

On Canonical Forms for Frequent Graph Mining

Christian Borgelt

Dept. of Knowledge Processing and Language Engineering
Otto-von-Guericke-University of Magdeburg
Universitätsplatz 2, 39106 Magdeburg, Germany
borgelt@iws.cs.uni-magdeburg.de

Abstract. In approaches to frequent graph mining that are based on growing subgraphs into a set of graphs, one of the core problems is how to avoid redundant search. A powerful technique to overcome this problem is a canonical description of a graph, which uniquely identifies it, and a corresponding test. This paper introduces a family of canonical forms that are based on systematic ways to construct spanning trees. I show that the canonical form used in gSpan [14] is a member of this family, and that MoSS/MoFa [1, 3] is implicitly based on a different member, which I make explicit and exploit in the same way as in gSpan.

1 Introduction

In recent years there emerged an intense and still growing interest in the problem how to find common subgraphs in a database of (attributed) graphs, that is, subgraphs that appear with a user-specified minimum frequency. For this task—which has applications in, for example, biochemistry, web mining, and program flow analysis—several algorithms have been proposed. Some of them rely on principles from inductive logic programming and describe the graph structure by logical expressions [5]. However, the vast majority transfers techniques developed originally for frequent item set mining.¹ Examples of algorithms developed in this way include MolFea [10], FSG [11], MoSS/MoFa [1], gSpan [14], CloseGraph [15], FFMS [8], and Gaston [12]. A related approach is used in Subdue [4]. The basic idea of these approaches is to grow subgraphs into the graphs of the database, adding an edge and maybe a node in each step, counting the number of graphs containing each grown subgraph, and eliminating infrequent subgraphs.

While in frequent item set mining it is trivial to ensure that the same item set is checked no more than once in the search (using an arbitrary, but fixed global order of the items), in frequent subgraph mining it is one of the core problems how to avoid redundant search. Since the same subgraph can be grown in several different ways, adding the same nodes and edges in different orders, it is difficult to guarantee that each subgraph is considered only once. Although such multiple tests of the same subgraph do not invalidate the result, they can be devastating for the execution time of the algorithm. Therefore methods that rule out such redundant search are very important to make the algorithms efficient.

¹ See, for example, [6, 7] for details and references on frequent item set mining.

One of the most promising ideas to avoid redundant search is to define a canonical description of a (sub)graph. Together with a specific way of growing the subgraphs, such a canonical description can be used to check whether a given subgraph has been considered in the search before and thus need not be extended. This approach underlies the gSpan algorithm [14] and its extension CloseGraph [15]. In this paper I generalize the canonical form introduced in gSpan, thus arriving at a family of canonical descriptions of which the one used in gSpan is a special case. However, it is not the only usable one. I show that a competing algorithm called MoSS/MoFa [1, 3] is implicitly also based on a canonical description from this family, different from the gSpan one. By making this canonical form explicit, it can be exploited in MoSS/MoFa in the same way as in gSpan, leading to a significant improvement of the MoSS/MoFa algorithm.

2 Finding Frequent Subgraphs

For the following considerations it is important to review how a graph database is searched for frequent subgraphs: Given an initial node (for which all possibilities have to be tried), a subgraph is grown by adding an edge and, if necessary, a node in each step. In this stepwise extension process one requires that at least one node incident to an added edge must already be in the subgraph, thus restricting the search to connected subgraphs (which suffices for most applications). In its most basic form the search considers all possible extensions of the current subgraph by an edge and, if necessary, a node. (It will be shown later how the set of extensions can be reduced by exploiting a canonical description.)

Note that, as a consequence of the above, the search produces a numbering of the nodes in each subgraph: the index of a node simply reflects the step in which it was added. In the same way it produces an order of the edges—again the order in which they were added. Even more: the search builds a spanning tree of the subgraph, which is enhanced by additional edges.

3 Canonical Forms of Attributed Graphs

In this section I describe the family of canonical descriptions that is introduced in this paper, using the special cases employed in gSpan and MoSS/MoFa as examples and pointing out alternatives. How these canonical forms define an extension strategy and thus a search order is discussed in the next section.

3.1 General Idea

The core idea underlying the family of canonical forms introduced in this paper is to construct a code word that uniquely identifies a graph up to isomorphism and symmetry (i.e. automorphism). The characters of this code word describe the edges of the graph, in particular which nodes they connect. If the graph is attributed or directed, they also comprise information about the attribute and/or direction of the edge as well as the attributes of the incident nodes.

While it is straightforward to capture the latter information about an edge (i.e. attributes and edge direction) in some alphabet, how to describe the connection structure is not so obvious. For this, the nodes of the graph must be numbered (or more generally: endowed with unique labels), because we need a way to specify the source and the destination node of an edge. Unfortunately, there are many different ways of numbering the nodes of a graph, all of which may lead to different code words, because each numbering leads to a different specification of an edge (simply because the indices of the source and the destination node differ). In addition, the edges can be listed in different orders. How these two problems can be treated is described in the following two subsections: the different possible solutions give rise to different canonical forms.

However, given a (systematic) way of numbering the nodes of a (sub)graph and a sorting criterion for the edges, a canonical description is generally derived as follows: each numbering of the nodes yields a code word, which is the concatenation of the sorted edge descriptions (details about the form of such a code word are given in Section 3.4). The resulting list of code words is sorted lexicographically. Then the lexicographically smallest code word is the canonical description. (Note that the graph can be reconstructed from this code word.)

3.2 Constructing Spanning Trees

From the review of the search process in Section 2 it is clear that we can confine ourselves to numberings of the nodes of a (sub)graph that result from spanning trees, because no other node numberings will ever occur in the search. Even more: we can confine ourselves to a specific systematic way of constructing a spanning tree. The reason is that the basic search algorithm produces *all* spanning trees of a (frequent) (sub)graph, though usually in different branches of the search tree.² Since the extensions of a (sub)graph need to be checked only once, we may choose to form them only in the branch of the search tree, in which the spanning tree of the (sub)graph has been built in the chosen way.

The most straightforward systematic methods for constructing a spanning tree of a graph are, of course, depth-first and breadth-first search. Thus it is not surprising to see that gSpan uses the former to define its canonical form [14]. However, the latter (i.e., breadth-first search) may just as well be chosen as a basis for a canonical form. And indeed: as will turn out later, the (heuristically introduced) local extension order of the MoSS/MoFa algorithm [1, 3] can be justified from a canonical form that is based on such a breadth-first search tree. Thus MoSS/MoFa can be seen as implicitly based on this canonical form.

Other alternatives include a spanning tree construction that first visits all neighbors of a node (like breadth-first search), but then chooses the next node to extend in a depth-first search manner. (This may be seen as a variant of depth-first search.) However, in the following I confine myself to (standard) depth-first and breadth-first search trees to keep things simple. Nevertheless, it should be kept in mind that there is a variety of other possibilities one may explore.

² They occur in the same search tree node only if the graph exhibits some symmetry, i.e., if there exists an automorphism that is not the identity.

It should be noted that at this point there is no restriction on the order in which the neighbors of a node are visited in the search. Hence there is generally a large number of different spanning trees, even if the root node is fixed. As a consequence, choosing a method for constructing a spanning tree is not sufficient to avoid redundant search. Since there are usually several spanning trees of a (sub)graph that are constructed in the chosen way, there are several search tree branches that qualify for an extension of the (sub)graph. Although this freedom will be reduced below by exploiting edge and node attributes, it cannot be eliminated completely, since there are no local (i.e. node-specific) criteria that allow for an unambiguous decision in all cases ([1] gives an example). Therefore we actually need to construct and compare code words to avoid all redundancy.

3.3 Edge Sorting Criteria

Once we have a numbering of the nodes of the graph, we can set up the edge descriptions and sort them. In principle, the edge descriptions can be sorted using any precedence order of the properties of an edge (i.e. attribute of the edge and attributes and indices of the source and destination node). However, we can exploit the purpose for which the canonical form is intended to find appropriate sorting criteria. Recall that the search constructs different spanning trees of the same (sub)graph in different branches of the search tree. Each of these gives rise to a numbering of the nodes and thus a code word. In addition, recall that the canonical form is intended for confining the extensions of a (sub)graph to one branch of the search tree. Hence we need a way of checking whether the code word resulting from the node numbering in a search tree node is minimal or not: if it is, we descend into the search tree branch, otherwise we prune it.

In order to carry out this test, we could construct all other possible code words for the same (sub)graph and compare them to the one resulting from the node numbering in the current search tree node. However, such a straightforward approach is much too costly. Fortunately, it can be made much more efficient by the insight that the code words are compared lexicographically. Hence we may not need to know the full code words in order to decide which of them is lexicographically smaller—a prefix may suffice. This immediately gives rise to the idea to check all code words in parallel that share the same prefix. However, whether this is (easily) possible or not, depends on how we sort the edges.

Fortunately, for both canonical forms, depth-first and breadth-first search, there is a sorting criterion that yields such an order. The core idea is to define the order of the edges in such a way that they are sorted into the order in which they have been added to the (sub)graph in the search. This has two advantages: in the first place, we need no sorting to obtain the code word of the (sub)graph under consideration that results from the node numbering in the current search tree node. Since it is easiest to implement the search by always appending the added edge to a list of contained edges, this edge list already yields the code word. Secondly, we can carry out the search for alternative code words in basically the same way as the whole search for frequent subgraphs. Doing so makes it possible to compare the prefixes of the code words after the addition of every

single edge, thus making the test of a code word maximally efficient. Details about the comparisons are given below, after the exact form of the code words for the two canonical forms considered here are defined.

3.4 Code Words

In my definition of a code word I deviate slightly from the definition of gSpan [14], where code words are simple lists of edge descriptions, each of which comprises all information about the edge and the incident nodes. This deviation is triggered by the insight that it is not necessary to compare the attribute of the source node, except for the first edge that is added. In other words, we may precede the sequence of edge descriptions by a character that specifies the attribute of the root of the spanning tree, while at the same time we cancel the attribute of the source node from the following edge descriptions. Then the general forms of code words (as regular expressions with non-terminal symbols) are:

- Depth-First Search: $a(i_d[n - i_s]ba)^m$
- Breadth-First Search: $a(i_sba i_d)^m$

In these words a is a node attribute and b an edge attribute. n is the total number of nodes in the (sub)graph and m the total number of edges. i_s is the index of the source node and i_d the index of the destination node of an edge. (Source and destination of an edge are defined by the relation $i_s < i_d$, that is, the incident node with the smaller index is the source.) Parentheses are for grouping characters; each parenthesized sub-word stands for one edge. The exponent m means that there are m repetitions of the group of characters to which it is attached. Square brackets indicate one character, the value of which is computed by the expression inside them. That is, in the depth-first search code word the second character for each edge is the number $n - i_s$.

The describing properties of an edge are compared in the order in which they appear in the parenthesized expressions. All characters are compared ascendingly. (This explains the expression in square brackets; alternatively one may define that i_s has to be compared descendingly here.) Note that the parenthesized expressions are sorted (that is, the edge descriptions are sorted) and are concatenated afterwards to form the code word. It is easy to see that in this way the code word describes how the edges have been added in the search.

It should be noted that one may let spanning tree edges take absolute precedence over other edges. That is, the code word may start with the spanning tree edges, and only after all of them have been listed, the other edges (which lead to cycles) are listed. In my definition, however, there is no such distinction of edge types. The order of the edge descriptions in the code words is defined by the stated edge properties alone and thus spanning tree edges may be intermingled with edges closing cycles. The reason is that I want the edges to be in exactly the order in which they have been added to the (sub)graph in the search tree. However, one may also choose to find spanning trees first before closing cycles. As the ideas underlying the Gaston algorithm [12] suggest, there may be good reasons for adopting such a strategy, as it may speed up the search.

3.5 Checking for Canonical Form

After the code words are defined, the test whether a code word is a canonical description of a (sub)graph can be formally stated. The pseudocode below describes the procedure. w is the code word to be tested, $G = (V, E)$ is the corresponding (sub)graph. Each node $v \in V$ has an attribute $v.a$, which I assume to be coded as an integer, and an index field $v.i$, which is filled by the algorithm. Likewise each edge $e \in E$ has an attribute $e.a$, which again I assume to be coded as an integer, and a marker $e.i$, which is used to record whether it was visited. Since apart from node and edge attributes a code word contains only indices of nodes, it can thus be represented as an array of integers.

```
function isCanonical ( $w$ : array of int,  $G$ : graph) : boolean;
var  $v$  : node;           (* to traverse the nodes of the graph *)
     $e$  : edge;           (* to traverse the edges of the graph *)
     $x$  : array of node;  (* to collect the numbered nodes *)
begin
  forall  $v \in G.V$  do           (* traverse all nodes and *)
     $v.i := -1$ ;                 (* clear their indices *)
  forall  $e \in G.E$  do           (* traverse all edges and *)
     $e.i := -1$ ;                 (* clear their markers *)
  forall  $v \in G.V$  do begin     (* traverse the potential root nodes *)
    if  $v.a = w[0]$  then begin  (* if  $v$  is acceptable as a root node *)
       $v.i := 1$ ;  $x[0] := v$ ;    (* number and record the root node *)
      if not  $\text{rec}(w, 1, x, 1, 0)$  (* check the code word recursively and *)
      then return false;      (* abort if a smaller code word is found *)
       $v.i := -1$ ;             (* clear the node index again *)
    end
  end
  return true;                (* the code word is canonical *)
end
```

This function is the same, regardless of whether a depth-first or a breadth-first search canonical form is used. The difference lies only in the implementation of the function “rec”, mainly in the order in which edge properties are compared. Here I confine myself to the implementation for a breadth-first search. However, for a depth-first search the function can be implemented in a very similar way.

The basic idea of the recursion is to add one edge in each level of the recursion. The description of this edge is generated and if it already allows to decide whether the generated code word is larger or smaller (prefix test!), the recursion is terminated immediately. Only if the edge description coincides with the one found in the code word to check, the edge and the node at the other end (if necessary) are marked/numbered and the function is called recursively. Note that the loop over the edges incident to the node $x[i]$ in the pseudocode below assumes that the edges are considered in sorted order, that is, the edges with the smallest attribute are tested first, and among edges with the same attribute, they are considered in increasing order of the attribute of the destination node.

```

function rec (w: array of int, k: int, x: array of node, n: int, i: int) : boolean;
var d: node;                (* node at the other end of an edge *)
    j: int;                  (* index of destination node *)
    u: boolean;             (* flag for unnumbered destination node *)
    r: boolean;             (* buffer for a recursion result *)
begin
  if k ≥ length(w) return true;  (* full code word has been generated *)
  while i < w[k] do begin      (* check whether there is an edge with *)
    forall e incident to x[i] do (* a source node having a smaller index *)
      if e.i < 0 then return false;
      i := i + 1;                (* go to the next extendable node *)
    end
    forall e incident to x[i] (in sorted order) do begin
      if e.i < 0 then begin      (* traverse the unvisited incident edges *)
        if e.a < w[k + 1] then return false;  (* check the *)
        if e.a > w[k + 1] then return true;  (* edge attribute *)
        d := node incident to e other than x[i];
        if d.a < w[k + 2] then return false;  (* check destination *)
        if d.a > w[k + 2] then return true;  (* node attribute *)
        if d.i < 0 then j := n else j := d.i;
        if j < w[k + 3] then return false;  (* check destination *)
        if j = w[k + 3] then begin      (* node index *)
          e.i := 1; u := d.i < 0; (* mark edge and number node *)
          if u then begin d.i := j; x[n] := d; n := n + 1; end
          r := rec(w, k + 4, x, n, i); (* check recursively *)
          if u then begin d.i := -1; n := n - 1; end
          e.i := -1; (* unmark edge (and node) again *)
          if not r then return false;
        end
      end
    end
  return true; (* return that no smaller code word *)
end (* than w could be found *)

```

3.6 A Simple Example

In order to illustrate the code words defined above, Figure 1 shows a simple molecule (no chemical meaning attached; it has been constructed merely for illustration purposes). This molecule is represented as an attributed graph, in which each node stands for an atom and each edge for a bond between atoms. The nodes carry the chemical element of the corresponding atom as an attribute, the edges are associated with bond types. To the right of this molecule are two spanning trees for this molecule, both of which are rooted at the sulfur atom. The spanning tree edges are depicted as solid lines, the edges closing cycles as dashed lines. Spanning tree A was built with depth-first search, spanning tree B with breadth-first search and thus correspond to the two considered approaches.

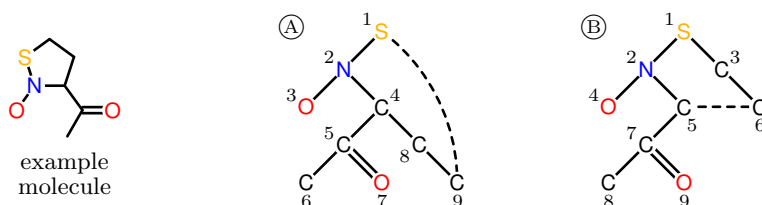


Fig. 1. An example fragment/molecule and two possible canonical forms: A – form based on a depth first search tree, B – form based on a breadth first search tree.

A: S 28-N 37-O 47-C 55-C 64-C 74=O 85-C 91-C 98-S
 1 2 2 4 5 5 4 8 1

B: S 1-N2 1-C3 2-O4 2-C5 3-C6 5-C6 5-C7 7-C8 7-O9

Fig. 2. Code words describing the two canonical forms shown in Figure 1. The second entries in the bond strings in A are $n - i_s$, where n is the number of nodes and i_s the index of the source node. The actual values of i_s are shown below the code.

If we adopt the precedence order $S \prec N \prec O \prec C$ for chemical elements (which is derived from the frequency of the elements in the molecule) and the order $- \prec =$ for the bond types, we obtain the two code words shown in Figure 2. It is easy to check that these two code words are actually minimal and thus represent the canonical description w.r.t. a depth-first and breadth-first search spanning tree, respectively. In this case it is particularly simple to check this, because the root of the spanning tree is fixed, as there is only one sulfur atom.

4 Restricted Extensions

Up to now canonical descriptions were only used to test whether a search tree branch corresponding to a (sub)graph has to be descended into or not *after* the (sub)graph has been constructed. However, canonical forms can also be used to restrict the possible extensions directly. The idea is that for certain extensions by an edge and maybe a node one can see immediately that they lead to a code word that is not minimal. Hence one need not construct and test the (sub)graph, but can skip the extension right way. For the two special cases I consider here (depth-first and breadth-first search spanning trees), the allowed extensions are:

- Depth First Search: *Rightmost Extension* [14]

Only nodes on the rightmost path of the spanning tree of the (sub)graph may be extended, and if the node is no leaf, it may be extended only by edges whose descriptions do not precede the description of the downward edge on the rightmost path. That is, the edge attribute must be no less than the attribute of the downward edge and if the edge attribute is identical,

the attribute of its destination node must be no less than the attribute of the downward edge's destination node. Edges between two nodes that are already in the (sub)graph must lead from a node on the rightmost path to the rightmost leaf (that is, the deepest node on the rightmost path). In addition, the index of the source node of such an edge must precede the index of the source node of an edge already incident to the rightmost leaf.

– Breadth First Search: *Maximum Source Extension* [1]

Only nodes having an index no less than the maximum source index of an edge already in the (sub)graph may be extended.³ If the node is the one having the maximum source index, it may be extended only by edges whose descriptions do not precede the description of any downward edge already incident to this node. That is, the attribute of the new edge must be no less than that of any downward edge, and if it is identical, the attribute of the new edge's destination node must be no less than the attribute of any corresponding downward edge's destination node (where corresponding means that the edge attribute is the same). Edges between two nodes already in the (sub)graph must start at an extendable node and must lead "forward", that is, to a node having a larger index.

In both cases it is easy to see that an extension violating the above rules leads to a (sub)graph description that is not in canonical form (that is, a numbering of the nodes of the (sub)graph that does not lead to the lexicographically smallest code word). This is easy to see, because there is a depth-first or breadth-first search numbering, respectively, *starting at the same root node*, which leads to a lexicographically smaller code word (which may or may not be minimal itself—all that matters here is that it is smaller than the one derived from the current node numbering of the (sub)graph). In order to find such a smaller word, one only has to consider the extension edge. Up to this edge, the construction of the code word is identical, but when it is added, its description precedes the description of the next edge in the code word of the unextended (sub)graph.

It is pleasant to see that the *maximum source extension*, which was originally introduced in the MoSS/MoFa algorithm based on heuristic arguments [1, 3], can thus nicely be justified and extended based on a canonical form.

As an illustration of the above extension rules, consider again the two search trees shown in Figure 1. In the depth-first search tree A atoms 1, 2, 4, 8, and 9 are extendable (rightmost extension). On the other hand, in the breadth-first search tree B atoms 7, 8, and 9 are extendable (maximum source extension). Other restrictions are, for example, that the nitrogen atom in A may not be extended by a bond to another oxygen atom, or that atom 7 in B may not be extended by a single bond. Tree A may not be extended by an edge between two nodes already in the (sub)graph, because the edge (1, 9) already has the smallest possible source (duplicate edges between nodes may be allowed, though). Tree B, however, may be extended by an edge between atoms 8 and 9.

³ Note that if the (sub)graph contains no edge, there can only be one node, and then, of course, this node may be extended without restriction.

5 Experimental Results

In order to test the pruning based on a breadth-first search canonical form, I extended the MoSS/MoFa implementation described in [3]. In order to compare the two canonical forms discussed in this paper, I also implemented a search based on rightmost extensions and a corresponding test for a depth-first search canonical form (that is, basically the gSpan algorithm [14]). This was done in such a way that only the functions explicitly referring to the canonical form are exchanged (extension generation and comparison and canonical form check), so that on execution the two algorithms share a maximum of the program code.

Figure 3 shows the results on the 1993 subset of the Index Chemicus (IC93) [9], Figure 4 shows the results on a data set consisting of 17 steroids. In all diagrams the grey solid line describes the results for a depth-first canonical form, the black solid line the results for a breadth-first canonical form. The diagram in the top left shows the number of generated and actually processed fragments (note that the latter, shown as a dashed grey line, is necessarily the same for both algorithms). The diagram in the top right shows the number of generated embeddings, the diagram in the bottom left the execution times.⁴

As can be seen from these diagrams, both canonical forms work very well (compared to an approach without canonical form pruning and an explicit removal of found duplicates, I observed speed-ups by factors between about 2.5 and more than 30). On the IC93 data the breadth-first search canonical form performs slightly better, needing about 10–15% less time. As the other diagrams show, this is mainly due to the lower numbers of fragments and embeddings that are generated. On the steroids data the depth-first search canonical form performs minimally better at low support values. Again this is due to a smaller number of generated fragments, which, however, is outweighed by a larger number of generated embeddings for higher support values.

6 Conclusions

In this paper I introduced a family of canonical forms of graphs that can be exploited to make frequent graph mining efficient. This family was obtained by generalizing the canonical form introduced in the gSpan algorithm [14]. While gSpan’s canonical form is defined with a depth-first search tree, my definition allows for any systematic way of obtaining a spanning tree. To show that this generalization is useful, I considered a breadth-first search spanning tree, which turned out to be the implicit canonical form underlying the MoSS/MoFa algorithm [1, 3]. Exploiting this canonical form in MoSS/MoFa in the same way as the depth-first search canonical form is exploited in gSpan leads to a considerable speed up of this algorithm. It is pleasant to see that based on this generalized canonical form, gSpan and MoSS/MoFa can nicely be described in the same general framework, which also comprises a variety of other possibilities.

⁴ Experiments were done with Sun Java 1.5.0.01 on a Pentium 4C@2.6GHz system with 1GB main memory running S.u.S.E. Linux 9.3.

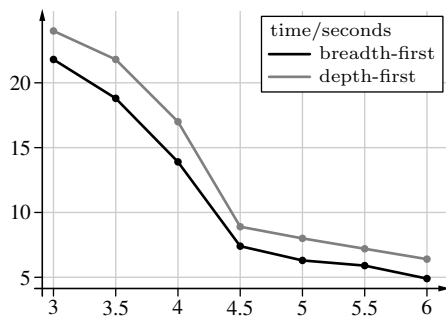
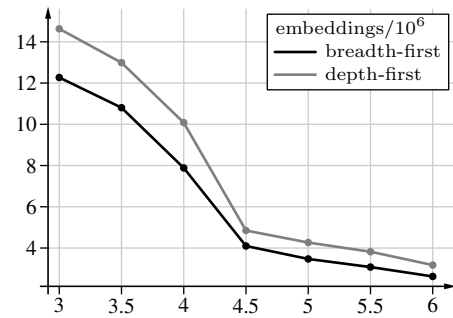
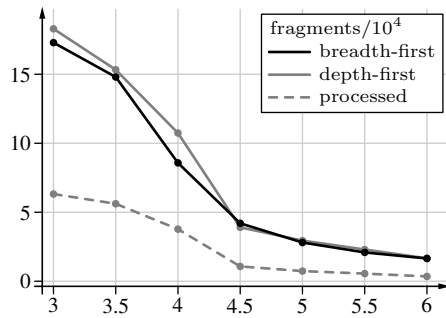


Fig. 3. Experimental results on the IC93 data. The horizontal axis shows the minimal support in percent. The curves show the number of generated and processed fragments (top left), number of generated embeddings (top right), and the execution time in seconds (bottom left) for the two canonical forms/extension strategies.

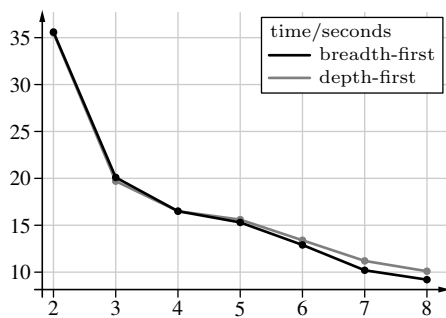
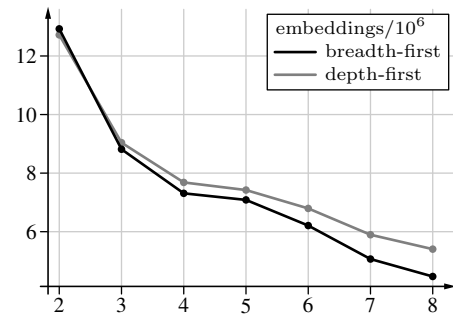
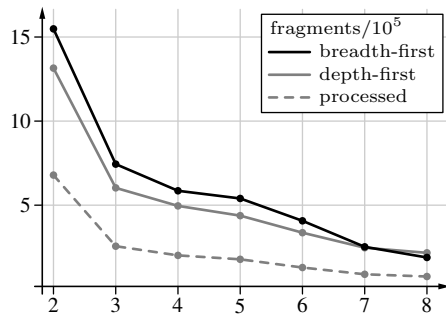


Fig. 4. Experimental results on the steroids data. The horizontal axis shows the absolute minimal support. The curves show the number of generated and processed fragments (top left), number of generated embeddings (top right), and the execution time in seconds (bottom left) for the two canonical forms/extension strategies.

References

1. C. Borgelt and M.R. Berthold. Mining Molecular Fragments: Finding Relevant Substructures of Molecules. *Proc. IEEE Int. Conf. on Data Mining (ICDM 2002, Maebashi, Japan)*, 51–58. IEEE Press, Piscataway, NJ, USA 2002
2. C. Borgelt, T. Meinl, and M.R. Berthold. Advanced Pruning Strategies to Speed Up Mining Closed Molecular Fragments. *Proc. IEEE Conf. on Systems, Man and Cybernetics (SMC 2004, The Hague, Netherlands)*, CD-ROM. IEEE Press, Piscataway, NJ, USA 2004
3. C. Borgelt, T. Meinl, and M.R. Berthold. MoSS: A Program for Molecular Substructure Mining. *Proc. Open Source Data Mining Workshop (OSDM 2005, at KDD2005, Chicago, IL)*, to appear.
4. D.J. Cook and L.B. Holder. Graph-Based Data Mining. *IEEE Trans. on Intelligent Systems* 15(2):32–41. IEEE Press, Piscataway, NJ, USA 2000
5. P.W. Finn, S. Muggleton, D. Page, and A. Srinivasan. Pharmacore Discovery Using the Inductive Logic Programming System PROGOL. *Machine Learning*, 30(2-3):241–270. Kluwer, Amsterdam, Netherlands 1998
6. B. Goethals and M. Zaki. *Proc. 1st IEEE ICDM Workshop on Frequent Itemset Mining Implementations (FIMI 2003, Melbourne, FL)*. CEUR Workshop Proceedings 90, Aachen, Germany 2003.
<http://www.ceur-ws.org/Vol-90/>
7. B. Goethals and M. Zaki. *Proc. 2nd IEEE ICDM Workshop on Frequent Itemset Mining Implementations (FIMI 2003, Melbourne, FL)*. CEUR Workshop Proceedings 126, Aachen, Germany 2003.
<http://www.ceur-ws.org/Vol-126/>
8. J. Huan, W. Wang, and J. Prins. Efficient Mining of Frequent Subgraphs in the Presence of Isomorphism. *Proc. 3rd IEEE Int. Conf. on Data Mining (ICDM 2003, Melbourne, FL)*, 549–552. IEEE Press, Piscataway, NJ, USA 2003
9. *Index Chemicus — Subset from 1993*. Institute of Scientific Information, Inc. (ISI). Thomson Scientific, Philadelphia, PA, USA 1993
10. S. Kramer, L. de Raedt, and C. Helma. Molecular Feature Mining in HIV Data. *Proc. 7th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD 2001, San Francisco, CA)*, 136–143. ACM Press, New York, NY, USA 2001
11. M. Kuramochi and G. Karypis. Frequent Subgraph Discovery. *Proc. 1st IEEE Int. Conf. on Data Mining (ICDM 2001, San Jose, CA)*, 313–320. IEEE Press, Piscataway, NJ, USA 2001
12. S. Nijssen and J.N. Kok. A Quickstart in Frequent Structure Mining can Make a Difference. *Proc. 10th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD2004, Seattle, WA)*, 647–652. ACM Press, New York, NY, USA 2004
13. T. Washio and H. Motoda. State of the Art of Graph-Based Data Mining. *SIGKDD Explorations Newsletter* 5(1):59–68. ACM Press, New York, NY, USA 2003
14. X. Yan and J. Han. gSpan: Graph-Based Substructure Pattern Mining. *Proc. 2nd IEEE Int. Conf. on Data Mining (ICDM 2003, Maebashi, Japan)*, 721–724. IEEE Press, Piscataway, NJ, USA 2002
15. X. Yan and J. Han. Closegraph: Mining Closed Frequent Graph Patterns. *Proc. 9th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD 2003, Washington, DC)*, 286–295. ACM Press, New York, NY, USA 2003