What is causal modeling, and why is causality necessary in social computing?

Shyama Sastha Krishnamoorthy Srinivasan[1]

[1]Department of Computer Science
Indraprastha Institute of Information Technology

Talk's Objectives and Takeaways

What is causal inference? Why should we care?

- Most machine learning algorithms depend on correlations.
- Correlations alone are a dangerous path to actionable insights.

Learn how to formulate and estimate causal effects.

- To evaluate the impact of social AI systems.
- To make underlying algorithms more robust to changes in data.

Apply causal inference methods to a practical problem

- Estimating the causal impact of a recommendation system.

**Most of the content you see here today is from Amit Sharma's causal inference tutorial.*

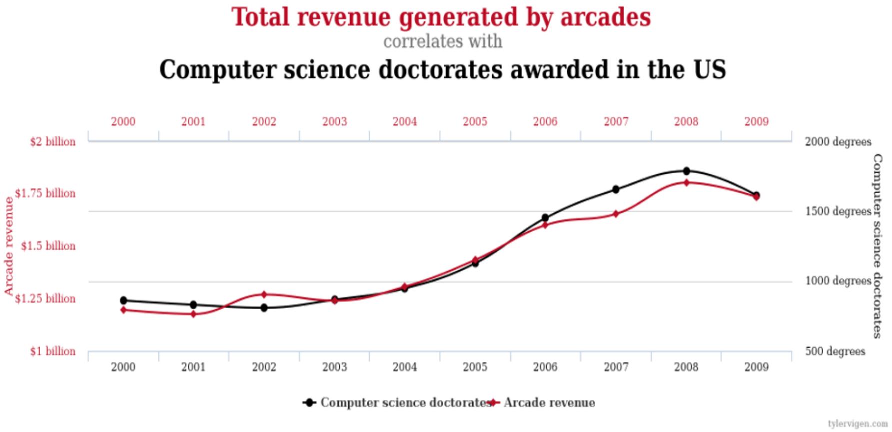

We have lots of data!

Agenda
Introduction
Causality
Why Causality?
Causal Inference Basics
Formulating a framework

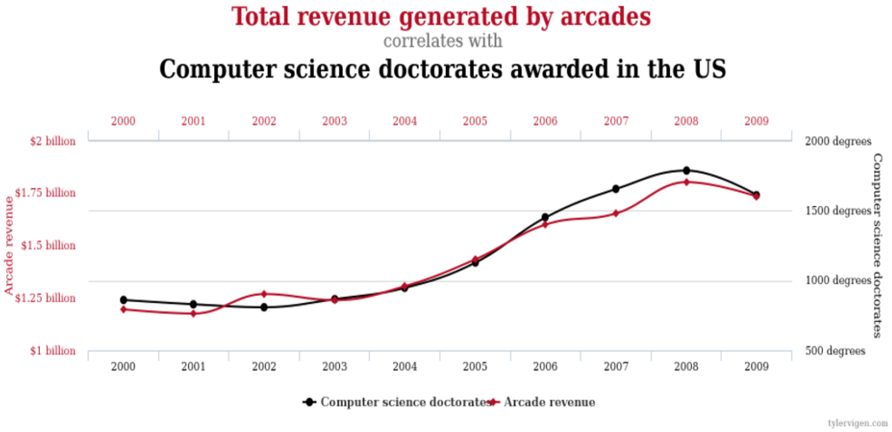

Observed data can be confounded, or even spurious.

A question that has attracted scholars for centuries

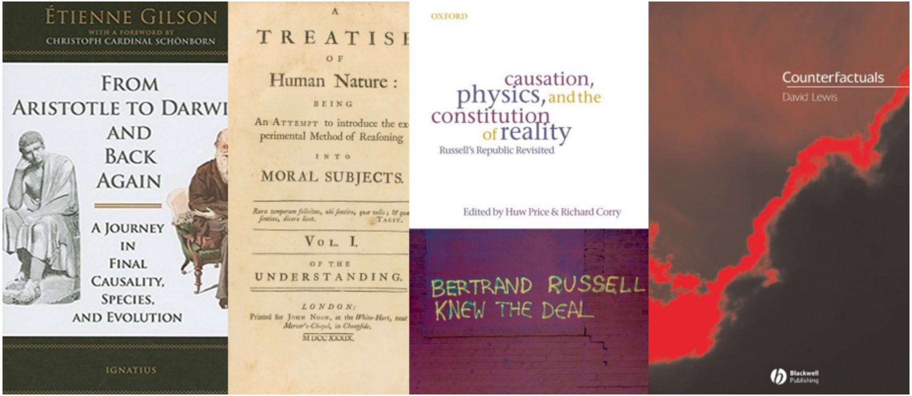

Largely a philosophical pursuit for many centuries.

"More has been learned about causal inference in the last few decades than the sum total of everything that had been learned about it in all prior recorded history"

— Gary King, Harvard University

A "causal revolution" is upon us.

— Judea Pearl, UCLA

What are they talking about?

No one knows what causality means

Hume, 18^{th} century

If you strike a match and it lights up, does the striking cause lighting? What if you repeat the experiment 100 times?

How do you know that striking *always* leads to light? How is it different from regularity or predictability?
Does causality even exist?

Still, everyone aims to find causal effects
Empirical causal inference is pragmatic, best-effort.
Concerns effect of actions **that generalize to all reasonable contexts.** For example, *If I strike a match, would it light up* ***(assuming that everything else stays the same)****?*

"A bag of tricks to produce knowledge!"
—Try the action multiple times
—Try controlling for the environment
—Somehow account for uncontrolled factors

A practical definition

Definition: X causes Y iff
changing X leads to a change in Y,
keeping everything else constant.

The **causal effect** is the magnitude by which Y is changed by a unit change in X.

Called the "interventionist" interpretation of causality.

*Interventionist definition

Bag of tricks ⇨ Powerful statistical frameworks

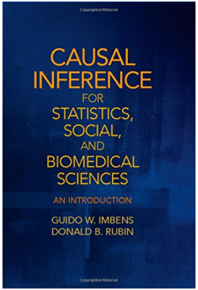

Potential Outcomes

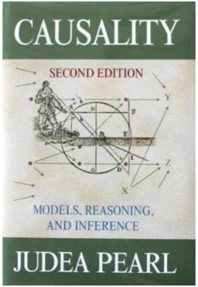

Bayesian Networks

Why should we care?

We have increasing amounts of data and highly accurate predictions. How is causal inference useful?

Predictive systems are everywhere!

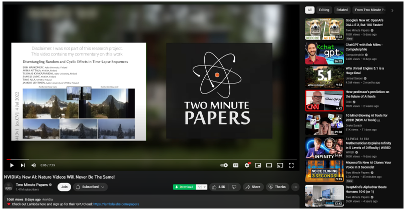

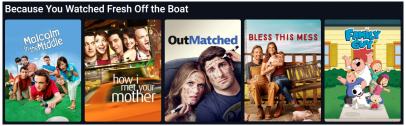

How do predictive systems work?

Aim: Predict future activity for a user.

We see data about their user profile and past activity.

E.g., for any user, we might see their age, gender, past activity and their social network.

From data to prediction

Use these correlations to make a predictive model.

Future Activity ->

f(number of friends, logins in past month)

From data to "actionable insights"

Number of friends can predict activity with high accuracy.
How do we increase activity of users?

Would increasing the number of friends increase people's activity on our system?
Maybe, may be not (!)

Different explanations are possible

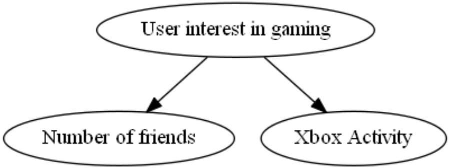

How do we know what causes what?

Decision: To increase activity, would it make sense to launch a campaign to increase friends?

Another example: Search Ads

Search engines uses ad targeting to show relevant ads.

Prediction model based on user's search query.

Search Ads have the highest click-through rate (CTR) in online ads.

Are search ads really that effective?

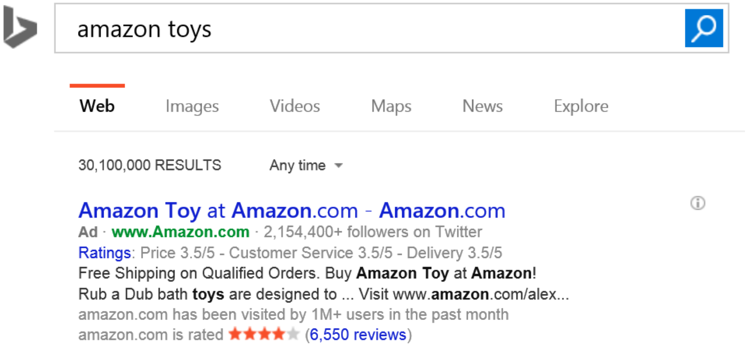

Ad targeting was highly accurate.

Blake-Tadelis-Noskos (2014)

But search results point to the same website

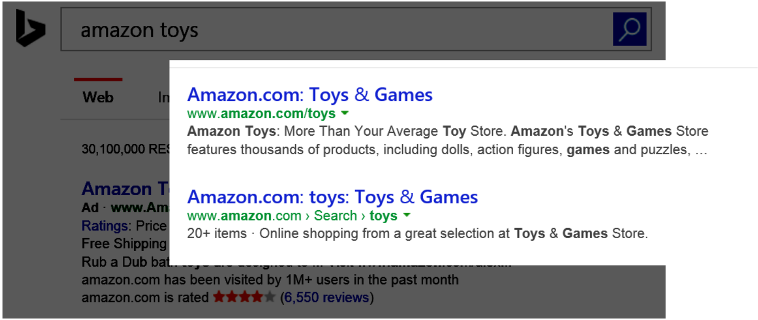

Counterfactual question: Would I have reached Amazon.com anyways, without the ad?

For example, let's take (ads shown) $\rightarrow$ (visit to toy website).

$x\%$ of ads shown are effective

$<x\%$ of ads shown are efffective

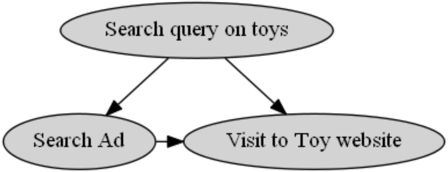

There can be many hidden causes for an action, some of which may be hard to quantify.

Estimating the impact of ads

Toys R Us designs new ads.

Big jump in clicks to their ads compared to past campaigns.
Were these ads more effective?

People always buy more toys in December

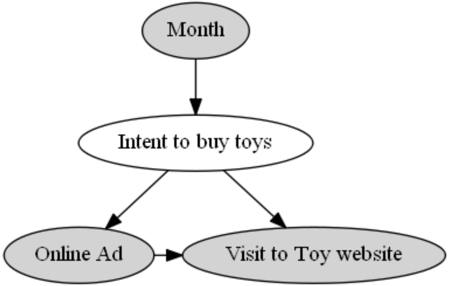

Misleading to compare ad campaigns with changing underlying demand!
Be mindful of hidden causes, or else we might overestimate causal effects.
(But) Ignoring hidden causes can also lead to completely wrong conclusions.

Comparing two algorithms: Old vs New

Two algorithms, A (production) and B (new) running on the system.

From system logs, collect data for 1000 sessions for each.

Measure Success Rate (SR).

Old Algorithm (A)	New Algorithm (B)
50/1000 **(5%)**	54/1000 **(5.4%)**

Is the New algorithm better?

Change in SR by income of people

So let us look at SR separately.

	Old algorithm (A)	New Algorithm (B)
CTR for Low-Activity users	10/400 (2.5%)	4/200 (2%)
CTR for High-Activity users	40/600 (6.6%)	50/800 (6.2%)
Total CTR	**50/1000 (5%)**	**54/1000 (5.4%)**

Is the old Algorithm better? *(The Simpson's paradox)*

Simpson (1951)

Answer (as usual): May be, may be not.

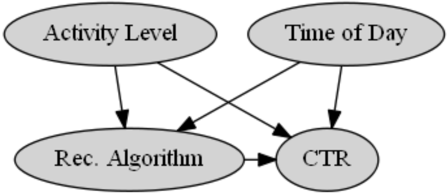

E.g., Algorithm A could have been shown at different times than B.
There could be other hidden causal variations.

Example: Simpson's paradox in Reddit

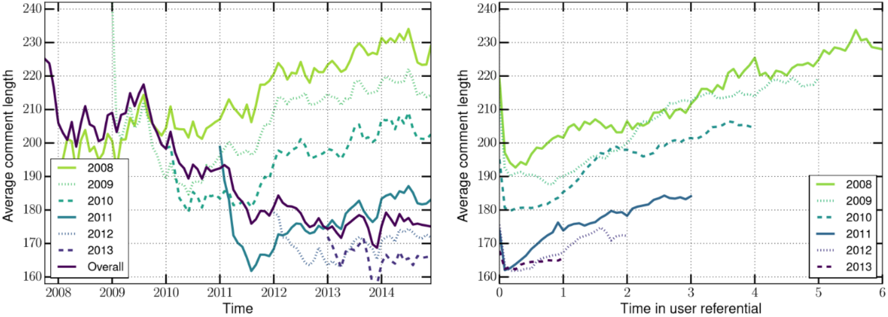

Average comment length decreases over time.
But for each yearly cohort of users, comment length increases over time.

Barbosa-Cosley-Sharma-Cesar (2016)

Making sense of such data can be too complex.

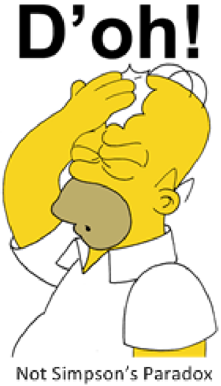

In search of an intervention and a counterfactual: A historical tour of causal inference

. . . And there was fire!

Action: Strike stones.
Outcome: Sparks? Fire?
Probably one of the oldest experiments in causal inference.
Intervention: Taking an action.

The idea of a "controlling" while taking action

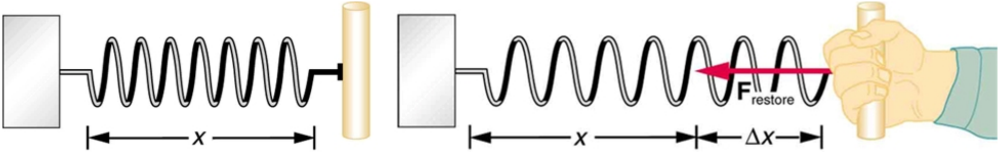

Hooke discovered that pulling a spring causes it to expand in length proportional to the amount of the force applied. [Hooke's law]
Only works if the other end is fixed, or controlled.
A common property of many physical experiments, leading to the notion of a "controlled" experiment.

Intervention: Taking an action while keeping other relevant factors constant.

What if you cannot intervene?

For centuries, controlled experiments have worked well for many physical experiments and do so even today.

However, they do not work in the messier life sciences or social sciences.

We needed another big idea for causality to branch out of physical experiments. The idea of a "counterfactual".

When you introduce a counterfactual and the system doesn't respond with the same outcome, it provides causality through counterfactual explanation.

Observed Data from the Real world

No data from the Counterfactual world

How do we systematically reason about and estimate the relationship between effects and their causes?

How do we formulate this?

Establishing causality requires two things: A specific language and a framework to process that language.

- A language for "complicated" questions.
  - Counterfactuals cannot be expressed by probability.
  - → Pearl's graphical model framework [Pearl 2009]
- Need an estimation framework that understands this language.
  - Methods for estimating causal quantity.
  - → Neyman-Rubin's potential outcome framework [Imbens-Rubin 2016]
- Using do-calculus

The do-calculus is an axiomatic system for replacing probability formulas containing the do operator with ordinary conditional probabilities.

Graphical Models: Express causal relationships visually

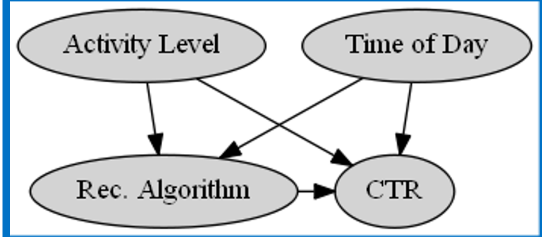

$ctr = f(alg, act, time)$

$alg = g(act, time)$

Edges represent *direct* causes.

Directed paths represent *indirect* causes.

Graphical Models: Express causal relationships as a Bayesian network

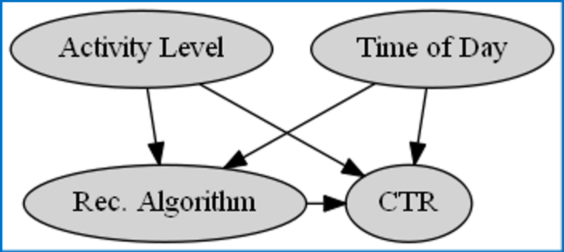

Markov condition: A node is independent of all other non-descendants given its parents.

Leads to factorization of joint probability.

$$P(G) = P(Activity)P(Time)\,P(Alg|Activity,Time)P(CTR|Alg,Time)$$

Graphical Models: Make assumptions explicit

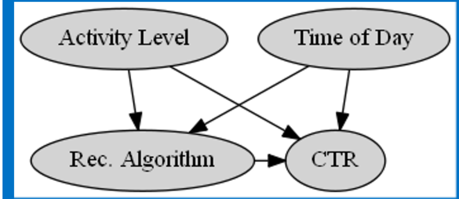

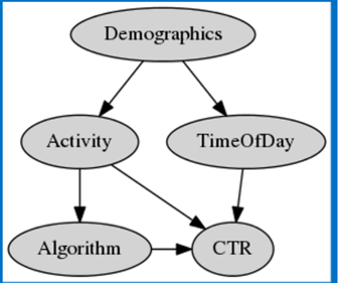

The graph encodes all causal assumptions.

Assumptions are the nodes and edges that are *missing*.

Example: Assumptions encoded in the graph

- *Activity* level of users affects which *Algorithm* they are shown and their overall *CTR*.
- *CTR* is different at different times of day.
- Unobserved *Demographics* of a user determine when they visit the Store, which also affects their *Activity* level, and in turn the *Algorithm* they are shown.

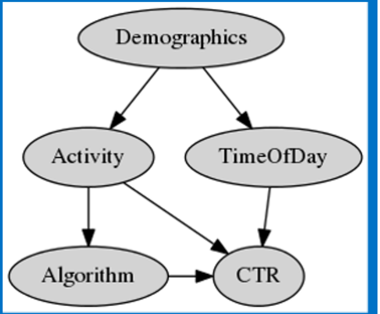

Graphical Models: A language for intervention and counterfactual

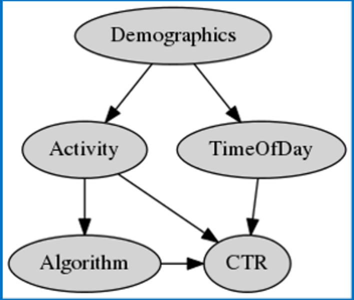

Intervention on a node: Acting to change the node exogenously by severing its ties to its parents.

$$P(CTR|\, do(Alg))$$

which is different from $P(CTR|Alg)$

Graphical Models: A language for intervention and counterfactual

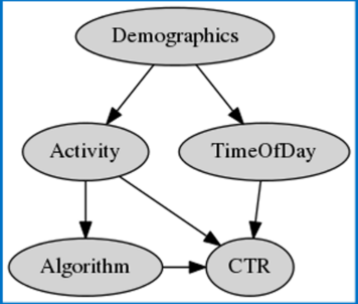

Counterfactual: The recommendations were always shown. But what would have happened if they were not?

$$P(CTR|do(\sim Alg), Alg)$$

Graphical models: Provide an mechanistic way of identifying a causal effect

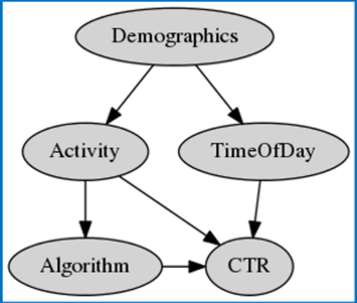

Appeals to the idea of controlling.

When we cannot control the environment, use *conditioning*.

Should we also restrict our comparison to people who come at the same times?

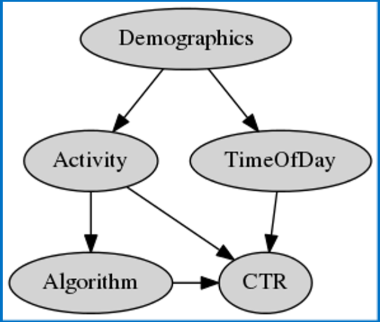

CTR does depend on ***Time of Day*** of a user's visit.

But the algorithm assigned does not change based on ***Time of Day***.

⇒ While ***CTR*** may be different at different times, any ***Algorithm*** is equally likely to be shown at any point in time.

Tricky to find correct variables to condition on. Fortunately, graphical models make it precise.

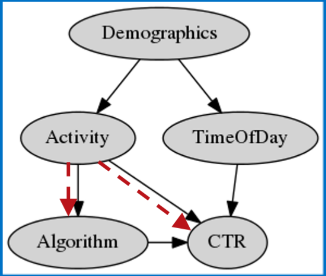

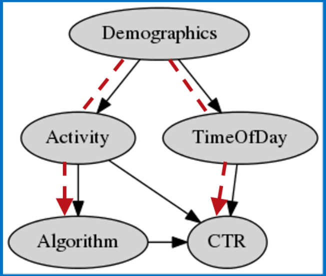

"Backdoor" paths: Look for (undirected) paths that point to both $Algorithm$ and CTR.

Backdoor criterion: Condition on enough variables to cover all backdoor paths

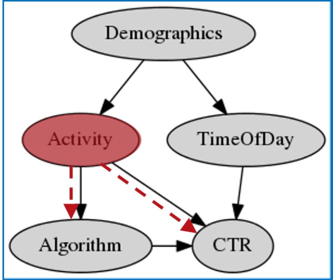

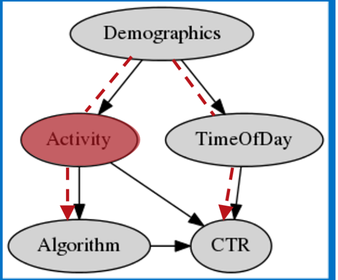

Identified Causal Effect:

$$P(CTR|do(Alg)) = \sum_{activity} P(CTR|Alg, Activity)P(Activity)$$

Identified Causal Effect:

$$P(CTR|do(Alg)) = \sum_{activity} P(CTR|Alg, Activity)P(Activity)$$

For complicated graphs, do-calculus provides a set of rules to automate the identification process.

Wait, but how to estimate this?

Potential Outcomes: Every variable has a counterfactual

Imagine Y_{Alg} *and* $Y_{\sim Alg}$

Need to estimate difference between the real and counterfactual world.

$$E\left[Y_{Alg}\right] - E\left[Y_{\sim Alg}\right]$$
$$= E[Y|do(Alg)] - E[Y|do(\sim Alg)]$$

This formulation has led to some powerful methods for estimating causal effect.

Equivalent to graphical models.

Potential Outcomes: Estimating an effect identified from the backdoor criterion

Imagine Y_{Alg} *and* $Y_{\sim Alg}$

Causal effect:

$$E\left[Y_{Alg}\right] - E\left[Y_{\sim Alg}\right] = E[Y|do(Alg)] - E[Y|do(\sim Alg)]$$
$$= E[Y|Alg, Activity] - E[Y|\sim Alg, Activity]$$

Can estimate using regression.

$$y = \alpha + \beta Alg + \gamma Activity + \epsilon$$

Valid if all effects are linear.

Unifying the two frameworks

Use graphical models and do-calculus for
- modeling the world
- identifying the causal effect

Use potential outcomes-based methods for
- estimating the causal effect

Best practices for Causal Inferencing

Three steps to causal inference:

- Model

  Use a causal graphical model.
  Make assumptions explicit even if it is cumbersome.

- Identify

  Use the graph to identify the specific effect, or check if desired identification strategy is valid (*BEFORE* looking at the data).

- Estimate

  Use any statistical method to estimate the identified effect.
  Best practices: Order of choosing a method for estimating causal effects:

  - ⇨ Randomization (explore-exploit methods for low number of output items)
  - ⇨ Natural experiments (multiple sources of natural experiments)
  - ⇨ Conditioning methods (using them as strong hints)

Now, with all this information, how do we go all the way to causal modeling of systems?

A practical definition for Causality

X causes **Y** *iff* changing **X** leads to a change in **Y** keeping everything else constant.

- But many variables(Xs) are present in problems that require ML/DL/AI.
- Causal discovery is identifying conditional independencies between variables.
- Causal inferencing is analogous to extrapolation after causal discovery.
- Causal reasoning is the process of explaining cause and effect.
- Causal modeling is conceptualizing causality for a complete system.

Content taken from Amit Sharma's presentation on causal inference tutorial.

Thank You!