Chapter 13

Proactive information retrieval

Samuel Kaski, Jarkko Salojärvi, Eerika Savia, and Kai Puolamäki

13.1 Introduction

Successful proactivity, i.e. anticipation, in varying contexts requires generalization from past experience. Generalization, on its part, requires suitable powerful (stochastic) models and a collection of data about relevant past history to learn the models.

The goal of the PRIMA (Proactive Information Retrieval by Adaptive Models of Users' Attention and Interests) project is to build statistical machine learning models that learn from the actions of people to model their intentions and actions. The models are used for disambiguating the users' vague commands and anticipating their actions.

Our application area is information retrieval, where we investigate to what extent the laborious explicit relevance feedback can be complemented or even replaced by implicit feedback derived from patterns of eye fixations and movements that exhibit both voluntary and involuntary signs of users intentions. Inference is supported by models of document collections and interest patterns of users.

PRIMA is a consortium with Complex Systems Computation Group, Helsinki Institute for Information Technology (Prof. Petri Myllymäki), and Center for Knowledge and Innovation Research (CKIR), Helsinki School of Economics (Doc. Ilpo Kojo). The project lasted for 2003 – 2005.

13.2 Implicit relevance feedback from eye movements

A promising new source of implicit feedback is eye movements measured during reading. During complex tasks such as reading, attention approximately lies on the location of the reader's gaze. Therefore the eye movements should contain information on the reader's interests. Deriving interest from a reading pattern is difficult however, since the signal is complex and very noisy, and interestingness or relevance is higly subjective and thus hard to define.

In a first feasibility study we constructed a controlled experimental setting in which it is known which documents are relevant [1]. The user was instructed to find an answer to a specific question, and was then shown a set of ten document titles (Fig. 13.1), of which four were known to be relevant (i.e. thay handled the same topic as the question), and one was the correct answer. Eye movements were recorded while the users were seeking the answer. Based on data gathered from eleven test subjects we then learned models that predict relevance using the measured eye movement patterns. The study was the first evidence (with statistical significance) that inferring relevance is possible [2].

The data was used in a Pascal NoE challenge "Inferring Relevance from Eye Movements", during March 2005 – October 2005. The challenge was organized in the form of a competition where the winner was the team that could most accurately predict relevance for a left-out test data set. The results of the challenge were reported in a workshop at Neural Information Processing Systems (NIPS) in December 2005.

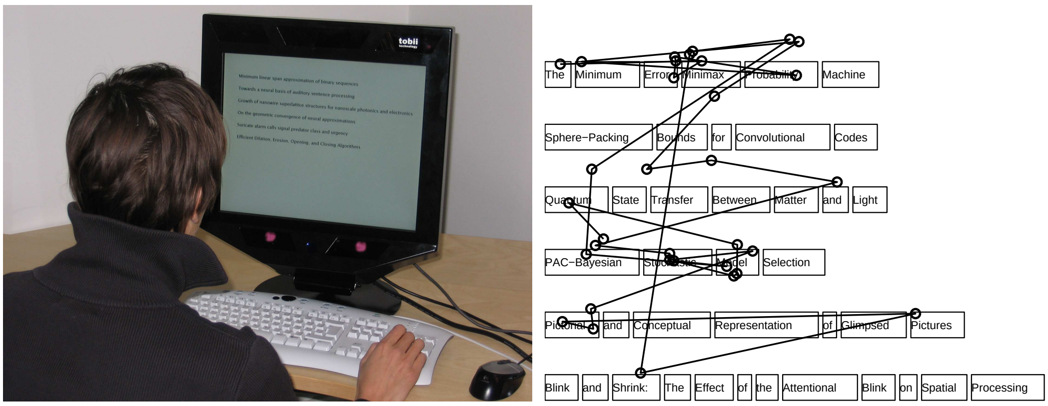

Figure 13.1: The experimental setup. Left: The eye movements of the user are being tracked with Tobii 1750 eye tracker. The tracker consists of an infra-red light source and a camera integrated into the monitor frame. Right: An example of an eye movement pattern during reading, plotted on the assignment. Lines connect successive fixations, denoted by circles (Matlab reconstruction). Each line contains one document title, and some of the titles are known to be relevant.

References

[1] Jarkko Salojärvi, Ilpo Kojo, Jaana Simola, and Samuel Kaski. Can relevance be inferred from eye movements in information retrieval? In *Proceedings of WSOM'03, Workshop on Self-Organizing Maps*, pages 261–266. Kitakyushu, Japan, 2003.

[2] Jarkko Salojärvi, Kai Puolamäki, and Samuel Kaski. Implicit relevance feedback from eye movements. In *Artificial Neural Networks: Biological Inspirations – ICANN 2005*, Lecture Notes in Computer Science 3696, pages 513–518, Berlin, Germany, 2005.

13.3 Collaborative Filtering: Inferring user interests from other available sources

Traditionally, user preferences have been predicted using so-called collaborative filtering methods, where the predictions are based on the opinions of similar-minded users. Collaborative filtering is needed when the task is to make personalized predictions but there is not yet sufficient amount of data about the user's personal interests [2]. Then the only possibility is to generalize over users, for instance by grouping them into like-minded user groups. The early methods were memory-based; predictions were made by identifying a set of similar users, and using their preferences fetched from memory. Model-based approaches are justified by the exponentially increasing time and memory requirements of the memory-based techniques. Recent work includes probabilistic and information-theoretic models. An interesting family of models are the latent component models, which have been successfully used in collaborative filtering. In these models, each user is assumed to belong to one or many latent user groups that explain her preferences. We went one step further and introduced a similar latent structure for the documents as well.

As a collaborative filtering system has to rely on the past experiences of the users, it will have problems when assessing new documents not seen yet by most of the users. To tackle this problem we have introduced a novel latent grouping model for predicting the relevance of a new document to a user [1]. The model assumes a latent group structure for both users and documents. See Figure 13.2.

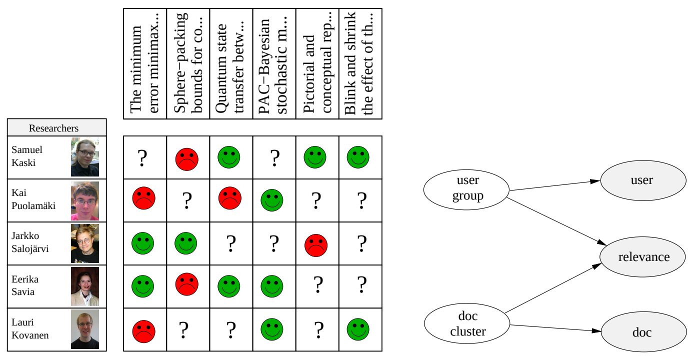

Figure 13.2: Left: An example of gathered data for collaborative filtering. Right: The model generalizes both over users and documents.

We compared the model against a state-of-the-art method, the User Rating Profile model, where only users have a latent group structure. We estimate both models by Gibbs sampling. The new method predicts relevance more accurately for new documents that have few known ratings. The reason is that generalization over documents then becomes necessary and hence the two-way grouping is profitable.

References

[1] E. Savia, K. Puolamäki, J. Sinkkonen, and S. Kaski. Two-way latent grouping model for user preference prediction. In F. Bacchus and T. Jaakkola, editors, *Uncertainty in Artificial Intelligence 21*, pages 518–525. AUAI Press, Corvallis, Oregon, 2005.

[2] Eerika Savia, Samuel Kaski, Ville Tuulos, and Petri Myllymäki. On text-based estimation of document relevance. In *Proc. IJCNN'04*, pages 3275–3280, 2004.

13.4 Application: Proactive information retrieval

We have studied a new task, proactive information retrieval by combining implicit relevance feedback and collaborative filtering [1]. We constructed a controlled experimental setting, in which the users tried to find interesting scientific articles by browsing their titles. The experimental setup was designed to resemble closely a real-world information retrieval scenario, where the user browses the output of, e.g., a web search engine in an attempt to find interesting documents. The test subjects rated the relevance of titles of scientific articles and eye movements were measured from a subset of the test subjects.

Implicit feedback was inferred from eye movement signals, with discriminative hidden Markov models estimated from existing data in which explicit relevance feedback was available. It produced a reasonable, though rather noisy, prediction of document relevance based on eye movement measurements.

We complemented the eye movement based prediction with a probabilistic collaborative filtering model that produced a quite robust document relevance prediction. The model was computed using Markov Chain Monte Carlo techniques.

In our scenario a database of user preferences was combined with the measured implicit relevance feedback, resulting in more accurate relevance predictions. We introduced a probabilistic mixture model [1, 2] that can be used to combine the predictions. The mixture model clearly outperformed a simple linear method and was found necessary for making use of several information sources, the quality of which varied. The best prediction accuracy still leaves room for improvement but shows that proactive information retrieval and combination of many sources of relevance feedback is feasible.

Our work provides the next step towards proactive information retrieval systems. A future extension is to supplement or replace the eye movements by other sources of implicit feedback.

References

[1] Kai Puolamäki, Jarkko Salojärvi, Eerika Savia, Jaana Simola, and Samuel Kaski. Combining eye movements and collaborative filtering for proactive information retrieval. In G. Marchionini, A. Moffat, J. Tait, R. Baeza-Yates, and N. Ziviani, editors, *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 146–153. ACM press, New York, NY, USA, 2005.

[2] Jarkko Salojärvi, Kai Puolamäki, and Samuel Kaski. On discriminative joint density modeling. In João Gama, Rui Camacho, Pavel Brazdil, Alípio Jorge, and Luís Torgo, editors, *Machine Learning: ECML 2005*, Lecture Notes in Artificial Intellligence 3720, pages 341–352, Berlin, Germany, 2005. Springer-Verlag.