PAC Style Guarantees for Doubly Robust Generalized Front-Door Estimator

Yuhao Wang[1] **Arnab Bhattacharyya**[1] **Jin Tian**[2] **N. V. Vinodchandran**[3]

[1]National University of Singapore
[2]Iowa State University
[3]University of Nebraska-Lincoln

Abstract

Doubly robust estimators present a promising methodology for estimating treatment effects in observational studies. This paper provides a finite sample analysis of the doubly robust estimators for both the back-door model (where treatment, outcome, and covariates are observed) and the generalized front-door model (which includes unmeasured confounding). Our approach establishes PAC-style guarantees of the deviation of the estimators in term of the divergence of probability distributions. These bounds demonstrate that minimizing the estimation error of the treatment effect in terms of Chi-square distance is crucial for minimizing the variance between true and estimated model.

1 INTRODUCTION

Estimating the causal effect of a treatment on an outcome is a fundamental task across empirical sciences [Winship and Morgan, 1999, Funk et al., 2011, Lechner et al., 2011, Lopez and Gutman, 2017, Fernández-Loría and Provost, 2022]. The problem of causal effect identification in the literature asks whether the causal effect of a treatment A on an outcome Y, $\Pr(y|do(a^*))$, can be computed from a combination of observational data and a causal diagram represented by a directed acyclic graph (DAG). Methods have been developed for solving the identification problem and its extensions including celebrated Pearl's do-calculus [Pearl, 1995] and complete identification algorithms [Tian and Pearl, 2003, Shpitser and Pearl, 2006, Huang and Valtorta, 2006, Lee et al., 2019, Correa et al., 2021]. These identification algorithms assume perfect knowledge of the observational distribution, and they express the target interventional distribution as a function of the observational distribution. In practice, however, one only has in hand estimates of the observational distribution, and the goal is to obtain the best possible approximation of the interventional distribution. Therefore, developing *robust* estimators for causal estimands is a problem of significant and sustained interest in the causal inference community [Jung et al., 2021, Xia et al., 2021, 2022, Bhattacharya et al., 2022, Jung et al., 2023].

One of the most popular methods for estimating causal effects is via covariate adjustment. If a set of covariates X satisfies the back-door (BD) criterion [Pearl, 1995] relative to (A, Y), then the causal effect of A on Y can be computed by the covariate adjustment formula

$$\Pr(y|do(a^*)) = \sum_x \Pr(y|a^*, x)\Pr(x) = \mathbb{E}_x[\Pr(y \mid a^*, x)],$$

also called the g-formula [Robins, 1986]. There exists an extensive literature on estimating the BD formula from finite samples including doubly robust estimators for addressing model misspecification [Bang and Robins, 2005, Robins et al., 2009, van der Laan and Gruber, 2012, Rotnitzky et al., 2017, Luedtke et al., 2017, Díaz et al., 2023] and double/debiased machine learning (DML) estimators [Robins et al., 1994, Van der Laan and Rose, 2011, Díaz and van der Laan, 2013, Benkeser et al., 2017, Kennedy et al., 2017, Chernozhukov et al., 2018, Rotnitzky and Smucler, 2019, Smucler et al., 2020, Colangelo and Lee, 2020].

Despite the popularity of covariate adjustment for causal effect estimation, its applications are limited to the settings where there are no unobserved confounders between X and Y. For example, in the causal diagram in Fig. 1(b), there exist no covariates to adjust for the confounding due to the unobserved confounder U. It turns out, another classical identification strategy, the front-door (FD) criterion [Pearl, 1995], is applicable in Fig. 1(b) to obtain the following (generalized) FD adjustment equation:

$$\Pr(y|do(a^*)) = \sum_{a,z,x} \Pr(z|a^*, x)\Pr(y|a, z, x)\Pr(a, x).$$

Glynn and Kashin [2017, 2018] discussed practical applications of FD adjustment for estimating causal effects. Fulcher et al. [2020] have developed a doubly robust (DR) estimator for estimating the generalized FD estimand from finite samples. Guo et al. [2023] extended the work of Fulcher et al. to provide targeted minimum loss based estimators (TMLEs) of the FD adjustment.

While Fulcher et al. [2020] provided asymptotic analysis and established asymptotic normality of their DR estimator, the finite sample behavior of their estimator is unknown. More broadly, while there have been previous works establishing finite sample guarantees for doubly robust estimators (e.g., Mou et al. [2022], Chernozhukov et al. [2023]), they have focused on Gaussian approximation and do not explicitly formulate the sample complexity in terms of standard notions of divergence between the estimate of the observational density and the true observational density.

In this paper, we provide a finite sample analysis of the doubly robust estimators for the BD and FD settings. We frame the finite sample complexity bounds in terms of divergence measures between the model distributions and the true distributions. In particular, our sample complexity bound is presented in terms of the natural measure of χ^2 distance and its generalizations. We first provide a PAC-style finite sample complexity bounds for the standard DR estimator for the BD adjustment. We then provide finite sample bound for Fulcher's DR estimator for the FD adjustment.

2 OUR RESULTS

Back-door adjustment We revisit the well-known double/debiased machine learning (DML) estimator for covariate adjustment in the BD setting [Robins and Rotnitzky, 1995, Chernozhukov et al., 2017] and analyze the mean-squared error in the finite sample setting. The novelty of our result is that we express the mean-squared error explicitly in terms of the errors in the estimates of the treatment and outcome distributions. These errors are formulated in terms of χ^2-divergence, roughly as follows:

Assumption 2.1. *Assume for all x, the following condition holds:*

$$\mathop{\mathbb{E}}_{Y(a^*,x)} Y^2 \le V, \quad \Pr[A(x)=a^*] \ge \mu, \textit{ and } \quad \Pr[\widehat{A}(x)=a^*] \ge \mu.$$

Theorem 2.2. *Under Assumption 2.1, for any $\varepsilon > 0$:*

$$\begin{aligned}&\Pr[|\psi_n - \psi^*| > \varepsilon]\\ &< \frac{1}{n\varepsilon^2}\mathcal{O}_{V,\mu}\left(1 + \mathop{\mathbb{E}}_{x}\chi^2\left(\widehat{Y}(a^*,x)\|Y(a^*,x)\right) + \mathop{\mathbb{E}}_{x}\chi^2\left(\widehat{A}(x)\|A(x)\right)\right)\\ &\quad + \frac{1}{\varepsilon^2}\mathcal{O}_{V,\mu}\left(\mathop{\mathbb{E}}_{x}\chi^2\left(\widehat{Y}(a^*,x)\|Y(a^*,x)\right)\cdot\chi^2\left(\widehat{A}(x)\|A(x)\right)\right)\end{aligned}$$

where $Y(a,x)$ is the conditional distribution of Y given $A=a, X=x$, and $A(x)$ is the conditional distribution of A given $X=x$, and their hatted versions are the model estimates. Hence, as the number of samples n for the doubly robust estimator goes to infinity, the remaining error is due to the mismatch measured in χ^2 between the model estimates and the truth, for the propensity and outcome distributions. Note that this is an expectation of the *product* of the two divergences, demonstrating the *mixed-bias* or *product rate* phenomenon of the doubly robust estimators [Chernozhukov et al., 2020]. Our bound can be used to get guidance on how to construct the estimators $\widehat{Y}$ and $\widehat{A}$; learning a distribution that minimizes a particular divergence is a question in *distribution learning*. For example, the problem of learning a distribution minimizing the χ^2 divergence was explicitly studied in [Kamath et al., 2015].

Front-door adjustment We extend our analysis to Fulcher's DR estimator for the generalized front door adjustment. Our formulation is inspired by a distribution learning framework. We frame the bounds in terms of divergences measures between the model distributions and the true distribution. In particular, our sample complexity bound is presented as the natural measure of χ^2 distance error in the estimates. The succinct form of our result can be gleaned from the statement below which we prove.

Assumption 2.3. *Assume $\forall a, z, x$,*

$$\begin{aligned}&\mathop{\mathbb{E}}_{Y(a,z,x)} Y^2 \le V,\ \Pr[Z(a,x)=z] \ge \mu_Z \cdot 1[\Pr[Z(a^*,x)=z] \ne 0],\\ &\textit{and} \quad \Pr[A(x)=a^*] \ge \mu_A\end{aligned}$$

Theorem 2.4. *Under Assumption 2.3*

$$\begin{aligned}&\Pr[|\psi_n - \psi^*| > \varepsilon]\\ &< \frac{1}{n\varepsilon^2}\mathcal{O}_{V,\mu_Z,\mu_A}\Bigg(1 + \mathop{\mathbb{E}}_{a,z,x}\chi^2\left(\widehat{Y}(a,z,x)\|Y(a,z,x)\right) + \mathop{\mathbb{E}}_{x}\chi^2\left(\widehat{A}(x)\|A(x)\right)\\ &+ \mathop{\mathbb{E}}_{x}\chi^2\left(\widehat{Z}(a^*,x)\|Z(a^*,x)\right) + \mathop{\mathbb{E}}_{a,x}\chi^2\left(\widehat{Z}(a,x)\|Z(a,x)\right)\Bigg)\\ &+ \frac{1}{\varepsilon^2}\mathcal{O}_{V,\mu_A,\mu_Z}\Bigg(\Bigg(\sqrt{\mathbb{E}_{a,z,x}\chi^4(\widehat{Y}(a,z,x)\|Y(a,z,x))} + \sqrt{\mathbb{E}_x\chi^4(\widehat{A}(x)\|A(x))}\\ &\quad + \sqrt{\mathbb{E}_x\chi^4(\widehat{Z}(a^*,x)\|Z(a^*,x))} + \sqrt{\mathbb{E}_{a,x}\chi^4(\widehat{Z}(a,x)\|Z(a,x))}\Bigg)^2\Bigg)\end{aligned}$$

Again, the mixed bias phenomenon can be seen in the part of the bound above that is independent of n.

3 CONCLUSION

In conclusion, we provide a finite sample analysis of doubly robust estimators for BD and FD settings, establishing PAC-style guarantees for the estimator's deviation based on divergence measures between model and true distributions. Our results highlight that minimizing the divergence of probability estimation error is essential for reducing the variance between true and estimated models.

References

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005. 1

David Benkeser, Marco Carone, MJ Van Der Laan, and PB Gilbert. Doubly robust nonparametric inference on the average treatment effect. *Biometrika*, 104(4):863–880, 2017. 1

Rohit Bhattacharya, Razieh Nabi, and Ilya Shpitser. Semiparametric inference for causal effects in graphical models with hidden variables. *Journal of Machine Learning Research*, 23:1–76, 2022. 1

Arnab Bhattacharyya, Sutanu Gayen, Saravanan Kandasamy, Ashwin Maran, and N. Variyam Vinodchandran. Learning and sampling of atomic interventions from observations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 842–853. PMLR, 2020. URL `http://proceedings.mlr.press/v119/bhattacharyya20a.html`. 5

Arnab Bhattacharyya, Sutanu Gayen, Saravanan Kandasamy, Vedant Raval, and N. V. Vinodchandran. Efficient interventional distribution learning in the PAC framework. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 7531–7549. PMLR, 2022. URL `https://proceedings.mlr.press/v151/bhattacharyya22a.html`. 5

V Chernozhukov, W Newey, and R Singh. A simple and general debiased machine learning theorem with finite sample guarantees. *Biometrika*, 110, 06 2022. doi:10.1093/biomet/asac033. 5

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, and Whitney Newey. Double/debiased/neyman machine learning of treatment effects. *American Economic Review*, 107(5):261–65, 2017. 2

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018. 1

Victor Chernozhukov, Whitney Newey, Rahul Singh, and Vasilis Syrgkanis. Adversarial estimation of riesz representers. *arXiv preprint arXiv:2101.00009*, 2020. 2

Victor Chernozhukov, Whitney K Newey, and Rahul Singh. A simple and general debiased machine learning theorem with finite-sample guarantees. *Biometrika*, 110(1):257–264, 2023. 2

Kyle Colangelo and Ying-Ying Lee. Double debiased machine learning nonparametric inference with continuous treatments. *arXiv preprint arXiv:2004.03036*, 2020. 1

Juan Correa, Sanghack Lee, and Elias Bareinboim. Nested counterfactual identification from arbitrary surrogate experiments. *Advances in Neural Information Processing Systems*, 34, 2021. 1

Iván Díaz and Mark J van der Laan. Targeted data adaptive estimation of the causal dose–response curve. *Journal of Causal Inference*, 1(2):171–192, 2013. 1

Iván Díaz, Nicholas Williams, Katherine L Hoffman, and Edward J Schenck. Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association*, 118(542):846–857, 2023. 1

Carlos Fernández-Loría and Foster Provost. Causal decision making and causal effect estimation are not the same... and why it matters. *INFORMS Journal on Data Science*, 1(1):4–16, 2022. 1

Isabel R Fulcher, Ilya Shpitser, Stella Marealle, and Eric J Tchetgen Tchetgen. Robust inference on population indirect causal effects: the generalized front door criterion. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1):199–214, 2020. 1, 2, 5

Michele Jonsson Funk, Daniel Westreich, Chris Wiesen, Til Stürmer, M Alan Brookhart, and Marie Davidian. Doubly robust estimation of causal effects. *American journal of epidemiology*, 173(7):761–767, 2011. 1

Adam N Glynn and Konstantin Kashin. Front-door difference-in-differences estimators. *American Journal of Political Science*, 61(4):989–1002, 2017. 1

Adam N Glynn and Konstantin Kashin. Front-door versus back-door adjustment with unmeasured confounding: Bias formulas for front-door and hybrid adjustments with application to a job training program. *Journal of the American Statistical Association*, 113(523):1040–1049, 2018. 1

Anna Guo, David Benkeser, and Razieh Nabi. Targeted machine learning for average causal effect estimation using the front-door functional. *arXiv preprint arXiv:2312.10234*, 2023. 2

Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, pages 217–224. AUAI Press, 2006. 1, 5

Yonghan Jung, Jin Tian, and Elias Bareinboim. Estimating identifiable causal effects on markov equivalence class through double machine learning. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 1

Yonghan Jung, Iván Díaz, Jin Tian, and Elias Bareinboim. Estimating causal effects identifiable from combination of observations and experiments. In *Proceedings of the 37th Neural Information Processing Systems*, 2023. URL `https://causalai.net/r97.pdf`. 1

Sudeep Kamath, Alon Orlitsky, Dheeraj Pichapati, and Ananda Theertha Suresh. On learning distributions from their samples. In *Conference on Learning Theory*, pages 1066–1100. PMLR, 2015. 2

Edward H Kennedy, Zongming Ma, Matthew D McHugh, and Dylan S Small. Nonparametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 79(4):1229, 2017. 1

Michael Lechner et al. The estimation of causal effects by difference-in-difference methods. *Foundations and Trends® in Econometrics*, 4(3):165–224, 2011. 1

Sanghack Lee, Juan D Correa, and Elias Bareinboim. General identifiability with arbitrary surrogate experiments. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2019. 1

Michael J Lopez and Roee Gutman. Estimation of causal effects with multiple treatments: a review and new ideas. *Statistical Science*, pages 432–454, 2017. 1

Alexander R Luedtke, Oleg Sofrygin, Mark J van der Laan, and Marco Carone. Sequential double robustness in right-censored longitudinal models. *arXiv preprint arXiv:1705.02459*, 2017. 1

Wenlong Mou, Martin J Wainwright, and Peter L Bartlett. Off-policy estimation of linear functionals: Non-asymptotic theory for semi-parametric efficiency. *arXiv preprint arXiv:2209.13075*, 2022. 2

J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–710, 1995. 1, 5

James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986. 1

James Robins, Lingling Li, Eric Tchetgen, and Aad W van der Vaart. Quadratic semiparametric von mises calculus. *Metrika*, 69:227–247, 2009. 1

James M Robins and Andrea Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90 (429):122–129, 1995. 2

James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994. 1

Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average treatment effect estimation in non-parametric causal graphical models. *arXiv preprint arXiv:1912.00306*, 2019. 1

Andrea Rotnitzky, James Robins, and Lucia Babino. On the multiply robust estimation of the mean of the g-functional. *arXiv preprint arXiv:1705.08582*, 2017. 1

Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-markovian causal models. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, page 1219. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006. 1

Ilya Shpitser and Judea Pearl. Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9:1941–1979, 2008. 5

Ezequiel Smucler, Facundo Sapienza, and Andrea Rotnitzky. Efficient adjustment sets in causal graphical models with hidden variables. *arXiv preprint arXiv:2004.10521*, 2020. 1

Jin Tian and Judea Pearl. On the identification of causal effects. Technical Report R-290-L, 2003. 1, 5

Mark J van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*, 8(1), 2012. 1

Mark J Van der Laan and Sherri Rose. *Targeted learning*. Springer, 2011. 1

Christopher Winship and Stephen L Morgan. The estimation of causal effects from observational data. *Annual review of sociology*, 25(1):659–706, 1999. 1

Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. *Advances in Neural Information Processing Systems*, 34, 2021. 1

Kevin Muyuan Xia, Yushu Pan, and Elias Bareinboim. Neural causal models for counterfactual identification and estimation. In *The Eleventh International Conference on Learning Representations*, 2022. 1

PAC Style Guarantees for Doubly Robust Generalized Front-Door Estimator

Yuhao Wang[1] **Arnab Bhattacharyya**[1] **Jin Tian**[2] **N. V. Vinodchandran**[3]

[1]National University of Singapore
[2]Iowa State University
[3]University of Nebraska-Lincoln

A RELATED WORK

Causal identification and estimation constitute a vast and crucial area of research within many fields including statistics, econometrics, and computer science. Despite its importance, the machine learning community has only recently begun to rigorously apply PAC (Probably Approximately Correct) style finite sample guarantees to causal inference. Double ML has emerged as a promising approach for the robust estimation of causal effects in observational studies. The integration of machine learning algorithms allows for handling high-dimensional data while maintaining robustness against model misspecification. Finite sample analysis of Doubbly Robust Back Door estimator has recently been developed by Chernozokov et al. in Chernozhukov et al. [2022]. Fulcher et al. extended these ideas by providing a doubly robust estimator for the generalized Front Door estimand from finite samples and provided asymptotic analysis Fulcher et al. [2020]. There has been some progress on establishing PAC style guarantee on Pearl's graphical causal identification model Pearl [1995]. The works by Bhattacharyya et al. [2020] and Bhattacharyya et al. [2022] are particularly notable, offering PAC guarantees within the framework of causal identification algorithms initially characterized by Tian and Pearl [2003], Shpitser and Pearl [2008], and Huang and Valtorta [2006]. The present work contributes to this important and evolving area by focusing on PAC-style guarantees on sample complexity for estimating treatment effects in observational studies. We present PAC style analysis of double robust Back Door and Generalized Front Door estimators by examining their performance in terms of distribution divergence, providing new insights into finite sample behavior in causal inference.

B MODEL DESCRIPTION

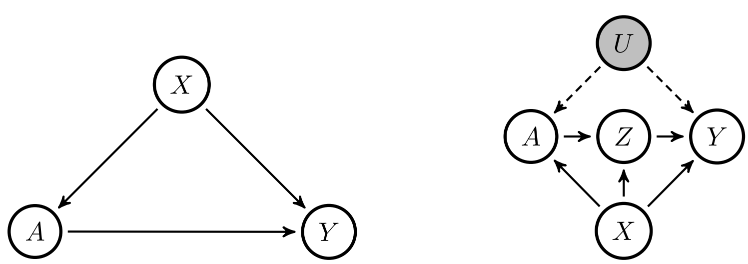

(a) G_1: Back-door (b) G_2: Front door

Figure 1: Causal inference with backdoor adjustment (a) and generalized front door adjustments (b)

B.1 BACK-DOOR ADJUSTMENT

For three variables A, X, Y in G_1 Fig. 1(a), represent the directed edges between these three variables are $X \to A$, $X \to Y$, $A \to Y$, where A denote the treatment, X denote the covariate, and Y represent the outcome.

We are given the following models: For any x, a random variable $\widehat{A}(x)$ which is supposed to be the model for the distribution of A conditioned on $X = x$. Besides, $\widehat{\mathcal{A}}(a; x)$ denote the conditional probability distribution $\Pr[\widehat{A}(x) = a]$. Similarly, for any x, and a, a random variable $\widehat{Y}(a, x)$ which is supposed to be the model for the distribution of Y conditioned on $A = a$, and $X = x$, and we apply $\widehat{\mathcal{Y}}(y; a, x)$ to represent the conditional probability distribution $\Pr[\widehat{Y}(a, x) = y]$.

We want to estimate the causal effect of A on Y, denote as:

$$\psi^* = \mathbb{E}[Y \mid \mathrm{do}(A = a^*)].$$

and ψ_n as the fimite sample estimator of ψ^*,

$$\psi_n = \frac{1}{n} \sum_{i=1}^{n} \phi(a_i, x_i, y_i).$$

where $\phi(a, x, y)$ can be estimated as follows:

$$\phi(a, x, y) = \mathop{\mathbb{E}}_{\widehat{Y}(a^*, x)} \widehat{Y} + \frac{1[a = a^*]}{\widehat{\mathcal{A}}(a; x)} \left(y - \mathop{\mathbb{E}}_{\widehat{Y}(a, x)} \widehat{Y} \right),$$

Our objective is, given finite samples, we want to deviate PAC-style guarantees with a high probability bound of the estimator on $\psi_n - \psi^*$ in term of divergence of probability distributions.

B.2 FRONT-DOOR ADJUSTMENT

As shown in Fig. 1(b) G_2, we have 4 observable: X, A, Z, Y. A represent the treatment, Z represent the mediator variable, X be a set of observed pre-exposure covariates, and Y represent the outcome. The directed edges are $X \to A$, $X \to Z$, $X \to Y$, $A \to Z$, and $Z \to Y$. There is also a bi-directed arrow between A and Y represent the effect of unobserved confounder. Note there is no direct arrow from A to Y.

We are given the following models:

- For any x, a random variable $\widehat{A}(x)$ which is supposed to be the model for the distribution of A conditioned on $X = x$.
- For any x, a, a random variable $\widehat{Z}(a, x)$ which is supposed to be the model for the distribution of Z conditioned on $A = a$ and $X = x$.
- For any x, a, and z, a random variable $\widehat{Y}(a, z, x)$ which is supposed to be the model for the distribution of Y conditioned on $A = a$, $Z = z$, and $X = x$.
- Let $\widehat{\mathcal{A}}(a; x)$ denote $\Pr[\widehat{A}(x) = a]$. Similar for $\widehat{\mathcal{Z}}(z; a, x)$ and $\widehat{\mathcal{Y}}(y; a, z, x)$.

Define the following quantity:

$$\begin{aligned}\phi(a, z, y, x) = \mathop{\mathbb{E}}_{\widehat{Z}(a^*, x)} \left[\mathop{\mathbb{E}}_{\widehat{Y}(a, \widehat{Z}, x)} \widehat{Y} \right] &+ \frac{1[a = a^*]}{\widehat{\mathcal{A}}(a; x)} \mathop{\mathbb{E}}_{\widehat{A}(x)} \left(\mathop{\mathbb{E}}_{\widehat{Y}(\widehat{A}, z, x)} \widehat{Y} - \mathop{\mathbb{E}}_{\widehat{Z}(a, x)} \left[\mathop{\mathbb{E}}_{\widehat{Y}(\widehat{A}, \widehat{Z}, x)} \widehat{Y} \right] \right) \\ &+ \left(y - \mathop{\mathbb{E}}_{\widehat{Y}(a, z, x)} \widehat{Y} \right) \cdot \frac{\widehat{\mathcal{Z}}(z; a^*, x)}{\widehat{\mathcal{Z}}(z; a, x)}\end{aligned}$$

Let

$$\psi_n = \frac{1}{n} \sum_{i=1}^{n} \phi(a_i, z_i, y_i, x_i)$$

for independent observations $(a_1, z_1, y_1, x_1), \ldots, (a_n, z_n, y_n, x_n)$. From above, $\mathbb{E}[\psi_n] = \psi^*$ whenever two of $\widehat{A}, \widehat{Z}, \widehat{Y}$ are correct.

Our goal is to bound the RMS error of the estimator on $\psi_n - \psi^*$ in term of divergence of probability distributions.