

A Multi-Purpose Ontology-Based Approach for Personalized Content Filtering and Retrieval

David Vallet, Iván Cantador, Miriam Fernández, and Pablo Castells
Universidad Autónoma de Madrid, Escuela Politécnica Superior
{david.vallet,ivan.cantador,miriam.fernandez,pablo.castells}@uam.es

Abstract

Personalized multimedia access aims at enhancing the retrieval process by complementing explicit user requests with implicit user preferences. We propose and discuss the benefits of the introduction of ontologies for an enhanced representation of the relevant knowledge about the user, the context, and the domain of discourse, as a means to enable improvements in the retrieval process and the performance of adaptive capabilities. We develop our proposal by describing techniques in several areas that exemplify the exploitation of the richness and power of formal and explicit semantics descriptions, and the improvements therein.

1. Introduction

Personalized multimedia access aims at enhancing the retrieval process by complementing explicit user requests with implicit user preferences, to better meet individual user needs [5]. Automatic user modeling and personalization has been a thriving area of research for nearly two decades, gaining significant presence in commercial applications around the mid-90's. Personalizing a content retrieval system involves considerable complexity, mainly because finding implicit evidence of user needs and interests through their behavior is not an easy task. This difficulty is often considerably increased by an imprecise and vague representation of the semantics involved in user actions and system responses, which makes it even more difficult to properly pair user interests and content descriptions. The ambiguity of terms used in this representation, the unclear relationships between them, their heterogeneity, especially in current ever-growing large-scale networked environments such as the WWW, often constitute a major obstacle for achieving an accurate personalization, e.g. when comparing user preferences to content items, or users among themselves.

In this paper we argue for the introduction of ontologies [8] as an enhanced representation of the relevant knowledge about the domain of discourse, about users, about contextual conditions, involved in the retrieval process, as a means to enable significant improvements in the performance of adaptive content retrieval services. We exemplify our point by describing the development of advanced features and enhancements in specific areas related to personalization where the ontology-based approach shows its benefit, including:

- Basic personalized content search and browsing based on user preferences.
- Dynamic contextualization of user preferences.
- Dynamic augmented social networking and collaborative filtering.

Domain ontologies and rich knowledge bases play a key role in the models and techniques that we propose in the above areas, as will be described in the sequel. The approaches presented in this paper share and exploit a common representation framework, thus obtaining multiple benefits from a shared single ontology-rooted grounding. Furthermore, it will be shown that modular semantic processing strategies, such as inference, graph processing, or clustering, over networked ontology concepts, may be reused and combined to serve multiple purposes.

The rest of the paper is organized as follows. Section 2 introduces the basic approach for the ontology-oriented representation of semantic user preferences, and its application to personalized content search and retrieval. Following this, section 3 describes an approach for the dynamic contextualization of semantic user preferences, and section 4 shows the extension of the techniques described in previous sections to multi-user environments, based on collaborative personalization strategies. Finally, some conclusions are given in section 5.

2. Ontology-based personalization for content retrieval

Most personalized retrieval techniques (e.g. collaborative filtering) keep and process long records of accessed documents by each user, in order to infer potential preferences for new documents (e.g. by finding similarities between documents, or between users). The data handled by these techniques have been rather low-level and simple: document IDs, text keywords and topic categories at most [9], [12]. The recent proposals and achievements towards the enrichment of multimedia content by formal, ontology-based, semantic descriptions open new opportunities for improvement in the personalization field from a new, richer representational level [2], [5]. We see the introduction of ontology-based technology in the area of personalization as a promising research direction [7]. Ontologies enable the formalization of user preferences in a common underlying, interoperable representation, whereby user interests can be matched to content meaning at a higher level, suitable for conceptual reasoning.

An ontology-based representation is richer, more precise, less ambiguous than a keyword-based model. It provides an adequate grounding for the representation of coarse to fine-grained user interests (e.g. interest for individual items such as a sports team, an actor, a stock value) in a hierarchical way, 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 personalization system to take advantage of. Moreover, an ontology-rooted vocabulary can be agreed and shared (or mapped) between different systems, or different modules of the same system, and therefore user preferences, represented this way, can be more easily shared by different players. For instance, a personalization framework may share a domain ontology with a knowledge-based content analysis tool that extracts semantic metadata from a/v content, conforming to the ontology [2]. On this basis, it is easier to build algorithms that match preference to content, through the common domain ontology.

In an ontology-based approach, semantic user preferences may be represented as a vector of weights (numbers from 0 to 1), representing the intensity of the user interest for each concept in a domain ontology [5]. If a content analysis tool identifies, for instance, a cat in a picture, and the user is known to like cats, the personalization module can find how the user may like the

picture by comparing the metadata of the picture, and the preferred concepts in the user profile. Based on preference weights, measures of user interest for content units can be computed, with which it is possible to discriminate, prioritize, filter and rank contents (a collection, a catalog section, a search result) in a personal way.

Furthermore, ontology standards backed by international consortiums (such as the W3C), and the corresponding available processing tools, support inference mechanisms that can be used to further enhance personalization, so that, for instance, a user interested in animals (superclass of cat) is also recommended pictures of cats. Inversely, a user interested in lizards, snakes, and chameleons can be inferred to be interested in reptiles with a certain confidence. Also, a user keen of Sicily can be assumed to like Palermo, through the transitive `locatedIn` relation. In fact, it is even possible to express complex preferences based on generic conditions, such as "athletes that have won a gold medal in the Olympic Games".

3. Contextual personalization

The shallowest consideration is sufficient to notice that human preferences are complex, variable and heterogeneous, and that not all preferences are relevant in every situation [13]. For instance, if a user is consistently looking for some contents in the Formula 1 domain, it would not make much sense that the system prioritizes some Formula 1 picture with a helicopter in the background just because the user happens to have a general interest for aircrafts. In other words, *in the context of* Formula 1, aircrafts are out of (or at least far from) context. Context is a difficult notion to grasp and capture in a software system, and the elements that can, and have been considered in the literature under the notion of context are manifold: user tasks and goals, computing platform, network conditions, social environment, physical environment, location, time, noise, external events, text around a word, visual context of a graphic region, to mention a few.

Complementarily to the ones mentioned, we propose a particular notion, for its tractability and usefulness in semantic content retrieval: that of *semantic runtime context*, which we define as the background themes under which user activities occur within a given unit of time. Using this notion, a finer, qualitative, context-sensitive activation of user preferences can be defined. Instead of a uniform level of personalization, user interests related to the context are prioritized, discarding the preferences that are out of focus. The problems to be addressed include how to represent

such context and determine it at runtime, and how the activation of user preferences should be related to it, predicting the drift of user interests over time. Our approach is based on a concept-oriented context representation, and the definition of distance measures between context and preferences as the basis for the dynamic selection of relevant preferences [13].

A runtime context is represented (is approximated) in our approach as a set of weighted concepts from the domain ontology. This set is obtained by collecting the concepts that have been involved, directly or indirectly, in the interaction of the user (e.g. issued queries and accessed items) with the system during a retrieval session. The context is built in such a way that the importance of concepts fades away with time (number of user requests back when the concept was referenced) by a decay factor.

Once a context is thus built, the contextual activation of preferences achieved by a computation of the semantic similarity between each user preference and the set of concepts in the context. In spirit, the approach consists of finding semantic paths linking preferences to context, where the considered paths are made of existing semantic relations between concepts in the domain ontology. The shorter, stronger, and more numerous such connecting paths, the more *in context* a preference is considered. The proposed techniques to find these paths use a form of *Constraint Spreading Activation* (CSA) strategy [6]. In fact, in the proposed approach a semantic expansion of both user preferences and the context takes place, during which the involved concepts are assigned preference weights and contextual weights, which decay as the expansion grows farther from the initial sets. This process can also be understood as finding a sort of fuzzy semantic intersection between user preferences and the semantic runtime context, where the final computed weight of each concepts represents the degree to which it belongs to each set (see Figure 1).

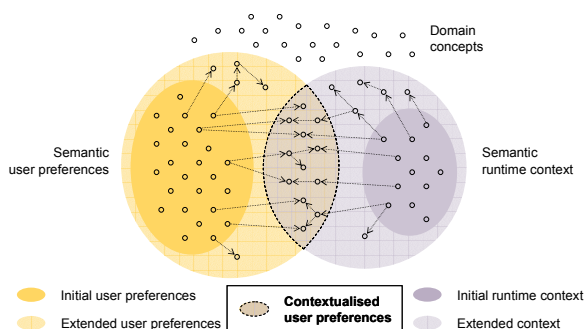


Figure 1. Contextual activation of semantic user preferences.

Finally, the perceived effect of contextualization is that user interests that are out of focus, under a given context, are disregarded, and only those that are in the semantic scope of the ongoing user activity (the “intersection” of user preferences and runtime context) are considered for personalization. The inclusion or exclusion of preferences is in fact not binary, but ranges on a continuum scale, where the contextual weight of a preference decreases monotonically with the semantic distance between the preference and the context.

Contextualized preferences can be understood as an improved, more precise, dynamic, and reliable representation of user preferences, and as such they can be used directly for the personalized ranking of content items and search results, as described in the previous section, or they can be input to any system that exploits this information in other ways, such as the one described in the next section.

4. Augmented social networking and collaborative filtering

When the system perspective is widened to take in contextual aspects of the user, it is often relevant to consider that in most cases the user does 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 STB), call for further research on group modeling, opening new problems and complexities. A variety of group-based personalization functionalities can be enabled by combining, comparing, or merging preferences from different users, where the expressive power and inference capabilities supported by ontology-based technologies can act as a fundamental piece towards higher levels of abstraction [3], [4].

4.1. Semantic group profiling

Group profiling can be understood under the explicit presence of a-priori given user groups, or as an activity that involves the automatic detection of implicit links between users by the system, in order to put users in contact with each other, or to help them benefit from each other’s experience. In the first view, collaborative applications may be required 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 demands on system outputs that each member of the groups may have. The question that arises is how can the system adapt itself to the group in such a way that each individual benefits from the results.

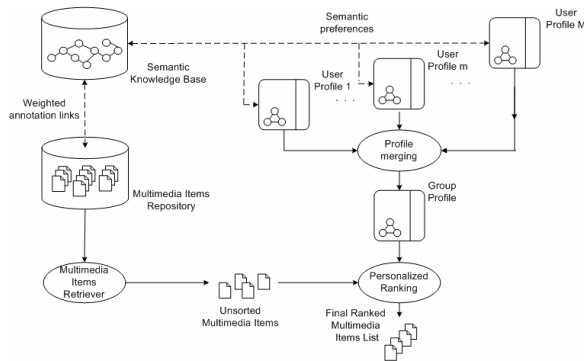


Figure 2. Group profiling by aggregation of individual user profiles.

We have explored the combination of the ontology-based profiles defined in section 2 to meet this purpose, on a per concept basis, following different strategies from social choice theory [11] for combining multiple individual preferences. In our approach, user profiles are merged to form a shared group profile, so that common content recommendations are generated according to this new profile (see Figure 2). Our preliminary experiments have shown that improved results can be obtained from the accuracy and expressivity of the ontology-based representation as proposed in this approach [3].

4.2. Semantic social networking

Even when explicit groups are not defined, users may take advantage of the experience of other users with common interests, without needing to know each other. The issue of finding hidden links between users based on the similarity of their preferences or historic behavior is not a new idea. In fact, this is the essence of the well-known collaborative recommender systems [1], where items are recommended to a certain user concerning those of her interests shared 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.

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 her similarities with other people when a recommendation has to be performed. Models of social networks partitioned into different common semantic layers can achieve more accurate and context-sensitive results.

The definition and generation of such models can be facilitated by a more accurate semantic description of user preferences, as supported by ontologies. A multi-layered approach to social networking can be developed by dividing user profiles into clusters of cohesive interests, so that several layers of social networks are found. This provides a richer model of interpersonal links, which better represents the way people find common interests in real life.

Taking advantage of the relations between concepts, and the (weighted) preferences of users for the concepts, we have defined a system that clusters the semantic space based on the correlation of concepts appearing in the preferences of individual users [4]. After this, user profiles are partitioned by projecting the concept clusters into the set of preferences of each user (see Figure 3). 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.




Figure 3. Multilayer generation of social links between users: 1) the initial sets of individual interests are expanded, 2) domain concepts are clustered based on the vector space of user preferences, and 3) users are clustered in order to identify the closest class to each user.

Multilayered social networks 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 addition to these possibilities, 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 group profiles enable efficient strategies for collaborative recommendation in real-time, by using the merged profiles as representatives of classes of users.

The degree of membership of the obtained subprofiles to the clusters, and the similarities among them, can be used to define social links to be exploited by collaborative filtering systems. We report early experiments with real subjects in [4], where the emergent social networks are applied to a variety of collaborative filtering models, showing the feasibility of the clustering strategy.

4.3. Semantic profile expansion for collaborative group profiling

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, the semantic preference spreading mechanism described in section 3 has proved highly useful for improving our group profiling techniques as well.

Previous experiments without the semantic spreading feature showed considerably poorer results. The profiles were very simple and the matching between the preferences of different users was low. Typically, the basic user profiles 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. Therefore, the extension is not only important for the performance of individual

personalization, but is essential for the clustering strategy described in the previous subsection.

In very open collaborative environments, it is also the case that not only direct evidence of user interests needs to be properly completed in their semantic context, but that they are not directly comparable with the input from other users in its initial form. If the environment is very heterogeneous, the potential disparity of vocabularies and syntax used by different users or subsystems pose an additional barrier for collaborative techniques. One of the major purposes for which ontologies are conceived is that of reflecting or achieving a consensus between different parties in a common knowledge space [8]. Therefore, they provide special-purpose facilities to ensure the required interoperability between semantic user spaces, and match descriptions that are syntactically different but semantically related.

5. Conclusion

In this paper we propose the introduction of ontologies as a key tool for moving beyond current state of the art in the area of personalization. We show ways in which ontology-driven representations can be used to improve the performance of different personalization techniques, and we describe a set of functionalities where ontologies are essential to achieve qualitative enhancements. In our approach, ontologies are used to model the domain of discourse in terms of which user interests, content meaning, retrieval context, and social relationships, can be described and analyzed.

The advantages of ontology-driven representations (expressiveness and precision, formal properties, inference capabilities, interoperability) enable further developments that exploit such capabilities, beyond the ones proposed here, on top of the basic personalization framework described in this paper.

A trade-off of our proposals is the cost and difficulty of building well-defined ontologies and populating large-scale knowledge bases, which is not addressed in this paper. Recent research on these areas is yielding promising results [10], in a way that any advancement on these problems can be played to the benefit of our proposed achievements.

6. Acknowledgements

The research leading to this document has received funding from the European Community's Sixth Framework (FP6-027685 – MESH), and the Spanish Ministry of Science and Education (TIN2005-06885). However, it reflects only the authors' views, and the

European Community is not liable for any use that may be made of the information contained therein.

7. References

- [1] M. Balabanovic and Y. Shoham, "Content-Based Collaborative Recommendation", *Communications of the ACM* 40(3), March 1997, pp. 66-72.
- [2] S. Bloehdorn, K. Petridis, C. Saathoff, N. Simou, V. Tzouvaras, Y. Avrithis, S. Handschuh, Y. Kompatsiaris, S. Staab, M. G. Strintzis, "Semantic Annotation of Images and Videos for Multimedia", *2nd European Semantic Web Conference (ESWC 2005)*. Springer Verlag Lecture Notes in Computer Science, Vol. 3532, 2005.
- [3] I. Cantador, P. Castells, and D. Vallet, "Enriching Group Profiles with Ontologies for Knowledge-Driven Collaborative Content Retrieval", *1st International Workshop on Semantic Technologies in Collaborative Applications (STICA 2006), at the 15th IEEE International Workshops on Enabling Technologies (WETICE 2006)*, Manchester, UK, June 2006.
- [4] I. Cantador and P. Castells, "Multilayered Semantic Social Network Modeling by Ontology-Based User Profiles Clustering: Application to Collaborative Filtering", *15th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2006)*, Podybrady, Czech Republic, October 2006, Springer Verlag Lecture Notes in Computer Science, to appear.
- [5] P. Castells, M. Fernández, D. Vallet, P. Mylonas, and Y. Avrithis, "Self-Tuning Personalized Information Retrieval in an Ontology-Based Framework", *1st International Workshop on Web Semantics (SWWS 2005)*, Agia Napa, Cyprus, July 2005. Springer Verlag Lecture Notes in Computer Science, Vol. 3762, pp. 977-986.
- [6] F. Crestani, "Application of Spreading Activation Techniques in Information Retrieval", *Artificial Intelligence Review* 11, 1997, pp. 453-482.
- [7] S. Gauch, J. Chaffee, and A. Pretschner, "Ontology-based personalized search and browsing", *WIAS Journal* 1(3-4), 2003, pp. 219-234.
- [8] T. R. Gruber, "A Translation Approach to Portable Ontology Specification", *Knowledge Acquisition* 5, 1993, pp. 199-220.
- [9] G. Jeh, J. Widom, "Scaling Personalized Web Search", *International World Wide Web Conference (WWW 2003)*, Budapest, Hungary, 2003, pp. 271-279.
- [10] A. Kiryakov, B. Popov, I. Terziev, D. Manov, D. Ognyanoff, "Semantic Annotation, Indexing, and Retrieval", *Journal of Web Semantics* 2(1), Elsevier, 2004, pp. 47-49.
- [11] J. Masthoff, "Group Modeling: Selecting a Sequence of Television Items to Suit a Group of Viewers", *User Modeling and User-Adapted Interaction* 14(1), 2004, pp. 37-85.
- [12] A. Micarelli, F. Sciarrone, "Anatomy and Empirical Evaluation of an Adaptive Web-Based Information Filtering System", *User Modelling and User-Adapted Interaction* 14(2-3), 2004, pp. 159-200.
- [13] D. Vallet, M. Fernández, P. Castells, P. Mylonas, and Y. Avrithis, "Personalized Information Retrieval in Context", *3rd International Workshop on Modeling and Retrieval of Context (MRC 2006) at the 21st National Conference on Artificial Intelligence (AAAI 2006)*, Boston, USA, July 2006.