

# 000 001 002 003 004 005 IMPROVING CAUSAL INFERENCE ROBUSTNESS VIA 006 REINFORCEMENT-GUIDED DIFFUSION MODELS 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026

## ABSTRACT

011 Estimating the Conditional Average Treatment Effect (CATE) is essential to per-  
012 sonalized decision-making in causal inference. However, in real-world prac-  
013 tices, CATE models often suffer degraded performance when faced with unknown  
014 distribution shifts between training and deployment environments. To tackle  
015 this challenge, we introduce **Causal Adversarial Reinforcement-guided Diffusion**  
016 (**CARD**), a model-agnostic framework that can be wrapped around any exist-  
017 ing CATE learner to improve its robustness against unknown distribution shifts.  
018 CARD formulates the CATE modeling process as a minimax game: a reinfor-  
019 cements learning agent guides a diffusion model to generate adversarial data aug-  
020 mentations that maximize the CATE learner’s loss, and then the learner is trained  
021 to minimize this worst-case loss, creating a principled robust optimization pro-  
022 cedure. The comprehensive experimental results demonstrate that CARD con-  
023 sistently improves the robustness of diverse CATE learners against challenging  
024 data corruptions, including measurement error, missing values, and unmeasured  
025 confounding, confirming its broad applicability and effectiveness.  
026

## 1 INTRODUCTION

027 Estimating the Conditional Average Treatment Effect (CATE) is a core problem in causal infe-  
028 rence, as it quantifies how an intervention would differentially affect subgroup-level (or approxi-  
029 mated individual-level) as a function of observed covariates, enabling personalized decision-making  
030 in various domains such as (Farrell, 2015; Chernozhukov et al., 2018; Kitagawa & Tetenov, 2018;  
031 Abadie et al., 2023), statistics (Wager & Athey, 2018; Li & Wager, 2022; Foster & Syrgkanis, 2023;  
032 Kennedy, 2023), clinical (Zhang et al., 2019; Qian et al., 2021; Bica et al., 2021; Kinyanui & Jo-  
033 hansson, 2022; Feuerriegel et al., 2024; Ma et al., 2025), and financial application (Li et al., 2023;  
034 Huang et al., 2023b; Fernández-Loría et al., 2023; Wu et al., 2025a). The practical of CATE mod-  
035 els often hinges on a crucial but fragile assumption: *external validity (or transportability)* (Pearl &  
036 Bareinboim, 2011; Bareinboim & Pearl, 2016). That is, causal conclusions derived from a source  
037 environment must remain valid when the model is deployed in a different target population.  
038

039 However, in real-world practice, this external validity assumption is often violated. CATE mod-  
040 els can suffer a degraded performance when confronted when unknown *distribution shifts* between  
041 the training and deployment environments present, which is often incurred by data imperfections  
042 (Kallus et al., 2018; Agarwal & Singh, 2021; Zhang et al., 2023), such as measurement error (Imai  
043 & Yamamoto, 2010; Battistin & Chesher, 2014; Kuroki & Pearl, 2014; Pei et al., 2019), missing val-  
044 ues (Rubin, 1976; Bang & Robins, 2005; Mohan et al., 2013; Yang et al., 2019; Mayer et al., 2020),  
045 and unmeasured confounding (Kallus et al., 2019; Ding et al., 2022; Oprescu et al., 2023; Xiao  
046 et al., 2024; Dorn et al., 2025). Such data imperfections signify structural discrepancies between the  
047 source and target domains, thereby rendering the direct transfer of causal conclusions invalid. We  
048 demonstrate this challenge with the following motivating example.  
049

050 **Motivating example.** Suppose a technical company trains a CATE model on its proprietary clean  
051 and well structured dataset, and intend to license it to a hospital system. The model must be val-  
052 idated on the hospital’s own Electronic Health Record (EHR) data, which constitutes an unseen  
053 target domain as it was unavailable during the initial training phase. The EHR data reflects a dif-  
ferent data generating process with systematic imperfections that induce a distribution shift from

054 the source data (Ruan et al., 2024): (i) measurement error, arising from device inaccuracies, re-  
 055 porting biases, or procedural variability; (ii) missing values, caused by privacy restrictions, legal  
 056 constraints, or transcription errors; and (iii) unmeasured confounders, such as socioeconomic sta-  
 057 tus, environmental exposures, or genetic predispositions, can trigger severe concept drift. These  
 058 uncertain data imperfections significantly hinder the generalization of the CATE model trained in  
 059 the source domain to the EHR data.

060 To tackle the distribution shift problem in causal inference, existing methods develop specialized  
 061 causal estimators for specific data shift types. Examples include CATE learners that aim to be robust  
 062 to covariate shift (Kern et al., 2024) or to concept drift (Zhang et al., 2024), and policy learning  
 063 methods that seek robustness under combined and unknown shifts (Kallus et al., 2022; Mu et al.,  
 064 2022; Si et al., 2023). Most of these studies are grounded in structural assumptions about the shift  
 065 or the estimator itself, and their advancements in robust causal learning motivates a complementary  
 066 question: *Can we develop a method that is capable to enhance the robustness of any existing CATE*  
 067 *learner to unknown distribution shifts without requiring additional structural assumptions?*

068 Inspired by this question, we propose **Causal Adversarial Reinforcement-guided Diffusion (CARD)**,  
 069 a model-agnostic framework that strengthens the robustness of *any* existing CATE learner without  
 070 redesigning its internal architecture. CARD frames robust CATE training as a minimax game be-  
 071 tween an adversarial generator and a CATE learner. In CARD, specifically, a reinforcement learning  
 072 agent guides a diffusion model to produce adversarial proxies that maximally challenge the CATE  
 073 learner, and the learner then adapts by minimizing the worst-case error over these generated aug-  
 074 mentations, yielding a principled robust optimization routine tailored to CATE estimation.

075 Our main **contributions** are threefold:

- 077 • We propose a novel model-agnostic framework, CARD, which can be flexibly integrated  
 078 with any existing CATE learner to enhance its robustness against a wide range of unknown  
 079 distribution shifts, without requiring additional structural assumptions or prior knowledge  
 080 of target information.
- 081 • To the best of our knowledge, we are the first to introduce a reinforcement-learning guided  
 082 diffusion model in causal inference literature. This might bring new possibilities for other  
 083 causal inference tasks, such as counterfactual generation (Yoon et al., 2018), dimension  
 084 reduction (Liu et al., 2024), and model evaluation (Athey et al., 2024), among others.
- 085 • We empirically demonstrate that CARD consistently improves the robustness of popular  
 086 CATE learners when deployed in challenging target data corruptions, involving measure-  
 087 ment error, missing values, and unmeasured confounding, confirming its reliability and  
 088 adaptability to real-world causal inference tasks.

## 089 2 RELATED WORK

090 **Data combination and external validity.** A central challenge in causal inference is to generalize  
 091 effects learned in the source dataset to a target population. This problem is formalized under ex-  
 092 ternal validity (or transportability) (Pearl & Bareinboim, 2011; Bareinboim & Pearl, 2016). Data  
 093 combination frameworks specify when and how evidence from multiple sources can be fused across  
 094 populations to identify causal quantities, explicitizing the role of distributional differences across  
 095 domains (Bareinboim & Pearl, 2016; Dahabreh & Hernán, 2019). When the transportability as-  
 096 sumption holds, combining data can improve the precision and efficiency of treatment effect estima-  
 097 tion (Hatt et al., 2022; Dahabreh et al., 2023; Huang et al., 2023a; Wu et al., 2025b; Rudolph et al.,  
 098 2025). In practice, however, the transportability assumption is often violated due to unobserved  
 099 heterogeneity between the source and target domains. To address this, statistical work develops sen-  
 100 sitivity analysis and partial identification tools for average treatment effects (Nie et al., 2021; Huang,  
 101 2024; Yadlowsky et al., 2022). On the modeling side, there is growing interest in conditional effect  
 102 estimation without assuming transportability. Several studies address latent confounding by com-  
 103 bining RCT and observational data, proposing methods such as the integrative R-learner (Wu &  
 104 Yang, 2022) and the MetaDebias neural network (Xiao et al., 2024).

105 **Causal inference under distribution shift.** A growing body of research has examined how to  
 106 make causal inference methods robust when the deployment distribution differs from the training

108 environment. Existing approaches can be broadly categorized into two lines. The first line focuses  
 109 on robust CATE estimation under specific types of shifts. For example, recent work addresses  
 110 covariate shift by controlling worst-case bias across target covariate distributions (Jeong & Namkoong,  
 111 2020) or by enforcing multi-accuracy constraints on CATE learners (Kern et al., 2024). Other studies  
 112 primarily focus on concept drift, for instance, by optimizing the CATE function within an un-  
 113 certainty set over convex combinations of multisite CATE functions under a known target covariate  
 114 distribution (Zhang et al., 2024). The second line of work emphasizes causal decision making via  
 115 robust optimization, which aims to learn treatment assignment rules that remain effective to unseen  
 116 confounding scenarios or target environments (Kallus & Zhou, 2021; Kallus et al., 2022; Mu et al.,  
 117 2022; Kido, 2022; Si et al., 2023; Shen et al.; Wang et al.; Hess et al., 2025). While powerful for de-  
 118 riving robust policies, these approaches are primarily designed for policy learning rather than CATE  
 119 estimation. This gap underscores the necessity of developing methods specifically tailored for gen-  
 120 eralizing CATE estimation to unseen target domains without requiring prior knowledge of covariate  
 121 distributions or treatment effect heterogeneity.

### 3 PROBLEM SETUP

124 This study is grounded in the potential outcome framework (Rubin, 1974; 2005). Let  
 125  $\{(X_i, A_i, Y_i)\}_{i=1}^n$  denote an observational sample of  $n$  i.i.d. units drawn from a *source* popula-  
 126 tion. For unit  $i$ ,  $X_i \in \mathcal{X} \subset \mathbb{R}^d$  is a  $d$ -dimensional pre-treatment covariate vector,  $A_i \in \{0, 1\}$  is a  
 127 binary treatment indicator, and  $\{Y_i^0, Y_i^1\}$  are the corresponding potential outcomes. The observed  
 128 (factual) outcome is  $Y_i = Y_i^{A_i}$ , and the unobserved (counterfactual) outcome is  $Y_i^{1-A_i}$ . Our target  
 129 estimand is the CATE, which captures the sub-population treatment heterogeneity:

$$\tau(x) := \mathbb{E}[Y^1 - Y^0 | X = x]. \quad (1)$$

130 Estimating  $\tau(x)$  from observational data presents a key challenge, due to the fundamental problem  
 131 of causal inference: for any unit, only one potential outcome can be observed. To identify the CATE  
 132 from observational data in the source domain, we rely on the following standard assumptions.

133 **Assumption 1** (SUTVA, Consistency, and Overlap). *For all units in the source domain, we have the*  
 134 *following assumptions: Consistency & SUTVA: The observed outcome for unit  $i$  receiving treatment*  
 135  *$a$  is the potential outcome  $Y^a$ , and potential outcomes of this unit are not affected by the treatment*  
 136 *assignments of other units. Overlap (Positivity): The probability of receiving treatment is bounded*  
 137 *away from 0 and 1 for all covariate profiles, i.e.,  $0 < P(A = 1 | X = x) < 1$  for all  $x \in \mathcal{X}$ . Internal*  
 138 *validity (Unconfoundedness): The treatment assignment is independent of the potential outcomes,*  
 139 *conditional on the observed covariates, i.e.,  $\{Y^1, Y^0\} \perp\!\!\!\perp A | X$ .*

#### 3.1 GENERALIZING CATE UNDER DISTRIBUTION SHIFT

140 A critical generalization challenge arises when an estimator  $\hat{\tau}(x)$ , trained on the source domain  
 141 ( $P_S$ ), must be deployed in an unseen target domain ( $P_T$ ). The CATE model’s performance in this  
 142 new domain is threatened by potential distribution shifts in two main types.

143 **Covariate shift.** The most common and well-studied type of distribution shift is covariate shift,  
 144 where the marginal distribution of covariates differs across domains, i.e.,  $P_S(X) \neq P_T(X)$ . Generalization  
 145 under this shift is made possible by the transportability assumption.

146 **Assumption 2** (External validity (Transportability)). *The conditional distribution of potential out-  
 147 comes given covariates is invariant across domains, i.e.,  $P_S(Y^a | X) = P_T(Y^a | X)$  for  $a \in \{0, 1\}$ .*

148 This assumption implies that the underlying causal mechanisms are stable across domains, and thus  
 149 the true CATE function is the same in both domains:  $\tau_S(x) = \tau_T(x)$ . Nevertheless, even with  
 150 transportability, CATE model performance can deteriorate when deploying  $\hat{\tau}$  in the target domain  
 151 due to the covariate distribution mismatch, which is a common issue in machine learning studies.

152 **Concept drift.** A more severe challenge arises from concept drift, where the transportability as-  
 153 sumption is violated, meaning  $P_S(Y^a | X) \neq P_T(Y^a | X)$ . As highlighted in our motivating exam-  
 154 ple, such drift is often caused by unmeasured confounders present only in the target domain, which  
 155 alter the treatment heterogeneity. Under concept drift, the true CATE function is no longer invariant  
 156 across domains, i.e.,  $\tau_S(x) \neq \tau_T(x)$ , making out-of-domain CATE estimation substantially more  
 157 challenging than the covariate-shift-only setting.

162 **4 METHOD**  
 163

164 In this section, we introduce our proposed framework, **Causal Adversarial Reinforcement-guided**  
 165 **Diffusion (CARD)**. We begin by formulating robust CATE estimation as a minimax optimization  
 166 problem. We then detail how a reinforcement learning agent guides a diffusion model to generate  
 167 adversarial proxies that realize this objective. Finally, we present the detailed training pipeline for  
 168 integrating CARD with any CATE estimator.  
 169

170 **4.1 A MINIMAX OBJECTIVE FOR ROBUST CATE ESTIMATION**  
 171

172 Given source data samples  $(X, A, Y) \sim P_S$ , a standard CATE learner  $f_\phi$  with parameters  $\phi$  is  
 173 trained by minimizing an objective  $\mathcal{L}^{\text{inf}}(\phi) = \mathbb{E}[\ell(X, A, Y; f_\phi)]$ , where  $\ell$  is the loss function asso-  
 174 ciated with the chosen meta-learner. To protect the estimator against unknown distribution shifts in  
 175 the target domain, we optimize  $f_\phi$  with a new objective  $\mathcal{L}^{\text{inf}}(\phi, Z)$  in a robust optimization manner:  
 176

$$\min_{\phi} \max_{Z \in \Omega} \mathcal{L}^{\text{inf}}(\phi, Z) := \mathbb{E}_{(X, A, Y) \sim P_S} [\ell(X \oplus Z, A, Y; f_\phi)], \quad (2)$$

178 where  $\Omega$  is an uncertainty set defining the space of learnable adversarial proxies. Conceptually,  
 179 solving this objective forces the inferencer  $f_\phi$  to be robust against the most harmful proxies in  
 180  $\Omega$ . While some causal inference studies formulate similar adversarial problems as distributionally  
 181 robust optimization (DRO), they often define  $\Omega$  based on statistical distances or strong structural  
 182 assumptions, as discussed in Section 2. Our key departure is that the inferencer  $f_\phi$  is trained on  
 183 covariates augmented by learned adversarial proxies, creating a more flexible robustness mechanism  
 184 without loss of original covariate information.  
 185

186 **4.2 ROBUST CATE ESTIMATION WITH CARD**  
 187

188 Instead of constraining the adversary to a predefined uncertainty set  $\Omega$  (e.g., a KL-ball (Kallus et al.,  
 189 2022; Si et al., 2023)), we design a framework that learns to generate worst-case proxies  $Z$  using a  
 190 score-based diffusion model guided by a reinforcement learning (RL) agent.

191 **Score-based diffusion.** A score-based diffusion model (Song et al.) consists of a forward process  
 192 that progressively adds noise to data  $Z_0$  over a time interval  $t \in [0, T]$ , governed by a stochastic  
 193 differential equation (SDE):  
 194

$$dZ = f(Z, t)dt + g(t)dW_t, \quad (3)$$

195 where  $f(Z, t)$  is the drift coefficient,  $g(t)$  is the diffusion coefficient, and  $W_t$  is a standard Wiener  
 196 process. The corresponding reverse process generates data by traversing time from  $T$  to 0. Using  
 197 the Fokker-Planck equation of the marginal density (Suh et al.), the reverse-time SDE is:  
 198

$$dZ = [f(Z, t) - g(t)^2 \nabla_z \log p_t(Z)] dt + g(t)d\bar{W}_t, \quad (4)$$

200 where  $\bar{W}_t$  is a Wiener process running backward from  $t = T$  to  $t = 0$ . The score  $\nabla_z \log p_t(Z)$  is  
 201 approximated by a neural network  $g_\theta(z, t)$ , pretrained with the standard score-matching objective:  
 202

$$\mathcal{L}^{\text{diff}}(\theta) = \mathbb{E}_{Z_0, Z_t \sim p_t(\cdot | Z_0), t \sim \mathcal{U}[\varepsilon, T]} [\lambda(t)^2 \|g_\theta(Z_t, t) - \nabla_z \log p_t(Z_t | Z_0)\|_2^2], \quad (5)$$

204 where  $\lambda(t) > 0$  weights time steps and  $\varepsilon > 0$  ensures numerical stability. In practice, the diffusion  
 205 model is often applied on a low-dimensional latent code  $Z$  obtained from an autoencoder.  
 206

207 **Reinforcement-guided adversarial generation.** To solve the inner maximization of Eqn. (2), we  
 208 frame the reverse diffusion process as a Markov Decision Process (MDP) (Black et al.) and use  
 209 an RL agent to steer the generation toward adversarial proxies. The objective is to guide the score  
 210 model  $g_\theta$  to maximize the expected cumulative reward along the denoising trajectory:  
 211

$$\mathcal{J}(\theta) = \mathbb{E} \left[ \sum_{t=1}^T \log g_\theta(Z_{t-1} | Z_t) \frac{G_t - \mu_G}{\sigma_G} \right], \quad \text{where } G_t = \sum_{k=t}^T \gamma^{k-t} \mathcal{L}^{\text{inf}}(\phi, Z). \quad (6)$$

212 Here,  $\mathcal{L}^{\text{inf}}(\phi, Z)$  is the immediate reward at step  $t$ ,  $G_t$  is the discounted return from step  $t$  with  
 213 discount factor  $\gamma \in (0, 1]$ , and the returns are standardized per-trajectory with the mean  $\mu_G$  and  
 214

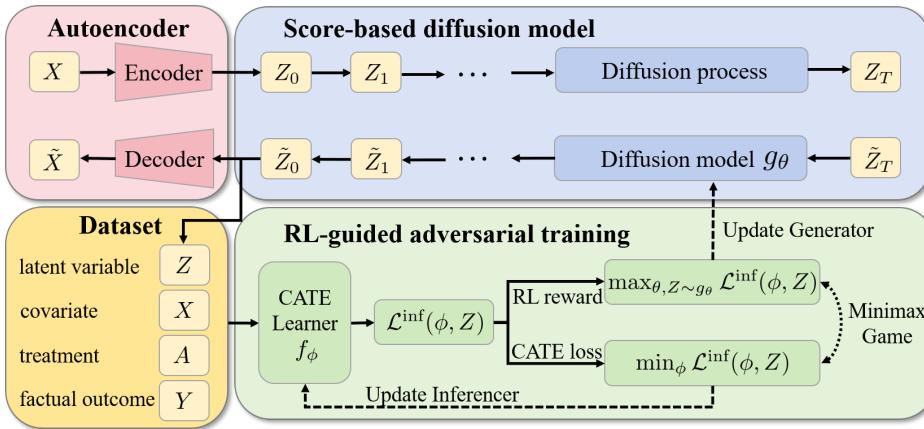


Figure 1: An overview of the proposed CARD training pipeline for robust CATE estimation.

standard deviation  $\sigma_G$ . To maintain generation quality and stabilize training, we combine this RL objective with the original score-matching loss, yielding the final objective for the generator:

$$\max_{\theta} \mathcal{L}^{\text{full}}(\theta) = \mathcal{J}(\theta) - \alpha \mathcal{L}^{\text{diff}}(\theta), \quad (7)$$

where  $\alpha > 0$  is a balancing hyperparameter. This fine-tuning process transforms the diffusion model from a simple data generator into a sophisticated adversary capable of generating worst-case proxies that approximate the solution to the inner maximization in our minimax objective (2).

**Algorithm of CATE learning with CARD.** The complete procedure of training CATE with CARD, which alternates between updating the generator and the inferencer, is outlined in Algorithm 1. We also illustrate the corresponding pipeline in Figure 1.

---

**Algorithm 1** CATE model training with CARD

---

**Require:** Source data  $(X, A, Y) \sim P_S$ , inferencer (base CATE learner)  $f_\phi$ , Autoencoder  $(\text{Enc}_\psi, \text{Dec}_\psi)$ , diffusion model  $g_\theta$ .

1: **Phase 1: Pre-training**

- 2: Train autoencoder on covariates  $X$  to learn a latent space  $Z$ .
- 3: Pre-train diffusion model  $g_\theta$  on latent representations  $Z = \text{Enc}_\psi(X)$  via Eqn. (5).
- 4: **Phase 2: CATE model training with minimax**

5: **for** each training epoch  $e = 1, \dots, E$  **do**

- 6: Sample  $Z_T \sim \mathcal{N}(0, I)$  and generate a denoising trajectory  $(Z_T, \dots, Z_0)$  using  $g_\theta$ .
- 7: For each sample, compute trajectory returns  $G_t$  via Eqn. (6).
- 8: Update generator parameters  $\theta$  by maximizing  $\mathcal{L}^{\text{full}}(\theta)$  from Eqn. (7).
- 9: Update inferencer parameters  $\phi$  by minimizing  $\mathcal{L}^{\text{inf}}(\phi, Z_0)$  from Eqn. (2).

**Ensure:** Generator  $g_\theta$  for  $Z_0$  generation, and inferencer  $f_\phi(X \oplus Z_0, A, Y)$  for CATE estimation.

---

## 5 EXPERIMENTS

### 5.1 EXPERIMENTAL SETUP

**CATE learners.** Our evaluation includes eight prominent CATE estimation methods, comprising five meta-learners and three specialized neural network models. The meta-learners represent a diverse set of strategies, including indirect-type (S-learner and T-Learner) and direct-type (X-learner (Künzel et al., 2019), DR-learner (Kennedy, 2023; Foster & Syrgkanis, 2023), and R-Learners (Nie & Wager, 2021)) approaches. The representation-based models consist of widely-recognized architectures: TARNNet & CFR-Wass (Shalit et al., 2017; Johansson et al., 2022) and DragonNet (Shi et al., 2019). The specific details of implementing these CATE learners with CARD are presented in Section A.1, and detailed parameter configurations for all models are provided in Appendix A.2.

Evaluating CATE estimators requires access to ground-truth treatment effects, which are unavailable in real-world data. Therefore, following established practice in causal inference research (Curth & Van der Schaar, 2021; Curth & Van Der Schaar, 2023; Huang et al., 2024), we employ a semi-synthetic data generating process with covariates collected from ACIC2016 dataset (Dorie et al., 2019). The dataset contains 4802 samples with  $d = 22$  continuous covariates. The treatment assignment  $A_i$  is generated from a Bernoulli distribution based on the covariates  $A_i|X_i \sim \text{Bern}(1/(1 + \exp(-(\beta'_T X_i))))$ . The potential outcome generation is based on additive interaction terms, with a complex quadratic heterogeneous treatment effects:

$$Y_i = \sum_j^d \beta'_j X_{i;j} + \sum_{j=1}^d \sum_{k=j}^d \beta'_{j,k} X_{i;j} X_{i;k} + A_i \sum_{j=1}^d \sum_{k=j}^d \gamma_{i,j} X_{i;j} X_{i;k} + \epsilon_i. \quad (8)$$

The coefficients are set as:  $\beta'_T \sim \text{Bern}(0.1)$ ,  $\beta'_j \sim \text{Bern}(0.5)$ ,  $\beta'_{j,k} \sim \text{Bern}(0.5)$ ,  $\gamma_{i,j} \sim \text{Bern}(0.1)$ , and the noise term  $\epsilon_i$  is sampled from  $\mathcal{N}(0, 0.1)$ . We repeat the above data generating process to generate 30 distinct datasets, each partitioned into training/validation/testing ratio of 49%/21%/30%.

**Distribution shift settings.** To evaluate the robustness of CARD, we introduce three types of distribution shifts exclusively in the test set:

- **Measurement error:** We simulate measurement error by adding Gaussian noise to the covariates of the test set, while the underlying data generating process remains unchanged. The observed covariates become  $X_i^{obs} = X_i + \mathcal{N}(0, \delta^2 I_d)$ , where  $\delta$  controls the shift level incurred by measurement error and varies across  $\{0.1, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$ .
- **Missing values:** This setting introduces missingness to the covariates of the test set, while the underlying data generating process remains unchanged. We apply a binary mask to the covariates, where each element is independently set to 0 with probability  $\rho$ . We use the MICE algorithm (Kallus et al., 2018) to enable CATE to be deployed on incomplete test data. The missingness rate  $\rho$  controls the shift level incurred by missing values and varies across  $\{0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$ .
- **Unmeasured confounding:** This scenario introduces a shift in the outcome generation mechanism for the test set. The observed covariates  $X_{obs}$  remain unchanged, but the potential outcomes are generated by a new model that includes both observed confounders  $X_{obs}$  and hidden confounders  $U$ , i.e.,  $X = (X_{obs}, U)$  in Eqn. (8). The unmeasured confounders are drawn from uniform distribution  $U_{i,j} \sim \mathcal{U}(-3, 3)$ , and the dimension of hidden confounders  $d^U$  is varied across  $\{1, 5, 10, 15, 20, 25, 30\}$ .

**Evaluation criteria.** We evaluate model performances using the Precision in Estimation of Heterogeneous Effect (PEHE) (Hill, 2011), a standard metric that measures the root mean squared error between the estimated and true CATE values, denoted by  $\epsilon_{\text{PEHE}}(\hat{\tau})$ . And we use  $\epsilon_{\text{RI}}(\hat{\tau})$  to denote the relative improvement of a base CATE learner  $\hat{\tau}$  when it is trained with CARD, i.e., the CATE learner with CARD  $\hat{\tau}^{\text{CARD}}$ , denoted by  $\epsilon_{\text{RI}}(\hat{\tau})$ .

$$\epsilon_{\text{PEHE}}(\tau) = \sqrt{\frac{1}{n} \sum_{i=1}^n (\tau(X_i) - \tau_{\text{true}}(X_i))^2},$$

where  $\tau$  and  $\tau_{\text{true}}$  are arbitrary CATE model and the ground-truth CATE function, respectively. To specifically quantify the benefit of our method, we also introduce the Relative Improvement (RI) in PEHE. This metric calculates the percentage reduction in PEHE achieved by applying our CARD framework to a base CATE estimator:

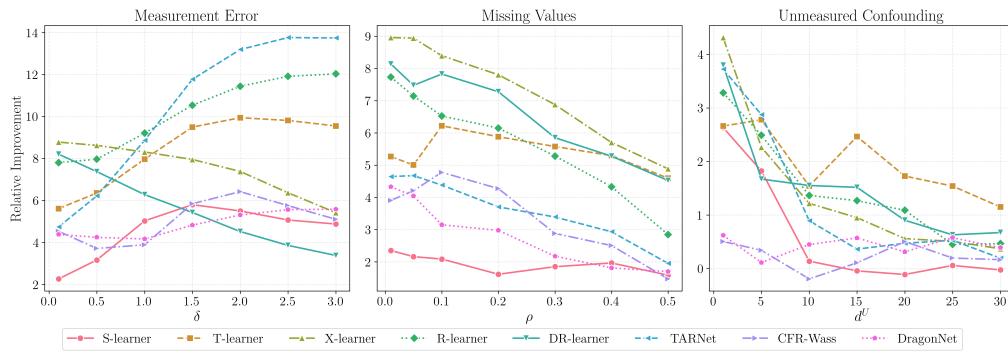
$$\epsilon_{\text{RI}}(\hat{\tau}) = \frac{\epsilon_{\text{PEHE}}(\hat{\tau}) - \epsilon_{\text{PEHE}}(\hat{\tau}^{\text{CARD}})}{\epsilon_{\text{PEHE}}(\hat{\tau})},$$

where  $\epsilon_{\text{PEHE}}(\hat{\tau})$  is the PEHE of the original CATE learner and  $\epsilon_{\text{PEHE}}(\hat{\tau}^{\text{CARD}})$  is the PEHE of the same estimator after being trained with CARD.

324

325  
326  
327  
Table 1: Comparison of **average PEHE** over 30 runs for various CATE learners, with and without  
the CARD framework, under three distribution shift scenarios: measurement error, missing values,  
and unmeasured confounding. **Bold** denotes the better results for each learner pair.

Settings	Measurement error (controlled by $\delta$ )						Missing values (controlled by $\rho$ )						Unmeasured confounding (controlled by $d^U$ )								
Bias level	0.1	0.5	1.0	1.5	2.0	2.5	3.0	0.01	0.05	0.1	0.2	0.3	0.4	0.5	1	5	10	15	20	25	30
S-learner	0.578	0.632	0.761	0.889	0.981	1.041	1.080	0.577	0.581	0.586	0.598	0.622	0.649	0.671	0.662	0.956	1.562	<b>1.883</b>	<b>2.139</b>	2.453	<b>2.862</b>
S+CARD	<b>0.565</b>	<b>0.612</b>	<b>0.722</b>	<b>0.838</b>	<b>0.927</b>	<b>0.988</b>	<b>1.028</b>	<b>0.564</b>	<b>0.569</b>	<b>0.574</b>	<b>0.589</b>	<b>0.611</b>	<b>0.636</b>	<b>0.661</b>	<b>0.644</b>	<b>0.939</b>	<b>1.560</b>	1.884	2.141	<b>2.452</b>	2.863
T-learner	1.105	1.283	1.833	2.669	3.653	4.705	5.792	1.096	1.085	1.072	1.056	1.065	1.057	1.014	1.240	1.420	1.902	2.187	2.396	2.693	3.069
T+CARD	<b>1.043</b>	<b>1.201</b>	<b>1.687</b>	<b>2.416</b>	<b>3.290</b>	<b>4.243</b>	<b>5.239</b>	<b>1.039</b>	<b>1.031</b>	<b>1.005</b>	<b>0.994</b>	<b>1.006</b>	<b>1.001</b>	<b>0.967</b>	<b>1.207</b>	<b>1.381</b>	<b>1.873</b>	<b>2.133</b>	<b>2.355</b>	<b>2.651</b>	<b>3.034</b>
X-learner	0.669	0.734	0.921	1.209	1.568	1.973	2.408	0.667	0.669	0.672	0.683	0.702	0.723	0.731	0.771	1.033	1.647	1.955	2.193	2.503	2.902
X+CARD	<b>0.610</b>	<b>0.670</b>	<b>0.844</b>	<b>1.113</b>	<b>1.453</b>	<b>1.848</b>	<b>2.277</b>	<b>0.608</b>	<b>0.609</b>	<b>0.615</b>	<b>0.629</b>	<b>0.654</b>	<b>0.681</b>	<b>0.695</b>	<b>0.738</b>	<b>1.010</b>	<b>1.627</b>	<b>1.936</b>	<b>2.180</b>	<b>2.490</b>	<b>2.891</b>
R-learner	0.825	0.889	1.082	1.379	1.741	2.143	2.567	0.822	0.819	0.815	0.813	0.818	0.820	0.818	0.847	1.053	1.746	2.051	2.267	2.549	2.962
R+CARD	<b>0.760</b>	<b>0.818</b>	<b>0.983</b>	<b>1.234</b>	<b>1.542</b>	<b>1.887</b>	<b>2.258</b>	<b>0.759</b>	<b>0.761</b>	<b>0.761</b>	<b>0.763</b>	<b>0.775</b>	<b>0.785</b>	<b>0.794</b>	<b>0.820</b>	<b>1.027</b>	<b>1.722</b>	<b>2.025</b>	<b>2.242</b>	<b>2.538</b>	<b>2.949</b>
DR-learner	0.755	0.861	1.169	1.632	2.193	2.810	3.460	0.751	0.747	0.746	0.752	0.766	0.782	0.783	0.900	1.142	1.701	2.001	2.237	2.546	2.939
DR+CARD	<b>0.693</b>	<b>0.797</b>	<b>1.095</b>	<b>1.544</b>	<b>2.094</b>	<b>2.701</b>	<b>3.343</b>	<b>0.690</b>	<b>0.691</b>	<b>0.687</b>	<b>0.697</b>	<b>0.721</b>	<b>0.741</b>	<b>0.747</b>	<b>0.872</b>	<b>1.123</b>	<b>1.675</b>	<b>1.971</b>	<b>2.217</b>	<b>2.530</b>	<b>2.920</b>
TARNet	0.675	0.776	1.080	1.559	2.164	2.843	3.564	0.670	0.670	0.671	0.683	0.708	0.731	0.731	0.842	1.109	1.612	1.903	2.173	2.499	2.881
TAR+CARD	<b>0.643</b>	<b>0.728</b>	<b>0.984</b>	<b>1.376</b>	<b>1.879</b>	<b>2.452</b>	<b>3.074</b>	<b>0.639</b>	<b>0.639</b>	<b>0.642</b>	<b>0.657</b>	<b>0.684</b>	<b>0.710</b>	<b>0.717</b>	<b>0.811</b>	<b>1.077</b>	<b>1.598</b>	<b>1.896</b>	<b>2.163</b>	<b>2.486</b>	<b>2.875</b>
CFR-Wass	0.707	0.765	0.965	1.336	1.840	2.409	3.009	0.693	0.694	0.705	0.715	0.728	0.746	0.744	0.771	1.040	<b>1.617</b>	1.936	2.187	2.498	2.898
CFR+CARD	<b>0.675</b>	<b>0.736</b>	<b>0.927</b>	<b>1.258</b>	<b>1.722</b>	<b>2.270</b>	<b>2.856</b>	<b>0.666</b>	<b>0.664</b>	<b>0.671</b>	<b>0.684</b>	<b>0.707</b>	<b>0.727</b>	<b>0.733</b>	<b>0.767</b>	<b>1.036</b>	1.620	<b>1.934</b>	<b>2.176</b>	<b>2.493</b>	<b>2.894</b>
DragonNet	0.689	0.780	1.058	1.502	2.059	2.682	3.343	0.686	0.685	0.683	0.695	0.716	0.736	0.738	0.832	1.084	1.640	1.945	2.193	2.512	2.910
Dragon+CARD	<b>0.659</b>	<b>0.747</b>	<b>1.014</b>	<b>1.429</b>	<b>1.950</b>	<b>2.533</b>	<b>3.156</b>	<b>0.656</b>	<b>0.657</b>	<b>0.662</b>	<b>0.674</b>	<b>0.700</b>	<b>0.723</b>	<b>0.725</b>	<b>0.827</b>	<b>1.083</b>	<b>1.633</b>	<b>1.934</b>	<b>2.186</b>	<b>2.498</b>	<b>2.899</b>

351  
352  
353  
354  
355  
Figure 2: Performance gains from CARD: Relative Improvement in **average PEHE** over 30 runs  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
399  
398  
397  
396  
395  
394  
393  
392  
391  
390  
389  
388  
387  
386  
385  
384  
383  
382  
381  
380  
379  
378  
377  
376  
375  
374  
373  
372  
371  
370  
369  
368  
367  
366  
365  
364  
363  
362  
361  
360  
359  
358  
357  
356  
355  
354  
353  
352  
351  
350  
349  
348  
347  
346  
345  
344  
343  
342  
341  
340  
339  
338  
337  
336  
335  
334  
333  
332  
331  
330  
329  
328  
327  
326  
325  
324  
323  
322  
321  
320  
319  
318  
317  
316  
315  
314  
313  
312  
311  
310  
309  
308  
307  
306  
305  
304  
303  
302  
301  
300  
299  
298  
297  
296  
295  
294  
293  
292  
291  
290  
289  
288  
287  
286  
285  
284  
283  
282  
281  
280  
279  
278  
277  
276  
275  
274  
273  
272  
271  
270  
269  
268  
267  
266  
265  
264  
263  
262  
261  
260  
259  
258  
257  
256  
255  
254  
253  
252  
251  
250  
249  
248  
247  
246  
245  
244  
243  
242  
241  
240  
239  
238  
237  
236  
235  
234  
233  
232  
231  
230  
229  
228  
227  
226  
225  
224  
223  
222  
221  
220  
219  
218  
217  
216  
215  
214  
213  
212  
211  
210  
209  
208  
207  
206  
205  
204  
203  
202  
201  
200  
199  
198  
197  
196  
195  
194  
193  
192  
191  
190  
189  
188  
187  
186  
185  
184  
183  
182  
181  
180  
179  
178  
177  
176  
175  
174  
173  
172  
171  
170  
169  
168  
167  
166  
165  
164  
163  
162  
161  
160  
159  
158  
157  
156  
155  
154  
153  
152  
151  
150  
149  
148  
147  
146  
145  
144  
143  
142  
141  
140  
139  
138  
137  
136  
135  
134  
133  
132  
131  
130  
129  
128  
127  
126  
125  
124  
123  
122  
121  
120  
119  
118  
117  
116  
115  
114  
113  
112  
111  
110  
109  
108  
107  
106  
105  
104  
103  
102  
101  
100  
99  
98  
97  
96  
95  
94  
93  
92  
91  
90  
89  
88  
87  
86  
85  
84  
83  
82  
81  
80  
79  
78  
77  
76  
75  
74  
73  
72  
71  
70  
69  
68  
67  
66  
65  
64  
63  
62  
61  
60  
59  
58  
57  
56  
55  
54  
53  
52  
51  
50  
49  
48  
47  
46  
45  
44  
43  
42  
41  
40  
39  
38  
37  
36  
35  
34  
33  
32  
31  
30  
29  
28  
27  
26  
25  
24  
23  
22  
21  
20  
19  
18  
17  
16  
15  
14  
13  
12  
11  
10  
9  
8  
7  
6  
5  
4  
3  
2  
1  
0
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
399  
398  
397  
396  
395  
394  
393  
392  
391  
390  
389  
388  
387  
386  
385  
384  
383  
382  
381  
380  
379  
378  
377  
376  
375  
374  
373  
372  
371  
370  
369  
368  
367  
366  
365  
364  
363  
362  
361  
360  
359  
358  
357  
356  
355  
354  
353  
352  
351  
350  
350  
349  
348  
347  
346  
345  
344  
343  
342  
341  
340  
340  
339  
338  
337  
336  
335  
334  
333  
332  
331  
330  
330  
329  
328  
327  
326  
325  
324  
324  
323  
322  
321  
320  
320  
319  
318  
317  
316  
315  
314  
313  
312  
311  
310  
310  
309  
308  
307  
306  
305  
304  
303  
302  
301  
301  
300  
300  
299  
298  
297  
296  
295  
294  
293  
292  
291  
291  
290  
290  
289  
288  
287  
286  
285  
284  
283  
282  
281  
281  
280  
280  
279  
278  
277  
276  
275  
274  
273  
273  
272  
272  
271  
271  
270  
270  
269  
268  
267  
266  
265  
264  
263  
263  
262  
262  
261  
261  
260  
260  
259  
258  
257  
256  
255  
254  
254  
253  
253  
252  
252  
251  
251  
250  
250  
249  
248  
247  
246  
245  
245  
244  
244  
243  
243  
242  
242  
241  
241  
240  
240  
239  
238  
237  
236  
235  
235  
234  
234  
233  
233  
232  
232  
231  
231  
230  
230  
229  
228  
227  
226  
225  
225  
224  
224  
223  
223  
222  
222  
221  
221  
220  
220  
219  
219  
218  
218  
217  
217  
216  
216  
215  
215  
214  
214  
213  
213  
212  
212  
211  
211  
210  
210  
209  
209  
208  
208  
207  
207  
206  
206  
205  
205  
204  
204  
203  
203  
202  
202  
201  
201  
200  
200  
199  
199  
198  
198  
197  
197  
196  
196  
195  
195  
194  
194  
193  
193  
192  
192  
191  
191  
190  
190  
189  
189  
188  
188  
187  
187  
186  
186  
185  
185  
184  
184  
183  
183  
182  
182  
181  
181  
180  
180  
179  
179  
178  
178  
177  
177  
176  
176  
175  
175  
174  
174  
173  
173  
172  
172  
171  
171  
170  
170  
169  
169  
168  
168  
167  
167  
166  
166  
165  
165  
164  
164  
163  
163  
162  
162  
161  
161  
160  
160  
159  
159  
158  
158  
157  
157  
156  
156  
155  
155  
154  
154  
153  
153  
152  
152  
151  
151  
150  
150  
149  
149  
148  
148  
147  
147  
146  
146  
145  
145  
144  
144  
143  
143  
142  
142  
141  
141  
140  
140  
139  
139  
138  
138  
137  
137  
136  
136  
135  
135  
134  
134  
133  
133  
132  
132  
131  
131  
130  
130  
129  
129  
128  
128  
127  
127  
126  
126  
125  
125  
124  
124  
123  
123  
122  
122  
121  
121  
120  
120  
119  
119  
118  
118  
117  
117  
116  
116  
115  
115  
114  
114  
113  
113  
112  
112  
111  
111  
110  
110  
109  
109  
108  
108  
107  
107  
106  
106  
105  
105  
104  
104  
103  
103  
102  
102  
101  
101  
100  
100  
99  
99  
98  
98  
97  
97  
96  
96  
95  
95  
94  
94  
93  
93  
92  
92  
91  
91  
90  
90  
89  
89  
88  
88  
87  
87  
86  
86  
85  
85  
84  
84  
83  
83  
82  
82  
81  
81  
80  
80  
79  
79  
78  
78  
77  
77  
76  
76  
75  
75  
74  
74  
73  
73  
72  
72  
71  
71  
70  
70  
69  
69  
68  
68  
67  
67  
66  
66  
65  
65  
64  
64  
63  
63  
62  
62  
61  
61  
60  
60  
59  
59  
58  
58  
57  
57  
56  
56  
55  
55  
54  
54  
53  
53  
52  
52  
51  
51  
50  
50  
49  
49  
48  
48  
47  
47  
46  
46  
45  
45  
44  
44  
43  
43  
42  
42  
41  
41  
40  
40  
39  
39  
38  
38  
37  
37  
36  
36  
35  
35  
34  
34  
33  
33  
32  
32  
31  
31  
30  
30  
29  
29  
28  
28  
27  
27  
26  
26  
25  
25  
24  
24  
23  
23  
22  
22  
21  
21  
20  
20  
19  
19  
18  
18  
17  
17  
16  
16  
15  
15  
14  
14  
13  
13  
12  
12  
11  
11  
10  
10  
9  
9  
8  
8  
7  
7  
6  
6  
5  
5  
4  
4  
3  
3  
2  
2  
1  
1  
0  
0

**CARD consistently enhances performance across diverse learners and shifts.** The

Table 2: Comparison of **worst-case PEHE** over 30 runs for various CATE learners, with and without the CARD framework, under three distribution shift scenarios: measurement error, missing values, and unmeasured confounding. Bold denotes the better results for each learner pair.

Settings	Measurement error (controlled by $\delta$ )						Missing values (controlled by $\rho$ )						Unmeasured confounding (controlled by $\eta^U$ )									
	Bias level	0.1	0.5	1.0	1.5	2.0	2.5	3.0	0.01	0.05	0.1	0.2	0.3	0.4	0.5	1	5	10	15	20	25	30
S-learner	1.020	1.058	1.173	1.448	1.746	1.961	2.118	1.023	1.024	1.032	1.043	1.071	1.105	1.128	0.942	1.282	1.897	2.083	2.296	2.525	<b>3.010</b>	
S+CARD	<b>1.024</b>	<b>1.040</b>	<b>1.108</b>	<b>1.157</b>	<b>1.401</b>	<b>1.581</b>	<b>1.706</b>	<b>1.026</b>	<b>1.029</b>	<b>1.037</b>	<b>1.050</b>	<b>1.067</b>	<b>1.090</b>	<b>1.119</b>	<b>0.826</b>	<b>1.123</b>	<b>1.880</b>	2.004	2.259	<b>2.498</b>	2.985	
T-learner	1.257	2.781	3.617	4.953	6.574	8.315	10.109	2.483	2.470	2.395	2.335	2.406	2.250	2.064	2.047	2.541	2.948	3.122	3.269	3.585	3.792	
T+CARD	<b>1.453</b>	<b>1.631</b>	<b>2.252</b>	<b>3.207</b>	<b>4.386</b>	<b>5.659</b>	<b>7.008</b>	<b>1.463</b>	<b>1.451</b>	<b>1.413</b>	<b>1.404</b>	<b>1.400</b>	<b>1.366</b>	<b>1.374</b>	<b>1.558</b>	<b>1.624</b>	<b>2.037</b>	<b>2.279</b>	<b>2.550</b>	<b>2.906</b>	<b>3.203</b>	
X-learner	1.447	1.478	1.688	2.096	2.737	3.707	4.729	1.441	1.445	1.381	1.396	1.400	1.373	1.306	1.247	1.424	2.283	2.285	2.489	2.762	3.070	
X+CARD	<b>1.294</b>	<b>1.310</b>	<b>1.260</b>	<b>1.622</b>	<b>1.344</b>	<b>1.722</b>	<b>2.105</b>	<b>1.288</b>	<b>1.295</b>	<b>1.245</b>	<b>1.258</b>	<b>1.264</b>	<b>1.239</b>	<b>1.200</b>	<b>0.973</b>	<b>1.177</b>	<b>2.165</b>	<b>2.382</b>	<b>2.595</b>	<b>3.623</b>	<b>3.021</b>	
R-learner	1.389	1.473	2.053	2.986	4.110	5.327	6.598	1.386	1.385	1.354	1.375	1.395	1.388	1.351	1.368	1.398	2.309	2.613	2.719	2.829	3.331	
R+CARD	<b>1.190</b>	<b>1.563</b>	<b>1.818</b>	<b>2.400</b>	<b>3.177</b>	<b>4.048</b>	<b>5.008</b>	<b>1.187</b>	<b>1.207</b>	<b>1.164</b>	<b>1.151</b>	<b>1.166</b>	<b>1.157</b>	<b>1.137</b>	<b>1.444</b>	<b>1.628</b>	<b>2.406</b>	<b>2.580</b>	<b>2.764</b>	<b>2.909</b>	<b>3.245</b>	
DR-learner	1.598	1.677	2.097	2.723	3.463	4.809	6.251	1.592	1.594	1.500	1.529	1.533	1.495	1.398	1.496	1.645	2.410	2.298	2.534	2.899	3.137	
DR+CARD	<b>1.485</b>	<b>1.588</b>	<b>1.966</b>	<b>2.517</b>	<b>3.240</b>	<b>4.160</b>	<b>5.462</b>	<b>1.479</b>	<b>1.481</b>	<b>1.360</b>	<b>1.369</b>	<b>1.392</b>	<b>1.351</b>	<b>1.314</b>	<b>1.382</b>	<b>1.548</b>	<b>2.330</b>	<b>2.232</b>	<b>2.463</b>	<b>2.831</b>	<b>3.060</b>	
TARNet	1.214	1.234	1.607	2.837	4.566	6.568	8.715	1.213	1.213	1.172	1.174	1.177	1.176	1.156	1.394	1.621	2.112	2.193	2.372	2.698	3.122	
TAR+CARD	<b>1.192</b>	<b>1.224</b>	<b>1.034</b>	<b>1.623</b>	<b>2.479</b>	<b>3.516</b>	<b>4.686</b>	<b>1.193</b>	<b>1.195</b>	<b>1.175</b>	<b>1.178</b>	<b>1.186</b>	<b>1.178</b>	<b>1.163</b>	<b>0.943</b>	<b>1.244</b>	<b>2.030</b>	<b>2.032</b>	<b>2.343</b>	<b>2.540</b>	<b>3.015</b>	
CFR-Wass	1.495	1.536	1.656	2.956	4.953	7.314	9.834	1.489	1.487	1.391	1.398	1.384	1.364	1.339	1.502	1.724	2.336	2.180	2.423	2.658	3.095	
CFR+CARD	<b>1.401</b>	<b>1.436</b>	<b>1.535</b>	<b>1.738</b>	<b>2.814</b>	<b>4.131</b>	<b>5.565</b>	<b>1.399</b>	<b>1.397</b>	<b>1.319</b>	<b>1.335</b>	<b>1.324</b>	<b>1.309</b>	<b>1.265</b>	<b>0.976</b>	<b>1.288</b>	<b>2.175</b>	<b>1.940</b>	<b>2.261</b>	<b>2.446</b>	<b>2.927</b>	
DragonNet	1.508	1.539	1.659	2.481	3.547	4.867	6.447	1.505	1.503	1.396	1.412	1.415	1.388	1.320	1.265	1.634	2.326	2.277	2.424	2.738	3.100	
Dragon+CARD	<b>1.269</b>	<b>1.295</b>	<b>1.361</b>	<b>2.422</b>	<b>3.376</b>	<b>5.429</b>	<b>7.072</b>	<b>1.262</b>	<b>1.265</b>	<b>1.220</b>	<b>1.221</b>	<b>1.225</b>	<b>1.205</b>	<b>1.172</b>	<b>1.526</b>	<b>1.740</b>	<b>2.145</b>	<b>2.011</b>	<b>2.238</b>	<b>2.587</b>	<b>3.059</b>	

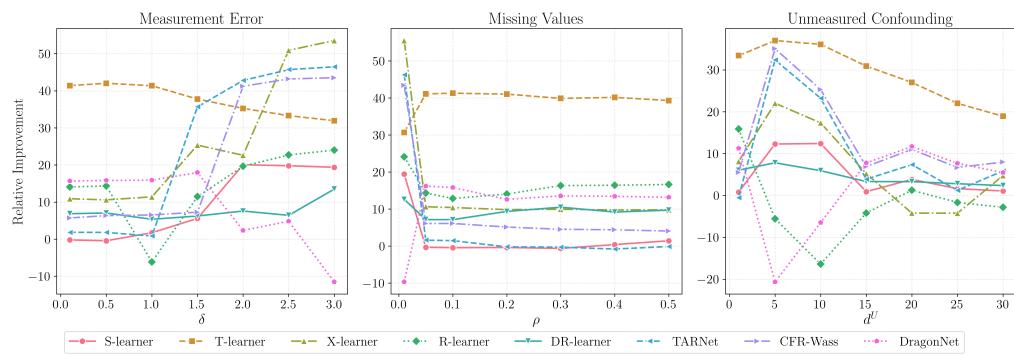


Figure 3: Performance gains from CARD: Relative Improvement in **worst-case PEHE** over 30 runs for various CATE learners under three distribution shift scenarios. The value is in percentage.

consistent gains under measurement error: relative improvements often exceed 10% for flexible learners such as TARNet and the R-learner at high noise levels, and TARNet’s RI rises monotonically with noise to a peak of approximately 14% at  $\delta = 2.5$ . This pattern is intuitive: small amounts of noise leave a learner near its clean optimum so adversarial augmentation gives modest gains, whereas larger noise exposes vulnerabilities that the RL-guided diffusion discovers and the learner then learns to resist. By contrast, missing-value corruptions yield smaller but stable improvements, with RI ranging from 2%-9%, likely because the pipeline applies imputation (MICE) (Kallus et al., 2018), which already reduces extreme covariate variation and therefore narrows the space of harmful yet realistic augmentations. Unmeasured confounding is the most challenging regime: relative gains are smaller (0%-4%) but remain practically important because they help preserve performance when treatment heterogeneity itself shifts. Overall, these results show that CARD offers stronger defense against measurement error and missing values, and it also delivers consistent, constructive gains when unseen shifts arise from latent confounders.

**The magnitude of CARD’s improvement is model-specific.** Results confirm that nearly every base learner benefits, and more interestingly, the magnitude of this improvement is heterogeneous. For instance, more flexible models like TARNet, R-learner, and X-learner are among the biggest beneficiaries, particularly under measurement error. This model-specific efficacy can be attributed to two main factors. First, learners possess different inductive biases (Curth & Van der Schaar, 2021). Models with more flexible function classes, such as TARNet, can better exploit the adversarial augmentations to learn more robust CATE functions. In contrast, simpler or heavily regularized learners may already exhibit some robustness, leaving less capacity for substantial improvement. Second, baseline vulnerability plays a key role. The T-learner, for example, which is highly susceptible to noise, still receives significant relative performance gains, suggesting that CARD effectively enhances resilience even when baseline errors are large.

432 5.2.2 WORST-CASE ROBUSTNESS ANALYSIS  
433434 In addition to the analysis of average PEHE, we also investigate whether CARD can improve the  
435 worst-case PEHE, which is a critical measure of model stability and robustness. Relevant results are  
436 reported in Table 2 and Figure 3.437 **CARD consistently enhances performance across diverse learners and shifts.** A key finding  
438 from Table 2 and Figure 3 is that CARD’s impact on a model’s worst-case performance is signif-  
439 icantly larger than its effect on average performance. Averaged across learners, the mean worst-  
440 case RI is substantial for measurement-error scenarios (around 21.1%), moderate for missing-value  
441 scenarios (around 11.5%), and smaller but nontrivial for unmeasured confounding (around 8.0%).  
442 Notably, the RI in worst-case PEHE is frequently three to five times greater than the improvement  
443 observed in the average-case. For instance, while CARD consistently improves the T-learner’s aver-  
444 age PEHE by approximately 10% in many high-noise settings, it enhances its worst-case PEHE by  
445 a massive 30-40% under the same conditions. This disparity reveals CARD’s primary mechanism:  
446 CARD not only shifts average behavior but also substantially reduces the tail risk.  
447448 **The magnitude of CARD’s improvement is context-dependent.** Similar to average-case results,  
449 we also find CARD’s capabilities are dependent on type of distribution shifts in worst-case. As  
450 shown in Figure 3, CARD delivers its most dramatic gains under measurement error, a scenario  
451 that often causes covariate shifts due to additive noise in standard models. Here, CARD slashes  
452 the worst-case PEHE of flexible learners CFR-Wass by 43.4% (from 9.834 to 5.565) and that of  
453 TARNNet by 46.2% (from 8.715 to 4.686) at the highest noise level ( $\delta = 3.0$ ). By contrast, missing-  
454 ness produces smaller but stable improvements, which is consistent with our previous observation.  
455 Under the more structured challenge of unmeasured confounding, the improvements, while smaller  
456 or sometimes negative, are still effective for enhancing the worst-case performance. For example,  
457 it reduces the T-learner’s worst-case PEHE by a substantial 20%-30% at the highest confounding  
458 dimension. However, this results also highlights a fundamental problem in causal identification:  
459 while data augmentations can significantly improve a model’s robustness, they cannot identify ora-  
460 cle causal information that is actually absent.461 **The magnitude of CARD’s improvement is model-specific.** The benefits of CARD are distinct  
462 across base learners. As illustrated in Figure 3, for instance, some learners exhibit large worst-case  
463 RI under measurement error: T-learner achieves worst-case RI averaged in all bias levels with about  
464 36.0%, and X-learner with about 32.8%. Others show modest gains, such as DragonNet (5.1%) and  
465 S-learner (12.2%). This heterogeneity can be attributed to two complementary factors: (i) learners  
466 with higher worst-case baseline PEHE have more room for improvement; and (ii) flexible neural  
467 architectures like TARNNet can leverage adversarial augmentations to learn more stable conditional  
468 effects. These two observations are aligned with previous average-case results. Interestingly, in a  
469 few cases (e.g., R-learner and DragonNet under certain hidden confounding levels), CARD produces  
470 marginally negative RI, which may be linked to their connections with targeted maximum likelihood  
471 estimation (TMLE), where CARD’s perturbations interact with the targeted nuisance components.472 6 CONCLUSION  
473474 In this work, we introduce CARD, a novel and model-agnostic framework that is capable to improve  
475 the robustness of any existing CATE learner to unknown distribution shift, without requiring prior  
476 knowledge or additional structural assumptions in the deployment domain. Rather than proposing  
477 a new CATE estimation algorithm, our primary goal is to investigate how reinforcement learning  
478 guided diffusion models can generate adversarial proxies that encourage the CATE learner to adapt  
479 and remain resilient to unseen distribution shifts. Experiments across diverse learners and distri-  
480 bution shift types show consistent gains from CARD, highlighting its potential effectiveness for  
481 real-world deployment. The limitation of this work lies in the computational complexity, a common  
482 challenge for diffusion models, as discussed in Section A.2. An interesting future research is the  
483 complexity improvement with recent acceleration techniques (Chen et al., 2024). Simultaneously,  
484 the success of this approach might open exciting future directions, including extending its applica-  
485 tion to other causal tasks related to generative modeling, such as counterfactual generation (Yoon  
et al., 2018), dimension reduction (Liu et al., 2024), and model evaluation (Athey et al., 2024).

486     **Reproducibility statement.** We defer the implementation details of using CARD to train CATE  
 487     in Appendix A.1. The uploaded code can be directly used to reproduce our experimental results.  
 488     Additionally, we list all the referred and required resources with an instruction file in supplementary.  
 489

490     REFERENCES  
 491

492     Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. When should you adjust  
 493     standard errors for clustering? *The Quarterly Journal of Economics*, 138(1):1–35, 2023.

494     Anish Agarwal and Rahul Singh. Causal inference with corrupted data: Measurement error, missing  
 495     values, discretization, and differential privacy. *arXiv preprint arXiv:2107.02780*, 2021.

496

497     Susan Athey, Guido W Imbens, Jonas Metzger, and Evan Munro. Using wasserstein generative  
 498     adversarial networks for the design of monte carlo simulations. *Journal of Econometrics*, 240(2):  
 499     105076, 2024.

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

502

503     Elias Bareinboim and Judea Pearl. Causal inference and the data-fusion problem. *Proceedings of  
 504     the National Academy of Sciences*, 113(27):7345–7352, 2016.

505

506     Erich Battistin and Andrew Chesher. Treatment effect estimation with covariate measurement error.  
 507     *Journal of Econometrics*, 178(2):707–715, 2014.

508

509     Ioana Bica, Ahmed M Alaa, Craig Lambert, and Mihaela Van Der Schaar. From real-world patient  
 510     data to individualized treatment effects using machine learning: current and future methods to  
 511     address underlying challenges. *Clinical Pharmacology & Therapeutics*, 109(1):87–100, 2021.

512

513     Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion  
 514     models with reinforcement learning. In *The Twelfth International Conference on Learning Rep-  
 515     resentations*.

516

517     Haoxuan Chen, Yinuo Ren, Lexing Ying, and Grant Rotskoff. Accelerating diffusion models with  
 518     parallel sampling: Inference at sub-linear time complexity. *Advances in Neural Information Pro-  
 519     cessing Systems*, 37:133661–133709, 2024.

520

521     Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney  
 522     Newey, and James Robins. Double/debiased machine learning for treatment and structural pa-  
 523     rameters, 2018.

524

525     Alicia Curth and Mihaela Van der Schaar. On inductive biases for heterogeneous treatment effect  
 526     estimation. *Advances in Neural Information Processing Systems*, 34:15883–15894, 2021.

527

528     Alicia Curth and Mihaela Van Der Schaar. In search of insights, not magic bullets: Towards de-  
 529     mystification of the model selection dilemma in heterogeneous treatment effect estimation. In  
 530     *International conference on machine learning*, pp. 6623–6642. PMLR, 2023.

531

532     Issa J Dahabreh and Miguel A Hernán. Extending inferences from a randomized trial to a target  
 533     population. *European journal of epidemiology*, 34(8):719–722, 2019.

534

535     Issa J Dahabreh, Sarah E Robertson, Lucia C Petito, Miguel A Hernán, and Jon A Steingrimsson.  
 536     Efficient and robust methods for causally interpretable meta-analysis: Transporting inferences  
 537     from multiple randomized trials to a target population. *Biometrics*, 79(2):1057–1072, 2023.

538

539     Sihao Ding, Peng Wu, Fuli Feng, Yitong Wang, Xiangnan He, Yong Liao, and Yongdong Zhang.  
 540     Addressing unmeasured confounder for recommendation with sensitivity analysis. In *Proceedings  
 541     of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 305–315,  
 542     2022.

543

544     Vincent Dorie, Jennifer Hill, Uri Shalit, Marc Scott, and Dan Cervone. Automated versus do-it-  
 545     yourself methods for causal inference: Lessons learned from a data analysis competition. *Statis-  
 546     tical Science*, 34(1):43–68, 2019.

540 Jacob Dorn, Kevin Guo, and Nathan Kallus. Doubly-valid/doubly-sharp sensitivity analysis for  
 541 causal inference with unmeasured confounding. *Journal of the American Statistical Association*,  
 542 120(549):331–342, 2025.

543 Max H Farrell. Robust inference on average treatment effects with possibly more covariates than  
 544 observations. *Journal of Econometrics*, 189(1):1–23, 2015.

545 Carlos Fernández-Loría, Foster Provost, Jesse Anderton, Benjamin Carterette, and Praveen Chandar.  
 546 A comparison of methods for treatment assignment with an application to playlist generation.  
 547 *Information Systems Research*, 34(2):786–803, 2023.

548 Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Ali-  
 549 cia Curth, Stefan Bauer, Niki Kilbertus, Isaac S Kohane, and Mihaela van der Schaar. Causal  
 550 machine learning for predicting treatment outcomes. *Nature Medicine*, 30(4):958–968, 2024.

551 Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *The Annals of Statistics*, 51  
 552 (3):879–908, 2023.

553 Tobias Hatt, Jeroen Berrevoets, Alicia Curth, Stefan Feuerriegel, and Mihaela van der Schaar. Com-  
 554 bining observational and randomized data for estimating heterogeneous treatment effects. *arXiv*  
 555 preprint *arXiv:2202.12891*, 2022.

556 Konstantin Hess, Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. Efficient and sharp  
 557 off-policy learning under unobserved confounding. *arXiv preprint arXiv:2502.13022*, 2025.

558 Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational*  
 559 *and Graphical Statistics*, 20(1):217–240, 2011.

560 Melody Huang, Naoki Egami, Erin Hartman, and Luke Miratrix. Leveraging population outcomes  
 561 to improve the generalization of experimental results: Application to the jtpa study. *The Annals*  
 562 *of Applied Statistics*, 17(3):2139–2164, 2023a.

563 Melody Y Huang. Sensitivity analysis for the generalization of experimental results. *Journal of the*  
 564 *Royal Statistical Society Series A: Statistics in Society*, 187(4):900–918, 2024.

565 Yiyuan Huang, Cheuk Hang Leung, Shumin Ma, Zhiri Yuan, Qi Wu, Siyi Wang, Dongdong Wang,  
 566 and Zhixiang Huang. Towards balanced representation learning for credit policy evaluation. In  
 567 *International Conference on Artificial Intelligence and Statistics*, pp. 3677–3692. PMLR, 2023b.

568 Yiyuan Huang, Cheuk H Leung, Siyi Wang, Yijun Li, and Qi Wu. Unveiling the potential of robust-  
 569 ness in selecting conditional average treatment effect estimators. *Advances in Neural Information*  
 570 *Processing Systems*, 37:135208–135243, 2024.

571 Kosuke Imai and Teppei Yamamoto. Causal inference with differential measurement error: Non-  
 572 parametric identification and sensitivity analysis. *American Journal of Political Science*, 54(2):  
 573 543–560, 2010.

574 Sookyo Jeong and Hongseok Namkoong. Robust causal inference under covariate shift via worst-  
 575 case subpopulation treatment effects. In *Conference on Learning Theory*, pp. 2079–2084. PMLR,  
 576 2020.

577 Fredrik D Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and rep-  
 578 resentation learning for estimation of potential outcomes and causal effects. *Journal of Machine*  
 579 *Learning Research*, 23(166):1–50, 2022.

580 Nathan Kallus and Angela Zhou. Minimax-optimal policy learning under unobserved confounding.  
 581 *Management Science*, 67(5):2870–2890, 2021.

582 Nathan Kallus, Xiaojie Mao, and Madeleine Udell. Causal inference with noisy and missing covari-  
 583 ates via matrix factorization. *Advances in neural information processing systems*, 31, 2018.

584 Nathan Kallus, Xiaojie Mao, and Angela Zhou. Interval estimation of individual-level causal effects  
 585 under unobserved confounding. In *The 22nd international conference on artificial intelligence*  
 586 *and statistics*, pp. 2281–2290. PMLR, 2019.

594 Nathan Kallus, Xiaojie Mao, Kaiwen Wang, and Zhengyuan Zhou. Doubly robust distributionally  
 595 robust off-policy evaluation and learning. In *International Conference on Machine Learning*, pp.  
 596 10598–10632. PMLR, 2022.

597

598 Edward H Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects.  
 599 *Electronic Journal of Statistics*, 17(2):3008–3049, 2023.

600

601 Christoph Kern, Michael P Kim, and Angela Zhou. Multi-accurate cate is robust to unknown co-  
 602 variate shifts. *Transactions on Machine Learning Research*, 2024.

603

604 Daido Kido. Distributionally robust policy learning with wasserstein distance. *arXiv preprint*  
 605 *arXiv:2205.04637*, 2022.

606

607 Newton Mwai Kinyanjui and Fredrik D Johansson. Adcb: An alzheimer’s disease simulator for  
 608 benchmarking observational estimators of causal effects. In *Conference on Health, Inference,  
 609 and Learning*, pp. 103–118. PMLR, 2022.

610

611 Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization  
 612 methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.

613

614 Sören R Küngel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heteroge-  
 615 neous treatment effects using machine learning. *Proceedings of the national academy of sciences*,  
 616 116(10):4156–4165, 2019.

617

618 Manabu Kuroki and Judea Pearl. Measurement bias and effect restoration in causal inference.  
 619 *Biometrika*, 101(2):423–437, 2014.

620

621 Haoxuan Li, Chunyuan Zheng, Peng Wu, Kun Kuang, Yue Liu, and Peng Cui. Who should be  
 622 given incentives? counterfactual optimal treatment regimes learning for recommendation. In  
 623 *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, pp.  
 624 1235–1247, 2023.

625

626 Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under  
 627 network interference. *The Annals of Statistics*, 50(4):2334–2358, 2022.

628

629 Qiao Liu, Zhongren Chen, and Wing Hung Wong. An encoding generative modeling approach  
 630 to dimension reduction and covariate adjustment in causal inference with observational studies.  
 631 *Proceedings of the National Academy of Sciences*, 121(23):e2322376121, 2024.

632

633 Yuchen Ma, Jonas Schweisthal, Hengrui Zhang, and Stefan Feuerriegel. A diffusion-based method  
 634 for learning the multi-outcome distribution of medical treatments. In *Proceedings of the 31st ACM  
 635 SIGKDD Conference on Knowledge Discovery and Data Mining* V. 2, pp. 2066–2077, 2025.

636

637 Imke Mayer, Erik Sverdrup, Tobias Gauss, Jean-Denis Moyer, Stefan Wager, and Julie Josse. Doubly  
 638 robust treatment effect estimation with missing attributes. *The Annals of Applied Statistics*, 14(3):  
 639 1409–1431, 2020.

640

641 Karthika Mohan, Judea Pearl, and Tian Jin. Missing data as a causal inference problem. In *Proceed-  
 642 ings of the neural information processing systems conference (nips)*, 2013.

643

644 Tong Mu, Yash Chandak, Tatsunori B Hashimoto, and Emma Brunskill. Factored dro: Factored  
 645 distributionally robust policies for contextual bandits. *Advances in Neural Information Processing  
 646 Systems*, 35:8318–8331, 2022.

647

648 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.  
 649 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

650

651 Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects.  
 652 *Biometrika*, 108(2):299–319, 2021.

653

654 Xinkun Nie, Guido Imbens, and Stefan Wager. Covariate balancing sensitivity analysis for extrapo-  
 655 lating randomized trials across locations. *arXiv preprint arXiv:2112.04723*, 2021.

648 Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-  
 649 learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. In An-  
 650 dreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan  
 651 Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume  
 652 202 of *Proceedings of Machine Learning Research*, pp. 26599–26618. PMLR, 23–29 Jul 2023.  
 653 URL <https://proceedings.mlr.press/v202/oprescu23a.html>.

654 Judea Pearl and Elias Bareinboim. Transportability of causal and statistical relations: A formal  
 655 approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, pp. 247–  
 656 254, 2011.

657

658 Zhuan Pei, Jörn-Steffen Pischke, and Hannes Schwandt. Poorly measured confounders are more  
 659 useful on the left than on the right. *Journal of Business & Economic Statistics*, 37(2):205–216,  
 660 2019.

661 Zhaozhi Qian, Yao Zhang, Ioana Bica, Angela Wood, and Mihaela van der Schaar. Syncwin: Treat-  
 662 ment effect estimation with longitudinal outcomes. *Advances in Neural Information Processing  
 663 Systems*, 34:3178–3190, 2021.

664

665 Xinrui Ruan, Jingshen Wang, Yingfei Wang, and Waverly Wei. Electronic medical records assisted  
 666 digital clinical trial design. In *International Conference on Artificial Intelligence and Statistics*,  
 667 pp. 2836–2844. PMLR, 2024.

668 Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies.  
 669 *Journal of educational Psychology*, 66(5):688, 1974.

670

671 Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

672

673 Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal  
 674 of the American statistical Association*, 100(469):322–331, 2005.

675 Kara E Rudolph, Nicholas T Williams, Elizabeth A Stuart, and Ivan Diaz. Improving efficiency in  
 676 transporting average treatment effects. *Biometrika*, pp. asaf027, 2025.

677 Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: general-  
 678 ization bounds and algorithms. In *International conference on machine learning*, pp. 3076–3085.  
 679 PMLR, 2017.

680

681 Yi Shen, Pan Xu, and Michael Zavlanos. Wasserstein distributionally robust policy evaluation and  
 682 learning for contextual bandits. *Transactions on Machine Learning Research*.

683

684 Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment  
 685 effects. *Advances in neural information processing systems*, 32, 2019.

686

687 Nian Si, Fan Zhang, Zhengyuan Zhou, and Jose Blanchet. Distributionally robust batch contextual  
 688 bandits. *Management Science*, 69(10):5772–5793, 2023.

689

690 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 691 Poole. Score-based generative modeling through stochastic differential equations. In *Inter-  
 692 national Conference on Learning Representations*.

693

694 Namjoon Suh, Xiaofeng Lin, Din-Yin Hsieh, Mehrdad Honarkhah, and Guang Cheng. Autodiff:  
 695 combining auto-encoder and diffusion model for tabular data synthesizing. In *NeurIPS 2023  
 696 Workshop on Synthetic Data Generation with Generative AI*.

697

698 Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using  
 699 random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

700

701 Jingyuan Wang, Zhimei Ren, Ruohan Zhan, and Zhengyuan Zhou. Distributionally robust policy  
 702 learning under concept drifts. In *Forty-second International Conference on Machine Learning*.

Anpeng Wu, Haoxuan Li, Chunyuan Zheng, Kun Kuang, and Kun Zhang. Classifying treatment  
 703 responders: Bounds and algorithms. In *Proceedings of the 31st ACM SIGKDD Conference on  
 704 Knowledge Discovery and Data Mining V. 1*, pp. 1611–1622, 2025a.

702 Lili Wu and Shu Yang. Integrative  $r$ -learner of heterogeneous treatment effects combining ex-  
 703 perimental and observational studies. In *Conference on Causal Learning and Reasoning*, pp.  
 704 904–926. PMLR, 2022.

705 Peng Wu, Shanshan Luo, and Zhi Geng. On the comparative analysis of average treatment ef-  
 706 fects estimation via data combination. *Journal of the American Statistical Association*, pp. 1–12,  
 707 2025b.

708 Yanghao Xiao, Haoxuan Li, Yongqiang Tang, and Wensheng Zhang. Addressing hidden confound-  
 709 ing with heterogeneous observational datasets for recommendation. *Advances in Neural Infor-  
 710 mation Processing Systems*, 37:130358–130383, 2024.

712 Steve Yadlowsky, Hongseok Namkoong, Sanjay Basu, John Duchi, and Lu Tian. Bounds on the con-  
 713 ditional and average treatment effect with unobserved confounding factors. *Annals of statistics*,  
 714 50(5):2587, 2022.

715 Shu Yang, Linbo Wang, and Peng Ding. Causal inference with confounders missing not at random.  
 716 *Biometrika*, 106(4):875–888, 2019.

718 Jinsung Yoon, James Jordon, and Mihaela Van Der Schaar. Ganite: Estimation of individualized  
 719 treatment effects using generative adversarial nets. In *International conference on learning rep-  
 720 resentations*, 2018.

721 Linying Zhang, Yixin Wang, Anna Ostropolets, Jami J Mulgrave, David M Blei, and George Hripc-  
 722 sak. The medical deconfounder: assessing treatment effects with electronic health records. In  
 723 *Machine Learning for Healthcare Conference*, pp. 490–512. PMLR, 2019.

724 Yi Zhang, Melody Huang, and Kosuke Imai. Minimax regret estimation for generalizing heteroge-  
 725 neous treatment effects with multisite data. *arXiv preprint arXiv:2412.11136*, 2024.

727 Zhiheng Zhang, Quanyu Dai, Xu Chen, Zhenhua Dong, and Ruiming Tang. Robust causal inference  
 728 for recommender system to overcome noisy confounders. In *Proceedings of the 46th International  
 729 ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2349–2353,  
 730 2023.

## 731 A APPENDIX

### 732 A.1 CATE ESTIMATION WITH CARD

736 We now detail the construction of a CATE learner under the CARD framework, leveraging the  
 737 observed samples  $\{(X_i, A_i, Y_i)\}_{i=1}^n$ . Since the CATE learner is trained on the training dataset, the  
 738 sample size  $n$  here corresponds to the size of the training sample. We denote  $n_t$  as the sample size  
 739 of the treatment group and  $n_c$  as that of the control group, with  $n = n_t + n_c$ . A key component  
 740 of the CARD framework is a time-dependent diffusion model  $g_\theta(z, t)$ , which takes as inputs a time  
 741 step  $t \sim \mathcal{U}[\epsilon, T]$  and a noise variable  $z \sim \mathcal{N}(0, 1)$ . The diffusion model's reverse process initiates at  
 742 time step  $T$  and progresses iteratively. Upon reaching time step 0, it generates the latent variables  $Z_0$ .

- 743 • **S-learner:** Let the predictors be  $(X, A)$  and the response be  $Y$ . We first initialize the  
 744 model  $\hat{\mu}(X, A)$  and then, under the CARD framework, employ  $\|Y - \hat{\mu}(X \oplus Z_0, A)\|_2^2$  as  
 745 both the loss function and reward function to co-optimize  $\hat{\mu}(X, A)$  and the score-based  
 746 diffusion model  $g_\theta$  through an alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we  
 747 obtain  $\hat{\tau}_S(X)$ :

$$748 \hat{\tau}_S(X) = \hat{\mu}(X \oplus Z_0, 1) - \hat{\mu}(X \oplus Z_0, 0).$$

- 749 • **T-learner:** Let the predictors be  $X^T$  (covariates in the treatment) and the response be  $Y^T$   
 750 (outcome in the treatment). Let the predictors be  $X^C$  (covariates in the control) and  
 751 the response be  $Y^C$  (outcome in the control). We first initialize the treatment outcome  
 752 model  $\hat{\mu}_1(X^T)$  and control outcome model  $\hat{\mu}_0(X^C)$ . Under the CARD framework, we then  
 753 employ  $\|Y^T - \hat{\mu}_1(X^T \oplus Z_0)\|_2^2 + \|Y^C - \hat{\mu}_0(X^C \oplus Z_0)\|_2^2$  as both the loss function and  
 754 reward function to co-optimize  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ , and the diffusion model  $g_\theta$  through an alternating  
 755 training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_T(X)$ :

$$756 \hat{\tau}_T(X) = \hat{\mu}_1(X \oplus Z_0) - \hat{\mu}_0(X \oplus Z_0).$$

756 • **X-learner:** First-step: Initialize  $\hat{\mu}_1(X)$  and  $\hat{\mu}_0(X)$  using the the above-mentioned procedure in T-learner. Let the predictors be  $X$  and the response be  $A$ . Initialize a propensity score model  $\hat{\pi}(X)$ . Second-step: Let the predictors be  $X^T$  and the response be  $\hat{\mu}_1(X^T) - Y^T$ . Let the predictors be  $X^C$  and the response be  $\hat{\mu}_0(X^C) - Y^C$ . Using these defined predictors and responses, we initialize the models  $\hat{\tau}_1(X^T)$  and  $\hat{\tau}_0(X^C)$ . Next, we utilize  $\|Y^T - \hat{\mu}_1(X^T \oplus Z_0)\|_2^2 + \|Y^C - \hat{\mu}_0(X^C \oplus Z_0)\|_2^2 + \text{CrossEntropyLoss}(X \oplus Z, A) + \|\hat{\mu}_1(X^T \oplus Z) - Y^T - \hat{\tau}_1(X^T \oplus Z_0)\|_2^2 + \|\hat{\mu}_0(X^C \oplus Z) - Y^C - \hat{\tau}_0(X^C \oplus Z_0)\|_2^2$  as both the loss function and reward function to co-optimize  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ ,  $\hat{\pi}$ ,  $\hat{\tau}_1$ ,  $\hat{\tau}_0$  and the diffusion model  $g_\theta$  through an alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_X(X)$ :

$$\hat{\tau}_X(X) = (1 - \hat{\pi}(X \oplus Z_0))\hat{\tau}_1(X^T \oplus Z_0) - \hat{\pi}(X \oplus Z_0)\hat{\tau}_0(X^C \oplus Z_0).$$

767 • **R-learner:** First-step: Let the predictors be  $X$  and the response be  $Y$ . Initialize a model  $\hat{\mu}(X)$  to approximate the conditional mean outcome  $\mathbb{E}[Y|X]$ . Initialize a propensity score model  $\hat{\pi}(X)$  using the the above-mentioned procedure in X-learner. Second-step: Compute the outcome residual  $\xi = Y - \hat{\mu}(X)$  and treatment residual  $\nu = T - \hat{\pi}(X)$ . We then initialize a model  $\hat{\tau}(X)$ . Under the CARD framework, we utilize  $\|Y - \hat{\mu}(X \oplus Z_0)\|_2^2 + \text{CrossEntropyLoss}(X \oplus Z_0, A) + \|\xi - \nu\hat{\tau}(X \oplus Z_0)\|_2^2$  as both the loss function and reward function to co-optimize  $\hat{\mu}$ ,  $\hat{\pi}$ ,  $\hat{\tau}$  and the diffusion model  $g_\theta$  through an alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_R(X)$ :

$$\hat{\tau}_R(X) = \hat{\tau}(X \oplus Z_0).$$

778 • **DR-learner:** First-step: Initialize  $\hat{\mu}_1(X)$  and  $\hat{\mu}_0(X)$  using the the above-mentioned procedure in T-learner. Initialize a propensity score model  $\hat{\pi}(X)$  using the the above-mentioned procedure in X-learner. Second-step: Construct surrogate of CATE using pseudo-outcomes with doubly robust (DR) formula:  $Y_{DR}^{0,1} = Y_{DR}^1 - Y_{DR}^0$ , where  $Y_{DR}^1 = \hat{\mu}_1(X) + \frac{T}{\hat{\pi}(X)}(Y - \hat{\mu}_1(X))$  and  $Y_{DR}^0 = \hat{\mu}_0(X) + \frac{1-T}{1-\hat{\pi}(X)}(Y - \hat{\mu}_0(X))$ . Using these defined predictors and responses, we initialize the models  $\hat{\tau}(X)$ . Next, we utilize  $\|Y^T - \hat{\mu}_1(X^T \oplus Z_0)\|_2^2 + \|Y^C - \hat{\mu}_0(X^C \oplus Z_0)\|_2^2 + \text{CrossEntropyLoss}(X \oplus Z_0, A) + \|Y_{DR}^{0,1} - \hat{\tau}(X \oplus Z_0)\|_2^2$  as both the loss function and reward function to co-optimize  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ ,  $\hat{\pi}$ ,  $\hat{\tau}$  and the diffusion model  $g_\theta$  through an alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_{DR}(X)$ :

$$\hat{\tau}_{DR}(X) = \hat{\tau}(X \oplus Z_0).$$

791 • **TARNet:** We first define the predictors as  $(X, A)$  and the response as  $Y$ , and construct 792 a representation model  $\hat{r}(X)$  to encode covariate information. The model architecture 793 incorporates two outcome heads:  $\hat{\mu}_1(\hat{r}(X))$  for the treatment group and  $\hat{\mu}_0(\hat{r}(X))$  for 794 the control group, which share the underlying representation  $\hat{r}(X)$  while learning separate 795 outcome estimates. Under the CARD framework, we employ the composite 796 function  $\|\hat{\mu}_1(\hat{r}(X \oplus Z_0)) - Y^T\|_2^2 + \|\hat{\mu}_0(\hat{r}(X \oplus Z_0)) - Y^C\|_2^2$  as both the loss function 797 and reward function to co-optimize  $\hat{r}(X)$ ,  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ , and the diffusion model  $g_\theta$  through an 798 alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_{TARNet}(X)$ :

$$\hat{\tau}_{TARNet}(X) = \hat{\mu}_1(\hat{r}(X \oplus Z_0)) - \hat{\mu}_0(\hat{r}(X \oplus Z)).$$

801 • **CFR\_WASS:** Initialize  $\hat{r}(X)$ ,  $\hat{\mu}_1(X)$  and  $\hat{\mu}_0(X)$  using the the above-mentioned 802 procedure in TARNet. Under the CARD framework, we employ the composite 803 function  $\|\hat{\mu}_1(\hat{r}(X \oplus Z_0)) - Y^T\|_2^2 + \|\hat{\mu}_0(\hat{r}(X \oplus Z_0)) - Y^C\|_2^2 + \text{IPMLoss}(X^T, X^C)$  as 804 both the loss function and reward function to co-optimize  $\hat{r}(X)$ ,  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ , and the diffusion 805 model  $g_\theta$  through an alternating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain 806  $\hat{\tau}_{CFR_{WASS}}(X)$ :

$$\hat{\tau}_{CFR-Wass}(X) = \hat{\mu}_1(\hat{r}(X \oplus Z_0)) - \hat{\mu}_0(\hat{r}(X \oplus Z)).$$

807 • **DragonNet:** Initialize  $\hat{r}(X)$ ,  $\hat{\mu}_1(X)$  and  $\hat{\mu}_0(X)$  using the the above-mentioned procedure 808 in TARNet. The model architecture incorporates three outcome heads:  $\hat{\mu}_1(\hat{r}(X))$  for the 809

810 treatment group,  $\hat{\mu}_0(\hat{r}(X))$  for the control group and  $\hat{\pi}(X)$  for the propensity score, which  
 811 share the underlying representation  $\hat{r}(X)$  while learning separate outcome estimates. Under  
 812 the CARD framework, we employ the composite function  $\|\hat{\mu}_1(\hat{r}(X \oplus Z_0)) - Y^T\|_2^2 +$   
 813  $\|\hat{\mu}_0(\hat{r}(X \oplus Z_0)) - Y^C\|_2^2 + \text{CrossEntropyLoss}(X \oplus Z_0, A)$  as both the loss function and  
 814 reward function to co-optimize  $\hat{r}(X)$ ,  $\hat{\mu}_1$ ,  $\hat{\mu}_0$ , and the diffusion model  $g_\theta$  through an alter-  
 815 nating training process. Using  $Z_0$  generated by  $g_\theta$ , we obtain  $\hat{\tau}_{\text{DragonNet}}(X)$ :

$$\hat{\tau}_{\text{DragonNet}}(X) = \hat{\mu}_1(\hat{r}(X \oplus Z_0)) - \hat{\mu}_0(\hat{r}(X \oplus Z_0)).$$

## 819 A.2 EXPERIMENTAL DETAILS AND HYPERPARAMETERS

821 **Implementation details.** All meta-learners in this work are implemented using neural network  
 822 architectures. Specifically, the S-learner, T-learner, X-learner, R-learner, and DR-learner share a  
 823 unified three-layer neural network structure, with each layer containing 200 neurons. In contrast,  
 824 TARNet, CFR-Wass, and DragonNet adopt a two-component architecture: a representation network  
 825 with three layers (200 neurons per layer) and a prediction layer with three layers (100 neurons per  
 826 layer).

828 **Hyperparameters.** All model training processes are conducted on a Dell 3640 workstation with  
 829 an Intel Xeon W-1290P 3.60GHz CPU and NVIDIA GeForce RTX 2080 Ti GPU. For optimizing  
 830 the CATE learner, we used the Adam optimizer with a learning rate of  $10^{-3}$  and weight decay  
 831 of  $10^{-4}$ . Model selection was based on the factual loss as the validation metric, with early stopping  
 832 implemented if no improvement was observed on the validation set for 20 consecutive epochs. In the  
 833 reinforcement fine-tuning phase, the AdamW optimizer was employed with a learning rate of  $2 \times$   
 834  $10^{-5}$ . The hyperparameters for fine-tuning varied by model type:

- 835 • For S-learner, T-learner, and DragonNet:  $\alpha = 0.8$  and fine-tuning frequency  $K = 10$ ;
- 836 • For X-learner, R-learner, and DR-learner:  $\alpha = 0.8$  and  $K = 2$ ;
- 837 • For TARNet and CFR\_WASS:  $\alpha = 0.1$  and  $0.8$ , with  $K = 10$  and  $5$  respectively.

840 Additionally, the imbalance loss coefficient for CFR-Wass and the BCE loss coefficient for Dragon-  
 841 Net were both set to 1.0. The discount factor  $\gamma$  was set to 0.99. The latent variable  $Z$  generated  
 842 by the diffusion model has a dimension half that of the covariate  $X$ .

844 **Model architecture.** The parameters of the autoencoder and score-based diffusion model largely  
 845 follow the default settings provided in (Suh et al.). Both models are trained for 10,000 epochs,  
 846 and the number of timesteps for the diffusion model is set to 50. The autoencoder adopts a multi-  
 847 layer perceptron (MLP) block-based architecture, with ReLU activation functions used in all hidden  
 848 layers. Its forward process is defined as:

$$\begin{aligned} 849 \text{MLPBlock}(X) &= \text{ReLU}(\text{Linear}(X)), \\ 850 Z &= \text{Linear}(\dots \text{MLPBlock}(X)), \\ 851 \tilde{X} &= \text{Linear}(\dots \text{MLPBlock}(Z)), \end{aligned} \tag{9}$$

853 where  $Z$  denotes the latent representation of the input  $X$ , and  $\tilde{X}$  is the reconstruction output of the  
 854 autoencoder.

856 Let  $t$  denote a timestep in the diffusion process, and SinTimeEmb represent the sinusoidal time  
 857 embedding proposed in (Nichol & Dhariwal, 2021). For any fixed  $t$ , the time embedding  $t^{\text{emb}}$  and  
 858 the processed input to the score network (denoted  $Z^{t\text{-emb}}$ ) are computed as:

$$\begin{aligned} 859 t^{\text{emb}} &= \text{LayerNorm}(\text{SiLU}(\text{Linear}(\text{SinTimeEmb}(t)))), \\ 860 Z^{t\text{-emb}} &= \text{LayerNorm}(\text{Linear}(Z_t)) + t^{\text{emb}}, \end{aligned} \tag{10}$$

863 where  $Z_t$  is the latent variable at timestep  $t$ , and the addition of  $t^{\text{emb}}$  injects timestep-aware information  
 864 into the latent input.

864 The time-dependent score network  $g$  is then constructed using MLP blocks with LayerNorm regu-  
 865 larization, and its calculation is given by:  
 866

$$\begin{aligned} 867 \text{MLPBlock}(Z^{t\text{-emb}}) &= \text{LayerNorm}(\text{ReLU}(\text{Linear}(Z^{t\text{-emb}}))), \\ 868 g(Z^{t\text{-emb}}, t) &= \text{Linear}(\dots \text{MLPBlock}(Z^{t\text{-emb}})). \end{aligned} \quad (11)$$

870 **Time complexity analysis.** We compare the time complexity between CATE learners trained with  
 871 the CARD framework and those trained with standard procedures. We assume the CATE learner is  
 872 trained with  $E$  epochs, resulting in a complexity of  $O(E)$ . In contrast, when training a CATE learner  
 873 using the CARD framework, each training epoch requires an additional  $T$  iterations for trajectory  
 874 generation in the diffusion model, leading to a time complexity of  $O(ET)$ . Thus, the improved  
 875 robustness of the CATE learner achieved via the CARD framework comes at the cost of increased  
 876 computational time, i.e., a tradeoff between model robustness and time cost. The practical users are  
 877 suggested to set  $T$  with early stop, and use new acceleration technique for training diffusion models.  
 878

### 879 A.3 ROLE OF LLM

880 In this paper, LLM was used to aid in writing and polish the texts. Importantly, we take full re-  
 881 sponsibility for the content of the manuscript, and we did not use LLM for idea generation, method  
 882 development, experimental coding. All research ideas, codes, experimental results, and experimen-  
 883 tal analysis are conducted by the authors. The contribution of LLM is only the linguistic quality  
 884 improvement.  
 885  
 886  
 887  
 888  
 889  
 890  
 891  
 892  
 893  
 894  
 895  
 896  
 897  
 898  
 899  
 900  
 901  
 902  
 903  
 904  
 905  
 906  
 907  
 908  
 909  
 910  
 911  
 912  
 913  
 914  
 915  
 916  
 917