A Multiple Imputation Approach for Optimal Treatment Decision Rules

Jenny Shen, Kristin Linn

Department of Biostatistics, Epidemiology, and Informatics, University of Pennsylvania

Introduction

Data-driven optimal treatment strategies can benefit individual patients, care providers, and other stakeholders by improving outcomes and lowering healthcare costs. An optimal treatment decision rule maximizes a population-level distributional summary such as the expected value of a clinical outcome. Guidance for estimating optimal decision rules in the presence of missing data is limited, and most existing methods rely on having a complete set of data that are observed. We propose a multiple imputation framework for estimating optimal decision rules for data with missing at random (MAR) missingness using simulations and discuss guidance for reproducible inference. These finds are applied to data from the randomized trial Social Incentives to Encourage Physical Activity and Understand Predictors (STEP UP), which compared multiple interventions aimed at increasing daily step counts among employees at a large professional services company.

Methods

- Assume the complete data from n trial participants of a single-stage randomized trial have the form $\{X_i, A_i, Y_i\}_{i=1}^n$, where $X_i \in \mathcal{X} \subseteq \mathbb{R}^p$ is baseline patient information measured prior to randomization, $A_i \in \{-1, 1\}$ is the treatment assignment for the i^{th} individual, and Y_i is the clinical outcome, which we assume has been coded so that larger values are clinically desirable.

- A decision rule is a function that takes in information about an individual and outputs a recommended treatment for that person, i.e. $d(x)$ is a map $d: \mathcal{X} \to \{-1, 1\}$.
- The Value of a decision rule $d(x)$ is $V(d) = E^d(Y)$.
- An optimal decision rule, $d^*(x)$, is defined such that $V(d^*) \geq V(d)$ for all $d(x)$.
- One way to estimate an optimal decision rule is by specifying a model (e.g., linear) for the Q-function, $Q(x, a) \triangleq E(Y \mid X, A)$.

- Two common estimators for the Value are the inverse-probability weighted (IPW) and augmented inverse probability-weighted estimator (AIPW). Let $C_d = \mathrm{I}\{A = d(X)\}$ indicate whether an individual's assigned treatment A coincides with the recommended treatment under d and $\pi_d(X) = P(C_d = 1 \mid X)$. Then, the IPW estimator of $V(d)$ is $\hat{V}_{IPW}(d) = \frac{1}{n}\sum_{i=1}^n \frac{C_{d,i} Y_i}{\pi_d(X_i; \hat{\gamma})}$, and $\hat{V}_{AIPW}(d) = \frac{1}{n}\sum_{i=1}^n \left\{ \frac{C_{d,i} Y_i}{\pi_d(X_i; \hat{\gamma})} - \frac{C_{d,i} - \pi_d(X_i; \hat{\gamma})}{\pi_d(X_i; \hat{\gamma})} \widehat{Q_n}(X_i, a_i = d(X_i)) \right\}$.
- For fixed d, these estimators are asymptotically normal with variances of known form.

- In settings where the data are MAR, multiple imputation (MI) is a method for handling missingness that provides unbiased results and valid inference. Given $r = 1, \ldots, R$ training imputations we obtain a model-averaged decision rule $\hat{d}^*_{MA}(x)$ and estimate its Value by averaging estimates from R independent test sets: $\hat{V}_{MA} = \frac{1}{r}\sum_{i=1}^R \hat{V}_r$.
- The average within-imputation variance is then estimated by $\hat{\sigma}^2_W = \frac{1}{R}\sum_{r=1}^R \hat{\sigma}^2_{Value,r}$, where is an estimate of var($\hat{V}_r$). The between-imputation variance is estimated by the sample variance of the R Value estimates: $\hat{\sigma}^2_B = \frac{1}{R-1}\sum_{r=1}^R [(\hat{V}_r - \hat{V}_{MA})^2]$. Applying Rubin's Rules, our final estimate of the variance of the Value is $\hat{\sigma}^2_V = \hat{\sigma}^2_B + \frac{1}{R}\hat{\sigma}^2_B + \hat{\sigma}^2_W$.
- We adapted von Hippel's (2018) two-stage quadratic rule for determining the number of imputations that provides replicability of standard error estimates.

- Data generation for simulation studies (n=300 split 1:1 into training/testing):
  - $X = (1, X_2)^T$ where X_2 is a five-feature vector generated from a multivariate normal distribution with exchangeable correlation of 0.1
  - $A \sim Uniform(-1, 1)$
  - $Y = \gamma^{\mathrm{T}} X + A(\phi^T X) + \epsilon$, where $\epsilon \sim N(0, 5)$

Methods (contd.)

- Induced MAR, targeting 30% missingness with 3 scenarios: missingness only in outcome Y, missingness in outcome Y and a covariate X4, and weak MAR with missingness in outcome Y and all covariates
- Multiple imputation via chained equations (MICE) was used for addressing missingness, specifying r=5, 10, 25, 50, and 100 imputations. Simulations were run with 1000 iterations for each scenario.

Simulation Results

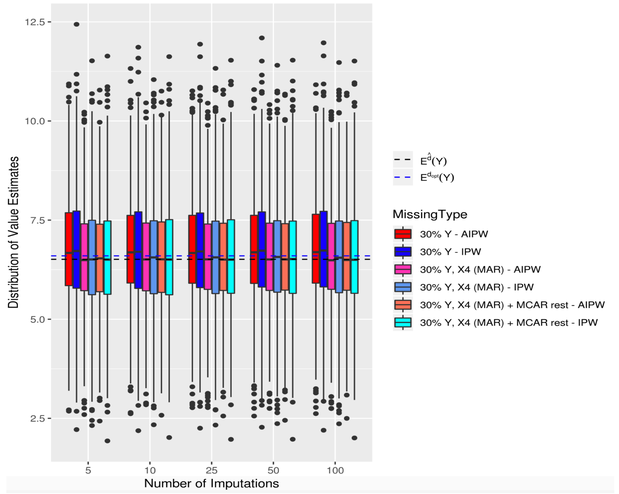

Figure 1. Distribution of Value estimates for differing MAR scenarios.

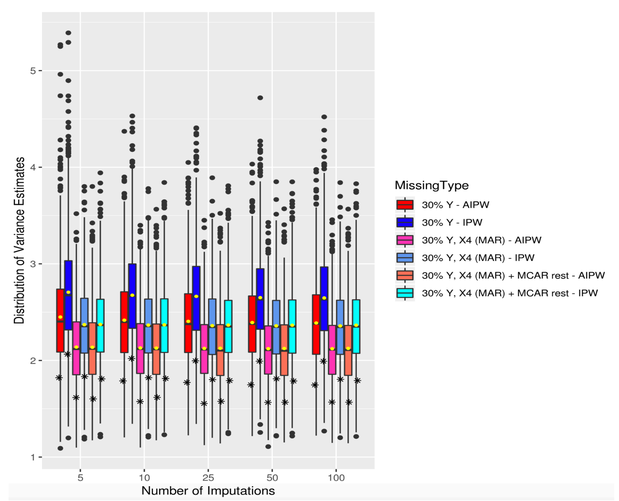

Figure 2. Distribution of variance estimates for differing MAR scenarios.

Real Data Analysis Results

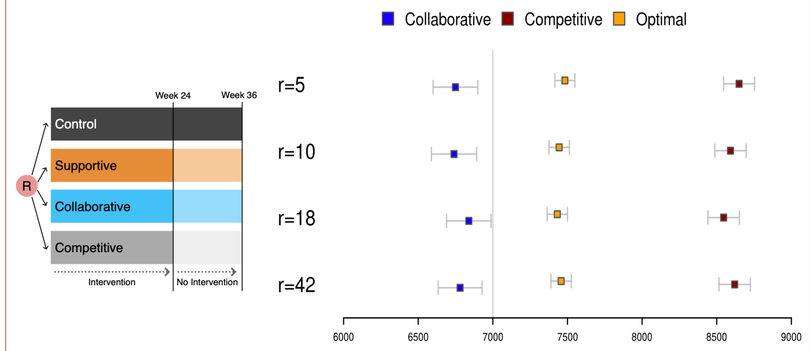

Figure 3. Diagram of STEP UP trial design (left) and estimates of Value with accompanying asymptotically normal 95% confidence intervals with the AIPW estimator for varying decision rules (right).

- Estimated a decision rule to recommend either collaborative or competitive using the AIPW estimator.
- Applied von Hippel's two-stage quadratic rule targeting a standard deviation of the standard error estimates of 0.5 across 100 replications.
- Using r=5 (r=10) imputations in the first stage of the procedure yielded r=42 (r=18) recommended imputations for achieving replicable standard errors.
- The estimated optimal decision rule based on Q-learning with linear models did not perform better than the AIPW estimate of the non-personalized rule that recommends everyone to the competitive intervention.

Conclusion

- For differing missingness scenarios and varying number of imputations, the AIPW estimator yielded higher estimates of Value and lower estimates of variance than its IPW counterpart.
- Variability of estimators of Value most notably decreased from r=5 to r=25 in our simulation studies, which suggests that defaulting to r=5 imputations would not have been sufficient for estimating Value in the presence of missing data.
- Differing recommended number of imputations from the two-stage quadratic rule procedure in the STEP UP analysis further demonstrate that concerns for reproducible standard errors can drive the choice of imputations when estimating optimal treatment decision rules.

References

- Chakraborty, B., Laber, E. B., & Zhao, Y. (2013). Inference for Optimal Dynamic Treatment Regimes Using an Adaptive m-Out-of-n Bootstrap Scheme. Biometrics, 69(3). https://doi.org/10.1111/biom.12052
- Chakraborty, B., Laber, E. B., & Zhao, Y. Q. (2014). Inference about the expected performance of a data-driven dynamic treatment regime. Clinical Trials, 11(4). https://doi.org/10.1177/1740774514537727
- Harrison, J. D. et al. Social incentives to encourage physical activity and understand predictors (STEP UP): Design and rationale of a randomized trial among overweight and obese adults across the United States. Contemp. Clin. Trials 80, (2019).
- Jakobsen, J.C., Gluud, C., Wetterslev, J. et al. When and how should multiple imputation be used for handling missing data in randomised clinical trials – a practical guide with flowcharts. BMC Med Res Methodol 17, 162 (2017). https://doi.org/10.1186/s12874-017-0442-1
- Patel, M. S. et al. Effectiveness of Behaviorally Designed Gamification Interventions with Social Incentives for Increasing Physical Activity among Overweight and Obese Adults Across the United States: The STEP UP Randomized Clinical Trial. JAMA Intern. Med. 179, (2019).
- Schomaker, M., Heumann, C. Bootstrap Inference When Using Multiple Imputation Statistics in Medicine, 37(14):2252-2266 http://dx.doi.org/10.1002/sim.7654
- von Hippel, P. T. (2018). How Many Imputations Do You Need? A Two-stage Calculation Using a Quadratic Rule. Sociological Methods and Research, 49(3). https://doi.org/10.1177/0049124117747303
- White, I. R., Royston, P., & Wood, A. M. (2011). Multiple imputation using chained equations: Issues and guidance for practice. Statistics in Medicine, 30(4). https://doi.org/10.1002/sim.4067