



HAL
open science

Enumeration of Unordered Forests

Florian Ingels, Romain Azaïs

► **To cite this version:**

| Florian Ingels, Romain Azaïs. Enumeration of Unordered Forests. 2020. hal-02511901v1

HAL Id: hal-02511901

<https://hal.science/hal-02511901v1>

Preprint submitted on 19 Mar 2020 (v1), last revised 13 Apr 2022 (v4)

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

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

ENUMERATION OF UNORDERED FORESTS

Florian Ingels
florian.ingels@inria.fr

Romain Azaïs
romain.azais@inria.fr

Laboratoire Reproduction et Développement des Plantes, Univ Lyon, ENS de Lyon, UCB Lyon 1, CNRS,
INRAE, Inria, F-69342, Lyon, France

Abstract

Reverse search is a convenient method for enumerating structured objects, that can be used both to address theoretical issues and to solve data mining problems. This method has already been successfully developed to handle unordered trees. If the literature proposes solutions to enumerate singletons of trees, we study in this article a more general, higher combinatorial problem, the enumeration of sets of trees – forests. By compressing each forest into a Directed Acyclic Graph (DAG), we develop a reverse search like method to enumerate DAG compressing forests. Remarkably, we prove that these DAG are in bijection with the row-Fishburn matrices, a well-studied class of combinatorial objects. In a second step, we derive our forest enumeration to provide algorithms for tackling two related problems : (i) the enumeration of “subforests” of a forest, and (ii) the frequent “subforest” mining problem. All the methods presented in this article enumerate each item uniquely, up to isomorphism.

keywords: Directed Acyclic Graph, Reverse Search, Unordered Trees, Enumeration, Forest

1 Introduction

1.1 Context of the work

Enumeration of trees is a long-term problem, where they were first counted in the 80’s [16]. The exhaustive enumeration of ordered and unordered trees was successfully tackled in the early 00’s by Nakano and Uno in [17, 18]. In the unordered case, an extension of the algorithm has been proposed to solve the problem of frequent substructure mining [1]. Moreover, in the field of machine learning, we have recently demonstrated that exhaustive enumeration of the subtrees of a tree makes it possible to design classification algorithms significantly more efficient than their counterpart without such enumeration [5].

Our ambition in this article is to take these two problems of enumeration – trees and subtrees – to a higher order, i.e. to enumerate sets of trees where only singletons have been considered in the literature so far. In other words, we are interested in the problem of enumerating forests of unordered trees, and then, given a tree or forest, to enumerate all its “subforests” – as forests of subtrees. The latter has already been discussed in the literature, but without consideration on isomorphism [20]. We emphasize here that we aim to enumerate these various items – forests and subforests – up to isomorphism.

Such an ambition immediately raises a number of obstacles. First of all, the trees are indeed unordered, but so are the sets of trees. For the former the literature has introduced the notion of the canonical form of a tree [18, 1], which is a unique ordered representation of an unordered tree. The enumeration therefore focuses only on these canonical trees. Unfortunately, if it is possible to order a set of vertices, there is no total order on the set of trees, to the best of our knowledge. In addition, forests belong to the powerset of trees and have an exponentially stronger combinatorics, making the enumeration problem that much trickier.

Enumeration problems are recurrent in many fields, notably combinatorial optimization and data mining. They involve the exhaustive listing of a subset of the elements of a search set (possibly all of them), e.g. graphs, trees or vertices of a simplex. Given the possibly high combinatorial nature of these elements, it is essential to adopt clever exploration strategies as opposed to brute-force enumeration, typically to avoid areas of the search set not belonging to the objective subset.

One proven way of proceeding is to provide the search set with an enumeration tree structure; starting from the root, the branches of the tree are explored recursively, eliminating those that do not address the problem. Based on this principle, we can notably mention the well-known “branch and bound” method in combinatorial optimization [15] and the gSpan algorithm for frequent subgraph mining in data mining [24]. Another of these methods is the so-called reverse search technique, which requires that the search set has a partial order structure, and which has solved a large number of enumeration problems since its introduction [2] until very recently [23]. Actually, the algorithms previously introduced in the literature to enumerate trees are based on this technique [17, 18, 1].

In the present paper, we restrict ourselves to reverse search methods, for which the following formalism is adapted from the one that can be found in [19, p. 45-51], and slightly differ from the original definition by Avis and Fukuda [2]. We refer the reader to these two references for further details.

Let (\mathcal{S}, \subseteq) be a partially order set, and $g : \mathcal{S} \rightarrow \{\top, \perp\}$ be a *property*, satisfying anti-monotonicity

$$\forall s, t \in \mathcal{S} : (s \subseteq t) \wedge g(t) \implies g(s).$$

The *enumeration problem* for the property g is the problem of listing all elements of $E_{\mathcal{S}}(g) = \{s \in \mathcal{S} : g(s) = \top\}$. An enumeration algorithm is an algorithm that returns $E_{\mathcal{S}}(g)$.

The reverse search technique relies on inverting a *reduction rule* $f : \mathcal{S} \setminus \emptyset \rightarrow \mathcal{S}$, where f satisfies the two properties of (i) *covering*: $\forall s \in \mathcal{S} \setminus \emptyset, f(s) \subset s$ and (ii) *finiteness*: $\forall s \in \mathcal{S} \setminus \emptyset, \exists k \in \mathbb{N}^*, f^k(s) = \emptyset$. Then, the *expansion rule* is defined as $f^{-1}(t) = \{s \in \mathcal{S} : f(s) = t\}$. This defines an enumeration tree rooted in \emptyset , and repeated call to f^{-1} can therefore enumerates all the elements of \mathcal{S} .

The reverse search algorithm is shown in Algorithm 1. $E_{\mathcal{S}}(g)$ can be obtained from the call of `REVERSESEARCH` $((\mathcal{S}, \subseteq), f^{-1}, g, \emptyset)$. As g is anti-monotone, if $g(s) = \perp$, then all elements $s \subseteq t$ also have $g(t) = \perp$, and therefore pruning the enumeration tree in s does not miss any element of $E_{\mathcal{S}}(g)$.

Algorithm 1: REVERSESEARCH

Input: $(\mathcal{S}, \subseteq), f^{-1}, g, s_0 \in \mathcal{S} - \text{s.t.}$
 $g(s_0) = \top$

- 1 **output** s_0 ;
 - 2 **for** $t \in \{s \in f^{-1}(s_0) | g(s) = \top\}$ **do**
 - 3 `REVERSESEARCH` $((\mathcal{S}, \subseteq), f^{-1}, g, t)$;
-

1.2 Precise formulation of the problem

A rooted tree T is a connected graph with no cycle such that there exists a unique vertex called the root, which has no parent, and any vertex different from the root has exactly one parent. Rooted trees are said unordered if the order between the sibling vertices of any vertex is not significant. As such, the set of children of a vertex v is considered as a multiset and denoted $\mathcal{C}(v)$. The leaves $\mathcal{L}(T)$ are all the vertices without children. The height of a vertex v of a tree T can be recursively defined as

$$\mathcal{H}(v) = \begin{cases} 0 & \text{if } v \in \mathcal{L}(T), \\ 1 + \max_{u \in \mathcal{C}(v)} \mathcal{H}(u) & \text{otherwise.} \end{cases} \quad (1)$$

The height $\mathcal{H}(T)$ of the tree T is defined as the height of its root. The outdegree of a vertex $v \in T$ is defined as $\deg(v) = \#\mathcal{C}(v)$; the outdegree of T is then defined as $\deg(T) = \max_{v \in T} \deg(v)$. The depth of a vertex v is the number of edges on the path from v to the root of the tree.

Two trees T_1 and T_2 are isomorphic if there exists a one-to-one correspondance ϕ between the vertices of the trees such that (i) $u \in \mathcal{C}(v)$ in $T_1 \iff \phi(u) \in \mathcal{C}(\phi(v))$ in T_2 and (ii) the roots are mapped together. For any vertex v of T , the subtree $T[v]$ rooted in v is the tree composed of v and all its descendants $\mathcal{D}(v)$. $\mathcal{S}(T)$ denotes the set of subtrees of T , which is the quotient set of $\{T[v] : v \in T\}$ by the tree isomorphism relation. In this article, we consider only unordered rooted trees that will simply be called trees in the sequel. The set of all trees is denoted \mathcal{T} .

As mentioned before, we are interested in this paper in the enumeration of *forests*. The literature acknowledges two definitions for a forest [6, p. 172]: (i) an undirected graph in which any two vertices are connected by at most one path or (ii) a disjoint union of trees. We adopt a slight variation of the latter one:

Definition 1.1. A set $\{T_1, \dots, T_n\}$ of trees is a forest if and only if

$$\forall i \neq j, T_i \notin \mathcal{S}(T_j). \quad (2)$$

We denote \mathcal{F} the set of all forests. Our goal is therefore to provide a reverse search like method that outputs \mathcal{F} . As already stated, this goal raises two major difficulties: firstly, the twofold unordered nature of forests (the set of trees and the trees themselves), and secondly, the higher combinatorics of forests compared to trees. While the latter problem is intrinsic, the main idea of this paper to address the former is to resort to the reduction of a forest into a Directed Acyclic Graph (DAG).

DAG reduction is a method meant to eliminate internal repetitions in the structure of trees and forests of trees. Beginning with [21], DAG representations of trees are also much used in computer graphics where the process of condensing a tree into a graph is called object instancing [11]. A precise definition of DAG reduction of trees, together with algorithms to compute it, are provided in [9], whereas one technique to extend those algorithms to forests is presented in [5, Section 3.2]. DAG reduction can be interpreted as the construction of the quotient graph of a forest by the tree isomorphism relation. However, in this paper, we provide the general idea of DAG reduction as a vertex coloring procedure.

Consider a forest $F = \{T_1, \dots, T_n\}$ to reduce. Each vertex of each tree is given a color such that if two distinct vertices u, v belonging respectively to T_i, T_j (not necessarily distinct) have the same

color, then $T_i[u]$ and $T_j[v]$ are isomorphic. Reciprocally, if two subtrees are isomorphic, their roots have to be identically colored. Let us denote $c(\cdot)$ the function that associates a color to any vertex. Then, we build a directed graph $D = (V, A)$ with as many vertices as colors used, i.e. $\#V = \#\text{Im}(c)$. For any two vertices u, v in the forest, if $u \in \mathcal{C}(v)$, then we create an arc $c(v) \rightarrow c(u)$ in D . Note that this definition implies that multiples arcs are possible in D , as if there exist $u, u' \in \mathcal{C}(v)$, for $v \in T$, such that $T[u]$ and $T[u']$ are isomorphic, then the arcs $c(v) \rightarrow c(u)$ and $c(v) \rightarrow c(u')$ are identical. The graph D is a DAG [9, Proposition 1], i.e. a connected directed graph without cycles. We refer to Figure 1 for an example of DAG reduction.

The DAG structure inherits of some properties of trees. For a vertex v in a DAG D , we will denote by $\mathcal{C}(v)$ the set of children of v . $\mathcal{H}(v)$ and $\text{deg}(v)$ are inherited as well. Similarly to trees, we denote by $D[v]$ the subDAG rooted in v composed of v and all its descendants $\mathcal{D}(v)$.

In this paper, $\mathfrak{R}(F)$ denotes the DAG reduction of F . It is crucial to notice that the function \mathfrak{R} is a one-to-one correspondence, which means that DAG reduction is a lossless compression algorithm. Since F fulfills condition (2), no tree of F is a subtree of another. If this were the case, say $T_i \in \mathcal{S}(T_j)$, then $\mathfrak{R}(T_i)$ would be a subDAG of $\mathfrak{R}(T_j)$, and therefore the numbers of roots in $\mathfrak{R}(F)$ would be strictly less than $\#F$. Since such a situation cannot occur, there are exactly as many roots in $\mathfrak{R}(F)$ as there are elements in F : no information is lost. In other words, F can be reconstructed from $\mathfrak{R}(F)$ and \mathfrak{R}^{-1} stands for the inverse function.

In the sequel, DAG compressing forests are called FDAG, to distinguish them from general directed acyclic graphs.

Since DAG compression is lossless, and since a forest can be reconstructed from its DAG reduction, it should be clear that enumerating all forests is equivalent to enumerating all FDAG. Yet, the latter approach has the merit of transforming set of trees into unique objects, which makes it possible, if able to design a canonical representation – like the trees in [18, 1], to get rid of the twofold unordered nature of forests, as claimed earlier. Indeed, any ordering of the vertices of the DAG induces an order on the roots of the DAG, and therefore on the elements of the forest, as well on the vertices of the trees themselves.

1.3 Aim of the paper

To the best of our knowledge, the enumeration of DAG has never been considered in the literature. The aim of this article is twofold, i.e (i) to open the way by presenting a reverse search algorithm enumerating FDAG, in Section 2, and (ii) to derive from it an algorithm for enumerating substructures in Section 3. The frequent pattern mining problem is a classical data mining problem – see [10] for a survey on that question – and we provides in Section 4 a slight variation of the algorithm of Section 3 to tackle this issue. In more details, our outline is as follows:




Figure 1: A forest F (left) and its DAG reduction (right). Roots of isomorphic subtrees are identically colored, as well as the corresponding vertex of the DAG. Red arrows indicate on the DAG the roots of the trees of F . Arcs with multiplicity greater than one are displayed with their multiplicity number.

- The first step is to introduce a canonical form for FDAG. For trees [18, Section 3], this consisted in associating an integer (its depth) to each vertex, and maximizing the sequence by choosing an appropriate ordering over the vertices. The notion of depth does not apply to FDAG, which forces us to find another strategy. DAG are characterized by the existence of a topological ordering [14], and we introduce in Subsection 2.2 a topological ordering that is unique if and only if a DAG compresses a forest. This canonical ordering is defined so that the sequence of children of the vertices is strictly increasing, where the multisets of children are ordered by the lexicographical order. In fact, these ordered multisets of children are considered as formal words, which brings us to a detour through the theory of formal languages in Subsection 2.1 to introduce useful results for the rest of the article. Compared to trees, we have here a first gain in complexity insofar as we maximize a sequence of words instead of a sequence of integers.
- The expansion rule used for trees [18, Section 4] is to add a new vertex in the tree as a child of some other vertex, so that the depth-sequence remains maximal. Consequently, a single arc is also added. On the other hand, for a FDAG, we want to be able to add either vertices or arcs independently. In Subsection 2.3, we define three expansion rules, reflecting the full spectrum of possible operations, so that the DAG obtained afterward is still a FDAG. Specifically, the branching rule allows to add an arc, where the elongation and widening rules add vertices at different height. Subsection 2.4 is primarily dedicated to showing (i) that the rules preserve canonicalness, and (ii) that a reduction rule can be derived from them; the resulting enumeration tree is then discussed.
- Notably, a bijection between FDAG and row-Fishburn matrices, a class of combinatorial objects much exploited in the literature – see [12, Section 2] for an overview, is shown in Theorem 2.11. The asymptotic behavior of these matrices being well known [13, 7], this allows us to derive from it the behavior of the enumeration tree. In return, since our bijection is constructive, the enumeration tree can be used to enumerate row-Fishburn matrices – and all the objects they are in bijection with – via the reverse search method. Remarkably, this bijection operates between two objects that, at first sight, have little in common.
- It has been mentioned above that the set of FDAG has a combinatorics in the order of the powerset of \mathcal{T} . For an enumeration algorithm to have any practical interest, it is necessary that the associated enumeration tree has a “reasonable” growth – with regard to the combinatorics of the explored space. This is the case for our algorithm since we prove, still in Subsection 2.4, that a FDAG with n vertices has a number of successors in the enumeration tree in the order of $\Theta(n)$. Finally, Subsection 2.5 concludes on enumeration by proposing sets of constraints that make the enumeration tree finite. Indeed, since the rules only allow to increase the height, degree or number of vertices, it is sufficient to set maximum values for some of these parameters to achieve this goal; however the combination of parameters has to be wisely chosen, as we show it.
- Since the structures we enumerate are forests, it is natural that the substructures we are interested in are “subforests”. A precise definition of the latter is given in Section 3, i.e. forests of subtrees, and are referred to as subFDAG. An algorithm to enumerate all subFDAG appearing in a FDAG is also provided. The frequent subFDAG mining problem is finally

addressed in Section 4. Concluding remarks concerning the implementation of our results in the Python library `treex` [3] are briefly mentioned at the end of the article.

2 Exhaustive enumeration of FDAG

In this section, we introduce our main result, that is, a reverse search algorithm for the enumeration of FDAG. As we will consider the multisets of children of vertices as formal words built on the alphabet formed by the set of vertices, we introduce in Subsection 2.1 some definitions and results on formal languages that will be useful for the sequel. We characterize unambiguously in Subsection 2.2 our objects of study, through the lens of topological orderings, defining a canonical topological ordering for DAG, that is unique if and only if a DAG compresses a forest of unordered trees, i.e. it is a FDAG – see Theorem 2.4. We then define three expansion rules that are meant to extend the structures of FDAG in Subsection 2.3, whereas Subsection 2.4 is dedicated (i) to prove that these rules preserve the canonicalness property and (ii) to discuss about the enumeration tree they define with respect to the reverse search technique. Notably, we exhibit a bijection – Theorem 2.11 – with a class of combinatorial objects from the literature, allowing us to obtain an asymptotic expansion of the growth of the tree. In addition, we show that any FDAG has a linear number of children in that tree in Theorem 2.13. Finally, options to constrain the enumeration tree – on maximum number of vertices, height or outdegree – and making it finite are proposed in Subsection 2.5.

2.1 Preliminary: a detour through formal languages

We present in this subsection some definitions and results on formal languages that will be useful for the sequel of Section 2.

Let \mathcal{A} be a totally ordered finite set, called alphabet, whose elements are called letters. A word is a finite sequence of letters of \mathcal{A} . The length of a word w is equal to its number of letters and is denoted $\#w$. There is a unique word with 0 letter called the empty word and denoted ϵ . The set of all words is denoted \mathcal{A}^* . Words can be concatenated to create a new word whose length is the sum of the lengths of the original words; ϵ is the neutral element of this concatenation operation.

The lexicographical order over \mathcal{A}^* , denoted $<_{\text{lex}}$, is defined as follows. Let $w_1 = a_0 \cdots a_p$ and $w_2 = b_0 \cdots b_q$ be two words, with $a_i, b_j \in \mathcal{A}$. If $\#w_1 = \#w_2$, then $w_1 <_{\text{lex}} w_2$ if and only if $\exists k \in \llbracket 0, p \rrbracket, a_i = b_i \forall i < k$ and $a_k < b_k$. Otherwise, let $m = \min(p, q)$; $w_1 <_{\text{lex}} w_2$ if and only if either (i) $a_0 \cdots a_m <_{\text{lex}} b_0 \cdots b_m$ or (ii) $a_0 \cdots a_m =_{\text{lex}} b_0 \cdots b_m$ and $m < q$ – that is, $p < q$. Note that, by convention, $\epsilon <_{\text{lex}} w$ for any word w .

Let $w \in \mathcal{A}^*$. We define the suffix-cut operator $SC(w)$, which removes the last letter of w :

$$SC(w) = \begin{cases} w' & \text{if } w = w'a \text{ with } a \in \mathcal{A} \text{ and } w' \in \mathcal{A}^*, \\ \epsilon & \text{otherwise.} \end{cases}$$

A *language* is a set of words satisfying some construction rules. We introduce hereafter two languages that will be useful in the sequel of the paper.

Definition 2.1. The language of decreasing words is defined as

$$\Lambda = \{w = a_0 \cdots a_m \in \mathcal{A}^* : a_i \geq_{\text{lex}} a_{i+1} \forall i \in \llbracket 0, m-1 \rrbracket\}.$$

Definition 2.2. Let $\bar{w} \in \Lambda$. The language of decreasing words bounded by \bar{w} is defined as

$$\Lambda^{\bar{w}} = \{w \in \Lambda : w >_{\text{lex}} \bar{w}\}.$$

Any word $w \in \Lambda^{\bar{w}}$ is said to be minimal if and only if $w \in \Lambda^{\bar{w}}$ but $\text{SC}(w) \notin \Lambda^{\bar{w}}$.

As an example, if $\mathcal{A} = \{0, 1, 2, 3\}$, then $\bar{w} = 211 \in \Lambda$, whereas $121 \notin \Lambda$. In addition, $\Lambda^{\bar{w}}$ contains words such as 31, 22, 21110, etc. 22 is a minimal word of $\Lambda^{\bar{w}}$ as $22 >_{\text{lex}} 211$ but $\text{SC}(22) = 2 <_{\text{lex}} 211$.

Our focus is now on the construction of the minimal words of $\Lambda^{\bar{w}}$. Let $\bar{w} = a_0 \cdots a_p$ and $w = b_0 \cdots b_q \in \Lambda^{\bar{w}}$. Taking into account that $w >_{\text{lex}} \bar{w}$ and that they both are decreasing words, there are only two possible cases:

- (i) w and \bar{w} share a common prefix $a_0 \cdots a_m$. Then $w = a_0 \cdots a_m b_{m+1} \cdots b_q$, and the word $a_0 \cdots a_m b_{m+1}$ is minimal by applying successive suffix-cut operations.
- (ii) w and \bar{w} do not share a common prefix. Necessarily $b_0 >_{\text{lex}} a_0$, and then the word b_0 is minimal by applying several suffix-cut operations.

From the above, we deduce a method for constructing all minimal word of $\Lambda^{\bar{w}}$. First, we partition \mathcal{A} into disjoint – potentially empty – subsets:

$$\mathcal{A}^0 = \{a \in \mathcal{A} : a >_{\text{lex}} a_0\},$$

$$\mathcal{A}^i = \{a \in \mathcal{A} : a_{i-1} \geq_{\text{lex}} a >_{\text{lex}} a_i \quad \forall i \in \{1, \dots, p\},$$

$$\mathcal{A}^{p+1} = \{a \in \mathcal{A} : a_p \geq_{\text{lex}} a\}.$$

It then follows that – empty \mathcal{A}^i 's not being considered,

- $\forall b \in \mathcal{A}^0$, the word b is minimal,
- $\forall b \in \mathcal{A}^i$ with $i \in \{1, \dots, p\}$, the word $a_0 \cdots a_{i-1} b$ is minimal,
- $\forall b \in \mathcal{A}^{p+1}$, the word $\bar{w} b$ is minimal.

As we partitioned \mathcal{A} , we have proved the following proposition.

Proposition 2.3. The number of minimal words of $\Lambda^{\bar{w}}$ is exactly $\#\mathcal{A}$.

As a follow-up of the example some lines ago, with $\mathcal{A} = \{0, 1, 2, 3\}$ and $\bar{w} = 211$, we apply the proposed method to find the minimal elements of $\Lambda^{\bar{w}}$. We partition \mathcal{A} into: $\mathcal{A}^0 = \{3\}$, $\mathcal{A}^1 = \{2\}$, $\mathcal{A}^2 = \emptyset$, $\mathcal{A}^3 = \{0, 1\}$. The four minimal words are therefore 3, 22, 2111 and 2110.

Although the previous result is completely general, if we require that $\mathcal{A} = \{0, \dots, n\}$, then the partition method described above can be rewritten into Algorithm 2. While this is not included in the pseudocode provided, note that the algorithm should return an empty list if $a_0 > n$, as in this case there would be no minimal word to look for.

Algorithm 2: MINIMALWORDS

Input: $\bar{w} = a_0 \cdots a_p$, $\mathcal{A} = \{0, \dots, n\}$

Output: All minimal words of $\Lambda^{\bar{w}}$

```

1 Set L to the empty list;
2 if  $a_0 < n$  then
3   for  $i \in \{a_0 + 1, \dots, n\}$  do
4     Add the word  $i$  to L
5 for  $k \in \{1, \dots, p\}$  do
6   if  $a_k < a_{k-1}$  then
7     for  $i \in \{a_k + 1, \dots, a_{k-1}\}$  do
8       Add the word  $a_0 \cdots a_{k-1} i$ 
        to L
9 for  $i \in \{0, \dots, a_p\}$  do
10  Add the word  $a_0 \cdots a_p i$  to L
11 return L

```

2.2 Canonical FDAG

FDAG are unordered objects, like the trees they compress, and therefore their enumeration requires to reflect this nature. In practice, finding a systematic way to order them makes it possible to design a simpler reduction rule, as done for trees [18], ignoring the combinatorics of permutations. The purpose of this subsection is to provide a unique way to order FDAG. We show that such an order exists in Theorem 2.4, unambiguously characterizing FDAG. The approach chosen is based on the notion of topological order.

Topological ordering Let D be a directed graph, where multiple arcs are allowed. A topological ordering on D is an ordering of the vertices of D such that for every arc uv from vertex u to vertex v , u comes after v in the ordering. Formally, $\psi : D \rightarrow \llbracket 0, \#D - 1 \rrbracket$ is a topological ordering if and only if ψ is bijective and $\psi(u) > \psi(v)$ for all $u, v \in D$ such that there exists at least one arc uv in D . A well known result establishes that D is a DAG if and only if it admits a topological ordering [14]. Nonetheless, when a topological ordering exists, it is in general not unique – see Figure 2. A reverse search enumeration of topological orderings of a given DAG can actually be found in [2, Section 3.5].




Figure 2: The DAG on the left admits five topological orderings, which are shown in the table.

Constrained topological ordering We aim to reduce the number of possible topological orderings of a DAG by constraining them. Let D be a DAG and ψ a topological ordering. Taking advantage of the vertical hierarchy of DAGs, our first constraint is

$$\forall (u, v) \in D^2, \quad \mathcal{H}(u) > \mathcal{H}(v) \implies \psi(u) > \psi(v). \quad (3)$$

Applying (3) to the topological orderings presented in Figure 2, ψ_5 must be removed, as $\psi_5(\text{light blue}) > \psi_5(\text{red})$ and $\mathcal{H}(\text{red}) > \mathcal{H}(\text{light blue})$.

For any vertex v , and any $u \in \mathcal{C}(v)$, by definition, $\mathcal{H}(v) > \mathcal{H}(u)$. Therefore, there can be no arcs between vertices at same height. Any arbitrary order on them leads to a different topological ordering. The next constraint we propose relies on the lexicographical order:

$$\forall (u, v) \in D^2, \quad \mathcal{H}(u) = \mathcal{H}(v) \text{ and } \mathcal{C}_\psi(u) >_{\text{lex}} \mathcal{C}_\psi(v) \implies \psi(u) > \psi(v), \quad (4)$$

where $\mathcal{C}_\psi(v)$ is the list $[\psi(v_i) : v_i \in \mathcal{C}(v)]$ sorted by decreasing order w.r.t. the lexicographical order. In other words, $\mathcal{C}_\psi(v)$ is a decreasing word – see Definition 2.1 – on the alphabet $\mathcal{A} = \llbracket 0, \#D - 1 \rrbracket$. Table 1 illustrates the behavior of (4) on the followed example of Figure 2.

The combination of those two constraints imposes uniqueness in all cases except when there exists $(u, v) \in D^2$ such that $\mathcal{C}_\psi(u) = \mathcal{C}_\psi(v)$ and $u \neq v$. It should be clear that if we impose the upcoming condition (5), such a pathological case can not occur.

$$\forall (u, v) \in D^2, \quad u \neq v \implies \mathcal{C}(u) \neq \mathcal{C}(v) \quad (5)$$

Upcoming Theorem 2.4 establishes that a DAG compresses a forest if and only if the topological order constrained by (3) and (4) is unique. In other words, an unambiguous characterization of FDAG is exhibited.

\circ	\bullet	$\color{magenta}\bullet$	(4)
$\mathcal{C}_{\psi_1}(\circ)$	11	10	\checkmark
$\mathcal{C}_{\psi_2}(\circ)$	00	10	\times
$\mathcal{C}_{\psi_3}(\circ)$	00	10	\checkmark
$\mathcal{C}_{\psi_4}(\circ)$	11	10	\times

Table 1: Application of (4) to the remaining topological orderings of Figure 2 that satisfy (3). As $\mathcal{C}_{\psi}(\bullet) = \mathcal{C}_{\psi}(\color{magenta}\bullet)$, we only need to consider vertices \bullet and $\color{magenta}\bullet$. As $\psi_i(\bullet) > \psi_i(\color{magenta}\bullet) \iff i \in \{1, 2\}$, the only orderings that are kept are ψ_1 and ψ_3 .

Theorem 2.4. *The following statements are equivalent:*

- (i) D fulfills (5),
- (ii) there exists a unique topological ordering ψ of D that satisfies both constraints (3) and (4),
- (iii) there exists a unique forest $F \in \mathcal{F}$ – cf. (2) – such that $D = \mathfrak{R}(F)$,

where \mathfrak{R} is the DAG reduction operation defined in Subsection 1.2.

Proof. (i) \iff (ii) follows from the above discussion. (iii) \implies (i) follows from the definition of \mathfrak{R} . Indeed, if there was two distinct vertices $(u, v) \in D^2$ with the same multiset of children, they would have been compressed as a unique vertex in the reduction. We now prove that (i) \implies (iii).

In the first place, if D fulfills (5), then D must admit a unique leaf, denoted $\mathcal{L}(D)$. Indeed, if there were two leaves l_1 and l_2 , we would have $\mathcal{H}(l_1) = \mathcal{H}(l_2) = 0$ but also $\mathcal{C}(l_1) = \mathcal{C}(l_2) = \emptyset$, which would violate (5). Let r_1, \dots, r_k be the vertices in D that have no parent. We define D_1, \dots, D_k as the subDAGs rooted respectively in r_1, \dots, r_k . Then, we define $T_i = \mathfrak{R}^{-1}(D_i)$ and $F = \{T_1, \dots, T_k\}$. The T_i 's are well defined as all vertices in D (consequently in D_i) have a different multiset of children, and therefore compress distinct subtrees – i.e. F fulfills (2), therefore $F \in \mathcal{F}$. Moreover, $D = \mathfrak{R}(F)$. $\color{red}\spadesuit$

In the sequel of the article, we shall only consider FDAG. Consequently, from Proposition 2.4, they admit a unique topological ordering ψ satisfying both constraints (3) and (4), called *canonical ordering*. Thus, for any FDAG D , the associated canonical ordering ψ will be implicitly defined. The vertices will be numbered accordingly to their ordering, i.e. $D = (v_0, \dots, v_n)$ with $\psi(v_i) = i$. Finally, as a consequence of constraints (3) and (4), note that D can be partitioned in subsets of vertices with same height, each of them containing only consecutive numbered vertices. Figure 3 provides an example of a FDAG and its canonical ordering.




Figure 3: A FDAG D (left) and its canonical ordering ψ (right). Vertices that are at the same height are enclosed in the table between the dashed lines. Red arrows indicates the roots of the trees of the forest that is compressed by D .

2.3 Expansion rules

Reverse search techniques implies finding reduction rules, and then inverse them. Equally, we will define instead three expansion rules, of which inverse will be reduction rules. An expansion rule takes a FDAG and create a new DAG, that is “expanded” in the sense of having either more vertices or more arcs. Expansion rules are built so that such a new DAG is still a FDAG, as it will be proved in the next Subsection 2.4. We begin with a preliminary definition.

Definition 2.5. Let D be a FDAG, with $D = (v_0, \dots, v_n)$. We define the two following alphabets

$$\begin{aligned}\mathcal{A}_= &= \{\psi(v) : v \in D, \mathcal{H}(v) = \mathcal{H}(v_n)\} = \{p+1, \dots, n\}, \\ \mathcal{A}_< &= \{\psi(v) : v \in D, \mathcal{H}(v) < \mathcal{H}(v_n)\} = \{0, \dots, p\},\end{aligned}$$

where $p \in \llbracket 0, n-1 \rrbracket$ and $\psi(\cdot)$ is the canonical ordering of D .

In other words, $\mathcal{A}_=$ contains the indices of all vertices that have the same height as the vertex with the highest index according to ψ , and $\mathcal{A}_<$ the indices of all vertices that have an inferior height. The FDAG presented in Figure 3 will serve as a guideline example all along this subsection. Here, we have $\mathcal{A}_= = \{4, 5\}$ and $\mathcal{A}_< = \{0, 1, 2, 3\}$.

The three expansion rules are now introduced. Let $D = (v_0, \dots, v_n)$. Each of these rules is associated with an explicit symbol, which may be used, when necessary, to designate the rule afterward. It is worth noting that all of these rules will operate according to the vertex of highest index, v_n .

Branching rule (\curvearrowright) This rule adds an arc between v_n and a vertex below. The end vertex of the new arc is chosen such that $\mathcal{C}_\psi(v_n)$ remains a decreasing word. In Figure 4, (\curvearrowright) is applied on our guideline example.

Definition 2.6. (\curvearrowright) Let $\mathcal{C}_\psi(v_n) = a_0 \dots a_m$. Choose $a_{m+1} \in \mathcal{A}_<$ such that $a_m \geq_{\text{lex}} a_{m+1}$ and add an arc between $\psi^{-1}(a_{m+1})$ and v_n .



Figure 4: Branching rule applied to the FDAG of Figure 3. As $\mathcal{C}_\psi(v_5) = 211$, the only letters a we can pick from $\mathcal{A}_< = \{0, 1, 2, 3\}$, satisfying $a \leq_{\text{lex}} 1$, are 0 and 1. The only two possible outcomes of (\curvearrowright) are the words (a) 2111 and (b) 2110.

Elongation rule (\uparrow) This rule adds a new vertex v_{n+1} such that $\mathcal{H}(v_{n+1}) = \mathcal{H}(v_n) + 1$. Consequently, the alphabets change and become $\mathcal{A}_= = \{n+1\}$ and $\mathcal{A}_< = \{0, \dots, n\}$. Note that after using this rule, it is not possible to ever add a new vertex at height $\mathcal{H}(v_n)$. See Figure 5 for an illustration of this rule on the guideline example.

Definition 2.7. (\uparrow) Add new vertex v_{n+1} such that $\mathcal{C}_\psi(v_{n+1}) = a_0 \in \mathcal{A}_=$.

Widening rule (\curvearrowleft) This rule adds a new vertex v_{n+1} at height $\mathcal{H}(v_n)$. The vertex is added with children that respects the canonicalness of the DAG, that is, such that $\mathcal{C}_\psi(v_{n+1}) >_{\text{lex}} \mathcal{C}_\psi(v_n)$ – as in condition (4). In other terms, denoting $\Lambda_<$ the language of decreasing words on alphabet $\mathcal{A}_<$, and with $\bar{w} = \mathcal{C}_\psi(v_n)$, $\mathcal{C}_\psi(v_{n+1})$ must be chosen in $\Lambda_<^{\bar{w}}$ – see Definition 2.2. However, this set is infinite, so we restrict $\mathcal{C}_\psi(v_{n+1})$ to be chosen among the minimal words of $\Lambda_<^{\bar{w}}$. It follows from the definition of suffix-cut operator $\text{SC}(\cdot)$ that, by inverting the said operator, the other words in $\Lambda_<^{\bar{w}}$ can be obtained by performing repeated (\curvearrowright) operations. Finally, this new vertex is added to $\mathcal{A}_=$.




Figure 5: Elongation rule applied to the FDAG of Figure 3. As $\mathcal{A}_= = \{4, 5\}$, there are only two choices leading to (a) $\mathcal{C}_\psi(v_6) = 4$ and (b) $\mathcal{C}_\psi(v_6) = 5$. The alphabets become $\mathcal{A}_< = \{0, \dots, 5\}$ and $\mathcal{A}_= = \{6\}$.

Definition 2.8. (\mathcal{V}) Add new vertex v_{n+1} such that

$$\mathcal{C}_\psi(v_{n+1}) \in \{w \in \Lambda_{<}^{\bar{w}} : w \text{ is a minimal word of } \Lambda_{<}^{\bar{w}}\}$$

with $\bar{w} = \mathcal{C}_\psi(v_n)$.

From Proposition 2.3 we now that such minimal words exist. We prove in the upcoming lemma that, as claimed, $\mathcal{H}(v_{n+1}) = \mathcal{H}(v_n)$.

Lemma 2.9. Any element of $\Lambda_{<}^{\bar{w}}$ defines a new vertex v_{n+1} such that $\mathcal{H}(v_{n+1}) = \mathcal{H}(v_n)$.

Proof. From the definition of $\mathcal{H}(\cdot)$ – (1), it suffices to prove that v_{n+1} admits at least one child at height $h = \mathcal{H}(v_n) - 1$. Let us denote b_0 and a_0 the first letter of, respectively, $\mathcal{C}_\psi(v_{n+1})$ and $\mathcal{C}_\psi(v_n)$. Denoting $v = \psi^{-1}(b_0)$ and $u = \psi^{-1}(a_0)$, we already know that $\mathcal{H}(u) = h - 1$ – as ψ respects (4) and $\mathcal{C}_\psi(v_n)$ is a decreasing word. Therefore, as by construction $\mathcal{C}_\psi(v_{n+1}) >_{\text{lex}} \mathcal{C}_\psi(v_n)$, either (i) $b_0 = a_0$ and therefore $v = u$, either (ii) $b_0 >_{\text{lex}} a_0$. In the latter, as ψ respects (3) and (4), $\mathcal{H}(v) \geq \mathcal{H}(u) = h - 1$. But, as $b_0 \in \mathcal{A}_<$, $\mathcal{H}(v) < \mathcal{H}(v_n) = h + 1$. In both cases, $\mathcal{H}(v) = h - 1$. \spadesuit

Figure 6 illustrates the use of the widening rule on the followed example. It should be noted that the possible outcomes of (\mathcal{V}) are obtained by using Algorithm 2, applied with $w = \mathcal{C}_\psi(v_n)$ and $p =$ with $\mathcal{A}_< = \{0, \dots, p\}$.




Figure 6: We apply (\mathcal{V}) to the FDAG of Figure 3. Here, $\mathcal{A}_< = \{0, 1, 2, 3\}$ and $\bar{w} = 211$. As seen in Subsection 2.1, the minimal words of $\Lambda_{<}^{\bar{w}}$ are 3, 22, 2111 and 2110. Therefore, there are 4 ways to add a new vertex v_6 via the widening rule, that are such that (a) $\mathcal{C}_\psi(v_6) = 3$, (b) $\mathcal{C}_\psi(v_6) = 22$, (c) $\mathcal{C}_\psi(v_6) = 2111$ or (d) $\mathcal{C}_\psi(v_6) = 2110$. Finally, we update $\mathcal{A}_=$ to be equal to $\{4, 5, 6\}$.

2.4 Discussion on the enumeration tree of FDAG

In this subsection, we prove that the expansion rules previously defined preserve the canonicalness property during the enumeration. We also introduce the reduction rule itself, and investigate some properties of the enumeration tree implied by the combination of our rules and the reverse search technique. First, we show that FDAG are in bijection with a set of particular matrices, whose combinatorial properties are known and give us access to an asymptotic expansion of the enumeration tree growth. Then, we prove that any FDAG has a linear number of successors¹ in this tree, which shows that its growth, although exponential, is reasonably controlled.

Proposition 2.10. *The expansion rules preserve the canonicalness property.*

Proof. Let $D = (v_0, \dots, v_n)$ be a FDAG. The proposition follows naturally from the definitions:

- (\curvearrowright) Let a be the letter added to $w = \mathcal{C}_\psi(v_n)$. As $wa >_{\text{lex}} w >_{\text{lex}} \mathcal{C}_\psi(v_{n-1})$, the ordering is unchanged.
- (\uparrow) The new vertex v_{n+1} is such that $\mathcal{H}(v_{n+1}) > \mathcal{H}(v_n)$, so condition (3) is still met.
- (\curvearrowleft) The new vertex v_{n+1} is chosen so that $\mathcal{H}(v_{n+1}) = \mathcal{H}(v_n)$ and $\mathcal{C}_\psi(v_{n+1}) >_{\text{lex}} \mathcal{C}_\psi(v_n)$, so condition (4) is also still met.

Therefore, any DAG obtained from D is still a FDAG. ✂

The reduction rule associated to the three expansion rule is presented in Algorithm 3. As mentioned in Subsection 1.1, through Algorithm 1 – here with $g(\cdot) = \top$, this reduction rule defines an enumeration tree, whose root is set to be the DAG with a unique leaf and no arcs. This root is denoted D_0 . If we denote E_k the set of all FDAG that are accessible from D_0 in exactly k steps in the tree – with $E_0 = \{D_0\}$, then Table 2 depicts the values of $\#E_k$ for the first nine values of k ².

k	0	1	2	3	4	5	6	7	8
$\#E_k$	1	1	3	12	61	380	2 815	24 213	237 348

Table 2: Number of DAGs accessible from D_0 in k steps in the enumeration tree.

Actually, the terms of Table 2 coincide with the first terms of OEIS sequence A158691³, which counts the number of *row-Fishburn matrices*, that are upper-triangular matrices with at least one nonzero entry in each row. The *size* of such a matrix is equal to the sum of its entries.

¹“successor” in the sense of “children in the enumeration tree”. We make the distinction to avoid confusion with the children denoted with $\mathcal{C}(\cdot)$.

²These numbers were obtained numerically (cf. “Implementation” at the end of the article).

³OEIS Foundation Inc. (2019), The On-Line Encyclopedia of Integer Sequences, <http://oeis.org/A158691>.

Algorithm 3: REDUCTIONRULE


Input: $D = (v_0, \dots, v_n); w = \mathcal{C}_\psi(v_n);$
 $w' = \mathcal{C}_\psi(v_{n-1})$

- 1 **if** v_n **is the only vertex of height** $\mathcal{H}(v_n)$ **then**
- 2 **if** $\#w = 1$ **then**
- 3 $(\uparrow)^{-1}$: Delete vertex v_n ;
- 4 **else**
- 5 $(\curvearrowright)^{-1}$: $w \leftarrow \text{SC}(w)$;
- 6 **else**
- 7 **if** w **is a minimal word of** $\Lambda_{<}^{w'}$ **then**
- 8 $(\curvearrowleft)^{-1}$: Delete vertex v_n ;
- 9 **else**
- 10 $(\curvearrowright)^{-1}$: $w \leftarrow \text{SC}(w)$;

$\text{SC}(\cdot)$ is the suffix-cut operator defined in Subsection 2.1.

Theorem 2.11. *There exists a bijection Φ between the set of FDAG and the set of row-Fishburn matrices, such that if D is a FDAG and $M = \Phi(D)$, then*

$$D \in E_k \iff \text{size}(M) = k.$$

Proof. The proof lies in Appendix A. 

This connection is to our advantage since Fishburn matrices (in general) are combinatorial objects widely explored in the literature as they are in bijection with many others – see [12, Section 2] for a general overview. Notably, the asymptotic expansion of the number of row-Fishburn matrices has been conjectured first by Jelínek [13] and then proved by Bringmann et al. [7].

Proposition 2.12 (Jelínek, Bringmann et al.). *As $k \rightarrow \infty$,*

$$\#E_k = k! \left(\frac{12}{\pi^2} \right)^k \left(\beta + O\left(\frac{1}{k}\right) \right)$$


with $\beta = \frac{6\sqrt{2}}{\pi^2} e^{\pi^2/24} = 1.29706861206\dots$

Given the overall combinatorics of FDAG, it is no surprise that the enumeration tree grows extremely fast. However, despite this combinatorial explosion, the growth of the tree is controlled, as proven by the upcoming Theorem 2.13.

Theorem 2.13. *Any FDAG $D = (v_0, \dots, v_n)$ has $\Theta(n)$ successors in the FDAG enumeration tree.*

Proof. We denote $\mathcal{C}_\psi(v_n) = a_0 \cdots a_m$. Depending on the rule chosen:

- (↔) a_{m+1} belongs to $\mathcal{A}_< = \{0, \dots, p\}$, so the maximum number of successors is at most $p + 1$, and at least 1, depending on the condition $a_m \geq_{\text{lex}} a_{m+1}$.
- (↑) The child of the new vertex is taken from $\mathcal{A}_= = \{p + 1, \dots, n\}$ so the number of successors is exactly $n - p$.
- (↘) Following Proposition 2.3, the number of successors is exactly $\#\mathcal{A}_< = p + 1$.

Combining everything, the number of successors is at least $n + 2$ and at most $n + p + 2 \leq 2n + 1$ (as $p \leq n - 1$, with equality for FDAG obtained just after using (↑) rule). 

In the previous proof, we have shown that the number of successors of a FDAG with n vertices is between $n + 1$ and $2n - 1$. Figure 7 illustrates that these boundaries are tight, on 1 000 randomly generated FDAG. The random generation of a FDAG is as follows. Starting from D_0 – the root, construct iteratively D_i as a successor of D_{i-1} in the enumeration tree, picked at random. We stop this random exploration at step k , and keep D_k . In Figure 7, we have generated 10 trajectories for each $k \in \{1, \dots, 100\}$.

To conclude, a fraction of the enumeration tree is shown in Figure 8, illustrating the path from the root D_0 to the FDAG of Figure 3. Unexplored branches are ignored, but are still indicated by their root.

2.5 Constraining the enumeration

In [18], the authors propose an algorithm to enumerate all trees with at most n vertices. They simply check whether the current tree has n vertices or not, and as their expansion rule adds one vertex at a time, they decide to cut a branch in the enumeration tree once they have reached n vertices. Similarly, adding a vertex to a tree can only increase its height or outdegree, so we can proceed in the same way to enumerate all trees with maximal height H and maximal outdegree d . Indeed, the number of trees satisfying those constraints is finite [4, Appendix D.2].

This property also holds in our approach : following one of the three expansion rules, we can only increase the height, outdegree or number of vertices of the FDAG. So, it makes sense to define similar constraints on the enumeration. However, for this constrained enumeration to generate a finite number of FDAG, constraints must be chosen wisely, as shown in the following proposition.

Proposition 2.14. *The FDAG enumeration tree is finite if at least one of those set of constraints is chosen:*

- (i) *maximum number of vertices n and maximum outdegree d ,*
- (ii) *maximum height H and maximum outdegree d .*

Proof. As (\curvearrowright) allows to add arcs indefinitely without changing the numbers of vertices, constraining on the maximum outdegree is mandatory in both cases. As the two others rules add vertices, constraining by the number of vertices leads to a finite enumeration tree – (i) is proved. To conclude, we only need to prove that (\curvearrowright) cannot be repeated an infinite number of times, i.e. there is only a finite number of new vertices that can be added at a given height, up to the maximum outdegree. This is achieved by virtue of the upcoming lemma.

Let $H > 2$ and $d \geq 1$. Let D be the FDAG constructed so that for each $0 \leq h \leq H$, D has the maximum possible number n_h of vertices of height h and with maximum outdegree d . Initial values are $n_0 = 1$ and $n_1 = d$.

Lemma 2.15. $\forall 2 \leq h \leq H, \quad n_h = \sum_{k=1}^d \binom{k + n_{h-1} - 1}{k} \binom{d - k + n_0 + \dots + n_{h-2}}{d - k}$.

Let $h \geq 2$ be fixed. To lighten the notation, let $n = n_{h-1}$ and $m = n_0 + \dots + n_{h-2}$. Let v be a vertex to be added at height h . For any vertex v_i at height $h - 1$, let x_i be the multiplicity of v_i in $\mathcal{C}(v)$ –

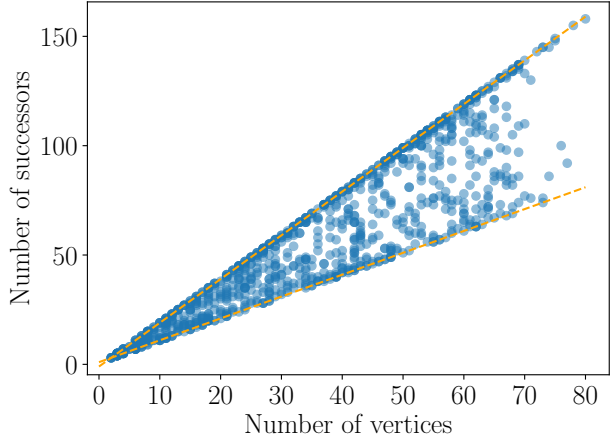


Figure 7: Numbers of successors of 1 000 random FDAG in the enumeration tree, according to their number of vertices. Orange lines have equations $y = n + 1$ and $y = 2n - 1$.

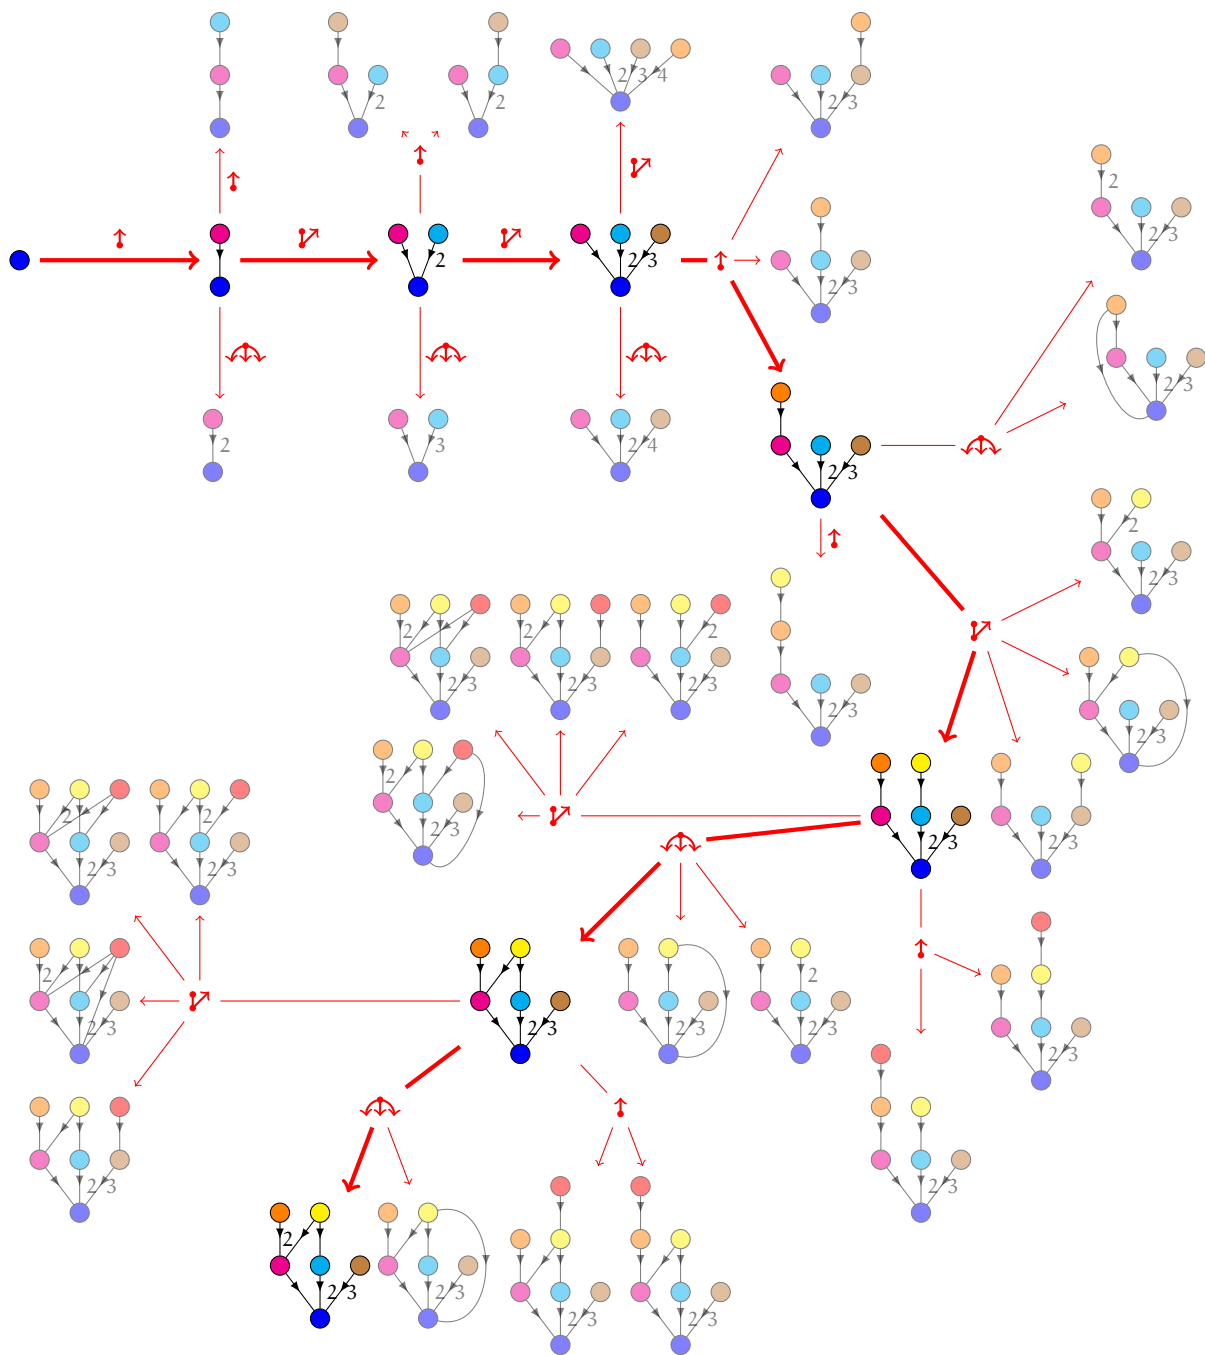



Figure 8: The path (in bold) in the FDAG enumeration tree leading to the FDAG of Figure 3. The unexplored branches are only displayed by their root, which are shown partially transparent. The order of insertion of the vertices of each FDAG is always the same, and follows the color code (in the order of insertion): \bullet , \bullet , \bullet , \bullet , \bullet , \bullet and \bullet . With respect to the canonical ordering, they are numbered 0 to 6 in the same order.

0 if $v_i \notin \mathcal{C}(v)$. Similarly, for any vertex v_h with $\mathcal{H}(v_j) \leq h - 2$, y_j is the multiplicity of v_j in $\mathcal{C}(v)$ – possibly 0. By definition of $\mathcal{H}(\cdot)$ – see (1), at least one x_i is non-zero. Therefore, there exist $k \in \llbracket 1, d \rrbracket$ such that:

$$\begin{aligned} x_1 + \cdots + x_n &= k \\ y_1 + \cdots + y_m &\leq d - k \end{aligned}$$

By virtue of the stars and bars theorem, for a fixed k , there are $\binom{k+n-1}{k}$ choices for variables x_i , and $\binom{d-k+m}{d-k}$ for variables y_j . Summing upon all values for k proves the claim. 

Remark 2.16. *In the constrained enumeration proposed in [18], all the trees with n vertices are the leaves of the enumeration tree. To get all trees with $n + 1$ vertices, it suffices to add to the enumeration all children of these leaves, i.e. trees obtained by adding a single vertex to them. This property – moving from one parameter value to the next by enumerating just one step further – does not hold anymore as soon as our set of constraints involve the maximum outdegree d , both for trees and FDAG. For instance, from a FDAG of height H , one can obtain FDAG of height $H + 1$ by using (\uparrow) once and repeating (\curvearrowright) up to $d - 1$ times.*

3 Enumeration of forests of subtrees

Once the reverse search scheme has been set up to enumerate a certain type of structure, it is natural to move to a finer scale by using the same scheme to enumerate substructures. However, the notion of “substructure” is not obvious to derive from the main structure, as several choices are possible – e.g. for trees one can think of subtrees [22, 5], subset trees [8], etc. From a practical point of view, the enumeration of substructures permits to solve the frequent pattern mining problem – which will be tackled in Section 4.

In this section we define forests of subtrees, which will be our substructures. Compressed as FDAG, these objects will be called subFDAG. We then address the problem of enumerating all subFDAG appearing in an FDAG D – similar as the one of enumerating all subtrees of a tree.


Forests of subtrees Similarly to forest being tuple of trees, *forests of subtrees* are tuple of subtrees, satisfying (2). Formally:

Definition 3.1. *Let F and f be two forests. f is a forest of subtrees of F if and only if*

$$\forall t \in f, \exists T \in F, t \in \mathcal{S}(T).$$

Forest of subtrees can be directly constructed from FDAG, as shown by the upcoming proposition. Let D be a FDAG, and V a subset of vertices of D .

Proposition 3.2. *If $\forall v \in V, \mathcal{C}(v) \subseteq V$, then V defines a FDAG Δ , such that $\mathfrak{R}^{-1}(\Delta)$ is a forest of subtrees of $\mathfrak{R}^{-1}(D)$.*

Proof. The demonstration is in two steps. (i) Remove from D the vertices that does not belong to V ; as there are no arcs that leave V by hypothesis, end up with a FDAG. Let us call Δ this FDAG. (ii) Let ρ be a root of Δ . By construction, ρ is also a vertex in D . Among all roots of D , there exists a root r such that $\rho \in \mathcal{D}(r)$. Therefore, $D[\rho]$ is a subDAG of $D[r]$, and then $t = \mathfrak{R}^{-1}(D[\rho])$ is a subtree of 

$T = \mathfrak{R}^{-1}(D[r])$ – with $T \in F = \mathfrak{R}^{-1}(D)$. As $\Delta[\rho]$ and $D[\rho]$ are isomorphic, $t \in f = \mathfrak{R}^{-1}(\Delta)$. Therefore we have proved that $\forall t \in f, \exists T \in F, t \in \mathcal{S}(T)$. ✍

We say that the FDAG Δ is a *subFDAG*⁴ of D . Figure 9 provides an example of such a construction.




Figure 9: Construction of a forest of subtrees from FDAG. (a) A FDAG D . The set V is circled in red. (b) The FDAG Δ (c) The forest f compressed by Δ . (d) The forest F compressed by D . One can spot that $t_1 \in \mathcal{S}(T_1)$, $t_2 \in \mathcal{S}(T_2)$ so f is a forest of subtrees of F , and Δ a subFDAG of D .

Enumeration of subFDAG We now solve the following enumeration problem: given a forest F , find all forests of subtrees of F . Equally, given a FDAG D , find all subFDAG of D . To address this, we make extensive use of the reverse search technique, adapting the one presented in Section 2.

Since a subFDAG is also a FDAG, it admits successors in the enumeration tree defined in Section 2. We are interested in those of these successors who are also subFDAG (if any). In fact, since a subFDAG can be defined from a set of vertices, all one has to do is determine which new vertex can be chosen to expand an existing subFDAG – corresponding to a (\uparrow) or (\uparrow') step. The covering of all added new arcs is implicit in this construction and corresponds to some steps of (\swarrow) .

Let Δ be a subFDAG of D and v its last inserted vertex – it is also the vertex with the largest ordering number in D . We denote by $S(\Delta)$ the set of all vertices $v' \in D$ that can be added to Δ to expand it to a new subFDAG. Let us call $S(\Delta)$ the *set of candidate vertices* of Δ . More precisely:

Lemma 3.3. $S(\Delta)$ is the set of vertices $v' \in D$ that satisfies both:

- (i) $\mathcal{C}(v') \subseteq \Delta$
- (ii) $\psi(v') > \psi(v)$

where $\psi(\cdot)$ is the canonical ordering of D .

Proof. (i) This condition is necessary so that $\Delta' = \Delta \cup \{v'\}$ fulfill the requirements for Proposition 3.2. (ii) This condition is necessary so that Δ' remains a FDAG. As $\psi(v') > \psi(v)$, either $\mathcal{H}(v') = \mathcal{H}(v) + 1$ – then it is a (\uparrow) step – or $\mathcal{H}(v') = \mathcal{H}(v)$ and $\mathcal{C}_\psi(v') >_{\text{lex}} \mathcal{C}_\psi(v)$ – for a (\uparrow') step. ✍

When $S(\Delta)$ is not empty, picking $s \in S(\Delta)$ ensure that $\Delta' = \Delta \cup \{s\}$ is a subFDAG of D . With respect to the enumeration tree of Section 2, Δ is an ancestor of Δ' – but not necessarily its parent, since the steps of (\swarrow) are implicit. Δ' is called an *heir* of Δ . We can in turn calculate $S(\Delta')$, by updating $S(\Delta)$:

⁴Not to be confused with *subDAG*, introduced in Subsection 1.2.

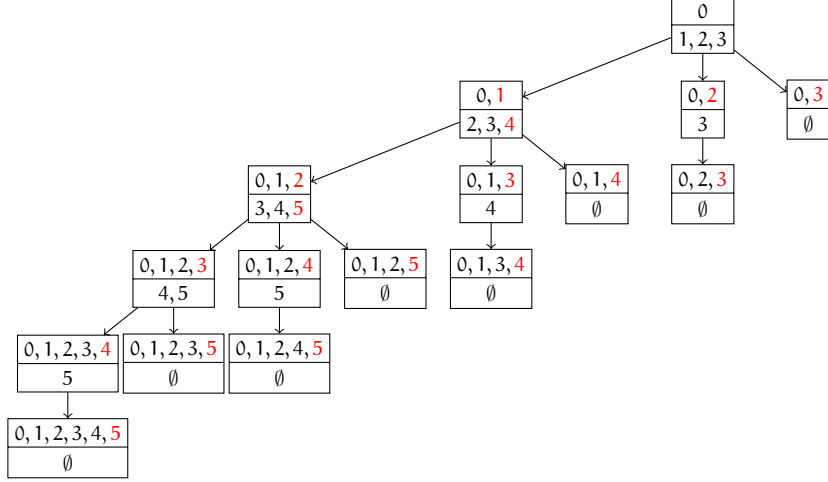


Figure 10: Enumeration tree of the subFDAG of the FDAG of Figure 3, using both Algorithm 1 and Algorithm 4. The indices of the vertices correspond to the canonical ordering defined in Figure 3. In each vertex, the upper part corresponds to the current subFDAG Δ whereas the lower part stands for the set $S(\Delta)$. Numbers in red indicate what changes for an heir compared to its parent.

(i) Remove from $S(\Delta)$ all vertices v' such that $\psi(s) > \psi(v')$. (ii) In D , look only after the vertices v' such that $s \in \mathcal{C}(v') \subseteq \Delta \cup \{s\}$ and add them to $S(\Delta')$.

If Δ' is an heir of Δ , then by removing the last inserted vertex of Δ' , one can retrieve Δ . This define a reduction rule f , and therefore an enumeration tree. Algorithm 4 is meant to construct the set $f^{-1}(\Delta)$. Applying Algorithm 1 together with it, and starting from $\Delta = \mathcal{L}(D)$ – in this case, $S(\Delta)$ is the set of parents of $\mathcal{L}(D)$ of height 1– permits to enumerate all subFDAG of D . Figure 10 provide an example by enumerating all subFDAG of the FDAG of Figure 3.

Algorithm 4: HEIRS

Input: $D, [\Delta, S(\Delta)]$

- 1 Set L to the empty list;
 - 2 **for** $s \in S(\Delta)$ **do**
 - 3 Let S' be a copy of $S(\Delta)$;
 - 4 $S' \leftarrow S' \setminus \{v' \in S' : \mathcal{C}_\psi(v') \leq_{\text{lex}} \mathcal{C}_\psi(s)\}$;
 - 5 $S' \leftarrow S' \cup \{v' \in D : s \in \mathcal{C}(v') \subseteq \Delta \cup \{s\}\}$;
 - 6 Add $[\Delta \cup \{s\}, S']$ to L ;
 - 7 **return** L
-

4 Frequent subFDAG mining problem

Using the reverse search formalism introduced in Subsection 1.1, the *frequent pattern mining problem* can be formulated as follow: from a dataset $X = \{s_1, \dots, s_n\}$ with $s_i \in \mathbb{S}$, and a fixed threshold σ , find all elements $s \in \mathbb{S}$ that satisfies $\text{freq}(s, X) \geq \sigma$, where $\text{freq}(\cdot)$ is a function, to be defined, that counts the frequency of appearance of s in the dataset X . This problem can be solved using Algorithm 1 with $g(s, X, \sigma) = (\text{freq}(s, X) \geq \sigma)$. The function g is trivially anti-monotone.

We emphasize here that each possible definition of “ s appears in X ” leads to a different data mining problem. The choice of this definition is therefore of prime importance. In particular, this choice should induce a way of calculating $\text{freq}(s, X)$ that reflects the specificity of the chosen reduction rule f , so that $\{s \in f^{-1}(s_0) | g(s) = \top\}$ can be constructed directly, instead of first generating $f^{-1}(s_0)$

and then filtering according to the value of g . Indeed, if g is too restrictive, and $f^{-1}(s_0)$ too large, one would have to enumerate objects that are not relevant to the problem, which is not desirable.

In this article, the problem we consider is the following: given a set of trees $X = \{T_1, \dots, T_n\}$, account for forests of subtrees that appear simultaneously in different T_i 's. In other words, if we denote \mathcal{F}_i the set of all forest of subtrees appearing in the forest formed by $\{T_i\}$, we are interested in the study of $\bigcap_{i \in I_\sigma} \mathcal{F}_i$ where $I_\sigma \subseteq \llbracket 1, n \rrbracket$, such that $\#I_\sigma \geq \sigma \cdot n$.

A first, naive strategy would be to first build the \mathcal{F}_i 's, e.g. by using Algorithm 4 on $\mathfrak{R}(T_i)$, and then construct $\bigcap_{i \in I_\sigma} \mathcal{F}_i$ for all possible choices of I_σ . Obviously, this approach has its weaknesses: (i) many subFDAG will be enumerated for nothing or in several copies, and (ii) it does not take into account that X is itself a forest. Our aim is to propose a variant of Algorithm 4 that, applied to $\mathfrak{R}(X)$, would enumerate only subFDAG appearing in the $\mathfrak{R}(T_i)$'s with a large enough frequency.

Given a forest $F = \{T_1, \dots, T_n\}$ and its DAG compression $D = \mathfrak{R}(F)$, we have to retrieve, for each vertex in D , their origin in the dataset, that is, from which tree they come from. This issue has already been addressed in a previous article [5, Section 3.3], and has led to the concept of *origin*. For any vertex $v \in D$, the origin of v is defined as

$$o(v) = \{i \in \llbracket 1, n \rrbracket : \mathfrak{R}^{-1}(D[v]) \in \mathcal{S}(T_i)\}.$$

In others words, $o(v)$ represents the set of trees for which $\mathfrak{R}^{-1}(D[v])$ is a subtree. We state in [5, Proposition 3.4] that origins can be iteratively computed in one exploration of D . The proof lies in the property that if $i \in o(v)$, then for all $v' \in \mathcal{D}(v)$, $i \in o(v')$ – as $\mathfrak{R}^{-1}(D[v']) \in \mathcal{S}(\mathfrak{R}^{-1}(D[v]))$.

Let Δ be a subFDAG of D . For Δ to compress a forest of subtrees of a tree T_i , it is necessary that $i \in o(v)$ for all $v \in V_0$. Therefore, the set of trees for which Δ compress a forest of subtrees – the *origin* of Δ , denoted $o(\Delta)$ – is equal to

$$o(\Delta) = \bigcap_{v \in V_0} o(v).$$

If $\Delta' = \Delta \cup \{s\}$ is an heir of Δ – as defined earlier, then $o(\Delta') = o(\Delta) \cap o(s)$. Algorithm 4 can therefore be refined so that Δ' should be ignored if $o(\Delta') = \emptyset$ – as Δ' does not anymore compresses any forest of subtrees actually present in the trees of F .

So far we neglected the threshold σ . We only want to keep subFDAG that appear in at least $\sigma\%$ of the data. If $\#o(\Delta)/\#F < \sigma$, then the successors of Δ are not investigated.

Indeed, as $o(\cdot)$ is a decreasing function, successors of Δ cannot exceed the threshold again. We can finally introduce Algorithm 5 that solves the frequent subFDAG mining problem for trees. With the notations of Subsection 1.1, this algorithm builds the set $\{\Delta' \in f^{-1}(\Delta) \mid \text{freq}(\Delta', F) \geq \sigma\}$, with $\text{freq}(\Delta', F) = \#o(\Delta')/\#F$. The set is also built directly, without any posterior filtering, which is suitable as discussed in Subsection 1.1.

Algorithm 5: FREQUENTHEIRS

Input: $D = \mathfrak{R}(F)$, $[\Delta, S(\Delta), o(\Delta)]$, σ

- 1 Set L to the empty list;
- 2 **for** $s \in S(\Delta)$ **do**
- 3 **if** $o(\Delta) \cap o(s) \neq \emptyset$ **and** $\#(o(\Delta) \cap o(s)) \geq \sigma \cdot \#F$ **then**
- 4 Let S' be a copy of $S(\Delta)$;
- 5 $S' \leftarrow S' \setminus \{v' \in S' : \mathcal{C}_\psi(v') \leq_{\text{lex}} \mathcal{C}_\psi(s)\}$;
- 6 $S' \leftarrow S' \cup \{v' \in D : s \in \mathcal{C}(v') \subseteq D_0 \cup \{s\}\}$;
- 7 Add $[\Delta \cup \{s\}, S', o(\Delta) \cap o(s)]$ to L ;
- 8 **return** L

We stated earlier that we wanted to avoid generating unnecessary or multiple copies of subFDAG, which is achieved with Algorithm 5. We now empirically study what we have gained from this, by comparing the use of Algorithm 5 on $D = \mathfrak{R}(F)$, with the use of Algorithm 4 on each $\mathfrak{R}(T_i)$. As in Subsection 2.4, we generated 1 000 random DAGs, 10 repetitions for each $k \in \{1, \dots, 100\}$, creating D_k as a random exploration in k steps of the enumeration tree defined in Section 2. We assume $D_k = \mathfrak{R}(f)$ where $f = \{\mathfrak{R}^{-1}(D_k[r]) : r \text{ root of } D_k\}$. For each D_k , we have computed the quotient

$$Q(D_k) = \frac{\text{number of subFDAG of } D_k \text{ enumerated via Algorithm 5}}{\sum_{r \text{ root of } D_k} \text{number of subFDAG of } D_k[r] \text{ enumerated via Algorithm 4}}$$

with parameter $\sigma = 0$ when using Algorithm 5. The results are provided in Figure 11. Despite a rather marked variability, there is a general trend of decreasing as the number of vertices increases. We obtain fairly low quotients, around 20%, quite quickly. Given the combinatorial explosion of the objects to be enumerated, such an advantage is of the greatest interest.

Implementation

The `treex` library for Python [3] is designed to manipulate rooted trees, with a lot of diversity (ordered or not, labeled or not). It offers options for random generation, visualization, edit operations, conversions to other formats, and various algorithms. The enumeration of Section 2 and the algorithms of Section 3 and 4 have been implemented as a module of `treex` so that the interested reader can manipulate the concepts discussed in this paper in a ready-to-use manner. Installing instructions and the documentation of `treex` can be found from [3].




Figure 11: Quotient $Q(D)$ according to the number of vertices of D . Here 1 000 random FDAG are displayed.

Acknowledgments

This work has been supported by the European Union’s H2020 project ROMI. The authors would like to thank Dr. Arnaud Mary for his helpful suggestions on the final draft of the article.

References

- [1] Tatsuya Asai, Hiroki Arimura, Takeaki Uno, and Shin-Ichi Nakano. Discovering frequent substructures in large unordered trees. In *International Conference on Discovery Science*, pages 47–61. Springer, 2003.
- [2] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1-3):21–46, 1996.
- [3] Romain Azaïs, Guillaume Cerutti, Didier Gemmerlé, and Florian Ingels. `treex`: a python package for manipulating rooted trees. *The Journal of Open Source Software*, 4, 2019.

- [4] Romain Azaïs, Jean-Baptiste Durand, and Christophe Godin. Approximation of trees by self-nested trees. In *ALENEX 2019 - Algorithm Engineering and Experiments*, pages 1–24, San Diego, United States, January 2019. URL: <https://hal.archives-ouvertes.fr/hal-01294013>, doi:10.10860.
- [5] Romain Azaïs and Florian Ingels. The weight function in the subtree kernel is decisive. *arXiv preprint arXiv:1904.05421*, 2019.
- [6] Edward A Bender and S Gill Williamson. *Lists, Decisions and Graphs*. S. Gill Williamson, 2010.
- [7] Kathrin Bringmann, Yingkun Li, and Robert C Rhoades. Asymptotics for the number of row-fishburn matrices. *European Journal of Combinatorics*, 41:183–196, 2014.
- [8] Michael Collins and Nigel Duffy. Convolution kernels for natural language. In *Advances in neural information processing systems*, pages 625–632, 2002.
- [9] Christophe Godin and Pascal Ferraro. Quantifying the degree of self-nestedness of trees: application to the structural analysis of plants. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 7(4):688–703, 2010.
- [10] Jiawei Han, Hong Cheng, Dong Xin, and Xifeng Yan. Frequent pattern mining: current status and future directions. *Data mining and knowledge discovery*, 15(1):55–86, 2007.
- [11] John C. Hart and Thomas A. DeFanti. Efficient antialiased rendering of 3-d linear fractals. *SIGGRAPH Comput. Graph.*, 25(4):91–100, July 1991. URL: <http://doi.acm.org/10.1145/127719.122728>, doi:10.1145/127719.122728.
- [12] Hsien-Kuei Hwang and Emma Yu Jin. Asymptotics and statistics on fishburn matrices and their generalizations. *arXiv preprint arXiv:1911.06690*, 2019.
- [13] Vít Jelínek. Counting general and self-dual interval orders. *Journal of Combinatorial Theory, Series A*, 119(3):599–614, 2012.
- [14] Arthur B Kahn. Topological sorting of large networks. *Communications of the ACM*, 5(11):558–562, 1962.
- [15] Ailsa H Land and Alison G Doig. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer, 2010.
- [16] Fionn Murtagh. Counting dendrograms: a survey. *Discrete Applied Mathematics*, 7(2):191–199, 1984.
- [17] Shin-Ichi Nakano. Efficient generation of plane trees. *Information Processing Letters*, 84(3):167–172, 2002.
- [18] Shin-ichi Nakano and Takeaki Uno. Efficient generation of rooted trees. *National Institute for Informatics (Japan), Tech. Rep. NII-2003-005E*, 8, 2003.
- [19] Sebastian Nowozin. *Learning with structured data: applications to computer vision*. PhD thesis, Berlin Institute of Technology, 2009.
- [20] Benno Schwikowski and Ewald Speckenmeyer. On enumerating all minimal solutions of feedback problems. *Discrete Applied Mathematics*, 117(1-3):253–265, 2002.
- [21] Ivan E. Sutherland. Sketchpad: A man-machine graphical communication system. In *Proceedings of the May 21-23, 1963, Spring Joint Computer Conference, AFIPS '63 (Spring)*, pages 329–346, New York, NY, USA, 1963. ACM. URL: <http://doi.acm.org/10.1145/1461551.1461591>, doi:10.1145/1461551.1461591.
- [22] S.V.N. Vishwanathan and Alexander J Smola. Fast kernels on strings and trees. *Advances on Neural Information Processing Systems*, 14, 2002.
- [23] Kazuaki Yamazaki, Toshiki Saitoh, Masashi Kiyomi, and Ryuhei Uehara. Enumeration of nonisomorphic interval graphs and nonisomorphic permutation graphs. *Theoretical Computer Science*, 806:310–322, 2020.
- [24] Xifeng Yan and Jiawei Han. gspan: Graph-based substructure pattern mining. In *2002 IEEE International Conference on Data Mining, 2002. Proceedings.*, pages 721–724. IEEE, 2002.

A A bijection between FDAG and row-Fishburn matrices

This section is dedicated to the proof of Theorem 2.11, which is in two steps. First, we recall the natural bijection between FDAG and their adjacency matrices; the latter are then put into bijection with the row-Fishburn matrices.

FDAG \leftrightarrow Reduced adjacency matrix Let $D = (v_0, \dots, v_n)$ be a FDAG constructed in k steps from D_0 in the enumeration tree defined in Subsection 2.4. The adjacency matrix of D is defined as $A = (A_{i,j})_{i,j \in [n,0]^2}$ where, if m is the multiplicity of v_j in $\mathcal{C}(v_i)$, then $A_{i,j} = m$ – possibly 0 if $v_j \notin \mathcal{C}(v_i)$. By construction of D , as v_n is the last inserted vertex, it has no parents, so $A_{n,\cdot}$ is a column of zeros; and as v_0 is a leaf, it has no children, so $A_{\cdot,0}$ is a row of zeros. We define the *reduced* adjacency matrix M as the matrix A deprived of this column and this row. Therefore, $M = (A_{i,j})_{i \in [n,1], j \in [n-1,0]}$. As a vertex can not be a parent to any vertex introduced after it, we have $A_{i,j} = 0$ for all $i \leq j$ – so that M is an upper-triangular matrix. In addition, as all vertices except v_0 have at least one child, there is at least one non-zero entry in each row of M . Therefore, M is a row-Fishburn matrix. However, we have no guarantee that this matrix verifies $\text{size}(M) = k$.

Reduced adjacency matrix \rightarrow Incremental adjacency matrix Let $D = (v_0, \dots, v_n)$ be a FDAG, and M its reduced adjacency matrix. Let M_i be the row of M corresponding to $\mathcal{C}_\psi(v_i)$. The incremental adjacency matrix \hat{M} is defined as:

$$\begin{cases} \hat{M}_1 = M_1 \\ \hat{M}_{i+1} = M_{i+1} \ominus M_i \end{cases}$$

where the \ominus operation is defined as follow: given two rows $a_0 \cdots a_n$ and $b_0 \cdots b_n$, then denoting $j = \min\{i : a_i \neq b_i\}$, and $c = a_j - b_j$,

$$\begin{array}{r} a_0 \cdots a_{j-1} \quad a_j \quad a_{j+1} \cdots a_n \\ \ominus b_0 \cdots b_{j-1} \quad b_j \quad b_{j+1} \cdots b_n \\ \hline = 0 \cdots 0 \quad c \quad a_{j+1} \cdots a_n \end{array} .$$

We claim that this new matrix \hat{M} is a row-Fishburn matrix of size k , if $D \in E_k$. Actually, since M was already a row-Fishburn matrix, we just have to check that the size is correct. Let us consider v_i and v_{i+1} . The vertex v_{i+1} has been constructed from v_i by using either (\uparrow) or (\uparrow') , and potentially several (\uparrow') after that – let us say $p \geq 0$ times. Therefore, if the claim is correct, the sum over \hat{M}_{i+1} should be exactly $p + 1$. Consider the operation by which v_{i+1} was added in the first place:

- (\uparrow) $\mathcal{C}_\psi(v_{i+1})$ is reduced to a single element a , such that $a >_{\text{lex}} \mathcal{C}_\psi(v_i)$. Therefore, the index j of the first non-zero coefficient of M_{i+1} is ahead of the one of M_i so that the coefficient c of \ominus is equal to the j -th coefficient of M_{i+1} minus zero. Since the (\uparrow') rule adds children to respect decreasing words, the p extra coefficients are added to the right of the j -th coefficient (including it) and therefore they are kept unchanged in the \ominus operation. Eventually, the sum over M_{i+1} is $p + 1$ and so is the sum over \hat{M}_{i+1} .

(\heartsuit) $\mathcal{C}_\psi(v_{i+1})$ is built from $\mathcal{C}_\psi(v_i)$ with Algorithm 2, and therefore they (i) share a common prefix, possibly empty and (ii) then differ by a single letter. The index of that letter in M_{i+1} corresponds to the index j defined in \ominus . Therefore, the coefficient \mathbf{c} is – before any (\heartsuit) – equal to one. The argument of (\heartsuit) letters being added to the right of j still hold and therefore the sum over \widehat{M}_{i+1} is also $p + 1$.

To conclude the proof, we have to exhibit the inverse function of the mapping we just defined. This will prove that this mapping is indeed a bijection, and then the theorem holds.

Incremental adjacency matrix \rightarrow Reduced adjacency matrix Let M and \widehat{M} be constructed as before. From \widehat{M} , we can define a matrix M' as:

$$\begin{cases} M'_1 = \widehat{M}_1 \\ M'_{i+1} = M'_i \oplus \widehat{M}_{i+1} \end{cases}$$

where the \oplus operation is defined as follow: given two words $a_0 \cdots a_n$ and $b_0 \cdots b_n$, then denoting $j = \min\{i : b_i \neq 0\}$, and $c = a_j + b_j$,

$$\begin{array}{r} \begin{array}{|c|c|c|} \hline a_0 \cdots a_{j-1} & a_j & a_{j+1} \cdots a_n \\ \hline \end{array} \\ \oplus \begin{array}{|c|c|c|} \hline b_0 \cdots b_{j-1} & b_j & b_{j+1} \cdots b_n \\ \hline \end{array} \\ \hline = \begin{array}{|c|c|c|} \hline a_0 \cdots a_{j-1} & \mathbf{c} & b_{j+1} \cdots b_n \\ \hline \end{array} \end{array} .$$

By construction, \oplus is the inverse operation of \ominus , so that we have the following lemma:

Lemma A.1. *The following properties hold:*

- $M_i \oplus (M_{i+1} \ominus M_i) = M_{i+1}$
- $(M_i \oplus \widehat{M}_{i+1}) \ominus M_i = \widehat{M}_{i+1}$

Therefore, $M = M'$.

The FDAG of Figure 3 is reproduced below to illustrates the stages of the proof. This FDAG is constructed in 7 steps, that are (in this order): (\uparrow), (\heartsuit), (\heartsuit), (\uparrow), (\heartsuit), (\heartsuit) and (\heartsuit). The matrices A , M and \widehat{M} are given in Figure 12. One can see that \widehat{M} is of size 7, as expected.

Remark A.2. *It should be noted that (general) Fishburn matrices, with at least one non-zero entry on each row and column, are in bijection with FDAG compressing forests made of a unique tree. Indeed, via the bijection above, as such FDAG have a unique root, it must be the last inserted vertex, and therefore, each column admits at least one non-zero entry (otherwise it would be another root).*

Remark A.3. *It is possible to enumerate row-Fishburn matrices by using the previous bijection and the FDAG enumeration tree together. Nevertheless, things are a little simpler in this case and the equivalent of the operations (\uparrow) and (\heartsuit) can be merged, giving two rules for matrix expansion:*

- (R1) Increase one coefficient to the (inclusive) left of the rightmost nonzero coefficient of the top row by 1.
- (R2) Increase the dimension of the matrix by 1 (to the left and top), all new coefficients set to zero. Set one coefficient of the top row to 1.

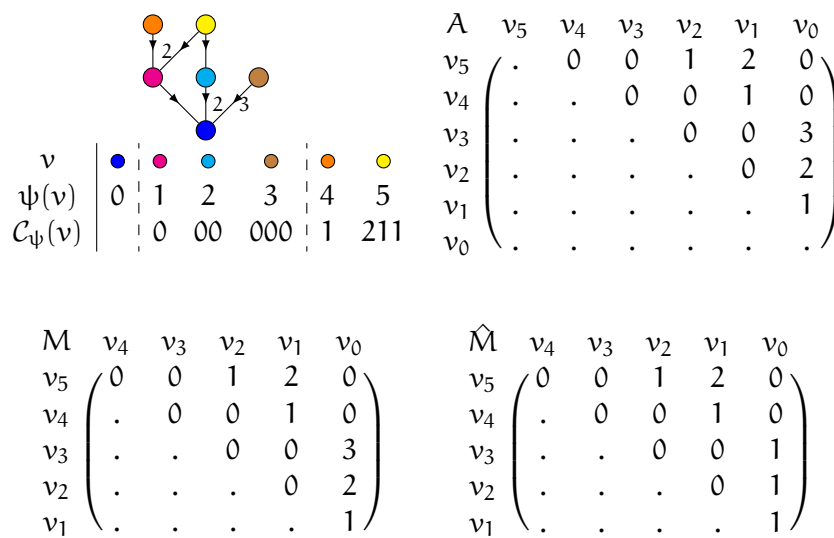


Figure 12: The FDAG of Figure 3 reproduced (top left), its adjacency matrix A (top right), its reduced adjacency matrix M (bottom left) and its incremental adjacency matrix \hat{M} (bottom right). Dots represent zeros corresponding to $A_{i,j}$ elements with $i \leq j$.