

I-CAM-UV: Integrating Causal Graphs over Non-Identical Variable Sets Using Causal Additive Models with Unobserved Variables

Hirofumi Suzuki¹, Kentaro Kanamori¹, Takuya Takagi¹, Thong Pham^{2, 5}
Takashi Nicholas Maeda^{3, 5}, Shohei Shimizu^{2, 4, 5}

¹Fujitsu Limited

²Faculty of Data Science, Shiga University

³Computer Centre, Gakushuin University

⁴SANKEN, The University of Osaka

⁵RIKEN AIP

suzuki-hirofumi@fujitsu.com, k.kanamori@fujitsu.com, takagi.takuya@fujitsu.com, thong-pham@biwako.shiga-u.ac.jp,
takashi.maeda@gakushuin.ac.jp, shohei-shimizu@ds.sanken.osaka-u.ac.jp

Abstract

Causal discovery from observational data is a fundamental tool in various fields of science. While existing approaches are typically designed for a single dataset, we often need to handle multiple datasets with non-identical variable sets in practice. One straightforward approach is to estimate a causal graph from each dataset and construct a single causal graph by overlapping. However, this approach identifies limited causal relationships because unobserved variables in each dataset can be confounders, and some variable pairs may be unobserved in any dataset. To address this issue, we leverage Causal Additive Models with Unobserved Variables (CAM-UV) that provide causal graphs having information related to unobserved variables. We show that the ground truth causal graph has structural consistency with the information of CAM-UV on each dataset. As a result, we propose an approach named I-CAM-UV to integrate CAM-UV results by enumerating all consistent causal graphs. We also provide an efficient combinatorial search algorithm and demonstrate the usefulness of I-CAM-UV against existing methods.

Introduction

Identifying causal relationships among variables is a fundamental task in various fields of science. While the most effective approach is a randomized controlled trial, such an approach is often difficult due to cost, ethical, and technical reasons. Thus, the mainstream is *causal discovery* from purely observational data (Spirtes, Glymour, and Scheines 2000): estimating causal relationships as a causal graph, often assumed to be a directed acyclic graph (DAG). Recent works in causal discovery span diverse fields, including environmental science (Runge et al. 2023; Fu et al. 2025b), biology (Fu et al. 2025a; Smith et al. 2025), and materials and drug sciences (Campomanes et al. 2014; Zhou et al. 2025).

In many of the causal discovery examples, the number of variables is relatively small—often around 10—because of limits in what can be measured or easily understood. On the other hand, most of these studies assume that there are

no unobserved confounders, even though such hidden common causes can exist even in small systems. This highlights the need for causal discovery methods that can handle unobserved confounding, especially in small to medium systems.

Moreover, in many situations, multiple datasets share common objectives but differ in measurement settings such as observed variables. If we can leverage such datasets, it is expected to identify more causal relationships over just using a single dataset. Thus, in this study, we consider causal discovery from multiple datasets with non-identical variable sets (Tillman and Eberhardt 2014; Mooij, Magliacane, and Claassen 2020).

Despite the usefulness and demand, existing approaches are typically designed for a single dataset, assuming no unobserved confounders, and only a few approaches tackle the above scenario. ION (Tillman, Danks, and Glymour 2009), IOD (Tillman and Spirtes 2011), and COMBINE (Triantafyllou, Tsamardinos, and Tollis 2010) construct partial ancestral graphs (PAGs) over the integrated variable set by utilizing conditional independence information. In contrast, CD-MiNi (Huang et al. 2020) constructs a full DAG by restricting the causal model to linear non-Gaussian cases. However, these approaches have the following limitations: PAGs contain uncertain information about causal relationships, and CD-MiNi is not suitable for non-linear cases.

Aiming to estimate full DAGs and handle a more expressive model, we focus on the Causal Additive Model (CAM) (Bühlmann, Peters, and Ernest 2014) that captures non-linear causal relationships. One straightforward approach is to estimate a causal graph from each dataset by CAM and integrate them into a single graph. However, this approach identifies limited causal relationships because some variables in each dataset can be confounded with each other. Moreover, if some variable pairs are unobserved in any dataset, their causal relationships can not be identified by direct estimation from the datasets.

To address this issue, we leverage the Causal Additive Model with Unobserved Variables (CAM-UV) (Maeda and Shimizu 2021) that allows for the presence of unobserved variables. CAM-UV provides structural information on causal graphs related to unobserved variables (unob-

served causal path and unobserved backdoor path). We show that the ground truth causal graph has structural consistency with the provided information of CAM-UV on each dataset. Because multiple different DAGs satisfy such characterization, we propose an approach to enumerate possible DAGs and name the approach Integrating CAM-UV (I-CAM-UV). Moreover, we show an efficient algorithm for I-CAM-UV and confirm its usefulness via experiments. Our contributions are summarized as follows:

- We propose an enumeration approach named I-CAM-UV to obtain integrated causal graphs of CAM-UV over multiple datasets with non-identical variable sets.
- We show that I-CAM-UV recovers the ground truth DAG when CAM-UV results have no enormous errors and the integrated variable set has no further confounders.
- We propose an efficient I-CAM-UV algorithm with a best-first strategy based on a monotonic inconsistency of DAGs and demonstrate its usefulness via experiments.

Related Work

While there are a variety of similar but slightly different datasets in the real world, almost all studies on causal discovery are specialized for a single dataset, and only a few studies tackle multiple datasets with non-identical variable sets (Tillman, Danks, and Glymour 2009; Triantafillou, Tsamardinos, and Tollis 2010; Tillman and Spirtes 2011; Huang et al. 2020). Our study challenges that field and aims to allow for richer and more practical causal analysis.

Causal discovery methods can be categorized into three approaches: *functional model-based* (Shimizu et al. 2006; Hoyer et al. 2009), *constraint-based* (Spirtes and Glymour 1991; Spirtes, Meek, and Richardson 1995), and *score-based* (Chickering 2002) methods. They are not mutually exclusive and can be combined. In this study, we focus on a non-linear functional model-based approach called the Causal Additive Model with Unobserved Variables (CAM-UV) (Maeda and Shimizu 2021). Causal Additive Model (CAM) (Bühlmann, Peters, and Ernest 2014) assumes non-linear causal functions where effects from each variable can be separated linearly, which is a special case of Additive Noise Model (ANM) (Hoyer et al. 2009). CAM-UV identifies causal relationships as much as possible and detects latent causal/backdoor paths containing unobserved variables. Leveraging the such ability of CAM-UV, we take a constraint-based approach to integrate CAM-UV results by a combinatorial search method.

CD-MiNi (Huang et al. 2020) takes a functional model-based approach over multiple datasets with non-identical variable sets, similar to our method. It assumes a linear non-Gaussian acyclic model (LiNGAM) (Shimizu et al. 2006) and utilizes a continuous optimization algorithm based on the overcomplete independent component analysis to learn a causal graph (Ding et al. 2019). While CD-MiNi handles linear causal relationships derived from LiNGAM, we handle non-linear causal relationships via CAM.

ION (Tillman, Danks, and Glymour 2009), IOD (Tillman and Spirtes 2011), and COmbINE (Triantafillou, Tsamardinos, and Tollis 2010) focus on constraint-based approaches




Figure 1: Illustration of unobserved causal path (UCP) and unobserved backdoor path (UBP).

such as PC and FCI over multiple datasets with non-identical variable sets. They consider partial ancestral graphs (PAGs) representing Markov equivalence classes (MECs) of causal graphs and enumerate all possible PAGs having structural consistency with PC or FCI results. On the other hand, we consider enumerating not PAGs but full DAGs. That is, while the existing methods output a PAG set implicitly representing DAGs of MECs, our method outputs a single explicit DAG set.

Preliminaries

Let $V = \{v_i\}_{i=1}^d$ be a variable set, and we consider causal relationships over V as a graph structure. Let $G^* = (V, A^*)$ be the ground truth causal graph, which is a DAG with directed edges $A^* \subset V^2$. Each $(v_i, v_j) \in A^*$ means v_i is a direct cause (parent) of v_j and let $P_j^* := \{v_i \mid (v_i, v_j) \in A^*\}$ be the set of parents of v_j .

We assume that the causal relationships are formulated by CAM: $v_i = \sum_{v_j \in P_i^*} f_j^{(i)}(v_j) + n_i$ (Bühlmann, Peters, and Ernest 2014) where each $f_j^{(i)}$ is a non-linear function, n_i is the noise term on v_i , and n_1, \dots, n_d are independent of each other. We have an algorithm to estimate G^* on CAM from a given dataset of V . However, in practice, unobserved variables $U \subset V$ may lead to incomplete estimation.

Let $P_j := P_j^* \setminus U$ and $Q_j := P_j^* \cap U$. CAM-UV is formulated as $v_i = \sum_{v_j \in P_i} f_j^{(i)}(v_j) + \sum_{v_k \in Q_i} f_k^{(i)}(v_k) + n_i$ (Maeda and Shimizu 2021). Given a dataset of $V \setminus U$, the estimation algorithm of CAM-UV aims to obtain a mixed graph $G = (V \setminus U, A, N)$ where N is an undirected edge set such that, if and only if $\{v_i, v_j\} \in N$ holds, the causal relationship between v_i and v_j is not identified due to the existence of an *unobserved causal path* (UCP) or *unobserved backdoor path* (UBP) between v_i and v_j on G^* .

Definition 1 (Unobserved Causal Path (UCP)). *A path from v_i to v_j is called a unobserved causal path iff its form is $v_i \rightarrow \dots \rightarrow v_k \rightarrow v_j$ where $v_k \in U$.*

Definition 2 (Unobserved Backdoor Path (UBP)). *A path between v_i and v_j is called a unobserved backdoor path iff its form is $v_i \leftarrow v_x \leftarrow \dots \leftarrow v_a \rightarrow \dots \rightarrow v_y \rightarrow v_j$ where $v_x, v_y \in U$. It is allowed to be $v_a = v_x$ and $v_a = v_y$.*

Figure 1 shows an illustration of UCP and UB. Figure 2 shows an example of CAM-UV result.




Figure 2: Example of CAM-UV. A UCP $v_4 \rightarrow v_5 \rightarrow v_6$ and a UBP $v_2 \leftarrow v_1 \rightarrow v_3$ are found. The dashed lines indicate undirected edges of the resulting mixed graph.




Figure 3: Example of inputs on a scenario we consider (left) and their simple overlapping (right).

Integrating CAM-UV

We handle a scenario to estimate a causal graph from multiple datasets with non-identical variable sets via CAM-UV. Given m datasets of different variable sets $V_1, \dots, V_m \subset V$ where $k \neq l \Leftrightarrow V_k \neq V_l$, we assume each variable set has a common variable with at least one other variable set. Since, if some datasets share common sample sets, it is better to combine them into a single dataset, we assume the datasets also differ from each other in their sample sets. Moreover, we assume all the datasets follow the same causal graph, but functional forms and data distributions may be different. Regarding variables not in V_k as unobserved for k -th dataset, CAM-UV provides a mixed graph $G_k = (V_k, A_k, N_k)$. Our goal is to estimate a causal graph over $\hat{V} = \bigcup_{k=1}^m V_k$ from the given CAM-UV results G_1, \dots, G_m .

The identified causal relationships on G_1, \dots, G_m can be naively overlapped into a single graph: Let $\hat{A} := \bigcup_{k=1}^m A_k$ and $\hat{G} := (\hat{V}, \hat{A})$. However, the causal relationships for some variable pairs are not identified yet (see Figure 3 for an example): $E_{\text{imp}} := (\bigcup_{k=1}^m N_k) \setminus \{\{v_i, v_j\} \mid (v_i, v_j) \in \hat{A}\}$ and $E_{\text{uno}} := \{\{v_i, v_j\} \subseteq \hat{V} \mid \forall k, \{v_i, v_j\} \not\subseteq V_k\}$. The causal relationships on E_{imp} are impossible to identify in any dataset due to the existence of a UCP/UBP. Moreover, since any variable pair in E_{uno} is not observed simultaneously in any dataset, it is even more challenging to identify their causal relationships.

Problem Definition for Ideal Situations

To address the above limitations, we propose a novel approach named *Integrating CAM-UV* (I-CAM-UV) that finds causal graphs satisfying consistency with respect to the given CAM-UV results. First, we define the consistency of causal graphs based on UCP/UBP. We then show that the ground truth causal graph satisfies the consistency.

Let $E := E_{\text{imp}} \cup E_{\text{uno}}$. For any $k \in \{1, \dots, m\}$, let $I_k := \{\{v_i, v_j\} \subseteq V_k \mid \{v_i, v_j\} \notin N_k\}$ be the variable pairs with the identified causal relationships on V_k . Let $\text{UP}_{G, V_k}(v_i, v_j)$




Figure 4: Example of I-CAM-UV on the input of Figure 3.




Figure 5: Example of I-CAM-UV resulting a unique DAG.

be the set of all the UCPs and UBPs between v_i and v_j on a graph G over \hat{V} where $U = \hat{V} \setminus V_k$.

Definition 3. Given CAM-UV results G_1, \dots, G_m , let \tilde{A} be a directed edge set obtained by assigning directions to or excluding edges in E . We call $\tilde{G} = (\hat{V}, \tilde{A} \cup \hat{A})$ consistent if and only if, for all $k \in \{1, \dots, m\}$, $\text{UP}_{\tilde{G}, V_k}(v_i, v_j) = \emptyset$ holds for any $\{v_i, v_j\} \in I_k$ and $\text{UP}_{\tilde{G}, V_k}(v_i, v_j) \neq \emptyset$ holds for any $\{v_i, v_j\} \in N_k$.

We here state what *ideal situations* are. Given CAM-UV results G_1, \dots, G_m , we are in an ideal situation if we have $\hat{V} = \bigcup_{k=1}^m V_k = V$, i.e., \hat{V} has no unobserved variables, and G_1, \dots, G_m have no estimation error.

Theorem 1. G^* is consistent under any ideal situation.

Theorem 1 implies that the ground truth causal graph corresponds to a consistent DAG. However, because there may exist multiple consistent DAGs, I-CAM-UV aims to solve the following problem.

Problem 1. Given CAM-UV results G_1, \dots, G_m , the problem is to enumerate all the consistent DAGs.

In other words, given a set of mixed graphs G_1, \dots, G_m , it implicitly represents a set of consistent DAGs \mathcal{G} , and we aim to explicitly enumerate them $\tilde{G}_1, \dots, \tilde{G}_{|\mathcal{G}|} \in \mathcal{G}$. From this perspective, I-CAM-UV might appear only to extract the implicitly represented DAGs. However, a crucial aspect of I-CAM-UV lies in its ability to orient unobserved variable pairs, E_{uno} , which allows for the identification of the causal relationships between variables where the direct estimation via observed datasets is difficult.

Example 1. When the input is the form shown in Figure 3, I-CAM-UV works as shown in Figure 4. Focusing on G_1 , we need $v_1 \rightarrow v_3$ or $v_1 \leftarrow v_3$ to obtain the UCP $v_1 \rightarrow v_3 \rightarrow v_4$ or the UBP $v_1 \leftarrow v_3 \rightarrow v_4$. Focusing on G_2 , we exclude $\{v_1, v_4\}$ to avoid the UBP $v_2 \leftarrow v_1 \rightarrow v_4$. A similar discussion holds when G_1 and G_2 are swapped. As a result, we obtain two consistent DAGs that have a directed edge between the unobserved variable pair $\{v_1, v_3\}$.

Example 2. In Figure 5, focusing on G_2 , while G_1 induces the UCP $v_1 \rightarrow v_3 \rightarrow v_4$, we still need the UCP $v_2 \rightarrow v_3 \rightarrow v_4$ or the UBP $v_2 \leftarrow v_3 \rightarrow v_4$. We fix $v_2 \rightarrow v_3$ to avoid




Figure 6: Examples of I-CAM-UV with relaxation.

the UCP $v_1 \rightarrow v_3 \rightarrow v_2$. We then exclude (v_2, v_4) to avoid the UBP $v_3 \leftarrow v_2 \rightarrow v_4$ with respect to G_1 . As a result, we obtain a unique consistent DAG which has a directed edge between the unobserved variable pair $\{v_2, v_3\}$.

Problem Relaxation for Realistic Situations

While I-CAM-UV works well under any ideal situations, in practical use cases, unobserved variables may still exist ($\hat{V} \subset V$) and CAM-UV may cause estimation errors. These may not yield any consistent graph. In fact, estimating CAM-UV is currently not complete in theory to obtain proper mixed graphs, although the first algorithm (Maeda and Shimizu 2021) is modified in the literature (Pham, Maeda, and Shimizu 2025) to improve accuracy. Therefore, we define a new problem with a relaxed graph consistency requirement by introducing an *inconsistency cost* which counts the number of variable pairs with inconsistent (non-)existence of UCP/UBP on each V_k .

Definition 4. For any graph \tilde{G} , to represent inconsistent variable pairs on \tilde{G} for each k , let $\bar{I}_k(\tilde{G}) := \{\{v_i, v_j\} \in I_k \mid \text{UP}_{\tilde{G}, V_k}(v_i, v_j) \neq \emptyset\}$ and $\bar{N}_k(\tilde{G}) := \{\{v_i, v_j\} \in N_k \mid \text{UP}_{\tilde{G}, V_k}(v_i, v_j) = \emptyset\}$. We define the inconsistency cost of \tilde{G} as $C(\tilde{G}) := \sum_{k=1}^m (|\bar{I}_k(\tilde{G})| + |\bar{N}_k(\tilde{G})|)$.

Problem 2. Given CAM-UV results G_1, \dots, G_m and the user parameter $b \in \mathbb{Z}_{\geq 0}$, the problem is to enumerate all the DAGs with the inconsistency cost less than or equal to $C^* + b$ where C^* is the minimum inconsistency cost of DAGs.

Example 3. In Figure 6-(a), we have the estimation error of $\{v_1, v_2\} \in N_2$ on CAM-UV results because the ground truth has no edge between v_1 and v_2 . This case gives three DAGs with the minimum inconsistency cost of 1. All the DAGs are inconsistent for $\{v_1, v_2\} \in N_2$ because no UCP/UBP between v_1 and v_2 will never occur. However, I-CAM-UV can recover the ground truth DAG as the rightmost DAG. Note that I-CAM-UV can not identify which of the enumerated graphs is the ground truth.

Example 4. In Figure 6-(b), we show an example such that I-CAM-UV can not recover the ground truth DAG even though we solve the relaxed problem. We have the estimation error of $(v_3, v_4) \notin A_1$ where $(v_3, v_4) \in A_1$ is true. Nevertheless, this case has three false DAGs with the inconsistency cost of 0. Moreover, even if we seek DAGs with non-zero inconsistency costs, the ground truth DAG does not appear because $(v_3, v_4) \notin \hat{A}$ and $\{v_3, v_4\} \notin E$.

Similar to Example 4, the ground truth DAG also cannot be recovered when CAM-UV estimates $(v_i, v_j) \in A_k$ where $(v_i, v_j) \notin A^*$. Therefore, considering I-CAM-UV can assign directions to or exclude only edges in E , we need the following condition to recover the ground truth DAG: The input of I-CAM-UV satisfies $\hat{A} \subseteq A^*$, $A^* \setminus \hat{A} \subseteq \{(v_i, v_j) \mid \{v_i, v_j\} \in E\}$, and $C(G^*) \leq C^* + b$.

That is, it is desirable that CAM-UV estimation algorithms make no false discovery and capture a trace of each true causal relationship, even if it forms a falsely estimated UCP/UBP. However, such desirable conditions are not always satisfied. Therefore, in practical settings, we only require that CAM-UV algorithms yield an acyclic overlapped graph \hat{G} . It is easily achieved by sharing estimated directed edges among all the datasets to avoid cycles. Under the above setting, we handle Problem 2 instead of Problem 1.

Algorithm

Now let us describe the proposed algorithm of I-CAM-UV to solve Problem 2. I-CAM-UV essentially requires the consideration of an exponential number of graphs via all the connection patterns on E , i.e., we assign the direction to or exclude each $\{v_i, v_j\} \in E$, and there are $3^{|E|}$ such patterns. Although the algorithm ignores cyclic graphs, there are still a huge number of DAGs to be considered. Therefore, we introduce the key idea to reduce the computational cost by focusing on the monotonicity of the inconsistency cost. That leads to designing a best-first search algorithm that enumerates DAGs in ascending order of the inconsistency cost.

Let the variable pairs $E := \{e_1, e_2, \dots, e_{|E|}\}$ be ordered. We can exhaustively search all the DAGs by sequentially processing the variable pairs. Let us denote a graph generated by adding a directed edge (v_i, v_j) into \tilde{G} as $\tilde{G}_{+(v_i, v_j)}$. We manage each search state as a pair (t, \tilde{G}) where t is the number of processed variable pairs and \tilde{G} is a candidate solution graph. First, we construct the overlapped DAG \hat{G} and push the initial state $(0, \hat{G})$ into the top of the search queue. After that, we iteratively pop a state (t, \tilde{G}) from the search queue to generate a finalized state $(|E|, \tilde{G})$ and the new states $(s+1, \tilde{G}_{+(v_i, v_j)})$ and $(s+1, \tilde{G}_{+(v_j, v_i)})$ for each $e_s = \{v_i, v_j\}$ where $s > t$. We push the generated states into the top of the search queue, ignoring cyclic graphs.

The above procedure seeks all the DAGs, not taking into account the inconsistency cost. On the other hand, we are interested in only the graphs with the inconsistency costs of less than or equal to $C^* + b$. Consequently, we have a chance to reduce the computation time by seeking limited DAGs. Intuitively, a faster computation is achieved by popping states in ascending order of the minimum inconsistency cost of the successor DAGs until a popped state has a cost greater than $C^* + b$. However, it is difficult to compute the minimum inconsistency cost of the successor DAGs. Therefore, we use a lower bound of the inconsistency cost as an alternative.

Definition 5. For any state (t, \tilde{G}) where $\tilde{G} = (\hat{V}, \hat{A} \cup \tilde{A})$, let $A_t := \bigcup_{e_s = \{v_i, v_j\} \in E, s > t} \{(v_i, v_j), (v_j, v_i)\}$ and $\tilde{G}[t] := (\hat{V}, \hat{A} \cup \tilde{A} \cup A_t)$. That is, $\tilde{G}[t]$ has the bi-directed edge for

Algorithm 1: I-CAM-UV

```

1: Compute  $\hat{G}$  and  $E \leftarrow E_{\text{imp}} \cup E_{\text{uno}}$ 
2: Set an edge order  $E = \{e_1, e_2, \dots, e_{|E|}\}$ 
3: Construct the heap  $\mathcal{Q}$  with the priority  $\tilde{C}$ 
4: Initialize  $\mathcal{Q} \leftarrow \{(0, \hat{G})\}$ ,  $\mathcal{G} \leftarrow \emptyset$ , and  $C^* \leftarrow \text{inf}$ 
5: while  $\mathcal{Q}$  is not empty do
6:   Pop the top state  $(t, \tilde{G})$  from  $\mathcal{Q}$ 
7:   if  $\tilde{C}(t, \tilde{G}) > C^* + b$  then break
8:   if  $t = |E|$  then
9:     if  $C^* = \text{inf}$  then  $C^* \leftarrow C(\tilde{G})$ 
10:     $\mathcal{G} \leftarrow \mathcal{G} \cup \{\tilde{G}\}$ 
11:   else
12:     for  $s = t + 1, \dots, |E|$  do
13:        $v_i, v_j \leftarrow e_s$ 
14:        $\tilde{G}_1 \leftarrow \tilde{G}_{+(v_i, v_j)}$  and  $\tilde{G}_2 \leftarrow \tilde{G}_{+(v_j, v_i)}$ 
15:       if  $\tilde{G}_1$  is a DAG then Push  $(s, \tilde{G}_1)$  into  $\mathcal{Q}$ 
16:       if  $\tilde{G}_2$  is a DAG then Push  $(s, \tilde{G}_2)$  into  $\mathcal{Q}$ 
17:     Push  $(|E|, \tilde{G})$  into  $\mathcal{Q}$ 
18: return  $\mathcal{G}$ 

```

each e_s where $s > t$. Note that UCP and UBP are well-defined even on cyclic graphs. We define the cost function $\tilde{C}(t, \tilde{G}) := \sum_{k=1}^m (|\bar{I}_k(\tilde{G})| + |\bar{N}_k(\tilde{G}[t])|)$.

Lemma 1. $C(\tilde{G}) = \tilde{C}(|E|, \tilde{G})$ holds for any state $(|E|, \tilde{G})$.

Theorem 2. For any state (t, \tilde{G}) and any its successor state (s, \tilde{G}') , $\tilde{C}(t, \tilde{G}) \leq \tilde{C}(s, \tilde{G}')$ holds.

According to Lemma 1, the cost function \tilde{C} can naturally evaluate the inconsistency cost of finalized DAGs. Moreover, according to Theorem 2, we have the monotonicity of the cost function \tilde{C} . We then propose to use the heap that prioritizes the states by ascending order of $\tilde{C}(t, \tilde{G})$. As a result, the algorithm forms a best-first search, enumerates the DAGs in ascending order of the inconsistency cost, and is expected to reduce the computation time by seeking heuristically limited DAGs. However, we have the remaining challenge to compute $\tilde{C}(t, \tilde{G})$. It is equal to computing $\bar{I}_k(\tilde{G})$ and $\bar{N}_k(\tilde{G}[t])$ for all k by searching UCPs/UBPs. Fortunately, we can construct polynomial-time algorithms to search a UCP/UBP on a given graph as follows.

Search a UCP Given a graph $G = (\hat{V}, A)$ and observed variables $W \subset \hat{V}$, we show an algorithm to search a UCP of $\{v_i, v_j\} \subseteq W$ on G . Let $T_i := \{v_k \in \hat{V} \setminus W \mid (v_k, v_i) \in A\}$ and $T_j := \{v_k \in \hat{V} \setminus W \mid (v_k, v_j) \in A\}$ be the unobserved variables of one hop before v_i and v_j respectively. A path from v_i (resp. v_j) to $v_k \in T_j$ not passing through v_j (resp. v_i) makes a UCP $v_i \rightarrow \dots \rightarrow v_k \rightarrow v_j$ (resp. $v_j \rightarrow \dots \rightarrow v_k \rightarrow v_i$). It is found in $O(|A|)$ time by a breadth-first search.

Search a UBP Given a graph $G = (\hat{V}, A)$ and observed variables $W \subset \hat{V}$, we show an algorithm to search a UBP of $\{v_i, v_j\} \subseteq W$ on G . We define T_i and T_j as in the case

of UCP. Moreover, let $\mathbf{1}_{v_i v_j}^G[u, w]$ indicates whether G has a path from u to w not passing through v_i and v_j . Let $S_i := \{v_a \in \hat{V} \mid \exists v_x \in T_i, \mathbf{1}_{v_i v_j}^G[v_a, v_x]\}$ and $S_j := \{v_a \in \hat{V} \mid \exists v_y \in T_j, \mathbf{1}_{v_i v_j}^G[v_a, v_y]\}$. We have a UBP between v_i and v_j if and only if $S_i \cap S_j \neq \emptyset$, i.e., there is a path $v_i \leftarrow v_x \leftarrow \dots \leftarrow v_a \rightarrow \dots \rightarrow v_y \rightarrow v_j$ where $v_a \in S_i \cap S_j$ and $v_x, v_y \in \hat{V} \setminus W$. We can compute S_i and S_j in $O(|A|)$ time by a breadth-first search.

Pseudocode Algorithm 1 summarizes the best-first search procedure of I-CAM-UV. The algorithm initializes the heap with the state consisting of the overlapped DAG (line 4), iteratively pops the top state (line 6), and checks the stop condition (line 7). If the algorithm continues, it adds the DAG of the current state into the solution set if $t = |E|$ (line 10), otherwise pushes the next states into the heap, computing their priority (line 12-17). Note that the algorithm computes C^* when the first solution DAG is obtained (line 9).

Computational Complexity The proposed I-CAM-UV algorithm has computational complexity depending on the number of searched states. The algorithm searches states in ascending order of \tilde{C} , stops when the popped state satisfies $\tilde{C}(t, \tilde{G}) > C^* + b$, and generates up to $2(|E| - t) + 1$ successors for each state. Thus, let Q be the number of the possible states with $\tilde{C}(t, \tilde{G}) \leq C^* + b$, and we have $O(Q|E|)$ searched states. On computing \tilde{C} , because we need to search a UCP/UBP for each variable pairs on all the datasets, the algorithm takes $O((|\hat{A}| + |E|) \sum_{k=1}^m |V_k|^2)$ time for each state where $|\hat{A}| + |E|$ is the maximum possible graph size. Therefore, let $R = (|\hat{A}| + |E|) \sum_{k=1}^m |V_k|^2$ for short, and the proposed best-first search algorithm via a heap takes $O(Q|E|(R + \log(Q|E|)))$ time.

Experiments

We conducted experiments to demonstrate the usefulness of I-CAM-UV. We posed the following four questions and answered each of them through experiments:

- Q1. How accurately does I-CAM-UV recover CAM-UV misses on observed variable pairs?
- Q2. How accurately does I-CAM-UV discover causal relationships on unobserved variable pairs E_{uno} .
- Q3. How many DAGs does I-CAM-UV enumerate, and how does the accuracy distribute on the DAGs?
- Q4. How realistically can I-CAM-UV operate in terms of computation times?

We implemented the core engine for the search procedure of I-CAM-UV in C++ and wrapped it with Cython. The finalized module of I-CAM-UV and other experimental modules were implemented by Python. The computer environment consisted of Ubuntu 24.04.1 LTS, Intel(R) Xeon(R) E-2174G CPU, and 64GB RAM.

Setup

Datasets We generated 100 synthetic datasets following CAM based on random graphs of Erdős–Rényi model (Er-

dos and Renyi 1960) with ten variables, the edge probability parameter 0.3, and some non-linear functions. The same setting is used in the existing CAM-UV literatures (Maeda and Shimizu 2021; Pham, Maeda, and Shimizu 2025). From each dataset, we constructed an instance by randomly resampling datasets with non-identical variable sets where the number of datasets was $m \in \{2, 3\}$ and the number of unobserved variables per dataset was $|U| \in \{3, 4\}$. In each instance, at least $|U|$ variable pairs with causal relationships were simultaneously observed (resp. unobserved). Moreover, each dataset was generated to have at least one common observed variable with one of the other datasets. The number of observations sampled per dataset was 1,000.

Competitors We set $b = 0$ for I-CAM-UV, i.e., we enumerated the DAGs with the minimum inconsistency cost. In addition, we compared I-CAM-UV with the following methods: (*CAM-UV-OVL*) We used the simple overlapped DAG of CAM-UV results as a baseline. (*PC-OVL*) We applied the PC algorithm to each dataset and overlapped the estimated directed edges. We took care to make the overlapped graph acyclic by ignoring later added edges if they make a cycle. (*Imputation*) Because unobserved variables can be treated as missing values, we combined the datasets into one by k -Nearest Neighbor imputation of $k = 5$. After the imputation, we applied CAM to the combined dataset with randomly resampling 1,000 observations to reduce the computational cost. (*CD-MiNi*) Although we focused on CAM, we directly applied CD-MiNi that estimates a LiNGAM on multiple datasets with non-identical variable sets.

Metrics Since I-CAM-UV outputs multiple DAGs, basic metrics cannot be used directly. Therefore, focusing on the frequency of edge existence on output DAGs, we modified true positive, false positive, and false negative, and called them MTP, MFP, and MFN. For a graph set \mathcal{G} and two different variables v_i and v_j , let c_{ij} be the number of graphs in \mathcal{G} having the directed edge (v_i, v_j) . With respect to the ground truth causal graph of the directed edge set A^* , we define $MTP = \sum_{(v_i, v_j) \in A^*} \frac{c_{ij}}{|\mathcal{G}|}$, $MFP = \sum_{(v_i, v_j) \notin A^*} \frac{c_{ij}}{|\mathcal{G}|}$, and $MFN = \sum_{(v_i, v_j) \in A^*} \frac{|\mathcal{G}| - c_{ij}}{|\mathcal{G}|}$. We then used the modified variations of recall $\frac{MTP}{MTP+MFN}$, precision $\frac{MTP}{MTP+MFP}$, and their harmonic mean (F1 score). All the modified metrics are equal to the original ones when the graph set is a singleton. Therefore, the modified metrics can be applied to all the competitors without any issues.

Results

Regarding Q1 and Q2 To answer Q1 and Q2, we individually evaluated the recall scores on observed and unobserved variable pairs. In addition, we also evaluated the overall precision and F1 scores. Figure 7 summarizes the results.

We found that the recall scores of I-CAM-UV were superior to all the competitors on both observed and unobserved variable pairs. On the other hand, I-CAM-UV decreased the precision scores from that of CAM-UV-OVL. Nonetheless, the overall F1 scores showed almost no difference between CAM-UV-OVL and I-CAM-UV. While the results support the ability of I-CAM-UV to discover causal relationships,

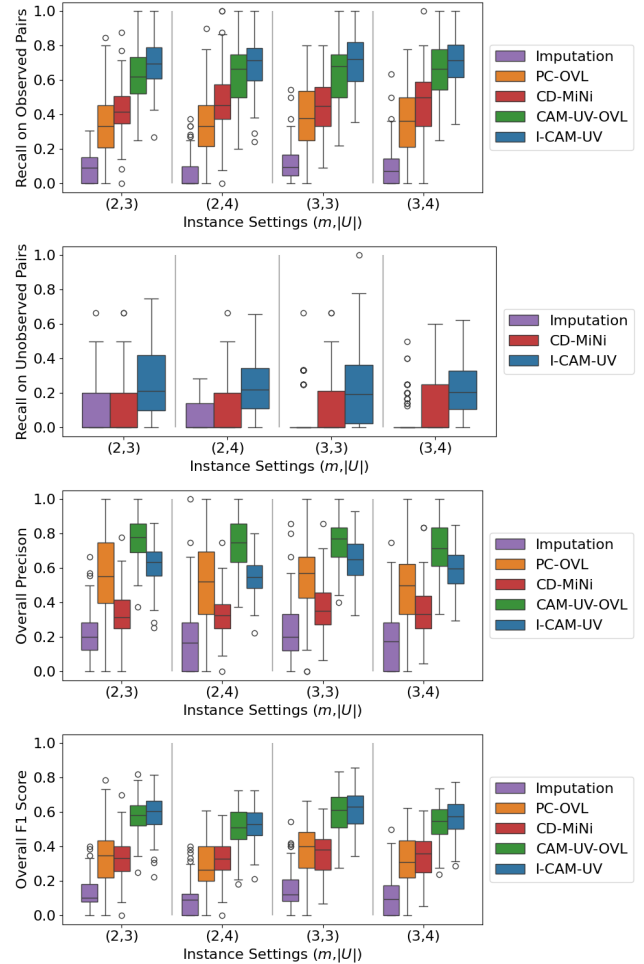


Figure 7: Accuracies on synthetic datasets.

which is challenging to identify by only analyzing observed samples, the results also indicate that I-CAM-UV outputs a similar number of false discoveries.

Consequently, conclusions regarding Q1 and Q2 are as follows: (A1 and A2) *I-CAM-UV has superior recall compared with the other methods to recover causal relationships missed by CAM-UV and discover causal relationships on unobserved variable pairs. However, we need to take into account that the output DAGs may contain a similar number of false discoveries.*

Regarding Q3 To answer Q3, we investigated the number of DAGs enumerated by I-CAM-UV and their individual accuracies. The left one of Figure 8 shows statistics on the number of DAGs enumerated by I-CAM-UV. We selected some instances with a moderate number of output DAGs and plotted their accuracy distributions as shown in Figure 9.

We found that the number of DAGs enumerated by I-CAM-UV could be both little and huge. Since the behavior of I-CAM-UV depends on CAM-UV results, it is difficult to predict the output size. Fortunately, the accuracy distribution seems to form a single cluster, that is, the enumerated DAGs




Figure 8: Number of DAGs enumerated by I-CAM-UV.

are mostly biased toward similar accuracies. These cluster-like accuracy distributions have been observed in many instances regardless the number of DAGs. In addition, as a similar result of the consequences leading to A1 and A2, the accuracy of each DAG enumerated by I-CAM-UV is often superior to the competitors in terms of the recall. Therefore, in practical applications, even if a huge number of DAGs are obtained by I-CAM-UV, a random sampling will narrow them down to a small alternative DAG set.

In conclusion regarding Q3, (A3) *while a wide range of numbers of DAGs are output by I-CAM-UV, almost all DAGs achieve similar accuracy. A random sampling approach may be effective for practical analysis even if a huge number of DAGs are obtained.*

Regarding Q4 To answer Q4, we report computation times of all the competitors on each setting in Figure 10. Note that the process of I-CAM-UV is only the enumeration of consistent graphs, and its former process obtaining CAM-UV results corresponds to CAM-UV-OVL. Thus, the plots of I-CAM-UV in Figure 10 show the computation times of only the enumeration process.

We found that the enumeration process of I-CAM-UV was fast enough compared with CAM-UV-OVL in many instances, and the total computation times were comparable with the other methods. That is, we confirmed that the best-first strategy of the I-CAM-UV algorithm is an efficient approach. However, some instances take significantly more computation time as outliers. Since the search space of I-CAM-UV depends on CAM-UV results, it is difficult to predict the computation time of I-CAM-UV.

In conclusion regarding Q4, (A4) *I-CAM-UV runs in a realistic computation time assuming a relatively sparse DAG with ten variables. Since I-CAM-UV is an exponential time algorithm in general, we should be cautious when handling more variables. Nonetheless, because the worst case has $3^{|E|}$ search states, we can easily apply I-CAM-UV to small E even if the number of variables is large.*

Conclusion

We proposed I-CAM-UV, enumerating integrated causal graphs as consistently as possible given CAM-UV results over non-identical variable sets. The algorithm is an efficient combinatorial search leveraging the monotonicity of the inconsistency cost of DAGs. We demonstrated that I-CAM-UV recovered causal relationships missed by CAM-

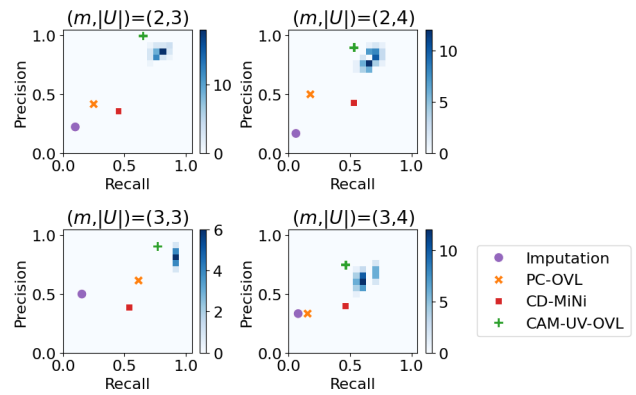


Figure 9: Accuracy distributions on specific instances of synthetic datasets. The heatmaps show the distributions of DAGs output by I-CAM-UV.




Figure 10: Computation times on synthetic datasets.

UV, discovered causal relationships on unobserved variable pairs, and output consistent DAGs of similar accuracies.

Limitations We think that I-CAM-UV has the following three limitations. (a) While tractable properties of UCP and UBP provide I-CAM-UV, extending the algorithm beyond CAM-UV is non-trivial. (b) Because we may obtain a huge number of DAGs by I-CAM-UV, human checks are not easy for all the individual DAGs. (c) The accuracy of I-CAM-UV is highly dependent on the accuracy of CAM-UV results. Based on these considerations, future work is as follows.

Future Work From a theoretical perspective, the following three points are of interest: (i-a) Can we construct an algorithm integrating causal graphs with unobserved variables via other methods than CAM-UV, such as RCD (Maeda and Shimizu 2020)? (i-b) How can we construct a Markov equivalence class-like notion for enumerated DAGs of I-CAM-UV? (i-b') What are the conditions under which a unique DAG is obtained by I-CAM-UV? On the other hand, from a practical perspective, the following three points are essential: (ii-b) Constructing a compressed and interpretable representation of enumerated DAGs. (ii-b') Developing an algorithm to evaluate the fitness of a given DAG over multiple datasets with non-identical variable sets. (ii-c) Making I-CAM-UV more practical by improving the accuracy of CAM-UV itself.

References

- Bühlmann, P.; Peters, J.; and Ernest, J. 2014. CAM: Causal additive models, high-dimensional order search and penalized regression. *Annals of Statistics*, 42(6): 2526–2556.
- Campomanes, P.; Neri, M.; Horta, B. A.; Röhrig, U. F.; Vanni, S.; Tavernelli, I.; and Rothlisberger, U. 2014. Origin of the spectral shifts among the early intermediates of the rhodopsin photocycle. *Journal of the American Chemical Society*, 136(10): 3842–3851.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov): 507–554.
- Ding, C.; Gong, M.; Zhang, K.; and Tao, D. 2019. Likelihood-free overcomplete ICA and applications in causal discovery. *Advances in Neural Information Processing Systems*, 32.
- Erδος, P.; and Renyi, A. 1960. On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 5: 17–61.
- Fu, X.; Mo, S.; Buendia, A.; Laurent, A. P.; Shao, A.; Alvarez-Torres, M. d. M.; Yu, T.; Tan, J.; Su, J.; Sagatelian, R.; Ferrando, A. A.; Ciccina, A.; Lan, Y.; Owens, D. M.; Palomero, T.; Xing, E. P.; and Rabadan, R. 2025a. A foundation model of transcription across human cell types. *Nature*, 637: 965–973.
- Fu, Z.; Yang, X.; Ma, Y.; Sun, Y.; and Wang, T. 2025b. Integrating explainable AI and causal inference to unveil regional air quality drivers in China. *Journal of Environmental Management*, 390: 126270.
- Hoyer, P. O.; Janzing, D.; Mooij, J.; Peters, J.; and Schölkopf, B. 2009. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems 21*, 689–696. Curran Associates Inc.
- Huang, B.; Zhang, K.; Gong, M.; and Glymour, C. 2020. Causal discovery from multiple data sets with non-identical variable sets. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 10153–10161.
- Maeda, T. N.; and Shimizu, S. 2020. RCD: Repetitive causal discovery of linear non-Gaussian acyclic models with latent confounders. In *Proc. 23rd International Conference on Artificial Intelligence and Statistics (AISTATS2010)*, volume 108 of *Proceedings of Machine Learning Research*, 735–745. PMLR.
- Maeda, T. N.; and Shimizu, S. 2021. Causal additive models with unobserved variables. In *Proc. 37th Conference on Uncertainty in Artificial Intelligence (UAI2021)*, 97–106. PMLR.
- Mooij, J. M.; Magliacane, S.; and Claassen, T. 2020. Joint Causal Inference from Multiple Contexts. *Journal of Machine Learning Research*, 21(99): 1–108.
- Pham, T.; Maeda, T. N.; and Shimizu, S. 2025. Causal Additive Models with Unobserved Causal Paths and Backdoor Paths. arXiv:2502.07646.
- Runge, J.; Gerhardus, A.; Varando, G.; Eyring, V.; and Camps-Valls, G. 2023. Causal inference for time series. *Nature Reviews Earth & Environment*, 4(7): 487–505.
- Shimizu, S.; Hoyer, P. O.; Hyvärinen, A.; and Kerminen, A. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct): 2003–2030.
- Smith, A.; Ghosh, D.; Tan, A.; Cheng, X.; and Daoutidis, P. 2025. Multi-scale causality in active matter. *Computers & Chemical Engineering*, 197: 109052.
- Spirtes, P.; and Glymour, C. 1991. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9: 67–72.
- Spirtes, P.; Glymour, C. N.; and Scheines, R. 2000. *Causation, prediction, and search*. MIT press.
- Spirtes, P.; Meek, C.; and Richardson, T. 1995. Causal Inference in the presence of latent variables and selection bias. In *Proc. 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI1995)*, 491–506.
- Tillman, R. E.; Danks, D.; and Glymour, C. 2009. Integrating locally learned causal structures with overlapping variables. In *Advances in Neural Information Processing Systems*, 1665–1672. Curran Associates Inc.
- Tillman, R. E.; and Eberhardt, F. 2014. Learning causal structure from multiple datasets with similar variable sets. *Behaviormetrika*, 41(1): 41–64.
- Tillman, R. E.; and Spirtes, P. 2011. Learning equivalence classes of acyclic models with latent and selection variables from multiple datasets with overlapping variables. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 3–15.
- Triantafillou, S.; Tsamardinos, I.; and Tollis, I. 2010. Learning causal structure from overlapping variable sets. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 860–867.
- Zhou, J.; Zhao, D.; Hao, Q.; Tu, S.; and Xu, L. 2025. Multi-Objective Structure-Based Drug Design Using Causal Discovery. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, PP(99): 1–12.