



**FACULTY
OF INFORMATION
TECHNOLOGY
CTU IN PRAGUE**

**DATA
SCIENCE
LABORATORY**

Introduction to Causality

AML Spring 2023



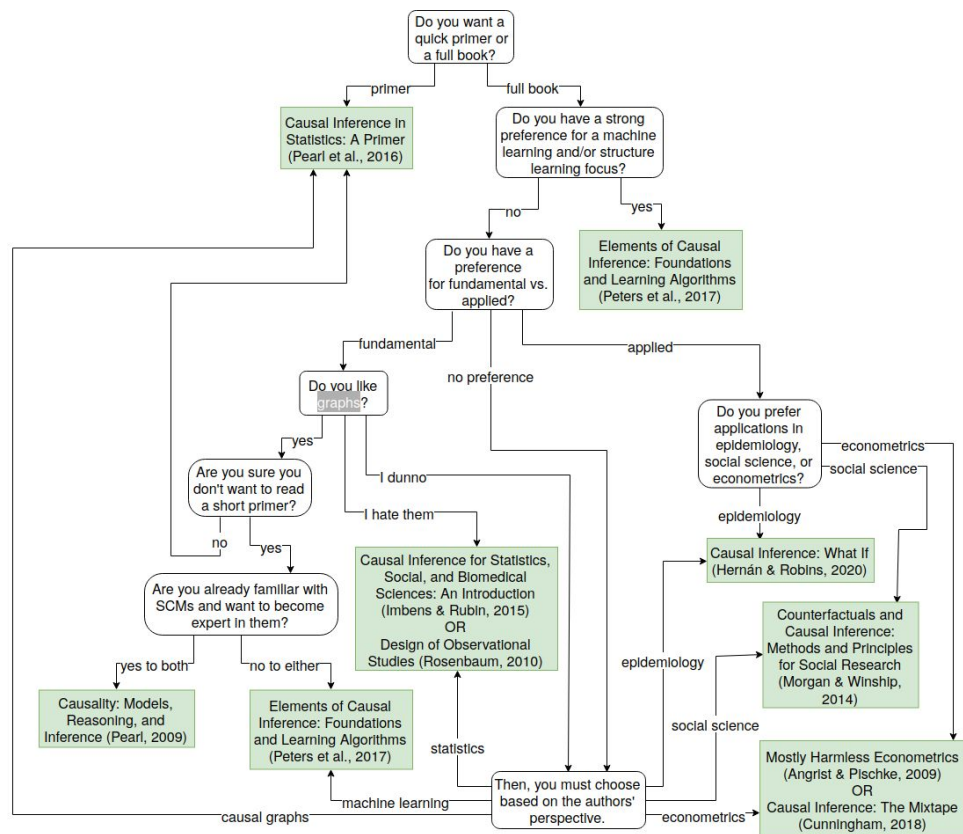
Causal Inference

- Inferring the effects of any treatment/policy/intervention/etc.
- Examples
 - Effect of treatment on a disease
 - Effect of climate change policy on emissions
 - Effect of social media on mental health
 - Many more (effect of X on Y)

Causality



- Concept that could be approached from various standpoints
- Used in field like
 - Econometrics
 - Social science
 - Epidemiology
 - Statistics
 - Machine Learning
 - (Multi-agent) Reinforcement Learning





Towards Causal Representation Learning

- Bengio et al., 2021
- Causal inference can help address important challenges in machine learning such as generalization, transfer learning, and data efficiency.
- Causal representation learning is a crucial problem for artificial intelligence and could unlock new capabilities in learning from data.
- Incorporating causality into machine learning models requires careful consideration of assumptions, limitations, and trade-offs.
- Combining causal inference techniques with machine learning models to improve generalization and transfer learning.
- Developing algorithms for causal representation learning that can handle complex data types such as images, audio, and video.
- Incorporating causality into reinforcement learning algorithms to improve the performance of agents in complex environments.



Simpson's Paradox

- Hypothetical disease with two possible treatments
- Table showing mortality rate

		Condition		
		Mild	Severe	Total
Treatment	A	15% (210/1400)	30% (30/100)	16% (240/1500)
	B	10% (5/50)	20% (100/500)	19% (105/550)

- Apparent paradox:
 - If condition is not know, treatment A is better
 - If condition is known, treatment B is better

Simpson's Paradox

Which treatment is better depends on the **causal** structure of the data




Figure 1.1: Causal structure of scenario 1, where condition C is a common cause of treatment T and mortality Y . Given this causal structure, treatment B is preferable.





Figure 1.2: Causal structure of scenario 2, where treatment T is a cause of condition C . Given this causal structure, treatment A is preferable.

Correlation is not causation



Association is Not Causation

- Correlation is meant statistical dependence
- Technically, it is measure of linear dependence, better term should be association
- Total association is not all or none, could be combination of
 - Spurious (correlation)
 - Confounding (hidden common cause)
 - Causal association




Figure 1.4: Causal structure, where drinking the night before is a common cause of sleeping with shoes on and of waking up with a headache.



Counterfactuals

- Alternatives scenarios that did not actually happened but could happen under different circumstances
- Humans as Counterfactual Reasoning Machines
 - Constantly evaluating alternative scenarios
 - Imagining outcomes of different actions
- Counterfactuals in Everyday Life
 - Informed decision-making based on "what-if" analysis
 - Learning from past experiences and mistakes
- Regret Minimization in Human Behavior
 - Comparing outcomes of taken and untaken actions
 - Guiding future decisions to minimize regret



Potential Outcomes Framework

- person has a headache and decides
 - Take a pill (treatment)
 - Not take a pill (control)
- Potential outcome: will headache persit
 - $Y_i(1)$ severity of headache hour after taking the pill
 - $Y_i(0)$ severity of headache hour after (not taking a pill)
- Individual treatment effect
 - $\tau_i = Y_i(1) - Y_i(0)$

$$\tau_i = Y_i(1) - Y_i(0)$$

- We can observe only one of outcomes $Y_i(1)$ and $Y_i(0)$
- How to compute the treatment effect?
 - Fundamental problem of causal inference



Average treatment effect

- To estimate the average causal effect of the pill, we can use a sample of individuals who took the pill and another sample of individuals who did not.
- Average treatment effect

$$\tau = \frac{1}{N_1} \sum_{i=1}^{N_1} Y_i(1) - \frac{1}{N_0} \sum_{i=1}^{N_0} Y_i(0)$$

- How to average question marks?

i	T	Y	$Y(1)$	$Y(0)$	$Y(1) - Y(0)$
1	0	0	?	0	?
2	1	1	1	?	?
3	1	0	0	?	?
4	0	0	?	0	?
5	0	1	?	1	?
6	1	1	1	?	?

Ignorability

- What makes it valid to calculate the ATE by taking the average of the $Y(0)$ column, ignoring the question marks, and subtracting that from the average of the $Y(1)$ column, ignoring the question marks?
- Ignoring the question is called ignorability
 - ignoring how people ended up selecting the treatment they selected
 - and just assuming they were randomly assigned their treatment;




Figure 2.2: Causal structure when the treatment assignment mechanism is ignorable. Notably, this means there's no arrow from X to T , which means there is no confounding.



Controlling for Confounding

- **Confounding Factor**
 - Variables that affect both the treatment and the outcome
 - Can lead to biased estimates of causal effects
- **Importance of Controlling for Confounding**
 - Obtain unbiased and accurate estimates of causal effects
 - Improve decision-making based on observational data
- **Methods to Control for Confounding**
 - Matching
 - Stratification



Control for confounding

- Matching
 - Attempt to create comparable groups of individuals who took the pill and those who did not
 - Match based on confounding variables (e.g., age, gender, baseline health)
 - nearest neighbor matching
 - directly pairs treated and control individuals based on their similarity in confounding variables,
- Stratification
 - divide the population into strata based on the confounding variables
 - Initial headache severity
 - estimate the causal effect within each stratum
 - combine these estimates to calculate the overall average causal effect
 - weighting the estimates by the proportion of individuals in each stratum.
 - divides the population into groups based on the values of confounding variables and estimates the causal effect within each group.



Social Influence in MARL


- Jacques et al.: Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning
- mechanism for achieving coordination and communication
- rewarding agents for having causal influence over other agents' actions
- causal influence is assessed using **counterfactual reasoning**
- At each timestep, an agent simulates alternate actions that it could have taken, and computes their effect on the behavior of other agents.
- Actions that lead to bigger changes in other agents' behavior are considered influential and are rewarded
- influence leads to enhanced coordination and communication in challenging social dilemma environments, dramatically increasing the learning curves of the deep RL agent



Sequential Social Dilemmas


- Can be thought of as analogous to spatially and temporally extended Prisoner's Dilemma-like games.
- The reward structure poses a dilemma because individual short-term optimal strategies lead to poor long-term outcomes for the group

Cleanup



- A public goods dilemma in which
- agents get a reward for consuming apples, but must use a cleaning beam to clean a river in order for apples to grow.
- While an agent is cleaning the river, other agents can exploit it by consuming the apples that appear.

Harvest



- A tragedy-of-the-commons dilemma
- apples regrow at a rate that depends on the amount of nearby apples.
- If individual agents employ an exploitative strategy by greedily consuming too many apples, the collective reward of all agents is reduced.

Multi-agent Reinforcement Learning



Definition 2.2 (Dec-POMDP) *Decentralised Partially Observable Multi-Agent Markov Decision process is It is a γ -tuple $\{S, \{A_i\}, T, R, \{\Omega_i\}, O, \gamma\}$, where S are states, $\{A_i\}$ is the joint action set, $T = P(s'|s, a)$ is the set of conditional transition probabilities between states, R is the reward function, $\{\Omega_i\}$ is the joint observation set, $O(s', a, o) = P(o|s', a)$ gives the conditional observation distribution, and $\gamma \in [0, 1]$ is the discount factor.*



Intrinsic motivation

- To stimulate agents to learn cooperative behavior introduce
- category of reward functions that allow learning of desired behaviors in a wide range of environments and tasks, sometimes even in the absence of environmental rewards.
- Social influence intrinsic motivation gives an agent k additional reward when it has causal influence on the actions of other agents.
- It adds a causal influence reward c_k^t to the agent's immediate environmental (extrinsic) reward e_t^k at time t :

$$r_t^k = \alpha e_t^k + \beta c_t^k.$$

Evaluation of social influence



To evaluate the causal influence of agent k on agent j at time t , agent j should be able to condition its action a_t^j on a_t^k , agent's k action at time t . Therefore, a_j can quantify the probability of the next step action as

$$p(a_t^j | a_k^t, s_t^j).$$

Then we can we can replace a_t^k by \tilde{a}_t^k , the counterfactual action, and compute a new next step probability

$$p(a_t^j | \tilde{a}_k^t, s_t^j).$$



Evaluation of social influence

By averaging the policy distribution from a sampling of several counterfactual actions, we would obtain the marginal policy of agent j :

$$p(a_t^j | s_t^j) = \sum_{\tilde{a}_t^k} p(a_t^j | \tilde{a}_t^k, s_t^j) p(\tilde{a}_t^k, s_t^j),$$


i.e. agent's j policy if it did not take into account actions of agent k .

The difference between agent's j marginal policy and the conditional policy of agent j after observing agent's k action is a degree of how agent k is causally influencing agent j .

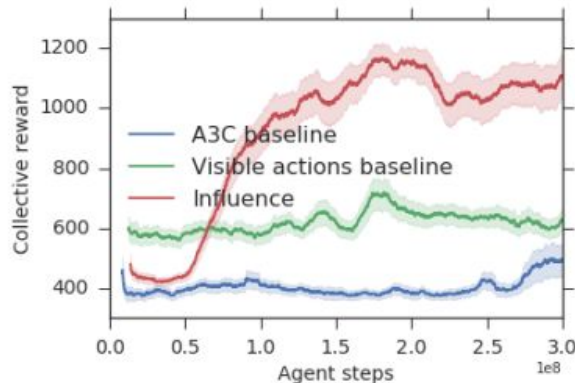
Therefore, the overall causal influence of agent k on all other agents is given by:

$$\begin{aligned} c_t^k &= \sum_{j=0, j \neq k}^N \left[D_{KL} \left[p(a_t^j | a_t^k, s_t^j) \parallel \sum_{\tilde{a}_t^k} p(a_t^j | \tilde{a}_t^k, s_t^j) p(\tilde{a}_t^k | s_t^j) \right] \right] \\ &= \sum_{j=0, j \neq k}^N [D_{KL} [p(a_t^j | a_t^k, s_t^j) \parallel p(a_t^j | s_t^j)]] , \end{aligned} \quad (4.1)$$

Effect of social influence



(a) *Cleanup*



(b) *Harvest*

Figure 1: Total collective reward obtained in Experiment 1. Agents trained with influence (red) significantly outperform the baseline and ablated agents. In Harvest, the influence reward is essential to achieve any meaningful learning.

Social influence




Figure 2: A moment of high influence when the purple influencer signals the presence of an apple (green tiles) outside the yellow influencee's field-of-view (yellow outlined box).

plementary Material).

Figure 2 shows a moment of high influence between the influencer and the yellow influencee. The influencer has chosen to move towards an apple that is outside of the ego-centric field-of-view of the yellow agent. Because the influencer only moves when apples are available, this signals to the yellow agent that an apple must be present above it which it cannot see. This changes the yellow agent's distribution over its planned action, $p(a_t^j | a_t^k, s_t^j)$, and allows the purple agent to gain influence. A similar moment occurs when the influencer signals to an agent that has been cleaning the river that no apples have appeared by staying still (see Figure 14 in the Sup-

- Agents continue to move and explore randomly while waiting for apples to spawn,
- The **influencer** only traverses the map when it is pursuing an apple, then stops. The rest of the time it stays still.
- The **influencer** agent learned to use its own actions as a binary code which signals the presence or absence of apples in the environment



Model of Other Agents

- Computing the causal influence reward requires knowing the probability of another agent's action given a counterfactual,
- Requires a centralized training approach in which agents could access other agents' policy network
- To relax this unrealistic assumption we equip each agent with its own internal Model of Other Agents (MOA).
- The MOA is trained to predict all other agents' next actions given their previous actions, and the agent's egocentric view of the state: $p(a_{t+1}|a_t, s_t)$.

Model of other agents




Figure 6: The Model of Other Agents (MOA) architecture learns both an RL policy π_e , and a supervised model that predicts the actions of other agents, a_{t+1} . The supervised model is used for internally computing the influence reward.

Causal Diagram




Figure 8: Causal diagrams of agent k 's effect on j 's action. Shaded nodes are conditioned on, and we intervene on a_t^k (blue node) by replacing it with counterfactuals. Nodes with a green background must be modeled using the MOA module. Note that there is no backdoor path between a_t^k and s_t in the MOA case, since it would require traversing a collider that is not in the conditioning set.

Next week



“The great question of the twenty-first century is going to be ‘Whose black box do you trust?’”

Wtf?

Tim O'Reilly

