

Automatic Recommendations for Machine-Assisted Multimedia Annotation: a Knowledge-Mining Approach

Mónica Díez and Paulo Villegas*

Telefónica I+D, Valladolid, Spain

Abstract. Recommender systems apply knowledge discovery techniques to help in finding associated information. In this paper, we investigate the use of association rule mining as an underlying technology for a recommender system aimed at improving the annotation process of multimedia news documents. The accuracy of these systems is very sensitive to the number of already annotated news items (the "cold-start" problem); ontology-based semantic relations are being used to alleviate this situation.

1 Introduction

To enable advanced query, browsing and content management capabilities, multimedia systems need to be able to provide semantic metadata about the content items they include. This semantic information is hard to obtain: automatic semantic extraction systems are still error-prone, and manual annotation is very costly to produce and can also suffer from lack of consistence between annotators. A number of current research lines strive to join the outcome of automatic knowledge discovery with manual intervention of an operator, achieving thus a semi-automatic approach that intends to merge the best of both worlds.

In this paper, we follow a slightly different and complementary approach, developed within the MESH project [1], which targets multimedia news management systems. Here the mixing of automatic and manual operations is done in a strictly forward fashion, by producing automatic recommendations of possible annotation keywords. This will enable an annotator to quickly select and confirm from among the recommended concepts and keywords, thus cutting down annotation time significantly. To create the suggested keywords we use as a basis the automatically extracted metadata, available context and past history.

2 Automatic Keyword Recommender

The metadata associated to a set of structured documents, such as a collection of semantic annotations of multimedia news documents, can be successfully related

* Work partially supported by the European Comission, under contract number FP6-027685 (MESH). We would also like to thank Pablo Castells and Miriam Fernández from Universidad Autónoma de Madrid for providing the customised IMDb datasets.

to systems containing transactions, like Web usage statistics or frequent shopping list mining: where these contain pages sequentially visited or sets of products bought at the same time, our data is composed of a set of disjoint (but maybe related) concepts linked as annotations to a media content item (a news video), extracted automatically via multimodal analysis of the media.

All these datasets can be viewed as transactions. Thus, our framework can be a good candidate to experiment with some of the technologies that have been developed within association mining for rule extraction.

Association mining is one of the important sub-fields in data mining, where rules that imply certain association relationships among a set of items in a transaction database are discovered. The efforts of most researchers focus on discovering rules in the form of implications between itemsets, which are subsets of items that have adequate supports (i.e. occur frequently enough).

The motivation for using association mining within the context of keyword recommendations is twofold. First, rules extracted from association mining can improve the automatic annotation process of news documents (through better understanding of the semantic concepts inferred from previous multimedia analysis); second, the estimation of further concepts can be considered by reusing annotation history, via the design of adaptive recommender systems.

A first prototype of an AKR (Automatic Keyword Recommender) component was thus developed and integrated in our news media framework, meant to be used for the management of multilingual selections of collections of news video items, related audio broadcasts and text excerpts.

Initial annotations are automatically produced from the content using statistical pattern recognition approaches and multimodal analysis based on modeled news ontologies. The keyword recommendation approach can then complement that automatic analysis by providing additional suggestions (which probably could never be produced by media analysis), to be validated by a human annotator through an specialised interface.

This first version was based on some of the principles from the Apriori algorithm [2], which will find, given an initial set of transactions, association rules that will predict the occurrence of an item based on the occurrences of other items in the transaction.

The pseudocode of this implementation is shown below in Table 1, where A_k denote the annotation used for filtering, R_k denote the k -recommendations returned by the algorithm, C_k denote the set of candidate k -itemsets, and F_k represents the set of frequent k -itemsets.

3 Performance analysis: the cold-start problem

To test and analyze the performance of the implemented keyword recommender we need a huge real-life transactional dataset made up of reasonable sets of keywords. Our own multimedia news framework does not contain yet a big enough set of annotated news, so we complemented the tests with a dataset from the

Table 1. Keyword recommendation generation of the KwdRecom algorithm.

```
1:  $k = 1$ 
2:  $F_k = \{ \text{Find } A_k | \text{support count} \geq \text{minsup threshold} \}$ 
3: do
4:    $k = k + 1$ 
5:    $C_k = \{ \text{Generate candidate itemsets for } (F_{k-1}) \}$ 
6:   for each transaction  $t$  (news document)
7:      $C_t = \{ \text{Identify all candidates} / C_k \in t \}$ 
8:     for each candidate itemset  $c$  in  $C_t$ 
9:       support( $c$ ) ++
10:    end for
11:  end for
12:  $F_k = \{ \text{Extract } C_t | \text{support count} \geq \text{minsup threshold} \}$ 
13:  $R_k = \{ \text{Collection of } k - \text{itemsets} | \text{annotation } A_{\text{pattern}} \text{ takes place} \}$ 
14:  $R_k = \{ \text{Eliminate from } R_k \text{ those annotations that are already present} \}$ 
15: while  $F_k$  is not empty
16: Return  $\{ \cup \{R_k\} \}$ 
```

Internet Movie Database (IMDb) [3], composed of a large list of movies with their associated set of explanatory keywords.

Experimental results carried out on this first version of the recommender showed that sometimes few or no recommendations are produced, because the correlation between the concepts involved is low. This problem decreases the overall efficiency and, much more important, the user confidence on the recommender. Although no recommender system can work without some initial information, the final efficiency is defined by the capacity to manage with a smaller number of examples.

This situation, commonly faced by recommender systems, is known as the cold-start problem [4], and appears whenever recommendations are required for new items for whom little or no information has yet been acquired. Our system cannot discover co-occurrence because there is not enough previously annotated semantic concepts upon which to base any correlations.

This fact has triggered our work to define a new and enhanced AKR, which can use other available information in addition to pure statistical co-occurrences. Two types of such additional information are currently being tried:

- One approach has been to look at the hierarchical structure of the annotated media, trying to find relationships between annotations across the hierarchy.
- The other is to look at the semantic structure of the annotations themselves, taking advantage of the fact that they belong to a pre-defined ontology. This is commented in the next section.

4 Ontologies: a good combination

Ontologies are defined by establishing relationships between linguistic terms that specify domain concepts at different abstraction levels, being structured as a set

of entities, attributes, relationships and axioms [5]. Thus, an ontology can provide a variety of concepts and its relationships from the working domain or scenario, defined from the news document context or the past history.

Taking into account the recommendations produced by the frequency itemset mining approach described in the first section, but trying to overcome its start-up problems, we are currently developing a mixed system that can take advantage of the knowledge defined in an ontology of the semantic concepts used in the annotation process. We are aiming both at lessening the cold-start problem and at an improved guarantee of the quality of the recommendation.

In this approach, an *is-a* hierarchy of relationships driven by the ontology entities is held, so that e.g. higher level entities could be used to infer broader concept interests, increasing the possibilities of finding a recommendation. Other types of relations as given by the ontology structure are also being investigated.

The confidence of the final recommendation is formulated through a correlation between the co-occurrence frequency of the concepts obtained as recommendations from the association mining algorithm, and the interest of those concepts based on the context defined for the news document being processed.

When a specific concept gets affected by the cold-start problem and no recommendations can be retrieved, some super-class entities or concepts derived from ontology-based relationships are used as input in the recommendation algorithm, broadening the activity field of the search. An exponentially decreasing fractional interest is used for those related classes and concepts.

5 Conclusions

In this paper, an algorithm for automatic suggestion of keywords is described, developed within an application for annotating multimedia news documents. Empirical evaluation of a first approach based on association-mining techniques showed that performance decreased in those cases in which little knowledge was available due to scarcity of previous annotations. In order to improve performance, ontologies are being used to make up for the lack of initial information.

References

1. MESH: Multimedia semantic syndication for enhanced news services. (<http://www.mesh-ip.eu>)
2. Agrawal, R., Srikant, R.: Fast algorithms for mining association rules in large databases. In Bocca, J.B., Jarke, M., Zaniolo, C., eds.: Proceedings of the 20th International Conference on Very Large Data Bases, VLDB, Santiago, Chile (1994) 487–499
3. IMDb: The internet movie database. (<http://www.imdb.com>)
4. Maltz, D., Ehrlich, K.: Pointing the way: active collaborative filtering. In: CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems, New York, NY, USA (1995) 202–209
5. Guarino, N., Poli, R.: Formal ontology in conceptual analysis and knowledge representation. International Journal of Human and Computer Studies **43**(5/6) (1995)