Causality and Autoencoders in the Light of Drug Repurposing for COVID-19

Caroline Uhler (MIT)

IST Lisbon

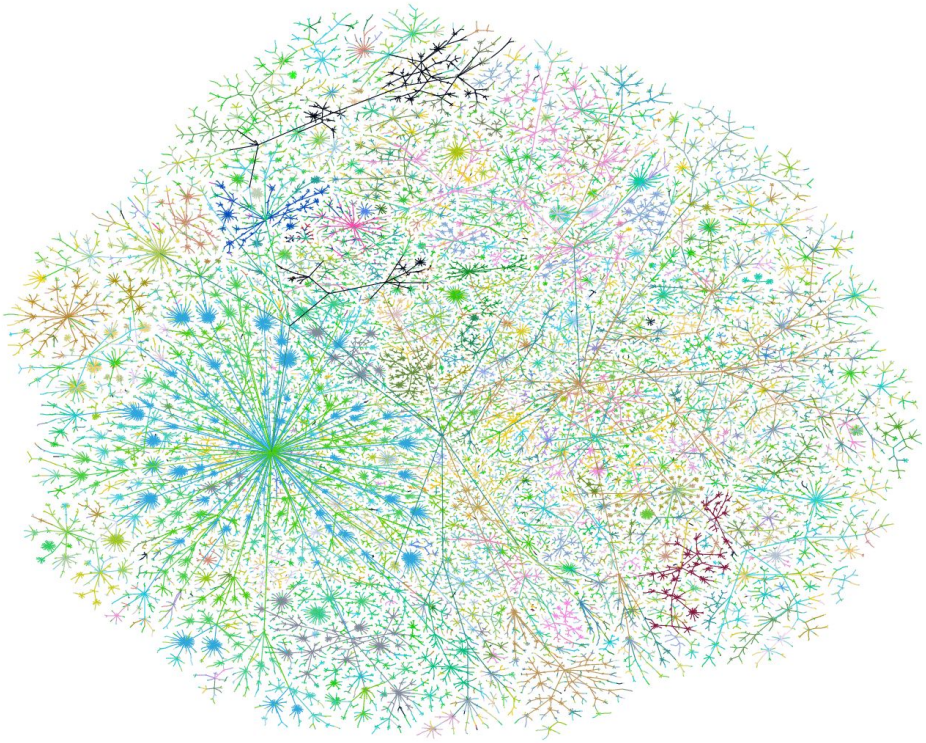

Drug development against COVID-19

- Given the urgency, most viable approach is drug repurposing
- Many drugs are inhibitors (bind to a protein so that it cannot perform its downstream role)
- Want to identify a drug that pushes the system back to normal state
- Available data: Drug signatures (screens with $\sim$1000 FDA approved drugs) and their targets, disease signatures, Protein-protein interaction networks (20,000 nodes, 200,000 edges)
- How to determine drug candidates for repurposing against a particular disease?

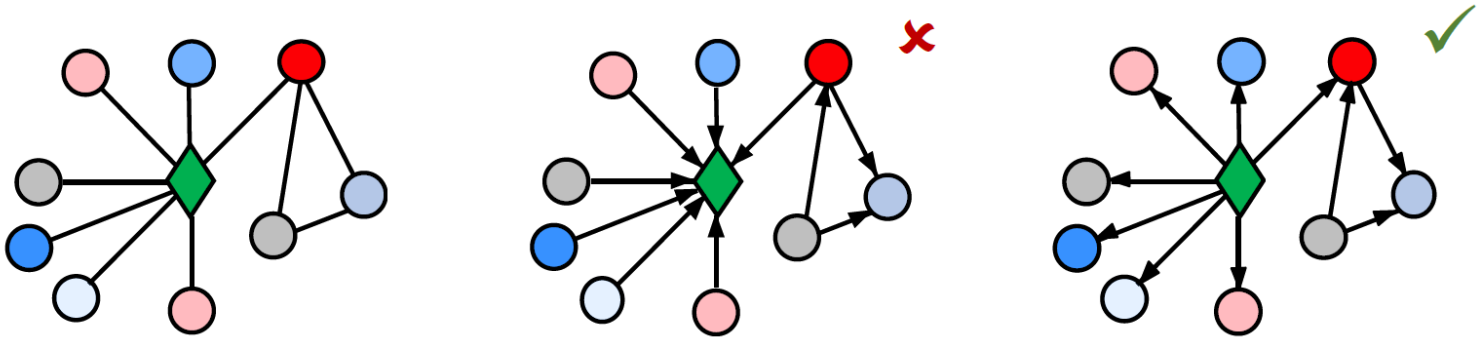

$\rightarrow$ **Learn causal graph!**

Overview - causal transport problems

Predicting the effect of an intervention:

1. Genomic interventions such as knockout experiments (few mostly known targets)

$\Rightarrow$ Theoretical and algorithmic framework for learning causal networks from observational and interventional data

2. Transport drug intervention (many unknown targets) to new cell type

$\Rightarrow$ Use inductive bias of autoencoders for synthetic interventions

Framework: Structural equation models

- Introduced by Sewell Wright in the 1920s
- Major contributions by Judea Pearl, Jamie Robins, Don Rubin, Peter Spirtes since 1970s
- Represent causal relationships by a directed acyclic graph (DAG)
- Each node is associated with a random variable; stochasticity is introduced by independent noise variables ϵ_i

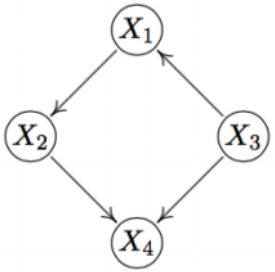

$$X_1 \leftarrow f_1(X_3, \epsilon_1)$$
$$X_2 \leftarrow f_2(X_1, \epsilon_2)$$
$$X_3 \leftarrow f_3(\epsilon_3)$$
$$X_4 \leftarrow f_4(X_2, X_3, \epsilon_4)$$

Framework: Structural equation models

- Introduced by Sewell Wright in the 1920s
- Major contributions by Judea Pearl, Jamie Robins, Don Rubin, Peter Spirtes since 1970s
- Represent causal relationships by a directed acyclic graph (DAG)
- Each node is associated with a random variable; stochasticity is introduced by independent noise variables ϵ_i

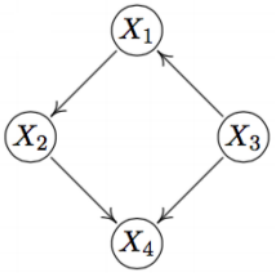

$$X_1 \leftarrow f_1(X_3, \epsilon_1)$$
$$X_2 \leftarrow f_2(X_1, \epsilon_2)$$
$$X_3 \leftarrow f_3(\epsilon_3)$$
$$X_4 \leftarrow f_4(X_2, X_3, \epsilon_4)$$

- Structural equation model also defines interventional distribution:
  - Intervention on X_2: $\mathrm{do}(X_2 = c)$
  - $p(X_3 \mid \mathrm{do}(X_4 = c)) = p(x_3) \neq p(x_3 \mid x_4)$, but $p(X_4 \mid \mathrm{do}(X_3 = c)) = p(x_4 \mid x_3) \neq p(x_4)$

From causal graphs to independence relations

- Missing edge (i, j) encodes conditional independence (CI) relation:

$$X_i \perp\!\!\!\perp X_j \mid X_{\text{ancestors}(i,j)\setminus\{i,j\}}$$

- Markov equivalence: different causal graphs can encode same CI relations and are generally indistinguishable from observational data

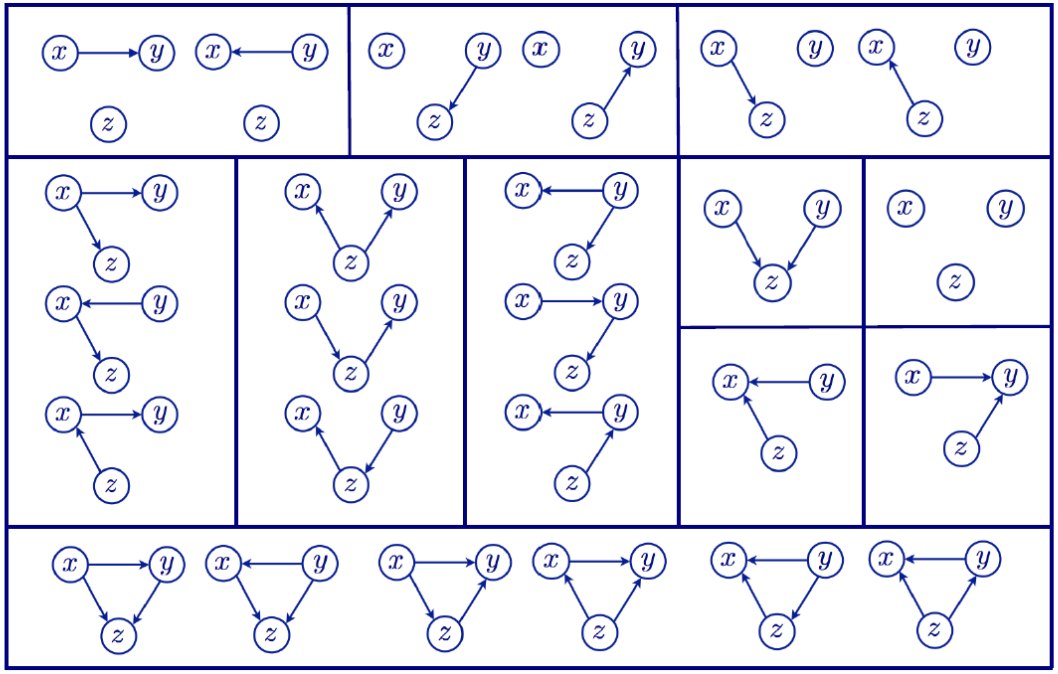

- Skeleton and immoralities (i→j←k) are identifiable

  Verma & Pearl, 1992

- Interventional Markov equivalence classes have been characterized

  Hauser & Buehlmann, JMLR 2012

  Yang, Katcoff & Uhler, ICML 2018

Permutation-based search

GES: Greedy search over Markov equivalence classes: [Chickering, 2012]

- Large search space
- No consistency guarantees in the presence of interventional data

Permutation-based search

GES: Greedy search over Markov equivalence classes: [Chickering, 2012]

- Large search space
- No consistency guarantees in the presence of interventional data

Idea: DAG defined by ordering of vertices (**permutation**) and **skeleton**

- For $p = 10$ search space is of size $10! = 3,628,800$ versus 10^{18}
- For each permutation π construct a DAG $G_\pi = (V, E_\pi)$ by

$$(i,j) \in E_\pi \iff X_i \not\perp\!\!\!\perp X_j \mid X_{\text{ancestors}_\pi(i,j)\setminus\{i,j\}}$$

Theorem (Uhler & Raskutti, Stat 2018)

Under weak conditions any sparsest DAG G_π is Markov equivalent to the true DAG (as sample size $n \to \infty$).

Greedy sparsest permutation (GSP) algorithm

edges in polytope of permutations (i.e., permutohedron) connect neighboring transpositions, e.g. $(3,1,4,2) - (3,4,1,2)$

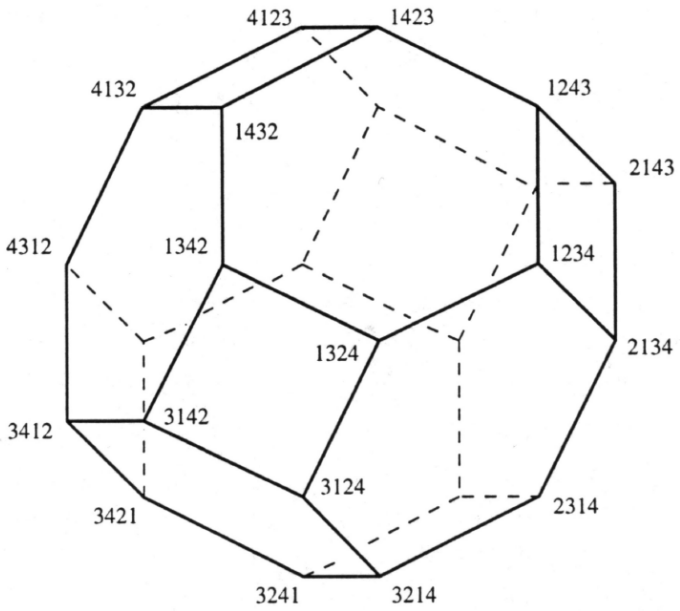

Theorem (Solus, Wang & U., 2018)

Greedy sparsest permutation (GSP) algorithm is consistent (as sample size $n \to \infty$), i.e., every local minimum is a global minimum.

Greedy sparsest permutation (GSP) algorithm

- Our Python package https://github.com/uhlerlab/causaldag has code for all methods, pre-processed perturb-seq data, etc.

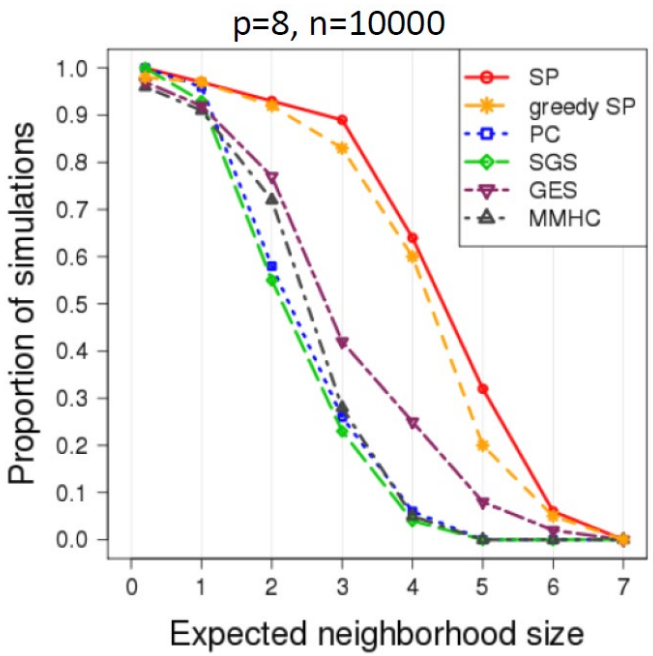

[Solus, Wang & Uhler, 2018]

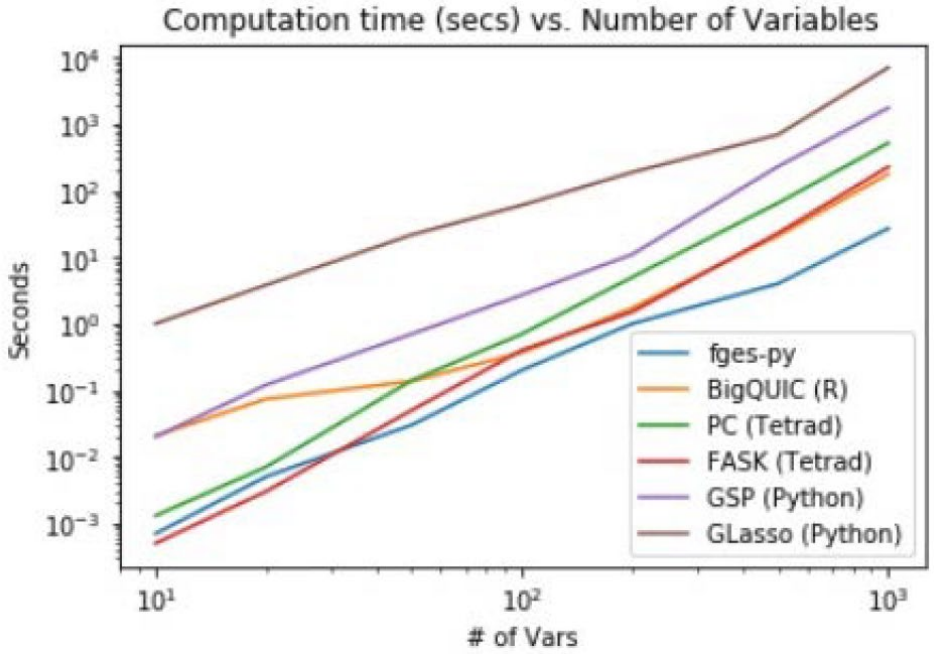

[Frederick Eberhardt: https:// www.slideshare.net/ SAMSI_Info/causal-inference-opening-workshop/ causal-discovery-in-neuroimaging-data-frederick-eberhardt-december-11-2019]

Learning from interventions and with latent variables

- **GIES:** perfect intervention adaptation of greedy search on space of Markov equivalence classes [Hauser & Bühlmann, 2012]
  - In general not consistent [Wang-Solus-Yang-Uhler, NIPS 2017]
- **IGSP:** interventional adaptation of GSP: provably consistent algorithm that can deal with interventional data
  - for hard interventions [Wang-Solus-Yang-Uhler, NIPS 2017]
  - for soft interventions [Yang-Katcoff-Uhler, ICML 2018]
  - for unknown intervention targets [Squires-Wang-Uhler, UAI 2020]
- **GSPo:** greedy search over posets to deal with latent confounders
  - sparsest poset is consistent [Bernstein-Saeed-Squires-Uhler, AISTATS 2020]
  - no consistency proof of greedy search yet

Overview

Predicting the effect of an intervention:

1. Genomic interventions such as knockout experiments (few mostly known targets)

$\Rightarrow$ Theoretical and algorithmic framework for learning causal networks from observational and interventional data

2. Transport drug intervention (many unknown targets) to new cell type

$\Rightarrow$ Use inductive bias of autoencoders for synthetic interventions

Style transfer and transporting causal effects

Latent space arithmetics for style transfer using autoencoders / GANs:

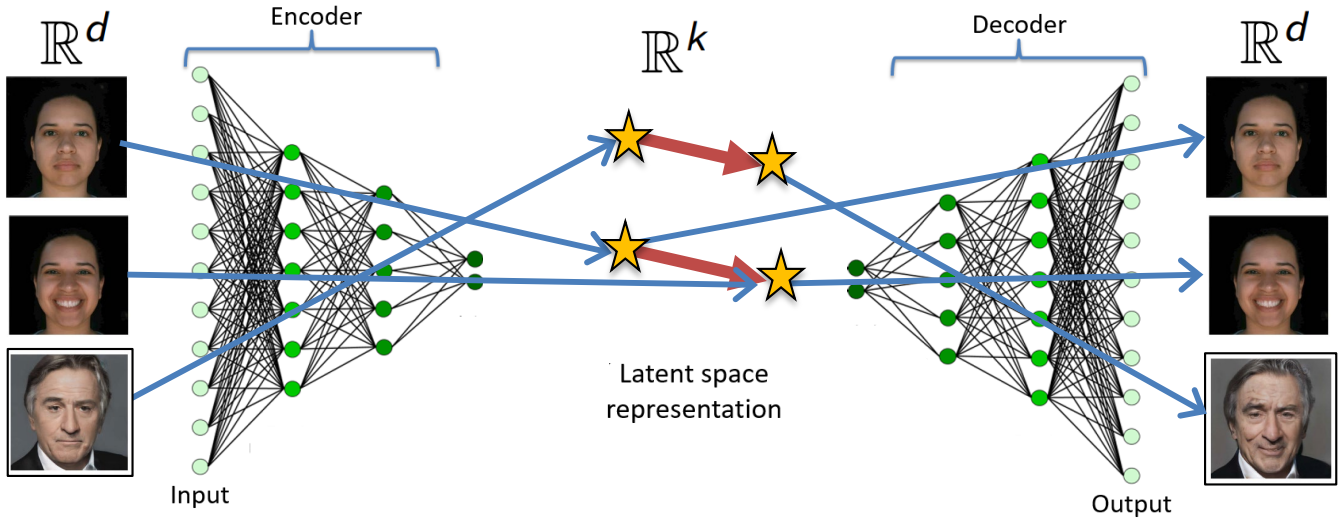

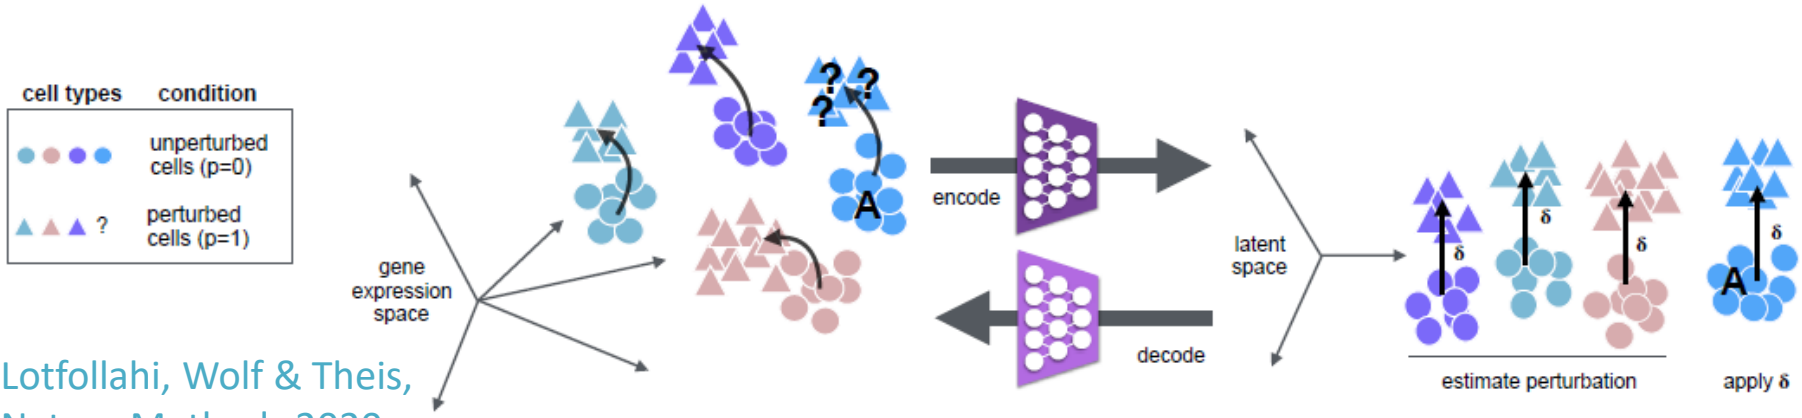

Lotfollahi, Wolf & Theis, Nature Methods 2020

Is this a general phenomenon? How does this fit in with work by Bareinboim, Pearl and co-authors on necessary and sufficient conditions for causal transportability?

Predicting the effect of a drug on a different cell type

- CMap: 1.2mio samples (1000-dim expression vectors), 1000s of perturbations (knockouts, overexpression, small molecules including ~800 FDA-approved drugs)

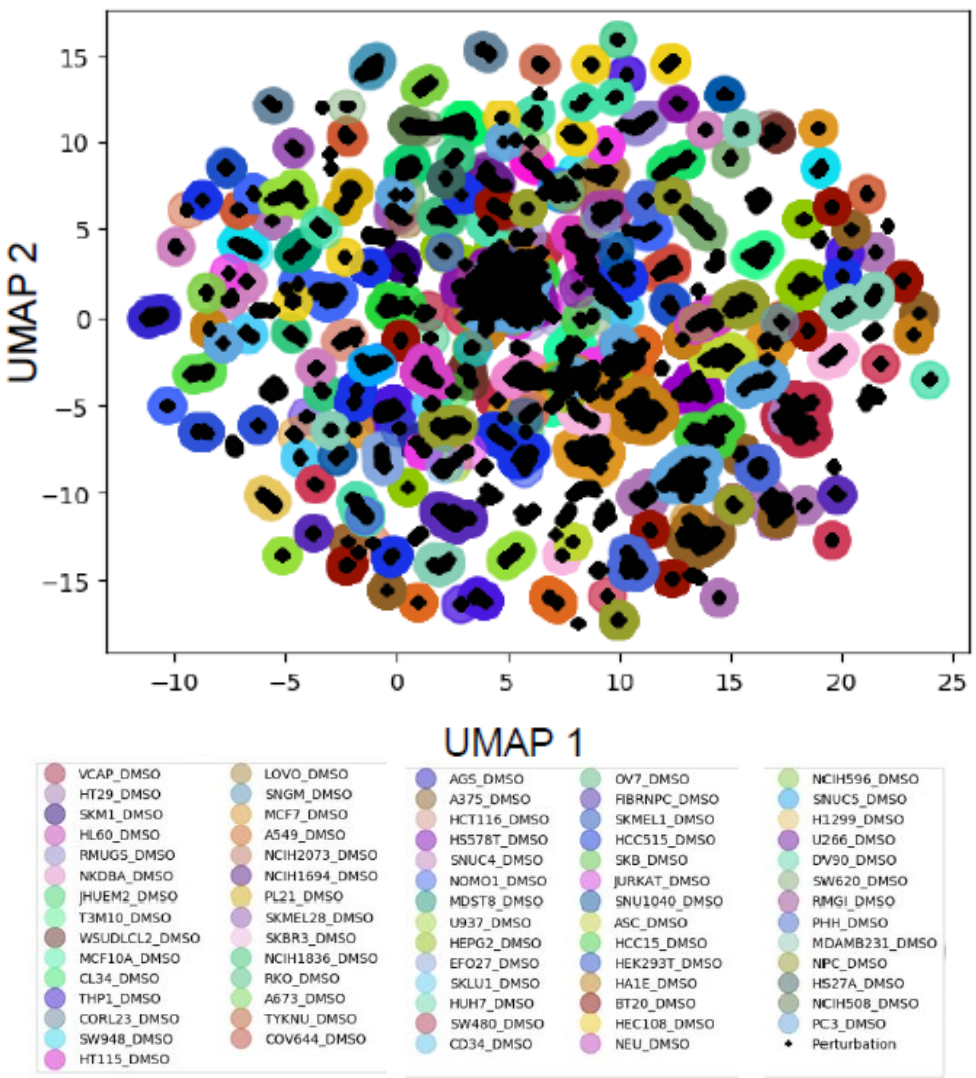

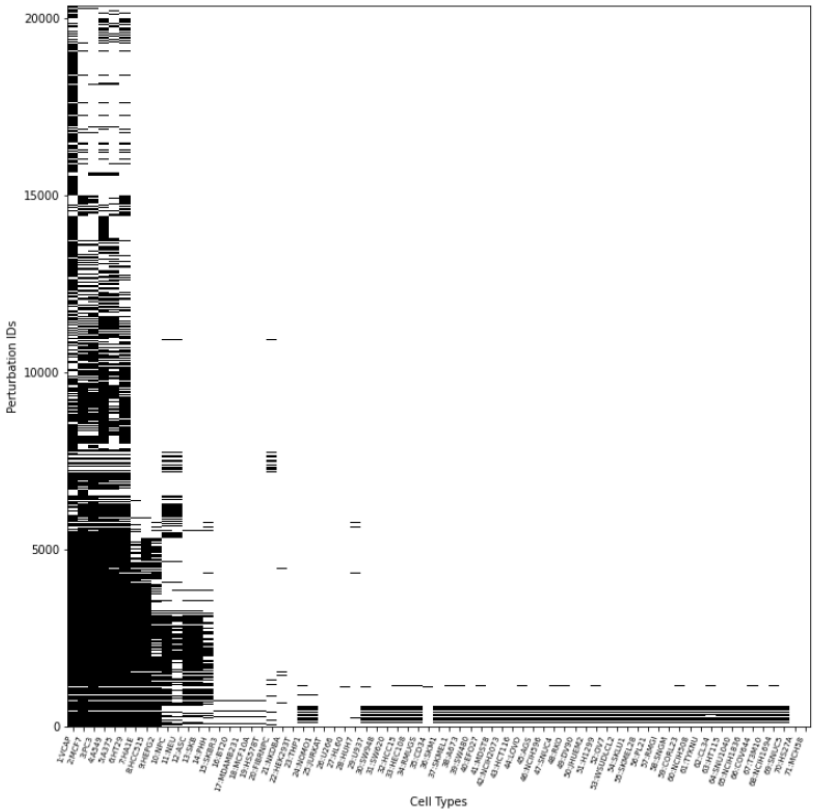

Overparameterized autoencoders align drug signatures

Under-parameterized autoencoder

Over-parameterized autoencoder

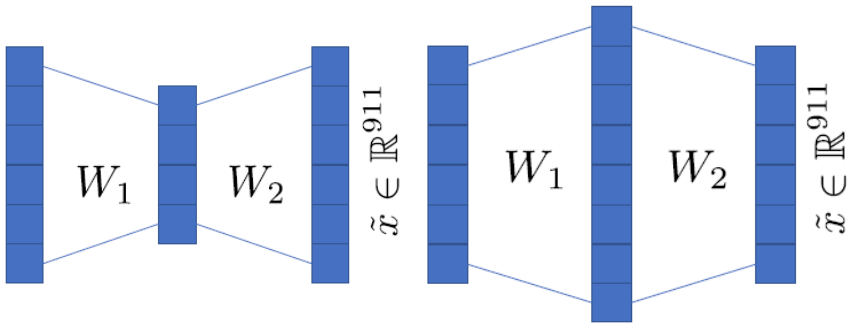

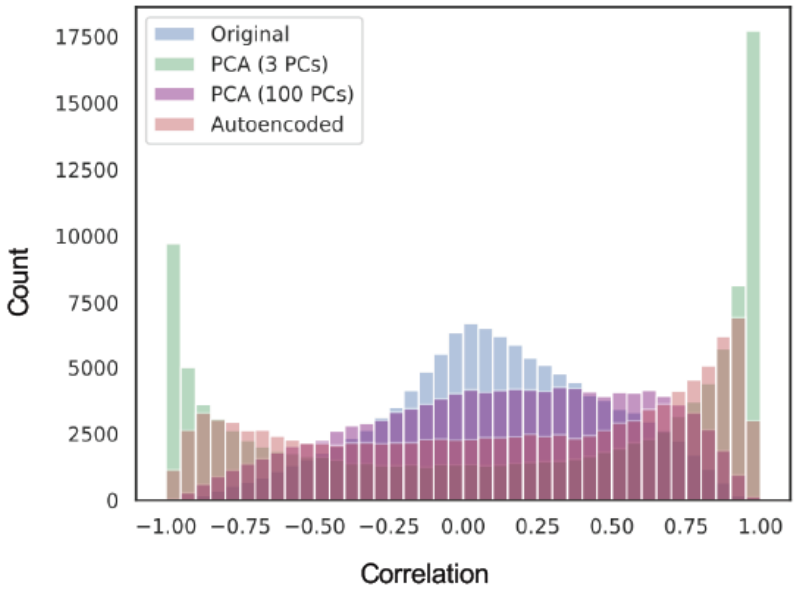

[Belyaeva, Cammarata, Radhakrishnan, Squires, Yang, Shivashankar & Uhler, arXiv: 2006.03735]

Correlation between drug signatures of A549 and MCF7 cells

Reconstruction accuracy

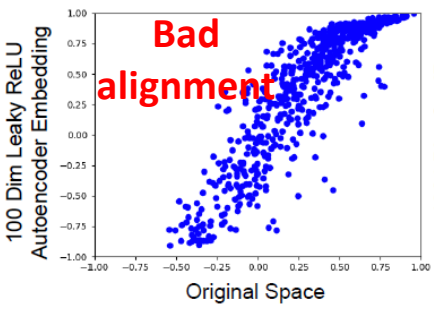

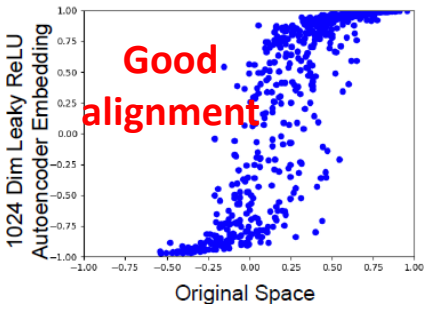

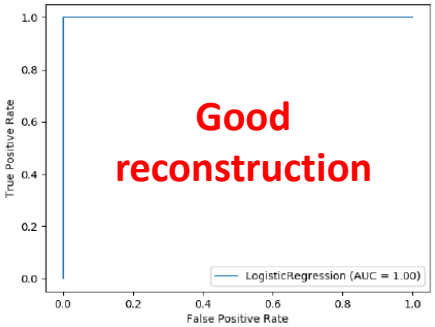

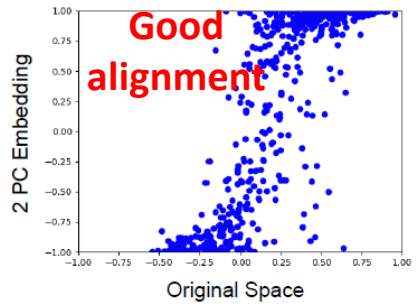

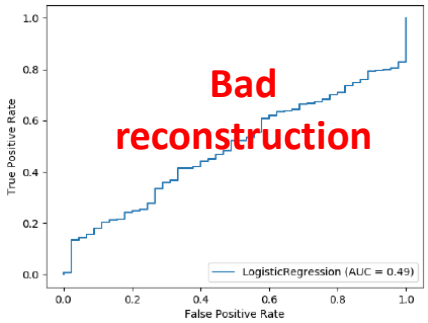

Inductive bias of over-parameterized autoencoders

- Given training examples $x^{(1)}, \ldots x^{(n)} \in \mathbb{R}^k$, $n < k$, autoencoders are typically trained using gradient descent initialized ≈ 0 to solve

$$\arg\min_{\psi \in \Psi} \sum_{i=1}^{n} \|\psi(x^{(i)}) - x^{(i)}\|_2^2$$

- Over-parameterized linear setting: solutions range from the identity map to the projection onto $\mathrm{span}(x^{(1)}, \ldots, x^{(n)})$

Inductive bias of over-parameterized autoencoders

- Given training examples $x^{(1)}, \ldots x^{(n)} \in \mathbb{R}^k$, $n < k$, autoencoders are typically trained using gradient descent initialized ≈ 0 to solve

$$\arg\min_{\psi \in \Psi} \sum_{i=1}^{n} \|\psi(x^{(i)}) - x^{(i)}\|_2^2$$

- Over-parameterized linear setting: solutions range from the identity map to the projection onto $\mathrm{span}(x^{(1)}, \ldots, x^{(n)})$

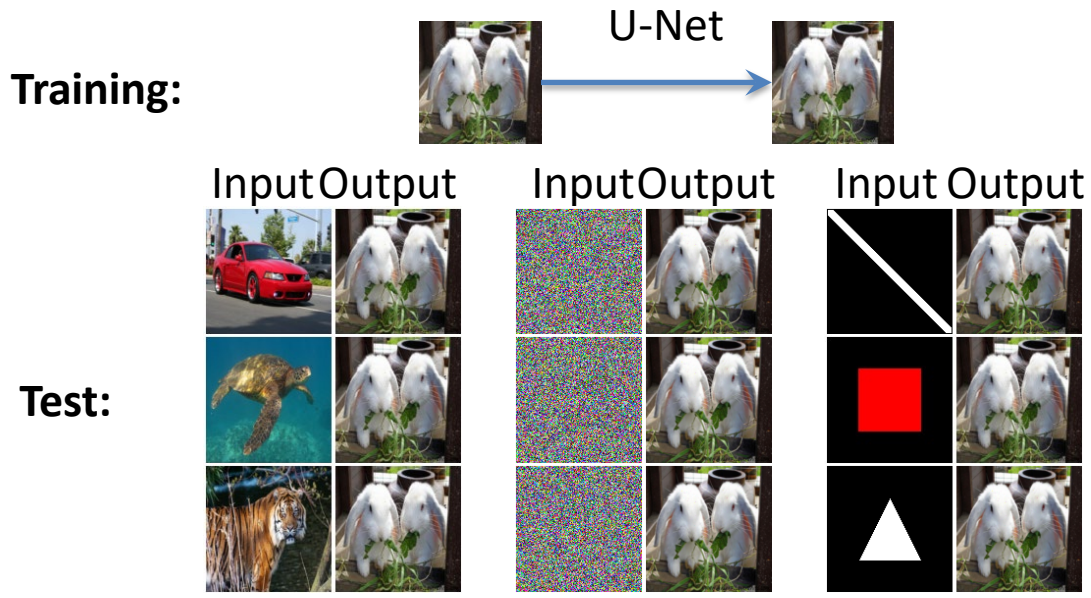

In the extreme case of n=1 the training example is memorized

Inductive bias of over-parameterized autoencoders

- Over-parameterized autoencoders have many ways to interpolate training data

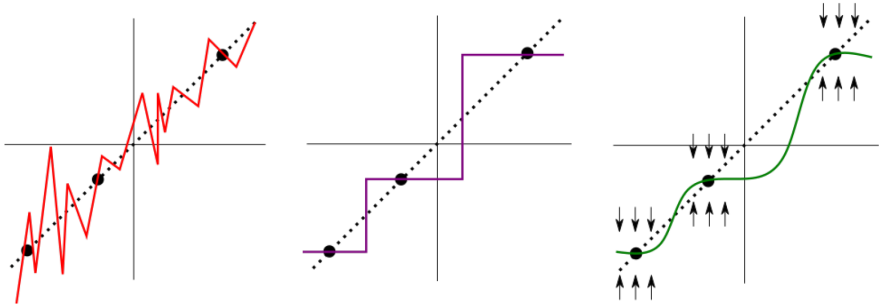

- View autoencoder as discrete dynamical system: $x_{t+1} = f(x_t)$

Noise Template				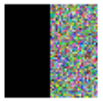		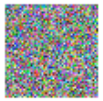
Recovery Rate	500/500 (No Noise)	499/500	421/500	266/500	233/500	1/500 (At Random)

- We proved that standard over-parameterized autoencoders (without additional regularizers) are self-regularizing: they learn maps, where training examples are attractive fixed points!

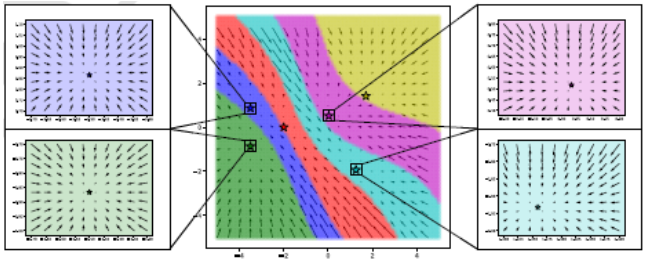

[Radhakrishnan, Belkin & Uhler, PNAS 2020]

Correlating disease and drug signatures

Correlate disease and drug signatures to obtain list of drug candidates

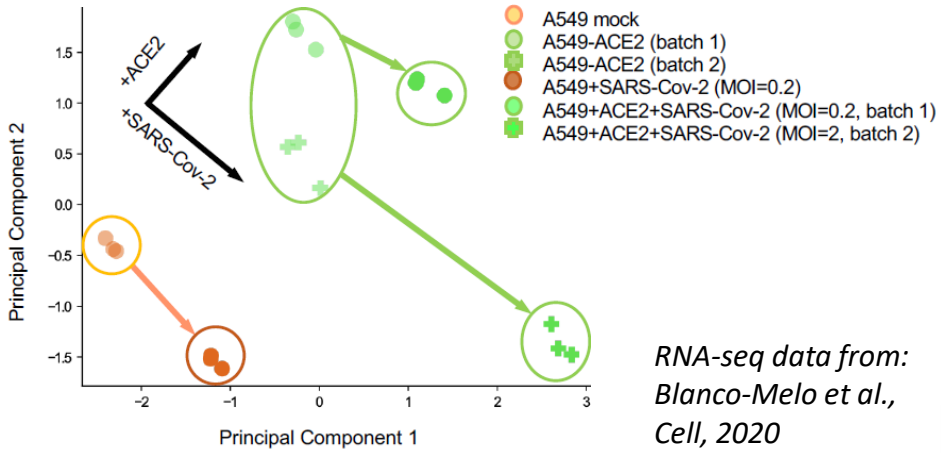

RNA-seq data from: Blanco-Melo et al., Cell, 2020

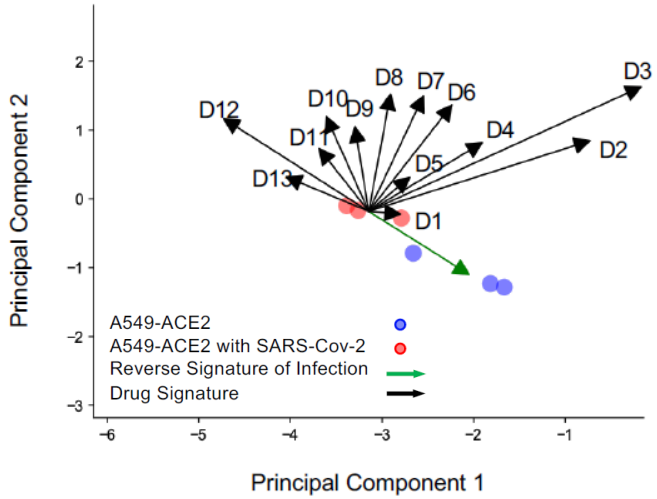

D1: chloroquine
D2: dantrolene
D3: propofol
D4: leflunomide
D5: estradiol
D6: miglitol
D7: escitalopram
D8: doxercalciferol
D9: calcipotriol
D10: salmeterol
D11: dicloxacillin
D12: nabumetone
D13: ketoconazole

gene	protein target	drug	correlation	affinity
ACVR2A	Activin receptor type-2A	dasatinib	0.88	6.68
AURKC	Aurora kinase C	erlotinib	0.87	6.22
		sorafenib	0.87	6.68
		sunitinib	0.87	6.66
		pazopanib	0.87	6.12
		nuvolitinib	0.87	5.06
		axitinib	0.88	8.89
BRSK1	Serine/threonine-protein kinase BRSK1	sunitinib	0.87	5.46
CDK17	Cyclin-dependent kinase 17	sorafenib	0.87	5.8
		sunitinib	0.87	5.92
EGFR	Epidermal growth factor receptor	dasatinib	0.88	7.1
		docetaxel	0.87	9.0862
		erlotinib	0.87	9.22
		imatinib	0.87	5.12
		sunitinib	0.87	6.07
		axitinib	0.88	5.64
		afatinib	0.86	10
		bosutinib	0.86	7.74
FGFR1	Fibroblast growth factor receptor 1	dasatinib	0.88	5.43
		imatinib	0.87	5
		sorafenib	0.87	5.6
		sunitinib	0.87	6.28
		pazopanib	0.87	6
		axitinib	0.88	6.42
FGFR3	Fibroblast growth factor receptor 3	dasatinib	0.88	5.41
		sorafenib	0.87	5.38
		sunitinib	0.87	6.54
		pazopanib	0.87	6.21
		axitinib	0.88	6.68
HDAC1	Histone deacetylase	vorinostat	0.87	8
		vorinostat	0.87	6.96
		vorinostat	0.87	8.89
		vorinostat	0.87	8
		belinostat	0.87	9.07
HSP90AA1	Heat shock protein HSP 90-alpha	tanespimycin	0.87	8.5229
		primaquine	0.87	8.2218
IRAK1	Interleukin-1 receptor-associated kinase 1	imatinib	0.87	5.92
		sunitinib	0.87	7.85
		pazopanib	0.87	5.23
		nuvolitinib	0.87	6.54
		axitinib	0.88	5.51
		afatinib	0.86	6.62
		bosutinib	0.86	6.22
PAK1	Serine/threonine-protein kinase PAK 1	bosutinib	0.86	5.64
PDE4B	Phosphodiesterase 4	milrinone	0.86	5.22
		vardenafil	0.86	5.35
RIPK1	Receptor-interacting serine/threonine-protein kinase 1	sunitinib	0.87	6.43
		pazopanib	0.87	6.59
		axitinib	0.88	5.6
RIPK2	Receptor-interacting serine/threonine-protein kinase 2	dasatinib	0.88	7.51
		erlotinib	0.87	6.39
		sorafenib	0.87	5.89
		pazopanib	0.87	6.24
		axitinib	0.88	5
		afatinib	0.86	5.57
		bosutinib	0.86	5.43
STK3	Serine/threonine-protein kinase 3	sunitinib	0.87	7.25
		axitinib	0.88	5.66
		bosutinib	0.86	6.43
		tofacitinib	0.88	5.37

Serine/threonine protein kinase

Receptor tyrosine kinase

Validating drug targets using a causal analysis

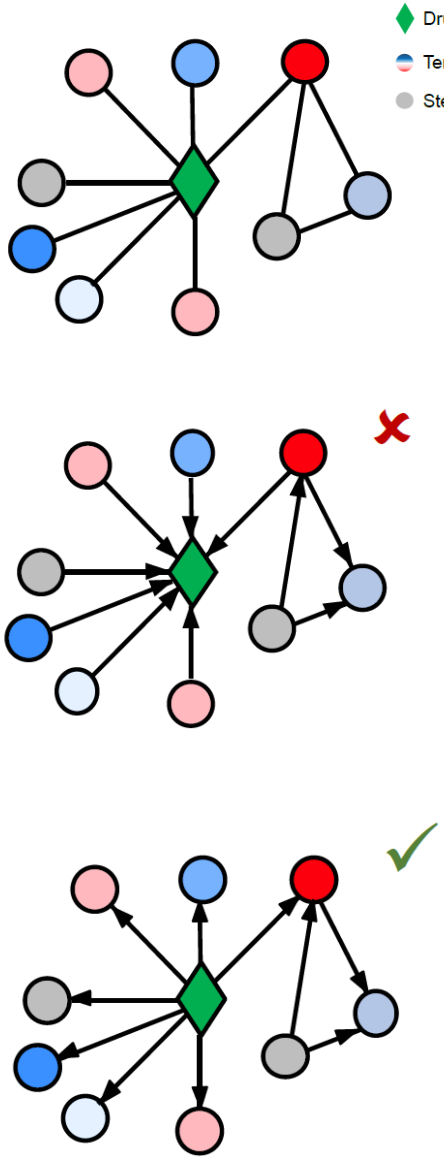

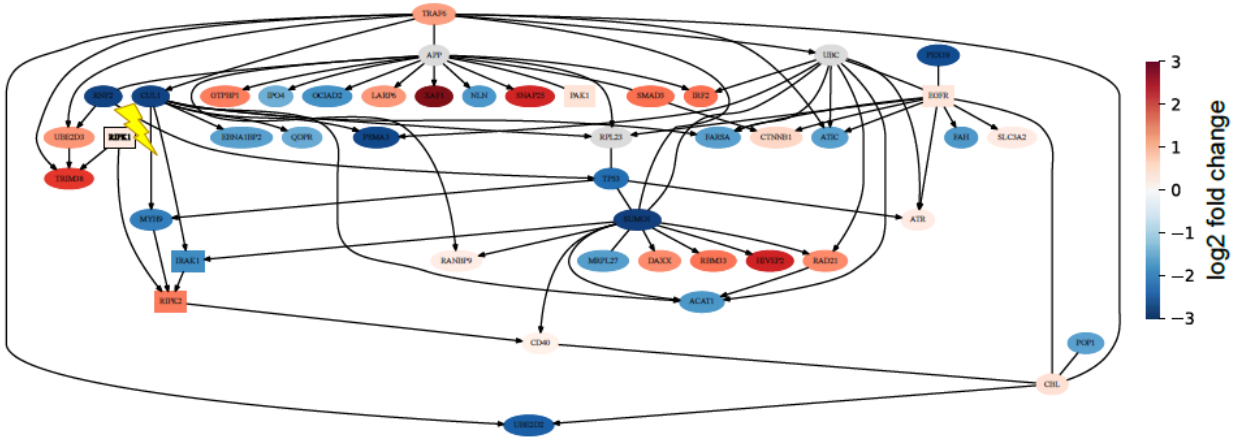

- RIPK1 has most downstream differentially expressed genes based on inferred (from single-cell RNA-seq data) causal graph in A549 cells and also in AT2 cells
- While role of other targets is similar, RIPK1 becomes peripheral in causal graph without taking ageing into account
- RIPK1 binds to SARS-CoV-2 proteins (*Gordon et al., Nature, 2020*)

Role of RIPK1 linking SARS-CoV-2 replication & ageing?

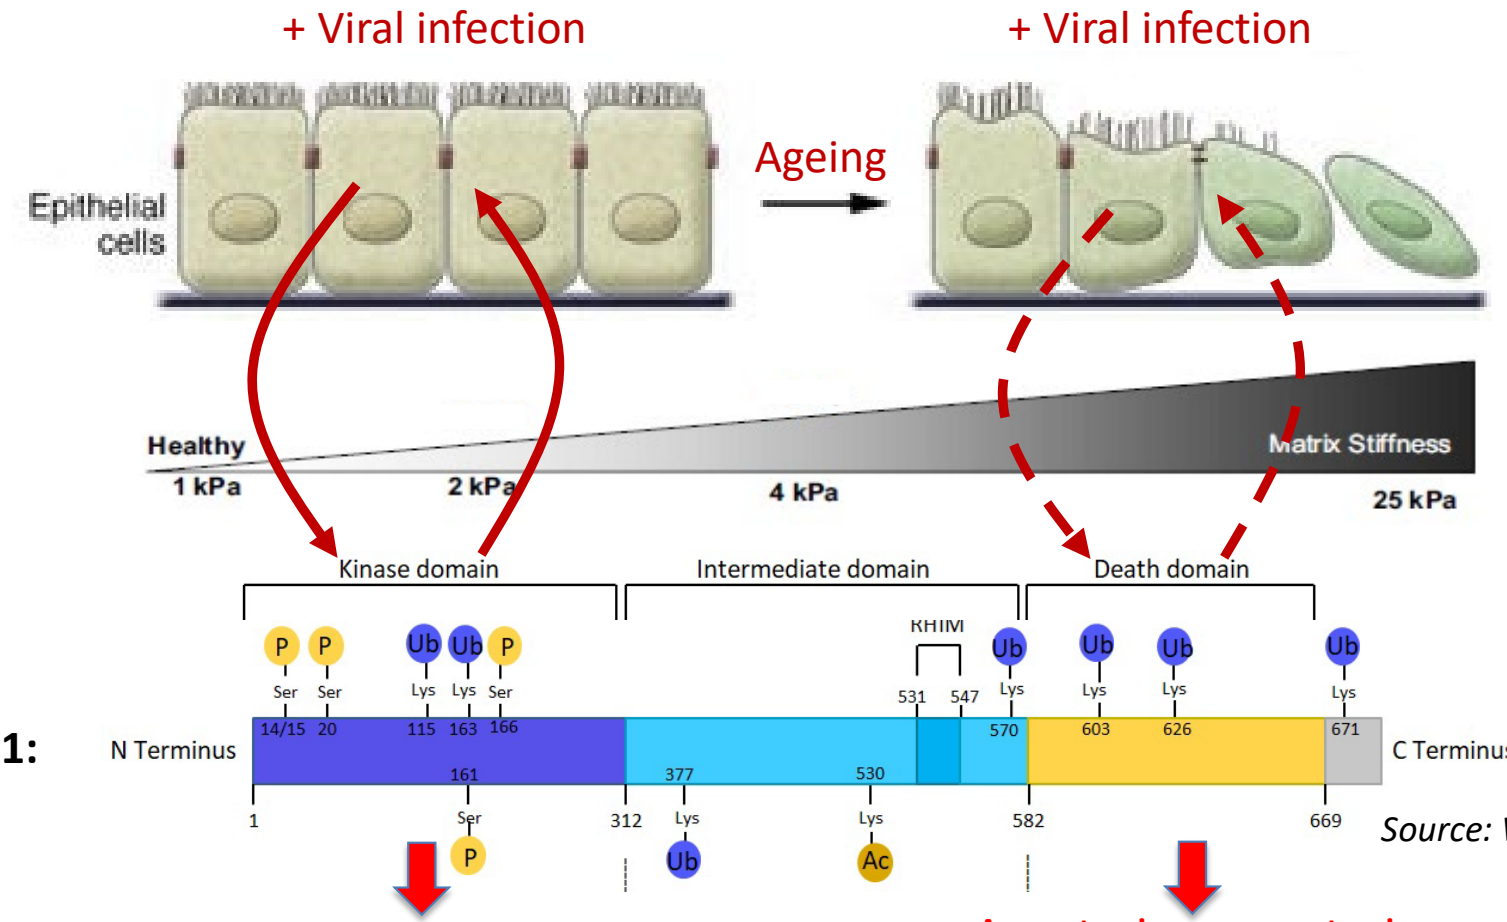

RIPK1:

Source: Wikipedia

[Uhler & Shivashankar, Nature Reviews (2020); Belyaeva, Cammarata, Radhakrishnan, Squires, Yang, Shivashankar & Uhler, arXiv: 2006.03735]

Conclusions

- Transporting between interventions: Developed a theoretical and algorithmic framework for integrating observational and interventional data for causal inference

- Transporting intervenion effects between populations: Over-parameterized autoencoders show implicit bias that may be of great interest for causal transportability

- A principled causal framework is critical for drug discovery

Belyaeva, Cammarata, Radhakrishnan, Squires, Yang, Shivashankar & Uhler, Nature Communications (in press)

Acknowledgments

PhD students:

- Raj Agrawal
- Anastasiya Belyaeva
- Louis Cammarata
- Saachi Jain
- Adityanarayanan Radhakrishnan
- Chandler Squires
- Karren Yang
- Jiaqi Zhang

MSc/undergraduate students

- Josh Amaniampong
- Sathwik Karnik
- Eshaan Nichani
- Neha Prasad
- George Stefanakis
- Annie Yun

Postdocs:

- Daniel Bernstein
- Jan-Christian Huetter
- Neriman Tokcan

Collaborators:

- Mikhail Belkin
- G.V. Shivashankar

Funding: