Causal Inference under Data Restrictions

by

Xiaoqing Tan

B.S., Sun Yat-sen University, 2018

Submitted to the Graduate Faculty of the

School of Public Health in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

UNIVERSITY OF PITTSBURGH SCHOOL OF PUBLIC HEALTH

This dissertation was presented

by

Xiaoqing Tan

It was defended on

July 25, 2022

and approved by

Lu Tang, PhD, Assistant Professor

Department of Biostatistics, University of Pittsburgh

Gong Tang, PhD, Professor

Department of Biostatistics, University of Pittsburgh

Chung-Chou H. Chang, PhD, Professor

Department of Medicine and Biostatistics, University of Pittsburgh

Emily Brant, MD, Assistant Professor

Department of Medicine, University of Pittsburgh

Zhengling Qi, PhD, Assistant Professor

School of Business, George Washington University

Copyright \bigodot by Xiaoqing Tan2022

Causal Inference under Data Restrictions

Xiaoqing Tan, PhD

University of Pittsburgh, 2022

This dissertation focuses on modern causal inference under uncertainty and data restrictions, with applications to neoadjuvant clinical trials, distributed data networks, and robust individualized decision making.

In the first project, we propose a method under the principal stratification framework to identify and estimate the average treatment effects on a binary outcome, conditional on the counterfactual status of a post-treatment intermediate response. Under mild assumptions, the treatment effect of interest can be identified. We extend the approach to address censored outcome data. The proposed method is applied to a neoadjuvant clinical trial and its performance is evaluated via simulation studies.

In the second project, we propose a tree-based model averaging approach to improve the estimation accuracy of conditional average treatment effects at a target site by leveraging models derived from other potentially heterogeneous sites, without them sharing subject-level data. To our best knowledge, there is no established model averaging approach for distributed data with a focus on improving the estimation of treatment effects. The performance of this approach is demonstrated by a study of the causal effects of oxygen therapy on hospital survival rate and backed up by comprehensive simulations.

In the third project, we propose a robust individualized decision learning framework with sensitive variables to improve the worst-case outcomes of individuals caused by sensitive variables that are unavailable at the time of decision. Unlike most existing work that uses mean-optimal objectives, we propose a robust learning framework via finding a newly defined quantile- or infimum-optimal decision rule. From a causal perspective, we also generalize the classic notion of (average) fairness to conditional fairness for individual subjects. The reliable performance of the proposed method is demonstrated through synthetic experiments and three real-data applications.

Public health significance: The dissertation addresses several aspects of causal infer-

ence: 1) identify principal stratum treatment effects; 2) enhance the estimation of treatment effects via heterogeneous data integration; 3) derive robust individualized decision rules considering worst-case scenarios. It has the potential to fundamentally improve the current practice in drug development and precision medicine.

Table of Contents

Preface				XX
1.0	Intro	oductio	on	1
	1.1	Challenges in Causal Inference under Data Restrictions		
		1.1.1	Data Collection Concerns in Traditional Clinical Trials	1
		1.1.2	Privacy Concerns in Distributed Data Networks	1
		1.1.3	Timeliness and Fairness Concerns in Decision Making	2
	1.2	Outlin	e and Contributions	2
2.0	Iden	tifying	g Principal Stratum Causal Effects Conditional on a Post-	
	treat	tment	Intermediate Response	5
	2.1	Introd	uction	5
	2.2	Relate	d Work	7
	2.3	The Principal Stratification Framework of Interest		
		2.3.1	Standard Setting for Neoadjuvant Studies	8
		2.3.2	Modeling a Counterfactual Outcome	9
	2.4	The P	roposed Method	10
		2.4.1	Key Identification Assumptions	10
		2.4.2	Identification of Model Parameters and Causal Estimands $\ . \ . \ .$	10
		2.4.3	Estimation of Causal Estimands	11
		2.4.4	Estimation of Model Parameters	13
		2.4.5	Consistency of Model Parameters and Causal Estimands $\ . \ . \ .$.	15
		2.4.6	Extension to Censored Data	15
	2.5	Simula	tion Studies	16
	2.6	Application to NSABP B-40 Trial		19
		2.6.1	B-40 Data Analysis	19
		2.6.2	Sensitivity of Initial Parameters in Optimization	20
		2.6.3	Comparisons to Sensitivity Analysis Method	21

	2.7	Discus	ssion and Future Work	22
3.0	AT	ree-ba	sed Model Averaging Approach for Personalized Treatment	
	Effe	ct Esti	imation from Heterogeneous Data Sources	24
	3.1	Introd	uction	24
	3.2	Relate	ed Work	26
		3.2.1	Heterogeneous Models	27
		3.2.2	Transportability	28
	3.3	A Tree	e-based Model Averaging Framework	28
		3.3.1	Causal Assumptions	29
		3.3.2	Model Ensemble	30
		3.3.3	Construction of Weights	32
		3.3.4	Interpretability of Weights	34
		3.3.5	Local Models: Obtaining $\hat{\tau}_k$	35
		3.3.6	Asymptotic Properties	35
	3.4	Simula	ation Studies	36
		3.4.1	Compared Estimators and Evaluation	37
		3.4.2	Estimation Performance	39
		3.4.3	Visualization of Information Borrowing	39
		3.4.4	Additional Simulations	41
	3.5	Exam	ple: A Multi-Hospital Data Network	42
	3.6	Discus	ssion	44
4.0	Rob	ust In	dividualized Decision Learning with Sensitive Variables	46
	4.1	Introd	uction	46
	4.2	Relate	ed Work	49
	4.3	Robust Decision Learning Framework with Sensitive Variables		51
		4.3.1	Preliminaries	51
		4.3.2	Robust Optimality with Sensitive Variables	53
		4.3.3	Identifying Vulnerable Subjects	55
		4.3.4	Estimation and Algorithm	55
		4.3.5	Extension to Multiple Discrete Sensitive Variables	57

	4.4	Numer	rical Studies	58
		4.4.1	Simulation Studies	59
		4.4.2	Real-data Applications	62
			4.4.2.1 Fairness in a Job Training Program	62
			4.4.2.2 Improvement of HIV Treatment	62
			4.4.2.3 Safe Resuscitation for Patients with Sepsis	63
			4.4.2.4 Results of Real-data Applications	64
	4.5	Discus	sion and Future Work	65
5.0	Sum	mary	and Future Work	67
App	pendi	x A. fo	or Chapter 2	69
	A.1	Estima	ation of $\Pr\{S_i(0) = 1 X_i = x\}$ and $\Pr\{S_i(1) = 1 X_i = x\}$	69
	A.2	Proofs	of Consistency of Model Parameters and Causal Estimands	71
	A.3	Calcul	ation of True Principal Stratum Causal Effects	74
App	pendi	x B. fo	or Chapter 3	75
	B.1	Relate	d Topics and Distinctions	75
	B.2	Proof	of Theorem 3.3.5	76
	B.3	Additi	onal Simulation Results	77
		B.3.1	Connection to Supervised Learning	77
		B.3.2	Various Sample Sizes in Local Sites	78
		B.3.3	Simulations under Observational Studies	78
		B.3.4	Covariate Dimensions	79
		B.3.5	Unequal Sample Size at Each Site	79
		B.3.6	Different Local Estimators	79
		B.3.7	Further Comparisons to Non-adaptive Ensemble	80
	B.4	Additi	onal Results for Data Application	94
	B.5	Real D	Data Access	96
App	pendi	x C. fo	or Chapter 4	97
	C.1	Additi	onal Literature Review	97
	C.2	Proofs	of Propositions	97
	C.3	Details	s on Modeling and Hyperparameter Tuning	99

C.4 Additional Simulations) 9
C.4.1 Different Quantile Criteria) 9
C.4.2 S as a Noise Variable \ldots 10)0
C.4.3 Violations of Causal Assumptions)0
C.5 Additional Information and Results for Real-data Applications 10 $$)5
C.5.1 Data Availability 10)5
C.5.2 Additional Background on the Sepsis Application $\ldots \ldots \ldots$)5
C.5.3 Visualizations $\ldots \ldots \ldots$)6
C.6 Causal and Decision Diagrams)7
Bibliography	12

List of Tables

1	Simulation results of the proposed method under three different parameter set-	
	tings and various sample sizes.	19
2	Sensitivity analysis for the estimated causal effect of bevacizumab in 3-year sur-	
	vival among those who would obtain pCR under chemotherapy plus bevacizumab.	22
3	MSE ratios of EF over LOC. As n increases, model averaging becomes more	
	powerful due to better estimation of τ_k , and is more pronounced when c is small.	38
4	Toy example setup of $E(Y X, S, A)$	48
5	Toy example results.	48
6	Simulation results for Example 1 and Example 2. Standard error in parenthesis.	61
7	Estimated objective and value of different IDRs for the three data applications.	
	Standard error in parenthesis. The outcome of each study is italicized	65
8	Simulation results for ratio between MSE of the estimator and MSE of LOC	
	(\mathbf{CT}) with a sample size of 500 at each site. A smaller number indicates larger	
	improvement over the local model. Estimators ending with "-oracle" makes use of	
	ground truth treatment effects. Our proposed methods ET and EF shows robust	
	performance in all settings whether or not using the information of ground truth	
	$ au_1(m{x})$	82
9	Hospital-level information of our analysis cohort in eICU database. Hospitals	
	are relabeled according to their average contribution to the estimation task at	
	hospital 1, the target site	95
10	Simulation results for Example 2 with continuous S using 0.1 quantile criterion	
	and 0.5 quantile criterion, respectively. Standard error in parenthesis. The pro-	
	posed RISE has more strengths when τ is small, as the algorithm aims to improve	
	the worst-outcome scenarios.	101

11	Simulation results for scenario when S is a noise variable. Vulnerable subjects	
	cannot be defined as S is not important in the example. The estimated objec-	
	tive and value function are similar for the compared IDRs, which indicates the	
	robustness of the proposed RISE.	102
12	Simulation results for Example 2 where the positivity assumption in Assump-	
	tion 4.3.2 is nearly violated. Standard error in parenthesis.	103
13	Simulation results for Example 2 where the unconfoundedness assumption in	
	Assumption 4.3.3 is violated. Standard error in parenthesis	104

List of Figures

1	Histogram of the causal estimates obtained from different initial values of model	
	parameters in the optimization process for 3-year EFS (Figure A) and 3-year	
	OS (Figure B), respectively. Except for some extreme initialization, most of the	
	causal estimates are the same or very close to the causal estimate calculated by	
	using zeros as initial parameters.	21
2	(a) Schema of the proposed algorithm. (b) Illustration of the augmented data	
	constructed from the estimation set of site 1	32
3	Box plots of MSE ratios of CATE estimators, respectively, over LOC, for (a)	
	discrete grouping and (b) continuous grouping across site. Different colors	
	imply different estimators, and x-axis, i.e., the value of c , differentiates the scale	
	of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA	
	performance is truncated due to large MSE ratios. The proposed ET and EF	
	achieve smaller MSE ratios compared to standard model averaging or ensemble	
	methods and are robust to heterogeneity across settings.	38
4	Visualization of simulation results under discrete grouping (a,b,c) and con-	
	tinuous grouping (d,e,f) when $c = 1$. (a) and (d) visualize the proposed ETs	
	where the site indicator and X_1 are selected as splitting variables, which is consis-	
	tent with the underlying data generation process. (b) and (e) show the predicted	
	treatment effects of the proposed EFs varying X_1 in each site, marginalized over	
	all other features. (b) is arranged according to the true grouping, odd sites ver-	
	sus even sites. The plot recovers the pattern of local and global heterogeneity.	
	(c) and (f) plot the interpretable model averaging weights in EFs over X_1 . The	
	weights of site 1 have a relatively large contribution to the weighted estimator	
	while models from other sites have different contributions for different X_1 de-	
	pending on their similarity in $\tau(\boldsymbol{x},k)$ to that in site 1. Corresponding ET and	
	EF show consistent patterns and recover the true grouping	40

- 5Application to estimating treatment effects of oxygen therapy on survival. (a) Expected survival of treatment decision following different estimators. The proposed EF shows the largest gain in improving survival rate, more promising than LOC and baseline. (b) Estimated treatment effects varying duration and BMI, two important features in the fitted EF. Patients with a BMI around 35, and a duration above 400 benefited the most. (c) Visualization of data-adaptive weights in the estimated EF varying BMI. Hospitals with a larger bed capacity tend to contribute more, the data of which might be more similar to hospital 1. 446 Box plots of the MSE ratios of CATE estimators, respectively, over LOC (\mathbf{CT}) and a sample size of **500** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings. 81 7 Box plots of the MSE ratios of CATE estimators, respectively, over LOC (\mathbf{CT}) and a sample size of **100** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings. 83
 - xiii

- 8 Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **1000** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.
- 9 Plots of the bias and MSE of EF-oracle varying sample site at each site for (a) discrete grouping and (b) continuous grouping across site, varying scale of global heterogeneity. Both bias and MSE reduces to zero as the sample size increases.

84

85

86

10 Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site under **observational design with a correctly specified propensity score model** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of *c*, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.

- Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site under **observational design with a misspecified propensity score model** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of *c*, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.
- 12 Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site, and covariate dimension of **20** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve closeto-zero MSE ratios with very small spreads in some settings.

87

- Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site, and covariate dimension of **50** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve closeto-zero MSE ratios with very small spreads in some settings.
- Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at site 1, and a sample size of **200** at other sites for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve closeto-zero MSE ratios with very small spreads in some settings.

89

Box plots of the MSE ratios of CATE estimators, respectively, over LOC (CF) and a sample size of 100 at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.

91

92

17	Box plots of the MSE ratios of CATE estimators, respectively, over LOC (\mathbf{CF})	
	and a sample size of 1000 at each site for (a) discrete grouping and (b) con-	
	tinuous grouping across site, respectively, varying scale of global heterogeneity.	
	Estimators ending with "-oracle" makes use of ground truth treatment effects.	
	Different colors imply different estimators, and x-axis, i.e., the value of c , differ-	
	entiates the scale of global heterogeneity. The red dotted line denotes an MSE	
	ratio of 1. MA performance is truncated due to large MSE ratios. The proposed	
	ET and EF achieve competitive performance compared to standard model aver-	
	aging or ensemble methods and are robust to heterogeneity across settings. Note	
	that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small	
	spreads in some settings.	93
18	Application to estimating treatment effects of oxygen therapy on survival with a	
	sample size weighting strategy. (a) Expected survival of treatment decision	
	following different estimators. (b) Estimated treatment effects varying duration	
	and BMI, two important features in the fitted EF. (c) Visualization of data-	
	adaptive weights in EF varying BMI.	96
19	Visualization for the job training program: SHAP variable importance plots for	
	decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked	
	by variable importance in descending order. Correlations between the feature	
	and their SHAP value are highlighted in color. The red color means a feature	
	is positively correlated with assigning treatment $A = 1$ and the blue indicates a	
	negative correlation.	108
20	Visualization for the ACTG175 dataset: SHAP variable importance plots for	
	decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked	
	by variable importance in descending order. Correlations between the feature	
	and their SHAP value are highlighted in color. The red color means a feature	
	is positively correlated with assigning treatment $A = 1$ and the blue indicates a	
	negative correlation.	109

- 22 (a) Causal diagram. Both X and S confound the effect of treatment A on outcome Y. The arrows represent the causal relationship between variables. Hence, both X and S should be used for decision making in general. (b) Decision diagram. Sensitive variables S is shown in a dotted circle as S may not be readily available at the time of decision making. We connect S and A with a dotted arrow to indicate that S is incorporated during model training of the decision rule, but not model deployment. Hence, the decision rule d only maps X to the treatment option, i.e., A = d(X). 111

Preface

This work, as with most research, is an effort in collaboration. First and foremost, I would like to express my deepest gratitude to my Ph.D. advisors Dr. Gong Tang and Dr. Lu Tang. This dissertation would not have been possible without their continuous guidance, support, and encouragement. Gong is an exemplary scholar, who is always hardworking and dedicated to research. I am constantly amazed by his sharpness on research and his deep insights about many different topics, ranging from medical science, applied math, to statistics, and computer science. His ability to persevere through difficulties and ask the right questions are abilities that I will continue to aspire to throughout my career. I am also fortunate enough to work with Lu, an inspirational advisor, who is incredibly generous with his ideas and time; and a caring mentor, who constantly offers his support and empathy. Our meetings and discussions have been a great source of inspiration and greatly shaped my way of approaching research problems. I am deeply indebted to him for devoting so much time and energy to mentoring me and guiding me through the transition from a student to a researcher. His passion for research have profoundly influenced me from both professional and personal perspectives.

I am also very grateful to have Dr. Joyce Chang, Dr. Zhengling Qi, and Dr. Emily Brant serving on my doctoral dissertation committee and providing me with invaluable comments. I am fortunate to collaborate with Joyce on the project of heterogeneous causal models integration. I am grateful for her training and guidance on causal inference, and I have learnt a lot about how to present research ideas, story telling, as well as critical thinking. I am grateful to Zhengling, who has helped me tremendously in developing the project of robust individualized decision learning. His constant support and humor make all the research meetings interesting and enjoyable. Zhengling's enthusiasm for research has left a lot of positive impacts on me. Emily has provided me with a wonderful opportunity to work on the exciting application of sepsis disease for interdisciplinary collaboration. Her domain expertise in medicine has been invaluable in leading me to find immense motivation for methodology development that has shaped my dissertation. Additionally, I would like to express my gratitude and appreciation to a number of wonderful researchers who I have been fortunate to collaborate with through my Ph.D. Specifically, I would like to thank Dr. Timothy Girard, Dr. David Samuels from Vanderbilt University, and others in their team. They have exposed me to the world of interesting applications of critical care medicine and provided me with numerous opportunities to apply meaningful and interpretable statistical tools in medical problems. I would like to thank Dr. Jiebiao Wang and Dr. Qi Yan from Columbia University for their guidance and support on our genetic side project. They have provided me with a great opportunity to discuss and learn causal inference problems under genetic context. I would like to thank the amazing research team I met when I interned at Eli Lilly and Company, including Dr. Shu Yang from North Carolina State University, Dr. Ilya Lipkovich, Dr. Douglas Faries, Dr. Wendy Ye, and Zbigniew Kadziola. My journey would not have been so rewarding without them, and I have learned a lot from our interactions.

I would also like to thank all my friends that I made and all the people that I met throughout the Ph.D. studies. My thanks also go to staff members at the Department of Biostatistics, University of Pittsburgh, who always patiently helped with my questions and warmly welcomed me into the office with big smiles.

Last but not least, I would like to thank my family, Yu Wang the duck, Kitty the cat, Larry the chinchilla, and Bibi the bunny for their unwavering support and unconditioned love.

1.0 Introduction

1.1 Challenges in Causal Inference under Data Restrictions

Modern statistical and machine learning methods are capable of capturing correlations between variables but often fail to inform us the causes behind. Causal inference, on the other hand, helps to understand the underlying data-generating process, which is critical when analyzing data from our contemporary world, particularly in fields such as healthcare, political science, and economics.

1.1.1 Data Collection Concerns in Traditional Clinical Trials

In traditional clinical trials, data collection often requires long years of follow-up, leading to problems such as patients' withdrawal and high cost of the study. Recently, in neoadjuvant clinical trials, early efficacy of a treatment is assessed first via an intermediate post-treatment response and the eventual efficacy is assessed via long-term outcomes such as survival. Although strongly associated with survival, this intermediate response has not been confirmed as a surrogate endpoint. To fully understand its clinical implication, it is important to establish causal estimands such as the causal effect in survival for patients who would obtain a certain intermediate response under treatment. In Chapter 2, driven by a recent neoadjuvant clinical trial, a method is developed under the principal stratification framework to identify and estimate the average treatment effects on the long-term outcome, conditional on the counterfactual status of the post-treatment intermediate response.

1.1.2 Privacy Concerns in Distributed Data Networks

In the modern context, new challenges arise in the research of causal inference due to data restrictions. Data privacy has become an important issue with the establishment of multiple distributed research networks in large scale studies (Fleurence et al., 2014; Hripcsak et al., 2015; Platt et al., 2018; Donohue et al., 2021). These distributed networks collect sensitive subject-level data and store them at individual research sites (e.g., hospitals). Effective statistical and machine learning approaches are hence needed to be developed to jointly analyze data across sites, without directly utilizing subject-level information. Chapter 3 introduces a tree-based model averaging approach to improve the estimation accuracy of conditional average treatment effects at a target site by leveraging models derived from other potentially heterogeneous sites, without them sharing subject-level data.

1.1.3 Timeliness and Fairness Concerns in Decision Making

There's been a growing concern around the timeliness and fairness of individualized decision making algorithms. For example, there may exist sensitive variables that are important to the intervention decision, but their inclusion in decision making is prohibited due to reasons such as delayed availability, fairness, or other concerns. Robust individualized decision rules that take into account the variation caused by the unavailability of these sensitive variables are needed. On the other hand, in most existing works such as Manski (2004); Qian and Murphy (2011), individualized decision rules aim to maximize the potential average performance. Consequently, certain groups may get unfairly or unsafely treated due to the heterogeneity in their response to the treatment. These problems can impact people's lives in direct and important ways like loan approvals or the length of a sentence in a court case. It is therefore an imperative task to develop fairness-aware decision learning methods. In Chapter 4, we propose a robust individualized decision learning framework with sensitive variables to improve the worst-case outcomes of individuals caused by sensitive variables that are unavailable at the time of decision.

1.2 Outline and Contributions

This section lists the chapters and corresponding contributions. Each chapter aims to be a self contained exposition on a specific topic; as a result, some introductory material for particular chapters are similar in scope. In the following of the dissertation, I develop several statistical and machine learning methods to address the aforementioned challenges in causal inference under uncertainty and data restrictions, with applications to neoadjuvant randomized trials, distributed data networks, and robust individualized decision making.

Chapter 2 concerns identifying and estimating causal effects that involve a post-treatment intermediate response in neoadjuvant randomized clinical trials. In neoadjuvant trials, early efficacy of a treatment is assessed via the binary pathological complete response (pCR) and the eventual efficacy is assessed via long-term clinical outcomes such as survival. Although pCR is strongly associated with survival, it has not been confirmed as a surrogate endpoint. To fully understand its clinical implication, it is important to establish causal estimands such as the causal effect in survival for patients who would achieve pCR under the new regimen. Under the principal stratification framework, previous studies focus on sensitivity analyses by varying model parameters in an imposed model on counterfactual outcomes. Under mild assumptions, we propose an approach to identify and estimate those model parameters using empirical data and subsequently the causal estimand of interest. We also extend our approach to address censored outcome data. The proposed method is applied to a recent clinical trial and its performance is evaluated via simulation studies. This chapter has been accepted for publication in the *Proceedings of the First Conference on Causal Learning and Reasoning (CLeaR'22*, Tan et al., 2022a).

Chapter 3 concerns improving the estimation accuracy of personalized treatment effects by leveraging models rather than subject-level data from heterogeneous data sources. Accurately estimating personalized treatment effects within a study site (e.g., a hospital) has been challenging due to limited sample size. Furthermore, privacy considerations and lack of resources prevent a site from leveraging subject-level data from other sites. We propose a tree-based model averaging approach to improve the estimation accuracy of conditional average treatment effects (CATE) at a target site by leveraging models derived from other potentially heterogeneous sites, without them sharing subject-level data. To our best knowledge, there is no established model averaging approach for distributed data with a focus on improving the estimation of treatment effects. Specifically, under distributed data networks, our framework provides an interpretable tree-based ensemble of CATE estimators that joins models across study sites, while actively modeling the heterogeneity in data sources through site partitioning. The performance of this approach is demonstrated by a real-world study of the causal effects of oxygen therapy on hospital survival rate and backed up by comprehensive simulation results. This chapter has been accepted for publication in the *Proceedings of* the 39th International Conference on Machine Learning (ICML'22, Tan et al., 2022b).

Chapter 4 introduces RISE, a robust individualized decision learning framework with <u>sensitive</u> variables, where sensitive variables are collectible data and important to the intervention decision, but their inclusion in decision making is prohibited due to reasons such as delayed availability or fairness concerns. A naive baseline is to ignore these sensitive variables in learning decision rules, leading to significant uncertainty and bias. To address this, we propose a decision learning framework to incorporate sensitive variables during *offline* training but do not include them in the input of the learned decision rule during model deployment. Specifically, from a causal perspective, the proposed framework intends to improve the worst-case outcomes of individuals caused by sensitive variables that are unavailable at the time of decision. Unlike most existing literature that uses mean-optimal objectives, we propose a robust learning framework via finding a newly defined quantile- or infimum-optimal decision rule. The reliable performance of the proposed method is demonstrated through synthetic experiments and three real-data applications. An earlier version of this work has been made available online (Tan et al., 2022c).

2.0 Identifying Principal Stratum Causal Effects Conditional on a Post-treatment Intermediate Response

2.1 Introduction

We have seen a major shift in the conduct of breast cancer clinical trials in recent years. Traditionally, breast cancer patients are randomly assigned to control or treatment after the primary surgery. Patients from the two groups are then followed over years for comparison of their long-term outcomes such as disease-free survival and overall survival. However, in recent years, there have been an increasing number of neoadjuvant trials where many of the systemic therapies are administered prior to the breast surgery (FDA, 2014).

The primary endpoint in neoadjuvant breast cancer clinical trials is pathological complete response (pCR), a binary indicator of absence of invasive cancer in the breast and auxiliary nodes (FDA, 2014). The rationale for using pCR is that efficacy of a treatment can be assessed at the time of surgery instead of the typical 5-10 years of follow-up on survival endpoints in the adjuvant setting. Strong association between pCR and survival has been well documented (Cortazar et al., 2014; Von Minckwitz et al., 2012; Song et al., 2021), making pCR an attractive candidate surrogate. In the latest guidance of the U.S. Food and Drug administration (FDA), pCR is accepted as an endpoint to support accelerated drug approvals, provided certain requirements are met (FDA, 2014). It is important to decipher the causal relationship among treatment, pCR, and survival in order to interpret the efficacy in survival when pCR is involved.

In the recently published National Surgical Adjuvant Breast and Bowel Project (NSABP) B-40 trial, patients with operable human epidermal growth factor receptor 2 (HER2)-negative breast cancer were randomly assigned to receive or not to receive bevacizumab along with their neoadjuvant chemotherapy regimens (Bear et al., 2012). The addition of bevacizumab significantly increased the rate of pCR (28.2% without bevacizumab vs. 34.5% with bevacizumab, p-value = 0.02). In terms of the long-term outcomes, patients on bevacizumab showed improvements in event-free survival (EFS) and overall survival (OS) compared to the control patients (EFS: hazard ratio 0.80, p-value = 0.06; OS: hazard ratio 0.65, p-value = 0.004) (Bear et al., 2015). Some investigators are interested in the comparison of survival between pCR patients in the treatment group and pCR patients in the control group. Such comparison, however, is problematic because these two groups of pCR patients are different and any direct comparison between them lacks causal interpretation.

Under the counterfactual framework (Rubin, 1974), potentially a patient has a pCR status after taking the control regimen and a pCR status after taking the treatment. Similarly, one can define counterfactual outcomes and causal effects in survival status (0/1) after a certain time period such as three years. The principal stratum framework proposed by Frangakis and Rubin (2002) can be used to describe causal effect in long-term outcomes (such as EFS) with an intermediate outcome (such as pCR) involved. Each principal stratum consists of subjects with the same pair of potential pCR status: the pCR status under the control regimen and the pCR status under the treatment regimen. One can then define the causal effect of treatment in EFS on each principal stratum.

In this chapter, we propose a method to identify and estimate principal stratum causal effects for a binary outcome and later extend our method for censored outcome data. The causal estimand of interest is the treatment efficacy in 3-year EFS and OS among patients who would achieve pCR under chemotherapy plus bevacizumab as in our motivating study, the NSABP B-40 trial. A model of counterfactual outcome given the observed data is imposed. Using some probabilistic arguments, we connect the model parameters with quantities that can be empirically estimated from the observed data. The resulting equations allow us to estimate the model parameters and subsequently the causal estimand of interest, and resolve the identifiability issue.

The remaining chapter is organized as follows. Section 2.2 presents related work in principal stratum causal effects. Section 2.3 introduces the standard data settings, causal estimands of interest, and a regression model in the context of a randomized neoadjuvant trial. In Section 2.4, we provide key assumptions for identification of the causal estimand and introduce the proposed method. In Section 2.5, we conduct a simulation study to assess the performance of our method in terms of bias and coverage of bootstrap confidence intervals. In Section 2.6, we apply the proposed method to the motivating NSABP B-40 study. We

conclude with a discussion of the proposed method and future work in Section 2.7.

2.2 Related Work

Frangakis and Rubin (2002) propose to split study population into principal strata. Each principal stratum is by definition independent of treatment assignment since it contains information on counterfactual, or potential outcomes rather than the observed outcome for a specific treatment assignment. One can then define treatment effects on each principal stratum. Additionally, any union of the basic principal strata would also be a valid principal stratum as it leads to comparisons among a common set of individuals. Gilbert et al. (2015) show the principal stratification framework is useful for evaluating whether and how treatment effects differs across subgroups characterized by the intermediate variable, thus being firmly associated with the utility of the treatment marker.

Identification of principal stratum causal effects is in general difficult. A major challenge is that we do not observe the individual membership of principal stratum because of its counterfactual nature (Gilbert and Hudgens, 2008; Wolfson and Gilbert, 2010). Under the principal stratification framework, Gilbert et al. (2003) propose to perform sensitivity analyses by varying model parameters in an imposed parametric model for counterfactual outcomes. Shepherd et al. (2006) and Jemiai et al. (2007) extend this sensitivity analyses approach by including baseline covariates in the model. These sensitivity analyses can provide researchers with a range of causal estimates under different values of the sensitivity parameters. In reality, however, it is often unclear what the plausible values are for these sensitivity parameters and the selected combinations may not be exhaustive. Li et al. (2010) and Zigler and Belin (2012) use Bayesian approaches to model the joint distribution of the counterfactual intermediate outcomes and long-term outcomes and incorporate prior information regarding non-identifiable associations. The lack of identifiability, however, still exists and is reflected by the over-coverage of confidence intervals in their simulation studies.

Principal stratum causal effects with regards to outcomes truncated by death are not identifiable without further assumptions (Zhang and Rubin, 2003; Kurland et al., 2009; Lee et al., 2010). Tchetgen Tchetgen (2014) identify causal effects by borrowing information from post-treatment risk factors of the intermittent outcome and the causal estimand may vary according to the selected risk factors. Instrumental variables are also introduced to provide information on the unobserved principal strata and the justification of that exclusion restriction assumption is often challenging (Ding et al., 2011; Wang et al., 2017).

All the above methods either fall into sensitivity analyses or require exclusion restriction assumptions. In this chapter, we propose a method to identify and estimate principal stratum causal effects under data settings as Shepherd et al. (2006) for a binary outcome and later extend our method to address issues of censored outcome data under mild assumptions. Identification of the causal effect is achieved with the bias minimal and the coverage probabilities close to the nominal levels.

2.3 The Principal Stratification Framework of Interest

2.3.1 Standard Setting for Neoadjuvant Studies

Consider a neoadjuvant breast cancer clinical trial where patients are randomized to two treatment groups. For subject i = 1, 2, ..., n, let $Z_i \in \{0, 1\}$ be the binary treatment assignment; $X_i \in \Gamma = \{0, 1, ..., K\}$ be a baseline discrete covariate. A continuous baseline variable X_i such as clinical tumor size, would be grouped into K + 1 categories based on scientific knowledge. We will discuss extensions to the scenarios with a continuous X_i in Section 2.7. Throughout this paper, we assume that the stable unit treatment value assumption (SUTVA) (Rubin, 1980) holds: the potential outcomes of any individual i are unrelated to the treatment assignment of other individuals. Then we can denote $S_i(Z_i) \in \{0, 1\}$ as a binary postrandomization intermediate response such as the pCR status for subject i under treatment Z_i (possibly counterfactual). And denote $Y_i\{Z_i, S_i(Z_i)\} = Y_i(Z_i) \in \{0, 1\}$ as a binary long-term outcome of interest such as the EFS status at 3-year after study entry for subject i under treatment Z_i (possibly counterfactual). For individual i, $\{Z_i, X_i, S_i(Z_i), Y_i(Z_i)\}$ represents the observed data of treatment assignment, baseline covariate, intermediate response and long-term outcome. If $Z_i = 0$, $\{S_i(0), Y_i(0)\}$ are observed and $\{S_i(1), Y_i(1)\}$ are counterfactual. If $Z_i = 1$, then $\{S_i(1), Y_i(1)\}$ are observed and $\{S_i(0), Y_i(0)\}$ are counterfactual. Thus for individual *i*, the complete counterfactual data would be $\{Z_i, X_i, S_i(0), S_i(1), Y_i(0), Y_i(1)\}$. Another important assumption is the monotonicity assumption: $S_i(0) \leq S_i(1)$ (Angrist et al., 1996), as in the motivating NSABP B-40 study, addition of bevacizumab led to improved pCR (Bear et al., 2012). We also assume for subject *i*, the treatment assignment Z_i is independent of X_i and the potential outcomes.

Under the principal stratification framework, denote the principal strata to be $E_{jk} = \{i : S_i(0) = j, S_i(1) = k\}, j, k = 0, 1$. The principal stratum causal effects of interest are

$$\theta_{jk} = \mathbb{E}\{Y_i(1) - Y_i(0) | i \in E_{jk}\}, \quad j, k = 0, 1.$$

Under the monotonicity assumption, the principal stratum E_{10} is empty. In the NSABP B-40 study, we are interested in the causal effect in $E_{01} \cup E_{11}$, those who would achieve pCR had they been treated with chemotherapy plus bevacizumab:

$$\theta = \mathbb{E}\{Y_i(1) - Y_i(0) | i \in E_{+1} = E_{01} \cup E_{11}\} = \mathbb{E}\{Y_i(1) - Y_i(0) | S_i(1) = 1\}.$$

Other principal stratum causal effects such as θ_{jk} can be estimated using a similar approach as we outline in Section 2.4.

2.3.2 Modeling a Counterfactual Outcome

In order to estimate the principal stratum causal effects, Gilbert et al. (2003) propose to use a logistic regression model for $\Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0)\}$ as

$$\Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0)\} = \operatorname{logit}^{-1}\{\beta_0 + \beta_1 Y_i(0)\}.$$

Shepherd et al. (2006) further extend the logistic regression by incorporating baseline covariates X_i as

$$\Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0), X_i = x\} = \text{logit}^{-1}\{\beta_0 + \beta_1 Y_i(0) + \beta_2 x\}$$
$$= \frac{\exp\{\beta_0 + \beta_1 Y_i(0) + \beta_2 x\}}{1 + \exp\{\beta_0 + \beta_1 Y_i(0) + \beta_2 x\}}.$$
(1)

Jemiai et al. (2007) consider a more general model framework:

$$\Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0), X_i = x\} = w[r(x) + g\{Y_i(0), x\}]$$

where $w(u) \equiv \{1 + \exp(-u)\}^{-1}$ and $g(\cdot, \cdot)$ is a known function. In the case of Shepherd et al. (2006), $g(u, v) = \beta_1 u$ with β_1 known. Jemiai et al. (2007) show that under the monotonicity assumption, inference could be made on θ for any fixed function g and sensitivity analyses could be performed by varying g.

2.4 The Proposed Method

2.4.1 Key Identification Assumptions

Identification of causal effects is achieved through two key assumptions. First, the monotonicity assumption: $S_i(0) \leq S_i(1)$ (Angrist et al., 1996). That is, a subject who responds under the control would respond if given the treatment. This monotonicity assumption could prove valuable (Bartolucci and Grilli, 2011) and can be justified in many scenarios that the additional therapy would help to improve the response. In the motivating NSABP B-40 study, addition of bevacizumab led to improved pCR (Bear et al., 2012). Second, a parametric model is used to describe the counterfactual response under the treatment for a control non-respondent. Both the future long-term outcome and a baseline covariate are predictors in this parametric model. It is required that the level of the covariates is at least of the same dimension of model parameters and the imposed linearity assumption is critical to identify and estimate those regression parameters. We will elaborate the second assumption in Section 2.4.2.

2.4.2 Identification of Model Parameters and Causal Estimands

As mentioned in Shepherd et al. (2006) and will be described in Section 2.4.4, when the parameters of model (1) are identified, the causal estimands can be identified.

Lemma 2.4.1. For any $x \in \Gamma = \{0, 1, ..., K\}$ and $y \in \{0, 1\}$, let $a_x = \Pr\{S(1) = 1 | S(0) = 0, X = x\}$ and $b_{xy} = \Pr\{Y(0) = y | S(0) = 0, X = x\}$. Let $\mathbf{a} = (a_0, a_1, ..., a_K)^T$ and $\mathbf{b}_y = (b_{y0}, b_{y1}, ..., b_{yK})^T$. Define $h_x(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) = a_x - \sum_{y=0}^1 b_{xy} \operatorname{logit}^{-1}\{\beta_0 + \beta_1 y + \beta_2 x\}$, and $H(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) = \{h_0(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1), ..., h_K(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1)\}^T$.

If rank $\{\partial H(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) / \partial \boldsymbol{\beta}\} = 3$, within the neighborhood of $\boldsymbol{\beta}$ there is a unique solution $\boldsymbol{\beta} = \psi(\mathbf{a}, \mathbf{b}_0, \mathbf{b}_1)$ such that $H\{\psi(\mathbf{a}, \mathbf{b}_0, \mathbf{b}_1), \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1\} = 0$.

Proof. For all $x \in \Gamma$, we have

$$a_x = \Pr\{S(1) = 1 | S(0) = 0, X = x\} = \sum_{y=0}^{1} \Pr\{S(1) = 1, Y(0) = y | S(0) = 0, X = x\}$$
$$= \sum_{y=0}^{1} \Pr\{Y(0) = y | S(0) = 0, X = x\} \Pr\{S(1) = 1 | Y(0) = y, S(0) = 0, X = x\}$$
$$= \sum_{y=0}^{1} b_{xy} \operatorname{logit}^{-1}(\beta_0 + \beta_1 y + \beta_2 x).$$

Hence, $H(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) = 0$ and $H(\cdot)$ is a smooth function of $\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0$, and \mathbf{b}_1 . By invoking the implicit function theorem, when $\operatorname{rank}(\partial H/\partial \boldsymbol{\beta}) = 3$, there exists a smooth function ψ such that $\boldsymbol{\beta} = \psi(\mathbf{a}, \mathbf{b}_0, \mathbf{b}_1)$ and $H\{\psi(\mathbf{a}, \mathbf{b}_0, \mathbf{b}_1), \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1\} = 0$.

The identifiability of model parameter $\boldsymbol{\beta}$ depends on the availability of $a_x = \Pr\{S(1) = 1 | S(0) = 0, X = x\}$ and $b_{xy} = \Pr\{Y(0) = y | S(0) = 0, X = x\}$, for $x \in \Gamma; y = 0, 1$. The linearity in X = x in model (1) also plays an important role. In general, when $\beta_2 \neq 0$ and $K \geq 2$, there are equal or more equations than the number of unknown parameters in $\boldsymbol{\beta}$, Lemma 2.4.1 would hold. In practice, given $(\mathbf{a}, \mathbf{b}_0, \mathbf{b}_1)$, one solves for $\boldsymbol{\beta}$ such that $H(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) = 0$. Then verify that rank $\{\partial H(\boldsymbol{\beta}, \mathbf{a}, \mathbf{b}_0, \mathbf{b}_1) / \partial \boldsymbol{\beta}\} = 3$ at the solution.

2.4.3 Estimation of Causal Estimands

The causal estimand of interest is

$$\theta = \mathbb{E}\{Y_i(1) - Y_i(0) | S_i(1) = 1\} = \mathbb{E}\{Y_i(1) | S_i(1) = 1\} - \mathbb{E}\{Y_i(0) | S_i(1) = 1\}.$$
 (2)

Because $\{Y_i(1), S_i(1)\}$ are observed for subjects in the treatment arm, $\Pr\{Y_i(1) = 1 | S_i(1) = 1\}$ can be estimated by

$$\widehat{\Pr}\{Y_i(1) = 1 | S_i(1) = 1\} = \frac{\sum_i \mathbb{1}\{Z_i = 1, S_i(1) = 1, Y_i(1) = 1\}}{\sum_i \mathbb{1}\{Z_i = 1, S_i(1) = 1\}}.$$
(3)

where $\mathbb{1}\{\cdot\}$ is the indicator function.

Meanwhile,

$$\Pr\{Y_i(0) = 1 | S_i(1) = 1\} = \frac{\Pr\{S_i(1) = 1, Y_i(0) = 1\}}{\Pr\{S_i(1) = 1\}}$$
$$= \frac{\sum_x \Pr\{S_i(1) = 1, Y_i(0) = 1 | X_i = x\} \cdot \Pr\{X_i = x\}}{\sum_x \Pr\{S_i(1) = 1 | X_i = x\} \cdot \Pr\{X_i = x\}}$$
(4)

In equation (4), $\Pr\{X_i = x\}$ can be estimated by $\widehat{\Pr}\{X_i = x\} = \sum_i \mathbb{1}(X_i = x)/n$ and

$$\Pr\{S_{i}(1) = 1, Y_{i}(0) = 1 | X_{i} = x\}$$

$$= \sum_{j=0}^{1} \Pr\{S_{i}(1) = 1, Y_{i}(0) = 1, S_{i}(0) = j | X_{i} = x\}$$

$$= \sum_{j=0}^{1} \Pr\{S_{i}(1) = 1, Y_{i}(0) = 1 | S_{i}(0) = j, X_{i} = x\} \cdot \Pr\{S_{i}(0) = j | X_{i} = x\}$$

$$= \sum_{j=0}^{1} \left[\Pr\{S_{i}(1) = 1 | S_{i}(0) = j, Y_{i}(0) = 1, X_{i} = x\} \cdot \Pr\{Y_{i}(0) = 1 | S_{i}(0) = j, X_{i} = x\} \cdot \Pr\{S_{i}(0) = j | X_{i} = x\} \right].$$
(5)

In equation (5), $\Pr\{Y_i(0) = 1 | S_i(0) = j, X_i = x\}, j = 0, 1$, can be estimated by

$$\widehat{\Pr}\{Y_i(0) = 1 | S_i(0) = j, X_i = x\} = \frac{\sum_i \mathbbm{1}\{Z_i = 0, S_i(0) = j, Y_i(0) = 1, X_i = x\}}{\sum_i \mathbbm{1}\{Z_i = 0, S_i(0) = j, X_i = x\}}.$$

By the monotonicity assumption, $\Pr\{S_i(1) = 1 | S_i(0) = 1, Y_i(0) = 1, X_i = x\} \equiv 1$. The estimation of $\Pr\{S_i(j) = 1 | X_i = x\}, j = 0, 1$, is described in Lemma 2.4.2.

Lemma 2.4.2. Under the monotonicity assumption, for any x, we denote

$$\widehat{q}_j(x) = \frac{\sum_i \mathbb{1}\{Z_i = j, S_i(j) = 1, X_i = x\}}{\sum_i \mathbb{1}\{Z_i = j, X_i = x\}}, \quad j = 0, 1;$$

the observed proportions of responders in the control group and the treatment group with X = x, respectively.

We use maximum likelihood estimation to estimate $Pr\{S_i(j) = 1 | X_i = x\}, j = 0, 1.$

- (a) when $\hat{q}_1(x) \ge \hat{q}_0(x)$, the maximum likelihood estimate of $\Pr\{S_i(j) = 1 | X_i = x\}$ is $\hat{q}_j(x)$, j = 0, 1;
- (b) when $\hat{q}_1(x) < \hat{q}_0(x)$, the maximum likelihood estimate of $\Pr\{S_i(j) = 1 | X_i = x\}$ is $\sum_i \mathbb{1}(S_i = 1, X_i = x) / \sum_i \mathbb{1}(X_i = x), \ j = 0, 1.$

In the second scenario, the estimates are the same as the pooled proportion of responders among patients with X = x. The proof of Lemma 2.4.2 is presented in Appendix A.1.

The last item in equation (4) needed for estimating the causal estimand is $\Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0) = 1, X_i = x\}$. Gilbert et al. (2003) and Shepherd et al. (2006) conduct sensitivity analyses by varying the values of the β in model (1). In Section 2.4.4, we will discuss how to estimate β using a probabilistic equation.

2.4.4 Estimation of Model Parameters

Let

$$G_L(x) = \Pr\{S_i(1) = 1 | S_i(0) = 0, X_i = x\}$$

$$G_R(x, y) = \Pr\{Y_i(0) = y | S_i(0) = 0, X_i = x\}$$

$$G_M(x, y; \beta) = \Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0) = y, X_i = x\}.$$

This leads to an equation system:

$$G_L(x) = \sum_{y=0}^{1} G_M(x, y; \boldsymbol{\beta}) \cdot G_R(x, y); x \in \Gamma$$

We can estimate $G_L(x)$ with the following empirical estimates from the observed data by

$$\widehat{G}_L(x) = \frac{\widehat{\Pr}\{S_i(0) = 0, S_i(1) = 1 | X_i = x\}}{\widehat{\Pr}\{S_i(0) = 0 | X_i = x\}}$$

where the numerator and the denominator are derived from Lemma 2.4.2. The details are presented in Appendix A.1.

Because $\{X_i, S_i(0), Y_i(0)\}$ are observed for subjects in the control arm, $G_R(x, y)$ can be estimated by

$$\widehat{G}_R(x,y) = \frac{\sum_i \mathbbm{1}\{Z_i = 0, S_i(0) = 0, Y_i(0) = y, X_i = x\}}{\sum_i \mathbbm{1}\{Z_i = 0, S_i(0) = 0, X_i = x\}}$$

With $\widehat{G}_L(x)$ and $\widehat{G}_R(x, y)$ estimated from the observed data and $G_M(x, y; \beta)$ specified as the regression model in equation (1), we have

$$\widehat{G}_L(x) = \sum_{y=0}^{1} G_M(x, y; \boldsymbol{\beta}) \cdot \widehat{G}_R(x, y); \quad x \in \Gamma$$
(6)

The number of unknown parameters β in system of equations (6) is three and the number of equations is (K + 1), for $X_i \in \Gamma = \{0, 1, \dots, K\}$. For (6), when K + 1 < 3, we cannot uniquely solve for β . When K + 1 = 3, the number of equations is the same as the number of unknown parameters and in general we can solve for β . When K + 1 > 3, there are more equations than the number of unknown parameters, and there are generally no exact solutions to the equation systems (6). In that case, we propose to estimate β by

$$\widehat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \sum_{x=0}^{K} \{ \widehat{G}_L(x) - \sum_{y=0}^{1} G_M(x,y;\boldsymbol{\beta}) \cdot \widehat{G}_R(x,y) \}^2$$
(7)

where $\widehat{G}_L(x)$, $\widehat{G}_R(x,y)$ and $G_M(x,y;\beta)$ are probabilities bounded between 0 and 1.

With β estimated, we can estimate the causal estimand θ via the procedure outlined in Section 2.4.3.

2.4.5 Consistency of Model Parameters and Causal Estimands

Here we provide the theoretical guarantee of our estimators β and θ . Let

$$Q_0^{(x)}(\boldsymbol{\beta}) = \{G_L(x) - \sum_{y=0}^{1} G_M(x, y; \boldsymbol{\beta}) \cdot G_R(x, y)\}^2, \quad x \in \Gamma; y = 0, 1$$
$$\tilde{Q}_0(\boldsymbol{\beta}) = \{Q_0^{(0)}(\boldsymbol{\beta}), Q_0^{(1)}(\boldsymbol{\beta}), \dots, Q_0^{(K)}(\boldsymbol{\beta})\}^T,$$
$$Q_n(\boldsymbol{\beta}) = \sum_{x=0}^{K} Q_n^{(x)}(\boldsymbol{\beta}) = \sum_{x=0}^{K} \{\widehat{G}_L(x) - \sum_{y=0}^{1} G_M(x, y; \boldsymbol{\beta}) \cdot \widehat{G}_R(x, y)\}^2$$

Theorem 2.4.3. Under the following conditions:

(a) β satisfies $Q_0^{(x)}(\beta) = 0, \forall x \in \Gamma = \{0, 1, \dots, K\}.$

- (b) rank $|\partial \tilde{Q}_0(\boldsymbol{\beta})/\partial \boldsymbol{\beta}| \ge \dim(\boldsymbol{\beta}).$
- (c) $\widehat{G}_L(x) \xrightarrow{p} G_L(x), \widehat{G}_R(x,y) \xrightarrow{p} G_R(x,y), \text{ as } n \to \infty, \forall x \in \Gamma; \forall y = 0, 1.$

Then $\widehat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} Q_n(\boldsymbol{\beta}) \xrightarrow{p} \boldsymbol{\beta}$ and the causal estimand $\widehat{\theta} \xrightarrow{p} \theta$ as $n \to \infty$.

The detailed proof of Theorem 2.4.3 is presented in Appendix A.2.

2.4.6 Extension to Censored Data

As in the motivating NSABP B-40 study, the long-term outcome Y_i may be subject to right censoring. For any time $T = t_0$ of interest, the binary counterfactual outcomes would be $\{Y_i(0; t_0), Y_i(1; t_0)\}$ and the causal estimand can be formulated as

$$\theta(t_0) = \mathbb{E}\{Y_i(1; t_0) - Y_i(0; t_0) | i \in E_{+1}\}$$

With Y_i subject to censoring, $\Pr\{Y_i(1;t_0) = 1 | i \in E_{+1}\}$ can be estimated by the Kaplan-Meier (KM) estimates at time $T = t_0$. The estimation is similar for other relevant quantities such as $\Pr\{Y_i(0;t_0) = 1 | S_i(0) = j, X_i = x\}$ in equation (5) under the scenario where $Y_i(Z_i)$ is always observed.
2.5 Simulation Studies

A simulation study is used to assess the performance of the proposed method. The setup is chosen to resemble the NSABP B-40 study by simulating treatment assignment, baseline tumor size category, binary pCR response status, and binary survival status, specifically:

$$\mathfrak{D} = [D_i = \{Z_i, X_i, S_i(0), S_i(1), Y_i(0), Y_i(1)\}, \ i = 1, \dots, n]$$

We simulate the subject-level data as follows. First, we simulate the categorical baseline tumor category X_i from a multinomial distribution with $\Pr\{X_i = x\} = 0.25, x \in \{0, 1, 2, 3\}$. Next, we simulate $S_i(0)$ given X_i from a Bernoulli distribution with $\Pr\{S_i(0) = 1 | X_i = x\} = p(x)$ with p(0), p(1), p(2), p(3) = 0.3, 0.25, 0.25, 0.2, respectively. We then simulate the survival status under control, $Y_i(0)$, with a Bernoulli draw with $\Pr\{Y_i(0) = 1 | S_i(0) = 0, X_i = x\} = 0.7, 0.65, 0.6, 0.55$ for x = 0, 1, 2, 3, respectively and $\Pr\{Y_i(0) = 1 | S_i(0) = 1, X_i = x\} = 0.84, 0.78, 0.72, 0.66$ for x = 0, 1, 2, 3, respectively. The choice of these numbers reflects a 20% improvement in 3-year EFS for respondents over nonrespondents under the control regimen.

Next, we simulate the conditional distribution $\{S_i(1)|S_i(0), Y_i(0), X_i\}$. For subjects with $S_i(0) = 1$ we set $S_i(1)$ to be 1 to enforce the monotonicity assumption. For subjects with $S_i(0) = 0$ we draw $S_i(1)$ from a Bernoulli distribution: $\Pr\{S_i(1) = 1|S_i(0) = 0, Y_i(0) = y, X_i = x\} = \text{logit}^{-1}\{\beta_0 + \beta_1 y + \beta_2 x\}$. We try different settings for $\boldsymbol{\beta} = (-3, -5, 0.2), (-5, -1, -2), \text{ and } (-7, 3, 0.2).$

We then simulate the survival status under treatment, $Y_i(1)$, according to the following probability distributions:

$$\begin{aligned} &\Pr\{Y_i(1) = 1 | S_i(0) = 0, S_i(1) = 0, Y_i(0) = 0\} = 0.5, \\ &\Pr\{Y_i(1) = 1 | S_i(0) = 0, S_i(1) = 0, Y_i(0) = 1\} = 0.6, \\ &\Pr\{Y_i(1) = 1 | S_i(0) = 0, S_i(1) = 1, Y_i(0) = 0\} = 0.85 \\ &\Pr\{Y_i(1) = 1 | S_i(0) = 0, S_i(1) = 1, Y_i(0) = 1\} = 0.9, \\ &\Pr\{Y_i(1) = 1 | S_i(0) = 1, S_i(1) = 1, Y_i(0) = 0\} = 0.85 \\ &\Pr\{Y_i(1) = 1 | S_i(0) = 1, S_i(1) = 1, Y_i(0) = 1\} = 0.9. \end{aligned}$$

These probabilities are chosen to make the 3-year EFS under treatment greater for those who would obtain pCR under treatment than those who would not, and have a greater 3-year EFS for those patients who would be event-free under control than those who would not be event-free under control. We set these probabilities to be independent of the baseline tumor size given the potential outcomes $\{S_i(0), S_i(1), Y_i(0)\}$.

Lastly we simulate the treatment assignment with equal probability for each arm as a Bernoulli draw with $\Pr\{Z_i = 0\}$ and $\Pr\{Z_i = 1\}$ both equal to 0.5 to ensure that independence between potential outcomes and treatment assignment. For the simulated data the true average causal effect for principal stratum $S_i(1) = 1$, $\mathbb{E}\{Y_i(1) - Y_i(0)|S_i(1) = 1\}$, can be calculated using the above parameters for simulations. The detailed calculations is given in Appendix A.3. Under the three parameter settings the true values of the causal estimands are $\theta=0.179$, 0.130, and 0.120, respectively. This means that under the three different settings, if the treatment was administered to all subjects who would achieve pCR under treatment there would be a 17.9%, 13.0%, 12.0% increment in survival respectively, within the time frame under consideration, than had all of them taken the control instead.

Under each parameter setting and a chosen sample size n=1000, 2000, or 4000, we simulate R=1000 replicates. A quasi-Newton method, the Broyden-Fletcher-Goldfarb-Shanno algorithm, is used for the optimization. We create B=500 bootstrap samples to obtain the 95% confidence interval for the causal estimates. Let $\hat{\theta}^{(r)}$ be the mean estimate among bootstrap samples from the r replicate, $r = 1, \ldots, R$.

We construct bootstrap confidence intervals to account for the variability introduced by estimating model parameters. We use the basic bootstrap CI, or the pivotal CI (Davison and Hinkley, 1997) for constructing CIs from bootstrap estimates. Let $\{\hat{\theta}^{(1)}, \hat{\theta}^{(2)}, \ldots, \hat{\theta}^{(B)}\}$ are the causal effect estimates from *B* bootstrap samples. Denote $\theta^*_{(1-\alpha/2)}$ and $\theta^*_{(\alpha/2)}$ as the $100(1 - \alpha/2)\%$ and $100(\alpha/2)\%$ of the bootstrap causal effect estimates. The $100(1 - \alpha)\%$ bootstrap confidence interval is given by $(2\hat{\theta} - \theta^*_{(1-\alpha/2)}, 2\hat{\theta} - \theta^*_{(\alpha/2)})$ where $\hat{\theta}$ is the estimate from the data.

We report the empirical bias, mean squared error (MSE), average length of 95% CIs, and

the coverage of those CIs, where

$$Bias(\widehat{\theta}) = R^{-1} \sum_{r=1}^{R} \{\widehat{\theta}^{(r)} - \theta\},$$

MSE($\widehat{\theta}$) = $R^{-1} \sum_{r=1}^{R} \{\widehat{\theta}^{(r)} - \theta\}^{2},$
95% CI width = $R^{-1} \sum_{r=1}^{R} |\widehat{\theta}^{(r)}_{U,0.05} - \widehat{\theta}^{(r)}_{L,0.05}|,$
95% CI coverage = $R^{-1} \sum_{r=1}^{R} \mathbb{1}\{\theta \in (\widehat{\theta}^{(r)}_{L,0.05}, \widehat{\theta}^{(r)}_{U,0.05})\}$

with $\widehat{\theta}_{L,0.05}^{(r)}$ and $\widehat{\theta}_{U,0.05}^{(r)}$ the lower bound and upper bound of the 95% bootstrap CIs of $\widehat{\theta}$ from the r^{th} simulated dataset. Table 1 shows the simulation results of the proposed method under three different parameter settings and various sample sizes. Our simulation results show the identification of causal effects is achieved with the bias negligible and the coverage probabilities close to the nominal levels.

Sample size	Empirical bias	MSE	95% CI width	95% CI coverage		
Setting 1: $\beta = (-3, -5, 0.2), \theta = 0.179$						
1000	-0.011	3.001e-3	0.206	0.952		
2000	-0.006	1.539e-3	0.155	0.955		
4000	-0.002	6.755e-4	0.116	0.962		
	Setting 2:	$\beta = (-5, -1)$, -2), θ =0.130			
1000	-6.011e-5	2.496e-3	0.185	0.943		
2000	9.358e-4	1.137e-3	0.130	0.948		
4000	1.086e-4	5.462 e-4	0.093	0.950		
Setting 3: $\beta = (-7, 3, 0.2), \theta = 0.120$						
1000	0.008	2.547e-3	0.194	0.955		
2000	0.006	1.319e-3	0.141	0.957		
4000	0.003	6.363e-4	0.100	0.953		

Table 1: Simulation results of the proposed method under three different parameter settings and various sample sizes.

2.6 Application to NSABP B-40 Trial

2.6.1 B-40 Data Analysis

Here we apply the proposed method to the NSABP B-40 study (Bear et al., 2012, 2015). Among the 1206 enrolled participants, 13 withdrew consent, 7 had missing data and 2 had had inoperable disease after chemotherapy. Another 15 patients did not have nodal assessment so their pCR status was not ascertained. We conduct our analysis among the rest 1169 patients. Our purpose is to estimate the causal treatment effect in 3-year EFS and OS among patients who would obtain a pCR had bevacizumab been added to their treatment regimen. KM estimates are used since there are 61 patients censored at 3 years.

To apply our method, the clinical tumor size is used as the baseline auxiliary covariate X. Patients are grouped into four nearly equal-sized groups: 2-3 cm, 3.1-4 cm, 4.1-6 cm and >6 cm, based on breast cancer expert knowledge. We code these four tumor size groups into $\{0, 1, 2, 3\}$, respectively. Among the 589 patients in the control arm, the proportions of those who achieved pCR in each patient group are 28%, 23%, 22% and 17%, respectively; among the 580 patients in the treatment arm, the proportions of those who achieved pCR are 31%, 26%, 25% and 27%, respectively. This does not violate the monotonicity assumption $S_i(0) \leq S_i(1)$. The 3-year long-term outcome status $Y_i = 1$ if the patient *i* survived within the first 3 years and 0 otherwise.

We calculate the 95% bootstrap confidence intervals from 500 bootstrap samples. The estimated causal treatment effect in 3-year EFS among those who would obtained pCR under treatment is $\hat{\theta}_{\text{EFS}} = 0.180 \ (95\% \text{ CI}=(0.056, 0.377))$ with $\hat{\beta} = (-1.797, -5.874, 0.285)$. The estimated causal treatment effect in 3-year OS among those who would obtained pCR under treatment is $\hat{\theta}_{\text{OS}} = 0.175 \ (95\% \text{ CI}=(0.062, 0.354))$ with $\hat{\beta} = (-1.85, -4.764, 0.289)$. For both scenarios, because 0 is outside of the 95% CIs, we would claim that the addition of bevacizumab improves 3-year EFS and OS among patients who would respond to neoadjuvant chemotherapy plus bevacizumab at a 95% confidence level.

2.6.2 Sensitivity of Initial Parameters in Optimization

For the real data application, the initial estimate $\beta_{init} = (\beta_0, \beta_1, \beta_2)$ is set at (0, 0, 0). To see the sensitivity of initial parameters, we try $9261 = 21 \times 21 \times 21$ different initial values of β_{init} , with β_0 , β_1 , and β_2 on the integer grids of $[-10, 10] \times [-10, 10] \times [-10, 10]$. The corresponding histograms of causal estimates in 3-year EFS and 3-year OS at convergence are presented in Figure 1. Our estimated model parameters $\hat{\beta}$ in Section 2.6.1 achieves the minimum loss of equation (7). Except for some extreme initialization such as (10,10,10), most of the $\hat{\theta}$ are the same or very close to the causal estimates calculated by using β_{init} = (0,0,0) as initial parameters. Therefore, we conclude that the causal estimand is not sensitive to the initial parameter settings in optimization. In practice, we suggest running optimization with various initial values and identify the right estimate.



Figure 1: Histogram of the causal estimates obtained from different initial values of model parameters in the optimization process for 3-year EFS (Figure A) and 3-year OS (Figure B), respectively. Except for some extreme initialization, most of the causal estimates are the same or very close to the causal estimate calculated by using zeros as initial parameters.

2.6.3 Comparisons to Sensitivity Analysis Method

We compare the performance of our method with that of the sensitivity analysis similar to Gilbert et al. (2003) and Shepherd et al. (2006). Recall that for $X = x \in \Gamma = \{0, 1, ..., K\}$, we have an equation system:

$$\widehat{G}_L(x) = \sum_{y=0}^1 G_M(x,y;\boldsymbol{\beta}) \cdot \widehat{G}_R(x,y); \quad x \in \Gamma = \{0, 1, \dots, K\}$$

where $G_M(x, y; \boldsymbol{\beta}) = \text{logit}\{\beta_0 + \beta_1 y + \beta_2 x\}$. In the sensitivity analysis we vary the value of β_1 from -7 to -3. Then for each category of x we define $\beta_x = \beta_0 + \beta_2 x$. Under this reparameterization we have only one unknown parameter, β_x , for each equation. We then solve for β_x for each equation independently and obtain the causal estimand subsequently. By varying values of β_1 around the estimated $\hat{\beta}_1$ from Section 2.6.1, the corresponding causal estimands in 3-year EFS and 3-year OS are presented in Table 2. The estimated causal effects in 3-year EFS vary from 0.159 to 0.181 with none of the 95% CIs including 0; the estimated causal effects in 3-year OS vary from 0.132 to 0.176 with none of the 95% CIs including 0. These intervals overlap a lot with the confidence intervals of real data. These results suggest the addition of bevacizumab may improve 3-year EFS and 3-year OS among patients who would respond to neoadjuvant chemotherapy plus bevacizumab.

Long-term survival		$\widehat{ heta}$	95% CI for $\widehat{\theta}$
EFS	-7	0.181	(0.025, 0.290)
	-6	0.180	(0.043, 0.289)
	-5	0.178	(0.040, 0.282)
	-4	0.172	(0.058, 0.272)
	-3	0.159	(0.065, 0.267)
OS	-7	0.176	(0.055, 0.278)
	-6	0.172	(0.067, 0.267)
	-5	0.166	(0.069, 0.267)
	-4	0.153	(0.066, 0.235)
	-3	0.132	(0.064, 0.200)

Table 2: Sensitivity analysis for the estimated causal effect of bevacizumab in 3-year survival among those who would obtain pCR under chemotherapy plus bevacizumab.

2.7 Discussion and Future Work

We have proposed a method under the principal stratification framework to estimate causal effects of a treatment on a binary long-term endpoint conditional on a post-treatment binary marker in randomized controlled clinical trials. We also extend our method to address censored outcome data. In our motivating study, we demonstrate the causal effect of the new regimen in the long-term survival for patients who would achieve pCR. Other principal stratum causal effects can be estimated in a similar fashion. Our approach can play an important role in a sensitivity analysis.

Identification of causal effects is achieved through two assumptions. First, a subject who responds under the control would respond if given the treatment. This monotonicity assumption could prove valuable (Bartolucci and Grilli, 2011) and can be justified in many scenarios that the additional therapy would help to improve the response. When the auxiliary variable X is discrete, we can identify and estimate $\Pr\{S(1) = 1 | S(0) = 0, X\}$ under the monotonicity assumption. Second, a parametric model is used to describe the counterfactual response under the treatment for a control non-respondent (Shepherd et al., 2006). Both the future long-term outcome and a baseline covariate are predictors in this parametric model. Shepherd et al. (2006) does not consider when the auxiliary X is discrete, the parameters of model (1) can be identified when the level of the discrete covariate is at least of the same dimension of model parameters. Instead they perform sensitivity analyses by varying the values of those model parameters in order to estimate the causal estimands. It is recognized that no diagnostic tool is available to verify the validity of this counterfactual model.

In the motivating dataset, we discretize a continuous baseline variable into several levels. In practice, the linearity assumption may not hold. We would consider a two-pronged approach: 1) to estimate $G_L(x)$ and $G_R(x, y)$ by nonparametric estimates such as spline or kernel density estimates for a univariate continuous X; 2) to use a more flexible model for the counterfactual response such as a logistic regression with natural cubic spline with fixed and even-spaced knots along the domain of X. For each given x, we can still use the same probabilistic argument to link those estimates and the model parameters. The objective function would be a weighted sum of the squared difference of those probabilistic estimates.

3.0 A Tree-based Model Averaging Approach for Personalized Treatment Effect Estimation from Heterogeneous Data Sources

3.1 Introduction

Estimating individualized treatment effects has been a hot topic because of its wide applications, ranging from personalized medicine, policy research, to customized marketing advertisement. Treatment effects of certain subgroups within the population are often of interest. Recently, there has been an explosion of research devoted to improving estimation and inference of covariate-specific treatment effects, or conditional average treatment effects (CATE) at a target research site (Athey and Imbens, 2016; Wager and Athey, 2018; Hahn et al., 2020; Künzel et al., 2019; Nie and Wager, 2021). However, due to the limited sample size in a single study, improving the accuracy of the estimation of treatment effects remains challenging.

Leveraging data and models from various research sites to conduct statistical analyses is becoming increasingly popular (Reynolds et al., 2020; Cohen et al., 2020; Berger et al., 2015). Distributed research networks have been established in many large scale studies (Fleurence et al., 2014; Hripcsak et al., 2015; Platt et al., 2018; Donohue et al., 2021). A question often being asked is whether additional data or models from other research sites could bring improvement to a local estimation task, especially when a single site does not have enough data to achieve a desired statistical precision. This concern is mostly noticeable in estimating treatment effects where sample size requirement is high yet observations are typically limited. Furthermore, information exchange between data sites is often highly restricted due to privacy, feasibility, or other concerns, prohibiting centralized analyses that pool data from multiple sources (Maro et al., 2009; Brown et al., 2010; Toh et al., 2011; Raghupathi and Raghupathi, 2014; DeShazo and Hoffman, 2015; Donahue et al., 2018; Dayan et al., 2021). One way to tackle this challenge is through model averaging (Raftery et al., 1997), where multiple research sites collectively contribute to the tasks of statistical modeling without sharing sensitive subject-level data. Although this idea has existed in supervised learning problems (Dai and Zhang, 2011; McMahan et al., 2017), to our best knowledge, there are no established model averaging approach and theoretical results on estimating CATE in a distributed environment. The extension is non-trivial because CATE is unobserved in nature, as opposed to prediction problems where labels are given.

This chapter focuses on improving the prediction accuracy of CATE concerning a target site by leveraging models derived from other sites where *transportability* (to be formally defined in Section 3.3.1, Pearl and Bareinboim, 2011; Stuart et al., 2011; Pearl and Bareinboim, 2014; Bareinboim and Pearl, 2016; Buchanan et al., 2018; Dahabreh et al., 2019) may not hold. Specifically, there may exist heterogeneity in treatment effects. In the context of our multi-hospital example, these are: 1) **local heterogeneity**: within a hospital, patients with different characteristics may have different treatment effects. This is the traditional notion of CATE; and 2) **global heterogeneity**: where the same patient may experience different treatment effects at different hospitals. The second type of heterogeneity is driven by site-level confounding, and hampers the transportability of models across hospital sites. We also note that these two types of heterogeneity may interact with each other in the sense that transportability is dependent on patient characteristics, which we will address.

We propose a model averaging framework that uses a flexible tree-based weighting scheme to combine learned models from sites that takes into account heterogeneity. The contribution of each learned model to the target site depends on subject characteristics. This is achieved by applying tree splittings (Breiman et al., 1984) at both the site and the subject levels. For example, effects of a treatment in two hospitals may be similar for female patients but not for male, suggesting us to consider borrowing information across sites only on selective subgroups. Our approach extends the classic model averaging framework (Raftery et al., 1997; Wasserman, 2000; Hansen, 2007; Yang, 2001) by allowing data-adaptive weights, which are interpretable in a sense that they can be used to lend credibility to transportability. For example, in the case of extreme heterogeneity where other sites merely contribute to the target, the weights can be used as a diagnostic tool to inform the decision against borrowing information.

Our primary contributions are summarized as follows. 1) We propose a model averaging scheme with interpretable weights that are adaptive to both local and global heterogeneity via tree-splitting dedicated to improving CATE estimation under distributed data networks. 2) We generalize model averaging techniques to study the transportability of causal inference. Causal assumptions with practical implications are explored to warrant the use of our approach. 3) We provide an extensive empirical evaluation of the proposed approach with a concrete real-data example on how to apply the method in practice. 4) Compared to other distributed learning methods, the proposed framework enables causal analysis without sharing subject-level data, is easy to implement, offers ease of operations, and minimizes infrastructure, which facilitates practical collaboration within research networks.

The remaining chapter is organized as follows. In Section 3.2, we present a general formulation of the problem and discuss related work on model averaging and data fusion. We describe the proposed method and assumptions in detail in Section 3.3. The performance of the proposed method is assessed by simulation experiments in Section 3.4 and illustrated through a multi-hospital electronic health data application for critical care medicine in Section 3.5 to estimate conditional treatment effects for oxygen therapy. We conclude the chapter in Section 3.6.

3.2 Related Work

There are two types of construct of a distributed database (Breitbart et al., 1986): homogeneous versus heterogeneous. For homogeneous data sources, data across sites are random samples of the global population. Recent modeling approaches (Lin and Zeng, 2010; Lee et al., 2017; McMahan et al., 2017; Battey et al., 2018; Jordan et al., 2019; Tang et al., 2020a; Wang et al., 2021) all assume samples are randomly partitioned, which guarantees identical data distribution across sites. The goal of these works is to improve overall prediction by averaging results from homogeneous sample divisions. The classic random effects meta-analysis (see, e.g., Whitehead (2002); Sutton et al. (2000); Borenstein et al. (2011) describes heterogeneity using modeling assumptions, but its focus mostly is still on global patterns.

3.2.1 Heterogeneous Models

In practice, however, there is often too much global heterogeneity in a distributed data network to warrant direct aggregation of models obtained from local sites. The focus shifts to improving the estimation of a target site by selectively leveraging information from other data sources. There are two main classes of approaches. The first class is based on comparison of the learned model parameters $\{\hat{\theta}_1, \ldots, \hat{\theta}_K\}$ from K different sites where for site k we adopt model $f_k(\mathbf{x}) = f(\mathbf{x}; \mathbf{\theta}_k)$ with subject features \mathbf{x} to approximate the outcome of interest Y. Clustering and shrinkage approaches are then used by merging data or models that are similar (Ke et al., 2015; Smith et al., 2017; Ma and Huang, 2017; Wang et al., 2020; Tang et al., 2020b). Most of these require the pooling of subject-level data. The second class of approaches falls in the model averaging framework (Raftery et al., 1997) with weights directly associated with the local prediction. Let site 1 be our target site, and the goal is to improve f_1 using a weighted estimator $f^*(\mathbf{x}) = \sum_{k=1}^{K} \omega_k f_k(\mathbf{x})$ with weights ω_k to balance the contribution of each model and $\sum_k \omega_k = 1$. It provides an immediate interpretation of usefulness of each data source. When the weights are proportional to the prediction performance of f_k on site 1, for example,

$$\omega_k = \frac{\exp\{-\sum_{i \in \mathcal{I}_1} (f_k(\boldsymbol{x}_i) - y_i)^2\}}{\sum_{\ell=1}^K \exp\{-\sum_{i \in \mathcal{I}_1} (f_\ell(\boldsymbol{x}_i) - y_i)^2\}},$$

with y_i being the observed outcome of subject *i* in site 1, indexed by \mathcal{I}_1 , the method is termed as the exponential weighted model averaging (EWMA). Several variations of ω_k can be found in Yang (2001); Dai and Zhang (2011); Yao et al. (2018); Dai et al. (2018). In general, separate samples are used to obtain the estimates of ω_k 's and f_k 's, respectively.

Here we focus on the literature review of model averaging. We note that our framework is also related to federated learning (McMahan et al., 2017). But the latter often involves iterative updating rather than a one-shot procedure, and could be hard to apply to nonautomated distributed research networks. Besides, it has been developed mainly to estimate a global prediction model by leveraging distributed data, and is not designed to target any specific site. We further discuss these approaches and other related research topics and their distinctions with model averaging in Appendix B.1.

3.2.2 Transportability

In causal inference, there is a lot of interest in identifying subgroups with enhanced treatment effects, targeting at the feasibility of customizing estimates for individuals (Athey and Imbens, 2016; Wager and Athey, 2018; Hahn et al., 2020; Künzel et al., 2019; Nie and Wager, 2021). These methods aim to estimate the CATE function $\tau(\mathbf{x})$, denoting the difference in potential outcomes between treatment and control, conditional on subject characteristics \boldsymbol{x} . To reduce uncertainty in estimation of personalized treatment effects, incorporating additional data or models are sought after. Pearl and Bareinboim (2011, 2014); Bareinboim and Pearl (2016) introduced the notion of transportability to warrant causal inference models be generalized to a new population. The issue of generalizability is common in practice due to the non-representative sampling of participants in randomized controlled trials (Cook et al., 2002; Druckman et al., 2011; Allcott, 2015; Stuart et al., 2015; Egami and Hartman, 2020). Progress on bridging the findings from an experimental study with observational data can be found in, e.g., Stuart et al. (2015); Kern et al. (2016); Stuart et al. (2018); Ackerman et al. (2019); Yang et al. (2020); Harton et al. (2021). See Tipton and Olsen (2018); Colnet et al. (2020); Degtiar and Rose (2021) and references therein for a comprehensive review. However, most methods require fully centralized data. In contrast, we leverage the distributed nature of model averaging to derive an integrative CATE estimator.

3.3 A Tree-based Model Averaging Framework

We first formally define the conditional average treatment effect (CATE). Let Y denote the outcome of interest, $Z \in \{0,1\}$ denote a binary treatment indicator, and X denote subject features. Correspondingly, let y, z and x denote their realizations. Using the potential outcome framework (Neyman, 1923; Rubin, 1974), we define CATE as $\tau(x) = E[Y^{(Z=1)} - Y^{(Z=0)} | \mathbf{X} = \mathbf{x}]$, where $Y^{(Z=1)}$ and $Y^{(Z=0)}$ are the potential outcomes under treatment arms Z = 1 and Z = 0, respectively. The expected difference of the potential outcomes is dependent on subject features \mathbf{X} . By the causal consistency assumption, the

observed outcome is $Y = ZY^{(Z=1)} + (1 - Z)Y^{(Z=0)}$.

Now suppose the distributed data network consists of K sites, each with sample size of n_k . Site k contains data $\mathcal{D}_k = \{y_i, z_i, \boldsymbol{x}_i\}_{i \in \mathcal{I}_k}$, where \mathcal{I}_k denotes its index set. Its CATE function is given by $\tau_k(\boldsymbol{x}) = E_k[Y^{(Z=1)} - Y^{(Z=0)} | \boldsymbol{X} = \boldsymbol{x}]$, where the expectation is taken over the data distribution in site k. Without loss of generality, we assume the goal is to estimate the CATE function in site 1, τ_1 .

3.3.1 Causal Assumptions

To ensure information can be properly borrowed across sites, we first impose the following idealistic assumptions, and then present relaxed version of Assumption 3.3.2. Let S be the site indicator taking values in $S = \{1, \ldots, K\}$ such that $S_i = k$ if $i \in \mathcal{I}_k$.

Assumption 3.3.1 (Unconfoundedness).

$$\{Y^{(Z=0)}, Y^{(Z=1)}\} \perp Z | \boldsymbol{X}, S;$$

Assumption 3.3.2 (Transportability).

$$\{Y^{(Z=0)}, Y^{(Z=1)}\} \perp S | \boldsymbol{X};$$

Assumption 3.3.3 (Positivity).

$$0 < P(S = 1 | \mathbf{X}) < 1$$
 and $0 < P(Z = 1 | \mathbf{X}, S) < 1$ for all \mathbf{X} and S .

Assumption 3.3.1 ensures treatment effects are unconfounded within sites so that $\tau_k(\boldsymbol{x})$ can be consistently identified. It holds by design when data are randomized controlled trials or when treatment assignment depends on \boldsymbol{X} . By this assumption, we have $\tau_k(\boldsymbol{x}) = E[Y|\boldsymbol{X} = \boldsymbol{x}, S = k, Z = 1] - E[Y|\boldsymbol{X} = \boldsymbol{x}, S = k, Z = 0]$. The equality directly results from the assumption. Assumption 3.3.2 essentially states that the CATE functions are transportable, i.e., $\tau_k(\boldsymbol{x}) = \tau_{k'}(\boldsymbol{x})$ for $k, k' \in \{1, \ldots, K\}$. See also Stuart et al. (2011), Buchanan et al. (2018) and Yang et al. (2020) for similar consideration. This assumption may not be satisfied due to heterogeneity across sites. In other words, site can be a confounder which prevents transporting of CATE functions across sites. Our method allows Assumption 3.3.2 to be violated and use model averaging weights to determine transportability. Explicitly, we consider a relaxed Assumption 3.3.4 to hold for a subset of sites that contains site 1.

Assumption 3.3.4 (Partial Transportability).

$$\{Y^{(Z=0)}, Y^{(Z=1)}\} \perp S_1 | \boldsymbol{X} \rangle$$

Here, S_1 takes values in $S_1 = \{k : \tau_k(\boldsymbol{x}) = \tau_1(\boldsymbol{x})\}$ and $\{1\} \subset S_1 \subset S$. We denote S_1 as the set of transportable sites with regard to site 1. Hence, transportability holds across some sites and specific subjects. In a special case in Section 3.4 where $S_1 = \{1\}$, bias may be introduced to by model averaging. However, our approach is still able to exploit the bias and variance trade off to improve estimation. Assumption 3.3.3 ensures that all subjects are possible to be observed in site 1 and all subjects in all sites are possible to receive either arm of treatment. The former ensures a balance of covariates between site 1 population and the population of other sites. Violation of either one may result in extrapolation and introduce unwanted bias to the ensemble estimates for site 1. This assumption is also used, e.g., in Stuart et al. (2011).

3.3.2 Model Ensemble

We consider an adaptive weighting of $\{\tau_1, \ldots, \tau_K\}$ by

$$\tau^*(\boldsymbol{x}) = \sum_{k=1}^{K} \omega_k(\boldsymbol{x}) \tau_k(\boldsymbol{x})$$
(8)

where τ^* is the weighted model averaging estimator. The weight functions $\omega_k(\boldsymbol{x})$'s are not only site-specific, but also depend on \boldsymbol{x} , and follow $\sum_{k=1}^{K} \omega_k(\boldsymbol{x}) = 1$. It measures the importance of τ_k in assisting site 1 when subjects with characteristics \boldsymbol{x} are of interest. We rely on each of the sites to derive their respective $\hat{\tau}_k$ from \mathcal{D}_k so that $\mathcal{D}_1, \ldots, \mathcal{D}_K$ do not need to be pooled. Only the estimated functions $\{\hat{\tau}_2, \ldots, \hat{\tau}_K\}$ are passed to site 1. We will describe the approaches to estimate $\hat{\tau}_k$ in Section 3.3.5.

A two-stage model averaging approach is proposed. We first split \mathcal{D}_1 , the data in the target site, into a training set and an estimation set indexed by $\{i \in \mathcal{I}_1^{(1)}\}$ and $\{i \in \mathcal{I}_1^{(2)}\}$, respectively. 1) Local stage: Obtain $\hat{\tau}_1$ from subjects in $\mathcal{I}_1^{(1)}$. Obtain $\hat{\tau}_k$ from local subjects

Algorithm 1	Tree-based	model	averaging	for	heterogeneous	data	sources

for k = 1 to K do \triangleright Loop through K sites. Can be run in parallel.Build a local model using site k data. Site 1 model uses its training set only.end forfor $i \in \mathcal{I}_1^{(2)}$ do \triangleright Loop through subjects in site 1 estimation set.for k = 1 to K do \triangleright Loop through K local models.Predict $\hat{\tau}_k(\boldsymbol{x}_i)$ using local model k. $D_{i,k} = [\boldsymbol{x}_i, k, \hat{\tau}_k(\boldsymbol{x}_i)].$ end forend for

Create augmented site 1 data $\mathfrak{D}_{aug,1}$ by concatenating $D_{i,k}$ vectors. $\widehat{\mathcal{T}}_{EF}(\boldsymbol{x},s) = \text{ENSEMBLEFOREST}(\mathfrak{D}_{aug,1}) \qquad \triangleright \text{ Or ENSEMBLETREE when } B = 1.$

in \mathcal{I}_k , $k = 2, \ldots, K$. These $\{\hat{\tau}_k\}_{k=1}^K$ are then passed to site 1 to get K predicted treatment effects for each subject in $\mathcal{I}_1^{(2)}$, resulting in an augmented data set as shown in Figure 2(b). 2) Ensemble stage: A tree-based ensemble model is trained on the augmented data by either an ensemble tree (ET) or an ensemble random forest (EF), with the predicted treatment effects from the previous stage, i.e., $\hat{\tau}_k(\boldsymbol{x}_i)$ as the outcome. The site indicator S of which local model is used as well as the subject features \boldsymbol{x}_i are fed into the ensemble model averaging estimator. Figure 2(a) illustrates a conceptual diagram of the proposed model averaging framework and structure of the augmented data. Note the idea of data augmentation has been used in, e.g., computer vision (Perez and Wang, 2017; Mo et al., 2020, 2021), statistical computing (van Dyk and Meng, 2001), and imbalanced classification (Chawla et al., 2002). Here the technique is being used to construct weights for model averaging, which will be discussed in the following paragraph. Algorithm 1 provides an algorithmic overview. Our method has been implemented as an R package ifedtree available on GitHub (https: //github.com/ellenxtan/ifedtree).



Figure 2: (a) Schema of the proposed algorithm. (b) Illustration of the augmented data constructed from the estimation set of site 1.

3.3.3 Construction of Weights

A tree-based ensemble is constructed to estimate the weighting functions $\{\omega_k\}_{k=1}^K$. Heterogeneity across sites is explained by including the site index into an augmented training set when building trees. An intuition of our approach is that sites that are split away from site 1 (by tree nodes) are ignored and the sites that fall into the same leaf node are considered homogeneous to site 1 hence contribute to the estimation of $\tau_1(\boldsymbol{x})$. A splitting by site may occur in any branches of a tree, resulting in an information sharing scheme across sites that is dependent on \boldsymbol{x} . We construct the ensemble by first creating an augmented data $\mathfrak{D}_{aug,1} = \{\boldsymbol{x}_i, k, \hat{\tau}_k(\boldsymbol{x}_i)\}_{i \in \mathcal{I}_1^{(2)}, k \in \mathcal{S}}$, for subjects in $\mathcal{I}_1^{(2)}$. The illustration of this augmented site 1 data is given in Figure 2(b). An ensemble is then trained on this data by either a tree or a random forest, with the estimated treatment effects $\hat{\tau}_k(\boldsymbol{x}_i)$ as the outcome, and a categorical site indicator of which local model is used along with all subject-level features as predictors, i.e., (\boldsymbol{x}_i, k) . We denote the resulting function as $\mathcal{T}(\boldsymbol{x}, s)$ which depends on both \boldsymbol{x} and site s, specifically, $\mathcal{T}_{\text{ET}}(\boldsymbol{x}, s)$ and $\mathcal{T}_{\text{EF}}(\boldsymbol{x}, s)$ for ensemble tree (ET) and ensemble forest (EF), respectively. Let $\mathcal{L}(\boldsymbol{x}, s)$ denote the final partition of the feature space by the tree to which the pair (\boldsymbol{x}, s) belongs. The ET estimate based on the augmented site 1 data can be derived by

$$\widehat{\mathcal{T}}_{\text{ET}}(\boldsymbol{x},s) = \{ |\{(i,k) : (\boldsymbol{x}_i,k) \in \mathcal{L}(\boldsymbol{x},s)\}_{i \in \mathcal{I}_1^{(2)},k \in \mathcal{S}} |\}^{-1} \sum_{\{(i,k) : (\boldsymbol{x}_i,k) \in \mathcal{L}(\boldsymbol{x},s)\}_{i \in \mathcal{I}_1^{(2)},k \in \mathcal{S}}} \widehat{\tau}_k(\boldsymbol{x}_i)$$

$$= \sum_{i \in \mathcal{I}_1^{(2)}} \sum_{k=1}^K \frac{\mathbb{1}\{(\boldsymbol{x}_i,k) \in \mathcal{L}(\boldsymbol{x},s)\}}{|\mathcal{L}(\boldsymbol{x},s)|} \widehat{\tau}_k(\boldsymbol{x}_i).$$
(9)

Intuitively, observations with similar characteristics $(\boldsymbol{x} \text{ and } \boldsymbol{x}')$ and from similar sites (s and s') are more likely to fall in the same partition region in the ensemble tree, i.e., $(\boldsymbol{x}, s) \in \mathcal{L}(\boldsymbol{x}', s')$ or $(\boldsymbol{x}', s') \in \mathcal{L}(\boldsymbol{x}, s)$. This resembles a non-smooth kernel where weights are $1/|\mathcal{L}(\boldsymbol{x}, s)|$ for observations that are within the neighborhood of (\boldsymbol{x}, s) , and 0 otherwise. The estimator borrows information from neighbors in the space of \boldsymbol{X} and S. The splits of the tree are based on minimizing in-sample MSE of $\hat{\tau}$ within each leaf and pruned by cross-validation over choices of the complexity parameter. Since a single tree is prone to be unstable, in practice, we use random forest to reduce variance and smooth the partitioning boundaries. By aggregating B ET estimates each based on a subsample of the augmented data, $\{\hat{T}^{(b)}\}_{b=1}^{B}$, an EF estimate can be constructed by

$$\widehat{\mathcal{T}}_{EF}(\boldsymbol{x},s) = \frac{1}{B} \sum_{b=1}^{B} \widehat{\mathcal{T}}^{(b)}(\boldsymbol{x},s)$$

$$= \sum_{i \in \mathcal{I}_{1}^{(2)}} \sum_{k=1}^{K} \lambda_{i,k}(\boldsymbol{x},s) \widehat{\tau}_{k}(\boldsymbol{x}_{i}), \qquad (10)$$
where $\lambda_{i,k}(\boldsymbol{x},s) = \frac{1}{B} \sum_{b=1}^{B} \frac{\mathbbm{1}\{(\boldsymbol{x}_{i},k) \in \mathcal{L}_{b}(\boldsymbol{x},s)\}}{|\mathcal{L}_{b}(\boldsymbol{x},s)|}.$

The form of $\widehat{\mathcal{T}}^{(b)}(\boldsymbol{x}, s)$ closely follows (9) but is based on a subsample of $\mathfrak{D}_{aug,1}$. The weights, $\lambda_{i,k}(\boldsymbol{x}, s)$, are similar to that in (9), and can be viewed as kernel weighting that defines an adaptive neighborhood of \boldsymbol{x} and s. We then obtain the model averaging estimates defined in (8) by fixing s = 1 such that $\widehat{\tau}^*_{\text{ET}}(\boldsymbol{x}) = \widehat{\mathcal{T}}_{\text{ET}}(\boldsymbol{x}, s = 1)$ or $\widehat{\tau}^*_{\text{EF}}(\boldsymbol{x}) = \widehat{\mathcal{T}}_{\text{EF}}(\boldsymbol{x}, s = 1)$. The weight functions $\{\omega_k(\boldsymbol{x})\}_{k=1}^K$ for $\widehat{\tau}^*(\boldsymbol{x})$ can be immediately obtained from the ET or EF by

$$egin{aligned} \widehat{ au}_{ ext{ET}}^*(oldsymbol{x}) &= \widehat{\mathcal{T}}_{ ext{ET}}(oldsymbol{x},1) = \sum_{k=1}^K \widehat{\omega}_k(oldsymbol{x}) \widehat{ au}_k(oldsymbol{x}), \ & ext{where} \ \ \widehat{\omega}_k(oldsymbol{x}) &= \sum_{i \in \mathcal{I}_1^{(2)}} rac{\mathbbm{1}\{(oldsymbol{x}_i,k) \in \mathcal{L}(oldsymbol{x},1)\}}{|\mathcal{L}(oldsymbol{x},1)|}; \ & ilde{ au}_{ ext{EF}}(oldsymbol{x}) &= \widehat{\mathcal{T}}_{ ext{EF}}(oldsymbol{x},1) = \sum_{k=1}^K \widehat{\omega}_k(oldsymbol{x}) \widehat{ au}_k(oldsymbol{x}), \ & ext{where} \ \ \widehat{\omega}_k(oldsymbol{x}) &= \sum_{i \in \mathcal{I}_1^{(2)}} \lambda_{i,k}(oldsymbol{x},1). \end{aligned}$$

It can be verified that $\sum_{k=1}^{K} \widehat{\omega}_k(\boldsymbol{x}) = 1$ for all \boldsymbol{x} . As our simulations in Section 3.4 show, $\widehat{\tau}^*$ improves the local functional estimate $\widehat{\tau}_1$. We set B = 2,000 throughout the paper. Tree and forest estimates are obtained by R packages rpart and grf, respectively.

3.3.4 Interpretability of Weights

The choice of tree-based models naturally results in such kernel weighting $w_k(\boldsymbol{x})$ (Athey et al., 2019), which are not accessible by other ensemble techniques. Such explicit and interpretable weight functions could deliver meaningful rationales for data integration. For example, under scenarios where there exists extreme global heterogeneity (as shown in Section 3.4 when c is large), $w_k(\boldsymbol{x})$ can be used as a diagnostic tool to decide which external data sources should be co-used. Weights close to 0 inform against model transportability, and they are adaptive to subject-level features \boldsymbol{x} so that decisions can be made based on the subpopulations of interest.

3.3.5 Local Models: Obtaining $\hat{\tau}_k$

Estimate of $\tau_k(\mathbf{x})$ at each local site must be obtained separately before the ensemble. Our proposed ensemble framework can be applied to a general estimator of $\tau_k(\mathbf{x})$. For each site, the local estimate could be obtained using different methods. Recently, there has been many work dedicated to the estimation of individualized treatment effects (Athey and Imbens, 2016; Wager and Athey, 2018; Hahn et al., 2020; Künzel et al., 2019; Nie and Wager, 2021). As an example, we consider using the causal tree (CT) (Athey and Imbens, 2016) to estimate the local model at each site. CT is a non-linear learner that *(i)* allows different types of outcome such as discrete and continuous, and can be applied to a broad range of real data scenarios; *(ii)* can manage hundreds of features and high order interactions by construction; *(iii)* can be applied to both experimental studies and observational studies by propensity score weighting or doubly robust methods. CT is implemented in the R package **causalTree**. We also explore another estimating option for local models in Appendix B.3.

3.3.6 Asymptotic Properties

We provide consistency guarantee of the proposed estimator $\widehat{\mathcal{T}}_{\text{EF}}$ for the true target τ_1 . Assuming point-wise consistent local estimators are used for $\{\tau_k\}_{k=1}^K$, EF with subsampling procedure described in Appendix B.2 is consistent.

Theorem 3.3.5. Suppose the subsample used to build each tree in an ensemble forest is drawn from different subjects of the augmented data and the following conditions hold:

- (a) Bounded covariates: Features X_i and the site indicator S_i are independent and have a density that is bounded away from 0 and infinity.
- (b) Lipschitz response: the conditional mean function $\mathbb{E}[\mathcal{T}|\mathbf{X} = \mathbf{x}, S = 1]$ is Lipschitzcontinuous.
- (c) Honest trees: trees in the random forest use different data for placing splits and estimating leaf-wise responses.

Then $\widehat{\mathcal{T}}_{EF}(\boldsymbol{x},1) \xrightarrow{p} \tau_1(\boldsymbol{x})$, for all \boldsymbol{x} , as $\min_k n_k \to \infty$. Hence, $\widehat{\tau}^*_{EF}(\boldsymbol{x}) \xrightarrow{p} \tau_1(\boldsymbol{x})$.

The conditions and a proof of Theorem 3.3.5 is given in Appendix B.2. To demonstrate the consistency properties of our methods, we add in Appendix B.3 oracle versions of ET and EF estimators, denoted as ET-oracle and EF-oracle, which use the ground truth of local models $\{\tau_k\}_{k=1}^K$ in estimating $\{\widehat{\omega}_k\}_{k=1}^K$. This removes the uncertainty in local models. The remaining uncertainty only results from the estimation of the ensemble weights, and we see both oracle estimators achieve minimal MSE. Section 3.4 gives a detailed evaluation of the finite sample performance.

3.4 Simulation Studies

Monte Carlo simulations are conducted to assess the proposed methods. We specify $m(\boldsymbol{x},k)$ as the conditional outcome surface and $\tau(\boldsymbol{x},k)$ as the conditional treatment effect for individuals with features \boldsymbol{x} in site k. The treatment propensity is specified as $e(\boldsymbol{x}) = \Pr(Z = 1 | \boldsymbol{X} = \boldsymbol{x})$. The potential outcomes can be written as $Y_i = m(\boldsymbol{X}_i, S_i) + \{Z_i - e(\boldsymbol{X}_i)\}\tau(\boldsymbol{X}_i, S_i) + \epsilon_i$, following notations in Robinson (1988); Athey and Imbens (2016); Wager and Athey (2018); Nie and Wager (2021). The mean function is $m(\boldsymbol{x}, k) = \frac{1}{2}x_1 + \sum_{d=2}^{4} x_d + (x_1 - 3) \cdot c \cdot U_k$, and the treatment effect function is specified as

$$\tau(\mathbf{x}, k) = \mathbb{1}\{x_1 > 0\} \cdot x_1 + (x_1 - 3) \cdot c \cdot U_k,$$

where $z = 0, 1, U_k$ denotes the global heterogeneity due to site-level confounding, controlled by a scaling factor c, and $\epsilon_i \sim N(0, 1)$. Features follow $\mathbf{X}_i \sim N(\mathbf{0}, \mathbf{I}_D)$, where D = 5, and are independent of ϵ_i . The simulation setting within each site (with k fixed) is motivated by designs in Athey and Imbens (2016). Features in τ are determinants of treatment effect while those in m but not in τ are prognostic only. The data are generated under a distributed data networks. We assume there are K = 20 sites in total, each with a sample size n = 500. In our main exposition, we consider an experimental study design where treatment propensity is $e(\mathbf{x}) = 0.5$, i.e., individuals are randomly assigned to treatment and control. Variations of the settings above are discussed, with results presented in Appendix B.3. Two types for global heterogeneity are considered by the choice of U_k . For **discrete** grouping, we assume there are two underlying groups among the K sites $U_k \sim Bernoulli(0.5)$. Specifically, we assume odd-index sites and even-index sites form two distinct groups $\mathcal{G}_1 = \{1, 3, \ldots, K-1\}$; $\mathcal{G}_2 = \{2, 4, \ldots, K\}$ such that $U_{k \in \mathcal{G}_1} = 0$ and $U_{k \in \mathcal{G}_2} = 1$. Sites from similar underlying groupings have similar treatment effects and mean effects, while sites from different underlying groupings have different treatment effects and mean effects. For **continuous** grouping, we consider $U_k \sim Unif[0, 1]$. We vary the scales of the global heterogeneity under the discrete and continuous cases, respectively, with c taking values $c \in \{0, 0.6, 1, 2\}$. A c = 0 implies all data sources are homogeneous. In other words, Assumption 3.3.2 is satisfied when c = 0 but not when c > 0.

3.4.1 Compared Estimators and Evaluation

The proposed approaches ET and EF are compared with several competing methods. **LOC:** A local CT estimator that does not utilize external information. It is trained on \mathcal{I}_1 only, combining training and estimation sets. **MA:** A naive model averaging method with weights $\omega_k^{\text{MA}} = 1/k$. This approach assumes models are homogeneous. **EWMA:** We consider a modified version of EWMA that can be used for CATE. We obtain an approximation of $\tau_1(\boldsymbol{x})$ by fitting another local model using the estimation set of site 1, denoted by $\tilde{\tau}_1(\boldsymbol{x})$. Its weights are given by

$$\omega_k^{\text{EWMA}} = \frac{\exp\{-\sum_{i \in \mathcal{I}_1^{(2)}} (\widehat{\tau}_k(\boldsymbol{x}_i) - \widetilde{\tau}_1(\boldsymbol{x}_i))^2\}}{\sum_{\ell=1}^K \exp\{-\sum_{i \in \mathcal{I}_1^{(2)}} (\widehat{\tau}_\ell(\boldsymbol{x}_i) - \widetilde{\tau}_1(\boldsymbol{x}_i))^2\}}.$$

STACK: A stacking ensemble, which is a linear ensemble of predictions of several models (Breiman, 1996). To our end, we regress $\tilde{\tau}_1(\boldsymbol{x})$ on the predictions of the estimation set in site 1 from each local model, $\{\hat{\tau}_1(\boldsymbol{x}), \ldots, \hat{\tau}_k(\boldsymbol{x})\}$. The stacking weights are not probabilistic hence not directly interpretable. We report the empirical mean squared error (MSE) of these methods over an independent testing set of sample size $n_{te} = 2000$ from site 1. $\text{MSE}(\hat{\tau}) = n_{te}^{-1} \sum_{i=1}^{n_{te}} \{\hat{\tau}(\boldsymbol{x}_i) - \tau_1(\boldsymbol{x}_i)\}^2$. Each simulation scenario is repeated for 1000 times. Experiments are performed on a 6-core Intel Xeon CPU E5-2620 v3 2.40GHz equipped with 64GB RAM.



Figure 3: Box plots of MSE ratios of CATE estimators, respectively, over LOC, for (a) discrete grouping and (b) continuous grouping across site. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve smaller MSE ratios compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings.

		c = 0	c = 0.6	c = 1	c = 2
Discrete grouping	n = 100	0.57	0.59	0.61	0.59
	n = 500	0.12	0.17	0.17	0.16
	n = 1000	0.07	0.12	0.12	0.13
Continuous grouping	n = 100	0.54	0.59	0.63	0.69
	n = 500	0.11	0.24	0.31	0.34
	n = 1000	0.08	0.17	0.21	0.26

Table 3: MSE ratios of EF over LOC. As n increases, model averaging becomes more powerful due to better estimation of τ_k , and is more pronounced when c is small.

3.4.2 Estimation Performance

Figure 3 shows the performance of the proposed estimators and the competing estimators, using LOC as the benchmark. The proposed ET and EF show the best performance in terms of the mean and variation of MSE among other estimators when c > 0, and comparable to equal weighting MA when c = 0. Although, a forest is more stable than a tree in practice, both ET and EF give similar results because the true model is relatively simple and can be accurately estimated by a single ensemble tree under the given sample size.

Although asymptotically consistency, under finite sample, bias exists in local models and leads to biased model averaging estimates. While explicit quantification of bias and variance remains challenging due to extra uncertainty carried forward from the local estimates, we demonstrated that the proposed estimators can improve upon the local models under small sample size via Table 3. It shows the MSE ratio of EF over LOC as a measure of gain resulting from model averaging by varying n = 100,500,1000. The decrease in MSE ratio as n increases, regardless of the choice of c, is consistent with our asymptotic results in Theorem 3.3.5. This is due to a bias-and-variance trade-off in the ensemble that ensures a small MSE, which remains smaller than that in LOC despite varying n. It also shows our method is robust to the existence of local uncertainty.

3.4.3 Visualization of Information Borrowing

Figure 4 visualizes the proposed ET and EF. In (a) and (d), the site indicator and X_1 appear as splitting variables in the ETs, which is consistent with the data generation process. The estimated treatment effect (b) and (e) reveals the pattern of transportability across sites and with respect to X_1 . Panels (c) and (f) plot the model averaging weights in EFs over X_1 . Site 1 has a relatively large contribution to the weighted estimator while models from other sites have different contributions at different values of X_1 depending on their similarity in $\tau(\boldsymbol{x}, k)$ to that in site 1. Corresponding ET and EF show consistent patterns.





Figure 4: Visualization of simulation results under **discrete grouping** (a,b,c) and continuous grouping (d,e,f) when c = 1. (a) and (d) visualize the proposed ETs where the site indicator and X_1 are selected as splitting variables, which is consistent with the underlying data generation process. (b) and (e) show the predicted treatment effects of the proposed EFs varying X_1 in each site, marginalized over all other features. (b) is arranged according to the true grouping, odd sites versus even sites. The plot recovers the pattern of local and global heterogeneity. (c) and (f) plot the interpretable model averaging weights in EFs over X_1 . The weights of site 1 have a relatively large contribution to the weighted estimator while models from other sites have different contributions for different X_1 depending on their similarity in $\tau(\mathbf{x}, k)$ to that in site 1. Corresponding ET and EF show consistent patterns and recover the true grouping.

3.4.4 Additional Simulations

The detailed results of these additional simulations are included in Appendix B.3.

1) Connection to supervised learning. The uniqueness of averaging $\tau_k(\boldsymbol{x})$ as opposed to supervised learning that averages prediction models $f_k(\boldsymbol{x})$ is that the outcome of $f_k(\boldsymbol{x})$ is immediately available. In our case, an additional estimation step is needed to construct the model averaging weights. We provide a comparison among estimators that utilize the ground truth $\{\tau_k(\boldsymbol{x})\}_{k=1}^K$ (denoted as "-oracle") when computing ensemble weights. This mimics the case of supervised learning where weights are based on observed outcomes. Oracle methods achieve smaller MSE ratios; the pattern is consistent with Table 3.

2) Simulation under observational studies. We also consider the treatment generation mechanism under an observational design. Specifically, the propensity is given as $e(\mathbf{x}) = \exp(0.6x_1)$. We consider both a correctly specified propensity model using a logistic regression of Z on X_1 and a misspecified propensity model with a logistic regression of Z on all \mathbf{X} . In general, the proposed estimators obtain the best performance with similar results as in Figure 3. With the correctly specified propensity score model, the local estimator is consistent in estimating $\tau_k(\mathbf{x})$, the proposed framework is valid. When the propensity model is misspecified, extra uncertainty is carried forward from the local estimates, but the proposed estimators can still improve upon LOC. This is due to a bias-and-variance trade-off that leads to small MSE, which remains smaller than the local models.

3) Covariate dimensions. Besides D = 5, we consider other choices of covariate dimension including D = 20, 50. With a higher dimension, the MSE ratio between the proposed estimates and LOC estimates increases but the same pattern across methods persists.

4) Unequal sample size at each site. In the distributed date network, different sites may have a different sample size n_k . Those with a smaller sample size may not be representative of their population, leading to an uneven level of precision for local causal estimates. We consider a simulation setting where site 1 has a sample size of $n_1 = 500$ while other site n_2, \ldots, n_K has a sample size of 200. Results show that the MSE ratio between the proposed estimates and LOC estimates increases compared to the scenario where the sample size in all sites are 500. However, the proposed estimators still enjoy the most robust performance. This also shows our method is robust to the existence of local uncertainty.

5) Different local estimators. We stress that other consistent estimators could be used as the local model. Options such as causal forest (Wager and Athey, 2018) are explored varying the sample size at local sites. Similar performance is observed as in Figure 3.

6) Further comparisons to non-adaptive ensemble. Here we provide a brief discussion of the implications of the proposed method and how it differs from non-adaptive methods such as stacking. Although unrealistic, when the true weights are non-adaptive, the performance may be similar. Plus, our learned weights can be used to examine adaptivity, as shown in Figure 4(c,f) and Figure 5(c). Stacking is shown to be more robust than non-adaptive model averaging in case of model misspecification. See discussion in Clarke (2003). Our additional simulation results show that in case of a large global heterogeneity, as c increases, the heterogeneity across sites gets larger, reducing the influence of important covariates on heterogeneity, hence the weights become more non-adaptive. However, the proposed methods still enjoy a comparable performance to STACK, which further indicates the robustness of the proposed methods.

3.5 Example: A Multi-Hospital Data Network

Application with contextual insights is provided based on an analysis of the eICU Collaborative Research Database, a multi-hospital database published by Philips Healthcare (Pollard et al., 2018). The analysis is motivated by a recent retrospective study that there is a higher survival rate when SpO₂ is maintained at 94-98% among patients requiring oxygen therapy (van den Boom et al., 2020), not "the higher the better". We use the same data extraction code to create our data. We consider SpO₂ within this range as treatment (Z = 1) and outside of this range as control (Z = 0). A total of 7,022 patients from 20 hospitals, each with at least 50 patients in each treatment arm, are included with a randomly selected target (hospital 1). Hospital-level summary information is provided in Appendix B.4. Patient-level features include age, BMI, sex, Sequential Organ Failure Assessment (SOFA) score, and duration of oxygen therapy. The outcome is hospital survival (Y = 1) or death (Y = 0). Figure 5 visualizes the performance of EF-based estimated effect of oxygen therapy setting on in-hospital survival. CT is used as the local model with propensity score modeled by a logistic regression. Figure 5(a) shows the propensity score-weighted **average survival** for those whose received treatment is consistent with the estimated decision. Specifically, the expected reward is given by

$$\frac{\sum_{i} Y_i 1(Z_i = Z_i^{est}) / \pi(Z_i, \boldsymbol{X}_i)}{\sum_{i} 1(Z_i = Z_i^{est}) / \pi(Z_i, \boldsymbol{X}_i)},$$

where $Z_i^{est} = 1(\hat{\tau} > 0)$ denotes the estimated treatment rule and $\pi(Z_i, \mathbf{X}_i)$ is the probability of receiving the actual treatment. We provide expected reward for the 1) observed treatment assignment (baseline), 2) LOC-based rule, and 3) EF-based rule. The treatment rule based on our method can increase mean survival by 3% points compared to baseline, and is more promising than LOC.

In the fitted EF, the hospital indicator is the most important, explaining about 50% of the decrease in training error. Figure 5(b) shows the estimated CATE varying two important features, BMI and oxygen therapy duration. Patients with BMI around 36 and duration above 400 show the most benefit from oxygen therapy in the target SpO₂ range. Patients with BMI between 20 and 30 and duration around 200 may not benefit from such alteration. Figure 5(c) visualizes the data-adaptive weights $\omega_k(\boldsymbol{x})$ in the fitted EF with respect to BMI for different models, while holding other variables constant. The weights of hospital 1 are quite stable while models from other sites may have different contribution to the weighted estimator for different values of BMI. Judging from hospital information in Appendix B.4, hospitals with a larger bed capacity tend to be similar to hospital 1, and are shown to provide larger contributions.

In this distributed research network, different hospitals have a different sample size. For sensitivity analysis, we consider a weighting strategy to adjust for the sample size of site k. Results show similar patterns as in Figure 5. Detailed results are provided in Appendix B.4. The real-data access is provided in Appendix B.5.



Figure 5: Application to estimating treatment effects of oxygen therapy on survival. (a) Expected survival of treatment decision following different estimators. The proposed EF shows the largest gain in improving survival rate, more promising than LOC and baseline. (b) Estimated treatment effects varying duration and BMI, two important features in the fitted EF. Patients with a BMI around 35, and a duration above 400 benefited the most. (c) Visualization of data-adaptive weights in the estimated EF varying BMI. Hospitals with a larger bed capacity tend to contribute more, the data of which might be more similar to hospital 1.

3.6 Discussion

We have proposed an efficient and interpretable tree-based model averaging framework for enhancing treatment effect estimation at a target site by borrowing information from potentially heterogeneous data sources. We generalize standard model averaging scheme in a data-adaptive way such that the generated weights depend on subject-level features. This work makes multi-site collaborations and especially treatment effect estimation more practical by avoiding the need to share subject-level data. Our approach extends beyond causal inference to estimating a general $f(\mathbf{x})$ from heterogeneous data.

Unlike in classic model averaging where prediction performance can be assessed against observed outcomes or labels, treatment effects are not directly observed. While our approach is guaranteed to be consistent under randomized studies, the weights are estimated based on expected treatment effects, hence relying on Assumption 3.3.1 (unconfoundedness) to hold. It may be a strong assumption in observational studies with unmeasured confounding.

4.0 Robust Individualized Decision Learning with Sensitive Variables

4.1 Introduction

Recently, there has been a widespread interest in developing methodology for individualized decision rules (IDRs) based on observational data. When deriving IDRs, some collectible data are important to the intervention decision, while their inclusion in decision making is prohibited due to reasons such as delayed availability or fairness concerns. For example, sensitive characteristics of subjects regarding their income, sex, race and ethnicity may not be appropriate to be used directly for decision making due to fairness concerns. In the medical field especially for patients in severe life-threatening conditions such as sepsis, timely bedside intervention decisions have to be made before lab measurements are ordered, assayed and returned to the attending physicians. However, due to the delayed availability of lab results, most of the decisions are made with great uncertainty and bias due to partial information at hand. We define *sensitive variables* as variables whose inclusion into decision rules is prohibited. The formal definition of sensitive variables will be given in Section 4.3.

In this work, we propose RISE (Robust Individualized decision learning with SEnsitive variables), a robust IDR framework to improve the outcome of individuals when there are informative yet sensitive variables that are either not available or prohibited from using during IDR deployment. To achieve this, we propose to estimate the optimal IDR by optimizing a quantile- or infimum-based objective, respectively, for continuous or discrete sensitive variables. Our idea falls along the lines of work that considers algorithmic fairness (Dwork et al., 2012) while extending it to the setting of causal inference (Rubin, 2005) in the sense that decisions are driven by causality rather than a general utility function. We show in our empirical analyses that this leads to fairer and safer real-life decisions with little sacrifice of the overall performance. This optimization problem is then shown to be equivalent to a weighted classification problem where most existing machine learning classifiers can be readily applied.

Assuming that a large outcome value is preferable, optimal IDRs are traditionally de-

rived through maximizing the mean outcome of the sample population. In this paper, we are interested in a specific yet broadly applicable setting of learning that involves sensitive variables. We consider offline learning where sensitive variables are collected and *can be used in training the IDRs*, but they *cannot be used as input in the resulting IDRs*. This is a setting commonly considered in the fairness and privacy-related literature for classification (e.g., Kamiran and Calders (2012)), but not from a causal standpoint. When there exist important variables that are simply left out from training, the estimated IDR will be biased. This bias can be removed if all important variables are used during training, which we will show in Section 4.3.1 a mean-optimal approach. The optimal action maximizes the mean outcome where the mean is taken over the sensitive variables, conditioning on other variables. This method, however, has no control of the disparity in sensitive variables. Subjects with different sensitive values may report large outcome differences, hence unfairly or unsafely treated. Therefore, objective functions with robustness guarantees for sensitive variables are preferred, since they offer protection to subjects in the lower tail of the outcome distribution with regards to the sensitive values.

For illustration, we consider a toy example with binary actions, $A \in \{-1, 1\}$. We note that the decision can only be made based on the variable X whereas S is a sensitive variable. The setup is shown in Table 4 and the oracle values under the mean-optimal rule and RISE ¹ are given in Table 5. The detailed setup can be found in Section 4.4.1 under Example 1. We consider *vulnerable subjects* as those with low outcome values, as highlighted in red in Table 4 (A full definition is given in Section 4.3.2). For $X \leq 0.5$, the mean-optimal rule would assign action A = 1 as it tries to achieve the largest average reward across S = 0 and S = 1. Recall that S is not available at the time of decision. However, this action results in great harm for subjects with S = 1 as they could get the worst expected outcome of 0. On the contrary, RISE improves the worst-case outcome by assigning A = -1, protecting the vulnerable subjects. Likewise, for X > 0.5, the mean-optimal rule assigns A = -1 while the RISE assigns A = 1 protecting those with S = 0 that could have experienced an outcome of 5. Compared to the mean-optimal rule, the proposed rule achieves a larger reward among

¹For the mean-optimal rule, overall average reward is calculated by (30 + 0 + 5 + 27)/4 = 15.5, reward among vulnerable subjects is calculated by (0 + 5)/2 = 2.5; for RISE, overall average reward is calculated by (11 + 13 + 15 + 13)/4 = 13, reward among vulnerable subjects is calculated by (13 + 15)/2 = 14.

vulnerable subjects while maintaining a comparable overall reward.

	$X \le 0.5$		_	X > 0.5			
	S = 0	S = 1		S = 0	S = 1		
A = -1	11	13		5	27		
A = 1	30	0		15	13		

Table 4: Toy example setup of E(Y|X, S, A).

Table 5: Toy example results.

	Average reward		
	Overall	Vulnerable	
Mean-optimal rule	15.5	2.5	
RISE	13	14	

Main contributions. The main contributions of the proposed framework are summarized as follows. *Methodology-wise*, 1) we propose a novel framework, RISE, to handle sensitive variables in causality-driven decision making. Robustness is introduced to improve the worst-case outcome caused by sensitive variables, and as a result, it reduces the outcome variation across subjects. The latter is directly associated with fairness and safety in decision making. To the best of our knowledge, we are among the first to propose a robust-type fairness criterion under causal inference. 2) We introduce a classification-based optimization framework that can easily leverage most existing classification tools catered to different functional classes, including state-of-the-art random forest, boosting, or neural network models. *Application-wise*, 3) the consideration of sensitive variables in decision learning is important to applications in policy, education, healthcare, etc. Specifically, we illustrate the application of RISE using three real-world examples from fairness and safety perspectives where robust decision rules are needed, across which we have observed robust performance of the proposed approach. From a fairness perspective, we consider a job training program where age is considered as a sensitive variable. From a safety perspective, we consider two applications to healthcare where lab measurements are considered as sensitive variables.

The remaining chapter is organized as follows. Section 4.2 discusses related work. Section 4.3 describes the proposed RISE framework in detail. The performance of the proposed framework is evaluated by simulation studies and applied to three real-data applications in Section 4.4. We conclude and discuss future work in Section 4.5.

4.2 Related Work

Our work focuses on individualized decision rules, which aim at assigning treatment decision based on subject characteristics. Existing methods for deriving IDRs include model-based methods such as Q-learning (Watkins and Dayan, 1992; Murphy, 2003; Moodie et al., 2007; Chakraborty et al., 2010; Goldberg and Kosorok, 2012; Song et al., 2015) and A-learning (Robins et al., 2000; Murphy, 2005), model-free policy search methods such as Robins et al. (2008); Orellana et al. (2010a,b); Zhang et al. (2012); Zhao et al. (2012, 2015), and contextual bandit methods Bietti et al. (2021); Li et al. (2011). In Appendix C.1, we provide additional literature review on general IDRs under causal settings. Fairness, safety and robustness are topics of interest that extend well beyond causal inference. In the following, we provide a review of these areas, with focus given to work related to causal inference and IDRs.

Fairness and safety in IDRs. The consideration of fairness and safety in machine learning has seen an explosion of interest in the past few years Dwork et al. (2012); Varshney (2016); Barocas et al. (2019); Nabi and Shpitser (2018); Hashimoto et al. (2018); Chouldechova and Roth (2020); Mehrabi et al. (2021); Pessach and Shmueli (2022), especially for solving classification and regression problems to help derive decisions that are not only accurate but also fair. In these work, sensitive variables are also referred to as sensitive, protected, or auxiliary attributes. We extend the definition of sensitive variables to include delayed information that is not available at deployment as it is also suitable for this framework.

Among earlier work, preprocessing (Kamiran and Calders, 2012; Feldman et al., 2015; Creager et al., 2019; Sattigeri et al., 2019) and inprocess training approaches (Beutel et al., 2017; Hashimoto et al., 2018; Lahoti et al., 2020) consider disentangling the input X from a known or unknown sensitive variable S so that the transformed X does not contain any information related to S. Due to the causal nature of IDRs, effect of IDRs cannot be estimated consistently when an informative S is left out and the resulting rule is suboptimal. This follows from the classic argument that any unmeasured confounding (i.e., S), if not accounted for, would lead to bias. Similar issues persist in contextual bandits (Joseph et al., 2016; Patil et al., 2020). Makar et al. (2022) considers reducing the impact of auxiliary variables on prediction under distributional shift. Although it is motivated from a causal idea, its main focus is still on prediction. Inside the causal framework, Zhang and Bareinboim (2018); Nabi et al. (2019) extend fairness from prediction to policy learning using causal graphical models by incorporating fairness constraints. Chen et al. (2022) considers counterfactual fairness that seeks to achieve conditional independence of the decisions via data preprocessing. Despite earlier efforts in bringing fairness into the causal framework, most of these approaches only ensure mean zero disparity in S but do not have robustness guarantees in the sense that the variance of the disparity in S is not controlled. Besides, most examples consider a single categorical sensitive variable, but not multiple or continuous ones.

Robustness in IDRs. Recently the statistical literature has witnessed a growing interest in developing robust methods for estimating IDRs. They introduce robustness into the objective function by using quantile-optimal treatment regimes or mean-optimal treatment regimes under certain constraints to improve the gain of individuals at the lower tail of the reward spectrum (Wang et al., 2018a,b; Qi et al., 2022, 2019; Fang et al., 2022). In particular, Wang et al. (2018a); Qi et al. (2022) propose to estimate quantile- or tail-optimal treatment regimes. Wang et al. (2018b) studies the mean-optimal treatment regime under a constraint to control for the average potential risk. Qi et al. (2019) proposes a decision-rulebased optimized covariates dependent equivalent for tasks of individualized decision making. Fang et al. (2022) considers mean and quantile objectives simultaneously by maximizing the average value with the guarantee that its tail performance exceeds a prespecified threshold. Robustness, in their sense, pertains to the outcome distribution subject to the sampling error. When sensitive variables are present, we consider instead the robustness of the outcome distribution subject to the uncertainty due to sensitive variables, providing a more targeted way of ensuring robustness, which is directly related to fairness and safety. Compared to algorithms based on explicit fairness constraints (for example Zafar et al. (2017); Zhang et al. (2018) in classification and Zhang and Bareinboim (2018); Chen et al. (2022) in causal inference) that seek to remove the disparity across different values of S, our method reduces the variance of disparity across S. In addition, constraint-based approaches typically require specialized optimization procedures whereas our approach presents an elegant and systematic way for optimization. To our knowledge, we are the first few to consider decision fairness via a robust objective under the causal framework.

4.3 Robust Decision Learning Framework with Sensitive Variables

4.3.1 Preliminaries

We let random variables be represented by upper-case letters, and their realizations be represented by lower-case letters. Suppose there are n independent subjects sampled from a given population. For subject i, let $A_i \in \{-1, 1\}$ denote a binary treatment assignment and Y_i denote the corresponding outcome. Without loss of generality, we assume a larger value of outcome is desirable. Under the potential outcomes framework (Rubin, 1978; Splawa-Neyman et al., 1990), let $Y_i(-1)$ be the potential outcome had the subject been assigned to control and $Y_i(1)$ be the potential outcome had the subject been assigned to treatment. Let $X_i \in \mathbb{X}$ be the feature vector and, for now, S_i be a single sensitive variable. Extension to multiple sensitive variables is presented in Section 4.3.5. We consider $S \in \mathbb{S}$ where $\mathbb{S} = \{1, \ldots, K\}$ if S is discrete and $\mathbb{S} = \mathbb{R}$ if S is continuous.

We define sensitive variables that are important to the intervention decision, but their inclusion in decision making is prohibited. Formally, consider variables X and S that are both
available during model training and are both determinants of conditional average treatment effect (Rubin, 1974). While X and S may be both involved in training, the derived decision rule $d(\cdot)$ precludes the input of S due to sensitive concerns. Hence, the derived IDR is only a function with the form $d(X) : \mathbb{X} \to \mathbb{A}$. Following the above definition, we consider an offline learning framework where sensitive variables are collected and can be used in obtaining the IDRs, but they cannot be used in the resulting IDRs. A causal diagram and a decision diagram are provided in Appendix C.6.

Assumption 4.3.1 (Consistency).

$$Y = Y(-1)\mathbb{1}(A = -1) + Y(1)\mathbb{1}(A = 1)$$

Assumption 4.3.2 (Positivity).

$$0 < Pr(A = 1|X, S) < 1$$

Assumption 4.3.3 (Unconfoundedness).

$$\{Y(-1), Y(1)\} \perp A | \{X, S\} \text{ and } \{Y(-1), Y(1)\} \not\perp A | X$$

Assumption 4.3.1 is the standard consistency assumption in causal inference and Assumption 4.3.2 states that every subject has a nonzero probability of getting the treatment. Assumption 4.3.3 states that given X and S, the potential outcomes are independent of the treatment assignments. Besides, the unconfoundedness does not hold when only X is conditional, signifying the importance of S. Under causal settings, Assumption 4.3.3 ensures that treatment effects cannot be non-parametrically identifiable without S (Neyman, 1923; Rubin, 1974; Holland, 1986; Imbens and Angrist, 1994; Pearl, 2009). Approaches such as disentanglement of X from S under supervised learning settings mentioned in Section 4.2 will introduce bias towards estimating the IDR.

Before introducing the proposed method, we first discuss two kinds of approaches to deal with sensitive variables under causal settings. The first kind is naive approaches that omit sensitive variables. When S is not available for future deployment, a naive approach is to maximize $E_X\{E(Y|X, A = d(X))\}$ over d using (X, A, Y) during the training procedure. This approach will introduce bias in the estimation of potential outcomes and lead to a suboptimal IDR due to the unmeasured confounder S. It is thus important that one includes S into the training procedure. For example, if we consider the value function framework (i.e., expected outcome) used by most existing works such as Manski (2004); Qian and Murphy (2011), we can show that

$$E\{Y(d)\} = E_{X,S} [E(Y(d)|X,S)] = E_X [E_{S|X} \{E(Y(d)|X,S)\}]$$
(11)
= $E_X [E_{S|X} \{E(Y|X,S,A = d(X))\}] \neq E_X [E(Y|X,A = d(X))],$

where the third equality in (11) holds by Assumptions in Section 4.3.1 and the last inequality also indicates the naive approaches without using S will in general fail. Then one valid approach is the mean-optimal approach that uses the sensitive variables. That is, to maximize $E_X[E_{S|X}{E(Y|X, S, A = d(X))}]$ over d using (X, S, A, Y). The optimal IDR under this criterion is, for every $X \in \mathbb{X}$,

$$\tilde{d}(X) \in \operatorname{sgn}(E_{S|X}\{E(Y|X, S, A=1)\} - E_{S|X}\{E(Y|X, S, A=-1)\}),$$

which guarantees to find the treatment that maximizes the conditional expected outcome given each X by averaging out the effect of the sensitive variable S. The mean-optimal approaches, however, fail to control the disparities across realizations of the sensitive variables due to the integration over S, which may lead to unsatisfactory decisions to certain subgroups, as illustrated in the toy example in Section 4.1.

4.3.2 Robust Optimality with Sensitive Variables

Driven by the limitation of existing approaches, our goal is to derive a robust decision rule that maximizes the worst-case scenarios of subjects when some sensitive information is not available at the time of deploying the decision rule. Specifically, our robust decision learning framework draws decisions based on individuals' available characteristics summarized in the vector X without the sensitive variable S, while improving the worst-case outcome of subjects in terms of the sensitive variable in the population. Formally, given a collection \mathbb{D} of all treatment decision rules depending only on X, the proposed RISE approach estimates the following IDR, which is defined as

$$d^* \in \arg\max_{d \in \mathbb{D}} E_X \big[G_{S|X} \{ E(Y|X, S, A = d(X)) \} \big], \tag{12}$$

where $G_{S|X}(\cdot)$ could be chosen as some risk measure for evaluating E(Y|X, S, A = d(X))for each $S \in \mathbb{S}$. Examples include variance, conditional value at risk, quantiles, etc. In this paper, we consider $G_{S|X}$ as the conditional quantiles (for a continuous S) or the infimum (for a discrete S) over \mathbb{S} .

Specifically, for a discrete S, $G_{S|X}$ is consider as an infimum operator of E(Y|X, S, A = d(X)) over S. We thus aim to find

$$d^* \in \arg\max_{d \in \mathbb{D}} E_X \big[\inf_{s \in \mathbb{S}} \{ E(Y|X, S = s, A = d(X)) \} \big],$$

where inf is the infimum taken with respect to E(Y|X, s, A = d(X)) over $s \in S$. This implies that for a given X, $d^*(X)$ assigns the treatment that yields the best worst-case scenario among all possible values of S for every $X \in X$, or equivalently,

$$d^*(X) \in \mathrm{sgn}(\inf_{s \in \mathbb{S}} \{ E(Y|X, S = s, A = 1) - \inf_{s \in \mathbb{S}} \{ E(Y|X, S = s, A = -1) \}).$$

For a continuous S, we consider $G_{S|X}\{E(Y|X, S, A = d(X))\}$ as $Q_{S|X}^{\tau}\{E(Y|X, S, A = d(X))\}$, which is the τ -th quantile of $\{E(Y|X, S, A = d(X))\}$ and $\tau \in (0, 1)$ is the quantile level of interest. Specifically, $Q_{S|X}^{\tau}\{E(Y|X, S, A = d(X))\} = \inf\{t : F(t) \geq \tau\}$ with F denoting the conditional distribution function of E(Y|X, S, A = d(X)) given X and d. Note the randomness behind E(Y|X, S, A = d(X)) given X and d is fully determined by the sensitive variable S. Then optimal IDR under this criterion is defined as

$$d^* \in \arg\max_{\mathbb{D}} E_X \left[Q_{S|X}^{\tau} \{ E(Y|X, S, A = d(X)) \} \right].$$

This implies that for a given X, $d^*(X)$ assigns a treatment that yields the largest τ -th quantile of the outcome over the distribution related to S, or equivalently,

$$d^*(X) \in \mathrm{sgn}(\{Q^\tau_{S|X}\{E(Y|X,S,A=1)\} - Q^\tau_{S|X}\{E(Y|X,S,A=-1)\})$$

We let $\tau = 0.25$ throughout the paper and suppress τ for simplicity. Results on varying the value of τ is provided in Appendix; see Section 4.4.1 for details.

4.3.3 Identifying Vulnerable Subjects

Our RISE framework provides a natural way to define vulnerable groups. Specifically, for a discrete S, if $\inf_{S} \{ E(Y|X, S, A = 1) \} > \inf_{S} \{ E(Y|X, S, A = 0) \}$, then $\arg \inf_{S} \{ E(Y|X, S, A = 0) \}$ is vulnerable given X, otherwise is $\arg \inf_{S} \{ E(Y|X, S, A = 1) \}$. In other words, the vulnerable subjects are those in the worst-off group that needs protection. Similarly, for a continuous S, if $Q_{S} \{ E(Y|X, S, A = 1) \} > Q_{S} \{ E(Y|X, S, A = 0) \}$, then the set $\{ S : E(Y|X, S, A = 0) \le Q_{S} \{ E(Y|X, S, A = 0) \} \}$ defines the vulnerable subjects given X, otherwise this group is defined as $\{ S : E(Y|X, S, A = 1) \le Q_{S} \{ E(Y|X, S, A = 1) \} \}$.

4.3.4 Estimation and Algorithm

Here we provide a transformation of the proposed RISE from an optimization problem to a weighted classification problem. There are several advantages to this conversion: 1) The optimization problem defined in (12) involves a nonsmooth and nonconvex objective function that could lead to computational challenges. 2) With multiple powerful statistical and machine learning toolbox to choose from, a classification problem can be more readily solved in practice. Hyperparameter tuning and model selection could be conducted to further boost performance. 3) Compared to a direct optimization of (12), a classification-based optimizer allows the use of off-the-shelf software packages that can be tailored to different functional classes or incorporate different properties such as model sparsity.

Proposition 4.3.4. Maximizing the objective function in (12) is equivalent to maximizing

$$E_X \{ \mathbb{1}(d(X) = 1) [G_{S|X} \{ E(Y|X, S, A = 1) \} - G_{S|X} \{ E(Y|X, S, A = -1) \}] \}$$

With Proposition 4.3.4 and a proper estimator of the outcome model E(Y|X, S, A) using our training data $\mathcal{D}_n = \{X_i, S_i, A_i, Y_i\}_{i=1}^n$, we replace the expectation of Y_i by its estimate \hat{Y}_i and solve the following problem.

$$\arg\max_{d\in\mathbb{D}} n^{-1} \sum_{i=1}^{n} [\mathbb{1}(d(x_i) = 1) \{g_1(x_i) - g_2(x_i)\}],$$
(13)

where $g_1(x_i) = G_{s|x}\{\hat{Y}_i(x_i, s, a_i = 1)\}$ and $g_2(x_i) = G_{s|x}\{\hat{Y}_i(x_i, s, a_i = -1)\}$. We have the following proposition to address noncontinuity in (13) and transform it into a classification problem. Define \mathbb{F} as a class of all measurable functions over \mathbb{X} .

Proposition 4.3.5. Let f(x) to be a smooth function. Maximizing the empirical objective in (13) is equivalent to a weighted classification problem of minimizing over $f \in \mathbb{F}$,

$$n^{-1} \sum_{i=1}^{n} \mathbb{1}[\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)|,$$
(14)

with features x_i , the true label sgn $\{g_1(x_i) - g_2(x_i)\}$, and the sample weight $|g_1(x_i) - g_2(x_i)|$, for subject i, i = 1, ..., n.

With Proposition 4.3.5, we have transformed the optimization problem (12) into a weighted classification problem (14) where for subject *i* with features x_i , the true label is $\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\}$ and the sample weight is $|g_1(x_i) - g_2(x_i)|$. The estimated optimal decision rule by (14) is then given by $\hat{d}(x) = \operatorname{sgn}\{\hat{f}(x)\}$.² The proof of Proposition 4.3.4 and Proposition 4.3.5 is presented in Appendix C.2.

Algorithm 2 provides an algorithmic overview. The inner expectation E(Y|X, S, A) can be modeled as $\hat{Y}(X, S, A)$ using a twin model separated by the treatment and control groups. For a continuous S, we propose to estimate $G(X, A) = Q_{S|X,A}\{E(Y|X, S, A)\}$ via a quantile regression of \hat{Y} on X but without S. For a discrete S, we propose to obtain an estimate of $G(X, A) = \inf_{S} \{E(Y|X, S, A)\}$ by finding the minimum among $\{E(Y|X, S = 1, A), \ldots, E(Y|X, S = K, A)\}$. The estimated decision rule can then be obtained from the weighted classification. In our implementation, neural networks are used to fit models in the training data sets. The details on modeling and hyperparameter tuning via cross-validations are given in Appendix C.3. A Python package **rise** based on neural networks is built. Note that the model choices are flexible.

²If $\hat{f}(x) = 0$, assign a random treatment.

Algorithm 2 RISE (Robust individualized decision learning with sensitive variables)

Input Training data $\mathcal{D}_n = \{Y_i, A_i, X_i, S_i\}_{i=1}^n$

Output Estimated decision rule \hat{d}

- 1: for i = 1 to n do
- 2: $\hat{Y}_i(x_i, s_i, a_i) \leftarrow \text{Model } E(Y|X, S, A = a) \text{ using } \mathcal{D}_n \text{ with } a = 1 \text{ and } a = -1, \text{ respectively.}$
- 3: **if** S is continuous **then**

4:
$$g_1(x_i) \leftarrow \text{Model } Q_{S|X,A}\{E(Y|X, S, A = a)\}$$
 via quantile regressions of $\hat{Y}_i(x_i, s_i, a_i)$
on x_i , for \mathcal{D}_n with $a = 1$.

5:
$$g_2(x_i) \leftarrow \text{Model } Q_{S|X,A}\{E(Y|X, S, A = a)\}$$
 via quantile regressions of $\tilde{Y}_i(x_i, s_i, a_i)$
on x_i , for \mathcal{D}_n with $a = -1$.

6: end if

7: **if** S is discrete **then**

8:
$$g_1(x_i) \leftarrow \text{Compute inf}_{s \in \mathbb{S}} \{ Y_i(x_i, s, a_i = 1) \}.$$

$$g_2(x_i) \leftarrow \text{Compute inf}_{s \in \mathbb{S}} \{ \hat{Y}_i(x_i, s, a_i = -1) \}.$$

10: end if

9

11: end for

12: $\hat{d} \leftarrow$ Build a weighted classification model with features x_i , label $\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\}$, and sample weight $|g_1(x_i) - g_2(x_i)|$ for $1 \le i \le n$.

13: Return d

4.3.5 Extension to Multiple Discrete Sensitive Variables

For multiple discrete sensitive variables, similar estimation procedure can be conducted as outlined in Section 4.3.4. Suppose there are L discrete sensitive variables, i.e., $S = \{S_1, S_2, \ldots, S_L\}$. The inner expectation $E(Y|X, S_1, \ldots, S_L, A)$ can be obtained with a twin model of Y on X and all S for each treatment level. The infimum over S is obtained by finding the minimum iterating space of possible parameter values for each sensitive variable. See Section 4.4.2 for an example of using multiple discrete sensitive variables. We will discuss in Section 4.5 the challenges and future work related to scenarios with multiple continuous sensitive variables or that with a mixture of continuous and discrete sensitive variables.

4.4 Numerical Studies

In this section, we perform extensive numerical experiments to investigate the merit of robustness of the proposed framework via simulations and three real-data applications. The results demonstrate that the proposed rules achieve a robust objective with sensitive variables unavailable at the time of decision while maintaining comparable mean outcomes.

For comparison, we consider the naive and mean-optimal approaches described in Section 4.3.1, which correspond to different choices of $G(\cdot)$ functions. The naive decision rule that simply disregard information of S, denoted as **Base**, can be formulated in our optimization framework of (12) by letting G(X, A) = E(Y|X, A). The IDR can be estimated directly by fitting a model of Y on X in each treatment arm. The resulting IDR is not sensitive variables-aware and is biased due to confounding, as discussed. Another IDR that resembles traditional mean-optimal decision rules, denoted as **Exp**, can be formulated as G(X, S, A) = E(Y|X, S, A). This can be obtained by training a classification model without S, i.e., only using X, after obtaining an outcome model for the inner expectation E(Y|X, S, A). Note that this approach is not robust to extreme behaviors in S. The modeling approaches described in Appendix C.3 applies to here. We also include the *double robust* (Chernozhukov et al., 2018) versions of Base and Exp, respectively, by adapting Policytree (PT, Sverdrup et al., 2020; Athey and Wager, 2021), the latest state-of-the-art policy learning method for maximizing the expected values. The two new methods are termed **PT-Base** and **PT-Exp**.

We consider the following evaluation metrics. 1) Objective: the quantile objective is estimated and reported for a continuous S and the infimum objective is for a discrete S. The objective, when $\tau < 0.5$, (here $\tau = 0.25$) represents the value of the "low performers" among all possible value of S under a given d. 2) Value: the value function, or expected reward used by the most existing methods, such as Manski (2004); Qian and Murphy (2011), is defined as $V(d) = E\{Y(d)\}$. It represents the "average performers". For randomized trials, an unbiased estimator of V(d) is given by $\hat{V}(d) = \{\sum_{i=1}^{T} Y_i \mathbb{1}(A_i = \hat{d}) / \pi(A_i, \mathbf{X}_i)\} / \{\sum_{i=1}^{T} \mathbb{1}(A_i = \hat{d}) / \pi(A_i, \mathbf{X}_i)\}$ (Murphy et al., 2001), where T is the sample size of the test data. For observational studies, the value is estimated with $\hat{V}(d) = T^{-1} \sum_{i=1}^{T} \hat{Y}_i(x_i, s_i, a_i = \hat{d})$. We report the metrics among all subjects and among the potential vulnerable subgroup, respectively. For simulation, we consider training data and the testing data, respectively, with sample sizes of 5,000. For real-data applications, we consider a 80-20 split of the dataset into a training data and a testing data. Continuous covariates are standardized before the estimation. All results are based on 100 replications.

4.4.1 Simulation Studies

Example 1. Here we provide the detail for the simulation of the motivating example introduced in Section 4.1. The outcome is generated using the following model: $Y_i = \mathbb{1}(X_i > 0.5)\{5+10\mathbb{1}(A_i=1)+22S_i-24\mathbb{1}(A_i=1)S_i\}+\mathbb{1}(X_i \leq 0.5)\{11+19\mathbb{1}(A_i=1)+2S_i-32\mathbb{1}(A_i=1)S_i\}+\epsilon_i$, where the covariate $X_i \sim Unif[0,1]$, treatment assignment $A_i \sim Bernoulli(0.5)$, and the noise $\epsilon_i \sim N(0,1)$. For a discrete type S, the sensitive variable $S_i \sim Bernoulli(0.5)$. For a continuous type S, S_i is generated from a mixture of beta distributions, Beta(4,1) and Beta(1,4), with equal mixing proportions.

Example 2. We generate the outcome Y using the following model: $Y_i = \{0.5 + \mathbb{1}(A_i = 1) + \exp(S_i) - 2.5S_i\mathbb{1}(A_i = 1)\}\{1 + X_{i1} - X_{i2} + X_{i3}^2 + \exp(X_{i4})\} + \{1 + 2\mathbb{1}(A_i = 1) + 0.2\exp(S_i) - 3.5S_i\mathbb{1}(A_i = 1)\}\{1 + 5X_{i1} - 2X_{i2} + 3X_{i3} + 2\exp(X_{i4})\} + \epsilon$, where $X_{ij} \sim U(0, 1)$, $j = 1, \ldots, 6, A$ satisfies $\log\{P(A_i = 1|X_i)/P(A = 0|X_i)\} = 0.6(-S_i + X_{i1} - X_{i2} + X_{i3} - X_{i4} + X_{i5} - X_{i6})$, and $\epsilon_i \sim N(0, 1)$. For a continuous type S, S_i is generated from a mixture of beta distributions, Beta(4, 1) and Beta(1, 4), with equal mixing proportions; for a discrete type S, we consider a binary S_i that satisfies $\log\{P(S_i = 1|X_i)/P(S_i = 0|X_i)\} = -2.5 + 0.8(X_{i1} + X_{i2} + X_{i3} + X_{i4} + X_{i5} + X_{i6})$.

Table 6 summarizes the performance of the proposed IDRs compared to the mean criterion for Example 1 and Example 2. The proposed RISE achieves the largest objectives and improves the value among vulnerable subjects, while maintaining comparative overall values. As demonstrated in the toy example introduced in Section 4.1, we expect that RISE helps improve the value among the vulnerable subjects while maintaining a comparable overall value. As for the objective, intuitively, the proposed rule is expected to have a larger objective. We also point out that there is no direct relationship between the objective among all subjects versus the objective among vulnerable subjects. For example, using the toy example with setup in Table 4, and limiting to subjects with $X \leq 0.5$ only, S = 1 is vulnerable and is assigned A = -1 by the proposed RISE. The objective among S = 1 is 13 but the objective among both S = 0 and S = 1 is 12 = (11 + 13)/2, which is smaller than that among the vulnerable group. In other words, by protecting the vulnerable subjects, the proposed rule may lead to an increase in the outcome of the vulnerable group, and the gain may result in a higher outcome than the overall mean outcome. PT-Exp method tends to show the best improvement in terms of the overall value, as the doubly robust-based estimators tend to reduce variance in value estimation. However, PT-Exp is shown to have minimal benefits for vulnerable subjects. RISE still shows the largest gain in the objective among vulnerable subjects among all compared methods.

In the appendix, we consider for a continuous S different quantile criteria $\tau = 0.1$ and 0.5, respectively, to test the robustness of the proposed RISE. Results show that when τ is small, there is more strength in the proposed method, as the algorithm aims to improve the worst-outcome scenarios. The proposed RISE has the largest gain in objective and value among vulnerable subjects when τ is 0.1, and has similar performance as the compared approaches when τ is 0.5. Also, we consider a scenario where S is not involved in the data generation of Y, i.e., Assumption 4.3.3 is simplified as $\{Y(-1), Y(1)\} \perp A | X$. The estimated objective and value function are similar across all compared approaches, which indicates the robustness of RISE. Finally, we study the performances of our method when Assumption 4.3.2 is nearly violated or Assumption 4.3.3 is violated. Similar patterns have been observed that the proposed RISE achieves the largest objectives and improves the value among vulnerable subjects, while maintaining comparative overall values. The details can be found in Appendix C.4.

Example	Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
1	Disc.	Base	7.03(0.03)	$7.01 \ (0.04)$	14.3 (0.05)	7.92(0.06)
		Exp	6.39(0.03)	6.39(0.04)	$14.4 \ (0.05)$	7.14(0.06)
		RISE	12.0 (0.01)	12.0(0.01)	$13.0\ (0.01)$	14.0 (0.01)
		PT-Base	2.66(0.02)	2.65(0.02)	$15.4 \ (0.05)$	2.58(0.02)
		PT-Exp	2.62(0.02)	2.62(0.02)	15.5 (0.05)	2.55 (0.02)
	Cont.	Base	9.12 (0.03)	9.14(0.04)	14.5 (0.08)	8.25 (0.11)
		Exp	8.75(0.03)	8.75(0.04)	14.6 (0.08)	7.58(0.06)
		RISE	10.3 (0.02)	10.3 (0.03)	13.9(0.04)	10.3 (0.06)
		PT-Base	6.71(0.03)	6.72(0.03)	$15.3 \ (0.05)$	4.52 (0.02)
		PT-Exp	6.68(0.02)	6.67(0.02)	15.4 (0.05)	4.47(0.02)
2	Disc.	Base	7.79(0.02)	8.66(0.03)	19.4 (0.04)	11.4 (0.06)
		Exp	9.12(0.03)	$10.1 \ (0.03)$	19.5(0.04)	$14.4 \ (0.05)$
		RISE	13.5(0.01)	14.0 (0.01)	17.4(0.02)	22.1 (0.02)
		PT-Base	7.19(0.03)	$7.77 \ (0.03)$	$19.0\ (0.05)$	$9.71 \ (0.05)$
		PT-Exp	8.30 (0.02)	9.03(0.03)	$19.1 \ (0.04)$	$12.2 \ (0.05)$
	Cont.	Base	9.89(0.02)	9.87(0.03)	17.6(0.02)	9.09 (0.04)
		Exp	11.1 (0.02)	11.1 (0.02)	17.8(0.02)	12.2 (0.04)
		RISE	13.7 (0.02)	13.7 (0.02)	$17.0\ (0.01)$	18.9(0.03)
		PT-Base	9.30(0.02)	9.29(0.03)	$18.0\ (0.03)$	7.61 (0.04)
		PT-Exp	9.41(0.02)	$9.41 \ (0.02)$	18.1 (0.02)	7.92(0.04)

Table 6: Simulation results for Example 1 and Example 2. Standard error in parenthesis.

4.4.2 Real-data Applications

We present three real-data examples to showcase the robust performance of RISE. These applications consider either fairness or safety in the context of policy (LaLonde, 1986) and healthcare (Hammer et al., 1996; Seymour et al., 2016) where sensitive variables exist. The information of data access is provided in Appendix C.5.

4.4.2.1 Fairness in a Job Training Program

To illustrate the implication of the proposed method from a fairness perspective, we consider the National Supported Work (NSW) program (LaLonde, 1986) for improving personalized recommendations of a job training program on increasing incomes. This program intended to provide a 6 to 18-month training for individuals in face of economic and social problems such as former drug addicts and juvenile delinquents. The original experimental dataset consists of 185 individuals who received the job training program (A = 1) and 260 individuals who did not (A = -1). The baseline covariates are age, years of schooling, race (1 = African Americans or Hispanics, 0 = others), married (1 = yes, 0 = no), high school diploma (1 = yes, 0 = no), earning in 1974, and earning in 1975. The outcome variable is the earning in 1978. In the exploratory analysis using causal forest (Wager and Athey, 2018), we observe that age may play an important role in the causal effect of the job training program on the long-term post-market earning. In the following data example we use age as the sensitive variable S and other baseline covariates as X. The earnings in years 1974, 1975, and 1978 are transformed by taking the logarithm of the earning plus one.

4.4.2.2 Improvement of HIV Treatment

To illustrate the implication of the proposed method from a safety perspective when there is delayed information, we consider the ACTG175 dataset among HIV positive patients (Hammer et al., 1996). The original study considers a total of 2,139 patients who were randomly assigned into four treatment groups. In this data application, we focus on finding the optimal IDRs between two treatments: zidovudine combined with didanosine (A = -1) and zidovudine combined with zalcitabine (A = 1). The total number of patients receiving these two treatments is 1,046. The baseline covariates we consider are age, weight, CD4 Tcell amount at baseline, hemophilia (1 = yes, 0 = no), homosexual activity (1 = yes, 0 = no), Karnofsky score, history of intravenous drug use (1 = yes, 0 = no), gender (1 = male, 0 =female), CD8 T-cell amount at baseline, race (1 = non-Caucasian, 0 = Caucasian), number of days of previously received antiretroviral therapy, use of zidovudine in the 30 days prior to treatment initiation (1 = yes, 0 = no), and symptomatic indicator (1 = symptomatic, 0 = asymptomatic). The outcome variable is the CD4 T-cell amount at 96 ± 5 weeks from the baseline. We consider CD8 T-cell amount at baseline as the sensitive variable. The response of CD8 T-cell among HIV positive patients has not been fully understood (Boppana and Goepfert, 2018). Clinically, it is plausible that only CD4 is measured in clinical visits where treatments are based on, hence CD8 might not be measured and not used in decision making. As our exploratory analysis using causal forest shows, CD8 T-cell amount may play an important part in the treatment effect of the outcome.

4.4.2.3 Safe Resuscitation for Patients with Sepsis

For this application, we apply the proposed method to treating sepsis, a life-threatening disease. This application intends to provide an example to apply our method with multiple categorical sensitive variables in the scenario where there is missing yet important information at the time of decision making. We apply the proposed method to a sepsis study from the University of Pittsburgh Medical Center (UPMC). The original study cohort includes 30,687 patients with Sepsis-3 (Seymour et al., 2016) within 6 hours of hospital arrival from 14 UPMC hospitals between 2013 and 2017. For our data analysis, we consider X to be baseline patient characteristics 4 hours before sepsis onset, which includes patient demographics of age, gender (1 = male, 0 = female), race (1 = Caucasian, 0 = others), and weight, and vital signs of usage of mechanical ventilation (1 = yes, 0 = no), respiratory rate, temperature, intravenous fluids (1 = yes, 0 = no), Glasgow Coma Scale score, platelets, blood urea nitrogen, white blood cell counts, glucose, creatinine. We consider two sensitive variables, lactate and Sequential Organ Failure Assessment (SOFA) score 4 hours before sepsis onset.

Lactate and SOFA score have been two important indicators of sepsis severity (Howell et al., 2007; Krishna et al., 2009; Shankar-Hari et al., 2016). Different from the baseline patient demographics or common vital signs that are typically obtained at the admission of patients, SOFA score combines performance of several organ systems in the body (Seymour et al., 2016), which requires additional calculation and cannot be obtained directly. Lactate labs measures the level of lactic acid in the blood (Andersen et al., 2013) and are less common in routine examination, which could be delayed in ordering. Hence, their measurements are obtained retrospectively after treatment decisions have been made and are not available at times of decision. We discretize lactate level at clinically meaningful value of 2 mmol/L (Shankar-Hari et al., 2016), and SOFA score at value of 6 for analysis (Vincent et al., 1996; Ferreira et al., 2001). The treatment option is whether the patient took any vasopressors during the first 24 hours after sepsis onset. The outcome is patient survival at day 90. The analysis cohort contains 6,539 patients in total. We are interested in making decision about whether to treat patients with vasopressors in the first 24 hours after sepsis onset given the measurements of lactate and SOFA are not available at the time of decision making. Additional rationale and background on this example is provided in Appendix C.5.

4.4.2.4 Results of Real-data Applications

Table 7 presents the performance of various IDRs on the three applications. As expected, RISE has the largest objective as well as value among vulnerable subjects. The patterns are similar to that in the synthetic experiments in Section 4.4.1. In applications to the job training data and the sepsis study, results show that RISE has a larger value among all subjects than other IDRs. This is possible when there are more gains in the vulnerable subjects than other subjects, which further demonstrate the superiority of the proposed approach in improving worst-case outcomes caused by sensitive variables. We provide visualizations by Shapley additive explanations (SHAP) (Lundberg and Lee, 2017) for RISE and Exp, respectively, in Appendix C.5 about feature importance in the final classification models to help interpret important covariates in making the decisions. The SHAP approach provides united values to describe the correlation between each feature and the predicted decision rule, respectively (Lundberg and Lee, 2017). Overall, the direction of correlations is similar for RISE and Exp, but their ranking of feature importance may be different.

Table 7: Estimated objective and value of different IDRs for the three data applications. Standard error in parenthesis. The outcome of each study is italicized.

Dataset	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
	Base	5.26(0.04)	5.28(0.05)	6.32(0.05)	6.33(0.07)
NICIN	Exp	5.22(0.04)	5.24(0.05)	$6.37 \ (0.05)$	6.37 (0.07)
	RISE	5.43 (0.04)	5.44 (0.04)	6.42(0.04)	6.42 (0.06)
log(income+1)	PT-Base	4.97(0.04)	5.08(0.06)	6.40(0.03)	6.38(0.05)
	PT-Exp	5.03(0.04)	$5.11 \ (0.05)$	6.43 (0.03)	6.40(0.05)
	Base	336.9(1.65)	338.1 (2.23)	350.5(1.86)	357.5(2.24)
	Exp	337.5(1.65)	338.9(1.80)	351.9(1.95)	359.1 (2.21)
ACIGI75	RISE	351.5 (1.67)	351.2 (1.80)	351.8(1.88)	363.1 (2.19)
CD4 T-cell amount	PT-Base	299.7(1.01)	299.5(1.91)	356.9(1.72)	350.7(2.54)
	PT-Exp	$300.1\ (0.99)$	299.9(1.83)	357.1 (1.55)	352.7(2.61)
	Base	0.752(0.001)	0.721 (0.001)	0.965(0.001)	0.905 (0.002)
C	Exp	$0.752 \ (0.001)$	$0.721 \ (0.002)$	$0.966\ (0.001)$	$0.908 \ (0.002)$
Sepsis	RISE	$0.771 \ (0.001)$	0.735 (0.001)	0.972(0.001)	0.923 (0.002)
survival rate	PT-Base	$0.753\ (0.001)$	0.642(0.002)	$0.981 \ (0.001)$	$0.848\ (0.003)$
	PT-Exp	$0.764\ (0.001)$	$0.644\ (0.002)$	0.984 (0.001)	$0.875\ (0.003)$

4.5 Discussion and Future Work

We have proposed RISE, a robust decision learning framework with a novel quantile-, or infimum-optimal treatment objective intended to improve the worst-case scenarios of individuals when decisions with uncertainty are needed to be made with sensitive yet important information missing. Our approach can be applied to a board of applications, including but not limited to politics, education, healthcare, etc. For multiple continuous sensitive variables, the estimated decision rule by RISE can be easily obtained by fitting similar quantile regression as described in Section 4.3.4. For a mixture of continuous and discrete sensitive variables, the estimated rule can be obtained by taking infimum over the discrete ones, then obtaining quantile over the continuous ones. However, challenges remain in finding vulnerable subjects described in Section 4.3.2 under these settings as it may be computationally difficult to find a vulnerable set of S when they are multi-dimensional. Another future work includes extension of the current binary treatment option to a multi-treatment option. It is also worth mentioning that our work can be naturally extended to the scenario where there exist unmeasured confounders. As long as the conditional average outcome given observed covariates can be identified (e.g., via instrumental variables (Wang and Tchetgen Tchetgen, 2018)), our method can be applied.

5.0 Summary and Future Work

This dissertation is motivated by challenges in causal inference under practical data restrictions. Traditional randomized clinical trials typically require long years of data collection, which leads to loss of follow-up, poor compliance, or other issues and subsequently affects the estimation of treatment effects. In the more recent neoadjuvant trials, on the other hand, the efficacy of a treatment can be estimated early with an intermediate posttreatment response. However, the clinical implication of this intermediate post-treatment response has not yet been understood. This idea, along with real data from a neoadjuvant clinical trial, has motivated the development of methods in Chapter 2. In the modern context, new challenges arise with growing concerns such as data privacy and operational feasibility in distributed research networks, and the timeliness and fairness of individualized decision rules. Driven by these concerns, we develop the privacy-protecting method for improving the estimation of conditional average treatment effects in Chapter 3 and the fairness-aware decision learning framework in Chapter 4 with board applications, including but not limited to politics, education, healthcare, etc. The three proposed methods in this dissertation respectively address important challenges in causal inference, which includes: identification of principal stratum treatment effects, enhancement of treatment effect estimation via heterogeneous data integration, and derivation of robust individualized decision rules. Each of the methods can be further improved and extended along in their own framework and settings as have been discussed in each of the chapters.

One promising direction for future research is to generalize the toolbox of privacyprotecting analytic and data-sharing methods and to construct a unified framework that is applicable to a broader range of problems pertaining to modern distributed data networks. It is of great interest to borrow the strengths of each of the proposed methods, Chapter 3 in particular, pairing with new methodologies developed by others to develop, test, and distribute open-source statistical software packages, maximizing value of methodology research to practical applications.

Driven by the recent initiatives of collaboratory distributed research networks, another

promising direction and natural extension of current work, Chapter 4 in particular, is to develop a generalizable recommendation system for treatment that is robust to population heterogeneity across multiple sites. A more robust treatment recommendation system that jointly take into account the population heterogeneity due to observed and/or unobserved confounding can enhance treatment gain from across all sites, hence more general and widely applicable.

Last but not least, it is of great interest to investigate the unconfoundedness assumption that is typically adopted in causal inference for observational studies, as mentioned in Chapter 3 and Chapter 4. Unconfoundedness is a strong and untestable assumption for observational studies where unmeasured confounding could exist. It is therefore important to stress how this assumption breaks down when there is unobserved confounding. For example, little has been understood about how bad things could get and when do things cancel out in cases of assumption violations. We would like to address these important issues by leveraging the knowledge accumulated through the development of this dissertation.

Appendix A for Chapter 2

A.1 Estimation of
$$\Pr\{S_i(0) = 1 | X_i = x\}$$
 and $\Pr\{S_i(1) = 1 | X_i = x\}$

We use the maximum likelihood approach to estimate $\Pr\{S_i(0) = 0 | X_i = x\}$, $\Pr\{S_i(0) = 1 | X_i = x\}$ and $\Pr\{S_i(1) = 1 | X_i = x\}$. Let

$$E_{jkx} = \{i : S_i(0) = j, S_i(1) = k | X_i = x\}, \quad j, k = 0, 1, x \in \Gamma$$

be the principal stratum under each category X = x. Because of the monotonicity assumption, E_{10x} is empty. Let

$$p_{jkx} = \Pr\{E_{jkx}\} = \Pr\{S_i(0) = j, S_i(1) = k | X_i = x\}, \ j, k = 0, 1, x \in \Gamma$$

Therefore, $p_{00x} + p_{01x} + p_{11x} = 1$ for all $x \in \Gamma$. For each x, $\Pr\{E_{jkx}\}$ can be estimated from the observed data $\{Z_i, X_i, S_i(Z_i), i = 1, 2, ..., n\}$ via maximum likelihood. Let N_{zsx} be the total number of subjects with Z = z, S(Z) = s and baseline category x with $\sum_{Z;S=0,1;X} N_{zsx} = n$. Then the likelihood function for $(p_{00x}, p_{01x}, p_{11x})$ is given by

$$L(p_{00x}, p_{01x}, p_{11x}|N_{00x}, N_{01x}, N_{10x}, N_{11x}) \propto f(N_{zsx})$$

$$\propto \Pr\{S(0) = 0|X = x\}^{N_{00x}} \cdot \Pr\{S(0) = 1|X = x\}^{N_{01x}}$$

$$\cdot \Pr\{S(1) = 0|X = x\}^{N_{10x}} \cdot \Pr\{S(1) = 1|X = x\}^{N_{11x}}$$

$$= (p_{00x} + p_{01x})^{N_{00x}} \cdot p_{11x}^{N_{01x}} \cdot p_{00x}^{N_{10x}} \cdot (p_{01x} + p_{11x})^{N_{11x}} \text{ (by monotonicity assumption)}$$

$$= (1 - p_{11x})^{N_{00x}} \cdot p_{11x}^{N_{01x}} \cdot p_{00x}^{N_{10x}} \cdot (1 - p_{00x})^{N_{11x}}$$

$$= (1 - p_{11x})^{N_{00x}} \cdot p_{11x}^{N_{01x}} \cdot (1 - p_{+1x})^{N_{10x}} \cdot p_{+1x}^{N_{11x}}$$

1) When $N_{00x} \cdot N_{11x} \ge N_{01x} \cdot N_{10x}$, the resulting MLEs for $(p_{00x}, p_{01x}, p_{11x})$ are given by

$$\widehat{p}_{00x} = \widehat{\Pr}\{S_i(0) = 0, S_i(1) = 0 | X_i = x\} = 1 - \widehat{p}_{+1x}$$
$$= \frac{N_{10x}}{N_{10x} + N_{11x}} = \frac{\sum_i \mathbb{1}(Z_i = 1, S_i(1) = 0, X_i = x)}{\sum_i \mathbb{1}(Z_i = 1, X_i = x)}$$

$$\widehat{p}_{11x} = \widehat{Pr}\{S_i(0) = 1, S_i(1) = 1 | X_i = x\}$$

$$= \frac{N_{01x}}{N_{00x} + N_{01x}} = \frac{\sum_i \mathbb{1}(Z_i = 0, S_i(0) = 1, X_i = x)}{\sum_i \mathbb{1}(Z_i = 0, X_i = x)}$$

$$\widehat{p}_{01x} = \widehat{Pr}\{S_i(0) = 0, S_i(1) = 1 | X_i = x\} = 1 - \widehat{p}_{00x} - \widehat{p}_{11x}$$

Obviously for each $x \in \Gamma$, \hat{p}_{00x} is the proportion of non-respondents in the treatment arm with X = x; \hat{p}_{11x} is the proportion of respondents in the control arm with X = x.

2) When $N_{00x} \cdot N_{11x} < N_{01x} \cdot N_{10x}$, $\hat{p}_{11x} = \hat{p}_{+1x}$. The likelihood function is given by

$$L(p_{00x}, p_{01x}, p_{11x} | N_{00x}, N_{01x}, N_{10x}, N_{11x})$$

= $(1 - p_{11x})^{N_{00x}} \cdot p_{11x}^{N_{01x}} \cdot (1 - p_{11x})^{N_{10x}} \cdot p_{11x}^{N_{11x}}$
= $(1 - p_{11x})^{N_{00x} + N_{10x}} \cdot p_{11x}^{N_{01x} + N_{11x}}$

The resulting MLEs for $(p_{00x}, p_{01x}, p_{11x})$ are given by

$$\begin{aligned} \widehat{p}_{01x} &= \widehat{Pr}\{S_i(0) = 0, S_i(1) = 1 | X_i = x\} = 0\\ \widehat{p}_{00x} &= \widehat{Pr}\{S_i(0) = 0, S_i(1) = 0 | X_i = x\} = \frac{N_{+0x}}{N_{++x}} = \frac{\sum_i \mathbb{1}(S_i = 0, X_i = x)}{\sum_i \mathbb{1}(X_i = x)}\\ \widehat{p}_{11x} &= \widehat{Pr}\{S_i(0) = 1, S_i(1) = 1 | X_i = x\} = \frac{N_{+1x}}{N_{++x}} = \frac{\sum_i \mathbb{1}(S_i = 1, X_i = x)}{\sum_i \mathbb{1}(X_i = x)} \end{aligned}$$

Then \hat{p}_{00x} is the proportion of non-respondents among all subjects with X = x; \hat{p}_{11x} is the proportion of respondents among all subjects with X = x.

A.2 Proofs of Consistency of Model Parameters and Causal Estimands

Here we show our estimator $\hat{\beta}$ is a consistent estimator for β . We first show that $\hat{\beta}$ can be considered as an extremum estimator as defined by Hayashi (2000). Then we prove that the conditions set forth by Hayashi (2000) for consistency of an extremum estimator are satisfied by our estimator. Then by Slutsky's theorem, the causal estimand $\hat{\theta}$ is a consistent estimator for θ .

Definition A.2.1 (Extremum Estimator). An estimator $\hat{\eta}$ is an extremum estimator if there is a function $Q_n(\eta)$ such that (Hayashi, 2000)

$$\widehat{\eta} = \arg\max_{\eta} Q_n(\eta); \quad \eta \in H.$$

One example of an extremum estimator is the maximum likelihood estimator where

$$Q_n(\eta) = \prod_{i=1}^n f(x_i|\eta).$$

Here we minimize the objective function,

$$Q_{n}(\boldsymbol{\beta}) = \sum_{x=0}^{K} Q_{n}^{(x)}(\boldsymbol{\beta})$$

= $\sum_{x=0}^{K} \{ \widehat{G}_{L}(x) - \sum_{y=0}^{1} G_{M}(x, y; \boldsymbol{\beta}) \cdot \widehat{G}_{R}(x, y) \}^{2}; x \in \Gamma$

which is equivalent to maximizing $-Q_n(\boldsymbol{\beta})$. Therefore $\hat{\boldsymbol{\beta}}$ is an extremum estimator.

Let

$$Q_0(\boldsymbol{\beta}) = \sum_{x=0}^{K} Q_0^{(x)}(\boldsymbol{\beta}); \ x \in \Gamma$$

where $Q_0^{(x)}(\boldsymbol{\beta}) = \{G_L(x) - \sum_{y=0}^{1} G_M(x,y;\boldsymbol{\beta}) \cdot G_R(x,y)\}^2$. We present sufficient conditions for the existence of a unique local minimizer of $Q_0(\boldsymbol{\beta})$ in Lemma A.2.2.

Lemma A.2.2. There exists a unique local minimizer $\boldsymbol{\beta}_0$ for $Q_0(\boldsymbol{\beta})$ if: (a) $Q_0^{(x)}(\boldsymbol{\beta}_0) = 0, \forall x \in \Gamma = \{0, 1, 2, \dots, K\}.$

(b)
$$\operatorname{rank} \left| \frac{\partial \tilde{Q}_0(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} \right|_{\boldsymbol{\beta} = \boldsymbol{\beta}_0} \ge \dim(\boldsymbol{\beta}) \text{ where } \tilde{Q}_0(\boldsymbol{\beta}) = \{Q_0^{(0)}(\boldsymbol{\beta}), Q_0^{(1)}(\boldsymbol{\beta}), \dots, Q_0^{(K)}(\boldsymbol{\beta})\}^T.$$

Proof. From (a) we have that $\boldsymbol{\beta}_0$ minimizes $Q_0(\boldsymbol{\beta})$ since $Q_0(\boldsymbol{\beta}) \ge 0$, $\forall \boldsymbol{\beta}$ and $Q_0(\boldsymbol{\beta}_0) = 0$.

Then from (b) and the Implicit Function Theorem, there exists a unique function $g\{G_L(x), G_R(x, y)\}$ such that $g\{G_L(x), G_R(x, y)\} = \beta_0$, in the neighborhood of $\{G_L(x), G_R(x, y)\}$ where $\{G_L(x), G_R(x, y)\} = [G_L(x), G_R(x, y); x \in \{0, 1, ..., K\}, y = 0, 1]$. Thus, β_0 is a unique local minimizer for $Q_0(\beta)$.

The proof of Theorem 3.3.5 is given as below.

Proof. From Proposition 7.1 in (Hayashi, 2000): an extremum estimator $\hat{\eta}$ is a consistent estimator for η if there is a function $Q_0(\eta)$ satisfying the following two conditions:

- (I) Identification: $Q_0(\eta)$ is uniquely maximized on H at $\eta_0 \in H$.
- (II) Uniform convergence: $Q_n(\cdot)$ converges uniformly in probability to $Q_0(\cdot)$.

The condition (I) is satisfied according to Lemma A.2.2. To show that the condition (II) is satisfied here, let

$$Q_{n}(\boldsymbol{\beta}) = \sum_{x=0}^{K} Q_{n}^{(x)}(\boldsymbol{\beta})^{2}$$

= $\sum_{x=0}^{K} \{ \widehat{G}_{L}(x) - \sum_{y=0}^{1} G_{M}(x, y; \boldsymbol{\beta}) \cdot \widehat{G}_{R}(x, y) \}^{2}$
 $Q_{0}(\boldsymbol{\beta}) = \sum_{x=0}^{K} Q_{0}^{(x)}(\boldsymbol{\beta})^{2}$
= $\sum_{x=0}^{K} \{ G_{L}(x) - \sum_{y=0}^{1} G_{M}(x, y; \boldsymbol{\beta}) \cdot G_{R}(x, y) \}^{2}.$

From

$$|Q_{n}(\boldsymbol{\beta}) - Q_{0}(\boldsymbol{\beta})| = |\sum_{x=0}^{K} Q_{n}^{(x)}(\boldsymbol{\beta})^{2} - \sum_{x=0}^{K} Q_{0}^{(x)}(\boldsymbol{\beta})^{2}|$$

$$\leq \sum_{x=0}^{K} |Q_{n}^{(x)}(\boldsymbol{\beta})^{2} - Q_{0}^{(x)}(\boldsymbol{\beta})^{2}|$$

$$= \sum_{x=0}^{K} |Q_{n}^{(x)}(\boldsymbol{\beta}) - Q_{0}^{(x)}(\boldsymbol{\beta})| \cdot |Q_{n}^{(x)}(\boldsymbol{\beta}) + Q_{0}^{(x)}(\boldsymbol{\beta})|$$

$$\leq \sum_{x=0}^{K} 2 \cdot |Q_{n}^{(x)}(\boldsymbol{\beta}) - Q_{0}^{(x)}(\boldsymbol{\beta})|, \quad x \in \Gamma$$

because $0 \leq |Q_n^{(x)}(\boldsymbol{\beta})| \leq 1$ and $0 \leq |Q_0^{(x)}(\boldsymbol{\beta})| \leq 1$, each of which is a difference of two probability estimates.

Therefore,

$$\begin{aligned} |Q_{n}(\boldsymbol{\beta}) - Q_{0}(\boldsymbol{\beta})| \\ &\leq \sum_{x=0}^{K} 2 \cdot \left\{ |\widehat{G}_{L}(x) - G_{L}(x)| + \sum_{y=0}^{1} G_{M}(x, y; \boldsymbol{\beta}) \cdot |\widehat{G}_{R}(x, y) - G_{R}(x, y)| \right\} \\ &\leq \sum_{x=0}^{K} 2 \cdot \left\{ |\widehat{G}_{L}(x) - G_{L}(x)| + \sum_{y=0}^{1} |\widehat{G}_{R}(x, y) - G_{R}(x, y)| \right\} \end{aligned}$$
(15)

because $G_M(x, y; \boldsymbol{\beta})$ is a probability bounded between 0 and 1.

Since $\widehat{G}_L(x)$ and $\widehat{G}_R(x,y)$ are either sample proportions or their ratios,

$$\widehat{G}_L(x) \xrightarrow{p} G_L(x), \text{ as } n \to \infty$$
$$\widehat{G}_R(x, y) \xrightarrow{p} G_R(x, y), \text{ as } n \to \infty$$

As $\widehat{G}_L(x)$ and $\widehat{G}_R(x,y)$ do not involve β , from (15) we have

$$Q_n(\boldsymbol{\beta}) \stackrel{p}{\Longrightarrow} Q_0(\boldsymbol{\beta}), \text{ as } n \to \infty$$

where $\stackrel{p}{\Longrightarrow}$ denotes uniform convergence in probability. This confirms condition (II) and completes the proof of $\hat{\boldsymbol{\beta}} \stackrel{p}{\to} \boldsymbol{\beta}$ as $n \to \infty$.

Because the causal estimate $\hat{\theta}$ is a continuously differentiable function of $\hat{\beta}$ and relevant sample proportions, by Slutsky's theorem, $\hat{\theta} \xrightarrow{p} \theta$ as $n \to \infty$.

A.3 Calculation of True Principal Stratum Causal Effects

For the simulated data, the true average causal effect for principal stratum $S_i(1) = 1$ can be calculated by

$$\mathbb{E}\{Y_i(1) - Y_i(0) | S_i(1) = 1\} = \mathbb{E}\{Y_i(1) = 1 | S_i(1) = 1\} - \mathbb{E}\{Y_i(0) = 1 | S_i(1) = 1\}$$
$$= \frac{\Pr\{Y_i(1) = 1, S_i(1) = 1\} - \Pr\{Y_i(0) = 1, S_i(1) = 1\}}{\Pr\{S_i(1) = 1\}}$$

where

$$\Pr\{S_i(1) = 1\} = \sum_x \left\{ \Pr\{S_i(0) = 1 | X_i = x\} \cdot \Pr\{X_i = x\} \right.$$
$$\left. + \sum_y \left[\Pr\{X_i = x\} \cdot \Pr\{S_i(0) = 0 | X_i = x\} \right.$$
$$\left. \cdot \Pr\{Y_i(0) = y | S_i(0) = 0, X_i = x\} \right.$$
$$\left. \cdot \Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0) = y, X_i = x\} \right] \right\}$$

$$\Pr\{Y_i(0) = 1, S_i(1) = 1\} = \sum_x \left[\Pr\{X_i = x\} \cdot \Pr\{S_i(0) = 1 | X_i = x\} \right.$$
$$\cdot \Pr\{Y_i(0) = 1 | S_i(0) = 1, X_i = x\}$$
$$+ \Pr\{X_i = x\} \cdot \Pr\{S_i(0) = 0 | X_i = x\}$$
$$\cdot \Pr\{Y_i(0) = 1 | S_i(0) = 0, X_i = x\}$$
$$\cdot \Pr\{S_i(1) = 1 | S_i(0) = 0, Y_i(0) = 1, X_i = x\} \right]$$

$$\begin{aligned} \Pr\{Y_i(1) = 1, S_i(1) = 1\} &= \sum_x \sum_y \left[\Pr\{X_i = x\} \cdot \Pr\{S_i(0) = 1 | X_i = x\} \\ &\cdot \Pr\{Y_i(0) = y | S_i(0) = 1, X_i = x\} \\ &\cdot \Pr\{Y_i(1) = 1 | Y_i(0) = y, S_i(0) = 1, X_i = x\} \\ &+ \Pr\{X_i = x\} \cdot \Pr\{S_i(0) = 0 | X_i = x\} \\ &+ \Pr\{Y_i(0) = y | S_i(0) = 0, X_i = x\} \\ &\cdot \Pr\{Y_i(0) = y | S_i(0) = 0, Y_i(0) = y, X_i = x\} \\ &\cdot \Pr\{Y_i(1) = 1 | S_i(0) = 0, S_i(1) = 1, Y_i(0) = y, X_i = x\} \end{aligned}$$

Appendix B for Chapter 3

B.1 Related Topics and Distinctions

In Section 3.2, we focused on the literature review of model averaging for ease of exposition, because the most innovated part of our method is motivated directly from this class of work. Here we clarify the main differences among model averaging, meta-analysis, federated learning, as well as super learner.

Model averaging: a convex averaging of models via model-specific weights (Raftery et al., 1997; Yang, 2001; Dai and Zhang, 2011; Yao et al., 2018; Dai et al., 2018). The extension of the weights from scalars to functions provides the best motivation for our approach.

Meta-analysis: classic in the way that it describes the site-level heterogeneity using modeling assumptions (Whitehead, 2002; Sutton et al., 2000), rather than a more datadriven approach such as tree models. It can be either frequentist or Bayesian, the latter of which tends to be more useful under limited sample sizes. However, its main interest is typically the overall effect rather than the site-level heterogeneity, which is usually modeled by a nuisance parameter (Borenstein et al., 2011; Riley et al., 2011; Tan et al., 2018; Röver and Friede, 2020).

Federated learning: originated from the field of computer science (McMahan et al., 2017), federated learning is a collaborative learning procedure that ensures data privacy by exchanging model parameters only. Federated learning methods often involves iterative updating (Fallah et al., 2020; Cho et al., 2021; Smith et al., 2017; Yang et al., 2019), rather than a one-shot procedure, which could be hard to apply to nonautomated distributed research networks. It has been developed mainly to estimate a global prediction model by leveraging distributed data (Li et al., 2020; Kairouz et al., 2019; Zhao et al., 2018; Hard et al., 2018), and is not designed to target any specific site.

Super learner: an ensemble of multiple statistical and machine learning models (van der Laan et al., 2007). It learns an optimal weighted average of those candidate models by minimizing the cross-validated risk, and assigns higher weights to more accurate models (Polley and van der Laan, 2010). The final prediction on an independent testing data is the weighted combination of the predictions of those models. Super learner has been showed empirically to improve treatment effect estimation via the modeling of propensity score in observational studies (Pirracchio et al., 2015; Wyss et al., 2018; Shortreed et al., 2019; Ju et al., 2019; Tan et al., 2022d).

Mixture of experts: an ensemble learning technique that decomposes a task into multiple subtasks with domain knowledge, followed by using multiple expert models to handle each subtask. A gating model is then used to decide which expert to use to make future prediction (Masoudnia and Ebrahimpour, 2014). It differs from other ensemble methods typically in that often only a few experts will be selected for predictions, rather than combining results from all experts (Masoudnia and Ebrahimpour, 2014).

B.2 Proof of Theorem 3.3.5

The proof of Theorem 3.3.5 closely follows arguments given in Wager and Athey (2018). Suppose the subsamples for building each tree in an ensemble forest are drawn from different subjects in the augmented site 1 data. Specifically, in one round of EF, we draw m samples from the augmented data, where m is less than the rows in the augmented data, i.e., $m < (n_1 \cdot K)$. By randomly picking m unique subjects from site 1 and then randomly picking a site indicator k out of K sites for each of the m subjects. The resulted m subsamples should not be from the same subject and are hence independent and identically distributed. As long as $m < n_1$, we can ensure that all the subsamples are independent. In practice, when the ratio of n_1/K is relatively large, the probability of obtaining samples from the same subject is small.

Assume that subject features X_i and the site indicator S_i are independent and have a density that is bounded away from 0 and infinity. Suppose moreover that the conditional mean function $\mathbb{E}[\mathcal{T}|X = x, S = k]$ is Lipschitz continuous. We adopt the honesty definition in Athey and Imbens (2016) when building trees in a random forest. Honest approaches separate the training sample into two halves, one half for building the tree model, and another half for estimating treatment effects within the leaves (Athey and Imbens, 2016). Following Definitions 1-5 and Theorem 3.1 in Wager and Athey (2018), the proposed estimator $\widehat{\mathcal{T}}_{\text{EF}}(\boldsymbol{x}, 1)$ is a consistent estimator of the true treatment effect function $\tau_1(\boldsymbol{x})$ for site 1.

B.3 Additional Simulation Results

B.3.1 Connection to Supervised Learning

Similar to ET-oracle and EF-oracle whose weights are built on the ground truth CATE functions τ_k 's, we also consider for EWMA and STACK under a similar hypothetical setting. Specifically, we assume the true τ_1 is known and use it to compute the weights. This version of EWMA estimator is denoted as EWMA-oracle and its weight is given by

$$\omega_k^{\text{EWMA-oracle}} = \frac{\exp\{-\sum_{i \in \mathcal{I}_1^{(2)}} (\widehat{\tau}_k(\boldsymbol{x}_i) - \tau_1(\boldsymbol{x}_i))^2\}}{\sum_{\ell=1}^K \exp\{-\sum_{i \in \mathcal{I}_1^{(2)}} (\widehat{\tau}_\ell(\boldsymbol{x}_i) - \tau_1(\boldsymbol{x}_i))^2\}}.$$

Similarly, the corresponding linear stacking approach, denoted as STACK-oracle, regresses the ground truth $\tau_1(\boldsymbol{x})$ on the predictions of the estimation set in site 1 from each local model, $\{\hat{\tau}_1(\boldsymbol{x}), \ldots, \hat{\tau}_k(\boldsymbol{x})\}$. We compare the proposed model averaging estimators with the local estimator, MA, two versions of modified EWMA, as well as two versions of the linear stacking approach. We present simulation results using CT as the local model and the sample size at local sites to be n = 500. Figure 6 presents the performance of the proposed estimators along with other competing estimators. Each series of boxes corresponds to a different strength of global heterogeneity c. Table 8 reports the ratio between MSE of the estimator and MSE of the local model in terms of average and standard deviation of MSE, respectively, over 1000 replicates. Our proposed estimators ET and EF shows the best performance overall in terms of the mean and variation of MSE among the estimators without using the information of ground truth $\tau_1(\boldsymbol{x})$. Comparing with ET, EF has a slightly smaller MSE when c is large, which is expected because forest models tend to be more stable and accurate than a single tree. ET-oracle achieves minimal MSE for low and moderate degrees of heterogeneity while EF-oracle has the minimal MSE under all settings. The local estimator (LOC) in general shows the largest MSE compared to other estimators, as it does not leverage information from other sites. By borrowing information from additional sites, variances are greatly reduced, resulting in a small MSE of ensemble estimators. MA that naively adopts the inverse of sample size as weights performs well under low levels of heterogeneity, but suffers from a huge MSE with large variation as c increases. EWMA estimators perform slightly better and are more stable than LOC and MA. EWMA-oracle has better performance than EWMA in all settings as the information of true CATE is used for weight construction. STACK estimators performs better than EWMA estimators. Similarly, STACK-oracle performs better than STACK in all settings. STACK-oracle, with ground truth $\tau_1(\mathbf{x})$ available, outperforms ET and EF when there exists a moderate to high level of heterogeneity across sites.

B.3.2 Various Sample Sizes in Local Sites

We provide detailed simulation results varying n (100, 500, 1000) with CT as the local model. Figure 7 and Figure 8 show box plots of simulation results with a sample size of 100 and 1000, respectively, at each site. Our proposed methods ET and EF show robust performance in all settings. ET-oracle and EF-oracle achieve close-to-zero MSE with very small spreads in some settings. Figure 9 shows plots of the bias and MSE of EF-oracle varying sample size at each site (n = 100, 500, 1000). As the sample size increases, both bias and MSE of EF-oracle reduce to zero. Consistency of EF-oracle can be shown via simulation when perfect estimates are obtained from local models. Meanwhile, our proposed method greatly reduce MSE by selectively borrowing information from multiple sites.

B.3.3 Simulations under Observational Studies

We also consider the treatment generation mechanism under an observational design. Specifically, the propensity is given as $e(\mathbf{x}) = \exp((0.6x_1))$. We consider both a correctly specified propensity model using a logistic regression of Z on X_1 and a misspecified propensity model with a logistic regression of Z on all \mathbf{X} . Figure 10 and Figure 11 show box plots of simulation results. In general, the proposed estimators obtain the best performance with similar results are obtained as in the Figure 3. With the correctly specified propensity score model, the local estimator is consistent in estimating $\tau_k(\boldsymbol{x})$, the proposed framework is valid. When the propensity model misspecified, extra uncertainty is carried forward from the local estimates, but the proposed estimators can improve upon the local models. This is due to a bias-and-variance trade-off that guarantees small MSE in prediction, which remains smaller than those from local estimators.

B.3.4 Covariate Dimensions

We consider various choices of covariate dimensions besides D = 5. Specifically, we also try D = 20 and D = 50. Figure 12 and Figure 13 show box plots of simulation results. With a higher dimension of variables, the MSE ratio between the proposed estimates and LOC estimates increases than that in the scenario with a small dimension.

B.3.5 Unequal Sample Size at Each Site

In the distributed date network, different sites may have a different sample size n_k . Those with a smaller sample size may not be representative of their population, leading to an uneven level of precision for local causal estimates. We consider a simulation setting where site 1 has a sample size of $n_1 = 500$ while other site n_2, \ldots, n_K has a sample size of 200. Figure 14 shows box plots of simulation results. Results show that the MSE ratio between the proposed estimates and LOC estimates increases compared to the scenario where the sample size in all sites are 500. However, the proposed estimators still enjoy the most robust performance via bias-and-variance trade-off. This also shows our method is robust to the existence of local uncertainty.

B.3.6 Different Local Estimators

We explore another option for the local model using the causal forest (CF) (Wager and Athey, 2018) varying the sample size at local sites. A causal forest is a stochastic averaging of multiple causal trees (Athey and Imbens, 2016), and hence is more powerful in estimating treatment effects. In each tree of the causal forest, MSE of treatment effect is used to select the feature and cutoff point in each split (Wager and Athey, 2018). CF is implemented in the R packages grf. Figure 15, Figure 16, and Figure 17 show box plots of simulation results with a sample size of 100, 500, and 1000, respectively, at each site. Our proposed methods ET and EF show robust performance in all settings regardless of the use of information of the ground truth $\tau_1(\boldsymbol{x})$.

B.3.7 Further Comparisons to Non-adaptive Ensemble

We provide simulation results to compare the proposed methods to the non-adaptive method STACK. Consider the following setting where the heterogeneity is continuous and nonlinear: $\tau(\boldsymbol{x}, k) = \mathbb{1}\{x_1 > 0\} \cdot x_1 + (x_1 - 3) \cdot (U_k)^c$, with $U_k \sim Unif[0,3]$, $\boldsymbol{X}_i \sim N(\mathbf{0}, \boldsymbol{I}_5)$, and c = (1, 2, 3, 4). As c increases, the heterogeneity across sites gets larger, reducing the influence of x_1 on heterogeneity, hence the weights become more non-adaptive. For c = (1, 2, 3, 4), the one-SD ranges of MSE ratios of EF over STACK are [0.73,0.82], [0.86,0.87], [0.99,1.04], [0.87,1.07], respectively. When c is relatively small, the proposed EF has a smaller MSE compared to STACK. As c increases, the performance of EF is similar to that of STACK, in the case of a large global heterogeneity. This further indicates the robustness of the proposed methods.



Figure 6: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.

Table 8: Simulation results for ratio between MSE of the estimator and MSE of LOC (**CT**) with a sample size of **500** at each site. A smaller number indicates larger improvement over the local model. Estimators ending with "-oracle" makes use of ground truth treatment effects. Our proposed methods ET and EF shows robust performance in all settings whether or not using the information of ground truth $\tau_1(\boldsymbol{x})$.

	Discrete grouping				Continuous grouping				
Estimator	c = 0	c = 0.2	c = 0.6	c = 1	c = 0	c = 0.2	c = 0.6	c = 1	
Ratio of average of MSEs over 1000 replicates									
MA	0.09	0.91	2.4	9.87	0.08	0.32	0.65	1.78	
EWMA	0.57	0.62	0.61	0.62	0.56	0.65	0.7	0.77	
EWMA-oracle	0.42	0.5	0.49	0.5	0.42	0.49	0.53	0.59	
STACK	0.44	0.45	0.44	0.45	0.45	0.45	0.48	0.54	
STACK-oracle	0.06	0.04	0.04	0.04	0.06	0.06	0.06	0.07	
ET	0.12	0.17	0.16	0.16	0.13	0.24	0.29	0.37	
ET-oracle	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	0.08	0.1	0.07	
EF	0.1	0.13	0.13	0.13	0.1	0.19	0.25	0.3	
EF-oracle	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	0.06	0.06	0.05	
Ratio of standard deviation of MSEs over 1000 replicates									
MA	0.15	0.35	0.76	3.05	0.14	0.24	0.38	0.81	
EWMA	0.61	0.65	0.67	0.66	0.58	0.65	0.69	0.75	
EWMA-oracle	0.46	0.52	0.54	0.54	0.44	0.52	0.55	0.6	
STACK	0.47	0.46	0.47	0.47	0.45	0.49	0.52	0.6	
STACK-oracle	0.1	0.08	0.08	0.08	0.09	0.11	0.12	0.14	
ET	0.18	0.23	0.22	0.22	0.18	0.26	0.32	0.43	
ET-oracle	0.02	0.03	0.02	0.02	0.02	0.06	0.07	0.07	
EF	0.17	0.19	0.19	0.2	0.17	0.23	0.29	0.39	
EF-oracle	0.03	0.03	0.03	0.03	0.03	0.06	0.07	0.08	



Figure 7: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **100** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 8: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **1000** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 9: Plots of the bias and MSE of **EF-oracle** varying sample site at each site for (a) discrete grouping and (b) continuous grouping across site, varying scale of global heterogeneity. Both bias and MSE reduces to zero as the sample size increases.



Figure 10: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site under **observational design with a correctly specified propensity score model** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 11: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (CT) and a sample size of 500 at each site under observational design with a misspecified propensity score model for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.


Figure 12: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site, and covariate dimension of **20** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 13: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CT**) and a sample size of **500** at each site, and covariate dimension of **50** for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 14: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (CT) and a sample size of 500 at site 1, and a sample size of 200 at other sites for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 15: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CF**) and a sample size of **100** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 16: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CF**) and a sample size of **500** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.



Figure 17: Box plots of the MSE ratios of CATE estimators, respectively, over LOC (**CF**) and a sample size of **1000** at each site for (a) discrete grouping and (b) continuous grouping across site, respectively, varying scale of global heterogeneity. Estimators ending with "-oracle" makes use of ground truth treatment effects. Different colors imply different estimators, and x-axis, i.e., the value of c, differentiates the scale of global heterogeneity. The red dotted line denotes an MSE ratio of 1. MA performance is truncated due to large MSE ratios. The proposed ET and EF achieve competitive performance compared to standard model averaging or ensemble methods and are robust to heterogeneity across settings. Note that ET-oracle and EF-oracle achieve close-to-zero MSE ratios with very small spreads in some settings.

B.4 Additional Results for Data Application

In real-life applications, hospitals may have different sample sizes n_k that may affect the accuracy of the estimation of τ_k . Table 9 shows hospital-level information for the 20 hospitals where the number of patients across sites varies. Information includes the region of the U.S. where the hospital is located, whether it is a teaching hospital, the bed capacity, and the number of patients within the hospital.

Hospitals with a smaller sample size may not be representative of the population, leading to an uneven level of precision for local causal estimates. To account for different sample sizes at each hospital, we consider a basic weighting strategy where we add weights to each observation $\hat{\tau}_k(\boldsymbol{x})$ in the augmented site 1 data adjusting for the sample size of site k. The weights are defined as $\eta_k(\boldsymbol{x}) = K n_k \{\sum_{j=1}^K n_j\}^{-1}$.

Figure 18 visualizes the performance of oxygen therapy on hospital survival with the weighting strategy adopted. CT is used as the local model with propensity score modeled by a logistic regression. Figure 18(a) shows the propensity score-weighted average survival for those whose received treatment is consistent with the estimated decision. Treatment rule based on our method can increase survival by 4%, more promising than the EF estimates without the weighting strategy and the LOC and the baseline. The weighting strategy takes account into the unequal sample size among the hospital network, and assign weights based on precision of local estimates.

In the fitted EF, the hospital indicator remains the most important, explaining about 48% of the decrease in training error. Figure 18(b) shows the estimated CATEs varying two important features, BMI and oxygen therapy duration. Patients with BMI between 36 and 40 and duration above 400 show the most benefit from oxygen therapy in the target SpO₂ range. Patients with BMI between 20 and 30 and duration between 100 and 400 may not benefit from such alteration. The treatment estimates are similar to that in Figure 5(b) Figure 18(c) visualizes the proposed model averaging scheme with data-adaptive weights $\omega_k(\boldsymbol{x})$ in the fitted EF with respect to BMI for different models, while holding other variables constant. The weights of hospital 1 are quite stable while models from other sites may have different contribution to the weighted estimator for different values of BMI. Similar to Figure 5(c),

Table 9: Hospital-level information of our analysis cohort in eICU database. Hospitals are relabeled according to their average contribution to the estimation task at hospital 1, the target site.

Hospital	Number of	Number of	Number of	Bed	Teaching	Domion
site	patients	$\operatorname{control}$	treated	capacity	status	Region
1	477	205	272	≥ 500	False	South
2	297	109	188	≥ 500	True	West
3	163	58	105	≥ 500	True	Midwest
4	222	58	164	≥ 500	False	South
5	659	165	494	≥ 500	True	Midwest
6	305	174	131	≥ 500	False	South
7	347	109	238	≥ 500	True	Midwest
8	523	162	361	≥ 500	False	South
9	210	78	132	Unknown	False	Unknown
10	379	161	218	≥ 500	True	Midwest
11	234	70	164	≥ 500	True	Midwest
12	747	185	562	≥ 500	True	Northeast
13	464	129	335	≥ 500	True	South
14	474	229	245	≥ 500	False	South
15	166	64	102	100 - 249	False	Midwest
16	388	94	294	≥ 500	False	Midwest
17	435	240	195	≥ 500	True	South
18	200	55	145	250 - 499	False	South
19	183	52	131	250 - 499	False	West
20	149	71	78	250 - 499	False	South

hospitals with a larger bed capacity tend to be similar to hospital 1, and are shown to provide larger contributions. In general, the weighting strategy helps further improve the expected survival rate. The patterns in each subfigure are similar to Figure 5, which indicates the robustness of our proposed estimators. We do stress that improvements to the weighting strategy for different sample sizes at each site are needed. A strategy considering both treatment proportion as well as covariate distributions across sites may further enhance the data-adaptive model averaging estimator.



Figure 18: Application to estimating treatment effects of oxygen therapy on survival with a **sample size weighting strategy**. (a) Expected survival of treatment decision following different estimators. (b) Estimated treatment effects varying duration and BMI, two important features in the fitted EF. (c) Visualization of data-adaptive weights in EF varying BMI.

B.5 Real Data Access

Although the eICU data used in our application example cannot be shared subject to the data use agreement, access can be individually requested at https://eicu-crd.mit.edu/gettingstarted/access/.

Appendix C for Chapter 4

C.1 Additional Literature Review

Typical model-based methods for deriving IDRs include Q-learning such as (Watkins and Dayan, 1992; Murphy, 2003; Moodie et al., 2007; Chakraborty et al., 2010; Goldberg and Kosorok, 2012; Song et al., 2015) and A-learning such as (Robins et al., 2000; Murphy, 2005) where a model of responses is imposed and the optimal decision rule is obtained by optimizing value function derived from the model. Model-based methods posit a model of responses given observed covariates and treatment assignments, and obtain the optimal IDR by optimizing the corresponding value function derived from the model. Q-learning optimizes the corresponding value function derived from a parametric model of responses given observed covariates and treatment assignments, and it results in an optimal decision rule. A-learning is a semiparametric method, which derives from a model that directly describes the difference between treatments, with the baseline remaining unspecified. On the other hand, model-free methods such as Robins et al. (2008); Orellana et al. (2010a,b); Zhang et al. (2012); Zhao et al. (2012, 2015) assign values to actions simply through trial and error without pre-specifying a model. Besides, contextual bandit methods (see Bietti et al. (2021) and references therein) test out different actions and automatically learn which one has the most rewarding outcome for a given situation. Other methods include (Robins, 2004; Moodie et al., 2009; Cai et al., 2011; Henderson et al., 2010; Thall et al., 2002; Imai and Ratkovic, 2013; Huang et al., 2015; Tao and Wang, 2017). See (Chakraborty et al., 2010; Chakraborty and Moodie, 2013; Laber et al., 2014; Kosorok and Moodie, 2015) and references therein for a comprehensive review on general IDRs under causal settings.

C.2 Proofs of Propositions

Here we show proofs of Propositions 4.3.4 and 4.3.5 in Section 4.3.4.

Proof of Proposition 4.3.4. We observe that to maximize the objective function in (12) is equivalent to maximizing

$$\begin{split} &E_X \Big[G_{S|X} \{ E(Y|X, S, A = d(X)) \} |X \Big] \\ &= E_X \Big[G_{S|X} \{ E(Y|X, S, A = 1) \} \mathbb{1}(d(X) = 1) \\ &+ G_{S|X} \{ E(Y|X, S, A = -1) \mathbb{1}(d(X) = -1) \} \Big] \\ &= E_X \Big\{ \mathbb{1}(d(X) = 1) [G_{S|X} \{ E(Y|X, S, A = 1) \} - G_{S|X} \{ E(Y|X, S, A = -1) \}] \\ &+ G_{S|X} \{ E(Y|X, S, A = -1) \} \Big\} \\ &\propto E_X \Big\{ \mathbb{1}(d(X) = 1) [G_{S|X} \{ E(Y|X, S, A = 1) \} - G_{S|X} \{ E(Y|X, S, A = -1) \}] \Big\}. \end{split}$$

Proof of Proposition 4.3.5. Let $d(x) = \operatorname{sgn}\{f(x)\}\)$, by this transformation, we consider the following objective on a smooth function f(x),

$$\begin{aligned} \arg\max_{d\in\mathbb{D}} \frac{1}{n} \sum_{i=1}^{n} \left\{ \mathbbm{1}(d(x_i) = 1)[g_1(x_i) - g_2(x_i)] \right\} \\ &= \arg\max_f \frac{1}{n} \sum_{i=1}^{n} \mathbbm{1}[\operatorname{sgn}\{f(x_i)\} = 1] \cdot [g_1(x_i) - g_2(x_i)] \\ &= \arg\min_f \frac{1}{n} \sum_{i=1}^{n} \mathbbm{1}\{1 \cdot f(x_i) < 0\} \cdot [g_1(x_i) - g_2(x_i)] \\ &= \arg\min_f \frac{1}{n} \sum_{i=1}^{n} \mathbbm{1}[\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)|. \end{aligned}$$

The sign of the estimated f above is a d to (13).

Hence, the proposed classification-based objective is to minimize

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{1}[\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\} \cdot f(x_i) < 0] \cdot |g_1(x_i) - g_2(x_i)|.$$

To this point, we have transformed the optimization problem (12) into a weighted classification problem where for subject i with features x_i , the true label is $\operatorname{sgn}\{g_1(x_i) - g_2(x_i)\}$ and the sample weight is $|g_1(x_i) - g_2(x_i)|$.

C.3 Details on Modeling and Hyperparameter Tuning

In our implementation, neural networks with mean or quantile losses are used to fit the models with hyperparameters tuned via a 5-fold cross validation in the training data sets. Specifically, implemented in TensorFlow (Abadi et al., 2016), neural networks with mean squared loss is used to model E(Y|X, S, A) separated by the control arm and the treatment arm, respectively. For continuous S, to model $Q_{S|X,A}{E(Y|X, S, A)}$, neural networks with quantile loss is used with a prespecified τ , for the control arm and the treatment arm, respectively. In the final weighted classification model, neural networks with cross-entropy loss is used. Note that the model choices here are flexible. One can perform model selection if they would like to.

Hyperparameter tuning helps prevent overfitting and is essential in machine learning methods or other black-box algorithms such as neural networks. In our implementation, the optimal hyperparameters are obtained via a 5-fold cross validation in the training data sets. Specifically, we consider the number of hidden layers (1, 2, and 3 layers), the number of hidden units in each layer (256, 512, and 1024 nodes), activation function (RELU, Sigmoid, and Tanh), optimizer (Adam, Nadam, and Adadelta), dropout rate (0.1, 0.2, and 0.3), number of epochs (50, 100, and 200), and batch size (32, 64, and 128).

C.4 Additional Simulations

C.4.1 Different Quantile Criteria

For the quantile criteria, we also consider $\tau = 0.1$ and 0.5, respectively. Table 10 presents the simulation results for Example 2 with continuous S using 0.1 quantile criterion and 0.5 quantile criterion, respectively. Results show that when τ is small, there is more strength in the proposed method, as the algorithm aims to improve the worst-outcome scenarios. The proposed RISE has the largest gain in objective and value among vulnerable subjects when τ is 0.1, and has similar performance as the compared approaches when τ is 0.5.

C.4.2 S as a Noise Variable

We generate the outcome Y using the following model where S is not involved: $Y = \mathbb{1}(X_1 \leq 0.5)\{8 + 12\mathbb{1}(A = 1) + 16\exp(X_2) - 26\mathbb{1}(A = 1)X_2\} + \mathbb{1}(X_1 > 0.5)\{13 + 3\mathbb{1}(A_i = 1) + 2\exp(X_2) - 8\mathbb{1}(A = 1)X_2\} + \epsilon$, where $X_j \sim U(0,1)$, $j = 1, 2, A \sim Bernoulli(0.5)$, and $\epsilon \sim N(0,1)$. For continuous S, $S = \exp{it}\{-2.5(1 - X_1 - X_2)\}$; for discrete S, we consider a binary S that satisfies $\log\{P(S = 1|X)/P(S = 0|X)\} = -2.5(1 - X_1 - X_2)$. Table 11 summarizes the performance of the proposed IDRs compared to the mean criterion for Example 2. The estimated objective and value function are similar for the compared IDRs, which indicates the robustness of the proposed RISE.

C.4.3 Violations of Causal Assumptions

To further test the robustness of the proposed RISE, we consider scenarios where the causal ssumptions in Section 4.3.1 may not hold. To test the violation of positivity assumption in Assumption 4.3.2, using the same setting as in Example 2, we consider an extreme propensity score, or the probability of being treated given X and S. Specifically, we let A satisfy $\log\{P(A_i = 1|X_i)/P(A = 0|X_i)\} = -1.2(-S_i + X_{i1} - X_{i2} + X_{i3} - X_{i4} + X_{i5} - X_{i6})$. To test the unconfoundedness assumption in Assumption 4.3.3, a random normal noise, $e \sim N(0, 1)$ is added to X_1 in the setting of Example 2. The simulation results are presented in Table 12 and Table 13 respectively.

Table 10: Simulation results for Example 2 with continuous S using 0.1 quantile criterion and 0.5 quantile criterion, respectively. Standard error in parenthesis. The proposed RISE has more strengths when τ is small, as the algorithm aims to improve the worst-outcome scenarios.

Type of S	au	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Cont.	0.1	Base	7.93(0.03)	7.92(0.03)	17.7(0.02)	8.64 (0.07)
		Exp	8.88 (0.05)	8.85(0.05)	17.8(0.02)	10.6 (0.12)
		RISE	13.8(0.01)	13.7 (0.02)	16.9(0.01)	20.9(0.03)
		PT-Base	6.97(0.02)	6.95(0.02)	17.9(0.03)	6.65 (0.04)
		PT-Exp	7.11(0.02)	7.08(0.03)	18.0(0.03)	6.96 (0.05)
Cont.	0.5	Base	17.3 (0.04)	17.2 (0.04)	17.7(0.02)	23.8(0.19)
		Exp	$17.2 \ (0.03)$	17.4 (0.03)	17.8(0.02)	22.1 (0.17)
		RISE	17.4(0.04)	17.4 (0.04)	17.8(0.02)	24.0 (0.22)
		PT-Base	17.3 (0.05)	$17.3\ (0.05)$	18.0(0.03)	23.9(0.26)
		PT-Exp	17.4 (0.05)	17.4 (0.05)	18.1 (0.03)	24.0 (0.25)

Table 11: Simulation results for scenario when S is a noise variable. Vulnerable subjects cannot be defined as S is not important in the example. The estimated objective and value function are similar for the compared IDRs, which indicates the robustness of the proposed RISE.

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	27.5(0.03)	-	27.5(0.06)	_
	Exp	27.5(0.03)	-	27.5(0.06)	-
	RISE	27.5 (0.03)	-	27.5 (0.06)	-
	PT-Base	27.5 (0.02)	-	27.5(0.03)	-
	PT-Exp	27.5 (0.02)	-	27.5(0.03)	-
Cont.	Base	27.2(0.04)	-	27.3(0.07)	-
	Exp	27.2 (0.04)	-	27.3(0.07)	-
	RISE	27.2 (0.04)	-	27.3(0.07)	-
	PT-Base	27.2 (0.04)	-	27.3(0.07)	-
	PT-Exp	27.2 (0.04)	-	27.3(0.07)	-

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	$10.0 \ (0.03)$	$11.1 \ (0.03)$	19.3 (0.02)	16.1 (0.04)
	Exp	8.80 (0.03)	9.77(0.04)	19.5 (0.02)	13.6(0.04)
	RISE	13.5(0.01)	14.0 (0.01)	17.3(0.01)	22.0 (0.02)
	PT-Base	9.88(0.03)	$10.7 \ (0.04)$	18.9(0.02)	$15.4 \ (0.05)$
	PT-Exp	8.42(0.03)	9.14(0.04)	$19.1 \ (0.02)$	$12.4 \ (0.05)$
Cont.	Base	11.5 (0.03)	11.5 (0.04)	17.5 (0.03)	$13.1 \ (0.04)$
	Exp	$10.4 \ (0.04)$	$10.4 \ (0.05)$	17.8(0.04)	$10.3 \ (0.05)$
	RISE	14.3 (0.01)	14.3 (0.02)	16.9(0.01)	20.4 (0.02)
	PT-Base	11.0 (0.04)	$10.9 \ (0.04)$	17.7 (0.02)	11.8(0.04)
	PT-Exp	9.63(0.03)	$9.61 \ (0.03)$	18.0 (0.02)	8.38 (0.03)

Table 12: Simulation results for Example 2 where the positivity assumption in Assumption 4.3.2 is nearly violated. Standard error in parenthesis.

Type of S	IDR	Obj. (all)	Obj. (vulnerable)	Value (all)	Value (vulnerable)
Disc.	Base	7.65(0.04)	8.44 (0.05)	19.3 (0.03)	11.1 (0.06)
	Exp	8.94 (0.05)	$9.91 \ (0.06)$	19.4 (0.02)	13.9(0.06)
	RISE	13.5(0.01)	14.0 (0.01)	$17.4\ (0.01)$	22.1 (0.02)
	PT-Base	6.84(0.03)	7.35(0.04)	18.9(0.03)	$8.91 \ (0.05)$
	PT-Exp	7.95~(0.05)	8.62(0.06)	$19.1 \ (0.03)$	$11.4 \ (0.06)$
Cont.	Base	9.58(0.03)	9.58(0.03)	17.9(0.02)	8.33(0.05)
	Exp	$10.2 \ (0.04)$	$10.2 \ (0.04)$	17.8(0.02)	9.83(0.06)
	RISE	14.2 (0.01)	14.1 (0.02)	16.9(0.01)	20.1 (0.03)
	PT-Base	9.27(0.02)	9.26(0.03)	17.9(0.02)	$7.51 \ (0.03)$
	PT-Exp	9.34(0.02)	9.34(0.03)	18.0 (0.02)	7.72(0.03)

Table 13: Simulation results for Example 2 where the unconfoundedness assumption in Assumption 4.3.3 is violated. Standard error in parenthesis.

C.5 Additional Information and Results for Real-data Applications

C.5.1 Data Availability

The access of the job training dataset (LaLonde, 1986) is available at https://users. nber.org/~rdehejia/data/.nswdata2.html. The access of the ACTG175 dataset (Hammer et al., 1996) is available from the R package speff2trial. The third dataset (Seymour et al., 2016) is not publicly available. All real data used in the chapter are deidentified with no personal information.

C.5.2 Additional Background on the Sepsis Application

Sepsis is leading cause of acute hospital mortality and commonly results in multi-organ dysfunction among ICU patients (Sakr et al., 2018; Onyemekwu et al., 2022). Clinically, treatment decisions for sepsis patients are needed to be made within a short period of time due to the rapid deterioration of patient conditions. Lactate and the Sequential Organ Failure Assessment (SOFA) score have been two important indicators of sepsis severity and has been found to be more useful for predicting the outcome of sepsis than other clinical vitals and comorbidity scores (Howell et al., 2007; Krishna et al., 2009; Shankar-Hari et al., 2016; Machicado et al., 2021). Typically, information of baseline patient characteristics such as age, gender, race, and weight, and common vital signs such as usage of mechanical ventilation, respiratory rate, temperature, intravenous fluids, Glasgow Coma Scale score, platelets, blood urea nitrogen, white blood cell counts, glucose, and creatinine are obtained at the admission of patients. On the other hand, SOFA score combines performance of several organ systems in the body such as neurologic, blood, liver, and kidney (Seymour et al., 2016; Liu et al., 2019) and cannot be obtained directly. Lactate labs measures the level of lactic acid in the blood (Andersen et al., 2013) and are less common in routine examination, which could be delayed in ordering. Hence, their information may not be available by the time of treatment decision due to multiple reasons including doctors' delayed ordering, long laboratory processing time, or the rapid deterioration of development of sepsis, which poses tremendous difficulties for early diagnosis and treatment decisions within a short time. According to the new definition

of Sepsis-3 (Shankar-Hari et al., 2016), a serum lactate level >2 mmol/L is considered to be in critical conditions and is highly likely to indicate a septic shock. Also, a SOFA score greater than 6 has been associated with a higher mortality (Vincent et al., 1996; Ferreira et al., 2001).

C.5.3 Visualizations

Here we provide visualizations of features that are important in the estimated decision rules for the three real-data applications in Section 4.4.2. The Shapely additive explanations (SHAP) (Lundberg and Lee, 2017) is considered to be a united approach to explaining the predictions of any machine learning or black-box models. Figure 19, Figure 20, and Figure 21 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively, for the three real-data applications. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment A = 1 and the blue indicates a negative correlation. Overall, the direction of correlation is similar for RISE and Exp, but their ranking of feature importance may be different.

Fairness in a job training program. Figure 19 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that whether having a high school diploma and income in 1974 are two important features in the variable important plot by RISE, while incomes in 1974 and in 1975 are important by Exp. It seems that being no degree and low income in 1974 has a higher chance of assigning A = 1 (to receive the job training program) by RISE, while low income in 1974 and but a higher income in 1975 may be associated with assigning A = 1 by Exp.

Improvement of HIV treatment. Figure 20 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that age and CD4 T-cell counts are two important features in the variable important plot by RISE, while weight and number of days of previously received antiretroviral therapy are important by Exp. It seems that being of a younger age and high CD4 T-cell count has a higher chance of assigning A = 1 (zidovudine combined with didanosine) by RISE, while being of a larger weight and few days of previously received antiretroviral therapy may be associated with assigning the treatment by Exp.

Safe resuscitation for patients with sepsis. Figure 21 presents the SHAP variable importance plots in the final weighted classification model by RISE and Exp, respectively. We observe that Glasgow Coma Scale score, age, and platelets appears to be important features in both the plot by RISE and that by Exp. Other important features in the plot by RISE include temperature and blood urea nitrogen, where in the plot by Exp, respiratory rate and white blood cell counts are of top importance. Being in a low temperature with a high blood urea nitrogen tends to be predicted as A = 1 (to assign vasopressors) by RISE while being of higher respiratory rate with high white blood cell counts tends to be predicted as A = 1 by Exp.

C.6 Causal and Decision Diagrams

We provide a causal diagram and a decision diagram under consideration, respectively, in Figure 22. As the causal diagram shows, both X and S confound the effect of treatment A on outcome Y. The arrows represent the causal relationship between variables. Hence, both X and S should be used for decision making in general. On the other hand, in the decision diagram under our setting, sensitive variables S is shown in a dotted circle as S may not be readily available at the time of decision making. We connect S and A with a dotted arrow to indicate that S is incorporated during model training of the decision rule, but not model deployment. Hence, the decision rule d only maps X to the treatment option, i.e., A = d(X).



Figure 19: Visualization for the job training program: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment A = 1 and the blue indicates a negative correlation.



Figure 20: Visualization for the ACTG175 dataset: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment A = 1 and the blue indicates a negative correlation.



Figure 21: Visualization for the sepsis data: SHAP variable importance plots for decision rules RISE (a) and Exp (b), respectively. Covariates (X) are ranked by variable importance in descending order. Correlations between the feature and their SHAP value are highlighted in color. The red color means a feature is positively correlated with assigning treatment A = 1 and the blue indicates a negative correlation.



Figure 22: (a) Causal diagram. Both X and S confound the effect of treatment A on outcome Y. The arrows represent the causal relationship between variables. Hence, both X and S should be used for decision making in general. (b) Decision diagram. Sensitive variables S is shown in a dotted circle as S may not be readily available at the time of decision making. We connect S and A with a dotted arrow to indicate that S is incorporated during model training of the decision rule, but not model deployment. Hence, the decision rule d only maps X to the treatment option, i.e., A = d(X).

Bibliography

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). {TensorFlow}: A system for {Large-Scale} machine learning. In 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pages 265–283.
- Ackerman, B., Schmid, I., Rudolph, K. E., Seamans, M. J., Susukida, R., Mojtabai, R., and Stuart, E. A. (2019). Implementing statistical methods for generalizing randomized trial findings to a target population. *Addictive Behaviors*, 94:124–132.
- Allcott, H. (2015). Site selection bias in program evaluation. The Quarterly Journal of Economics, 130(3):1117–1165.
- Andersen, L. W., Mackenhauer, J., Roberts, J. C., Berg, K. M., Cocchi, M. N., and Donnino, M. W. (2013). Etiology and therapeutic approach to elevated lactate levels. *Mayo Clinic Proceedings*, 88(10):1127–1140.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455.
- Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. Proceedings of the National Academy of Sciences, 113(27):7353–7360.
- Athey, S., Tibshirani, J., Wager, S., et al. (2019). Generalized random forests. Annals of Statistics, 47(2):1148–1178.
- Athey, S. and Wager, S. (2021). Policy learning with observational data. *Econometrica*, 89(1):133–161.
- Bareinboim, E. and Pearl, J. (2016). Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352.
- Barocas, S., Hardt, M., and Narayanan, A. (2019). *Fairness and Machine Learning*. fairml-book.org. http://www.fairmlbook.org.
- Bartolucci, F. and Grilli, L. (2011). Modeling partial compliance through copulas in a principal stratification framework. *Journal of the American Statistical Association*, 106(494):469–479.
- Battey, H., Fan, J., Liu, H., Lu, J., and Zhu, Z. (2018). Distributed testing and estimation under sparse high dimensional models. *Annals of Statistics*, 46(3):1352–1382.
- Bear, H. D., Tang, G., Rastogi, P., Geyer Jr, C. E., Liu, Q., Robidoux, A., Baez-Diaz, L., Brufsky, A. M., Mehta, R. S., Fehrenbacher, L., et al. (2015). Neoadjuvant plus adjuvant

bevacizumab in early breast cancer (NSABP B-40 [NRG Oncology]): secondary outcomes of a phase 3, randomised controlled trial. *The Lancet Oncology*, 16(9):1037–1048.

- Bear, H. D., Tang, G., Rastogi, P., Geyer Jr, C. E., Robidoux, A., Atkins, J. N., Baez-Diaz, L., Brufsky, A. M., Mehta, R. S., Fehrenbacher, L., et al. (2012). Bevacizumab added to neoadjuvant chemotherapy for breast cancer. *New England Journal of Medicine*, 366(4):310–320.
- Berger, M. L., Lipset, C., Gutteridge, A., Axelsen, K., Subedi, P., and Madigan, D. (2015). Optimizing the leveraging of real-world data to improve the development and use of medicines. *Value in Health*, 18(1):127–130.
- Beutel, A., Chen, J., Zhao, Z., and Chi, E. H. (2017). Data decisions and theoretical implications when adversarially learning fair representations. arXiv preprint arXiv:1707.00075.
- Bietti, A., Agarwal, A., and Langford, J. (2021). A contextual bandit bake-off. Journal of Machine Learning Research, 22(133):1–49.
- Boppana, S. and Goepfert, P. (2018). Understanding the CD8 T-cell response in natural HIV control. *F1000Research*, 7.
- Borenstein, M., Hedges, L. V., Higgins, J. P., and Rothstein, H. R. (2011). Introduction to Meta-Analysis. John Wiley & Sons.
- Breiman, L. (1996). Stacked regressions. *Machine Learning*, 24(1):49–64.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and Regression Trees.* CRC Press.
- Breitbart, Y., Olson, P. L., and Thompson, G. R. (1986). Database integration in a distributed heterogeneous database system. In 1986 IEEE Second International Conference on Data Engineering, pages 301–310. IEEE.
- Brown, J. S., Holmes, J. H., Shah, K., Hall, K., Lazarus, R., and Platt, R. (2010). Distributed health data networks: a practical and preferred approach to multi-institutional evaluations of comparative effectiveness, safety, and quality of care. *Medical Care*, 48(6):S45–S51.
- Buchanan, A. L., Hudgens, M. G., Cole, S. R., Mollan, K. R., Sax, P. E., Daar, E. S., Adimora, A. A., Eron, J. J., and Mugavero, M. J. (2018). Generalizing evidence from randomized trials using inverse probability of sampling weights. *Journal of the Royal Statistical Society. Series A*,(*Statistics in Society*), 181(4):1193.
- Cai, T., Tian, L., Wong, P. H., and Wei, L. (2011). Analysis of randomized comparative clinical trial data for personalized treatment selections. *Biostatistics*, 12(2):270–282.
- Chakraborty, B. and Moodie, E. E. (2013). *Statistical Methods for Dynamic Treatment Regimes*. Springer.

- Chakraborty, B., Murphy, S., and Strecher, V. (2010). Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research*, 19(3):317–343.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. (2002). Smote: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357.
- Chen, H., Lu, W., Song, R., and Ghosh, P. (2022). On learning and testing of counterfactual fairness through data preprocessing. *arXiv preprint arXiv:2202.12440*.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68.
- Cho, Y. J., Wang, J., Chiruvolu, T., and Joshi, G. (2021). Personalized federated learning for heterogeneous clients with clustered knowledge transfer. arXiv preprint arXiv:2109.08119.
- Chouldechova, A. and Roth, A. (2020). A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5):82–89.
- Clarke, B. (2003). Comparing bayes model averaging and stacking when model approximation error cannot be ignored. *Journal of Machine Learning Research*, 4(Oct):683–712.
- Cohen, J. A., Trojano, M., Mowry, E. M., Uitdehaag, B. M., Reingold, S. C., and Marrie, R. A. (2020). Leveraging real-world data to investigate multiple sclerosis disease behavior, prognosis, and treatment. *Multiple Sclerosis Journal*, 26(1):23–37.
- Colnet, B., Mayer, I., Chen, G., Dieng, A., Li, R., Varoquaux, G., Vert, J.-P., Josse, J., and Yang, S. (2020). Causal inference methods for combining randomized trials and observational studies: a review. arXiv preprint arXiv:2011.08047.
- Cook, T. D., Campbell, D. T., and Shadish, W. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Houghton Mifflin Boston, MA.
- Cortazar, P., Zhang, L., Untch, M., Mehta, K., Costantino, J. P., Wolmark, N., Bonnefoi, H., Cameron, D., Gianni, L., Valagussa, P., et al. (2014). Pathological complete response and long-term clinical benefit in breast cancer: the ctneobc pooled analysis. *The Lancet*, 384(9938):164–172.
- Creager, E., Madras, D., Jacobsen, J.-H., Weis, M., Swersky, K., Pitassi, T., and Zemel, R. (2019). Flexibly fair representation learning by disentanglement. In *International Conference on Machine Learning*, pages 1436–1445. PMLR.
- Dahabreh, I. J., Robertson, S. E., Tchetgen, E. J., Stuart, E. A., and Hernán, M. A. (2019). Generalizing causal inferences from individuals in randomized trials to all trial-eligible individuals. *Biometrics*, 75(2):685–694.

- Dai, D., Han, L., Yang, T., and Zhang, T. (2018). Bayesian model averaging with exponentiated least squares loss. *IEEE Transactions on Information Theory*, 64(5):3331–3345.
- Dai, D. and Zhang, T. (2011). Greedy model averaging. Advances in Neural Information Processing Systems, 24:1242–1250.
- Davison, A. C. and Hinkley, D. V. (1997). Bootstrap Methods and their Application. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Dayan, I., Roth, H. R., Zhong, A., Harouni, A., Gentili, A., Abidin, A. Z., Liu, A., Costa, A. B., Wood, B. J., Tsai, C.-S., et al. (2021). Federated learning for predicting clinical outcomes in patients with covid-19. *Nature Medicine*, 27(10):1735–1743.
- Degtiar, I. and Rose, S. (2021). A review of generalizability and transportability. arXiv preprint arXiv:2102.11904.
- DeShazo, J. P. and Hoffman, M. A. (2015). A comparison of a multistate inpatient ehr database to the hcup nationwide inpatient sample. *BMC Health Services Research*, 15(1):1–8.
- Ding, P., Geng, Z., Yan, W., and Zhou, X.-H. (2011). Identifiability and estimation of causal effects by principal stratification with outcomes truncated by death. *Journal of the American Statistical Association*, 106(496):1578–1591.
- Donahue, M., Bouhaddou, O., Hsing, N., Turner, T., Crandall, G., Nelson, J., and Nebeker, J. (2018). Veterans health information exchange: successes and challenges of nationwide interoperability. In AMIA Annual Symposium Proceedings, volume 2018, page 385. American Medical Informatics Association.
- Donohue, J. M., Jarlenski, M. P., Kim, J. Y., Tang, L., Ahrens, K., Allen, L., Austin, A., Barnes, A. J., Burns, M., Chang, C.-C. H., et al. (2021). Use of medications for treatment of opioid use disorder among us medicaid enrollees in 11 states, 2014-2018. *The Journal* of the American Medical Association, 326(2):154–164.
- Druckman, J. N., Greene, D. P., Kuklinski, J. H., and Lupia, A. (2011). *Cambridge handbook* of experimental political science. Cambridge University Press.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012). Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, pages 214–226.
- Egami, N. and Hartman, E. (2020). Elements of external validity: Framework, design, and analysis. *Design, and Analysis (June 30, 2020)*.
- Fallah, A., Mokhtari, A., and Ozdaglar, A. (2020). Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. Advances in Neural Information Processing Systems, 33:3557–3568.

- Fang, E. X., Wang, Z., and Wang, L. (2022). Fairness-oriented learning for optimal individualized treatment rules. *Journal of the American Statistical Association*, 0(0):1–14.
- FDA (2014). Guidance for industry. pathological complete response in neoadjuvant treatment of high-risk early-stage breast cancer: use as an endpoint to support accelerated approval.
- Feldman, M., Friedler, S. A., Moeller, J., Scheidegger, C., and Venkatasubramanian, S. (2015). Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 259–268.
- Ferreira, F. L., Bota, D. P., Bross, A., Mélot, C., and Vincent, J.-L. (2001). Serial evaluation of the sofa score to predict outcome in critically ill patients. *The Journal of the American Medical Association*, 286(14):1754–1758.
- Fleurence, R. L., Curtis, L. H., Califf, R. M., Platt, R., Selby, J. V., and Brown, J. S. (2014). Launching PCORnet, a national patient-centered clinical research network. *Journal of the American Medical Informatics Association*, 21(4):578–582.
- Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. Biometrics, 58(1):21–29.
- Gilbert, P. B., Bosch, R. J., and Hudgens, M. G. (2003). Sensitivity analysis for the assessment of causal vaccine effects on viral load in hiv vaccine trials. *Biometrics*, 59(3):531–541.
- Gilbert, P. B., Gabriel, E. E., Huang, Y., and Chan, I. S. (2015). Surrogate endpoint evaluation: Principal stratification criteria and the prentice definition. *Journal of causal inference*, 3(2):157–175.
- Gilbert, P. B. and Hudgens, M. G. (2008). Evaluating candidate principal surrogate endpoints. *Biometrics*, 64(4):1146–1154.
- Goldberg, Y. and Kosorok, M. R. (2012). Q-learning with censored data. Annals of Statistics, 40(1):529.
- Hahn, P. R., Murray, J. S., Carvalho, C. M., et al. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056.
- Hammer, S. M., Katzenstein, D. A., Hughes, M. D., Gundacker, H., Schooley, R. T., Haubrich, R. H., Henry, W. K., Lederman, M. M., Phair, J. P., Niu, M., et al. (1996). A trial comparing nucleoside monotherapy with combination therapy in hiv-infected adults with cd4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine*, 335(15):1081–1090.
- Hansen, B. E. (2007). Least squares model averaging. *Econometrica*, 75(4):1175–1189.

- Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., and Ramage, D. (2018). Federated learning for mobile keyboard prediction. arXiv preprint arXiv:1811.03604.
- Harton, J., Segal, B., Mamtani, R., Mitra, N., and Hubbard, R. (2021). Combining realworld and randomized control trial data using data-adaptive weighting via the on-trial score. arXiv preprint arXiv:2108.08756.
- Hashimoto, T., Srivastava, M., Namkoong, H., and Liang, P. (2018). Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, pages 1929–1938. PMLR.
- Hayashi, F. (2000). *Econometrics*. Princeton University Press.
- Henderson, R., Ansell, P., and Alshibani, D. (2010). Regret-regression for optimal dynamic treatment regimes. *Biometrics*, 66(4):1192–1201.
- Holland, P. W. (1986). Statistics and causal inference. Journal of the American statistical Association, 81(396):945–960.
- Howell, M. D., Donnino, M., Clardy, P., Talmor, D., and Shapiro, N. I. (2007). Occult hypoperfusion and mortality in patients with suspected infection. *Intensive Care Medicine*, 33(11):1892–1899.
- Hripcsak, G., Duke, J. D., Shah, N. H., Reich, C. G., Huser, V., Schuemie, M. J., Suchard, M. A., Park, R. W., Wong, I. C. K., Rijnbeek, P. R., et al. (2015). Observational health data sciences and informatics (ohdsi): opportunities for observational researchers. *Studies* in Health Technology and Informatics, 216:574.
- Huang, X., Choi, S., Wang, L., and Thall, P. F. (2015). Optimization of multi-stage dynamic treatment regimes utilizing accumulated data. *Statistics in Medicine*, 34(26):3424–3443.
- Imai, K. and Ratkovic, M. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470.
- Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.
- Jemiai, Y., Rotnitzky, A., Shepherd, B. E., and Gilbert, P. B. (2007). Semiparametric estimation of treatment effects given base-line covariates on an outcome measured after a post-randomization event occurs. *Journal of the Royal Statistical Society: Series B* (Statistical Methodology), 69(5):879–901.
- Jordan, M. I., Lee, J. D., and Yang, Y. (2019). Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681.
- Joseph, M., Kearns, M., Morgenstern, J. H., and Roth, A. (2016). Fairness in learning: Classic and contextual bandits. *Advances in Neural Information Processing Systems*, 29.

- Ju, C., Combs, M., Lendle, S. D., Franklin, J. M., Wyss, R., Schneeweiss, S., and van der Laan, M. J. (2019). Propensity score prediction for electronic healthcare databases using super learner and high-dimensional propensity score methods. *Journal of Applied Statistics*, 46(12):2216–2236.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2019). Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977.
- Kamiran, F. and Calders, T. (2012). Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33.
- Ke, Z. T., Fan, J., and Wu, Y. (2015). Homogeneity pursuit. Journal of the American Statistical Association, 110(509):175–194.
- Kern, H. L., Stuart, E. A., Hill, J., and Green, D. P. (2016). Assessing methods for generalizing experimental impact estimates to target populations. *Journal of Research on Educational Effectiveness*, 9(1):103–127.
- Kosorok, M. R. and Moodie, E. E. (2015). Adaptive treatment strategies in practice: planning trials and analyzing data for personalized medicine. SIAM.
- Krishna, U., Joshi, S. P., and Modh, M. (2009). An evaluation of serial blood lactate measurement as an early predictor of shock and its outcome in patients of trauma or sepsis. *Indian Journal of Critical Care Medicine*, 13(2):66.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165.
- Kurland, B. F., Johnson, L. L., Egleston, B. L., and Diehr, P. H. (2009). Longitudinal data with follow-up truncated by death: match the analysis method to research aims. *Statistical Science*, 24(2):211.
- Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E., and Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, 8(1):1225–1272.
- Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., Wang, X., and Chi, E. (2020). Fairness without demographics through adversarially reweighted learning. Advances in Neural Information Processing Systems, 33:728–740.
- LaLonde, R. J. (1986). Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, pages 604–620.
- Lee, J. D., Liu, Q., Sun, Y., and Taylor, J. E. (2017). Communication-efficient sparse regression. *The Journal of Machine Learning Research*, 18(1):115–144.

- Lee, K., Daniels, M. J., and Sargent, D. J. (2010). Causal effects of treatments for informative missing data due to progression/death. *Journal of the American Statistical Association*, 105(491):912–929.
- Li, L., Chu, W., Langford, J., and Wang, X. (2011). Unbiased offline evaluation of contextualbandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 297–306.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60.
- Li, Y., Taylor, J. M., and Elliott, M. R. (2010). A bayesian approach to surrogacy assessment using principal stratification in clinical trials. *Biometrics*, 66(2):523–531.
- Lin, D.-Y. and Zeng, D. (2010). On the relative efficiency of using summary statistics versus individual-level data in meta-analysis. *Biometrika*, 97(2):321–332.
- Liu, G., Tan, X., Dang, C., Tan, S., Xing, S., Huang, N., Peng, K., Xie, C., Tang, X., and Zeng, J. (2019). Regional shape abnormalities in thalamus and verbal memory impairment after subcortical infarction. *Neurorehabilitation and Neural Repair*, 33(6):476–485.
- Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30.
- Ma, S. and Huang, J. (2017). A concave pairwise fusion approach to subgroup analysis. Journal of the American Statistical Association, 112(517):410–423.
- Machicado, J. D., Gougol, A., Tan, X., Gao, X., Paragomi, P., Pothoulakis, I., Talukdar, R., Kochhar, R., Goenka, M. K., Gulla, A., et al. (2021). Mortality in acute pancreatitis with persistent organ failure is determined by the number, type, and sequence of organ systems affected. United European Gastroenterology Journal, 9(2):139–149.
- Makar, M., Packer, B., Moldovan, D., Blalock, D., Halpern, Y., and D'Amour, A. (2022). Causally motivated shortcut removal using auxiliary labels. In *International Conference* on Artificial Intelligence and Statistics, pages 739–766. PMLR.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246.
- Maro, J. C., Platt, R., Holmes, J. H., Strom, B. L., Hennessy, S., Lazarus, R., and Brown, J. S. (2009). Design of a national distributed health data network. *Annals of Internal Medicine*, 151(5):341–344.
- Masoudnia, S. and Ebrahimpour, R. (2014). Mixture of experts: a literature survey. Artificial Intelligence Review, 42(2):275–293.

- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35.
- Mo, S., Tan, X., Xia, J., and Ren, P. (2020). Towards improving spatiotemporal action recognition in videos. arXiv preprint arXiv:2012.08097.
- Mo, S., Xia, J., Tan, X., and Raj, B. (2021). Point3D: tracking actions as moving points with 3D CNNs. In 32nd British Machine Vision Conference 2021, BMVC 2021, Online, November 22-25, 2021, page 259. BMVA Press.
- Moodie, E. E., Platt, R. W., and Kramer, M. S. (2009). Estimating response-maximized decision rules with applications to breastfeeding. *Journal of the American Statistical Association*, 104(485):155–165.
- Moodie, E. E., Richardson, T. S., and Stephens, D. A. (2007). Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455.
- Murphy, S. A. (2003). Optimal dynamic treatment regimes. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 65(2):331–355.
- Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24(10):1455–1481.
- Murphy, S. A., van der Laan, M. J., Robins, J. M., and Group, C. P. P. R. (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423.
- Nabi, R., Malinsky, D., and Shpitser, I. (2019). Learning optimal fair policies. In International Conference on Machine Learning, pages 4674–4682. PMLR.
- Nabi, R. and Shpitser, I. (2018). Fair inference on outcomes. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 32.
- Neyman, J. (1923). Sur les applications de la thar des probabilities aux experiences agaricales: Essay des principle. Excerpts reprinted (1990) in English. *Statistical Science*, 5:463–472.
- Nie, X. and Wager, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319.
- Onyemekwu, C., Tan, X., Gamboa, J., Moghbeli, K., Potter, K., Prendergast, N., Tiberio, P., Billings, F., Ely, E., Hulgan, T., Pandharipande, P., Samuels, D., Tang, G., and Girard, T. (2022). Associations Between Plasma Cell-Free Mitochondrial DNA, Delirium and Coma During Sepsis, pages A3700–A3700. American Thoracic Society.

- Orellana, L., Rotnitzky, A., and Robins, J. M. (2010a). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part i: main content. *The International Journal of Biostatistics*, 6(2).
- Orellana, L., Rotnitzky, A., and Robins, J. M. (2010b). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part ii: proofs of results. *The International Journal of Biostatistics*, 6(2).
- Patil, V., Ghalme, G., Nair, V., and Narahari, Y. (2020). Achieving fairness in the stochastic multi-armed bandit problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5379–5386.
- Pearl, J. (2009). Causality. Cambridge University Press.
- Pearl, J. and Bareinboim, E. (2011). Transportability of causal and statistical relations: A formal approach. In *Twenty-fifth AAAI Conference on Artificial Intelligence*.
- Pearl, J. and Bareinboim, E. (2014). External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595.
- Perez, L. and Wang, J. (2017). The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621.
- Pessach, D. and Shmueli, E. (2022). A review on fairness in machine learning. ACM Computing Surveys (CSUR), 55(3):1–44.
- Pirracchio, R., Petersen, M. L., and van der Laan, M. (2015). Improving propensity score estimators' robustness to model misspecification using super learner. *American Journal of Epidemiology*, 181(2):108–119.
- Platt, R., Brown, J. S., Robb, M., McClellan, M., Ball, R., Nguyen, M. D., and Sherman, R. E. (2018). The FDA sentinel initiative—an evolving national resource. New England Journal of Medicine, 379(22):2091–2093.
- Pollard, T. J., Johnson, A. E., Raffa, J. D., Celi, L. A., Mark, R. G., and Badawi, O. (2018). The eICU collaborative research database, a freely available multi-center database for critical care research. *Scientific Data*, 5:180178.
- Polley, E. C. and van der Laan, M. J. (2010). Super learner in prediction. U.C. Berkeley Division of Biostatistics Working Paper Series.
- Qi, Z., Cui, Y., Liu, Y., and Pang, J.-S. (2019). Estimation of individualized decision rules based on an optimized covariate-dependent equivalent of random outcomes. SIAM Journal on Optimization, 29(3):2337–2362.
- Qi, Z., Pang, J.-S., and Liu, Y. (2022). On robustness of individualized decision rules. Journal of the American Statistical Association, pages 1–15.

- Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. Annals of Statistics, 39(2):1180.
- Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92(437):179–191.
- Raghupathi, W. and Raghupathi, V. (2014). Big data analytics in healthcare: promise and potential. *Health Information Science and Systems*, 2(1):1–10.
- Reynolds, M., Christian, J., Mack, C., Hall, N., and Dreyer, N. (2020). Leveraging real-world data for covid-19 research: challenges and opportunities. *Journal of Precision Medicine*, 6:1–6.
- Riley, R. D., Higgins, J. P., and Deeks, J. J. (2011). Interpretation of random effects metaanalyses. *BMJ*, 342.
- Robins, J., Orellana, L., and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*, 27(23):4678–4721.
- Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics*, pages 189–326. Springer.
- Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560.
- Robinson, P. M. (1988). Root-n-consistent semiparametric regression. *Econometrica: Jour*nal of the Econometric Society, pages 931–954.
- Röver, C. and Friede, T. (2020). Dynamically borrowing strength from another study through shrinkage estimation. *Statistical Methods in Medical Research*, 29(1):293–308.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The* Annals of Statistics, pages 34–58.
- Rubin, D. B. (1980). Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593.
- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. Journal of the American Statistical Association, 100(469):322–331.
- Sakr, Y., Jaschinski, U., Wittebole, X., Szakmany, T., Lipman, J., Namendys-Silva, S. A., Martin-Loeches, I., Leone, M., Lupu, M.-N., Vincent, J.-L., and ICON Investigators (2018). Sepsis in intensive care unit patients: worldwide data from the intensive care over nations audit. Open Forum Infectious Diseases, 5(12):ofy313.

- Sattigeri, P., Hoffman, S. C., Chenthamarakshan, V., and Varshney, K. R. (2019). Fairness gan: Generating datasets with fairness properties using a generative adversarial network. *IBM Journal of Research and Development*, 63(4/5):3–1.
- Seymour, C. W., Liu, V. X., Iwashyna, T. J., Brunkhorst, F. M., Rea, T. D., Scherag, A., Rubenfeld, G., Kahn, J. M., Shankar-Hari, M., Singer, M., et al. (2016). Assessment of clinical criteria for sepsis: for the third international consensus definitions for sepsis and septic shock (sepsis-3). The Journal of the American Medical Association, 315(8):762–774.
- Shankar-Hari, M., Phillips, G. S., Levy, M. L., Seymour, C. W., Liu, V. X., Deutschman, C. S., Angus, D. C., Rubenfeld, G. D., Singer, M., et al. (2016). Developing a new definition and assessing new clinical criteria for septic shock: for the third international consensus definitions for sepsis and septic shock (sepsis-3). The Journal of the American Medical Association, 315(8):775–787.
- Shepherd, B. E., Gilbert, P. B., Jemiai, Y., and Rotnitzky, A. (2006). Sensitivity analyses comparing outcomes only existing in a subset selected post-randomization, conditional on covariates, with application to hiv vaccine trials. *Biometrics*, 62(2):332–342.
- Shortreed, S. M., Cook, A. J., Coley, R. Y., Bobb, J. F., and Nelson, J. C. (2019). Challenges and opportunities for using big health care data to advance medical science and public health. *American Journal of Epidemiology*, 188(5):851–861.
- Smith, V., Chiang, C.-K., Sanjabi, M., and Talwalkar, A. S. (2017). Federated multi-task learning. Advances in Neural Information Processing Systems, 30.
- Song, N., Tan, X. E., Wang, Y., Kim, R. S., Bandos, H., Tang, G., Mamounas, E., Geyer, C. E., Rastogi, P., Jacobs, S. A., et al. (2021). Association of pCR and the 8-gene signature: NRG Oncology/NSABP B-41. *Cancer Research*, 81(13_Supplement):532–532.
- Song, R., Wang, W., Zeng, D., and Kosorok, M. R. (2015). Penalized q-learning for dynamic treatment regimens. *Statistica Sinica*, 25(3):901.
- Splawa-Neyman, J., Dabrowska, D. M., and Speed, T. (1990). On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, pages 465–472.
- Stuart, E. A., Ackerman, B., and Westreich, D. (2018). Generalizability of randomized trial results to target populations: design and analysis possibilities. *Research on Social Work Practice*, 28(5):532–537.
- Stuart, E. A., Bradshaw, C. P., and Leaf, P. J. (2015). Assessing the generalizability of randomized trial results to target populations. *Prevention Science*, 16(3):475–485.
- Stuart, E. A., Cole, S. R., Bradshaw, C. P., and Leaf, P. J. (2011). The use of propensity scores to assess the generalizability of results from randomized trials. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 174(2):369–386.
- Sutton, A. J., Abrams, K. R., Jones, D. R., Jones, D. R., Sheldon, T. A., and Song, F. (2000). *Methods for meta-analysis in medical research*, volume 348. Wiley Chichester.
- Sverdrup, E., Kanodia, A., Zhou, Z., Athey, S., and Wager, S. (2020). policytree: Policy learning via doubly robust empirical welfare maximization over trees. *Journal of Open Source Software*, 5(50):2232.
- Tan, X., Abberbock, J., Rastogi, P., and Tang, G. (2022a). Identifying principal stratum causal effects conditional on a post-treatment intermediate response. In Schölkopf, B., Uhler, C., and Zhang, K., editors, *Proceedings of the First Conference on Causal Learning* and Reasoning, volume 177 of Proceedings of Machine Learning Research, pages 734–753. PMLR.
- Tan, X., Chang, C.-C. H., Zhou, L., and Tang, L. (2022b). A tree-based model averaging approach for personalized treatment effect estimation from heterogeneous data sources. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 21013–21036. PMLR.
- Tan, X., Qi, Z., Seymour, C. W., and Tang, L. (2022c). RISE: Robust individualized decision learning with sensitive variables. *ICML 2022 Workshop on Responsible Decision Making* in Dynamic Environments.
- Tan, X., Ross, C. A., Miller, M., and Tang, X. (2018). Changepoint analysis of putamen and thalamus subregions in premanifest huntington's disease. In 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), pages 531–535.
- Tan, X., Yang, S., Ye, W., Faries, D. E., Lipkovich, I., and Kadziola, Z. (2022d). When doubly robust methods meet machine learning for estimating treatment effects from realworld data: A comparative study. arXiv preprint arXiv:2204.10969.
- Tang, L., Zhou, L., and Song, P. X.-K. (2020a). Distributed simultaneous inference in generalized linear models via confidence distribution. *Journal of Multivariate Analysis*, 176:104567.
- Tang, X., Xue, F., and Qu, A. (2020b). Individualized multidirectional variable selection. Journal of the American Statistical Association, pages 1–17.
- Tao, Y. and Wang, L. (2017). Adaptive contrast weighted learning for multi-stage multitreatment decision-making. *Biometrics*, 73(1):145–155.
- Tchetgen Tchetgen, E. J. (2014). Identification and estimation of survivor average causal effects. *Statistics in Medicine*, 33(21):3601–3628.
- Thall, P. F., Sung, H.-G., and Estey, E. H. (2002). Selecting therapeutic strategies based on efficacy and death in multicourse clinical trials. *Journal of the American Statistical Association*, 97(457):29–39.

- Tipton, E. and Olsen, R. B. (2018). A review of statistical methods for generalizing from evaluations of educational interventions. *Educational Researcher*, 47(8):516–524.
- Toh, S., Platt, R., Steiner, J., and Brown, J. (2011). Comparative-effectiveness research in distributed health data networks. *Clinical Pharmacology & Therapeutics*, 90(6):883–887.
- van den Boom, W., Hoy, M., Sankaran, J., Liu, M., Chahed, H., Feng, M., and See, K. C. (2020). The search for optimal oxygen saturation targets in critically ill patients: Observational data from large ICU databases. *Chest*, 157(3):566–573.
- van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1).
- van Dyk, D. A. and Meng, X.-L. (2001). The art of data augmentation. *Journal of Computational and Graphical Statistics*, 10(1):1–50.
- Varshney, K. R. (2016). Engineering safety in machine learning. In 2016 Information Theory and Applications Workshop (ITA), pages 1–5. IEEE.
- Vincent, J.-L., Moreno, R., Takala, J., Willatts, S., De Mendonça, A., Bruining, H., Reinhart, C., Suter, P., and Thijs, L. G. (1996). The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. On behalf of the working group on sepsis-related problems of the european society of intensive care medicine. *Intensive Care Medicine*, 22(7):707–710.
- Von Minckwitz, G., Untch, M., Blohmer, J.-U., Costa, S. D., Eidtmann, H., Fasching, P. A., Gerber, B., Eiermann, W., Hilfrich, J., Huober, J., et al. (2012). Definition and impact of pathologic complete response on prognosis after neoadjuvant chemotherapy in various intrinsic breast cancer subtypes. *Journal of Clinical Oncology*, 30(15):1796–1804.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.
- Wang, L. and Tchetgen Tchetgen, E. (2018). Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 80(3):531–550.
- Wang, L., Zhou, X.-H., and Richardson, T. S. (2017). Identification and estimation of causal effects with outcomes truncated by death. *Biometrika*, 104(3):597–612.
- Wang, L., Zhou, Y., Song, R., and Sherwood, B. (2018a). Quantile-optimal treatment regimes. *Journal of the American Statistical Association*, 113(523):1243–1254.
- Wang, Y., Fu, H., and Zeng, D. (2018b). Learning optimal personalized treatment rules in consideration of benefit and risk: with an application to treating type 2 diabetes patients with insulin therapies. *Journal of the American Statistical Association*, 113(521):1–13.

- Wang, Y., Jang, B., and Hero, A. (2020). The sylvester graphical lasso (syglasso). In International Conference on Artificial Intelligence and Statistics, pages 1943–1953. PMLR.
- Wang, Y., Klein, N., Morley, S., Jordanova, V., Henderson, M., Biswas, A., and Lawrence, E. (2021). Tributarypca: Distributed, streaming pca for in situ dimension reduction with application to space weather simulations. In 2021 7th International Workshop on Data Analysis and Reduction for Big Scientific Data (DRBSD-7), pages 33–39. IEEE.
- Wasserman, L. (2000). Bayesian model selection and model averaging. Journal of Mathematical Psychology, 44(1):92–107.
- Watkins, C. J. and Dayan, P. (1992). Q-learning. Machine Learning, 8(3):279–292.
- Whitehead, A. (2002). *Meta-analysis of controlled clinical trials*, volume 7. John Wiley & Sons.
- Wolfson, J. and Gilbert, P. (2010). Statistical identifiability and the surrogate endpoint problem, with application to vaccine trials. *Biometrics*, 66(4):1153–1161.
- Wyss, R., Schneeweiss, S., van der Laan, M., Lendle, S. D., Ju, C., and Franklin, J. M. (2018). Using super learner prediction modeling to improve high-dimensional propensity score estimation. *Epidemiology*, 29(1):96–106.
- Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2):1–19.
- Yang, S., Zeng, D., and Wang, X. (2020). Elastic integrative analysis of randomized trial and real-world data for treatment heterogeneity estimation. arXiv preprint arXiv:2005.10579.
- Yang, Y. (2001). Adaptive regression by mixing. Journal of the American Statistical Association, 96(454):574–588.
- Yao, Y., Vehtari, A., Simpson, D., and Gelman, A. (2018). Using stacking to average bayesian predictive distributions (with discussion). *Bayesian Analysis*, 13(3):917–1007.
- Zafar, M. B., Valera, I., Rogriguez, M. G., and Gummadi, K. P. (2017). Fairness constraints: Mechanisms for fair classification. In Artificial Intelligence and Statistics, pages 962–970. PMLR.
- Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018.
- Zhang, B. H., Lemoine, B., and Mitchell, M. (2018). Mitigating unwanted biases with adversarial learning. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, pages 335–340.
- Zhang, J. and Bareinboim, E. (2018). Fairness in decision-making—the causal explanation formula. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

- Zhang, J. L. and Rubin, D. B. (2003). Estimation of causal effects via principal stratification when some outcomes are truncated by "death". *Journal of Educational and Behavioral Statistics*, 28(4):353–368.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. (2018). Federated learning with non-iid data. arXiv preprint arXiv:1806.00582.
- Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118.
- Zhao, Y.-Q., Zeng, D., Laber, E. B., Song, R., Yuan, M., and Kosorok, M. R. (2015). Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*, 102(1):151–168.
- Zigler, C. M. and Belin, T. R. (2012). A bayesian approach to improved estimation of causal effect predictiveness for a principal surrogate endpoint. *Biometrics*, 68(3):922–932.