



Steps Towards Causal Formal Concept Analysis

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo Napoli

► To cite this version:

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo Napoli. Steps Towards Causal Formal Concept Analysis. International Journal of Approximate Reasoning, 2022, 142, pp.338-348. hal-03183755v1

HAL Id: hal-03183755

<https://hal.science/hal-03183755v1>

Submitted on 28 Mar 2021 (v1), last revised 25 Jan 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Steps Towards Causal Formal Concept Analysis

Alexandre Bazin, Miguel Couceiro, Marie-Dominique Devignes, Amedeo Napoli

LORIA, Université de Lorraine – CNRS – Inria, Nancy, France
`firstname.name@loria.fr`

Abstract. Efficiently discovering causal relations from data and representing them in a way that facilitates their use is an important problem in science that has received much attention. In this paper, we propose an adaptation of the formal concept analysis formalism to the discovery and representation of causal relations. We show that formal concept analysis structures and algorithms are well-suited to this problem.

1 Introduction

Studying causal relations is at the heart of science [24]. Sometimes, experimental protocols are thought of in order to manipulate a potential cause (an object, state or process) in such a way that its effects can be inferred. This produces what we call interventional data. Often, the objects cannot be manipulated directly so we have to rely on purely observational data. Studying causality from data involves two related tasks:

- discovering the causal structure (is there a causal relation between these two sets of variables?)
- inferring causal effects (how does the cause causes the effect?)

In this paper, we are interested in the discovery of the causal structure in observational data.

Inferring causal relations from observational data is a difficult problem that has attracted much attention. Most approaches are based on probability theory [19, 22, 30, 31, 33] but algorithmic approaches based on approximations of Kolmogorov complexity are gaining traction [6, 7, 5, 27, 29]. The set of causal relations forms the causal structure. It can be presented to humans to help them understand the situation described by the data and make decisions [10] or it can be used to automatically select important variables in the data [38]. Most of the works on the topic focus on univariate causal inference, in which the causal relations are between two variables x and y . The causal structure is then usually represented by a directed acyclic graph called the causal diagram. In practice, effects often require the interaction of multiple causes to appear. We then talk of multivariate causal relations. In the multivariate case, in which causal relations can be between sets X and Y of variables, the equivalent of causal diagrams have yet to be properly studied.

Formal concept analysis [12] is a lattice-theoric mathematical framework that allows for the representation of the underlying structure of data. In its original form, it can only be applied on binary data. Extensions of the framework have been proposed to apply formal concept analysis to more complex data such as numerical data [11]. In this paper, we propose a new variation of the formal concept analysis framework that uses causal inference in order to

- represent the causal structure of the variables in a dataset with formal concept analysis structures
- have formal concept analysis work on different kinds of data, such as numerical data, without explicit or implicit binarisation

In Section 2, we provide a brief overview of existing causal inference methods and present the notions of formal concept analysis that are important for this work. In Section 3, we discuss properties of causal relations and use them to build a closure operator. In Section 4, we discuss the use of formal concept analysis structures for the representation of causal relations and provide algorithms for computing them. In Section 5, we present an example of our proposed approach applied to the well-known public dataset Iris. Finally, in Section 6, we discuss the usefulness of our approach for computing and representing causal relations.

2 Preliminaries

2.1 Notation

In this paper, as a rule of thumb, important sets and structures are denoted by calligraphed letters (\mathcal{A}), (sub)sets are denoted by upper case letters (A) and individual elements are denoted by lower case letters (a). Sets of elements are mostly written in the usual form ($\{a, b, c, d, e\}$) but can sometimes be written in a simplified form ($abcde$) in figures for legibility’s sake.

We use $X \xrightarrow{c} Y$ to denote the fact that X causes Y .

2.2 Causal Inference

In the univariate case, the structure of the causal relation is represented by a directed acyclic graph (DAG), called the *causal diagram* or *causal network*, in which the vertices are variables and the edges denote a causal relation between a cause and its effect. Learning a causal diagram [15] from observational data has been the object of much interest, whether it be to learn a whole diagram [34, 35] or just parts of it [13]. Many approaches first try to learn the skeleton of the diagram (is there a causal relation between these two variables?) then to orient the edges (which of these variables is the cause?).

Inferring the direction of a causal relation between x and y is a difficult task and assumptions are usually made. The most common one is the absence of confounders [23, 36, 39], i.e. variables z that cause both x and y . Traditional

approaches are based on additive noise models [19, 20]. They assume that, if x causes y , the value of y is a function of the value of x plus some additive noise, i.e.

$$y = f(x) + n$$

where n is independent of x . Under this assumption, state of the art approaches [19, 22, 30, 31, 33] make use of the idea that the marginal distribution of the cause, $P(x)$, and the conditional distribution of the effect given the cause, $P(y|x)$, should be independent.

In contrast, algorithmic approaches are based on the algorithmic Markov condition [23] and the idea that it is algorithmically easier to compute the effect from the cause than the other way around. The Kolmogorov complexity $K(s)$ of a finite binary string s is the length of the shortest program for which a universal Turing machine outputs s and halts. In this setting, it is said that x causes y when $K(P(x)) + K(P(y|x)) \ll K(P(y)) + K(P(x|y))$. As the Kolmogorov complexity is not computable, approximations of it are used to infer causal directions in Boolean [6, 7] and numerical [5, 27, 29] data.

Probabilistic [8, 21, 40] and algorithmic [37] approaches have also been proposed for the multivariate case, when causal relations are between sets of variables X and Y . However, this has received much less attention than the univariate case. In the experimental example presented in Section 5, we use ERGO [37], an algorithmic approach for inferring causal directions in the multivariate case by approximating the Kolmogorov complexity using cumulative and Shannon entropy.

2.3 Formal Concept Analysis

Definition 1. (FORMAL CONTEXT)

A formal context is a triple $(\mathcal{G}, \mathcal{M}, \mathcal{R})$ in which \mathcal{G} is a set of objects, \mathcal{M} is a set of attributes and $\mathcal{R} \subseteq \mathcal{G} \times \mathcal{M}$ is a binary relation between objects and attributes. We say that an object g is described by an attribute m when $(g, m) \in \mathcal{R}$.

	a	b	c	d	e
1	×	×			
2		×	×	×	
3		×		×	×
4	×		×		
5				×	×

Fig. 1. Crosstable representing a formal context with five objects ($\{1, 2, 3, 4, 5\}$) and five attributes $\{a, b, c, d, e\}$.

Associated to a formal context $(\mathcal{G}, \mathcal{M}, \mathcal{R})$ are two *derivation operators* \cdot' such that

$$\begin{aligned} \cdot' : 2^{\mathcal{G}} &\mapsto 2^{\mathcal{M}} \\ O' &= \{a \in \mathcal{M} \mid \forall o \in O, (o, a) \in \mathcal{R}\} \\ \cdot' : 2^{\mathcal{M}} &\mapsto 2^{\mathcal{G}} \\ A' &= \{o \in \mathcal{G} \mid \forall a \in A, (o, a) \in \mathcal{R}\} \end{aligned}$$

The derivation operators form a Galois connection and so both compositions \cdot'' are *closure operators* on, respectively, \mathcal{G} and \mathcal{M} . We say that a formal context is *object-reduced* when, for all objects o , there are no other objects p and q such that $\{o\}' = \{p\}' \cap \{q\}'$. We define *attribute-reduced* contexts analogously.

Definition 2. (FORMAL CONCEPT)

Let $\mathcal{K} = (\mathcal{G}, \mathcal{M}, \mathcal{R})$ be a formal context. A pair $(G, M) \in 2^{\mathcal{G}} \times 2^{\mathcal{M}}$ is called a formal concept of \mathcal{K} if and only if $M = G'$ and $G = M'$. In this case, G is called the *extent* and M the *intent* of the concept.

Formal concepts represent classes of objects or, equivalently, of attributes that are in the data. The set of all the formal concepts of a formal context together with the partial order induced by the inclusion relation on either the extents or intents forms a complete lattice called the *concept lattice* of the formal context.

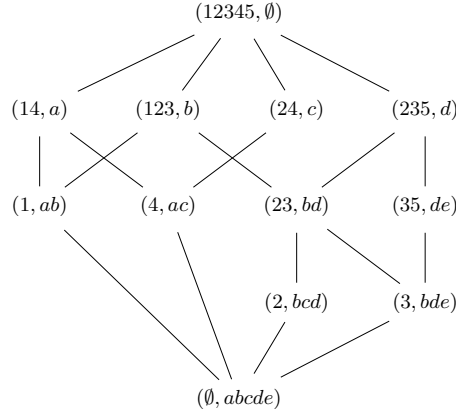


Fig. 2. Concept lattice of the formal context depicted in Fig. 1.

Definition 3. (IMPLICATIONS) Let $\mathcal{K} = (\mathcal{G}, \mathcal{M}, \mathcal{R})$ be a formal context. An implication is a pair (A, B) of attribute sets, often written in rule form $A \rightarrow B$. An implication $A \rightarrow B$ holds in \mathcal{K} when $B' \subseteq A'$, i.e. all the objects described by the attributes in A are also described by the attributes in B . Let $\mathcal{I}_{\mathcal{K}}$ denote the set of all the implications that hold in \mathcal{K} .

Implications represent regularities in the descriptions of the objects. For instance, in the formal context depicted in Figure 1, the implications $\{b, c\} \rightarrow \{d\}$ and $\{e\} \rightarrow \{d\}$ hold while $\{a\} \rightarrow \{b\}$ does not. The number of implications in a formal context $(\mathcal{G}, \mathcal{M}, \mathcal{R})$ can grow exponentially with $|\mathcal{G}|$, $|\mathcal{M}|$ and $|\mathcal{R}|$. Hence, even relatively small formal contexts can give rise to an intractable number of implications [26]. Whether it be to present to a human analyst or to use in a computation, two approaches to reduce the implication set have been explored: with or without loss of information. When loss of information is acceptable, one can use the many interestingness measures that have been developed over the years [14]. When all the information has to be preserved, one has to use implication bases.

Definition 4. (IMPLICATIONS BASES) *Let \mathcal{K} be a formal context. An implication set $I \subseteq \mathcal{I}_{\mathcal{K}}$ is an implication base of \mathcal{K} if all the implications in $\mathcal{I}_{\mathcal{K}}$ can be derived from those of I through Armstrong's axioms:*

- If $Y \subseteq X$, then $X \rightarrow Y$
- If $X \rightarrow Y$, then $X \cup Z \rightarrow Y \cup Z$ for all Z
- If $X \rightarrow Y$ and $Y \rightarrow Z$, then $X \rightarrow Z$

As implication bases allow for the derivation of all the implications that hold, they contain the same information as the whole implication set. Different implication bases, with different properties, have been studied [16, 32, 4]. Here, we present the two best-known.

Definition 5. (LOGICAL CLOSURE) *Let I be an implication set. The logical closure $I(X)$ of an attribute set X by I is the smallest superset of X such that $(A \rightarrow B \in I \text{ and } A \subseteq I(X))$ implies $B \subseteq I(X)$.*

When I is an implication base, $I(A) = A''$ for all $A \subseteq \mathcal{M}$. If \mathcal{K} is the context depicted in Figure 1, then $\mathcal{I}_{\mathcal{K}}(\{e\}) = \{d, e\}$.

Definition 6. (LOGICAL PSEUDO-CLOSURE) *Let I be an implication set. The logical pseudo-closure $I^{\square}(X)$ of an attribute set X by I is the smallest superset of X such that $(A \rightarrow B \in I \text{ and } A \subset I^{\square}(X))$ implies $B \subseteq I^{\square}(X)$.*

Note that the only difference between the logical closure and pseudo-closure is the strictness of the inclusion of premises. The logical pseudo-closure is a closure operator. If \mathcal{K} is the context depicted in Figure 1, then $\mathcal{I}_{\mathcal{K}}^{\square}(\{e\}) = \{e\}$ because no implications in which the premise is a proper subset of $\{e\}$ holds.

Definition 7. (PSEUDO-INTENT) *Let \mathcal{K} be a formal context. An attribute set P is a pseudo-intent if $P \neq P''$ and $P = \mathcal{I}_{\mathcal{K}}^{\square}(P)$.*

The attribute sets closed under the logical pseudo-closure by the set of implications are either intents or pseudo-intents. For this reason, computing the set of pseudo-intents can be done by computing all the sets that are closed under $\mathcal{I}_{\mathcal{K}}^{\square}(\cdot)$ [1, 2, 25, 9].

Definition 8. (CANONICAL BASE) *The canonical base (or Duquenne-Guigues base [16]) of a formal context is*

$$\{P \rightarrow P'' \mid P \text{ is a pseudo-intent}\}$$

The canonical base is the smallest implication base of a formal context in terms of number of implications. As such, when the number of implications is a concern, the canonical base constitutes the optimal compression of the information contained in the implications. In the formal context depicted in Figure 1, the canonical base is

- $\{e\} \rightarrow \{d, e\}$
- $\{a, d\} \rightarrow \{a, b, c, d, e\}$
- $\{b, c\} \rightarrow \{b, c, d\}$
- $\{c, d\} \rightarrow \{b, c, d\}$
- $\{b, c, d, e\} \rightarrow \{a, b, c, d, e\}$

Definition 9. (PROPER PREMISE) *Let \mathcal{K} be a formal context and a an attribute. An attribute set X is a proper premise of a if $X \rightarrow \{a\}$ holds, $X \neq \{a\}$, and, for all $Y \subset X$, $Y \rightarrow \{a\}$ does not hold.*

A proper premise is a minimal attribute set that implies another attribute.

Definition 10. (PROPER PREMISES BASE) *The base of proper premises [32] of a formal context is*

$$\{X \rightarrow X'' \mid X \text{ is a proper premise}\}$$

The base of proper premises is the implication base with the smallest premises. As such, it constitutes the optimal compression of the information contained in the implications when the size of the premises is a concern. In the formal context depicted in Figure 1, the base of proper premises is

- $\{e\} \rightarrow \{d, e\}$
- $\{b, c\} \rightarrow \{b, c, d\}$
- $\{c, d\} \rightarrow \{b, c, d\}$
- $\{a, d\} \rightarrow \{a, b, c, d, e\}$
- $\{c, e\} \rightarrow \{a, b, c, d, e\}$
- $\{a, e\} \rightarrow \{a, b, c, d, e\}$
- $\{a, b, c\} \rightarrow \{a, b, c, d, e\}$

Note that minimal pseudo-intents are also necessarily proper premises.

3 Causal Closure

Works on modelling or discovering causal relations usually make a number of assumptions on the nature of the relation: univariate or multivariate, absence of confounders, linearity of the function in the Additive Noise Models... The present

work is no exception. In this section, we present the different properties that we suppose the causal relations have. These properties are intuitively acceptable and make the causal relations behave like implications.

In the remainder of this paper, we shall use the following set of causal relations, supposedly existing in a dataset with five variables, as a running example:

- v_1 causes v_2 : $\{v_1\} \xrightarrow{c} \{v_2\}$
- v_3 causes v_4 : $\{v_3\} \xrightarrow{c} \{v_4\}$
- v_1 and v_4 together cause v_5 : $\{v_1, v_4\} \xrightarrow{c} \{v_5\}$

Let X , Y and Z be sets of variables. Then, the following five properties hold:

$$X \xrightarrow{c} X \quad (1)$$

A set of variables causes itself.

$$X \xrightarrow{c} Y \wedge Y \xrightarrow{c} Z \Rightarrow X \xrightarrow{c} Z \quad (2)$$

Causality is transitive. Indeed, if X directly causes Y and Y directly causes Z , one can say that X indirectly causes Z .

$$X \xrightarrow{c} Y \Rightarrow X \xrightarrow{c} Z, \forall Z \subseteq Y \quad (3)$$

If X causes a set of variables Y , it causes its subsets. Most importantly, this implies that X causes $\{y\}$ for all y in Y .

$$X \xrightarrow{c} Y \Rightarrow X \cup Z \xrightarrow{c} Y \quad (4)$$

All supersets of X cause its effects. Note that this is not incompatible with possible notions of *negation* as \xrightarrow{c} only denotes the presence of a causal relation and not the causal effect itself, which can be different between $X \xrightarrow{c} Y$ and $X \cup Y \xrightarrow{c} Y$.

$$X \xrightarrow{c} Y \wedge X \xrightarrow{c} Z \Rightarrow X \xrightarrow{c} Y \cup Z \quad (5)$$

Causing two sets of variables means causing them together.

Properties 1 and 3 imply that if $Y \subseteq X$, then $X \xrightarrow{c} Y$. Properties 1, 3 and 4 imply that if $X \xrightarrow{c} Y$, then $X \cup Z \xrightarrow{c} Y \cup Z$. As such, and since property 2 is the transitivity, these five properties imply Armstrong's axioms (see Definition 4). Note that the transitivity of causality, while intuitive, is disputed [28] and counterexamples have been presented. Halpern [17] identified conditions under which causality is transitive.

Proposition 1. *Let X be a variable set. The operator*

$$\xi(X) = \bigcup_{X \xrightarrow{c} Y} Y$$

is a closure operator.

Proof. Because of property 1, we have that $X \subseteq \xi(X)$ so $\xi(\cdot)$ is extensive. Because of properties 1, 2 and 3, we have that $X \subseteq Y \Rightarrow \xi(X) \subseteq \xi(Y)$ so $\xi(\cdot)$ is monotone. Because of properties 1 and 2, we have that $\xi(\xi(X)) = \xi(X)$ so ξ is idempotent. As such, ξ is a closure operator.

We call $\xi(X)$ the *causal closure* of X . The causal closure of X is X plus all its direct and indirect effects. When all the causal relations are known, computing the causal closure can be done by starting from X and adding the effects until a fixpoint is reached. However, when faced with a dataset from which the causal relations have to be inferred at great cost, one cannot test whether $X \xrightarrow{c} Y$ for all possible Y . Fortunately, property 3 ensures that the computation of $\xi(X)$ can be simplified.

Proposition 2. *Let X be a variable set and y a variable. Then,*

$$\xi(X) = \{y \mid X \xrightarrow{c} \{y\}\}$$

Proof. This follows from the fact that $X \xrightarrow{c} Y \Rightarrow X \xrightarrow{c} Z, \forall Z \subseteq Y$.

Hence, computing $\xi(X)$ directly from a dataset requires $|\mathcal{V} \setminus X|$ causal tests.

4 Reconstructing the FCA Trinity with Causality

As we have shown in the previous section, causal relations give rise to a closure operator ξ . Let \mathcal{V} be the set of variables in a dataset. Now, let us suppose that \mathcal{K}_c is a formal concept which set of attributes is \mathcal{V} and in which the implication $X \rightarrow Y$ holds if and only if $X \xrightarrow{c} Y$. Clearly, as causal relations respect Armstrong's axioms, $\mathcal{I}_{\mathcal{K}_c}(X) = \xi(X)$ for all X and any implication base of \mathcal{K}_c allows for the derivation of all the causal relations. Thus, by using the closure operator ξ as if it were the closure operator induced by the (unknown) formal context \mathcal{K}_c , it is possible to apply existing formal concept analysis algorithms to the computation of various representations of the causal relations in a dataset.

Formal concept analysis offers many structures to represent the content of the closure operator and, thus, of the causal relations: the formal context, the concept lattice, the AOC-poset, the implication bases... All of them provide their unique perspective but some are easier to handle or compute given only the closure operator as a black box. In this section, we go through these structures and discuss both their usefulness as representations of causal relations and their ease of use.

4.1 The Causal Canonical Base

The canonical base of \mathcal{K}_c is the cardinality-wise smallest implication base of \mathcal{K}_c . As such, it corresponds to the cardinality-wise smallest set of causal relations that allow for the derivation of all the causal relations in the dataset. We call it the *causal canonical base* of the dataset. It is a set of multivariate causal relations

of the form $X \rightarrow \xi(X)$ (analog to the implications $X \rightarrow X''$ in classical FCA) in which X contains all the variables it causes except for those specified by the rule $X \rightarrow \xi(X)$ itself. In our running example, the causal canonical base is

- $\{v_1\} \rightarrow \{v_2\}$
- $\{v_3\} \rightarrow \{v_4\}$
- $\{v_1, v_2, v_4\} \rightarrow \{v_5\}$

As the smallest rule-based representation of causal relations, the causal canonical base is most useful when the number of rules is a concern: when visualising the structure or when storing it for computer use. This base is also arguably the easiest representation to compute as existing algorithms for computing pseudo-closed sets make use of the closure operator as a black box. For instance, the NEXT CLOSURE (Algorithm 1) can be used.

Algorithm 1: Causal NEXT CLOSURE for implications

Input: Dataset with variables \mathcal{V}

Output: The causal canonical base of the dataset

```

1 begin
2    $I = \emptyset;$ 
3    $V = \emptyset;$ 
4   while  $V \neq \mathcal{V}$  do
5     if  $V \neq \xi(V)$  then
6        $I = I \cup \{V \xrightarrow{c} \xi(V)\}$ 
7      $V = \text{Next}(V)$ 
8 return  $I$ 
```

4.2 Causal Proper Premises

The base of proper premises of \mathcal{K}_c is the implication base with the smallest premises. As such, it is made of rules of the form $X \rightarrow \xi(X)$ in which X is a *minimal*, or sufficient, cause of the variables $y \in \xi(X) \setminus X$. We call this base the *causal sufficiency base*. In our running example, the causal sufficiency base is

- $\{v_1\} \rightarrow \{v_2\}$
- $\{v_3\} \rightarrow \{v_4\}$
- $\{v_1, v_3\} \rightarrow \{v_5\}$
- $\{v_1, v_4\} \rightarrow \{v_5\}$

Sufficient causes are particularly important because they allow us to pinpoint what can be acted on to predict or modify the future. For instance, if the association of high blood pressure and sedentary lifestyle is found to be a sufficient cause of the development of a disease, we know that correcting either of those

will certainly help prevent the disease. The causal sufficiency base is therefore interesting as it readily contains all the sufficient causes of all the variables. However, it is more difficult to compute than the causal canonical base. Indeed, computing proper premises is usually seen as computing the minimal transversals of hypergraphs constructed from the formal context [32]. In our case, we only have a closure operator and the formal context is unknown. Thus, computing the proper premises of the causal relation requires the use of a well-known FCA algorithm tailor-made for when the underlying formal context is unknown: ATTRIBUTE EXPLORATION [12].

ATTRIBUTE EXPLORATION is an algorithm for computing an implication base of a formal context that is only known to an expert. It works by repeatedly asking the expert whether an implication holds and, if not, to provide a counterexample in the form of an object which description invalidates the implication. The algorithm thus produces an implication base and a formal context that contains the same information as the unknown one. In our case, we use the version proposed in [32] for proper premises and the expert is played by the causal closure operator ξ as illustrated in Algorithm 2.

Algorithm 2: Causal ATTRIBUTE EXPLORATION.

Input: Dataset with variables \mathcal{V}
Output: The causal sufficiency base of the dataset

```

1 begin
2    $I = \emptyset$ ;
3    $\mathcal{G} = \emptyset$ ;
4    $\mathcal{R} = \emptyset$ ;
5    $K = (\mathcal{G}, \mathcal{V}, \mathcal{R})$ ;
6   for  $a \in \mathcal{V}$  do
7      $\mathcal{E} = H_{K,a}^{\mathcal{C}}$ ;
8      $T = \{\emptyset\}$ ;
9     while there exists  $E \in \mathcal{E}$  do
10       $T = \min(T \vee \{\{a\} \mid a \in E\})$ ;
11       $\mathcal{E} = \mathcal{E} \setminus \{E\}$ ;
12      while there exists  $Q \in T$  with  $L \not\models \{Q \rightarrow Q''\}$  do
13        if  $Q'' \subseteq \xi(Q)$  then
14           $I = I \cup \{Q \rightarrow Q''\}$ ;
15        else
16           $\mathcal{G} = \mathcal{G} \cup \{o\}$ ;
17           $\mathcal{R} = \mathcal{R} \cup \{(o, x) \mid x \in \xi(Q)\}$ ;
18          if  $a \notin o'$  then
19             $\mathcal{E} = \mathcal{E} \cup \{M \setminus o'\}$ ;
20  return  $I$ 

```

We suppose that we initially have no information about the unknown formal context, i.e. the causal relations, so we start with an empty implication set I and an empty formal context $K = (\emptyset, \mathcal{V}, \emptyset)$ (lines 2–5). The algorithm then enumerates implications $Q \rightarrow Q''$ that hold in K and checks whether $Q \xrightarrow{c} Q''$ through computing $\xi(Q)$. If $Q \rightarrow Q''$ does not hold in the underlying context because Q does not cause Q'' , the counterexample is a new object o such that $o' = Q'' \setminus \{q\}$ with $q \in Q'' \setminus I(Q)$ such that Q does not cause $\{q\}$ in the dataset. The proper premises of an attribute a in the context K are the minimal transversals of the hypergraph

$$H_{K,a}^{\neq} = \{\overline{o'} \setminus \{a\} \mid a \notin o'\}.$$

where $\overline{o'} = \mathcal{M} \setminus o'$. Thus, the implications are enumerated as follows. For each attribute a (line 6), the algorithm computes $H_{K,a}^{\neq}$ (line 7) and then incrementally computes its set of minimal transversals T using Berge’s multiplication algorithm [3] (lines 9–11). For each transversal, and thus each proper premise Q (line 12), the validity of $Q \rightarrow \xi(Q)$ is tested (line 13). If the causal relation holds, $Q \rightarrow Q''$ is added to the set of implications I (lines 13–14). If it does not hold, a new object o is created, which description is $\xi(Q)$ (lines 15–19).

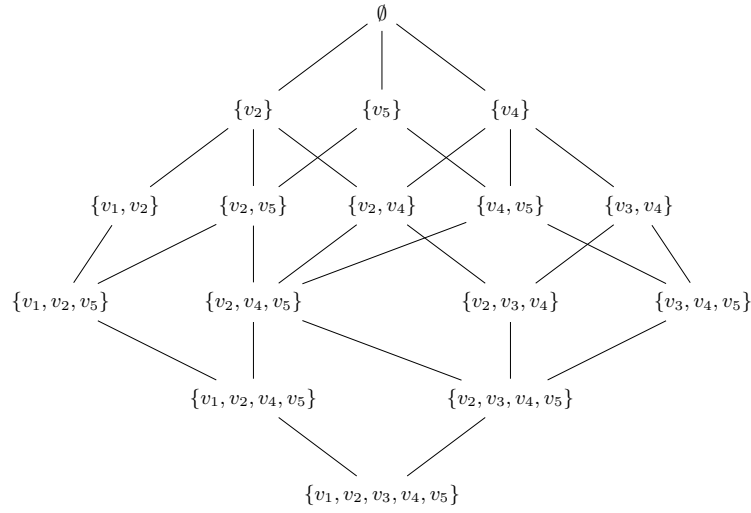
4.3 Causal Intents Lattice and Context

The lattice formed by the sets of variables closed under ξ ordered by the inclusion relation is another representation of the causal structure. Closed sets C are the maximal elements of their equivalence classes C_{\equiv} that contain all the variable sets that have the same direct and indirect causes as C . We call them *causal intents*. The causal intents lattice corresponding to our running example is depicted in Fig. 3. The set $\{v_3, v_4\}$ is a causal intent as it is the maximal element of the equivalence class that contains $\{v_3\}$ and $\{v_3, v_4\}$. The fact that $\{v_3\}$ belongs to the same equivalence class as $\{v_3, v_4\}$ means that v_3 causes v_4 . As such, the causal relations can be read from the causal intents lattice. However, despite the similarities between the graphical representations of causal intents lattices and causal diagrams [15], the causal relations do not correspond to edges in the lattice.

The formal contexts underlying the closure operator ξ and of which the causal intents lattice is the intents lattice are called the *causal contexts*. The causal context which objects correspond exactly to the \wedge -irreducible elements of the causal intents lattice is called the *reduced causal context*. The causal context of our running example is depicted in Fig 4.

5 Experimental Example

Computing the causal structure representations discussed in this paper requires the ξ operator. In practice, it is impossible to correctly infer all the causal relations from an observational dataset so we have to rely on approximations provided by the various approaches presented in Section 2.2.

**Fig. 3.** Causal intents lattice of the running example.

v_1	v_2	v_3	v_4	v_5
×	×		×	×
	×	×	×	×
×	×			×
	×	×	×	
		×	×	×
×	×			

Fig. 4. The reduced causal context of the running example.

In our experiments, we considered that a variable set X causes a variable y if and only if X and $\{y\}$ are sufficiently correlated and $X \xrightarrow{c} \{y\}$ is inferred by ERGO, the multivariate causal relation inference approach proposed in [37]. The correlation between the variable sets X and $\{y\}$, denoted by $\text{corr}(X, \{y\})$, is measured using Linear Canonical Correlation Analysis [18] and both sets are deemed sufficiently correlated when $\text{corr}(X, \{y\}) > 0.8$.

The Iris dataset is one of the best-known datasets for pattern recognition and machine learning. It describes 150 flowers of the *iris* genus by the values of five variables *sepal_length*, *sepal_width*, *petal_length*, *petal_width* and *class*. The first four are numerical variables while the class is nominal with three possible values.

In the Iris dataset, both the causal canonical base and the causal sufficiency base contain the same rules:

- $\{\text{petal_width}\} \rightarrow \{\text{petal_length}, \text{class}\}$
- $\{\text{petal_length}\} \rightarrow \{\text{class}\}$

$$- \{sepal_length\} \rightarrow \{petal_length, petal_width, class\}$$

We observe that *class* is caused by *petal_width*, *petal_length* and *sepal_length* but not by *sepal_width*. Additionally, *petal_length* appears to be a direct cause of *class*, itself caused by *petal_width*, itself caused by *sepal_length*.

The corresponding reduced causal context is depicted in Fig. 5 and the causal intents lattice in Fig. 6. The causal relations can be read from the lattice, e.g. $\{petal_length\} \xrightarrow{c} \{class\}$ can be inferred from the fact that $\{petal_length\}$ belongs to the same equivalence class as $\{petal_length, class\}$.

	<i>sepal_length</i>	<i>sepal_width</i>	<i>petal_length</i>	<i>petal_width</i>	<i>class</i>
	×		×	×	×
		×	×	×	×
		×	×		×
		×			×
		×			

Fig. 5. The reduced causal context of the Iris dataset.

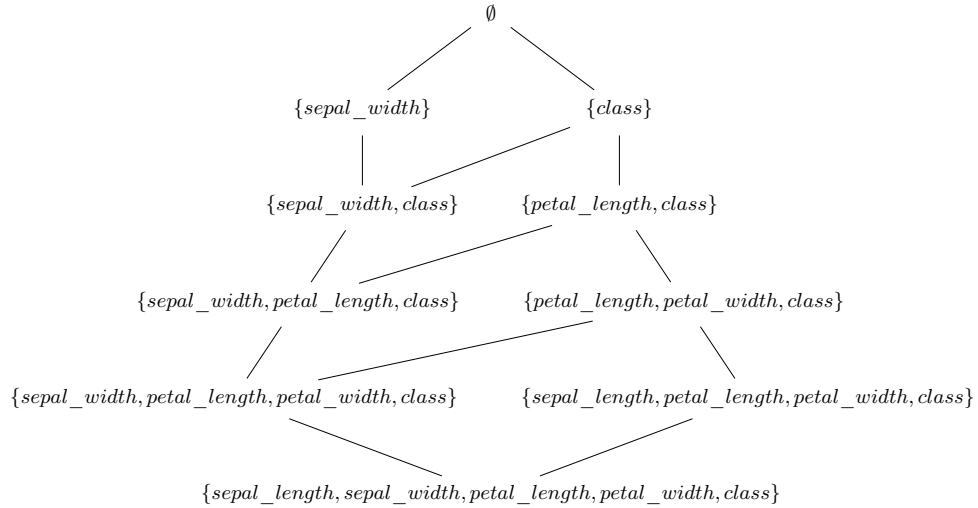


Fig. 6. Causal intents lattice of the Iris dataset.

6 Discussion and Conclusion

In this paper, we proposed a method to adapt the formal concept analysis framework to the computation and representation of causal relations between the

variables in datasets. Multivariate causal inference is a very difficult problem for which the current state of the art approaches only manage to obtain satisfactory results in specific settings. As our proposed method uses causal inference as a black box, any existing or future approaches for this problem can be plugged in. This allows it to work on any type of data for which causal inference is or will be studied, such as Boolean or numerical data. This, once again, highlights the flexibility of formal concept analysis.

A number of choices have been made in this work. The properties of causality discussed in Section 3 are necessary for us to obtain a closure operator but some of them, such as the transitivity, do not necessarily hold. Without a closure operator, the approach is not usable. For the experimental example, we chose to use a combination of ERGO and linear canonical correlation analysis for inferring causal relations. Using different approaches for measuring correlation and orienting the causal direction could have produced different causal relations and thus a different lattice and a different causal context.

The performances of approaches for univariate causal inference are evaluated on real or synthetic datasets [30] for which the true causal relations are known. In the multivariate case, there are very few datasets with ground truths and they only contain a single cause and a single effect. This is clearly not enough to evaluate an approach made for representing multiple multivariate causal relations. Furthermore, it is not clear how to generate synthetic data according to multivariate causal relations. In the general case, loop-like structures such as $\{\{a\} \xrightarrow{c} \{b\}, \{b\} \xrightarrow{c} \{a, c\}\}$ are admissible and existing causal inference approaches are not made with such relations in mind. Simplifying the causal structure is also risking removing the need for multivariate causal relations.

Acknowledgements

This work was supported partly by the french PIA project « Lorraine Université d’Excellence », reference ANR-15-IDEX-04-LUE.

References

1. Konstantin Bazhanov and Sergei Obiedkov. Optimizations in computing the Duquenne–Guigues basis of implications. *Annals of mathematics and artificial intelligence*, 70(1):5–24, 2014.
2. Alexandre Bazin and Jean-Gabriel Ganascia. Computing the Duquenne–Guigues basis: An algorithm for choosing the order. *International Journal of General Systems*, 45(2):57–85, 2016.
3. Claude Berge. *Hypergraphs*, North-Holland Mathematical Library. 1989.
4. Karell Bertet and Bernard Monjardet. The multiple facets of the canonical direct unit implicational basis. *Theoretical Computer Science*, 411(22-24):2155–2166, 2010.

5. Patrick Blöbaum, Dominik Janzing, Takashi Washio, Shohei Shimizu, and Bernhard Schölkopf. Cause-effect inference by comparing regression errors. In *International Conference on Artificial Intelligence and Statistics*, pages 900–909. PMLR, 2018.
6. Kailash Budhathoki and Jilles Vreeken. MDL for causal inference on discrete data. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 751–756. IEEE, 2017.
7. Kailash Budhathoki and Jilles Vreeken. Origo: causal inference by compression. *Knowledge and Information Systems*, 56(2):285–307, 2018.
8. Zhitang Chen, Kun Zhang, and Laiwan Chan. Nonlinear causal discovery for high dimensional data: A kernelized trace method. In *2013 IEEE 13th International Conference on Data Mining*, pages 1003–1008. IEEE, 2013.
9. Felix Distel and Barış Sertkaya. On the complexity of enumerating pseudo-intents. *Discrete Applied Mathematics*, 159(6):450–466, 2011.
10. Ellery Eells. *Rational decision and causality*. Cambridge University Press, 2016.
11. Bernhard Ganter and Sergei O Kuznetsov. Pattern structures and their projections. In Harry S. Delugach and Gerd Stumme, editors, *Conceptual Structures: Broadening the Base, 9th International Conference on Conceptual Structures, ICCS 2001, Stanford, CA, USA, July 30-August 3, 2001, Proceedings*, volume 2120 of *Lecture Notes in Computer Science*, pages 129–142. Springer, 2001.
12. Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer - Verlag, 1999.
13. Tian Gao and Qiang Ji. Local causal discovery of direct causes and effects. *Advances in Neural Information Processing Systems*, 28:2512–2520, 2015.
14. Liqiang Geng and Howard J Hamilton. Interestingness measures for data mining: A survey. *ACM Computing Surveys (CSUR)*, 38(3):9–es, 2006.
15. Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
16. Jean-Louis Guigues and Vincent Duquenne. Familles minimales d’implications informatives résultant d’un tableau de données binaires. *Mathématiques et Sciences humaines*, 95:5–18, 1986.
17. Joseph Y Halpern. Sufficient conditions for causality to be transitive. *Philosophy of Science*, 83(2):213–226, 2016.
18. Harold Hotelling. Relations between two sets of variates. *Biometrika*, 1936.
19. Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21:689–696, 2008.
20. Shoubo Hu, Zhitang Chen, Vahid Partovi Nia, Laiwan Chan, and Yanhui Geng. Causal inference and mechanism clustering of a mixture of additive noise models. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 5212–5222, 2018.
21. D Janzing, P Hoyer, and B Schölkopf. Telling cause from effect based on high-dimensional observations. In *27th International Conference on Machine Learning (ICML 2010)*, pages 479–486. Omnipress, 2010.
22. Dominik Janzing, Joris Mooij, Kun Zhang, Jan Lemeire, Jakob Zscheischler, Povilas Daniušis, Bastian Steudel, and Bernhard Schölkopf. Information-geometric approach to inferring causal directions. *Artificial Intelligence*, 182:1–31, 2012.
23. Dominik Janzing and Bernhard Schölkopf. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory*, 56(10):5168–5194, 2010.

24. Samantha Kleinberg. *Why: A guide to finding and using causes*. "O'Reilly Media, Inc.", 2015.
25. Sergei O Kuznetsov. On the intractability of computing the Duquenne–Guigues base. *Journal of Universal Computer Science*, 10(8):927–933, 2004.
26. Sergei O Kuznetsov and Sergei Obiedkov. Counting pseudo-intents and #P-completeness. In *Formal Concept Analysis*, pages 306–308. Springer, 2006.
27. Alexander Marx and Jilles Vreeken. Identifiability of cause and effect using regularized regression. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 852–861, 2019.
28. Neil McDonnell. Transitivity and proportionality in causation. *Synthese*, 195(3):1211–1229, 2018.
29. Osman Mian, Alexander Marx, and Jilles Vreeken. Discovering fully oriented causal networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
30. Joris M Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *The Journal of Machine Learning Research*, 17(1):1103–1204, 2016.
31. Jonas Peters, Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.
32. Uwe Ryssel, Felix Distel, and Daniel Borchmann. Fast algorithms for implication bases and attribute exploration using proper premises. *Annals of Mathematics and Artificial Intelligence*, 70(1):25–53, 2014.
33. Eleni Sgouritsa, Dominik Janzing, Philipp Hennig, and Bernhard Schölkopf. Inference of cause and effect with unsupervised inverse regression. In *Artificial intelligence and statistics*, pages 847–855. PMLR, 2015.
34. Peter Spirtes, Clark Glymour, Richard Scheines, Stuart Kauffman, Valerio Aimale, and Frank Wimberly. Constructing bayesian network models of gene expression networks from microarray data. 2000.
35. Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT press, 2000.
36. Xiaohai Sun, Dominik Janzing, and Bernhard Schölkopf. Causal reasoning by evaluating the complexity of conditional densities with kernel methods. *Neuro-computing*, 71(7-9):1248–1256, 2008.
37. Jilles Vreeken. Causal inference by direction of information. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 909–917. SIAM, 2015.
38. Kui Yu, Xianjie Guo, Lin Liu, Jiuyong Li, Hao Wang, Zhaolong Ling, and Xindong Wu. Causality-based feature selection: Methods and evaluations. *ACM Computing Surveys (CSUR)*, 53(5):1–36, 2020.
39. K Zhang and A Hyvärinen. On the identifiability of the post-nonlinear causal model. In *25th Conference on Uncertainty in Artificial Intelligence (UAI 2009)*, pages 647–655. AUAI Press, 2009.
40. Jakob Zscheischler, Dominik Janzing, and Kun Zhang. Testing whether linear equations are causal: a free probability theory approach. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 839–848, 2011.