Computers in Human Behavior 50 (2015) 239–251

Contents lists available at ScienceDirect

Computers in Human Behavior

journal homepage: www.elsevier.com/locate/comphumbeh

Research Report

Harnessing the power of social bookmarking for improving tag-based recommendations

Georgios Pitsilis [a,*], Wei Wang [b]

[a] Computer Science Research, Athens, Greece
[b] School of Computer and Information Technology, Beijing Jiaotong University, Beijing, China

ARTICLE INFO

Article history:

Keywords:
Recommender systems
Collaborative tagging
Clustering
Affinity propagation
citeUlike
Taxonomy

ABSTRACT

Social bookmarking and tagging has emerged a new era in user collaboration. *Collaborative Tagging* allows users to annotate content of their liking, which via the appropriate algorithms can render useful for the provision of product recommendations. It is the case today for tag-based algorithms to work complementary to rating-based recommendation mechanisms to predict the user liking to various products. In this paper we propose an alternative algorithm for computing personalized recommendations of products, that uses exclusively the tags provided by the users. Our approach is based on the idea of using the semantic similarity of the user-provided tags for clustering them into groups of similar meaning. Afterwards, some measurable characteristics of users' *Annotation Competency* are combined with other metrics, such as user similarity, for computing predictions. The evaluation on data used from a real-world collaborative tagging system, *citeUlike*, confirmed that our approach outperforms the baseline *Vector Space* model, as well as other state of the art algorithms, predicting the user liking more accurately.

1. Introduction

Collaborative tagging, a web-based service that is representative of the new Web 2.0 technology, allows users to store and share various kinds of web resources, such as news, blogs, and photos into social data repositories. Resources are stored into self-emerging structures called *Folksonomies*, in the form of posts that combine (a) an identifier of the resource, (b) the user who posted it and (c) a set of tags. Many web-based resource sharing and publishing services, like youtube,[1] flickr,[2] and Amazon[3] have already adopted such model, allowing user-generated tags to facilitate user information search. The concept of using tags for on-line annotation of objects, also known as *Social Bookmarking* or *Collaborative Tagging*, constitutes tags as a novel source of information. Although the use of tags has been found very convenient for managing and organizing people's digital material, from the research perspective it seems to have attracted much interest in Recommender Systems (RS) in the recent years, with literature rapidly expanding.

Despite Collaborative Filtering (CF) algorithms being the most adopted techniques for Recommender Systems, the increasing popularity of collaborative tagging systems pushed towards to tags being integrated into the process of recommendation production. Mechanisms which employ the tags alone for computing item recommendations are less common, Jaschke, Marinho, Hotho, Schmidt-thieme, and Stumme (2007), while wherever numeric ratings are additionally provided, they are used complementary to tags for computing item recommendations, Wei, Hsu, and Lee (2011). Relying exclusively on the user-provided tags for computing recommendations, it requires that such information is utilized in the best way for achieving satisfactory quality of predictions. This is the case for digital publication services, like *flickr*, and in general for social networking services, which they provide no-mechanism for evaluation of published content by the users based on numeric ratings.

As a matter of fact, traditional CF models that are based on numeric ratings do not take *Context* into consideration. As opposed to numeric ratings, the semantic information contained in the tags is further exploitable. In addition, tag-words can be classified into hierarchical ordered systems, called *Taxonomies*, structured upon the natural relationships between their elements. Measurements like, *Semantic Distance* and *Relatedness* between tags are computable using the *Taxonomies*. Knowing such distance can prove very useful when needing to group similar tags together, and in many cases, grouping can help to overcome issues like *Polysemy* or *Synonymy* of tags. *Polysemy* exists because a tag might have

* Corresponding author.
E-mail address: georgios.pitsilis@gmail.com (G. Pitsilis).
[1] www.youtube.com.
[2] www.flickr.com.
[3] www.amazon.com.

http://dx.doi.org/10.1016/j.chb.2015.03.045
0747-5632/

multiple related meanings. *Synonymy* exists when different tags share the same or similar meaning, something that is quite common in users' exercise on labeling the same item. In that respect, users express a personal style on the task of labeling, which differs from one another.

While exploring the information hidden on tags for improving recommendations has already been a topic for investigation by the research community in the past, however, the special properties of the users labeling exercise, have not been studied yet. In this paper we attempt to capture the users personal style on the task of annotation with the notion of *Annotation Competency* we introduce. Furthermore, along with giving a useful insight into the *Annotation Competency* of the users, we attempt to utilize the power of *Taxonomies* through tag clustering.

The rest of the paper is organized as follows: Section 2 is referred to related work in the field and in Section 3 we explain our motivation. In Section 4 we reason about the idea of tag clustering in more detail and describe our proposed algorithm. Section 5 is referred to the evaluation tests we performed and the results received, and finally in Section 6 we present our conclusions.

2. Related knowledge

Based on the existing literature, a simple taxonomy of the tag recommender systems could be as depicted in Fig. 1 and it is explained as follows: We can distinguish two major types of algorithms, (a) *tag recommendation algorithms* and (b) *tag-oriented resource recommenders*. The first category comprises solutions aiming to ease the process of annotation, by providing personalized recommendations of tags to users about specific items, Lipczak and Milios (2010), Peng and Zeng (2010) and Symeonidis, Nanopoulos, and Manolopoulos (2008). *Tag recommendation* serves also other purposes, like consolidating the vocabulary across the users and reminding what the resource is about. As Sood, Hammond, Owsley, and Birnbaum (2007) point out, "tag-recommendations fundamentally change the tagging process form generation to recognition". Mechanisms that belong to the *Tag-recommendation* category can either exist as part of a larger concept for resource recommendation, or they can stand as independent services, enabling the social network applications with features of automated annotation of various kinds, Li and Wang (2008) and Moxley, Kleban, and Manjunath (2008). *FolkRank* by Hotho, Jaschke, Schmitz, and Stumme (2006) is an also interesting algorithm which belongs to this category. FolkRank computes PageRank vectors from the tri-party graph of a *Folksonomy* to improve tag recommendations.

The *tag-oriented* category regards prediction models exclusively for resource recommendations, which can be further divided into two sub-categories. For better reference we will label these categories: *Tag-assisted CF* and *Tag-based CF*. In *Tag-assisted CF* models, the computation of recommendations require both item rating values and tags to be provided as input. On the other hand, *Tag-based CF* comprises those models in which recommendations can be computed using the tags alone. The former category has been more explored than the latter one, hence the more literature available, Li, Liang, Nayak, and Xu (2008), Peng, Zeng, Zhao, and Wang (2010), Zeng and Li (2008), Tso-sutter, Marinho, and Schmidt-thieme (1995) and Nakamoto, Nakajima, Miyazaki, and Uemura (2007).

Similarly to all RS algorithms, the *Tag-oriented* type is also distinguished into *Memory-based* and *Model-based*. It is interesting to note that the majority of the models which belong to the *Tag-oriented* category are of *Memory-based* type. As such, the computation of recommendations is carried out in two steps. First, is computed the necessary similarity correlations, as imposed by

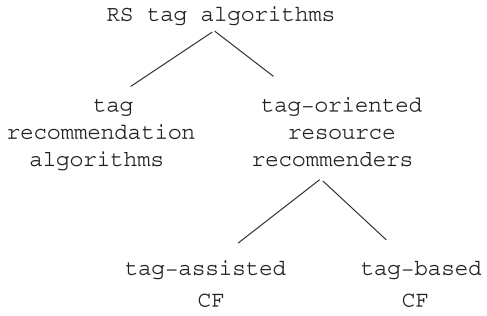

Fig. 1. Taxonomy of algorithms.

the neighborhood-based CF mechanism for user-based or item-based CF, for identifying the top-N most similar neighbors. Next, the item recommendations are worked out using the similarities derived from the first step. It is worth mentioning the works by Li et al. (2008), Zeng and Li (2008) and Peng et al. (2010), in which there has been an effort to user and tag similarities to be combined together into a single expression of similarity, while the work by Parra and Brusilovsky (2009), employs a type of tag-based similarity. Tso-sutter et al. (1995), introduced a method of fusing user-based with item-based CF, treating the users' tags as additional data. Another work by Sen, Vig, and Riedl (2009) is referred to the concept of adding social tagging features to a typical users-ratings based RS, to enhance the quality of recommendations.

Nevertheless, in this work our interest is focused especially on the *Tag-based CF* model, mainly for two reasons. First, because it seems quite reasonable for the user-provided tags to exist as the only available source of information in a RS, and second, the less constraints imposed by the *Tag-based* CF model makes it suitable to a wider range of applications.

For instance, we will refer to one key model from the literature for the *Tag-based CF* type. In the work by Peng and Zeng (2009b), each tag is seen as a distinct topic, while the liking of a user to an item is regarded as the probability of this user experiencing that item. The numeric value of this probability is computed by summing the transition probability over all tags used for annotating this item. The formula they introduced for computing the probability $p(i|u)$ that a user u would like to experience item i is given as follows:

$$p(i|u) = \sum_{t \in T} p(t|u) \cdot p(i|t) \tag{1}$$

where T is the set of tags used by user u, and $p(t|u)$ is the probability that user u chooses tag t for item annotation, $p(i|t)$ is the conditional probability of experiencing item i, when tag t is given. The intuition behind their formula can be phrased as follows: The liking of a user u to item i is highly related to the probability that a particular tag is chosen by that user, as well as the popularity for this tag to be used for annotating item i. In reality, the value of $p(i|u)$ is more or less dependent on the vocabulary used by a user in his/her annotation exercises, something that the existing models do not take into account yet.

We should also point out that for the above reason the computation of recommendations using Eq. (1) becomes inefficient because the vocabularies used for the annotation tasks differ from one user to another. As such, the likelihood for $p(i|u)$ to be computable is strongly affected, requiring a significant overlap to exist between the language elements used by various users. That translates to a serious limitation for such model to work on sparse data.

Another evidence that supports our argument that, the way it is, the computation of recommendations is inefficient, is the fact that, most systems that belong to the above two *Tag-oriented* categories, regard the relationship between the available sources of data as a