Information Retrieval by Inferring Implicit Queries from Eye Movements

David R. Hardoon, John Shawe-Taylor
Computer Science Dept.
University College London
Gower Street, London
UK, WC1E 6BT

Antti Ajanki, Kai Puolamäki, Samuel Kaski
Helsinki Institute for Information Technology
Laboratory of Computer and Information Science
Helsinki University of Technology
P.O. Box 5400, FI-02015 TKK, Finland

Abstract

We introduce a new search strategy, in which the information retrieval (IR) query is inferred from eye movements measured when the user is reading text during an IR task. In training phase, we know the users' interest. We learn a predictor which links eye movements related to a term to the role of that term in the query. Assuming this predictor is universal with respect to the users' interests, it can also be applied to infer the implicit query when we have no prior knowledge of the users' interests. The result of an empirical study is that it is possible to learn the implicit query from a small set of read documents, such that relevance predictions for a large set of unseen documents are ranked significantly better than by random guessing.

1 INTRODUCTION

Current information retrieval (IR) systems rely mostly on explicit, typed queries, combined with explicit feedback telling the system which of the search results were relevant. The relevance feedback is used to refine the query, and the search converges iteratively towards more relevant documents. The standard web search engines are simplified versions of this scheme; they take advantage of the large scale which allows inferring general interest of documents from data.

A main disadvantage of the traditional IR paradigm is that formulating good textual queries is a difficult and challenging task, even for experienced users (Turpin & Scholer, 2006). The task is made even more difficult by the fact that the true interests of the users are often ambiguous even for the users themselves, and therefore the query may be difficult to formulate explicitly. It would be ideal if the system could infer the interests of the users while they work, and then have some suggestions readily available when the users ask for help. We call this task *proactive information retrieval*.

A proactive information retrieval system would additionally solve the problem that giving explicit feedback is laborious. Such a system would use *implicit feedback* to infer relevance. Click-stream data is one form of implicit feedback. While useful and often readily available, it offers limited information about a users interests and intentions (Joachims, Granka, Pan, Hembrooke, & Gay, 2005; Kelly & Teevan, 2003).

In this work, we propose to use implicit feedback from the observation of eye movements as an alternative source to infer the users' intentions. We combine the eye movements with the textual content of the documents in a novel way: *we use the eye movements to formulate an IR query*, which is then used to rank unseen documents with respect to their relevance to the current interests of the user.

While the traditional IR task is to find relevant documents given a query typed in by the user, we address the following tasks: (i) Construct a query from eye movements alone. (ii) Construct a query by combining information from implicit relevance feedback from eye movements *and* explicit relevance feedback.

We devise a controlled experimental setting, in which test subjects read through short text snippets searching for documents related to a given topic. During the reading the users' eye movements are recorded with an eye tracking system. We extract term-specific eye movement features for each document the user reads which are then used to predict importance of the term for the search task. The set-up is an extended version of earlier work (Puolamäki, Salojärvi, Savia, Simola, & Kaski, 2005).

We assume that there is a link between relevance or interest and eye movements, and that this link can be learned by observing the users' behaviour in search tasks where the ground truth (the true interest of the user) is known. Our main assumption is that this link between interest and eye movements is, to a reasonable extent at least, independent of the actual topic. Hence, we propose to construct an IR query for a previously unseen topic solely by observing the eye movements of the user during an IR task finding documents on that topic.

This work is a feasibility study on whether there is any truth in this assumption, and whether it would be realistic to use eye movements to formulate queries in information retrieval. We assume that no explicit relevance feedback is available.

As another case study, we investigate whether combining the eye movements with document features would help in the standard IR task where explicit relevance feedback is available for a subset of the documents.

More formally, we work with a bag-of-words (BOW) representation of the documents. We use *term-specific eye movement features*, denoted collectively by $\mathbf{e}_t$ for term t, to predict the parameters of our query, denoted collectively by $\mathbf{w}$, as $w_t = f_\lambda(\mathbf{e}_t, \mathbf{s}_t)$, where $\mathbf{s}_t$ are possible query-independent parameters associated with the term t (e.g., document frequency of term t). The $f_\lambda()$ could in principle be any predictor with *term-independent* parameters specified by λ. We then use a query function $g_\mathbf{w}(d)$, where d is a BOW representation of a document, to rank unseen documents with respect to their predicted interest to the user.

The parameters λ of the predictor $f_\lambda()$ are learned in the training phase, where we know the ground truth (true interest of the user). Note that following our assumptions laid out above, we assume that the functional form of the predictor $f_\lambda()$ is independent of the actual interest of the user. Hence, we can use it to formulate query parameters $\mathbf{w}$ for previously unseen topics.

In the following we choose the query function $g_\mathbf{w}()$ to be a Support Vector Machine (SVM) based model, with term-specific parameters $\mathbf{w}$, as we believe them to be most suitable for the proposed task. As a predictor $f_\lambda()$, which gives the parameters of the SVM model, we use some standard linear and non-linear regressors. As this is the first study of this kind, and we aim for robustness, we purposely use standard state-of-the-art machine learning methodologies. Developing methods tailored for the task is left for future work.

The paper is laid out as follows; In Section 2 we discuss the usage of eye movements in information retrieval as well as previous work. Section 3 provides a description of the data and how it was acquired. We discuss our proposed models in Section 4. In Section 5 we elaborate on the experiment set up while in Section 6 the results are given. Our final remarks and discussion are given in Section 7.

2 EYE MOVEMENTS IN INFORMATION RETRIEVAL

Use of eye movements in IR is a relatively new approach. Maglio, Barrett, Campbell, and Selker (2000); Maglio and Campbell (2003) introduced a prototype attentive agent application which monitors eye movements while the user views web pages, in order to determine whether the user is reading or just browsing. If reading is detected, more information of the topic is sought and displayed. The feasibility of the application was not however experimentally verified.

The eye movements were first used in an information retrieval task in (Salojärvi, Kojo, Simola, & Kaski, 2003; Salojärvi, Puolamäki, & Kaski, 2005). Discriminative hidden Markov models were applied to estimate the relevance of lines of read text, and the performance of the method was verified in a controlled experiment. A competition was subsequently set up, where the participants competed in predicting relevance based on the eye movements (Puolamäki & Kaski, 2006).

(Puolamäki et al., 2005) introduced a prototype information retrieval system, which used relevance information combined with collaborative filtering to seed out relevant scientific articles. This earlier prototype did not use the textual content of the documents at all.

3 DATA DESCRIPTION

The data set consists of about 500 Wikipedia documents from 25 different categories, which were used as search topics. The training corpus is as depicted in Table 1. The Wikipedia documents were truncated to fit to the screen (11 lines, about a dozen words on each). We made sure that the contents of each truncated document was sufficient for inferring its topic by manually inspecting all documents. For testing the accuracy we have another data set of 244 non-truncated documents, 10 documents for each category, except for *Natural disasters* which has only 4 documents. The documents in the testing corpus were not used in any way during the training of the model. For the BOW representation, we define the dictionary to be the set of stemmed words occurring in the training corpus. Some frequent 'stop' words like 'of' and 'the' are omitted from the dictionary.

In the experiments users were shown ten documents and they were asked, after they had viewed each document, to identify whether the document was rele-

Table 1: Summary of the Training Corpus.

Search Topic	Number of Documents
Astronomy	23
Ball games	23
Cities	13
Court systems	23
Dinosaurs	17
Education	22
Elections	22
Family	18
Film	21
Government	21
Internet	23
Languages	22
Literature	23
Music	16
Natural disasters	21
Olympics	22
Optical devices	23
Postal system	23
Printing	23
Sculpture	20
Space exploration	23
Speeches	23
Television	23
Transportation	23
Writing systems	17

vant to a search topic they had been given beforehand (Fig. 1). Once the users had assessed the relevance of the document, they pressed any button, after which they reported the relevance. The next document showed up immediately. On average, half of the documents shown during each topic search were relevant. Each user repeated the task ten times, with different search topic and documents in each round. The user group consisted of ten post-graduate and senior researchers.

While the users were reading, their eye gaze position on the screen was recorded with a Tobii 1750 eye tracker. Tobii measures eye positions 50 times per second by illuminating both eyes with infrared LED and measuring the reflected light. The system is fairly robust to head movements. The eye tracker was calibrated in the beginning of the experiment for each user.

The gaze direction is an indicator of the focus of attention, since accurate viewing is possible only in the central fovea area (1–2 degrees of visual angle). The correspondence is not one-to-one, however, since the attention can be shifted without moving the eyes. The eye movement trajectory is traditionally divided into fixations, during which the eye is fairly motionless, and saccades, rapid eye movements from one fixation to another.

We extracted 21 feature vectors, described in (Salojärvi, Puolamäki, Simola, et al., 2005), from the recorded eye movement data for each fixation. Fixation points were found using the default values of parameters recommended for text-only stimuli by Tobii: if the recorded gaze points stayed inside a 20 pixel area (about 0.5 degrees of visual angle) for at least 40 ms, they were considered as one fixation. Every fixation was mapped to the closest word, unless the fixation occurred well outside any text, in which case it was discarded.

4 MODELS USED

Our main task is to formulate an IR query, using the eye movements as the only feedback signal. The query need not however be understandable by humans; in fact, it suffices to formulate the query in such a way that it can be used by a relevance predictor to predict relevance for new documents.

What we will do is to create a *parameter vector* for a relevance predictor, using eye movements as inputs. For the training data, we need to know the correct solution: for this purpose we use the parameter vector, referred to as the *ideal weights*, of a predictor trained to classify according to known relevance labels, which are available for the learning data. We need to estimate the ideal parameter vector and to construct a *regressor* that tries to predict the ideal weights given the eye movements.

Furthermore, we study with a straightforward *combination of the text and eye movements* whether the eye movements help in the relevance prediction task when explicit relevance feedback is also available.

4.1 SUPPORT VECTOR MACHINES

We use support vector machines (SVM) for two tasks; to compute *ideal weights* and to predict relevance of unseen documents, and to combine eye movement and textual features in the IR task.

4.1.1 Ideal Weights

We use an SVM[1] for each search topic, to predict the relevance labels (two classes: relevant or not) given by the user. The input is the BOW representation of the document. The SVM has one weight for each term, and these weights are used as the ground truth or *ideal weights*.

The ideal weights are computed using one SVM per search topic against all other search topics. This procedure produces a weight value for each term in the

[1] All SVM's mentioned use the default setting of $C = 1$ and a linear kernel.

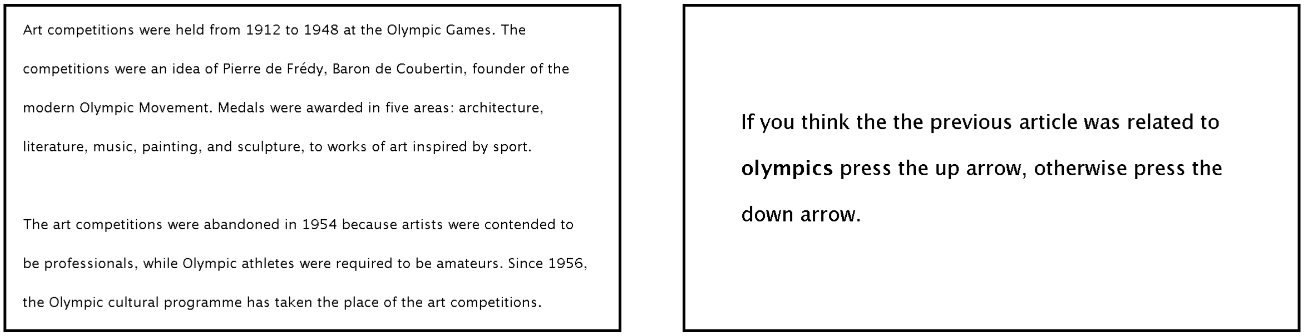

Figure 1: A sample document (on the left; it consists of two paragraphs) and the screen which the user gets after pressing a button (on the right).

dictionary; the weight represents the term's "fit" to the given search topic. We compute the SVM weights for all the non-truncated documents but restrict the dictionary to that of the truncated documents as there will never be eye movement features on terms that do not appear in the truncated documents. These SVM weights are the base in our work and we consider them to be the optimal baseline performance achievable.

We use a similar linear model, with weights given by the regressor described later in subsection 4.2, to predict the relevance of the unseen documents.

4.1.2 Combining Eye Movement and Explicit Feedback

We are interested in observing whether using eye movement features in conjunction with document features would improve the classification accuracy in comparison with only using the document features. This task differs from the one above in that now explicit relevance feedback is assumed available.

In our initial assumption we consider the measured eye movement information on the documents to be a second representation of the document.

When two views of the same phenomenon are available kernel Canonical Correlation Analysis (KCCA) (Hardoon, Szedmak, & Shawe-Taylor, 2004) has been shown to be an effective pre-processing step that can improve the performance of classification algorithms such as the SVM. The SVM-2K (Farquhar, Hardoon, Meng, Shawe-Taylor, & Szedmak, 2006) is a method that combines these two stages of learning (KCCA followed by SVM) into a single optimisation.

4.2 REGRESSION

We assume that there is a link between eye movements and the relevance of words within a document. This relevance can be thought of as the "fit" of the words in the documents to the search topic, which could be thought of as the *ideal weight* of a term for a given search topic. Our proposed model aims to learn this mapping, within each search topic, from the eye measurements on a word to its associated ideal weight value. The learned mapping can then be used to infer a new weight vector for new documents and new interests, which in turn is used as a classifier on new unseen documents.

In addition to a linear regressor, we use a sparse dual Partial Least Squares (PLS) approach (Dhanjal, Gunn, & Shawe-Taylor, 2006). This method uses a general framework for feature extraction based on kernel PLS (Rosipal & Trejo, 2001) deflation method. Here the projections are selected according to criteria that are useful for targeting them towards a particular task. In this task we choose the combination of features with maximal covariance. These are then used for least squares regression.

We use a Gaussian kernel with the sparse-KPLS where the σ parameter is computed, per search topic, using 10-fold cross validation on the training data.

5 EXPERIMENTAL SET-UP

We apply the Term Frequency Inverse Document Frequency (TFIDF) (Salton & McGill, 1983) on the documents in order to increase the weighting of terms that occur frequently within a document but infrequently across the corpus. If our corpus is given by D and dictionary by T, then the TFIDF for term $t \in T$ in document $d \in D$ is given by

$$\mathrm{TFIDF}(d,t) = n_{dt} \log \frac{|D|}{|\{\delta \in D \mid n_{\delta t} > 0\}|},$$

where n_{dt} denotes the number of occurrences of term t in document d. The TFIDF representation is used throughout the experiments.

5.1 EYE MOVEMENTS ONLY

Our main goal is to study whether it is possible to discriminate between the categories using a weight vector inferred only from the corresponding eye movement features. The set-up for this task is as follows; for each search topic we do the following:

- For each search topic we collect all the eye movements from the 10 viewed documents (we want both the documents labelled positive and negative) for each user viewing that topic and extract for each viewed word the set of word-specific features described in Section 3. These will form the inputs to our regressor, for each topic with corresponding outputs the weights of the corresponding words in the ideal weight SVM vector for that topic as discussed in Section 4.1.1[2].

- Leaving each topic out in turn, we use the data associated with all the other topics to train a regressor. This is accomplished by applying least squares regression to learn the corresponding regression coefficients.

- Using the learnt regressor we compute a weight vector for the left out search topic from eye movement features of the viewed terms. The features for reoccurring words are averaged. Zero is assigned to terms that had no fixations. We refer to the inferred weight vector classifier as W_i when the regressor is linear, and as $W_i(x)$ with the nonlinear regression, as described in subsection 4.2, with x being the number of feature directions used in the KPLS.

Figure 2 shows a graphical representation of the predicted weights for one document and one search topic.

5.2 COMBINING IMPLICIT AND EXPLICIT FEEDBACK

The classifiers that is trained with explicit feedback is created as follows; for each search topic select all truncated documents presented to the users who have been assigned that search topic, and using the labelling given by the users train an SVM. This gives a classifier for the explicit feedback condition that we refer to as SVM$_i$.

Next, we extract the eye movement features from the positive and negative documents presented to the same set of subjects. We ignore the eye movements on terms that do not belong to the dictionary.

[2]We normalise the weights in the 2-norm.

The SVM-2K$_i$ is used on the same training set of documents but now with the eye movement features providing a second representation of the truncated documents. For each document, we concatenate the eye movement features for words occurring in that document. We refer to this classifier as SVM-2K$_i$. In other words, SVM-2K$_i$, takes into account both the textual content and the eye movements for that document.

Note that for this case the eye movements features are not associated with individual words unlike in the earlier experiment. Furthermore in the testing phase there are no eye movements features available so only the text is taken into account (they are an average for word reoccurrence).

6 RESULTS

Evaluation criteria. The SVM predictor is used to rank the new documents in the order of expected relevance. Specifically, we compute the measure of performance by first ranking the 244 unseen documents in the test set according to their predicted relevance to the user, the document predicted most relevant having a ranking of one, and the document predicted least relevant having a ranking of 244. The *average precision* is then given by

$$PREC = \frac{1}{R}\sum_{i=1}^{R}\frac{i}{r_i}, \tag{1}$$

where R is the number of positive examples in the test set of 244 documents, and r_i are rankings of the positive examples, ordered such that $r_i < r_{i+1}$. If the classifier is working optimally, that is, we predict the positive examples to have rankings $1, \ldots, R$, i.e., $r_i = i$, we obtain a precision of 100%. Notice that for conciseness and in a slight abuse of terminology, in the following we use *precision* to refer to the average precision defined above.

Baseline models. The *random model* ranks documents uniformly at random. The median of the precision for the random model is obtained by taking the median of the precisions given by equation (1) with respect to all possible rankings of the positive examples. For 4 positive examples out of 244 documents, we obtain the median precision of $PREC_{RAND} = 2.37\%$, and for 10 positive examples we obtain the median precision of $PREC_{RAND} = 5.11\%$. The *mean* precisions for random model are 3.69% and 6.10%, respectively (overall 6%). We present the median rather than expected precision for the random model as we apply the Fisher Sign Test for a probabilistic estimate of our performance, this makes the assumption that the probability of drawing a positive (or negative) sample is one

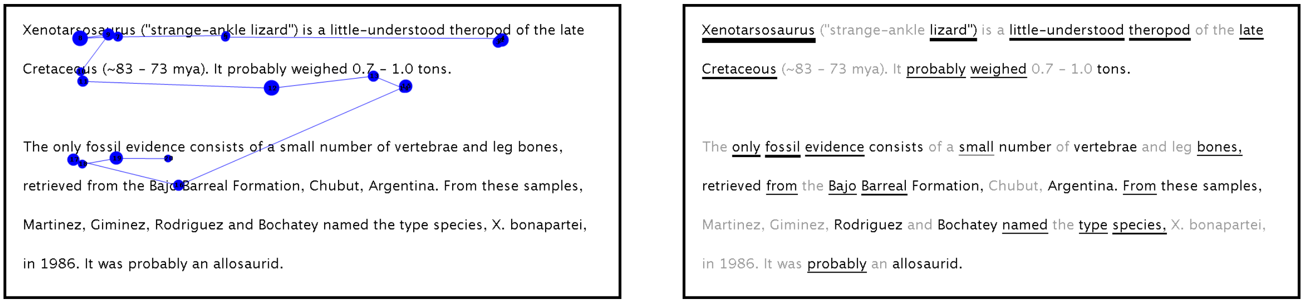

Figure 2: Sample plot of saccades (lines) and fixations (dots) on a document (on the left) and term weights inferred from eye movements on all documents in the Dinosaurs category (on the right). The magnitude of the inferred weight is shown by the thickness of the underlining. The words which do not appear in the dictionary are shown in light grey.

half. A random permutation will do on 50% of the time better than the median and 50% of the time worse.

The upper expected limit of performance is given by the *ideal weights*, denoted by SVM in the results in Table 2.

Eye movements only -models. The results are shown in Table 2[3]. The leftmost W_i column is for linear regression while the remaining $W_i(\cdot)$ columns are the non-linear regression models.

All regression models outperform the random classifier ($P < 0.01$, Fisher Sign Test). For example, only 5 search topics out of 25 in W_i perform worse than the median precision for random guessing. This is quite a strong result considering the complexity of the task.

We are inferring relevance of a new document based solely on eye movements from previous document viewings, and no other feedback. Several search topics achieve a much better than random median precision with Natural Disasters leading with 35.30% followed by Electronics with 16.48%. We believe these to be extremely promising results as they show that we are able to learn implicit information from eye movements for the task of search topic document relevance.

It is interesting to observe that the *Astronomy* and *Cities* search topics achieve a higher precision with the linear regression model. Despite these two results it is apparent that the non-linear approach outperforms the linear one across all selections of feature direction and in overall performance the sparse KPLS with 61 feature directions outperforms that of the others.

The 5 categories that perform worse than the median random precision are: *Dinosaurs, Internet, Literature, Sculpture* and *Writing systems*. It is reasonable to assume that the poor performance for *Internet* and *Literature* are due to the poor performance of the ideal SVM classifier. Since we are training the eye movements to those ideal weights it is unreasonable to expect good classification when inferring those weights from eye movements. One possible reason for poor performance in these categories might be that users may have read through most of the document instead of just finding the relevance, perhaps because the user was interested in the topic. This kind of reading behaviour would not emphasise the interesting words and would make it impossible to learn the regressor.

Eye movements combined with explicit relevance feedback and text content. It is not entirely surprising that we are able to observe that the SVM_i and SVM-$2\mathrm{K}_i$ models in Table 2 are always able to outperform random guessing. Observe that our initial assumption that using eye movements with the document features improves overall performance, although once we analyse the individual results on a search topic basis, we find it is not always the case that using eye movements improves the results. While it is hard to make an assessment of individual search topics the results imply that, due to the overall improvement in precision, using eye movements with document features improves the results.

We find that the SVM_i and SVM-$2\mathrm{K}_i$ classifiers actually outperform our baseline SVM on two search topics, i.e., the *Olympics* and *Sculpture* for the SVM_i and SVM-$2\mathrm{K}_i$, respectively.

7 DISCUSSION

We addressed the extremely hard task of constructing a query in an information retrieval task, without being given either an explicit query or explicit relevance feedback. Only eye movement measurements for a small

[3] Some $W(\cdot)_i$ results are missing due to insufficient time to complete computations.

Table 2: The precision for various predictors and search topics, in percent. Larger precision is better. The baseline models are random guessing, for which the median precision is shown, together with an SVM classifier used to train the regressor. The baseline models provide expected lower and upper limits for the expected performance of the predictors. The largest precision for a given search topic and class of predictors is shown in boldface. The topics for which the best W_i classifier has achieved a sub-random result are shown in italics. All models outperform random guessing ($P < 0.01$, Fisher Sign Test).

	Baseline		Eye movements only				Expl. feedb.	Impl.&Expl. feedback
	Random	SVM	W_i	$W_i(21)$	$W_i(41)$	$W_i(61)$	SVM$_i$	SVM-2K$_i$
Astronomy	5.11	58.92	**10.09**	9.84	9.98	10.02	36.35	**37.48**
Ball games	5.11	98.09	9.59	11.95	12.36	**12.83**	65.12	**78.57**
Cities	5.11	91.20	**9.94**	9.62	9.92	9.70	64.26	**64.86**
Court systems	5.11	67.55	8.68	9.25	9.47	**10.26**	59.23	**60.48**
Dinosaurs	5.11	100.00	*3.20*	*3.37*	*3.37*	**3.45**	**98.33**	89.75
Education	5.11	73.16	5.55	5.80	6.36	**6.88**	32.03	**41.12**
Elections	5.11	74.48	14.35	**16.48**	15.89	15.93	68.49	**68.87**
Family	5.11	77.12	8.33	10.74	**10.84**	10.76	50.35	**67.97**
Film	5.11	70.33	6.07	7.01	7.68	**8.30**	41.78	**48.32**
Government	5.11	60.10	6.10	**6.86**	6.66	6.80	**27.64**	25.59
Internet	5.11	37.69	*3.47*	**3.80**	*3.78*	*3.77*	11.72	**12.69**
Languages	5.11	97.14	6.03	6.55	**6.84**	6.81	**94.30**	92.90
Literature	5.11	33.77	*3.65*	*3.69*	*3.75*	**3.82**	**16.94**	15.61
Music	5.11	78.16	7.23	8.13	7.86	**8.25**	58.81	**69.24**
Natural disasters	2.37	100.00	33.84	34.30	34.89	**35.30**	**100.00**	91.67
Olympics	5.11	87.38	7.16	8.76	9.52	**9.59**	**100.00**	82.25
Optical devices	5.11	75.73	11.19	11.02	**11.23**	10.63	63.73	**63.85**
Postal system	5.11	77.83	5.57	5.97	5.99	**6.20**	**46.29**	44.20
Printing	5.11	78.11	5.93	**8.67**	8.40	8.02	**55.01**	51.35
Sculpture	5.11	72.92	*4.51*	*4.94*	**5.11**	*4.99*	72.78	**76.07**
Space exploration	5.11	67.00	14.52	16.49	**17.74**	16.92	62.64	**64.95**
Speeches	5.11	85.97	10.01	10.36	10.55	**10.72**	45.27	**46.94**
Television	5.11	69.01	6.07	**6.88**	6.78	6.66	**29.77**	25.41
Transportation	5.11	52.02	16.51	16.66	**16.73**	16.71	**34.86**	26.13
Writing systems	5.11	78.67	*3.61*	**3.62**	*3.60*	*3.52*	**35.49**	32.01
Average		74.49	8.85	9.62	9.81	**9.87**	54.85	**58.71**

set of viewed snippets, and the text content of the snippets was available. This is a prototype of a task where the intent or interests of the user are inferred from implicit feedback signals, and used to anticipate the users actions.

We were able to learn a "universal predictor of relevance predictors" from a collected database of queries, their relevant and irrelevant documents, and the corresponding eye movements. The predictions performed better than chance on new queries. There is ample room for improvement in the prediction percentages, our best model gives an mean precision of 9.87% as opposed to 6% of the random classifier, but nonetheless the feasibility study was successful; it is possible to extract some useful information from eye movements.

We further experimented with a model, where the textual content of the documents and explicit relevance feedback given by the user (the document is/is not in the search topic) was taken into account. As expected, the explicit feedback improved the precision significantly to 54.85%. Our results show that also in this scenario, taking the eye movement into account we can further improve the precision by about 4%.

We conclude that in constructing a query, eye movements provide an implicit feedback channel. As expected, the feedback obtained from the eye movements is less informative than relevance feedback typed in by the user, but nonetheless this implicit feedback can be exploited. In practical applications all available feedback channels, in addition to the eye movements, should of course be utilized; the practical implication of this study is that if eye movement data is cheaply available it might be a good idea to include it as well.

In this work, we used relatively standard state-of-the-art machine learning methods. An obvious direction for future research would be to develop a tailored method to obtain a query from eye movements, optimised as one single model.

Acknowledgements

AA and SK belong to Adaptive Informatics Research Centre, a centre of excellence of the Academy of Finland. The authors thank Wray Buntine from the Complex Systems Computation Group at University of Helsinki for providing the Wikipedia data and Craig Saunders from the ISIS Research Group at University of Southampton for many fruitful discussions. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778. This publication only reflects the authors views.

References

Dhanjal, C., Gunn, S. R., & Shawe-Taylor, J. (2006). Sparse feature extraction using generalised partial least squares. In *Proceedings of the IEEE international workshop on machine learning for signal processing* (pp. 27–32).

Farquhar, J. D. R., Hardoon, D. R., Meng, H., Shawe-Taylor, J., & Szedmak, S. (2006). Two view learning: Svm-2k, theory and practice. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems 18* (pp. 355–362). Cambridge, MA: MIT Press.

Hardoon, D. R., Szedmak, S. R., & Shawe-Taylor, J. R. (2004). Canonical correlation analysis: An overview with application to learning methods. *Neural Computation*, *16*(12), 2639–2664.

Joachims, T., Granka, L., Pan, B., Hembrooke, H., & Gay, G. (2005). Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 154–161).

Kelly, D., & Teevan, J. (2003). Implicit feedback for inferring user preference: a bibliography. In *ACM SIGIR forum* (pp. 18–28).

Maglio, P. P., Barrett, R., Campbell, C. S., & Selker, T. (2000). SUITOR: An attentive information system. In *IUI-2000: International conference on intelligent user interfaces* (pp. 169–176).

Maglio, P. P., & Campbell, C. S. (2003). Attentive agents. *Communications of the ACM*, *46*(3), 47–51.

Puolamäki, K., & Kaski, S. (Eds.). (2006, May). *Proceedings of the NIPS 2005 workshop on machine learning for implicit feedback and user modeling.* Otaniemi, Finland. (http://www.cis.hut.fi/inips2005/)

Puolamäki, K., Salojärvi, J., Savia, E., Simola, J., & Kaski, S. (2005). Combining eye movements and collaborative filtering for proactive information retrieval. In *Proceedings of SIGIR 2005, twenty-eighth annual international ACM SIGIR conference on research and development in information retrieval* (pp. 146–153). ACM.

Rosipal, R., & Trejo, L. J. (2001). Kernel partial least squares regression in reproducing kernel hilbert space. *Journal of Machine Learning Research*, *2*, 97–123.

Salojärvi, J., Kojo, I., Simola, J., & Kaski, S. (2003, September). Can relevance be inferred from eye movements in information retrieval? In *Proceedings of WSOM'03, workshop on self-organizing maps* (pp. 261–266). Hibikino, Kitakyushu, Japan: Kyushu Institute of Technology.

Salojärvi, J., Puolamäki, K., & Kaski, S. (2005). Implicit relevance feedback from eye movements. In W. Duch, J. Kacprzyk, E. Oja, & S. Zadrozny (Eds.), *Artificial neural networks: Biological inspirations – ICANN 2005* (pp. 513–518). Berlin, Germany: Springer-Verlag.

Salojärvi, J., Puolamäki, K., Simola, J., Kovanen, L., Kojo, I., & Kaski, S. (2005, March). *Inferring relevance from eye movements: Feature extraction* (Tech. Rep. No. A82). Helsinki University of Technology, Publications in Computer and Information Science. (http://www.cis.hut.fi/eyechallenge2005/)

Salton, G., & McGill, M. J. (1983). *Introduction to modern information retrieval.* New York, NY, USA: McGraw-Hill, Inc.

Turpin, A., & Scholer, F. (2006). User performance versus precision measures for simple search tasks. In *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 11–18).