

# A Rule-Based Flickr Tag Recommendation System

Luca Cagliero, Alessandro Fiori, and Luigi Grimaudo

**Abstract** Personalized tag recommendation focuses on helping users find desirable keywords (tags) to annotate Web resources based on both user profiles and main resource characteristics. Flickr is a popular online photo service whose resource sharing system significantly relies on annotations. However, recommending tags to a Flickr user who is annotating a photo is a challenging task as the lack of a controlled tag vocabulary makes the annotation history collection very sparse.

This chapter presents a novel rule-based personalized tag recommendation system to suggest additional pertinent tags to partially annotated resources. Rules represent potentially valuable correlations among tag sets. Intuitively, the system should recommend tags highly correlated with the previously annotated tags. Unlike previous rule-based approaches, a WordNet taxonomy is used to drive the rule mining process and discover rules, called generalized rules, that may contain either single tags or their semantically meaningful aggregations. The use of generalized rules in tag recommendation makes the system (1) more robust to data sparsity and (2) able to capture different viewpoints of the analyzed data. Experiments demonstrate the usefulness of generalized rules in recommending additional tags for real photos published on Flickr.

## 1 Introduction

Social networks and online communities allow creating and managing annotations to categorize and index the resources published by the community users. Tags are keywords that provide meaningful descriptors of the Web resources (e.g., bookmarks, photo, academic articles). Their usefulness has been recently demonstrated in a number of research contexts, e.g., Web content indexing [3],

---

L. Cagliero (✉) • A. Fiori • L. Grimaudo  
Politecnico di Torino. Corso Duca degli Abruzzi, 24 10129 Torino, Italy  
e-mail: [luca.cagliero@polito.it](mailto:luca.cagliero@polito.it); [alessandro.fiori@polito.it](mailto:alessandro.fiori@polito.it); [luigi.grimaudo@polito.it](mailto:luigi.grimaudo@polito.it)

multimedia data retrieval [10], and enterprise Web searches [12]. When dealing with multimedia content (e.g., images or videos) for which the amount of available textual content is limited, tags play an even more important role in Web searching and browsing.

The popularity of Web tagging sites (e.g., [Delicious](#) [11], Flickr [14], and [Zoomr](#) [36]) has prompted the need of novel and more effective recommendation systems to support users in resource annotation by suggesting novel and pertinent tags. Recommendations may be either *personalized*, i.e., dependent on the user who is tagging the resource, or *collective* (user-independent). A significant research effort has been devoted to addressing personalized tag recommendation for Flickr photos [15, 25, 29]. However, the lack of a controlled vocabulary from which tags could be selected during the annotation process makes the set of previously assigned annotations very sparse. Thus, results provided by the mostly used information retrieval or data mining techniques may become unreliable. Indeed, the problem of personalized Flickr tag recommendation is a challenging task.

This chapter presents a novel personalized Flickr tag recommendation system that suggests additional tags to partially annotated Flickr photos. To this aim, it discovers strong generalized association rules from the personal and the community-based sets of past annotations and exploits them to support the process of tag recommendation. An association rule [2] is an implication  $A \Rightarrow B$ , where  $A$  and  $B$  are itemsets (sets of items) named, respectively, as rule antecedent and consequent. In the context of tag recommendations, any item is associated with a distinct tag assigned by a user, while each transaction belonging to the source data is associated with a specific annotation made by user to a given photo and is composed of a set of tags. A rule  $A \Rightarrow B$  may be used to recommend one or more tags contained in  $B$  if the given photo has already been annotated with the tags in  $A$ . A WordNet taxonomy is used to define a hierarchy of is-a or is-part-of aggregations built over the tags occurring in the source data. For instance, based on the WordNet taxonomy, a tag (e.g., *Rome*) may be generalized as its corresponding geographical aggregation *Italy*, while *Europe* may be considered a generalization of both of them. Items relative to aggregated values (e.g., *Italy*) are also called generalized items and denote high-level tag generalizations. The generalization level of a node (i.e., a tag or an aggregation) indicates the length of the path on the taxonomy from the node to a leaf. For instance, in the above-mentioned example, *Italy* has generalization level 1 while *Europe* has level 2. This chapter proposes to exploit WordNet taxonomies to drive the rule mining process and discover rules  $A \Rightarrow B$ , called generalized rules [30], in which itemsets  $A$  and  $B$  may include either single tags (items) or their aggregations (generalized items). To make the generalized rule extraction problem tractable in real-life cases, only a subset of all the possible generalized rules is usually extracted. Selected generalized rules  $A \Rightarrow B$  are characterized by the following properties: (a) the observed frequency of occurrence of the itemset  $A \cup B$  in the analyzed data, called support, is above a given threshold and (b) the conditional probability of occurrence  $P(B|A)$  in the source data, called confidence, is above a given threshold. A (generalized) rule that satisfies (a) is said to be *frequent*. Differently, a (generalized) rule that satisfies both (a) and (b) is said to be *strong*. The use of

generalized rules may allow effective coping with sparse data collections as well as provides different viewpoints of the analyzed data, as shown in the following with the help of a toy example.

Consider, as running example, a photo of the Colosseum, the famous Roman amphitheater situated in the center of the city of Rome (Italy). An example of not generalized association rule may be  $Rome \rightarrow Colosseum$ , where *Rome* and *Colosseum* are tag examples. If the user has already annotated the photo with *Rome*, *Colosseum* is an example of subsequent tag to recommend. However, if the collection is very sparse, the rule is likely to be infrequent in the collection of the past annotations, and thus it is not extracted. The use of a taxonomy that generalizes the tag *Rome* as the corresponding state *Italy* may allow the extraction of a generalized rule  $Italy \rightarrow Colosseum$  that suggests the same annotation while considering a higher-level viewpoint the latter tag correlation.

To select the most relevant tags to recommend, two distinct rule sets are generated: (1) a personalized rule set, which includes the generalized rules extracted from the past annotations made by the user to which the recommendation is targeted, and (2) a community-based rule set, which includes the generalized rules mined from the past annotations made by the community. Tags contained in the consequents of the selected rules are ranked based on the confidence value of the corresponding rules. The ranking process is driven by a newly proposed metrics that weighs differently the rules extracted from the personal and community-based collections. Experiments, reported in Sect. 4, show that the best tag recommendation performance were achieved when giving higher importance to the confidence of the rules extracted from the personal annotations than those mined from the community-based annotations.

The effectiveness of the proposed system has been validated on a real photo collection retrieved from Flickr. The use of generalized rules allows improving the performance of the state-of-the-art approaches.

This chapter is organized as follows. Section 2 overviews the most relevant related works concerning tag recommendation and generalization rule mining. Section 3 presents the framework of the proposed recommender system and describes its main blocks. Section 4 assesses the effectiveness of the system in providing personalized tag recommendations, while Sect. 5 draws conclusions and presents future developments of this work.

## 2 Previous Work

A recommender system helps users find desirable products or services by analyzing user interests and behaviors. Overviews of the most recently proposed recommendation systems are given in [1, 21, 26]. In the last years, a relevant effort has been devoted to the development of novel and more effective tag recommendation systems. This chapter specifically addresses the issue of personalized tag recommendation by means of generalized association rules. In the following, we present

and compare the main state-of-the-art works concerning tag recommendation (see Sect. 2.1) and generalized association rule mining (see Sect. 2.2) with the proposed approach.

## 2.1 Tag Recommendation

The popularity of social networks and online communities (e.g., [Delicious](#) [11], [Flickr](#) [14], and [Zoomr](#) [36]) has increased the attention to the problem of recommending Web resource annotations, i.e., the tags. More specifically, tag recommendation is focused on suggesting pertinent tags to users who are annotating a Web resource. The suggestion may be either personalized, i.e., dependent on the user who is annotating the resources or not, i.e., exclusively based on the collective knowledge.

Several approaches have been proposed to address personalized tag recommendation. For instance, content-based filtering methods [9,22] focus on recommending tags that are similar to those that a user annotated in the past (or is annotating in the present). They commonly analyze the characteristics of the recommended tags to generate detailed user profiles. Tag relevance and user similarity are commonly evaluated by exploiting information retrieval or data mining techniques. In [22], tags for [Delicious](#) bookmarks are recommended by evaluating the cosine similarity among tags and by considering both the cases in which prior tag information is available or not. Differently, in [9], the authors present an application for large-scale automatic generation of personalized annotation tags. They propose an algorithm, named P-TAG, that automatically extracts personalized keywords as tags from bookmarked Web page contents to generate personalized Web document recommendations. Tags are selected based on their relevance to the textual content of the target Web page as well as to the documents residing on the surfer's Desktop.

The use of collaborative filtering approaches in personalized tag recommendation has been addressed in [20, 23, 28]. They collect and analyze a large amount of information on user behaviors, activities, or preferences to predict what users will like based on their similarity to other user features. To this aim, they commonly rely on the assumption that similar users share similar tastes. For instance, in [23], the authors address post tag recommendation by combining, similar to [28], a collaborative filtering method with information retrieval techniques for evaluating the similarities between posts, users, and tags. A hybrid approach that combines a collaborative filtering method with a content-based analysis is proposed in [20]. This system tunes its parameters based on the user feedback to better suit the user preferences. Differently, the combined usage of collaborative filtering and graph-based indexing algorithms is addressed in [18, 33]. In particular, in [18], a user-resource-tag (URT) graph is analyzed by means of an ad hoc indexing strategy derived from the popular PageRank algorithm [7], while in [33], singular value decomposition (SVD) methods are applied to reduce the sparsity of the generated graphs. In [15], an interactive approach to Flickr tag recommendation is proposed.

Suggested tags are first selected according to the set of previously assigned tags based on co-occurrence measures. Next, based on the suggestion, the candidate set is narrowed down to make the suggestion more specific. To overcome challenges of co-occurrence and graph-based measures due to the sparsity of the analyzed data, this chapter proposes to exploit associations at different abstraction levels.

A parallel research issue has been devoted to addressing the problem of collective tag recommendation [17, 19, 29]. The most commonly used approaches are based on co-occurrence measures. For instance, authors in [29] propose a Flickr tag recommendation system that analyzes tag co-occurrences in the collective past annotation collection to suggest additional tags to partially annotated resources. Authors in [25] extend the previous approach to the context of personalized recommendation by combining the knowledge coming from different contextual layers (i.e., personal, collective, and group levels). Differently, the system presented in [19] specifically tackles the cold start problem, i.e., the annotation of not previously annotated resources, by using latent dirichlet allocation (LDA). Unlike [19], this chapter specifically addresses, similar to [25, 29], the task of tag recommendation to partially annotated Flickr photos by using generalized rules instead of traditional rules or co-occurrence measures. Authors in [17] reformulate the task of content-based tag recommendation as a (supervised) classification problem. Using page text, anchor text, surrounding hosts, and available tag information as training data, they build a classifier for each tag they want to predict. The main drawback of the proposed approach is that the overall training time may become very high when the cardinality of the considered tags increases. In the same work, the use of association rules in tag recommendation has been also proposed. Unlike [17], this chapter proposes to overcome the limitations of traditional association rules in coping with sparse data collections by aggregating tags at different abstraction levels according to the given generalization hierarchies.

## 2.2 *Generalized Association Rule Mining*

Association rule mining is a widely used exploratory data mining technique, introduced in [2] in the context of market basket analysis, to discover valuable correlations among data. To focus on rules that are relatively strong, i.e., the ones that frequently occur in the source data and hold in most cases, the mining phase is commonly driven by two main rule quality indexes, i.e., the support and the confidence indexes. However, in some cases, this approach is not effective in discovering relevant data recurrences due to the excessive level of detail of the hidden information. Generalized rules have been first introduced in [30] to address rule mining in the presence of taxonomy. By evaluating a taxonomy built over the data items, items are aggregated into higher-level (generalized) concepts. Each generalized itemset is a high-level representation of a set of “lower level” itemsets according to the given taxonomy. The first generalized association rule mining algorithm [30] generates itemsets by considering, for each item, all its parents in

a taxonomy. Hence, candidate frequent itemsets are generated by exhaustively evaluating the taxonomy. To reduce the mining complexity, several optimizations have been proposed (e.g., [4, 16, 24, 31, 32]). Furthermore, the application of generalized itemsets or rules in different application contexts has been recently investigated as well (e.g., network traffic analysis [4, 6], context-aware systems [5, 8]). Unlike any previously mentioned approaches, this chapter proposes to exploit generalized rules to accomplish the personalized tag recommendation task.

### 3 The Rule-Based Recommendation System

This chapter presents a novel personalized tag recommendation system. Given a photo and a set of user-defined tags, the system proposes novel pertinent tags based on both the personal user preferences, i.e., the tags already annotated by the same user, and the community-based knowledge, i.e., the annotations provided by the other users. Its main architectural blocks are shown in Fig. 1. A brief description of each block follows.

**Tag set data representation.** This block aims at making the history collection of the previously assigned tags suitable for the rule mining process. The tag set is tailored to a transactional data format, where each transaction corresponds to an

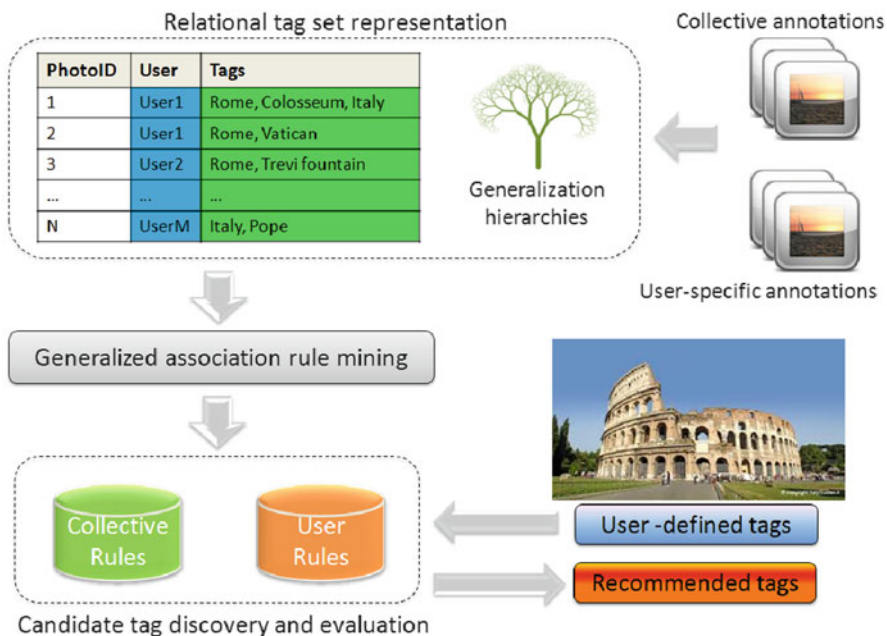


Fig. 1 The recommendation system architecture

annotation performed by a user to a given photo and includes the corresponding set of assigned tags. To allow generalizing tags as the corresponding high-level categories, a set of hierarchies of aggregations (i.e., the generalization hierarchies) built over the analyzed tags is derived from the WordNet [30] lexical database.

**Generalized association rule mining.** This block focuses on discovering high-level correlations, in the form of generalized association rules, from the transactional representation of the past annotation collection. The tag generalization hierarchies are exploited to aggregate tags at higher abstraction levels. Two distinct generalized rule sets are generated: (1) a personalized rule set, which includes the generalized rules extracted from the past annotations made by the user to which the recommendation is targeted and (2) a community-based rule set, which includes the generalized rules mined from the collective knowledge (i.e., the photo annotations of the other users).

**Candidate tag discovery and evaluation.** Given a photo and a set of tags already assigned by the user, this block aims at generating a ranked list of suggested tags. To this aim, the selection of the tags pertinent to the previously annotated ones is driven by the personalized and community-based rules.

This section is organized as follows. Section 3.1 formally states the recommendation task addressed by this chapter, while Sects. 3.2, 3.3, and 3.4 thoroughly describe the main blocks of the recommendation system separately.

### 3.1 Problem Statement

Given a set of photos  $P$ , a set of tags  $T$ , and a set of users  $U$ , the ternary relation  $X = P \times T \times U$  represents the user-specific assignments of tags in  $T$  to photos in  $P$ . We denote as  $\mathcal{T}(p_i, u_j) \subseteq T$  the set of tags assigned by the user  $u_j$  to an arbitrary photo  $u_j$ , where  $u_j \in U$  and  $p_i \in P$ . It could be defined as follows:

$$\mathcal{T}(p_i, u_j) = \pi_t \sigma_{p_i, u_j} X \quad (1)$$

where  $\pi$  and  $\sigma$  are the commonly used projection and selection primitive operators of the relational algebra [13].

To discriminate between past assignments made by user  $u_j$  and those made by the other users, the ternary relation  $X$  may be partitioned as follows:

$$X(u_j) = \pi_t \sigma_{u_j} X \quad (2)$$

$$X(\neg u_j) = \pi_t \sigma_{U \setminus u_j} X \quad (3)$$

Given a set  $\tau(p_i, u_j)$  of user-defined tags and the personal and community-based knowledge  $X(u_j)$  and  $X(\neg u_j)$ , the Flickr personalized tag recommendation task addressed by this chapter focuses on suggesting to user  $u_j$  new tags in  $T \setminus \tau(p_i, u_j)$  for a photo  $p_i$ .



### 3.2 Tag Set Data Representation

Flickr [14] is an online photo-sharing system whose resources are commonly annotated by the system users. The use of association rule mining techniques is particularly suitable for discovering correlations hidden in large real-life tag annotation collections [17]. This chapter investigates the use of an established data mining technique, i.e., generalized association rule mining [30], in recommending tags to partially annotated Flickr photos. Since, in real cases, photo annotations are commonly unsuitable for being directly analyzed by means of data mining algorithms, a preprocessing step is needed.

The collection of past Flickr photo annotations is tailored to a transactional data format. A transactional dataset is a set of transactions, where each transaction is a set of items of arbitrary size. To map the tag set to a transactional data format, the annotations made by a user to a given photo are considered as a transaction composed of the set of (not repeated) assigned tags. More formally, given a ternary relation  $X$  (Cf. Sect. 3.1), any tag set  $\tau(p_i, u_j)$  generated from  $X$  is considered as a transaction. For instance, if the user  $u_j$  assigns to the photo  $p_i$  the tags *Colosseum* and *Rome*, the corresponding transaction becomes  $\{Colosseum, Rome\}$ . The transactional dataset  $D$  is the set of all distinct  $\tau(p_i, u_j)$  occurring in  $X$ , i.e., the full list of the past annotations.

Given a user  $u_j$  to which the personalized tag recommendation is targeted, the transactional dataset  $D$  is partitioned between the annotations made by  $u_j$  and not denoted as  $D(u_j)$  and  $D(\neg u_j)$  are generated. The separate analysis of  $D(u_j)$  and  $D(\neg u_j)$  allows the discovery of both user-specific and collective recurrences, in the form of generalized rules.

To enable the process of generalized rule mining process, a WordNet taxonomy composed of set of hierarchies of aggregations (generalizations) over the tag set  $T$  is built. The WordNet lexical database [30] is queried to retrieve the most relevant semantic relationships holding between a tag in  $T$  and any other term. More specifically, hyponyms (i.e., is-a-subtype-of relationships) and meronyms (is-part-of relationships) are considered. All the terms that belong to these relationships are generalizations of the original tag. For instance, consider again the tag *Rome*. If the following semantic relationship is retrieved from the WordNet database

`<Rome> <is-part-of> <Italy>`

then the term *Italy* is selected as the upper level generalization (aggregation) of the tag *Rome*. Tag generalizations may be further aggregated into high-level categories. For instance, the semantic relationship `< Italy > < is - part - of > < Europe >` prompts the selection of *Europe* as generalization of *Italy* and *Rome*. The generalization level of a tag (or a tag aggregation) is defined as the length of the path on the taxonomy hierarchy from the corresponding node to a leaf. Recalling the previous example, *Europe* has generalization level 2 as the path from *Europe* to *Rome* has length two. Differently, *Rome* has level 0 because it is already a leaf of the taxonomy. Notice that the generalization relationship of two tags (or tag sets)



holds even if they are not characterized by consecutive generalization levels. For instance, Europe is considered one of the possible generalizations of Rome as well. The generalization of an itemset (i.e., a tag set) is defined as the maximum among the levels of its items (tags). For instance, since  $\{Colosseum, Italy\}$  is composed of one item with level 0 and one item of level 1, its generalization level is 1.

### 3.3 Generalized Association Rule Mining

This block focuses on discovering high-level correlations among tags, in the form of generalized association rules, from the transactional representations of the tag sets  $D(u_j)$  and  $D(\neg u_j)$ . Strong association rules represent implications among tags or tag sets that frequently occur and almost hold in the source data [2]. More specifically, an association rule is an implication  $A \rightarrow B$ , where  $A$  and  $B$  are itemsets (i.e., sets of data items). In the transactional representation of the tag set, items are tags in  $T$  associated with any photo included in the collection.

In the context of tag recommendation, generalized association rules [30] are association rules that may include either tags or their high-level aggregations, also denoted as generalized items. By considering the taxonomy built over the tag set (Cf. Sect. 3.2), any concept that aggregates one or more tags in  $T$  at a higher level of generalization is considered as a semantically meaningful tag aggregation. For instance, consider again the semantic relationship  $\langle Rome \rangle \langle is-part-of \rangle \langle Italy \rangle$ . If *Rome* is a tag that occurs in the analyzed data, *Italy* is an example of tag aggregation (generalized item). Similarly, generalized itemsets are itemsets including at most one aggregation (e.g.,  $\{Colosseum, Italy\}$ ). Generalized itemsets are characterized by a notable quality index, i.e., the support, which is defined in terms of the itemset coverage with respect to the analyzed data. A generalized itemset  $I$  covers a given transaction  $d \in D$  if all its (possibly generalized) items  $x \in I$  are either included in  $d$  or ancestors (generalizations) of items  $i \in d$ . Given a transaction dataset  $D$  and a (generalized) itemset  $I$ , the support of  $I$  is given by the ratio between the number of transactions  $d \in D$  covered by  $I$  and the cardinality of  $D$ .

The concept of generalized association rule extends the traditional association rules to the case in which they may include either generalized or not generalized itemsets. A generalized association rule is represented in the form  $A \rightarrow B$ , where  $A$  and  $B$  are two (generalized) itemsets that are named, respectively, as the body and the head of the rule. Similarly,  $A$  and  $B$  are also denoted as rule antecedent and consequent. Generalized association rule extraction is driven by rule support and confidence quality indexes. The support of a generalized rule is defined as the observed frequency of occurrence of  $A \cup B$  in the source dataset. The confidence of a rule  $A \rightarrow B$  is the conditional probability of occurrence of the generalized itemset  $B$  given  $A$  and represents the strength of the implication. For instance, the generalized association rule  $\{Colosseum \rightarrow Italy\}$  characterized by support equal to 10% and confidence equal to 88% states that the tag *Colosseum* co-occurs with

the tag generalization *Italy* in 10% of the transactions (photo annotations) of the collection and the implication holds in 88% of the cases.

The generalized association rule mining task is usually accomplished by means of a two-step process [30]: (1) generalized itemset mining, driven by a minimum support threshold *minsup* and (2) generalized association rule generation, starting from the set of previously extracted itemsets, driven by a minimum confidence threshold *minconf*. A generalized association rule is said to be *strong* if it satisfies both *minsup* and *minconf* [2].

Given a set of generalization hierarchies built over the tags in  $T$ , a minimum support threshold *minsup*, and a minimum confidence threshold *minconf*, the generalized rule mining process is performed on  $D(u_j)$  and  $D(\neg u_j)$  separately. More specifically, given a photo  $p_i$ , a user  $u_j$ , and a set of user-specific tags  $\tau(p_i, u_j)$ , the main idea behind our approach is to treat strong high-level correlations related to annotations made by the user  $u_j$  differently from that made by the other users. To this aim, two distinct rule sets are generated: (1) a personalized rule set  $R_{D(u_j)}$ , which includes the strong generalized rules extracted from the past annotations made by the user to which the recommendation is targeted and (2) a community-based rule set  $R_{D(\neg u_j)}$ , which includes all the strong generalized rules mined from the past annotations made by the other users.

To perform generalized rule mining from the tag history collections, we exploit our more efficient implementation of the Cumulate algorithm [30]. However, different algorithms may be easily integrated as well.

### 3.4 Candidate Tag Discovery and Evaluation

Given a photo  $p_i$ , a set of user-defined tags  $\tau(p_i, u_j)$  already assigned by user  $u_j$ , and the sets  $R_{D(u_j)}$  and  $R_{D(\neg u_j)}$  of generalized rules mined, respectively, from  $D(u_j)$  and  $D(\neg u_j)$ , this block entails the selection and the ranking of the tags to recommend to  $u_j$  for  $p_i$ . In the following, we discuss how to tackle the candidate tag selection and ranking problems separately.

#### 3.4.1 Candidate Tag Selection

The selection step focuses on identifying additional tags, pertinent to the user-specified tag set  $\tau(p_i, u_j)$ , based on the previously generated rule sets. To guarantee the pertinence of the candidate tags, for each photo  $p_i$  only the subset of the personalized and community-based rules including tags in  $\tau(p_i, u_j)$  in their antecedent are considered. More specifically, the candidate tag selection step exclusively considers the strong generalized rules in  $R_{D(u_j)}$  and  $R_{D(\neg u_j)}$  whose (1) rule antecedent exactly covers, at any level of abstraction, the tag set  $\tau(p_i, u_j)$  (or any of its subsets) and (2) rule consequent includes an arbitrary set of not generalized

**Table 1** Generalized rules used for recommending to user  $u_j$  tags subsequent to *Rome*

ID	Generalized rule	Support (%)	Confidence (%)
Annotations made by user $u_j$ ( $R_{D(u_j)}$ )			
1	$\{Rome\} \rightarrow \{Colosseum\}$	25	100
2	$\{Rome\} \rightarrow \{History\}$	17	85
Annotations made by the other users ( $R_{D(\neg u_j)}$ )			
3	$\{Rome\} \rightarrow \{Colosseum, Gladiator\}$	15	95
4	$\{Italy\} \rightarrow \{Roman Age\}$	23	80

items (tags). The coverage of a tag in  $\tau(p_i, u_j)$  may be due to the presence in the rule antecedent of either an exact matching (i.e., the same tag) or one of its higher-level generalizations. Any rule that does not fulfill the above-mentioned constraints is not considered in the subsequent ranking process. The set of tags that occur in the selected rule consequents is chosen as set of candidate recommendable tags.

Consider, for instance, a photo  $p_i$  annotated by user  $u_j$  with the tag *Rome*. In Table 1 is reported the selection of generalized rules, fulfilling the above-mentioned constraints, that has been taken from the set of rules mined from the personalized collection  $D(u_j)$  and community-based one  $D(\neg u_j)$ . In this example, we exploit the generalization hierarchies described in Sect. 3.2, and we enforce, respectively, a minimum support threshold equal to 15 % and a minimum confidence threshold equal to 50 %.

Readers could notice that any selected rule must have (1) as rule antecedent, either the user-specified tag *Rome* or its generalization *Italy*, and (2) as rule consequent, an arbitrary set of (not generalized) tags. Tags occurring in the rule consequents include the potentially relevant tags to recommend. Recalling the previous example, the set  $C$  of candidate tags is  $\{Colosseum, History, Gladiator, Roman Age\}$ . Notice that a single rule may include one or more candidate tags (e.g., *Colosseum* and *Gladiator* co-occur in the consequent of the rule (3)).

The generalization process prevents the discarding of potentially relevant knowledge. In fact, it allows also recommending tags contained in the consequent of rules having as antecedent a generalization of the user-defined tags. For instance, the generalized rule  $Italy \rightarrow Roman Age$  mined from  $R_{D(\neg u_j)}$  suggests to recommend the tag *Roman Age* subsequently to *Rome*. Indeed, even if the rule  $Rome \rightarrow Roman Age$  is infrequent with respect to the minimum support threshold in  $R_{D(\neg u_j)}$  (possibly because of the sparsity of the personal annotation collection), the co-occurrence between *Roman Age* and *Rome* does not remain hidden.

Consider now the case in which the set of user-specified tags  $\tau(p_i, u_j)$  is  $\{Rome, Roman Empire\}$ . Rules including as antecedent the tag set  $\{Rome, Roman Empire\}$  or any of its subsets belonging to any abstraction level (e.g., *Rome, Italy*) are deemed worth considering in the selection of the candidate tags. For instance,  $Italy, Roman Empire \rightarrow Roman Age$  may be considered to recommend the tag *Roman Age* as well.

### 3.4.2 Candidate Tag Ranking

The last but not the least task in tag recommendation is the ranking of the candidate tags in  $C$  to recommend to  $u_j$  for  $p_i$ . The tag ranking should reflect (1) their significance with respect to the user-defined tags, (2) their relevance according to the personal user preferences, and (3) their relevance based on the collective knowledge.

To take the correlation with the previously annotated tags into account, we propose a ranking strategy that evaluates the candidate tags in terms of the interestingness of the rules in  $R_{D(u_j)}$  and  $R_{D(\neg u_j)}$  from which they have been selected. Generalized rule interestingness is evaluated in terms of its confidence index value [2], i.e., the rule strength in the analyzed dataset (see Sect. 3.3). Based on the assumption that personal recommendations frequently assigned by user  $u_j$  might be weighted differently from that made by the other users, we evaluate the contribution of each rule set separately and then we properly combine the resulting scores.

More formally, let  $c \in C$  be an arbitrary candidate tag and  $R_{D(u_j)}^c \subseteq R_{D(u_j)}$  and  $R_{D(\neg u_j)}^c \subseteq R_{D(\neg u_j)}$  be, respectively, the subsets of rules in  $R_{D(u_j)}$  and  $R_{D(\neg u_j)}$  whose antecedent covers  $c$  (at any level of abstraction). The ranking score of  $c$  is defined as follows:

$$\text{rankscore}(c) = \lambda \cdot \frac{\sum_{r_{u_j} \in R_{D(u_j)}^c} \text{conf}(r_{u_j})}{|R_{D(u_j)}^c|} + (1 - \lambda) \cdot \frac{\sum_{r_{\neg u_j} \in R_{D(\neg u_j)}^c} \text{conf}(r_{\neg u_j})}{|R_{D(\neg u_j)}^c|} \quad (4)$$

where  $\lambda \in [0,1]$  is a user-provided algorithm parameter.

Relatively speaking, when  $\lambda > 0.5$  the impact of the confidence value of the rules occurring in  $R_{D(u_j)}^c$  is higher than that in  $R_{D(\neg u_j)}^c$ , i.e., preferences given by user  $u_j$  are deemed more significant than those given by the other users. Oppositely, in case  $\lambda < 0.5$  user  $u_j$  preferences are averagely penalized. An analysis of the impact of  $\lambda$  on the performance of the proposed recommendation system is reported in Sect. 4.


The recommendation system returns the set  $C$  of selected tags sorted by the ranking score reported in Eq. 4.

## 4 Experimental Results

We performed a set of experiments addressing the following issues: (1) a performance comparison between our system and a set of recently proposed approach, (2) the impact of the generalization process on the recommendation process, and (3) the analysis of the recommendation system parameters.

This section is organized as follows. Section 4.1 describes the characteristics of the photo collection exploited in the experimental evaluation. Section 4.2 describes the experimental design and introduces the evaluation metrics adopted for the performance evaluation. Section 4.3 compares the results achieved by our system

**Fig. 2** Portion of an example generalization hierarchy built over the photo collection tags



with that achieved by the system presented in [25] and a baseline version of our approach that does not exploit generalized rules. Finally, Sect. 4.4 analyzes the impact of the system parameters on the recommendation performance.

#### 4.1 Photo Collection

To evaluate the performance of our approach we retrieved, by means of the Flickr APIs, 2,300 real photos, each one annotated with at least 5 tags by a set of 30 users. The selected photos were chosen based on a series of high-level geographical topics, i.e., *New York*, *San Francisco*, *London*, and *Vancouver*. By following the strategy described in Sect. 3.2, a set of generalization hierarchies is derived from the WordNet lexical database over the collected photo tags. A portion of one of the generated generalization hierarchies is reported in Fig. 2.

To evaluate the effectiveness of our system in coping with heterogeneous photo annotations, the considered photo collection is ensured to be unevenly distributed among the analyzed upper level tag categories.

#### 4.2 Experimental Design

Since our system retrieves a ranked list of pertinent additional tags based on the extracted frequent generalized rules, we defined the tag recommendation task as a ranking problem. Given a photo  $p_i$  and a set of user-defined tags  $\tau(p_i, u_j)$ ,

the system has to recommend tags that describe the photo based on both the user-specific and the collective past annotations. To perform personalized recommendation, from the whole photo collection, the user-specific annotations made by 10 users who annotated at least 15 photos are considered separately. The above selection allows making the statistical evaluation of our recommendation system reliable. Once a user-specific annotation subset is selected, the rest of the collection is considered as the collective set. For each analyzed user collection, the evaluation process performs an hold-out train-test validation, i.e., the user-specific collection is partitioned in a training set, including the 75 % of the whole annotations, whereas the remaining part is chosen as test set. To evaluate the additional tag recommendation performance of our system, for each test photo, two random tags are selected as initial (user-specified) tag set and the recommended tag list is compared with the held-out test tags. A recommended tag is judged as correct if it is present in the held-out set. Since the held-out tags need not be the only tags that could be assigned to the photo, the evaluation method actually gives a lower bound on the system performance.

To evaluate the performance of both our recommendation system and its considered competitors, we exploited three standard information retrieval metrics, previously adopted in [25, 29] in the context of additional Flickr tag recommendation. The selected measures are deemed suitable for evaluating the system performance at different aspects. Let  $Q$  be the set of relevant tags, i.e. the tags really assigned by the user to the test photo, and  $C$  the tag set recommended by the system under evaluation. The adopted evaluation measures are defined as follows.

**Mean reciprocal rank (MRR).** This measure captures the ability of the system to return a relevant tag (i.e., a held-out tag) at the top of the ranking. The measure is averaged over all the photos in the testing collection and is computed by:

$$\text{MRR} = \max_{q \in Q} \frac{1}{c_q} \quad (5)$$

where  $c_q$  is the rank achieved by the relevant tag  $q$ .

**Success at rank  $k$  (S@ $k$ ).** This measure evaluates the probability of finding a relevant tag among the top- $k$  recommended tags. It is averaged over all the test photos and is defined as follows:

$$S@k = \begin{cases} 1 & \text{if } q \in C_k, \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where  $q \in Q$  is a relevant tag and  $C_k$  is the set of the top- $k$  recommended tags.

**Precision at rank  $k$  (P@ $k$ ).** This metric evaluates the percentage of relevant tags among the set of retrieved ones. The measure, averaged over all test photos, is defined as follows:

$$P@k = \frac{|Q \cap C_k|}{|Q|} \quad (7)$$

For any evaluated measure, the estimates on each test photo are averaged over ten runs, where, within each run, a different (randomly generated) held-out tag set ranking is considered.

### 4.3 Performance Comparison

The aim of this section is twofold. Firstly, it experimentally demonstrates the effectiveness of our system against a state-of-the-art approach. Secondly, it evaluates the impact of the generalization process on the recommendation performance. To achieve these goals, we compared the performance of our system, in terms of the evaluation metrics described in Sect. 4.2, with (1) a recently proposed personalized Flickr tag recommendation system [25] and (2) a baseline version of our approach, which does not exploit generalized rules.

The system presented in [25] is a personalized recommender system that proposes additional photo tags pertinent to a number of different user contexts, among which the personal and the collective ones. The system generates a list of recommendable tags based on a probabilistic co-occurrence measure for each context. Then, it aggregates the results achieved within each context in a final recommended list by exploiting the Borda count group consensus function [34]. To the best of our knowledge, it is the most recent work proposed on the topic of personalized additional tag recommendation. To perform a fair comparison, we evaluated the performance of our implementation of the approach presented in [25] (denoted as *probabilistic prediction* in the following) when coping with the combination of the collective and the personalized contexts.

To demonstrate the usefulness of generalized rules in tag recommendation, we also compared the performance of our system with that of a baseline version, which exploits traditional (not generalized) association rules [2] solely. More specifically, the baseline method performs the same steps of the proposed approach, while disregarding the use of tag generalizations in discovering significant tag associations (see Sect. 3.4.1).

To test the performance of our approach, we consider as standard configuration the following setting: minimum support threshold  $minsup = 30\%$ , minimum confidence threshold  $minconf = 40\%$ , and  $\lambda = 0.75$ . A more detailed analysis of the impact of these parameters on the recommendation system performance is reported in Sect. 4.4. Even for the baseline version of our system we tested several support and confidence threshold values. For the sake of brevity, in the following we select as representative the configuration that achieved the best results in terms of MMR measure, i.e., minimum support threshold equal to 30% and minimum confidence threshold equal to 40%.



**Table 2** Performance comparison in terms of S@1, P@1, S@5, P@5, and MRR metrics

System	S@1–P@1	S@5	P@5	MRR
<i>GR-TAG</i>	0.7392	0.8695	0.5696	0.7935
<i>Baseline</i>	0.7174*	0.8261*	0.4957*	0.7572*
<i>rule-based</i>				
<i>Probabilistic prediction</i>	0.6956*	0.8478	0.4435*	0.7681*

Statistically relevant worsening in the comparisons between our system and the other approaches is starred

The achieved results are summarized in Table 2. In particular, the success and the precision at ranks 1 and 5 (i.e., S@1, P@1, S@5, and P@5, respectively) as well as the mean reciprocal rank (MRR) achieved by both our system (named *GR-TAG*) and all the tested competitors are reported. The selected ranks ( $k$ ) for the precision at rank  $k$  and the success at rank  $k$  are chosen analogously to what was previously done in [25, 29]. To validate the statistical significance of the achieved performance improvements, the student t-test has been adopted [27] by using as  $p$ -value 0.05. Significant worsening in the comparisons between our system and the other tested competitors are starred in Table 2. For each tested measure, the result(s) of the best system(s) is written in boldface.

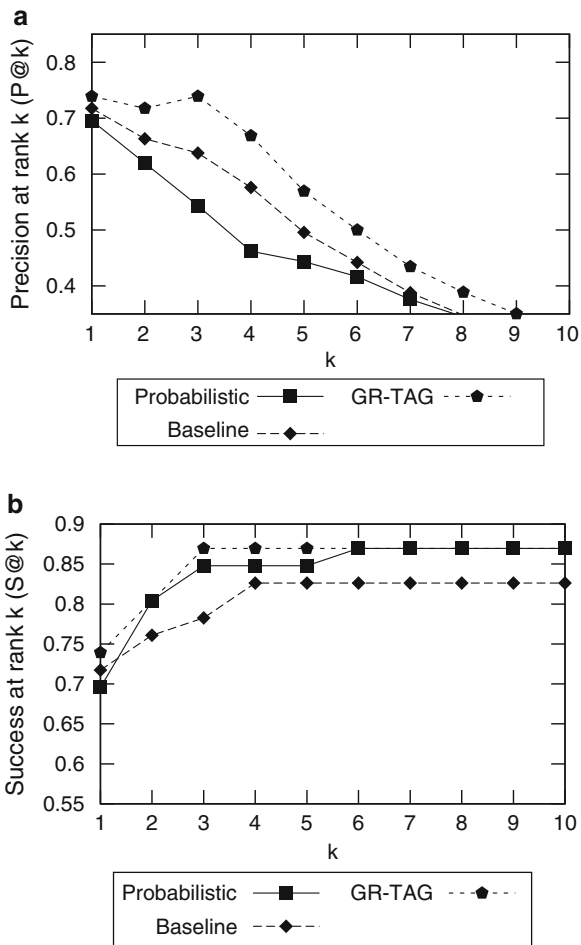
Our recommendation system outperforms both its baseline version and *probabilistic prediction* in terms of all the tested measures. The performance improvement with respect to the baseline version is always statistically significant, while, for *probabilistic prediction*, is significant for MRR, S@1, P@1, P@5. To have a more deep insight into the achieved results, in Fig. 3a and b we also plot the variation of, respectively, the precision and the success by varying  $k$  in the range [1, 10]. Results show that, when increasing the rank value, our system and all the other competitors worsen their performance in terms of precision at rank  $k$ , while averagely perform better in terms of success until reaching a steady state value. Our approach performs best for any value of  $k$  in terms of precision (see Fig. 3b) and for  $k = 1, 3, 4, 5$  in terms of success (see Fig. 3a), while it performs as good as *Probabilistic prediction* in terms of success for the other values of  $k$ .

In summary, results show that our approach averagely selects the most suitable recommendable tags at the top of the ranking and precisely identify the potential user interests.

#### 4.4 Parameter Analysis

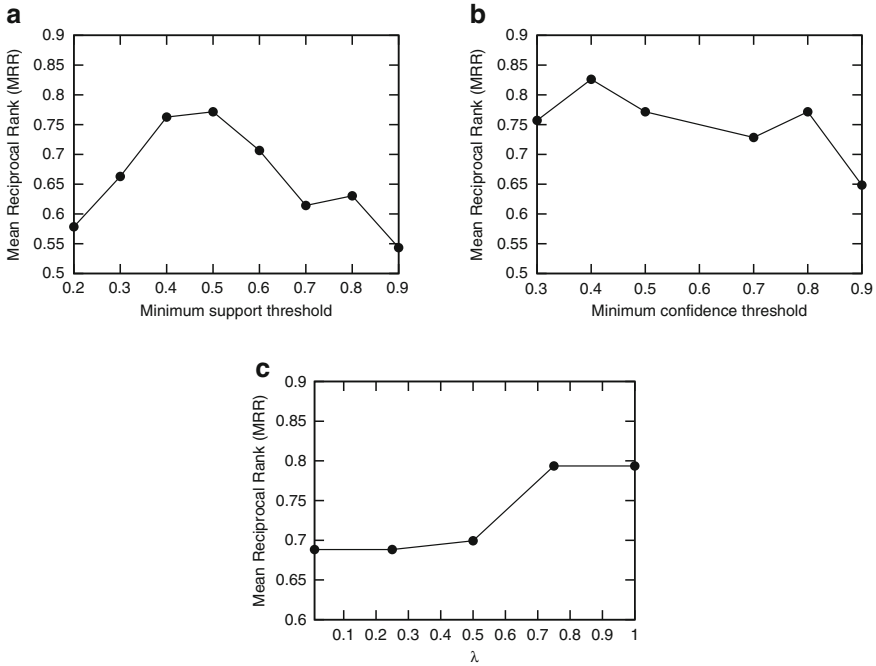
We also analyzed the impact of the system parameters on the performance of the tag recommendation process. In Fig. 4, we plot the average MRR estimate achieved by our *GR-TAG* system by (1) varying the support threshold and by setting the minimum confidence threshold *minconf* to 40% and  $\lambda$  to 0.75 (see Fig. 4a), (2) varying the confidence threshold and by setting the minimum support threshold

**Fig. 3** Performance comparison by varying the reference rank  $k$ . **(a)** Precision at rank  $k$  ( $P@k$ ) and **(b)** success at rank  $k$  ( $S@k$ )



$minsup$  to 30% and  $\lambda$  to 0.75 (see Fig. 4b), and (3) varying the lambda parameter by setting the minimum support threshold  $minsup$  to 30% and the minimum confidence threshold  $minconf$  to 40% (see Fig. 4c).

The support threshold relevantly affects the quality of the tag recommendation. When higher support thresholds (e.g., 70%) are enforced, the percentage of not generalized rules is quite limited (e.g., around 18% of the user-specific rule set mined from the training photo collection, described in Sect. 4.1) and many informative rules (generalized and not) are discarded. Oppositely, when low-support thresholds (e.g., 20%) are enforced, many low-level tag associations (e.g., around 3.5% of the user-specific rule set from the same training data) become frequent and, thus, are extracted by our system. However, the high sparsity of the analyzed tag collections still left some of the most peculiar associations hidden. Aggregating tags into high-level categories allows achieving the best balancing between specialization and generalization of the discovered associations.



**Fig. 4** GR-TAG performance analysis. (a) Impact of the support threshold on MRR estimate.  $minconf = 40\%$ .  $\lambda = 0.75$ , (b) impact of the confidence threshold on P@1 and S@1 estimates.  $minsup = 30\%$ .  $\lambda = 0.75$ , and (c) impact of  $\lambda$  on the MRR estimate.  $minsup = 30\%$ .  $minconf = 40\%$

The confidence threshold may also significantly affect the system performance. By enforcing very low confidence threshold values (e.g., 20%), a large amount of (possibly misleading) low-confidence rules is selected. Indeed, the quality of the rule-based model at the top of which the recommendation system is built worsens. Differently, when increasing the confidence threshold, a more selective pruning of the low-quality rules may allow significantly enhancing the system performance. Finally, when enforcing very high confidence thresholds (e.g., 90%), the rule pruning selectivity becomes too high to allow dealing with a considerable amount of interesting rules.

Finally, we also analyzed the impact of the parameter  $\lambda$  on the achieved MRR. Similarly trends were achieved by using the other tested measures. The value of  $\lambda$  discriminates between the contribution of personal and collective knowledge. More specifically, when  $\lambda < 0.5$ , rules extracted from the community-based history of past annotations are deemed more significant than that discovered from the personal annotation set. Indeed, the recommendation process becomes less personalized, and the knowledge about the personal user interests is partially ignored. Differently, when setting  $\lambda > 0.5$ , tags mainly referable to the personalized rule set are deemed the most relevant ones for tag recommendation. Results show that, as expected, the

proposed system performs significantly better when the recommendation is more personalized, i.e., when user preferences are considered more relevant than the community-based annotations.

## 5 Conclusions and Future Work

In this chapter, we addressed the issue of recommending additional tags to partially annotated Flickr photos by exploiting both the personalized and the collective knowledge. We propose a rule-based recommendation system that also considers associations at higher abstraction levels, i.e., the generalized rules, to counteract the effect of data sparsity on the recommendation performance. A set of experiments performed on a real Flickr photo collection show the effectiveness of the proposed approach.

As pointed out by our work, the integration of the personalized and the collective annotations may effectively improve the quality of the recommended tags. To enrich the background knowledge related to the users and the community, we plan to integrate the analysis of the user-generated content coming from social networks and online communities in the tag recommendation system. Furthermore, we will also address, as future work, the integration in the proposed system of efficient disk-based indexing strategies to store and retrieve very large pattern collections.

## References

1. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Proceedings of the 2008 ACM conference on Recommender Systems, RecSys '08, pp. 335–336. ACM, New York (2008). doi:<http://doi.acm.org/10.1145/1454008.1454068>, <http://doi.acm.org/10.1145/1454008.1454068>
2. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: ACM SIGMOD Record, vol. 22, pp. 207–216. ACM, New York (1993)
3. Bao, S., Xue, G., Wu, X., Yu, Y., Fei, B., Su, Z.: Optimizing web search using social annotations. In: Proceedings of the 16th International Conference on World Wide Web, WWW '07, pp. 501–510. ACM, New York (2007). DOI <http://doi.acm.org/10.1145/1242572.1242640>, <http://doi.acm.org/10.1145/1242572.1242640>
4. Baralis, E., Cagliero, L., Cerquitelli, T., D'Elia, V., Garza, P.: Support driven opportunistic aggregation for generalized itemset extraction. In: IEEE Conference of Intelligent Systems, pp. 102–107 (2010)
5. Baralis, E., Cagliero, L., Cerquitelli, T., Garza, P., Marchetti, M.: Cas-mine: providing personalized services in context-aware applications by means of generalized rules. *Knowl. Inf. Syst.* **28**(2), 283–310 (2011)
6. Baralis, E., Cagliero, L., Cerquitelli, T., Garza, P.: Generalized association rule mining with constraints. *Inf. Sci.* **194**, 68–84 (2012)
7. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. In: Proceedings of the Seventh International Conference on World Wide Web 7, pp. 107–117. Elsevier Science Publishers B. V. Amsterdam, The Netherlands (1998)

8. Cagliero, L.: Discovering temporal change patterns in the presence of taxonomies. *IEEE Trans. Knowl. Data Eng.* **99** (2011, Preprints). doi:<http://doi.ieeecomputersociety.org/10.1109/TKDE.2011.233>
9. Chirita, P.A., Costache, S., Nejd, W., Handschuh, S.: P-tag: large scale automatic generation of personalized annotation tags for the web. In: Proceedings of the 16th International Conference on World Wide Web, WWW '07, pp. 845–854. ACM, New York (2007). doi:<http://doi.acm.org/10.1145/1242572.1242686>, <http://doi.acm.org/10.1145/1242572.1242686>
10. Datta, R., Ge, W., Li, J., Wang, J.Z.: Toward bridging the annotation-retrieval gap in image search. *IEEE MultiMed.* **14**, 24–35 (2007). doi:10.1109/MMUL.2007.67, <http://dl.acm.org/citation.cfm?id=1435658.1436725>
11. Delicious: Delicious. Website (2012). <http://delicious.com>. Accessed 28 Oct 2012
12. Dmitriev, P.A., Eiron, N., Fontoura, M., Shekita, E.: Using annotations in enterprise search. In: Proceedings of the 15th International Conference on World Wide Web, WWW '06, pp. 811–817. ACM, New York (2006). doi:<http://doi.acm.org/10.1145/1135777.1135900>, <http://doi.acm.org/10.1145/1135777.1135900>
13. Elmasri, R., Navathe, S.B.: Fundamentals of Database Systems, 5th edn. Addison Wesley, Boston (2006). <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike04-20&path=ASIN/0321369572>
14. Flickr: Flickr Website (2012). <http://www.flickr.com>. Accessed 28 Oct 2012
15. Garg, N., Weber, I.: Personalized, interactive tag recommendation for flickr. In: Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys '08, pp. 67–74. ACM, New York (2008). doi:<http://doi.acm.org/10.1145/1454008.1454020>, <http://doi.acm.org/10.1145/1454008.1454020>
16. Han, J., Fu, Y.: Mining multiple-level association rules in large databases. *IEEE Trans. Knowl. Data Eng.* **11**(5), 798–805 (2002)
17. Heymann, P., Ramage, D., Garcia-Molina, H.: Social tag prediction. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08, pp. 531–538. ACM, New York (2008). doi:<http://doi.acm.org/10.1145/1390334.1390425>, <http://doi.acm.org/10.1145/1390334.1390425>
18. Jäschke, R., Marinho, L., Hotho, A., Lars, S.T., Gerd, S.: Tag recommendations in folksonomies. In: Proceedings of the 11th European Conference on Principles and Practice of Knowledge Discovery in Databases, PKDD 2007, pp. 506–514. Springer, Berlin/Heidelberg (2007). doi:[http://dx.doi.org/10.1007/978-3-540-74976-9\\_52](http://dx.doi.org/10.1007/978-3-540-74976-9_52), [http://dx.doi.org/10.1007/978-3-540-74976-9\\_52](http://dx.doi.org/10.1007/978-3-540-74976-9_52)
19. Krestel, R., Fankhauser, P., Nejd, W.: Latent dirichlet allocation for tag recommendation. In: Proceedings of the Third ACM Conference on Recommender Systems, RecSys '09, pp. 61–68. ACM, New York (2009). doi:<http://doi.acm.org/10.1145/1639714.1639726>, <http://doi.acm.org/10.1145/1639714.1639726>
20. Lipczak, M., Milios, E.: Efficient tag recommendation for real-life data. *ACM Trans. Intell. Syst. Technol.* **3**(1), 2:1–2:21 (2011). doi:10.1145/2036264.2036266, <http://doi.acm.org/10.1145/2036264.2036266>
21. Lops, P., Gemmis, M., Semeraro, G.: Content-based recommender systems: state of the art and trends. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) *Recommender Systems Handbook*, pp. 73–105. Springer, New York/London (2011)
22. Lu, Y.T., Yu, S.L., Chang, T.C., Hsu, J.Y.j.: A content-based method to enhance tag recommendation. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI'09, pp. 2064–2069. Morgan Kaufmann Publishers, San Francisco (2009). <http://dl.acm.org/citation.cfm?id=1661445.1661775>
23. Mishne, G.: Autotag: a collaborative approach to automated tag assignment for weblog posts. In: Proceedings of the 15th International Conference on World Wide Web, WWW '06, pp. 953–954. ACM, New York (2006). doi:<http://doi.acm.org/10.1145/1135777.1135961>, <http://doi.acm.org/10.1145/1135777.1135961>
24. Pramudiono, I., Kitsuregawa, M.: Fp-tax: tree structure based generalized association rule mining. In: Proceedings of the 9th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery, p. 63. ACM, New York (2004)

25. Rae, A., Sigurbjörnsson, B., van Zwol, R.: Improving tag recommendation using social networks. In: *Adaptivity, Personalization and Fusion of Heterogeneous Information*, RIAO '10, pp. 92–99. Le centre de Haute Etudes Internationales d'Informatique Documentaire, Paris (2010). <http://dl.acm.org/citation.cfm?id=1937055.1937077>
26. Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.): *Recommender Systems Handbook*. Springer, New York/London (2011)
27. Sanderson, M., Zobel, J.: Information retrieval system evaluation: effort, sensitivity, and reliability. In: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05*, pp. 162–169. ACM, New York (2005). doi:<http://doi.acm.org/10.1145/1076034.1076064>, <http://doi.acm.org/10.1145/1076034.1076064>
28. Sarwar, B., Karypis, G., Konstan, J., Reidl, J.: Item-based collaborative filtering recommendation algorithms. In: *Proceedings of the 10th International Conference on World Wide Web, WWW '01*, pp. 285–295. ACM, New York (2001). doi:<http://doi.acm.org/10.1145/371920.372071>, <http://doi.acm.org/10.1145/371920.372071>
29. Sigurbjörnsson, B., van Zwol, R.: Flickr tag recommendation based on collective knowledge. In: *Proceedings of the 17th International Conference on World Wide Web, WWW '08*, pp. 327–336. ACM, New York (2008). doi:<http://doi.acm.org/10.1145/1367497.1367542>, <http://doi.acm.org/10.1145/1367497.1367542>
30. Srikant, R., Agrawal, R.: Mining generalized association rules. In: *Proceedings of the International Conference on Very Large Data Bases*, pp. 407–419. Morgan Kaufmann, San Francisco (1995)
31. Srikant, R., Vu, Q., Agrawal, R.: Mining association rules with item constraints. In: *Proceedings of the Conference on Knowledge Discovery and Data Mining*, vol. 97, pp. 67–73. AAAI Press, Palo Alto, California (1997)
32. Sriphaew, K., Theeramunkong, T.: A new method for finding generalized frequent itemsets in generalized association rule mining. In: *Seventh International Symposium on Computers and Communications*, pp. 1040–1045. IEEE, Washington, DC (2002)
33. Symeonidis, P., Nanopoulos, A., Manolopoulos, Y.: Tag recommendations based on tensor dimensionality reduction. In: *Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys '08*, pp. 43–50. ACM, New York (2008). doi:<http://doi.acm.org/10.1145/1454008.1454017>, <http://doi.acm.org/10.1145/1454008.1454017>
34. van Erp, M., Schomaker, L.: Variants of the borda count method for combining ranked classifier hypotheses. In: *Proceedings of the 7th International Workshop on Frontiers in Handwriting Recognition, Amsterdam*, pp. 443–452 (2000)
35. WordNet: WordNet Lexical Database (2012). <http://wordnet.princeton.edu>. Accessed 28 Oct 2012
36. Zoomr: Zoomr Website (2012). <http://www.zoomr.com>. Accessed 28 Oct 2012