TECHNICAL REPORT
R-140
Mar, 2026

Relational Structural Causal Models

Adiba Ejaz [1] **Elias Bareinboim** [1]

Abstract

An artificial intelligence must have a model of its environment that is *causal*, supporting reasoning about interventions and counterfactuals, and also *combinatorial*, supporting generalization to unseen combinations of objects. In this work, we formally study when and how such a model can be learned. We develop *relational structural causal models*, extending structural causal models (Pearl, 2009) to settings where objects and their relations vary. First, we show how answers to not only causal but also observational queries about unseen combinations of objects can not be identified without further assumptions. To enable such identification—including in the presence of unobserved confounding—we define *relational causal graphs* and derive symbolic identification criteria. Finally, we propose *relational neural causal models*, a provably correct approach that outperforms non-relational baselines on simulated traffic scenes with varying cars, signals, and pedestrians.

1. Introduction

Behind a Rube Goldberg machine is a sequence of simple mechanisms. A ball rolls down a ramp, tipping a weight, pulling a string, swinging a hammer, and striking a gong. Predicting what happens next and why requires a model of how these bodies interact. This is precisely what *world models* aim to provide for AI systems to learn efficiently and generalize across environments (Ha & Schmidhuber, 2018; LeCun, 2022; Gurnee & Tegmark, 2024; Richens & Everitt, 2024; Vafa et al., 2024). In this work, we consider two important problems that such a model must address.

The first problem is that of representing objects and composing them via relations (Battaglia et al., 2018; Lake et al., 2017; Tenenbaum et al., 2011; Chollet, 2019). Downstream of such representations is the ability to answer questions about unseen combinations of objects, e.g., a new Rube Goldberg machine with an added ramp. Such combinatorial structure arises in many domains. Robots must reason about varying types of objects and their spatial relations to navigate and manipulate the world (Li et al., 2019; Wang et al., 2025; Locatello et al., 2020); language is generated from unbounded combinations of nouns related by verbs (Chomsky, 1965); and biological systems are naturally described in terms of interacting proteins, metabolites, and cells (Barabási et al., 2011; Veličković, 2023; Regev et al., 2017). The generality of this problem has inspired active research into relational and object-centric machine learning (Koller & Friedman, 2009; Getoor & Taskar, 2007; Veličković et al., 2018; Kipf & Welling, 2017; Zambaldi et al., 2018) aimed at such combinatorial generalization.

The second problem is that of answering counterfactual questions: what if the weight were lighter, the string were cut, or the ramp angle were changed in our Rube Goldberg machine? A common view is that such questions cannot be answered from observations of the environment alone, requiring either interventions, or *causal* inductive biases, or often both (Pearl, 2009; Pearl & Mackenzie, 2018; Bareinboim et al., 2022; Schölkopf et al., 2021; Bareinboim, 2025). In our case, evidence for this point of view comes from the weaknesses of relational machine learning. Despite strong in-distribution predictive performance, relational and non-relational methods alike can still exploit correlations that are unstable under interventions or distribution shift (de Haan et al., 2019; Park et al., 2021; Fan et al., 2022; Wu et al., 2022; Vo et al., 2025). They do not guarantee answers to interventional or counterfactual questions.

Causal machine learning methods aim to answer such questions in the contexts of decision-making, generative modeling, fairness, and more (Schölkopf, 2022; Kaddour et al., 2025; Xia et al., 2021; Plečko & Bareinboim, 2024; Bareinboim et al., 2021; Pan & Bareinboim, 2024; 2025). However, many of these results do not easily lend themselves to combinatorial generalization because they rely on a fixed causal graph over a fixed set of variables. In domains such as autonomous driving, where traffic scenes differ in how many signals, pedestrians and cars appear and how they relate, causal methods must make assumptions that simplify these relations. They often assume that objects are unrelated (i.i.d.) or that the relational structure is fixed. Causal relational learning (Lee & Honavar, 2016; Ahsan et al., 2022b; Maier et al., 2010; Salimi et al., 2020) and methods for causal inference from non-i.i.d. data (Rubin, 1990; Sobel,

[1]Causal Artificial Intelligence Lab, Columbia University, New York, USA. Correspondence to: Adiba Ejaz <adiba.ejaz@cs.columbia.edu>.

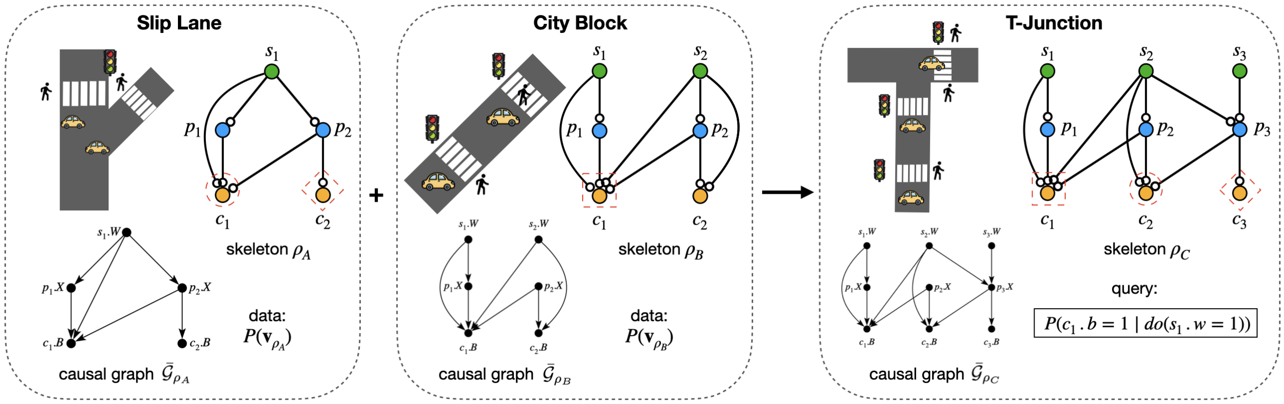

Figure 1. A schematic for the problem of relational identification across varying traffic scenes, following the schema in Ex. 1. Each panel shows: (i) a relational skeleton ρ representing a particular combination of signals (s), pedestrians (p), and cars (c), (ii) the corresponding causal graph $\bar{\mathcal{G}}_\rho$, and (iii) the available data or the query of interest. The goal is to use data from the *source* skeletons ρ_A and ρ_B to answer a query about the *target* skeleton ρ_B. In any ρ, an edge (s_i, p_j) indicates that signal s_i controls pedestrian p_j; an edge (s_i, c_j) indicates that signal s_i controls car c_j; and an edge (p_i, c_j) indicates that pedestrian p_i is in the path of car c_j. A bubble marks the second tuple element. Cars with similar relational neighborhoods (i.e., number of related pedestrians and signals) across the skeletons are circumscribed by similar shapes (dotted red outline; Table 1).

2006; Hudgens & Halloran, 2008; Ogburn & VanderWeele, 2014; Weinstein & Blei, 2024; Zhang et al., 2022b; Guo et al., 2024) address important pieces of the puzzle. However, they often preclude unobserved confounders, and do not address the problem of generalizing across combinations where both object counts and their relations can vary.

Contributions. We develop causal models for object-relational settings, enabling causal inference across varying combinations of objects. More specifically, our contributions are as follows.

1. **Relational SCMs.** In Sec. 3, we formalize how different combinations of objects can be unified by the same data-generating process: a relational structural causal model (Def. 3.1). Based on this formalization, we prove the limits of observational and causal inference about seen and unseen combinations of objects without further assumptions.
2. **Graphical identification.** In Sec. 4.1, we introduce a graphical model approach for the task of *relational identification* (Def. 4.2) across combinations of objects. We show when and how existing causal inference tools can be used for this task.
3. **Relational neural causal models.** We develop *relational neural causal models* (Def. 5.1) which form the basis of a sound and complete neural approach for relational identification in practice.

Experiments with simulated traffic scenes (Sec. 6, Sec. E) support our findings. We give an extended discussion of related literature in Sec. A, as well as proofs (including proof sketches) and further results in Sec. D.

2. Preliminaries

Notation. Capital letters (X) denote variables, dom(X) denotes their domains, small letters (x) denote values in their domains, and bold letters denote sets of variables ($\mathbf{X}$) and their values ($\mathbf{x}$). $P(\mathbf{X})$ denotes the probability distribution over a set of variables $\mathbf{X}$. We consistently use $P(\mathbf{x})$ to abbreviate probabilities $P(\mathbf{X} = \mathbf{x})$.

Structural causal models. Our framework extends that of *structural causal models* (SCMs) (Pearl, 2009; Bareinboim, 2025), a formalism for data-generating processes. An SCM $\mathcal{M}$ is a four-tuple $\mathcal{M} = \langle \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ where $\mathbf{V}$ and $\mathbf{U}$ are sets of endogenous (observed) and exogenous (unobserved) variables respectively. $\mathcal{F}$ is a set of mechanisms: each $V \in \mathbf{V}$ takes the value $f_V(\mathbf{pa}_V, \mathbf{u}_V)$, a function of the values of its endogenous and exogenous parents, $\mathbf{Pa}_V \subseteq \mathbf{V}$ and $\mathbf{U}_V \subseteq \mathbf{U}$, respectively. $P(\mathbf{U})$ is a joint distribution over $\mathbf{U}$; as in prior work (Zhang et al., 2022b; Xia et al., 2021), we assume the variables in $\mathbf{U}$ are jointly independent, although a given U may affect more than one V.

Every SCM induces a *causal graph*, constructed as follows: (1) add a vertex for every $V \in \mathbf{V}$ (2) add an edge $V_i \to V_j$ for every $V_i, V_j \in \mathbf{V}$ if $V_i \in \mathbf{Pa}_{\mathbf{V_j}}$ (3) add a dashed bidirected edge between V_i, V_j if $\mathbf{U}_{V_i} \cap \mathbf{U}_{V_j} \neq \emptyset$. See Sec. B for additional background.

Objects and relations. We build on the entity-relationship (ER) model (Ullman & Widom, 2002). A *relational schema* is a 3-tuple $\mathcal{S} = \langle \mathcal{E}, \mathcal{R}, \mathcal{A} \rangle$ where $\mathcal{E}$ is a set of entity (or object) types; $\mathcal{R}$ a set of relation types over $\mathcal{E}$; and $\mathcal{A}$ a set of observed attribute types $O.A$ for types $O \in \mathcal{E} \cup \mathcal{R}$. A *relational skeleton* ρ of $\mathcal{S}$ is a finite set of ground entities and relations o of the specified types $O \in \mathcal{E} \cup \mathcal{R}$. We write

$\rho(O)$ for the set of instances o in ρ of type O.

Example 1 (Relational schema and skeleton for traffic scene). A simple relational schema for traffic scenes would be

$$\mathcal{E} = \{\text{Signal (Sig)},\ \text{Car (Car)},\ \text{Pedestrian (Ped)}\}$$
$$\mathcal{R} = \{\text{Ctrl(Sig, Ped)}, \text{Ctrl(Sig, Car)}, \text{Path(Ped, Car)}\}$$
$$\mathcal{A} = \{\text{Sig}.W, \text{Ped}.X, \text{Car}.B\},$$

with all attributes binary-valued: $\text{Sig}.W \in \{1, 0\}$ denotes walk/drive; $\text{Ped}.X \in \{1, 0\}$ cross/wait; and $\text{Car}.B \in \{1, 0\}$ brake/go. Ctrl(Sig, Ped) (resp. Ctrl(Sig, Car)) indicates that a signal controls a pedestrian (resp. car), and Path(Ped, Car) indicates that the pedestrian is in the car's path. Figure 1 shows three skeletons for three traffic scenes, e.g., in ρ_A (slip lane), one signal s_1 controls pedestrians p_1, p_2 and car c_1 (but not c_2). □

3. Defining and Characterizing Relational SCMs

In this section, we introduce relational structural causal models (RSCMs) with the goal of specifying a data-generating process that underlies varying combinations of objects.

3.1. Defining Relational SCMs

An RSCM generalizes a standard SCM in two ways. First, different types of objects carry different attributes, and hence different sets of variables in an RSCM. Second, an attribute of one object may affect that of another only when the two objects stand in a particular relation. Following previous work on relational modeling (Koller & Friedman, 2009), we capture such contingent dependencies using *relational constraints* (Def. B.1).

Example 2 (Relational constraints for traffic scene). In Ex. 1, consider a signal Sig, pedestrian Ped, and car Car. In skeleton ρ_A, the constraint ϕ : Ctrl(Sig, Car) is true for Sig $= s_1$ and Car $= c_1$ but not Car $= c_2$. In ρ_B, the constraint ϕ' : Path(Ped, Car) is true for Ped $= p_1$ and Car $= c_1$ but not Car $= c_2$. □

An RSCM specifies one mechanism per attribute (e.g., for whether a car brakes). Its output can depend on the attributes of related objects (e.g., the crossing states of all pedestrians in the car's path), possibly via aggregation (Def. B.2).

Definition 3.1 (Relational structural causal model (RSCM)). A *relational structural causal model* (RSCM) is a 5-tuple $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$, where $\mathcal{S} = \langle \mathcal{E}, \mathcal{R}, \mathcal{A} \rangle$ is a relational schema; $\mathbf{V}$ is a set of endogenous variables $O.A$ for each attribute $O.A$ in $\mathcal{A}$; $\mathcal{F}$ is a set of mechanisms $f_{O.A}$ for each variable $O.A$ in $\mathbf{V}$; $\mathbf{U}$ is a set of exogenous variables $O.U$ tied to objects $O \in \mathcal{E} \cup \mathcal{R}$; and $P(\mathbf{U})$ is a probability distribution over $\mathbf{U}$ factorizing as $P(\mathbf{U}) = \prod_{O.U \in \mathbf{U}} P(O.U)$. Each mechanism $f_{O.A}$ has the form

$$O.A \leftarrow f_{O.A}(\mathbf{pa}_{O.A}, \mathbf{u}_{O.A}, \mathbf{pa}^r_{O.A}, \mathbf{u}^r_{O.A}).$$

Here, $\mathbf{Pa}_{O.A} \subseteq \mathbf{V}$ and $\mathbf{U}_{O.A} \subseteq \mathbf{U}$ are *non-relational parents* comprising attributes of the same object instance. On the other hand, $\mathbf{Pa}^r_{O.A}$ and $\mathbf{U}^r_{O.A}$ are *relational parents*. Each endogenous relational parent is a tuple $(\mathbf{W}, \phi, \text{AGG})$, where $\mathbf{W} \subseteq \mathbf{V}$ are variables belonging to some type $T \in \mathcal{E} \cup \mathcal{R}$; ϕ is a relational constraint over entities associated with O and T; and AGG is an optional list of aggregators for each $T.W \in \mathbf{W}$. Exogenous relational parents are analogous.

Example 3 (RSCM for traffic scene). Continuing Ex. 1 and 2, we define an RSCM for traffic scenes. The endogenous variables are $\mathbf{V} = \{\text{Sig}.W,\ \text{Ped}.X,\ \text{Car}.B\}$. The exogenous variables $\mathbf{U}$ are $\text{Sig}.U_W \sim \mathcal{B}(0.3)$, $\text{Ped}.U_X \sim \mathcal{B}(0.4)$ and $\text{Car}.U_B \sim \mathcal{B}(0.2)$, capturing unobserved factors such as a pedestrian's intent to cross or a driver's alertness. The mechanisms are

$$\text{Sig}.W \leftarrow \text{Sig}.U_W,$$
$$\text{Ped}.X \leftarrow \text{Ped}.U_X \oplus \bigwedge_{\text{Ctrl(Sig,Ped)}} \text{Sig}.W, \text{ and}$$
$$\text{Car}.B \leftarrow \text{Car}.U_B \oplus \Bigg(\bigvee_{\text{Ctrl(Sig,Car)}} \text{Sig}.W \vee \bigvee_{\text{Path(Ped,Car)}} \text{Ped}.X \Bigg).$$

For example, Ped.X has the non-relational parent Ped.U_X and relational parent $(\{\text{Sig}.W\}, \text{Ctrl(Sig, Ped)}, \wedge)$. Intuitively, $f_{\text{Ped}.X}$ makes a pedestrian cross when all controlling signals are in the 'walk' state, up to noise (e.g., the pedestrian does not intend to cross). Similarly, $f_{\text{Car}.B}$ makes a car brake when any controlling signal says 'walk' or any pedestrian in its path is crossing, again up to noise (e.g., the driver is not alert). See Ex. 9 for an extended example. □

Note how mechanisms in an RSCM differ from those in an SCM. Since the number of objects satisfying a constraint (e.g., pedestrians in a car's path) can vary across skeletons, each $f_{O.A}$ in an RSCM must accept multisets of varying size, while in an SCM, $f_{O.A}$ accepts a fixed-size input. In practice, relational learning often uses permutation-invariant *aggregators* (Def. B.2) such as mean, sum, max, majority, attention pooling, etc. to implement functions on sets.

An RSCM $\mathcal{M}$ may additionally be *Markovian*.

Definition 3.2 (RSCM Markovianity). We say an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ is *ρ-Markovian* if for each variable $O.A \in \mathbf{V}$, the set of exogenous relational parents $\mathbf{U}^r_{O.A}$ is empty. We say $\mathcal{M}$ is *Markovian* if it is ρ-Markovian and no two variables $O.A, T.B \in \mathbf{V}$ share a non-relational exogenous parent.

An RSCM can be instantiated for any skeleton ρ. It induces a standard *ground* RSCM (Def. B.4) with a ground variable $o.A$ for each attributes $O.A$ and instance $o \in \rho(O)$. The function determining $o.A$ substitutes the relational parents in $f_{O.A}$ with ground variables $t.W$ where o and t stand in the required relation.

Example 4 (Ground RSCM for traffic scene). For the RSCM $\mathcal{M}$ in Ex. 3 and skeleton ρ_A in Fig. 1, the ground RSCM $\mathcal{M}_{\rho_A} = \langle \mathbf{V}_{\rho_A}, \mathbf{U}_{\rho_A}, \mathcal{F}_{\rho_A}, P(\mathbf{U}_{\rho_A})\rangle$ is as follows.

$$\mathbf{V}_{\rho_A} = \{s_1.W,\ p_1.X,\ p_2.X,\ c_1.B,\ c_2.B\}$$
$$\mathbf{U}_{\rho_A} = \{s_1.U_W,\ p_1.U_X,\ p_2.U_X, c_1.U_B,\ c_2.U_B\}$$
$$s_1.W \leftarrow s_1.U_W$$
$$p_1.X \leftarrow p_1.U_X \oplus \bigwedge\{s_1.W\}$$
$$p_2.X \leftarrow p_2.U_X \oplus \bigwedge\{s_1.W\}$$
$$c_1.B \leftarrow c_1.U_B \oplus \Big(\bigvee\{s_1.W\} \vee \bigvee\{p_1.X, p_2.X\}\Big)$$
$$c_2.B \leftarrow c_2.U_B \oplus \Big(\bigvee\emptyset \vee \bigvee\{p_2.X\}\Big)$$

with $s_1.U_W \sim \mathcal{B}(0.3)$; $p_1.U_X,\ p_2.U_X \sim_{\text{iid}} \mathcal{B}(0.4)$; and $c_1.U_B,\ c_2.U_B \sim_{\text{iid}} \mathcal{B}(0.2)$. $\mathcal{M}_{\rho_A}$ describes the generative process for various traffic scenes with the structure ρ_A. □

We assume, throughout, that for any skeleton ρ, the ground RSCM $\mathcal{M}_\rho$ is recursive (or acyclic, Def. B.3).[1]

3.2. Limits of Learning Relational SCMs

In most domains, the true data-generating process, or RSCM, is unknown (Pearl, 2009; Bareinboim et al., 2022; Bareinboim, 2025). What we observe instead is data from many skeletons, each with its own combination of objects and relations. In this section, we consider what can be learned from such data about the true RSCM, and what this implies for unseen relational structures.

A ground RSCM $\mathcal{M}_\rho$ induces observational, interventional, and counterfactual distributions over $\mathbf{V}_\rho$ (Def. B.5).

Example 5 (RSCM distributions for traffic scene). Consider $\mathcal{M}_\rho$ in Ex. 4. The observational query $P(s_1.W = 1)$ is the probability that signal s_1 says 'walk'. The interventional query $P(c_1.B = 1 \mid do(p_1.X = 1))$ is the probability that car c_1 brakes when pedestrian p_1 crosses under intervention, irrespective of the signal (e.g., by an officer). The counterfactual

$$P\big(c_1.B_{s_1.W=1} = 1 \,\big|\, s_1.W = 0,\ p_1.X = 1,\ c_1.B = 0\big)$$

asks: in a scene where the signal was 'drive', p_1 crossed, and c_1 did not brake, what is the probability that c_1 *would have* braked had s_1 been set to 'walk'? □

The classic challenge in causal inference is inferring causal effects from observational data, assuming a fixed skeleton. However, when the skeleton varies, we show that even the *observational* distribution for an unseen skeleton is not learnable without further assumptions.

Theorem 3.3 (Impossibility of observational inference across skeletons)**.** *Consider a schema $\mathcal{S}$, source skeletons $\rho_1, \dots, \rho_l$, and target skeleton $\rho_\star$. Then, for any RSCM $\mathcal{M}$ over $\mathcal{S}$, there exists another RSCM $\mathcal{M}'$ over $\mathcal{S}$ such that $\mathcal{M}$ and $\mathcal{M}'$ agree on observational distributions $P(\mathbf{v}_{\rho_k})$ for every source skeleton ρ_k but disagree on the observational distribution $P(\mathbf{v}_{\rho_\star})$ of the target skeleton.*

Say we want to learn an interventional distribution for an unseen target skeleton $\rho_\star$. Thm. 3.3 already limits our ability to do this, since even the observational distribution $P(\mathbf{v}_{\rho_\star})$, a prerequisite of existing causal inference methods, may not be identified by the source data. An independently interesting question, however, is: when we *do* know some distributions for the target skeleton, e.g., $P(\mathbf{v}_{\rho_\star})$, does this suffice to identify other interventional (or counterfactual) queries in the target? We give a negative answer.[2]

Theorem 3.4 (Impossibility of causal inference within a skeleton)**.** *Consider a schema $\mathcal{S}$ where at least one entity or relation type has more than one observed attribute. For any relational SCM $\mathcal{M}$ over $\mathcal{S}$ and skeleton ρ, there exists another relational SCM $\mathcal{M}'$ over $\mathcal{S}$ such that $\mathcal{M}$ and $\mathcal{M}'$ agree on the observational distribution $P(\mathbf{v}_\rho)$ but disagree on some interventional distribution over $\mathbf{V}_\rho$.*

Thms. 3.3 and 3.4 hold even when the relational structure is known. This suggests the need for assumptions about the causal structure, in addition to the relational structure. We consider this in the next section.

4. Relational Identification

In the previous section, we showed that without further assumptions, even the observational distribution for unseen combinations of objects cannot be identified. In this section, we develop a graphical model approach to overcome this impossibility.

[1]This is weaker than requiring the template RSCM itself to be acyclic. For instance, an RSCM may specify that for cars C and C', $\mathsf{Car}.B$ affects $C'.B$ if C' is behind C. This appears cyclic at the template level; however, in any grounding $\mathcal{M}_\rho$, two cars cannot both be behind each other, and so $\mathcal{M}_\rho$ is acyclic. We implement such an RSCM in Exp. 6.2.

[2]It may seem that a negative answer follows immediately from the causal hierarchy theorem (CHT) (Bareinboim et al., 2022). However, the proof of the CHT relies on being able to construct an *arbitrary* SCM $\mathcal{M}'$ that matches the true SCM $\mathcal{M}$ on the given distribution(s) but not on the query. We are in a stricter setting where $\mathcal{M}'$ must share exogenous distributions and functions across objects of the same type, i.e., be a grounding of an RSCM.

4.1. Defining Relational Causal Graphs

First, we extend causal graphs to include relational constraints (Pearl, 2009; Koller & Friedman, 2009).

Definition 4.1 (Relational causal graph)**.** An RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ induces a *relational causal graph* $\mathcal{G}$ constructed as follows.

- **Non-relational subgraph.** For each object type $O \in \mathcal{E} \cup \mathcal{R}$, let $\mathcal{G}$ contain nodes for each variable $O.A \in \mathbf{V}$, a directed edge $O.B \to O.A$ for any $O.B \in \mathbf{Pa}_{O.A}$, and a dashed bidirected edge $O.A \leftrightarrow O.B$ for any $O.B \in \mathbf{V}$ such that $\mathbf{U}_{O.A} \cap \mathbf{U}_{O.B} \neq \emptyset$, annotated with the constraint $O = O'$.
- **Relational subgraph.** For each variable $O.A \in \mathbf{V}$ and each relational parent $R = (\mathbf{W}, \phi, \text{AGG}) \in \mathbf{Pa}^r_{O.A}$, let $\mathcal{G}$ contain a *relational node* $O.R$ and an edge $O.R \to O.A$. For each $T.W \in \mathbf{W}$, add an edge $T.W \to O.R$ annotated with ϕ and AGG. Finally, for any $T.B \in \mathbf{V}$ such that $O.A$ and $T.B$ have exogenous relational parents $(\mathbf{W}_1, \phi_1, \text{AGG}_1)$ and $(\mathbf{W}_2, \phi_2, \text{AGG}_2)$ respectively such that some $Z.U \in \mathbf{W}_1 \cap \mathbf{W}_2$, add a dashed bidirected edge $O.A \leftrightarrow T.B$ annotated with the constraint $\exists Z : \phi_1 \wedge \phi_2$, or append the constraint to an existing $O.A \leftrightarrow T.B$ edge.

Fig. 2 shows the graph for the traffic RSCM (Ex. 3). Like an RSCM, a relational causal graph $\mathcal{G}$ can be instantiated for any skeleton to yield a *ground graph* $\mathcal{G}_\rho$ (Def. B.6). The relational nodes in $\mathcal{G}_\rho$ can be *marginalized* (Def. B.7, Fig. 1) to yield a standard causal graph $\bar{\mathcal{G}}_\rho$.

Causal graphs can be used to encode assumptions about the space of possible RSCMs. They help circumvent the impossibility results of Sec. 3, enabling *relational identification*.

Definition 4.2 (Relational counterfactual identification)**.** Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeletons $\rho_1, \ldots, \rho_l$, source distributions $\mathbb{P} = \{\{P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))\}_{j=1}^{m_k}\}_{k=1}^{l}$, and target skeleton $\rho_\star$. Let $P(\mathbf{y}_* \mid \mathbf{x}_*)$ be a target query with $\mathbf{Y}_*, \mathbf{X}_* \subseteq \mathbf{V}_{\rho_\star}$.

We say $P(\mathbf{y}_* \mid \mathbf{x}_*)$ is *relationally identifiable* from $\mathcal{G}$ and $\mathbb{P}$ if for any RSCMs $\mathcal{M}, \mathcal{M}'$ consistent with $\mathcal{G}$ agreeing on the source data, so that for every ρ_k and $j = 1, \ldots, m_k$,

$$P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) > 0,$$

they also agree on the query:

$$P^{\mathcal{M}_{\rho_\star}}(\mathbf{y}_* \mid \mathbf{x}_*) = P^{\mathcal{M}'_{\rho_\star}}(\mathbf{y}_* \mid \mathbf{x}_*).$$

Otherwise, the query is *relationally non-identifiable*.

A special case of the above definition is when all skeletons (source and target) are *isomorphic* to each other (Def. B.1. This is the task of same-skeleton identification. When the target skeleton is non-isomorphic to every source skeleton, we refer to the task as cross-skeleton identification.

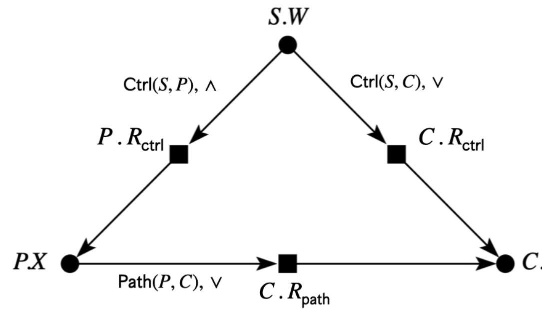

Figure 2. A relational causal graph (Def. 4.1) for the traffic RSCM in Ex. 3. Sig.W denotes the state of a signal, Ped.X whether a pedestrian crosses, and Car.B whether a car brakes. A signal affects a pedestrian or car only if it controls them (Ctrl), and a pedestrian affects a car only if they are in the car's path (Path). The relational nodes represent aggregated values (or, and) of the related objects.

Example 6 (Relational identification across skeletons). Continuing Ex. 1, suppose we are given the graph $\mathcal{G}$ in Fig. 2. Let the target skeleton be ρ_B (Fig. 1) and consider the query $P^{\rho_B}(c_1.B = 1 \mid do(s_1.W = 1))$, the effect of setting signal s_1 to 'walk' on whether car c_1 brakes. A non-relational identification task for this query would assume all available distributions are from the same skeleton ρ_B. On the other hand, relational identification can ask whether the same query is answerable from a distribution for a different skeleton, e.g., $P(\mathbf{v}_{\rho_A})$.

Crucially, the target query is not interchangeable with $P^{\rho_A}(c_1.B = 1 \mid do(s_1.W = 1))$: in ρ_A, s_1 controls two pedestrians in the path of c_1, whereas in ρ_B, s_1 controls only one pedestrian in the path of c_1. So, the contribution of $s_1.W$ need not be the same. The target quantity is also not interchangeable with $P^{\rho_B}(c_2.B = 1 \mid do(s_1.W = 1))$, since s_1 has no effect on c_2. A non-relational query such as $P(\text{Car}.B = 1 \mid do(\text{Sig}.W = 1))$, the usual target of causal inference, hides this variation. □

4.2. Identification Machinery

We now introduce tools for relational identification.

Observational inference. We first ask when observational distributions can be transferred across skeletons, addressing the limitation in Thm. 3.3. Say a ground variable $o.A$ is 'unconfounded' if there are no bidirected edges incident to it in the ground graph $\mathcal{G}_\rho$.

Theorem 4.3 (Observational identification across skeletons)**.** *Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeleton ρ, and target skeleton $\rho_\star$. Let $o.A$ be an* *unconfounded variable in* $\mathbf{V}_{\rho_\star}$. *The conditional* $P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A})$ *is relationally identifiable from* $\mathcal{G}$ *and* $P(\mathbf{v}_\rho)$ *if there exists a source instance* $o' \in \rho$ *such that* $o'.A$ *is unconfounded and* $\text{dom}(\mathbf{Pa}^r_{o.A}) \subseteq \text{dom}(\mathbf{Pa}^r_{o'.A})$. *In this case,* $P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}) = P(o'.a \mid \mathbf{pa}_{o'.A}, \mathbf{pa}^r_{o'.A})$.

If $\mathcal{G}$ is Markovian, (1) holds automatically, allowing us to recover the full target observational distribution $P(\mathbf{v}_{\rho_\star})$ if the support condition is satisfied for each instance (Cor. D.4).

Example 7 (Observational identification across skeletons). Continuing Ex. 1, let ρ_A and ρ_C be the source and target skeletons respectively, and $\mathcal{G}$ (Fig. 2) the relational causal graph. The conditional $P^{\rho_C}(c_2.b \mid s_2.w, p_2.x, p_3.x)$ in the target is identifiable from $P(\mathbf{v}_{\rho_A})$ and $\mathcal{G}$. This is because $c_1.B$ in $\mathbf{V}_{\rho_A}$ and $c_2.B$ in $\mathbf{V}_{\rho_C}$ are both unconfounded, and affected by exactly two pedestrians and one signal (and therefore have the same parent domains). In particular, writing multisets of values $\bar{w} = \{s_2.w\}$ and $\bar{x} = \{p_2.x, p_3.x\}$,

$$\begin{aligned} &P^{\rho_C}(c_2.b \mid s_2.w, p_2.x, p_3.x) \\ &= P^{\rho_C}(c_2.b \mid \{s_2.W\} = \bar{w}, \{p_2.X, p_3.X\} = \bar{x}) \\ &= P^{\rho_A}(c_1.b \mid \{s_1.W\} = \bar{w}, \{p_1.X, p_2.X\} = \bar{x}). \end{aligned}$$

Consider a more complex query $P^{\rho_C}(c_1.b \mid s_1.w, s_2.w, p_2.x, p_2.x)$ in the target. There is no car in the source controlled by two signals. If the parent domains are multisets of values, no car in the source meets the support condition for this query. However, note the aggregation constraint in $\mathcal{G}$: $\mathsf{Car}.B$ only depends on its controlling signals via the aggregate $\vee$. Then, letting $\bar{w} = \vee(s_1.w, s_2.w)$ and $\bar{x} = \vee(p_2.x, p_3.x)$, we have

$$\begin{aligned} &P^{\rho_C}(c_1.b \mid s_1.w, s_2.w, p_2.x, p_2.x) \\ &= P^{\rho_C}(c_1.b \mid \vee(s_1.W, s_2.W) = \bar{w}, \vee(p_1.X, p_2.X) = \bar{x}) \\ &= P^{\rho_A}(c_1.b \mid \vee(s_1.W) = \bar{w}, \vee(p_1.X, p_2.X) = \bar{x}), \end{aligned}$$

rendering the query identifiable. □

Causal inference. For the task of same-skeleton identification, we show how the *ctf-calculus* (Correa & Bareinboim, 2025), a recent generalization of do-calculus (Pearl, 2009), can be used to show identifiability.

Proposition 4.4 (Sufficient condition for same-skeleton relational identification). *Consider a schema* $\mathcal{S}$, *relational causal graph* $\mathcal{G}$, *skeleton* ρ, *and family of interventional distributions* $\mathbb{P}$ *over* $\mathbf{V}_\rho$. *If* $P(\mathbf{y}_* \mid \mathbf{x}_*)$ *is identifiable via ctf-calculus from the marginalized ground graph* $\bar{\mathcal{G}}_\rho$ *and* $\mathbb{P}$, *then it is also relationally identifiable from* $\mathcal{G}$ *and* $\mathbb{P}$.

As a corollary, we recover the well-known backdoor adjustment formula (Pearl, 2009) in the relational setting.

Corollary 4.5 (Relational backdoor adjustment). *Consider a schema* $\mathcal{S}$, *relational causal graph* $\mathcal{G}$, *skeleton* ρ, *and observational distribution* $P(\mathbf{v}_\rho)$. *Let* $P(\mathbf{y} \mid do(\mathbf{x}))$ *be some query with* $\mathbf{X}, \mathbf{Y} \subseteq \mathbf{V}_\rho$ *and let* $\mathbf{Z} \subseteq \mathbf{V}_\rho$ *be such that*

1. $\mathbf{Z}$ *contains no descendants of* $\mathbf{X}$ *in the marginalized ground graph* $\bar{\mathcal{G}}_\rho$, *and*
2. $\mathbf{Z}$ *blocks every path between* $\mathbf{X}$ *and* $\mathbf{Y}$ *that contains an arrow into* $\mathbf{X}$ *in* $\bar{\mathcal{G}}_\rho$.

Then, $P(\mathbf{y} \mid do(\mathbf{x}))$ *is relationally identifiable from* $\mathcal{G}$ *and* $P(\mathbf{v}_\rho)$ *as*

$$P(\mathbf{y} \mid do(\mathbf{x})) = \sum_{\mathbf{z}} P(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) P(\mathbf{z})$$

Example 8 (Same-skeleton relational identification). Continuing Ex. 1, let $\rho = \rho_A$ be the skeleton of interest, $\mathcal{G}$ (Fig. 2) the graph, and $P(c_1.B \mid do(p_1.X))$ the query. In the marginalized ground graph $\bar{\mathcal{G}}_{\rho_A}$ (Fig. 1), $\mathbf{Z} = \{s_1.W, p_2.X\}$ is a valid backdoor adjustment set for the query. This gives

$$\begin{aligned} &P(c_1.b \mid do(p_1.x)) = \\ &\sum_{p_2.x, s_1.w} P(c_1.b \mid p_1.x, p_2.x, s_1.2) \cdot P(p_2.x, s_1.w) \end{aligned}$$

Next, we turn to non-identifiability. Many practical queries are within-instance: they ask how intervening on an individual's treatment affects that same individual's outcome when individuals interact (e.g., showing an online ad to one user in a social network and measuring that user's purchases). We show that if a within-instance query is standardly non-identifiable, then it is also relationally non-identifiable.

Proposition 4.6 (Necessary condition for within-instance relational identification). *Consider a schema* $\mathcal{S}$, *relational causal graph* $\mathcal{G}$, *source skeletons* $\rho_1, \ldots, \rho_l$ *with available interventional distributions* $\mathbb{P}$, *and a target skeleton* $\rho_\star$. *Let* $o \in \rho_\star$ *be a target instance and consider a counterfactual query* $P(\mathbf{y}_\star \mid \mathbf{x}_\star)$ *with* $\mathbf{Y}_\star, \mathbf{X}_\star \subseteq \mathbf{V}_o$, *the attributes of* o.

Let the restriction $\mathbb{P}|_O$ *be as follows. For each source skeleton* ρ_k, *each distribution* $P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) \in \mathbb{P}$, *and each object* $o' \in \rho_k(O)$, *include* $P(\mathbf{v}_{\rho_k, o'} \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_{\rho_k, o'}))$ *in* $\mathbb{P}|_O$, *with instance identifiers omitted. Let* $\mathcal{G}_o$ *be the induced subgraph of the marginalized ground graph* $\bar{\mathcal{G}}_{\rho_\star}$ *on* $\mathbf{V}_o$ *with instance identifiers omitted.*

If $P(\mathbf{y}_\star \mid \mathbf{x}_\star)$ *is non-identifiable via ctf-calculus from* $\mathbb{P}|_O$ *and* $\mathcal{G}_o$, *then it is relationally non-identifiable from* $\mathcal{G}$ *and* $\mathbb{P}$.

See Ex. C.2 for an application of the above result.

This section developed symbolic criteria for relational identification across a range of settings. We next introduce a neural approach that generalizes these criteria and provides a practical route to identification.

5. Relational Neural Causal Models

Neural causal models (NCMs) (Xia et al., 2021; 2023) parameterize an SCM with neural networks. We adopt the same idea in the relational setting, developing a neural RSCM constrained by a given graph to enable relational identification from graph and data.

A neural RSCM contains one neural network for each variable $O.A$; this network is reused for all ground $o.A$. The given graph $\mathcal{G}$ constrains which inputs each network may use. In the traffic example, this means one network for Car.B is applied to every car's $c.B$, taking as input that car's non-relational parents (e.g., $c.U_B$) and relational parents (e.g., the counts of $s.W$ for controlling signals and Ped.X for in-path pedestrians).

Assumptions. We assume throughout this section that observed attributes are discrete and finite, that $\mathcal{G}$ is ρ-Markovian (no unobserved confounding between different instances), and that each relational-parent multiset has bounded size (see Sec. D.3).

Definition 5.1 ($\mathcal{G}$-Constrained Relational Neural Causal Model ($\mathcal{G}$-RNCM)). Consider a schema $\mathcal{S}$ and a relational causal graph $\mathcal{G}$. A $\mathcal{G}$-RNCM $\mathcal{N} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$ is an RSCM constructed as follows.

1. $\mathbf{V}$ contains a variable $O.A$ for every non-relational node $O.A$ in $\mathcal{G}$.
2. For every object type $O \in \mathcal{E} \cup \mathcal{R}$, and every maximal bidirected clique $\mathbf{C}$ in $\mathcal{G}$ over variables belonging to type O, $\mathbf{U}$ contains a variable $O.U_{\mathbf{C}} \sim \mathcal{U}([0,1])$.[3]
3. For each $O.A \in \mathbf{V}$, the mechanism $f_{O.A}$ is a feed-forward neural network
$$O.A \leftarrow f_{O.A}(\mathbf{pa}_{O.A}, \mathbf{u}_{O.A}, \mathbf{pa}^r_{O.A}).$$
Above, $\mathbf{Pa}_{O.A}$ consists of variables $O.B \in \mathbf{V}$ with a directed edge $O.B \rightarrow O.A$ in $\mathcal{G}$; $\mathbf{U}_{O.A}$ consists of variables $O.U_{\mathbf{C}}$ such that $O.A$ is in the clique $\mathbf{C}$; and $\mathbf{Pa}^r_{O.A}$ consists of multisets (or aggregates as specified in $\mathcal{G}$) of variables $\mathbf{W}$ such that there is an edge $O.R \rightarrow O.A$ in $\mathcal{G}$ where $O.R$ is a relational node with $R = (\mathbf{W}, \phi, \text{AGG})$.

Note that since $\mathcal{G}$ is ρ-Markovian, the set $\mathbf{U}^r_{O.A}$ is empty for each $O.A$ in the definition above. We show that $\mathcal{G}$-RNCMs are expressive enough to represent any RSCM consistent with $\mathcal{G}$, under our stated assumptions.

Theorem 5.2 (Expressivity of RNCMs). *Consider a relational schema $\mathcal{S}$. For every RSCM $\mathcal{M}$ over $\mathcal{S}$ inducing relational causal graph $\mathcal{G}$, there exists a $\mathcal{G}$-RNCM $\mathcal{N}$ such that for every skeleton ρ, the ground RSCMs $\mathcal{M}_\rho$ and $\mathcal{N}_\rho$ induce the same counterfactual distributions over $\mathbf{V}_\rho$.*

The expressivity of RNCMs lays the groundwork for causal identification via neural methods. In particular, we can now train the parameters of an RNCM to fit the source data while minimizing or maximizing the query on our target skeleton; this procedure is given in Alg. 1. As a corollary of Thm. 5.2, we can derive that Alg. 1 is sound and complete for the task of relational identification (Cor. D.6).

[3] See Sec. B.2 for the definition of a maximal bidirected clique.

Algorithm 1 RelationalNeuralID

Input: schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source data $\mathcal{D} = \left\{\left(\rho_k,\ \{P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))\}_{j=1}^{m_k}\right)\right\}_{k=1}^{l}$, target skeleton $\rho_\star$, query $P(\mathbf{y}_* \mid \mathbf{x}_*)$

$\hat{M} \leftarrow \mathcal{G}$-RNCM

$\hat{\theta}_l \leftarrow \arg\min_{\theta \in \Theta(\hat{M})} P^{\hat{M}_{\rho_\star}(\theta)}(\mathbf{y}_* \mid \mathbf{x}_*)$ subject to $\forall k, j$

$$P^{\hat{M}_{\rho_k}(\theta)}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))$$

$q_l \leftarrow P^{\hat{M}_{\rho_\star}(\hat{\theta}_l)}(\mathbf{y}_* \mid \mathbf{x}_*)$

$\hat{\theta}_r \leftarrow \arg\max_{\theta \in \Theta(\hat{M})} P^{\hat{M}_{\rho_\star}(\theta)}(\mathbf{y}_* \mid \mathbf{x}_*)$ subject to $\forall k, j$

$$P^{\hat{M}_{\rho_k}(\theta)}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))$$

$q_r \leftarrow P^{\hat{M}_{\rho_\star}(\hat{\theta}_r)}(\mathbf{y}_* \mid \mathbf{x}_*)$

if $q_l = q_r$ **then**
  return q_l
else
  return `FAIL`
end if

While Alg. 1 is stated as taking distributions for inputs, in practice, and in our implementation, each distribution is observed only through finitely many samples. The equality constraints are replaced by approximate constraints induced by a maximum-likelihood objective, adapting standard NCM training procedures to allow for parameter sharing across instances (Xia et al., 2021; 2023) (Appendix E).

6. Experiments

6.1. Estimation accuracy across traffic scenes

We evaluate how well RNCMs can estimate identifiable queries on both seen and unseen skeletons using Alg. 2, a modification of Alg. 1 that fits only one RNCM to match the available data. We use the traffic schema in Ex. 1 and graph $\mathcal{G}$ in Fig. 2, but omit the aggregators, creating a more challenging task.[4]

Setup. We train on four source settings from Fig. 1: three use a single skeleton for training (ρ_A, ρ_B or ρ_C) and one uses a pair of skeletons (ρ_A and ρ_C). For each source skeleton, we generate $n = 10^4$ observational samples and fit three models: (i) a $\mathcal{G}$-RNCM, (ii) NCM-J, a flat

[4] We use a histogram of counts of the different values in the domain of the relational parents. For discrete variables, this is a sufficient statistic for the multiset of relational parent values.

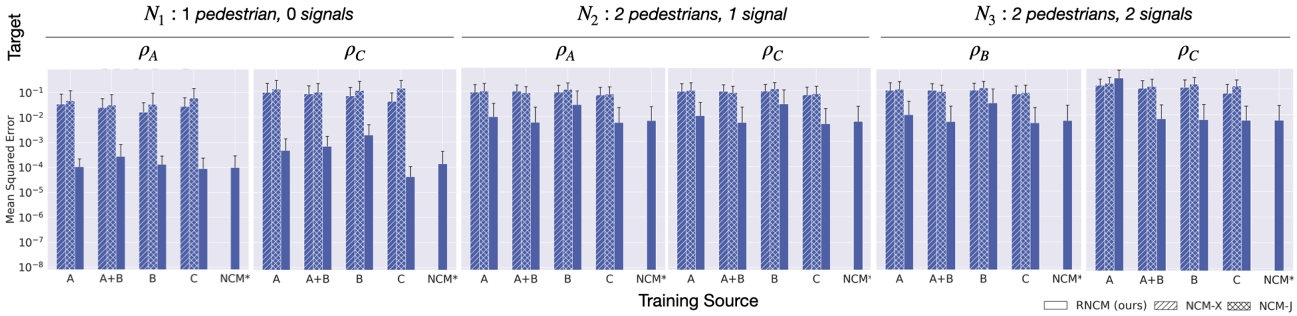

Figure 3. Estimation accuracy of RNCMs for causal effects across the traffic scenes in Fig. 1 (Exp. 6.1). The target query is the probability that a car c brakes when all signals controlling c are intervened to 'walk' and all pedestrians in its path are intervened to 'cross'. Each panel corresponds to a different target car c, i.e., a choice of target skeleton and car neighborhood (Table 1). The x-axis indicates the training source, and the y-axis reports $\log(\text{MSE})$ to ground truth, averaged over 10 seeds (**lower is better**). RNCMs (solid) consistently outperform non-relational NCM baselines (textured)—even when the baselines are trained directly on the target, while the RNCM is not. In identifiable cases, RNCMs often match the gold-standard NCM*, which is trained with target structure and target data. Performance degrades primarily in non-identifiable settings (e.g., training on ρ_A and evaluating on neighborhood N_3).

NCM trained on a relational join of triples $(s.W, p.X, c.B)$ where $\mathsf{Ctrl}(s,p)$, $\mathsf{Ctrl}(s,c)$, and $\mathsf{Path}(p,c)$ hold, and (iii) NCM-X, a flat NCM trained on the Cartesian product of all $(s.W, p.X, c.B)$ triples, ignoring relations. We also train a gold-standard NCM* directly on data from the target skeleton. Then, we evaluate each model on the question: what is the probability that car c will brake given that all the signals controlling c are set to 'walk' and all the pedestrians in its path are set to 'cross'? For three car neighbourhoods (Table 1) appearing in two target skeletons each, this results in six queries total; Fig. 3 reports mean squared error (MSE) to ground truth for each query, averaged over 10 random seeds.

Comparison with baselines. Across all identifiable source–target settings, the $\mathcal{G}$-RNCM (solid bars) outperforms the flat NCM baselines (textured bars), often by 10x or more. Notably, even when the $\mathcal{G}$-RNCM is trained on a source skeleton distinct from the target, it outperforms flat baselines trained directly on the target. For instance, an RNCM trained on source skeleton ρ_A outperforms NCM-X and NCM-J trained on ρ_C when estimating the query for neighborhood N_1 in target ρ_C. This highlights the importance of relational structure in determining causal effects. Moreover, in identifiable cases, RNCMs trained on sources distinct from the target are often strong enough to match the gold-standard NCM* trained directly on the target. The main exception is training on ρ_A and testing on N_3 in ρ_B and ρ_C, which we discuss below.

Neighborhoods	# Signals	# Pedestrians	Seen in
N_1	1	0	ρ_A, ρ_C
N_2	1	2	ρ_A, ρ_C
N_3	2	2	ρ_B, ρ_C

Table 1. Relational neighborhoods of cars in Fig. 1 and Exp. 6.1, with number of signals that control the car and pedestrians in the car's path.

Role of training source. When identifiability conditions hold, an RNCM trained on a source distinct from the target often matches the accuracy of an RNCM trained directly on the target. For example, for neighbourhood N_2 in target skeleton ρ_C, RNCMs trained on ρ_A and ρ_C achieve similar MSE, though ρ_A is a smaller, less dense skeleton, corroborating Thm. 4.3 and Prop. 4.4.

We find two exceptions to this pattern. First, training on ρ_A underperforms on N_3: in ρ_A, every car is controlled by at most one signal, so it fails the support condition (Thm. 4.3) for N_3. Accordingly, RNCMs trained on ρ_A perform markedly worse on N_3 than RNCMs trained on sources where this configuration occurs (e.g., ρ_B and ρ_C). Second, RNCMs trained on ρ_B perform slightly worse on N_2 than sources that contain the exact N_2 configuration. Although ρ_B satisfies the support condition, having matched neighborhoods in the source appears to reduce estimation error in finite samples. In both cases, combining ρ_A and ρ_B in training recovers performance, illustrating how RNCMs can integrate complementary sources to improve accuracy on an unseen target.

6.2. Identification accuracy across traffic scenes

We evaluate how well RNCMs are able to decide when a causal effect is identifiable, following Alg. 1.

Setup. We use a schema with one entity type Car (C) and one relation $\mathsf{Behind}(\mathsf{Car}_1, \mathsf{Car}_2)$. Each car has two observed attributes $\mathsf{Car}.X$ and $\mathsf{Car}.Y$. The source skeleton ρ has two cars c_1, c_2 and a single relation $\mathsf{Behind}(c_1, c_2)$. The target skeleton $\rho_\star$ has three cars c_1, c_2, c_3 with relations $\mathsf{Behind}(c_1, c_2)$, $\mathsf{Behind}(c_2, c_3)$, and $\mathsf{Behind}(c_1, c_3)$. We consider two causal graphs, the relational bow $\mathcal{G}_{\text{bow}}$ and the relational IV $\mathcal{G}_{\text{iv}}$ (Figs. 4, E.5.1). We follow a similar data-generation set-up as in Exp. 6.1, but using a majority aggregator.

For each graph, we consider two target queries on $\rho_\star$:

$$Q_d := P^{\rho_\star}(c_3.Y \mid do(c_2.X))$$
$$Q_s := P^{\rho_\star}(c_3.Y \mid do(c_3.X))$$

where Q_d is an effect across different cars and Q_s is an effect within the same car. To test identifiability, we run Alg. 1 and report the resulting max–min gap in Fig. 4.

Results. Though the target car c_3 in $\rho_\star$ has a larger relational neighborhood (two cars behind it) than any car in the source ρ, $\mathcal{G}$-RNCMs correctly assess identifiability in the target, matching the theory.

In the non-relational IV and bow graphs, the causal effect $P(y \mid do(x))$ is non-identifiable from purely observational data (Pearl, 2009). By Prop. 4.6, the same conclusion carries over to the within-car query Q_s. Consistent with this, RNCMs trained to minimize vs. maximize Q_s remain far apart (orange curve in Fig. 4). In contrast, for both graphs we show that the cross-car query Q_d on $\rho_\star$ *is* identifiable from source data $P(\mathbf{v}_\rho)$ (Prop. E.2). Accordingly, the RNCM max–min gap for Q_d collapses to (approximately) zero (solid blue curve in Fig. 4), indicating that Alg. 1 correctly certifies identifiability in this setting.

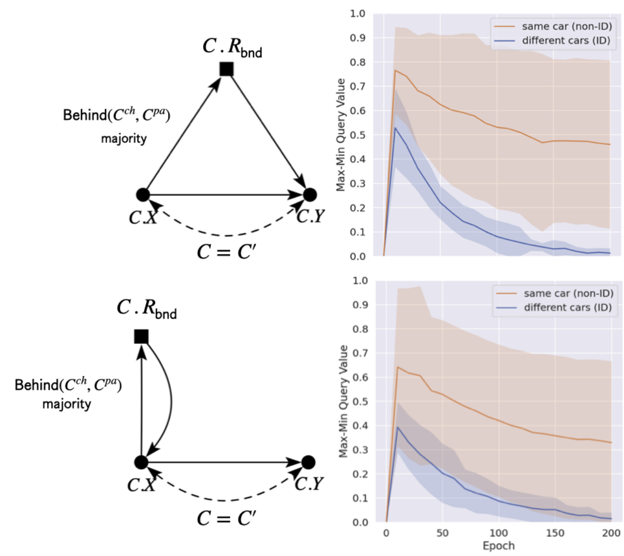

Figure 4. Identification accuracy of RNCMs trained and evaluated on distinct traffic scenes (Exp. 6.2). **(left)** Relational causal graphs $\mathcal{G}_{\text{bow}}$ (top) and $\mathcal{G}_{\text{iv}}$ (bottom). Within each car C, $\mathsf{Car}.X$ affects and is confounded with $\mathsf{Car}.Y$ in both graphs. Across cars, if C^{pa} is behind C^{ch}, then $C^{pa}.X$ affects $C^{ch}.Y$ in $\mathcal{G}_{\text{bow}}$ and $C^{ch}.X$ in $\mathcal{G}_{\text{iv}}$. **(right)** Average max–min gap produced by RelationalNeuralID (Alg. 1) versus training epochs. Shaded regions denote the 25th and 75th percentile across 10 random seeds. The gap collapses toward zero for the identifiable different-car query (blue) but remains large for the non-identifiable same-car query (orange).

7. Conclusions

In this paper, we introduced *relational structural causal models* (RSCMs; Def. 3.1), a generalization of SCMs to object-relational domains and characterized limits of learning in this setting (Thm. 3.3, Thm. 3.4). To enable relational identification (Def. 4.2) both within and across relational skeletons, we proposed *relational causal graphs* (Def. 4.1) and proved when and how such identification is possible (Thm. 4.3, Cor. D.4, Prop. 4.4, 4.6). Finally, we developed *relational neural causal models* (RNCMs; Def. 5.1, Alg. 1) for identification that are provably correct (Thm. 5.2, Cor. D.6) and empirically outperform existing neural-causal baselines (Sec. 6). We hope our work provides a foundation for causal reasoning in applications with both fixed interaction structures such as social networks as well as varying combinations of interacting objects such as autonomous driving and robotic manipulation.

Acknowledgements

This research is supported in part by the NSF, ONR, AFOSR, DoE, Amazon, JP Morgan, and The Alfred P. Sloan Foundation. We thank Kevin Xia and Yushu Pan for their helpful comments.

Impact Statement

This paper presents results that advance the field of machine learning, bringing together the areas of causal inference and relational learning. In particular, we contribute methodological foundations for causal inference in object-relational settings and validate our methods on simulated traffic scenes. If adopted in practice, our work could enable better generalization in settings such as autonomous driving and robotics.

References

Ahsan, R., Arbour, D., and Zheleva, E. Relational causal models with cycles: Representation and reasoning. In Schölkopf, B., Uhler, C., and Zhang, K. (eds.), *Proceedings of the First Conference on Causal Learning and Reasoning*, volume 177 of *Proceedings of Machine Learning Research*, pp. 1–18. PMLR, 11–13 Apr 2022a. URL `https://proceedings.mlr.press/v177/ahsan22a.html`.

Ahsan, R., Arbour, D., and Zheleva, E. Relational causal models with cycles: Representation and reasoning. In Schölkopf, B., Uhler, C., and Zhang, K. (eds.), *Proceedings of the First Conference on Causal Learning and Reasoning*, volume 177 of *Proceedings *of Machine Learning Research*, pp. 1–18. PMLR, 11–13 Apr 2022b. URL https://proceedings.mlr.press/v177/ahsan22a.html.

Arbour, D., Garant, D., and Jensen, D. Inferring network effects from observational data. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pp. 715–724, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939791. URL https://doi.org/10.1145/2939672.2939791.

Bach, S. H., Broecheler, M., Huang, B., and Getoor, L. Hinge-loss markov random fields and probabilistic soft logic. *J. Mach. Learn. Res.*, 18(1):3846–3912, January 2017. ISSN 1532-4435.

Baek, J., Wu, Y.-F., Singh, G., and Ahn, S. Dreamweaver: Learning compositional world models from pixels. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=e5mTvjXG9u.

Barabási, A.-L., Gulbahce, N., and Loscalzo, J. Network Medicine: A Network-based Approach to Human Disease. *Nature reviews. Genetics*, 12(1):56–68, January 2011. ISSN 1471-0056. doi: 10.1038/nrg2918. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC3140052/.

Bareinboim, E. *Causal Artificial Intelligence: A Roadmap for Building Causally Intelligent Systems*. 2025. URL https://causalai-book.net/. Draft version (October 11, 2025).

Bareinboim, E., Lee, S., and Zhang, J. An introduction to causal reinforcement learning. *arXiv preprint arXiv:2101.06498*, 2021.

Bareinboim, E., Correa, J. D., Ibeling, D., and Icard, T. On pearl's hierarchy and the foundations of causal inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pp. 507–556. Association for Computing Machinery, New York, NY, USA, 1st edition, 2022.

Battaglia, P., Hamrick, J. B. C., Bapst, V., Sanchez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G. E., Vaswani, A., Allen, K., Nash, C., Langston, V. J., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational inductive biases, deep learning, and graph networks. *arXiv*, 2018. URL https://arxiv.org/pdf/1806.01261.pdf.

Bhattacharya, R., Malinsky, D., and Shpitser, I. Causal inference under interference and network uncertainty. In Adams, R. P. and Gogate, V. (eds.), *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pp. 1028–1038. PMLR, 22–25 Jul 2020. URL https://proceedings.mlr.press/v115/bhattacharya20a.html.

Blei, D. M. Probabilistic topic models. 55, 2012.

Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/1cecc7a77928ca8133fa24680a88d2f9-Paper.pdf.

Bunescu, R. and Mooney, R. J. Collective information extraction with relational markov networks. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, ACL '04, pp. 438–es, USA, 2004. Association for Computational Linguistics. doi: 10.3115/1218955.1219011. URL https://doi.org/10.3115/1218955.1219011.

Buntine, W. L. Operations for learning with graphical models. *J. Artif. Int. Res.*, 2(1):159–225, December 1994. ISSN 1076-9757. doi: 10.1613/jair.62. URL https://doi.org/10.1613/jair.62.

Chen, T., Kanatsoulis, C., and Leskovec, J. RelGNN: Composite message passing for relational deep learning. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 8296–8312. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/chen25ad.html.

Chollet, F. On the measure of intelligence, 2019. URL https://arxiv.org/abs/1911.01547.

Chomsky, N. *Aspects of the Theory of Syntax*. The MIT Press, 50 edition, 1965. ISBN 9780262527408. URL http://www.jstor.org/stable/j.ctt17kk81z.

Codd, E. F. A relational model of data for large shared data banks. *Commun. ACM*, 13(6):377–387, June 1970. ISSN 0001-0782. doi: 10.1145/362384.362685. URL https://doi.org/10.1145/362384.362685.

Correa, J. D. and Bareinboim, E. Counterfactual graphical models: Constraints and inference. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=Z1qZoHa6ql.

Cotta, L., Bevilacqua, B., Ahmed, N., and Ribeiro, B. Causal lifting and link prediction. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 479(2276):20230121, 08 2023. ISSN 1364-5021. doi: 10.1098/rspa.2023.0121. URL https://doi.org/10.1098/rspa.2023.0121.

Cui, Z., Henrickson, K., Ke, R., Pu, Z., and Wang, Y. Traffic graph convolutional recurrent neural network: A deep learning framework for network-scale traffic learning and forecasting, 2019. URL https://arxiv.org/abs/1802.07007.

Cussens, J. Stochastic logic programs: sampling, inference and applications. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, UAI'00, pp. 115–122, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 1558607099.

Dai, H., Dai, B., and Song, L. Discriminative embeddings of latent variable models for structured data. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pp. 2702–2711. JMLR.org, 2016.

de Finetti, B. Funzione caratteristica di un fenomeno aleatorio. *Atti della R. Accademia Nazionale dei Lincei, Ser. 6. Memorie, Classe di Scienze Fisiche, Matematiche e Naturali 4*, pp. 251–299, 1931.

de Haan, P., Jayaraman, D., and Levine, S. *Causal confusion in imitation learning*. Curran Associates Inc., Red Hook, NY, USA, 2019.

De Raedt, L., Kimmig, A., and Toivonen, H. Problog: a probabilistic prolog and its application in link discovery. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, pp. 2468–2473, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

Derrow-Pinion, A., She, J., Wong, D., Lange, O., Hester, T., Perez, L., Nunkesser, M., Lee, S., Guo, X., Wiltshire, B., Battaglia, P. W., Gupta, V., Li, A., Xu, Z., Sanchez-Gonzalez, A., Li, Y., and Velickovic, P. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, CIKM '21, pp. 3767–3776, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384469. doi: 10.1145/3459637.3481916. URL https://doi.org/10.1145/3459637.3481916.

Du, Y., Durkan, C., Strudel, R., Tenenbaum, J. B., Dieleman, S., Fergus, R., Sohl-Dickstein, J., Doucet, A., and Grathwohl, W. Reduce, reuse, recycle: compositional generation with energy-based diffusion models and mcmc. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

Duan, X., He, Y., Tajwar, F., Chen, W., Salakhutdinov, R., and Schneider, J. State combinatorial generalization in decision making with conditional diffusion models, 2025. URL https://openreview.net/forum?id=PH7ja3T0vN.

Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, pp. 2224–2232, Cambridge, MA, USA, 2015. MIT Press.

Dwivedi, V. P., Jaladi, S., Shen, Y., López, F., Kanatsoulis, C. I., Puri, R., Fey, M., and Leskovec, J. Relational graph transformer, 2026. URL https://arxiv.org/abs/2505.10960.

Falcon, W. and Cho, K. A framework for contrastive self-supervised learning and designing a new approach, 2020. URL https://arxiv.org/abs/2009.00104.

Fan, S., Wang, X., Mo, Y., Shi, C., and Tang, J. Debiasing graph neural networks via learning disentangled causal substructure. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

Feng, F., Lippe, P., and Magliacane, S. Learning interactive world model for object-centric reinforcement learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=E0cjqfM55C.

Ferraro, S., Mazzaglia, P., Verbelen, T., and Dhoedt, B. FOCUS: Object-centric world models for robotic manipulation. In *Intrinsically-Motivated and Open-Ended Learning Workshop @NeurIPS2023*, 2023. URL https://openreview.net/forum?id=RoQbZRv1zw.

Fey, M., Hu, W., Huang, K., Lenssen, J. E., Ranjan, R., Robinson, J., Ying, R., You, J., and Leskovec, J. Position: Relational deep learning - graph representation learning on relational databases. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 13592–13607. PMLR, 21–27 Jul

2024. URL https://proceedings.mlr.press/v235/fey24a.html.

Friedman, N., Getoor, L., Koller, D., and Pfeffer, A. Learning probabilistic relational models. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'99, pp. 1300–1307, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.

Gao, H., Pei, J., and Huang, H. Conditional random field enhanced graph convolutional neural networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pp. 276–284, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362016. doi: 10.1145/3292500.3330888. URL https://doi.org/10.1145/3292500.3330888.

Gelman, A. and Hill, J. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press, 2006.

Getoor, L. and Taskar, B. *Introduction to Statistical Relational Learning*. The MIT Press, 08 2007. ISBN 9780262256230. doi: 10.7551/mitpress/7432.001.0001. URL https://doi.org/10.7551/mitpress/7432.001.0001.

Getoor, L., Friedman, N., Koller, D., and Taskar, B. Learning probabilistic models of link structure. *J. Mach. Learn. Res.*, 3(null):679–707, March 2003. ISSN 1532-4435.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pp. 1263–1272. JMLR.org, 2017.

Gori, M., Monfardini, G., and Scarselli, F. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pp. 729–734 vol. 2, 2005. doi: 10.1109/IJCNN.2005.1555942.

Guo, S., Zhang, C., Mohan, K., Huszár, F., and Schölkopf, B. Do finetti: On causal effects for exchangeable data. In Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 127317–127345. Curran Associates, Inc., 2024. doi: 10.52202/079017-4044. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/e656933e4e9626f5ada42349bcadffb8-Paper-Conference.pdf.

Gurnee, W. and Tegmark, M. Language models represent space and time. In Kim, B., Yue, Y., Chaudhuri, S., Fragkiadaki, K., Khan, M., and Sun, Y. (eds.), *International Conference on Representation Learning*, volume 2024, pp. 2483–2503, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/0a6059857ae5c82ea9726ee9282a7145-Paper-Conference.pdf.

Ha, D. and Schmidhuber, J. Recurrent world models facilitate policy evolution. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/2de5d16682c3c35007e4e92982f1a2ba-Paper.pdf.

Hajij, M., Zamzmi, G., Papamarkou, T., Miolane, N., Guzmán-Sáenz, A., Ramamurthy, K. N., Birdal, T., Dey, T. K., Mukherjee, S., Samaga, S. N., Livesay, N., Walters, R., Rosen, P., and Schaub, M. T. Topological deep learning: Going beyond graph data, 2023. URL https://arxiv.org/abs/2206.00606.

Hamaguchi, T., Oiwa, H., Shimbo, M., and Matsumoto, Y. Knowledge transfer for out-of-knowledge-base entities: a graph neural network approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, pp. 1802–1808. AAAI Press, 2017. ISBN 9780999241103.

Hamilton, W. L., Ying, R., and Leskovec, J. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 1025–1035, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

Hamrick, J. B., Allen, K. R., Bapst, V., Zhu, T., McKee, K. R., Tenenbaum, J. B., and Battaglia, P. W. Relational inductive bias for physical construction in humans and machines, 2018. URL https://arxiv.org/abs/1806.01203.

Heckerman, D., Meek, C., and Koller, D. Probabilistic models for relational data. Technical Report MSR-TR-2004-30, March 2004. URL https://www.microsoft.com/en-us/research/publication/probabilistic-models-for-relational-data/.

Hudgens, M. G. and Halloran, M. E. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008. doi: 10.1198/

016214508000000292. URL https://doi.org/10.1198/016214508000000292. PMID: 19081744.

Hwang, G., Choi, J., Cho, H., and Kang, M. MAGANet: Achieving combinatorial generalization by modeling a group action. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 14237–14248. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/hwang23b.html.

Jensen, D., Burroni, J., and Rattigan, M. Object conditioning for causal inference. In Adams, R. P. and Gogate, V. (eds.), *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pp. 1072–1082. PMLR, 22–25 Jul 2020. URL https://proceedings.mlr.press/v115/jensen20a.html.

Jeong, H., Ejaz, A., Tian, J., and Bareinboim, E. Testing causal models with hidden variables in polynomial delay via conditional independencies. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(25): 26813–26822, April 2025. ISSN 2159-5399. doi: 10.1609/aaai.v39i25.34885. URL http://dx.doi.org/10.1609/aaai.v39i25.34885.

Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C. L., and Girshick, R. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning, 2016. URL https://arxiv.org/abs/1612.06890.

Jothimurugan, K., Hsu, S., Bastani, O., and Alur, R. Robust subtask learning for compositional generalization. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

Kaddour, J., Lynch, A., Liu, Q., Kusner, M. J., and Silva, R. Causal machine learning: A survey and open problems. *Found. Trends Optim.*, 9(1–2):1–247, August 2025. ISSN 2167-3888. doi: 10.1561/2400000052. URL https://doi.org/10.1561/2400000052.

Kersting, K. and Raedt, L. D. Bayesian logic programs, 2001. URL https://arxiv.org/abs/cs/0111058.

Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=SJU4ayYgl.

Koller, D. and Friedman, N. *Probabilistic Graphical Models: Principles and Techniques*. Adaptive computation and machine learning. MIT Press, 2009. ISBN 9780262013192. URL https://books.google.co.in/books?id=7dzpHCHzNQ4C.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253, 2017. doi: 10.1017/S0140525X16001837.

LeCun, Y. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022.

Lee, J., Lee, Y., Kim, J., Kosiorek, A., Choi, S., and Teh, Y. W. Set transformer: A framework for attention-based permutation-invariant neural networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3744–3753. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/lee19d.html.

Lee, S. and Honavar, V. On learning causal models from relational data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), Mar. 2016. doi: 10.1609/aaai.v30i1.10417. URL https://ojs.aaai.org/index.php/AAAI/article/view/10417.

Li, R., Jabri, A., Darrell, T., and Agrawal, P. Towards practical multi-object manipulation using relational reinforcement learning. In *arXiv preprint arXiv:XXXX*, 2019.

Li, Y., Yu, R., Shahabi, C., and Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SJiHXGWAZ.

Liang, Q., Qian, D., Ziyin, L., and Fiete, I. R. Compositional generalization via forced rendering of disentangled latents. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=rkHCHI5H5W.

Lin, Y., Wang, A. S., Undersander, E., and Rai, A. Efficient and interpretable robot manipulation with graph neural networks. *IEEE Robotics and Automation Letters*, 7(2): 2740–2747, 2022. doi: 10.1109/LRA.2022.3143518.

Liu, N., Li, S., Du, Y., Torralba, A., and Tenenbaum, J. B. Compositional visual generation with composable diffusion models. In Avidan, S., Brostow, G., Cissé, M., Farinella, G. M., and Hassner, T. (eds.), *Computer Vision – ECCV 2022*, pp. 423–439, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-19790-1.

Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., and Kipf, T. Object-centric learning with slot attention. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 11525–11538. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/8511df98c02ab60aea1b2356c013bc0f-Paper.pdf.

Maier, M., Taylor, B., Oktay, H., and Jensen, D. Learning causal models of relational domains. *Proceedings of the AAAI Conference on Artificial Intelligence*, 24(1):531–538, Jul. 2010. doi: 10.1609/aaai.v24i1.7695. URL https://ojs.aaai.org/index.php/AAAI/article/view/7695.

Mendez, J. A., van Seijen, H., and EATON, E. Modular lifelong reinforcement learning via neural composition. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=5XmLzdslFNN.

Montero, M. L., Ludwig, C. J., Costa, R. P., Malhotra, G., and Bowers, J. The role of disentanglement in generalisation. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=qbH974jKUVy.

Mosbach, M., Ewertz, J. N., Villar-Corrales, A., and Behnke, S. SOLD: Slot object-centric latent dynamics models for relational manipulation learning from pixels. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=XOUpHJPYRX.

Nakano, A., Suzuki, M., and Matsuo, Y. Interaction-based disentanglement of entities for object-centric world models. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=JQc2VowqCzz.

Ogburn, E. L. and VanderWeele, T. J. Causal Diagrams for Interference. *Statistical Science*, 29(4):559 – 578, 2014. doi: 10.1214/14-STS501. URL https://doi.org/10.1214/14-STS501.

Okawa, M., Lubana, E. S., Dick, R. P., and Tanaka, H. Compositional abilities emerge multiplicatively: Exploring diffusion models on a synthetic task. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=frVo9MzRuU.

Orbanz, P. and Roy, D. M. Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37 (2):437–461, February 2015. ISSN 0162-8828. doi: 10.1109/TPAMI.2014.2334607. URL https://doi.org/10.1109/TPAMI.2014.2334607.

Oñoro, D., Niepert, M., García-Durán, A., González-Sánchez, R., and López-Sastre, R. Representation learning for visual-relational knowledge graphs. 09 2017. doi: 10.48550/arXiv.1709.02314.

Pan, Y. and Bareinboim, E. Counterfactual image editing. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 39087–39101. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/pan24a.html.

Pan, Y. and Bareinboim, E. Counterfactual image editing with disentangled causal latent space. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=u2Lgi4NIe7.

Papamarkou, T., Birdal, T., Bronstein, M. M., Carlsson, G. E., Curry, J., Gao, Y., Hajij, M., Kwitt, R., Lio, P., Di Lorenzo, P., Maroulas, V., Miolane, N., Nasrin, F., Natesan Ramamurthy, K., Rieck, B., Scardapane, S., Schaub, M. T., Veličković, P., Wang, B., Wang, Y., Wei, G., and Zamzmi, G. Position: Topological deep learning is the new frontier for relational learning. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 39529–39555. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/papamarkou24a.html.

Papillon, M., Sanborn, S., Mathe, J., Cornelis, L., Bertics, A., Buracas, D., J Lillemark, H., Shewmake, C., Dinc, F., Pennec, X., and Miolane, N. Beyond euclid: an illustrated guide to modern machine learning with geometric, topological, and algebraic structures. *Machine Learning: Science and Technology*, 6(3):031002, aug 2025. doi: 10.1088/2632-2153/adf375. URL https://doi.org/10.1088/2632-2153/adf375.

Park, J., Seo, Y., Liu, C., Zhao, L., Qin, T., Shin, J., and Liu, T.-Y. Object-aware regularization for addressing causal confusion in imitation learning. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc. ISBN 9781713845393.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library, 2019. URL https://arxiv.org/abs/1912.01703.

Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2nd edition, 2009.

Pearl, J. and Mackenzie, D. *The Book of Why*. Basic Books, New York, 2018.

Plečko, D. and Bareinboim, E. Causal fairness analysis: A causal toolkit for fair machine learning. *Found. Trends Mach. Learn.*, 17(3):304–589, January 2024. ISSN 1935-8237. doi: 10.1561/2200000106. URL https://doi.org/10.1561/2200000106.

Qu, M., Bengio, Y., and Tang, J. GMNN: Graph Markov neural networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5241–5250. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/qu19a.html.

Ranjan, R., Hudovernik, V., Znidar, M., Kanatsoulis, C., Upendra, R., Mohammadi, M., Meyer, J., Palczewski, T., Guestrin, C., and Leskovec, J. Relational transformer: Toward zero-shot foundation models for relational data, 2025. URL https://arxiv.org/abs/2510.06377.

Raposo, D., Santoro, A., Barrett, D., Pascanu, R., Lillicrap, T., and Battaglia, P. Discovering objects and their relations from entangled scene representations, 2017. URL https://arxiv.org/abs/1702.05068.

Regev, A., Teichmann, S. A., Lander, E. S., Amit, I., Benoist, C., Birney, E., Bodenmiller, B., Campbell, P., Carninci, P., Clatworthy, M., Clevers, H., Deplancke, B., Dunham, I., Eberwine, J., Eils, R., Enard, W., Farmer, A., Fugger, L., Göttgens, B., Hacohen, N., Haniffa, M., Hemberg, M., Kim, S., Klenerman, P., Kriegstein, A., Lein, E., Linnarsson, S., Lundberg, E., Lundeberg, J., Majumder, P., Marioni, J. C., Merad, M., Mhlanga, M., Nawijn, M., Netea, M., Nolan, G., Pe'er, D., Phillipakis, A., Ponting, C. P., Quake, S., Reik, W., Rozenblatt-Rosen, O., Sanes, J., Satija, R., Schumacher, T. N., Shalek, A., Shapiro, E., Sharma, P., Shin, J. W., Stegle, O., Stratton, M., Stubbington, M. J. T., Theis, F. J., Uhlen, M., van Oudenaarden, A., Wagner, A., Watt, F., Weissman, J., Wold, B., Xavier, R., Yosef, N., and Human Cell Atlas Meeting Participants. The Human Cell Atlas. *eLife*, 6:e27041, December 2017. ISSN 2050-084X. doi: 10.7554/eLife.27041.

Richardson, M. and Domingos, P. Markov logic networks. *Mach. Learn.*, 62(1–2):107–136, February 2006. ISSN 0885-6125. doi: 10.1007/s10994-006-5833-1. URL https://doi.org/10.1007/s10994-006-5833-1.

Richens, J. and Everitt, T. Robust agents learn causal world models. In Kim, B., Yue, Y., Chaudhuri, S., Fragkiadaki, K., Khan, M., and Sun, Y. (eds.), *International Conference on Representation Learning*, volume 2024, pp. 15786–15817, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/44a2b9f7bf9aec3f1fa333ad875b0ee0-Paper-Conference.pdf.

Robinson, J., Ranjan, R., Hu, W., Huang, K., Han, J., Dobles, A., Fey, M., Lenssen, J. E., Yuan, Y., Zhang, Z., He, X., and Leskovec, J. Relbench: A benchmark for deep learning on relational databases. In Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 21330–21341. Curran Associates, Inc., 2024. doi: 10.52202/079017-0672. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/25cd345233c65fac1fec0ce61d0f7836-Paper-Datasets_and_Benchmarks_Track.pdf.

Rosenbaum, P. R. Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477):191–200, 2007. doi: 10.1198/016214506000001112. URL https://doi.org/10.1198/016214506000001112.

Rubin, D. B. [On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9.] Comment: Neyman (1923) and Causal Inference in Experiments and Observational Studies. *Statistical Science*, 5(4):472 – 480, 1990. doi: 10.1214/ss/1177012032. URL https://doi.org/10.1214/ss/1177012032.

Sakai, Y., Kamigaito, H., and Watanabe, T. Revisiting compositional generalization capability of large language models considering instruction following ability. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 31219–31238, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1508. URL https://aclanthology.org/2025.acl-long.1508/.

Salimi, B., Parikh, H., Kayali, M., Getoor, L., Roy, S., and Suciu, D. Causal relational learning. In *Pro-*

ceedings of the 2020 ACM SIGMOD International Conference on Management of Data, SIGMOD '20, pp. 241–256, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367356. doi: 10.1145/3318464.3389759. URL https://doi.org/10.1145/3318464.3389759.

Santoro, A., Raposo, D., Barrett, D. G., Malinowski, M., Pascanu, R., Battaglia, P., and Lillicrap, T. A simple neural network module for relational reasoning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/e6acf4b0f69f6f6e60e9a815938aa1ff-Paper.pdf.

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *Trans. Neur. Netw.*, 20(1):61–80, January 2009a. ISSN 1045-9227. doi: 10.1109/TNN.2008.2005605. URL https://doi.org/10.1109/TNN.2008.2005605.

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. Computational capabilities of graph neural networks. *Trans. Neur. Netw.*, 20(1):81–102, January 2009b. ISSN 1045-9227. doi: 10.1109/TNN.2008.2005141. URL https://doi.org/10.1109/TNN.2008.2005141.

Schölkopf, B. *Causality for Machine Learning*, pp. 765–804. Association for Computing Machinery, New York, NY, USA, 1 edition, 2022. ISBN 9781450395861. URL https://doi.org/10.1145/3501714.3501755.

Schott, L., von Kügelgen, J., Träuble, F., Gehler, P., Russell, C., Bethge, M., Schölkopf, B., Locatello, F., and Brendel, W. Visual representation learning does not generalize strongly within the same domain, 2022. URL https://arxiv.org/abs/2107.08221.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. doi: 10.1109/JPROC.2021.3058954.

Sherman, E. and Shpitser, I. Identification and estimation of causal effects from dependent data. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/024677efb8e4aee2eaeef17b54695bbe-Paper.pdf.

Singla, P. and Domingos, P. Entity resolution with markov logic. In *Proceedings of the Sixth International Conference on Data Mining*, ICDM '06, pp. 572–582, USA, 2006. IEEE Computer Society. ISBN 0769527019. doi: 10.1109/ICDM.2006.65. URL https://doi.org/10.1109/ICDM.2006.65.

Sobel, M. E. What do randomized studies of housing mobility demonstrate? *Journal of the American Statistical Association*, 101(476):1398–1407, 2006. doi: 10.1198/016214506000000636. URL https://doi.org/10.1198/016214506000000636.

Sommestad, T., Ekstedt, M., and Johnson, P. A probabilistic relational model for security risk analysis. *Computers Security*, 29(6):659–679, 2010. ISSN 0167-4048. doi: https://doi.org/10.1016/j.cose.2010.02.002. URL https://www.sciencedirect.com/science/article/pii/S0167404810000209.

Song, Y., Lee, D., and Kim, G. Compositional conservatism: A transductive approach in offline reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=HRkyLbBRHI.

Spalević, S., Veličković, P., Kovačević, J., and Nikolić, M. Hierarchical protein function prediction with tail-gnns, 2020. URL https://arxiv.org/abs/2007.12804.

Spiegelhalter, D. J. Bayesian graphical modelling: A case-study in monitoring health outcomes. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 47(1):115–133, 01 2002. ISSN 0035-9254. doi: 10.1111/1467-9876.00101. URL https://doi.org/10.1111/1467-9876.00101.

Spirtes, P., Glymour, C. N., and Scheines, R. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition, 2000.

Sutton, C. and McCallum, A. An introduction to conditional random fields for relational learning. In *Introduction to Statistical Relational Learning*. The MIT Press, 08 2007. ISBN 9780262256230. doi: 10.7551/mitpress/7432.003.0006. URL https://doi.org/10.7551/mitpress/7432.003.0006.

Sypetkowski, M., Wenkel, F., Poursafaei, F., Dickson, N., Suri, K., Fradkin, P., and Beaini, D. On the scalability of GNNs for molecular graphs. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=klqhrq7fvB.

Taskar, B., Abbeel, P., and Koller, D. Discriminative probabilistic models for relational data. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, UAI'02, pp. 485–492, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc. ISBN 1558608974.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331 (6022):1279–1285, 2011. doi: 10.1126/science.1192788. URL https://www.science.org/doi/abs/10.1126/science.1192788.

Ullman, J. and Widom, J. *A First Course in Database Systems*. An Alan R. Apt book. Prentice Hall, 2002. ISBN 9780131225206. URL https://books.google.com/books?id=B5hOQAAACAAJ.

Vafa, K., Chen, J. Y., Rambachan, A., Mullainathan, S., and Kleinberg, J. Evaluating the world model implicit in a generative model. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385.

Veerapaneni, R., Co-Reyes, J. D., Chang, M., Janner, M., Finn, C., Wu, J., Tenenbaum, J., and Levine, S. Entity abstraction in visual model-based reinforcement learning. In Kaelbling, L. P., Kragic, D., and Sugiura, K. (eds.), *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pp. 1439–1456. PMLR, 30 Oct–01 Nov 2020. URL https://proceedings.mlr.press/v100/veerapaneni20a.html.

Veličković, P. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology*, 79:102538, April 2023. ISSN 0959-440X. doi: 10.1016/j.sbi.2023.102538. URL https://www.sciencedirect.com/science/article/pii/S0959440X2300012X.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJXMpikCZ.

Vo, A., Nguyen, K.-N., Taesiri, M. R., Dang, V. T., Nguyen, A. T., and Kim, D. Vision language models are biased, 2025. URL https://arxiv.org/abs/2505.23941.

Wang, Z., Wang, K., Zhao, L., Stone, P., and Bian, J. Dyn-o: Building structured world models with object-centric representations. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=b2u1yrTwFK.

Weinstein, E. N. and Blei, D. M. Hierarchical causal models, 2024. URL https://arxiv.org/abs/2401.05330.

Wu, Y.-X., Wang, X., Zhang, A., He, X., and seng Chua, T. Discovering invariant rationales for graph neural networks. In *ICLR*, 2022.

Wu, Z., Dvornik, N., Greff, K., Kipf, T., and Garg, A. Slotformer: Unsupervised visual dynamics simulation with object-centric models. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=TFbwV6I0VLg.

Xia, K., Lee, K.-Z., Bengio, Y., and Bareinboim, E. The causal-neural connection: expressiveness, learnability, and inference. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc. ISBN 9781713845393.

Xia, K. M., Pan, Y., and Bareinboim, E. Neural causal models for counterfactual identification and estimation. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=vouQcZS8KfW.

Xu, Z., Tresp, V., Yu, K., Yu, S., and Kriegel, H.-P. Dirichlet enhanced relational learning. In *Proceedings of the 22nd International Conference on Machine Learning*, ICML '05, pp. 1004–1011, New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1595931805. doi: 10.1145/1102351.1102478. URL https://doi.org/10.1145/1102351.1102478.

Yang, H., Lu, H., Lam, W., and Cai, D. Exploring compositional generalization of large language models. In Cao, Y. T., Papadimitriou, I., Ovalle, A., Zampieri, M., Ferraro, F., and Swayamdipta, S. (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pp. 16–24, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-srw.3. URL https://aclanthology.org/2024.naacl-srw.3/.

Zaheer, M., Kottur, S., Ravanbhakhsh, S., Póczos, B., Salakhutdinov, R., and Smola, A. J. Deep sets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 3394–3404, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E., Shanahan, M., Langston, V., Pascanu, R., Botvinick, M., Vinyals, O., and Battaglia, P. Relational deep reinforcement learning, 2018. URL https://arxiv.org/abs/1806.01830.

Zečević, M., Dhami, D. S., Veličković, P., and Kersting, K. Relating graph neural networks to structural causal models, 2021. URL https://arxiv.org/abs/2109.04173.

Zhang, C., Mohan, K., and Pearl, J. Causal inference with non-iid data using linear graphical models. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 13214–13225. Curran Associates, Inc., 2022a. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/5573c63e8a89e32086e5c71cf0cc8fe4-Paper-Conference.pdf.

Zhang, C., Mohan, K., and Pearl, J. Causal inference with non-iid data using linear graphical models. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 13214–13225. Curran Associates, Inc., 2022b. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/5573c63e8a89e32086e5c71cf0cc8fe4-Paper-Conference.pdf.

Zhang, Y., Chen, X., Yang, Y., Ramamurthy, A., Li, B., Qi, Y., and Song, L. Efficient probabilistic logic reasoning with graph neural networks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=rJg76kStwH.

Zhao, C., Tan, Z., Ma, P., Li, D., Jiang, B., Wang, Y., Yang, Y., and huan liu. Is chain-of-thought reasoning of LLMs a mirage? a data distribution lens. In *First Workshop on Foundations of Reasoning in Language Models*, 2025. URL https://openreview.net/forum?id=o2AoLPIjle.

Zhao, L., Kong, L., Walters, R., and Wong, L. L. Toward compositional generalization in object-oriented world modeling. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 26841–26864. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/zhao22b.html.

Zhou, D., Schärli, N., Hou, L., Wei, J., Scales, N., Wang, X., Schuurmans, D., Cui, C., Bousquet, O., Le, Q., and Chi, E. Least-to-most prompting enables complex reasoning in large language models, 2023. URL https://arxiv.org/abs/2205.10625.

Appendices

A. Background and Related Works

Much of machine learning is 'Euclidean' (Papillon et al., 2025). It operates on data represented as coordinates in a high-dimensional space. Such representations are also described as 'flat', 'variable-based', 'feature-vector', 'attribute-value,' and 'propositional' (Getoor & Taskar, 2007; Koller & Friedman, 2009). Statistical learning guarantees assume that the data are values of independent and identically distributed (i.i.d.) random variables. In this section, we discuss in which settings it is useful to relax this assumption, and present methods for such settings, causal and otherwise.

A.1. Relational Data

A.1.1. What is Relational Data?

There is no single definition of what counts as relational data. Broadly construed, it is data that is not flat. We describe, below, what various literatures mean by the term 'relational.'

In database theory, it refers to data organized under the relational model in a relational database (Codd, 1970; Ullman & Widom, 2002). A vast amount of enterprise, government, and academic data is stored in relational databases.

In statistical relational learning, 'relational' concerns domains modeled not as a fixed set of random variables, but rather structured spaces consisting of many types of objects (entities) related in different ways (Getoor et al., 2003; Koller & Friedman, 2009). Both entities and relations may have attributes. For example, to model inheritance across family trees, each tree consists of its own set of individuals, each with their own attributes (Koller & Friedman, 2009). These individuals have varying relations: Mother-of, Father-of, Married-to, Sibling-of. Both the number of individuals and their relations vary across family trees. Since individuals are related, inferences about one individual may be made based on observations of another, e.g., a child's risk of a disease based on their parents'.

In deep learning, 'relational' refers not only to a type of input data but also to a type of learning problem that involves reasoning about objects in various relations with one another. For example, while images can be encoded in a fixed Euclidean space, the task of answering questions like 'are there any rubber things that have the same size as the yellow metallic cylinder?' for a given image is considered relational (Santoro et al., 2017); intermediate representations of the data for this task might be structured as objects and relations. Relational reasoning involves composing together different objects (entities) via relations following certain rules (Battaglia et al., 2018).

A.1.2. How is Relational Data Represented?

We give some standard representations in increasing order of generality.

Sequences and sets (without explicit relations) At the least structured end, an instance is represented as a finite set of entities (objects) with attributes. Their relations are not specified explicitly, covering cases where relations are absent, unobserved, or deliberately abstracted away. A common modeling assumption is that entity indices are arbitrary, so the distribution (and the model) should be invariant to permutations of entities. For such exchangeable sequences, de Finetti's theorem implies they can be viewed as i.i.d. draws conditional on a latent variable (de Finetti, 1931; Orbanz & Roy, 2015). For sets, modern architectures (Zaheer et al., 2017; Lee et al., 2019; Locatello et al., 2020) enforce permutation invariance (or equivariance) directly while constructing representations. Examples include object-centric representations of scenes (a set of detected objects), multi-agent systems (a set of agents with state vectors), and point clouds (a set of points in space). A limitation of set-based representations is they do not take relational structure as an explicit input. However, they may be used to learn the relations between objects, e.g., by attention between all pairs.

Matrices and tensors (special cases of graphs). Matrices arise when there is a single relation between two types of entities, one indexed by rows and the other by columns, e.g., user–item interactions in recommender system. This can be viewed as a *bipartite graph*. The matrix entries are edge attributes (ratings, clicks, links), and missing entries correspond to absent/unobserved edges. Tensors extend this to multi-way relations (e.g., user $\times$ item $\times$ context) or to multiple relation types. The gain relative to sets is that some interaction structure is now explicit (who is connected to whom), but the limitation is that the representation presupposes a small number of entity types and a small number of relations that can be aligned to axes; many relational domains do not naturally factor into a single rectangular array.

Relational databases (tables). A relational database contains one or more tables ('relations') with attributes as columns and records as rows. Each attribute has a data type, and each record specifies a value of that type. The primary key is an attribute that uniquely identifies each row, whereas the foreign keys are attributes that connect a given row to other rows in other tables. Relational databases can be encoded as heterogenous graphs (Fey et al., 2024), which we explain next.

Graphs and heterogeneous graphs (explicit interaction structure). A graph input consists of explicit entities (nodes) and relations (edges), possibly with attributes on both and possibly with multiple node/edge types. This representation is natural for domains such as social networks (people connected by friendship/follow), transportation networks (intersections connected by road segments), molecular graphs (atoms connected by bonds), or citation/knowledge graphs (papers or entities connected by typed relations). It is often used in relational neural methods that compute by passing information along edges (Gilmer et al., 2017; Battaglia et al., 2018). Compared to set representations, graphs commit to a notion of locality (neighbors) and hence constrain which entities can directly influence which others; compared to matrices, graphs permit arbitrary sparsity patterns, multiple relation types, and heterogeneous entities without forcing them into a single array.

More general non-Euclidean spaces. Finally, some domains require modeling more than just edges or pairwise relations between objects. Higher-order relations (hyper-edges) and interactions between edges, triangles, and cliques are often useful in domains spanning physical systems, traffic forecasting, drug discovery, and more (Papillon et al., 2025; Papamarkou et al., 2024; Hajij et al., 2023). Structures such as hypergraphs, simplicial complexes, sheaves, and combinatorial complexes are used to represent such data in the active field of topological deep learning.

A.2. Probabilistic Models for Relational Data

Much of the probabilistic modeling literature in the 20th century operated on flat, i.i.d. data. For instance, Bayesian networks and related graphical models (Pearl, 2009; Koller & Friedman, 2009)captured dependencies among attributes of a fixed-dimensional random vector. An important exception was the tradition of hierarchical modeling—including multi-level and cross-classified models—for non-i.i.d. data (Gelman & Hill, 2006; Orbanz & Roy, 2015). However, these models typically involve two or fewer entity and relation types and simple, homogeneous relational neighborhoods (i.e., sequence- or matrix-structured data).

A new literature emerged in the 1990s, often grouped under *statistical relational learning* (SRL), that shifted from attribute-value representations of data to object-relational ones (Koller & Friedman, 2009; Getoor & Taskar, 2007). The goal of these methods was to enable probabilistic inference in object-relational settings; for e.g., infer a patient's risk of a disease based on that of their family members', or a user's rating for a movie based on that of their friends. To do so, these methods define a template-level model for a particular schema, laying out probabilistic independencies that could be grounded for any relational instances.

Plate models. Plate models represent objects as plates, and relations as intersections between these plates (Buntine, 1994; Spiegelhalter, 2002). Their primary use is in modeling domains with repeated measurements, encoding parameter sharing for objects in the same plate. For example, in a recommender system, one may have one plate for users, and another for movies; the intersection of these plates may define a variable 'rating'. This variable exists for every (user, movie) pair. Plate models are widely used in applied Bayesian statistics (**?**Blei, 2012), but do not express dependencies that are contingent on relational constraints, and assume simple, pre-defined relational neighbourhoods.[5] For example, a plate model can express that a user's rating for a movie depends on their preferences and the movie's quality. However, they can not express that this rating additionally depends on the user's friends' ratings for that movie.

Directed relational models. More general directed relational graphical models such as probabilistic relational models (PRMs, also known as relational Bayesian networks) (Friedman et al., 1999; Getoor et al., 2003; Koller & Friedman, 2009) and the directed acyclic probabilistic entity-relationship model (DAPER) (Heckerman et al., 2004; Getoor & Taskar, 2007) address the limitation of plate models by defining dependencies using first-order constraints and allowing variable-sized neighbourhoods. Application areas include citation networks and web hyperlinks (including link prediction and topic classification) (**?**), medical diagnosis (Xu et al., 2005), and IT security risk analysis (Sommestad et al., 2010). Nevertheless, like Bayesian networks, they specify factorizations of observational distributions but do not provide a causal semantics for

[5] While this is true of plate models as typically used in the literature, (Heckerman et al., 2004) define a generalized plate model that can express constraints.

interventions and counterfactuals.

Other approaches. In our work, we build on the directed graphical model approach for object-relational domains. However, we note there are a number of other prominent approaches to SRL. Undirected models address the limitation of directed models in capturing cyclic or symmetric dependencies (e.g., feedback loops, dynamical systems, dependencies such as "these individuals are likely to enjoy the same movies"). They include relational Markov networks (Taskar et al., 2002), Markov logic networks (Richardson & Domingos, 2006), and conditional random fields (Sutton & McCallum, 2007). Probabilistic inference tends to be more difficult in undirected models than directed models; still, undirected models have proved useful for entity resolution in citation networks (e.g., is author o the same as author o'?) (Singla & Domingos, 2006); for classifying webpages (Taskar et al., 2002); and for extracting knowledge from unstructured text (Bunescu & Mooney, 2004). In addition, given that causation is directed and asymmetric, such undirected models do not suffice to represent causality in relational settings (Maier et al., 2010). Besides graphical models, probabilistic logic programs constitute another (and in some cases, equivalent) approach, and have been applied to link prediction in heterogenous biological networks (Kersting & Raedt, 2001; Cussens, 2000; Bach et al., 2017; De Raedt et al., 2007).

A.3. Relational Deep Learning

Relational deep learning combines and extends the tradition of graphical models with deep neural networks for scalable inference. Unlike SRL, it is not necessarily probabilistic; like SRL, it operates on non-Euclidean data, most commonly graphs and databases. Its goal is to design architectures with 'relational inductive biases' that leverage relational information for prediction and promote *combinatorial generalization*, understood as generalization across combinations of objects (Battaglia et al., 2018).

Graph neural networks. Graph neural networks (GNNs) are perhaps the most popular architecture for relational deep learning (Gori et al., 2005; Scarselli et al., 2009a;b; Hamilton et al., 2017; Kipf & Welling, 2017; Veličković et al., 2018; Veličković, 2023). Given a graph with node (entity) features and edge (relation) features, message passing architectures iteratively update each node representation by aggregating messages from its neighbors. GNNs have been applied to molecular property prediction (Duvenaud et al., 2015; Gilmer et al., 2017; Sypetkowski et al., 2024), physical reasoning and interaction networks in vision and control (Raposo et al., 2017; Zambaldi et al., 2018; Hamrick et al., 2018), knowledge graphs (Bordes et al., 2013; Oñoro et al., 2017; Hamaguchi et al., 2017), and spatiotemporal forecasting on transportation networks (Li et al., 2018; Cui et al., 2019; Derrow-Pinion et al., 2021). GNNs also provide a computational interface between SRL and modern representation learning, since many SRL domains can be compiled into heterogeneous graphs. A number of works have integrated undirected SRL methods into GNN architectures (Dai et al., 2016; Gao et al., 2019; Spalević et al., 2020; Qu et al., 2019; Zhang et al., 2020).

Causality and GNNs. An important distinction to make is that the graphs used as input to GNNs are not causal graphs. They are relational graphs, with nodes depicting entities and edges depicting relations between then. Causal graphs, on the other hand, are naturally formulated over features (of nodes or edges). As such GNNs, like the previously discussed SRL approaches, lack the architecture and guaranties for predicting the effects of interventions and counterfactuals, as distinct from observations. For instance, changing a node feature (a movie's log line) or adding an edge (recommending a movie to a user) are not modeled as do-interventions. (Cotta et al., 2023) take an important step towards causal prediction using graph embeddings for the task of link prediction. (Zečević et al., 2021) propose GNNs for causal inference; their method, restricted to settings without unobserved confounding, is equivalent to neural causal models (Xia et al., 2021) under this restriction. However, it applies only to i.i.d. attribute-value data, thus leaving open the relational setting.

Relational deep learning on databases. A distinct and increasingly active direction studies deep learning directly on relational databases, with the goal of avoiding manual, error-prone feature engineering to flatten relational data (Fey et al., 2024; Robinson et al., 2024). These works model typically model databases as heterogenous graphs to leverage the power of graph neural networks. They fall into roughly two categories: models are trained on and applicable to a fixed schema (wu:; Chen et al., 2025; Dwivedi et al., 2026) versus 'foundation models' that are trained on and applicable to diverse schemas (Wang et al., 2025; Ranjan et al., 2025).

A.4. Combinatorial and Compositional Generalization

The focus of our work echoes the problem of *combinatorial generalization* studied in artificial intelligence. There is no agreed-upon definition of combinatorial generalization. Sometimes, it is used interchangeably with the term *compositional generalization* (e.g., in (Liu et al., 2022)). However, compositional generalization may also involve the composition of functions or tasks (e.g., subtask reinforcement learning (Mendez et al., 2022; Jothimurugan et al., 2023) and compositional instruction following or skill composition in LLMs (Yang et al., 2024; Zhao et al., 2025; Sakai et al., 2025; Zhou et al., 2023)).

Combining features vs combining objects. We distinguish between two types of combinatorial generalization. First, *feature-combination generalization* holds the underlying object-relational structure fixed but varies the combinations of features of these objects. A canonical example is attribute binding: training on blue circles and red squares and testing on blue squares. Second, *object-combination generalization* varies the underlying set of objects and relations itself (often including the number of objects), e.g., in going from 2 stacked blocks to 3 stacked blocks, or from small interaction graphs to larger ones. In this work, we study the latter type of combinatorial generalization.

Computer vision. In computer vision, combinatorial generalization is often studied for visual scenes with varying objects, attributes, and relations (Okawa et al., 2023; Hwang et al., 2023). CLEVR is a paradigmatic dataset for combinatorial generalization in visual-question answering (Johnson et al., 2016). Combinatorial generalization is also a goal of image generation. (Liu et al., 2022; Du et al., 2023) develop an approach that composes pre-trained diffusion models, explicitly enforce compositionality during inference using logical operators, instead of relying on implicit learning. This approach, alongside several other works (Schott et al., 2022; Montero et al., 2021; Liang et al., 2025), challenges a common view that disentangled representation learning is sufficient for combinatorial generalization.

Decision-making and world models. For tasks such as robotic manipulation and autonomous driving, combinatorial generalization is unavoidable since scenes naturally vary in the number of objects and their relations (Cui et al., 2019; Derrow-Pinion et al., 2021; Lin et al., 2022). (Zambaldi et al., 2018) introduce an approach for relational deep reinforcement learning that uses attention over entity representations, showing improved generalization to more complex instances than those seen during training. (Duan et al., 2025) introduce a formal definition of 'out-of-combination' generalization in the decision-making context that assumes a fixed number of objects and requires generalization to regions out of the *support* of the state space training distribution. While their diffusion-based approach shows promising zero-shot generalization for this task, both the task definition and their approach assume that all objects (or 'base elements') are seen during training (Duan et al., 2025, Sec 3.3), thus ruling out varying numbers of objects. (Song et al., 2024) solve a similar problem as (Duan et al., 2025) using an approach that maps unseen states to the closest state seen during training. An active recent literature on object-oriented (as opposed to monolithic) *world models* aims to learn object-centric representations of pixel data (Nakano et al., 2023; Wu et al., 2023; Baek et al., 2025; Ferraro et al., 2023; Veerapaneni et al., 2020; Wang et al., 2025; Mosbach et al., 2025; Feng et al., 2025; Zhao et al., 2022), for instance, by decomposing the latent space into 'slots' for different objects and sometimes explicitly modeling relations between objects. While these methods are evaluated on out-of-distribution generalization tasks, their performance on unseen combinations of objects remains relatively understudied.

A.5. Relational Causal Models

The intersection of causal and relational modeling is relatively understudied. In the graphical framework of causality, most works focus on relational causal discovery–the task of learning a relational causal graph from data. Works on relational causal inference–answering causal queries from graph and data–are few in number, and deal with special types of relations, as we describe below.

Relational causal discovery. (Maier et al., 2010) provide the first algorithm for causal discovery over data stored in relational databases. They use the DAPER model to encode conditional independencies in relational data, and extend the PC algorithm (Spirtes et al., 2000) to the relational setting. (Lee & Honavar, 2016) build on this DAPER framework, coining the term 'relational causal model' for the DAPER model, and providing more efficient and informative algorithms for relational causal discovery; (Ahsan et al., 2022a) extend this to include cyclic dependencies. However, the DAPER model is a probabilistic model, offering a compact encoding of conditional independencies in relational data (as Bayesian networks do for flat data). It lacks a causal semantics for interventions and counterfactuals, just as Bayesian networks do. This is precisely what motivated the formulation of structural causal models (SCMs) and causal/counterfactual Bayesian networks

(Pearl, 2009; Bareinboim et al., 2022; Bareinboim, 2025; Correa & Bareinboim, 2025). Therefore, these works leave open the grounding of causality in relational domains, and, therefore, the definition and inference of causal queries. Additionally, all of them address only those conditional independence structures that can be represented using graphs with directed edges, leaving out settings with unobserved confounding (often represented via bidirected edges) (Bareinboim et al., 2022; Jeong et al., 2025).

Causal inference under interference. Causal inference under interference can be seen as a special case of relational causal inference where all objects and relations are of the same type (Sobel, 2006; Rosenbaum, 2007; Ogburn & VanderWeele, 2014; Bhattacharya et al., 2020; Hudgens & Halloran, 2008; Zhang et al., 2022a; Sherman & Shpitser, 2018). The interference literature assumes a fixed number of 'units' (or objects) and a fixed interaction structure between them: this captures settings such as people in a given neighborhood, students in a given school, or patients in a given hospital. The query of interest is typically an aggregated causal effect across all units (e.g., an average direct effect, an average spillover effect, or a global average treatment affect). Unlike DAPER, however, these models do not necessarily enforce parameter-sharing across units, e.g., (Zhang et al., 2022a) study interference for linear models without enforcing mechanism while allowing different coefficients for different neighbors (violating permutation-invariance). Additionally, works in graphical causality under interference assume the absence of unobserved confounding, an assumption violated in many real-world settings.

Relational causal inference. Relational causal inference generalizes inference under interference, allowing heterogeneous objects and relations (Arbour et al., 2016; Jensen et al., 2020; Salimi et al., 2020; Weinstein & Blei, 2024). (Jensen et al., 2020) and (Weinstein & Blei, 2024) provide sound methods for causal inference in plate models, showing how 'object-conditioning' can lead to greater identifiability even in the presence of unobserved confounding. However, these models capture only a subset of the restricted types of relations (not including first-order constraints) expressible in plate models. (Arbour et al., 2016) are the first to generalize causal inference to the entity-relationship model, giving a backdoor criterion for identifying interventional queries from abstract ground graphs (Maier et al., 2010). They assume a ground relational network for a particular relational skeleton as input. As such, they do not define a template-level causal model which can be instantiated on different skeletons, tying them together. (Salimi et al., 2020) define a template-level formalism with 'causal rules' that define first-order constraints for when one variable affects another, and give a similar backdoor adjustment criterion. However, these causal rules do not specify the *mechanism* by which this effect unfolds, and thus resemble the coarse-grained information encoded in a causal graph.

Firstly, neither (Arbour et al., 2016) nor (Salimi et al., 2020) provide a mechanism-level definition of relational causal models, defining interventions only using the truncated Markov factorization (Pearl, 2009). As such, their set-up is limited to modeling interventions in settings without unobserved confounding, and does not provide semantics or identification criteria for counterfactuals. Secondly, both works assume that for causal identification in a given skeleton, observational data from that skeleton is available; they do not address the problem of cross-skeleton inference. Finally, while (Arbour et al., 2016) allows the effect of one variable on another to depend on the exact relation satisfied, (Salimi et al., 2020) does not; for example, if both friends' and family members' vaccination statuses affect a given individual's infection risk, they are assumed to affect it in the same way. While possibly beneficial for estimation in practice, this assumption is quite restrictive in real-world settings.

B. Further Definitions

Graphical terminology. Consider a graph $\mathcal{G} = (\mathbf{V}_\mathcal{G}, \mathbf{E}_\mathcal{G})$ with directed and bidirected edges. For a node O in $\mathcal{G}$, the *graphical parents* of O are the set of nodes $\{Y : Y \to X \in \mathbf{E}_\mathcal{G}\}$. The descendants of O are nodes Y such that there is a directed path $X \rightsquigarrow Y$ in $\mathcal{G}$. For a set of nodes $\mathbf{X} \subseteq \mathbf{V}$, the *induced subgraph* $\mathcal{G}_\mathbf{X}$ is the graph containing nodes $\mathbf{X}$ with an edge V, W between $V, W \in \mathbf{X}$ if and only if this edge is present in $\mathcal{G}$. A *bidirected clique* in $\mathcal{G}$ is a set of nodes every pair of which is connected by a bidirected edge in $\mathcal{G}$. Such a clique is *maximal* if there is no larger bidirected clique in which it is strictly contained.

B.1. Relational Structural Causal Models

We use a modified definition of relational constraints from (Heckerman et al., 2004, Def. 6).

Definition B.1 (Relational constraint)**.** Consider a relational schema $\mathcal{S} = \langle \mathcal{E}, \mathcal{R}, \mathcal{A} \rangle$. Let $\mathcal{X} \subseteq \mathcal{E}$ be a set of entity types. A relational constraint $\phi(\mathcal{X})$ is a first-order expression whose atoms are relationship symbols in $\mathcal{R}$ (alongside pre-defined constraints such as equality) and whose only free variables range over the space of ground entities of the types in $\mathcal{X}$.

RSCMs are defined using permutation-invariant functions on multiset inputs (Zaheer et al., 2017). Often, in relational learning, such mutlisets are summarized using *aggregators* (Koller & Friedman, 2009; Getoor & Taskar, 2007). We define such aggregators below.

Definition B.2 (Aggregator)**.** Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets, and let $\mathcal{X}^{\text{multiset}}$ be the set of all finite multisets over $\mathcal{X}$. An *aggregator* is a function

$$\text{AGG} : \mathcal{X}^{\text{multiset}} \to \mathcal{Y}$$

that maps a finite multiset of elements from $\mathcal{X}$ to an element of $\mathcal{Y}$.

Note that since aggregators take (multi-)sets as input, they do not depend on any ordering over the elements of their input.

We also define what it means for two relational skeletons to be considered the 'same'.

Definition B.1 (Skeleton isomorphism)**.** An isomorphism between skeletons ρ and ρ' over a given schema $\mathcal{S}$ is a bijection $\pi : \rho \to \rho'$ on entities and relations that (a) preserves types and (b) preserves relations between entities. We write $\rho \cong \rho'$ if such a π exists.

In Prop. D.1, we show how RSCM-induced distributions are invariant to skeleton isomorphism.

Next, we define what it means for an RSCM to be acyclic, or recursive.

Definition B.3 (Recursive RSCM)**.** An RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ is said to be *recursive* (or *acyclic*) if its relational causal graph $\mathcal{G}$ contains no directed cycles. Equivalently, there exists a topological ordering of the variables $\mathbf{V}$ such that for every attribute $O.A \in \mathbf{V}$, the structural function $f_{O.A}$ only takes as input (relational or non-relational) variables that precede $O.A$ in this ordering.

For a given skeleton ρ, an RSCM $\mathcal{M}$ induces a *ground RSCM*, a standard SCM with functions and exogenous distributions shared across variables of the same type.

Definition B.4 (Ground relational structural causal model)**.** Given an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ and a relational skeleton ρ following the schema $\mathcal{S}$, a ground relational structural causal model of $\mathcal{M}$ for ρ is an SCM $\mathcal{M}_\rho = \langle \mathbf{V}_\rho, \mathbf{U}_\rho, \mathcal{F}_\rho, P(\mathbf{U}_\rho) \rangle$ with

1. ground endogenous variables $\mathbf{V}_\rho = \{o.A \mid O.A \in \mathbf{V}, o \in \rho(O),$
2. ground exogenous variables $\mathbf{U}_\rho = \{o.U \mid O.U \in \mathbf{U}, o \in \rho(O)\}$ with distributions $o.U \sim P(O.U)$ given by $P(\mathbf{U})$, and
3. ground mechanisms $\mathcal{F}_\rho$, with $f_{o.A}$ obtained by substituting relational parents $(\mathbf{W}, \phi, \text{AGG})$ in $f_{O.A}$ with (the specified aggregate of) the set of ground variables $\{t.W \mid T.W \in \mathbf{W}, t \in \rho(T), \phi(o,t)$ holds in $\rho\}$. This yields a function $f_{o.A}(\mathbf{pa}_{o.A}, \mathbf{u}_{o.A}, \mathbf{pa}^r_{o.A}, \mathbf{u}^r_{o.A})$ with $\mathbf{Pa}_{O.A} \subseteq \mathbf{V}_\rho$ and $\mathbf{U}_{O.A} \subseteq \mathbf{U}_\rho$.

Note that a ground RSCM can be viewed as a standard SCM with typed variables and function/noise sharing across variables of the same type. The standard mechanism for each relational mechanism $f_{o.A}(\mathbf{pa}_{o.A}, \mathbf{u}_{o.A}, \mathbf{pa}^r_{o.A}, \mathbf{u}^r_{o.A})$ is the same, but with the argument signature $f_{o.A}(\mathbf{pa}_{o.A} \cup \mathbf{pa}^r_{o.A}, \mathbf{u}_{o.A} \cup \mathbf{u}^r_{o.A})$. A ground RSCM $\mathcal{M}_\rho$ is said to be recursive if it is recursive viewed as a standard RSCM (Pearl, 2009).

Next, we define the various observational, interventional, and counterfactual distributions that a ground RSCM induces. Note the similarity to standard SCMs (Pearl, 2009; Bareinboim, 2025), since ground RSCMs are simply standard SCMs.

Definition B.5 (RSCM-induced distributions)**.** Fix an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$, a relational skeleton ρ following the schema $\mathcal{S}$, and a corresponding ground RSCM $\mathcal{M}_\rho = \langle \mathbf{V}_\rho, \mathbf{U}_\rho, \mathcal{F}_\rho, P(\mathbf{U}_\rho) \rangle$. For any counterfactual events $\mathbf{Y}_\mathbf{x}, \dots, \mathbf{Z}_\mathbf{w}$ over $\mathbf{V}_\rho$, $\mathcal{M}_\rho$ induces the distribution

$$\begin{aligned} &P^{\mathcal{M}_\rho}(\mathbf{y}_\mathbf{x}, \dots, \mathbf{z}_\mathbf{w}) \\ &:= \sum_{\mathbf{u}_\rho} \mathbf{1}[\mathbf{Y}_\mathbf{x}(\mathbf{u}_\rho) = \mathbf{y}, \dots, \mathbf{Z}_\mathbf{w}(\mathbf{u}_\rho) = \mathbf{z}] P(\mathbf{u}_\rho) \end{aligned}$$

where the value $\mathbf{Y}_\mathbf{x}(\mathbf{u})$ is obtained by the standard SCM semantics over $\mathcal{F}_\rho$.

A distribution over ground variables $\mathbf{V}_\rho$ is said to be *observational*, *interventional*, or *counterfactual* depending on the form of the counterfactual event(s) it assigns probability to.

1. If none of the events involve interventions, i.e., each event is of the form $\mathbf{Y}$ (equivalently $\mathbf{Y}_\emptyset$), then $P^{\mathcal{M}_\rho}$ is *observational*. In this case we write $P^{\mathcal{M}_\rho}(\mathbf{y})$ for $\mathbf{Y} \subseteq \mathbf{V}_\rho$ or $P^{\mathcal{M}_\rho}(\mathbf{v}_\rho)$.

2. If all events share the same intervention and contain no nested interventions, i.e., all events are of the form $\mathbf{Y}_\mathbf{x}$ for a single intervention $\mathbf{X}{=}\mathbf{x}$, then $P^{\mathcal{M}_\rho}$ is *interventional*. In this case we write $P^{\mathcal{M}_\rho}(\mathbf{y} \mid do(\mathbf{x}))$ or $P^{\mathcal{M}_\rho}(\mathbf{y}_\mathbf{x})$, and similarly $P^{\mathcal{M}_\rho}(\mathbf{v}_\rho \mid do(\mathbf{x}))$.

3. If the distribution involves at least two events $\mathbf{Y}_\mathbf{x}, \mathbf{Z}_\mathbf{w}$ such that $\mathbf{X}$ and $\mathbf{W}$ are distinct variables or $\mathbf{x}$ and $\mathbf{w}$ are distinct values, then it is *counterfactual*. In this case we write $P^{\mathcal{M}_\rho}(\mathbf{y}_\mathbf{x}, \ldots, \mathbf{z}_\mathbf{w})$.

B.2. Relational Identification

Definition B.6 (Ground relational causal graph)**.** Consider a relational causal graph $\mathcal{G}$ over schema $\mathcal{S}$ and a skeleton ρ. The *ground relational causal graph* $\mathcal{G}_\rho$ is constructed as follows.

- For each node $O.V$ in $\mathcal{G}$ (relational or otherwise), and each instance $o \in \rho(O)$, include a node $o.V$ in $\mathcal{G}_\rho$.
- For each edge $O.B \to O.A$ in $\mathcal{G}$ where $O.A$ is a non-relational node, include an edge $o.B \to o.A$ in $\mathcal{G}_\rho$.
- For each edge $O.B \leftrightarrow T.B$ in $\mathcal{G}$ annotated with constraint ϕ, include an edge $o.A \leftrightarrow t.B$ in $\mathcal{G}_\rho$ for each $o \in \rho(O), t \in \rho(T)$ such that $\phi(o,t)$ holds.
- For each edge $T.W \to O.R$ in $\mathcal{G}$ where $O.R$ is a relational node with $R = (\mathbf{W}, \phi, \text{AGG})$, include edges $t.W \to o.R$ for every $y \in \rho(T)$ such that $\phi(o,t)$ holds.

Definition B.7 (Marginalized ground relational causal graph)**.** Fix $\mathcal{G}$ and ρ, and let $\mathcal{G}_\rho$ be the ground relational causal graph (Def. B.6). The *marginalized ground relational causal graph* $\bar{\mathcal{G}}_\rho$ is the graph on node set $\mathbf{V}_\rho$ obtained by marginalizing out all ground relational role nodes. Equivalently, start from $\mathcal{G}_\rho$, delete every relational node $o.R$, and for each deleted node add directed edges $t.W \to o.A$ whenever $\mathcal{G}_\rho$ contained $t.W \to o.R \to o.A$. Bidirected edges among nodes in $\mathbf{V}_\rho$ are inherited unchanged.

C. Further Examples

C.1. Extended Traffic Example

In this section, we consider an extended version of our running example (Ex. 1).

Example C.1 (Extended relational schema)**.** We extend the schema in Ex. 1 with an attribute $\mathsf{Ped}.V$ for whether they a pedestrian is visible and $\mathsf{Ped}.A$ for whether they look alert. An extended relational schema for traffic scenes would be

$$\begin{aligned}
\mathcal{E} &= \{\mathsf{Signal\ (Sig)},\ \mathsf{Car\ (Car)},\ \mathsf{Pedestrian\ (Ped)}\} \\
\mathcal{R} &= \{\mathsf{Ctrl}(\mathsf{Sig}, \mathsf{Ped}), \mathsf{Ctrl}(\mathsf{Sig}, \mathsf{Car}), \mathsf{Path}(\mathsf{Ped}, \mathsf{Car})\} \\
\mathcal{A} &= \{\mathsf{Sig}.W, \mathsf{Ped}.V, \mathsf{Ped}.A, \mathsf{Ped}.X, \mathsf{Car}.B\},
\end{aligned}$$

with all attributes binary-valued: $\mathsf{Sig}.W \in \{1,0\}$ denotes walk/drive; $\mathsf{Ped}.V \in \{1,0\}$ visible/not; $\mathsf{Ped}.A \in \{1,0\}$ alert/not;$\mathsf{Ped}.X \in \{1,0\}$ cross/wait; and $\mathsf{Car}.B \in \{1,0\}$ brake/go. $\mathsf{Ctrl}(\mathsf{Sig}, \mathsf{Ped})$ (resp. $\mathsf{Ctrl}(\mathsf{Sig}, \mathsf{Car})$) indicates that a signal controls a pedestrian (resp. car), and $\mathsf{Path}(\mathsf{Ped}, \mathsf{Car})$ indicates that the pedestrian is in the car's path.

Since the object and relation types are the same as in Ex. 1, the three relational skeletons shown in Fig. 1 can also be viewed as skeletons of this new schema. We give an example RSCM for this schema that departs from Ex. 3 in two ways: (i) it includes unobserved confounding, and (ii) it includes a relational parent consisting of more than one variable, illustrating why $\mathbf{W}$ in Def. 3.1 is a set and not a single variable.

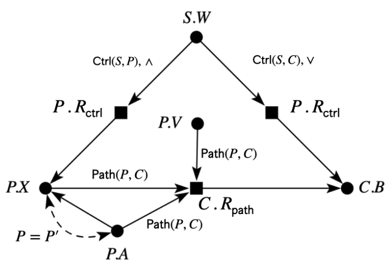

(a) Relational causal graph $\mathcal{G}$ for extended traffic RSCM in Ex 9.

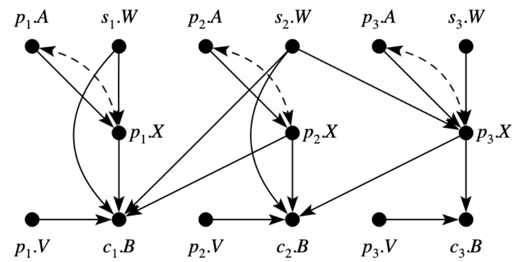

(b) Marginalized ground relational causal graph $\bar{\mathcal{G}}_{\rho_C}$ for skeleton ρ_C in Fig. 1.

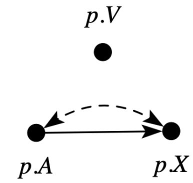

(c) Induced subgraph of $\bar{\mathcal{G}}_{\rho_C}$ on attributes of pedestrian p_1 (with instance identifiers omitted) for Ex. C.2.

Figure C.1.1. Relational causal graphs for the extended traffic examples.

Example 9 (RSCM for traffic scene). The endogenous variables are $\mathbf{V} = \{\mathsf{Sig}.W,\ \mathsf{Ped}.V,\ \mathsf{Ped}.,\ \mathsf{Ped}.X,\ \mathsf{Car}.B\}$. The exogenous variables $\mathbf{U}$ capturing unobserved factors (e.g., a pedestrian's intent to cross or a driver's alertness) are $\mathsf{Sig}.U_W \sim \mathcal{B}(0.3)$, $\mathsf{Ped}.U_{XA} \sim \mathcal{B}(0.4)$, $P, U_V \sim \mathcal{B}(0.8)$, and $\mathsf{Car}.U_B \sim \mathcal{B}(0.2)$. The mechanisms are

$$\begin{aligned}
\mathsf{Sig}.W &\leftarrow \mathsf{Sig}.U_W, \\
\mathsf{Ped}.V &\leftarrow \mathsf{Ped}.U_V \\
\mathsf{Ped}.A &\leftarrow \mathsf{Ped}.U_{XA} \\
\mathsf{Ped}.X &\leftarrow \mathsf{Ped}.U_{XA} \oplus \bigwedge_{\mathsf{Ctrl(Sig,Ped)}} \mathsf{Sig}.W, \\
\mathsf{Car}.B &\leftarrow \mathsf{Car}.U_B \oplus \left(\bigvee_{\mathsf{Ctrl(Sig,Car)}} \mathsf{Sig}.W \vee \bigvee_{\mathsf{Path(Ped,Car)}} (\mathsf{Ped}.X \vee \mathsf{Ped}.A) \wedge \mathsf{Ped}.V \right).
\end{aligned}$$

For each pedestrian, $\mathsf{Ped}.A$ and $\mathsf{Ped}.X$ are confounded by the pedestrian's unobserved intent $\mathsf{Ped}.U_{XA}$. $\mathsf{Car}.B$ has a relational parent $(\{\mathsf{Sig}.W\}, \mathsf{Ctrl}(\mathsf{Sig}, \mathsf{Car}), \vee)$ as in Ex. 3. It also has the relational parent $(\{\mathsf{Ped}.V, \mathsf{Ped}.A, \mathsf{Ped}.X\}, \mathsf{Path}(\mathsf{Ped}, \mathsf{Car}), \bot)$ which means: for each pedestrian p in the path of the car, the mechanism for $\mathsf{Car}.B$ may jointly use that pedestrian's triple $(\mathsf{Ped}.V, \mathsf{Ped}.A, \mathsf{Ped}.X)$. The relational causal graph for this RSCM is shown in Fig. C.1.1a.

It is important that $\{\mathsf{Ped}.V, \mathsf{Ped}.A, \mathsf{Ped}.X\}$ appears as a single relational parent rather than as three separate relational parents $(\{\mathsf{Ped}.V\}, \mathsf{Path}(\mathsf{Ped}, \mathsf{Car}))$, $(\{\mathsf{Ped}.A\}, \mathsf{Path}(\mathsf{Ped}, \mathsf{Car}))$, and $(\{\mathsf{Ped}.X\}, \mathsf{Path}(\mathsf{Ped}, \mathsf{Car}))$. This is because the braking condition depends on a within-pedestrian interaction: $\exists p$ in path $: \mathsf{Ped}.V \wedge (\mathsf{Ped}.A \vee \mathsf{Ped}.X)$, i.e., a given pedestrian must be visible and either crossing or alert to trigger braking. If we aggregate $\mathsf{Ped}.V, \mathsf{Ped}.A$, and $\mathsf{Ped}.X$ separately across pedestrians, we lose information about whether these properties are true of the same pedestrian. In general,

$$\bigvee_{\mathsf{Path(Ped,Car)}} (\mathsf{Ped}.X \vee \mathsf{Ped}.A) \wedge \mathsf{Ped}.V \neq \left(\bigvee_{\mathsf{Path(Ped,Car)}} \mathsf{Ped}.X \vee \bigvee_{\mathsf{Path(Ped,Car)}} \mathsf{Ped}.A \right) \wedge \bigvee_{\mathsf{Path(Ped,Car)}} \mathsf{Ped}.V$$

The right-hand side can be true even when no single pedestrian satisfies both conditions—for instance, one pedestrian is visible but not crossing/alert, while another is crossing/alert but not visible. The left-hand side is false in that situation.

This illustrates why Def. 3.1 allows a **relational parent to contain a set of variables**: it lets the mechanism represent interactions among attributes of the same related object. □

Next, we illustrate an application of Prop. 4.6 to show non-identifiability in this example.

Example C.2 (Relational non-identifiability using Prop. 4.6)**.** Continuing Ex. 9, consider the causal diagram $\mathcal{G}$ in Fig. C.1.1a, the source skeleton $\rho = \rho_A$, and the target skeleton $\rho_\star = \rho_C$ from Fig. 1. Say we have as input source distributions $\mathbb{P} = \{P(\mathbf{v}_\rho), P(\mathbf{v}_\rho \mid do(p_1.x)), P(\mathbf{v}_\rho \mid do(p_1.x, p_1.a))\}$, and we are interested in the query $P^{\rho_\star}(p_1.x \mid do(p_1.a))$ in $\rho_\star$, the causal effect of pedestrian p_1's alertness on whether or not they cross. Notice how in the marginalized ground graph

$\bar{\mathcal{G}}_{\rho_C}$ (Fig. C.1.1b) we see a bow-graph (Pearl, 2009) structure over $p_1.A$ and $p_1.X$. Standard identification theory usually suggests that in this case, $P^{\rho_\star}(p_1.x \mid do(p_1.a))$ is not identifiable from $\mathcal{G}_{\rho_C}$ and $P(\mathbf{v}_{\rho_\star})$. We show, using Prop. 4.6, that it is also not relationally identifiable from $\mathbb{P}$ and $\mathcal{G}$.

Following the notation of Prop. 4.6, our query concerns attributes of the instance $x = p_1$ in target $\rho_\star$, where $\mathbf{V}_{p_1} = \{p_1.V, p_1.X, p_1.A\}$.

The induced subgraph of $\bar{\mathcal{G}}_{\rho_C}$ on $\mathbf{V}_{p_1}$ (with instance identifiers omitted) is given in Fig. C.1.1c. From $\mathbb{P}$, we construct the restriction $\mathbb{P}|_P$ for instances p_1, p_2 in ρ of type P (Pedestrian) as follows. Note that $\mathbf{V}_{p_1} = \{p_1.V, p_1.X, p_1.A\}$ and $\mathbf{V}_{p_2} = \{p_2.V, p_2.X, p_2.A\}$.

1. $P(\mathbf{v}_\rho) \in \mathbb{P}$
   - p_1 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_1}) = P(\mathbf{v}_{p_1})$
   - p_2 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_2}) = P(\mathbf{v}_{p_2})$
2. $P(\mathbf{v}_\rho \mid do(p_1.x))$
   - p_1 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_1} \mid do(\{p_1.x\} \cap \mathbf{v}_{p_1})) = P(\mathbf{v}_{p_1} \mid do(p_1.x))$
   - p_2 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_2} \mid do(\{p_1.x\} \cap \mathbf{v}_{p_2})) = P(\mathbf{v}_{p_2})$
3. $P(\mathbf{v}_\rho \mid do(p_1.x, p_1.a))$
   - p_1 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_1} \mid do(\{p_1.x, p_1.a\} \cap \mathbf{v}_{p_1})) = P(\mathbf{v}_{p_1} \mid do(p_1.x, p_1.a))$
   - p_2 gives $P(\mathbf{v}_\rho \cap \mathbf{v}_{p_2} \mid do(\{p_1.x, p_1.a\} \cap \mathbf{v}_{p_2})) = P(\mathbf{v}_{p_2})$

Omitting identifiers, we get the restriction $\mathbb{P}|_P = \{P(\mathsf{Ped}.v, \mathsf{Ped}.X, \mathsf{Ped}.a), P(\mathsf{Ped}.v, \mathsf{Ped}.X, \mathsf{Ped}.a \mid do(\mathsf{Ped}.X)), P(\mathsf{Ped}.v, \mathsf{Ped}.X, \mathsf{Ped}.a \mid do(\mathsf{Ped}.X, \mathsf{Ped}.a)\}$ and the query $P(\mathsf{Ped}.X \mid do(\mathsf{Ped}.a))$. By counterfactual calculus, since the subgraph $\mathcal{G}_P$ in Fig. C.1.1c contains a bow-structure over $\mathsf{Ped}.A$ and $\mathsf{Ped}.X$, the query $P(\mathsf{Ped}.X \mid do(\mathsf{Ped}.a))$ is non-identifiable from $\mathcal{G}_P$ and $\mathbb{P}|_P$. Then, by Prop. 4.6, the original query $P^{\rho_\star}(p_1.x \mid do(p_1.a))$ is non-identifiable from $\mathbb{P}$ and $\mathcal{G}$.

D. Further Results and Proofs

D.1. Proofs for Sec. 3

The following proposition justifies how two isomorphic skeletons induce the 'same' counterfactual distributions over variables.

Proposition D.1 (Isomorphism-invariance of RSCM distributions)**.** *Consider an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$ and relational skeletons ρ, ρ' isomorphic under a mapping π. Then, for any counterfactual events $\mathbf{Y_x}, \ldots, \mathbf{Z_w}$ over $\mathbf{V}_\rho$,*

$$P^{\mathcal{M}_\rho}(\mathbf{y_x}, \ldots, \mathbf{z_w}) = P^{\mathcal{M}_{\rho'}}(\pi(\mathbf{y})_{\pi(\mathbf{x})}, \ldots, \pi(\mathbf{z})_{\pi(\mathbf{w})})$$

where $\pi(o.A) = \pi(o).A$ extends π to ground variables $o.A \in \mathbf{V}_\rho$.

Proof. Recall from Def. B.5 that

$$P^{\mathcal{M}_\rho}(\mathbf{y_x}, \ldots, \mathbf{z_w}) = \sum_{\mathbf{u}_\rho} \mathbf{1}[\mathbf{Y_x}(\mathbf{u}_\rho) = \mathbf{y}, \ldots, \mathbf{Z_w}(\mathbf{u}_\rho) = \mathbf{z}]\; P(\mathbf{u}_\rho),$$

and

$$P^{\mathcal{M}_{\rho'}}(\pi(\mathbf{y})_{\pi(\mathbf{x})}, \ldots, \pi(\mathbf{z})_{\pi(\mathbf{w})}) = \sum_{\mathbf{u}_{\rho'}} \mathbf{1}[\pi(\mathbf{Y})_{\pi(\mathbf{x})}(\mathbf{u}_{\rho'}) = \pi(\mathbf{y}), \ldots, \pi(\mathbf{Z})_{\pi(\mathbf{w})}(\mathbf{u}_{\rho'}) = \pi(\mathbf{z})]\; P(\mathbf{u}_{\rho'}).$$

We prove the desired equality by a change of variables $\mathbf{u}_{\rho'} = \pi(\mathbf{u}_\rho)$. The isomorphism $\pi : \rho \to \rho'$ induces a bijection on the exogenous assignments $\mathbf{u}_\rho \mapsto \pi(\mathbf{u}_\rho)$. Therefore, we re-index the second sum by writing $\mathbf{u}_{\rho'} = \pi(\mathbf{u}_\rho)$.

$$P^{\mathcal{M}_{\rho'}}(\pi(\mathbf{y})_{\pi(\mathbf{x})}, \ldots, \pi(\mathbf{z})_{\pi(\mathbf{w})}) = \sum_{\pi(\mathbf{u}_\rho)} \mathbf{1}[\pi(\mathbf{Y})_{\pi(\mathbf{x})}(\pi(\mathbf{u}_\rho)) = \pi(\mathbf{y}), \ldots, \pi(\mathbf{Z})_{\pi(\mathbf{w})}(\pi(\mathbf{u}_\rho)) = \pi(\mathbf{z})]\; P(\pi(\mathbf{u}_\rho)).$$

We claim that for any intervention assignment $\mathbf{x}$ and any exogenous assignment $\mathbf{u}_\rho$,

$$\pi(\mathbf{Y})_{\pi(\mathbf{x})}(\pi(\mathbf{u}_\rho)) = \pi(\mathbf{y}) \;\text{ in } \mathbf{M}_{\rho'} \iff \mathbf{Y}_{\mathbf{x}}(\mathbf{u}_\rho) = \mathbf{y} \;\text{ in } \mathbf{M}_\rho$$

and similarly for other counterfactual events $\mathbf{Z}_{\mathbf{w}}$.

To see this, fix a ground variable $o.A \in \mathbf{V}_\rho$. By Def. B.4, each relational parent $(\mathbf{W}, \phi, \text{AGG})$ in the template mechanism $f_{O.A}$ is instantiated in $\mathcal{M}_\rho$ as the multiset

$$\{\, t.W \mid T.W \in \mathbf{W},\; t \in \rho(T),\; \phi(o,t) \text{ holds in } \rho \,\},$$

and analogously in $\mathcal{M}_{\rho'}$. Since π is a skeleton isomorphism, it preserves relations and hence satisfaction of constraints:

$$\phi(o,t) \text{ holds in } \rho \iff \phi(\pi(o), \pi(t)) \text{ holds in } \rho'.$$

Hence the structural function (or intervened constant) for $\pi(o).A$ in $\mathcal{M}_{\rho'}$ is exactly the π-renaming of the structural function for $o.A$ in $\mathcal{M}_\rho$. This proves the claim, so that for every $\mathbf{u}_\rho$,

$$\begin{aligned}&\mathbf{1}[\pi(\mathbf{Y})_{\pi(\mathbf{x})}(\pi(\mathbf{u}_\rho)) = \pi(\mathbf{y}), \ldots, \pi(\mathbf{Z})_{\pi(\mathbf{w})}(\pi(\mathbf{u}_\rho)) = \pi(\mathbf{z})] \\ &\quad = \mathbf{1}[\mathbf{Y}_{\mathbf{x}}(\mathbf{u}_\rho) = \mathbf{y}, \ldots, \mathbf{Z}_{\mathbf{w}}(\mathbf{u}_\rho) = \mathbf{z}].\end{aligned}$$

It remains to show that for every assignment of values $\mathbf{u}_\rho$,

$$P(\mathbf{u}_\rho) \;\text{ in } \mathcal{M}_\rho = P(\pi(\mathbf{u}_\rho)) \;\text{ in } \mathcal{M}_{\rho'}$$

For each entity/relation type O, recall that the RSCM $\mathcal{M}$ specifies for each exogenous variable $O.U \in \mathbf{U}$ a distribution $O.U \sim P(O.U)$. By definition of the ground RSCM $\mathcal{M}_\rho$, for each $o \in \rho(O)$, $o.U \sim P(O.U)$. Since π preserves types, we also have $\pi(o).U \sim P(O.U)$ by definition of the ground RSCM $\mathcal{M}_{\rho'}$. Therefore, the above equality follows. □

Theorem 3.3 (Impossibility of observational inference across skeletons). Consider a schema $\mathcal{S}$, source skeletons $\rho_1, \ldots, \rho_l$, and target skeleton $\rho_\star$. Then, for any RSCM $\mathcal{M}$ over $\mathcal{S}$, there exists another RSCM $\mathcal{M}'$ over $\mathcal{S}$ such that $\mathcal{M}$ and $\mathcal{M}'$ agree on observational distributions $P(\mathbf{v}_{\rho_k})$ for every source skeleton ρ_k but disagree on the observational distribution $P(\mathbf{v}_{\rho_\star})$ of the target skeleton.

Proof idea. We will prove this by constructing a relational constraint $\phi_\star$ that evaluates to true only on skeletons isomorphic to the given $\rho_\star$. So, given $\mathcal{M}$, we will construct another SCM $\mathcal{M}'$ that is almost identical to $\mathcal{M}$, except that it has different behaviour when $\phi_\star$ is true. For example, such a constraint for ρ_A given in Ex. 1 would be:

$$\begin{aligned}\phi_A :\; &\exists \text{ Signal } S_1,\; \text{Pedestrian } P_1, P_2,\; \text{Car } C_1, C_2 \text{ such that } (P_1 \neq P_2) \wedge (C_1 \neq C_2)\\ &\wedge (\forall \text{ Signal } S, S = S_1)\\ &\wedge (\forall \text{ Pedestrian } P, P = P_1 \vee P = P_2)\\ &\wedge (\forall \text{ Car } C, C = C_1 \vee C = C_2)\\ &\wedge \mathsf{Ctrl}(S_1, P_1) \wedge \mathsf{Ctrl}(S_1, P_2)\\ &\wedge \mathsf{Ctrl}(S_1, C_1) \wedge \neg\mathsf{Ctrl}(S_1, C_2)\\ &\wedge \mathsf{Path}(P_1, C_1) \wedge \mathsf{Path}(P_1, C_2)\\ &\wedge \mathsf{Path}(P_2, C_1) \wedge \neg\mathsf{Path}(P_2, C_2)\end{aligned}$$

Proof. Since $\rho_\star$ is a finite relational skeleton, there exists a first-order formula $\phi_\star$ that is true for a given skeleton ρ if and only if $\rho \cong \rho_\star$. Such a $\phi_\star$ is constructed as follows. For each entity/relation type O and instance o of O in $\rho_\star$, introduce one existentially quantified variable. Check that each of these variables (for a given type) are distinct. Introduce a universally quantified variable of type O, and check that it is equal to atleast one of these o-variables. Finally, check that every relation R in the schema $\mathcal{S}$ holds on exactly those instance variables for which it is true in $\rho_\star$, and no others. By construction, $\phi_\star$ is true only on skeletons isomorphic to $\rho_\star$.

Having constructed $\phi_\star$, consider the given RSCM $\mathcal{M}$ and skeletons $\rho_1, \ldots, \rho_l$. Let $\mathcal{M}'$ be the same as $\mathcal{M}$, with the following changes. First, $\mathcal{M}'$ contains, for some arbitrary entity or relation type O with an observed attribute $O.A \in \mathbf{V}$, an additional exogenous variable $O.U$ with the same domain as $O.A$, so that $\mathbf{U}' = \mathbf{U} \cup \{O.U\}$. Second, $\mathcal{M}'$ has a function $f'_{O.A}$ as follows:

$$f'_{O.A}(\mathbf{pa}_{O.A}, \mathbf{u}'_{O.A}, \mathbf{pa}^r_{O.A}, \mathbf{u}^r_{O.A}) = \begin{cases} u_{O.A} & \phi_\star(X) \\ f_{O.A}(\mathbf{pa}_{O.A}, \mathbf{u}'_{O.A} \setminus \{O.u\}, \mathbf{pa}^r_{O.A}, \mathbf{u}^r_{O.A}) & \neg\phi_\star(X) \end{cases}$$

As a result, since $f'_{O.A}$ in $\mathcal{M}'$ is equal to $f_{O.A}$ in $\mathcal{M}$ whenever $\phi_\star$ is false, $\mathcal{M}'$ and $\mathcal{M}$ will induce the same observational distributions on skeletons $\rho_1, \ldots, \rho_l \ncong \rho_\star$. On the skeleton $\rho_\star$, however, the distribution $P(O.U)$ in $\mathcal{M}'$ can be chosen depending on $P^{\mathcal{M}}(\mathbf{v}_{\rho_\star})$ so that the different behaviour of $f'_{O.A}$ and $f_{O.A}$ results in different observational distributions of $\mathcal{M}'$ and $\mathcal{M}$. □

Remark D.1. The proof of Thm. 3.3 above does not rely on unobserved confounding between variables (be it of the same entity/relation instance, or across such instances). As such, Thm. 3.3 holds even if we restrict to the class of Markovian RSCMs.

Theorem 3.4 (Impossibility of causal inference within a skeleton). Consider a schema $\mathcal{S}$ where at least one entity or relation type has more than one observed attribute. For any relational SCM $\mathcal{M}$ over $\mathcal{S}$ and skeleton ρ, there exists another relational SCM $\mathcal{M}'$ over $\mathcal{S}$ such that $\mathcal{M}$ and $\mathcal{M}'$ agree on the observational distribution $P(\mathbf{v}_\rho)$ but disagree on some interventional distribution over $\mathbf{V}_\rho$.

Proof idea. We will construct an SCM $\mathcal{M}'$ such that for an entity/relation type O with observed attributes $O.A$ and $O.B$, the function $f'_{O.B}$ in $\mathcal{M}'$ can 'detect' that $O.A$ has been intervened on (for the same instance). Then, $f'_{O.B}$ different behaviour than $f_{O.B}$ when $O.A$ is under intervention, but the same behaviour otherwise. The fact that $O.A$ and $O.B$ belong to the same instance allows us to implement such 'detection', since the constraints in $f_{O.A}$ that hold for an instance o will also hold for o in $f'_{O.B}$.

Proof. By assumption, there exists some entity/relation type O such that $O.A, O.B \in \mathbf{V}$. First, assume WLOG, that $O.B \notin \mathbf{Pa}_{O.A}$ in $\mathcal{M}$. Define $\mathcal{M}'$ to be the same as $\mathcal{M}$, but with two modifications.

First, $\mathcal{M}'$ contains an additional exogenous variable $O.U$ with the same domain as $O.B$, so that $\mathbf{U}' = \mathbf{U} \cup \{O.U\}$. Second, $\mathcal{M}'$ has a function $f'_{O.B}$ as follows:

$$\begin{aligned} & f'_{O.B}(\mathbf{pa}_{O.B} \cup \mathbf{pa}_{O.A} \cup \{O.a\}, \mathbf{u}'_{O.B} \cup \mathbf{u}_{O.A}, \mathbf{pa}^r_{O.B} \cup \mathbf{pa}^r_{O.A}, \mathbf{u}^r_{O.B} \cup \mathbf{u}^r_{O.A}) \\ &= \begin{cases} f_{O.B}(\mathbf{pa}_{O.B}, \mathbf{u}'_{O.B} \setminus \{O.U\}, \mathbf{pa}^r_{O.B}, \mathbf{u}^r_{O.B}) & \textbf{if } O.a = f_{O.A}(\mathbf{pa}_{O.A}, \mathbf{u}'_{O.A} \setminus \{O.u\}, \mathbf{pa}^r_{O.A}, \mathbf{u}^r_{O.A}) \\ O.u & \text{otherwise} \end{cases} \end{aligned}$$

Above, $f'_{O.B}$ has an extended parent set for $O.B$ that takes in all the parents (endogenous and exogenous, non-relational and relational) of $O.A$. Since $O.A$ and $O.B$ belong to the same instance, for any relational parent $(\mathbf{W}, \phi) \in \mathbf{Pa}^r_{O.A} \cup \mathbf{U}^r_{O.A}$, and any skeleton ρ, $\phi(o, t)$ will hold in $f'_{O.B}$ in ρ iff it holds in $f_{O.A}$ in ρ.

In the observational regime, the condition $O.a = f_{O.A}(\ldots)$ is always true; therefore, $f'_{O.B}$ is exactly equal to $f_{O.B}$, and $\mathcal{M}$ and $\mathcal{M}'$ induce the same $P(\mathbf{v}_{\rho_\star})$. However, this is not the case under intervention. Consider an intervention that sets $o_1.A = a$ for some instance o_1 of O in ρ. There is some assignment $\mathbf{u}_\rho$ to the exogenous variables such that the valuation $o_1.A(\mathbf{u}_\rho) \neq a$. Under this assignment, the condition $o_1.a = f_{O.A}(\ldots)$ fails, and thus $o_1.B \leftarrow x_1.u$. The probability $P(O.U)$ in $\mathcal{M}'$ can be chosen such that $P^{\mathcal{M}'_\rho}(o_1.B = b \mid do(o_1.A = a)) = P((o_1.U = b) \neq P^{\mathcal{M}_\rho}(o_1.B = b \mid do(o_1.A = a))$ for some b in the domain of $O.B$.

□

Remark D.2. The proof of Thm. 3.4 above does not rely on unobserved confounding between variables of different entity/relation instances. As such, Thm. 3.4 holds even if we restrict to the class of ρ-Markovian RSCMs.

D.2. Proofs for Sec. 4.1

First, we will show how if an RSCM $\mathcal{M}$ induces by Def. 4.1 a causal graph $\mathcal{G}$, then for any skeleton ρ, the ground RSCM $\mathcal{M}_\rho$, when viewed as a standard SCM, induces (by the definition in (Sec. 2)) a causal graph that is equal to the marginalized ground graph $\bar{\mathcal{G}}_\rho$ (Def. B.7).

Lemma D.3. *Consider a relational schema $\mathcal{S}$, skeleton ρ, and RSCM $\mathcal{M}$ inducing the relational causal graph $\mathcal{G}$. Then, the ground RSCM $\mathcal{M}_\rho$ induces the marginalized ground graph $\bar{\mathcal{G}}_\rho$.*

Proof. Let $\mathcal{M}$, $\mathcal{G}$, ρ, and $\bar{\mathcal{G}}_\rho$ be as given in the statement of the lemma. Let $\mathcal{G}'$ be the causal graph induced by $\mathcal{M}_\rho$ viewed as a standard SCM, according to Sec. 2). $\mathcal{G}'$ and $\bar{\mathcal{G}}_\rho$ contain the same nodes by construction, since relational nodes have been marginalized from $\mathcal{G}_\rho$ to get $\bar{\mathcal{G}}_\rho$. We will show $\mathcal{G}' = \bar{\mathcal{G}}_\rho$ by showing that every edge in $\bar{\mathcal{G}}_\rho$ is also in $\mathcal{G}'$ and vice-versa.

Within-instance edges. Fix a type $O \in \mathcal{E} \cup \mathcal{R}$ and an instance $o \in \rho(O)$. For any variable $o.A, o.B \in \mathbf{V}_\rho$, the endogenous variables in $\mathcal{M}_\rho$,

1. Every within-instance edge in $\mathcal{G}'$ is also an edge in $\bar{\mathcal{G}}_\rho$.

$$
\begin{aligned}
& o.B \in \mathbf{Pa}_{o.A} \text{ in } \mathcal{M}_\rho \\
\implies & O.B \in \mathbf{Pa}_{O.A} \in \mathcal{M} && (\mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho\text{, Def. B.4}) \\
\implies & O.B \to O.A \in \mathcal{G} && (\text{by construction of } \mathcal{G} \text{ from } \mathcal{M}\text{, Def. 4.1}) \\
\implies & o.B \to o.A \in \mathcal{G}_\rho && (\text{by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}\text{, Def. B.6}) \\
\implies & o.B \to o.A \in \bar{\mathcal{G}}_\rho && (\text{by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho\text{, Def. B.7})
\end{aligned}
$$

$$
\begin{aligned}
& \exists o.U \in \mathbf{U}_\rho, o.U \in \mathbf{U}_{o.A} \cap \mathbf{U}_{o.B} \\
\implies & O.U \in \mathbf{U}_{O.A} \cap \mathbf{U}_{O.B} \text{ in } \mathcal{M} && (\text{since } \mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho) \\
\implies & O.A \overset{O=O'}{\leftrightarrow} O.B \in \mathcal{G} && (\text{by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & o.A \leftrightarrow o.B \in \mathcal{G}_\rho && (\text{by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & o.A \leftrightarrow o.B \in \bar{\mathcal{G}}_\rho && (\text{by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho)
\end{aligned}
$$

2. Every within-instance edge in $\bar{\mathcal{G}}_\rho$ is also an edge in $\mathcal{G}'$.

$$
\begin{aligned}
& o.B \to o.A \in \bar{\mathcal{G}}_\rho \\
\implies & o.B \to o.A \in \mathcal{G}_\rho && (\text{within-instance edges preserved by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho) \\
\implies & O.B \to O.A \in \mathcal{G} && (\text{by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & O.B \in \mathbf{Pa}_{O.A} \text{ in } \mathcal{M} && (\text{by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & o.B \in \mathbf{Pa}_{o.A} \text{ in } \mathcal{M}_\rho && (\text{since } \mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho)
\end{aligned}
$$

$$
\begin{aligned}
& o.B \leftrightarrow o.A \in \bar{\mathcal{G}}_\rho \\
\implies & o.B \leftrightarrow o.A \in \mathcal{G}_\rho && (\text{within-instance edges preserved by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho) \\
\implies & O.B \overset{O=O'}{\leftrightarrow} O.A \in \mathcal{G} && (\text{by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & \exists O.U \in \mathbf{U}, O.U \in \mathbf{U}_{O.A} \cap \mathbf{U}_{O.B} \text{ in } \mathcal{M} && (\text{by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & \exists o.U \in \mathbf{U}_\rho, o.U \in \mathbf{U}_{o.A} \cap \mathbf{U}_{o.B} \text{ in } \mathcal{M}_\rho && (\mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho)
\end{aligned}
$$

Cross-instance edges. Fix types $O, T \in \mathcal{E} \cup \mathcal{R}$ and non-identical instances $o \in \rho(O), t \in \rho(T)$. For any variables $o.A, t.B \in \mathbf{V}_\rho$,

1. Every cross-instance edge in $\mathcal{G}'$ is also an edge in $\bar{\mathcal{G}}_\rho$.

$$
\begin{aligned}
& t.B \in \mathbf{Pa}_{o.A} \text{ in } \mathcal{M}_\rho \\
\implies & \exists R = (\mathbf{W}, \phi, \text{AGG}) \text{ with } T.B \in \mathbf{W}, R \in \mathbf{Pa}^r_{O.A} \in \mathcal{M} \text{ such that } \phi(o,t) \quad (\mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho, \text{Def. B.4}) \\
\implies & T.B \overset{\phi,\text{AGG}}{\to} O.R \to O.A \in \mathcal{G} \quad \text{(by construction of } \mathcal{G} \text{ from } \mathcal{M}, \text{Def. 4.1)} \\
\implies & t.B \overset{\phi,\text{AGG}}{\to} o.R \to o.A \in \mathcal{G}_\rho \quad \text{(by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}, \text{Def. B.6)} \\
\implies & o.B \to o.A \in \bar{G}_\rho \quad \text{(by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho)
\end{aligned}
$$

$$
\begin{aligned}
& \exists z.U \in \mathbf{U}_\rho, z.U \in \mathbf{U}_{o.A} \cap \mathbf{U}_{o.B} \\
\implies & \exists R_1 = (\mathbf{W_1}, \phi_1, \text{AGG}_1) \in \mathbf{U}^r_{O.A} \text{ and } R_2 = (\mathbf{W_2}, \phi_2, \text{AGG}_2) \in \mathbf{U}^r_{T.B} \text{ in } \mathcal{M} \\
& \text{with } Z.U \in \mathbf{W_1} \cap \mathbf{W_2} \text{ and } \phi_1(o,z) \wedge \phi_2(t,z) \quad (\mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho) \\
\implies & O.A \overset{\exists Z:\phi_1(O,Z)\wedge\phi_2(T,Z)}{\leftrightarrow} T.B \in \mathcal{G} \quad \text{(by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & o.A \leftrightarrow o.B \in \mathcal{G}_\rho \quad \text{(by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & o.A \leftrightarrow o.B \in \bar{\mathcal{G}}_\rho \quad \text{(by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho \text{, which preserves bidirected edges)}
\end{aligned}
$$

2. Every cross-instance edge in $\bar{\mathcal{G}}_\rho$ is also an edge in $\mathcal{G}'$.

$$
\begin{aligned}
& t.B \to o.A \in \bar{\mathcal{G}}_\rho \\
\implies & t.B \to o.R \to o.A \in \mathcal{G}_\rho \text{ for some } R = (\mathbf{W}, \phi, \text{AGG}) \quad \text{(by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho) \\
\implies & T.B \overset{\phi,\text{AGG}}{\to} O.R \to O.A \in \mathcal{G} \quad \text{(by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & R \in \mathbf{Pa}^r_{O.A} \text{ in } \mathcal{M} \quad \text{(by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & t.B \in \mathbf{Pa}^r_{o.A} \text{ in } \mathcal{M}_\rho \quad \text{(since } \mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho)
\end{aligned}
$$

$$
\begin{aligned}
& t.B \leftrightarrow o.A \in \bar{\mathcal{G}}_\rho \\
\implies & o.B \leftrightarrow o.A \in \mathcal{G}_\rho \quad \text{(bidirected edges preserved by construction of } \bar{\mathcal{G}}_\rho \text{ from } \mathcal{G}_\rho) \\
\implies & O.B \overset{\phi}{\leftrightarrow} O.A \in \mathcal{G} \text{ for some } \phi \text{ such that } \phi(o,t) \quad \text{(by construction of } \mathcal{G}_\rho \text{ from } \mathcal{G}) \\
\implies & \exists Z.U \in \mathbf{U}, z \in \rho(Z), R_1 = (\mathbf{W_1}, \phi_1, \text{AGG}_1) \in \mathbf{U}^r_{O.A} \text{ and } R_2 = (\mathbf{W_2}, \phi_2, \text{AGG}_2) \in \mathbf{U}^r_{T.B} \text{ in } \mathcal{M} \\
& \text{with } Z.U \in \mathbf{W_1} \cap \mathbf{W_2} \text{ and } \phi_1(o,z) \wedge \phi_2(t,z) \quad \text{(by construction of } \mathcal{G} \text{ from } \mathcal{M}) \\
\implies & \exists z.U \in \mathbf{U}_\rho,\ z.U \in \mathbf{U}_{o.A} \cap \mathbf{U}_{t.B} \text{ in } \mathcal{M}_\rho \quad \text{(since } \mathcal{M}_\rho \text{ grounds } \mathcal{M} \text{ on } \rho)
\end{aligned}
$$

□

Corollary D.1. *Consider a relational schema $\mathcal{S}$, skeleton ρ, and RSCM $\mathcal{M}$ inducing the relational causal graph $\mathcal{G}$. Then, the ground RSCM $\mathcal{M}_\rho$ induces counterfactual distributions that satisfy all counterfactul equality constraints encoded in $\bar{\mathcal{G}}_\rho$.*

Proof. By Lemma D.3, we have that $\mathcal{M}_\rho$ induces $\bar{\mathcal{G}}_\rho$. (Xia et al., 2023, Lemma 1) shows that if a standard SCM $\mathcal{M}$ induces a causal diagram $\mathcal{G}$, then its induced distributions satisfy all counterfactual equality constraints encoded in $\mathcal{G}$. The result follows. □

Theorem 4.3 (Observational identification across skeletons). Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeleton ρ, and target skeleton $\rho_\star$. Let $o.A$ be an unconfounded variable in $\mathbf{V}_{\rho_\star}$. The conditional $P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A})$ is relationally identifiable from $\mathcal{G}$ and $P(\mathbf{v}_\rho)$ if there exists a source instance $o' \in \rho$ such that $o'.A$ is unconfounded and $\text{dom}(\mathbf{Pa}^r_{o.A}) \subseteq \text{dom}(\mathbf{Pa}^r_{o'.A})$. In this case, $P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}) = P(o'.a \mid \mathbf{pa}_{o'.A}, \mathbf{pa}^r_{o'.A})$.

Proof. Consider any two RSCMs $\mathcal{M}, \mathcal{M}'$ compatible with $\mathcal{G}$ such that $P^{\mathcal{M}_\rho}(\mathbf{v}_\rho) = P^{\mathcal{M}'_\rho}(\mathbf{v}_\rho)$. We need to show that

$$P^{\mathcal{M}_{\rho_\star}}(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}) = P^{\mathcal{M}'_{\rho_\star}}(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}).$$

Consider the $o' \in \rho$ given by the assumption. Since $\mathcal{M}, \mathcal{M}'$ agree on the observational distribution over $P(\mathbf{v}_\rho)$, this implies that

$$P^{\mathcal{M}_\rho}(o'.a \mid \mathbf{pa}_{o'.A}, \mathbf{pa}^r_{o'.A}) = P^{\mathcal{M}'_\rho}(o'.a \mid \mathbf{pa}_{x].A}, \mathbf{pa}^r_{o'.A})$$

It suffices, then, to show that for all values $o.a = o'.a, \mathbf{pa}_{o.A} = \mathbf{pa}_{o'.A}$, and $\mathbf{pa}^r_{o.A} = \mathbf{pa}^r_{o'.A}$ (with the latter comparison made for values in the support of $\mathbf{pa}^r_{o.A}$).

$$P^{\mathcal{M}_\rho}(o'.a \mid \mathbf{pa}_{o'.A}, \mathbf{pa}^r_{o'.A}) = P^{\mathcal{M}_{\rho_\star}}(o.A \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}) \tag{1}$$

and

$$P^{\mathcal{M}'_\rho}(o'.a \mid \mathbf{pa}_{o'.A}, \mathbf{pa}^r_{o'.A}) = P^{\mathcal{M}'_{\rho_\star}}(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}). \tag{2}$$

Let the structural equation for $o.A$ in $\mathcal{M}_{\rho_\star}$ be $o.A \leftarrow f_{O.A}(\mathbf{pa}_{o.A},\ \mathbf{u}_{o.A},\ \mathbf{pa}^r_{o.A})$, where $\mathbf{U}_{o.A}$ are the exogenous parents of $o.A$ in $\mathcal{M}_{\rho_\star}$. By condition (1) of unconfoundedness, $o.A$ shares no bidirected edge with any variable in $\bar{\mathcal{G}}_\rho$. By Prop. D.3, this implies that $\mathbf{U}_{o.A} \perp\!\!\!\perp (\mathbf{Pa}_{o.A}, \mathbf{Pa}^r_{o.A})$. Therefore, for any values $a, \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}$,

$$\begin{aligned}
P^{\mathcal{M}_{\rho_\star}}(o.a \mid \mathbf{pa}_{o.A},\ \mathbf{pa}^r_{o.A}) &= \sum_{\mathbf{u}_{o.A}} P(o.a \mid \mathbf{u}_{o.A}, \mathbf{pa}_{o.A},\ \mathbf{pa}^r_{o.A}) P(\mathbf{u}_{o.A} \mid \mathbf{pa}_{o.A},\ \mathbf{pa}^r_{o.A}) \\
&= \sum_{\mathbf{u}_{o.A}} P(o.a \mid \mathbf{u}_{o.A}, \mathbf{pa}_{o.A},\ \mathbf{pa}^r_{o.A}) P(\mathbf{u}_{o.A}) \qquad (\mathbf{U}_{o.A} \perp\!\!\!\perp (\mathbf{Pa}_{o.A}, \mathbf{Pa}^r_{o.A})) \\
&= \sum_{\mathbf{u}_{o.A}} \mathbf{1}[f_{O.A}(\mathbf{pa}_{o.A}, \mathbf{u}_{o.A}, \mathbf{pa}^r_{o.A}) = a] P(\mathbf{u}_{o.A})
\end{aligned}$$

Analogously, for $o'.A$ in $\mathcal{M}_\rho$, we get

$$P^{\mathcal{M}_\rho}(o'.a \mid \mathbf{pa}_{o'.A},\ \mathbf{pa}^r_{o'.A}) = \sum_{\mathbf{u}_{o'.A}} \mathbf{1}[f_{O.A}(\mathbf{pa}_{o'.A}, \mathbf{u}_{x].A}, \mathbf{pa}^r_{o'.A}) = a] P(\mathbf{u}_{o'.A})$$

Since $\mathcal{M}_\rho$ and $\mathcal{M}_{\rho_\star}$ are both groundings of the same RSCM $\mathcal{M}$, $\mathbf{U}_{o'.A}$ and $\mathbf{U}_{o.A}$ have the same domains, and whenever $\mathbf{u}_{o'.A} = \mathbf{u}_{o.A}$, we also have $P(\mathbf{u}_{o'.A}) = P(\mathbf{u}_{o.A})$. Additionally, whenever $\mathbf{pa}_{o.A} = \mathbf{pa}_{o'.A}$ and $\mathbf{pa}^r_{o.A} = \mathbf{pa}^r_{o'.A}$, because the mechanism $f_{O.A}$ is shared, we also have $\mathbf{1}[f_{O.A}(\mathbf{pa}_{o'.A}, \mathbf{u}_{x].A}, \mathbf{pa}^r_{o'.A}) = a] = \mathbf{1}[f_{O.A}(\mathbf{pa}_{o.A}, \mathbf{u}_{o.A}, \mathbf{pa}^r_{o.A}) = a]$. This proves the equality in Eq.(1). A similar calculation for the groundings of $\mathcal{M}'$ proves the equality in Eq.(2). □

Corollary D.4 (Observational identification across skeletons - Markovian)**.** *Consider a schema $\mathcal{S}$ and a Markovian relational causal graph $\mathcal{G}$. Given source skeletons $\rho_1, \ldots, \rho_k$ and a target skeleton $\rho_\star$, $P(\mathbf{v}_{\rho_\star})$ is identifiable from the distributions $\{P(\mathbf{v}_{\rho_k})\}_{k=1}^l$ and $\mathcal{G}$ assuming the support condition of Thm. 4.3 is met for every variable $o.A \in \mathbf{V}_{\rho_\star}$ by some ρ_k.*

Proof. A Markovian graph $\mathcal{G}$ contains no bidirected edges. Therefore, condition (1) of Thm. 4.3 is met for every $o.A \in \mathbf{V}_{\rho_\star}$ by some $o'.A$ in some ρ_k. Additionally, since $\mathcal{G}$ is Markovian, for any RSCM $\mathcal{M}$ compatible with $\mathcal{G}$, the ground RSCM $\mathcal{M}_{\rho_\star}$ is also Markovian. Since, for each $o.A \in \mathbf{V}_{\rho_\star}$, $P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A})$ is identifiable from some $P(\mathbf{v}_{\rho_k})$ by Thm. 4.3, so is the joint $P(\mathbf{v}_{\rho_\star})$ by the by the Markov factorization (Bareinboim, 2025, Thm. 2.4.1) and the compatibility of $P(\mathbf{v}_{\rho_\star})$ and the marginalized ground graph $\bar{\mathcal{G}}_\rho$.

$$P(\mathbf{v}_{\rho_\star}) = \prod_{o.a \in \mathbf{v}_{\rho_\star}} P(o.a \mid \mathbf{pa}_{o.A}, \mathbf{pa}^r_{o.A}).$$

Above, $\mathbf{Pa}_{o.A}, \mathbf{Pa}^r_{o.A} \subseteq \mathbf{V}_\rho$ are the graphical parents of $o.A$ in $\bar{\mathcal{G}}_\rho$ containing variables within the same instance and from different instances respectively. □

Proposition 4.4 (Sufficient condition for same-skeleton relational identification). Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, skeleton ρ, and family of interventional distributions $\mathbb{P}$ over $\mathbf{V}_\rho$. If $P(\mathbf{y}_* \mid \mathbf{x}_*)$ is identifiable via ctf-calculus from the marginalized ground graph $\bar{\mathcal{G}}_\rho$ and $\mathbb{P}$, then it is also relationally identifiable from $\mathcal{G}$ and $\mathbb{P}$.

Proof. Assume that the given $P(\mathbf{y}_* \mid \mathbf{x}_*)$ is identifiable via ctf-calculus from the marginalized ground graph $\bar{G}_\rho$ and $\mathbb{P}$. By the soundness of ctf-calculus (Correa & Bareinboim, 2025), this implies that for all (standard) SCMs $\mathcal{N}, \mathcal{N}'$ consistent with $\bar{G}_\rho$ and agreeing on $\mathbb{P}$, $\mathcal{N}, \mathcal{N}'$ also agree on $P(\mathbf{y}_* \mid \mathbf{x}_*)$. We need to show that for any two RSCMs $\mathcal{M}, \mathcal{M}'$ over $\mathcal{S}$ consistent with $\mathcal{G}$, if the ground RSCMs $\mathcal{M}_\rho$ and $\mathcal{M}'_\rho$ agree on $\mathbb{P}$, then they also agree on $P(\mathbf{y}_* \mid \mathbf{x}_*)$. By Lemma D.3, $\mathcal{M}_\rho$ and $\mathcal{M}'_\rho$ induce the marginalized ground graph $\bar{G}_\rho$. Since $\mathcal{M}_\rho$ and $\mathcal{M}'_\rho$ are a subset of the space of standard NCMs over $\mathbf{v}_\rho$, the result follows. □

Proposition 4.6 (Necessary condition for within-instance relational identification). Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeletons $\rho_1, \ldots, \rho_l$ with available interventional distributions $\mathbb{P}$, and a target skeleton $\rho_\star$. Let $o \in \rho_\star$ be a target instance and consider a counterfactual query $P(\mathbf{y}_\star \mid \mathbf{x}_\star)$ with $\mathbf{Y}_\star, \mathbf{X}_\star \subseteq \mathbf{V}_o$, the attributes of o.

Let the restriction $\mathbb{P}|_O$ be as follows. For each source skeleton ρ_k, each distribution $P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) \in \mathbb{P}$, and each object $o' \in \rho_k(O)$, include $P(\mathbf{v}_{\rho_k,o'} \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_{\rho_k,o'}))$ in $\mathbb{P}|_O$, with instance identifiers omitted. Let $\mathcal{G}_o$ be the induced subgraph of the marginalized ground graph $\bar{\mathcal{G}}_{\rho_\star}$ on $\mathbf{V}_o$ with instance identifiers omitted.

If $P(\mathbf{y}_\star \mid \mathbf{x}_\star)$ is non-identifiable via ctf-calculus from $\mathbb{P}|_O$ and $\mathcal{G}_o$, then it is relationally non-identifiable from $\mathcal{G}$ and $\mathbb{P}$.

Proof. Fix an instance $o \in \rho_\star(O)$ and a query $P(\mathbf{y}_* \mid \mathbf{x}_*)$ where $\mathbf{Y}_*, \mathbf{X}_* \subseteq \mathbf{V}_{\rho,x}$ as described.

Assume that $P(\mathbf{y}_* \mid \mathbf{x}_*)$ is not identifiable via ctf-calculus from $\mathcal{G}_o$ and the restriction $\mathbb{P}|_O$. Note that here, and in the remainder of the result, we ignore instance identifiers in $\mathbb{P}|_O$ and $\mathcal{G}_o$, simply considering within-instance attributes as in the standard SCM setting.

Then, by the completeness of ctf-calculus (Correa & Bareinboim, 2025), there exist two (standard) SCMs $\mathcal{N}, \mathcal{N}'$ consistent with $\mathcal{G}_o$that agree on all distributions in $\mathbb{P}|_O$ but disagree on $P(\mathbf{y}_* \mid \mathbf{x}_*)$. Using $\mathcal{N}, \mathcal{N}'$, we will construct two RSCMs $\mathcal{M}, \mathcal{M}'$ that (when grounded) agree on $\mathbb{P}$ but not on $P(\mathbf{y}_* \mid \mathbf{x}_*)$.

Construct $\mathcal{M}$ as follows. Let $\mathcal{M}$ contain the endogenous variables given in $\mathcal{G}$.

1. First, consider attributes of type O in $\mathcal{M}$. For each exogenous variable U in $\mathcal{N}$, let $\mathcal{M}$ contain exogenous variable $O.U$. For attributes $O.A$ belonging to type O, let the function determining $O.A$ in $\mathcal{M}$ be the same as that determining $o.A$ in $\mathcal{N}$. In particular, $\mathbf{Pa}^{\mathcal{M}}_{O.A} = \mathbf{Pa}^{\mathcal{N}}_{o.A}$ and $\mathbf{U}^{\mathcal{M}}_{O.A} = \mathbf{U}^{\mathcal{N}}_{o.A}$. $O.A$ has no non-relational parents (endogenous or exogenous) in $\mathcal{M}$.

2. Next, consider attributes of type $Y \neq X$ in $\mathcal{M}$. For each attribute $T.B$, let $\mathcal{M}$ contain an exogenous variable $T.U_B \sim \mathcal{B}(0.5)$. Define the function $f_{T.B} : T.B \leftarrow T.U_B$. In other words, for each type Y, each attribute $T.B$ is determined by an independent fair coin fliPed.

Define $\mathcal{M}'$ similarly but using functions from $\mathcal{N}$. Then, for any skeleton ρ, the groundings $\mathcal{M}_\rho, \mathcal{M}'_\rho$ consist of 'copies' of $\mathcal{N}$ and $\mathcal{N}'$ for every instance of type O, and coin flips for all other variables. Importantly, in both $\mathcal{M}$ and $\mathcal{M}'$, there are no relational effects; all instances are independent in any grounding.

We need to show that $\mathcal{M}$ and $\mathcal{M}'$ are consistent with $\mathcal{G}$. Since $\mathcal{M}$ and $\mathcal{M}'$ contain no relational effects, it suffices to show that they are consistent with non-relational edges in $\mathcal{G}$. Since attributes of type $Y \neq X$ are all independent by construction, $\mathcal{M}$ and $\mathcal{M}'$ are consistent with any edges in $\mathcal{G}$ incident to such attributes. For attributes of type O, note that for any instances o', o'' of type O in any skeleton ρ, the induced subgraphs $\mathcal{G}_{o'}$ and $\mathcal{G}_o$ are the same.

Next, we show that $\mathcal{M}$ and $\mathcal{M}'$ agree on $\mathbb{P}$. Consider some distribution $P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})$ in $\mathbb{P}$. Then,

$$
\begin{aligned}
P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_{\rho k} \mid do(\mathbf{x}_{k,j}) &= \prod_{T \in \mathcal{E} \cup \mathcal{R}} \prod_{t \in \rho_k(T)} p(\mathbf{v}_y \mid do(\mathbf{x}_{k,j})) \\
&\qquad \text{(no relational effects in } \mathcal{M} \implies \text{all instances independent in } \mathcal{M}_{\rho_k}) \\
&= \prod_{T \in \mathcal{E} \cup \mathcal{R}} \prod_{t \in \rho_k(T)} P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y), do(\mathbf{x}_{k,j} \setminus \mathbf{v}_y)) \\
&= \prod_{T \in \mathcal{E} \cup \mathcal{R}} \prod_{t \in \rho_k(T)} P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y)) \qquad \text{(all instances independent in } \mathcal{M}_{\rho_k}) \\
&= \left(\prod_{T \in \mathcal{E} \cup \mathcal{R} \setminus X} \prod_{t \in \rho_k(T)} P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y)) \right) \cdot \prod_{o \in \rho_k(O)} P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_x \mid do(\mathbf{x}_{k,j} \cap \mathbf{x}_y)) \\
&= \left(\prod_{T \in \mathcal{E} \cup \mathcal{R} \setminus X} \prod_{t \in \rho_k(T)} P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y)) \right) \cdot \prod_{o \in \rho_k(O)} P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_x \mid do(\mathbf{x}_{k,j} \cap \mathbf{x}_y)) \\
&\qquad \text{(variables of all but } X\text{-type objects are coin flips in } \mathcal{M}, \mathcal{M}') \\
&= \left(\prod_{T \in \mathcal{E} \cup \mathcal{R} \setminus X} \prod_{t \in \rho_k(T)} P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y)) \right) \cdot \prod_{o \in \rho_k(O)} P^{\mathcal{N}}(\mathbf{v}_x \mid do(\mathbf{x}_{k,j} \cap \mathbf{x}_y)) \\
&\qquad \text{(by construction of } \mathcal{M}) \\
&= \left(\prod_{T \in \mathcal{E} \cup \mathcal{R} \setminus X} \prod_{t \in \rho_k(T)} P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_y \mid do(\mathbf{x}_{k,j} \cap \mathbf{v}_y)) \right) \cdot \prod_{o \in \rho_k(O)} P^{\mathcal{N}'}(\mathbf{v}_x \mid do(\mathbf{x}_{k,j} \cap \mathbf{x}_y)) \\
&\qquad \text{(since } \mathcal{N}, \mathcal{N}' \text{ agree on } \mathbb{P}_x) \\
&= P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_{\rho k} \mid do(\mathbf{x}_{k,j}) \qquad \text{(by a symmetric derivation)}
\end{aligned}
$$

Finally, since $\mathcal{N}$ and $\mathcal{N}'$ disagree on $P(\mathbf{y}_* \mid \mathbf{x}_*)$, so do $\mathcal{M}_\rho$ and $\mathcal{M}'_\rho$.

□

D.3. Proofs for Sec. 5

The proofs of this section resemble that of the expressvity and correctness of NCM training in (Xia et al., 2023), adapted to enforce parameter sharing for variables of the same type and permutation-invariance for relational parents.

Assumptions. We assume, that in any domain of interest, the true RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ is as follows.

(I) $\mathcal{M}$ is ρ-Markovian.

(II) All variables in $\mathbf{V}$ have discrete, finite domains.

(III) There is a non-negative integer D such that for any skeleton ρ, the grounding $\mathcal{M}_\rho$ has variables $o.A$ with relational parent multisets of size at most D.

D.3.1. Discrete RSCMs

First, we introduce *discrete RSCMs*, following (Zhang et al., 2022b).

Definition D.1 (Discrete RSCM)**.** An RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$ is said to be a discrete RSCM if all variables in $\mathbf{V}$ are discrete and finite, and all variables in $\mathbf{U}$ are discrete.

We will show that the space of discrete RSCMs is equally expressive as that of all RSCMs satisfying assumptions (I)-(III). This licenses the assumption, without loss of generality, that the exogenous variables have discrete and finite domains.

We adapt certain definitions from (Zhang et al., 2022b).

Definition D.5 (Equivalence classes (Zhang et al., 2022b, Def. A.1))**.** For an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$, for every $O.A \in \mathbf{V}$, let functions in $\text{dom}_{\mathbf{Pa}_{O.A}} \times \text{dom}_{\mathbf{Pa}^r_{O.A}} \mapsto \text{dom}(O.A)$[6] be ordered by $\{h^{(i)}_{O.A} \mid i \in I_{O.A}\}$, where $I_{O.A} = \{1, \dots, m_{O.A}\}$ and $m_{O.A} = |\text{dom}_{\mathbf{Pa}_{O.A}} \times \text{dom}_{\mathbf{Pa}^r_{O.A}} \mapsto \text{dom}(O.A)|$. An *equivalence class* $\mathcal{U}^{(i)}_{O.A}$ for function $h^{(i)}_{O.A}$, $i = 1, \dots, m_{O.A}$, is a subset of $\text{dom}(\mathbf{U}_{O.A})$ such that

$$\mathcal{U}^{(i)}_{O.A} \;=\; \left\{ \mathbf{u}_{O.A} \in \text{dom}(\mathbf{U}_{O.A}) \;\middle|\; f_{O.A}(\cdot, \mathbf{u}_{O.A}) = h^{(i)}_{O.A} \right\}. \tag{26}$$

Definition D.6 (Canonical Partition (Zhang et al., 2022b, Def. A.2))**.** For an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$, $\{\mathcal{U}^{(i)}_{O.A} \mid i \in I_{O.A}\}$ is the *canonical partition* over the exogenous domain $\text{dom}(\mathbf{U}_{O.A})$ for every $O.A \in \mathbf{V}$.

Corollary D.2. *For any RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$, for each $O.A \in \mathbf{V}$, the function $f_{O.A} \in \mathcal{F}$ can be decomposed as:*

$$f_{O.A}(\mathbf{pa}_{O.A}, \mathbf{pa}^r_{O.A}, \mathbf{u}_{O.A}) \;=\; \sum_{i \in I_{O.A}} h^{(i)}_{O.A}(\mathbf{pa}_{O.A}, \mathbf{pa}^r_{O.A}) \mathbf{1}_{\,\mathbf{u}_{O.A} \in \mathcal{U}^{(i)}_{O.A}}. \tag{27}$$

Proof. This is an immediate consequence of Lemma A.3 in (Zhang et al., 2022b), considering the union of relational and non-relational parents. □

Corollary D.3. *Consider an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$. For any relational skeleton ρ, let the ground RSCM be $\mathcal{M}_\rho = \langle \mathbf{V}_\rho, \mathbf{U}_\rho, \mathcal{F}_\rho, P(\mathbf{U})_\rho\rangle$. Let the indexing set $\mathbf{I}$ be $\mathbf{I} = \bigtimes_{O.A \in \mathbf{V}} \bigtimes_{o \in \rho(O)} \mathbf{I}_{O.A}$, i.e., a product of indexing sets across the different types of variables, and for instances of each type. Then, for any variables $\mathbf{Y}, \dots, \mathbf{Z}, \mathbf{X}, \dots, \mathbf{W} \subseteq \mathbf{V}$,*

$$\begin{aligned} &P^{\mathcal{M}_\rho}(\mathbf{y}_{\mathbf{x}}, \dots \mathbf{z}_{\mathbf{w}}) \\ &= \sum_{\mathbf{i} \in \mathbf{I}} \mathbf{1}_{\mathbf{Y}_{\mathbf{x}}(\mathbf{i}) = \mathbf{y}, \dots, \mathbf{Z}_{\mathbf{w}}(\mathbf{i}) = \mathbf{z}} \cdot P\left(\bigcap_{O.A \in \mathbf{V}, o \in \rho(O)} \mathcal{U}^{(i)}_{O.A} \right) \qquad \text{(where } i \text{ is the index for } o.A \text{ in } \mathbf{i}) \end{aligned}$$

where $\mathbf{Y}_{\mathbf{x}}(\mathbf{i}) = \{Y_{\mathbf{x}}(\mathbf{i}) = y \mid o.B \in \mathbf{Y}\}$ and each $o.B_{\mathbf{x}}(\mathbf{i})$ is recursively computed as

$$o.B_{\mathbf{x}}(\mathbf{i}) = \begin{cases} \mathbf{x}_{o.B} & \text{if } o.B \in \mathbf{X} \\ h^{(i)}_{O.B}((\mathbf{Pa}_{o.B})_{\mathbf{x}}(\mathbf{i})) & \text{otherwise} \end{cases}$$

Proof. This follows from (Zhang et al., 2022b, Lemma A.4), noting that each instance $o.A$ is determined by the same function $f_{O.A}$ of its parents in $\mathcal{M}_\rho$. □

Corollary D.4. *Consider an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$ inducing relational causal graph $\mathcal{G}$. For any relational skeleton ρ, let the indexing set be $\mathbf{I}$ as in Cor. D.3. For a given type $X \in \mathcal{E} \cup \mathcal{R}$, $\mathcal{C}(\mathcal{G})_X = \{\mathbf{C} \subseteq \mathbf{V} \mid C$ is a bidirected connected component in $\mathcal{G}$ containing some $O.A\}$. Then, for the ground RSCM $\mathcal{M}_\rho$,*

$$P\left(\bigcap_{O.A \in \mathbf{V}, o \in \rho(O)} \mathcal{U}^{(i)}_{O.A} \right) = \prod_{X \in \mathcal{E} \cup \mathcal{R}} \prod_{o \in \rho(O)} \prod_{\mathbf{C} \in \mathcal{C}(\mathcal{G})_X} P(\bigcap_{O.A \in \mathbf{C}} \mathcal{U}^{i}_{O.A}).$$

Proof. This follows by a similar argument as in (Zhang et al., 2022b, Lemma A.5), noting that since $\mathcal{M}$ is ρ-Markovian, exogenous variables are independent across different instances o, and identically distributed across different instances o of the same type O. □

Corollary D.5 (Discrete RSCM expressiveness)**.** *Consider an RSCM $\mathcal{M} = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U})\rangle$ satisfying assumptions (I)-(III) and inducing relational causal graph $\mathcal{G}$. Then, there exists a discrete RSCM $\mathcal{M}'$ consistent with $\mathcal{G}$ such that for any skeleton ρ, the ground RSCMs $\mathcal{M}_\rho$ and $\mathcal{M}'_\rho$ agree on all counterfactual distributions over $\mathbf{V}_\rho$.*

Proof. This follows from (Zhang et al., 2022b, Lemma A.6), Cor. D.4, and Cor. D.3.

□

[6]Here, $\text{dom}_{\mathbf{Pa}^r_{O.A}}$ contains multisets of size $\leq D$ for each relational parent, per assumption (II).

D.3.2. NEURAL RSCMS

Theorem 5.2 (Expressivity of RNCMs). Consider a relational schema $\mathcal{S}$. For every RSCM $\mathcal{M}$ over $\mathcal{S}$ inducing relational causal graph $\mathcal{G}$, there exists a $\mathcal{G}$-RNCM $\mathcal{N}$ such that for every skeleton ρ, the ground RSCMs $\mathcal{M}_\rho$ and $\mathcal{N}_\rho$ induce the same counterfactual distributions over $\mathbf{V}_\rho$.

Proof. Fix an RSCM $\mathcal{M}^*$ over $\mathcal{S}$ inducing relational causal graph $\mathcal{G}$. By the discrete RSCM expressiveness corollary above, there exists a discrete RSCM $\mathcal{M}' = \langle \mathcal{S}, \mathbf{V}, \mathbf{U}', \mathcal{F}', P(\mathbf{U}')\rangle$ consistent with $\mathcal{G}$ such that for every skeleton ρ, the ground rscms $\mathcal{M}^*_\rho$ and $\mathcal{M}'_\rho$ agree on all counterfactual distributions over $\mathbf{V}_\rho$. Thus, it suffices to construct a $\mathcal{G}$-RNCM $\mathcal{N}$ that agrees with $\mathcal{M}'_\rho$ on all counterfactuals, for every ρ.

We follow the proof strategy of (Xia et al., 2023): (i) represent the discrete exogenous distribution using uniform $\mathcal{U}([0,1])$ noise via a neural inverse probability integral transform [Lemma 5](Xia et al., 2021), and (ii) represent each mechanism $f_{O.A}$ using a multi-layer perceptron (MLP).

Since $\mathcal{M}'$ is ρ-Markovian (Assumption (I)), bidirected edges occur only among attributes of the same instance o. Fix a type $X \in \mathcal{E} \cup \mathcal{R}$, and let $\mathcal{C}(\mathcal{G})_X = \{\mathbf{C}_1, \ldots, \mathbf{C}_k\}$ be the set of bidirected maximal cliques among $\{O.A : O.A \in \mathbf{V}\}$ in $\mathcal{G}$. We construct $\mathcal{N}$ as in Def. 5.1 by introducing, for each $\mathbf{C} \in \mathcal{C}(\mathcal{G})_X$, an exogenous variable $\hat{U}_\mathbf{C} \sim \mathcal{U}([0,1])$. For each clique $\mathbf{C}$, let $\mathbf{U}'_\mathbf{C}$ denote the tuple of discrete exogenous variables in $\mathcal{M}'$ that are the shared exogenous parents of the variables in $\mathbf{C}$ (within an instance of type O). Since $\mathbf{U}'_\mathbf{C}$ has a finite domain, Lemma 5 of (Xia et al., 2021) provides an MLP $g_\mathbf{C}$ such that $g_\mathbf{C}(\hat{U}_\mathbf{C})$ has the same distribution as $\mathbf{U}'_\mathbf{C}$.

Next, consider some variable $O.A \in \mathbf{V}$. Under Assumptions (II)–(III), the input space dom$(\mathbf{Pa}_{O.A}) \times$ dom$(\mathbf{Pa}^r_{O.A}) \times$ dom$(\mathbf{U}'_{O.A})$ is finite (the relational-parent domain is finite because each multiset has size at most D). Hence, by Lemma 4 of (Xia et al., 2021), there exists an MLP $h_{O.A}$ that agrees with $f'_{O.A}$ for each possible input.

Composing MLPs $h_{O.A}$ with every MLP $g_\mathbf{C}$ such that $O.A \in \mathbf{C}$ yields a function $\hat{f}_{O.A}$: dom$(\mathbf{Pa}_{O.A}) \times$ dom$(\mathbf{Pa}^r_{O.A}) \times$ dom$(\hat{\mathbf{U}}_{O.A}) \rightarrow$ dom$(O.A)$.

For any skeleton ρ, the ground RSCM $\mathcal{N}_\rho$ shares the same function $\hat{f}_{O.A}$ across all instances $o \in \rho(O)$; this is also the case in $\mathcal{M}_\rho$. Therefore, for every ground variable $o.A \in \mathbf{V}_\rho$, $\mathcal{M}_\rho$ and $\mathcal{N}_\rho$ share the same distribution over exogenous parents and the same mechanism. It follows that $\mathcal{M}'_\rho$ and $\mathcal{M}^*_\rho$ agree on all counterfactual distributions. □

Definition D.2 (Data-dependent relational counterfactual identification)**.** Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeletons $\rho_1, \ldots, \rho_l$, source distributions $\mathbb{P} = \{\{P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) > 0\}_{j=1}^{m_k}\}_{k=1}^{l}$, and target skeleton $\rho_\star$. Let $P(\mathbf{y}_* \mid \mathbf{x}_*)$ be a target query with $\mathbf{Y}_\star, \mathbf{X}_\star \subseteq \mathbf{V}_{\rho_\star}$.

We say $P(\mathbf{y}_\star \mid \mathbf{x}_\star)$ is *relationally identifiable* from $\mathcal{G}$ and $\mathbb{P}$ if for any RSCMs $\mathcal{M}, \mathcal{M}'$ consistent with $\mathcal{G}$ agreeing on the source data, so that for every ρ_k and $j = 1, \ldots, m_k$,

$$P^{\mathcal{M}_{\rho_k}}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P^{\mathcal{M}'_{\rho_k}}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))$$

they also agree on the query:

$$P^{\mathcal{M}_{\rho_\star}}(\mathbf{y}_\star \mid \mathbf{x}_\star) = P^{\mathcal{M}'_{\rho_\star}}(\mathbf{y}_\star \mid \mathbf{x}_\star).$$

Otherwise, the query is *relationally non-identifiable dependent on the data*.

Corollary D.6 (Correctness of RelationalNeuralID (Alg. 1))**.** *Consider a schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source skeletons $\rho_1, \ldots, \rho_l$, source distributions $\mathbb{P} = \{\{P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) > 0\}_{j=1}^{m_k}\}_{k=1}^{l}$, and target skeleton $\rho_\star$. Let $P(\mathbf{y}_* \mid \mathbf{x}_*)$ be a target query with $\mathbf{Y}_\star, \mathbf{X}_\star \subseteq \mathbf{V}_{\rho_\star}$. Let Q be the output of* RelationalNeuralID *(Alg. 1) given these inputs. Then, Q is not* `FAIL` *if and only if the query $P(\mathbf{y}_* \mid \mathbf{x}_*)$ is relationally identifiable dependent on the data from $\mathcal{G}$ and $\mathbb{P}$, in which case $Q = P(\mathbf{y}_* \mid \mathbf{x}_*)$.*

Proof. Fix a relational causal diagram $\mathcal{G}$. By Thm. 5.2, for every RSCM $\mathcal{M}$ consistent with $\mathcal{G}$, there exists a $\mathcal{G}$-RNCM agreeing with $\mathcal{M}$ on counterfactual distributions for every possible skeleton ρ. Additionally, every $\mathcal{G}$-RNCM is consistent with $\mathcal{G}$ by construction. Therefore, minimizing / maximizing the query Q in the space of RSCMs consistent with $\mathcal{G}$ and inducing the given $\mathbb{P}$ is equivalent to minimizing / maximizing the query Q in the space of $\mathcal{G}-$ RSCMs inducing the given $\mathbb{P}$. The correctness of Alg. 1 thus follows from the definition of data-dependent relational identification (Def. D.2).

□

Algorithm 2 RelationalNeuralEstimation

Input: schema $\mathcal{S}$, relational causal graph $\mathcal{G}$, source data $\mathcal{D} = \left\{\left(\rho_k,\ \{P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))\}_{j=1}^{m_k}\right)\right\}_{k=1}^{l}$, target skeleton $\rho_\star$, query $P(\mathbf{y}_* \mid \mathbf{x}_*)$
$\hat{M} \leftarrow \mathcal{G}$-RNCM
$\hat{\theta} \leftarrow \theta \in \Theta(\hat{M})$ subject to $\forall k, j\ \ P^{\hat{M}_{\rho_k}(\theta)}(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j})) = P(\mathbf{v}_{\rho_k} \mid do(\mathbf{x}_{k,j}))$
$q \leftarrow P^{\hat{M}_{\rho_\star}(\hat{\theta}_l)}(\mathbf{y}_* \mid \mathbf{x}_*)$
return q

D.4. Additional Algorithms

Estimating identifiable queries. Alg. 2 provides an algorithm for point-estimation of a given query from graph and data. It modifies Alg. 1 by simply fitting a single $\mathcal{G}$-RNCM to the available data, and outputting the value of the query for that RNCM. Since Alg. 2, it is correct only assuming that the given query is relationally identifiable from the given graph and distributions.

E. Experimental Details

E.1. Compute and Implementation

All implementations are in Python, adapted from the Neural Causal Models codebase (Xia et al., 2023). Neural Causal Models are built using PyTorch (Paszke et al., 2019) and trained using PyTorch Lightning (Falcon & Cho, 2020). We trained our models on a single NVIDIA H100 GPU on a shared compute cluster with 2x Intel Xeon Platinum 8480+ CPUs (112 cores total, 224 threads) at up to 3.8 GHz, and 210 MiB L3 cache

E.2. Data Generation

We generated all synthetic data using *regional canonical models* (Xia et al., 2023, Def. 11) adapted to share functions across variables of the same type. Concretely, given a ρ-Markovian relational causal graph $\mathcal{G}$, we introduce one exogenous variable $U_{\mathbf{C}} \sim \mathcal{U}([0,1])$ for each maximal bidirected clique $\mathbf{C}$ in $\mathcal{G}$, to represent within-instance latent confounding encoded by $\mathcal{G}$. For each endogenous variable $O.A$, we then sample a random deterministic function

$$f_{O.A} : \mathrm{dom}(\mathbf{Pa}_{O.A}) \times \mathrm{dom}(\mathbf{Pa}^r_{O.A}) \times \mathrm{dom}(\mathbf{u}_{O.A}) \to \mathrm{dom}(O.A)$$

according to the canonical construction of Xia et al. (2023), where $\mathbf{u}_{O.A} = \{U_{\mathbf{C}} : O.A \in \mathbf{C}\}$.

For each relational parent $(\mathbf{W}, \phi, \mathrm{AGG}) \in \mathbf{Pa}^r_{O.A}$, we set its domain to be the set of histograms over $\mathrm{dom}(\mathbf{W})$ induced by multisets of size at most D; in our experiments $D = 5$, which upper-bounds the relational-parent multiset sizes in Fig. 1. Note that since we assume discrete endogenous domains, the count is a sufficient statistic for a multiset of $\mathbf{w}$-values.

For each trial, we sample a random regional canonical model as above. Given a skeleton ρ, to generate data from $P(\mathbf{v}_\rho)$, we instantiate variables $o.A$ and $o.U_{\mathbf{C}}$ for every instance o, and reuse the same mechanism for every variable of the same type. Due to the expressivity of canonical models (Zhang et al., 2022b; Xia et al., 2023), this avoids bias in data-generation induced by choosing a particular parametric model.

E.3. Model Architecture and Training

We refer readers to the comprehensive Appendix B.6 in (Xia et al., 2023) for details on the MLE-NCM architecture and training procedure. In addition to the clique-level noise variables (Def. 5.1), an MLE NCM additionally includes a Gumbel noise variable for each attribute.

Given a relational causal graph $\mathcal{G}$, we initialize a module for each mechanism $f_{O.A}$ to get an RNCM $\hat{\mathcal{M}}$. This is reused for every instance $o.A$ across source and target skeletons. To train to maximize a query $P(\mathbf{y}_* \mid \mathbf{x}_*)$ on a target $v_{\rho_\star}$ given data $\mathcal{D}$ from sources $\rho_1, \ldots, \rho_k$, where each source ρ_k has data $\{\mathbf{v}_{\rho_\mathbf{k}, \mathbf{z_{j,k}}}{}^{(i)}\}_{i=1}^{n_{j,k}}$ comprising $n_{j,k}$ datapoints from $j = 1, \ldots, m_k$ regimes, we use the following modified MLE-NCM loss.

Neighborhood	Target skeleton	Query
N_1	ρ_A	$P(c_2.B = 1 \mid do(p_2.X = 1)$
N_1	ρ_C	$P(c_3.B = 1 \mid do(p_3.X = 1)$
N_2	ρ_A	$P(c_1.B = 1 \mid do(s_1.W = 1, p_1.X = 1, p_2.X = 1)$
N_2	ρ_C	$P(c_2.B = 1 \mid do(s_2.W = 1, p_2.X = 1, p_3.X = 1)$
N_3	ρ_B	$P(c_1.B = 1 \mid do(s_1.W = 1, s_2.W = 1, p_1.X = 1, p_2.X = 1)$
N_3	ρ_C	$P(c_1.B = 1 \mid do(s_1.W = 1, s_2.W = 1, p_1.X = 1, p_2.X = 1)$

Table E.4.1. Target queries used Exp. 6.1 and Fig. 3.

$$L(\hat{\mathcal{M}}, \mathcal{D}) = \sum_{k=1}^{l} \sum_{j=1}^{m_k} \frac{1}{n_{j,k}} \sum_{i=1}^{n_{j,k}} -\log P^{\hat{\mathcal{M}}_{\rho_k}}(\mathbf{v}^{(i)}_{\rho_k, \mathbf{z}_{j,k}}) - \lambda \log P^{\hat{\mathcal{M}}_{\rho_\star}}(\mathbf{y} \mid \mathbf{x}_*) \tag{3}$$

Here, λ is a parameter that decreases during training. To minimize the query, we replace the $\lambda \log P^{\hat{\mathcal{M}}_{\rho_\star}}(\mathbf{y} \mid \mathbf{x}_*)$ term with $\lambda(1 - \log P^{\hat{\mathcal{M}}_{\rho_\star}}(\mathbf{y} \mid \mathbf{x}_*))$. In our estimation experiments (Exp. 6.1), we set $\lambda = 0$ and train only one RNCM.

All modules contain 2 hidden layers with 128 neurons each, and were trained for 200 epochs (often converging earlier) with a learning rate of 10^{-3} and a batch size of 1000 datapoints.

E.4. RNCM Estimation Experiments: Details and Extension

Specification of queries. In Table E.4.1, we give the the exact target queries used in Exp. 6.1.

E.5. RNCM Identification Experiments: Details

In Fig. E.5.1, we show the marginalized ground causal graphs for the source skeleton ρ and target skeleton $\rho_\star$ of both the relational bow and the relational ID graphs.

We also prove the identifiability result for the different-car query in both cases.

Proposition E.1. *Given source skeleton ρ and target skeleton $\rho_\star$ as in Exp. 6.2, and causal graph $\mathcal{G}_{\mathsf{bow}}$ (Fig. E.5.1, top left), the distribution $P^{\rho_\star}(c_3.y)$ is relationally identifiable from $P(\mathbf{v}_\rho)$ and $\mathcal{G}_{\mathsf{bow}}$.*

Proof. Consider any RSCM $\mathcal{M}$ consistent with $\mathcal{G}_{\mathsf{bow}}$. Since $\mathcal{G}_{\mathsf{bow}}$ is ρ-Markovian, so is $\mathcal{M}$. Let $\mathbf{U}_X$ (resPed., $\mathbf{U}_Y$) denote the non-relational exogenous variables affecting $\mathsf{Car}.X$ (resPed., $\mathsf{Car}.Y$), and $\mathbf{U}'_X = \mathbf{U}'_X \setminus \mathbf{U}'_Y$ Then, in the ground RSCM $\mathcal{M}_{\rho_\star}$,

$$\begin{aligned}
P^{\mathcal{M}_{\rho_\star}}(c_3.y) &= \sum_{c_3.x, c_3.\mathbf{u}_Y}^{\mathcal{M}_{\rho_\star}} P^{\mathcal{M}_{\rho_\star}}(c_3.y \mid c_3.x, c_3.\mathbf{u}_Y) P^{\mathcal{M}_{\rho_\star}}(c_3.x \mid c_3.\mathbf{u}_Y) P^{\mathcal{M}_{\rho_\star}}(c_3.\mathbf{u}_Y) \\
&= \sum_{c_3.x, c_3.\mathbf{u}_Y} \mathbf{1}_{f_{\mathsf{Car}.Y}(c_3.x, \mathbf{u}_Y) = c_3.y} \sum_{c_3.\mathbf{u}'_X} \mathbf{1}_{f_{\mathsf{Car}.X}(\mathbf{u}_X) = c_3.x} P^{\mathcal{M}_{\rho_\star}}(c_3.\mathbf{u}_Y, c_3.\mathbf{u}'_X) \\
&\qquad \text{($c_3.Y$ has no relational parents)} \\
&= \sum_{c_3.x, c_3.\mathbf{u}_Y} \mathbf{1}_{f_{\mathsf{Car}.Y}(c_3.x, \mathbf{u}_Y) = c_3.y} \sum_{c_3.\mathbf{u}'_X} \mathbf{1}_{f_{\mathsf{Car}.X}(\mathbf{u}_X) = c_3.x} P^{\mathcal{M}_{\rho}}(c_2.\mathbf{u}_Y, c_2.\mathbf{u}'_X) \\
&\qquad \text{(exogenous distributions are shared across c in $\rho, \rho_\star$)} \\
&= \sum_{c_2.x, c_2.\mathbf{u}_Y} \mathbf{1}_{f_{\mathsf{Car}.Y}(c_2.x, \mathbf{u}_Y) = c_2.y} \sum_{c_2.\mathbf{u}'_X} \mathbf{1}_{f_{\mathsf{Car}.X}(\mathbf{u}_X) = c_2.x} P^{\mathcal{M}_{\rho}}(c_2.\mathbf{u}_Y, c_2.\mathbf{u}'_X) \\
&\qquad \text{(functions $f_{\mathsf{Car}.X}, f_{\mathsf{Car}.Y}$ are shared across c in $\rho, \rho_\star$)} \\
&= P^{\mathcal{M}_\rho}(c_2.y) \qquad \text{(by a symmetric derivation)}
\end{aligned}$$

This proves that for any RSCMs $\mathcal{M}, \mathcal{M}'$ consistent with $\mathcal{G}_{\mathsf{bow}}$, we have $P^{\mathcal{M}_\rho}(c_2.y) = P^{\mathcal{M}_{\rho_\star}}(c_3.y)$ and $P^{\mathcal{M}'_\rho}(c_2.y) = P^{\mathcal{M}'_{\rho_\star}}(c_3.y)$. If $\mathcal{M}, \mathcal{M}'$ agree on the source data, then $P^{\mathcal{M}_\rho}(c_2.y) = P^{\mathcal{M}'_\rho}(c_2.y)$, and we are done. □

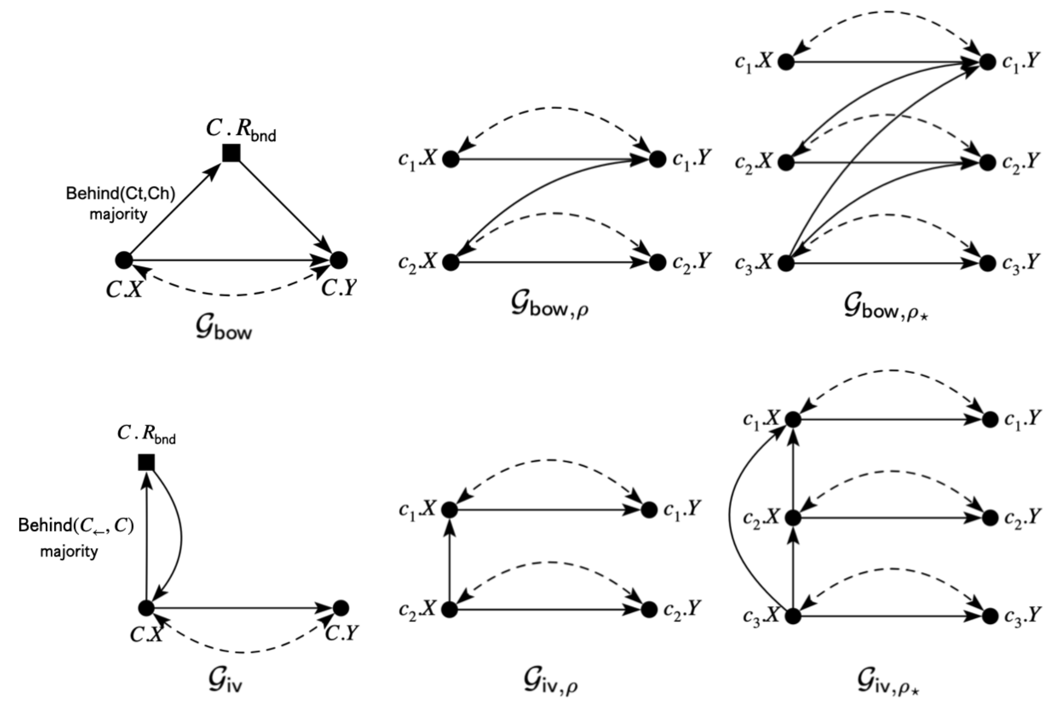

Figure E.5.1. Marginalized ground graphs corresponding to $\mathcal{G}_{\text{bow}}$ (top) and $\mathcal{G}_{\text{iv}}$ (bottom) for the source ρ (middle) and target $\rho_\star$ (right) in Exp. 6.2.

Proposition E.2. *Given source skeleton ρ and target skeleton $\rho_\star$ as in Exp. 6.2, and causal graph $\mathcal{G}_{\text{bow}}$ (Fig. E.5.1, top left), the query $P^{\rho_\star}(c_3.y \mid do(c_2.x))$ is relationally identifiable from $P(\mathbf{v}_\rho)$ and $\mathcal{G}_{\text{bow}}$.*

Proof. There is no directed path from $c_2.X$ to $c_3.Y$ in $\mathcal{G}_{\text{bow},\rho_\star}$ (Fig. E.5.1, top right). By do-calculus and Prop. 4.4, $P^{\rho_\star}(c_3.y \mid do(c_2.x)) = P^{\rho_\star}(c_3.y)$. By Prop. E.1, $P^{\rho_\star}(c_3.y)$ is relationally identifiable, and we are done. □

An identical derivation shows that $P^{\rho_\star}(c_3.y \mid do(c_2.x))$ is also relationally identifiable from $P(\mathbf{v}_\rho)$ and $\mathcal{G}_{\text{iv}}$.