

NewsFlash: Adaptive TV News Delivery on the Web

Alan Haggerty, Ryen W. White and Joemon M. Jose

Department of Computing Science,
University of Glasgow,
Glasgow, Scotland. G12 8QQ.

{haggerty, ryen, jj@dcs.gla.ac.uk}

Abstract. In this paper we present NewsFlash, an adaptive search system that assists a searcher to efficiently search a library of stored TV news reports. The system automatically summarises the closed-caption subtitles embedded in the TV broadcasts and selects appropriate sentences to best describe report content in respect to the searcher's query. During interaction the system selects useful terms from these summaries and uses these terms to update the display and potentially update a stored searcher profile. We evaluate the worth of our approach with real searchers and realistic information seeking scenarios. A novel means of testing the worth of a permanent profile of searchers' general interests is also proposed. Our results show that the adaptive techniques we propose can work well in multimedia search environments.

1 Introduction and Motivation

We are all active consumers of information. However, our thirst for knowledge can soon turn into a glut if we are not careful about the amount of information we endeavour to consume. Information overload [8] is a well-recognised, common problem and to tackle it we must be selective in how we choose to allocate our cognitive resources and manage our time. Television news is perceived as a means of reducing this overload and bringing only the important stories to the attention of viewers. The advent of 24-hour TV news channels has dramatically increased the available information and in doing so has placed increasing demands on viewers to filter out irrelevant stories and make optimal use of their time. When attempting to locate stories of interest, viewers must often examine many that may not match their current information need or general interests.

Multimedia Information Retrieval (MMIR) systems are a way of helping searchers locate information of interest from large multimedia corpora. The use of many such systems requires searchers to explicitly devise queries that represent their needs in a language understood by the search system.

However, queries are only an approximate, or 'compromised' information need [13], and may fall short of the description necessary to infer relevant media objects (e.g. images, videos). The problems searchers have with expressing their information needs have been acknowledged [6]. These difficulties are compounded by the difficulty of providing increasingly better ranked results based solely on the initial query.

Consequently, search systems need to offer robust, reliable methods for query modification.

Relevance Feedback (RF) has become a popular method for automatically improving a system's representation of a searcher's information need. Searchers view a number of media objects and provide explicit feedback on which objects are relevant. The RF system is then able to formulate a query more attuned to the actual information need of the searcher and provide results that pertain more closely to this need (i.e. similar to those marked as relevant). Adaptive search systems unobtrusively monitor search behaviour and remove the need for the searcher to explicitly indicate which objects are relevant or formulate revised queries. They develop and enhance their knowledge of searcher needs incrementally from inferences made about their interaction and use this knowledge to help searchers in their seeking.

In this paper we present *NewsFlash*, a prototype adaptive search system for online TV news that automatically offers searchers alternatives and extensions to their initial query based on their interaction. The system allows searchers to access large quantities of news in an efficient way, using story detection to combine news reports from different broadcasts, hence reducing redundancy. The system monitors the ephemeral interaction of the searcher, and based on their actions, offers them support in their seeking. The fundamental assumption behind the adaptive approach we adopt is that a searcher's actions are driven by their information need and that actions reflect needs as needs influence actions.

The system operates on two temporal levels, combining a permanent, modifiable profile, unique to the searcher, with adaptive interface technologies and the selection of potential query expansion terms based on recent interaction. In response to searcher interaction, the system can reorder a list of news stories and recommend terms to be added to a searcher's permanent profile.

Systems that personalise the interaction with video catalogues are becoming increasingly popular [4,5]. Such systems use stored searcher profiles and require the searcher to enter and reformulate queries explicitly. This relies on the searcher being able to adequately express their information needs, something that they may not always be able to do. The use of relevance feedback for video retrieval is also popular, and attempts to tackle this problem, yet still relies on the explicit relevance assessments of searchers to automatically reformulate the initial query using, for example, content-based techniques (e.g. colour features, histograms) [2,9]. This approach is effective but onerous, and as a result searchers may be unwilling to provide such feedback.

Fischlár-News [12], an extension to the Fischlár [7] video search system, facilitates access to large volumes of daily evening TV news. Fischlár-News has been operational for two years and offers searchers access to a large database of stored news, indexed with the closed-caption subtitles of the programs. The system requires the searcher to enter a textual query and returns a list of potentially interesting stories. The system does not offer any assistance in refining or enhancing searchers' initial queries and no means of storing the general interests of particular searchers. The aim of this paper is not to present an alternative to well-developed technologies such as Fischlár-News, but to focus on a particular aspect of the overall search experience that

may prove helpful to searchers and that systems such as Fischlár-News currently lack, namely *adaptive querying*.

To evaluate the worth of our approach we solicit real searchers, and present them with realistic information seeking scenarios. We compare searcher effectiveness between a baseline and systems that implement varying degrees of adaptivity.

In the remainder of this paper we describe the NewsFlash system and the adaptive methods, Section 2, and the evaluation methodology in Section 3. The results of our evaluation are presented in Section 4, are discussed in Section 5 and we conclude in Section 6.

2 NewsFlash

The NewsFlash adaptive search system was developed to meet a perceived need for TV viewers to be able to search through hours of video in the minimum of time, allowing them to view the stories they want without having to view the news programmes in their entirety.

The system is divided into two main components; *Indexing* (responsible for the off-line processing of the video) and *Search* (responsible for the online searcher interaction). In this paper we focus on the *Search* components and in particular the adaptive querying aspects of the search interface. However, it may be helpful for context and later reference to briefly describe the subcomponents that combine to form the *Indexing* part of the system.

We split the processing and retrieval of the videos for performance reasons alone. The overheads involved in processing the images were too great to practically process video ‘on the fly’ at query time.

2.1 Indexing

The Indexing component handles all aspects of video segmentation. There are five parts of this component (Fig. 1). The parts operate in the order illustrated in the figure.




Figure 1: Indexing components

The shot detection mechanism uses the comparison of colour histograms to detect shots in a video. It takes two frames from the video and compares the overall level of each colour within each frame. If the difference exceeds a certain threshold, then it is inferred that the two frames originated in different shots. In NewsFlash, we initially grab 100 keyframes, evenly distributed throughout the video. The shot detection then

passes through this set, calculating the exact location of the shot transitions. The method used is similar to that in [3].

The video then needs to be categorised, an essential requirement of the search process. NewsFlash carries out text searches on the soundtrack of the video (encoded in the closed-caption subtitles broadcast with the news). These subtitles take the form of a text file with metadata and the text of the soundtrack and associated timestamp. The subtitle alignment is carried out after the shot detection, since the boundaries between shots will be known and timestamps from subtitles and boundaries can be aligned. Captions that are deemed to be within the shot are added to the text of that shot.

The shot detection often results in a series of fractured clips that individually do not form a coherent story. NewsFlash merges shots to present one video per story. The videos are clustered by analysing the textual content of the shots and if they are sufficiently similar they are deemed to be part of the same story and the shots and their text are merged together.

To allow the searcher to view search results in the web-based tool, a medium of presentation was developed where the searcher is shown a thumbnail of the video. This means that some still images of the video report had to be stored to allow the searcher to quickly preview the story. The system stores the first frame and the middle frame of the story, to allow searcher's to preview what happens as the story unfolds. The middle frame is shown in place of the first frame on thumbnail mouseover. Thumbnail generation is not query dependent, so the same thumbnail is shown regardless of the query entered.

An inverted index is constructed of the textual content of the video to allow fast and efficient searching. The text of the video is represented internally using a standard inverted index and postings list.

2.2 Searching

The main functionality of the NewsFlash system is embodied in its adaptive search mechanism. Terms are entered explicitly by the searcher, or chosen implicitly by the system's adaptive components and acted on by the search mechanism. The result of these searches is a series of ranked stories with associated textual summaries and still images that illustrate them. The interface presented in response to a submitted query is shown in Fig. 2.

The interface contains a number of features, the top 12 stories, ranked in descending order of relevance to the searcher's initial query (left to right, top to bottom), the video playback panel (for playing the selected video), the query reformulation panel and the profile search panel. A searcher views a video by clicking on the thumbnail for that story. When the searcher moves the mouse over the thumbnail they are shown the middle shot of the story and a four sentence summary. The summary is biased toward the initial query of the searcher and potentially consists of four query-relevant sentences.

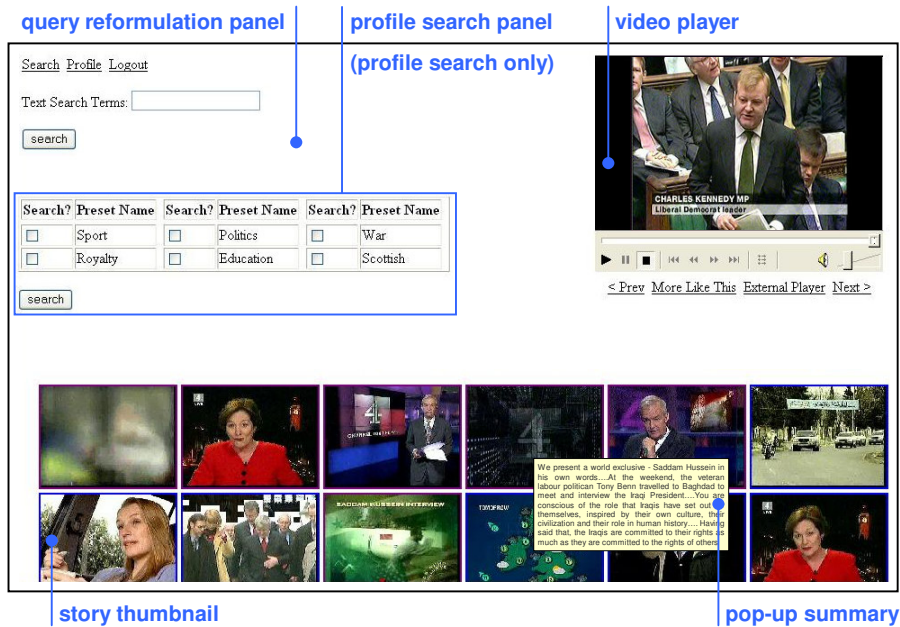


Figure 2: NewsFlash results interface

The system supports two main forms of search; a *plain text search* and a *profile search*. The plain text search removes stopwords from the searcher's explicit textual query and the remaining terms are checked against the entries in the inverted index. The stories are scored using a best-match *tf.idf* weighting scheme and the top 12 results are presented to the searcher, ranked in descending order of their relevance score. In the profile search, NewsFlash uses a stored profile, unique to each searcher, to form the query. The *topics* within each profile reflect the general interests of the searcher. The *terms* resident in each topic are the specifics of the general interests, and these terms are used to rank stories depending on the profile selected. The profile search therefore brings stories that are of general interest to the attention of the searcher. Stories on topics that not of interest are pushed towards to bottom of the list. The higher ranked stories form a potentially relevant information space in which to begin deeper investigation. More details on this feature are given in Section 2.4.

NewsFlash is an adaptive search system that expands the initial query in light of relevance information gleaned from searcher interaction. Searchers may have problems in adequately expressing their information needs and the terms selected by the system are those that NewsFlash perceives are of the highest utility for their current search. The next section describes how these terms are automatically selected.

2.3 Term selection

When the searcher passes the mouse pointer over a thumbnail they are presented with a query-specific textual summary of the story that shot describes. The searcher then has the option of clicking the video (to watch it play) or moving the pointer away. In our approach we regard the playing of a video as a strong indication of relevance. On *every such indication* we select possible expansion terms from the summary of the story just played and use these terms to reorder the list of stories.

The playing implies interest in the content of the video. When the searcher makes the choice to play, she does so having seen the first shot of the video, the middle shot of the video and a four sentence summary of the content of the video subtitles, biased towards the searcher’s initial query. It is assumed that these three things are sufficiently indicative to allow the searcher to make an sound assessment of potential relevance.

To rank possible expansion terms, NewsFlash uses the *wpq* algorithm [11] shown in Equation 1. For any term t , where N is the total number of summaries, n_t is the number of summaries containing t , R is total number of relevant summaries and r_t is the total number of relevant summaries that contain t . R and r_t are based on relevance assessments and are therefore prone to increase as searchers interact with the systems. Possible expansion terms only come from the document summary generated by the system, and it is assumed that N is set to equal the total number of stories that are detected, and therefore ready to be retrieved by the system.

$$wpq_t = \log \left(\frac{(r_t + 0.5)(N - n_t - R + 0.5)}{(n_t - r_t + 0.5)(R - r_t + 0.5)} \right) \times \left(\frac{r_t}{R} - \frac{n_t - r_t}{N - R} \right)$$

Equation 1: *wpq* formula

All terms are ranked according to their score and the new query representation – the original terms plus the new expansion terms – is presented to the searcher and then used to reorder the list of stories. At this point it is possible for stories to leave the top 12 shots displayed, as the ranking of stories previously resident outside the top 12 is boosted by the inclusion of the query expansion terms. The expansion terms selected will have a high discriminatory power between relevant and non-relevant stories and will therefore promote those that are relevant, whilst demoting those that are not. The reordering uses the top six terms plus the searcher’s original query, and as mentioned earlier, uses a best match *tf.idf* weighting scheme to rank the stories. The searcher has the option to undo the effects of any reordering operation. The system reorders *then* offers the option to undo, rather than simply offering the option to reorder. The authors feel that it is better to show searchers the output of the action and let them decide on the value of the action, rather than let them rely on the *perceived value of the potential action*, a judgment that may be tainted by previous experience.

We fully acknowledge the uncertainty surrounding the use of implicit evidence and as a result, we assign the six expansion terms a reduced weight in the new query. In our system, the expansion terms are given a weight half that of their original *tf.idf*

weights. Equation 2 shows this, where q_1 is the expanded query, q_0 is the initial query, θ is the normalising factor (set depending on the reliability of the implicit evidence) and q_{e_i} is an expansion term¹.

$$q_1 = q_0 + \theta(q_{e_1}, \dots, q_{e_n})$$

Equation 2: Term re-weighting

The decision to set $n = 6$ was based on the authors' previous experience with similar systems [14,15], where this number of expansion terms has proven sufficient.

The effect of the scoring is cumulative. Stories that start with a low initial ranking can 'bubble up' to the top of the list in response to a series of expanded queries that match the story content. All expansion term scores are reset to zero when the searcher enters a new query and new result set is generated. The adaptive component therefore only operates in the interaction *between* searcher-controlled query iterations.

2.4 Searcher profile

Whilst the modified query created automatically by the system during interaction is only temporary, and disappears when the searcher submits a new query, the system also provides a means of personalisation that allows a searcher to express their long-term general interests. NewsFlash supports the creation of a searcher profile, unique to each searcher. Within these profiles are general topics that can be created by the searcher, and within the topics are keywords, either entered explicitly, or *recommended* by the system using the methods described in the previous section. When the system reorders the list of stories, it also displays the top six expansion terms to the searcher, who can then decide which terms to store in their profile and which to discard. Fig. 3 shows an example of what is shown on the results interface directly above the list of thumbnails. In this figure the searcher had previously viewed stories on the latest Iraq war.

We recommend adding the words: borders , arrive , wisdom , planes , centre , raf , to profile "war"

Figure 3: Profile term recommendation

Unlike the story reordering described in the previous section, the searcher is given complete control over which terms are added to their profile.

There are two temporal states in the system, and a term is resident in one state at any particular time. The states are a *transient*, temporary set of terms selected based on their estimated worth (wpq score) for each search session and a *permanent*, stored set of terms for each topic in the searcher's profile that traverses multiple search sessions. Terms can move between states, although searchers have control over which terms make this move (Fig. 4).

¹ It is worth noting that *wpq* is only used to *rank* potential expansion terms, and plays no role in expansion term *weighting*.

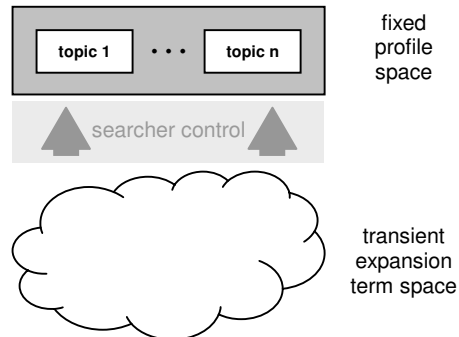


Figure 4: Temporal term states in NewsFlash

NewsFlash can make recommendations about terms to add to the profile based on videos the searcher has watched. The system therefore only recommends potential terms, it does not update the profile without searcher consent.

3 Evaluation

In this section we describe the evaluation of NewsFlash. In particular we detail the three versions of the NewsFlash system used for this evaluation, the participants involved, the tasks undertaken, the techniques used to evaluate the searcher profile and the experimental methodology employed.

3.1 Research Hypotheses

We endeavour to test whether (H_1) the reordering of the list of stories and (H_2) the addition of terms to a permanent profile, leads to increased searcher satisfaction and more effective searching.

3.2 Systems

Three systems were used in the evaluation, each with a different level of adaptive functionality. As the versions of the NewsFlash system were specifically developed for the evaluation, the level of functionality depended directly on the research hypotheses. System 1 was our experimental baseline, was not adaptive, and therefore did not carry out any operations automatically on the searcher's behalf. System 2 chose terms from summaries viewed by the searcher and used potentially useful non-query terms to reorder the list of shots (for H_1). System 3 (effectively the complete NewsFlash system) reordered the list of stories, as in System 2, but also recommended potentially useful terms to add to the searcher's saved profile under an appropriate topic heading (for H_2). The topic headings were unique to each searcher.

3.3 Subjects

We recruited 9 subjects, who were all male, and had an average age of 21.7 years. The subjects were all educated to graduate level and used computers and search tools frequently. Subjects were also regular viewers of television news programmes. 6 out of the 9 subjects watched Channel 4 news (the source of the test data) on a regular basis.

3.4 Tasks

In our evaluation each subject was asked to complete three search tasks. Each search task was placed within a simulated work task situation, [1]. This technique asserts that subjects should be given search scenarios that reflect real-life search situations and allow the searcher to make personal assessments on what constitutes relevant material. An example task is shown in the appendix.

3.5 Evaluating Searcher Profiles

System 3 allowed the searcher to save some or all of the terms recommended by the system for use in later queries. These terms are stored in the searcher's profile under the appropriate topic headings. For example, the terms "Gulf Iraq Saddam" may be put under the topic heading War. To cater for contextual ambiguities, the searcher had the choice of which terms are retained in the profile.

To test the effectiveness of the searcher profile (in terms of applicability in subsequent search sessions) the subjects were invited to return the day after their first session, and attempt a second task on System 3 only. Systems 1 and 2 did not implement the profile updating, so there was no need to test this feature on these systems. The follow-up task was varied according to the subject's level of success at the original task (attempted on the first day). If the subject was successful, they were asked to attempt a task that is a continuation of the original one (i.e. they are assumed to know the answer to the original task), if they do not succeed they re-attempt the original one. The follow-up task for the example given in Section 3.4 is shown in the appendix.




Figure 5: Allocation of the follow-up task

Fig. 5 shows the allocation of the follow-up task (t_1) depending on the success (S) or failure (F) of the original task (t_0).

3.6 Data Set

The data set for this evaluation was the Channel 4 television news² from two consecutive days, the 25th and 26th of February 2003. This comprised of 2 hours of footage. The small data set was used to test our approach and does not reflect the length of video that NewsFlash could adequately handle.

3.7 Methodology

The experiment followed a within-subjects, repeated measures design (i.e. each participant completed 3 tasks in total, one on each of the 3 search systems). Tasks and systems were allocated according to a Greco-Latin square design. To evenly distribute fatigue and learning effects we rotated both the order in which tasks and systems were presented across participants. The negative effect of task-bias minimised by this rotation. Each subject was given 10 minutes to complete each task, although the subjects could terminate the search early if they felt they had completed the task. The time restriction ensured consistency between subjects.

The subjects were welcomed and given a short tutorial on the features of the three systems being tested. We also collected background data on aspects such as the subjects' experience and training in online searching. After this, subjects were introduced to tasks and systems according to the experimental design. Subjects were instructed to attempt the task to the best of their ability and write their answer on a sheet provided. A search was seen to be successful if the searcher felt they had succeeded in their performance of the task.

When they completed a search, the subjects were asked to complete questionnaires regarding various aspects of the search. We used semantic differentials, Likert scales and open-ended questions to collect this data.

4 Results and Analysis

In this section we present preliminary results from our system evaluation. In particular, we focus on results pertinent to each of our two research questions: the worth of the adaptive document reordering and the worth of the searcher profile. Tests for statistical significance will be given where appropriate with $p \leq .05$, unless otherwise stated. S_1 , S_2 and S_3 denote System 1, System 2 and System 3 respectively. We ana-

² Channel 4 is a UK terrestrial television channel.

lyse aspects of the results individually. M is used to denote the mean, and 5-point scales are used throughout, with a value of 1 reflecting more agreement.

4.1 Task Success

Task success was measured from the perspective of the searcher and therefore searchers decided whether they had completed a task. This was felt to be representative of most real-life search situations³ and fitted well with the use of simulated work task situations. The differences in levels of task success between the systems were marked. When using the experimental baseline (System 1), searcher's completed 80% of their tasks, with 63% on System 2 (reordering only) and 53% on System 3 (complete NewsFlash). The differences between all systems are not significant with paired T -tests (T_{16}), the significance values are shown in Table 1.

Table 1. Inter-system comparison of task completion rate

Comparison	Task completion rate	T -value	Significance
S_1 vs. S_2	80% vs. 67%	.57	.426
S_1 vs. S_3	80% vs. 53%	1.34	.130
S_2 vs. S_3	67% vs. 53%	.53	.473

4.2 Reordering of Stories

The stories were reordered every time the system received a positive relevance indication (i.e. playing of a video). This reordering shaped the result set to reflect the perceived current information needs of the searcher. To test whether the reordering was effective, subjects were asked to indicate on a Likert scale how helpful this operation was. The results showed that the results significantly differed from the median value (i.e. 3, $M = 2.286$) using a paired T -test ($T_{16} = 2.38$, $p = .032$). Subjects felt that the reordering of the stories by the system on their behalf helped, rather than hindered their seeking.

4.3 Searcher Profiles

NewsFlash maintained a stored profile of searcher's general interests over a period of time. To test this, searchers participated over two days. The task they attempted on the second day was dependent on their level of success with the task on the first day. In this regard, participants were asked to rate, on a Likert scale, the worth of being able to store terms during the original task and reuse them for the follow-up task. The results significantly differed from median ($M = 1.714$, $T_{16} = .98$, $p = .176$). This per-

³ See Reid [10] for a counterexample.

haps shows that the searcher profiles can indeed be useful, assuming the follow-up tasks are in some way related to the original task. Through dividing the profile into a number of topics, and allowing searchers to place terms within these topics we cater for the potential diversity of long-term needs and intentions. Ultimately the searcher has control over which topic is used in NewsFlash at a certain time, so the terms chosen for the profile search will relate to their current *general* interests. The system will suggest *specific* additions to the profile based on searcher interaction.

4.4 Worth of Expansion Terms

When selecting terms on behalf of the searcher it is important that these terms are of use in their search. To test this aspect of the NewsFlash system, participants were asked to rate (using a semantic differential) whether the terms added to their original query were useful *always*, *occasionally* or *never*. The responses were mixed ($M = 3$) and not significantly different from the median value with a paired T -test ($T_{16} = .65$, $p = .375$). The results seem promising if we allow for the uncertainty inherent in implicit techniques of this nature and the relative simplicity of the methodology used to select expansion terms.

5 Discussion

Selecting worthwhile terms on behalf of searchers relies on an ability to predict their information needs to a very fine level of granularity. The approach used in NewsFlash to detect searcher needs is coarse-grained and as such is likely to produce a certain number of erroneous terms (i.e. not all terms in a summary will be relevant). We acknowledge this uncertainty in Equation 2, assigning a reduced weight to the terms selected implicitly. It is envisaged that as the system's ability to perceive information needs improves through modifications to our approach, the normalisation weight θ will tend to one, and the terms selected by the system will play a more important role in the new query.

As is evident in the previous section, many of the differences in the results presented were not statistically significant. This is perhaps due to the small number of experimental participants in this preliminary study. The most disconcerting result is that whilst subjects liked the query expansion, it appears to have reduced their perception of task success. This could be for a number of reasons, perhaps the most likely is the unfamiliarity of NewsFlash's adaptive features. This is supported by increased levels of task success on the follow-up task, over the original task. Despite the pre-experiment tutorial, the original task was the first time searchers had an opportunity to attempt a 'real' task in a time-constrained context. We posit that as searchers become more familiar with the adaptive features they will become more effective in their seeking.

The reordering of the stories and the recommendation of terms for addition to the permanent searcher profile were well received by subjects. This is a promising result,

the experimental baseline was strict (i.e. used the same search interface as the adaptive systems), so such any responses that suggested performance gains or increased levels of searcher satisfaction are indeed worthwhile.

During the experiments, we also logged task completion time and number of query iterations for each of the three systems. There were no significant differences between any of the systems for each of these measures. Finding that the adaptive system performed just as well as the baseline showed that it did not hinder searchers, but it did not help them sufficiently (in this regard at least) to reduce task completion times by a significant level.

Due to the small sample size, it was thought that the most interesting results may come from the informal comments of searchers as they searched. These comments proved promising, searchers liked the complete NewsFlash system substantially more than the other systems. The negative comments concerned small interface issues, although one participant did raise what was thought to be a particularly valid concern. He suggested that NewsFlash “*could do with more descriptions of what was happening*”. This is an important point, if adaptive systems are going to work on behalf of the searcher, it seems reasonable that they explain their actions rather than detaching themselves completely. At present, adaptive search systems assume a ‘black box’ approach to assisting those they are meant to help. Explanations may open this box, help engender trust in NewsFlash’s actions and perhaps bridge the gap between searcher and system.

6 Conclusions

In this paper we have presented NewsFlash, an adaptive querying system for searching online repositories of stored TV news footage. Unlike similar systems developed to search such corpora, our approach implicitly refines and enhances searchers’ initial queries based on their interaction. The system adapts to the needs of searchers without requiring the searcher to explicitly define these needs or changes in them. We proposed a means of extracting and weighting terms for implicit query expansion during the interaction with the results of a multimedia search system. We use this query for two operations; reordering the list of stories and suggesting terms to be added to a searcher’s unique profile.

We carried out a preliminary evaluation of the system with real searchers and realistic information seeking scenarios and adopted a novel means of testing the worth of stored searcher profiles. The results, whilst not being statistically significant, appeared promising. Searchers felt that terms selected by the system were useful and storing terms in the searcher profile was helpful for future searches. The results also indicated adaptive interface support may be of most use to novice searchers who typically have problems conceptualising their information needs.

Systems that are able to adapt to the needs of searchers can remove the cognitive burden and expense of term selection and query reformulation and offer them additional support in their seeking. Systems of this nature have the potential to improve

the search experience for struggling searchers everywhere. Future versions of the NewsFlash system will focus on better identifying information needs and explaining its actions to those who use it.

Appendix

Original Task

You are considering travelling abroad, you have however been getting strange letters through the post, informing you that you are defaulting on credit card accounts you don't have. You are concerned that someone may have stolen your identity.

Try to find if there are any similar cases in the news today.

Follow-up Task

Derek Bond's family were interviewed after his release. What was their opinion on the actions of the British government in this case?

References

- [1] Borlund, P. 'Experimental components for the evaluation of interactive information retrieval systems'. *Journal of Documentation*, **56**(1). pp. 71-90. 2000.
- [2] Chen, J-Y., Taskiran, C., Albiol, A., Bouman, C. and Delp, E. 'ViBE: A Video Indexing and Browsing Environment'. *Proceedings of the SPIE Conference on Multimedia Storage and Archiving Systems*. pp. 148-164. 1999.
- [3] Girgensohn, A. and Boreczky, J. 'Time-constrained Keyframe Selection Technique'. *Proceedings of the IEEE International Conference on Multimedia Computing and Systems. IEEE Computer Society*. pp. 756-761. 1999.
- [4] Haas, H., Bolle, R., Dimitrova, N., Javenski, A. and Zimmerman, J. 'Personalized News through Content Augmentation and Profiling'. *Proceedings of the IEEE International Conference on Image Processing*. 2001.
- [5] Jasinschi, R.S., Dimitrova, N., McGee, T., Agnihotri, L. and Zimmerman, J. 'Video Scouting: An Architecture and System for the Integration of Multimedia Information and Personal TV Applications'. *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*. 2001.
- [6] Koenemann, J. and Belkin, N.J. 'A case for interaction: A study of interactive information retrieval behaviour and effectiveness'. *Proceedings of the Conference on Human Factors in Computing Systems*. pp. 205-212. 1996.

- [7] Lee H, Smeaton A.F, O'Toole C, Murphy N, Marlow S and O'Connor N. 'The Físchlár Digital Video Recording, Analysis and Browsing System'. *Proceedings of RIAO 2000 - Content-based Multimedia Information Access*. 2000.
- [8] Nelson, Mark R. 'We Have the Information You Want, But Getting It Will Cost You: Being Held Hostage by Information Overload'. *Crossroads*, **1**(1). 1994.
- [9] Nilback, W., Hafner, J.L., Breuel, T., Ponceleon, D. 'Updates to the QBIC System'. *Proceedings of the SPIE Conference on Multimedia Storage and Archiving Systems* pp. 150-161. 1997.
- [10] Reid J. 'A new, task-oriented paradigm for information retrieval: implications for evaluation of information retrieval systems'. *Proceedings of the 3rd International Conference on Conceptions of Library and Information Science*. pp. 97-108. 1999.
- [11] Robertson, S.E., 'On Term Selection for Query Expansion'. *Journal of Documentation*, **46**(4). pp 359-364. 1990.
- [12] Smeaton A.F., Lee H, O'Connor N, Marlow S. and Murphy N. 'TV News Story Segmentation, Personalisation and Recommendation'. *Proceedings of the AAAI 2003 Spring Symposium on Intelligent Multimedia Knowledge Management*. 2003.
- [13] Taylor, R.S. 'Question-negotiation and information seeking in libraries'. *College and Research Libraries*, **29**. pp. 178-194. 1968.
- [14] White, R.W., Ruthven, I. and Jose, J.M. 'The use of implicit evidence for relevance feedback in web retrieval'. *Proceedings of the 24th BCS-IRSG European Colloquium on IR Research*. pp. 93-109. 2002.
- [15] White, R.W., Ruthven, I. and Jose, J.M. 'Finding Relevant Web Documents using Top Ranking Sentences: An Evaluation of Two Alternative Schemes'. *Proceedings of the 25th Annual ACM SIGIR Conference on Research and Development in Information Retrieval*. pp. 57-64. 2002.