Practical Algorithms for Orientations of Partially Directed Graphical Models

Author names withheld

Editor: Under Review for CLeaR 2023

Abstract

In observational studies, the true causal model is typically unknown and needs to be estimated from available observational and limited experimental data. In such cases, the learned causal model is commonly represented as a partially directed acyclic graph (PDAG), which contains both directed and undirected edges indicating uncertainty of causal relations between random variables. The main focus of this paper is on the maximal orientation task, which, for a given PDAG, aims to orient the undirected edges maximally such that the resulting graph represents the same Markov equivalent DAGs as the input PDAG. This task is a subroutine used frequently in causal discovery, e. g., as the final step of the celebrated PC algorithm. Utilizing connections to the problem of finding a consistent DAG extension of a PDAG, we derive faster algorithms for computing the maximal orientation by proposing two novel approaches for extending PDAGs, both constructed with an emphasis on simplicity and practical effectiveness.

Keywords: Causal graphical models, Directed acyclic graphs, Markov equivalence, Consistent extension, Maximal orientation, Meek rules

1. Introduction

The development of probabilistic graphical models enables a mathematically sound language to handle uncertainty in a coherent and compact way (Spirtes et al., 2000; Pearl, 2009; Koller and Friedman, 2009; Elwert, 2013). They also provide scientists with intuitive tools for causal analysis and currently receive substantial attention in epidemiology, sociology and other disciplines. Moreover, the graphical modeling approach allows for the use of computational methods that have enabled significant progress towards automated causal inference and causal structure discovery.

Certainly, one of the most prominent graphical models is the *directed acyclic graph* (DAG), whose edges encode direct causal influences between the random variables of interest. In practice, however, the underlying true DAGs are unknown and from observational or limited experimental data they can only be inferred to a certain degree of uncertainty. In such cases, the learned causal model is given, typically, as a *partially directed acyclic graph* (PDAG) which contains both directed and undirected edges. Such a PDAG represents a class of *Markov equivalent* DAGs encoding the same statistical properties and its undirected edges indicate which directed edges may vary across DAGs of the class (Verma and Pearl, 1990; Meek, 1995; Andersson et al., 1997).

In this paper, we investigate *orientations* of PDAG models – primitive tasks used to solve more complex problems of causal analysis. Our goal is to provide simple and effective algorithms that improve performance of downstream tasks, e.g., in causal structure learning, active learning or causal effect estimations. The main focus of our study is on the *maximal orientation* problem, the goal of which is to orient a maximal number of undirected edges such that the resulting graph, called *maximally oriented PDAG* (MPDAG), represents the same DAGs as G. It is well known that

© 2023 .

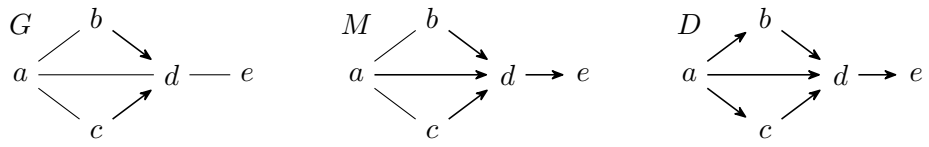

Figure 1: A PDAG G with a maximal orientation M. DAG D shows a consistent extension of G.

the MPDAG can be obtained from the PDAG by closing it under the celebrated Meek rules (Meek, 1995). For an example PDAG G and its MPDAG M, see Fig. 1.

The primary significance of the MPDAG model, including *completed PDAGs* (CPDAGs, also called *essential graphs*) (Andersson et al., 1997) is that it represents Markov equivalent DAGs in an elegant and convenient way. It is commonly used for solving many important problems as, e.g., counting and sampling Markov equivalent DAGs (He et al., 2015; Wienöbst et al., 2021), estimating causal effects (Maathuis et al., 2009; van der Zander and Liśkiewicz, 2016; Perković et al., 2017), or learning causal models (Chickering, 2002), where CPDAGs represent the states of the search space. Consequently, orienting a given PDAG maximally, is a frequently used primitive in causal inference and discovery, perhaps most prominently arising in constraint-based causal structure learning, as, e. g., the final step in the PC algorithm (Spirtes et al., 2000; Kalisch and Bühlman, 2007) and its modifications. Similarly, several score-based algorithms, e.g., (Pellet and Elisseeff, 2008) based on a generic feature-selection approach, rely on this task as well. Another important example are algorithms for active learning which, beyond observational, use experimental (interventional) data to resolve orientation ambiguities. In searching for an optimal strategy, typical algorithms iteratively construct *interventional essential graphs*, which are obtained by orienting subsequent PDAGs maximally (He and Geng, 2008; Hauser and Bühlmann, 2012; Shanmugam et al., 2015; Squires et al., 2020). Furthermore, methods enumerating all possible total effects or estimating bounds on the effects in a Markov equivalence class need a subroutine to compute the maximal orientations (Maathuis et al., 2009; Guo and Perkovic, 2021).

While the algorithms for computing maximal orientations used in practice resort to repeatedly applying the Meek rules, the best theoretical methods are based on two other important primitives of causal analysis: (i) the *consistent extension* of a PDAG to a DAG and (ii) computing the CPDAG representation of the graphs Markov equivalent to a given DAG. This was already mentioned by Chickering (1995) and generalized to instances with background knowledge by Wienöbst et al. (2021). Using the clever algorithm of Chickering (1995), the second task (ii) can be solved in linear time. On the other hand, the computational complexity of the first problem, which is to orient all undirected edges such that no new v-structures arises, is significantly higher. Hence, developing effective practical methods for finding consistent extensions of PDAGs can be a key building block in the efficient computation of maximal orientations.

Previous Work. The study on extendability of PDAGs has been initiated by Verma and Pearl (1992), who provided a method to find a consistent DAG extension in time $O(n^4m)$, where n denotes the number of nodes and m the number of edges. At the same time, Dor and Tarsi (1992) proposed a faster method of time complexity $O(n^4)$ which is conceptually simple, easy to implement, and widely used in practice so far. In 2021, Wienöbst et al. proposed a new algorithm for the problem which runs in time $O(n^3)$. Simultaneously, the authors showed that, under a computational intractability assumption, the cubic algorithm is optimal.

Despite its importance, relatively little attention has been paid to the algorithmic aspects of the maximal orientation problem so far. A commonly used approach relies on the direct, iterative application of the four Meek rules until none can be applied anymore. This approach is, however, computationally expensive, and its worst-case run time is $O(n^4 m)$, where n and m denote the number of vertices and edges of the graph. In 1995, Chickering proposed a sophisticated algorithm, avoiding the direct application of the Meek rules, which firstly extends a PDAG to a DAG and next computes the corresponding CPDAG directly from the DAG. Since a DAG can be transformed to the corresponding CPDAG in linear time, the total running time of the algorithm is dominated by the time needed for extension. More recently, Wienöbst et al. (2021) generalized this approach, while maintaining the same time complexity, to maximally orient a PDAG to an MPDAG.

Our contributions. The contributions of this paper are twofold: on the one hand, we give a practical evaluation of methods for the consistent extension problem, contributing two simple yet effective algorithms for this task, on the other hand, we aim to illustrate the strengths of using consistent DAG extensions in the computation of the maximal orientation of a PDAG. While the latter approach has been proposed theoretically in (Chickering, 1995; Wienöbst et al., 2021), it has, to the best of our knowledge, not found widespread use in practice, as most software packages such as `pcalg` (Kalisch et al., 2012) and `causaldag` (Squires, 2018) rely on algorithms, which apply the Meek rules directly. Utilizing consistent DAG extensions instead, yields a better worst-case complexity and is superior in practice as we demonstrate in our experiments.

2. Preliminaries

We consider partially directed graphs $G = (V_G, E_G, A_G)$, in which the pairs of vertices $x, y \in V$ are connected by undirected edges $u - v$ (given in E_G) or directed edges $u \rightarrow v$ (also called arcs, given in A_G). We restrict ourselves to graphs, where at most one edge exists between any pair $x, y \in V$ and if there is such an edge, we call x and y adjacent, denoted by $x \sim y$. The arcs $u \rightarrow v$ and $u \leftarrow v$ have different *orientations* and *orienting* an undirected edge $u - v$ means replacing it with an arc. Vertex x is called a parent of y if $x \rightarrow y \in G$, a child of y if $x \leftarrow y \in G$ and a sibling of y if $x - y \in G$. The sets of parents, children and siblings of vertex v are denoted by $Pa(v)$, $Ch(v)$, $Si(v)$. By $Ne(v) = Pa(v) \cup Ch(v) \cup Si(v)$ we denote the set of neighbors of v and the degree of v counts the number of its neighbors $|Ne(v)|$. Edges, which have v as endpoint, are called *incident* to v. For a set $S \subseteq V$, the induced subgraph $G[S]$ contains all edges from G with both endpoints in S.

A partially directed graph is called acyclic (or partially directed acyclic graph, PDAG for short) if it does not contain a directed cycle, that is, a sequence of distinct vertices $(c_1, c_2, \ldots, c_k)$, with $k \geq 3$, and edges $c_i \rightarrow c_{i+1}$ for $i \in \{1, \ldots, k-1\}$, and $c_k \rightarrow c_1$. A PDAG without any undirected edges ($E = \emptyset$) is called a directed acyclic graph (DAG). There is a linear ordering (called *topological ordering*) of the vertices of every DAG such that $u \rightarrow v$ if u comes before v in the ordering.

Definition 1 *A DAG D is a* consistent extension *of a PDAG G if*

1. *D and G have the same vertex set and $Ne(v)_G = Ne(v)_D$ for all vertices v,*
2. *every directed edge in G is also in D, i. e., $A_G \subseteq A_D$,*
3. *for all edges $u - v$ in E_G, there is $u \rightarrow v$ in A_D or $v \leftarrow u$ in A_D, and*
4. *for all $u, v, w \in V$, the induced subgraph $u \rightarrow v \leftarrow w$ is in G iff it is in D.*

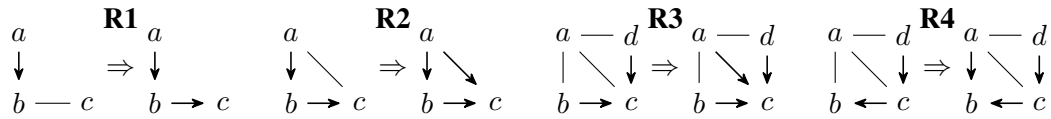

Figure 2: The four Meek rules that are used to characterize MPDAGs (Meek, 1995).

The set of consistent extensions of a PDAG G is denoted by $[G]$. We remark that an induced subgraph of the form $u \rightarrow v \leftarrow w$ is also called a v-structure. Hence, every $D \in [G]$ has the same v-structures as G. If $[G] \neq \emptyset$, we call G *extendable*. Some undirected edges in a PDAG G might be oriented the same way in all of its consistent extensions. The graph where these undirected edges are replaced by the corresponding invariant arcs is called the *maximal orientation* of G, denoted by $\mathrm{MPDAG}(G)$, which stands for *maximally-oriented PDAG*. It is a well-known fact that $\mathrm{MPDAG}(G)$ can be computed by repeatedly applying the four Meek rules (Meek, 1995), which are shown in Fig. 2, to G until none applies anymore. In this work, we consider the following two closely connected computational problems:

Problem 1 EXT

Input: *PDAG* G.
Output: $D \in [G]$ *or* $\perp$ *if* $[G] = \emptyset$.

Problem 2 MAX-ORIENT

Input: *PDAG* G.
Output: $\mathrm{MPDAG}(G)$.

For MAX-ORIENT, we assume that G is extendable, else G has no causal interpretation and the definition of $\mathrm{MPDAG}(G)$ is without meaning. For EXT, the notion of a *potential-sink* is of central importance. It asserts that all incident edges can be oriented towards a vertex v without introducing a new v-structure or changing the orientation of an arc. Formally, vertex v in G is called a potential-sink if (i) there is no arc $v \rightarrow x$ directed outward from v (i. e., $Ch(v) = \emptyset$) and (ii) for every sibling $y \in Si(v)$, vertex y is adjacent to all other neighbors of v. For example, vertex e is the only potential-sink in the graph G shown in Fig. 1.

The following fact states that a consistent extension of a graph G can be obtained by iteratively removing potential-sinks from G. This strategy is known as the Dor-Tarsi algorithm.

Fact 1 (Dor and Tarsi, 1992) *Let G_0 be an extendable PDAG and let G_i, $1 \leq i \leq n$, be obtained by removing potential-sink v_i and its incident edges from G_{i-1}. It holds that $(v_n, v_{n-1}, \ldots, v_1)$ is a topological ordering of a consistent extension of G_0.*

In particular, it is true that every G_i has a potential-sink due to the observation that removing a vertex and its incident edges from an extendable graph, again yields an extendable graph (extendability is closed under taking subgraphs).

These notions are exemplified in Fig. 1. The figure shows that the consistent extension D can be obtained by iteratively identifying a potential-sink (at the start, e is the only potential-sink), orienting its incident undirected edges towards it, and continuing on the induced subgraph over the remaining vertices. In this particular case, d would be the potential-sink found afterward, followed by b and c (in arbitrary order), and finally a.

The graph M is the maximal orientation of G. It contains the arc $d \rightarrow e$ as $d \leftarrow e$ cannot occur in any consistent extension of G (it creates a new v-structure that is not in G, see Meek rule R1) and $a \rightarrow d$ follows from Meek rule R3, whereas $a - b$ and $a - c$ are oriented in both directions in different consistent extensions of D and are consequently undirected in M.

3. Two New Simple Algorithms for Extendability

As discussed above, algorithms for extending PDAGs play a key role in the maximal orientation task. While, from a theoretical point of view, previous results suggest that it is likely not possible to further improve the asymptotic run time of extendability algorithms, as Wienöbst et al. (2021) gave a conditional $O(n^3)$ lower bound for combinatorial algorithms, from a practical perspective, the current algorithms are either extremely simple, as Dor-Tarsi's approach, or very sophisticated with considerable practical overhead, as the methods proposed by Wienöbst et al. (2021). In this work, we are searching for a middle ground and give two novel approaches to extend PDAGs, with the main focus lying on simplicity and effectiveness. The first approach is a direct modification of the Dor-Tarsi algorithm, which greatly improves its practical performance. The second one gives a practical $O(n^3)$ algorithm, which is, however, conceptually simpler than the one presented by Wienöbst et al. (2021) and has significantly less overhead.

3.1. Dor Tarsi with Degree-Heuristic

The Dor-Tarsi algorithm has an asymptotic run time of $O(n^4)$ and is computationally simplistic. It iteratively identifies a potential-sink and then removes it and its incident edges from the graph. Hence, the whole algorithm consists of n iterations, each amounting to the task of finding a potential-sink. Naively, a potential-sink can be found in time $O(n^3)$ by going through all vertices and checking for each vertex whether it satisfies the potential-sink property. Implemented this way, each vertex is repeatedly tested for being a potential-sink throughout the algorithm (until it eventually becomes a potential-sink and is subsequently removed from the graph). While this brute-force method appears to be severely worse than more clever approaches, which store and update relevant information such as a list of all potential-sinks in sophisticated data structures, there are two notable advantages of this approach: (i) whenever a potential-sink is found, the loop over all vertices can be exited immediately, and (ii) the check for potential-sinkness can be cancelled once a single missing edge violates the potential-sink property.

Hence, the first observation of this work, discussed in more detail throughout the experiments in Section 4, is that in many cases, a naive implementation of the Dor-Tarsi algorithm performs quite well in practice (in particular in cases where many potential-sinks cause frequent early exits of the for-loop and also for graphs in which the test for potential-sinkness usually fails early), being competitive with more subtle approaches. Building on the surprising effectiveness of the Dor-Tarsi method, a heuristic refinement is given in Algorithm 1. The idea is to go through the vertices in order of increasing degree. Using this heuristic leads to fewer cost in case the loop is exited early as the iteration in line 3 always starts with the "cheapest" vertices, i. e., those with minimal degree (checking for potential-sinkness results in worst-case costs quadratic in the number of neighbors $O(|Ne(v)|^2)$). Fig. 3 exemplifies this advantage.

It demonstrates that checking potential-sinkness for the low-degree vertices first has, on the one hand, lower cost per vertex and, on the other hand, is often more likely to succeed early (for example, vertices with degree one and no outgoing edges are always potential-sinks). We note that some overhead is induced by this approach because the degree of the vertices has to be continuously updated. However, this cost is mostly tolerable as our experiments presented in Section 4 confirm and the heuristic yields a simple and practical improvement over the standard Dor-Tarsi algorithm. Clearly, a heuristic is not always optimal and worst-case examples can be constructed where the algorithm reaches its asymptotic $O(n^4)$ run time.

Algorithm 1: A heuristic implementation of Dor-Tarsi iterating over the vertices in order of increasing degree.

input : A PDAG $G = (V, E)$.
output: $D \in [G]$.

```
D := (V, ∅)
repeat n times
    for v ∈ V in increasing degree in G do
        if v is a potential-sink then
            Remove v and its incident edges from G.
            Add arcs {(u, v) | u ∈ Si(v) ∪ Pa(v)} to D.
            break
        end
    end
    if no potential-sink has been found then
        return ⊥
    end
end
return D
```

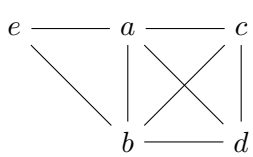

#Adj. tests for the first potential-sink

DT:	$b \overset{?}{\sim} c$, $b \overset{?}{\sim} d$, $b \overset{?}{\sim} e$, $c \overset{?}{\sim} d$, $c \overset{?}{\sim} e$	ps(a) ✗
	$a \overset{?}{\sim} c$, $a \overset{?}{\sim} d$, $a \overset{?}{\sim} e$, $c \overset{?}{\sim} d$, $c \overset{?}{\sim} e$	ps(b) ✗
	$a \overset{?}{\sim} b$, $a \overset{?}{\sim} d$, $b \overset{?}{\sim} d$	ps(c) ✓
DTH:	$a \overset{?}{\sim} b$	ps(e) ✓

Figure 3: A PDAG and the adjacency tests required by Dor-Tarsi (DT) and the heuristic adaption (DTH) for finding the first potential-sink. We assume that DT iterates over the vertices in alphabetical order. DT checks vertices a and b first, which are no potential-sinks. DTH starts with the lowest-degree vertex e. This reduces the cost of testing potential-sinkness and increases the chances of finding a potential-sink early.

Fact 2 *There are instances on which the Dor-Tarsi heuristic yields a run time of* $\Omega(n^4)$.

Proof Consider graph $G = (V, E)$ with $V = C_v \cup C_w \cup C_{vw} \cup \{v\} \cup \{w\}$, where C_{vw} consists of k vertices and C_v and C_w consist of $2k$ vertices each. The edge set is constructed as

$$E = \{\{x, y\} \mid x, y \in C_L \text{ with } L \in \{v, w, vw\}\} \cup (\{v, w\} \times C_{vw}) \cup (\{v\} \times C_v) \cup (\{w\} \times C_w).$$

In words, the cliques C_v, C_w are fully connected to v and w, respectively, and C_{vw} is fully connected to both v and w. Moreover, v and w are non-adjacent. The graph is illustrated in Fig. 4.

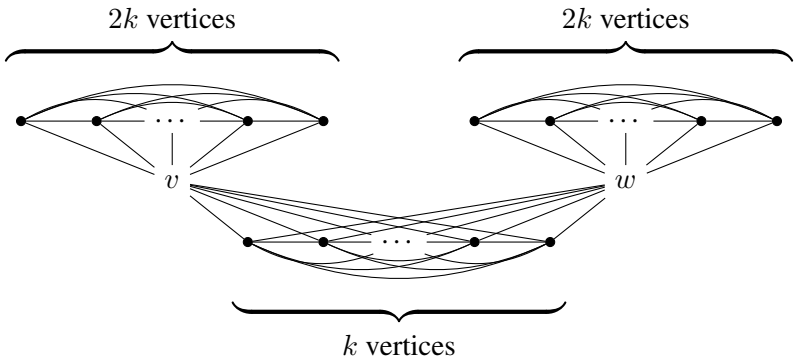

Figure 4: An example instance where the heuristic exhibits run time $O(n^4)$.

Dor-Tarsi in combination with the degree-heuristic always iterates over the vertices in C_{vw} with initial degree $k+1$ first. None of these vertices are, however, potential-sinks. Consequently, $O(k)$ vertices are checked before a potential-sink is found. Moreover, testing for potential-sinkness only fails for one pair of neighbors, that is, for v and w, causing the loop over the pairs of neighbors to require $\Omega(k^2)$ iterations before exiting with high probability. Note that k is in $O(n)$. Overall, the steps undertaken by the heuristic yield a run time of $O(n^4)$. ■

3.2. Dor Tarsi with Improved Worst-Case Complexity

In the previous section, a direct and effective heuristic improvement of the Dor-Tarsi algorithm is presented. In this section, we aim to give another simple adaptation of Dor-Tarsi, but this time with an $O(n^3)$ worst-case run time, thus matching the lower-bounds of the algorithm given in (Wienöbst et al., 2021).

If we take a look at the counterexample from Fig. 4 again, we observe that Algorithm 1 repeatedly iterates over the same neighborhoods when searching for a potential-sink, i. e., there is no information stored between iterations. It is clearly desirable to avoid this repeated computational effort. Previously, improvements were made by constructing elaborate data structures, which maintain the set of potential-sinks throughout the course of the algorithm (Wienöbst et al., 2021). Such data structures, however, induce a significant practical overhead and, in particular, demand an expensive initialization step at the start of the algorithm, even before searching for the first potential-sink.

To address these issues, our proposed approach iterates over the neighbors of every vertex v exactly once, even though potential-sinkness of v might be checked multiple times during the course of the algorithm. After the neighbors of a vertex v have been iterated over, all tuples of neighbors violating the potential-sink property of v are stored in the set $B[v]$, which is afterward continuously updated every time vertices are removed from the graph. The sets $B[v]$ are computed lazily, that is, only at the time the potential-sinkness of v is checked for the first time and thus, there is no pre-computation involved. In case the set of violating neighbors $B[v]$ is empty or becomes empty by removing other vertices, v satisfies the potential-sink property, which is due to the observation that throughout the course of the algorithm no *new* potential-sinkness violations are incurred.

The whole approach is described in detail in Algorithm 2. We combine it with the heuristic from the previous section, which is generally favorable to use. Algorithm 2 combines the advantages of

Algorithm 2: An adaptation of Dor-Tarsi with worst-case complexity $O(n^3)$.

input : A PDAG $G = (V, E)$.
output: $D \in [G]$.

1 $D := (V, \emptyset)$
2 $B :=$ array of n initially empty sets
3 $C :=$ bitvector of length n initialized with $false$
4 **while** *there are vertices left in* G **do**
5 **foreach** $v \in V$ *in increasing degree in* G **do**
6 **if** $C[v] = false$ *and* $Ch(v) = \emptyset$ **then**
7 $B[v] := B[v] \cup \{(u, u') \mid u \in Si(v), u' \in Si(v) \cup Pa(v) \wedge u \neq u' \wedge u \not\sim_G u'\}$
8 $C[v] := true$
9 **end**
10 **if** $C[v] = true$ *and* $B[v] = \emptyset$ **then**
11 Remove v and its incident edges from G.
12 Add arcs $\{(u, v) \mid u \in Si(v) \cup Pa(v)\}$ to D.
13 Remove v from all sets in B containing a tuple where v is included.
14 **break**
15 **end**
16 **end**
17 **if** *no potential-sink has been found* **then**
18 **return** $\perp$
19 **end**
20 **end**
21 **return** D

both DT and WBL by, one the one hand, storing information from previous iterations to ensure an $O(n^3)$ worst-case run time while, on the other hand, keeping the overhead to a minimum and, in particular, not relying on an initialization step.

Theorem 1 *Algorithm 2 is implementable with expected worst-case complexity* $O(n^3)$.

Proof We assume that edge deletion and adjacency tests are supported in constant time and iterating over the neighbors of v has cost $O(N(v))$. Using hash tables to store the neighbors of a vertex, these run times are reached *in expectation*. The graph representation is discussed in more detail in the appendix (Section A).

The computation of neighbors violating the potential-sink property in line 7 can clearly be implemented in $O(n^2)$ by looping over all pairs of neighbors of v and checking adjacency. Due to the flag $C[v]$, the neighborhood for each vertex v is visited exactly once and thus, line 7 is executed $O(n)$ times, yielding a total run time of $O(n^3)$ for searching potential-sinks. The removal of a potential-sink in lines 10 to 15 runs in time $O(n^2)$ as v has at most n neighbors for which edges are removed from G and added to D and there are at most $O(n^2)$ tuples in the sets in B where v is included. By storing all tuples where v is included separately, we can iterate over them and remove each tuple from the corresponding set in $O(1)$. As every vertex is removed at most once

from the graph, the removal is repeated $O(n)$ times, yielding a run time of $O(n^3)$ for the removal of all vertices. Consequently, the run time of the whole algorithm is bounded by $O(n^3)$. ■

The correctness of the algorithm follows immediately from the correctness of the Dor-Tarsi algorithm (Fact 1) as our proposed modification only avoids repeated visits to neighbors but does not skip any neighbor check. As a final remark, we note that Algorithm 2 is also implementable using a combination of linked lists and an adjacency matrix to represent the neighbors of a vertex v, as proposed by (Wienöbst et al., 2021). Using such a representation, instead of collecting all neighbors of v violating the potential-sink property in line 7, one could stop as soon as the first pair of neighbors violating the property is found and store a pointer to the violating neighbor, allowing to start the iteration over the neighbors of v the next time at the lastly visited neighbor[1]. However, using linked lists instead of hash sets to represent adjacencies has no impact on the worst-case complexity and is not significantly faster than the hashed data structure. More details about the comparison between these two representations can be found in Appendix A.

4. Evaluation of Extension Algorithms

In this section, we conduct an experimental evaluation of the algorithms discussed in the previous sections to demonstrate their practical effectiveness. We compare the Dor-Tarsi algorithm (DT), its heuristic refinement (DTH), its adaptation with improved worst-case complexity (DTIC), and the algorithm given by Wienöbst et al. (2021) (WBL). All algorithms are implemented in Julia (Bezanson et al., 2017). We present the results for random PDAGs, which are generated by (i) creating a random DAG D, (ii) replacing all directed edges not participating in a v-structure by an undirected edge, and (iii) orienting between two and five (randomly chosen) undirected edges according to D, resulting in an extendable PDAG. The initial DAG D is generated by creating a random undirected graph and afterward using a random permutation of its vertices as a topological ordering according to which the edges are then oriented. We vary the number of vertices in the input graphs by setting $n = 128, 256, \ldots, 8192$ and the number of edges is set to $m = 3 \cdot n, \log_2(n) \cdot n$. For each choice of parameters, we generate ten instances and then run every algorithm ten times on each instance. The run times reported in this section and in upcoming sections are averages over all runs on those ten instances. Further experimental results for scale-free PDAGs and chordal graphs are given in Appendix B.

The results are shown in Fig. 5. In the left plot, the run times on sparse graphs with $m = 3 \cdot n$ edges are illustrated. For small graphs, there are no visible differences among the run times of the algorithms. However, with an increasing number of vertices, it becomes clear that WBL performs best, which is quite expected, as it was shown to have linear-time $O(m)$ for sparse graphs (formalized as constant-degeneracy graphs in (Wienöbst et al., 2021)). The run time of DT increases rapidly as the number of vertices grows, while both DTH and DTIC exhibit a slightly smaller increase. Although DTIC does not outperform WBL, it shows a significant improvement compared to DT.

For slightly denser graphs ($m = \log_2(n) \cdot n$, shown in the right plot), WBL yields a weaker performance than the other algorithms. This can be explained by the fact that WBL relies on a initialization step, which requires an iteration of every pair of neighbors for each vertex in the graph and thereby induces costs of $O(n^3)$ for dense graphs. Those costs emerge even before searching for the first potential-sink and they are particularly high as during the initialization loops cannot be

1. Such a strategy is not possible for hashed data structures due to potential rehashing.

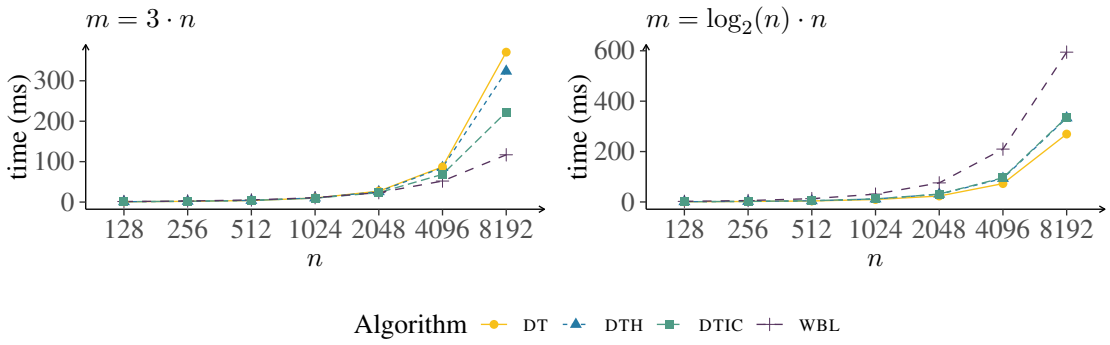

Figure 5: Run times of the algorithms DT, DTH, DTIC, and WBL on randomly generated PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

exited early. All of the three other approaches do allow for early exits in loops, which turns out to be a non-negligible advantage because the graph is getting smaller and smaller during the course of the algorithm, meaning that the high initialization effort of WBL dominates the run time. For some input graphs, DT is even slightly faster than DTH and DTIC, which indicates that both DTH and DTIC induce some overhead by continuously maintaining the vertices in sorted order by their degree.

Overall, DTIC provides the most stable performance across different graph sizes and graph densities, verifying that it combines the advantages of DT and WBL. DT is often a solid choice in practice but its performance depends heavily on the order of checking vertices for potential-sinkness and thus, combining its advantages with a guaranteed worst-case run time of $O(n^3)$ yields a promising algorithm for practical applications.

5. Application to Maximal Orientations

Building on the results in the previous sections, we are able to demonstrate an immediate application of extension algorithms. As shown by Chickering (1995) and Wienöbst et al. (2021), it is possible to utilize extension algorithms when computing the maximal orientation of a PDAG G. In practice, it is common to implement the step from PDAG to MPDAG the "direct way", that is, by repeatedly applying the Meek rules in a while-loop. We argue that using consistent extensions will not only give desirable worst-case guarantees, but also performs remarkably well in practice.

5.1. How to Use Extendability for the Computation of Maximal Orientations

Applying the Meek rules repeatedly to a PDAG G yields its maximal orientation. A direct implementation, which loops over the given graph repeatedly is computationally expensive, leading to worst-case run times such as $O(n^4 \cdot m)$[2]. Clearly, it is preferable to traverse the graph only a single time, while deciding which edges should be oriented. The crucial observation is that this can be achieved by utilizing a topological ordering of a consistent extension of G.

2. $O(m)$ edges might be oriented successively and naively checking the applicability of the Meek rules results in costs of $O(n^4)$.

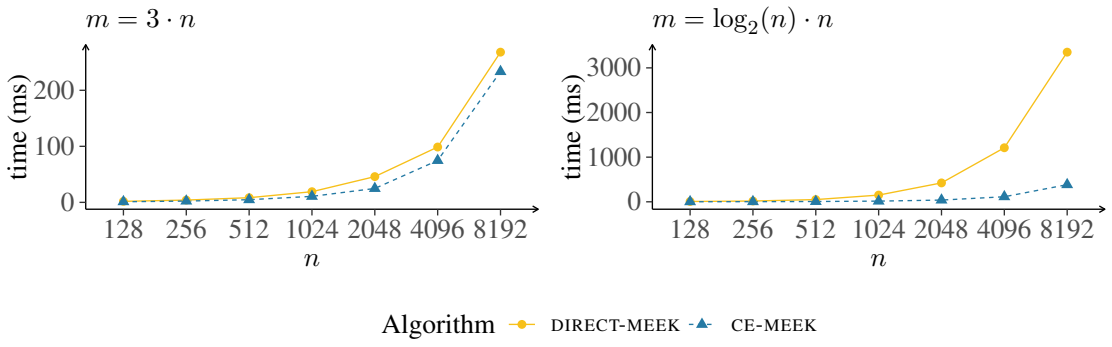

Figure 6: Run times of the algorithms DIRECT-MEEK and CE-MEEK on randomly generated PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

More precisely, when computing the maximal orientation of a PDAG, one can distinguish between two situations: (i) the result of applying the Meek rules yields a CPDAG, such as in the final phase of the PC algorithm, and (ii) in case of additional background knowledge, the resulting graph is not necessarily a CPDAG, but a maximally oriented PDAG (i. e., an MPDAG).

In case (i), the CPDAG can be obtained without applying the Meek rules by extending the PDAG into a DAG and afterward computing the corresponding CPDAG directly from the DAG. Indeed, for this second step from DAG to CPDAG, Chickering (1995) gave a linear-time algorithm[3]. In case (ii), starting with the CPDAG obtained as in (i), that is, by performing Chickering's DAG-to-CPDAG algorithm on a consistent extension D, further orienting all edges which are directed in G (i. e., the background knowledge edges), and afterward applying the Meek rules in a single iteration over the vertices of the graph in order of a topological ordering of D yields the maximal orientation (Wienöbst et al., 2021). While this procedure appears quite involved, we demonstrate that the extension step is actually the most expensive part and the subsequent steps are comparably cheap. Moreover, we show that this strategy is significantly faster than applying the Meek rules directly. From a theoretical point of view, these approaches guarantee an $O(n^3)$ worst-case run time, which is a significant improvement as well.

5.2. Experimental Comparison with the Direct Application of the Meek Rules

In this section, we compare the two approaches to maximally orient a given PDAG in an experimental evaluation. More specifically, we compare the direct application of Meek's rules (DIRECT-MEEK) to the above mentioned approach originally introduced by Chickering (1995) and generalized to arbitrary PDAGs by Wienöbst et al. (2021), which utilizes the topological ordering of a consistent extension (CE-MEEK). The input PDAGs are the same as in Section 4, and further results for scale-free PDAGs are provided in Appendix C.

The run times of the two approaches on random PDAGs are presented in Fig. 6. Not surprisingly, CE-MEEK outperforms DIRECT-MEEK on every input graph. While the difference in their run times

3. The algorithm proceeds by checking for each edge whether it should be undirected. Although not utilizing the Meek rules, the algorithm makes heavy use of the topological ordering provided by the consistent extension.

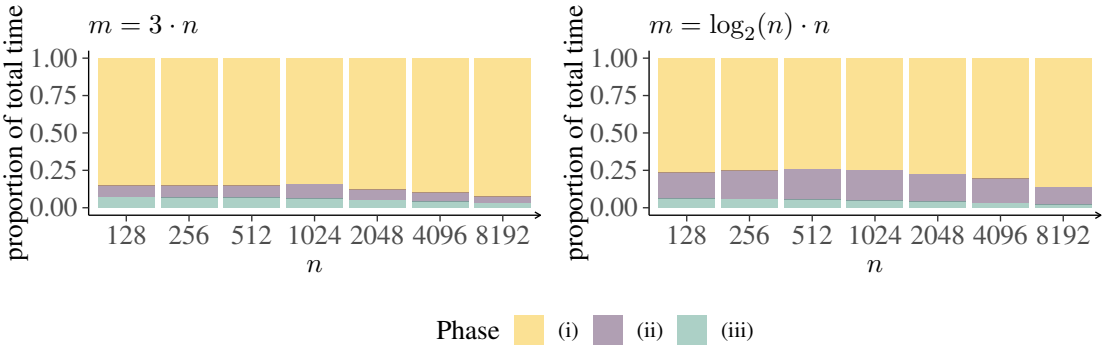

Figure 7: Proportions of the total run time of CE-MEEK for the three phases on randomly generated PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

is relatively small on sparse input graphs (left), it increases drastically as n grows on denser input graphs (right) [4]. In addition to the plain comparison of run times, we also analyzed the time spent by CE-MEEK on the different phases of the algorithm. More specifically, we measured the time for (i) computing a consistent extension, (ii) finding the corresponding CPDAG, and (iii) applying the Meek rules in a single iteration over the vertices. The results for the same input graphs as before are depicted in Fig. 7. Each bar is divided into the proportions of the total run time for the three phases, i. e., adding together the proportions of the three phases equals 100 percent of the run time. For both sparse graphs (left) and slightly denser graphs (right), the majority of the time (mostly more than 75 percent) is spent on phase (i), demonstrating that in order to obtain fast algorithms to compute maximal orientations, it is crucial to have algorithms solving the extension problem efficiently.

6. Conclusions

In this paper, we demonstrate the effectiveness of utilizing consistent extensions for the task of maximally orienting a PDAG. We started by revisiting the extension problem, presenting two new approaches to efficiently compute consistent DAG extensions in practical applications. The first approach (DTH) refines the widespread Dor-Tarsi algorithm by the employment of a simple heuristic, which reduces the computational effort by prioritizing low-cost vertices throughout its iterations. The second approach (DTIC) stores additional information between iterations to avoid duplicate iterations and thereby matches the worst-case complexity of the WBL algorithm, which achieves the conditional lower bound of $O(n^3)$ for the extension problem. In a practical evaluation, we show that DTIC exhibits the most stable performance overall, combining the advantages of the other approaches. Based on those insights and results, we highlight an important application – the maximal orientation of PDAGs, a procedure ubiquitous in causal discovery. We demonstrate experimentally that utilizing consistent extensions yields highly reliable and effective algorithms for this task, which outperform the direct use of the Meek rules currently used most commonly.

4. We use DTIC to compute the consistent extension. For sparse graphs, the use of WBL can give a significant speedup, which would in turn translate to CE-MEEK.

References

Réka Albert and Albert-László Barabási. Statistical Mechanics of Complex Networks. *Reviews of Modern Physics*, 74:47–97, 2002.

Steen A Andersson, David Madigan, and Michael D Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2):505–541, 1997.

Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, 59:65–98, 2017.

David Maxwell Chickering. A Transformational Characterization of Equivalent Bayesian Network Structures. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 87–98. Morgan Kaufmann Publishers Inc., 1995.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.

Dorit Dor and Michael Tarsi. A Simple Algorithm to Construct a Consistent Extension of a Partially Oriented Graph. Technicial Report 185, Cognitive Systems Laboratory, UCLA, 1992.

Felix Elwert. Graphical causal models. In *Handbook of Causal Analysis for Social Research*, Handbooks of Sociology and Social Research, pages 245–273. Springer, 2013.

Richard Guo and Emilija Perkovic. Minimal enumeration of all possible total effects in a markov equivalence class. In *International Conference on Artificial Intelligence and Statistics*, pages 2395–2403. PMLR, 2021.

Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012.

Yang-Bo He and Zhi Geng. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research*, 9(Nov):2523–2547, 2008.

Yangbo He, Jinzhu Jia, and Bin Yu. Counting and exploring sizes of Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 16(79):2589–2609, 2015.

Markus Kalisch and Peter Bühlman. Estimating high-dimensional directed acyclic graphs with the pc-algorithm. *Journal of Machine Learning Research*, 8(3), 2007.

Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H Maathuis, and Peter Bühlmann. Causal inference using graphical models with the R package pcalg. *Journal of statistical software*, 47:1–26, 2012.

Daphne Koller and Nir Friedman. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009. ISBN 978-0-262-01319-2.

Marloes H Maathuis, Markus Kalisch, and Peter Bühlmann. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, 37(6A):3133–3164, 2009.

Christopher Meek. Causal Inference and Causal Explanation with Background Knowledge. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 403–410. Morgan Kaufmann Publishers Inc., 1995.

Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition, 2009.

Jean-Philippe Pellet and André Elisseeff. Using markov blankets for causal structure learning. *Journal of Machine Learning Research*, 9(7), 2008.

Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. Complete graphical characterization and construction of adjustment sets in Markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18:220:1–220:62, 2017.

Oylum Şeker, Pinar Heggernes, Tınaz Ekim, and Z. Caner Taşkın. Linear-Time Generation of Random Chordal Graphs. In *Proccedings of the 10th International Conference on Algorithms and Complexity (CIAC-17)*, pages 442–453. Springer International Publishing, 2017.

Karthikeyan Shanmugam, Murat Kocaoglu, Alexandros G. Dimakis, and Sriram Vishwanath. Learning causal graphs with small interventions. In *Processing of the 28th Annual Conference on Neural Information Processing Systems, NIPS '15*, pages 3195–3203, 2015.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition, 2000.

Chandler Squires. `causaldag`*: creation, manipulation, and learning of causal models*, 2018. URL https://github.com/uhlerlab/causaldag.

Chandler Squires, Sara Magliacane, Kristjan Greenewald, Dmitriy Katz, Murat Kocaoglu, and Karthikeyan Shanmugam. Active structure learning of causal DAGs via directed clique trees. In *Advances in Neural Information Processing Systems*, volume 33, pages 21500–21511, 2020.

Benito van der Zander and Maciej Liśkiewicz. Separators and adjustment sets in Markov equivalent DAGs. In *Proceedings of the 30th Conference on Artificial Intelligence, AAAI '16*, pages 3315–3321, 2016.

Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence, UAI'90*, pages 255–270, 1990.

Thomas Verma and Judea Pearl. An algorithm for deciding if a set of observed independencies has a causal explanation. In *Proc. of the Conference on Uncertainty in Artificial Intelligence, UAI 1992*, pages 323–330. Elsevier, 1992.

Marcel Wienöbst, Max Bannach, and Maciej Liśkiewicz. Polynomial-time algorithms for counting and sampling Markov equivalent DAGs. In *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI 2021*, volume 35, pages 12198–12206. AAAI Press, 2021.

Marcel Wienöbst, Max Bannach, and Maciej Liśkiewicz. Extendability of Causal Graphical Models: Algorithms and Computational Complexity. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI-21)*, pages 1248–1257. PMLR, 2021.

Appendix A. Comparison of Graph Representations

Choosing an appropriate graph representation is crucial to obtain effective extension algorithms. In particular, two important operations that need to be executed fast to solve the extension problem efficiently are adjacency tests and the removal of vertices and edges from the graph. In this section, we compare two different graph representations to support carrying out these operations efficiently. The first representation utilizes a hashed data structure to represent adjacencies in a graph and a the second one employs a combination of linked lists and an adjacency matrix for this purpose.

The graph representation used throughout the course of this paper keeps track of three sets storing the neighbors (ingoing, outgoing, and undirected) for each vertex in the graph. To access elements in these sets in expected constant time, hashing is applied. Representing the neighbors of every vertex with the help of hash sets leads to a simple, yet practical graph representation. However, iterating over hash sets (e. g., during the potential-sink check) is suboptimal, as the entire hash table has to be iterated over. Furthermore, as outlined in Section 3.2, a representation based on linked lists allows DTIC to exit its loops earlier, that is, as soon as a neighbor violating the potential-sink property is found, which is not possible if neighbors are stored in a hashed data structure.

To analyze the impact of these aspects on the performance of DTIC and the other algorithms, we also implemented a graph representation using an adjacency matrix and linked lists to store adjacencies in the graph, which was proposed by (Wienöbst et al., 2021). More precisely, there are three linked lists to store the ingoing, outgoing, and undirected neighbors, respectively, for each vertex. The adjacency matrix contains a pointer to the corresponding linked list entry for every edge in the graph. Adjacency tests run in constant time on the adjacency matrix and the removal of vertices from the graph runs in constant time as well, as the corresponding pointer in the matrix enables an access to any neighbor in the linked list in $O(1)$. The main drawback of this approach is clearly the large memory requirement of the $O(n^2)$ adjacency matrix. Allocating the corresponding memory is also time-consuming.

We found that the usage of linked lists instead of hash sets does not provide significant improvements for the run times of the extension algorithms. In most of the evaluated scenarios, the contrary is the case. A direct comparison of implementations using hash sets and linked lists can be found in Table 1 where the input graphs and settings are the same as in Section 4.

Clearly, the usage of linked lists appears to be at a disadvantage for the instances considered in this work. However, we remark that, in some cases, the usage of linked lists is beneficial. For example, on chordal graphs (see Appendix B for more details), DT is faster on graphs with a linked list implementation compared to graphs using a hashed data structure. In Appendix B, we also introduce scale-free PDAGs, for which the comparison of hash sets and linked lists yields the same relations as in Table 1. In conclusion, our experiments demonstrate that the linked list representation introduces overhead, which does not pay off for the generally sparse graphs considered in this work and usually occuring in practice.

Appendix B. Further Experimental Results for Extendability

To complement the experimental results for randomly generated PDAGs presented in Section 4, we evaluate the algorithms DT, DTH, DTIC, and WBL on scale-free PDAGs and chordal graphs.

Scale-free graphs are graphs whose degree distribution is a power law distribution, i. e., there are few vertices with a high degree while most of the vertices have a rather small degree. The scale-free PDAGs are generated in a similar fashion as the graphs from Section 4, that is, the general

n	DT	DT-LL	DTH	DTH-LL	DTIC	DTIC-LL	WBL	WBL-LL
128	0.83	0.14	0.88	0.19	0.94	0.20	1.29	0.26
256	1.80	0.41	1.89	0.54	1.97	0.58	2.60	0.68
512	3.89	1.60	4.07	1.88	4.13	2.03	5.21	2.15
1024	9.96	6.76	10.19	7.62	9.72	8.28	10.88	9.53
2048	27.01	39.55	25.66	40.40	23.32	42.86	23.35	50.90
4096	87.13	149.06	85.38	190.61	68.50	201.69	51.98	236.52
8192	370.79	606.02	323.64	608.67	222.52	654.45	116.92	871.23

n	DT	DT-LL	DTH	DTH-LL	DTIC	DTIC-LL	WBL	WBL-LL
128	0.87	0.19	0.97	0.28	1.04	0.29	2.35	0.46
256	1.86	0.55	2.11	0.81	2.23	0.84	5.61	1.21
512	3.92	1.97	4.71	2.86	4.93	2.98	13.09	4.30
1024	9.62	9.14	11.90	11.68	12.14	12.08	30.98	15.93
2048	23.60	48.02	30.77	54.55	31.38	56.30	75.68	72.03
4096	72.80	175.78	93.31	236.88	96.00	244.46	209.80	323.04
8192	269.53	685.58	333.44	752.79	336.09	785.27	594.51	1159.83

Table 1: Comparison of the algorithms DT, DTH, DTIC, and WBL with a linked lists implementation (-LL) on random PDAGs with $m = 3 \cdot n$ edges (above) and $m = \log_2(n) \cdot n$ edges (below).

procedure of (i) creating a random DAG D, (ii) replacing all directed edges not participating in a v-structure by an undirected edge, and (iii) orienting between two and five (randomly chosen) undirected edges according to D stays the same but the initial DAG D is generated in a different way. More precisely, the initial DAG is generated by first creating a random undirected scale-free graph using the Barabási-Albert model (Albert and Barabási, 2002) and afterward constituting a random permutation of its vertices as a topological ordering according to which the edges are then directed. All parameter choices are identical to those in Section 4.

To generate the chordal graphs, we apply the random subtree intersection method introduced by Şeker et al. (2017). Chordal graphs provide an interesting addition to PDAGs as they are fully undirected and extendable by definition, thus demanding quite some computational effort to extend the graph. We set $n = 128, 256, \ldots, 8192$ again and use $k = 3, \log_2(n)$ as a parameter for the random subtree intersection method. The parameter k determines the average size of the random subtrees used to generate the chordal graph and thus, controls the number of edges in the graph. However, the exact number of edges slightly differs between instances.

The results for scale-free PDAGs can be found in Fig. 8. Generally, we can observe the same patterns as in Fig. 5, i. e., WBL is the fasted for sparse graphs (left) but does not perform well for slightly denser graphs (right) while the opposite holds for DT. In total, DTIC yields the most stable performance again. Note that, although the overall relations stay identical, the run times are slightly higher in Fig. 8 than in Fig. 5, showing that scale-free PDAGs require a little more computational effort than PDAGs whose edges are distributed at random.

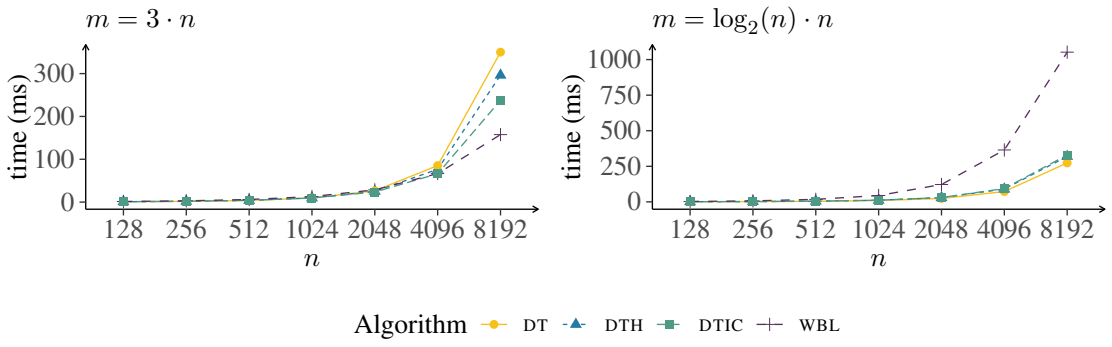

Figure 8: Run times of the algorithms DT, DTH, DTIC, and WBL on randomly generated scale-free PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

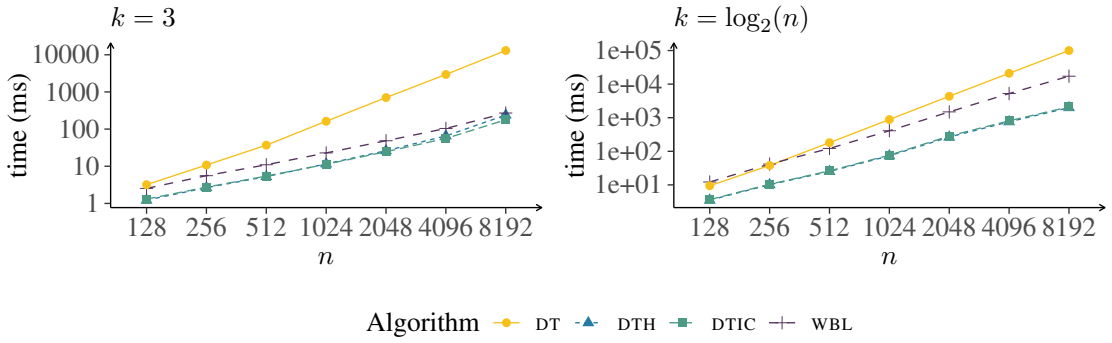

Figure 9: Run times of the algorithms DT, DTH, DTIC, and WBL on randomly generated chordal graphs with $k = 3$ (left) and $k = \log_2(n)$ (right).

Further, Fig. 9 displays the run times of DT, DTH, DTIC, and WBL on chordal graphs. Note that the y-axes are log-scaled in both plots. Clearly, DT is inferior to all other algorithms for both sparse graphs (left) and slightly denser graphs (right). Eventually, the run time of DT becomes almost 100 times slower than the run times of the other algorithms at $n = 8192$ and $k = 3$. WBL handles sparse graphs roughly as good as DTH and DTIC but cannot keep pace for slightly denser graphs. The results indicate that DTH and DTIC yield the best performance on graphs where more effort is necessary to compute a consistent extension.

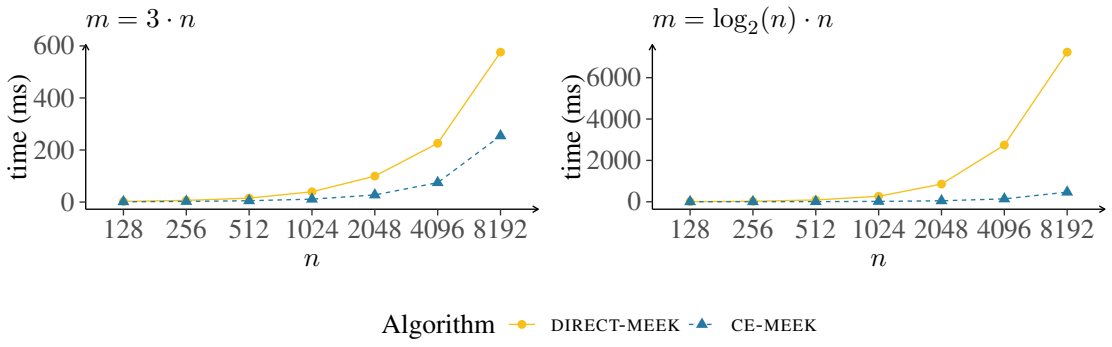

Figure 10: Run times of the algorithms DIRECT-MEEK and CE-MEEK on randomly generated scale-free PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

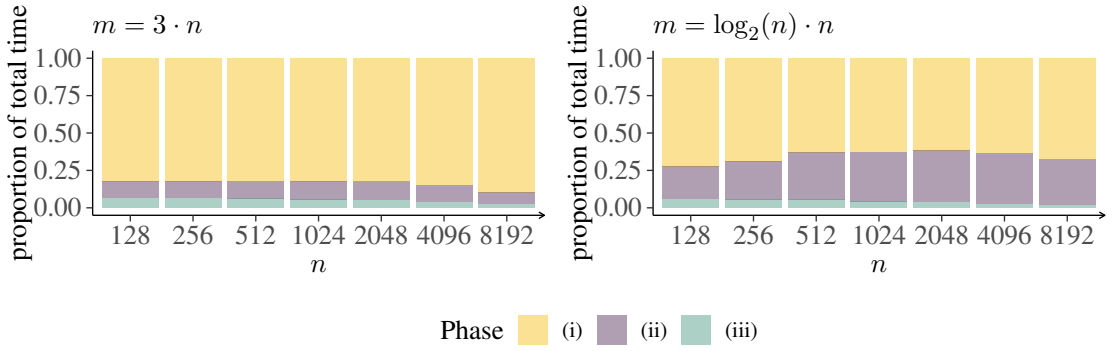

Figure 11: Proportions of the total run time of CE-MEEK for the three phases on randomly generated scale-free PDAGs with $m = 3 \cdot n$ edges (left) and $m = \log_2(n) \cdot n$ edges (right).

Appendix C. Further Experimental Results for Maximal Orientations

For the sake of completeness, we also give further experimental results for the algorithms DIRECT-MEEK and CE-MEEK in addition to the results presented in Section 5.2. We report the results for scale-free PDAGs which are identical to those from the previous section[5].

Fig. 10 contains the run times of DIRECT-MEEK and CE-MEEK on scale-free PDAGs. As in Fig. 6, CE-MEEK is superior to DIRECT-MEEK both on sparse graphs (left) and on slightly denser graphs (right). The lead of CE-MEEK is now even larger than it was in Fig. 6. Overall, handling scale-free PDAGs takes more computational effort compared to handling PDAGs having their edges distributed at random.

5. Chordal graphs are obviously not interesting in the setting of maximal orientations as they contain only undirected edges and hence, no Meek rule is applicable on them.

Furthermore, a visualization of the time spent by CE-MEEK on the different phases of the algorithm for scale-free input PDAGs is given in Fig. 11. While the plot for sparse graphs (left) looks similar to the left plot from Fig. 7, the plot on the right exhibits greater proportions of phase (ii) than the right plot in Fig. 7, showing that phase (ii) (i. e., finding the corresponding CPDAG to the DAG computed in the first phase) requires more effort on scale-free PDAGs than on PDAGs having their edges distributed at random. However, phase (i) (i. e., computing a consistent DAG extension) still requires the majority of the total run time in every scenario.