

# The Físchlár-News-Stories System: Personalised Access to an Archive of TV News

Alan F. Smeaton(+), Cathal Gurrin, Hyowon Lee, Kieran McDonald, Noel Murphy,  
Noel E. O'Connor, Derry O'Sullivan(\*), Barry Smyth(\*) and David Wilson

Centre for Digital Video Processing  
Dublin City University  
Glasnevin, Dublin 9,  
Ireland

(\*) Smart Media Institute  
Department of Computer Science  
University College Dublin  
Belfield, Dublin 4, Ireland.

(+) Contact author: Alan Smeaton@DCU.ie

## Abstract

The “Físchlár” systems are a family of tools for capturing, analysis, indexing, browsing, searching and summarisation of digital video information. Físchlár-News-Stories, described in this paper, is one of those systems, and provides access to a growing archive of broadcast TV news. Físchlár-News-Stories has several notable features including the fact that it automatically records TV news and segments a broadcast news program into stories, eliminating advertisements and credits at the start/end of the broadcast. Físchlár-News-Stories supports access to individual stories via calendar lookup, text search through closed captions, automatically-generated links between related stories, and personalised access using a personalisation and recommender system based on collaborative filtering. Access to individual news stories is supported either by browsing keyframes with synchronised closed captions, or by playback of the recorded video. One strength of the Físchlár-News-Stories system is that it is actually used, in practice, daily, to access news. Several aspects of the Físchlár systems have been published before, but in this paper we give a summary of the Físchlár-News-Stories system in operation by following a scenario in which it is used and also outlining how the underlying system realises the functions it offers.

## 1. Introduction to the Físchlár Systems

The Centre for Digital Video Processing at Dublin City University has been building, testing, deploying and evaluating systems which provide content access to archives of digital video information for several years. The systems we build are known under the generic title of “Físchlár”, and each addresses a need for content-based access to reasonably-sized archives of digital video. These content-based operations include searching (Smeaton, 2004), browsing (Lee and Smeaton, 2002) and summarisation (Sadlier *et al.*, 2003) and we have tried where possible to deploy our work in operational settings. Thus the Físchlár-TV system allowed recording, browsing and playback of broadcast TV programmes for staff and students on the University campus, including student residences; the Físchlár-Nursing system allows students and Faculty in our School of Nursing to browse and playback educational videos related to their syllabus for Nursing courses (Gurrin *et al.*, 2004); the Físchlár-TREC systems we developed in 2001, 2002 and 2003 provided tailored access to video collections as part of the annual the TRECVID exercise (Smeaton *et al.*, 2003) and the Físchlár-News-Stories system, the focus of this paper, provides access to a growing archive of main evening TV news.

The Físchlár-News-Stories System brings together a range of technologies which we have developed in-house and combines them into a single system. These technologies include at least the following: automatic video capture and digitisation, automatic segmentation of video into discrete shots and the identification of keyframes, segmentation of TV News into independent news stories, automatic detection of advertisements, speech/music discrimination, shot clustering, anchorperson detection, interfaces for rapid browsing of video keyframes, text (closed caption) searching, interfaces for explicit user feedback on news stories, automatic creation of content-based news story-story links, and the development of personalisation and recommender techniques for news stories based on sparse recommendation data.

Work on the development of the Físchlár-News-Stories System is at the stage where the system is deployed and in daily use by people as part of their work and lifestyle. A formal evaluation of the completed system is ongoing and will be reported elsewhere at a later stage. In this paper we present a description of the functionality of the full system where this functionality has evolved over time and is based on needs analysis. We believe that Físchlár-News-Stories provides a set of functions which are suited to the needs of modern technology-aware users who are mobile, have multiple points of access to the Internet and most importantly wish to remain informed of news stories as they develop in their own home country, no matter where they are. Such a user will have a lifestyle which does not allow for regular daily access to local and international news at a fixed place and time because of work, travel or leisure commitments and so the Físchlár-News-Stories System provides personalised, tailored access to recent and archive news, any time, and from any place.

The remainder of this paper is organised as follows. In the next section we present a step-by-step scenario for the use of the Físchlár-News-Stories System illustrating all the functions of the system as seen from a user's perspective. In section 3 we outline the architecture of Físchlár-News-Stories as a means to explain and illustrate how such functionality is obtained. Following that we present an analysis of some related work and in a concluding section we give an outline of our future work in this area.

## 2. Usage Scenario for Físchlár-News-Stories

Firstly, let us introduce our hypothetical user. John is a regular user of the Físchlár-News-Stories System. He is a research staff member in a University, technically aware and competent, and interested in current affairs especially the recent developments in Iraq, and he usually watches the RTE1<sup>1</sup> 9 o'clock main news on TV every evening. He sometimes needs to travel abroad for a few days on business where he can access the internet away from home. As a subscribed user of the system, every morning John reads an automated email alert with the highlights of the previous night's news and a link to the system for full access (see Fig 1). On the morning in question John looks at the first three story summaries, each with a keyframe image and a text summary, and wants to see what other stories were on last night's news besides these. So he clicks the link which opens a new web browser on his desktop. On this browser screen he is presented with a calendar on the left with the most recent date being yesterday, 18<sup>th</sup> September 2003. John selects this date and that day's news stories are displayed. John can see that there were a total of 12 news stories last night, and he browses the automatically prepared summaries, each composed of the anchorperson's face with an illustrative story image projected behind her, the first few lines of what the anchorperson said and the length of the story (see Fig 2.). John's attention is attracted to the 3<sup>rd</sup> story on wounded US soldiers in Baghdad so he clicks on the story link to allow him to see the full story. A detailed summary of the story is then

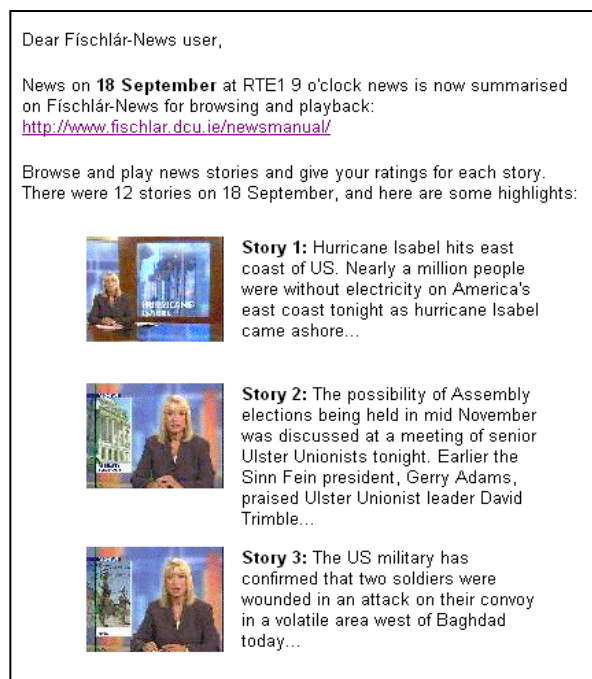


Fig. 1: Email Alert from the Físchlár-News-Stories System

<sup>1</sup> RTE1 is the main state-sponsored broadcast TV station in Ireland.

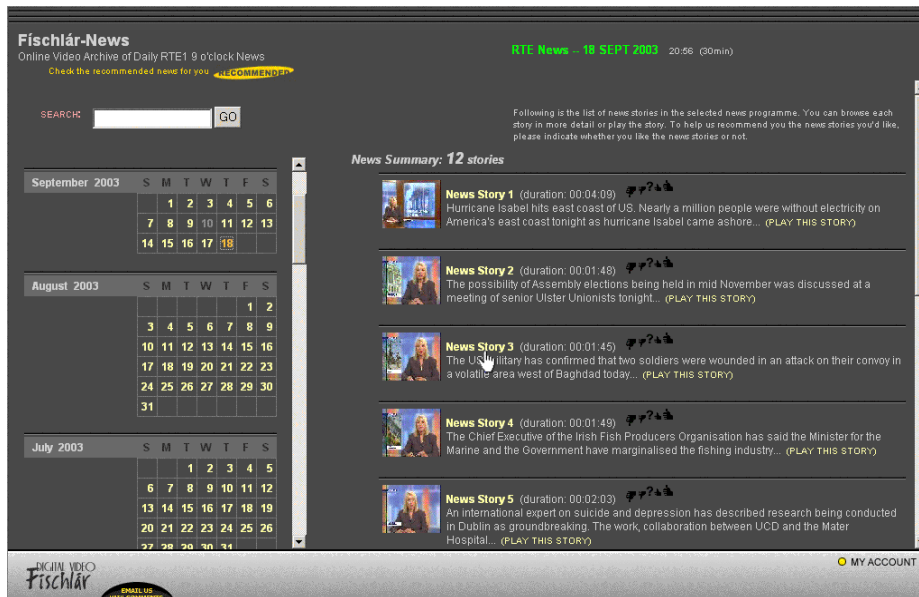


Fig 2: Físchlár-News-Stories listing of daily stories

displayed, showing automatically extracted keyframes and teletext taken from that story (see background part of Fig 3). John decides from the keyframes that this is a story he wants to watch and clicks the first image to watch the full story. A video player window opens on-screen and starts playing the story from the start. The summary in the background shows that this story is 1 minute 45 seconds long and John doesn't mind watching. He could watch the news story on full screen mode, but decides not to so he can see on his taskbar if he has any incoming emails ( Fig. 3).

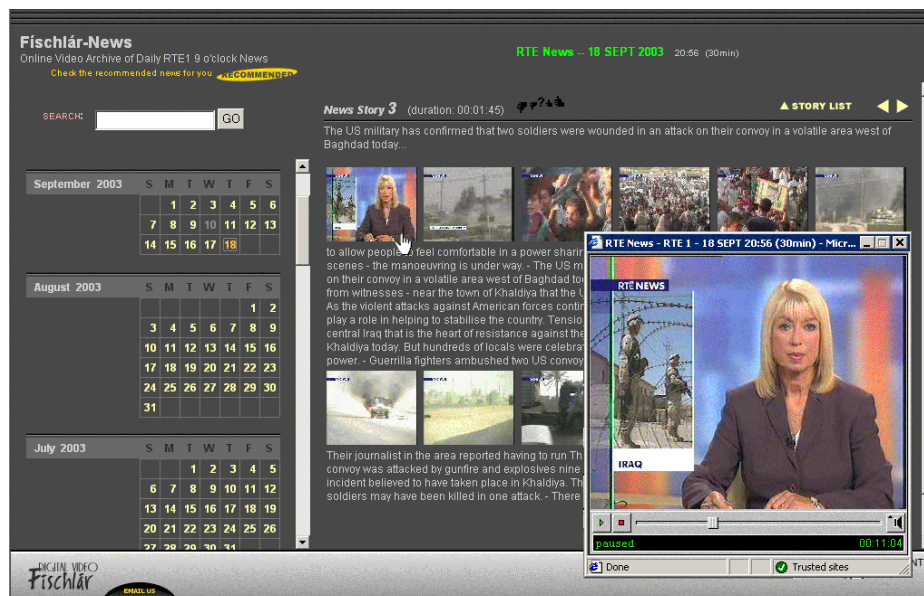


Fig. 3: Details of a full story in detail with video playback window on

After watching the story, John clicks on the "love it" icon, one of five different thumbs-up / thumbs-down icons, to indicate to Físchlár-News-Stories that he is really interested in this story so that in future the system will recommend similar stories for him. The thumbs-up icon highlights his preference for this particular story type and from now on this story is registered as an example of a preferred story.

John then wants to see a couple of other stories that followed the wounded US soldiers story so he clicks on the next story button. The next story of the day (yesterday) is about the new fees to be charged on the fishing industry by the Irish Fish Producers Organisation, and John only browses the keyframes with their synchronised text without actually playing the video, and clicks on the “dislike it” icon to indicate that he is not really interested in such a story (Fig. 4).

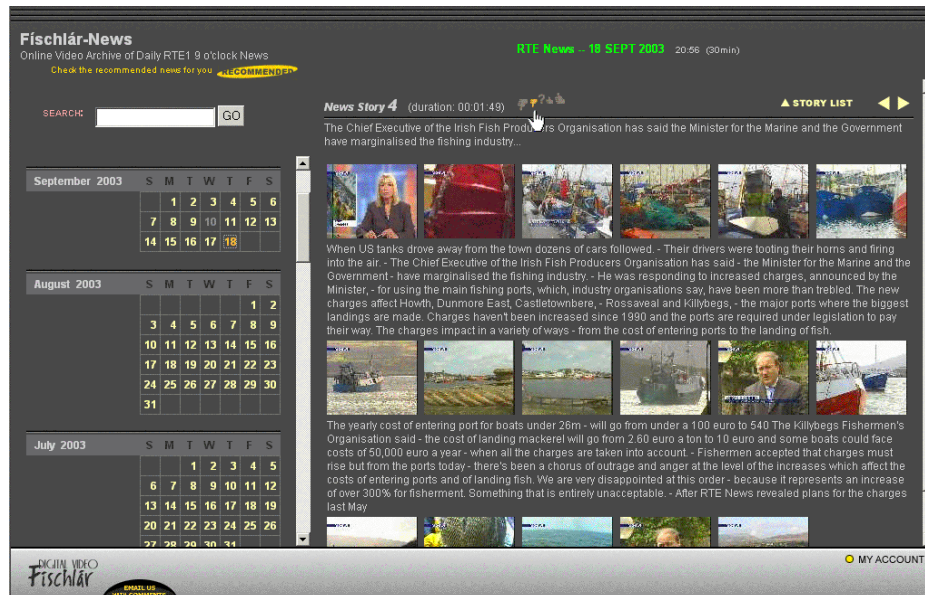


Fig. 4: User clicks thumbs-down icon to indicate dislike of a story about the Fishing Industry

He clicks on the next story which is about medical research to determine if a person can be identified as potentially susceptible to suicide attempts based on information from brain scans. John is very interested in this story, so clicks on the “love it” icon and then plays the video of a few interesting parts from this story.

John then clicks on the “story list” button to go back to the news summary, and notices the highlighted thumbs-up icons on the three stories he has browsed and played. He knows these thumbs icons will be remembered by the system and be used for later recommendation.

Later that day, John heard a colleague mention Paul Bremer’s role in the Iraqi conflict and he realises he either missed or has forgotten this story so he would like to find news stories about what Paul Bremer said and did in the Iraqi conflict. He comes back to Físchlár-News-Stories and types in “Bremer” in the text search box and clicks on the GO button. His search finds 4 news stories from different dates spanning May, July and August 2003, as shown in Fig. 5. Each story bears the date link to which John can jump to see other stories on that day.

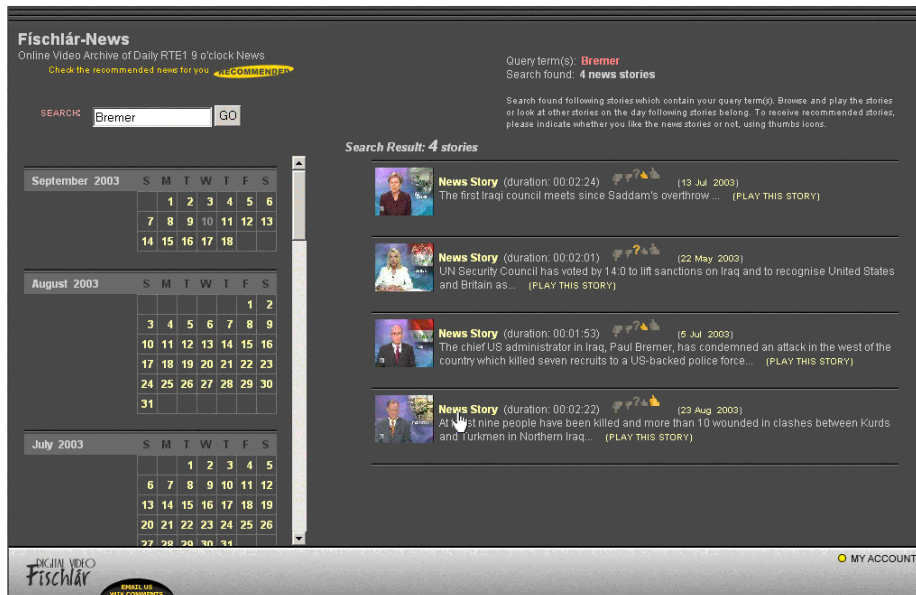


Fig. 5: Results of Text Search through News Archive

John briefly looks at the search results, and the 4<sup>th</sup> story dated 23 August mentioning 9 people killed in Northern Iraq catches his attention. He must have seen this story on the system before because the thumbs-up icon (“love it”) had previously been chosen by him and is highlighted. He clicks on that story to see its detail which is presented as before (Fig. 6). John reads the text mentioning Paul Bremer stating the bombing may have been carried out by Islamic terrorists, and he is reminded of that event.

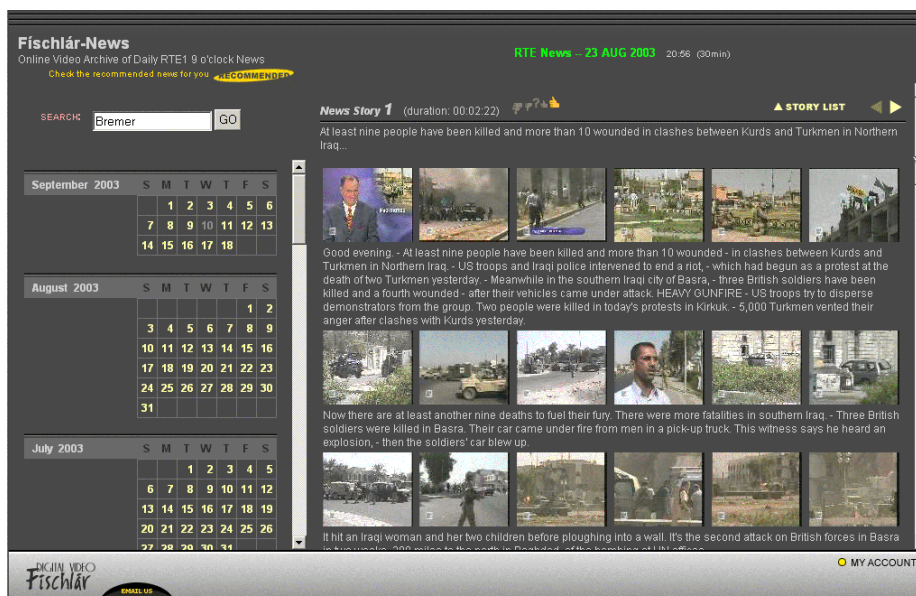


Fig. 6: Story Details

John scrolls down the story, looking at keyframes and reading text. Below the story detail, a list of **Related Stories** is presented. The stories on this list come from other dates in the news collection and they are stories that are related to the currently selected story on Bremer’s Islam terrorist statement, based on content overlap. The related stories are automatically selected at the time of story browsing so they can include the most recent additions to the news archive.

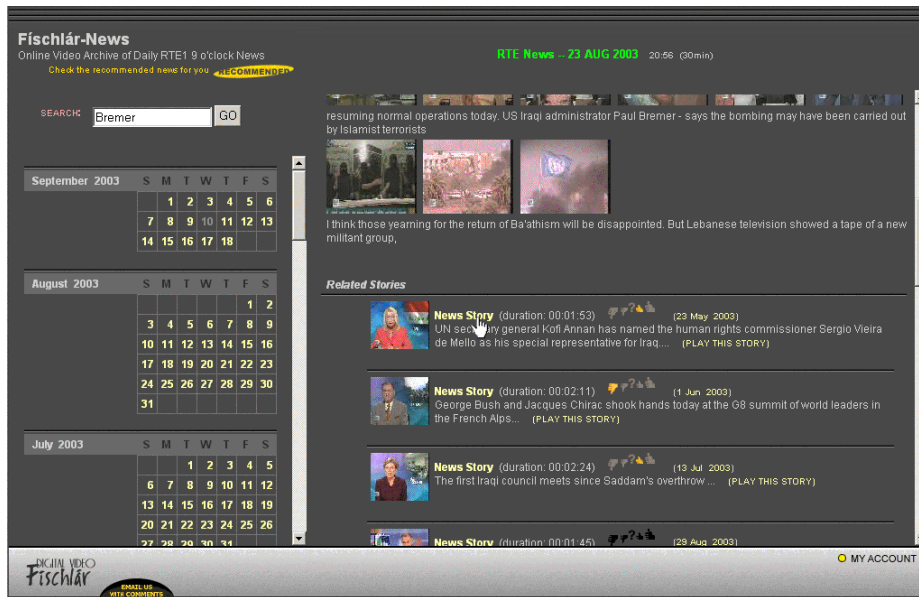


Fig. 7: List of Related Stories

John browses the related stories (Fig. 7) and is interested in the first one dated 23 May about the special representative to Iraq assigned by Kofi Anan, which he seems to have indicated as “like it” based on the “thumbs-up” icon assigned to it. He wants to check the story in detail and clicks on the story, now diverting from the initial interest in Paul Bremer. The story detail about Kofi Anan’s assignment of a special representative to Iraq is displayed and John plays some parts of the story, and finds the story very interesting.

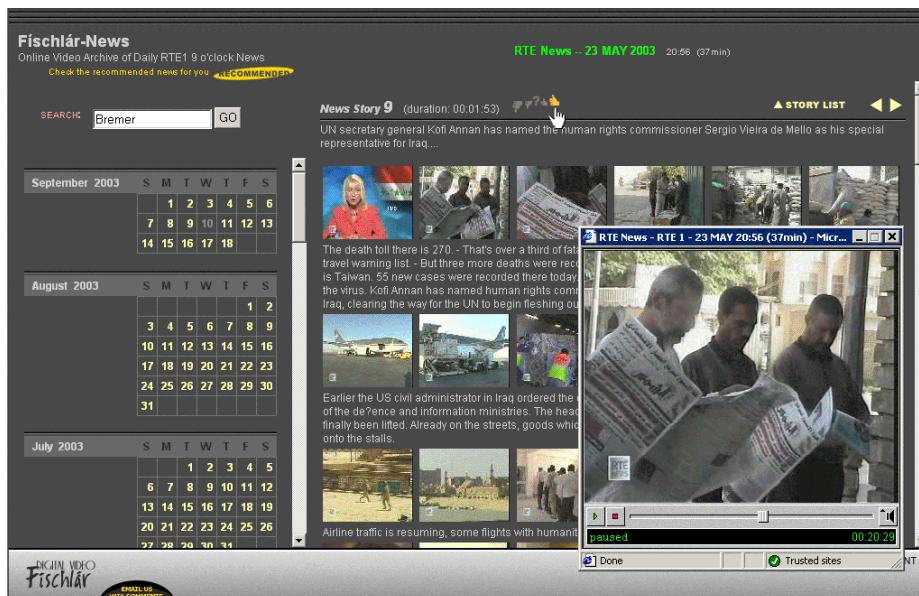


Fig. 8: Playback and Preference Indication of Related Story

He changes his previous thumbs-up preference icon into “love it” (Fig. 8).

John closes the player window, and scrolls down to see if there are any other related stories he might want to see. He notices a story labelled with a yellow “Recommendation Missed” icon (Fig. 9). This is a story from 3 July, which John has missed and is one that Físchlár-News-Stories recommends to him based on his past interests and the interests of others who use the system.

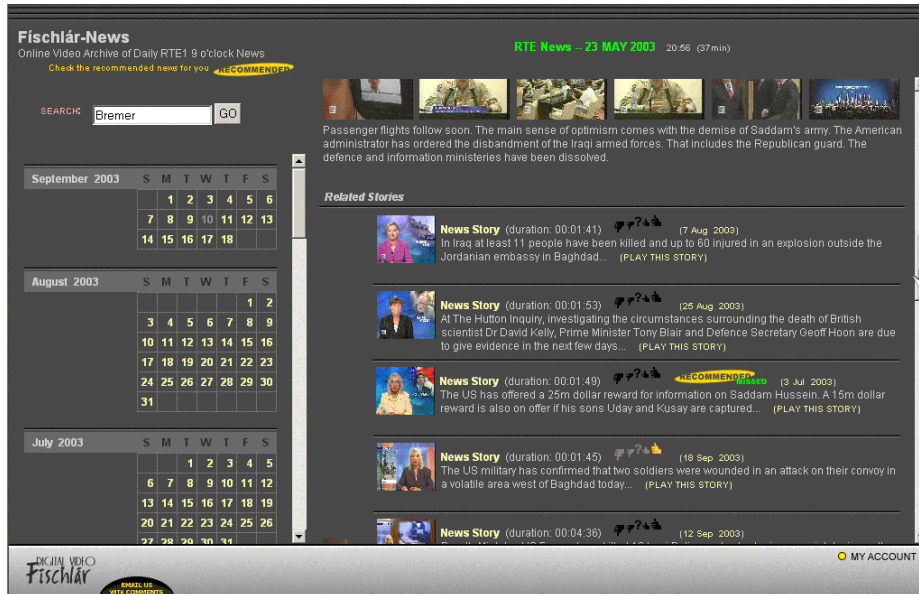


Fig. 9: Recommended Stories

John clicks on that recommended story and the recommended story's details are presented in the same format as any other story details. The story is about a 25 million dollar reward for information on the whereabouts of Saddam Hussein, which John find interesting and had heard about from a friend but hadn't seen it on the news – he obviously missed this news story in July, and he's glad to have found the actual broadcast news about this topic.

Having found this Recommended story as useful for him, John wants to know what other Recommended stories Físchlár-News-Stories has for him so he clicks on the "Recommendation" button on the top left of the screen, which will present the list of Físchlár's Recommended stories, personalised for John. These Recommended stories are selected based on his thumbs-up/thumbs-down icon preference indications from the last few months. He hopes the list will show some of the stories he missed before, or wants to see again. He also hopes this will save his searching effort, especially in the case of stories he would be interested in but has missed so far.

On the right side of the screen, the list of Recommended stories are displayed (Fig. 10). John hasn't seen any of these stories from the system before (the reason why they have been recommended by Físchlár), but notices these are of interest to him.

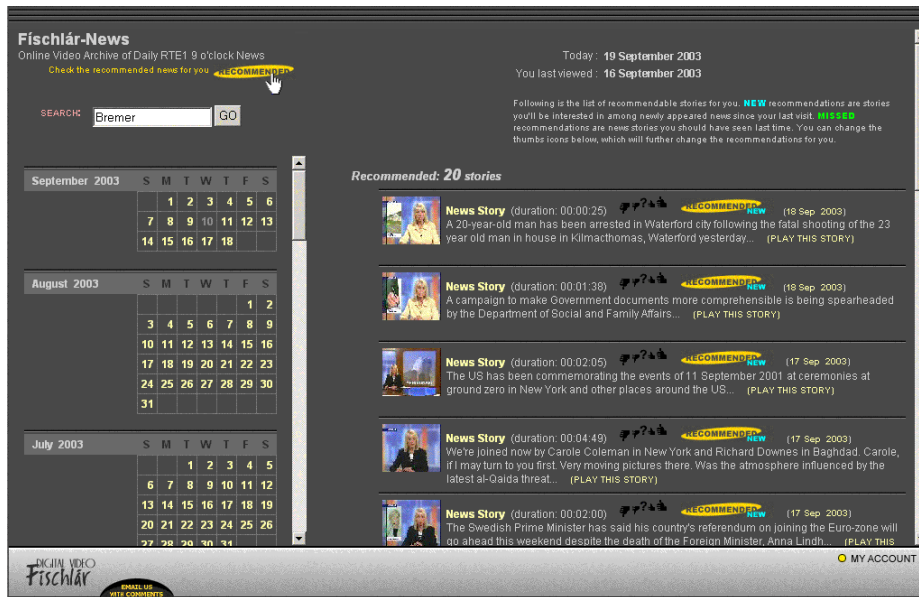


Fig. 10: List of Recommended Stories

John notices the third story in the list on commemorating 11 September, and plays that story. He continues browsing and playing news stories that have been recommended for about 15 more minutes.

When John is without his laptop, he sometimes checks recent recommendations from Físchlár-News-Stories with his XDA, a mobile device for accessing the internet over GPRS networks. He turns it on, opens a web browser, and visits the already-bookmarked Físchlár-News-Stories site. Físchlár-News-Stories detects John is using a mobile device, and displays the Físchlár-News-Stories version designed for such a mobile device (Fig. 11). A simple feature of Físchlár-News-Stories on the XDA is to present a list of recommended stories up front and allow him to browse these stories and play them (Gurrin *et al.*, 2003).



Fig. 11: Físchlár-News-Stories on a Mobile Device

The usage scenario described above is not fabricated. All screenshots are real and while hypothetical John does not exist, the system does and all the described functionality is present and working and



used. Físchlár-News-Stories is updated daily and automatically as soon as post-processing on the captured video is completed. Físchlár-News-Stories is hands-free and has been used by the authors of this paper from their desktops, their homes, and while travelling from hotel rooms, internet cafes, businesses and Universities. From a technical perspective, the only difficulties that arise with usage are to do with streaming video through firewalls which means that outside of our University environment, users usually have to view stories through keyframes and text only.

### **3. Overview of the Architecture and Functions of the Físchlár-News-Stories System**

Físchlár-News-Stories automatically records the main evening news from RTE1, the Irish national broadcast TV station. Every evening we digitally record this 30-minute broadcast, with a few minutes “padding” either side of the news to cater for late/early starts due to previous program over-runs. Before describing the architecture of our system it is worth outlining the structure of the RTE1 main evening news.

The news program begins with theme music and animated graphics followed by a voiceover of the anchorperson giving a 5 to 10 second “taster” of two or three of the main stories with outside or archive footage, or sometimes a still image as background. These tasters are then followed by the main theme music and graphic introduction with credits and we are then into the news program proper. The program is composed of between 10 and 15, but sometimes even 20 individual stories of anything from 30 seconds to 5 minutes in duration. Each story is introduced by the anchorperson with a studio shot and projected onto the background behind the anchorperson is some graphic related to the story. After the in-studio portion, the story is followed by some non-studio footage, perhaps an interview with a reporter, or a reporter voice-over some recorded footage, or sometimes some live footage, or some archive material. Sometimes the story returns to the anchorperson for further details, and then we move to the next story. There is a break in the middle of the news for advertisements with the ad break bounded by theme music with animated graphics. Towards the end of the news, especially at the weekends, we have sports news presented by a second anchorperson where the individual stories report sporting events and are much shorter in duration than the main news. The news program finishes up with the weather forecast, a weather person in front of a weather map followed by weathermap graphic with 3-day forecast.

The RTE1 evening news is fairly typical of the structure of TV news worldwide and the format is familiar and comfortable to viewers. Occasionally, in Spring or Autumn, the producers of the TV news will “spice-up” the broadcast by changing the studio décor and altering the presentation by changing the camera angle as it frames the anchorperson and this is done to remove what could become a tired look to the broadcast. These changes are cosmetic and the structure of the broadcast remains the same.

Once we digitise the news broadcast into MPEG-1 format, the broadcast is then subjected to automatic shot boundary detection as the first level of structuring each news program. Video consists of a series of still frames, usually 25 or 30 per second, which when displayed rapidly give the illusion of motion plus an audio stream, synchronised with the visual component. Video programs, including movies, sit-coms, documentaries, and TV news, are almost always composed of a series of shots, defined as a single camera motion in time, and within a shot we can have camera movement (pan, zoom, tilt) as well as object motion where objects in the video move. Shots are appended together to make programs either by joining them back-to-back where the last frame of one shot is followed directly by the first frame of the following shot (a hard cut), or by using a gradual transition (GT) from one shot to the next. Examples of GTs include fade to black and then from black into the second shot, or fade-in/fade-out, or wipes. The automatic detection of shot cuts is regarded as a first step in structuring video and usually involves some kind of comparison of adjacent frames in order to see where the major changes, indicative of a hard cut, lie.

In previous work (Browne *et al.*, 2000) we have developed several techniques for shot boundary detection using frame-frame similarity over a window of frames and with an adaptive threshold depending on the program type. We evaluated the performance of this as part of the first TRECVID in 2001 where performance figures obtained were over 90% precision and recall for hard cuts, and somewhat less for gradual transitions. However, as TV news is a live program there are generally no gradual transitions, except during the advertisement break, and so we are happy with the performance of our shot bound detection algorithm and we apply a variation to automatically structure the recording of the TV news in Físchlár-News-Stories, into shots. As part of this process, we also choose one keyframe or still image from each shot as a representative of the shot content for shot browsing.

Navigating video by browsing at the level of shots can be achieved with some degree of satisfaction when the task is to locate and retrieve shots, as in the TRECVID search task (TRECVID, 2003), but navigating archives of TV news at the level of shots is not satisfactory since users generally want to locate news *stories* rather than single shots and this is the type of information need that Físchlár-News-Stories addresses. Thus it is essential that we also structure the news broadcast into individual news stories, automatically. Story segmentation in Físchlár-News-Stories is achieved by performing several types of analysis on the digitised video, as follows:

- We cluster together all keyframes representing shots from each news broadcast. This is based on computing keyframe-keyframe similarity using known and established techniques taken from image retrieval, principally based around measuring similarity between colour and edge histograms (O'Connor *et al.*, 2001) where these are calculated globally, for the whole frame, as well as for regions within the frame. Keyframes representing shots taken from the same cluster will be visually similar and will tend to be taken from the same news story. For example, a story featuring an anchorperson followed by an outside broadcast which cuts between an interviewer, an interviewee, and some footage of a crowd protesting at a rally, will have shots of the crowd in the one cluster, and the interviewer and interviewee in other clusters.
- One of these clusters will have an average keyframe-keyframe similarity much higher than the others and this will most likely be a cluster of anchorperson shots which are visually very similar. Knowing where anchorperson shots are will help with story segmentation.
- We apply a speech-music discrimination algorithm to the audio which will help to delineate the exact beginning of the news broadcast, as well as the beginning/end of the advertisement break.
- We also apply a technique we have developed to detect individual advertisements, reported in (Sadlier *et al.*, 2002).
- Finally, we use the lengths of the shots as clues for story bound segmentation, namely the fact that outside broadcasts and footage video tends to have shorter shot lengths than the in-studio broadcasts.

We combine the outputs from these analyses together using a trained support vector machine (SVM) to determine story bounds where the SVM has been trained on the manual tagging of story bounds from several weeks of news broadcasts. In 2003, story bound segmentation formed one of the four tasks which was evaluated as part of the annual TRECVID exercise (TRECVID, 2003), the others being shot bound detection, feature extraction, and searching for shots. We submitted our algorithm for evaluation but were disappointed with the results obtained in this formal setting where we achieved a best of 31% recall and 45% precision. and a forthcoming publication shows far better results (O'Hare *et al.*, 2004). For the present time, the automatic segmentation is manually checked for accuracy every day and if needed, story bounds are adjusted.

As well as capturing the audio and video from each news broadcast, we also automatically capture closed captions, also known as teletext. Closed captions correspond to the dialogue spoken during TV broadcasts. For post-produced programs (i.e. anything not transmitted live) this is accurate and synchronised well with the audio and video. For live broadcasts such as sports commentary, chat shows, and TV news, a stenographer types in the dialogue in real time and depending on the skill of the individual the resulting text can be accurate, correctly punctuated and correctly capitalised. If the

news reader speaks fast or the stenographer is slow, then the resulting text can have spelling errors, can omit phrases or sentences in order so as not to fall behind, and the presentation of the text can lag behind the audio-visual signal by several seconds. However, closed captions are becoming ubiquitous in US and many European countries, and provide the basis for much of the content-access we have developed in Físchlár-News-Stories.

Having the closed captions means that for each news story we have the audio and video, plus the text of what is spoken, plus an anchorperson keyframe with content-indicating background image, plus the keyframes from the non-studio materials. We have considered performing a more detailed analysis of the video to help with story retrieval such as image matching in order to identify repeated transmission of the same footage, indicating related stories, or incorporating a face recognition algorithm in order to automatically recognise people regularly in the TV news, but we decided against that as for news story retrieval the content that really matters is the subject matter and the closed captions give us this information.

Having gathered much information from the analysis of a news broadcast we store this information in a database for faster access. In order to represent a news story we have the anchorperson keyframe and we use the first few lines of closed captions as the summary. We also have the set of keyframes and the closed caption text and this means that when the system displays an individual story we can display keyframes and text synchronised.

We provide 4 different ways for users to access individual news stories as described in section 2. Browsing stories by calendar date is straightforward and uses story metadata and story bounds to identify stories associated with each date. For content-based search based on a text query we use a simple IR search engine to rank stories. We have not incorporated support for phrase retrieval and this is something that we should probably do. When a story's detail is displayed, we use the closed caption text from that story as a query against the closed caption archive and we display summaries of the 10 top-ranked stories. There are issues raised by the fact that the story bounds for audio-visual will be offset to different degrees to the story bounds for the closed captions due to the variable amount of time lagging that the live closed captioning suffers from. However, these potential difficulties have not materialised and the generation of links to related stories based on the approach we have taken seems to work well. The choice of displaying the top-10 ranked stories is a subjective guess and we should perhaps be more selective in deciding how many related stories to show, and we will discuss this point later.

There are many places during a user's interaction with Físchlár-News-Stories where we can capture explicit user feedback and ratings on a given news story using a 5-point scale. These ratings are used as input to a collaborative filtering system which can recommend news stories to users based on ratings from other users. In addition to explicit feedback we also capture implicit feedback through logging when users actually play news video and we use this as a positive indicator of preference and interest for such news stories. One of the characteristics of the TV news domain that makes it challenging for a personalisation and recommender system is the fact that there is a requirement for recommendations on new content (today's news stories) even before it has been rated by any users. In addition, we have a situation where the user vs. stories ratings matrix is very sparse as most users will not rate many stories. To work around these constraints we can use story-story similarity as well as user-story ratings in order to predict which stories, as yet unrated by any user, are likely to be of interest to different users. Also, because we track individual users and their history of access, we are able to divide up recommendations for each user into those which are "unseen" as they have been added to the archive since that user's last session with the system, and those which are "missed", namely, stories which are recommended and have a transmission date previous to the user's last session with the system. Anecdotal evidence to date suggests that users find it helpful to make this distinction among stories recommended by the system and we are presently conducting a formal, larger-scale user test of this.

The internal architecture of all Físchlár systems is that descriptors from analysis are formatted according to the MPEG-7 standard and user interaction with the system generates XML documents which are transformed to HTML using XSL stylesheets. This means that we can generate interfaces to the Físchlár-News-Stories system from a number of different devices without having to change the system at all, just by generating an XSL stylesheet for each supported device. We have stylesheets for a desktop computer, for an XDA GPRS-enabled mobile PDA and for a Compaq iPAQ PDA. On the mobile devices the stylesheet filters out much of the functionality such as searching and related links and we allow access via calendar story listing, and personalised recommendations (Lee *et al.*, 2002). For playback from mobile devices we have transcoded the recorded video from MPEG-1 into RealVideo for playback on the mobile platform, and we use an MPEG-1 video streamer for playback onto the desktop platform. The generic architecture for Físchlár-News-Stories is shown in Figure 12.

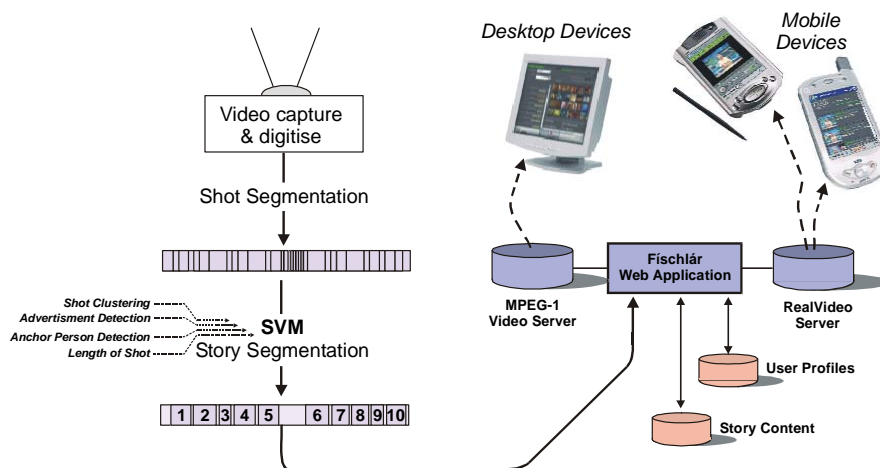


Fig. 12: Architecture of Físchlár-News-Stories

At the time of writing, eleven months of recorded daily RTÉ1 news is available online within the Físchlár-News-Stories archive in addition to another 2 years which we have digitised but not made available. The archive is made available to University staff and students, and is accessible from any computer lab, the University library and student residences from within the campus.

#### 4. Comparison to Related Work

Any analysis of the strengths and weaknesses of Físchlár-News-Stories should include a comparison to similar systems and the foremost system which also provides content access to an archive of TV news is Informedia, developed at Carnegie Mellon University (Hauptmann and Witbrock, 1998) and (Christel *et al.*, 1999). Informedia and Informedia II, provide access to over 1,500 hours of broadcast TV and radio news content (compared to our 250 hours). Informedia includes speech recognition, captures closed captions and recognises text appearing on screen and applies natural language processing techniques to segment and index video as well as provide intelligent video search, navigation and retrieval. Informedia concentrates less on the whole news story as the discrete information unit and provides summaries for story segments, which could be whole stories, via headlines, filmstrip story boards and video skims. Informedia supports retrieval of these video paragraphs or story segments/summaries based on text-text matching using speech recognition, closed captions or text overlaid on the video. Informedia also does face recognition and speaker voice identification and can support retrieval based on known faces. It also interprets proper names, places, dates and times at the indexing stage as well and tries to resolve ambiguities associated with these.

One of Informedia's greatest strengths is its browsing interface which leverages the large amount of video interpretation and analysis done at indexing time. For example, the geographic location based

on place recognition from selected video segments can be used to find other video segments which mention the same or close locations; the recognition of a newsworthy face can be used as a link to other video segments where the same face appears; a filmstrip summary of a group of video segments from different periods in time can be plotted on a timeline to see the relationship between segments and so video summaries can extend across single video documents. Finally, a user's query can be multidimensional in that it may combine image elements, video clips, text and speech.

What Informedia offers is a far more mature system than Físchlár, where the developers have had opportunity to integrate advanced techniques from computer vision which we have not done yet. While we have concentrated on stories as the unit, Informedia concentrates on smaller units of video information and Informedia is more multi-modal, supporting cross-media retrieval while Físchlár is based on text only. The amount of video, audio and text analysis performed in the Informedia system is far greater than in Físchlár and the video summarisation appears to be better. To our advantage, however, we offer story-based navigation which is more suited to the user who wishes to retrieve news rather than news shots. By focussing on the story as the unit we have also been able to personalise content and to include user preference feedback and recommendations which would not have been possible where the unit of information was smaller than a story.

Another contemporary system which provides online access to broadcast TV news is the ANSES system developed at Imperial College London (Pickering *et al.*, 2003). Here, the main evening news from the BBC is digitised and broken into stories, just as in Físchlár-News-Stories. Closed captions are also digitised and much effort is put into analysing the text in order to currently identify organisations, person names, location names, and dates. From this, a text-based search is supported which allows a user to specify constraints on the search based on the identified categories. Físchlár-News-Stories and ANSES share the common concept of the story as the unit of information but the browsing interface, recommender system and story-story linking in Físchlár-News-Stories makes it a more developed system.

One of the four tasks in the annual TRECVID exercise in 2003 (Smeaton *et al.*, 2003) and (TRECVID, 2003) was to provide shot-level retrieval from a corpus of over 60 hours of CNN and ABC news and there were several groups, including CMU and Imperial College, to develop or modified existing video retrieval systems to complete this task. However, while they did provide video retrieval from an archive of TV news, the task was to locate shots and not news stories or even news items and thus systems which took part in the 2003 TRECVID cannot really be compared to Físchlár-News-Stories except indirectly.

The main strength of Físchlár-News-Stories is that it offers a way for users to search, digitally, for information content to which they previously had no access at all, unless they had recorded the TV news themselves. Thus a part of the novelty and advantage of Físchlár-News-Stories is in the fact that the TV news video is available to users for playback at any time, and part of the advantage is the set of functions of the Físchlár-News-Stories system itself. Personal video recorders (PVRs) like TiVo and SKY+ are becoming more commonplace yet these do not offer any degree of sophistication in browsing TV news. From a technical perspective, Físchlár-News-Stories is robust and works well and shortcomings of the system in terms of functionality not offered.

## **5. Conclusions and Future Work**

Físchlár-News-Stories is still a work in progress and while it already has a good deal of functionality, there are a range of other features we would like to develop and incorporate as future work. While Físchlár-News-Stories generates links to related stories these are based on a 1-step lookahead computed at presentation time. This is a kind of local optimisation with no global view of the archive or the content links that exist beyond the similarity between the current story, and the rest of the archive. Obviously some kind of clustering of related news stories would address this concern but this

would have to be done in real-time which has already been shown in another application as possible (Smeaton *et al.*, 1998).

There is a definite need for some more global examination of the links between related stories in the TV news archive. By clustering together related news stories this would open up the possibility of generating summaries of groups of stories as opposed to summaries of individual stories. This has been investigated in (Ide *et al.*, 2003). who report a mechanism whereby such information can be presented but who have yet to deploy this in a real TV news setting. Generating summaries of groups of news stories would allow us to incorporate elimination of redundancy in TV news which is both a good and a bad thing as one user's redundancy is another user's alternative view !

Besides presenting a user with access to TV news from a single news source, we would also like to include topical news from other news sources, namely from other TV channels and from online newspapers. A single news event can be reported differently by different TV stations and can be reported in greater detail in newspapers and other online sources and users frequently track interesting news topics over multiple news sources. This could be achieved quite simply and crudely by including hot links to one of the many news portals which when clicked generate searches to those portals. A more sophisticated approach would be to poll such news portals as each new TV news story is generated, gather information on related stories at that time and then create clusters of related stories.

One thing we have not addressed in the current Físchlár-News-Stories system is the shifting nature of user preferences and interests over time. Even in our worked scenario in section 2 of this paper we encountered a situation where our hypothetical user changed a preference rating for a previously viewed story from a rating assigned in the past. It may be that user ratings need to be timestamped and older ratings weighted less importantly than more recent weightings. This is a problem faced by all personalisation and recommender systems in all application domains.

Finally, the major outstanding task is in our work on the Físchlár-News-Stories system is to carry out a user evaluation. Some of the individual components of the system have been evaluated and benchmarked, such as the shot boundary detection, story boundary detection, story browsing via keyframes and the techniques used to make recommendations but the package as a whole will be evaluated from a user perspective and this work is ongoing.

### **Acknowledgements:**

The support of the Informatics Directorate of Enterprise Ireland is gratefully acknowledged. This material is based upon work supported by the IST programme of the EU in the project IST-2000-32795 SCHEMA. Part of this material is based on works supported by Science Foundation Ireland under Grant No. 03/IN.3/I361

### **References**

- 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 IMVIP 2000 - Irish Machine Vision and Image Processing Conference*, Belfast, Northern Ireland, 31 August - 2 September 2000.
- Christel, M., Warmack, A., Hauptmann, A., Crosby, S. (1999). Adjustable Filmstrips and Skims as Abstractions for a Digital Video Library. *IEEE Advances in Digital Libraries Conference 1999*, Baltimore, MD. pp. 98-104, May 19-21, 1999
- Gurrin, C., Browne, P., Smeaton, A.F., Lee, H., McDonald, 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* February 16-18, 2004 Innsbruck, Austria

- Gurrin, C., Smeaton, A.F., Lee, H., Mc Donald, K., Murphy, N., O'Connor, N. and Marlow, S. (2003). Mobile Access to the Físchlár-News Archive. In *Proceedings of Mobile HCI 2003 - 5th International Symposium on Human Computer Interaction with Mobile Devices and Services, Workshop on Mobile and Ubiquitous Information Access*, Udine, Italy, 8-11 September 2003.
- 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*, Santa Barbara, CA, 22-24 April, 1998.
- 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 Spring Symposium on Intelligent Multimedia Knowledge Management*, Stanford University, 24-26 March, 2003.
- 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*, Vol.2 No.4, May 2002.
- Lee, H. and Smeaton, A.F. (2002). Searching the Físchlár-News Archive on a Mobile Device. In: *Proceedings of ACM SIGIR 2002 Workshop on Mobile Personal Information Retrieval*, Tampere, Finland, August 2002.
- O'Connor, N., Czirik, C., Deasy, S., Marlow, S., Murphy, N. and Smeaton, A.F. (2001). News Story Segmentation in the Físchlár Video Indexing System. In: *Proceedings of ICIP 2001*, Thessaloniki, Greece, 7-10 October 2001.
- O'Hare, N., Smeaton, A.F., Czirik, C., O'Connor, N., and Murphy, N. (2004). A Generic News Story Segmentation System and its Evaluation. To appear in *Proceedings of ICASSP 2004 - IEEE International Conference on Acoustics, Speech, and Signal Processing*, Montreal, Quebec, Canada, 17-21 May 2004.
- O'Sullivan, D., Wilson, D.C. and Smyth, B. (2003). Preserving Recommender Accuracy and Diversity in Sparse Datasets. In *Proceedings of the 16th International FLAIRS Conference*, Florida, USA, 2003.
- Pickering, M., Wong, L., and Ruger, S. (2003). ANSES: Summarisation of News Video. In: *Proceedings of CIVR-2003*, University of Illinois, IL, USA, July 24-25, 2003.
- Sadlier, D., Marlow, S., O'Connor, N., and Murphy, N. (2002). Automatic TV Advertisement Detection from MPEG Bitstreams. *Journal of the Pattern Recognition Society*, 35(12), 2002.
- Sadlier, D., O'Connor, N., Marlow, S., and Murphy, N. (2003). A Combined Audio-Visual Contribution to Event Detection in Field Sports Broadcast Video. Case Study: Gaelic Football. In *Proceedings of ISSPIT'03 - IEEE International Symposium on Signal Processing and Information Technology*, Darmstadt, Germany, 14-17 December 2003.
- Smeaton, A.F. (2004). Indexing, Browsing and Searching of Digital Video. In *ARIST - Annual Review of Information Science and Technology*, Vol. 38, Chapter 8, pp 371-407, 2004.
- Smeaton, A.F., Burnett, M., Kellely, F. and Quinn, G. (1998). An Architecture for Efficient Document Clustering and Retrieval on a Dynamic Collection of Newspaper Texts. In *Proceedings of the 20th BCS-IRSG Colloquium*, Grenoble, France, Springer-Verlag Workshops in Computing, April 1998.
- Smeaton, A.F., Kraaij, W., and Over, P. (2003). TRECVID 2003 - An Introduction. In *Proceedings of TRECVID 2003 - Text REtrieval Conference TRECVID Workshop*, Gaithersburg, Maryland, 17-18 November 2003.
- TRECVID (2003). Guidelines for the TREC Video Track. <http://www-nlpir.nist.gov/projects/t01v/>. Last visited December 2003.