

User Evaluation of Físchlár-News: An Automatic Broadcast News Delivery System

HYOWON LEE, ALAN F. SMEATON, and NOEL E. O'CONNOR

Dublin City University

and

BARRY SMYTH

University College Dublin

Abstract: Technological developments in content-based analysis of digital video information are seeing much progress, with ideas for fully-automatic systems being proposed and now demonstrated. Yet because we do not yet have robust, operational video retrieval systems that could be deployed and used by people the usual HCI practise of conducting a usage study and an informed iterative system design is thus not possible. Físchlár-News is one of the first line of automatic, content-based broadcast news analysis and archival systems that process broadcast news video to allow users to search, browse and play it in an easy to use manner using a conventional web browser. The system incorporates a number of state-of-the-art research components, some of which are not yet considered as mature technology, yet it has been built to be robust enough to be deployed to users who are interested in access to daily news, throughout a university campus. In this paper we report and discuss a user evaluation study conducted with 16 users, each of whom used the system freely for a 1 month period. Results from a detailed qualitative analysis are presented, looking at collected questionnaires, incident diaries and interaction log data. The findings suggest our users used the system in conjunction with their other news update methods such as watching TV news at home and browsing online news websites at their workplace, the major concerns being up-to-datedness and coverage of the news content. They tried to accommodate the system to fit their established web browsing habits, and they found local news contents and being able to play self-contained news stories on their desktop as a major value of the system. Our study also resulted in a detailed wishlist of new features which will help in further development of our and others' systems.

Categories and Subject Descriptors: H.5.1 [Multimedia Information Systems]: Evaluation/Methodology

General Terms: Design, Experimentation, Human Factors

Additional Keywords and Phrases: User evaluation, usage analysis, content-based video retrieval

Authors' addresses: H. Lee, A.F. Smeaton, N.E. O'Connor: Adaptive Information Cluster and Centre for Digital Video Processing, Dublin City University, Glasnevin, Dublin 9, Ireland; email: hlee@computing.dcu.ie, asmeaton@computing.dcu.ie, oconnorn@eeng.dcu.ie; B. Smyth: Adaptive Information Cluster, School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland; email: barry.smyth@ucd.ie

1. INTRODUCTION

It is clear that early promises of easy access to the right information at the right time are beginning to ring hollow as users struggle to locate relevant information quickly and easily. For example, conservative estimates of the Web's current size speak about its 10 billion documents and a growth rate that tops 60 terabytes of new information per day [Roush 2004]. To put this into perspective, in 2000 the entire World-Wide Web consisted of just 21 terabytes of information. Now it grows by 3 times this figure every single day [Lyman and Varian, 2003]. This makes for an extremely challenging information retrieval environment, but only reflects one small part of a much larger information access problem. As we begin to consider richer forms of non-text media (video, audio etc.) the challenge becomes even greater. In particular, recent market developments are now beginning to usher in the next generation of convergent media devices that combine traditional PC and home entertainment technologies. The result is the availability of much larger stores of rich-media content, but once again we are faced with the significant challenge of how best to provide users with effective access to this content.

In this paper we focus on the domain of TV news and how users can be helped to access broadcast news content through a mixture of searching, browsing and recommendation interfaces. Broadcast news is an excellent domain of study for a variety of reasons. For a start, broadcast news is commonplace and familiar, with most major TV channels devoting at least 2 hours of programming to news per day; of course a small number of dedicated new channels deliver far more than this, but represent an extreme. In addition, most people have a keen interest in the news and many will take the time to watch news programmes at least once per day. However, for most of us, our interest extends to just a portion of the stories that are aired on any particular day. Moreover, during the course of some extended period of time we are likely to be interested in tracking the development of certain news threads. These observations motivated our work on the Físchlár-News system, as described in this paper. Físchlár-News provides users with access to archived daily news broadcasts through a range of enhanced searching and browsing facilities. In addition, by profiling the news consumption habits of individual users, Físchlár-News is able to proactively recommend news stories based on the collective preferences of a community of like-minded individuals. Físchlár-News achieves this by automatically segmenting new broadcasts into indexable story units and by combining video and audio knowledge with closed-caption textual information in order to provide a platform for search, browsing and recommendation.

The remainder of this paper is organised as follows. In the next section we discuss the broadcast news domain in more detail, highlighting key challenges that serve as a context for our own work. Section 3 describes the components of a generic architecture for broadcast news content delivery, focusing in particular on certain key components including: video analysis for shot-boundary detection and keyframe extraction; news story segmentation and linking; user profiling and story recommendation; story visualization and interaction management. In Section 4 we outline the Físchlár-News system followed, in Section 5, by the results of an extended live-user study.

2. BACKGROUND CONTEXT

Recent years have seen dramatic changes in the TV sector on a number of fronts. The advent of Digital TV services has offered consumers a greater range of channels and programming content in addition to a host of new interactive services. A typical user might have up to 200 channels of content available per day, for example, offering in excess of 20,000 hours of content per week if we assume an average of 15 hours programming per day per channel. At the same time new types of devices have become available for providing access to this content, with the TV and video recorder being complemented by personal video recorders (PVRs) such as TiVo and PC-based devices such as WinTV. These provide consumers with access to potentially hundreds of hours of video content at any given time and, as a result, consumers are faced with new challenges when it comes to accessing the right content at the right time, quickly and easily. It is one thing for a user to be able to quickly locate a particular past episode of the new season of "24", which their PVR has been reliably recording over the past 3 months, but it is an entirely different matter when it comes to helping this user to quickly locate the particular scene where the show's maverick agent, Jack Bauer, discovers that he is once again forced to choose between his country and his family. How might our user locate this specific scene from the 400 or so minutes of content recorded so far? Fast-forwarding through hours of content is unlikely to be satisfactory and we believe that a combination of interactive search, scene-based navigation and intelligent recommendation is required.

Transforming broadcast TV into a content repository that can be searched, browsed and replayed is far from a trivial task. Raw video content must be analysed, segmented and indexed in order to identify meaningful units of content that can be separately recommended to, and reconstituted for, any given user, according to their personal preferences and information priorities. To meet this challenge requires bringing together state-of-the-art techniques from disparate areas of research

including signal processing (video and audio analysis), information retrieval, topic tracking, HCI, user profiling and recommender systems, to name but a few.

Recent advances reported in the literature serve to highlight the significant progress that has been made in many of these areas in recent times. For example, information retrieval researchers have begun to move beyond the world of text-based retrieval and into image and video retrieval scenarios. The annual TRECVideo activity is an excellent example of the critical mass that has developed in the area of video retrieval research [Smeaton et al. 2004a]. Each year TRECVideo sees dozens of research groups from around the world benchmarking their video retrieval systems on a controlled video corpus. These systems provide a showcase for the latest image analysis, feature extraction, segmentation, and video indexing and retrieval techniques. Indeed one of the TRECVideo tasks concerns the evaluation of interactive search systems operating in the broadcast news domain, with systems evaluated in terms of their ability to successfully locate relevant clips from a standardised news archive. The task also allows researchers to explore a wide variety of interface and interaction paradigms and progress to date in the news domain points optimistically towards a range of future applications scenarios including next-generation home-video and photo management systems and CCTV (Closed-Circuit Television) retrieval systems.

The TV news domain is the main focus of the evaluation study described in this paper. Specifically, we use the Físchlár-News system that allows its users to access daily broadcast TV news content through their web browsers. Físchlár-News incorporates a variety of fully-automatic, content-based video analysis and indexing techniques to provide users with a range of enhanced browsing and search capabilities. Físchlár-News has been deployed within the Dublin City University campus for the past 36 months and is regularly used by students and staff for study, research and entertainment purposes. This deployment has provided a fertile environment for the large-scale testing and evaluation of the system's various features, especially from a performance and usability perspective. In this paper we describe an in-depth analysis of the usage of Físchlár-News by a selection of 16 regular users over a 7-week period, in order better understand the way in which real users would interact with the system with a view to evaluating its usability and performance over time.

3. COMPONENTS OF A FULLY-AUTOMATIC, CONTENT-BASED TV NEWS DELIVERY SERVICE

A key objective of our research in recent years was the development of Físchlár-News, a scalable broadcast news retrieval platform that incorporates automatic segmentation, indexing, retrieval and recommendation features to provide end-users with fast and efficient access to timely news stories via a desktop or mobile interfaces. Although there are a range of similar projects elsewhere, none appears to offer the same breadth and depth of features. For example, many researchers focus on particular aspects of processing TV news, such as story segmentation [Arlandis et al. 2005], content indexing, video retrieval [Duygulu et al. 2004; Amir et al. 2004] or story recommendation [Billsus and Pazzani 1999], rather than on the development of a unified new retrieval platform.

For the purpose of our research we propose that such a unified platform should include the following key components, combined according to the architecture presented in Figure 1.

1. *Shot Boundary Detection & Keyframe Extraction* – to segment broadcast news into individual camera shots, enabling further structuring of the news content and other operations
2. *News Story Segmentation & Linking* – to identify news story units in the broadcast news and relate them to other identified stories by their similarity
3. *News Story Recommendation / Alerting* – to automatically highlight noteworthy news stories to individual users by recognising each user's topical interests
4. *News Information Visualisation & Interaction Provision* – to present news stories and their inter-relationships in a way that is understandable and easy to interact with
5. *Integration & Deployment* – to combine all the above components and be robust enough to be used by a large number of users at any time

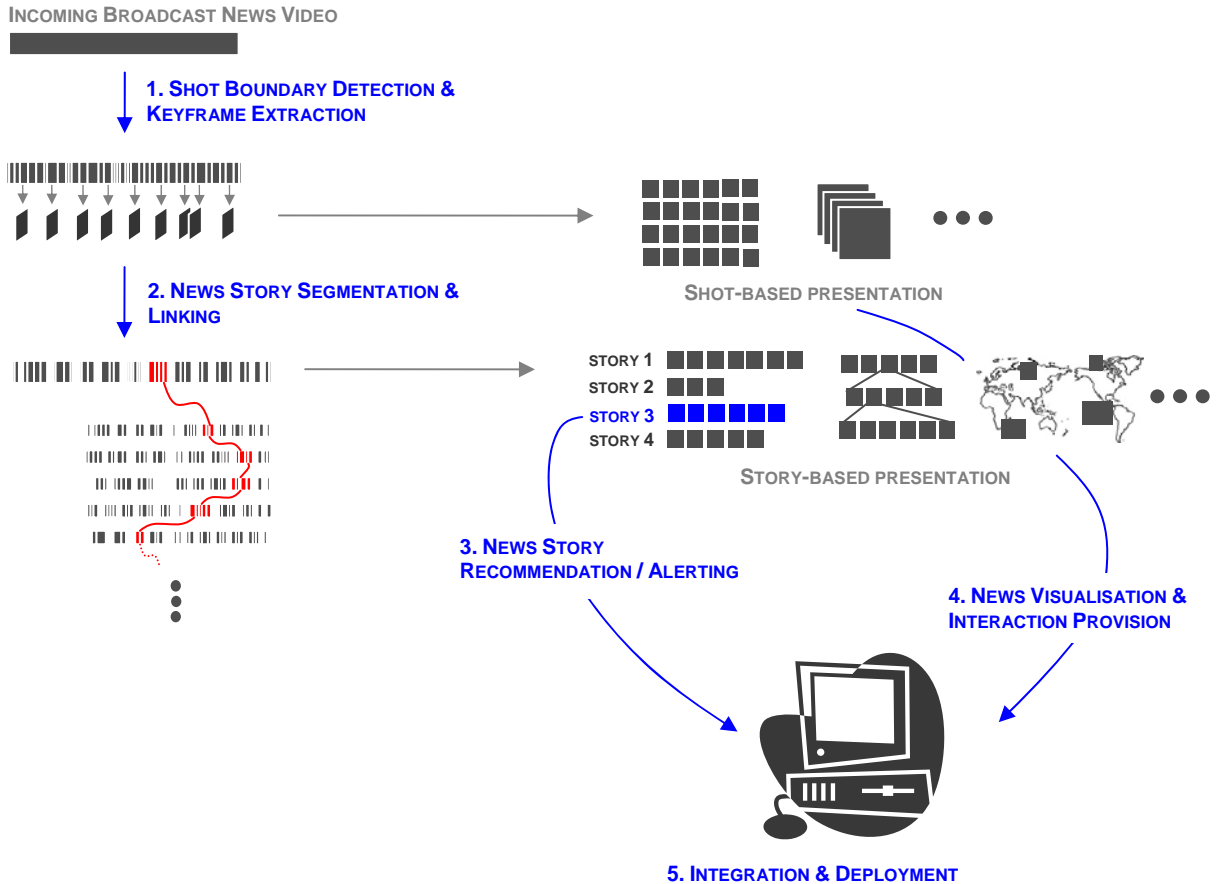


Fig. 1. Necessary components of a deployable, content-based, fully-automatic news video delivery service

The current state of video retrieval research sees these components at different levels of technological maturity. For example, under certain conditions the shot boundary detection problem can be considered a “solved problem” – hard-cut boundaries can be detected with a precision of 95%. Also recent results on the more difficult task of news story segmentation are reporting very promising results [Kraaij *et al.* 2004]. In this section we review the state-of-the-art in each of these five main components, before describing the Físchlár-News system in detail in the section that follows.

3.1 Shot Boundary Detection & Keyframe Extraction

Shot boundary detection (SBD) and keyframe extraction are essential precursors to video indexing and retrieval, and in this section we briefly review the current state-of-the-art; for a more detailed exploration see Lee and Smeaton [2002]. A video *shot* refers to the basic unit of retrieval – a short, coherent video sequence that serves as the starting point for the semantic analysis and structuring of content – and relates directly to camera shooting boundaries within the video sequence. Most SBD approaches involve measuring the visual similarity between adjacent or near-adjacent video frames and if these are visually similar, within some threshold value, then it is likely that they belong to the same shot. Camera and/or object motion will cause adjacent frames to be slightly dissimilar while shot cuts will generally cause a noticeable increase in this dissimilarity. Sudden changes in shots or hard-cuts are relatively easy to identify. For example, in TRECVID 2004, 17 of the 33 participating groups evaluated SBD functionality and many reported precision levels in excess of 90% for the detection of hard-cut boundaries [Kraaij *et al.* 2004]. More challenging shot boundaries include those that are masked by common shot transition effects such as dissolves, fades or wipes, and for these more difficult tasks precision levels of 70-80% were reported.

Once a video sequence has been segmented into a sequence of shots, representative keyframes can be extracted for each shot as keyframes serve as a visual shot summary. What constitutes a *good* keyframe and how to recognise a good keyframe remains an open question and current techniques generally rely on a variety of heuristic methods. For example, the simplest approach often involves selecting the first, middle, or last frame in the shot. More sophisticated techniques, however, can take account of in-shot camera movement by selecting the frame where camera movement stops.

Reliable and accurate shot boundary detection and keyframe extraction techniques have led to the development of a number of systems that focus on providing users with more effective sequence navigation and shot browsing features, as opposed to the full gamut of video retrieval features contemplated above. For example, Zhang *et al.* [1995] describe a system in which shot-level keyframes are hierarchically presented to allow the user to interactively drill up and down through large amounts of video data. Físchlár-TV, an earlier version of the Físchlár-News system which is the focus of this paper [Lee and Smeaton 2002] incorporates shot-level keyframe presentation methods with navigation features to allow browsing of broadcast TV content. Boreczky *et al.* [2000] propose an interactive comic book style presentation of keyframes in which the size of a keyframe is varied according to its level of importance. In turn, selecting a keyframe results in the presentation of a new set of related keyframes in a similar format, thus allowing the user to access different levels of keyframe detail. It is also worth noting that a number of commercial systems have utilised SBD and keyframe browsing techniques; e.g., VideoLogger (<http://www.virage.com>) and Screening Room (<http://www.excalib.com>).

3.2 News Story Segmentation & Linking

Identifying shot boundaries and keyframes is just the start of the analysis task when it comes to effectively utilising structure within news broadcasts. Ultimately the goal is to segment an individual broadcast into a set of constituent stories and then to link these stories (both within and between broadcasts) according to their topics. Automatically segmenting broadcasts into individual stories is more challenging than shot boundary detection. While an individual story may correspond to an individual shot sequence, in which case the shot boundaries will correspond to the story boundaries, more often than not individual stories will span a number of shots.

When it comes to segmenting news broadcasts we can profitably take advantage of the fact that most news broadcasts follow a generic well-defined structure. For example, a broadcast news programme will usually begin with a short summary of the day's headlines, followed by a sequence of individual stories. In turn, each story will begin with the anchor person introducing the story before moving to an on-site reporter or perhaps a studio interview. The anchor may wrap-up before moving on to the next story and the camera view of the anchor person is likely to be consistent across the broadcast with a standard shot change to the on-site reporter etc. In other words, there is often a well-defined grammar behind the structure of a news broadcast, and while this grammar will vary from broadcaster to broadcaster (the ABC news structure may differ from the BBC's for example), knowledge of the generic structure can greatly aid the story segmentation task. For example, knowing that each story will be introduced by the news anchor means that we can detect the start of a news story by recognising a shot which contains the anchor; and knowing that the anchor shot is always taken from a particular viewpoint (with a consistent background scene) can usefully constrain this "anchor detection" task. This approach will work well when the story boundary is aligned with a specific type of shot boundary (the anchor shot) but breaks down when the story boundary occurs within a shot sequence. However, in this case it may be possible to take advantage of news transcript text and apply linguistic techniques to detect within-shot story boundaries. For example the work of [Merlino *et al.* 1997] looks at the correlation of audio/visual cues, such as blank frames, and closed-caption analysis. In contrast [Gauch *et al.* 1999] proposed an initial story segmentation based on SBD followed by a merging of shots into coherent story segments based on audio energy-level analysis techniques. Others have looked at alternative and combination techniques for shot merging. For instance, [Haggerty *et al.* 2003; Pickering *et al.* 2003; Ide *et al.* 2003] all focus on the use of text-based analysis and topic-distance methods using closed-caption information whereas [Hauptmann and Witbrock 1998; Christel *et al.* 2002] used a combination of face detection, optical character recognition, speech recognition and close-caption analysis methods.

In TRECVID 2004, 8 participating groups worked on a story segmentation task and evaluated their systems using 64 hours of broadcast news content, containing 3,105 ground-truth stories (stories that we know as correctly segmented so that we can evaluate a system's performance against it). Segmentation precision ranged between 10% and 83% highlighting the difficulty of the task. Various techniques that have been tried in TRECVID are summarised in [Chua *et al.* 2004] and [Kraaij *et al.* 2004]. How important the accuracy of story segmentation is will depend on the way a system utilises the story units for supporting story-based searching/browsing features to its users. For example, if the primary way of accessing news stories is by sequential story navigation by one story after the next in a visually abstracted format, seeing an end part of a previous story at the beginning part of the next story may not be a very big problem to a user and she may tolerate and get around it in her understanding of the news content. On the other hand, if the primary way of accessing is by searching and a story is moved away from the rest of its broadcast sequence and presented in isolation, then the importance of correct segmentation becomes more important. An incorrectly segmented story retrieved from a user's search where the user expects an anchorperson to appear and start introducing the usual story summary but it actually starts playing somewhere in the middle, will result in frustration when the user finds out that the playback of a story finished in the middle and there is no way to find the lost half to continue watching the rest. A few seconds of lost keyframes or non-relevant keyframes might not be critical for a user browsing a large number of keyframes within a story,

but an interface that presents different stories retrieved from different broadcasts will become intolerable if the user starts noticing they do not start or finish at the correct story boundaries.

Once a broadcast has been segmented into separate stories, these stories should be linked to related stories in order to provide the type of added functionality that is likely to be valuable to end users. Ultimately, users are likely to be interested in tracking related stories over a period of time. Such features are now commonplace on many text-based news Web sites, but often the story linking is achieved through a mixture of manual and semi-automatic methods. Providing similar features in the context of a fully automatic video-based system introduces significant challenges. The story linking task is related to the more traditional text-based topic detection and tracking (TDT) tasks (see [Allan *et al.* 1998]), which involve the use of a variety of linguistic techniques to perform topic detection and tracking across a corpus of documents. For example the TDT-2 activity [Allan *et al.* 1998] applied linguistic analysis techniques, such as lexical chaining [Morris and Hirst 1991; Pickering *et al.* 2003] to a standard corpus of 60,000 ground-truth news stories. Video news story segmentation represents an interesting variation on the TDT tasks. In the case where only a pure video and audio stream is available it is, in many ways a more challenging problem than text-based TDT, mainly because the current generation of image and audio analysis techniques do not match the performance of text-based methods. However, it is interesting to note that in many situations the video and audio streams may be accompanied by textual transcripts (e.g. closed-caption text). In these situations, video, audio and text analysis techniques can be combined to good effect [Kraaij *et al.* 2004].

3.3 News Story Recommendation / Alerting

The ability to accurately segment and index individual news stories quickly leads to an explosion of time-sensitive content, and alerting users to this content is yet another challenge that should be addressed. In recent years, the traditional query-based search engine approach – whereby users search for content items by providing explicit queries – has been supplemented by new forms of information access. In particular, the field of recommender systems combines ideas from information retrieval and filtering, artificial intelligence and user modeling in order to provide users with a more proactive and intelligent form of information access. In short, recommender systems are designed to respond to an individual user's personal information preferences by monitoring their information accesses over time in order to develop comprehensive models of their tastes and preferences. In the news domain this implies monitoring their news consumption in order to track the stories and topics that a particular user is most interested in during a given period of time, and then using this information to prioritise similar on-topic stories for the future.

Generally speaking there are two basic approaches to recommending information items – be it news stories, TV programmes, books or music – to a user: *content-based* and *collaborative* recommendation. The former is grounded in more traditional approaches to content analysis and retrieval and relies on the availability of rich content descriptions for the items that are to be recommended. For instance, in the news domain this would typically require individual stories to be associated with meta-data descriptions to classify their contents according to such features as *news category* (e.g. sports, politics, business etc), *relevant parties* (i.e., the individuals involved), *time*, *date* and *location* information etc. Individual user profiles would be represented using similar features and the task of identifying a relevant news story to a specific user then becomes one of matching stories to profiles in order to identify the top-matching stories for recommendation. While such content-based approaches have the potential to work well, they do rely on comprehensive content descriptions which may not be always available. Moreover, these methods have a tendency to limit recommendation diversity as they continue to recommend stories that are always similar to those that the user has liked in the past, thus precluding the sort of serendipitous suggestions that the average user might come to expect of an intelligent news assistant.

Collaborative recommendation techniques, or *collaborative filtering* techniques represent an alternative recommendation strategy that does away with the need for item-based content descriptions. Instead user profiles are made up of *item ratings*, which may have been explicitly provided by the users or implicitly learned by analysing their behaviour. For example, a profile might indicate that a particular user likes news stories A, B and C (all of which have been rated highly) but dislikes stories X, Y and Z (all of which have been rated poorly). Then, when it comes to recommending new stories to a target user, a collaborative filtering system will first identify a set of *recommendation partners* – users that have similar rating patterns to the target – and then prioritise recent stories that have been enjoyed by these partners but not yet been seen by the target user. The collaborative filtering approach has proved successful in a variety of domains including movie recommendation [Cosley *et al.* 2003], TV programme recommendation [O'Sullivan *et al.* 2004], and even joke recommendation [Goldberg *et al.* 2001], and similar approaches have enjoyed some considerable commercial success in the guise of Amazon's basic recommendation system.

There has been some related recommendation work in the news domain. For example, the NewsDude news filtering system is a recommender system that suggests text-based news stories that the user might like to read [Billsus and Pazzani 1999] by using supervised machine learning techniques to induce a classifier that can discriminate between items likely to

be of interest to the user and those likely to be uninteresting. Alternatively, in the broadcast news domain, the underlying concept of personalising news presentation has been addressed in a limited capacity by some news segmentation systems [Merlino *et al.* 1997; Faudemay *et al.* 1998; Haggerty *et al.* 2003]. For example, [Merialdo *et al.* 1999] looks at the idea of dynamically generating a storyboard of selected news stories from broadcast TV news for individual users by analysing the news story contents relative to learned user profiles [Merialdo *et al.* 1999].

In the context of our own research, and the focus on this paper, automatic news story recommendation is seen as a core component of our news delivery framework. In later sections we will describe how collaborative recommendation techniques have been integrated into the Físchlár-News system, and while the true utility of these recommendation components can only be fully appreciated in the context of a large-scale user study, our preliminary results do highlight their potential value. We have also incorporated an element of content-based recommendations through the automatic generation of content links from each story to a set of related stories. In the true sense this isn't a recommendation system since it is based upon a *single* story recommendation rather than recommendations from a *set* of stories, but it does provide useful links from each story to others.

3.4 News Visualisation and Interaction Provisioning

The manner in which news stories are presented to users is another important research issue that is likely to impact on the usability and utility of any interactive news information system. For example, collections of news stories or events must be presented in such a way as to facilitate browsing and threading of story content. For the most part, the current generation of news websites adopt the familiar newspaper layout for presenting news stories, with the screen typically divided into areas filled with headlines and brief descriptions with some photos, and links to the detailed article on each story. While this approach has the advantage of familiarity it comes at a functionality cost when it comes to the provision of more sophisticated modes of interaction and browsing. As an alternative, [Swan and Allan 2000] describe a more novel approach to presentation which employs the concept of an explicit timeline for the presentation of related stories. [Frey *et al.* 2001] focus on the importance of story topics and describe a clustering technique for visualising clusters of topically related stories in a way that also highlights the frequency of new content (breaking stories) within each topic cluster. The BreakingStory system [Fitzpatrick *et al.* 2003] focuses on the dynamic nature of news content more explicitly and presents a visual summary of evolving news content by analysing term frequency information from more than 100 news sites on a daily basis.

In general, the ability to visualise trends within the story-space of current news is likely to be a very important feature in any interactive news system. [Ide *et al.* 2004] describe the Topic Browser system, which proposes a user-interface that is well-adapted to the browsing of threaded stories along a timeline. The Interactive Map [Christel *et al.* 2000] emphasises locality rather than time by presenting evolving stories on a geographic map. Time and locality are combined to good effect in the work of [Christel *et al.* 2002] which presents news content in the form of a dynamically assembled video collage, a summary of images, videos and text of news content presented according to a map and timeline for each user query. Finally, the work of [Rennison 1994] describes the Galaxy of News experiment, which visualises news themes in a highly interactive thematic space, where a user can navigate towards or away from a particular theme in a highly interactive fashion. This type of interaction paradigm highlights the physicality of the news information space and may provide a powerful and intuitive navigation facility.

The above examples represent the state of the art in news content visualisation, and although they are currently implemented as demonstration systems or concepts they nevertheless serve as useful indicators of what is to come in this important area. Within our own work we will describe a less ambitious approach to news story visualisation, but one that can be deployed as part of a large-scale application.

3.5 Integration and Deployment

From a deployment perspective the Web is the most obvious deployment platform for our interactive news service, although as digital TV technology matures it is likely to provide an alternative platform with many advantages when it comes to issues such as video delivery, for example. Of course Web-based news services are commonplace and range from the traditional human-edited services to more recent fully-automatic services such as Newsblaster [McKeown *et al.* 2002] and Google News (<http://news.google.ie>), both of which generate news summaries automatically from a wide variety of original sources. These web services are predominantly text-based and the prospect of a video-based news service introduces an entirely new level of technical challenge when it comes to the indexing and retrieval of video content. In addition, a more mature level of technical integration is required if existing streaming services are to prove as easy to use as existing text-based browsing services; today, many users continue to be frustrated by bandwidth-related delays and configuration problems when it comes to viewing video content through their browsers.

Within the Centre for Digital Video Processing at Dublin City University, we have been engaged in a programme of basic research and systems integration and deployment to provide a variety of prototype video-based browsing systems to a wide range of users. Físchlár-TV [Lee and Smeaton 2002] and Físchlár-Nursing [Gurrin *et al.* 2004] are two early examples of this work. The former provides students with access to recording and playback services, and a host of advanced browsing services, for general broadcast TV content. The latter focuses on the provision of video-based training material for nursing students with a similar level of information browsing and access. Both of these deployments have served as valuable experimental testbeds for the design, development and deployment of large-scale video-based systems. Físchlár-News, the focus of this paper, represents our largest deployment and has the advantage of a range of next-generation search, browsing and recommendation features. In the following sections we will describe this system in detail along with the results of a recent live-user evaluation study which has provided valuable feedback regarding the views, opinions and expectations of these users.

4. FÍSCHLÁR-NEWS

Físchlár-News is one of the family of Físchlár Digital Video Systems built and deployed since 1999, and it shares its system architecture with Físchlár-TV [Lee and Smeaton 2002], Físchlár-Nursing [Gurrin *et al.* 2004], and various Físchlár systems developed for TRECVid [Gaughan *et al.* 2004; Browne *et al.* 2003; Cooke *et al.* 2004]. To give an idea of the usage of different Físchlár systems, during 29 April 2003 – 23 June 2004 (1 year and 2 months), a total of 16,323 visits were made by 2,466 unique registered users, of which 4,739 visits by 364 unique registered users accessed Físchlár-News.

Físchlár-News captures the daily 9 o'clock main evening news from the Irish national channel RTÉ1 (approximately 20-30 minutes long) into MPEG-1 format, along with a Closed-Caption transcript that comes via an analogue teletext signal. Once the video and Closed-Captions are captured, Físchlár-News processes them by analysing content, and here we briefly describe this process in terms of five component technologies mentioned previously.

Shot boundary detection & keyframe extraction – this has been an enabling basic technique of all Físchlár systems and we have developed different algorithms for SBD including colour histogram based, edge detection based and a technique based on comparisons among the MPEG-1 stream macroblocks, as well as combinations of these methods in order to enhance detection accuracy [O'Toole *et al.* 1999; Browne *et al.* 2000]. Físchlár-News uses a colour histogram based method for SBD, following which keyframes are extracted by calculating the average colour histogram across each entire shot and selecting the one frame from the shot that is most similar to this.

News story segmentation & linking – Físchlár-News uses face detection, shot clustering based on content similarity, the length of each shot, and an activity measure for each shot, as input to a SVM (Support Vector Machine) which combines these evidences to result in detection of news story boundaries. The method is not tailored to a particular broadcast convention and is thus generalisable across different news broadcasters, and we have tested this method on CNN and ABC broadcast news and this worked to an accuracy level similar to the segmentation of RTÉ1 news [O'Hare *et al.* 2004]. Because the story segmentation is never 100% accurate, and because accurate segmentation is critical to the whole system as discussed earlier, once the day's news has been segmented automatically we correct it manually if required. The correction points for our daily inspection are the boundaries of stories in which either two stories are merged as one, or one story segmented into two or three, which happens in the system. This is done to ensure our users are not hampered by an imperfect story segmentation and this is the only human intervention in the entire capture and analysis process.

News story recommendation / alerting – Físchlár-News uses the ClixSmart personalisation engine (<http://www.changingworlds.com>), previously applied to TV programme recommendations for users of Físchlár-TV [O'Sullivan *et al.* 2004]. Users of Físchlár-News explicitly indicate whether or not they like a particular story by rating it using a 5-scale thumbs-up and thumbs-down icons displayed beside each story. During the period 29 April – 23 June 2004, corresponding to the 7 weeks of the trial period of this user evaluation, the total number of story ratings from all 364 registered users was 17,191. Currently only explicit feedback is used to identify user preferences. The engine runs once a day early in the morning, using the previous day's user ratings to calculate and generate a list of recommended news stories for each user. In the Físchlár-News interface only twenty recommended stories are presented to user, with confidence determined by the personalisation engine. With the current number of users and the average number of ratings, it takes 2-3 days for a newly registered user to start receiving recommended news stories, although this also depends on that new user's amount of initial ratings.

News information visualisation & interaction provision – the Físchlár-News interface provides a calendar on the left side of the screen. Selecting one of the dates will display that day’s news stories on the right side of the screen. Each story is represented by an anchorperson with a small poster of the topic of the story, the duration of the story, the 5-scale thumbs-up and -down icon for story rating, and the first 2 lines of Closed-Caption text to provide a preview of the story. Wherever a news story is presented, this form of story representation is consistently used. A user can further browse within a story using a shot-level keyframe browsing interface or can simply playback the story. Search can be done by typing in words in the text box on the top left and the result will be displayed as a ranked list of matched stories on the right side of the screen, as shown in Figure 2.

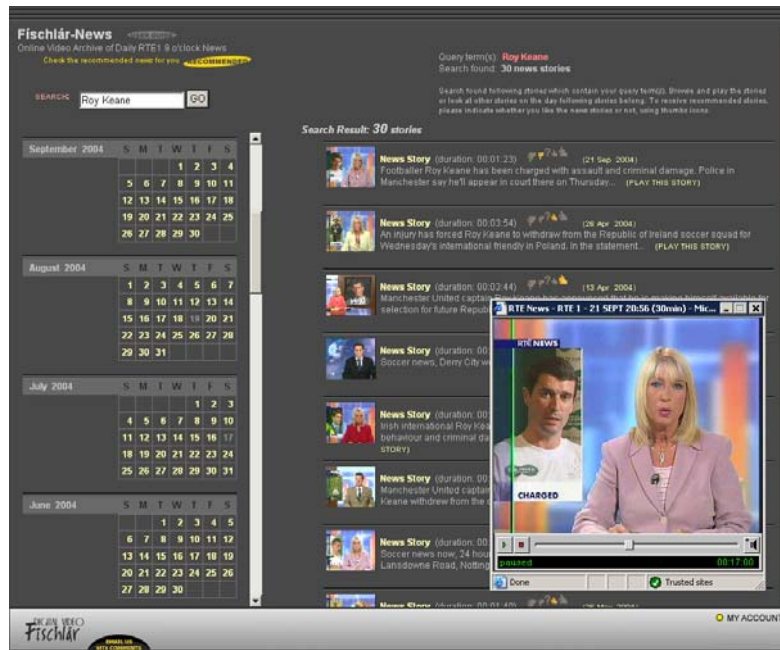


Fig. 2. The web-based user-interface to Físchlár-News, showing the result of a search with the query “Roy Keane” and playback of the first story.

Clicking on the bright yellow oval button labelled “Recommended” just above the search button will display a list of recommended stories on the right side of the screen for the currently logged in user. The user can play the story by clicking on the “Play this story” link, or play any part of the story in the story detail view (not shown in Figure 2) by clicking on a keyframe. A pop-up window will appear and start streamed playback from the beginning of a story or from the shot the user clicked (shown in Figure 2).

Integration & deployment – the Físchlár-News system architecture is based on an XML publishing framework. Any request from the user through the web interface will result in the collection of required pieces of data from different databases to be displayed as a result of the user request, generated internally as XML-formatted data, which is then transformed by a set of appropriate XSLT stylesheets into HTML. Because the transformation results in standard HTML and JavaScript, any user with a conventional web browser can access the system. Streamed playback is by an Oracle Video Server, supporting up to 200 simultaneous streams. The entire capture, analysis and indexing system is robust enough for operation without any human intervention, except, as mentioned above, for correcting the story segmentation result if it has failed to segment accurately.

5. USAGE STUDY OF FÍSCHLÁR-NEWS

With the Físchlár-News system deployed and used by our on-campus users, we conducted a usage study to assess real usage patterns of, and to find out about the level of user acceptance of, the system. The focus of the study was to gain a qualitative understanding of the realistic circumstances and contexts in which use of Físchlár-News occurs and the way it integrates with people’s daily activities. Our approach was highly qualitative and ethnographic, in line with similar

studies such as O'Brien *et al.* [1999], Lee [2000], Perry *et al.* [2001] and Petersen *et al.* [2002] in which the emphases were to understand real use, the development of use over time and the context in which a new technology is used when introduced.

5.1 Evaluation Methodology

We conducted a user evaluation of Físchlár-News with 21 participants during 24 March – 12 May (1 month and 3 weeks) in 2004, after which the data from 5 participants was discarded due to their low use of the system, resulting in a total of 16 participants' usage data. Full details of the way we have developed the evaluation methodology and procedure can be found in [Lee *et al.* 2005], and in this section we concisely explain the procedure taken.

Twenty five research students in Dublin City University were introduced to the system and features of the system were demonstrated to them in a hands-on way, including the search, browsing, linking and calendar features. Each person was invited to participate in a 1-month trial, after which a short briefing on the trial period was given. Twenty one people were recruited as participants and were each given a diary in which they were asked to write down a record of the circumstances of their use and any incidents or any thoughts they had during the period. As an incentive they were told that if they filled more than half of the diary by the end of the trial period they would receive a small gift. A pre-trial questionnaire was administered and from that point on it was completely up to the participant to access and make use of Físchlár-News. Interaction logging started 2 hours after this first visit was finished.

During the 1-month period, each participant received a daily email from Físchlár-News, with the previous night's top 3 news stories summarised and a link to the Físchlár-News website. Each participant was also informally visited 3 - 4 times during the trial to ensure they were not experiencing technical difficulties. Participants wrote diary entries during the trial period and all their interactions were automatically logged by the system.

One day before the end of the trial each participant was visited and a post-trial questionnaire was administered. Figure 3 shows the trial period for each of the participants, including 2 pilot participants both of whom started their trial on 4 March and ended 2 April on (1 week after the start of the main participants). The grey dotted lines indicate the 5 participants who used the system less than 3 times during the trial and were discarded from the analysis.




Fig. 3. The 1-month trial period by 21 participants

Once the post-trial questionnaire was collected, it was officially the end of the trial for that participant. Within 2 weeks after the end, a follow-up email questionnaire was administered to clarify some of the comments the participants wrote in their diaries and questionnaires and also to finally thank them for their participation and contribution. In summary, for each of the participants the following data was obtained:

- Pre-trial questionnaire (first visit)
- Diary (during 1-month trial)
- Interaction log data (during the trial)
- Post-trial questionnaire (at the end of the trial)
- Follow-up email questionnaire (1 week after the end of the trial)

5.2 Analysis and Discussion

At the end of the trial period, the 16 participants' usage data was collected. The participants were composed of 9 male and 7 female, and aged between 20 and 34. Most of them watched TV news 2-5 times a week. All of them were very familiar with the Internet, doing web searching and browsing daily and frequently checking news on a few specific news websites they like to use. To help understand our participants' use and opinions, which will be discussed in detail below, Table I

shows the participants' characteristics in terms of their TV watching, Internet browsing and their main source of news, as obtained from the pre-trial questionnaire and thus reported by participants themselves. The comment column in the table explains some additional characteristics we consider helpful in further understanding the circumstances for user comments later in the section.

To give a rough idea on the amount and kinds of data obtained through this study, Table II summarises the obtained interaction log data and diary comments. We found that our participants' diary comments and their comments in the post-trial questionnaire can overlap, but also tend to cover different aspects of use. While specific news topics and ideas for improvements were frequently mentioned in their diaries, their opinions on commonly used features such as the story overview were more often assessed in the post-trial questionnaire (where each feature was explicitly asked about). This needs to be kept in mind when looking at the numbers in Table II(b).

Table I. Participant characteristics for TV watching, Internet access and source of news

User ID (Gender)	TV watching	TV News watching	Main source of news	Web searching / browsing	Web searching / browsing for news	Search / browse news by	Comment
User 3 (M) Age: 25-29	30min-1hr /day	2-5 times /week	Internet – www.breakingnews.ie www.independent.ie news.google.com	Daily	Daily	Incident	Loves watching football and generally sports, except rugby; Listens to Internet radio news
User 4 (F) Age: 20-24	< 30min /day	1 time /week	Internet – www.unison.ie	Daily	2-5 times /week	Category (as in unison.ie)	
User 5 (F) Age: 20-29	30min-1hr /day	Daily	Internet - www.independent.ie www.independent.co.uk	Daily	1 time /week	Incident	Watches RTE1 news at 9pm nearly daily, and if home late, TV3 news at 11pm.
User 7 (M) Age: 30-34	30min-1hr /day	2-5 times /week	Internet - www.lefigaro.fr	Daily	Daily	Incident	French nationality; Uses news alert email services
User 8 (M) Age: 30-34	A few hours /week (mostly on weekend)	2-5 times /week	Internet - www.spiegel.de www.stern.de	Daily	Daily	Person / Incident / Category	German nationality
User 10 (F) Age: 20-24	1-2 hr /day	2-5 times /week	Reading newspaper; watching TV news; Internet – Irish Independent (www.independent.ie)	Daily	Daily	Category	Watches 9pm RTE1 news nearly daily; Uses Irish Independent online daily; Likes sports news in general
User 11 (F) Age: 20-24	1-2hr /day	2-5 times /week	Internet - news.bbc.co.uk	Daily	< 1 time /month	Incident	Uses BBC News website every morning before starting work; Interested in news her own town (townX)
User 12 (F) Age: 20-24	> 2 hr /day	Daily	Watching TV; Internet – www.irishexaminer.com, news.bbc.co.uk	Daily	Daily	Any current daily issues	Watches TV3 news every morning before going to work; Uses Internet news sites during work; Keen on current affairs
User 13 (M) Age: 20-24	30min-1hr /day	2-5 times /week	Reading newspaper, Watching TV news, Listening to radio news; Internet – www.independent.ie	Daily	< 1 time /month	Incident	Uses variety of news sources together, including Internet nearly daily; Mostly interested in up-to-date news
User 14 (M) Age: 30-34	A few hours /week (mostly on weekend)	2-5 times /week	Reading newspaper; Watching TV news; Internet – various sites	Daily	Daily	Person / Date / Incident / Category (in any way appropriate)	Highly technically-oriented user – generally very particular about his preferences on web browsing
User 16 (M) Age: 25-29	A few hours /week (mostly on weekend)	2-5 times /week	Internet – news.bbc.co.uk Irish Times online www.independent.ie	Daily	2-5 times /week	Mainly by headlines, sometimes clicking on related stories	Highly technically-oriented – web browser short cut keys and multiple windows while waiting for loading, etc.
User 18 (F) Age: 25-29	1-2hr /day	2-5 times /week	Reading newspaper; Internet – www.unison.ie	Daily	< 1 time /month	Category	Checks unison.ie website for various online news papers, nearly daily
User 19 (M) Age: 20-24	1-2hr /day	2-5 times /week	Watching TV news (Sky News almost daily)	Daily	2-5 times /week	Incident	Watches TV news daily for current affairs update, and online news sites for technology-related updates
User 20 (F) Age: 20-24	30min-1hr /day	2-5 times /week	Internet – www.unison.ie www.sky.com/skynews www.rte.ie	Daily	Daily	Category, followed by Incident	Says she loves news in general
User 21 (M) Age: 20-24	1-2hr /day	2-5 times /week	Internet - www.cnn.com news.bbc.co.uk www.dallasnews.com	2-5 times /week	2-5 times /week	Category	US nationality; Accesses news websites from his own homepage where he made links to a lot of news sites, mostly US general/sports news sites.
User 22 (M) Age: 25-29	1-2hr /day	1 time /week	Radio (especially sports news)	Daily	1 time /week	Person	Not a frequent consumer of any news, but main interest is sports; Listens to radio for general/sports news updates while working and at home

Table II. Summary of the access log data and diary-keeping by the participants

(a) Log data summary⁺

	User 3	User 4	User 5	User 7	User 8	User 10	User 11	User 12	User 13	User 14	User 16	User 18	User 19	User 20	User 21	User 22	Avg. (user)
Total # visits ⁺⁺	9	14	6	8	19	7	21	4	5	9	18	6	10	7	3	4	9.4
# Days visited (# Days story played)	9 (8)	13 (10)	6 (5)	7 (6)	14 (12)	6 (2)	17 (12)	4 (4)	5 (3)	8 (7)	12 (11)	6 (6)	7 (5)	6 (5)	3 (2)	4 (4)	7.9 (6.4)
Avg. time spent /visit (min)	17.9	9.6	17.1	6.3	16.7	8.6	13.3	3.2	5.5	9.4	8.5	14.5	29.9	33.8	12.1	15.7	13.9
# Story detail viewed	25	26	37	3	31	5	26	6	3	23	10	14	8	21	2	11	15.7
# Searched	5	7	9	0	3	3	6	2	0	2	1	0	17	11	0	6	4.5
Recommen- dation # story rating	62	15	8	2	2	6	34	0	0	0	3	40	51	51	0	53	20.4
# Recom. Page viewed	14	6	8	4	5	4	11	0	0	0	6	5	3	7	1	6	5
# Playback	22	24	15	7	71	13	46	5	4	27	29	28	44	19	6	16	23.5

⁺ This data was obtained by analysing system log data during the trial period and calculated with a commercial log data analysis tool.

⁺⁺ The log analysis setting was that a visit session is terminated when a user is inactive for more than 20 minutes.

(b) Diary summary*

	User 3	User 4	User 5	User 7	User 8	User 10	User 11	User 12	User 13	User 14	User 16	User 18	User 19	User 20	User 21	User 22	Avg. (user)	
# Pages written	11	20	9	3	4	14	18	3	2	1	7	10	11	16	5	8	8.9	
Expression of experience**	Positive	5	7	7	1	10	11		1	2	3	1	2	1	1	3	3.9	
	Negative	2	11	6	2	9	4	12	3	4		2	10	9	8	8	6	6.4
Features mentioned***	Story detail		3	2		2	2			1				1			1.9	
	Search	2	3	3		2	1	3	1			1		2	3		2.1	
	Recom.	4	4	1	1		1	7		1		3	6	1	1		4	2.8
	Playback	5	6	2	1	3	3	5	1		1	5	3	2	3	2	3	3
	Related stories	1	2	2		1						1	3	3	2	1	1	1.7
"Accessed because I missed TV news"		2	2			5	4	1			3	1		3		1	2.4	
"Heard about a news so accessed"	3		1				1				3			1			1.8	
Other reason for access / no access	5	5		1	1	2	5	1	1					3		1	2.5	
Up-to-date vs. Archive mentioned		1			1	2	4		1							1	1.7	
Coverage & other news source mentioned	7	6	1	1		7	7	2			1					1	3.7	
"Story breakdown overview useful"		1				4	3										2.7	
Improvement idea	1	7		2	1	1		1	2		1	3	3		10	3	2.9	
News topics mentioned	Sports	5	2	2		3	4		1		4			2		8	3.4	
	Politics	1	1	2					1			1	4	2			1.7	
	Crime and terrorism	3	3	1		1	1							7			2.7	
	Business, Tech, Ent.		1											2		1	1.3	
	Local issue****	1	1	1				3	1			2	1	1	11	1	2.3	
	Accidents	1						1				1	2		1		1.2	

* Diary summary was calculated by going through all entries in the 16 collected diaries and categorising each topic. A count was made whenever a user mentioned the feature or topic, regardless of whether s/he used it or not at that point. The count in the diary summary is only meant to serve as a rough guide of the kinds of topics covered by each of the participants.

** In Expression of experience, we counted as 'positive' when a user wrote down a sentence in a positive tone, e.g. mentioning a successful location of a story ("managed to find it and watched..."), expressing the usefulness of a feature ("I found this very useful...") or a positive emotion ("I am happy with the quality of..."). We counted as 'negative' when a sentence carried a negative tone, e.g. mentioning an unsuccessful task ("but all of it was a repeat of what I had seen, so I didn't play any of them"), an intention not to use a particular feature ("I don't think I would rely on this alone for some time yet") or a complaint about a feature ("Too many pictures and it is too small").

*** The counting of features was based on whenever a user mentioned s/he used it, both when a feature was the main topic of the entry, or when it was used to describe something else, or when s/he simply described that she used it.

**** Local issues mentioned are: indoor smoking ban in Ireland (2), Irish EU presidency (1), St. Patrick's Day parade (1), Beverly Flynn (Irish politician) expulsion from political party (3), Irish citizenship referendum (2), May Day disturbance (2), AIB Bank (1), LUAS tram trial in Dublin (1), murder incident in townX (2), Dublin bus crash accident (1), abandoned baby in a house (1), IRA matter (1), an Irish judge convicted of possession of child pornography (3), Dublin elevation (vectorial elevation event in Dublin city centre in Spring 2004) (1).

Figure 4 shows the access patterns of individual participants throughout their respective trial periods, with the x-axis showing the daily progression and the y-axis showing the number of access on that day.

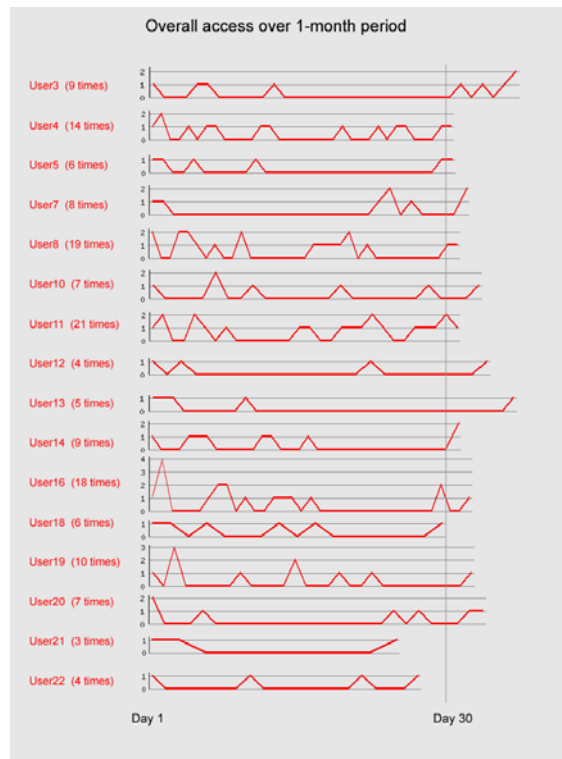


Fig. 4. Access throughout 1-month period (x-axis=day; y-axis=number of access)

In the rest of this section we present our findings in five categories: purpose of use, news archival vs. up-to-date service, news coverage, linking to related stories, and user expectations and system image.

5.2.1 PURPOSE OF USE

Each of our participants accessed Físchlár-News in his/her own time and for his/her own reasons, fitting in access to the system sometime during their own work hours, which makes this evaluation more realistic and meaningful than a laboratory-based user evaluation would permit. Strictly speaking, even though this situation might be considered artificial because we were asking users to participate in an evaluation, the nature of news delivery in Físchlár-News tended to blend itself easily to participants' daily routines, as one participant reacted:

"Initially I only used Físchlár because I was taking part in this evaluation. However, using it has become part of my routine, so I can see myself browsing through every week or so in the future, or maybe more often." [User11:Diary:28/04 – 28 days after the trial started]

Because the system is an archive of national TV news, there was a connection with participants' watching of TV. Most of our participants watched TV news at least 2-5 times a week (see Table I), and Físchlár-News was often used as a secondary support for TV news watching: if they missed news on TV on the previous evening or previous weekend, the following day or following week they would look for stories on Físchlár-News.

"I was late home last night and saw no news so I watched Físchlár news when I came in this morning." [User4:Diary:27/04]

"... there was a big game yesterday which my team lost. I didn't see the news last night and so was interested to catch RTE reports on the game." [User5:Diary:26/04]

"I didn't watch much news over the bank holiday weekend so I logged on to Físchlár and browsed the news stories for the 10th and 11th april. I found it very useful as I could quickly catch up on everything." [User10:Diary:12/04]

At least 9 participants mentioned in their diary, 6 of them more than once, that they accessed the system to catch up the TV news that they missed the previous night or during the weekend (see Table II (b)), and this was amongst the most frequent reasons users had for using the system.

On the other hand, because of watching TV news directly, there were occasions in which users did not feel the need to access the site the following day because they had watched the news on TV. The fact that Físchlár-News records and archives daily as opposed to TV in which one has to follow the fixed timetable of the broadcaster, a user can access the archive as many times as s/he wants, as a few participants commented:

"17:00 - Did not log on to Físchlár today as I didn't have time. I quickly glanced at the headlines of the Irish Independent online but did not get much information. It is good that I can use Físchlár tomorrow to catch up on today's news." [User10:Diary:14/04 – 15 days since start]

"Good to see news at any time of the day. News items can be watched several times. Older news items can be easily accessed." [User8:Post-Trial Q]

"It allows an access to video news at different times of the day. This makes people independent from the time tables of TV channels – it is video on demand, which goes towards more freedom and individual self-determination." [User7:Pre-Trial Q]

PLAYING NEWS VIDEO CLIPS

A big incentive for accessing Físchlár-News was the fact that participants could watch the news video on their desktops. As the interaction log data showed, streamed playback of news story clips was the most frequently used feature of Físchlár-News during the trial (total 376 times; average 23.5 times /person, see Table II (a)). In almost all cases, a participant who accessed Físchlár-News played video clips, often many times in one visit (compare # Days visited, # Days story played and # Playback in Table II(a)). Not playing a video clip while visiting the site was more of an exception, as one participant commented in her diary:

"Traffic was bad on the way into college so I was a bit later than planned and I have a busy day today. Because of this I decided to browse through the headlines on Físchlár, but not read any of the full text or play any of the stories." [User11:Diary:29/04]

The log data shows the above user played stories mostly 1-5 times whenever she visited the system (with an exception of playing 17 times on the 2nd day of the trial), and has used playback 12 days out of the total 17 days she accessed the system.

The use of playback was one of the main activities conducted throughout the 1-month period, unlike the usage of some other features such as the Closed-Caption based searching feature which was used as a once-off by most participants; this will be discussed in more detail later. The reason for this frequent use of story playback can be found in numerous questionnaires and diaries:

"...I had read about the story already (Iraq car bombing) but I wanted to see some of the video and interviews. It might seem strange to want to see the video of the bombings, but you get a much better idea from one video than all the newspapers descriptions..." [User3:Diary:22/04]

"My main interest was to get video footage to accompany news stories that I had read..." [User3:Post-Trial Q]

"...Then watched the story about the lights in the sky [Dublin Elevation for celebrating EU expansion]. I had heard about it but for a story like that it really is nice to see the pictures [and playing]..." [User22:Diary:23/04 – after playing news story 7 on 22 April about vectorial elevation lights in Dublin city centre]

"I did find it very good to be able to view [play] the highlights of the match" [User10:Diary:12/04]

"...Its nice to be able to view [play] the latest goals. But with most [other] news stories the pictures are not important..." [User22:Post-Trial Q]

The usefulness of being able to play news stories in the news archive is something that had been noted from participants' reactions at the time of introduction and pre-trial questionnaire, but after the trial we saw definite usage of this feature as opposed to what people *thought* they would find useful from their first impression. Indeed, playback is one feature that a newspaper or many of our participants' favourite commercial news websites (for example, BBC News and Irish Independent) did not feature as a facility for viewers, and thus being able to play a short self-contained news story in Físchlár-News was seen as a unique benefit and used frequently during the trial. Watching sports scenes in sports stories was the main interest of a few of our participants. Sports goal scenes are something where a lengthy textual description or

still photos cannot compete with a few seconds of video footage, and those interested in sports often used Físchlár-News to play highlights from the previous nights' games. Some participants comments are:

"Having the footage of the goals is great. I read a couple of soccer pages regularly, but they don't have video footage so this is handy." [User3:Diary: 17/04]

"Look [watched] at it for weekends goals - good..." [User13:Diary:05/04]

"[8 days after start] Checked my recommended news. Found a story about soccer in it. Watched the goals." [User22:Diary:15/04; sports fan only watching sports news]

While most users played to watch some particular stories or particular part of stories, some of our participants had a habit of playing news in the background while browsing other parts of Físchlár-News or working on something else, and just listening to the stories as the following participants explicitly mentioned:

"...often will leave the news running when I return to work and just listen... ..I've often minimised the video window and just listened." [User22:Post-Trial Q]

During informal visits and discussions during the trial, we noticed there were more cases of participants continuing to play video while doing something else, although not all of them noted this in their diaries or questionnaires. Sometimes this was not the participant's intention but because s/he did not expect Físchlár-News to continue playing beyond the end of a story. Currently when an end of a story is reached on playback, it continues on to the next story. Related to this, frequent use of playback throughout the 1-month trial also resulted in our participants wishing to have more elaborate and advanced features *on* the playback window, to make their viewing experience more enjoyable and efficient. Some of the more frequently voiced playback-related features our participants wanted include:

- To have buttons to jump between news stories
- To display Closed-Caption text scrolling at the bottom of the player, so that it is possible to play with the sound turned off
- To display text banner at the bottom of the player, on today's other breaking news
- To be able to play the headline of the day (usually half a minute at the beginning of a news broadcast)
- To display elapsed time (from the beginning of the story)

These enhancement ideas are mostly understandable, some of them a result of oversight at design time, some due to the difficulty of implementation, and some simply not conceived at the system development stage. Regarding our participants' wishes, by the end of the trial we had compiled a long wishlist of small and large enhancement/refinement ideas for each of the features provided by Físchlár-News. This could be a useful starting point for future development of similar systems, and is listed in Section 5.2.6. Considering the importance of playback, we can consider a more playback-centric browsing interface (such as Windows Media Player or RealNetworks RealPlayer), in which all other searching, browsing and recommendation features are secondary to playback.

SUMMARY

We found that the common purpose of our participants' access to the system was to look at TV news stories that they missed on the previous night or the previous weekend or they heard from somewhere else that they want to check up on. Playing and watching TV news stories was an important incentive for accessing the system, something that other news media (or news websites for that matter) does not offer. Some advanced features that could enhance their playing/watching experience were voiced and some of these will be incorporated in our future versions of the system. What advanced features in playback would be useful in the context of TV news video archive will be a subject of another, more focused user study perhaps starting with a laboratory-based experiment followed by a longer-term trial with such an interface.

5.2.2 NEWS ARCHIVAL VS. UP-TO-DATE SERVICE

Físchlár-News is a daily archive of the broadcast TV news with the facility to search, browse, play and trace related news stories in the archive. As Físchlár-News archives the daily 9pm news, the latest news one can see on most days is the *previous night's news stories*. Although this serves users interested in searching for some previously broadcast news or in looking at last night's missed TV news, its up-to-datedness cannot be compared with the commercial news websites where the update frequency is near-immediate. Indeed many of our participants frequently accessed news websites such as the Irish Independent website and the BBC News website, as seen in Table I.

The value of Físchlár-News as a daily news archive as opposed to an up-to-date news delivery service has been recognised as better than other news websites by our participants, as the following comments reveal:

"Físchlár-News website is more of an news archival search system as opposed to an uptodate online news system... ..as a simple, terse news archival system, it delivers exactly what it promises and does it well." [User14:Post-Trial Q]

"... the fact that every broadcast was archived meant that I didn't have to worry about missing anything." "...its archive of video stories and the ability to find related stories are impressive and very useful..." [User19:Post-Trial Q]

"Sky News, RTE.ie. They are more up to date. But Físchlár is better for going through stories + watching what bits you want." [User20:Post-Trial Q]

Because of the strength of the archival aspect of the system, some participants wanted Físchlár-News to archive other current affairs-related TV programmes they regularly watched on TV. Also, being able to *play* old news was considered a very important feature of Físchlár-News. On the other hand, our participants expected Físchlár-News to also be an up-to-date service as other commercial news websites, in which case they were often disappointed:

"I am listening to the radio over the internet – this gives me the current news in detail so I don't feel I would gain anything from using Físchlár to view yesterday news stories" [User10:Diary:6/4 – 7 days since start]

"... I'm interested in reading news as soon as it happens, rather than hearing about it the next day." [User11:Diary:08/04 – 8 days since start] "[After looking at Físchlár-News for Maradona story and others:] I will probably look at the BBC news website this morning as stories will probably have been updated during the night." [User11:Diary:20/04 – 20 days since start] "...As it is a news archive it does not provide update during the day, which is a disadvantage for me." [user11:Diary:28/04 – 28 days since start]

"How often does one look at old news" [User12:Pre-Trial Q] "Useful only to people who are interested in old news." [User12:Post-Trial Q]

Numerous complaints were made about this aspect of the system, the fact that its news update interval is 1 day regardless of whether or not there was breaking news. Considering the fact that all our participants were research students and their main interest in accessing news-related sites is to inform them about up-to-date news, it is understandable that Físchlár-News seems to be unable to satisfy them in this area.

THE USE OF SEARCHING

The low usage of Closed-Caption based searching is worth exploring as searching was done total only 72 times (average 4.5 times/participant) with 4 participants not using it even once during the trial period (see Table II (a)). Also many of these 72 searches were the same topic by the one participant experimenting with search results. For example, User19 did a total of 17 searches during his 1 month period, but 12 of them were during one session, trying different spellings and combinations of search terms. The interaction log data on the use of the search feature by each of the 16 participants over the 1-month period is shown in Figure 5.

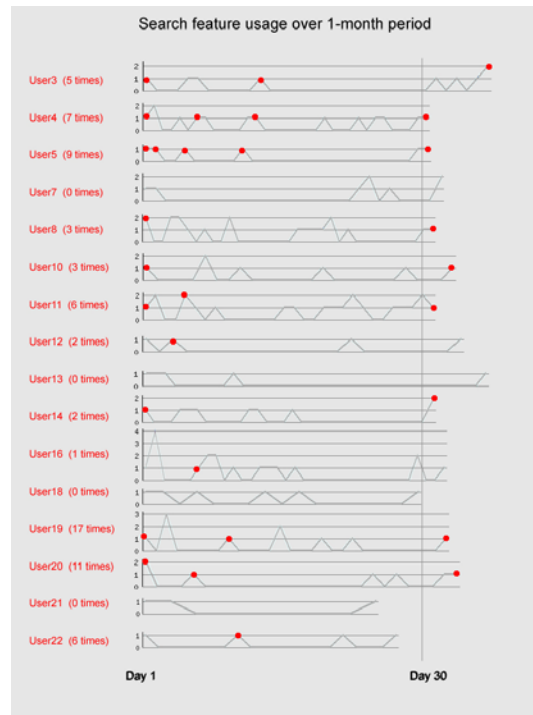


Fig. 5. The dates when search was used over 1-month period

The dots in Figure 5 are when each participant did searching in a session, and search could be multiple times in one session. Most of the times participants used searching were at the very beginning and at the very end, rarely in the middle of the trial. When they used it, they used a few times in bursts, with some variations of similar queries to test the search result. It was less of a phenomenon of making use of the feature for their own purposes, but partly as an experimental search to see results (just after the trial started), and partly searching on the last day of the trial when the post-trial questionnaire was administered.

Some of the reasons why searching was so underused can be found in the participants' comments. One is related to the search facility and search result: with the small amount of searching conducted there were relatively many complaints about the search result not being optimally presented, a summary of which is as follows:

- *Order of search result* – when a user clicks on the SEARCH button, the search result is presented as a list of stories ordered based on estimated relevance. Many participants wanted to see the order of stories chronologically (by date), as the following comments show:

"...it would have been useful to rank news stories by date. I wanted to view the latest Rugby news" [User4:Diary:25/03; searched 7 times total]

"...sort results by relevance OR chronologically, etc." [User19:Pre-Trial Q]

"When using the search function, it might be useful to permit other sortings of the matches other than the default 'by relevance'. For instance, I did a search for news stories on the recent situation in Iraq, but was instead presented with a mixture of stories from last year's conflict and the current al-Sadr standoff. An option to sort chronologically by ascending/descending date order would be invaluable in such situations, allowing easier tracking of the long-term development of current news items." [User19:Diary:19/04]

- *Search within a range of date, week, a page of results, etc.* – participants wanted to search within the area that is displayed currently or be able to specify the date/week range for searching. Currently search is done against all stories in the archive:

"Perhaps you should provide a 'range of time' option for the search" [User8:Diary:25/03]

"A 'search within page' checkbox would be nice for searching the news stories of a particular day... an 'advanced search' form which allowed users to search for a keyword in stories within a particular date range..." [User21:Diary:07/04]

In addition to the above, participants mentioned other specific wishes for enhanced search features such as allowing searching for a particular journalist or reporter, and allowing phrase search.

Although these shortcomings in the current search feature are part of the reason for low usage of search, probably a more fundamental reason can be traced from the following participants' comments:

"I didn't find the need to use this feature [searching] too often as I had been checking the system most days." [User11:Post-Trial Q; searched 6 times, 4 of which was 'townX' – she's from townX and wanted to check some murder cases happened there]

"... I was searching for some information on sports results and could not find a relative [relevant] story. I did find it in the end but through browsing myself and not through the search facility." [User10:Diary:12/04; searched 3 times]

"I heard about the train crash in North Korea on the radio this morning, on my way to college. I hadn't seen the news yesterday, so I used fischlar news to find out more. RTE led with the story so I didn't have to search." [User16:Diary:23/04; searched 1 time total]

In most cases, our participants *did not need* the search facility, mainly because they were mostly interested in recent news (such as yesterday or last weekend) thus the browsing by calendar feature was enough to satisfy their needs. Some were interested in searching for more related stories, but *related other latest stories* rather than the *older news stories about the same event* (more on the participants' expectations on the related stories will be discussed in Section 5.2.3 on news coverage). This point brought us back to the up-to-datedness of the system. On the whole, what our participants wanted was the latest news update service in which they could quickly get the latest updates during short gaps in their work, rather than a comprehensive news archive facility in which they can search through.

This leads us into consider the possibility of a more targeted approach in system development: instead of trying to provide an up-to-date service, we can target a system for a user group such as media or journalism students who need to efficiently search for video footage of news events. Had we recruited our participants from such a user group, probably we would not have received as many negative comments on the limited up-to-datedness of the system, participants would have used the system to help their study (rather than casual 10-minute news updates during lunch break), perhaps mainly using the search feature. If we are to satisfy users who want up-to-date news, the system should be strengthened in that direction in some way also. For example, in order to be able to constantly keep users up-to-date with new news, the system needs input from TV news video sources more frequently, perhaps an hourly processing and archiving of news video from dedicated news channels such as Sky One or CNBC News. Once strengthened in this way, the system should be able to cater for both kinds of needs, and we believe the core functionality of Físchlár-News should be delivered as a single system that can seamlessly support both needs. While this approach would not be too far off the direction of our current approach to the system-based research and development, more engineering in terms of scalability, processing speed and storage capacity would become an important concern and would required additional resources.

SUMMARY

Físchlár-News's one-news-per-day interval of update makes it less suitable for satisfying users at work who want quick updates of very recent news. The system's major values as easily browsable, searchable, and playable video on the web, will be more utilised when targeted at a user group who come to the system with a mindset similar to when a user comes to a Google search engine to search for something from a large archive of the web. A new system that has all of Físchlár-News's constituent technology with the aim of an up-to-date news service will need to focus on the technical scalability implications to more frequently process broadcast TV news from a dedicated news channel.

5.2.3 NEWS COVERAGE

Físchlár-News archives only the daily 9 o'clock news from the Irish national news broadcaster RTÉ1, thus in terms of news coverage it represents all the benefits and drawbacks of the RTÉ1 news. RTÉ1 news covers important events that happen both domestically and internationally, but tends to focus on the Irish national events both small and large. The value of Físchlár-News was seen when our participants were interested in local news, as the following participant mentioned in his diary:

"It was good to have the coverage of the smoking ban. This was better than the newspaper coverage because it had interviews in pubs with customers, I knew one of the pubs in Inchicore [a town in Dublin, Ireland]. This coverage is the kind I don't get in the Times, etc." [User3:Diary:29/03]

In 2004, the Irish government banned all smoking in public buildings including pubs, and this was a controversial issue debated extensively as the launch date approached. While other countries' news media also dealt with the Irish smoking ban, Ireland's own perspective and the practical details of the implementation and consequences would be something only an Irish broadcaster could truly portray. RTÉ1 news showed some representative people's views by interviewing owners and customers of local pubs and this was not found in other, more generally popular news websites such as BBC News or Irish Independent websites. Being able to play interview scenes with local people in them seemed to make playback even more valuable. For another example, User11's comment:

"I usually look up BBCNews.com through the day... BBC can often not be local enough for me. A murder is being investigated in the town I am from [townX], so I decided to browse [on Físchlár-News] last night's news to see... There was, and I watched the story, as well as rating it [positively]." [User11:Diary:16/04] "... I found it [Físchlár-News] useful... BBC News.com generally doesn't have news specific to Ireland." [User11:Post-Trial Q]

News stories showing an interview with worried factory workers at a factory closure announcement, the streets and houses shown in the murder investigation in townX, an Irish singer's preparation for the Eurovision Song Contest, shots from St. Patrick's Day parade in Cork, a landslide scene in a small town in Galway, a vectorial elevation event in the Dublin city centre all became valuable news clips because of their local value. The list of local news explicitly mentioned in our participants' diaries is at the bottom of Table II (b). Local sports news was also a valuable aspect of Físchlár-News's local news coverage:

"Heard late last night about Roy Keane [Irish football player] might be coming back to play for Ireland. When I got in this morning the first story on the news-highlights mail was about Roy Keane, which reminded me to go read about it." [User16:Diary:14/04]

"Checked last night's news on Roy Keane. Under related stories there was no more stories about Keane, so I did a further search. I found another story about Roy Keane's return to Ireland and watched that." [User22:Diary:15/04; played 14 April story, then 13 April story both on Roy Keane return to Ireland team]

"Heard on the radio that Mayo ladies team were being banned from football because of a dispute over jerseys. Checked the news [on Físchlár-News] to find out more, and watched a report about it." [User16:Diary:16/04; played 15 April story on Mayo team jersey dispute]

Thirteen of the 16 participants being Irish nationals, it is understandable that they value the Irish-related and local news content and make use of them. On the other hand, the Irish-centric news coverage was a point of dissatisfaction for some participants. An RTÉ1 9 o'clock news broadcast is usually about 20-30 minutes long and contains about 10-15 stories (no matter how many events actually happened on the day). This means that other potentially interesting news of the day (especially international news) are not mentioned in the broadcast, and the less important stories and especially sports news are at the mercy of available time, and are often omitted completely when there is not enough time to fit in due to other more important stories.

"Stuck to 9pm news only & Anne Doyle's [major anchorperson in RTÉ1 news] news can be a bit one tracked sometimes and I like a general overview of all national news... Much broader coverage on unison.ie [Irish online news website]" [User4:Post-Trial Q]

"It would be good to be able to get different coverages by different news channel for the same story..." [User7:Diary:02/04]

The perceived limitation on news coverage by some of our participants (along with the limited up-to-datedness as discussed already) was severe, and turned their main news source back to the BBC News and Irish Independent websites, or at best using Físchlár-News only as complementary news support along with the other news sources. As a consequence of perceiving this limited news coverage, our participants wished Físchlár-News archived other TV news programmes or news websites:

"...should look at ways of expansion of news, maybe joining news stories from different broadcasters..." [User13:Pre-Trial Q]

"... if possible can you record all 'prime-time' and 'Questions & Answers' show [other current affair programmes] – from an archival point of view, they are a lot more useful than football matches!... Needless to say, the more news sources Físchlár can incorporate the better." [User14:Post-Trial Q]

"...some integration with online news resources would be useful. ...it could utilise its web-based nature to a greater extent by linking each story to other online news resources, offering alternative viewpoints, editorials, etc." [User19:Post-Trial Q]

We can identify a tension with respect to news coverage with users sometimes wanting local news only, and the same users also wanting a broader coverage at other times. Expanding a news archive system so that it integrates across news sources would seem to address this issue if the actual news stories presented can then be categorised into "local" and "broader."

SUMMARY

Although Físchlár-News adds value to the RTÉ1 TV news in terms of *access*, it inherits all the benefits and limitations of the RTÉ1 broadcast news in terms of its *contents*. It indicates that a database with poor quality content will not be of much use even with highly refined retrieval features. Our participants' comments revealed that the perceived value and the usefulness of the system was heavily influenced by the system's news video content – positively when local news showed its value to our local participants, and negatively when coverage was limited to national issues. When there are so many useful resources and services on the web nowadays (albeit all manually indexed and authored), the true value of the automatic video indexing processes does not necessarily stand up unless it leverages what people can so easily avail of from the web. Incorporating multiple news sources (TV broadcast videos and news websites' information) would help ease this problem, but this will involve not only technical challenges but also copyright and social issues.

5.2.4 LINKING TO RELATED STORIES

Once a user views a news story's detailed view, they are also presented with a list of the 10 most related stories to the currently viewed story. From there the user can further jump to one of these 10 related stories to follow the threading of related news stories. As automatically linking news stories was one of the major features of Físchlár-News, we were interested to find more about our participants' usage and view of this feature. Related to the issue of broader news coverage which our participants wished, some participants wanted to see further "related" links to other resources for each presented story. In general we can differentiate "related news stories" into three different categories: (i) stories about the same event but appearing on different dates in the archive; (ii) stories about the same event but from different sources; (iii) similar but actually different stories about different events.

(i) *News stories about the same event that appeared on different dates* - this group of stories can be "linked" to each other, forming a thread of an event development over time. Visualisation of news story threading [Ide *et al.* 2003] is based on this category of related stories. Físchlár-News automatically generates the list of Related Stories based on the similarity between Closed-Captions of the stories in the whole archive, thus as long as two stories contain enough of same terms they are related. This includes stories in the past as well as today's other stories. Many commercial news websites that focus on up-to-date delivery of news do not have this concept of threading older stories with recent ones, because they do not archive previously reported stories. Some of our participants, when looking at the "Related Stories" of Físchlár-News, did not want to see older stories about the same event, because mostly they were already familiar with the background of the story. This is understandable considering most of our participants generally keep up-to-date with current news by watching TV news and checking up web-based news services during daytime as Table I showed. These participants thus complained that the list on the Related Stories was not up-to-date:

"... I was just interested in getting the most up to date news. Generally related news seemed to be older, so I had less interest in it."
[User11:Follow-up Q]

"[seldom used related stories because] I was interested in the most recent part of the story." [User12:Follow-up Q]

"I think these less useful, as they tend to link to old stories, compared to the BBC site, which tends to link to key concepts mentioned in the story." [User16:Follow-up Q]

On the other hand, being able to track a particular event from today back to the previous days was viewed positively and used by some participants:

"...Also the search method for related stories is very effective. I looked a the news story re. an abandoned baby & every case that I can remember in the last year was reported in related stories – very impressive!" [User4:Diary:19/04]

"I also noticed a story about Beverley Flynn's ejection from Fianna Fail. Since I did not know the full details of why she had been ejected, I used the 'Related stories' feature to track the events leading up to it; it proved very useful..." [User19:Diary:06/05]

"It's great to be able to miss a couple of days worth of news and not have to trawl through each day to find what went on. If something is still topical that happened a few days ago, you can easily pull up the previous reports... Generally [when I used it] it was if I read a story and hadn't previously heard much about this topic or felt I needed more background to fully understand what was going on."
[User4:Follow-up Q]

(ii) *Same news stories from different sources* - this refers to a group of stories about the same unique event, but produced by different broadcasters and reporters, and thus likely to have different points of view and different footage. Físchlár-News currently includes only one source, RTÉ1 news, and thus does not support this type of related stories. This is

related to some of our participants' earlier comments on broadening the coverage by incorporating other sources (Section 5.2.3), as well as the following comments:

"It would be good to be able to get different coverages by different news channel for the same story. Once a story is viewed, links to other channels could be followed." [User7:Diary:02/04]

"...some external links to the stories would be great maybe to newspapers, business websites, etc." [User13:Pre-Trial Q]

Wishing to see links to more resources and views from different news sources on the same events, a technique to crawl and collect other news websites' coverage and make a list of link sets could be utilised similar to Google News, so that users have ready links to other sources. The above comments are also in accordance with our participants' expectations on Físchlár-News as an up-to-date news delivery service rather than a searchable archive.

(iii) *News stories that are about different events but similar in nature* - for example, different events involving the same person/company/country, different stories about the same town, or house fires in different regions of the country. This group of stories could provide interesting new light on a topic or help discover a new pattern of incidents, which is not usually the case in conventional news media. "Related Stories" in Físchlár-News, based on simply analysing Closed-Captions of individual stories, tends to find and present these potentially related news stories that a user might find interesting as well as the first category above. This kind of opportunistic or serendipitous value of Related Stories in Físchlár-News was witnessed in our participants' comments, too:

"... I often found myself wandering through several layers of related stories which were more interesting than the news stories of that particular day." [User3:Follow-up Q]

"...I played the story on the boat accident off the Canary Islands from April 17th. This led me to a related story from 22nd June 2003 on a row reporting the Italian government about their immigration policy..." [User18:Diary:22/04]

Overall, without separately presenting the 3 different categories but presenting in a single list with 10 stories under "Related Stories" in Físchlár-News, our participants did reveal a tension between sometimes wanting the feature, and sometimes not, but overall they had a positive view on it and actually made use of it, as some of their comments reveal:

"A related news facility is a really useful feature to have. It has to remain unobtrusive so that a user can ignore it if they wish to. I have no real idea at the moment on how to improve, but I made a lot of use of the facility." [User3: Follow-up Q]

"...my main interest [of using Físchlár-News] lay in tracking the progression of news stories in which I had an interest over the past few months." [User19:Post-Trial Q]

"...I would first use the calendar navigation or the text search to locate it, and if necessary would then use the 'Related stories' feature to track its development over the past few weeks/months." [User19:Post-Trial Q]

We administered the follow-up questionnaire to our participants to specifically ask whether they were aware of or used any related stories feature with their own favourite news websites. Only 1 participant was aware of such a feature and actually used it to follow the related stories within the news site, another 2 participants were aware of such a feature but said they rarely used it. Given the overall experiences of these participants with online news websites, navigating related stories is probably still a novel feature with which the user behaviour will start co-evolving once such a feature becomes more commonly available and people recognise its usefulness. We know now that clarifying different kinds of "Related Stories" in some way would reduce some of our participants' dissatisfaction and perhaps lead to more use. For example, we can separate the 3 different kinds of related stories by labelling them differently, or by listing them in different areas on the screen with different terminology. In this way, a user expecting only the latest news stories will not try to look at the thread of news stories linking to past stories but the list of other similar stories for serendipitous browsing, while another user investigating the background or development of a particular event could directly look into "previous stories on this event" list, avoiding the stories that are about completely different events involving a few related names or companies.

SUMMARY

Our participants made use of the Related Stories feature to varying degrees and had varying opinions on it. We have identified 3 different types of possible Related Stories in news, and tried to see their characteristics through our participants' comments. By explicitly differentiating the different types of Related Stories, we will be able to enhance further the usability of this feature by detecting and responding to users' intentions more immediately. We see this feature as potentially useful and worth promoting the use by providing more visibility, thus shaping and observing new user behaviour in news story browsing.

5.2.5 USER EXPECTATIONS AND SYSTEM IMAGE

In theory it takes 30 minutes to watch 30 minutes of broadcast TV news though there are techniques to speed this including fast-forwarding, and skipping of portions. In the case of our users, many advantages over simply having to watch the news for 30 minutes were noted and highlighted. Most notably, the fact that the broadcast news is presented as story units, segmented as a result of the story segmentation process - was one of the major points of praise as the following comments from participants show:

"The story list provides a concise overview of the news items." [User8:Post-Trial Q]

"I find it useful that you can read a little of each story and can play those ones which might be of interest. Sometimes there might be aspects of the news that I'm not interested in. By using Físchlár, I can skip over those sections and just get details of the stories which are of interest to me." [User10:Diary:08/04]

"[Story list] Allows quick perusal of all news items on a given day; useful as an overview. ...[High] quality of the news breakdowns and summaries..." [User19:Post-Trial Q]

When participants considered Físchlár-News as a utility for broadcast TV news, there were many things they had to say about the added value from the story segmentation process. On the other hand, our users also considered Físchlár-News as an "alternative to commercial news websites," those websites which they use on a daily basis to get the latest updates on news of the moment. They wanted to see in Físchlár-News those features available on commercial news sites, notably presenting news stories in broad categories:

"[in online papers] I like the way you can choose from different sections ie. National, Europe or International news." [User4:Diary:20/04]

"Link to different specialised areas such as sport, entertainment, etc." [User10:Post-Trial Q]

Automatically classifying segmented news stories into broad categories such as National, International, Health, Entertainment, Sports, etc. would be possible by analysing the Closed-Caption text of each story, but this had not been conceived for implementation during system development stage. While we think it is possible to do this, some errors would be introduced (e.g., a story on "Business" found in "National" category).

Físchlár-News was designed to be operational without human intervention 24/7 and most of the time it performs its automatic processes without much problem. However, as mentioned in Sections 2 and 3, some of the processes are not at the stage of development where the results could be claimed as 100% accurate, and these individual processes are active research areas within our team and among many other research groups around the world. Because of this, inevitably some errors and imprecision in system performance was from time to time introduced during the trial period, and these errors were often spotted by our participants and pointed out and noted in their feedback. Our participants did not tolerate system errors, expecting Físchlár-News to work as error-free as the carefully handcrafted and manually maintained commercial news websites such as BBC News and the Irish Independent sites.

Along with the fact that Físchlár-News is updated only daily and its coverage is limited to RTÉ1 news (as discussed earlier), minor technical problems that were introduced due to the automatically processed and maintained nature of the system were spotted and noted:

- Closed-Caption text occasionally appears illegible, with jumbled-up characters due to the fact that Físchlár-News captures analogue TV signals and sometimes the weak signal caused by environmental and weather conditions resulted in this jumbled-up text
- Closed-Caption mis-alignment with the video where the text is a few seconds behind the video due to the latency with stenographers who type in the Close-Caption in real-time
- Slow page loading: displaying the story detail page requires the computation of a Related Stories listing, which queries the Closed-Caption of the story against all stories in the archive at loading time, and thus as the archive becomes larger, this takes longer to compute

STORY RECOMMENDATION AS NEWS DELIVERY

We were also interested in finding out how our participants found the news story recommendation feature in Físchlár-News. From the briefing and introduction before the trial period started, most of our participants saw the potential usefulness of the recommendations and voiced their interest in making use of it. However, the actual usage frequency of this feature was low:

- Total story ratings (by clicking on the thumbs-up and -down icon): **327** times (average 20.4 times / person)
- Recommendation viewed **80** times (average 5 times / person)

Of all story ratings, 73% (238 ratings out of 327) were done by 6 participants, each rating about 30-60 times (average 39.7 ratings / person). The rest of the participants did not rate stories often and did not view their recommendation often. The reasons for the low usage of this feature, in spite of the positive initial reaction by the participants, can be found by tracing comments in the diary entries and post-trial questionnaires. According to User11's diary (who rated stories 34 times, and viewed his recommendation page 11 times):

[After 5 days into the trial] "I continued to rate stories to improve my recommendations... I don't think I would rely on the recommendations alone for some time yet." **[User11:Diary:06/04]**

[8 days into trial] "Recommendations appeared for me today for the first time. I would be interested in viewing probably half of the stories recommended..." **[User11:Diary:09/04]**

[10 days into trial] "As there was rugby on over the weekend, I checked my recommendations because I had positively rated rugby stories last week. A rugby story appeared, but so did a golf one, even though I negatively rated golf stories last week..." **[User11:Diary:15/04]**

[26 days into trial] "I had a quick look at the stories recommended for me. The first few of these stories don't interest me at all really, so I doubt I'll look at my recommendations or rate stories too often from now on." **[User11:Diary:27/04]**

The explicit preference indication (by clicking on the thumbs icon) is something extra that a user needs to do and thus is often felt as a burden or an extraneous task while interacting with the system for some purpose. If this effort does not seem to pay off, the user becomes disappointed and stops using it.

According to User18's comments (who rated stories 40 times and viewed his recommendation page 5 times):

[After 2 days into the trial] "Before I logged in today, I first went to my recommendations section to see if it contained much stuff. In particular, I was hoping it would point to story from last night's news, but everything that came up was before March 10th..." **[User18:Diary:07/04]**

[9 days into trial] "Today I went directly to my recommendations list but, surprisingly, the most recent stories in it was still from April 8th and nothing new has been added... It is disappointing when personal effort goes into customising a system but this effort does not produce visible results." **[User18:Diary:15/04]**

"In theory, it would seem useful, but I didn't find this documents [stories] recommended to me very interesting." **[User18:Post-Trial Q; used only 5 times, though spread out]**

Tracing these participants' diaries and questionnaire comments revealed that their expectations on useful recommendations was not satisfied especially when they had to provide extra input by explicitly indicating their preference for stories, and by the end they were disappointed and stopped using the feature. The relatively small number of users of Físchlár-News during the trial period could be the culprit for the poor recommendation results, as the recommendation engine is based partly on collaborative filtering which requires a number of users continue to rate stories.

For other participants who did not rate many stories during the trial, we found that they forgot to rate stories while browsing/playing news:

"I also occasionally used the recommended feature but not very often as I kept forgetting to rate stories as it was then less useful for me. [It would be nice] to have a reminder like a pop-up about the recommended feature after accessing a story" **[User5:Post-Trial Q; rated 8 times]**

"I think I'm quite unlikely to remember to rate clips, so I don't think my recommendations will be that good & as a result I probably won't use it that much." **[User16:Diary:20/04; rated 3 times]**

The nature of explicit feedback (clicking on the 5-scale thumbs-up and thumbs-down icons) means that users require to put effort to consciously indicate their preference whenever they encounter a news story. Even though the thumbs icons are highly visible and placed whenever a story is presented, any extra action not directly related to the user's immediate interest or task can be easily neglected and forgotten and this is one of the negative consequences of an explicit feedback mechanism. The suggestion by User5 above of having a "pop-up message to remind the rating after accessing a story" would not be a desirable strategy to promote more rating, but we see her point and the need to resolve this problem in some way.

As mentioned in Section 4, the system's preference rating scale is simply the degree of preference in a 5-point scale, excluding any further dimensions such as the preference on the topic or this particular story, or the informativeness of the story, or whether this story was useful to him/her. Our participants simply rated when they had positive experience with the story, and in their diaries this was mostly expressed as "I liked the story so I gave thumbs up." At the beginning of the trial it was an issue for several participants as to exactly what kind of stories they should rate, to which our reply was "rate positively the kinds of stories that you would like to be recommended in the future." Some diary comments reveal why they were rating the stories:

"I'm interested in rugby, but not all that interested in other sports. So, I decided to positively rate some rugby stories and negatively rate some other sport stories." [User11:D:08/04]

"... Then checked yesterday's news. Roy Keane was in the news again! Gave it thumbs up. Marked all other stories thumbs down, I wasn't interested in watching them." [User22:D:27:04]

The story recommendation of Físchlár-News did not satisfy some of our participants partly because they found it did not accurately perform in a way they expected thus confirming the findings of another study that users are sensitive to the inaccuracy of recommendation [Cosley *et al.* 2003] and partly because the interaction mechanism to make it work well did not fit in smoothly with their main interest in using Físchlár-News and this did not generate enough of a database of story ratings. However we did have some of our participants who found the recommendations accurate and useful:

"I found it worked in both recommending interesting stories and filtering out stories I specifically didn't want. I made an effort to get the feature to not recommend certain stories [rugby stories] and it seemed to work well." [User3:Post-Trial Q]

[16 days into trial] "Checked my recommended news. Found one story that really interested me and I gave it the thumbs up symbol." [User22:Diary:23/04]

[29 days into trial] "I looked for the first time at recommended news stories I wanted to see how well the system had "represented my interests"! I was quite impressed by the selection of stories presented." [User4:Diary:26/04]

[13 days into trial] "I found most of the stories there were of interest to me. This is a useful feature." [User5:Diary:08/04]

There is much room for improvement in the quality of the recommendation in Físchlár-News, to make it more relevant and useful. Firstly, we need to devise ways to get the users more frequently rate the stories. These could be more prominent and visible thumbs-up/-down icons or re-organising the interface in such a way that the whole interaction with the system is centred around the rating. More proactive schemes such as a reward for certain number of ratings can also be adopted, and this would require visible advertisement of this on the interface with a counting display of how many times the user has rated so far. Going further towards an explicit user indication method, we can consider having a user-configurable keyword list to be used to help recommendation similar to those email news reminder services that ask the user to tick the subjects s/he is interested in, to be used for filtering incoming news for that user. Exactly how much preference data is required for more accurate recommendation and whether users will be willing to supply is an open question and would need further study with a larger number of users, but there are examples of recommender systems in comparable domains, including TV programs [Smyth and Cotter 2000], movies [Cosley *et al.* 2003] and text-based news articles [Billsus and Pazzani 1999], in which people did explicitly rate the items to the point of effective recommendation. Secondly, to get more preference data, we can consider employing an implicit feedback mechanism. As found in Section 5.2.1, our participants used the playback feature very frequently and this was central to their usage of looking at news stories. Thus it is conceivable to regard a playback of a story as "I like this story." Considering the relative amount of playback initiated by our participants, this could prove a good way to add more user preference. The story detail view has also been accessed frequently, and thus could be considered for possibly useful implicit evidence to help add more preference data. Although implicit feedback is known to be less accurate in obtaining user preferences, O'Sullivan *et al.* [2003] showed that in a web-based TV recording/playback system the implicit ratings from playback and recording request were just as accurate as explicit rating.

SUMMARY

Most of our participants had one or more favourite news websites they looked up to get news updates, and thus regarded Físchlár-News, when introduced, as a similar type of web site, wanting to have news website-like features particularly a categorisation of news stories into international, business, entertainment or sports. Our system lacked this feature, but the organised look of story listing of the day (as a result of story segmentation) was rated highly. We think automatic categorisation of stories with the accuracy of automatic story segmentation would enhance the perceived usefulness of the system.

The automatic story recommendation feature, although initially perceived as useful, turned out to be less and less used due to imprecise output and the lack of feedback data. Making the story recommendation more useful will require re-evaluation of various interaction design and strategic aspects of the system.

5.2.6 SUMMARY

As mentioned in previous sections, Físchlár-News is a system that turns transient broadcast news into a tangible, easy-to-access the information entity in an automated manner, and to the best of our knowledge there is no similar system deployed and used by people today. Deploying this system allowed us to conduct a detailed user evaluation, which would be one of the first user studies on a system of this kind. The analysis of usage by our participants revealed usability and utility issues about the system's various elements: coverage of the news, usefulness of playback, archival value, limited up-to-datedness compared to other news websites, and yielded a long wishlist for possible feature improvements. Our participants were probably typical of any other "users" of technology products, mostly interested in continuing and satisfying their news update habits that they have been doing in their daily lives. The fact that our participants used the system for a month, complained about it, praised it, and voiced their wishes for improved features is a sign of approval that Físchlár-News as it is can be integrated into people's daily lives.

Now that web access and usage is ubiquitous, having a web interface to Físchlár-News meant easy access by anybody from his/her own desk or home. However, our observations highlighted to us that due to this very reason, a web interface needs to accommodate the established web browsing habits of users. For example, opening another web browser window while waiting for a page to load was how our User16 and User21 cope with slow page loading in their day-to-day web browsing, thus the prescribed interaction stages enforced in Físchlár-News, following the standalone application-like interface in which clicking on a link usually makes changes to one part of the screen, frustrated them. Trying to get a new service or product penetrating into people's everyday lives requires the understanding of the details of their current everyday lives where the service or product is to be integrated.

Sometimes what we thought as negligible and did not even regard as a system error blew up as "major" problem such as the playback window continuing to the next story rather than stopping at the end of a story; what we did not consider as potentially useful features for the system were voiced strongly as highly desirable features such as bookmarking or sending a story to a friend; the system we regarded as radically different and advanced from any other existing system was constantly compared to the manually-maintained news websites our users used. Our experience with this study showed the well-emphasised gap between what developers think and what users think, however we tried to focus on the end-users as we were developing the system.

In terms of the scale of the study, we think the number of studied users, although small compared to many other lab-based experiments, was about as far as we could manage in a 7-week long field study. With more participants, we might have collected a greater amount of usage information and possibly a greater variety of comments, but most likely at the cost of missing out the detailed considerations of each participant's situation and context which is what we wanted to focus on in this study.

SUMMARY OF PARTICIPANTS' WISHES

Our study has resulted in a large number of user comments from questionnaires and incident diaries. Although the main purpose of collecting participants' comments was to obtain information on general and specific situations, contexts, and the value of Físchlár-News when it was put into real use by real users and in this way to try to understand the system and its use, there were a large number of "wishes" or ideas for a better system voiced by the participants as a by-product of this study. While many of them have been mentioned in the context of the larger themes, we have extracted a summary of this wishlist from our participants. This wishlist is a usability enhancement inventory which our team will use to refine and improve the current Físchlár-News's features, and also can be considered as a useful starting point for future researchers/developers in designing a similar system.

Better coverage

- Include more/other TV news (RTÉ1's 1pm & 6pm news, news from BBC1, RTÉ2 and TV3)
- Include other current affairs TV programmes (Questions & Answers and Prime Time, i.e. current affairs chat shows)
- Include information from news/business websites
- Include sources from other countries' broadcasters
- Include other language options

More up-to-date

- Record more frequently

Search feature

- Allow a date range to be specified in searching
- Allow search results to be presented by date
- Query syntax simplification (more like Google)
- Allow search by journalist/reporter

Related stories feature

- Order by date
- Group the same events together (also by month)
- List links to stories from other sources, e.g. websites
- Visualise the threading of event development

Playback feature

- Play just 1 story and then stop
- Use buttons to jump to previous/next stories
- Play the news headline of the day
- Play the news headline's audio on selecting a day from the calendar
- Display elapsed time (from the beginning of the story)
- Display Closed-Caption text scrolling at the bottom of the video playback
- Display the text banner from today's breaking news at the bottom of the playback

Service-oriented features

- Categorise stories into broad themes (business, international, entertainment, sports, etc.)
- Allow multiple windowing
- Send a news story to a friend
- Bookmark a news story

6. CONCLUSION

It is difficult to conduct a long-term, out-of-the-lab, user evaluation of an experimental system in which the constituent technology is not all yet sufficiently mature. In our study, story segmentation had to be checked by a human after automatic segmentation has been done, to be able to deliver accurately segmented news stories to our users, in a way “cheating” our users into believing that the system is completely automatic in its entirety. Even with this we experienced instances where because some underlying components did not perform perfectly, the study results magnified any imperfections with the system.

In the post-trial questionnaire, we asked our participants whether they would like to continue to use the system. Thirteen participants replied that they wanted to continue, and three replied they did not want to. With our qualitative approach to understand participants' usage within their contexts, a statement such as “13 out of 16 said positive” is not in itself of much significance. However, after observing and interacting with each of our 16 participants for a 1 month period, we did get many useful insights and “13 out of 16” did make sense to us in some way.

In this time of technological fervour and constant development it is easy to forget the basic fact that these systems are for the betterment of our lives at work and play. On the other hand, it is probably the technology sector that can best address the other side of the whole picture. From this point of view, in our user evaluation of Físchlár-News we tried to focus on the user side of the developed system by thinking not about what the system can do but about what our users can do, the mindset being to move forward to the “new computing” [Shneiderman 2002].

With the recognition that currently used usability engineering methods for new interactive systems tend to take only a snapshot of the very first impression and usage ignoring any understanding of the context, the development of use over time and longer-term implications, we believe our qualitative assessment by visiting each user at his/her desk and getting them to use our system for a 1-month period was most appropriate at this technological stage for this system. Taking the

context of usage as the centre of understanding and designing the products accordingly is the core concept of the “new usability” rightly addressed by Thomas and Macredie [2002].

We believe the real breakthrough and advancement in content-based video technology will come when it goes hand and hand with usability concerns of products. We hope that researchers working in this field will pay more attention to the outcomes of their work and possible usage of their systems by carrying out this type of deployment effort and usage monitoring. The guesswork about fully-automated future technology products and their usability can be reduced, and future directions can be more grounded by human needs. Hopefully we will see more examples of deployed efforts with experimental systems that will result in people being more exposed to these new breeds of automatic systems, as we learned from our users’ experiences, and from *our* experience in having real users of our end-to-end automatic content-based broadcast news delivery system.

ACKNOWLEDGEMENTS

The support of the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged. Part of this work was funded by Science Foundation Ireland under grant 03/IN.3/I361. The authors also would like to thank all participants in this study, and members of CDVP who rated news stories to make recommendation more accurate, and other users of the Físchlár-News system for contributing to the usage on a larger scale. We are also grateful for the very constructive comments and feedback offered by the reviewers of this paper whose insights have helped to improve it.

REFERENCES

- ALLAN, J., PAPKA, R., AND LAVRENKO, V. 1998. On-line new event detection and tracking. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR’98)*, Melbourne, Australia, August 1998, 37-45.
- AMIR, A., IYENGAR, G., LIN, C.-Y., NAPHADE, M., NATSEV, A., NETI, C., NOCK, H.J., SMITH, J.R., AND TSENG, B. 2004. Multimodal video search techniques: late fusion of speech-based retrieval and visual content-based retrieval. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2004)*, Montreal, Quebec, Canada, May 2004, Vol 3, 1048-51.
- ARLANDIS, J., OVER, P., AND KRAAIJ, W. 2005. Boundary error analysis and categorization in the TRECVID news story segmentation task. In *Proceedings of the 4th International Conference on Image and Video Retrieval (CIVR 2005)*, Singapore, July 2005, 103-112.
- BILLSUS, D., AND PAZZANI, M. 1999. A Hybrid User Model for News Story Classification. In *Proceedings of the 7th International Conference on User Modeling*. Banff, Canada, June 1999, 99-108.
- BORECKZY, J., GIRGENSOHN, A., GOLOVCHINSKY, G., AND UCHIHASHI, S. 2000. An interactive comic book presentation for exploring video. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2000)*, The Hague, The Netherlands, April 2000, 185-192.
- BROWNE, P., SMEATON A.F., MURPHY, N., O’CONNOR, N., MARLOW, S., AND BERRUT, C. 2000. Evaluating and Combining Digital Video Shot Boundary Detection Algorithms. In *Proceedings of the Irish Machine Vision and Image Processing Conference (IMVIP 2000)*, Belfast, Northern Ireland, September 2000.
- BROWNE, P., CZIRJEK, C., GAUGHAN, G., GURRIN, C., JONES, G., LEE, H. MARLOW, S., MCDONALD, K., MURPHY, N., O’CONNOR, N., O’HARE, N., SMEATON, A.F., AND YE, J. 2003. Dublin City University Video Track Experiments for TREC 2003. In *Proceedings of the Text REtrieval Conference TRECVID Workshop (TRECVID 2003)*, Gaithersburg, Maryland, November 2003.
- CHRISTEL, M., HAUPTMANN, A., WACTLAR, H., AND NG, T. 2002. Collages as dynamic summaries for news video. In *Proceedings of the 10th ACM International Conference on Multimedia (MM’02)*, Juan-les-Pins, France, December 2002, 561-569.
- CHRISTEL, M., OLLIGSCHLAEGER, A., AND HUANG, C. 2000. Interactive maps for a digital video library. *IEEE Multimedia* 7(1), 60-67.
- CHUA, T-S., CHANG, S-F., CHAISORN, L., AND HSU, W. 2004. Story boundary detection in large broadcast news video archives: techniques, experience and trends. In *Proceedings of the 12th Annual ACM International Conference on Multimedia (MM’04)*, New York, NY, October 2004, 656-659.
- COOKE, E., FERGUSON, P., GAUGHAN, G., GURRIN, C., JONES, G., LE BORGNE, H., LEE, H., MARLOW, S., MC DONALD, K., MCHUGH, M., MURPHY, N., O’CONNOR, N., O’HARE, N., ROTHWELL, S., SMEATON, A.F., AND WILKINS, P. 2004. TRECVID 2004 Experiments in Dublin City University. In *Proceedings of the Text REtrieval Conference TRECVID Workshop (TRECVID 2004)*, Gaithersburg, MD.
- COSLEY, D., LAM, S., ALBERT, I., KONSTAN, J., AND RIEDL, J. 2003. Is seeing believing? How recommender interfaces affect users’ opinions. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2003)*, Ft. Lauderdale, FL, April 2003, 585-592.
- DUYGULU, P., PAN, J., AND FORSYTH, D. 2004. Towards auto-documentary: tracking the evolution of news stories. In *Proceedings of the 12th Annual ACM International Conference on Multimedia (MM’04)*, New York, NY, October 2004, 820-827.
- FAUDEMAY, P., DURAND, G., SEYRAT, C. AND TONDRE, N. (1998). Indexing and retrieval of multimedia objects at different levels of granularity. In *Proceedings of SPIE Vol. 3527, Multimedia Storage and Archiving Systems III*, Boston, MA, November 1998, 112-121.

- FITZPATRICK, J., REFFELL, J., AND AYDELOTT, M. 2003. BreakingStory: visualizing change in online news. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2003) Extended abstracts*, Ft. Lauderdale, FL, April 2003, 900-901.
- FREY, D., GUPTA, R., KHANDELWAL, V., LAVRENKO, V., LEUSKI, A., AND ALLAN, J. 2001. Monitoring the news: a TDT demonstration system. In *Proceedings of the 1st International Human Language Technologies Conference (HLT'01)*, San Diego, CA, March 2001, 351-355.
- GAUCH, J., GAUCH, S., BOUIX, S., AND ZHU, X. 1999. Real time video scene detection and classification. *Information Processing and Management*, 35, 381-400.
- GAUGHAN, G., SMEATON, A.F., GURRIN, C., LEE, H., AND MC DONALD, K. 2003. Design, Implementation and Testing of an Interactive Video Retrieval System. In *Proceedings of the 5th ACM SIGMM International Workshop on Multimedia Information Retrieval (MIR 2003)*, Berkeley, CA, November 2003, 23-30.
- GOLDBERG, K., ROEDER, T., GUPTA, D., AND PERKINS, C. 2001. Eigentaste: a constant time collaborative filtering algorithm. *Information Retrieval Journal*, 4(2), 133-151.
- GURRIN, C., BROWNE, P., SMEATON, A.F., LEE, H., MC DONALD, K., AND MACNEELA, P. 2004. Físchlár-Nursing, Using Digital Video Libraries to Teach Processes to Nursing Students. In *Proceedings of the IASTED International Conference on Web-Based Education (WBE 2004)*, Innsbruck, Austria, February 2004, 111-116.
- HAGGERTY, A., WHITE, R., AND JOSE, J. 2003. NewsFlash: adaptive TV news delivery on the web. In *Proceedings of the 1st International Workshop on Adaptive Multimedia Retrieval (AMR2003)*, Hamburg, Germany, September 2003, 72-86.
- HAUPTMANN, A., AND WITBROCK, M. 1998. Story segmentation and detection of commercials in broadcast news video. In *Proceedings of the Advances in Digital Libraries Conference (ADL'98)*, Santa Barbara, CA, April 1998.
- HERLOCKER, J., KONSTAN, J., TERVENN, L., AND RIEDL, J. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1), 5-53.
- HURST, W., GOTZ, G., AND JARVERS, P. 2004. Advanced user interfaces for dynamic video browsing. In *Proceedings of the 12th Annual ACM International Conference on Multimedia (MM2004)*, New York, NY, October 2004, 742-743.
- IDE, I., MO, H., KATAYAMA, N., AND SATOH, S. 2003. Topic-based structuring of a very large-scale news video corpus. In *Proceedings of the AAAI-03 Spring Symposium on Intelligent Multimedia Knowledge Management*, Palo Alto, CA, March 2003.
- IDE, I., MO, H., KATAYAMA, N., AND SATOH, S. 2004. Topic threading for structuring a large-scale news video archive. In *Proceedings of the 3rd International Conference on Image and Video Retrieval (CIVR2004)*, Dublin, Ireland, July 2004, 123-131.
- KELLY, D., AND TEEVAN, J. 2003. Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37(2), 18-28.
- KRAAIJ, W., SMEATON, A.F., AND OVER, P. 2004. TRECVID 2004 – An overview. In *Proceedings of the Text REtrieval Conference TRECVID Workshop (TRECVID 2004)*, Gaithersburg, MD, November 2004.
- LEE, H., AND SMEATON, A.F. 2002. Designing the user-interface for the Físchlár Digital Video Library. *Journal of Digital Information, Special Issue on Interactivity in Digital Libraries*, 2(4).
- LEE, H., SMEATON, A.F., AND SMYTH, B. 2005. User evaluation outside the lab: the trial of Físchlár-News. In *Proceedings of the 5th International Conference on Conceptions of Library and Information Science, Workshop on Evaluating User Studies in Information Access*, Glasgow, U.K., June 2005, 1-12.
- LEE, W.O. 2000. Introducing Internet Terminals to the home: interaction between social, physical, and technological spaces. *Proceedings of the 14th Annual Conference on Human-Computer Interaction (HCI 2001)*, 119-132.
- LYMAN, P., AND VARIAN, H. 2003. How Much Information? Retrieved from <http://www.sims.berkeley.edu/how-much-info-2003/>.
- GUILLEMOT, M., WELLNER, P., GATICA-PEREZ, D., AND ODOBEZ, J.M. 2003. A hierarchical keyframe user interface for browsing video over the Internet. In *Proceedings of the 9th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2003)*, Zurich, Switzerland,, September 2003.
- MARCHIONINI, G. 2004. From information retrieval to information interaction. In *Proceedings of the 26th European Conference on Information Retrieval (ECIR'04)*, Sunderland, U.K., April 2004.
- MERLINO, A., MOREY, D., AND MAYBURY, M. 1997. Broadcast news navigation using story segmentation. In *Proceedings of the 5th ACM International Conference on Multimedia (MM'97)*, Seattle, WA, November 1997, 381-391.
- MERIALDO, B., LEE, K-T., LUPARELLO, D., AND ROUDAIRE, J. 1999. Automatic construction of personalized TV news programs. In *Proceedings of the 7th ACM International Conference on Multimedia (MM'99)*, Orlando, FL, October 1999, 323-331.
- MCKEOWN, K., BARZILAY, R., EVANS, D., HATZIVASSILOGLU, V., KLAVANS, J., NENKOVA A., SABLE, C., SCHIFFMAN, B., AND SIGELMAN, S. 2002. Tracking and summarizing news on a daily basis with Columbia's Newsblaster. In *Proceedings of the International Human Language Technologies Conference (HLT'02)*, San Diego, CA, March 2002.
- MCNEE, S.M., LAM, S.K., GUETZLAFF, C., KONSTAN, J.A., AND RIEDL, J. 2003a. Confidence displays and training in recommender systems. In *Proceedings of the 9th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2003)*, Zurich, Switzerland,, September 2003, 176-183.
- MCNEE, S.M., LAM, S.K., KONSTAN, J.A., AND RIEDL, J. 2003b. Interfaces for eliciting new user preferences in recommender systems. In *Proceedings of the 9th International Conference on User Modeling (UM'03)*, Johnstown, PA, June 2003, 178-188.
- MORRIS, J., AND HIRST, G. 1991. Lexical cohesion computed by thesaural relations as an indicator of the structure of text. *Computational Linguistics*, 17(1), 21-48.

- NG, T., CHRISTEL, M., HAUPTMANN, A., AND WACTLAR, H. 2003. Collages as dynamic summaries of mined video content for intelligent multimedia knowledge management. In *Proceedings of the AAAI-03 Spring Symposium on Intelligent Multimedia Knowledge Management*, Palo Alto, CA, March 2003.
- O'BRIEN, J., RODDEN, T., ROUNCFIELD, M., AND HUGHES, J. 1999. At home with the technology: an ethnographic study of a set-top-box trial. *ACM Transactions on Computer-Human Interaction*, 6(3), 282-308.
- O'HARE, N., SMEATON, A.F., CZIRJEK, C., O'CONNOR, N., AND MURPHY, N. 2004. A generic news story segmentation system and its evaluation. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2004)*, Montreal, Quebec, Canada, May 2004, 1028-1031.
- O'SULLIVAN, D., WILSON, D., AND SMYTH, D. 2003. Explicit vs. implicit profiling – a case-study in electronic programme guides. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI 03)*, Acapulco, Mexico, 9-15 August 2003.
- O'SULLIVAN, D., SMYTH, B., WILSON, D., MCDONALD, K., AND SMEATON, A.F. 2004. Improving the quality of the personalized electronic program guide. *User Modeling and User-Adapted Interaction*, 14(1), 5-36.
- O'TOOLE, C., SMEATON A.F., MURPHY, N., AND MARLOW, S. 1999. Evaluation of Automatic Shot Boundary Detection on a Large Video Test Suite. In *Proceedings of The Challenge of Image Retrieval: 2nd UK Conference on Image Retrieval (CIR'99)*, Newcastle, U.K., February 1999.
- PERRY, M., O'HARA, K., SELLEN, A., BROWN, B., AND HARPER, R. 2001. Dealing with mobility: understanding access anytime, anywhere. *ACM Transactions on Computer-Human Interaction*, 8(4), 323-347.
- PETERSEN, M., MADSEN, K., AND KJAER, A. 2002. The usability of everyday technology - emerging and fading opportunities. *ACM Transactions on Computer-Human Interaction*, 9(2), 74-105.
- PICKERING, M., WONG, L., AND RUGER, S. 2003. ANSES: summarisation of news video. In *Proceedings of the 2nd International Conference on Image and Video Retrieval (CIVR2003)*, Urbana, IL, July 2003.
- RADEV, D., GLAIR-GOLDENSOHN, S., ZHANG, Z., AND RAGHAVAN, R. 2001. NewsInEssence: a system for domain-independent real-time news clustering and multi-document summarization. In *Proceedings of the 1st International Human Language Technologies Conference (HLT'01)*, San Diego, California.(demo session paper), March 2001.
- RENNISON, E. 1994. Galaxy of news: an approach to visualizing and understanding expansive news landscapes. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology (UIST'94)*, Marina del Rey, CA, November 1994, 3-12.
- ROUSH, W. 2004. Search Beyond Google. *MIT Technology Review*, 34-45.
- SHNEIDERMAN, B. 2003. *Leonardo's Laptop : Human Needs and the New Computing Technologies*. MIT Press, Cambridge, MA.
- SMEATON A.F. 2004. Indexing, Browsing and Searching of Digital Video. *Annual Review of Information Science and Technology (ARIST)*, 38, Chapter 8, 371-407.
- SMEATON A.F., KRAAIJ, W., AND OVER, P. 2004a. The TREC Video Retrieval Evaluation (TRECVID): A Case Study and Status Report. In *Proceedings of RIAO 2004 - Coupling Approaches, Coupling Media and Coupling Languages for Information Retrieval*, Avignon, France, April 2004.
- SMEATON, A.F., GURRIN, C., LEE, H., MC DONALD, K., MURPHY, N., O'CONNOR, N., O'SULLIVAN, D., SMYTH, B., AND WILSON, D. 2004b. The Fischlär-News-Stories System: Personalised Access to an Archive of TV News. In *Proceedings of RIAO 2004 - Coupling Approaches, Coupling Media and Coupling Languages for Information Retrieval*, Avignon, France, April 2004.
- SMYTH, B., AND COTTER, P. 2000. A personalized television listings service. *Communications of the ACM*, 43(8), 107-111.
- SWAN, R., AND ALLAN, J. 2000. TimeMine: visualizing automatically constructed timelines. In *Proceedings of the 23rd Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR2000)*, Athens, Greece, July 2000, 393.
- SWEARINGEN, K., AND SINHA, R. 2001. Beyond algorithms: an HCI perspective on recommender systems. In *Proceedings of the 24th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR2001), Workshop on Recommender Systems*, New Orleans, LA, September 2001.
- SWEARINGEN, K., AND SINHA, R. 2002. Interaction design for recommender systems. In *Proceedings of the Conference on Designing Interactive Systems: Process, Practices, Methods, and Techniques (DIS2002)*, London, U.K., June 2002.
- THOMAS, P., AND MACREDIE, R. 2002. Introduction to the new usability. *ACM Transactions on Computer-Human Interaction*, 9(2), Special Issue on the New Usability.
- ZHANG, H., LOW, C., SMOLIAR, S., AND WU, J. 1995. Video parsing, retrieval and browsing: an integrated and content-based solution. In *Proceedings of 3rd ACM International Conference on Multimedia (MM '95)*, San Francisco, CA, November 1995, 503-512.