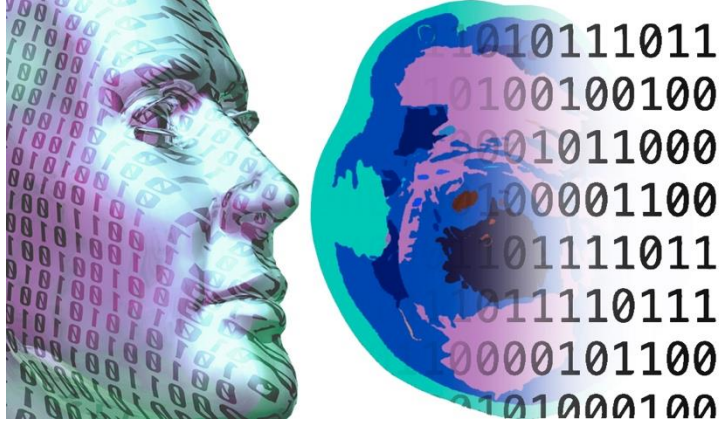

Learning and characterizing cellular dynamics with Neural and Graph ODE Networks (Part 2)

Smita Krishnaswamy

Department of Computer Science, Department of Genetics

Program for Applied Math, Program for Computational Biology & Bioinformatics

Program for Interdepartmental Neuroscience

Wu Tsai Institute Center for Neurocomputation and Machine Intelligence

Learning and modulating dynamics

Learn complex cellular dynamics

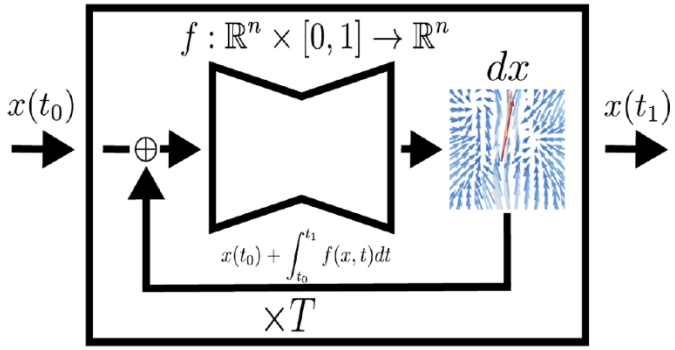

TrajectoryNet
MIOflow

Build networks,
Find targets

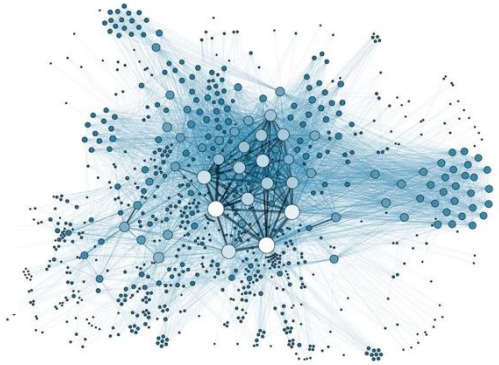

Granger
RITINI

Phenoscapes

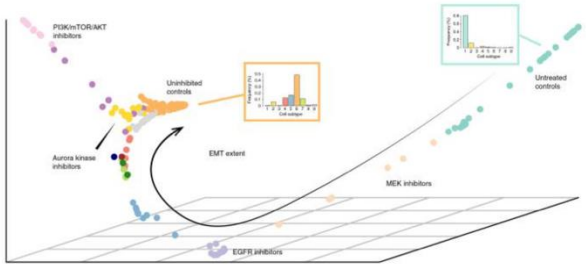

Diffusion EMD
Trellis

Manifold Interpolating flows

(Huguet, Magruder, et al. NeurIPS 2022, Sun, Kuchroo, Tong et al., preprint, Sun et al. In preparation)

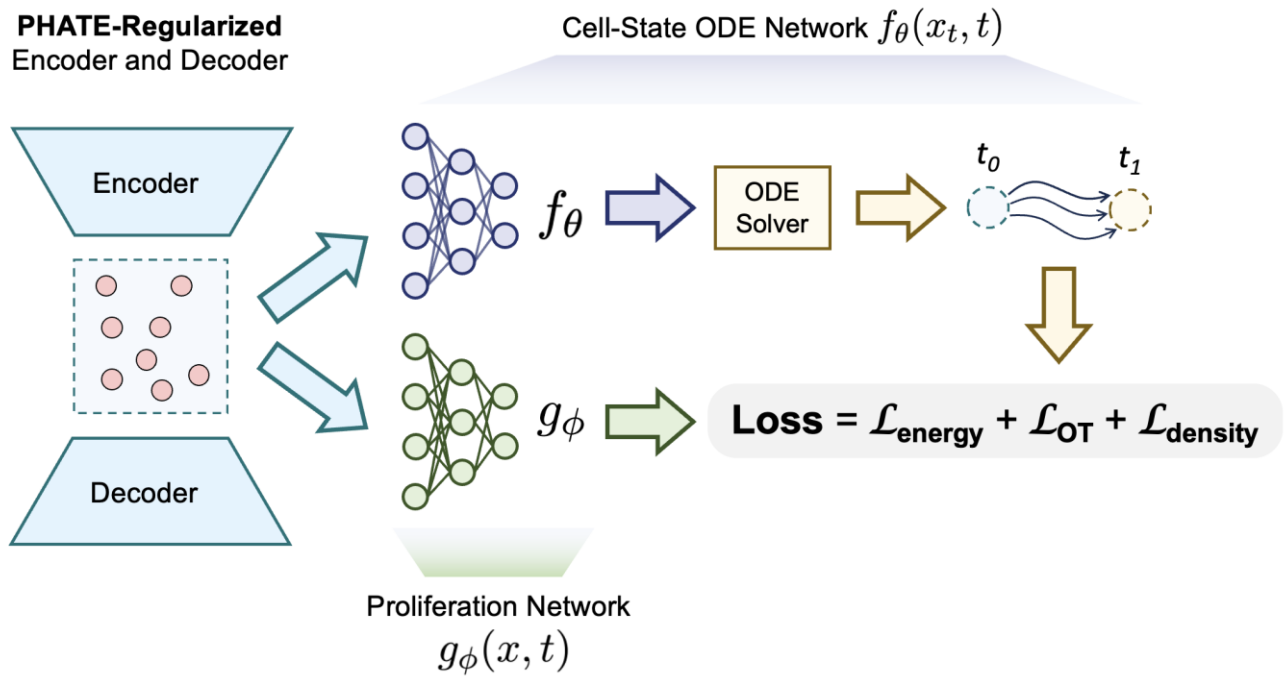

Xingzhi Sun
Guillaume Huguet
Alex Tong
Manik Kuchroo
Sumner Magruder
Guy Wolf

Main idea: Flow data in a manifold embedding penalized by optimal transport

Dynamic optimal transport of cells

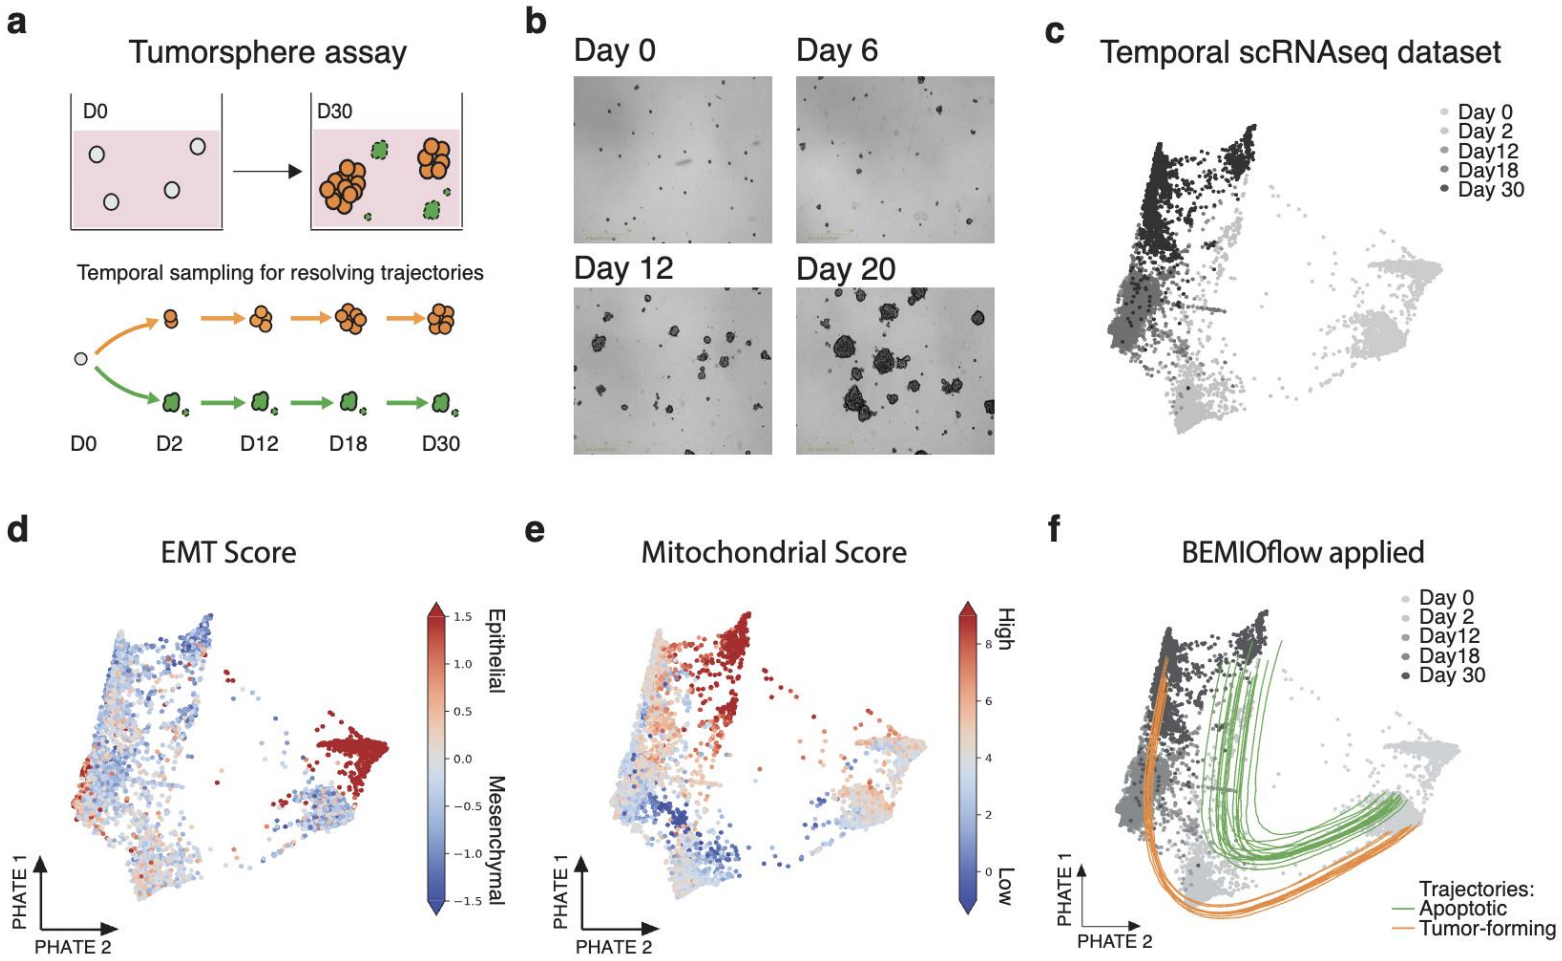

Cell-of-Origin Analysis (Cancer stem cells)

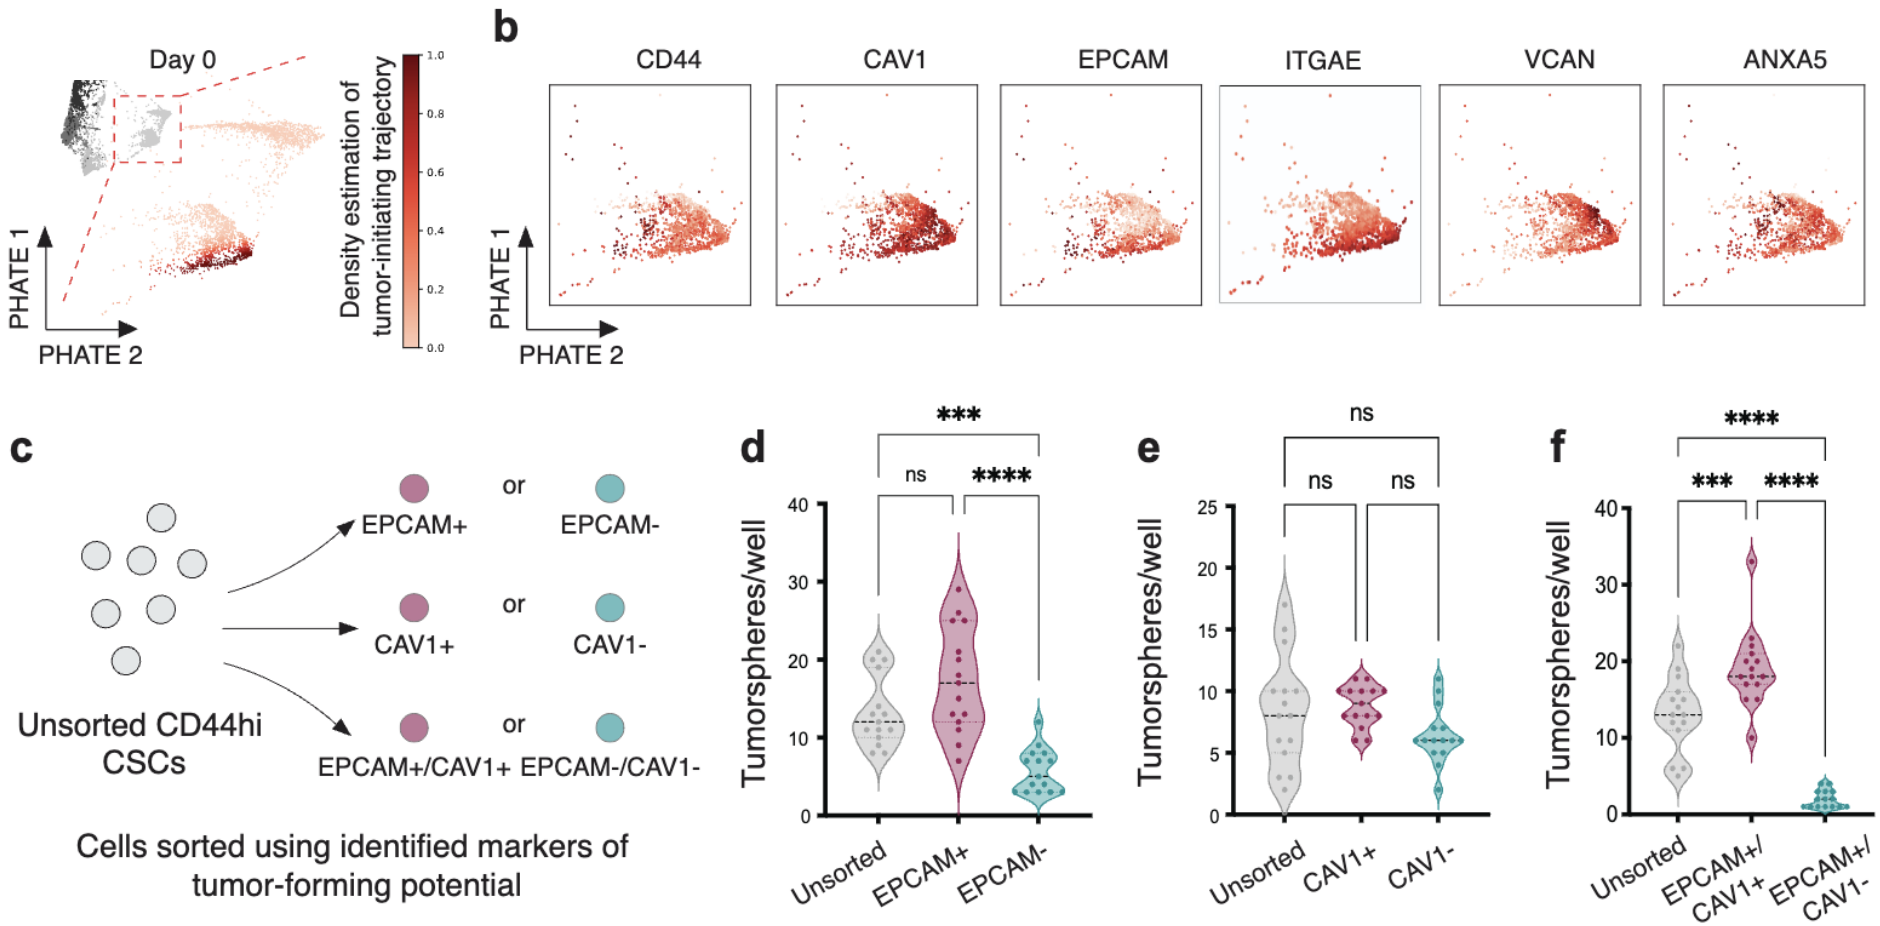

Timed-gene expression modules

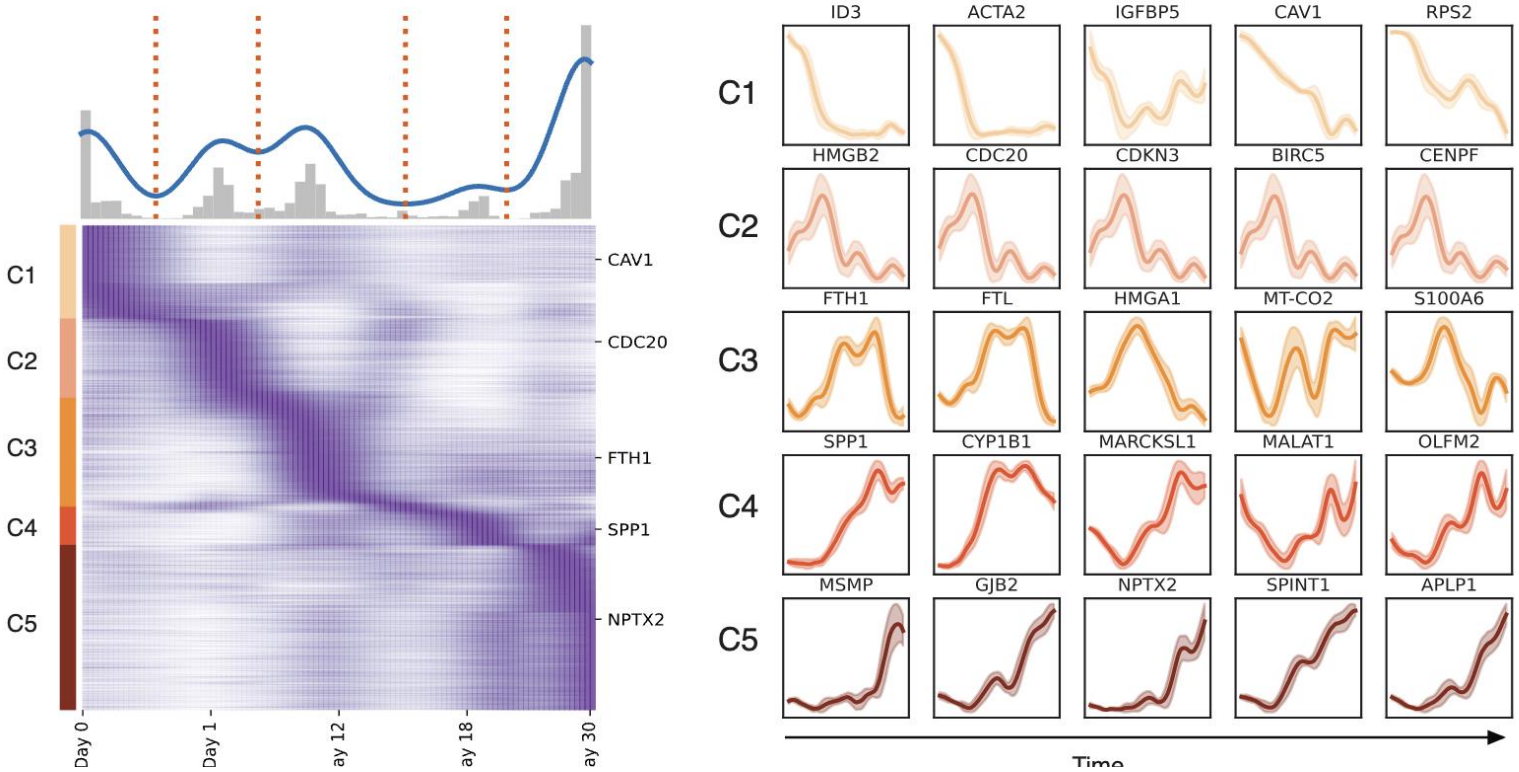

Granger Causality

Granger Causality

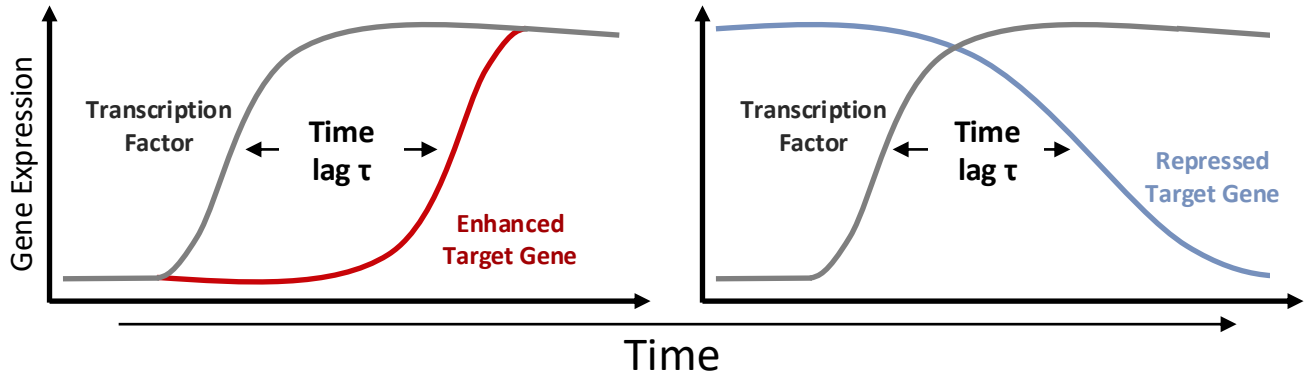

Granger causality checks that

1. Cause happens prior to effect
2. The cause has unique information about the effect

$$\mathbb{P}[Y(t+1) \in A \mid \mathcal{I}(t)] \neq \mathbb{P}[Y(t+1) \in A \mid \mathcal{I}_{-X}(t)]$$

Computing Granger Causality

3. Perform an F-test

Test whether the coefficients $\gamma_1, \gamma_2, \ldots, \gamma_p$ are jointly zero.

- Null hypothesis H_0: X does **not** Granger-cause Y (i.e., $\gamma_i = 0$)
- Alternative hypothesis H_1: At least one $\gamma_i \neq 0$

$$F = \frac{(RSS_{\text{restricted}} - RSS_{\text{unrestricted}})/p}{RSS_{\text{unrestricted}}/(T - 2p - 1)}$$

Where:

- RSS = residual sum of squares
- T = number of observations

If the F-statistic is large and the p-value is small, reject H_0 → X Granger-causes Y.

Granger causality creates complex transcriptional networks

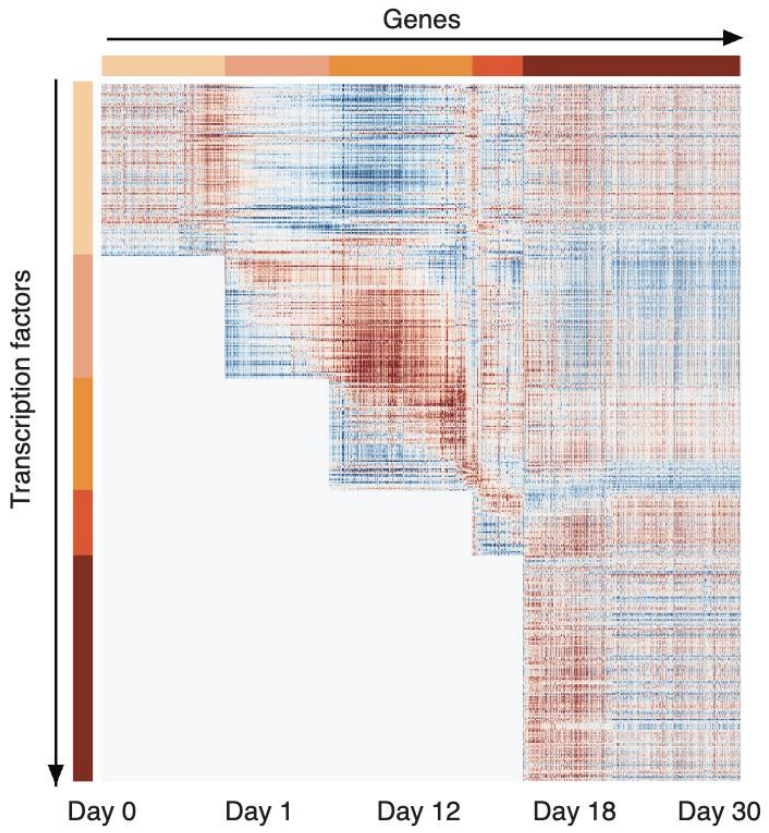

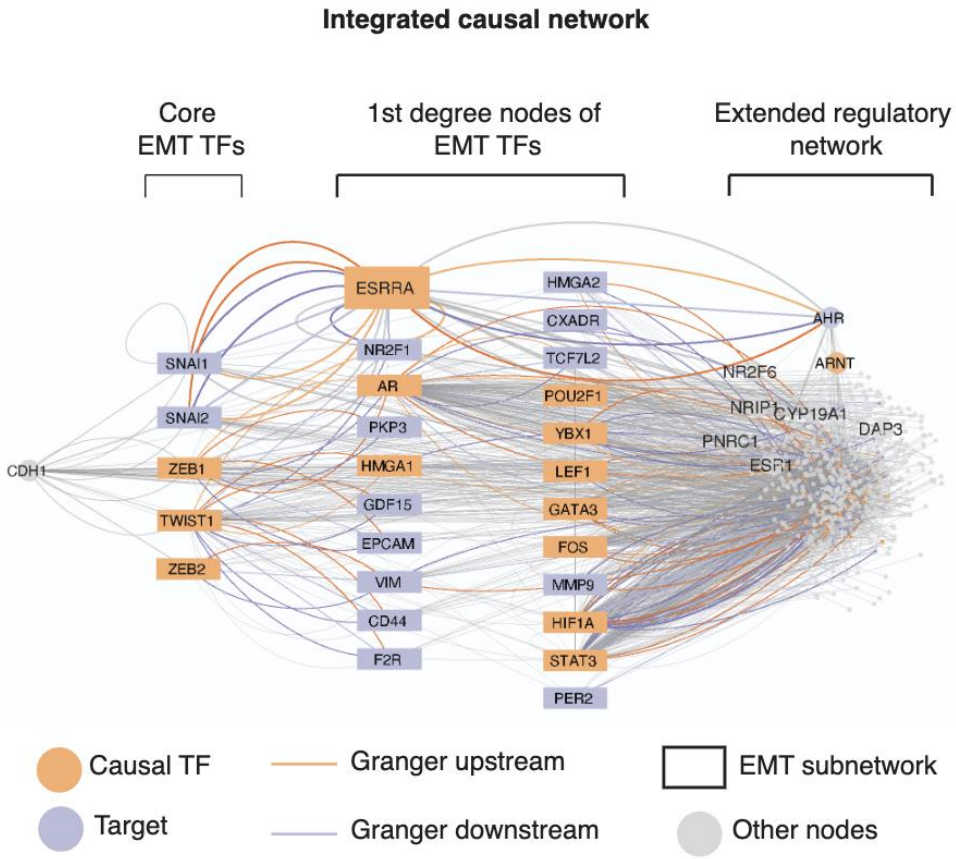

Experimental validation

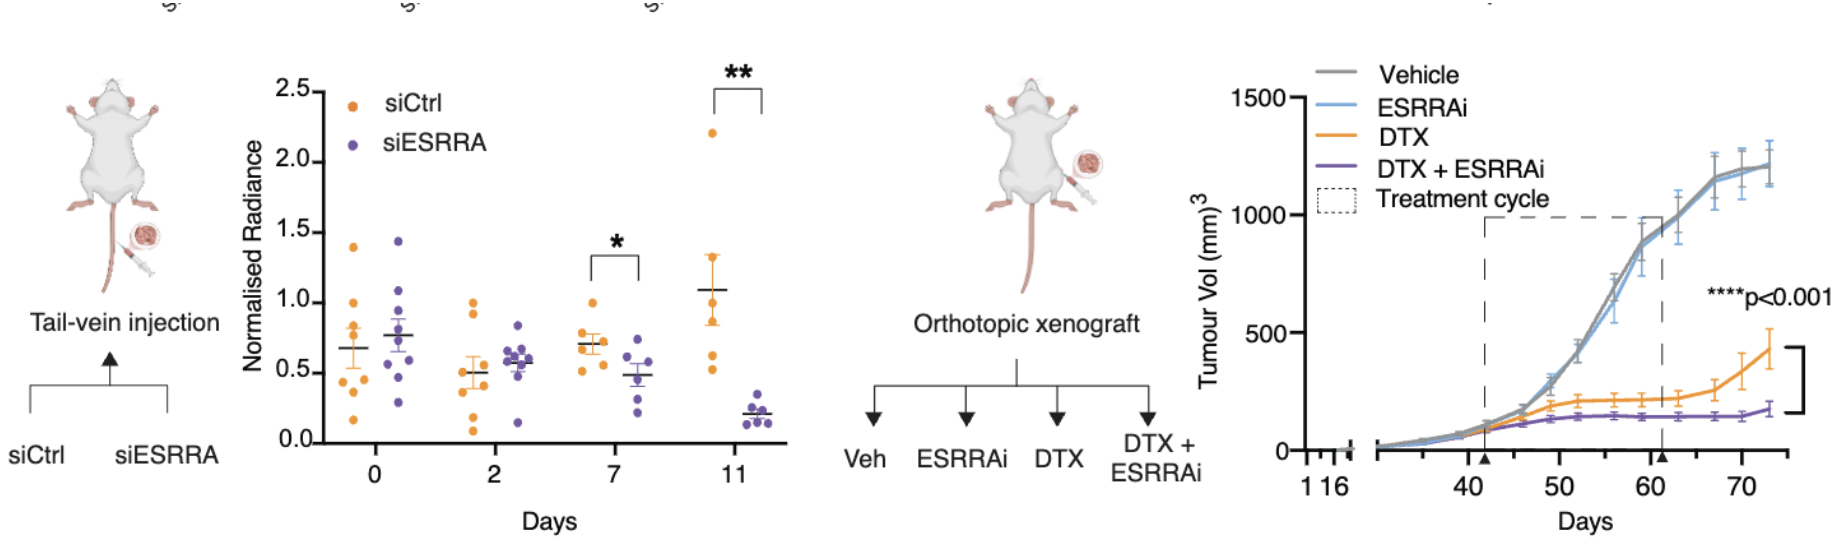

Regulatory Rewiring

- Cells undergo dynamic rewiring of the gene regulatory network
- Signals that are “off” in certain states deactivate parts of the network
- Epigenetic and other changes can turn on/off connections

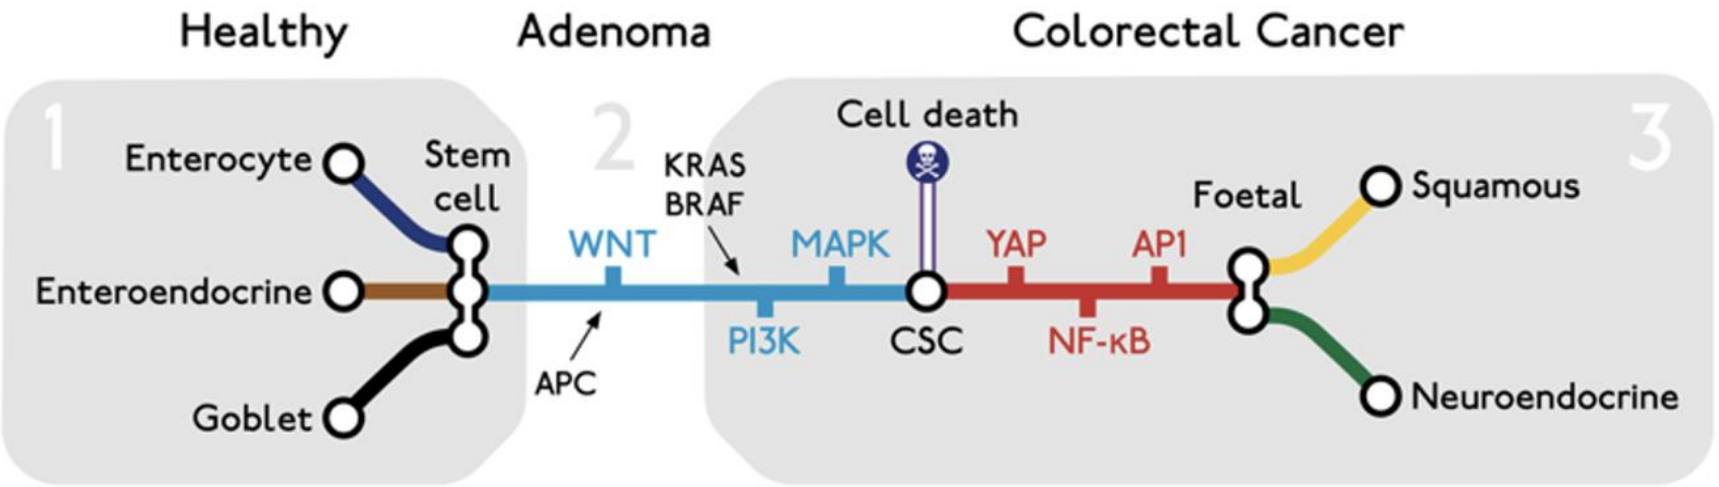

Modulating the rewiring can restore treatment sensitivity

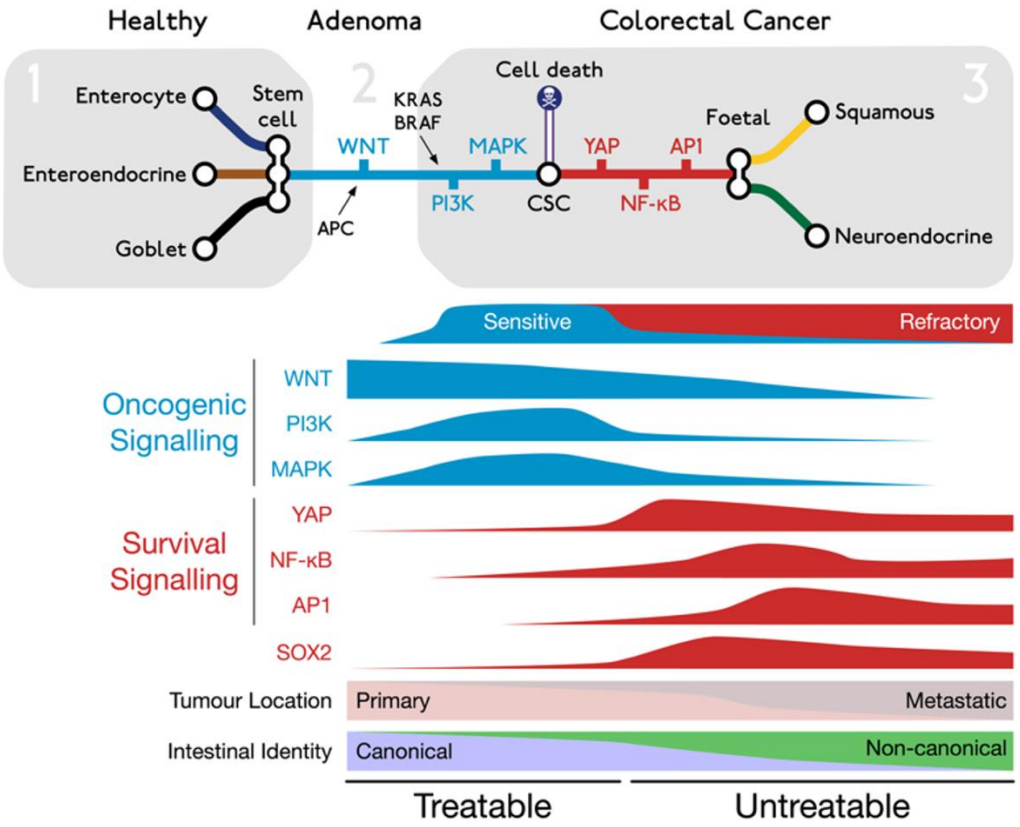

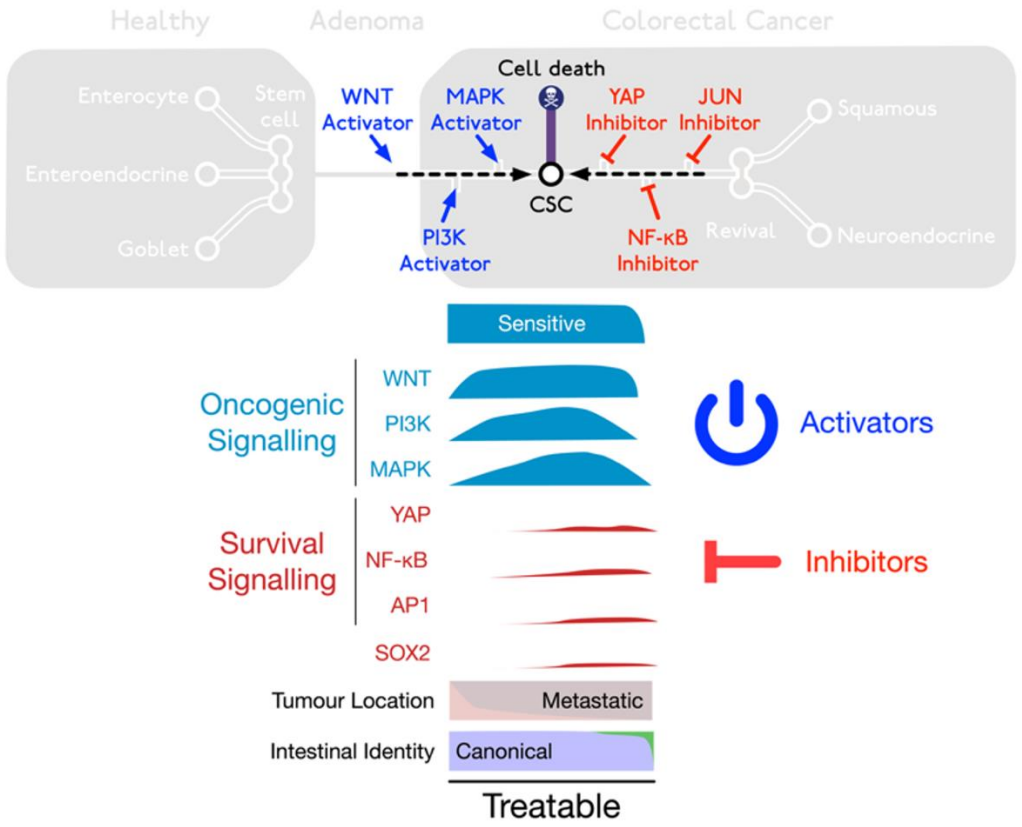

Proposed work for REWIRE-Cancer CRUK Project

Revealing dynamic temporal trajectories and underlying regulatory networks with *Cflows*

Xingzhi Sun[1*], Shabarni Gupta[2,3*] Alexander Tong[1,4,5*], Manik Kuchroo[6,7*], Dhananjay Bhaskar[1*], Chen Liu[1], Aarthi Venkat[8], Beatriz P. San Juan[2,3], Laura Rangel[2,3], Vanina Rodriguez[2,3], John G. Lock[9], Christine L. Chaffer[2,3§], Smita Krishnaswamy[1,5,10§]

Beta cells are essential drivers of pancreatic ductal adenocarcinoma development

Cathy C. Garcia[1,2,3,12], Aarthi Venkat[4,12], Daniel C. McQuaid[1,2,5], Sherry Agabiti[1,2], Alex Tong[6]; Rebecca L. Cardone[7,8], Rebecca Starble[1], Akin Sogunro[1,2,5], Jeremy B. Jacox[2,9,10], Christian F. Ruiz[1,2], Richard G. Kibbey[7,8,9], Smita Krishnaswamy[1,4,6,13], and Mandar Deepak Muzumdar[1,2,3,5,9,10,11,13,14]

Regulatory Temporal Interaction Network Inference (RiTINI)

(Bhaskar, Magruder, et al. 2023, Learning On Graphs (LOG))

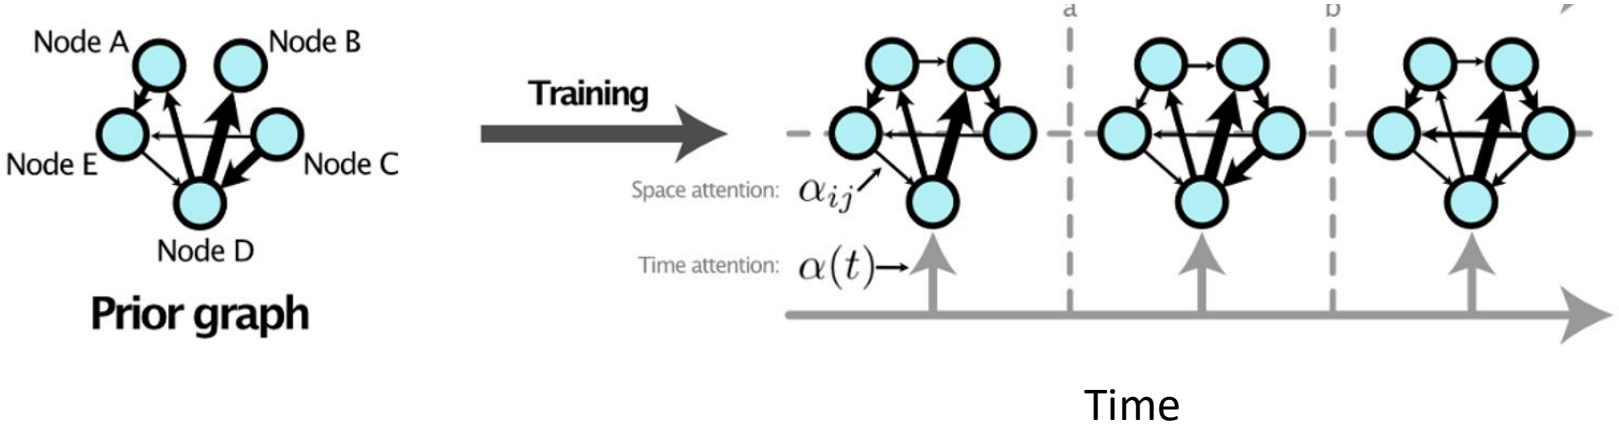

Dhananjay Bhaskar
Siddharth Viswanath
Sumner Magruder
Edward De Brouwer
Frederik Wenkel
Guy Wolf

Main idea: Learn a dynamic interaction graph that matches the regulatory network using graph attention

Dynamic Regulatory Network Inference

- Genes dynamically change their expression patterns in response to altered interactions
- Can we create a generative dynamics model that follows this same principle?
- **Caveat:** there could be many graphs that match the same dynamics

GRN inference from single-cell gene expression trajectories

RiTINI : “Regulatory Temporal Interaction Network Inference”

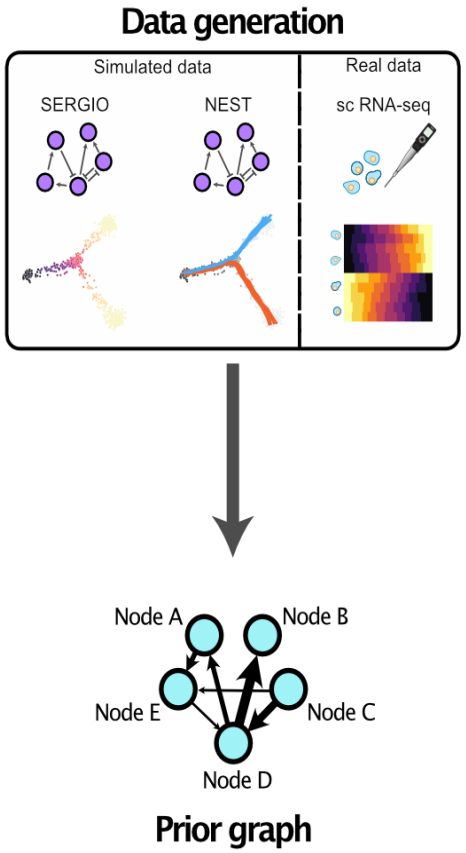

GRN inference from single-cell gene expression trajectories

RiTINI : “Regulatory Temporal Interaction Network Inference”

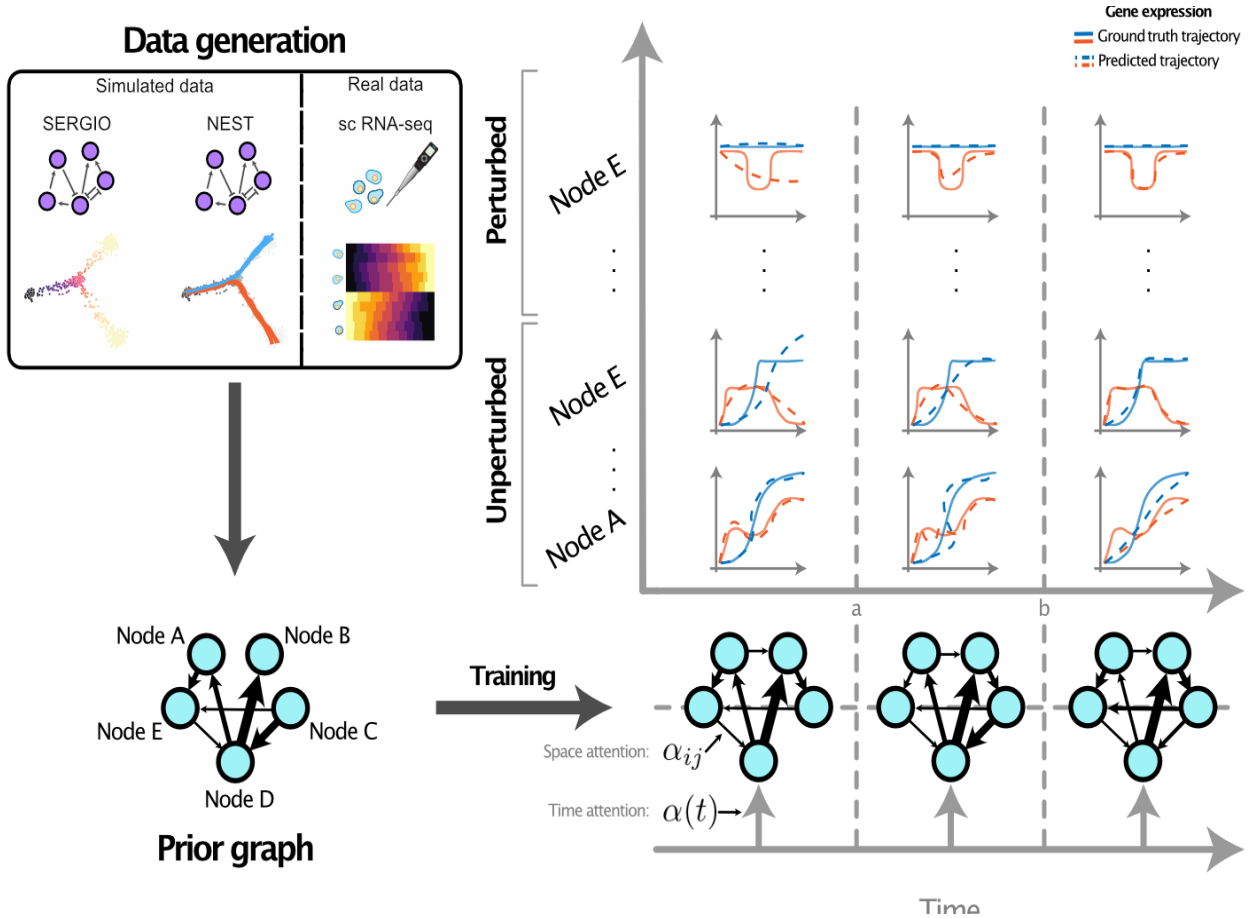

GRN inference from single-cell gene expression trajectories

RiTINI : “Regulatory Temporal Interaction Network Inference”

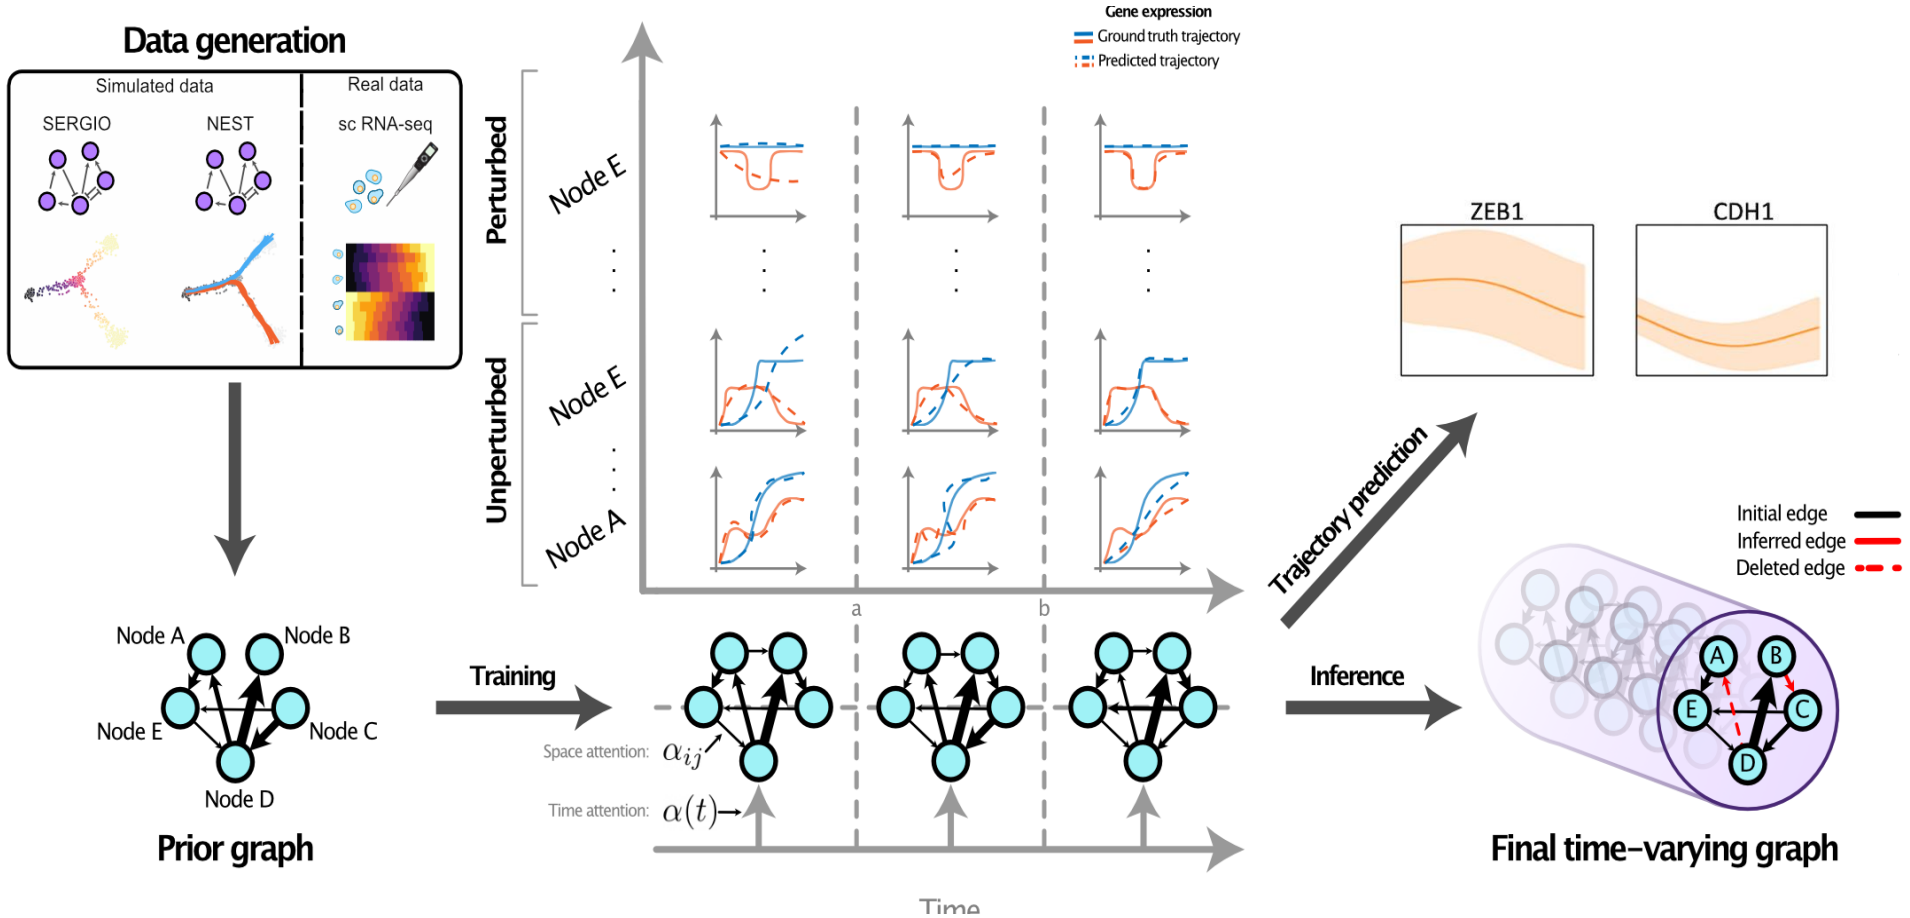

Prior graph and identifiability

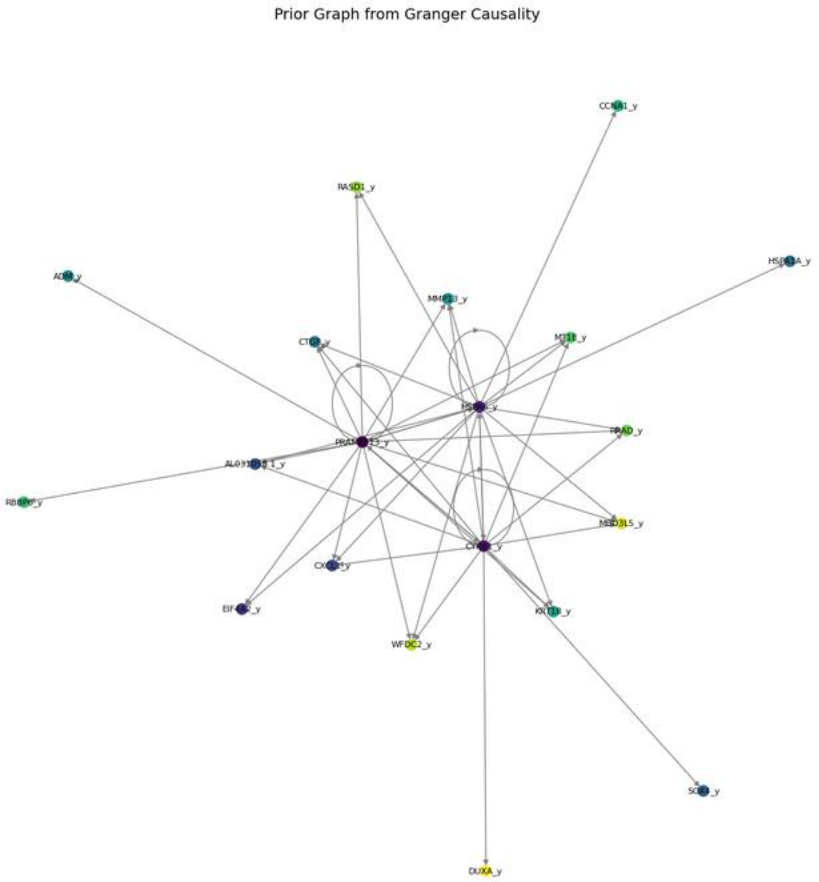

- We use granger causality as a prior graph
- We could prior knowledge networks
- This helps deal with the identifiability issue: many graphs match the same dynamics
- Final loss function will penalize for deviating from prior
- Hyperparameter for this has to be tuned, otherwise the graphs don't change

Bio-inspired Attention

$$g_i'(t) = \sum_{\delta \in [0,\tau]} l_\delta \cdot \Big(\alpha_{ii}(t) W X(v_i, t-\delta) + \sum_{j \in \mathcal{N}(i)} \alpha_{ij}(t) W X(v_j, t-\delta) \Big)$$

Time lag

Autocorrelation

Attention-over neighbors

Attention-over time

$$\alpha_{ij}(t) = \mathrm{Att}_\phi \Big(\bigcup_{\delta \in [0,\tau]} X(v_i, t-\delta) \circ X(v_j, t-\delta) \Big)$$

Time traces

Attention network

1. Attention network gives time-varying graph
2. Supports hysteresis, time attention can look at a specific point in time or over an interval
3. Outputs instantaneous derivative

Graph-ODE prediction

$$\frac{d}{dt} g_i(t) = f_\theta(g_i'(t), t)$$

The graph attention network outputs a derivative which is then integrated with an ODE solver:

$$\hat{X}(v_i, t) = \text{ODESolve}(f_\theta, g_i'(t_0), t_0, t_1) \quad \forall t \in (t_0, t_1]$$

The final loss function also enforces closeness to a prior graph $\mathcal{E}_\mathcal{P}$ and sparsity to tackle in lack of identifiability inherent in this problem:

$$\mathcal{L}(t) = \underbrace{\sum_i \|\hat{X}(v_i, t) - X(v_i, t)\|}_{\text{MSE between ground truth and predicted dynamics}} + \lambda_1 \underbrace{\sum_{i,j} \|\alpha_{ij}(t) - \mathcal{E}_\mathcal{P}\|_F}_{\text{closeness to prior}} + \underbrace{\lambda_2 |\alpha(t)|_1}_{\text{sparsity}}$$

Training with Perturbations

- "Mild" perturbations elucidate network structure
- If a small timed perturbation is applied then its effects propagate through the network
- Heavier perturbations can rewire network into new network and therefore may not give much insight on the original network!

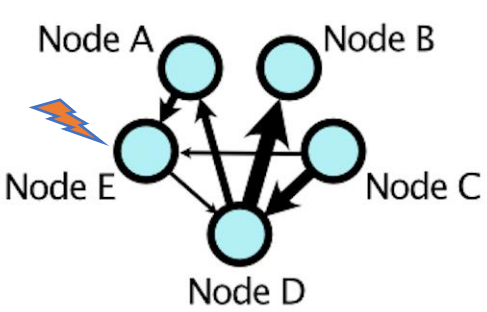

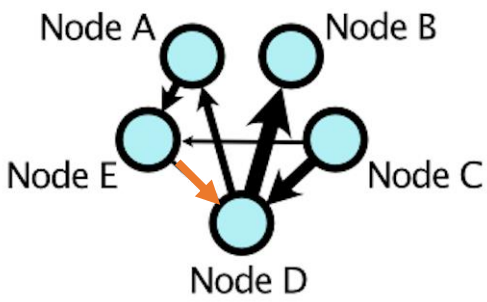

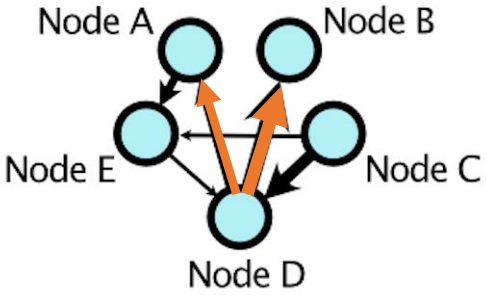

Validated with synthetic data

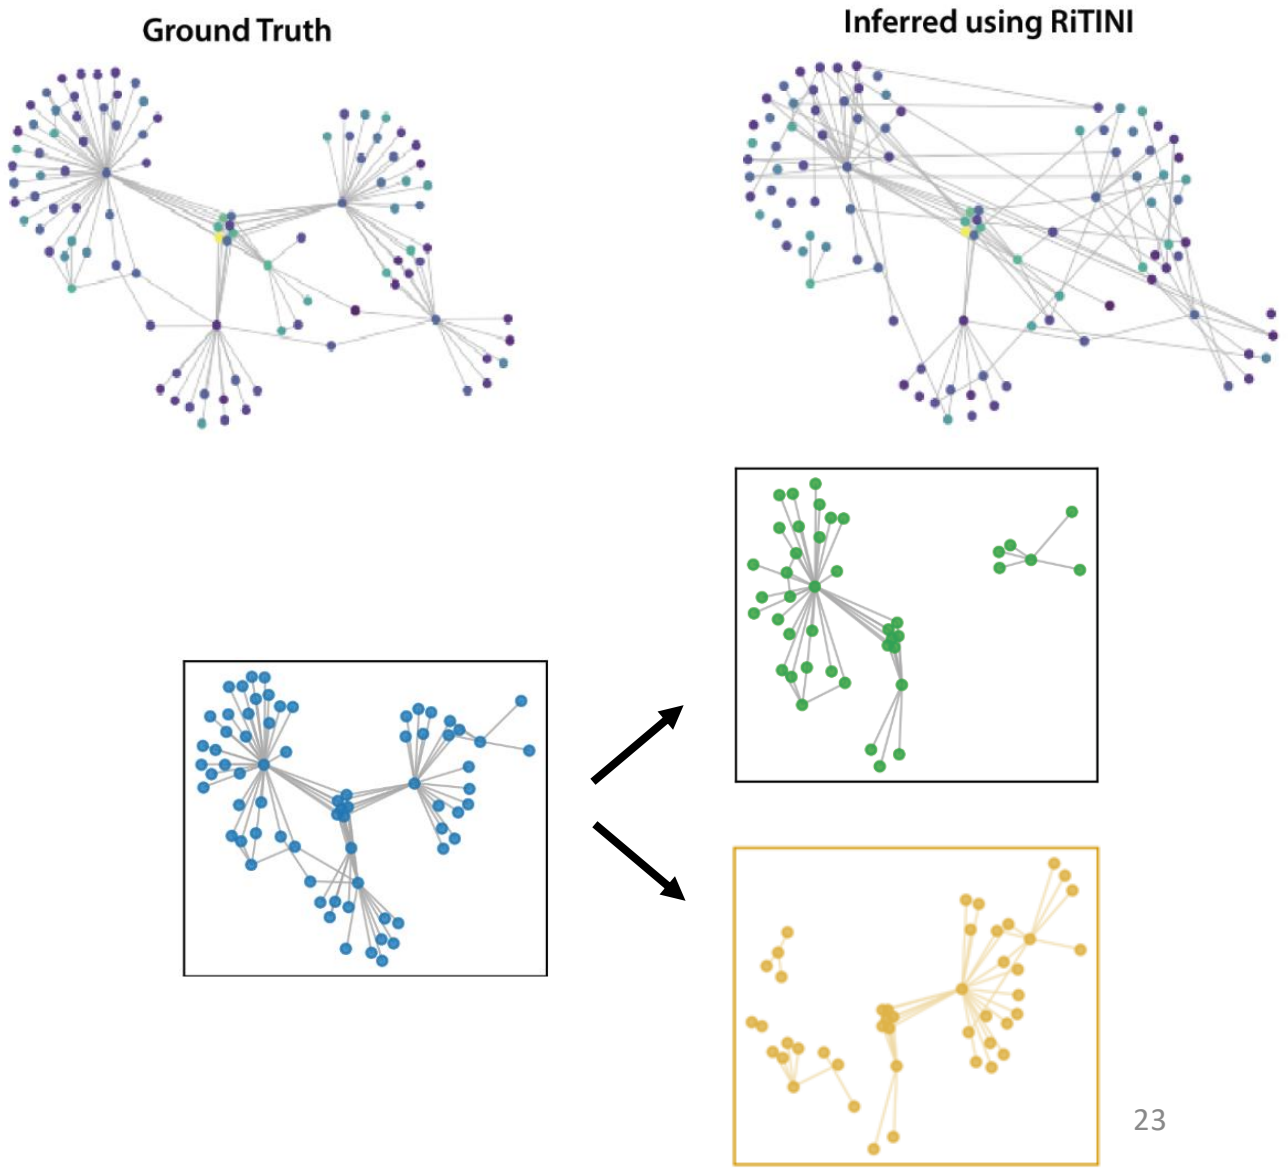

Validation using SERGIO

Method	Dataset						
	Dynamical System	Neuronal Network (NEST)			Gene Regulatory Network (SERGIO)		
$(\lvert\mathcal{V}\rvert, \lvert\mathcal{E}\rvert) =$	$(5, 5)$	$(40, 78)$	$(50, 126)$	$(75, 308)$	$(100, 137)$	$(150, 329)$	$(200, 507)$
GC [41]	2.6 ± 1.0	28.8 ± 3.3	39.2 ± 2.3	75.4 ± 9.2	51.2 ± 3.3	109.0 ± 6.4	158.8 ± 12.6
OCE [42]	4.4 ± 1.6	72.8 ± 3.5	109.6 ± 4.2	255.6 ± 7.3	138.6 ± 3.5	293.4 ± 2.9	449.8 ± 1.1
PC [43, 44]	4.4 ± 1.4	75.8 ± 2.1	117.8 ± 3.9	284.6 ± 3.2	140.4 ± 3.9	317.2 ± 3.7	495.6 ± 6.5
mTE [45]	4.6 ± 1.7	64.2 ± 3.5	100.0 ± 7.6	232.2 ± 7.1	126.4 ± 2.4	261.0 ± 2.2	397.4 ± 8.8
mMI [46]	1.8 ± 0.7	20.6 ± 5.0	$\mathbf{34.6 \pm 5.2}$	82.0 ± 3.9	51.2 ± 3.3	99.8 ± 4.0	162.8 ± 6.2
NRI [47]	$\mathbf{0.5 \pm 0.1}$	$\mathbf{18.3 \pm 4.2}$	43.7 ± 8.6	94.1 ± 3.9	72.1 ± 6.2	106.6 ± 5.4	219.8 ± 13.4
DCRNN [48]	2.2 ± 0.4	44.7 ± 3.8	81.3 ± 9.6	117.0 ± 13.8	158.14 ± 8.6	303.79 ± 12.4	508.25 ± 23.6
GTS [49]	0.8 ± 0.3	23.5 ± 1.3	76.2 ± 11.9	181.7 ± 24.2	215.4 ± 13.8	347.2 ± 19.3	481.8 ± 7.0
NIR [50]	1.3 ± 0.1	22.5 ± 2.8	39.4 ± 1.9	106.2 ± 4.7	62.7 ± 3.2	86.3 ± 2.8	159.2 ± 11.6
RiTINI	2.2 ± 1.6	25.4 ± 2.8	35.2 ± 2.3	$\mathbf{69.6 \pm 7.7}$	$\mathbf{44.6 \pm 6.2}$	$\mathbf{83.6 \pm 4.2}$	$\mathbf{128.0 \pm 4.3}$
RiTINI (w/o hysteresis)	0.6 ± 0.4	22.4 ± 1.7	38.0 ± 6.9	91.3 ± 5.2	63.6 ± 8.2	118.3 ± 9.3	205.7 ± 8.3
RiTINI (w/o neural ODE)	1.7 ± 0.3	48.1 ± 3.0	72.4 ± 8.8	129.3 ± 9.4	114.2 ± 3.1	168.0 ± 12.6	329.7 ± 22.5

Table 1: Mean and standard deviation of the graph edit distance between the inferred graph and the ground truth, across 5 different simulations with perturbations (lower is better).

Embryonic Stem Cell Differentiation

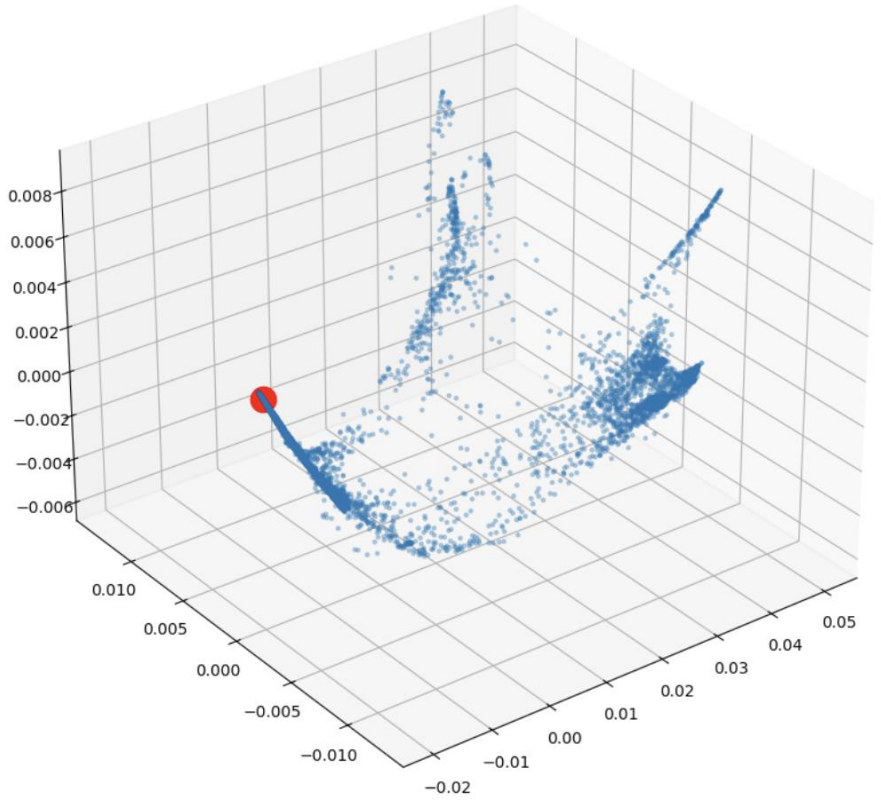

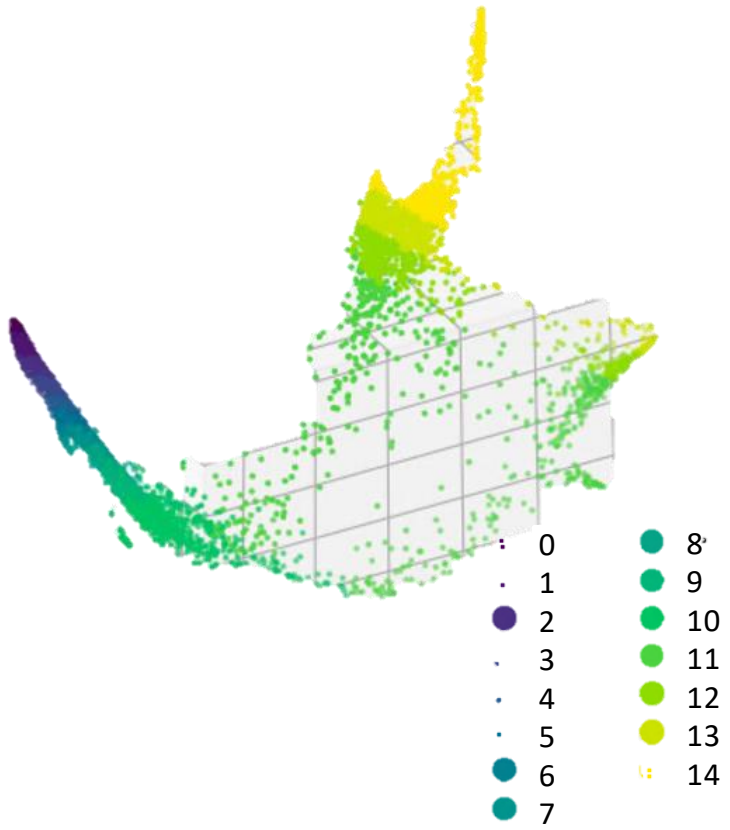

Branch Segmentation - Embryonic Stem Cell Differentiation

ESC

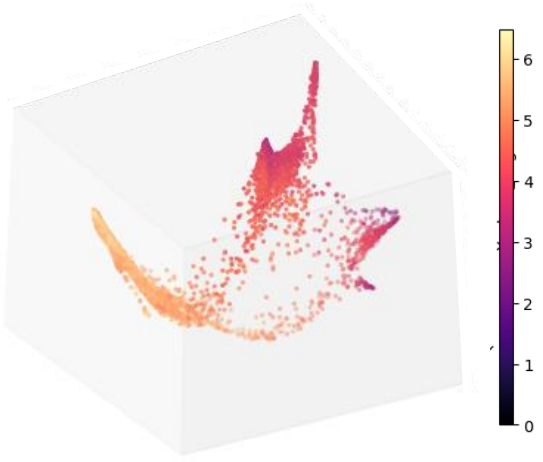

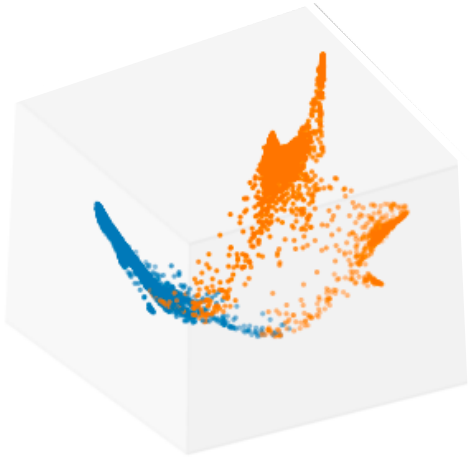

Neural Progenitors

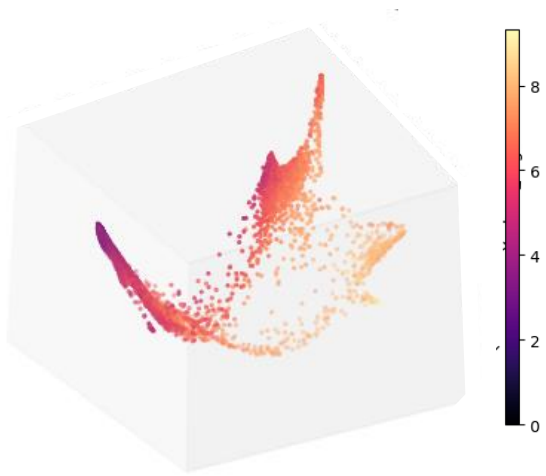

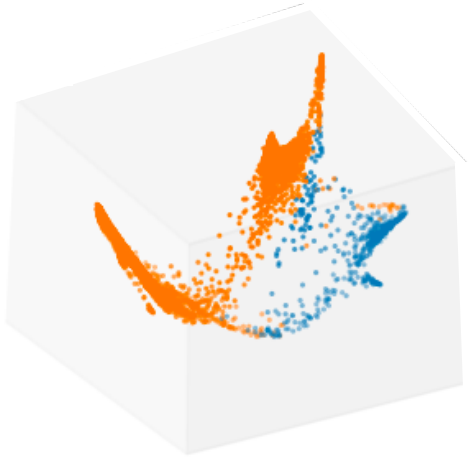

Neural Crest

MIOFlow Trajectory Inference - Embryonic Stem Cell Differentiation

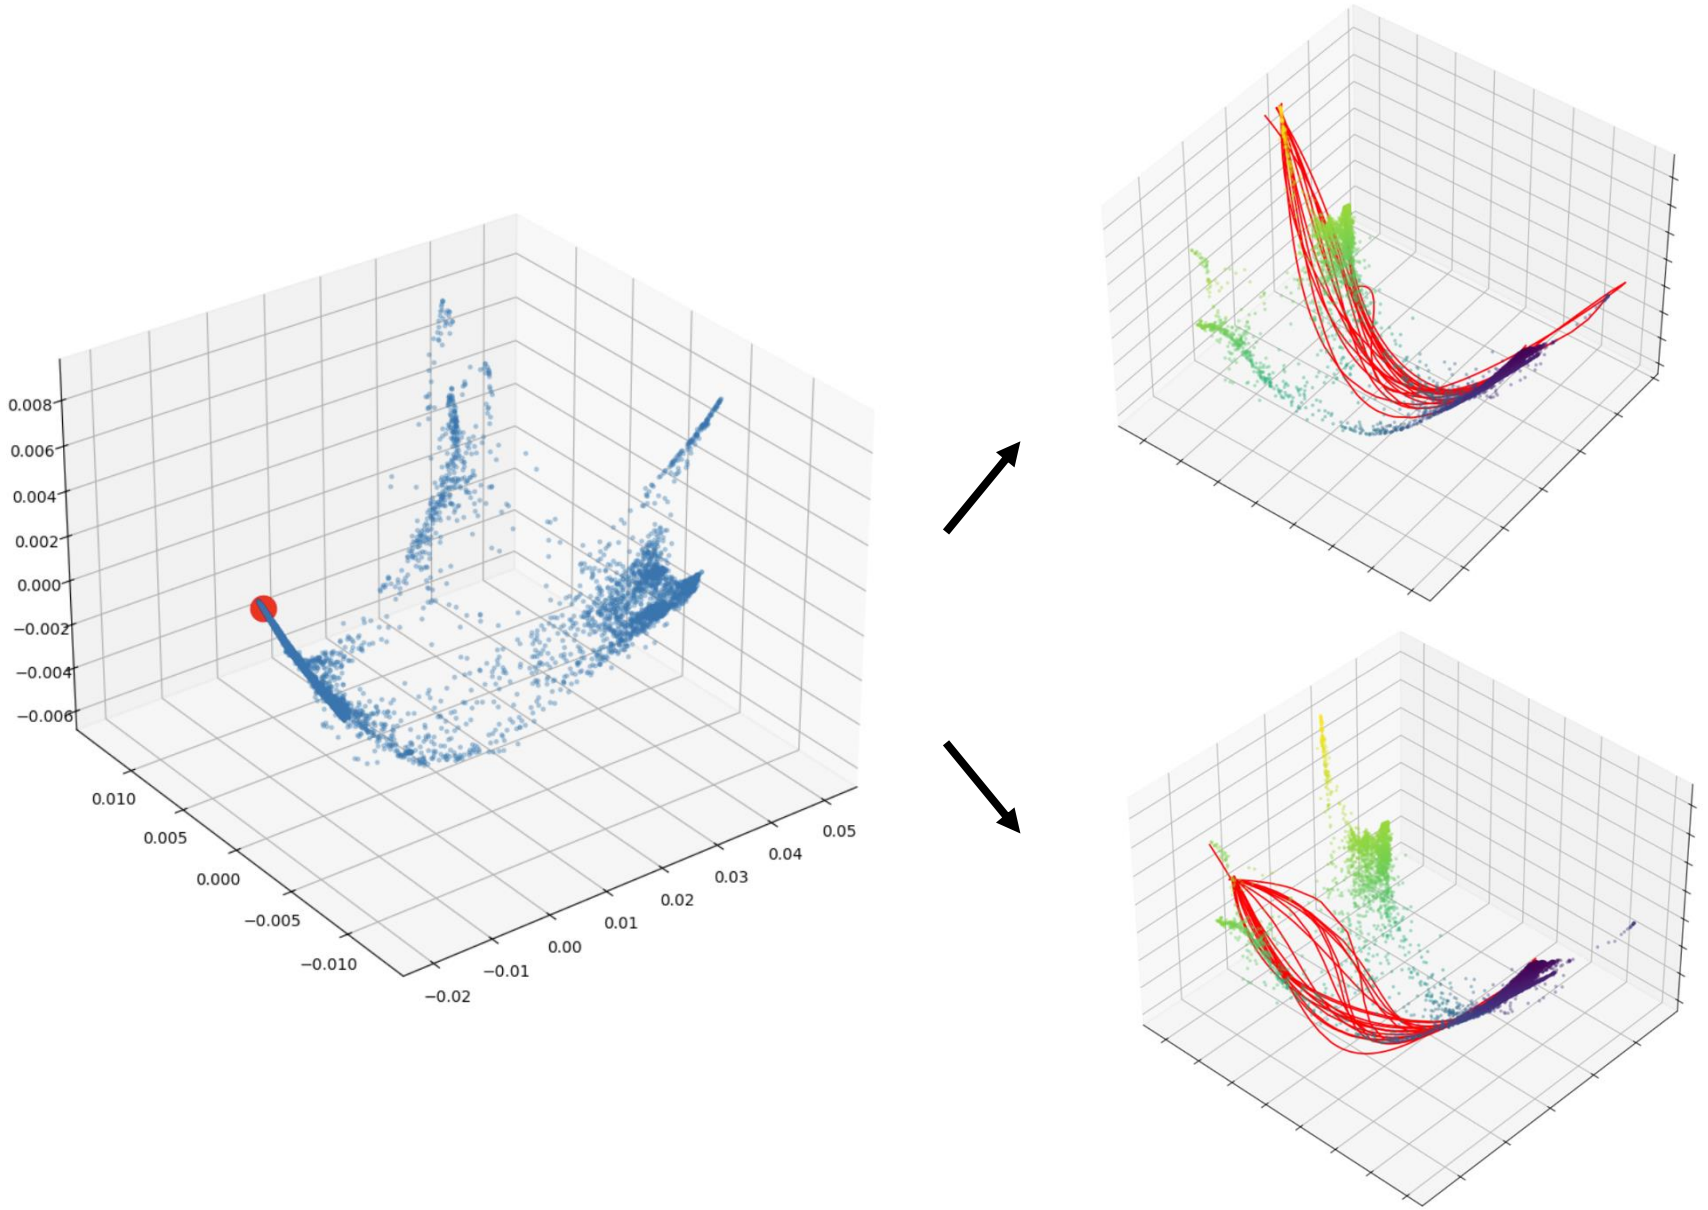

Inferred GRN

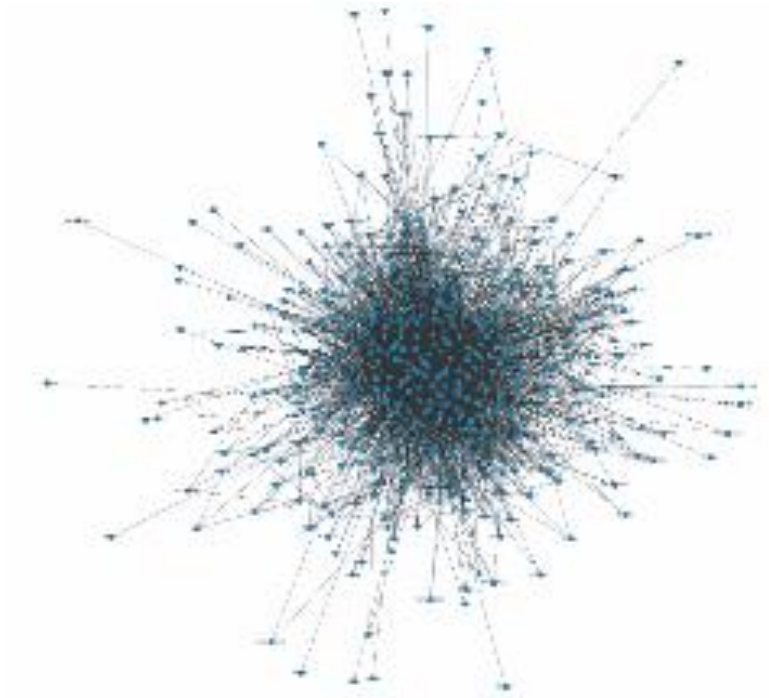

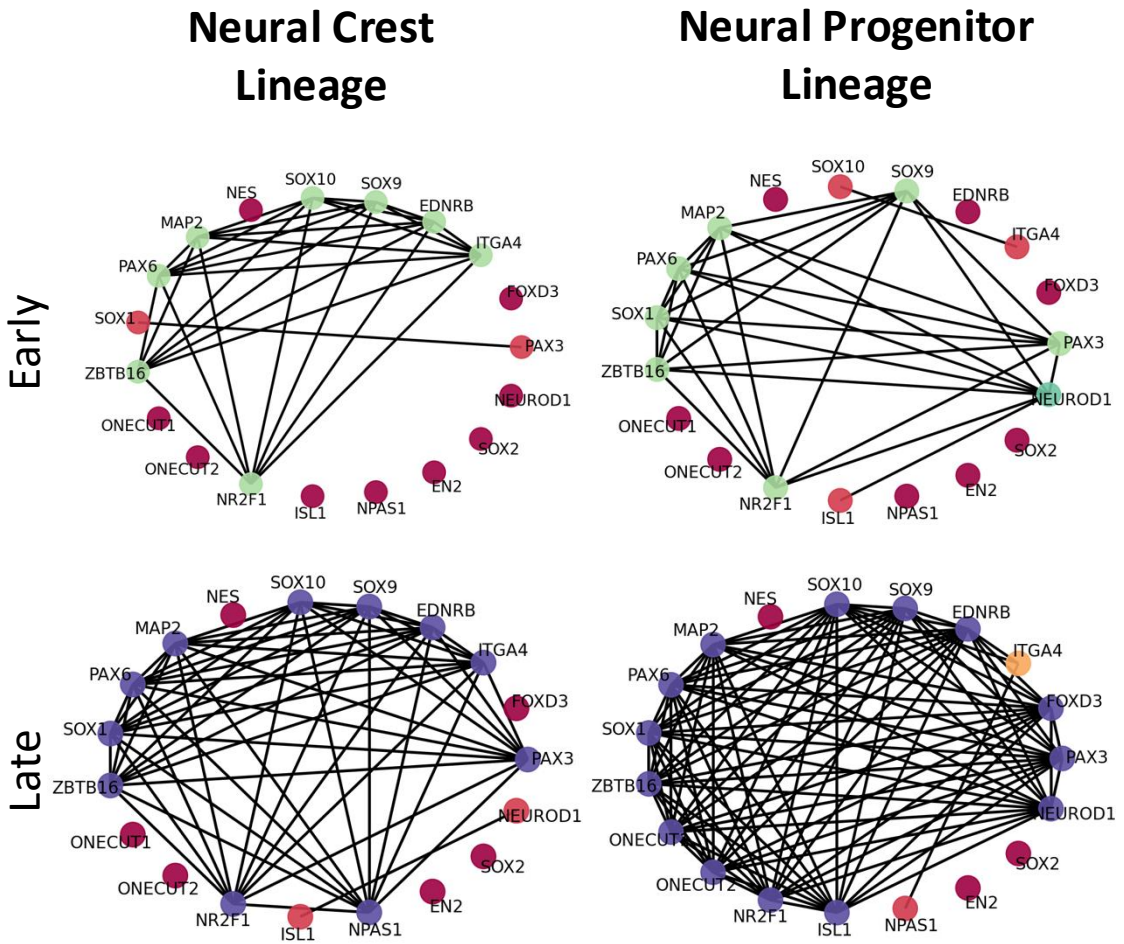

In-silico perturbations

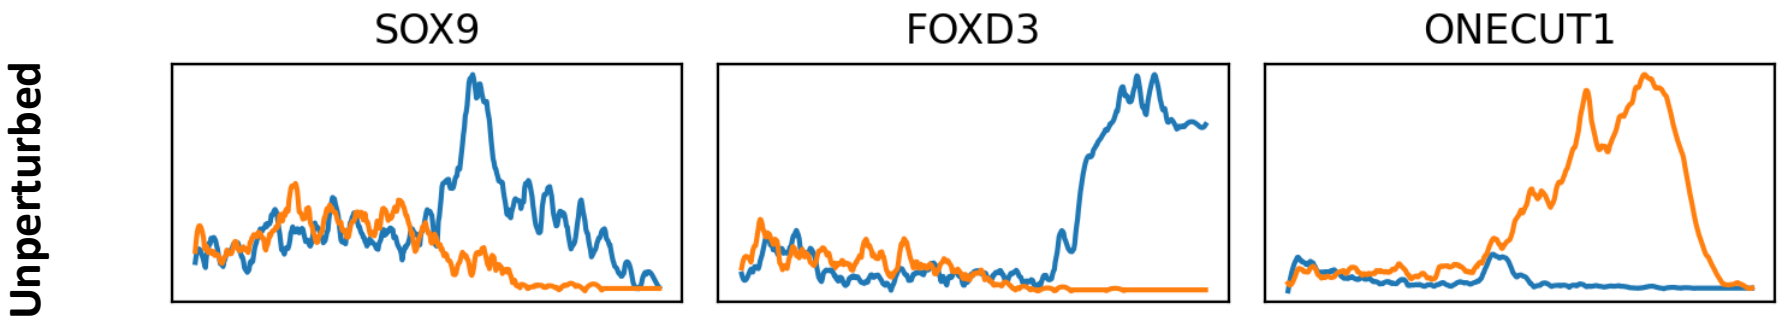

Neural Crest
Neural Prog.
knockout onset

In-silico perturbations

In-silico perturbations

Dev Biol. 2011 Dec 1;360(1):173-85. doi: 10.1016/j.ydbio.2011.09.019. Epub 2011 Sep 22.

Tfap2a and Foxd3 regulate early steps in the development of the neural crest progenitor population

Wen-Der Wang [1], David B Melville, Mercedes Montero-Balaguer, Antonis K Hatzopoulos, Ela W Knapik

Heterodimerization of TFAP2 pioneer factors drives epigenomic remodeling during neural crest specification

Megan Rothstein and Marcos Simoes-Costa[1]

Author Affiliations

- * Corresponding author; email: simoescosta@cornell.edu

ELSEVIER

Developmental Biology

Volume 304, Issue 1, 1 April 2007, Pages 338-354

Redundant activities of Tfap2a and Tfap2c are required for neural crest induction and development of other non-neural ectoderm derivatives in zebrafish embryos

Wei Li [a], Robert A. Cornell [a b]

Multiple perturbation-testing

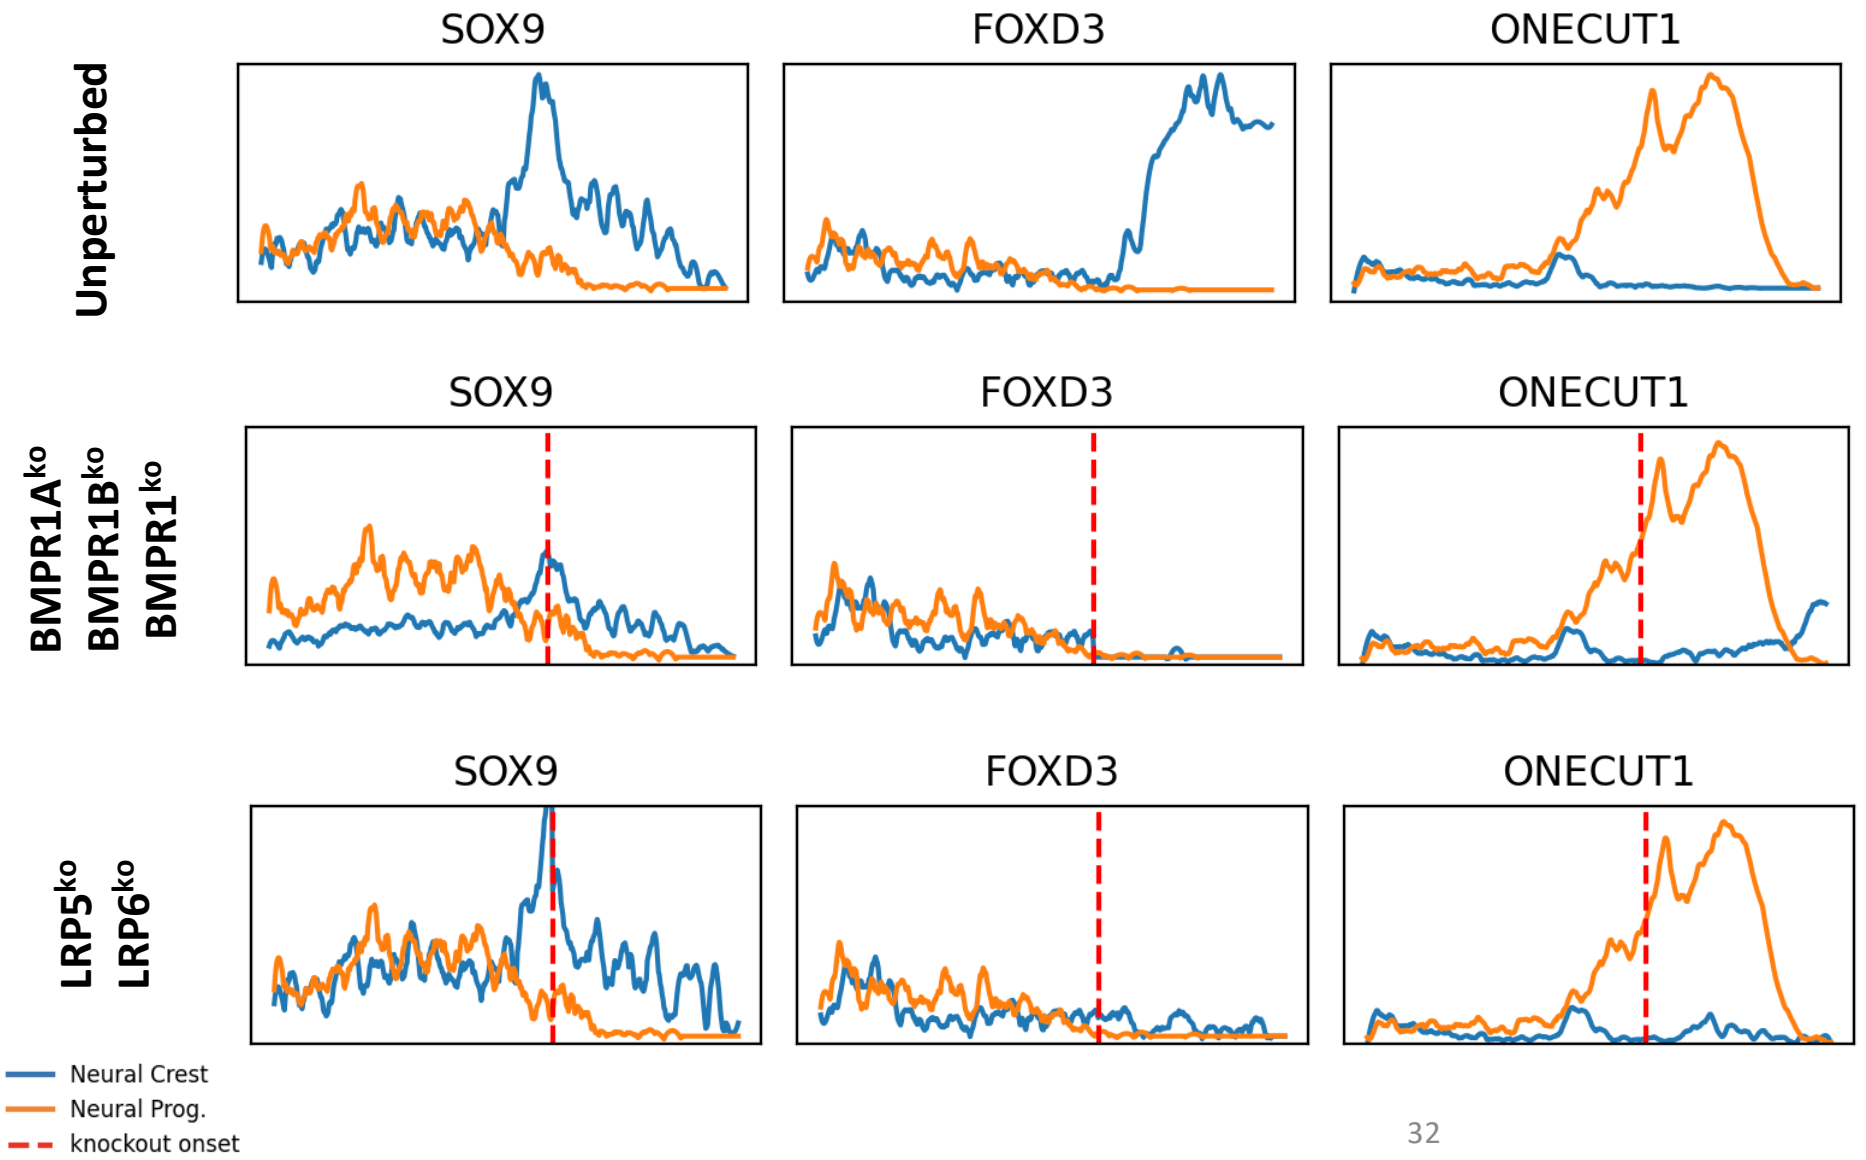

Applications to neural activity data

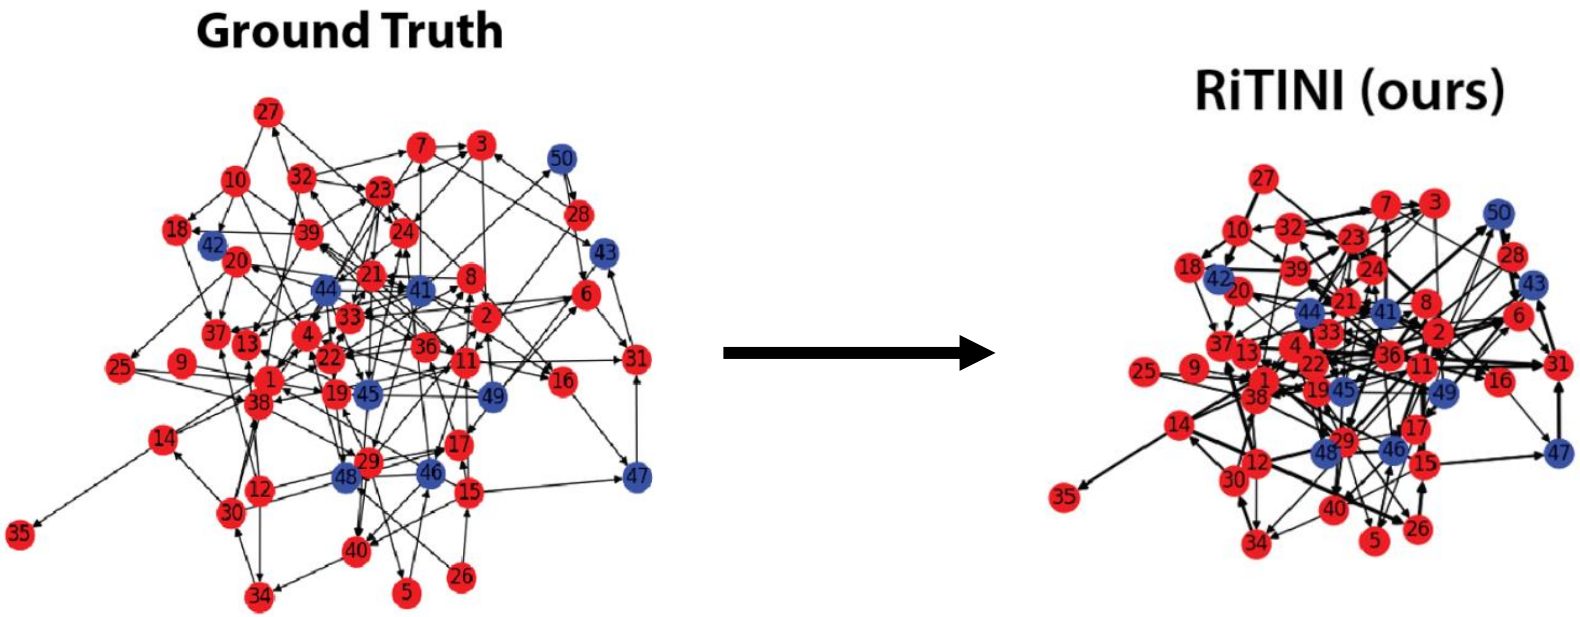

Method	Dataset						
	Dynamical System	Neuronal Network (NEST)			Gene Regulatory Network (SERGIO)		
$(\lvert\mathcal{V}\rvert, \lvert\mathcal{E}\rvert) =$	$(5, 5)$	$(40, 78)$	$(50, 126)$	$(75, 308)$	$(100, 137)$	$(150, 329)$	$(200, 507)$
GC [41]	2.6 ± 1.0	28.8 ± 3.3	39.2 ± 2.3	75.4 ± 9.2	51.2 ± 3.3	109.0 ± 6.4	158.8 ± 12.6
OCE [42]	4.4 ± 1.6	72.8 ± 3.5	109.6 ± 4.2	255.6 ± 7.3	138.6 ± 3.5	293.4 ± 2.9	449.8 ± 1.1
PC [43, 44]	4.4 ± 1.4	75.8 ± 2.1	117.8 ± 3.9	284.6 ± 3.2	140.4 ± 3.9	317.2 ± 3.7	495.6 ± 6.5
mTE [45]	4.6 ± 1.7	64.2 ± 3.5	100.0 ± 7.6	232.2 ± 7.1	126.4 ± 2.4	261.0 ± 2.2	397.4 ± 8.8
mMI [46]	1.8 ± 0.7	20.6 ± 5.0	$\mathbf{34.6 \pm 5.2}$	82.0 ± 3.9	51.2 ± 3.3	99.8 ± 4.0	162.8 ± 6.2
NRI [47]	$\mathbf{0.5 \pm 0.1}$	$\mathbf{18.3 \pm 4.2}$	43.7 ± 8.6	94.1 ± 3.9	72.1 ± 6.2	106.6 ± 5.4	219.8 ± 13.4
DCRNN [48]	2.2 ± 0.4	44.7 ± 3.8	81.3 ± 9.6	117.0 ± 13.8	158.14 ± 8.6	303.79 ± 12.4	508.25 ± 23.6
GTS [49]	0.8 ± 0.3	23.5 ± 1.3	76.2 ± 11.9	181.7 ± 24.2	215.4 ± 13.8	347.2 ± 19.3	481.8 ± 7.0
NIR [50]	1.3 ± 0.1	22.5 ± 2.8	39.4 ± 1.9	106.2 ± 4.7	62.7 ± 3.2	86.3 ± 2.8	159.2 ± 11.6
RiTINI	2.2 ± 1.6	25.4 ± 2.8	35.2 ± 2.3	$\mathbf{69.6 \pm 7.7}$	$\mathbf{44.6 \pm 6.2}$	$\mathbf{83.6 \pm 4.2}$	$\mathbf{128.0 \pm 4.3}$
RiTINI (w/o hysteresis)	0.6 ± 0.4	22.4 ± 1.7	38.0 ± 6.9	91.3 ± 5.2	63.6 ± 8.2	118.3 ± 9.3	205.7 ± 8.3
RiTINI (w/o neural ODE)	1.7 ± 0.3	48.1 ± 3.0	72.4 ± 8.8	129.3 ± 9.4	114.2 ± 3.1	168.0 ± 12.6	329.7 ± 22.5

Table 1: Mean and standard deviation of the graph edit distance between the inferred graph and the ground truth, across 5 different simulations with perturbations (lower is better).

Inferring Dynamic Regulatory Interaction Graphs From Time Series Data With Perturbations

Dhananjay Bhaskar, Daniel Sumner Magruder, Matheo Morales, Edward De Brouwer, Aarthi Venkat, Frederik Wenkel, James Noonan, Guy Wolf, Natalia Ivanova, Smita Krishnaswamy

Proceedings of the Second Learning on Graphs Conference, PMLR 231:22:1-22:21, 2024.

Creating a landscape of perturbations and cancer conditions

ML on high dimensional clouds

- We now collect datasets as we once used to collect data points

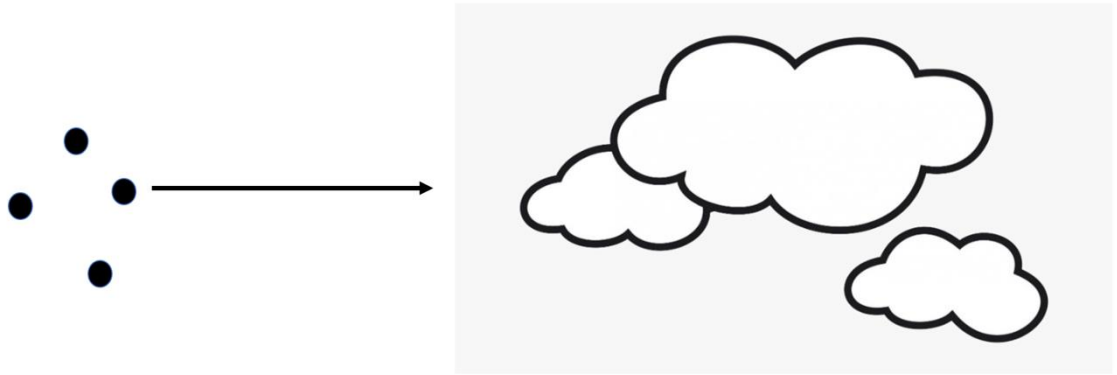

- We need ways of reasoning about these point clouds
  - want to understand the state space

Perturbation screens

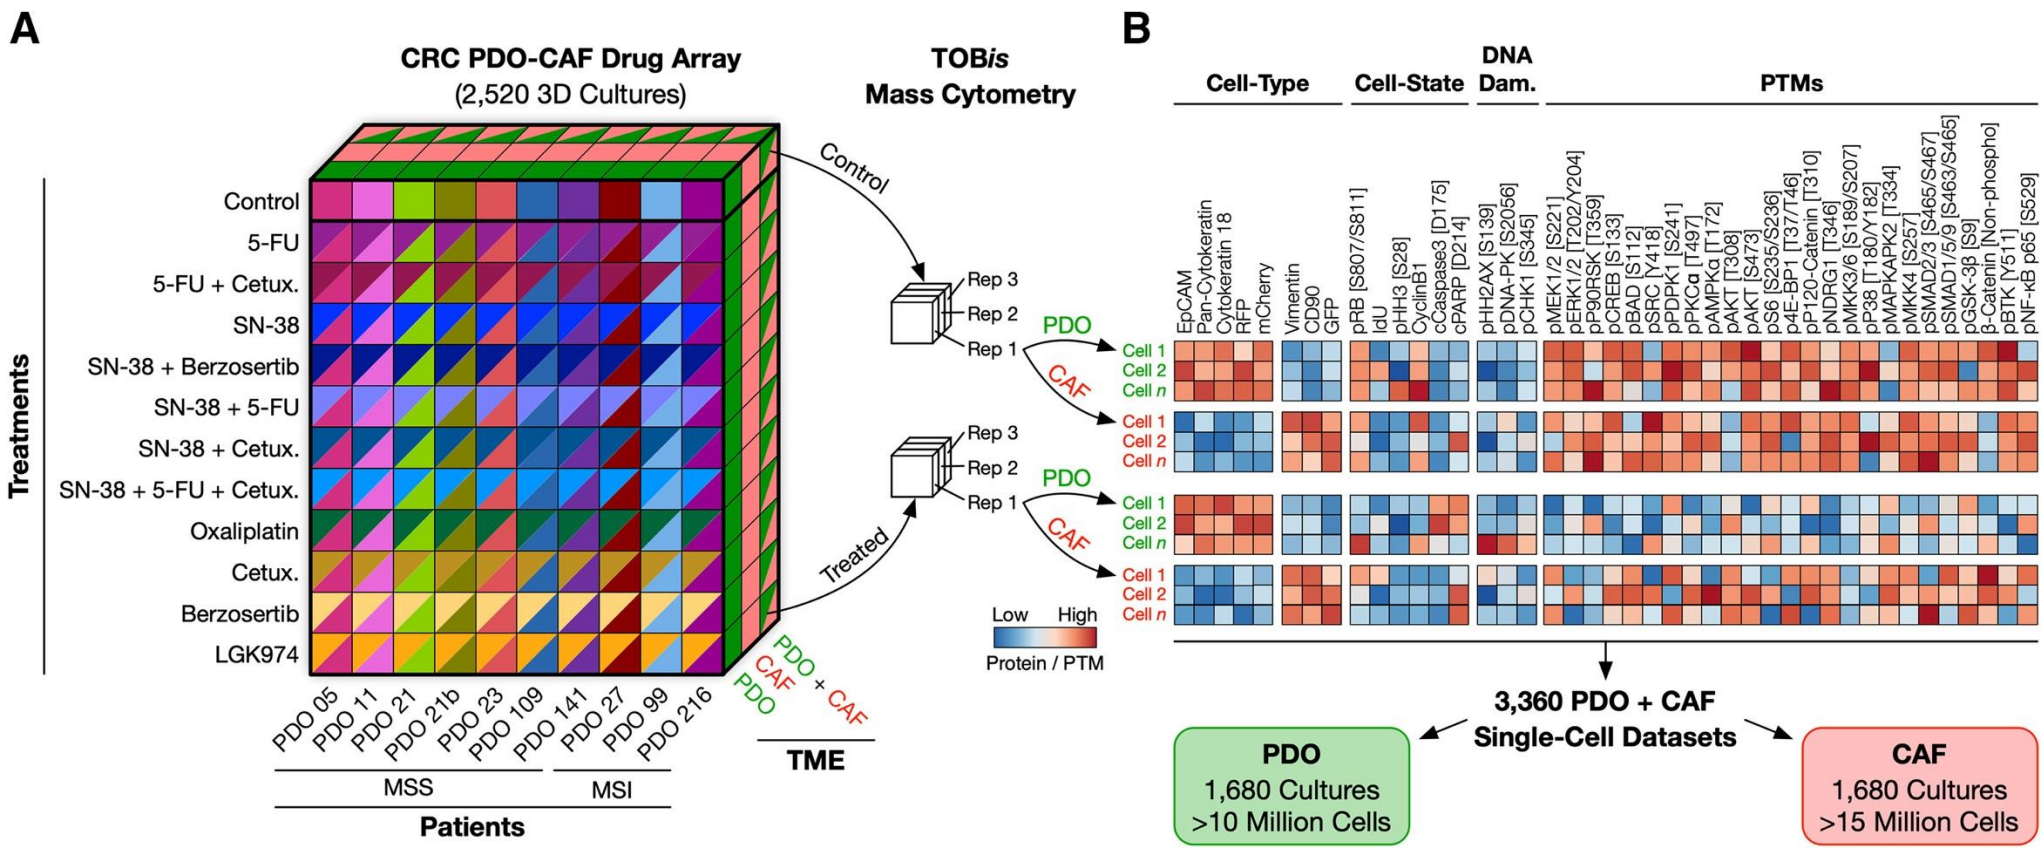

Trellis Method

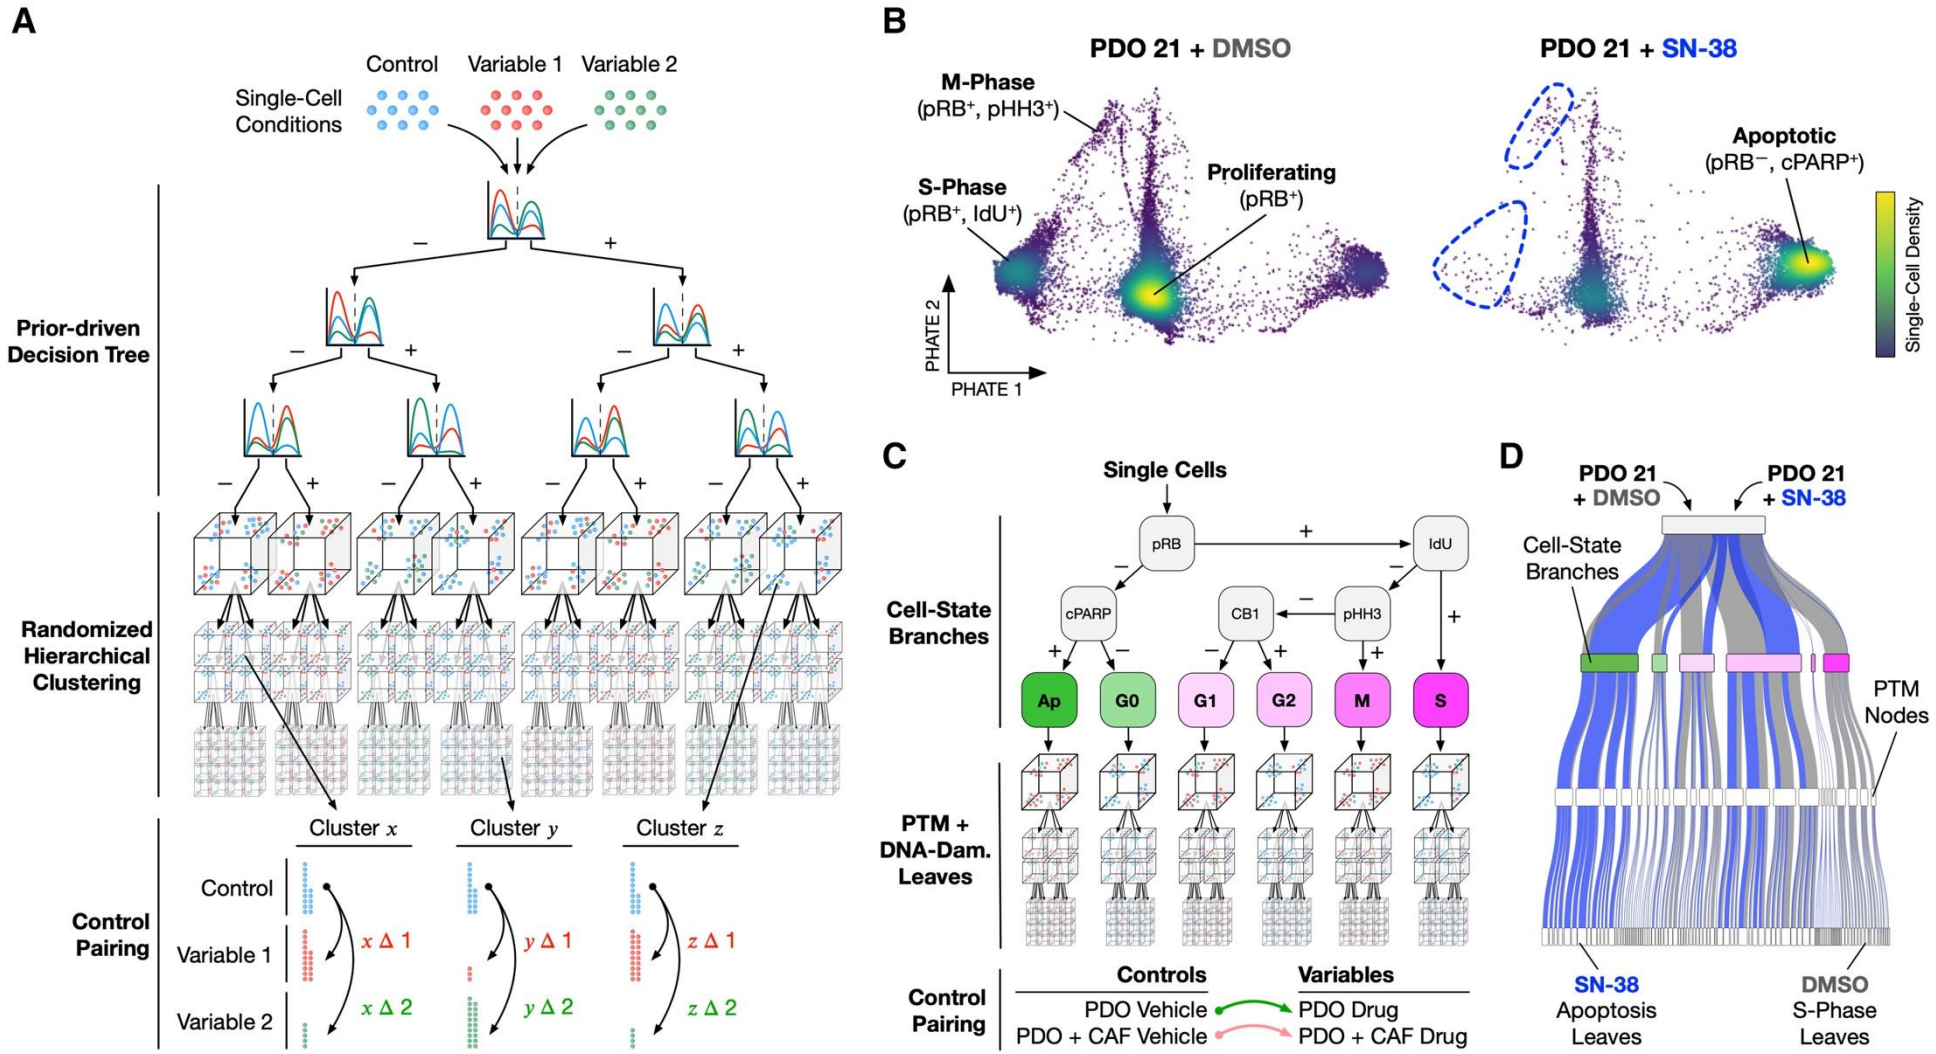

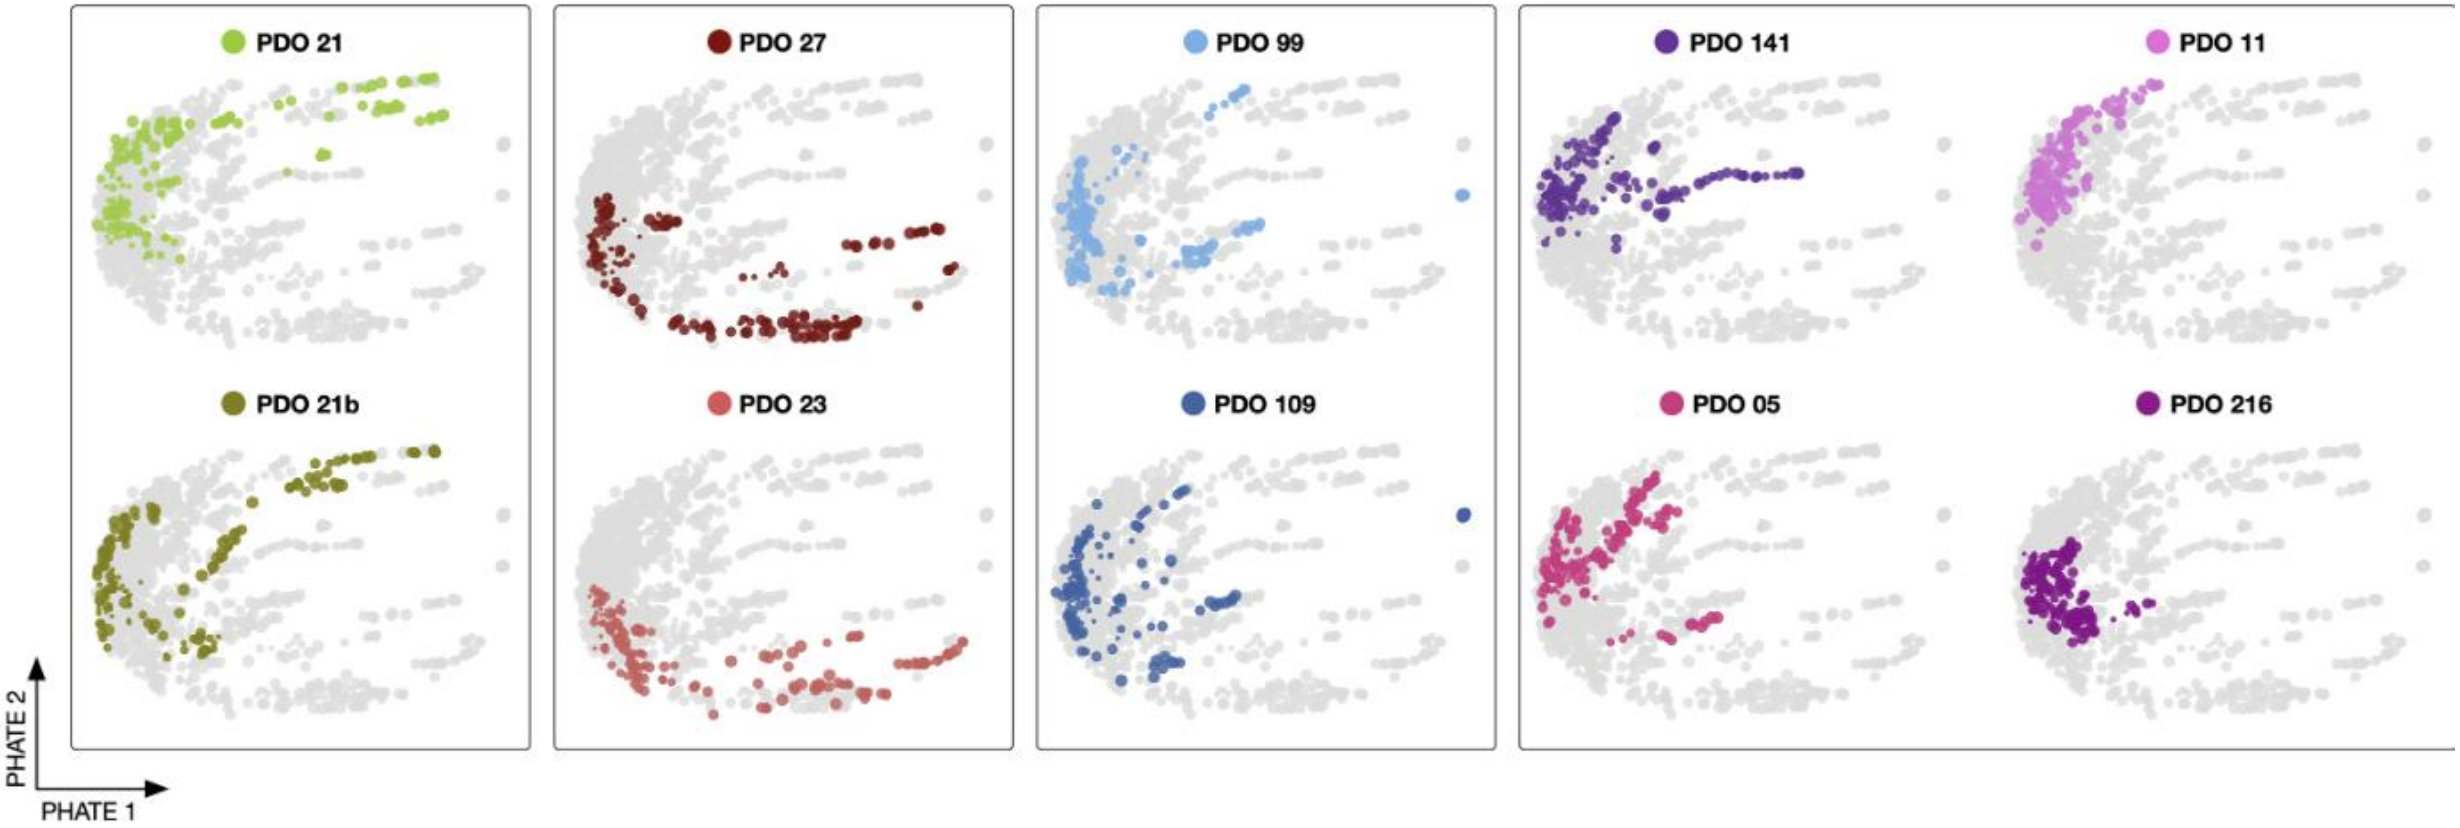

Each dot here is an entire tumor and it is in a map with other tumors

This static map can be zoomed in to give “google map directions” with an approach like Mioflow

DiffusionEMD

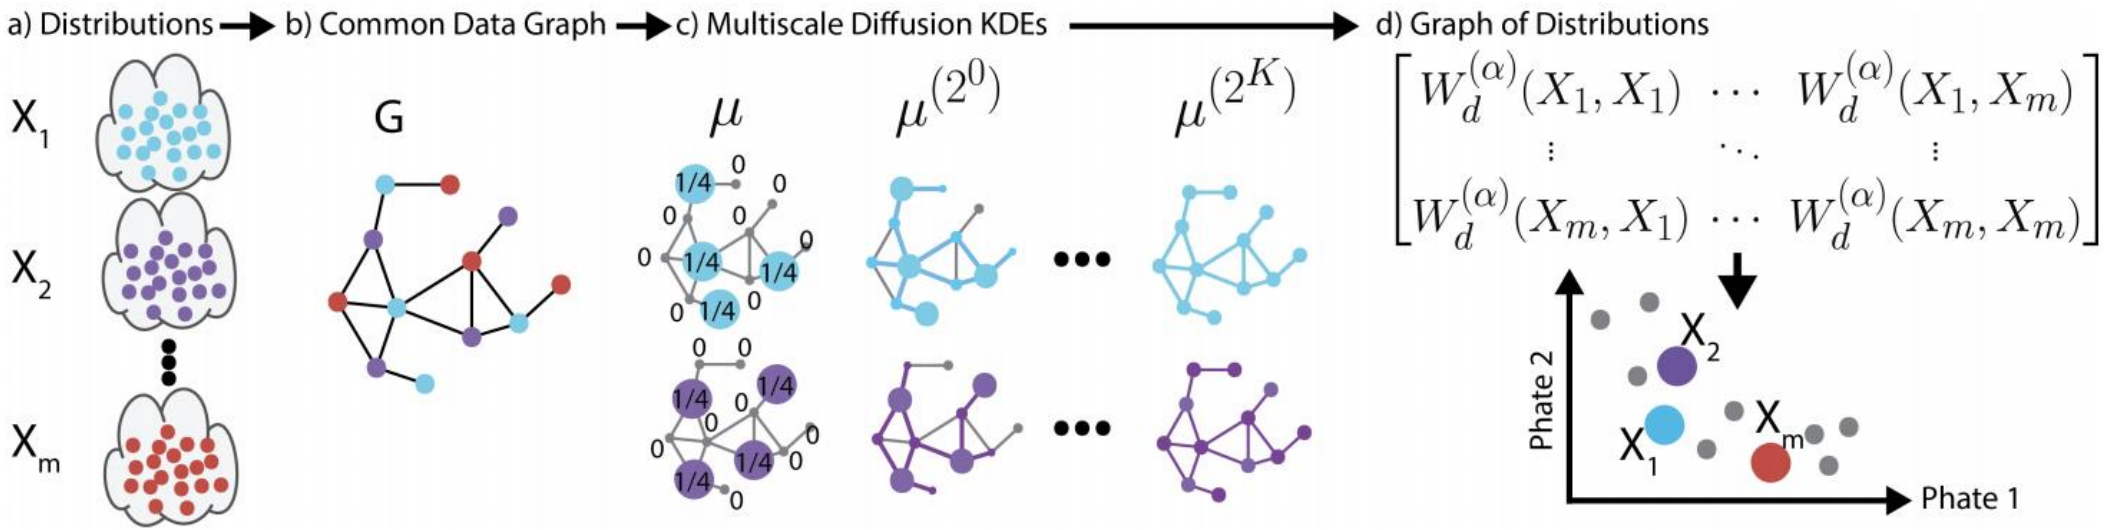

- Each perturbation is defined by distribution of cellular states
- DiffusionEMD measures distances between perturbation distributions on a common graph by using multiscale diffusion density estimates

Tong et al. ICML, 2021

Mapping the landscape of all perturbations

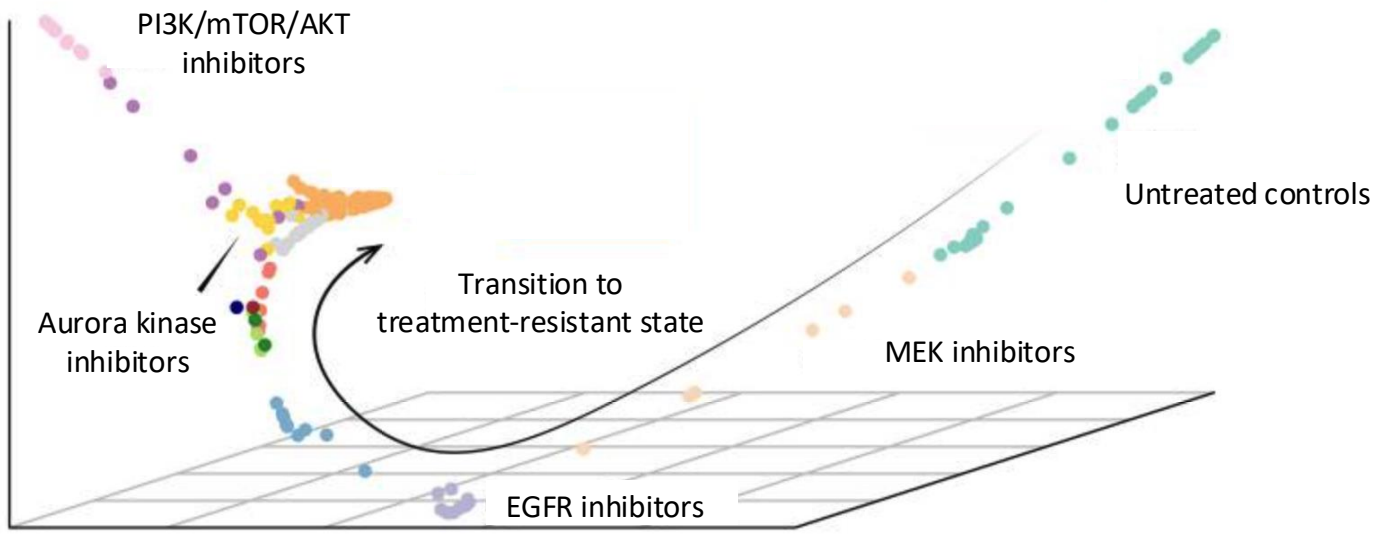

- Each point here is a whole dataset
- Distances between points are optimal transport distances
- We usually compute these using our fast methods like Diffusion EMD [Tong et al. ICML 2021] or Geodesic Sinkhorn [Huguet, IEEE MLSP 2023]

RESOURCE · Volume 186, Issue 25, P5606-5619.E24, December 07, 2023 · *Open Access* Download Full Issue

Trellis tree-based analysis reveals stromal regulation of patient-derived organoid drug responses

María Ramos Zapatero [1,11] · Alexander Tong [2,3,4,11] · James W. Opzoomer [1] · … · Daniel Hochhauser [6] · Smita Krishnaswamy [2,7,8,9,10,12] · Christopher J. Tape [1,12,13] … Show more

Affiliations & Notes · Article Info · Linked Articles (2)

Diffusion Earth Mover's Distance and Distribution Embeddings

Alexander Tong [*1] **Guillaume Huguet** [*2 3] **Amine Natik** [*2 3] **Kincaid MacDonald** [4] **Manik Kuchroo** [5 1] **Ronald R. Coifman** [4] **Guy Wolf** [†2 3] **Smita Krishnaswamy** [†5 1]

Article | Published: 13 January 2020

Uncovering axes of variation among single-cell cancer specimens

William S. Chen, Nevena Zivanovic, David van Dijk, Guy Wolf, Bernd Bodenmiller & Smita Krishnaswamy

Nature Methods **17**, 302–310 (2020) | Cite this article

STAGED

(Rocha et al., Graph Signal Processing 2025, MLCB 2025)

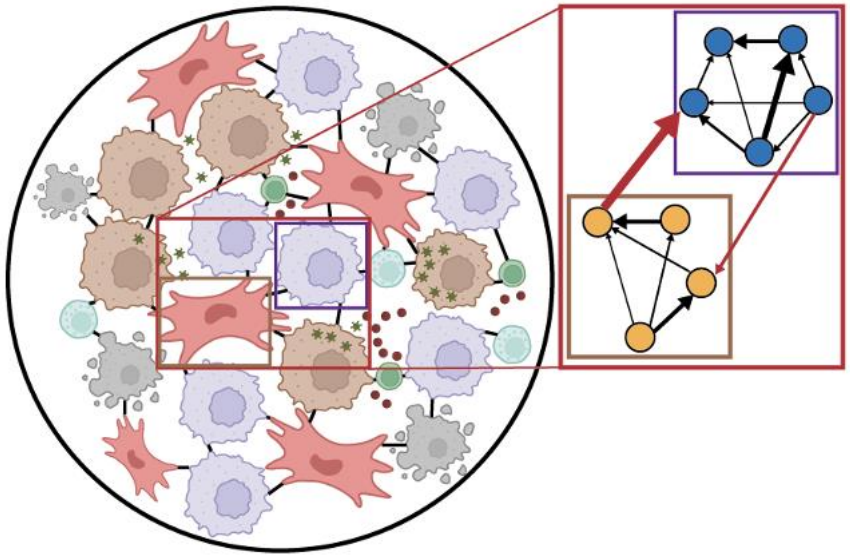

Main idea: A multi-agent neural network for learning cellular dynamics

Joao Rocha
Ke Xu
Xingzhi Sun
Ananya Krishna
Dhananjay Bhaskar
Blanche Mungeon
Morgan Craig
Mark Gerstein

Motivation: Cells aren't loners

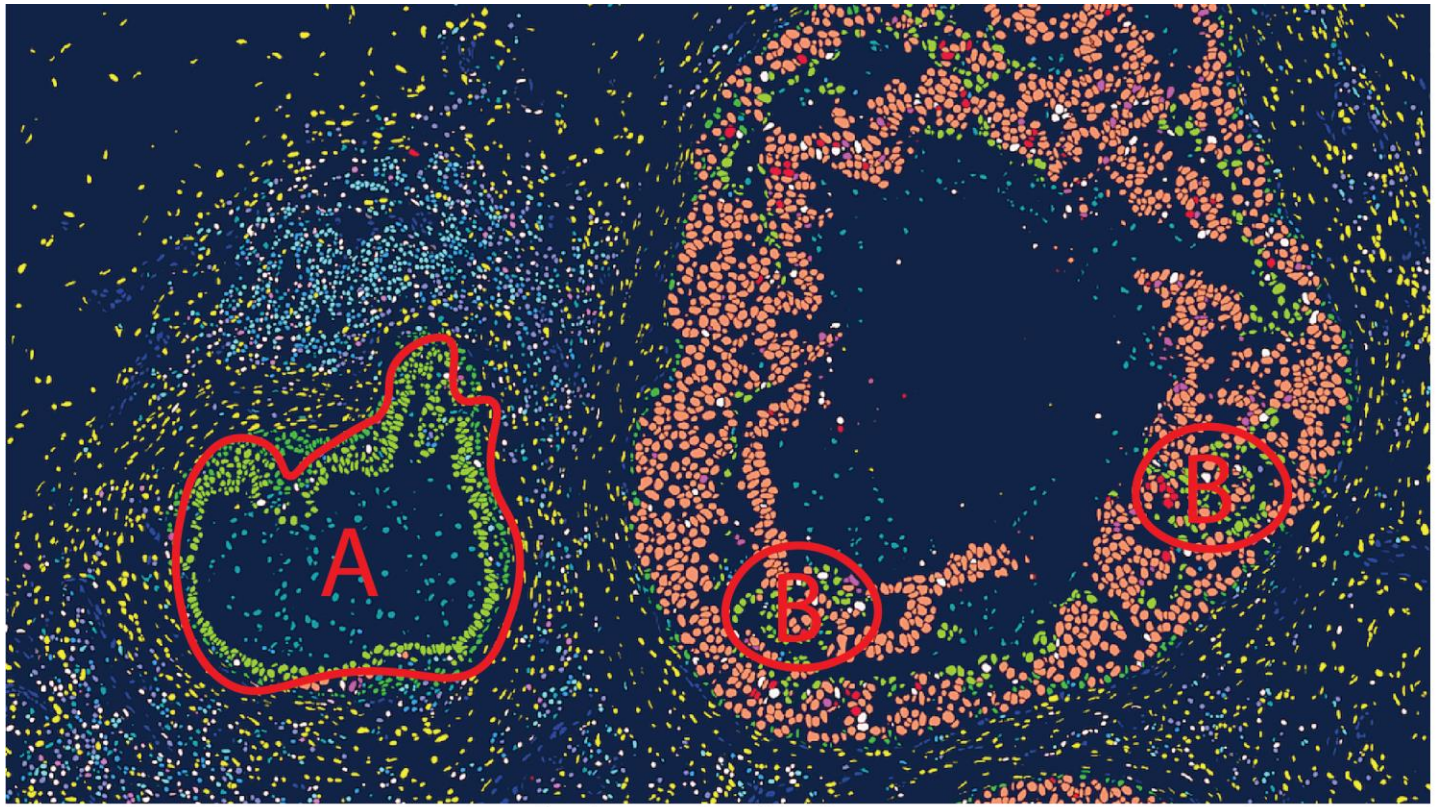

https://ostr.ccr.cancer.gov/emerging-technologies/spatial-biology/xenium/ Accessed: Mar. 04, 2025

- **Spatial Transcriptomics**: Gives information on both the space and gene-level.
- **Cells** are agents interacting via biochemical signals.
- **Microenvironment** drives biological processes.

Spatial Temporal Agent-Based Graph Evolution Dynamics (STAGED)

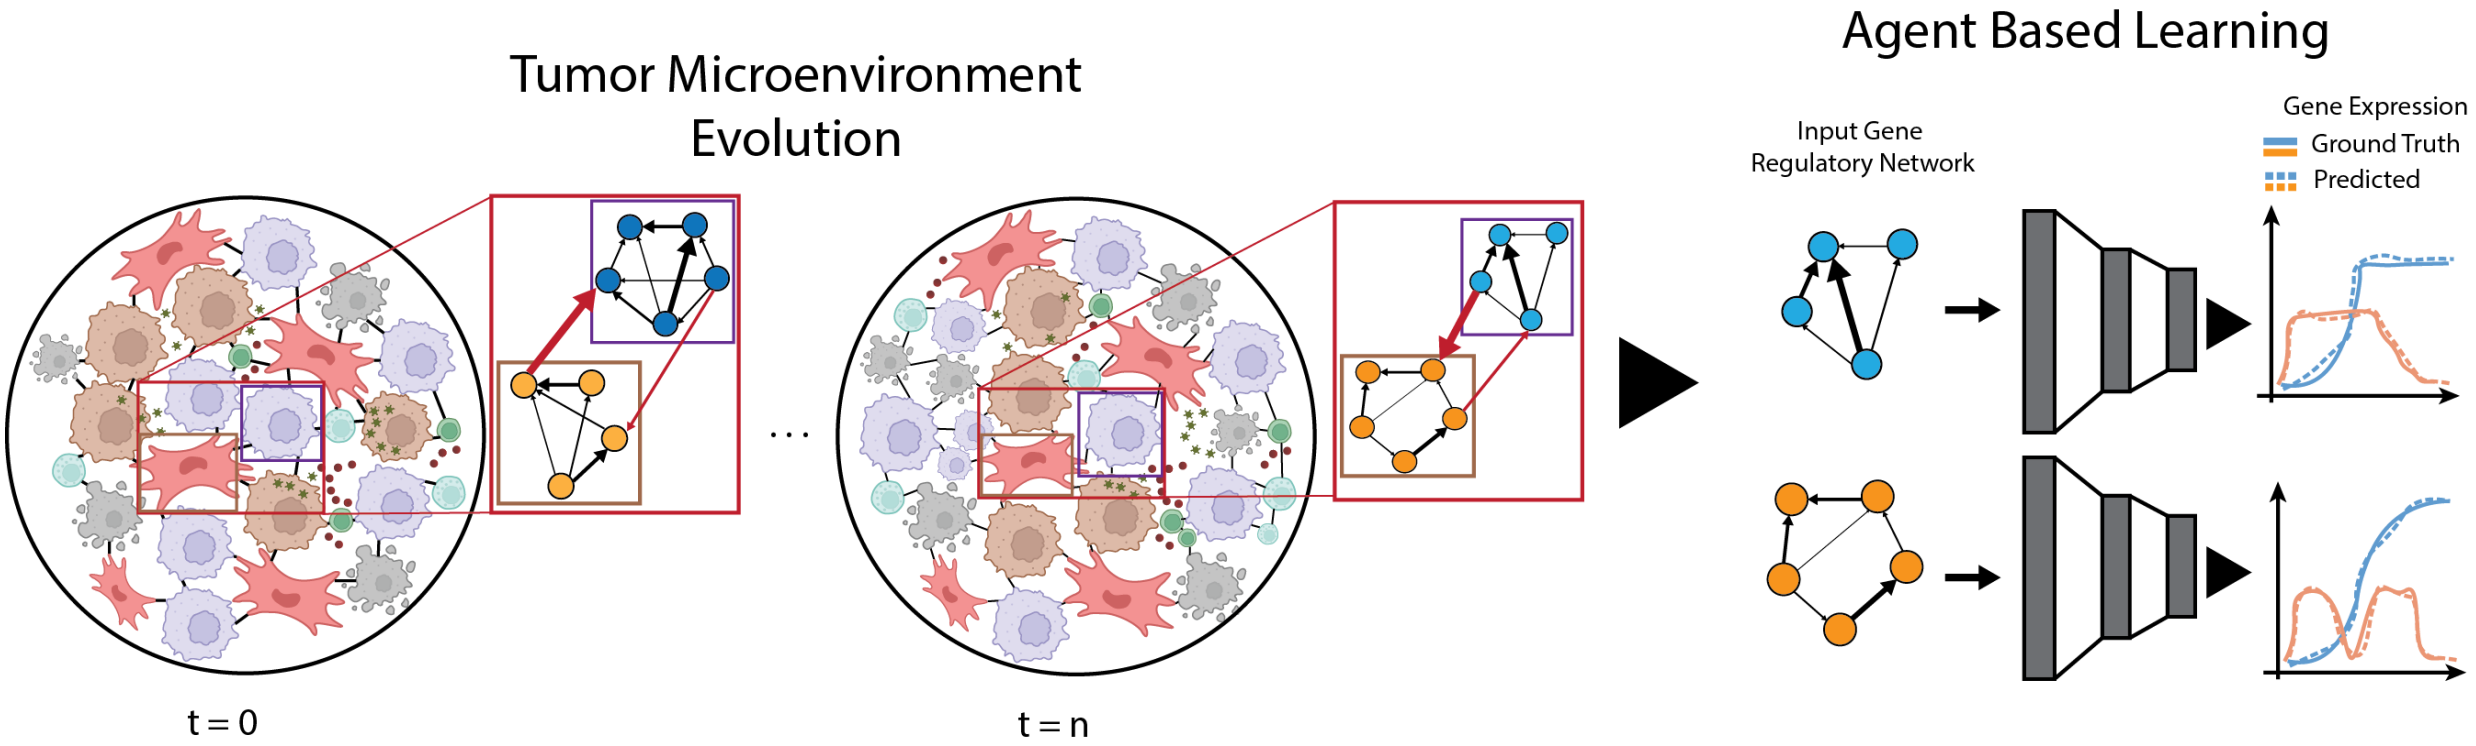

- **Level 1**: Cell-cell communication (agents talking)
- **Level 2**: Gene regulatory networks (within each cell)

Network of graph attention ODEs

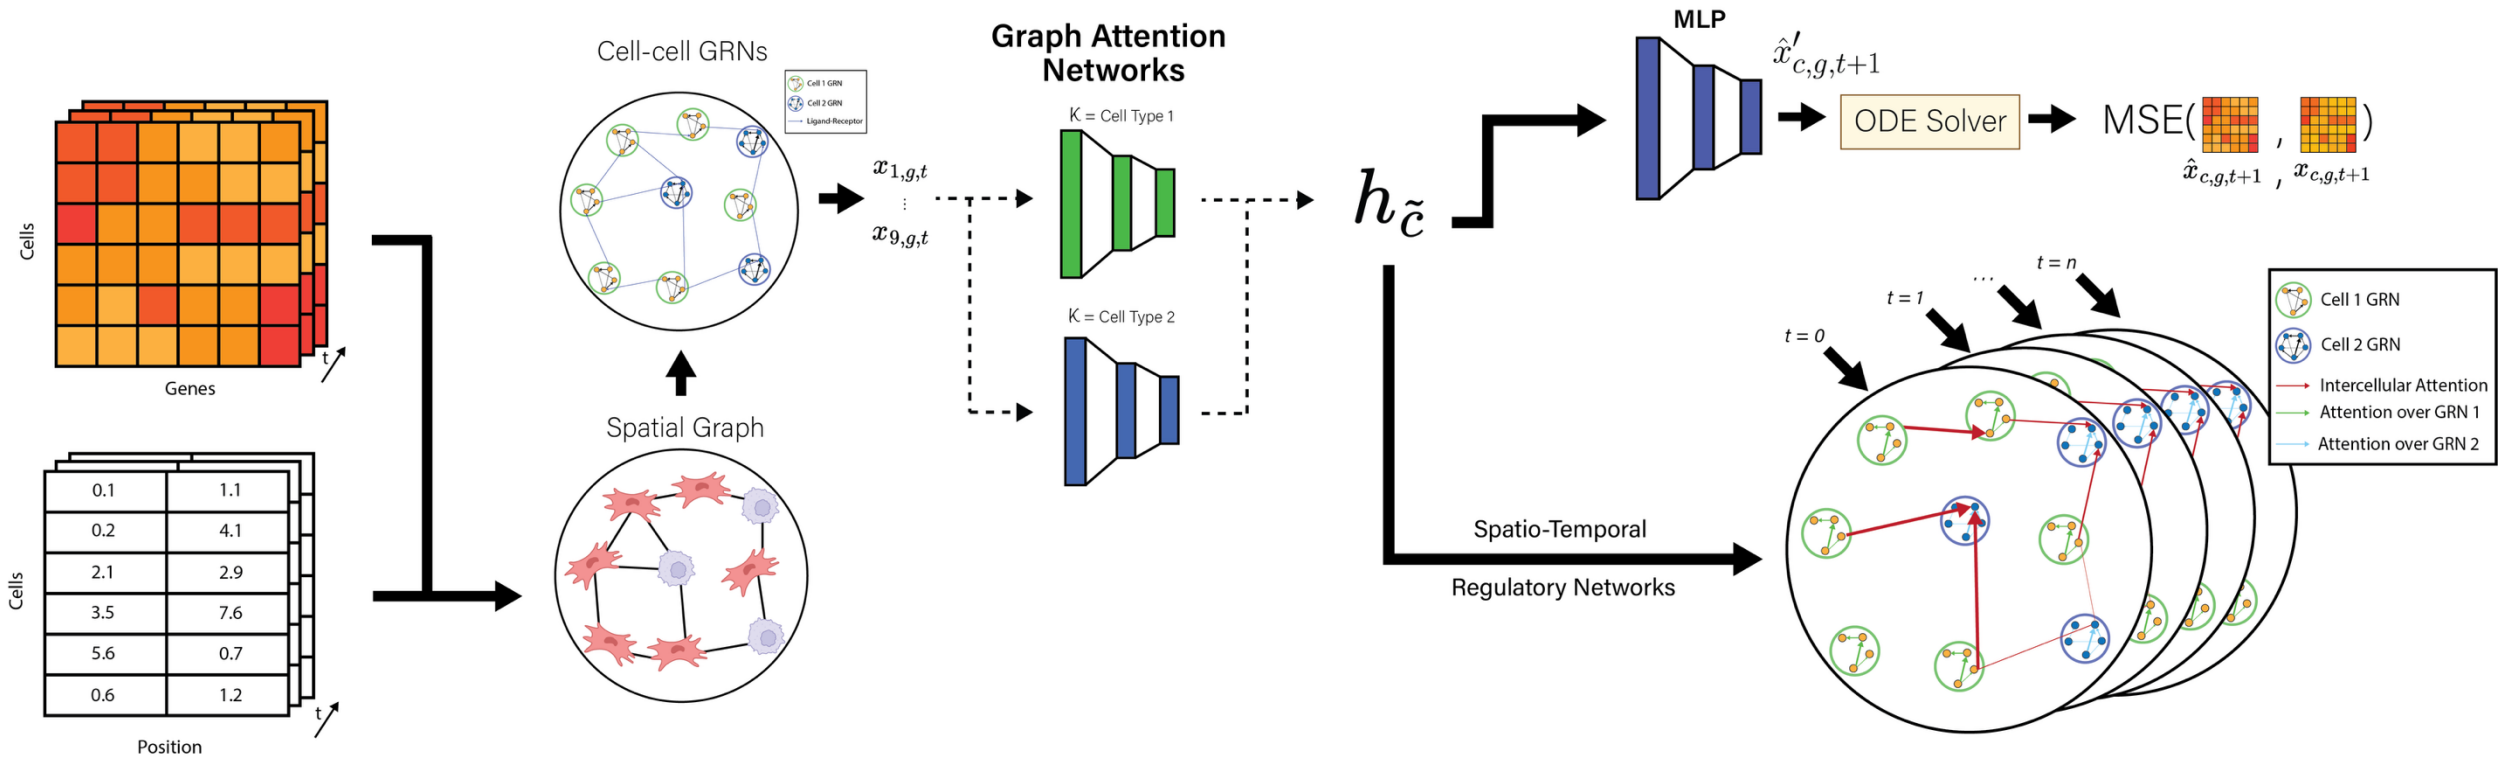

Spatial neighborhoods → Attention mechanism → Gene dynamics

Each agent in **STAGED** is a cell-type–specific model composed of a graph attention network (GAT) and a neural ODE.

Agent-based simulation of glioblastoma

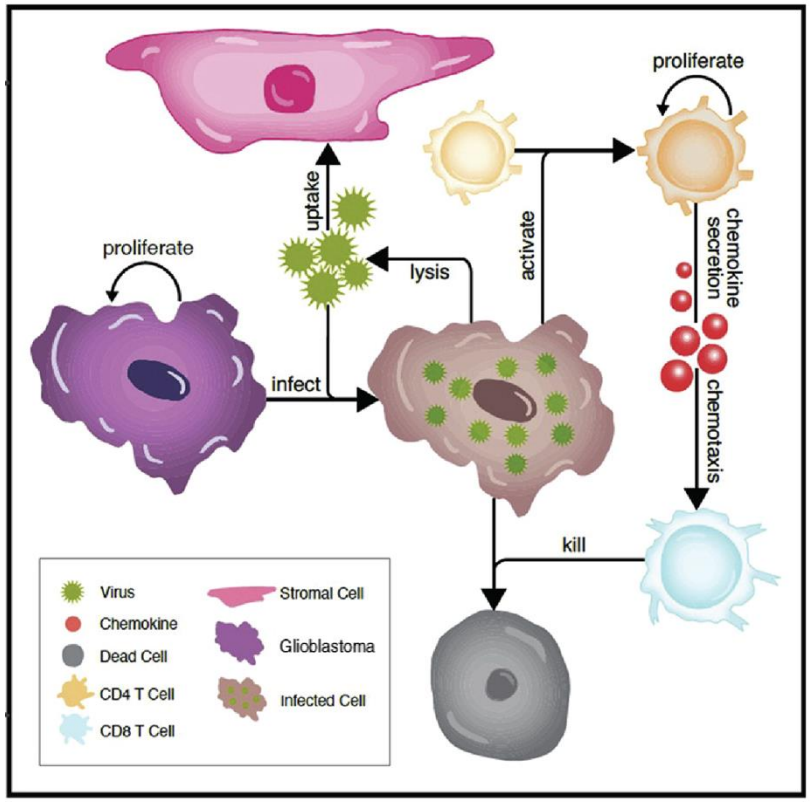

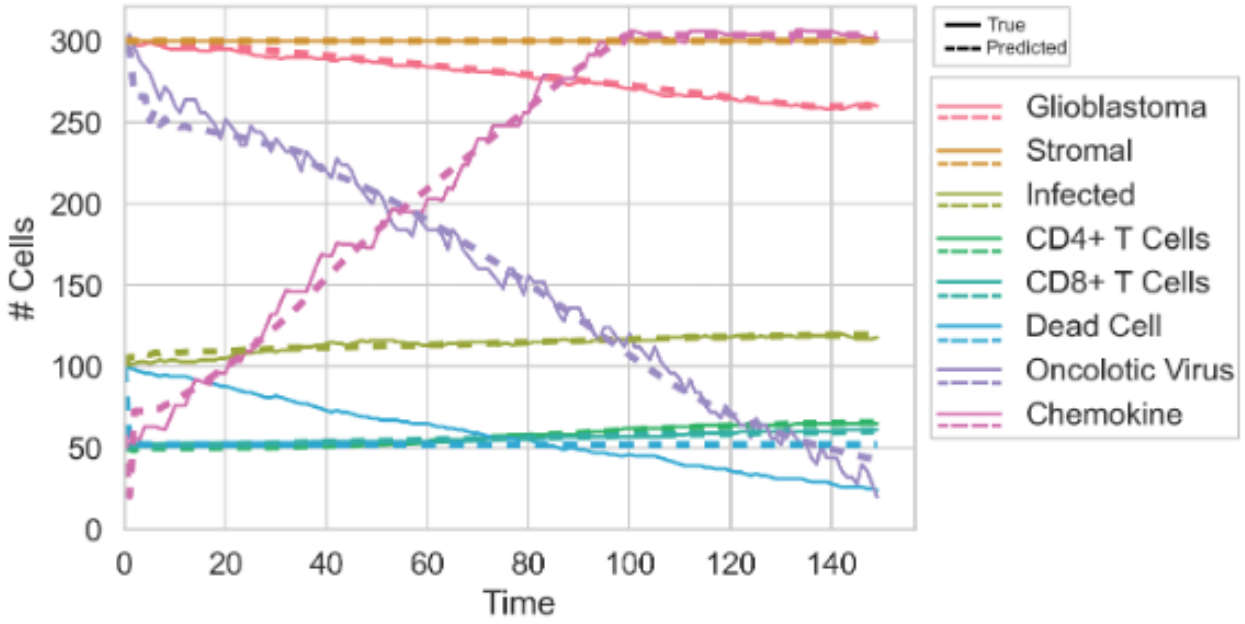

A simplified version of proliferation rates, death, and total count features.

STAGED: A Multi-Agent Neural Network for Learning Cellular Interaction Dynamics

João F. Rocha[*§], Ke Xu[*§], Xingzhi Sun[*§], Ananya Krishna[*], Dhananjay Bhaskar[*], Blanche Mongeon[†‡], Morgan Craig[†‡], Mark Gerstein[*], Smita Krishnaswamy[*]

Future work

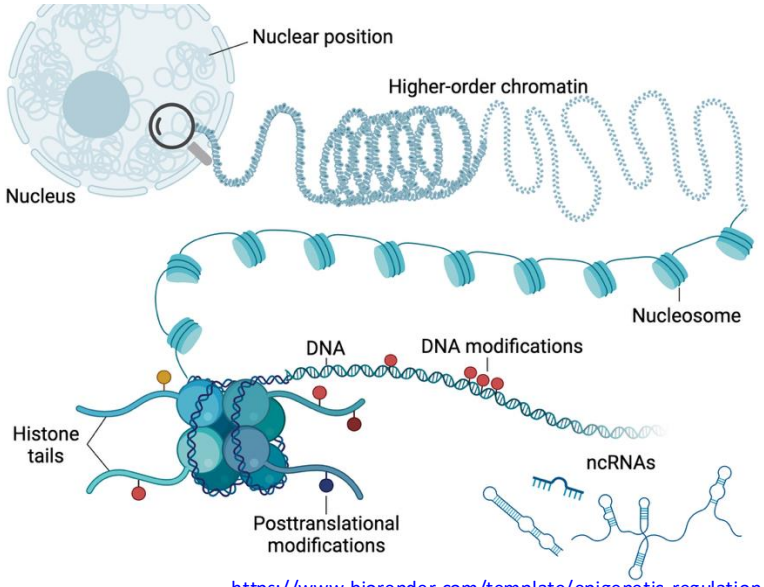

https://www.biorender.com/template/epigenetic-regulation

- Modeling epigenetic regulation

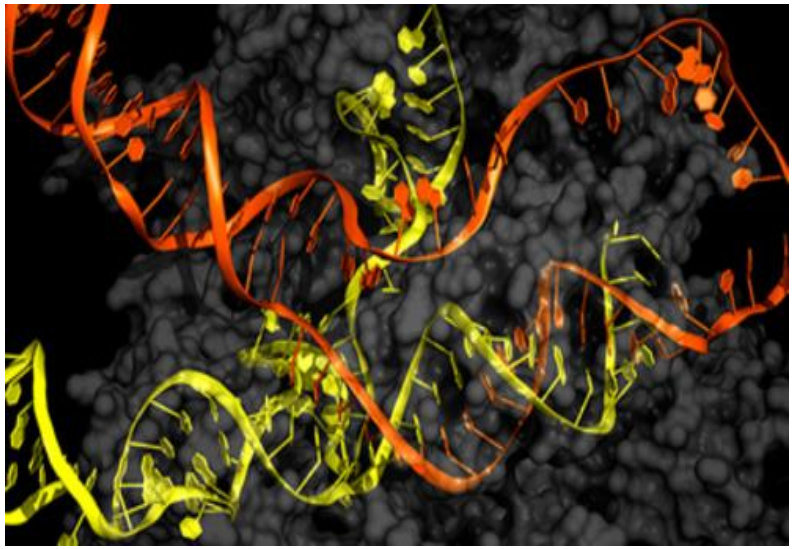

https://www.broadinstitute.org/genetic-perturbation-platform

- Simulating perturbations such as gene knockouts, receptor inhibition

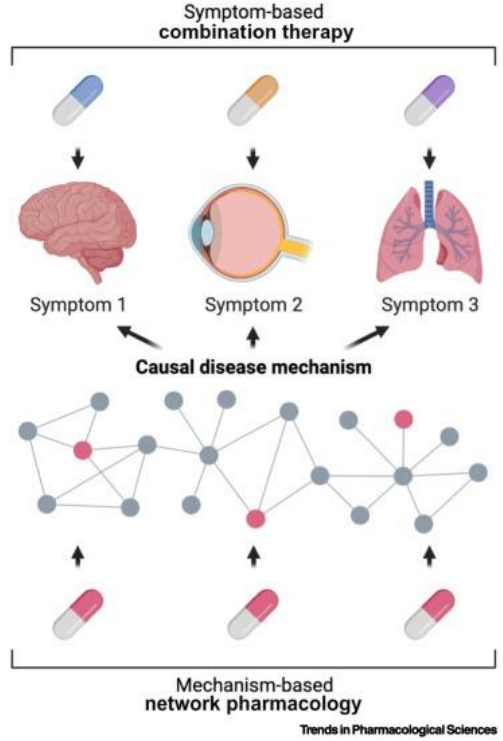

- Microenvironmental remodeling to investigate causal mechanisms in development and disease

Krishnaswamy Lab Resources

Website

https://www.krishnaswamylab.org

GitHub

github.com/KrishnaswamyLab

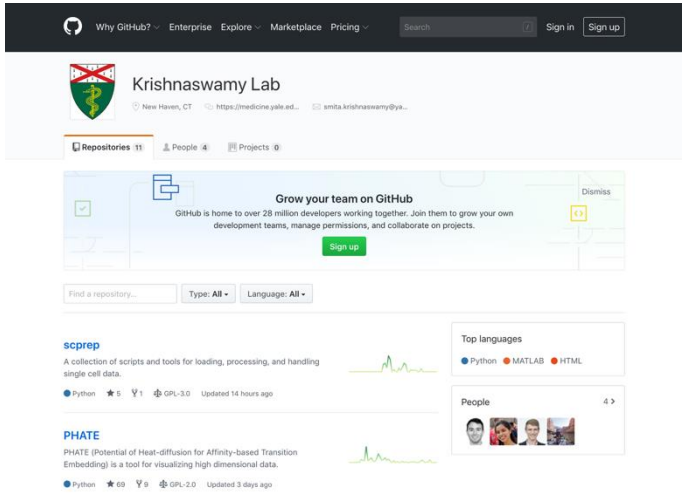

Twitter

@krishnaswamylab