FACULTY OF SCIENCES

Causal Inference in Psychology and Applied Health Sciences, Ghent, June 2024

An Introduction to Causal Machine Learning

Stijn Vansteelandt

Ghent University, Belgium

GHENT UNIVERSITY

Introduction

GHENT UNIVERSITY

Evaluating treatment effects

- Evaluation of the effect of a treatment A on an outcome Y is commonly based on contrasts
$$E(Y^1 - Y^0)$$
of the expected outcome with (Y^1) versus without (Y^0) treatment.
- In observational studies, this demands adjustment for potentially high-dimensional confounders.
- Two popular approaches are standardisation and inverse probability weighting.

Standardisation

To estimate the mean outcome under treatment,

- train a prediction model for outcome in the treated, using confounders;
- use this to predict outcome for all;
- average these predictions.
- The use of machine learning is increasingly popular.

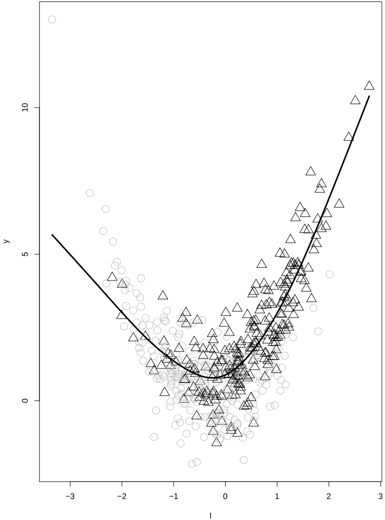

WHY MACHINE LEARNING?

- Model misspecification is likely,
  and difficult to diagnose
  when treated and untreated subjects have limited overlap.
- The analysis can be made more objective
  by pre-specifying the machine learning algorithms.
  - In contrast, the human process of building a model is time-consuming
    and even more black box; pre-specifying it is difficult.
- If a more statistical approach is deemed preferable,
  then stacking statistical and machine learners allows one to do at least as good.

But…

GHENT UNIVERSITY

Two caveats

Caveat 1: no valid uncertainty margins

machine learning 'easily' produces estimates,
but we have 'no clue' how precise these are...

- Even sample splitting or the bootstrap does not work.

(e.g. Samworth, 2011)

Caveat 2: plug-in bias

plugging machine learning predictions into a statistical analysis,
typically induces plug-in bias.

- The bias-variance tradeoff is so heavily optimized towards minimal prediction error, that machine learning algorithms underperform when used for other purposes.
- It leads to biased estimates, p-values and confidence intervals.

WHAT IS PLUG-IN BIAS?

- Plug-in bias is the result of oversmoothing in the range of the data where predictions are needed,
- or due to mistakenly throwing out important confounders.

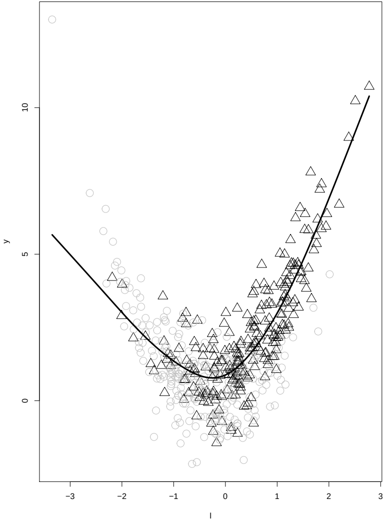

DEBIASED MACHINE LEARNING

GHENT UNIVERSITY

A BIT OF HISTORY...

- Foundations for a solution have been laid in the 80's - 90's.

  (e.g. Pfanzagl, 1982; Bickel et al., 1998; Newey, 1990; Robins and Rotnitzky, 1995; van der Vaart, 1991)

- van der Laan made use of this theory
  to construct plug-in estimators based on machine learning,
  which he called Targeted Maximum Likelihood Estimators.

  (van der Laan and Rubin, 2008; van der Laan and Rose, 2014)

- His approach is now called targeted learning.
- Chernozhukov, Newey, Robins, ... popularised this theory,
  under weaker conditions by invoking sample splitting.

  (Robins et al., 2008; Chernozhukov et al., 2018)

- They refer to their approach as double / debiased machine learning.

Observational data

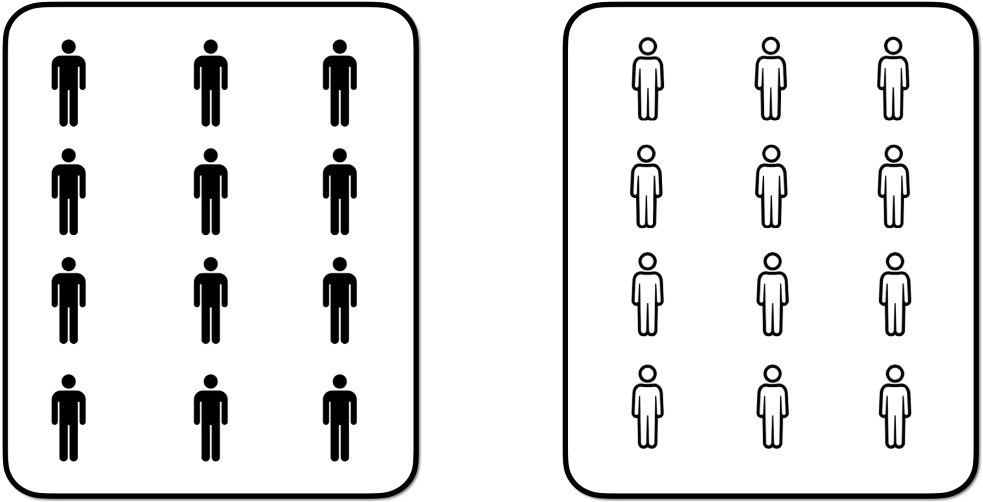

TRAIN IN TREATED, USING CONFOUNDERS

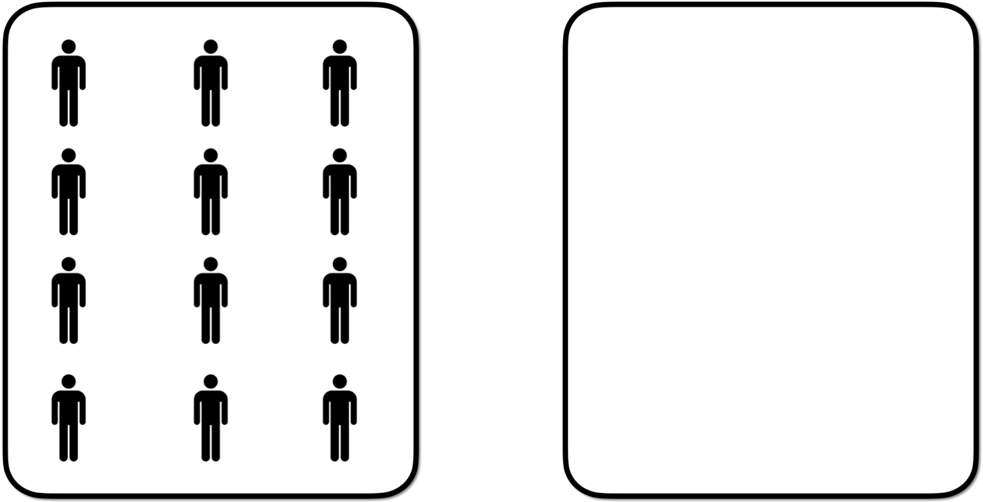

PREDICT OUTCOME ON TREATMENT FOR ALL

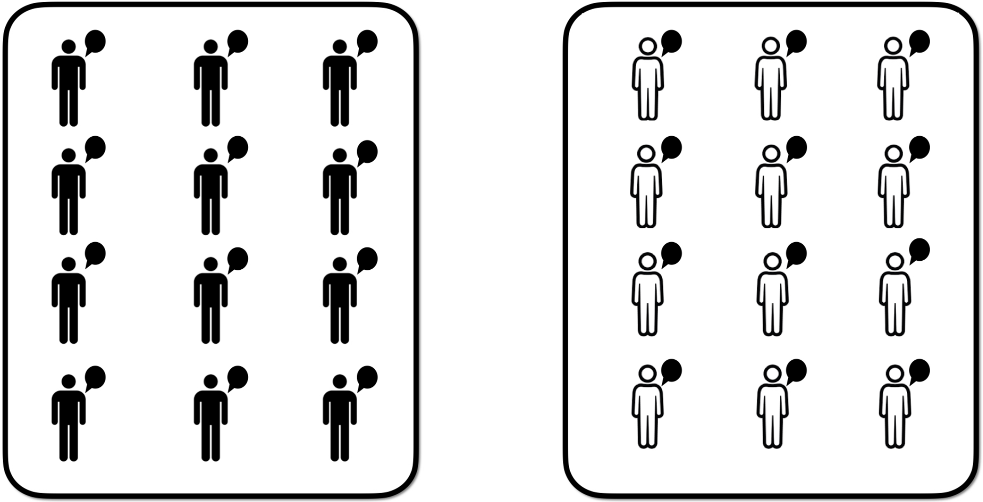

AVERAGE PREDICTED TREATMENT OUTCOME OVER ALL

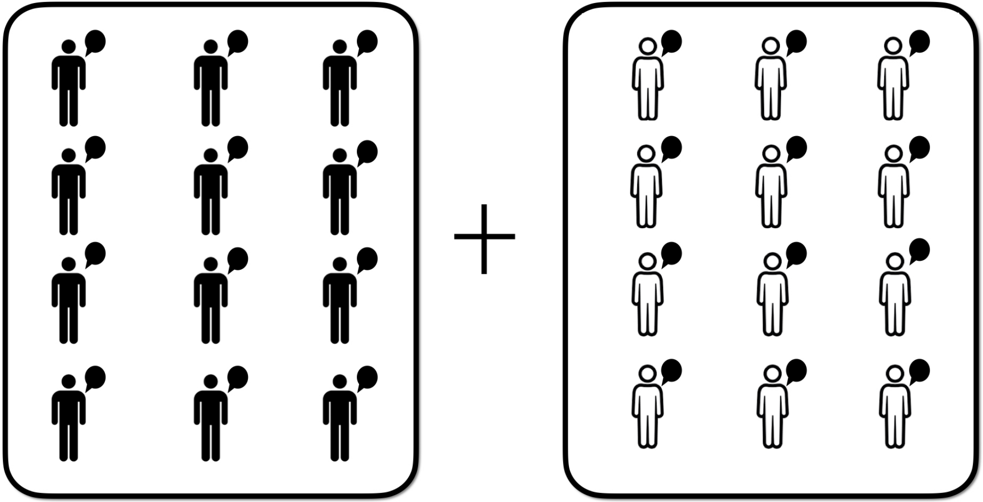

How to debias outcome mean on treatment?

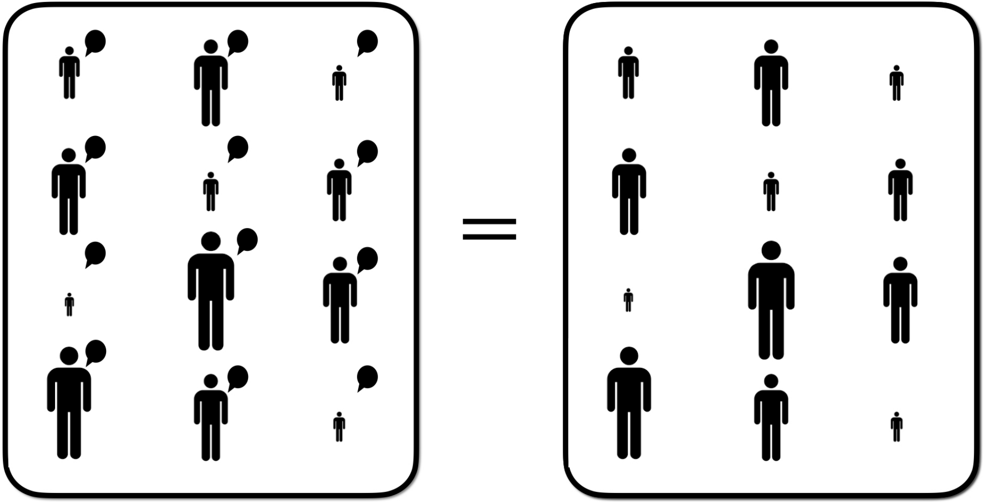

A SKETCH HOW TO DEBIAS OUTCOME MEAN ON TREATMENT

- To learn the amount of plug-in bias,
  we evaluate prediction errors in the treated,
  but weigh them (inversely to the propensity score) to approximate bias in the full sample.
- Debiased machine learning subtracts this bias from the estimate.
- Targeted learning updates predictions to be free of bias.
- Sample splitting is used to prevent overfitting bias.

TARGETED / DEBIASED LEARNING

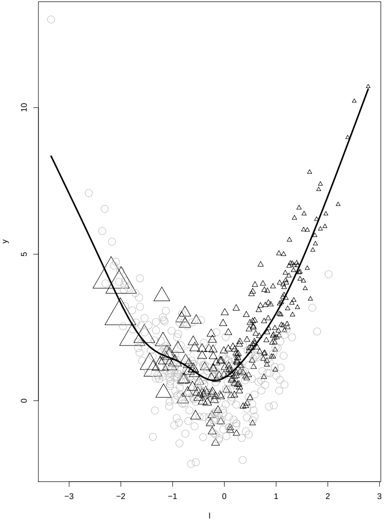

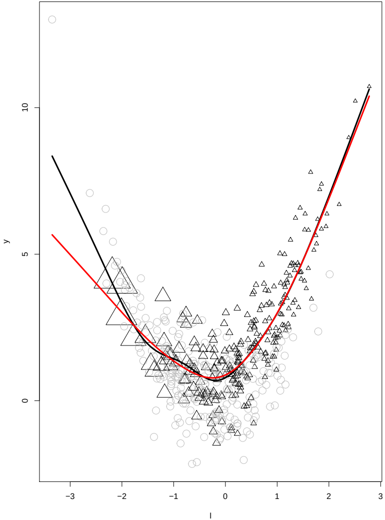

AN IMPRESSION FROM SIMULATION STUDIES

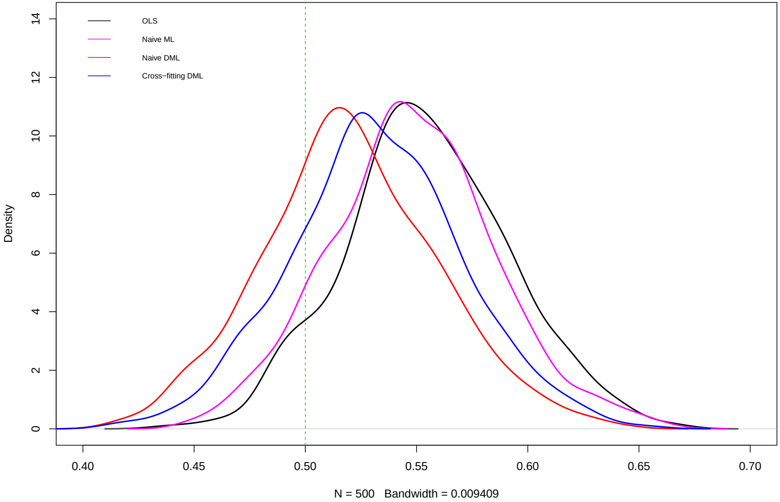

SUMMARY

GHENT
UNIVERSITY

SUMMARY

- Standard statistical analyses leave residual confounding bias
  due to model misspecification or the difficulty to pre-specify the analysis.
- Future lies in debiased / targeted learning.
- Enables one to be model-free,
  via flexible, automated, objective modeling.
- It delivers honest standard errors that acknowledge model uncertainty.
    - Debiased learning techniques are therefore important,
      even when parametric models with variable selection are used.

Summary

- Causal machine learning can be viewed as machine learning for evaluating treatment effects as opposed to prediction.
- This is much harder: we can compare predictions with observed outcomes, but cannot compare estimated with true treatment effects.
- Machine-learning based effect estimates must be de-biased, based on estimand's efficient influence curve.

  Hines, O., Dukes, O., Diaz-Ordaz K., and Vansteelandt, S. (2021). Demystifying statistical learning based on efficient influence functions. The American Statistician, 1-48.
- Most existing works have focused on the average effect of a binary treatment, leading to lack of flexibility and oversimplification.
- Assumption-lean modeling bridges traditional modeling with debiased machine learning.

  Vansteelandt, S., & Dukes, O. (2022). Assumption-lean inference for generalised linear model parameters (with discussion). Journal of the Royal Statistical Society - B, 84, 657-685.

Assumption-lean modeling

- For an exposure A and confounders L, consider the log-linear model

$$\log\{E(Y|A=a,L)\} = \alpha' L + \beta a$$

- We can relax this to a semi-parametric regression model

$$\log\{E(Y|A=a,L)\} = \alpha(L) + \beta a$$

- And even further to assumption-lean modeling

$$\log\{E(Y|A=a,L)\} = \alpha(L) + \beta(L) a$$

where we learn the mean and variance of $\beta(L)$
using debiased machine learning for estimation.

ASSUMPTION-LEAN MODELING ALGORITHM

1. Predict A based on L to obtain predictions $\hat{p}_i$.
2. Predict Y based on A and L to obtain predictions $\hat{Y}_i$.
3. Predict $\log\left(\hat{Y}\right)$ based on L to obtain predictions $\hat{q}_i$.
4. Calculate

$$\log\left(\hat{Y}_i\right) - \hat{q}_i + \frac{Y_i}{\hat{Y}_i} - 1$$

   and linearly regress it on $A_i - \hat{p}_i$ using least squares
   to obtain an estimate for β and a robust standard error.

Selected references

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. The Econometrics Journal, 21, C1–C68.

Hines, O., Dukes, O., Diaz-Ordaz K., and Vansteelandt, S. (2021). Demystifying statistical learning based on efficient influence functions. The American Statistician, 76, 292-304..

van der Laan, M. J., & Rose, S. (2011). Targeted learning: causal inference for observational and experimental data. Springer Science & Business Media.

Vansteelandt, S., & Dukes, O. (2022). Assumption-lean inference for generalised linear model parameters (with discussion). Journal of the Royal Statistical Society - B, 84, 657-685.

Vansteelandt, S. (2021). Statistical modelling in the age of data science. Observational Studies, 7, 217-228.

Slides: `users.ugent.be/~svsteela/`