Recitation 6

MLHC MIT

Hussein Mozannar

3/24/2023

Causal Inference vs Machine Learning

- In Supervised Learning (regular ML classification), we try to predict given patient features X, the most probable value of outcome Y:
  - Recall, we want to approximate E(Y|X)
- Causal Inference*, given patient X, if we give treatment T, predict what is the most probable outcome Y:
  - **Question**: why not E(Y|X,T)? i.e. cast this in supervised learning language with X'=(X,T)?

* simplified setup for this class

Causal Inference

- First of all, what is $E(Y|X,T)$?
  - In words, conditional probability distribution of Y given patient features X and treatment T
  - Assumes a data distribution over tuple (X, T, Y), i.e. a dataset where we observe patients who receive treatments and have certain outcomes
- Example: Y is if person dies (1) or not (0), T is a choice between two cancer drugs (T=1 is drug 1, T=0 is drug 0) and X is patient age (X=1 if age>35 and X=0 if age <35).
- We observe data from a hospital where patients above >35 receive treatment 1 and patients below <35 receive treatment 0

Example of Treatment Effects

- We want effect of drug 1 on patients below 35:
- -> Is this: $E(Y|X=0,T=1)$?
- Nonsensical! We never even see $(X=0,T=1)$ and we will never!
- What about treatment effect of the two drugs?
- -> Is this: $E(Y|T=1) - E(Y=1|T=0) = 1/3$
- Is drug 0 better? Or is it because drug 0 is given to young patients?
- We need a new language to talk about causal inference more than conditional probability!

X	Y	T
0	0	0
0	0	0
1	0	1
1	1	1
1	0	1

Problem Set 3: Problem 1

In this problem, we will present you with several free text scenarios. For each scenario, you must answer first whether or not causal inference is required in this scenario, and, if so, you must identify the relevant covariates (X), treatments (T), outcomes (Y), and any hidden confounders (H) that pose particular concern in this setting.

Problem Set 3: Problem 1

Example:

You notice that ice cream sales are correlated with drowning rates.

You decide to test whether ice cream sales cause drowning rates to go up.

You have a large dataset: for each city, for each month, # ice cream ordered, # hot dogs eaten, average income in city, and # drownings.

Problem Set 3: Problem 1

Answer:

- Yes this requires causal inference, because we care about causation rather than prediction.
- T = # ice cream ordered
- Y = # drownings
- X = # hot dogs, average income
- What are possible hidden confounders?
- -> Temperature of month

Causal Inference: Do-Calculus

- One formal way to write down the quantity we care about is:
- E(Y| X=0, do(T=1))

(This is equivalent in potential outcomes language to $E[Y_1|X=0]$)

- The do(T=1) implies a direct intervention where we go and set T=1 for patients below 35 and see the results.
- Similarly: E(Y| do(T=1)) – E(Y|do(T=0))

-> How do we compute E(Y|X=0, do(T=1)) ?

How to compute Treatment Effects?

- Ideally, we have a simulator of the world where we can go back in time, or simulate the future, to see the effect of the intervention.
- Sadly, we don't have a simulator, we only have a dataset of observations (X,Y,T) that we need to leverage!
- How can we (**rigorously**) compute treatment effects using only data?

Step 1: Drawing a DAG of the experiment

- Example 1: Y is if person dies (1) or not (0), T is a choice between two cancer drugs (T=1 is drug 1, T=0 is drug 0) and X is patient age (X=1 if age>35 and X=0 if age <35). Cancer drugs are given randomly to patients*.

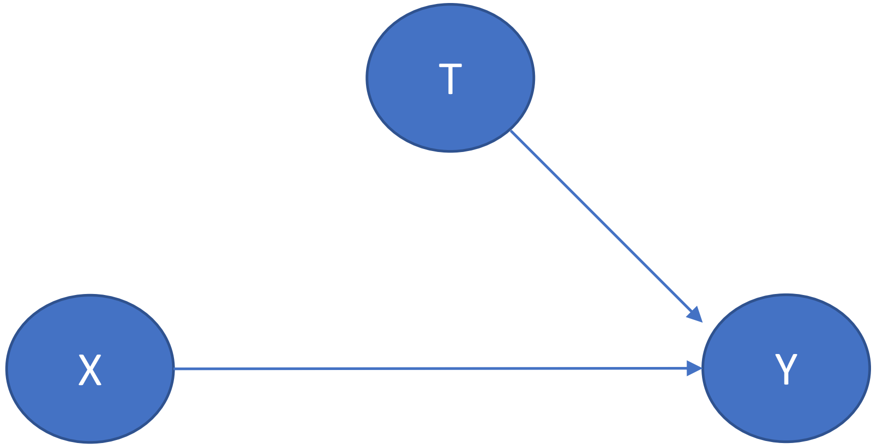

*Note that this is a Randomized Control Trial (RCT)

Step 2: Inferring relation between treatment and outcome (Problem 2)

In the last example, suppose we want to compute the average treatment effect:

$$E(Y|\ do(T=1)) - E(Y|do(T=0))$$

- From DAG we have: no factor influences both T and Y, thus if we look at our dataset (T, Y) it will be like we simulated for some people T=0 and others T=1 without changing anything else.
- Moreover, P(T=1) = 0.5, thus we observe both treatments
- Then our ATE is equivalent to:

$$E(Y|T=1) - E(Y|T=0)$$

Intuitively because there is no difference between people who receive drug 1 or drug 0. Formally this is Rule 2 of the do-calculus (an axion).

Step 1: Drawing a DAG of the experiment

- Example 2: Y is if person dies (1) or not (0), T is a choice between two cancer drugs (T=1 is drug 1, T=0 is drug 0) and X is patient age (X=1 if age>35 and X=0 if age <35). Let Z=1 if patient have arthritis. Drug 1 given to X=1, and Drug 0 given to X=0

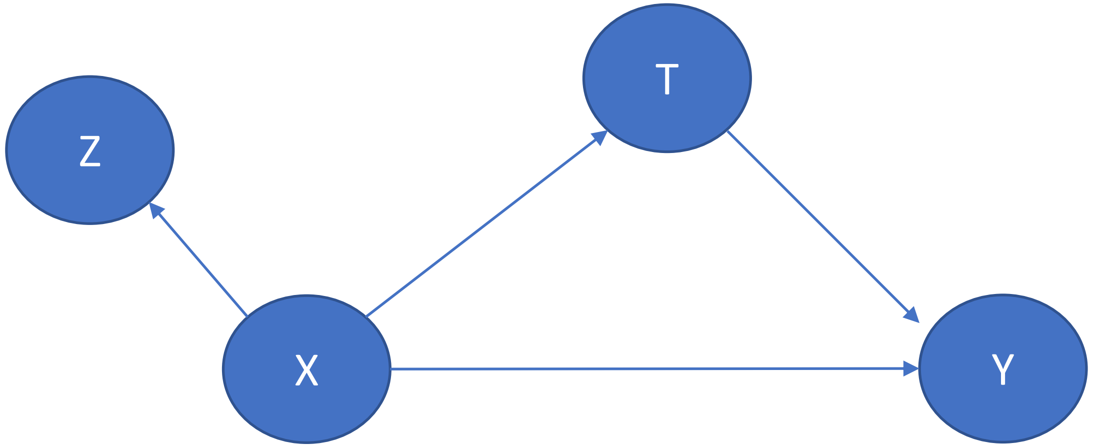

Step 2: Covariate Adjustment (Problem 3)

In the last example, suppose we want to compute the average treatment effect:

$$E(Y|\ do(T=1)) - E(Y|do(T=0)$$

- Can we do:

$$E(Y|T=1) - E(Y|T=0)$$

- No! there exists variable, namely X, that effects both T and Y!
- Solution: Covariate Adjustment!

Step 2: Covariate Adjustment

- **Backdoor criterion:** Find the set of variables S, such that no variable in S is a descendant of T, and S blocks every path between T and Y that contains an arrow into T.
- For example, all the **parents** of T satisfy the backdoor criterion.

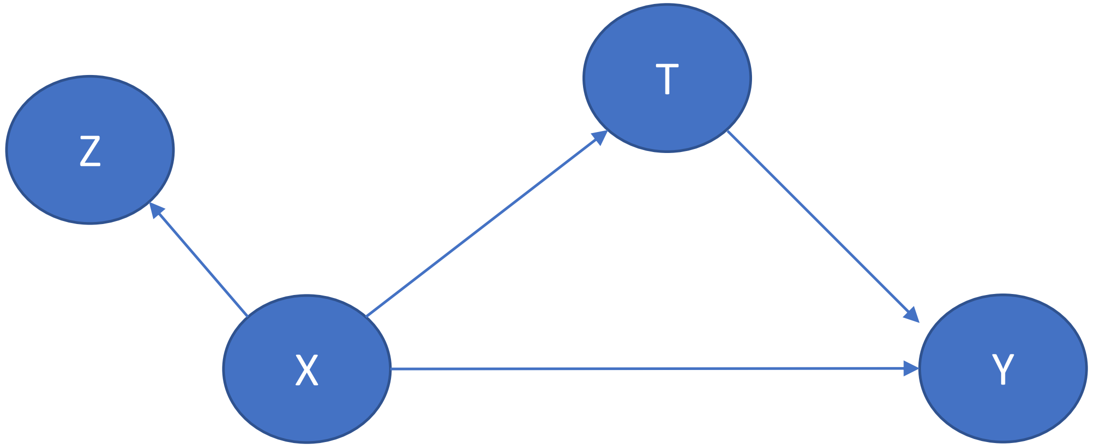

Step 2: Covariate Adjustment

- **Backdoor criterion:** Find the set of variables S, such that no variable in S is a descendant of T, and S blocks every path between T and Y that contains an arrow into T.
- For example, all the **parents** of T satisfy the backdoor criterion.

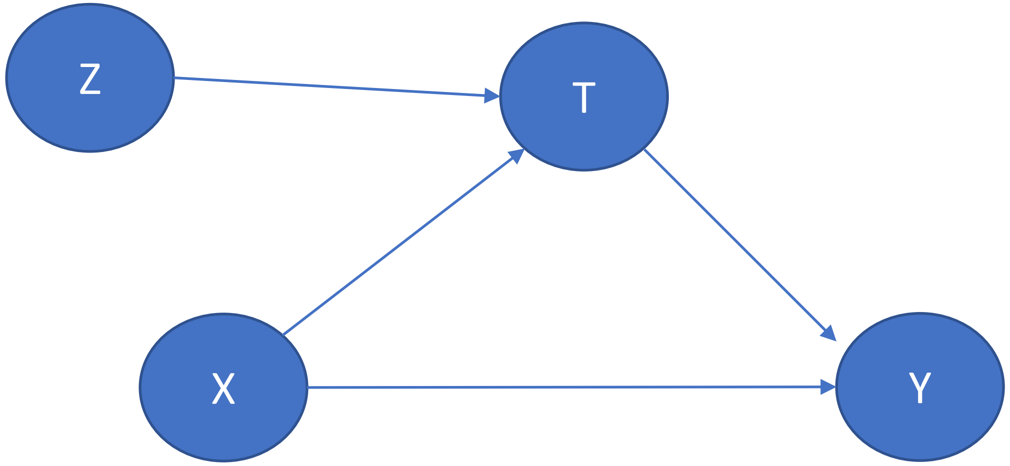

Step 2: Covariate Adjustment

- **Backdoor criterion:** Find the set of variables S, such that no variable in S is a descendant of T, and S blocks every path between T and Y that contains an arrow into T.
- For example, all the **parents** of T satisfy the backdoor criterion, or common ancestors of T and Y

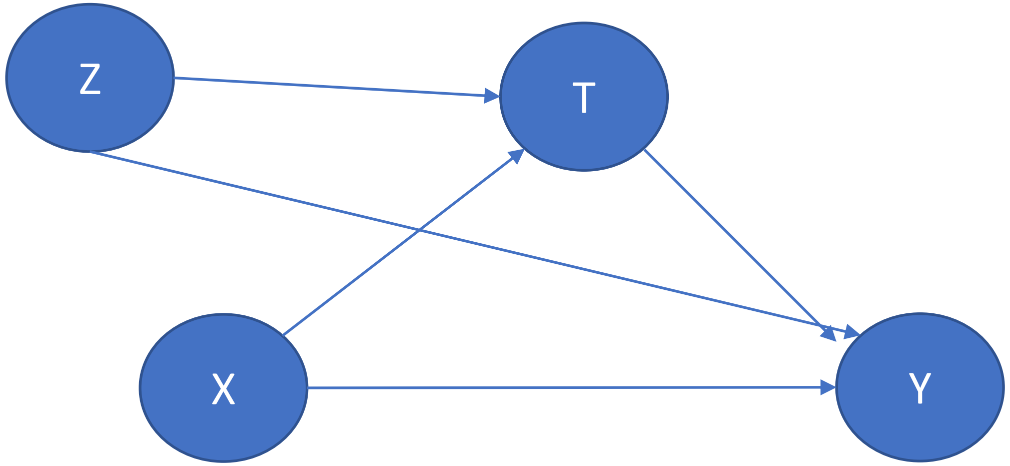

Step 2: Covariate Adjustment

- Now if we want to compute:
- E[Y|X,Z,do(T=1)] this is the same as E[Y|X,Z,T=1]
- Why? Because given (X,Z), now the assignment of T is random! There is no other factor that affects the pair (T,Y) when we condition on (X,Z)

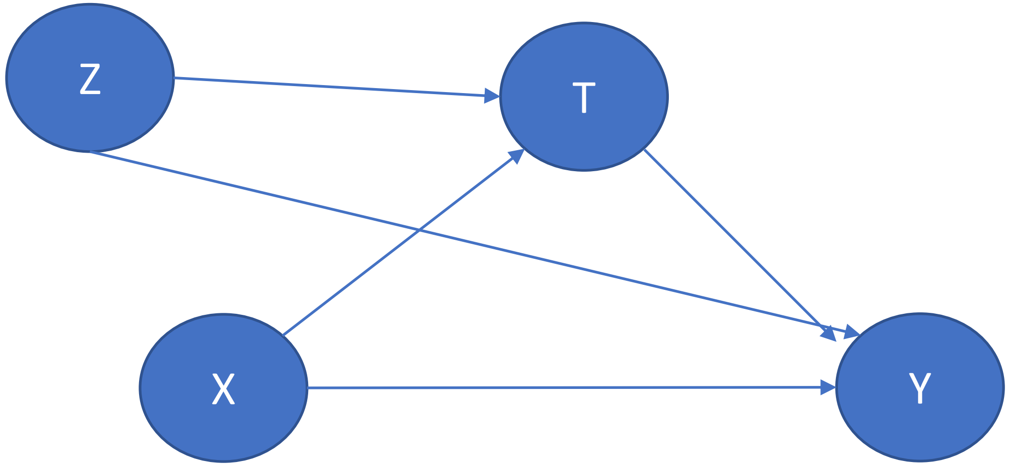

Step 2: Covariate Adjustment

Rule:

$$E[Y|do(T=1)] = \sum_s E[Y|T=1, S=s]P(S=s)$$

We also need: 0<P(T=1|S=s)<1 for all s (common support)

Step 2: Covariate Adjustment (Problem 3)

we want to compute ATE:

$$E(Y|\ do(T=1)) - E(Y|do(T=0)$$

- Condition on X,Z per last slide:

$E[Y|do(T=1)] = E_{X,Z}[\ E[Y|T=1,X,Z]]$ (iterated expectation)

Intuition behind Covariate Adjustment

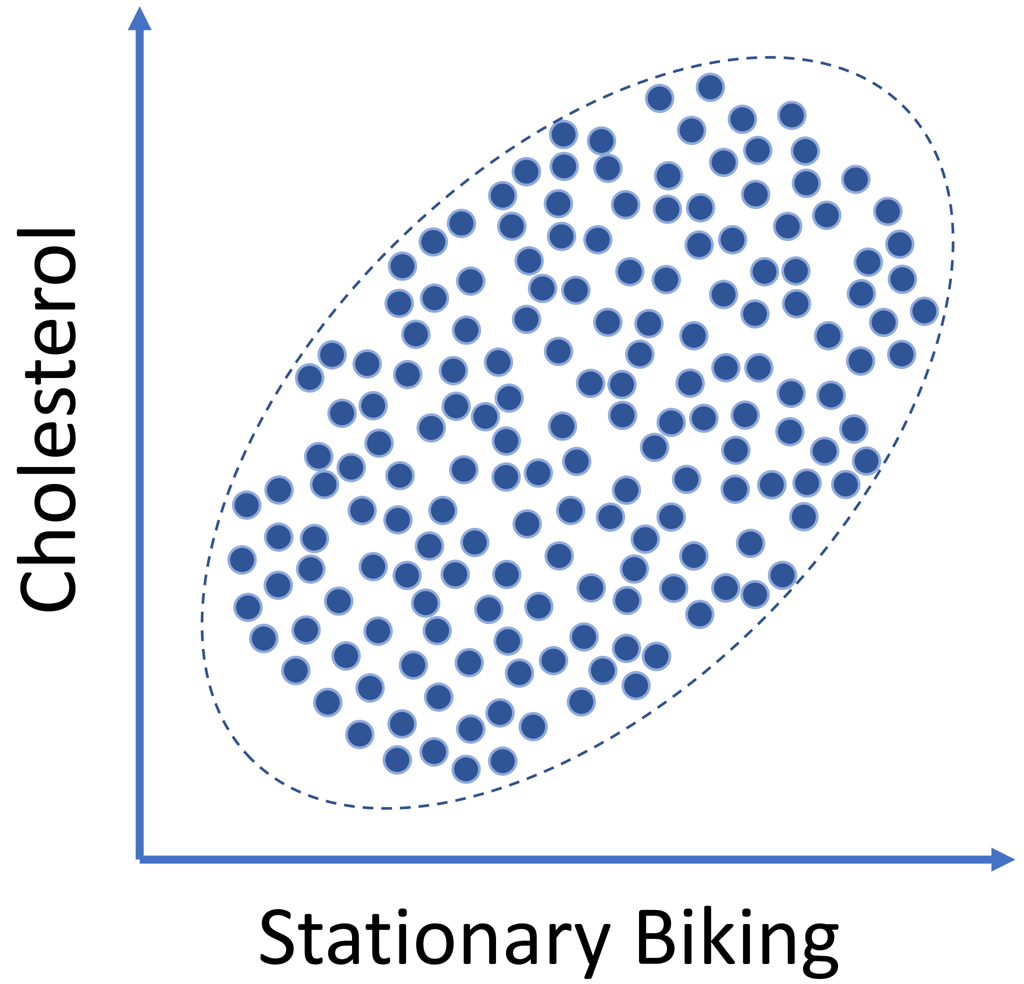

Intuition behind Covariate Adjustment

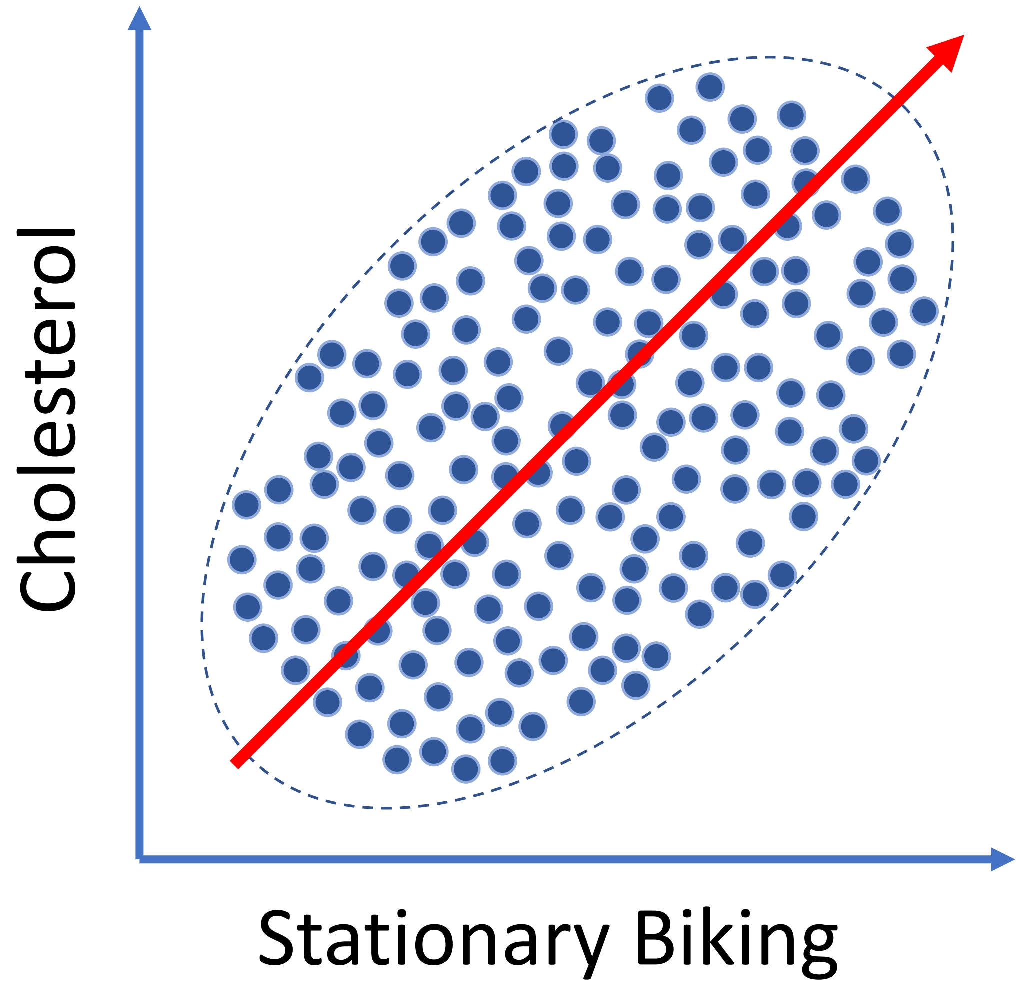

Intuition behind Covariate Adjustment

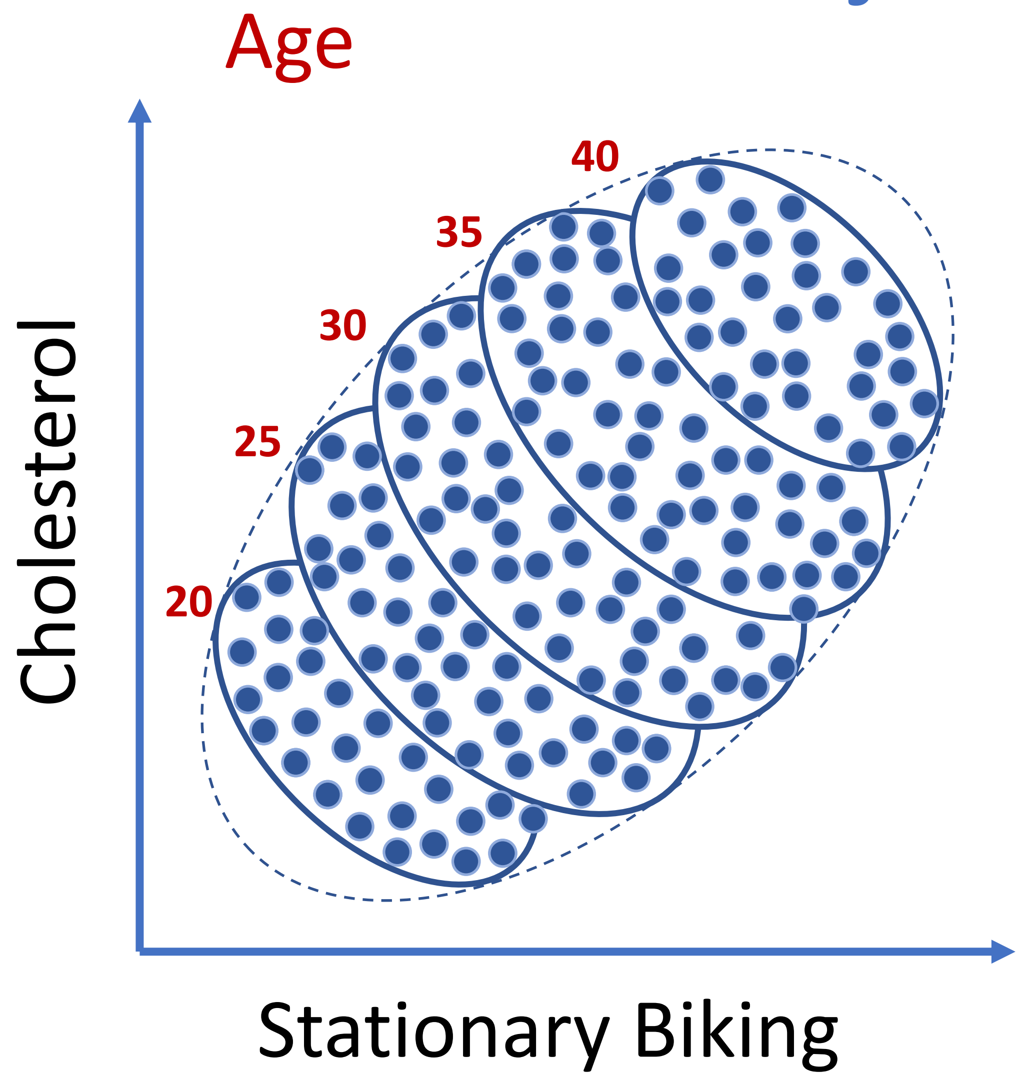

Intuition behind Covariate Adjustment

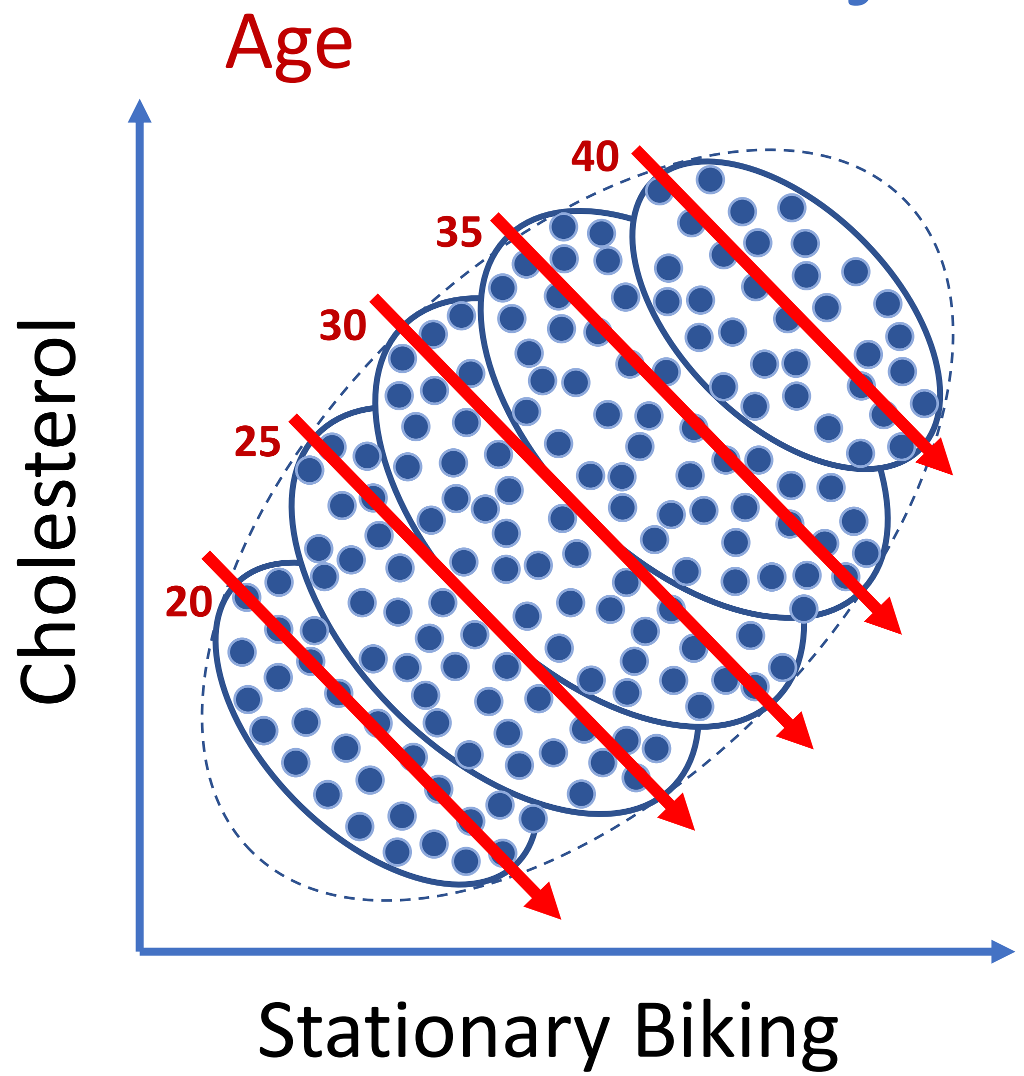

Other Methods

- Covariate Adjustment is just one method, there are many more we very briefly covered in class:
  - Matching
  - Instrumental Variables
  - Propensity Scores
  - Regression Discontinuity Design
  - Synthetic Control

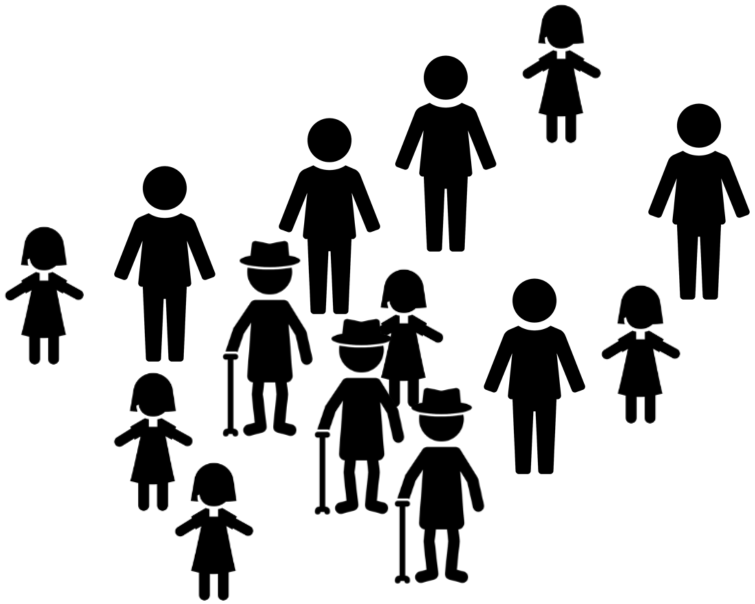

Avg Cholesterol = 200

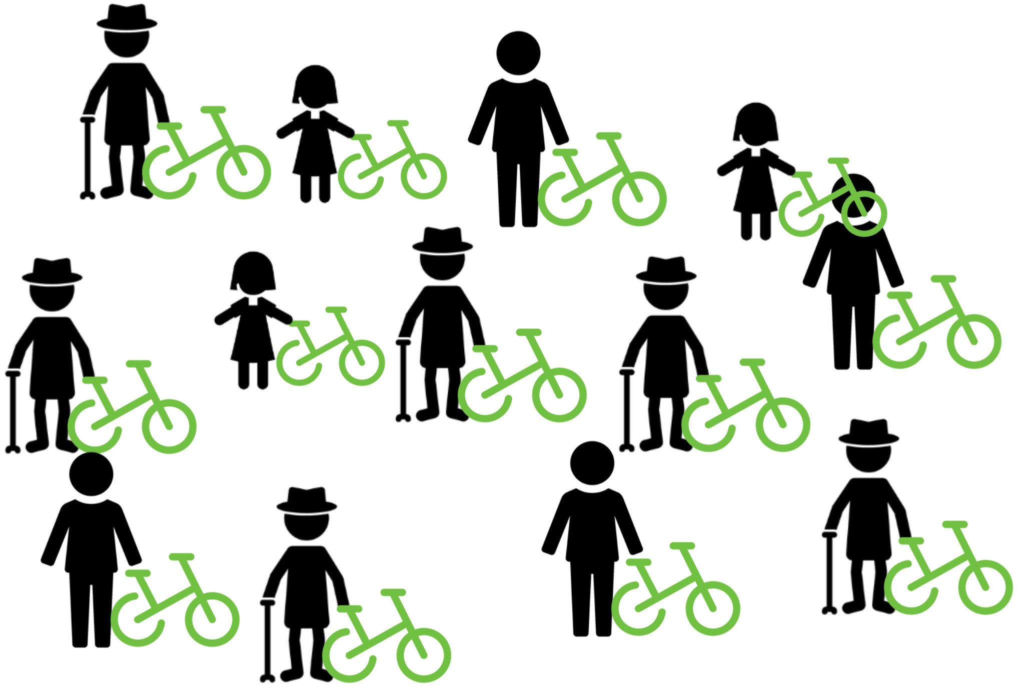

Avg Cholesterol = 206

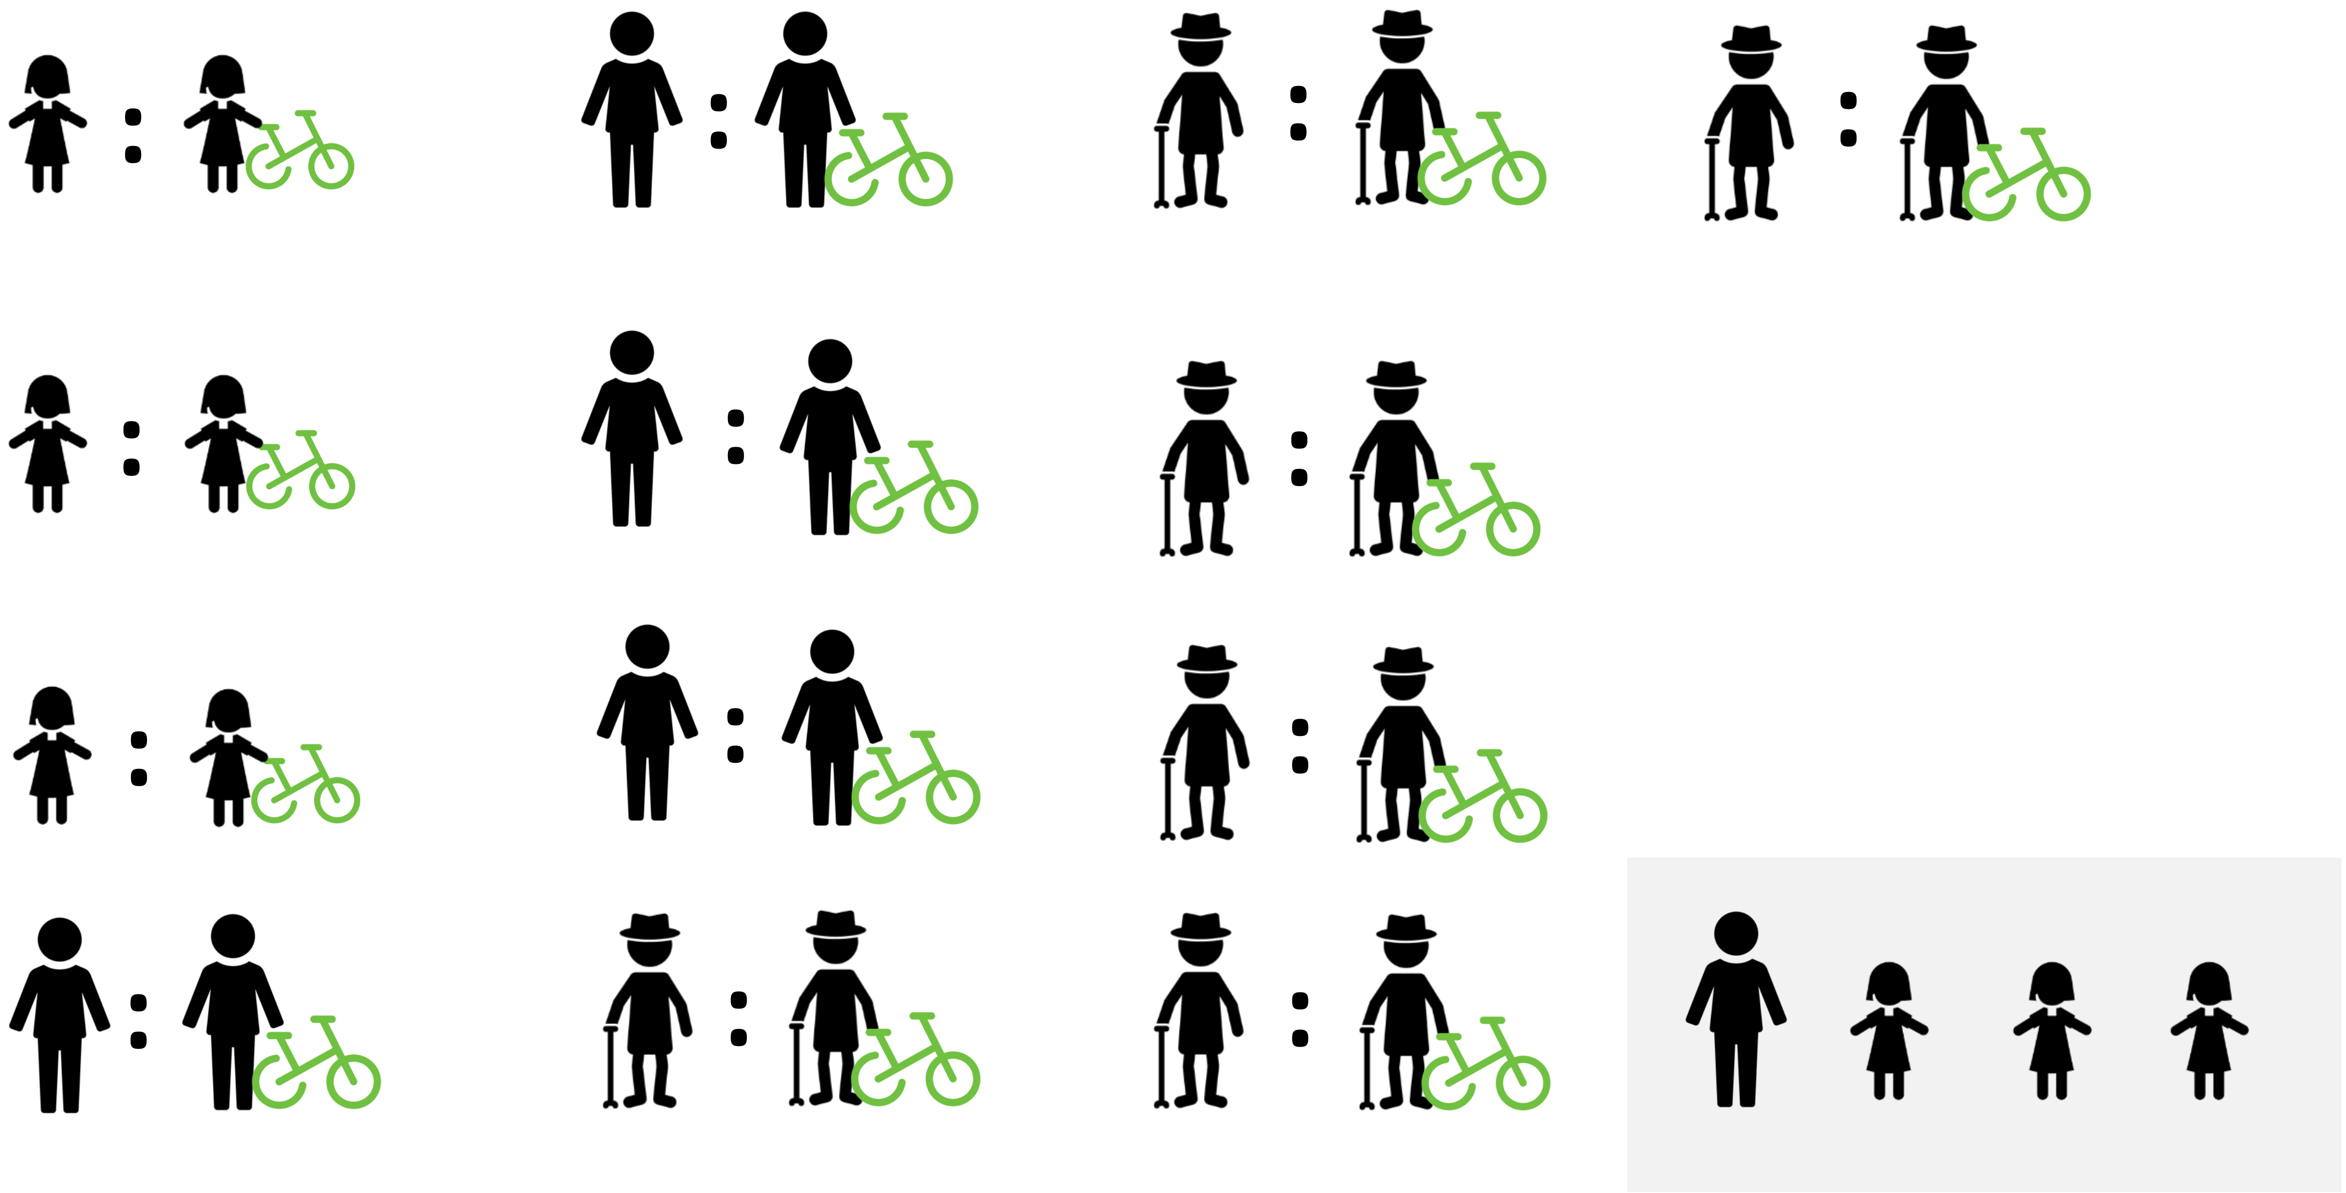

Matching

Identify pairs of treated and untreated individuals who are very similar or even identical to each other

Very similar ::= $Distance(X_i, X_j) < \epsilon$

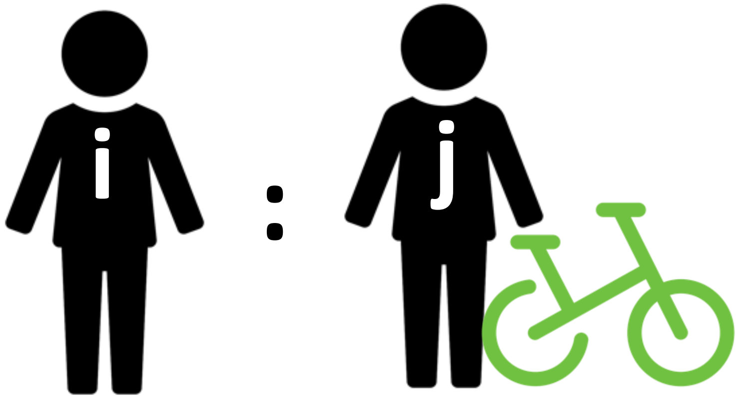

Paired individuals provide the counterfactual estimate for each other.

Average the difference in outcomes within pairs to calculate the *average-treatment-effect on the treated*

Instrumental Variables

- A variable correlated with treatment, but independent of outcome.