IMPERIAL

# Lessons from the Past for Smarter Decisions and Future Wins

Sonali Parbhoo
s.parbhoo@imperial.ac.uk
AI for Actionable Impact (AI4AI) Lab
Imperial College London

Microsoft Research Cambridge
July 11, 2024

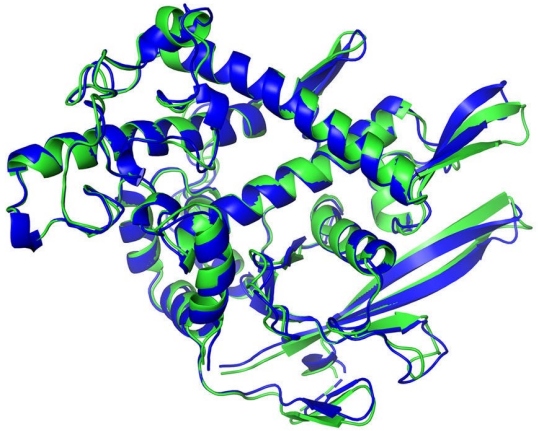
*Collaborators: Aaman Rebello, Haoyang Hong, Ioanna Papanikolaou, Leo Benac, Abhishek Sharma, Shengpu Tang, Omer Gottesman, Joseph Futoma, Shalmali Joshi, Taylor Killian, Stefan Bauer, Jenna Wiens, Marzyeh Ghassemi, Leo Celi, Finale Doshi-Velez.*

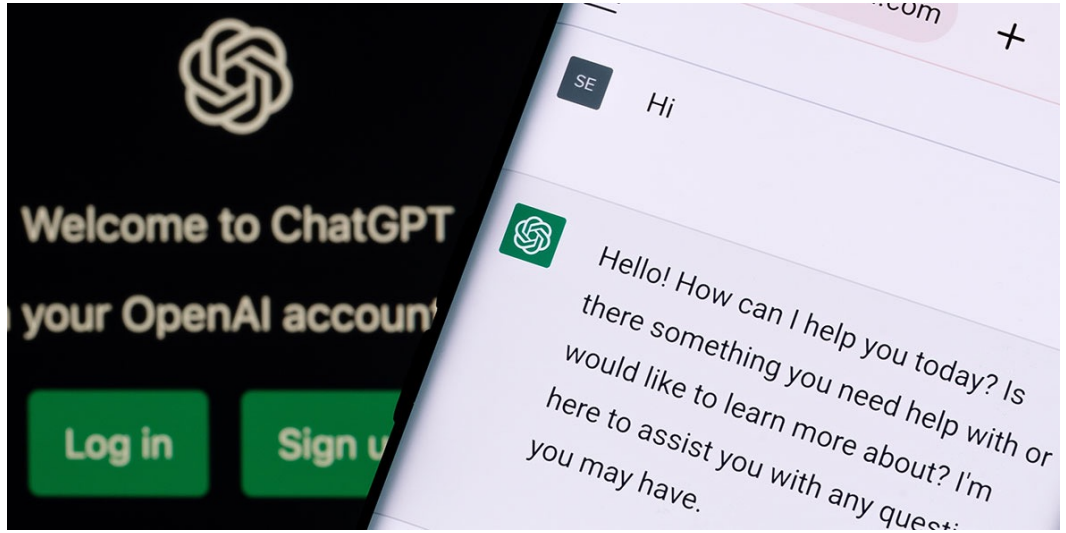
# Success in AI














# Making sequential decisions in healthcare is not so easy

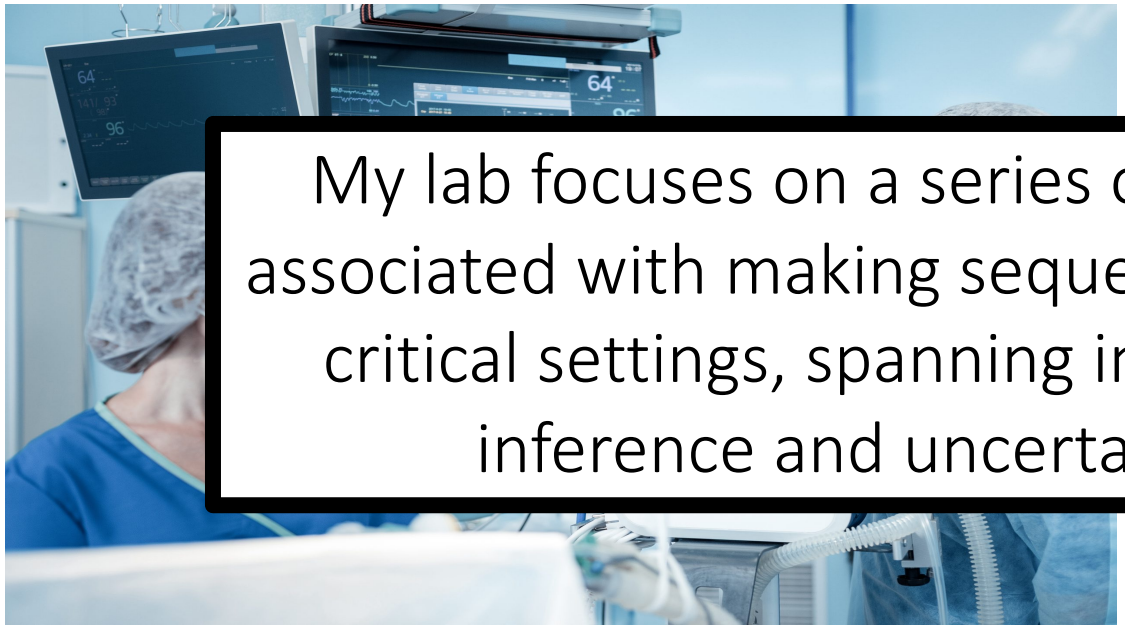


- No good simulator of human physiology, behaviour or learning
- Gathering small amounts of relevant data from EHRs
- Complex multi-objectives
- Safety, generalization critical
- Non-stationary settings
- Frequently subject to confounding




- Simulator of domain
- Enormous amounts of data to train on
- Can actively interact with environment and try out new strategy

# Making sequential decisions in healthcare is not so easy



My lab focuses on a series of technical questions associated with making sequential decisions in safety critical settings, spanning interpretability, causal inference and uncertainty estimation.




- No good simulator of human physiology, behaviour or learning
- Gathering small amounts of relevant data from EHRs
- Complex multi-objectives
- Safety, generalization critical
- Non-stationary settings
- Frequently subject to confounding

- Simulator of domain
- Enormous amounts of data to train on
- Can actively interact with environment and try out new strategy


Clinical Trial Design
Drug Discovery
Clinical Practice
Physicians
Life scientists
Patients
Therapy Scheduling
Reinforcement Learning
Statistical/Machine Learning
Causal Inference
Research
Interpretable ML
Bayesian Decision Theory

# Lesson 1: Data is often small, so we should identify how to make best use of it

Amidst the hype of big data and machine learning, most problems in fact suffer from issues relating to **small data:**

- *Scenario 1:  You want to build a translation system for a low-resource language. We may have a lot of information about Italian available online, but with Rhaeto-Romance languages the availability of usable data is more complicated.*

- *Scenario 2: Data distribution might model the world well, but you cannot collect all that information because it's too expensive or simply isn't commonly observed.*

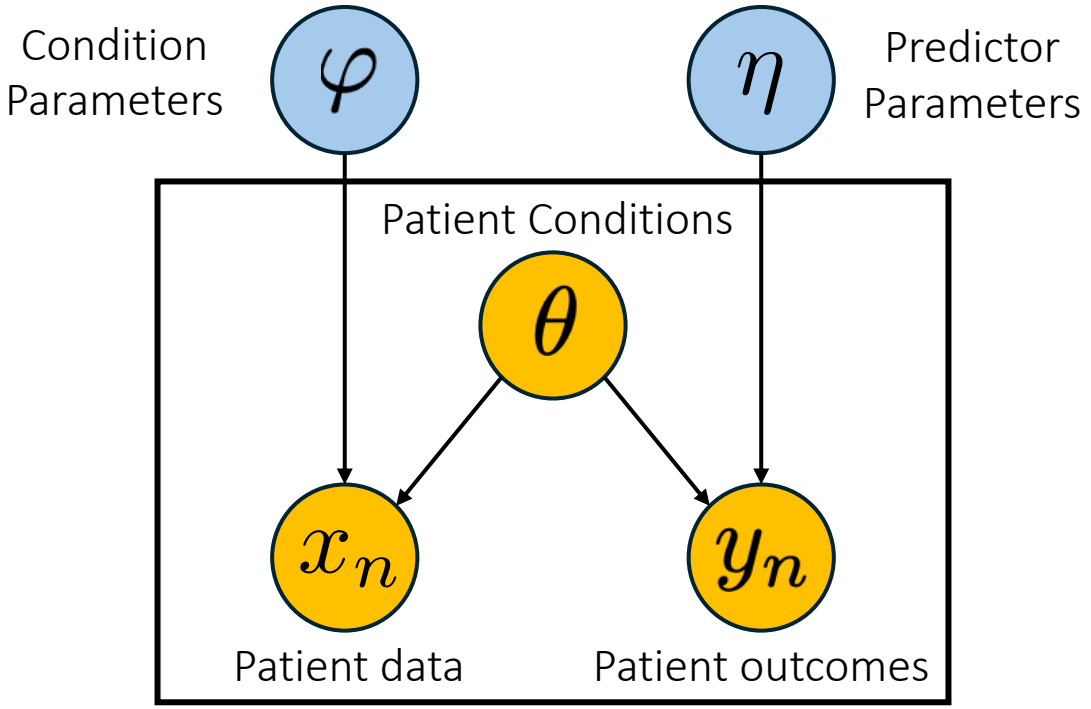
# Lesson 1: Data is often small, so we should identify how to make best use of it

In both Scenario 1 and 2, using a larger model for the data will not work or will be prone to **overfitting**.

If you are clever about it, in both these scenarios you can **build smaller models** which can be inspected and audited by domain experts, that are in fact performant

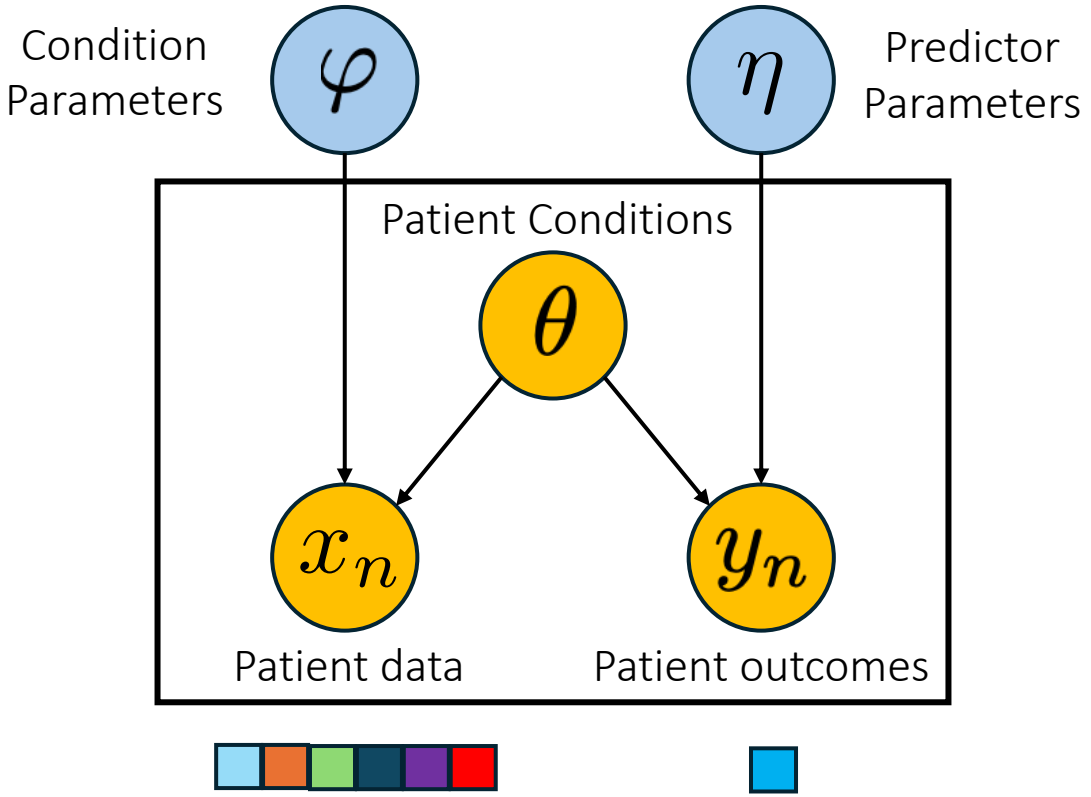
# Task-Focused Modeling

- Idea: Create small, inspectable generative models by using tasks to focus on **relevant signal**



# Task-Focused Modeling

- Idea: Create small, inspectable generative models by using tasks to focus on **relevant signal**



# A Task-Constrained Objective

$$\min_{\phi,\eta} -\sum_n \log p(x_n|\phi)$$

Explain the data as best you can

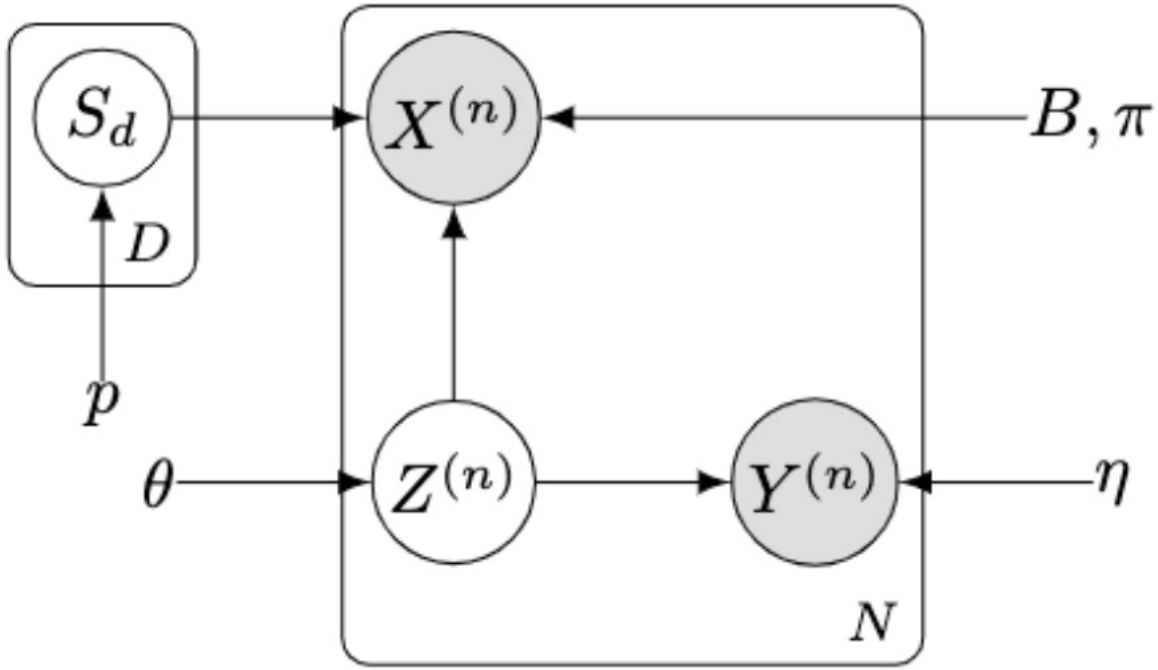
Subject to $$-\sum_n \log p(y_n|x_n,\phi,\eta) < L$$

While making good predictions

This objective captures our goal, but is tricky to optimize as it does not correspond to a valid graphical model

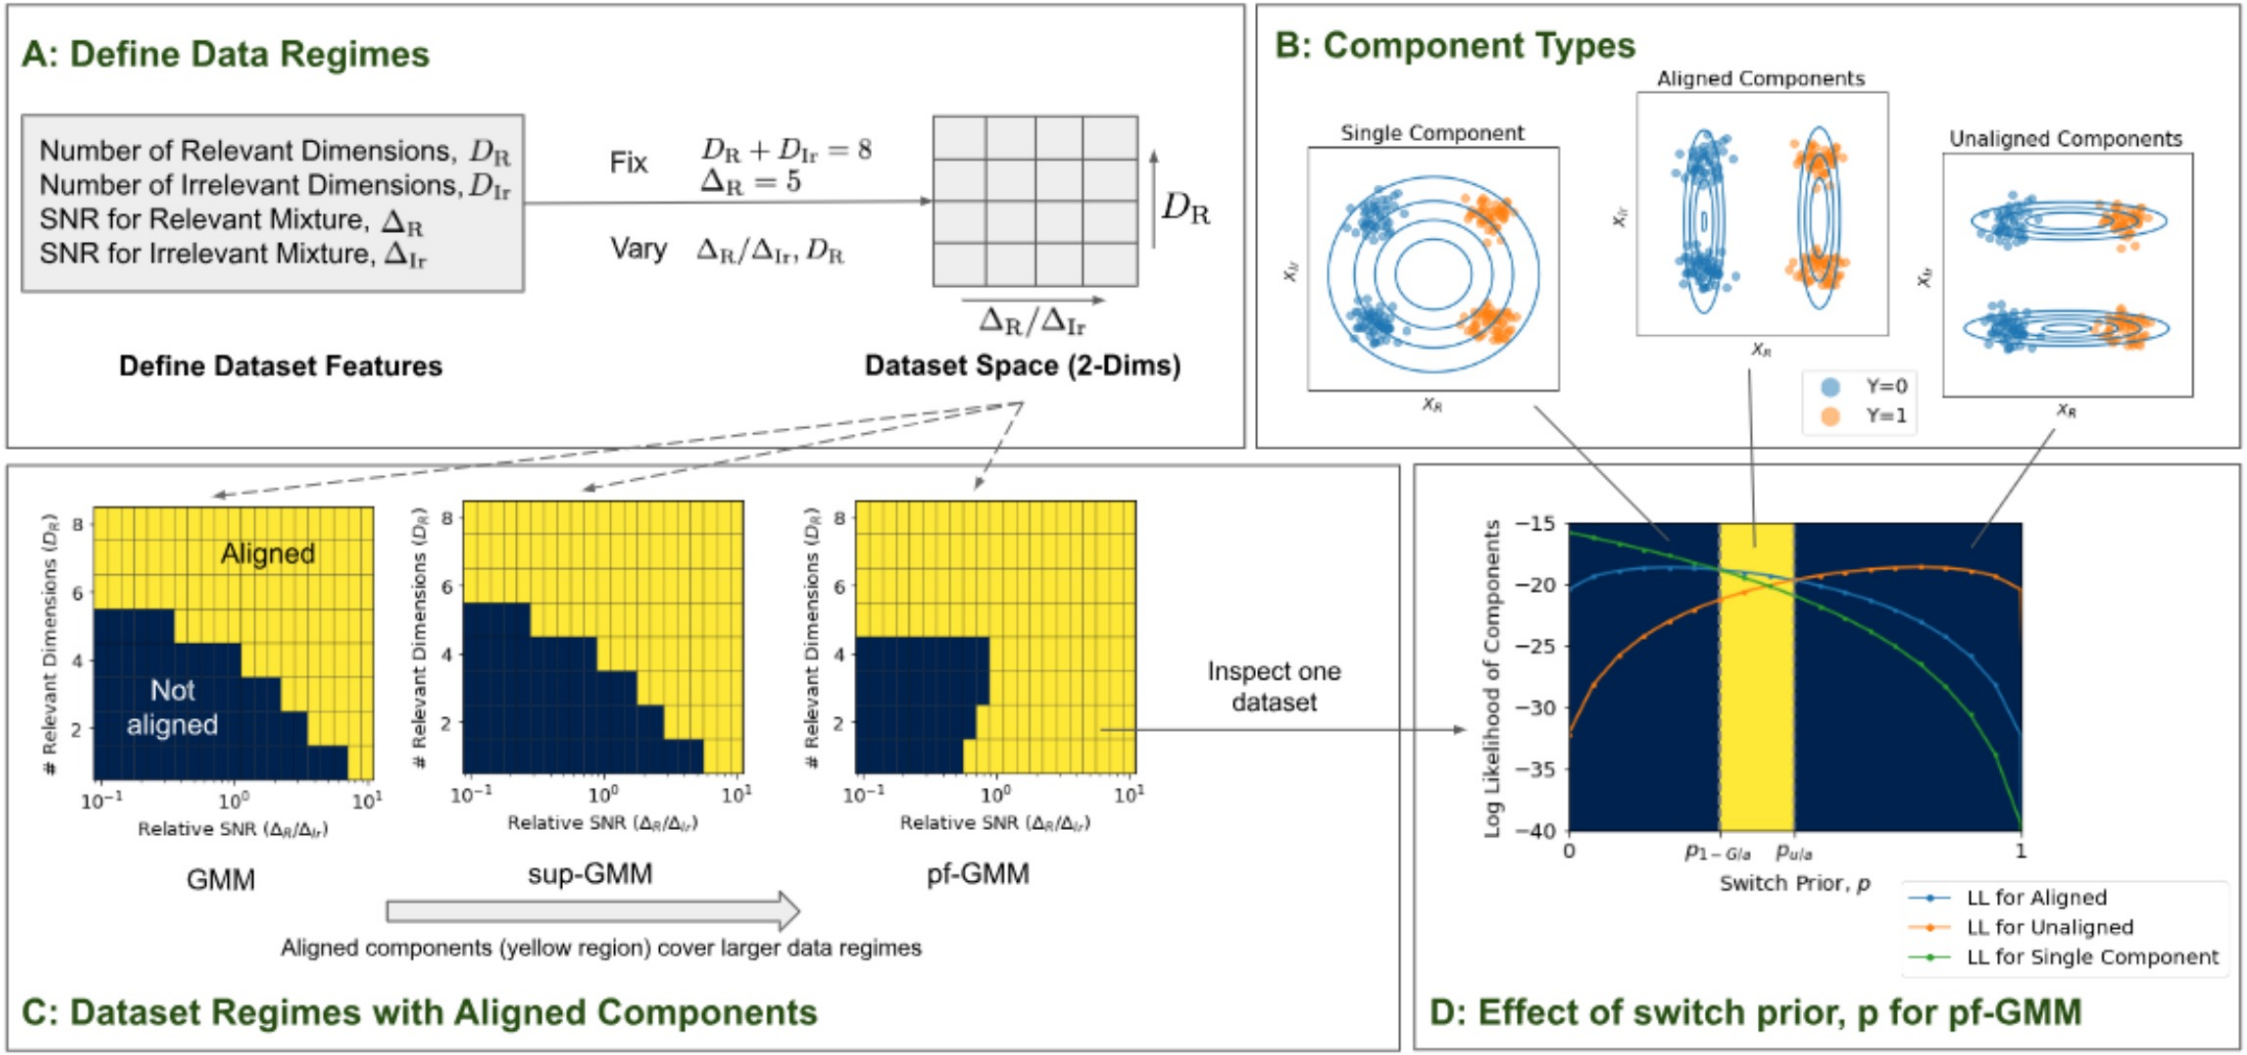
# Our Solution: **Task-focused modeling with simpler inference**

- An alternative to modelling the relationship between relevant and irrelevant dimensions is to assume some dimensions are interesting and others are not.




Task-Relevant Feature Selection with Prediction Focused Mixture Models – Sharma et al 2024, TMLR

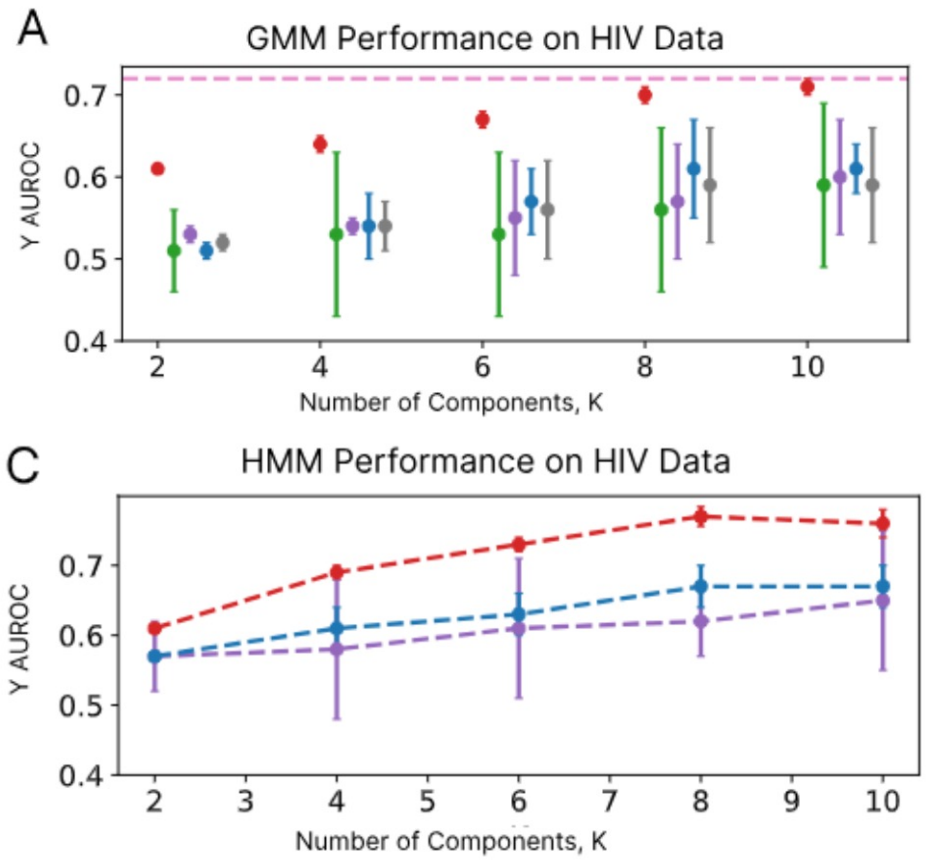
# Task-focused models are **versatile**



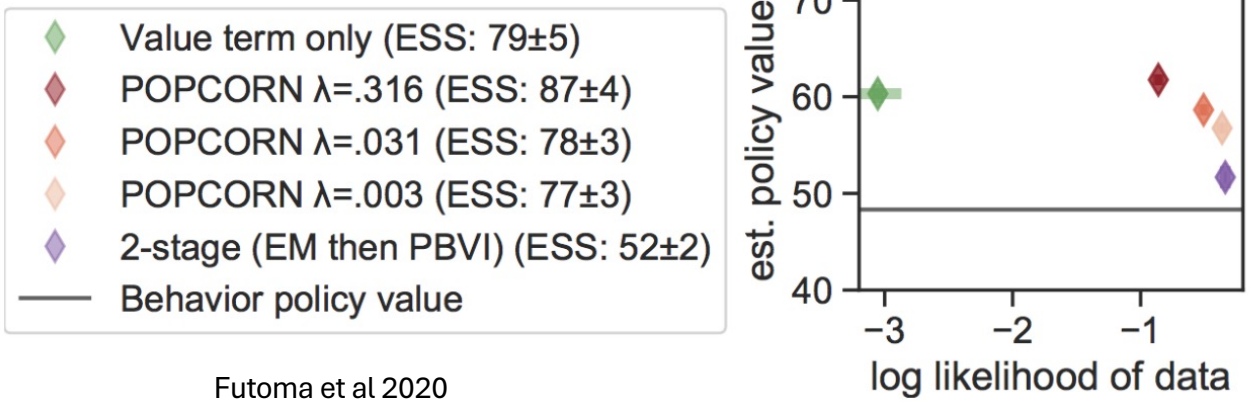

Task-Relevant Feature Selection with Prediction Focused Mixture Models – Sharma et al 2024, TMLR

# Task-focused models are **performant in real health and sequential settings**

HIV management with small, task-focused GMMs and HMMs




Hypotension management in the ICU with small, task-focused models



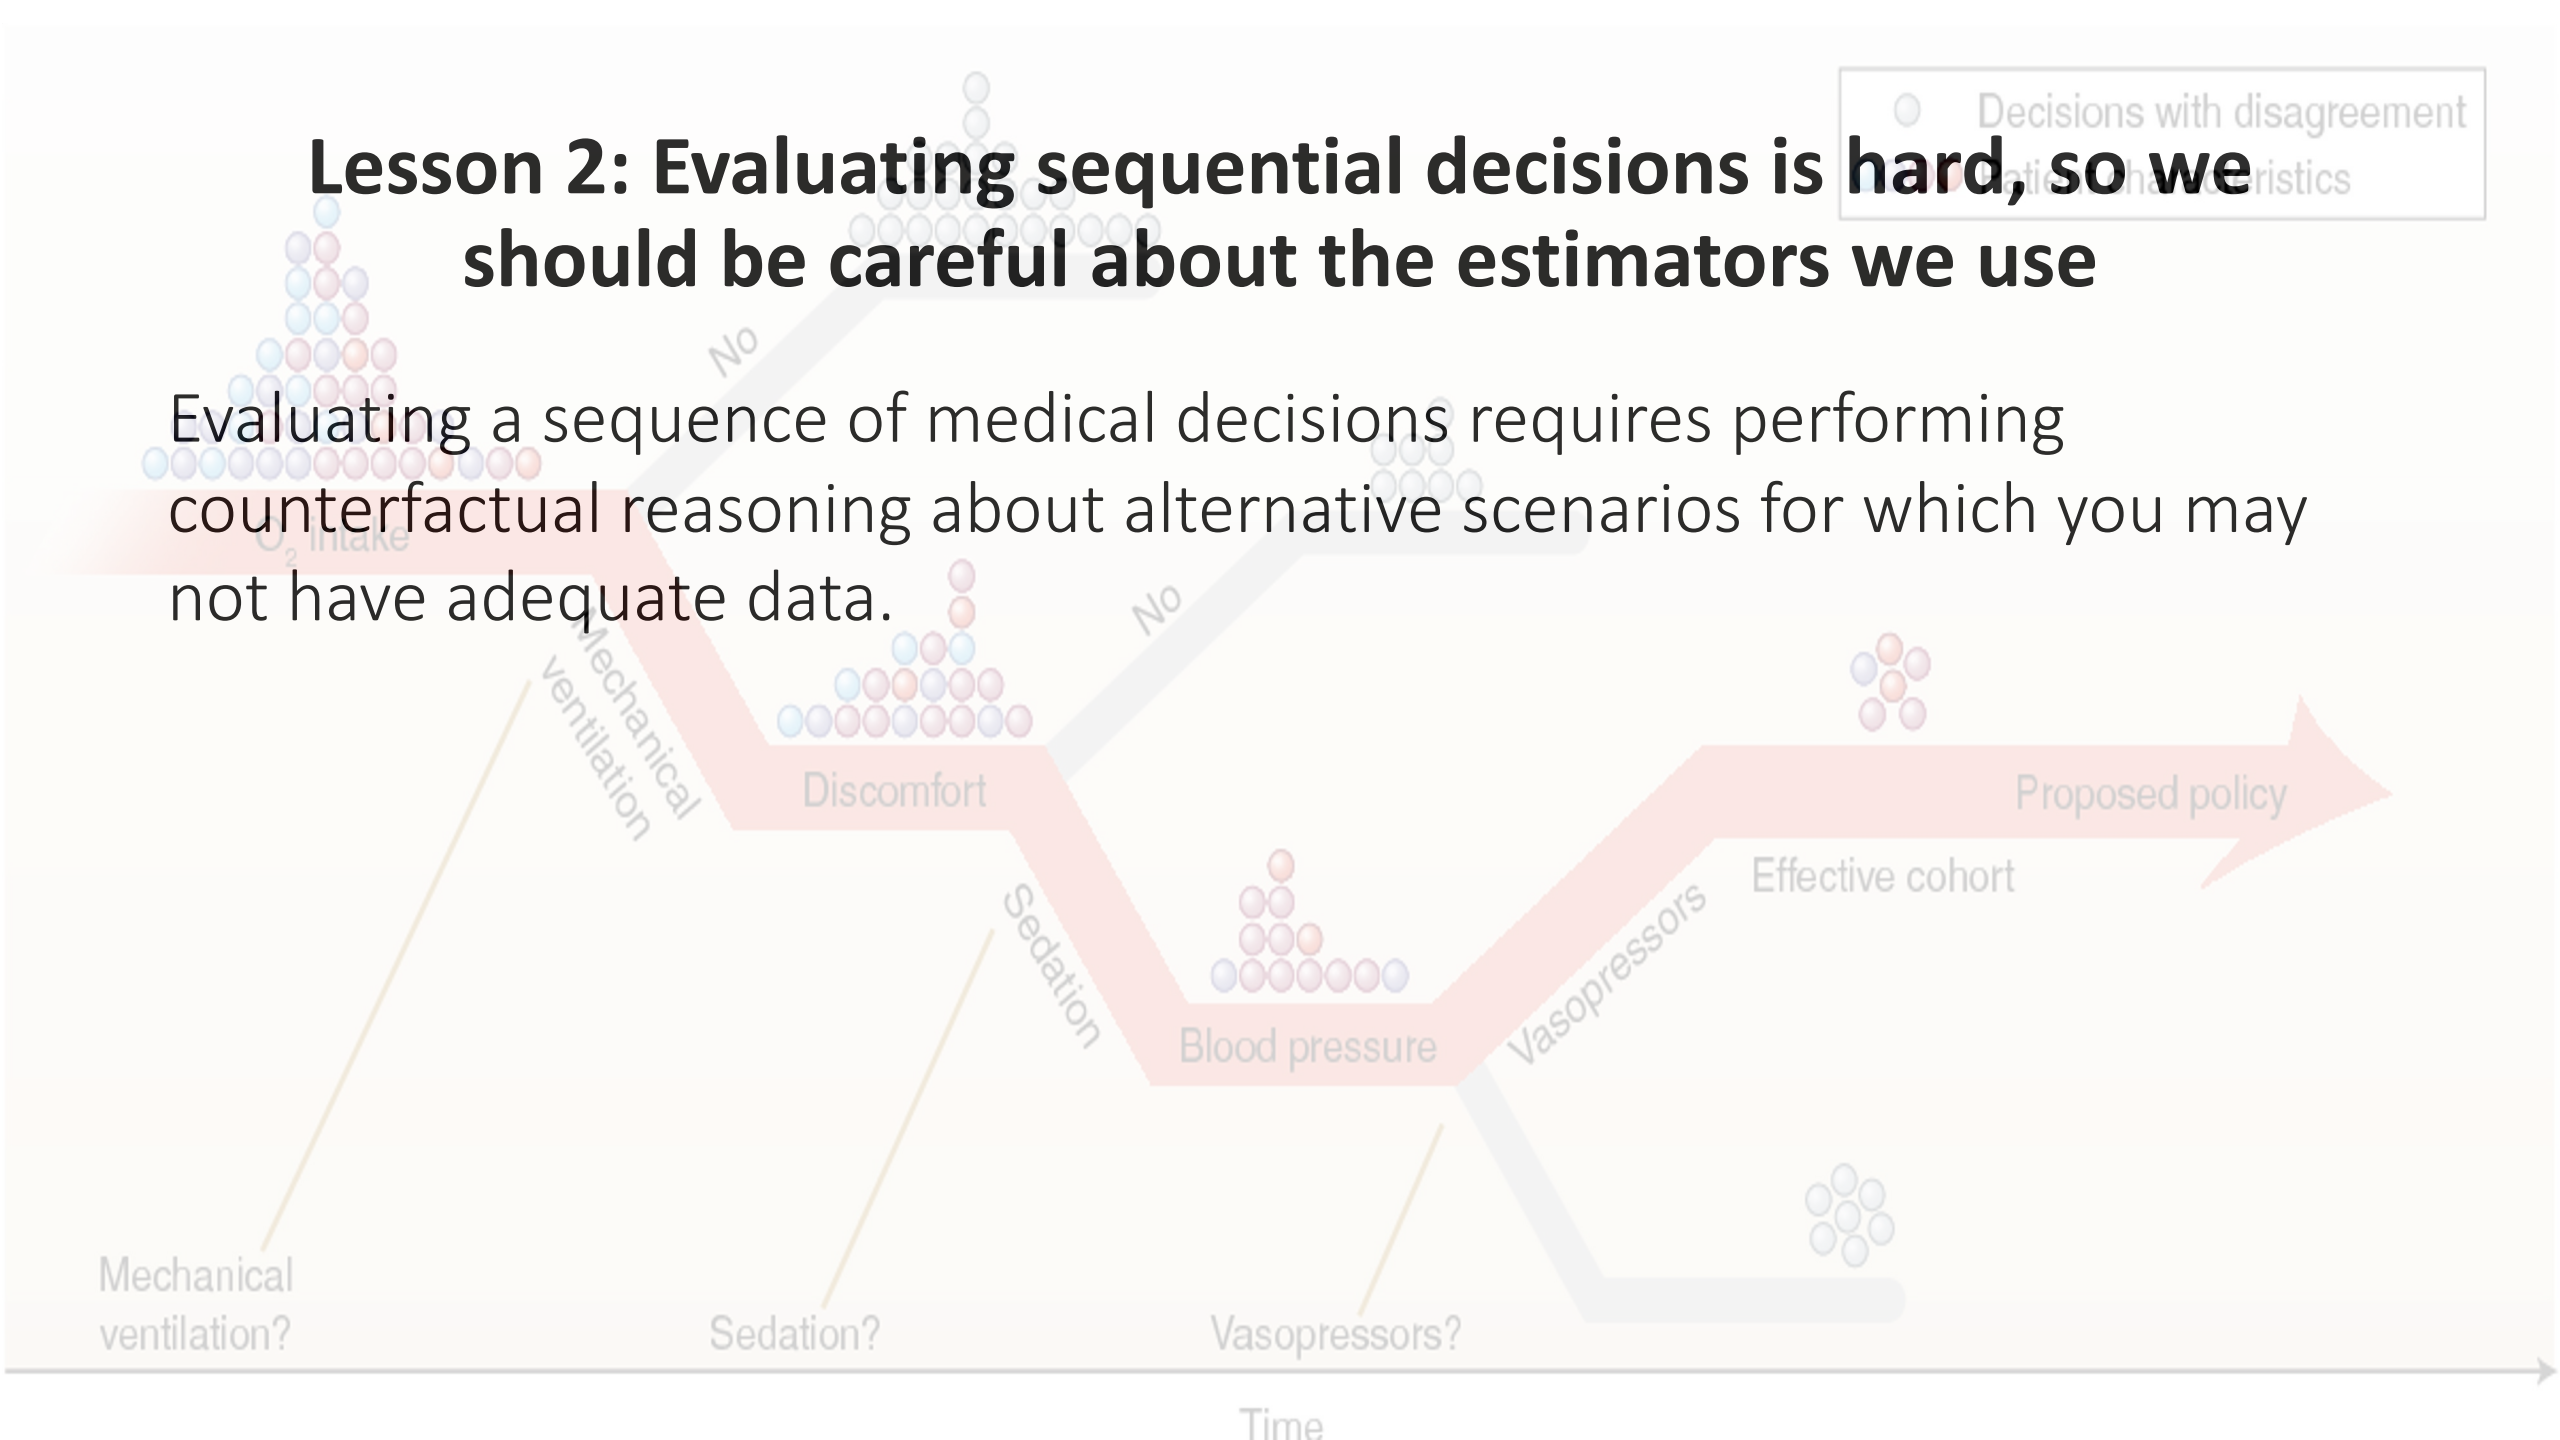

Futoma et al 2020

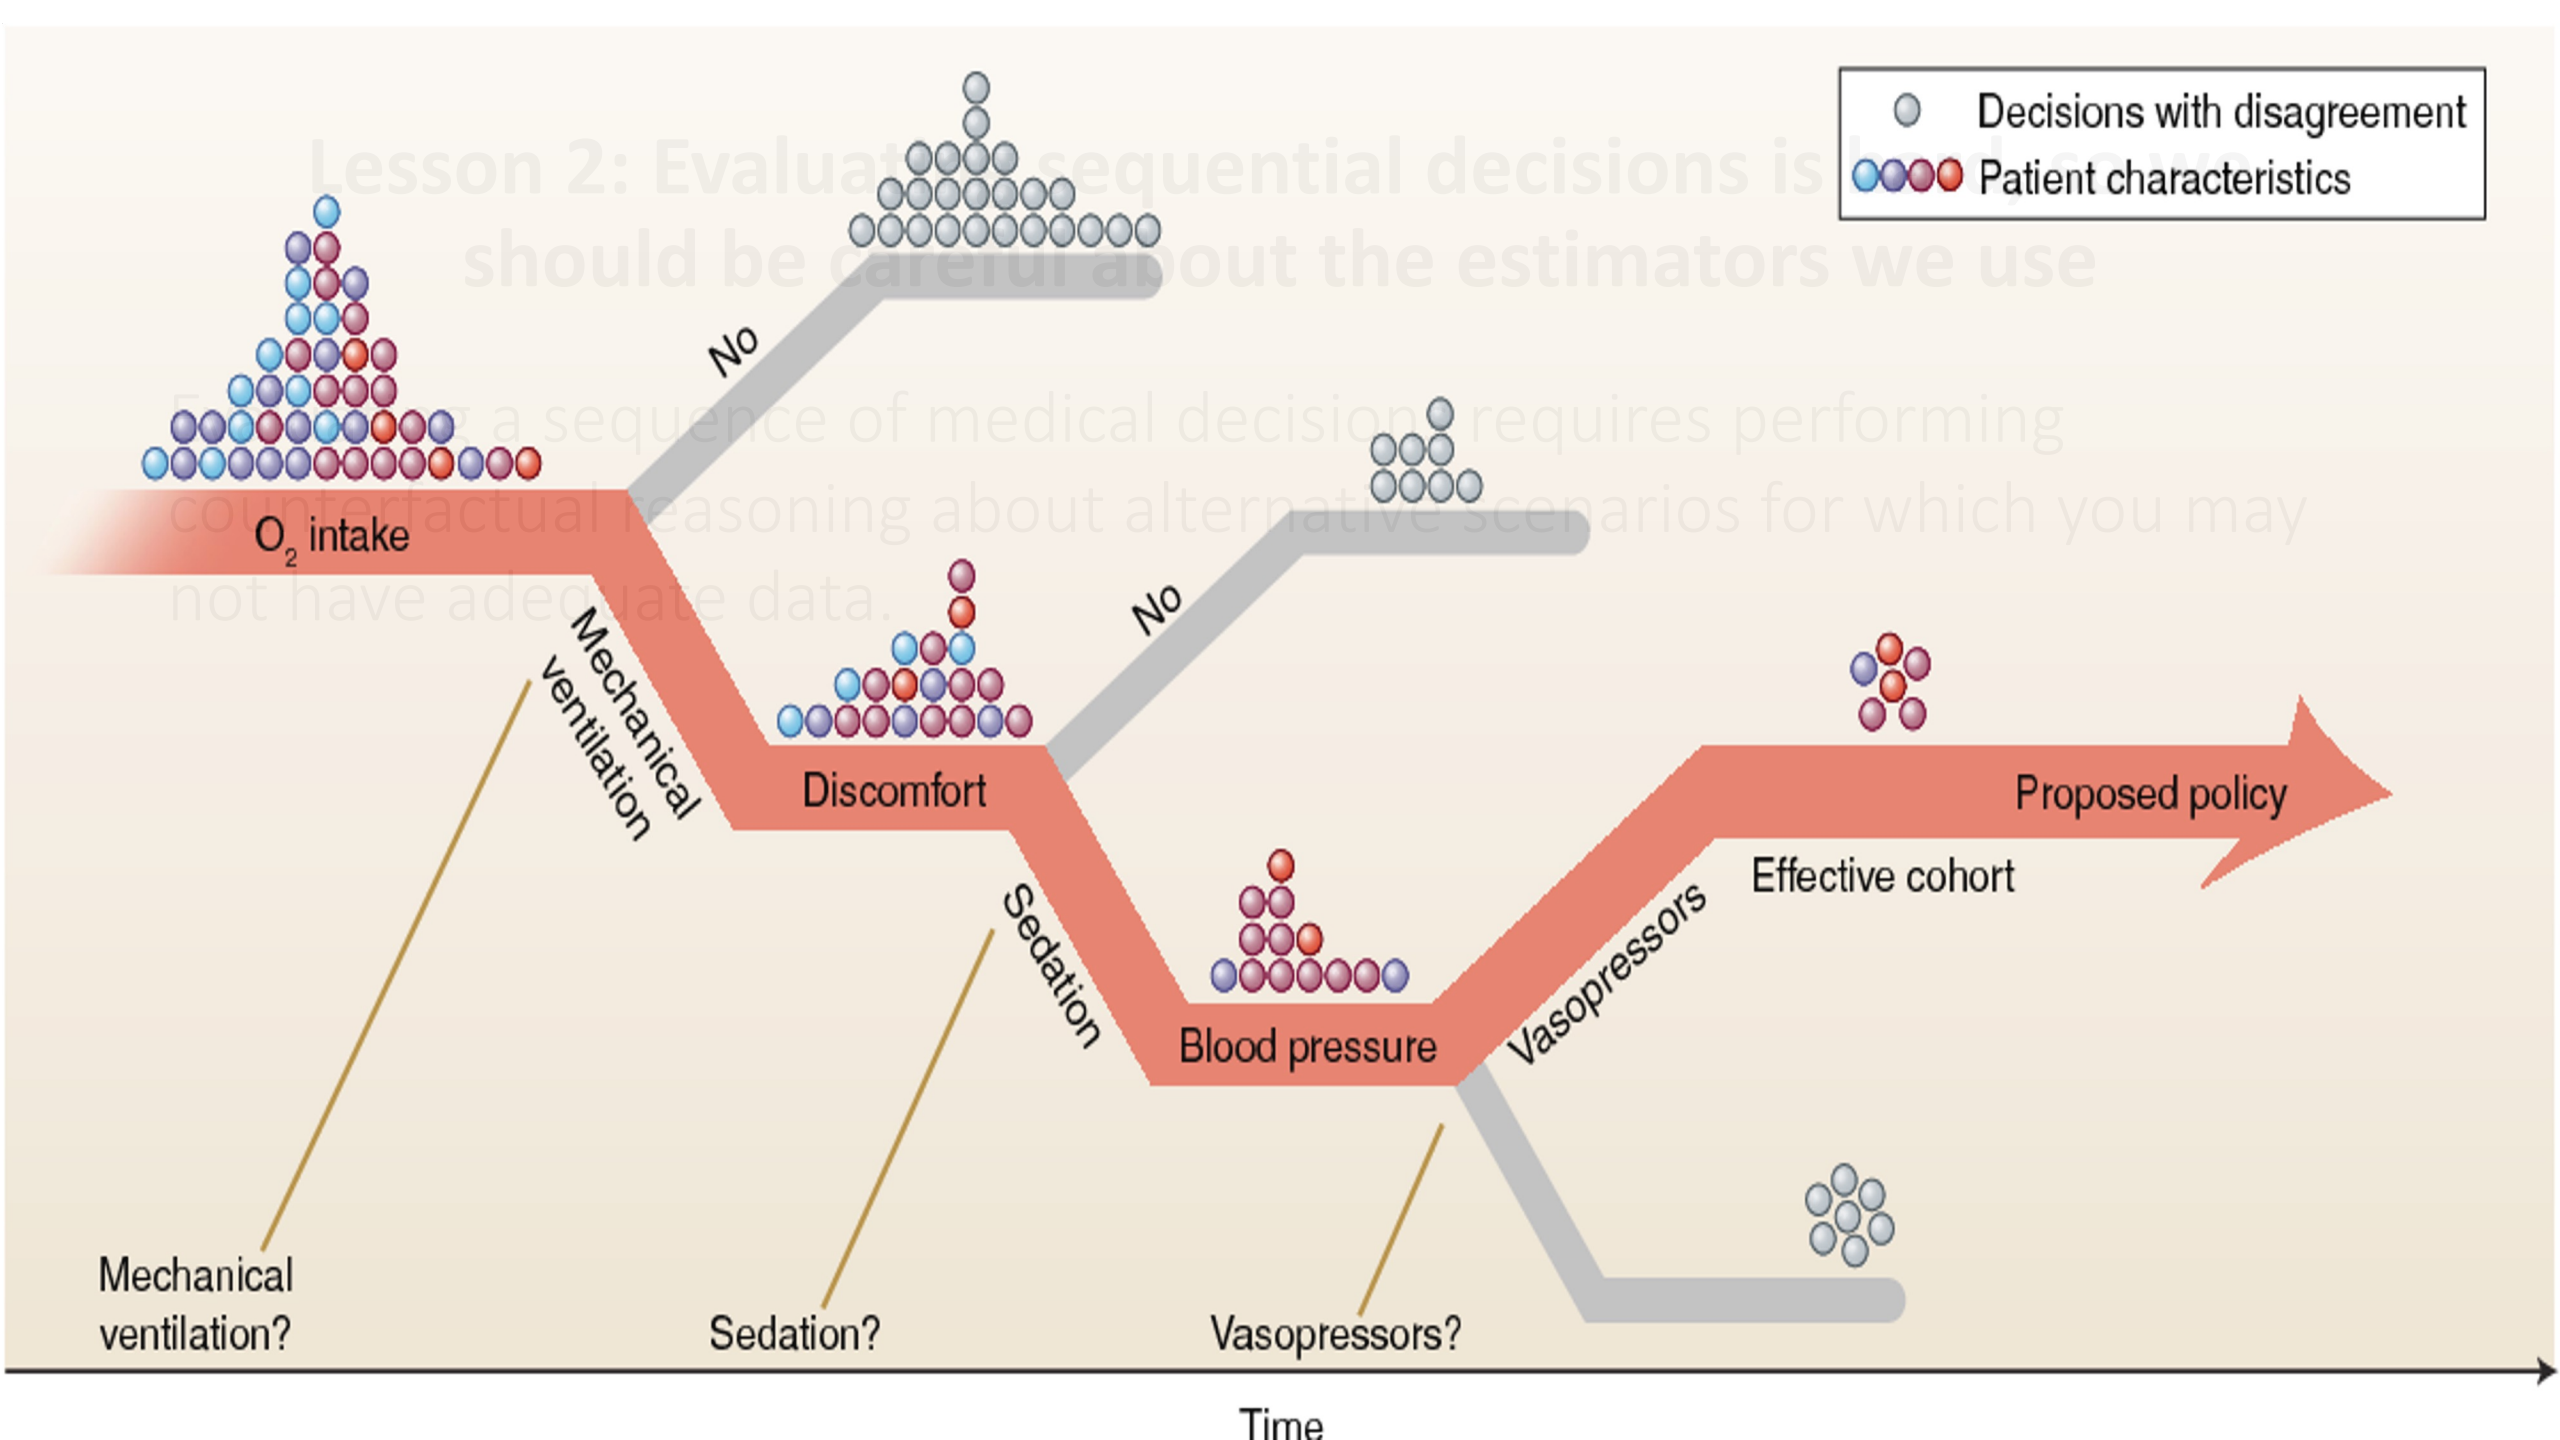
Similar success with smaller models
e.g. 5 states vs. 128 states

Task-Relevant Feature Selection with Prediction Focused Mixture Models – Sharma et al 2024, TMLR

# Lesson 2: Evaluating sequential decisions is hard, so we should be careful about the estimators we use

Evaluating a sequence of medical decisions requires performing counterfactual reasoning about alternative scenarios for which you may not have adequate data.




Decisions with disagreement
Patient characteristics
$O_2$ intake
No
Mechanical ventilation
Discomfort
No
Sedation
Blood pressure
Vasopressors
Effective cohort
Proposed policy
Mechanical ventilation?
Sedation?
Vasopressors?
Time

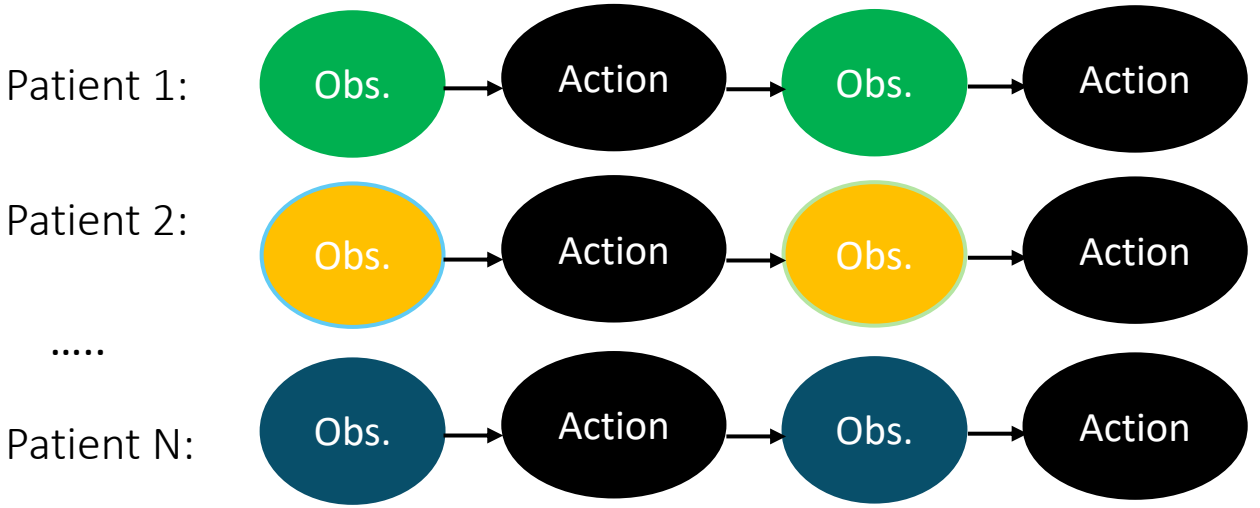
# Off-Policy Evaluation

Given data collected by some **behaviour policy** $\pi_b$, can we estimate the value of **another evaluation policy** $\pi_e$?

Clinician:



$\pi_b$




Evaluation Agent:

$\pi_e$

Reweight trajectories according to how much they overlap using:

$$\hat{V}_{\pi_e}^{PDIS} = \frac{1}{N}\sum_{n=1}^{N}\sum_{t=0}^{T}\gamma^t \cdot \rho_{0:t}^{(n)} \cdot r_t^{(n)} \qquad \rho_{0:T}^{(n)} = \prod_{t=0}^{T}\frac{\pi_e(a_{t+1}^{(n)},\ s_t^{(n)})}{\pi_b(a_{t+1}^{(n)},\ s_t^{(n)})}$$

# Our Insight: Rather than trying to match entire policies, match only the relevant groups of actions or subspaces

Find each unit's corresponding counterfactual 'twin' by considering only the **relevant subspaces** and check those outcomes.



Obama had he gone to law school



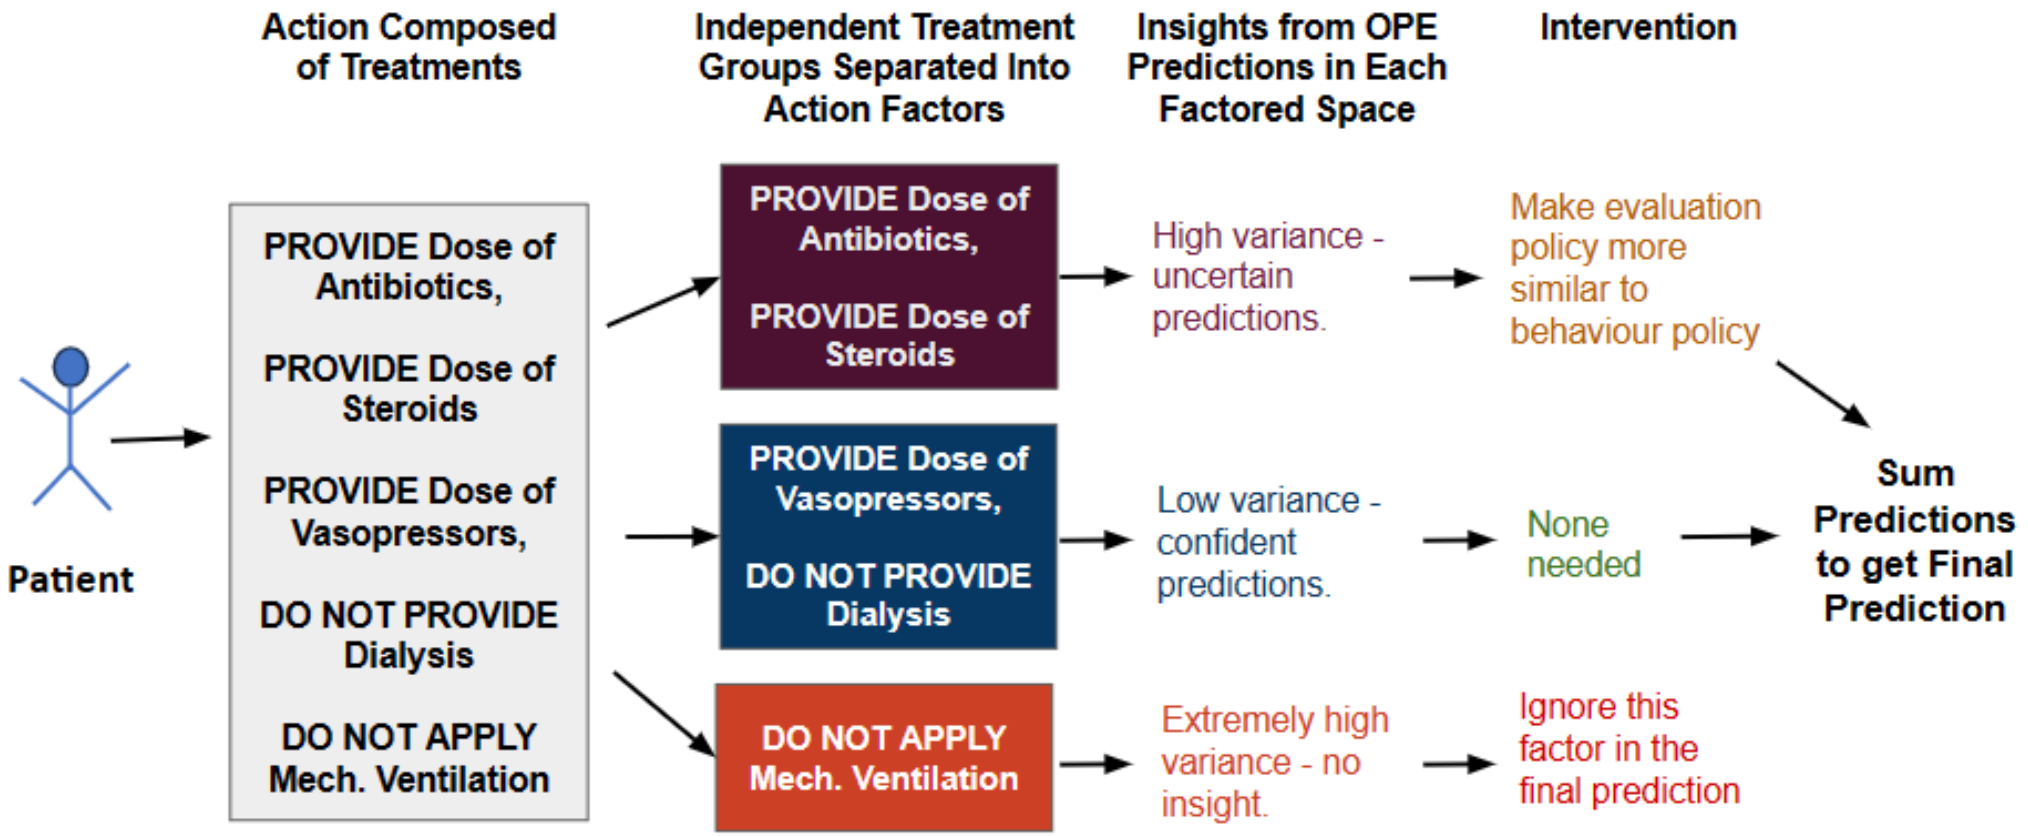
Obama had he gone to business school

vs



Other US presidents

# Our Solution: Decompose evaluation based on constituent factors that account for different subspaces




Rebello, A, Tang, S., Wiens, J and Parbhoo, S, "Leveraging Factored Action Spaces for Off-Policy Evaluation", in submission 2024.

# Our Solution: Decompose evaluation based on constituent factors that account for different subspaces

- **Decompose each action** into constituent factors. These may either be determined from using human expertise or learnt.

Intuition: It is easier to induce overlap across some factored dimensions than across an entire trajectory

**IS reweighting:**

$$\hat{V}_{\pi_e}^{PDIS} = \frac{1}{N}\sum_{n=1}^{N}\sum_{t=0}^{T}\gamma^t \cdot \rho_{0:t}^{(n)} \cdot r_t^{(n)}$$

$$\rho_{0:t}^{(n)} = \prod_{t'=0}^{t}\frac{\pi_e(a_{t'+1}^{(n)},\ s_{t'}^{(n)})}{\pi_b(a_{t'+1}^{(n)},\ s_{t'}^{(n)})}$$

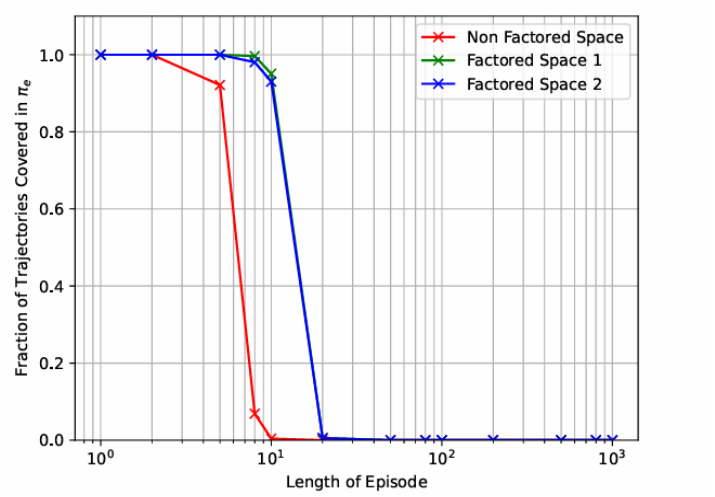
**Decomposed IS reweighting:**

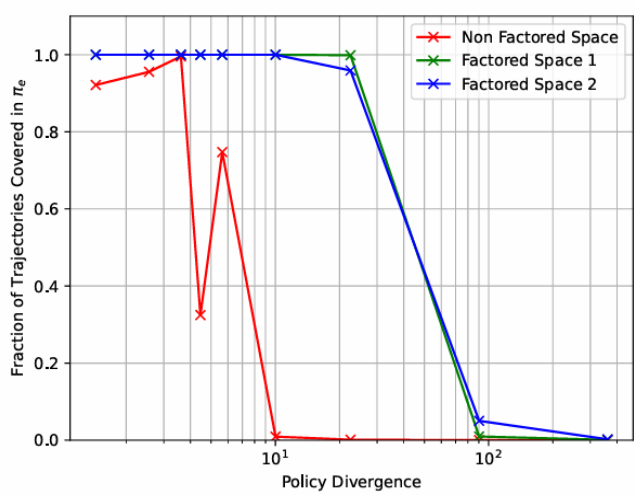
$$\hat{Q}_{\pi_e}^{DecIS} = \sum_{d=1}^{D}\frac{1}{N}\sum_{n=1}^{N}\sum_{t=0}^{T}\gamma^t \cdot \rho_{0:T}^{(n),d} \cdot r_t^{(n),d}$$

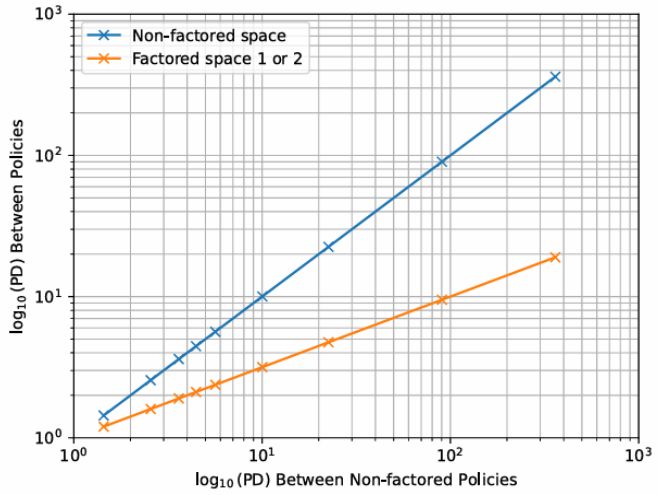
$$\rho_{0:T}^{(n),d} = \prod_{t=0}^{T}\frac{\pi_e^d(a_t^{(n),d}|z_t^{(n),d})}{\pi_b^d(a_t^{(n),d}|z_t^{(n),d})}$$

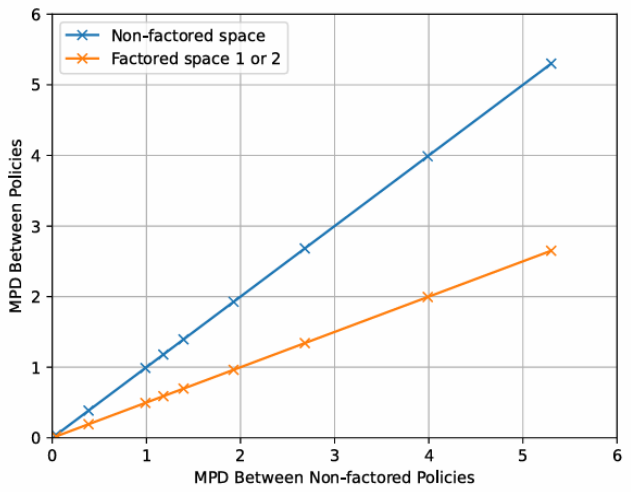
Rebello, A, Tang, S., Wiens, J and Parbhoo, S, “Leveraging Factored Action Spaces for Off-Policy Evaluation”, in submission 2024.

# Factored OPE estimators have better overlap



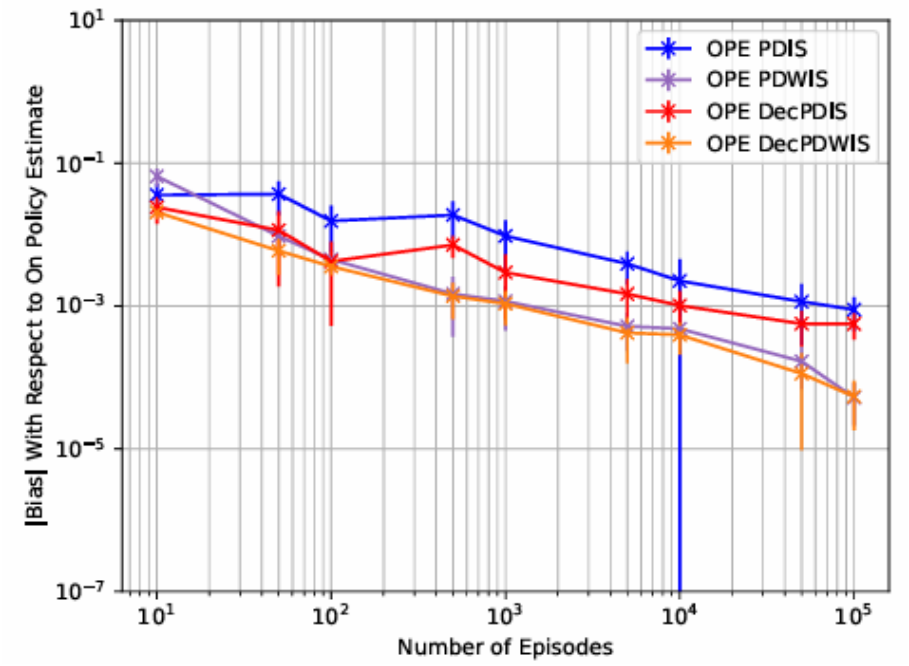



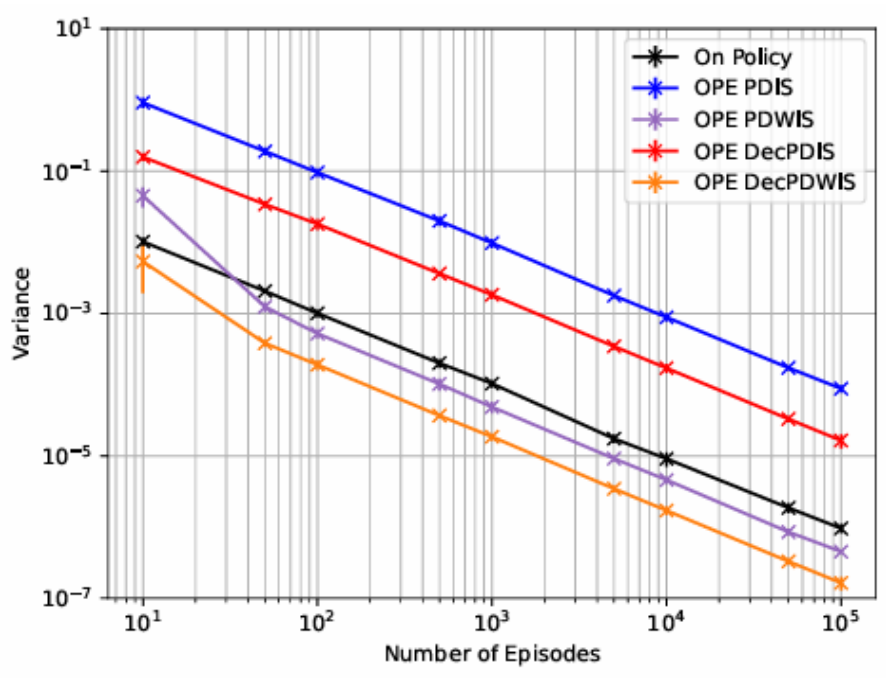




Rebello, A, Tang, S., Wiens, J and Parbhoo, S, “Leveraging Factored Action Spaces for Off-Policy Evaluation”, in submission 2024.

# Factored estimators have lower variance and bias

By grouping patients by relevant action factors, we obtain better statistical guarantees and performance with less data.

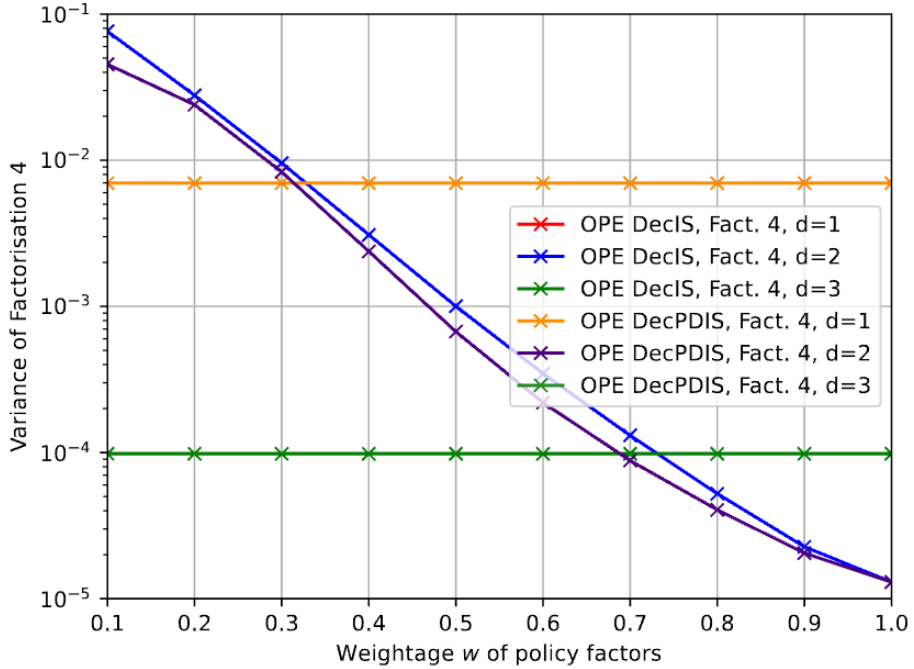


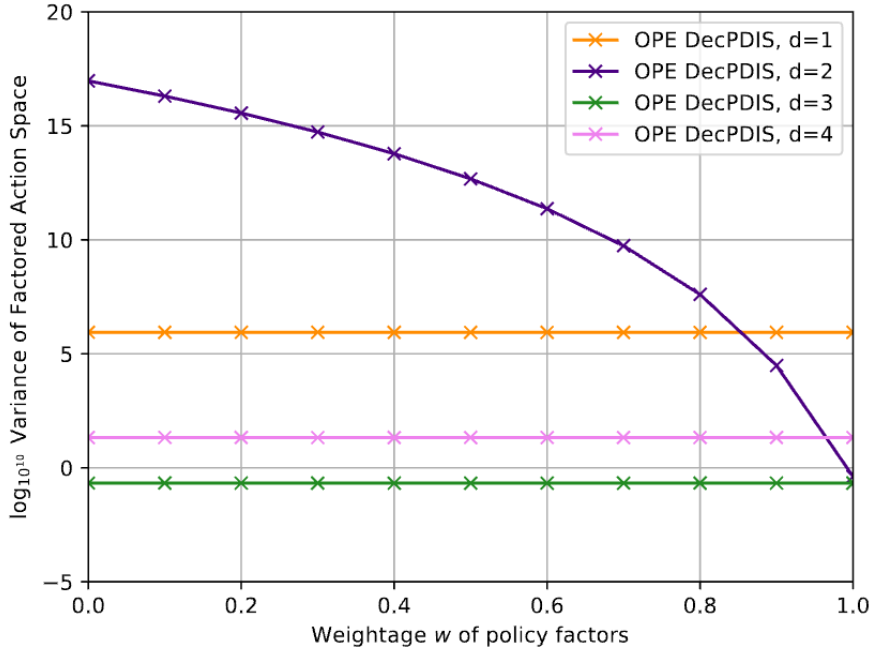




Rebello, A, Tang, S., Wiens, J and Parbhoo, S, “Leveraging Factored Action Spaces for Off-Policy Evaluation”, in submission 2024.

# Factors can be intervened on to remedy high variance OPE

We can isolate those subproblems that contribute to the highest variance in OPE and intervene on those factors to remedy the issues. We might choose to ignore those factors or specifically elicit human expertise for those factors accordingly.




Rebello, A, Tang, S., Wiens, J and Parbhoo, S, "Leveraging Factored Action Spaces for Off-Policy Evaluation", in submission 2024.

# Lesson 3: Causality matters, especially in sequential decision-making

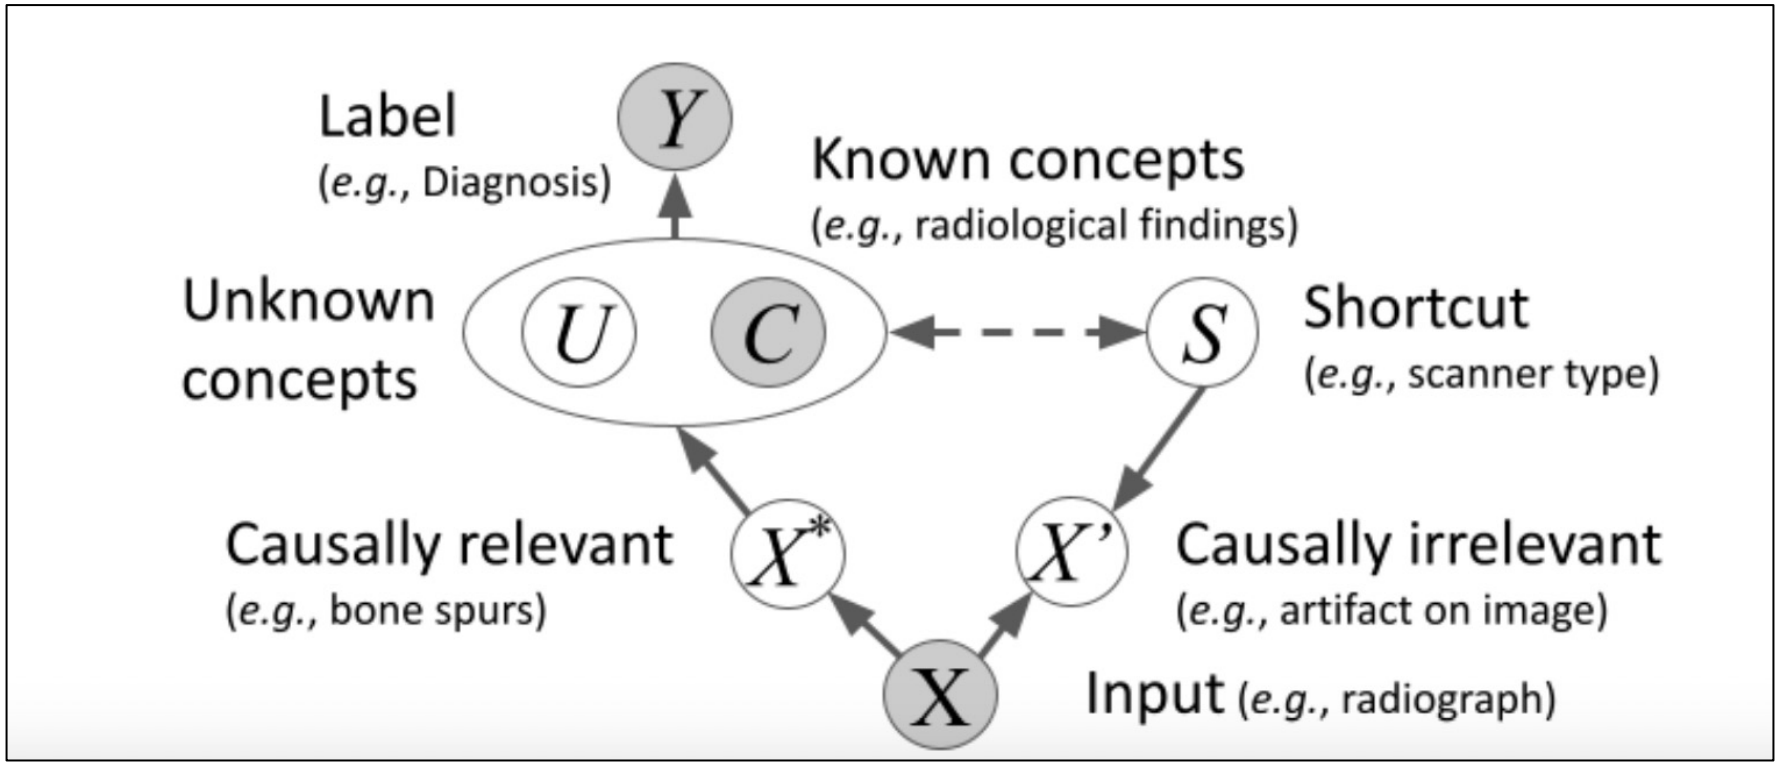
Both learning task-focused representations and factored OPE methods require identifying **relevant information** or subgroups for performing decision-making and evaluation.

There are also several tools for representation learning e.g. S-learners, T-Learners, X-Learners and a long list of causal representation learning methods for estimating treatment effects. These methods are thought to **implicitly learn causally relevant information** for predicting treatment effects.

*Should we care?*

# Lesson 3: Causality matters, especially in sequential decision-making

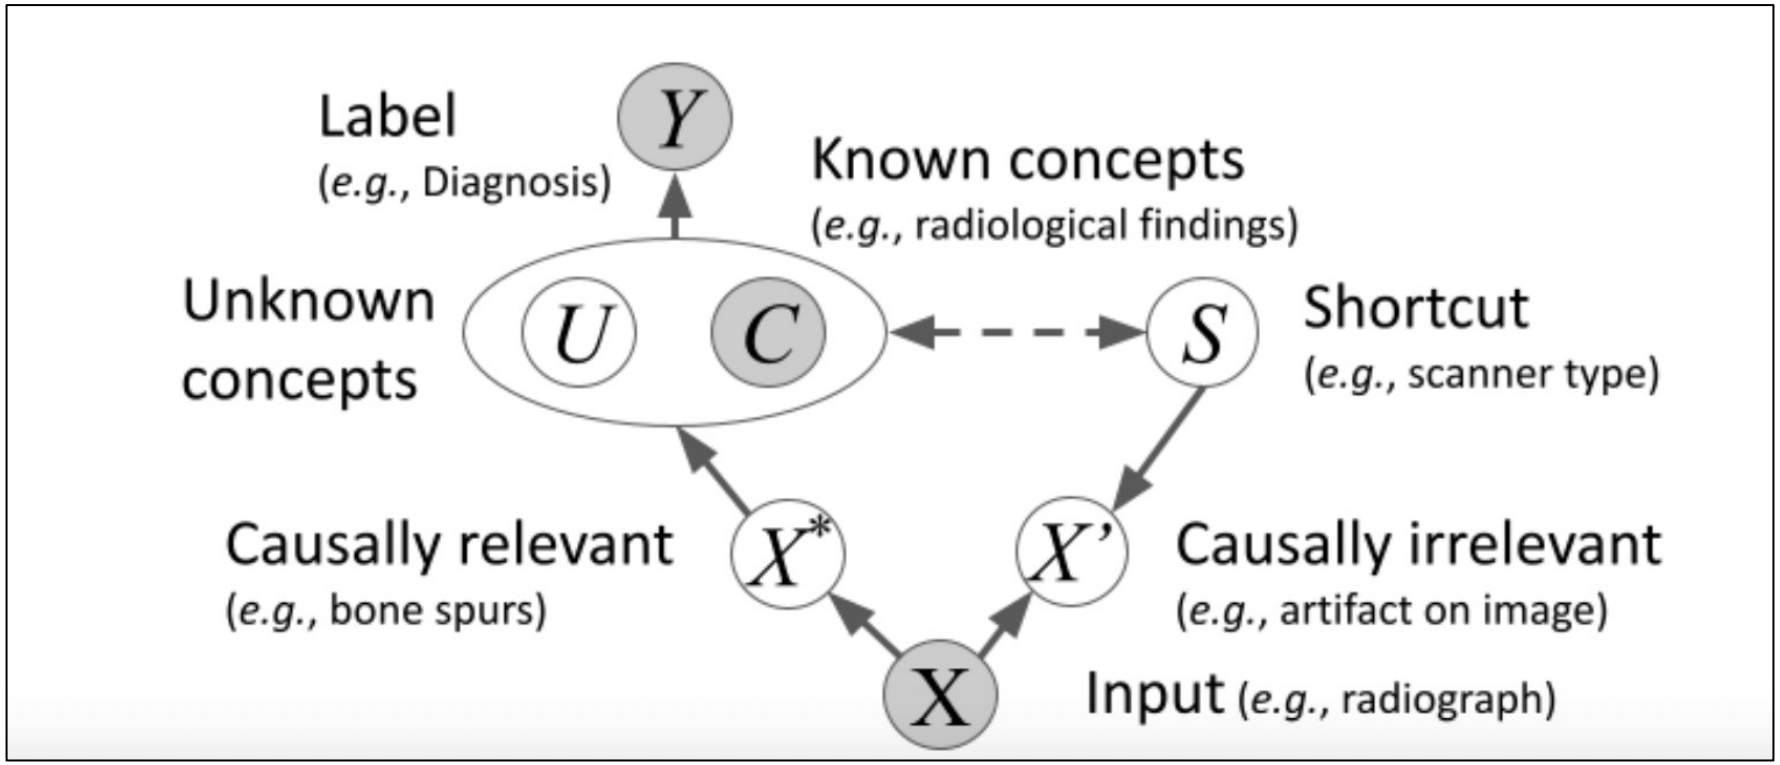
- What happens if **'relevant' information is confounded**?




Learning Concept Credible Models for Mitigating Shortcuts – Wang, Jabbour, Makar, Sjoding, Wiens, NeurIPS 2023

# Lesson 3: Causality matters, especially in sequential decision-making

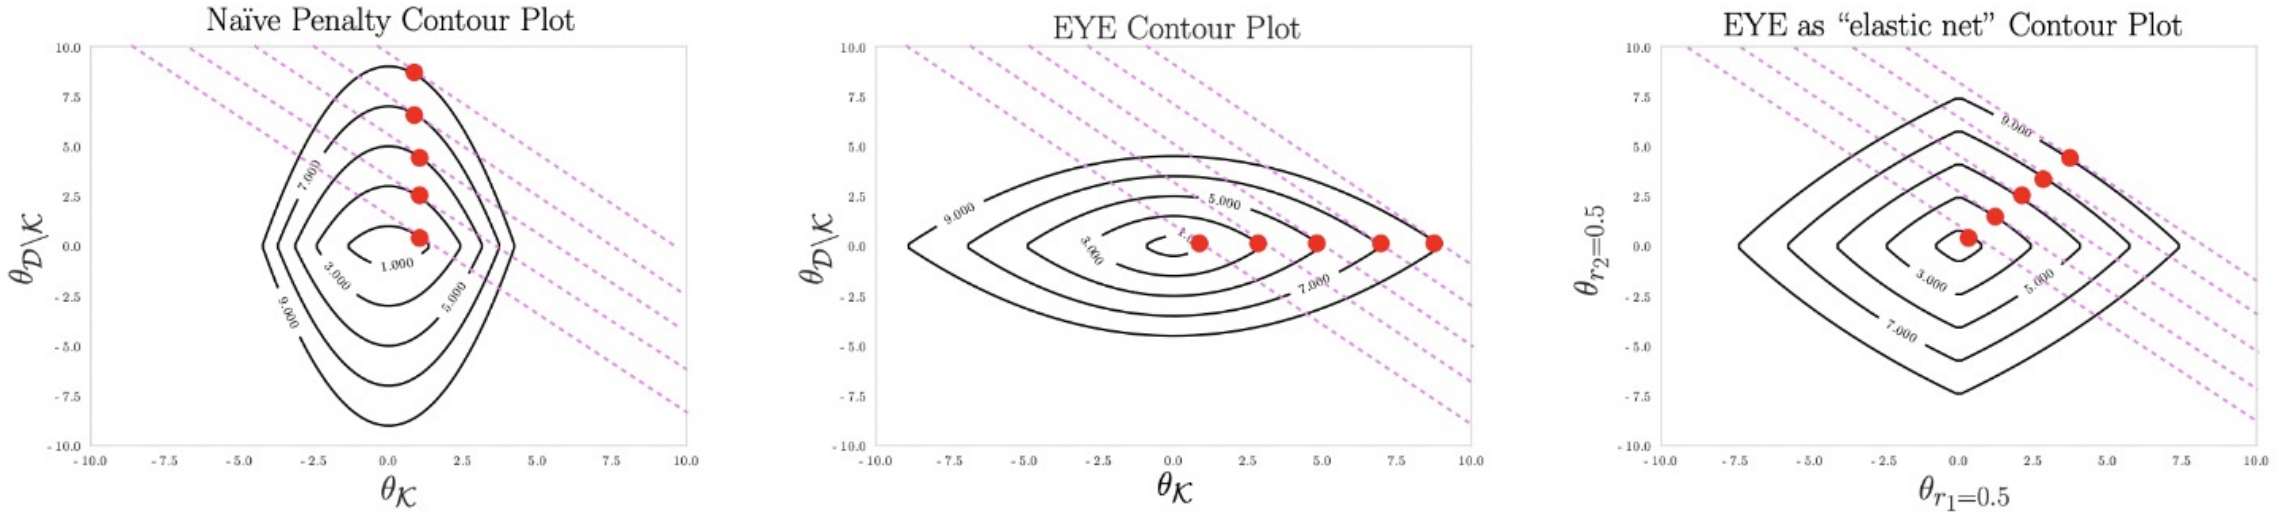
- What happens if **'relevant' information is confounded**?




Learning Concept Credible Models for Mitigating Shortcuts – Wang, Jabbour, Makar, Sjoding, Wiens, NeurIPS 2022

# Can regularization help?

Various regularizers exist to downweight the contributions of potential shortcuts via C and U.

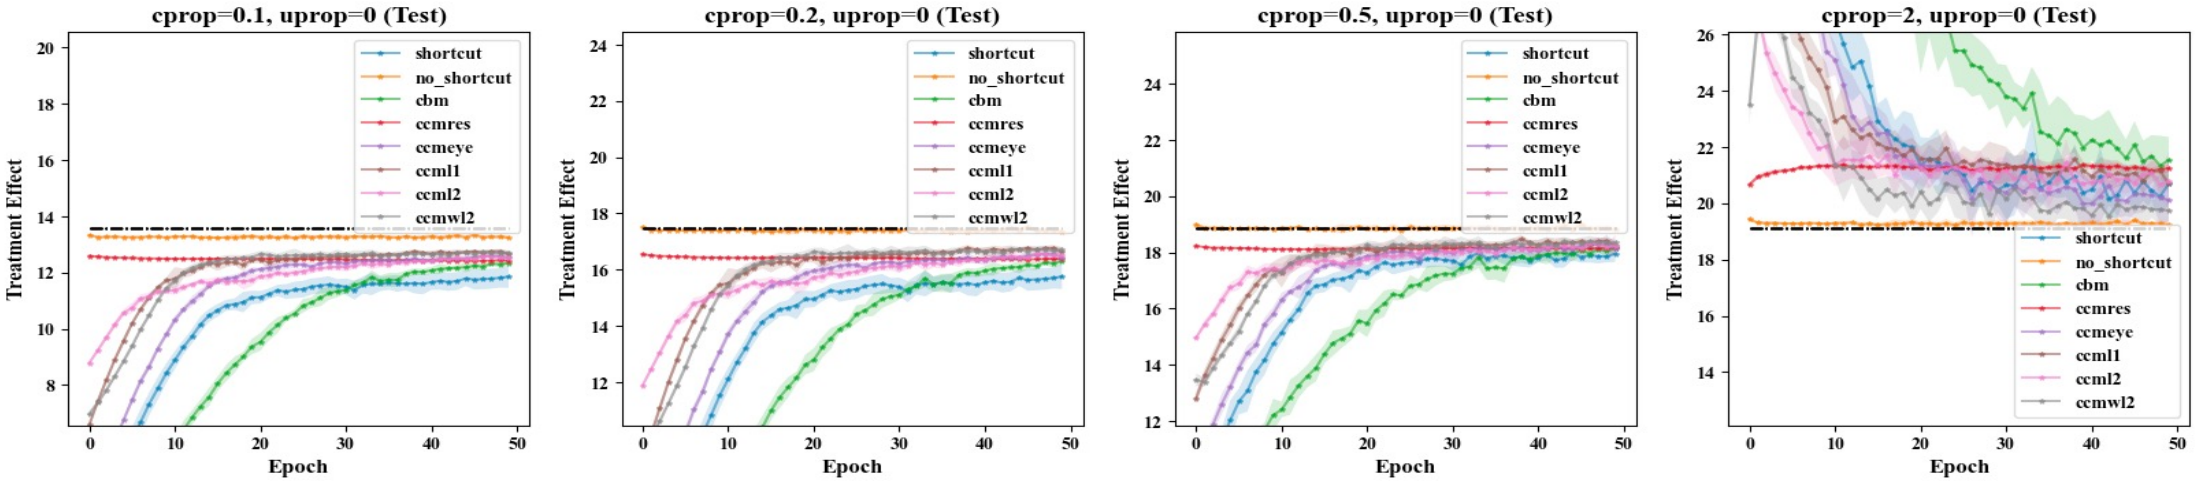



$$||\theta_x||_1 + \sqrt{||\theta_x||_1^2 + ||\theta_c||_2^2}$$

Learning Concept Credible Models for Mitigating Shortcuts – Wang, Jabbour, Makar, Sjoding, Wiens, NeurIPS 2022

# Can regularization help?

- In summary, somewhat! More often than not though, you end up overadjusting the weights which induces further bias downstream.



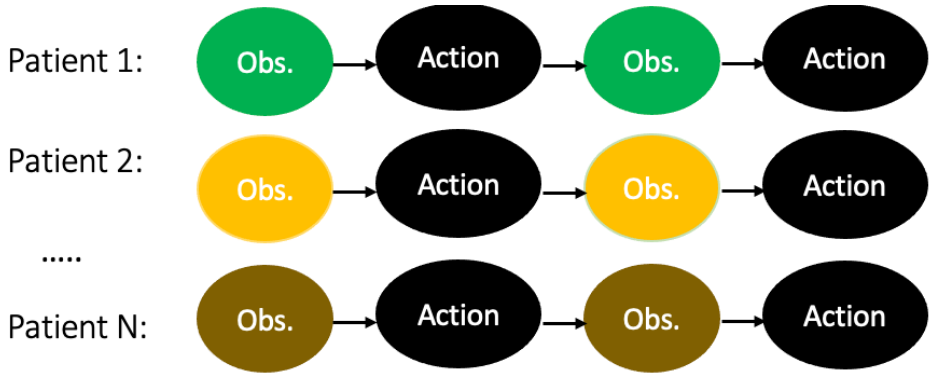

- Similar results apply for other causal representation learning methods
- Thus existing ‘causal’ representation methods alone may not always be adequate for inferring the true treatment effect.

Pitfalls of Using CausalML in Healthcare: the failure of existing approaches in the presence of confounding – Yang, Papanikoulaou, Wiens, Parbhoo. In preparation, 2024.

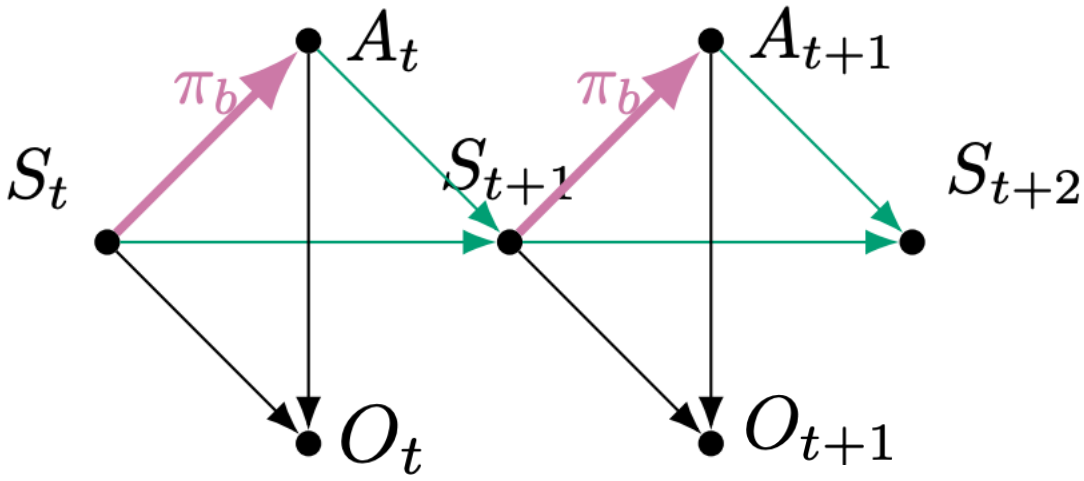
# From Evaluation to Causally-Aware Evaluation

Reweight trajectories according to how much they overlap using:

$$\rho_n = \prod_t \frac{\pi_e(a_{tn}|s_{tn})}{\pi_b(a_{tn}|s_{tn})}$$



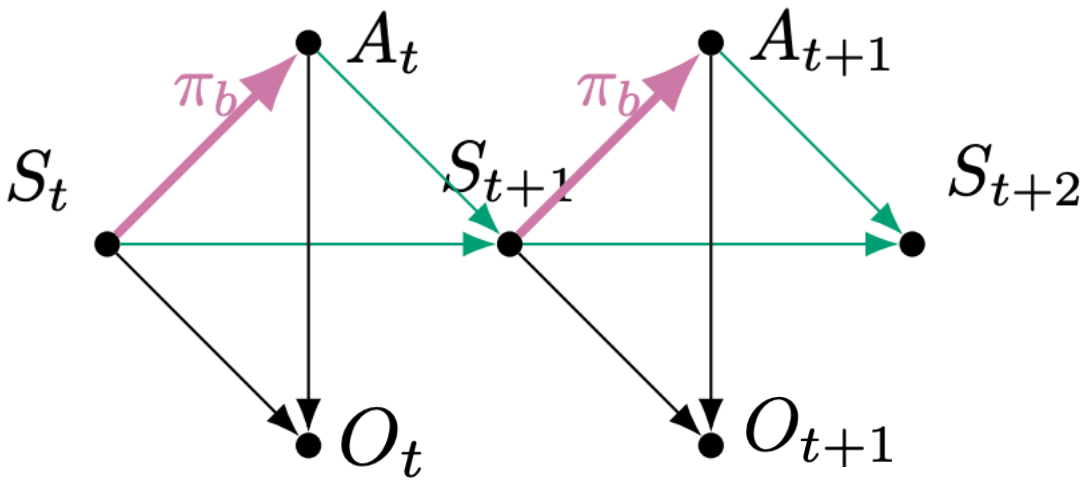

Standard OPE **ignores** the data-generating mechanism resulting in states, actions and observations.

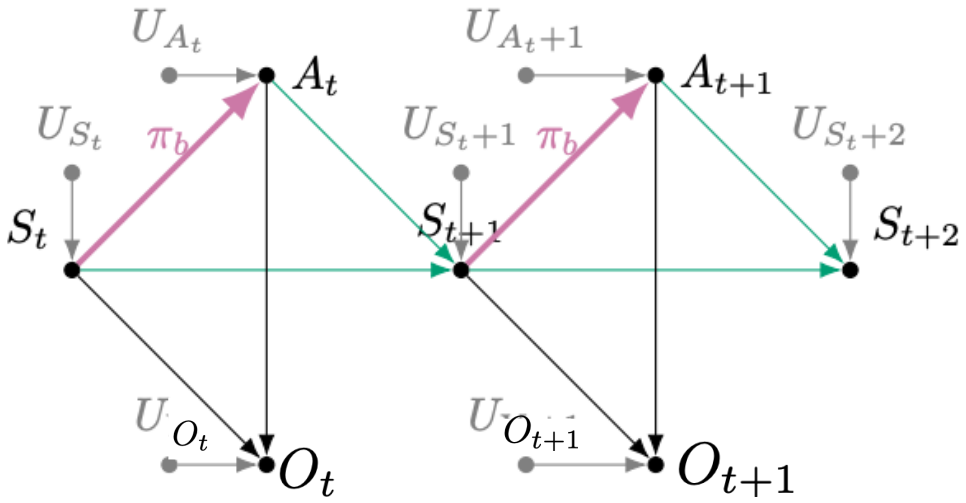



But the data-generating mechanism is important to understand when an OPE estimate makes sense and who it is applicable to

# From Evaluation to Causally Aware Evaluation

The data-generating mechanism is important to understand when an OPE estimate makes sense and who it is applicable to

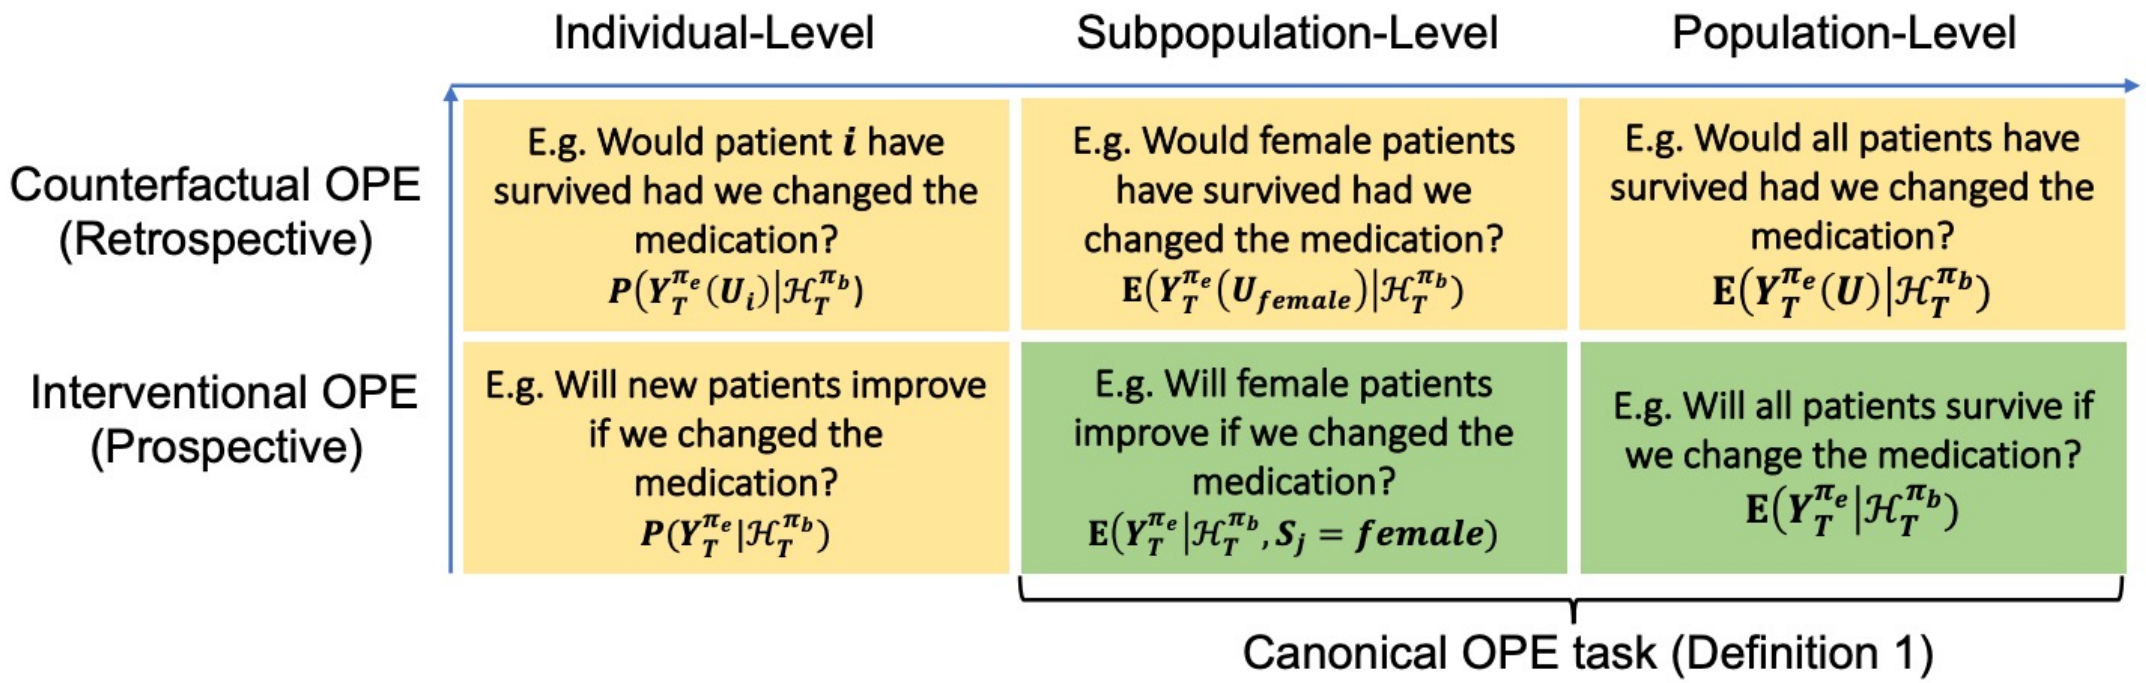



Parbhoo et al, Generalizing OPE from a Causal Perspective, arXiv pre-print, 2022.

# From Evaluation to Causally Aware Evaluation

Specifically, retrospective and prospective queries **require different estimands**:




Parbhoo et al, Generalizing OPE from a Causal Perspective, arXiv pre-print, 2022.

# Summary

- Small models are helpful in dealing with small data **and isolating relevant information** for downstream decision-making
- Evaluation should **tailored to the application and the nature of the query for careful validation.**
- Factored action spaces are recommended for use in OPE when:
  - Data is limited
  - When the problem is clearly decomposable into independent sub-problems
  - Isolating certain action features can greatly improve the variance and MSE
- In general, however **causality is crucial** for ensuring we adjust for biases in the correct way
- Existing causal representation learning methods **fail to do so** appropriately.
- Please feel free to reach out at: s.parbhoo@imperial.ac.uk