Mediating expert knowledge and visitor interest in art work recommendation

Leendert van Maanen
Department of Artificial Intelligence, University of Groningen
Grote Kruisstraat 2/1, 9712 TS Groningen, the Netherlands
leendert@ai.rug.nl

Abstract

In this paper, we will present an outline for an online recommender system for art works. The system, termed Virtual Museum Guide, will take the interest that visitors of an online museum express into account in recommending suitable art works, as well as the relationships that exist between art works in the collection. To keep the Virtual Museum Guide similar to a human museum guide, we based its design on principles from research on human memory. This way, the Virtual Museum Guide can 'remember' which is the most suitable art work to present, based on its perception of the visitor's interests and its knowledge of the works of art.

1 Introduction

With the advent of online information presentation, cultural heritage institutions are starting to make their collections available online. Many museums already have websites displaying digital reproductions of part of their collection. Some of these online repositories are annotated, making it possible to search for specific art works: For example, the website of the Amsterdam Rijksmuseum in The Netherlands [Rijksmuseum, 2007] is driven by an ontology on art and artists.

With the online presentation of cultural heritage content, new issues arise. While one of the advantages of digitalization and online presentation is the greater accessibility of cultural heritage [e.g., because of better search capabilities, Van Ossenbruggen *et al.*, 2007], one of the drawbacks is that there is less control over what is presented to an individual visitor. Cultural heritage institutions have as one of their aims to educate people on history and culture, which becomes harder to realize once the contents of their collection is accessible from anywhere; They can no longer cater the individual interests of museum visitors while maintaining coherence in the presented information. Besides the decreased control that cultural heritage institutions experience, finding interesting art works in an online museum poses a problem. Just like in a real museum, most online museum visitors are not aware of their specific interests or of the exact contents of the museum's collection [Bell, 2002]. Instead, they only have a general impression of what they want to see and what is available. This makes it difficult to adjust the presentation of the art works to the visitors' personal interests.

Consider the example of a professional, educated museum guide, touring a party of interested visitors through a museum. The guide can (and has to) select information on the art works from her extensive knowledge that relates to the personal interests of the party, and can choose which art work to present next from the collection on display. To reproduce a similar personal experience in an online setting, personal interests as well as relationships between art works have to be known. A successful recommender system for the cultural heritage domain should incorporate both issues mentioned above: On the one hand, it should take care of the educational role of a cultural heritage institution, and on the other hand it should provide an enjoyable and personalized experience.

1.1 Overview

In this paper, we will present an online recommender system that presents art works from the Amsterdam Rijksmuseum collection. In our approach we will try to model the way a human museum guide will behave while touring a visitor through a museum. The assumption is that if the recommender system mimics the behavior of the museum guide, we will have a successful recommender system. In order to achieve this, we will ground the structure of the recommender system in cognitive theories on how human declarative memory works [Anderson *et al.*, 2004].

To stress the analogy with a museum guide touring a group of visitors through a museum, we termed the system the Virtual Museum Guide (VMG). The VMG combines the relationships that art works have to each other with the personal interests of the visitor to arrive a suitable art recommendations. We will first give an overview of the most important aspects of the system, and then discuss each aspect in more detail.

In the system we will present here, the art works presented online are accompanied by sets of key words that indicate what are the interesting aspects of the art work. As these key words are provided by the museum's art experts, expert knowledge on the art works and their interrelations are contained therein. We have applied statistical inference tools from natural language research [Landauer *et al.*, 1998] to infer how the art works relate to each other (details will be provided in the implementation section below). This way, all art works are related to each other with an association value indicating the relevance of one art work for another. This structure can be thought of as a *semantic* or *spreading activation network* [Collins and Loftus, 1975; Quillian, 1968].

Based on the visitor's feedback on presented art works, the guide generates hypotheses on the visitor's interest. For the system presented here we opted for the use of an explicit interest indicator using an *Interesting* and a *Not*

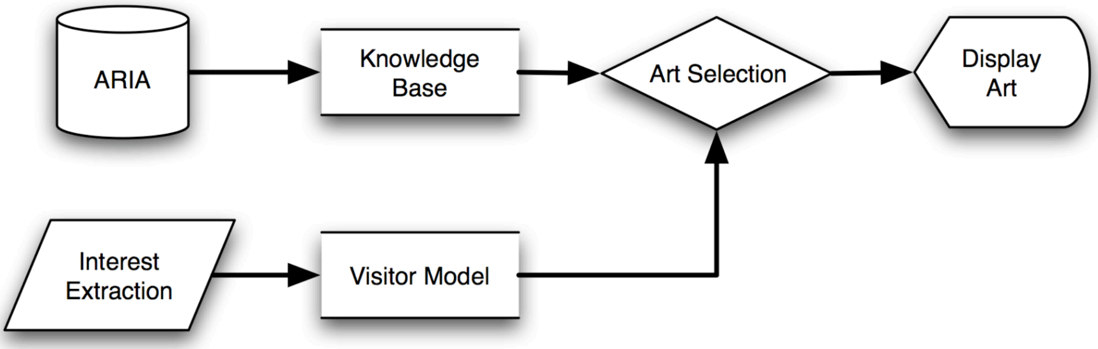

Figure 1. A flowchart of the Virtual Museum Guide.

interesting button. The interest hypotheses are represented as declarative facts that are stored in the VMG's memory.

Each time a user indicates interest in an art work by clicking one of the two interest-buttons, a new art work will be selected by computing the most relevant and interesting art work *given the current context*. First, the visitor's interest in the already visited art works will be assessed. Using a spreading activation algorithm (described in more detail below) a combined measure of interest and relevance will be computed.

2 ACT-R

The Virtual Museum Guide's memory is based on a formal theory of human cognition called ACT-R [Anderson *et al.*, 2004]. A major part of ACT-R is its model of declarative memory functioning [Anderson and Milson, 1989; Anderson and Schooler, 1991], and we will apply this approach in the context of a recommender system.

The key insight here is that human memory is optimally adapted to deal with information that has been presented in the past [Anderson and Milson, 1989; Anderson and Schooler, 1991]. Following this line of reasoning, the way information is represented in memory may also be optimal for storing a model of a person's interactions with information presented in the past.

Anderson and Schooler [1991] demonstrated that for each declarative fact stored in memory, the probability that that piece of information will be needed in the immediate future reflects the history of usage of that piece. That is, information that has been presented recently is more likely to be needed again than items that have been presented in the more distant past. Also, information that has been presented more frequently is more likely to be needed again. In ACT-R, the probability that information will be needed in the immediate future is represented by a quantity called *activation*. The declarative memory representation consists of small pieces of declarative knowledge, called chunks, that together represent a person's long-term memory. Each chunk has an activation value and associations with other chunks.

The two environmental observations (recency and frequency) have crystallized [Anderson *et al.*, 2004] into the following activation equation:

$$B_i = \ln\left[\sum_{j=1}^{n} \frac{1}{\sqrt{t_j}}\right] \qquad \text{(Equation 1)}$$

B_i represents the base-level activation of a chunk (indicated by the index i). The equation captures the effect of frequency of presentation by summing over multiple presentations, and the effect of recency of presentation by dividing by the square root of each presentation time lag (represented by t_j), that is, the time since the presentation of the chunk. This equation has been used in numerous studies predicting memory retrieval effects, both for theoretical purposes [e.g., Anderson *et al.*, 1998; Van Maanen and Van Rijn, *in press*] and for application-based research [e.g., Pirolli, 2005; Van Maanen *et al.*, 2006].

Besides the frequency and the recency with which memory facts are encountered, also the contexts in which they are encountered adds to their activation. This is determined by the likelihood that two facts have co-occurred in the past [Anderson and Milson, 1989]. The likelihood that one fact needs to be retrieved from memory is predicted by the recent retrieval from memory of another fact, and this prediction is based on how often it has been accurate in the past.

3 Virtual Museum Guide

Figure 1 presents a flowchart of the Virtual Museum Guide. We start out with extraction of a Resource Description Framework (RDF) specification of each art work from the online ARIA (Amsterdam Rijksmuseum Inter-Actief) repository, which can be inspected at http://media.cwi.nl/sesame/[1]. The RDF specification is transformed to an associative network structure, called the Knowledge Base. The Knowledge Base contains the knowledge the VMG has on the art works and their interrelations. Besides the knowledge on the art works, the VMG forms hypotheses on the interests of the visitors. These are extracted from the visitors' behavior and stored in a Visitor Model. Based on both knowledge sources, the VMG selects a suitable art work and displays it for the visitor, together with a little background information on the art work.

3.1 Knowledge Base

A human museum guide might present two similar art works right after each other, for instance because they are painted by the same artist. Therefore, the similarity be-

[1] To inspect the RDF repository, select *Topia's RDF Aria for Sesame* in the drop-down menu and slect one of the read actions. More information on how to query this repository can be found on openRDF.org [openRDF, 2007].

tween two art works might also represent the likelihood of two art works co-occurring. The similarity between art works will be represented by an associative network structure [Collins and Loftus, 1975; Quillian, 1968], in which the association values indicate similarity: The stronger the association between two art works, the more similar they are considered to be. The associative values in the associative network are based on the frequency statistics of the key words that occur in the RDF specifications of the art works. The general idea is that two sets of key words that greatly overlap might be considered similar to one another. This can be thought of as two feature-vectors that lie close together in a highly dimensional space. We applied Latent Semantic Analysis (LSA) [Deerwester *et al.*, 1990; Landauer and Dumais, 1997; Landauer *et al.*, 1998] on the scaled frequency vectors representing the art works. For this we used the standard TF-IDF weighting scheme [Salton and McGill, 1983], which scales the frequency of terms in a document by the number of documents in which the terms occur. In the VMG's knowledge base, this means that the frequency of the key words in all RDF specifications is taken into account.

At this point it is important to note that LSA is more than just a correlation of frequency counts [Deerwester *et al.*, 1990; Landauer and Dumais, 1997; Landauer *et al.*, 1998]. Instead, LSA depends on a mathematical analysis (singular value decomposition) that is capable of a higher-order inference. For example, let's assume that the specification of Rembrandt's *The Night Watch* contains the key word *claire-obscure,* and Gerard van Honthorst's *The Merry Fiddler* contains the key word *caravaggists*. LSA is capable of inferring that these two art works share a feature (namely, the technique used) if these key words co-occur in similar contexts. For example, a specification of Dirck van Baburen's *Prometheus Being Chained by Vulcan* might mention the key words *caravaggists, light, dark,* and *contrast*, and Rembrandt's *Ecce Homo* might contain the key words *light, dark, contrast*, and *claire-obscure*. In a sense, LSA estimates the likelihood that the word *claire-obscure* would occur in the specification of *The Merry Fiddler*, and the likelihood that *caravaggists* would occur in the specification of *The Night Watch*. For a more detailed, but still non-technical introduction to Latent Semantic Analysis, the reader is referred to Landauer, Foltz, and Laham [1998].

After constructing the semantic space using LSA, the similarity between two art works is computed by calculating the cosine between their representing vectors [Salton *et al.*, 1975]. The cosine between two feature-vectors represents their angle. This indicates how much they deviate in the semantic space, with small angles representing greater similarity than large angles.

3.2 Visitor Model

In the current setup of our virtual museum, visitors encounter each art work only once. Therefore, we simplified Equation 1 to

$$B_i = \ln\left[\frac{1}{\sqrt{t_i}}\right] \qquad \text{(Equation 2)}$$

in which t_i is the time stamp of the presentation of an art work (i). Note that there is no reason why this simplification will remain in future versions of the VMG, since both Equation 1 and Equation 2 provide an activation value. Figure 2 presents an example of the dynamics of two interest hypotheses in the current setup on the VMG.

For the purpose of designing a virtual museum guide we are not really interested in the probability that a piece of information will be needed, but rather in the likelihood that user interest in an art work extends to interest in similar art works. Thus, in this case, the base level activation represents the activation of the interest hypotheses that the VMG has about the visitor's interest in the already presented art works. Note that the activation of the interest hypotheses decays over time, indicating that the VMG forgets interests that have been expressed in the past. Although this seems as if the VMG incorporates a negative value of the human museum guide (forgetting), we believe that this is actually a positive effect. Museum visitors might switch the main topic of their interests, either because their knowledge on the art works has increased, or simply because they want to be presented with something different. The VMG can cope with this by forgetting hypotheses formed in the past.

A virtual museum visitor can express her interest in the art work that is currently being presented by clicking the *interest* button on the interface. Also, she can express disinterest by clicking a *disinterest* button. A positive interest expression will increase the likelihood that similar art is of interest for this visitor, whereas a negative interest expression will decrease the likelihood that similar art works are of interest. Therefore, the activation value of an art work is determined by positive spreading activation from art works in which interest is expressed, and negative spreading activation from art works in which disinterest is expressed.

3.3 Art Selection

The selection of art works depends on a weighted scheme of visitor interest. For each art work, the spreading activation from already visited art is computed. Art that is rated as uninteresting spreads negative activation; art that is rated as interesting spreads positive activation. Because of

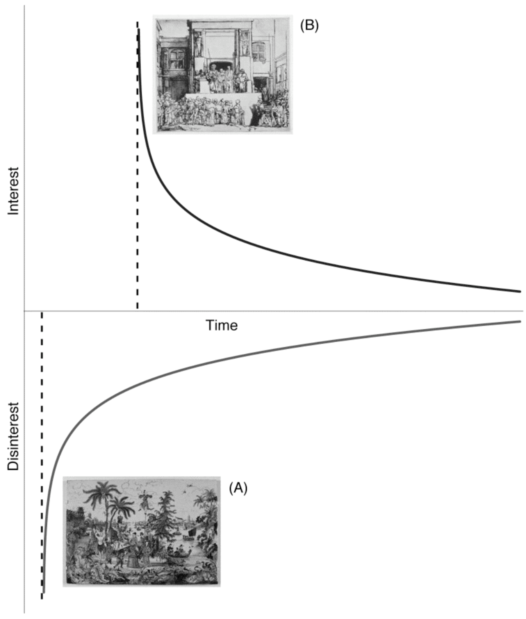

Figure 2. An example of the dynamics of the interest hypotheses. The visitor has expressed disinterest in art work A and interest in art work B. The activation of the hypothesis facts decays over time, decreasing the influence the hypotheses have on the art selection mechanism.

the inclusion of the recency component in the activation equation discussed above, the influence of recently presented art work is higher than the influence of art work presented longer ago. The spreading activation is scaled according to the similarity between art works. Thus, art works that are highly similar spread relatively more activation towards each other. These considerations result in the following equation (Equation 3), in which A_i represents the activation of a certain art work i, B_j represents the hypothesis on visitor interest in already presented art works (j), and S_{ji} represents the similarity between art works i and j.

$$A_i = \sum_j B_j S_{ji} \qquad \text{(Equation 3)}$$

This equation represents how suitable an art work will be to present to the current visitor given what the system knows from the visitor's interest, and the relations that exist between art works. Since B_j can be either a positive value or a negative value (depending on the VMG's hypothesis on the visitor's (dis)interests), art works that were considered uninteresting decrease the activation of related art works, while art works that were considered interesting increase the activation of related work. Thus, the resulting activation of an art work will be high if a visitor expressed interest in related art work, and did not expressed disinterest in related art work. Similarly, the activation will be low (that is, negative), if a visitor only expressed disinterest in related art work. The art work with the highest activation will be selected next for presentation.

After an art work has been selected for presentation, a web page will be generated that contains a digital reproduction of the art work under consideration and some information on the art work. These snippets of information are taken from the Rijksmuseum database, so it is ensured that the information is correct and relevant to the art work.

4 Discussion and Conclusion

This section will wrap up the paper by contrasting our approach to already existing tools for personalized information presentation in the museum domain. Also, we will indicate the future directions of our work.

4.1 Related Work

The key features of the VMG are the combination of the spreading activation network structure of the knowledge base combined with the decaying level of visitor interest. Also, the generation of the knowledge base using Latent Semantic Analysis is an important aspect, as well as the dynamic generation of web content.

Although most of these features have been applied in previous information presentation tools for the museum domain, the combination we apply is, to our knowledge, unique. Also, most other applications focus on the presentation aspects of dynamically generated content, especially in the context of a real, non-virtual, museum, where the mobility of the visitors poses specific challenges for the presentation of information [e.g., Hatala and Wakkary, 2005; Stock *et al.*, 2007; for a review see Raptis *et al.*, 2005]. A third obvious difference between related work and our approach is that while most applications focus on the personalized presentation of *background information* with an artefact, personalization in the VMG involves the selection of the museum artefacts themselves. In this section, we will discuss two systems that seem to be most similar to ours in the key features we have identified for the VMG. That is, both systems – ec(h)o [Hatala and Wakkary, 2005] and PEACH [Stock *et al.*, 2007] – are constructed around a conceptual network, in which selection of concepts is mediated by expressed visitor interests.

Similar to VMG, PEACH [Stock *et al.*, 2007] also adopts an activation based network. Since PEACH's main output modality is video, the nodes in the network represent video segments, and the edges represent semantic relations between these video segments. Interest expressed in one video segment propagates as activation through the network to all related other segments, and new information will be presented based on the activation values of all video segments. This seems to be a similar approach as the VMG deploys, although the level of semantic relatedness is less fine-grained, due to the Latent Semantic Analysis performed on the edges of the VMG associative network.

PEACH also differs from the VMG in the temporal aspects of the relevance feedback. Visitor's expressed interest in a video segment in PEACH does not extend to another artefact, but only applies to the current art work. Therefore, decay of visitor interest values is unneccesary. Since the VMG is intended for the dynamic selection of art works, visitor interest must extend to other art works.

Just like the VMG, ec(h)o [Hatala and Wakkary, 2005] uses a conceptual ontology as a knowledge base. In ec(h)o, the ontology is based on the Conceptual Reference Model [Crofts *et al.*, 2003] which is specifically developed for cultural heritage concepts. Selection of information is subsequently established by reasoning over the relationships in the ontology. ec(h)o also has a decay mechanism to ensure that more recent interests are more important than older ones. The mechanism implemented in ec(h)o is however not time-based (as is the decay mechanism of the VMG), but rather the interest values of concepts are normalized such that the highest value stays under a certain upper bound. An advantage of that approach could be that a longer visit to an art work does not result in 'forgetting' of interests, which is a side-effect of the way interest decay is modeled in the VMG.

The ec(h)o system differs from the VMG and PEACH in the way relevance feedback can be expressed. Were VMG and PEACH adopt an explicit strategy in which interest as well as disinterest can be expressed, ec(h)o presents the user with three small audio snippets, from which the visitor can choose. The assumption is that the visitor chooses the audio fragment that is the most interesting to her. As a result of this design choice, visitors cannot express disinterest. Moreover, they have to base their decision on a small snippet of the actual information, and could well change their minds after they hear all the information. In this sense, ec(h)o does not really incorporate a relevance feedback mechanism.

4.2 Conclusion and Future Work

An analysis of the selection of art works provided to the author suggests that the VMG is capable of recommending art works that relate to those recently indicated as interesting. Also, the VMG's knowledge base correctly relates items that seem similar upon visual inspection of the art work and the description. However, since the author is no art specialist, we are planning two evaluation studies. First, we are planning a study in which art experts can assess the relationships between art works that are present in the VMG's knowledge base. Secondly, a user study to measure how visitors of the online museum respond to personalized recommendations of art works is currently conducted. This user study can be visited at the website of our institute: www.ai.rug.nl/cogmod [AI, 2007].

In the current setup, we opted to generate interest hypothesis by having visitors press one of two interest-buttons. However, the way visitor feedback is provided can be very diverse, ranging from the simple button press to more unobtrusive methods, including the time spent observing the art work [e.g., Claypool *et al.*, 2001] or even eye gaze analysis [e.g., Van Maanen *et al., submitted*]. In a future version of the Virtual Museum Guide, we plan to incorporate less obtrusive methods to infer visitor interest. This will also include a more gradient sense of the likelihood that a visitor is interested. We will implement this by adding a parameter to Equation 2 that can control the impact of each presentation of information.

Using principles from cognitive science, we were able to implement a working system that recommends art works, based on both visitor interests and expert knowledge on the relations between the art works. In the context of a museum, both aspects are important. Because of the educational role of museums, recommending art work is more than mapping visitor interest on the museum's collection. The museum needs to ensure that the resulting sequence of art works is coherent and transfers (part of) the museum's message. It seems that the Virtual Museum Guide ensures both aspects in art work recommendation.

Acknowledgements

The art works in Figure 2 are from the collection Rijksmuseum Amsterdam. This research is supported by the Netherlands Organisation for Scientific Resarch (ToKeN/I^{2}RP project grant no. 634.000.002). Hedderik van Rijn, Chris Janssen, and three anonymous reviewers are acknowledged for valuable comments on earlier versions of this paper.

References

[AI, 2007] Cognitive Modeling, University of Groningen (www.ai.rug.nl/cogmod). Retrieved 4-7-2007.

[Anderson *et al.*, 2004] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y. An integrated theory of the mind. *Psychological Review, 111*(4), 1036-1060, 2004.

[Anderson *et al.*, 1998] Anderson, J. R., Bothell, D., Lebiere, C., and Matessa, M. An integrated theory of list memory. *Journal of Memory and Language, 38*(4), 341-380, 1998.

[Anderson and Milson, 1989] Anderson, J. R., and Milson, R. Human memory: An adaptive perspective. *Psychological Review, 96*(4), 703-719, 1989.

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

[Bell, 2002] Bell, G. *Making sense of museums: The museum as 'cultural ecology'*: Intel Labs, 2002.

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

[Collins and Loftus, 1975] Collins, A. M., and Loftus, E. F. A spreading activation theory of semantic processing. *Psychological Review, 82*(6), 407-428, 1975.

[Crofts *et al.*, 2007] Crofts, R., Doerr, M, and Gill. T. The CIDOC conceptual reference model: A standard for communicating cultural contents. *Cultivate Interactive, 9*, 2003.

[Deerwester *et al.*, 1990] Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. Indexing by latent semantic analysis. *Journal of the American Society for Information Science, 41*(6), 391-407, 1990.

[Hatala and Wakkary, 2005] Hatala, M., and Wakkary, R. Ontology-based user modeling in an augmented audio reality system for museums. *User Modeling and User-Adapted Interaction, 15*(3-4), 339-380, 2005.

[Landauer and Dumais, 1997] Landauer, T. K., and Dumais, S. T. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*(2), 211-240, 1997.

[Landauer *et al.*, 1998] Landauer, T. K., Foltz, P. W., and Laham, D. An introduction to latent semantic analysis. *Discourse Processes, 25*(2-3), 259-284, 1998.

[openRDF, 2007] openRDF.org (www.openrdf.org). Retrieved 29-8-2007.

[Pirolli, 2005] Pirolli, P. Rational analyses of information foraging on the web. *Cognitive Science, 29*(3), 343-373, 2005.

[Quillian, 1968] Quillian, M. R. Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 216-270). MIT Press, Cambridge, MA, 1968.

[Raptis *et al.*, 2005] Raptis, D., Tselios, N., and Avouris, N. Context-based design of mobile applications for museums: A survey of existing practices, *Mobile-HCI'05* (pp. 153-160). Salzburg, Austria: ACM.

[Rijksmuseum, 2007] Rijksmuseum Amsterdam, national museum for art and history (www.rijksmuseum.nl). Retrieved 4-7-2007.

[Salton and McGill, 1983] Salton, G., and McGill, M. *Introduction to modern information retrieval*. McGraw-Hill, New York, 1983.

[Salton *et al.*, 1975] Salton, G., Wong, A., and Yang, C. S. Vector-space model for automatic indexing. *Communications of the ACM, 18*(11), 613-620, 1975.

[Sesame, 2007] Sesame@media.cwi.nl, public repositories (http://media.cwi.nl/sesame/). Retrieved 21-6-2007.

[Stock *et al.*, 2007] Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Kruger, A., Kruppa, M., et al. Adaptive, intelligent presentation of information for the museum visitor in peach. *User Modeling and User-Adapted Interaction, 17*(3), 257-304, 2007.

[Van Maanen *et al.*, 2006] Van Maanen, L., Borst, J. P., Janssen, C. P., and Van Rijn, H. Memory structures as user models, *13th Annual ACT-R Workshop*. Pittsburgh, PA, 2006.

[Van Maanen and Van Rijn, *in press*] Van Maanen, L., and Van Rijn, H. An accumulator model of semantic interference. *Cognitive Systems Research,* in press.

[Van Maanen *et al.*, *submitted*] Van Maanen, L., Van Rijn, H., and Janssen, C. P. Eye-gaze based interest awareness for adaptive multimedia art presentations, submitted.

[Van Ossenbruggen *et al*., 2007] Van Ossenbruggen, J., Amin A., Hardman L., Hildebrand M., Van Assem M., Omelayenko B., Schreiber G., Tordai A., De Boer, V., Wielinga B., Wielemaker J., De Niet, M., Taekema J., Van Orsouw, M.-F., and Teesing A.. Searching and Annotating Virtual Heritage Collections with Semantic-Web Techniques. In *Proceedings of Museums and the Web 2007*, San Francisco, CA, March, 2007.