

Faster Algorithms for Constructing a Concept (Galois) Lattice

Vicky Choi

Department of Computer Science, Virginia Tech, USA.
vchoi@cs.vt.edu

Abstract. In this paper, we present a fast algorithm for constructing a concept (Galois) lattice of a binary relation, including computing all concepts and their lattice order. We also present two efficient variants of the algorithm, one for computing all concepts only, and one for constructing a frequent closed itemset lattice. The running time of our algorithms depends on the lattice structure and is faster than all other existing algorithms for these problems.

1 Introduction

Formal Concept Analysis (FCA) [14] has found many applications since its introduction. As the size of datasets grows, such as data generated from high-throughput technologies in bioinformatics, there is a need for efficient algorithms for constructing concept lattices. The input of FCA consists of a triple $(\mathcal{O}, \mathcal{M}, \mathcal{I})$, called *context*, where \mathcal{O} is a set of objects, \mathcal{M} is a set of attributes, and \mathcal{I} is a binary relation between \mathcal{O} and \mathcal{M} . In FCA, the context is structured into a set of *concepts*. The set of all concepts, when ordered by set-inclusion, satisfies the properties of a *complete lattice*. The lattice of all concepts is called *concept* [24] or *Galois* [9] lattice. When the binary relation is represented as a bipartite graph, each concept corresponds to a maximal bipartite clique (or maximal biclique). There is also a one-one correspondence of a closed itemset [34] studied in data mining and a concept in FCA. The one-one correspondence of all these terminologies – concepts in FCA, maximal bipartite cliques in theoretical computer science (TCS), and closed itemsets in data mining (DM) – was known, e.g. [3, 34]. There is extensive work of the related problems in these three communities, e.g. [2]–[8] in TCS, [10]–[23] in FCA, and [25]–[36] in DM. In general, in TCS, the research focuses on efficiently enumerating all maximal bipartite cliques (of a bipartite graph); in FCA, one is interested in the lattice structure of all concepts; in DM, one is often interested in computing frequent closed itemsets only.

Time complexity. Given a bipartite graph, it is not difficult to see that there can be exponentially many maximal bipartite cliques. For problems with potentially exponential (in the size of the input) size output, in their seminal paper [6], Johnson et al introduced several notions of *polynomial time* for algorithms for these problems: *polynomial total time*, *incremental polynomial time*, *polynomial delay time*. An algorithm runs in polynomial total time if the time is bounded by a polynomial in the size of the input and the size of the output. An algorithm runs in incremental polynomial time if the time

required to generate a successive output is bounded by the size of input and the size of output generated thus far. An algorithm runs in polynomial delay time if the generation of each output is only polynomial in the size of input. It is not difficult to see that polynomial delay is stronger than incremental polynomial (namely an algorithm with polynomial delay time is also running in incremental polynomial), which is stronger than polynomial total time. polynomial delay algorithm, we can further distinguish if the space used is polynomial or exponential in the input size.

Previous work. Observe that the maximal bipartite clique (MBC) problem is a special case of the maximal clique problem in a general graph. Namely, given a bipartite graph $G = (V_1, V_2, E)$, a maximal bipartite clique corresponds to a maximal clique in $\tilde{G} = (V_1 \cup V_2, \tilde{E})$ where $\tilde{E} = E \cup (V_1 \times V_1) \cup (V_2 \times V_2)$. Consequently, any algorithm for enumerating all maximal cliques in a general graph, e.g., [8, 6], also solves the MBC problem. In fact, the best known algorithm in enumerating all maximal bipartite cliques, which was proposed by Makino and Uno [7] that takes $O(\Delta^2)$ polynomial delay time where Δ is the maximum degree of G , was based on this approach. The fact that the set of maximal bipartite cliques constitutes a lattice was not observed in the paper and thus the property was not utilized for the enumeration algorithm.

In FCA, much of research has been devoted to study the properties of the lattice structure. There are several algorithms, e.g. [19, 23, 18], that construct the lattice, i.e. computing all concepts together with its lattice order. There are also some algorithms that compute only concepts, e.g. [21, 14]. (We remark that the idea of using a total *lectical order* on concepts Ganter’s algorithm [14] is also used in [6, 7] for enumerating maximal (bi)cliques.) See [16] for a comparison studies of these algorithms. The best polynomial total time algorithm was by Nourine and Raynaud [19] with $O(nm|\mathcal{B}|)$ time and $O(n|\mathcal{B}|)$ space, where $n = |\mathcal{O}|$ and $m = |\mathcal{M}|$ and \mathcal{B} denote the set of all concepts. This algorithm can be easily modified to run in $O(mn)$ incremental time [20]. Observe that the space of total size of all concepts is needed if one is to keep the entire structure explicitly. There were several other algorithms, e.g.[14, 18], all run in $O(n^2m)$ polynomial delay. There is another algorithm [23] that is based on divide-and-conquer approach, but the analytical running time of the algorithm is unknown as it is difficult to analyze.

There are several algorithms in data mining for computing frequent closed itemsets, such as CHARM(-L) [35, 36], and CLOSET(+) [29, 32]. To our best knowledge, the algorithm with theoretical analysis running time was given in [3] with $O(m^2n)$ incremental polynomial running time, where $n = |\mathcal{O}|$ and $m = |\mathcal{M}|$.

Our Results. In this paper, by making use of the lattice structure of concepts, we present a simple and fast algorithm for computing all concepts together with its lattice order. The main idea of the algorithm is that given a concept, when all of its successors are considered together (i.e. in a batch manner), they can be efficiently computed. We compute concepts in the Breadth First Search (BFS) order – the ordering given by BFS traversal of the lattice. When computing the concepts in this way, not only do we compute all concepts but also we identify all successors of each concept. Another idea of the algorithm is that we make use of the concepts generated to dynamically update the adjacency relations. The running time of our algorithm is $O(\sum_{a \in \text{ext}(C)} |\text{cnbr}(a)|)$

polynomial delay for each concept C (see Section 2 for related background and terminology), where $\text{cnbr}(a)$ is the reduced adjacency list of a . Our algorithm is faster than the best known algorithms for constructing a lattice because the algorithm is faster than a basic algorithm that runs in $O(\sum_{a \in \text{ext}(C)} |\text{nbr}(a)|)$, where $|\text{nbr}(a)|$ is number of attributes adjacent to the object a , and this basic algorithm is already as fast as the current best algorithms for the problem.

We also present two variants of the algorithm: one is computing all concepts only and another is constructing the frequent closed itemset lattice. Both algorithms are faster than the current start-of-the-art program for these problems.

Outline. The paper is organized as follows. In Section 2, we review some background and notation on FCA. In Section 3, we describe some basic properties of concepts that we use in our lattice-construction algorithm. In Section 4, we first describe the high level idea of our algorithm. Then we describe how to efficiently implement the algorithm. In Section 5, we describe two variants of the algorithm. One is for computing all concepts only and another is for constructing a frequent closed itemset lattice. We conclude with discussion in Section 6.

2 Background and Terminology on FCA

In FCA, a triple $(\mathcal{O}, \mathcal{M}, \mathcal{I})$ is called a *context*, where $\mathcal{O} = \{g_1, g_2, \dots, g_n\}$ is a set of n elements, called *objects*; $\mathcal{M} = \{1, 2, \dots, m\}$ is a set of m elements, called *attributes*; and $\mathcal{I} \subseteq \mathcal{O} \times \mathcal{M}$ is a binary relation. The context is often represented by a *cross-table* as shown in Figure 1. A set $X \subseteq \mathcal{O}$ is called an *object set*, and a set $J \subseteq \mathcal{M}$ is called an *attribute set*. Following the convention, we write an object set $\{a, c, e\}$ as *ace*, and an attribute set $\{1, 3, 4\}$ as *134*.

For $i \in \mathcal{M}$, denote the adjacency list of i by $\text{nbr}(i) = \{g \in \mathcal{O} : (g, i) \in \mathcal{I}\}$. Similarly, for $g \in \mathcal{O}$, denote the adjacency list of g by $\text{nbr}(g) = \{i \in \mathcal{M} : (g, i) \in \mathcal{I}\}$.

Definition 1. The function $\text{attr} : 2^{\mathcal{O}} \rightarrow 2^{\mathcal{M}}$ maps a set of objects to their common attributes: $\text{attr}(X) = \bigcap_{g \in X} \text{nbr}(g)$, for $X \subseteq \mathcal{O}$. The function $\text{obj} : 2^{\mathcal{M}} \rightarrow 2^{\mathcal{O}}$ maps a set of attributes to their common objects: $\text{obj}(J) = \bigcap_{j \in J} \text{nbr}(j)$, for $J \subseteq \mathcal{M}$.

It is easy to check that for $X \subseteq \mathcal{O}$, $X \subseteq \text{obj}(\text{attr}(X))$, and for $J \subseteq \mathcal{M}$, $J \subseteq \text{attr}(\text{obj}(J))$.

Definition 2. An object set $X \subseteq \mathcal{O}$ is closed if $X = \text{obj}(\text{attr}(X))$. An attribute set $J \subseteq \mathcal{M}$ is closed if $J = \text{attr}(\text{obj}(J))$.

The composition of obj and attr induces a *Galois connection* between $2^{\mathcal{O}}$ and $2^{\mathcal{M}}$. Readers are referred to [14] for properties of the Galois connection.

Definition 3. A pair $C = (A, B)$, with $A \subseteq \mathcal{O}$ and $B \subseteq \mathcal{M}$, is called a *concept* if $A = \text{attr}(B)$ and $B = \text{obj}(A)$.

For a concept $C = (A, B)$, by definition, both A and B are closed. The object set A is called the *extent* of C , written as $A = \text{ext}(C)$, and the attribute set B is called the

intent of C , and written as $B = \text{int}(C)$. The set of all concepts of the context $(\mathcal{O}, \mathcal{M}, \mathcal{I})$ is denoted by $\mathcal{B}(\mathcal{O}, \mathcal{M}, \mathcal{I})$ or simply \mathcal{B} when the context is understood.

Let (A_1, B_1) and (A_2, B_2) be two concepts in \mathcal{B} . Observe that if $A_1 \subseteq A_2$, then $B_2 \subseteq B_1$. We order the concepts in \mathcal{B} by the following relation \prec :

$$(A_1, B_1) \prec (A_2, B_2) \iff A_1 \subseteq A_2 (B_2 \subseteq B_1).$$

It is not difficult to see that the relation \prec is a partial order on \mathcal{B} . In fact, $\mathcal{L} = \langle \mathcal{B}, \prec \rangle$ is a complete lattice and it is known as the *concept* or *Galois* lattice of the context $(\mathcal{O}, \mathcal{M}, \mathcal{I})$. For $C, D \in \mathcal{B}$ with $C \prec D$, if for all $E \in \mathcal{B}$ such that $C \prec E \prec D$ implies that $E = C$ or $E = D$, then C is called the *successor*¹ (or *lower neighbor*) of D , and D is called the *predecessor* (or *upper neighbor*) of C . The diagram representing an ordered set (where only successors/predecessors are connected by edges) is called a *Hasse diagram* (or a line diagram). See Figure 1 for an example of the line diagram of a Galois lattice.

For a concept $C = (\text{ext}(C), \text{int}(C))$, $\text{ext}(C) = \text{obj}(\text{int}(C))$ and $\text{int}(C) = \text{attr}(\text{ext}(C))$. Thus, C is uniquely determined by either its extent, $\text{ext}(C)$, or by its intent, $\text{int}(C)$. We denote the concepts restricted to the objects \mathcal{O} by $\mathcal{B}_{\mathcal{O}} = \{\text{ext}(C) : C \in \mathcal{B}\}$, and the attributes \mathcal{M} by $\mathcal{B}_{\mathcal{M}} = \{\text{int}(C) : C \in \mathcal{B}\}$. For $A \in \mathcal{B}_{\mathcal{O}}$, the corresponding concept is $(A, \text{attr}(A))$. For $J \in \mathcal{B}_{\mathcal{M}}$, the corresponding concept is $(\text{obj}(J), J)$. The order \prec is completely determined by the inclusion order on $2^{\mathcal{O}}$ or equivalently by the reverse inclusion order on $2^{\mathcal{M}}$. That is, $\mathcal{L} = \langle \mathcal{B}, \prec \rangle$ and $\mathcal{L}_{\mathcal{M}} = \langle \mathcal{B}_{\mathcal{M}}, \supseteq \rangle$ are order-isomorphic. We have the property that $(\text{obj}(Z), Z)$ is a successor of $(\text{obj}(X), X)$ in \mathcal{L} if and only if Z is a successor of X in $\mathcal{L}_{\mathcal{M}}$. Since the set of all concepts is finite, the lattice order relation is completely determined by the covering (successor/predecessor) relation. Thus, to construct the lattice, it is sufficient to compute all concepts and identify all successors of each concept.

3 Basic Properties

In this section, we describe some basic properties of the concepts on which our lattice construction algorithms are based.

Proposition 1. *Let C be a concept in $\mathcal{B}(\mathcal{O}, \mathcal{M}, \mathcal{I})$. For $i \in \mathcal{M} \setminus \text{int}(C)$, if $E_i = \text{ext}(C) \cap \text{nbr}(i)$ is not empty, E_i is closed. Consequently, $(E_i, \text{attr}(E_i))$ is a concept.*

Proof. For $i \in \mathcal{M} \setminus \text{int}(C)$, suppose that $E_i = \text{ext}(C) \cap \text{nbr}(i)$ is not empty. We will show that $\text{obj}(\text{attr}(E_i)) = E_i$. Since $E_i \subseteq \text{obj}(\text{attr}(E_i))$, it remains to show that $\text{obj}(\text{attr}(E_i)) \subseteq E_i$. By definition, $\text{obj}(\text{int}(C) \cup \{i\}) = (\bigcap_{j \in \text{int}(C)} \text{nbr}(j)) \cap \text{nbr}(i) = \text{ext}(C) \cap \text{nbr}(i) = E_i$. Thus, $(\text{int}(C) \cup \{i\}) \subseteq \text{attr}(\text{obj}(\text{int}(C) \cup \{i\})) = \text{attr}(E_i)$. Consequently, $\text{obj}(\text{attr}(E_i)) \subseteq \text{obj}(\text{int}(C) \cup \{i\}) = E_i$. \square

Example. Consider the concept $C = (abcd, \emptyset)$ of context in Figure 1, we have $E_1 = abc$, $E_2 = bd$, $E_3 = ac$, $E_4 = bd$.

¹ Some authors called this as immediate successor.




Fig. 1. (a) A context $(\mathcal{O}, \mathcal{M}, \mathcal{I})$ with $\mathcal{O} = \{a, b, c, d\}$ and $\mathcal{M} = \{1, 2, 3, 4\}$. The cross \times indicates a pair in the relation \mathcal{I} . (b) The corresponding Galois/concept lattice. (c) $\text{Child}(abcd, \emptyset) = \{(abc, 1), (bd, 24), (ac, 3)\}$; $\text{Child}(abc, 1) = \{(ac, 13), (b, 124)\}$.

3.1 Defining the equivalence classes

For a closed attribute set $X \subset \mathcal{M}$, denote the set of remaining attributes $\{i \in \mathcal{M} \setminus X : \text{obj}(X) \cap \text{nbr}(i) \neq \emptyset\}$ by $\text{res}(X)$. Consider the following equivalence relation \sim on $\text{res}(X)$: $i \sim j \iff \text{obj}(X) \cap \text{nbr}(i) = \text{obj}(X) \cap \text{nbr}(j)$, for $i \neq j \in \text{res}(X)$.

Let S_1, \dots, S_t be the equivalence classes induced by \sim , i.e. $\text{res}(X) = S_1 \cup \dots \cup S_t$, and $\text{obj}(X) \cap \text{nbr}(i) = \text{obj}(X) \cap \text{nbr}(j)$ for any $i \neq j \in S_k$, $1 \leq k \leq t$. We denote the set $\{S_1, \dots, S_t\}$ by $\text{AttrChild}(X)$. We call S_j the sibling of S_i for $j \neq i$. For convenience, we will write $X \cup S_i$ by XS_i . When there is no confusion, we abuse the notation by writing $X \cup \text{AttrChild}(X) = \{XS : S \in \text{AttrChild}(X)\}$. Note that by definition, $\text{obj}(XS_k) = \text{obj}(X) \cap \text{obj}(S_k) = \text{obj}(X) \cap \text{nbr}(i)$ for some $i \in S_k$. We denote the pairs $\{(\text{obj}(XS_1), XS_1), \dots, (\text{obj}(XS_t), XS_t)\}$ by $\text{Child}(\text{obj}(X), X)$.

Recall that $\mathcal{L} = \langle \mathcal{B}, \prec \rangle$ and $\mathcal{L}_{\mathcal{M}} = \langle \mathcal{B}_{\mathcal{M}}, \supseteq \rangle$ are order-isomorphic. We have the property that $(\text{obj}(Y), Y)$ is a successor of $(\text{obj}(X), X)$ in \mathcal{L} if and only if Y is a successor of X in $\mathcal{L}_{\mathcal{M}}$. For each $S \in \text{AttrChild}(X)$, we call XS a child of X and X a parent of XS . By the definition of the equivalence class, for each Z that is a successor of X , there exists a $S \in \text{AttrChild}(X)$ such that $Z = XS$. That is, if Z is a successor of X , Z is a child of X .

Let $\text{Succ}(X)$ denote all the successors of X , then we have $\text{Succ}(X) \subseteq X \cup \text{AttrChild}(X)$. However, not every child of X is a successor of X . For the example in Figure 1, $\text{AttrChild}(\emptyset) = \{1, 24, 3\}$, where 1 and 24 are successors of \emptyset but 3 is not. $\text{Succ}(\emptyset) = \{1, 24\} \subset \text{AttrChild}(\emptyset)$; while $\text{AttrChild}(1) = \{24, 3\}$, $\text{Succ}(a) = \{124, 13\} = 1 \cup \text{AttrChild}(1)$. Similarly, if P is a predecessor of X , then P is parent of X but it is not necessary that every parent of X is a predecessor of X .

Note that for $S \in \text{AttrChild}(X)$, if $XS \in \text{Succ}(X)$, then by definition XS is closed. It is easy to check that the converse is also true. Namely, if XS is closed, then $XS \in \text{Succ}(X)$. In other words, we have the following proposition.

Proposition 2. $\text{Succ}(X) = \{XS : XS \text{ is closed}, S \in \text{AttrChild}(X)\}$.

3.2 Characterizations of Closure

By definition, an attribute set X is closed if $\text{obj}(\text{attr}(X)) = X$. In the following we give two characterizations for an attribute set being closed based on its relationship with its siblings.

Proposition 3. *For $S \in \text{AttrChild}(X)$, XS is not closed if and only if there exists $T \in \text{AttrChild}(X)$, $T \neq S$, such that $\text{obj}(XS) \subset \text{obj}(XT)$. Furthermore, for all $T \in \text{AttrChild}(X)$ with $\text{obj}(XS) \subset \text{obj}(XT)$, there exists $S' \in \text{AttrChild}(XT)$ such that $S \subseteq S'$, $\text{obj}(XS) = \text{obj}(XTS')$ and $XS \subset XTS'$.*

Proof. If XS is not closed, by definition, there exists $i \in \text{res}(X) \setminus S$ such that $i \in \text{attr}(\text{obj}(XS))$. As $\text{AttrChild}(X)$ is a partition of $\text{res}(X)$, there exists a $T \in \text{AttrChild}(X)$ such that $i \in T$, and thus $\text{obj}(XT) = \text{obj}(X) \cap \text{nbr}(i) \supset \text{obj}(XS)$.

Conversely, suppose there exists $T \in \text{AttrChild}(X)$ such that $\text{obj}(XS) \subset \text{obj}(XT)$. Then $\text{attr}(\text{obj}(XS)) \supseteq XTS$. That is, $XS \subset XTS \subseteq \text{attr}(\text{obj}(XS))$, which implies XS is not closed.

Suppose that $\text{obj}(XS) \subset \text{obj}(XT)$ with $T \in \text{AttrChild}(X)$. For $i \in S$, $\text{obj}(XT) \cap \text{nbr}(i) = \text{obj}(XT) \cap \text{obj}(X) \cap \text{nbr}(i) = \text{obj}(XT) \cap \text{obj}(XS) = \text{obj}(XS)$. Thus, there exists $S' \in \text{AttrChild}(XT)$ such that $S \subseteq S'$, $\text{obj}(XS) = \text{obj}(XTS')$. Since X, S, T are disjoint, $XS \subset XTS \subseteq XTS'$. \square

Based on the first part of this proposition (first characterization), we can test if XS is closed, for $S \in \text{AttrChild}(X)$, by using *subset testing* of its object set against its siblings' object set. Namely, XS is closed if and only if $\text{obj}(XS)$ is not a proper subset of its siblings' object set. In our running example in Figure 1, 3 is not closed because its object set $\text{obj}(3) = ac$ is a proper subset of the object set of its sibling, $\text{obj}(1) = abc$.

In general, subset testing operations are expensive. We, however, can make use of the second part of the proposition (second characterization) for testing closure using set exact matching operations instead of subset testing operations. This is because if we process the children in the decreasing order of their object-set size, we can test the closure of XS by comparing its size against the size of the attribute set (if exists) of $\text{obj}(XS)$. Namely, we first search if $\text{obj}(XS)$ exists by a set exact matching operation. If it does not, then XS is closed. Otherwise, if the size of the existing attribute set of $\text{obj}(XS)$ is greater than $|XS|$, then XS is not closed. In our running example, 3 is not closed because $\text{obj}(3) = ac$ has a larger attribute set 13.

4 Algorithm: Constructing a Concept/Galois Lattice

In this section, we first describe the algorithm in general terms, independent of the implementation details. We then show how the algorithm can be implemented efficiently.

4.1 High Level Idea

Recall that constructing a concept lattice includes generating all concepts and identifying each concept's successors.

Our algorithm starts with the top concept $(\mathcal{O}, \text{attr}(\mathcal{O}))$. We process the concept by computing all its *successors*, and then recursively process each successor by either the Depth First Search (DFS) order — the ordering obtained by DFS traversal of the lattice — or Breadth First Search (BFS) order. According to Proposition 2, successors of a concept can be computed from its children. Let $C = (\text{obj}(X), X)$ be a concept. First, we compute all the children $\text{Child}(C) = \{(\text{obj}(XS), XS) : S \in \text{AttrChild}(X)\}$. Then for each $S \in \text{AttrChild}(X)$, we check if XS is closed. If XS is closed, $(\text{obj}(XS), XS)$ is a successor of C . Since a concept can have several predecessors, it can be generated several times. We check its existence to make sure that each concept is processed once and only once. The pseudo-code of the algorithm based on BFS is shown in Algorithm 1.

Algorithm 1 CONCEPT-LATTICE CONSTRUCTION – BFS

```

1: Compute the top concept  $C = (\mathcal{O}, \text{attr}(\mathcal{O}))$ ;
2: Initialize a queue  $Q = \{C\}$ ;
3: Compute  $\text{Child}(C)$ ;
4: while  $Q$  is not empty do
5:    $C = \text{dequeue}(Q)$ ;
   Let  $X = \text{int}(C)$  and suppose  $\text{AttrChild}(X) = \langle S_1, S_2, \dots, S_k \rangle$ ;
6:   for  $i = 1$  to  $k$  do
7:     if  $XS_i$  is closed then
8:       Denote the concept  $(\text{obj}(XS_i), XS_i)$  by  $K$ ;
9:       if  $K$  does not exist then
10:        Compute  $\text{Child}(K)$ ;
11:        Enqueue  $K$  to  $Q$ ;
12:       end if
13:     end if
14:   end for
15: end while

```

4.2 Implementation

The efficiency of the algorithm depends on the efficient implementation of processing a concept that include three procedures: (1) computing $\text{Child}()$; (2) testing if an attribute set is closed; (3) testing if a concept already exists.

First, we describe how to compute $\text{Child}(\text{obj}(X), X)$ in $O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$ time, using a procedure, called SPROUT, described in the following lemma.

Lemma 1. For $(\text{obj}(X), X) \in \mathcal{B}$, it takes $O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$ to compute $\text{Child}(\text{obj}(X), X)$.

Proof. Let $\text{res}(X) = \cup_{a \in \text{obj}(X)} \text{nbr}(a) \setminus X$. For each $i \in \text{res}(X)$, we associate it with a set E_i (which is initialized as an empty set). For each object $a \in \text{obj}(X)$, we scan through each attribute i in its neighbor list $\text{nbr}(a)$, append a to the set E_i . This step

takes $O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$. Next we collect all the sets $\{E_i : i \in \text{res}(X)\}$. We use a trie to group the same object set: search E_i in the trie; if not found, insert E_i into the trie with $\{i\}$ as its attribute set, otherwise we append i to E_i 's existing attribute set. This step takes $O(\sum_{i \in \text{res}(X)} |E_i|) = O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$. Thus, this procedure, called $\text{SPROUT}(\text{obj}(X), X)$, takes $O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$ time to compute $\text{Child}(\text{obj}(X), X)$. \square

For $S \in \text{AttrChild}(X)$, we test if XS is closed based on the second characterization in Proposition 3. For this method to work, it requires processing the children $\text{Child}(\text{obj}(X), X)$ in the decreasing order of their object-set size. Suppose $\text{AttrChild}(X) = \{S_1, \dots, S_k\}$ where $|\text{obj}(XS_1)| \geq |\text{obj}(XS_2)| \geq \dots \geq |\text{obj}(XS_k)|$. We process S_{i-1} before S_i . If XS_{i-1} is closed, we also compute its children $\text{Child}(\text{obj}(XS_{i-1}), XS_{i-1})$. Now to test if XS_i is closed, we check if $\text{obj}(XS_i)$ exists. If it does not, then XS_i is closed. Otherwise, we compare $|XS_i|$ against the size of the existing attribute set of $\text{obj}(XS_i)$. If $|XS_i|$ is not smaller, then XS_i is closed otherwise it is not. To efficiently search $\text{obj}(XS_i)$, we use a trie (with hashing over each node) to store the object sets of concepts generated so far and it takes linear time to search and insert (if not exists) an object set. That is, it will take $O(|\text{obj}(XS_i)|)$ time to check if XS_i is closed. The total time it takes to check if all children are closed is $O(\sum_{i=1}^k |\text{obj}(XS_i)|)$.

Recall that a concept $C = (\text{obj}(X), X)$ is uniquely determined by its extent $\text{obj}(X)$ or its intent X . Therefore, we can store either the object sets or the attribute sets generated so far in a trie, and then test the existence of C by testing the existence of $\text{obj}(X)$ or X . Since searching the object sets are needed in testing the closure of an attribute set as described above, the cost of testing the existence $\text{obj}(X)$ comes for free.

Note that $\sum_{a \in \text{obj}(X)} |\text{nbr}(a)| > \sum_{i=1}^k |\text{obj}(XS_i)| \cdot |S_i|$. Hence, the time it takes to process a concept is dominated by the procedure SPROUT , in $O(\sum_{a \in \text{obj}(X)} |\text{nbr}(a)|)$ time. If we can reduce the sizes of the adjacency lists ($|\text{nbr}(\cdot)|$), we can reduce the running time of the algorithm. Note that this basic algorithm is already as fast as any existing algorithm for constructing a concept lattice (or computing all concepts only that takes $O(\Delta^2)$ time where Δ is the maximum size of adjacency lists).

In the following we describe how to dynamically update the adjacency lists that will reduce the sizes of adjacent lists, and thus improve the running time of the algorithm.

Further Improvement: Dynamically Update Adjacency Lists. Consider a concept $C = (\text{obj}(X), X)$, the object sets of all descendants of C are all subsets of $\text{obj}(X)$. To compute the descendants of C , it suffices to consider the objects with restriction to $\text{obj}(X)$. For $S \in \text{AttrChild}(X)$, by definition, all attributes in S have the same adjacency lists when restricting to $\text{obj}(X)$. That is, for all $i \neq j \in S$, $\text{nbr}(i) \cap \text{obj}(X) = \text{nbr}(j) \cap \text{obj}(X) (= \text{obj}(XS))$. In other words, for all $a \in \text{obj}(X)$, $i \in \text{nbr}(a) \Leftrightarrow j \in \text{nbr}(a)$, for all $i, j \in S$, i.e., the adjacent list of a either contains all elements in S or no element in S . Therefore, we can reduce the sizes of adjacent lists of objects by representing all attributes in S by a single element. For example in Figure example2, we can use a single element 16 to represent the two attributes 1 and 6, and 35 to represent 3 and 5. In doing so, we reduce the size of adjacency list of b from

5 elements $\{1, 3, 4, 5, 6\}$ to three elements $\{16, 35, 4\}$. We call the reduced adjacency lists the condensed adjacency lists. Denoted the condensed adjacent list by $\text{cnbr}()$. The set of condensed adjacency lists corresponds to a reduced cross-table. For example, the reduced cross table of $\text{Child}(abcde, \emptyset)$ of the above example is shown in Figure 2.

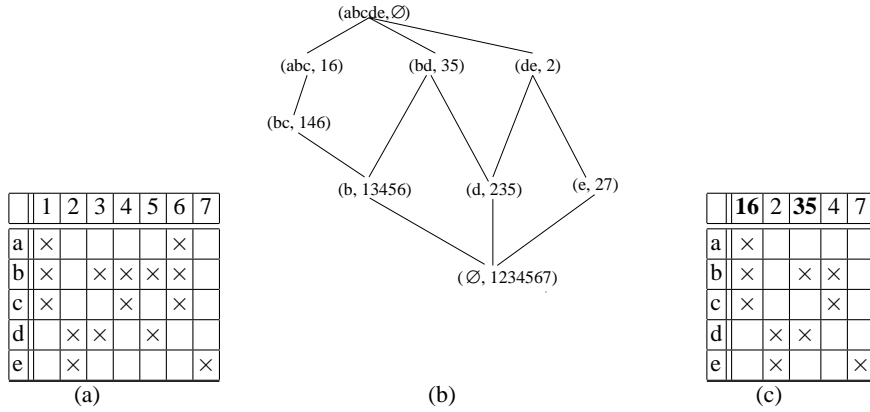


Fig. 2. (a) A context. (b) The corresponding concept lattice. (c) Reduced cross-table of $\text{Child}(abcde, \emptyset)$ of the context.

In order to use the condensed adjacency lists in procedure *SPROUT*, we need to process our concepts in BFS order and it requires one extra level, i.e. in a two-level manner. More specifically, for a concept $C = (\text{obj}(X), X)$, we first compute all its children $\text{Child}(C)$. Then we dynamically update the adjacency lists by representing the attributes in each child of C with one single element. We then use these condensed adjacency lists to process each child of C . That is, instead of using the global adjacency lists, when processing $(\text{obj}(XS), XS)$, we use the condensed adjacency lists of its parent. It takes $O(\sum_{S \in \text{AttrChild}(X)} |\text{obj}(XS)|)$ for C to generate its condensed adjacency lists $\text{cnbr}()$ (see Algorithm 3 in the Appendix for the pseudo-code). And the time for the procedure *SPROUT* is $O(\sum_{a \in \text{obj}(X)} |\text{cnbr}(a)|)$ (see Algorithm 2 in the Appendix for the pseudo-code). Notice that $\sum_{a \in \text{obj}(X)} |\text{cnbr}(a)| > \sum_{S \in \text{AttrChild}(X)} |\text{obj}(XS)|$, the time for updating the adjacency lists is subsumed by the time required for procedure *SPROUT*. Therefore, our new running time is $O(\sum_{a \in \text{obj}(X)} |\text{cnbr}(a)|)$ for each concept $(\text{obj}(X), X)$. See Algorithm 4 for the pseudo-code and Figure 3 for a step-by-step illustration of the algorithm.

5 Variants of The Algorithm

For some applications, one is not interested in the entire concept lattice. In the following, we will describe how to modify our algorithm to solve two special cases: enumerating all concepts only and constructing a frequent closed itemset lattice.

5.1 Algorithm 2: Computing All Concepts or Maximal Bipartite Cliques

If one is interested in computing all the concepts and not in their lattice order, as in enumerating all maximal bicliques studied in [7]. We can easily modify our algorithm to give an even faster algorithm for this purpose. This is because in our algorithm, each concept is generated many times, more precisely, at least number of its predecessors times. For example in Figure 3, $(d, 235)$ is generated twice, one by each of its predecessor. However, when we need all concepts only, we do not need regenerate the concepts again and again. This can be easily accomplished by considering the right siblings only in the procedure SPROUT, i.e. changing the line 3 to for $i \in \text{nbr}(a)$ AND $i > s$ do, while the other parts of the algorithm remain the same. Depending on the lattice structure, this can significantly speed up the algorithm as the number of siblings is decreasing in a cascading fashion. A more careful analysis is needed for the running time of this algorithm.

5.2 Algorithm 3: Constructing a Closed Itemset Lattice

In data mining, one is interested in large concepts, i.e. $(\text{obj}(X), X)$ where $|\text{obj}(X)|$ is larger than a threshold. Although our algorithm can naturally be modified to construct such a closed itemset lattice: we stop processing a concept when the size of its object set is less than the given threshold, where objects correspond to transactions and attributes correspond to items. Theoretically, when the memory requirement is not a concern, our algorithm is faster than all other existing algorithms (including the state-of-art program CHARM-L) for constructing such a frequent closed itemset lattice. However, in practice, for large data sets (as those studied in data mining), the data structure – a trie on objects (transactions) – requires huge memory and this may threaten the algorithm’s practical efficiency. However, it is not difficult to modify our algorithm so that a trie on attributes (items) instead is used. Recall that a trie on objects are required in two steps of our algorithm: testing the closure of an attribute set and testing the existence of a concept. As noted above, the existence of a concept can also be tested on its intent (i.e. attributes), thus we can use a trie on attributes for testing the existence of a concept. To avoid using a trie on objects for testing the closure of an attribute set, we can use the first characterization in Proposition 3 instead, that is, we test the closure of an attribute set by using *subset testing* of its object set against its siblings’ object set, as described in Section 3. Further, we can employ the practically efficient technique *diffset* as in CHARM(-L) for both our SPROUT procedure and subset testing operations. We are testing the performance of the diffset based implementation on the available benchmarks and the results will be reported elsewhere.

6 Discussion

Our interest in FCA stems from our research in microarray data analysis [1]. We have implemented an not yet optimized version of our algorithm (with less than 500 effective lines in C++). The program is very efficient for our applications, in which our data consists of about 10000 objects and 29 attributes. It took less than 1 second for the program to produce the concept lattice (about 530 vertices/concepts and 1500 edges) in a Pentium IV 3.0GHz computer with 2G memory running under Fedora 2 linux OS.

The program is available upon request at this point and will release to the public in the near future.

As FCA finds more and more applications, especially in bioinformatics, efficient algorithms for constructing concept/Galois lattices are much needed. Our algorithm is faster than the existing algorithms for this problem, nevertheless, it seems to have much room to improve.

Acknowledgment

We would like to thank Reinhard Laubenbacher for introducing us FCA. We thank Yang Huang for his participation in this project.

References

1. V. Choi, Y. Huang, V. Lam, D. Potter, R. Laubenbacher, K. Duca. Using Formal Concept Analysis for Microarray Data Comparison. *DIMACS Workshop on Clustering Problems in Biological Networks*. Piscataway, New Jersey, May 9–11, 2006. Manuscript in preparation.
2. J. Abello, A. Pogel, L. Miller. Breadth first search graph partitions and concept lattices. *J. of Universal Computer Science*, 10(8), p934–954, 2004.
3. E. Boros, V. Gurvich, L. Khachiyan, K. Makino. On maximal frequent and minimal infrequent sets in binary matrices. *Annals of Mathematics and Artificial Intelligence*, 39, p211–221, 2003. (STACS 2002)
4. D. Eppstein. Arboricity and bipartite subgraph listing algorithm. *Information Processing Letters*, 51(4), p207–211, 1994.
5. T. Kashiwabara, S. Masuda, K. Nakajima, T. Fujisawa. Generation of maximal independent sets of a bipartite graph and maximum cliques of a circular-arc graph. *J. Algorithms*, 13, p161–174, 1992.
6. D.S. Johnson, M. Yannakakis, C.H. Papadimitriou. On generating all maximal independent sets. *Information Processing Letters*, 27, p119–123, 1988.
7. K. Makino, T. Uno. New algorithms for enumerating all maximal cliques. *Proc. 9th Scand. Workshop on Algorithm Theory (SWAT 2004)*, p260–272. Springer Verlag, Lecture Notes in Computer Science 3111, 2004.
8. S. Tsukiyama, M. Ide, H. Ariyoshi, and I. Shirakawa. A new algorithm for generating all the maximal independent sets *SIAM Journal on Computing*, 6(3), p505–517, 1977.
9. M. Barbut and B. Montjardet. *Ordre et Classifications: Algebre et combinatoire*. Hachette, 1970.
10. F. Baklouti, R.E. Grarvy. A fast and general algorithm for Galois lattices building. *J. of Symbolic Data Analysis*, 3(1), p19–31, 2005. www.icons.rodan.pl/publications/
11. J. Besson, C. Robardet, J-F. Boulicaut. Constraint-based mining of formal concepts in transactional data. *Proceedings of the 8th Pacific-Asia Conference on Knowledge Discovery and Data Mining PaKDD04*, 2004. Springer-Verlag LNCS 3056, pp. 615–624.
12. A Formal Concept Analysis Homepage. <http://www.upriss.org.uk/fca/fca.html>
13. B. Ganter, G. Stumme, R. Wille (eds.). *Formal Concept Analysis: Foundations and Applications*. *Lecture Notes in Computer Science*, vol 3626, Springer, 2005.
14. B. Ganter, R. Wille. *Formal Concept Analysis: Mathematical Foundations*. Springer Verlag, 1996 (Germany version), 1999 (English version).
15. S.O. Kuznetsov. Complexity of learning in context lattices from positive and negative examples. *Discrete Applied Mathematics*, 142, p111–125, 2004.
16. S.O. Kuznetsov and S.A. Obedkov. Comparing performance of algorithms for generating concept lattices. *Journal of Experimental and Theoretical Artificial Intelligence*, 14(23), p189–216, 2002.

17. S.O. Kuznetsov. On computing the size of a lattice and related decision problems. *Order*, 18, p313–321, 2001.
18. C. Lindig. Fast concept analysis. Gerhard Stumme ed. *Working with Conceptual Structures - Contributions to ICCS 2000*, Shaker Verlag, Aachen, Germany, 2000.
19. L. Nourine, O. Raynaud. A fast algorithm for building lattices. *Information Processing Letters*, 71, p199–204, 1999.
20. L. Nourine, O. Raynaud. A fast incremental algorithm for building lattices. *J. of Experimental and Theoretical Artificial Intelligence*, 14, p217–227, 2002.
21. E.M. Norris. An algorithm for computing the maximal rectangles in a binary relation. *Revue Roumaine de mathematiques Pures et Appliquees*, 23(2), p243–250, 1978.
22. P. Valtchev, R. Missaoui, R. Godin, M. Meridji. Generating frequent itemsets incrementally: two novel approaches based on Galois lattice theory. *Journal of Experimental and Theoretical Artificial Intelligence*, 14(2-3), p115–142, 2002.
23. P. Valtchev, R. Missaoui, P. Lebrun. A partition-based approach towards constructing Galois (concept) lattices. *Discrete Mathematics*, 256(3), p801–829, 2002.
24. R. Wille. Restructuring the lattice theory: An approach based on hierarchies of concepts. In I. Rival, editor, *Ordered sets*, pages 445–470, Dordrecht-Boston, 1982, Reidel.
25. F. Afrati, A. Gionis, H. Mannila. Approximating a collection of frequent sets. *Proc. 10th ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, p12–19, 2004.
26. G. Cong, K.L. Tan, A.K.H. Tung, F. Pan. Mining frequent closed patterns in microarray data. *Proc. 4th IEEE International Conference on Data Mining*, p363–366, 2004.
27. T. Calders, C. Rigotti, J.F. Boulicaut. A survey on condensed representations for frequent sets. *Constrained-based Mining*, Springer, 3848, 2005.
28. N. Pasquier, Y. Bastide, R. Taouil, L. Lakhal. Efficient mining of association rules using closed itemset lattices. *Information Systems*, 24(1), p25–46, 1999.
29. J. Pei, J. Han, R. Mao. CLOSET: An efficient algorithm for mining frequent closed itemsets. *Proc. ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, p21–30, 2000.
30. F. Riout, J.F. Boulicaut, B. Cremilleux, J. Besson. Using transposition for pattern discovery from microarray data. *Proc. 8th ACM SIGMOD workshop on Research issues in Data Mining and Knowledge Discovery*, p73–79, 2003.
31. G. Stumme, R. Taouil, Y. Bastide, N. Pasquier, L. Lakhal. Fast Computation of Concept Lattices using data mining techniques. *Proc. 7th International Workshop on Knowledge Representation meets Databases*, p129–139, 2000.
32. J. Wang, J. Han, J. Pei. CLOSET+: Searching for the best strategies for mining frequent closed itemsets. *Proc. 9th ACM SIGKDD international conference on Knowledge Discovery and Data mining*, p236–245, 2003.
33. G. Yang. The complexity of mining maximal frequent itemsets and maximal frequent patterns. *Proc. 10th ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, p344–353, 2004.
34. M.J. Zaki, M. Ogihara. Theoretical foundations of association rules. *Proc. 3rd ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, p1–7, 1998.
35. M.J. Zaki, C.-J. Hsiao. CHARM: An efficient algorithm for closed association rule mining. *Proc. 2nd SIAM International Conference on Data Mining*, 2002.
36. M.J. Zaki, C.-J. Hsiao. Efficient Algorithms for mining closed itemsets and their lattice structure. *IEEE Trans. Knowledge and Data Engineering*, 17(4), p462–478, 2005.

Appendix

Algorithm 2 SPROUT

Input: s , content and nbr
 $(\text{obj}(X), X)$ is the s th child of G . Let $K = \{1, \dots, k\}$ be all the children of G .
Output: $\text{Child}(\text{obj}(X), X) = \{(\text{obj}(XS_i), XS_i) : 1 \leq i \leq t\}$

- 1: For each $i \in K$, set $C_i = \emptyset$.
- 2: **for** $a \in C$ **do**
- 3: **for** $i \in \text{nbr}(a) \setminus \{s\}$ **do**
- 4: Append a to C_i ;
- 5: **end for**
- 6: **end for**

The following takes $O(\sum_{i \in K} |C_i|) = O(\sum_{a \in C} |\text{nbr}(a)|)$ time.

- 7: Initialize a local trie T_C over objects;
- 8: **for** $i \in K$ **do**
- 9: **if** C_i does not exist in T_C **then**
- 10: Insert C_i into T_C ;
- 11: $S_i = \text{content}(i)$;
- 12: **else**
- 13: Merge S_i with $\text{content}(i)$;
- 14: **end if**
- 15: **end for**
- 16: Output all the pairs in T_C : $\{(\text{obj}(XS_j), XS_j) : 1 \leq j \leq t\}$.

Algorithm 3 CONDENSEADJACENTLISTS

Input: $\text{Child}(C) = \{(\text{obj}(XS_i), XS_i) : 1 \leq i \leq t\}$
Output: $\text{content}(i)$ for $i = 1 \dots t$, and new adjacency lists, $\text{nbr}(a)$, $a \in \text{obj}(X)$

- 1: For each $a \in \text{obj}(X)$, $\text{nbr}(a) = \emptyset$;
- 2: **for** $i = 1$ to t **do**
- 3: $\text{content}(i) = S_i$;
- 4: for each $a \in \text{obj}(XS_i)$, append i to $\text{nbr}(a)$;
- 5: **end for**


Algorithm 4 CONCEPT-LATTICE CONSTRUCTION – 2-LEVEL BFS

```


1: Compute the top concept  $C = (\mathcal{O}, \text{attr}(\mathcal{O}))$ ;
2: Initialize a queue  $Q = \{C\}$ ;
3: Initialize a trie  $T$  for the object set  $\mathcal{O}$ ;
4:  $\text{content}(i) = \{i\}$  for  $i \in \mathcal{M}$ ;
5:  $\text{Child}(C) = \text{SPROUT}(0, \text{content}, \text{nbr})$ ;
6: while  $Q$  is not empty do
7:    $C = \text{dequeue}(Q)$ ;
8:   Sort the pairs in  $\text{Child}(C)$  according to its extent size in decreasing order:
     ( $\text{obj}(XS_i), XS_i$ ),  $1 \leq i \leq k$ .
9:    $(\text{content}, \text{nbr}) = \text{CONDENSEADJACENTLISTS}(\text{Child}(C))$ ;
10:  for  $i = 1$  to  $k$  do
11:    Search  $\text{obj}(XS_i)$  in  $T$ ;
     Denote  $(\text{obj}(XS_i), XS_i)$  by  $K$ ; ▷  $K$  is not necessary a concept.
12:    if  $\text{obj}(XS_i)$  does not exist then
13:      Insert  $\text{obj}(XS_i)$  into  $T$ , and associate it with the attribute set  $XS_i$ ;
14:      Identify  $K$  as the successor of  $C$ ;
15:       $\text{Child}(K) = \text{SPROUT}(i, \text{content}, \text{nbr})$ ;
16:      Enqueue  $K$  into  $Q$ ;
17:    else if the attribute set associate with  $\text{obj}(XS_i)$  is not greater than  $XS_i$  then
18:      Identify  $K$  as the successor of  $C$ ;
19:    end if
20:  end for
21: end while

```


(1) Sprout($abcde, \emptyset$) :




(2) Sprout($abc, 16$)




(3) Eliminate $(bc, 4)$ as it is not closed




(4) Sprout($bd, 35$)




(5) Sprout($de, 2$)



(6) Sprout($bc, 146$)



(7) Eliminate $(b, 1356)$ as it is not closed



(8) Sprout($b, 13456$), ($d, 235$), ($e, 27$)




Fig. 3. Step by step illustration of the 2-level BFS lattice construction algorithm. The context and the corresponding lattice are shown in Figure 2.