On the Construction of a Probabilistic Model for Assessing Users' Beliefs in an Information Environment

Eric Schwarzkopf
German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3
66123 Saarbruecken, Germany
Eric.Schwarzkopf@dfki.de

Abstract: We are developing an interaction environment for information portals in which users can create their own personal knowledge bases to store and define relationships between retrieved items. In this paper, we present our approach for combining these personal knowledge bases to a knowledge base for the entire user community that is to be used as part of an automatic recommendation system. We utilize a probabilistic model for assessing how strongly a single user's definition (or lack thereof) of a relationship should influence the system's belief in the relationship. The latter is revised using Pearl's method of uncertain evidence.

1 Introduction

An important requirement of an information system is to assist a user in efficiently identifying resources relevant to her current information need. Automatic recommendation systems are a valuable tool towards fulfilling this requirement. But how to acquire the data required to provide high quality recommendations? One approach is to elicit this data from the user community itself: By allowing users to rank items or capturing for each user certain interactions with the information system (such as downloading a paper) that are assumed to indicate the endorsement of a resource.

We are currently developing an interaction environment for scientific information portals [4] that includes a component for creating personal knowledge bases: Each user can create her own collection of information items and define relationships between those items, where the set of possible relations is predefined. For example, a user can add the fact that *"Document A elaborates on document B"* to her knowledge base by defining an `elaborates on` relationship between A and B. Utilizing this data in a recommendation system should prove to be useful.

Our goal is to combine the personal knowledge bases of a user community to a community knowledge base, which could then be used, for example, to assist members of that community to identify documents that `elaborate on` a given document. In the following, we assume that we have already identified a community of users who agree on whether a specific relationship exists between two given items, that is, we assume there is a *canon-*

ical knowledge base c for the group of users we are interested in. Hence, our task is to construct c from the set of personal knowledge bases. This is trivial if we assume that a user defines a relationship $r(A, B)$ between items A and B in her knowledge base if and only if $r(A, B)$ is valid in c, but sadly, this is rather unrealistic: Whether a user defines $r(A, B)$ depends on whether she recognizes the relationship, cares about it, and actually utilizes her knowledge base as intended, among other things. That is, we cannot assume that a user who does not define a relationship believes that this relationship is invalid. In order to cope with the uncertain evidence, we use a probabilistic model to assessing an individual user's belief in the validity of a relationship given the user's observed interactions with her knowledge base. These assessments are then combined to assess the community's belief in the respective relationship. A detailed description of this approach is presented in section 3.

We are currently implementing the approach described in this paper. In section 4, we discuss a preliminary evaluation, forming the first step towards this implementation. Section 5 provides our conclusions and directions for future work.

2 Related Work

Uncertain reasoning is applied widely in the user modeling community. Here are only two among many related works:

Horvitz et al. [2] describe the Lumière project in which probabilistic models for inferring a software user's need for assistance were developed. This is one of the most well known projects applying Bayesian reasoning in the domain of user modeling.

Jameson [3] provides a survey about methods for handling uncertainty in user modeling. Although published in 1996, the presented approaches and issues are still current today.

3 The Approach

Imagine there is an observer assigned to each user who assess the probability of a relationship $r(A, B)$ being valid based only on his observations of the user's interaction with her personal knowledge base. From time to time, he reports his then current assessment to his director, who is responsible for providing an approximation of the user communty's canonical model. Given the new assessments of his minions, the director revises his belief in the validity of $r(A, B)$ in the canonical model. We chose this arrangement as a model for our approach because it has some beneficial properties:

- Observers are independent from one another and there is only a loose coupling between an observer and the director. This maps very well to our system architecture, which employs a number of fat clients connected via a possibly low-bandwidth connection to an information portal.

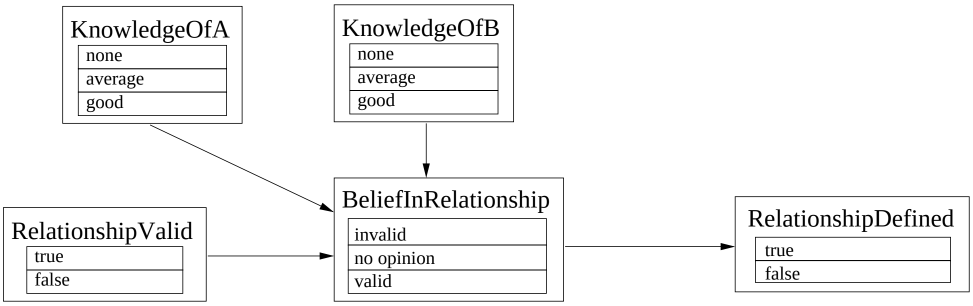

Figure 1: Structure of a Bayesian network representing the initial probabilistic model of an observer. Rectangles represent probabilistic variables and their possible values

- The expensive process of collecting and analyzing observations to generate an assessment is separated from the comparatively cheap combination of the assessments.
- Because each observer is assigned to a single user, the observer can adapt to his user in order to provide more accurate assessments.

In our approach, the observer is represented by a probabilistic model, depicted in figure 3 by a Bayesian network. This model reflects the assumed causal dependencies between the user's belief in the validity of a relationship $r(A, B)$ and whether she defines $r(A, B)$ in her knowledge base, as well as between the validity of $r(A, B)$ in the canonical model, her knowledge of A, and her knowledge of B and her belief in $r(A, B)$. Our observer component assesses the user's knowledge of A and B and adapts its belief in whether $r(A, B)$ is valid in the canonical model based on this assessment and whether the user defined $r(A, B)$ or not. This belief is then sent to the component representing the director.

The director component manages a database of beliefs in the validities of relationships in respect to the canonical model. The belief in the validity of a relationship $r(A, B)$ is revised when an observer changes its respective assessment or provides such an assessment for the first time. This revision is performed by utilizing Pearl's method of virtual evidence as generalized by Chan et al. [1]. In our concrete case, this results in the following revision formular, where $P(r(A, B))$ denotes the current belief, $P'(r(A, B))$ the revised belief, and $O(r(A, B))$ the new belief supplied by an observer:

$$P'(r(A,B)) = P(r(A,B)) \times \frac{O(r(A,B))}{O(r(A,B))\times P(r(A,B)) + (1-O(r(A,B)))\times(1-P(r(A,B)))}$$

This method of revision has two important properties:

- Iterated belief revisions commute. Thus, an observer can retract formerly supplied evidence at any time by simply sending "inverse" evidence $(1-O(r(A, B)))$ without any special processing required from the director.
- Evidences must not be influenced by any background knowledge already incorporated into the belief which is being revised (the *"Nothing else considered"* interpretation of evidence). For our system, this means in particular that an observer does

not take anything the director learned about the canonical model into account: The only flow of evidence is from the observer to the director.

But is there a good reason to expect that the described approach can be applied successfully in a real-world setting? And is it necessary for an observer to adapt to its user to provide accurate data for the director? As a first step towards answering these questions, we performed tests with the simplest instantiation of our approach: Observers that only take the definition of relationships into account (i.e., the `RelationDefined` node in figure 3 is the only observed node) and do not adapt to their users. Thus, the probabilistic model was used only as a tool to derive two constant belief values to be sent to the director if a user defined or did not define a relationship, respectively.

4 A Preliminary Evaluation

The first question we were concerned with was whether we can expect our approach to work at all. To test this, we decided on the properties of a possible user community in order to create the observers' probabilistic model, and used this model to generate several sets of sample data for our evaluation. In this community, it was expected that a relationship would be defined by about 19% of the users if it were valid, and by about 5% otherwise. This resulted in a belief of about 0.8 in the validity of a relationship if its definition was observed, and about 0.46 otherwise. We created several data sets for a valid or invalid relationship and 1000 or 100 observations from the probability distribution corresponding to the test community. With an initial belief of 0.5 (reflecting ignorance) in the validity of the relationship in the canonical model, we were able to determine accurately whether the relationship was valid or not for all sample sets.

Can we expect similar results for other communities of which we have a perfect model? We repeated the evaluation with two other models we derived from our original one by changing the probability of users believing in the validity of a relationship providing true positive (valid relationship defined) or false positive (invalid relationship defined) feedback: Our second evaluation community was expected to provide 14% true positive and 4% false positive feedback, our third 8% and 2%, respectively. We selected this property because we think that it is among the most critical in our application domain: People who believe that a relationship is invalid or do not have an opinion about the relationship are not very likely to define it, and only a small number of people who believe in the relationship might actually decide to define it. This provides us with only very little true positive evidence. We expect this lack of true positive feedback to increase with the number of relations available to the users, so the 19% true positive percentage of the first community could be too optimistic. For both the second and the third community, sample sizes of 1000 led once again to accurate results, but there was a noticeable decline in accuracy with the samples of size 100. This was to be expected because we decreased the ratio between true positive and false positive evidence, thus requiring more evidence to gain confidence in a conclusion.

Is the decline of accuracy when there are only 100 samples of relevance? Of course we aim for an accurate result with as few samples as possible, and in the case of information items that are not well known, expecting to receive at least 100 samples might be too optimistic. So how could we improve the accuracy when only a few samples are available? Since in the current evaluations we use a perfect model of the community, there is not much sense in trying to somehow improve it, but we could adapt it to each individual user and relationship. We repeated the evaluations with observers which have perfect knowledge of the users' knowledge of the information items. The resulting beliefs were only slightly more accurate than without this adaptation, so it does not seem to make sense to spend effort in improving an oberserver's assessment in its user's knowledge in models similar to those used in our evaluation.

An oberserver's belief in a relationship is 0.79 using the second and about 0.76 using the third model. The differences to the original weight of 0.8 are quite modest, which gives rise to the question of how robust the models are in respect to the properties of the community. To get an indication of this robustness, we applied each of our models to the samples generate from the two other models. Here, the resulting accuracy is pretty poor with sample sizes of 100, unsurprisingly decreasing with the difference in the models (the initial model cannot identify the valid relationship with any sample generated from the third model). Applying a simple adaption as described above did not have any significant effects. This means that we should not rely on guessed values for the probabilistic model, but adapt our initial model to the actual behavior of the user community.

5 Conclusions

At the time of writing, we were in the process of identifying a probabilistic model for determining how strongly the definition of a relationship by an individual user should be weighted when the community's belief in this relationship is being assessed. Our simple evaluation indicates that guessing a presumably sensible probabilistic model is unlikely to lead to good results. Thus, in the next development steps, we will focus on how to adapt the initially deployed model to the actual user community.

6 Acknowledgments

I thank Anthony Jameson for his ongoing guidance and support. The descriped work is being funded by the German Ministry of Education, Research, Science, and Technology under grant 08C5940 as part of the SemIPort project.

References

[1] H. Chan and A. Darwiche. On the Revision of Probabilistic Beliefs using Uncertain Evidence. In *Proceedings of the 18th International Conference on Artificial Intelligence (IJCAI)*, 2003.

[2] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumière project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. In G. Cooper and S. Moral, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference*, pages 256–265. Morgan Kaufmann, 1998.

[3] A. Jameson. Numerical Uncertainty Management in User and Student Modeling: An Overview of Systems and Issues. *User Modeling and User-Adapted Interaction*, 5(3-4):193–251, 1996.

[4] E. Schwarzkopf. Enhancing the Interaction with Information Portals. In *Proceedings of the 9th International Conference on Intelligent User Interfaces (IUI)*, pages 322–324. ACM Press, 2004.