
Nonparametric Identifiability of Causal Representations from Unknown Interventions

Julius von Kügelgen^{1,2}, Michel Besserve¹, Liang Wendong¹, Luigi Gresele¹, Armin Kekić¹,
Elias Bareinboim³, David M. Blei³, and Bernhard Schölkopf¹

¹Max Planck Institute for Intelligent Systems, Tübingen, Germany

²Department of Engineering, University of Cambridge, United Kingdom

³Columbia University, USA

Abstract

We study causal representation learning, the task of inferring latent causal variables and their causal relations from high-dimensional functions (“mixtures”) of the variables. Prior work relies on weak supervision, in the form of counterfactual pre- and post-intervention views or temporal structure; places restrictive assumptions, such as linearity, on the mixing function or latent causal model; or requires partial knowledge of the generative process, such as the causal graph or intervention targets. We instead consider the general setting in which both the causal model and the mixing function are *nonparametric*. The learning signal takes the form of multiple datasets, or environments, arising from *unknown interventions* in the underlying causal model. Our goal is to identify both the ground truth latents and their causal graph up to a set of ambiguities which we show to be irresolvable from interventional data. We study the fundamental setting of two causal variables and prove that the observational distribution and one perfect intervention per node suffice for identifiability, subject to a genericity condition. This condition rules out spurious solutions that involve fine-tuning of the intervened and observational distributions, mirroring similar conditions for nonlinear cause-effect inference. For an arbitrary number of variables, we show that at least one pair of distinct perfect interventional domains per node guarantees identifiability. Further, we demonstrate that the strengths of causal influences among the latent variables are preserved by all equivalent solutions, rendering the inferred representation appropriate for drawing causal conclusions from new data. Our study provides the first identifiability results for the general nonparametric setting with unknown interventions, and elucidates what is possible and impossible for causal representation learning without more direct supervision.

1 Introduction

Causal representation learning (CRL) seeks to describe high-dimensional, low-level observations through a small number of interpretable, causally-related latent variables [107, 109]. In doing so, its goal is to combine the strengths of classical causal inference with those of modern machine learning. A causal model represents an entire family of distributions arising from *interventions* on a system of variables [10, 94, 99]. This provides a principled way for reasoning about distribution shifts, which facilitates out-of-distribution generalization and planning [9, 63, 95, 101]. However, causal models require that most (or at least some) relevant causal variables are directly observed. While reasonable in domains such as economics [7], social [90] or biomedical science [48, 57], this assumption has challenged the application of causal methodology to complex and high-dimensional data [84]. Machine learning, on the other hand, has proven successful at learning useful “representations”—latent vectors generating the observables via some nonlinear map—of high-dimensional data such as images, video, or text [16, 20, 76, 105]. However, most methods rely on independent and identically distributed (i.i.d.) data and only extract *associational* information. As a consequence, they often fail under distribution shifts and do not generalize beyond the training distribution [106], as exhibited by their reliance on *spurious correlations* [13, 86] or their vulnerability to *adversarial examples* [40, 118].




Figure 1: *(Left)* **Data-Generating Process for Causal Representation Learning (CRL)**. Observations \mathbf{X} are generated by applying a nonlinear mixing function f to a set of causal latent variables $\mathbf{V} = \{V_1, \dots, V_n\}$, which are related through a structural causal model (SCM) with independent exogenous variables $\mathbf{U} = \{U_1, \dots, U_n\}$; illustration by Ana Martín Larrañaga. *(Right)* **Comparison to Structure Learning**. The causal direction between two unconfounded *observed* variables (top) is uniquely identified from a single intervention [37]; for CRL (bottom), this is not the case, as the nonlinear mixing introduces additional ambiguity due to spurious representations. Shaded nodes are observed, white ones unobserved, and interventions highlighted as red diamonds.

Identifiability. It has been argued that to address these shortcomings, we ought to learn representations endowed with causal model semantics. However, a major challenge to this goal is that *different representations can explain the same data equally well*. Strictly simpler representation learning tasks, such as disentanglement [14] or independent component analysis (ICA) [53], are already *non-identifiable* in general [54, 82]. Identifiability studies are thus required to characterize additional assumptions under which the desired latent variables can be provably recovered [4, 17, 21, 31, 42, 43, 45, 46, 51, 52, 55, 56, 65, 68, 69, 74, 75, 83, 89, 100, 111, 113, 126, 138]. In CRL, we need not only identify the latents, but also the causal graph encoding their relations. Even in the fully observed case, this task of causal discovery, or structure learning, is very challenging: the graph can only be recovered up to Markov equivalence [114] based only on (observational) i.i.d. data [115], meaning that the direction of some edges cannot be determined. For CRL, the task gets strictly harder. For instance, if the causal latents V_i in Fig. 1 (*Left*) form a valid representation, then replacing them by the *independent* exogenous U_i might be considered an equally valid alternative.

What sets causal representations apart? A crucial feature of causal variables is that they are the ones on which *interventions* are defined and whose relations we are interested in [94]. Causal discovery and CRL thus often rely on non-i.i.d. data linked to interventions on the underlying causal variables [59, 71]. Unless all variables are subject to intervention, however, some fundamental differences between the fully observed and representation learning settings in the level of ambiguity in the graph remain [116], as illustrated in Fig. 1 (*Right*). In a sense, the non-identifiabilities of representation and structure learning combine, and both need to be addressed in conjunction.

Problem Setting. We study the general *nonparametric* CRL problem (§ 2.1) in which both the causal mechanisms and the mixing function are completely unconstrained. Our goal (§ 2.2) is to identify the latent causal variables up to element-wise nonlinear rescaling and their graph up to isomorphism (Defn. 2.6). As motivated above, doing so without further supervision requires access to interventional data. To this end, we consider learning from heterogeneous data from multiple related domains, or *environments*, that arise from interventions in a shared underlying causal model (§ 2.3).

Contributions. Our main contributions are theoretical in nature (§ 3). First, we establish the minimality of the targeted equivalence class (Prop. 3.1) in the sense that its ambiguities cannot be resolved from interventional data. We then present our main identifiability results. For the case of two latent causal variables, we show that an observational environment and one for each perfect intervention on either variable suffice (Thm. 3.2)—provided that the intervened and unintervened mechanisms are not “fine-tuned” to each other, which we formalize in the form of a *genericity condition* (3.2). For any number of latent causal variables, we prove that access to pairs of environments corresponding to two distinct perfect interventions on each node guarantees identifiability (Thm. 3.4). We then question how to use or interpret causal representations (§ 4), and show that certain quantities, such as the strengths of causal influences among variables, are preserved by all equivalent solutions (Thm. 4.2). We sketch possible learning objectives (§ 5), and empirically investigate training different generative models (§ 6), finding that only those based on the correct causal structure attain the best fit and identify the ground truth. We conclude by discussing limitations and extensions of our work (§ 7).

Related Work on Multi-Environment CRL. Most closely related are recent identifiability studies which also leverage multiple environments arising from single node interventions [5, 22, 116, 121], thus mirroring in different interventional setups the result of Brehmer et al. [18] based on counterfactual, multi-view data, which is harder to obtain. Squires et al. [116] provide results for linear causal models and linear mixing; Ahuja et al. [5] consider nonlinear causal models and polynomial mixings, subject to additional constraints on the latent support [124]; and Varici et al. [121] employ a score-based approach for nonlinear causal models and linear mixing. The concurrent work of Buchholz et al. [22] extends the results of Squires et al. [116] to general nonlinear mixings and linear Gaussian causal models. Liu et al. [81] leverage recent advances in nonlinear ICA [55, 64] to identify a linear Gaussian causal model with context-dependent weights and nonlinear mixing from sufficiently diverse environments. A more extensive discussion of other related works and a detailed comparison of existing results with the present study are provided in Appx. A and Tab. 1.

2 Nonparametric Causal Representation Learning

In this section, we describe the considered problem setting and state our main assumptions. First, we specify the assumed data generating process (§ 2.1) and learning task in the form of a target identifiability class (§ 2.2). We then demonstrate the hardness of our task from i.i.d. data or imperfect interventions and use this to motivate a multi-environment approach with perfect interventions (§ 2.3).

Notation. We use upper-case X for random variables and lower-case x for their realizations. Bold uppercase \mathbf{X} denotes random vectors and lowercase \mathbf{x} their realizations. We assume throughout that all distributions P_X possess densities $p(x)$ with respect to (w.r.t.) the Lebesgue measure. We denote the pushforward of P_X through a measurable function f by $f_*(P_X)$, and write $[n] = \{1, \dots, n\}$.

2.1 Data Generating Process

The assumed data generating process consists of a latent causal model and a mixing function, see Fig. 1 (left). For the former, we build on the structural causal model (SCM) framework [94, 99].

Definition 2.1 (Latent SCM). Let $\mathbf{V} = \{V_1, \dots, V_n\}$ denote a set of causal “endogenous” variables, with each V_i taking values in \mathbb{R} , and let $\mathbf{U} = \{U_1, \dots, U_n\}$ denote a set of mutually independent “exogenous” random variables. The latent SCM consists of a set of structural equations

$$\{V_i := f_i(\mathbf{V}_{\text{pa}(i)}, U_i)\}_{i=1}^n. \quad (2.1)$$

where $\mathbf{V}_{\text{pa}(i)} \subseteq \mathbf{V} \setminus \{V_i\}$ are the direct causes, or causal parents, of V_i , and f_i are deterministic functions; and a fully factorized joint distribution P_U over the exogenous variables. The associated causal diagram G , a directed graph with vertices \mathbf{V} and edges $V_j \rightarrow V_i$ iff. $V_j \in \mathbf{V}_{\text{pa}(i)}$, is assumed acyclic.

By acyclicity, recursive substitution of the assignments in (2.1) yields the reduced form $\mathbf{V} = f_{\text{RF}}(\mathbf{U})$. The SCM thus induces a unique distribution P_V over the endogenous variables, given by the pushforward of P_U via (2.1), that is, $P_V = f_{\text{RF}*}(P_U)$. By construction, P_V is Markovian w.r.t. the causal graph G [94, 99], meaning that its density obeys the causal Markov factorization:

$$p(v_1, \dots, v_n) = \prod_{i=1}^n p_i(v_i \mid \mathbf{v}_{\text{pa}(i)}). \quad (2.2)$$

We place the following additional assumption on the distribution P_V induced by the SCM.

Assumption 2.2 (Faithfulness). The *only* conditional independence relations satisfied by P_V are those implied by $\{V_i \perp\!\!\!\perp \mathbf{V}_{\text{nd}(i)} \mid \mathbf{V}_{\text{pa}(i)}\}_{i \in [n]}$, where $\mathbf{V}_{\text{nd}(i)}$ denotes the non-descendants of V_i in G .

Asm. 2.2 ensures a one-to-one correspondence between (conditional) independence in P_V and graphical separation in G and is a standard assumption in causal discovery [114]. Faithfulness rules out cancellations of influences along different paths, which occurs with probability zero for random path-coefficients [120]. It can thus also be viewed as a minimality or genericity assumption.

In contrast to classical causal inference, we assume that both the exogenous variables \mathbf{U} and the endogenous causal variables \mathbf{V} are unobserved. Instead, we will only have access to d -dimensional nonlinear mixtures \mathbf{X} of \mathbf{V} . We therefore make the following additional assumption.

Assumption 2.3 (Known n). The number n of latent causal variables is known.

Next, we specify the relationship between the unobserved causal variables \mathbf{V} and the observed \mathbf{X} .

Definition 2.4 (Mixing function). The observations \mathbf{X} are deterministically generated from \mathbf{V} by applying a mixing function $f : \mathbb{R}^n \rightarrow \mathbb{R}^d$ to \mathbf{V} , that is, $\mathbf{X} := f(\mathbf{V})$.

The terminology and setting of a deterministic mixing is rooted in the nonlinear ICA literature [56]. In deep generative models, it is commonly relaxed by considering additive noise in Defn. 2.4 [64, 89]. For representation learning scenarios, we are particularly interested in the case $n \ll d$. To allow for recovery of \mathbf{V} from \mathbf{X} , we assume that f is invertible, which is a standard assumption for identifiability.

Assumption 2.5 (Diffeomorphic mixing). f is a diffeomorphism¹ onto its image $\text{Im}(f) = \mathcal{X} \subseteq \mathbb{R}^d$.

2.2 Learning Target: The CRL Identifiability Class \sim_{CRL}

Our goal is to infer the underlying latent causal variables $\mathbf{V} = f^{-1}(\mathbf{X})$ and their causal relations. We therefore consider the true unmixing function f^{-1} and the causal graph G our joint learning target: f^{-1} informs us how to map observations \mathbf{X} to causal variables \mathbf{V} , and G tells us how to factorise the implied joint $p(\mathbf{v})$ into the causal mechanisms $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ from (2.2). Given only observations of \mathbf{X} , this is a challenging task since neither \mathbf{V} nor G are directly observed or known a priori.

When is a candidate solution (h, G') that satisfies a given learning objective (such as maximizing the likelihood, possibly subject to additional constraints) guaranteed to match the ground truth (f^{-1}, G) ? This is the subject of identifiability studies and the main focus of our work. A statistical model $\mathcal{P} = \{p_\theta : \theta \in \Theta\}$ with parameter space Θ is identifiable if the mapping $\theta \mapsto p_\theta$ is injective [77]. For representation learning tasks, full identifiability is often not attainable, as there are some fundamental ambiguities that cannot be resolved. One therefore typically instead considers identifiability up to an appropriately chosen equivalence class in the model space [2, 64, 126].

For the assumed data generating process (§ 2.1), the order of the causal variables is arbitrary, since \mathbf{V} is unobserved. We can therefore assume without loss of generality (w.l.o.g.) that the V_i 's are partially ordered w.r.t. G , that is, $V_i \rightarrow V_j \implies i < j$.² Learning G thus reduces to inferring whether the edges $\{V_1, \dots, V_{i-1}\} \rightarrow V_i$ exist for $i = 2, \dots, n$. The only remaining permutation ambiguity arises from permutations π that preserve the partial order: for example, if G is given by $V_1 \rightarrow V_3 \leftarrow V_2$, the order of V_1 and V_2 cannot be uniquely determined without further assumptions. Moreover, the scaling of the causal variables is also arbitrary: any invertible element-wise transformation can be undone as part of f . We therefore define the desired identifiability class through the following equivalence relation.³

Definition 2.6 (\sim_{CRL} -identifiability). Let \mathcal{H} be a space of unmixing functions $h : \mathcal{X} \rightarrow \mathbb{R}^n$ and let \mathcal{G} be the space of DAGs over n vertices. Let \sim_{CRL} be the equivalence relation on $\mathcal{H} \times \mathcal{G}$ defined as

$$(h_1, G_1) \sim_{\text{CRL}} (h_2, G_2) \iff (h_2, G_2) = (\mathbf{P}_{\pi^{-1}} \circ \phi \circ h_1, \pi(G_1)) \quad (2.3)$$

for some element-wise diffeomorphism $\phi(\mathbf{v}) = (\phi_1(v_1), \dots, \phi_n(v_n))$ of \mathbb{R}^n and a permutation π of $[n]$ such that $\pi : G_1 \mapsto G_2$ is a graph isomorphism and \mathbf{P}_π the corresponding permutation matrix.

Remark 2.7. When G has no edges, any permutation is admissible and \sim_{CRL} reduces to the standard notion of identifiability up to permutation and element-wise reparametrisation of nonlinear ICA [56].

The ground truth (f^{-1}, G) is identified up to \sim_{CRL} by a given learning objective if any candidate solution (h, G') satisfies $(h, G') \sim_{\text{CRL}} (f^{-1}, G)$. We seek to discover suitable conditions that ensure this.

2.3 Multi-Environment Data

Given only a single dataset of i.i.d. observations from $P_{\mathbf{X}}$, there is no hope for \sim_{CRL} -identifiability. Even for observed \mathbf{V} (i.e., with $n = d$ and known $f = \text{id}$), G can only be identified up to Markov equivalence [114]. With unknown mixing f , the degree of observational non-identifiability gets even worse: for example, by using the reduced form of the SCM (§ 2.1) we can express \mathbf{X} in terms of the latent exogenous variables \mathbf{U} via $\mathbf{X} = f(\mathbf{V}) = f \circ f_{\text{RF}}(\mathbf{U})$ [107]. This gives rise to a ‘‘spurious ICA solution’’ $(f_{\text{RF}}^{-1} \circ f^{-1}, G_{\text{ICA}})$ where G_{ICA} denotes the empty graph with independent components. Due to the non-identifiability of nonlinear ICA [54, 82], however, we cannot even learn the composition $f \circ f_{\text{RF}}$, let alone separate it into its constituents f and f_{RF} to isolate the intermediate causal variables \mathbf{V} .

Motivated by these challenges to identifiability from i.i.d. data, we instead consider learning from *multiple environments* $e \in \mathcal{E}$. That is, we assume access to heterogenous data from multiple distinct distributions $P_{\mathbf{X}}^e$. Environments can arise, for example, from different experimental settings or

¹A diffeomorphism is a bijective function f such that both f and f^{-1} are continuously differentiable.

²If they were not in such order to begin with, we could apply an appropriate permutation $\tilde{\pi}$ to \mathbf{V} and incorporate the inverse permutation into the unknown mixing f without affecting $\mathbf{X} = f(\mathbf{V}) = (f \circ \tilde{\pi}^{-1})(\tilde{\pi}(\mathbf{V}))$ [116].

³ \sim_{CRL} satisfies symmetry and transitivity because permutations and element-wise functions commute.




Figure 2: **Multi-Environment Setup with Single-Node Perfect Interventions and Shared Mixing Function.** Illustration of the considered multi-environment setup for $n = 2$ causal variables $\mathbf{V} = \{V_1, V_2\}$ with graph G given by $V_1 \rightarrow V_2$, shared mixing function f , and environments $\mathcal{E} = \{e_0, e_1, e_2\}$, corresponding to the observational setting (e_0) and perfect stochastic interventions on V_1 (in e_1) and V_2 (in e_2). The learnt unmixing function, or encoder, is denoted by h , and the inferred latent representation by $\mathbf{Z} = h(\mathbf{X})$. The corresponding inferred latent distributions $Q_{\mathbf{Z}}^{e_i} = h_*(P_{\mathbf{X}}^{e_i})$ are Markovian w.r.t. the candidate graph G' (here, equal to G). Since the intervention targets are not known, they may in principle differ in $Q_{\mathbf{Z}}^e$ as shown here. However, as we prove in Thm. 3.2, such misalignment is only possible if a certain genericity condition (3.2) is violated.

correspond to broader contexts such as climate or time. Previous work has shown that this setting can, in principle, provide useful causal learning signals [5, 8, 19, 32, 34, 47, 50, 59, 70–72, 81, 88, 97, 98, 101, 103, 108, 116, 119, 121]. However, multi-environment data is not necessarily useful if the corresponding distributions $P_{\mathbf{X}}^e$ are allowed to differ in arbitrary ways. What makes this setting interesting is the assumption that certain parts of the causal generative process are shared across environments.

Here, we assume that all environments share the same invariant mixing function and underlying SCM, and that any distribution shifts arise from interventions on some of the causal mechanisms.⁴ General interventions can be modelled in SCMs by replacing some of the equations in (2.1) by new assignments $V_i := \tilde{f}_i(\mathbf{V}_{\text{pa}(i)}, \tilde{U}_i)$, resulting in “intervened” mechanisms $\tilde{p}_i(v_i | \mathbf{v}_{\text{pa}(i)})$ replacing the corresponding conditionals $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ in (2.2) [29, 36, 94]. We summarise this as follows.

Assumption 2.8 (Shared mixing and mechanisms). Each environment e shares the same mixing f ,

$$P_{\mathbf{X}}^e = f_*(P_{\mathbf{V}}^e)$$

and each $P_{\mathbf{V}}^e$ results from the same SCM through an intervention on a subset of mechanisms $\mathcal{I}^e \subseteq [n]$:

$$p^e(\mathbf{v}) = p^e(v_1, \dots, v_n) = \prod_{i \in \mathcal{I}^e} p_i^e(v_i | \mathbf{v}_{\text{pa}(i)}) \prod_{j \in [n] \setminus \mathcal{I}^e} p_j(v_j | \mathbf{v}_{\text{pa}(j)}).$$

Importantly, the intervention targets \mathcal{I}^e are not assumed to be known.

An intervention on some V_i that fully removes the influence from its parents $\mathbf{V}_{\text{pa}(i)}$ is referred to as *perfect*, otherwise as *imperfect*. It has been shown that imperfect interventions are generally insufficient for full identifiability [18], even in the linear case [116]. This is intuitive: if arbitrary imperfect interventions were allowed, including ones which preserve $f_i(\mathbf{V}_{\text{pa}(i)}, \cdot)$ and only replace U_i with some new \tilde{U}_i , then the spurious ICA solution $(f_{\text{RF}}^{-1} \circ f^{-1}, G_{\text{ICA}})$ should be indistinguishable from the ground truth. In line with prior work [5, 18, 116, 121], we therefore assume perfect interventions.

Assumption 2.9 (Perfect interventions). For all $e \in \mathcal{E}$ and $i \in \mathcal{I}^e$, we have $p^e(v_i | \mathbf{v}_{\text{pa}(i)}) = p^e(v_i)$.

Based on the principle of independent causal mechanisms [99, 108], the *sparse mechanism shift hypothesis* [97, 109] posits that distribution changes tend to manifest themselves in a sparse or local way in the causal factorization. In this spirit, we will assume single-node (“atomic”) interventions, $|\mathcal{I}^e| = 1$, for our main results, as also required for existing results [5, 18, 116, 121].

As motivated in § 2.2, given data from $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$ we consider candidate solutions of the form (h, G') where h is an unmixing function, or encoder, which maps observations \mathbf{X} to the inferred latents $\mathbf{Z} = h(\mathbf{X})$, and G' is causal graph capturing the relations among the Z_i . The corresponding distributions of the inferred latents are thus given by the push-forward

$$Q_{\mathbf{Z}}^e = h_*(P_{\mathbf{X}}^e) = (h \circ f)_* P_{\mathbf{V}}^e$$

The multi-environment setup with unknown single-node perfect interventions is illustrated in Fig. 2.

⁴One could also say that there exists only one environment and each dataset corresponds to an intervention.

3 Identifiability Theory

We start by showing that identifiability up to \sim_{CRL} is, in fact, the best we can hope for when learning from interventional data, without any more direct forms of supervision. To this end, we state the following result, which is presented more formally and proven in Appx. B.

Proposition 3.1 (Minimality of \sim_{CRL} ; informal). *Let \mathbf{Z} be any representation that is \sim_{CRL} equivalent to \mathbf{V} , with $G' = \pi(G)$ the associated DAG. Then for any intervention on $V_{\mathcal{I}} \subseteq \mathbf{V}$ in G , there exists an equally sparse intervention on $\mathbf{Z}_{\pi(\mathcal{I})} \subseteq \mathbf{Z}$ in G' inducing the same observed distribution on \mathbf{X} .*

We now present our identifiability results. We first study the most fundamental bivariate case with two latent causal variables V_1 and V_2 . This can loosely be seen as the CRL analogue of the widely studied cause-effect problem ($X \rightarrow Y$ or $Y \rightarrow X$?) in classical structure learning [87, 99].

Theorem 3.2 (Bivariate identifiability up to \sim_{CRL} from one perfect stochastic intervention per node). *Suppose that we have access to multiple environments $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$ generated as described in § 2 under Asms. 2.2, 2.5, 2.8 and 2.9 with $n = 2$. Let (h, G') be any candidate solution such that the inferred latent distributions $Q_{\mathbf{Z}}^e = h_*(P_{\mathbf{X}}^e)$ of $\mathbf{Z} = h(\mathbf{X})$ and the inferred mixing function h^{-1} satisfy the above assumptions w.r.t. the candidate causal graph G' . Assume additionally that*

- (A1) all densities p^e and q^e are continuously differentiable and fully supported on \mathbb{R}^n ;
- (A2) we have access to a known observational environment e_0 and one single node perfect intervention for each node, with unknown targets: there exist $n+1$ environments $\mathcal{E} = \{e_i\}_{i=0}^n$ such that $\mathcal{I}^{e_0} = \emptyset$ and for each $i \in [n]$ we have $\mathcal{I}^{e_i} = \{\pi(i)\}$ for an unknown permutation π of $[n]$;
- (A3) for all $i \in [n]$, the intervened mechanisms $\tilde{p}_i(v_i)$ differ from the corresponding base mechanisms $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ everywhere, in the sense that

$$\forall \mathbf{v} : \quad \frac{\partial}{\partial v_i} \frac{\tilde{p}_i(v_i)}{p_i(v_i | \mathbf{v}_{\text{pa}(i)})} \neq 0; \quad (3.1)$$

- (A4) (“**genericity**”) the base and intervened mechanisms are not fine-tuned to each other, in the sense that there exists a continuous function $\varphi : \mathbb{R}^+ \rightarrow \mathbb{R}$ for which

$$\mathbb{E}_{\mathbf{v} \sim P_{\mathbf{V}}^{e_0}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 | v_1)} \right) \right] \neq \mathbb{E}_{\mathbf{v} \sim P_{\mathbf{V}}^{e_1}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 | v_1)} \right) \right] \quad (3.2)$$

Then the ground truth is identified in the sense of Defn. 2.6, that is, $(f^{-1}, G) \sim_{\text{CRL}} (h, G')$.

Proof sketch (full proof in Appx. C.2). Consider $V_1 \rightarrow V_2$ (the proof for $V_1 \nrightarrow V_2$ is similar). Let $\psi = f^{-1} \circ h^{-1} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $\mathbf{V} = \psi(\mathbf{Z})$.⁵ By the change of variable formula, for all \mathbf{z}

$$q^e(\mathbf{z}) = p^e(\psi(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (3.3)$$

where $(\mathbf{J}_{\psi}(\mathbf{z}))_{ij} = \frac{\partial \psi_i}{\partial z_j}(\mathbf{z})$ denotes the Jacobian of ψ . We consider two separate cases, depending on whether the intervention targets in q^{e_i} for $e_i \in \{e_1, e_2\}$ match those in p^{e_i} (Case 1) or not (Case 2).

Case 1: Aligned Intervention Targets. According to Asm. 2.8 and (A2), (3.3) applied to the known observational environment e_0 and the interventional environments e_1, e_2 leads to the system of equations:

$$q_1(z_1)q_2(z_2 | z_{\text{pa}(2;G')}) = p_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_1(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (3.4)$$

$$\tilde{q}_1(z_1)q_2(z_2 | z_{\text{pa}(2;G')}) = \tilde{p}_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_1(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (3.5)$$

$$q_1(z_1)\tilde{q}_2(z_2) = p_1(\psi_1(\mathbf{z})) \tilde{p}_2(\psi_2(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (3.6)$$

where $z_{\text{pa}(2;G')}$ denotes the parents of z_2 in G' . Taking quotients of (3.5) and (3.4) yields

$$\frac{\tilde{q}_1}{q_1}(z_1) = \frac{\tilde{p}_1}{p_1}(\psi_1(\mathbf{z})) \xrightarrow{\frac{\partial}{\partial z_2}} 0 = \left(\frac{\tilde{p}_1}{p_1} \right)'(\psi_1(\mathbf{z})) \frac{\partial \psi_1}{\partial z_2}(\mathbf{z}) \xrightarrow{(A3)} \frac{\partial \psi_1}{\partial z_2}(\mathbf{z}) = 0. \quad (3.7)$$

Thus $V_1 = \psi_1(Z_1)$ and $q_1(z_1) = p_1(\psi_1(z_1)) \left| \frac{\partial \psi_1}{\partial z_1}(z_1) \right|$. Substitution into (3.6) yields

$$\tilde{q}_2(z_2) = \tilde{p}_2(\psi_2(z_1, z_2)) \left| \frac{\partial \psi_2}{\partial z_2}(z_1, z_2) \right| \quad (3.8)$$

⁵By Asm. 2.5, f, h , and thus also $h \circ f$ are diffeomorphisms. Hence, ψ is well-defined and also diffeomorphic.

where we have used that, according to (3.7), \mathbf{J}_ψ is triangular. According to (3.8), for all z_1 , the mapping $z_2 \mapsto \psi_2(z_1, z_2)$ is measure preserving for \tilde{q}_2 and \tilde{p}_2 . By Lemma C.1 [18, § A.2, Lemma 2], it follows that ψ_2 must be constant in z_1 .⁶ Hence, ψ is an element-wise function. Finally, $G = G'$ follows from faithfulness (Asm. 2.2), for otherwise $(V_1, V_2) = (\psi_1(Z_1), \psi_2(Z_2))$ would be independent.

Case 2: Misaligned Intervention Targets. If $G' \neq G$, a similar argument to Case 1 (with the roles of z_1 and z_2 interchanged) also yields a contradiction to faithfulness. This leaves $G = G'$. Writing down the system of equations similar to (3.4)–(3.6), and then taking ratios of e_1 and e_2 with e_0 yields

$$\frac{\tilde{q}_1(z_1)}{q_1(z_1)} = \frac{\tilde{p}_2(\psi_2(\mathbf{z}))}{p_2(\psi_2(\mathbf{z}) \mid \psi_1(\mathbf{z}))} \quad \text{and} \quad \frac{\tilde{q}_2(z_2)}{q_2(z_2 \mid z_1)} = \frac{\tilde{p}_1(\psi_1(\mathbf{z}))}{p_1(\psi_1(\mathbf{z}))}. \quad (3.9)$$

These conditions highlight the misalignment in intervention targets (see Fig. 2). Unlike in Case 1, they do not directly imply that some elements of \mathbf{J}_ψ need to be zero, that is \mathbf{Z} can be arbitrarily mixed w.r.t. \mathbf{V} . However, (3.9) imposes constraints on the form of ψ that, by exploiting the invariance of q_1 across e_0 and e_1 while p_1 changes to \tilde{p}_1 , can ultimately be shown to only be satisfied if the two expectations in (3.2) are equal for all continuous φ . However, such fine-tuning is ruled out by (A4). \square

Remark 3.3. The main difficulty of the proof is that (3.9) may, in principle, hold when $(p, \tilde{p}, q, \tilde{q})$ and ψ are completely unconstrained. This does not arise in prior work if the intervention targets are known [79, 80, 125] (Case 1), or the densities or mixing are parametrically constrained [3, 116, 121].

On the Genericity Condition (A4). The condition in (3.2) contrasts expectations of the same quantity w.r.t. the observational distribution $P_{\mathbf{V}}^{e_0}$ and the interventional distribution $P_{\mathbf{V}}^{e_1}$ corresponding to an intervention on V_1 that turns p_1 into \tilde{p}_1 . The shared argument, on the other hand, is a function of the ratio between the intervened mechanism \tilde{p}_2 and its base mechanism p_2 . While the two expectations are always equal for linear φ , other choices imply non-trivial constraints. For instance, $\varphi(y) = y^2$ yields

$$\int (\tilde{p}_1(v_1) - p_1(v_1)) \int \frac{\tilde{p}_2(v_2)^2}{p_2(v_2 \mid v_1)} dv_2 dv_1 \neq 0.$$

Since $p_1 \neq \tilde{p}_1$ by assumption (A3), we argue that (A4) should generally hold for randomly chosen $(p_1, p_2, \tilde{p}_1, \tilde{p}_2)$ and can only be violated if they are fine-tuned to each other. It can thus be viewed as encoding some notion of genericity—in line with the principle of independent causal mechanisms [15, 43, 61, 99, 108], but also involving the intervened mechanisms. Interestingly, related genericity conditions also arise in the study of nonlinear cause-effect inference from observational data, where identifiability is often obtained up to a set of pathological (“fine-tuned”) spurious solutions satisfying a partial differential equation involving the original mechanisms [49, 58, 117, 134]. Further, we note that φ in (3.2) can also be thought of as a *witness function of genericity*, similar to witness functions in kernel-based two sample and independence testing [44].

Next, we provide our identifiability result for an arbitrary number of causal variables.

Theorem 3.4 (Identifiability up to \sim_{CRL} from two paired perfect stochastic interventions per node). *Suppose that we have access to multiple environments $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$ generated as described in § 2 under Asms. 2.2, 2.3, 2.5, 2.8 and 2.9. Let (h, G') be any candidate solution such that the inferred latent distributions $Q_{\mathbf{Z}}^e = h_*(P_{\mathbf{X}}^e)$ of $\mathbf{Z} = h(\mathbf{X})$ and the inferred mixing function h^{-1} satisfy the above assumptions w.r.t. the candidate causal graph G' . Assume additionally that*

- (A1) all densities p^e and q^e are continuously differentiable and fully supported on \mathbb{R}^n ;
- (A2') we have access to at least one pair of single-node perfect interventions per node, with unknown targets: there exist $m \geq n$ known pairs of environments $\mathcal{E} = \{(e_j, e'_j)\}_{j=1}^m$ such that for each $i \in [n]$ there exists some unknown $j \in [m]$ for which $\mathcal{I}^{e_j} = \mathcal{I}^{e'_j} = \{i\}$;
- (A3') for all $i \in [n]$, the intervened mechanisms $\tilde{p}_i(v_i)$ and $\tilde{\tilde{p}}_i(v_i)$ differ everywhere, in the sense that

$$\forall v_i : \left(\frac{\tilde{\tilde{p}}_i}{\tilde{p}_i} \right)'(v_i) \neq 0; \quad (3.10)$$

Then the ground truth is identified in the sense of Defn. 2.6, that is, $(f^{-1}, G) \sim_{\text{CRL}} (h, G')$.

⁶This step is where the assumption of perfect interventions (Asm. 2.9) is leveraged: the conclusion would not hold for arbitrary imperfect interventions for which (3.8) would involve $\tilde{q}_2(z_2 \mid z_1)$ and $p_2(\psi_2(z_1, z_2) \mid \psi_1(z_1))$.

Proof sketch (full proof in Appx. C.3). By considering ratios between e_j and e'_j , taking partial derivatives w.r.t. z_l , and using assumptions (A3'), we can identify a subset $\mathcal{E}_n \subseteq \mathcal{E}$ of exactly n pairs of environments corresponding to distinct intervention targets in p (for otherwise ψ cannot be invertible). For $(e_i, e'_i) \in \mathcal{E}_n$, w.l.o.g. fix the intervention targets in p to $\mathcal{I}^{e_i} = \mathcal{I}^{e'_i} = \{i\}$ and let π be a permutation of $[n]$ such that $\pi(i)$ denotes the inferred intervention target in q that by (A2') is shared across (e_i, e'_i) . By the same argument as before, we must have $V_i = \psi_i(Z_{\pi(i)})$, that is, ψ is a permutation composed with an element-wise function. It remains to show that π is a graph isomorphism, that is, $V_i \rightarrow V_j$ in $G \iff Z_{\pi(i)} \rightarrow Z_{\pi(j)}$ in G' . We prove \implies ; the other direction is analogous. Suppose for a contradiction that there exist (i, j) such that $V_i \rightarrow V_j$ in G , but $Z_{\pi(i)} \not\rightarrow Z_{\pi(j)}$ in G' . Consider e_i in which there are perfect interventions on $Z_{\pi(i)}$ and \tilde{V}_i . For $Z_{\pi(k)} \in \mathbf{Z}_{\text{pa}(\pi(j); G')}$, let $\tilde{V}_k = \psi_k(Z_{\pi(k)})$ and denote $\tilde{\mathbf{V}} = \cup_k \tilde{V}_k \subset \mathbf{V} \setminus \{V_i, V_j\}$. Since $Z_{\pi(i)}$ and $Z_{\pi(j)}$ are d-separated by $\mathbf{Z}_{\text{pa}(\pi(j); G')}$ in the post-intervention graph $G'_{\tilde{\mathbf{Z}}_{\pi(i)}}$ with arrows pointing into $Z_{\pi(i)}$ removed [94], it follows by Markovianity of q w.r.t. G' that $Z_{\pi(i)} \perp\!\!\!\perp Z_{\pi(j)} \mid \mathbf{Z}_{\text{pa}(\pi(j); G')}$ in Q^{e_i} . By applying the corresponding diffeomorphic functions ψ_i , it follows from Lemma C.2 in Appx. C.1 that $V_i \perp\!\!\!\perp V_j \mid \tilde{\mathbf{V}}$ in $P_{\tilde{\mathbf{V}}}^{e_i}$. This violates faithfulness (Asm. 2.2) of $P_{\mathbf{V}}$ to G since V_i and V_j are d-connected in $G_{\tilde{\mathbf{V}}}$. \square

(A2') states that we know that a *pair of datasets* corresponds to two distinct interventions on the same underlying variable, even though we may not know the exact target of the intervention. This situation could arise, for example, if both datasets are collected under the same experimental setup but with varying experimental parameters. We stress that this is different from data consisting of *pairs of views* $(\mathbf{x}, \tilde{\mathbf{x}})$ sharing the values of some variables, which is *counterfactual* in nature [4, 18, 122]. One of the main challenges for our analysis (compared to a counterfactual multi-view setting) thus stems from the lack of correspondences across observations from different datasets. We also stress that a purely observational environment is not needed in this case, cf. (A2) in Thm. 3.2.

(A3') states that the intervention mechanisms are distinct in that their ratio is strictly monotonic, similar to (3.1) in (A3). This is a slightly stronger version of the assumption of *interventional discrepancy* proposed by Wendong et al. [125],⁷ which has been shown to be necessary for identifiability even if the graph G is known. For Gaussian $\tilde{p}_i, \tilde{p}'_i$, (A3') is satisfied, for example, by a shift in mean. In the proofs, this is used to show that ψ must be an element-wise function, see (3.7). Intuitively, if $\tilde{p}_i = \tilde{p}'_i$ in some open set for more than one i , then the underlying variables can be nonlinearly mixed by a measure-preserving automorphism within this set without affecting the observed distributions [125].

4 Interpreting Causal Representations

Suppose that we succeed in identifying \mathbf{V} and G up to \sim_{CRL} (Defn. 2.6). How can we use or interpret such a causal representation? Since the scale of the variables is arbitrary (§ 2.2), we clearly cannot predict the exact outcomes of interventions. We therefore seek causal quantities that are preserved by the irresolvable ambiguities of \sim_{CRL} . A prime candidate for this are interventional causal notions defined in terms of information theoretic quantities [28] and in particular the KL divergence D_{KL} .

Definition 4.1 (Causal influence; based on Defn. 2 of Janzing et al. [62]). Let $P_{\mathbf{V}}$ be Markovian w.r.t. a DAG G with vertices \mathbf{V} . For any $V_i \rightarrow V_j$ in G , the *causal influence of V_i on V_j* is given by

$$\mathfrak{C}_{i \rightarrow j}^{P_{\mathbf{V}}} := D_{\text{KL}}(P_{\mathbf{V}} \parallel P_{\mathbf{V}}^{i \rightarrow j}), \quad \text{where } p_j^{i \rightarrow j}(v_j \mid \mathbf{v}_{\text{pa}(j) \setminus \{i\}}) = \int p_j(v_j \mid \mathbf{v}_{\text{pa}(j)}) p_i(v_i) dv_i$$

and $P_{\mathbf{V}}^{i \rightarrow j}$ is the interventional distribution arising from replacing the j^{th} mechanism by $p_j^{i \rightarrow j}$.⁸

The following result, proven in Appx. C.4, states that the causal influences are invariant to reparametrisation and equivariant to permutations, the two irresolvable ambiguities of the \sim_{CRL} equivalence class.

Theorem 4.2 (Preservation of causal influences under \sim_{CRL}). Let $P_{\mathbf{V}}$ be Markovian w.r.t. G , let π be a graph isomorphism of G , and let ϕ be an element-wise diffeomorphism. Let $\mathbf{Z} = \mathbf{P}_{\pi^{-1}} \circ \phi(\mathbf{V})$ and denote its induced distribution by $Q_{\mathbf{Z}}$. Then for any $V_i \rightarrow V_j$ in G we have $\mathfrak{C}_{i \rightarrow j}^{P_{\mathbf{V}}} = \mathfrak{C}_{\pi(i) \rightarrow \pi(j)}^{Q_{\mathbf{Z}}}$.

Thm. 4.2 implies that the strength of causal relations among variables in the inferred graph carry meaning. They can thus be used to uncover changes to the latent causal mechanisms underlying different experimental datasets, for example, to gain scientific insights when combined with domain knowledge.

⁷Cf. the *interventional regularity* assumption of Varici et al. [121, Asm. 2] which instead considers partial derivatives w.r.t. the parents and is related to c-faithfulness [59, Defn. 7].

⁸Intuitively, this intervention captures the process of removing the arrow $V_i \rightarrow V_j$ in G and ‘‘feeding’’ the conditional $p(v_j \mid \mathbf{v}_{\text{pa}(j)})$ with an independent copy of V_i , distributed according to its marginal, see [62] for details.

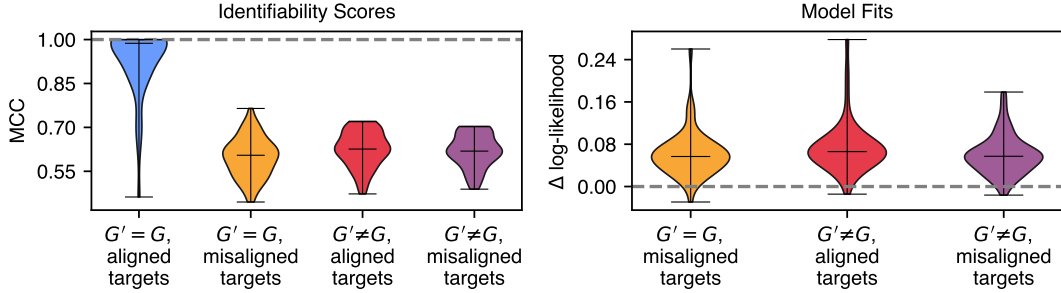


Figure 3: **Empirical Comparison of Correctly and Incorrectly Specified Normalizing Flow-Based Models.** For $n = 2$ latent causal variables with graph G given by $V_1 \rightarrow V_2$, we compare a generative model based on the correct causal graph $G' = G$ and intervention targets (blue) to other generative models assuming the wrong graph $G' \neq G$ or misaligned intervention targets (yellow, red, purple). We show mean correlation coefficients (MCCs) between the learned and ground truth latents (*Left*) and the difference in validation model log-likelihood between the well-specified and misspecified models (*Right*). Each violin plot is based on 50 different ground truth data generating processes; the horizontal lines indicate the minimum, median and maximum values.

5 Learning Objectives

While our main focus is on studying identifiability, our theoretical insights also suggest approaches to learning causal representations from finite interventional datasets sampled from $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$. The main idea is to fit the data in a way that preserves the sparsity of interventions [97, 109] by employing the same (un)mixing function and sharing mechanisms across environments (Asm. 2.8). We sketch two approaches, which, according to our theory, should both asymptotically identify the ground truth up to \sim_{CRL} if the set of available environments \mathcal{E} is sufficiently diverse (and the other assumptions hold).

- *Autoencoder Framework:* Jointly learn an encoder h , a graph G' , and intervention targets \mathcal{I}^e such that the encoded latents $\mathbf{Z} = h(\mathbf{X})$ can be used to reconstruct the observed \mathbf{X} across all environments, while ensuring all but the intervened mechanisms are shared. Using E as an environment indicator, the latter corresponds to the constraint $Z_i \perp\!\!\!\perp E \mid \mathbf{Z}_{\text{pa}(i; G')}$ for $i \notin \mathcal{I}^e$, implementable, for example, through a suitable conditional independence test [39, 93, 135].
- *Generative Modelling Approach:* Fit a base generative model (G', p, f) and intervention models $(\mathcal{I}^e, \{p_i^e\}_{i \in \mathcal{I}^e})_{e \in \mathcal{E}}$ by maximizing the likelihood of the multi-environment data. For example, given a candidate graph G' and candidate intervention targets $\{\mathcal{I}^e\}_{e \in \mathcal{E}}$, learn the base and intervened mechanisms and mixing; then pick the graph and intervention targets that achieve the best fit.

6 Experiments

Setup. We experimentally pursue the second, generative approach. Specifically, we model the mixing function generating \mathbf{X} from \mathbf{Z} as a *normalizing flow* [30, 92] with different environment-specific base distributions, determined by the underlying causal graph, intervention targets, and (learned) base and intervened mechanisms. Here, we focus on the bivariate case with two ground-truth causal latent variables $V_1 \rightarrow V_2$. According to Thm. 3.2, this setting should be identifiable from three environments: an observational one and one perfect intervention on each of V_1 and V_2 . Our goal is to verify this empirically in light of finite data and optimization issues. We fix the observation dimension to $d = 2$ to facilitate exact likelihood training of the normalizing flows, and fit a separate generative model for each choice of graph and intervention targets.⁹ The base and intervened mechanisms are linear Gaussian and the mixing function a three-layer MLP, see Appx. D.1 for implementation details.

Results. The results are summarized in Fig. 3. Our main findings are two-fold: First, we observe that, in the majority of cases, the well-specified model attains the highest held-out log-likelihood, as shown in Fig. 3 (*Right*). This suggests that the likelihood of otherwise comparable generative models can act as a reliable criterion to select the correct causal graph and intervention targets. Second, we find that the ground truth latent causal variables are approximately identified up to element-wise rescaling (MCC values close to one) by the correctly specified model and not by any other model, as shown in Fig. 3 (*Left*). This indicates that recovering the correct graph and targets is not only sufficient but also necessary for reliable identification of the causal representation. Taken together, these findings are consistent with our notion of \sim_{CRL} -identifiability (Defn. 2.6) and Thm. 3.2.

⁹This can be viewed as a nested approach in which the inner loop corresponds to the method for Causal Component Analysis (CauCA) of Wendong et al. [125], and the outer loop to a search over G' and $\{\mathcal{I}^e\}_{e \in \mathcal{E}}$. Code to reproduce our experiments is available at: <https://github.com/akekic/causal-component-analysis>.

In Appx. D.2, we perform an additional experiment with $n = 3$ variables, nonlinear latent SCMs, fixed causal order without graph search, and one interventional environment per node, thus assaying violations of assumption (A2') in the context of Thm. 3.4. Generally, we find that a correct choice of intervention targets can be selected based on model likelihood, leading to approximate identification of the causal variables, even when only the causal order is given, see Fig. 4 and Appx. D.2 for details.

7 Conclusions and Discussion

The world is full of domain shifts and different environments. Often, we cannot pin down what exactly differs between two domains, but it may reasonably be modelled as a change in some causal mechanisms. While prior work relied on harder-to-obtain counterfactual data or parametric constraints for identifiability, our study demonstrates that interventional data can be sufficient—even in the nonparametric case where the spaces of mixing functions and mechanisms are infinite dimensional. Our results can be considered a step towards justifying the use of expressive machine learning methods for learning interpretable causal representations from high-dimensional experimental data in situations where parametric assumptions are undesirable, e.g., for complex systems in physics, biology, or medicine.

Weaker Notions of Identifiability. In the present work, we have focused on the strongest notion of identifiability that is achievable in a nonparametric setting (Defn. 2.6). However, subject to the available data and assumptions, identifiability up to \sim_{CRL} will not always be possible. In this case, weaker notions of identifiability are of interest. For example, we may not be able to uniquely recover variables that are not targeted by interventions [5, 79, 122], or only recover groups of variables up to (non-)linear mixing [5, 121, 122] and the graph up to transitive closure [116] if interventions are imperfect, or soft. A precise characterization of weaker notions of *nonparametric* identifiability from different types of *interventions* (cf. [73] for a temporal, semi-parametric setting) is an interesting direction for future work.

Known vs. Unknown Intervention Targets. When intervention targets can be considered known appears to be a more nuanced concept in CRL than in a fully observed setting, see also [125, §E] for an extended discussion. Recall that we assume w.l.o.g. that $V_1 \preceq \dots \preceq V_n$ and only consider graphs respecting this ordering (§ 2.2), see also [116, Remark 1]. The intervention targets are then unknown w.r.t. the pre-imposed causal ordering. This is a key aspect that makes our setting more realistic, but also substantially complicates the analysis (see (3.9) and Remark 3.3). Similar to [121], for Thm. 3.2 we require a set of *exactly* n environments (one intervention on each node).¹⁰ However, we relax this requirement to mere coverage (“at least 1”) in Thm. 3.4 as shown for linear causal models in [22, 116].

Identifiability From One Intervention per Node for Any n . We conjecture that Thm. 3.2 can be generalized to $n > 2$, subject to a suitably adjusted set of genericity conditions involving several intervened and base mechanisms, akin to (3.2) in the bivariate case. The main challenge to such a generalization appears to be combinatorial, as there are $n! - 1$ ways of misaligning intervention targets across p and q . In Thm. 3.4, we sidestep this issue by assuming pairs of environments. Thus, while two single-node perfect interventions are sufficient, we do not believe this to be necessary.

On the Assumption of Known n and Its Relation to Markovianity. Asm. 2.3 relates to the notions of *causal sufficiency* or *Markovianity* in classical causal inference, which correspond to the assumption of independent U_i in Defn. 2.1, implying the causal Markov factorization (2.2) [94, 99, 114]. With *unobserved V* and *unknown n* , the notion of “unobserved confounders” gets blurred, since one can, in principle, always construct a causally sufficient system by increasing n and adding any causes of two or more endogenous variables to \mathbf{V} . Asm. 2.3 then states that the minimum number of variables required to do so is known.¹¹ However, this can lead to very large systems, which may in turn challenge the assumption of an invertible mixing (Asm. 2.5). Extensions of our analysis to unknown n , non-Markovian, or non-invertible models constitute an interesting direction for future investigations.

Practicality. The method explored in § 6 requires searching over graphs and intervention targets, which gets intractable even for moderate n . Simultaneously learning an (un)mixing function, causal graph, intervention targets, and mechanisms is challenging. Further work is needed to make methods for nonparametric CRL from multi-environment data more practical, e.g., by exploring the proposed autoencoder framework with a continuous parametrisation of graph [136, 137] and targets [60].

¹⁰By dropping the assumption of a fixed ordering and considering all DAGs, in this case one could also call V_i the variable intervened upon in environment i , and then consider the targets “known” in this sense.

¹¹Techniques for estimating the intrinsic dimensionality n of the observation manifold $\mathcal{X} \subseteq \mathbb{R}^d$ or methods rooted in Bayesian nonparametrics could provide a means of relaxing this assumption.

Acknowledgments and Disclosure of Funding

We thank Simon Buchholz, Jiaqi Zhang, and the anonymous reviewers for helpful discussions, comments, and suggestions.

This work was supported by the Tübingen AI Center and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy, EXC number 2064/1, Project number 390727645. L.G. was supported by the VideoPredict project, FKZ: 01IS21088.

References

- [1] J. Adams, N. Hansen, and K. Zhang. Identification of partially observed linear causal models: Graphical conditions for the non-gaussian and heterogeneous cases. *Advances in Neural Information Processing Systems*, 34:22822–22833, 2021. [Cited on p. 20.]
- [2] K. Ahuja, J. Hartford, and Y. Bengio. Properties from mechanisms: an equivariance perspective on identifiable representation learning. In *International Conference on Learning Representations*, 2021. [Cited on p. 4 and 20.]
- [3] K. Ahuja, D. Mahajan, V. Syrgkanis, and I. Mitliagkas. Towards efficient representation identification in supervised learning. In *First Conference on Causal Learning and Reasoning*, 2021. [Cited on p. 7 and 20.]
- [4] K. Ahuja, J. S. Hartford, and Y. Bengio. Weakly supervised representation learning with sparse perturbations. In *Advances in Neural Information Processing Systems*, volume 35, pages 15516–15528, 2022. [Cited on p. 2 and 8.]
- [5] K. Ahuja, D. Mahajan, Y. Wang, and Y. Bengio. Interventional causal representation learning. In *International Conference on Machine Learning*, pages 372–407. PMLR, 2023. [Cited on p. 3, 5, 10, and 21.]
- [6] T. V. Anand, A. H. Ribeiro, J. Tian, and E. Bareinboim. Effect identification in cluster causal diagrams. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, 2023. [Cited on p. 20.]
- [7] J. D. Angrist and J.-S. Pischke. *Mostly harmless econometrics: An empiricist’s companion*. Princeton University Press, 2009. [Cited on p. 1.]
- [8] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz. Invariant risk minimization. *arXiv preprint 1907.02893*, 2019. [Cited on p. 5 and 20.]
- [9] E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 2016. [Cited on p. 1.]
- [10] E. Bareinboim, J. D. Correa, D. Ibeling, and T. Icard. On Pearl’s hierarchy and the foundations of causal inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, page 507–556. Association for Computing Machinery, 2022. [Cited on p. 1 and 20.]
- [11] S. Beckers and J. Y. Halpern. Abstracting causal models. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 33, pages 2678–2685, 2019. [Cited on p. 20.]
- [12] S. Beckers, F. Eberhardt, and J. Y. Halpern. Approximate causal abstractions. In *Uncertainty in Artificial Intelligence*, pages 606–615. PMLR, 2020. [Cited on p. 20.]
- [13] S. Beery, G. Van Horn, and P. Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 456–473, 2018. [Cited on p. 1.]
- [14] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. [Cited on p. 2.]
- [15] M. Besserve, N. Shajarisales, B. Schölkopf, and D. Janzing. Group invariance principles for causal generative models. In *International Conference on Artificial Intelligence and Statistics*, pages 557–565. PMLR, 2018. [Cited on p. 7.]

- [16] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003. [Cited on p. 1.]
- [17] J. Brady, R. S. Zimmermann, Y. Sharma, B. Schölkopf, J. von Kügelgen, and W. Brendel. Provably learning object-centric representations. In *International Conference on Machine Learning*, 2023. [Cited on p. 2.]
- [18] J. Brehmer, P. De Haan, P. Lippe, and T. Cohen. Weakly supervised causal representation learning. In *Advances in Neural Information Processing Systems*, volume 35, pages 38319–38331, 2022. [Cited on p. 3, 5, 7, 8, 20, 21, 23, and 25.]
- [19] P. Brouillard, S. Lachapelle, A. Lacoste, S. Lacoste-Julien, and A. Drouin. Differentiable causal discovery from interventional data. In *Advances in Neural Information Processing Systems*, volume 33, pages 21865–21877, 2020. [Cited on p. 5.]
- [20] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020. [Cited on p. 1.]
- [21] S. Buchholz, M. Besserve, and B. Schölkopf. Function classes for identifiable nonlinear independent component analysis. In *Advances in Neural Information Processing Systems*, 2022. [Cited on p. 2.]
- [22] S. Buchholz, G. Rajendran, E. Rosenfeld, B. Aragam, B. Schölkopf, and P. Ravikumar. Learning linear causal representations from interventions under general nonlinear mixing. In *Advances in Neural Information Processing Systems*, 2023. [Cited on p. 3, 10, and 21.]
- [23] R. Cai, F. Xie, C. Glymour, Z. Hao, and K. Zhang. Triad constraints for learning causal structure of latent variables. In *Advances in Neural Information Processing Systems*, volume 32, pages 12883–12892, 2019. [Cited on p. 20 and 21.]
- [24] K. Chalupka, P. Perona, and F. Eberhardt. Visual causal feature learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 181–190. AUAI Press, 2015. [Cited on p. 20.]
- [25] K. Chalupka, F. Eberhardt, and P. Perona. Multi-level cause-effect systems. In *Artificial Intelligence and Statistics*, pages 361–369. PMLR, 2016.
- [26] K. Chalupka, F. Eberhardt, and P. Perona. Causal feature learning: an overview. *Behaviormetrika*, 44(1):137–164, 2017. [Cited on p. 20.]
- [27] P. Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994. [Cited on p. 20.]
- [28] T. M. Cover. *Elements of information theory*. John Wiley & Sons, 1999. [Cited on p. 8.]
- [29] A. P. Dawid. Influence diagrams for causal modelling and inference. *International Statistical Review*, 70(2):161–189, 2002. [Cited on p. 5.]
- [30] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios. Neural spline flows. In *Advances in Neural Information Processing Systems*, volume 32, pages 7511–7522, 2019. [Cited on p. 9 and 33.]
- [31] C. Eastwood, A. L. Nicolicioiu, J. Von Kügelgen, A. Kekić, F. Träuble, A. Dittadi, and B. Schölkopf. DCI-ES: An extended disentanglement framework with connections to identifiability. In *The Eleventh International Conference on Learning Representations*, 2022. [Cited on p. 2.]
- [32] C. Eastwood, A. Robey, S. Singh, J. von Kügelgen, H. Hassani, G. J. Pappas, and B. Schölkopf. Probable domain generalization via quantile risk minimization. In *Advances in Neural Information Processing Systems*, 2022. [Cited on p. 5 and 20.]
- [33] C. Eastwood, S. Singh, A. L. Nicolicioiu, M. Vlastelica, J. von Kügelgen, and B. Schölkopf. Spuriousity didn’t kill the classifier: Using invariant predictions to harness spurious features. In *Advances in Neural Information Processing Systems*, 2023. [Cited on p. 20.]

- [34] D. Eaton and K. Murphy. Exact bayesian structure learning from uncertain interventions. In *Artificial Intelligence and Statistics*, pages 107–114. PMLR, 2007. [Cited on p. 5.]
- [35] F. Eberhardt. Green and grue causal variables. *Synthese*, 193(4):1029–1046, 2016. [Cited on p. 20.]
- [36] F. Eberhardt and R. Scheines. Interventions and causal inference. *Philosophy of science*, 74(5):981–995, 2007. [Cited on p. 5.]
- [37] F. Eberhardt, C. Glymour, and R. Scheines. N-1 experiments suffice to determine the causal relations among n variables. *Innovations in Machine Learning: Theory and Applications*, pages 97–112, 2006. [Cited on p. 2.]
- [38] J. Eriksson and V. Koivunen. Identifiability, separability, and uniqueness of linear ICA models. *IEEE Signal Processing Letters*, 11(7):601–604, 2004. [Cited on p. 20.]
- [39] K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems*, volume 20, pages 489–496. Curran Associates, Inc., 2007. [Cited on p. 9.]
- [40] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations*, 2015. [Cited on p. 1.]
- [41] L. Gresele, G. Fissore, A. Javaloy, B. Schölkopf, and A. Hyvarinen. Relative gradient optimization of the Jacobian term in unsupervised deep learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 16567–16578, 2020. [Cited on p. 33.]
- [42] L. Gresele, P. K. Rubenstein, A. Mehrjou, F. Locatello, and B. Schölkopf. The incomplete rosetta stone problem: Identifiability results for multi-view nonlinear ICA. In *Uncertainty in Artificial Intelligence*, pages 217–227. PMLR, 2020. [Cited on p. 2.]
- [43] L. Gresele, J. von Kügelgen, V. Stimper, B. Schölkopf, and M. Besserve. Independent mechanism analysis, a new concept? In *Advances in Neural Information Processing Systems*, volume 34, pages 28233–28248, 2021. [Cited on p. 2 and 7.]
- [44] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. [Cited on p. 7.]
- [45] H. Hälvä and A. Hyvarinen. Hidden markov nonlinear ica: Unsupervised learning from nonstationary time series. In *Conference on Uncertainty in Artificial Intelligence*, pages 939–948. PMLR, 2020. [Cited on p. 2.]
- [46] H. Hälvä, S. Le Corff, L. Lehericy, J. So, Y. Zhu, E. Gassiat, and A. Hyvarinen. Disentangling identifiable features from noisy data with structured nonlinear ica. In *Advances in Neural Information Processing Systems*, volume 34, pages 1624–1633, 2021. [Cited on p. 2.]
- [47] C. Heinze-Deml, J. Peters, and N. Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2), 2018. [Cited on p. 5.]
- [48] M. A. Hernán and J. M. Robins. *Causal inference: What if*. Boca Raton: Chapman & Hall/CRC., 2020. [Cited on p. 1.]
- [49] P. Hoyer, D. Janzing, J. M. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems*, volume 21, pages 689–696, 2008. [Cited on p. 7.]
- [50] B. Huang, K. Zhang, J. Zhang, J. D. Ramsey, R. Sanchez-Romero, C. Glymour, and B. Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020. [Cited on p. 5.]
- [51] A. Hyvarinen and H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. In *Advances in Neural Information Processing Systems*, pages 3765–3773, 2016. [Cited on p. 2.]

- [52] A. Hyvärinen and H. Morioka. Nonlinear ICA of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pages 460–469. PMLR, 2017. [Cited on p. 2.]
- [53] A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000. [Cited on p. 2.]
- [54] A. Hyvärinen and P. Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks*, 12(3):429–439, 1999. [Cited on p. 2 and 4.]
- [55] A. Hyvärinen, H. Sasaki, and R. Turner. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 859–868, 2019. [Cited on p. 2 and 3.]
- [56] A. Hyvärinen, I. Khemakhem, and H. Morioka. Nonlinear independent component analysis for principled disentanglement in unsupervised deep learning. *arXiv preprint arXiv:2303.16535*, 2023. [Cited on p. 2 and 4.]
- [57] G. W. Imbens and D. B. Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015. [Cited on p. 1.]
- [58] A. Immer, C. Schultheiss, J. E. Vogt, B. Schölkopf, P. Bühlmann, and A. Marx. On the identifiability and estimation of causal location-scale noise models. In *40th International Conference on Machine Learning*, 2023. [Cited on p. 7.]
- [59] A. Jaber, M. Kocaoglu, K. Shanmugam, and E. Bareinboim. Causal discovery from soft interventions with unknown targets: Characterization and learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 9551–9561, 2020. [Cited on p. 2, 5, and 8.]
- [60] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2017. [Cited on p. 10.]
- [61] D. Janzing, J. Mooij, K. Zhang, J. Lemeire, J. Zscheischler, P. Daniušis, B. Steudel, and B. Schölkopf. Information-geometric approach to inferring causal directions. *Artificial Intelligence*, 182:1–31, 2012. [Cited on p. 7.]
- [62] D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *The Annals of Statistics*, 41(5):2324 – 2358, 2013. [Cited on p. 8.]
- [63] A.-H. Karimi, J. von Kügelgen, B. Schölkopf, and I. Valera. Towards causal algorithmic recourse. In *xxAI - Beyond Explainable AI. Lecture Notes in Computer Science*, volume 13200, pages 139–166. Springer, 2022. [Cited on p. 1.]
- [64] I. Khemakhem, D. Kingma, R. Monti, and A. Hyvärinen. Variational autoencoders and nonlinear ICA: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pages 2207–2217, 2020. [Cited on p. 3 and 4.]
- [65] I. Khemakhem, R. Monti, D. Kingma, and A. Hyvärinen. ICE-BeeM: Identifiable conditional energy-based deep models based on nonlinear ICA. In *Advances in Neural Information Processing Systems*, volume 33, pages 12768–12778, 2020. [Cited on p. 2.]
- [66] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015. [Cited on p. 34.]
- [67] B. Kivva, G. Rajendran, P. Ravikumar, and B. Aragam. Learning latent causal graphs via mixture oracles. In *Advances in Neural Information Processing Systems*, volume 34, pages 18087–18101, 2021. [Cited on p. 20 and 21.]
- [68] B. Kivva, G. Rajendran, P. Ravikumar, and B. Aragam. Identifiability of deep generative models without auxiliary information. In *Advances in Neural Information Processing Systems*, volume 35, pages 15687–15701, 2022. [Cited on p. 2.]
- [69] D. A. Klindt, L. Schott, Y. Sharma, I. Ustyuzhaninov, W. Brendel, M. Bethge, and D. Paiton. Towards nonlinear disentanglement in natural data with temporal sparse coding. In *International Conference on Learning Representations*, 2020. [Cited on p. 2.]

- [70] M. Kocaoglu, K. Shanmugam, and E. Bareinboim. Experimental design for learning causal graphs with latent variables. In *Advances in Neural Information Processing Systems*, volume 31, pages 7021–7031, 2017. [Cited on p. 5.]
- [71] M. Kocaoglu, A. Jaber, K. Shanmugam, and E. Bareinboim. Characterization and learning of causal graphs with latent variables from soft interventions. In *Advances in Neural Information Processing Systems*, volume 33, pages 14369–14379, 2019. [Cited on p. 2.]
- [72] D. Krueger, E. Caballero, J.-H. Jacobsen, A. Zhang, J. Binas, D. Zhang, R. Le Priol, and A. Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, pages 5815–5826. PMLR, 2021. [Cited on p. 5 and 20.]
- [73] S. Lachapelle and S. Lacoste-Julien. Partial disentanglement via mechanism sparsity. *arXiv preprint arXiv:2207.07732*, 2022. [Cited on p. 10 and 20.]
- [74] S. Lachapelle, P. Rodriguez, Y. Sharma, K. E. Everett, R. Le Priol, A. Lacoste, and S. Lacoste-Julien. Disentanglement via mechanism sparsity regularization: A new principle for nonlinear ICA. In *Conference on Causal Learning and Reasoning*, pages 428–484. PMLR, 2022. [Cited on p. 2 and 20.]
- [75] S. Lachapelle, T. Deleu, D. Mahajan, I. Mitliagkas, Y. Bengio, S. Lacoste-Julien, and Q. Bertrand. Synergies between disentanglement and sparsity: Generalization and identifiability in multi-task learning. In *International Conference on Machine Learning*, pages 18171–18206. PMLR, 2023. [Cited on p. 2.]
- [76] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436, 2015. [Cited on p. 1.]
- [77] E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer Science & Business Media, 2006. [Cited on p. 4.]
- [78] M. S. Lewicki and T. J. Sejnowski. Learning overcomplete representations. *Neural computation*, 12(2):337–365, 2000. [Cited on p. 20.]
- [79] P. Lippe, S. Magliacane, S. Löwe, Y. M. Asano, T. Cohen, and S. Gavves. Citris: Causal identifiability from temporal intervened sequences. In *International Conference on Machine Learning*, pages 13557–13603. PMLR, 2022. [Cited on p. 7, 10, and 20.]
- [80] P. Lippe, S. Magliacane, S. Löwe, Y. M. Asano, T. Cohen, and E. Gavves. Causal representation learning for instantaneous and temporal effects in interactive systems. In *The Eleventh International Conference on Learning Representations*, 2023. [Cited on p. 7 and 20.]
- [81] Y. Liu, Z. Zhang, D. Gong, M. Gong, B. Huang, A. v. d. Hengel, K. Zhang, and J. Q. Shi. Identifying weight-variant latent causal models. *arXiv preprint arXiv:2208.14153*, 2022. [Cited on p. 3, 5, and 20.]
- [82] F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *36th International Conference on Machine Learning*, pages 4114–4124. PMLR, 2019. [Cited on p. 2 and 4.]
- [83] F. Locatello, B. Poole, G. Rätsch, B. Schölkopf, O. Bachem, and M. Tschannen. Weakly-supervised disentanglement without compromises. In *International Conference on Machine Learning*, pages 6348–6359. PMLR, 2020. [Cited on p. 2.]
- [84] D. Lopez-Paz, R. Nishihara, S. Chintala, B. Schölkopf, and L. Bottou. Discovering causal signals in images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2017)*, pages 58–66, 2017. [Cited on p. 1.]
- [85] C. Lu, Y. Wu, J. M. Hernández-Lobato, and B. Schölkopf. Invariant causal representation learning for out-of-distribution generalization. In *International Conference on Learning Representations*, 2022. [Cited on p. 20.]

- [86] C. Mao, K. Xia, J. Wang, H. Wang, J. Yang, E. Bareinboim, and C. Vondrick. Causal transportability for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7521–7531, 2022. [Cited on p. 1.]
- [87] J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *The Journal of Machine Learning Research*, 17(1):1103–1204, 2016. [Cited on p. 6.]
- [88] J. M. Mooij, S. Magliacane, and T. Claassen. Joint causal inference from multiple contexts. *The Journal of Machine Learning Research*, 21(1):3919–4026, 2020. [Cited on p. 5.]
- [89] G. E. Moran, D. Sridhar, Y. Wang, and D. Blei. Identifiable deep generative models via sparse decoding. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. [Cited on p. 2 and 4.]
- [90] S. L. Morgan and C. Winship. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, 2014. [Cited on p. 1.]
- [91] K. Muandet, D. Balduzzi, and B. Schölkopf. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pages 10–18, 2013. [Cited on p. 20.]
- [92] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research*, 22(1):2617–2680, 2021. [Cited on p. 9 and 33.]
- [93] J. Park, U. Shalit, B. Schölkopf, and K. Muandet. Conditional distributional treatment effect with kernel conditional mean embeddings and u-statistic regression. In *International Conference on Machine Learning*, pages 8401–8412. PMLR, 2021. [Cited on p. 9.]
- [94] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, NY, 2nd edition, 2009. [Cited on p. 1, 2, 3, 5, 8, 10, and 31.]
- [95] J. Pearl and E. Bareinboim. External validity: From do-calculus to transportability across populations. *Statistical Science*, pages 579–595, 2014. [Cited on p. 1.]
- [96] J. Pearl and D. Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018. [Cited on p. 20.]
- [97] R. Perry, J. von Kügelgen, and B. Schölkopf. Causal discovery in heterogeneous environments under the sparse mechanism shift hypothesis. In *Advances in Neural Information Processing Systems*, 2022. [Cited on p. 5 and 9.]
- [98] J. Peters, P. Bühlmann, and N. Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012, 2016. [Cited on p. 5.]
- [99] J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT Press, 2017. [Cited on p. 1, 3, 5, 6, 7, and 10.]
- [100] G. Roeder, L. Metz, and D. Kingma. On linear identifiability of learned representations. In *International Conference on Machine Learning*, pages 9030–9039, 2021. [Cited on p. 2.]
- [101] M. Rojas-Carulla, B. Schölkopf, R. Turner, and J. Peters. Invariant models for causal transfer learning. *Journal of Machine Learning Research*, 19(1):1309–1342, 2018. [Cited on p. 1 and 5.]
- [102] K. Roth, M. Ibrahim, Z. Akata, P. Vincent, and D. Bouchacourt. Disentanglement of correlated factors via hausdorff factorized support. In *The Eleventh International Conference on Learning Representations*, 2023. [Cited on p. 21.]
- [103] D. Rothenhäusler, N. Meinshausen, P. Bühlmann, and J. Peters. Anchor regression: Heterogeneous data meet causality. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(2):215–246, 2021. [Cited on p. 5.]

- [104] P. Rubenstein, S. Weichwald, S. Bongers, J. Mooij, D. Janzing, M. Grosse-Wentrup, and B. Schölkopf. Causal consistency of structural equation models. In *33rd Conference on Uncertainty in Artificial Intelligence*, pages 808–817. Curran Associates, Inc., 2017. [Cited on p. 20.]
- [105] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015. [Cited on p. 1.]
- [106] B. Schölkopf. Causality for machine learning. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 765–804, 2022. [Cited on p. 1.]
- [107] B. Schölkopf and J. von Kügelgen. From statistical to causal learning. In *Proceedings of the International Congress of Mathematicians*, 2022. [Cited on p. 1 and 4.]
- [108] B. Schölkopf, D. Janzing, J. Peters, E. Sgouritsa, K. Zhang, and J. Mooij. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning*, pages 459–466, 2012. [Cited on p. 5 and 7.]
- [109] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. [Cited on p. 1, 5, 9, and 20.]
- [110] X. Shen, F. Liu, H. Dong, Q. Lian, Z. Chen, and T. Zhang. Weakly supervised disentangled generative causal representation learning. *Journal of Machine Learning Research*, 23:1–55, 2022. [Cited on p. 20.]
- [111] R. Shu, Y. Chen, A. Kumar, S. Ermon, and B. Poole. Weakly supervised disentanglement with guarantees. In *International Conference on Learning Representations*, 2019. [Cited on p. 2.]
- [112] R. Silva, R. Scheines, C. Glymour, P. Spirtes, and D. M. Chickering. Learning the structure of linear latent variable models. *Journal of Machine Learning Research*, 7(2), 2006. [Cited on p. 20.]
- [113] P. Sorrenson, C. Rother, and U. Köthe. Disentanglement by nonlinear ICA with general incompressible-flow networks (gin). In *International Conference on Learning Representations*, 2019. [Cited on p. 2.]
- [114] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. The MIT Press, 2001. [Cited on p. 2, 3, 4, 10, and 20.]
- [115] C. Squires and C. Uhler. Causal structure learning: a combinatorial perspective. *Foundations of Computational Mathematics*, pages 1–35, 2022. [Cited on p. 2.]
- [116] C. Squires, A. Seigal, S. Bhate, and C. Uhler. Linear causal disentanglement via interventions. In *40th International Conference on Machine Learning*, 2023. [Cited on p. 2, 3, 4, 5, 7, 10, 21, and 36.]
- [117] E. V. Strobl and T. A. Lasko. Identifying patient-specific root causes with the heteroscedastic noise model. *arXiv preprint arXiv:2205.13085*, 2022. [Cited on p. 7.]
- [118] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations*, 2014. [Cited on p. 1.]
- [119] J. Tian and J. Pearl. Causal discovery from changes. In *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence*, pages 512–522, 2001. [Cited on p. 5.]
- [120] C. Uhler, G. Raskutti, P. Bühlmann, and B. Yu. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics*, pages 436–463, 2013. [Cited on p. 3.]
- [121] B. Varici, E. Acarturk, K. Shanmugam, A. Kumar, and A. Tajer. Score-based causal representation learning with interventions. *arXiv preprint arXiv:2301.08230*, 2023. [Cited on p. 3, 5, 7, 8, 10, and 21.]

- [122] J. von Kügelgen, Y. Sharma, L. Gresele, W. Brendel, B. Schölkopf, M. Besserve, and F. Locatello. Self-supervised learning with data augmentations provably isolates content from style. In *Advances in Neural Information Processing Systems*, volume 34, pages 16451–16467, 2021. [Cited on p. 8, 10, 20, and 21.]
- [123] Y. Wald, A. Feder, D. Greenfeld, and U. Shalit. On calibration and out-of-domain generalization. In *Advances in Neural Information Processing Systems*, volume 34, pages 2215–2227, 2021. [Cited on p. 20.]
- [124] Y. Wang and M. I. Jordan. Desiderata for representation learning: A causal perspective. *arXiv preprint arXiv:2109.03795*, 2021. [Cited on p. 3, 20, and 21.]
- [125] L. Wendong, A. Kekić, J. von Kügelgen, S. Buchholz, M. Besserve, L. Gresele, and B. Schölkopf. Causal component analysis. In *Advances in Neural Information Processing Systems*, 2023. [Cited on p. 7, 8, 9, 10, and 20.]
- [126] Q. Xi and B. Bloem-Reddy. Indeterminacy in generative models: Characterization and strong identifiability. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206. PMLR, 2023. [Cited on p. 2 and 4.]
- [127] F. Xie, R. Cai, B. Huang, C. Glymour, Z. Hao, and K. Zhang. Generalized independent noise condition for estimating latent variable causal graphs. In *Advances in Neural Information Processing Systems*, volume 33, pages 14891–14902, 2020. [Cited on p. 20 and 21.]
- [128] F. Xie, B. Huang, Z. Chen, Y. He, Z. Geng, and K. Zhang. Identification of linear non-gaussian latent hierarchical structure. In *International Conference on Machine Learning*, pages 24370–24387. PMLR, 2022. [Cited on p. 20 and 21.]
- [129] M. Yang, F. Liu, Z. Chen, X. Shen, J. Hao, and J. Wang. CausalVAE: Disentangled representation learning via neural structural causal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9593–9602, 2021. [Cited on p. 20.]
- [130] W. Yao, Y. Sun, A. Ho, C. Sun, and K. Zhang. Learning temporally causal latent processes from general temporal data. In *International Conference on Learning Representations*, 2021. [Cited on p. 20.]
- [131] W. Yao, G. Chen, and K. Zhang. Temporally disentangled representation learning. In *Advances in Neural Information Processing Systems*, 2022. [Cited on p. 20.]
- [132] F. M. Zennaro. Abstraction between structural causal models: A review of definitions and properties. *arXiv preprint arXiv:2207.08603*, 2022. [Cited on p. 20.]
- [133] F. M. Zennaro, M. Drávucz, G. Apachitei, W. D. Widanage, and T. Damoulas. Jointly learning consistent causal abstractions over multiple interventional distributions. In *2nd Conference on Causal Learning and Reasoning*, 2023. [Cited on p. 20.]
- [134] K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *25th Conference on Uncertainty in Artificial Intelligence*, pages 647–655. AUAI Press, 2009. [Cited on p. 7.]
- [135] K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *27th Conference on Uncertainty in Artificial Intelligence*, pages 804–813. AUAI Press, 2011. [Cited on p. 9.]
- [136] X. Zheng, B. Aragam, P. Ravikumar, and E. Xing. DAGs with no tears: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018. [Cited on p. 10.]
- [137] X. Zheng, C. Dan, B. Aragam, P. Ravikumar, and E. Xing. Learning sparse nonparametric dags. In *International Conference on Artificial Intelligence and Statistics*, pages 3414–3425. PMLR, 2020. [Cited on p. 10.]
- [138] R. S. Zimmermann, Y. Sharma, S. Schneider, M. Bethge, and W. Brendel. Contrastive learning inverts the data generating process. In *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021. [Cited on p. 2.]

Appendices

Table of Contents

A	Extended Discussion of Related Work	20
A.1	Comparison of Related Identifiability Results	20
B	Proof of Minimality of the CRL Equivalence Class \sim_{CRL}	22
C	Identifiability Proofs	23
C.1	Auxiliary Lemmata	23
C.2	Proof of Thm. 3.2	24
C.3	Proof of Thm. 3.4	29
C.4	Proof of Thm. 4.2	31
D	Experimental Details and Additional Results	33
D.1	Experimental Details for § 6	33
D.2	Additional Results: Learning Nonlinear Latent SCMs from Partial Causal Order	34
E	Discussion of the Role of Our Assumptions	36

A Extended Discussion of Related Work

Prior work on causal representation learning with general nonlinear relationships (both among latents and between latents and observations) and without an explicit task or label typically relies on some form of *weak supervision*. One example of weak supervision is “multi-view” data consisting of tuples of related observations. von Kügelgen et al. [122] consider counterfactual pairs of observations arising from imperfect interventions through data augmentation, and prove identifiability for the invariant non-descendants of intervened-upon variables. Brehmer et al. [18] also use counterfactual pre- and post-intervention views and show that the latent SCM can be identified given all single-node perfect stochastic interventions. Another type of weak-supervision is temporal structure [2], possibly combined with nonstationarity [130, 131], interventions on known targets [79, 80], or observed actions inducing sparse mechanism shifts [73, 74, 109]. Other works use more explicit supervision in the form of annotations of the ground truth causal variables or a known causal graph [110, 125, 129].

A different line of work instead approaches causal representation learning from the perspective of causal discovery in the presence of latent variables [114]. Doing so from *purely observational i.i.d. data* requires additional constraints on the generative process, such as restrictions on the graph structure or particular parametric and distributional assumptions, and typically leverages the identifiability of linear ICA [27, 38, 78]. For *linear, non-Gaussian* models, Silva et al. [112] show that the causal DAG can be recovered up to Markov equivalence if all observed variables are “pure” in that they have a unique latent causal parent. Cai et al. [23] and Xie et al. [127, 128] extend this result to identify the full graph given two pure observed children per latent, and Adams et al. [1] provide sufficient and necessary graphical conditions for full identification. For *discrete* causal variables, Kivva et al. [67] introduce a similar “no twins” condition to reduce the task of learning the number and cardinality of latents and the Markov-equivalence class of the graph to mixture identifiability.

Other lines of work have investigated the relationship between causal models at different levels of coarse-graining or abstraction [6, 11, 12, 24–26, 35, 104, 132, 133], or learning invariant representations in a supervised setting [3, 8, 32, 33, 72, 85, 91, 123, 124], often for domain generalization.

A.1 Comparison of Related Identifiability Results

To complement the presentation of related multienvironment CRL works in § 1, we provide a structured overview of and comparison with existing identifiability results for causal representation learning in Tab. 1. The table categorizes work along different dimensions. First, we make a distinction based on the type of data (observational, interventional, or counterfactual) results rely on (colour coded in green, yellow, and red, respectively). These are also referred to as different rungs in the “ladder of causation” [96] or layers in the Pearl Causal Hierarchy [10]. Second, we categorize work depending on the assumptions placed on the latent causal model and the mixing function. As can be seen, works relying solely on observational data often require restrictive graphical assumptions on the mixing function. On the other hand, access to much more informative counterfactual data has allowed identification even for nonparametric causal models and mixing functions. Our work can be viewed as a step towards addressing the lack of nonparametric identifiability results in the interventional regime.

We emphasize that Tab. 1 is not exhaustive: certain relevant works did not easily fit into our categorization or the causal representation learning framework adopted in the present work. For example, not listed are works that leverage temporal structure [2, 73, 74, 79, 80], rely on heterogenous data and distribution shifts which are not directly expressed in terms of or linked to interventions [81, 130, 131], allow for edges from observed to latent variables [1], or require more direct supervision [110, 129].

Table 1: Comparison of Existing Identifiability Results for Causal Representation Learning. All of the listed works assume invertibility (or injectivity) of the mixing function, as well as causal sufficiency (Markovianity) for the causal latent variables. Most or all of the listed results require additional technical assumptions, and may provide additional results, which we omit for sake of readability; see the references for more details.

Work	Layer	Causal Model	Mixing Function	Main Identifiability Result
Cai et al. [23], Xie et al. [127, 128]	observational	linear, non-Gaussian	linear with non-Gaussian noise s.t. each V_i has 2 pure (obs. or unobs.) children	number of latents + G
Kivva et al. [67]	observational	discrete, nonparametric	identifiable mixture model s.t. obs. children of $V_i \not\subseteq$ obs. children of V_j	number, cardinality & dist. of discrete latents + G up to Markov equivalence
Ahuja et al. [5, Thm. 4]	observational	nonlinear w. independent support [102, 124]	finite-degree polynomial	V up to permutation, shift, & linear scaling
Squires et al. [116, Thms. 1 & 2]	interventional	linear	linear	G and V up to partial-order preserving permutations from obs. dist. & all single-node <i>perfect</i> interventions
Squires et al. [116, Thm. 1]	interventional	linear	linear	G up to transitive closure from obs. dist. & all single-node <i>imperfect</i> interventions
Varici et al. [121, Thm. 16]	interventional	nonparametric	linear	G and V up to partial-order preserving permutations from obs. dist. & all single-node <i>perfect</i> interventions
Ahuja et al. [5, Thm. 2]	interventional	nonparametric	finite-degree polynomial	V up to permutation, shift, and linear scaling from all single-node <i>perfect hard</i> interventions
Buchholz et al. [22]	interventional	linear Gaussian	nonparametric	G and V up to permutation from obs. dist. & all single-node <i>perfect</i> interventions
This Work (Thm. 3.2)	interventional	nonparametric	nonparametric	for $n = 2$: G and V up to \sim_{CRL} from all single-node <i>perfect</i> interventions, subject to genericity (3.2)
This Work (Thm. 3.4)	interventional	nonparametric	nonparametric	G and V up to \sim_{CRL} from two distinct, paired single-node <i>perfect</i> interventions per node
von Kügelgen et al. [122]	counterfactual	nonparametric	nonparametric	block of non-descendants $V_{\text{nd}(\mathcal{I})}$ up to invertible function from fat-hand <i>imperfect</i> interventions on $V_{\mathcal{I}}$
Brehmer et al. [18]	counterfactual	nonparametric	nonparametric	G and V up to \sim_{CRL} from all single-node <i>perfect</i> interventions

B Proof of Minimality of the CRL Equivalence Class \sim_{CRL}

First, let us recall the main statements from § 2.2.

Definition 2.6 (\sim_{CRL} -identifiability). Let \mathcal{H} be a space of unmixing functions $h : \mathcal{X} \rightarrow \mathbb{R}^n$ and let \mathcal{G} be the space of DAGs over n vertices. Let \sim_{CRL} be the equivalence relation on $\mathcal{H} \times \mathcal{G}$ defined as

$$(h_1, G_1) \sim_{\text{CRL}} (h_2, G_2) \iff (h_2, G_2) = (\mathbf{P}_{\pi^{-1}} \circ \phi \circ h_1, \pi(G_1)) \quad (\text{B.1})$$

for some element-wise diffeomorphism $\phi(\mathbf{v}) = (\phi_1(v_1), \dots, \phi_n(v_n))$ of \mathbb{R}^n and a permutation π of $[n]$ such that $\pi : G_1 \mapsto G_2$ is a graph isomorphism and \mathbf{P}_{π} the corresponding permutation matrix.

Proposition 3.1 (Minimality of \sim_{CRL} ; informal). *Let \mathbf{Z} be any representation that is \sim_{CRL} equivalent to \mathbf{V} , with $G' = \pi(G)$ the associated DAG. Then for any intervention on $\mathcal{V}_{\mathcal{I}} \subseteq \mathcal{V}$ in G , there exists an equally sparse intervention on $\mathcal{Z}_{\pi(\mathcal{I})} \subseteq \mathcal{Z}$ in G' inducing the same observed distribution on \mathcal{X} .*

We now restate this result more formally.

Proposition B.1 (Minimality of \sim_{CRL}). *Let $(h, G') \sim_{\text{CRL}} (f^{-1}, G)$ with π denoting the graph isomorphism mapping G to G' (i.e., a permutation that preserves the partial topological order of G). Let $\mathbf{Z} = h(\mathbf{X})$ be the inferred representation with distribution $Q_{\mathbf{Z}} = h_*(P_{\mathbf{X}})$ Markov w.r.t. G' and associated density q . Let $\mathcal{I}^e \subseteq [n]$ denote a set of intervention targets, and consider an intervention that changes $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ to some intervened mechanism $\tilde{p}_i(v_i | \mathbf{v}_{\text{pa}(i)})$ for all $i \in \mathcal{I}^e$, giving rise to the interventional distributions $P_{\mathbf{V}}^e$ and $P_{\mathbf{X}}^e = f_*(P_{\mathbf{V}}^e)$. Then there exist appropriately chosen $\tilde{q}_{\pi(i)}(z_{\pi(i)} | \mathbf{z}_{\text{pa}(\pi(i), G')})$ for $i \in \mathcal{I}^e$ such that the resulting interventional distribution $Q_{\mathbf{Z}}^e$ gives rise to the same observed distributions, that is, $P_{\mathbf{X}}^e = h_*^{-1}(Q_{\mathbf{Z}}^e)$.*

Proof. Since $(h, G') \sim_{\text{CRL}} (f^{-1}, G)$, by Defn. 2.6 we have

$$\mathbf{Z} = \mathbf{P}_{\pi^{-1}} \circ \phi(\mathbf{V}) \quad (\text{B.1})$$

for some element-wise diffeomorphism ϕ with inverse $\psi = \phi^{-1}$. Then (B.1) implies that for all $i \in [n]$

$$V_i = \psi_i(Z_{\pi(i)}) \quad (\text{B.2})$$

According to (B.2), each conditional in the Markov factorization of $Q_{\mathbf{Z}}$ is given in terms of p by

$$q_{\pi(i)}(z_{\pi(i)} | \mathbf{z}_{\text{pa}(\pi(i), G')}) = p_i \left(\psi_i(z_{\pi(i)}) | \psi_{\text{pa}(i)}(\mathbf{z}_{\text{pa}(\pi(i), G')}) \right) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right| \quad (\text{B.3})$$

where we have used the change of variables in (B.2) and the fact that $\pi(\text{pa}(i)) = \text{pa}(\pi(i), G')$ since $\pi : G \mapsto G'$ is a graph isomorphism.

Consider an intervention that changes $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ to some intervened mechanism $\tilde{p}_i(v_i | \mathbf{v}_{\text{pa}(i)})$ for all $i \in \mathcal{I}^e$. Denote the corresponding interventional joint distribution by $P_{\mathbf{V}}^e$ with joint density p^e given by

$$p^e(\mathbf{v}) = \prod_{i \in \mathcal{I}^e} \tilde{p}_i(v_i | \mathbf{v}_{\text{pa}(i)}) \prod_{j \in [n] \setminus \mathcal{I}^e} p_j(v_j | \mathbf{v}_{\text{pa}(j)}) \quad (\text{B.4})$$

Denote by $Q_{\mathbf{Z}}^e = (\mathbf{P}_{\pi^{-1}} \circ \phi)_*(P_{\mathbf{V}}^e)$ the corresponding distribution over \mathbf{Z} with joint density given by q^e

$$q^e(\mathbf{z}) = p^e(\psi \circ \mathbf{P}_{\pi}(\mathbf{z})) |\det \mathbf{J}_{\psi \circ \mathbf{P}_{\pi}}(\mathbf{z})| \quad (\text{B.5})$$

$$= \prod_{i \in \mathcal{I}^e} \tilde{p}_i \left(\psi_i(z_{\pi(i)}) | \psi_{\text{pa}(i)}(\mathbf{z}_{\text{pa}(\pi(i), G')}) \right) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right| \prod_{j \in [n] \setminus \mathcal{I}^e} q_{\pi(j)}(z_{\pi(j)} | \mathbf{z}_{\text{pa}(\pi(j), G')}) , \quad (\text{B.6})$$

where we have used (B.3), (B.4), and the fact that \mathbf{J}_{ψ} is diagonal.

By defining

$$\tilde{q}_{\pi(i)}(z_{\pi(i)} | \mathbf{z}_{\text{pa}(\pi(i), G')}) := \tilde{p}_i \left(\psi_i(z_{\pi(i)}) | \psi_{\text{pa}(i)}(\mathbf{z}_{\text{pa}(\pi(i), G')}) \right) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right| \quad (\text{B.7})$$

we finally arrive at

$$q^e(\mathbf{z}) = \prod_{i \in \mathcal{I}^e} \tilde{q}_{\pi(i)} \left(z_{\pi(i)} \mid \mathbf{z}_{\text{pa}(\pi(i); G')} \right) \prod_{j \in [n] \setminus \mathcal{I}^e} q_{\pi(j)} \left(z_{\pi(j)} \mid \mathbf{z}_{\text{pa}(\pi(j); G')} \right). \quad (\text{B.8})$$

This shows that any intervention on $\{V_i\}_{i \in \mathcal{I}^e} \subseteq \mathbf{V}$ which replaces

$$\left\{ p_i(v_i \mid \mathbf{v}_{\text{pa}(i)}) \right\}_{i \in \mathcal{I}^e} \mapsto \left\{ \tilde{p}_i(v_i \mid \mathbf{v}_{\text{pa}(i)}) \right\}_{i \in \mathcal{I}^e}, \quad (\text{B.9})$$

can equivalently be captured by an intervention on $\{Z_{\pi(i)}\}_{i \in \mathcal{I}^e} \subseteq \mathbf{Z}$ which replaces

$$\left\{ q_{\pi(i)} \left(z_{\pi(i)} \mid \mathbf{z}_{\text{pa}(\pi(i); G')} \right) \right\}_{i \in \mathcal{I}^e} \mapsto \left\{ \tilde{q}_{\pi(i)} \left(z_{\pi(i)} \mid \mathbf{z}_{\text{pa}(\pi(i); G')} \right) \right\}_{i \in \mathcal{I}^e}. \quad (\text{B.10})$$

with \tilde{q}_i defined according to (B.7). \square

C Identifiability Proofs

C.1 Auxiliary Lemmata

Lemma C.1 (Lemma 2 of Brehmer et al. [18]). *Let $A = C = \mathbb{R}$ and $B = \mathbb{R}^n$. Let $f : A \times B \rightarrow C$ be differentiable. Define two differentiable measures p_A on A and p_C on C . Let $\forall b \in B$, $f(\cdot, b) : A \rightarrow C$ be measure-preserving, in the sense that the pushforward of p_A is always p_C . Then $f(\cdot, b)$ is constant in b on B .*

Proof. See Appendix A.2 of Brehmer et al. [18]. \square

Lemma C.2 (Preservation of conditional independence under invertible transformation.). *Let X and Y be continuous real-valued random variables, and let \mathbf{Z} be a continuous random vector taking values in \mathbb{R}^n . Suppose that (X, Y, \mathbf{Z}) have a joint density w.r.t. the Lebesgue measure. Let $f : \mathbb{R} \rightarrow \mathbb{R}$, $g : \mathbb{R} \rightarrow \mathbb{R}$, and $h : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be diffeomorphisms. Then $X \perp\!\!\!\perp Y \mid \mathbf{Z} \implies f(X) \perp\!\!\!\perp g(Y) \mid h(\mathbf{Z})$.*

Proof. Denote by $p(x, y, \mathbf{z})$ the density of (X, Y, \mathbf{Z}) . Then $X \perp\!\!\!\perp Y \mid \mathbf{Z}$ implies that for all (x, y, \mathbf{z}) , p can be factorized as follows:

$$p(x, y, \mathbf{z}) = p_{\mathbf{z}}(\mathbf{z}) p_x(x \mid \mathbf{z}) p_y(y \mid \mathbf{z}). \quad (\text{C.1})$$

Let $(A, B, C) = (f(X), g(Y), h(\mathbf{Z}))$, and write $\tilde{f} = f^{-1}$, $\tilde{g} = g^{-1}$, and $\tilde{h} = h^{-1}$.

Then (A, B, C) also has a density $q(a, b, \mathbf{c})$, which for all (a, b, \mathbf{c}) is given by the change of variable formula:

$$q(a, b, \mathbf{c}) = p \left(\tilde{f}(a), \tilde{g}(b), \tilde{h}(\mathbf{c}) \right) \left| \frac{d\tilde{f}}{da}(a) \frac{d\tilde{g}}{db}(b) \det \mathbf{J}_{\tilde{h}}(\mathbf{c}) \right| \quad (\text{C.2})$$

$$= p_{\mathbf{z}} \left(\tilde{h}(\mathbf{c}) \right) \left| \det \mathbf{J}_{\tilde{h}}(\mathbf{c}) \right| p_x \left(\tilde{f}(a) \mid \tilde{h}(\mathbf{c}) \right) \left| \frac{d\tilde{f}}{da}(a) \right| p_y \left(\tilde{g}(b) \mid \tilde{h}(\mathbf{c}) \right) \left| \frac{d\tilde{g}}{db}(b) \right| \quad (\text{C.3})$$

where in (C.2) we have used the fact that $(X, Y, \mathbf{Z}) \mapsto (f(X), g(Y), h(\mathbf{Z}))$ has block-diagonal Jacobian, and in (C.3) that p factorises as in (C.1). Next, define

$$q_{\mathbf{c}}(\mathbf{c}) := p_{\mathbf{z}} \left(\tilde{h}(\mathbf{c}) \right) \left| \det \mathbf{J}_{\tilde{h}}(\mathbf{c}) \right|, \quad (\text{C.4})$$

$$q_a(a \mid \mathbf{c}) := p_x \left(\tilde{f}(a) \mid \tilde{h}(\mathbf{c}) \right) \left| \frac{d\tilde{f}}{da}(a) \right|, \quad (\text{C.5})$$

$$q_b(b \mid \mathbf{c}) := p_y \left(\tilde{g}(b) \mid \tilde{h}(\mathbf{c}) \right) \left| \frac{d\tilde{g}}{db}(b) \right|. \quad (\text{C.6})$$

Since $p_{\mathbf{z}}$, p_x , and p_y are valid densities (non-negative and integrating to one), so are $q_{\mathbf{c}}$, q_a , and q_b . Substitution into (C.3) yields for all (a, b, \mathbf{c}) ,

$$q(a, b, \mathbf{c}) = q_{\mathbf{c}}(\mathbf{c}) q_a(a \mid \mathbf{c}) q_b(b \mid \mathbf{c}), \quad (\text{C.7})$$

which implies that $A \perp\!\!\!\perp B \mid C$, concluding the proof. \square

C.2 Proof of Thm. 3.2

Theorem 3.2 (Bivariate identifiability up to \sim_{CRL} from one perfect stochastic intervention per node). *Suppose that we have access to multiple environments $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$ generated as described in § 2 under Asms. 2.2, 2.5, 2.8 and 2.9 with $n = 2$. Let (h, G') be any candidate solution such that the inferred latent distributions $Q_{\mathbf{Z}}^e = h_*(P_{\mathbf{X}}^e)$ of $\mathbf{Z} = h(\mathbf{X})$ and the inferred mixing function h^{-1} satisfy the above assumptions w.r.t. the candidate causal graph G' . Assume additionally that*

- (A1) *all densities p^e and q^e are continuously differentiable and fully supported on \mathbb{R}^n ;*
- (A2) *we have access to a known observational environment e_0 and one single node perfect intervention for each node, with unknown targets: there exist $n+1$ environments $\mathcal{E} = \{e_i\}_{i=0}^n$ such that $\mathcal{I}^{e_0} = \emptyset$ and for each $i \in [n]$ we have $\mathcal{I}^{e_i} = \{\pi(i)\}$ for an unknown permutation π of $[n]$;*
- (A3) *for all $i \in [n]$, the intervened mechanisms $\tilde{p}_i(v_i)$ differ from the corresponding base mechanisms $p_i(v_i | \mathbf{v}_{\text{pa}(i)})$ everywhere, in the sense that*

$$\forall \mathbf{v} : \quad \frac{\partial}{\partial v_i} \frac{\tilde{p}_i(v_i)}{p_i(v_i | \mathbf{v}_{\text{pa}(i)})} \neq 0; \quad (3.1)$$

- (A4) (“**genericity**”) *the base and intervened mechanisms are not fine-tuned to each other, in the sense that there exists a continuous function $\varphi : \mathbb{R}^+ \rightarrow \mathbb{R}$ for which*

$$\mathbb{E}_{\mathbf{v} \sim P_{\mathbf{V}}^{e_0}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 | v_1)} \right) \right] \neq \mathbb{E}_{\mathbf{v} \sim P_{\mathbf{V}}^{e_1}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 | v_1)} \right) \right] \quad (3.2)$$

Then the ground truth is identified in the sense of Defn. 2.6, that is, $(f^{-1}, G) \sim_{\text{CRL}} (h, G')$.

Proof. From the assumption of a shared mixing f and shared encoder h across all environments, we have that

$$\mathbf{Z} = h(\mathbf{X}) = h(f(\mathbf{V})) = h \circ f(\mathbf{V}). \quad (\text{C.8})$$

Let $\psi = f^{-1} \circ h^{-1} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that

$$\mathbf{V} = \psi(\mathbf{Z}).$$

By Asm. 2.5, f , h , and thus also $h \circ f$ are diffeomorphisms. Hence, ψ is well-defined and also diffeomorphic.

By the change of variable formula, for all e and all \mathbf{z} the density $q^e(\mathbf{z})$ is given by

$$q^e(\mathbf{z}) = p^e(\psi(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (\text{C.9})$$

where $(\mathbf{J}_{\psi}(\mathbf{z}))_{ij} = \frac{\partial \psi_i}{\partial z_j}(\mathbf{z})$ denotes the Jacobian of ψ .

We now consider two separate cases, depending on whether the intervention targets in q^{e_i} for $e_i \in \{e_1, e_2\}$ match those in p^{e_i} (Case 1) or not (Case 2).

Case 1: Aligned Intervention Targets. According to Asm. 2.8 and (A2), (C.9) applied to the known observational environment e_0 and the interventional environments e_1, e_2 leads to the system of equations:

$$q_1(z_1)q_2(z_2 | z_{\text{pa}(2; G')}) = p_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_{\text{pa}(2)}(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (\text{C.10})$$

$$\tilde{q}_1(z_1)q_2(z_2 | z_{\text{pa}(2; G')}) = \tilde{p}_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_{\text{pa}(2)}(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (\text{C.11})$$

$$q_1(z_1)\tilde{q}_2(z_2) = p_1(\psi_1(\mathbf{z})) \tilde{p}_2(\psi_2(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| \quad (\text{C.12})$$

where $z_{\text{pa}(2; G')}$ denotes the parents of z_2 in G' , and $\text{pa}(2)$ denotes the parents of V_2 in G .

Note that neither side of the previous equations can be zero due to the full support assumption¹² (A1) and ψ being diffeomorphic implying the determinant is non-zero.

¹²This can also be relaxed to fully supported on a Cartesian product of intervals.

Taking quotients of (C.11) and (C.10), yields

$$\frac{\tilde{q}_1}{q_1}(z_1) = \frac{\tilde{p}_1}{p_1}(\psi_1(\mathbf{z})). \quad (\text{C.13})$$

Next, we take the partial derivative w.r.t. z_2 on both sides and use the chain rule to obtain:

$$0 = \left(\frac{\tilde{p}_1}{p_1} \right)' (\psi_1(\mathbf{z})) \frac{\partial \psi_1}{\partial z_2}(\mathbf{z}). \quad (\text{C.14})$$

Now, by assumption (A3), the first term on the RHS of (C.14) is non-zero everywhere. Hence,

$$\forall \mathbf{z} : \frac{\partial \psi_1}{\partial z_2}(\mathbf{z}) = 0. \quad (\text{C.15})$$

It follows that ψ_1 is, in fact, a scalar function, and

$$V_1 = \psi_1(Z_1). \quad (\text{C.16})$$

Since ψ is a diffeomorphism, ψ_1 must also be diffeomorphic. Hence, by the change of variable formula applied to (C.16), the marginal density $q_1(z_1)$ is given by

$$q_1(z_1) = p_1(\psi_1(z_1)) \left| \frac{\partial \psi_1}{\partial z_1}(z_1) \right|. \quad (\text{C.17})$$

Further, since \mathbf{J}_ψ is triangular due to (C.15), its determinant is given by

$$|\det \mathbf{J}_\psi(\mathbf{z})| = \left| \frac{\partial \psi_1}{\partial z_1}(z_1) \frac{\partial \psi_2}{\partial z_2}(z_1, z_2) \right|. \quad (\text{C.18})$$

Substituting (C.17) and (C.18) into (C.12) yields (after cancellation of equal terms):

$$\tilde{q}_2(z_2) = \tilde{p}_2(\psi_2(z_1, z_2)) \left| \frac{\partial \psi_2}{\partial z_2}(z_1, z_2) \right|. \quad (\text{C.19})$$

The expression in (C.19) implies that, for all z_1 , the mapping $\psi_2(z_1, \cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is measure preserving for the differentiable \tilde{q}_2 and \tilde{p}_2 . By Lemma C.1 (Lemma 2 of Brehmer et al. [18, § A.2]), it then follows that ψ_2 must, in fact, be constant in z_1 , that is

$$\forall \mathbf{z} : \frac{\partial \psi_2}{\partial z_1}(\mathbf{z}) = 0. \quad (\text{C.20})$$

Note that this last step is where the assumption of perfect interventions (Asm. 2.9) is leveraged: the conclusion would not hold for arbitrary imperfect interventions for which (3.8) would involve $\tilde{q}_2(z_2 | z_1)$ and $p_2(\psi_2(z_1, z_2) | \psi_1(z_1))$.

Hence, we have shown that ψ is an element-wise function:

$$\mathbf{V} = (V_1, V_2) = \psi(\mathbf{Z}) = (\psi_1(Z_1), \psi_2(Z_2)). \quad (\text{C.21})$$

Finally, since ψ is a diffeomorphism, (C.21) implies that

$$V_1 \perp\!\!\!\perp V_2 \iff Z_1 \perp\!\!\!\perp Z_2. \quad (\text{C.22})$$

It then follows from the faithfulness assumption (Asm. 2.2) that we also must have $G = G'$.

This concludes the proof of Case 1, as we have shown that $(h, G') \sim_{\text{CRL}} (f^{-1}, G)$ with $G' = \pi(G) = G$ (π being the identity permutation) and $h \circ f = \psi^{-1} =: \phi$ an element-wise diffeomorphism.

Case 2: Misaligned Intervention Targets. Writing down the system of equations similar to (C.10)–(C.12), but for the case with misaligned intervention targets across p and q yields:

$$q_1(z_1)q_2(z_2 | z_{\text{pa}(2;G')}) = p_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_{\text{pa}(2)}(\mathbf{z})) |\det \mathbf{J}_\psi(\mathbf{z})| \quad (\text{C.23})$$

$$\tilde{q}_1(z_1)q_2(z_2 | z_{\text{pa}(2;G')}) = p_1(\psi_1(\mathbf{z})) \tilde{p}_2(\psi_2(\mathbf{z})) |\det \mathbf{J}_\psi(\mathbf{z})| \quad (\text{C.24})$$

$$q_1(z_1)\tilde{q}_2(z_2) = \tilde{p}_1(\psi_1(\mathbf{z})) p_2(\psi_2(\mathbf{z}) | \psi_{\text{pa}(2)}(\mathbf{z})) |\det \mathbf{J}_\psi(\mathbf{z})|. \quad (\text{C.25})$$

Taking ratios of e_1 and e_2 with e_0 yields

$$\frac{\tilde{q}_1(z_1)}{q_1} = \frac{\tilde{p}_2(\psi_2(\mathbf{z}))}{p_2(\psi_2(\mathbf{z}) \mid \psi_{\text{pa}(2)}(\mathbf{z}))} \quad (\text{C.26})$$

$$\frac{\tilde{q}_2(z_2)}{q_2(z_2 \mid z_{\text{pa}(2;G')})} = \frac{\tilde{p}_1(\psi_1(\mathbf{z}))}{p_1} . \quad (\text{C.27})$$

We separate the remainder of the proof of Case 2 into different subcases depending on G and G' : as we will see, we can use a similar reasoning as in Case 1, except for the case where both G and G' are missing no edge.

Case 2a: $V_1 \not\rightarrow V_2$ in G , that is $\text{pa}(2) = \emptyset$. Then we can proceed similarly to Case 1. First, we take the partial derivative of (C.26) w.r.t. z_2 to arrive at:

$$0 = \left(\frac{\tilde{p}_2}{p_2} \right)' (\psi_2(\mathbf{z})) \frac{\partial \psi_2}{\partial z_2}(\mathbf{z}) . \quad (\text{C.28})$$

Using (A3), this implies that ψ_2 does not depend on Z_2 , that is, $V_2 = \psi_2(Z_1)$.

Next, we again write $q(z_1)$ in terms of $p_2(\psi_2(z_1))$ using the univariate change of variable formula, substitute into (C.25), cancel the corresponding terms, and arrive at:

$$\tilde{q}_2(z_2) = \tilde{p}_1(\psi_1(z_1, z_2)) \left| \frac{\partial \psi_1}{\partial z_2}(z_1, z_2) \right| \quad (\text{C.29})$$

Lemma C.1 applied to $\psi_1(z_1, \cdot)$ which preserves \tilde{q}_2 and \tilde{p}_1 for all z_1 shows that ψ_1 is constant in Z_1 , that is

$$\mathbf{V} = (V_1, V_2) = \psi(\mathbf{Z}) = (\psi_1(Z_2), \psi_2(Z_1)) . \quad (\text{C.30})$$

Since $V_1 \perp\!\!\!\perp V_2$ by the assumption of Case 2a, it follows from the invertible element-wise reparametrisation above that $Z_1 \perp\!\!\!\perp Z_2$ and hence, by faithfulness, $Z_1 \not\rightarrow Z_2$ in G' .

Finally, note that there is no partial order on the empty graph and so $G' = \pi(G) = G$ and $\mathbf{Z} = \mathbf{P}_{\pi^{-1}} \cdot \psi^{-1}(\mathbf{V})$ where π is the nontrivial permutation of $\{1, 2\}$.

Case 2b: $V_1 \rightarrow V_2$ in G , that is $\text{pa}(2) = \{1\}$. If $G' \neq G$, that is $Z_1 \not\rightarrow Z_2$ in G' , then the same argument as in Case 2a, this time starting by taking the partial derivative of (C.27) w.r.t. z_1 , can be used to reach the same conclusion in (C.30). However, this contradicts faithfulness since $V_1 \not\perp\!\!\!\perp V_2$ in G .

Hence, we must have $G' = G$, and the following two equations must hold for all \mathbf{z} :

$$\frac{\tilde{q}_1(z_1)}{q_1(z_1)} = \frac{\tilde{p}_2(\psi_2(\mathbf{z}))}{p_2(\psi_2(\mathbf{z}) \mid \psi_1(\mathbf{z}))} \quad (\text{C.31})$$

$$\frac{\tilde{q}_2(z_2)}{q_2(z_2 \mid z_1)} = \frac{\tilde{p}_1(\psi_1(\mathbf{z}))}{p_1(\psi_1(\mathbf{z}))} \quad (\text{C.32})$$

The remainder of the proof consists of exploring the implications of (C.32) and (C.31), ultimately resulting in a violation of the genericity condition (A4).

To ease notation, define the following auxiliary functions:

$$a(z_1) := \frac{\tilde{q}_1(z_1)}{q_1(z_1)} , \quad (\text{C.33})$$

$$b(\mathbf{v}) := \frac{\tilde{p}_2(v_2)}{p_2(v_2 \mid v_1)} , \quad (\text{C.34})$$

$$c(\mathbf{z}) := \frac{\tilde{q}_2(z_2)}{q_{2|1}(z_2 \mid z_1)} , \quad (\text{C.35})$$

$$d(v_1) := \frac{\tilde{p}_1(v_1)}{p_1(v_1)} . \quad (\text{C.36})$$

With this, (C.31) and (C.32) take the following form:

$$a(z_1) = b(\psi(\mathbf{z})). \quad (\text{C.37})$$

$$c(\mathbf{z}) = d(\psi_1(\mathbf{z})), \quad (\text{C.38})$$

Next, define the following maps:

$$\kappa : \mathbf{z} \mapsto \begin{bmatrix} a(z_1) \\ c(\mathbf{z}) \end{bmatrix} \quad (\text{C.39})$$

$$\rho : \mathbf{v} \mapsto \begin{bmatrix} b(\mathbf{v}) \\ d(v_1) \end{bmatrix} \quad (\text{C.40})$$

Then, (C.37) and (C.38) together imply that

$$\kappa = \rho \circ \psi. \quad (\text{C.41})$$

Recalling that by (A1) all densities are continuously differentiable, the Jacobians of κ and ρ are given by:

$$\mathbf{J}_\kappa(\mathbf{z}) = \begin{bmatrix} a'(z_1) & 0 \\ \frac{\partial c}{\partial z_1}(\mathbf{z}) & \frac{\partial c}{\partial z_2}(\mathbf{z}) \end{bmatrix}, \quad (\text{C.42})$$

$$\mathbf{J}_\rho(\mathbf{v}) = \begin{bmatrix} \frac{\partial b}{\partial v_1}(\mathbf{v}) & \frac{\partial b}{\partial v_2}(\mathbf{v}) \\ d'(v_1) & 0 \end{bmatrix}, \quad (\text{C.43})$$

and the corresponding determinants are given by

$$|\det \mathbf{J}_\kappa(\mathbf{z})| = \left| a'(z_1) \frac{\partial c}{\partial z_2}(\mathbf{z}) \right| \neq 0 \quad (\text{C.44})$$

$$|\det \mathbf{J}_\rho(\mathbf{v})| = \left| d'(v_1) \frac{\partial b}{\partial v_2}(\mathbf{v}) \right| \neq 0 \quad (\text{C.45})$$

where the inequalities for all \mathbf{z} follow since, by assumption (A3), the derivatives of ratios of intervened and original mechanisms are non-vanishing everywhere:

$$a'(z_1) \neq 0 \neq \frac{\partial c}{\partial z_2}(\mathbf{z}) \quad \text{and} \quad d'(v_1) \neq 0 \neq \frac{\partial b}{\partial v_2}(\mathbf{v}), \quad (\text{C.46})$$

This implies that the following families of maps are continuously differentiable, monotonic, and invertible,

$$a : z_1 \mapsto a(z_1), \quad (\text{C.47})$$

$$b_{v_1} : v_2 \mapsto b(v_1, v_2), \quad (\text{C.48})$$

$$c_{z_1} : z_2 \mapsto c(z_1, z_2), \quad (\text{C.49})$$

$$d : v_1 \mapsto d(v_1), \quad (\text{C.50})$$

with continuously differentiable inverses

$$a^{-1} : w_1 \mapsto a^{-1}(w_1), \quad (\text{C.51})$$

$$b_{v_1}^{-1} : w_1 \mapsto b_{v_1}^{-1}(w_1), \quad (\text{C.52})$$

$$c_{z_1}^{-1} : w_2 \mapsto c_{z_1}^{-1}(w_2), \quad (\text{C.53})$$

$$d^{-1} : w_2 \mapsto d^{-1}(w_2). \quad (\text{C.54})$$

This implies that ρ and κ are valid diffeomorphisms onto their image and their inverses are given by:

$$\kappa^{-1} : \mathbf{w} \mapsto \begin{bmatrix} a^{-1}(w_1) \\ c_{a^{-1}(w_1)}^{-1}(w_2) \end{bmatrix}, \quad (\text{C.55})$$

$$\rho^{-1} : \mathbf{w} \mapsto \begin{bmatrix} d^{-1}(w_2) \\ b_{d^{-1}(w_2)}^{-1}(w_1) \end{bmatrix}. \quad (\text{C.56})$$

Since $\mathbf{V} = \psi(\mathbf{Z})$, by (C.41) we have

$$\mathbf{W} := \rho(\mathbf{V}) = \rho \circ \psi(\mathbf{Z}) = \kappa(\mathbf{Z}). \quad (\text{C.57})$$

Denote the distributions of \mathbf{W} by $R_{\mathbf{W}}$ and its density by $r(\mathbf{w})$. Since for all e , we have

$$P_{\mathbf{V}}^e = \psi_*(Q_{\mathbf{Z}}^e) \quad (\text{C.58})$$

it follows from (C.57) that

$$R_{\mathbf{W}}^e = \rho_*(P_{\mathbf{V}}^e) = \kappa_*(Q_{\mathbf{Z}}^e). \quad (\text{C.59})$$

This provides two different ways of applying the change of variable formula to compute $r(\mathbf{w})$.

First, we consider the pushforward of $Q_{\mathbf{Z}}^{e_0}$ by κ :

$$r(\mathbf{w}) = q\left(\kappa^{-1}(\mathbf{w})\right) \left| \det \mathbf{J}_{\kappa^{-1}}(\mathbf{w}) \right| \quad (\text{C.60})$$

$$= q_1\left(a^{-1}(w_1)\right) q_2\left(c_{a^{-1}(w_1)}^{-1}(w_2) \mid a^{-1}(w_1)\right) \left| \frac{d}{dw_1} a^{-1}(w_1) \frac{d}{dw_2} c_{a^{-1}(w_1)}^{-1}(w_2) \right| \quad (\text{C.61})$$

By integrating this joint density with respect to w_2 , we obtain the following expression for the marginal $r_1(w_1)$:

$$r_1(w_1) = \left| \frac{d}{dw_1} a^{-1}(w_1) \right| q_1\left(a^{-1}(w_1)\right) \int q_2\left(c_{a^{-1}(w_1)}^{-1}(w_2) \mid a^{-1}(w_1)\right) \left| \frac{d}{dw_2} c_{a^{-1}(w_1)}^{-1}(w_2) \right| dw_2. \quad (\text{C.62})$$

By the diffeomorphic change of variable $z_2 = c_{a^{-1}(w_1)}^{-1}(w_2)$,¹³ this can be written as

$$r_1(w_1) = \left| \frac{d}{dw_1} a^{-1}(w_1) \right| q_1\left(a^{-1}(w_1)\right) \int q_2\left(z_2 \mid a^{-1}(w_1)\right) dz_2 \quad (\text{C.63})$$

$$= \left| \frac{d}{dw_1} a^{-1}(w_1) \right| q_1\left(a^{-1}(w_1)\right) \quad (\text{C.64})$$

Next, we carry out the same calculation for the pushforward of $P_{\mathbf{V}}^{e_0}$ by ρ :

$$r(\mathbf{w}) = p\left(\rho^{-1}(\mathbf{w})\right) \left| \det \mathbf{J}_{\rho^{-1}}(\mathbf{w}) \right| \quad (\text{C.65})$$

$$= p_1\left(d^{-1}(w_2)\right) p_2\left(b_{d^{-1}(w_2)}^{-1}(w_1) \mid d^{-1}(w_2)\right) \left| \frac{d}{dw_2} d^{-1}(w_2) \frac{d}{dw_1} b_{d^{-1}(w_2)}^{-1}(w_1) \right|, \quad (\text{C.66})$$

leading to the marginal

$$r_1(w_1) = \int p_1\left(d^{-1}(w_2)\right) p_2\left(b_{d^{-1}(w_2)}^{-1}(w_1) \mid d^{-1}(w_2)\right) \left| \frac{d}{dw_1} b_{d^{-1}(w_2)}^{-1}(w_1) \right| \left| \frac{d}{dw_2} d^{-1}(w_2) \right| dw_2 \quad (\text{C.67})$$

$$= \int p_1(v_1) p_2\left(b_{v_1}^{-1}(w_1) \mid v_1\right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dv_1, \quad (\text{C.68})$$

where the second line is obtained by the diffeomorphic change of variable $v_1 = d^{-1}(w_2)$.

Equating the two expressions for $r(w_1)$ in e_0 in (C.68) and (C.64), we obtain for all w_1 :

$$\left| \frac{d}{dw_1} a^{-1}(w_1) \right| q_1\left(a^{-1}(w_1)\right) = \int p_1(v_1) p_2\left(b_{v_1}^{-1}(w_1) \mid v_1\right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dv_1. \quad (\text{C.69})$$

Applying the same approach to the environment in which V_1 is intervened upon changing p_1 to \tilde{p}_1 while Z_2 is intervened upon leaving q_1 invariant, yields for all w_1 :

$$\left| \frac{d}{dw_1} a^{-1}(w_1) \right| q_1\left(a^{-1}(w_1)\right) = \int \tilde{p}_1(v_1) p_2\left(b_{v_1}^{-1}(w_1) \mid v_1\right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dv_1. \quad (\text{C.70})$$

¹³Note that: $\int q_2(z_2(w_2)) \left| \frac{dz_2}{dw_2} \right| dw_2 = \int q_2(z_2) dz_2$.

Finally, by equating (C.69) and (C.70), we arrive at the following expression which must hold for all w_1 :

$$\int p_1(v_1) p_2 \left(b_{v_1}^{-1}(w_1) \mid v_1 \right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dw_1 = \int \tilde{p}_1(v_1) p_2 \left(b_{v_1}^{-1}(w_1) \mid v_1 \right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dw_1 \quad (\text{C.71})$$

which we can rewrite as

$$\int (\tilde{p}_1(v_1) - p_1(v_1)) p_2 \left(b_{v_1}^{-1}(w_1) \mid v_1 \right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dw_1 = 0. \quad (\text{C.72})$$

Multiplying by any continuous function $\varphi(w_1)$, integrating w.r.t. w_1 and applying the diffeomorphic change of variable $v_2 = b_{v_1}^{-1}(w_1)$, this can be expressed as:

$$0 = \int \varphi(w_1) \int (\tilde{p}_1(v_1) - p_1(v_1)) p_2 \left(b_{v_1}^{-1}(w_1) \mid v_1 \right) \left| \frac{d}{dw_1} b_{v_1}^{-1}(w_1) \right| dv_1 dw_1 \quad (\text{C.73})$$

$$= \int \int \varphi(b_{v_1}(v_2)) (\tilde{p}_1(v_1) - p_1(v_1)) p_2(v_2 \mid v_1) dv_2 dv_1 \quad (\text{C.74})$$

$$= \int \int \varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 \mid v_1)} \right) (\tilde{p}_1(v_1) - p_1(v_1)) p_2(v_2 \mid v_1) dv_2 dv_1 \quad (\text{C.75})$$

where we have resubstituted the expression for $b_{v_1}(v_2)$ in the last line.

Equivalently, this can be written as: *for any continuous function φ ,*

$$\mathbb{E}_{v \sim P_V^{\varepsilon_0}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 \mid v_1)} \right) \right] = \mathbb{E}_{v \sim P_V^{\varepsilon_1}} \left[\varphi \left(\frac{\tilde{p}_2(v_2)}{p_2(v_2 \mid v_1)} \right) \right]. \quad (\text{C.76})$$

However, the genericity condition (A4) precisely rules this out, since the above equality must be violated for at least one φ , concluding this last case.

To sum up, all cases either lead to a contradiction, or imply the conclusion that $(f^{-1}, G) \sim_{\text{CRL}} (h, G')$, concluding the proof. \square

C.3 Proof of Thm. 3.4

Theorem 3.4 (Identifiability up to \sim_{CRL} from two paired perfect stochastic interventions per node). *Suppose that we have access to multiple environments $\{P_{\mathbf{X}}^e\}_{e \in \mathcal{E}}$ generated as described in § 2 under Asms. 2.2, 2.3, 2.5, 2.8 and 2.9. Let (h, G') be any candidate solution such that the inferred latent distributions $Q_{\mathbf{Z}}^e = h_*(P_{\mathbf{X}}^e)$ of $\mathbf{Z} = h(\mathbf{X})$ and the inferred mixing function h^{-1} satisfy the above assumptions w.r.t. the candidate causal graph G' . Assume additionally that*

- (A1) *all densities p^e and q^e are continuously differentiable and fully supported on \mathbb{R}^n ;*
- (A2') *we have access to at least one pair of single-node perfect interventions per node, with unknown targets: there exist $m \geq n$ known pairs of environments $\mathcal{E} = \{(e_j, e'_j)\}_{j=1}^m$ such that for each $i \in [n]$ there exists some unknown $j \in [m]$ for which $\mathcal{I}^{e_j} = \mathcal{I}^{e'_j} = \{i\}$;*
- (A3') *for all $i \in [n]$, the intervened mechanisms $\tilde{p}_i(v_i)$ and $\tilde{\tilde{p}}_i(v_i)$ differ everywhere, in the sense that*

$$\forall v_i : \quad \left(\frac{\tilde{\tilde{p}}_i}{\tilde{p}_i} \right)'(v_i) \neq 0; \quad (3.10)$$

Then the ground truth is identified in the sense of Defn. 2.6, that is, $(f^{-1}, G) \sim_{\text{CRL}} (h, G')$.

Proof. First, we show that we can extract from the $m \geq n$ available pairs of environments a suitable subset \mathcal{E}_n of exactly n pairs, containing one pair of interventional environments for each node.

Let $\mathcal{E}_n \subseteq \mathcal{E}$ be a subset of n pairs of environments which are assumed to correspond to distinct targets in the model q , and suppose for a contradiction that this is not actually the case for the ground truth p (i.e., there are duplicate and missing interventions w.r.t. p). Then there must be two pairs of environments $(e_a, e'_a), (e_b, e'_b) \in \mathcal{E}_n$, both corresponding to interventions on some V_i in p , but which

are modelled as interventions on distinct nodes Z_j and Z_k with $j \neq k$ in q . We show that this implies that V_i must simultaneously be a deterministic function of only Z_j and only Z_k . Similar to the proof of Thm. 3.2, we obtain the following equations,

$$\frac{\tilde{q}_j}{\tilde{q}_j}(z_j) = \frac{\tilde{p}_i}{\tilde{p}_i}(\psi_i(\mathbf{z})), \quad (\text{C.77})$$

$$\frac{\tilde{q}_k}{\tilde{q}_k}(z_k) = \frac{\hat{p}_i}{\hat{p}_i}(\psi_i(\mathbf{z})). \quad (\text{C.78})$$

By taking partial derivatives w.r.t. z_l and applying assumption (A3'), we find that

$$\frac{\partial \psi_i}{\partial z_l} = 0 \quad \forall l \neq j, \quad (\text{C.79})$$

$$\frac{\partial \psi_i}{\partial z_l} = 0 \quad \forall l \neq k. \quad (\text{C.80})$$

Since $j \neq k$, this implies that $\partial \psi_i / \partial z_l = 0$ for all l which contradicts invertibility of ψ . Thus, by contradiction, we find that \mathcal{E}_n must contain exactly one pair of intervention per node also w.r.t. p . For the remainder of the proof, we only consider \mathcal{E}_n .

W.l.o.g., for any $(e_i, e'_i) \in \mathcal{E}_n$ we now fix the intervention targets in p to $\mathcal{I}^{e_i} = \mathcal{I}^{e'_i} = \{i\}$ and let π be a permutation of $[n]$ such that $\pi(i)$ denotes the inferred intervention target in q that by (A2') is shared across (e_i, e'_i) . (We will show later that not all permutations are admissible, but only ones that preserve the partial order of G .)

The first part of the proof is similar to Case 1 in the proof of Thm. 3.2. Consider the densities in environments e_i and e'_i , which are related through the change of variable formula by:

$$\tilde{q}_{\pi(i)}(z_{\pi(i)}) \prod_{j \in [n] \setminus \{\pi(i)\}} q_j(z_j | \mathbf{z}_{\text{pa}(j; G')}) = \tilde{p}_i(\psi_i(\mathbf{z})) \prod_{j \in [n] \setminus \{i\}} p_j(\psi_j(\mathbf{z}) | \psi_{\text{pa}(j)}(\mathbf{z})) |\det \mathbf{J}_\psi(\mathbf{z})|, \quad (\text{C.81})$$

$$\tilde{q}_{\pi(i)}(z_{\pi(i)}) \prod_{j \in [n] \setminus \{\pi(i)\}} q_j(z_j | \mathbf{z}_{\text{pa}(j; G')}) = \tilde{p}_i(\psi_i(\mathbf{z})) \prod_{j \in [n] \setminus \{i\}} p_j(\psi_j(\mathbf{z}) | \psi_{\text{pa}(j)}(\mathbf{z})) |\det \mathbf{J}_\psi(\mathbf{z})|, \quad (\text{C.82})$$

where $\mathbf{Z}_{\text{pa}(j; G')} \subseteq \mathbf{Z} \setminus \{Z_j\}$ denotes the parents of Z_j in G' .

Taking the quotient of the two equations yields

$$\frac{\tilde{q}_{\pi(i)}}{\tilde{q}_{\pi(i)}}(z_{\pi(i)}) = \frac{\tilde{p}_i}{\tilde{p}_i}(\psi_i(\mathbf{z})). \quad (\text{C.83})$$

Next, for any $j \neq \pi(i)$, taking partial derivatives w.r.t. z_j on both sides yields

$$0 = \left(\frac{\tilde{p}_i}{\tilde{p}_i} \right)'(\psi_i(\mathbf{z})) \frac{\partial \psi_i}{\partial z_j}(\mathbf{z}). \quad (\text{C.84})$$

By assumption (A3'), the first term on the RHS is non-zero everywhere. Hence, (C.84) implies

$$\forall j \neq \pi(i), \forall \mathbf{z} : \frac{\partial \psi_i}{\partial z_j}(\mathbf{z}) = 0 \quad (\text{C.85})$$

from which we can conclude that

$$V_i = \psi_i(Z_{\pi(i)}) \quad (\text{C.86})$$

for all $i \in [n]$. That is, ψ is the composition of the permutation π with an element-wise reparametrisation.

It remains to show that π must, in fact, be a graph isomorphism, which is equivalent to the statement

$$V_i \rightarrow V_j \text{ in } G \iff Z_{\pi(i)} \rightarrow Z_{\pi(j)} \text{ in } G'. \quad (\text{C.87})$$

(\implies) Suppose for a contradiction that there exist (i, j) such that $V_i \rightarrow V_j$ in G , but $Z_{\pi(i)} \not\rightarrow Z_{\pi(j)}$ in G' .

The main idea is to demonstrate that the lack of such direct arrow implies a certain conditional independence which, by faithfulness, would contradict the unconditional dependence of V_i and V_j .

Consider environment e_i in which there are perfect interventions on $Z_{\pi(i)}$ and V_i , which has the effect of removing all incoming arrows to $Z_{\pi(i)}$ and V_i in the respective post-intervention graphs $G'_{\overline{Z_{\pi(i)}}}$ and $G'_{\overline{V_i}}$.

As a result of this and the lack of direct arrow by assumption, any d-connecting path between $Z_{\pi(i)}$ and $Z_{\pi(j)}$ must enter the latter via $\mathbf{Z}_{\text{pa}(\pi(j); G')}$ [94].

It then follows from Markovianity of q w.r.t. G' that the following holds in $Q_{\mathbf{Z}}^{e_i}$:

$$Z_{\pi(i)} \perp\!\!\!\perp Z_{\pi(j)} \mid \mathbf{Z}_{\text{pa}(\pi(j); G')}. \quad (\text{C.88})$$

We now consider the corresponding implication for $P_{\mathbf{V}}^{e_i}$. Define

$$\tilde{\mathbf{V}} = \left\{ V_k = \psi_k \left(Z_{\pi(k)} \right) : Z_{\pi(k)} \in \mathbf{Z}_{\text{pa}(\pi(j); G')} \right\} \subseteq \mathbf{V} \setminus \{V_i, V_j\}, \quad (\text{C.89})$$

and note that by assumption, $Z_{\pi(i)} \notin \mathbf{Z}_{\text{pa}(\pi(j); G')}$ and hence $V_i \notin \tilde{\mathbf{V}}$.

By applying the corresponding diffeomorphic functions ψ_i from (C.86) to (C.88), it follows from Lemma C.2 that

$$V_i \perp\!\!\!\perp V_j \mid \tilde{\mathbf{V}} \quad (\text{C.90})$$

in $P_{\mathbf{V}}^{e_i}$. However, this violates faithfulness (Asm. 2.2) of $P_{\mathbf{V}}$ to G since V_i and V_j are d-connected in $G_{\overline{V_i}}$.

Thus, by contradiction, we must have $Z_{\pi(i)} \rightarrow Z_{\pi(j)}$ in G' .

(\Leftarrow) Now, suppose for a contradiction that there exist (i, j) such that $Z_{\pi(i)} \rightarrow Z_{\pi(j)}$ in G' , but $V_i \not\rightarrow V_j$ in G .

By the same argument as before, we find that

$$V_i \perp\!\!\!\perp V_j \mid \mathbf{V}_{\text{pa}(j)} \quad (\text{C.91})$$

in $P_{\mathbf{V}}^{e_i}$, and thus by Lemma C.2

$$Z_{\pi(i)} \perp\!\!\!\perp Z_{\pi(j)} \mid \tilde{\mathbf{Z}} \quad (\text{C.92})$$

in $Q_{\mathbf{Z}}^{e_i}$ where

$$\tilde{\mathbf{Z}} = \left\{ Z_{\pi(k)} : V_k \in \mathbf{V}_{\text{pa}(j)} \right\} \subseteq \mathbf{Z} \setminus \{Z_{\pi(i)}, Z_{\pi(j)}\}.$$

However, this contradicts faithfulness of $Q_{\mathbf{Z}}$ to G' . Hence, we must have that $V_i \rightarrow V_j$ in G .

This shows that π must be a graph isomorphism, thus concluding the proof. \square

C.4 Proof of Thm. 4.2

Theorem 4.2 (Preservation of causal influences under \sim_{CRL}). *Let $P_{\mathbf{V}}$ be Markovian w.r.t. G , let π be a graph isomorphism of G , and let ϕ be an element-wise diffeomorphism. Let $\mathbf{Z} = \mathbf{P}_{\pi^{-1}} \circ \phi(\mathbf{V})$ and denote its induced distribution by $Q_{\mathbf{Z}}$. Then for any $V_i \rightarrow V_j$ in G we have $\mathfrak{C}_{i \rightarrow j}^{P_{\mathbf{V}}} = \mathfrak{C}_{\pi(i) \rightarrow \pi(j)}^{Q_{\mathbf{Z}}}$.*

Proof. First, recall that according to Defn. 4.1,

$$\mathfrak{C}_{i \rightarrow j}^{P_{\mathbf{V}}} := D_{\text{KL}}(P_{\mathbf{V}} \parallel P_{\mathbf{V}}^{i \rightarrow j}), \quad (\text{C.93})$$

where $P_{\mathbf{V}}^{i \rightarrow j}$ denotes the interventional distribution obtained by replacing $p_j(v_j \mid \mathbf{v}_{\text{pa}(j)})$ with

$$p_j^{i \rightarrow j}(v_j \mid \mathbf{v}_{\text{pa}(j) \setminus \{i\}}) = \int_{\mathbf{V}_i} p_j(v_j \mid \mathbf{v}_{\text{pa}(j)}) p_i(v_i) dv_i. \quad (\text{C.94})$$

Writing out the KL divergence and noting that all terms except the interved mechanism j cancel inside the log, we obtain

$$\mathfrak{E}_{i \rightarrow j}^{P_{\mathbf{V}}} = \int_{\mathcal{V}} \log \left(\frac{p_j(v_j | \mathbf{v}_{\text{pa}(j)})}{\int_{\mathcal{V}_i} p_j(v_j | \mathbf{v}_{\text{pa}(j)}) p_i(v_i) dv_i} \right) p(\mathbf{v}) d\mathbf{v}. \quad (\text{C.95})$$

and similarly

$$\mathfrak{E}_{\pi(i) \rightarrow \pi(j)}^{Q_{\mathbf{Z}}} = \int_{\mathcal{Z}} \log \left(\frac{q_{\pi(j)}(z_{\pi(j)} | \mathbf{z}_{\text{pa}(\pi(j); G')})}{\int_{\mathcal{Z}_{\pi(i)}} q_{\pi(j)}(z_{\pi(j)} | \mathbf{z}_{\text{pa}(\pi(j); G')}) q_{\pi(i)}(z_{\pi(i)}) dz_{\pi(i)}} \right) q(\mathbf{z}) d\mathbf{z}. \quad (\text{C.96})$$

Since $\mathbf{Z} = \mathbf{P}_{\pi^{-1}} \circ \phi(\mathbf{V})$, we have $V_i = \psi_i(Z_{\pi(i)})$ for all $i \in [n]$ where $\psi = \phi^{-1}$.

Thus, by the change of variable formula, and using the fact that $\pi(\text{pa}(i)) = \text{pa}(\pi(i); G')$ since $\pi : G \mapsto G'$ is a graph isomorphism, we have for all $i \in [n]$:

$$q_{\pi(i)}(z_{\pi(i)} | \mathbf{z}_{\text{pa}(\pi(i); G')}) = p_i(\psi_i(z_{\pi(i)}) | \psi_{\text{pa}(i)}(\mathbf{z}_{\text{pa}(\pi(i); G')})) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right|, \quad (\text{C.97})$$

as well as for the marginal density

$$q_{\pi(i)}(z_{\pi(i)}) = p_i(\psi_i(z_{\pi(i)})) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right|, \quad (\text{C.98})$$

and

$$q(\mathbf{z}) = p(\psi \circ \mathbf{P}_{\pi}(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})|. \quad (\text{C.99})$$

Substitution into the expression for $\mathfrak{E}_{\pi(i) \rightarrow \pi(j)}^{Q_{\mathbf{Z}}}$ yields:

$$\mathfrak{E}_{\pi(i) \rightarrow \pi(j)}^{Q_{\mathbf{Z}}} = \int_{\mathcal{Z}} \log \left(\frac{p_j(\psi_j(z_{\pi(j)}) | \psi_{\text{pa}(j)}(\mathbf{z}_{\text{pa}(\pi(j); G')}))}{\int_{\mathcal{Z}_{\pi(i)}} p_j(\psi_j(z_{\pi(j)}) | \psi_{\text{pa}(j)}(\mathbf{z}_{\text{pa}(\pi(j); G')})) p_i(\psi_i(z_{\pi(i)})) \left| \frac{d\psi_i}{dz_{\pi(i)}}(z_{\pi(i)}) \right| dz_{\pi(i)}} \right) \quad (\text{C.100})$$

$$p(\psi \circ \mathbf{P}_{\pi}(\mathbf{z})) |\det \mathbf{J}_{\psi}(\mathbf{z})| d\mathbf{z}. \quad (\text{C.101})$$

$$= \int_{\mathcal{V}} \log \left(\frac{p_j(v_j | \mathbf{v}_{\text{pa}(j)})}{\int_{\mathcal{V}_i} p_j(v_j | \mathbf{v}_{\text{pa}(j)}) p_i(v_i) dv_i} \right) p(\mathbf{v}) d\mathbf{v} \quad (\text{C.102})$$

$$= \mathfrak{E}_{i \rightarrow j}^{P_{\mathbf{V}}}. \quad (\text{C.103})$$

where the second to last line follows by integration by substitution, applied to both integrals. \square

D Experimental Details and Additional Results

In this appendix, we describe the experiments presented in § 6 in more details (Appx. D.1), and present additional results (Appx. D.2).

D.1 Experimental Details for § 6

Synthetic Data Generating Process. We consider linear Gaussian latent SCMs of the form

$$V_1 := U_1, \quad V_2 := \alpha V_1 + U_2, \quad (\text{D.1})$$

with standard normal U_1 and U_2 . As a mixing function, we use a three-layer multilayer perceptron (MLP),

$$f = \sigma \circ \mathbf{W}_3 \circ \sigma \circ \mathbf{W}_2 \circ \sigma \circ \mathbf{W}_1 \quad (\text{D.2})$$

where $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3 \in \mathbb{R}^{2 \times 2}$ are invertible weight matrices, and σ is an element-wise invertible nonlinear leaky-tanh activation function used in [41]:

$$\sigma(x) = \tanh(x) + 0.1x. \quad (\text{D.3})$$

To compute averages of our results over multiple runs, we construct different ground truth data generating processes as follows. We generate different latent SCMs by drawing α uniformly from $[-10, -2] \cup [2, 10]$. (We exclude $(-2, 2)$ to avoid sampling near unfaithful models.) We generate the corresponding mixing functions by uniformly sampling each element of the weight matrices, $(\mathbf{W}_k)_{ij} \sim U(0, 1)$. (To avoid the sampled weight matrices being too close to singular, we reject and resample if $|\det \mathbf{W}_k| < 0.1$.)

Interventional Environments. In line with Thm. 3.2, for each choice of latent SCM and mixing function, we generate three environments: one observational environment and one interventional environment for each perfect single-node intervention. For $i = 1, 2$, we model a perfect intervention on V_i by removing the influence of the parent variables and changing the exogenous noise by shifting its mean up or down. Specifically, we replace the corresponding assignment in (D.1) by

$$V_i := \tilde{U}_i, \quad \text{where } \tilde{U}_i \sim \mathcal{N}(m_i, 1) \quad (\text{D.4})$$

where the mean m_i of the shifted Gaussian noise is fixed per environment and sampled uniformly from $\{\pm 2\}$.

We label the observational environment as $e = 0$ and the environment arising from intervention on V_i by $e = i$ for $i = 1, 2$. Samples from p^e are then generated by sampling latents \mathbf{v} from the respective (un)intervened SCM and then applying the mixing function.

Model Architecture. We use normalizing flows [92] to model observations \mathbf{x} as the result of an invertible, differentiable transformation g of some latent (noise) variable \mathbf{z} ,

$$\mathbf{x} = g(\mathbf{z}). \quad (\text{D.5})$$

We apply a series of L such transformations $g^l : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $g = g^L \circ \dots \circ g^1$ which we refer to as *flow layers*. We use Neural Spline Flows [30] for the invertible transformation, with a 3-layer feedforward neural network with hidden dimension 128 and permutation in each flow layer and $L = 12$ layers. The transformations g, g^1, \dots, g^L have learnable parameters (the weights and biases of the neural networks), which we omit to simplify notation.

Typically, simple distributions such as a uniform or isotropic Gaussian are used as base distribution $q(\mathbf{z})$ in normalizing flows. Here, we instead choose a base distribution that encodes information about the latent SCM. Specifically, we model the base mechanism as

$$q_1(z_1) = \mathcal{N}(\mu_1, \sigma_1^2), \quad q_2(z_2 | z_1) = \mathcal{N}(\hat{\alpha}z_1, \sigma_2^2), \quad q_2(z_2) = \mathcal{N}(\mu_2, \hat{\sigma}_2^2) \quad (\text{D.6})$$

and the intervened mechanism as

$$\tilde{q}_1(z_1) = \mathcal{N}(\tilde{\mu}_1, \tilde{\sigma}_1^2), \quad \tilde{q}_2(z_2) = \mathcal{N}(\tilde{\mu}_2, \tilde{\sigma}_2^2). \quad (\text{D.7})$$

Candidate Graphs and Intervention Targets. We train a separate normalizing-flow based model for each choice of candidate graph G' and inferred intervention targets. For the bivariate case with $n = 2$, this gives rise to four models, depending on whether G' matches G or not, and whether the intervention targets are aligned or misaligned w.r.t. the ground truth intervention targets. To model the setting $G' \neq G$ in which Z_1 and Z_2 are assumed independent, we use $q_2(z_2)$ in place of $q_2(z_2 | z_1)$ in (D.6). If the intervention targets are aligned, we use \tilde{q}_i instead of q_i in $e = i$ for $i = 1, 2$. Else, if they are misaligned, we use \tilde{q}_2 instead of q_2 in $e = 1$ and \tilde{q}_1 instead of q_1 in $e = 2$. By multiplying the respective mechanisms, we thus obtain three environment-specific joint base distributions $q^e(\mathbf{z})$ for $e = 0, 1, 2$.

Learning Objective. Given multi-environment data, the parameters $\mu_1, \sigma_1, \hat{\alpha}, \sigma_2, \mu_2, \hat{\sigma}_2, \tilde{\mu}_2, \tilde{\sigma}_2, \tilde{\mu}_1$ and $\tilde{\sigma}_1$ are jointly learned with the parameters of the invertible transformations g^l by maximising the log-likelihood of the data under our model, which is given by:

$$\sum_{e \in \mathcal{E}} \mathbb{E}_{\mathbf{x} \sim p^e(\mathbf{x})} [\log p_{\text{model}}^e(\mathbf{x})] = \sum_{e \in \mathcal{E}} \mathbb{E}_{\mathbf{x} \sim p^e(\mathbf{x})} [\log q^e(h(\mathbf{x})) + \log |\det \mathbf{J}_h(\mathbf{x})|] \quad (\text{D.8})$$

where the encoder $h := g^{-1}$ is the inverse of the normalizing flow which is readily available by construction; and where the expectations are empirical averages over the respective datasets in practice.

Training and Model Selection Details. Each environment comprises a total of 200k data points. We use the ADAM optimizer [66] with cosine annealing learning rate scheduling, starting with a learning rate of 5×10^{-3} and ending with 1×10^{-7} . We train the model for 200 epochs with a batch size of 4096. We split the dataset into 70% for training, and 15% for validation and held-out test data, each sampled randomly across all environments. For each drawn data generating process, we train three versions of each model with different random initializations and select the one with the highest validation log likelihood at the end of training for evaluation.

Evaluation Metrics. We evaluate the trained models w.r.t. *mean correlation coefficient* (MCC) on held-out data and *log-likelihood* on validation data (for model selection).

- The MCC measures the extent to which there is a one-to-one correspondence between the ground truth latents V_i and (a permuted version of) the inferred latents $Z_i = h_i(\mathbf{X})$. Its maximum value of one indicates a perfect correlation between the two. MCC is thus a proxy measure for the level of identifiability up to permutation and invertible reparametrisation. We report MCC based on Pearson (linear) correlation, though we found the results based on Spearman (nonlinear monotonic) correlation to be almost identical.
- The log-likelihood, on the other hand, measures how well a model explains or fits the data. Since the ground truth is typically unknown, a reasonable procedure when training multiple models is to select the one that attains the highest likelihood. For this reason, we report the difference in log-likelihood between misspecified models (ones assuming a wrong graph or intervention targets) to the correctly specified model. Whenever this difference is larger than zero, the correct model fits the data better and would thus be selected.

D.2 Additional Results: Learning Nonlinear Latent SCMs from Partial Causal Order

In this subsection, we present an additional experiment, in which we extend the setting investigated in § 6 and Appx. D.1 along the following axes.

- We fit generative models over three instead of two variables, corresponding to the setting of Thm. 3.4.
- The ground-truth SCM is now given by nonlinear mechanisms with non-additive, non-Gaussian noise.
- The generative model, including the learnt mechanisms, is now fully nonlinear.
- Despite Thm. 3.4 formally requiring two environments per single-node intervention, we only provide one interventional environment per node.
- Rather than searching over candidate graphs, we only fix the causal order and fit the reduced form of the SCM (see § 2.1) with a second normalizing flow.

Below, we describe these differences in more detail.

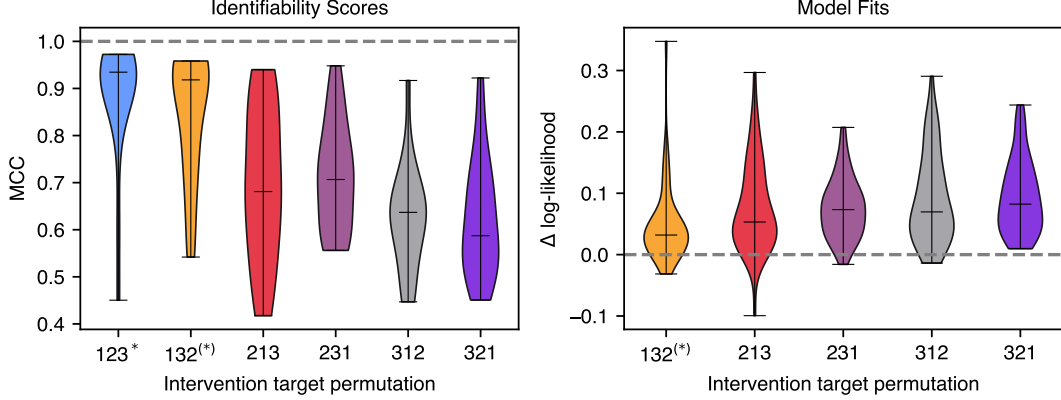


Figure 4: **Comparison of Correctly and Incorrectly Specified Models for $V_2 \leftarrow V_1 \rightarrow V_3$ with Fixed Causal Order and Nonlinear SCM.** Each violinplot corresponds to one setting where the intervention target labels are permuted. The blue plot (123^{*}) is the setting with correct intervention target labels. The yellow plot (132^(*)) has the targets for the two children V_2 and V_3 permuted, which also corresponds to a correct causal ordering and should thus be considered equivalent. We show mean correlation coefficients (MCCs) between the learned and ground truth latents (*Left*) and the difference in validation model log-likelihood between the well-specified (blue) and misspecified models (*Right*). Each violin plot is based on 20 different ground truth data generating processes; the horizontal lines indicate the minimum, median and maximum values.

Three-Variable Graph. The *unknown* ground truth graph is given by

$$V_2 \leftarrow V_1 \rightarrow V_3. \quad (\text{D.9})$$

This is consistent with the partial ordering $V_1 \preceq V_2 \preceq V_3$, which is assumed for all models a priori w.l.o.g., see § 2.2. Note that, due to the encoding of causal structure in the nonparametric model explained below, we only iterate over different permutations of the intervention targets and not over latent graph configurations. Due to the causal order implied by the graph (D.9), the permutations (1, 2, 3) (no permutation) and (1, 3, 2) (permutation of the two effects) are equivalent since the latter also implies the correct causal ordering.

Nonlinear, Non-Gaussian SCM. The mechanisms in the ground-truth SCM are now given by

$$V_i := \beta f_i^{\text{loc}}(\mathbf{V}_{\text{pa}(i)}) + f_i^{\text{scale}}(\mathbf{V}_{\text{pa}(i)})U_i \quad (\text{D.10})$$

for all i , where the location and scale functions $f_i^{\text{loc}}, f_i^{\text{scale}} : \mathbb{R}^{|\text{pa}(i)|} \rightarrow \mathbb{R}$ are parameterized by random 3-layer neural networks (sampled as the random mixing function in (D.2)) and the noise variables are Gaussian, $U_i \sim \mathcal{N}(0, 1)$. The factor β controls the influence of the parent variables relative to the exogenous noise. As β increases, variables tend to become more dependent, as also the mean shifts as a function of the parent variables. We set $\beta = 10$ for the experiments shown in Fig. 4.

Nonparametric Latent SCM. We use a second normalizing flow to learn a *reduced form of the latent SCM* via the transformation $g^{\text{SCM}} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ mapping an exogenous noise variable ϵ to the latent variable \mathbf{z} ,

$$\mathbf{z} = g^{\text{SCM}}(\epsilon). \quad (\text{D.11})$$

The distribution of the exogenous noise variable ϵ as well as the distribution of the intervened mechanisms $\tilde{q}_i(z_i)$ for $i = 1, 2, 3$ is fixed and standard (isotropic) Gaussian. The flow layers in g^{SCM} have an *upper triangular Jacobian* and thus allow us to encode assumptions about the causal graph: by passing the variables in topological order, which we can assume w.l.o.g., we ensure that an exogenous noise variable ϵ_i can only influence endogenous variables in \mathbf{z} that are descendants of z_i . The learned weights of the flow layers then implicitly encode which endogenous variables are connected. Therefore, only different choices of the permutations of the intervention targets need to be considered as candidate models. We use a similar architecture based on Neural Spline Flows. However, we omit permutation layers, which would violate the topological order of the variables.

Results. In Fig. 4, we present identifiability scores and model fits for both well-specified and misspecified models (corresponding to different intervention target choices). Notably, we observe that the well-specified model (in blue) or its equivalent (in yellow) yield the highest log-likelihood in the majority of cases, as depicted in Fig. 4 (Right). This demonstrates that, even in this nonparametric setting without fully specified graph, the log-likelihood remains a reliable criterion for selecting the correct intervention targets. Fig. 4 (Left) shows that the selected models (blue or yellow) approximately identify the ground-truth latent variables up to element-wise rescaling, whereas other choices lead to much lower MCCs.

It is worth noting that, compared to the parametric setting investigated in § 6 and Fig. 3, the nonparametric setting appears to be more challenging (as expected), as there is a less pronounced distinction between well-specified and misspecified models, both in terms of identifiability scores and model fits. Moreover, future work is needed to parse the implicitly learned causal relationships in the transformation g^{SCM} in (D.11): since only the (pre-imposed) causal order is specified, in practice, g^{SCM} may learn to use additional or fewer edges than in the true graph G .

E Discussion of the Role of Our Assumptions

Below, we summarize the rationale and intuition behind each assumption:

- Asm. 2.2 is required to rule out degenerate cases (cancellation along different paths) in which variables are (conditionally) independent despite being causally related. It is a standard assumption in classical causal discovery, and therefore also needed in CRL to recover the causal graph.
- Asm. 2.3 is required to know how many latent variables we are looking for. It is a standard assumption in identifiable representation learning (that is often made implicitly). However, it may be dropped when suitable techniques for estimating the intrinsic dimensionality of \mathcal{X} can be employed.
- Asm. 2.5 is needed for the mapping between latents and observations to be invertible in the first place. Without it, full recovery of the causal variables (up to CRL equivalence) is infeasible. This assumption is also standard for the simpler problem of nonlinear ICA.
- Asm. 2.8 is a characterisation of our generative setup. Sharing of some mechanisms and the mixing function is needed for the multi-environment setting to provide useful additional information: if everything may change across environments, the datasets can only be analysed in isolation, running into the non-identifiability of CRL from iid data.
- Asm. 2.9 and (A2) / (A2') are needed since with imperfect interventions or interventions not on all nodes, identifiability is not achievable even in the linear setting as shown by Squires et al. [116].
- Asm. (A1) is a technical assumption needed for our analysis. It is not strictly necessary (it can also be relaxed to fully supported on a Cartesian product of intervals) but substantially eases the readability and accessibility of the proof, without a major impact on the main causal aspects of the problem setup.
- Asm. (A3) / (A3') is needed to avoid spurious solutions based on applying a measure preserving transformation on a part of the domain unaffected by the intervention.
- Asm. (A4) is needed to rule out a fine-tuning of the ground-truth generating process that are possible due to fully non-parametric nature of the setup, see also Remark 4.2 and the following paragraph.