

Subsea: an efficient heuristic algorithm for subgraph isomorphism

V. Lipets · N. Vanetik · E. Gudes

Received: 14 November 2007 / Accepted: 12 May 2009 / Published online: 29 May 2009
Springer Science+Business Media, LLC 2009

Abstract We present a novel approach to the problem of finding all subgraphs and induced subgraphs of a (target) graph which are isomorphic to another (pattern) graph. To attain efficiency we use a special representation of the pattern graph. We also combine our search algorithm with some known bisection algorithms. Experimental comparison with other algorithms was performed on several types of graphs. The comparison results suggest that the approach provided here is most effective when all instances of a subgraph need to be found.

Keywords Graph algorithms · Subgraph isomorphism · Heuristic · Bisection

1 Introduction

Subgraph isomorphism is an important and very general form of pattern matching that finds practical application in areas such as pattern recognition and computer vision, computer-aided design, image processing, graph grammars, graph transformation, bio computing, search operation in chemical structural formulae database, and numerous others. Theoretically, subgraph isomorphism is a common generalization of many important graph problems including finding Hamiltonian paths, cliques, matching,

Responsible editor: Charles Elkan.

V. Lipets · N. Vanetik (✉) · E. Gudes
Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel
e-mail: orlovn@cs.bgu.ac.il

V. Lipets
e-mail: lipets@cs.bgu.ac.il

E. Gudes
e-mail: ehud@cs.bgu.ac.il

girth, and shortest paths. Subgraph isomorphism is also a basic component of any graph mining algorithm, an area which has become important in many data mining applications of unstructured and complex data. Graph mining algorithms often require finding not one but all subgraphs of the database graph isomorphic to a given small graph in order to compute the measure of statistical significance (also called 'support') of that small graph in the database. The subgraph isomorphism problem is generally NP-complete (see [Garey and Johnson 1979](#)) and computationally difficult to solve. But because of its wide applicability much effort was invested in finding algorithms and heuristics which can reduce the total search effort. Since the number of subgraphs isomorphic to a given small graph in a graph database can also be exponential in the size of the database, support computation is the bottleneck of all graph mining algorithms. Graph mining is an important and fast growing field of data mining with many applications such as mining biochemical structures, program flow analysis, mining XML structures and Web communities.

The most common technique to establish a subgraph isomorphism is based on backtracking in a search tree. In order to prevent the search tree from growing unnecessarily large, different refinement procedures are used. Best known is the one by Ullman (see [Ullmann 1976](#)). Cordella et al. have suggested another general graph matching algorithm in [Sansone et al. \(2004\)](#). This algorithm is tailored for dealing with large graphs without using information about the topology of the graphs to be matched. It employs a recursive algorithm which grows a set of partial sub-graphs until the isomorphic sub-graph is found. The growing is done in such a way that reduces considerably the search space by providing pruning rules, and a dynamic ordering method for node matching. (The idea of pruning also appears in our algorithm as will be seen later.) The algorithm has become quite popular and the authors supply executables which can be used for evaluation ([Cordella et al. 1999](#)).

On the other hand, a lot of work has been done to solve this problem in polynomial time for specific families of graphs, for example papers [Dessmark et al. \(2000\)](#), [Eppstein \(1999\)](#), [Lingas and Sysło \(1989\)](#) and [Matula \(1978\)](#). [Alon et al. \(1995\)](#) has presented a novel randomized method—color-coding—for finding simple paths and cycles of a specified length k , and other small subgraphs, within a given graph G .

The essential part of the recent research on subgraph isomorphism algorithms has been based on heuristic search techniques and was described in [Akinniyi et al. \(1986\)](#), [Cheng and Huang \(1981\)](#), [Cortadella and Valiente \(2000\)](#), [Larrosa and Valiente \(2000\)](#), [Messmer and Bunke \(1996\)](#). Some of the techniques work better in some special cases. For example, [Akinniyi et al. \(1986\)](#) use a tree searching procedure over the projections of the implicit product of the two graphs and utilize the minimum number of neighbors of the projected graphs to detect infeasible subtrees. [Cheng and Huang \(1981\)](#) use bitwise parallelism during the resolution process even though a sequential computer is used. [Cortadella and Valiente \(2000\)](#) use a representation of relations and graphs by Boolean functions, which allows handling the combinatorial explosion in the case of small pattern graphs and large target graphs. [Larrosa and Valiente \(2000\)](#) use the notion of neighborhood constraints. In particular, an efficient algorithm is described by [Foggia et al. \(2001\)](#). This algorithm uses a set of feasibility rules to significantly reduce the computational cost of the matching process.

Two additional algorithms have been published recently. [Krissinel and Henrick \(2004\)](#) describe an algorithm that recursively enumerates all possible mappings of subgraphs of the two graphs. They improve the computational complexity of the algorithm by limiting the search area for possible matching. This is done by maintaining a static vertex matching matrix (VMM). The VMM stores all previous matchings, which the algorithm uses to determine if a certain mapping can or succeeded to cannot obtain a result with a sufficient size. If the mapping is present in the matrix, and is known to produce a result that is too small, the algorithm will simply skip calculating with that mapping and continue to the next one, thus saving computing the entire recursive branch.

[Batz \(2006\)](#) describes a heuristic algorithm for finding a graph-subgraph mapping on directed labeled graphs, based on building a “plan graph” that helps model all available mappings and finds the mapping that takes the shortest time to build. First, a “search plan” is defined as an iterative mapping between two graphs built by atomic (“primitive”) operations. on the primitive operations is defined, based on information from the host graph. This paper gives a heuristic algorithm of finding an optimal selection and ordering of the “search plan” that will give a matching between the two graphs. According to the analysis done by this paper, the algorithm runs in almost linear time when very few labels are repeated. However, cases where the graph is not labeled, or is labeled poorly (few different labels for many components) cause the algorithm to find more expensive search plans which take more time to calculate.

There are other variations of the subgraph isomorphism problem. An algorithm for subgraph isomorphism detection from a set of a priori known model graphs to an input graph that is given on-line is presented in [Bunke and Messmer \(2000\)](#). The approach is based on a compact representation of the model graphs that is computed offline. A related problem is that of *Graph matching* where exact isomorphism is not required but some similarity measure is defined. One such measure is based on the “edit distance” [Berretti et al. \(2004\)](#), but other measures are possible. A recent paper by [Sammoud et al. \(2005\)](#) describes an algorithm for finding an approximation to the maximal subgraph matching between two graphs. First, the authors define a function that measures the similarity of two graphs with respect to a certain mapping, and a function that measures the “desirability” of a single vertex matching for maximizing the similarity factor (a score function). They use these functions to iteratively build the best mapping (or close to it). mapping using a greedy algorithm. This greedy algorithm works by choosing the most desirable vertex matching available to it and adding it to its original graph mapping. The algorithm ends after the score function of the mapping no longer increases, or after a maximal number of iterations is reached.

Also, in some papers it is proposed to go a step further by introducing multi-vertex matchings, where a vertex in one graph may be matched with a set of vertices of the other graph [Ambauen et al. \(2003\)](#), [Boeres et al. \(2004\)](#), [Champin and Solnon \(2003\)](#).

Recently, the area of *graph mining* and especially finding frequent graph patterns, has become an important research area with applications in mining and discovering frequent patterns in semi-structured data such as Web and XML data [Chen et al. \(1998\)](#), [Wang and Liu \(1998\)](#), [Lin et al. \(1998\)](#), chemical compounds data such as [Dehaspe et al. \(1998\)](#), [Nijssen and Kok \(2004a,b\)](#), or biological data [Pennec and Ayache \(1998\)](#). All the algorithms for frequent graph mining require heavy use of

subgraph isomorphism as a basic step; therefore, their performance is critically dependent on the efficiency of the subgraph isomorphism algorithm. Examples of graph mining algorithms which make heavy use of subgraph isomorphism can be found in [Kuramochi and Karypis \(2001\)](#), [Vanetik et al. \(2002\)](#) and [Yan and Han \(2002\)](#).

Recall that the subgraph isomorphism problem is NP-complete in general. However, for any target graph with n vertices and any pattern graph with k vertices, where k is fixed, both the enumeration and decision problems can easily be solved in polynomial $O(n^k)$ time, and for some patterns an even better bound might be possible. However, for the general subgraph isomorphism problem, nothing better than the naive $O(n^k)$ bound is known. Thus one is led to the problem of finding heuristics for reducing the expected complexity of a subgraph isomorphism algorithm for a given fixed pattern graph.

In this paper, we propose an efficient heuristic method for finding all subgraphs and induced subgraphs of a target graph which are isomorphic to another connected pattern graph of a fixed size. The proposed algorithm is presented for non-directed labeled graphs but it has general validity, since no constraints are imposed on the topology of the pattern and the target graphs, and the method can be easily extended to the directed graphs case as well. To attain the efficiency we use a special representation of the pattern graph, which gives preference to checking the more “dense” parts first, thus obtaining “negative” answers as early as possible. We also combine our search algorithm with some known bisection algorithms and use them recursively which eases the search in large target graphs.

A special property of our algorithm is the fact that it finds *all* instances of the subgraph isomorphism and not just the first one. Our motivation was graph mining in the *single* graph-setting case, see e.g., [Kuramochi and Karypis \(2004\)](#), [Gudes et al. \(2006\)](#). The presented algorithm is particularly useful when the target graph is much larger than the pattern graph, a case which is common in many applications, such as mining the Web or Social networks [Cai et al. \(2005\)](#). Recently the interest in mining all occurrences in a single large graph has increased. An example biological application was presented by [Zhang et al. \(2009\)](#). The following is a citation from [Zhang et al. \(2009\)](#): “For example, biological networks (protein-protein interaction networks, and gene regulatory networks, etc.) are often much larger than the graphs used in previous exact indexing methods. A single biological network may contain thousands or tens of thousands of vertices. Biologists may want to find all the occurrences of a particular pattern (subgraph), e.g., protein type A interacts with protein type B and C, and protein type C interacts with protein type A and D. In different occurrences of the pattern, the exact proteins involved may be different since multiple proteins may share the same type. As a result, all occurrences of a particular pattern need to be retrieved.”

To evaluate our approach we implemented our algorithm and compared it to one of the most popular subgraph isomorphism algorithms, the one by [Ullmann \(1976\)](#), and to the algorithm of Cordella et al. [Sansone et al. \(2004\)](#), using latest publicly available implementations at [Sansone \(2001\)](#). Given a graph and a database graph, both of these algorithms are capable of finding and reporting all subgraph isomorphisms the given graph and the database graph, either as induced graph or not. The main conclusions of the evaluation is that our Subsea algorithm outperforms both these algorithms in finding all the instances of a subgraph in a graph, while finding a first instance of a

subgraph in a graph may be slower than either one of them. Thus Subsea is particularly attractive for graph mining in a single transaction setting case Kuramochi and Karypis (2004), Zhang et al. (2009).

The rest of this paper is structured as follows: Sect. 2 is devoted To definitions and notations, and also contains the high-level outline of the algorithm. Sections 3, 4, and 5 present the necessary auxiliary algorithms; in Sect. 3 we describe some well known bisection algorithms, in Sect. 4 we give some heuristic methods to represent a “small” pattern graph, and Sect. 5 describes our search strategy. In sect. 6 we present our main algorithm to find all occurrences of the given subgraph in the target graph and which uses the above auxiliary algorithms. In Sect. 7, we present the experimental comparison of our algorithm with the algorithms presented in Ullmann (1976) and Sansone et al. (2004), and we show the effectiveness of the proposed heuristic method. We conclude and propose further research in Sect. 8.

2 Definition and notations

2.1 General definitions

A graph $G = (V, E)$ is called *vertex-labeled* (or simply *labeled*) if a mapping $l : V \rightarrow \mathbb{N}$ is given. $l(v)$ is called a label of a vertex v . Throughout this paper we always speak of a labeled graph unless otherwise specified. Two graphs which contain the same number of vertices with the same labels connected in the same way are said to be isomorphic. Formally, two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ are *isomorphic*, denoted by $G_1 \cong G_2$, if there is a (label-preserving) bijection $\varphi : V_1 \rightarrow V_2$ such that, for every pair of vertices $v_i, v_j \in V_1$, $(v_i, v_j) \in E_1$ if and only if $(\varphi(v_i), \varphi(v_j)) \in E_2$. Bijection φ is said to be an *isomorphism* between two graphs. An isomorphism which maps $v \in V$ into $v' \in V$ we denote by $(v \rightarrow v')$ -isomorphism.

A graph G' is a subgraph of a given graph G if vertices and edges of G' form subsets of the vertices and edges of G . If G' is a subgraph of G , then G is said to be a *supergraph* of G' . We denote this relationship by $G' \subseteq G$.

A graph $G_1 = (V_1, E_1)$ is *isomorphic to a subgraph* of a graph $G_2 = (V_2, E_2)$ if there exists a subgraph of G_2 , say G'_2 , such that $G_1 \cong G'_2$. In this case the corresponding bijection between vertices of G_1 and G'_2 is said to be a *subisomorphism* between two graphs. Note that for labeled graphs, this bijection must be label-preserving.

An *induced subgraph* is a subset of the vertices of a graph G together with all edges whose endpoints are both in this subset. Formally, let G be a graph and $V' \subset V(G)$. We call the graph $G' = (V', E(G) \cap \{(u, v) | u, v \in V'\})$ the *subgraph of G induced by V'* and we denote it by $G(V')$. The relationship between G' and G in this case we denote by \sqsubseteq .

Thus, an induced subgraph isomorphism is an isomorphism with an induced subgraph of a given graph, i.e., a graph $G_1 = (V_1, E_1)$ is isomorphic to an induced subgraph of a graph $G_2 = (V_2, E_2)$ if there exists an induced subgraph of G_2 , say G'_2 , such that $G_1 \cong G'_2$. In this case a corresponding bijection between vertices of G_1 and G'_2 is said to be an *induced subisomorphism* between two graphs.

Note that for induced subgraph isomorphism all edges of G_1 must be mapped into edges of G'_2 , but also all non-edges of G_1 must be mapped to non-edges of G'_2 . Thus without loss of generality we assume that the number of non-edges of the pattern graph is larger than the number of edges (for the induced subgraph isomorphism problem only), otherwise we can solve the inverse graph problem (recall that the inverse graph is obtained by replacing non-edges by edges). An *automorphism* of a graph $G = (V, E)$ is a graph isomorphism with itself.

The neighborhood of a vertex v in graph G , denoted by $N_G(v)$, is the set of vertices in G that are adjacent to v , i.e., $N_G(v) = \{u \in V \mid (u, v) \in E\}$. For any $e \in E(G)$, we define $G - e = (V(G), E(G) \setminus \{e\})$.

A *cut* in G is a partition of V into two disjoint sets, say (A, \bar{A}) . For $A \subseteq V$, denote by $e(A, \bar{A})$ the set of all $A - \bar{A}$ edges:

$$e(A, \bar{A}) = \{(u, v) \in E \mid u \in A, v \notin A\}.$$

The *size* of the cut (A, \bar{A}) , denoted by $c(A, \bar{A})$, is the number of edges having exactly one vertex in A and the other in \bar{A} , namely $|e(A, \bar{A})|$.

The *minimum bisection* is a cut (A, \bar{A}) minimizing $c(A, \bar{A})$ over all sets with A of size $\lceil |V|/2 \rceil$. For arbitrary graphs G , the problem of determining the minimum bisection is NP-hard [Garey and Johnson \(1979\)](#).

For any graph $G = (V, E)$ we assume that there exists a total order on the set of vertices V which enables us to compare any two vertices formally for each $u \neq v \in V$ either $v < u$ or $u < v$. We use the notation $w = \langle w_1, w_2, \dots, w_k \rangle$ to represent a sequence of k elements and denote the i -th element by $w[i]$ for all $i \in \{1, 2, \dots, k\}$. Two sequences of numbers can be compared according to the lexicographical order.

In order to test bisection algorithms, we need a way to generate graphs with known best bisections. We use the following graph generation methods:

The probability distribution $G_{n,p,r}$ is commonly used in the literature on minimum bisection because the minimum bisection of a generated graph is known with high probability, and the simplicity of the distribution leads to the possibility of determining analytically the probability of success of some algorithms.

$G_{n,p,r}$ is a probability distribution on graphs with vertex set $\{1, 2, \dots, n\}$ in which the presence of each possible edge is independent, with probability p for edges within $\{1, 2, \dots, n/2\}$ or $\{n/2 + 1, \dots, n\}$ and probability $r < p$ for other edges. If $p - r$ is sufficiently large, then with high probability, the minimum bisection is the obvious one.

Throughout this paper we use notations of “*small*” and “*large*” graphs, denoted by G_S and G_L respectively. We assume that $|E(G_S)| \ll |E(G_L)|$, where $E(G_S)$ is the pattern graph and $E(G_L)$ is the target graph. The number of edges of a “*large*” graph was usually three or more orders of magnitude more than that of the “*small*” graph.

2.2 Outline of the Subsea algorithm

The algorithm, called **Subsea**, relies on the fact that the target graph is much larger than the pattern graph. Therefore it uses *bisection* to decompose the “*large*” graph and search for isomorphisms in the bisected parts (recursively). Clearly, if the bisection

is minimal the amount of time which is spent on checking edges belonging to the bisection is also small, and the work can proceed on the two decomposed parts which are considerably smaller. This, of course, relies on the fact that the target graphs are not highly connected and that the bisection is not very large, something which is true in most of the applications we mentioned in the Introduction. The isomorphism check itself uses a canonical numbering scheme called *traverse history* constructed in such a way that the isomorphism check will find non-matching subgraphs as early as possible. Since the isomorphism check can start at any pair of adjacent vertices, a traverse history may be generated at most twice for every edge in the pattern graph and, in the worst case, $2e$ traverse histories are generated for e edges.

The main steps of the algorithm are as follows:

1. In a preprocessing step generate all the traverse histories of the pattern graph.
2. Decompose the target graph by finding an approximate minimum bisection (heuristically).
3. Check all possible isomorphisms using edges belonging to the bisection.
4. Apply step 2 again recursively on the two parts of the bisected graph until the target graph becomes comparable in size to the “small” graph.

Section 3 describes steps 1 and 2. Section 4 describes the isomorphism check using traverse histories. Section 5 describes the subgraph search algorithm and Sect. 6 discusses the entire algorithm.

There are two main reasons why Subsea performs as it does. First, traverse histories for the small graph are generated once in the pre-processing stage of the algorithm and the small graph is never traversed later during the search. We ‘pay’ in running time for that, but since the number of database subgraphs isomorphic to given graph can be exponential in the size of the database, a lot of running time is saved later. This is also the reason our approach has shown an advantage over dynamic re-ordering implemented in the algorithm of Sansone et al. (2004). The second reason is that bisecting the large graph each time decreases search space dramatically each time it is applied.

3 Bisection algorithms

In this section we describe two well-known approximation algorithms for finding a minimum bisection of a given graph G . Note that since the minimum bisection is only a tool to decompose the problem efficiently, we are not required to find the actual minimum bisection (which is a hard problem), but it is enough to provide an approximation for it.

We ignore vertex labels in this construction, because minimum bisection size depends only on the number of edges and not on vertex labels.

3.1 Black holes bisection algorithm

This simple method which we devised works well on graphs with a very thin minimum bisection, and has the advantage of running very fast.

Given a graph $G = (V, E)$, the algorithm runs as follows. Initialize $B_1 = B_2 = \emptyset$. These are the black holes. Choose uniformly at random an edge from $V \setminus (B_1 \cup B_2)$ to B_1 , and add the first endpoint to B_1 . If no such edge exists, choose uniformly at random among all vertices in $V \setminus (B_1 \cup B_2)$ for a vertex to add to B_1 . Do the same for B_2 . Repeat until $|B_1 \cup B_2| = |V|$ (See Algorithm 3.1).

The motivation for this algorithm is that, assuming that the black holes are currently contained in opposite sides of a minimal bisection, we are likely to add to each hole a vertex from the correct side because there will be more edges from this side.

Algorithm 3.1 BLACK HOLES BISECTION

Input: Graph $G = (V, E)$

Output: Cut (B, \bar{B}) of V which approximates a minimal bisection

- 1: $B_1 \leftarrow B_2 \leftarrow \emptyset$
- 2: $B_0 \leftarrow V \setminus (B_1 \cup B_2)$
- 3: **repeat**
- 4: *Add2Hole*(1)
- 5: *Add2Hole*(2)
- 6: **until** $B_0 = \emptyset$
- 7: **return** (B_1, \bar{B}_1)

Algorithm 3.1 :: Procedure *Add2Hole*(i)

- 1: **if** $B_0 = \emptyset$ **return**
- 2: $E_0 \leftarrow \{(u, v) : u \in B_i, v \in B_0\}$
- 3: **if** $E_0 \neq \emptyset$ **then**
- 4: **choose randomly** $e = (u, v) \in E_0$ with $v \in B_0$
- 5: **else**
- 6: **choose randomly** $v \in B_0$
- 7: **end if**
- 8: $B_i \leftarrow B_i \cup \{v\}$
- 9: $B_0 \leftarrow B_0 \setminus \{v\}$

Let G be a graph chosen at random according to the distribution $G_{n,p,r}$. The probability that randomized black holes will find the minimum bisection on G is $(p/p + r)^{n-2}/2$.

In fact, in the first step, we will simply be picking two vertices uniformly at random; the probability that they will lie on opposite sides of the canonical bisection is $\frac{1}{2}$ (or close to it). Assume that after k steps $H_1 \in P_1$ and $H_2 \in P_2$, where P_1 and P_2 are the two halves of the canonical bisection. The expected number of edges from H_1 to $P_1 \setminus H_1$ is $pk(n/2 - k)$, while the expected number from H_2 to $P_2 \setminus H_2$ is $rk(n/2 - k)$. Hence the probability that we pick an edge from the correct side is $p/(p + r)$. Since we must make $n - 2$ such correct choices, our overall probability of success is as stated.

Obviously this success rate becomes quite bad as n becomes large if p and r are constant. But consider the case where $r = c/n$ for some constant c , or equivalently

where the number of edges across the minimum bisection is $O(n)$. Now the probability of success is greater than

$$\frac{1}{2} \left(1 - \frac{r}{p+r} \right)^n > \frac{1}{2} \left(1 - \frac{r}{p} \right)^n = \frac{1}{2} \left(1 - \frac{c}{np} \right)^n$$

This last expression approaches the non-zero constant, say q , as n approaches infinity. Hence by running randomized black holes $1/q$ times, a number depending only on c and p , we achieve a success rate of at least $1/e$.

3.2 Simple greedy bisection method

The obvious greedy algorithm for the graph bisection problem consists of starting with any bisection (B, \bar{B}) of V (in particular, we can use a bisection obtained by Black Holes algorithm) and computing a new bisection by swapping the pair of elements $x \in B$, $y \in \bar{B}$ which maximizes the gain (number of edges in (B, \bar{B}) before the swap minus the number after the swap). This process is repeated until the maximum gain is less than zero or until the maximum gain is zero and another heuristic has determined that it is time to stop swapping “zero gain” pairs. The issue of breaking ties is resolved by choosing the pair to swap uniformly at random from the set of pairs for which the gain is maximum.


To efficiently determine the cost of a swap, we store for each vertex $x \in B$ its internal cost $I(x)$, which is the number of edges $(x, x') \in E$ with $x' \in B$, and its external cost $E(x)$, which is the number of edges $(x, y) \in E$ with $y \in \bar{B}$. This can be done in $O(m+n)$ time, and can be updated after swapping x and y in $O(|N_G(x)| + |N_G(y)|)$ time. The gain in swapping x and y is then $gain = E(x) - I(x) + E(y) - I(y) - 2w(x, y)$, where $w(x, y) = 1$ if $(x, y) \in E$ and $w(x, y) = 0$ otherwise. We can therefore determine the best pair to swap in $O(n^2)$ time by simply running through all $(n/2)^2$ candidate pairs, keeping track of the leaders for the random choice at the end.

4 Traverse history

In this section we give two heuristic methods to represent a “small” pattern graph. Each such representation enumerates the vertices of the pattern graph in a particular order. This order will determine the order in which the isomorphism check is done. In the experiments that we carried out, these simple methods essentially improved our search technique by reducing the run time of the corresponding search algorithm (when compared to the case when some induced traverse history was chosen randomly). The motivation for using these representations will become clearer in Sect. 5.

Let $d : V \rightarrow \mathbf{N}$ be a numbering of vertices of graph G . Let l_i denote the label of the vertex that has number i in numbering d , i.e., $l_i := l(v)$, $d(v) = i$; let $N_i := \{d(u) < i : u \in N_G(v), d(v) = i\}$. The sequence $(l_1, N_1), (l_2, N_2), \dots, (l_{|V|}, N_{|V|})$ is called a *traverse history* of graph G induced by numbering d . In particular, we say that traverse history started on vertices v_1, v_2, \dots, v_i , if $d(v_i) = i$, where $i \leq |V|$. Informally, N_i is the set of adjacent vertices to i with numbering smaller than i . The

Fig. 1 Traverse history of graph G (labels omitted)



traverse history $\langle (l_1, N_1), (l_2, N_2), \dots, (l_{|V|}, N_{|V|}) \rangle$ is *connected* if $N_i \neq \emptyset$ for each $2 \leq i \leq |V|$. See Fig. 1 for an example of N_i on given graph. For traverse history H , $H[i]$ denotes (l_i, N_i) , $H[i].l$ denotes l_i and $H[i].N$ denotes N_i .

We have two heuristic methods, Algorithm 4.1—the DFS method, and 4.2—the “black holes” method, for finding a traverse history for “small” pattern graph with given starting vertices. Algorithm 6.1 finds such traversals for all pairs of starting with adjacent vertices. This is needed since the starting edge of the isomorphism check may be arbitrary. The traverses obtained in this stage are used then as an input for the search algorithm which is described in Sect. 5.

4.1 The DFS approach

Our first approach is based on a modification of the well known Depth-First Search (DFS) algorithm which provides a general technique for traversing a graph (traversing a graph means visiting all of its vertices in some systematic order). Recall that the DFS traversing is not deterministic, i.e., for any graph G a number of traverses is possible. We extend the traversing strategy by some heuristic rules, to provide a “fastest” return to the visited nodes.

In Algorithm 4.1 we present our modification of the *DFS* algorithm. This algorithm produces a traverse history of a connected graph G starting at given vertices v_1, v_2 . The same as the DFS algorithm, Algorithm 4.1 visits all vertices of G in some schematic order, thereby producing the numbering of vertices of $V(G)$. Note that the obtained traverse history is induced by the above numbering.

Algorithm 4.1 TRAVERSE HISTORY

Input: Graph $G = (V, E)$, starting vertices $v_1, v_2 \in V$, with $(v_1, v_2) \in E$

Output: Traverse history H started on $v_1, v_2 \in V$.

- 1: **for** all $v \in V$ **do**
 - 2: $d(v) \leftarrow 0$
 - 3: **end for**
-

Algorithm 4.1 CONTINUED

```

4:  $vtime \leftarrow etime \leftarrow 1$ 
5:  $Visit(v_1)$ 
6: return  $H$ 

```

Algorithm 4.1 :: Procedure $Visit(v)$

```

1:  $d(v) \leftarrow vtime$ 
2:  $H[vtime ++] = (l(v), \{0 < d(u_1) \leq \dots \leq d(u_m) : u_1, \dots, u_m \in N_G(v)\})$ 
3: if  $v = v_1$  then  $Visit(v_2)$ 
4:  $N_0 \leftarrow \{u \in N_G(v) : d(u) = 0\}$ 
5: while  $N_0 \neq \emptyset$  do
6:   choose  $w \in N_0$  with lexicographically minimal  $EstimateNext(w, v)$  pair
7:   if  $d(w) = 0$  then  $Visit(w)$ 
8:    $N_0 \leftarrow N_0 \setminus \{w\}$ 
9: end while

```

Algorithm 4.1 :: Procedure $EstimateNext(w, v)$

```

1:  $S \leftarrow \{w\}$ 
2:  $len \leftarrow 1$ 
3: repeat
4:    $N_S \leftarrow \cup_{z \in S} N_{G-(v,w)}(z)$ 
5:    $p = |\{y \in N_S : d(y) > 0\}|$ 
6:   if  $p > 0$  then return  $\langle len, -p \rangle$ 
7:    $len ++$ 
8:    $S \leftarrow S \cup N_S$ 
9: until  $N_S = \emptyset$ 
10: return  $\langle \infty, -|S| \rangle$ 

```

The main goal of this approach is the choice of the next vertex to visit. To implement our strategy, we define Breadth first search (BFS) based procedure $EstimateNext$, which helps to choose a next vertex to visit, by finding for any unvisited vertex (adjacent to the last visited vertex) an estimation pair using the following heuristic method, which prefers to visit first nodes which have high “proximity” to the current node: Let v be the last visited vertex and let $w \in N_G(v)$, be an unvisited vertex adjacent to v . We distinguish between two cases (illustrated in Fig. 2):

- (a) there exists a path in $G - (v, w)$ started in w and leading to a visited vertex,
- (b) such a path does not exist.

The first coordinate of the estimation pair is at least 1 in case (a) and equal to ∞ in case (b). This provides highest priority to case (a) and lowest to (b), since the minimal value returned by $EstimateNext$ procedure is chosen by the lexicographical order. Actually, in case (a) the first coordinate is the length of the shortest path (paths) leading to the visited vertices and the second coordinate equals the number of visited vertices

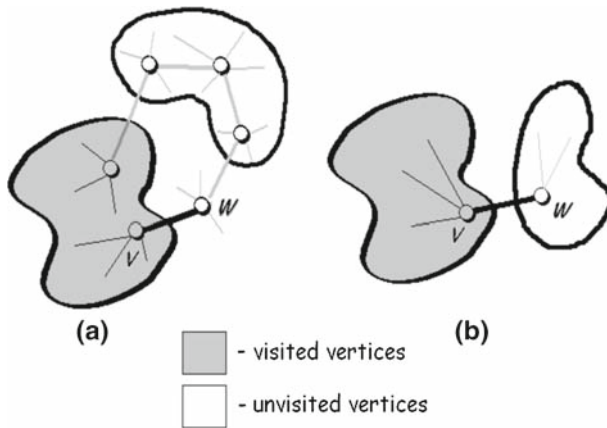


Fig. 2 Estimation cases

that the above path (paths) is leading to (multiplied by -1). This forces the choice of the next vertex which belongs to the shortest path (paths) leading to the maximal number of visited vertices. In case (b) the second coordinate equals to the size of the connected component of $G - (v, w)$ containing w (multiplied by -1). This forces the choice of the next vertex which belongs to the smallest connected non-visited component, when the above path leading to the visited vertices does not exist.

Remark In fact, the run time $O((|E| + |V|)^2)$ of Algorithm 4.1 can be essentially improved. But since we apply this algorithm on the “small” graph only at the preprocessing stage, it is not necessary, and we present the simplest version here.

4.2 Black holes approach

In this section we give another approach for finding induced traverse history for “small” pattern graph. In the experiments that we carried out, this method essentially improved the run time when the main algorithm was applied on randomly generated graphs.

For any traverse history $H = \langle (l_1, N_1), (l_2, N_2), \dots, (l_{|V|}, N_{|V|}) \rangle$ of graph G we define a *traverse integrality* as follows:

$$I_G(H) = \sum_{i=1}^{|V|} (|V| + 1 - i) |N_i|.$$

Actually this value gives a measure about the number of edges of subgraphs of some given graph G induced by the first $i = 1, 2, \dots, V(G)$ vertices of any numbering d (Fig. 3).

We provide a simple (and very fast) randomized method for finding the induced traverse history with the largest (or the smallest) traverse integrality. This method is very similar to the Black Holes Bisection algorithm, presented in Sect. 3.

Given a graph $G = (V, E)$ and starting vertices $v_1, v_2 \in V$, the algorithm runs as follows: Initialize $B = \{v_1, v_2\}$. This is the black hole. Choose uniformly at random an




Fig. 3 Traverse integrality

edge from $V \setminus B$ to B , and add the first endpoint to B (in the Case when we are seeking the smallest integrality, choose Randomly a vertex in $\{v \in \bar{B} : N_G(v) \cap B \neq \emptyset\}$ and add it to B). Enumerate v and for each vertex in $N_G(v) \cap B$ add $d(u)$ to $N_{|d(v)|}$. Repeat until $|B| = |V|$.

The motivation for this algorithm is the same as for the bisection algorithm: assuming that the black holes are currently contained in B , we are likely to add to a hole a vertex with the largest number of neighbors in B , thereby increasing the corresponding $|N_i|$ value.

Algorithm 4.2 HOLE INTEGRATION

Input: Graph $G = (V, E)$, starting vertices $v_1, v_2 \in V, (v_1, v_2) \in E$

Output: Traverse history H started on $v_1, v_2 \in V$.

- 1: **for** all $v \in V$ **do**
 - 2: $d[v] \leftarrow 0$
 - 3: **end for**
 - 4: $d(v_1) = 1$
 - 5: $d(v_2) = 2$
 - 6: $B \leftarrow \{v_1, v_2\}$
 - 7: **while** $|B| \neq |V|$ **do**
 - 8: $Add2HoleAndVisit()$
 - 9: $I \leftarrow I + etime$
 - 10: **end while**
 - 11: **return** (A_H, I)
-

Algorithm 4.2 :: Procedure *Add2HoleAndVisit*

- 1: **choose randomly** $e = (u, v) \in e(B, \bar{B})$, where $v \in \bar{B}$
 - 2: $d(v) = vtime$
 - 3: $H[vtime ++] = (l(v), \{d(u) : u \in N_G(v) \cap B\})$
-

Note that since initially $B = v_1, v_2$ and each vertex added to B has a neighbor in B , the sequence of the set obtained by Algorithm 4.2 is connected traverse history of graph G started on vertices $v_1, v_2 \in V$.

Remark The problem of finding traverse history with largest integrality is not well studied. We suppose that it is NP-complete, and we will try to prove this fact in the future. In our approach only the “small” graph is used as an input for this algorithm, thus we can apply it a sufficiently large number of times to attain a high probability that largest (smallest) integrality is found.

5 Search technique


In this section we start with an algorithm which searches for the isomorphic subgraphs by given traverse history. In fact, the above algorithm seeks for subgraphs satisfying the condition of the following obvious lemma.

Lemma 5.1 *Let $H_1 = \langle (l_1, N_1), (l_2, N_2), \dots, (l_{|V_1|}, N_{|V_1|}) \rangle$ be a traverse history of graph $G_1 = (V_1, E_1)$ with labeling l , induced by numbering d_1 of V_1 . A subgraph (induced subgraph) G'_2 of a graph $G_2 = (V_2, E_2)$ with labeling m is isomorphic to G_1 if $|V_2| \geq |V_1|$ and there exists $H_2 = \langle (m_1, M_1), (m_2, M_2), \dots, (m_{|V_2|}, M_{|V_2|}) \rangle$, a traverse history of graph G_2 induced by some numbering d_2 of V_2 such that $l_i = m_i$ and $N_i \subseteq M_i (N_i = M_i)$ for each $1 \leq i \leq |V_1|$. Moreover $d_2^{-1} \circ d_1$ is an isomorphism between G_1 and $G'_2 \subseteq G_2 (G'_2 \subseteq G_2)$.*

In particular, graph G_1 is $(v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -isomorphic to subgraph (induced subgraph) of G_2 if $v_1, v_2 \in V_1$ and $v'_1, v'_2 \in V_2$ are starting vertices of H_1 and H_2 , respectively. \square

Our search technique is based on the DFS method and is depicted in Algorithm 5.2. As an input it receives “large” target graph $G_L = (V_L, E_L)$, starting vertices $v_1, v_2 \in V_L$ and traverse history H of a “small” pattern graph G_S , previously obtained by Algorithms 4.1 or 4.2. It finds all subgraphs (induced subgraphs) of $G_L (v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -isomorphic to G_S , where v'_1, v'_2 are the first two vertices of the traverse history H . An example of such an isomorphism is given in Fig. 4. All the labels are assumed to be the same in this example and are therefore omitted from traverse histories. Here, G_S is $(v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -isomorphic to two subgraphs of G_L , induced by vertices v'_1, v'_2, v'_3 and v'_1, v'_2, v'_4 , respectively. The traverse history of G_L shown in the picture starts with v'_1 and v'_2 , and therefore both these isomorphisms are discovered by simple comparison of the traverse histories.

Fig. 4 Estimation cases



The extension of the above algorithm for the induced subisomorphism is the use of the “Black” edges, namely some edge is colored to “Black” instead of being removed from the graph by the main algorithm described in the next section. Note that only if we finished seeking for all subgraphs $(v_i \rightarrow v'_1, v_j \rightarrow v'_2)$ -isomorphic to G_S where $i, j = 1, 2 \dots, |V_S|$, then the edge (v'_1, v'_2) can just be removed from the graph. However, in an induced subgraph isomorphism case, this edge can be used as non-edge after being removed. Thus, in each stage of Algorithm 5.2 we verify that the found subgraphs do not contain “Black” edges.

Notation: in algorithms throughout the rest of the chapter we write “/ * ... */” for the induced subisomorphism case.

Algorithm 5.2 SEARCH TRAVERSE

Input: Graph $G_L = (V_L, E_L)$, starting vertices $v'_1, v'_2 \in V_L$, traverse history H of G_S , /* set of “Black” edges */

Output: All subgraphs /* or induced subgraphs */ of $G_L (v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -isomorphic to G_S , where v_1, v_2 are starting vertices of H /* and do not contain “Black” edges */

```

1: for all  $v' \in V_L$  do
2:    $g(v') \leftarrow 0$ 
3: end for
4:  $g(v'_1) \leftarrow 1$ 
5:  $g(v'_2) \leftarrow 2$ 
6: return  $SearchVisit(\{v'_1, v'_2\}, \emptyset)$ 

```

Algorithm 5.2 :: Procedure $SearhVisit(V', E')$

```

1:  $vtime \leftarrow |V'|$ 
2:  $v' = g^{-1}(vtime)$ 
3: if  $H[vtime].N \not\subseteq \{g(u') > 0 : (v', u') \in E\}$  or  $H[vtime].l \neq l(v')$  then return
   false
4: /* if  $|H[vtime]| \neq |\{g(u') > 0 : (v', u') \in E\}|$  or  $H[vtime].l \neq l(v')$  then
   return false */
5:  $E' \leftarrow E' \cup \{(u', v' : d(u') \in H[vtime].N)$ 
6: if  $|H| = vtime$  then return  $\{(V', E')\}$ 
7: choose  $v \in H[vtime + 1].N$ 
8:  $L \leftarrow \cap \{N_{G_L}(u) : u \in H[vtime + 1].N\}$ 
9:  $S \leftarrow \emptyset$ 
10: for each  $w \in L$  do
11:   if  $g(w) = 0$  /* and  $color(v, w) \neq \text{“Black”}$  */ then
12:      $g(w) \leftarrow vtime$ 
13:      $S \leftarrow S \cup SearchVisit(V' \cup \{w\}, E')$ 
14:      $g(w) \leftarrow 0$ 
15:   end if
16: end for
17: return  $S$ 

```

The main idea of this algorithm is to perform recursively all possible traversals (starting on given vertices $v_1, v_2 \in V_L$) on the subgraphs of the target graph. Each such traverse visiting vertices in a certain order proceed by numbering the visited vertices $g : V' \rightarrow \mathbf{N}$. In each stage of the procedure we verify that the condition of Lemma 5.1 holds for a given traverse history H and for traverse history H' induced by g .

To implement this strategy we use a *SearchVisit* procedure. Logically it can be divided into two parts: in the first part we are trying to check that the condition of Lemma 5.1 holds (lines 3–4), namely that $H[i].N \subseteq H'[i].N$ and $l_i = l'_i$ (for induced subgraph isomorphism $H[i].N = H'[i].N$ and $l_i = l'_i$) for $i \in \{1, 2, \dots, |V_S|\}$. If this condition holds, we extend E' by the corresponding (not “Black” for induced subgraph isomorphism case) edges. In the case when the set of vertices V' is extended to the number of vertices of the pattern graph (namely to $|H|$), then isomorphic subgraph (induced subgraph) $G' = (V', E')$ is found (line 6). Otherwise, in the second part we are trying to extend V' by adding each unvisited vertex adjacent to the one of the visited vertices (recall that H is connected traverse history, then each newly visited vertex has a visited neighbor), thereby finding all possible subgraphs of G_L isomorphic to G_S (line 12).

Actually, the above procedure will work “faster” if it does not enter the second recursive part. Except for the case when the subisomorphism is found, it happens only when there is no edge between two visited vertices when the corresponding edge exists in the pattern graph (or in addition there is an edge between two visited vertices when the corresponding edge does not exist in the pattern graph for the induced subisomorphism case). If this mismatch is discovered in the earlier stage of Algorithm 5.2, then its run time is reduced. This observation explains the importance of using the heuristics presented in Algorithms 4.1 and 4.2, since this technique forces the above checking to be done as soon as possible, thereby decreasing the expected run time.

Remark Obviously, each isomorphic subgraph is found by Algorithm 5.2 exactly k times, where k is the number of $(v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -automorphisms of graph G_S . Therefore the modification of Algorithm 5.2 to just count the number of such subgraphs of G_L which are isomorphic to G_S can be obtained as follows:

- *SearchVisit* procedure should return 1 instead of (V', E') in the case when isomorphic subgraph is found.
- The value returned by Algorithm 5.2 should be divided by k . (Note that the corresponding number of automorphisms can be also calculated in precomputational stage by Algorithm 5.2).

The traverse history with the smallest integrality approach can improve the run time of the main search algorithm for the induced subisomorphism problem in following cases:

- The number of edges of the “large” target graph is close to the number of non-edges (recall that we assume that the number of non-edges of the “large” pattern graph is larger than number of edges)

- The number of edges of the “small” pattern graph is significantly smaller than the number of non-edges.

6 Subsea: subgraph isomorphism algorithm

In this section we combine all the previously defined algorithms to achieve the main goal: subgraph (induced subgraph) isomorphism algorithm.

First, we present an algorithm which collects all traverse histories of a “small” pattern graph in a precomputation stage. Then we present the main algorithm, which finds all the subgraphs (or induced subgraphs) of a given target graph G_L isomorphic to a pattern graph G_S . Note that since our patterns are small, the space required to store all the traverse histories is not prohibitive.

6.1 Precomputation stage

A pair of vertices $\langle v_1, v_2 \rangle \in V^2$ of graph G we will call *redundant* if there exists an $(v_1 \rightarrow v'_1, v_2 \rightarrow v'_2)$ -automorphism of G such that $\langle v_1, v_2 \rangle > \langle v'_1, v'_2 \rangle$.

Algorithm 6.1 finds a corresponding traverse history for each non-redundant pair of adjacent vertices. Note that each edge of the pattern graph may derive 0, 1, or 2 traverse histories.

Algorithm 6.1 ALL TRAVERSE HISTORIES

Input: Graph $G = (V, E)$

Output: Set of traverse histories of G

- 1: $A \leftarrow \emptyset$
 - 2: **for** each $\langle v_1, v_2 \rangle \in V^2$ such that $(v_1, v_2) \in E$ **do**
 - 3: **run** Algorithm 4.1 on G, v_1, v_2 to obtain traverse history H_{v_1, v_2}
 - 4: **if** $!IsRedundant(v_1, v_2)$ **then** $A \leftarrow A \cup \{H_{v_1, v_2}\}$
 - 5: **end for**
 - 6: **return** A
-

Algorithm 6.1 :: Procedure $IsRedundant(v_1, v_2)$

- 1: **for** each $\langle v'_1, v'_2 \rangle \in V^2$ such that $\langle v'_1, v'_2 \rangle \in E$ and $\langle v'_1, v'_2 \rangle < \langle v_1, v_2 \rangle$ **do**
 - 2: **run** Algorithm 5.2 on $G, v'_1, v'_2, H_{v_1, v_2}$ to obtain set S of graphs
 - 3: **if** $S \neq \emptyset$ **then return true else return false**
 - 4: **end for**
-

We use Algorithm 5.2 to find redundant edges, which also can be used to look for the corresponding automorphisms.

This approach enables us to minimize the number of stored traversals, when a set of automorphisms of G is non-empty, thereby reducing the running time of the main search algorithm.

6.2 Main algorithm

In general our technique can be described as follows:

1. Find the traverse history (using Algorithm 6.1) for each non-redundant pair of adjacent vertices of the pattern graph.
2. Divide vertices of a given “large” target graph into two parts using the bisection methods of Sect. 3.
3. For each edge with endpoints in distinct parts of the obtained bisection, find the set of all subgraphs (or induced subgraphs) containing this edge and isomorphic to a given pattern graph (Algorithm 5.2 is used).

After performing these steps, we continue to apply, in recursive manner, the same approach on two subgraphs of G_L induced by the two parts of bisection. We stop when we get a graph with fewer vertices than the pattern graph.

Algorithm 6.2 ALL SUBGRAPH ISOMORPHISMS

Input: Pattern graph $G_S = (V_S, E_S)$; target graph $G_L = (V_L, E_L)$; the set A of all non-redundant traverse histories on graph G_S .

Output: All subgraphs (induced subgraphs) of G_L which are isomorphic to G_S

- 1: Apply Algorithm 6.1 on G_S to obtain set of traverse histories A of G_S .
 - 2: **return** $SubIso(A, G_L)$
-

Algorithm 6.2 :: Procedure $SubIso(A, G_L)$

- 1: **if** $|V_S| > |V_L|$ **then return** \emptyset
 - 2: Apply several times Algorithm 3.1 on G_S to obtain (B, \bar{B}) – bisection of vertices of G_L
 - 3: **for** each $e = (v_1, v_2) \in (B, \bar{B})$ **do**
 - 4: **for** each $H \in A$ **do**
 - 5: Apply Algorithm 5.2 on G_L, v_1, v_2, H to obtain S – a set of subgraphs of G
 - 6: $G \leftarrow G - e$ /* or $color(e) \leftarrow \text{“Black”}$ */
 - 7: **end for**
 - 8: **end for**
 - 9: Eliminate duplicates from S
 - 10: **return** $S \cup SubIso(G_L(B)) \cup SubIso(G_L(\bar{B}))$
-

Theorem Algorithm 6.2 applied on (G_L, G_S) finds all subgraphs of target graph G_L which are isomorphic to pattern graph G_S .

Proof The correctness is assured by Lemma 5.1. Completeness is proved by induction in the number of vertices of the target graph. If $|V_L| < |V_S|$, then, obviously, no subgraph of G_L exists isomorphic to G_S .

Assume that $|V_L| = k \geq |V_S|$, where $k \in \mathbf{N}$. Let $G' = (E', V')$ be a subgraph of G_L isomorphic to G_S (if such subgraph does not exist, then we are done). Let (B, \bar{B})

be a partition of V_L obtained after applying Algorithm 3.1. If $E' \cap e(B, \bar{B}) = \emptyset$ then, obviously G' is a subgraph of $G_L(B)$ or $G_L(\bar{B})$ and is found by recursion. Otherwise, let $e' = (v'_1, v'_2) \in E'$ be an edge that was removed before any other edge in E' . Since $G' \cong G_S$ thereby exists an isomorphism φ which maps V' to V_S . Put $u_1 = \varphi(v'_1)$ and $u_2 = \varphi(v'_2)$, where $u_1, u_2 \in V_S$. Let w_1, w_2 be a smallest pair of vertices of G_S such that there exists an automorphism of G_S which maps u_1 to w_1 and u_2 to w_2 . Then, clearly,

- there is an isomorphism between G' and G_S which maps v'_1 to w_1 and v'_2 into w_2 ,
- $\langle w_1, w_2 \rangle$ is not a redundant pair and consequently an induced traverse history of G_S starting in w_1, w_2 is contained in A .

Thus applying Algorithm 6.2 on v'_1 and v'_2 and on traverse history starting on w_1, w_2 we obtain all subgraphs of G_L and isomorphic to G_S , such that v_1 and v_2 mapped to w_1 and w_2 , respectively. Since we assume, that e' was an edge that was removed before any other edge in E' , all other edges of G' are not removed when the above algorithm is applied and G' is found. \square

The proof for the case of induced subgraph isomorphism is quite similar.

6.3 Duplicate counting problem

Algorithm 6.2 actually finds a copy of G_S for every automorphism of G_S . We eliminate duplicates from the set of G_S instances in Procedure *SubIso* before the set is returned. This elimination is relatively easy since all we have to do is to compare vertex sets of instances.

If we aim to only obtain a count of instances rather than the instance set, we need to return the size of set $S \cup \text{SubIso}(G_L(B)) \cup \text{SubIso}(G_L(\bar{B}))$ instead of the set itself in Procedure *SubIso*. This size needs to be divided by the size of automorphism group of G_S . This size can be computed during Traverse History construction, thus adding no computational overhead to the algorithm.

7 Experimental results

7.1 Settings

The Subsea algorithm was implemented in C++, on an Intel Pentium 4 CPU 2.4 GHz running Windows 2000 and with 512 MB of main memory.

We performed three sets of experiments: experiments on random synthetic graphs, experiments on real-life database called [NCI \(2008\)](#) and experiments on benchmark graphs used by [Sansone et al. \(2004\)](#).

For the synthetic case, We generated random graphs of sizes 100–5000 nodes with various numbers of edges and labels to serve as database graphs. We used two types of random graphs: graphs with large diameter (which are not expander-like) and Erdős–Renyi graphs (which are expander-like, see [Erdős and Renyi 1959](#)) where each edge in a graph exists with a probability p . Labels distribution was uniformly random in

part of experiments, and normal in another part. We have also used unlabeled cliques, stars and lines for the experiments.

7.2 Algorithms

We compared our algorithm's results on these graphs with the results of Ullman's algorithm (see [Ullmann 1976](#)) and (see [Sansone et al. 2004](#)) algorithm. Both algorithms are publicly available on the Web (see [Sansone 2001](#)). In all experiments we use the sub-algorithm of these algorithms that find all distinct subgraphs isomorphic to a given subgraph.

7.3 Detailed comparison

In all of the charts, the running time of all the algorithms is specified in seconds. Bar charts are used to present the results and names of the algorithms (Subsea, Ullman and VF2 for Cordella et al. algorithm) appear above their respective bars.

7.3.1 Present subgraphs

Here, we have compared Subsea, Ullman's and Cordella et al. algorithms on various random and real-life database graphs and subgraphs that have instances in these database graphs (sometimes, quite large number of instances). The X axis shows the number of nodes in the large graph. All randomly generated graphs have average node degree 3 or more. We use log-scale for the Y axis (running time) in order to see small running times better.

The first series of experiments refers to random database graphs with 10 and 5 labels and having uniform label distribution. Figures 5 and 6 show how the three algorithms perform when a subgraph has 15 and 100 nodes respectively and the database graph contains 10 labels. For a subgraph of 15 nodes, Ullman's and Cordella et al. algorithms did not finish in many of the cases (denoted by missing bars). In Fig. 6, we show times on Subsea algorithm only because the other two algorithms did not finish for either database graph within 5 min (we counted pure application time, system time in these cases exceeded 30 min). We see that the Subsea algorithm has an advantage over the other two algorithms.

Figure 7 shows how the three algorithms perform when a subgraph has 15 nodes and the database graph has five different labels; the label distribution on both graphs is uniform. Smaller number of labels causes the number of instances to rise dramatically and the problem of finding all subgraph instances quickly becomes unfeasible. Again, we see that Subsea algorithm performs better in these cases.

The second series of experiments was performed on labeled random Erdős–Renyi with 10 labels and normal label distribution (see Figs. 8, 9 and 10). As expected, normal label distribution implies more appearances of a subgraph in the database graph and thus larger search space. We included in charts only the cases where at least one of the algorithms finishes. For smaller subgraphs (Fig. 8), we see an advantage to




Fig. 5 Subgraph with 15 nodes and 10 labels, uniform label distribution

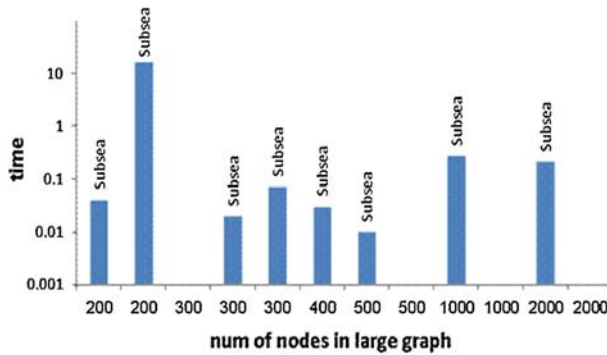


Fig. 6 Subgraph with 100 nodes and 10 labels, uniform label distribution

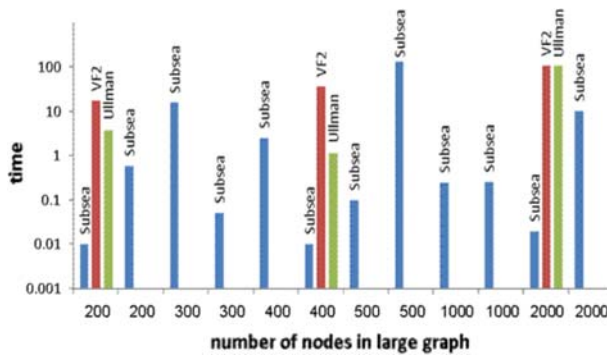


Fig. 7 Subgraph on 15 nodes and 5 labels, uniform label distribution

the Subsea algorithm, but in other cases there are some subgraphs instances where Ullman’s algorithm performs better.

The third series of experiments have been made on unlabeled graphs of specific structure in an attempt to understand the behavior of the Subsea algorithm. These cases, for which the number of subgraph instances can be easily found combinatorially, are in fact very hard for every algorithm tested. Double-digit running time (in seconds) occurs when the number of subgraph instances in a database graph reaches hundreds




Fig. 8 Subgraph of 15 nodes, normal label distribution




Fig. 9 Subgraph of 50 nodes, normal label distribution




Fig. 10 Subgraph of 100 nodes, normal label distribution

of thousands. The horizontal axis in the charts shows which small graphs were tested on the same database graph (*star i* denotes a star with *i* leaves, *line i* denotes a simple path on *i* nodes, *clique i* denotes a complete graph on *i* nodes and *cycle i* denotes a simple cycle on *i* nodes). Figure 11 gives runtimes for the case when both database and subgraphs are stars, with the database graph having 100 nodes. Since Subsea algorithm

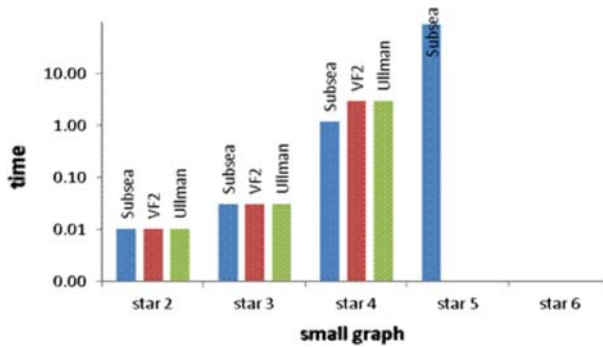


Fig. 11 The database graph is an unlabeled star with 100 nodes




Fig. 12 The database graph is an unlabeled line with 200 nodes

uses divide-and-conquer approach based on cuts, and a well-chosen cut in the center of a star stops the recursion very quickly, we see that Subsea outperforms the other algorithms. In Fig. 12, both the database and subgraphs are lines, and the database graph has 200 nodes. In this case, the Subsea recursion is deep (but not exponential), and still it performs better than other algorithms. For an unlabeled line of 2000 nodes and larger subgraphs, both Ullman's and VF2 algorithms did not finish, while Subsea showed reasonable results (see Fig. 13). Finally, Fig. 14 shows how the algorithms behave when a database graph is an unlabeled complete graph and a subgraph is either a line, a cycle or a star. Complete graph is the worst case for any approach, and we see a clear dependence between the feasibility of a case and the number of nodes in a subgraph. It is not surprising that all algorithms fail when a subgraph size increases. However, since Subsea is based on finding-a-good-cut technique (which unlabeled clique definitely lacks), it performs worse than Ullman's and VF2 algorithms.

Our final series of experiments was performed on a real-life database obtained from National Cancer Institution (see [NCI 2008](#)). We tested all three algorithms on subsets of the database graph, because the complete graph is too large; we present in Table 1 Subsea results only since other algorithms did not finish in any of the cases.




Fig. 13 Unlabeled line on 2000 nodes as a large graph

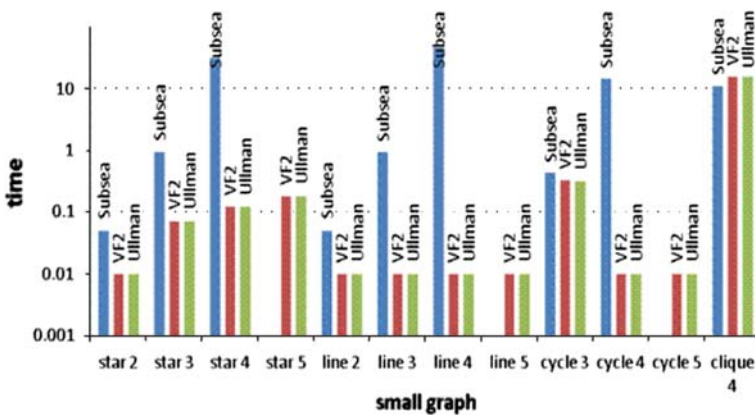


Fig. 14 Complete graph on 50 nodes as a large graph

7.3.2 Absent subgraphs

We have compared all three algorithms on a series of cases where a subgraph in question is not present in the database graph. This is an important domain since in graph mining candidate subgraphs that are not present in database must be eliminated as quickly as possible. An absent subgraph detection can be also viewed as the time required to find the first instance of a subgraph in a database graph. In our charts *X* axis shows the number of nodes in a subgraph and the *Y* axis denoted the time in seconds on logarithmic scale. Distribution of labels on the nodes of database graphs and small graphs in these tests is uniformly random.

Figure 15 shows the behavior of the three algorithms on large-diameter database graph with 1000 nodes and 10 labels. The chart in Fig. 16 shows the behavior of the three algorithms on large-diameter database graph with 2000 nodes and 10 labels. We see that Subsea runs faster than Ullman’s algorithm on these graphs, but slower than the VF2 algorithm. However, the number of cases where the latter algorithm was stuck is quite large, while Subsea provides an answer in most of the cases.

Table 1 NCI database tests

Large graph (nodes)	Small graph (nodes)	Subsea time (s)	Num of matches
532	20	0.23	172
532	80	0.26	696
532	50	0.22	172
999	50	0.95	257
999	50	0.61	257
999	54	0.78	257
999	41	0.86	1158
999	50	0.85	1157
999	50	0.69	1158
1000	90	0.64	290
1000	20	0.61	290
1000	70	0.73	1221
1000	99	0.81	1221
1000	71	0.62	290
2000	22	3.18	642
2000	90	3.43	2605

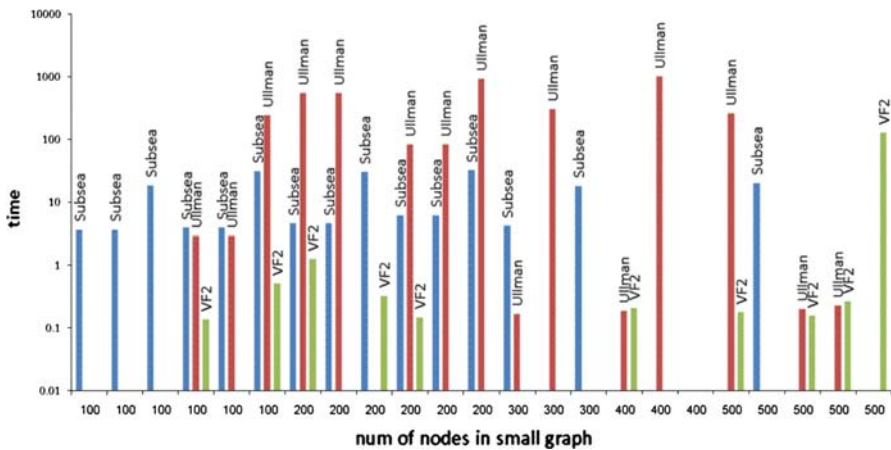


Fig. 15 Database graph with 1000 nodes

7.4 Benchmark graphs

For the last series of experiments we used the ARG Database (see [SIVALab of the University of Naples 2003](#)) by the Intelligent Systems and Artificial Vision Laboratory (SIVALab) of the University of Naples. We tested Subsea and VF2 algorithms on selection of unlabeled graphs from this database. We have run these experiments on cluster

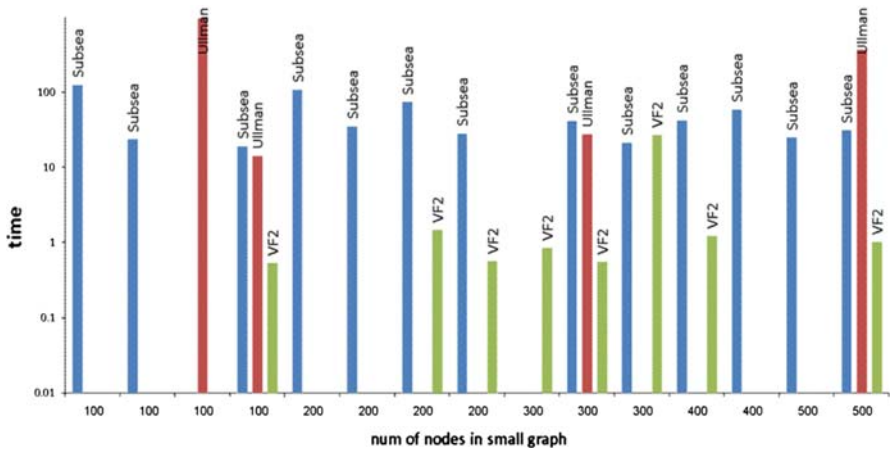


Fig. 16 Database graph with 2000 nodes

Table 2 Comparison of Subsea and VF2 on ARG random graphs and 2D meshes

Nodes	Random graphs			2D meshes			
	$\eta = 0.01$	$\eta = 0.05$	$\eta = 0.1$	Regular	$\rho = 0.2$	$\rho = 0.4$	$\rho = 0.6$
Up to 60 nodes	Subsea	Subsea	VF2	VF2	VF2	Subsea	Subsea
60–80 nodes	Subsea	Subsea	Subsea	VF2	Subsea	Subsea	Subsea
80–1000 nodes	Subsea	Subsea	Subsea	Subsea	Subsea	Subsea	Subsea

consisting of four Intel SMP servers with 2 Dual Core Xeon 5140 2.33GHz processors and 4G RAM. Table 2 summarizes obtained results and shows which algorithm ran faster for various types of graphs. For a random graph, η denotes the probability with which the edge is added and for a 2-dimensional mesh ρ indicates the ratio of extra edges added. Figures 17, 18 and 19 show a detailed comparison of runtimes of both algorithms on sample graphs of various size. On graph with bounded degree, VF2 runs faster than Subsea (see Fig. 19).

Detailed comparison shows that Subsea runs faster on randomly generated graphs and two-dimensional meshes, while VF2 shows better times on bounded degree graphs. In general, as the graph becomes larger and the number of isomorphic subgraph grows, the Subsea advantage becomes more prominent.

Finally, Figs. 20 and 21 show how Subsea scales when graph density, graph size and number of labels change. In all these experiments, single small graph with 10 nodes is used. Figure 20 demonstrates how the runtime of Subsea changes when the number of edges of the large graph with 100 nodes increases. When large graph becomes denser, the number of isomorphisms rises to hundreds of thousands. Figure 21 shows the dependency of Subsea runtime on the large graph size when the density remains the same. Figure 22 demonstrates how Subsea scales when the number of labels increases in the same dense database graph of 100 nodes and 1200 edges.

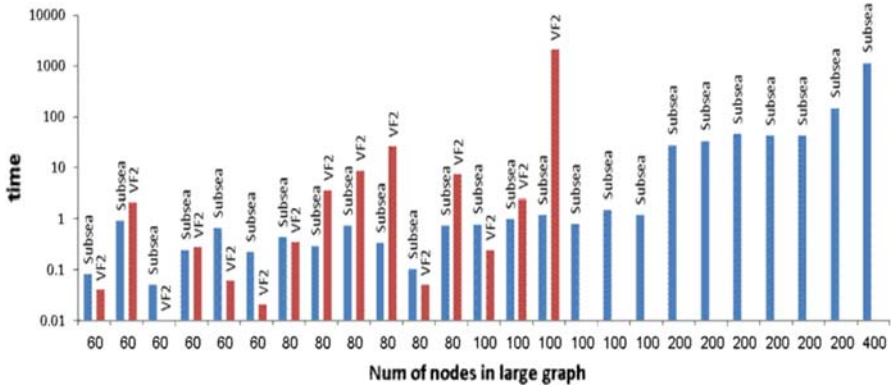


Fig. 17 Randomly generated graphs with $\eta = 0.01$

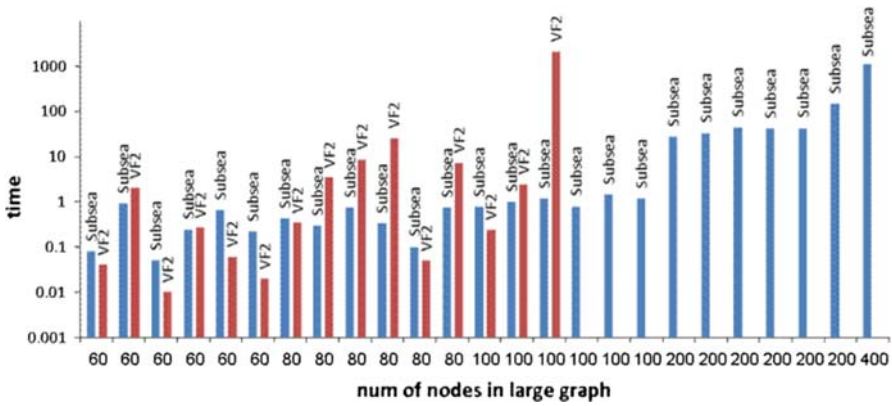


Fig. 18 Randomly generated graphs with $\eta = 0.1$

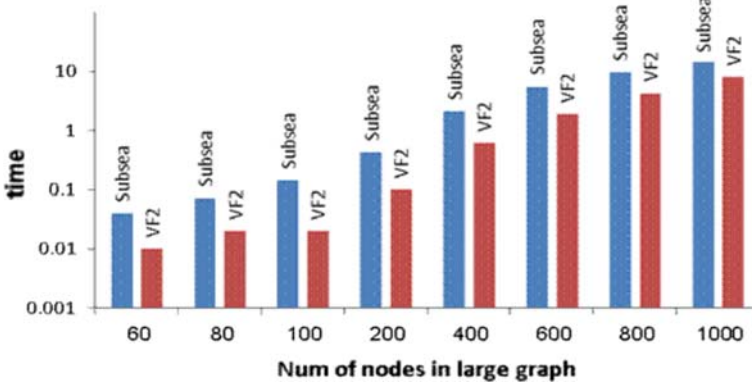


Fig. 19 Bounded-valence graphs with valence 3

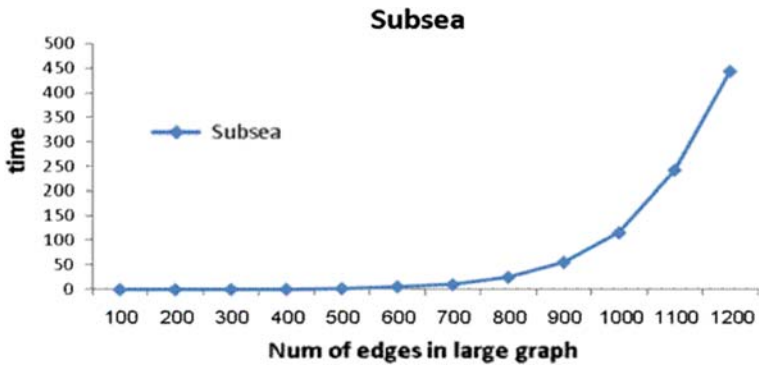


Fig. 20 Subsea runtimes for database graphs of increasing density

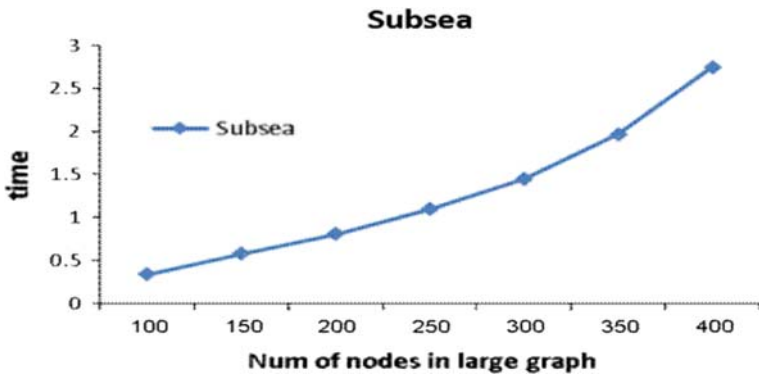


Fig. 21 Subsea runtimes for database graphs of increasing size and the same density

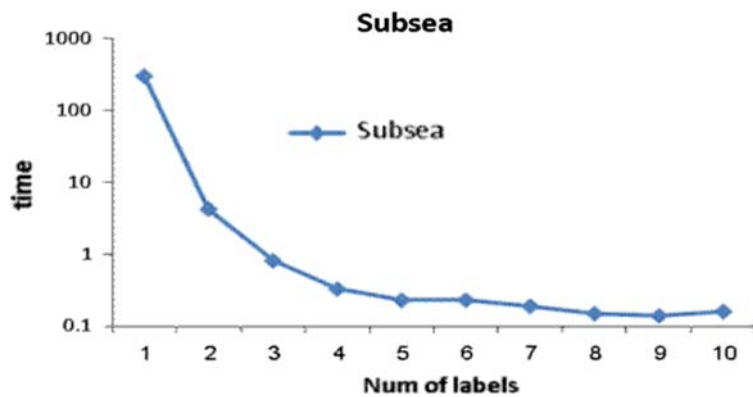


Fig. 22 Subsea runtimes for database graphs of increasing label number

As expected, running times rise when density and graph size increase and drops dramatically when the number of labels increases.

7.5 Summary

There is one distinct observation from all the above experiments. Subsea outperforms the other algorithm when there exist many isomorphic instances of the small graph within the database graph. We feel that the main reason is the bisection approach. Subsea starts the search only from “small” cuts, while in other algorithms the entire graph is searched for the next instance.

8 Conclusions

We presented a new heuristic for finding all isomorphisms of a source pattern graph (or induced source graph) in a target graph. The new approach is based on a bisection algorithm and on a search strategy which is directed by an original traverse history notion. The experimental results show advantage of our approach over algorithms of [Ullmann \(1976\)](#) and [Sansone et al. \(2004\)](#) when all the instances of a subgraph need to be found. The conclusion from the above experiments is that Subsea is much better than the two other algorithms when multiple instances of a subgraph are searched for. The reason is that in Subsea in every bisection we test for multiple instances and never search that part of the graph later. When a single instance of a subgraph needs to be found, the algorithms of [Ullmann \(1976\)](#) and [Sansone et al. \(2004\)](#) have an advantage over Subsea.

The proposed algorithm has general validity, since no constraints are imposed on the topology of the pattern and the target graphs, and the method can be easily be extended to the case of directed graphs as well. An extension to the case of *edge-labeled* graphs is less trivial and will require some effort.

Another advantage of our algorithm is that the bisection procedure leads itself naturally to parallel computation. We plan to pursue this topic in the future. An important goal is to enable current researchers to compare their codes with each other, in hopes of identifying the more effective of the recent algorithmic innovations that have been proposed. Although many implementations do not include the most sophisticated speed-up tricks and thus may not be able to compete on speed with highly tuned alternatives, we may be able to gain insight into the effectiveness of various algorithmic ideas by comparing codes with similar levels of internal optimization. It will also be interesting to compare the run times of the best current optimization codes with those of the more complicated heuristics on instances that both can handle.

Acknowledgments We thank anonymous reviewers for their useful comments on our paper. We thank Lior Becker and Arie Ohana for doing extraordinary job implementing the Subsea algorithm and making it even more efficient than we anticipated. We also thank Rosa Shlayzer for conducting the experiments. We thank Fraenkel Center for Computer Science and Paul Ivanier Center for Robotics and Production Management for partially supporting this work.

References

- Akinniyi FA, Wong AKC, Stacey DA (1986) A new algorithm for graph monomorphism based on the projections of the product graph. *Trans Systems, Man and Cybernetics* SMC-16:740–751
- Alon N, Yuster R, Zwick U (1995) Color-coding. *Electronic colloquium on computational complexity (ECCC)*, 1(009), 1994. Full paper appears in *J ACM* 42(4):44–856
- Ambauen R, Fischer S, Bunke H (2003) Graph edit distance with node splitting and merging, and its application to diatom identification. *Lecture notes in computer science*, vol 2726, pp 259–264
- Bunke H, Messmer BT (2000) Efficient subgraph isomorphism detection: a decomposition approach. *IEEE Trans Knowl Data Eng* 12:307–323
- Batz GV (2006) An optimization technique for subgraph matching strategies. Technical report 2006–2007, Universitat Karlsruhe, Faculty of Informatik
- Berretti S, Del Bimbo A, Pala P (2004) A graph edit distance based on node merging. In: *Proceedings of CIVR2004*, pp 464–472
- Boeres M, Ribeiro C, Bloch I (2004) A randomized heuristic for scene recognition by graph matching. In: *Proceedings of WEA 2004*, pp 100–113
- Cai D, Shao Z, He X, Yan X, Han J (2005) Mining hidden community in heterogeneous social networks. In: *Proceedings of the 3rd international workshop on link discovery*, pp 58–65
- Champin P, Solnon C (2003) Measuring the similarity of labeled graphs. *Conference on case-based reasoning (ICCBR)*, pp 100–113
- Chen MS, Park JS, Yu PS (1998) Efficient data mining for path traversal patterns. *IEEE Trans Knowl Data Eng* 10(2):209–221
- Cheng JK, Huang TS (1981) A subgraph isomorphism algorithm using resolution. *Pattern Recognit* 13:371–379
- Cordella LP, Foggia P, Sansone C, Vento M (1999) Performance evaluation of the vf graph matching algorithm. In: *Proceedings of ICIAP 1999*, pp 1172–1177
- Cortadella J, Valiente G (2000) A relational view of subgraph isomorphism. In: *Proceedings of fifth international seminar on relational methods in computer science*, pp 45–54
- Dehaspe L, Toivonen H, King RD (1998) Finding frequent substructures in chemical compounds. In: *Proceedings of the 4th international conference on knowledge discovery and data mining (KDD-98)*, pp 30–36
- Dessmark A, Lingas A, Proskurowski A (2000) Faster algorithms for subgraph isomorphism of k -connected partial k -trees. *Algorithmica* 27(3):337–347
- Eppstein D (1999) Subgraph isomorphism in planar graphs and related problems. *J Graph Algorithms & Appl* 3(3):1–27
- Erdoős P, Renyi A (1959) On random graphs. I. *Publicationes Mathematicae* 6:290–297
- Foggia P, Sansone C, Vento M (2001) An improved algorithm for matching large graphs. In: *3rd IAPR-TC15 workshop on graph-based representations. Lecture Notes in Computer Science*, Ischia
- Garey MR, Johnson DS (1979) *Computers and Intractability: A guide to the theory of NP-completeness*. W. H. Freeman
- Gudes E, Shimony SE, Vanetik N (2006) Discovering frequent graph patterns using disjoint paths. *IEEE Trans Knowl Data Eng* 18(11):1441–1456
- Krissinel EB, Henrick K (2004) Common subgraph isomorphism detection by backtracking search. *Softw Prac Exp* 34(6):591–607
- Kuramochi M, Karypis G (2001) Frequent subgraph discovery. In: *Proceedings of ICDM 2001*
- Kuramochi M, Karypis G (2004) Finding frequent patterns in a large sparse graph. In: *Proceedings of SDM 2004*
- Larrosa J, Valiente G (2000) Graph pattern matching using constraint satisfaction. In: *Proceedings of joint APPLIGRAPH/GETGRATS workshop of graph transformation systems*, pp 189–196
- Lin X, Liu Ch, Zhang Y, Zhou X (1998) Efficiently computing frequent tree-like topology patterns in a web environment. In: *Proceedings of 31st International conference on technology of object-oriented land systems*
- Lingas A, Syslo MM (1989) A polynomial-time algorithm for subgraph isomorphism of two-connected series-parallel graphs. Technical report LiTH-IDA-R-89-05, Linköping University, Department of Computer and Information Science
- Matula DW (1978) Subtree isomorphism in $O(n^{5/2})$. *Ann Discrete Math* 2:91–106

- Messmer BT, Bunke H (1996) Subgraph isomorphism detection in polynomial time on preprocessed model graphs. In: Proceedings of 2nd Asian conference on computer vision. Springer, pp 373–382
- NCI (2008) Database of interacting proteins. Technical report. Available at <http://dtp.nci.nih.gov>
- Nijssen S, Kok JN (2004a) Frequent graph mining and its application to molecular databases. In: Proceedings of the IEEE international conference on systems, man and cybernetics, SMC 2004, Den Haag, Netherlands, October 10–13, 2004. IEEE Press
- Nijssen S, Kok JN (2004b) A quickstart in frequent structure mining can make a difference. In: Kohavi R, Gehrke J, DuMouchel W, Ghosh J (eds) Proceedings of the 10th ACM SIGKDD international conference on knowledge discovery and data mining, KDD2004, Seattle, USA, August 22–25, 2004, pp 647–652. ACM Press
- Pennec X, Ayache N (1998) A geometric algorithm to find small but highly similar 3d substructures in proteins. *Bioinformatics* 14(6):516–522
- Sammoud O, Solnon C, Ghédira K (2005) Ant algorithm for the graph matching problem. In: Proceedings of EvoCOP 2005, pp 213–223
- Sansone C (2001) Graph database library for subgraph isomorphism. Technical report. Available at <http://amalfi.dis.unina.it/graph/>
- Sansone C, Cordella LP, Foggia P, Vento M (2004) A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Trans Pattern Anal Mach Intell* 26(10):1367–1372
- SIVALab of the University of Naples (2003) Arg graph database. Technical report. Available at <http://amalfi.dis.unina.it/graph/>
- Ullmann JR (1976) An algorithm for subgraph isomorphism. *J Assoc Comput Mach* 23:31–42
- Vanetik N, Gudes E, Shimony SE (2002) Computing frequent graph patterns from semistructured data. In: Proceedings of ICDM 2002, pp 458–465
- Wang K, Liu H (1998) Discovering typical structures of documents: a road map approach. In: Proceedings of SIGIR 1998, pp 146–154
- Yan X, Han J (2002) gSpan: graph-based substructure pattern mining. In: Proceedings of ICDM 2002, pp 721–724
- Zhang S, Li S, Yang J (2009) Gaddi: distance index based subgraph matching in biological networks. In: Proceedings of EDBT conference 2009