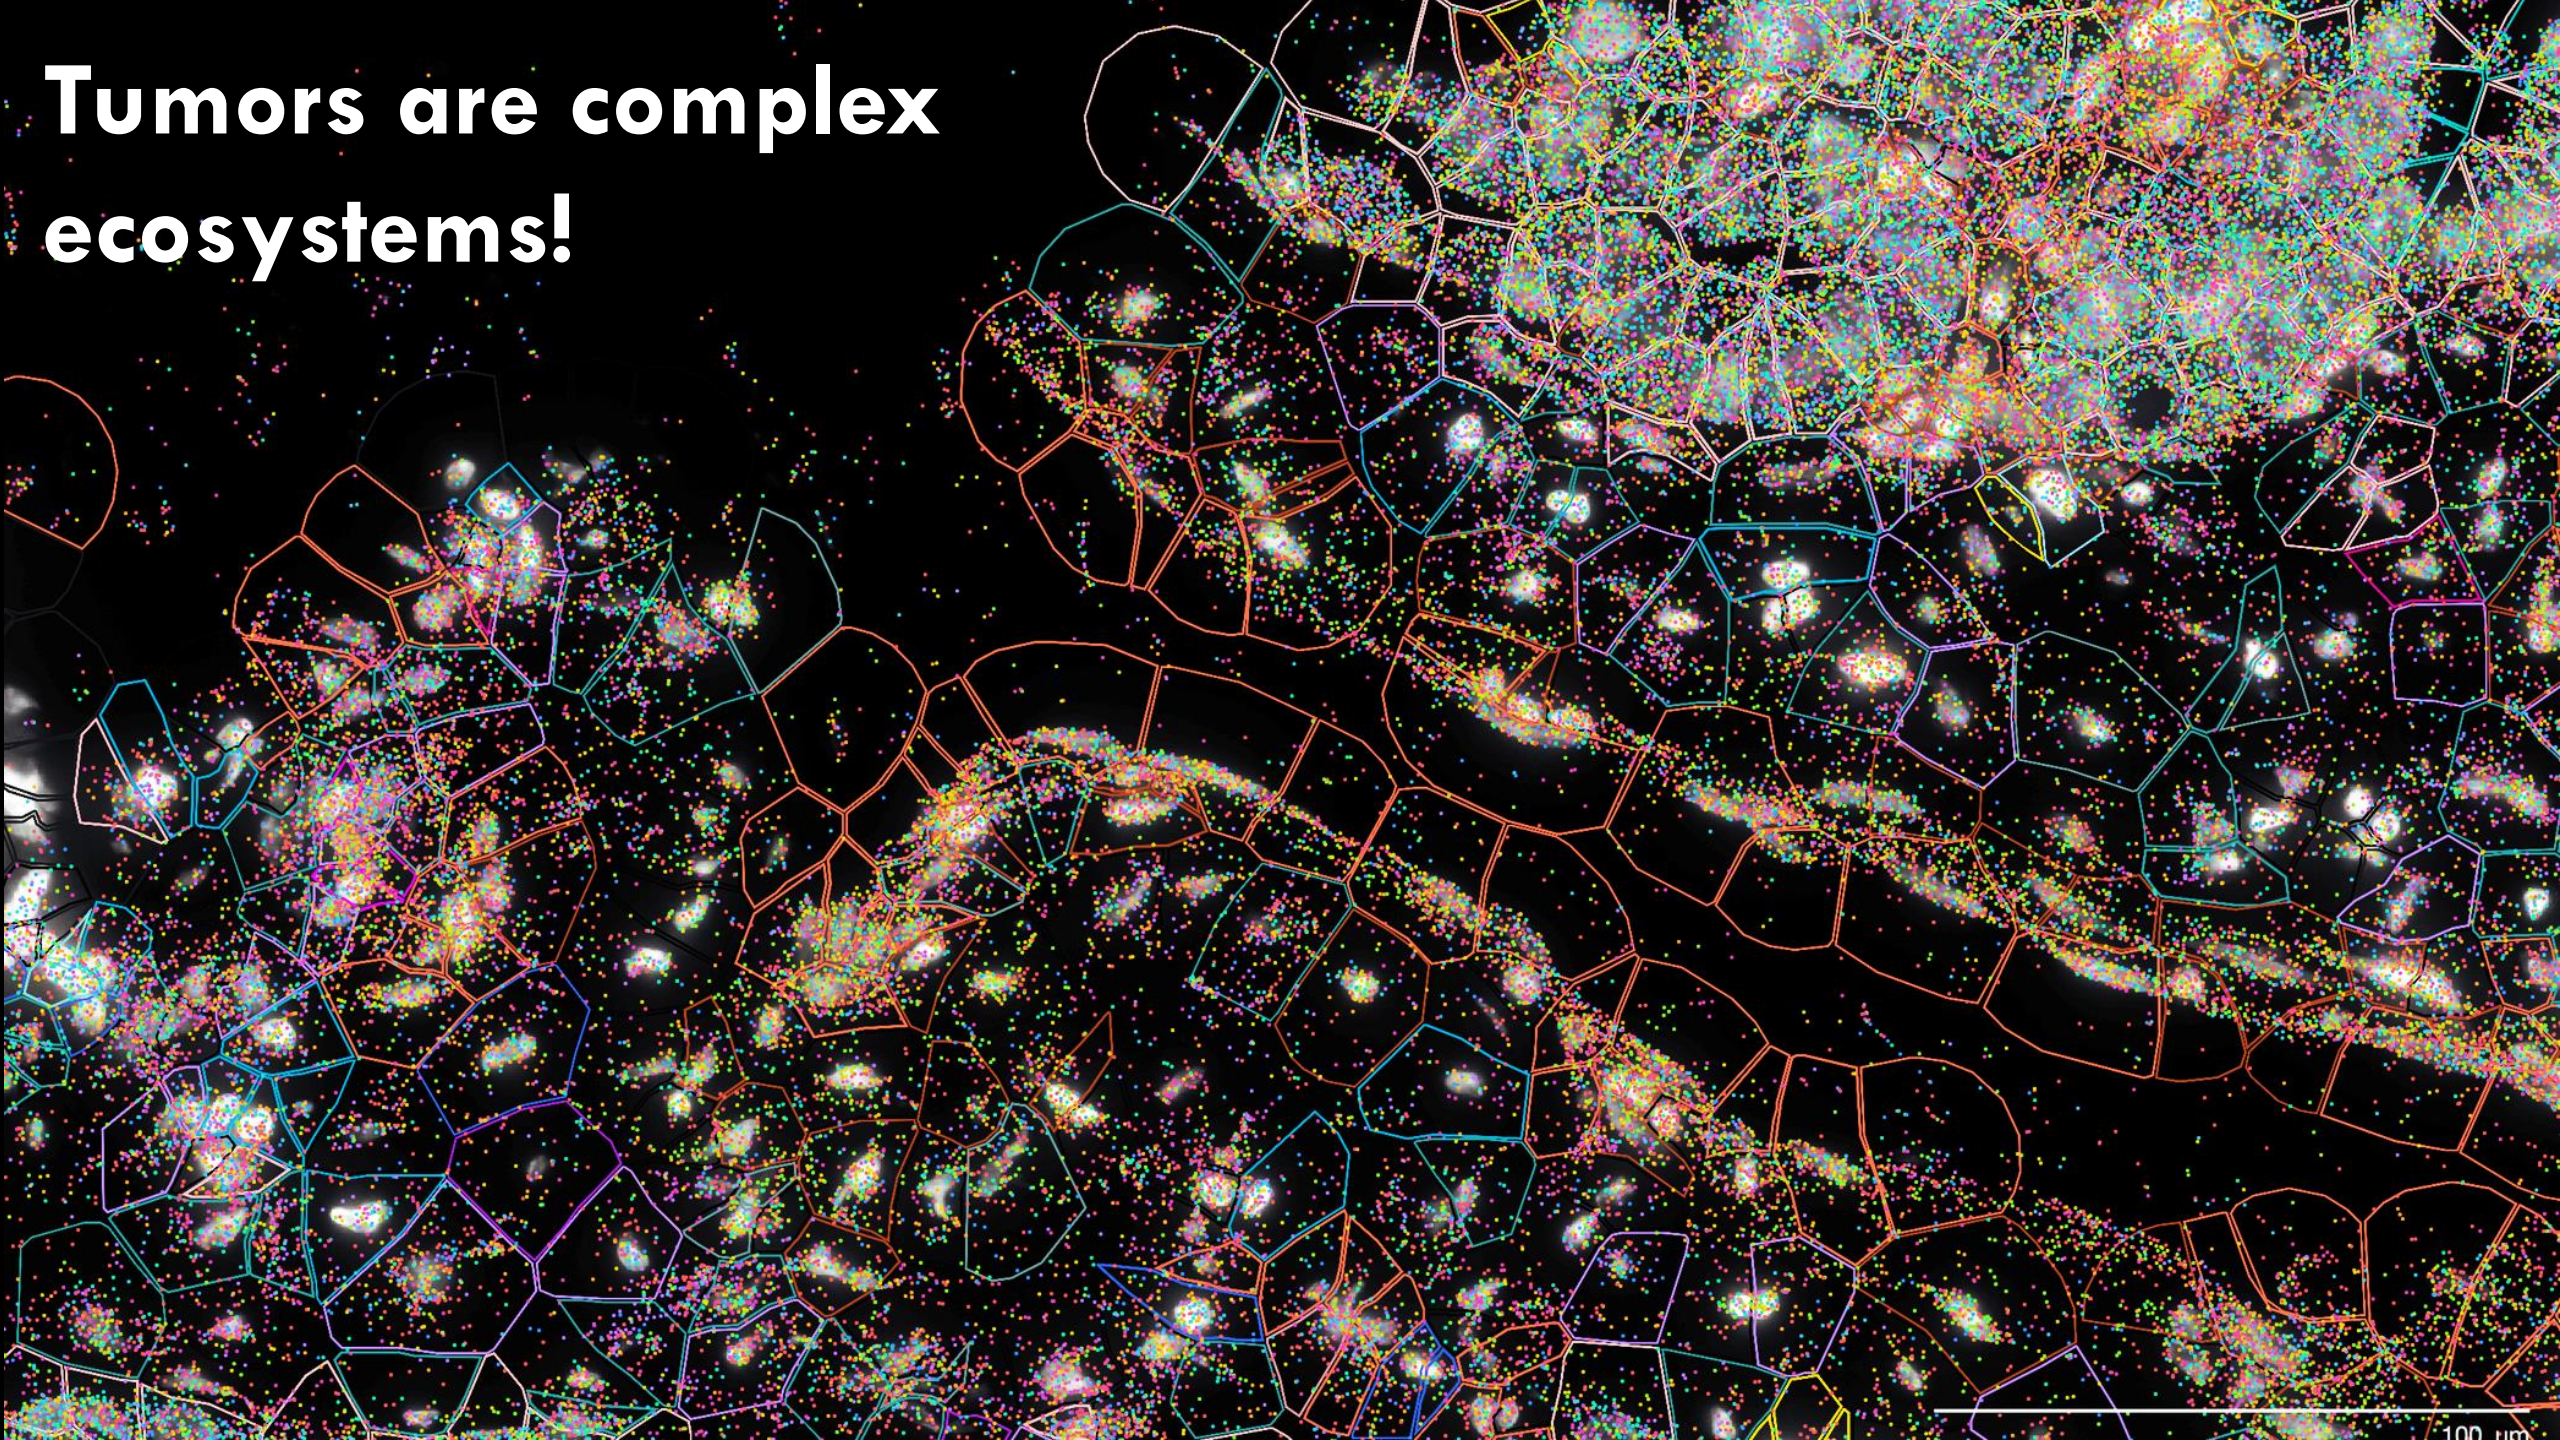

The Bottleneck: Scale, Cost, and Complexity

10X Xenium-Breast Cancer

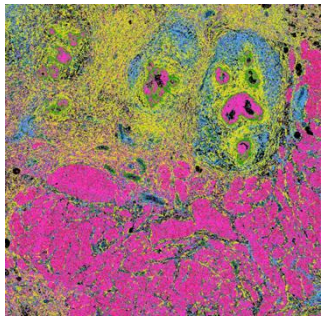

Nanostring CosMx –Lung

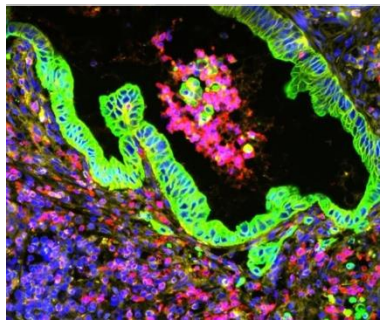

CODEX – Lymph node

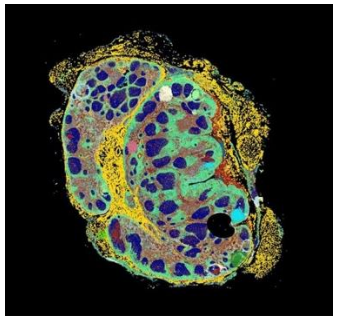

IMC - Human Placenta

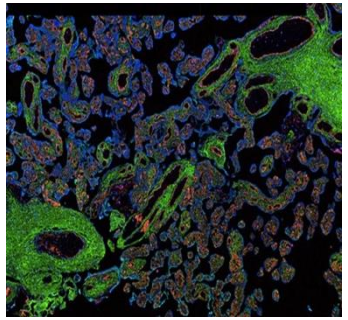

Images from https://hubmapconsortium.org, 10Xgenomics.com and Nanostring.com

Spatial omics – multiplexed imaging:

- ✓ Many genes/proteins simultaneously
- x High cost and long turnaround
- x Specialized equipment
- x Tissue-destructive
- x Small capture area

Immunohistochemistry (IHC)

Traditional histopathology:

- ✓ Fast & cheap
- ✓ Standardized protocols
- x Unaligned images (no multiplexing)

Virtual staining

Source domain

Target domain

Zhu *et al.*, ICCV, 2017.

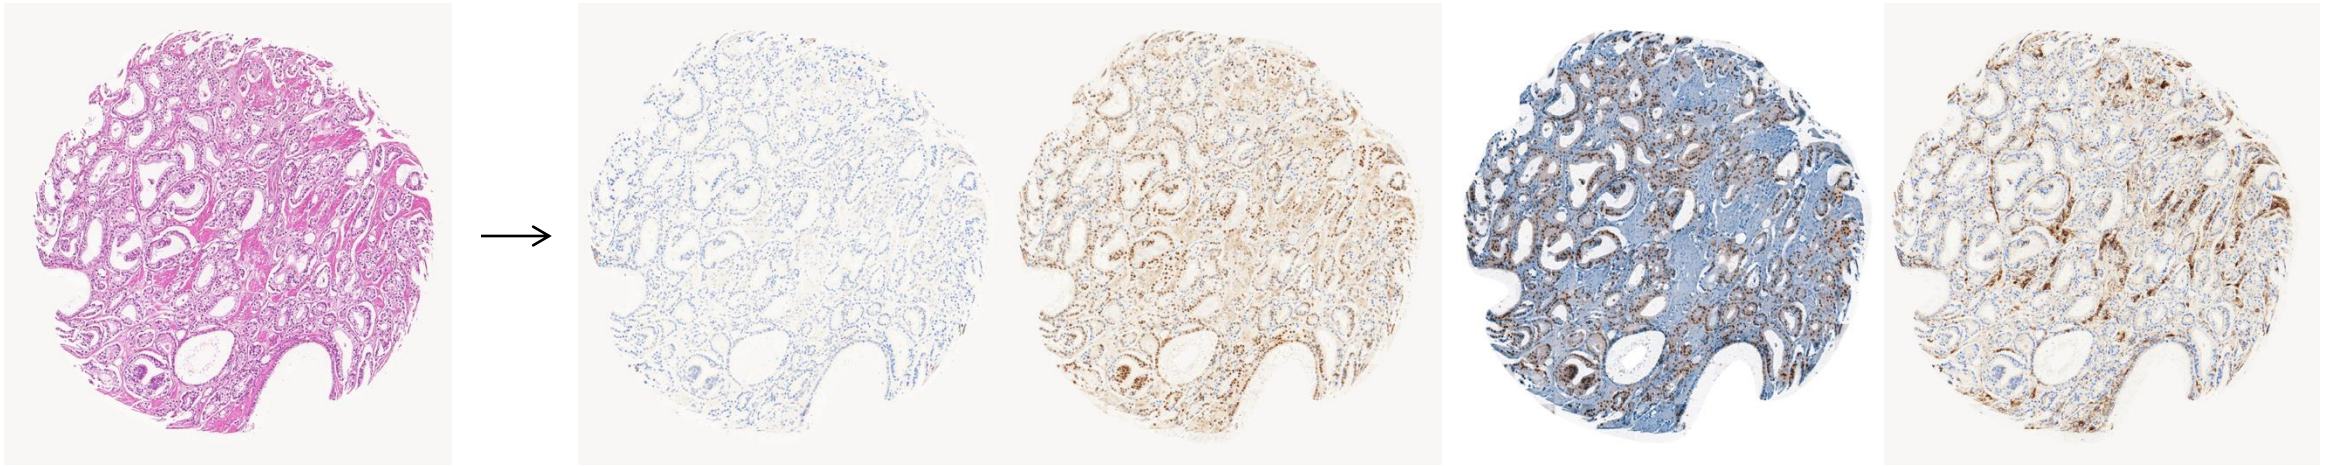

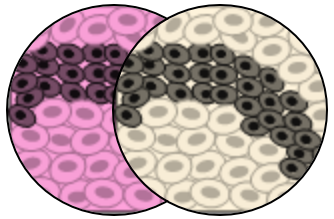

Virtual Multiplexer

Multiplexed tumor profiling with generative AI

Pushpak Pati

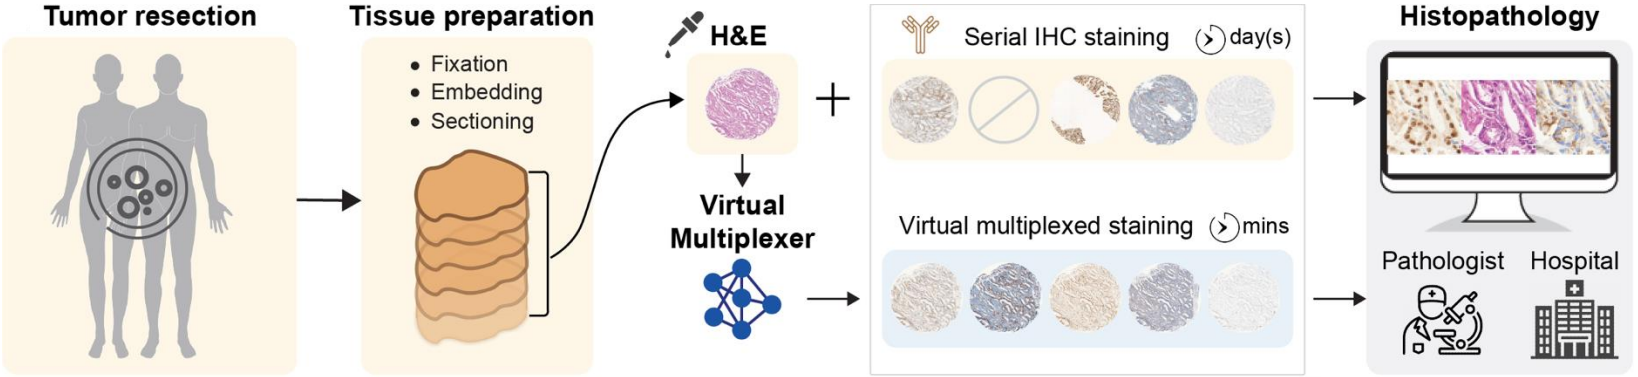

Virtual staining:

✓ Fast & cheap

✓ Standardized protocols

Pati *et al.*, *Nature Machine Intelligence*, **6**, 1077–1093 (2024)

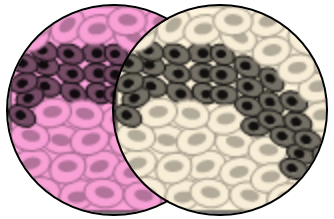

Virtual Multiplexer

Multiplexed tumor profiling with generative AI

Pushpak Pati

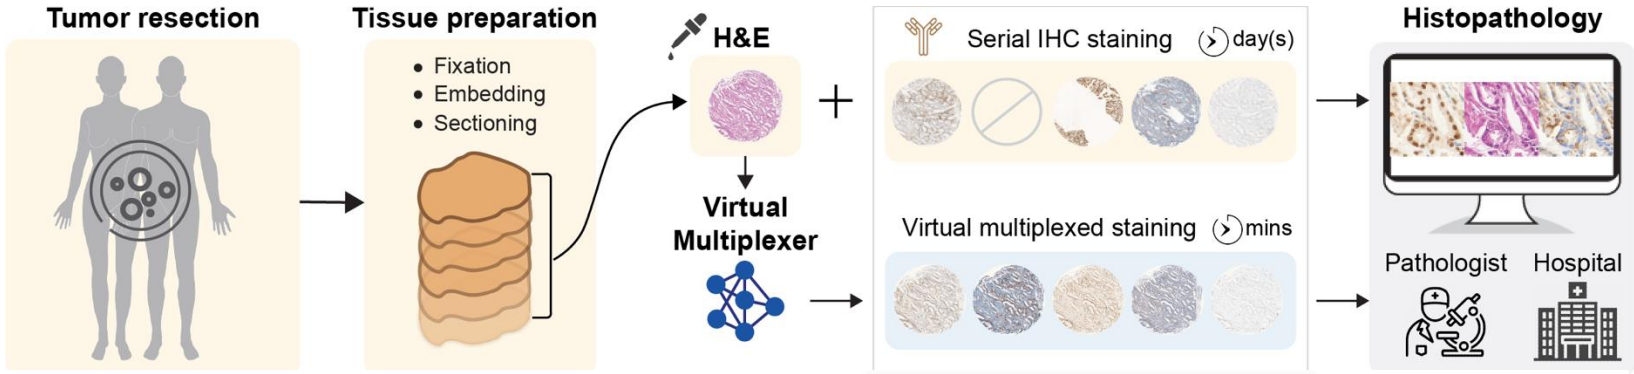

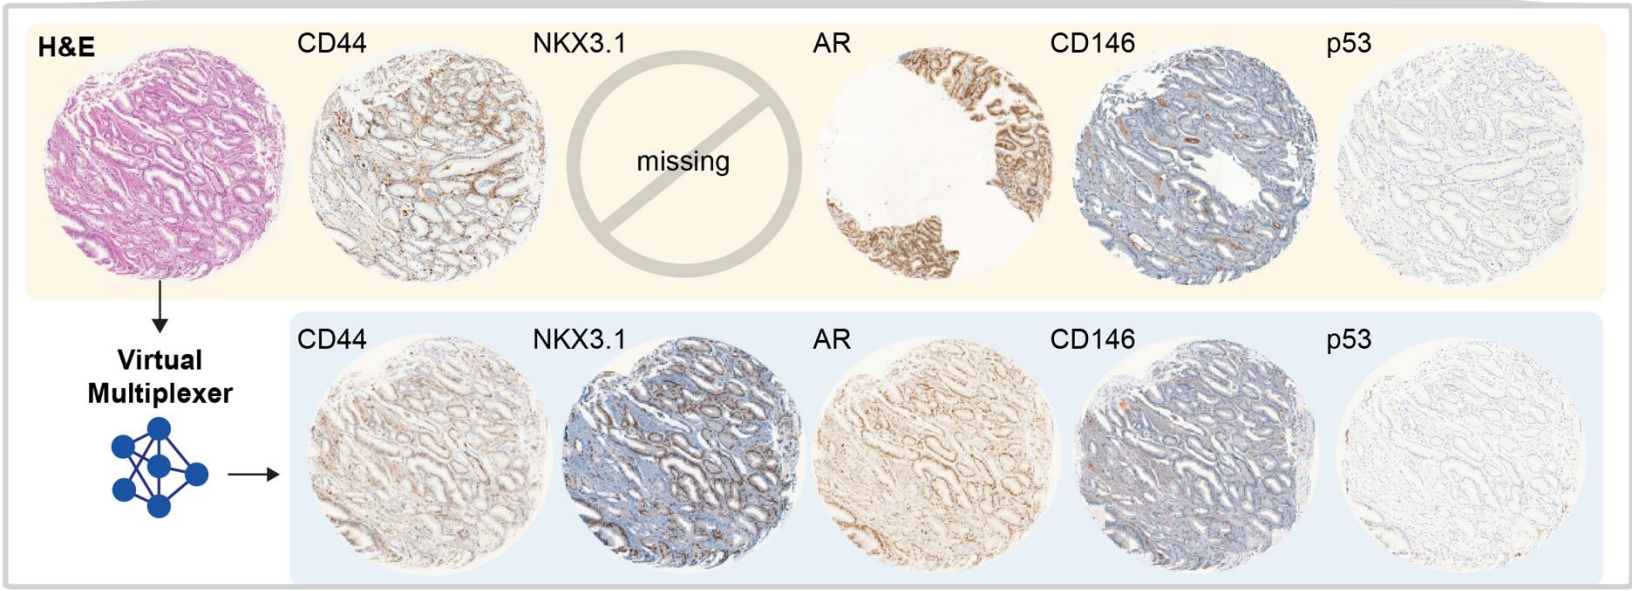

Virtual staining:

✓ Fast & cheap

✓ Standardized protocols

✓ Generates realistic virtual stains for several markers

✓ Unpaired data → paired, virtually multiplexed data

Pati *et al.*, *Nature Machine Intelligence*, **6**, 1077–1093 (2024)

Unpaired image-to-image translation

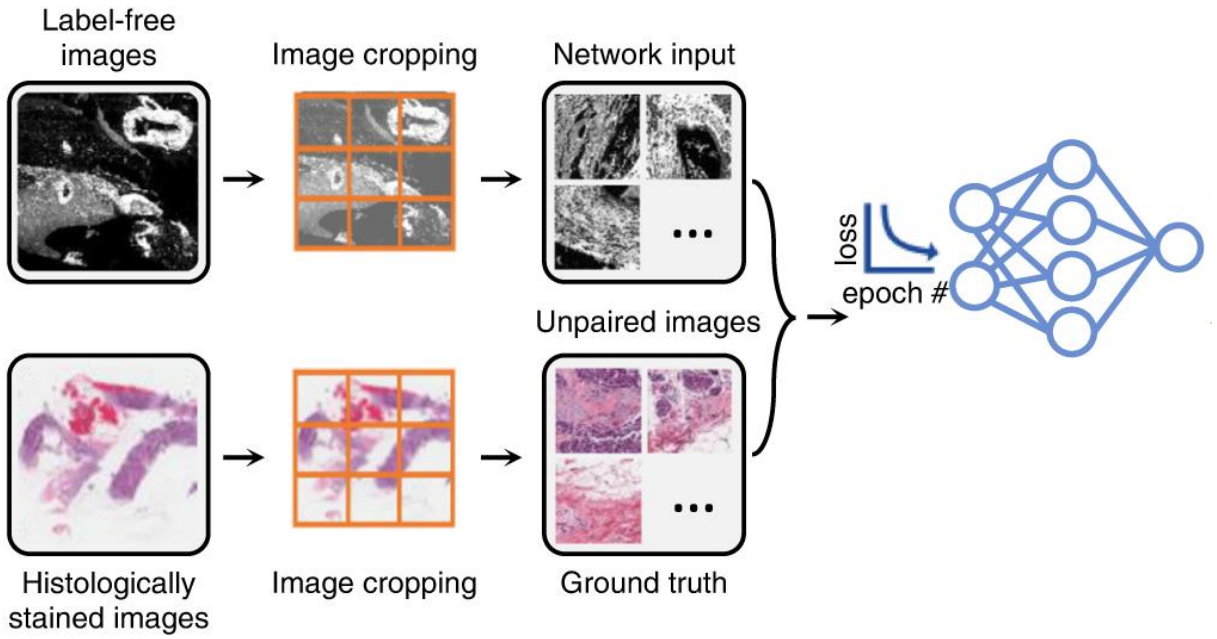

Bai *et al.*, 2023.

Unpaired image-to-image translation

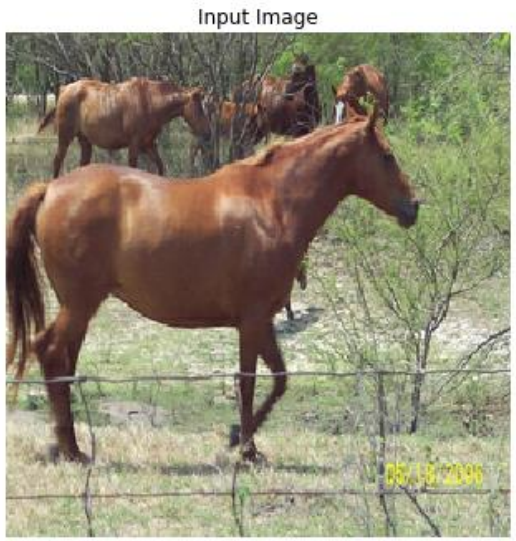

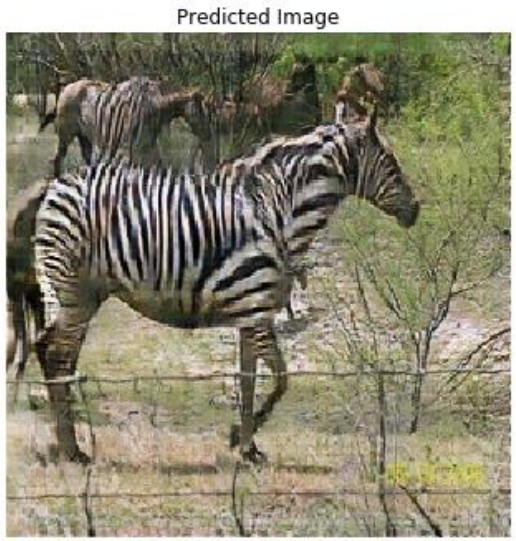

https://www.tensorflow.org/tutorials/generative/cyclegan

Unpaired image-to-image translation

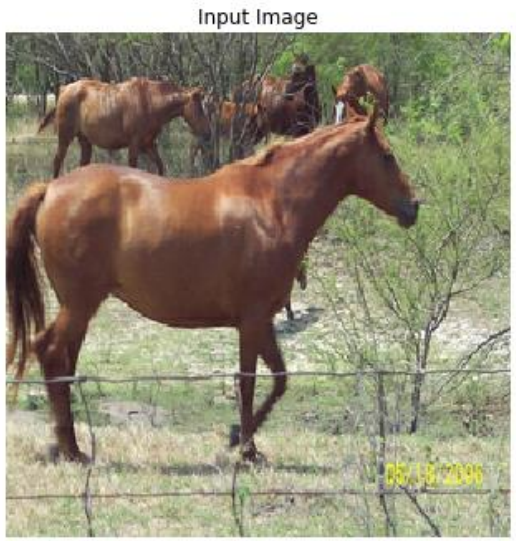

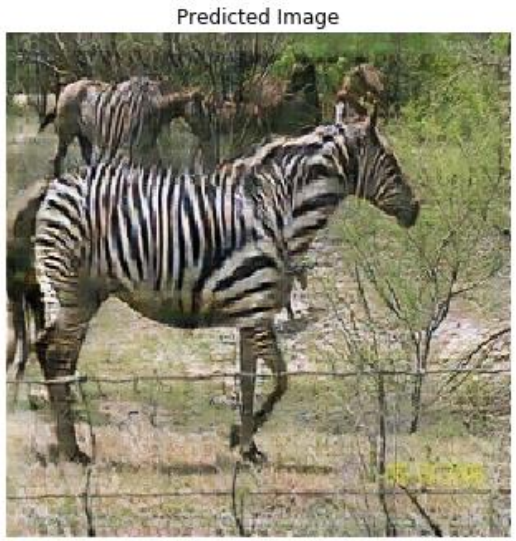

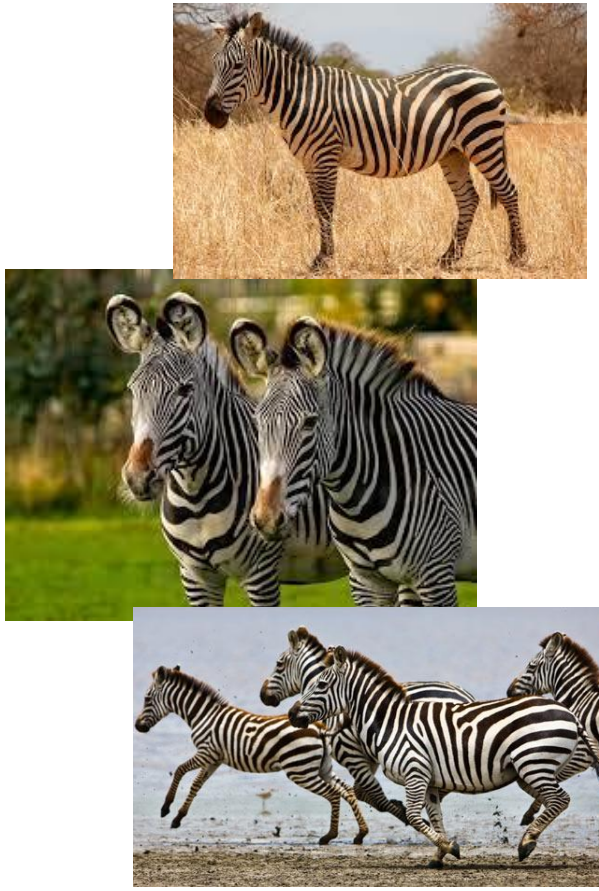

https://www.tensorflow.org/tutorials/generative/cyclegan

Cycle consistency loss

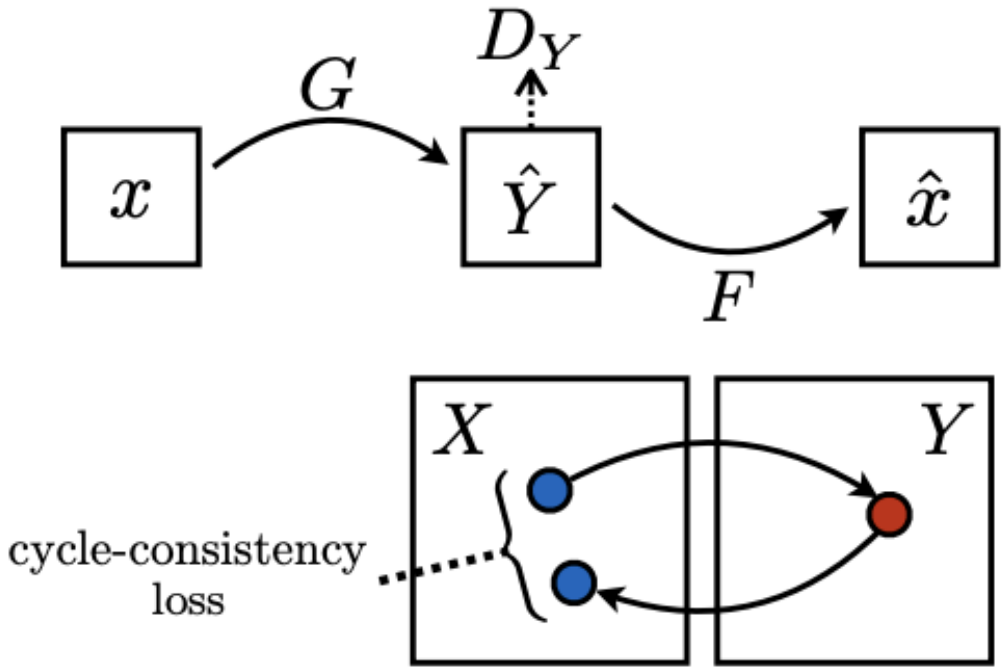

Zhu *et al.*, ICCV 2017

Cycle consistency loss

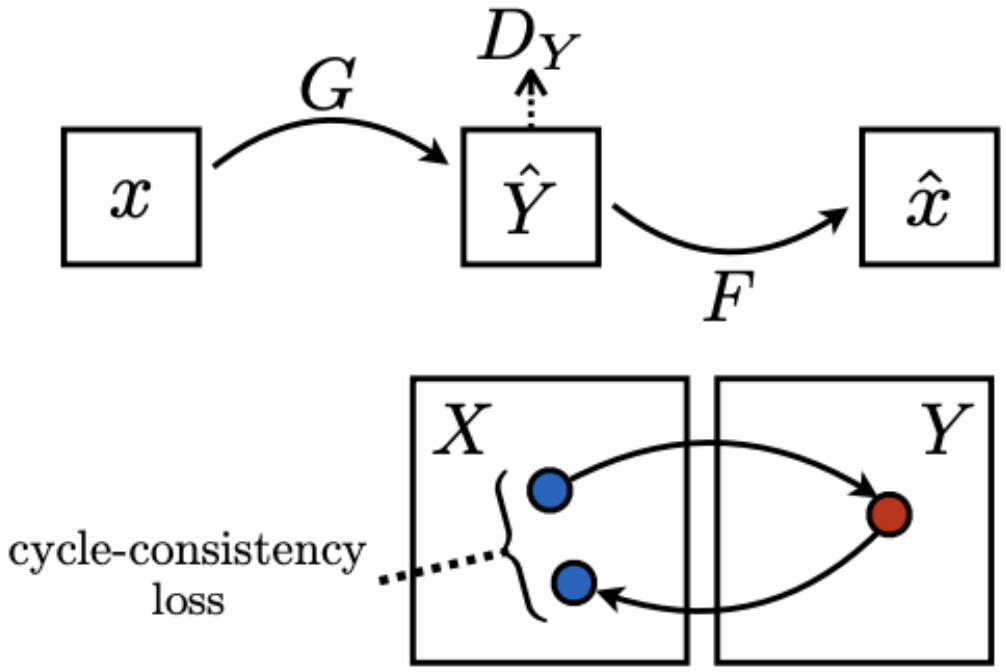

Zhu *et al.*, ICCV 2017

https://github.com/junyanz/CycleGAN

VirtualMultiplexer: model architecture

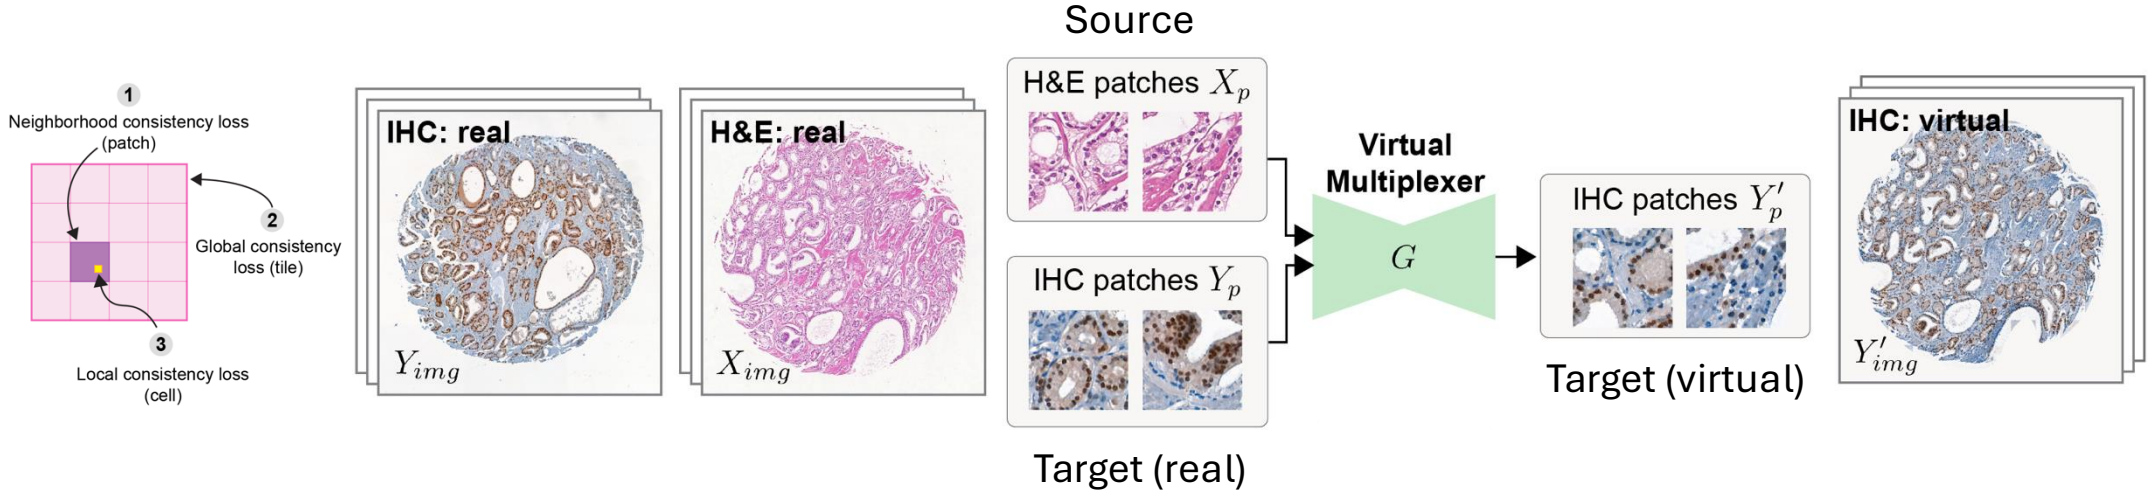

Pati *et al.*, *Nature Machine Intelligence*, **6**, 1077–1093 (2024)

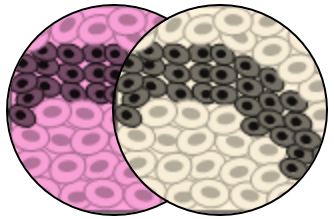

VirtualMultiplexer: model architecture

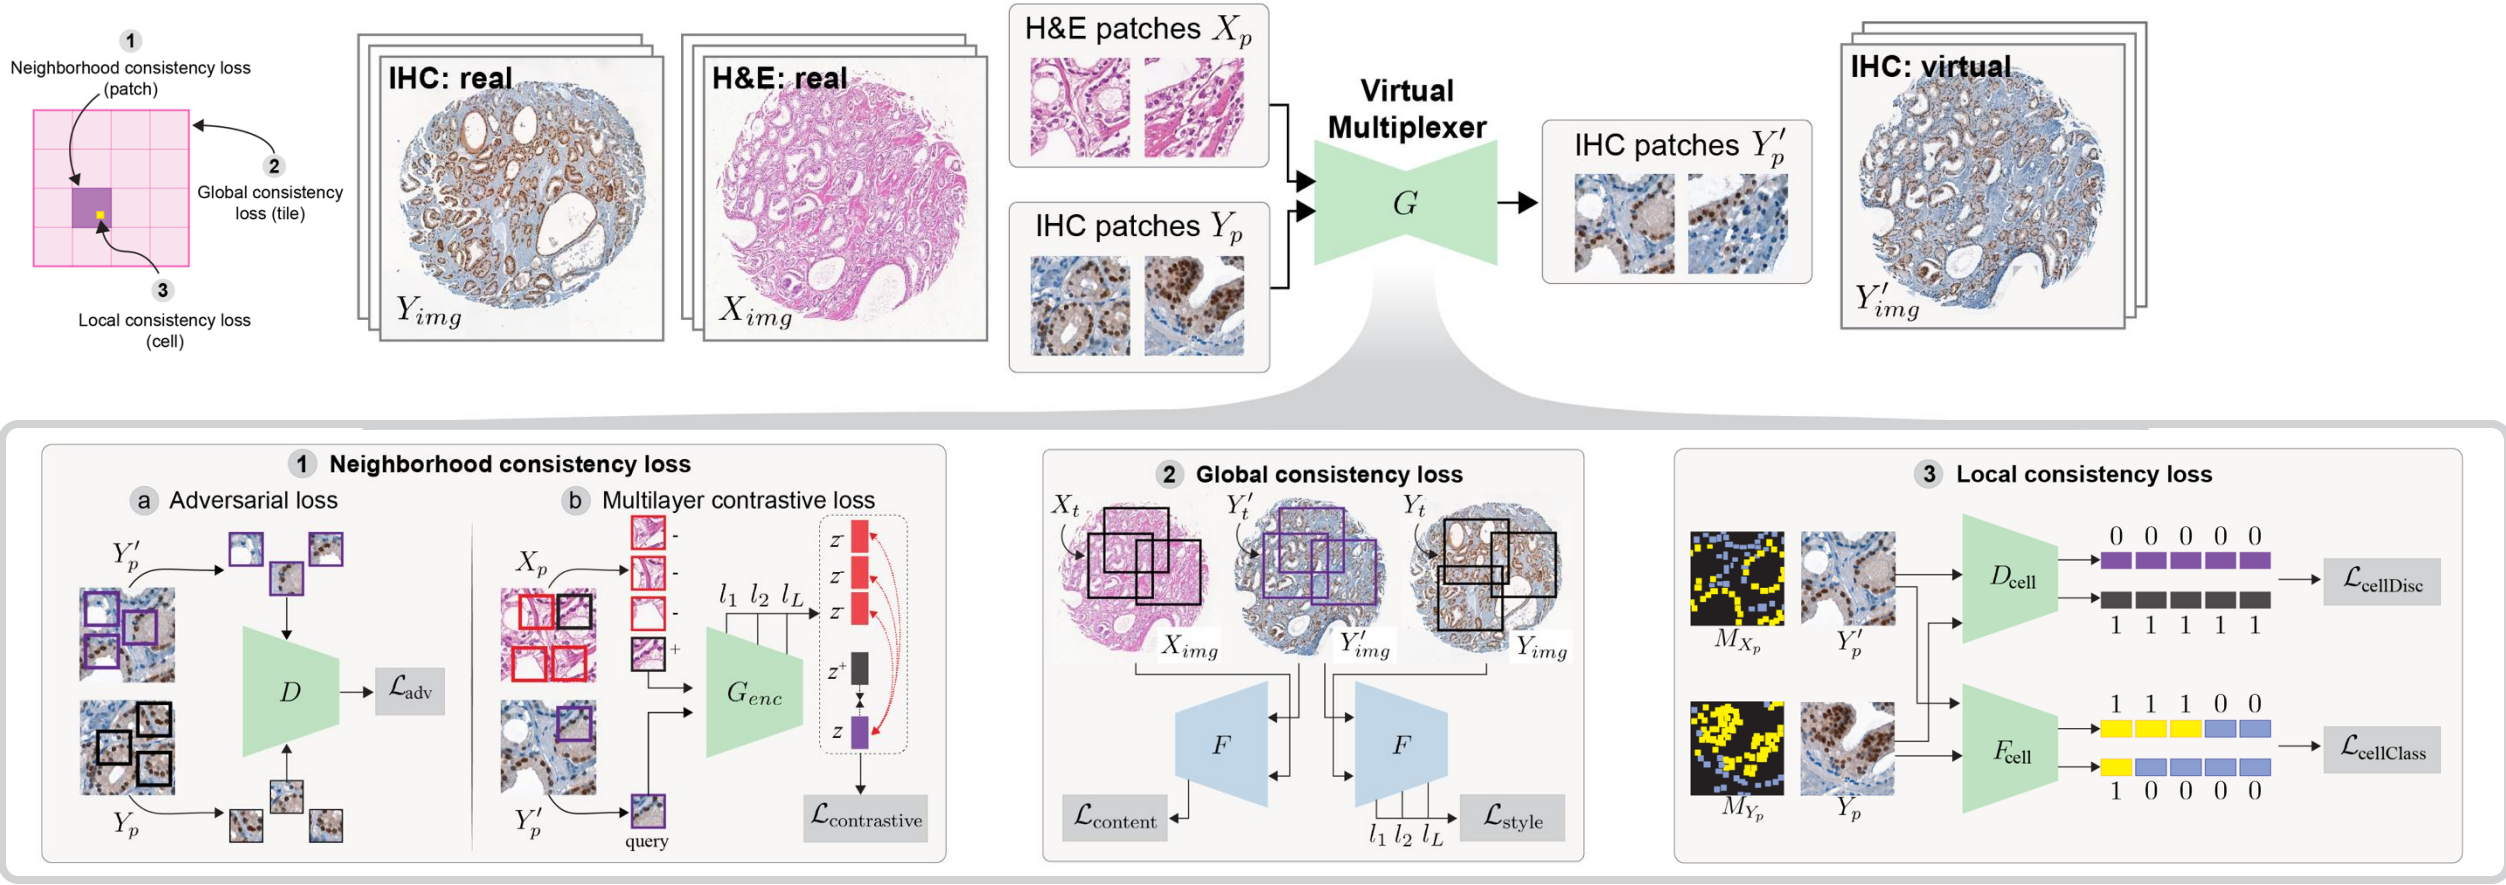

Notations: Frozen model | Trainable model | Real/Virtual | Positive/Negative cells

G : Encoder | D : Patch discriminator | D_{cell} : Cell discriminator | F : Feature extractor | F_{cell}: Cell classifier | $l_1\ l_2\ l_L$: Layers

X_p : Real H&E patches | Y_p, Y'_p : Real/Virtual IHC patches | $X_{img}, Y_{img}, Y'_{img}$: Images | X_t, Y_t, Y'_t : Tiles

Effect of different losses

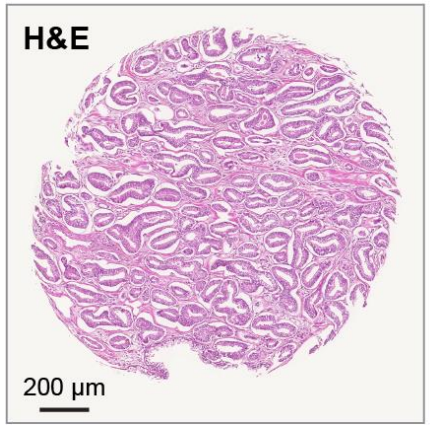

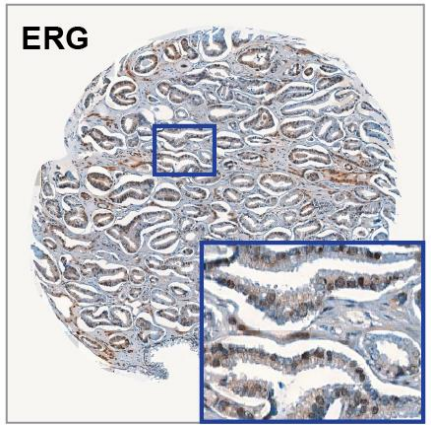

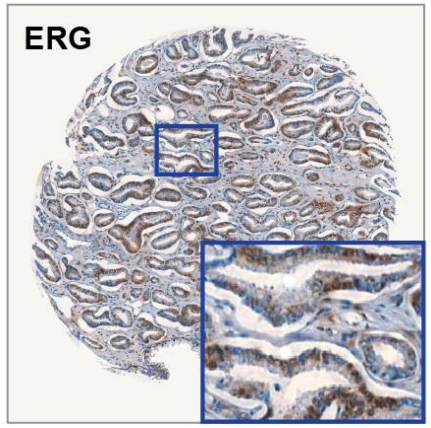

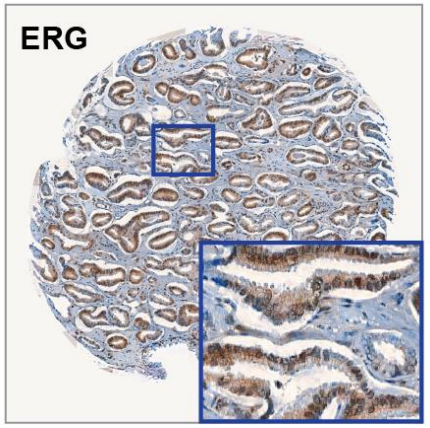

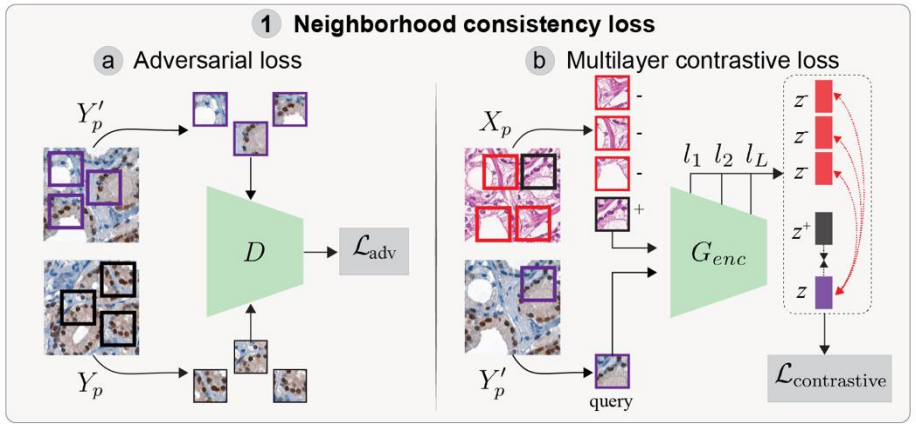

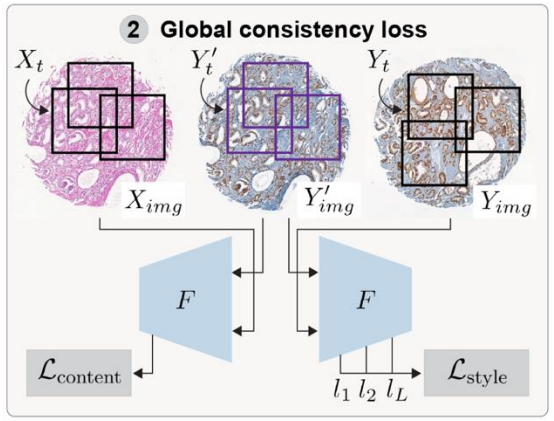

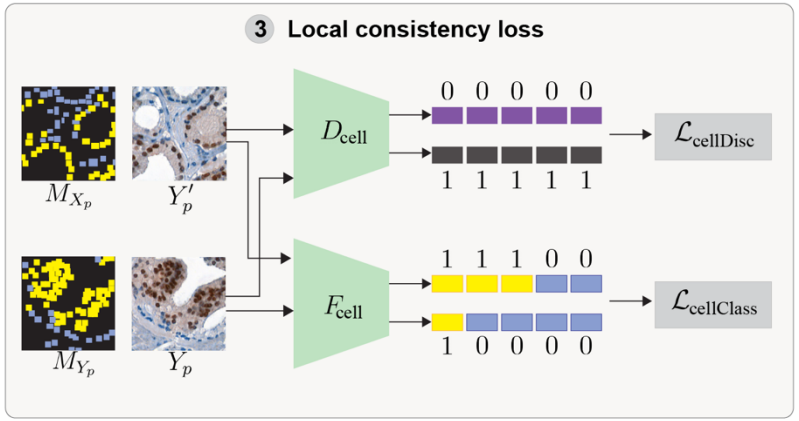

Pati *et al.*, *Nature Machine Intelligence*, **6**, pages1077–1093 (2024)

Performance assessment

Real H&E

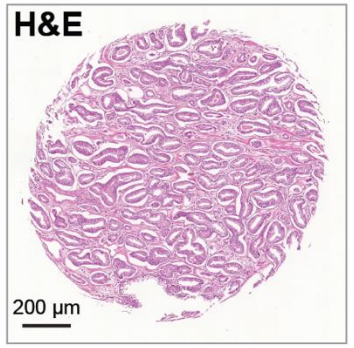

Real IHC

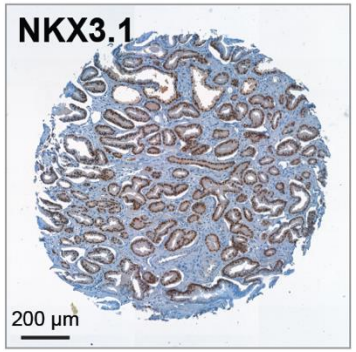

Virtual IHC

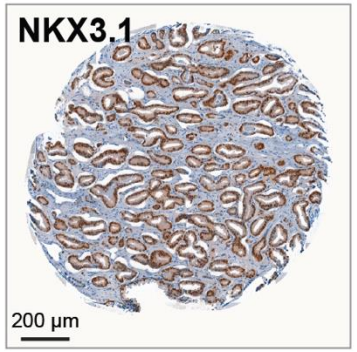

Fréchet Inception Distance

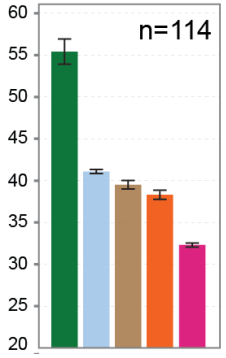

Visual Turing Test

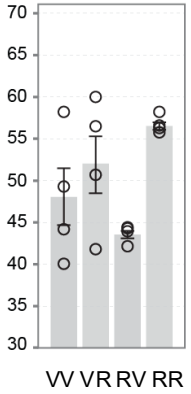

Staining quality

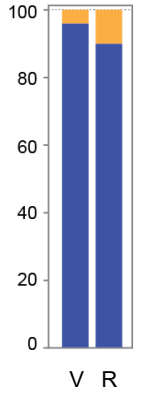

Unpaired S2S translation methods

- CycleGAN
- CUT
- CUT+KIN
- AI-FFPE
- VirtualMultiplexer

Visual Turing test

VV: Virtual as Virtual
VR: Virtual as Real
RV: Real as Virtual
RR: Real as Real

Staining quality

- Acceptable
- Background
- Border
- Not acceptable

EMPaCT prostate cancer TMA*:

- 210 patients (4 cores / patient)
- 6 markers (NKX3.1, AR, CD44, CD146, p53, ERG)

- The VirtualMultiplexer **outperformed all other methods** in all 6 IHC markers in terms of FID score.
- Visual Turing test **close to random guess**: average sensitivity/ specificity of 52.1% / 54.1% across all six markers
- Virtual staining quality was **on par or higher than real** staining quality for 4/6 markers

*https://zenodo.org/records/10066853

Picasso

Van Gogh

Monet

Generalization on OOD WSIs

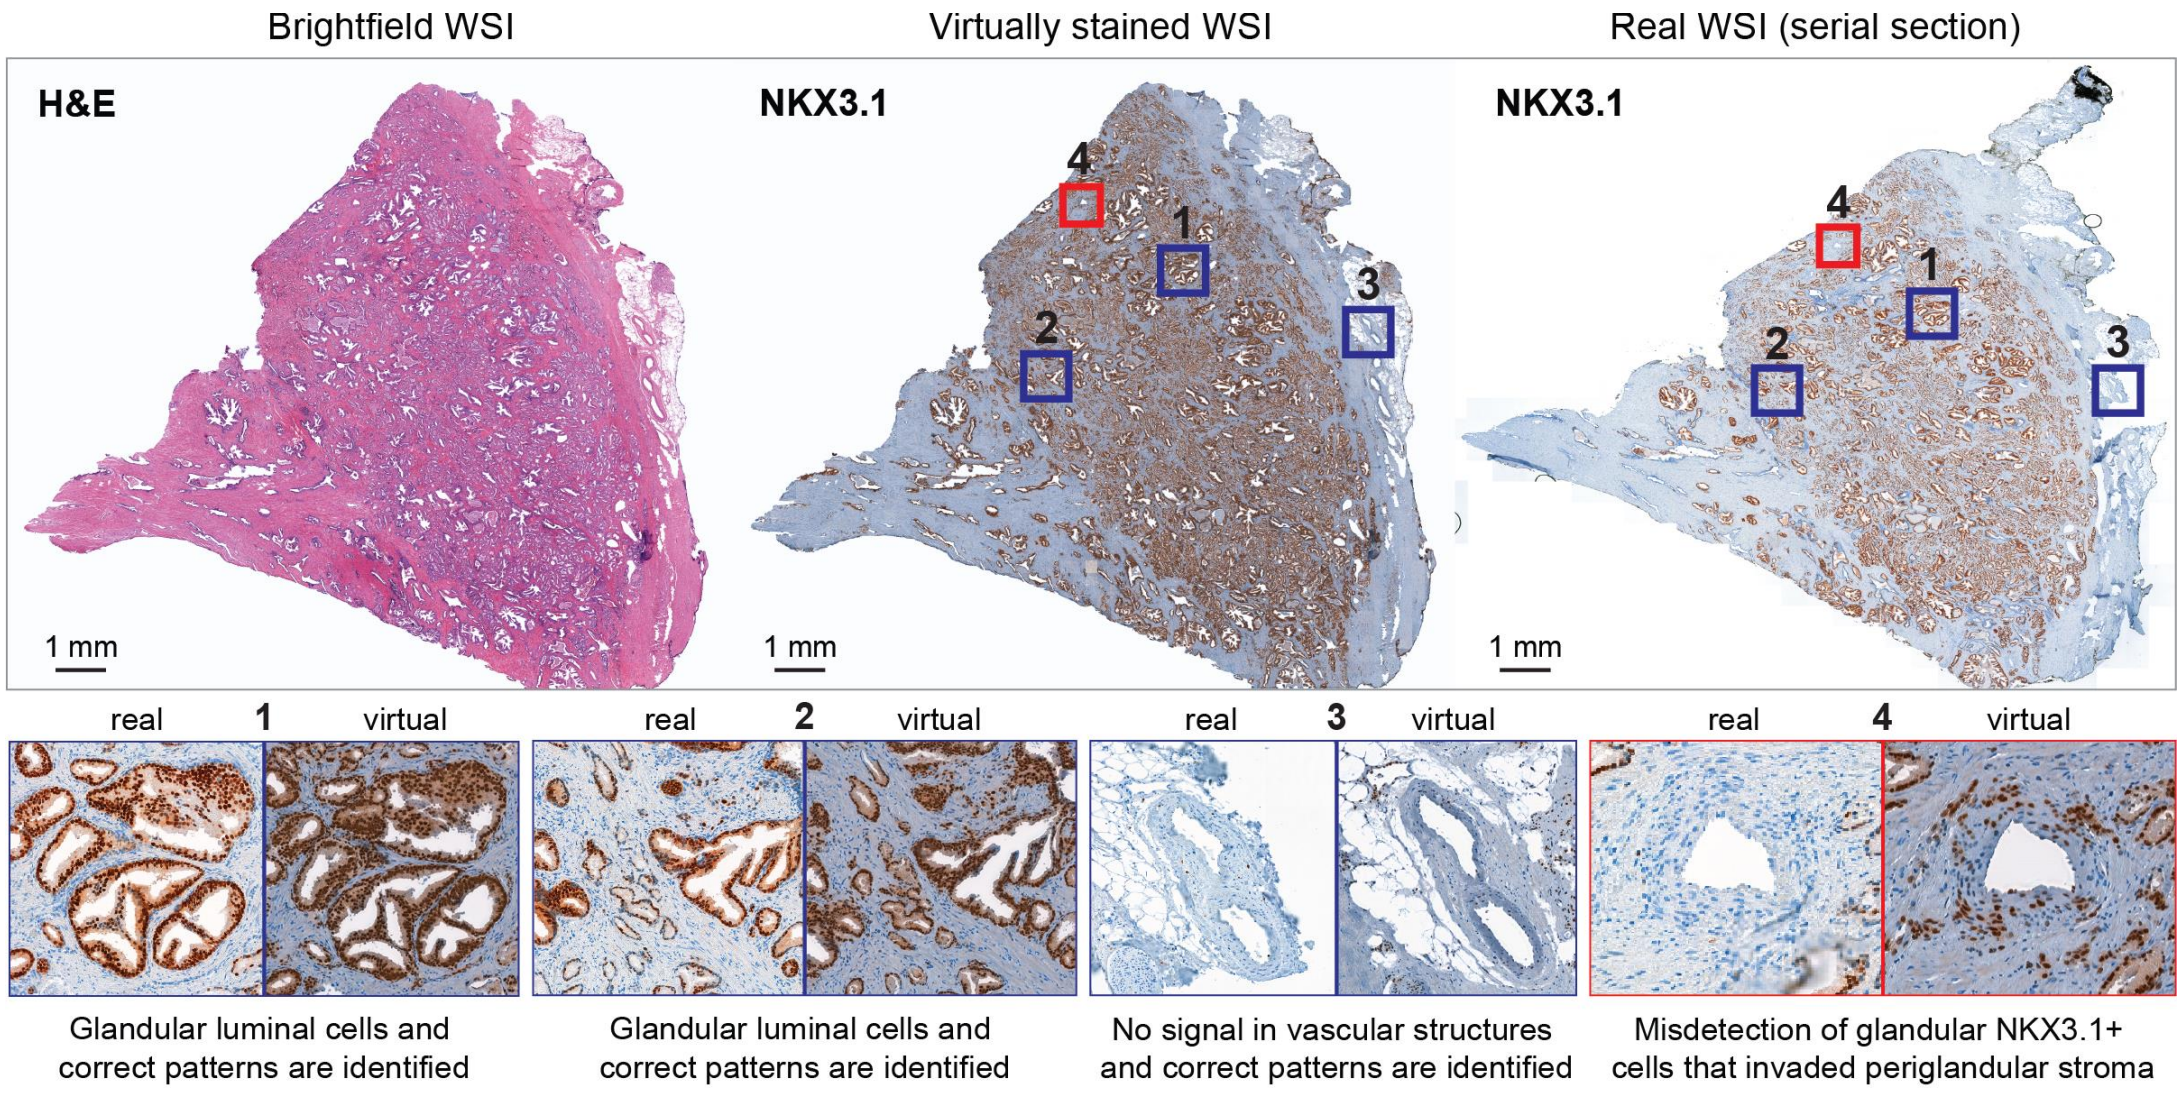

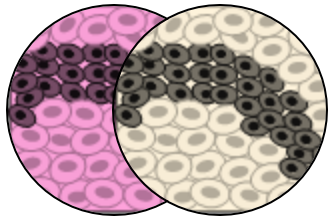

Improved clinical predictions

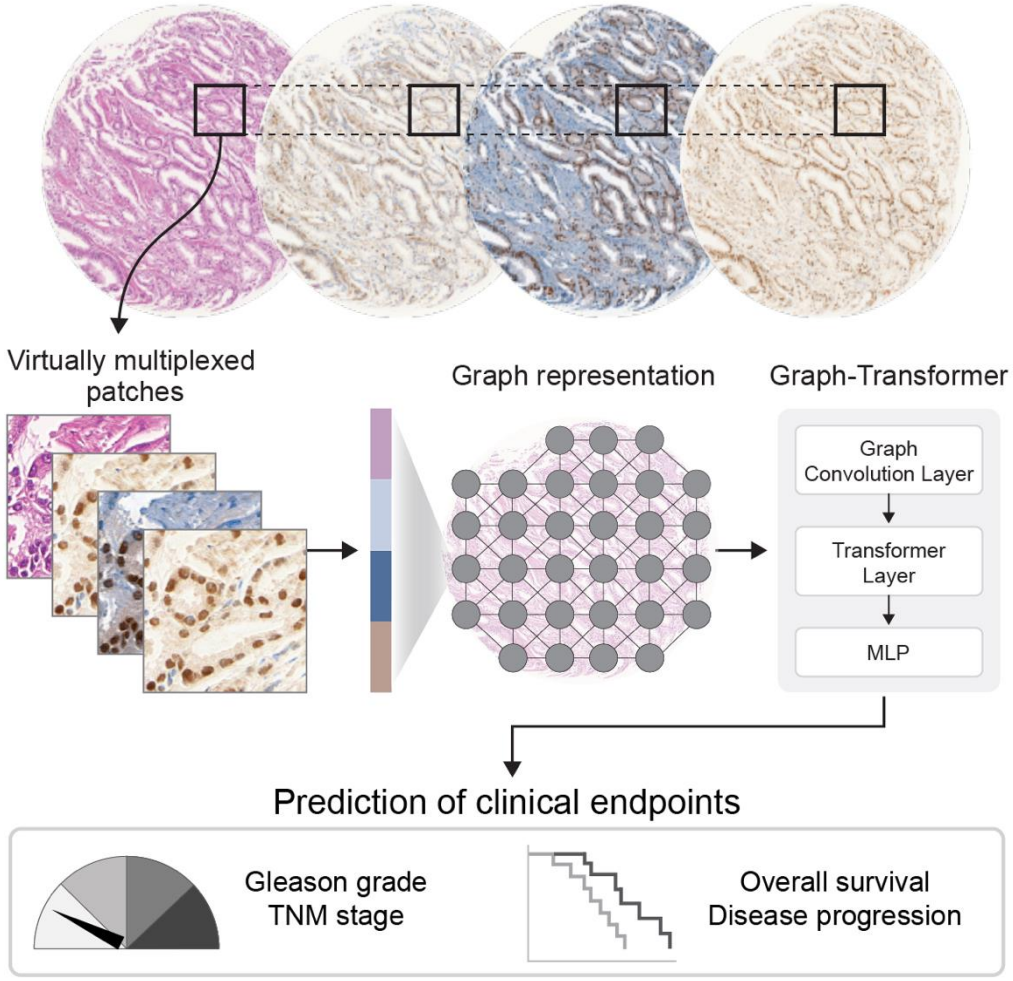

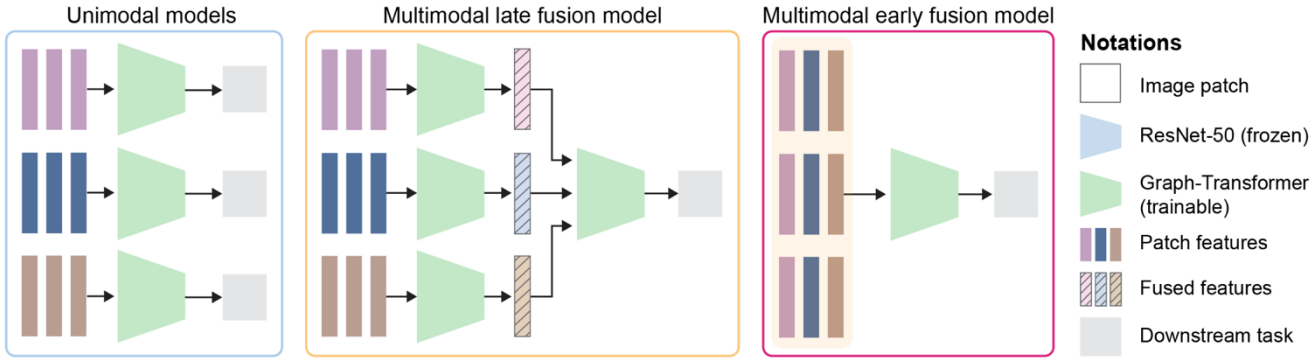

Improved clinical predictions

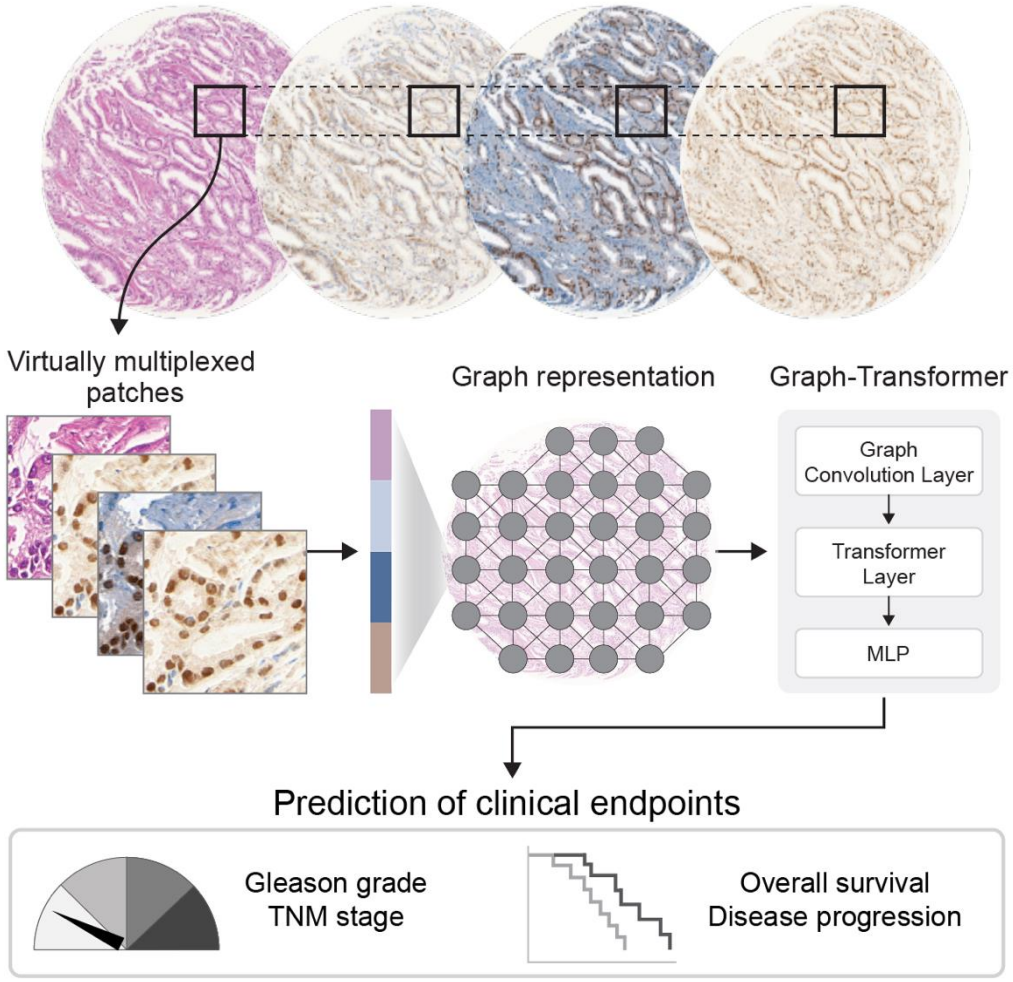

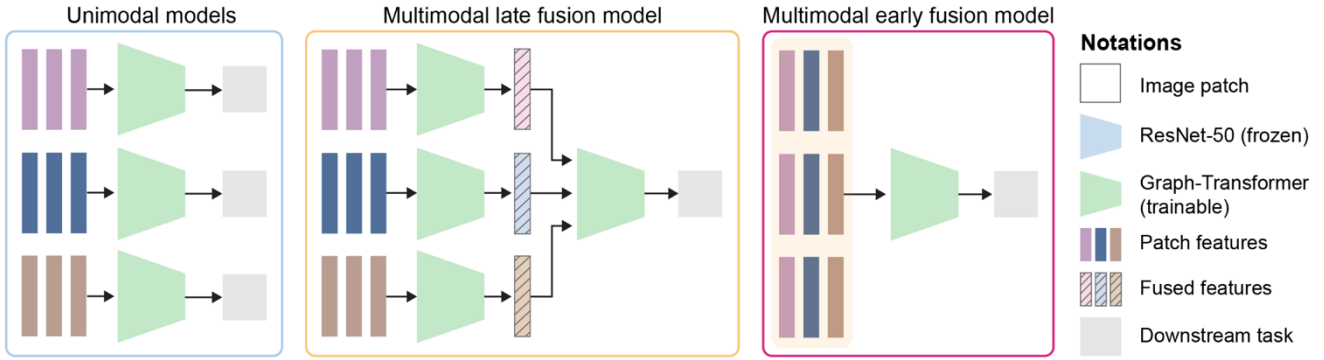

EMPaCT (results on test data)

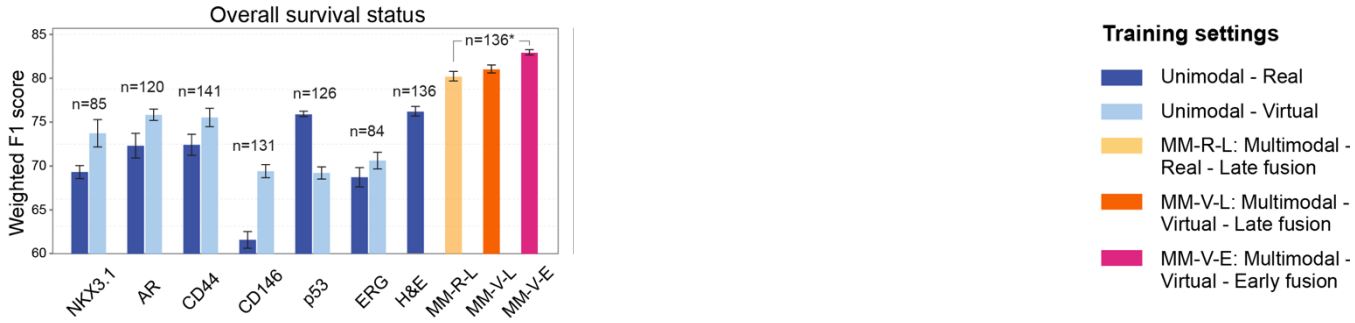

Transfer to other datasets - PANDA

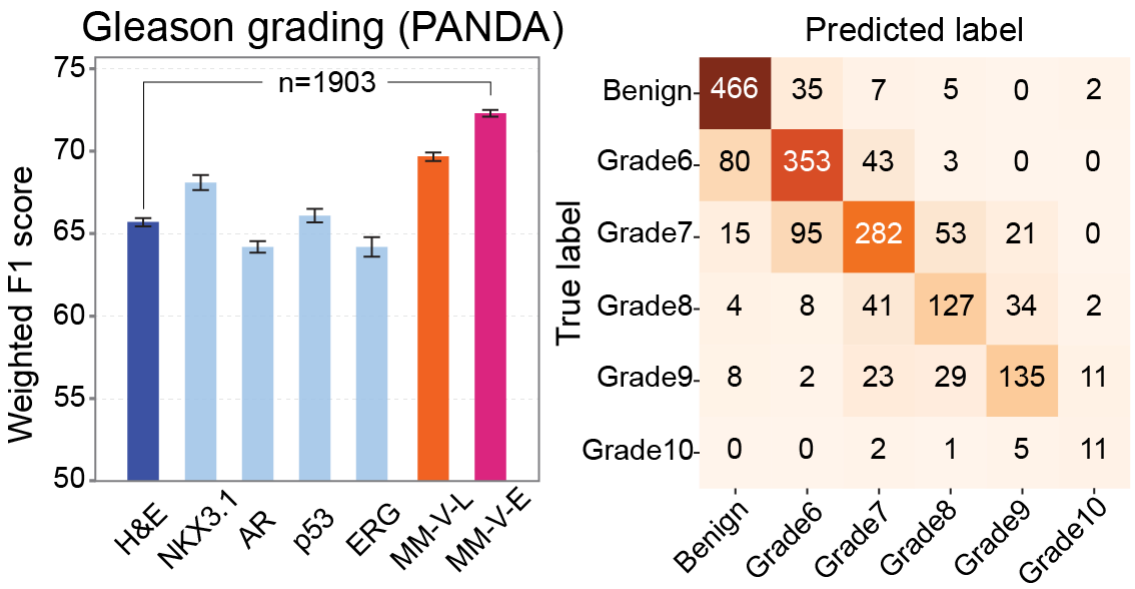

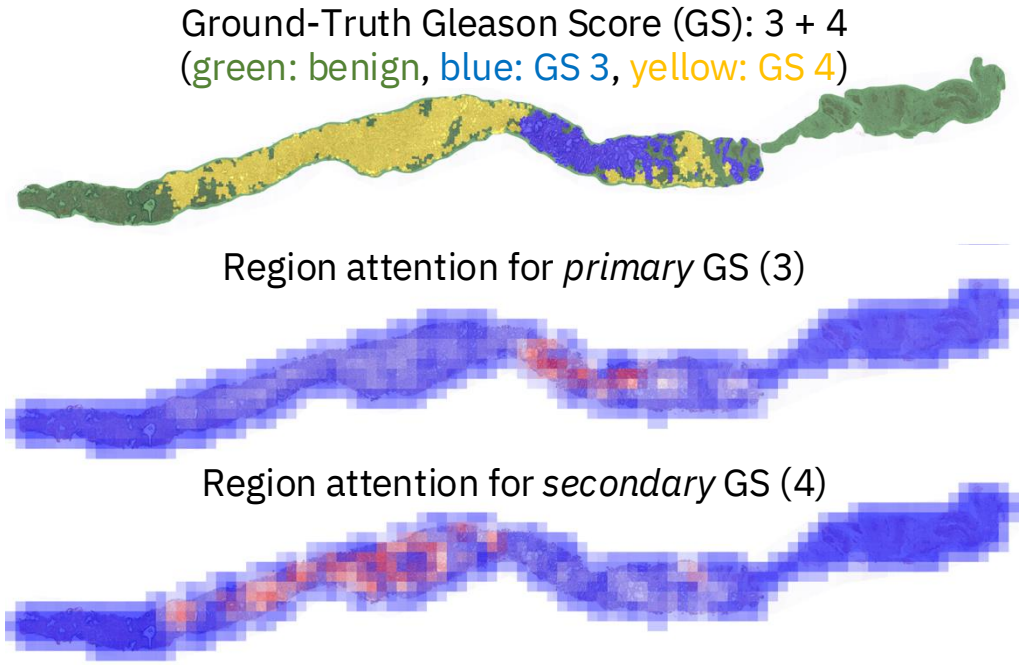

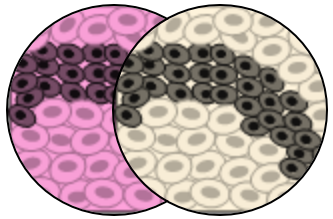

Draw me a cell?

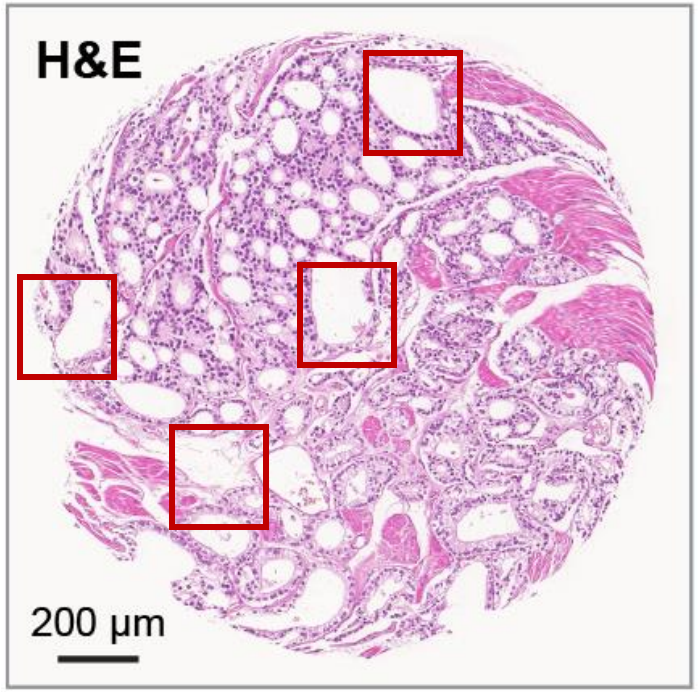

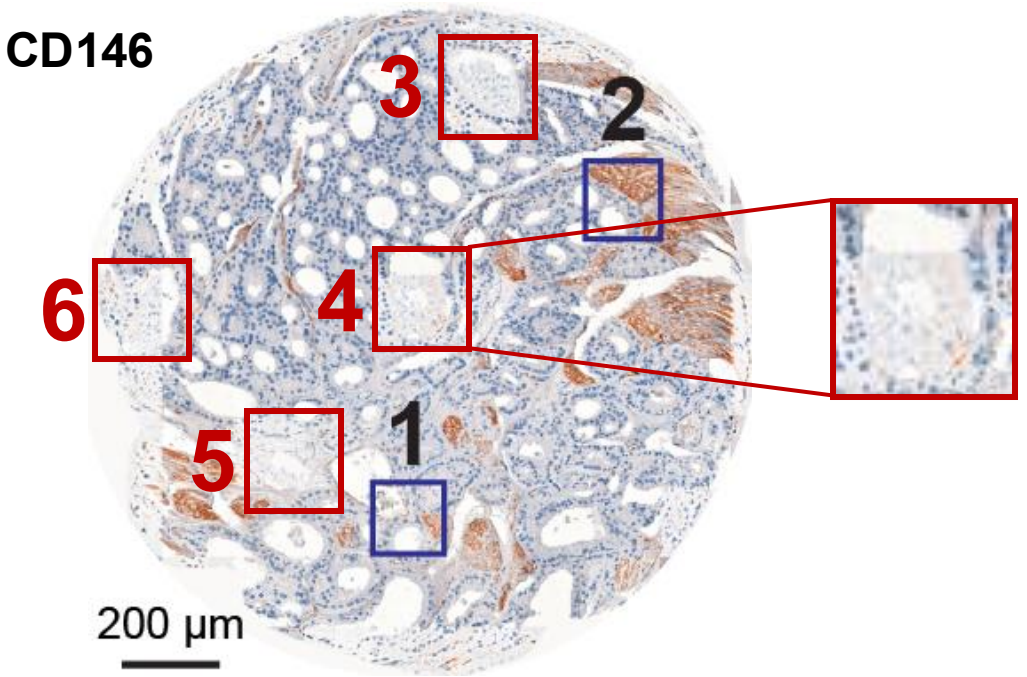

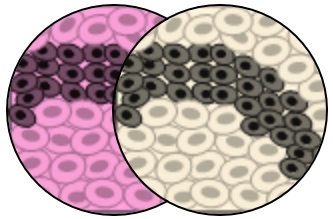

Hallucination detection

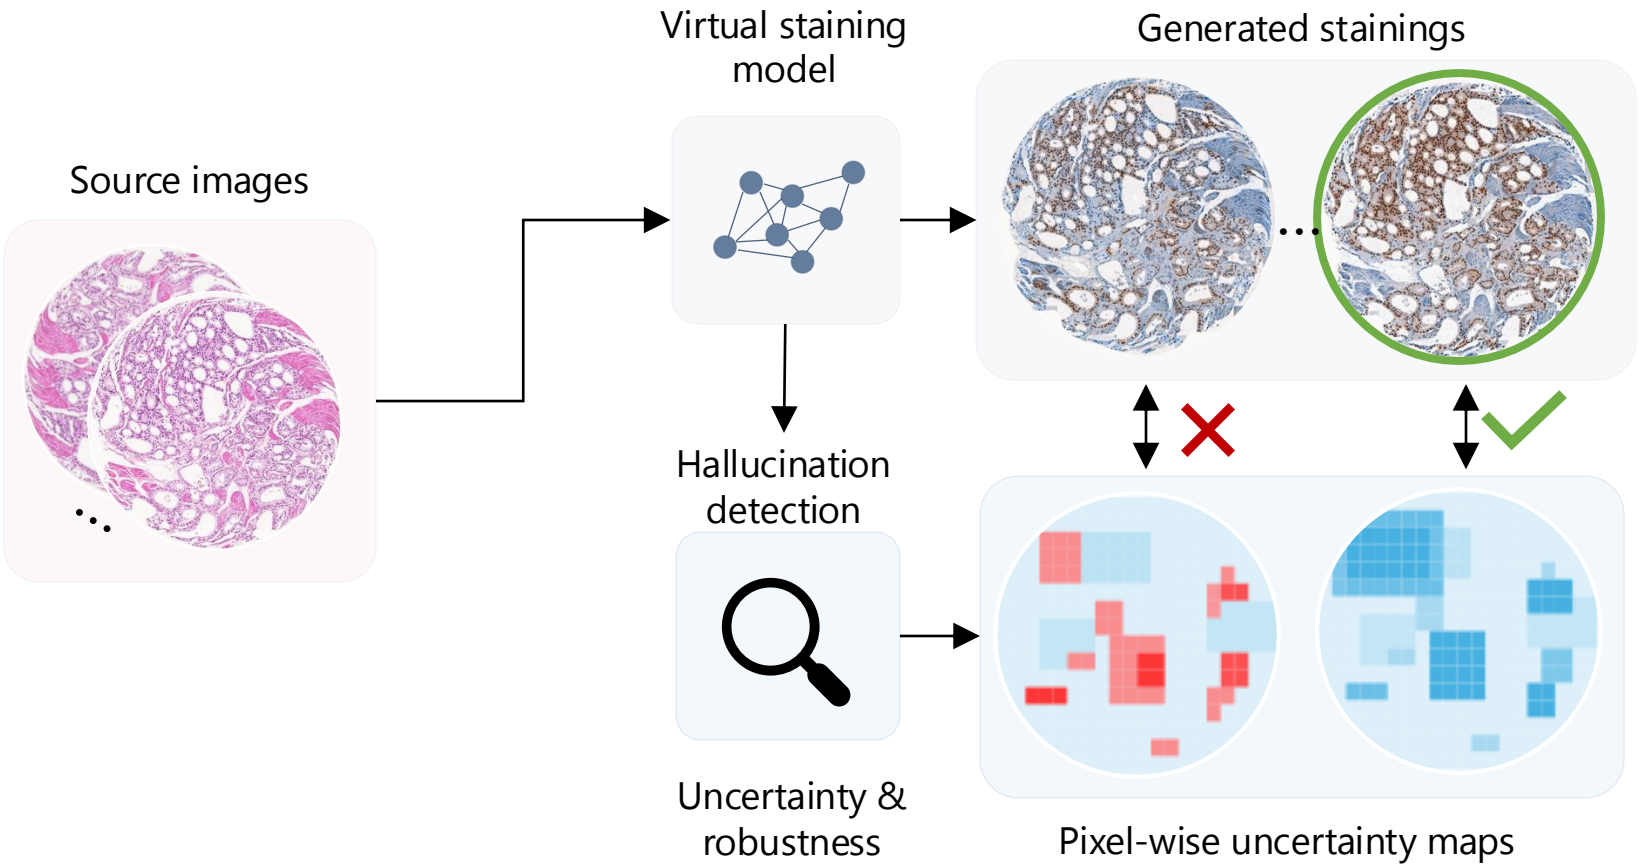

Virtual Multiplexer

Multiplexed tumor profiling with generative AI

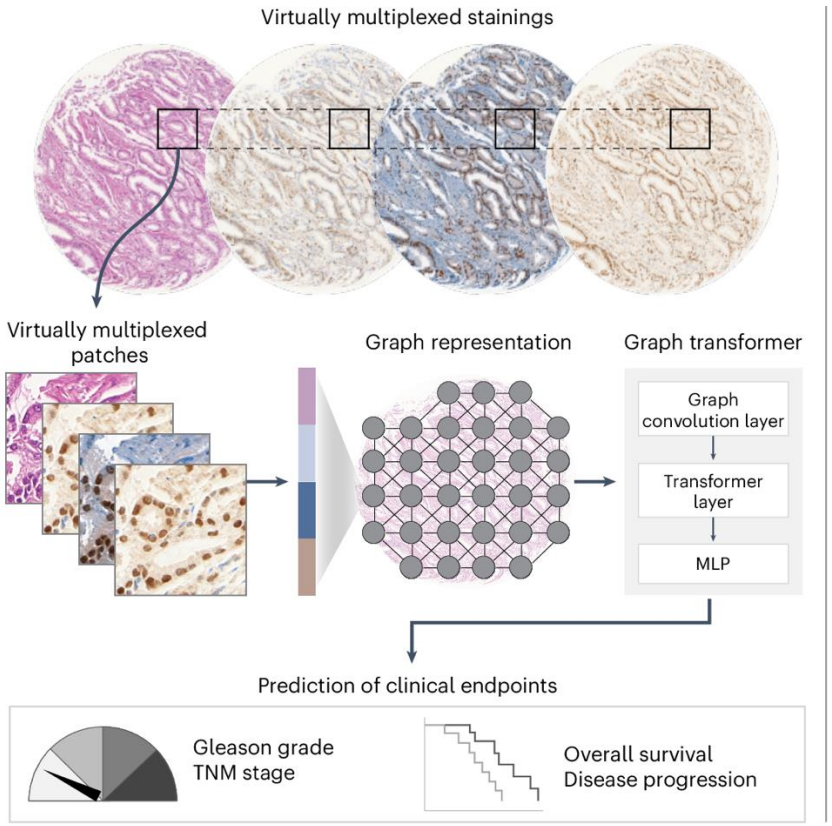

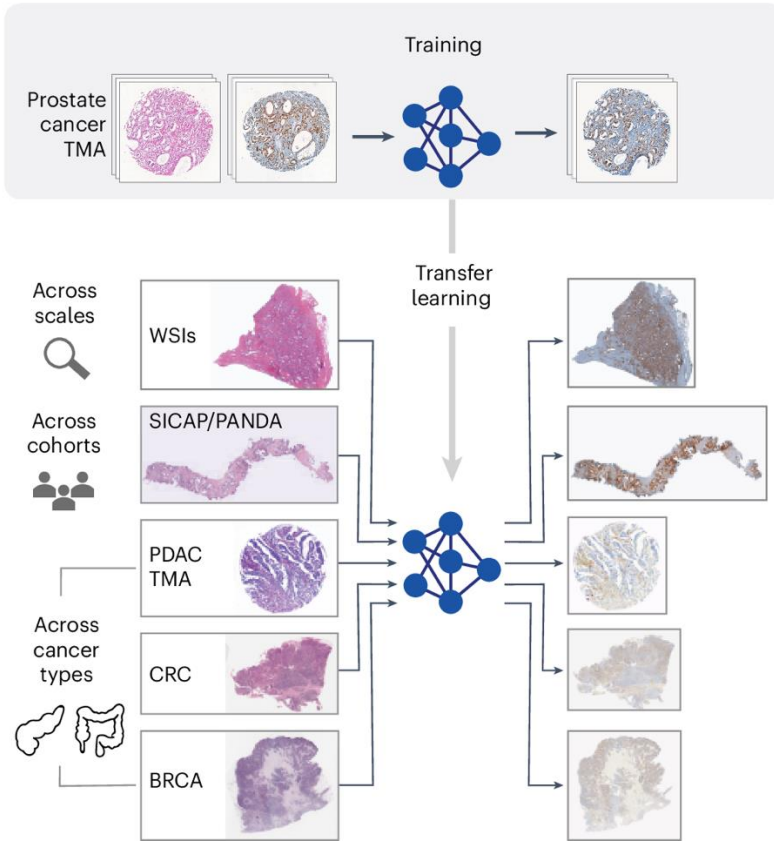

- ✓ Consistently **improved the prediction** of clinically relevant endpoints
- ✓ Able to **generalize** to unseen cohorts and cancer types

Towards **AI-assisted histopathology**:

- Data inpainting, sample imputation, harmonizing datasets, **experimental design**
- Translation to cutting-edge **spatial omics** technologies (e.g., spatial transcriptomics)

Pati *et al.*, *Nature Machine Intelligence*, **6**, 1077–1093 (2024)

An end-to-end AI framework for Inflammatory Skin Diseases

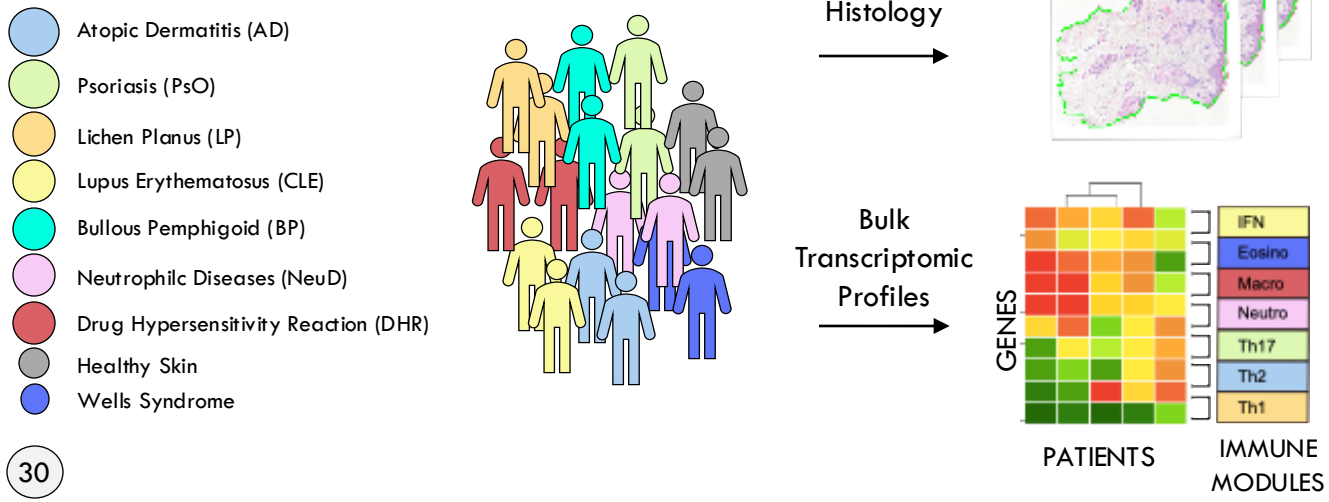

? Can we extract H&E representations that capture tissue heterogeneity?

? Can we use them to train AI models to predict molecular characteristics or diagnostic labels?

? How interpretable and clinically understandable are the representations learned by these models?

Melissa Ensmenger

Collaboration with Prof. Raphael Gottardo (BDSC) and Prof. Michel Gilliet (Department of Dermatology, CHUV)

Team members: Antoine Girardin, Jeremy Di Domizio, Hugo Cometto

An end-to-end AI framework for Inflammatory Skin Diseases

① Segment and patch WSI H&E

② Extract patch embeddings with pre-trained foundation models

③ Train MIL aggregator + head for prediction tasks

④ Explain prediction by visualizing patch attention scores

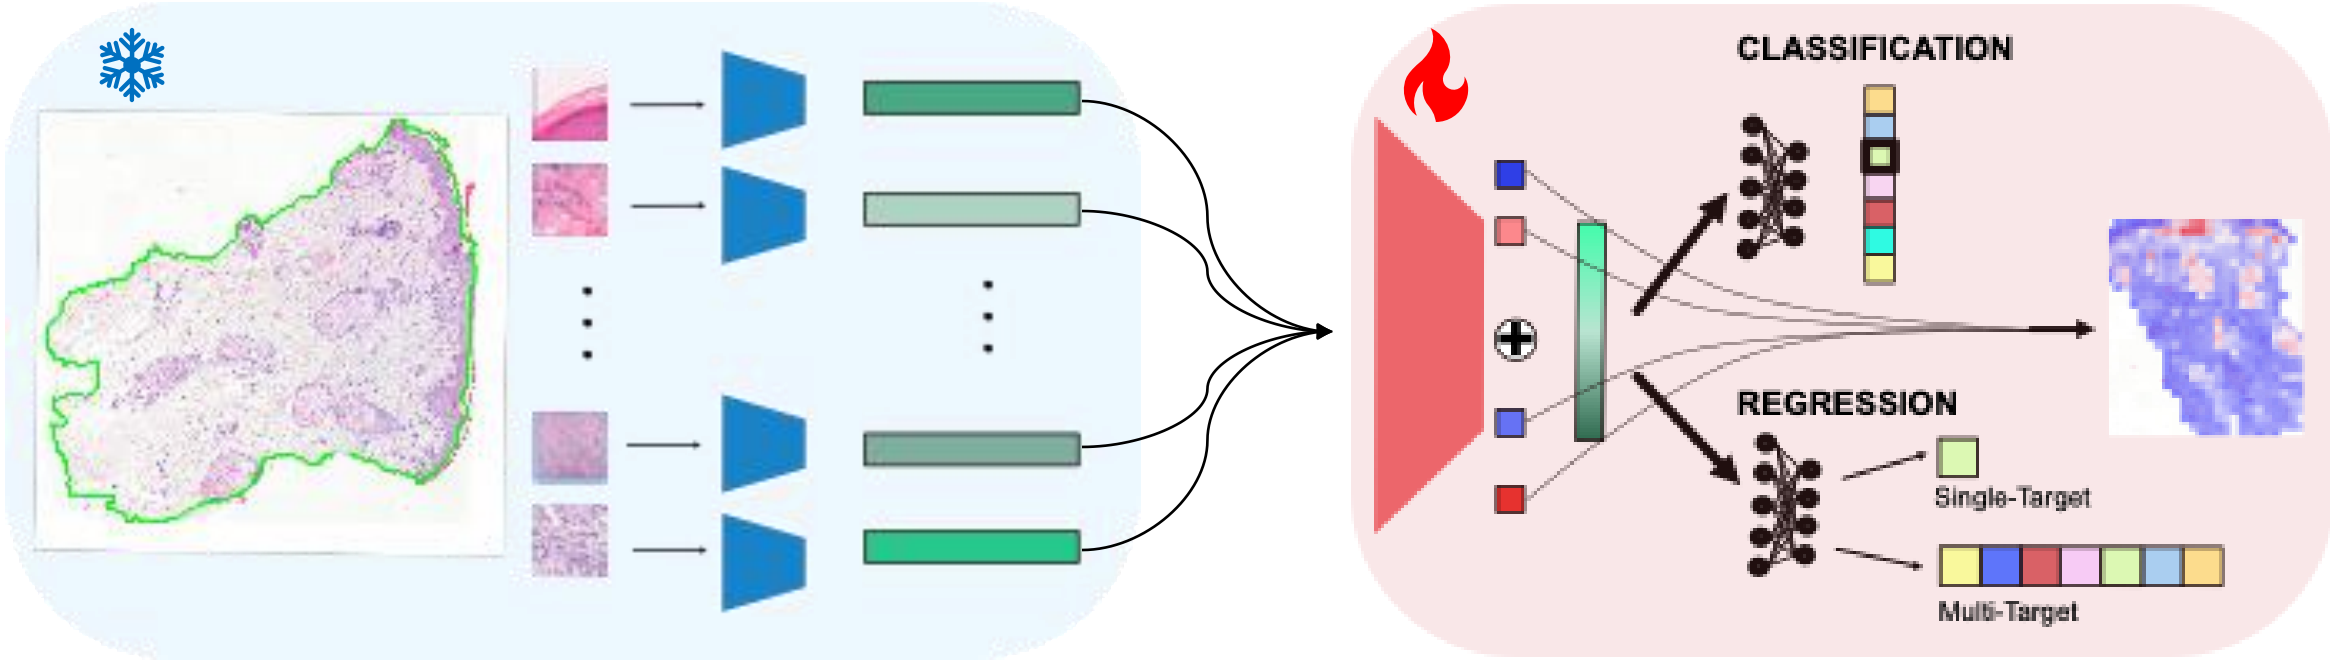

Patch-level Foundation Model
Hoptimus-0[6] – UNI-v2[4] – Virchow-2[8] – CONCH-v1.5[5] – Hibou-L[7]

Patch embeddings

Frozen weights

MIL Aggregator
Attention-based MIL module (ABMIL)

Prediction Head
2-layer MLP

Trainable parameters

Patient embedding

Attention Scores
= weights for patch aggregation

An end-to-end AI framework for Inflammatory Skin Diseases

① Segment and patch WSI H&E

② Extract patch embeddings with pre-trained foundation models

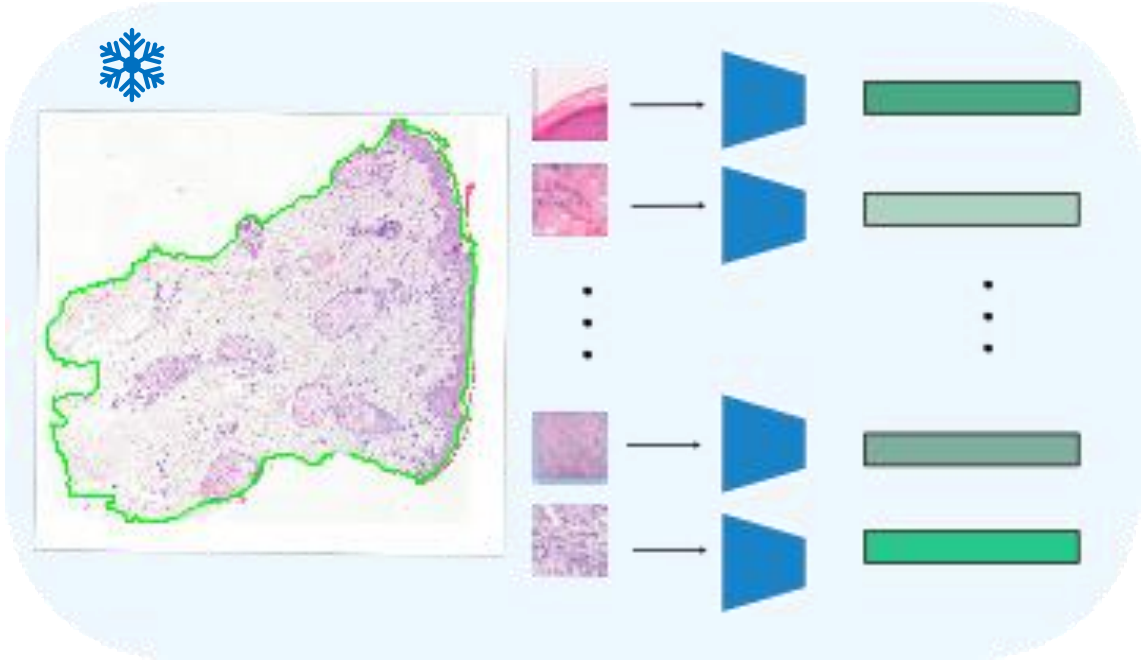

Patch-level Foundation Model

Hoptimus-0[6] – UNI-v2[4] – Virchow-2[8] – CONCH-v1.5[5] – Hibou-L[7]

Patch embeddings

Frozen weights

FM representations capture different layers of the skin as well as specialized morphological structures and technical artifacts.

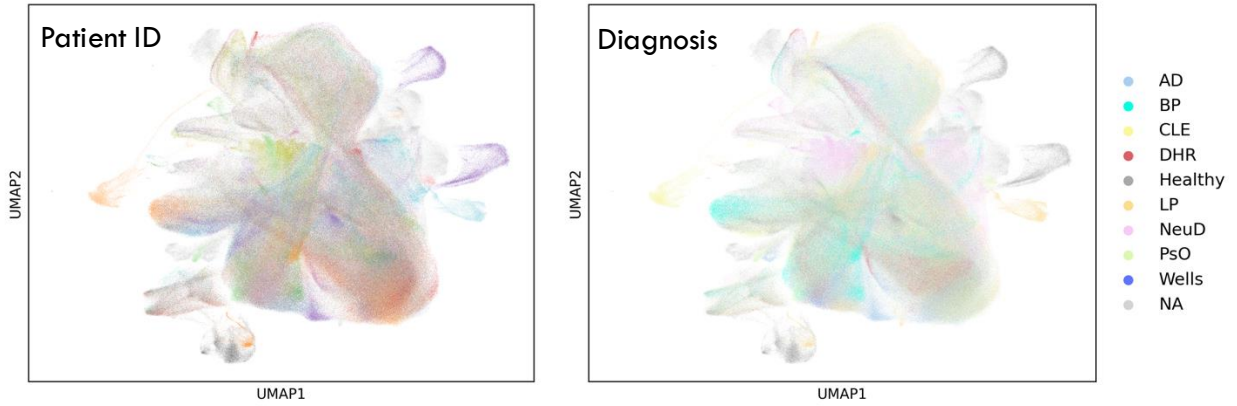

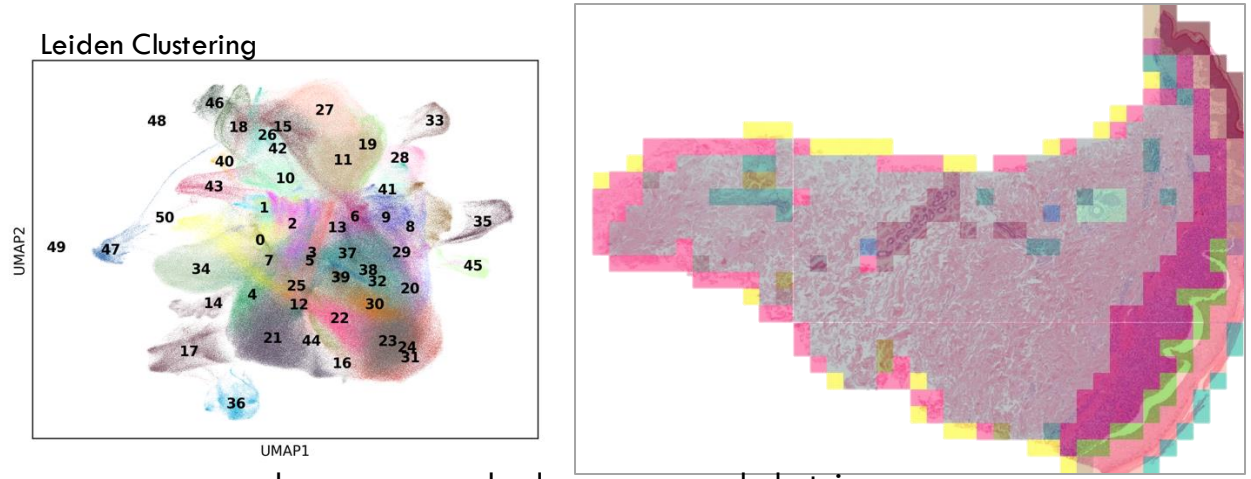

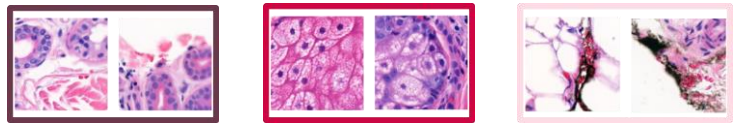

An end-to-end AI framework for Inflammatory Skin Diseases

① Segment and patch WSI H&E

② Extract patch embeddings with pre-trained foundation models

③ Train MIL aggregator + head for prediction tasks

④ Explain prediction by visualizing patch attention scores

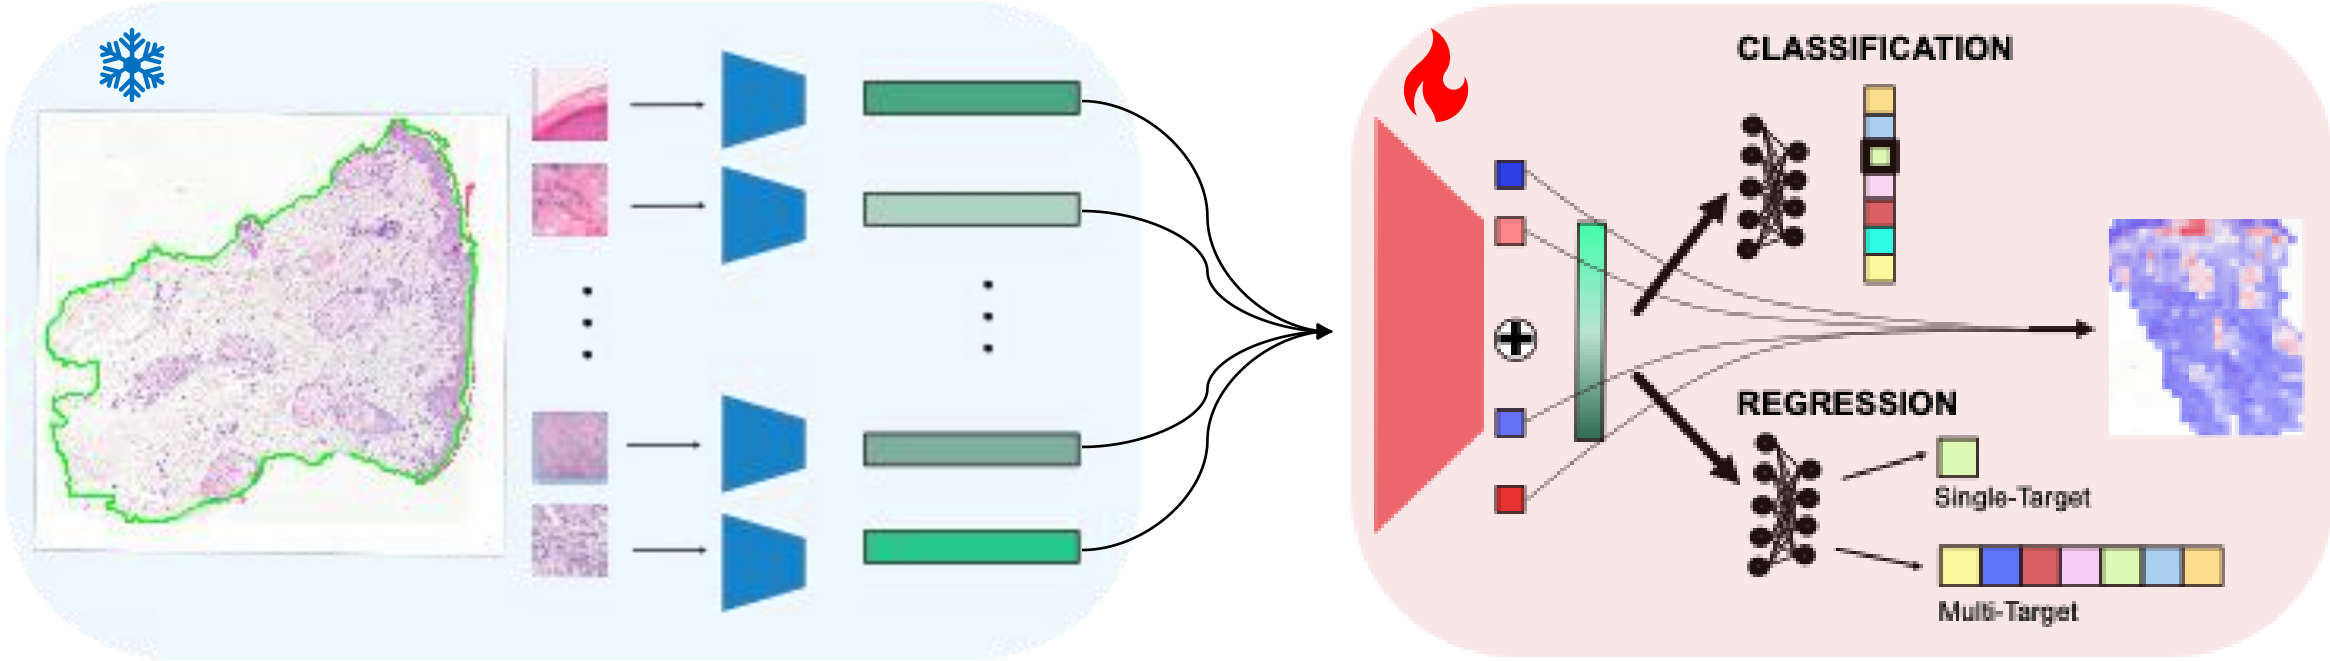

Patch-level Foundation Model
Hoptimus-0[6] – UNI-v2[4] – Virchow-2[8] – CONCH-v1.5[5] – Hibou-L[7]

Patch embeddings

Frozen weights

MIL Aggregator
Attention-based MIL module (ABMIL)

Prediction Head
2-layer MLP

Trainable parameters

Patient embedding

Attention Scores
= weights for patch aggregation

An end-to-end AI framework for Inflammatory Skin Diseases

H&E patient embeddings

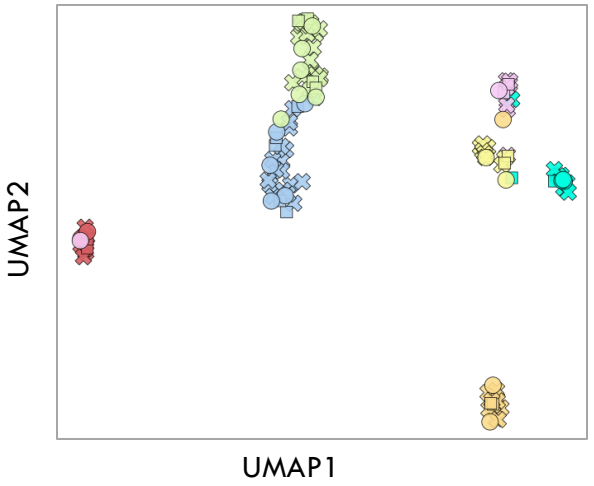

Gene expression patient embeddings

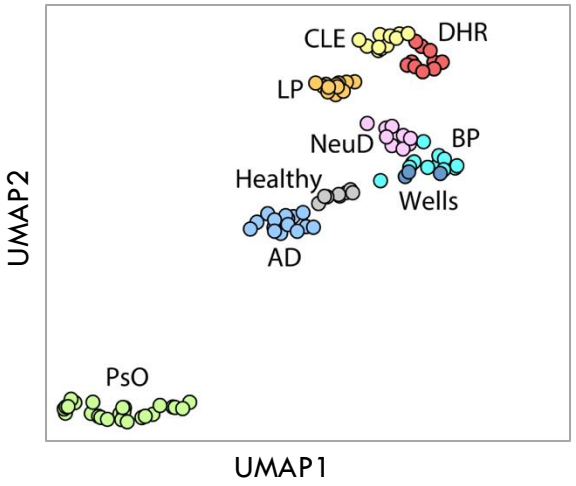

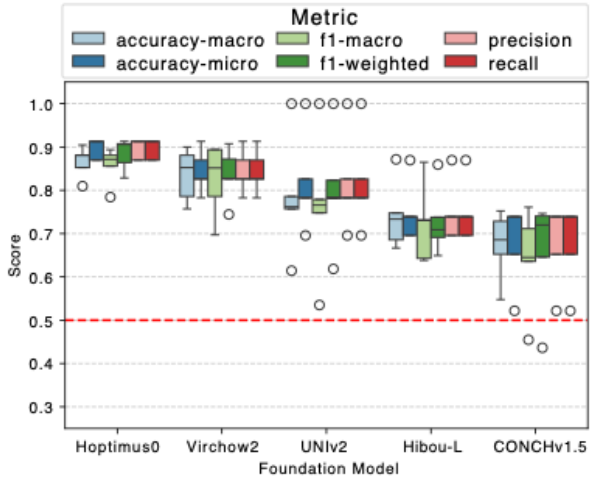

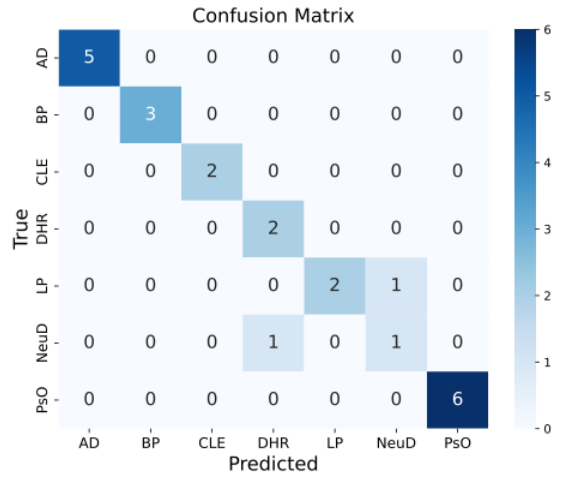

③ Train MIL aggregator + head for prediction tasks

④ Explain prediction by visualizing patch attention scores

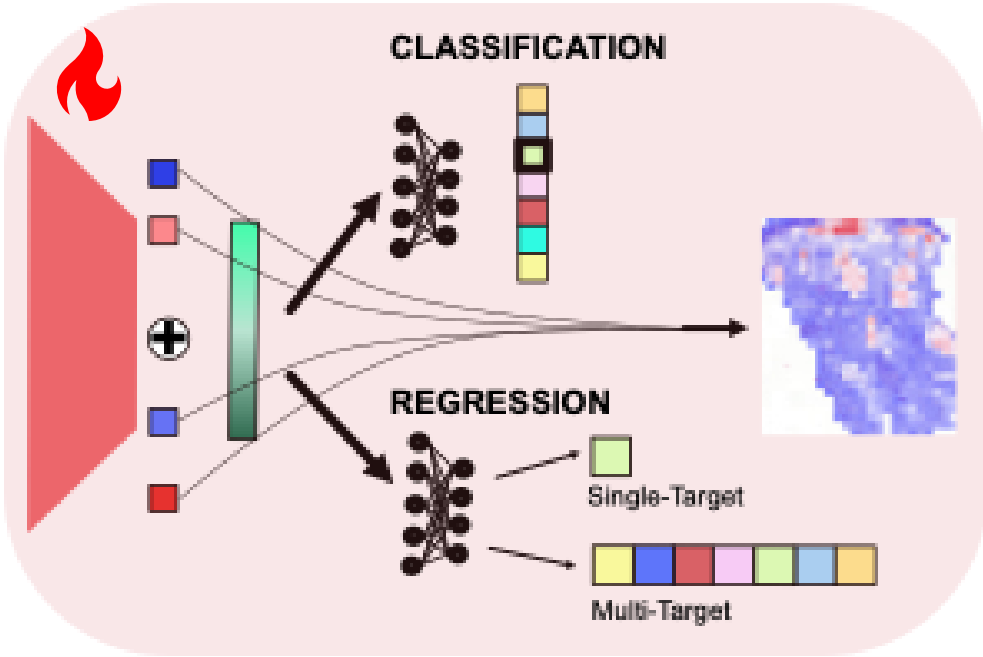

MIL Aggregator
Attention-based MIL module (ABMIL)

Prediction Head
2-layer MLP

Trainable parameters

Patient embedding

Attention Scores
= weights for patch aggregation

An end-to-end AI framework for Inflammatory Skin Diseases

Our framework can predict joint immune pathway activity from tissue morphology (H&E)…

Comparison of **5-fold cross-validation Spearman correlations** for pathway activity prediction between **multi-target** and **single-target models***
*trained on Hoptimus0 patch embeddings

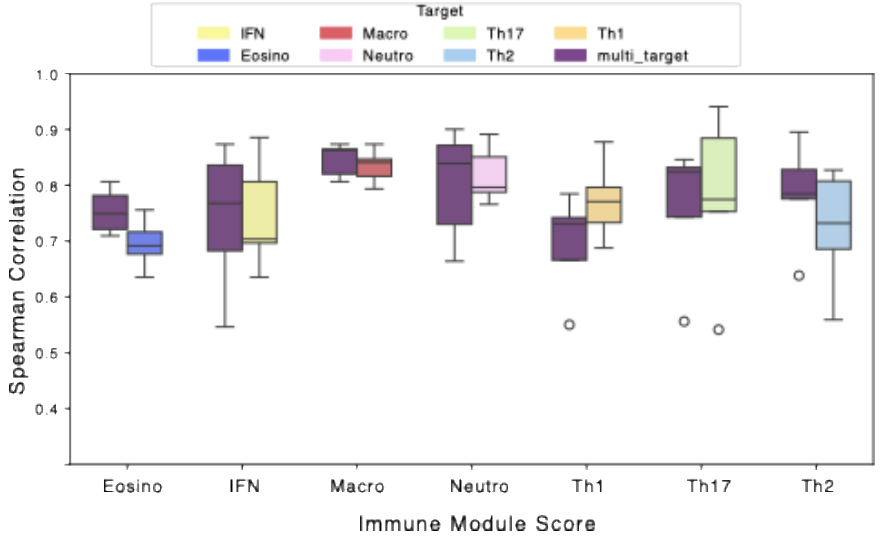

Predicted (y-axis) vs. **reference** (x-axis) pathway activity for the **multi-objective model** (example test set of specific split)

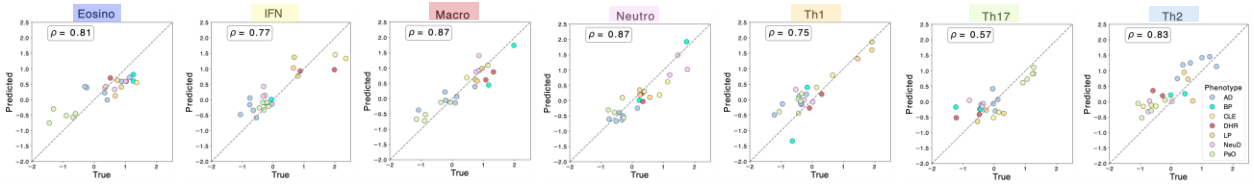

③ Train MIL aggregator + head for prediction tasks

④ Explain prediction by visualizing patch attention scores

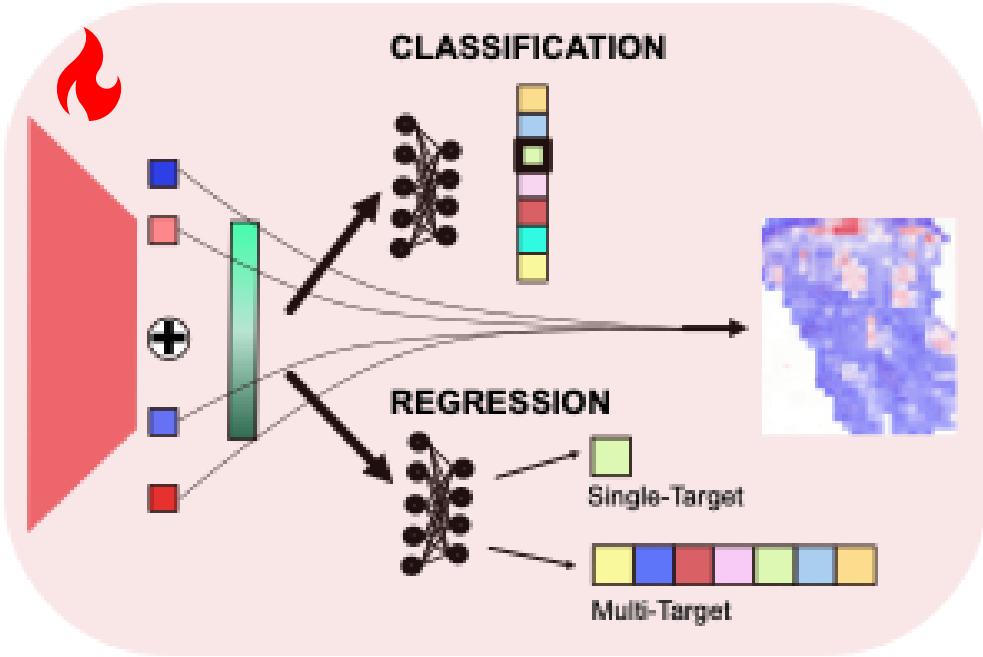

MIL Aggregator
Attention-based MIL module (ABMIL)

Prediction Head
2-layer MLP

Trainable parameters

Patient embedding

Attention Scores
= weights for patch aggregation

An end-to-end AI framework for Inflammatory Skin Diseases

... and map it back to specific tissue regions using attention scores.

Patch attention scores as a proxy for pathway activity (prediction)

Pathway activity measured by **spatial transcriptomics** (Visium HD; consecutive slide) (ground truth)

Patient 1 (CLE)

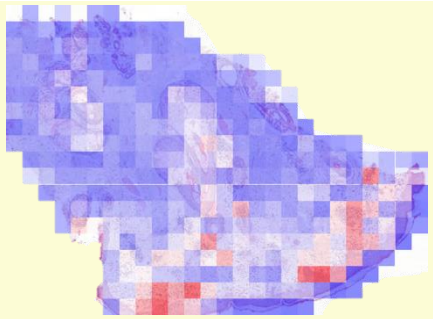

Predicted IFN-activity : 0.79

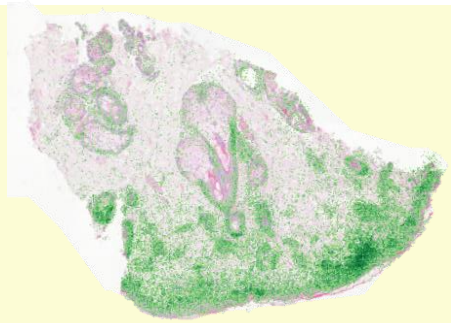

Measured IFN-activity : 1.34

Patient 2 (LP)

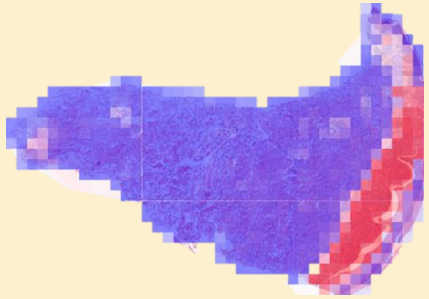

Predicted Th1-activity : 1.60

Measured Th1-activity : 1.75

③ Train MIL aggregator + head for prediction tasks

④ Explain prediction by visualizing patch attention scores

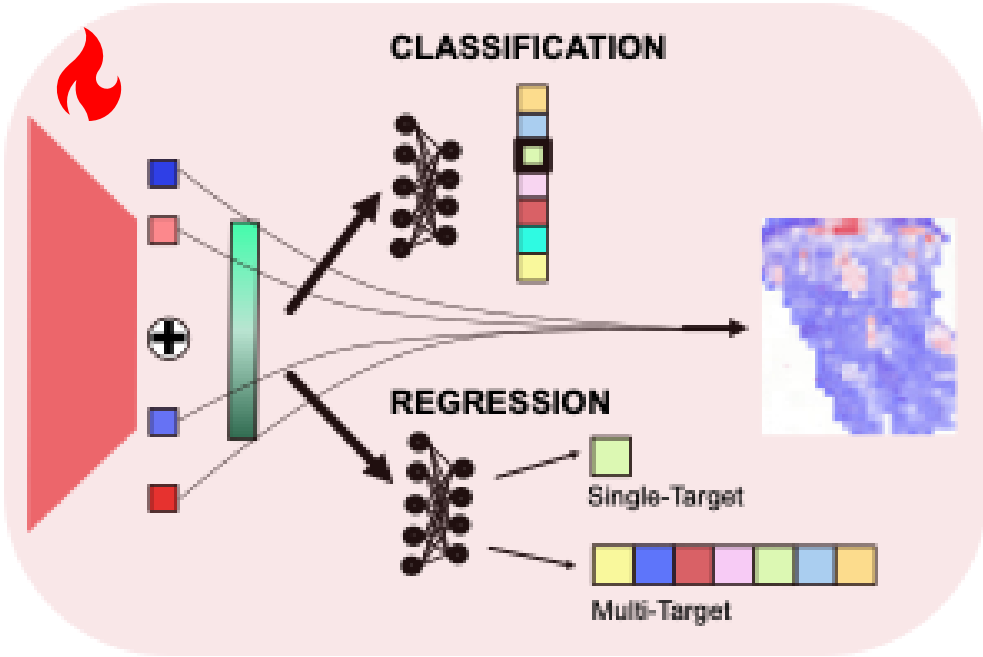

MIL Aggregator
Attention-based MIL module (ABMIL)

Patient embedding

Prediction Head
2-layer MLP

Attention Scores
= weights for patch aggregation

Trainable parameters

Represent

Histopathology

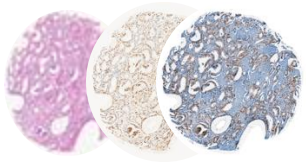

Spatial omics

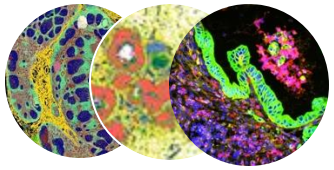

Clinical data

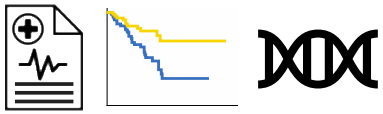

Interpret

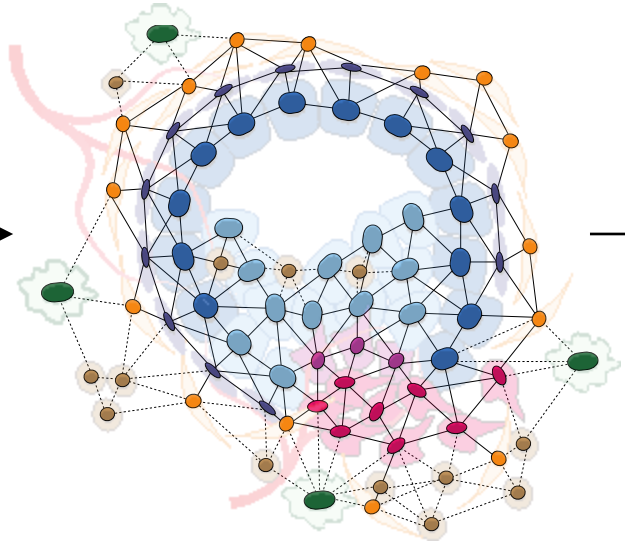

Perturb

Perturbations

Biological discovery

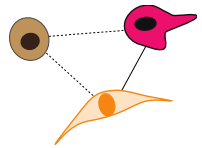

Treatment selection

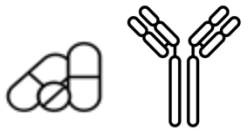

How can we learn multimodal **representations** from spatial omics?

How can we **interpret** these representations of the TME?

How can we predict the effects of drug **perturbations**?

Modeling the TME with **Graph Representation Learning**

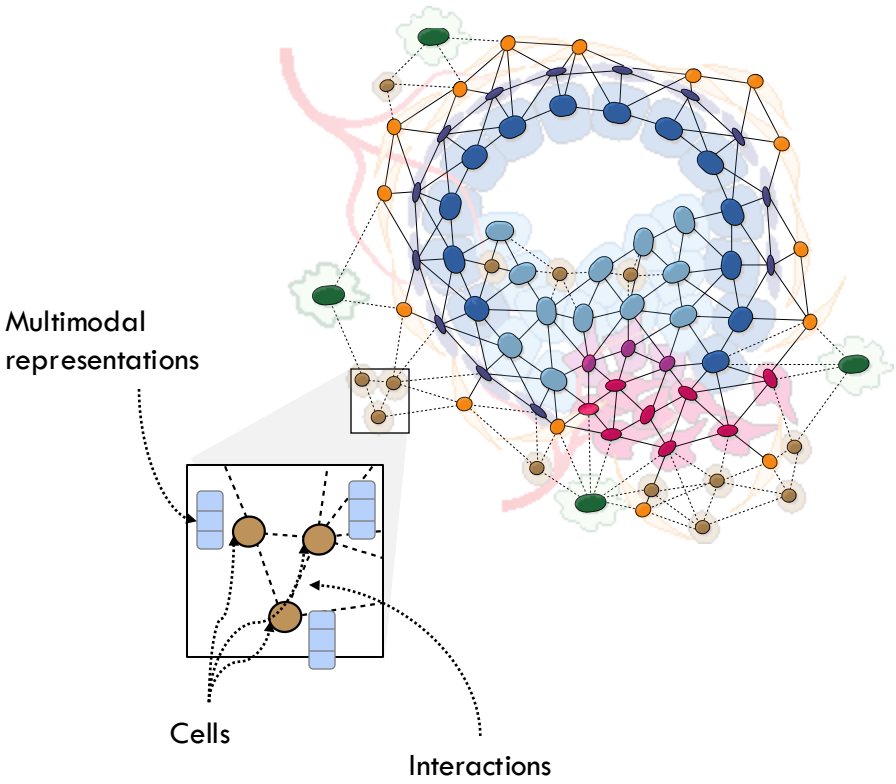

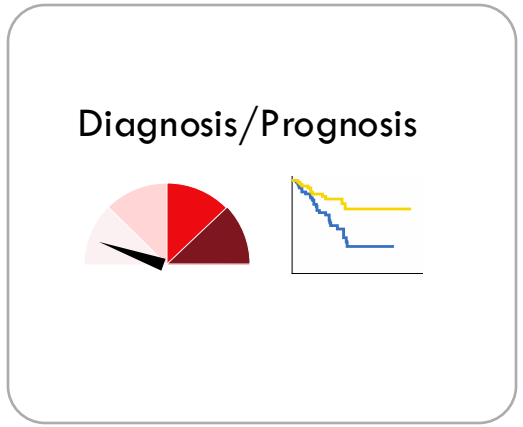

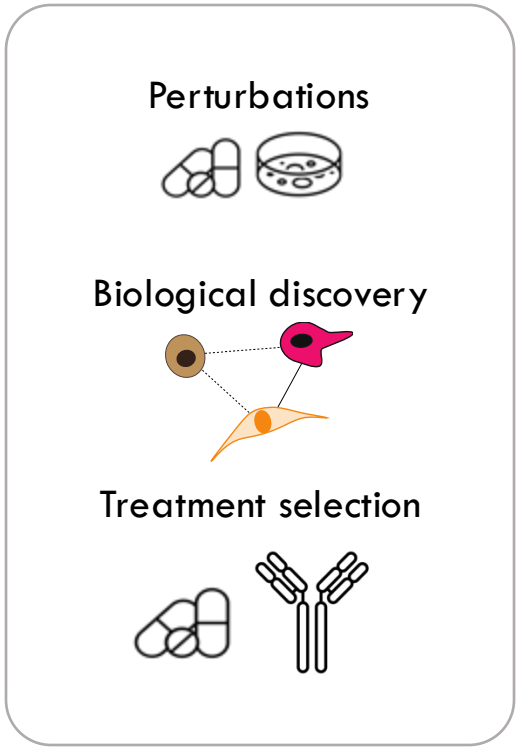

Modeling the TME with **Graph Representation Learning**

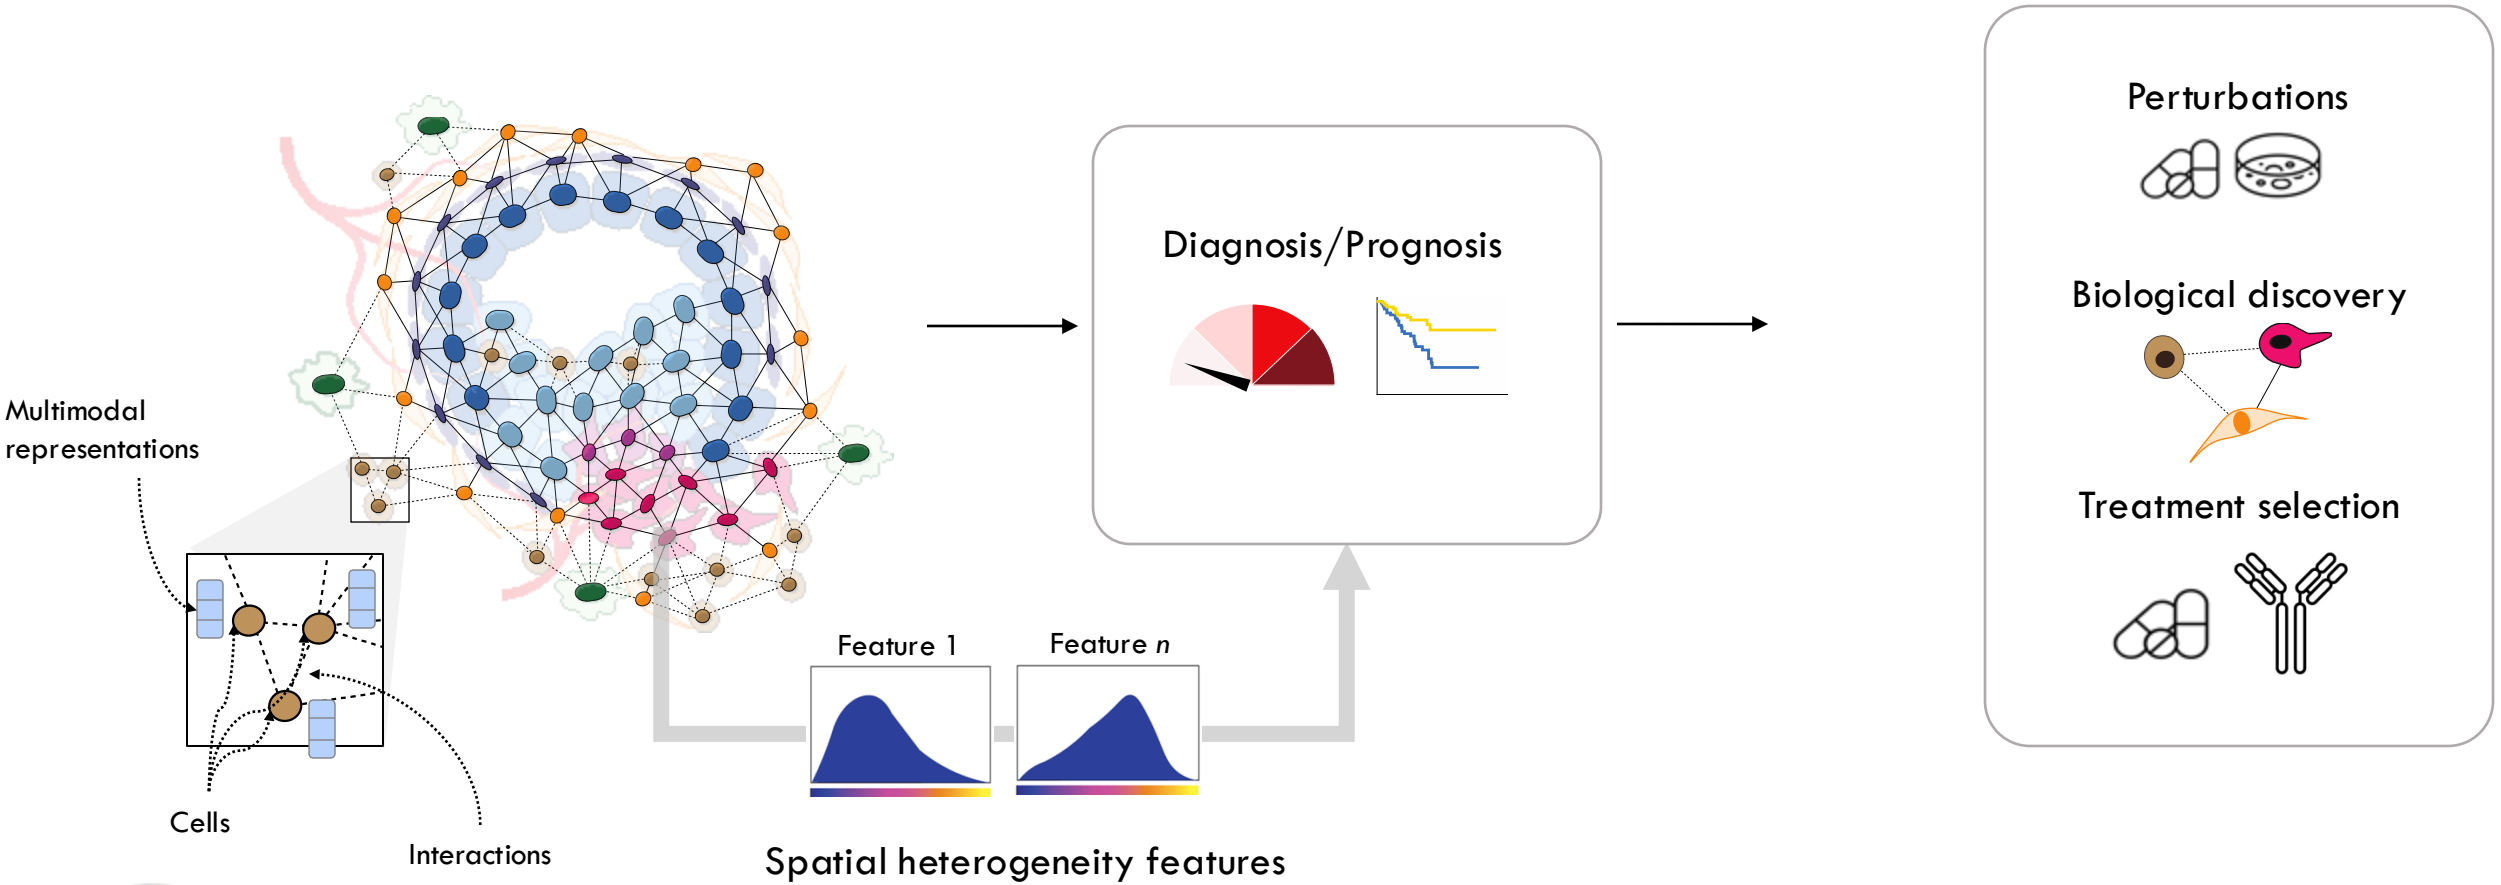

Adriano Martinelli

ATHENA

Martinelli and Rapsomaniki, *Bioinformatics* (2022)

https://github.com/AI4SCR/ATHENA

Modeling the TME with **Graph Representation Learning**

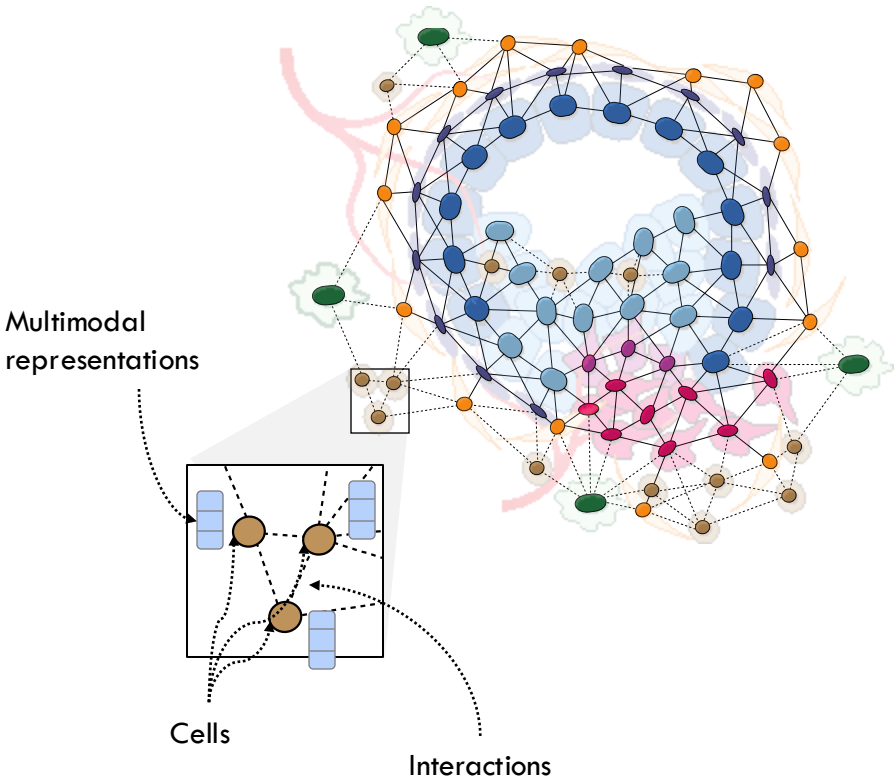

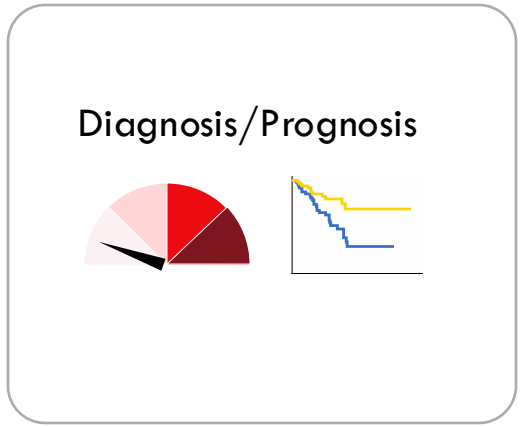

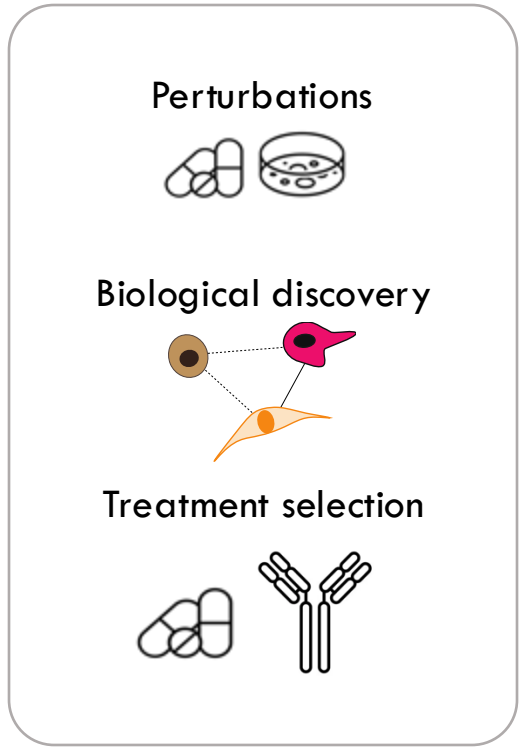

Can we build explainable Graph ML models that learn diectly from data?

Post-hoc explainers are not consistent

Publicly available NSCLC IMC dataset (Cords et al., 2024).

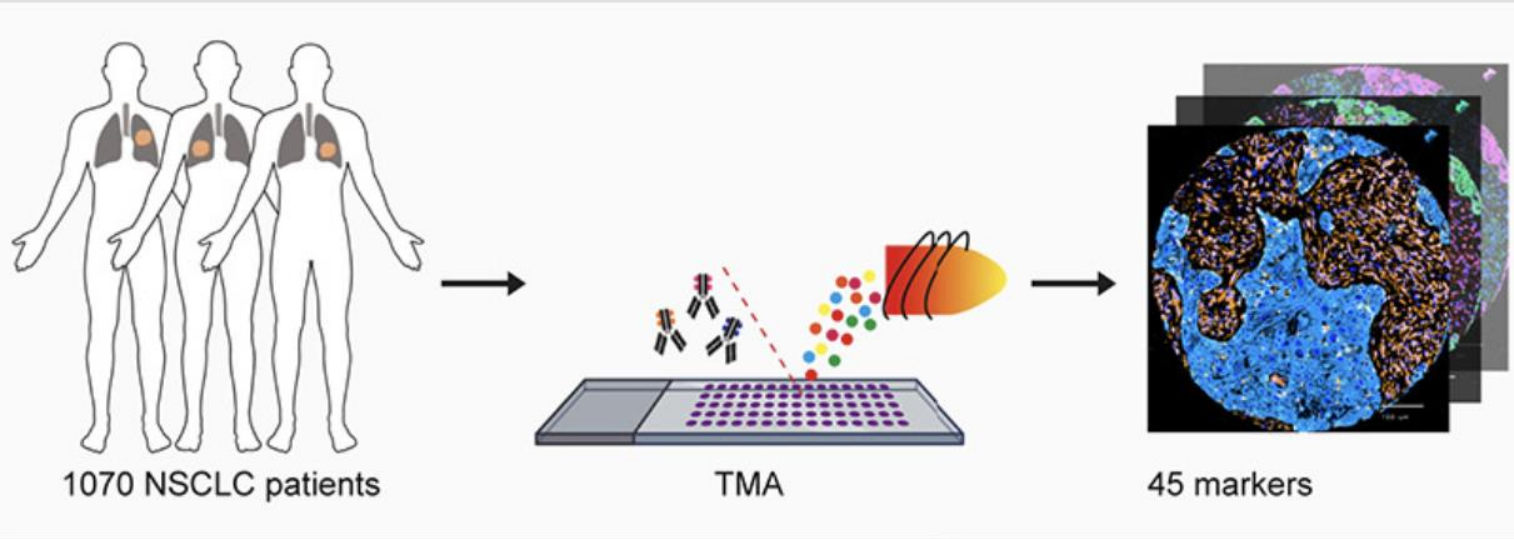

58% LUAD, 36% LUSC

Predict LUAD vs. LUSC

	metric	
Model	F1 weighted	F1 macro
MLP	0.641 ± 0.005	0.635 ± 0.013
GCN	0.715 ± 0.007	0.704 ± 0.004
GAT	$\mathbf{0.760 \pm 0.016}$	$\mathbf{0.752 \pm 0.015}$
GIN	0.723 ± 0.013	0.710 ± 0.029

Theo Maffei

Post-hoc explainers are not consistent

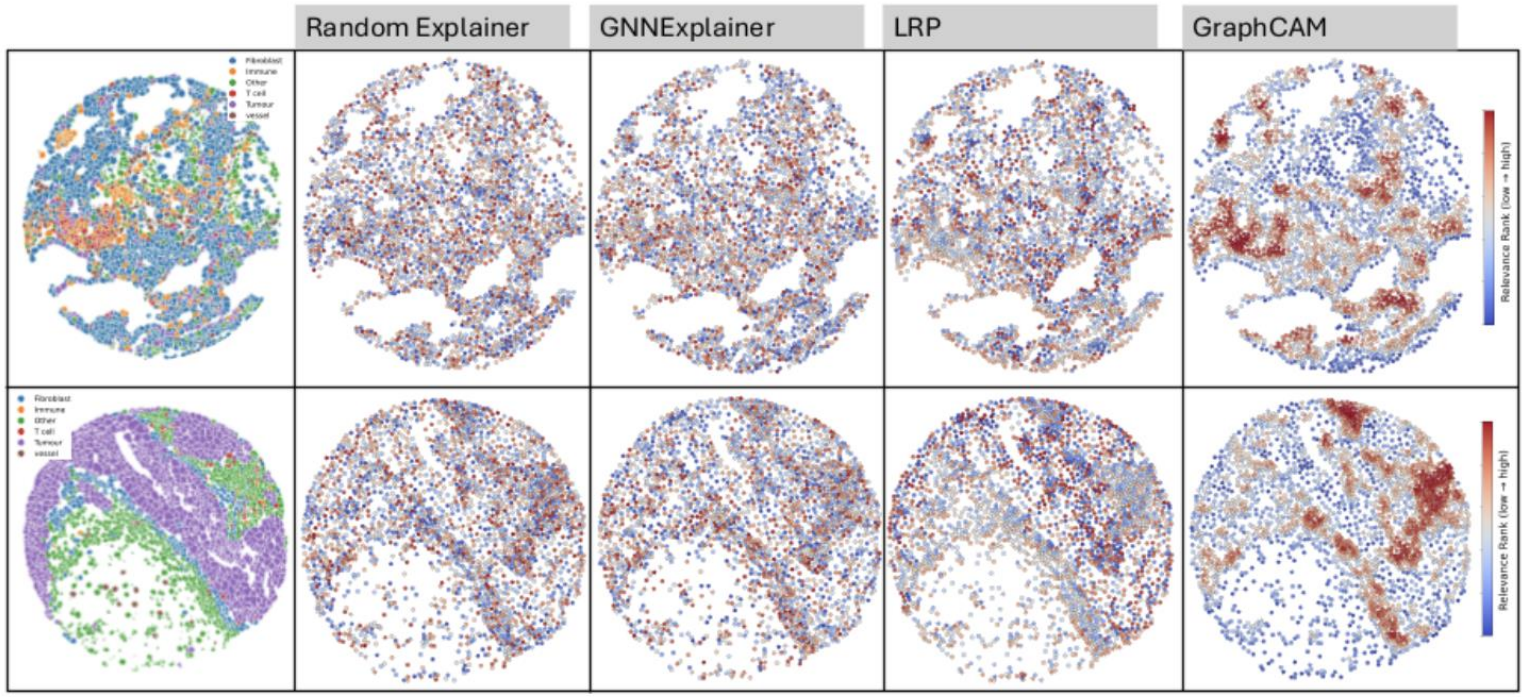

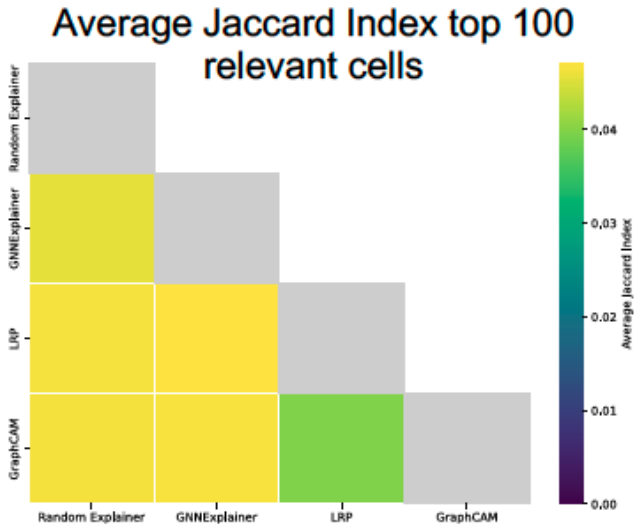

Theo Maffei

Although GNN-based models achieves an F1-weighted score of 0.76, post hoc explainers exhibit limited agreement

ProteinPNet: Prototypical Part Networks for Concept Learning in Spatial Proteomics

Louis McConnell

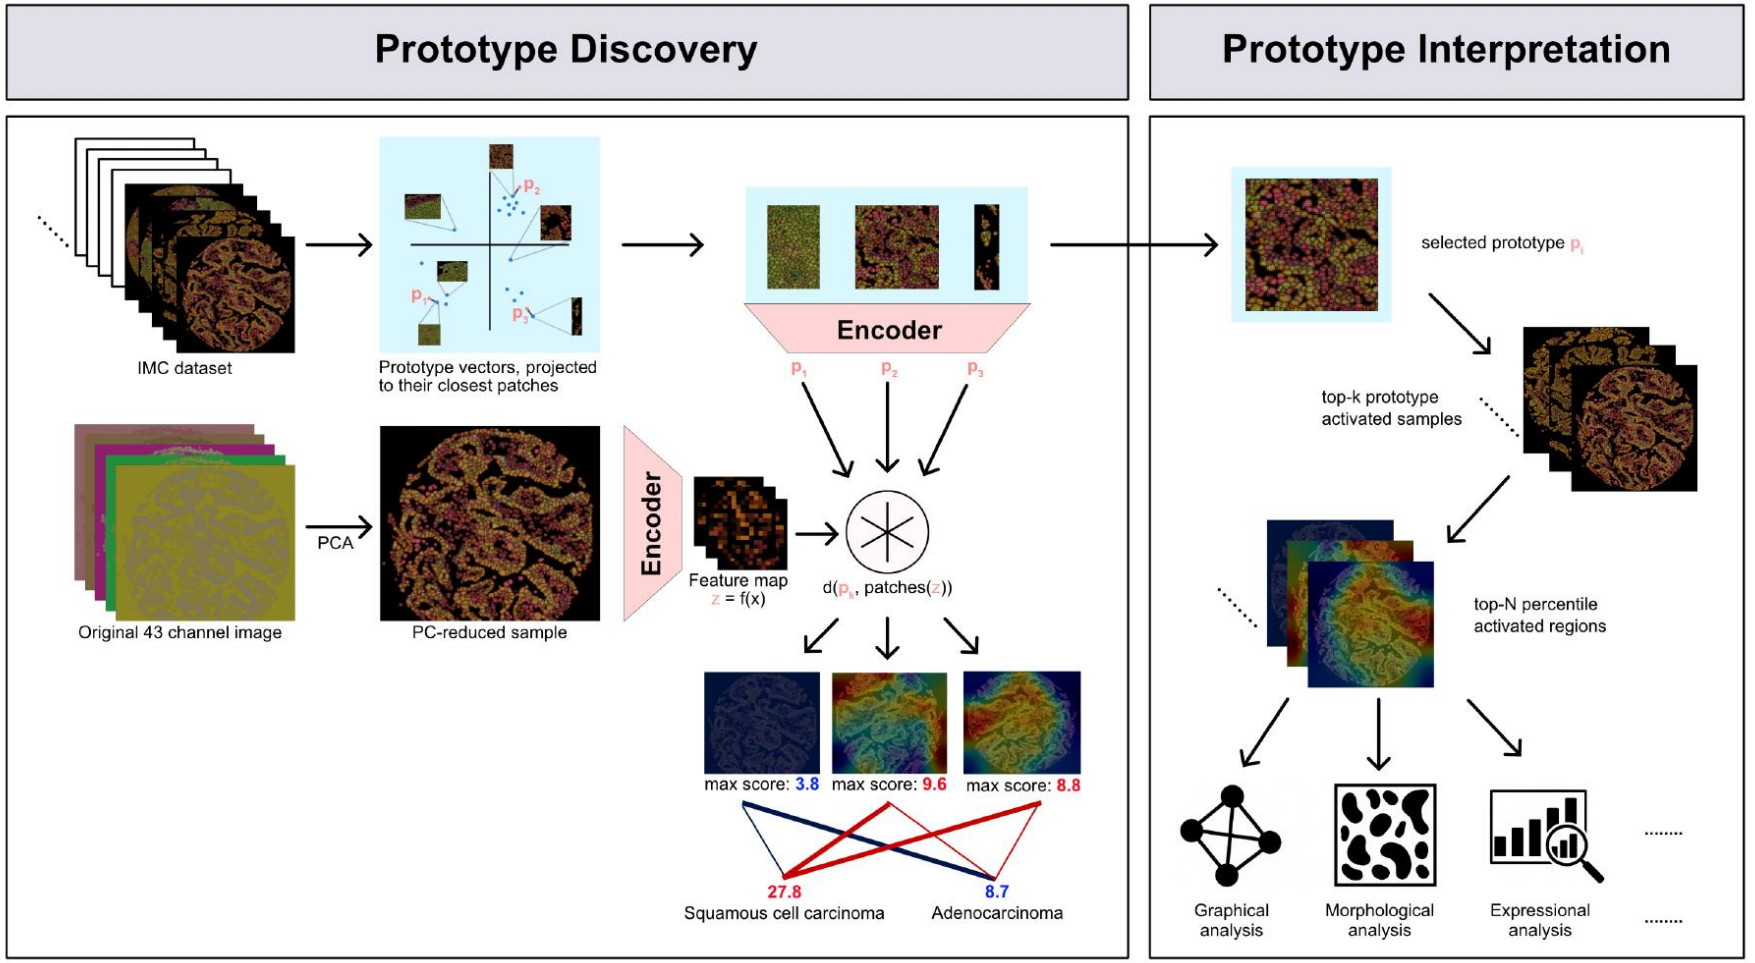

McConnel, Sun, Maffei, Gottardo, Rapsomaniki, Imageomics NeurIPS workshop (to appear)

ProteinPNet: Prototypical Part Networks for Concept Learning in Spatial Proteomics

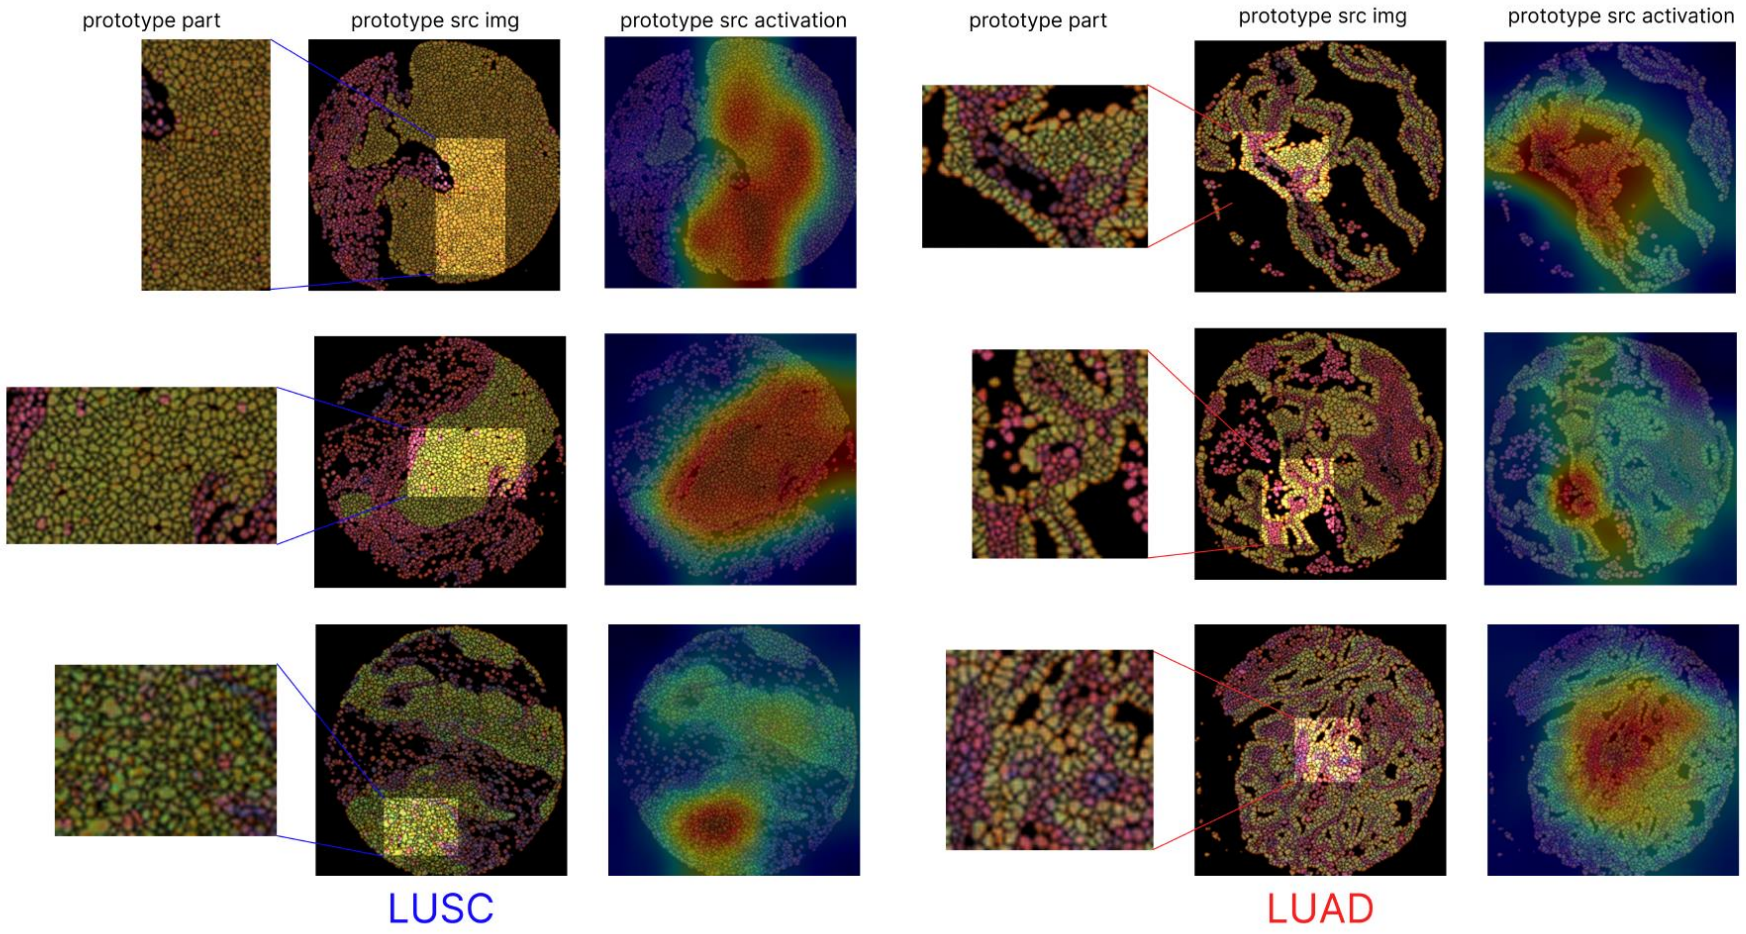

McConnel, Sun, Maffei, Gottardo, Rapsomaniki, Imageomics NeurIPS workshop (to appear)

From TME complexity to interpretable concepts

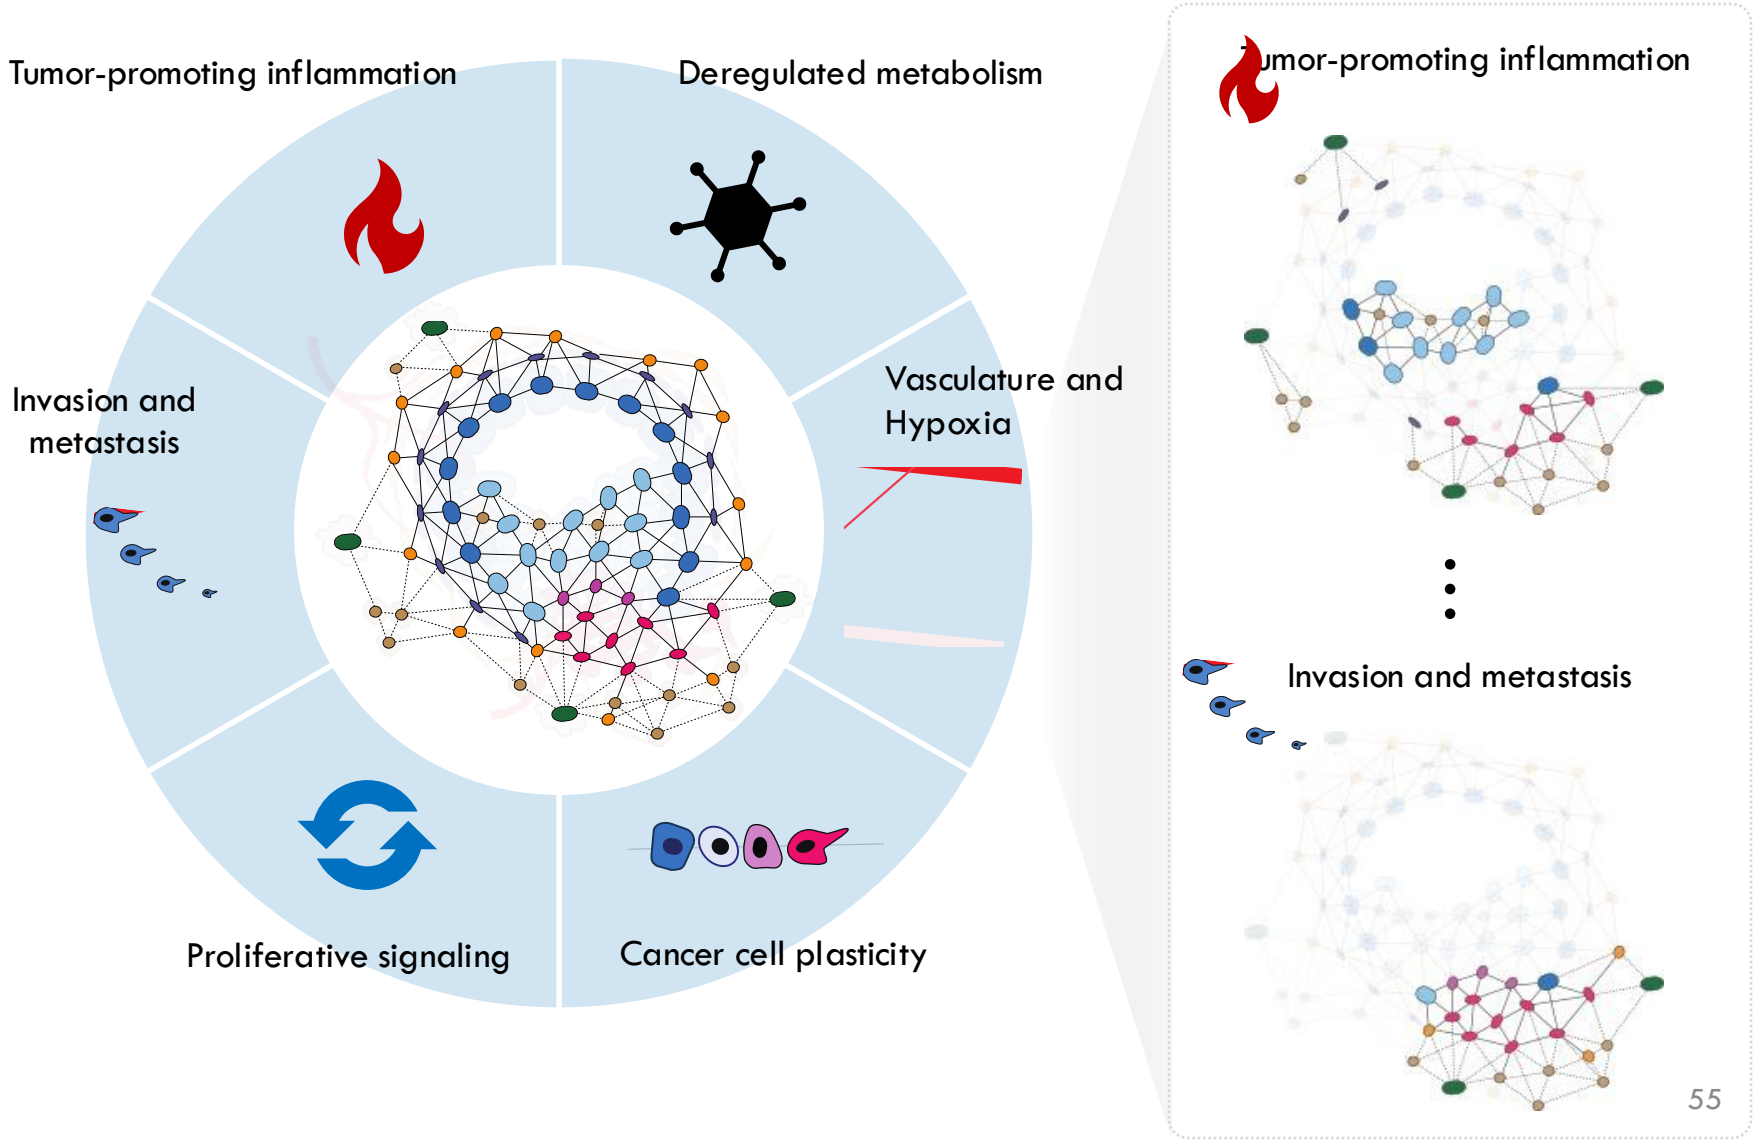

From TME complexity to interpretable concepts

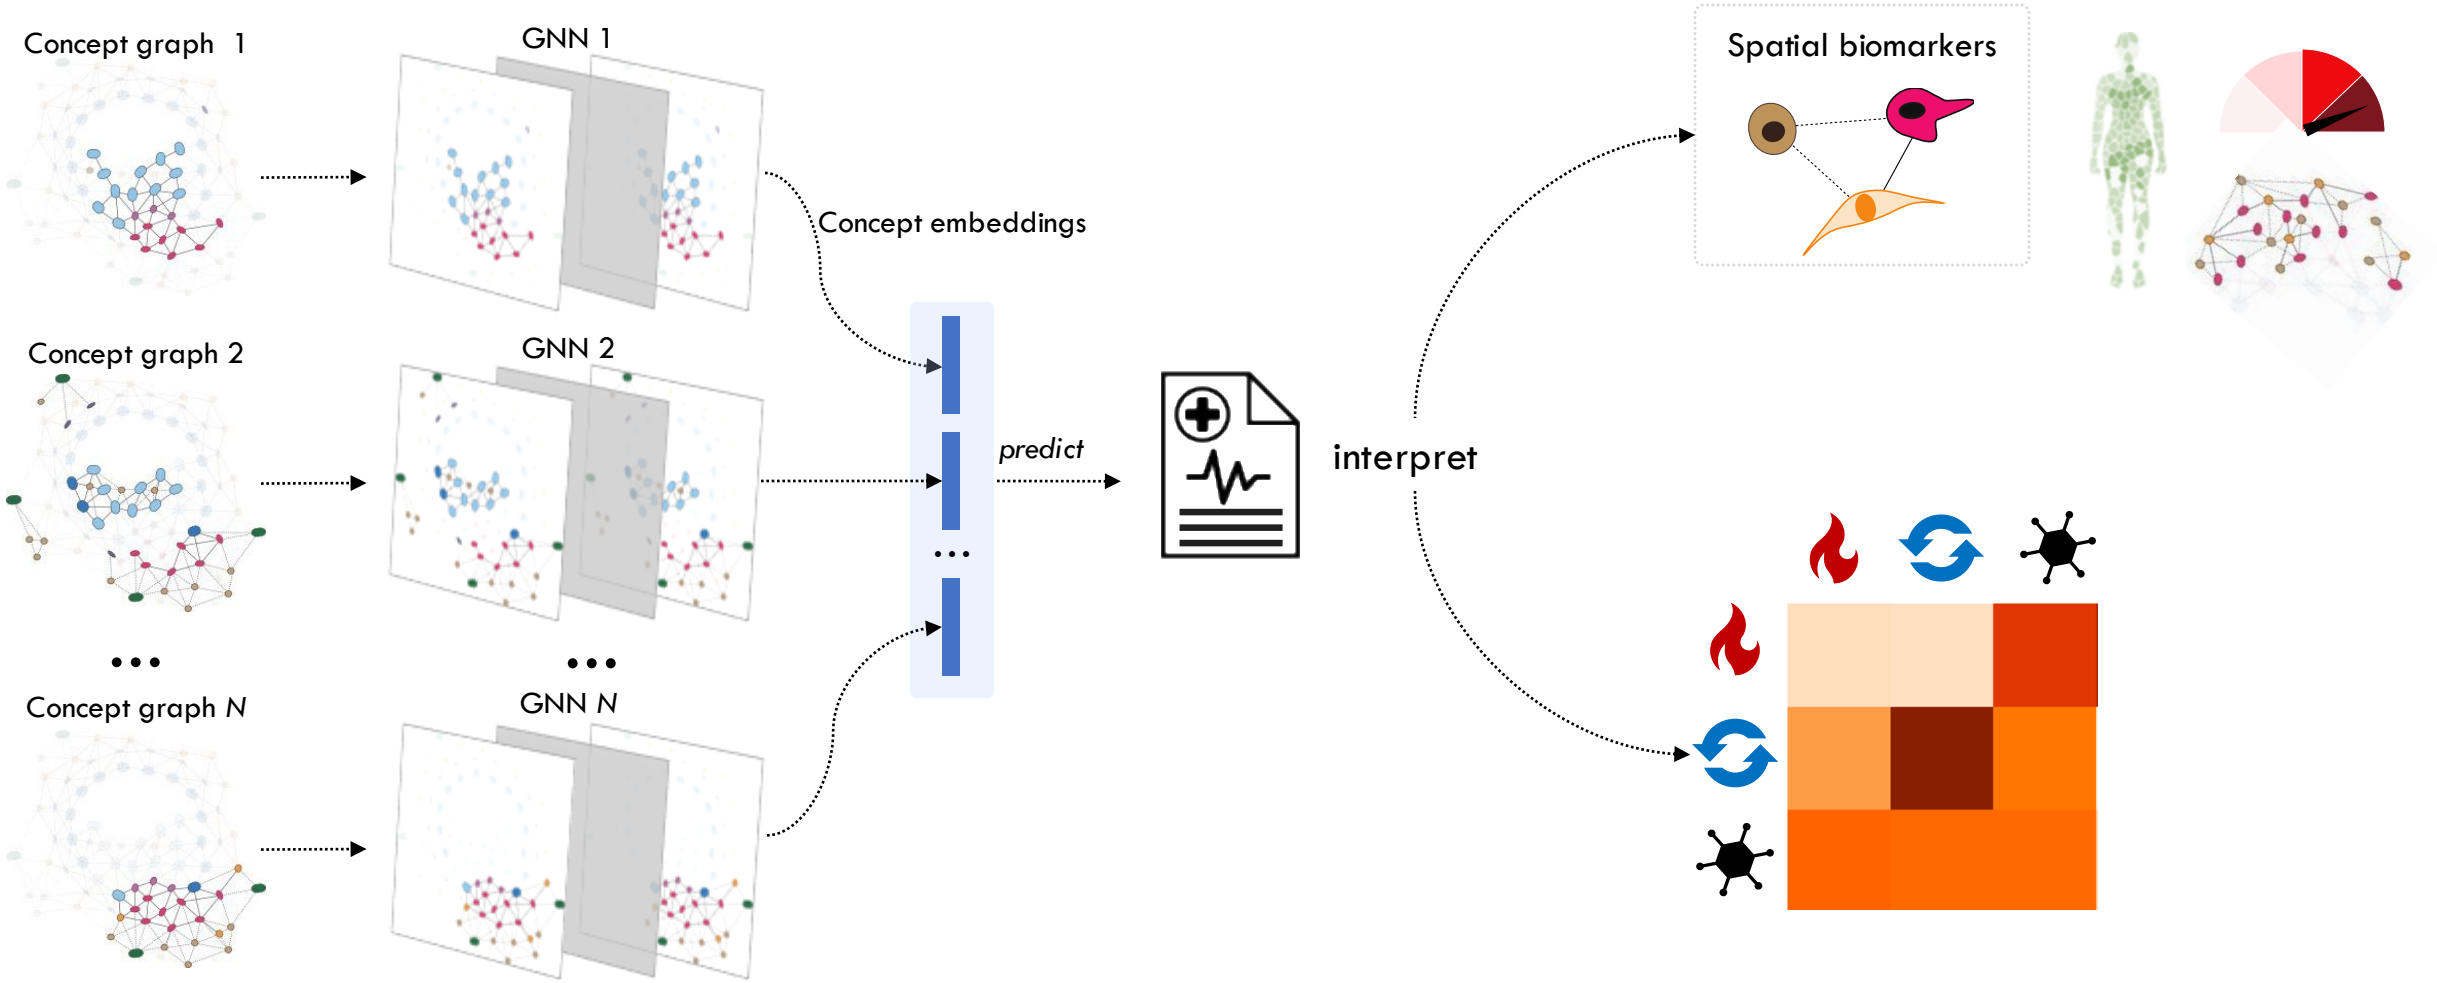

MOSAIC: the world's largest spatial multiomic dataset in oncology

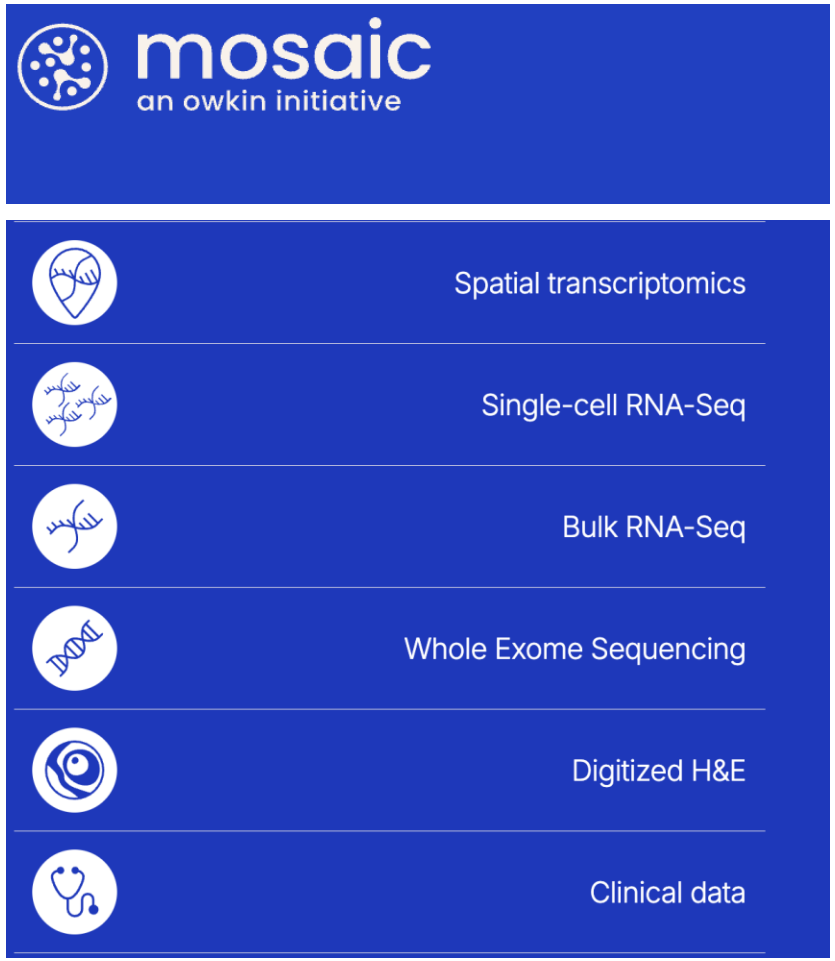

9	Cancer indications	2,055	Total patients included
6	Data modalities	1,833	Spatial omics data generated
100	Clinical variables per patient	1,449	Single-cell omics generated

OWKIN
Owkin uses artificial intelligence (AI) to find the right treatment for every patient.

Uniklinikum Erlangen
Since 1815, Erlangen University hospital has operated at the frontier of healthcare innovation.

CHUV Centre hospitalier universitaire vaudois
The CHUV, since 2019 ranked among the best 15 hospitals worldwide by Newsweek, is one of the five Swiss universities.

GUSTAVE ROUSSY CANCER CAMPUS GRAND PARIS
Leading Cancer Center in Europe, ranked third best oncology hospital in the world, according to Newsweek in 2023.

CHARITÉ UNIVERSITÄTSMEDIZIN BERLIN
The Charité is the medical faculty of the Humboldt University and the Free University of Berlin.

University of Pittsburgh
A top-ranked, public institution in the U.S.

Collaboration with Prof. Raphael Gottardo (BDSC), Dr. Krisztian Homicsko (Dept. of Oncology, CHUV), Prof. Laurence de Laval (Dept. Pathology, CHUV)
Team members: Sari Issa, Theo Maffei, Spencer Watson, Jonathan Bac

Represent

Histopathology

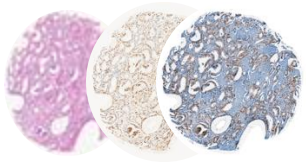

Spatial omics

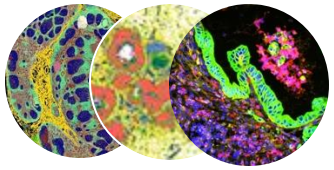

Clinical data

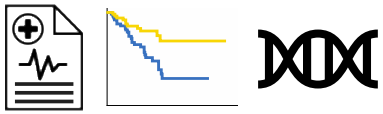

Interpret

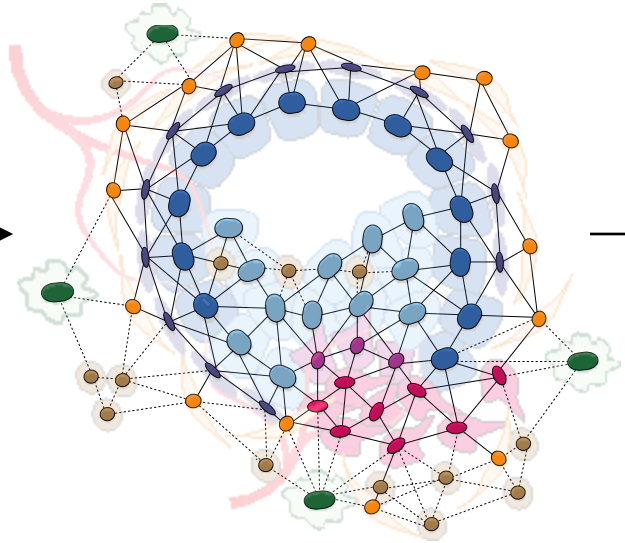

Perturb

Perturbations

Biological discovery

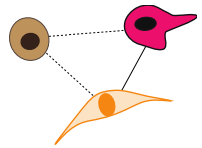

Treatment selection

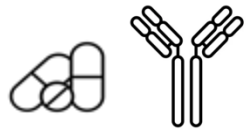

How can we learn multimodal **representations** from spatial omics?

How can we **interpret** these representations of the TME?

How can we predict the effects of drug **perturbations**?

High correlations of drug response metrics between cell lines of distinct origins

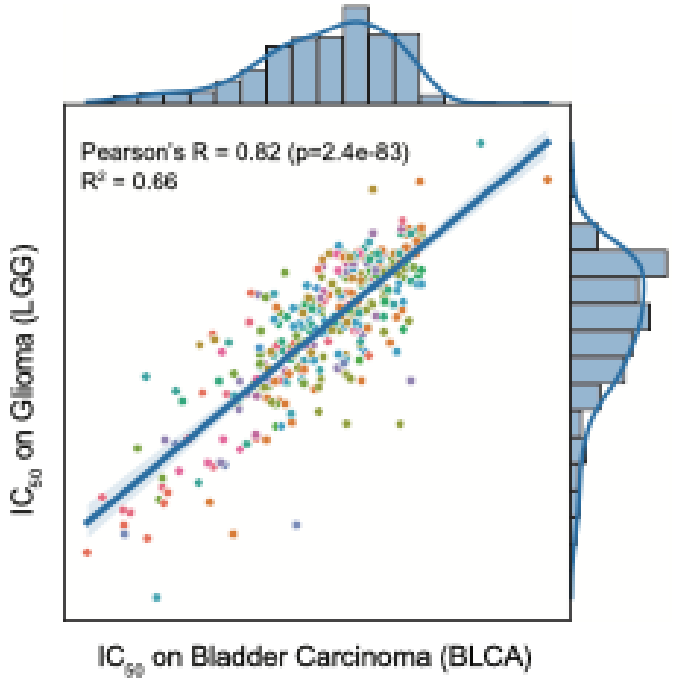

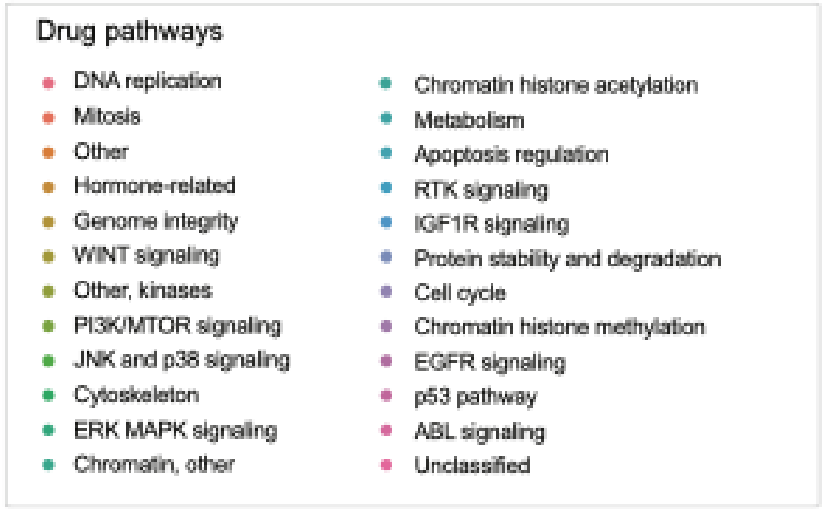

UNIVERSITÄT BERN

Katia Ovchinnikova
Marianna Kruithof-de Julio

IBM **Research** Zurich
Jannis Born

Ovchinnikova *et al.*, npj Precision Oncology, 2024

Drug response is heavily affected by the inherent **potency** or **toxicity** of each drug **independently** of the cell line/organoid it was tested on

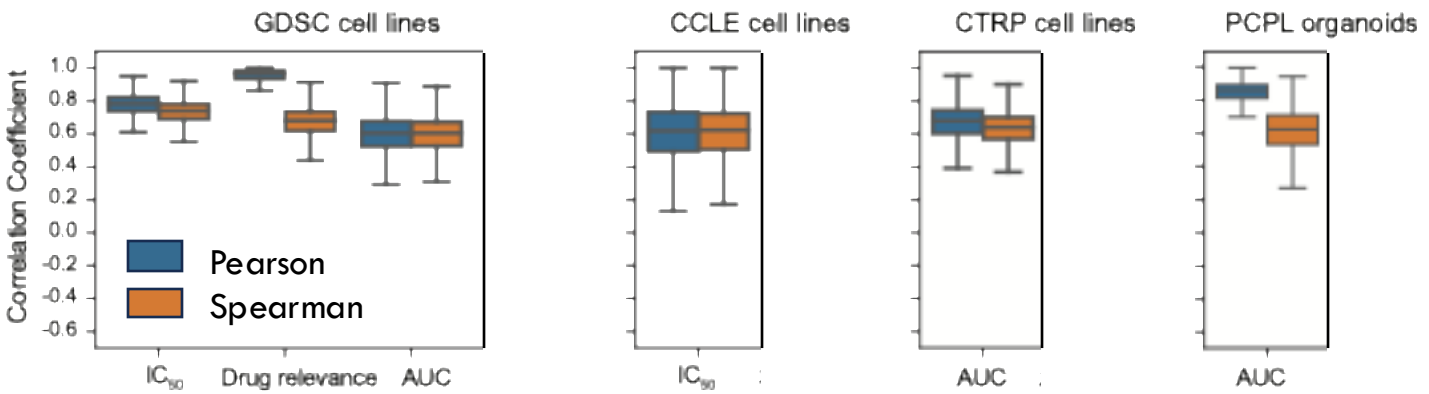

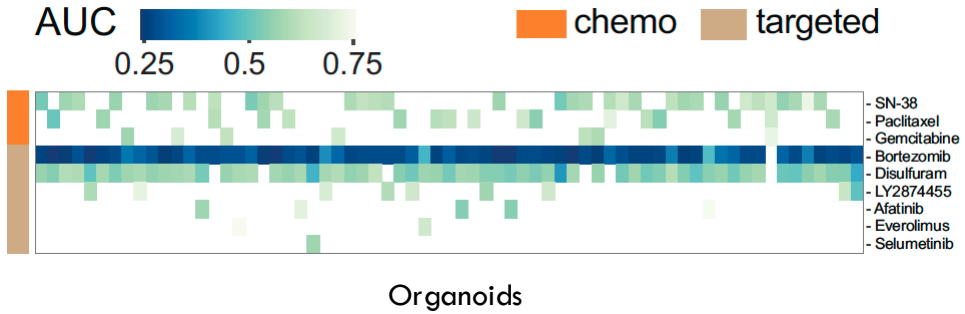

Ovchinnikova *et al.*, npj Precision Oncology, 2024

Z-score normalization for all values separately for each drug removes the drug-specific bias

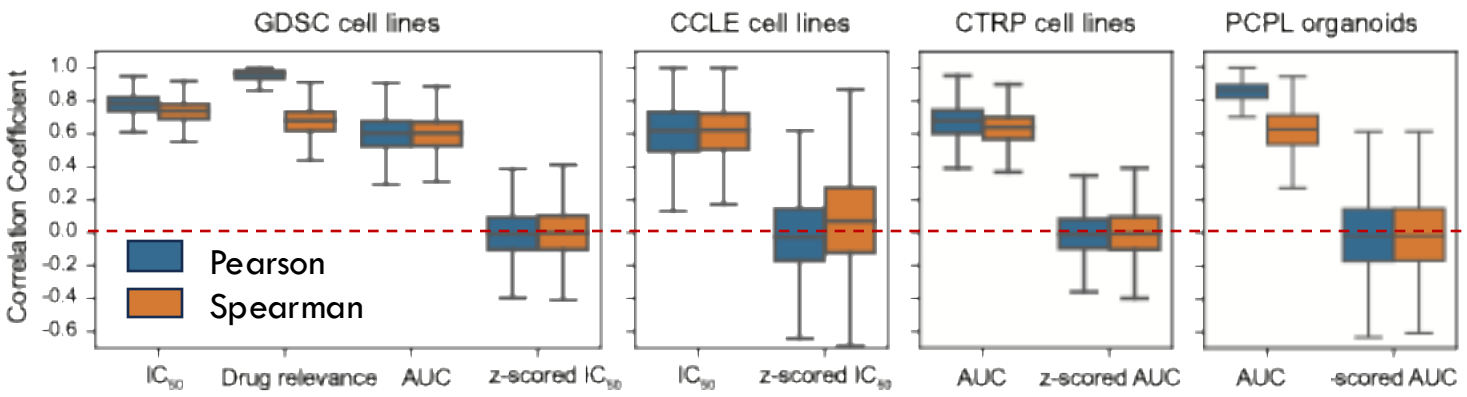

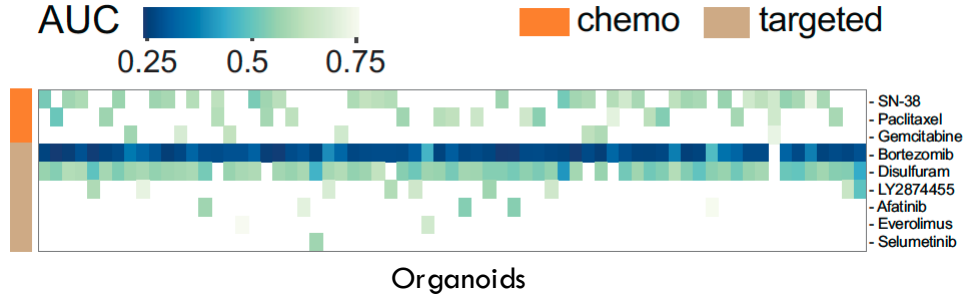

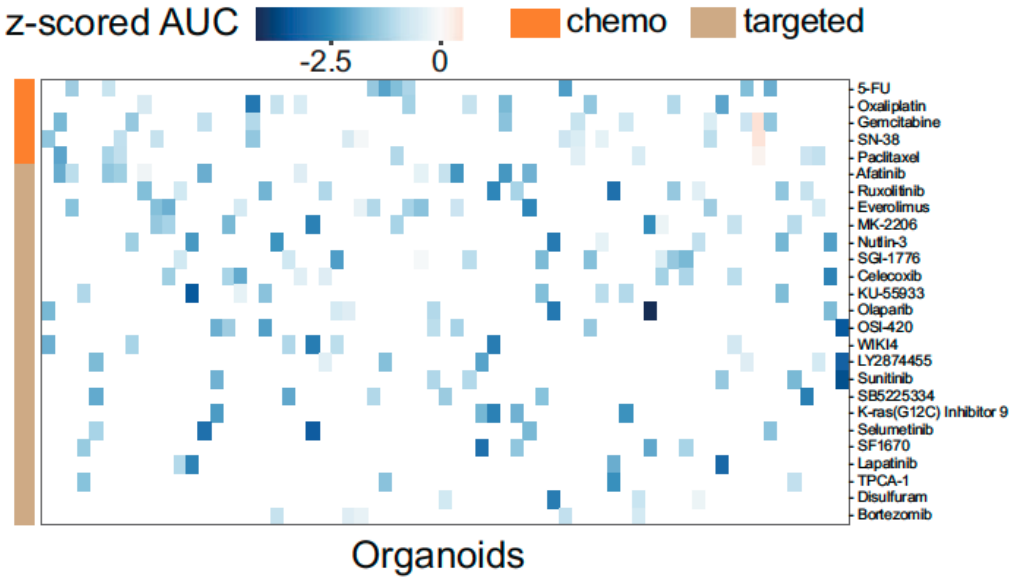

Ovchinnikova *et al.*, npj Precision Oncology, 2024

AI models are heavily affected by these issues!

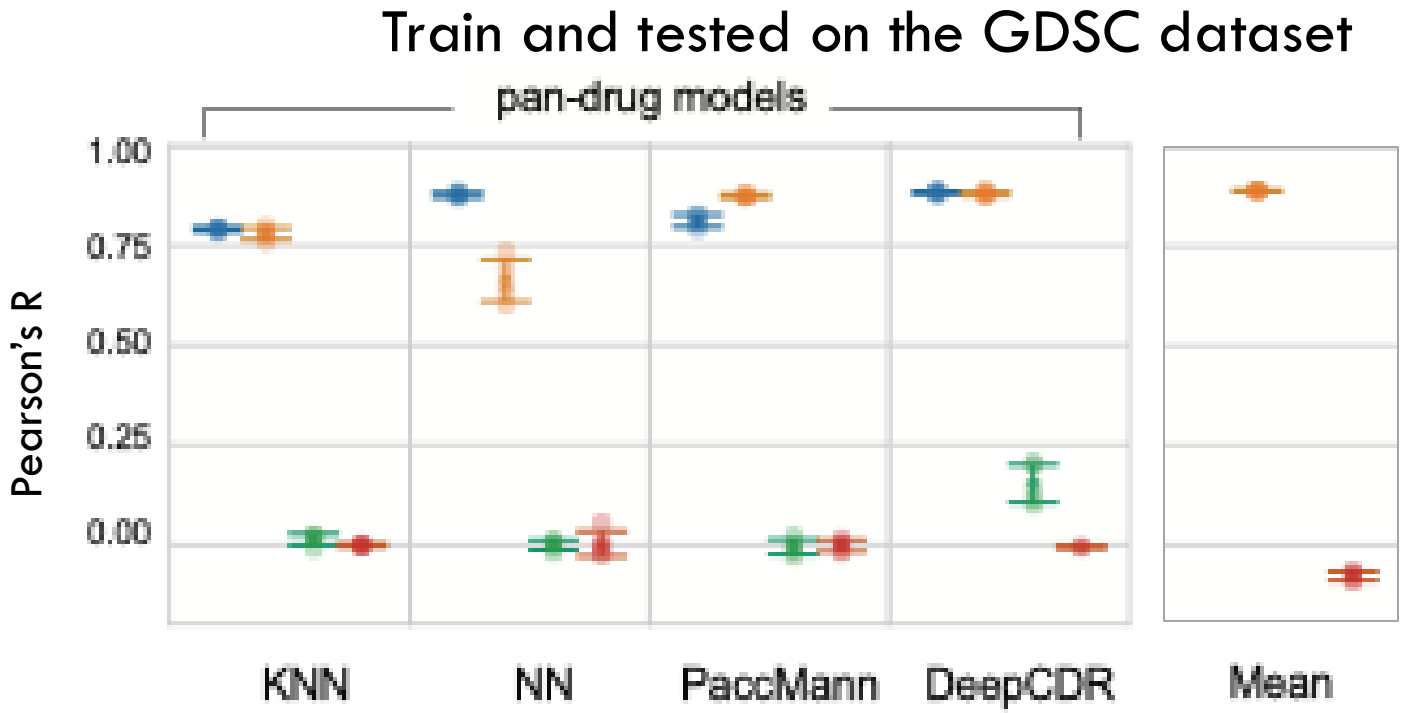

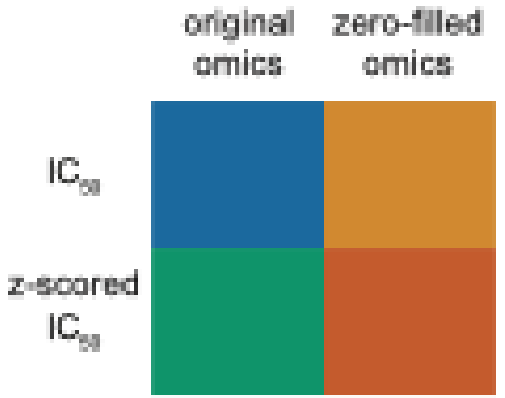

Ovchinnikova *et al.*, npj Precision Oncology, 2024

Predicting perturbation effects **at the single-cell level**

- **Challenge:** scRNA-seq measurements are destructive, no direct pairing between cells
- **Input:** a dataset of *before* & *after* cell states coupled to perturbations $\left(p_i(\mu_i, \nu_i)\right)$

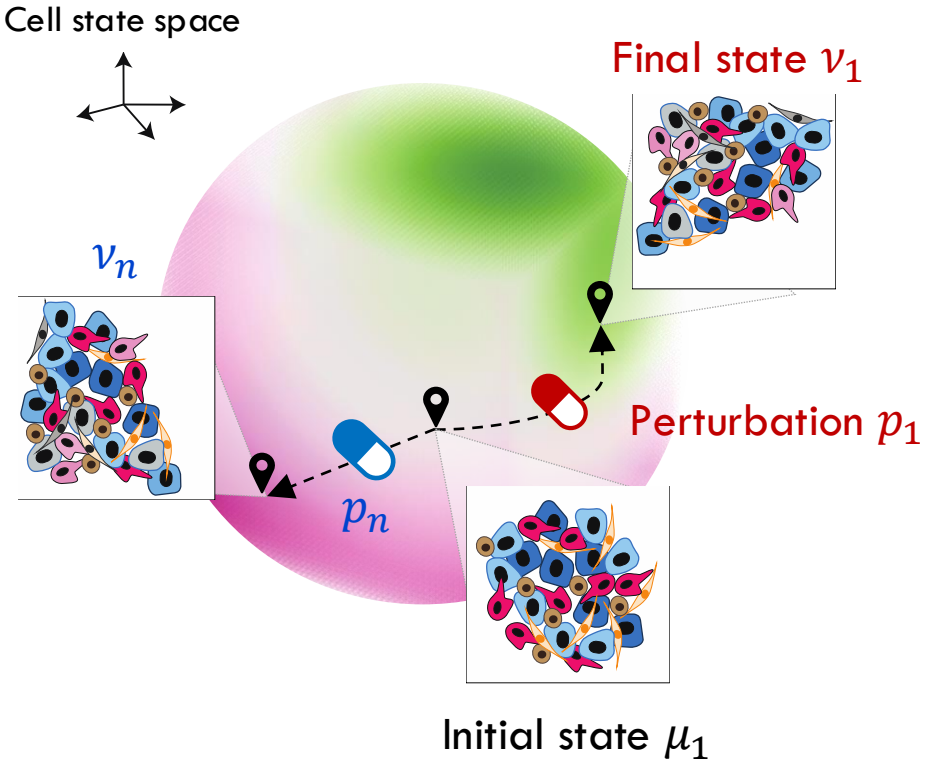

Predicting perturbation effects **at the single-cell level**

- **Challenge:** scRNA-seq measurements are destructive, no direct pairing between cells
- **Input:** a dataset of *before* & *after* cell states coupled to perturbations $\left(p_i(\mu_i, \nu_i)\right)$
- **The problem:** Given an unseen initial state $\boldsymbol{\mu_{new}}$, predict the final state $\boldsymbol{\hat{\nu}_1}$ after known perturbation $\boldsymbol{p_1}$ or predict the final state $\boldsymbol{\hat{\nu}_{new}}$ after unseen perturbation $\boldsymbol{p_{new}}$

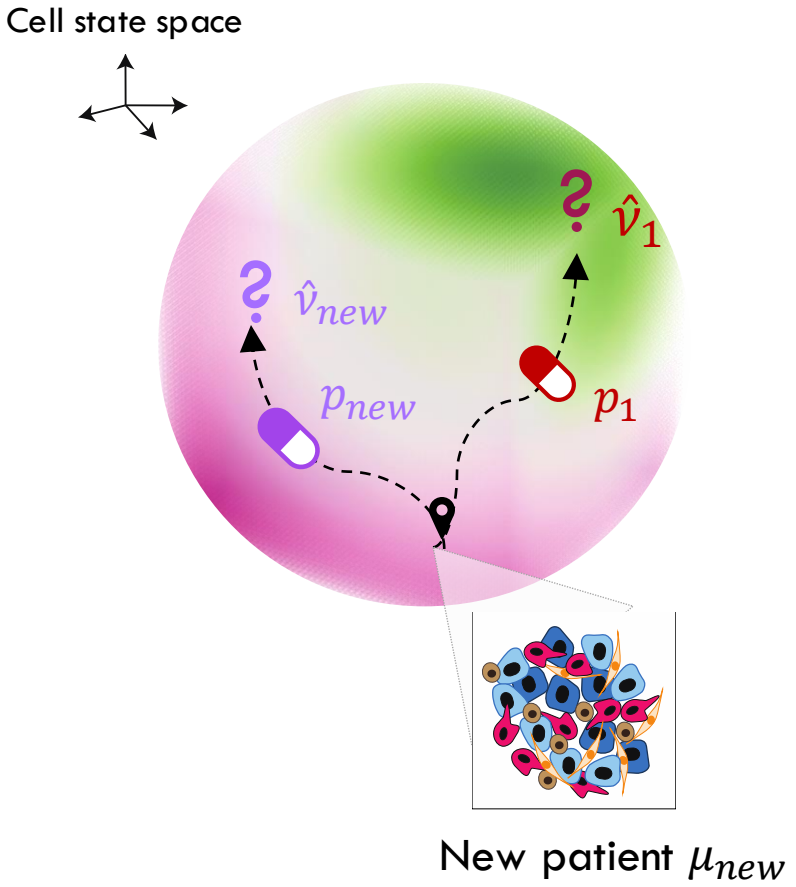

Predicting perturbation effects **at the single-cell level**

- **CMonge***: Use conditional Optimal Transport (OT) to learn transportation maps M_k conditioned on covariates (e.g., dosage, drug, or cell type)

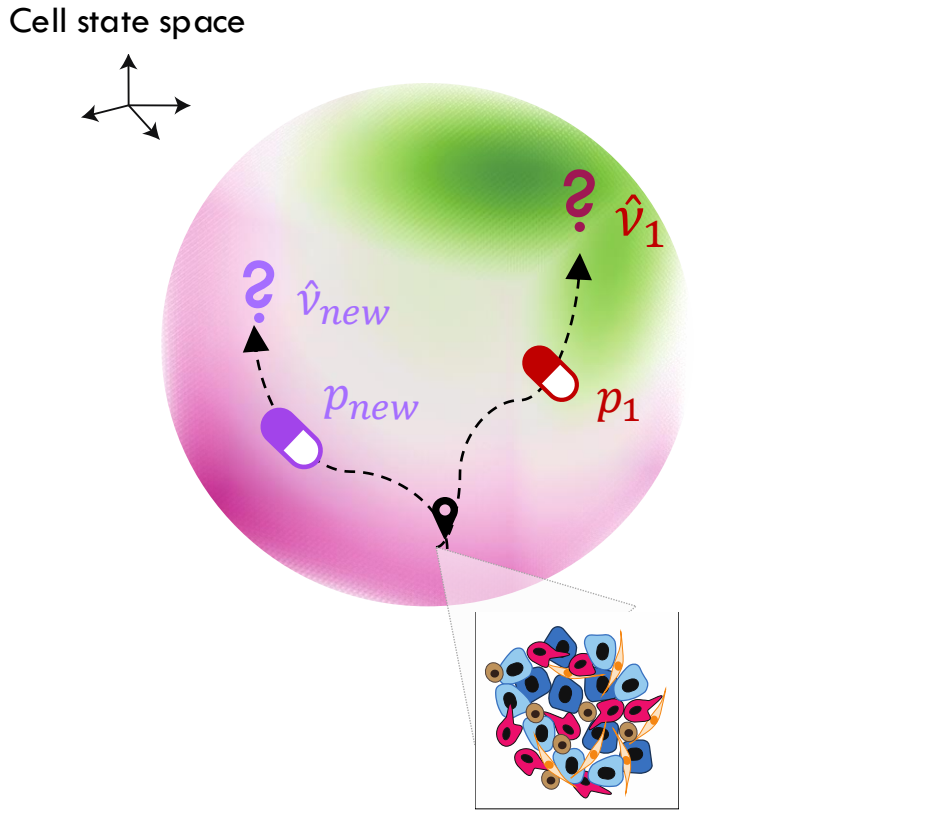


Jannis Born


Benedek Harsanyi


Alice Driessen

*Driessen *et al.* (under review) arXiv:2504.08328

Modeling responses to CAR T cell therapy

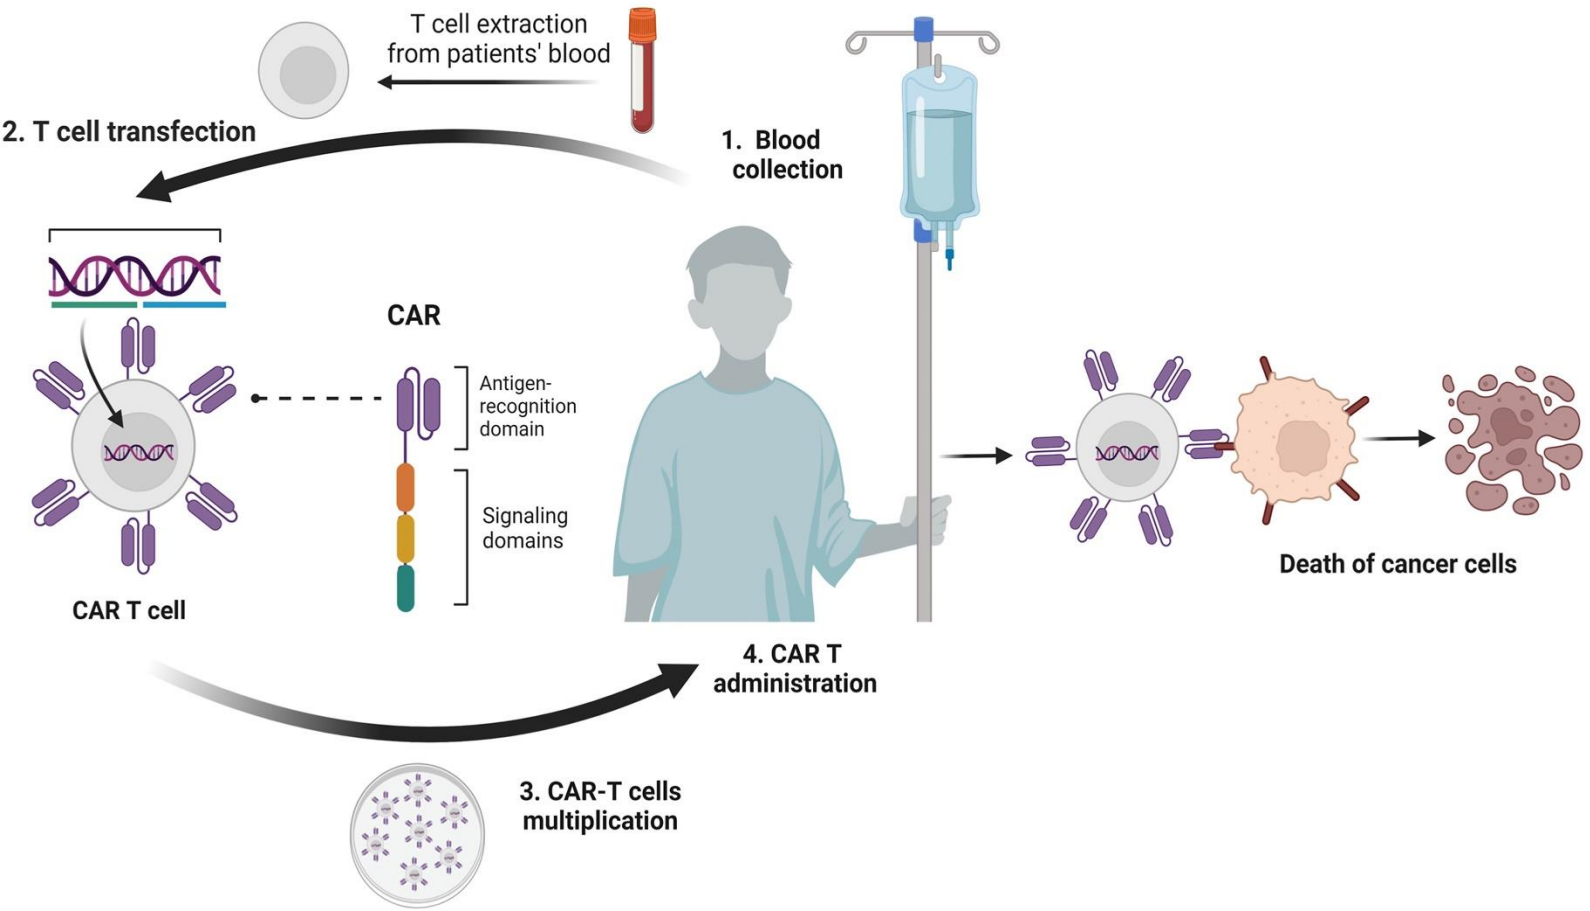

Bui, Thuy Anh et al. eBioMedicine 2024

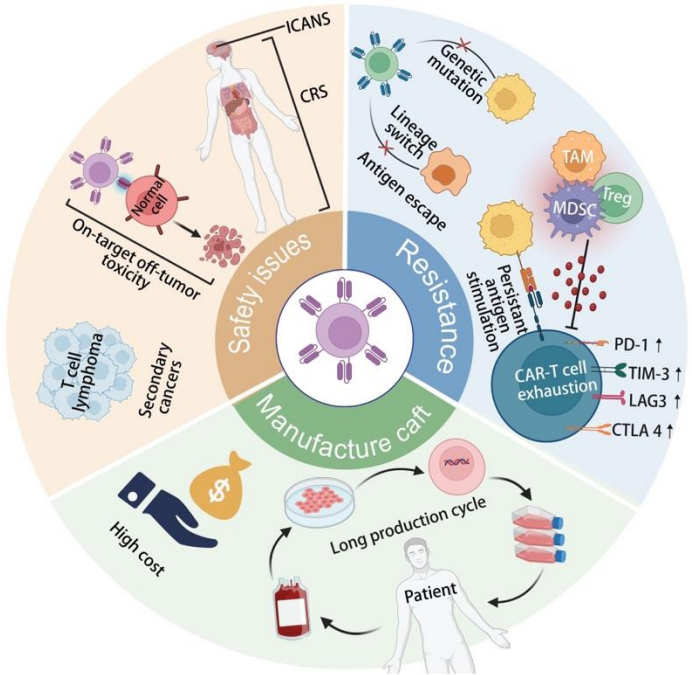

Li Ying et al. Frontiers in Immunology 2024

Modeling responses to CAR T cell therapy

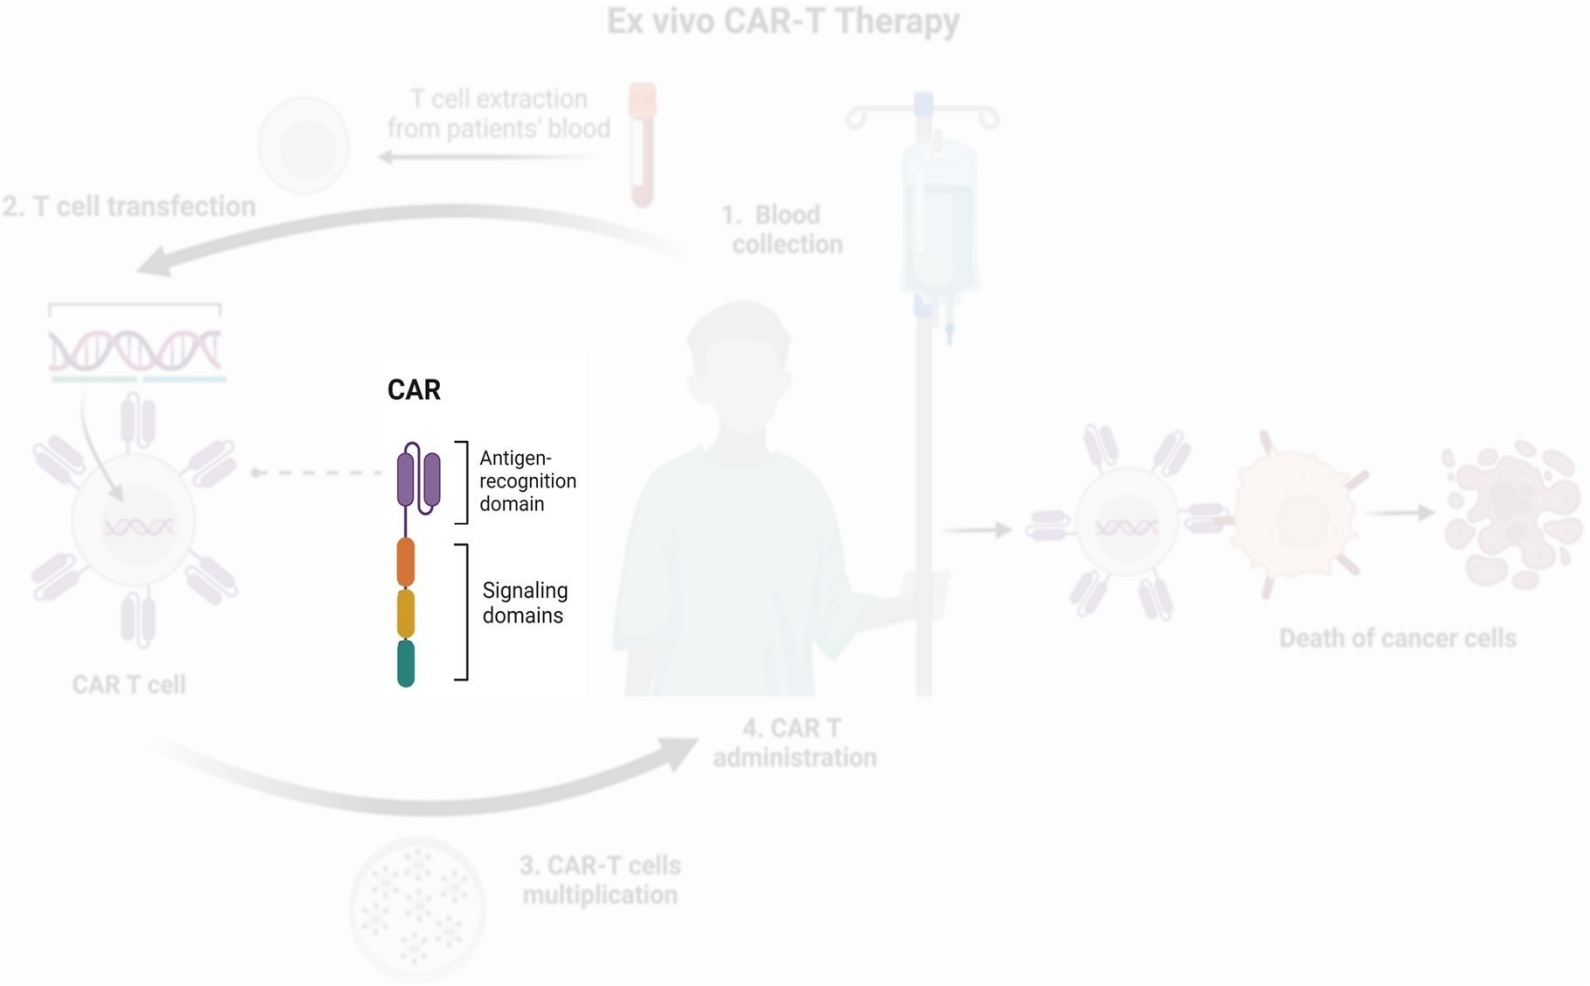

Bui, Thuy Anh et al. eBioMedicine 2024

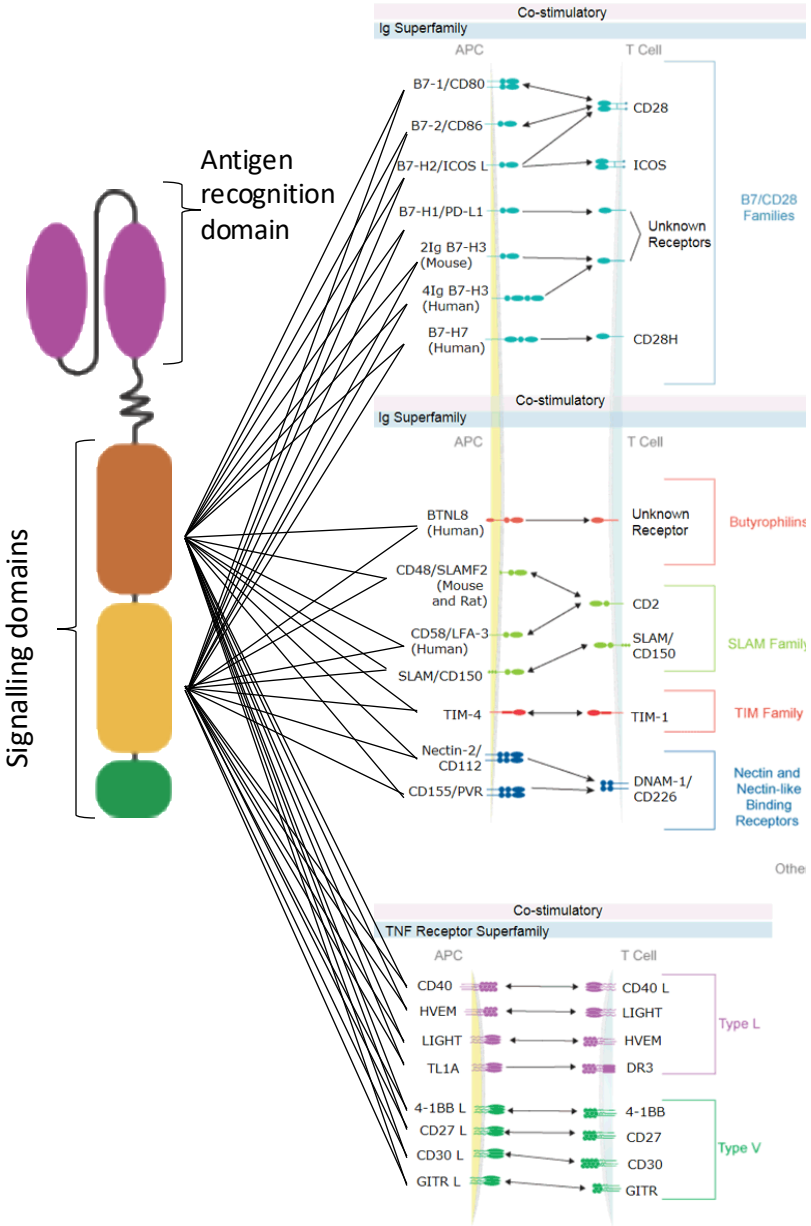

>100'000 possible CARs

>100'000 possible CARs

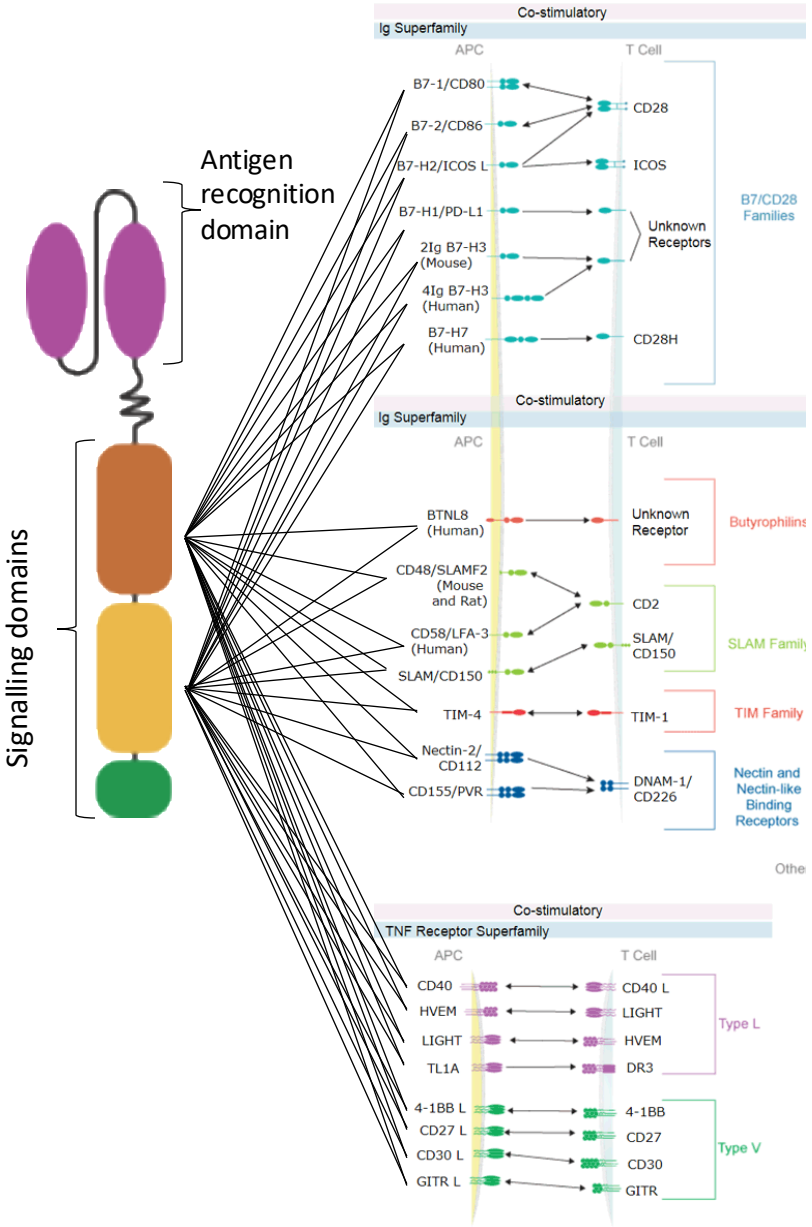

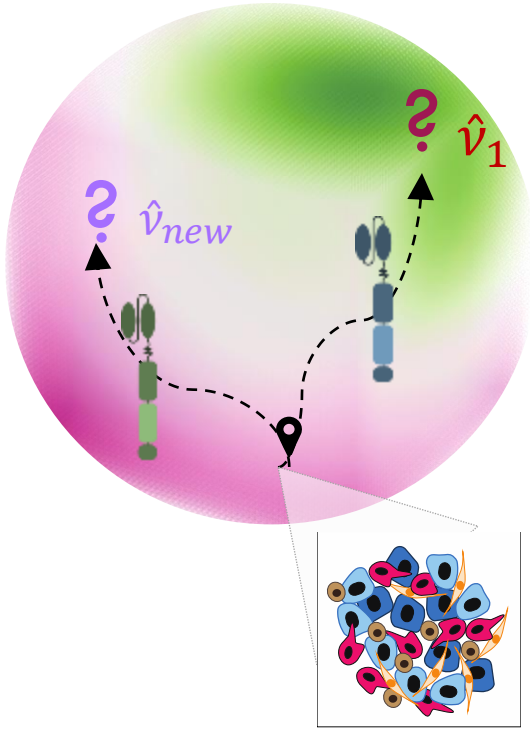

Modeling responses to CAR T cell therapy

Alice Driessen

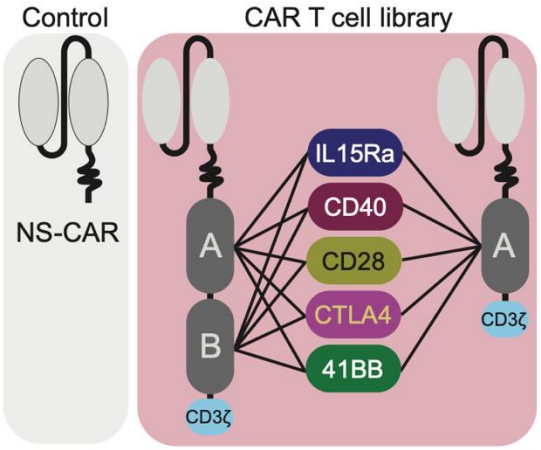

30 different CAR designs

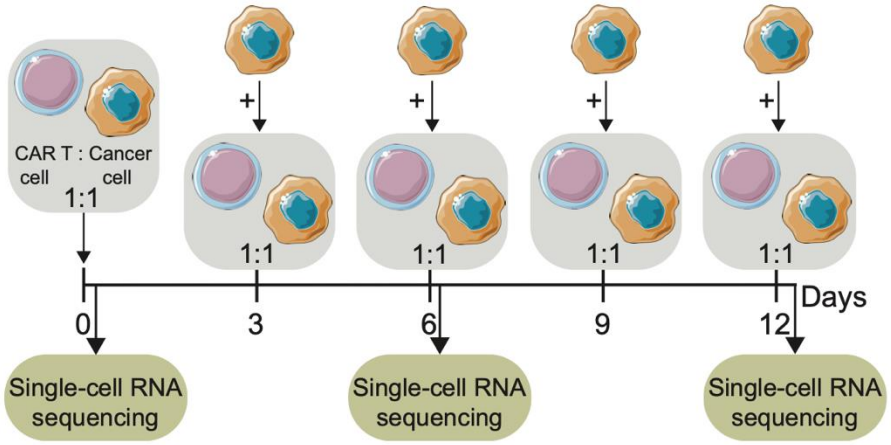

Collaboration with Prof. Sai Reddy, ETHZ BSSE

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

Modeling responses to CAR T cell therapy

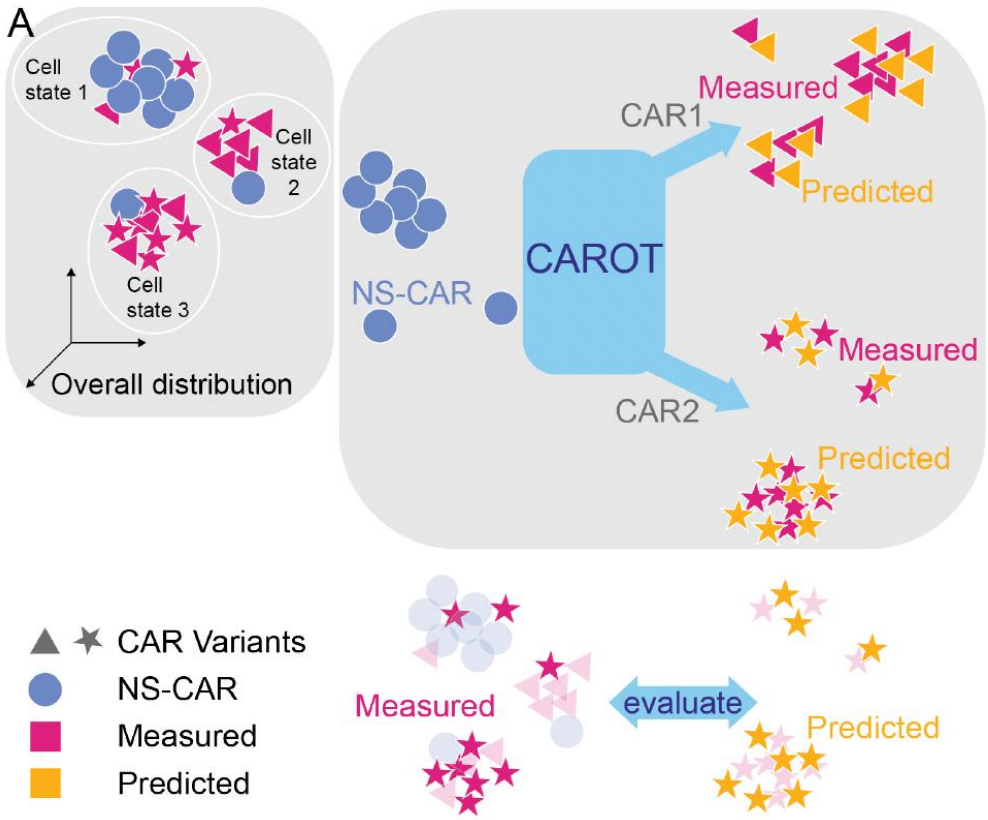

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

Modeling responses to CAR T cell therapy

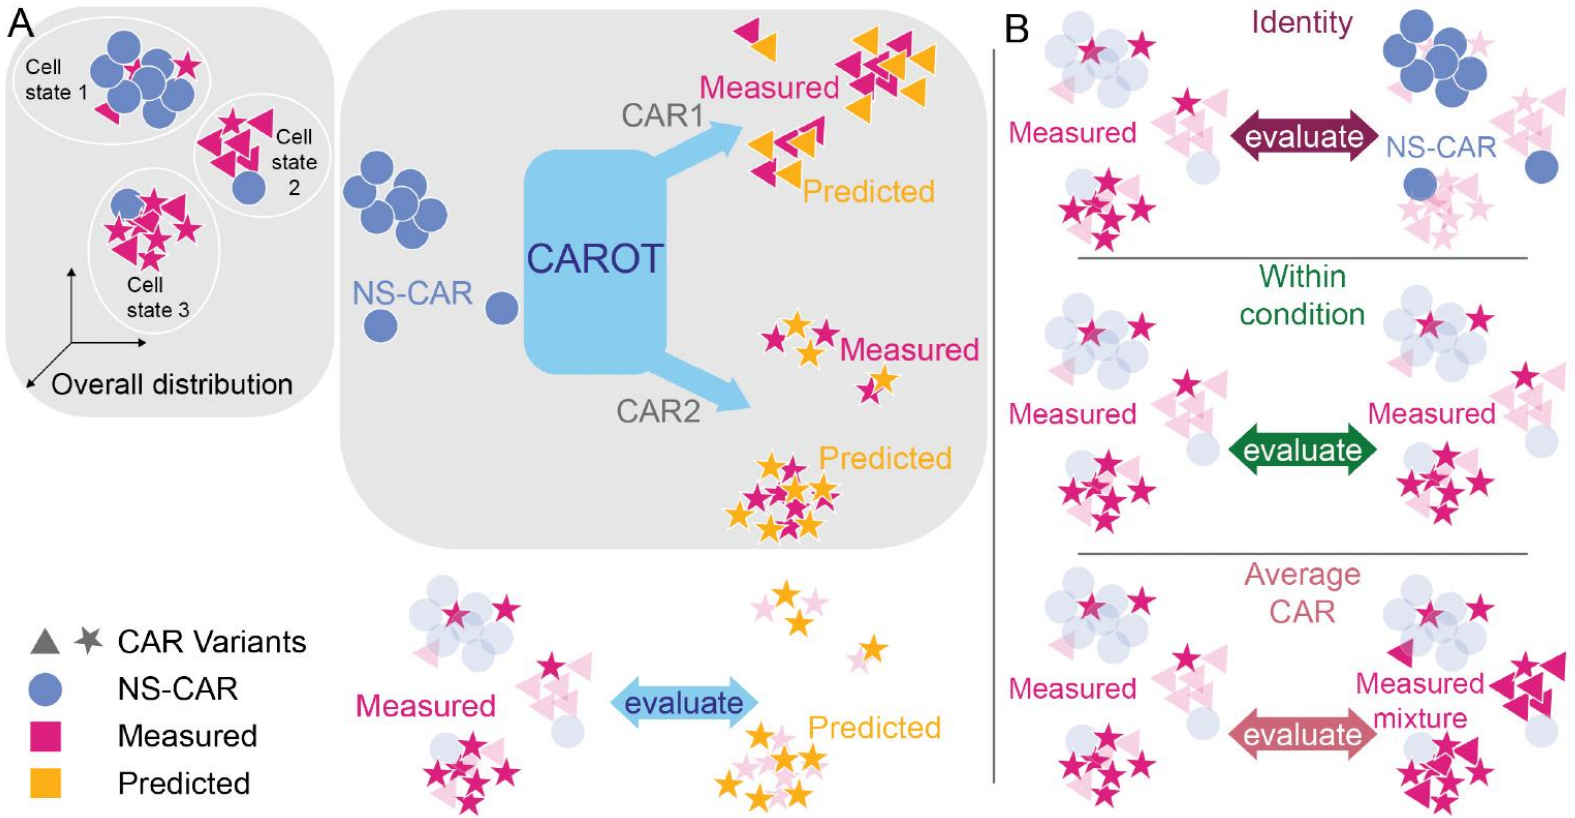

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

One CAROT model *per CAR*

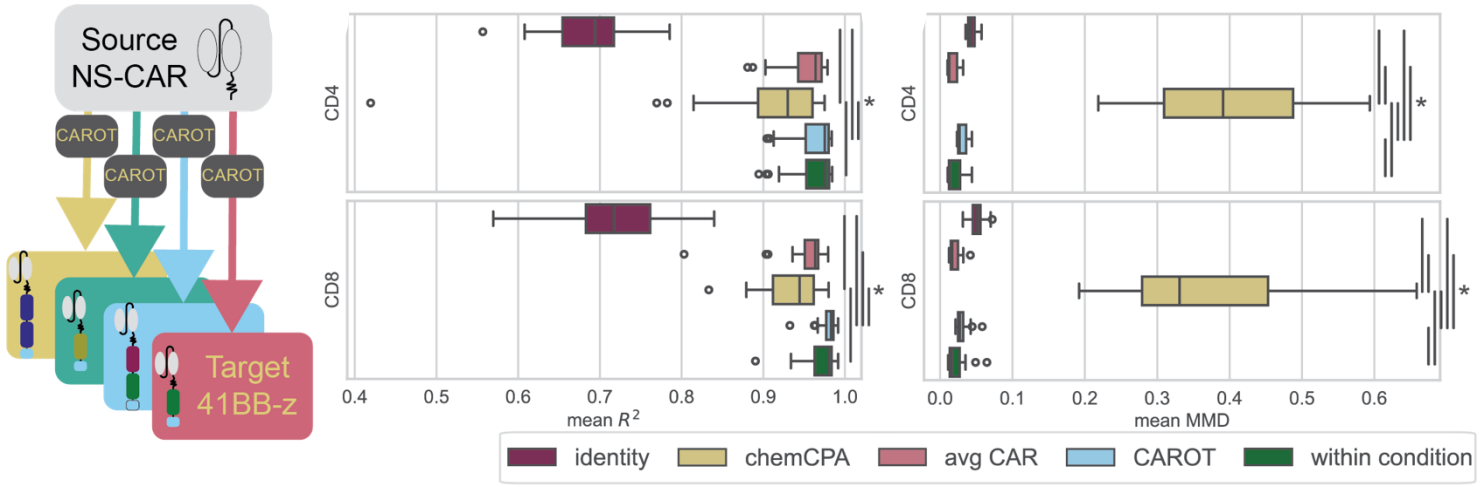

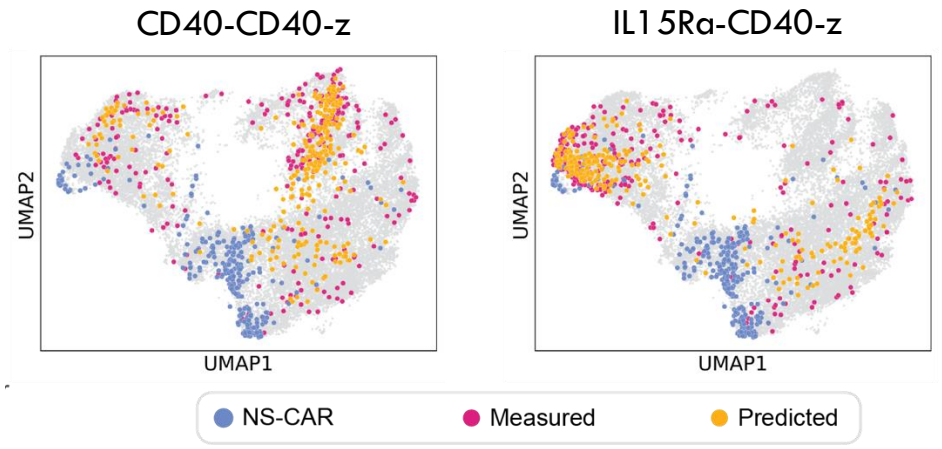

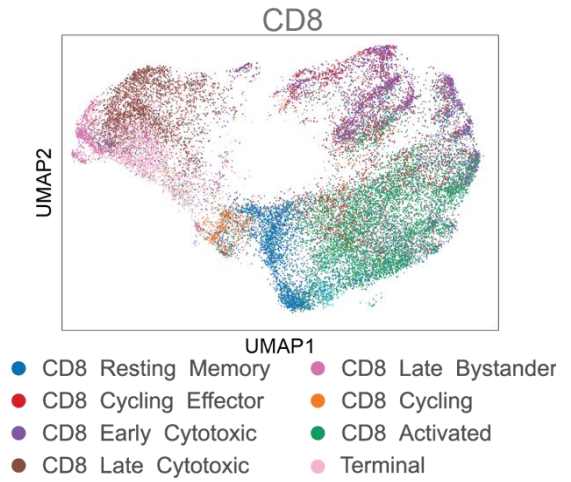

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

One CAROT model *across all CARs*

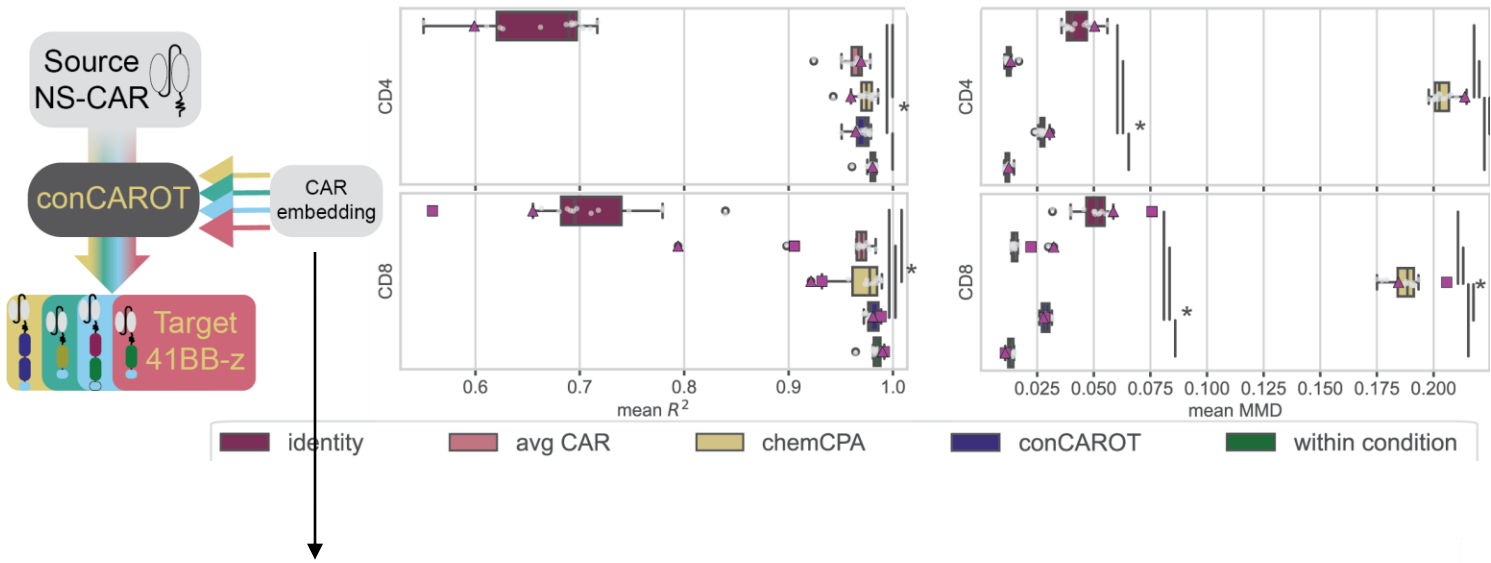

- ESM2 CAR embedding
- Predict novel CAR designs

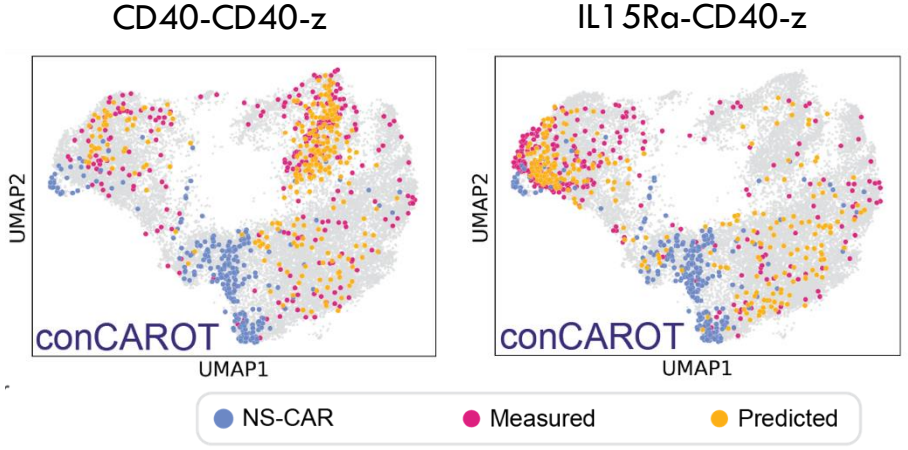

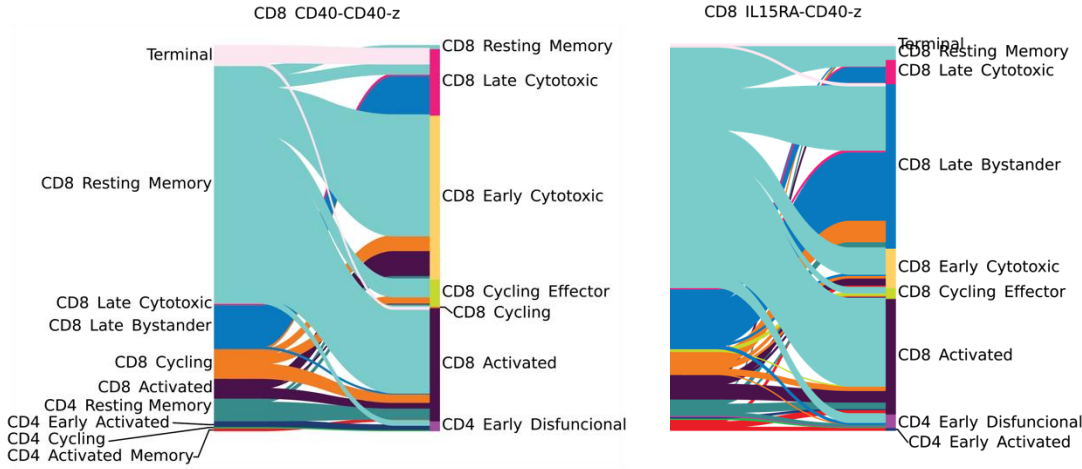

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

One CAROT model *on unseen CARs*

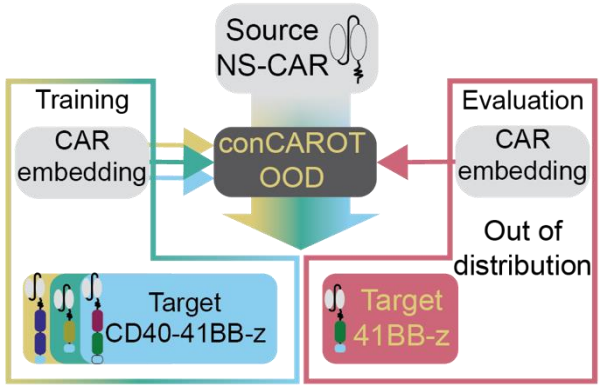

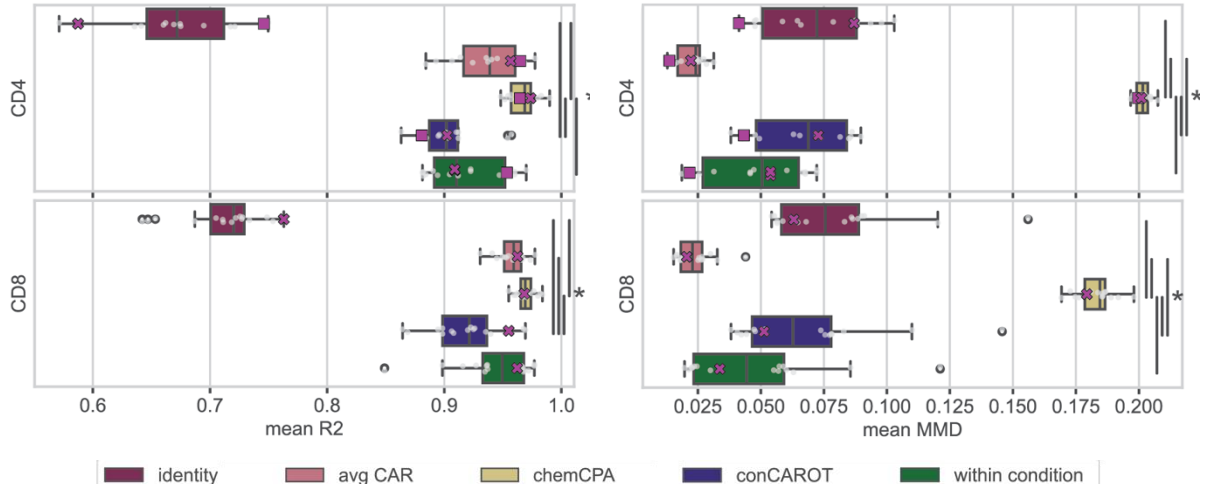

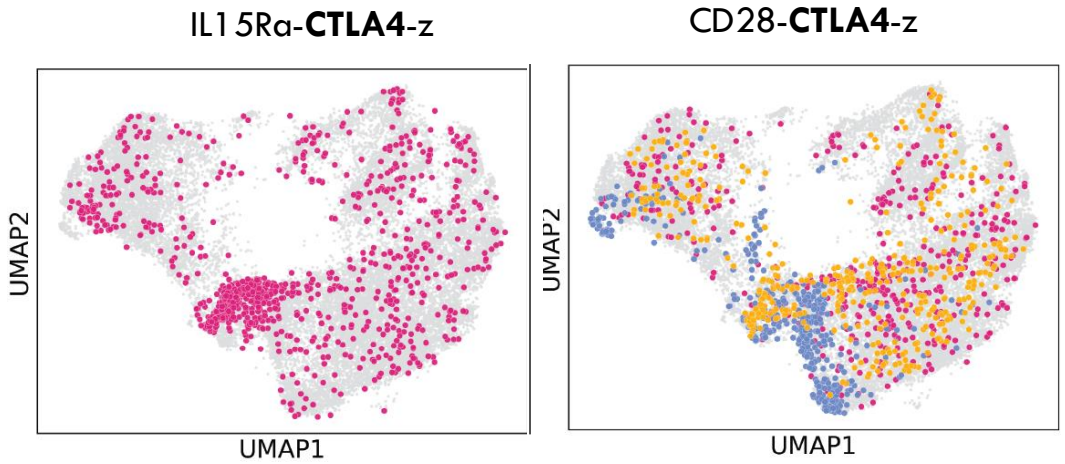

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

CAROT: modeling response to CAR T cell therapy with conditional OT

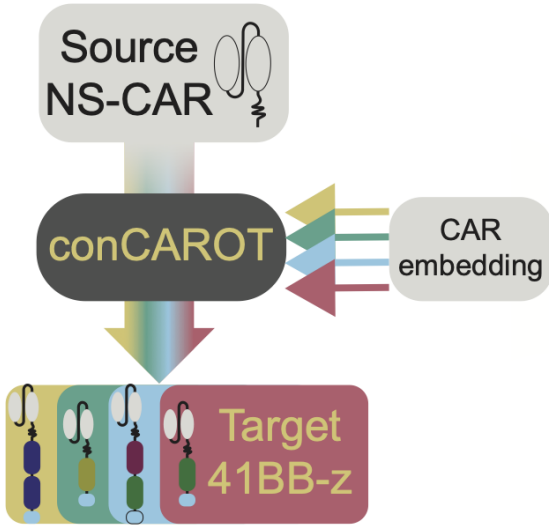

- First single-cell generative AI CAR T cell model
- Conditional model on par with 1-1 model
- ESM2 CAR embedding: Predict novel CAR designs

Driessen et al., NeurIPS AIXDrug Spotlight, https://doi.org/10.1101/2024.11.11.622906

Represent

Histopathology

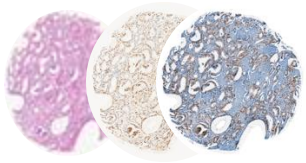

Spatial omics

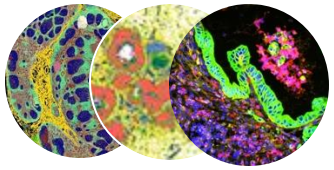

Clinical data

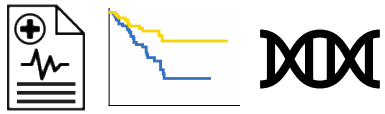

- **MatchCLOT**: Gossi *et al.*, *Briefings in Bioinformatics*, 2023
- Precision Dermatology: Predicting molecular signatures from H&E (in prep)
- Upcoming: Multimodal Foundation models for oncology (in prep)

Interpret

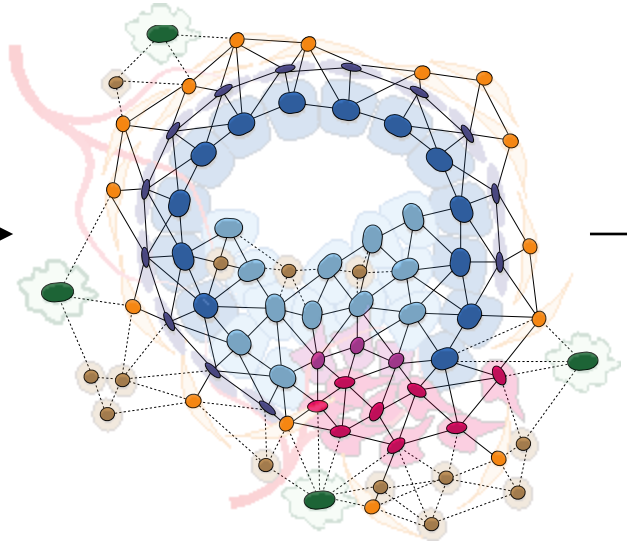

- **ATHENA**: Martinelli & Rapsomaniki, *Bioinformatics*, 2022
- ProteinPNet (McConnel, Imageomics NeurIPS, to appear)
- Concept learning and discovery (in prep)

Perturb

Perturbations

Biological discovery

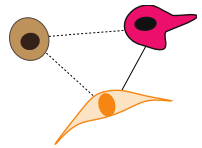

Treatment selection

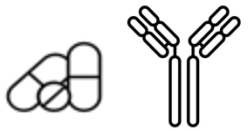

- **Ovchinnikova et al.**, *npj Precision Oncology*, 2024
- **CMonge**: Driessen *et al.*, ICML MLXGen Workshop, arxiv, 2024
- **CAROT**: Driessen et al., bioRxiv 2024.11.11.622906

Agentic AI: the new frontier?

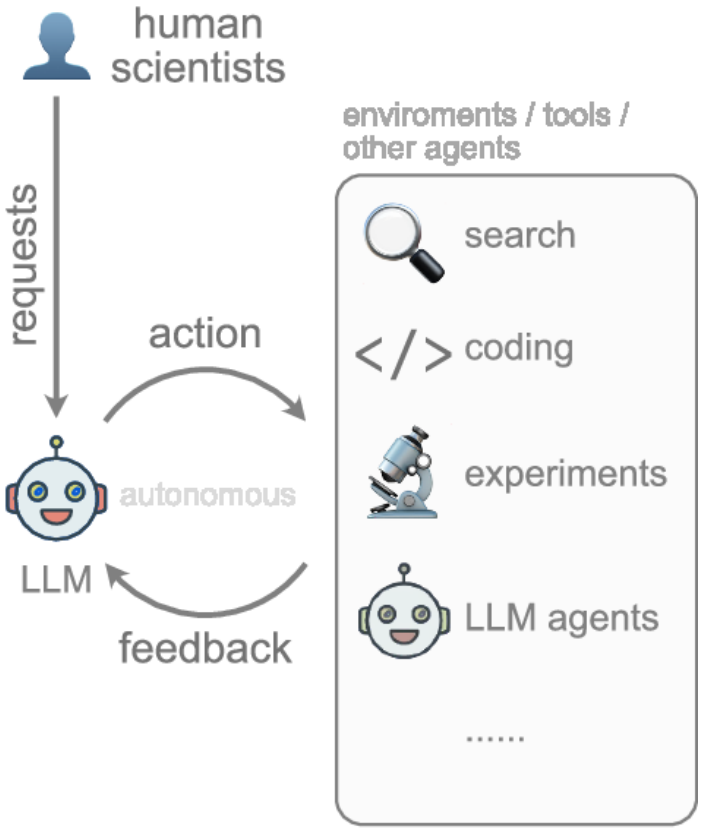

FutureHouse

Automating scientific discovery.

We're a non-profit building AI agents to automate research in biology and other complex sciences.

Sam Rodriques
@SGRodriques

Today, we're announcing the first major discovery made by our AI Scientist with the lab in the loop: a promising new treatment for dry AMD, a major cause of blindness.

Our agents generated the hypotheses, designed the experiments, analyzed the data, iterated, even made figures for the paper. The resulting manuscript is a first-of-a-kind in the natural sciences, in which everything that needed to be done to write the paper was done by AI agents, apart from actually conducting the physical experiments in the lab and writing the final manuscript. We are also introducing Robin, the first multi-agent system that fully automates the in-silico components of scientific discovery, which made this discovery. This is the first time that we are aware of that hypothesis generation, experimentation, and data analysis have been joined up in closed loop, and is the beginning of a massive acceleration in the pace of scientific discovery that will be driven by these agents. We will be open-sourcing the code and data next week.

businesswire
A BERKSHIRE HATHAWAY COMPANY

Business Wire

Owkin Launches K Navigator, a Ground-breaking Agentic Co-pilot to Speed up Breakthroughs in Biomedical Research by 20x

AI/ML for Biomedicine

https://github.com/AI4SCR/

marianna.rapsomaniki@unil.ch

Raphael Gottardo & Krisztian Homicsko
Jonathan Bac, Daria Buszta, Roberto Colotti, Estella Dong, Antoine Girardin, Bernd Illing, Senbai Kang, Jieran Sun, Spencer Watson

Department of Dermatology:
Prof. Michel Guillet, Dr. Jeremy Di Domizio, Dr. Hugo Cometto

IBM **Research** Zurich
Pushpak Pati
Alice Driessen
Jannis Born

Robert Berke, Fei Tang

Sai Reddy & team

University of Bern
Marianna Kruithof-de Julio
Sofia Karkampouna
Francesco Bonollo
Katia Ovchinnikova
Martin Wartenberg

University of Vienna
Eva Comperat

Funding

Swiss National Science Foundation

CHAN ZUCKERBERG INITIATIVE

krebsforschung schweiz
recherche suisse contre le cancer
ricerca svizzera contro il cancro
swiss cancer research

OWKIN

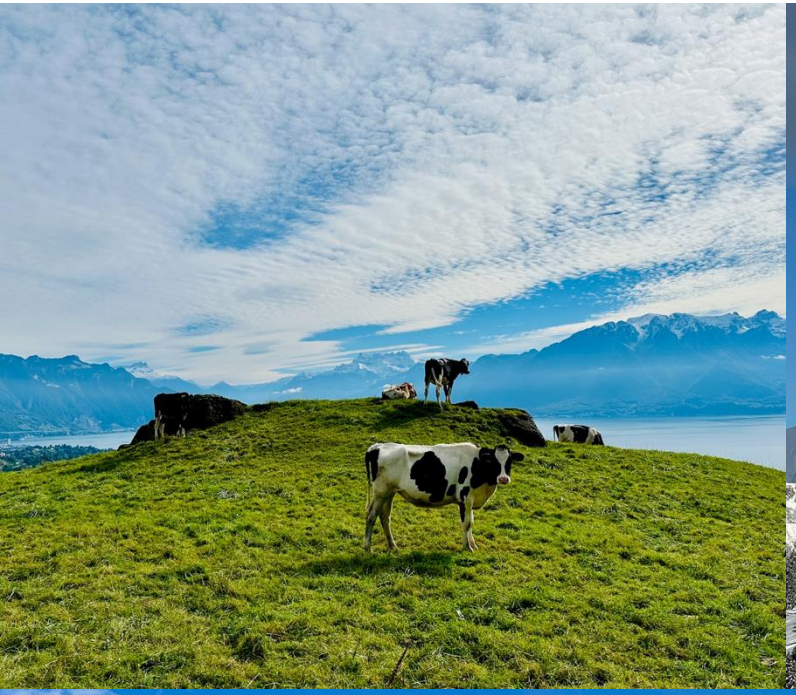

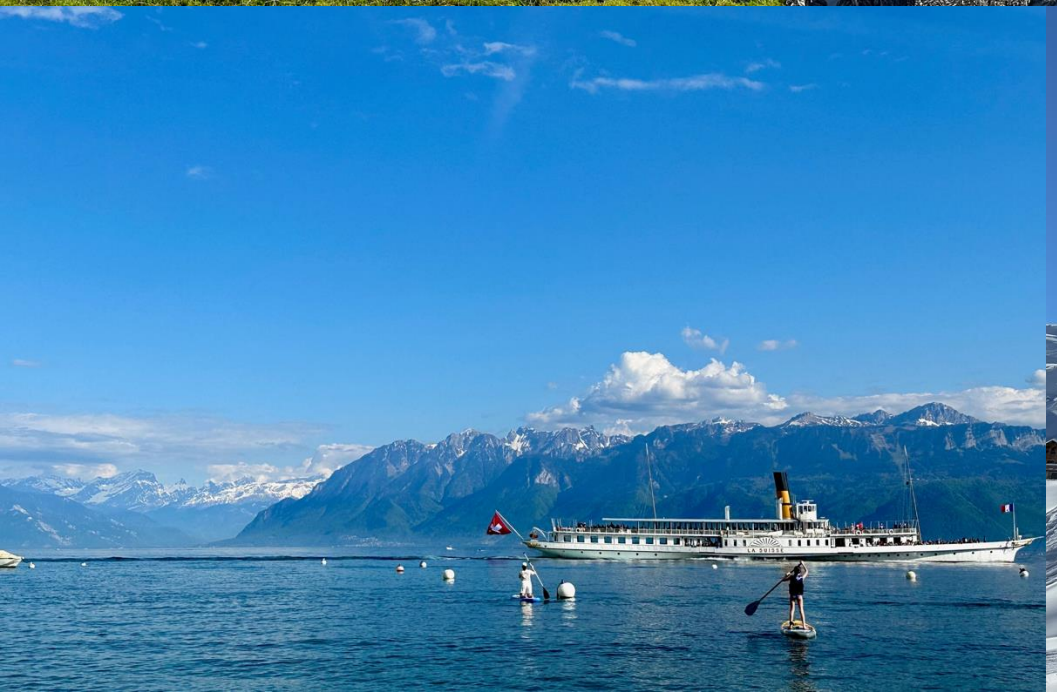

marianna.rapsomaniki@unil.ch