Memory Structures as User Models

Leendert van Maanen (leendert@ai.rug.nl)
Jelmer Borst (jpborst@ai.rug.nl)
Chris Janssen (C.Janssen@ai.rug.nl)
Hedderik van Rijn (D.H.Van.Rijn@rug.nl)
Artificial Intelligence, University of Groningen
Grote Kruisstraat 2/1
9712 TS Groningen, The Netherlands

Introduction

The role of information increases. Both for individuals as for society as a whole, handling information has become a tremendously important aspect of daily live. Simultaneously, the amount of available information increases as well. Given this current information overload (Brusilovsky & Tasso, 2004), research into personalization and recommender systems seems necessary. Applications that limit the amount of information presented to a user by selecting only relevant information would be extremely useful.

Relevant information could be filtered by creating a personal profile of a user, and subsequently selecting information that fits the constraints of that profile. We refer to such a profile as a user model (Brusilovsky & Tasso, 2004). The user model could be explicitly created by presenting a user with a questionnaire on her interests, and using the answers to the questionnaire as a model of that user's interests. A drawback of this approach is that it takes time for a questionnaire to be completed, and the user is thus presented with even more information than before. In addition, in many situations users find it hard to explicate their interests, or their interests may change over time, making it hard to infer their interest using a questionnaire. Therefore, implicit inference of user interests should be applied, for instance using eye movements (Van Maanen *et al.*, 2006) or mouse clicks (Claypool *et al.*, 2001).

ACT-R's declarative memory structure might prove useful for maintaining these personal profiles. ACT-R proposes that chunks in declarative memory are characterized by activation, a quantity that reflects how likely it is that a chunk will be needed in the immediate future (Anderson & Schooler, 1991). The level of activation depends on the history of usage of a chunk (base-level activation), and a component reflecting the influence of the current context on a chunk's activation (spreading activation, Anderson & Lebiere, 1998). The spreading activation component is a weighed sum of the activation of associated chunks, with the weights being the strengths of association. The chunks in ACT-R's declarative memory module form a semantic network structure, in which the edges represent spreading activation between chunks.

The strengths of association can be determined by looking at the frequency of co-occurrences of chunks. If two words frequently co-occur, the presence of one word can be regarded as a predictor for the presence of the other word. However, if a word co-occurs with many different words (such as for instance determiners), than the predictive value of that word is less (Posterior Strength Equation, Anderson & Lebiere, 1998).

These strengths of association may also reflect individual interests. As an example, consider the case of a sports fan reading the newspaper: For her, reading a newspaper will usually involve reading the sports section. Therefore, chunks representing sports related notions and chunks representing the newspaper co-occur more frequently for a sports fan than for a non-sports fan. In ACT-R, a higher strength of association would thus be created between newspaper chunks and sports related chunks for sports fans than for non-sports fans.

Image Recommender System

This feature of ACT-R's associative strength learning mechanism can be exploited to create personalized applications. Searching images on the internet is a typical domain in which personalization is useful, because image search based on one key word generally results in very diverse search results. For instance, searching for the key word *apple* results in images of fruit and images of computers, and searching for the key word *mouse* results in images of rodents or images of computer equipment. Using ACT-R's declarative memory structure, we have developed a recommender system that expands search queries for image search. The Image Recommender System functions as follows.

The user can issue a query to an online image search engine (we used Yahoo! Search SDK), which returns a series of images. By clicking on an image, the user can indicate interest in that particular image. Each time the user indicates interest in an image, the website that contains the image is parsed, and the words are harvested. The assumption is that the words on the websites visited by a user represent not only the content of the websites, but are also indicative of the content of the images on these websites.

Because the user only (or at least generally) visits websites that are of interest to her, the words on these websites also reflect her interests. Spreading activation between these words is calculated using the Posterior Strength Equation. To reduce the computational load, high-frequent words in the semantic network are excluded. These words will probably not influence the recommendations that the system will give, because they likely co-occur with many other words, resulting in low spreading activation. The words that are excluded are for instance determiners or pronouns. Also to reduce the computational load, only the ten most frequent words on a webpage plus the search query are used to calculate strengths of activation, because, low frequent words on a webpage are less indicative of the contents of a webpage than higher frequent words. Again, these words would have low spreading activation. It should be noted that there is no principled reason for these implementation choices, but are only intended to make the size of the semantic network incorporated in the Image Recommender System feasible.

Every new query triggers a retrieval from declarative memory and provides an opportunity to train the strengths of association. The query is stored as a goal chunk, which, being in the focus of attention, spreads activation towards all associated chunks. Given the individualized strengths of association, different chunks might be retrieved for individual users: The chunk with the highest activation will be retrieved, which differs for individual users. The retrieved chunk is the chunk that is the most associated with the goal chunk (search query). That is, the word represented by that chunk occurred most frequently in the context of the query key word. Since the frequency of co-occurrence is determined by the mouse clicking behavior of the user, the retrieved chunk also represents the most likely notion of interest for the current user in the current context. The retrieved chunk is used to expand the search query. In another application, it could be involved in some other personalized task component

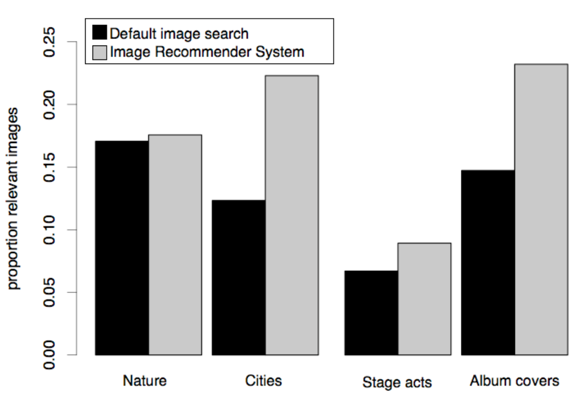

Figure 1. Proportion of relevant images returned by the standard search engine (Without Association) and returned by the Image Recommender System (With Association), for four different image categories and two different query sets.

The Image Recommender System was tested in two experiments. In the first experiment, we performed a series of searches and counted the number of relevant hits, with and without expanding the query. We performed searches for images of 38 European countries, and selected images from a specific category. In one condition, we only selected images that depicted natural scenes, whereas in a second condition, we only selected images that depicted cities. Semantic networks were formed based on these selections, and afterwards we searched for the same 38 European country names, but this time with expanding the query using the Image Recommender System. Searches were performed with queries that were expanded with one of the two most associated items. We did the same experiment with image searches for 14 pop band names, and selected images representing stage acts of these bands and album covers, respectively. We found that in all categories recommending a related key word based on the declarative memory user model increased the number of relevant images, as is depicted in Figure 1.

In the second experiment, we searched for the same word using two different semantic networks. We used the *Nature* and *Cities* networks for this test. Figure 2 shows the image results for the search query *picture*. As can be seen, the recommender system based on the *Nature* semantic network gives different results than the recommender system based on the *Cities* network. The *Nature* recommender system suggested the terms *Lofoten*, an archipelago near the Norwegian coast, and *Reine*, a small fishing village on one of the Lofoten islands. The *Cities* recommender system suggested *Nicosia* and *Nuernberg*, two European cities. Similar results were obtained for the key words *view*, *photo*, *country*, and *time*.

Because during the training of the semantic networks European countries were used as queries, it is not surprising that all recommended terms relate to Europe. However, because of the specific choices made when training the *Nature* and *Cities* semantic networks, the Image Recommender System, expands new queries differently for the users modeled by these declarative memory structures.

Discussion

An issue in our tests is the relatively small size of the declarative memories. Because the initial period in which the network of associations was trained was relatively short, the network size never exceeded 8,000 unique entries and no more than 30,000 words were parsed. Therefore, the system has not reached a stable configuration in which always appropriate recommendations can be made. It could be that some words are strongly associated, because at the web sites visited these words co-occur, although these web sites are not representative of the normal contexts of these

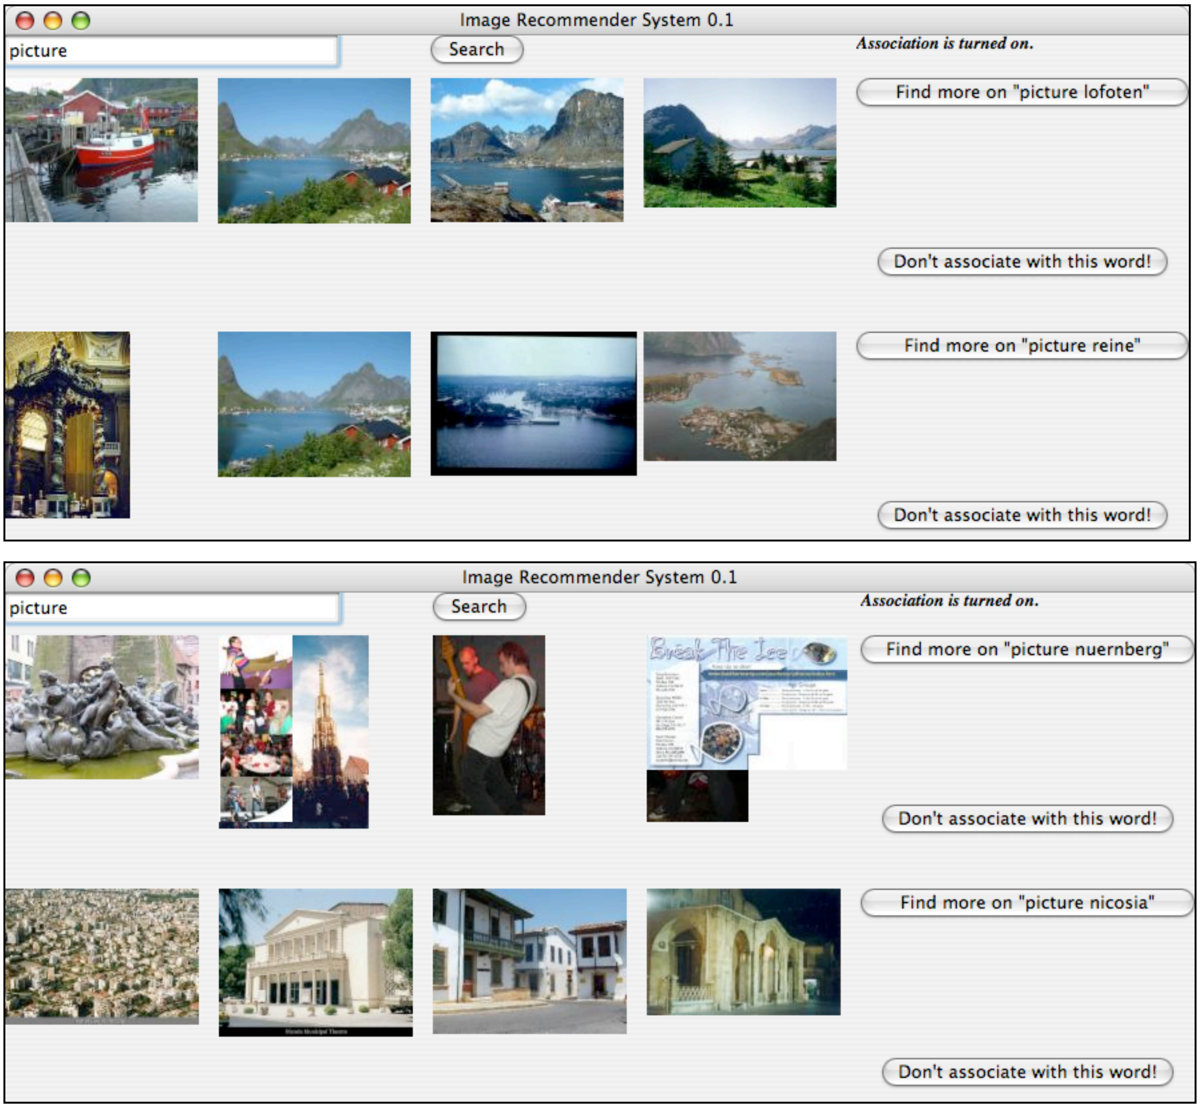

Figure 2. Image search for the key word *picture* using the *Nature* semantic network (top) and the *Cities* semantic network (bottom). The recommender system that uses the *Nature* network expands the query with the key words *Lofoten* and *Reine*, and mainly finds images with natural scenes. The recommender system that uses the *Cities* network expands the query with the key words *Nuernberg* and *Nicosia*, finding images of buildings.

words. In those cases, inappropriate recommendations will be made.

In addition, because of the limited network size, some words that are highly frequent will not be eliminated, but instead will be used for expanding the query. We expect that these issues will resolve if a larger training period is allowed.

In the Image Recommender System we developed, we only relied on strength of association for recommending possibly interesting chunks. The strengths of association can be regarded as reflecting the user's long-term interests, because the strengths of association only change slowly. The short-term interests of a user might be incorporated by including the base-level activation into the equation. If a chunk is recently attended, for instance because the word represented by that chunk has recently been used in a search query, the base-level activation of that chunk has been increased. An increased base-level activation means that the likelihood of being retrieved has also been increased. In this enhanced Image Recommender System, retrieval of the chunk will depend on the strengths of association – based on the long-term interests of the user – and on the base-level activation – reflecting the short-term interests of the user.

Conclusion

A dynamically updated declarative memory structure, consisting of a semantic network of chunks connected by strengths of association, might serve as a model of interest of an individual user. This model subsequently can be used to limit the amount of information presented to a user to a relevant subset. A typical domain of application is (of course) web search, but all situations that involve high information load (Brusilovsky & Tasso, 2004) might benefit from applying ACT-R's declarative memory principles to personalization research.

References

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2*(6), 396-408.

Brusilovsky, P., & Tasso, C. (2004). Special issue on user modeling for web information retrieval - preface. *User Modeling and User-Adapted Interaction, 14*(2-3), 147-157.

Claypool, M., Brown, D., Le, P., & Waseda, M. (2001). Inferring user interest. *IEEE Internet Computing, 5*(6), 32-39.

Van Maanen, L., Janssen, C., & Van Rijn, H. (2006). *Personalization of a virtual museum guide using eye-gaze.* Paper presented at the CogSci 2006, Vancouver, BC.

Yahoo! Search SDK. [software]: Yahoo! Inc.