1-1-2010

Mining Frequent Generalized Patterns for Web Personalization in the presence of Taxonomies

P. Giannikopoulos
University of Peloponnese, Greece

I. Varlamis
University of Athens, Greece

Magdalini Eirinaki
San Jose State University, magdalini.eirinaki@sjsu.edu

Follow this and additional works at: https://scholarworks.sjsu.edu/computer_eng_pub

Part of the [Computer Engineering Commons](#)

Recommended Citation

P. Giannikopoulos, I. Varlamis, and Magdalini Eirinaki. "Mining Frequent Generalized Patterns for Web Personalization in the presence of Taxonomies" *International Journal of Data Warehousing and Mining* (2010): 58-76. <https://doi.org/DOL:10.4018/jdwm.2010090804>

This Article is brought to you for free and open access by the Computer Engineering at SJSU ScholarWorks. It has been accepted for inclusion in Faculty Publications by an authorized administrator of SJSU ScholarWorks. For more information, please contact scholarworks@sjsu.edu.

INTERNATIONAL JOURNAL OF DATA WAREHOUSING AND MINING

January-March 2010, Vol. 6, No. 1

Table of Contents

SPECIAL ISSUE: ECAI'08 WORKSHOP ON MINING SOCIAL DATA

GUEST EDITORIAL PREFACE

i ECAI'08 Workshop on Mining Social Data

Alexandros Nanopoulos, University of Hildesheim, Germany

Grigorios Tsoumakas, Aristotle University of Thessaloniki, Greece

RESEARCH ARTICLES

1 Investigating the Properties of a Social Bookmarking and Tagging Network

Ralitsa Angelova, Max Planck Institut für Informatik, Germany

Marek Lipczak, Dalhousie University, Canada

Evangelos Milios, Dalhousie University, Canada

Paweł Pralat, Dalhousie University, Canada

20 The Dynamics of Content Popularity in Social Media

Symeon Papadopoulos, Aristotle University of Thessaloniki, Greece, and Informatics & Telematics Institute, Greece

Athena Vakali, Aristotle University of Thessaloniki, Greece

Ioannis Kompatsiaris, Informatics & Telematics Institute, Greece

**38 Detecting Trends in Social Bookmarking Systems:
A del.icio.us Endeavor**

Robert Wetzker, Technische Universität Berlin, Germany

Carsten Zimmermann, University of San Diego, USA

Christian Bauchhage, University of Bonn and Fraunhofer IAIS, Germany

**58 Mining Frequent Generalized Patterns for Web Personalization in the Presence of
Taxonomies**

Panagiotis Giannikopoulos, University of Peloponnese, Greece

Iraklis Varlamis, Harokopio University of Athens, Greece

Magdalini Eirinaki, San Jose State University, USA

Mining Frequent Generalized Patterns for Web Personalization in the Presence of Taxonomies

Panagiotis Giannikopoulos, University of Peloponnese, Greece

Iraklis Varlamis, Harokopio University of Athens, Greece

Magdalini Eirinaki, San Jose State University, USA

ABSTRACT

The Web is a continuously evolving environment, since its content is updated on a regular basis. As a result, the traditional usage-based approach to generate recommendations that takes as input the navigation paths recorded on the Web page level, is not as effective. Moreover, most of the content available online is either explicitly or implicitly characterized by a set of categories organized in a taxonomy, allowing the page-level navigation patterns to be generalized to a higher, aggregate level. In this direction, the authors present the Frequent Generalized Pattern (FGP) algorithm. FGP takes as input the transaction data and a hierarchy of categories and produces generalized association rules that contain transaction items and/or item categories. The results can be used to generate association rules and subsequently recommendations for the users. The algorithm can be applied to the log files of a typical Web site; however, it can be more helpful in a Web 2.0 application, such as a feed aggregator or a digital library mediator, where content is semantically annotated and the taxonomic nature is more complex, requiring us to extend FGP in a version called FGP+. The authors experimentally evaluate both algorithms using Web log data collected from a newspaper Web site.

Keywords: Apriori, FP-Growth, Frequent Pattern Mining, Generalized Association Rule Mining, GP-Close

INTRODUCTION

The role of recommendations is very important in everyday transactions. When

buying a product, or reading a newspaper article, one would like to have recommendations on related items. To achieve this, recommendation engines first build a predictive model, by discovering itemsets

DOI: 10.4018/JDWM.2010090804

or item sequences with high support among users. Recommendations are subsequently generated by matching new transaction patterns to the predictive model. Most current approaches in Web personalization consider that a Web site consists of a finite number of Web pages and build their predictive models based on this assumption (Mobasher, 2007). The Web, however, is a continuously evolving environment and this assumption does no longer hold. News portals are typical examples of this situation since they update their content on a regular basis. As a result, the traditional usage-based approach that takes as input the navigation paths recorded on the Web page level is not as effective. Since most predictive models are based on frequent itemsets, the more recent a page is, the more difficult it is to become part of the recommendation set; at the same time, such pages are more likely to be of interest for the average user. This problem can be addressed by generalizing the page-level navigation patterns to a higher, aggregate level (Eirinaki et. al. 2003; Mobasher, 2007).

In this work, we present the FGP algorithm, to address the aforementioned problem. The FGP algorithm is in essence the result of the modification and combination of two algorithms that have been proposed in different contexts. The first one, FP-Growth (Han et. al. 2004), is given a database of user transactions that comprise one or more unordered items (itemsets) and a minimum support threshold. The algorithm processes the transaction database and mines the complete set of frequent itemsets (whose frequency surpasses the threshold). FP-Growth considers the support of each item in the set to be equal to one. In this work, we extend the algorithm so that it assigns different weights to every

item in the set depending on its importance in the transaction. We should note that the FP-Growth algorithm does not consider any relation between items in the database. This, however, is not the case in the Web, where items in a Web site are (conceptually) hierarchically organized. This intrinsic characteristic of the Web can be tackled by the second algorithm, GP-Close (Jiang and Tan, 2006; Jiang et. al., 2007). GP-Close considers a hierarchical organization of all items in the transaction database and uses this information to produce generalized patterns. The two algorithms are very efficient and solve many of the problems of pattern mining, such as the costly generation of candidate sets and the over-generalization of rules.

The FGP algorithm works efficiently in the case of Web sites that have a well-defined underlying hierarchy of topics, such as news portals. Many Web 2.0 sites, however, present a more complex underlying structure. For instance, feed aggregators summarize and present content that is collected from multiple sources. In such sites, the content is not necessarily classified into predefined categories (Inform 2007), being described by user-defined tags instead. This collaborative tagging process results into folksonomies (Voss 2007; P. Heymann and H. Garcia-Molina 2006) that differentiate from the traditional top-down taxonomies. The more complex structure of folksonomies, the use of plurals, the synonym polysemy and specificity of tagging raise new issues for the recommendation engines. In this context, we propose an extension of the FGP algorithm, named FGP+ that takes a more composite topic hierarchy as input, and supports multiple category assignments per topic.

In brief, the contributions of our work are outlined in what follows:

- We modify and combine the forces of FP-Growth and GP-Close in one efficient generalized pattern mining algorithm, named FGP which:
 - extends the frequent-pattern tree, the main structure of the FP-Growth algorithm, to include weight information about items, thus producing a weighted FP-Tree (*WFP-Tree*)
 - addresses the problem of continuously updated content by using the WFP-Tree and the taxonomic information related to the Web site's content as input to GP-Close, and generates generalized recommendations
- We present an extension of FGP, named FGP+ that supports an extended taxonomy of topic categories and/or multiple category assignments per item. The extended algorithm directly addresses the special characteristics of social networking applications.
- We experimentally evaluate our approach using data collected from a newspaper Web site.

The paper is organized as follows. First, we provide an overview of the related research in the area of pattern and association rule mining, as well as in the area of personalizing news sites. Furthermore, we briefly describe the fundamentals of the FP-Growth and GP-Close algorithms, and we present the details of the FGP algorithm. Moreover, we motivate the need for an extension of the basic algorithm, and convey our solution, as well as the experimental

results. We conclude with our plans for future work.

RELATED WORK

There exist numerous approaches that address the problem of personalizing a Web site. An extensive overview can be found in (Mobasher, 2007). In this paper, we overview those that are more similar to ours with regards to: a) the personalization of news sites and b) the abstraction of the generated patterns using a hierarchy.

Both research projects (Antonellis et. al., 2006; Banos et. al., 2006; Katakis et. al. 2008, Gabrilovich et. al., 2004) and commercial sites, such as Spotback (<http://spotback.com>) and Topix (<http://www.topix.net>), have attempted to address the need of personalizing the content of a news site according to users' preferences. Most of those approaches, however, are based on the preference information explicitly provided by the users. However, users' interests change from time to time. In the existence of this concept-drift issue (Tsymbal, 2004; Katakis et. al., 2008), either Web users should continuously update their preferences, or the system will eventually fail to present useful, personalized recommendations. We can see that this is a situation analogous to the cold-start problem, which appears when a system should make predictions in the absence of any transaction history. The cold-start problem has been addressed mainly in the context of collaborative filtering systems (Lam et. al., 2008; Schein et. al., 2002), by creating hybrid recommender systems that take into account both the content of the site and the user ratings or profiles. When there is not adequate user-based informa-

tion, similarities between the content can be used to make predictions.

The idea of integrating the content in the recommendation process has also been addressed by generalizing the page-level navigation patterns to a higher, aggregate level, with the aid of a topic hierarchy. In a previous work, we have proposed the mapping of all user sessions to the topics of a hierarchy (Eirinaki et. al., 2003). These generalized sessions were then used as input to the Apriori algorithm (Agrawal and Srikant, 1994), in order to generate category-based recommendations. Oberle et. al. (2003) proposed a similar framework for semantic Web sites, where the content was annotated using an ontology. This framework focused on Web mining instead of personalization tasks. In (Middleton et. al., 2004) an approach focusing on recommending academic research papers was proposed. The authors mapped the user profiles as well as the research papers to ontology terms, and used these data as input to a collaborative filtering recommender.

Considering several shortcomings of collaborative filtering, such as data sparsity and lack of scalability (Mobasher, 2007), we opted for an association rule mining algorithm as the core of the personalization process. Compared to Apriori or its extensions, namely, AprioriTid and AprioriHybrid (Agrawal and Srikant, 1994), the FP-Growth algorithm is more efficient in that it does not generate candidate itemsets, but adopts a pattern-fragment growth method instead. Moreover, we use the topic hierarchy as an inherent component of our algorithm, and adapt the GP-Close mechanism in order to produce generalized recommendations taking as input hierarchical, as well as complex taxonomies. In contrast to existing techniques that recom-

mend either pages or categories, FGP and its extension generate frequent itemsets consisting of both of them. Thus, it supports the generation of recommendations that include a combination of pages and page categories.

FP-GROWTH AND GP-CLOSE

The FP-Growth Algorithm

The details of the FP-Growth algorithm can be found in the related bibliography (Han et. al. 2004). In what follows we present an overview of the algorithm using a running example. This same example is employed in order to demonstrate the differences between FP-Growth and our algorithm, FGP.

In the first step, FP-Growth scans the transaction database, finds all frequent items (minimum support is 3 in our example) and orders them in descending frequency order. In a second database scan, the FP-Tree is constructed. Each transaction is mapped to a path in the FP-Tree. For the items already in the tree, the count of the respective nodes in the path is updated, whereas new nodes are added for the remaining items. For items belonging to more than one frequent itemsets, all their appearances in the tree are linked. An index table containing all frequent items sorted in descending global frequency order, points to the first appearance of each item in the FP-Tree. The FP-Tree resulting from the transaction database of Table 1 is shown in Figure 1.

As proven by Han et al. (2004), the FP-Tree is adequate for mining frequent patterns and can replace the transaction database. In order to compute the support of a k-itemset, FP-Growth scans the tree in

Table 1. A sample transaction database

TID	Itemset	Ordered frequent items (min freq=3)
100	f, a, c, a, d, g, i, a, m, c, p	f, c, a, m, p
200	a, b, c, f, c, l, a, m, o	f, c, a, b, m
300	b, f, h, j, o, f	f, b
400	b, c, k, s, p, c, b	c, b, p
500	a, c, f, c, e, l, f, p, m, n, a	f, c, a, m, p

Figure 1. The steps of constructing an FP-tree

order to find the less frequent items in the set. The items in the path from the root to the item under examination form the *conditional pattern base* of the item and their support equals the support of the item under examination (count adjustment). Table 2 contains the conditional pattern base for the FP-Tree of Figure 1.

The GP-Close Algorithm

The GP-Close algorithm takes as input a transaction database *DB* and a taxonomy *T*, containing all items of *DB*. Using a minimum support threshold, it generates a tree called *closure enumeration tree* (CET) that contains all the frequent generalized itemsets. The children of a node in the

CET expand their parent itemset by adding one item.

The first step of the algorithm is to locate all frequent 1-itemsets and generate all their frequent generalizations by looking up to *T*. After sorting them in a support increasing manner, it gradually expands them to *n*-itemsets, by combining smaller sets and updating support count. Two pruning techniques prevent from exploring unnecessary combinations: the Subtree pruning and the Child-Closure pruning. The details of the algorithm and an explanation of the pruning techniques are available in (Jiang and Tan, 2006).

Table 2. The conditional pattern base and FP-tree

Item	Conditional pattern base	Conditional FP-Tree
p	{fcam:2, cb:1}	{c:3} p
m	{fca:2, fcab:1}	{f:3, c:3, a:3} m
b	{fca:1, f:1, c:1}	{}
a	{fc:3}	{f:3, c:3} a
c	{f:3}	{f:3} c
f	{}	{}

THE FGP ALGORITHM

In order to demonstrate how the FGP algorithm functions we use the running example introduced in the previous principal Section. Consider that all items in the transaction database of Table 1 are articles in a news site and that the taxonomy of topics depicted in Figure 2 exists for this site (numbers correspond to topic ids, and letters to article ids). Without loss of generality we assume that each article belongs to a single topic. We show how the algorithm is extended to handle multiple category assignments in the following main Section.

Pre- Processing: Item Weighting

We should point out that the information we store in the FP-Tree differs from that of

the original implementation. In the original paper (Han et al. 2004), each transaction identifier (TID) stores only one occurrence for each node. However, in the case of Web log files, a user might visit a Web page more than once during a session. Repetitiveness signifies the importance of a page for a specific user, thus the input format is modified to include *<pageID, weight, support>* triplets, instead of merely pageID information.

Although the importance of a page in a session depends on the number of repetitive visits, its importance in the database is related to the number of distinct sessions it appears in. Thus, analogously to term weighting in document collections (i.e. $tf*idf$), we consider the weight of a page in a session to be the number of its appearances

Figure 2. The taxonomy of items

in the session divided by the total number of page hits in the session (page frequency) and the support of a page to be the number of sessions that contain this page (inverse session frequency).

The FGP Algorithm

The FGP algorithm takes as input a transaction database (as in Table 3) and a taxonomy (as in Figure 2) and constructs a set of generalized association rules as follows:

1. Scan the transaction database and construct the WFP-Tree
2. Find frequent 1-itemsets using the WFP-Tree
3. Create frequent generalized 1-itemsets using the hierarchy
 - a. Sort 1-itemsets in increasing support order
 - b. Prune Children: While creating the generalization tree prune 1-item generalizations that have support equal to a frequent 1-itemset already in the tree
4. Combine 1-itemsets to generate the complete generalized itemsets tree
 - a. Prune subtrees: If a n -itemset A can be subsumed by an identified k -itemset B already in the tree with $n \subset k$ and $\text{support}(A) = \text{support}(B)$

then A and its corresponding subtree is pruned.

In what follows, we demonstrate the implementation of the FGP algorithm on the WFP-Tree and put light on the details of support counting, tree generation and pruning, using the running example introduced before.

Construction of the WFP-Tree

In order to construct the *WFP-Tree*, the transaction database is parsed and the support and weight for each individual page in a transaction is calculated. The algorithm then aggregates the weights of the remaining page ids and stores a reference to the header table. The transactions are stored in decreasing weight order. The final result for the database in Table 3 is depicted in Figure 3. The *WFP-Tree* can replace the original transaction database in the remaining steps of the algorithm.

Discovering Frequent 1-Itemsets and their Generalizations

The header table, which accompanies the *WFP-Tree*, contains a reference to every page in the tree. This table, along with the taxonomy, is used as input in order to find *frequent 1-itemsets* and produce

Table 3. The Web log entries grouped by session

TID	Session items (PID, hits)	Total hits/session
100	(a,3), (c,2), (f,1), (d,1), (g,1), (i,1), (m,1), (p,1)	11
200	(a,2), (c,2), (b,1), (f,1), (l,1), (m,1), (o,1)	9
300	(f,2), (b,1), (h,1), (j,1), (o,1)	6
400	(b,2), (c,2), (k,1), (s,1), (p,1)	7
500	(a,2), (f,2), (c,2), (e,1), (l,1), (p,1), (m,1), (n,1)	11

The result of this processing for Table 1 is depicted in Table 3, which is consequently mapped to the *WFP-Tree*.

Figure 3. The weighted FP-tree

the corresponding *frequent generalized 1-itemsets*. These itemsets are, in essence, the frequently visited categories in the database.

Since categories correspond to more than one page, in order to find the total weight for each category (internal node in the taxonomy tree), FGP finds all the corresponding pages (leaf nodes) in the taxonomy tree. It subsequently processes the index file, from bottom to top, in order to locate all the appearances of the leaf nodes in the WFP-Tree and sum their weights.

For computing the support of a topic (i.e. the number of transactions that contain at least one page from this topic), FGP examines all appearances of the corresponding pages in the WFP-Tree. The transactions that contain many pages from the same topic are counted only once in the support of the latter.

For example, the support for category 11 is computed based on pages j, b and a. First, the algorithm aggregates the appearances of j (1), which is lower in the header table, then of b (1+1 + 1-1, due to j) and consequently those of a (3-1 since b has been added). The total support for category 11 is consequently 5, which corresponds to the number of transactions that contain at least one of {j, b, a}. The weight of 11 is 1.42, which is the sum of the weights of j, b and a.

Pruning 1-Itemset Generalizations

In this step all 1-itemsets and their generalizations that do not have high support (e.g. support < 3 in our example) are being pruned.

Furthermore, in order to avoid the combinatorial explosion of GP-Close when

it searches for all frequent n -itemsets, FGP also prunes those frequent 1-item generalizations that have the same support as one of their specializations. For example, the support of category 37, comprising pages f and g, is 4, which is equal to the support of f. As a result, the generalization of 37 is pruned from the final tree and so do all the combinations of 37.

In order to prune the frequent 1-item generalizations the algorithm sorts all frequent 1-itemsets in increasing support order. If a generalization has the same support as one of its specializations, then it is pruned from the closure enumeration tree. The first level of the tree containing the frequent generalized 1-itemsets is shown in Figure 4.

Discovering Frequent K -Itemsets

FGP incrementally combines the frequent 1-itemsets to generate larger sets. After computing their support and weight, the sets that do not meet the minimum support requirements are pruned. The support for the itemset K is computed over the WFP-Tree as follows:

Suppose that L_z is the set of all leaf nodes for item z . If z is a page then $L_z = \{z\}$.

1. construct $LS = \{L_z\} : \forall z \in K$
support $_K = 0$
2. for $L_1 \in LS$, the first set of pages in LS

3. $\forall i \in L_1$ find i_{ALL} : all appearances of i in WFP-Tree
4. $\forall i_x \in i_{ALL}$ if contain(subnodes(i_x), $LS - L_1$)
then support $_K = \text{support}_K + \text{supportlast}$

where the method contain() parses the list of subnodes of i_x until at least a page from all the sets in $(LS - L_1)$ is found, and supportlast is the support of the last page checked. If the end of a subnodes list has been reached without finding a page for every set then supportlast=0.

To provide an example, we calculate the support of $K = \{f, 24\}$. We first construct $LS = \{\{f\}, \{p, l\}\}$. We check all appearances of f and search for either p or l in the subnode lists. The support for K is 1 (the support of left shaded l in Figure 3) + 2 (the support of the leftmost occurrence of p in the WFP-tree) + 0 (the rightmost f does not contain p or l in its node list). A support of 3 is above the minimum threshold in our example, so $\{f, 24\}$ is a frequent 2-itemset. The weight of this itemset is the aggregate of the weights of all WFP-Tree nodes involved in the support counting, which means f and l in the leftmost branch (0.38+0.11) and p and l (0.18+0.09) in the second leftmost sub-branch (which shares f as an ancestor). The total weight for K is 0.76.

Pruning Redundant Subtrees

It is obvious that certain combinations will be pruned due to insufficient support. For example, a scan in the WFP-Tree of Figure

Figure 4. Frequent generalized 1-itemsets

3 gives to $\{m,p\}$ a support of 2, which is below the specified threshold. Thus, $\{m,p\}$ and its subtree are directly pruned. All the 2-itemsets generated from $\{m\}$ are listed in Figure 5.

A second pruning strategy is applied in this step. According to this, when a k -itemset has equal support to a $(k+1)$ -itemset and is a subset of this itemset then it is a subsumed one and can be pruned. For example, the shaded node *a* in Figure 5 is pruned. This strategy further reduces the possible combinations that need to be checked in the next expansion step.

The complete expansion of the first 1-itemset results in pruning most of the n -itemsets created ($n > 1$). Figure 6 illustrates the result of this expansion, where all shaded nodes are pruned. Expansion continues with the remaining 1-itemsets.

When the tree of sets cannot be further expanded, each node in the tree is exported as a *frequent k-itemset*, which can be used to generate recommendations. For example, similarly to association rules-based recommendations, the recommender system can find the k -itemset that is more similar to the current user's navigation, by comparing the $k-1$ items (pages or categories) to the current visit. The system can then recommend the k -th item, if it is a page (e.g. a news article), or the most popular/recent pages belonging to the k -th item, if it is a category.

THE FGP+ ALGORITHM

The FGP algorithm can be used for generating recommendations in sites where the categories are organized in a hierarchy and each page is characterized by a single category. However, in the case of feed aggregators or digital library mediators, the underlying connectivity of categories is more complex (i.e. taxonomies are enhanced with "related category" links and categories have multiple direct ancestors) and pages (or items in general) are assigned to multiple categories. In what follows, we propose FGP+, which is an extension of the FGP algorithm that addresses the aforementioned issues.

Major Differences Between a Newspaper Site and an Aggregator

In the newspaper world, things are quite simple: there is a collection of articles, each one belonging to a single category only, and a hierarchy of categories, strictly defined by the administrator, so that every node has one parent at most. Nevertheless, this is not the case in Web 2.0 sites where both the content and the category "tagging" is user-controlled. Even when the users select from a list of predefined tags to assign to their articles (e.g. blog posts), they may choose more than one per item. Moreover, those tags may belong to inner categories of the taxonomy (e.g. page *i* in Figure 7, which is assigned to categories 2 and 13).

Figure 5. Creating the 2-itemsets for the first 1-itemset

Figure 6. Expansion of the first 1-itemset and subtree pruning

Figure 7. The modified taxonomy of Figure 2, in which it is allowed for a node to have more than one parents

In several cases the tags assigned to items are different but share similar meaning. Such related or synonym categories can be shown in the taxonomy with horizontal relations (e.g. the relation between categories 3 and 17 in Figure 7). This implies that we need a method to map words to their synonyms in the taxonomy.

Principal Extensions to the FGP Algorithm

The modifications in the taxonomy do not affect the original transaction database and the structure of the WFP-Tree. The taxonomy tree, however, has to be enhanced with additional relations allowing each node to have more than one parent. Thus a *taxonomy graph* should be created, which will maintain its directed acyclic nature. Additionally, each category can have a set of synonyms associated with it, in order for

the mining algorithm to address cases where different users describe the same content using different category tags.

The construction of the extended tree structure and especially the computation of support and weight for a composite itemset, comprising both pages and categories, is not a straightforward task. In fact, even in case we consider no “related topic” edges for the taxonomy graph, the latter should be constructed very cautiously so as to take into account multiple parent assignments for the same internal or leaf node.

Moreover, when we further extend our data structure, allowing each node to contain a list of synonyms associated with it, things become really challenging. For example, Figure 8 presents a taxonomy enhanced with synonymous/related categories information. Furthermore, both leaf and inner categories have instances (pages, denoted with letters).

In what follows, we discuss the different situations that need to be handled during the construction of the extended closure enumeration tree and the actions that are taken for the support and weight computation (the relations between nodes are depicted in the taxonomy graph of Figure 8).

- *Compute the support and weight of categories containing items that have a parent/child relationship* (e.g. {18, 182} in Figure 8). In this case we can either, a) ignore the subclass relation in the creation of the closure enumeration tree, i.e. regard the two categories as unrelated, or b) consider that either the parent or the child is supported when an instance of the child category is found (e.g. k). In our implementation, we follow the first alternative and ignore the existence of the child nodes, as far as the parent one is concerned, when computing the support and weight. For example, the support and weight computation for the pair {18, 182} of the closure enumeration tree of Figure 8 occurs independently for the two nodes, taking the sessions containing p and m, or p and k but not those that contain m and k but not p, into account.
- *Handle nodes having descendant/ancestor relationships* (e.g. {1, 111} in Figure 8), which is in essence a generalization of the parent/child relationship. Thus, we follow the same solution as before. For example, we will not take into consideration the items tagged with 111 during the

Figure 8. The taxonomy graph now supports a list of synonyms for each node, whereas the instances (words from actual tags) do not necessarily belong to the same level

support and weight computation of category 1.

- *Handle synonym/related relationships* (e.g. nodes {1} and {2} in Figure 8). The corresponding values for support and weight should be aggregated over all items belonging to synonym categories. For example, the values of support and weight for any of the two categories {1} and {2} should be the aggregate of the corresponding values of instances j and f .
- *Handle synonyms in descendant/ancestor relations* (e.g. {2, 182} in Figure 8). We are currently facing the problem of support and weight computation for nodes, which, despite seeming irrelevant, possess an ancestor/descendant relationship indeed. In order to identify such relations, we must first replace each instance of a synonym by its principal term. For example, node {2} should be replaced by {1} which is easily identified as an ancestor of 182.

In order to deal with the utilization of synonym tags by the users (e.g. in blog posts), we can keep a list of synonyms associated with each principal category (for example in Figure 8, {1} is a principal category, whereas {2} corresponds to a synonymous one). This list should be used in

order to replace all synonyms by their principal categories, prior to applying FGP+. In this way, the articles belonging to the synonymous category will be considered as members of the principal one, too.

EXPERIMENTAL EVALUATION

Performance Testing

In order to evaluate the performance of the FGP and FGP+ algorithms we use the Web log files of a news site (www.reporter.gr) collected over a 31 days' period (during August 2006). The log files were cleaned, preprocessed, and sessionized, based on the assumption that the Web pages viewed by a user within half an hour belong to the same session. The log files were transformed into a transaction database as the one shown in Table 3. Each page in the Web site belongs to a topic and the hierarchy of topics was used as input to our algorithm. Table 4 shows the statistics of our log file set.

When no pruning is used, the time needed for the creation of the Closure Enumeration Tree (CET) is 21.04 seconds and the tree contains 1707 rules on average, against only 17.2 seconds and 281 rules obtained by applying the two pruning techniques (Child-Closure pruning and Subtree pruning) to FGP. This shows that the two

Table 4. Log files processing statistics

Total number of files	31
Avg num of hits per day	8708
Avg num of sessions per day	882
Avg session length (in page hits)	8.5
Avg num of k-item sets per day (FP-Growth)	7
Avg num of generalized k-item sets per day	56

pruning strategies avoid redundancies and accelerate the tree creation.

In order to evaluate FGP+, we use the same set of log files, but this time we take into account the tag information assigned by the site owners. Although in this data set the tag information is centrally controlled by the administrators, it has some useful features that allow us to test the extension of the FGP algorithm, since a) multiple tags are assigned per article, and b) the tags correspond to topics in the aforementioned hierarchy, in which it is allowed for a node to have more than one parents. As a result, we are able to use FGP+ and find frequent generalized itemsets comprising both articles and tags.

Validity of Results of the FGP Algorithm

The output of the FGP algorithm is a set of frequent k -itemsets, each one associated with a weight and a support score. A recommendation engine can use these frequent k -itemsets against Web usage patterns: when a user's pattern matches the $(k-1)$ items in the set, then the k -th item is suggested to the user, as a recommended hyperlink. The recommendation is considered successful if the user clicks on the hyperlink. Furthermore, if this element corresponds to a category, the n most recent articles belonging to it are recommended, thus providing a solution to the cold-start problem (Lam et. al., 2008; Schein et. al., 2002). We could even propose the $(k-r)$ items in the pattern, provided that the user has requested the rest r , where r is a system parameter.

We measure the accuracy of the recommendations generated by FGP as follows: we produce frequent k -itemsets by applying FGP to the log file of a certain day and

evaluate the rules against the Web log file of the following day. We repeat the same process for every pair of consecutive days and find the average values, performing in essence a 30-fold cross-validation. We validate the itemsets produced from a day's logs only against the logs of the subsequent day, since the life of article ids in the logs is short and rules containing solely article ids will have limited support. In our experiments, we do not make use of the support and weight information of itemsets when counting for sessions matching a set of elements.

We define the *session coverage* (SC) of a set of rules (frequent k -itemsets) measured against a set of sessions as the number of sessions that match at least one rule in the set divided by the total number of sessions, as shown in formula (1):

$$SC = \frac{validSessions}{allSessions} \quad (1)$$

The results of our experiments are depicted in Figure 9. The horizontal axis corresponds to the day used for generating the frequent k -itemsets, whereas the vertical axis shows the percentage of sessions that match at least one rule (session coverage). The results in Figure 9 show that the coverage of the generalized itemsets is larger than that of the page-level ones. The average coverage for the generalized itemsets produced by FGP is almost 29% (dashed line in Figure 9), and it lowers into 5.4% when page-level itemsets are only used (thin line in Figure 9).

In a second set of experiments, we count the total number of rules being matched, for those sessions that satisfy at least one rule for both FGP and FP-Growth. The average values per day are shown in Figure

Figure 9. Session coverage (24 hours)

Figure 10. Valid itemsets per session (24 hours)

10 (dashed and thin line for FGP and FP-Growth respectively). We should point out that the number of matching rules is strongly related to the size of the recommendation set, since the more rules that are matched, the more recommendations will be provided to the end-users. As shown in Figure 10, the average number of rules, produced by FGP, that match a session for the complete dataset is approximately 37, while the value for FP-Growth is almost 2.

After the initial evaluation, we proceeded in evaluating our algorithm by utilizing the half-day (i.e. 12 hours) log as a training set (i.e. for rule generation) and the

next half as a test set (i.e. for evaluation). This approach makes sense in continuously updated sites, such as news sites, where new articles are published every few hours and visitors tend to read the most recent of them. The results, shown in Figures 11 and 12, prove this necessity. This is also an indication that in a recommendation engine, the rule database should be constantly refreshed to capture the readers' shift of interest.

More specifically, as depicted in Figure 11, the average coverage for the generalized itemsets produced by FGP increases to almost 32,3% (dashed line in the graph), while when page-level itemsets are only

Figure 11. Session coverage (12 hours) *Figure 12. Valid itemsets per session (12 hours)*

used (thin line) the coverage raises but is still smaller (hardly reaches an average of 10.5%). Similarly, as shown in Figure 12, the average number of rules produced by FGP that match at least one session for the complete dataset almost doubles (raises to 63), whereas the value for FP-Growth remains 2.

Working with Enhanced Hierarchies

In order to test the ability of FGP+ to work with more complex taxonomies, we used the tag information and repeated the same

set of experiments using the 24 and 12 hour log files. The results of FGP+ are illustrated with a thick line in Figures 9 to 12.

Table 5 summarizes the average values for the three techniques. According to these results, the average coverage for the generalized itemsets produced by FGP+ raises to 36.3% when full day logs are used and to 39% when we use half day logs.

Table 5. Summary of average values for the three algorithms

Average values		FP-Growth	FPG	FPG+
24 hours	Coverage (%)	5.48	29.18	36.30
	Rules per session	1.82	37.70	44.02
12 hours	Coverage (%)	10.48	32.29	39.02
	Rules per session	1.92	62.82	73.88

CONCLUSION & FUTURE WORK

In this article, we presented the FGP algorithm, which takes as input a database of transactions consisting of items that are organized in a taxonomy, as well as the taxonomy itself, and produces a set of frequent k-itemsets comprising items and/or categories from the hierarchy. The set consists of all itemsets above a minimum support threshold and their generalizations but omits redundant generalizations. In the current implementation we modified and combined two state-of-the-art algorithms: FP-Growth for frequent itemset creation and GP-Close for itemset generalization and pruning of redundancies. The proposed algorithm, as well as its extension named FGP+, is capable to deal with taxonomies of various levels of complexity, ranging from simple ones (i.e. taxonomy tree of a newspaper site) to more complicated ones (i.e. taxonomy graph of a feed aggregator), allowing each node to have more than one parents. FGP+ also handles multiple category assignment and a list of synonyms for each concept. The performance evaluation of FGP and FGP+ has shown that they produce many useful itemsets, while avoiding redundancies.

An extensive evaluation of FGP+ against more Web log data sets is in our next plans. We also plan to perform a user-based

evaluation by implementing a recommendation engine on top of the FGP algorithm and use it on a news feed aggregator.

REFERENCES

- Agrawal, R., & Srikant, R. (1994). Fast Algorithms for Mining Association Rules in Large Databases. In J. B. Bocca, M. Jarke, & C. Zaniolo (Eds), *Proceedings of the 20th International Conference on Very Large Data Bases*. San Francisco, CA: Morgan Kaufmann Publishers.
- Antonellis, I., Bouras, C., & Pouloupoulos, V. (2006). Personalized News Categorization Through Scalable Text Classification. In *Proceedings of the 8th Asia-Pacific Web Conference- Frontiers of WWW Research and Development, Lecture Notes in Computer Science: Vol. 3841*. New York: Springer-Verlag.
- Banos, E., Katakis, I., Bassiliades, N., Tsoumakas, G., & Vlahavas, I. (2006). *PersoNews: A Personalized News Reader Enhanced by Machine Learning and Semantic Filtering*. Paper presented at the 5th International Conference on Ontologies, DataBases, and Applications of Semantics (ODBASE 2006), Montpellier, France.
- Eirinaki, M., Vazirgiannis, M., & Varlamis, I. (2003). SEWeP: using site semantics and a taxonomy to enhance the Web personalization process. In *Proceedings of the 9th ACM SIG-KDD International Conference on Knowledge Discovery and Data Mining* (pp. 99-108). New York: ACM Press.

- Gabrilovich, E., Dumais, S., & Horvitz, E. (2004). Newsjunkie: providing personalized newsfeeds via analysis of information novelty. In *Proceedings of the 13th International Conference on World Wide Web* (pp. 482-490). New York: ACM Press.
- Han, J., Pei, J., Yin, Y., & Mao, R. (2004). Mining Frequent Patterns without Candidate Generation: A Frequent-Pattern Tree Approach. *Data Mining and Knowledge Discovery*, 8, 53-87. doi:10.1023/B:DAMI.0000005258.31418.83
- Heymann, P., & Garcia-Molina, H. (2006). *Collaborative creation of communal hierarchical taxonomies in social tagging systems* (Preliminary Technical Report). InfoLab, Stanford, 2006. Retrieved September, 26, 2008, from: <http://heymann.stanford.edu/taghierarchy.html>.
- <https://www.cs.tcd.ie/publications/tech-reports/reports.04/TCD-CS-2004-15.pdf>.
- Inform Inc. (2008), *Inform's Essential Technology Platform*. Retrieved September, 26, 2008, from: <http://www.inform.com/contents/pdf/informwhitepaper.pdf>
- Jiang, T., & Tan, A. H. (2006). Mining RDF Metadata for Generalized Association Rules: Knowledge Discovery in the Semantic Web era. In *Proceedings of the 15th International Conference on World Wide Web* (pp. 951-952). New York: ACM Press.
- Jiang, T., & Wang, K. (2007). Mining Generalized Associations of Semantic Relations from Textual Web Content. *IEEE Transactions on Knowledge and Data Engineering*, 19(2), 164-179. doi:10.1109/TKDE.2007.36
- Katakis, I., Tsoumakas, G., Banos, E., Bassiliades, N., & Vlahavas, I. (2008). An adaptive personalized news dissemination system. *Journal of Intelligent Information Systems*. Springer. DOI - 10.1007/s10844-008-0053-8
- Katakis, I., Tsoumakas, G., & Vlahavas, I. (2008). An Ensemble of Classifiers for coping with Recurring Contexts in Data Streams. In M. Ghallab, C. Spyropoulos, N. Fakotakis, & N. Avouris (Eds.), *Proceedings of the 18th European Conference on Artificial Intelligence* (pp. 763-764), Amsterdam: IOS Press.
- Lam, X., Vu, T., Le, T., & Duong, A. (2008). *Addressing cold-start problem in recommendation systems*. In *Proceedings of the 2nd international Conference on Ubiquitous Information Management and Communication* (pp. 208-211). New York: ACM Press.
- Middleton, S., Shadbolt, N., & De Roure, D. (2004). Ontological User Profiling in Recommender Systems. *ACM Transactions on Information Systems*, 22(1), 54-88. doi:10.1145/963770.963773
- Mobasher, B. (2007). Data Mining for Personalization. In P. Brusilovsky, A. Kobsa, & W. Nejdl, (Eds.), *Lecture Notes in Computer Science: Vol. 4321. The Adaptive Web: Methods and Strategies of Web Personalization* (pp. 90-135). Berlin-Heidelberg: Springer.
- Oberle, D., Berendt, B., Hotho, A., & Gonzalez, J. (2003). Conceptual User Tracking. In *Proceedings of the Atlantic Web Intelligence Conference, Lecture Notes in Computer Science: Vol. 2663*. Springer.
- Schein, A., Popescul, A., & Lyle, H. (2002). Methods and Metric for Cold-Start Recommendations. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 253-260). ACM.
- Tsatsaronis, G., Varlamis, I., & Vazirgiannis, M. (2008). Word Sense Disambiguation with Semantic Networks. In *Proceedings of the 11th International Conference on Text, Speech and Dialogue*, Brno, Czech Republic.

Tsymbal, A. (2004). *The problem of concept drift: definitions and related work* (Technical Report). University of Dublin. Retrieved September 26, 2008, from

Voss, J. (2007). *Tagging, Folksonomy & Co - Renaissance of Manual Indexing?* In *Proceedings of the International Symposium of Information Science* (pp. 234–254).

Panagiotis Giannikopoulos is a PhD candidate at the Department of Computer Science and Technology, University of Peloponnese, Greece, after receiving his BSc from the latter. His research interests focus on data mining, frequent-pattern mining (both centralized and distributed), generalized association rule mining, information retrieval, as well as information filtering.

Iraklis Varlamis is a lecturer at the Department of Informatics and Telematics of Harokopio University of Athens. He received his PhD in computer science from Athens University of Economics and Business, Greece. He collaborates with DB-NET (<http://www.db-net.aueb.gr/>) and WIM (<http://wim.aueb.gr>) research groups. His research interests focus in the area of data-mining and especially in the use of semantics in web mining. He has published several articles in international journals and conferences, concerning web document clustering, the use of semantics in web link analysis and web usage mining, word sense disambiguation using thesauruses etc

Magdalini Eirinaki is an assistant professor at the computer engineering department, San Jose State University, California. She received the PhD degree in informatics from Athens University of Economics and Business. Her research interests cover the areas of Web mining and recommender systems and, in particular, on personalization of web sites and search engines, interactive database exploration, and mining of social networks. She has published several papers in refereed journals and international conference proceedings in the above areas.

CALL FOR ARTICLES

International Journal of Data Warehousing and Mining

An official publication of the Information Resources Management Association!

Prospective authors are invited to submit manuscripts for consideration for publication in the *International Journal of Data Warehousing and Mining*. IJDWM publishes original material concerned with all aspects of data warehousing and mining. The journal invites both conceptual and empirical manuscripts of high quality not currently under review by another publication.

MISSION:

The IJDWM aims to publish and disseminate knowledge on an international basis in the areas of data warehousing and data mining. It is published multiple times a year, with the purpose of providing a forum for state-of-the-art developments and research, as well as current innovative activities in data warehousing and mining. In contrast to other journals, this journal focuses on the integration between the fields of data warehousing and data mining, with emphasis on the applicability to real world problems. The journal is targeted at both academic researchers and practicing IT professionals.

COVERAGE/MAJOR TOPICS:

- data models
- data structures
- data mart
- mining databases
- and more! See www.igi-global.com/ijdw

ISSN 1548-3924

eISSN 1548-3932

Published quarterly

**Full submission guidelines
available at:
<http://www.igi-global.com>**

**All submissions should be emailed to:
david.tanar@infotech.monash.edu.au**

For Full Submission Guidelines, please turn to the back of this journal or visit the IGI Global website at www.igi-global.com.

Ideas for Special Theme Issues may be submitted to the Editor-in-Chief.

Please recommend this publication to your librarian. For a convenient easy-to-use library recommendation form, please visit: <http://www.igi-global.com/ijdw> and click on the "Library Recommendation Form" link along the left margin.