A POV-Based User Model: From Learning Preferences To Learning Personal Ontologies

Francesco Osborne[1]
[1]Dept. of Computer Science, University of Torino, 10149 Torino, Italy
osborne@di.unito.it

Abstract. In recent years a variety of ontology based recommender systems, which make use of a domain ontology to characterize the user model, have shown to be very effective. There are however some open issues with this approach, such as: 1) the creation of an ontology is an expensive process; 2) the ontology seldom take into account the perspectives of target user communities; 3) different groups of users may have different domain conceptualizations; 4) the ontology is usually static and not able to learn automatically new semantic relationships or properties. To address these points, I propose an approach to automatically build multiple personal ontology views (POVs) from user feedbacks, tailored to specific user groups and exploited for recommendation purpose via spreading activation techniques.

Keywords: Ontology-based User Modelling, Ontology-based Recommender Systems, Spreading Activation, Personal Ontology, Ontology Learning.

1 Introduction

In recent years a variety of ontology-based recommender systems, which make use of a domain ontology to characterize the user model, have proved to be very effective. The usual implementation of ontology-based user modelling is as an overlay on a domain model represented as an ontology. For example Middleton et al [1] exploit the user feedback on research papers and use the relationships between classes to infer other topics of interest. Many of these works use spreading activation techniques [2] to propagate the interest on a certain item on to similar concepts in the ontology. This is the case of Sieg et al [1] who implement this solution by propagating preferences in an ontology treated as a semantic network. These approaches are usually effective in alleviating the cold start problem, e.g. the situation at the beginning of an interaction with a user, when the system does not have enough data to provide an appropriate adaptation.

There is however some issues associated with most today ontology-based approaches. First of all, ontology crafting is often an expensive process that requires a variety of different expertise, and this is particularly true for very large ontologies describing complex domains and having a high number of classes. Ontologies are usually crafted by domain experts in cooperation with ontology engineers; however different experts may have different opinions on the conceptualization of a domain, and thus need to find some middle of the road compromise. More important, the expert view of a domain can be different from the common user view. Since the

relationship between concepts is to be used to produce recommendations, the ontology should instead try to mirror as best as possible the personal view of a user.

A **personal ontology** [4], also called Personal Ontology View [5] (**POV**), is an ontology that is actually tailored on the prospective and the specific domain view of a user. POVs showed to be effective in assisting tasks like web navigation and search, allowing the user to classify items according to her/his own mental categories [6]. For example a group of users may have a personal ontology of the beverage domain that is structured according to the social context in which drinks are consumed, whereas another group may prefer a classification based on their alcoholic scale or sweetness.

Using spreading activation techniques on an ontology that is too different from the personal view of the user may mislead and fail to predict the right preferences: thus, the need for spreading activation able to exploit the user own personal ontology, or a good enough approximation.

A further problem is posed by the fact that a user may alternate her/his prospective according to the social group she/he is in or the context. To find a way to integrate different perspectives by the same user instead than choosing only one is another interesting challenge.

2 Goals And Objectives

The ontology used in ontology-based user modelling [5] is usually a static knowledge base that is exploited to support various techniques for learning user preferences associated to classes or instances, but which is not able to evolve by itself. A more appealing perspective might be to consider the ontology, and in particular the semantic relationships between concepts, as a dynamic structure that can be learned, adjusted and adapted to the needs of a user or group of users. The learning process may become bidirectional: ontologies help to handle user feedbacks by allowing to propagate interests to similar concepts and these information are returned to the learning module that supports the ontology evolution and adaptation.

Haase et al [7] propose an approach for suggesting the users some adaptations to the domain ontology, mostly in terms of which topic is to be excluded from the personal view. The learning process can however go much further than accepting or rejecting topics or relationships in an existing ontology. A learning module should also be able to discover original concepts, relationships and properties by analysing patterns in the user behaviour. As an example, given an ontology that classifies cheese, we may discover that users who like the Italian cheese “Gorgonzola” tend to like also “Blue Danish” more often then one might expect on the basis of the semantic relationship expressed by the original ontology, in which they are two sibling classes with no apparent common property. Both of these two products are however blue mould cheeses; the original ontology should either include this characteristic or be adapted to accommodate the specific group of users who think that this relationship is significant. To this purpose either a common super-class “Blue Mould Cheese” or another property can be added.

The discussed issues suggest some interesting research questions:

Q1: How can the creation of ontologies for user modelling be made simpler and less expensive?

Q2: How can the user perspectives be integrated within the domain expert opinions in the crafting process?

Q3: Is it possible to automatically learn the user POVs from the feedback?

Q4: What is the best way to handle different mental models of users that would call for different personal ontology views?

Q5: How can the spreading activation techniques be adapted to work in an environment with multiple POVs that can be updated on the run?

To investigate these research questions I propose a POV-based user model that should be able to:

1) **Learn semantic relationships** and relevant **properties** from user feedbacks and use them to craft, refine or adapt a domain ontology (**Q1**, **Q2, Q3**);
2) **Cluster the users in terms of different mental models** and assign to each group a POV derived from the main ontology (**Q4**);
3) **Use a novel spreading activation** technique that is able to propagate user feedback on multiple evolving personal ontologies (**Q5**).

The suggested POV-based user model is an overlay on the user specific POV that will keep being refined as the system learns more about the user and other users similar to her/him. It is also possible to theorize a POV-based user model that allows the user to belong to more than one group and is hence able to handle her/his multiple POVs. Of course, this requires a technique flexible enough to compromise among different views and/or to choose the one deemed most helpful for a given task or context.

3 Present Achievements And Future Directions

To address **Q1, Q2** and **Q3**, I am presently working on an algorithm called *Klink UM* (Klink for User Modelling). It is a modified version of *Klink* [8], an algorithm which uses a combination of statistical and machine-learning techniques to mine semantic relationships among research areas by exploiting the co-occurrences between keywords associated to scientific papers. Preliminary trials showed that such an approach could be adapted to take as input a set of user ratings on items associated with keywords/categories and use this rating patterns to infer semantic relationships among them. While the task of Klink is limited to building a complete ontology from scratch, Klink UM will also be able to receive as input an initial ontology and proceed to mould it according to the perspective inferred by user ratings. Thus Klink UM can be used both to semi-automatically craft a domain ontology (**Q1**) and to refine it, by integrating the users point of view (**Q2**). If fed with the ratings of a specific group it will yield a specific POV tailored to it (**Q3**). A preliminary evaluation on the very first version of Klink UM yielded good results in creating ontologies similar to the human crafted ones. An interesting challenge will be to add property support, e.g. the ability to detect how a certain property is important for a given group of users and to detect possible implicit properties to be manually inserted.

The identification of groups of users with different mental models of the domain (**Q4**) will be based on a novel similarity distance that takes in consideration the closeness of the various perceptions of the relationships between the elements of the domain. So, a user who estimates a set of elements as conceptually very near when they are linked by semantic relationship R whereas very distant when they have

different values for property P, would be clustered with users who share the same attitude and not with users who instead consider property P not very discriminant.
I plan to use this metric with a clusterization algorithm to find groups of users with a similar view of the domain and return their feedbacks to Klink UM to build their POV. I think that a clustering approach might be easier and more realistic than building a specific and probably redundant POV for every single user in the system. However a user may also adopt different mental models according to the context and the social group with which she/he is in a certain moment: thus my longer-term plan is to permit a user to be associated with more than one cluster with which she/he is compatible (non-exclusive clustering), e.g. to allow her/him to refer to different POVs.

Finally, to address **Q5**, I am working on a spreading activation technique that takes in consideration 1) the nature of semantic relationships, 2) the properties and their values, and 3) the similarity of the classes according to user feedback. The stepping-stone has been my recent work on propagation techniques that use the topic distance in a conceptual hierarchy [9] and on a conceptual similarity distance based on common or different properties of two classes [10]. The main challenge for the future will be to develop an original approach capable of maintaining a reliable and coherent user model even when using multiple POVs, with all of them able to be continuously enriched and refined by new evidence.

References

1. Middleton, S.E., Shadbolt, N.R., De Roure, D.C.: Ontological user profiling in recommender systems. ACM Transactions on Information Systems. (2004)
2. Salton, G., Buckley, C.: On the use of spreading activation methods in automatic information. In: Proceedings of the 11th annual international ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 1988, pages 147–160, Grenoble, France. (1988)
3. Sieg, A., Mobasher, B., Burke, R.: Web search personalization with ontological user profiles. In: Proceeding of the 16th ACM Conference on Information and Knowledge Management. CIKM '07. ACM, pp. 525–534. (2007)
4. Huhns, M. N., Stephens, L. M.: Personal Ontologies, IEEE Internet Computing 3 (5), pp. 85 – 87. (1999)
5. Sosnovsky, S., Dicheva, D.: Ontological technologies for user modelling, International Journal of Metadata, Semantics and Ontologies 5(1): 32-71. (2010)
6. Chaffee, J., Gauch, S.: Personal ontologies for web navigation In: Proceedings of CIKM 2000: Proceedings of the ninth international conference on Information and knowledge management, pp. 227-234. (2000)
7. Haase, P., Hotho, A., Schmidt-Thieme, L., Sure, Y.: Collaborative and Usage-driven Evolution of Personal Ontologies. In: Proceedings of 2nd European Semantic Web Conference, Heraklion, Greece. Springer. (2005)
8. Osborne, F., Motta, E.: Mining Semantic Relations between Research Areas. In: Proceedings of the 11th International Semantic Web Conference. Boston, USA. (2012)
9. Cena, F., Likavec, S., Osborne, F.: Propagating user interests in ontology-based user model. In: Proceedings of Advances in Artificial Intelligence, AI*IA 2011, volume 6934 of LNCS, pages 299– 311. Springer-Verlag. (2011)
10. Cena, F., Likavec, S., Osborne, F.: Property-based interest propagation in ontology-based user model. In: Proceedings of the 20th Conference on User Modeling, Adaptation and Personalization. Montreal, Canada. (2012)