


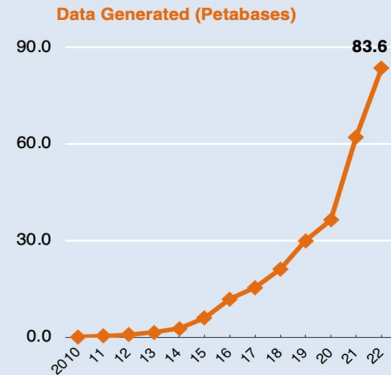
Causal Representation Learning in the context of perturbation screens and cell state transitions

Caroline Uhler (MIT & Broad Institute)



Converging 2 revolutions of 21st century


Broad Genomics, by the numbers



Similar to Twitter


Netflix movie corpus: 60pb

Total data currently under management at Broad: ~100pb



Biological technologies

- Medical imaging
- Single cell genomics
- Cryo-EM



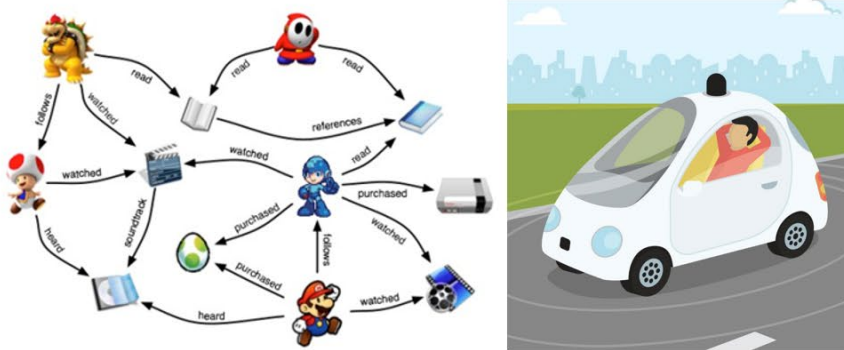
- Machine learning (ML) & AI
- Cloud computing

Data technologies

New era of research at the interface of the life and data sciences

Current ML methods are not sufficient


Key drivers of ML today:



Emphasis on prediction accuracy

In Biology:

Ultimate goal is to discover natural laws and underlying mechanisms



We need to bridge the gap from predictive to causal modeling




- **Biology has genetic & chemical tools that allow perturbational screens on a scale unmatched by any other field**
- **Novel AI theory is needed to make use of this data and get at causality**
- **Biology has the potential to inspire new and profound avenues of AI research**



ERIC AND WENDY
SCHMIDT CENTER
AT BROAD INSTITUTE

Cancer Immunotherapy Data Science Challenge

Immunotherapy: Boost own immune system to kill cancer cells



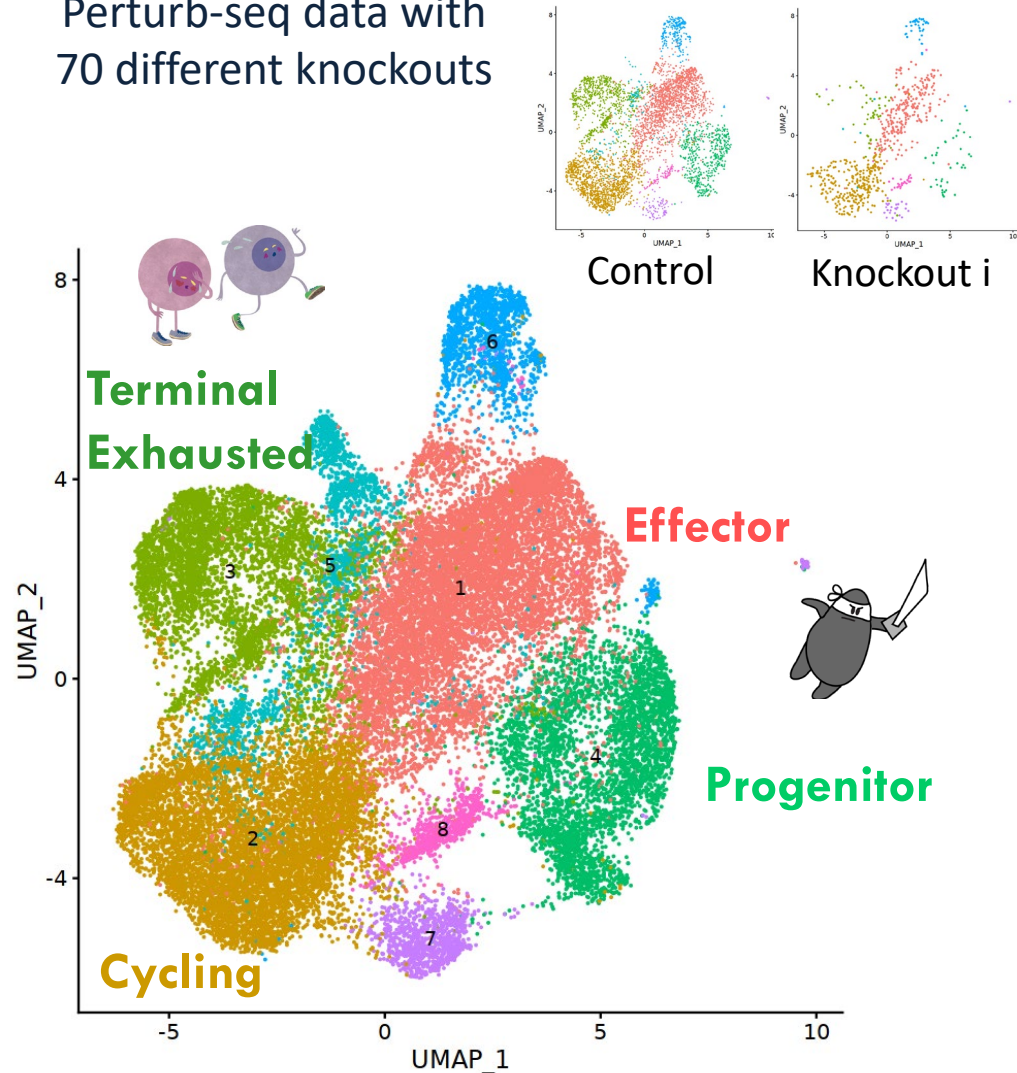
Competition: Make immunotherapy more effective for more cancers

Identify knockouts to move T-cells away from the terminal exhausted state

Jan 9-Feb 3, 2023


Validation in mouse, prizes of \$50k

Perturb-seq data with 70 different knockouts




<https://go.topcoder.com/schmidtcentercancerchallenge>

Causal imputation problems in single-cell biology



Transport to new contexts



Transport to new perturbations

How to think of causal variables in images?
Can multi-modality help?

Over-parameterized neural networks

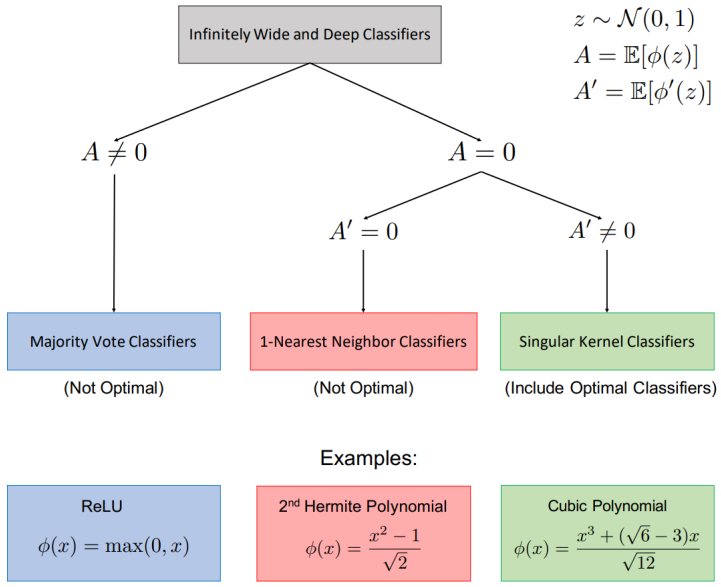
- Deep neural networks can generalize while interpolating the training data

Belkin et al., PNAS 2019

- Infinitely wide neural networks converge to the neural tangent kernel

Jacot et al., NeurIPS 2018

- Neural tangent kernel with specific activation function is Bayes optimal for classification



Examples:

ReLU

$$\phi(x) = \max(0, x)$$

2nd Hermite Polynomial

$$\phi(x) = \frac{x^2 - 1}{\sqrt{2}}$$

Cubic Polynomial


$$\phi(x) = \frac{x^3 + (\sqrt{6} - 3)x}{\sqrt{12}}$$

$$z \sim \mathcal{N}(0, 1)$$

$$A = \mathbb{E}[\phi(z)]$$

$$A' = \mathbb{E}[\phi'(z)]$$

Autoencoder:



Sanchez-Lengeling et al., Science 361, 360–365 (2018)

Over-parameterized neural networks


- Deep neural networks can generalize while interpolating the training data

Belkin et al., PNAS 2019

- Infinitely wide neural networks converge to the neural tangent kernel

Jacot et al., NeurIPS 2018


- Neural tangent kernel with specific activation function is Bayes optimal for classification




$z \sim \mathcal{N}(0, 1)$
 $A = \mathbb{E}[\phi(z)]$
 $A' = \mathbb{E}[\phi'(z)]$

Radhakrishnan et al., arXiv:2204.14126

Autoencoder:




Sanchez-Lengeling et al., Science 361, 360-365 (2018)




Over-parameterized autoencoders learn maps that are **contractive at training examples**

Radhakrishnan, Belkin & Uhler, PNAS 2020

Causal imputation problems in single-cell biology



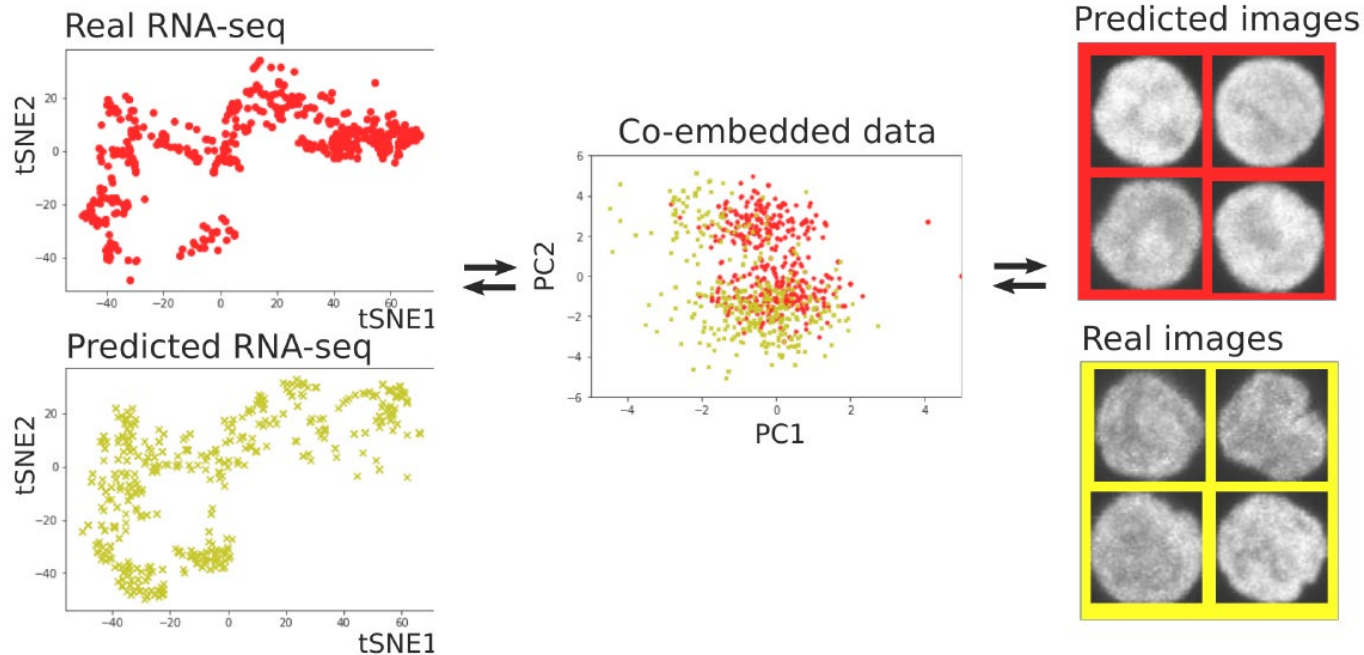
Transport to new contexts



Transport to new perturbations

How to think of causal variables in images?
Can multi-modality help?

Idea 1: Multi-modal learning to discover causal feature



Yang et al., Nature Communications 2021


**Representation learning as a tool for causal feature discovery
by learning integrated latent spaces:**

Causal features should be invariant to modality in which they are measured!


Invariant prediction for causal inference: Peters, Buehlmann, Meinshausen
Invariant risk minimization: Arjovsky, Bottou, Gulrajani, Lopez-Paz

Causal feature learning: Chalupka, Perona, Eberhardt
Disentanglement: Schoelkopf, Bengio,...


Over-parameterization & spatial transcriptomics




STARmap data: single-cell gene expression plus spatial coordinates from mouse Alzheimer model



Over-parameterized (4 mice brains)




Standard autoencoder

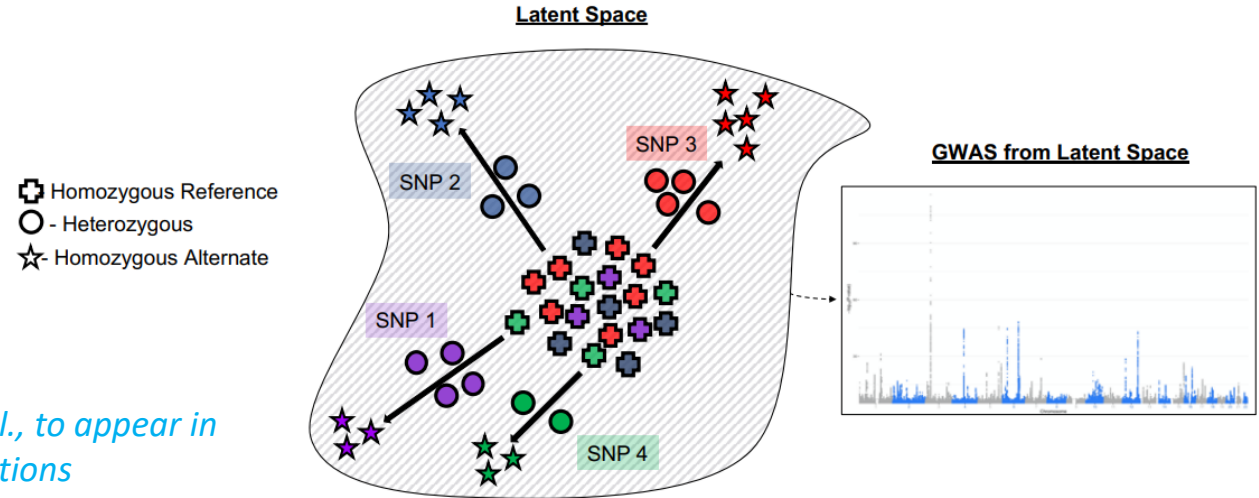


Unsupervised GWAS in multi-modal latent space

Cross-Modal Training Pipeline




GWAS and Clustering of Cross-modal Latent Space



Learning latent causal graph

$$Z = AZ + \varepsilon$$

linear additive causal structural equation model with non-Gaussian noise



$$X^e = G^e \cdot \begin{pmatrix} Z_{\mathcal{L}} \\ Z_{I_e} \end{pmatrix}$$

full column rank

Allows application of the results on identifiability of linear ICA

For each shared latent node there is at least one child in every domain


Theorem: The number of shared latent nodes and the joint domain distribution is identifiable. If there are no edges between the shared and domain-specific latent components and each shared latent node has at least 1 pure child, then also the shared latent graph is identifiable.

Sturma, Squires, Drton & Uhler, arXiv:2302.00993


Similar identifiability results under interventions on latent nodes when only having access to a single modality

Squires, Seigal, Bhate & Uhler, arXiv:2211.16467

Causal imputation problems in single-cell biology




Transport to new contexts



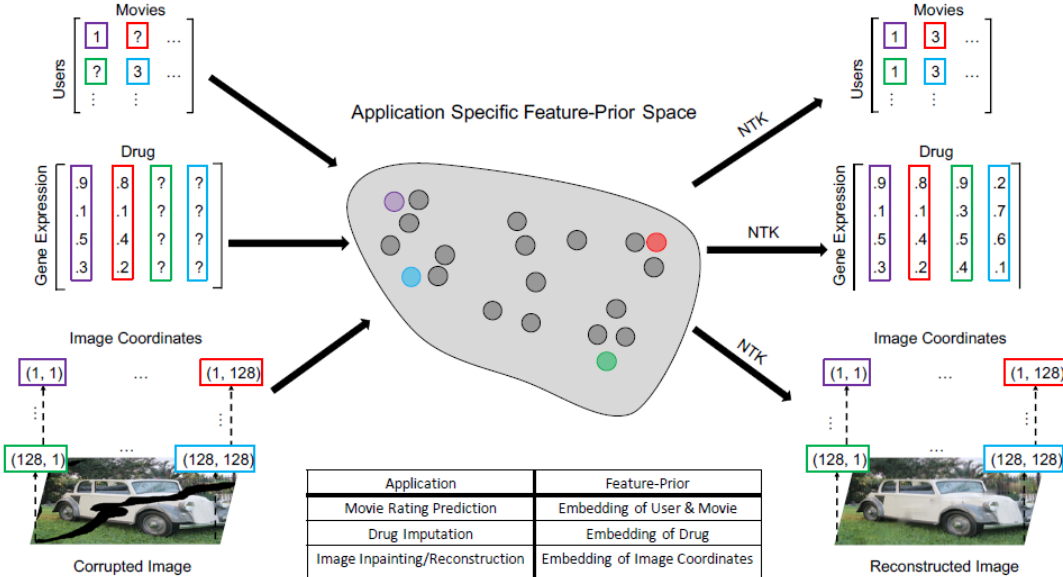
Transport to new perturbations

How to think of causal variables in images?
Can multi-modality help?


Idea 2: Matrix completion using neural tangent kernel



We built an NTK framework for matrix completion that can make use of feature priors on rows and columns



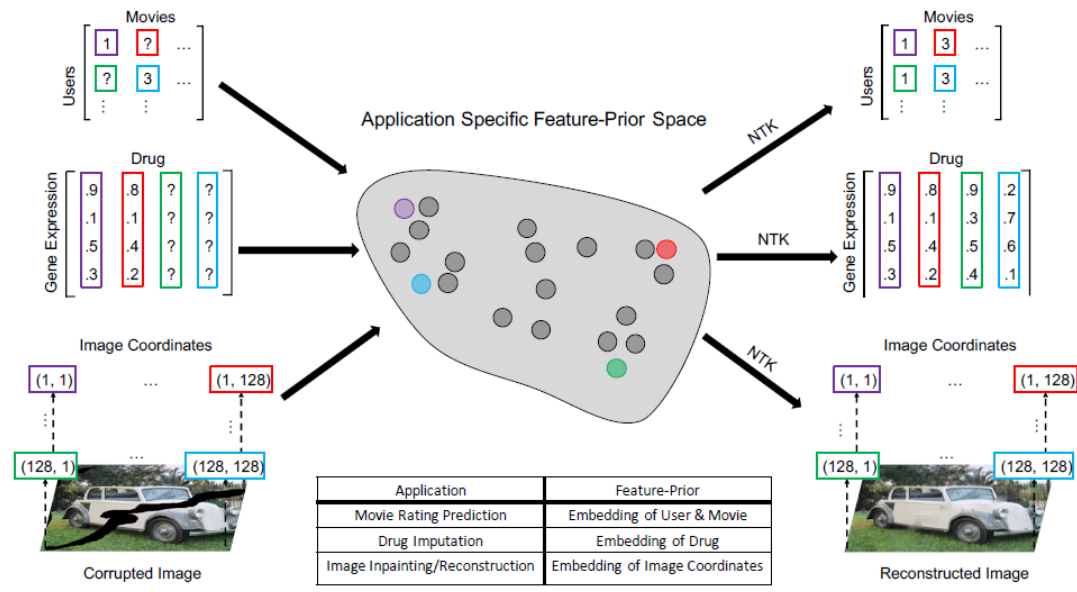
Idea 2: Matrix completion using neural tangent kernel



Predict the effect of a new perturbation / context combination?

Standard low-rank matrix completion approaches won't work!

We built an NTK framework for matrix completion that can make use of feature priors on rows and columns



CMap (Full Dataset)

Evaluation Metric*	Mean Over Cell Type (Naïve Baseline)	FaLRTC (Liu et al. 2013)	DNPP (Hodos et al. 2018)	NTK (Ours)
Pearson r	0.374 ± 0.0004	0.545 ± 0.0003	0.556 ± 0.0003	0.572 ± 0.0002
Mean R ²	0.134 ± 10 ⁻⁵	0.286 ± 0.0003	0.296 ± 0.0004	0.320 ± 0.0002
Mean Cosine Similarity	0.371 ± 10 ⁻⁵	0.536 ± 0.0004	0.541 ± 0.0004	0.554 ± 0.0002


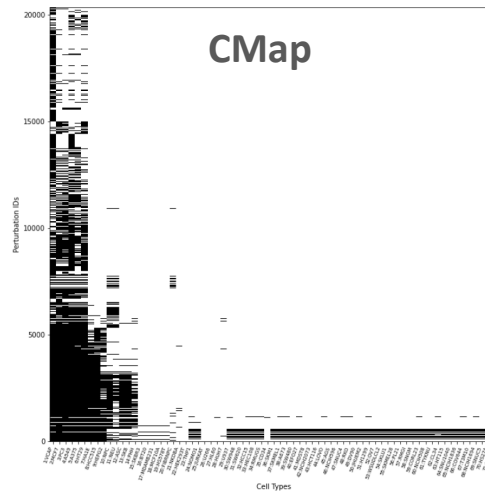
CMap (Sparse Regime)

Evaluation Metric*	Mean Over Cell Type (Naïve Baseline)	FaLRTC (Liu et al. 2013)	DNPP (Hodos et al. 2018)	NTK (Ours)
Pearson r	0.450	0.544	0.538	0.573
Mean R ²	0.197	0.285	0.278	0.324
Mean Cosine Similarity	0.448	0.536	0.532	0.565

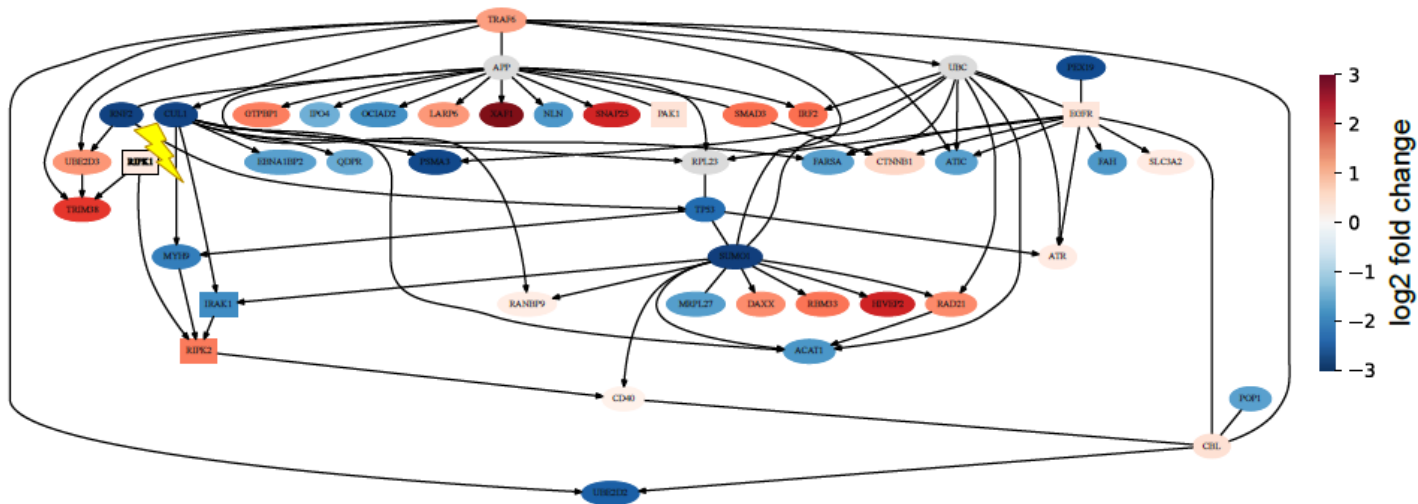
*Higher is better, with a maximum of 1.

Target identification in the context of COVID-19

Transport effect of perturbations from CMap to SARS-CoV-2 infected A549 cells and find drug that is most anticorrelated with disease signature



Causal structure discovery to identify putative causal drug targets:



Biomedical sciences are uniquely suited not only to being one of the greatest beneficiaries of research machine learning but also one of the greatest sources of inspiration for it.

- ❖ Developed a theoretical and algorithmic framework for integrating and translating between different data modalities, observational and interventional
- ❖ Autoencoders are not only extremely useful for data integration and translation, but also for studying the theoretical properties of neural networks
- ❖ Over-parameterization leads to remarkable self-regularization properties and computational gains
- ❖ If we are able to predict the effect of unseen perturbations, we can build active framework for optimal intervention design to induce desired distribution shift

Acknowledgments

PhD students:

- Louis Cammarata
- Daniel Paysan
- **Adityanarayanan Radhakrishnan**
- Hannah Schlueter
- **Chandler Squires**
- Yitong Tseo
- Chenyu Wang
- Jiaqi Zhang
- Xinyi Zhang

MSc/undergraduate students

- Alvaro Ribot Barrado
- Cathy Cai
- Emily Liu
- Nten Nyiam
- Ishika Shah


Postdocs:

- Matthew Amodio
- Salil Bhate
- Wengong Jin
- Kiran Shiragur

Collaborators:

Misha Belkin (UCSD), Mathias Drton (TUM)
GV Shivashankar (ETH Zurich)

Funding:



ERIC AND WENDY
SCHMIDT CENTER

AT BROAD INSTITUTE



SIMONS
FOUNDATION

