

Subgraph Support in a Single Large Graph

Mathias Fiedler and Christian Borgelt

European Center for Soft Computing

c/ Gonzalo Gutiérrez Quirós s/n, 33600 Mieres, Spain

Email: mail@mathias-fiedler.info, christian.borgelt@softcomputing.es

Abstract—Defining the support (or frequency) of a subgraph is trivial when a database of graphs is given: it is simply the number of graphs in the database that contain the subgraph. However, if the input is one large graph, an appropriate support definition is much more difficult to find. In this paper we study the core problem, namely overlapping embeddings of the subgraph, in detail and suggest a definition that relies on the non-existence of equivalent ancestor embeddings in order to guarantee that the resulting support is anti-monotone. We prove this property and describe a method to compute the support defined in this way.

I. INTRODUCTION

In recent years frequent subgraph mining has received intense and still growing attention, since it has a wide and constantly expanding range of applications areas, which include biochemistry, web mining, and program flow analysis. As a consequence, several frequent subgraph mining algorithms have been developed. Some of them rely on principles from inductive logic programming and describe the graph structure by logical expressions [6]. However, the vast majority transfers techniques that were originally developed for frequent item set mining. Examples include MolFea [9], FSG [10], MoSS/MoFa [1], gSpan [14], CloseGraph [15], FFSM [7], and Gaston [12]. A related, but slightly different approach is used in Subdue [3].

Most existing work in frequent subgraph mining considers a database of (attributed) graphs as input. In this setting it is straightforward to define the support (or frequency) of a subgraph: it is simply the number of graphs in the given database that contain the subgraph. If, however, the input is a single large graph (which is the case we consider here), it is surprisingly difficult to find a support definition that has certain desirable properties. Until recently there was only one suggestion, which is used in the algorithms in [11]: the support of a subgraph is the size of a maximum independent node set of the overlap graph of the embeddings of the subgraph (details of this definition are reviewed in Sections III to V).

However, as we argue in this paper, this approach has the drawback that it can sometimes be too restrictive: there are certain cases in which it allows us to count at most one of two embeddings, even though counting both does not have any harmful effects. In addition, a proper proof was lacking that the support defined in this way is anti-monotone. By studying the core problem of overlapping embeddings in detail, we arrive at a modified support definition and also provide a clear proof that both types of subgraph support are anti-monotone.

Finally, we briefly look at a recently presented alternative support definition [4], which is based on the minimum number of different images of a node, and compare its properties.

II. FREQUENT SUBGRAPH MINING

In order to fix the algorithmic setting, we briefly review the core principles of frequent subgraph mining. The search is usually restricted to connected subgraphs, since this reduces the search space considerably and suffices for most applications. The search grows all possible (connected) subgraphs, starting from a single node and adding an edge and maybe a node (if it is not yet in the subgraph) in each step. For each grown subgraph the support is computed (for example, as the number of database graphs that contain the subgraph) and infrequent subgraph are eliminated (where infrequent means that the support does not reach a user-specified threshold).

Since with this basic procedure the same subgraph can be grown in several ways, namely by adding its nodes and edges in different orders, a fundamental problem is how to avoid redundant search. The predominant method for this is so-called *canonical form pruning*, which is based on the definition of a canonical code word of a graph that uniquely identifies it up to automorphisms. Together with a specific way of growing subgraphs (and thus of building code words), a canonical form can be used to check whether a subgraph has already been considered and thus can be pruned from the search tree [2]. The exact way in which the search space (that is, the subgraph (semi-)lattice) is traversed does not matter much. We use a depth-first search, since this has advantages w.r.t. memory consumption, but a breadth-first approach is also feasible.

It should be noted that the above description does not fix a specific definition of the support of a subgraph. However, it exploits a fundamental property, namely that the support is anti-monotone: the fact that infrequent subgraphs are eliminated from the search presupposes that no supergraph of a subgraph can have a higher support than the subgraph itself. Otherwise it would not be possible to prune infrequent subgraphs, because we would run the risk to miss frequent subgraphs. Since support-based pruning is essential for the efficiency of the mining algorithm, it is of vital importance to ensure that the support is defined in such a way that it is anti-monotone.

III. EMBEDDINGS

The fundamental concept underlying any type of support definition for subgraphs is that of an *embedding* of the subgraph into the input graph(s). It is formally defined by the notion of a *subgraph isomorphism*, which we consider in the context of *labeled* or *attributed (simple) graphs*.

Definition 1: A *labeled* or *attributed graph* is a triple $G = (V, E, l)$, where V is the set of vertices, $E \subseteq V \times V -$

$\{(v, v) \mid v \in V\}$ is the set of edges, and $l : V \cup E \rightarrow L$ is a function that assigns labels from the set L to nodes and edges.

Definition 2: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs. A *subgraph isomorphism* of S to G is an injective function $f : V_S \rightarrow V_G$ satisfying $\forall v \in V_S : l_S(v) = l_G(f(v))$ and $\forall (u, v) \in E_S : (f(u), f(v)) \in E_G \wedge l_S((u, v)) = l_G((f(u), f(v)))$.

Every subgraph isomorphism of a graph S to the input graph(s) defines one *embedding* of S (we use *embedding* as a synonym for *subgraph isomorphism*). Note that two different embeddings may refer to the same nodes and edges in G or, more generally, may share a subset of the nodes and edges.

Definition 3: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs, f_1 and f_2 two subgraph isomorphisms of S to G , and let $V_i = \{v \in V_G \mid \exists u \in V_S : v = f_i(u)\}$, $E_i = \{e \in E_G \mid \exists (u, v) \in E_S : e = (f_i(u), f_i(v))\}$, $i = 1, 2$. That is, let V_i and E_i be the f_i -images (in G) of the nodes and edges of S . f_1 and f_2 are called *overlapping*, written $f_1 \circledast f_2$, iff $V_1 \cap V_2 \neq \emptyset$. f_1 and f_2 are called *equivalent*, written $f_1 \circ f_2$, iff $V_1 = V_2$ and $E_1 = E_2$. Finally, f_1 and f_2 are called *identical*, written $f_1 \equiv f_2$, iff $\forall v \in V_S : f_1(v) = f_2(v)$.

Note that two identical subgraph isomorphisms are actually the same subgraph isomorphism and thus define only one embedding. Note also that identical subgraph isomorphisms are necessarily equivalent, which in turn must be overlapping. The inverse inclusions, however, do not hold in general.

IV. SUBGRAPH SUPPORT AND OVERLAP GRAPHS

In terms of the notions defined in the preceding section the support of a graph S in a given database of graphs is simply the number of graphs into which there exists an embedding of S . However, if the input is a single large graph, this definition is not particularly useful: since there either is an embedding of S into the input graph or not, the support of any graph is either 1 or 0. Even worse: if the output is restricted to closed subgraphs (which are subgraphs no supergraph of which has the same support), the mining result is always the input graph (or its connected components) and thus entirely useless.

The most intuitive definition of the support of a subgraph in the single graph setting would be to simply count its embeddings. However, this definition is not feasible, since it is not anti-monotone: the support of a supergraph of a subgraph S can exceed the support of S , as is demonstrated by the example shown in Figure 1. (In order to keep things simple, only nodes are labeled, while edges are assumed to be unlabeled, or rather seen as all labeled with the same label, so that it can be neglected.) The graph consisting of a single node labeled A has one embedding into the input graph B-A-B and thus support 1. However, the graph B-A-B, obviously a supergraph of the single node graph, has two embeddings (which are equivalent and differ only in the way in which the nodes labeled B are mapped). Hence it has support 2.

The same example also shows that we cannot fix this problem by counting at most one out of a set of equivalent embeddings (which would be a simple and straightforward amendment, since equivalent embeddings are, in an intuitive




Fig. 1. The number of embeddings of a subgraph is not anti-monotone.




Fig. 2. Overlap graph of the subgraph B-A-B w.r.t. the graph B-A-B-A-B.

sense, the “same” embedding): the intermediate subgraph A-B also has two embeddings—which are not equivalent, because the node labeled B is mapped to different nodes. Hence it also has support 2, thus violating anti-monotony.

A closer look at this example suggests that the problem is caused by overlapping embeddings (see Definition 3), since the two embeddings of the subgraph A-B overlap on the node labeled A, which, if considered alone, has only one embedding. Therefore it is plausible to consider the overlap graph of the embeddings of a subgraph S and derive the support of S from this overlap graph. This is the basis of the approach used in the algorithms presented in [11].

Definition 4: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and let V_O be the set of all embeddings (that is, subgraph isomorphisms) of S into G . The *overlap graph* of S w.r.t. G is the graph $O = (V_O, E_O)$, which has the set V_O of embeddings as its node set and the edge set $E_O = \{(f_1, f_2) \mid f_1, f_2 \in V_O \wedge f_1 \neq f_2 \wedge f_1 \circledast f_2\}$.

An example of an overlap graph for the subgraph B-A-B w.r.t. the graph B-A-B-A-B is shown in Figure 2. It contains four nodes, one for each of the four possible embeddings. Since in this case every embedding overlaps with every other (because they share at least the node labeled B in the middle of the input graph), the overlap graph is a complete graph.

V. MAXIMUM INDEPENDENT SET SUPPORT

Since overlaps of embeddings appear to be harmful, it is plausible to count at most one of each pair of overlapping embeddings. This leads directly to the idea to define the support of a subgraph as the size of a maximum independent node set of the overlap graph of its embeddings.

Definition 5: Let $G = (V, E)$ be a graph with node set V and edge set $E \subseteq V \times V - \{(v, v) \mid v \in V\}$. An *independent node set* of G is a set $I \subseteq V$ with $\forall u, v \in I : (u, v) \notin E$. I is a *maximum independent node set* iff it is an independent node set and for all independent node sets J of G it is $|I| \geq |J|$.

Note that a maximum independent set need not be unique: several independent node sets may have the same (maximum) size. For example, for the overlap graph shown in Figure 2, any single node is a maximum independent node set.

Definition 6: Let $O = (V_O, E_O)$ be the overlap graph of the embeddings of a labeled graph $S = (V_S, E_S, l_S)$ into a labeled graph $G = (V_G, E_G, l_G)$. Then the *maximum independent set support* (or *MIS-support* for short) of the graph S w.r.t. the graph G is the size of a maximum independent node set of O .

For example, the MIS-support of the graph B-A-B w.r.t. the graph B-A-B-A-B is 1 (see the overlap graph in Figure 2).

The approach of [11] is essentially based on the MIS-support of a graph. The only difference is that [11] defines that two embeddings overlap if they share an edge, while we defined that they overlap if they share a node. Although this leads to minor differences, the basic properties are the same.¹

However, a fundamental problem of [11] is that no proof is given that MIS-support (with either definition of overlap) is anti-monotone. For this, [11] refers to [13], which, however, actually considers a different property (which is difficult to relate to MIS-support) and itself refers to a (hard to obtain) Master’s thesis for the details of its proof. Therefore we provide here a clear and fairly straightforward proof, which can be transferred directly to our own support definition.

Definition 7: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and let f_S be an embedding of S into G . An embedding f_T of a proper subgraph $T = (V_T, E_T, l_T)$ of S (that is, $V_T \subset V_S$, $E_T = (V_T \times V_T) \cap E_S$, and $l_T \equiv l_S|_{V_T \cup E_T}$) is called a *T-ancestor* of f_S iff $f_T \equiv f_S|_{V_T}$, that is, if the embedding f_T coincides with f_S on the node set V_T of T .

We now consider two very simple observations about T -ancestors of embeddings, which form the basis of our proof.

Observation 1: For given G, S, T and f_S the T -ancestor f_T of the embedding f_S is uniquely defined.

This is obvious, since f_T is only the restriction of f_S to the node set V_T of T and thus there is no possibility to choose.

Observation 2: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and let $T = (V_T, E_T, l_T)$ be a proper subgraph of S . Furthermore, let f_1 and f_2 be two (non-identical, but possibly equivalent) embeddings of S into G . f_1 and f_2 overlap if there exist overlapping T -ancestors f'_1 and f'_2 of the embeddings f_1 and f_2 , respectively.

This is obvious, since f'_1 and f'_2 are restrictions of f_1 and f_2 to a subset of the nodes of S . Hence, if f'_1 and f'_2 overlap, so must f_1 and f_2 . Note, however, that f_1 and f_2 may overlap even if f'_1 and f'_2 do not overlap: they may overlap on (a subset of) $V_S - V_T$. Hence the reverse implication does not hold.

Theorem 1: MIS-support is anti-monotone.

Proof: We have to show that the MIS-support of a subgraph S w.r.t. a graph G cannot exceed the MIS-support of any (non-empty) proper subgraph T of S . To do so we consider an arbitrary independent node set I_S of the overlap O_S graph of S w.r.t. G . This node set induces a subset I_T of the nodes of the overlap graph O_T of an (arbitrary, but fixed) subgraph T of S ,

¹It has to be noted, though, that defining embeddings as overlapping only if they share an edge makes it possible that they can share several nodes without being considered as overlapping (which may be somewhat unintuitive). It also rules out the possibility to consider single node subgraphs, since such an approach does not rule out violations of anti-monotony occurring in the step from single node subgraphs to subgraphs with two nodes.




Fig. 3. Not all overlaps of embeddings are harmful.

which consists of the (uniquely defined, see Observation 1) T -ancestors of the nodes in I_S . It is $|I_S| = |I_T|$, because no two nodes in I_S can have the same T -ancestor: if they did, they would overlap on this ancestor, which would give rise to an edge between them (see Observation 2), contradicting the presupposition that I_S is an independent node set. With a similar argument we obtain that I_T is an independent node set of the overlap graph O_T : if two nodes of I_T were connected by an edge in O_T , the corresponding two nodes in I_S would have to be connected (since there is a pair of overlapping ancestors, see Observation 2), again contradicting the presupposition that I_S is an independent node set. As a consequence, since I_S is arbitrary, every independent node set of O_S induces an independent node set of O_T of the same size. Therefore the maximum independent node set of O_T must be at least as large as the maximum independent node set of O_S . \square

VI. HARMFUL OVERLAP SUPPORT

The definition of MIS-support sees any overlap as harmful: at most one of two overlapping embeddings may be counted. However, there are pairs of embeddings for which the overlap is clearly harmless, so that there is no reason why they should not both be counted. An example is shown in Figure 3. Even though the embeddings overlap on the node labeled A in the middle of the input graph, they do not have an embedding of a graph with a single node labeled A as a common ancestor. Neither of the two nodes labeled A is mapped to the same node in both embeddings. Hence the two embeddings cannot be constructed from the same embedding of a single node labeled A, which is then extended in corresponding ways.

In other words: the two embeddings have to be built from two different embeddings of a single node labeled A. Therefore the fact that they overlap does not destroy the anti-monotony of the support. There were two embeddings for any ancestors of these embeddings and thus the support has always been 2. As a consequence, it can be 2 for the subgraph A-B-C-A without harming anti-monotony. This observation gives rise to a more sophisticated support definition.

Definition 8: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and f_1 and f_2 two subgraph isomorphisms of S to G . f_1 and f_2 are called *harmfully overlapping* (or *H-overlapping* for short), written $f_1 \blacklozenge f_2$, iff they are equivalent (see Definition 3) or there exists a (non-empty) proper subgraph T of S , so that the T -ancestors f'_1 and f'_2 of f_1 and f_2 , respectively, are equivalent.

Note that this definition refers to an arbitrary (non-empty) subgraph T of S . However, the search for frequent subgraphs




Fig. 4. Harmful overlap graphs of embeddings and ancestor relations.

is usually restricted to connected substructures. As a consequence, only connected ancestors can actually produce counterexamples to anti-monotony. Therefore we should restrict the definition to connected subgraphs T in this case.

The support of a subgraph can now be defined in direct analogy to maximum independent set support.

Definition 9: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and let V_H be the set of all embeddings (that is, subgraph isomorphisms) of S into G . The *harmful overlap graph* of S w.r.t. G is the graph $H = (V_H, E_H)$, which has the set V_H of embeddings as its node set and the edge set $E_H = \{(f_1, f_2) \mid f_1, f_2 \in V_H \wedge f_1 \not\equiv f_2 \wedge f_1 \blacktriangleright f_2\}$.

Definition 10: Let $H = (V_H, E_H)$ be the harmful overlap graph of the embeddings of a labeled graph $S = (V_S, E_S, l_S)$ into a labeled graph $G = (V_G, E_G, l_G)$. Then the *harmful overlap support* (or *HO-support* for short) of the graph S w.r.t. G is the size of a maximum independent node set of H .

Theorem 2: HO-support is anti-monotone.

Proof: The proof is exactly the same as for MIS-support (see Theorem 1), since Observations 1 and 2 also hold for HO-support and they were all that was needed for the proof. \square

An illustration of the relationships between embeddings of a subgraph and their ancestors is shown in Figure 4. The graphs A, B, A-B, and B-A-B are embedded into the input graph B-A-B-A-B shown at the top. The grey circles are the nodes of harmful overlap graphs for the different subgraphs and correspond to the embeddings shown to the left or to the right of them. Arrows connect ancestors to descendant embeddings (for simplicity, transitive ancestor relations, which are implied, are not shown). Note how any two nodes that are connected by an edge in the square-shaped graph at the bottom (the harmful overlap graph for B-A-B) have equivalent ancestors, while those that are not connected do not.

Note that the HO-support of the graph B-A-B w.r.t. the graph B-A-B-A-B is 2, as there are two maximum independent node sets, each consisting of the two nodes at the ends of one diagonal of the harmful overlap graph. MIS-support, however, is only 1 (recall Figure 2), as a simple overlap graph contains these diagonals. These diagonals are not present in the harmful overlap graph, because the corresponding overlaps are not harmful (there are no equivalent ancestor embeddings).

It should be clear that HO-support is never less than MIS-support, since harmful overlap is the weaker concept and thus gives rise to fewer edges in the resulting overlap graph. Therefore, if the same minimum support threshold is used, it cannot be that fewer frequent subgraphs are found with HO-support than with MIS-support. We can rather expect that the number of frequent subgraphs is larger. That the difference can actually be surprisingly large will be seen in Section IX.

VII. SUBGRAPH SUPPORT COMPUTATION

It is easy to check whether two embeddings overlap, hence computing MIS-support is fairly simple. (Note, though, that finding a maximum independent node set is NP-complete. What is easy is building the overlap graph.) However, building a harmful overlap graph is not quite as simple, since it is more difficult to check whether two embeddings overlap harmfully. Fortunately, there is still a fairly simple procedure, which is based on trying to construct a subgraph $S_E = (V_E, E_E, l_E)$ that yields equivalent ancestors of two given embeddings f_1 and f_2 . The core idea of this procedure is that for such a subgraph S_E the mapping $g : V_E \rightarrow V_E$ with $v \mapsto f_2^{-1}(f_1(v))$, where f_2^{-1} is the inverse of f_2 , must be a bijective mapping. More generally, g must be an *automorphism* of S_E , that is, a subgraph isomorphism of S_E to itself. This view of equivalent ancestors leads to the following test procedure:

Input: Two (different) embeddings f_1 and f_2 of a labeled graph $S = (V_S, E_S, l_S)$ into a labeled graph $G = (V_G, E_G, l_G)$.

Output: Whether f_1 and f_2 overlap harmfully.

- 1) Form the sets $V_1 = \{v \in V_G \mid \exists u \in V_S : v = f_1(u)\}$ and $V_2 = \{v \in V_G \mid \exists u \in V_S : v = f_2(u)\}$.
- 2) Form the sets $W_1 = \{v \in V_S \mid f_1(v) \in V_1 \cap V_2\}$ and $W_2 = \{v \in V_S \mid f_2(v) \in V_1 \cap V_2\}$.
- 3) If $V_E = W_1 \cap W_2 = \emptyset$, return *false*, otherwise return *true*.

Obviously, V_E , provided it is not empty, is the node set of a subgraph S_E of S that induces equivalent ancestors of f_1 and f_2 . On the other hand, any node $v \in V_S - V_E$ cannot contribute to such ancestors, because it is either mapped to a node on which f_1 and f_2 do not overlap (and thus $g(v)$ is not defined) or because there is no original w.r.t. g , that is, no node u with $g(u) = v$ (so that g is not bijective for v). Hence V_E is a maximal set of nodes for which g is a bijection. Therefore: if $V_E \neq \emptyset$, the embeddings f_1 and f_2 overlap harmfully and thus are connected by an edge in the harmful overlap graph. Otherwise any overlap is harmless and we need no edge.

Note that the edge set E_E of the considered subgraph S_E can be found by an analogous construction:

- 1) Form the sets $E_1 = \{(v_1, v_2) \in E_G \mid \exists (u_1, u_2) \in E_S : (v_1, v_2) = (f_1(u_1), f_1(u_2))\}$ and $E_2 = \{(v_1, v_2) \in E_G \mid \exists (u_1, u_2) \in E_S : (v_1, v_2) = (f_2(u_1), f_2(u_2))\}$.
- 2) Let $F_1 = \{(v_1, v_2) \in E_S \mid (f_1(v_1), f_1(v_2)) \in E_1 \cap E_2\}$ and $F_2 = \{(v_1, v_2) \in E_S \mid (f_2(v_1), f_2(v_2)) \in E_1 \cap E_2\}$.
- 3) Let $E_E = F_1 \cap F_2$.

This construction of E_E ensures that the mapping g as defined above is actually an automorphism of the subgraph S_E .




Fig. 5. A simple example of harmful overlap without identical images.

Note, though, that the above procedure assumes that the subgraph S_E , regardless of its structure, always induces equivalent ancestors of two embeddings f_1 and f_2 . However, if the search for frequent subgraphs is restricted to connected subgraphs, an unconnected subgraph S_E should not be seen as giving rise to equivalent ancestors, since then S_E is not considered in the search. Unfortunately, if the subgraph S_E is not connected, there may still be a connected subgraph of S_E that induces equivalent ancestors. Hence we need an extended test.

For this extended test we can exploit the following property of the bijective mapping g we defined above (reminder: $\forall v \in V_E : g(v) = f_2^{-1}(f_1(v))$): let $S_C = (V_C, E_C, l_C)$ be an arbitrary (but fixed) connected component of the subgraph S_E and let $W = \{v \in V_C \mid g(v) \in V_C\}$ (that is, let W be the set of nodes of the connected component S_C that are mapped into S_C itself). Then it is either $W = \emptyset$ or $W = V_C$. (Intuitively: a connected component is either mapped completely to itself or completely to another connected component.)

We prove this property by contradiction: suppose that there is a connected component S_C for which $W \neq \emptyset$ and $W \neq V_C$. Then we can choose two nodes $v_1 \in W$ and $v_2 \in V_C - W$. These two nodes are connected by a path in S_C , since S_C is a connected component. On this path there must be an edge (v_a, v_b) with $v_a \in W$ and $v_b \in V_C - W$. (Along the path there must be (at least) one transition from a node in W to a node not in W , since its start is in W and its end is not in W .) We know that $(v_a, v_b) \in E_E$ (since $E_C \subseteq E_E$) and thus can infer that $(g(v_a), g(v_b)) \in E_E$ (since g is an automorphism of S_E). Since $g(v_a) \in V_C$ (because $v_a \in W$), it follows $g(v_b) \in V_C$, because $g(v_b)$ is connected to $g(v_a)$. However, this implies $v_b \in W$ (due to the construction of W), contradicting $v_b \in V_C - W$, which implies $v_b \notin W$.

With this property the test whether there is a connected subgraph that induces equivalent ancestors becomes very simple: if there is such a subgraph at all, the connected component S_C of S_E containing it must, as a whole, be such a subgraph as well (otherwise we would have $W \neq \emptyset$ and $W \neq V_C$). Whether a connected component S_C is such a subgraph can be determined by checking for an arbitrary node v of S_C whether $g(v)$ is in S_C or not. If it is, then we know that $W = V_C$ and that g , restricted to the nodes of S_C , is an automorphism of S_C . Hence S_C gives rise to equivalent ancestors. If, on the other hand, $g(v)$ is not in S_C , it must be $W = \emptyset$ and neither S_C nor any of its subgraphs induce equivalent ancestors.

The test can be further optimized by the following simple insight: obviously two embeddings f_1 and f_2 overlap harmfully if $\exists v \in V_S : f_1(v) = f_2(v)$. (There are, of course, also other cases; Figure 5 shows an example. The subgraph inducing equivalent ancestors can be arbitrarily complex even if $\forall v \in V_S : f_1(v) \neq f_2(v)$.) As this test can be performed

very quickly, it should be the first step. This has the additional advantage that afterwards we can neglect isolated nodes (since the only way in which the mapping g can be an automorphism for an isolated node v is $f_1(v) = f_2(v)$), i.e., we can confine our checks to nodes of V_E that are incident to at least one edge in E_E . Therefore the optimized test for harmful overlap is:

Input: Two (different) embeddings f_1 and f_2 of a labeled graph $S = (V_S, E_S, l_S)$ into a labeled graph $G = (V_G, E_G, l_G)$.
Output: Whether f_1 and f_2 overlap harmfully.

- 1) If $\exists v \in S : f_1(v) = f_2(v)$, return *true*.
- 2) Form the edge set E_E as described above and the (reduced) node set $V_R = \{v \in V_S \mid \exists u \in V_S : (v, u) \in E_E\}$.
- 3) Let $S_C^i = (V_C^i, E_C^i)$, $1 \leq i \leq n$, be the connected components of $S_R = (V_R, E_E)$. If $\exists i; 1 \leq i \leq n : \exists v \in V_C^i : f_2^{-1}(f_1(v)) \in V_C^i$, return *true*, otherwise return *false*.

VIII. MINIMUM NUMBER OF NODE IMAGES

Only very recently another definition of the support of a subgraph in a single large graph was proposed, which is based on the minimum number of different images of a node [4].

Definition 11: Let $G = (V_G, E_G, l_G)$ and $S = (V_S, E_S, l_S)$ be two labeled graphs and let F be the set of all subgraph isomorphisms of S to G . Then the *minimum number of node images support* (or *MNI-support* for short) of S w.r.t. G is defined as $\min_{v \in V_S} |\{u \in V_G \mid \exists f \in F : f(v) = u\}|$.

As examples consider again Figures 1 and 2. In the former, the MNI-support is 1 for all shown subgraphs, because there is only one node in the input graph to which the node labeled A is mapped. In the latter, the MNI-support of the graph B-A-B is 2: even though there are three images for either node labeled B, there are only two images for the node labeled A. Hence in these two cases MNI-support coincides with HO-support. However, MNI-support can also be greater than HO-support as the example in Figure 5 demonstrates: while the HO-support of the graph A-A-B w.r.t. the graph B-A-A-B is 1, its MNI-support is 2. More generally, it is fairly easy to show that MNI-support is anti-monotone and an upper bound for HO-support (if the overlap definition is based on nodes).

Obviously, MNI-support has the highly desirable advantage that it avoids the costly maximum independent set computations needed for both MIS-support and HO-support. However, it has the disadvantage that it counts both of certain equivalent embeddings, which is certainly somewhat counterintuitive. An example is the MNI-support of the graph A-A w.r.t. the graph B-A-A-B (cf. Figure 5): since there are two possible images for both nodes labeled A, the MNI-support is 2, even though in an intuitive sense the graph A-A occurs only once.

IX. EXPERIMENTAL RESULTS

We implemented all described support definitions as part of the MoSS program², which is written in Java. This implementation supports an exact computation of a maximum

²MoSS is available for download under the Gnu Lesser (Library) Public License at <http://www.borgelt.net/moss.html>




Fig. 6. Experimental results on the IC93 data.




Fig. 7. Experimental results on the Tic-Tac-Toe win data.

independent node set of an overlap graph as well as a greedy heuristic algorithm, which only very rarely yields results that differ much from the exact results, but is considerably faster.

We applied this implementation to several data sets. Often MIS-, HO-, and MNI-support gave the same results, which, however, was mainly due to a lack of overlaps of embeddings. Two examples in which MIS-, HO-, and MNI-support exhibit clearly different behavior are the IC93 dataset [8], which we interpreted for this application as a single large graph, of which each of the 1283 molecules forms a connected component, and the Tic-Tac-Toe win dataset, which is part of the SUBDUE datasets [5] and consists of 626 connected components.

Results on these datasets are shown in Figures 6 and 7, respectively, which depict the number of frequent subgraphs (*not* restricted to closed subgraphs) over absolute minimum support. In both diagrams the solid black line refers to HO-support, the solid grey line to MIS-support, and the dashed grey line to MNI-support. In Figure 6 we also included the number of subgraphs found with standard support (number of database graphs that contain the subgraph, dotted grey line). While for the Tic-Tac-Toe dataset the increases in the numbers of frequent subgraphs are moderate (around 5%), the difference of the three support types on the IC93 dataset is surprisingly large: with HO-support, there are up to 50%, with MNI-support even up to 80% more frequent subgraphs than with MIS-support. This behavior is mainly due to heavily overlapping molecular fragments with several carbon atoms.

X. CONCLUSION

In this paper we studied how to define and compute the support of a subgraph of a single large input graph. We reviewed the MIS-support of [11] and provided a clear and simple proof that it is anti-monotone. In addition, we argued that MIS-support is sometimes too restrictive and suggested an alternative definition (HO-support), which allows embeddings to overlap in certain harmless ways. Only harmful overlaps,

which have the effect that the resulting support violates the anti-monotone condition, prevent that both of two overlapping embeddings are counted for the support of a subgraph. Our experiments show that HO-support can sometimes lead to considerably higher support values than MIS-support.

An important point is that HO-support is a kind of upper bound for the support of a subgraph in a single large graph if the support is to be anti-monotone and if of two equivalent embeddings at most one should be counted (otherwise MNI-support is a larger bound). It only prevents that two embeddings are both counted if they can result from extending equivalent ancestor embeddings (of which at most one was counted). Although there are certain exceptions, these are degenerate cases, which, as it seems to us, are safe to neglect.

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. On Canonical Forms for Frequent Graph Mining. *Proc. 3rd Int. Workshop on Mining Graphs, Trees and Sequences (MGTS'05, Porto, Portugal)*, 1–12. ECML/PKDD 2005 Organization Committee, Porto, Portugal 2005
- [3] 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
- [4] B. Bringmann and S. Nijssen. What is Frequent in a Single Graph? *Proc. 5th Int. Workshop on Mining and Learning with Graphs (MLG 2007, Florence, Italy)*, 183–186. University of Florence, Italy 2007
- [5] D.J. Cook and L.B. Holder. SUBDUE datasets. <http://cygnus.uta.edu/subdue/download.htm>
- [6] 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
- [7] 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
- [8] *Index Chemicus — Subset from 1993*. Institute of Scientific Information, Inc. (ISI). Thomson Scientific, Philadelphia, PA, USA 1993 <http://www.thomsonscientific.com/products/indexchemicus/>
- [9] 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
- [10] 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
- [11] M. Kuramochi and G. Karypis. Finding Frequent Patterns in a Large Sparse Graph. *Proc. 4th SIAM Int. Conf. on Data Mining (SDM 2004, Lake Buena Vista, FL)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA 2004
- [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] N. Vanetik, E. Gudes, and S.E. Shimony. Computing Frequent Graph Patterns from Semistructured Data. *Proc. IEEE Int. Conf. on Data Mining (ICDM 2002, Maebashi, Japan)*, 458–465. IEEE Press, Piscataway, NJ, USA 2002
- [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