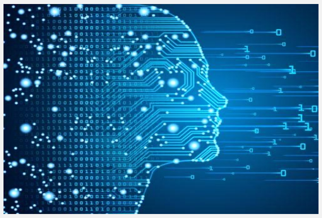

Principles of Causal Inference

Vasant G. Honavar

Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence
Professor of Data Sciences, Informatics, Computer Science and Engineering, Bioinformatics & Genomics, Public Health Sciences and Neuroscience
Center for Artificial Intelligence Foundations and Scientific Applications
Institute for Computational and Data Sciences
Huck Institutes of the Life Sciences
Clinical and Translational Sciences Institute
Northeast Big Data Hub
Pennsylvania State University
vhonavar@psu.edu
http://faculty.ist.psu.edu/vhonavar
http://ailab.ist.psu.edu

Graphical Causal Models

Motivating Examples

- Democracy and GDP
  - You can get (ok) measures of democracy and GDP growth for every country in the world.
  - So you observe these measures for all countries.
  - Can you tell from these data whether democracy has a positive *effect* on GDP growth?
    - No. Maybe something else determines both democracy and GDP growth, and we cannot measure it.
- Cardiovascular health and brushing teeth:
  - Suppose we find that there is a high correlation between brushing teeth regularly and low incidence of heart disease.
  - Can you conclude that not brushing teeth is a cause of heart disease?

Confounding

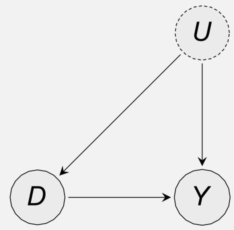

- You find $P(Y|D = d) \neq P(Y|D = d')$. Why?
- Intuition: Measured association between *D* and *Y* consists of
  - causal effect of *D* on *Y* and
  - confounding due to *U*
- Formalize intuition? How can we generally “read off” dependencies from causal graphs?
- Sometimes our causal assumptions can be tested!

Motivation for d-separation

- Step 1. Assuming a specific graph, which dependencies/correlations would we see in the data?
- Step 2. We will think about what "causal effect" actually means.
- Step 3. We will try to equate causal effects with population quantities ("causal inference")

Causal Inference

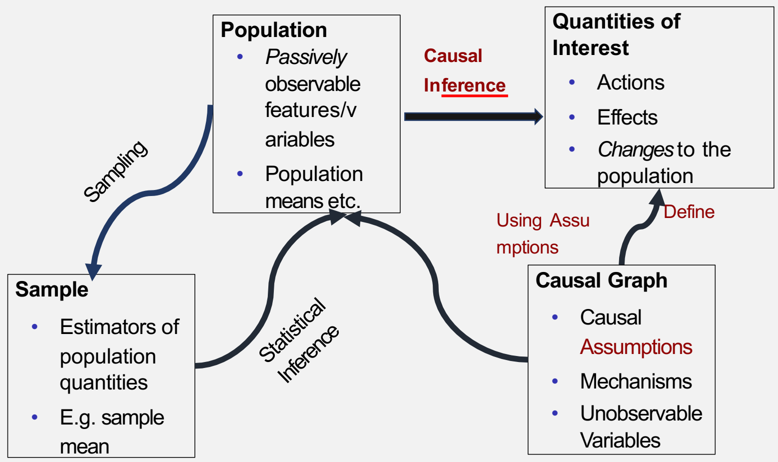

Causal Graphs

- Three elements:
  - Variables (nodes)
  - Edges: **Possible direct** causal effects (we will make this more precise)
  - Missing edges: Strong **assumption** about **absent** causal effects
- Causal graphs are **D**irected **A**cyclic **G**raphs
  - Edges are directed
  - No directed cycles, so no variable causing itself indirectly
- Semantics: Every node is independent of its non descendents given its parents (we will make this more precise)

DAG Terminology

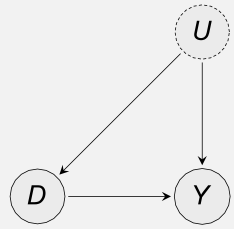

- Descendants of variables = children, grandchildren…
- Ancestors of variables = parents, grandparents…
- A **path** is a sequence of neighboring arrows, without crossing a variable more than once
- Direction of the arrow does not matter (when it does, we will say directed paths)
- What are the edges and paths in this graph that start from *D*?

  $D \rightarrow Y$, $D \leftarrow U$, $D \leftarrow U \rightarrow Y$, $D \rightarrow Y \leftarrow U$

Causal DAGs

- Nodes are independent of non-descendants given their parents

d-separation:

- a graph theoretic criterion for reading independence statements
- can be computed in linear time (in the number of edges)

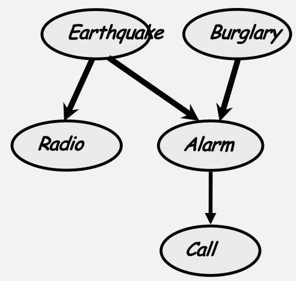

Three Basic Paths

- Aside from direct connections $A \to B$, only three different paths involving three variables possible in DAGs:
  - $A \to B \to C$ (**chain of mediation**)
  - $A \leftarrow B \to C$ (**common cause/fork**)
  - $A \to B \leftarrow C$ (**collider**)
- Strategy: Understand how different paths lead to (conditional) (in)dependence between A and B
- Of course, $A \to B$ means that $A \not\perp\!\!\!\perp B$
- That is, A is not independent of B)

$a \perp\!\!\!\perp b \mid c$

$a \not\perp\!\!\!\perp b \mid c$

Chain of Mediation

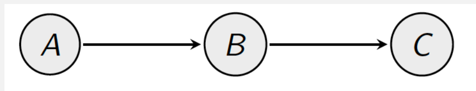

- Rumor A on a given day ($P(A = 1) = 0.5$),
- Person B knows the rumor sometimes ($B = 1$),
- Person B sometimes spreads rumor to person C ($C = 1$).
- B or C do not invent rumors beyond A
- You measure A, B, C for multiple distinct days.
  - Will C be informative about A?
    - Yes, when C knows a rumor, a rumor A definitely occurred
    - $P(A = 1|C = 1) = 1 > P(A = 1) = 0.5$ or $A \not\perp C$

- What happens to this dependence when we only look at days when $B = 1$?

Chain of Mediation

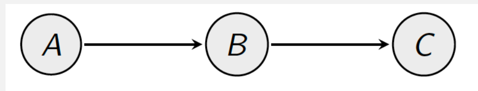

Example:

- You measure A, B, C for multiple distinct days.
  - $A \not\perp\!\!\!\perp C$
- What happens to this dependence when we only look at days where $B = 1$?
- $P(A = 1 | B = 1) = 1$ because B is truthful and does not invent rumors.
- We are sure there was a rumor when we know $B = 1$
- So $P(A = 1 | B = 1) = P(A = 1 | B = 1, C) \implies A \perp\!\!\!\perp C | B$
  (i.e., A is independent of C given B)

Chain of Mediation

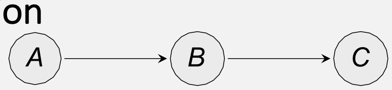

- In chain of mediation: $A \not\perp\!\!\!\perp C$, but $A \perp\!\!\!\perp C | B$
- We say this path is **open** unconditionally, but conditional on the middle node it is **blocked**
- As in "blocking the information flow"
- Note that $P(A|C = 1) > P(A)$, so that C predicts A, while the causal influence actually flows along $A \rightarrow B \rightarrow C$.

Common Cause

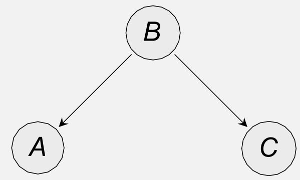

- Example: B denotes Season; A, daily ice cream sales in the city; and C, daily number of drownings in the city
- Question: Does knowing C help you predict A?
  - B is a common cause of both A and C
  - A and C are correlated
  - Knowing the number of drownings is high, implies probably $B =$ summer, and that means ice cream sales C are relatively high
  - So $P(A) \neq P(A | C)$, or $A \not\perp C$

Common Cause

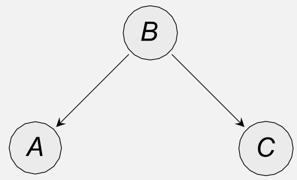

- Example: B denotes Season; A, daily ice cream sales in the city; and C, daily number of drownings in the city. $A \not\perp\!\!\!\perp C$
- What if we look at a subset of the data where season is held constant, e.g. $B =$ summer?
  - During summer, you see still variation in ice cream sales and drownings
  - But once you fix B, changes in B cannot influence A or C.
  - $A \perp\!\!\!\perp C | B$

Common Cause

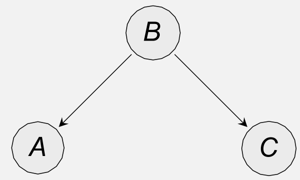

- So in the common cause/fork graph, $A \not\perp\!\!\!\perp C$ but $A \perp\!\!\!\perp C | B$
- Unconditionally, the path is **open**. Conditional on the middle node, it is **blocked**
- This is exactly like in the chain of mediation, but different “causal story”

Collider

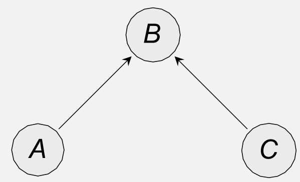

- Example: *A* and *C* denote independent coin flips (results are 0 or 1).
  - Suppose *B* is the sum of A and B, so 0, 1, or 2
- Can you predict *A* when you know *C*?
  - No! Because $A \perp\!\!\!\perp C$
- But suppose you know one coin *A* = 1 and *B* = 1. Can you then predict the other coin *C*?
  - Yes, it HAS to be 0
- So here $A \perp\!\!\!\perp C$, but $A \not\perp\!\!\!\perp C | B$

Descendants of Colliders

- A and C are independent coin flips,
- B is the sum of A and B
- $D = 1$ if $B = 2$, 0 otherwise

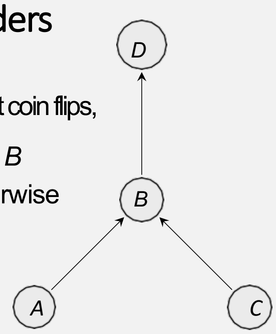

- You know $A = 1$ and $D = 0$. Can you predict C?
- Yes, same reasoning: C has to be 0, otherwise D would be 1
- So $A \not\perp\!\!\!\perp C \mid D$: Conditioning on descendants of colliders has same qualitative consequences as conditioning on colliders themselves

Conditioning on Colliders: Visualization

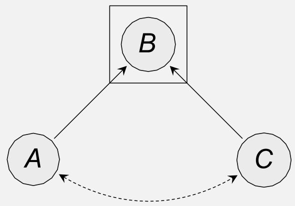

- Sometimes helpful: Boxed variable indicates we condition on it
- Bi-directed arc between the start and end node indicates associations created by conditioning on collider
- **Does not indicate cyclic causation**; rather equivalent to an additional common cause of the nodes connected by the bi-directed arc
- Is then treated as a normal path

d-separation

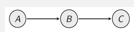

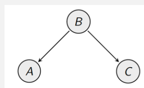

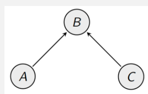

- Chain of mediation: Path is open unconditionally, but blocked conditional on the middle node. $A \not\perp\!\!\!\perp C$ but $A \perp C|B$.
- Common cause/fork: Path is open unconditionally, but blocked conditional on the middle node. $A \not\perp\!\!\!\perp C$ but $A \perp\!\!\!\perp C|B$.
- Collider: Path is blocked unconditionally, but open conditional on the middle node or one of its descendants. $A \perp\!\!\!\perp C$ but $A \not\perp\!\!\!\perp C|B$.
- What if there are multiple, longer paths between A and C? Will A and C be (conditionally) independent? **d-separation** gives the answer

d-separation

- A path p is blocked by a set of nodes Z if and only if
  1. p contains a chain of nodes $A \rightarrow B \rightarrow C$ or a fork $A \leftarrow B \rightarrow C$ such that the middle node B is in Z (i.e., B is conditioned on), or
  2. p contains a collider $A \rightarrow B \leftarrow C$ such that the collision node B is not in Z, and no descendant of B is in Z
- If Z blocks every path between two nodes X and Y, then X and Y are **d-separated**, conditional on Z, and thus are independent conditional on Z
- We sometimes write $(X \perp\!\!\!\perp Y|Z)_G$, "$Z$ d-separates X from Y in graph G"
- $(X \perp\!\!\!\perp Y|Z)_G \rightarrow X \perp\!\!\!\perp Y|Z$ (**testable implication** of the graph)
- "d-separation" = "directional separation" (in directed graphs)
- Path p may be very long, but as long as you block sub-path, you block the whole path
- X, Y, Z may contain multiple variables

d-separation?

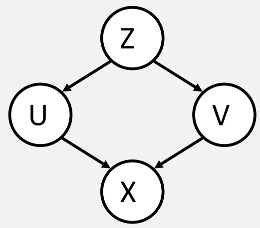

- $X \perp\!\!\!\perp Z$? No.
- $X \perp\!\!\!\perp Z \mid U$? No
- $X \perp\!\!\!\perp Z \mid U, V$? Yes

d-separation

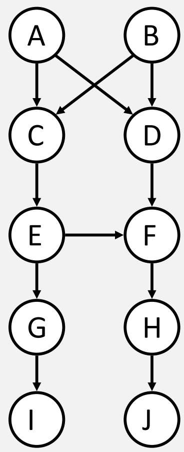

$C \perp\!\!\!\perp D$?

$C \perp\!\!\!\perp D \mid A$?

$C \perp\!\!\!\perp D \mid A, B$?

$I(C \perp\!\!\!\perp D \mid A, B, J)$?

d-separation

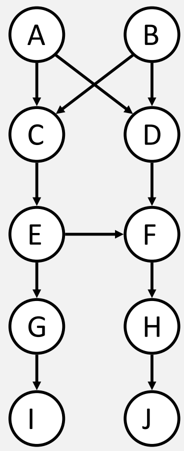

C ⫫ D? – No
C ⫫ D | A? – No
C ⫫ D | A, B? – Yes
I(C ⫫ D | A, B, J)? – No

d-separation: Example

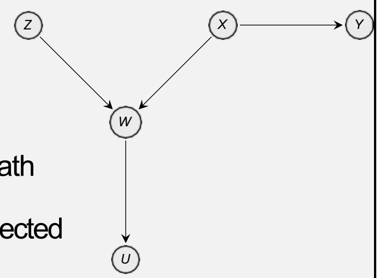

- *Z* and *Y* d-separated?
  - Yes: Collider *W* blocks only path
- What happens when I condition on *U*?
  - Just like conditioning on *W* : Opens path $Z \leftarrow W \rightarrow X \rightarrow Y$
  - Conditional on *U*, *Z* and *Y* are d-connected
- What happens when I condition on *U, X*?
- Definition: Path blocked if middle node of chain/fork is conditioned **or** collider not conditioned **or both**
  - So *U, X* d-separates *Z* and *Y*

d-separation

- Theorem [Verma & Pearl, 1998]: If a set of evidence variables *E* *d*-separates *X* and *Z* in an SCM then $X \perp\!\!\!\perp Z \mid E$
- *d*-separation can be computed in linear time using a depth-first search like algorithm.
- *d*-separation can be used to test whether finding out about the value of one variable might give us any additional hints about some other variable, given what we already know.

Causal Graphs as Structural Equations

Causal Graphs as Structural Equations

- What are U_A and U_B ?
  - **Unobserved** factors that causally influence the observables (coin flips in this example), and **structural errors**
  - U_A and U_B are random variables, and hence so are the observed variables

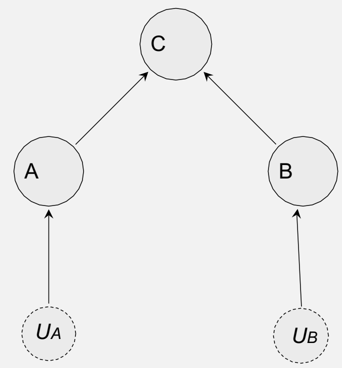

Causal Graphs as Structural Equations

- $U_A \perp\!\!\!\perp U_B$
- $A = f_A(U_A)$
- $B = f_B(U_B)$
- $C = f_C(A, B)$, e.g., $C = A + B$
- A structural causal model describes
  - Our qualitative beliefs about nature assigns values to variables of interest in a domain of study
  - Causal assumptions
- $E(A) = E\big(f_A(U_A)\big)$

  $= \sum_{u_A}(U_A = u_A)P(f_A\,(u_A))$

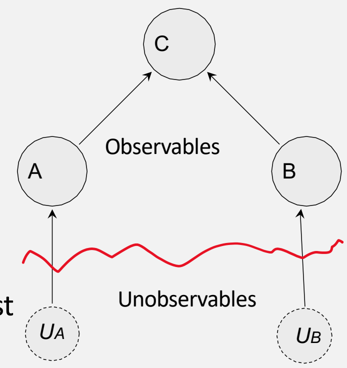

Structural Causal Models

Three main ingredients:

- Observed variables V
- Unobserved variables U
  - Unobserved causes of V
  - U may contain **infinitely** many variables (U_1, U_2,…)
  - U describe unobserved causes of any relevant feature of a unit
  - We usually do **not make distributional assumptions** on U (and hence V)
- Structural functions f for each observable in V
  - When we do not specify f the form of f, we say that functions are **nonparametric**
- V = endogenous variables (explained in the model)
- U = exogenous variables (not explained in the model)

SCM provide a language for expressing counterfactuals

- $do(X)$ denotes intervention on X
- Intervention on X has the effect of removing all incoming links into X (or eliminating all direct causes of X)

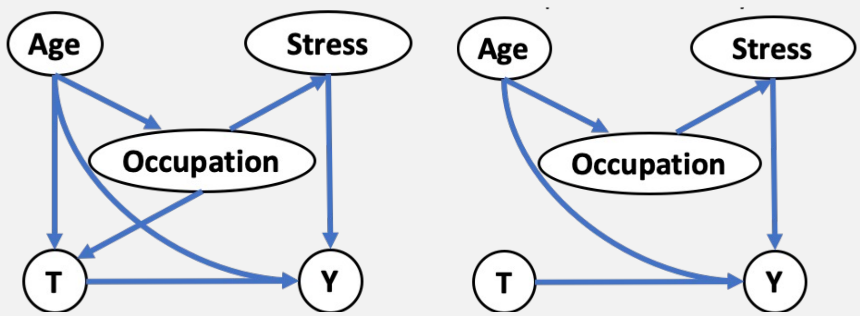

Connecting the SCM and Joint Probability Distribution

- Under some assumptions (Causal Markov Condition) an SCM represents a factorization of the joint probability distribution over the observables:

$$P(X_1, X_2, ..., X_n) = \prod_{i=1}^{n} P(X_i \mid DirectCauses(X_i))$$

- The above equation specifies the full joint probability distribution over the model variables.
- More on this later

What is a Direct Cause?

- The direct causes of X_i are the variables which will change the distribution of X_i as we vary them while holding everything else unchanged

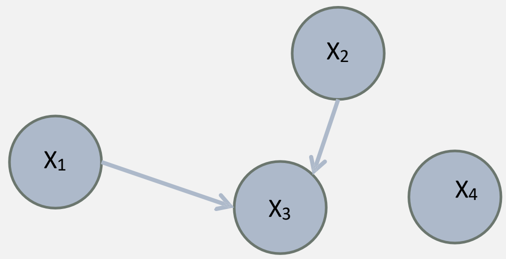

$P(X_3 = x_3 \mid do(X_1 = x_1), do(X_2 = x_2), do(X_4 = x_4)) \neq$
$P(X_3 = x_3 \mid do(X_1 = x_1'), do(X_2 = x_2), do(X_4 = x_4))$

$P(X_3 = x_3 \mid do(X_1 = x_1), do(X_2 = x_2), do(X_4 = x_4)) =$
$P(X_3 = x_3 \mid do(X_1 = x_1), do(X_2 = x_2), do(X_4 = x_4'))$

Causal Graphs as Structural Equations

- $U_A \perp\!\!\!\perp U_B$
- $A = f_A(U_A)$
- $B = f_B(U_B)$
- $C = f_C(A, B)$, e.g., $C = A + B$
- A structural causal model describes
  - Our qualitative beliefs about nature assigns values to variables of interest in a domain of study
  - Causal assumptions

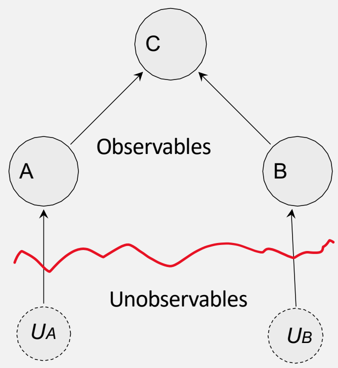

- Suppose $C = A + B$.
- Intuitively, what's the **causal effect** of A on C?
- $C^{a=1} - C^{a=0} = (1 + B) - (0 + B) = 1$

Linear Structural Models

- Structural causal model is NOT a regression.
- It is not an algebraic equation
- It is a causal model: LHS is caused by RHS!
  - Rearranging $C = A + B$ to obtain $A = C - B$ does not make sense!
- However,
  - You can use observed B and observed C to predict A (perfectly) $E[A|B,C] = C - B$
  - Even perfect regression fit does not tell you anything about causation!
- When we use an equation, we need to state whether it is structural causal model or a regression
- We will use
  - $Y = f_Y(\ldots)$ for structural models and
  - $E[Y|D] = f(D)$ for regressions

Linear Structural Equations

- $C = A + B$ is a special case of a linear structural model $Y = \alpha + \beta D + \epsilon_Y$
- This is NOT a regression. A regression describes $E[Y|D]$: The mean of Y given **observations** of D
- A structural model is a mechanism for the generation of Y, and predicts Y when you **control** D, and can be represented using a causal graph
- A regression is associational, you observe D, predict Y
- The regression error is, by construction, independent of D
- The structural error ϵ_Y **may** be independent of D, **if there is no variable that influences** Y **that also influences** D (clear from graph!)

Parametric versus Nonparametric Structural Equations

- In general, causal graphs are representations of
  - non-parametric structural equations (as general as possible)
- Causal graphs based on parametric e.g., *linear* structural equations are very restrictive, but easier to understand
- We will often use linear models to highlight basic ideas and then (try to) generalize them to nonparametric models

Causal Effects as Interventions

- Suppose $C = A + B$.
- How can we obtain the causal effect of A on C?
- By intervening on A independent of other variables and comparing $C^{a=1}$ with $C^{a=0}$
- $C^{a=1} - C^{a=0} = 1$

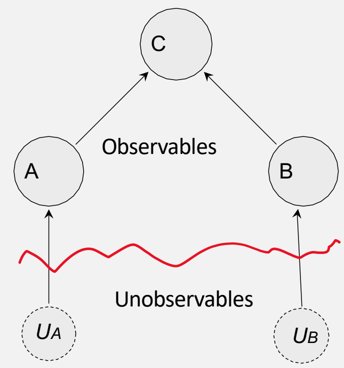

Causal Effects as Interventions

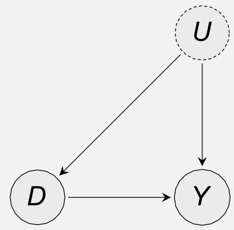

- What structural equation does the above model describe?
- $D = f_D(U)$
- $Y = f_Y(D, U)$
- What happens to the causal graph and the structural equations when we intervene on D i.e., "do" $D = 1$?

Causal Effects as Interventions

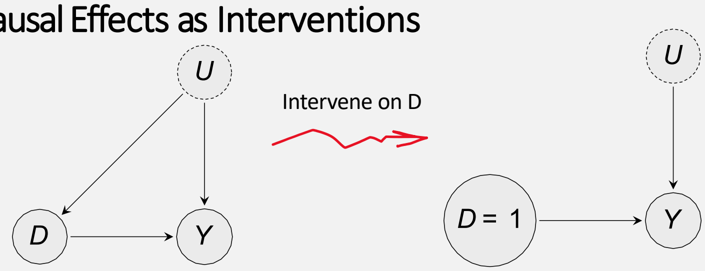

- $D = f_D(U)$
- $Y = f_Y(D, U)$

- $D = 1$
- $Y = f_Y(1, U)$

Causal Effects as Interventions

- $D = 1$
- $Y = f_Y(1, U)$
- This Y under the intervention is a function of U (so differs across units because U may vary across units)
- The mean of Y under the intervention $do(D = 1)$ is averaged over U:

$$E[Y \mid do(D = 1)] = \sum_u f_Y(1, u) P(U = u)$$

If D is a binary explanatory or "treatment" variable, we call

$$E[Y \mid do(D = 1)] - E[Y \mid do(D = 0)]$$

the **causal effect of D on Y**

Causal Effects as Interventions

- $E[Y|do(D=1)] - E[Y|do(D=0)]$ is the **causal effect of** D **on** Y
- $E[Y|do(D=1)]$ is the average outcome if one forces $D=1$ for all individuals
- Correlation is not causation
$$E[Y|D=1] - E[Y|D=0] \neq E[Y|do(D=1)] - E[Y|do(D=0)]$$
- Observation is not intervention
- Seeing is not the same as doing!

Causal Effects as Interventions

- $E[Y \mid do(D = 1)] - E[Y \mid do(D = 0)]$ is the **causal effect of D on Y**
- We refer to learning

  $$E[Y \mid do(D = 1)] - E[Y \mid do(D = 0)]$$

  as the identifying the **causal effect of D on Y**
- To "identify" something with something else is to assert (with justification that the two things are equal

Identifying causal effects

- We refer to learning

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

as the identifying the **causal effect of D on Y**

- Two approaches to identify $E[Y|do(D = d)]$:
  - Intervene in the real world
    - Intervene on D independently of other variables (e.g., conduct a randomized experiment): $\forall d\ do(D = d)$
    - Observe the resulting interventional outcomes
    - Calculate the causal effect of D on Y
  - Under identifiability assumptions (SUTVA) try to (uniquely) equate a causal effect of interest with a function of the population distribution $P(Y, D, X)$, which we observe passively without intervening

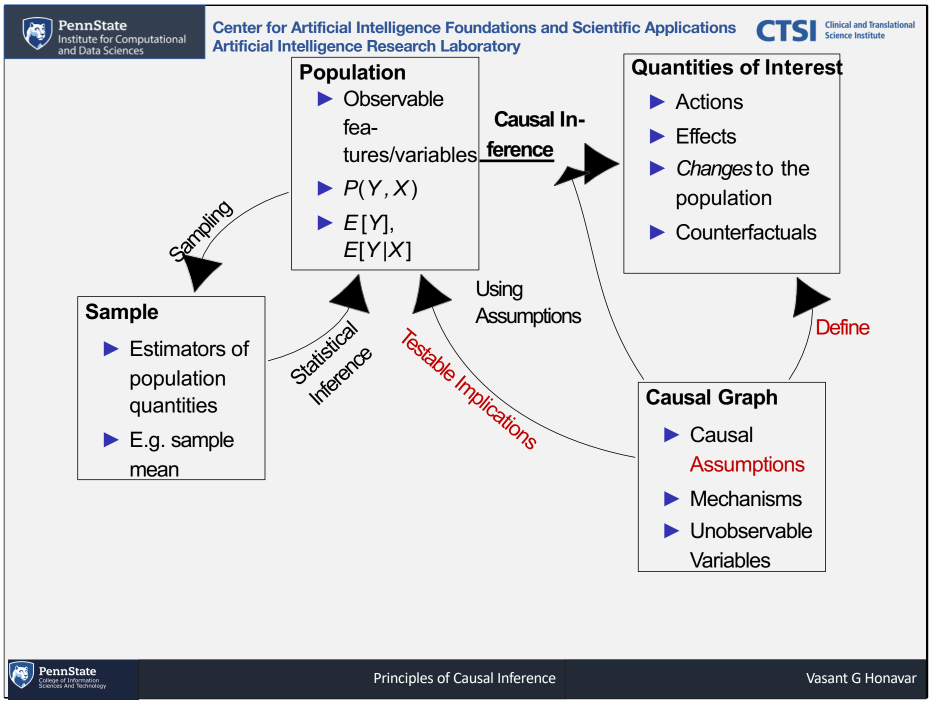

PennState Institute for Computational and Data Sciences
Center for Artificial Intelligence Foundations and Scientific Applications
Artificial Intelligence Research Laboratory
CTSI Clinical and Translational Science Institute
Population
Observable features/variables
$P(Y, X)$
$E[Y]$, $E[Y|X]$
Causal Inference
Quantities of Interest
Actions
Effects
Changes to the population
Counterfactuals
Sampling
Sample
Estimators of population quantities
E.g. sample mean
Statistical Inference
Using Assumptions
Testable Implications
Define
Causal Graph
Causal Assumptions
Mechanisms
Unobservable Variables
PennState College of Information Sciences And Technology
Principles of Causal Inference
Vasant G Honavar

Structural Causal Models – The Story so far

- Directed acyclic graphs or causal graphs
  - Three canonical path types
  - $A \rightarrow B \rightarrow C$
  - $A \leftarrow B \rightarrow C$
  - $A \rightarrow B \leftarrow C$
- d-separation: Variables Z d-separates variables X and Y if Z blocks every path between X and Y
- d-separation implies conditional independence
  - If Z d-separates X from Y in a causal graph G that is, $(X \perp\!\!\!\perp Y \mid Z)_G \rightarrow X \perp\!\!\!\perp Y \mid Z$
  - (Note the overloading of $\perp\!\!\!\perp$)
  - d-separation is testable from data using suitable independence tests

Structural causal models – the story so far

- Causal graphs specify a set of structural equations or a structural causal model (SCM)
- SCM causally connect observable variables in V with other observable variables and / or unobservable variables in U ("error" terms) via structural functions f
- f specify causal mechanisms that describe how nature assigns values to observable variables based on the values of other variables
- Structural equations are not regressions, which are purely predictive
- Structural causal models can be used to specify causal effects in terms of interventions $do(D = d)$ a minimal intervention on only D, independent of other variables, by setting it to some value d
- Average causal effect of (binary) D on Y is given by

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

Exercise

- Consider a linear causal model given by

$$Z = \alpha_Z + U_Z$$
$$D = \alpha_D + \beta_{ZD} Z + U_D$$
$$Y = \alpha_Y + \beta_{DY} D + U_Y$$

where U_Z, U_D, U_Y, WLOG are assumed to have zero mean.

- Draw the corresponding linear structural causal model, assuming that the exogenous variables U_Z, U_D, U_Y are independent
- Calculate the causal effect of D on Y and of Z on Y

Exercise

$$Z = \alpha_Z + U_Z$$
$$D = \alpha_D + \beta_{ZD} Z + U_D$$
$$Y = \alpha_Y + \beta_{DY} D + U_Y$$

where U_Z, U_D, U_Y, are independent and have zero mean.

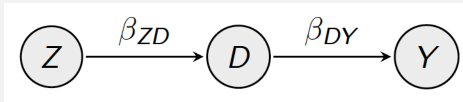

Show that:

$$E\left[D \,\middle|\, do(Z = 1)\right] - E\left[D \,\middle|\, do(Z = 0)\right] = \beta_{ZD}$$
$$E\left[Y \,\middle|\, do(D = 1)\right] - E\left[Y \,\middle|\, do(D = 0)\right] = \beta_{DY}$$
$$E\left[Y \,\middle|\, do(Z = 1)\right] - E\left[Y \,\middle|\, do(Z = 0)\right] = \beta_{ZD} \cdot \beta_{DY}$$

Exercise

$$Z = \alpha_Z + U_Z$$
$$D = \alpha_D + \beta_{ZD} Z + U_D$$
$$Y = \alpha_Y + \beta_{DY} D + U_Y$$

$U_Z \perp\!\!\!\perp U_D$ but $U_D \not\!\perp\!\!\!\perp U_Y$

- Draw the structural causal model
- What are the testable implications?
- What is the causal effect of Z on Y?

$$E[Y \mid do(Z = 1)] - E[Y \mid do(Z = 0)]$$

Exercise

- Consider the structural causal model shown
- Write down the structural equations
- Calculate the causal effect of Z on Y

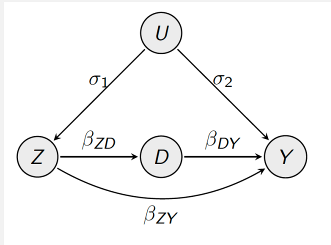

Exercise

$$Z = \alpha_Z + U_Z$$
$$D = \alpha_D + \beta_{ZD} Z + U_D$$
$$Y = \alpha_Y + \beta_{DY} D + U_Y$$

$U_Z \perp\!\!\!\perp U_D$ but $U_D \not\perp\!\!\!\perp U_Y$

- Draw the structural causal model
- What are the testable implications?
- What is the causal effect of Z on Y?

$$E[Y \mid do(Z = 1)] - E[Y \mid do(Z = 0)]$$

Path Tracing in Linear Graphs

- If a graph represents a linear SCM, where additionally all variables are assumed to be normally distributed with mean 0 and variance 1, then to find $E[Y|Z = 1] - E[Y|Z = 0]$
  - List all **open (not blocked)** paths between Z and Y
  - Multiply all path/structural coefficients (= causal effects) along a given path, and sum up the results
- Conditional causal effects are a bit more involved
  - In this course, we will use an approximate solution:
  - For $E[Y|Z = 1, X = x] - E[Y|Z = 0, X = x]$, if X does not open up additional paths between Z and Y, do the above, but **only across paths that are not blocked conditional on X**

Path Tracing in Linear Graphs: Colliders

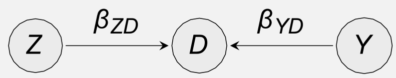

- What's $E[Y|Z = 1] - E[Y|Z = 0]$?
- Zero
- Why?

Path Tracing in Linear Graphs: Colliders

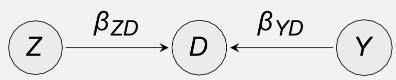

- $E[Y|Z = 1] - E[Y|Z = 0] = 0$
- Let's say β_{ZD} and β_{DY} are positive.
- Is $E[Y|Z = 1, D = 1] - E[Y|Z = 0, D = 1]$ positive?
  - Look at units with same $D = 1$, but different Z.
  - If you have $Z = 0$ but still $D = 1$, that must be because Y makes up for lack of Z
- So mean difference is **negative**
- $E[Y|Z = 1, D = 1] - E[Y|Z = 0, D = 1] \approx -\beta_{ZD} \cdot \beta_{DY}$
- We will return to this. See section 3.8, Pearl, Glymour, Jewell (2016) and Pearl (2013): "Linear Models: A Useful Microscope for Causal Analysis"

Estimating causal effect in the presence of confounding

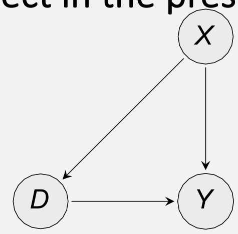

- Confounder is any variable that impacts both "treatment" and "outcome"
- Which paths does the association between D and Y consist of?
  - Causal effect of D on Y and
  - Confounding due to X
- We want to estimate $E\,[Y\,|do(D\ =\ d)]$
- How?

Estimating causal effect in the presence of confounding

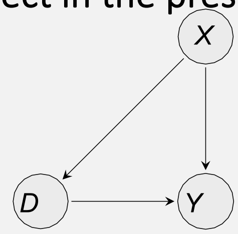

- How can we obtain $E\,[Y\,|do(D\ =\ d)]$
- Intervene on D, i.e., $do(D = d\,)$ independently of all other variables
- If you cannot intervene on D, find control variables that can be used to
  - block all “non-causal” paths between D and Y
  - while leaving open all causal paths between D and Y
  - without opening up any “non-causal” paths (colliders…) between D and Y

Estimating causal effect in the presence of confounding

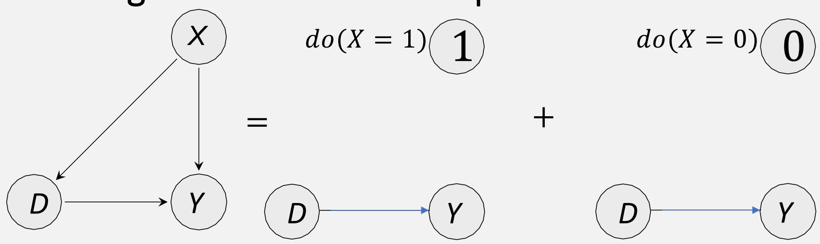

- Suppose we control for X
- We block the non-causal paths between D and Y without eliminating the causal paths or introducing any non-causal paths
- Now we can estimate the causal effect of D on Y separately from observational data with $X=0$ and with $X=1$ and take a weighted average of the two effects where the weights correspond to $P(X=1)$ and $P(X=0)$

The Back-Door Criterion

- Given an ordered pair of variables (D, Y) in a DAG G, a set of variables X satisfies the backdoor criterion relative to (D, Y) if
  - no node in X is a descendant of D, and
  - X blocks every path between D and Y that contains an arrow into D
- Ordered pair because we are interested in the causal effect if D on Y
- A path that starts with an arrow into D is called a **back-door path**
  - Blocking back-door paths makes sure we block "bad" paths
  - Not conditioning on descendants of D ensures that we leave all "good" paths open and that we do not open up new bad paths
- Applicable for any DAG, and hence non-parametric, distribution-free

The Back-Door Criterion: Example

- Suppose we want the causal effect of interest D on Y .
- Which variables do we need to adjust for?
  - *X*, in order to block back-door path *D* ← X → *M* → *Y*
- Does (X, V) also satisfy the back-door criterion?
  - Yes, blocks only back-door path, no descendant
- May we condition on M?
  - No, M is a descendant of D

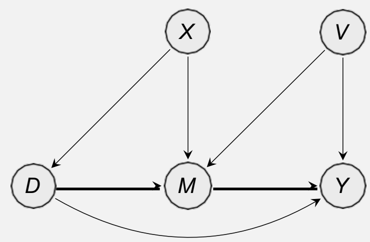

Backdoor criterion: Example

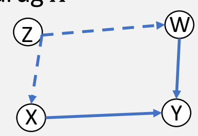

- Suppose we want to estimate the causal effect of a drug X on recovery Y
- X, Y, W are observed, Z is not.
- Is there any unobserved confounder?
- Yes, Z is an unobserved confounder
- How do we de-confound the causal effect X on Y?
- Look for an observed variable that satisfies the backdoor criterion
- W is such a variable – it blocks the backdoor path X ⍰ $Z \rightarrow W \rightarrow Y$, W is not a descendent of X
- Upon adjusting for W we have

$$P(Y = y|do(X = x)) = \sum_w P(Y = y|X = x, W = w)P(w)$$

- Hence, the causal effect of X on Y is identifiable from observational data

Using the Back-Door Criterion for Identification

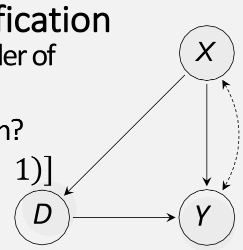

- Bi-directed arc is additional unobserved confounder of X and Y and hence D and Y
- Does X fulfill the BDC wrt D and Y in this graph?
- We want $E[Y|do(D=1)] - E[Y|do(D=1)]$
- We measure $P(Y, D, X)$.
- We somehow need to condition on X
- Question: How can we re-express $E[Y|do(D=1)]$ as something that is conditional on X without making additional assumptions?

- Law of Iterated Expectations!

$E[Y|do(D=1)] =$

$$\sum_x E[Y|do(D=1), X=x]\, P(X=x|do(D=1))$$

Using the Back-Door Criterion for Identification

$$E[Y\,|do(D=1)] =$$
$$\sum_x E[Y\,|do(D=1), X=x]\cdot P(X=x\,|do(D=1))$$

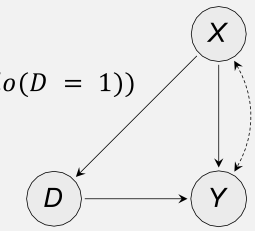

- What is $P(X=x\,|do(D=1))$?
- D does not affect X, so $P(X=x\,|do(D=1)) = P(X=x)$

$$E[Y\,|do(D=1)] =$$
$$\sum_x E[Y\,|do(D=1), X=x]\cdot P(X=x)$$

Using the Back-Door Criterion for Identification

$$E[Y \mid do(D = 1)] = \sum_x E[Y \mid do(D = 1), X = x] \cdot P(X = x)$$

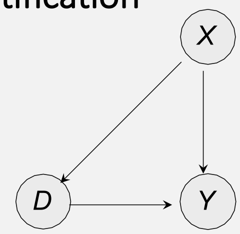

- How to get rid of the other $do(D = 1)$?
- Conditional on X , observing $D = 1$ is the same as **do**ing $D = 1$, at least with respect to Y

$$E[Y \mid do(D = 1)] = \sum_x E[Y \mid D = 1, X = x] \cdot P(X = x)$$

Hence the causal effect of D on Y is given by

$$\sum_x (E[Y \mid D = 1, X = x] - E[Y \mid D = 0, X = x])P(X = x)$$

Estimation

Causal effect of D on Y is given by

$$\sum_x (E\,[Y \mid D = 1, X = x] - E[Y \mid D = 0, X = x])P(X = x)$$

- With population data:
  - Compute x -specific difference in means, then compute weighted average of those x –specific differences, using $P(X = x)$
- With sample:
  - One-on-one matching. For every unit in sample with $X = x$ and $D = 1$, find a matching person with $X = x$, but $D = 0$.
  - Compute pair-wise difference in Y .
  - Take their mean.

The Back-Door Criterion

- Given an ordered pair of variables (D, Y) in a DAG G, a set of variables X satisfies the backdoor criterion relative to (D, Y) if
  - no node in X is a descendant of D, and
  - X blocks every path between D and Y that contains an arrow into D
- Ordered pair because we are interested in the causal effect if D on Y
- A path that starts with an arrow into D is called a **back-door path**
  - Blocking back-door paths makes sure we block "bad" paths
  - Not conditioning on descendants of D ensures that we leave all "good" paths open and that we do not open up new bad paths
- Applicable for any DAG, and hence non-parametric, distribution-free

Back-door criterion

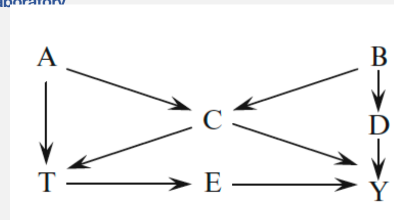

- Which nodes satisfy the back-door criterion for causal effect of T on Y?
  - {A, C}
  - {B, C}
  - {C, D}
  - {A, B, C}
  - {A, C, D}
  - {B, C, D}
  - {A, B, C, D}

Back-door criterion Dissected

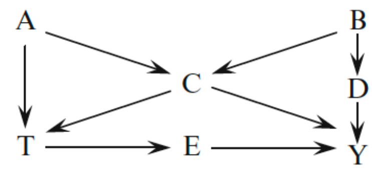

- All backdoor paths start with an arrow into the treatment variable
- This implies they can be blocked by conditioning on parents of the treatment, if they are observed
- Parents of treatment (POT) criterion: Any causal effect of T on Y can be identified by conditioning on all of the parents of T if they are observed.
- Sets of nodes meeting POT criterion:
  - {A, C}

Back-door criterion Dissected

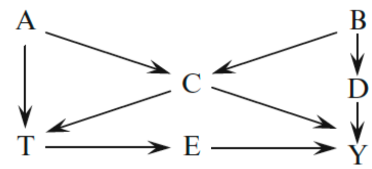

- Parents of the outcome (POO) criterion: If no backdoor path shares a node with any causal path (other than T and Y), then conditioning on all parents of the outcome Y (if they are observed) that do not lie on a causal path from T to Y identifies the total causal effect of T on Y.
- Do these sets of nodes meet the POO criterion?
  - {C,D} Yes
  - {C, D, E} No. (Why?)

Back-door criterion Dissected

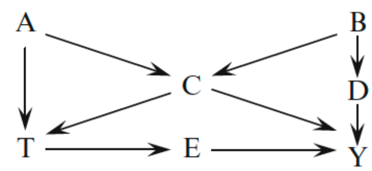

- All unconditionally open backdoor paths must contain a variable that is a joint direct or indirect cause of treatment T and outcome, Y.
- Joint Ancestor (JAN) Criterion: Conditioning exclusively on all joint ancestors of T and Y identifies the total causal effect of T on Y. (Must avoid conditioning on additional variables)
- Can you find a set of nodes meet the JAN criterion?
  - {A,B,C} Yes

Back-door criterion Dissected

- Identifiability using the backdoor criteria or any of its special cases (POT, POO, JAN) assume that there exist at least one set of covariates that satisfy the criteria are indeed observed.
- Backdoor criterion implies ignorability (of Rubin's potential outcomes framework) and identifies a set of covariates that when controlled and adjusted for, ensure ignorability when the causal graph is correctly specified (meaning all of the relevant variables are included, and no edges that should be present are omitted).

Back-door criterion Dissected

- What if the complete structure of the DAG is unknown?
- Confounder selection criterion (COS) is implied by the backdoor criterion
- Even if the DAG is only partially specified, if there is a set of observed covariates that meets the backdoor criterion (i.e., if we are willing to assume that the unobserved variables do not influence who gets treated), then it suffices to condition on all observed pretreatment covariates that either cause treatment, outcome, or both (VanderWeele and Shpitser, 2011).

Confounding in the language of causal calculus

- Confounder is any factor that makes $P(Y|X) \neq P(Y|do(X))$
- To de-confound two variables X and Y
  - We need to block all non-causal paths between X and Y without perturbing any causal paths
    - A backdoor path is any path from X to Y that starts with an arrow pointing into X
    - X and Y will be de-confounded if we block every such backdoor path
    - If we do this by controlling for some variables Z, we need to make sure that no member of Z is a descendent of X on a causal path
  - That is all there is to it!

Confounding through the lens of causal calculus

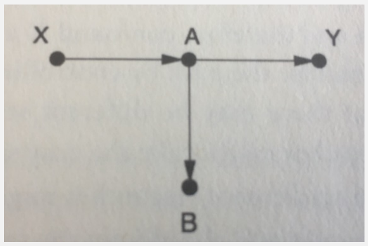

Source: The book of Why, Pearl and Mackenzie

What do we need to control for in order to de-confound X and Y?

- Nothing!
  - There is no backdoor path into X
  - A, B are descendants of X (and hence should not be controlled for)

However,

- B passes a classical epidemiological definition of confounding
- But if we control for B, we introduce confounding rather than eliminating it!

Confounding through the lens of causal calculus

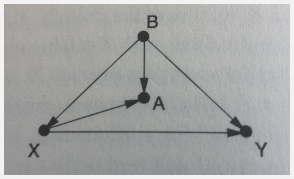

Source: The book of Why, Pearl and Mackenzie

What do we need to control for in order to de-confound X and Y?

- There is a backdoor path X ←B →Y
- We can block it only by blocking B
- If B is observable, we are all set
- If B is unobservable
    - We cannot control for it, so there is no way we can de-confound X and Y, so there is no way to estimate the causal effect of X on Y without running a RCT
    - Current statistical practice would advocate controlling for A, a proxy of B – but this only partially eliminates the confounding bias and introduces a collider biasl

Confounding through the lens of causal calculus

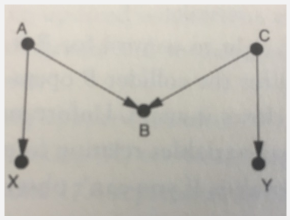

Source: The book of Why, Pearl and Mackenzie

What do we need to control for in order to de-confound X and Y?

- There is a backdoor path X ←A →B ←C → Y which is already blocked by B
- Some of the correlation based statistical definitions of confounding would identify B as a confounder!
- B becomes a confounder when we control for it!
- Example
  - B – Seatbelt use, X – Smoking, A – Attitude towards societal norms, C – Attitude towards safety and health related measures, Y – lung cancer
  - A 2006 study found B to be correlated with both X and Y

Confounding through the lens of causal calculus

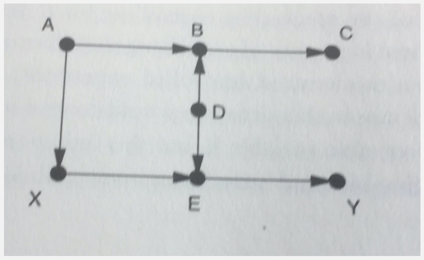

Source: The book of Why, Pearl and Mackenzie

What do we need to control for in order to de-confound X and Y?

- A, B, C, D are pre-treatment variables, X is the treatment
- The only backdoor path X ←A → B ←D→ E→Y is already blocked by the collider B, so no need to control for anything!
- Standard statistical practice would be to control for B and C
- Reinforced by Rubin (2009): "To avoid conditioning on some observed covariates … is non-scientific ad hockery"
- Controlling for B and C introduces confounding (unless we control for A or D as well)

Confounding in the language of causal calculus

- Confounder is any factor that makes $P(Y \mid X) \neq P(Y \mid do(X))$
- To de-confound two variables X and Y
  - We need to block all non-causal paths between X and Y without perturbing any causal paths
    - A backdoor path is any path from X to Y that starts with an arrow pointing into X
    - X and Y will be de-confounded if we block every such backdoor path
    - If we do this by controlling for some variables Z, we need to make sure that no member of Z is a descendent of X on a causal path
  - That is all there is to it!

Confounding through the lens of causal calculus

- Wrong! Rubin (2009): "To avoid conditioning on some observed covariates … is non-scientific ad hockery"
- Wrong! A major 2007 paper in the Journal of the American Medical Association advises investigators to condition on variables that are predictive of treatment assignment without regard to whether they are predictive of outcome.

Confounding and causal models

- If we can identify and measure the confounders, we can control for them
- But as Pearl's work has shown, standard epidemiological and statistical criteria for identifying confounders are flawed
- Both false positive and false negative confounders can yield misleading conclusions
- Causal calculus and tools based on graph theoretic criteria like d-separation provide effective methods for identifying the confounders (and only the confounders)

Review exercises – Back door criterion

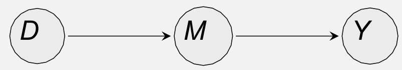

- Which set of variables in this graph satisfy the BDC with respect to the causal effect of D on Y?
- The empty set!
- $E[Y \mid do(D = 1)] - E[Y \mid do(D = 0) = E[Y \mid D = 1] - E[Y \mid D = 0]$ (correlation is causation)
  - No paths into D - just like we intervened on it
  - But you may have learned in statistics that…
- "M correlates with D and Y, so you need to control for it. Otherwise, you have omitted-variable bias"
- Bad idea: Conditional on M, D and Y are *d*-separated!
- Montgomery et al. 2018 AJPS estimate that 50% of political science studies suffer from this problem (of controlling for post-treatment variables)

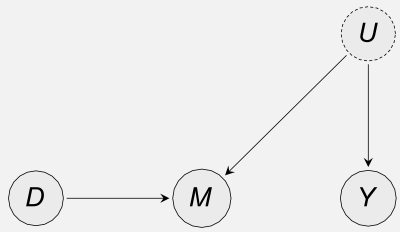

- Which set of variables in this graph satisfy the BDC wrt effect of D on Y?
- The empty set - no controls necessary
- $E[Y|do(D = 1)] - E[Y|do(D = 0)]$
  $= E[Y|D = 1] - E[Y|D = 0]$.
- What is $E[Y|D]$?
- $E[Y|D] = E[Y]$ by d-separation.
- Correct estimator equals
  - $E[Y] - E[Y] = 0$. Which is also clear from the graph.

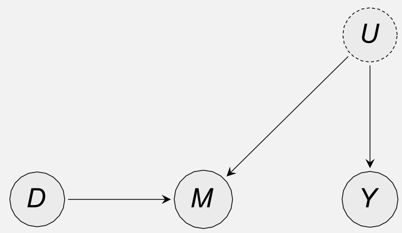

- "M correlates with D and Y. You may have learned in stats that you need to control for M; or else suffer from omitted variable bias"
- Bad idea: Conditional on M, D and Y are d-connected! Collider!
- $E[Y|D = 1, M = m] \neq E[Y|D = 1]$

Post-Treatment Variables: General Case

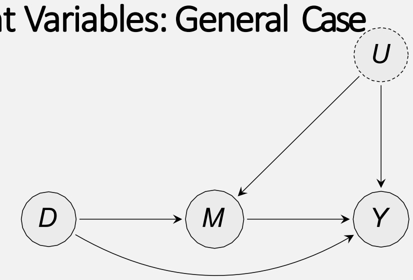

- This graph applies to situations where there are no back-door paths into *D*. Perhaps via randomization, or you block them by conditioning on *X* (not shown).
- Conditioning on *M* is forbidden by the BDC and will have two consequences:
  - You block a causal path, which you do not want
  - You open up a non-causal path, which you do not want
- This introduces bias, and it can go in any direction

Post-Treatment Variables: Remarks

- Although it is intuitively clear using causal graphs, the fact that conditioning on the descendants of the treatment may actually introduce bias is not well-known
- Usually not mentioned in textbooks that do not use causal graphs
- Even if mentioned, not really explained (see for example “Mostly Harmless Econometrics”, section on “Bad Control”)
- What is somewhat better known is “selection bias” is also often related to post-treatment variables

Why does a music record get into the *Rolling Stone*?

- Schmutz 2005: Why do music records get into the *Rolling Stone 500 Greatest Albums of All Time*?
- Compare these 500 records to 1,200 additional, also successful records (e.g. no. 1 records)
- Result: Strong *negative* association between being a no. 1 record and being included in *Rolling Stone* list
- Maybe *Rolling Stone* journalists are snobby and disregard commercial success?
- Or it's **selection bias**...

Source: Elwert and Winship 2014, "Endogenous Selection Bias"

Why does a music record get into the *Rolling Stone*?

- If Schmutz 2005 had sampled randomly, inclusion in data $S = 1$ would have only been driven by a coin flip
- Sample is conditional on $S = 1$, but S is *d*-separated from everything, so could be ignored

Why does a music record get into the *Rolling Stone*?

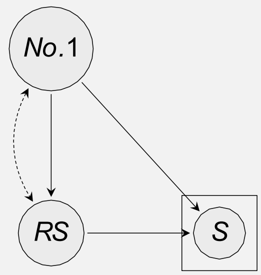

- Schmutz 2005 did not sample from general population; included units depending on *No.*1 and *RS* variable
- Opens *No.*1 → *S* ← *RS* path
- Creates non-causal association between *No.*1 and *RS* even without confounders

Motherhood and wages

- Can you study the causal relationship between motherhood and wages offered by employers?
  - Selection bias (wages are reported only for employed women)
  - Mothers' choice to be employed may be influenced by the wages offered

Selection Bias

- Similar problems may occur whenever units are sampled based on some success (or failure) measure
- This is essentially what every business school's "case studies" do
- If we are interested in the causal effect of some factor on "success", sample everyone, not only the successful
- However, sometimes sample selection is hard to avoid (e.g. motherhood-wage example)
- Solutions are possible that use parametric assumptions on the structural functions (e.g. linearity) or distributional assumptions on the errors (e.g. normality)
  - Work of James Heckman (Nobel-laureate in Economics)

Recap of Terminology

- Causal graphs are our **assumptions**
- Sometimes, they have **testable implications**, via d-separation of variables
- We observe $P(Y, X, D)$ ("observables"); so we also observe
  - $P(X)$ and $E[Y|X]$ etc.
- Unless we actually are in a situation where we have resources to intervene, we don't observe $E[Y|do(D)]$
- The process of getting from $E[Y|do(D)]$ to something like $\sum_x E[Y|D, X = x]P(X = x)$ using our assumptions is called **identification**

Plan

- Covariate-Specific Effects: Definition and Identification
- Contrast covariate-specific effects with multiple interventions; identification in easy case
- Multiple Interventions: Complicated case, identification under linearity assumption
- To find solution to nonparametric version, introduce "do-Calculus"

Covariate-specific Effects

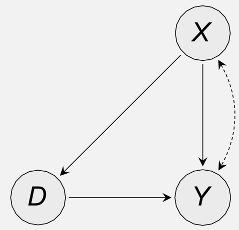

- What is $E[Y|do(D = 1), X = x]$?
- It's the effect of setting $D = 1$ for those units with $X = x$
- Covariate-specific effect
- **Effect heterogeneity**:
  - $E[Y|do(D = 1), X = x] - E[Y|do(D = 0), X = x]$ may differ for different x! In fact, almost always will (X "moderates" effect of D on Y)

Covariate-specific Effects: Examples

- Messages D, socio-economic characteristics X, turnout Y (Imai/Strauss 2011)
  - Limited budget for messages D, which people (X) should you target as to maximize turnout?
- X ethnic heterogeneity in a village, D size of vote district, Y electoral result (candidate with extreme preferences, educated candidate)
- Beath et al. 2016:
  - When vote districts are small you elect an extremist who bargains hard for your ethnically homogeneous borough
  - If the voting districts are large, you tend to elect a candidate that represents the preferences of the electorate at large

Covariate-specific Effects: Identification

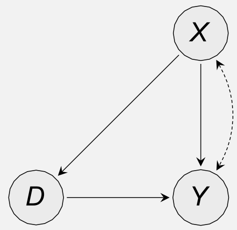

- When we used the BDC, we first wrote $E[Y|do(D=1)] =$

$$\sum_x E[Y|do(D=1), X=x]P(X=x|do(D=1))$$

- What were the next two steps?
  - $P(x=x|do(D=1)) = P(X=x)$ because if X fulfills BDC, it contains no descendants of D
  - $E[Y|do(D=1), X=x] = E[Y|D=1, X=x]$:
  - Conditional on X, doing D is like observing D

Covariate-specific Effects: Identification

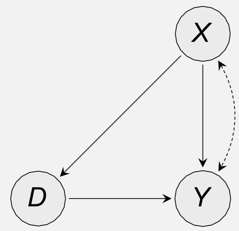

- So we have already proven that
  - $E[Y|do(D = 1), X = x] = E[Y|D = 1, X = x]$ if X fulfills BDC
- More general: X-specific effect identified if some set (X, Z) fulfills BDC (e.g. if X alone does not).
- So for X -specific effect, you always condition on X , don't average over X

Covariate-specific Effects vs. Causal Interactions

- $E[Y|do(D=1), X=x] - E[Y|do(D=0), X=x]$ will be usually different from
  $E[Y|do(D=1), do(X=x)] - E[Y|do(D=0), do(X=x))]$
- Just like $E[Y|do(D=1)]$ will be usually different from
  $E[Y|D=1]$
- "Doing" two or more variables: "multiple interventions"
- If $E[Y|do(D=1), do(X=x)] - E[Y|do(D=0), do(X=x))]$
  varies for different x, then D and X "causally" interact
- Sending messages to low-income people will affect their turnout differently than sending messages and **increasing** their income!
- The distinction between covariate-specific effects/effect heterogeneity and causal interaction gets totally lost in traditional statistics

The Back-Door Criterion for Multiple Interventions

- Given an ordered pair of **sets** of variables (D, Y) in a DAG G, a set of variables X satisfies the backdoor criterion relative to (D, Y) if
  - no node in X is a descendant of D, and
  - X blocks every path between D and Y in $G_{\underline{D}}$
- D is a set, so $D = (D_1, D_2 \ldots)$
- Otherwise, nothing changes!

Multiple Interventions: Where the BDC fails

- Let's first analyze this graph:
  - D_1 affects X , D_2 affects Y
  - X affects D_2 and Y
  - D_2 affects Y
- Example:
  - Medical treatment D_1 at time 1.
  - X health status after a while.
  - If healthy, stop treatment. If not, give treatment D_2 at time 2.
  - Then check health Y again

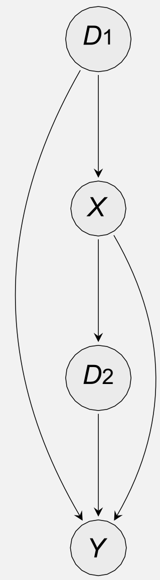

Source: Biostatistics, James Robins

Time-Varying Treatments: Problem with BDC

- Question: What's the effect of a given "strategy" (D_1, D_2) on Y?
- Put differently, can we identify the joint direct effect of (D_1, D_2)?
- Which variables do we need to adjust for?

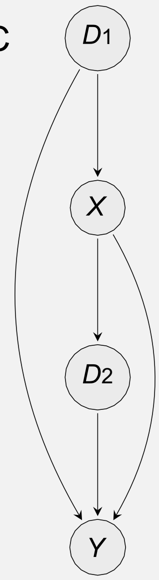

Time-Varying Treatments: Problem with BDC

- To get at the effect on Y of interventions D_1 and D_2:
  - You need to adjust for X, which is a common cause of D_2 and Y
- But X is also descendant of D_1 in G, so
  - X does not fulfill BDC!
  - Conditioning on X would block part of the effect of D_1 we are actually interested in!
    - "Post-treatment confounder"
- Is there a way around this?
- Not if all we have is BDC

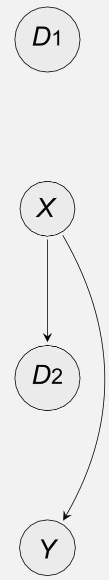

Time-Varying Treatment: Linear Version

- If the SCM of this graph is **linear**, which paths constitute the effect of *D*1 on *Y* if *D*2 is fixed?
  - $D_1 \rightarrow Y$ and $D_1 \rightarrow X \rightarrow Y$

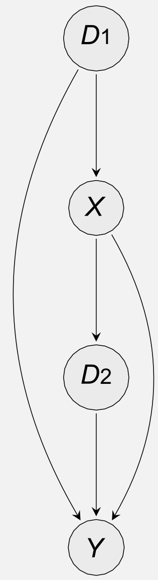

Time-Varying Treatment: Linear Version

- Write down the formula for the effect of $D1$ on Y if $D2$ is fixed using this graph
  - $\beta_{D_1Y} + \beta_{XY} \cdot \beta_{D_1X}$
- How would you estimate β_{D_1Y} and β_{XY}?
  - Regression of Y on D_1, X and D_2
- And β_{D_1X}?
  - Regression of X on D_1
- Then just multiply estimates of β_{XY} and β_{D_1X} and add estimate of β_{D_1Y}
- This is an example of **two-stage estimation**

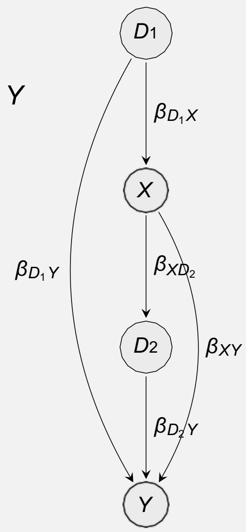

Do-Calculus

Do-calculus is for causal inference what Newton's laws of motion are for classical physics

Structural Causal Models: The Story So Far

- Causal conclusions require causal assumptions
- Structural causal models encode causal assumptions
- Causal assumptions have testable implications – conditional independence relations (via d-separation)
- Causal effects are defined in terms of interventions
  - Average causal effect of (binary) D on Y is given by $E[Y|do(D = 1)] - E[Y|do(D = 0)]$
  - We observe (samples from) $P(Y, X, D)$ and hence we can obtain $P(X)$ and $E[Y|X]$ etc.
  - Unless we have the resources and ability to experiment, we seldom observe $P(Y|do(D))$ and hence can't use it to obtain $E[Y|do(D)]$

Structural Causal Models: The Story So Far

- Causal effects are defined in terms of interventions
  - Average causal effect of (binary) D on Y is given by $E[Y|do(D=1)] - E[Y|do(D=0)]$
  - We observe (samples from) $P(Y, X, D)$ and hence we can obtain $P(X)$ and $E[Y|X]$ etc.
  - Unless we have the resources and ability to experiment, we seldom observe $P(Y|do(D))$ and hence can't use it to obtain $E[Y|do(D)]$
- Identification of causal effects from observational data entails reducing $E[Y|do(D)]$ to an expression that is free of $do()$, e.g., $\sum_x E[Y|D,X]P(X=x)$ using the causal assumption encoded in the causal graph
- Once such reduction is done, $E[Y|do(D)]$ can be estimated from observational data

Structural Causal Models: The Story So Far

- Identification of causal effects from observational data entails reducing $E\,[Y\;|do(D)]$ to an expression that is free of $do(\,)$, e.g., $\sum_x E[Y|D,X]P(X=x)$ using the causal assumptions encoded in the causal graph
- Once such reduction is done, $E\,[Y\;|do(D)]$ can be estimated from observational data
- In some cases, such identification is trivial. In other cases, it is not
- Primary challenge: observed or unobserved confounders
  - If we know the confounders X, and they are observed, we can adjust for them $E\,[Y\;|do(D)] = \sum_x E[Y|D,X]P(X=x)$
  - How do we know which confounders to adjust for?
  - Confounders are precisely those variables which make $P(Y|do(D)\;\neq P(Y|D)$

- Unless we actually are in a situation where we have resources to intervene, we don't observe $E[Y|do(D)]$

Structural Causal Models: The Story So Far

- If we know the confounders X, and they are observed, we can adjust for them $E[Y \mid do(D)] = \sum_x E[Y|D,X]P(X=x)$
- How do we know which confounders to adjust for?
- Confounders are precisely those variables which make $P(Y|do(D) \neq P(Y|D)$
- Backdoor criterion allows us to identify the confounders
- A path that starts with an arrow into D is called a **back-door path**
    - Blocking back-door paths makes sure we block bad, i.e., non-causal, paths
    - Not conditioning on descendants of D ensures that we leave all good, i.e., causal, paths open and that we do not open up new bad paths

Structural Causal Models: The Story So Far

- Backdoor criterion (BDC) allows us to identify the confounders
- BDC implies
  - Parents of treatment criterion
  - Parents of outcome criterion
  - Joint ancestors (of treatment and outcome)
- Is BDC powerful enough to identify all causal effects that are identifiable from any causal graph?
  - No!
  - More formally, BDC is sound, but not complete for identifiability of causal effects from causal graphs
- Is there a general algorithm that we can use to identify any causal effect that is identifiable from a causal graph?

Identifying causal effects

Surgeon General (1964):

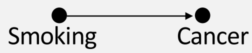

$P(c \mid do(s)) \approx P(c \mid s)$

Tobacco Industry (and Ron Fisher)

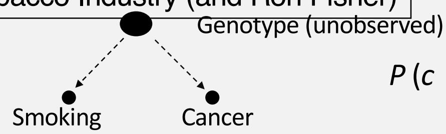

$P(c \mid do(s)) = P(c)$

Combined:

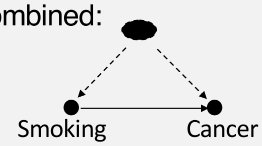

$P(c \mid do(s))$ - not identifiable

The Front-Door Criterion (FDC)

Suppose we assume:

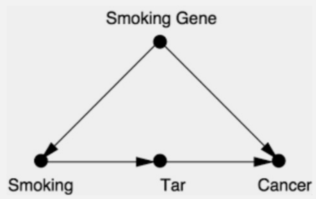

Causal assumptions:

- Smoking gene (G) causes both smoking (S) and Lung Cancer (C)
- Smoking (S) causes lung Tar (T)
- Tar (T) causes Lung Cancer (C)
- Causal effect of interest is

$$E[C|do(S = 1)] - E[C|do(S = 0)]$$

Suppose

- We have collected observational data on S, T, C for a set of individuals
- We cannot collect data for G because we do not know if a smoking gene exists

The Front-Door Criterion (FDC)

Suppose we assume:

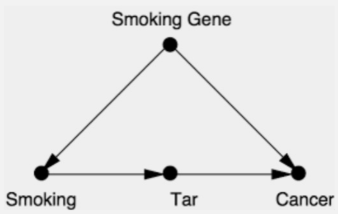

Causal assumptions:

- Smoking gene (G) causes both smoking (S) and Lung Cancer (C)
- Smoking (S) causes lung Tar (T)
- Tar (T) causes Lung Cancer (C)
- Causal effect of interest is $E[C \mid do(S = 1)] - E[C \mid do(S = 0)]$

- Are there any confounders?
- Yes. G is the only confounder
- Are there any variables that satisfy the back-door criterion with respect to (S, C) ?

The Front-Door Criterion (FDC)

Suppose we assume:

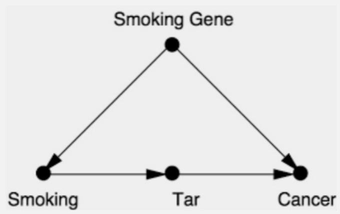

Causal assumptions:

- Smoking gene (G) causes both smoking (S) and Lung Cancer (C)
- Smoking (S) causes lung Tar (T)
- Tar (T) causes Lung Cancer (C)
- Causal effect of interest is

$E[C|do(S = 1)] - E[C|do(S = 0)]$

- Are there any variables that satisfy the back-door criterion with respect to (S, C)?
- S, T and C are not candidates
- What about G?
- G would satisfy the backdoor criterion if it were observed!
- But it is not!

The Front-Door Criterion (FDC)

Suppose we assume:

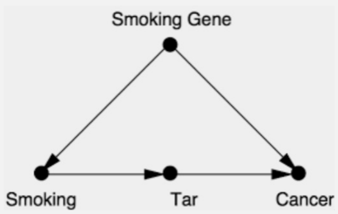

Causal assumptions:

- Smoking gene (G) causes both smoking (S) and Lung Cancer (C)
- Smoking (S) causes lung Tar (T)
- Tar (T) causes Lung Cancer (C)
- Causal effect of interest is

$$E[C \mid do(S = 1)] - E[C \mid do(S = 0)]$$

- We cannot use the back door to adjust for G
- Is there another way?

The Front-Door Criterion (FDC)

Suppose we assume: $E[C\,|do(S = 1)] - E[C\,|do(S = 0)]$

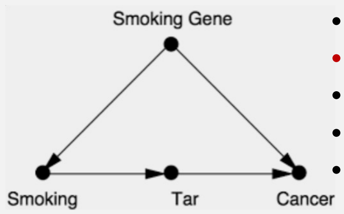

- We cannot use BDC to adjust for G
- Is there another way?
- We can get the causal effect of S on T
- We can get the causal effect of T or C
- And combine them to get the causal effect of S on C

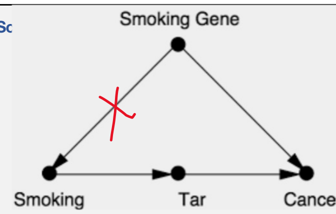

The Front-Door Criterion (FDC)

$E[C\,|do(S\ =\ 1)] - E[C\,|do(S\ =\ 0)]$

- We can get the causal effect of S on T
- Why?
- When we condition on S, There is no unblocked backdoor path from S to C because $S \leftarrow G \rightarrow C \leftarrow T$ is already blocked by the collider C
- We can observe $P(T|S=1) - P(T|S=0)$ to get the causal effect of S on T

The Front-Door Criterion (FDC)

$E[C\,|do(S\ =\ 1)]\ -\ E[C\,|do(S\ =\ 0)]$

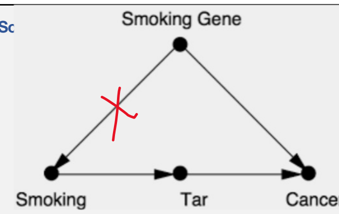

- We can get the causal effect of T on C
- How?
- We can block the backdoor path into T which is $T \leftarrow S \leftarrow$ G $\rightarrow$ C by adjusting for S
- We can get $P(C|do(T = 1))\ -\ P(S|do(T = 0))$ using the backdoor adjustment formula

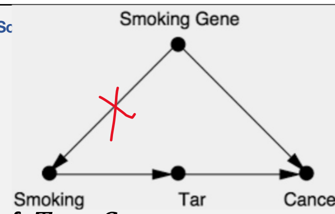

The Front-Door Criterion (FDC)

$E[C\,|do(S\ =\ 1)]\ -\ E[C\,|do(S\ =\ 0)]$

- We have the causal effect of S on T and of T on C
- Can we use these to get the causal effect of S on C?
- Cancer can come about in 2 ways: $T = 1$ or $T = 0$
- If we $do(S = 1)$, the probabilities of these states are $P(T = 1|do(S = 1))$ and $P(T = 0|do(S = 1))$
- If we $do(S = 0)$, they are $P(T = 1|do(S = 0))$ and $P(T = 0|do(S = 0))$
- If $T = 0$, the probability of cancer is $P(C|\ T = 0)$
- If $T = 1$, the probability of cancer is $P(C|\ T = 1)$
- We can compute $P(C|do(S))$ by weighting the two scenarios according to their respective probabilities under $do(S)$
- We can then get $E[C\,|do(S\ =\ 1)]\ -\ E[C\,|do(S\ =\ 0)]$

Center for Artificial Intelligence Foundations and Sc
Artificial Intelligence Research Laboratory

The Front-Door Criterion (FDC)

$E[C\,|do(S\ =\ 1)]\ -\ E[C\,|do(S\ =\ 0)]$

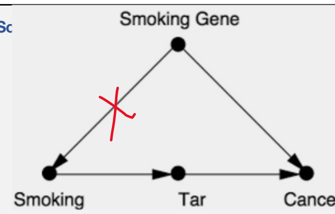

- What did we do?
- To obtain the causal effect of S on C, we adjusted for S and T which lie on the front-door path from S to C

$$P(C|do(S)) = \sum_t P(T = t, S) \sum_s P(C|S = s, T = t)P(S = s)$$

- There is do on the LHS but no do on the RHS!
- G, the unobserved confounder does not appear in the RHS
- If the causal graph shown is an accurate model of causal mechanism of cancer, the controversy about whether and to what extent smoking causes cancer could have been answered by an observational study that measured S, T, and C

Front-Door Criterion

- Glynn and Kashin (2018) studied the effect of job training services on earnings (without assuming a SCM)
- Data from RCT compared with an observational study
- They did not use a causal diagram but compared RCT with results of back-door and front-door adjustments
- Motivation is unobserved – back door criterion can't be applied, but Glynn and Kashin did with other potential confounders like age, gender, race, education
- Strictly speaking, front-door criterion can also not be applied exactly – because of the direct link from Motivation to Showed up
- Adjustment using FDC beats that using BDC in terms of agreement with RCT
- Study shows the power of FDC

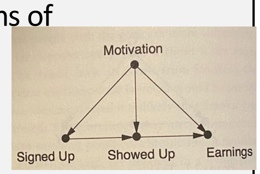

Glynn, Adam N., and Konstantin Kashin. "Front-door versus back-door adjustment with unmeasured confounding: Bias formulas for front-door and hybrid adjustments with application to a job training program." *Journal of the American Statistical Association* 113, no. 523 (2018): 1040-1049.

Causal Interactions: Identification

- In this graph, is
  $E[Y \mid do(D=1), do(M=m)] - E[Y \mid do(D=0), do(M=m)]$
  identified? Let's check $G_{\underline{DM}}$
  - Deleting front-door arrows from D eliminates back-door path $M \leftarrow D \rightarrow Y$!

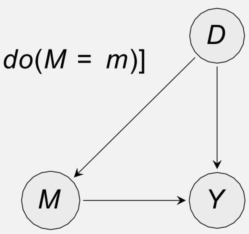

By our usual reasoning, therefore

$$E[Y \mid do(D=1), do(M=m)] - E[Y \mid do(D=0), do(M=m)]$$
$$= E[Y \mid D=1, M=m] - E[Y \mid D=0, M=m]$$

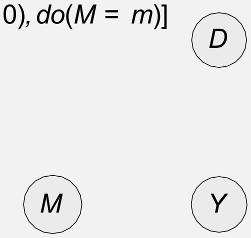

Causal Interactions: Identification

- In this simple case,
  - $E[Y|do(D = 1), do(M = m)] - E[Y|do(D = 0), do(M = m)]$
    $= E[Y|D = 1, M = m] - E[Y|D = 0, M = m]$
- Look at average outcome difference of units with same M but different D
- **Same estimator as for M-specific effect of D**
- Adjustment for post-treatment variable M!
- In our graph, M transmits part of $do(D)$. If you $do(M)$, you kill this connection
- In such situations,
  - $E[Y|do(D = 1), do(M = m)] - E[Y|do(D = 0), do(M = m)]$ is also called the **controlled direct effect** (CDE) of D (with respect to M).
- Controlled because mediator M is intervened upon
- BDC easy to apply

Structural Causal Models: The Story So Far

- Backdoor criterion (BDC) allows us to adjust for confounders
- BDC is not powerful enough to identify all causal effects that are identifiable from any causal graph.
- Front-door criterion allows us (under some conditions when there are unobserved confounders) to identify causal effects that cannot be identified using BDC
  - BDC is sound, but not complete for identifiability of causal effects from causal graphs
  - So is FDC
- Is there a general algorithm that we can use to identify any causal effect that is identifiable from a causal graph?

The *do*-Calculus

- Are there some simple rules which you can apply to any DAG in order to check whether and how any causal effect – based covariate specific, joint, etc. - can be identified?
  - The do-calculus! (Judea Pearl)
  - Perhaps the most important body of work in causal inference
- Three rules/laws/theorems:
  - Insertion/Deletion of Observations
  - Action/Observation Exchange
  - Insertion/Deletion of Actions
- "Observation" = conditioning on variable
- "Action" = *do*-ing variable

Identifying causal effects

- Postulate a causal graph (causal assumptions)
- See how the estimand (causal effect of interest) can be written as a function of the postulated causal graph
- Check whether this function can be calculated from observations (using do calculus)

Causal identifiability

- Given a causal Bayesian network G, we say that a causal effect $P(Y|do(X))$ is identifiable when $P(Y|do(X))$ can be computed using only the joint distribution over the observable variables