

A Survey of Frequent Subgraph Mining Algorithms

Chuntao Jiang, Frans Coenen and Michele Zito

The University of Liverpool, Department of Computer Science, Ashton Building, Ashton Street, Liverpool, L69 3BX, UK

E-mail: cjiang@csc.liv.ac.uk, frans@csc.liv.ac.uk, michele@csc.liv.ac.uk

Abstract

Graph mining is an important research area within the domain of data mining. The field of study concentrates on the identification of frequent subgraphs within graph data sets. The research goals are directed at: (i) effective mechanisms for generating candidate subgraphs (without generating duplicates) and (ii) how best to process the generated candidate subgraphs so as to identify the desired frequent subgraphs in a way that is computationally efficient and procedurally effective. This paper presents a survey of current research in the field of frequent subgraph mining, and proposed solutions to address the main research issues.

1 Introduction

The primary goal of data mining is to extract statistically significant and useful knowledge from data (Chen et al. 1996; Han & Kamber 2006). The data of interest can take many forms: vectors, tables, texts, images, and so on. Data can also be represented by various means. Structured data and semi-structured data are naturally suited to graph representations. To give one example, if we consider protein-protein interaction networks (a common application area for graph mining) these can be represented in a graph format such that the vertexes indicate genes, and the directed or undirected edges indicate physical interactions or functional associations (Alm & Arkin 2003). Because of the ease with which structured and semi-structured data can be represented in graph formats, there has been much interest in the mining of graph data (often referred to as graph based data mining or graph mining). A number of popular research sub-domains of graph mining are listed in Table 1.

Table 1 Popular graph mining research sub-domains

Frequent subgraph mining (Cook & Holder 1994,2000; Inokuchi et al. 2000; Yan & Han 2002)
Correlated graph pattern mining (Ke et al. 2007; Ke et al. 2009; Ozaki & Ohkawa 2008)
Optimal graph pattern mining (Yan et al. 2008; Fan et al. 2008)
Approximate graph pattern mining (Kelley et al. 2003; Sharan et al. 2005; Chen et al. 2007a)
Graph pattern summarization (Xin et al. 2006; Chen et al. 2008)
Graph classification (Huan et al. 2004; Kudo et al. 2004; Deshpande et al. 2005)
Graph clustering (Flake et al. 2004; Huang & Lai 2006; Newman 2004)
Graph indexing (Shasha et al. 2002; Yan et al. 2004)
Graph searching (Yan et al. 2005b; Yan et al. 2006; Chen et al. 2007b)
Graph kernels (Gärtner et al. 2003; Kashima et al. 2003; Borgwardt & Kriegel 2005)
Link mining (Chakrabarti et al. 1999; Kosala & Blockeel 2000; Getoor & Diehl 2005; Liu 2008)
Web structure mining (kleinberg 1998; Brin & Page 1998)
Work-flow mining (Greco et al. 2005)
Biological network mining (Hu et al. 2005)

Frequent Subgraph Mining (FSM) is the essence of graph mining. The objective of FSM is to extract all the frequent subgraphs, in a given data set, whose occurrence counts are above a specified threshold. Figure 1 presents an overview of the the domain of FSM in terms of the number of significant FSM algorithms that have been proposed over the period 1994 to the present. From the figure we can see periods of activity in the early 90s (coinciding with the introduction of the concept of data mining) followed by another period of activity from 2002 to 2007. No “new” algorithms have been introduced over the past few years, indicating that the field is reaching maturity, although there has been much work focussed on variations of existing algorithms.

Other than the research activity associated with FSM the importance of FSM is also reflected in its many areas of its application. Figure 1(b) presents an overview of the application domain of FSM in terms of the number of FSM algorithms reported in the literature and the specific application domain at which they have been directed. From the figure it can be seen that three application domains (chemistry, web, and biology) dominated the usage of FSM algorithms.

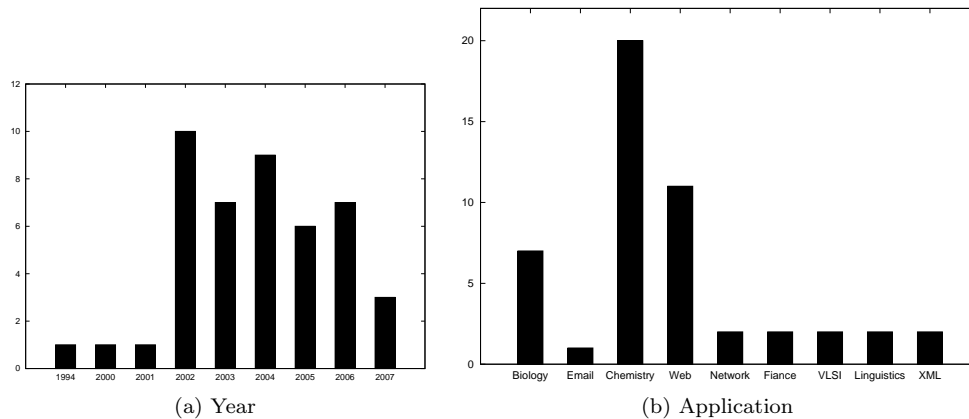


Figure 1 The distribution of the most significant FSM algorithms with respect to year of introduction and application domain

The straightforward idea behind FSM is to “grow” candidate subgraphs, in either a breadth first or depth first manner (candidate generation), and then determine if the identified candidate subgraphs occur frequently enough in the graph data set for them to be considered interesting (support counting). The two main research issues in FSM are thus how to efficiently and effectively (i) generate the candidate frequent subgraphs and (ii) determine the frequency of occurrence of the generated subgraphs. Effective candidate subgraph generation requires that the generation of duplicate or superfluous candidates is avoided. Occurrence counting requires repeated comparison of candidate subgraphs with subgraphs in the input data, a process known as *isomorphism* checking. FSM, in many respects, can be viewed as an extension of Frequent Itemset Mining (FIM) popularised in the context of association rule mining (see for example Agrawal & Srikant 1994). Consequently, many of the proposed solutions to addressing the main research issues effecting FSM are based on similar techniques found in the domain of FIM. For example the downward closure property associated with itemsets has been widely adopted with respect to candidate subgraph generation.

In this paper the authors present a survey of the current “state of the art” of FSM. With reference to the literature we can identify many different types of mining strategies, with respect to many different types of graph, to produce many different kinds of patterns. So as to impose some form of order to the domain of FSM we have focused on the nature of FSM algorithms; categorising such algorithms according to: (i) candidate generation strategy, (ii) the mechanism for traversing the search space and (iii) the occurrence counting process. To further facilitate

understanding of the field of FSM we distinguish between frequent subtree mining and the much general domain of frequent subgraph mining. The rest of this paper is organised as follows. We begin in Section 2 by introducing some formal definitions and terminology; followed, in Section 3, with a generic overview of the FSM process. In Section 4 and 5 we then consider current frequent subtree and subgraph mining algorithms respectively. A brief summary is provided at the end of both of these two sections. Finally, Section 6 presents some conclusions and future directions.

2 Formalism

There are two separate problem formulations for FSM: (i) *graph transaction based FSM* and (ii) *single graph based FSM*. In graph transaction based FSM, the input data comprises a collection of medium-size graphs called *transactions*. Note that the term “transaction” is borrowed from the field of Association Rule Mining (Agrawal & Srikant 1994). In single graph based FSM the input data, as the name implies, comprises one very large graph.

A subgraph g is considered to be frequent if its *occurrence count* is greater than some predefined threshold value. The occurrence count for a subgraph is usually referred to as its *support*, and the consequently the threshold is referred to as the *support threshold*. The support of g may be computed using either *transaction-based counting* or *occurrence-based counting*. Transaction-based counting is only applicable to graph transaction based FSM, while occurrence-based counting may be applied to either transaction based FSM or single graph based FSM. However, occurrence-based counting is typically used with single graph based FSM.

In transaction-based counting the support is defined by the number of graph transactions that g occurs in, one count per transaction regardless of whether g occurs once or more than once in a particular graph transaction. Thus, given a database $\mathcal{G} = \{G_1, G_2, \dots, G_T\}$ consisting of a collection of graph transactions, and a support threshold $\sigma (0 < \sigma \leq 1)$; then the set of graph transactions where a subgraph g occurs is defined by $\delta_{\mathcal{G}}(g) = \{G_i | g \subseteq G_i\}$. Thus, the *support* of g is defined as:

$$sup_{\mathcal{G}}(g) = |\delta_{\mathcal{G}}(g)|/T \quad (1)$$

where $|\delta_{\mathcal{G}}(g)|$ denotes the cardinality of $\delta_{\mathcal{G}}(g)$ and T the number of graphs (transactions) in \mathcal{G} . Therefore, g is *frequent* if and only if $sup_{\mathcal{G}}(g) \geq \sigma$. In occurrence-based counting we simply count up the number of occurrences of g in the input set.

Transaction-based counting offers the advantage that the well-known *Downward Closure Property*¹ (DCP) can be employed to significantly reduce the computation overhead associated with candidate generation in FSM. In the case of occurrence-based counting, either an alternative frequency measure, which maintains the DC property, must be established; or some heuristics adopted to keep the computation as inexpensive as possible. There are a variety of support measures (Vanetik 2002; Kuramochi & Karypis 2004c,2005; Vanetik et al. 2006) that may be adopted for single graph based FSM, these will be discussed further in Section 5.1.2.

2.1 Preliminary definitions

Generally speaking, a *graph* is defined to be a set of vertexes (nodes) which are interconnected by a set of edges (links) (Gibbons 1985). The graphs used in FSM are assumed to be *labelled simple graphs*². In the following paragraphs a number of widely used definitions, used later in this paper, are introduced.

Labelled Graph: A labelled graph can be represented as $G(V, E, L_V, L_E, \varphi)$, where V is a set of vertexes, $E \subseteq V \times V$ is a set of edges; L_V and L_E are sets of vertex and edge labels

¹If a graph is frequent, then all of its subgraphs will also be frequent

²A *simple graph* is an un-weighted and un-directed graph with no loops and no multiple links between any two distinct nodes (Gibbons 1985; West 2000).

respectively; and φ is a label function that defines the mappings $V \rightarrow L_V$ and $E \rightarrow L_E$. G is *(un)directed* if $\forall e \in E$, e is an (un)ordered pair of vertexes. A *path* in G is a sequence of vertexes which can be ordered such that two vertexes form an edge if and only if they are consecutive in the list (West 2000). G is *connected*, if it contains a path for every pair of vertexes in it and *disconnected* otherwise. G is *complete* if each pair of vertexes is joined by an edge and G is *acyclic* if it contains no cycle.

Subgraph: Given two graphs $G_1(V_1, E_1, L_{V_1}, L_{E_1}, \varphi_1)$ and $G_2(V_2, E_2, L_{V_2}, L_{E_2}, \varphi_2)$, G_1 is a subgraph of G_2 , if G_1 satisfies: (i) $V_1 \subseteq V_2$, and $\forall v \in V_1, \varphi_1(v) = \varphi_2(v)$, (ii) $E_1 \subseteq E_2$, and $\forall (u, v) \in E_1, \varphi_1(u, v) = \varphi_2(u, v)$. G_1 is an **induced subgraph** of G_2 , if G_1 further satisfies: $\forall u, v \in V_1, (u, v) \in E_1 \Leftrightarrow (u, v) \in E_2$, in addition to the above conditions. G_2 is also a supergraph of G_1 (Iokuchi et al. 2002; Huan et al. 2003).

Graph Isomorphism: A graph $G_1(V_1, E_1, L_{V_1}, L_{E_1}, \varphi_1)$ is isomorphic to another graph $G_2(V_2, E_2, L_{V_2}, L_{E_2}, \varphi_2)$, if and only if a bijection $f: V_1 \rightarrow V_2$ exists such that: (i) $\forall u \in V_1, \varphi_1(u) = \varphi_2(f(u))$, (ii) $\forall (u, v) \in E_1 \Leftrightarrow (f(u), f(v)) \in E_2$, (iii) $\forall (u, v) \in E_1, \varphi_1(u, v) = \varphi_2(f(u), f(v))$. The bijection f is an isomorphism between G_1 and G_2 . A graph G_1 is **subgraph isomorphic** to a graph G_2 , if and only if there exists a subgraph $g \subseteq G_2$ such that G_1 is isomorphic to g (Huan et al. 2003). In this case g is called an **embedding** of G_1 in G_2 .

Lattice: Given a database \mathcal{G} , a lattice is a structural form used to model the search space for finding frequent subgraphs, where each vertex represents a connected subgraph of the graph in \mathcal{G} (Thomas et al. 2006). The lowest vertex depicts the empty subgraph and the vertexes at the highest level depict the graphs in \mathcal{G} . A vertex p is a parent of the vertex q in the lattice, if q is a subgraph of p , and q is different from p by exactly one edge. The vertex q is a child of p . All the subgraphs of each graph $G_i \in \mathcal{G}$ which occur in the database are present in the lattice and every subgraph occurs only once in it .

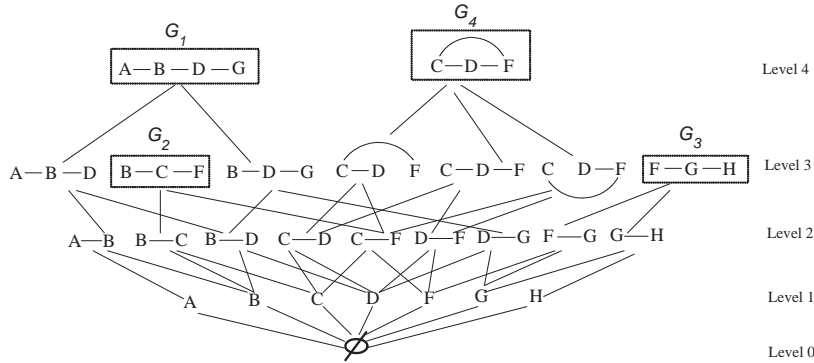


Figure 2 $Lattice(\mathcal{G})$ (Figure based on a similar figure presented in (Thomas et al. 2006))

Example: given a graph data set $\mathcal{G} = \{G_1, G_2, G_3, G_4\}$, the corresponding $Lattice(\mathcal{G})$, is given in Figure 2. In the figure, the lowest vertex ϕ represents the empty subgraph, and the vertexes at the highest level correspond to G_1, G_2, G_3 , and G_4 . The parents of the subgraph $B-D$ are subgraphs $A-B-D$ (joining the edge $A-B$) and $B-D-G$ (joining the edge $D-G$). Similarly, subgraphs $B-C$ and $C-F$ are the children of the subgraph $B-C-F$. \diamond

Free Tree: An undirected graph that is connected and acyclic (Chi et al. 2004; Chi et al. 2004a).

Labelled Unordered Tree: A labelled unordered tree (an unordered tree, for short) is a directed acyclic graph denoted as $T(V, \phi, E, v_r)$, where V is a set of vertexes of T ; ϕ is a labelling function, such that $\forall v_i \in V, \phi(v_i) \rightarrow v_i$; $E \subseteq V \times V$ is a set of edges of T ; and v_r is

distinguished vertex called *root* of T . For $\forall v_i \in V$, there is a unique path $(v_r, v_1, v_2, \dots, v_i)$ from the root v_r to v_i (Asai et al. 2002; Asai et al. 2003). If a vertex v_i is on the path from the root to the vertex v_j , then v_i is an *ancestor* of v_j , and v_j is a *descendant* of v_i . For each edge $(v_i, v_j) \in E$, v_i is the *parent* of v_j , and v_j is a *child* of v_i . Vertexes that share the same parent are *siblings*. The *size* of T is defined to be the number of vertexes in T . A vertex without any child is a *leaf* vertex; otherwise it is an *intermediate* vertex. The *right most path* of T is the path from the root vertex to the *rightmost leaf*. The *depth(level)* of a vertex is the *length of the path*³ from the root to that vertex. The *degree* of a vertex v , denoted by $degree(v)$, is the number of edges incident to it (West 2000; Chi et al. 2004; Chi et al. 2004a; Tan et al. 2005).

Labelled Ordered Tree: A labelled ordered tree⁴ (an ordered tree, for short) is a labelled unordered tree but with a left-to-right ordering imposed among the children of each vertex (Asai et al. 2002; Asai et al. 2003; Chi et al. 2004).

Bottom-up Subtree: Given a rooted tree $T(V, \phi, E, v_r)$ (ordered or unordered), $\vec{T}(\vec{V}, \vec{\phi}, \vec{E}, \vec{v}_r)$ is a bottom-up subtree of T if and only if: (i) $\vec{V} \subseteq V$, (ii) $\vec{E} \subseteq E$, (iii) the labelling of \vec{V} and \vec{E} in T is preserved in \vec{T} , (iv) $\forall v \in V$, if $v \in \vec{V}$ then all descendants of v must also be in \vec{V} and (v) if T is ordered, then the left-to-right ordering among the siblings in T should be preserved in \vec{T} (Chi et al. 2004; Valiente 2002).

Induced Subtree: Given a labelled tree $T(V, \phi, E, v_r)$ (free tree or unordered tree or ordered tree), $\vec{T}(\vec{V}, \vec{\phi}, \vec{E}, \vec{v}_r)$ is an induced subtree of T , if and only if (1) $\vec{V} \subseteq V$; (2) $\vec{E} \subseteq E$; (3) the labelling of \vec{V} and \vec{E} in T is preserved in \vec{T} ; (4) if defined for ordered trees, the left-to-right ordering among the siblings in \vec{T} should be a sub-ordering of the corresponding vertexes in T (Chi et al. 2004; Tan et al. 2006).

Embedded Subtree: Given a labelled tree $T(V, \phi, E, v_r)$, $\vec{T}(\vec{V}, \vec{\phi}, \vec{E}, \vec{v}_r)$ is an embedded subtree of T , if and only if: (i) $\vec{V} \subseteq V$, (ii) $\forall v \in \vec{V}, \vec{\phi}(v) = \phi(v)$, (iii) $\forall (u, v) \in \vec{E}$ such that u is the parent of v , u is an ancestor of v in T and (iv) in the case of ordered trees, $\forall (u, v) \in \vec{V}, preorder(u) < preorder(v)$ in \vec{T} if and only if $preorder(u) < preorder(v)$ in T , where the pre-order of a vertex is its index in the tree according to the pre-order traversal⁵.

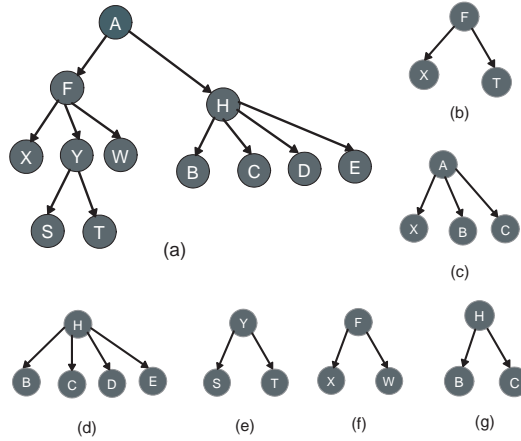


Figure 3 Different types of trees

³The *length of a path* is equivalent to the number of edges in the path

⁴A labelled ordered tree, in graph theory, is also called a *rooted plane tree* (West 2000).

⁵A *preorder traversal* is where a sequence of operations are performed recursively as follows: visit the root first; and then do a preorder traversal of each of the subtrees of the root one-by-one in the order given (Preiss 1998)

Table 2 Categorisation of exact matching (sub) graph isomorphism testing algorithms

Algorithms	Main Techniques	Matching Types
Ullmann	backtracking + look ahead function	graph & subgraph isomorphism
SD	distance matrix + backtracking	graph isomorphism
Nauty	group theory + canonical labelling	graph isomorphism
VF	DFS strategy + feasibility rules	graph & subgraph isomorphism
VF2	VF's rationale + advanced data structures	graph & subgraph isomorphism

To summarise the above Figure 3 gives some examples of bottom-up subtrees, induced subtrees, and embedded subtrees. In the figure: tree (a) on the left represents a data tree, trees (d) and (e) are two bottom-up subtrees of (a), trees (f) and (g) are two induced subtrees of (a), and trees (b) and (c) are two embedded subtrees of (a). The relationship among these three types of subtrees can be denoted as: bottom-up subtree \subseteq induced subtree \subseteq embedded subtree.

2.2 Graph isomorphism detection

The kernel of FSM is *(sub)graph isomorphism detection*. Graph isomorphism is neither known to be solvable in polynomial time nor *NP*-complete, while subgraph isomorphism, where we wish to establish whether a subgraph is wholly contained within a supergraph, is known to be *NP*-complete (Garey & Johnson 1979). When restricting the graphs to trees, (sub)graph isomorphism detection becomes (sub)tree isomorphism detection. Tree isomorphism detection can be solved in a linear time (see algorithm proposed in Hopcroft & Tarjan (1972)). Faster subtree isomorphism detection algorithms, with worst case time complexity of $\mathcal{O}(k^{1.5}n)$, were proposed by Matula (1978) and Chung (1987), and further improved upon by Shamir & Tsur (1999) in $\mathcal{O}(\frac{k^{1.5}}{\log k}n)$ time (k and n are the sizes of the subtree and the tree to be searched in terms of the number of vertexes).

Subgraph isomorphism detection is fundamental to FSM. A significant number of "efficient" techniques have been proposed, all directed at reducing, as far as possible, the computational overhead associated with subgraph isomorphism detection. Subgraph isomorphism detection techniques can be roughly categorized as being either: exact matching (Ullmann 1976; Schmidt & Druffel 1976; McKay 1981; Cordella et al. 1998; Cordella et al. 2001) or error tolerant matching (Shapiro & Haralick 1981; Bunke & Allerman 1983; Christmas et al. 1995; Messmer & Bunke 1998). Most FSM algorithms adopt exact matching. A categorisation of the main exact matching subgraph isomorphism detection algorithms is presented in Table 2. In Table 2, column two indicates the main methods employed to carry out the isomorphism detection, and column three indicates whether the isomorphism detection algorithm applies to graph isomorphism, or subgraph isomorphism.

With reference to Table 2, Ullmann's algorithm employs a backtracking procedure with a look-ahead function to reduce the size of the search space (Ullmann 1976). The SD algorithm, in turn, utilizes a distance matrix representation of a graph with a backtracking procedure to reduce the search (Schmidt & Druffel 1976). The Nauty algorithm (McKay 1981) uses group theory to transform graphs to be matched into a canonical form so as to provide for more efficient and effective graph isomorphism checking. However, it has been noted (Conte et al. 2004) that the construction of the canonical forms can lead to exponential complexity in the worst case. Although Nauty was regarded as the fastest graph isomorphism algorithm by Conte et al. (2004), Miyazaki (1997) demonstrated that there exists some categories of graphs which required exponential time to generate the canonical labelling. The VF (Cordella et al. 1998) and VF2 (Cordella et al. 2001) algorithms use a Depth First Search (DFS) strategy, assisted by a set of feasibility rules to prune the search tree. VF2 is an improved version of VF, that explores the search space more effectively so that the matching time and the memory consumption are significantly reduced. In Foggia et al. (2001) a detailed experimental analysis of these five algorithms is provided to indicate that

none of the existing algorithms is completely superior to the others. In general, VF2 was found to give best performance with respect to the size and the type of graphs to be matched.

3 Overview of FSM

This section provides a generic overview of the process of FSM. It is widely accepted that FSM techniques can be divided into two categories: (i) Apriori-based approaches, and (ii) pattern growth-based approaches. These two categories are similar in spirit to counterparts found in Association Rule Mining (ARM), namely the Apriori algorithm (Agrawal & Srikant 1994) and FP-growth algorithm (Han et al. 2000) respectively. The Apriori-based approach proceeds in a generate-and-test manner using a Breadth First Search (BFS) strategy to explore the subgraph lattice of the given database. Therefore, before considering $(k + 1)$ subgraphs, this approach has to first consider all k subgraphs. The pattern growth-based adopts a DFS strategy as depicted where, for each discovered subgraph g , the subgraph is extended recursively until all frequent supergraphs of g are discovered (Han & Kamber 2006). The distinction between the two approaches is illustrated in In Figure 4.

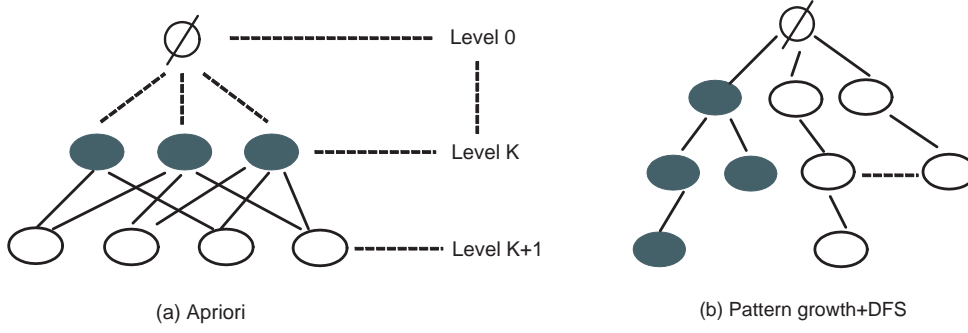


Figure 4 Two types of search space. Note that the subgraph lattice is shown “upside-down”. Vertices corresponding to graphs with fewer edges are displayed at the top of the picture in each case.

Algorithm 3.1: Apriori-based approach

Input: \mathcal{G} = a graph data set, σ = minimum support
Output: F_1, F_2, \dots, F_k , a set of frequent subgraphs of cardinality 1 to k

```

1  $F_1 \leftarrow$  detect all frequent 1 subgraphs in  $\mathcal{G}$ 
2  $k \leftarrow 2$ 
3 while  $F_{k-1} \neq \emptyset$  do
4    $F_k \leftarrow \emptyset$ 
5    $C_k \leftarrow$  candidate-gen( $F_{k-1}$ )
6   foreach candidate  $g \in C_k$  do
7      $g.count \leftarrow 0$ 
8     foreach  $G_i \in \mathcal{G}$  do
9       if subgraph-isomorphism( $g, G_i$ ) then
10         $g.count \leftarrow g.count + 1$ 
11      end
12    end
13    if  $g.count \geq \sigma|\mathcal{G}| \wedge g \notin F_k$  then
14       $F_k = F_k \cup g$ 
15    end
16  end
17   $k \leftarrow k + 1$ 
18 end

```

The basic Apriori-based algorithm is presented in 3.1. In line 5 all frequent $(k - 1)$ subgraphs are used to generate k subgraph candidates. If any of the $k - 1$ candidate subgraphs are not

frequent, then the DCP (see Section 2) can be used to safely prune the candidates. Most existing FSM approaches adopt an iterative pattern mining strategy where each iteration can be divided into two phases: (i) candidate generation (line 5 in Algorithm. 3.1) and (ii) support computation (lines 6-12 in Algorithm. 3.1). Generally, research on FSM focuses on these two phases using a variety of techniques. Since it is harder to address subgraph isomorphism detection, more research effort is directed at how to efficiently generate subgraph candidates. Because subtree isomorphism detection can be solved in $\mathcal{O}(\frac{k^{1.5}}{\log k}n)$ time, the computational complexity is reduced within the context of FSM. Therefore the survey presented in this paper makes a distinction between frequent subgraph mining and frequent subtree mining. In the rest of this paper we will continue to use the acronym FSM to mean both frequent subgraph and subtree mining; and the acronyms FGM and FTM to indicate frequent subgraph and subtree mining respectively where a distinction is required.

Before considering specific subgraph and subtree mining algorithms in detail (Sections 4 and 5), techniques used to represent graphs and trees will first be considered. The aim here is to represent graphs and trees in such a manner that subgraphs can be enumerated efficiently so as to facilitate the desired FSM.

3.1 Canonical representations

The simplest mechanism whereby a graph structure can be represented is by employing an *adjacency matrix* or *adjacency list*. Using an adjacency matrix the rows and columns represent vertexes, and the intersection of row i and column j represents a potential edge connecting the vertexes v_i and v_j . The value held at intersection $\langle i, j \rangle$ typically indicates the number of links from v_i to v_j . However, the use of adjacency matrices, although straightforward, does not lend itself to isomorphism detection, because a graph can be represented in many different ways depending on how the vertexes (and edges) are enumerated (Washio & Motoda 2003). With respect to isomorphism testing it is therefore desirable to adopt a consistent labelling strategy that ensures that any two identical graphs are labelled in the same way regardless of the order in which vertexes and edges are presented (i.e. a *canonical* labelling strategy).

A canonical labelling strategy defines a unique code for a given graph (Read & Corneil 1977; Fortin 1996). Canonical labelling facilitates isomorphism checking because it ensures that if a pair of graphs are isomorphic, then their canonical labellings will be identical (Kuramochi & Karypis 2001). One simple way of generating a canonical labelling is to flatten the associated adjacency matrix by concatenating rows or columns to produce a code comprising a list of integers with a minimum (or maximum) lexicographical ordering imposed. To further reduce the computation resulting from the permutations of the matrix, canonical labellings are usually compressed, using what is known as a *vertex invariant scheme* (Read & Corneil 1977), that allows the content of an adjacency matrix to be partitioned according to the vertex labels. Various canonical labelling schemes have been proposed, some of the more significant are described in this subsection.

Minimum DFS Code (M-DFSC): There are a number of variants of DFS encodings, but essentially each vertex is given a unique identifier generated from a DFS traversal of a graph (DFS subscripting). Each constituent edge of the graph in the DFS code is then represented by a 5-tuple: (i, j, l_i, l_e, l_j) , where i and j are the vertex identifiers, l_i and l_j are the labels for the corresponding vertexes, and l_e is the label for the edge connecting the vertexes. Based on the DFS lexicographic order, the M-DFSC of a graph g can be defined as the canonical labelling of g (Yan & Han 2002). The DFS codes for the left-most branch and the right-most branch of the example graph in Figure 5(c) are $\{(0, 1, a, 1, b), (1, 2, b, 1, e), (2, 3, e, 1, f), (3, 4, f, 1, c), (4, 2, c, 1, e)\}$ and $\{(0, 9, a, 1, d), (9, 10, d, 1, f), (10, 11, f, 1, g), (11, 9, g, 1, d)\}$ respectively.

Canonical Adjacency Matrix (CAM): Given an adjacency matrix M of a graph g , an encoding of M can be obtained by the sequence obtained from concatenating the lower

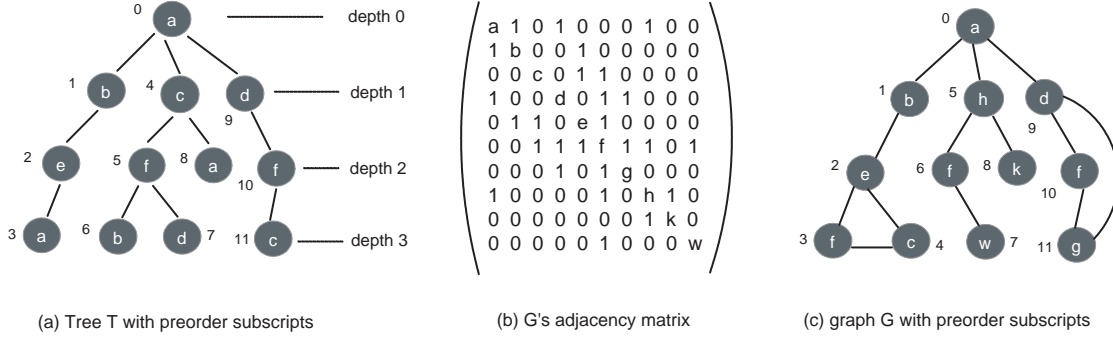


Figure 5 Graph examples to illustrate the canonical representations discussed in Section 3.1 (for ease of illustration, all edge labels are assumed to be the same and represented by “1”).

(or upper) triangular entries of M , including entries on the diagonal. Since different permutations of the set of vertexes correspond to different adjacency matrices, the canonical (CAM) form of g is defined as the maximal (or minimal) encoding. The adjacency matrix from which the canonical form is generated defines the *Canonical Adjacency Matrix* or CAM (Inokuchi et al. 2000,2002; Kuramochi & Karypis 2001; Huan et al. 2003). The encoding for the example graph given in Figure 5(c), represented by the matrix in Figure 5(b), is thus $\{a1b00c100d0110e0001100001011011000101g000010000010h10000000001k0000001000w\}$.

The above two schemes are applicable to any simple undirected graph. However, it is easier to define a canonical labelling for trees than graphs because trees have an inherent structure associated with them. There also exist more specific schemes that are uniquely focused on trees. Among these, DFS-LS and DLS are directed at rooted ordered trees, BFCS and DFCS are used for rooted unordered trees. Each of these will be briefly described below.

DFS Label Sequence (DFS-LS): Given a labelled ordered tree T , the labels of $\forall v_i \in V$ are added to a string S , during a DFS traversal of T . Whenever backtracking occurs a unique symbol, such as “-1” or “\$” or “/”, is added to S (Zaki 2002; Zaki 2005a; Tan et al. 2006). The DFS-LS code for the example tree given in Figure 5(a) is thus $\{abea$$$cfbdaa$$$dfc$$$\}$.

Depth-Label Sequence (DLS): Given a labelled ordered tree T , depth-label pairs comprising the depth and label $\forall v_i \in V$, $(d(v_i), l(v_i))$, are added to a string S , during a DFS traversal of T . The depth-label sequence of T is defined as $S = \{(d(v_1), l(v_1)), \dots, (d(v_k), l(v_k))\}$ (Asai et al. 2002; Wang et al. 2004a). The DLS code for the example tree given in Figure 5(a) is $\{(0, a), (1, b), (2, e), (3, a), (1, c), (2, f), (3, b), (3, d), (2, a), (1, d), (2, f), (3, c)\}$.

Breadth-First Canonical String (BFCS): For a labelled ordered tree, every vertex label is added to a string S , by traversing the tree in a BFS manner. Additionally, a “\$” symbol is used to partition the families of siblings, and a “#” symbol to indicate the end of the string encoding. “\$” is considered to be lexicographically before “#” and both of them order greater than any other vertex and edge labels. Given an unordered tree T , different ordered trees with corresponding BFS string encodings can be produced by imposing different orderings on the children of the intermediate vertexes. The BFCS of T is the lexicographically minimal of these encodings, and the corresponding rooted ordered tree defines the breadth-first canonical form (BFCF) of T (Chi et al. 2005). BFCS’s variants can be found in Chi et al. (2003,2004c). Thus, the BFS string encoding of the example tree given in Figure 5(a) is abcdefafabd$$c\#$.

Depth-First Canonical String (DFCS): Similar to the BFCS but using DFS. The depth-first string encoding, for a labelled ordered tree, labels each vertex by traversing the tree in

a DFS manner. The DFCS of a unordered tree T is then the minimal of all the possible DFS encodings, according to the lexicographical ordering. The corresponding rooted ordered tree defines the depth-first canonical form (DFCF) of T (Chi et al. 2005). DFCS's variants can also be found in Chi et al. (2003,2004b). The DFS string encoding of the example tree given in Figure 5(a) is $abea$$$cfbd$$$a$$$dfc$$$#$.

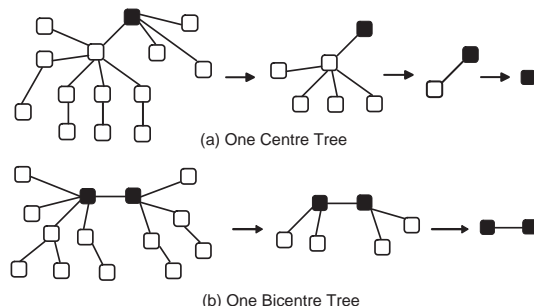


Figure 6 An example of two types of free trees

Canonical Representation of Free Trees: Free trees do not have roots. In this case a unique representation for a free tree is usually constructed by selecting one vertex or a pair of vertices as the root(s). The procedure starts with removing all leaf vertexes and their incident edges recursively until a single vertex or two adjacent vertexes are left. In the first case, the remaining vertex is called the centre, and a rooted unordered tree is obtained with the centre as the root. The procedure is displayed in Figure 6(a). In the second case, the pair of remaining vertexes are called the bi-centre; a pair of rooted unordered trees are obtained with the bi-centre as the roots (along with an edge connecting the two roots). The procedure is illustrated in Figure 6(b). A pair of trees is thus ordered so that the root of the smaller one is chosen as the root of the whole tree (Chi et al. 2003; Rückert & Kramer 2004). After obtaining rooted unordered trees, any canonical representations for rooted unordered trees (see above) can be employed to represent the free trees.

3.2 Candidate generation

As noted earlier, candidate generation is an essential phase in FSM. How to systematically generate candidate subgraphs without redundancy (i.e. each subgraph should be generated only once) is a key issue. Many FSM algorithms can be characterized by the strategy adopted for candidate generation. A number of the most significant are briefly described below. Since a significant proportion of strategies employed in FTM is interwoven with those employed in FGM, no clear distinction can be made between candidate generation strategies in terms of FTM and FGM, i.e. strategies initially proposed for (say) FGM are equally applicable to FTM, and vice versa.

3.2.1 Level-wise join

The level-wise join strategy was introduced by Kuramochi & Karypis (2001). Basically, a $(k + 1)$ subgraph⁶ candidate is generated by combining two frequent k subgraphs which share the same $(k - 1)$ subgraph. This common $(k - 1)$ subgraph is referred to as a *core* for these two frequent k subgraphs. The main issue concerning this strategy is that one k subgraph can have at most k different $(k - 1)$ subgraphs and the joining operation may generate many redundant candidates. In Kuramochi & Karypis (2004a) this issue was addressed by limiting the $(k - 1)$ subgraphs to the two $(k - 1)$ subgraphs with the smallest and the second smallest canonical labels. By carrying

⁶ k refers to the expansion unit for growing the candidate subtrees which can be expressed in terms of vertexes, edges.

out this adapted join operation, the number of duplicate candidates generated was significantly reduced. Other algorithms that adopted this strategy, and its variants, are AGM (Inokuchi et al. 2000), DPMine (Vanetik et al. 2002; Gudes et al. 2006), and HSIGRAM (Kuramochi & Karypis 2005), these will be discussed later.

3.2.2 Rightmost path expansion

Rightmost path expansion is the most common candidate generation strategy, it generates $(k + 1)$ -subtrees from frequent k -subtrees by adding vertexes only to the rightmost path of the tree (Asai et al. 2002; Zaki 2002; Asai et al. 2003; Nijssen & Kok 2003). In Figure 7(a), “RMB” denotes the rightmost branch, which is the path from the root to the rightmost leaf ($k - 1$), and a new vertex k is added by attaching it to any vertexes along the RMB. An enumeration DAG (directed acyclic graph) using rightmost expansion is a tree with a root ϕ , where each node is a subtree pattern. A node S is linked by another node T if and only if T is a rightmost expansion of S . Every 1-subtree is a rightmost expansion of the root ϕ and every $(k + 1)$ -subtree is a rightmost expansion of the k -subtree. Hence, all subtree patterns can be enumerated by traversing in either BFS or DFS manner (Asai et al. 2002). Figure 7(b) shows a part of an enumeration DAG grown by rightmost path expansion. Each square in the figure represents a vertex in the tree. An enumeration DAG (sometimes also simplified as an *enumeration tree*) is used to illustrate how a set of patterns is completely enumerated in a search problem. Enumeration DAGs have been used extensively in Association Rule Mining (Bayardo 1998; Agarwal et al. 2001); and subsequently, in a variety of ways, by many subtree mining algorithms (Asai et al. 2002; Nijssen & Kok 2003; Asai et al. 2003; Chi et al. 2004a; Chi et al. 2005).

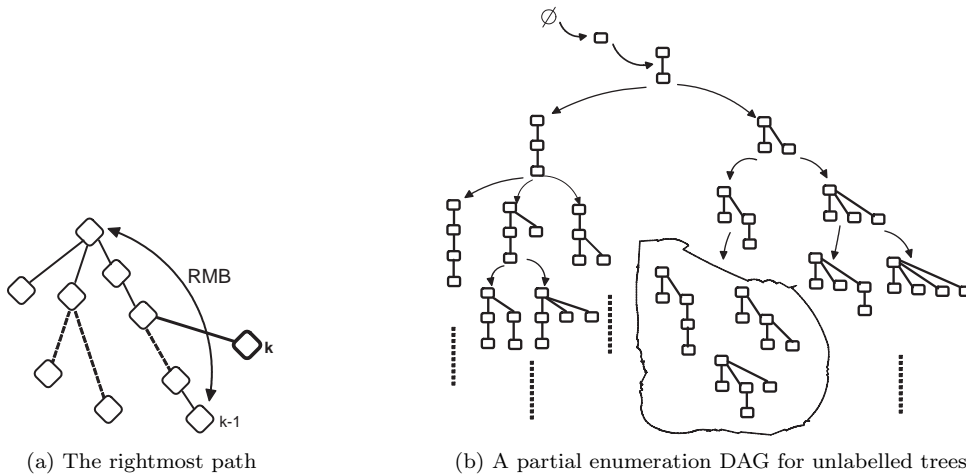


Figure 7 An illustration of rightmost path expansion

3.2.3 Extension and join

The extension and join strategy was first proposed by Huan et al. (2003), and later used by Chi et al. (2004a). It employed a BFCF representation; whereby a leaf at the bottom level of a BFCF tree is defined as a “leg”. For a node “ V_n ” in an enumeration tree, if the height of the BFCF tree corresponding to “ V_n ” is assumed to be h , all children of “ V_n ” can be obtained by either of the following two operations:

- (a) **Extension Operation:** Adding a new leg at the bottom level of the BFCF tree yields a new BFCF with height $h + 1$.
- (b) **Join Operation:** joining “ V_n ” and one of its sibling yields a new BFCF with height h .

3.2.4 Equivalence class based extension

Equivalence class based extension (Zaki 2002,2005) is founded on a DFS-LS representation for trees. Basically, a $(k + 1)$ -subtree is generated by joining two frequent k -subtrees. The two k -subtrees must be in the same *equivalence class* $[C]$ ⁷. An equivalence class consists of the class prefix encoding, and a list of members. Each member of the class can be represented as a (l, p) pair, where l is the k -th vertex label and p is the depth-first position of the k -th vertex's parent. It is verified, in Zaki (2002), that all potential $(k + 1)$ -subtrees with the prefix $[C]$ of size $(k - 1)$ can be generated by joining each pair of members of the same equivalent class $[C]$.

3.2.5 Right-and-left tree join

The right-and-left tree join strategy was proposed by Hido & Kawano (2005). It essentially uses the rightmost leaf (see 2.1) and leftmost leaf⁸ of the tree to generate candidates in a BFS manner. Let $lml(T)$ denote the leftmost leaf of T and $Right(T)$ the right most tree obtained by removing $lml(T)$; and let $rml(T)$ denote the rightmost leaf and $Left(T)$ the left tree obtained by removing $rml(T)$. Given two trees s and t where $Right(s) = Left(t)$, their right-and-left tree join is defined as: $join(s, t) = s \cup rml(t) = lml(s) \cup t$. A diagram depicting this join operation is given in Figure 8.

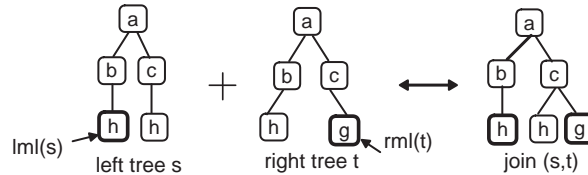


Figure 8 An illustration of right-and-left tree join

Among these candidate generation strategies the level-wise join and the extension and join are directed at FGM, and all others at FTM.

4 Frequent subtree mining algorithms

The previous section considered the joint issues of representation (canonical forms) and candidate generation, in terms of both trees and graphs. In this section a number of prominent FTM algorithms are reviewed. FTM has attracted a great deal of research interest in areas such as: network IP multicast⁹ routing (Cui et al. 2005), web usage mining (Zaki 2005b), computer vision (Liu & Geiger 1999), XML mining (Zaki & Aggarwal 2003; Tan et al. 2005), bio-informatics (Hein et al. 1996; Rückert & Kramer 2004; Zhang & Wang 2006), and so on. The attraction of frequent subtree mining is that subgraph isomorphism detection becomes subtree isomorphism detection, which can be solved in $\mathcal{O}(\frac{k^{1.5}}{\log k}n)$ time (Shamir & Tsur 1999). In addition the structure of trees may be usefully employed to simplify the overall mining process.

The FTM algorithms discussed in this section have been categorized as in Table 3, according to the nature of the trees that the FTM algorithm is directed at: (i) unordered trees, (ii) ordered trees, (iii) free trees, or (iv) hybrid trees (any combinations of (i), (ii) and (iii)). The algorithms are also categorizations according to the nature of the subtrees to be output (*maximal subtrees*, *closed subtrees*, *induced subtrees* or *embedded subtrees*), and the nature of the support metrics employed (transaction-based counting, denoted by T_c ; or occurrences-based counting, denoted by

⁷In Zaki (2002), two k -subtrees T_1, T_2 are in the same “prefix” equivalence class if and only if they share the same encoding up to the $(k - 1)$ -th vertex.

⁸The leftmost leaf of the tree, is the first leaf vertex in the DFS traversal of that tree (Hido & Kawano 2005).

⁹*IP multicast*: a method for building multicast trees at the Internet Protocol layer so as to send packets to multiple receivers in a single transmission (Paul 1998)

Table 3 Categorisation of common frequent subtree mining algorithms

	<i>Maximal</i>	<i>Closed</i>	<i>Induced</i>	<i>Embedded</i>	T_c	O_c
<i>Unordered tree mining</i>						
TreeFinder	*			*	*	
uFreqT			*		*	
cousinPair				*	*	
RootedTreeMiner			*		*	
SLEUTH				*	*	
<i>Ordered tree mining</i>						
FREQT			*		*	
TreeMiner				*	*	
Chopper				*	*	
XSpanner				*	*	
AMIOT			*		*	
IMB3-Miner			*	*		*
TRIPS			*		*	
TIDS			*		*	
<i>Free tree mining</i>						
FreeTreeMiner			*		*	
FTMiner			*		*	
F3TM			*		*	
CFFTTree		*	*		*	
<i>Hybrid tree mining</i>						
CMTreeMiner	*	*	*		*	
HybridTreeMiner			*		*	

O_c). For an alternative review of FTM algorithms readers may like to refer to Chi et al. (2004), who provide a theoretical foundation and performance study of a representative collection of FTM algorithms proposed prior to 2004.

4.1 Unordered tree mining

Labelled unordered trees are often used to model structural data, two popular areas of application are the analysis of chemical compounds and the hyper-link structure of the web (Asai et al. 2003).

Unordered tree FTM tend to use DLS-LS, DLS or BFCS to represent the trees (as described in Subsection 3.1). An often cited example of a DLS-DS based algorithm is the SLEUTH Algorithm (Zaki 2005a). SLEUTH is founded on earlier work directed at the FTM of other types of tree. This algorithm uses *scope-lists* to compute the support. Zaki et al. (2005) considered two extension mechanisms for candidate generation: (i) *class-based extension* and (ii) *canonical extension*. Using class-based extension not all candidates generated by this mechanism necessarily adhere to the desired canonical form, consequently it is necessary to check each candidate subtree to ensure that it is in canonical form. Alternatively, canonical extension can be applied only to canonical frequent subtrees that have a known frequent edge, however this results in many infrequent but canonical candidates. As noted by Zaki et al. (2005a), there is a trade-off between using the two extension mechanisms. Experiments conducted by Zaki et al. (2005a) demonstrated that using class-based extension is more efficient than canonical extension.

An established example of an unordered tree FTM algorithm that uses the DLS representation is uFreqT (Nijssen & Kok 2003). At the candidate generate phase, uFreqT uses the rightmost path expansion technique to generate candidates. At the support counting phase a tree mapping algorithm, to determine the frequency of the current pattern, is translated into a more computationally efficient maximum bipartite matching algorithm. In order to facilitate this support counting, uFreqT maintains a data structure in which to store all potential mappings for the vertexes on the rightmost path and pointers to the parent mappings.

Chi et al. (2005) presented an algorithm, RootedTreeMiner, founded on a BFCS encoding. Unlike SLEUTH or uFreqT, RootedTreeMiner is directed at finding only frequent induced subtrees. Thus, at the candidate generation phase, the range of allowable vertexes at a given position can be computed beforehand. At the support counting phase, an *occurrence list* is first built for each discovered subtree t . This list records the identifier (ID) of each graph transaction in the tree dataset that contains t , along with the mapping between the vertex indexes in t and those in the transaction. Using this occurrence list, the support of t equates to the number of elements in the list which have distinct IDs.

The above all employ exact matching techniques. An example of an unordered tree FTM that employs inexact matching is TreeFinder (Termier et al. 2002). Treefinder employs an Apriori-based approach, using an ancestor-descendant relationships, to mine embedded subtrees. Algorithms that employ inexact matching are, of course, not guaranteed to discover the complete set of frequent subtrees, however, they tend to be very efficient.

Some unordered tree FTM algorithms are directed at specific applications and can use features of these applications to enhance the efficiency of the operation of the algorithm. For example, Shasha et al. (2004) presented an unordered tree FTM algorithm, cousinPair, for application to phylogeny. They defined an interesting pattern as being a “cousin pair”, a pair of vertexes satisfying some cousin distance and minimum occurrence threshold. By using such constraints, interesting patterns were mined from a tree database. The objective here was to get a better understanding of the evolutionary history of species. The obvious disadvantage of algorithms such as cousinPair is that they are not generally applicable.

4.2 Ordered tree mining

In contrast to unordered tree mining, the ordering inherent in ordered trees can be used to introduce efficiencies with respect to subtree generation and subtree isomorphism testing. Candidate subtrees are typically grown using rightmost path extension or equivalence class based extension. For example Asai et al. (2002) use rightmost expansion with respect to their FREQT algorithm. In addition only the rightmost leaf occurrences of the patterns are saved so as to make the support counting more efficient. Asai et al. modelled semi-structure data (namely Web pages) using a labelled ordered tree to evaluate FREQT.

The advantage of rightmost expansion with respect to ordered trees is that the generation of duplicate candidate sets can be avoided. Hido & Kawano (2005) noted that enumeration using rightmost path expansion, as adopted by FREQT and other FSM algorithms, tends to generate many non-frequent candidates thus resulting in unnecessary support counting. Consequently Hido & Kawano introduced an algorithm, AMIOT, to utilize a new enumeration scheme to reduce the number of non-frequent candidates while maintaining the advantage offered by right most expansion. This scheme, *right-and-left tree join*, guaranteed that the set of subtree candidates was always a subset of that achieved by the enumeration using rightmost path expansion. The performance of AMIOT, with respect to both synthetic data and XML data, demonstrated that it was scalable and performed faster than FREQT. However, the memory usage of AMIOT is larger than that of FREQT, due to the nature of the BFS strategy used by AMIOT.

Zaki (2002) proposed a FTM algorithm, TreeMiner, that uses equivalence class based extension (coupled with a DFS-LS representation). The notion of *scope-lists*, later also employed in SLEUTH (see above) was also developed to facilitate fast support counting. Unlike FREQT and AMIOT, TreeMiner is directed at discovering frequent embedded subtrees. The performance of the algorithm was compared with a *base algorithm*, PatternMatcher, which employed a BFS strategy. The experimental results demonstrated that TreeMiner outperformed PatternMatcher when applied to real data. However, the pruning technique adopted by TreeMiner is not as efficient as that used by PatternMatcher given a low support threshold. TreeMiner is a frequently referenced FTM algorithm.

Wang et al. (2004a) also proposed an algorithm, Chopper, to mine frequent embedded subtrees from tree data sets, but used a DLS representation. Chopper first uses a revised PrefixSpan (Pei et al. 2001) to mine frequent sequential patterns. Then the tree database is again scanned, with reference to the discovered sequential patterns, to generate candidate patterns and support counts. The two processes of sequential pattern mining and subtree pattern verification are separated in Chopper, and thus an additional computational overhead is incurred. In order to improve the efficiency of Chopper, the XSpanner algorithm was subsequently produced to integrate the sequential pattern mining into the process of subtree pattern verification. Using projected database techniques, XSpanner grew larger frequent subtrees from smaller ones starting from one vertex. Both Chopper and XSpanner outperform TreeMiner when the support threshold is below 5%. However, XSpanner was found to be more stable than Chopper when the support threshold was further reduced.

IMB3-Miner (Tan et al. 2006) is also directed at frequent embedded subtree mining (from ordered tree datasets), but uses a parameter with which to control the *level of embedding*¹⁰. When the level of embedding is equal to 1, the discovered frequent subtrees are induced subtrees. Thus, by adjusting the embedding level, the algorithm can be used to mine both induced and embedded subtrees. By combining an *Embedding List* data structure with the *TMG* enumeration strategy (a specialized rightmost path expansion strategy), IMB3-Miner guarantees that candidate subtrees are generated without duplication. Furthermore, an occurrence list is stored for each generated subtree to speed-up support counting. Unlike the foregoing, instead of using T_c , O_c is employed to calculate the support of patterns. It has been experimentally demonstrated that IMB3-Miner achieves higher performance and scalability than TreeMiner and FREQT. The usage of O_c , instead of T_c , is typically adopted where the repetition and order of the patterns are important.

Tatikonda et al. (2006) proposed the TRIPS and TIDS algorithms to mining induced or embedded subtrees in a database of rooted ordered trees. TRIPS uses präfer sequencing¹¹ and the leftmost path¹² of the pattern as the extension position. TIDS uses DFS sequencing and rightmost path extension. The support computation for both algorithms employs an *embedding list*, an array based structure, to facilitate the recursive generation of the patterns. There is a trade-off between the cost of maintaining the embedding lists and the efficiency of the support computation, when the number of distinct vertex labels is low compared to the total number of vertexes in the dataset. Experiments demonstrated that both TRIPS and TIDS performed better than TreeMiner in terms of execution time and memory usage on both synthetic and real data sets. Both TRIPS and TIDS were found to be scalable when the database size increased and able to mine large databases even when using low support threshold values.

4.3 Free tree mining

Free tree mining algorithms, as the name suggests, are directed at the discovery frequent subtrees in collections of free trees. An early example is FreeTreeMiner (Chi et al. 2003) where a *self-join* operation was used for candidate subtree generation and a subtree isomorphism algorithm or support computation (Chung 1987). Experiments demonstrated that FreeTreeMiner can handle large real data well with a large range of support values; however, it was not found to be scalable when the size of the maximal frequent subtrees was increased, due to the exponential growth of the subtrees.

Similar work conducted by Rückert and Kramer defined a canonical representation for labelled free trees (Rückert and Kramer 2004). This was embedded in a free tree mining algorithm, FTMiner. The algorithm extended more than one vertex at each recursive step during the

¹⁰The *level of embedding* is defined as “the length of path between two vertexes that form an ancestor-descendant relationship” (Tan et al. 2006).

¹¹The präfer sequence (Prüfer 1918) of a labelled tree on n vertexes is a unique $(n - 2)$ length sequence which can be formed by an iterative algorithm (Tatikonda et al. 2006).

¹²The path from the root to the left most leaf is the *leftmost path* (Tatikonda et al. 2006).

candidate generation phase. It also adopted the concept of an *extension table*, which is a data structure for storing all the extensions for a subtree pattern along with the set of graph transactions containing the pattern. Utilizing this extension table, the algorithm not only kept track of the frequency of each subtree pattern, but also gathered information required for the extension of the current pattern, thus reducing significantly the number of database scans. Experiments on a large scale database suggest that the algorithm is able to mine frequent patterns among a collection of more than 37,330 chemical compounds using support threshold of 2%.

With the focus mainly on reducing the cost of candidate generation, Zhao & Yu (2006) presented the F3TM free tree mining algorithm. The algorithm introduced the idea of an *extension frontier* to define the positions (vertexes) for growing frequent subtrees in the candidate generation phase, and uses automorphism-based pruning and canonical pruning techniques to enhance the efficiency of candidate generation. Compared with other free tree mining algorithms, performance studies indicated that F3TM was more efficient than FTMiner and FreeTreeMiner with respect to a chemical database of 42,390 compounds. CFFTree (Zhao & Yu 2007) is an extension of F3TM directed at closed FTM. CFFTree employs a mechanism called *safe position pruning* to grow subtrees only on safe positions, thus introducing efficiencies when deciding which branch of the enumeration tree to prune. In addition, CFFTree employs *safe label pruning* to grow subtrees only on the vertexes with labels lexicographically less than the new “adding vertex”, this then serves to remove some unnecessary enumeration. The reported evaluation of CFFTree demonstrated that it outperformed its base algorithm, F3TM, using post-processing when finding closed patterns.

4.4 Hybrid tree mining

Hybrid tree mining algorithms are directed at more generic tree formats. As such they can be categorized as being directed at either: (i) unordered or free trees, or (ii) ordered or unordered trees. An example of the first is HybridTreeMiner, an example of the second is CMTreeMiner.

HybridTreeMiner (Chi et al. 2004a) uses a BFCS representation. In the enumeration tree, each node represents an unordered tree in BFCF. For a node v in the enumeration tree, all the children of v may be generated using either an extension or a join operation. The join operation is applied to pairs of sibling nodes with a height (depth) of h , resulting in a BFCF tree with the same height. The extension operation is applied by extending a new leaf at the bottom level of the BFCF tree with a height of h , resulting in a BFCF tree with a height of $(h + 1)$. This hybrid enumeration strategy was further extended to handle the free tree case. Reported experimental results demonstrated that HybridTreeMiner was faster than FreeTreeMiner, and that its memory usage was also much less than that required by FreeTreeMiner.

CMTreeMiner was introduced to mine both *closed* and *maximal* frequent subtrees in collections of labelled ordered or unordered trees (Chi et al. 2004b). By using pruning and heuristic techniques the enumeration tree was grown only on the branches that can potentially produce closed or maximal frequent subtrees, thus avoiding the computational overhead associated with finding all frequent subtrees. The advantage offered by CMTreeMiner is that it directly mines closed and maximal frequent subtrees without first generating all frequent subtrees. Experimental results showed that: (i) for an ordered tree database CMTreeMiner outperformed FREQT, and (ii) for an unordered tree database CMTreeMiner ran faster than HybridTreeMiner.

4.5 Summary of frequent subtree mining algorithms

From the foregoing it can be seen that many different methods, techniques and strategies have been proposed to achieve FTM. From the perspective of applications, these algorithms can be divided into three main domains:

- (a) **Web access Analysis:** Examples are SLEUTH, RootedTreeMiner, TreeMiner, IMB3-Miner, Chopper, XSpanner, TRIPS, TIDS, CMTreeMiner, and HybridTreeMiner.

Table 4 Summary of popular FTM algorithms and their candidate generation and support computation mechanisms

<i>Algorithm</i>	<i>Candidate Generation</i>	<i>Support Computation</i>
TreeFinder	Apriori itemset generation	clustering techniques
uFreqT	rightmost path expansion	maximum bipartite matching
SLEUTH	equivalence class extension	scope-lists
cousinPair	cousin distance	lookup table
RootedTreeMiner	enumeration tree	occurrence list
FREQT	rightmost path expansion	occurrence list
TreeMiner	equivalence class extension	scope list join
Chopper		
XSpanner	n/a	n/a
AMIOT	right-and-left tree join	occurrence list
IMB3-Miner	TMG	occurrence list
TRIPS	leftmost path extension	hash table
TIDES	rightmost path extension	hash table
FreeTreeMiner	self-join	subtree isomorphism
FTMiner	extension tables	support sets
F3TM		
CFFTree	enumeration tree + extension frontier	Ullmann's backtracking algorithm
CMTreeMiner	enumeration tree	n/a
HybridTreeMiner	extension + join	occurrence list

- (b) **IP multicast Analysis:** Examples are FreeTreeMiner and CMTreeMiner.
- (c) **Chemical Compound Analysis:** Examples are FreeTreeMiner, FTMiner, F3TM, CFFTree, and HybridTreeMiner.

From the perspective of the traversing strategy employed in the search space, FTM algorithms can be categorized into two groups:

- (a) **BFS strategy:** The BFS strategy has the advantage of performing full pruning; which, however, requires significant memory usage. Examples are RootedTreeMiner, AMIOT, FreeTreeMiner, and HybridTreeMiner.
- (b) **DFS strategy:** The DFS strategy has the disadvantage of weak pruning. However, the memory usage is smaller than that required for BFS. Examples are uFreqT, SLEUTH, FREQT, TreeMiner, IMB3-Miner, TIDS, FTMiner, and CMTreeMiner.

Table 4 lists the main techniques used for candidate generation and support counting with respect to the algorithms described in this section. Generally, each frequent subtree mining algorithm has its strengths and weaknesses. There is no universally applicable frequent subtree mining algorithm. In terms of FTM efficiency the following techniques are considered to offer the best performance:

- DFS sequence and its variants for tree representation.
- DFS strategy for traversing the search space.
- Enumeration tree growth with rightmost path expansion in candidate generation phase.
- Occurrence list for support counting.

Examples of algorithms containing at least three of these techniques are SLEUTH, FREQT, TreeMiner, and IMB3-Miner. Among these, FREQT, and TreeMiner are usually chosen as base algorithms for comparison with others. TreeMiner is an Apriori-like FTM algorithm, while FREQT is a rightmost path expansion style algorithm. These two styles represent two streams within the realm of FTM. Although subtree isomorphism can be solved in $\mathcal{O}(\frac{k^{1.5}}{\log k}n)$ time, very few frequent subtree mining algorithms adopted it directly for support counting; occurrence lists are more frequently adopted. The main reason for this is that occurrence list counting is much more straightforward to implement.

5 Frequent subgraph mining algorithms

As was indicated in Figure 1(b), FGM algorithms find substantial application in chemical informatics and biological network analysis. There are a variety of FGM algorithms reported in the literature. As in the case of FTM, candidate generation and support counting are key issues. Since subgraph isomorphism detection is known to be *NP*-complete, a significant amount of research work has been directed at various approaches for effective candidate generation. The mechanism employed for candidate generate is the most significant distinguishing feature of such algorithms. An exploration of current well-known frequent subgraph mining algorithms is provided in this section. Interested readers should note that a good review of the theoretical foundation of FGM, prior to 2003, can be found in Washio & Motoda (2003). A more recent review of mining frequent patterns including: itemsets, subsequence, and subgraphs; appears in Han et al. (2007).

For discussion purposes, the FGM algorithms examined in this section are categorized into “general purpose” and “pattern dependent” FGM. The distinction is that in the latter case the nature of the patterns to be discovered is in some way specialized or limited because of the nature of the application domain (e.g. we are only interested in subgraphs satisfying some specific constraints). Consequently, knowledge of the nature of these special patterns allows for a reduction of the search space.

5.1 General purpose frequent subgraph mining

In this subsection a number of general purpose FGM algorithms are considered. To aid the discussion the algorithms are categorized according to three criteria: (i) the completeness of the search (exact search or inexact search), (ii) the type of input (transactions graphs or one single graph), and (iii) the search strategy (BFS or DFS).

5.1.1 Inexact FGM

Inexact search based FGM algorithms use an approximate measure to compare the similarity of two graphs, i.e. any two subgraphs are not required to be entirely identical to contribute to the support count, instead a subgraph may contribute to the support count for a candidate subgraph if it is in some sense similar to the candidate. Inexact search is of course not guaranteed to find all frequent subgraphs, but the nature of the approximate graph comparison often leads to computational efficiency gains. There are only a few examples of inexact frequent subgraph mining algorithms in the literature. However, one frequently quoted example is the SUBDUE algorithm (Cook & Holder 1994,2000). SUBDUE uses the minimum description length principle to compress the graph data; and a heuristic beam search method, that makes use of background knowledge, to narrow down the search space. Although the application of SUBDUE shows some promising results in domains such as image analysis and CAD circuit analysis, the scalability of the algorithm is an issue, i.e. the run time does not increase linearly with the size of the input graph. Furthermore, SUBDUE tends to discover only a small number of patterns.

Another inexact search based FGM algorithm is GREW (Kuramochi & Karypis 2004b). However, GREW is directed at finding connected subgraphs which have many vertex-disjoint embeddings¹³, in single large graphs. GREW uses a heuristic based approach that is claimed to be scalable, because it employs ideas of *edge contraction* and *graph rewriting*. GREW deliberately underestimates the frequency of each discovered subgraph in an attempt to reduce the search space. Experiments on four benchmark data sets showed that GREW significantly outperformed SUBDUE with respect to: runtime, number of patterns found, and size of patterns found.

To the best knowledge of the authors, the two most recent inexact search based FGM algorithms are gApprox (Chen et al. 2007a) and RAM (Zhang & Yang 2008). The gApprox algorithm uses the notion of an upper-bound for support counting, and an approximation measure

¹³Two embeddings in a graph G are *vertex-disjoint*, if they do not share any vertexes in G .

to discover frequent approximately connected subgraphs in very large networks. Empirical studies based on protein-protein interaction networks indicated that gApprox is efficient and that the discovered patterns were biological meaningful. RAM is founded on a formal definition of frequent approximate patterns in the context of biological data represented as graphs, where the edge information tended to be inaccurate. Reported experiments showed that RAM can discover some important patterns which can not be found by exact search based mining algorithms.

5.1.2 Exact FGM

Exact FGM algorithms are much more common than inexact search based FGM algorithms. They can be applied in the context of graph transaction based mining or single graph based mining. A fundamental feature for exact search based algorithms is that the mining is complete, i.e. the mining algorithms are guaranteed to find all frequent subgraphs in the input data. As noted in Kuramochi & Karypis (2004b), such complete mining algorithms perform efficiently only on sparse graphs with a large amount of labels for vertexes and edges. Due to this completeness restriction, these algorithms undertake extensive subgraph isomorphism comparison, either explicitly or implicitly, resulting in a significant computational overhead.

We will commence the discussion of exact FGM algorithms by considering graph transaction based FGM, the mining of collections of relatively small graphs; single graph based FGM will be considered at the end of this subsection. With respect to graph transaction mining, the algorithms can be divided into two groups: BFS and DFS, according to the traversing strategy adopted. BFS tends to be more efficient in that it allows for the pruning of infrequent subgraphs (at the cost of high I/O and memory usage) at an early stage in the FGM process, whereas DFS requires less memory usage (in exchange for less efficient pruning). We will consider the BFS algorithms first.

As in the case of Association Rule Mining algorithms, such as Apriori (Agrawal & Srikant 1994), BFS based FGM algorithms utilize the DCP, i.e. a $(k + 1)$ subgraph can not be frequent if its immediate parent k subgraph is not frequent. Using BFS the complete set of k candidates is processed before moving on to the $(k + 1)$ candidates, where k refers to the expansion unit for growing the candidates, which can be expressed in terms of vertexes, edges, or disjoint paths. Four well-established exact FGM algorithms are itemized below:

- AGM (Inokuchi et al. 2000) is a well established algorithm used to identify frequent induced subgraphs. AGM uses an adjacency matrix to represent graphs and a level-wise search to discover frequent subgraphs. AGM assumes that all vertexes in a graph are distinct. The evaluation of AGM on chemical carcinogenesis data demonstrated that it was more efficient than an inductive logic programming based approach combined with a level-wise search. AGM discovers not only connected subgraphs, but also unconnected subgraphs with several isolated graph components. A more efficient version of AGM, called AcGM, has also been developed to mine only frequent connected subgraphs (Inokuchi et al. 2002). AcGM uses the same principles and graph representation as AGM. Experimental results indicate that AcGM is significantly faster than AGM and FSG (see below). Inokuchi et al. have further extended their original work to mine frequent induced subgraphs from general graph databases that can contain directed (or undirected), labelled (or unlabelled) graphs and even loops (Inokuchi et al. 2003).
- FSG (Kuramochi & Karypis 2001, 2004a) is directed at finding all frequent connected subgraphs. FSG uses the BFS strategy to grow candidates whereby pairs of identified frequent k subgraphs are joined to generate $(k + 1)$ subgraphs. FSG uses a canonical labelling method for graph comparison and computes the support of the patterns using a vertical *transaction list* data representation, which has been used extensively in FTM. Experiments show that FSG does not perform well when graphs contain many vertexes and edges that have identical

labels because the *join* operation used by FSG allows multiple automorphism¹⁴ of single or multiple cores¹⁵.

- The FSG algorithm is directed at graph databases consisting of a two dimensional arrangement of vertexes and edges in each graph (sometimes referred to as topological graphs). However, in chemical compound analysis users are often interested in graphs that have coordinates associated with the vertexes in two-or three-dimensional space (sometimes referred to as geometric graphs). gFSG (Kuramochi & Karypis, 2002) extends the FSG algorithm to discover frequent geometric subgraphs with some degree of tolerance among geometric graph transactions. The extracted geometric subgraphs are rotation, scaling and translation invariant. gFSG shares the approach of candidate generation with FSG. In order to speed up the computation of geometric isomorphism, a number of topological properties and geometric transform invariants are used in the matching process. In the process of support counting, geometric transform invariants (such as an edge-angle list¹⁶), and transaction lists are used to facilitate the computation. Experimental evaluation was performed using a chemical database with more than 20,000 chemical compounds to show that gFSG operated well with low support values and scaled linearly with respect to data size.
- DPMine (Vanetik et al. 2002; Gudes et al. 2006) uses edge-disjoint paths as the expansion units for candidate generation. Use of a large expansion unit reduces the number of candidates that are generated. DPMine firstly identifies all frequent paths, secondly it finds all subgraphs with two paths, and thirdly merges pairs of frequent subgraphs with $(k - 1)$ paths, which have $(k - 2)$ paths in common, in order to obtain subgraphs with k paths. Experimental results indicated that the support computation was the most significant contributor to the overall computation time. Gudes et al. (2006) also suggested that reducing the support computation overhead is more important than reducing the candidate generation computation overhead. (DPMine can operate on both graph transaction based and single graph based data.)

FGM algorithms that adopt a DFS strategy tend to need less memory because they traverse the lattice of all possible frequent subgraphs in a DFS manner. Five well-known example algorithms are listed below:

- MoFa (Borgelt and Berthold, 2002) is directed at mining frequent connected subgraphs describing molecules. The algorithm stores the embedding list of previously found subgraphs and the extension operation is restricted only to these embeddings. MoFa also uses structural pruning and background knowledge to reduce support computation. However MoFa still generates many duplicates, resulting in unnecessary support computation.
- gSpan (Yan & Han 2002) uses a canonical representation, M-DFSC, to uniquely represent each subgraph. The algorithm uses DFS lexicographic ordering to construct a tree-like lattice over all possible patterns, resulting in a hierarchical search space called a DFS code tree. Each node of this search tree represents a DFS code. The $(k + 1)$ -th level of the tree has nodes which contain DFS codes for k subgraphs. The k subgraphs are generated by one edge expansion from the k -th level of the tree. This search tree is traversed in a DFS manner and all subgraphs with non-minimal DFS codes are pruned so that redundant candidate generations are avoided. Instead of keeping the embedding list, gSpan only preserves the transaction list for each discovered pattern; subgraph isomorphism detection only operates on the graphs within the list. In comparison with embedding list based algorithms, the gSpan algorithm saves on memory usage. Experiments show that gSpan outperforms FSG by an order of magnitude. gSpan is arguably the most frequently cited FSM algorithm.

¹⁴*Automorphism* is a graph isomorphism to itself via a non-identity mapping.

¹⁵In the candidate generation phase, a *core* is a common $(k - 1)$ subgraph shared by two frequent k subgraphs. Two frequent k subgraphs are eligible for joining only if they contain the same core (Kuramochi & Karypis 2001).

¹⁶An edge-angle list is a multi-set where each element represents the angle formed by two distinct edges sharing the same end points.

- ADI-Mine (Wang et al., 2004b) addresses the issue of mining large disk-based graph data sets. ADI-Mine uses a general index structure, *ADI*. Experiments have indicated that ADI-Mine can mine graph data sets with one million graphs, while gSpan could only mine databases with 300,000 graphs.
- FFSM (Huan et al. 2003) is directed at graphs that are large and dense with a small number of labels. For example, protein structure mining. FFSM adopted the CAM representation. Thus, a tree-like structure, a suboptimal CAM tree, was constructed to include all possible patterns. Each node in this suboptimal CAM tree can be enumerated by either a join or an extension operation. FFSM records embedding lists for each discovered pattern to avoid explicit subgraph isomorphism testing in the support counting phase. Performance evaluation, using several chemical data sets, indicated that FFSM outperformed gSpan.
- GASTON integrates frequent path, subtree, and subgraph mining into one algorithm, due to the observation that most frequent sub-structures in molecular databases are free trees (by Nijssen & Kok, 2004). The algorithm provided a solution by splitting up the frequent subgraph mining process into path mining, then subtree mining, and finally subgraph mining. Consequently the subgraph mining is only invoked when needed. Thus, GASTON operates best when the graphs are mainly paths or trees, because the more expensive subgraph isomorphism testing is only encountered in the subgraph mining phase. GASTON records the embedding list so as to grow only patterns that actually appear; thus saving on unnecessary isomorphism detection. Experiments show that GASTON is at a competitive level with a wide range of other FGM algorithms.

Because of the diversity of FGM algorithms, it is difficult to enumerate the strong and weak points of the various algorithms. However, Wörlein et al. (2005) presented a detailed comparison of four DFS-based miners: MoFa, gSpan, FFSM, and GASTON, with respect to their performance on various chemical data sets. In the experiments, they found that the use of embedding lists did not offer significant gains at the expense of memory usage. They also confirmed that using canonical representations for duplicate detection required less computation than explicit subgraph isomorphism detection. By utilizing the two main distinguishing features of molecular data, “symmetries in molecules” and “non-uniform frequency distribution of atom and bond types”, Jahn and Kramer (2005) optimized the performance of gSpan with respect to the mining molecular databases.

We complete this subsection by considering *Single graph* based exact FGM algorithms where the frequency of a pattern is determined by occurrence based counting (we have already noted that DPMine can operate on both graph transaction based and single graph based data). A fundamental issue regarding single graph based mining is how to define the support of the pattern. The DCP, often used to prune the search space when using transaction based counting, does not hold in the case of occurrence based counting. Thus, occurrence based support measures that satisfy the DCP are desirable. One well established type of occurrence based support measure that maintains the DCP is founded on the concept of *overlap graphs*¹⁷. By building an overlap graph for each pattern, the occurrence based support measure is defined as the size of the Maximum Independent Set (MIS) of vertices in the overlap graph. The MIS measure was first introduced in Vanetik (2002) and Kuramochi & Karypis (2004c,2005). In Vanetik et al. (2006), the formal definition was provided together with proofs for the sufficient and necessary conditions required for occurrence based support measures to maintain the DCP. Their work was further extended to introduce a new occurrence based support measure, which maintained the DCP, and was computable in polynomial time (Calders et al. 2008).

Kuramochi and Karypis (2004c, 2005) proposed two algorithms: HSIGRAM and VSIGRAM to find all frequent subgraphs in a large sparse graph. These two algorithms used BFS and DFS

¹⁷An *overlap graph* for a given pattern with a set of all embeddings (occurrences) is a constructed graph where each vertex represents a non-identical embedding of the pattern; two vertices are connected if the corresponding embeddings overlap (Kuramochi & Karypis 2005).

strategies respectively and the support of each pattern was determined by the overlap graph based MIS measure (Vanetik et al. 2006). Several variations of the MIS measures, including exact and approximate MIS measures, were implemented. Experiments demonstrated that both algorithms scaled well when mining large graphs, although VSIGRAM was faster than HISGRAM. The reason for the performance advantage of the VSIGRAM algorithm is that it keeps track of the embeddings of the frequent subgraphs along the DFS path, resulting in less subgraph isomorphism checking. In comparison with SUBDUE, the results indicated that SUBDUE performed worse than both HSIGRAM and VSIGRAM; SUBDUE tends to focus on small subgraphs with high frequency and consequently tends to miss significant patterns. Kuramochi and Karypis' work was further extended by Schreiber and Schwöbbermeyer (2005) to mine frequent patterns of a given size, but considering alternative frequency concepts. This frequency based algorithm, FPF, was applied to two different biological networks to discover network motifs¹⁸. Surprisingly, a comparison of the number of frequent patterns found using the alternative frequency concepts demonstrated that the frequency of a pattern alone was not sufficient to identify network motifs, and that it was not clear whether frequent patterns could play functional roles in the biological network.

5.2 Pattern dependent frequent subgraph mining

In FSM users are usually interested in a certain type of pattern, rather than the complete set of patterns, i.e. some subset of the set of all frequent subgraphs. Such “special patterns” are diagnosed according to their topology and/or some specific constraint (limitation) on the nature of the patterns. Pattern dependent FGM algorithms can be grouped according to nature of the patterns they are directed at: (i) relational patterns, (ii) maximal and closed patterns, (iii) cliques and (iv) other constrained patterns. Each is discussed in more detail below.

5.2.1 Relational pattern mining

Relational graphs are suitable for modelling large scale networks such as biological or social networks. Yan et al. (2005a) indicated that relational pattern mining has three features which serve to differentiate it from general purpose frequent subgraph mining: (i) the data has distinct vertex labels, (ii) the data comprises very large graphs, and (iii) a focus on frequent patterns with certain *connectivity constraints* (e.g. the minimum degree¹⁹ of a pattern). Thus relational graph mining aims to identify all frequent patterns displaying a specified connectivity constraint.

CLOSECUT and SPLAT, both proposed by Yan et al. (2005a), are directed at mining (closed) frequent subgraphs with connectivity constraints. CLOSECUT uses a pattern growth approach to integrating connectivity constraints, together with graph condensation and decomposition techniques. The SPLAT algorithm uses a pattern reduction approach to integrating the graph decomposition technique. Experiments indicated that CLOSECUT performed better than SPLAT on patterns with low connectivity when using a high support threshold value; however SPLAT performed better than CLOSECUT on patterns with high connectivity when using a low support threshold value. The results, with respect to biological data, showed that both algorithms could find interesting patterns with strong biological meanings.

5.2.2 Mining maximal and closed patterns

The number of possible frequent subgraphs increases exponentially with the size of the graph, i.e., for a frequent k -graph, the number of its frequent subgraphs can be as large as 2^k . In (Yan & Han 2003) it was observed that about 1,000,000 frequent graph patterns were generated from 422 chemical compounds (using a support threshold of 5%); amongst these many were

¹⁸*Network Motifs* are defined as “patterns of interconnections occurring in complex networks at numbers that are significantly higher than those in randomized networks” (Milo et al. 2002).

¹⁹The *minimum degree* of a pattern g is the minimum of the degree of v , for all $v \in V(g)$ (Yan et al. 2005a).

found to be structurally repetitive. Therefore, both *closed* and *maximal* FGM approaches have been proposed as mechanisms to limit the number of frequent subgraphs generated. These approaches are discussed further in this subsection. The following notation is used, MFS denotes the set of maximal frequent subgraphs, CFS the set of closed frequent subgraphs, and FS the set of all frequent subgraphs in the graph database. Thus, $MFS \subseteq CFS \subseteq FS$.

Let $MFS = \{g | g \in FS \wedge \neg(\exists h \in FS \wedge g \subset h)\}$. The task of maximal frequent subgraph mining is to find all graphic patterns that belong to MFS . Maximal frequent subgraphs encode the maximal common structures, in the case of biological networks these are deemed to be the most interesting patterns (e.g. Koyutürk et al. 2004). However, the frequency of non-maximal subgraphs is not produced. Two example maximal FGM algorithms are SPIN and MARGIN.

SPIN (Huan et al. 2004) is a spanning tree based frequent subgraph mining algorithm designed to discover only maximal frequent subgraphs with the intention of reducing the overall computation cost. The concept of *tree-based equivalence classes* is introduced by the notion of a *canonical spanning tree*²⁰. In SPIN, the graph partitioning method utilizes such tree-based equivalence classes together with three pruning techniques. The algorithm has two main phases: (i) identification of all frequent subtrees within the input data using an appropriate frequent subtree mining algorithm, and (ii) the detection of all frequent subgraphs whose canonical spanning tree is isomorphic to each discovered frequent subtree. The desired maximal frequent subgraphs are generated by optimized post-processing. The performance of SPIN has been compared with gSpan and FFSM. Results demonstrated that SPIN displayed significantly better performance than gSpan and FFSM with respect to both synthetic and chemical data.

MARGIN (Thomas et al. 2006) was founded on the observation that the set of potential maximally frequent subgraphs is included in the set of frequent k subgraphs that have infrequent $(k + 1)$ supergraphs. Consequently the search space of MARGIN is significantly reduced by pruning the lattice around the set of potential maximally frequent subgraphs. The set of candidates is recursively discovered by the core algorithm, *ExpandCut*, and the maximal frequent subgraphs are then found by the post-processing operation. Experimental results showed that MARGIN was computationally faster than gSpan when applied to some databases. However, the efficiency of MARGIN largely relies on the initial *cut*²¹.

Let $CFS = \{g | g \in FS \wedge \neg(\exists h \in FS \wedge g \subset h \wedge sup(g) = sup(h))\}$. The task of closed frequent subgraph mining is to find all patterns that belong to CFS . These closed patterns have some biological meaning, because generally, a biochemist is only interested in the largest structures with certain properties (Fischer & Meinel 2004). CLOSECUT and SPLAT are two example closed FGM algorithms that have already been considered (see Subsection 5.2.1). Another example is CloseGraph (Yan & Han 2003), which is founded on gSpan. CloseGraph uses an equivalent occurrence based early termination to prune the search space. For the case where the early termination fails and can not be applied, the detection of the failure of early termination is implemented. Experimental results demonstrated that CloseGraph performed better than gSpan and FSG.

5.2.3 Mining cliques

A clique (or quasi-clique) is a subset of one subgraph with a fixed topology. The first algorithm directed at detecting cliques was proposed by Harary & Ross (1957). Since then many more algorithms have been devised directed at a variety of the clique detection problems (Bomze et al. 1999; Gutin 2004). More recently it has been found that discovering frequent cliques from a set of graph transactions is useful in domains such as communication, finance, and bio-informatics. Example applications where the mining of cliques, or quasi-cliques, has been applied includes:

²⁰A *canonical spanning tree* of a graph is defined as the lexicographically maximal spanning tree of the graph (Huan et al. 2004).

²¹A *cut* between two nodes in the lattice is defined as an ordered pair (p, c) where node p represents the parent of node c and p is infrequent while c is frequent (Thomas et al. 2006).

community mining (Abello et al. 2002), gene expression mining (Pei et al. 2005), and the discovery of highly correlated stocks from stock market graphs (Wang et al. 2006). General purpose FGM algorithms can be used to discover such “special patterns”, however the computation can be made more efficient if the special properties of cliques is taken into account. Two examples clique mining algorithms, CLAN and Cocain, are discussed in the following paragraphs.

CLAN (Wang et al. 2006) is directed at mining frequent closed cliques²² from large dense graph databases. The algorithm utilized the properties of the clique structure to facilitate clique or sub-clique isomorphism testing by introducing a canonical representation of a clique. Wang et al. also devised several pruning techniques to effectively reduce the search space. The experimental results showed that CLAN can efficiently mine large and dense graph data sets. However, the reported evaluation only used high support threshold values, and scalability was demonstrated using only a small and sparse graph data set.

Extending the work of CLAN, Zeng et al. introduced a general form of clique mining algorithm, Cocain (Zeng et al. 2006), to mine closed γ -quasi-cliques²³ from large and dense graph data sets. In Cocain, cliques are required to satisfy a user specified parameter, γ . Cocain utilized the properties of quasi-cliques to prune the search space, combined with a closure checking scheme to speed up the discovery process. However, the reported evaluation of Cocain was only directed at US stock market data.

5.2.4 Constrained pattern mining

The main idea of user constraint based frequent pattern mining is to integrate constraints into the mining process in order to prune the search space. Zhu et al. (2007) presented a framework, called gPrune, to incorporate various constraints into the frequent subgraph mining process. In gPrune, the search spaces of both data and patterns were examined and a new concept, *pattern-inseparable-data-anti-monotonicity* (D-anti-monotonicity), introduced to support effective pruning of the search space. However, an empirical study showed that the benefit of such D-anti-monotonicity pruning was coupled tightly with the speed of the corresponding constraint measure function. Furthermore, experiments indicated that the effectiveness of integrating constraints into the FSM process is influenced by many aspects, including the properties of the data and the pruning cost. Constraint based mining algorithms therefore need to take into account the trade-off between the pruning cost and any potential benefit.

5.3 Summary

Table 5 summarizes the approaches to canonical representation, candidate generation and support counting used by the FGM algorithms described in this section. With respect to the FGM algorithms listed in the Table SUBDUE, AGM, FSG, MoFa, gSpan, FFSM, and GASTON are the most frequently cited. Among these algorithms, SUBDUE is used more widely than others. However, one frequently quoted disadvantage of SUBDUE is that the algorithm tends to find only small sizes patterns, consequently it may miss interesting larger patterns. AGM and FSG are two representative BFS-based miners. MoFa is a specialized miner for molecular data and is able to mine directed graphs. FFSM and GASTON can not be used in the context of directed graphs; while gSpan, with some minor changes, can accommodate directed graphs.

One common feature for the majority of algorithms in the table is that the search space is usually modelled as a tree-like lattice over all possible patterns, which are ordered lexicographically. Each node in the lattice represents a pattern, and the relationship between patterns at levels $(k + 1)$ and k will only differ by one vertex or edge (i.e. there is a “parent-child” relation). Search strategies therefore comprise a traversal of the lattice and the storing all patterns that

²²Let $V(g)$ denote the set of vertexes in a graph g , a subset $s \subseteq V(g)$ is a *clique* if the subgraph induced on s is a complete graph.

²³A γ -*quasi-clique* ($0 \leq \gamma \leq 1$) is a k subgraph ($k \geq 1$), g , where $\forall v \in V(g)$, $degree(v) \geq \lceil \gamma(k - 1) \rceil$ (Zeng et al. 2006).

Table 5 Summary of popular FGM algorithms and their candidate generation and support computation mechanisms

<i>Algorithm</i>	<i>Representation</i>	<i>Candidate Generation</i>	<i>Support Computation</i>
AGM/AcGM	CAM	level-wise join	database scan
FSG	CAM	level-wise join	transaction list
gFSG	n/a	level-wise join	edge-angle list transaction list hybrid
DPMine	n/a	level-wise join	n/a
MoFa	n/a	extension	embedding list
gSpan	M-DFSC	rightmost path extension	transaction list
ADI-Mine	M-DFSC	rightmost path extension	transaction list
FFSM	CAM	join + extension	embedding list
GASTON	n/a	path,tree, and graph enumeration	embedding list
HSIGRAM	CAM	level-wise join	various MIS measures
VSIGRAM	CAM	extension	various MIS measures
FPF	CAM	extension	MIS measure
DPMine	n/a	level-wise join	n/a
CLOSECUT	M-DFSC	rightmost path extension	transaction list
SPLAT	n/a	n/a	n/a
SPIN	n/a	join	embedding set
MARGIN	n/a	ExpandCut	n/a
CloseGraph	M-DFSC	rightmost path extension	transaction list
CLAN	vertex label sequence	DFS-based extension	n/a
Cocain	vertex label sequence	DFS-based extension	n/a
gPrune	M-DFSC	rightmost path extension	transaction list

satisfy some threshold. Either a BFS or DFS strategy can be used to traverse the lattice. BFS strategy based miners offer the advantage over DFS strategy based miners that they can obtain a much “tighter” upper bound for the support of k subgraphs from the support associated with the complete set of identified $(k - 1)$ subgraphs. Knowledge of this upper bound can be employed to limit the number of candidate subgraphs that are generated. DFS strategy based miners typically derive an upper bound for K candidates based only on a single $k - 1$ parent frequent subgraph.

As indicated in Wörlein et al. (2005), an efficient FGM algorithm usually displays three distinct features:

- **Restrictive extension:** The extension of a subgraph is valid only when the extension exists in the graphs within the subgraph’s occurrence list. Examples of such operations are rightmost path expansion as used by gSpan, and rightmost path extension as used by MoFa.
- **Efficient candidate generation:** This operation is achieved by using a canonical graph representation. Such representation can facilitate the filtering out of candidate duplicates before performing graph isomorphism testing. The two main canonical representations are: (i) CAM, used by AGM, FSG, and FFSM; and (ii) M-DFSC used by gSpan.
- **Essential subgraph isomorphism:** When computing the support of a pattern a trade-off needs to be sought between using explicitly subgraph isomorphism and keeping embeddings of the pattern. Examples of keeping embeddings are FFSM and GASTON, and instances of using subgraph isomorphism are FSG and gSpan.

Although distributed algorithms can offer a distinct advantage with respect to excessively large databases, very few researchers have used such algorithms for FGM. One example, proposed by Fatta & Berthold (2005), is an extension of MoFa to accommodate the distributed computation of mining frequent patterns with respect to data sets representing large molecular compounds.

6 Discussion and conclusion

A view of the “state of the art” of current FSM, referencing especially those algorithms most frequently referred to in the literature, has been presented. The most computationally expensive aspects of FSM algorithms are candidate generation and support computation; with the latter being the most computationally expensive. Broadly the distinguishing feature of the mining algorithms considered in this survey is how they efficiently address candidate generation and support counting.

With reference to the literature many different mining strategies have been proposed with respect to many different types of graph, to produce many different kinds of patterns. So as to impose some structure on the wide range of FSM algorithms featured in the literature we have adopted a categorisation whereby FSM algorithms are considered according to that: (i) candidate generation strategy, (ii) search strategy and (iii) approach to frequency counting. Generally FTM algorithms can not be directly applied to graphs, while FGM algorithms can be applied to both graphs and trees. FTM and FGM algorithms have been developed differently for different purposes. Thus in this survey we describe these two types of algorithms separately. For FTM algorithms common applications are web usage and XML mining; while FGM algorithms tend to be directed at chem- and bio-informatics. Although there are abundant research publications on FGM applications many important issues remain to be addressed.

Firstly, can we discover a compact and meaningful set of frequent subgraphs instead of a complete set of frequent subgraphs? A lot of research effort has been directed at reducing the resultant set of frequent subgraphs; for example the use of maximal frequent subgraphs, closed frequent subgraphs, approximate frequent subgraphs and discriminative frequent subgraphs. However, there is no clear understanding of what kind of frequent subgraphs are the most compact and representative for any given application. In many cases the resultant set of frequent subgraphs are too large to be analysed individually and many of the identified frequent subgraphs are often found to be structurally repetitive. Research work focusing on how to significantly reduce the size of the resultant set of frequent subgraphs is much in demand.

Secondly, can we achieve better classification using frequent subgraph based classifiers than other approaches? Can we integrate *feature selection* techniques deeply into the frequent subgraph mining process and directly identify the most discriminative subgraphs which are effective for classification? There is still much room for researchers to utilize classic data mining techniques and integrate them into the FSM process.

Thirdly, as many researchers have noted, exact frequent subgraphs are not very helpful with respect to many real application. Can we therefore devise more efficient algorithms to generate *approximate frequent subgraphs*? Little work has been conducted in the context of approximate frequent subgraphs mining with the notable exception of the well-known SUBDUE algorithm.

Finally, in domains like: document image classification, work-flow mining, social network mining, single graph based mining, and so on; there is still a lot of work that can be done to improve the mining task. There is always a trade-off between the combinatorial complexity of FSM algorithms and the utility of the frequent subgraphs discovered by them. Much work is needed to circumvent this issue. Is the *frequency* of a subgraph really a good measure for discovering interesting subgraphs? Can we devise other interestingness metrics for subgraph discovery, rather than adopting those from the domain of Association Rule Mining?

References

- Abello, A., Resende, M.G.C. and Sundarsky, S. 2002. Massive Quasi-Clique Detection, In *Proceedings of the 5th Latin America Symposium on Theoretical Informatics*, 598–612.
- Agrawal, R. and Srikant, R. 1994. Fast Algorithm for Mining Association Rules, In *Proceedings of the 20th International Conference on Very Large Databases (VLDB)*, pp.487–499, Morgan Kaufmann.
- Agrawal, R.C., Aggarwal, C.C. and Prasad, V.V.V. 2001. A Tree Projection Algorithm for Generation of Frequent Itemsets, *Journal of Parallel and Distributed Computing* **61**(3), 350–371.

- Alm, E. and Arkin, A.P. 2003. Biological Networks, *Current Opinion in Structural Biology* **13**(2), 193–202.
- Aldous, J.M. and Wilson, R.J. 2000. *Graphs and Applications, an Introductory Approach*, Springer.
- Asai, T., Abe, K., Kawasoe, S., Arimura, H., Satamoto, H. and Arikawa, S. 2002. Efficient Substructure Discovery from Large Semi-Structured Data, In *Proceedings of the 2nd SIAM International Conference on Data Mining*, 158–174.
- Asai, T., Arimura, H., Uno, T. and Nakano, S. 2003. Discovering Frequent Substructures in Large Unordered Trees, In *Proceedings of the 6th International Conference on Discovery Science*, 47–61.
- Bayardo Jr., R.J. 1998. Efficiently Mining Long Patterns from Databases, In *Proceedings of the 1998 International Conference on Management of Data*, 85–93.
- Borgelt, C. and Berthold, M. 2002. Mining Molecular Fragments: Finding Relevant Substructures of Molecules, In *Proceedings of International Conference on Data Mining*, 211–218.
- Borgwardt, K.M. and Kriegel, H.P. 2005. Shortest-Path Kernels on Graphs, In *Proceedings of the 2005 International Conference on Data Mining*, 74–81.
- Bomze, I.M., Budinich, M., Pardalos, P.M. and Pelillo, M. 1999. The Maximum Clique Problem, *Handbook of Combinatorial Optimization*, Kluwer Academic Publishers **4**, 1–74.
- Brin, S. and Page, L. 1998. The Anatomy of a Large-scale Hyper-textual Web Search Engine, In *Proceedings of the 7th International World Wide Web Conference*, 107–117.
- Bunke, H. and Shearer, K. 1998. A Graph Distance Metric based on the Maximal Common Subgraph, *Pattern Recognition Letters* **19**, 225–259.
- Bunke, H. and Allerman, G. 1983. Inexact Graph Matching for Structural Pattern Recognition, *Pattern Recognition Letters* **1**(4), 245–253.
- Calders, T., Ramon, J. and van Dyck, D. 2008. Anti-monotonic Overlap-graph Support Measures, In *Proceedings of the Eighth IEEE International Conference on Data Mining*, 73–82.
- Chakrabarti, S., Dom, B., Gibson, D., Kleinberg, J., Kumar, R., Raghavan, P., Rajagopalan, S. and Tomkins, A. 1999. Mining the Link Structure of the World Wide Web, *IEEE Computer* **32**(8), 60–67.
- Chen, C., Yan, X., Zhu, F. and Han, J. 2007a. gApprox: Mining Frequent Approximate Patterns from a Massive Network, In *Proceedings of the 7th IEEE International Conference on Data Mining*, 445–450.
- Chen, C., Yan, X., Yu, P.S., Han, J., Zhang, D. and Gu, X. 2007b. Towards Graph Containment Search and Indexing, *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB'07)*, 926–937.
- Chen, M.S., Han, J. and Yu, P.S. 1996. Data Mining: An Overview from Database Perspective, *IEEE Transaction on Knowledge and Data Engineering* **8**, 866–883.
- Chen, C., Lin, C.X., Yan, X. and Han, J. 2008. On Effective Presentation of Graph Patterns: A Structural Representative Approach, In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, 299–308.
- Chi, Y., Yang, Y., Xia, Y. and Muntz, R.R. 2003. Indexing and Mining Free Trees, In *Proceedings of the 2003 IEEE International Conference on Data Mining*, 509–512.
- Chi, Y., Nijssen, S., Muntz, R. and Kok, J. 2004. Frequent Subtree Mining An Overview, *Fundamenta Informaticae, Special Issue on Graph and Tree Mining* **66**(1-2), 161–198.
- Chi, Y., Yang, Y., Xia, Y. and Muntz, R.R. 2004a. HybridTreeMiner: An Efficient Algorithm for Mining Frequent Rooted Trees and Trees using Canonical Forms, In *Proceedings of the 16th International Conference on Scientific and Statistical Database Management*, 11–20.
- Chi, Y., Yang, Y., Xia, Y. and Muntz, R.R. 2004b. CMTreeMiner: Mining both Closed and Maximal Frequent subtrees, In *Proceedings of the 8th Pacific Asia Conference on Knowledge Discovery and Data Mining*, 63–73.
- Chi, Y., Yang, Y., Xia, Y. and Muntz, R.R. 2005. Canonical Forms for labelled Trees and Their Applications in Frequent subtree Mining. *Journal of Knowledge and Information Systems* **8**(2), 203–234.
- Christmas, W.J., Kittler, J. and Petrou, M. 1995. Structural Matching in Computer Vision using Probabilistic Relaxation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(8), 749–764.
- Chung, J.C. 1987. $O(n^{2.5})$ Time Algorithm for Subgraph Homeomorphism Problem on Trees, *Journal of Algorithms* **8**, 106–112.
- Conte, D., Foggia, F., Sansone, C. and Vento, M. 2004. Thirty Years of Graph Matching in Pattern Recognition, *International Journal of Pattern Recognition and Artificial Intelligence* **18**(3), 265–298.
- Cook, D.J. and Holder, L.B. 1994. Substructure Discovery Using Minimum Description Length and Background Knowledge, *Journal of Artificial Intelligence Research* **1**, 231–255.
- Cook, D.J. and Holder, L.B. 2000. Graph-based Data Mining, *IEEE Intelligent Systems* **15**(2), 32–41.
- Cordella, L.P., Foggia, P., Sansone, C., Tortorella, F. and Vento, M. 1998. Graph Matching: A Fast Algorithm and its Evaluation, In *Proceedings of the 14th Conference on Pattern Recognition*, 1582–1584.

- Cordella, L.P., Foggia, P., Sansone, C. and Vento, M. 2001. An Improved Algorithm for Matching Large Graphs, In *Proceedings of the 3rd IAPR-TC15 Workshop on Graph-based Representation in Pattern Recognition*, 149–159.
- Cormen, T.H., Leiserson, C.E., Rivest, R.L. and Stein, C. 2001. *Introduction to Algorithms*, 2nd Edition, MIT Press and McGraw-Hill.
- Cui, J.H., Kim, J., Maggiorini, D., Boussetta, K. and Gerla, M. 2005. Aggregated Multicast-a Comparative Study, *Cluster Computing* **8**(1), 15–26.
- Diestel, R. 2000. *Graph Theory*, Springer-Verlag.
- Dehaspe, L., Toivonen, H. and King, R.D. 1998. Finding Frequent Substructures in Chemical Compounds, In *Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining (KDD'98)*, 30–36.
- Deshpande, M., Kuramochi, M., Wale, N., and Karypis, G. 2005. Frequent Sub-Structure-based Approach for Classifying Chemical Compounds, *IEEE Transactions on Knowledge and Data Engineering* **17**(8), 1036–1050.
- Fan, W., Zhang, K., Cheng, H., Gao, J., Yan, X., Han, J., Yu, P.S. and Verscheure, O. 2008. Direct Mining of Discriminative and Essential Frequent Patterns via Model-based Search Tree, In *Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 230–238, Las Vegas, USA.
- Fatta, G.D. and Berthold, M.R. 2005. High Performance Subgraph Mining in Molecular Compounds, In *Proceedings of the 2005 International Conference on High Performance Computing and Communications (HPCC'05)*, 866–877.
- Fischer, I. and Meinel, T. 2004. Graph based Molecular Data Mining - An Overview, In *Proceedings of the 2004 IEEE International Conference on Systems, Man and Cybernetics*, 4578–4582.
- Flake, G., Tarjan, R. and Tsioutsouluklis, K. 2004. Graph Clustering and Minimum Cut Trees, *Internet Mathematics* **1**, 385–408.
- Fortin, S. 1996. *The Graph Isomorphism Problem*, Technical Report, no. TR06-20, The University of Alberta.
- Foggia, P., Genna, R. and Vento, M. 2001. A Performance Comparison of Five Algorithms for Graph Isomorphism, In *Proceedings of the 3rd IAPR-TC15 Workshop on Graph-based Representation in Pattern Recognition*, 188–199.
- Freeman, L. 1979. Centrality in Social Networks Conceptual Clarification, *Social Networks* **1**(3), 215–239.
- Garey, M.R. and Johnson, D.S. 1979. *Computers and Intractability - A Guide to the Theory of NP-Completeness*. W.H. Freeman And Company. New York.
- Gärtner, T., Flach, P. and Wrobel, S. 2003. On Graph Kernels: Hardness Results and Efficient Alternatives, In *Proceedings of the 16th Annual Conference on Learning Theory (COLT'03)*, 129–143.
- Getoor, L. and Diehl, C. 2005. Link Mining: A Survey, *ACM SIGKDD Explorations Newsletter* **7**(2), 3–12.
- Gibbons, A. 1985. *Algorithmic Graph Theory*, Cambridge University Press.
- Greco, G., Guzzo, A., Manco, G., Pontieri, L. and Saccà, D. 2005. Mining Constrained Graphs: the Case of Workflow Systems, *Constraint based Mining and Inductive Databases*, 155–171, LNCS, Springer.
- Gudes, E., Shimony, S.E. and Vanetik, N. 2006. Discovering Frequent Graph Patterns using Disjoint Paths, *IEEE Transaction on Knowledge and Data Engineering* **18**(11), 1441–1456.
- Gutin, G. 2004. 5.3 Independent Sets and Cliques, In *Gross, J.L. and Yellin, J. Handbook of Graph Theory, Discrete Mathematics & Its Applications*, CRC Press, 389–402.
- Han, J., Pei, J. and Yin, Y. 2000. Mining Frequent Patterns without Candidate Generation, In *Proceedings of ACM SIGMOD International Conference on Management of Data*, 1–12.
- Han, J., Cheng, H., Xin, D. and Yan, X. 2007. Frequent Pattern Mining: Current Status and Future Directions, *Journal of Data Mining and Knowledge Discovery* **15**(1), 55–86.
- Han, J. and Kamber, M. 2006. *Data Mining Concepts and Techniques*, 2nd Edition, San Francisco: Morgan Kaufmann.
- Harary, F. and Ross, I.C. 1957. A Procedure for Clique Detection using the Group Matrix, *Sociometry* **20**(3), 205–215.
- Hein, J., Jiang, T., Wang, L. and Zhang, K. 1996. On the Complexity of Comparing evolutionary trees. *Discrete Applied Mathematics* **71**(1-3), 153–169.
- Hido, S. and Kawano, H. 2005. AMIOT: Induced Ordered Tree Mining in Tree-structured Databases, In *Proceedings of the 5th IEEE International Conference on Data Mining*, 170–177.
- Hopcroft, J.E. and Tarjan, R.E. 1972. Isomorphism of Planar Graphs, In *Miller, R.E. and Thatcher, J.W., Complexity of Computer Computations*, 131–152.
- Hu, H., Yan, X., Huang, Y., Han, J. and Zhou, X. 2005. Mining Coherent Dense Subgraphs across Massive Biological Networks for Functional Discovery, *Bioinformatics* **21**(1), 213–221.
- Huan, J., Wang, W. and Prins, J. 2003. Efficient Mining of Frequent Subgraph in the Presence of Isomorphism, In *Proceedings of the 2003 International Conference on Data Mining*, 549–552.

- Huan, J., Wang, W., Prins, J. and Yang, J. 2004a. SPIN: Mining Maximal Frequent Subgraphs from Graph Databases, In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 581–586.
- Huang, X. and Lai, W. 2006. Clustering Graphs for Visualization via Node Similarities, *Visual Language and Computing* **17**, 225–253.
- Inokuchi, A., Washio, T. and Motoda, H. 2000. An Apriori-based Algorithm for Mining Frequent Substructures from Graph Data, In *Proceedings of the 4th European Conference on Principles and Practice of Knowledge Discovery in Databases*, 13–23.
- Inokuchi, A., Washio, T., Nishimura, K. and Motoda, H. 2002. A Fast Algorithm for Mining Frequent Connected Subgraphs, Technical Report RT0448, IBM Research, Tokyo Research Laboratory, Japan.
- Inokuchi, A., Washio, T. and Motoda, H. 2003. Complete Mining of Frequent Patterns from Graphs: Mining Graph Data, *Journal of Machine Learning*, 50(3), 321–354.
- Jahn, K. and Kramer, S. 2005. Optimizing gSpan for Molecular Datasets, In *Proceedings of the 3rd International Workshop on Mining Graphs, Trees and Sequences*, 509–523.
- Kashima, H., Tsuda, K. and Inokuchi, A. 2003. Marginalized Kernels Between labelled Graphs, In *Proceedings of the 20th International Conference on Machine Learning (ICML'03)*, 321–328.
- Kelley, B., Sharan, R., Karp, R., Sittler, E., Root, D., Stockwell, B. and Tdeker, T. 2003. Conserved Pathways within Bacteria and Yeast as Revealed by Global Protein Network Alignment. In *Proceedings of the National Academy of Science of the United States of America (PNAS'03)* **100**(20), 11394–11399.
- Ke, Y., Cheng, J. and Ng, W. 2007. Correlated Search in Graph Databases, In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 390–399.
- Ke, Y., Cheng, J. and Yu, J. 2009. Efficient Discovery of Frequent Correlated Subgraph Pairs. In *Proceedings of the 9th IEEE International Conference on Data Mining*, 239–248.
- Kleinberg, J.M. 1998. Authoritative Sources in a Hyper-linked Environment, In *Proceedings of ACM-SIAM Symposium Discrete Algorithms*, 668–677.
- Kosala, R. and Blockeel, H. 2000. Web Mining Research: A Survey, *ACM SIGKDD Explorations Newsletter* **2**(1), 1–15.
- Koyutürk, M., Grama, A. and Szpankowski, W. 2004. An Efficient Algorithm for Detecting Frequent Subgraphs in Biological Networks, *Journal of Bioinformatics* **20**(1), 200–207.
- Kudo, T., Maeda, E. and Matsumoto, Y. 2004. An Application to Boosting to Graph Classification, In *Proceedings of the 8th Annual Conference on Neural Information Processing Systems*, 729–736.
- Kuramochi, M. and Karypis, G. 2001. Frequent Subgraph Discovery, In *Proceedings of International Conference on Data Mining*, 313–320.
- Kuramochi, M. and Karypis, G. 2002. Discovering Frequent Geometric Subgraphs, In *Proceedings of the IEEE International Conference on Data Mining*, 258–265.
- Kuramochi, M. and Karypis, G. 2004a. An Efficient Algorithm for Discovering Frequent Subgraphs, *IEEE Transactions on Knowledge and Data Engineering* **16**(9), 1038–1051.
- Kuramochi, M. and Karypis, G. 2004b. GREW-A Scalable Frequent Subgraph Discovery Algorithm, In *Proceedings of the 4th IEEE International Conference on Data Mining*, 439–442.
- Kuramochi, M. and Karypis, G. 2004c. Finding Frequent Patterns in a Large Sparse Graph, In *Proceedings of the SIAM International Conference on Data Mining*, 345–356.
- Kuramochi, M. and Karypis, G. 2005. Finding Frequent Patterns in a Large Sparse Graph, *Data Mining and Knowledge Discovery* **11**(3), 243–271.
- Liu, B. 2008. *Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data*, Springer.
- Liu, T.L. and Geiger, D. 1999. Approximate Tree Matching and Shape Similarity, In *Proceedings of 7th International Conference on Computer Vision*, 456–462.
- Matula, D.W. 1978. Subtree Isomorphism in $\mathcal{O}(n^{5/2})$, *Annals of Discrete Mathematics* **2**, 91–106.
- McKay, B.D. 1981. Practical Graph Isomorphism, *Congressus Numerantium* **30**, 45–87.
- Messmer, B.T. and Bunke, H. 1998. A New Algorithm for Error-Tolerant Subgraph Isomorphism Detection, *IEEE Transaction on Pattern Analysis and Machine Intelligence* **20**(5), 493–504.
- Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D. and Alon, U. 2002. Network Motifs: Simple Building Blocks of Complex Networks, *Science* **298**(5594), 824–827.
- Miyazaki, T. 1997. The Complexity of McKay's Canonical labelling Algorithm, *Groups and Computation II, DIMACS Series Discrete Mathematics Theoretical Computer Science*, American Mathematical Society **28**, 239–256.
- Newman, M.E.J. 2004. Detecting Community Structure in Networks, *The European Physical Journal B - Condensed Matter and Complex Systems* **38**(2), 321–330.
- Nijssen, S. and Kok, J.N. 2003. Efficient Discovery of Frequent Unordered Trees, In *Proceedings of the 1st International Workshop on Mining Graphs, Trees and Sequences*, 55–64.
- Nijssen, S. and Kok, J.N. 2004. A Quickstart in Frequent Structure Mining can Make a Difference, In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 647–652.

- Ozaki, T. and Ohkawa, T. 2008. Mining Correlated Subgraphs in Graph Databases, In *Proceedings of the 12th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'08)*, 272–283.
- Paul, S. 1998. *Multicasting on the Internet and Its Applications*, Kluwer Academic Publishers.
- Pei, J., Han, J., Mortazavi-Asl, B., Pinto, H., Chen, Q., Dayal, U. and Hsu, M.C. 2001. PrefixSpan: Mining Sequential Patterns Efficiently by Prefix-Projected Pattern Growth, In *Proceedings of 12th IEEE International Conference on Data Engineering (ICDE 01)*, 215–224, Heidelberg, Germany.
- Pei, J., Jiang, D. and Zhang, A. 2005. On Mining Cross-Graph Quasi-Cliques, In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 228–238, Chicago, USA.
- Preiss, B.R. 1998. *Data Structures and Algorithms with Object-Oriented Design Patterns in C++*, Wiley.
- Prüfer, H. 1918. Neuer Beweis eines Satzes über Permutationen, *Archiv für Mathematik und Physik* **27**, 742–744.
- Raedt, L.D., and Kramer, S. 2001. The Levelwise Version Space Algorithm and its Application to Molecular Fragment Finding. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence* **2**, 853–859.
- Read, R.C. and Corneil, D.G. 1977. The Graph Isomorph Disease, *Journal of Graph Theory* **1**, 339–363.
- Rückert, U. and Kramer, S. 2004. Frequent Free Tree Discovery in Graph Data, In *Proceedings of Special Track on Data Mining, ACM Symposium on Applied Computing*, 564–570.
- Russel, S. and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*, 2nd Edition, Prentice Hall, Upper Saddle River, New Jersey.
- Schmidt, D.C. and Druffel, L.E. 1976. A Fast Backtracking Algorithm to Test Directed Graphs for Isomorphism using Distance Matrices, *Journal of the ACM* **23**(3), 433–445.
- Schreiber, F. and Schwöbbermeyer, H. 2005. Frequency Concepts and Pattern Detection for the analysis of motifs in Networks, *Transactions on Computational Systems Biology* **3**, 89–104.
- Shamir, R. and Tsur, D. 1999. Faster Subtree Isomorphism, *Journal of Algorithms* **33**(2), 267–280.
- Shapiro, L.G., and Haralick, R.M. 1981. Structural Descriptions and Inexact Matching, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **3**, 504–519.
- Shasha, D., Wang, J. and Zhang, S. 2004. Unordered Tree Mining with Applications to Phylogeny, In *Proceedings of the 20th International Conference on Data Engineering (ICDE 04)*, 708–719.
- Shasha, D., Wang, J.T.L. and Giugno, R. 2002. Algorithms and Applications of Tree and Graph Searching, In *Proceedings of the Twenty-first ACM SIGMOD-SIGACT-SIGART Symposium on Principles on Database Systems*, 39–52.
- Sharan, R., Suthram, S., Kelley, R., Kuhn, T., McCuine, S., Uetz, P., Sittler, T., Karp, R. and Ideker, T. 2005. Conserved Patterns of Protein Interaction in Multiple Species, In *Proceedings of the National Academy of Science of the United States of America (PNAS'05)* **102**(6), 1974–1979.
- Tan, H., Dillon, T.S., Hadzic, F., Chang, E. and Feng, L. 2006. IMB3-Miner: Mining Induced/Embedded subtrees by Constraining the Level of Embedding, In *Proceedings of the 8th Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 450–461.
- Tan, H., Dillon, T.S., Feng, L., Chang, E. and Hadzic, F. 2005. X3-Miner: Mining Patterns from XML Database, In *Proceedings of the 6th International Data Mining*, 287–297.
- Tatikonda, S., Parthasarathy, S. and Kurc, T. 2006. Trips and Tides: New Algorithms for Tree Mining, In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, 455–464.
- Termier, A., Rousset, M.C. and Sebag, M. 2002. Treefinder: a First Step Towards XML Data Mining, In *Proceedings of the 2002 IEEE International Conference on Data Mining*, 450–457.
- Thomas, L.T., Valluri, S.R. and Karlapalem, K. 2006. MARGIN: Maximal Frequent Subgraph Mining, In *Proceedings of the 6th International Conference on Data Mining (ICDM 06)*, 1097–1101, Hong Kong.
- Tsuda, K. and Kudo, T. 2006. Clustering Graphs by Weighted Substructure Mining. In *Proceedings of the 23rd International Conference on Machine Learning (ICML'06)*, 953–960.
- Ullmann, J.R. 1976. An Algorithm for Subgraph Isomorphism. *Journal of the ACM* **23**(1), 31–42.
- Valiente, G. 2002. *Algorithms on Trees and Graphs*, Springer.
- Vanetik, N., Gudes, E. and Shimony, S.E. 2002. Computing Frequent Graph Patterns from Semi-structured Data, In *Proceedings of the 2nd International Conference on Data Mining*, 458–465.
- Vanetik, N. 2002. *Discovery of Frequent Patterns in Semi-structured Data*, Department of Computer Science, Ben Gurion University.
- Vanetik, N., Shimony, S.E. and Gudes, E. 2006. Support Measures for Graph Data, *Journal of Data Mining and Knowledge Discovery* **13**(2), 243–260.
- Wang, C., Hong, M., Pei, J., Zhou, H., Wang, W. and Shi, B. 2004a. Efficient Pattern-Growth Methods for Frequent Tree Pattern Mining, In *Proceedings of the 8th Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 441–451.
- Wang, C., Wang, W., Pei, J., Zhu, Y. and Shi, B. 2004b. Scalable Mining of Large Disk-based Graph Databases, In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 316–325.

- Wang, J., Zeng, Z. and Zhou, L. 2006. CLAN: an Algorithm for Mining Closed Cliques from Large Dense Graph Databases, In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–802, Philadelphia, USA.
- Washo, T. and Motoda, H. 2003. State of the Art of Graph-based Data Mining, *SIGKDD Explorations* **5**, 59–68.
- West, D.B. 2000. *Introduction to Graph Theory*, 2nd Edition, Prentice Hall.
- Wörlein, M., Meinel, T., Fischer, I., and Philippsen, M. 2005. A Quantitative Comparison of the Subgraph Miners MoFa, gSpan, FFSM and Gaston, In *Proceedings of the 9th European Conference on Principles and Practice of Knowledge Discovery in Databases*, 392–404, Porto, Portugal.
- Xin, D., Cheng, H., Yan, X. and Han, J. 2006. Extracting Redundancy Aware Top K Patterns, In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 444–453.
- Yan, X. and Han, J.W. 2002. gSpan: Graph-based Substructure pattern mining, In *Proceedings of International Conference on Data Mining*, 721–724.
- Yan, X. and Han, J. 2003. CloseGraph: Mining Closed Frequent Graph Patterns, In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 286–295, Washington D.C., USA.
- Yan, X., Yu, P.S. and Han, J. 2004. Graph Indexing: a Frequent Structure-based Approach, In *Proceedings of ACM-SIGMOD International Conference on Management of Data*, 335–346, Paris, France.
- Yan, X., Zhou, X. and Han, J. 2005a. Mining Closed Relational Graphs with Connectivity Constraints, In *Proceeding of the 11th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 324–333.
- Yan, X., Yu, P.S. and Han, J. 2005b. Sub-Structure Similarity Search in Graph Databases, In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, 766–777.
- Yan, X., Zhu, F., Han, J. and Yu, P.S. 2006. Searching Substructures with Superimposed Distance, In *Proceedings of the 22nd International Conference on Data Engineering*, 88–97.
- Yan, X., Cheng, H., Han, J. and Yu, P.S. 2008. Mining Significant graph patterns by leap search, In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, 433–444, Vancouver, Canada.
- Zaki, M.J. 2002. Efficiently Mining Frequent Trees in a Forest, In *Proceedings of SIGKDD 2002*, 71–80, ACM.
- Zaki, M.J. and Hsiao, C.J. 2002. CHARM: An Efficient Algorithm for Closed Itemset Mining, In *Proceedings of the 2nd SIAM International Conference on Data Mining*, 457–473.
- Zaki, M.J. and Aggarwal, C.C. 2003. XRules: An Effective Structural Classifier for XML Data, In *Proceedings of the 2003 International Conference on Knowledge Discovery and Data Mining*, 316–325.
- Zaki, M.J. 2005a. Efficiently Mining Frequent Embedded Unordered Trees, *Fundamenta Informaticae* **66**(1-2), 33–52.
- Zaki, M.J. 2005b. Efficiently Mining Frequent Trees in a Forest: Algorithms and Applications, *IEEE Transactions on Knowledge and Data Engineering* **17**(8), 1021–1035.
- Zeng, Z., Wang, J., Zhou, L. and Karypis, G. 2006. Coherent Closed Quasi-Clique Discovery from Large Dense Graph Databases, In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–802, Philadelphia, USA.
- Zhang, S. and Wang, J.T.L. 2006. Mining Frequent Agreement Subtrees in Phylogenetic Databases, In *Proceedings of the 6th SIAM International Conference on Data Mining*, 222–233.
- Zhang, S. and Yang, J. 2008. RAM: Randomized Approximate Graph Mining, In *Proceedings of the 20th International Conference on Scientific and Statistical Database Management*, 187–203.
- Zhao, P. and Yu, J. 2006. Fast Frequent Free Tree Mining in Graph Databases, In *Proceedings of the 6th IEEE International Conference on Data Mining Workshop*, 315–319.
- Zhao, P. and Yu, J. 2007. Mining Closed Frequent Free Trees in Graph Databases, In *Proceedings of the 12th International Conference on Database Systems for Advanced Applications*, 91–102, Thailand.
- Zhu, F., Yan, X., Han, J. and Yu, P.S. 2007. gPrune: A Constraint Pushing Framework for Graph Pattern Mining, In *Proceedings of Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 388–400.