From "What" to "Why" Towards Causal Deep Learning

Nan Rosemary Ke[1,3*], [1]Mila [2]University of Montreal [3]Polytechnique Montreal

POLYTECHNIQUE MONTRÉAL UNIVERSITÉ D'INGÉNIERIE

Mila

Université de Montréal

Motivation

- ▷ **Systematic generalization**: major challenge for deep neural networks:
  - **Robust** to irrelevant changes.
  - **Fast adaptation** to meaningful changes.
- ▷ **Causal models** can help with systematic generalization.
  - Model distribution shifts through the notion of interventions (**robustness**).
  - Reusing modules within with little or no further training (**Fast Adaptation**).

List of works

1. Causal structure search as continuous optimization.
2. Amortized learning of neural causal representations.
3. Essential ingredients for causal induction in model-based RL.
4. Learning causally correct partial models for model-based RL.

Neural causal representations

Causal structure representations: `arxiv:2008.09301`
Most classic causal literature

- Learns a discrete graph structure

We propose to use neural networks:

- Learn a **fully continuous structure representation**.
- Propose an **evaluation method** for the continuous structure.

Model

- **Attentive relational model**: captures causal relationships.
- **Sequence model accumulates information.**
- **Graph decoder**: outputs the estimated adjacency matrix of the learned model, evaluating the continuous representation.

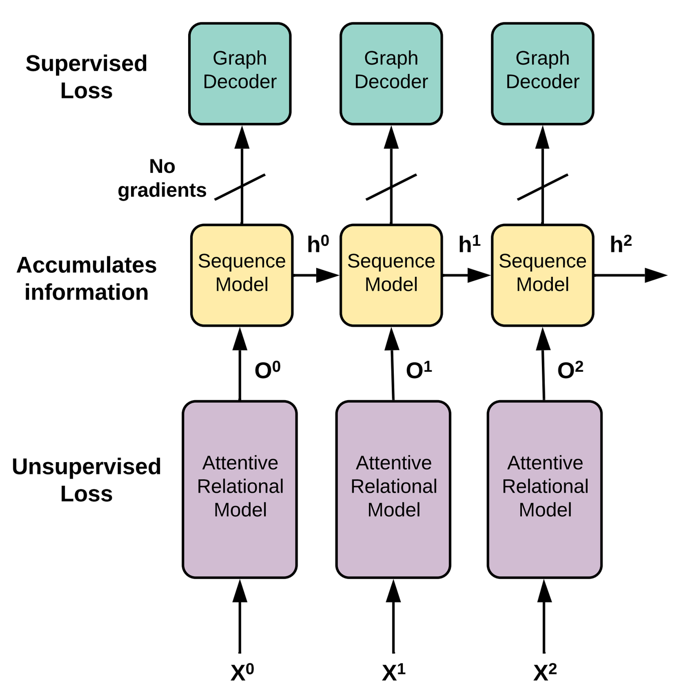

Causal structure search as continuous optimization

Causal structure discovery: (`arxiv:1910.01075`)
Most classic causal discovery methods:

- Treats it as a discrete search problem ==> Super exponential

We propose to use neural networks:

- Learn adjacency matrix (as parameters) of the graph.
- Treats causal discovery as a continuous optimization problem.

Our model

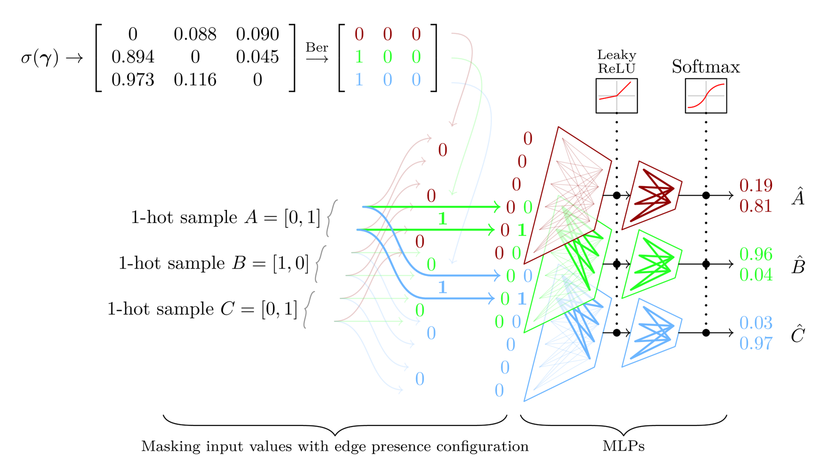

(a) Model architecture for $M = 3, N = 2$

- Use N neural networks to represent causal graph with N variables
- Each neural network models:
  Who are the direct causal parents (**Structural parameters**)
  What is the relationship between them (**Functional parameter**)

Method Overview

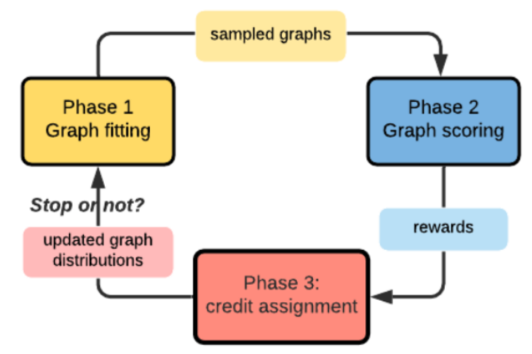

(a) Method overview

- Phase 1: Graph fitting on observational data
- Phase 2: Graph scoring on interventional data (predict interventions)
- Phase 3: Credit assignment to structural parameters

Comparison to baselines

Method	Asia	Sachs	collider	chain	jungle	collider	full
M	8	11	8	13	13	13	13
Zheng et al. [10]	14	22	18	39	22	24	71
Yu et al. [11]	10	19	7	14	16	12	77
Heinze-Deml et al. [48]	8	17	7	12	12	7	28
Peters et al. [33]	5	17	2	2	8	2	16
Eaton and Murphy [49]	0	OOM	7	OOM	OOM	OOM	OOM
Proposed Method (SDI)	0	6	0	0	0	0	7

(a) Comparisons to baselines, measured in Structured Hamming Distance (smaller is better)

Causal induction in model-based RL

Model-based reinforcement learning as causal learning
World model describes outcomes of actions (interventions) and hence is a causal model.
Classic RL environments (framework):

- Unclear what the underlying causal structure is.
- Often difficult to control different aspects of the environment.
- Difficult to evaluate if model has learned the causal structure.

We propose:

- A suite of interactive environments and causal experiments to evaluate the learned causal knowledge.
- We identify two forms of inductive bias related to causality—factorizable representations and sparse predictive associations.

Physics Environment
Environment models simple physics:

- Heavier objects can push lighter objects, transitions are deterministic.
- Models an acyclic-tournament graph.

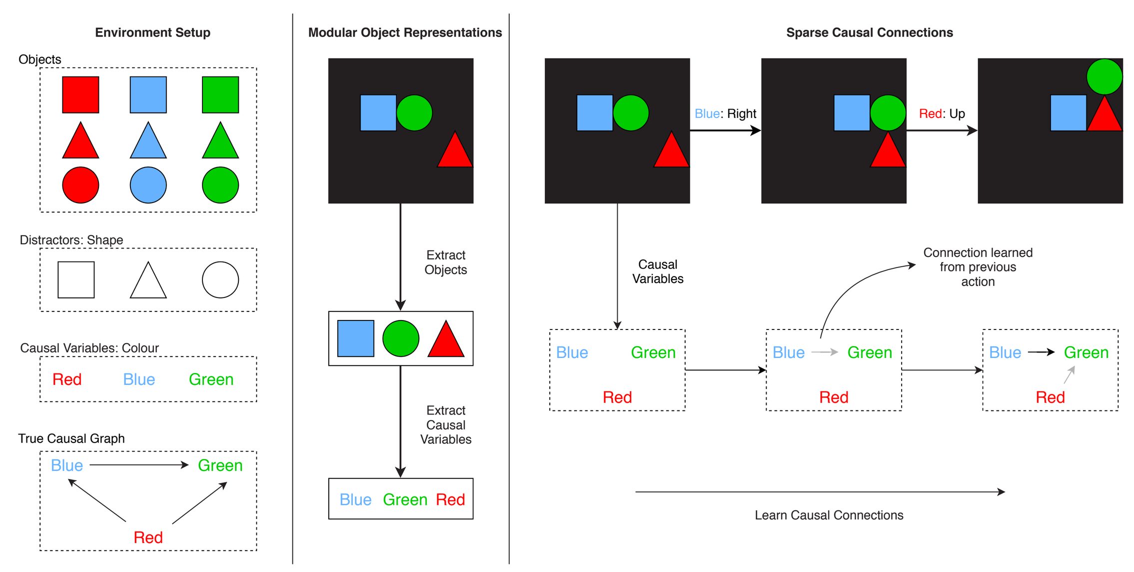

(a) Physics environment

Chemistry environment

- Can model arbitrary causal graph.
- Transitions can be determinstic or stochastic.

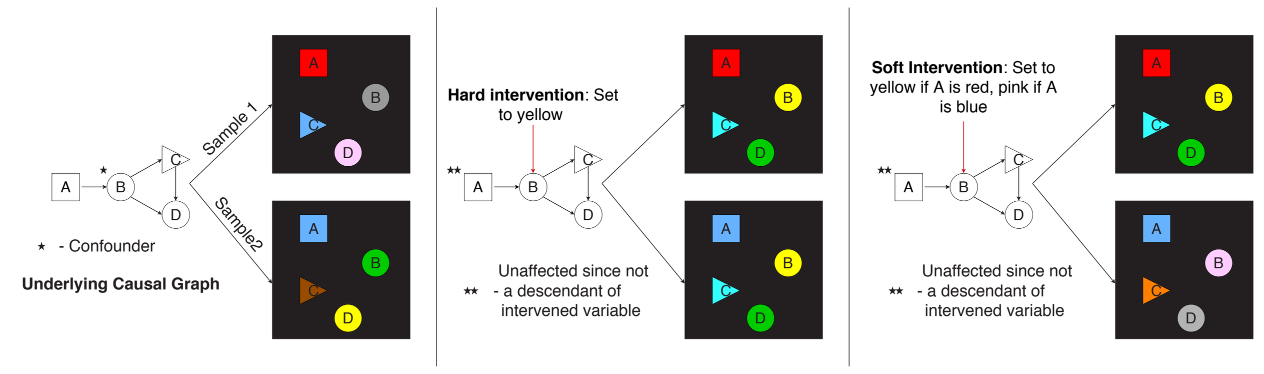

(a) Chemistry environment

Evaluation

- Likelihood and ranking loss.
- Reinforcement learning tasks.

Findings:

- Modular representation helps to learn more complicated graphs.
- Sparse predictive associations helps to learn causal structures.

Causally correct partial models

`arxiv:2002.02836`

Model-based RL

- Models the intervention outcomes in the world (environment).
- Often are partial models.

Partial Models

- Do not model the world at every step, models $P(s_t + k | s_t, a_t, a_t + 1, a_t + 2, \ldots)$.
- Partial models are causality incorrect, environment is an unobserved confounder.
- They model the observational conditional $p(r_2|s_0, a0, a_1)$.
- They should be modeling the interventional conditional

$$p(r_2|s_0, do(a0), do(a_1)) = \sum_{s_1} p(r_2|s_1, a1)p(r_1|s_0, a0)$$

.

- Partial models are biased, they do not know if reward is caused by environment or the action.

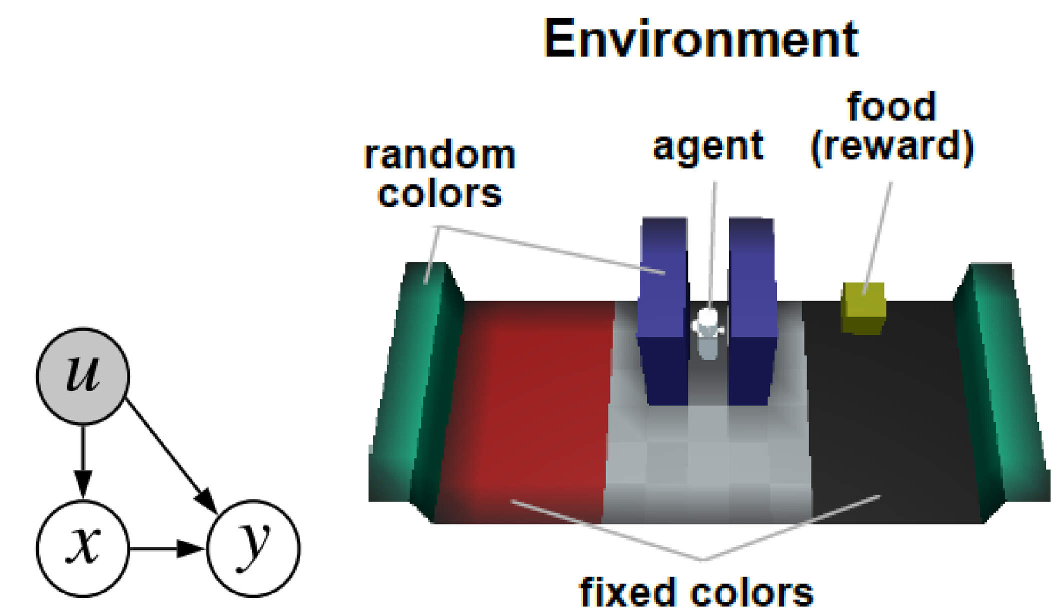

Figure 5: Left: The environment is an unobserved confounder for partial world models. Right: Partially observed environment, agent gets rewarded for getting the food.

Backdoor adjustment
Variable Z blocks all paths from U (env) to X (action)

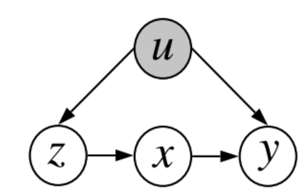

Figure 6: Modeling and conditioning on backdoor variable Z adjust for the unobserved confounder bias.

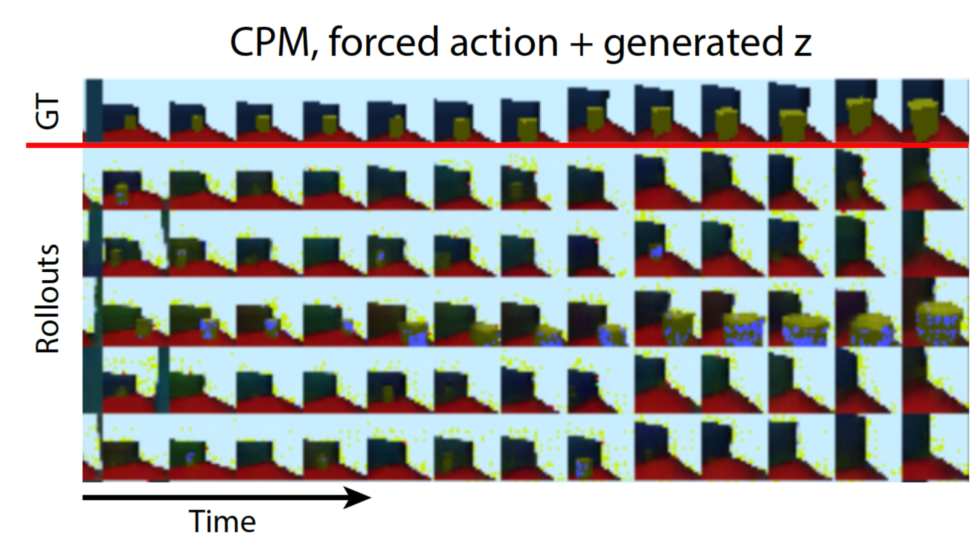

Figure 7: Partial models are corrected, different rollouts (actions) give different (reward) outcomes.