Targeted Maximum Likelihood Estimation: *Evaluation of the effects of longitudinal interventions including dynamic regimes*

Maya Petersen

Graduate Program in Biostatistics
University of California, Berkeley

March 15 2019

Motivation: Global Burden of HIV

- High HIV prevalence in Sub-Saharan Africa
- Limited financial and human resources

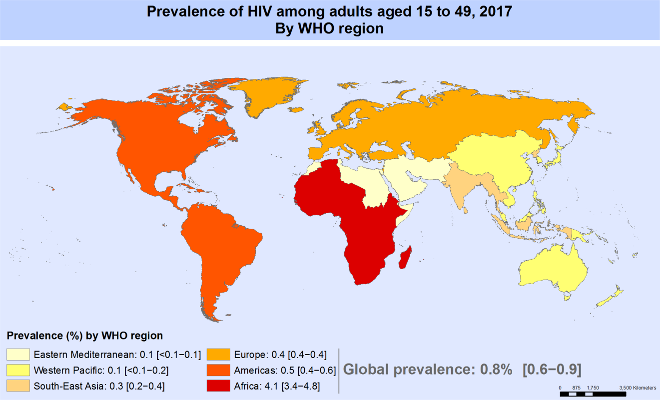

The boundaries and names shown and the designations used on this map do not imply the expression of any opinion whatsoever on the part of the World Health Organization concerning the legal status of any country, territory, city or area or of its authorities, or concerning the delimitation of its frontiers or boundaries. Dotted and dashed lines on maps represent approximate border lines for which there may not yet be full agreement.

Data Source: World Health Organization
Map Production: Information Evidence and Research (IER)
World Health Organization

World Health Organization

Retention in HIV Care in East Africa

Background

- Loss to HIV care is common in Sub-Saharan Africa
- Loss to care (retention failure) is associated with high mortality

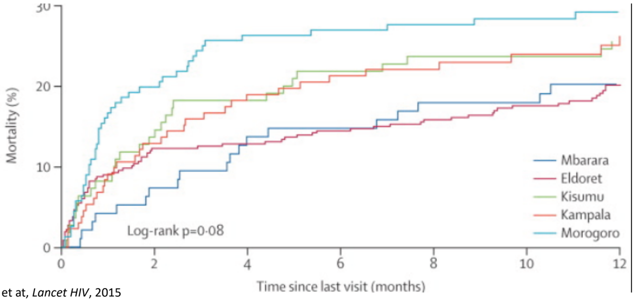

Geng et at, *Lancet HIV*, 2015

Outline: Case studies of causal inference methods to improve retention in HIV care in East Africa

1. Example 1: Effect of nurse-based triage on retention in HIV care (Tran et al., 2016)
   - The Causal Roadmap: Review of TMLE for point treatment effects
   - Extension to longitudinal interventions- LTMLE
     - Implementation choices
     - Data and simulation results
     - Challenges and ongoing work
2. Example 2: Adaptive behavioral interventions to improve retention in HIV care (Petersen et al., 2016)
   - LTMLE to evaluate dynamic regimes (adaptive treatment strategies)
     - Effects of longitudinal dynamic regimes
     - Estimating optimal dynamic regimes

Example 1. Low Risk Express Care (LREC)

- LREC: Task-shifting HIV care for clinically stable "low risk" patients from clinicians to nurses
  - USAID- AMPATH partnership; IeDEA- East Africa
  - Implemented in 15 clinics in Kenya 2007-2008
- **Impact of enrollment into LREC on loss-to-follow up/death?**
  - Clinical cohort data: Subset of eligible "low risk" patients enrolled at varying (non-random) times following eligibility

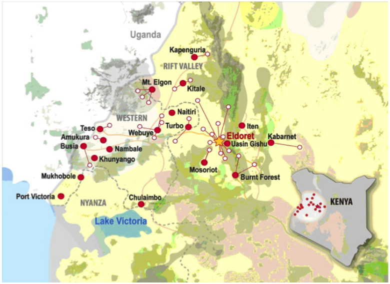

The Causal Roadmap

1. Specify Causal Question
   - As a parameter of counterfactual distributions
2. Specify Observed Data and Statistical Model
   - Statistical Model: Set of possible observed data distributions
3. Identify
   - Translate causal parameter into parameter of observed data distribution (estimand)
   - Under explicit casual assumptions (expressed in language of graphs or counterfactuals)
4. Estimate
   - Estimand + Statistical Model= Statistical Estimation Problem
   - Multiple estimators: IPTW, parametric G-computation, Double robust (including TMLE)
   - Different estimators $\rightarrow$ different statistical properties

see, e.g. Petersen and van der Laan (2014)

Causal Question: Point Treatment Example

1. Time scale
   - 90 day time scale; Baseline: First date eligible for LREC
2. Intervention (a.k.a. exposure or treatment): A
   - A: Indicator of immediate enrollment in LREC program
3. Counterfactual outcomes: $Y(a)$
   - $Y(1)$: Counterfactual retention status at 18 months under immediate enrollment
   - $Y(0)$: Counterfactual retention status at 18 months under deferred enrollment
4. Target Causal Parameter: Ex. $\mathbb{E}[Y(1) - Y(0)]$:
   - Difference in proportion lost to care if **all enrolled immediately** vs. **all deferred enrollment**
   - Focus here on $\mathbb{E}[Y(a)]$

Specify Observed Data and Statistical Model

- Observed Data: n=15,225 i.i.d. copies of $O_i = (W_i, A_i, Y_i) \sim P_0$
  - Baseline covariates W: age, sex, CD4 pre-ART, urban/rural,...
  - Treatment A: Indicator of immediate enrollment in LREC
  - Outcome Y: Lost to care at 18 months (death=fail)
- Statistical Model $\mathcal{M}$: $P_0 \in \mathcal{M}$
  - Model should reflect real knowledge: large enough to contain the true P_0
  - Probability distribution P of O can be factorized as:

$$P(O) = P(W)P(Y|A, W)P(A|W)$$

  - Often: Model places restrictions, if any, only on $P(A|W)$ propensity score or treatment mechanism

Identify

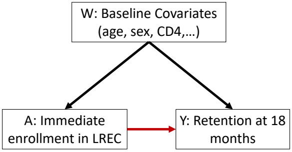

1. Randomization assumption $Y(a) \perp\!\!\!\perp A|W$
   - Baseline covariates sufficient to control for confounding
   - Holds if W blocks all backdoor paths $A \rightarrow Y$ (eg, Pearl (1995))
2. Positivity: $P(A = a|W) > 0$ for $a \in \{0, 1\}$
   - Ex. Violation if sickest patients never enroll immediately

Under these assumptions, can express casual parameter as a statistical parameter (estimand)

$$\mathbb{E}[Y(a)] = \sum_w \mathbb{E}[Y|A = a, W = w]P(W = w)$$

Estimate

Three general classes of estimator:

1. Propensity score-based
   - For example, Inverse Probability of Treatment Weighted (see eg., Robins and Rotnitzky (1992); Hernán et al. (2006))
2. Outcome Regression-based
   - For example, Parametric G-computation (see eg, Robins (1986))
3. Double robust
   - For example, Targeted Maximum Likelihood (see eg, van der Laan and Rose (2011))

Inverse Probability of Treatment Weighting (IPTW)

- Estimate the treatment mechanism: $P(A|W)$:
  - Ex. probability of immediate enrollment given baseline covariates
  - Classically: Based on parametric regression model (eg logistic regression)
    - Susceptible to bias due to model mis-specification
- IPTW Estimator

$$I\hat{PTW} = \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbb{I}(A_i = a)Y_i}{\hat{P}(A|W_i)}$$

  - or stabilized counterpart
  - $\hat{P}(A|W)$ is estimated propensity score
- Additional Limitations:
  - High variance
  - Unstable/biased in settings of strong confounding

Parametric G-computation

- Estimate the outcome regression: $\mathbb{E}(Y|A,W)$
  - Ex: Probability lost to care given enrollment and covariates
  - Based on parametric regression model (eg. logistic regression)
    - Susceptible to bias due to model misspecification
- Marginal distribution of W estimated using the empirical distribution
- Parametric G computation Estimator:

$$G\hat{comp} = \frac{1}{n}\sum_{i=1}^{n}\hat{\mathbb{E}}(Y|a,W_i)$$

  - $\hat{\mathbb{E}}(Y|a,W_i)$ is estimated outcome regression

Targeted Maximum Likelihood Estimation- Motivation & Overview

- Machine Learning (e.g. Super Learning) to generate an initial (^{0}) estimate of the outcome regression $\hat{\mathbb{E}}^0(Y|A,W)$
  - Avoid bias due to mis-sepcified parametric models
    - Could just 'plug-in" resulting estimate:

$$\frac{1}{n}\sum_{i=1}^{n}\hat{\mathbb{E}}^0(Y|a,W_i)$$

    - But... not good for inference (95% CI, p values...)
- **Instead: TMLE updates initial estimate** of outcome regression $\hat{\mathbb{E}}^0(Y|A,W)$ to obtain targeted estimate $\hat{\mathbb{E}}^*(Y|A,W)$
- Targeting step uses estimate of propensity score $\hat{P}(A|W)$ to provide opportunity to
  - reduce asymptotic bias if initial $\hat{\mathbb{E}}^0(Y|A,W)$ not consistent
  - reduce finite sample bias
  - reduce variance

A brief introduction to Super Learning

“Ensemble” Machine Learning approach (van der Laan et al., 2007; Breiman, 1996)

- Competition of algorithms
  - Parametric regression models
  - Data-adaptive (ex. Random forest, Neural nets)
- Best team wins
  - Convex combination of algorithms
- Performance judged on independent data: V-fold cross validation (Internal data splits)
  - Partition the data into “folds”
  - Fit each algorithm on the training set
  - Evaluate its performance on the validation set

Ex: 10-fold cross-validation

- Rotate through the folds
- Average performance estimates across the folds
- Choose the algorithm (or "team") with the best performance

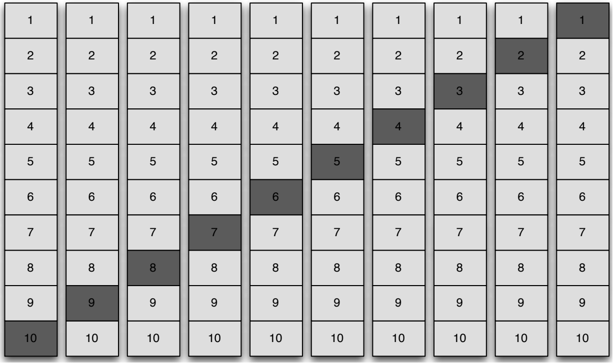

TMLE Algorithm for $\mathbb{E}[Y(a)]$

1. Obtain initial estimate of the outcome regression: $\hat{\mathbb{E}}^0(Y|A,W)$
2. Target (update) the initial estimate (logit scale)

$$\hat{\mathbb{E}}^*(Y|A,W) = \hat{\mathbb{E}}^0(Y|A,W) + \hat{\epsilon}$$

   - Maximum likelihood to fit ϵ: Logistic regression of Y on intercept, using $\hat{\mathbb{E}}^0(Y|A,W)$ as offset and weights $\frac{\mathbb{I}(A=a)}{\hat{P}(A|W)}$
   - Update model constructed to ensure that fitting ϵ solves the efficient influence curve (EIC) estimating equation (confers double robustness)
3. Plug in ("targeted") estimate of outcome regression:

$$T\hat{M}LE = \frac{1}{n}\sum_{i=1}^{n} \hat{\mathbb{E}}^*(Y|a,W_i)$$

Targeted Maximum Likelihood Estimation: Properties

- Double Robust
  - Consistent if either $\mathbb{E}(Y|A, W)$ or $P(A|W)$ estimated consistently
- Efficient
  - Lowest (asymptotic) variance among reasonable estimators if both $\mathbb{E}(Y|A, W)$ AND $P(A|W)$ estimated consistently at reasonable rates
- Can incorporate Machine Learning
  - To estimate $\mathbb{E}(Y|A, W)$ AND $P(A|W)$ while maintaining valid statistical inference (meaningful p values and confidence intervals)
  - Not a guarantee- still need estimators of these quantities to converge fast enough
- Substitution (aka "plug in") Estimator
  - Improved robustness to sparse data compared to estimating equation alternatives

Example code: ltmle R package
TMLE for point treatment: $\mathbb{E}[Y(0)]$

Schwab et al. (2013); link to ltmle vignette

```
> data
           W A Y
1 -1.2070657 0 1
2  0.2774292 0 0
3  1.0844412 1 1
> r <- ltmle(data,Anodes = "A",Ynodes = "Y",abar = 0)

> summary(r)
Estimator:  tmle
ltmle(data = data,Anodes = "A",Ynodes = "Y",abar = 0)
   Parameter Estimate:  0.50682
    Estimated Std Err:  0.0075484
              p-value:  <2e-16
    95% Conf Interval: (0.49203, 0.52162)
```

Example code: ltmle R package
IPTW and G-comp for point treatment: $\mathbb{E}[Y(0)]$

Schwab et al. (2013); link to ltmle vignette

```
> summary(r, estimator = "iptw")
Estimator:  iptw
Call:
ltmle(data = data,Anodes = "A",Ynodes = "Y",abar = 0)

   Parameter Estimate:  0.50285
    Estimated Std Err:  0.0082819
              p-value:  <2e-16
    95% Conf Interval: (0.48662, 0.51908)

> ltmle(data,Anodes = "A",Ynodes = "Y",abar = 0,
        gcomp = TRUE)
GCOMP Estimate:  0.5038029
```

Example code: ltmle R package
TMLE for point treatment: $\mathbb{E}[Y(1) - Y(0)]$

Schwab et al. (2013); link to ltmle vignette

```
> r <- ltmle(data, Anodes = "A",
             Ynodes = "Y", abar = list(1, 0))
> summary(r)
Estimator:  tmle

Additive Treatment Effect:
   Parameter Estimate:  0.19383
    Estimated Std Err:  0.010055
              p-value:  <2e-16
    95% Conf Interval: (0.17412, 0.21354)

Relative Risk:
   Parameter Estimate:  1.3824
  Est Std Err log(RR):  0.017493
              p-value:  <2e-16
    95% Conf Interval: (1.3358, 1.4307)
```

Beyond single time point static interventions...

Extending the roadmap to more complex causal questions

1. Effects of multiple interventions
   - longitudinal interventions
2. Effects of adaptive interventions
   - dynamic regimes

Beyond single time point static interventions...

Extending the roadmap to more complex causal questions

1. **Effects of multiple interventions**
   - **Longitudinal interventions**
2. Effects of adaptive interventions
   - dynamic regimes

Effects of multiple interventions

- **Motivating causal question:** Effect of enrollment into LREC on retention?
- **Effect of a single time point treatment:**
  Ex. $\mathbb{E}[Y(1) - Y(0)]$: Difference in retention (loss to care) if all eligible did vs. did not enroll immediately in LREC
  - Effect of a decision or action at a single time point
  - **But wait...** Counterfactual $Y(0)$: Retention status if did not enroll immediately (first 90 days)
    - **Could have enrolled after 90 days...**
- What if we want to know about the effect of enrolling immediately versus **never** enrolling?
  - Requires intervention at multiple time points: don't enroll in first 90 days *or* in second 90 days *or*...

Longitudinal Observed Data

Longitudinal data are also more complex

- Discrete time scale: 90 days (clinic visit interval)
  - $t = 0, \ldots, 6$ (18 months)
- **Covariates** W_t:
  - **Baseline:** age, sex, CD4 pre-ART, urban/rural,...
  - **Time-varying:** recent and nadir CD4, ART regimen, adherence, TB, pregnancy, ...
- **Outcome** Y_t: Indicator lost to care (or died) by t
- **Exposure** E_t: Indicator enrolled in LREC program by t
- **Right censoring** C_t: Indicator transferred to clinic with no LREC program (or database closure) by t

Notation: Longitudinal Observed Data

- **"Non-intervention" nodes:** $L_t = (Y_t, W_t)$
  - $\bar{L}_t = L_0, \ldots, L_t$
- **"Intervention" nodes:** $A_t = (E_t, C_t)$
  - $\bar{A}_t = A_0, \ldots, A_t$
  - Censoring treated as an additional "intervention" node: evaluate effect of enrollment in the absence of censoring
- We observe $n = 15,225$ i.i.d. copies of

$$O = (L_0, A_0, \ldots, L_5, A_5, L_6) = (\bar{L}_6, \bar{A}_5) \sim P_0$$

Statistical Model

- Probability distribution P of O can be factorized as:

$$P(O) = \prod_{t=0}^{6} P(L_t|\bar{L}_{t-}, \bar{A}_{t-}) \prod_{t=0}^{5} P(A_t|\bar{L}_t, \bar{A}_{t-})$$

- Statistical model places restrictions, if any, only on treatment mechanism

Target Causal Parameter: Multiple interventions

- **Intervention-Specific Mean**: Probability lost to care by 18 months ($t = 6$) if
  - Never enrolled in LREC and censoring prevented: $\mathbb{E}[Y_6(\bar{e} = 0, \bar{c} = 0)] \equiv \mathbb{E}[\mathbf{Y_6}(\bar{\mathbf{0}})]$
  - Enrolled immediately in LREC and censoring prevented: $\mathbb{E}[(Y_6(\bar{e} = 1, \bar{c} = 0)] \equiv \mathbb{E}[\mathbf{Y_6}(\bar{\mathbf{1}})]$

- **Average Treatment Effect**: Difference in probability lost to care by 18 months if enrolled immediately vs never enrolled (and censoring prevented):
  - $\mathbb{E}[Y_6(\bar{1}) - Y_6(\bar{0})]$

Identification for longitudinal treatments

Causal graph (simplified for illustration)

Causal Effects of Interest

- Including effects mediated by interim health

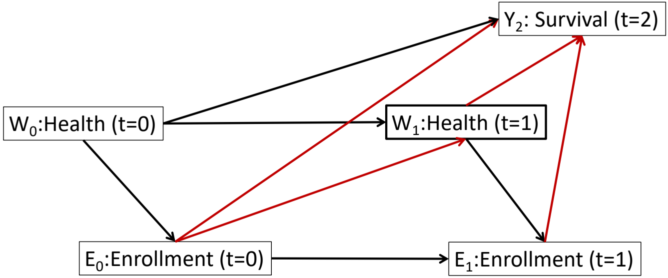

The challenge of time-dependent confounding

Causal Effects of Interest

- Including effects mediated by interim health

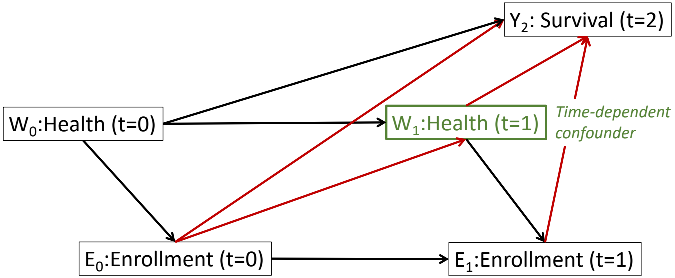

- Covariates needed to block back-door paths are affected by earlier exposure (see, e.g. Robins (1989))

Identification Assumptions (1)

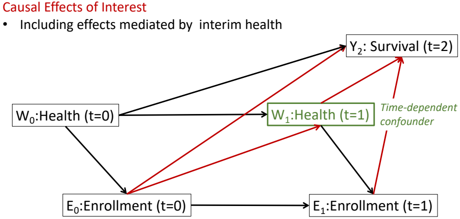

- **Sequential randomization** (Robins, 1989)

$$Y_6(\bar{a}) \perp\!\!\!\perp A_t \big| \bar{L}_t, \bar{A}_{t-1} : t = 0, \ldots, 5$$

  - Apply back door criteria to each intervention node in sequence (Pearl and Robins, 1995)

Identification assumptions (2)

- **Positivity**

$$P(A_t = a_t | \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_t) > 0, t = 0, \ldots, 5$$

for all regimes of interest
  - Ex: Needs to hold for $\bar{a} \in \{(\bar{e} = 1, \bar{c} = 0), (\bar{e} = 0, \bar{c} = 0)\}$
- Example: Positivity violation
  - Patients who lose eligibility have zero probability of enrolling
  - Regimes such as "enroll two time points after eligibility" would not be supported

Longitudinal G computation formula

- Under sequential randomization and positivity, the intervention-specific mean outcome is identified as (Robins, 1986):

$$E(Y_6(\bar{a})) = \sum_{\bar{l}_5} \left(E(Y_6 | \bar{A}_5 = \bar{a}_5, \bar{L}_5 = \bar{l}_5) \prod_{t=0}^{5} P(l_t | \bar{A}_{t-1} = \bar{a}, \bar{L}_{t-1} = \bar{l}_{t-1}) \right)$$

- Analog to point treatment, uses expectation of outcome conditional on exposure and confounding covariate *history*
  - Ex. Probability of loss to care by 18 months given uncensored, never enrolled, and full covariate history
- Because some of these covariate values affected by earlier exposure, now need to "standardize" to a different distribution of covariates
  - The "post-intervention" covariate distribution
  - Ex. the values the time-varying covariates would have had if never censored and never enrolled

Parametric G-computation

- Estimate the components of the longitudinal G-computation formula directly
  - Non-intervention factors of the likelihood: Conditional distributions (densities) of non-intervention covariates given the past

$$P(O) = \prod_{t=0}^{6} P(L_t|\bar{L}_{t-}, \bar{A}_{t-}) \prod_{t=0}^{5} P(A_t|\bar{L}_t, \bar{A}_{t-})$$

- Classically, based on parametric regression models
  - Susceptible to bias due to model mis-specification

(Robins, 1986)

Inverse Probability of Treatment Weighting (IPTW)

- Estimate the treatment mechanism
  - Treatment mechanism: Conditional probability of exposure and censoring given the past

$$P(O) = \prod_{t=0}^{6} P(L_t | \bar{L}_{t-}, \bar{A}_{t-}) \prod_{t=0}^{5} P(A_t | \bar{L}_t, \bar{A}_{t-})$$

  - Ex. For each time point ($t = 0, ...5$), estimate
    - Probability enroll in LREC given not already enrolled, uncensored, and past covariates
    - Probability remain uncensored given enrollment history, previously uncensored, and past covariates
- Based on parametric regression models
  - Susceptible to bias due to model mis-specification
- Data-adaptive/Super Learning methods
  - Challenges for inference

(Robins and Rotnitzky, 1992; Hernán et al., 2006)

Alternative representation of the longitudinal G-computation formula

- Can rewrite longitudinal G-comp formula using **iterated conditional expectations (ICE)** (Robins, 2000; Bang and Robins, 2005):

$$\mathbb{E}\left[\ldots\left[\mathbb{E}\left[\mathbb{E}\left[Y_6|\bar{L}_5,\bar{A}_5=\bar{a}_5\right]|\bar{L}_4,\bar{A}_4=\bar{a}_4\right]\right]\ldots\right]$$

- Basis for alternative parametric G-computation and double robust estimators
- **Advantage:** Lower dimensional set of "non-intervention factors"
  - Series of conditional expectations vs. conditional densities
  - Easier to estimate well

ICE G-computation Estimator

Parametric regression models to estimate series of conditional expectations (nested outcome regressions) (Robins, 2000; Bang and Robins, 2005)

1. Estimate inner most conditional expectation ($t = 6$)
   - Regress outcome Y_6 on past $(\bar{A}_5, \bar{L}_5)$
   - Generate predicted values by evaluating at $\bar{A}_5 = \bar{a}_5$
2. Estimate next conditional expectation ($t = 5$)
   - Use predicted values from prior step as new "outcome"
   - Regress on past $(\bar{A}_4, \bar{L}_4)$
   - Generate predicted values by evaluating at $\bar{A}_4 = \bar{a}_4$
3. Repeat for $t = 4, ..., 1$
4. Take empirical mean

Longitudinal TMLE

Properties

- **Double robust:** Consistent if either ICEs or treatment mechanism: $\prod_{t=0}^{5} P(A_t|\bar{L}_t, \bar{A}_{t-1})$ estimated consistently
- **Efficient** in semiparametric statistical model if both estimated consistently (at reasonable rates)
- Can incorporate **Machine Learning**: But care needed- more coming up...
- **Substitution estimator:** i.e. 'plug-in" estimator; may perform better in sparse data settings

Robins (2000); Bang and Robins (2005); Robins et al. (2007); van der Laan and Gruber (2012)

TMLE Algorithm for $\mathbb{E}[Y(\bar{a})]$

Analog to the ICE G-comp estimator, with two differences

1. Can generate initial estimate of each conditional expectation (i.e. iterated outcome regression) using machine learning
2. Before fitting the next conditional expectation, update the initial fit
   - Approach analogous to single time point TMLE
   - Update uses an inverse propensity score-based weight
   - Confers double robustness properties

L-TMLE Algorithm: Example

1. For inner-most conditional expectation ($t = 6$): $\mathbb{E}\left[Y_6|\bar{L}_5, \bar{A}_5 = \bar{a}_5\right]$
    1. Generate initial estimate
        - Using Super Learning
    2. Update initial estimate (as for single point)
        - Use MLE to fit an intercept only logistic regression
        - Initial fit as offset
        - Using weights $\mathbb{I}(\bar{A}_5 = \bar{a}_5)/\prod_{j=0}^{5} \hat{P}(A_j|\bar{L}_j, \bar{A}_{j-1})$
        - Treatment mechanism can be estimated using Super Learning
2. Repeat for next conditional expectation ($t = 5$)...
    1. Generate initial fit using predicted value from prior step as "outcome"
    2. Update, using weight $\mathbb{I}(\bar{A}_4 = \bar{a}_4)/\prod_{j=0}^{4} \hat{P}(A_j|\bar{L}_j, \bar{A}_{j-1})$
3. Repeat for $t = 4, ..., 1$
4. Take empirical mean

ltmle R package: Effects of multiple interventions

Syntax

- "Anodes": treatment or exposure nodes
  - LREC Example: Enrollment in LREC: E_t, $t = 0, ..., 5$
- "Cnodes": Indicator of right censoring
  - LREC Example: Transfer to new clinic by time t: C_t, $t = 0, ..., 5$
- "Lnodes": Time varying covariates
  - LREC Example: CD4 count, etc. at time t: W_t, $t = 1, ..., 5$
- "Ynodes": Outcome or outcomes
  - LREC Example: Indicator lost to care by t: Y_t, $t = 1, ..., 6$

Example R code: Estimation of $\mathbb{E}[Y(\bar{0})]$ in ltmle package


```
> head(data)
           W A1          L A2 Y
1 -1.3435214  0 -1.4164248  0 0
2  0.6217756  1  1.0621048  1 1
3  0.8008747  1  0.2808690  1 0
4 -1.3888924  0 -0.8677043  0 0
5 -0.7143569  1 -0.9064954  1 0
6 -0.3240611  1  0.7103158  0 0

> ltmle(data, Anodes = c("A1", "A2"), Lnodes = "L",
        Ynodes = "Y", abar = c(0, 0))
TMLE Estimate:  0.5128132
```

Example R code: Estimation $\mathbb{E}[Y(\bar{a})]$ for $\bar{a} = (1, 0)$ in ltmle package, with censoring


```
> head(data)
           W A1          C         L A2  Y
1  1.3514112  1   censored        NA NA NA
2  0.1854795  1   censored        NA NA NA
3  0.4315265  0 uncensored 0.1251185  0  0
4 -0.1906075  1   censored        NA NA NA
5 -0.9715509  1 uncensored 0.3115363  1  0
6  0.7680671  1 uncensored 0.6744166  0  1
#set all A1 to 1, set all A2 to 0,
#set C to uncensored, use glm
> ltmle(data, Anodes = c("A1", "A2"), Cnodes = "C",
        Lnodes = "L", Ynodes = "Y", abar = c(1, 0))
TMLE Estimate:  0.4704012
```

link to ltmle vignette

Example R code: Estimation $\mathbb{E}[Y(\bar{a})]$ for $\bar{a} = (1, 0)$ in ltmle package, using SuperLearner


```
#set all A1 to 1, set all A2 to 0,
#set C to uncensored, use default SuperLearner library
> ltmle(data, Anodes=c("A1", "A2"), Cnodes = "C",
        Lnodes="L", Ynodes="Y", abar = c(1, 0),
        SL.library = "default")
TMLE Estimate:  0.4692075
```

link to ltmle vignette

Example R code: Additive Treatment Effect and Relative Risk


```
> result <- ltmle(data, Anodes=c("A1", "A2"), Cnodes = "C",
     Lnodes="L", Ynodes="Y", abar=list(c(1, 0), c(1, 1)))
> summary(result)
Treatment Estimate:
   Parameter Estimate:  0.42744
    Estimated Std Err:  0.086301
              p-value:  7.3109e-07
    95% Conf Interval: (0.2583, 0.59659)

Control Estimate:
   Parameter Estimate:  0.29593
    Estimated Std Err:  0.046223
              p-value:  1.5315e-10
    95% Conf Interval: (0.20533, 0.38653)

Additive Treatment Effect:
   Parameter Estimate:  0.13151
    Estimated Std Err:  0.097835
              p-value:  0.17887
    95% Conf Interval: (-0.06024, 0.32327)

Relative Risk:
   Parameter Estimate:  1.4444
  Est Std Err log(RR):  0.25507
              p-value:  0.14943
    95% Conf Interval: (0.87614, 2.3812)
```

link to ltmle vignette

Challenge: Estimation of treatment mechanism and outcome regressions

Need that initial fits of the outcome regressions not be too overfit

- Internal sample splitting approaches relax this (Zheng and van der Laan, 2011)
- Not implemented in ltmle package (yet!)
- **Be careful of default in package**
  - Default: logistic regression (glm) with all past variables as main terms
  - If using a parametric model for treatment mechanism and outcome regressions, specify carefully and consider *a priori* reduction in adjustment variables
    - Ex. Background knowledge (eg most recent values of time- varying covariates)
    - Ex. Marginal association with the outcome

Challenge: Estimation of treatment mechanism and outcome regressions

- DR estimators make it possible to use machine-learning approaches to estimate treatment mechanism and outcome regressions
  - Doesn't guarantee they will work well enough
- If using Super Learning (or other machine learning) to estimate treatment mechanism and outcome regressions, need estimates to converge to truth fast enough
  - If can estimate treatment mechanism with a correctly specified parametric model (e.g. an RCT), then just need estimators of outcome regressions to be consistent
  - Remains a challenge in high dimensional data
    - Some progress on this front: Highly Adaptive LASSO (van der Laan, 2017)
- Choose your machine learning library carefully
  - see eg, Schomaker et al. (2018); Tran et al. (2010, 2016)

Challenge: "Practical" positivity violations

- Poor support for regime (treatment history) of interest
  - Ex: $\prod_{t=0}^{5} P(A_t = \bar{a}_t | \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_t)$ is small
  - Problem increases with increasing number of time points
  - Ex: Small probability of not enrolling given healthy at each time point $\rightarrow$ product can get very small
- Can lead to both bias and underestimates of variance (see eg Petersen et al. (2012, 2014); Tran et al. (2010))

Some (partial) responses (defaults in ltmle package)

1. Use a substitution estimator (G-computation, TMLE)
   - But a challenge for all estimators
2. Use robust variance estimator (Tran et al., 2018)
   - "blows up" when confidence intervals become unreliable
3. Bound estimated propensity score away from 0

Simulations: In care survival if never enroll

- Correctly specified parametric models to estimate iterated outcome regressions and treatment mechanism
  - Positivity violations increase with increasing time points
- Choice of estimator can make a difference

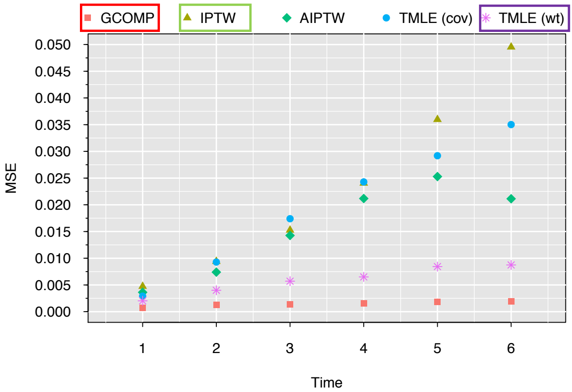

Real Data: Effect of Low risk Express Care

- TMLE+Super Learning to estimate propensity scores and outcome regressions
- Results: LREC enrollment appears to improves retention outcomes
  - Results consistent with better control of confounding by TMLE if patients who become sick less likely to enroll

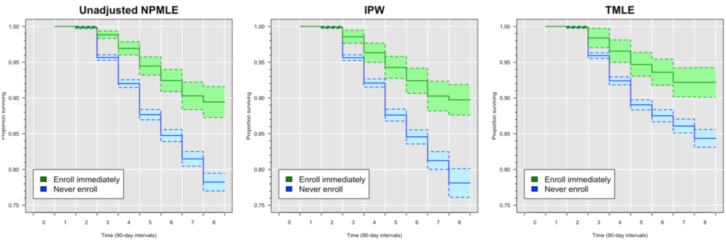

Beyond multiple time point static interventions...

Extending the roadmap to more complex causal questions

1. Effects of multiple interventions
   - longitudinal interventions
2. **Effects of adaptive interventions**
   - **Longitudinal dynamic regimes**

HIV+ persons face diverse barriers to retention

- Structural (eg. transport too expensive)
- Psycho-social (eg. patient-clinic interactions)
- Medical (e.g. too sick to travel to clinic)

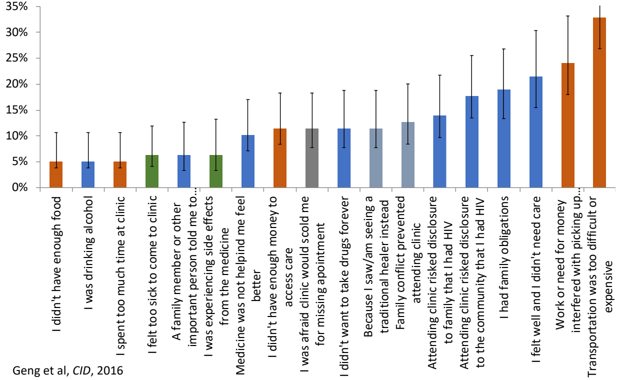

Geng et al, *CID*, 2016

Dynamic regimes to optimize retention in HIV care

Motivation:

- **Several behavioral interventions with proven efficacy** (compared to standard-of-care):
  - SMS text messages: reminders and support
  - Travel Vouchers: small conditional cash incentives for on-time visits
  - Peer Navigators: relationship-based support for overcoming barriers to care
- Hypothesis: Any **one-size-fits-all approach** will be
  - **Inefficient** - many patients will do well with no intervention
  - **Sub-optimally effective** - failing to help some in need by assigning them an intervention less likely to work for them

Precision Medicine/Public Health: The challenge

- **Objective:** Improve effectiveness and efficiency by targeting interventions based on individual characteristics
- **Dynamic regime:** A rule for assigning and modifying an intervention based on evolving individual characteristics
- **Ex. Target causal parameters:**
  - Expected outcome under a specific longitudinal regime
    - *Mean outcome if all subjects had followed a given rule?*
  - Optimal dynamic regime
    - *What rule would result in best mean outcome if all subjects followed it?*
  - Expected outcome under optimal regime
    - *Mean outcome if all subjects followed optimal rule (compared to some alternative)?*

ADAPT-R Trial: Adaptive strategies to improve retention in HIV Care

- Sequential Multiple Assignment Randomized Trial (NCT02338739; PIs: Geng, Petersen)
- 1800 HIV patients initiating ART in Kenya
- Objective: Develop and evaluate adaptive treatment strategies (aka "dynamic regimes") to optimize retention in HIV care

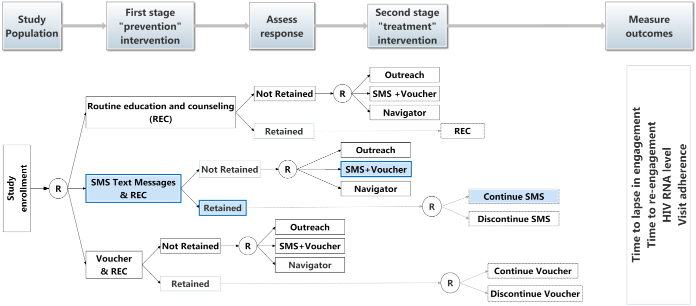

ADAPT-R Trial: Data and Model

- Data
  - **Baseline covariates** L_0:
    - V: Wealth
    - S_0: Patient satisfaction with care
  - 1^{st}**-line intervention** A_0: SMS, Voucher, Education
  - **Time-varying covariates** L_1:
    - Y_1: "Retention failure," 14 days late for visit
    - S_1: Updated satisfaction with care
  - 2^{nd}**-line Intervention** A_1:
    - If fail ($Y_1 = 1$): SMS+Voucher, Navigator, Outreach
    - If don't fail ($Y_1 = 0$): continue or stop 1^{st}-line
  - **Outcome** Y_2: Viral failure at year 2
- Statistical model makes assumptions only on g
  - Randomization: $g_0(A_0|L_0)$ and $g_0(A_1|L_0, A_0, L_1)$ known

Target Parameter: Regime-specific mean outcome

- **Decision rule:** $d_t(\bar{l}_t)$ assigns an "intervention" value a_t based on observed past at time t
- **Dynamic regime:** set of rules, one for each time point $d = (d_0, d_1 ...) \in \mathcal{D}$
  - ADAPT-R: Simple example of a rule d:
    - SMS at ART start

$$d_0 : A_0 = SMS$$

    - If 14 days late, escalate to Peer Navigator (Nav), otherwise stop SMS

$$d_1 : \text{If } Y_1 = 1 \text{ then } A_1 = Nav, \text{ else } A_1 = stop$$

- **Regime-Specific Mean** $\mathbb{E}(Y_2(d))$: Counterfactual probability of viral failure if followed rule d

Identification and Estimation

- **Identification assumptions**: Analogous to longitudinal "static" regime
  1. Sequential randomization
  2. Positivity
  - Both hold by design in sequentially randomized trials
- **Estimators**:Analogous to longitudinal "static" regime
  1. G-computation (including ICE version)
  2. IPTW
  3. LTMLE
  - Simply evaluate for treatment and covariate history that correspond to regime of interest
    - $\bar{A} = \bar{d}(\bar{L})$

Example R code: Estimate of $\mathbb{E}[Y(d)]$ for a simple data structure and regime d in ltmle package

Data: $W, A1, L, A2, Y$
Dynamic regime d of interest is:

- Always treat at time 1 (A1 $=$ 1)
- Treat at at time 2 (A2 $=$ 1) if L $>$ 0

```
> abar <- matrix(nrow=n, ncol=2)
> abar[, 1] <- 1
> abar[, 2] <- L > 0
> ltmle(data, Anodes=c("A1", "A2"),
        Lnodes="L", Ynodes="Y", abar=abar)
TMLE Estimate:  0.3061747
```

link to ltmle vignette

Optimal rule for assigning retention interventions?

- **Optimal Regime:** $d^{opt} \in \mathcal{D}$ that minimizes $E(Y_2(d))$
  - $E(Y_2(d))$: Probability fail at year 2 under rule d
- **Option 1:** Estimate $E(Y_2(d))$ for each d
  (e.g. Zhao and Laber (2014))
  - Requires each rule $d \in \mathcal{D}$ be supported
- **Option 2:** Dynamic Marginal Structural Working Model
  (Robins, 1999; Van der Laan and Petersen, 2007)
  - Lower dimensional summary of how $E(Y_2(d))$ varies as a function of d
  - Possibly conditional on a subset of baseline covariates V

Example: Marginal Structural Working Model

- Consider limited set $\mathcal{D}$ based on satisfaction threshold θ
  - $d_0^\theta(S_0)$: If $S_0 > \theta$ then Voucher, else SMS
  - $d_1^\theta(S_1)$:
    - If $Y_1 = 0$ then stop 1^{st}-line
    - If $Y_1 = 1$ and $S_1 > \theta$ then Voucher+SMS, else Navigator
  - $E(Y_2(\theta))$: Expected outcome under rule d^θ
- **Optimal threshold θ?**
  - **Does optimal threshold differ depending on wealth V?**
- Pose following working model $m_\beta(\theta, V)$ for $E_0(Y_2(\theta)|V)$:

$$m_\beta(\theta, V) = expit(\beta_0 + \beta_1\theta + \beta_2\theta^2 + \beta_3 V + \beta_4\theta V)$$

  - Working model-specific optimal regime given V:

$$\theta^*(V) \equiv \arg\min_\theta m_\beta(\theta, V) = \frac{\beta_1}{2\beta_2} - \frac{\beta_4}{2\beta_2}V$$

ADAPT-R: Marginal Structural Working Model

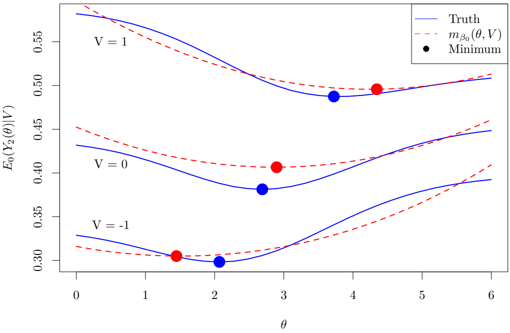

- Misspecified Model → Working model-specific optimal $\theta^*(V)$ may differ from true optimal θ^{opt}
  - Data adaptive estimation of the MSM (Petersen et al., 2016)

Ex. Estimators of Longitudinal Dynamic Marginal Structural Model Parameters

Analogous classes of estimator:

1. IPTW (Robins, 1999; Van der Laan and Petersen, 2007)
2. DRICE (Robins, 2000; Bang and Robins, 2005):
   - Double robust and semiparametric efficient
   - Uses sequential regression methodology
   - Defined as solution to estimating equation
3. LTMLE (Petersen et al., 2014)
   - Double robust and semiparametric efficient
   - Substitution estimator

- Implementation more complex
- Implemented in `ltmle` R package

Outcome under estimated optimal regime

- Estimate of β in MSM gives estimate of
  - Optimal threshold

$$\begin{aligned}\theta^*(V) \quad &\equiv \arg\min_{\theta} m_\beta(\theta, V) \\ &= \frac{\beta_1}{2\beta_2} - \frac{\beta_4}{2\beta_2} V \\ &= \alpha_0 + \alpha_1 V\end{aligned}$$

- Inference on expected outcome under optimal threshold $E(Y(\theta^*(V)))$ (Zhang et al., 2013)
  - Simply construct confidence interval for $E(Y(\theta))$, plugging in estimated optimal rule $\theta_n^*(V)$, and ignoring that it was estimated

Simulation: Covariate adjustment with TMLE reduces variance

- All estimators unbiased with good 95% CI coverage (Petersen et al., 2016)

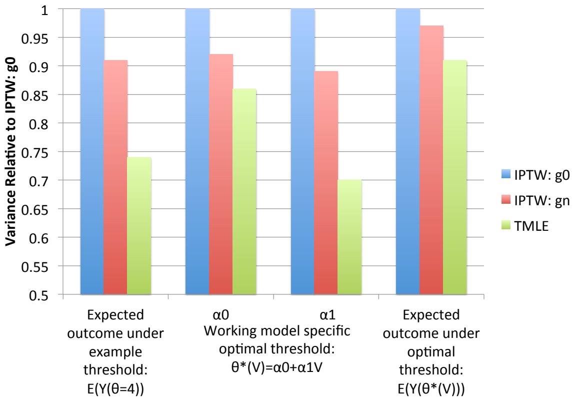

Summary: Longitudinal Dynamic Regimes

- Double Robust ICE Estimators incl. TMLE
  - Available for **regime specific mean, MSM parameters, optimal regime, and expected outcome under optimal regime**
  - Observational data: Reduce bias and variance
  - Sequentially randomized trials: Reduce variance
- Practical positivity violations
  - Ubiquitous in longitudinal data
  - Despite partial solutions: still a major concern
- Optimal dynamic regime (within a restricted class)
  - Directly or using marginal structural working model
  - Inference on both the optimal rule and expected outcome under optimal rule

`ltmle` R package

- Causal effect estimation with multiple intervention nodes
  - Intervention-specific mean under longitudinal static and dynamic interventions
  - Static and dynamic marginal structural working models
  - Controlled Direct Effects
- General longitudinal data structures
  - Repeated measures outcomes (including survival)
  - Right censoring
  - Hierarchical data
- Estimators
  - IPTW
  - ICE G-comp (no inference)
  - TMLE
- Options include nuisance parameter estimation via glm regression formulas or calling `SuperLearner()`

Acknowledgements

- Mark van der Laan, UC Berkeley
- Low risk Express Care
  - International Epidemiological Databases to Evaluate AIDS-East Africa (IeDEA-EA); AMPATH Eldoret, Kenya
  - Drs. Linh Tran, Constantin Yiannoutsos, Kara Wools Kaloustian, Abraham Siika, Sylvester Kimaiyo
  - The AMPATH Patients
- Adaptive Interventions to Prevent and Treat Lapses in Retention (AdaPT-R)
  - Drs. Elvin Geng, Thomas Odeny
- `ltmle` R package
  - Joshua Schwab
  - International Epidemiological Databases to Evaluate AIDS-Southern Africa (IeDEA-SA); Dr. Matthias Egger
- Sponsors
  - National Institutes of Health
  - President's Emergency Plan for AIDS Relief (PEPFAR)

References I

H. Bang and J.M. Robins. Doubly-robust estimation in missing data and causal inference models. *Biometrics*, 61:962 – 972, 2005.

L. Breiman. Bagging predictors. *Machine Learning*, 24:123 – 140, 1996.

M A Hernán, E Lanoy, D Costagliola, and J M Robins. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic & Clinical Pharmacology & Toxicology*, 98:237 – 242, 2006.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4): 669–688, 1995.

Judea Pearl and James M Robins. Probabilistic evaluation of sequential plans from causal models with hidden variables. In *UAI*, volume 95, pages 444–453. Citeseer, 1995.

M Petersen, J. Schwab, E Geng, and M van der Laan. Evaluation of longitudinal dynamic regimes with and without marginal structural working models. In Moodie E and Kosorok M, editors, *Dynamic Treatment Regimes in Practice: Planning Trials and Analyzing Data for Personalized Medicine.*, chapter 10, pages 157–186. ASA-SIAM, 2016.

References II

Maya L Petersen and Mark J van der Laan. Causal models and learning from data: integrating causal modeling and statistical estimation. *Epidemiology (Cambridge, Mass.)*, 25(3):418, 2014.

M.L. Petersen, K.E. Porter, S. Gruber, Y. Wang, and M.J. Van der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21:31 – 54, 2012.

M.L. Petersen, J. Schwab, S. Gruber, N. Blaser, M. Schomaker, and M. van der Laan. Targeted maximum likelihood estimation for dynamic and static marginal structural working models. *Journal of Causal Inference*, 2(2):DOI: 10.1515/jci–2013–0007, 2014.

J. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, pages 297–331. Springer, 1992.

James Robins. The control of confounding by intermediate variables. *Statistics in medicine*, 8(6):679–701, 1989.

References III

J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393 – 1512, 1986.

J.M. Robins. *Marginal Structural Models versus Structural Nested Models as Tools for Causal Inference*, volume 116 of *IMA*, pages 95 – 134. Springer, New York, NY, 1999.

J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association on Bayesian Statistical Science, 1999*, pages 6–10, 2000.

JM Robins, M Sued, Q Lei-Gomez, and A. Rotnitsky. Comment: Performance of double-robust estimators when "inverse probability" weights are highly variable. *Statistical Science*, 22(4):544–559, 2007.

M Schomaker, MA Luque-Fernandez, V Leroy, and M Davies. Using longitudinal targeted maximum likelihood estimation in complex settings with dynamic interventions. *arXiv preprint arXiv:1802.05005*, 2018.

References IV

J. Schwab, S. Lendle, M. Petersen, and M. van der Laan. *ltmle: Longitudinal Targeted Maximum Likelihood Estimation*, 2013. R package version 0.9.3, http://cran.r-project.org/web/packages/ltmle/.

L Tran, C Yiannoutsos, K Wools-Kaloustian, A Siika, M van der Laan, and M Petersen. Double robust efficient estimators of longitudinal treatment effects: Comparative performance in simulations and a case study. *International Journal of Biostatistics*, page To appear, 2010.

L Tran, Constantin T Yiannoutsos, B Musick, K Wools-Kaloustian, A Siika, S Kimaiyo, M van der Laan, and M Petersen. Evaluating the impact of a hiv low-risk express care task-shifting program: a case study of the targeted learning roadmap. *Epidemiologic Methods*, 5(1):69–91, 2016.

Linh Tran, Maya Petersen, Joshua Schwab, and Mark J van der Laan. Robust variance estimation and inference for causal effect estimation. *arXiv preprint arXiv:1810.03030*, 2018.

M J Van der Laan and M L Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *The International Journal of Biostatistics*, 3, 2007.

References V

Mark van der Laan. A generally efficient targeted minimum loss based estimator based on the highly adaptive lasso. *The international journal of biostatistics*, 13(2), 2017.

M.J. van der Laan and S. Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1):Article 8, 2012.

MJ van der Laan and S Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.

M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007.

Baqun Zhang, Anastasios A. Tsiatis, Eric B. Laber, and Marie Davidian. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694, 2013. doi: 10.1093/biomet/ast014. URL http://biomet.oxfordjournals.org/content/100/3/681.abstract.

References VI

Y. Zhao and E.B. Laber. Estimation of optimal dynamic treatment regimes. *Clinical Trials*, 11:400 – 407, 2014.

W Zheng and M van der Laan. Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning: Causal Inference for Observational and Experimental Data.*, chapter 27, pages 459–474. Springer, New York, 2011.