Single Proxy Identifiability of Causal Effects

under unobserved confounding and a known error mechanism

Silvan Vollmer
Department of Mathematics
University of Copenhagen

Sebastian Weichwald, Niklas Pfister

arXiv paper - April, 2026

KØBENHAVNS UNIVERSITET

Motivation

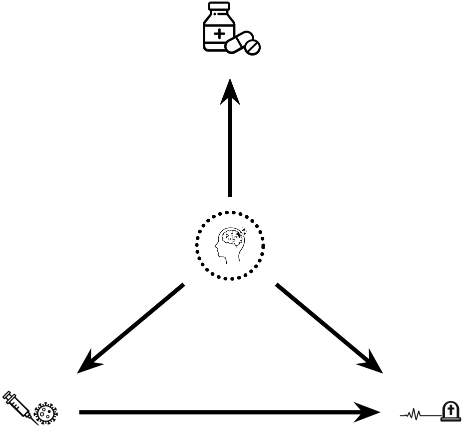

Jackson et al. (2005); Wyss et al. (2022)

Introduction

treatment $X \in \mathcal{X} \subseteq \mathbb{R}^p$
outcome $Y \in \mathcal{Y} \subseteq \mathbb{R}$
confounder $U \in \mathcal{U} \subseteq \mathbb{R}^k$
proxy $W \in \mathcal{W} \subseteq \mathbb{R}^d$

$p, k, d \in \mathbb{N}$

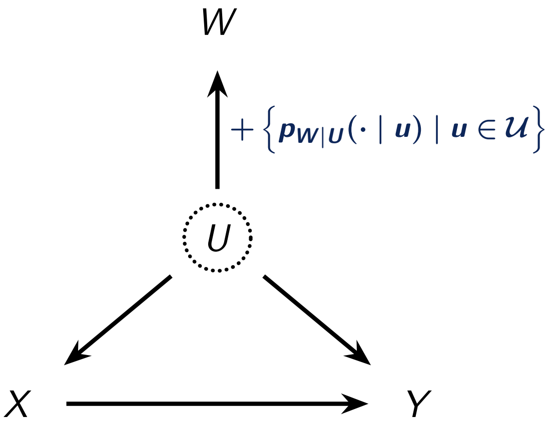

Introduction

We identify
causal effects
under **unobserved confounding**
and a **known error mechanism**.

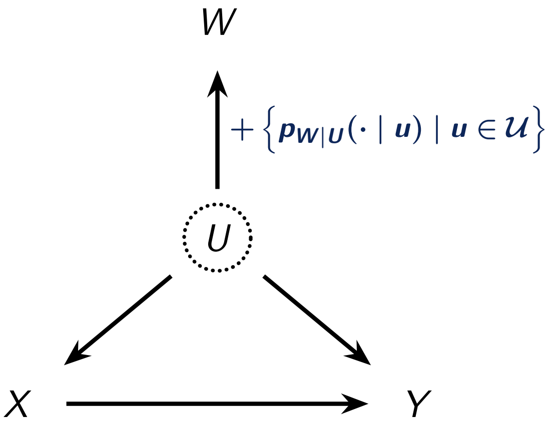

Pearl (2010); Kuroki and Pearl (2014)

Introduction

We identify causal effects under unobserved confounding and a **known error mechanism**.

differential privacy

U.S. Census Bureau (2021)

Overview

1. Model and causal function
2. Identifiability
3. Method
4. Application

1. Model

We assume a **Proxy-Confounded Structural Causal Model** (**PC-SCM**).

$$M : \begin{cases} U & \leftarrow N_U \\ W & \leftarrow f_W(U, E) \\ X & \leftarrow f_X(U, N_X) \\ Y & \leftarrow f_Y(U, X, N_Y) \end{cases}$$

N_U, E, N_X, N_Y jointly independent

$$\implies \quad W \perp\!\!\!\perp (X, Y) \mid U$$

- joint distribution P^M has a density p^M
- interventional distribution $P_Y^{M;\, do(X:=x)}$ for each $do(X = x)$
- observe n i.i.d. copies of (X, Y, W) generated by M_0
- error mechanism $\{p_{W|U}^{M_0}(\cdot \mid u) \mid u \in \mathcal{U}\}$ is known

1. Causal function

Definition (causal function)

The **causal function** $\theta^{M_0} : \mathcal{X} \to \mathbb{R}^q$ is defined for all $x \in \mathcal{X}$ as

$$\theta^{M_0}(x) := \mathbb{E}_Y^{M_0;do(X:=x)}[Y].$$

- binary treatment $X \in \{0, 1\}$:

$$\text{ATE} = \theta^{M_0}(1) - \theta^{M_0}(0)$$

- Adjusting for the measurement W does not suffice.

2. Identifiability

The causal function is identifiable if it is the same across all PC-SCMs that induce the same observational features.

2. Identifiability

The causal function is identifiable if it is the same across all PC-SCMs that induce the same **observational features**.

- observational features: $\mathcal{F}^M := \{p^M_{W,X,Y}\} \cup \{p^M_{W|U}(\cdot \mid u) \mid u \in \mathcal{U}\}$

2. Identifiability

The causal function is identifiable if it is the same across **all PC-SCMs that induce the same observational features**.

- observational features: $\mathcal{F}^M := \{p^M_{W,X,Y}\} \cup \{p^M_{W|U}(\cdot \mid u) \mid u \in \mathcal{U}\}$
- for all observationally equivalent PC-SCMs $M \in \mathcal{M}(M_0)$, it holds that $\mathcal{F}^M = \mathcal{F}^{M_0}$

2. Identifiability

The causal function is identifiable if it is the same across all PC-SCMs that induce the same observational features.

- observational features: $\mathcal{F}^M := \{p^M_{W,X,Y}\} \cup \{p^M_{W|U}(\cdot \mid u) \mid u \in \mathcal{U}\}$
- for all observationally equivalent PC-SCMs $M \in \mathcal{M}(M_0)$, it holds that $\mathcal{F}^M = \mathcal{F}^{M_0}$

Definition (identifiability)

θ^{M_0} is **identifiable** if for all models $M \in \mathcal{M}(M_0)$, it holds that

$$\theta^M \equiv \theta^{M_0}.$$

2. Identifiability

The causal function is identifiable if the error mechanism is

(i) **L_1-complete**, that is, for all $\delta \in L_1(\mathcal{U})$

$$\int_{\mathcal{U}} p^{M_0}_{W|U}(\cdot \mid u)\delta(u)\,\mu(du) \equiv 0 \qquad \Longrightarrow \qquad \delta \equiv 0 \quad \text{a.e.}$$

or

(ii) **L_∞-complete**, that is, for all $\delta \in L_\infty(\mathcal{U})$

$$\int_{\mathcal{U}} p^{M_0}_{W|U}(\cdot \mid u)\delta(u)\,\mu(du) \equiv 0 \qquad \Longrightarrow \qquad \delta \equiv 0 \quad \text{a.e.}$$

and for all $M \in \mathcal{M}(M_0)$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ it holds that $p^M_{U|X,Y}(\cdot \mid x, y) \in L_\infty$.

2. Discrete confounders and proxies

Theorem (discrete)

The error mechanism is L_1-complete and L_∞-complete if

(i) (**sufficient dimension**) we have a discrete confounder $U \in \mathcal{U} = \{u_1, \ldots, u_k\}$ and a proxy $W \in \mathcal{W} = \{w_1, \ldots, w_d\}$ with $d \geq k$ and

(ii) (**full rank error mechanism**) there exists a subset $\mathcal{W}_r := \{w'_1, \ldots, w'_r\} \subseteq \mathcal{W}$ with cardinality $r \in \mathbb{N}$ and $k \leq r \leq d$ such that the matrix

$$\begin{pmatrix} p^{M_0}_{W|U}(w'_1 \mid u_1) & \cdots & p^{M_0}_{W|U}(w'_1 \mid u_k) \\ \vdots & \ddots & \vdots \\ p^{M_0}_{W|U}(w'_r \mid u_1) & \cdots & p^{M_0}_{W|U}(w'_r \mid u_k) \end{pmatrix} \in [0,1]^{r \times k}$$

has full column rank.

2. Continuous confounders and proxies

Theorem (continuous)

Assume that the SPICE Assumption holds.
Then, the error mechanism $\{p^{M_0}_{W|U}(\cdot \mid u) \mid u \in \mathbb{R}^k\}$ is L_∞-complete.

2. Continuous confounders and proxies

Assumption (SPICE)

(i) (**sufficient dimension**) The proxy $W \in \mathcal{W} \subseteq \mathbb{R}^d$ has at least the dimension of the confounder $U \in \mathcal{U} \subseteq \mathbb{R}^k$, that is, $d \geq k$ and

(ii) (**additive noise**) for a matrix $A \in \mathbb{R}^{d \times k}$ with full column rank and a random variable $E \in \mathcal{E} \subseteq \mathbb{R}^d$, we have

$$W = AU + E \quad \text{with} \quad U \perp\!\!\!\perp E$$

and

(iii) (**density**) the distribution of E has a density $p_E^{M_0}$ on $\mathbb{R}^d$ and

(iv) (**non-vanishing Fourier**) the Fourier transform of $p_E^{M_0}$ has no zeros.

3. Method: SPICE-Net

we know for all $(u, w) \in \mathcal{U} \times \mathcal{W}$ that
$p^{M_0}_{W|U}(w \mid u) = p^{M_0}_{E}(w - Au)$

1.

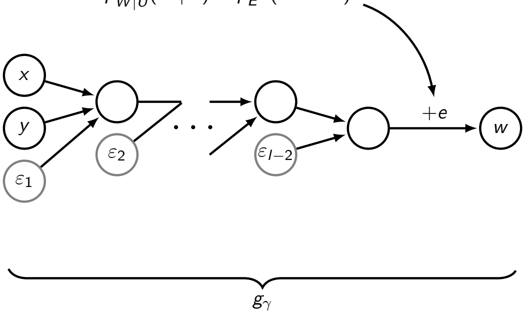

optimise weights γ by minimising energy loss Shen and Meinshausen (2025)

2. causal function estimator

guarantees

- $g^*_\gamma - E \sim P^{M_0}_{\widetilde{A}U|(X,Y)=(x,y)}$ for almost all $(x, y) \in \mathcal{X} \times \mathcal{Y}$
- adjusting for $\widetilde{A}U$ is fine

4. Simulation

Gaussian

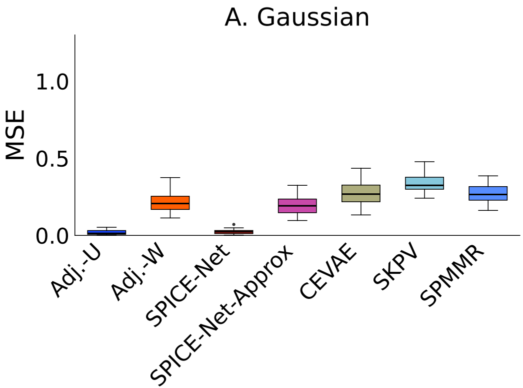

$$
\begin{aligned}
U &\quad \mathcal{N}(0,1) \\
f_W &\quad U + \mathcal{N}(0,1) \\
f_X &\quad U + \mathcal{N}(0,1) \\
f_Y &\quad U + X + \mathcal{N}(0,1)
\end{aligned}
$$

Louizos et al. (2017)
Xu and Gretton (2025)

4. Simulation

Binary

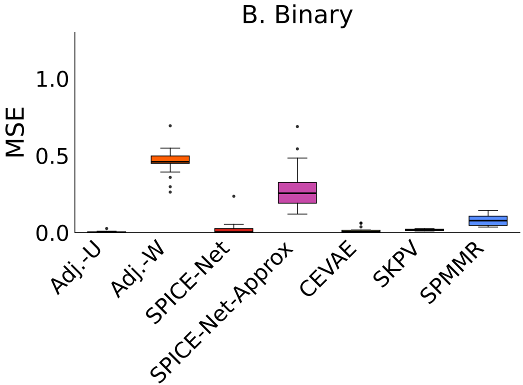

$$
\begin{aligned}
U &\quad \mathcal{N}(0,1) \\
f_W &\quad U + \mathcal{N}(0,1) \\
f_X &\quad \mathbb{1}\{(U + \mathcal{N}(0,1)) > 0\} \\
f_Y &\quad U + X + \mathcal{N}(0,1)
\end{aligned}
$$

4. Simulation

Exponential

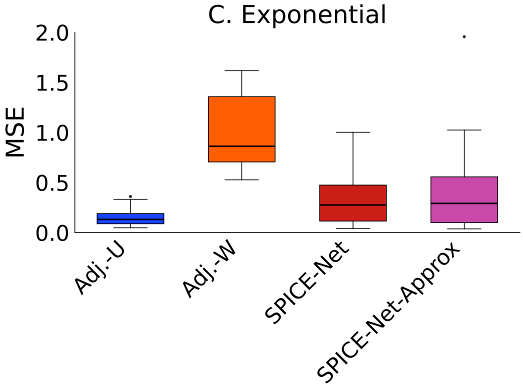

U	$\mathrm{Exp}(1)$
f_W	$U + \mathrm{Exp}(1)$
f_X	$U + \mathrm{Exp}(1)$
f_Y	$X^2 + UX + \mathrm{Exp}(1)$

4. Simulation

High-dimensional

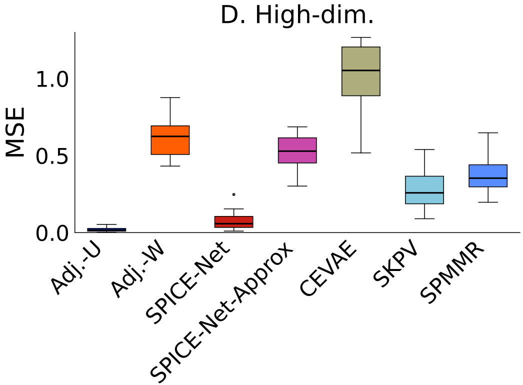

$$U \quad \mathcal{N}_2(0, I)$$
$$f_W \quad AU + E$$
$$f_X \quad BU + \mathcal{N}(0,1)$$
$$f_Y \quad BU + X + \mathcal{N}(0,1)$$

with

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 1 \end{bmatrix},$$

$$E \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2.0 & -0.3 & 0.5 \\ -0.3 & 1.5 & 0.4 \\ 0.5 & 0.4 & 1.8 \end{bmatrix} \right)$$

4. Real-world data

Causal Chamber® Gamella et al. (2025)

A Light Tunnel Mk2

B Chamber diagram

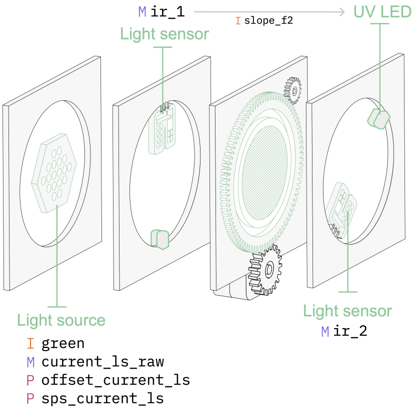

C Ground-truth graph

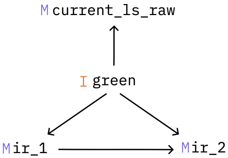

4. Real-world data

Causal Chamber® Gamella et al. (2025)

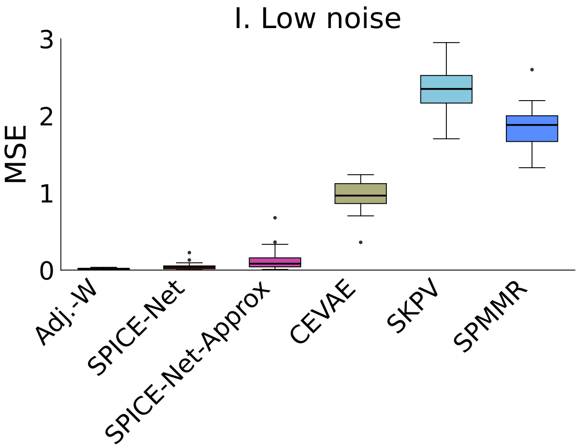

4. Real-world data

Causal Chamber® Gamella et al. (2025)

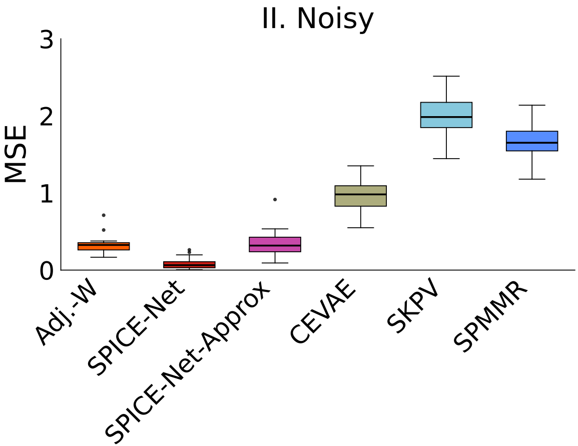

4. Conclusion

1. Model

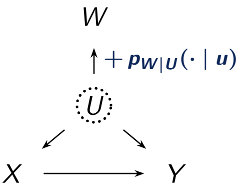

identifiability

$$\mathcal{F}^{M} = \mathcal{F}^{M_0} \;\Rightarrow\; \theta^{M} \equiv \theta^{M_0}$$

2. SPICE

The error mechanism is complete if

- sufficient proxy dimension $d \geq k$
- $W = AU + E$ with $U \perp\!\!\!\perp E$
- non-vanishing Fourier of $p_E^{M_0}$

3. SPICE-Net

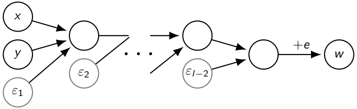

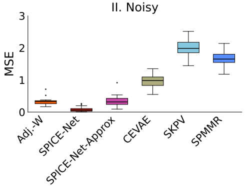

References

Juan L. Gamella, Jonas Peters, and Peter Bühlmann. Causal chambers as a real-world physical testbed for ai methodology. *Nature Machine Intelligence*, 7:107–118, 2025.

Lisa A. Jackson, Jennifer C. Nelson, Patti Benson, Kathleen M. Neuzil, Robert J. Reid, Bruce M. Psaty, Susan R. Heckbert, Eric B. Larson, and Noel S. Weiss. Functional status is a confounder of the association of influenza vaccine and risk of all cause mortality in seniors. *International Journal of Epidemiology*, 35(2):345–352, 2005.

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

Christos Louizos, Uri Shalit, Joris M. Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pages 6449–6459, Long Beach, California, 2017.

Judea Pearl. On measurement bias in causal inference. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 425–432, Santa Catalina Island, California, 2010.

Xinwei Shen and Nicolai Meinshausen. Engression: extrapolation through the lens of distributional regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 87(3):653–677, 2025.

U.S. Census Bureau. Disclosure avoidance for the 2020 census: An introduction. Technical report, U.S. Government Publishing Office, Washington, D.C., 2021.

Richard Wyss, Chen Yanover, Tal El-Hay, Dimitri Bennett, Robert W. Platt, Andrew R. Zullo, Grammati Sari, Xuerong Wen, Yizhou Ye, Hongbo Yuan, Mugdha Gokhale, Elisabetta Patorno, and Kueiyu J. Lin. Machine learning for improving high-dimensional proxy confounder adjustment in healthcare database studies: An overview of the current literature. *Pharmacoepidemiology and Drug Safety*, 31(9):932–943, 2022.

Liyuan Xu and Arthur Gretton. Kernel single proxy control for deterministic confounding. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, pages 3736–3744, Mai Khao, Thailand, 2025.