Machine Learning for Healthcare

6.7930, HST.956

Lecture 12: Causal Inference Part 3

David Sontag

Reminder: Causal inference

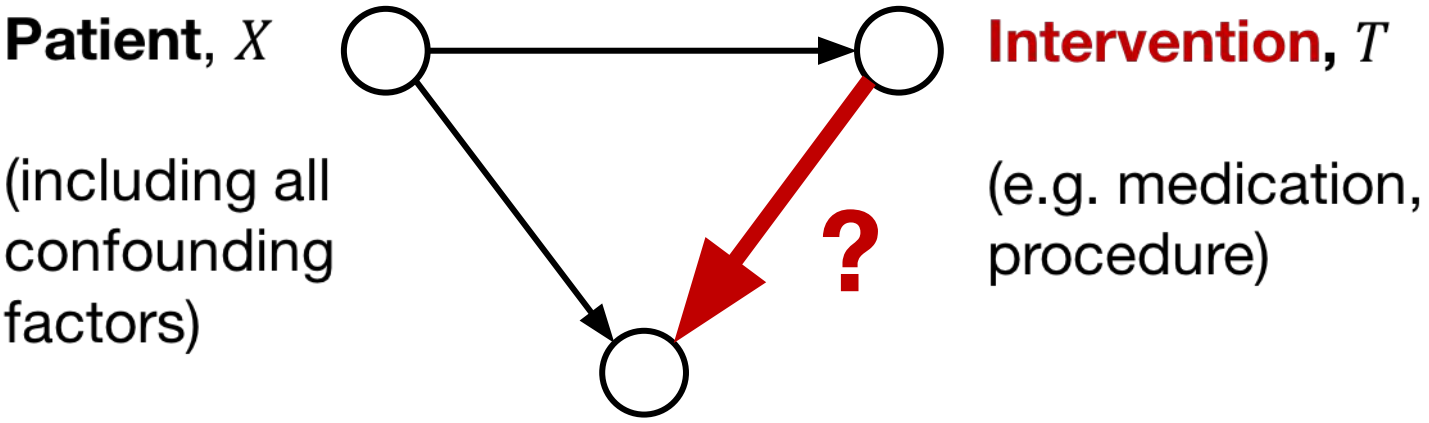

High dimensional

Observational data

Reminder: Causal inference

- Two approaches to use machine learning for causal inference
  - Predict outcome given features and treatment – i.e., E[Y | X, T] – then use to impute counterfactuals (*covariate adjustment*)
  - Predict treatment using features (*propensity score*) – i.e., Pr(T|X) – then use to reweight outcomes

Consistency of estimates depend on:

- Causal graph being correct (i.e., no unobserved confounding)
- Identifiability of causal effect (i.e., overlap or correctly specified model)

Same ideas can be used to evaluate *policies* using observational data

- Suppose someone gave us a policy $\pi(l)$ that outputs a_1 vs a_2

  Example: which antibiotic to prescribe?

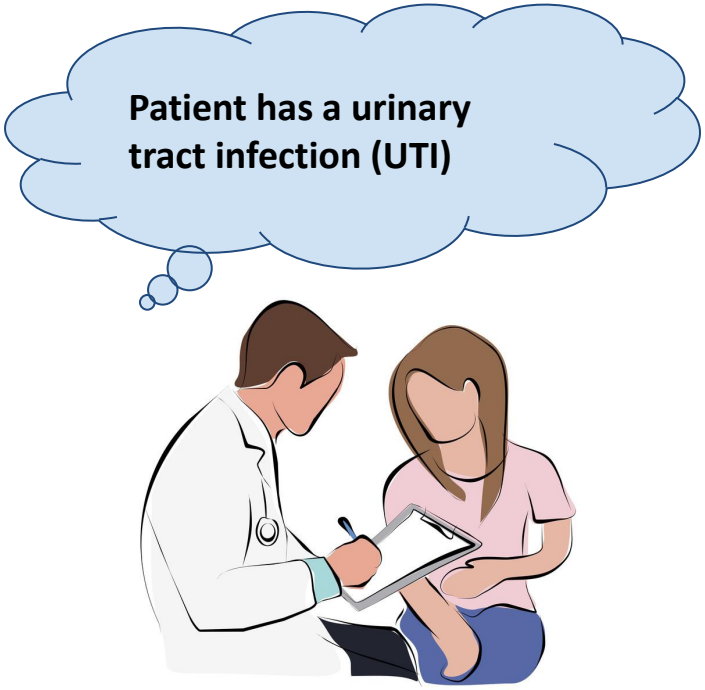

Affects 1 in 2 women during lifetime; 3rd most common cause for antibiotic treatment

[Kanjilal et al., A decision algorithm to promote outpatient antimicrobial stewardship for uncomplicated urinary tract infection. *Science Translational Medicine*, 2020.]

Same ideas can be used to evaluate *policies* using observational data

- Suppose someone gave us a policy $\pi(l)$ that outputs a_1 vs a_2

Example: which antibiotic to prescribe?

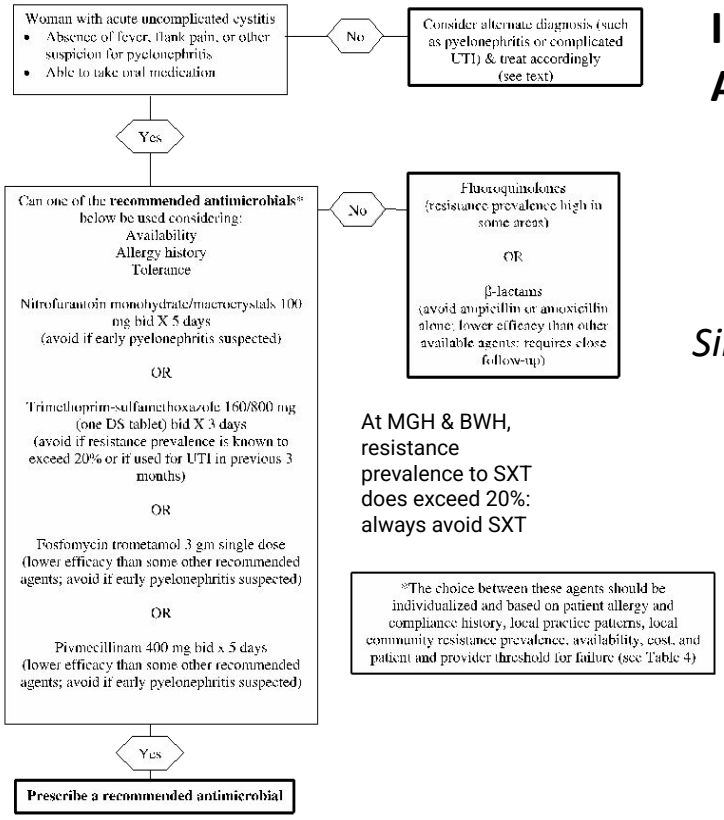

Infectious Disease Society of America (IDSA) guidelines

At MGH & BWH, resistance prevalence to SXT does exceed 20%: always avoid SXT

Simplifies to

Resistance or exposure to NIT in past 90 days?

No → Prescribe NIT (Nitrofurantoin)

Yes → Prescribe CIP (Ciprofloxacin)

[Gupta et al., *Clinical Infections Diseases*, 2011.]

Same ideas can be used to evaluate *policies* using observational data

- Suppose someone gave us a policy $\pi(l)$ that outputs a_1 vs a_2
- How do we evaluate it?
- We give two approaches, one based on potential outcomes and the other based on propensity scores
- In both cases, we have to first consider the causal graph that underlies the *observational data*

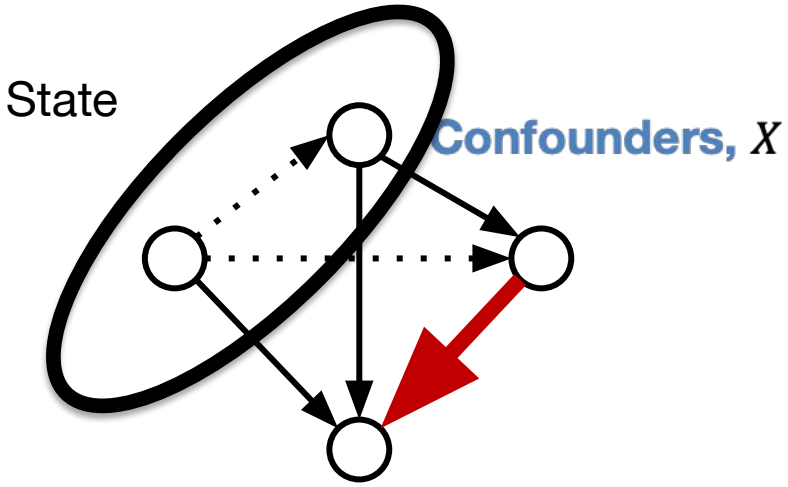

Switched notation to what's more typically used in RL
action *A*: Treatment *T*
reward *R*: Outcome *Y*

Evaluating policies using potential outcomes

- First, use machine learning to obtain a model that can predict potential outcomes (we need ignorability, overlap)

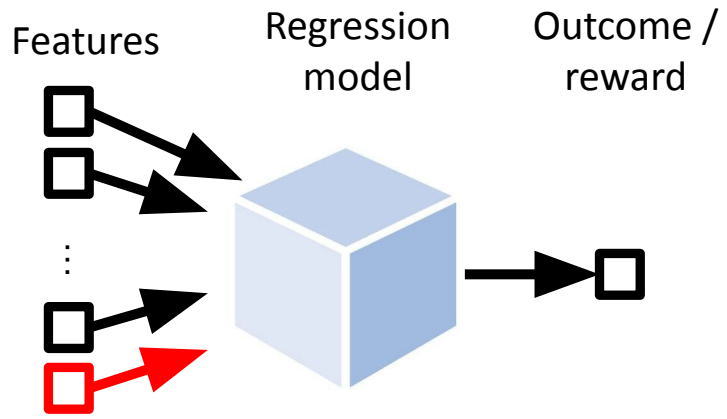

- Then, use this model to impute policy outcomes:

$$\hat{Q}(\pi) = \frac{1}{n} \sum_{i=1}^{n} f(l_i, x_i, \pi(l_i))$$

Evaluating policies using inverse propensity scores

- 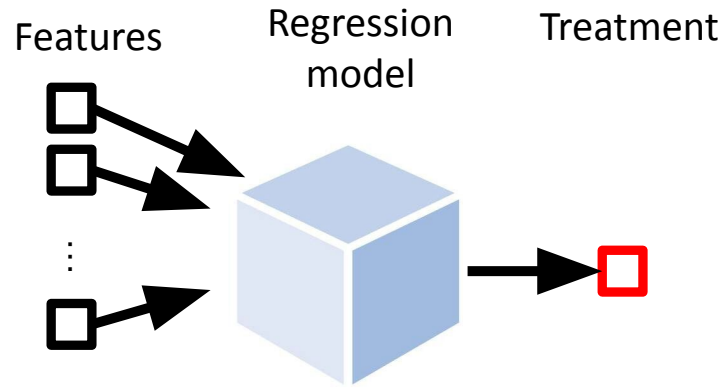

$$\hat{Q}^{IPW}(\pi) = \frac{1}{n}\sum_{i=1}^{n}\frac{1[a_i = \pi(l_i)]}{\hat{p}(a_i \mid s_i)}R_i$$

Aside: is this the right goal? What if we wanted to control worst-case reward instead of average?

Learning policies from observational data

- Consider our first estimator: $\hat{Q}(\pi) = \frac{1}{n} \sum_{i=1}^{n} f(l_i, x_i, \pi(l_i))$

- Create data set $\{(l_i, o_i)\}$ where

$$o_i = \arg\max_A f(l_i, x_i, A)$$ Notice relationship to CATE

- Use an (interpretable) ML algorithm to fit this new dataset
- The resulting policy may be a much simpler function than *f*!

(Makar, Swaminathan, Kiciman. A distillation approach to data efficient individual treatment effect estimation. AAAI, 2019)

Does gastric bypass surgery prevent onset of diabetes?

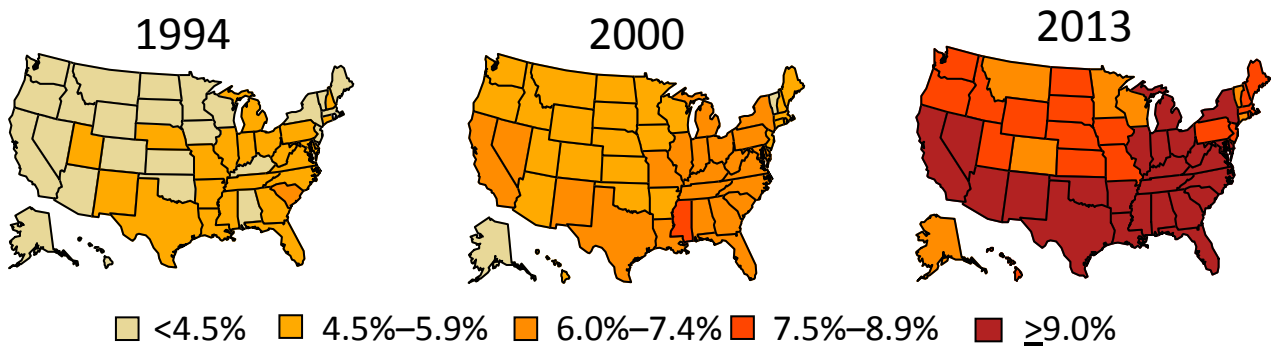

- Gastric bypass surgery is the highest negative weight (9th most predictive feature)
  - Does this mean it would be a good intervention?
- Yes, *if*….
  - Interpret ‘gastric bypass surgery’ feature as T
  - Interpret all the other features as X; assume they all include all relevant confounders and do not include anything post-treatment
  - True potential outcome function is linear

What is the likelihood this patient, with breast cancer, will survive 5 years?

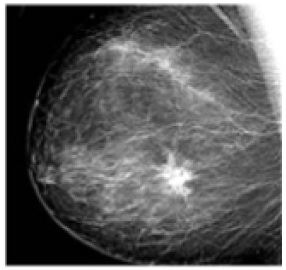

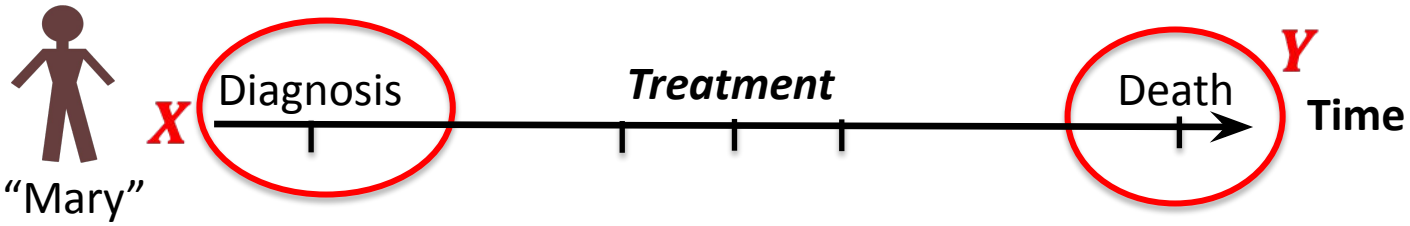

A long survival time may be because of treatment!

- Group into K categories of treatment strategies T (one of which might be "no treatment")
- Gather data on confounding factors C that might influence both treatment decision and outcome
- Learn f(X,C,T) to predict Y (survival time)
- Assess overlap* by looking at p(X,C|T) or p(T|X,C)
- Predict survival under a specific treatment regime *k* using f(X,C,*k*)
- Will survive 5 years when treated *optimally* if $\max_k$ f(X,C, *k*) > 5

* See, e.g., Oberst, Johansson, Wei, Gao, Brat, Sontag, Varshney. Characterization of Overlap in Observational Studies, Conference on Artificial Intelligence and Statistics (AI-STATS), 2020.

A path to personalized medicine

- Clinical practice: Clinicians make (a series of) treatment decision(s) over the course of a patient's disease or disorder
  - Key decision points in the disease process
  - Could be a fixed schedule, a milestone in the disease process, or an event necessitating a decision
  - Several treatment options at each decision point
- Thus: treatment in practice involves **sequential decision-making** based on accruing information

(Marie Davidian, *An Introduction to Dynamic Treatment Regimes*)

Dynamic treatment regime

- Sequential decision rules, each corresponding to a key decision point
- Each rule tells us treatment to be given from among the available options based on the accrued information on the patient to that point
- Taken together, the rules define an algorithm for making treatment decisions
- *Dynamic* because the treatment action can vary depending on the accrued information

(Marie Davidian, *An Introduction to Dynamic Treatment Regimes*)

Example: ADHD therapy

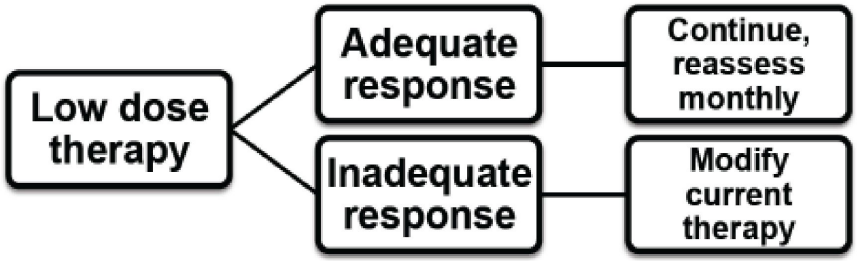

- Decision 1: Low-dose therapy – 2 options: medication or behavior modification
- Subsequent monthly decisions:
  - Responders: Continue initial therapy
  - Non-responders – 2 options: add the other therapy or increase dose of current therapy
- Objective: maximize *end-of-school-year performance*

(Material from Marie Davidian, *An Introduction to Dynamic Treatment Regimes*; example from Susan Murphy)

Example: ADHD therapy

- This is a dynamic treatment strategy because of the decision when to stop

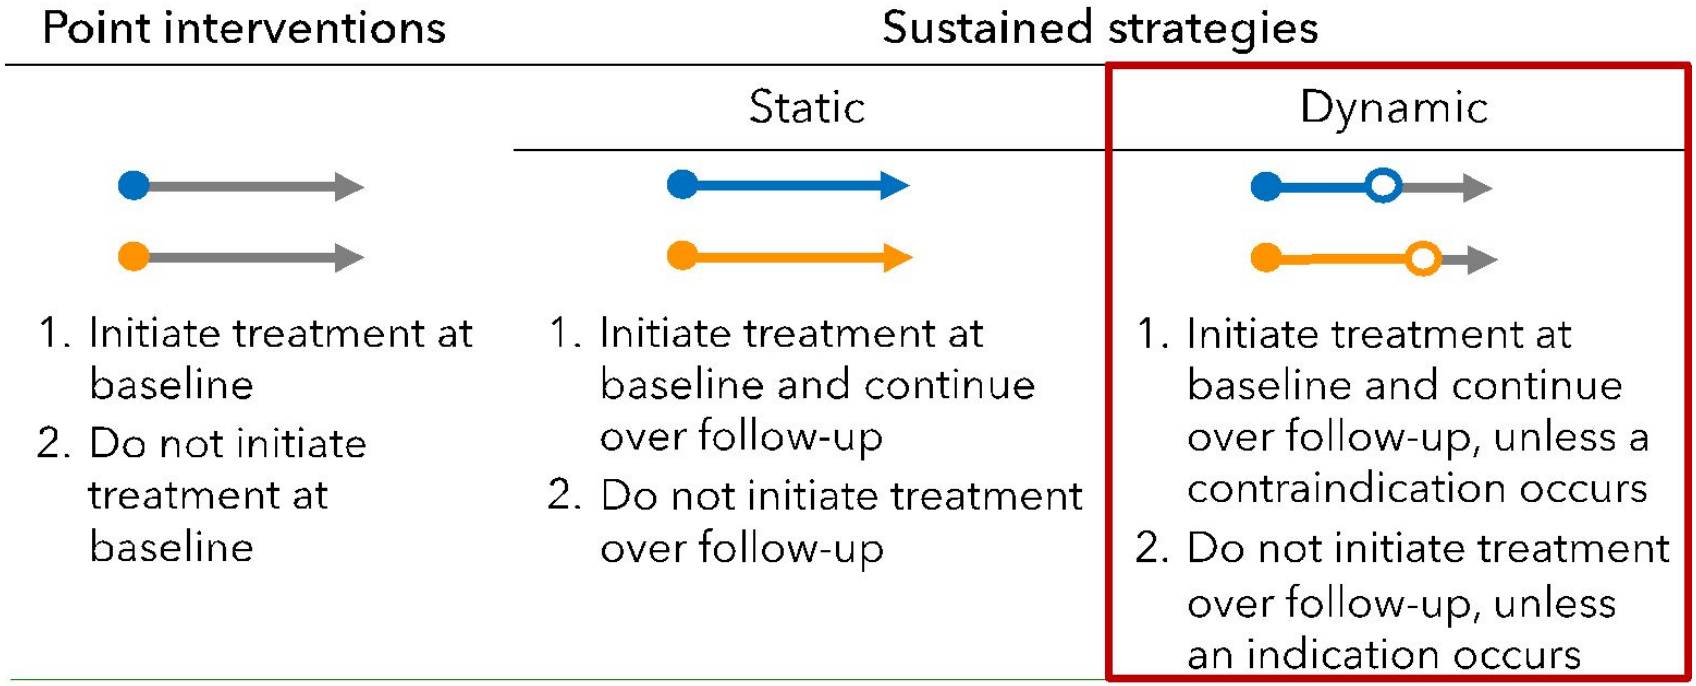

(Material from Marie Davidian, *An Introduction to Dynamic Treatment Regimes*; example from Susan Murphy)

Example: First-line treatment for multiple myeloma

- Decision 1: Induction chemotherapy (options C_1, C_2)
- Decision 2:
  - Maintenance treatment for patients who *respond* (options M_1, M_2)
  - Start a different cancer treatment for those who don't respond (options S_1, S_2)
- Objective: maximize *survival time*
- Example rules for decision 1:
  - C_1: If "age < 65 and in excellent physical health", give bortezomib, lenalidomide, dexamethasone chemotherapy followed by autologous stem cell transplant. Otherwise, treat with daratumumab, bortezomib, melphalan, & prednisone.
  - C_2: treat everyone with daratumumab, bortezomib, melphalan, & prednisone

(Marie Davidian, *An Introduction to Dynamic Treatment Regimes*)

Example: First-line treatment for multiple myeloma

- Which is the best treatment regime (policy)?
- Evaluate each of the following 8 dynamic regimes:
  1. Give C_1 followed by (M_1 if response, S_1 if no response)
  2. Give C_1 followed by (M_1 if response, S_2 if no response)
  3. Give C_1 followed by (M_2 if response, S_1 if no response)
  4. Give C_1 followed by (M_2 if response, S_2 if no response)
  5. Give C_2 followed by (M_1 if response, S_1 if no response)
  6. Give C_2 followed by (M_1 if response, S_2 if no response)
  7. Give C_2 followed by (M_2 if response, S_1 if no response)
  8. Give C_2 followed by (M_2 if response, S_2 if no response)

- Goal: evaluate the average *outcome* if all patients in the population were to follow each regime

(Marie Davidian, *An Introduction to Dynamic Treatment Regimes*)

Warm up: Evaluating dynamic treatment regimes

- 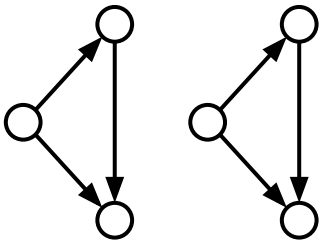

Warm up: Evaluating dynamic treatment regimes

- 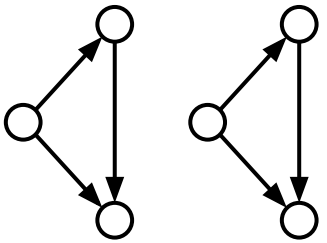

Evaluating dynamic treatment regimes

- Notice that the same estimator *does not* make sense when, e.g., S_2 depends on A_1
- The distribution of states S_2 will be affected by the policy's choice of actions A_1
  - Cannot use the observational distribution

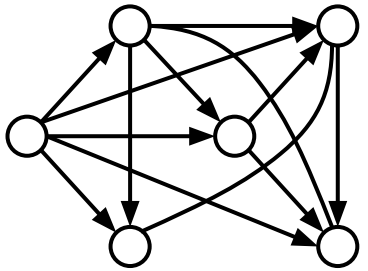

Evaluating dynamic treatment regimes: parametric G-formula

1. **Fit parametric regression models** for confounders and death at each follow-up time *t* as a function of treatment and covariate history among those under follow-up at time *t*
2. **Monte Carlo simulation** to generate a 10,000-person population under each strategy by sampling with replacement from the original study population (to estimate the standardized cumulative risk under a given strategy)
3. **Repeat in 500 bootstrap samples** to obtain 95% confidence intervals (CIs)

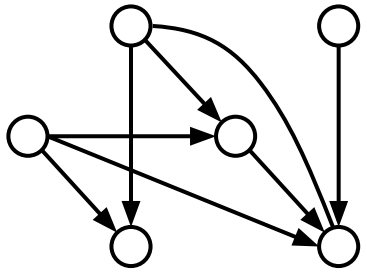

Concern: Errors may compound; also, may be insufficient data for any one time step.

[James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. Mathematical Modelling, 1986.

For recent work, see: Rui Li et al., G-Net: a Recurrent Network Approach to G-Computation for Counterfactual Prediction Under a Dynamic Treatment Regime. Proceedings of Machine Learning Research 158:282–297, 2021.]

Many more ideas and methods

- Doubly robust estimators that combine both regression and IPW
- Natural experiments & regression discontinuity
- Instrumental variables
- Sensitivity analyses

Many more ideas and methods – Natural experiments

- Does stress during pregnancy affect later child development?
- Confounding: genetic, mother personality, economic factors…
- Natural experiment: the Cuban missile crisis of October 1962. Many people were afraid a nuclear war is about to break out.
- Compare children who were in utero during the crisis with children from immediately before and after

Many more ideas and methods – Instrumental variables

- Informally: a variable which affects treatment assignment but not the outcome
- Example: are private schools better than public schools? Which students would benefit the most?
- Confounding: different student population, different teacher population
- Can’t force people which school to go to

Many more ideas and methods – Instrumental variables

- Informally: a variable which affects treatment assignment but not the outcome
- Example: are private schools better than public schools? Which students would benefit the most?
- Can't force people which school to go to
- Can *randomly* give out vouchers to some children, giving them an opportunity to attend private schools
- The voucher assignment is the instrumental variable

Estimation using an instrumental variable

Goal: estimation in setting where there are unobserved confounders, *U*, not captured in *X*

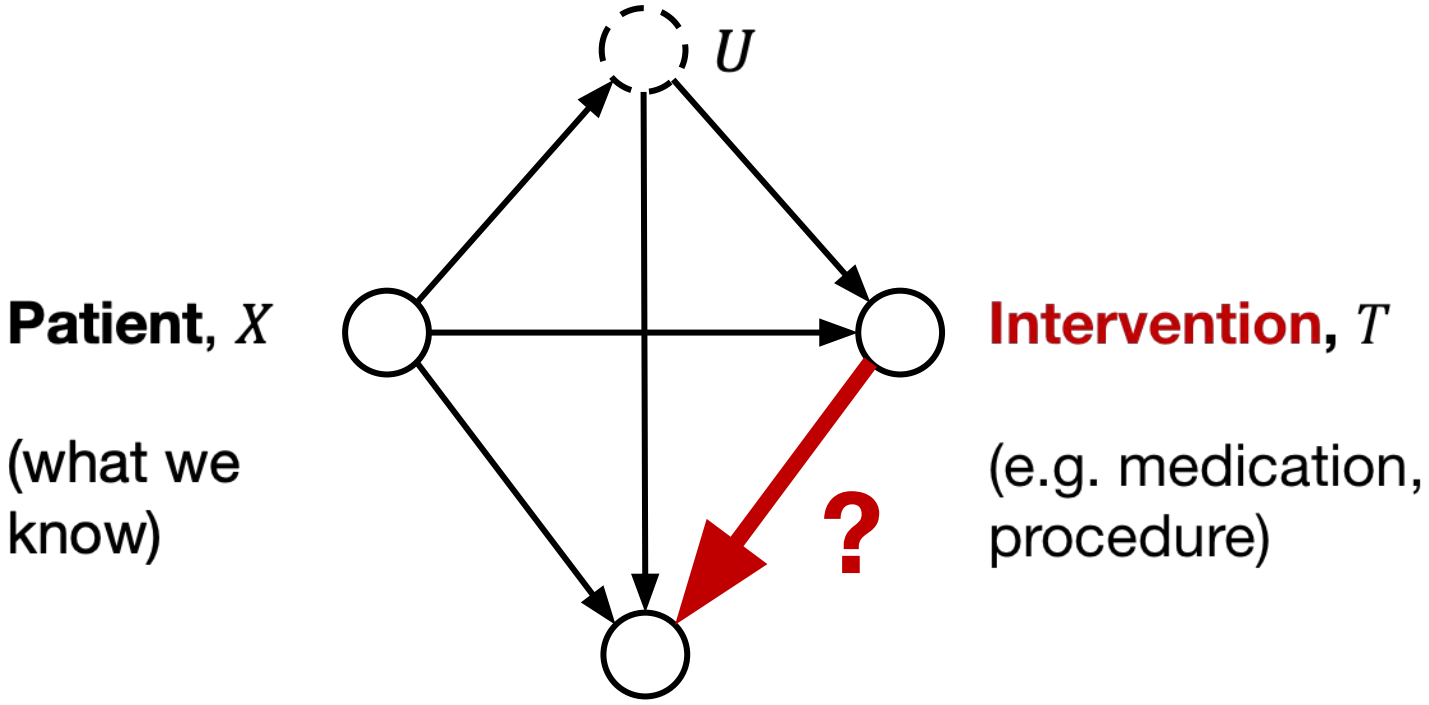

Estimation using an instrumental variable

First, assume no patient covariates (with this, we will only be able to estimate ATE not CATE)

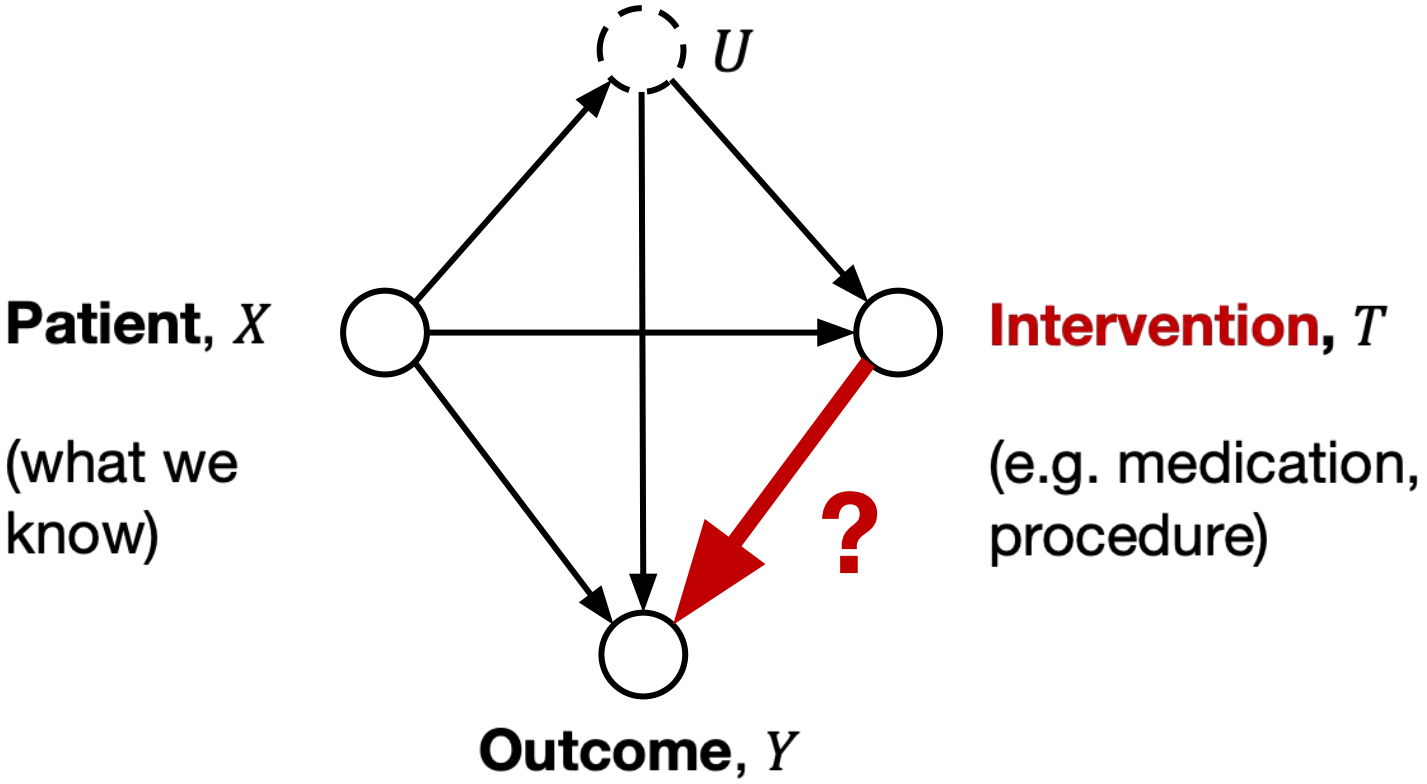

Estimation using an instrumental variable

First, assume no patient covariates (with this, we will only be able to estimate ATE not CATE)

Note: this is without loss of generality (since U could include all of X)

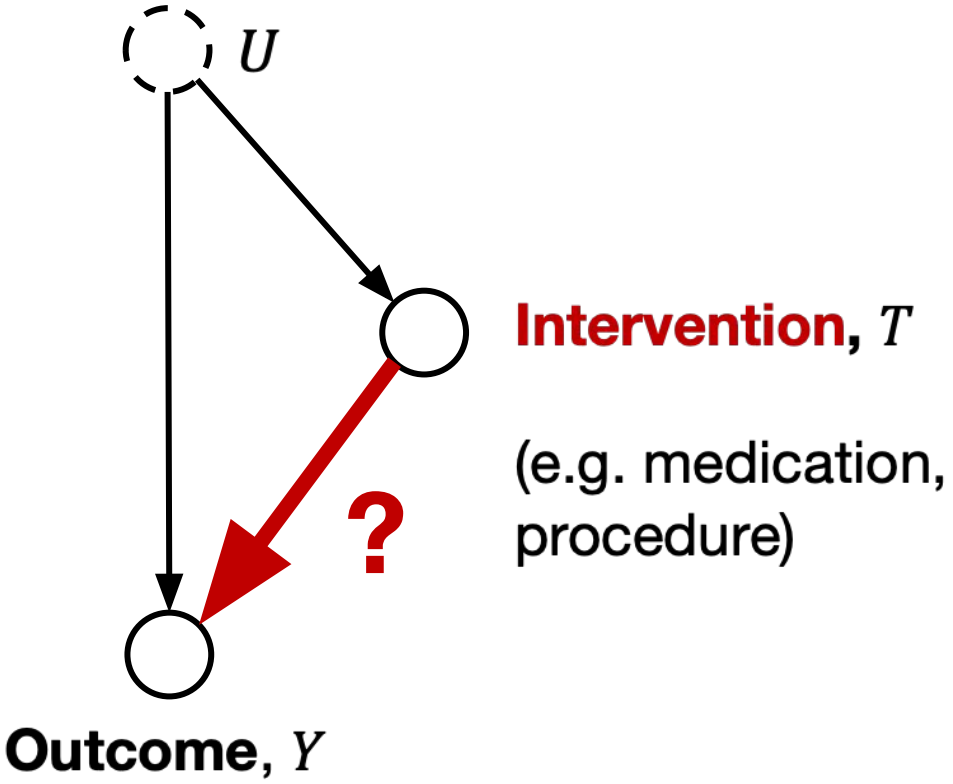

Estimation using an instrumental variable

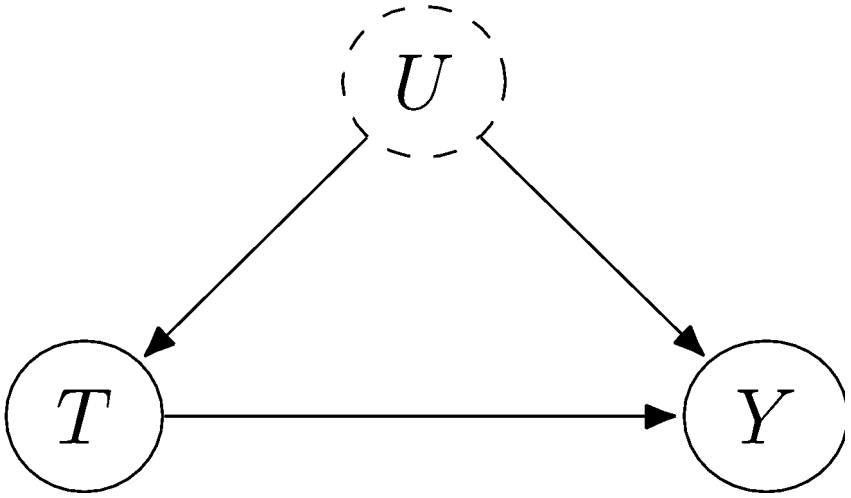

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Estimation using an instrumental variable

Instrument (e.g., voucher)

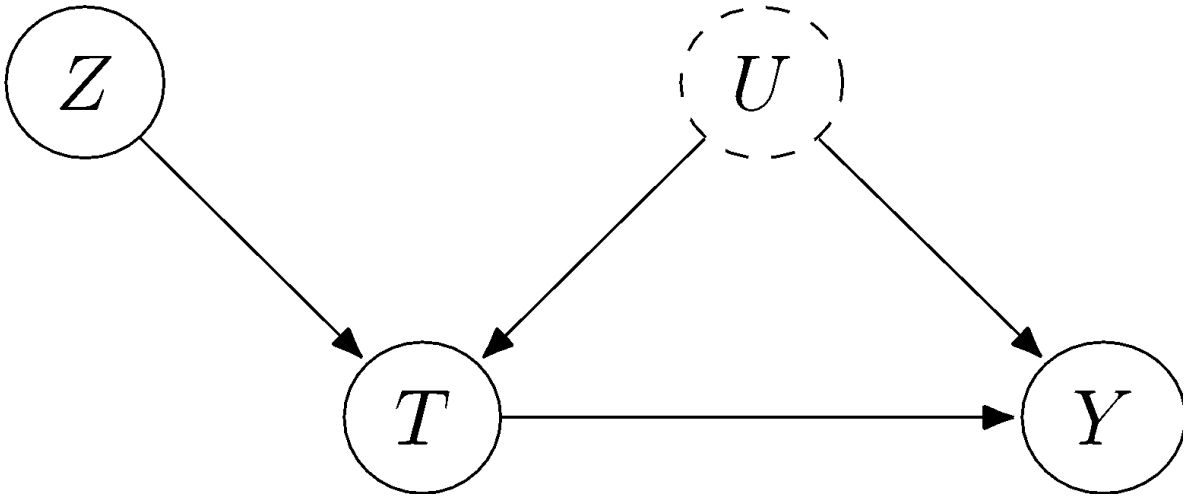

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Assumption 1: Relevance

Z has a causal effect on T

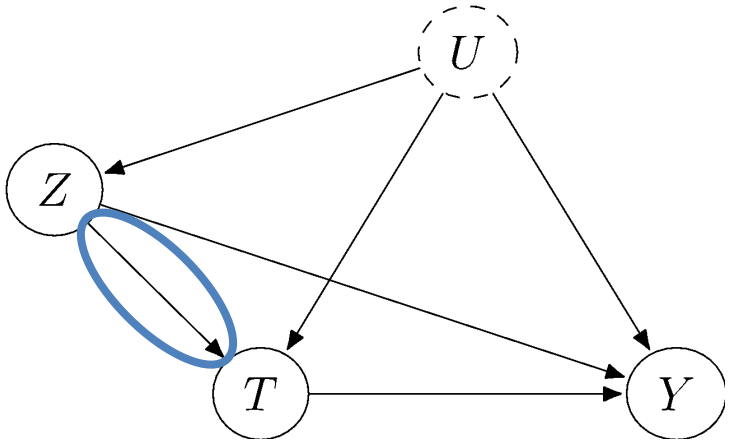

What is an Instrument?

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Assumption 2: Exclusion Restriction

The causal effect of Z on Y is fully mediated by T

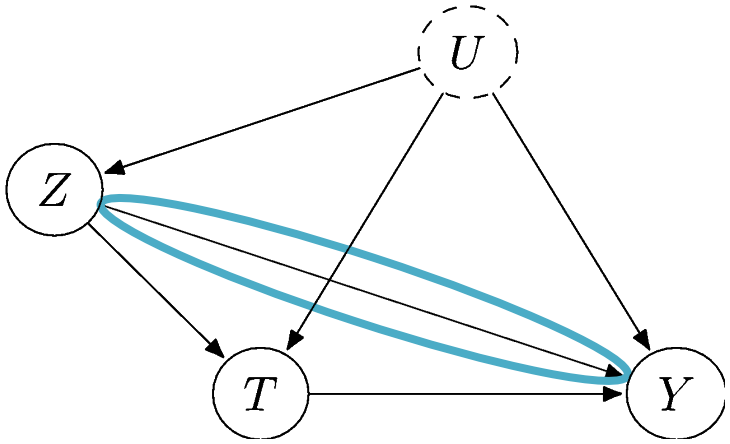

What is an Instrument?

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Assumption 3: Instrumental Unconfoundedness

Z is unconfounded (in the setting of no *X*, this simply means *U* and *Z* are independent)

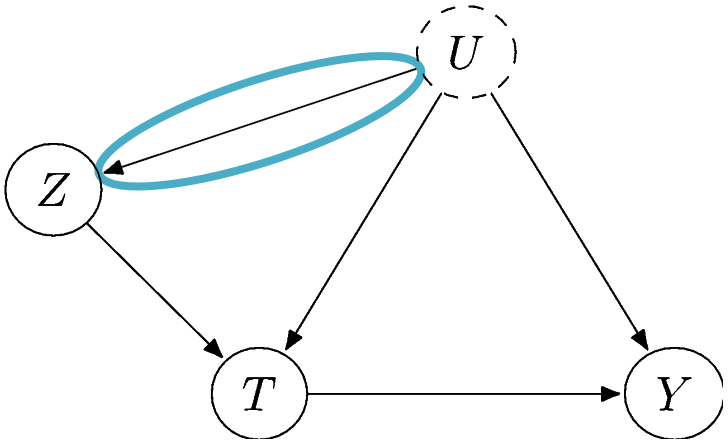

What is an Instrument?

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Warm-up: linear potential outcome, no X

Assume potential outcomes given by the linear model,

$$Y_t(U) = \alpha_u U + \delta \cdot t + \epsilon_t, \quad \mathbb{E}[\epsilon_t] = 0$$

Z doesn't appear because of the exclusion restriction assumption

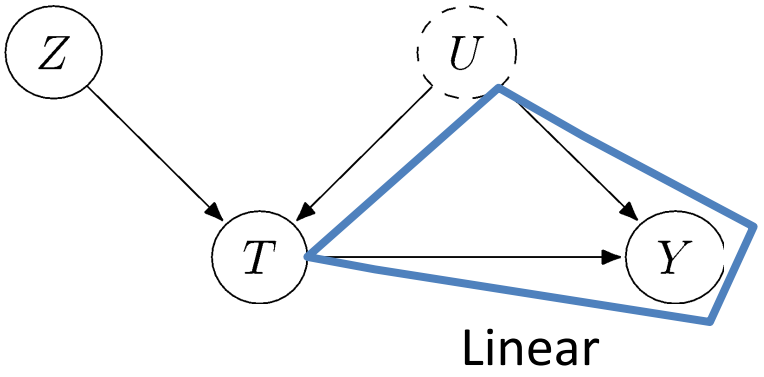

Linear

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Warm-up: linear potential outcome, no *X*

$$
\begin{aligned}
&\mathbb{E}[Y \mid Z=1] - \mathbb{E}[Y \mid Z=0] \\
&\quad = \mathbb{E}[\delta T + \alpha_u U \mid Z=1] - \mathbb{E}[\delta T + \alpha_u U \mid Z=0] && \text{(exclusion restriction and linear outcome assumptions)} \\
&\quad = \delta\left(\mathbb{E}[T \mid Z=1] - \mathbb{E}[T \mid Z=0]\right) + \alpha_u\left(\mathbb{E}[U \mid Z=1] - \mathbb{E}[U \mid Z=0]\right) \\
&\quad = \delta\left(\mathbb{E}[T \mid Z=1] - \mathbb{E}[T \mid Z=0]\right) + \alpha_u\left(\mathbb{E}[U] - \mathbb{E}[U]\right) && \text{(instrumental unconfoundedness assumption)} \\
&\quad = \delta\left(\mathbb{E}[T \mid Z=1] - \mathbb{E}[T \mid Z=0]\right)
\end{aligned}
$$

$$
\delta = \frac{\mathbb{E}[Y \mid Z=1] - \mathbb{E}[Y \mid Z=0]}{\mathbb{E}[T \mid Z=1] - \mathbb{E}[T \mid Z=0]}
$$

(non-zero due to relevance assumption)

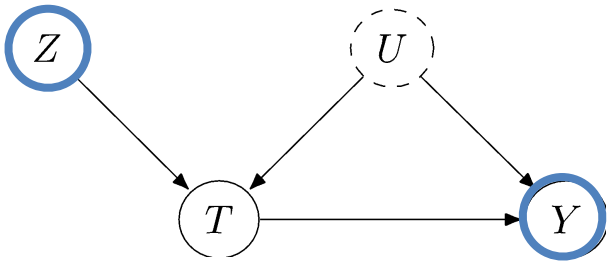

$$
Y_t(U) = \alpha_u U + \delta \cdot t + \epsilon_t
$$

(Slides adapted from Brady Neal's Introduction to Causal Inference class)

Estimation using (conditional) instruments

Assume potential outcomes given by:

$$Y_T(x, U) = \delta(x)T + g(x, U) + \epsilon_T$$

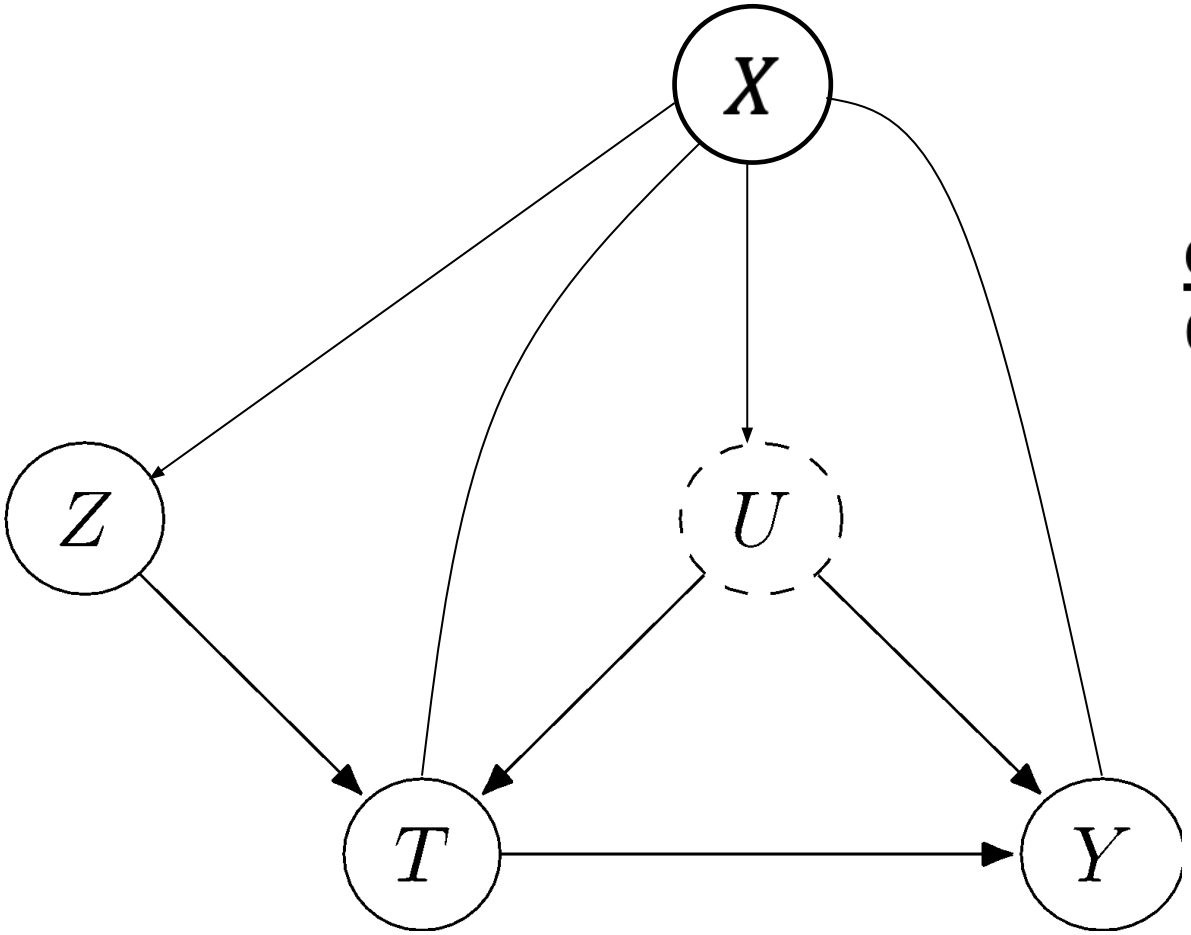

Goal: estimate

$$\text{CATE}(x) = \delta(x)$$

Estimation using (conditional) instruments

Assume potential outcomes given by:

$$Y_T(x, U) = \delta(x)T + g(x, U) + \epsilon_T(x)$$

Theorem: $$\text{CATE}(x) = \delta(x) = \frac{\mathbb{E}[Y|Z=1,x] - \mathbb{E}[Y|Z=0,x]}{p(T=1 \mid Z=1,x) - p(T=1 \mid Z=0,x)}$$

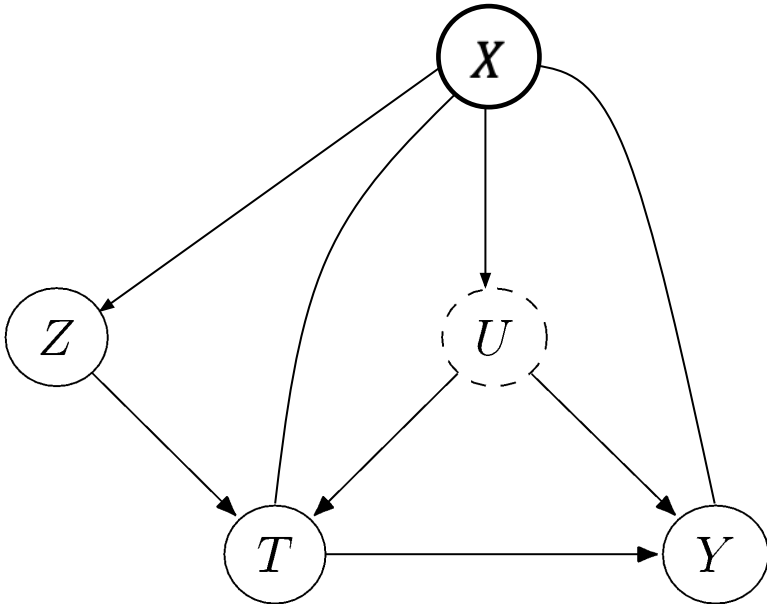

Assume

$\mathbb{E}[\epsilon_0 | x] = 0$

$\mathbb{E}[\epsilon_1 | x] = 0$

Summary

- Close connection between causal inference and off-policy evaluation
- Same ideas can be used to evaluate dynamic treatment regimes when there are *multiple timepoints / actions*
- Instrumental variables can be used to estimate ATE and CATE when there is unobserved confounding

References

- Introduction to causal inference from a machine learning perspective by Brady Neal, 2020.
  - Section 8.2: Sensitivity Analysis
  - Chapter 9: Instrumental Variables

  (See also the many references within for both recent literature and where these methods were originally introduced.)
- Syrgkanis et al., Machine Learning Estimation of Heterogeneous Treatment Effects with Instruments, NeurIPS 2019.
- Boominathan et al., Treatment Policy Learning in Multiobjective Settings with Fully Observed Outcomes, KDD 2020.

Additional references

- Chakraborty & Moodie, Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine. Springer, 2013
- O. Gottesman, F. Johansson, M. Komorowski, A. Faisal, D. Sontag, F. Doshi-Velez, L. Celi. Guidelines for reinforcement learning in healthcare. Nature medicine, 2019
- Li et al., G-Net: a Recurrent Network Approach to G-Computation for Counterfactual Prediction Under a Dynamic Treatment Regime. Proceedings of Machine Learning Research 158:282–297, 2021
- Hua, Mei, Zohar, Giral, Xu. Personalized Dynamic Treatment Regimes in Continuous Time: A Bayesian Approach for Optimizing Clinical Decisions with Timing. Bayesian Analysis. Advance Publication, 1-30, 2021