

Extracting Multilayered Communities of Interest from Semantic User Profiles: Application to Group Modelling and Hybrid Recommendations

Iván Cantador¹, Pablo Castells

Departamento de Ingeniería Informática, Escuela Politécnica Superior,

Universidad Autónoma de Madrid, 28049, Madrid, Spain

Abstract

A Community of Interest is a specific type of Community of Practice. It is formed by a group of individuals who share a common interest or passion. These people exchange ideas and thoughts about the given passion. However, they are often not aware of their membership to the community, and they may know or care little about each other outside of this clique. This paper describes a proposal to automatically identify Communities of Interest from the tastes and preferences expressed by users in personal ontology-based profiles. The proposed strategy clusters those semantic profile components shared by the users, and according to the clusters found, several layers of interest networks are built. The social relations of these networks might then be used for different purposes. Specifically, we outline here how they can be used to model group profiles and make semantic content-based collaborative recommendations.

Keywords: communities of practice, communities of interest, ontology, user profile, group modeling, content-based collaborative filtering

¹ Corresponding author. Telephone: +34 91 497 2293. Fax: +34 91 497 2235.

E-mail address: ivan.cantador@uam.es (I. Cantador).

1. Introduction

The rapid development, spread and convergence of information and communication technologies and infrastructures in the last two decades, now pervading all aspects of businesses, personal spheres, and our everyday lives, are giving rise to new and unforeseen ways of inter-personal connection, communication, and collaboration. Virtual communities, computer-supported social networks, and collective interaction support technologies are indeed starting to proliferate in increasingly sophisticated ways, opening new research opportunities on social group analysis, modeling, and exploitation.

In this scenario, Communities of Practice (CoP) have been defined as groups of people who get involved in a process of collective work in a shared domain of human endeavour (Wenger, 1998): a community of scientists investigating a specific problem, a group of engineers working on similar projects, a clique of students having a discussion about a common subject, etc. These people collaborate over a period of time, sharing ideas and experiences in order to find solutions and build innovations for a particular practice.

However, it is very often the case that the membership to a community is unknown or unconscious. In many social applications, a person describes his interests and knowledge in a personal profile to find people with similar ones, but he is not aware of the existence of other (directly or indirectly) related interests and knowledge that might be useful to find those people. Furthermore, depending on the context of application or situation, a user can be interested in different topics and groups of people. In both cases, a strategy to automatically identify CoP might be very beneficial (Alani, O'Hara & Shadbolt, 2002).

The issue of finding hidden links between users based on the similarity of their preferences or historic behaviour is not a new idea. In fact, this is the essence of the well-known collaborative recommender systems (Adomavicius & Tuzhilin, 2005; Linden, Smith & York, 2003; Sarwar et al., 2001), where items are recommended to a specific user based on his shared interests with other users, or according to opinions, comparatives, and ratings of items given by similar users. However, in typical approaches, the comparison between users and items is done globally, in such a way that partial, but strong and useful similarities may be missed. For instance, two people may have a highly coincident taste in cinema, but a very divergent one in sports. The opinions of these people on movies could be highly valuable for each other, but risk to be ignored by many collaborative recommender systems, because the global similarity between the users might be low.

Communities of Interest (CoI) are a particular case of CoP, which have been defined as a group of people who share a common interest or passion. They exchange ideas and thoughts about the given passion, creating a self-organizing commune where they come back frequently and remain for extended periods. In this paper, we propose a novel approach towards building emerging multilayered CoI by analyzing the individual motivations and preferences of users, described in ontology-based user profiles, and broken into potentially different areas of personal interest. Like in previous approaches (Liu, Maes, & Davenport, 2006), our method builds and compares profiles of user interests for semantic topics and specific concepts in order to find similarities among users. But in contrast to prior work, we divide the user profiles into clusters of cohesive interests, and based on this, several layers of CoI are found. This provides a richer model of interpersonal links, which better represents the way people find common interests in real life.

Our approach is based on an ontological representation of the domain of discourse where user interests are defined (Castells, Fernández & Vallet, 2007). The ontological space takes the shape of a semantic network of interrelated domain concepts, and the user profiles are initially described as weighted lists measuring the user interests for those concepts. Taking advantage of the relations between concepts, and the (weighted) preferences of users for the concepts, our system clusters the semantic space based on the correlation of concepts appearing in the preferences of individual users. After this, user profiles are partitioned by projecting the concept clusters into the set of preferences of each user. Then, users can be compared on the basis of the resulting subsets of interests, in such a way that several, rather than just one, (weighted) links can be found between two users.

The identified multilayered CoI are potentially useful for many purposes. For instance, users may share preferences, items, knowledge, and benefit from each other's experience in focused or specialized conceptual areas, even if they have very different profiles as a whole. Such semantic subareas need not be defined manually, as they emerge automatically with our proposed method. Users may be recommended items or direct contacts with other users for different aspects of day-to-day life.

In recommendation environments there is an underlying need to distinguish different layers within the interests and preferences of the users. Depending on the current context, only a specific subset of the segments (layers) of a user profile should be considered in order to establish his similarities with other people when a recommendation has to be performed. Models of CoI partitioned at different common semantic layers can enable more accurate and context-sensitive results in recommender processes. Thus, as an applicative development of our automatic semantic clustering and CoI building methods, in this paper we propose and evaluate empirically several

content-based collaborative filtering models that retrieve information items according to a number of real user profiles and within different contexts.

Furthermore, our two-way space clustering, which finds clusters of users based on the clusters of concepts found in a first pass, offers a reinforced partition of the user space that can be exploited to build group profiles for sets of related users. These groups enable an efficient strategy for collaborative recommendation in real-time, by using the merged profiles as representatives of classes of users. To this end, we adapt and test several user profile merging techniques based on social choice theory (Masthoff, 2004), and we show the results of an empirical evaluation to assess which of them are more appropriate for collaborative content retrieval.

The rest of the paper has the following structure. Section 2 summarizes past works on community of practice identification and social collaborative filtering recommendations that are relevant for our proposal. Section 3 describes the ontology-based knowledge representation, upon which our personalised content retrieval processes described in section 4 are built. The proposed clustering technique to build the multi-level relations between users is presented in Section 5. The exploitation of the derived communities of interest to enhance group modeling and content-based collaborative filtering is explained in Sections 6 and 7. Both sections also describe a simple example and early experiments with real subjects and user profiles where the techniques are tested. Finally, some discussions are given in Section 8.

2. State of the Art

In social systems, the profile of a user is mainly composed of his relationships with others, and possible additional information about these relationships: reliability, frequency, context, etc. Connected one to another, the users form graphs of social links,

named in the literature as social networks (Wasserman & Faust, 1994). In these graphs, users' relationships with others are usually described explicitly or can be discovered directly from different sources of information, such as address books, IRC contact lists, or e-mail message boxes. Thus, for example, text classification techniques can be applied to e-mails in order to contextualize and define the topic of relationships, while co-citation of people in web pages can be used to build a social network. Indeed, there have been recently proposed approaches that automatically collect the above and other kind of social network information from the Web in order to apply methods of Semantic Network Analysis (SNA) for the study of online communities (Mika, 2005).

For modelling the social profile of a user, the relationships between users can also be formalised using ontologies. The Friend-Of-A-Friend (FOAF) ontology is one of the most popular in this area. It aims to create a network of machine-readable pages describing people, the links between them and the things they create and do. FOAF is a technology that makes it easier to share and use information about people, their activities and their resources (e.g., photos, calendars, weblogs), to transfer information between web sites, and to automatically extend, merge and re-use it online.

Flink (Mika, 2005b) is a system for the extraction, aggregation and visualization of online social networks. It employs semantic technologies for reasoning with personal information extracted from a number of electronic information sources including web pages, emails, publication archives and FOAF profiles. Extending the traditional bipartite model of ontologies (concepts and instances) with the social dimension leads to a tripartite model of the Semantic Web, namely the layer of communities and their relations, the layer of semantics (ontologies and their relations) and the layer of content items and their relations (the hypertext Web). The application of this representation is

demonstrating in (Mika, 2005c) showing how community-based semantics emerges from this model through a process of graph transformation.

ONTOCOPI (Alani, O'Hara & Shadbolt, 2002) is another tool for discovering Communities of Practice (Wenger, 2000), CoP, by analysing ontologies of a given relevant domain of discourse. It attempts to disclose informal CoP relations by identifying patterns in the relations represented in ontologies, and traversing the ontology from instance to instance via selected relations. Performing experiments to determine particular CoP from an academic ontology, the authors show how the alteration of the weights applied to the ontology's relations affect the structure of the identified CoP.

Up to date, one of the most significant uses of social relations and CoP is the implementation of social collaborative filtering strategies. The most popular collaborative filtering implementations require either a critical mass of referenced resources or a lot of active users. Recent collaborative recommendation solutions are based on finding referrals with expertise on the given domain of discourse. FOAFRealm (Kruk & Decker, 2005) is a user profile distributed management system based on the FOAF metadata. It enables collaboration among people in order to develop effective information retrieval. In the system, users' managed collections are exploited to provide a collaborative filtering strategy that makes use of the social network maintained by the users themselves. Apart of the explicit FOAF friendship relations, the framework controls the access to personal resources, giving different weights of votes during negotiations and specifying the maximum length of the path between different people.

Another ontological approach to user profiling within recommender systems is presented in (Middleton, Roure & Shadbolt, 2004). Working on the problem of recommending on-line academic research papers, the authors present two systems,

Quickstep and Foxtrot, which create user profiles monitoring the behaviour of the users and gathering relevance feedback from them. The obtained profiles are represented in terms of a research paper topic ontology. Research papers are classified using ontological classes, and the proposed collaborative recommendation algorithms suggest documents seen by similar people on their current topics of interest. In this scenario, ontological inference is shown to ease user profiling, external ontological knowledge seems to successfully improve the recommendations, and the profile visualization is used to enhance profiling accuracy.

In (Golbeck & Mannes, 2006), a novel approach for inferring relationships using provenance information and trust annotations in Semantic Web-based social networks is presented. A recommender application, FilmTrust (Golbeck & Hendler, 2006), combines the computed trust values with the provenance of other annotations to personalise the website. The FilmTrust system uses trust to compute personalised recommended movie ratings and to order reviews. The results obtained with FilmTrust illustrate the success that can be achieved using the proposed method. The authors show that the obtained recommendations are more accurate than other techniques when the user's opinions about a film are divergent from the average.

In addition to explicit social relations, recent researches focus their attention in finding implicit relations among people, according to personal tastes, interests and preferences. Hence, for example, the work (Liu, Maes, & Davenport, 2006) presents a theory and implementation of 'taste fabrics', a semantic mining approach to the modeling and computation of personal tastes for different topics of interests. The taste fabric affords a flexible representation of a user in taste-space, enabling a keyword-based profile to be 'relaxed' into a spreading activation (Cohen & Kjeldsen, 1987; Crestani & Lee, 2000) pattern on the taste fabric. An evaluation of taste-based recommendation using the taste

fabric implementation shows that it compares favourably to classic collaborative filtering recommendation methods, and whereas collaborative filtering is an opaque mechanism, recommendation using taste fabrics can be effectively visualized, thus enhancing transparency and user trust.

3. Ontology-based Knowledge Representation

In contrast to other approaches in personalised content retrieval, our approach makes use of explicit user profiles (as opposed to e.g. sets of preferred documents). Working within an ontology-based personalisation framework (Castells et al., 2005), user preferences are represented as vectors $\mathbf{u}_m = (u_{m,1}, u_{m,2}, \dots, u_{m,K})$ where the weight $u_{m,k} \in [0,1]$ measures the intensity of the interest of user $u_m \in U$ for concept $c_k \in O$ (a class or an instance) in the domain ontology O , K being the total number of concepts in the ontology. Similarly, the items $d_n \in D$ in the retrieval space are assumed to be described (annotated) by vectors $\mathbf{d}_n = (d_{n,1}, d_{n,2}, \dots, d_{n,K})$ of concept weights, in the same vector-space as user preferences. Based on this common logical representation, measures of user interest for content items can be computed by comparing preference and annotation vectors, and these measures can be used to prioritize, filter and rank contents (a collection, a catalog, a search result) in a personal way. Figure 1 shows our twofold-space ontology-based knowledge representation, in which M and N are respectively the number of users and items registered in the system.

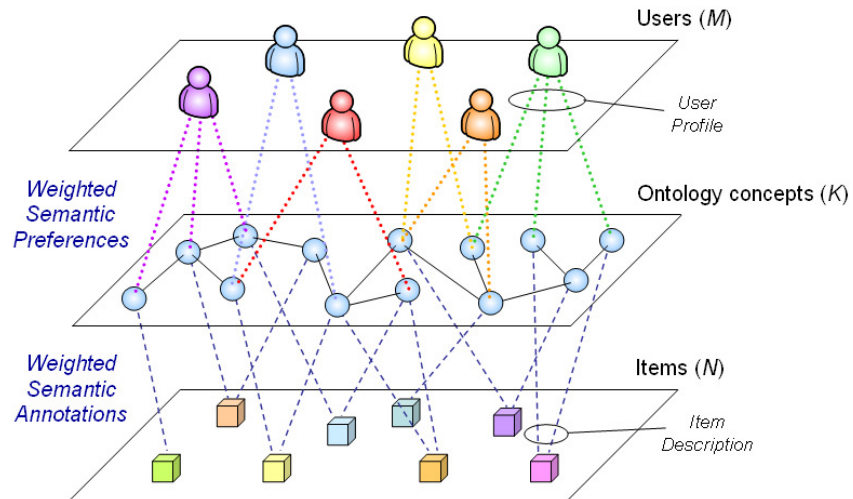


Figure 1. Ontology-based item description and user profile representations

The ontology-based representation is richer and less ambiguous than a keyword-based or item-based model. It provides an adequate grounding for the representation of coarse to fine-grained user interests (e.g. interest for items such as a sports team, an actor, a stock value), and can be a key enabler to deal with the subtleties of user preferences. An ontology provides further formal, computer-processable meaning on the concepts (who is coaching a team, an actor's filmography, financial data on a stock), and makes it available for the personalisation system to take advantage of. Furthermore, ontology standards, such as RDF¹ and OWL², support inference mechanisms that can be used to enhance personalisation, so that, for instance, a user interested in *animals* (superclass of *cat*) is also recommended items about *cats*. Inversely, a user interested in *lizards*, *chameleons* and *snakes* can be inferred with certain confidence to be interested in *reptiles*. Also, a user keen on *Spain* can be assumed to like *Madrid*, through the *locatedIn* transitive relation. These characteristics are exploited in our personalised retrieval model.

4. Personalised Semantic Content Retrieval

¹ <http://www.w3.org/TR/rdf-primer/>

² <http://www.w3.org/TR/owl-ref/>

Our ontology-based retrieval framework assumes the availability of a corpus D of items (texts, multimedia documents, etc.), annotated by domain concepts (instances or classes) from an ontology-based knowledge base O . The knowledge base is implemented using any ontology representation language for which appropriate processing tools (query and inference engines, programming APIs) are available. In our semantic search model, D rather than O is the final search space.

Our retrieval model (wrapped by the ‘Item retrieval’ component in Figure 2) works in two phases. In the first one, a formal ontology-based query (e.g. in RDQL¹) is issued by some form of query interface (e.g. NLP-based) which formalizes a user information need. The query is processed against the knowledge base using any desired inference or query execution tool, outputting a set of ontology concept tuples that satisfy the query. From this point, the second retrieval phase is based on an adaptation of the classic vector-space Information Retrieval model (Baeza-Yates & Ribeiro-Neto, 1999), where the axes of the vector space are the concepts of O , instead of text keywords. Like in the classic model, in ours the query and each item are represented by vectors \mathbf{q} and \mathbf{d} , so that the degree of satisfaction of a query by an item can be computed by the cosine measure:

$$\text{sim}(d, q) = \cos(\mathbf{d}, \mathbf{q}) = \frac{\mathbf{d} \cdot \mathbf{q}}{\|\mathbf{d}\| \|\mathbf{q}\|}$$

The problem, of course, is how to build the \mathbf{d} and \mathbf{q} vectors. For more details, see (Castells et al., 2005). Here we obviate this issue, and continue explaining our content retrieval process with its personalisation phase (component ‘Personalised Ranking’ in Figure 2).

Our personalisation framework is built as an extension of the ontology-based retrieval model. It shares the concept-based representation proposed for retrieval, and the

¹ <http://www.w3.org/Submission/RDQL/>

expressiveness of ontologies to define user interests on the basis of the same concept space that is used to describe contents.

In the framework, the semantic preferences of a user are represented as a vector $\mathbf{u} = (u_1, u_2, \dots, u_K) \in [0, 1]^K$ of concept weights, where for each domain concept $c_k \in \mathcal{O}$, $u_k \in [0, 1]$ represents the intensity of the user interest for c_k . With respect to other approaches, where user interests are described in terms of preferred documents, words, or categories, here an explicit conceptual representation brings all the advantages of ontology-based semantics, such as reduction of ambiguity, formal relations and class hierarchies. Our representation can also be interpreted as fuzzy sets defined on the sets of concepts, where the degree of membership of a concept to a preference corresponds to the degree of preference of the user for the concept.

Once a semantic profile of user preferences is obtained, either automatically and/or refined manually, our notion of content retrieval is based on the definition of a matching algorithm that provides a personal relevance measure $pref(d, u)$ of an item d for a user u . This measure is set according to the semantic preferences of the user, and the semantic annotations of the item. The procedure for matching d and u is based again on a cosine function for vector similarity computation:

$$pref(d, u) = \cos(\mathbf{d}, \mathbf{u}) = \frac{\mathbf{d} \cdot \mathbf{u}}{\|\mathbf{d}\| \|\mathbf{u}\|}$$

In order to bias the result of a search (the ranking) to the preferences of the user, the above measure has to be combined with the query-based score without personalisation $sim(d, q)$ defined previously, to produce a combined ranking (Castells et al., 2005).

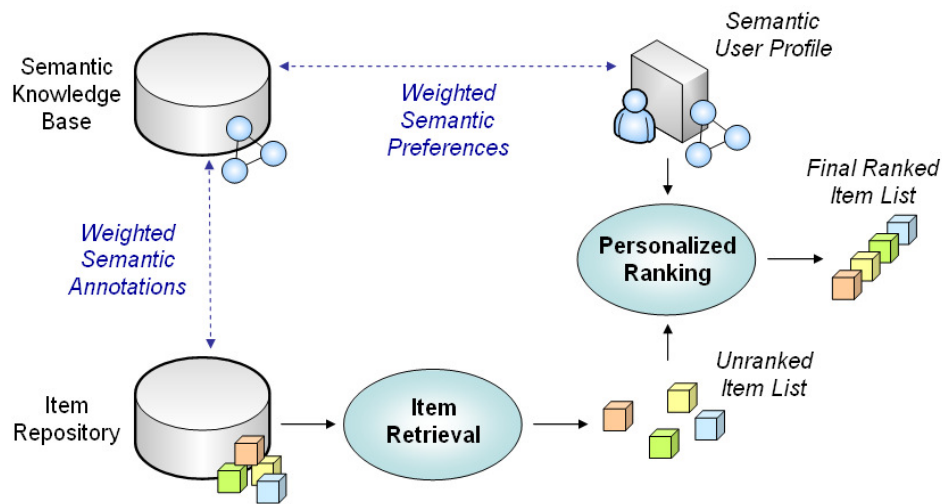


Figure 2. Architecture of the personalised semantically annotated item retrieval process

In real scenarios, user profiles tend to be very scattered, especially in those applications where user profiles have to be manually defined. Users are usually not willing to spend time describing their detailed preferences to the system, even less to assign weights to them, especially if they do not have a clear understanding of the effects and results of this input. On the other hand, applications where an automatic preference learning algorithm is applied tend to recognize the main characteristics of user preferences, thus yielding profiles that may entail a lack of expressivity. To overcome this problem, we propose a semantic preference spreading mechanism, which expands the initial set of preferences stored in user profiles through explicit semantic relations with other concepts in the ontology (Figure 3). Our approach is based on Constrained Spreading Activation (CSA) strategies (Cohen & Kjeldsen, 1987; Crestani & Lee, 2000). The expansion is self-controlled by applying a decay factor to the intensity of preference each time a relation is traversed, and taking into account constraints (threshold weights) during the spreading process.




Figure 3. Preference expansion of a semantic user profile

Thus, the system outputs ranked lists of content items taking into account not only the initial preferences of the current user, but also a semantic spreading mechanism through the user profile and the domain ontology.

We have conducted several experiments showing that the performance of the personalisation system is considerably poorer when the spreading mechanism is not enabled. Typically, the basic user profiles without expansion are too simple. They provide a good representative sample of user preferences, but do not reflect the real extent of user interests, which results in low overlaps between the preferences of different users. Moreover, the extension is not only important for the performance of individual personalisation, but is essential for the clustering strategy described in the next section.

5. Multilayered Semantic Communities of Interest

In social communities, it is commonly accepted that people who are known to share a specific interest are likely to have additional connected interests. For instance, people who share interests in travelling might be also keen on topics related in photography, gastronomy or languages. In fact, this assumption is the basis of most recommender system technologies. We assume this hypothesis here as well, in order to cluster the concept space in groups of preferences shared by several users.

We propose to exploit the links between users and concepts to extract relations among users and derive semantic social networks according to common interests. Analyzing the structure of the domain ontology and taking into account the semantic preference weights

of the user profiles we shall cluster the domain concept space generating groups of interests shared by several users. Thus, those users who share interests of a specific concept cluster will be connected in the network, and their preference weights will measure their degree of membership to each cluster. Specifically, a vector $\mathbf{c}_k = (c_{k,1}, c_{k,2}, \dots, c_{k,M})$ is assigned to each concept vector c_k present in the preferences of at least one user, where $c_{k,m} = u_{m,k}$ is the weight of concept c_k in the semantic profile of user u_m . Based on these vectors a classic hierarchical clustering strategy (Duda, Hart & Stork, 2001) is applied. The obtained clusters (Figure 4) represent the groups of preferences (topics of interests) in the concept-user vector space shared by a significant number of users.




Figure 4. Semantic concept clustering based on the shared interests of the users

Once the concept clusters are created, each user can be assigned to a specific cluster. The similarity between a user's preferences $\mathbf{u}_m = (u_{m,1}, u_{m,2}, \dots, u_{m,K})$ and a cluster C_q is computed by:

$$\text{sim}(\mathbf{u}_m, C_q) = \frac{\sum_{c_k \in C_q} u_{m,k}}{|C_q|} \quad (1)$$

where c_k represents the concept that corresponds to the $u_{m,k}$ component of the user preference vector, and $|C_q|$ is the number of concepts included in the cluster. The

clusters with highest similarities are then assigned to the users, thus creating groups of users with shared interests (Figure 5).




Figure 5. Groups of users obtained from the semantic concept clusters

Furthermore, the concept and user clusters can be used to find emergent, focused semantic Communities of Interest (CoI). The preference weights of the user profiles, the degrees of membership of the users to each cluster, and the similarity measures between clusters are used to find relations between two distinct types of social items: individuals and groups of individuals.

Taking into account the concept clusters, user profiles are partitioned into semantic segments. Each of these segments corresponds to a concept cluster, and represents a subset of the user interests that is shared by the users who contributed to the clustering process. By thus introducing further structure in user profiles, it is now possible to define relations among users at different levels, obtaining a multilayered network of users. Figure 6 illustrates this idea. The image on the left represents a situation where four user clusters are obtained. Based on them (images on the right), user profiles are partitioned in four semantic layers. On each layer, weighted relations among users are derived, building up different semantic Communities of Interest.

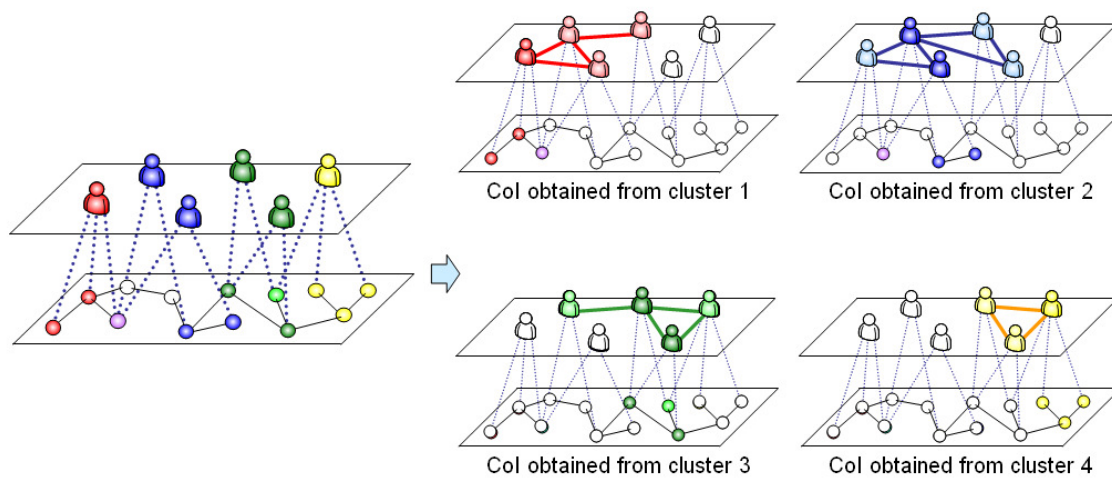


Figure 6. Multilayered semantic Communities of Interest built from the obtained clusters

The resulting semantic CoI have many potential applications. For example, they can be exploited to the benefit of content-based collaborative filtering recommendations, not only because they establish similarities between users, but also because they provide powerful means to focus on different semantic contexts for different information needs. The design of information retrieval models in this direction is explored in Section 7. Additionally, the identified user clusters can be utilized for group profile modeling. In the next section, we propose several user profile merging strategies that attempt to build group profiles that reflect human voting criteria when a choosing of an item has to be made taking into consideration the interests and preferences of a collective.

6. Group Profiles for Content Retrieval

Recently, a number of domains have been identified in which personalisation has a great potential impact, such as news, education, advertising, tourism or e-commerce. It may encompass large range of personal characteristics. Among them, user interest for topics or concepts (directly observed, or indirectly, via user behaviour monitoring followed by system inference) is one of the most useful in many domains, and widely studied e.g. in the user modeling and personalisation research community. However, while the creation

and exploitation of individual models of user preferences and interests have been largely explored in the field, group modeling - combining individual user models to model a group - has not received the same attention (Ardissono et al., 2003; McCarthy & Anagnost, 1998; O'Connor et al., 2001).

It is very often the case that users do not work in isolation. Indeed, the proliferation of virtual communities, computer-supported social networks, and collective interaction (e.g. several users in front of a Set-top Box), call for further research on group modeling, opening new problems and complexities. Collaborative applications should be able to adapt to groups of people who interact with the system. These groups may be quite heterogeneous, e.g. age, gender, intelligence and personality influence on the perception and complacency with the system outputs each member of the groups may have. Of course, the question that arises is how a system can adapt itself to a group of users, in such a way that each individual enjoys or even benefits from the results.

Though explicit group preference modeling has been addressed to a rather limited extent, or in an indirect way in prior work in the computing field, the related issue of social choice (also called group decision making, i.e. deciding what is best for a group given the opinions of individuals) has been studied extensively in economics, politics, sociology, and mathematics (Pattanaik, 1971; Taylor, 1995). The models for the construction of a social welfare function in these works are similar to the group modeling problem we put forward here.

Other areas in which social choice theory has been studied are meta-search, collaborative filtering, and multi-agent systems. In meta-search, the ranking lists produced by multiple search engines need to be combined into one single list, forming the well-known problem of rank aggregation in Information Retrieval. In collaborative filtering, preferences of a group of individuals have to be aggregated to produce a

predicted preference for somebody outside the group. In multi-agent systems, agents need to take decisions that are not only rational from an individual's point of view, but also from a social point of view.

In this work, we study the feasibility of applying strategies, based on social choice theory (Masthoff, 2004), for combining multiple individual preferences in the personalisation framework explained in Section 4, and using the semantic CoI obtained with the user clustering strategy described in Section 5. Several authors have tackled the problem combining, comparing, or merging content-item based preferences from different members of a group. We propose to exploit the expressive power and inference capabilities supported by ontology-based technologies. As we explained before (Section 3), user preferences are gathered in ontology semantic concept-based user profiles. Combining a set of these profiles, the framework retrieves personalised ranked lists of items and shows them in a graphical interface according to the interests and preferences of the members of the group. The mechanism to apply the above strategies in the retrieval process is shown in Figure 7.

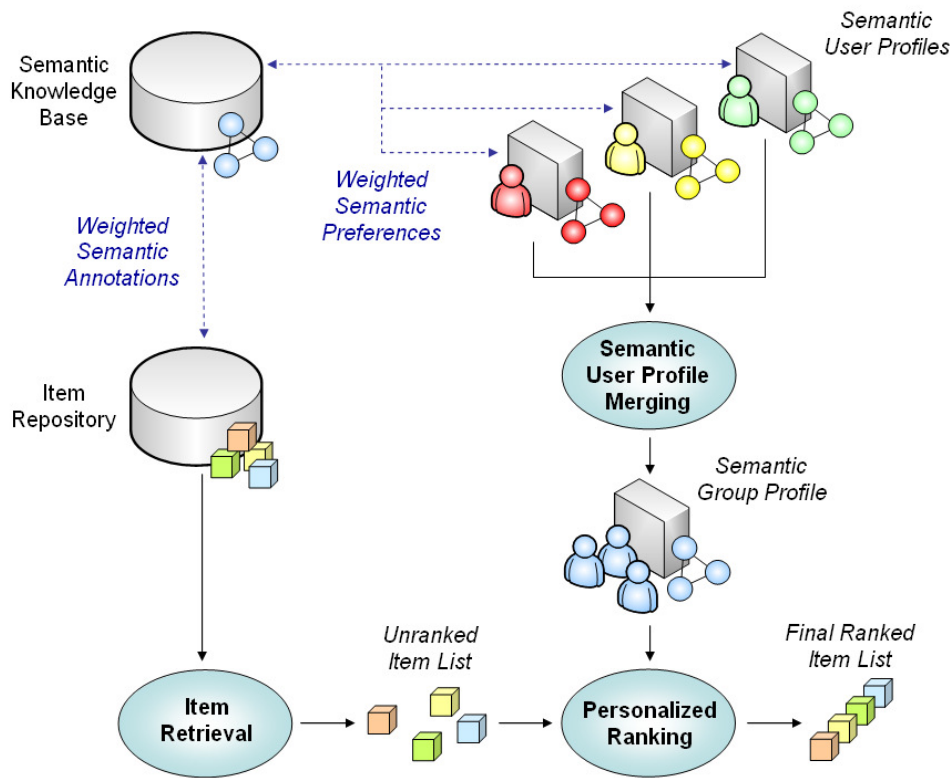


Figure 7. Architecture of the group profile-based semantic annotated item retrieval process

With the combination of several profiles using the considered group modeling strategies we seek to establish how humans create an optimal multimedia item ranked list for a group, and how they measure the satisfaction of a given item list. The theoretical and empirical experiments performed will demonstrate the benefits of using semantic user preferences and exhibit which semantic user profiles combination strategies could be appropriate for a collaborative environment.

In the next two subsections we describe the studied user profile merging strategies and the experiments done to evaluate their feasibility in our information retrieval model.

6.1. Group Modeling Strategies

In (Masthoff, 2004), the author discusses several techniques for combining individual user models to adapt to groups. Considering a list of TV programs and a group of viewers, she investigates how humans select a sequence of items for the group to watch,

how satisfied people believe they would be with the sequence chosen by the different strategies, and how their satisfactions correspond with that predicted by a number of satisfaction functions. These are the three questions we wanted to investigate using semantic user profiles.

In this scenario, because of we have explored the combination of ontology-based user profiles, instead of user rating lists, we had to slightly modify the original techniques described in (Masthoff, 2004). For instance, due to item preference weights have to belong to the range $[0,1]$, the weights obtained for a certain group profile must be normalized after applying the techniques. The following are brief descriptions of the ten selected strategies.

Additive Utilitarian Strategy. Preference weights from all the users of the group are added, and the larger the sum the more influential the preference is for the group. Note that the resulting group ranking will be exactly the same as that obtained taking the average of the individual preference weights.

Multiplicative Utilitarian Strategy. Instead of adding the preference weights, they are multiplied, and the larger the product the more influential the preference is for the group. This strategy could be self-defeating: in a small group the opinion of each individual will have too much large impact on the product. Moreover, in our case it is advisable not to have null weights because we would lose valued preferences. So, if this situation happens, we change the weight values to very small ones (e.g. 10^{-3}).

Borda Count. Scores are assigned to the preferences according to their weights in a user profile: those with the lowest weight get zero scores, the next one up one point, and so on. When an individual has multiple preferences with the same weight, the averaged sum of their hypothetical scores are equally distributed to the involved preferences.

Copeland Rule. Being a form of majority voting, this strategy sorts the preferences according to their Copeland index: the difference between the number of times a preference beats (has higher weights) the rest of the preferences and the number of times it loses.

Approval Voting. A threshold is considered for the preferences weights: only those weights greater or equal than the threshold value are taken into account for the profile combination. A preference receives a vote for each user profile that has its weight surpassing the established threshold. The larger the number of votes the more influential the preference is for the group. In the experiments the threshold will be set to 0.5.

Least Misery Strategy. The weight of a preference in the group profile is the minimum of its weights in the user profiles. The lower weight the less influential the preference is for the group. Thus, a group is as satisfied as its least satisfied member. Note that a minority of the group could dictate the opinion of the group: although many members like a certain item, if one member really hates it, the preferences associated to it will not appear in the group profile.

Most Pleasure Strategy. It works as the Least Misery Strategy, but instead of considering for a preference the smallest weights of the users, it selects the greatest ones. The higher weight the more influential the preference is for the group.

Average Without Misery Strategy. As the Additive Utilitarian Strategy, this one assigns a preference the average of the weights in the individual profiles. The difference here is that those preferences which have a weight under a certain threshold (we used 0.25) will not be considered.

Fairness Strategy. The top preferences from all the users of the group are considered. We have decided to select only the $L/2$ best ones, where L is the number of preferences not assigned to the group profile yet. From them, the preference that least

misery causes to the group (that from the worst alternatives that has the highest weight) is chosen for the group profile with a weight equal to 1. The process continues in the same way considering the remaining $L-1$, $L-2$, etc. preferences and uniformly diminishing to 0 the further assigned weights.

Plurality Voting. This method follows the same idea of the Fairness Strategy, but instead of selecting from the $L/2$ top preferences the one that least misery causes to the group, it chooses the alternative which most votes have obtained.

Some of the above strategies, e.g. the Multiplicative and the Least Misery ones, apply penalties to those preferences that involve dislikes from few users. As mentioned before, this fact can be dangerous, as the opinion of a minority would lead the opinion of the group. If we assume users have common preferences, the effect of this disadvantage will be obviously weaker. Indeed, our multilayer CoI identification algorithm described in Section 5 finds individual profiles with preferences shared by the users in more or less degree.

6.2. Experiments

Two different sets of experiments have been done for this work. The first one will try to find the group modeling strategy that best fits the human way of selecting items when personal tastes of a group have to be considered. We shall try to establish the strategy that most satisfaction offers to the members of the group. The second one tackles the problem in the opposite direction. Given a group modeling strategy, we shall try to determine how to measure the satisfaction the strategy offers to the group.

The scenario of the experiments was the following. A set of twenty four pictures was considered. For each picture several semantic-annotations were taken, describing their

topics (at least one of *beach*, *construction*, *family*, *vegetation*, and *motor*) and the degrees (real numbers in $[0,1]$) of appearance these topics have on the picture. Twenty subjects participated in the experiments. They were Computer Science Ph.D. students of our department. They were asked in all experiments to think about a group of three users with different tastes. In decreasing order of preference (i.e., progressively smaller weights): a) User₁ liked *beach*, *vegetation*, *motor*, *construction* and *family*, b) User₂ liked *construction*, *family*, *motor*, *vegetation* and *beach*, and c) User₃ liked *motor*, *construction*, *vegetation*, *family* and *beach*.

In the following, we describe in detail the experiments done and expose the results and conclusions obtained from them.

Optimal ranking according to human subjects on behalf of a group of users

We have defined two distances that measure the existing difference between two given ranked item lists. The goal is to determine which group modeling strategies give ranked lists closest to those empirically obtained from several subjects.

Consider D as the set of items stored and retrieved by the system. Let $\tau_{sub} \in [0,1]^N$ the item ranked list for a given subject and let $\tau_{str} \in [0,1]^N$ the item ranked list for a specific combination strategy, where N is the number of items stored by the system. We use the notation $\tau(d)$ to refer the position of the item $d \in D$ in the ranked list τ . The first defined distance between these two ranked lists is defined as follows:

$$d_1(\tau_{sub}, \tau_{str}) = \sum_{d \in D} |\tau_{sub}(d) - \tau_{str}(d)| \quad (2)$$

This expression basically sums the differences between the positions of each item in the subject and strategy ranked lists. Thus, the smaller the distance the more similar the ranked lists.

The distance might represent a good measure of the disparity between the user preferences and the ranked list obtained from a group modeling strategy. However, in typical information retrieval systems, where many items are retrieved for a specific query, a user usually takes into account only the first top ranked items. In general, he will not browse the entire list of results, but stop at some top n in the ranking. We propose to more consider those items that appear before the n -th position of the strategy ranking and after the n -th position of the subject ranking, in order to penalize more those of the top n items in the strategy ranked list that are not relevant for the user.

With these ideas in mind, the following could be a valid approximation for our purposes:

$$d(\tau_{sub}, \tau_{str}) = \sum_{n=1}^N P(n) \frac{1}{n} \sum_{d \in D} |\tau_{sub}(d) - \tau_{str}(d)| \cdot \chi_n(d, \tau_{sub}, \tau_{str})$$

where $P(n)$ is the probability that the user stops browsing the ranked item list at position n , and

$$\chi_n(d, \tau_{sub}, \tau_{str}) = \begin{cases} 1 & \text{if } \tau_{str}(d) \leq n \text{ and } \tau_{sub}(d) > n \\ 0 & \text{otherwise} \end{cases}.$$

Again, the smaller the distance the more similar the ranked lists.

The problem here is how to define the probability $P(n)$. Although an approximation to the distribution function for $P(n)$ can be taken e.g. by interpolation of data from a statistical study, we simplify the model fixing $P(10) = 1$ and $P(n) = 0$ for $n \neq 10$, assuming that users are only interested in those multimedia items shown in the screen at first time after a query. Our second distance is defined as follows:

$$d_2(\tau_{sub}, \tau_{str}) = \frac{1}{10} \sum_{d \in D} |\tau_{sub}(d) - \tau_{str}(d)| \cdot \chi_{10}(d, \tau_{sub}, \tau_{str}) \quad (3)$$

Observing the twenty four pictures, and taking into account the preferences of the three users belonging to the group, the subjects were asked to make an ordered list of the pictures. With the obtained lists we measured the distance d_2 with respect to the ranked lists given by the group modeling strategies. The average results are shown in Figure 8. From the figure, it seems that strategies like *Borda Count* and *Copeland Rule*, which do not depend on certain thresholds or parameters, give lists more similar to those manually created by the subjects, and strategies such as *Average Without Misery* and *Plurality Voting* obtained the greatest distances.

This deduction is founded on an empirical point of view. To obtain more theoretical results we also compared the strategies lists against those obtained using semantic user profiles. Surprisingly, they are very similar to the empirical ones. They agree with the strategies that seem to be more or less adequate for group modeling.




Figure 8. Average distance d_2 between the ranked lists obtained with the combination strategies, and the lists created by the subjects and the lists retrieved using the individual semantic user profiles

Human-measured satisfaction for a content ranking on behalf of a group of users

In the previous experiments we tried to find which group modeling strategies generate ranked list most similar to those established by humans and those created from our ontology-based user profiles. The idea behind this search is the assumption that the more similar a ranked list is to that generated from a user profile, the most pleasure causes to the user. In this section, we seek the same goal, but directly trying to measure the satisfaction each strategy provides. This time, the top ten ranked items from each strategy with all the combination methods were presented to the subjects. Then they were asked to decide the degree of satisfaction each list offers to each of the three users in the group. Four different satisfaction levels were used: very *satisfied*, *satisfied*, *unsatisfied* and very *unsatisfied*, corresponding to four, three, two and one vote respectively. The normalized sums of the obtained votes for each strategy are shown in Figure 9.

Once more, a theoretical foundation is needed. In (Masthoff, 2004), three satisfaction functions are presented: a) linear addition satisfaction, b) quadratic addition satisfaction, and, c) quadratic addition minus misery satisfaction. Here, we only study the first one. The quadratic forms are not applicable to our lists because their ratings take values in $[0,1]$, instead of being natural numbers. The way the linear addition satisfaction function measures the pleasure a strategy gives to a specific user is the following. For the n top items of the ranked list τ_{str} , the weights or ratings assigned to these items in the user ranked list are added, and finally normalized:

$$\frac{\sum_{d:\tau_{str}(d)\leq n} r_{sub}(d)}{\sum_{d\in D} r_{sub}(d)}$$

In order to be consistent with the empirical experiments, we established $n=10$. Note that it is necessary for our system to use normalization. The values of the rankings are

skewed within the strategies: some of them are close to 0 and others provide uniform distributed weights in $[0,1]$. Thus absolute satisfactions values can not be considered.




Figure 9. Subject Average Satisfaction and User Normalized Linear Addition Satisfaction

As it can be seen from the figure, the normalized linear addition satisfaction might be a good approximation to real satisfaction values. The satisfaction levels are relatively similar to those obtained from the subjects, especially in the *Plurality Voting*, where both empirical and theoretical satisfactions are the worst of all the studied strategies.

7. Content-based Collaborative Recommendations

Collaborative filtering applications adapt to groups of people who interact with the system, in a way that single users benefit from the experience of other users with which they have certain traits or interests in common. User groups may be quite heterogeneous, and it might be very difficult to define the mechanisms for which the system adapts itself to the groups of users, in such a way that each individual enjoys or even benefits from the results. Furthermore, once the user association rules are defined, an efficient search

for neighbours among a large user population of potential neighbours has to be addressed. This is the great bottleneck in conventional user-based collaborative filtering algorithms. Item-based algorithms attempt to avoid these difficulties by exploring the relations among items, rather than the relations among users. However, the item neighbourhood is fairly static and do not allow to easily apply personalised recommendations or inference mechanisms to discover potential hidden user interests.

We believe that exploiting the relations of the underlying CoI which emerge from the users' interests, and combining them with semantic item preference information can have an important benefit in collaborative filtering recommendation. Using our semantic multilayered CoI proposal explained in Section 5, we present here two recommender models that generate ranked lists of items in different scenarios taking into account the obtained links between users. The first model (that we shall label as UP) is based on the semantic profile of the user to whom the ranked list is delivered. This model represents the situation where the interests of a user are compared to other interests in a social network. The second model (labelled NUP) outputs ranked lists disregarding the user profile. This can be applied in situations where a new user does not have a profile yet, or when the general preferences in a user's profile are too generic for a specific context, and do not help to guide the user towards a very particular, context-specific need. Additionally, we consider two versions for each model: a) one that generates a unique ranked list based on the similarities between the items and all the existing semantic clusters, and, b) one that provides a ranking for each semantic cluster. Thus, we shall study four different retrieval strategies, UP (profile-based), UP- q (profile-based, considering a specific cluster C_q), NUP (no profile), and NUP- q (no profile, considering a specific cluster C_q).

The four strategies are formalized next. In the following, for a user profile \mathbf{u}_m , an information object vector \mathbf{d}_n , and a cluster C_q , we denote by \mathbf{u}_m^q and \mathbf{d}_n^q the projections of the corresponding concept vectors onto cluster C_q , i.e. the k -th component of \mathbf{u}_m^q and \mathbf{d}_n^q is $u_{m,k}$ and $d_{n,k}$ respectively if $c_k \in C_q$, and 0 otherwise.

Model UP. The semantic profile of a user u_m is used by the system to return a unique ranked list. The preference score of an item d_n is computed as a weighted sum of the indirect preference values based on similarities with other users in each cluster. The sum is weighted by the similarities with the clusters, as follows:

$$pref(d_n, u_m) = \sum_q nsim(d_n, C_q) \sum_i nsim_q(u_m, u_i) \cdot sim_q(d_n, u_i) \quad (4)$$

where:

$$sim(d_n, C_q) = \frac{\sum_{c_k \in C_q} d_{n,k}}{\|\mathbf{d}_n\| \sqrt{|C_q|}}, \quad nsim(d_n, C_q) = \frac{sim(d_n, C_q)}{\sum_i sim(d_n, C_i)}$$

are the single and normalized similarities between the item d_n and the cluster C_q ,

$$sim_q(u_m, u_i) = \cos(\mathbf{u}_m^q, \mathbf{u}_i^q) = \frac{\mathbf{u}_m^q \cdot \mathbf{u}_i^q}{\|\mathbf{u}_m^q\| \cdot \|\mathbf{u}_i^q\|}, \quad nsim_q(u_m, u_i) = \frac{sim_q(u_m, u_i)}{\sum_j sim_q(u_m, u_j)}$$

are the single and normalized similarities at layer q between users u_m and u_i , and

$$sim_q(d_n, u_i) = \cos(\mathbf{d}_n^q, \mathbf{u}_i^q) = \frac{\mathbf{d}_n^q \cdot \mathbf{u}_i^q}{\|\mathbf{d}_n^q\| \cdot \|\mathbf{u}_i^q\|}$$

is the similarity at layer q between item d_n and user u_i .

The idea behind this first model is to compare the current user interests with those of the others users, and, taking into account the similarities among them, weight all their complacencies about the different items. The comparisons are done for each concept

cluster measuring the similarities between the items and the clusters. We thus attempt to recommend an item in a double way. First, according to the item characteristics, and second, according to the connections among user interests, in both cases at different semantic layers.

Model UP- q . The preferences of the user are used by the system to return one ranked list per cluster, obtained from the similarities between users and items at each cluster layer. The ranking that corresponds to the cluster for which the user has the highest membership value is selected. The expression is analogous to equation (4), but it does not include the term that connects the item with each cluster C_q .

$$pref_q(d_n, u_m) = \sum_i nsim_q(u_m, u_i) \cdot sim_q(d_n, u_i) \quad (5)$$

where q maximizes $sim(u_m, C_q)$.

Analogously to the previous model, this one makes use of the relations among the user interests, and the user satisfactions with the items. The difference here is that recommendations are done separately for each layer. If the current semantic cluster is well identified for a specific item, we expect to achieve better precision/recall results than those obtained with the overall model.

Model NUP. The semantic profile of the user is ignored. The ranking of an item d_n is determined by its similarity with the clusters, and the similarity of the item with the profiles of the users within each cluster. Since the user does not have connections to other users, the influence of each profile is averaged by the number of users M .

$$pref(d_n, u_m) = \frac{1}{M-1} \sum_q nsim(d_n, C_q) \sum_{i \neq m} sim_q(d_n, u_i) \quad (6)$$

Designed for situations in which the current user profile has not yet been defined, this model uniformly gathers all the user complacencies about the items at different semantic layers. Although it would provide worse precision/recall results than the models UP and

UP- q , this one might be fairly suitable as a first approach to recommendations previous to manual or automatic user profile constructions.

Model NUP- q . The preferences of the user are ignored, and one ranked list per cluster is delivered. As in the UP- q model, the ranking that corresponds to the cluster the user is most close to is selected. The expression is analogous to equation (6), but it does not include the term that connects the item with each cluster C_q .

$$pref_q(d_n, u_m) = \frac{1}{M-1} \sum_{i \neq m} sim_q(d_n, u_i) \quad (7)$$

This last model is the most simple of all the proposals. It only measures the users' complacencies with the items at the layers that best fit them, representing thus a kind of item-based collaborative filtering system.

7.1. An Example

For testing the proposed strategies and models a simple experiment has been set up. A set of twenty user profiles are considered. Each profile is manually defined considering six possible topics: *animals*, *beach*, *construction*, *family*, *motor* and *vegetation*. The degree of interest of the users for each topic is shown in Table 1, ranging over *high*, *medium*, and *low* interest, corresponding to preference weights close to 1, 0.5, and 0.

Table 1. Degrees of interest of users for each topic, and expected user clusters to be obtained

	<i>Motor</i>	<i>Construction</i>	<i>Family</i>	<i>Animals</i>	<i>Beach</i>	<i>Vegetation</i>	Expected Cluster
<i>User1</i>	High	High	Low	Low	Low	Low	1
<i>User2</i>	High	High	Low	Medium	Low	Low	1
<i>User3</i>	High	Medium	Low	Low	Medium	Low	1
<i>User4</i>	High	Medium	Low	Medium	Low	Low	1
<i>User5</i>	Medium	High	Medium	Low	Low	Low	1
<i>User6</i>	Medium	Medium	Low	Low	Low	Low	1
<i>User7</i>	Low	Low	High	High	Low	Medium	2
<i>User8</i>	Low	Medium	High	High	Low	Low	2
<i>User9</i>	Low	Low	High	Medium	Medium	Low	2
<i>User10</i>	Low	Low	High	Medium	Low	Medium	2
<i>User11</i>	Low	Low	Medium	High	Low	Low	2
<i>User12</i>	Low	Low	Medium	Medium	Low	Low	2
<i>User13</i>	Low	Low	Low	Low	High	High	3
<i>User14</i>	Medium	Low	Low	Low	High	High	3
<i>User15</i>	Low	Low	Medium	Low	High	Medium	3
<i>User16</i>	Low	Medium	Low	Low	High	Medium	3
<i>User17</i>	Low	Low	Low	Medium	Medium	High	3
<i>User18</i>	Low	Low	Low	Low	Medium	Medium	3
<i>User19</i>	Low	High	Low	Low	Medium	Low	1
<i>User20</i>	Low	Medium	High	Low	Low	Low	2

As it can be seen from the table, the six first users (1 to 6) have *medium* or *high* degrees of interests in *motor* and *construction*. For them it is expected to obtain a common cluster, named cluster 1 in the table. The next six users (7 to 12) share again two topics in their preferences. They like concepts associated with *family* and *animals*. For them a new cluster is expected, named cluster 2. The same situation happens with the next six users (13 to 18); their common topics are *beach* and *vegetation*, an expected cluster named cluster 3. Finally, the last two users have noisy profiles, in the sense that they do not have preferences easily assigned to one of the previous clusters. However, it is comprehensible that User19 should be assigned to cluster 1 because of his high interests in *construction* and User20 should be assigned to cluster 2 due to his high interests in *family*.

Table 2 shows the correspondence of concepts to topics. Note that user profiles do not necessarily include all the concepts of a topic. As mentioned before, in real world

applications it is unrealistic to assume profiles are complete, since they typically include only a subset of all the actual user preferences.

Table 2. Initial concepts for each of the six considered topics

Topic	Concepts
<i>Motor</i>	Vehicle, Motorcycle, Bicycle, Helicopter, Boat
<i>Construction</i>	Construction, Fortress, Road, Street
<i>Family</i>	Family, Wife, Husband, Daughter, Son, Mother, Father, Sister, Brother
<i>Animals</i>	Animal, Dog, Cat, Bird, Dove, Eagle, Fish, Horse, Rabbit, Reptile, Snake, Turtle
<i>Beach</i>	Water, Sand, Sky
<i>Vegetation</i>	Vegetation, Tree (instance of Vegetation), Plant (instance of Vegetation), Flower (instance of Vegetation)

We have tested our method with this set of twenty user profiles, as explained next. First, new concepts are added to the profiles by the CSA strategy mentioned in Section 4, enhancing the concept and user clustering that follows. The applied clustering strategy is a hierarchical procedure (Duda, Hart & Stork, 2001) based on the Euclidean distance to measure the similarities between concepts, and the average linkage method to measure the similarities between clusters. During the execution, $K - 1$ (with K the total number of distinct concepts stored in the user profiles) clustering levels were obtained, and a stop criterion to choose an appropriate number of clusters would be needed. In our case, the number of expected clusters is three so the stop criterion was not necessary. Table 3 summarizes the assignment of users to clusters, showing their corresponding similarities values. It can be shown that the obtained results completely coincide with the expected values presented in Table 1. All the users are assigned to their corresponding clusters. Furthermore, the users' similarities values reflect their degrees of belonging to each cluster.

Table 3. User clusters and associated similarity values between users and clusters. The maximum and minimum similarity values are shown in bold and italics respectively

Cluster	Users						
1	<i>User1</i>	<i>User2</i>	<i>User3</i>	<i>User4</i>	<i>User5</i>	<i>User6</i>	<i>User19</i>
	0.522	0.562	0.402	0.468	0.356	0.218	<i>0.194</i>
2	<i>User7</i>	<i>User8</i>	<i>User9</i>	<i>User10</i>	<i>User11</i>	<i>User12</i>	<i>User20</i>
	0.430	0.389	0.374	0.257	0.367	<i>0.169</i>	<i>0.212</i>
3	<i>User13</i>	<i>User14</i>	<i>User15</i>	<i>User16</i>	<i>User17</i>	<i>User18</i>	
	0.776	0.714	0.463	0.437	0.527	<i>0.217</i>	

Once the concept clusters have been automatically identified and each user has been assigned to a specific cluster, we apply the information retrieval models presented in the previous section. A set of twenty four pictures was considered as the retrieval space. Each picture was annotated with (weighted) semantic metadata describing what the image depicts using a domain ontology. Observing the weighted annotations, an expert rated the relevance of the pictures for the twenty users of the example, assigning scores between 1 (totally irrelevant) and 5 (very relevant) to each picture, for each user. We show in Table 4 the final concepts obtained and grouped in the semantic Constrained Spreading Activation and concept clustering phases. Although most of the final concepts do not appear in the initial user profiles, they are very important in further steps because they help in the construction of the clusters. In the next subsection we include a study about the influence of the CSA in realistic empirical experiments.

Table 4. Concepts assigned to the obtained user clusters classified by semantic topic

Cluster	Concepts
1	MOTOR: Vehicle, Racing-Car, Tractor, Ambulance, Motorcycle, Bicycle, Helicopter, Boat, Sailing-Boat, Water-Motor, Canoe, Surf, Windsurf, Lift, Chair-Lift, Toboggan, Cable-Car, Sleigh, Snow-Cat CONSTRUCTION: Construction, Fortress, Garage, Road, Speedway, Racing-Circuit, Short-Oval, Street, Wind-Tunnel, Pier, Lighthouse, Beach-Hut, Mountain-Hut, Mountain-Shelter, Mountain-Villa
2	FAMILY: Family, Wife, Husband, Daughter, Son, Mother-In-Law, Father-In-Law, Nephew, Parent, 'Fred' (instance of Parent), Grandmother, Grandfather, Mother, Father, Sister, 'Christina' (instance of Sister), Brother, 'Peter' (instance of Brother), Cousin, Widow ANIMALS: Animal, Vertebrates, Invertebrates, Terrestrial, Mammals, Dog, 'Tobby' (instance of Dog), Cat, Bird, Parrot, Pigeon, Dove, Parrot, Eagle, Butterfly, Fish, Horse, Rabbit, Reptile, Snake, Turtle, Tortoise, Crab
3	BEACH: Water, Sand, Sky VEGETATION: Vegetation, 'Tree' (instance of Vegetation), 'Plant' (instance of Vegetation), 'Flower' (instance of Vegetation)

The four different models are finally evaluated by computing their average precision/recall curves for the users of each of the three existing clusters. Figure 10 shows the results.

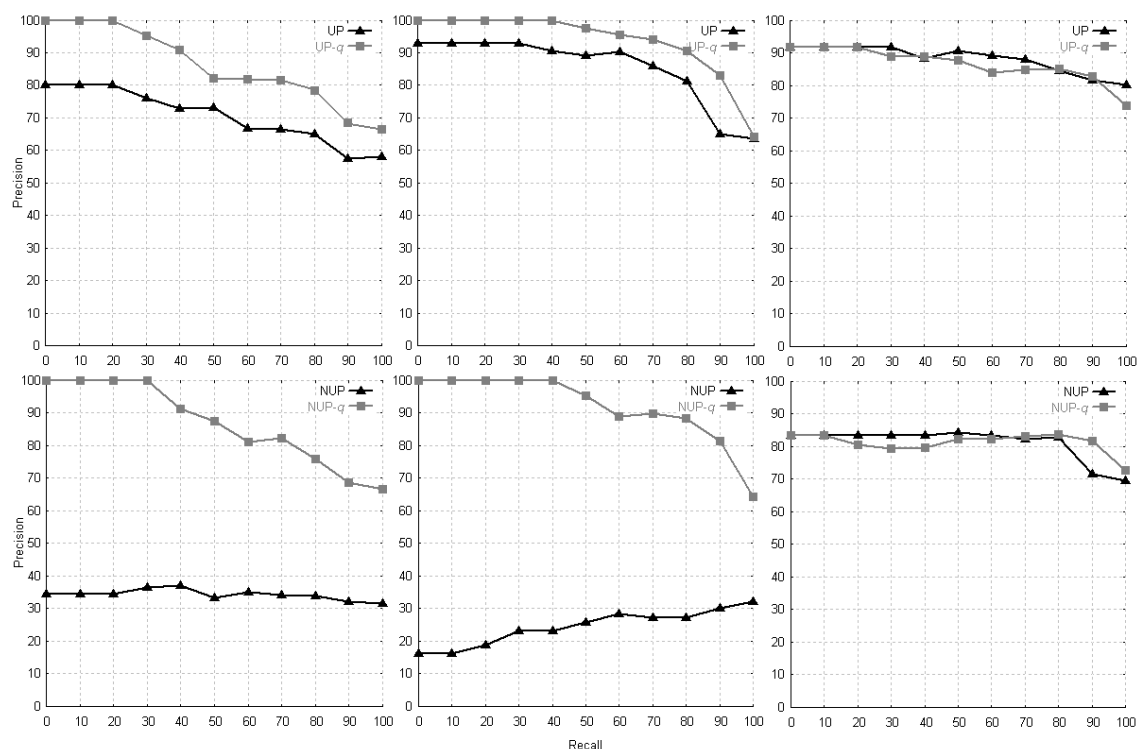


Figure 10. Average precision vs. recall curves for users assigned to cluster 1 (left), cluster 2 (centre) and cluster 3 (right). The graphics on top show the performance of the UP and UP-*q* models. The ones below correspond to the NUP and NUP-*q* models

Two conclusions can be inferred from the results: a) the version of the models that returns ranked lists according to specific clusters (UP- q and NUP- q) outperforms the one that generates a unique list, and, b) the models that make use of the relations among users in the social networks (UP and UP- q) result in significant improvements with respect to those that do not take into account similarities between user profiles.

7.2. Experiments

We have performed an experiment with real subjects in order to evaluate the effectiveness of our proposed recommendation models. Following the ideas exposed in the simple example of the previous section, the experiment was setup as follows.

The set of twenty four pictures used in the example was again considered as the retrieval space. As mentioned before, each picture was annotated with semantic metadata describing what the image depicts, using a domain ontology including six certain topics: *animals*, *beach*, *construction*, *family*, *motor* and *vegetation*. A weight in $[0,1]$ was assigned to each annotation, reflecting the relative importance of the concept in the picture. Twenty graduate students of our department participated in the experiment. They were asked to independently define their weighted preferences about a list of concepts related to the above topics and existing in the pictures semantic annotations. No restriction was imposed on the number of topics and concepts to be selected by each of the students. Indeed, the generated user profiles showed very different characteristics, observable not only in their joint interests, but also in their complexity. Some students defined their profiles very thoroughly, while others only annotated a few concepts of interest. This fact was obviously very appropriate for the experiment done. In a real scenario where an automatic preference learning algorithm will have to be used, the obtained user profiles would include noisy and incomplete components that will hinder the clustering and recommendation mechanisms.

Once the twenty user profiles were created, we run our method. After the execution of the semantic preference spreading procedure, the domain concept space was clustered according to similar user interests. In this phase, because our strategy is based on a hierarchical clustering method, various clustering levels (representable by the corresponding dendrogram) were found, expressing different compromises between complexity, described in terms of number of concept clusters, and compactness, defined by the number of concepts per cluster or the minimum distance between clusters. In Figure 11 we graph the minimum inter-cluster distance against the number of concept clusters.




Figure 11. Minimum inter-cluster distance at different concept clustering levels

A stop criterion has then to be applied in order to determine what number of clusters should be chosen. In this case, we shall use a rule based on the *elbow criterion*, which says you should be choose a number of clusters so that adding another cluster does not add sufficient information. We are interested in a clustering level with a relative small number of clusters and which does not vary excessively the inter-cluster distance with respect to previous levels. Therefore, attending to the figure, we will focus on clustering levels with $R = 4,5,6$ clusters, corresponding to the angle (elbow) in the graph. Table 5 shows the users that most contributed to the definition of the different concept cluster, and their corresponding similarities values.

Table 5. User clusters and associated similarity values between users and clusters obtained at concept clustering levels $R = 4, 5, 6$

R	Cluster	Users										
4	1	User01	User02	User05	User06	User19						
		0.388	0.370	0.457	0.689	0.393						
	2											
	3	User03	User04	User07	User09	User12	User15	User16	User18			
		0.521	0.646	0.618	0.209	0.536	0.697	0.730	0.461			
	4	User08	User10	User11	User13	User14	User17	User20				
		0.900	0.089	0.810	0.591	0.833	0.630	0.777				
5	1	User03	User07									
		0.818	0.635									
	2											
	3	User04	User09	User12	User16	User18						
			0.646	0.209	0.536	0.730	0.461					
	4	User01	User02	User05	User06	User15	User19					
		0.395	0.554	0.554	0.720	0.712	0.399					
	5	User08	User10	User11	User13	User14	User17	User20				
		0.900	0.089	0.810	0.591	0.833	0.630	0.777				
6	1	User6										
		0.818										
	2											
	3	User18										
			0.481									
		4	User02	User05	User06	User19						
		0.554	0.554	0.720	0.399							
	5	User08	User13	User11	User17	User20						
		0.900	0.591	0.810	0.630	0.777						
	6	User01	User04	User07	User09	User10	User12	User14	User15	User16		
		0.786	0.800	0.771	0.600	0.214	0.671	0.857	0.829	0.814		

It has to be noted that not all the concept clusters have assigned user profiles. However, there are semantic relations between users within a certain concept cluster, independently of being associated to other clusters or the number of users assigned to the cluster. For instance, at clustering level $R = 4$, we obtained the weighted semantic relations plotted in Figure 12. Representing the semantic CoI of the users, the diagrams of the figure describe the similarity terms $sim_q(u_i, u_j)$, $i, j \in \{1, 20\}$ (see equations (4) and (5)). The colour of each cell depicts the similarity values between two given users: the dark and light grey cells indicate respectively similarity values greater and lower than 0.5, while the white ones mean no existent relation. Note that a relation between

two certain users with a high weight does not necessary implicate a high interest of both for the concepts on the current cluster. What it means is that they interests agree at this layer. They could really like it or they might hate its topics.




Figure 12. Symmetric user similarity matrices at layers 1, 2, 3 and 4 between user profiles u_i and u_l ($i, l \in \{1, 20\}$) obtained at clustering level $R=4$. Dark and light grey cells represent respectively similarity values greater and lower than 0.5. White cells mean no relation between users

Table 6 shows the concept clusters obtained at clustering level $R=4$. We have underlined those general concepts that initially did not appear in the profiles and were in the upper levels of the domain ontology. Inferred from our preference spreading strategy, these concepts do not necessary define the specific semantics of the clusters, but help to build the latter during the clustering processes.

Table 6. Concept clusters obtained at clustering level $R=4$

Cluster	Concepts
1	ANIMALS: Rabbit CONSTRUCTION: <u>Construction</u> , Speedway, Racing-Circuit, Short-Oval, Garage, Lighthouse, Pier, Beach-Hut, Mountain-Shelter, Mountain-Villa, Mountain-Hut, MOTOR: <u>Vehicle</u> , Ambulance, Racing-Car, Tractor, Canoe, Surf, Windsurf, Water-Motor, Sleigh, Snow-Cat, Lift, Chair-Lift, Toboggan, Cable-Car
2	ANIMALS: <u>Organism</u> , <u>Agentive-Physical-Object</u> , Reptile, Snake, Tortoise, Sheep, Dove, Fish, Mountain-Goat, Reindeer CONSTRUCTION: <u>Non-Agentive-Physical-Object</u> , <u>Geological-Object</u> , <u>Ground</u> , <u>Artifact</u> , Fortress, Road, Street FAMILY: <u>Civil-Status</u> , Wife, Husband MOTOR: <u>Conveyance</u> , Bicycle, Motorcycle, Helicopter, Boat, Sailing-Boat
3	ANIMALS: <u>Animal</u> , <u>Vertebrates</u> , <u>Invertebrates</u> , <u>Terrestrial</u> , <u>Mammals</u> , Dog, ‘Tobby’ (instance of Dog), Cat, Horse, Bird, Eagle, Parrot, Pigeon, Butterfly, Crab BEACH: Water, Sand, Sky VEGETATION: <u>Vegetation</u> , ‘Tree’ (instance of Vegetation), ‘Plant’ (instance of Vegetation), ‘Flower’ (instance of Vegetation)
4	FAMILY: <u>Family</u> , Grandmother, Grandfather, Parent, Mother, Father, Sister, Brother, Daughter, Son, Mother-In-Law, Father-In-Law, Cousin, Nephew, Widow, ‘Fred’ (instance of Parent), ‘Christina’ (instance of Sister), ‘Peter’ (instance of Brother)

Some conclusions can be drawn from this experiment. Cluster 1 contains the majority of the most specific concepts related to *construction* and *motor*, showing a significant correlation between these two topics of interest. Checking the profiles of the users associated to the cluster, we observed they overall have medium-high weights on the concepts of these topics. Cluster 2 is the one with more different topics and general concepts. In fact, it is the cluster that does not have assigned users in Table 6 and does have the most weakness relations between users in Figure 12. It is also notorious that the concepts ‘wife’ and ‘husband’ appear in this cluster. This is due to these concepts were not be annotated in the profiles by the subjects, who were students, not married at the moment. Cluster 3 is the one that gathers all the concepts about *beach* and *vegetation*. The subjects who liked vegetation items also seemed to be interested in beach items. It also has many of the concepts belonging to the topic of *animals*, but in contrast to cluster 2, the annotations were for more common and domestic animals. Finally, cluster 4 collects the majority of the *family* concepts. It can be observed from the user profiles that a number of subjects only defined their preferences in this topic.

Finally, as we did in the example described previously, we evaluate the proposed retrieval models computing their average precision/recall curves for the users of each of the existing clusters. In this case we calculate the curves at different clustering levels ($R = 4, 5, 6$), and we only consider the models UP and UP- q because they make use of the relations among users in the social networks, and offer significant improvements with respect to those that do not take into consideration similarities between user profiles. Figure 13 exposes the results.

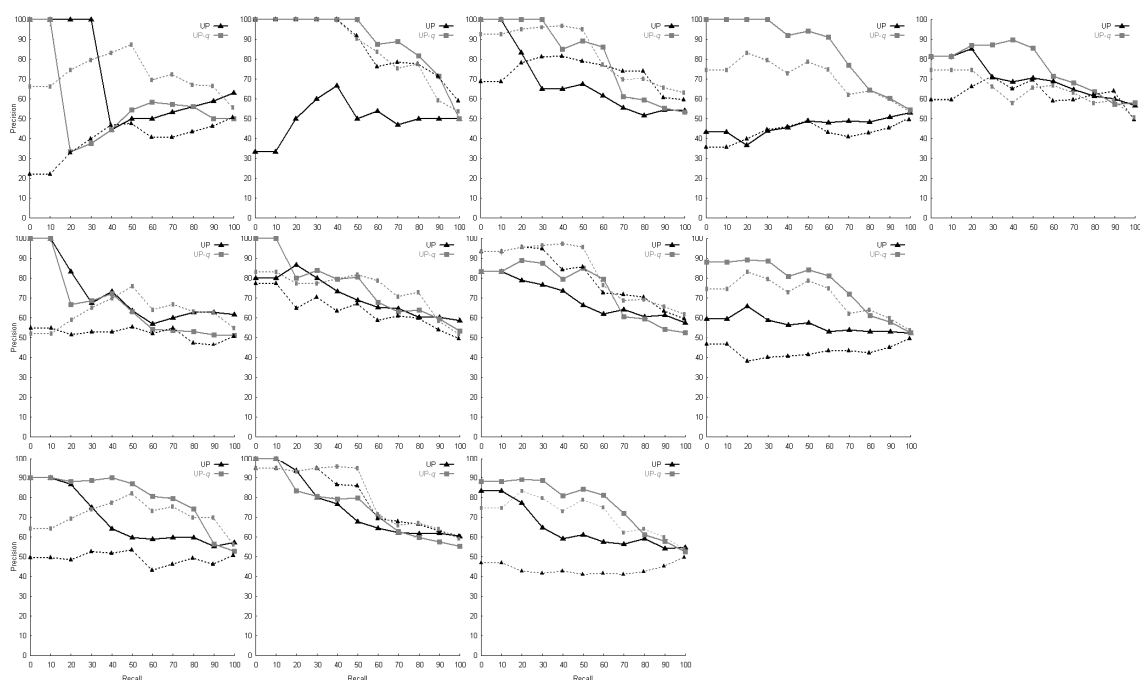


Figure 13. Average precision vs. recall curves for users assigned to the user clusters obtained with the UP (black lines) and UP- q (grey lines) models at levels $R=6$ (graphics on the top), $R=5$ (graphics in the middle), and $R=4$ (graphics on the bottom) concept clusters. For both models, the dotted lines represent the results achieved without semantic preference spreading

Again, the version UP- q , which returns ranked lists according to specific clusters, outperforms the version UP, which generates a unique list assembling the contributions of the users in all the clusters. Obviously, the more clusters we have, the better performance is achieved. The clusters tend to have assigned fewer users and seem more similar to the individual profiles. However, it can be seen that very good results are obtained with only

three clusters. Additionally, for both models, we have plotted with dotted lines the curves achieved without spreading the user semantic preferences. Although more statistically significant experiments have to be done in order to make founded conclusions, it can be pointed out that our clustering strategy performs better when it is combined with the CSA algorithm, especially in the UP- q model. This fact let give us preliminary evidences of the importance of spreading the user profiles before the clustering processes.

8. Discussion

In this work, we have presented an approach to the automatic identification of semantic Communities of Interest according to ontology-based user profiles. Taking into account the semantic preferences of several users we cluster the ontology concept space, obtaining common topics of interest. With these topics, preferences are partitioned into different layers.

The degree of membership of the obtained sub-profiles to the clusters, and the similarities among them, are used to define links that can be exploited by group modeling techniques and collaborative filtering recommendations. Early experiments with real subjects have been done applying the emergent CoI to a variety of group modeling and content-based collaborative filtering strategies showing the feasibility of our clustering strategy. However, more sophisticated and statistically significant experiments need to be performed in order to properly evaluate the models. We have planned to implement a web-based recommender system that will allow users to easily define their profiles, see their semantic relations with other people, and evaluate the existing items and recommendations given by the system. Thus, we expect to enlarge the repositories of items and user profiles, and improve our empirical studies.

Our implementation of the applied clustering strategy was a hierarchical procedure based on the Euclidean distance to measure the similarities between concepts, and the average linkage method to measure the similarities between clusters. Of course, several aspects of the clustering algorithm have to be investigated in future work using noisy user profiles, such as the type of clustering, the distance measure between two concepts, the distance measure between two clusters, the stop criterion that determines what number of clusters should be chosen, and the similarity measure between given clusters and user profiles; we have used a measure considering the relative size of the clusters, but we have not taken into account what proportion of the user preferences is being satisfied by the different concept clusters. Moreover, we have to study efficient clustering strategies based on Latent Semantic Analysis (Deerwester et al., 1990; Landauer, Foltz & Laham, 1998) and/or co-clustering (George & Merugu, 2005).

We are also aware of the need to test our approach in combination with automatic user preference learning techniques in order to investigate its robustness to imprecise user interests, and the impact of the accuracy of the ontology-based profiles on the correct performance of the clustering processes. An adequate acquisition of the concepts of interest and their further classification and annotation in the ontology-based profiles will be crucial to the correct performance of the clustering processes.

Acknowledgements

This work was supported by the Spanish Ministry of Science and Innovation (TIN2008-06566-C04-02) and the Ministry of Industry, Tourism and Commerce (CENIT-2007-1012).

References

Adomavicius, G., and Tuzhilin, A. (2005). Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering* 17(6), 734-749.

Alani, H., O'Hara, K., and Shadbolt, N. (2002). ONTOCOPI: Methods and Tools for Identifying Communities of Practice. *Proceedings of the 17th IFIP World Computer Congress* (pp. 225-236). Montreal, Canada.

Ardissono, L., Goy, A., Petrone, G., Segnan, M., and Torasso, P. (2003). INTRIGUE: Personalised Recommendation of Tourist Attractions for Desktop and Handset Devices. *Applied Artificial Intelligence* 17(8-9), 687-714.

Baeza-Yates, R., and Ribeiro-Neto, B. (1999). Modern Information Retrieval. *Addison-Wesley*.

Castells, P., Fernández, M., and Vallet, D. (2007). An Adaptation of the Vector-Space Model for Ontology-Based Information Retrieval. *IEEE Transactions on Knowledge and Data Engineering* 19(2), pp. 261-272.

Castells, P., Fernández, M., Vallet, D., Mylonas, P., and Avrithis, Y. (2005). Self-Tuning Personalised Information Retrieval in an Ontology-Based Framework. *Proceedings of the 1st IFIP WG 2.12 & WG 12.4 International Workshop on Web Semantics* (pp. 977-986). Agia Napa, Cyprus.

Cohen, P. R., and Kjeldsen, R. (1987). Information Retrieval by Constrained Spreading Activation in Semantic Networks. *Information Processing and Management* 23(4), 255-268.

Crestani, F., and Lee, P. L. (2000) Searching the Web by Constrained Spreading Activation. *Information Processing and Management* 36(4), 585-605.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by Latent Semantic Analysis. *Journal of the Society for Information Science* 41(6), 391-407.

Duda, R. O., Hart, P., and Stork, D. G. (2001). Pattern Classification. *John Wiley*.

George, T., and Merugu, S. (2005). A Scalable Collaborative Filtering Framework based on Co-Clustering. *Proceedings of the 5th IEEE Conference on Data Mining* (pp. 625-628). Houston, TX, USA.

Golbeck, J., and Hendler, J. (2006). FilmTrust: Movie recommendations using trust in web-based social networks. *Proceedings of the 3rd IEEE Consumer Communications and Networking Conference* (pp. 282-286). Las Vegas, NV, USA.

Golbeck, J., and Mannes, A. (2006). Using Trust and Provenance for Content Filtering on the Semantic Web. *Proceedings of the Workshop on Models of Trust on the Web, at the 15th World Wide Web Conference*. Edinburgh, UK.

Kruk, S. R., and Decker, S. (2005). Semantic Social Collaborative Filtering with FOAFRealm. *Proceedings of the 1st Semantic Desktop Workshop, at the 4th International Semantic Web conference*. Galway, Ireland.

Landauer, T. K., Foltz, P. W., and Laham, D. (1998). Introduction to Latent Semantic Analysis. *Discourse Processes* 25.

Linden, G., Smith, B., and York, J. (2003). Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing* 7(1), 76-80.

Liu, H., Maes, P., and Davenport, G. (2006). Unravelling the Taste Fabric of Social Networks. *International Journal on Semantic Web and Information Systems* 2(1), 42-71.

Masthoff, J. (2004). Group Modeling: Selecting a Sequence of Television Items to Suit a Group of Viewers. *User Modeling and User-Adapted Interaction* 14(1), 37-85.

McCarthy, J. F., and Anagnost, T. D. (1998). MusicFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work* (pp. 363-372). Seattle, WA, USA.

Middleton, S. E., Roure, D. D., and Shadbolt, N. R. (2004). Ontology-based Recommender Systems. *Handbook on Ontologies* (pp. 477-498). Springer Verlag.

Mika, P. (2005). Ontologies are us: A Unified Model of Social Networks and Semantics. *Proceedings of the 4th International Semantic Web Conference* (pp. 522-536). Galway, Ireland.

Mika, P. (2005b). Flink: Semantic Web Technology for the Extraction and Analysis of Social Networks. *Journal of Web Semantics* 3(2-3), 211-223.

Mika, P. (2005c). Social Networks and the Semantic Web: The Next Challenge. *IEEE Intelligent Systems* 20(1), 80-93.

O'Connor, M., Cosley, D., Konstan, J. A., and Riedl, J. (2001). PolyLens: A Recommender System for Groups of Users. *Proceedings of the 7th European Conference on Computer Supported Cooperative Work* (pp. 199-218). Bonn, Germany.

Pattanaik, P. K. (1971). Voting and Collective Choice. *Cambridge University Press*.

Sarwar, B. M., Karypis, G., Konstan, J., and Riedl, J. (2001). Item-Based Collaborative Filtering Recommendation Algorithms. *Proceedings of the 10th International World Wide Web Conference* (pp. 285-295). Hong Kong, China.

Taylor, A. (1995). Mathematics and Politics: Strategy, Voting, Power and Proof. *Springer Verlag*.

Wasserman, S., and Faust, K. (1994). Social Network Analysis: Methods and Applications. *Cambridge University Press*.

Wenger, E (1998). *Communities of Practice: Learning, Meaning and Identity.*
Cambridge University Press.

Wenger, E. (2000). *Communities of Practice: The Key to Knowledge Strategy.*
Knowledge and Communities. Butterworth-Heinemann.