Collaborative Searching for Video Using the Físchlár System and a DiamondTouch Table

Alan F. Smeaton   Colum Foley   Cathal Gurrin   Hyowon Lee

Sinéad McGivney

Adaptive Information Cluster and Centre for Digital Video Processing,

Dublin City University, Glasnevin, Dublin 9, Ireland.

Alan.Smeaton@DCU.ie

Abstract

Físchlár-DT is one of a family of systems which support interactive searching and browsing through an archive of digital video information. Previous Físchlár systems have used a conventional screen, keyboard and mouse interface, but Físchlár-DT operates with using an horizontal, multi-user, touch sensitive tabletop known as a DiamondTouch. We present the Físchlár-DT system partly from a systems perspective, but mostly in terms of how its design and functionality supports collaborative searching. The contribution of the paper is thus the introduction of Físchlár-DT and a description of how design concerns for supporting collaborative search can be realised on a tabletop interface.

1 Introduction

In this paper we are concerned with developing an horizontal, interactive, multi-user, tabletop interface to a video search and browsing system. Within our research group, we have amassed expertise in building systems which analyse, index and provide searching and browsing on video archives. These have been built, tested and deployed for a variety of applications under the generic name of Físchlár. In this paper we describe the construction of another Físchlár system called Físchlár-DT which users access using a DiamondTouch tabletop device.

This paper is thus a systems paper in that we describe an operational system, its architecture, design and functionality. However rather than describing Físchlár-DT from a purely systems perspective we present it in terms of its design and how it supports *collaborative* searching among users. The rest of the paper is organised as follows. In the next section we give a summary of related work on collaborative searching and then summarise the Físchlár system developed for a conventional user interface. We then out-

line the DiamondTouch hardware and DiamondSpin software toolkit which we used. Following that we discuss what we believe to be the most important aspects of design for supporting collaborative search and in Section 6 we present the Físchlár-DT system in terms of how it realises those design features. Our experimental evaluation of Físchlár-DT in terms of the annual TRECVid benchmarking exercise is then outlined and we conclude the paper with a plan for future work.

2 Related Work

Collaborative concerns in the context of *information searching* have been discussed previously in several areas and there has been a small but significant amount of related work which addresses this. Hansen and Jarvelin [10] observed various collaborative activities happening in a patent office when users search for patent-related information, both mediated by physical or electronic form or by human. They noted that collaborative activities are an important characteristic of the information searching process in task-based information retrieval. A novel collaborative system in [2] used a query formulation process whereby the participants discussed and re-arranged physical tokens on subject trays on a table. As they re-arrange the tokens, the system automatically retrieved documents with the arranged tokens as a formulated query, and participants viewed them and continued to discuss and re-arrange the tokens to refine the query. Collaborative information searching was also explored in [25], noting how people continue to work collaboratively even though the single-user interfaces hinder collaborative information searching. Various facets of collaborative searching were discussed including instructional collaboration (between a subject librarian and a user) vs. mutual (among library users themselves), product-related (communication on searched documents) vs. process-related (communication on how to search). A

collaborative interface was developed [24] that helped users to explain their search history and get assistance from a subject librarian.

The Pond [21] is an informal information seeking environment in which users can search for their favourite music using a tabletop device and predefined query tags. The use of audio is an important feature of The Pond as the system associates particular sounds with user actions. Personal Digital Historian [18] is a collaborative application running on a DiamondTouch table, providing search features for a digital photo collection. *Who* and *What* views show thumbnail size photos linked with a similar and related person or things; *When* view shows a timeline from which the users can tap to see more photos from a particular period in time; *Where* view shows a geographical worldmap from which the users can tap on a location to see photos taken at that place. This browsing-oriented interface uses various implicit AND and OR boolean combinations when some of the views are used together.

3 Físchlár Systems for Video Retrieval

As part of our research into browsing and searching of digital video content, we develop digital video archive systems that we collectively refer to under the name Físchlár. Several systems have been developed, including: Físchlár-TV, a searchable and browsable archive of TV programs [13]; Físchlár-News, a searchable archive of more than two years of TV news stories [20]; Físchlár-Nursing, a browsable video library developed specifically for the classroom environment [7] and Físchlár-TREC, a set of experimental retrieval systems developed annually [4, 3, 6].

To benchmark the effectiveness of our techniques, we take part in the annual TRECVid benchmarking activity, the goal of which is to promote progress in content-based retrieval from digital video archives using an open, metrics-based evaluation. TRECVid, run by the National Institute of Standards and Technology in the U.S., was first introduced in 2001 and by 2005 has 62 participating groups from research labs throughout the world. Participants can partake in several tasks, though in this paper we are concerned with the task of interactive searching in which system evaluation is based on the performance of human users interacting with a video retrieval and browsing system. In order to support evaluation across participating groups, TRECVid provides all participating groups with a source video corpus (150 hours of broadcast TV news in 2005), along with 24 queries (topics) to be run against the corpus. An example topic could be "Find shots zooming in on the U.S. Capitol building in Washington DC".

The interactive video retrieval system we employed for TRECVid in 2004 [4] has an XML-based architecture using MPEG-7 compliant video descriptions internally, with
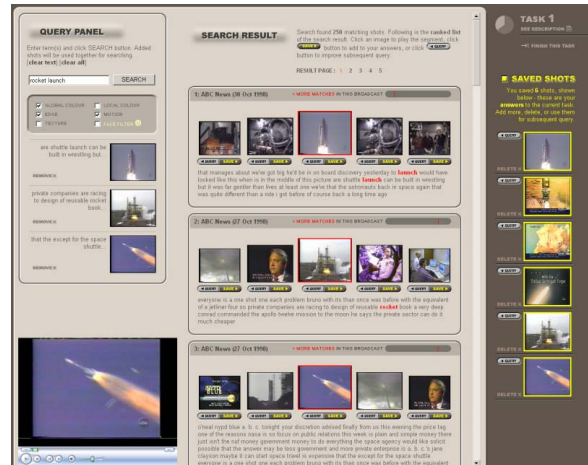

Figure 1. Físchlár-TRECVid 2004 web-based user-interface: for a single user

a web-based interface. Search and retrieval is based on both text and image evidence which has been automatically extracted from the source video. The unit of retrieval used for TRECVid experiments is the shot, which is extracted from the video by an automatic process called Shot Boundary Detection. We represent each shot with a single key image (a keyframe) that represents the content of that shot. For each keyframe, we extract MPEG-7 features (two colour-based, texture and edge) and construct parallel indices for each of the four visual features. This allows us, for any given keyframe or group of keyframes, to identify the most similar keyframes from the video source (in rank order) for one, all or any combinations of the MPEG-7 keyframe features.

The textual aspect of our video retrieval is based on converting the spoken words from the source video into text and associating the spoken text with the shot from which it came. The result of processing the text and image evidence, is that Físchlár-TREC supports a user searching for video content using text queries (like Google), or image queries using example keyframes, or even a combination of both.

In our previous TRECVid experiments (from 2001 until 2004) the approach we have taken has been to let a single user evaluate the performance of the system for each of the 24 topics (See Figure 1). In 2005, we are participating in TRECVid, using a similar architecture to Físchlár-TREC2004 and similar search engines, providing our underlying video retrieval service. However instead of a keyboard/mouse/screen interface we are using a Diamond-Touch tabletop from MERL as a front-end in order to support collaborative searching for video. Our reasons for doing this are to explore the potential of a horizontal tabletop interface for supporting video searching, and secondly to

explore how this environment can be used for collaborative searching by a team of users searching together.

4 The DiamondTouch and DiamondSpin

4.1 The DiamondTouch Tabletop

DiamondTouch [5] is a tabletop input device developed by Mitsubishi Electric Research Labs (MERL) as a research prototype. The device supports multi-user, collaborative face to face interaction for up to four people. The Diamond-Touch is commonly used in conjunction with a projector to render a PC's display directly onto its surface, enabling users to directly manipulate objects on screen. When a user touches the tabletop surface a circuit is completed which runs from the transmitter in the table through the user to a receiver in the users chair and finally back to the transmitter, thus each touch on the input surface can be associated with a particular user. There are two versions of the system, the DT88 with a 79cm diagonal and DT107 with a 107cm diagonal and both operate with a 4:3 aspect ratio. The Físchlár-DT system described in this paper uses the DT107 DiamondTouch tabletop.

The surface consists of rows and columns of antennae in a diamond arrangement, with each row or column connected in one direction whilst isolated in the other. This multilayered arrangement ensures maximum surface area whilst minimising interference and also allows Diamond-Touch to be debris tolerant; objects left on the tabletop surface do not interfere with users' normal interaction. To enable user identification, DiamondTouch requires reasonable electrical isolation between users, i.e. if two users were to touch this inhibits the system's ability to uniquely identify a particular input, however social norms of 'personal space' have been sufficient to keep interference between users manageable.

4.2 DiamondSpin API

DiamondSpin [19] is a Java software toolkit designed by researchers at MERL for the proficient prototyping and experimentation of applications on interactive, multi-user, collaborative displays. It investigates interaction techniques for Single Display Groupware that enables concurrent multi-user touch-based manipulation of documents (in our case keyframes taken from the video footage). This API may be used for system development on tabletops which are rectangular, circular or octagonal.

The toolkit provides a real-time polar to Cartesian conversion engine, which allows the random positioning and orientation of keyframes around the surface of the DiamondTouch tabletop. The motivation for developing an interface using a polar coordinate system is that a polar


Figure 2. Object rotation and orientation

system does not have a principle direction for displayed keyframes, unlike a Cartesian coordinate system. Images are instead displayed in relation to one and only one meaningful centre and all sub-areas of the interface must be aware of where this centre is at all times.

DiamondSpin is composed of two key concepts:

1. Translation of the origin of the Cartesian display, which is generally at the top left or bottom left corner, to the centre of the tabletop

2. Each keyframe on the tabletop has 3 degrees of freedom - d, α and β. These support the tabletop-specific handling of keyframes.

Figure 2 shows an object displayed at a distance d to the centre O at an angle α. The object's β value can then be used to rotate it around its own centre. DiamondSpin uses multiple threads to handle multiple simultaneous user-input actions. A single thread for each user handles the user interactions with tabletop components while another handles the repainting of the user interface, showing changes made to the surface.

The DiamondSpin toolkit is implemented in Java 2D classes with JAI (Java Advanced Imaging) and JMF (Java Media Framework). The software must handle extremely demanding user activities (such as rotation of the entire tabletop screen). The technique for handling this in the DiamondSpin architecture is to provide a multi-layer representation (in this case, four layers) with a multi-depth, multi-repaint structure. The bottom layer (Layer 3), is comprised of non-interactive elements e.g. a background image. Layer 2 components have the potential to become active. At least one active component, which currently receive user input actions such as the rotation of documents/keyframes are contained in Layer 1. Finally, Layer 0 contains rotation-sensitive components e.g. menu bars. This structural design enables the selective refresh of parts of the tabletop surface, allowing the application to use only the layers it requires.

5 Collaborative Interaction Design for Searching on a Tabletop

Before designing a specific scheme for collaborative interaction, we start by considering what are the major issues that we wish to explore in the system to be developed. This results in a better understanding and enumeration of possible schemes which we can use to implement the actual interface that supports collaboration. Here we consider three important elements in collaborative interaction design for a tabletop search system which we have identified from available literature on tabletop design, groupware, Computer Supported Cooperative Work (CSCW) and our own experiences in developing Físchlár-DT. These are:

1. Task allocation for each collaborator;

2. The degree of group awareness;

3. The degree of coordination policy;

Each of these will be briefly described in the context of video searching on a tabletop.

5.1 Division of Labour: Task Allocation for Each Collaborator

Providing a collaborative interface requires careful consideration of the way the overall tasks should be allocated to each user in the collaboration. In the case of searching for video in TRECVid 2005, this concern can be summarised as the following question: *given the task of finding as many relevant video shots as possible in a given time, how should this task be divided into sub-tasks which can be allocated to each user in a way that ensures collaborative benefit ?*

Different kinds of collaboration in information searching have been observed in a library setting, for example *joint search* where a group of users work around a single terminal and *coordinated search* where a group of users work on two or more adjacent terminals discussing what they are doing [25]. Other collaboration styles include working in parallel, working sequentially in tightly coupled activities, and working independently [15].

In the context of searching task, we can use a typical stages of action model in information search [11] as shown on the left of Figure 3. On the right side of Figure 3, we illustrate a few possible allocation schemes among two users, including User A conducting the first half of the task then passing the result to User B to conduct the rest of the search (Task division 1); Task division 2, where both users are conducting all stages in parallel, either by working with separate (duplicate) tools for each user or by sharing a single set of tools. Task division 3 depicts a hybrid collaboration in which both users start working together in clarifying and
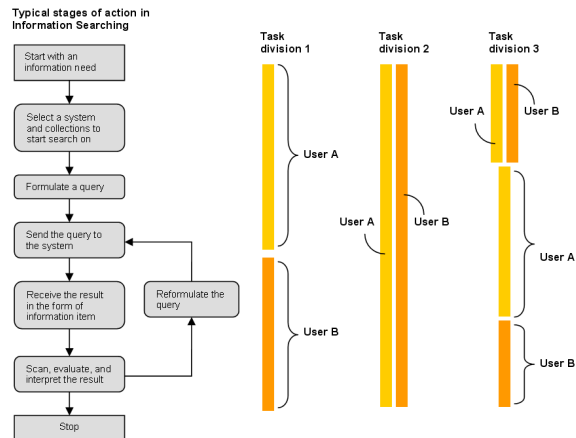

Figure 3. Stages of action in information searching and some possible task allocations schemes


Figure 4. The degree of group awareness and some techniques

agreeing on the exact information need, then working separately in the actual execution of searching.

5.2 The Degree of Group Awareness

Whether using a distributed (remote) or co-located setting, the ability to be aware of what the other user is doing is an important consideration in collaborative interaction design. However, compared to a distributed setting, in a co-located tabletop setting the group awareness comes more naturally since the collaborators are physically close around the table. For example, In User A's peripheral vision, she constantly informs herself of how actively User B is working or how frequently tapping on something.

In a tabletop setting, more gestural or physical/direct manipulations (e.g. dragging) tend to allow better group awareness than more symbolic manipulations (e.g. menu item selection, button clicking, short-cut keys), but with a reduced individual efficiency or power [8], and this is a trade-off the designer should decide on. There are vari-

Figure 5. The degree of software coordination policy and some techniques

ous techniques that can help enhance and maintain group awareness to different degrees. Figure 4 illustrates this in the spectrum between having no explicit awareness feature and having strong awareness cues, including techniques in between. In the case of a distributed setting, provision of awareness techniques is more crucial as there are no naturally occurring awareness cues. Example techniques showing a short history of the cursor's past movement or tracing [9, 1] or a semi-transparent/summarised menu appearing when opened by a remote user can help maintain awareness [8]. In the case of a co-located tabletop, simply by bringing users close together is already an advantage in terms of raising awareness. Example techniques to further help awareness are distinctive sound effects when an action is triggered, exaggerated animation (such as blowing up and exploding when an object is deleted), or limiting user actions by dragging [8].

5.3 The Degree of User Coordination

An important concern for tabletop interaction design is the degree of coordination policy enforced on the collaborators. Having more than one user on the shared tabletop means that access to the interface resources (documents, widgets and other controls) can be potentially conflicting, for example two users wanting to view the same object at the same time. User A may want to verbally ask permission for the use of an object that is near User B. While frequent social "mishaps" have been observed in collaborative tasks [14], people tend to naturally partition their workspaces so that interference and conflicts will be minimised [23].

We consider a spectrum of usage coordination policy, from leaving the conflict resolution entirely to the users (thus relying on social protocols among collaborators, for example, User A asks User B "can I have a look at this document, please?") to enforcing a strict software-level coordination policy (for example, the system does not respond when User A attempts dragging an object very near User B to himself). Figure 5 illustrates this, with some techniques for coordination in the middle of the spectrum. The illustrated coordination techniques include: a user can use an object only when nobody else is selecting or touching it

[14]; automatically re-orienting objects to the nearby user's side, which implicitly discourages the other user from using that object [12]; provision of a personal area on the tabletop for each user where the other user cannot access and of a group area where users share access [16, 17]; more strict and explicit rules such as the lower ranked user cannot access the higher ranked user's objects, or a user explicitly granting an access to an object to the other user [14].

In the next section, we describe the interaction design of Físchlár-DT, with its design decisions informed by the above three concerns (task allocation, group awareness, and coordination policy). Some of these three decision factors have become the focus of our comparative experiments, and in the following section we describe one specific set of decisions we made, our *baseline* interface.

6 Físchlár-DT: Interaction Design and System Description

In designing Físchlár-DT, we based our decisions on the three considerations described in Section 5, within the support and constraints of the DiamondSpin software toolkit used for implementation. Here we briefly describe the Físchlár-DT interface, then we describe the design decisions made by considering the allocation the task, group awareness, and coordination policy, with which the interface was strengthened to enhance its collaborative aspects.

6.1 Overall Interaction of Físchlár-DT

The purpose of Físchlár-DT is to allow its collaborators to quickly search for video shots, and the system is currently being used in user experiments for interactive search as part of TRECVid 2005. Físchlár-DT supports two users sitting opposite each other around the DiamondTouch table. Initially a query panel is presented on the empty table. We decided to provide only one query box per team as we wanted to see how groups would react to this constraint, for example would one user become the dominant searcher. To type in query terms, the user uses a virtual keyboard (the keyboard itself can be also dragged and rotated) to type in query terms. Once the 'Enter' key is pressed, the keyboard disappears and the user taps on the 'Search' button. Search terms are then matched against the Automatic Speech Recognised (ASR) text of the video shot database and the top 20 matched shots are retrieved. Each of the retrieved shots is represented by a keyframe from the shot. As a result of retrieval, the table is populated with 20 miniaturised keyframes for the users to browse (see Figure 6). Users can select keyframes by touching any part of the image and once selected images can be dragged, resized and rotated. Keyframes can be dragged by selecting an image and moving across the tabletop as desired. Keyframes can

Figure 6. Table with top 20 matched keyframes, with contextual menu open

Figure 7. Middle of a search with saved keyframes pilled on the bottom-right corner

be resized by selecting and dragging the red icons on the corner of the image. Rotation of keyframes is achieved by selecting the purple icon on the left corner and dragging in the direction of the required rotation. Each keyframe can be double-tapped to bring up a contextual menu (see Figure 6) on which four possible actions are displayed: Save, Remove, Find Similar, and Play. To examine the content of a video shot, the user can tap on 'Play', which will play that shot on a separate, dedicated monitor beside the table. If the user decides that the shot is relevant, she brings up the contextual menu and selects 'Save', which will stamp a bright yellow 'SAVED' label on the keyframe. If the user decides a keyframe is not relevant, she can select 'Remove' on the menu, which will remove the keyframe from the table and for the remainder of the current search.

An important feature of the system is 'Find Similar': once selected, the system searches the database for other keyframes that are similar to that keyframe (based on content similarity using the four MPEG-7 features outlined in Section 3), and the table will be populated with the top 20 most similar keyframes. Since selecting 'Find Similar' means 20 more keyframes are added to the table, users are encouraged to clear up the table before this action. Thus, the two collaborators work on the 20 keyframes to play, remove and save, and when ready, retrieve another 20 keyframes to play, remove and save, and so on. Non-relevant keyframes are removed while relevant keyframes remain on the table with the yellow 'SAVED' label on each. Figure 7 shows the table after 10 minutes of the search task, on which the collaborators decided to pile the saved keyframes on one corner (bottom right) to free up their workspace for browsing.

6.2 Design Considerations for Collaboration

Now we describe the design decisions we made for each of the three collaborative interaction considerations from the previous section, applied to the baseline system. One set of design decisions on the collaborative aspects made up our baseline interface, and other decisions will make different interfaces which will be used to conduct a comparative experiment (see Section 7 for details).

6.2.1 The Task Allocation for Each Collaborator

For our baseline system, we decided to provide a workspace in close collaboration, with two users working on the same task at the same time at all stages of search (Task division 2 in Figure 3). A clear division of stages and sequential allocation of sub-task to each user is more suitable when the boundaries between the stages are clear-cut. In modern interactive IR systems, the user's searching process tends to be highly flexible with rigorous retrieval result visualisation, provision of relevance feedback and instant full-document viewing at any time during the search. Querying and browsing are mixed together, with most features visible and executable at any stage, allowing the user to refine her information need by browsing during search process. In terms of allocation of task, we therefore allowed the collaborators to do everything at the same time, without any rigid stage-by-stage division. Only one set of available resources were used on the table to enforce collaboration between the two users, thus requiring them to work together rather than in parallel. This meant that the other two concerns (group awareness and coordination policy) were to be more carefully considered and decided.

Consider two users working closely together on the conventional desktop interface of Físchlár-TRECVid2004 (Section 3) designed for a single user. We expect to see users competing for control (mouse and keyboard), one dominant user, with the passive user having a lack of attention and suffering frustration, the typical phenomenon when two users are working on a single-user desktop as observed in [22] [25]. Two users working closely together on Físchlár-DT, on the other hand, are expected not to suffer from these phenomena even though only one set of interface resources are available to them as the tabletop interface allows both users to work with the system simultaneously.

6.2.2 The Degree of Group Awareness

One of the DiamondSpin toolkit's major provisions is its draggable objects on the table: any keyframe on the table can be dragged and moved around with a finger, allowing more natural and physical manipulation of objects that helps enhance and maintain group awareness among the collaborators. Saved keyframes are dragged by a user to the side of the table so that the next set of retrieved keyframes will be populated on the central area of the table without overlapping the existing keyframes. In addition, we assigned distinctive sound effects for each of the major user actions and system responses, to enhance group awareness, including:

- Save a relevant keyframe - a distinctive stamping sound is heard immediately after selecting the 'Save' menu item, and the bright yellow 'SAVED' mark appears on the keyframe.

- Remove a non-relevant keyframe - a squashing sound is heard and the keyframe immediately disappears.

- Play a shot - when a user selects 'Play' in the contextual menu on the keyframe, a low 'click' sound is heard before the shot starts playing.

- Find similar keyframes - when a user selects 'Find Similar' a 'click' sound is heard before the set of matched keyframes appear on the table. When keyframes appear there is a distinctive sound (below).

- Displaying search result - finally, when a new set of keyframes are populated on the table, a distinctive sound is heard indicating the appearance of new information on the table.

6.2.3 The Degree of Coordination Policy

For the baseline system, we relaxed the software coordination rule and thus no explicit policy is enforced in accessing the table areas, widgets and keyframes: the whole table space and all elements on the table have equal access rights for both users. This encourages users to rely on social conventions and dialogue to manage conflicts or interference with each other. DiamondSpin by default re-orientates documents on the table to the user's side, thus providing an implicit personal area on each user's side though the other user is allowed to drag a keyframe from the other user's side.

These design considerations described here help us to see the context of our decisions regarding collaborative aspects and to understand possible consequences of the design decisions we made in terms of its collaborative features.

7 Experiments

To test out the ideas outlined in the earlier parts of this paper we are carrying out a series of interactive user experiments using Físchlár-DT as part of the TRECVid 2005 evaluation. The TRECVid schedule, which all participants adhere to, means that topics will be distributed to all participating groups in August with results of on-site searches due back to NIST for manual assessment in September and performance results will be available in November. We will use 24 multimedia topics (each topic consists of text, images and perhaps a video clip to describe the information need) and we will give each pair of users 10 minutes to start and complete the task of locating as many shots among the 60 hours of broadcast TV news content index as they can find within the time limit. Because we are using similar underlying image and text search engines as we used in TRECVid 2004, where we used a conventional desktop interface, the relative performances of the two approaches to video search can be compared, though not directly since the dataset for 2004 is different to that being used in 2005.

However, testing a Físchlár-DT vs. a conventional interface does not really explore the issues of user collaboration we discussed earlier. For example, there are different techniques that implement the degree of group awareness as we discussed in Section 5. We will use 2 variations of the baseline Físchlár-DT, one with emphasis on increased awareness amongst collaborators through interaction driven mostly by dragging and strong sound effects, and another with less awareness but increased efficiency and power for individual users through stronger symbolic actions such as contextual menu-driven interaction. With these variations we can then compare the overall effectiveness of the techniques by looking at how they affect performance and we can do this because we will be running all 24 topic searches each for 16 different users and that provides us with the statistical base to make such evaluations. Also, as part of our interactive evaluation we are also video recording all subject searches and using a video annotation tool to annotate the recordings of the search with a record of the human-human interactions, which we can then correlate with retrieval performance.

8 Conclusions

In this paper we have introduced the Físchlár-DT system, which has a DiamondTouch interface to an underlying video search system. We presented the Físchlár system with its conventional interface as we have developed in previous work, as well as the version we have developed specifically for use on the DiamondTouch. In describing this we outlined the typical user interaction and how Físchlár-DT supports a collaborative search process between two users. The main part of the paper, however, concentrated on how we can incorporate aspects of collaborative interaction design into a tabletop interface for video searching and the main contribution was a description of how Físchlár-DT addresses issues of designing a user interaction to support collaborative search for video.

To test the design decisions we incorporated into Físchlár-DT we outlined our experiments which we are carrying out as part of the annual TRECVid benchmarking activity, during which we will evaluate some of these design features. Subsequently, and using an annotated video recording of the search processes we will be able to examine the correlation between retrieval performance of the users' searches and their actual interaction experiences.

At the time of writing the system as described has been developed and tested, and we are carrying out these user experiments. We will have performance results in October.

Acknowledgements

References

[1] M. Ashdown and P. Robinson. Escritoire: a personal projected display. *IEEE Multimedia*, 12(1):34–42, March 2005.

[2] A. Blackwell, M. Stringer, E. Toye, and J. Rode. Tangible interface for collaborative information retrieval. In *ACM CHI2004*, pages 1473–1476, 2004.

[3] P. Browne, C. Czirjek, G. Gaughan, C. Gurrin, G. Jones, H. Lee, S. Marlow, K. McDonald, N. Murphy, N. O'Connor, N. O'Hare, A. Smeaton, and J. Ye. DCU video track experiments for TREC'03. In *TRECVid2003*, 2003.

[4] E. Cooke, P. Ferguson, G. Gaughan, C. Gurrin, G. Jones, H. LeBorgne, H. Lee, S. Marlow, K. McDonald, M. McHugh, N. Murphy, N. O'Connor, N. O'Hare, S. Rothwell, A. Smeaton, and P. Wilkins. TRECVid 2004 Experiments in Dublin City University. In *TRECVid2004*, 2004.

[5] P. Dietz and D. Leigh. DiamondTouch: a multi-user touch technology. In *ACM UIST'01*, pages 219–226, 2001.

[6] G. Gaughan, A. Smeaton, C. Gurrin, H. Lee, and K. McDonald. Design, implementation & testing of an interactive video retrieval system. In *ACM MIR'03*, pages 23–30, 2003.

[7] C. Gurrin, P. Browne, A. Smeaton, H. Lee, K. McDonald, and P. MacNeela. Físchlár-Nursing, using digital video libraries to teach processes to nursing students. In *IASTED'04*, pages 111–116, 2004.

[8] C. Gutwin and S. Greenberg. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *ACM CSCW'98*, pages 207–216, 1998.

[9] C. Gutwin and R. Penner. Improving interpretation of remote gestures with telepointer traces. In *ACM CSCW'02*, pages 49–57, 2002.

[10] P. Hansen and K. Jarvelin. Collaborative information retrieval in an information-intensive domain. *IP&M*, 41(5):1101–1119, September 2005.

[11] M. Hearst. User interfaces and visualization. In *Modern information retrieval, Chapt. 10*, pages 257–323, 1999.

[12] R. Kruger, S. Carpendale, S. Scott, and S. Greenberg. How people use orientation on tables: comprehension, coordination and communication. In *ACM Group'03*, pages 369–378, 2003.

[13] H. Lee and A. Smeaton. Designing the user interface for the Físchlár digital video library. *JoDI*, 2(4), May 2002.

[14] M. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond "social protocols": multi-user coordination policies for co-located groupware. In *ACM CSCW'04*, pages 262–265, 2004.

[15] S. Scott, K. Grant, and R. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW'03*, pages 159–178, 2003.

[16] S. Scott, M. Sheelagh, M. Carpendale, and K. Inkpen. Territoriality in collaborative tabletop workspaces. In *ACM CSCW'04*, pages 294–303, 2004.

[17] C. Shen, K. Everitt, and K. Ryall. UbiTable: impromptu face-to-face collaboration on horizontal interactive surfaces. In *UbiComp 2003*, pages 281–288, 2003.

[18] C. Shen, N. Lesh, and F. Vernier. Personal Digital Historian: Story sharing around the table. *interactions*, 10(2):15–22, March 2003.

[19] C. Shen, F. Vernier, C. Forlines, and M. Ringel. DiamondSpin: an extensible toolkit for around-the-table interaction. In *ACM CHI'04*, pages 167–174, 2004.

[20] A. Smeaton, C. Gurrin, H. Lee, K. McDonald, N. Murphy, N. O'Connor, D. O'Sullivan, B. S. B, and D. Wilson. The Físchlár-News-Stories system: personalised access to an archive of tv news. In *RIAO 2004*, pages 3–17, 2004.

[21] O. Ståhl, A. Wallberg, J. Söderberg, J. Humble, L. E. Fahlén, A. Bullock, and J. Lundberg. Information exploring using the pond. In *CVE'02*, pages 72–79, 2002.

[22] J. Stewart, E. Raybourn, B. Bederson, and A. Druin. When two hands are better than one: enhancing collaboration using single display groupware. In *ACM CHI'98*, pages 287–288, 1998.

[23] E. Tse, J. Histon, S. Scott, and S. Greenberg. Avoiding interference: how people use spatial separation and partitioning in SDG workspaces. In *ACM CSCW'04*, pages 252–261, 2004.

[24] M. Twidale and D. Nichols. Designing interfaces to support collaboration in information retrieval. *Interacting with Computers*, 10(2):177–193, May 1998.

[25] M. Twidale, D. Nichols, and C. Paice. Browsing is a collaborative process. *IP&M*, 33(6):761–783, 1997.