

ESTIMATING INDIVIDUAL TREATMENT EFFECTS UNDER UNOBSERVED CONFOUNDING USING BINARY INSTRUMENTS

Dennis Frauen^{1,2} & Stefan Feuerriegel^{1,2}

¹ Munich Center for Machine Learning

² LMU Munich

ABSTRACT

Estimating conditional average treatment effects (CATEs) from observational data is relevant in many fields such as personalized medicine. However, in practice, the treatment assignment is usually confounded by unobserved variables and thus introduces bias. A remedy to remove the bias is the use of instrumental variables (IVs). Such settings are widespread in medicine (e.g., trials where the treatment assignment is used as binary IV). In this paper, we propose a novel, multiply robust machine learning framework, called MRIV, for estimating CATEs using binary IVs and thus yield an unbiased CATE estimator. Different from previous work for binary IVs, our framework estimates the CATE directly via a pseudo-outcome regression. (1) We provide a theoretical analysis where we show that our framework yields multiple robust convergence rates: our CATE estimator achieves fast convergence even if several nuisance estimators converge slowly. (2) We further show that our framework asymptotically outperforms state-of-the-art plug-in IV methods for CATE estimation, in the sense that it achieves a faster rate of convergence if the CATE is smoother than the individual outcome surfaces. (3) We build upon our theoretical results and propose a tailored deep neural network architecture called MRIV-Net for CATE estimation using binary IVs. Across various computational experiments, we demonstrate empirically that our MRIV-Net achieves state-of-the-art performance. To the best of our knowledge, our MRIV is the first multiply robust machine learning framework tailored to estimating CATEs in the binary IV setting.

1 INTRODUCTION

Conditional average treatment effects (CATEs) are relevant across many disciplines such as marketing (Varian, 2016) and personalized medicine (Yazdani & Boerwinkle, 2015). Knowledge about CATEs provides insights into the heterogeneity of treatment effects, and thus helps in making potentially better treatment decisions (Frauen et al., 2023).

Many recent works that use machine learning to estimate causal effects, in particular CATEs, are based on the assumption of unconfoundedness (Alaa & van der Schaar, 2017; Lim et al., 2018; Melnychuk et al., 2022a;b). In practice, however, this assumption is often violated because it is common that some confounders are not reported in the data. Typical examples are income or the socioeconomic status of patients, which are not stored in medical files. If the confounding is sufficiently strong, standard methods for estimating CATEs suffer from confounding bias (Pearl, 2009), which may lead to inferior treatment decisions.

To handle unobserved confounders, instrumental variables (IVs) can be leveraged to relax the assumption of unconfoundedness and still compute reliable CATE estimates. IV methods were originally developed in economics (Wright, 1928), but, only recently, there is a growing interest in combining IV methods with machine learning (see Sec. 3). Importantly, IV methods outperform classical CATE estimators if a sufficient amount of confounding is not observed (Hartford et al., 2017). We thus aim at estimating CATEs from observational data under unobserved confounding using IVs.

In this paper, we consider the setting where a single binary instrument is available. This setting is widespread in personalized medicine (and other applications such as marketing or public policy) (Bloom et al., 1997). In fact, the setting is encountered in essentially all observational or randomized studies with observed non-compliance (Imbens & Angrist, 1994). As an example, consider a randomized controlled trial (RCT), where treatments are randomly assigned to patients and their outcomes are observed. Due to some potentially unobserved confounders (e.g., income, education), some patients refuse to take the treatment initially assigned to them. Here, the treatment assignment serves as a binary IV. Moreover, such RCTs have been widely used by public decision-makers, e.g., to analyze the effect of health insurance on health outcome (see the so-called *Oregon health insurance experiment*) (Finkelstein et al., 2012) or the effect of military service on lifetime earnings (Angrist, 1990).

We propose a novel machine learning framework (called MRIV) for estimating CATEs using binary IVs. Our framework takes an initial CATE estimator and nuisance parameter estimators as input to perform a pseudo-outcome regression. Different to existing literature, our framework is **multiply robust**¹, i.e., we show that it is consistent in the union of three different model specifications. This is different from existing methods for CATE estimation using IVs such as Okui et al. (2012), Syrgkanis et al. (2019), or plug-in estimators (Bargagli-Stoffi et al., 2021; Imbens & Angrist, 1994).

We provide a theoretical analysis, where we use tools from Kennedy (2022) to show that our framework achieves a multiply robust convergence rate, i.e., our MRIV converges with a fast rate even if several nuisance parameters converge slowly. We further show that, compared to existing plug-in IV methods, the performance of our framework is asymptotically superior. Finally, we leverage our framework and, on top of it, build a tailored deep neural network called MRIV-Net.

Contributions: (1) We propose a novel, multiply robust machine learning framework (called MRIV) to learn the CATE using the binary IV setting. To the best of our knowledge, ours is the first that is shown to be multiply robust, i.e., consistent in the union of three model specifications. For comparison, existing works for CATE estimation only show double robustness (Wang & Tchetgen Tchetgen, 2018; Syrgkanis et al., 2019). (2) We prove that MRIV achieves a multiply robust convergence rate. This is different to methods for IV settings which do not provide robust convergence rates (Syrgkanis et al., 2019). We further show that our MRIV is asymptotically superior to existing plug-in estimators. (3) We propose a tailored deep neural network, called MRIV-Net, which builds upon our framework to estimate CATEs. We demonstrate that MRIV-Net achieves state-of-the-art performance.

2 PROBLEM SETUP

Data generating process: We observe data $\mathcal{D} = (x_i, z_i, a_i, y_i)_{i=1}^n$ consisting of $n \in \mathbb{N}$ observations of the tuple (X, Z, A, Y) . Here, $X \in \mathcal{X}$ are observed confounders, $Z \in \{0, 1\}$ is a binary instrument, $A \in \{0, 1\}$ is a binary treatment, and $Y \in \mathbb{R}$ is an outcome of interest. Furthermore, we assume the existence of unobserved confounders $U \in \mathcal{U}$, which affect both the treatment A and the outcome Y . The causal graph is shown in Fig. 1.

Applicability: Our proposed framework is widely applicable in practice, namely to all settings with the above data generating process. This includes both (1) observational data and (2) RCTs with non-compliance. For (1), observational data is commonly encountered in, e.g., personalized medicine. Here, modeling treatments as binary variables is consistent with previous literature on causal effect estimation and standard in medical practice (Robins et al., 2000). For (2), our setting is further encountered in RCTs when the instrument Z is a randomized treatment assignment but individuals do not comply with their treatment assignment. Such RCTs have been extensively used by public decision-makers, e.g., to analyze the effect of health insurance on health outcome (Finkelstein et al., 2012) or the effect of military service on lifetime earnings (Angrist, 1990).

We build upon the potential outcomes framework (Rubin, 1974) for modeling causal effects. Let $Y(a, z)$ denote the potential outcome that




Figure 1: Underlying causal graph. The instrument Z has a direct influence on the treatment A , but does not have a direct effect on the outcome Y . Note that we allow for unobserved confounders for both Z - A (dashed line) and A - Y (given by U).

¹For a detailed introduction to multiple robustness and its importance in treatment effect estimation, we refer to (Wang & Tchetgen Tchetgen, 2018), Section 4.5.

would have been observed under $A = a$ and $Z = z$. Following previous literature on IV estimation (Wang & Tchetgen Tchetgen, 2018), we impose the following standard IV assumptions on the data generating process.

Assumption 1 (Standard IV assumptions (Wang & Tchetgen Tchetgen, 2018; Wooldridge, 2013)). We assume: (1) *Exclusion*: $Y(a, z) = Y(a)$ for all $a, z \in \{0, 1\}$, i.e., the instrument has no direct effect on the patient outcome; (2) *Independence*: $Z \perp\!\!\!\perp U \mid X$; (3) *Relevance*: $Z \not\perp\!\!\!\perp A \mid X$, (iv) *The model includes all A–Y confounder*: $Y(a) \perp\!\!\!\perp (A, Z) \mid (X, U)$ for all $a \in \{0, 1\}$.

Assumption 1 is standard for IV methods and fulfilled in practical settings where IV methods are applied (Angrist, 1990; Angrist & Krueger, 1991; Imbens & Angrist, 1994). Note that Assumption 1 does not prohibit the existence of unobserved Z – A confounders. On the contrary, it merely prohibits the existence of unobserved counfounders that affect all Z , A , and Y simultaneously, as it is standard in IV settings (Wooldridge, 2013). A practical and widespread example where Assumption 1 is satisfied are randomized controlled trials (RCTs) with non-compliance (Imbens & Angrist, 1994). Here, the treatment assignment Z is randomized, but the actual relationship between treatment A and outcome Y may still be confounded. For instance, in the *Oregon health insurance experiment* (Finkelstein et al., 2012), people were given access to health insurance (Z) by a lottery with aim to study the effect of health insurance (A) on health outcome (Y) (Finkelstein et al., 2012). Here, the lottery winners needed to sign up for health insurance and thus both Z and A are observed. **Objective:** In this paper, we are interested in estimating the *conditional average treatment effect* (CATE)

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

If there is no unobserved confounding ($U = \emptyset$), the CATE is identifiable from observational data (Shalit et al., 2017). However, in practice, it is often unlikely that all confounders are observable. To account for this, we leverage the instrument Z to identify the CATE. We state the following assumption for identifiability.

Assumption 2 (Identifiability of the CATE (Wang & Tchetgen Tchetgen, 2018)). At least one of the following two statements holds true: (1) $\mathbb{E}[A \mid Z = 1, X, U] - \mathbb{E}[A \mid Z = 0, X, U] = \mathbb{E}[A \mid Z = 1, X] - \mathbb{E}[A \mid Z = 0, X]$; or (2) $\mathbb{E}[Y(1) - Y(0) \mid X, U] = \mathbb{E}[Y(1) - Y(0) \mid X]$.

Example: Assumption 1 holds when the function $f(a, X, U) = \mathbb{E}[Y(a) \mid X, U]$ is additive with respect to a and U , e.g., $f(a, X, U) = g(a, X) + h(U)$ for measurable functions h and g . This implies that no unobserved confounder affects the outcome through a path which is also affected by the treatment. For example, with patient income as unobserved confounder, the treatment should not affect the (future) patient income.

Under Assumptions 1 and 2, the CATE is identifiable (Wang & Tchetgen Tchetgen, 2018). It can be written as

$$\tau(x) = \frac{\mu_1^Y(x) - \mu_0^Y(x)}{\mu_1^A(x) - \mu_0^A(x)} = \frac{\delta_Y(x)}{\delta_A(x)}, \quad (2)$$

where $\mu_i^Y(x) = \mathbb{E}[Y \mid Z = i, X = x]$ and $\mu_i^A(x) = \mathbb{E}[A \mid Z = i, X = x]$. Even if Assumption 2 does not hold, all our results in this paper still hold for the quantity on the right-hand side of Eq. (2). In certain cases, this quantity still allows for interpretation: If no unobserved Z – A confounders exist, it can be interpreted as conditional version of the *local average treatment effect* (LATE) (Imbens & Angrist, 1994) under a monotonicity assumption. Furthermore, under a no-current-treatment-value-interaction assumption, it can be interpreted as conditional *treatment effect on the treated* (ETT) (Wang & Tchetgen Tchetgen, 2018).² This has an important implication for our results: If Assumption 2 does not hold in practice, our estimates still provide conditional LATE or ETT estimates under the respective assumptions because they are based on Eq. (2). If Assumption 2 does hold, all three – i.e., CATE, conditional LATE, and ETT – coincide (Wang & Tchetgen Tchetgen, 2018).

3 RELATED WORK

Machine learning methods for IV: Only recently, machine learning has been integrated into IV methods. These are: Singh et al. (2019) and Xu et al. (2021a) generalize 2SLS by learning complex

²The conditional LATE measures the CATE for individuals which are part of the complier subpopulation, i.e., for whom $A(Z = 1) > A(Z = 0)$. The conditional ETT measures the CATE for treated individuals.

feature maps using kernel methods and deep learning, respectively. Hartford et al. (2017) adopts a two-stage neural network architecture that performs the first stage via conditional density estimation. Bennett et al. (2019) leverages moment conditions for IV estimation. However, the aforementioned methods are not specifically designed for the binary IV setting but, rather, for multiple IVs or treatment scenarios. In particular, they impose stronger assumptions such as additive confounding in order to identify the CATE. Note that additive confounding is a special case of our Assumption 2. Moreover, they do not have robustness properties. In the binary IV setting, current methods proceed by estimating $\mu_i^Y(x)$ and $\mu_i^A(x)$ separately, before plugging them in Eq. 2 (Imbens & Angrist, 1994; Angrist et al., 1996; Bargagli-Stoffi et al., 2021). As a result, these suffer from plug-in bias and do *not* offer robustness properties.

Doubly robust IV methods: Recently, doubly robust methods have been proposed for IV settings: Kennedy et al. (2019) propose a pseudo regression estimator for the local average treatment effect using continuous instruments, which has been extended to conditional effects by Semenova & Chernozhukov (2021). Furthermore, Singh & Sun (2019) use a doubly robust approach to estimate average compiler parameters. Finally, Ogburn et al. (2015) and Syrgkanis et al. (2019) propose doubly robust CATE estimators in the IV setting which both rely on doubly robust parametrizations of the uncentered efficient influence function (Okui et al., 2012). However, none of these estimators has been shown to be multiply robust in the sense that they are consistent in the union of more than two model specifications (Wang & Tchetgen Tchetgen, 2018).

Multiply robust IV methods: Multiply robust estimators for IV settings have been proposed only for average treatment effects (ATEs) (Wang & Tchetgen Tchetgen, 2018) and optimal treatment regimes (Cui & Tchetgen, 2021) but not for CATEs. In particular, Wang & Tchetgen Tchetgen (2018) derive a multiply robust parametrization of the efficient influence function for the ATE. However, there exists no method that leverages this result for CATE estimation. We provide a detailed, technical comparison of existing methods and our framework in Appendix G.

Doubly robust rates for CATE estimation: Kennedy (2022) analyzed the doubly robust learner in the standard (non-IV) setting and derived doubly robust convergence rates. However, Kennedy’s result is not applicable in the IV setting, because we use the multiply robust parametrization of the efficient influence function from Wang & Tchetgen Tchetgen (2018). In our paper, we rely on certain results from Kennedy, but use these to derive of a multiply robust rate. In particular, this required the derivation of the bias term for a larger number of nuisance parameters (see Appendix B).

Research gap: To the best of our knowledge, there exists no method for CATE estimation under unobserved confounding that has been shown to be *multiply robust*. To fill this gap, we propose MRIV: a *multiply robust* machine learning framework tailored to the binary IV setting. For this, we build upon the approach by Kennedy (2022) to derive robust convergence rates, yet this approach has not been adapted to IV settings, which is our contribution.

Table 1: Key methods for causal effect estimation with IVs and associated robustness properties

	Estimand	ATE	ITE
Doubly robust		Okui et al. (2012)	Syrgkanis et al. (2019)
Multiply robust		Wang & Tchetgen Tchetgen (2018)	MRIV (ours)

4 MRIV FOR ESTIMATING CATES USING BINARY INSTRUMENTS

In the following, we present our MRIV framework for estimating CATEs under unobserved confounding (Sec. 4.1). We then derive an asymptotic convergence rate for MRIV (Sec. 4.2) and finally use our framework to develop a tailored deep neural network called MRIV-Net (Sec. 4.4).

4.1 FRAMEWORK

Motivation: A naïve approach to estimate the CATE is to leverage the identification result in Eq. (2). Assuming that we have estimated the nuisance components $\hat{\mu}_i^Y$ and $\hat{\mu}_i^A$ for $i \in \{0, 1\}$, we can simply plug them into Eq. (2) to obtain the so-called (plug-in) Wald estimator $\hat{\tau}_W(x)$ (Wald, 1940).

However, in practice, the true CATE curve $\tau(x)$ is often simpler (e.g., smoother, more sparse) than its complements $\mu_i^Y(x)$ or $\mu_i^A(x)$ (Künzel et al., 2019). In this case, $\hat{\tau}_W(x)$ is inefficient because it models all components separately, and, to address this, our proposed framework estimates τ directly using a pseudo-outcome regression.

Overview: We now propose MRIV. MRIV is a two-stage meta learner that takes any base method for CATE estimation as input. For instance, the base method could be the Wald estimator from Eq. (2), any other IV method such as 2SLS, or a deep neural network (as we propose in our MRIV-Net later in Sec. 4.4). In Stage 1, MRIV produces nuisance estimators $\hat{\mu}_0^Y(x)$, $\hat{\mu}_0^A(x)$, $\hat{\delta}_A(x)$, and $\hat{\pi}(x)$, where $\hat{\pi}(x)$ is an estimator of the propensity score $\pi(x) = \mathbb{P}(Z = 1 | X = x)$. In Stage 2, MRIV estimates $\tau(x)$ directly using a pseudo outcome \hat{Y}_{MR} as a regression target.

Given an arbitrary initial CATE estimator $\hat{\tau}_{\text{init}}(x)$ and nuisance estimates $\hat{\mu}_0^Y(x)$, $\hat{\mu}_0^A(x)$, $\hat{\delta}_A(x)$, and $\hat{\pi}(x)$, we define the pseudo outcome

$$\hat{Y}_{\text{MR}} = \left(\frac{Z - (1 - Z)}{\hat{\delta}_A(X)} \right) \left(\frac{Y - (\hat{\mu}_0^Y(X) + \hat{\tau}_{\text{init}}(X)(A - \hat{\mu}_0^A(X)))}{Z \hat{\pi}(X) + (1 - Z)(1 - \hat{\pi}(X))} \right) + \hat{\tau}_{\text{init}}(X). \quad (3)$$

Algorithm 1: MRIV

Input: data (X, Z, A, Y) , initial CATE estimator $\hat{\tau}_{\text{init}}(x)$
 // Stage 1: Estimate nuisance components
 $\hat{\pi}(x) \leftarrow \hat{\mathbb{E}}[Z | X = x]$,
 $\hat{\mu}_0^Y(x) \leftarrow \hat{\mathbb{E}}[Y | X = x, Z = 0]$,
 $\hat{\mu}_0^A(x) \leftarrow \hat{\mathbb{E}}[A | X = x, Z = 0]$
 $\hat{\delta}_A(x) \leftarrow \hat{\mathbb{E}}[A | X = x, Z = 1] - \hat{\mathbb{E}}[A | X = x, Z = 0]$
 // Stage 2: pseudo-outcome regression
 $\hat{Y}_{\text{MR}} \leftarrow$
 $\left(\frac{Z - (1 - Z)}{\hat{\delta}_A(X)} \right) \left(\frac{Y - A \hat{\tau}_{\text{init}}(X) - \hat{\mu}_0^Y(X) + \hat{\mu}_0^A(X) \hat{\tau}_{\text{init}}(X)}{Z \hat{\pi}(X) + (1 - Z)(1 - \hat{\pi}(X))} \right) +$
 $\hat{\tau}_{\text{init}}(X)$
 $\hat{\tau}_{\text{MRIV}}(x) \leftarrow \hat{\mathbb{E}}[\hat{Y}_{\text{MR}} | X = x]$

The pseudo outcome \hat{Y}_{MR} in Eq. (3) is a multiply robust parameterization of the (uncentered) efficient influence function for the average treatment effect $\mathbb{E}_X[\tau(X)]$ (see the derivation in (Wang & Tchetgen Tchetgen, 2018)). Once we have obtained the pseudo outcome \hat{Y}_{MR} , we regress it on X to obtain the Stage 2 MRIV estimator $\hat{\tau}_{\text{MRIV}}(x)$ for $\tau(x)$. The pseudocode for MRIV is given in Algorithm 1. MRIV can be interpreted as a way to remove plug-in bias from $\hat{\tau}_{\text{init}}(x)$ (Curth et al., 2020). Using the fact that \hat{Y}_{MR} is a multiply robust parametrization of the efficient influence function, we derive a multiple robustness property of $\hat{\tau}_{\text{MRIV}}(x)$.

Theorem 1 (multiple robustness property). *Let $\hat{\mu}_0^Y(x)$, $\hat{\mu}_0^A(x)$, $\hat{\delta}_A(x)$, $\hat{\pi}(x)$, and $\hat{\tau}_{\text{init}}(x)$ denote estimators of $\mu_0^Y(x)$, $\mu_0^A(x)$, $\delta_A(x)$, $\pi(x)$, and $\tau(x)$, respectively. Then, for all $x \in \mathcal{X}$, it holds that $\mathbb{E}[\hat{Y}_{\text{MR}} | X = x] = \tau(x)$, if least one of the following conditions is satisfied: (1) $\hat{\mu}_0^Y = \mu_0^Y$, $\hat{\mu}_0^A = \mu_0^A$, and $\hat{\tau}_{\text{init}} = \tau$; or (2) $\hat{\pi} = \pi$ and $\hat{\delta}_A = \delta_A$; or (3) $\hat{\pi} = \pi$ and $\hat{\tau}_{\text{init}} = \tau$.*

The equalities in Theorem 1 are meant to hold almost surely. Consistency of $\hat{\tau}_{\text{MRIV}}(x)$ is a direct consequence: If either the nuisance estimators in (1), (2), or (3) converge to their oracle estimands, $\hat{\tau}_{\text{MRIV}}(x)$ will converge to the true CATE. As a result, our MRIV framework is *multiply robust* in the sense that our estimator, $\hat{\tau}_{\text{MRIV}}(x)$, is consistent in the union of three different model specifications. Importantly, this is different from *doubly robust* estimators which are only consistent in the union of two model specifications (Wang & Tchetgen Tchetgen, 2018). Our MRIV is directly applicable to RCTs with non-compliance: Here, the treatment assignment is randomized and the propensity score $\pi(x)$ is known. Our MRIV framework can be thus adopted by plugging in the known $\pi(x)$ into the pseudo outcome in Eq. (3). Moreover, $\hat{\tau}_{\text{MRIV}}(x)$ is already consistent if either $\hat{\tau}_{\text{init}}(x)$ or $\hat{\delta}_A(x)$ are.

4.2 THEORETICAL ANALYSIS

We derive the asymptotic bound on the convergence rate of MRIV under smoothness assumptions. For this, we define s -smooth functions as functions contained in the Hölder class $\mathcal{H}(s)$, associated with Stone’s minimax rate (Stone, 1980) of $n^{-2s/(2s+p)}$, where p is the dimension of \mathcal{X} .

Assumption 3 (Smoothness). We assume that (1) the nuisance component $\mu_0^Y(\cdot)$ is α -smooth, $\mu_0^A(\cdot)$ is β -smooth, $\pi(\cdot)$ is γ -smooth, and $\delta_A(\cdot)$ is δ -smooth; (2) all nuisance components are estimated with their respective minimax rate of $n^{\frac{-2k}{2k+p}}$, where $k \in \{\alpha, \beta, \gamma, \delta\}$; and (3) the oracle CATE $\tau(\cdot)$ is η -smooth and the initial CATE estimator $\hat{\tau}_{\text{init}}$ converges with rate $r_\tau(n)$. We provide a rigorous definition in Appendix D.

Assumption 3 for smoothness provides us with a way to quantify the difficulty of the underlying nonparametric regression problems. Similar assumptions have been imposed for asymptotic analysis of previous CATE estimators in (Kennedy, 2022; Curth & van der Schaar, 2021). They can be replaced with other assumptions such as assumptions on the level of sparsity of the CATE components. We also provide an asymptotic analysis under sparsity assumptions (see Appendix C).

Assumption 4 (Boundedness). We assume that there exist constants $C, \rho, \tilde{\rho}, \epsilon, K > 0$ such that for all $x \in \mathcal{X}$ it holds that: (1) $|\mu_i^Y(x)| \leq C$; (2) $|\delta_A(x)| = |\mu_1^A(x) - \mu_0^A(x)| \geq \rho$ and $|\hat{\delta}_A(x)| \geq \tilde{\rho}$; (3) $\epsilon \leq \hat{\pi}(x) \leq 1 - \epsilon$; and (4) $|\hat{\tau}_{\text{init}}(x)| \leq K$.

Assumptions 4.1, 4.3, and 4.4 are standard and in line with previous works on theoretical analyses of CATE estimators (Curth & van der Schaar, 2021; Kennedy, 2022). Assumption 4.2 ensures that both the oracle CATE and the estimator are bounded. Violations of Assumption 4.2 may occur when working with “weak” instruments, which are IVs that are only weakly correlated with the treatment. Using IV methods with weak instruments should generally be avoided (Li et al., 2022). However, in many applications such as RCTs with non-compliance, weak instruments are unlikely to occur as patients’ compliance decisions are generally correlated with the initial treatment assignments.

We state now our main theoretical result: an upper bound on the oracle risk of the MRIV estimator. To derive our bound, we leverage the sample splitting approach from (Kennedy, 2022). The approach in (Kennedy, 2022) has been initially used to analyze the DR-learner for CATE estimation under unconfoundedness and allows for the derivation of robust convergence rates. It has later been adapted to several other meta learners (Curth & van der Schaar, 2021), yet not for IV methods.

Theorem 2 (Oracle upper bound under smoothness). *Let \mathcal{D}_ℓ for $\ell \in \{1, 2, 3\}$ be independent samples of size n . Let $\hat{\tau}_{\text{init}}(x)$, $\hat{\mu}_0^Y(x)$, and $\hat{\mu}_0^A(x)$ be trained on \mathcal{D}_1 , and let $\hat{\delta}_A(x)$ and $\hat{\pi}(x)$ be trained on \mathcal{D}_2 . We denote \hat{Y}_{MR} as the pseudo outcome from Eq. (3) and $\hat{\tau}_{\text{MRIV}}(x) = \hat{\mathbb{E}}_n[\hat{Y}_{\text{MR}} | X = x]$ as the pseudo-outcome regression on \mathcal{D}_3 for some generic estimator $\hat{\mathbb{E}}_n[\cdot | X = x]$ of $\mathbb{E}[\cdot | X = x]$.*

We assume that the second-stage estimator $\hat{\mathbb{E}}_n$ yields the minimax rate $n^{-\frac{2\eta}{2\eta+p}}$ and satisfies the stability assumption from Kennedy (2022), Proposition 1 (see Appendix B). Then, under Assumptions 1-4 the oracle risk is upper bounded by

$$\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim n^{\frac{-2\eta}{2\eta+p}} + r_\tau(n) \left(n^{\frac{-2\gamma}{2\gamma+p}} + n^{\frac{-2\delta}{2\delta+p}} \right) + n^{-2\left(\frac{\alpha}{2\alpha+p} + \frac{\gamma}{2\gamma+p}\right)} + n^{-2\left(\frac{\beta}{2\beta+p} + \frac{\gamma}{2\gamma+p}\right)}.$$

Proof. See Appendix B. The proof provides a more general bound which depends on the pointwise mean squared errors of the nuisance parameters (Lemma 2). \square

Recall that the first summand of the lower bound in Eq. (2) is the minimax rate for the oracle CATE $\tau(x)$ which cannot be improved upon. Hence, for a fast convergence rate of $\hat{\tau}_{\text{MRIV}}(x)$, it is sufficient if either: (1) $r_\tau(n)$ decreases fast and α, β are large; (2) γ and δ are large; or (3) $r_\tau(n)$ decreases fast and γ is large. This is in line with the multiply robustness property of MRIV (Theorem 1) and means that MRIV achieves a fast rate even if the initial or several nuisance estimators converge slowly.

Improvement over $\hat{\tau}_{\text{init}}(x)$: From the bound in Eq. (2), it follows that $\hat{\tau}_{\text{MRIV}}(x)$ improves on the convergence rate of the initial CATE estimator $\hat{\tau}_{\text{init}}(x)$ if its rate $r_\tau(n)$ is lower bounded by

$$r_\tau(n) \gtrsim n^{\frac{-2\eta}{2\eta+p}} + n^{-2\left(\frac{\alpha}{2\alpha+p} + \frac{\gamma}{2\gamma+p}\right)} + n^{-2\left(\frac{\beta}{2\beta+p} + \frac{\gamma}{2\gamma+p}\right)}. \quad (4)$$

Hence, our MRIV estimator is more likely to improve on the initial estimator $\hat{\tau}_{\text{init}}(x)$ if either (1) γ is large or (2) α and β are large. Note that the margin of improvement depends also on the size of γ and δ , i.e., on the smoothness of $\pi(x)$ and $\delta_A(x)$. In fact, this is widely fulfilled in practice. For example, the former is fulfilled for RCTs with non-compliance, where $\pi(x)$ is often some known, fixed number $p \in (0, 1)$.

4.3 MRIV VS. WALD ESTIMATOR

We compare $\hat{\tau}_{\text{MRIV}}(x)$ to the Wald estimator $\hat{\tau}_W(x)$. First, we derive an asymptotic upper bound.

Theorem 3 (Wald oracle upper bound). *Assume that $\mu_1^Y(x), \mu_0^Y(x)$ are α -smooth, $\mu_1^A(x), \mu_0^A(x)$ are β -smooth, and are estimated with their respective minimax rate. Let $\hat{\delta}_A(x) = \hat{\mu}_1^A(x) - \hat{\mu}_0^A(x)$ satisfy Assumption 4. Then, the oracle risk of the Wald estimator $\hat{\tau}_W(x)$ is bounded by*

$$\mathbb{E} \left[(\hat{\tau}_W(x) - \tau(x))^2 \right] \lesssim n^{-\frac{2\alpha}{2\alpha+p}} + n^{-\frac{2\beta}{2\beta+p}}. \quad (5)$$

Proof. See Appendix B. \square

We now consider the MRIV estimator $\hat{\tau}_{\text{MRIV}}(x)$ with $\hat{\tau}_{\text{init}} = \hat{\tau}_{\text{W}}(x)$, i.e., initialized with the Wald estimator (under sample splitting). Plugging the Wald rate from Eq. (5) into the Eq. (2) yields

$$\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim n^{-\frac{2\eta}{2\eta+p}} + n^{-2\left(\frac{\alpha}{2\alpha+p} + \frac{\delta}{2\delta+p}\right)} + n^{-2\left(\frac{\beta}{2\beta+p} + \frac{\delta}{2\delta+p}\right)} + n^{-2\left(\frac{\alpha}{2\alpha+p} + \frac{\gamma}{2\gamma+p}\right)} + n^{-2\left(\frac{\beta}{2\beta+p} + \frac{\gamma}{2\gamma+p}\right)}.$$

For $\alpha = \beta = \gamma = \delta$, the rates of $\hat{\tau}_{\text{MRIV}}(x)$ and $\hat{\tau}_{\text{W}}(x)$ reduce to $\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim n^{-\frac{2\eta}{2\eta+p}} + n^{-\frac{4\alpha}{2\alpha+p}}$ and $\mathbb{E} \left[(\hat{\tau}_{\text{W}}(x) - \tau(x))^2 \right] \lesssim n^{-\frac{2\alpha}{2\alpha+p}}$. Hence, $\hat{\tau}_{\text{MRIV}}(x)$ outperforms $\hat{\tau}_{\text{W}}(x)$ asymptotically for $\eta > \alpha$, i.e., when the CATE $\tau(x)$ is smoother than its components, which is usually the case in practice (Künzel et al., 2019). For $\eta = \alpha$, the rates of both estimators coincide. Hence, we should expect MRIV to improve on the Wald estimator in real-world settings with large sample size.

4.4 MRIV-NET

Based on our MRIV framework, we develop a tailored deep neural network called MRIV-Net for CATE estimation using IVs. Our MRIV-Net produces both an initial CATE estimator $\hat{\tau}_{\text{init}}(x)$ and nuisance estimators $\hat{\mu}_0^Y(x)$, $\hat{\mu}_0^A(x)$, $\hat{\delta}_A(x)$, and $\hat{\pi}(x)$.

For MRIV-Net, we choose deep neural networks for the nuisance components due to their predictive power and their ability to learn complex shared representations for several nuisance components. Sharing representations between nuisance components has been exploited previously for CATE estimation, yet only under unconfoundedness (Shalit et al., 2017; Curth & van der Schaar, 2021). Building shared representations is more efficient in finite sample regimes than estimating all nuisance components separately as they usually share some common structure.

In MRIV-Net, not all nuisance components should share a representation.

Recall that, in Theorem 2, we assumed that (1) $\hat{\tau}_{\text{init}}(x)$, $\hat{\mu}_0^Y(x)$, and $\hat{\mu}_0^A(x)$; and (2) $\hat{\delta}_A(x)$ and $\hat{\pi}(x)$ are trained on two independent samples in order to derive the upper bound on the oracle risk. Hence, we propose to build two separate representations Φ_1 and Φ_2 , so that (i) Φ_1 is used to learn the parameters (1), and (ii) Φ_2 is used to learn the parameters (2). This ensures that the nuisance estimators (1) share minimal information with nuisance estimators (2) even though they are estimated on the same data (cf. (Curth & van der Schaar, 2021)).

The architecture of MRIV-Net is shown in Fig. 2. MRIV-Net takes the observed covariates X as input to build the two representations Φ_1 and Φ_2 . The first representation Φ_1 is used to output estimates $\hat{\mu}_1^Y(x)$, $\hat{\mu}_0^Y(x)$, $\hat{\mu}_1^A(x)$, and $\hat{\mu}_0^A(x)$ of the CATE components. The second representation Φ_2 is used to output estimates $\tilde{\mu}_1^A(x)$, $\tilde{\mu}_0^A(x)$, and $\hat{\pi}(x)$. MRIV-Net is trained by minimizing an overall loss

$$\mathcal{L}(\theta) = \sum_{i=1}^n \left[(\hat{\mu}_{z_i}^Y(x_i) - y_i)^2 + \text{BCE}(\hat{\mu}_{z_i}^A(x_i), a_i) + \text{BCE}(\tilde{\mu}_{z_i}^A(x_i), a_i) + \text{BCE}(\hat{\pi}(x_i), z_i) \right], \quad (6)$$

where θ denotes the neural network parameters and BCE is the binary cross entropy loss. After training MRIV-Net, we obtain the $\hat{\tau}_{\text{init}}(x) = \frac{\hat{\mu}_1^Y(x) - \hat{\mu}_0^Y(x)}{\hat{\mu}_1^A(x) - \hat{\mu}_0^A(x)}$ and obtain the nuisance estimators $\hat{\mu}_0^Y(x)$, $\hat{\mu}_0^A(x)$, $\hat{\delta}_A(x) = \tilde{\mu}_1^A(x) - \tilde{\mu}_0^A(x)$ and $\hat{\pi}(x)$. Then, we perform the pseudo regression (Stage 2) of MRIV to obtain $\hat{\tau}_{\text{MRIV}}(x)$.

Implementation: Details on the implementation, the network architecture and hyperparameter tuning are in Appendix I. We perform both the training of MRIV-Net and the pseudo-outcome regression on the full training data. Needless to say, MRIV-Net can be easily adopted for sample splitting or cross-fitting procedures as in Chernozhukov et al. (2018), namely, by learning separate networks for each representation Φ_1 and Φ_2 . In our experiments, we do not use sample splitting or cross-fitting, as this can affect the performance in finite sample regimes. Of note, our choice is consistent with previous work (Curth & van der Schaar, 2021). In Appendix K we report results using cross-fitting.

5 COMPUTATIONAL EXPERIMENTS

5.1 SIMULATED DATA

In causal inference literature, it is common practice to use simulated data for performance evaluations (Bica et al., 2020a; Curth & van der Schaar, 2021; Hartford et al., 2017). Simulated data offers the




Figure 2: Architecture of MRIV-Net.

crucial benefit that it provides ground-truth information on the counterfactual outcomes and thus allows for direct benchmarking against the oracle CATE.

Data generation: We generate simulated data by sampling the oracle CATE $\tau(x)$ and the nuisance components $\mu_i^Y(x)$, $\mu_i^A(x)$, and $\pi(x)$ from Gaussian process priors. Using Gaussian processes has the following advantages: (1) It allows for a fair method comparison, as there is no need to explicitly specify the nuisance components, which could lead to unwanted inductive biases favoring a specific method; (2) the sampled nuisance components are non-linear and thus resemble real-world scenarios where machine learning methods would be applied; and, (3) by sampling from the prior induced by the Matérn kernel (Rasmussen & Williams, 2008), we can control the smoothness of the nuisance components, which allows us to confirm our theoretical results from Sec. 4.2. For a detailed description of our data generating process, we refer to Appendix E.

Baselines: We compare our MRIV-Net with state-of-the-art IV baselines. Details regarding baselines and nuisance parameter estimation are in Appendix G. Note that many of the baselines do not directly aim at CATE estimation but rather at counterfactual outcome prediction. We nevertheless use these methods as baselines and, for this, obtain the CATE by taking the difference between the predictions of the factual and counterfactual outcomes.

Performance evaluation: For all experiments, we use a 80/20 split as training/test set. We calculate the root mean squared errors (RMSE) between the CATE estimates and the oracle CATE on the test set. We report the mean RMSE and the standard deviation over five data sets generated from random seeds.

Results: Table 2 shows the results for all baselines. Here, the DR-learner does not improve the performance of TAR-Net, which is reasonable as both the DR-learner and TARNet assume unconfoundedness and are thus biased in our setting. Our MRIV-Net outperforms all baselines. Our MRIV-Net also achieves a smaller standard deviation. For additional results, we refer to Appendix J.

We further compare the performance of two different meta-learner frameworks – DRIV (Syrgkanis et al., 2019) and our MRIV – across different base methods. The results are in Table 3. The nuisance parameters are estimated using feed forward neural networks (DRIV) or TAR-

Table 2: Performance comparison: our MRIV-Net vs. existing baselines.

Method	$n = 3000$	$n = 5000$	$n = 8000$
(1) STANDARD ITE			
TARNet (Shalit et al., 2017)	0.76 ± 0.14	0.70 ± 0.12	0.69 ± 0.17
TARNet + DR (Shalit et al., 2017; Kennedy, 2022)	0.78 ± 0.10	0.66 ± 0.09	0.70 ± 0.10
(2) GENERAL IV			
2SLS (Wooldridge, 2013)	1.22 ± 0.23	0.79 ± 0.37	1.12 ± 0.29
KIV (Singh et al., 2019)	1.54 ± 0.53	1.18 ± 1.14	3.80 ± 4.71
DFIV (Xu et al., 2021a)	0.43 ± 0.11	0.40 ± 0.21	0.46 ± 0.54
DeepIV (Hartford et al., 2017)	0.96 ± 0.30	0.28 ± 0.09	0.23 ± 0.04
DeepGMM (Bennett et al., 2019)	0.95 ± 0.38	0.37 ± 0.09	0.42 ± 0.14
DMLIV (Syrgkanis et al., 2019)	1.92 ± 0.71	0.92 ± 0.41	1.14 ± 0.24
DMLIV + DRIV (Syrgkanis et al., 2019)	0.41 ± 0.12	0.22 ± 0.04	0.21 ± 0.06
(3) WALD ESTIMATOR (WALD, 1940)			
Linear	1.06 ± 0.63	0.62 ± 0.22	0.81 ± 0.34
BART	0.95 ± 0.30	0.63 ± 0.33	0.88 ± 0.28
MRIV-Net (ours)	0.26 ± 0.11	0.15 ± 0.03	0.13 ± 0.03

Reported: RMSE for base methods (mean ± standard deviation). Lower = better (best in bold)

Table 3: Base model with different meta-learners (i.e., none, DRIV, and our MRIV).

Base methods	Meta-learners			$n = 3000$			$n = 5000$			$n = 8000$		
	None	DRIV	MRIV (ours)	None	DRIV	MRIV (ours)	None	DRIV	MRIV (ours)			
(1) STANDARD ITE												
TARNet (Shalit et al., 2017)	0.76 ± 0.14	0.31 ± 0.05	0.34 ± 0.13	0.70 ± 0.12	0.17 ± 0.06	0.17 ± 0.05	0.69 ± 0.17	0.21 ± 0.04	0.16 ± 0.04			
(2) GENERAL IV												
2SLS (Wooldridge, 2013)	1.22 ± 0.23	0.40 ± 0.11	0.31 ± 0.08	0.79 ± 0.37	0.17 ± 0.09	0.19 ± 0.05	1.12 ± 0.29	0.21 ± 0.05	0.16 ± 0.02			
KIV (Singh et al., 2019)	1.54 ± 0.53	0.40 ± 0.10	0.39 ± 0.11	1.18 ± 1.14	0.20 ± 0.08	0.17 ± 0.06	3.80 ± 4.71	0.31 ± 0.18	0.28 ± 0.19			
DFIV (Xu et al., 2021a)	0.43 ± 0.11	0.26 ± 0.05	0.27 ± 0.07	0.40 ± 0.21	0.18 ± 0.09	0.16 ± 0.04	0.46 ± 0.54	0.21 ± 0.06	0.18 ± 0.05			
DeepIV (Hartford et al., 2017)	0.96 ± 0.30	0.27 ± 0.03	0.26 ± 0.05	0.28 ± 0.09	0.18 ± 0.08	0.18 ± 0.05	0.23 ± 0.04	0.21 ± 0.03	0.16 ± 0.03			
DeepGMM (Bennett et al., 2019)	0.95 ± 0.38	0.40 ± 0.15	0.36 ± 0.13	0.37 ± 0.09	0.24 ± 0.12	0.16 ± 0.05	0.42 ± 0.14	0.21 ± 0.03	0.17 ± 0.03			
DMLIV (Syrgkanis et al., 2019)	1.92 ± 0.71	0.41 ± 0.12	0.37 ± 0.11	0.92 ± 0.41	0.22 ± 0.05	0.16 ± 0.05	1.14 ± 0.24	0.21 ± 0.06	0.18 ± 0.05			
(3) WALD ESTIMATOR (WALD, 1940)												
Linear	1.06 ± 0.63	0.42 ± 0.15	0.38 ± 0.14	0.62 ± 0.22	0.19 ± 0.09	0.25 ± 0.09	0.81 ± 0.34	0.19 ± 0.09	0.18 ± 0.04			
BART	0.95 ± 0.30	0.48 ± 0.14	0.46 ± 0.12	0.63 ± 0.33	0.26 ± 0.13	0.20 ± 0.07	0.88 ± 0.28	0.31 ± 0.08	0.29 ± 0.04			
MRIV-Net\w network only (ours)	0.39 ± 0.13	0.35 ± 0.12	0.26 ± 0.11	0.31 ± 0.04	0.19 ± 0.13	0.15 ± 0.03	0.26 ± 0.06	0.18 ± 0.08	0.13 ± 0.03			

Reported: RMSE (mean ± standard deviation). Lower = better (best improvement over none meta-learner in bold)

Nets with either binary or continuous outputs (MRIV). Our MRIV improves over the variant without any meta-learner framework across all base methods (both in terms of RMSE and standard deviation). Furthermore, MRIV is clearly superior over DRIV. This demonstrates the effectiveness of our MRIV across different base methods (note: MRIV with an arbitrary base model is typically superior to DRIV with our custom network from above). MRIV-Net is overall best. We also performed additional experiments where we used semi-synthetic data and cross-fitting approaches for both meta-learners (see Appendix J and K).

Table 4: Ablation study.

Method	$n = 3000$	$n = 5000$	$n = 8000$
MRIV-Net\w network only	0.39 ± 0.13	0.31 ± 0.04	0.26 ± 0.06
MRIV-Net\w single repr.	0.28 ± 0.12	0.21 ± 0.04	0.32 ± 0.10
MRIV-Net (ours)	0.26 ± 0.11	0.15 ± 0.03	0.13 ± 0.03

Reported: RMSE (mean ± standard deviation). Lower = better (best in bold)

Ablation study: Table 4 compares different variants of our MRIV-Net. These are: (1) MRIV but network only; (2) MRIV-Net with a single representation for all nuisance estimators; and (3) our MRIV-Net from above. We observe that MRIV-Net is best. This justifies our proposed network architecture for MRIV-Net. Hence, combing the result from above, our performance gain must be attributed to both our framework and the architecture of our deep neural network.

Robustness checks for unobserved confounding and smoothness: Here, we demonstrate the importance of handling unobserved confounding (as we do in our MRIV framework). For this, Fig. 3 plots the results for our MRIV-Net vs. standard




Figure 3: Results over different levels of confounding α_U . Shaded area shows standard deviation.

CATE without customization for confounding (i.e., TARNet with and without the DR-learner) over over different levels of unobserved confounding. The RMSE of both TARNet variants increase almost linearly with increasing confounding. In contrast, the RMSE of our MRIV-Net only marginally. Even for low confounding regimes, our MRIV-Net performs competitively.

Fig. 4 varies the smoothness level. This is given by α of $\mu_i^Y(\cdot)$ (controlled by the Matérn kernel prior). Here, the performance decreases for the baselines, i.e., DeepIV and our network without MRIV framework. In contrast, the performance of our MRIV-Net remains robust and outperforms the baselines. This confirms our theoretical results from above. It thus indicates that our MRIV framework works best when the oracle CATE $\tau(x)$ is smoother than the nuisance parameters $\mu_i^Y(x)$.




Figure 4: Results over different levels of smoothness α of $\mu_i^Y(\cdot)$, sample size $n = 8000$. Larger $\alpha =$ smoother. Shaded areas show standard deviation.

5.2 CASE STUDY WITH REAL-WORLD DATA

Setting: We demonstrate effectiveness of our framework using a case study with real-world, medical data. Here, we use medical data from the so-called *Oregon health insurance experiment* (OHIE) (Finkelstein et al., 2012). It provides data for an RCT with non-compliance: In 2008, $\sim 30,000$ low-income, uninsured adults in Oregon were offered participation in a health insurance program by a lottery. Individuals whose names were drawn could decide to sign up for health insurance. After a period of 12 months, in-person interviews took place to evaluate the health condition of the respective participant.




Figure 5: Results on real-world medical data. The results for MRIV-Net, our neural network architecture without the MRIV framework, and TARNet are in Fig. 5.

Results: Our MRIV-Net estimates larger causal effects for an older age. In contrast, TARNet does not estimate positive CATEs even for an older age. Even though we cannot evaluate the estimation quality on real-world data, our estimates seem reasonable in light of the medical literature: the benefit of health insurance should increase with older age. This showcases that TARNet may suffer from bias induced by unobserved confounders. We also report the results for DRIV with DMLIV as base method, and observe that in contrast to MRIV-Net, the corresponding CATE does not vary much between ages. Interestingly, both our MRIV-Net estimate a somewhat smaller CATE for middle ages (around 30–50 yrs). In sum, the findings from our case study are of direct relevance for decision-makers in public health (Imbens & Angrist, 1994), and highlight the practical value of our framework. We performed further experiments on real-world data which are reported in Appendix L.

In our analysis, the lottery assignment is the instrument Z , the decision to sign up for health insurance is treatment A , and an overall health score is the outcome Y . We also include five covariates X , including age and gender. For details, we refer to Appendix F. We first estimate the CATE function and then report the treatment effect heterogeneity w.r.t. age and gender, while

Reproducibility: The codes for reproducing the experimental results can be found at <https://github.com/DennisFrauen/MRIV-Net>.

REFERENCES

- Ahmed M. Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task Gaussian processes. In *NeurIPS*, 2017.
- Joshua D. Angrist. Lifetime earnings and the vietnam era draft lotter: Evidence from social security administrative records. *The American Economic Review*, 80(3):313–336, 1990.
- Joshua D. Angrist and Alan B. Krueger. Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4):979–1014, 1991.
- Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- Falco J. Bargagli-Stoffi, Kristof de Witte, and Giorgio Gnecco. Heterogeneous causal effects with imperfect compliance: A bayesian machine learning approach. *Annals of Applied Statistics*, 2021.
- Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. In *NeurIPS*, 2019.
- Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. In *ICLR*, 2020a.
- Ioana Bica, Ahmed M. Alaa, and Mihaela van der Schaar. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In *ICML*, 2020b.
- Howard S. Bloom, Larry L. Orr, Stephen H. Bell, George Cave, Fred Doolittle, Winston Lin, and Johannes M. Bos. The benefits and costs of JTPA title ii-a programs: Key findings from the national job training partnership act study. *Journal of Human Resources*, 32(32):549–586, 1997.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- Jean Chesson. A non-central multivariate hypergeometric distribution arising from biased sampling with application to selective predation. *Journal of Applied Probability*, 13(4):795–797, 1976.
- Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- Yifan Cui and Eric Tchetgen Tchetgen. A semiparametric instrumental variable approach to optimal treatment regimes under endogeneity. *Journal of the American Statistical Association*, 116(553):126–137, 2021.
- Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semiparametric proximal causal inference. *arXiv preprint*, 2020.
- Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *AISTATS*, 2021.
- Alicia Curth, Ahmed M. Alaa, and Mihaela van der Schaar. Estimating structural target functions using machine learning and influence functions. *arXiv preprint*, arXiv:2008.06461, 2020.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *ICLR*, 2018.
- Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, and Katherine Baicker. The oregon health insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics*, 127(3):1057–1106, 2012.

- Dennis Frauen, Tobias Hatt, Valentyn Melnychuk, and Stefan Feuerriegel. Estimating average causal effects from patient trajectories. In *AAAI*, 2023.
- Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep IV: A flexible approach for counterfactual prediction. In *ICML*, 2017.
- Guido W. Imbens and Joshua D. Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.
- Andrew Jesson, Sören Mindermann, Yarin Gal, and Uri Shalit. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In *ICML*, 2021.
- Nathan Kallus, Xiaojie Mao, and Angela Zhou. Interval estimation of individual-level causal effects under unobserved confounding. In *AISTATS*, 2019.
- Edward H. Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint*, 2022.
- Edward H. Kennedy, Scott A. Lorch, and Dylan S. Small. Robust causal inference with continuous instruments using the local instrumental variable curve. *Journal of the Royal Statistical Society: Series B*, 81(1):121–143, 2019.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences (PNAS)*, 116(10):4156–4165, 2019.
- Chunxiao Li, Cynthia Rudin, and Tyler H. McCormick. Rethinking nonlinear instrumental variable models through prediction validity. *Journal of Machine Learning Research*, 23:1–55, 2022.
- Bryan Lim, Ahmed M. Alaa, and Mihaela van der Schaar. Forecasting treatment responses over time using recurrent marginal structural networks. In *NeurIPS*, 2018.
- Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Causal transformer for estimating counterfactual outcomes. In *ICML*, 2022a.
- Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Normalizing flows for interventional density estimation. *arXiv preprint*, arXiv:2209.06203, 2022b.
- Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.
- Elizabeth L. Ogburn, Andrea Rotnitzky, and James M. Robins. Doubly robust estimation of the local average treatment effect curve. *Journal of the Royal Statistical Society: Series B*, 77(2):373–396, 2015.
- Ryo Okui, Dylan S. Small, Zhiqiang Tan, and James M. Robins. Doubly robust instrumental variable regression. *Statistica Sinica*, 22(1):173–205, 2012.
- Judea Pearl. *Causality*. Cambridge University Press, New York City, 2009. ISBN 9780521895606.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, Cambridge, Mass., 3. print edition, 2008. ISBN 9780262182539.
- James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974.
- Vira Semenova and Victor Chernozhukov. Debaised machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal*, 24(2):264–289, 2021.

- Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *ICML*, 2017.
- Rahul Singh and Liyang Sun. Double robustness for complier parameters and a semiparametric test for complier characteristics. *arXiv preprint*, 2019.
- Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. In *NeurIPS*, 2019.
- Charles J. Stone. Optimal rates of convergence for nonparametric estimators. *Annals of Statistics*, 8(6), 1980.
- Vasilis Syrgkanis, Victor Lei, Miruna Oprescu, Maggie Hei, Keith Battocchi, and Greg Lewis. Machine learning estimation of heterogeneous treatment effects with instruments. In *NeurIPS*, 2019.
- Hal R. Varian. Causal inference in economics and marketing. *Proceedings of the National Academy of Sciences (PNAS)*, 113(27):7310–7315, 2016.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Abraham Wald. The fitting of straight lines if both variables are subject to error. *Annals of Mathematical Statistics*, 11(3):284–300, 1940.
- Guihua Wang, Jun Li, and Wallace J. Hopp. An instrumental variable forest approach for detecting heterogeneous treatment effects in observational studies. *Management Science*, 2021.
- Linbo Wang and Eric J. Tchetgen Tchetgen. Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. *Journal of the Royal Statistical Society: Series B*, 80(3):531–550, 2018.
- Yixin Wang and David M. Blei. The blessings of multiple causes. *Journal of the American Statistical Association*, 114(528):1574–1596, 2019.
- Jeffrey M. Wooldridge. *Introductory Econometrics: A modern approach*. Routledge, 2013. ISBN 9781136586101.
- Phillip G. Wright. *The tariff on animal and vegetable oils*. Macmillan, New York, 1928.
- Liyuan Xu, Yutian Chen, Siddarth Srinivasan, Nando de Freitas, Arnaud Doucet, and Arthur Gretton. Learning deep features in instrumental variable regression. In *ICLR*, 2021a.
- Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. In *NeurIPS*, 2021b.
- Yun Yang and Surya T. Tokdar. Minimax-optimal nonparametric regression in high dimensions. *The Annals of Statistics*, 43(2):652–674, 2015.
- Azam M. Yazdani and Eric Boerwinkle. Causal inference in the age of decision medicine. *Journal of Data Mining in Genomics & Proteomics*, 6(1), 2015.
- Jinsung Yoon, James Jordon, and Mihaela van der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *ICLR*, 2018.
- Yao Zhang, Alexis Bellot, and Mihaela van der Schaar. Learning overlapping representations for the estimation of individualized treatment effects. In *AISTATS*, 2020.

A EXTENDED RELATED WORK

CATE methods without unconfoundedness: Various machine learning methods for estimating CATEs *without* unobserved confounding have been proposed in recent literature (Alaa & van der Schaar, 2017; Curth & van der Schaar, 2021; Künzel et al., 2019; Lim et al., 2018; Shalit et al., 2017; Wager & Athey, 2018; Yoon et al., 2018; Zhang et al., 2020). To remove plug-in bias, the DR-learner performs a second stage regression on the uncentered influence function of the average treatment effect (Kennedy, 2022). However, under unobserved confounding, all of these methods are biased (see Appendix G). As a result, this hampers their performance in our setting.

Non-IV methods for unobserved confounding: There is a rich literature for causal effect estimation under unobserved confounding, which does not assume the existence of instrumental variables. Methods include deconfounding methods (Wang & Blei, 2019; Bica et al., 2020b), proxy learning methods (Cui et al., 2020; Xu et al., 2021b), and causal sensitivity analysis (Kallus et al., 2019; Jesson et al., 2021)..

Classical IV methods: IV methods address the problem of unobserved confounding by exploiting the variance in treatment and outcome induced by the instruments. Traditionally, two-stage least squares (2SLS) has been used for estimating causal effects (Wright, 1928; Angrist & Krueger, 1991). 2SLS was originally developed in economics, and follows a two-stage procedure: it performs a first stage regression of treatment A on the instrument Z , and then uses the fitted values for a second stage regression to predict the outcome Y . Several nonparametric methods have been developed in econometric to generalize 2SLS in order to account for non-linearities within the data (Newey & Powell, 2003; Wang et al., 2021), yet these are limited to low-dimensional settings.

B PROOFS

We start by deriving an auxiliary Lemma. That is, we derive an explicit expression for the Stage 2 oracle pseudo-outcome regression $\mathbb{E}[\hat{Y}_{\text{MR}} | X = x]$ of MRIV.

Lemma 1.

$$\begin{aligned} & \mathbb{E}[\hat{Y}_{\text{MR}} | X = x] \\ &= \frac{\pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} (\mu_1^Y(x) - \mu_1^A(x) \hat{\tau}_{\text{init}}(x)) + \frac{(1 - \pi(x))}{\hat{\delta}_A(x)(1 - \hat{\pi}(x))} (\mu_0^A(x) \hat{\tau}_{\text{init}}(x) - \mu_0^Y(x)) \\ & \quad + \frac{\hat{\mu}_0^A(x) \hat{\tau}_{\text{init}}(x) - \hat{\mu}_0^Y(x)}{\hat{\delta}_A(x)} \left(\frac{\pi(x)}{\hat{\pi}(x)} - \frac{1 - \pi(x)}{1 - \hat{\pi}(x)} \right) + \hat{\tau}_{\text{init}}(x) \end{aligned} \quad (7)$$

Proof.

$$\begin{aligned} & \mathbb{E}[\hat{Y}_{\text{MR}} | X = x] \tag{8} \\ &= \pi(x) \mathbb{E} \left[\frac{Y - A \hat{\tau}_{\text{init}}(X) - \hat{\mu}_0^Y(X) + \hat{\mu}_0^A(X) \hat{\tau}_{\text{init}}(X)}{\hat{\delta}_A(X) \hat{\pi}(X)} \middle| X = x, Z = 1 \right] \\ & \quad + (1 - \pi(x)) \mathbb{E} \left[\frac{Y - A \hat{\tau}_{\text{init}}(X) - \hat{\mu}_0^Y(X) + \hat{\mu}_0^A(X) \hat{\tau}_{\text{init}}(X)}{\hat{\delta}_A(X) (1 - \hat{\pi}(X))} \middle| X = x, Z = 0 \right] + \hat{\tau}_{\text{init}}(x) \end{aligned} \quad (9)$$

$$\begin{aligned} &= \frac{\pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} (\mu_1^Y(x) - \mu_1^A(x) \hat{\tau}_{\text{init}}(x) - \hat{\mu}_0^Y(x) + \hat{\mu}_0^A(x) \hat{\tau}_{\text{init}}(x)) \\ & \quad + \frac{1 - \pi(x)}{\hat{\delta}_A(x) (1 - \hat{\pi}(x))} (\mu_0^Y(x) - \mu_0^A(x) \hat{\tau}_{\text{init}}(x) - \hat{\mu}_0^Y(x) + \hat{\mu}_0^A(x) \hat{\tau}_{\text{init}}(x)) + \hat{\tau}_{\text{init}}(x) \end{aligned} \quad (10)$$

Rearranging the terms yields the desired result. \square

B.1 PROOF OF THEOREM 1 (MULTIPLE ROBUSTNESS PROPERTY)

We use Lemma 1 to show that under each of the three conditions it follows that $\mathbb{E}[\hat{Y}_{\text{MR}} | X = x] = \tau(x)$.

1.

$$\begin{aligned} & \mathbb{E}[\hat{Y}_{\text{MR}} | X = x] \tag{11} \\ &= \frac{\pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} (\mu_1^Y(x) - \mu_1^A(x) \tau(x) + \mu_0^A(x) \tau(x) - \mu_0^Y(x)) \\ & \quad + \frac{(1 - \pi(x))}{\hat{\delta}_A(x) (1 - \hat{\pi}(x))} (\mu_0^A(x) \tau(x) - \mu_0^Y(x) - \mu_0^A(x) \tau(x) + \mu_0^Y(x)) + \tau(x) \end{aligned} \quad (12)$$

$$= \frac{\pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} (\delta_Y(x) - \delta_Y(x)) + \tau(x) = \tau(x). \tag{13}$$

2.

$$\mathbb{E}[\hat{Y}_{\text{MR}} | X = x] = \frac{(\mu_1^Y(x) - \mu_1^A(x) \hat{\tau}_{\text{init}}(x))}{\delta_A(x)} + \frac{(\mu_0^A(x) \hat{\tau}_{\text{init}}(x) - \mu_0^Y(x))}{\delta_A(x)} + \hat{\tau}_{\text{init}}(x) \tag{14}$$

$$= \frac{\delta_Y(x) - \hat{\tau}_{\text{init}}(x) \delta_A(x)}{\delta_A(x)} + \hat{\tau}_{\text{init}}(x) = \tau(x). \tag{15}$$

3.

$$\mathbb{E}[\hat{Y}_{\text{MR}} | X = x] = \frac{(\mu_1^Y(x) - \mu_1^A(x)) \tau(x)}{\hat{\delta}_A(x)} + \frac{(\mu_0^A(x) \tau(x) - \mu_0^Y(x))}{\hat{\delta}_A(x)} + \tau(x) \quad (16)$$

$$= \frac{\delta_Y(x)}{\hat{\delta}_A(x)} - \tau(x) \frac{\delta_A(x)}{\hat{\delta}_A(x)} + \tau(x) = \tau(x) \quad (17)$$

B.2 PROOF OF THEOREM 2 (CONVERGENCE RATE OF MRIV)

To prove Theorem 2, we need an additional assumption on the second stage regression estimator $\hat{\mathbb{E}}_n$. We refer to Kennedy (2022) (Proposition 1) for a detailed discussion on this assumption.

Assumption 5 (From Proposition 1 of Kennedy (2022)). Let Y_{MR} be the corresponding oracle to the estimated pseudo-outcome \hat{Y}_{MR} . We assume that the pseudo-regression estimator $\hat{\mathbb{E}}_n$ satisfies

$$\frac{\hat{\mathbb{E}}_n[\hat{Y}_{\text{MR}} | X = x] - \hat{\mathbb{E}}_n[Y_{\text{MR}} | X = x] - \hat{\mathbb{E}}_n[\hat{Y}_{\text{MR}} - Y_{\text{MR}} | X = x]}{\sqrt{\mathbb{E} \left[\left(\hat{\mathbb{E}}_n[Y_{\text{MR}} | X = x] - \mathbb{E}[Y_{\text{MR}} | X = x] \right)^2 \right]}} \xrightarrow{p} 0 \quad (18)$$

and

$$\mathbb{E} \left[\hat{\mathbb{E}}_n[\hat{r}(X) | X = x]^2 \right] \asymp \mathbb{E} [\hat{r}(x)^2], \quad (19)$$

where $r(x) = \mathbb{E}[\hat{Y}_{\text{MR}} | X = x] - \tau(x)$

To prove Theorem 2, we derive a more general bound the depends on the pointwise mean squared errors of the nuisance estimators. Theorem 2 follows immediately by applying Assumption 3.

Lemma 2. *Consider the setting described in Theorem 2. Then,*

$$\begin{aligned} & \mathbb{E} \left[(\hat{\tau}_{\text{init}}(x) - \tau(x))^2 \right] \quad (20) \\ & \lesssim \mathcal{R}(x) + \mathbb{E} \left[(\hat{\tau}_{\text{init}}(x) - \tau(x))^2 \right] \left(\mathbb{E} \left[\left(\hat{\delta}_A(x) - \delta_A(x) \right)^2 \right] + \mathbb{E} \left[(\hat{\pi}(x) - \pi(x))^2 \right] \right) \\ & \quad + \mathbb{E} \left[(\hat{\pi}(x) - \pi(x))^2 \right] \left(\mathbb{E} \left[\left(\hat{\mu}_0^Y(x) - \mu_0^Y(x) \right)^2 \right] + \mathbb{E} \left[\left(\hat{\mu}_0^A(x) - \mu_0^A(x) \right)^2 \right] \right). \quad (21) \end{aligned}$$

Proof. Let Y_{MR} be the corresponding oracle to \hat{Y}_{MR} and define $\tilde{\tau}_{\text{MRIV}}(x) = \hat{\mathbb{E}}_n[Y_{\text{MR}} | X = x]$. Using Assumption 5, we can apply Proposition 1 of Kennedy (2022) and obtain

$$\mathbb{E} \left[(\hat{\tau}_{\text{init}}(x) - \tau(x))^2 \right] \lesssim \mathcal{R}(x) + \mathbb{E} [\hat{r}(x)^2], \quad (22)$$

where $\mathcal{R}(x) = \mathbb{E} \left[(\tilde{\tau}_{\text{MR}}(x) - \tau(x))^2 \right]$ is the oracle risk of the second stage regression. We can apply Lemma 1 to obtain

$$\begin{aligned} \hat{r}(x) &= \frac{\pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} (\mu_1^Y(x) - \mu_1^A(x) \hat{\tau}_{\text{init}}(x)) + \frac{(1 - \pi(x))}{\hat{\delta}_A(x) (1 - \hat{\pi}(x))} (\mu_0^A(x) \hat{\tau}_{\text{init}}(x) - \mu_0^Y(x)) \\ &\quad + \frac{\hat{\mu}_0^A(x) \hat{\tau}_{\text{init}}(x) - \hat{\mu}_0^Y(x)}{\hat{\delta}_A(x)} \left(\frac{\pi(x)}{\hat{\pi}(x)} - \frac{1 - \pi(x)}{1 - \hat{\pi}(x)} \right) + \hat{\tau}_{\text{init}}(x) - \tau(x) \end{aligned} \quad (23)$$

$$\begin{aligned} &= \left(\frac{\mu_1^Y(x) - \mu_0^Y(x)}{\hat{\delta}_A(x)} \right) \frac{\pi(x)}{\hat{\pi}(x)} + \frac{\mu_0^Y(x) - \hat{\mu}_0^Y(x)}{\hat{\delta}_A(x)} \left(\frac{\pi(x)}{\hat{\pi}(x)} - \frac{1 - \pi(x)}{1 - \hat{\pi}(x)} \right) + (\hat{\tau}_{\text{init}}(x) - \tau(x)) \\ &\quad + \left(\frac{(\mu_0^A(x) - \mu_1^A(x)) \hat{\tau}_{\text{init}}(x)}{\hat{\delta}_A(x)} \right) \frac{\pi(x)}{\hat{\pi}(x)} + \frac{(\hat{\mu}_0^D(x) - \mu_0^D(x)) \hat{\tau}_{\text{init}}(x)}{\hat{\delta}_A(x)} \left(\frac{\pi(x)}{\hat{\pi}(x)} - \frac{1 - \pi(x)}{1 - \hat{\pi}(x)} \right) \end{aligned} \quad (24)$$

$$\begin{aligned} &= \frac{\delta_Y(x) \pi(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} + \frac{(\mu_0^Y(x) - \hat{\mu}_0^Y(x)) (\pi(x) - \hat{\pi}(x))}{\hat{\delta}_A(x) \hat{\pi}(x) (1 - \hat{\pi}(x))} + (\hat{\tau}_{\text{init}}(x) - \tau(x)) \\ &\quad - \frac{\delta_A(x) \pi(x) \hat{\tau}_{\text{init}}(x)}{\hat{\delta}_A(x) \hat{\pi}(x)} + \frac{(\hat{\mu}_0^A(x) - \mu_0^A(x)) \hat{\tau}_{\text{init}}(x) (\pi(x) - \hat{\pi}(x))}{\hat{\delta}_A(x) \hat{\pi}(x) (1 - \hat{\pi}(x))} \end{aligned} \quad (25)$$

$$\begin{aligned} &= \frac{(\pi(x) - \hat{\pi}(x))}{\hat{\delta}_A(x) \hat{\pi}(x) (1 - \hat{\pi}(x))} [(\mu_0^Y(x) - \hat{\mu}_0^Y(x)) + (\hat{\mu}_0^A(x) - \mu_0^A(x)) \hat{\tau}_{\text{init}}(x)] \\ &\quad + (\hat{\tau}_{\text{init}}(x) - \tau(x)) + \frac{\pi(x) \delta_A(x)}{\hat{\pi}(x) \hat{\delta}_A(x)} (\tau(x) - \hat{\tau}_{\text{init}}(x)) \end{aligned} \quad (26)$$

$$\begin{aligned} &= \frac{(\pi(x) - \hat{\pi}(x))}{\hat{\delta}_A(x) \hat{\pi}(x) (1 - \hat{\pi}(x))} [(\mu_0^Y(x) - \hat{\mu}_0^Y(x)) + (\hat{\mu}_0^A(x) - \mu_0^A(x)) \hat{\tau}_{\text{init}}(x)] \\ &\quad + (\tau(x) - \hat{\tau}_{\text{init}}(x)) (\delta_A(x) - \hat{\delta}_A(x)) \pi(x) + (\tau(x) - \hat{\tau}_{\text{init}}(x)) (\pi(x) - \hat{\pi}(x)) \hat{\delta}_A(x). \end{aligned} \quad (27)$$

Applying the inequality $(a + b)^2 \leq 2(a^2 + b^2)$ together with Assumption 4 and the fact that $\pi(x) \leq 1$ yields

$$\begin{aligned} \hat{r}(x)^2 &\leq \frac{4}{\epsilon^4 \rho^2} (\pi(x) - \hat{\pi}(x))^2 \left[(\mu_0^Y(x) - \hat{\mu}_0^Y(x))^2 + (\hat{\mu}_0^A(x) - \mu_0^A(x))^2 K^2 \right] \\ &\quad + 4 (\tau(x) - \hat{\tau}_{\text{init}}(x))^2 (\delta_A(x) - \hat{\delta}_A(x))^2 + 4 (\tau(x) - \hat{\tau}_{\text{init}}(x))^2 (\pi(x) - \hat{\pi}(x))^2. \end{aligned} \quad (28)$$

By setting $\tilde{K} = \max\{K, 1\}$, we obtain

$$\begin{aligned} \hat{r}(x)^2 &\leq \frac{4\tilde{K}^2}{\epsilon^4 \rho^2} \left((\pi(x) - \hat{\pi}(x))^2 \left[(\mu_0^Y(x) - \hat{\mu}_0^Y(x))^2 + (\hat{\mu}_0^A(x) - \mu_0^A(x))^2 + (\hat{\tau}_{\text{init}}(x) - \tau(x))^2 \right] \right. \\ &\quad \left. + (\tau(x) - \hat{\tau}_{\text{init}}(x))^2 (\delta_A(x) - \hat{\delta}_A(x))^2 \right). \end{aligned} \quad (29)$$

Applying expectations on both sides yields the results, because $(\hat{\pi}(x), \hat{\delta}_A(x)) \perp\!\!\!\perp (\hat{\mu}_0^Y(x), \hat{\mu}_0^A(x), \hat{\tau}_{\text{init}}(x))$ due to sample splitting. \square

B.3 PROOF OF THEOREM 3 (CONVERGENCE RATE OF THE WALD ESTIMATOR)

Proof. We define $\tilde{C} = \max\{C, 1\}$ and obtain the upper bound

$$(\hat{\tau}_W(x) - \tau(x))^2 \tag{30}$$

$$= \left(\frac{(\hat{\mu}_1^Y(x) - \mu_1^Y(x))\delta_A(x) + (\mu_0^Y(x) - \hat{\mu}_0^Y(x))\delta_A(x) + (\delta_A(x) - \hat{\delta}_A(x))\delta_Y(x)}{\delta_A(x)\hat{\delta}_A(x)} \right)^2 \tag{31}$$

$$\leq \frac{4\tilde{C}^2}{\rho^2\hat{\rho}^2} \left[(\hat{\mu}_1^Y(x) - \mu_1^Y(x))^2 + (\hat{\mu}_0^Y(x) - \mu_0^Y(x))^2 + (\delta_A(x) - \hat{\delta}_A(x))^2 \right] \tag{32}$$

$$\leq \frac{8\tilde{C}^2}{\rho^2\hat{\rho}^2} \left[(\hat{\mu}_1^Y(x) - \mu_1^Y(x))^2 + (\hat{\mu}_0^Y(x) - \mu_0^Y(x))^2 + (\hat{\mu}_1^A(x) - \mu_1^A(x))^2 + (\hat{\mu}_0^A(x) - \mu_0^A(x))^2 \right], \tag{33}$$

where we used the inequality $(a + b)^2 \leq 2(a^2 + b^2)$ several times. Taking expectations and applying the smoothness assumptions yields the result. \square

C THEORETICAL ANALYSIS UNDER SPARSITY ASSUMPTIONS

In Sec. 4.2, we analyzed MRIV theoretically by imposing smoothness assumptions on the underlying data generating process. In particular, we derived a multiple robust convergence rate and showed that MRIV outperforms the Wald estimator if the oracle CATE is smoother than its components. In this section, we derive similar results by relying on a different set of assumptions. Instead of using smoothness, we make assumptions on the level of sparsity of the CATE components. This assumption is often imposed in high-dimensional settings ($n < p$) and is in line with previous literature on analyzing CATE estimators (Curth & van der Schaar, 2021; Kennedy, 2022).

In the following, we say a function $f(x)$ is k -sparse, if it is linear in $x \in \mathbb{R}^p$ and it only depends on $k < \min\{n, p\}$ predictors. (Yang & Tokdar, 2015) showed, that in this case the minimax rate of $f(x)$ is given by $\frac{k \log(p)}{n}$. The linearity assumption can be relaxed to an additive structural assumption, which we omit here for simplicity. In the following, we replace the smoothness conditions in Assumption 3 with sparsity conditions.

Assumption 6 (Sparsity). We assume that (1) the nuisance components $\mu_i^Y(\cdot)$ are α -sparse, $\mu_i^A(\cdot)$ and $\delta_A(\cdot)$ are β -sparse, and $\pi(\cdot)$ is δ -sparse; (2) all nuisance components are estimated with their respective minimax rate of $\frac{k \log(p)}{n}$, where $k \in \{\alpha, \beta, \delta\}$; and (3) the oracle CATE $\tau(\cdot)$ is γ -sparse and the initial CATE estimator $\hat{\tau}_{\text{init}}$ converges with rate $r_\tau(n)$.

We restate now our result from Theorem 3 for MRIV using the sparsity assumption.

Theorem 4 (MRIV upper bound under sparsity). *We consider the same setting as in Theorem 2 under the sparsity assumption 6. If the second-stage estimator $\hat{\mathbb{E}}_n$ yields the minimax rate $\frac{\gamma \log(p)}{n}$ and satisfies Assumption 5, the oracle risk is upper bounded by*

$$\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim \frac{\gamma \log(p)}{n} + r_\tau(n) \frac{(\beta + \delta) \log(p)}{n} + \frac{(\alpha + \beta) \delta \log^2(p)}{n^2}.$$

Proof. Follows immediately from Lemma 2 by applying Ass- 6. \square

Again, we obtain a multiple robust convergence rate for MRIV in the sense that MRIV achieves a fast rate even if the initial estimator or several nuisance estimators converge slowly. More precisely, for a fast convergence rate of $\hat{\tau}_{\text{MRIV}}(x)$, it is sufficient if either: (1) $r_\tau(n)$ decreases fast and δ is small; (2) $r_\tau(n)$ decreases fast and α and β are small; or (3) all α , β , and δ are small.

We derive now the corresponding rate for the Wald estimator.

Theorem 5 (Wald oracle upper bound). *Given estimators $\hat{\mu}_i^Y(x)$ and $\hat{\mu}_i^A(x)$. Let $\hat{\delta}_A(x) = \hat{\mu}_1^A(x) - \hat{\mu}_0^A(x)$ satisfy Assumption 4. Then, under Assumption 6 the oracle risk of the Wald estimator $\hat{\tau}_W(x)$ is bounded by*

$$\mathbb{E} [(\hat{\tau}_W(x) - \tau(x))^2] \lesssim \frac{(\alpha + \beta) \log(p)}{n} \quad (34)$$

Proof. Follows immediately from the proof of Theorem 3, i.e., from Eq.(30) by applying Ass- 6. \square

If $\alpha = \beta = \delta$, we obtain the rates

$$\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim \frac{\gamma \log(p)}{n} + \frac{\alpha^2 \log^2(p)}{n^2} \quad \text{and} \quad \mathbb{E} [(\hat{\tau}_W(x) - \tau(x))^2] \lesssim \frac{\alpha \log(p)}{n}, \quad (35)$$

which means that $\hat{\tau}_{\text{MRIV}}(x)$ outperforms $\hat{\tau}_W(x)$ for $\gamma < \alpha$, i.e., if the oracle CATE is more sparse than its components.

D MATHEMATICAL DETAILS REGARDING ASSUMPTION 3

In this section, we briefly state the formal definitions of the convergence rates in Assumption 3. We follow Stone (1980). Let θ be a parameter and $T(\theta)$ some target functional that we want to estimate.

Definition 1. A sequence of estimators $\hat{T}_n(\theta)$ of a functional $T(\theta)$ converges with rate $r_T(n)$ if

$$\lim_{c \rightarrow 0} \liminf_{n \in \mathbb{N}} \sup_{\theta \in \Theta} \mathbb{P}_\theta(|\hat{T}_n(\theta) - T(\theta)| > cr_\theta(n)) = 0. \quad (36)$$

Definition 2. A rate $r_T(n)$ is called an upper bound to the rate of convergence if for all estimators $\hat{T}_n(\theta)$ it holds for all $c > 0$ that

$$\liminf_{n \in \mathbb{N}} \sup_{\theta \in \Theta} \mathbb{P}_\theta(|\hat{T}_n(\theta) - T(\theta)| > cr_\theta(n)) > 0. \quad (37)$$

and

$$\lim_{c \rightarrow 0} \liminf_{n \in \mathbb{N}} \sup_{\theta \in \Theta} \mathbb{P}_\theta(|\hat{T}_n(\theta) - T(\theta)| > cr_\theta(n)) = 1. \quad (38)$$

$r_T(n)$ is called optimal if it is both achievable and an upper bound.

Stone (1980) showed that for a nonparametric regression problem with an η -smooth regression function, the optimal rate of convergence is $n^{-\frac{2\eta}{2\eta+p}}$, where p is the dimension of the covariate space.

E SIMULATED DATA

In the following, we describe how we simulate synthetic data for the experiments in Sec. 5.1 from the main paper. As mentioned therein, we simulate the CATE components from Gaussian processes using the prior induced by the Matern kernel (Rasmussen & Williams, 2008)

$$K_{\ell,\nu}(x_i, x_j) = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{\ell} \|x_i - x_j\|_2 \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\ell} \|x_i - x_j\|_2 \right), \quad (39)$$

where $\Gamma(\cdot)$ is the Gamma function and $K_\nu(\cdot)$ is the modified Bessel function of second kind. Here, ℓ is the length scale of the kernel and ν controls the smoothness of the sampled functions.

We set $\ell = 1$ and sample functions $\delta_Y \sim \mathcal{GP}(0, K_{\ell,\gamma})$, $\mu_0^Y \sim \mathcal{GP}(0, K_{\ell,\alpha})$, $f_1 \sim \mathcal{GP}(0, K_{\ell,\beta})$, $f_0 \sim \mathcal{GP}(0, K_{\ell,\beta})$ and $g \sim \mathcal{GP}(0, K_{\ell,\beta})$. Then, we define $\mu_1^Y = \delta_Y + \mu_0^Y$, $\mu_1^A = 0.3 \cdot \sigma \circ f_1 + 0.7$, $\mu_0^A = 0.3 \cdot \sigma \circ f_0$, $\delta_A = \mu_1^A - \mu_0^A$, $\mu_0^Y = c_0 \delta_A$, and $\pi = \sigma \circ g$. Finally, we set the oracle CATE to

$$\tau = \frac{\mu_1^Y - \mu_0^Y}{\mu_1^A - \mu_0^A} = \frac{\delta_Y}{\delta_A}. \quad (40)$$

Note that we can create a setup where the CATE τ is smoother than its components by using a small α/β ratio. An example is shown in Fig. 6.




Figure 6: Gaussian process simulation for $\alpha = 1.5$ and $\beta = 50$.

In the following, we describe how we generate data the (X, Z, A, Y) using the CATE components $\mu_i^Y(x)$, $\mu_i^A(x)$, and $\pi(x)$. We begin by sampling n observed confounder $X \sim \mathcal{N}(0, 1)$, unobserved confounders $U \sim \mathcal{N}(0, 0.2^2)$, and instruments $Z \sim \text{Bernoulli}(\pi(X))$. Then, we obtain treatments via

$$A = Z \mathbb{1}\{U + \epsilon_A > \alpha_1(X)\} + (1 - Z) \mathbb{1}\{U + \epsilon_A > \alpha_0(X)\} \quad (41)$$

with indicator function $\mathbb{1}$, noise $\epsilon_A \sim \mathcal{N}(0, 0.1^2)$, and $\alpha_i(X) = \Phi^{-1}(1 - \mu_i^A(X)) \sqrt{0.1^2 + 0.2^2}$, where Φ^{-1} denotes the quantile function of the standard normal distribution. Finally, we generate the outcomes via

$$Y = A \left(\frac{(\mu_1^A(X) - 1)\mu_0^Y(X) - \mu_0^A(X)\mu_1^Y(X) + \mu_1^Y(X)}{\delta_A(X)} \right) \quad (42)$$

$$+ (1 - A) \left(\frac{\mu_1^A(X)\mu_0^Y(X) - \mu_0^A(X)\mu_1^Y(X)}{\delta_A(X)} \right) + \alpha_U U + \epsilon_Y, \quad (43)$$

where $\epsilon_Y \sim \mathcal{N}(0, 0.3^2)$ is noise and $\alpha_U > 0$ is a parameter indicating the level of unobserved confounding. This choice of A and Y in Eq. (41) and Eq. (42), respectively, implies that $\tau(x)$ is indeed the CATE, i. e., it holds that $\tau(x) = \mathbb{E}[Y(1) - Y(0) | X = x]$.

Lemma 3. *Let (X, Z, A, Y) be sampled from the the previously described procedure. Then, it holds that*

$$\mu_i^A(x) = \mathbb{E}[A \mid Z = i, X = x] \quad \text{and} \quad \mu_i^Y(x) = \mathbb{E}[Y \mid Z = i, X = x]. \quad (44)$$

Proof. The first claim follows from

$$\mathbb{E}[A \mid Z = i, X = x] = \mathbb{P}(U + \epsilon_A > \alpha_i(x)) = 1 - \Phi(\Phi^{-1}(1 - \mu_i^A(x))) = \mu_i^A(x), \quad (45)$$

because $U + \epsilon_A \sim \mathcal{N}(0, \sqrt{0.1^2 + 0.2^2})$. The second claim follows from

$$\mathbb{E}[Y \mid Z = i, X = x] = \mu_i^A(x) \left(\frac{(\mu_1^A(x) - 1)\mu_0^Y(x) - \mu_0^A(x)\mu_1^Y(x) + \mu_1^Y(x)}{\delta_A(x)} \right) \quad (46)$$

$$+ (1 - \mu_i^A(x)) \left(\frac{\mu_1^A(x)\mu_0^Y(x) - \mu_0^A(x)\mu_1^Y(x)}{\delta_A(x)} \right) \quad (47)$$

$$= \frac{\mu_i^Y(x)\delta_A(x)}{\delta_A(x)} = \mu_i^Y(x). \quad (48)$$

□

F OREGON HEALTH INSURANCE EXPERIMENT

The so-called *Oregon health insurance experiment*³ (OHIE) (Finkelstein et al., 2012) was an important RCT with non-compliance. It was intentionally conducted as large-scale effort among public health to assess the effect of health insurance on several outcomes such as health or economic status. In 2008, a lottery draw offered low-income, uninsured adults in Oregon participation in a Medicaid program, providing health insurance. Individuals whose names were drawn could decide to sign up for the program.

In our analysis, the lottery assignment is the instrument Z , the decision to sign up for the Medicaid program is the treatment A , and an overall health score is the outcome Y . The outcome was obtained after a period of 12 months during in-person interviews. We use the following covariates X : age, gender, language, the number of emergency visits before the experiment, and the number of people the individual signed up with. The latter is used to control for peer effects, and it is important to include this variable in our analysis as it is the only variable influencing the propensity score (see below). We extract $\sim 10,000$ observations from the OHIE data and plot the histograms of all variables in Fig. 7. We can clearly observe the presence of non-compliance within the data, because the ratio of treated / untreated individuals is much lower than the corresponding ratio for the treatment assignment.

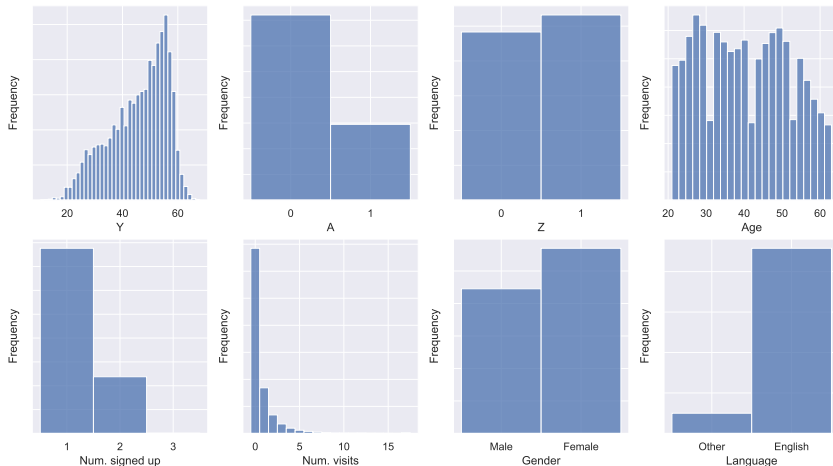


Figure 7: Histograms of each variable in our sample from OHIE.

The data collection in the OHIE was done follows: After excluding individuals below the age of 19, above the age of 64, and individuals with residence outside of Oregon, 74,922 individuals were considered for the lottery. Among those, 29,834 were selected randomly and were offered participation in the program. However, the probability of selection depended on the number of household members on the waiting list: for instance, an individual who signed up with another person was twice as likely to be selected. From the 74,922 individuals, 57,528 signed up alone, 17,236 signed up with another person, and 158 signed up with two more people on the waiting list. Thus, the probability of being selected conditional on the number of household members on the waiting list follows the multivariate version of Wallenius’ noncentral hypergeometric distribution (Chesson, 1976).

Propensity score: We computed the propensity score as follows. To account for the Wallenius’ noncentral hypergeometric distribution, we use the R package *BiasedUrn* to calculate the propensity score $\pi(x) = \mathbb{P}(Z = 1 \mid X = x)$. We obtained

$$\pi(x) = \begin{cases} 0.345, & \text{if individual } x \text{ signed up alone,} \\ 0.571, & \text{if individual } x \text{ signed up with one more person,} \\ 0.719, & \text{if individual } x \text{ signed up with two more people.} \end{cases} \quad (49)$$

³Data available here: <https://www.nber.org/programs-projects/projects-and-centers/oregon-health-insurance-experiment>

During the training of both MRIV and DRIV, we use the calculated values from Eq. (49) for the propensity score.

G DETAILS FOR BASELINE METHODS

In this section, we give a brief overview on the baselines which we used in our experiments. We implemented: (1) CATE methods for unconfoundedness: **TARNet** (Shalit et al., 2017) and TARNet combined with the **DR-learner** (Kennedy, 2022); (2) general IV methods, i.e., IV methods developed for IV settings with multiple or continuous instruments and treatments: **2SLS** (Wright, 1928), kernel IV (**KIV**) (Singh et al., 2019), **DFIV** (Xu et al., 2021a), **DeepIV** (Hartford et al., 2017), **DeepGMM** (Bennett et al., 2019), **DMLIV** (Syrkkanis et al., 2019), and DMLIV combined with **DRIV** (as described in (Syrkkanis et al., 2019)); (3) the (plug-in) Wald estimator using **linear models** and Bayesian additive regression trees (**BART**) (Chipman et al., 2010). Of note, the DR-learner assumes unconfoundedness, which is why we only combine it TARNet in our experiments. In the following, we provide details regarding methods and implementation.

G.1 CATE METHODS FOR UNCONFOUNDEDNESS

Many CATE methods assume *unconfoundedness*, i.e., that all confounders are observed in the data. Formally, the unconfoundedness assumption can be expressed in the potential outcomes framework as

$$Y(1), Y(0) \perp\!\!\!\perp A \mid X. \quad (50)$$

Under unconfoundedness, the CATE is identified as

$$\tau(x) = \mu_1(x) - \mu_0(x) \quad \text{with} \quad \mu_i(x) = \mathbb{E}[Y \mid A = i, X = x]. \quad (51)$$

Methods that assume unconfoundedness proceed by estimating $\mu_i(x) = \mathbb{E}[Y \mid A = i, X = x]$ from Eq. (51). However, if unobserved confounders U exist, it follows that

$$\tau(x) = \mathbb{E}[Y \mid A = 1, X = x, U] - \mathbb{E}[Y \mid A = 0, X = x, U] \neq \mu_1(x) - \mu_0(x), \quad (52)$$

which means that estimators that assume unconfoundedness are generally biased. Nevertheless, we include two baselines that assume unconfoundedness into our experiments: TARNet (Shalit et al., 2017) and the DR-learner (Kennedy, 2022).

TARNet (Shalit et al., 2017): TARNet (Shalit et al., 2017) is a neural network that estimates the CATE components $\mu_i(x)$ from Eq. 51 by learning a shared representation $\Phi(x)$ and two potential outcome heads $h_i(\Phi(x))$. We train TARNet by minimizing the loss

$$\mathcal{L}(\theta) = \sum_{i=1}^n L(h_{a_i}(\Phi(x_i, \theta_\Phi), \theta_{h_i}), y_i), \quad (53)$$

where $\theta = (\theta_{h_1}, \theta_{h_0}, \theta_\Phi)$ denotes the model parameters and L denotes squared loss if Y is continuous or binary cross entropy loss if Y is binary.

Note regarding balanced representations: In (Shalit et al., 2017), the authors propose to add an additional regularization term inspired from domain adaptation literature, which forces TARNet to learn a balanced representation $\Phi(x)$, i.e., that minimizes the distance the treatment and control group in the feature space. They showed that this approach leads to minimization of a generalization bound on the CATE estimation error if the representation is invertible.

In our experiments, we refrained from learning balanced representations because minimizing the regularized loss from (Shalit et al., 2017) does not necessarily result in an invertible representation and thus may even harm the estimation performance. For a detailed discussion, we refer to (Curth & van der Schaar, 2021). Furthermore, by leaving out the regularization, we ensure comparability between the different baselines. If balanced representations are desired, the balanced representation approach could also be extended to MRIV-Net, as we also build MRIV-Net on learning shared representations.

DR-learner (Kennedy, 2022): The DR-learner (Kennedy, 2022) is a meta learner that takes arbitrary estimators of the CATE components μ_i and the propensity score $\pi(x) = \mathbb{P}(A = 1 \mid X = x)$ as input and performs a pseudo-outcome regression by using the pseudo outcome

$$\hat{Y}_{\text{MR}} = \left(\frac{A}{\hat{\pi}(X)} - \frac{1-A}{1-\hat{\pi}(X)} \right) Y + \left(1 - \frac{A}{\hat{\pi}(X)} \right) \hat{\mu}_1(X) - \left(1 - \frac{1-A}{1-\hat{\pi}(X)} \right) \hat{\mu}_0(X). \quad (54)$$

In our experiments, we use TARNet as base method to provide initial estimators $\hat{\mu}_i(X)$. We further learn propensity score estimates $\hat{\pi}(X)$ by adding a separate representation to TARNet as done in (Shalit et al., 2017).

G.2 GENERAL IV METHODS

2SLS (Wright, 1928): 2SLS (Wright, 1928) is a linear two-stage approach. First, the treatments A are regressed on the instruments Z and fitted values \hat{A} are obtained. In the second stage, the outcome Y is regressed on \hat{A} . We implement 2SLS using the scikit-learn package.

KIV (Singh et al., 2019): Kernel IV (Singh et al., 2019) generalizes 2SLS to nonlinear settings. KIV assumes that the data is generated by

$$Y = f(A) + U, \quad (55)$$

where U is an additive unobserved confounder and f is some unknown (potentially nonlinear) structural function. KIV then models the structural function via

$$f(a) = \mu^t \psi(a) \quad \text{and} \quad \mathbb{E}[\psi(A) \mid Z = z] = V \phi(z), \quad (56)$$

where ψ and ϕ are feature maps. Here, kernel ridge regressions instead of linear regressions are used in both stages to estimate μ and V .

Following (Singh et al., 2019) we use the exponential kernel (Rasmussen & Williams, 2008) and set the length scale to the median inter-point distance. KIV does not provide a direct way to incorporate the observed confounders X . Hence, we augment both the instrument and the treatment with X , which is consistent with previous work (Bennett et al., 2019; Xu et al., 2021a). We also use two different samples for each stage as recommended in (Singh et al., 2019).

DFIV (Xu et al., 2021a): DFIV (Xu et al., 2021a) is a similar approach KIV in generalizing 2SLS to nonlinear setting by assuming Eq. (55) and Eq. (56). However, instead of using kernel methods, DFIV models the features maps ψ_{θ_A} and ϕ_{θ_Z} as neural networks with parameters θ_A and θ_Z , respectively. DFIV is trained by iteratively updating the parameters θ_A and θ_Z . The authors also provide a training algorithm that incorporates observed confounders X , which we implemented for our experiments. During training, we used two different datasets for each of the two stages as described in the paper.

DeepIV (Hartford et al., 2017): DeepIV (Hartford et al., 2017) also assumes additive unobserved confounding as in Eq. (55), but leverages the identification result (Newey & Powell, 2003)

$$\mathbb{E}[Y \mid X = x, Z = z] = \int h(a, x) dF(a \mid x, z), \quad (57)$$

where $h(a, x) = f(a, x) + \mathbb{E}[U \mid X = x]$ is the target counterfactual prediction function. DeepIV estimates $F(a \mid x, z)$, i.e., the conditional distribution function of the treatment A given observed covariates X and instruments Z , by using neural networks. Because we consider only binary treatments, we simply implement a (tunable) feed-forward neural network with sigmoid activation function. Then, DeepIV proceeds by learning a second stage neural network to solve the inverse problem defined by Eq. (57).

DeepGMM (Bennett et al., 2019): DeepGMM (Bennett et al., 2019) adopts neural networks for IV estimation inspired by the (optimally weighted) Generalized Method of Moments. The DeepGMM estimator is defined as the solution of the following minimax game:

$$\hat{\theta} \in \arg \min_{\theta \in \Theta} \sup_{\tau \in \mathcal{T}} \frac{1}{n} \sum_{i=1}^n f(z_i, \tau)(y_i - g(a_i, \theta)) - \frac{1}{4n} \sum_{i=1}^n f^2(z_i, \tau)(y_i - g(a_i, \tilde{\theta}))^2, \quad (58)$$

where $f(z_i, \cdot)$ and $g(a_i, \cdot)$ are parameterized by neural networks. As recommended in (Bennett et al., 2019), we solve this optimization via adversarial training with the Optimistic Adam optimizer (Daskalakis et al., 2018), where we set the parameter $\tilde{\theta}$ to the previous value of θ .

DMLIV (Syrkanis et al., 2019): DMLIV (Syrkanis et al., 2019) assumes that the data is generated via

$$Y = \tau(X)A + f(X) + U, \quad (59)$$

where τ is the CATE f some function of the observed covariates. First, DMLIV estimates the functions $q(X) = \mathbb{E}[Y \mid X]$, $h(Z, X) = \mathbb{E}[A \mid Z, X]$, and $p(X) = \mathbb{E}[A \mid X]$. Then, the CATE is learned by minimizing the loss

$$\mathcal{L}(\theta) = \sum_{i=1}^n (y_i - \hat{q}(x_i) - \hat{\tau}(x_i, \theta)(\hat{h}(z_i, x_i) - \hat{p}(x_i)))^2, \quad (60)$$

where $\hat{\tau}(X, \cdot)$ is some model for $\tau(X)$. In our experiments, we use (tunable) feed-forward neural networks for all estimators.

DRIV (Syrngkanis et al., 2019): DRIV (Syrngkanis et al., 2019) is a meta learner, originally proposed in combination with DMLIV. It requires initial estimators for $q(X)$, $p(X)$, $\pi(X) = \mathbb{E}[Z | X = x]$, and $f(X) = \mathbb{E}[A \cdot Z | X = x]$ as well as an initial CATE estimator $\hat{\tau}_{\text{init}}(X)$ (e.g., from DMLIV). The CATE is then estimated by a pseudo regression on the following doubly robust pseudo outcome:

$$\hat{Y}_{\text{DR}} = \hat{\tau}_{\text{init}}(X) + \frac{(Y - \hat{q}(X) - \hat{\tau}_{\text{init}}(X)(A - \hat{p}(X))Z - \hat{\pi}(X))}{\hat{f}(X) - \hat{p}(X)\hat{r}(X)}. \quad (61)$$

We implement all regressions using (tunable) feed-forward neural networks.

Comparison between DRIV vs. MRIV: There are two key differences between our paper and (Syrngkanis et al., 2019): (i) In contrast to DRIV, we showed that our MRIV is multiply robust. (ii) We derive a multiple robust convergence rate, while the rate in (Syrngkanis et al., 2019) is not robust with respect to the nuisance rates.

Ad (i): Both MRIV and DRIV perform a pseudo-outcome regression on the efficient influence function (EIF) of the ATE. The key difference: DRIV uses the doubly robust parametrization of the EIF from (Okui et al., 2012), whereas we use the multiply robust parametrization of the EIF from (Wang & Tchetgen Tchetgen, 2018)⁴. Hence, our MRIV frameworks extends DRIV in a non-trivial way to achieve multiple robustness. Thus, our estimator is consistent in the union of *three* different model specifications.⁵

Ad (ii): Here, we compare the convergence rates from DRIV and our MRIV and, thereby, show the strengths of our MRIV. To this end, let us assume that the pseudo regression function is γ -smooth and that we use the same second-stage estimator \hat{E}_n with minimax rate $n^{-\frac{2\gamma}{2\gamma+p}}$ for both DRIV and MRIV. If the nuisance parameters $q(X)$, $p(X)$, $f(X)$, and $\pi(X)$ are α -smooth and further are estimated with minimax rate $n^{\frac{-2\alpha}{2\alpha+p}}$, Corollary 4 from (Syrngkanis et al., 2019) states that DRIV converges with rate

$$\mathbb{E} \left[(\hat{\tau}_{\text{DRIV}}(x) - \tau(x))^2 \right] \lesssim n^{\frac{-2\gamma}{2\gamma+p}} + n^{\frac{-4\alpha}{2\alpha+p}}.$$

In contrast, MRIV assumes estimation of the nuisance parameters $\mu_0^Y(x)$ with rate $n^{\frac{-2\alpha}{2\alpha+p}}$, $\mu_0^A(x)$ and $\delta_A(x)$ with rate $n^{\frac{-2\beta}{2\beta+p}}$, and $\pi(x)$ with rate $n^{\frac{-2\delta}{2\delta+p}}$. If the initial estimator $\hat{\tau}_{\text{init}}(x)$ converges with rate $r_\tau(n)$, our Theorem 2 yields the rate

$$\mathbb{E} \left[(\hat{\tau}_{\text{MRIV}}(x) - \tau(x))^2 \right] \lesssim n^{\frac{-2\gamma}{2\gamma+p}} + r_\tau(n) \left(n^{\frac{-2\beta}{2\beta+p}} + n^{\frac{-2\delta}{2\delta+p}} \right) + n^{-2\left(\frac{\alpha}{2\alpha+p} + \frac{\delta}{2\delta+p}\right)} + n^{-2\left(\frac{\beta}{2\beta+p} + \frac{\delta}{2\delta+p}\right)}.$$

If all nuisance parameters converge with the same minimax rate of $n^{\frac{-2\alpha}{2\alpha+p}}$, the rates of DRIV and our MRIV coincide. However, different to DRIV, our rate is additionally multiple robust in spirit of Theorem 1. This presents a crucial strength of our MRIV over DRIV: For example, if δ is small (slow convergence of $\hat{\pi}(x)$), our MRIV still with fast rate as long as α and β are large (i.e., if the other nuisance parameters are sufficiently smooth).

G.3 WALD ESTIMATOR

Finally, we consider the Wald estimator (Wald, 1940) for the binary IV setting. More precisely, we estimate the CATE components $\mu_i^Y(x)$ and $\mu_i^A(x)$ separately and plug them into

$$\tau(x) = \frac{\hat{\mu}_1^Y(x) - \hat{\mu}_0^Y(x)}{\hat{\mu}_1^A(x) - \hat{\mu}_0^A(x)}. \quad (62)$$

We consider two versions of the Wald estimator:

⁴For a detailed discussion on multiple robustness and the importance of the EIF parametrization, we refer to (Wang & Blei, 2019), Section 4.5.

⁵On a related note, a similar, important contribution of developing multiply robust method was recently made for the average treatment effect. Here, the estimator of (Okui et al., 2012) was extended by the estimator of (Wang & Tchetgen Tchetgen, 2018) to allow for multi robustness. Yet, this different from our work in that it focuses on the average treatment effect, while we study the conditional average treatment effect in our paper.

Linear: We use linear regressions to estimate the $\mu_i^Y(x)$ and logistic regressions to estimate the $\mu_i^A(x)$.

BART: We use Bayesian additive regression trees (Chipman et al., 2010) trees to estimate the $\mu_i^Y(x)$ and random forest classifier to estimate the $\mu_i^A(x)$.

H VISUALIZATION OF PREDICTED CATES

We plot the predicted CATEs for the different baselines and MRIV-Net in Fig. 8 (for $n = 3000$). As expected, the linear methods (2SLS and linear Wald) are not flexible enough to provide accurate CATE estimates. We also observe that the curve of MRIV-Net without MRIV is quite wiggly, i.e., the estimator has a relatively large variance. This variance is reduced when the full MRIV-Net is applied. As a result, curve is much smoother. This is reasonable because MRIV does not estimate the CATE components individually, but estimates the CATE directly via the Stage 2 pseudo-outcome regression. Overall, this confirms the superiority of our proposed framework.

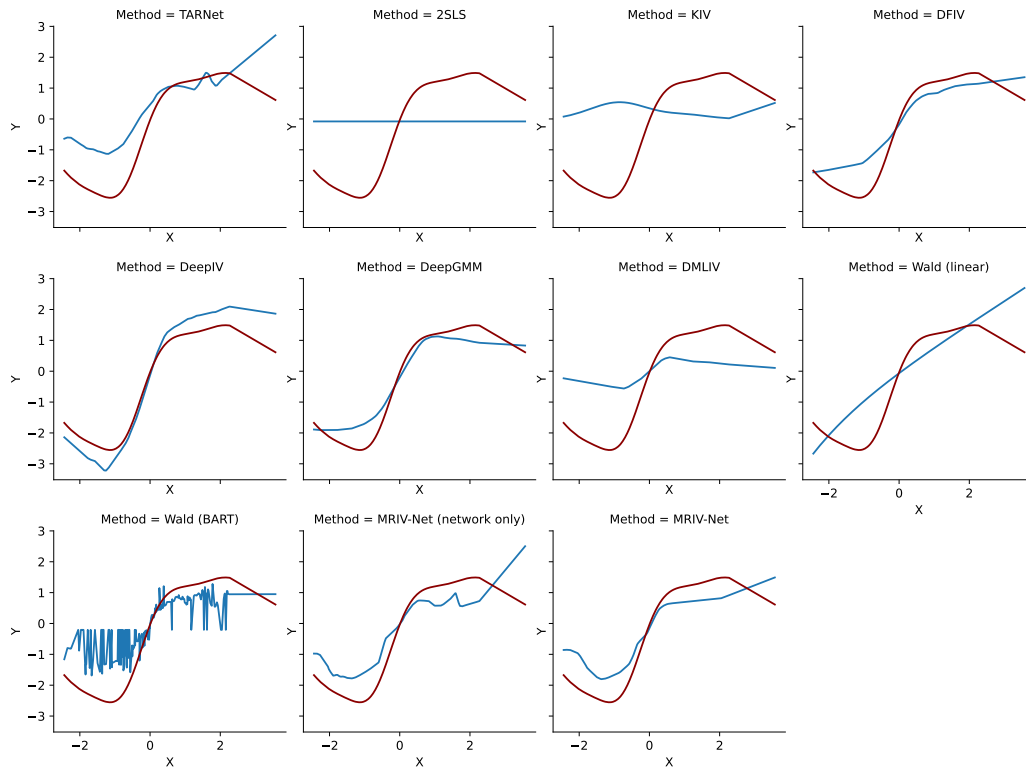


Figure 8: Predicted CATEs (blue) and oracle CATE (red) for different baselines.

I IMPLEMENTATION DETAILS AND HYPERPARAMETER TUNING

Implementation details for deep learning models: To make the performance of the deep learning models comparable, we implemented all feed-forward neural networks (including MRIV-Net) as follows: We use two hidden layers with RELU activation functions. We also incorporated a dropout layer for each hidden layer. We trained all models with the Adam optimizer (Kingma & Ba, 2015) using 100 epochs. Exceptions are only DFIV and DeepGMM, where we used 200 epochs for training, accounting for slower convergence of the respective (adversarial) training algorithms. For DeepGMM, we further used Optimistic Adam (Daskalakis et al., 2018) as in the original paper.

Training times: We report the approximate times needed to train the deep learning models on our simulated data with $n = 5000$ in Table 5. For training, we used an AMD Ryzen Pro 7 CPU. Compared to DMLIV and DRIV, the training of MRIV-Net is faster because only a single neural network is trained.

Table 5: Training times for deep learning models (in seconds).

TARNet	TARNet + DR	DFIV	DeepIV	DeepGMM	DMLIV	DMLIV + DRIV	MRIV-Net
~10.62	~28.57	~164.98	~30.21	~17.31	~74.98	~91.12	~32.20

Hyperparameter tuning: We performed hyperparameter tuning for all deep learning models (including MRIV-Net), KIV, and the BART Wald estimator on all datasets. For all methods except KIV and DFIV, we split the data into a training set (80%) and a validation set (20%). We then performed 40 random grid search iterations and chose the set of parameters that minimized the respective training loss on the validation set. In particular, the tuning procedure was the same for all baselines, which ensures that the performance gain of MRIV-Net is due to the method itself and not due to larger flexibility. Exceptions are only KIV and DFIV, for which we implemented the customized hyperparameter tuning algorithms proposed in (Singh et al., 2019) and (Xu et al., 2021a) to ensure consistency with prior literature. For the meta learners (DR-learner, DRIV, and MRIV), we first performed hyperparameter tuning for the base methods and nuisance models, before tuning the pseudo-outcome regression neural network by using the input from the tuned models. The tuning ranges for the hyperparameter are shown in Table 6. These include both the hyperparameter rangers shared across all neural networks and the model-specific hyperparameters. For reproducibility purposes, we publish the selected hyperparameters in our GitHub project as `.yaml` files.

Table 6: Hyperparameter tuning ranges.

MODEL	HYPERPARAMETER	TUNING RANGE
Feed-forward neural networks (Shared parameter ranges for all deep learning baselines)	Hidden layer size(es)	$p, 5p, 10p, 20p, 30p$ (simulated data) $p, 3p, 5p, 8p, 10p$ (OHIE)
	Learning rate	0.0001, 0.0005, 0.001, 0.005, 0.01
	Batch size	64, 128, 256
	Dropout probability	0, 0.1, 0.2, 0.3
KIV	λ (Ridge penalty first stage)	5, 6, 7, 8, 9, 10, 12
	ξ (Ridge penalty second stage)	5, 6, 7, 8, 9, 10, 12
DFIV	λ_1 (Ridge penalty first stage)	0.0001, 0.001, 0.01, 0.1 (simulated data) 0.01, 0.05, 0.1 (OHIE)
	λ_2 (Ridge penalty second stage)	0.0001, 0.001, 0.01, 0.1 (simulated data) 0.01, 0.05, 0.1 (OHIE)
DeepGMM	λ_f (learning rate multiplier)	0.5, 1, 1.5, 2, 5
Wald (BART)	Number of trees (BART)	20, 30, 40, 50
	Number of trees (Random forest classifier)	20, 30, 40, 50

p = network input size

Hyperparameter robustness checks: We also investigate the robustness of MRIV-Net with respect to hyperparameter choice. To to this, we fix the optimal hyperparameter constellation for our simulated data for $n = 3000$ and perturb the hidden layer sizes, learning rate, dropout probability, and batch size. The results are shown in Fig. 9. We observe that the RMSE only changes marginally when perturbing the different hyperparameters, indicating that our method is to a certain degree robust against hyperparameter misspecification. Furthermore, our results indicate that the performance improvement

of MRIV-Net over the baselines observed in our experiments is not due to hyperparameter tuning, but to our method itself.




Figure 9: Robustness checks for different hyperparameters of MRIV-Net.

J RESULTS FOR SEMI-SYNTHETIC DATA

In the main paper, we evaluated MRIV-Net both on synthetic and real-world data. Here, we provide additional results by constructing a semi-synthetic dataset on the basis of OHIE. It is common practice in causal inference literature to use semi-synthetic data for evaluation, because it combines advantages of both synthetic and real-world data. On the one hand, the real-world data part ensures that the data distribution is realistic and matches those in practice. On the other hand, the counterfactual ground-truth is still available, which makes it possible to measure the performance of CATE methods.

We construct our semi-synthetic data as follows: First, we extract the covariates $X \in \mathbb{R}^5$ and instruments $Z \in \{0, 1\}$ of our OHIE dataset from Sec. F. Then, we construct the treatment components $\mu_i^A(x)$ via

$$\mu_1^A(X) = 0.3 \cdot \sigma(X_1) + 0.7 \quad \text{and} \quad \mu_0^A(X) = 0.3 \cdot \sigma(X_1), \quad (63)$$

where X_1 is the (standardized) age and $\sigma(\cdot)$ is the sigmoid function. The outcome components are constructed via

$$\mu_1^Y(X) = 0.5X_1^2 + \sum_{i=2}^5 X_i^2 \quad \text{and} \quad \mu_0^Y(X) = -0.5X_1^2 + \sum_{i=2}^5 X_i^2. \quad (64)$$

We then sample treatments A and outcomes Y as in Eq. (41) and Eq. (42). Lemma 3 ensures that $\mu_i^Y(X) = \mathbb{E}[Y | Z = i, X]$ and $\mu_i^A(X) = \mathbb{E}[A | Z = i, X]$.

Given the above, the oracle CATE becomes

$$\tau(X) = \frac{X_1^2}{0.7}. \quad (65)$$

Note that $\tau(X)$ is sparse in the sense that it only depends on age, while the outcome components depend on all five covariates. Following our theoretical analysis in Sec. C, MRIV-Net should thus outperform methods that aim at estimating the components directly. This is confirmed in Table 7, where we show the results for all baselines and MRIV-Net on the semi-synthetic data. Indeed, we observe that MRIV-Net outperforms all other baselines, confirming both the superiority of our method as well as our theoretical results under sparsity assumptions from Sec. C.

Table 7: Results for semi-synthetic data.

Method	$n = 3000$	$n = 5000$	$n = 8000$
(1) STANDARD ITE			
TARNet (Shalit et al., 2017)	1.66 ± 0.11	1.58 ± 0.07	1.57 ± 0.11
TARNet + DR (Shalit et al., 2017; Kennedy, 2022)	1.31 ± 0.28	1.22 ± 0.37	1.12 ± 0.15
(2) GENERAL IV			
2SLS (Wooldridge, 2013)	1.34 ± 0.06	1.31 ± 0.03	1.32 ± 0.02
KIV (Singh et al., 2019)	1.97 ± 0.10	1.92 ± 0.05	1.93 ± 0.05
DFIV (Xu et al., 2021a)	1.67 ± 0.44	1.63 ± 0.47	1.45 ± 0.17
DeepIV (Hartford et al., 2017)	1.24 ± 0.26	0.99 ± 0.22	0.84 ± 0.19
DeepGMM (Bennett et al., 2019)	1.39 ± 0.03	1.37 ± 0.16	1.18 ± 0.16
DMLIV (Syrgkanis et al., 2019)	2.12 ± 0.10	2.09 ± 0.09	2.02 ± 0.11
DMLIV + DRIV (Syrgkanis et al., 2019)	1.22 ± 0.10	1.18 ± 0.19	1.00 ± 0.08
(3) WALD ESTIMATOR (WALD, 1940)			
Linear	1.42 ± 0.24	1.28 ± 0.07	1.32 ± 0.07
BART	1.48 ± 0.24	1.29 ± 0.04	1.06 ± 0.13
MRIV-Net (network only)	1.11 ± 0.15	0.84 ± 0.14	0.95 ± 0.21
MRIV-Net (ours)	0.71 ± 0.24	0.75 ± 0.18	0.78 ± 0.26

Reported: RMSE (mean ± standard deviation). Lower = better (best in bold)

K RESULTS FOR CROSS-FITTING

Here, we repeat our experiments from the main paper but now make use of *cross-fitting*. Recall that, in Theorem 2, we assume that the nuisance parameter estimation and the pseudo-outcome regression are performed on three independent samples. We now address this through *cross-fitting*. To this end, our aim is to show that our proposed MRIV framework is again superior.

For MRIV, we proceeded as follows: We split the sample \mathcal{D} into three equally sized samples \mathcal{D}_1 , \mathcal{D}_2 , and \mathcal{D}_3 . We then trained $\hat{\tau}_{init}(x)$, $\hat{\mu}_0^Y(x)$, and $\hat{\mu}_0^A(x)$ on \mathcal{D}_1 , $\hat{\delta}_A(x)$ and $\hat{\pi}(x)$ on \mathcal{D}_2 , and performed the pseudo-outcome regression on \mathcal{D}_3 . Then, we repeated the same training procedure two times, but performed the pseudo-outcome regression on \mathcal{D}_2 and \mathcal{D}_1 . Finally, we averaged the resulting three CATE estimators. For DRIV, we implemented the cross-fitting procedure described in (Syrkkanis et al., 2019). For the DR-learner, we followed (Kennedy, 2022).

The results are in Table 8. Importantly, the results confirm the effectiveness of our proposed MRIV. Overall, we find that our proposed MRIV outperforms DRIV for the vast majority of base methods when performing cross-fitting. Furthermore, MRIV-Net is highly competitive even when comparing it with the cross-fitted estimators. This shows that our heuristic to learn separate representations instead of performing sample splits works in practice. In sum, the results confirm empirically that our MRIV is superior.

Table 8: Results for base methods with different meta-learners (i.e., DRIV, and our MRIV) using cross-fitting and results for MRIV-Net without cross-fitting.

Meta-learners		$n = 3000$		$n = 5000$		$n = 8000$	
		DRIV	MRIV (ours)	DRIV	MRIV (ours)	DRIV	MRIV (ours)
(1) STANDARD ITE							
TARNet (Shalit et al., 2017)		0.30 ± 0.02	0.36 ± 0.16	0.18 ± 0.06	0.16 ± 0.03	0.21 ± 0.08	0.13 ± 0.04
TARNet + DR-learner (Shalit et al., 2017; Kennedy, 2022)			0.85 ± 0.11		0.66 ± 0.08		0.67 ± 0.12
(2) GENERAL IV							
2SLS (Wooldridge, 2013)		0.42 ± 0.11	0.33 ± 0.09	0.20 ± 0.07	0.23 ± 0.11	0.24 ± 0.10	0.14 ± 0.02
KIV (Singh et al., 2019)		0.47 ± 0.18	0.45 ± 0.15	0.20 ± 0.06	0.19 ± 0.08	0.22 ± 0.04	0.15 ± 0.03
DFIV (Xu et al., 2021a)		0.35 ± 0.05	0.28 ± 0.09	0.22 ± 0.10	0.18 ± 0.08	0.24 ± 0.12	0.16 ± 0.04
DeepIV (Hartford et al., 2017)		0.38 ± 0.09	0.44 ± 0.16	0.20 ± 0.07	0.19 ± 0.07	0.20 ± 0.08	0.12 ± 0.02
DeepGMM (Bennett et al., 2019)		0.42 ± 0.09	0.42 ± 0.16	0.19 ± 0.04	0.19 ± 0.07	0.22 ± 0.06	0.13 ± 0.02
DMLIV (Syrkkanis et al., 2019)		0.44 ± 0.09	0.46 ± 0.16	0.21 ± 0.04	0.19 ± 0.07	0.21 ± 0.05	0.14 ± 0.02
(3) WALD ESTIMATOR (WALD, 1940)							
Linear		0.47 ± 0.23	0.36 ± 0.12	0.24 ± 0.05	0.20 ± 0.08	0.22 ± 0.05	0.15 ± 0.02
BART		0.43 ± 0.12	0.39 ± 0.12	0.14 ± 0.05	0.13 ± 0.05	0.23 ± 0.08	0.15 ± 0.02
MRIV-Net w network only (ours)		0.35 ± 0.12	0.26 ± 0.11	0.19 ± 0.13	0.15 ± 0.03	0.18 ± 0.08	0.13 ± 0.03

Reported: RMSE (mean ± standard deviation). Lower = better (best in bold)

L FURTHER EXPERIMENTAL RESULTS ON REAL-WORLD DATA

In Section 5.2, we estimated the ITE on the OHIE data and visualized the treatment heterogeneity with respect to *age* and *gender*. In this section, we provide additional results and also visualize the heterogeneity with respect to age as well as *additional covariates*. These are: (1) the number of emergency visits a patient has in its history before signing up for the lottery and (2) the language spoken by the patient (English or other). We fixed the gender to “female”.

For (1), we plot the estimated ITE for three different age groups over the number of emergency visits. The results are shown in Fig. 10.




Figure 10: Estimated treatment effects for different age and number of emergency visits.

We observe that all methods tend to estimate a larger effect for individuals who had more emergency visits in their patient history. However, the IV methods (in particular our MRIV-Net) estimate a much larger effect for patients with many visits. In contrast to the other methods, MRIV-Net also estimates larger effects for older than for younger patients. The results provided by MRIV-Net seem intuitive, as older patients with a history of emergency visits should be exposed to higher health-related risks, thus benefiting from health insurance. The fact that TARNet consistently estimates small (and even negative) effects could be an indicator of bias due to unobserved confounding.

For (2), we plot the estimated ITE for three different age groups over the spoken language. The results are shown in Fig. 11.

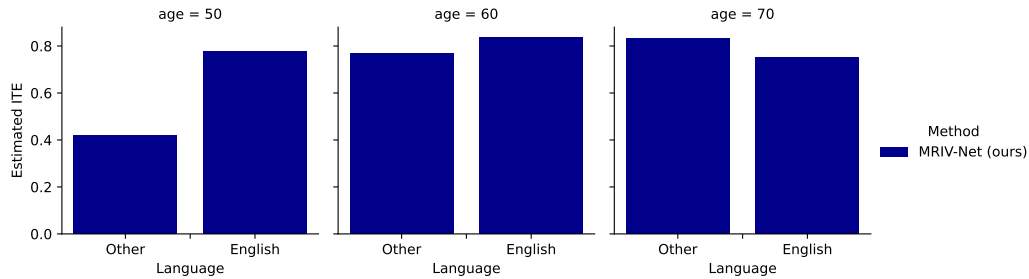


Figure 11: Estimated treatment effects for different age and language

For patients of age 50, our MRIV-Net estimates a higher effect for the English-speaking patients. Interestingly, for older patients, the estimated effect increases also for non-English speaking patients.