
Typing assumptions improve identification in causal discovery

Philippe Brouillard
ServiceNow
Mila, Université de Montréal

Perouz Taslakian
ServiceNow

Alexandre Lacoste
ServiceNow

Sébastien Lachapelle
Mila, Université de Montréal

Alexandre Drouin
ServiceNow

Abstract

Causal discovery from observational data is a challenging task to which an exact solution cannot always be identified. Under assumptions about the data-generative process, the causal graph can often be identified up to an equivalence class. Proposing new realistic assumptions to circumscribe such equivalence classes is an active field of research. In this work, we propose a new set of assumptions that constrain possible causal relationships based on the nature of the variables. We thus introduce *typed directed acyclic graphs*, in which variable types are used to determine the validity of causal relationships. We demonstrate, both theoretically and empirically, that the proposed assumptions can result in significant gains in the identification of the causal graph.

1 Introduction

Can the temperature of a city alter its altitude (Peters et al., 2017)? Can a light bulb change the state of a switch? Can the brakes of a car be activated by their indicator light (de Haan et al., 2019)? Chances are, you did not need to think very hard to answer these questions, since you intuitively understand the implausibility of causal relationships between certain *types of entities*. This form of prior knowledge has been shown to play a key role in causal reasoning (Griffiths et al., 2011; Schulz & Gopnik, 2004; Gopnik & Sobel, 2000). In fact, in the absence of evidence (e.g., data), humans tend to reason inductively and use domain knowledge to generalize known causal relationships to new, similar, entities (Kemp et al., 2010).

Nonetheless, the elucidation of causal relationships often goes beyond human intuition. The abundance of large-scale scientific endeavours to understand the causes of diseases (IKGP, 2010) or natural phenomena (Runge et al., 2019) are good examples. In such cases, computational methods for *causal discovery* may help reveal causal relationships based on patterns of association in data (see Heinze-Deml et al. (2018) for a review). The most common setting consists of representing causal relationships as a directed acyclic graph where vertices correspond to variables of interest and edges indicate causal relationships. Additional assumptions, like the *faithfulness* condition, are then made to enable reasoning about graph structures based on conditional independences in the data. While these enable data-driven causal discovery, the underlying causal graph can only be identified up to its *Markov equivalence class* (Peters et al., 2017), which can often be very large (He et al., 2015) thus leaving many edges unoriented.

Inspired by how humans use types to reason about causal relationships, this work explores if prior knowledge about the nature of the variables can help reduce the size of such equivalence classes.

Building on the theoretical foundations of causal discovery in directed acyclic graphs, we propose a new theoretical framework for the case where *variables are labeled by type*. Such types can be attributed based on prior knowledge, e.g., via a domain expert. We then impose assumptions on how types can interact with each other, which constrains the space of possible graphs and leads to reduced equivalence classes. We show, both theoretically and empirically, that when such assumptions hold in the data, significant gains in the identification of causal relationships can be made.

Contributions:

- We propose a new set of assumptions for causal discovery, based on variable types (Section 4).
- We present a simple algorithm to incorporate these assumptions in causal discovery (Section 5).
- We prove theoretical results that guarantee convergence of the equivalence class to a singleton, i.e., identification, when the number of vertices tends to infinity and the number of types is fixed (Section 6).
- We present an empirical study that supports our theoretical results (Section 7).

This is a work in progress. At the end of each section that presents a contribution, we include a short *work in progress* paragraph to highlight extensions that we are currently exploring.

2 Problem formulation

Causal graphical models. In this work, we adopt the framework of causal graphical models (CGM) (Peters et al., 2017). Let $X = (X_1, \dots, X_d)$ be a random vector with distribution P_X . Let $G = (V, E)$ be a directed acyclic graph (DAG) with vertices $V = \{v_1, \dots, v_d\}$. Each vertex $v_i \in V$ is associated to variable X_i in G , and a directed edge $(v_i, v_j) \in E$ represents a direct causal relationship from X_i to X_j . We assume that P_X can be factorized according to G , that is,

$$p(x_1, \dots, x_d) = \prod_{j=1}^d p(x_j | pa_j^G),$$

where pa_j^G denotes the parents¹ of X_j in G . From this graph, it is possible to answer causal questions (e.g., via do-calculus (Pearl, 1995)). However, in many situations, the structure of G is unknown and must be inferred from data.

Causal discovery. The task of causal discovery consists of learning the structure of G based on observations from P_X . Some assumptions are required to make this possible. By adopting the CGM framework, we assume: (i) *causal sufficiency*, which states there is no unobserved variable that causes more than one variable in X and (ii) the *causal Markov property*, which states that $X_i \perp\!\!\!\perp_G X_j | Z \implies X_i \perp\!\!\!\perp_{P_X} X_j | Z$, where Z is a set composed of variables in X , $X_i \perp\!\!\!\perp_G X_j | Z$ indicates that X_i and X_j are d -separated by Z in G , and $X_i \perp\!\!\!\perp_{P_X} X_j | Z$ indicates that X_i and X_j are independent conditioned on Z . Additionally, we assume (iii) *faithfulness*, which states that $X_i \perp\!\!\!\perp_{P_X} X_j | Z \implies X_i \perp\!\!\!\perp_G X_j | Z$. Hence, conditional independences in the data can be used to learn about d -separations in G .

Equivalence classes. Even with these assumptions, G can only be recovered up to a *Markov equivalence class* (MEC), which is a set containing all the DAGs that can represent the same distributions as G . The MEC is often characterized graphically using an *essential graph* or *Completed Partially Directed Acyclic Graph* (CPDAG), which corresponds to the union of all Markov equivalent DAGs (Andersson et al., 1997). In some cases, e.g., for sparse graphs, the size of the MEC can be huge (He & Yu, 2016; He et al., 2015), significantly limiting inference about the direction of edges in G . Hence, it is a problem of key importance to find new realistic assumptions to shrink the equivalence class.

There have been a wealth of approaches to this problem. For instance, some have made progress by including data collected under intervention (Hauser & Bühlmann, 2012), making assumptions about the functional form of causal relationships (e.g. Peters et al. (2014); Peters & Bühlmann (2014);

¹We are referring to the corresponding vertex v_j in G .




Figure 1: **(a)** Representation of t-edges orientations, where colors represent the different types a , b , and c . **(b)** Representation of a t-DAG that is consistent and follows the orientation of the t-edges in (a). **(c)** Representation of a t-DAG that is not consistent: the red edge $c_2 \rightarrow b_1$ is not consistent with $b_1 \rightarrow c_1$, and the $c_2 \rightarrow c_1$ edge violates Definition 3.

Shimizu et al. (2006)), or including background knowledge on the direction of edges Meek (1995). In this work, we propose an alternative approach, based on background knowledge, where types are attributed to variables and the interaction between types is constrained.

3 Related works

The inclusion of background knowledge in causal discovery aims to reduce the size of the solution space by adding or ruling-out causal relationships based on expert knowledge. Several forms of background knowledge have been proposed, which place various levels of burden on the expert. Below, we review those most relevant to our work (see Constantinou et al. (2021) for a review).

Hard background knowledge. This type of background knowledge is “hard” in the sense that it *must be respected* in the inferred graph structures. Previous works have considered: sets of forbidden and known edges (Meek, 1995), a known ordering of the variables (Cooper & Herskovits, 1992), partial orderings of the variables (Andrews, 2020; Scheines et al., 1998), and ancestral constraints (Li & Beek, 2018; Chen et al., 2016). Among these, partial orderings (or *tiered background knowledge*) are the most similar to our contribution. In this setting, it is assumed that an expert partitions the variables into sets called *tiers*, and orders the tiers such that variables in later tier cannot cause variables in an earlier tier. In contrast, while we require the expert to partition variables into sets (by type), we do not assume that an ordering is known *a priori*. In the context of causal discovery, the possible inter-type interactions would be learned with the graph structure.

Soft background knowledge. A setting similar to ours, where the type of each variable dictates its possible causal relations, was presented in Mansinghka et al. (2012). They propose a Bayesian method to use this prior knowledge in causal discovery. Their work shows the benefits of such priors, but they do not investigate this space of graphs and their properties w.r.t. to structure identifiability.

Realism of assumptions. Interestingly, several recent works applying causal discovery to real-world problems have used expert knowledge that is compatible with our proposed framework. For example, in their work on Alzheimer’s disease, Shen et al. (2020) claim “edges from biomarkers or diagnosis to demographic variables are prohibited” and “edges among demographic variables are prohibited”; clearly reasoning about relationships between types of variables. Similarly, the work of Flores et al. (2011) outlines an application of tiered background knowledge in a medical case study. Converting this setting to ours simply involves considering each tier as a variable type. Hence, it is clear that the *typing assumptions*, and associated theoretical results, that we propose are realistic and that they are applicable in practice.

4 Typed directed acyclic graphs

Our work builds on two fundamental structures: *typed* directed acyclic graphs (t-DAG), which are essentially DAGs with typed vertices; and *t-edges*, which are sets of edges relating vertices of given types. Formal definitions follow.

Definition 1 (t-DAG). A t-DAG is a tuple (V, E) with a mapping $T : V \rightarrow \mathcal{T}$ such that (V, E) forms a DAG and $T(v_i) \in \mathcal{T}$ is the type of vertex $v_i \in V$, with $\mathcal{T} = \{t_1, \dots, t_k\}$.

Definition 2 (t-edge). A t-edge is a set of edges between vertices of given types in a t-DAG. More formally, $E(t_1, t_2) = \{(v_i, v_j) \in E \mid T(v_i) = t_1, T(v_j) = t_2\}$ for any pair of types $t_1, t_2 \in \mathcal{T}$.

For example, the graphs illustrated in Fig. 1 (b) and (c) are t-DAGs where colors represent types and the set $E(a, c) = \{(a_1, c_1), (a_1, c_2)\}$ is a t-edge between types a and c .

4.1 Assumptions on type interactions

We now introduce the *type consistency* assumption, which constrains the possible interactions between variables of different types.

Definition 3 (Consistent t-DAG). A consistent t-DAG is a t-DAG where for every t-edge $E(t_1, t_2) \neq \emptyset$, we have that $E(t_2, t_1) = \emptyset$. We refer to this additional constraint as *type consistency*.

For conciseness, we denote $E(t_1, t_2) \neq \emptyset$ as $t_1 \xrightarrow{t} t_2$. Note that this definition implies that there are no t-edges between vertices of the same type in a consistent t-DAG.

In Fig. 1 (b), we present an example of a consistent t-DAG. In contrast, the t-DAG shown in Fig. 1 (c) is not consistent: the t-edge $E(c, b)$ (purple to white) contains the edge (c_2, b_1) , while the reverse t-edge, $E(b, c)$, is not empty since it contains (b_1, c_1) . Notice how, the orientation of all t-edges, summarized in Fig. 1 (a), fully determines the orientation of edges in a consistent t-DAG.

Note that alternative assumptions could have been considered. For instance, we could have assumed that t-edges form a DAG (i.e. the types have a partial ordering). However, the assumptions considered here are less restrictive and, as we demonstrate later, lead to interesting results.

4.2 Equivalence classes for consistent t-DAGs

We define the equivalence classes MEC and t-MEC as the set of DAGs and the set of consistent t-DAGs that are Markov equivalent, respectively.

Definition 4 (MEC). The MEC of a t-DAG D_t is $M(D_t) = \{D' \mid D' \sim D_t\}$ where “ \sim ” denotes Markov equivalence.

Definition 5 (t-MEC). The t-MEC of a consistent t-DAG D_t is $M_t(D_t) = \{D'_t \mid D'_t \overset{t}{\sim} D_t\}$ where “ $\overset{t}{\sim}$ ” denotes Markov equivalence limited to consistent t-DAGs with the same mapping T as D_t .

To represent an equivalence class, we can use an *essential graph*, which corresponds to the union of equivalent DAGs.

Definition 6 (Essential graph). The essential graph D^* associated to the consistent t-DAG D_t is:

$$D^* := \cup_{D \in M(D_t)} D$$

The union over graphs is defined as the union of their vertices and edges: $G_1 \cup G_2 := (V_1 \cup V_2, E_1 \cup E_2)$. Also, if $(v_i, v_j) \in G$ and its opposite $(v_j, v_i) \in G$, then the edge is considered to be undirected.

Definition 7 (t-Essential graph). The t-essential graph D_t^* associated to the consistent t-DAG D_t is

$$D_t^* := \cup_{D \in M_t(D_t)} D$$

4.3 Space of consistent t-DAGs and size of t-MEC

We consider some statements that can directly be made about the space of consistent t-DAGs and the size of t-MEC with respect to their non-typed counterparts.

By definition, the space of consistent t-DAGs (with k types and d vertices), $\mathcal{T}\text{-DAG}(k, d)$, is included in the space of DAGs with the same number of nodes and types, $\mathcal{DAG}(k, d)$. Moreover, as stated in Proposition 1, for a number of types smaller than the number of vertices, the space of consistent t-DAGs is strictly smaller than that of DAGs.

Proposition 1. $\mathcal{T}\text{-DAG}(k, d) \subset \mathcal{DAG}(k, d)$ for $k < d$, where d is the number of vertices and k is the number of types.

Notice that, in the limit case $k = d$, i.e. every vertex has a different type, the t-DAG is always consistent and $\mathcal{T}\text{-DAG}(k, d) = \mathcal{DAG}(k, d)$.




Figure 2: **(a)** The two-type fork structure. In this illustration, a_1 and a_2 are of one type (purple) and b_1 is of another type (orange). **(b)** An orientation rule that can be used in combination with Meek (1995)'s rules to orient two-type forks.

The t-essential graph is identical to the essential graph except for the additional edges which can be oriented thanks to the type consistency assumption. This is summarized in Proposition 2.

Proposition 2. *Let D_t^* and D^* be, respectively, the t-essential and essential graphs of an arbitrary consistent t-DAG D_t . Then, $D_t \subseteq D_t^* \subseteq D^*$.*

From the t-essential graph, we can easily derive an upper bound on the size of the t-MEC based on the number of undirected edges.

Proposition 3 (Upper bound on the size of the t-MEC). *For any consistent t-DAG D_t , we have $|M_t(D_t)| \leq 2^u$, where u is the number of undirected t-edges in the t-essential graph of D_t .*

From this bound, we can also directly conclude that if the t-essential graph contains no undirected t-edges, $|M_t(D_t)| = 1$. In other words, D_t is identified.

Work in progress The type consistency assumption proposed above does not allow for edges to exist between variables of the same type. We are currently working on a relaxed version that allows such edges. Concretely, this involves minor modifications to Definition 3, Proposition 1, and Proposition 3. Notably, Proposition 3 needs to include an additional term to account for undirected edges between variables of the same type. Conclusions such as the fact that $\mathcal{T}\text{-DAG}(k, d) \subset \text{DAG}(k, d)$ (Proposition 1) will still hold, with only minor changes to the conditions.

5 An algorithm to find the t-essential graph

To make practical use of the above definitions, we seek an algorithm that can recover the t-essential graph based on a set of observational data and variable types attributed by a domain expert. Following previous work, it would be intuitive to make use of Meek (1995)'s rules as follows:

1. Recover the essential graph by applying an algorithm like PC (Spirtes et al., 2000) to the data.
2. Enforce type consistency: If there exists an oriented edge between any pair of variables with types $t_1, t_2 \in \mathcal{T}$ in the graph of the previous step, then conclude $t_1 \xrightarrow{t} t_2$, i.e., orient all edges between these types.
3. Apply Meek (1995)'s rules to propagate the edge orientations derived in Step (2) (see their Section 2.1.2).
4. Repeat from Step (2) until the graph is unchanged.

This algorithm is sound, i.e., all edges that it orients are also oriented in the t-essential graph. However, it is not complete — some edges that should be oriented in the t-essential graph will remain unoriented.

To see this, consider the following simple example of a consistent three-vertex t-DAG that we call the *two-type fork* (illustrated in Fig. 2 a). The essential graph for this t-DAG obtained at Step (1), would be completely undirected since it contains no immoralities. This would result in a case where none of Meek (1995)'s rules are applicable and thus, Step (2) would not orient any edges. The algorithm would therefore stop and return a fully undirected graph. However, according to the following proposition, the two-type fork should have been oriented.

Proposition 4. *If a consistent t-DAG D_t contains vertices a_1, a_2, b_1 with types $T(a_1) = T(a_2) = a$ and $T(b_1) = b$, with edges $a_1 \leftarrow b_1 \rightarrow a_2$, then the t-edge $b \xrightarrow{t} a$ is directed in the t-essential graph, i.e., the direction of causation between types b and a is known.*

The proof follows the argument that if the true orientation were $a \xrightarrow{t} b$, this would create an immorality, which in turn would orient the edges in the essential graph. Because we assume type consistency (Definition 3), the only other possible orientation is $b \xrightarrow{t} a$.

Nevertheless, adding a rule to orient such structures in Step (2), as illustrated in Fig. 2 b), is not sufficient to obtain a complete algorithm. In Appendix B, we show more complex cases (not local and involving multiple t-edges) that are not covered, even with this additional orientation rule.

It remains an open question whether one could design a polynomial-time algorithm to find the t-essential graph as Meek (1995) and Andersson et al. (1997) did for essential graphs. For now, we propose the following algorithm, which has a time complexity exponential in the number of types:

1. Run the previously described (incomplete) algorithm, adding the two-type fork case to the orientation rules in Step (2).
2. Enumerate all t-DAGs that can be produced by orienting edges in the graph resulting from the previous step. Reject any inconsistent t-DAG.
3. Take the union of all these t-DAGs (see Definition 7) to obtain the t-essential graph.

In practice, the non-polynomial time complexity was found to be non-prohibitive in our experiments.

Work in progress The algorithm presented in this section assumes that the practitioner has access to the essential graph and then further orients edges using type consistency. What if the true essential graph is unknown? One could use a causal discovery algorithm, but for finite sample sizes, the inferred essential graph could contain errors, which may contradict type consistency. We are thus developing a causal discovery algorithm (with standard consistency results) that is guaranteed to produce a valid t-essential graph.

6 Identification for random graphs

In this section, we explore the benefits in identification that result from using variable typing in a class of graphs generated at random based on a process inspired by the Erdős-Rényi random graphs model Erdős & Rényi (1959).

Assume we are given a set of k types t_1, \dots, t_k , probabilities $p_1, \dots, p_k \in (0, 1]^k$ of observing each type s.t. $p_1 + \dots + p_k = 1$, and a t -interaction matrix $A \in [0, 1]^{k \times k}$ where each cell (i, j) defines the probability p_{ij} that a variable of type t_i is a direct cause of a variable of type t_j . As per Definition 3 (type consistency), we impose that if $p_{ij} > 0$, then $p_{ji} = 0$.

Definition 8. *(Random sequence of growing t-DAG) We define a random sequence of t-DAGs $(D_t^n)_{n=0}^\infty$ with $D_t^n = (V^n, E^n)$ and $|V^n| = n$, such that $D_t^0 = (\emptyset, \emptyset)$. Each new t-DAG D_t^n in the sequence is obtained from D_t^{n-1} as follows: Create a new vertex v_n and sample its type from a categorical distribution with probabilities p_1, \dots, p_k ; Let $V^n = V^{n-1} \cup \{v_n\}$; To obtain E^n , for every vertex $v_i \in V^{n-1}$, add the edge (v_i, v_n) to E^{n-1} with probability $p_{T(v_i), T(v_n)}$.*

Our main theorem below states that, as we add more vertices to a random sequence of t-DAGs, we get closer to identifiability. We defer the proof to Appendix A.5.

Theorem 1. *Let $(D_t^n)_{n=0}^\infty$ be a random sequence of growing t-DAGs as defined in Definition 8. Then the size of the t-MEC converges to 1 exponentially fast, i.e., there exist a C and n_0 such that for all $n > n_0$,*

$$P(|M_t(D_t^n)| > 1) < Ce^{-rn},$$

where $r = -\frac{1}{2} \max(\ln(1 - p_j p_{ij}), \ln(1 - p_i))$.

To give an intuition for the proof, recall that the t-MEC shrinks every time we orient a t-edge. In addition, Proposition 3 tells us that we can orient a t-edge if we observe a two-type fork structure. We

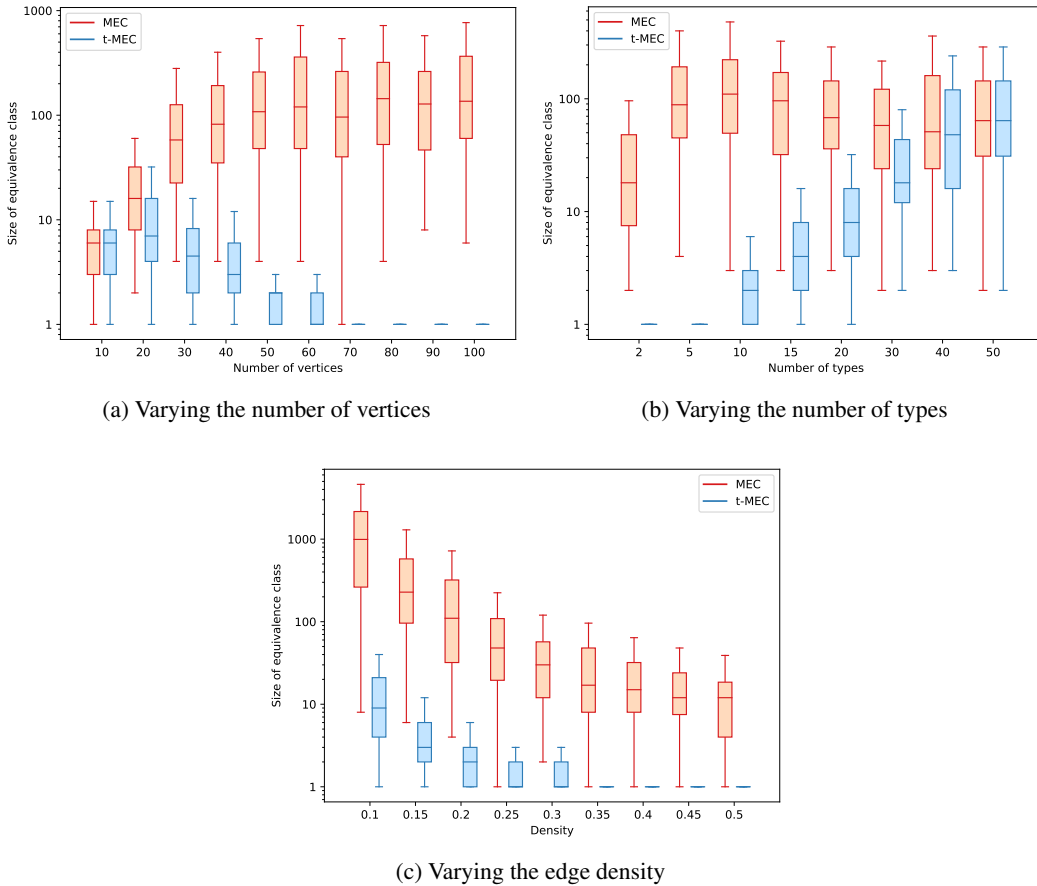


Figure 3: Size of the equivalence classes (MEC and t-MEC) w.r.t. various graph properties. a) Number of vertices: the t-MEC shrinks with the number of vertices, while the MEC does not. b) Number of types: The t-MEC grows with the number of types, since less constraints are placed on the structure of the graphs. The size of the MEC is *mostly* constant. c) Edge density (i.e., probability of connectivity p): The MEC and t-MEC both shrink as connectivity increases.

thus argue that, as we add more vertices, the probability of observing a two-type fork (Proposition 4) for arbitrary type pairs converges to 1. This argument relies on the fact that the number of types k remains constant throughout as the random t-DAG grows.

Work in progress Could we get a result similar to Theorem 1 if we modified Definition 3 to allow edges between variables of the same type? Preliminary empirical results (not shown) suggest that the size of the t-MEC does converge to something much smaller than the size of the MEC. We are making progress towards a formal result that explains our empirical observations, which we will include in future revisions of this work.

7 Experiments

We perform multiple experiments to understand how the size of MECs and t-MECs compare ². For a given t-DAG, the size of the MEC is obtained by finding its CPDAG (without considering types) and enumerating all DAGs that do not introduce a cycle or add an immorality. For the same t-DAG, the size of the t-MEC is obtained by applying the algorithm described in Section 5.

²The code used to perform these experiments is available at <https://github.com/ElementAI/typed-dag>.

The t-DAGs that we consider are randomly generated according to the process described at Definition 8. Let k be the number of types in the t-DAG. We attribute uniform probability to each type, i.e., $p_i = 1/k, \forall i \in \{1, \dots, k\}$. The t-interaction matrix A is defined as follows. For each pair of types (t_i, t_j) , s.t., $i \neq j$, the direction of the t-edge is sampled randomly with uniform probability and we use a fixed probability of interaction p , which controls the graph density. For example, if the direction $t_i \xrightarrow{t} t_j$ is sampled, then $A_{ij} = p$ and $A_{ji} = 0$. In what follows, unless otherwise specified, the number of vertices is 50, the number of types is 10, and the probability of interaction is $p = 0.2$.

In Fig. 3a, 3b, and 3c, the size of the equivalence classes are compared with respect to the number of vertices, the number of types, and the density, respectively. All boxplots are calculated over 100 random consistent t-DAGs. First, in Fig. 3a we see that as the number of vertices increases (and the number of types remains constant), the size of the t-MEC converges to 1, as demonstrated in Section 6. In contrast, the size of the MEC first increases and then remains near constant. Notice how the size of the MEC and the t-MEC are identical when the number of vertices equals the number of types; this is because type consistency does not constrain the graph structure. Second, in Fig. 3b, as the number of types increases, the size of the t-MEC increases. This is expected, because as the number of types approaches the number of vertices, type consistency imposes less structural constraints. Further, notice that the size of the MEC changes with the number of types, even though it is agnostic to type consistency. This is because t-DAGs with fewer types (e.g., 2) are more likely to contain immoralities (see Proposition 1), leading to smaller MECs. Third, in Fig. 3c, as the density increases, the size of the MEC and the t-MEC both decrease. This is in line with the observations of He et al. (2015).

In summary, all our experiments indicate that the size of the t-MEC is smaller or equal to that of the MEC for random t-DAGs. The difference is particularly striking when the number of types is low and the number of vertices is high. Of particular interest are the results shown in Fig. 3a, as they provide empirical evidence for the correctness of Theorem 1.

Work in progress As stated in Section 5, we are developing a causal discovery algorithm to learn consistent t-DAGs from data. Future revisions of this work will include structure learning experiments on simulated and real datasets, and a comparison to state-of-the-art causal discovery algorithms that do not exploit type consistency as part of their graph learning process.

8 Discussion

In this work, we addressed an important problem in causal discovery: it is often impossible to identify the causal graph precisely, due to the size of its Markov equivalence class. This is particularly true for sparse graphs, where the size of the MEC grows super-exponentially with the number of vertices (He et al., 2015). In this sense, we proposed a new type of assumption based on variable types, which we formalized as typed directed acyclic graphs (t-DAGs). Our theoretical and empirical results clearly demonstrate that there exists conditions in which our variable-typing assumptions greatly shrink the size of the equivalence class. Hence, when such assumptions hold in the data, gains in identification are to be expected.

We note that the new assumptions that we introduce can be used in conjunction with other strategies to shrink the size of the equivalence class, such as considering interventions (Hauser & Bühlmann, 2012), hard background knowledge on the presence/absence of edges (Meek, 1995), or functional-form assumptions (Peters et al., 2014; Peters & Bühlmann, 2014; Shimizu et al., 2006). Moreover, our assumptions could be used with methods that estimate treatment effects from equivalence classes, such as IDA (Perkovic et al., 2017), to improve their accuracy.

We believe that this work may stimulate new advances at the intersection of machine learning and causality (Schölkopf et al., 2021; Schölkopf, 2019). In fact, machine learning algorithms excel at classification, and thus it may be interesting to explore a setting where the variable types are learned based on some features. Type assignments could be learned in parallel with causal discovery, using recent methods for differentiable causal discovery (Brouillard et al., 2020; Zheng et al., 2018). This may further reduce the burden on the human expert in cases where types are hard to assign. As an example, consider the task of learning causal models of gene regulatory networks. One could train a model to assign types to genes based on their categorization in the gene ontology (Gene Ontology Consortium, 2004) or on features of their DNA sequence.

Additionally, an interesting future direction would be to use our typing assumptions to perform causal discovery on multiple graphs at once, i.e., *multi-task causal discovery*. In fact, assume that we are given data for multiple groups of variables that correspond to disjoint systems (no interactions across groups), but that share similar types. It would be possible to use type consistency (Definition 3) to propagate t-edge orientations across graphs.

In conclusion, the results reported in this work show that considering typing assumptions has the potential to improve identification in causal discovery. However, we barely scratched the surface of what is possible. Future work will include extensive experiments to put our theoretical work into practice on real-world datasets and will further explore the aforementioned directions.

Acknowledgements

The authors are grateful to Assya Trofimov, David Berger, and Jean-Philippe Reid for helpful comments and suggestions. Sébastien Lachapelle is supported by an IVADO Excellence PhD scholarship.

References

- 1000 Genomes Project Consortium (1KGP) and others. A map of human genome variation from population-scale sequencing. *Nature*, 467(7319):1061, 2010.
- Andersson, S. A., Madigan, D., Perlman, M. D., et al. A characterization of markov equivalence classes for acyclic digraphs. *Annals of statistics*, 25(2):505–541, 1997.
- Andrews, B. On the completeness of causal discovery in the presence of latent confounding with tiered background knowledge. In *The 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 4002–4011. PMLR, 2020.
- Brouillard, P., Lachapelle, S., Lacoste, A., Lacoste-Julien, S., and Drouin, A. Differentiable causal discovery from interventional data. In *Advances in Neural Information Processing Systems 33*, 2020.
- Chen, E. Y., Shen, Y., Choi, A., and Darwiche, A. Learning bayesian networks with ancestral constraints. In *Advances in Neural Information Processing Systems 29*, pp. 2325–2333, 2016.
- Constantinou, A. C., Guo, Z., and Kitson, N. K. Information fusion between knowledge and data in bayesian network structure learning. *arXiv preprint arXiv:2102.00473*, 2021.
- Cooper, G. F. and Herskovits, E. A bayesian method for the induction of probabilistic networks from data. *Machine learning*, 9(4):309–347, 1992.
- de Haan, P., Jayaraman, D., and Levine, S. Causal confusion in imitation learning. In *Advances in Neural Information Processing Systems 32*, pp. 11693–11704, 2019.
- Erdős, P. and Rényi, A. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- Flores, M. J., Nicholson, A. E., Brunskill, A., Korb, K. B., and Mascaro, S. Incorporating expert knowledge when learning bayesian network structure: a medical case study. *Artificial intelligence in medicine*, 53(3):181–204, 2011.
- Gene Ontology Consortium. The gene ontology (go) database and informatics resource. *Nucleic acids research*, 32(suppl_1):D258–D261, 2004.
- Gopnik, A. and Sobel, D. M. Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child development*, 71(5):1205–1222, 2000.
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., and Gopnik, A. Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*, 35(8):1407–1455, 2011.

- Hauser, A. and Bühlmann, P. Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 13(1):2409–2464, 2012.
- He, Y. and Yu, B. Formulas for counting the sizes of markov equivalence classes of directed acyclic graphs. *arXiv preprint arXiv:1610.07921*, 2016.
- He, Y., Jia, J., and Yu, B. Counting and exploring sizes of markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 16(1):2589–2609, 2015.
- Heinze-Deml, C., Maathuis, M. H., and Meinshausen, N. Causal structure learning. *Annual Review of Statistics and Its Application*, 5:371–391, 2018.
- Kemp, C., Goodman, N. D., and Tenenbaum, J. B. Learning to learn causal models. *Cognitive Science*, 34(7):1185–1243, 2010.
- Li, A. and Beek, P. Bayesian network structure learning with side constraints. In *International Conference on Probabilistic Graphical Models*, pp. 225–236. PMLR, 2018.
- Mansinghka, V., Kemp, C., Griffiths, T., and Tenenbaum, J. Structured priors for structure learning. *arXiv preprint arXiv:1206.6852*, 2012.
- Meek, C. Causal inference and causal explanation with background knowledge. *arXiv preprint arXiv:1302.4972*, 1995.
- Pearl, J. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- Perkovic, E., Kalisch, M., and Maathuis, M. H. Interpreting and using cpdags with background knowledge. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2017.
- Peters, J. and Bühlmann, P. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- Peters, J., Mooij, J. M., Janzing, D., and Schölkopf, B. Causal discovery with continuous additive noise models. *The Journal of Machine Learning Research*, 15(1):2009–2053, 2014.
- Peters, J., Janzing, D., and Schölkopf, B. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Runge, J., Bathiany, S., Bollt, E., Camps-Valls, G., Coumou, D., Deyle, E., Glymour, C., Kretschmer, M., Mahecha, M. D., Muñoz-Marí, J., et al. Inferring causation from time series in earth system sciences. *Nature communications*, 10(1):1–13, 2019.
- Scheines, R., Spirtes, P., Glymour, C., Meek, C., and Richardson, T. The tetrad project: Constraint based aids to causal model specification. *Multivariate Behavioral Research*, 33(1):65–117, 1998.
- Schölkopf, B. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- Schulz, L. E. and Gopnik, A. Causal learning across domains. *Developmental psychology*, 40(2):162, 2004.
- Shen, X., Ma, S., Vemuri, P., and Simon, G. Challenges and opportunities with causal discovery algorithms: application to alzheimer’s pathophysiology. *Scientific reports*, 10(1):1–12, 2020.
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., and Kerminen, A. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct):2003–2030, 2006.
- Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. *Causation, prediction, and search*. MIT press, 2000.
- Verma, T. S. and Pearl, J. On the equivalence of causal models. *arXiv preprint arXiv:1304.1108*, 1990.

Zheng, X., Aragam, B., Ravikumar, P., and Xing, E. P. Dags with NO TEARS: continuous optimization for structure learning. In *Advances in Neural Information Processing Systems 31*, pp. 9492–9503, 2018.

A Proofs

A.1 Proof of Proposition 1

$\mathcal{T}\text{-DAG}(k, d) \subset \mathcal{DAG}(k, d)$ for $k < d$, where d is the number of vertices and k is the number of types.

Proof. By definition, a t-DAG is a DAG, so we know that $\mathcal{T}\text{-DAG}(k, d) \subseteq \mathcal{DAG}(k, d)$.

Since $k < d$, by the pigeonhole principle, at least two vertices have the same type. We claim that two vertices of the same type cannot be linked by an edge. To show the claim is true, let D_t be a consistent t-DAG, and assume without loss of generality that $v_1, v_2 \in V$ and $T(v_1) = T(v_2) = a \in \mathcal{T}$. For the sake of contradiction, assume further that $a \xrightarrow{t} a$. Then we have that:

$$E(a, a) = \{(v_i, v_j) \in E \mid T(v_i) = a, T(v_j) = a\}.$$

This means D_t contains both edges $v_1 \rightarrow v_2$ and $v_1 \leftarrow v_2$, contradicting the fact that D_t is consistent.

Therefore, when $k < d$, the space of consistent t-DAGs will exclude at least some DAGs that are in the space of DAGs, giving us the strict subset $\mathcal{T}\text{-DAG}(k, d) \subset \mathcal{DAG}(k, d)$. \square

A.2 Proof of Proposition 2

Let D_t^* and D^* be, respectively, the t-essential and essential graphs of an arbitrary consistent t-DAG D_t . Then, $D_t \subseteq D_t^* \subseteq D^*$.

Proof. It is clear that $D_t \subseteq D_t^*$ and $D_t \subseteq D^*$ since D_t^* and D^* are obtained by undirecting edges from D_t . Also, since enforcing type consistency can only orient more edges in D^* , we have that $D_t^* \subseteq D^*$. \square

A.3 Proof of Proposition 3

For any consistent t-DAG D_t , we have $|M_t(D_t)| \leq 2^u$, where u is the number of undirected t-edges in D_t^* , the t-essential graph of D_t .

Proof. A t-essential graph is the union of consistent t-DAGs. First, note that we do not have to consider every edge of a consistent t-DAG independently since, by consistency, we have that all the edges included in a t-edge of D_t will always take the same orientation. Thus, if a t-edge is undirected in D_t^* , it means that there exists at least one consistent t-DAG in $M_t(D_t)$ for each orientation of the t-edge. Since each of the u undirected t-edges can take on two directions, there are 2^u possible combinations. Note that this is only an upper bound — some of these orientations are not part of the equivalence class, since they create either a cycle or new immoralities not present in D_t . \square

A.4 Proof of Proposition 4

If a consistent t-DAG D_t contains vertices a_1, a_2, b_1 with types $T(a_1) = T(a_2) = a$ and $T(b_1) = b$, with edges $a_1 \leftarrow b_1 \rightarrow a_2$, then the t-edge $b \xrightarrow{t} a$ is directed in the t-essential graph, i.e., the direction of causation between types b and a is known.

Proof. To prove the statement we show that among all possible orientations $b \xrightarrow{t} a$, $b \xleftarrow{t} a$, and $b \leftrightarrow^t a$ of the t-edge, the last two are not valid.

For the sake of contradiction, first assume $b \xleftarrow{t} a$ is directed in the t-essential graph of D_t . This means that there exists a consistent t-DAG D_1 , having t-edge $b \xleftarrow{t} a$, that is Markov equivalent to D_t . Recall that two graphs are Markov equivalent if and only if they have the same skeleton and the same immoralities (Verma & Pearl, 1990). Given that $b \xleftarrow{t} a$, then D_1 has the structure $a_1 \rightarrow b_1 \leftarrow a_2$, which forms an immorality. But since D_t does not contain this immorality, this contradicts the fact that D_1 is Markov equivalent to D_t .

Now, suppose that $b \xleftrightarrow{t} a$ is not directed in the t-essential graph of D_t . This means that there exist two consistent t-DAGs D_1 and D_2 that are Markov equivalent to D_t , having the t-edge orientations $b \xleftrightarrow{t} a$ and $b \xrightarrow{t} a$, respectively. As per the argument in the previous case, the existence of D_1 leads to a contradiction.

Therefore, the only possible orientation for the types a and b is $b \xrightarrow{t} a$. \square

A.5 Proof of Theorem 1

Let $(D_t^n)_{n=0}^\infty$ be a random sequence of growing t-DAG as defined in Definition 8. Then the size of the t-MEC converges to 1 exponentially fast, i.e., there exist a C and n_0 such that for all $n > n_0$,

$$P(|M_t(D_t^n)| > 1) < Ce^{-rn},$$

where $r = -\frac{1}{2} \max(\ln(1 - p_j p_{ij}), \ln(1 - p_i))$.

Proof. To demonstrate this, first recall that the t-MEC shrinks every time we orient a t-edge (Proposition 3). In addition, Proposition 4 tells us that we can orient a t-edge if we observe a two-type fork structure. Therefore, proving the theorem is equivalent to showing that, as n goes to infinity, and for an arbitrary pair of types t_i and t_j with $p_{ij} > 0$, the probability of observing a two-type fork structure converges to 1. Our proof relies on the fact that as n grows, the number of types k remains constant.

Case $k = 1$: Let $D_t^n = (V^n, E^n)$ be a random t-DAG. If D_t^n has vertices of only one type (i.e., $k = 1$), then D_t^n is a set of disconnected vertices; hence $|M_t(D_t^n)| = 1$ and the theorem holds.

Case $k \geq 2$: Let F_{ij} represent the event of observing a two-type fork structure of type t_i and t_j . Our aim is to bound this probability from below and show that it converges exponentially fast to 1 as $n \rightarrow \infty$. We split our search into two steps. Let $Y_i^{n/2}$ be the number of vertices of type t_i after observing the first $n/2$ vertices in the topological ordering induced by D_t^n . Then, we have:

$$P\left(Y_i^{n/2} \geq 1\right) = 1 - \mathcal{B}_{p_i}^{n/2}(0), \quad (1)$$

where $\mathcal{B}_q^n(k)$ is the Binomial distribution of k success out of n trials with probability q . In the second step, we search for at least two vertices of type t_j that are direct causes of v_i , the first vertex of type t_i found in the first step. Let $W_{ij}^{n/2}$ be the count of such vertices among the last $n/2$ vertices, and let $q_{ij} = p_j p_{ij}$ be the probability of sampling a vertex of type t_j and connecting it to v_i . Then, we have:

$$P\left(W_{ij}^{n/2} \geq 2\right) = 1 - \mathcal{B}_{q_{ij}}^{n/2}(0) - \mathcal{B}_{q_{ij}}^{n/2}(1) \quad (2)$$

Combining these two steps, we have:

$$\begin{aligned} P(F_{ij}|n) & \geq P\left(Y_i^{n/2} \geq 1\right) P\left(W_{ij}^{n/2} \geq 2\right) \\ & = \left(1 - \mathcal{B}_{p_i}^{n/2}(0)\right) \left(1 - \mathcal{B}_{q_{ij}}^{n/2}(0) - \mathcal{B}_{q_{ij}}^{n/2}(1)\right) \\ & \geq 1 - \mathcal{B}_{q_{ij}}^{n/2}(0) - \mathcal{B}_{q_{ij}}^{n/2}(1) - \mathcal{B}_{p_i}^{n/2}(0). \end{aligned} \quad (3)$$

The last inequality arises from the fact that the remaining terms are strictly positive. Now, we look at the probability of this event not happening:

$$\begin{aligned} P(\neg F_{ij}|n) & \leq \mathcal{B}_{q_{ij}}^{n/2}(0) + \mathcal{B}_{q_{ij}}^{n/2}(1) + \mathcal{B}_{p_i}^{n/2}(0) \\ & = (1 - q_{ij})^{n/2} + (1 - p_i)^{n/2} + \frac{n}{2} q_{ij} (1 - q_{ij})^{n/2-1}. \end{aligned} \quad (4)$$

Using $\alpha := (1 - q_{ij})$, $\beta := (1 - p_i)$, and $m = n/2$, we can rewrite as follows:

$$\begin{aligned} P(\neg F_{ij}|n) & \leq \alpha^m + \beta^m + m(1 - \alpha)\alpha^{m-1} \\ & = \beta^m + \alpha^m \left(1 + m \frac{(1-\alpha)}{\alpha}\right) \\ & = e^{m \ln \beta} + e^{m \ln \alpha} \left(1 + m \frac{(1-\alpha)}{\alpha}\right) \end{aligned} \quad (5)$$




Figure 4: An example where the algorithm would not orient some edges oriented in the t-essential graph. In this case, the orientation of the t-edge is forced by the fact that the reverse orientation would either create a cycle or a new immorality. **(a)** The original t-DAG, **(b)** the algorithm output (which is supposed to be equal to the t-essential graph), **(c)** the ground-truth t-essential graph

Since $\beta \in (0, 1)$, and $\alpha \in (0, 1)$, this converges to 0 as $m \rightarrow \infty$. Now, we look at the 2 different convergence rates for α and β .

Rate 1. Assume $\ln \beta > \ln \alpha$, there exist constants C_0 and m_0 such that for all $m > m_0$:

$$\begin{aligned} e^{m \ln \beta} + e^{m \ln \alpha} \left(1 + m \frac{(1-\alpha)}{\alpha}\right) &\leq C_0 e^{m \ln \beta} \\ \Rightarrow e^{m(\ln \alpha - \ln \beta)} &\leq \frac{C_0 - 1}{1 + m \frac{1-\alpha}{\alpha}}. \end{aligned} \quad (6)$$

Since $\ln \alpha - \ln \beta < 0$, this is true for $C_0 = 2$ and m_0 sufficiently large.

Rate 2. Assume that $\ln \beta \leq \ln \alpha$, there exist constants C_1 and m_1 such that for all $m > m_1$:

$$\begin{aligned} e^{m \ln \beta} + e^{m \ln \alpha} \left(1 + m \frac{(1-\alpha)}{\alpha}\right) &\leq C_1 e^{m \frac{1}{2} \ln \alpha} \\ \Rightarrow e^{m(\ln \beta - \frac{1}{2} \ln \alpha)} + e^{m \frac{1}{2} \ln \alpha} \left(1 + m \frac{(1-\alpha)}{\alpha}\right) &\leq C_1, \end{aligned}$$

Since $\ln \beta - \frac{1}{2} \ln \alpha < 0$, and $\ln \alpha < 0$, this is true for $C_1 = 1$ and m_1 sufficiently large.

Combining the two rates. By combining the 2 rates, and the fact that $m = n/2$, we have that $P(|M_t(D_t^n)| > 1)$ converges exponentially fast to 0 with rate r as follows:

$$\begin{aligned} r &\geq \frac{1}{2} \min(-\frac{1}{2} \ln \alpha, -\ln \beta) \\ &= -\frac{1}{2} \max(\ln(1 - p_i p_{ij}), \ln(1 - p_i)). \end{aligned}$$

□

B Additional counterexamples

In this section, we give two additional examples where the algorithm without the enumeration presented in Section 5 would not orient some edges that are oriented in the t-essential graph.

Our first example is interesting because it shows that in order to decide the orientation of a t-edge sometimes several t-edges (possibly not local) have to be considered simultaneously. The second counterexample shows that looking only for the direct parent or child of a t-edge is not always sufficient.

The first example is presented in Fig. 4. Note that vertices denoted by the same letter have the same type. The algorithm orients the t-edge $a \xrightarrow{t} c$ since one of its edges in the t-DAG is part of an immorality. All other edges are unoriented because they are not covered by any rules. However, in the t-essential graph (see Fig. 4 c) the t-edge $b \xrightarrow{t} c$ is oriented. To see why this is the case, consider the four possible orientations of the t-edges $a \xleftrightarrow{t} b$ and $b \xleftrightarrow{t} c$ (recall that an orientation cannot create a cycle or a new immorality):

1. $a \xrightarrow{t} b, b \xrightarrow{t} c$ possible.




Figure 5: A second example where the algorithm would not orient some edges oriented in the t-essential graph. In this case, the orientation of the t-edge is forced by the fact that the reverse orientation would create a new immorality. **(a)** The original t-DAG, **(b)** the algorithm output (which is supposed to be equal to the t-essential graph), **(c)** the ground-truth t-essential graph

2. $a \xrightarrow{t} b, b \xleftarrow{t} c$ impossible (creates an immorality that is not present in the original t-DAG).
3. $a \xleftarrow{t} b, b \xrightarrow{t} c$ possible.
4. $a \xleftarrow{t} b, b \xleftarrow{t} c$ impossible (creates a cycle).

In the two configurations that are possible, the t-edge $b \xleftrightarrow{t} c$ is always oriented as $b \xrightarrow{t} c$. Thus, this is an essential edge that should have been recovered by the algorithm.

The second example is presented in Fig. 5. The dashed line between a_1 and a_2 represents a path that does not contain oriented edges in the t-essential graph. Thus, the t-DAG does not contain any immorality. Without loss of generality, let us consider the dashed line as a chain $a_1 \leftarrow c_1 \leftarrow c_2 \leftarrow a_2$. The algorithm does not orient any t-edges because they are not covered by any rules. However, in the t-essential graph (see Fig. 5 c) the t-edge $a \xrightarrow{t} b$ is oriented. Consider the impossible orientation $a \xleftarrow{t} b$. By construction, the t-DAG contains no immorality. Thus, let us orient the edges of the chain as $a_1 \leftarrow c_1 \leftarrow c_2 \leftarrow a_2$ or $a_1 \rightarrow c_1 \rightarrow c_2 \rightarrow a_2$. In both cases, a new immorality is created (respectively, $b_1 \rightarrow a_1 \leftarrow c_1$ and $c_2 \rightarrow a_2 \leftarrow b_2$) leading to a contradiction. Thus, the t-edge has to be oriented as $a \xrightarrow{t} b$.