Implicit Search Trails for Video Recommendation

*Martin Halvey *Frank Hopfgartner *†David Vallet *Joemon Jose

*Department of Computing Science,
University of Glasgow,
Glasgow, Scotland, United Kingdom.

†Universidad Autónoma de Madrid,
Escuela Politécnica Superior Ciudad Universitaria de Cantoblanco, 28049 Madrid, Spain.

{ halvey, hopfgarf, jj} @ dcs.gla.ac.uk, david.vallet@uam.es

ABSTRACT
In this demo paper we demonstrate our approach and system for using implicit actions involved in video search to provide recommendations to users. The goal of this system is to improve the quality of the results that users find, and in doing so, help users to explore a large and difficult information space and help them consider search options that they may not have considered otherwise. Results of a user evaluation show that this approach achieves all of these goals.

Categories and Subject Descriptors
H.5.1 Multimedia Information Systems, H.5.3 Group and Organization Interfaces

General Terms
Human Factors.

Keywords
Video, search, collaborative, community, feedback, recommender, user studies.

INTRODUCTION
People now build their own digital libraries from materials created through digital cameras and camcorders, and use a number of systems to place this material on the web, as well as store them as their own personal collection. However, the systems that currently exist to organise and retrieve theses videos are not sufficient to deal with large and rapidly growing volumes of video. Current state of the art systems rely on using annotations provided by users, methods that use the low level features available in the videos or on an existing representation of concepts associated with the retrieval tasks. None of these methods are sufficient to overcome the problems associated with video search. In order to alleviate some of these problems we have developed a video retrieval system that uses the implicit actions involved in previous user searches to help and inform future users of the system, through recommendations. Our system does not require users to alter their normal searching behaviour, provide annotations or provide any other supplementary feedback. Furthermore, we are able to exploit the rich spectrum of actions that can be collected from a video retrieval system. We achieve this outcome by utilising the available information of past user interactions such as launching a query, playing a video or navigating through it, to name a few. The user actions are represented as links of a graph where nodes are the user queries and interacted documents. We exploit this implicit feedback representation in order to provide recommendations during a search or browsing process . The use of this system can result in a number of desirable outcomes for users. In particular, it helps improving the user search performance; it can aid user exploration of the collection and can also increase user satisfaction with their search and their search results.

APPROACH
For the implementation of our recommendation model based on user actions we have opted for a graph based model. There are two main desired properties that we require for this graph of implicit user actions. The first property is the representation of all of the user interactions with the system. This allows us to fully exploit all of the interactions to provide rich recommendations. The second property is the aggregation of implicit information from multiple sessions and users into a single representation, thus facilitating the analysis and exploitation of past implicit information. To achieve these properties we draw inspiration from previous graph based models implemented by White et al and Craswell and Szummer . We take the concept of trails from White et al.; however unlike White et al. we do not limit the possible recommended documents to those documents that are at the end of the search trail. This is because we believe that during an interactive search the documents that most of the users with similar interaction sequences interacted with are the documents that could be most relevant for recommendation, not just the final document in the search trail. Thus a major difference between our search trail and that of White et al. is that ours is a more complex representation. Similar to Craswell and Szummer , our approach represents queries and documents in the same graph, however unlike their approach, where the clicked documents are linked directly to the query node; we represent the whole interaction sequence. The approach from Crasswell and Szummer does not represent search trails; their approach is based on finding correlations query-clicked document. We use search trails because we want to recommend potentially important documents that are part of the interaction sequence and not just at the end of interaction sequences. Another difference between our approach and previous work is that we take into consideration other types of implicit feedback actions, related to multimedia search, e.g. length of play time, browsing keyframes etc., as well as click through data. This additional data allows us to provide a richer representation of user actions and potentially better recommendations. Overall our representation exploits a greater range of user interactions in comparison with other approaches ,

this results in a more full representation of a wide range of user actions that may facilitate better recommendations. In addition while these other approaches , have been successful in other domains they have not been applied to video search. A full description of our approach can be found in Vallet et al . In the following section we describe the implementation of this approach and our interface for presenting the recommendations.

INTERFACE

A screenshot of our system is shown in Figure 1 and can be divided into three main panels. The search panel (A) is where users formulate and carry out their searches. Users enter a text based query in the search panel (A) to begin their search. The users are presented with text based recommendations for search queries that they can use to enhance their search (b). The users are also presented with recommendations of video shots that might match their search criteria (a). The result panel is where users can view the search results (B). Users can mark results as being relevant or irrelevant by using a sliding bar (c). In the result panel additional information about each video shot can be retrieved. Hovering the mouse tip over a video keyframe, will result in that keyframe being highlighted, along with neighbouring keyframes and any text associated with the highlighted keyframe (d). The playback panel (C) is for viewing video shots (g). As a video is playing it is possible to view the current keyframe for that shot (e), any text associated with that keyframe (f) and the neighbouring keyframes. Users can interact with a video as they can on a normal media player, and also make relevance judgements about the keyframe (h).

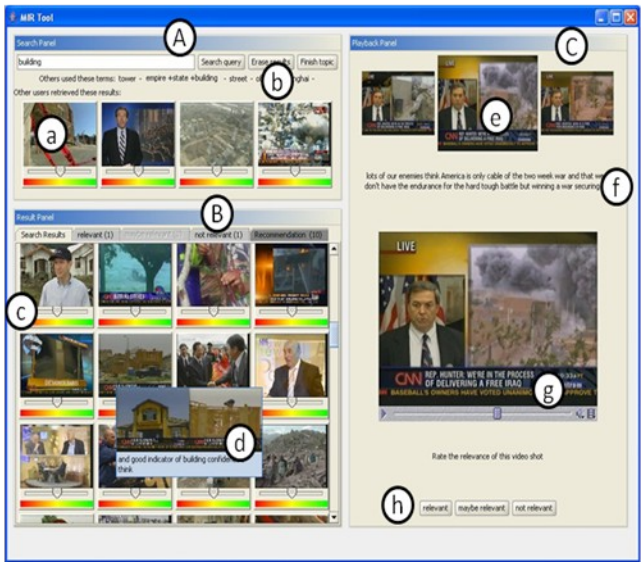

Figure 1: Interface for video retrieval system

EVALUATION AND CONCLUSION

In order to measure the effectiveness of our proposed approach we conducted a user-centred evaluation. The recommendation system was compared with a baseline system by 24 users, each user carried out the 4 most difficult topics (i.e. the topics with the lowest MAP) for the TRECVID 2006 evaluation. Each user carried out 2 topics with the baseline system and 2 topics with the recommender system. Full details of the experimental set up can be found in Hopfgartner et al. .

There are a number of conclusions that can be made about our system based on the experimental evaluation. To measure task performance we measured P@N and MAP values, it was found that the results retrieved by users using the recommendation system outperforms the baseline system, and that this difference is statistically significant. This demonstrates that the performance of users of the recommendation system will improve with the use of recommendations based on implicit feedback. From analysing user logs it was found that different users were able to explore the collection to a greater extent and find more relevant videos using the recommendation system. This finding was further validated in by our analysis of user questionnaires where the users gave an indication that the recommendation system helped them to explore the collection. The participants indicated in their post task questionnaires that the system that provided recommendations helped them to explore the task and find aspects of the task that they otherwise would not have considered, in comparison with the baseline system. It was also found in these questionnaires that the users have a definite preference for the recommendation system.

Follow up simulation based evaluations found that our approach is scalable. In fact the relevance pool consistently gained performance as more users were added to the implicit graph . In conclusion, the results of the evaluation, for our system have highlighted the promise of this approach to alleviate the major problems that users have while searching for multimedia.

ACKNOWLEDGMENTS

This research work was partially supported by the EC under contracts: K-Space (FP6-027026) and SEMEDIA (FP6-045032).

REFERENCES

Craswell, N. and Szummer, M., Random walks on the click graph. In Proc. SIGIR 2007, ACM Press (2007), 239-246.

Hopfgartner, F., Vallet, D., Halvey, M., and Jose, J. Search Trails using User Feedback to Improve Video Search. In Proc ACM MM 2008.

Vallet, D., Hopfgartner, F., and Jose, J. Use of Implicit Graph for Recommending Relevant Videos: A Simulated Evaluation. In Proc ECIR 2008

White, R., Bilenko, M. and Cucerzan, S., Studying the use of popular destinations to enhance web search interaction. In Proc. SIGIR 2007, ACM Press (2007), 159-166.