PROJECT MUSE®

Causaltoolbox—Estimator Stability for Heterogeneous Treatment Effects

Sören R. Künzel, Simon J. S. Walter, Jasjeet S. Sekhon

Observational Studies, Volume 5, Issue 2, 2019, pp. 105-117 (Article)

Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/obs.2019.0005>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/793360/summary>

Causaltoolbox—Estimator Stability for Heterogeneous Treatment Effects

Sören R. Künzel*

*Department of Statistics
University of California, Berkeley*

srk@berkeley.edu

Simon J. S. Walter*

*Department of Statistics
University of California, Berkeley*

sjswalter@berkeley.edu

Jasjeet S. Sekhon

*Department of Political Science & Department of Statistics
University of California, Berkeley*

sekhon@berkeley.edu

Abstract

Estimating heterogeneous treatment effects has become increasingly important in many fields: for example, they are required to select a personalized treatment for a patient, which may be a life or death decision. Recently, a variety of procedures relying on different assumptions have been suggested for estimating heterogeneous treatment effects. Unfortunately, there are no compelling approaches that allow identification of the procedure that has assumptions that hew closest to the process generating the data set under study and researchers often select one arbitrarily. This approach risks making inferences that rely on incorrect assumptions and gives the experimenter too much scope for p -hacking. A single estimator will also tend to overlook patterns other estimators could have picked up. We believe that the conclusion of many published papers might change had a different estimator been chosen and we suggest that practitioners should evaluate many estimators and assess their similarity when investigating heterogeneous treatment effects. We demonstrate this by applying 28 different estimation procedures to an emulated observational data set; this analysis shows that different estimation procedures may give starkly different estimates. We also provide an extensible R package which makes it straightforward for practitioners to follow our recommendations.

Keywords: Heterogeneous treatment effects, conditional average treatment effect, X-learner, joint estimation.

1. Introduction

Heterogeneous Treatment Effect (HTE) estimation is now a mainstay in many disciplines, including personalized medicine (Henderson et al., 2016; Powers et al., 2018), digital experimentation (Taddy et al., 2016), economics (Athey and Imbens, 2016), political science (Green and Kern, 2012), and statistics (Tian et al., 2014). Its prominence has been driven by a combination of the rise of big data, which permits the estimation of fine-grained heterogeneity, and recognition that many interventions have heterogeneous effects, suggesting that

*. These authors contributed equally to this work.

much can be gained by targeting only the individuals likely to experience the most positive response. This increase in interest amongst applied statisticians has been accompanied by a burgeoning methodological and theoretical literature: there are now many methods to characterize and estimate heterogeneity; some recent examples include Hill (2011); Athey and Imbens (2015); Künzel et al. (2017); Wager and Athey (2017a); Nie and Wager (2017). Many of these methods are accompanied by guarantees suggesting they possess desirable properties when specific assumptions are met, however, verifying these assumptions may be impossible in many applications; so, practitioners are given little guidance for choosing the best estimator for a particular data set. As an alternative to verifying these assumptions we suggest practitioners construct a large family of HTE estimators and consider their similarities and differences.

Treatment effect estimation contrasts with prediction, where researchers can use cross-validation (CV) or a validation set to compare the performance of different estimators or to combine them in an ensemble. This is infeasible for treatment effect estimation because of the fundamental problem of causal inference: we can never observe the treatment effect for any individual unit directly, so we have no source of truth to validate or cross-validate against. Partial progress has been made in addressing this problem; for example, Athey and Imbens (2015) suggest using the transformed outcome as the truth, a quantity equal in expectation to the individual treatment effect and Künzel et al. (2017) suggests using matching to impute a quantity similar to the unobserved potential outcome. However, even if there were a reliable procedure for identifying the estimator with the best predictive performance, we maintain that using multiple estimates can still be superior, because the best performing method or ensemble of methods may perform well in some regions of the feature space and badly in others; using many estimates simultaneously may permit identification of this phenomenon. For example, researchers can construct a worst-case estimator that is equal to the most pessimistic point estimate, for each point in the feature space, or they can consider stability (Yu, 2013) to assess whether one can trust estimates for a particular subset of units.

2. Methods

2.1 Study setting

The data set we analyzed was constructed for the Empirical Investigation of Methods for Heterogeneity Workshop at the 2018 Atlantic Causal Inference Conference. The organizers of the workshop: Carlos Carvalho, Jennifer Hill, Jared Murray, and Avi Feller used the National Study of Learning Mindsets, a randomized controlled trial in a probability sample of U.S. public high schools, to simulate an observational study. The organizers did not disclose how the simulated observational data were derived from the experimental data because the workshop was intended to evaluate procedures for analyzing observational studies, where the mechanism of treatment assignment is not known *a priori*.

2.2 Measured variables

The outcome was a measure of student achievement; the treatment was the completion of online exercises designed to foster a learning mindset. Eleven covariates were available for

each student: four are specific to the student and describe the self-reported expectations for success in the future, race, gender and whether the student is the first in the family to go to college; the remaining seven variables describe the school the student is attending measuring urbanicity, poverty concentration, racial composition, the number of pupils, average student performance, and the extent to which students at the school had fixed mindsets; an anonymized school id recorded which students went to the same school.

2.3 Notation and estimands

For each student, indexed by i , we observed a continuous outcome, Y_i , a treatment indicator, Z_i , that is 1 if the student was in the treatment group and 0 if she was in the control group, and a feature vector X_i . We adopt the notation of the Neyman-Rubin causal model: for each student we assume there exist two potential outcomes: if a student is assigned to treatment we observe the outcome $Y_i = Y_i(1)$ and if the student is assigned to control we observe $Y_i = Y_i(0)$. Our task was to assess whether the treatment was effective and, if so whether the effect is heterogeneous. In particular, we are interested in discerning if there is a subset of units for which the treatment effect is particularly large or small.

To assess whether the treatment is effective, we considered the average treatment effect,

$$\text{ATE} := \mathbb{E}[Y_i(1) - Y_i(0)],$$

and to analyze the heterogeneity of the data, we considered average treatment effects for a selected subgroup S ,

$$\mathbb{E}[Y_i(1) - Y_i(0) | X_i \in S],$$

and the Conditional Average Treatment Effect (CATE) function,

$$\tau(x) := \mathbb{E}[Y_i(1) - Y_i(0) | X_i = x].$$

2.4 Estimating average effects

Wherever we computed the ATE or the ATE for some subset, we used four estimators. Three were based on the `CausalGAM` package of Glynn and Quinn (2017). This package uses generalized additive models to estimate the expected potential outcomes, $\hat{\mu}_0(x) := \hat{\mathbb{E}}[Y_i(0) | X_i = x]$, and $\hat{\mu}_1(x) := \hat{\mathbb{E}}[Y_i(1) | X_i = x]$, and the propensity score: $\hat{e}(x) := \hat{\mathbb{E}}[Z_i | X_i = x]$. With these estimates we computed the inverse probability weighting (IPW) estimator,

$$\hat{\text{ATE}}_{\text{IPW}} := \frac{1}{n} \sum_{i=1}^n \left(\frac{Y_i Z_i}{\hat{e}_i} - \frac{Y_i (1 - Z_i)}{1 - \hat{e}_i} \right),$$

the regression estimator,

$$\hat{\text{ATE}}_{\text{Reg}} := \frac{1}{n} \sum_{i=1}^n [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)],$$

and the augmented inverse probability weighted (AIPW) estimator,

$$\hat{\text{ATE}}_{\text{AIPW}} := \frac{1}{2n} \sum_{i=1}^n \left(\frac{[Y_i - \hat{\mu}_0(X_i)] Z_i}{\hat{e}_i} + \frac{[\hat{\mu}_1(X_i) - Y_i][1 - Z_i]}{1 - \hat{e}_i} \right).$$

We also used the `Matching` package of Sekhon (2011) to construct a matching estimator for the ATE. Matches were required to attend the same school as the student to which they were matched and to be assigned to the opposite treatment status. Amongst possible matches satisfying these criteria we selected the student minimizing the Mahalanobis distance on the four student specific features.

2.5 Characterizing heterogeneous treatment effects

In any data set there might be some units where estimators significantly disagree; when this happens, we should not trust any estimate unless we understand why certain estimates are unreasonable for these units.¹ Instead of simply reporting an estimate that is likely to be wrong, we should acknowledge that conclusions cannot be drawn and more data or domain specific knowledge is needed. Figure 1 demonstrates this phenomenon arising in practice. It shows the estimated treatment effect for ten subjects corresponding to 28 CATE estimators (these estimates arise from the data analyzed in the remainder of this paper). Some of these estimators may have better generalization error than others. However, a reasonable analyst could have selected any one of them. We can see that for five units the estimators all fall in a tight cluster, but for the remaining units, the estimators disagree markedly. This may be due to those units being in regions with little overlap, where the estimators overcome data scarcity by pooling information in different ways.

Since an estimate of the entire CATE function is hard to interpret and drawing statistically significant conclusions based on it is difficult, we decided to use our estimate of the CATE function to identify large subgroups with a markedly different average treatment effects and then form conclusions based on the differences in ATE for these subgroups.

To ensure the treatment effect estimates for the selected subgroups were valid, we divided the data into an **exploration set** and an equally sized **validation set**. We used the exploration set to identify subsets, which may behave differently. To do this, we trained all 28 CATE estimators on the exploration set and formulated hypotheses based on plots of the marginal CATE: for example, based on plots of the CATE estimates we might theorize that students in schools with more than 900 students have a much higher treatment effect than those in schools with less than 300 students. Next, we used the validation set to verify our findings by estimating the ATEs of each of the subgroups.

The exploration and validation sets were constructed by randomly associating schools (not individual students) with each set, so students who attended the same school were never split between the exploration and validation set. We adopted this procedure because it excludes the possibility of there being dependence between the exploration and validation sets if students who attend the same school influence each others' outcomes; it also mirrors the probability sampling approach used to construct the full sample and it means that we can argue that the estimand captured by evaluating our hypotheses on the validation set is the estimand corresponding to the population from which all schools were drawn.

1. A standard approach to capture uncertainty is to report the standard errors or confidence regions for a single estimator, and we recommend using this approach in addition to our approach. However, for CATE estimation, these methods can be misleading and cannot be trusted blindly. For example, in Appendix C of Künzel et al. (2017), the authors found that in regions without overlap, bootstrap confidence intervals were smaller than in regions with overlap. This suggests that estimates were more trustworthy in regions without overlap but, in fact, the opposite was true.

Figure 1: CATE estimation for ten units. For each unit, the CATE is estimated using 28 different estimators.

2.6 CATE estimators

We use several procedures to estimate the CATE and we give a brief overview of the procedures here; however, interested readers should consult the referenced papers for a complete exposition.

Many of the procedures can be classified as meta-learners: they modify base learners designed for standard non-causal prediction problems to estimate the CATE. This is advantageous because we can select a base learner that is designed to work well on the data we are analyzing.

1. The **T-Learner** is the most common meta-learner. Base learners are used to estimate the control and treatment response function separately, $\hat{\mu}_1(x) := \hat{\mathbb{E}}[Y_i(1)|X_i = x]$ and $\hat{\mu}_0(x) := \hat{\mathbb{E}}[Y_i(0)|X = x]$. The CATE estimate is then the difference between these two estimates, $\hat{\tau}^T(x) := \hat{\mu}_1(x) - \hat{\mu}_0(x)$.
2. The **S-Learner** uses one base learner to estimate the joint outcome function, $\hat{\mu}(x, z) := \hat{\mathbb{E}}[Y_i|X_i = x, Z_i = z]$. The predicted CATE is the difference between the predicted values when the treatment assignment indicator is changed from treatment to control, $\hat{\tau}^S(x) := \hat{\mu}(x, 1) - \hat{\mu}(x, 0)$.
3. The **MO-Learner** (Rubin and van der Laan, 2007; Walter et al., 2018) is a two stage meta-learner. It first uses the base learners to estimate the propensity score, $\hat{e}(x) := \hat{\mathbb{E}}[Z_i|X = x]$, and the control and treatment response functions. It then

defines the adjusted modified outcome as

$$R_i := \frac{Z_i - \hat{e}(x_i)}{\hat{e}(x_i)[1 - \hat{e}(x_i)]} \left(Y_i - \hat{\mu}_1(x_i)[1 - \hat{e}(x_i)] - \hat{\mu}_0(x_i)\hat{e}(x_i) \right).$$

An estimate of the CATE is obtained by using a base learner to estimate the conditional expectation of R_i given X_i , $\hat{\tau}^{MO}(x) := \hat{\mathbb{E}}[R_i|X_i = x]$.

4. The **X-Learner** (Künzel et al., 2017) also uses base learners to estimate the response functions and the propensity score. It then defines the imputed treatment effects for the treatment group and control group separately as $\tilde{D}_i^1 := Y_i(1) - \hat{\mu}_0(X_i)$ and $\tilde{D}_i^0 := \hat{\mu}_1(X_i) - Y_i(0)$. The two estimators for the CATE are obtained by using base learners to estimate the conditional expectation of the imputed treatment effects, $\hat{\tau}_1^X := \hat{\mathbb{E}}[\tilde{D}_i^1|X_i = x]$, and $\hat{\tau}_0^X := \hat{\mathbb{E}}[\tilde{D}_i^0|X_i = x]$. The final estimate is then a convex combination of these two estimators,

$$\hat{\tau}^X(x) := \hat{e}(x)\hat{\tau}_0^X(x) + (1 - \hat{e}(x))\hat{\tau}_1^X(x).$$

All of these meta-learners have different strength and weaknesses. For example, the T-Learner performs particularly well when the control and treatment response function are simpler than the CATE. The S-Learner performs particularly well when the expected treatment effect is mostly zero or constant. The X-Learner, on the other hand, has very desirable properties when either the treatment or control group is much larger than the other group.

Note, however, that all of these meta-learners need a base learners to be fully defined. We believe that tree-based estimators perform well on mostly discrete and low-dimensional data sets. Therefore, we use the `causalToolbox` package (Künzel et al., 2018) that implements all of these estimators combined with RF and BART.

Using two different tree estimators is desirable because CATE estimators based on BART perform very well when the data-generating process has some global structure (e.g., global sparsity or linearity), while random forest is better when the data has local structure that does not necessarily generalize to the entire space. However, to protect our analysis from biases caused by using tree-based approaches only, we also included methods based on neural networks. We followed Künzel et al. (2018) and implemented the S, T and X-Neural Network methods.

We also included tree-based learners that were not meta-learners that we believe would work well on this data set:

5. The **causal forest** algorithm (Wager and Athey, 2017b) is a generalization of the random forest algorithm to estimates the CATE directly. Similar to random forest, it is an ensemble of many tree estimators. Each of the tree estimators follows a greedy splitting strategy to generate leaves for which the CATE function is as homogeneous as possible. The final estimate for each tree for a unit with features x is the difference-in-means estimate of all units in the training set that fall in the same leaf as x .

6. The **R-Learner** (Nie and Wager, 2017) is a set of algorithms that use an approximation of the following optimization problem to estimate the CATE,

$$\arg \min_{\tau} \left\{ \frac{1}{n} \sum_{i=1}^n \left(\left(Y_i - \hat{\mu}^{(-i)}(X_i) \right) - \left(W_i - \hat{e}^{(-i)}(X_i) \right) \tau(X_i) \right)^2 + \Lambda_n(\tau(\cdot)) \right\}.$$

$\Lambda_n(\tau(\cdot))$ is a regularizer and $\mu^{(-i)}(x)$ and $\hat{e}^{(-i)}(X_i)$ are held-out predictions of $\mu(x) = \mathbb{E}[Y_i|X_i = x]$ and the propensity score, $e(x)$, respectively. There are several versions of the R-Learner; we have decided to use one that is based on XGBoost (Chen and Guestrin, 2016) and one that is based on RF.

Although we expected there would be school-level effects, and that both the expected performance of each student and the CATE would vary from school to school, it was not clear how to incorporate the school id. The two choices we considered were to include a categorical variable recording the school id, or to ignore it entirely. The former makes parameters associated with the six school-level features essentially uninterpretable because they cannot be identified separately from the school id; the second may lead to less efficient estimates because we are denying our estimation procedure the use of all data that was available to us. Because we do not want our inference to depend on this decision, we fit each of our estimators twice: once including school id as a feature, and once excluding it. We considered 14 different CATE estimation procedures; since each procedure was applied twice, a total of 28 estimators were computed.

3. Workshop Results

Our sample consisted of about 10,000 students enrolled at 76 different schools. The intervention was applied to 33% of the students. Pre-treatment features were similar in the treatment and control groups but some statistically significant differences were present. Specifically, a variable capturing self-reported expectations for success in the future had mean 5.22 (95% CI, 5.20-5.25) in the control group and mean 5.36 (5.33-5.40) in the treatment group. This meant students with higher expectations of achievement were more likely to be treated.

We assessed whether overlap held by fitting a propensity score model and we found that propensity score estimate for all students in the study was between 0.15 and 0.46 therefore, the overlap condition is likely to be satisfied.

3.1 Average treatment effects

The IPW, regression, and AIPW estimator yielded estimates identical up to two significant figures: 0.25 with 95% bootstrap confidence interval of (0.22, 0.27). The matching estimator gave a similar ATE estimate of 0.26 with confidence interval (0.23, 0.28).

The similarity of all the estimates we evaluated is reassuring, but we cannot exclude the possibility that the experiment is affected by an unobserved confounder that affects all estimators in a similar way. To address this, we characterize the extent of hidden bias required to explain away our conclusion. We conducted a sensitivity analysis for the matching estimator using the `sensetivitymv` package of Rosenbaum (2018). We found that

a permutation test for the matching estimator still finds a significant positive treatment effect provided the ratio of the odds of treatment assignment for the treated unit relative to the odds of treatment assignment for the control unit in each pair can be bounded by 0.40 and 2.52. This bound is not very large, and it is plausible that there exists an unobserved confounder that increases the treatment assignment probability for some unit by a factor of more than 2.52. More information about the treatment assignment mechanism would be required to conclude whether this extent of confounding exists.

3.2 Heterogeneous effects

The marginal distribution and partial dependence plots for the 28 CATE estimators as a function of school-level pre-existing mindset norms are shown on the left-hand side of Figure 2. There appears to be substantial heterogeneity present: students at schools with mindset norms lower than 0.15 may have a larger treatment effect than students at schools with higher mindset norms. However, the Figure suggests the conclusion is not consistent for all of the 28 estimators. A similar analysis of the feature recording the school achievement level is shown on the right-hand side of this Figure. Again, we appear to find the existence of heterogeneity: students with school achievement level near the middle of the range had the most positive response to treatment. On the basis of this figure, we identified thresholds of -0.8 and 1.1 for defining a low achievement level, a middle achievement level, and a high achievement level subgroup.

Figure 2: Marginal CATE and Partial Dependence Plot (PDP) of the CATE as a function of school-level pre-existing mindset norms and school achievement level.

We then used the validation set to construct ATE estimates for each of the subgroups. We found that students who attended schools where the measure of fixed mindsets was less than 0.15 had a higher treatment effect (0.31, 95% CI 0.26-0.35) than students where

the fixed mindset was higher (0.21, 0.17–0.26). Testing for equality of the ATE for these two groups yielded a p -value of 0.003. However, when we considered the subsets defined by school achievement level the differences were not so pronounced. Students at the lowest achieving schools had the smallest ATE estimate 0.19 (0.10–0.32); while students at middle and high-achieving schools had similar ATE estimates: 0.28 (0.24–0.32) and 0.24 (0.16–0.31) respectively. However, none of the pairwise difference between the three groups were significant.

4. Postworkshop results

During the workshop, other contributors found that the variable recording the urbanicity of the schools might explain some of the heterogeneity and so we will investigate that finding here. The left-hand side of Figure 3 shows the CATE as a function of urbanicity, and the right hand side of this figure shows the CATE as a function of the student's self-reported expectation of success. We formulated two hypotheses: students at schools with an urbanicity of 3 seemed to have a lower treatment effect than students at other schools; students with a self-reported evaluation of 4 might enjoy a higher treatment effect.

These hypotheses were obtained by only using the exploration set; to confirm or refute these hypotheses we used the validation set. The validation set confirmed the hypothesis that students at schools with an urbanicity of 3 had a lower treatment effect (0.16, 0.08–0.24) compared to students at schools with a different urbanicity (0.28, 0.25–0.31); however, we could not reject the null hypothesis of no difference for the subsets identified by the self-reported evaluation measure. The urbanicity test yielded a p -values of 0.008 and the self-reported evaluation test yielded a p -value of 0.56.

Figure 3: Marginal CATE and PDP of Urbanicity and self-reported expectations.

5. Discussion

5.1 The importance of considering multiple estimators

The results of our analysis confirm that point estimates of the CATE can differ markedly depending on subtle modelling choices; so, we confirm that an analyst’s discretion may be the deciding factor in whether and what kind of heterogeneity is found. As the methodological literature on heterogeneous treatment effect estimation continues to expand this problem will become more serious. To facilitate applying many estimation procedures we have authored an R package `causalToolbox` that provides a uniform interface for constructing many common heterogeneous treatment estimators. The design of the package makes it straightforward to add new estimators as they are proposed and gain currency.

Differences that arise in our estimation of the CATE function translate directly into suboptimal real-world applications of the treatment considered. To see this, we propose a thought experiment: suppose we wanted to determine the treatment for a particular student: a natural treatment rule is to allocate her to treatment if her estimated CATE exceeds a small positive threshold or withhold treatment if it is below the threshold. A CATE estimator may be chosen on the basis of personal preference or prior experience and it is likely that, for some experimental subjects, the choice of estimator will affect the CATE estimated to such an extent that it changes the treatment decision. This is particularly problematic in studies where analysts have a vested interest in a particular result and are working without a pre-analysis plan, as they should not have discretion to select a procedure that pushes the results in the direction they desire. On the other hand, if analysts consider a wide variety of estimators, as we recommend, and if most estimators agree for an individual, we can be confident that our decision for that individual is not a consequence of arbitrary modelling choices. Conversely, if some estimators predict a positive and some a negative response, we should reserve judgment for that unit until more conclusive data is available and admit that we do not know what the best treatment decision is.

5.2 Would we recommend the online exercises?

We find that the overall effect of the treatment is significant and positive. We were not able to identify a subgroup of units that had significant and negative treatment effect and we would therefore recommend the treatment for every student. In addition, our sensitivity analysis suggested that our findings would still hold provided the confounder is not too strong. Nonetheless, we cannot exclude the possibility that there is a strong confounder and we would need to better understand the assignment mechanism to eliminate this possibility. This issue deserves investigation because we saw that students who had higher expectations for success in the future were more likely to be in the treatment group. Uncovering the heterogeneity in the CATE function proved to be difficult. We found heterogeneity could be identified from school-level pre-existing mindset norms and urbanicity but in general we had limited power to detect heterogeneous effects. For example, experts believe that the heterogeneity might be moderated by pre-existing mindset norms and school-level achievement. For both covariates, we see that most CATE estimators produce estimates that are consistent with this theory. Domain experts also believe that there could be a “Goldilocks effect” where middle-achieving schools have the largest treatment effect. We are not able

to verify this statistically, but we do observe that most CATE estimators describe such an effect.

Acknowledgments

We thank Carlos Carvalho, Jennifer Hill, Jared Murray, and Avi Feller for organizing the Empirical Investigation of Methods for Heterogeneity Workshop and their valuable feedback. We thank Office of Naval Research (ONR) grant N00014-15-1-2367.

References

- Athey, S. and Imbens, G. W. (2015). Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050(5).
- Athey, S. and Imbens, G. W. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences of the United States of America*, 113(27):7353–60.
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA. ACM.
- Glynn, A. and Quinn, K. (2017). *CausalGAM: Estimation of Causal Effects with Generalized Additive Models*. R package version 0.1-4.
- Green, D. P. and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511.
- Henderson, N. C., Louis, T. A., Wang, C., and Varadhan, R. (2016). Bayesian analysis of heterogeneous treatment effects for patient-centered outcomes research. *Health Services and Outcomes Research Methodology*, 16(4):213–233.
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Künzel, S., Sekhon, J., Bickel, P., and Yu, B. (2017). Meta-learners for estimating heterogeneous treatment effects using machine learning. *arXiv preprint arXiv:1706.03461*.
- Künzel, S., Tang, A., Xie, L., Saarinen, T., Bickel, P., Yu, B., and Sekhon, J. (2018). *causalToolbox: Toolbox for Causal Inference with emphasis on Heterogeneous Treatment Effect Estimator*. R package version 0.0.1.000.
- Künzel, S. R., Stadie, B. C., Vemuri, N., Ramakrishnan, V., Sekhon, J. S., and Abbeel, P. (2018). Transfer learning for estimating causal effects using neural networks. *arXiv preprint arXiv:1808.07804*.
- Nie, X. and Wager, S. (2017). Learning objectives for treatment effect estimation. *arXiv preprint arXiv:1712.04912*.
- Powers, S., Qian, J., Jung, K., Schuler, A., Shah, N. H., Hastie, T., and Tibshirani, R. (2018). Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*.
- Rosenbaum, P. R. (2018). *sensitivitymv: Sensitivity Analysis in Observational Studies*. R package version 1.4.3.
- Rubin, D. and van der Laan, M. J. (2007). A doubly robust censoring unbiased transformation. *The international journal of biostatistics*, 3(1).

- Sekhon, J. S. (2011). Multivariate and propensity score matching software with automated balance optimization: The Matching package for R. *Journal of Statistical Software*, 42(7):1–52.
- Taddy, M., Gardner, M., Chen, L., and Draper, D. (2016). A nonparametric bayesian analysis of heterogenous treatment effects in digital experimentation. *Journal of Business & Economic Statistics*, 34(4):661–672.
- Tian, L., Alizadeh, A. A., Gentles, A. J., and Tibshirani, R. (2014). A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association*, 109(508):1517–1532.
- Wager, S. and Athey, S. (2017a). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*.
- Wager, S. and Athey, S. (2017b). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, (just-accepted).
- Walter, S., Sekhon, J., and Yu, B. (2018). Analyzing the modified outcome for heterogeneous treatment effect estimation. *Unpublished manuscript*.
- Yu, B. (2013). Stability. *Bernoulli*, 19:1484–1500.