NeurIPS 2018 Spotlight Presentation

Removing Hidden Confounding by Experimental Grounding

Nathan Kallus
Cornell
(↑ me)

Aahlad Manas Puli
NYU

Uri Shalit
Technion

Poster: Today 10:45AM–12:45PM @ Room 210 & 230 AB #2

Individual-level causal effects

- Patient Anna is diagnosed with type-II diabetes
  - Blood sugar: 8.7% A1C
  - Age: 51
  - Weight: 102kg
  - BMI: 35.3
  - ...

  Baseline covariates X
- **Q**: What first-line glucose control treatment to give? Insulin ($t = 1$) or Metformin ($t = 0$)?
  - Want to know the individual-level **causal** effect of treatment, *i.e.*, the *conditional average treatment effect (CATE)*

$$\tau(X) = \mathbb{E}[Y(1) - Y(0) \mid X]$$

    - $Y(t)$ = Anna's potential outcome under treatment t
  - Same Q in targeted advertising, public policy, ...

Large-scale observational data can help

Age	Weight	BMI	A1C	LDL	T	Y
49	106	31			Insulin	9
54	89	26			Metformin	7
43	130	38			Metformin	10
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Large-scale observational data can help

Age	Weight	BMI	A1C	LDL	T	Y
49	106	31			Insulin	9
54	89	26			Metformin	7
43	130	38			Metformin	10
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Fit $\omega(X) = \mathbb{E}[Y \mid X, T = 1] - \mathbb{E}[Y \mid X, T = 0]$ to the data

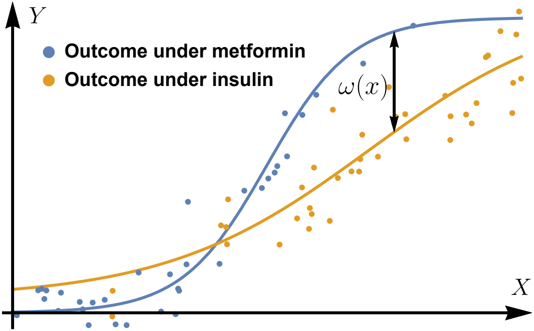

E.g.:
Wager & Athey '17 (CF),
Shalit et al. '17 (TARNet),
...

usually assume $\omega = \tau$
(no hidden confounding)

The problem with hidden confounding

- Hidden confounding $=$ hidden correlations between treatments and outcome idiosyncrasies
  - *E.g.*: healthier patients tend to get metformin
  - Confounding $\implies \omega \neq \tau$ ☹
  - To some extent always unavoidable in observational data

The problem with hidden confounding

- Hidden confounding = hidden correlations between treatments and outcome idiosyncrasies
  - *E.g.*: healthier patients tend to get metformin
  - Confounding $\implies \omega \neq \tau$ ☹
  - To some extent always unavoidable in observational data

Experimental data

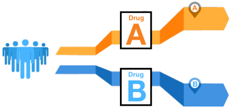

Unconfounded by design ☺
Limited generalizability ☹
Small samples ☹

Observational data

Confounded by default ☹
Covers population of interest ☺
Large samples ☺

Experimental grounding

- **Our Q**: How to use a small & limited experimental dataset to remove confounding errors in individual-level treatment effect estimates from a large observational dataset?

Experimental grounding

- **Our Q**: How to use a small & limited experimental dataset to remove confounding errors in individual-level treatment effect estimates from a large observational dataset?
- Outline of our method:
  - Fit $\hat{\omega}(X)$ using blackbox on observational data (*e.g.*, causal forest, TARNet, ...)
  - A **new way** to fit $\eta(X) = \omega(X) - \tau(X)$ across the observational and experimental datasets
  - Return the grounded estimate $\hat{\tau}(X) = \hat{\omega}(X) - \hat{\eta}(X)$

Experimental grounding

- **Our Q**: How to use a small & limited experimental dataset to remove confounding errors in individual-level treatment effect estimates from a large observational dataset?
- Outline of our method:
  - Fit $\hat{\omega}(X)$ using blackbox on observational data (*e.g.*, causal forest, TARNet, ...)
  - A **new way** to fit $\eta(X) = \omega(X) - \tau(X)$ across the observational and experimental datasets
  - Return the grounded estimate $\hat{\tau}(X) = \hat{\omega}(X) - \hat{\eta}(X)$
- Our theoretical guarantee: if η is parametric and $\hat{\omega}$ is consistent then $\hat{\tau}$ is consistent! ☺
  - Strictly weaker than assuming no confounding ($\eta = 0$)

Empirical results

- Estimate the effect of large vs small classrooms on first graders' test scores
  - Data from STAR experiment (Word et al. 1990)

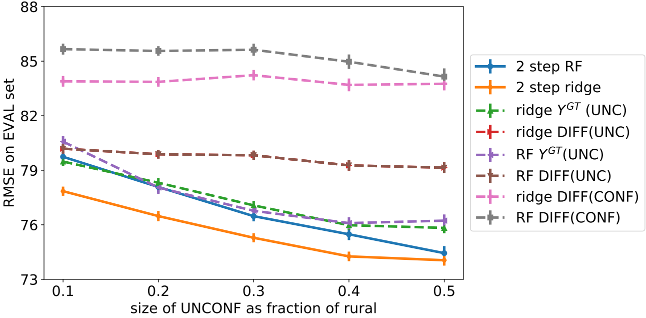

Thank you!

Poster: Today 10:45AM–12:45PM @ Room 210 & 230 AB #2