CausalSim: A Causal Framework for Unbiased Trace-Driven Simulation

Abdullah Alomar*, Pouya Hamadanian*, Arash Nasr-Esfahany*, Anish Agarwal, Mohammad Alizadeh, Devavrat Shah

Trace-Driven Simulation

- Simulates only one (or a few) components of the system where we want to experiment with a new *intervention.*
- (*Exogenous Trace Assumption*): A trace capturing other components' behavior is *replayed* to account for their effect on simulation.

✓ Faster and simpler than simulating all system components

✗ Hard to guarantee the exogenous trace assumption in traces collected from real-world systems, which could lead to biased simulation outcomes.

Example: Heterogeneous Server Load Balancing

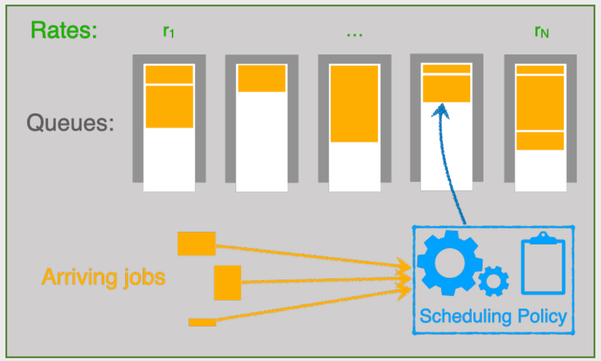

Job #	1	2	Source Algorithm: Shortest Queue
Server #	2	5	...
Runtime	120s	237s	
Job #	1	2	Target Algorithm: Power of 2
Server #	?	?	...
Runtime	?	?	Simulated

Unbiased Trace-Driven Simulation

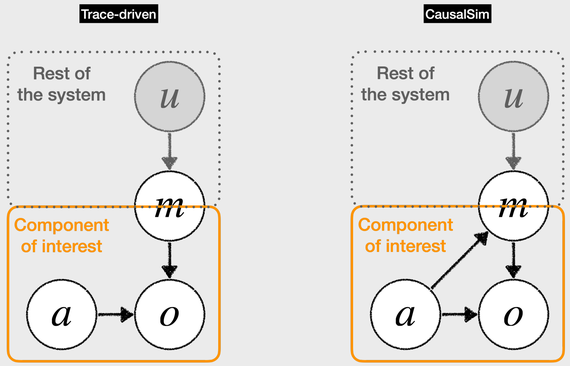

Estimating Counterfactual Traces with Matrix Completion

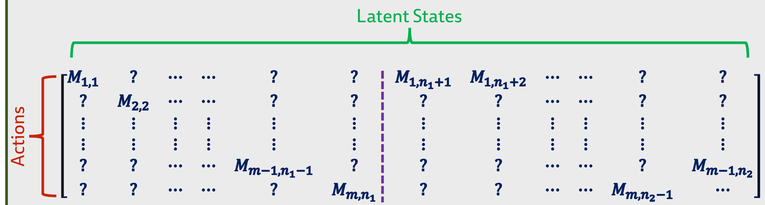

Generic Matrix Completion Assumptions:

✓ Low-Rank Structure

✗ Random Missingness Pattern

✗ Enough Number of Revealed Entries

RCT Property: Distribution of latent factors is the same for traces assigned to different algorithms.

$$\begin{bmatrix} a_1\beta_1 & \cdots & \cdots & a_1\beta_{n_1} & ? & ? & ? & ? \\ ? & ? & ? & ? & a_2\beta_{n_1+1} & \cdots & \cdots & a_2\beta_{n_2} \end{bmatrix}$$

Distribution of $\beta_1, \ldots, \beta_{n_1}$ ≈ Distribution of $\beta_{n_1+1}, \ldots, \beta_{n_2}$

$$\frac{\alpha_1}{\alpha_2} \approx \frac{\text{Average of observations in row 1}}{\text{Average of observations in row 2}}$$

Learning Architecture

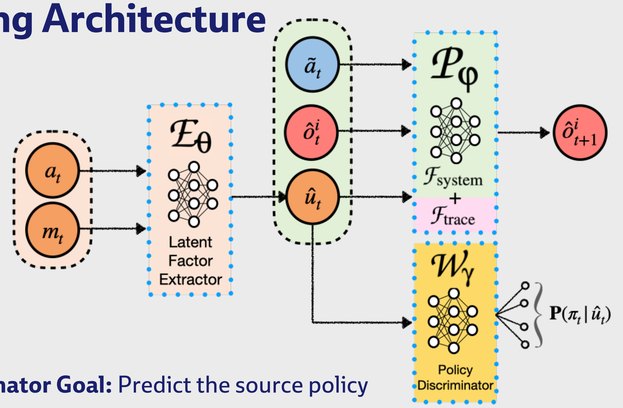

- **Discriminator Goal:** Predict the source policy

Evaluation: Video Streaming

Case Study: CausalSim in the Wild

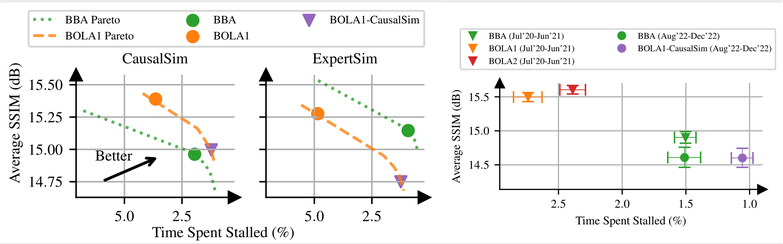

Simulation Accuracy