Targeted Learning of Marginal Effects

Alex Luedtke

A gentle introduction to targeted learning in RCTs: what, why and how?

June 29, 2021

Department of Statistics, University of Washington
Department of Biostatistics, University of Washington
Vaccine and Infectious Disease Division, Fred Hutch

Thanks to **Kelly Van Lancker** for organizing this interesting workshop!

Table of Contents

Marginal effect estimands

- Marginal effects aim to describe the impact of a treatment/intervention at an aggregate level.
- **Primary objectives** in clinical trials often concern marginal effects.
- Marginal effects of interest typically represent a contrast between what would have been expected to happen if all trial participants had received the treatment versus control.
  - Under randomization, this may, for example, correspond to a contrast between $E[Y|A = 1]$ and $E[Y|A = 0]$.
  - E.g., **relative risks** or **average treatment effects**.
- There are also methods for defining marginal effect estimands that **generalize** from the clinical trial population to another population of interest.
  - Though I won't focus on these methods today, everything I discuss also carries over to these more general marginal effect estimands.

Imputation strategy

- Typically, the disease risk on treatment and control, namely $E[Y|A=1]$ and $E[Y|A=0]$, are estimated via empirical means on the treatment and control arms.
- Stijn showed how an imputation approach could leverage baseline covariate information to better estimate these quantities.

Age	Trt	Y	Y^1	$\hat{P}^1$	Y^0	$\hat{P}^0$
40	1	1	1	0.8	?	**0.70**
50	1	0	0	0.6	?	**0.55**
60	1	1	1	0.7	?	**0.60**
50	0	0	?	0.7	0	**0.60**
30	0	1	?	0.6	1	**0.50**
40	0	0	?	0.5	0	**0.45**

- **Averaging** $\hat{P}^0$ gives an estimate of the disease risk on control.
  - And an analogous strategy can be used to estimate the disease risk on treatment.

How can this strategy be further improved?

- Improving the quality of the imputations should improve the estimates.
- The statistics and machine learning communities have developed many flexible strategies for predicting an outcome given covariates.
  - E.g., random forest, gradient boosting, generalized additive models, splines.

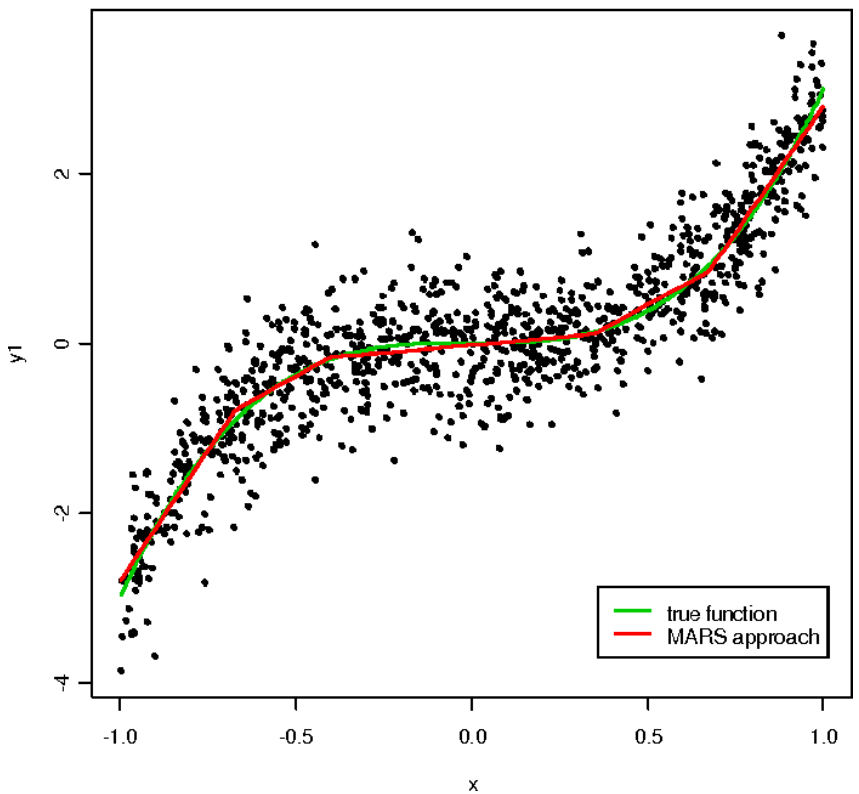

How can this strategy be further improved?

- **Ensemble methods** also exist that can optimally choose between parsimonious approaches, such as linear regression, and more flexible strategies (e.g., van der Laan et al. 2007).
  - See the `SuperLearner` package in R for an implementation of one such approach.

How can this strategy be improved?

Age	Trt	Y	Y^1	$\hat{P}^1$	Y^0	$\hat{P}^0$
40	1	1	1	0.8	?	**0.70**
50	1	0	0	0.6	?	**0.55**
60	1	1	1	0.7	?	**0.60**
50	0	0	?	0.7	0	**0.60**
30	0	1	?	0.6	1	**0.50**
40	0	0	?	0.5	0	**0.45**

- Kelly showed that, in some cases, the simple imputation strategy that estimates the disease risk on control by averaging $\hat{P}^0$ is **robust to model misspecification**.
  - E.g., this is true when $\hat{P}^0$ is obtained via a linear model or a logistic regression.
- BUT: when flexible approaches are used to obtain the imputations, the resulting estimator may be **overly biased**.
  - Consequently, confidence intervals based on these estimators may not have proper coverage.

Targeting initial imputations to reduce bias

Age	Trt	Y	Y^1	$\hat{P}^1$	Y^0	$\hat{P}^0$
40	1	1	1	0.8	?	0.70
50	1	0	0	0.6	?	0.55
60	1	1	1	0.7	?	0.60
50	0	0	?	0.7	**0**	**0.60**
30	0	1	?	0.6	**1**	**0.50**
40	0	0	?	0.5	**0**	**0.45**

- **Good news:** in a randomized trial, the bias can be estimated by simply comparing the mean of the **imputed outcomes** to that of the **actual observed outcomes** among the controls!

- Targeted learning is designed to modify the initial imputations so that the estimate of the bias is exactly zero.
  - Some slightly involved arguments show that this in fact works: the resulting imputation estimator is has negligible bias!

Benefits of targeted learning of marginal effects

- The **standard approach** to estimate the disease risk on control only makes use of data from controls.

- **Targeted learning** makes use of all available data.

  1) Imputation model is fitted using **baseline covariates** and **outcomes** on the control arm.
  2) Initial imputations are obtained by evaluating the fitted model on each participant's **baseline covariates**, regardless of their randomization arm.
  3) Control-arm **baseline covariates** and **outcomes** are then used to remove the bias from these initial imputations.

- As a consequence of more efficiently using the available data, targeted learning typically yields **more precise estimates** and **tighter confidence intervals** than do standard approaches.
  - This can make it possible to achieve a desired power with **smaller sample sizes**, resulting in **faster enrollment** or **fewer trial sites**.

Table of Contents

What changes when the outcome is a time-to-event?

- Targeted learning can also be used for time-to-event endpoints.
- In these cases, targeted learning can lead to **improved robustness**.
- Here I'll focus on estimating a **survival function**.
- Note: there hasn't been as much work in the targeted literature on estimating hazard ratios.
  - But there are some examples: e.g., Whitney et al. (2019).
  - Closely related covariate adjustment approaches can also be employed to estimate hazard ratios: e.g., see Lu and Tsiatis (2008).

What do standard analyses require?

- Standard analyses (e.g., Kaplan-Meier) require **independent censoring**, that is, that the survival time is independent of censoring time within each randomization arm.
- Under this condition, the **risk set is representative** of all individuals who haven't experienced the event by a certain time, and so can be used to impute their outcome:

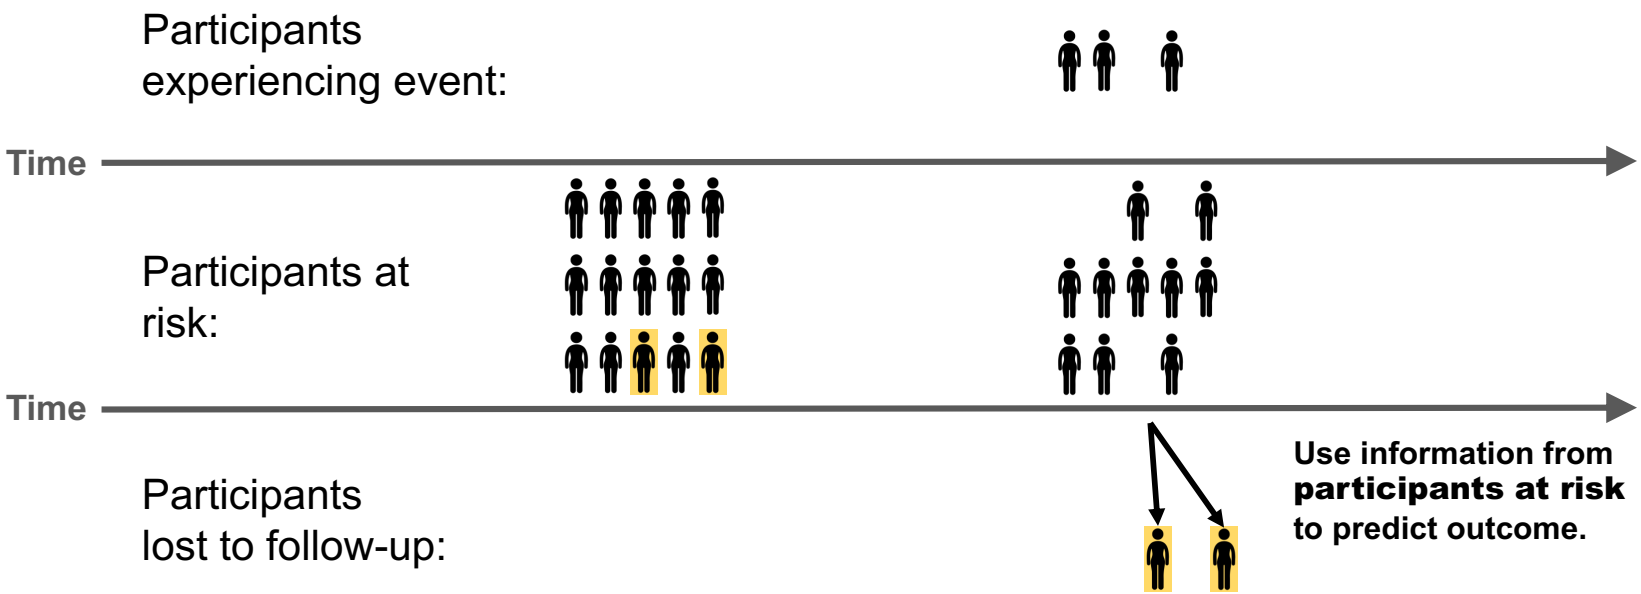

When will the independent censoring condition be violated?

- Independent censoring may be violated if, for example, participants at higher risk of the disease are also less likely to be lost to follow-up.
    - In this case, Kaplan-Meier would overestimate the disease risk.

What does targeted learning allow for?

- Targeted learning allows for a **conditionally independent censoring** condition that can often be more plausible.
  - This condition states that the survival time is independent of censoring time within each (randomization arm, baseline covariate) stratum.
- Under this condition, outcomes can be imputed within each (randomization arm, baseline covariate) stratum.

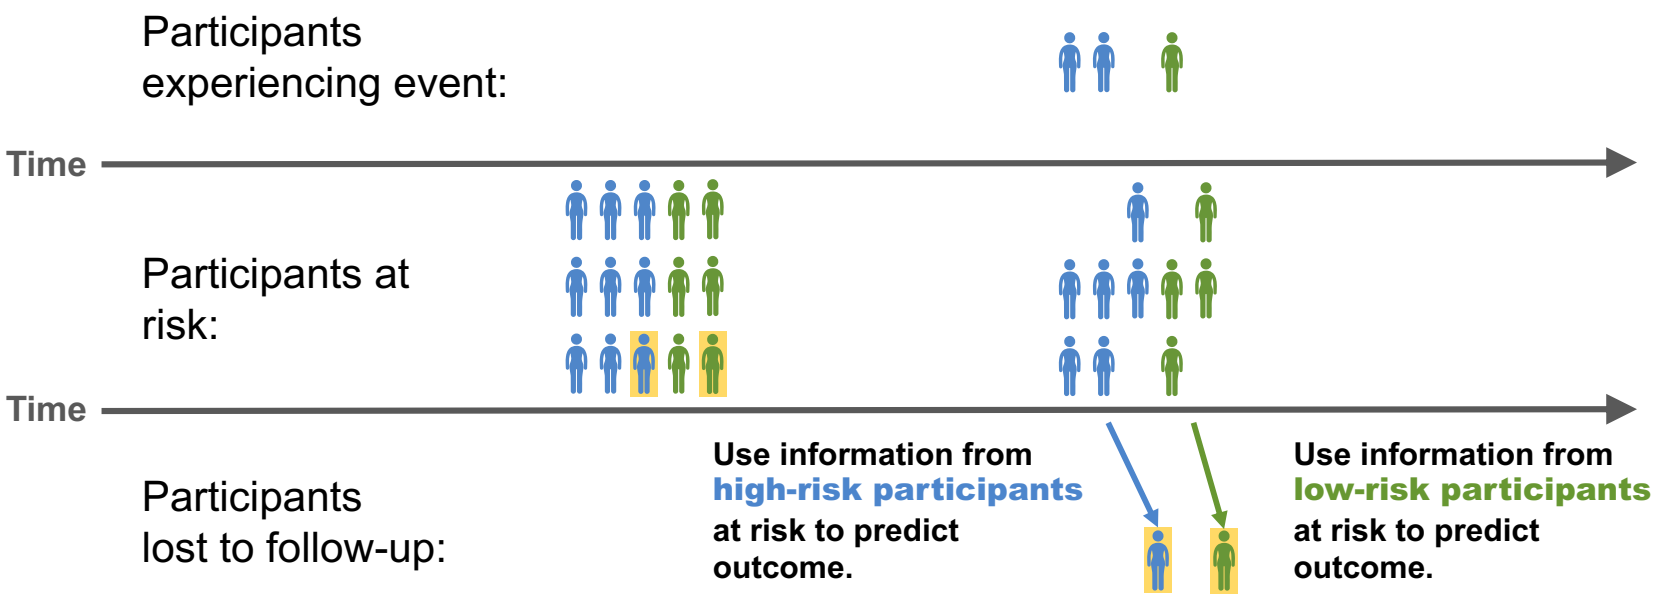

What about covariate adjustment for survival outcomes?

- Covariate adjustment for time-to-event outcomes works similarly as for non-time-to-event outcomes.
- When estimating the disease risk on control, covariates can be employed to predict what outcome participants on the treatment arm would have had if they had received control.

What are the advantages and disadvantages of targeted learning with time-to-event outcomes?

- Targeted learning estimators are **more robust** than standard approaches.
- Unlike in non-time-to-event settings, in the survival context, typical targeted learning estimators may or may not be more precise than standard approaches.
- Consider two extremes:

  Scenario 1: When covariates are **predictive of survival** and are **not predictive of censoring**, targeted learning estimators will be more precise.

  Scenario 2: When covariates are **predictive of censoring** and are **not predictive of survival**, standard approaches will be more precise.

- In intermediate cases, there is no clear ordering between the precision of these estimators.
- **Key takeaway:** Only adjust for covariates that may plausibly be predictive of the outcome!

Table of Contents

Example for non-time-to-event data

```
# Using implementation from tmle package
library(tmle)

# Generate toy data set
set.seed(1)
n = 100 # sample size
W = data.frame(W1=rnorm(n),W2=rnorm(n,1,1/2)) # covariates
A = sample(c(rep(0,n/2),rep(1,n/2))) # treatment
Y = rnorm(n) + A + 0.25*A*W$W1 # outcome

# Estimate average treatment effect E[Y|A=1]-E[Y|A=0]
out = tmle(Y,A,W,gform=A~1)
# Above, gform specifies model for Pr(A=1|W=w)

# Estimated ATE
out$estimates$ATE$psi
[1] 1.016946

# 95% confidence interval
out$estimates$ATE$CI
[1] 0.6569847 1.3769072
```

Software for time-to-event endpoints

- For time-to-event endpoints, the `survtmle` package in R implements the methods discussed today.

References

X. Lu and A. A. Tsiatis. Improving the efficiency of the log-rank test using auxiliary covariates. *Biometrika*, 95(3):679–694, 2008.

M. J. van der Laan, E. C. Polley, and A. E. Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6(1), 2007.

D. Whitney, A. Shojaie, and M. Carone. Comment: Models as (deliberate) approximations. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 34(4):591, 2019.