

On Discovery of Local Independence over Continuous Variables via Neural Contextual Decomposition

Author names withheld

Editor: Under Review for CLear 2023

Abstract

Conditional independence provides a way to understand causal relationships among the variables of interest. An underlying system may exhibit more fine-grained causal relationships especially between a variable and its parents, which will be called the local independence relationships. One of the most widely studied local relationships is Context-Specific Independence (CSI), which holds in the specific assignment of conditioned variables. However, its applicability is often limited since it does not allow continuous variables: data conditioned to the specific value of a continuous variable contains few instances, if not none, making it infeasible to test independence. In this work, we focus on the local independence relationship which holds in a specific set of joint assignments of parental variables. We aim to resolve a challenging problem of discovering such local independences of continuous variables by proposing a novel method, coined neural contextual decomposition (NCD), which partitions the joint outcome space and imposes each partition to induce the local independence via modeling a conditional distribution. We empirically demonstrated that the proposed method successfully discovers the ground truth local independences in both static and sequential settings.

Keywords: Context-Specific Independence, Local Independence, Causal Discovery

1. Introduction

Discovering the causal relationships in a system is an important and challenging problem in many areas of scientific research such as social science (Sobel, 1995), biology (Shipley, 2016), and economics (Angrist et al., 1996; Angrist and Pischke, 2008; Imbens and Rubin, 2015; Banerjee et al., 2016; Imbens, 2020). There have been many causal discovery algorithms finding causal relationships given observational data (Spirtes et al., 2000; Chickering, 2002; Hoyer et al., 2008; Shimizu et al., 2006; Zhang et al., 2011; Zheng et al., 2018; Zhu et al., 2019). These methods often exploit conditional independence either explicitly (constraint-based methods) or implicitly (score-based or gradient-based).

A system often exhibits more fine-grained relationships between a variable and its parents, i.e., local independence relationship. For example when pushing an object on the ground, it will move only when a force exceeds the friction which is determined by the mass and the ground. Context-specific independence (CSI) (Poole and Zhang, 2003; Boutilier et al., 2013) is an independence that holds in a *certain* conditioning value (i.e., *context*) as opposed to *any* value of the conditioned variables. It has been shown that such independence leads to more efficient probabilistic inference by exploiting the underlying local structure (Poole, 1998; Poole and Zhang, 2003; Gogate and Domingos, 2010; Van den Broeck et al., 2011; Dal et al., 2018). Further, it allows the identification of causal effects, which would not be possible without CSI relationships (Tikka et al., 2019; Robins et al., 2020).

Despite the fact that many real-world scenarios involve continuous variables, most prior works on local independence relationships assumed that variables are discrete. This is partly due to the




Figure 1: (a) Causal graph in Example 1. (b, c) CSI can represent the local independence of Y and X_2 , but not the other. In contrast, CSSI is able to represent both. (d) Augmented causal graph with PIV Z added in Example 4. (e, f) PIV representing CSSIs for each context set.

notion of CSI that is inherently suited for discrete variables as conditioning on the specific value of a continuous variable is infeasible, where the resulting subset of data would be practically empty. Hence, it is non-trivial to explore how local independence for continuous variables can be empirically discovered and utilized in probabilistic or causal inferential tasks.

Consider a system as an example consisting of three observed variables X_1 , X_2 , Y , and an unobserved one U . We describe the system using a structural causal model (SCM) (Pearl, 2009; Peters et al., 2017). Let X_1, X_2 follow a uniform distribution between 0 and 1 and U follow a standard Gaussian distribution. Further, let Y be determined as $Y = X_1 + U$ if $X_1 X_2 < 1/2$ and $Y = X_2 + U$ otherwise. Focusing on the observable variables, while Y depends on both X_1 and X_2 overall (Fig. 1(a)), (i) Y depends only on X_1 under a certain condition and (ii) on X_2 , otherwise. Since $X_1 < 1/2$ implies the condition $X_1 X_2 < 1/2$ and Y being the function of X_1 and U , it can be represented as CSI (Fig. 1(b)). On the other hand, CSI cannot capture the other local independence where Y being determined only by X_2 and U since Y and X_1 are always dependent given any assignment of X_2 .

Against this background, we propose context-set specific independence, a generalized notion of local independence where conditional independence holds in a certain *context set*, a set of joint assignments of the variables. We characterize context-set specific independence in order to uncover CSSIs for continuous variables, which is more challenging due to an infinite number of context sets. Our characterization leads to the design of NCD (Neural Contextual Decomposition) to find the partitions of the joint outcome space where each partition set induces a CSSI. Our approach is based on approximating a conditional distribution for each partition given a subset of parents. For instance, based on the above example, our method learns two partitions of the Cartesian product of the domains of X_1 and X_2 , where each partition is flexibly modeled as $\hat{P}(Y|X_1)$ and $\hat{P}(Y|X_2)$.

Our contributions are summarized as follows. (i) We introduce the notion of context-set specific independence (CSSI), a new class of local conditional independence relationships. This generalizes two well-known local independence, CSI and PCI, and permits continuous conditioned variables. (ii) We characterize context-set specific independence by relating possible CSSIs arisen from a given distribution. We examine a canonical representation of CSSIs and provide a sufficient condition under which a unique CSSIs exists. These characterization leads to defining contextual decomposition, a way to understand an entire group of CSSIs presenting in a given dataset. (iii) We devise neural contextual decomposition, a simple and effective method for finding a contextual decomposition without directly testing for individual local independences. It utilizes an auxiliary variable (partition indicator variable) used for training of a neural network for context decomposition. Empirical

evaluations on a synthetic dataset and Spriteworld demonstrate the effectiveness of the proposed method in recovering local independences.

2. Preliminaries

Throughout this paper, we use capital letters for random variables and small letters for their assignments. Bold letters denote the set of random variables or its assignments. Calligraphy letters can be the domain of corresponding variables or complex mathematical objects.

2.1. Structural Causal Model

We adopt structural causal model (SCM) (Pearl, 2009) as a causal framework to understand generated data. An SCM $\mathcal{M}$ is defined as a tuple $\langle \mathbf{V}, \mathbf{U}, \mathbf{F}, P(\mathbf{U}) \rangle$, where $\mathbf{V} = \{V_1, \dots, V_d\}$ and $\mathbf{U}$ are a set of endogenous and exogenous variables, respectively. $\mathbf{F} = \{f_1, \dots, f_d\}$ is a set of functions determining each endogenous variable, i.e., $V_j \leftarrow f_j(Pa(j), \mathbf{U}_j)$ where $Pa(j) \subseteq \mathbf{V} \setminus \{V_j\}$ is a parent of V_j and $\mathbf{U}_j \subseteq \mathbf{U}$. We may use variable V_i and its index i interchangeably. $\mathcal{V}_j$ is the outcome space of each variable V_j . In this work, we restrict to a Markovian model¹ and assume that an SCM satisfies a structural minimality (Pearl, 2009; Peters et al., 2017), which asserts that $j \in Pa(i)$ if and only if there exist two values for j that lead to different values for i , i.e., $f_i(v_j, \mathbf{v}_{Pa(i) \setminus \{j\}}, \mathbf{u}_i) \neq f_i(v'_j, \mathbf{v}_{Pa(i) \setminus \{j\}}, \mathbf{u}_i)$ where $v_j, v'_j \in \mathcal{V}_j$ and $\mathbf{v}_{Pa(i) \setminus \{j\}} \in \mathcal{V}_{Pa(i) \setminus \{j\}}$. An SCM $\mathcal{M}$ induces a causal graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$, which is a directed acyclic graph (DAG), where $\mathbf{V} = \{1, \dots, d\}$ and $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$ are the set of nodes and edges, respectively. If $i \in Pa(j)$ in $\mathcal{M}$, then a directed edge (i, j) is in $\mathbf{E}$. In the language of SCM, each node i corresponds to a random variable V_i and each edge (i, j) denotes a direct causal relationship from V_i to V_j .

2.2. Local Independence Relationship of Discrete Variables

As described earlier, in many cases, the causal mechanism exhibits a fine-grained independence relationships, which do not hold in general but only under certain conditions, i.e., *local independence*. To begin with, we first provide the definition of context-specific independence (CSI) and partial conditional independence (PCI). We adopt the definitions of CSI and PCI from prior works (Boutilier et al., 2013; Pensar et al., 2015, 2016) for the relationship between a variable and its parents. For the definitions of CSI and PCI, variables are assumed to be discrete. Let $\mathbf{X}$ be a non-empty set of parents of Y and $\mathbf{X}_A, \mathbf{X}_B$ be non-empty partitions of $\mathbf{X}$, i.e., $\mathbf{X}$ is disjoint union of $\mathbf{X}_A$ and $\mathbf{X}_B$.

Definition 1 (Context-Specific Independence (CSI)) Y is said to be contextually independent of $\mathbf{X}_B$ given the context $\mathbf{X}_A = \mathbf{x}_A$ if $P(y | \mathbf{x}_A, \mathbf{x}_B) = P(y | \mathbf{x}_A)$, holds for all $y \in \mathcal{Y}$ and $\mathbf{x}_B \in \mathcal{X}_B$ whenever $P(\mathbf{x}_A, \mathbf{x}_B) > 0$. This will be denoted by $Y \perp\!\!\!\perp \mathbf{X}_B | \mathbf{X}_A = \mathbf{x}_A$.

Definition 2 (Partial Conditional Independence (PCI)) Y is said to be partially conditionally independent of $\mathbf{X}_B$ in the domain $\mathcal{D}_B \subseteq \mathcal{X}_B$ given the context $\mathbf{X}_A = \mathbf{x}_A$ if $P(y | \mathbf{x}_A, \mathbf{x}_B) = P(y | \mathbf{x}_A, \mathbf{x}'_B)$, holds for all $y \in \mathcal{Y}$ and $\mathbf{x}_B, \mathbf{x}'_B \in \mathcal{D}_B$ whenever $P(\mathbf{x}_A, \mathbf{x}_B) > 0$ and $P(\mathbf{x}_A, \mathbf{x}'_B) > 0$. This will be denoted by $Y \perp\!\!\!\perp \mathbf{X}_B | \mathcal{D}_B, \mathbf{X}_A = \mathbf{x}_A$.

1. Conditional independence is statistical notion and would be irrelevant to whether the model is Markovian or semi-Markovian. However, Markovian models will allow clearer interpretation of context-specific independence with respect to the functional form of the target variable irrespective of regimes, i.e., observational or interventional.

PCI is a more fine-grained notion than CSI and, thus, PCI does not imply CSI: $P(y \mid \mathbf{x}_A, \mathbf{x}_B) = P(y \mid \mathbf{x}_A, \mathbf{x}'_B)$ for all $\mathbf{x}_B, \mathbf{x}'_B \in \mathcal{D}_B$ is not equivalent to $P(y \mid \mathbf{x}_A, \mathbf{x}_B) = P(y \mid \mathbf{x}_A)$.

3. Compactly Representing Local Independences of Continuous Variables

Our goal is to find such local independence relationships in a system which may contain continuous variables. The information about local independence relationships of continuous variables can also be leveraged in various tasks as demonstrated for the case of discrete variables. Given that a naive generalization of CSI or PCI to continuous variables seems practically implausible, we first define an alternative notion of local independence, which we call context-set specific independence (CSSI). Then, we characterize CSSI and describe how it can be equipped within the framework of SCM. All omitted proofs can be found in the Appendix.




Figure 2: Example 1.

3.1. Local Independences of Continuous Variables

Here, we focus on a causal mechanism of a target variable Y and its parents $\mathbf{X} = \{X_1, \dots, X_d\}$, i.e., SCM $\mathcal{M} = \langle \mathbf{V}, \mathbf{U}, \mathbf{F}, P(\mathbf{U}) \rangle$ where $\{X_1, \dots, X_d, Y\} \subseteq \mathbf{V}$ is a set of continuous random variables and $Y = f(\mathbf{X}, U)$. We denote $\mathbf{X}_A = \{X_i \mid i \in A\}$ where $A \subseteq Pa(Y) = \{1, \dots, d\}$. Similarly, $\mathbf{X}_{A^c} = \{X_i \mid i \in A^c\}$ where $A^c = Pa(Y) \setminus A$. Throughout the paper, we assume strictly positive densities. We first define a context-set specific independence (CSSI).

Definition 3 (Context-Set Specific Independence) Let $\mathbf{X} = \{X_1, \dots, X_d\}$ be a non-empty set of the parents of Y in a causal graph. and $\mathcal{E} \subseteq \mathcal{X}$ be an event with a positive probability. $\mathcal{E}$ is said to be a **context set** which induces **context-set specific independence (CSSI)** of $\mathbf{X}_{A^c}$ from Y if $p(y \mid \mathbf{x}_{A^c}, \mathbf{x}_A) = p(y \mid \mathbf{x}'_{A^c}, \mathbf{x}_A)$ holds for every $(\mathbf{x}_{A^c}, \mathbf{x}_A), (\mathbf{x}'_{A^c}, \mathbf{x}_A) \in \mathcal{E}$. This will be denoted by $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$.

CSSI extends CSI and PCI in a sense that the independence holds in a set of conditioned values in CSSI since they consider a point-based condition. We revisit the earlier example to illustrate how CSI, PCI, and CSSI (Defs 1 to 3) represents the local independence relationships in the system (insufficiency results for CSI and PCI in representing Example 1 is depicted in Appendix A.1.)

Example 1 Let X_1, X_2 be a uniform random variable defined on $[0, 1]$, s.t. $X_1, X_2 \sim \text{Unif}[0, 1]$ and U be an exogenous variable. Let Y be:

$$Y = \begin{cases} X_1 + U & \text{if } X_1 X_2 < 1/2, \\ X_2 + U & \text{if } X_1 X_2 \geq 1/2. \end{cases}$$

Let $\mathcal{E} = \{(x_1, x_2) \mid x_1 x_2 < 1/2\}$. Then, $Y \perp\!\!\!\perp X_2 \mid X_1, \mathcal{E}$ and $Y \perp\!\!\!\perp X_1 \mid X_2, \mathcal{E}^c$ hold. On the other hand, $Y \not\perp\!\!\!\perp X_1 \mid X_2 = c$ for any c .

As described earlier, we focus on the discovery of local independences (i.e., CSSI) of continuous variables and henceforth assume variables are continuous. In case where the variables are discrete, CSSI can be discovered by first (i) discovering PCI relationships and then (ii) integrating them (see Appendix A.1 for the detail). We now introduce important properties of CSSI.

Proposition 4 (CSSI Entailment) Suppose a CSSI relationship $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ holds. Then, the following CSSI relationships also hold:

$$(i) Y \perp\!\!\!\perp \mathbf{X}_{B^c} \mid \mathbf{X}_B, \mathcal{E} \quad \text{for any } B \supseteq A, \quad (ii) Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{F} \quad \text{for any } \mathcal{F} \subseteq \mathcal{E}.$$

Due to such implications, we are interested in CSSI with *minimal* conditioned variables.

Definition 5 A CSSI of the form $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ is said to be **regular** if there does not exist any $B \subsetneq A$ such that $Y \perp\!\!\!\perp \mathbf{X}_{B^c} \mid \mathbf{X}_B, \mathcal{E}$ holds. Then, we say $\mathbf{X}_A$ to be a **local parent set** of Y in $\mathcal{E}$.²

Among the CSSIs with the same context set, the regular CSSIs are the most *informative*, in a sense that the set of the conditioned variables cannot be further reduced. We say a CSSI **trivial** if $A = Pa(Y)$ (i.e., $A^c = \emptyset$ and $\mathbf{X} = \mathbf{X}_A$), since it trivially holds for any $\mathcal{E} \subseteq \mathcal{X}$. A trivial CSSI on $\mathcal{E} = \mathcal{X}$ is indeed regular. In general, regular CSSI may not be unique on the given context set. We show the non-uniqueness of local parent sets by an example below.




Figure 3: Example 2.

Example 2 Let $X_1, X_2, X_3 \sim \text{Unif}[0, 1]$ and U be an exogenous variable. Let Y be:

$$Y = \begin{cases} X_3 + U & \text{if } X_1, X_2 < 1/2, \\ X_3 + 2 \cdot U & \text{if } X_1, X_2 \geq 1/2, \\ X_1 + X_2 + X_3 + U & \text{otherwise.} \end{cases}$$

Let $\mathcal{E}_1 = \{\mathbf{x} \mid x_1, x_2 < 1/2\}$, $\mathcal{E}_2 = \{\mathbf{x} \mid x_1, x_2 \geq 1/2\}$. Then, regular CSSIs $Y \perp\!\!\!\perp X_{12} \mid X_3, \mathcal{E}_1$ and $Y \perp\!\!\!\perp X_{12} \mid X_3, \mathcal{E}_2$ hold. Also, the following regular CSSIs hold on $\mathcal{E}_1 \cup \mathcal{E}_2$:

$$Y \perp\!\!\!\perp X_1 \mid X_{23}, \mathcal{E}_1 \cup \mathcal{E}_2 \quad \text{and} \quad Y \perp\!\!\!\perp X_2 \mid X_{13}, \mathcal{E}_1 \cup \mathcal{E}_2.$$

In the example above, there exist two distinct regular CSSIs on $\mathcal{E}_1 \cup \mathcal{E}_2$. If the local parent set is not unique on the given context set, discovering local parent set would yield an inconsistent result. Thus, we explore a sufficient condition that guarantees the uniqueness of local parent set. We now provide an intersection property of CSSI, which will be used to derive a sufficient condition.

Proposition 6 (Intersection Property of CSSI) Suppose $\mathcal{E} \subseteq \mathcal{X}$ is convex. If CSSI relationships $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ and $Y \perp\!\!\!\perp \mathbf{X}_{B^c} \mid \mathbf{X}_B, \mathcal{E}$ hold, then $Y \perp\!\!\!\perp \mathbf{X}_{(A \cap B)^c} \mid \mathbf{X}_{A \cap B}, \mathcal{E}$ hold.

In Example 2, we cannot derive $Y \perp\!\!\!\perp X_{12} \mid X_3, \mathcal{E}_1 \cup \mathcal{E}_2$ from the two regular CSSIs on $\mathcal{E}_1 \cup \mathcal{E}_2$, since $\mathcal{E}_1 \cup \mathcal{E}_2$ is not convex. We now introduce a sufficient condition which guarantees the uniqueness of regular CSSIs given a context set $\mathcal{E}$.

Theorem 7 (Uniqueness of Local Parent Set) Suppose $\mathcal{E} \subseteq \mathcal{X}$ is convex. There exists a unique $A \subseteq Pa(Y)$ such that the regular CSSI relationship $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ holds.

Proof Suppose regular CSSIs $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ and $Y \perp\!\!\!\perp \mathbf{X}_{B^c} \mid \mathbf{X}_B, \mathcal{E}$ hold for some A and B . By Prop. 6 and the convexity of $\mathcal{E}$, $(Y \perp\!\!\!\perp \mathbf{X}_{(A \cap B)^c} \mid \mathbf{X}_{A \cap B}, \mathcal{E})$ holds true. Due to the minimality (i.e., regular), $A \cap B$ cannot be a proper subset of A nor B . Hence, $A = B = A \cap B$. ■

2. We use $\mathbf{pa}^{\mathcal{E}}$ and $Pa^{\mathcal{E}}(Y)$ to denote the local parent set of Y in $\mathcal{E}$, i.e., $\mathbf{pa}^{\mathcal{E}} = \mathbf{X}_A$ and $Pa^{\mathcal{E}}(Y) = A$.

This implies that one may devise a conditional independence test or causal discovery algorithm to find the (unique) local parent set on the convex subset of data (e.g., a rectangular $\mathcal{E} = [a_1, b_1] \times \dots \times [a_d, b_d] \in \mathbb{R}^d$ or a ball $\mathcal{E} = \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq r\}$). If the subset of data is non-convex, the local parent set *may* not be unique,³ and thus the results may not be valid without any additional assumptions. We provide more general condition for Prop. 6 and Thm. 7 in Appendix B.

4. Discovering Local Independences by Learning the Partition

We investigate the representation and discovery of *multiple* CSSIs in a system. To compactly represent CSSIs, we introduce a decomposition of the parents' outcome space, where each partition corresponds to a CSSI-inducing context set. Then, we transform the finding of such decomposition as a learning problem involving an auxiliary variable representing the decomposition. Finally, we develop a neural approach to learning the decomposition via the auxiliary variable.

4.1. Representing CSSIs in a system as a partition

We have seen through examples that a system may exhibit multiple CSSIs. We define a canonical notion of CSSI in order to examine a compact representation of CSSIs.

Definition 8 A CSSI relationship $Y \perp\!\!\!\perp \mathbf{X}_{A^c} \mid \mathbf{X}_A, \mathcal{E}$ is **canonical** if there does not exist any $\mathcal{F} \subseteq \mathcal{E}$ such that $Y \perp\!\!\!\perp \mathbf{X}_{B^c} \mid \mathbf{X}_B, \mathcal{F}$ and $B \subsetneq A$.

By definition, the canonicity of a CSSI implies its regularity. Canonical CSSIs are the most fine-grained ones since any subset of the context set also entails the same regular CSSI relationship. In Example 2, CSSI relationship $Y \perp\!\!\!\perp X_1 \mid X_{23}, \mathcal{E}_1 \cup \mathcal{E}_2$ is regular but not canonical since $Y \perp\!\!\!\perp X_{12} \mid X_3, \mathcal{E}_1$. If a trivial CSSI on $\mathcal{E} = \mathcal{X}$ is canonical, a system does not exhibit any CSSI relationship (e.g., Example 3 in Appendix D). We now define a compact representation of multiple CSSIs in a system.

Definition 9 (Contextual Decomposition) Let $\{\mathcal{E}_k\}_{k=0}^N$ be a partition of $\mathcal{X}$. $\{(\mathcal{E}_k, A_k)\}_{k=0}^N$ is a **contextual decomposition (CD)** if regular CSSI $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, \mathcal{E}_k$ holds for all k where $A_0 = Pa(Y)$ and $A_k \subsetneq Pa(Y)$ for all $k \geq 1$.⁴ We say a CD $\{(\mathcal{E}_k, A_k)\}_{k=0}^N$ is **canonical** if $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, \mathcal{E}_k$ is canonical for all k , and is **distinctive** if $A_i \neq A_j$ for all $i \neq j$.

$\mathcal{E}_0$ is a set to cover a remaining subset of $\mathcal{X}$ excluding non-trivial context sets. $(\mathcal{E}_0, A_0)$ corresponds to a trivial CSSI. For any causal system, $\{(\mathcal{X}, Pa(Y))\}$ is a CD, which we call **trivial CD**. When a system does not exhibit any CSSI relationship, a trivial CD is the only existing CD. In Example 2, both $\{(\mathcal{E}_0, \mathbf{X}_{123}), (\mathcal{E}_1, \mathbf{X}_3), (\mathcal{E}_2, \mathbf{X}_3)\}$ and $\{(\mathcal{E}_0, \mathbf{X}_{123}), (\mathcal{E}_1 \cup \mathcal{E}_2, \mathbf{X}_{23})\}$ are CDs where $\mathcal{E}_0 = (\mathcal{E}_1 \cup \mathcal{E}_2)^c$. The former is canonical, but the latter is not. Contextual decomposition can be viewed as the discretization of the joint outcome space, in contrast to discretizing each variables which is commonly used strategy to handle continuous variables (Nyman et al., 2017). As a canonical CSSI is a fine-grained notion of CSSI relationship, a canonical CD can also be viewed as a fine-grained CD. While a canonical CD may not be unique (e.g., Appendix A), the following theorem characterizes the shared property of the canonical CDs in a system.

3. Note that this does not imply that local parent set is *always* not unique on the non-convex subset of data, e.g., $Y \perp\!\!\!\perp X_2 \mid X_1, \mathcal{E}$ is the unique regular CSSI on non-convex set $\mathcal{E}$ in Example 1.

4. While the context sets have positive probability by its definition, we allow the case of $p(\mathcal{E}_0) = 0$.

Theorem 10 Let $\{(\mathcal{E}_i, A_i)\}_{i=0}^N$ and $\{(\mathcal{F}_j, B_j)\}_{j=0}^M$ be canonical CDs where $\mathcal{E}_i$ and $\mathcal{F}_j$ are open sets for all $i, j \geq 1$. Then, the following holds: (i) if $p(\mathcal{E}_i \cap \mathcal{F}_j) > 0$ then $A_i = B_j$, and (ii) for any $C \subseteq Pa(Y)$, $p(\mathcal{E}[C] \triangle \mathcal{F}[C]) = 0$ ⁵ where $\mathcal{E}[C] = \bigcup_{A_i=C} \mathcal{E}_i$ and $\mathcal{F}[C] = \bigcup_{B_j=C} \mathcal{F}_j$.

Roughly, any intersecting context sets (i.e., $p(\mathcal{E}_i \cap \mathcal{F}_j) > 0$) from different canonical CDs share the same local parent set (i.e., $A_i = B_j$). Also, for any $C \subseteq Pa(Y)$, the union of the context sets having the same local parent set C is identical for any canonical CSSI in a system (i.e., $p(\mathcal{E}[C] \triangle \mathcal{F}[C]) = 0$). We provide an example of canonical CDs to elaborate Thm. 10 in Example 5 (Appendix D). An interesting property of CDs is the union of the context sets having the same local parents set may no longer entails the same one, i.e., $\text{pa}^{\mathcal{E}_1 \cup \mathcal{E}_2} \neq \text{pa}^{\mathcal{E}_1} \cup \text{pa}^{\mathcal{E}_2}$ in general as shown in Example 2. For the case of distinctive canonical CDs, we have stronger uniqueness which directly follows from Thm. 10.

Corollary 11 (Uniqueness of Distinctive Canonical CD) Let $\{(\mathcal{E}_i, A_i)\}_{i=0}^N$ and $\{(\mathcal{F}_j, B_j)\}_{j=0}^M$ be distinctive canonical CDs where $\mathcal{E}_i$ and $\mathcal{F}_j$ are open sets for all $i, j \geq 1$. Then, the following hold: (i) $N = M$, and (ii) if $p(\mathcal{E}_i \cap \mathcal{F}_j) > 0$ then $A_i = B_j$ and $p(\mathcal{E}_i \triangle \mathcal{F}_j) = 0$.

In words, distinctive canonical CD is *unique* since any intersecting context sets (i.e., $p(\mathcal{E}_i \cap \mathcal{F}_j) > 0$) from different distinctive canonical CDs share the same local parent set (i.e., $A_i = B_j$), and their difference is negligible (i.e., $p(\mathcal{E}_i \triangle \mathcal{F}_j) = 0$).

4.2. Representing CD with an Auxiliary Variable

We now introduce a partition indicator variable, which will serve as a mapping between $\mathcal{X}$ and an arbitrary contextual decomposition. This variable will be useful in transforming the task of testing for CSSI to the task of learning such decomposition. We first show that a CD can be represented by introducing an auxiliary variable.

Proposition 12 (Expressing CSSI as CSI with Auxiliary Variable) Let $\{\mathcal{E}_k\}_{k=0}^N$ be a partition of $\mathcal{X}$. $\{(\mathcal{E}_k, A_k)\}_{k=0}^N$ is a contextual decomposition if and only if $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, Z = z_k$ holds for all $k \geq 1$, where Z is the deterministic random variable defined as $Z = z_k$ if $\mathbf{X} \in \mathcal{E}_k$ for all k .

In Example 1, we can introduce Z so that $Y \perp\!\!\!\perp X_2 \mid X_1, Z = z_1$ and $Y \perp\!\!\!\perp X_1 \mid X_2, Z = z_2$ where $Z = z_1$ if $\mathbf{X} \in \mathcal{E}$ and z_2 otherwise. We now formally define a partition indicator variable.

Definition 13 Let $\{\mathcal{E}_k\}_{k=0}^N$ be a partition of $\mathcal{X}$ and a random variable Z be defined as $Z = z_k$ if $\mathbf{X} \in \mathcal{E}_k$ for all $0 \leq k \leq N$. Variable Z is a **partition indicator variable (PIV)** if for all $k \geq 1$ there exists some $A_k \subsetneq Pa(Y)$ such that $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, Z = z_k$.

Prop. 12 implies that a PIV represents a corresponding CD. With PIV Z , each CSSI relationship $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, \mathcal{E}_k$ is equivalently expressed as $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, Z = z_k$. Thus, for $\mathbf{x} \in \mathcal{E}_k$,

$$p(y \mid \mathbf{x}) = p(y \mid \mathbf{x}, z_k) = p(y \mid \mathbf{x}_{A_k^c}, \mathbf{x}_{A_k}, z_k) = p(y \mid \mathbf{x}_{A_k}, z_k), \quad (1)$$

where the first equality holds by definition and the last equality holds since PIV entails CSI relationships. Therefore some of the parent variables (i.e., $\mathbf{X}_{A_k^c}$) would be ignored for modeling the conditional distribution of Y given $\mathcal{E}_k$. In contrast to CSI, which uses the value of a proper subset of $Pa(Y)$ ignoring the rest, the value of Z is determined by the whole $\mathbf{X}$, in general, and a subset of $Pa(Y)$ to be ignored (i.e., $\mathbf{X}_{A_k^c}$) could involve in determining the value of Z .

5. $\triangle$ is a symmetric difference, i.e., $A \triangle B = (A \setminus B) \cup (B \setminus A)$

4.3. Neural Contextual Decomposition

As described earlier, conditional independence tests can be used to discover the CSSI relationship on a particular subset of data. Further, assuming strictly positive densities, testing on a convex subset of data is sufficient to obtain a unique local parent set without any additional assumptions. However, it is generally infeasible to test on every possible subset to discover multiple CSSIs in a given system.

Against this background, we propose **Neural Contextual Decomposition (NCD)**, a neural approach to recovering *distinct* contextual decomposition from given data $P(\mathbf{V})$ via learning PIV Z . Recalling $Y \perp\!\!\!\perp \mathbf{X}_{A_k^c} \mid \mathbf{X}_{A_k}, Z = z_k$, we let $\mathcal{Z} \subseteq \{0, 1\}^d$ so that the local parent set A_k corresponds to the set of the indices of nonzero element of z_k , e.g., $Y \perp\!\!\!\perp X_2 \mid X_{13}, Z = (1, 0, 1)$.⁶ With PIV Z , we can write the conditional density as:

$$p(y \mid \mathbf{x}) = \sum_z p(y \mid \mathbf{x}, z) p(z \mid \mathbf{x}) = \sum_z p(y \mid \mathbf{x}_{A_z}, z) p(z \mid \mathbf{x}), \quad (2)$$

where A_z is A_k corresponding to z_k and $p(z \mid \mathbf{x}) = \delta(z = z_i)$ if $\mathbf{x} \in \mathcal{E}_i$. Our method models $p(y \mid \mathbf{x}, z)$ and $p(z \mid \mathbf{x})$.

Modeling $p(y \mid \mathbf{x}, z)$ and $p(z \mid \mathbf{x})$. Since $p(y \mid \mathbf{x}, z) = p(y \mid \mathbf{x}_{A_z}, z)$, some of the parent variables (i.e., $\mathbf{x}_{A_z^c}$) are redundant for modeling a conditional distribution. Therefore, our method approximates the conditional distribution $p(y \mid \mathbf{x}, z)$ as $\hat{p}(y \mid \mathbf{x}, z)$ and let a neural network f_θ takes $(\mathbf{x} \odot z, z)$ as an input where $\odot$ denotes an element-wise product and outputs the parameters of the estimator $\hat{p}(y \mid \mathbf{x}, z)$.⁷ For the input of f_θ , we simply concatenate $\mathbf{x} \odot z$ and z , e.g., if $z = (1, 0, 1)$ then f_θ takes $(\mathbf{x} \odot z, z) = (x_1, 0, x_3, 1, 0, 1)$ as an input. We emphasize that labeling (i.e., concatenation of z) is essential since $\mathbf{x}_{A_k^c}$ cannot be ignored without conditioning on z , i.e., $p(y \mid \mathbf{x}) = p(y \mid \mathbf{x}, z_k) = p(y \mid \mathbf{x}_{A_k}, z_k)$ holds for $\mathbf{x} \in \mathcal{E}_k$ but $p(y \mid \mathbf{x}) \neq p(y \mid \mathbf{x}_{A_k})$ and $p(y \mid \mathbf{x}, z_k) \neq p(y \mid \mathbf{x}_{A_k})$ in general. Moreover, without labeling it is unable to distinguish whether the zero entries of input is either masked or its assignment being zero. We also model the conditional distribution $p(z \mid \mathbf{x})$ as $\hat{p}(z \mid \mathbf{x})$, and let a neural network g_ϕ takes $\mathbf{x}$ as input and outputs the parameters of Bernoulli variables. That is, $g_\phi(\mathbf{x}) = (\pi_1, \dots, \pi_d)$ and $z_{(k)} \mid \mathbf{x} \sim \text{Bernoulli}(\pi_k)$ for all k where $z = (z_{(1)}, \dots, z_{(d)}) \in \{0, 1\}^d$.

Training objective. Our model becomes $p(y \mid \mathbf{x}; \theta, \phi) = \sum_z \hat{p}(y \mid \mathbf{x}, z) \hat{p}(z \mid \mathbf{x}) = \mathbb{E}_{\hat{p}(z \mid \mathbf{x})} \hat{p}(y \mid \mathbf{x}, z)$, and its maximum likelihood estimation is as follows:

$$\sup_{\theta, \phi} \mathbb{E}_{p(\mathbf{x}, y)} [\log \mathbb{E}_{\hat{p}(z \mid \mathbf{x})} \hat{p}(y \mid \mathbf{x}, z)]. \quad (3)$$

For the inner expectation, we use Monte Carlo approximation: $\frac{1}{N} \sum_{i=1}^N \hat{p}(y \mid \mathbf{x}, z^{(i)})$ where $z^{(i)} \sim \hat{p}(z \mid \mathbf{x})$. We also use Gumbel-Softmax estimator (Jang et al., 2016; Maddison et al., 2016) to enable the learning with neural networks. Final training objective is as follows:

$$\mathcal{L}(\theta, \phi; \mathcal{D}) = \sum_{(\mathbf{x}, y) \in \mathcal{D}} \left(\log \frac{1}{N} \sum_{i=1}^N \hat{p}(y \mid \mathbf{x}, z^{(i)}) \right), \quad (4)$$

where $z^{(i)}$ is sampled from $\hat{p}(z \mid \mathbf{x})$ and $\mathcal{D}$ is a dataset of samples, which could be a minibatch in practice, from the joint distribution entailed by an underlying SCM. Further details of the implementation of NCD is provided in Appendix E.3.

6. We assume that $Pa(Y) = \{1, \dots, d\}$ is identifiable and correctly discovered, e.g., by some causal discovery methods.

7. In our experiments, we model $\hat{p}$ as the Gaussian and let f_θ outputs the parameters of the Gaussian distribution.




Figure 4: Learned decision boundaries (partition) for epoch 5, 10, 15, and 20. Labels 0, 1, 2, and 3 denote $z = (0, 0), (0, 1), (1, 0),$ and $(1, 1)$, respectively.

5. Empirical Evaluation

In this section, we evaluate of our proposed method to discover contextual decomposition and CSSIs. We consider causal systems with varying type of functional model and complexity of a partition (Sec. 5.1) and with a more complex system reflecting the real-world physical dynamics (Sec. 5.2).

We compared our method with attention-based methods (Pitis et al., 2020) which aims to discover the local structure in a causal system of continuous variables, by computing the attention weights of the input $\mathbf{x}$. It first trains a transformer (Vaswani et al., 2017) to model the conditional distribution $p(y \mid \mathbf{x})$. Then, it uses the attention weights $(a_1, \dots, a_d)$ of the input $\mathbf{x}$ to discover the local independence relationship. Intuitively, low attention score a_j implies that the input variable X_j has weak influence to the target variable Y . We also compare with a mixture model, which applies a single attention layer on top of the multiple NNs. Local independence is discovered by first computing the approximation of the Jacobian of each NNs and then applying weighted sum with the learned attention score. For the experiments on the synthetic dataset, we let the number of the NNs be equal to the ground truth number of partition sets (i.e., oracle modeling).

For the implementation of $f_\theta(\mathbf{x} \odot z, z)$ which outputs the parameters of the Gaussian distribution $\hat{p}(y \mid \mathbf{x}, z)$ and $g_\phi(\mathbf{x})$ which outputs the parameters of the Bernoulli distributions, we used MLPs with 3 hidden layers and hidden units of 128. For all experiments, we set batch size to 1000 and train our method and the baselines for 100 epochs. All of our experiments were conducted with 3 different seeds, and a shaded area in an ROC plot represents a standard deviation. A detailed description of the baseline methods and further implementation details are provided in Appendix E.3.

5.1. Synthetic Data

We consider causal systems where distinctive CD exists. We generate synthetic datasets consisting of the samples $(\mathbf{X}, Y)$ where $\mathbf{X} = \{X_1, \dots, X_d\}$ is the parent variables of Y . We assume the parent variables are independent and follow the unit normal. The data generating process is as follows:

$$Y = f_i(\mathbf{X}_{A_i}, U) \quad \text{if} \quad \mathbf{x} \in \mathcal{E}_i, \quad \text{for} \quad 0 \leq i \leq N, \quad (5)$$

where U is an exogenous variable, $\{\mathcal{E}_k\}_{k=0}^N$ is a partition of $\mathcal{X}$, A_i is distinctive to each other (i.e., $A_i \neq A_j$ for all $i \neq j$). In our experiments, $d = 9$ and the exogenous variable follows the unit normal. We evaluate our method with varying types of the functions f_k , and how local parents sets A_k and partition sets $\mathcal{E}_k$ are defined. For the **functions** $\{f_k\}$, we consider nonlinear functions with




Figure 5: ROC curves for the ground truth local independence relationships on (top, middle) the synthetic dataset of (top) uniformly and (middle) non-uniformly distributed local parents sets with different types of boundaries and (bottom) on the Spriteworld.

(i) additive and (ii) non-additive noise models. We use randomly initialized neural networks for the nonlinear functions, i.e., we do not assume any specific family of distributions for f_k . For the partition $\{\mathcal{E}_k\}$, we consider the cases where the **boundaries** of the partition sets are (i) linear and (ii) nonlinear. In the case of linear boundaries, partition sets are defined by a linear function h , i.e., $x \in \mathcal{E}_k$ if $h(\mathbf{X}_{A_k}) = \max(h(\mathbf{X}_{A_0}), \dots, h(\mathbf{X}_{A_N}))$. For the nonlinear boundaries, we further control the complexity of the partition by considering both the cases when nonlinear boundaries are determined by the norm of $\mathbf{x}$ and some nonlinear function h . For the **local parent sets** $\{\mathbf{X}_{A_0}, \mathbf{X}_{A_1}, \mathbf{X}_{A_2}\}$, we consider two cases: (i) uniformly distributed as $\{\mathbf{X}_{123}, \mathbf{X}_{456}, \mathbf{X}_{789}\}$ and (ii) non-uniformly distributed as $\{\mathbf{X}, \mathbf{X}_{123}, \mathbf{X}_{456789}\}$. In our experiments, we use the different configurations of functions, local parent sets, and boundaries. For the main experiments, we consider the case of nonlinear functions with non-additive noise model and report the results under additive noise model in the appendix. We provide the details of the setups and implementations in Appendix E.1.

Experimental results. We illustrate our results in Fig. 5 for (top) the uniformly and (middle) non-uniformly distributed local parent sets, respectively. Our method successfully discovers local independence structures and outperforms the baselines in most of the cases. We again emphasize that the mixture model exploits the ground truth number of partition sets for its implementation.

Visualization of the predicted partition. We present plots in Fig. 4 illustrating how the neural network learns the partition throughout the training procedure based on a toy experiment with 2 parent variables $\mathbf{X} = \{X_1, X_2\}$. The quality of the discovered contextual decomposition is improved as the training goes on. Our model successfully discovered the local independence and the corresponding context sets. It shows that the data points within each partition are properly classified, while the model fails to classify some points close to the decision boundary.

5.2. Complex System Reflecting Real-world Dynamics

We examine our method on the modified Spriteworld environment (Watters et al., 2019; Pitis et al., 2020), which reflects real-world physics dynamics, to evaluate our method in a more complex environment. In a 2 dimensional space, there are n moving objects (i.e., sprites), which can collide with each other. Each i -th object corresponds to two 2-dimensional variables $X_{2i-1} = (x, y)_{(i)}$, $X_{2i} = (v_x, v_y)_{(i)}$ which represents the position and the velocity of the object. An agent can influence the objects through an action. Action corresponds to a 2-dimensional variable $X_{2n+1} = (a_x, a_y)$, representing some position in the 2D scene, and it influences the dynamics of the nearest object, thus there are interactions between the action and each object. In each time-step, the object and action variables influences the object variables of the next time-step. Formally, $\mathbf{X} = \{X_1, \dots, X_{2n+1}\}$ and $\mathbf{Y} = \{Y_1, \dots, Y_{2n}\}$ where $\mathbf{X}$ denotes the object and action variables and $\mathbf{Y}$ denotes the object variables of the next time-step. Further, there exist dependencies among the parent variables. Each Y_k is generated from $\mathbf{X}$ with an unknown dynamics function f_k , i.e., $Y_k = f_k(\mathbf{X}, U)$ where U is the exogenous variable. As a whole, all $\mathbf{X}$ are involved in generating each Y_k since the objects can collide with each other and the action influences to the objects. However, the interactions are sparse—the collision occurs only when objects are (i) close enough and (ii) heading to each other. Thus, the dynamics system exhibits numerous CSSI relationships. As shown in the bottom of Fig. 5, our method recovers the ground truth local structure on the varying number of objects. We provide further details of the dataset and the implementation in Appendix E.2.

6. Related Work

Local independence relationship. One of the most widely used local independence relationship is the notion of context-specific independence (CSI) (Boutilier et al., 2013; Poole, 1998). Pensar et al. (2015) proposed LDAG, a graph labeled with CSIs. In order to represent a more flexible local relationships, Pensar et al. (2016) introduced partial conditional independence (PCI), which is a generalization of CSI. Benefits of exploiting local independence relationships such as CSI has explored recently in the field of causal effect identification. Tikka et al. (2019) established the framework of causal calculus in the presence of CSI relationships along with LDAG and showed that it enables richer causal effect identification. Robins et al. (2020) proposed the algorithm for the identification of controlled direct effect leveraging CSI relationships. While our focus is on continuous variables in SCM, Nyman et al. (2017) also considered continuous variables following

Gaussian distributions but discretized the variables with certain intervals. Hence it is limited in terms of both the type of distribution and the conditioning set for local independence.

Local independence of continuous variables. Pitris et al. (2020) proposed attention-based methods to learn local causal structures. They trained a single-head transformer network and an adjacency matrix is then obtained by the product of softmax attention masks of each layer. Since it is trained *without* an inferred local causal graph, its prediction uses all of the values of the input and does not take any local independences into account for approximating conditional distribution. In contrast, our method learns *both* local causal graph within the parental relationship and conditional distribution simultaneously. Recently, Seitzer et al. (2021) proposed to estimate the conditional mutual information to discover the local independences. However, they assumed that the underlying ground truth conditional distribution is the Gaussian in order to compute the conditional mutual information. In addition, their method is restricted to a single edge, i.e., one of the parent variables and the target variable, thus not scalable and does not discover a CSSI relationships.

Neural network based causal discovery. Causal discovery attempts to reconstruct a ground truth causal graph through, e.g., conditional independences induced from the underlying system, where most of the methods fall into one of the three categories: constraint-based (Spirtes et al., 2000), score-based (Chickering, 2002), or hybrid (Wang et al., 2017) (see (Glymour et al., 2019) for an extensive review). In contrast, our goal is to discover the local independences especially within a variable and its parents. Masking the input instances with the learned mask (i.e., adjacency matrix) are widely used technique in prior works on neural network based causal discovery (Kalainathan et al., 2022; Ng et al., 2019; Brouillard et al., 2020), although our work has different focus. The key difference is that they learn the mask which is invariant to individual instances where there exists a single ground truth adjacency matrix. In our work, on the contrary, a mask is a function of the input variables which is modeled as PIV where there exist multiple valid contextual decompositions. For the same reason, they do not require labeling instances, while it is essential for valid approximation in our model.

7. Conclusion

It has been shown that local independence relationships (e.g., CSI) could be leveraged in many tasks. But most of the methods are restricted to discrete variables. To this end, we formalized the notion of context-set specific independence (CSSI) providing a more flexible representation of local independence, and further allowing both discrete and continuous variables. Along with CSSI, we introduced *local parent set* which allows CSSI to be represented in a causal graph and hence provides intuitive and compact description of local independence. For the discovery of CSSIs in a system, we focus on continuous variables which is challenging, since it is impractical to directly test every individual CSSIs due to the nature of continuous variables, and adopting prior methods to discover CSIs for discrete variables is not trivial. In this work, we tackled this challenge by finding the *contextual decomposition*, a partition of the joint outcome space where each partition is CSSI-inducing context-set. We presented NCD, effectively discovering such decomposition by augmenting an auxiliary partition indicator variable (PIV) which enables the gradient based training. While finding PIV is equivalent to find contextual decomposition which implies the discovery of CSSIs, experimental results showed that NCD successfully finds PIV and discovers CSSIs in a synthetic dataset and is also effective in a complex environment reflecting real-world dynamics.

References

- Joshua D Angrist and Jörn-Steffen Pischke. *Mostly harmless econometrics*. Princeton university press, 2008.
- Joshua D Angrist, Guido W Imbens, and Donald B Rubin. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455, 1996.
- Abhijit Vinayak Banerjee, Esther Duflo, and Michael Kremer. The influence of randomized controlled trials on development economics research and on development policy. *The State of Economics, The State of the World*, pages 482–488, 2016.
- Craig Boutilier, Nir Friedman, Moisés Goldszmidt, and Daphne Koller. Context-specific independence in bayesian networks. *CoRR*, abs/1302.3562, 2013.
- Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable causal discovery from interventional data. *arXiv preprint arXiv:2007.01754*, 2020.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- Giso H Dal, Alfons W Laarman, and Peter JF Lucas. Parallel probabilistic inference by weighted model counting. In *International Conference on Probabilistic Graphical Models*, pages 97–108. PMLR, 2018.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- Vibhav Gogate and Pedro Domingos. Exploiting logical structure in lifted probabilistic inference. In *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- Patrik O Hoyer, Dominik Janzing, Joris Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Proceedings of the 21st International Conference on Neural Information Processing Systems*, pages 689–696, 2008.
- Guido W Imbens. Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4):1129–79, 2020.
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Diviyani Kalainathan, Olivier Goudet, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Structural agnostic modeling: Adversarial learning of causal graphs. *Journal of Machine Learning Research*, 23(219):1–62, 2022.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

- Ignavier Ng, Zhuangyan Fang, Shengyu Zhu, Zhitang Chen, and Jun Wang. Masked gradient-based causal structure learning. *arXiv preprint arXiv:1910.08527*, 2019.
- Henrik Nyman, Johan Pensar, and Jukka Corander. Stratified gaussian graphical models. *Communications in Statistics-Theory and Methods*, 46(11):5556–5578, 2017.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Johan Pensar, Henrik Nyman, Timo Koski, and Jukka Corander. Labeled directed acyclic graphs: a generalization of context-specific independence in directed graphical models. *Data mining and knowledge discovery*, 29(2):503–533, 2015.
- Johan Pensar, Henrik Nyman, Jarno Lintusaari, and Jukka Corander. The role of local partial independence in learning of bayesian networks. *International journal of approximate reasoning*, 69:91–105, 2016.
- Jonas Peters. On the intersection property of conditional independence and its application to causal discovery. *Journal of Causal Inference*, 3(1):97–108, 2015.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Silviu Pitis, Elliot Creager, and Animesh Garg. Counterfactual data augmentation using locally factored dynamics. *Advances in Neural Information Processing Systems*, 33, 2020.
- David Poole. Context-specific approximation in probabilistic inference. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 447–454, 1998.
- David Poole and Nevin Lianwen Zhang. Exploiting contextual independence in probabilistic inference. *Journal of Artificial Intelligence Research*, 18:263–313, 2003.
- James M Robins, Thomas S Richardson, and Ilya Shpitser. An interventionist approach to mediation analysis. *arXiv preprint arXiv:2008.06019*, 2020.
- Maximilian Seitzer, Bernhard Schölkopf, and Georg Martius. Causal influence detection for improving efficiency in reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 22905–22918, 2021.
- Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- Bill Shipley. *Cause and correlation in biology: a user’s guide to path analysis, structural equations and causal inference with R*. Cambridge University Press, 2016.
- Michael E Sobel. Causal inference in the social and behavioral sciences. In *Handbook of statistical modeling for the social and behavioral sciences*, pages 1–38. Springer, 1995.
- Pater Spirtes, Clark Glymour, Richard Scheines, Stuart Kauffman, Valerio Aimale, and Frank Wimberly. Constructing Bayesian network models of gene expression networks from microarray data. 2000.

- Santtu Tikka, Antti Hyttinen, and Juha Karvanen. Identifying causal effects via context-specific independence relations. *Advances in Neural Information Processing Systems*, 32:2804–2814, 2019.
- Guy Van den Broeck, Nima Taghipour, Wannes Meert, Jesse Davis, and Luc De Raedt. Lifted probabilistic inference by first-order knowledge compilation. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence*, pages 2178–2185, 2011.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- Yuhao Wang, Liam Solus, Karren Dai Yang, and Caroline Uhler. Permutation-based causal inference algorithms with interventions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5824–5833, 2017.
- Nicholas Watters, Loic Matthey, Matko Bosnjak, Christopher P Burgess, and Alexander Lerchner. Cobra: Data-efficient model-based rl through unsupervised object discovery and curiosity-driven exploration. *arXiv preprint arXiv:1905.09275*, 2019.
- Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*, pages 804–813. AUAI Press, 2011.
- Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. DAGs with NO TEARS: Continuous optimization for structure learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. In *International Conference on Learning Representations*, 2019.