Causal Inference Lecture 2

Kenneth (Kenny) Joseph

University at Buffalo
Department of Computer Science and Engineering
School of Engineering and Applied Sciences

Pass phrase: Daphne Koller

Daphne Koller

文A 27 languages

Article Talk | Read Edit View history Tools

From Wikipedia, the free encyclopedia

Daphne Koller (Hebrew: דפנה קולר; born August 27, 1968) is an Israeli-American computer scientist. She was a professor in the department of computer science at Stanford University[4] and a MacArthur Foundation fellowship recipient.[1] She is one of the founders of Coursera, an online education platform. Her general research area is artificial intelligence[5][6] and its applications in the biomedical sciences.[7] Koller was featured in a 2004 article by *MIT Technology Review* titled "10 Emerging Technologies That Will Change Your World"[8] concerning the topic of Bayesian machine learning.[9][10]

Daphne Koller

Koller in 2019

Born	August 27, 1968 (age 54) Israel[citation needed]
Education	Hebrew University of Jerusalem (BSc, MSc) Stanford University (PhD)
Known for	Machine learning Graphical models MOOCs

Education [edit]

Koller received a bachelor's degree from the Hebrew University of Jerusalem in 1985, at the age of 17, and a master's degree from the same institution in 1986, at the age of 18.[11] She completed her PhD at Stanford in 1993 under the supervision of Joseph Halpern.[2]

Career and research [edit]

After her PhD, Koller did postdoctoral research at University of California, Berkeley from 1993 to 1995 under Stuart J. Russell,[12] and joined the faculty of the Stanford University computer science department in 1995. She was named a MacArthur Fellow in 2004. She was elected a member of the National Academy of Engineering in 2011 for contributions to representation, inference, and learning in probabilistic models with applications to robotics, vision, and biology. She was also elected a fellow of the American Academy of Arts and Sciences in 2014.

In April 2008, Koller was awarded the first ever $150,000 ACM-Infosys Foundation Award in Computing Sciences.[13]

She and Andrew Ng, a fellow Stanford computer science professor in the AI lab, founded Coursera in 2012. She served as the co-CEO with Ng, and then as president of contributions to online education by being named one of *New* *Time* magazine's 100 Most Influential People in 2012, and *Fa*

She left Coursera in 2016 to become chief computing officer lead Insitro, a drug discovery startup.[15]

Koller is primarily interested in representation, inference, lear applications to computer vision and computational biology.[16] Stanford University, Koller developed PhysiScore, which uses premature babies are likely to have health issues.[17]

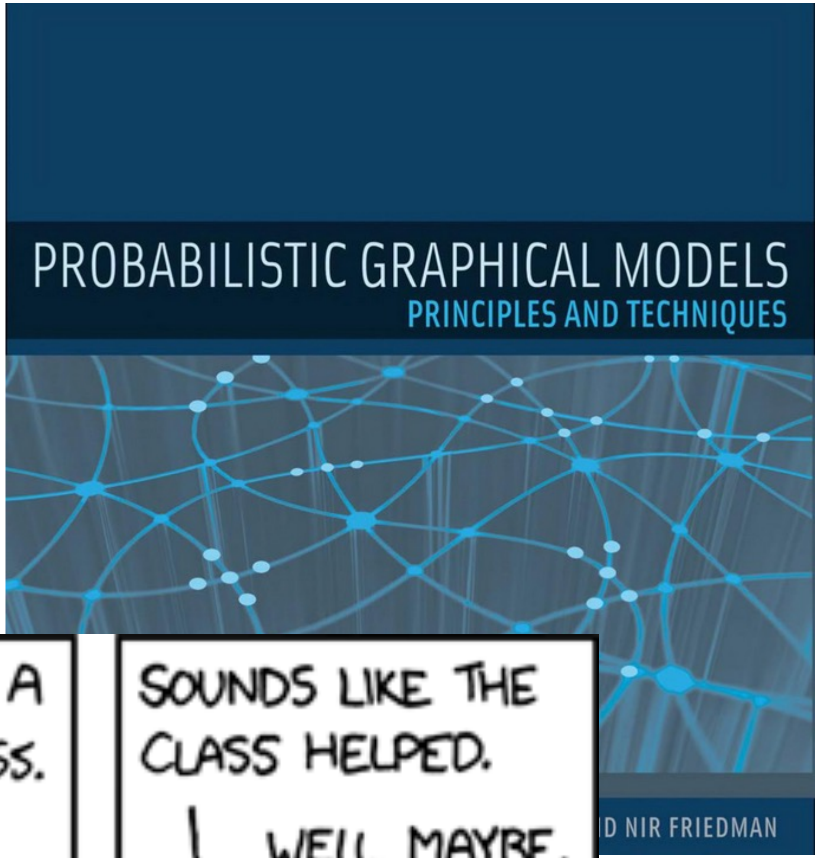

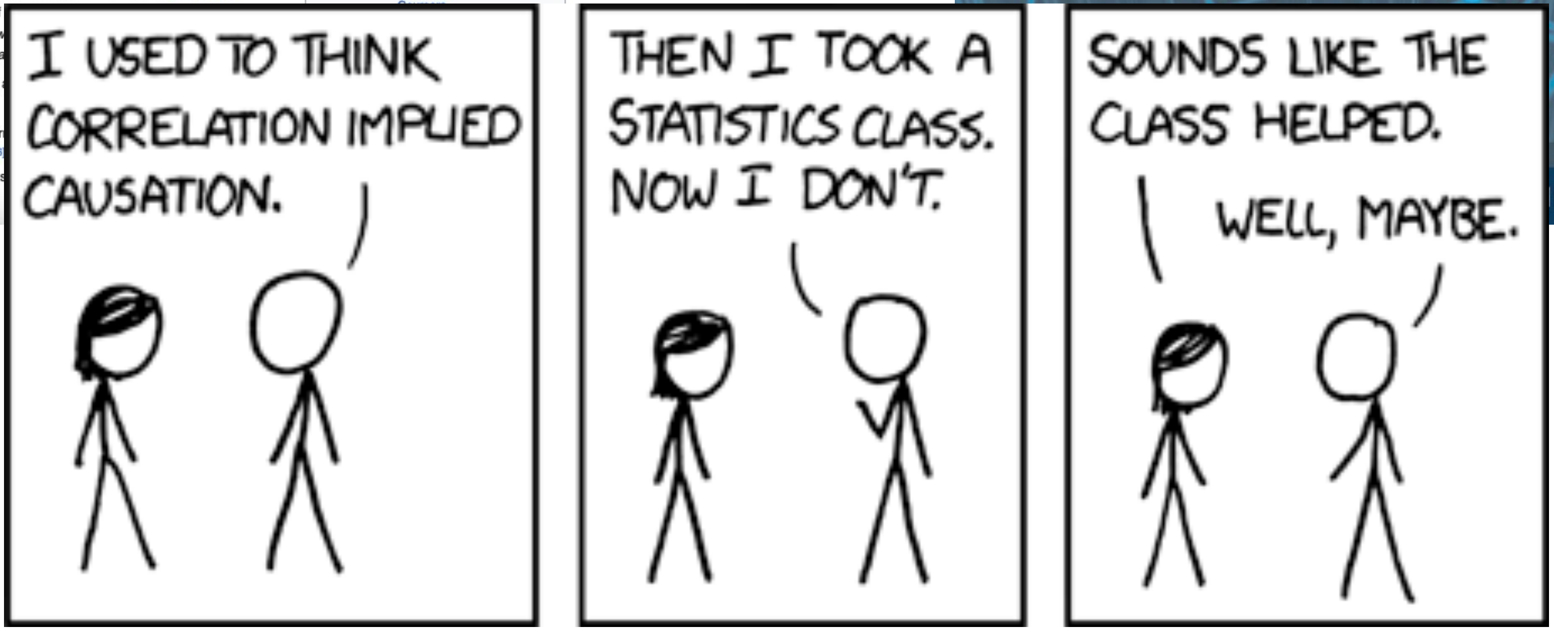

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

To help us with this, I'm going to introduce three things, and then tie them together:

1. Probability/stats
2. Causal inference
3. Directed Acyclic Graphs

First – the math

- Goal: give us the bare-bones quantitative language
  - [Also, good to remind ourselves of these concepts]
- Concepts
  - Univariate Stats
    - Probability, Random Variables (RVs), Probability of RVs
    - Expectation & Variance
    - Probability distributions
  - Multivariate Stats
    - Conditional Probability
    - Covariance and Correlation

Seeing Theory

https://seeing-theory.brown.edu/basic-probability/

Where we are at

- We now have the basic statistical tools to quantitatively describe probabilistic events in the world.
- And! We can think about the **relationship** between two random variables.

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Correlation vs. Causation

- Correlation: A **measure of the relationship between two variables**
- Causation
  - Informally: "if I change A, then B will also change"
  - Formally, two ways to think about this:
    - "Do calculus"
    - Potential outcomes
    - People argue about which of these is better. Most of that is semantics. We'll make use of both.
- Causal inference: the process of inferring the causal relationships between variables.

Correlation vs causation

https://tylervigen.com/spurious-scholar

- Can you summarize the difference between correlation and causation?
- Does correlation always mean causation?
- Does *no* correlation always mean *no* causation?

No correlation != No causation

- What does pressing down the gas pedal do to speed?
- What does going up/down a hill do to speed?
- If we didn't know that, what conclusions might we draw from a skilled driver?
- **Can you think of another example?**

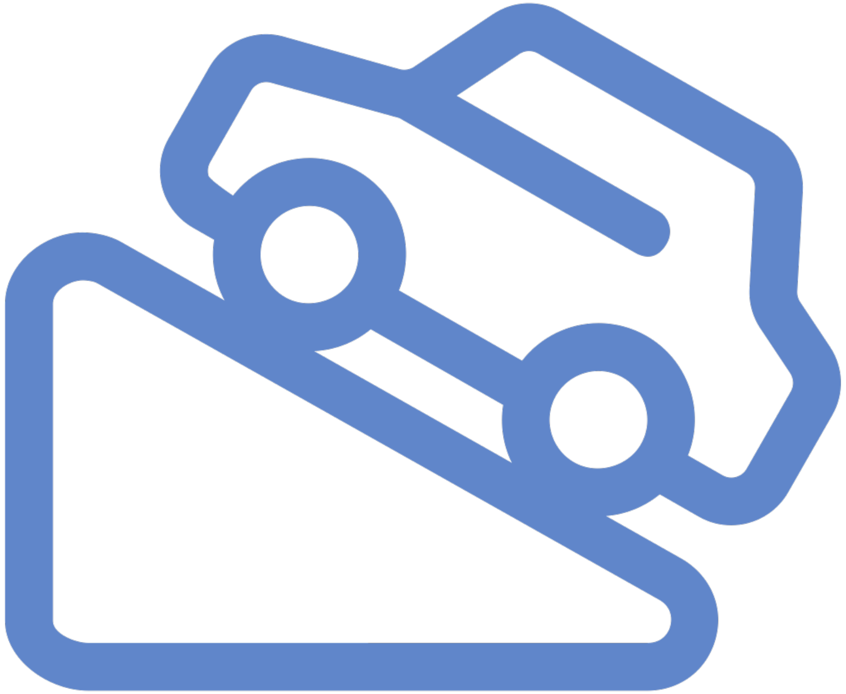

https://goodauthority.org/news/milton-friedmans-thermostat/

bq. Watch what happens on a really steep uphill bit of road. Watch what happens when the driver puts the pedal to the metal, and holds it there. Does the car slow down? If so, ironically, that confirms the theory that pressing down on the gas pedal causes the car to speed up! Because it means the driver knows he needs to press it down further to prevent the speed dropping, but can't. It's the exception that proves the rule.

Here, for example, a naive observer might take a particular campaign action, which is associated with the candidate's defeat, as evidence of incompetence by the campaign. It's not – it may be correlated only because it is the best thing that the campaign can do under particularly difficult external circumstances.

In general, correlations are fragile beings

https://www.youtube.com/watch?v=DbJyPELmhJc&ab_channel=AutodeskResearch

Introducing causal inference

What if we had a method that could **explicitly account for these challenges and help us think through *real causation*?**

Enter causal inference!

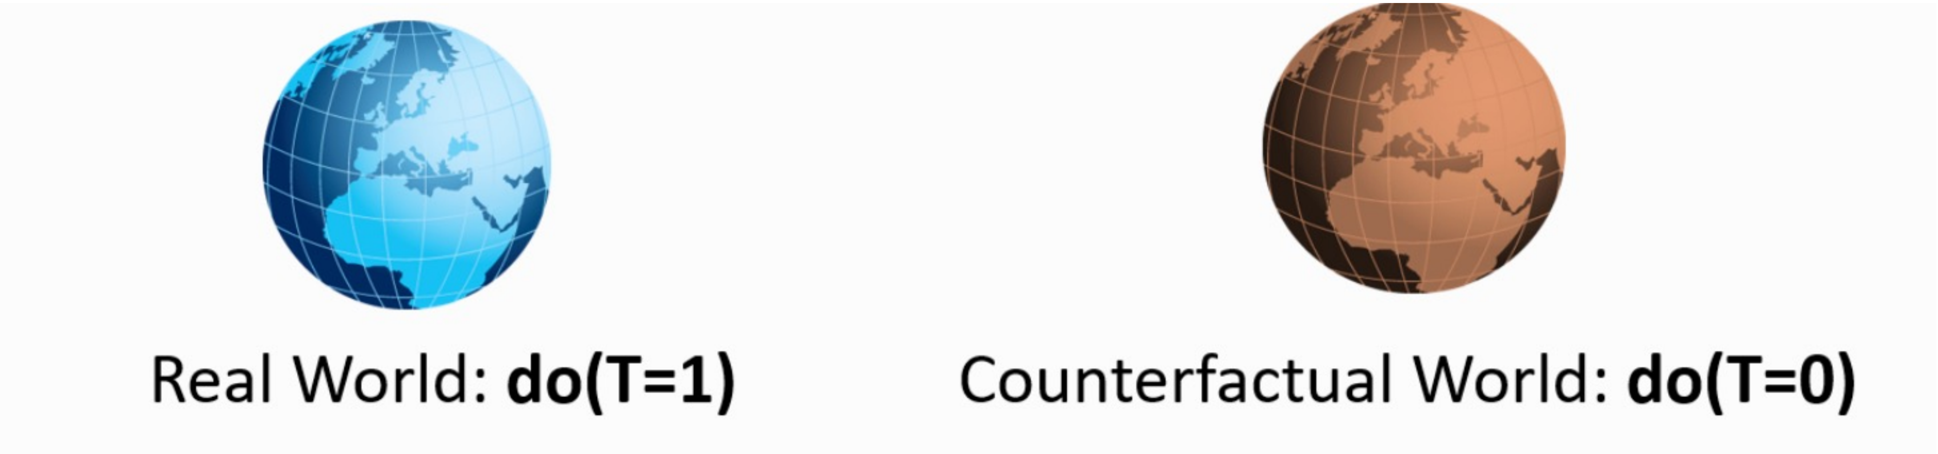

https://microsoft.github.io/dowhy/example_notebooks/tutorial-causalinference-machinelearning-using-dowhy-econml.html

Sounds great! But there's a catch...

I want to know if telling Steven he is going to get an A no matter what will make his participation better or worse

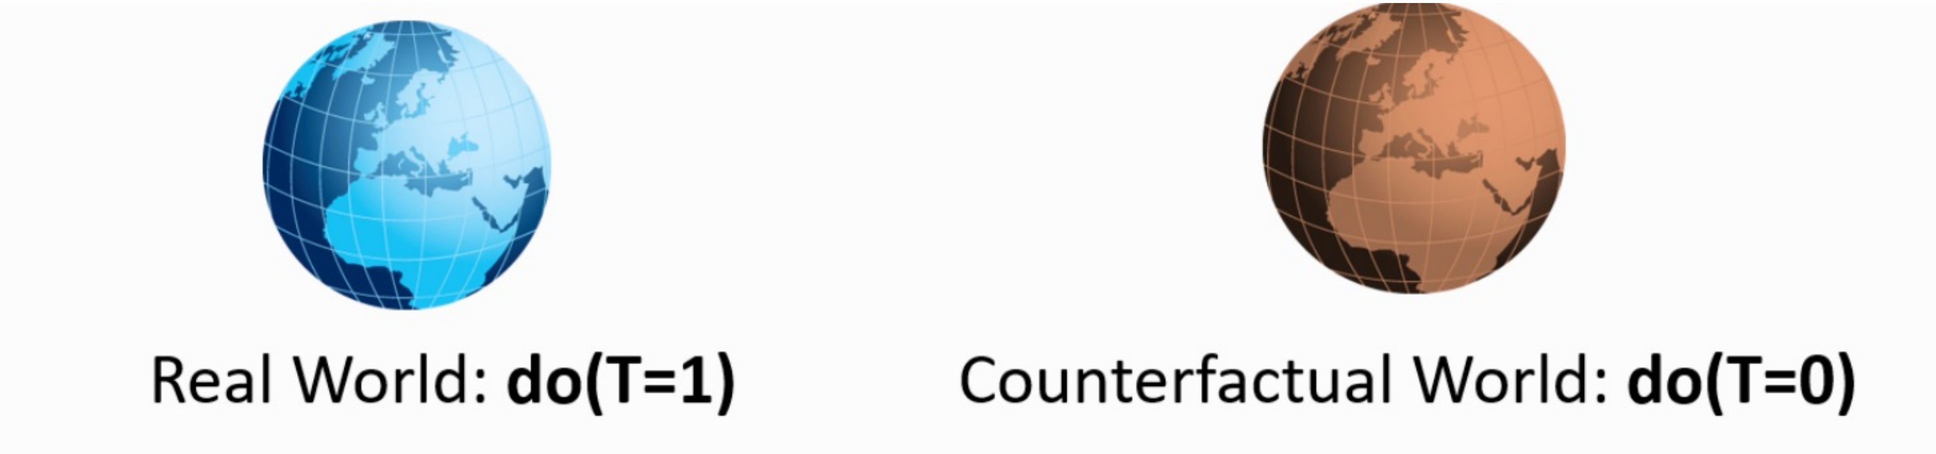

I can't both give Steven an A AND not give it to him!
This is the **Fundamental Problem of Causal Inference**

https://microsoft.github.io/dowhy/example_notebooks/tutorial-causalinference-machinelearning-using-dowhy-econml.html

Where are we at?

- Returning to our questions from last lecture
  - What is the goal of machine learning?
    - [Not necessarily one answer here]
  - What is causation?
  - Why do we care about causation?
  - When might we *not* care about causation?
  - How do we know something causes something else, and how do we prove it?

Causal inference vs. ML

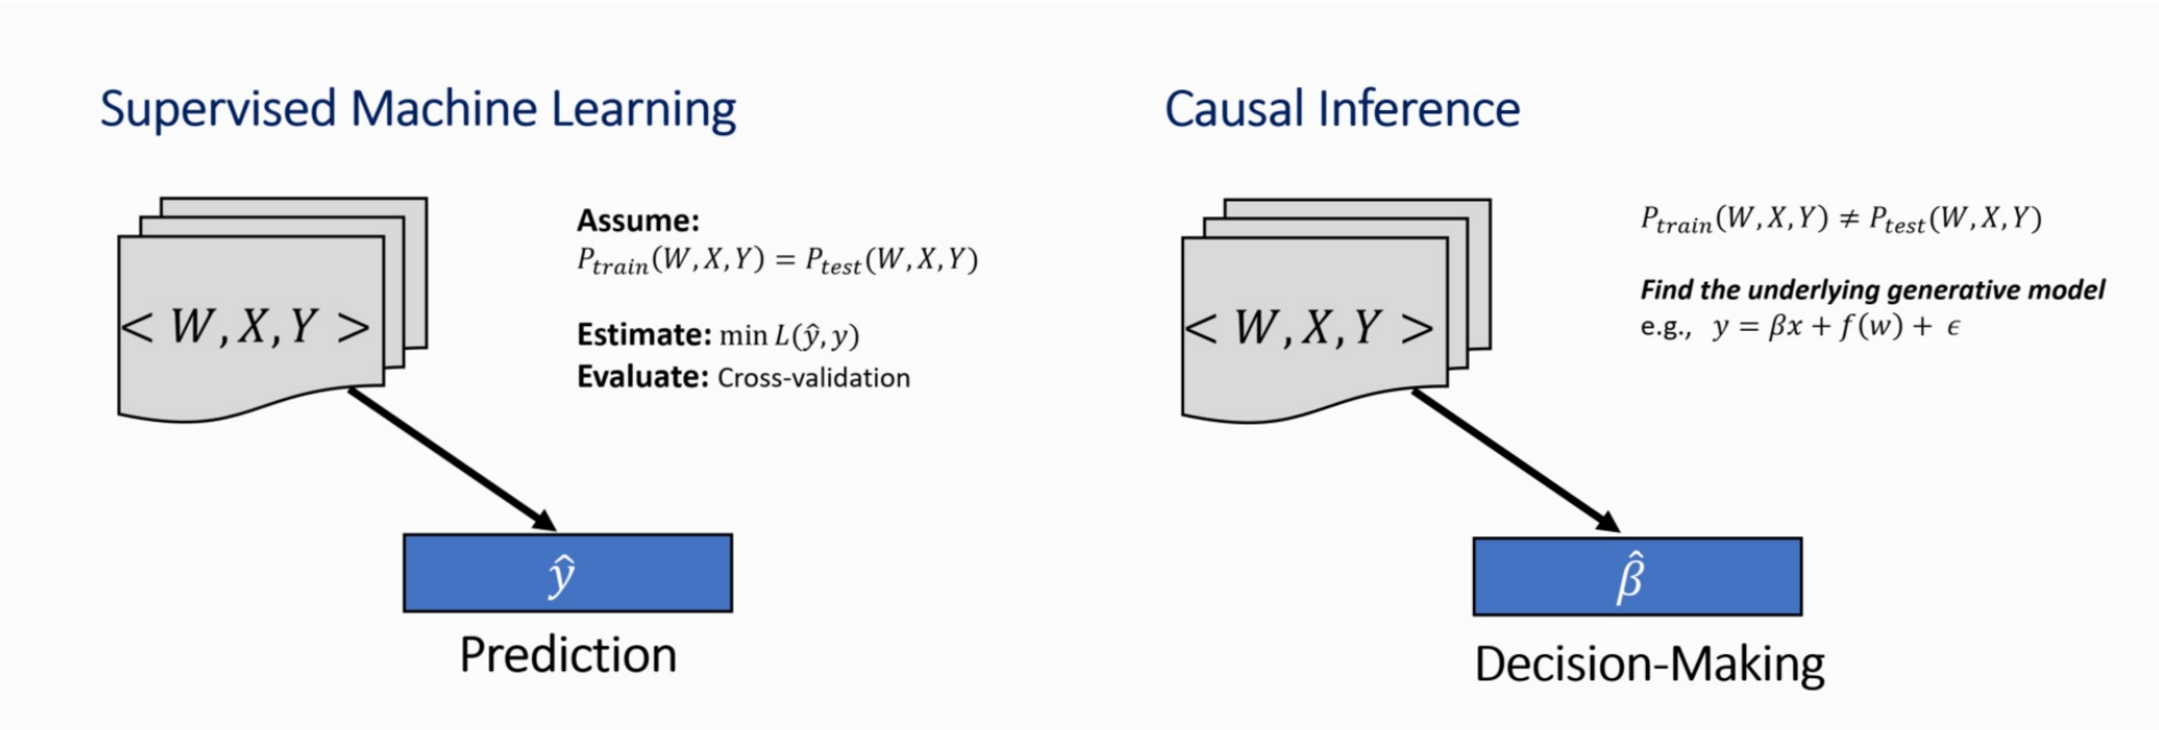

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Why causal inference is hard

I want to know if telling Steven he is going to get an A no matter what will make his participation better or worse

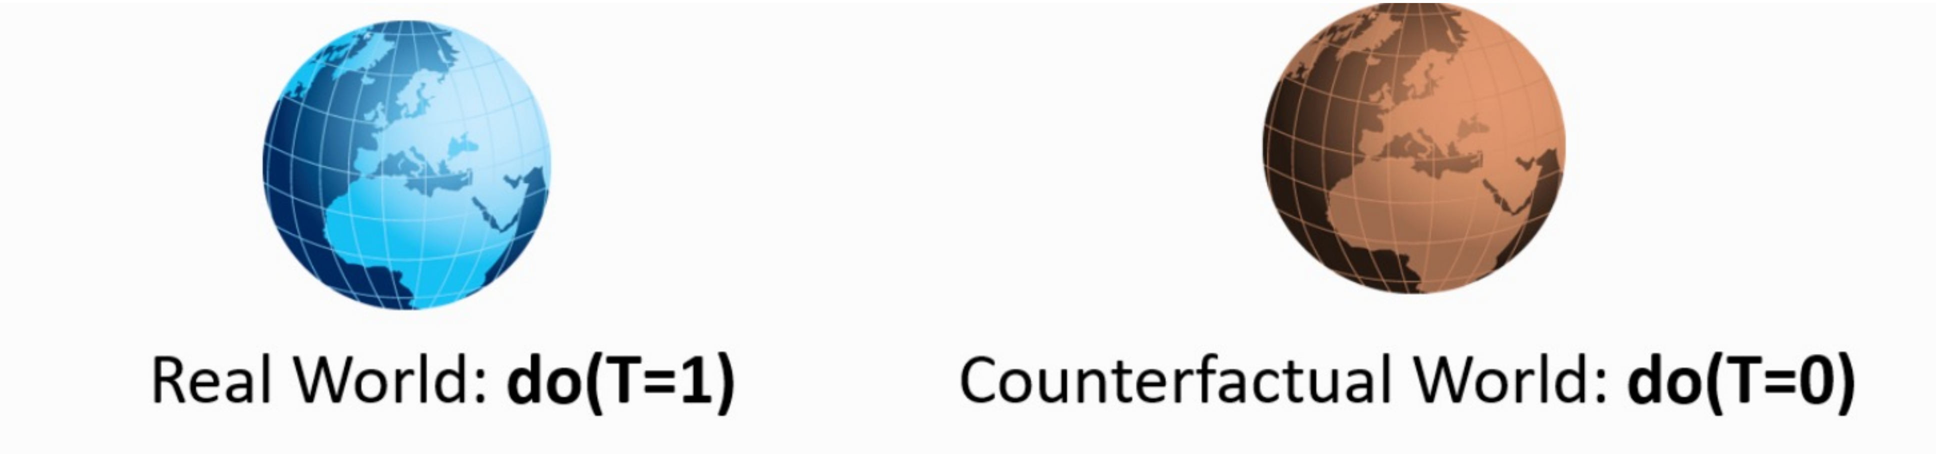

I can't both give Steven an A AND not give it to him!
This is the **Fundamental Problem of Causal Inference**

https://microsoft.github.io/dowhy/example_notebooks/tutorial-causalinference-machinelearning-using-dowhy-econml.html

What can we do, then?

1. Target averages/expectations instead of an individualized effect

$$E[Y|do(A=1)] - E[Y|do(A=0)]$$

2. Experiment!

Note: do operator indicates an intervention, in an RCT we intervene by randomly assigning treatment and control to comparable groups.

Overview of next 2-ish lectures

Three things you need for your project/to do good [computational social] science (IMO)

1. The ability to **understand/reason about** the social world
2. The ability to use that understanding to **design interventions** on the world
3. The ability to **conduct measurements** that assess the [potential] **effect** of that intervention

Problems with the simple story

- Sometimes, experimentation is unethical
  - I suspect half of you would be very angry if I gave you a placebo study guide ☺
- Other times, we might have wanted to experiment but simply couldn't, and are left with a bunch of observational data
- Since we do not always have access to experimental data, we rely on observational data for estimating causal effect. Wherein, $E[Y|do(A=1)] \neq E[Y|A=1]$; as treatment assignment and outcome might be effected by confounding elements, mediator, etc.

Uh-oh

- Now:
  - We have an idea how to evaluate an intervention, but when we can’t experiment, we have to **control for factors associated with both the intervention and the outcome**
  - Related: how do we decide on a treatment in the first place?
- Enter DAGs [and then, causal graphical models]. Informally, there are two “kinds” of probabilistic DAGs
  - Bayes Nets encode the factorizations of any joint probability distribution
  - Causal Graphical Models put explicit assumptions about causation into DAGs.
  - We’ll focus on the latter

Causal DAGs– a tool to explain causal reasoning

https://causalens.com/resources/white-papers/why-correlation-based-machine-learning-leads-to-bad-predictions/

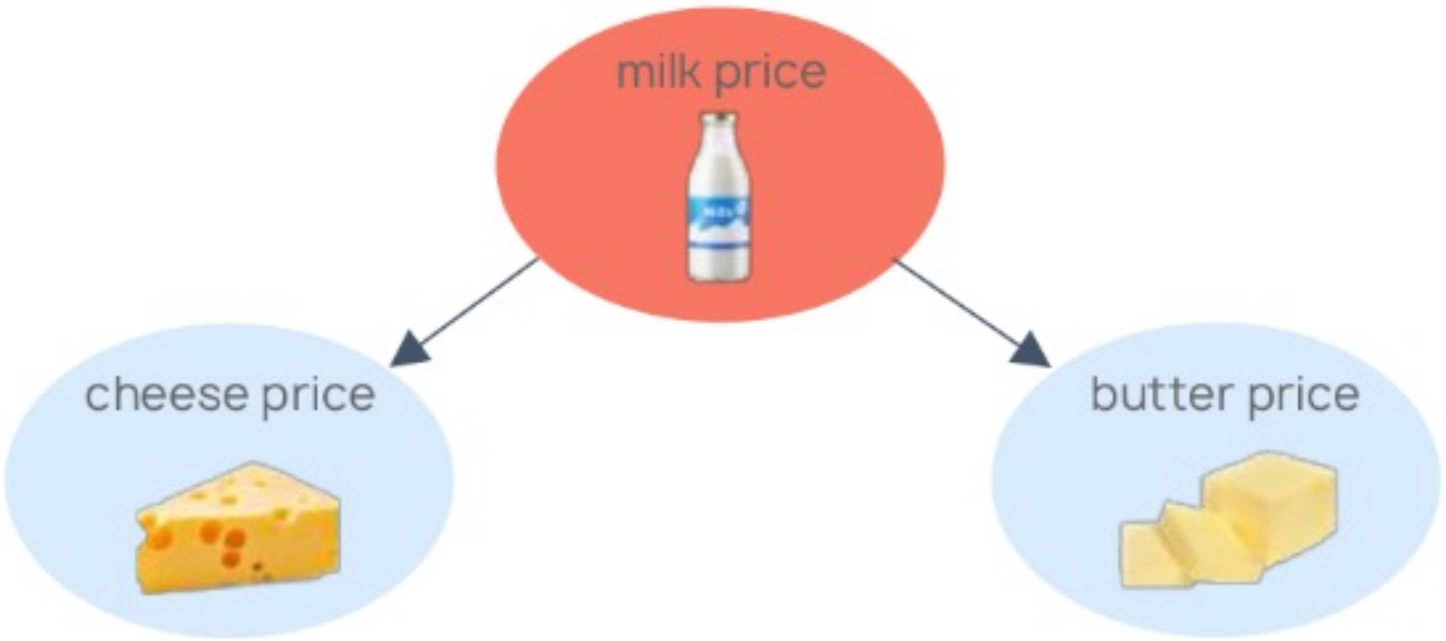

What is a causal graphical model?

- A **diagram** that helps us **explain** our **assumed causal** relationships between two things
  - Nodes are **random variables**
  - Edges are **causal relationships between RVs**
- Let's practice!
- Draw a DAG that represents the algorithmic bias scenario we saw at the end of last class.