CEVA: A Tool for Collaborative Video Analysis

Andy Cockburn and **Tony Dale**
Department of Computer Science
University of Canterbury
Christchurch, New Zealand
{andy, tony}@cosc.canterbury.ac.nz

ABSTRACT

Video protocol analysis is a standard technique in many research disciplines including human-computer interaction and computer supported cooperative work. It is notoriously time consuming, and a variety of single-user computer based tools have been developed to ease the task.

This paper examines collaborative video analysis. The motivation for groupware tools for video analysis is described, and the desirable features of such tools are identified. The design, implementation, and preliminary evaluation of a prototype synchronous groupware tool for video analysis, CEVA, are described.

Keywords

Collaborative video analysis, evaluation, groupware, user interfaces, design.

INTRODUCTION

Video protocol analysis is a popular and effective technique for understanding the interactions between people, and between people and machines. Analysing and transcribing video recordings is, however, notoriously time consuming. Neal [14], for instance, cites the "extreme tediousness of transcribing the data from the videotape into annotated transcripts," and Allen [1] reports that transcription "can take from two to ten hours per hour of tape at a minimum."

The aim of computer-supported tools for video analysis is to ease the analyst's task and to reduce the high ratio of analysis-time to observational session time. In a recent review, Sanderson, Scott, Johnston, Mainzer, Watanade, and James [20] refer to more than forty video analysis tools. Some of these tools partially support *asynchronous* collaboration between multiple analysts (by allowing the system state to be saved and reloaded at a later time), but none of them support simultaneous collaboration.

In contrast to the single-user design of video analysis tools, video analysis is often a collaborative activity. It is common, for instance, for research colleagues involved in HCI to collaboratively study a video, mutually identifying salient usability issues. Lewis, Mateas, Palmiter and Lynch [10] acknowledge the use of collaborative video analysis in their description of an ethnographic design process:

> "... we devised a coding scheme to help us parse the raw video tape. Three main roles, with specific coding duties for each, were defined for each of three observers. One observer recorded overall observations and impressions, the second interactions and the third events."

Similarly, in the commercial sector where usability specialists are becoming increasingly common, an interface development team may use a video recording of user errors to convince management that further development is necessary. Naturally, the users may also be invited to view and contextualise the video observations. At Bell Northern Research Labs, Kennedy [9] emphasises the role of *collaborative* video analysis, stating that video is "an essential communication tool."

This paper describes the design, implementation, and preliminary evaluation of CEVA, a prototype synchronous CollaborativE Video Analysis tool. The following section reviews related work on video analysis and on single-user video analysis tools. The review is used to derive a set of design goals for *collaborative* video analysis tools. We then detail CEVA's user interface, its collaborative facilities, and its implementation. CEVA's preliminary evaluation is described prior the conclusions which identify directions for further work.

VIDEO ANALYSIS: TECHNIQUES AND TOOLS

Video observational analysis is common across many research (and commercial) disciplines including sociology, psychology, ethnography, human-computer interaction, and computer supported cooperative work. The name given to video analysis techniques differs between disciplines, but the term "Exploratory Sequential Data Analysis," or ESDA [20], encapsulates the entire range of methods. In general,

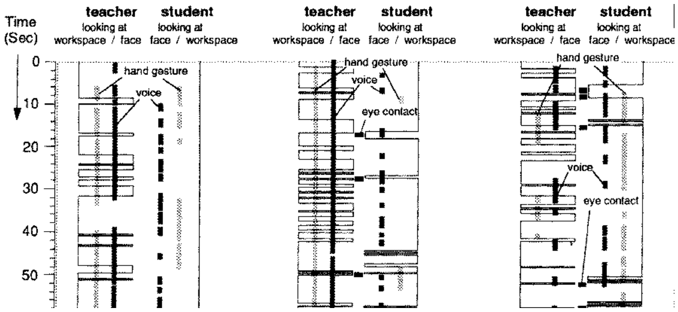

Figure 1: Analysis of Clearboard prototypes, extracted from Ishii, Kobayashi and Grudin (1992).

ESDA techniques preserve the sequential integrity of observed events, and involve an element of empirical data analysis. ESDA methods do not necessarily use video, but video is by far the most common medium for recording the observational session.

Beyond the assumed use of video, ESDA analytical techniques vary widely between research disciplines and between projects. The spectrum of ESDA analytical goals ranges from ethnographical immersion in the observational session to quantitative measures and event counts. ESDA methods and techniques are reviewed in Sanderson *et al.* [20] and in Neal [14].

ESDA methods often involve analysing concurrent *threads of activity* during an observational session. Trigg [26] used the term "data-streams" to describe "multiple chronologically organised representations of a single sequence of activity [which] can be edited, juxtaposed, and aligned." In CSCW research, threads of activity can be readily applied to analysis of collaborative activity. For example, Ishii, Kobayashi and Grudin [8] and Takeuchi and Naito [22] use data-streams to analyse styles of collaboration. Figure 1, extracted from Ishii *et al.* [8] shows an aligned representation of the concurrent activities of two people using the experimental system Clearboard. The threads of activity eloquently describe a combination of gaze-awareness (direction of sight), vocal comments, and gestural activities over a two minute period in three variants of the Clearboard system.

Video Analysis Tools

Reducing the high ratio of analysis-time to observational-session time has been a primary motivation for the development of computer supported tools for video analysis. Computers, however, offer many benefits to ESDA beyond reduced time and effort. These include configurable representations of the logged events, semi-automated search and query facilities, automated quantitative and statistical measures, powerful visualisation facilities, and linking between related events. Additionally, the emergence of synchronous groupware allows simultaneous analysis by multiple users. The potential benefits of simultaneous analysis will be discussed in the following section.

Most single user tools for video analysis offer similar functionality. Our design criteria for collaborative video analysis tools, described in the following section, has drawn extensively on previous work. These sources include the guidelines and design rationale for video analysis tools presented by Mackay [11] for her system EVA, by Trigg [26] with his Workplace project, by Harrison [7] for her system VANNA, and by Weber and Poon [27] with their system Marquee. Other important video analysis systems include CVideo [18], MacSHAPA [20], and VRSS [3]. For an extensive review of ESDA tools, see Sanderson *et al.* [20] and Posner & Baecker [17].

DESIGN CRITERIA FOR COLLABORATIVE VIDEO ANALYSIS

Ideally, computer tools for ESDA would be sufficiently generic to support the full range of multi-disciplinary tasks that analysts require. Sanderson *et al.* [20] note that this ideal is essentially unobtainable because of the designers' prior assumptions of analytical tasks. This section presents the design rationale that motivates and directs our work on collaborative video analysis. These design goals serve as a model for idealised video analysis tools. In particular, we focus on the rationale behind *collaborative* video analysis.

1. Collaborative (synchronous and asynchronous) video analysis

Mackay [11], among others, noted that video analysis is often a collaborative activity. Her design criteria for EVA stressed the importance of supporting the various views and interests of different analysts. To our knowledge, CEVA is the only video analysis tool to explicitly support collaborative activity.

Synchronous collaboration offers many possibilities to video analysts. In particular, two potential benefits that we wish to explore are synergy in collaborative analysis and parallel analysis, discussed below.

Synergistic Analysis

Collaborative video analysis offers a synergy of differing views and perspectives. The richness and fluidity of collaboration between video analysts' is vividly illustrated by Tang's videos, associated with his work reported in [23, 24]. The videos show a group of HCI professionals analysing a video of groups working together over a shared work surface. Although Tang's observations focus on the activities of those manipulating the work-surface, it is fascinating to study the collaboration between the analysts.

Without tools to directly support the analysts' collaboration there is a risk that these benefits of group participation will be lost after the session. A potential solution would be to have one person document all users' observations, but this reintroduces the original problem: that of one person transcribing multiple threads of activity.

Tools for video analysis should support and maintain the shared understanding that is dynamically constructed by collaborating analysts.

Parallel Analysis

Parallel activity by multiple analysts is likely to speed up video analysis[1]. In addition, multiple analysts are likely to observe a greater coverage of salient events: for instance, Nielsen [15] reports on improvements in usability analysis brought about by using multiple evaluators.

Video analysis often involves transcribing events on several concurrent threads of activity (see point 2 below). Simultaneous parallel collaboration allows one (or more) thread to be allocated to each user for concurrent analysis. Lewis *et al.* [10] describe precisely this style of parallel separation of analysis tasks within their ethnographic methodology for system design. The medium that they use for coding, however, is paper cards, and they state "it would be useful to have a computerised tool to assist in coding."

The methods and means of exploiting CEVA's collaborative facilities are briefly discussed later in the paper.

2. **Synchronised, multi-threaded event logging**

Video analysis tools should allow information from various channels of activity to be recorded in a variety of representations. Systems should assist visualisation of simultaneous events to ease comparisons between threads of activity.

3. **Animated, direct manipulation, interface**

Video is a dynamic medium. A dynamic and animated interface can reinforce the correlation between the dynamically changing events on the video and the synchronised events logged by the user [26, 17].

Additionally, we believe that a highly visual direct manipulation interface is necessary to reinforce the system's role as a "conversational prop" [2].

4. **Symbolic annotation and visualisation of data at various levels of granularity**

Video analysis requires many levels of detail when recording events and when browsing the logged transcript. Users should be able to tailor analysis tools to their needs: from detailed transcription at a high level of detail, to abstract representation of a large section of the transcript at a low-level of detail. In addition, symbolic representations are likely to assist visual browsing of the transcript.

Trigg (1989) proposed that further work on the Workplace system should include "zooming" of the transcript.

[1] We suspect that the total-person hours will remain the same or increase, but this is yet to be studied.

5. **Minimal intrusion on video playback**

Logging mechanisms should not intrude on the playback of the video. Each stop-start of the video interrupts the event logging process and, in group use, may disrupt the activity of colleagues.

Non-intrusive transcription mechanisms can also be used in real-time to analyse the original interaction while it is being recorded. Users can observe and log events "on the fly" rather than through the video. Mackay [11] notes that "annotation during the session can save time later when the researcher is ready to review and analyse the session." Weber and Poon [27] note that real-time annotation provides a skeletal structure for subsequent analysis.

6. **Quantitative analysis**

The retrieval of quantitative measures is often the primary reason for carrying out ESDA. In HCI, for example, quantitative values such as events counts and the total time spent in an event type can indicate the severity of usability problems. The quantitative measures recorded by the system should be directly transferable to standard packages for statistical tests, or preferably a statistical package should be integrated within the system.

7. **Event searching mechanisms**

Video analysis tools should support fast retrieval of the video sections associated with the logged transcript. Effective video retrieval mechanisms, such as those of VRSS [3], are desirable.

8. **Reordering of video segments**

By allowing sections of the tape to be reordered, users can make condensed collations of multiple occurrences of similar event patterns [11].

9. **Consumer level (inexpensive) hardware**

Some of the existing single-user tools for video analysis require moderately high set-up costs which restrict their use in low-budget institutions. We want CEVA to run on consumer-level VHS video equipment and on low-end computer hardware. Furthermore, CEVA's collaborative facilities should not prohibit its use as a stand-alone single-user system when multiple machines (or multiple analysts) are unavailable.

CEVA, discussed below, supports all of the points above except reordering of events (point 8), which is planned for future work.

CEVA, A COLLABORATIVE VIDEO ANALYSIS TOOL

CEVA is a prototype groupware tool that supports co-located simultaneous video analysis. It provides equivalent functionality to many single-user video analysis tools, but in addition it supports WYSIWIS[2] group video analysis. It is designed

[2] What you see is what I see [21].

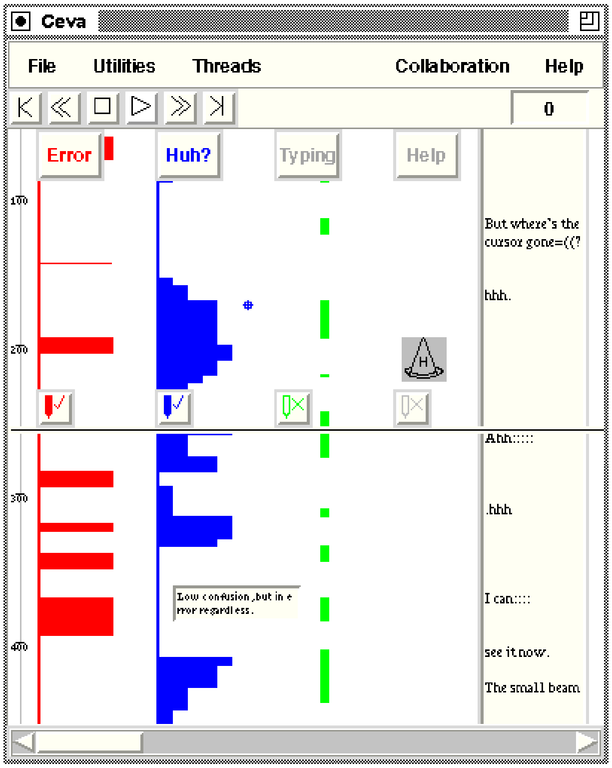

Figure 2: A collaborative analysis session with CEVA.

to work with consumer level VCR equipment and with base-level computer hardware.

In figure 2 CEVA is being used to analyse five threads of activity during a usability study of a help sub-system. The 'Error' thread represents the occurrence of user errors. The 'Huh?' thread shows an assessment of the user's level of confusion. The 'Typing' thread shows keyboard activity, and the 'Help' thread represents direct interaction with the help subsystem. The right hand side of the figure shows a text-transcription thread. CEVA's threads of activity are directly equivalent to Trigg's 'data-streams' [26].

CEVA's interface components and its groupware facilities are introduced below.

CEVA's User Interface

CEVA provides an animated direct manipulation interface that is synchronised with the video playback. Events can be logged on customisable threads of activity which may be graphical or textual, and on text annotations. The logged transcript of events can be saved to a file for later use, and a hardcopy of the transcript can be printed out.

The video control buttons at the top of the window allow control of a consumer-level VHS video cassette recorder (or several recorders) with accuracy of ± 1 frame. The tape is started by clicking the play icon or by a key-binding which allows the users to position the mouse over the appropriate logging button prior to starting the tape. As the video plays (on a separate screen) the user-specified event-threads scroll up towards the top of the window, dynamically revealing the events recorded by the users.

The 'frame-line' across the middle of the window correlates with the current video frame. Events on graphical threads are recorded by pressing the corresponding log-button, holding it for the duration of the event. In figure 2 the log-buttons associated with each of the four graphical threads are labelled 'Error', 'Huh?', 'Typing', and 'Help'. Text is directed at a specific text-thread or text annotation (there may be many) by clicking in the desired text widget prior to typing. The text-entry cursor moves to the line corresponding to the frame-line whenever one of the users (whose keyboard focus is at the text-widget) presses the return key. The intention is to ease the maintenance of the temporal relationship between the textual transcription and the current video frame.

The pencil icons on the frame line allow the users to control whether each graphical event thread is being logged. Clicking on a pencil toggles between a logging state and a non-logging state. In figure 2 the 'Error' and 'Huh?' threads are logging (shown by a filled pencil and a tick mark), and the 'Typing' and 'Help' threads are not (shown by an unfilled pencil and a cross mark). Non-logging states allow multiple passes through the tape without deleting previously logged information.

The log of events on any graphical thread can be changed by rewinding to the appropriate point, clicking the thread's pencil to turn logging on, restarting the tape, and logging over the old thread record. Textual transcriptions can be edited with normal text-editing procedures.

Types of Event Threads

Any number of text-threads and graphical threads can be added to CEVA's log. CEVA supports four types of graphical event threads:

1. Two-state line threads—useful for recording frequently occurring events that are either 'on' or 'off' (for example, the 'Error' thread in figure 2).

2. Multi-state line threads—useful for events that may be categorised at more than two-levels. Customisable key bindings allow users to record events at several discrete levels. The log-button records a maximum-level event. The 'Huh?' thread in figure 2 is a multi-state line thread.

3. Block threads—an alternative representation for two-state events (for example, the 'Typing' thread in figure 2).

4. Iconic threads—normally used to mark rarely occurring events. In figure 2 an iconic thread is used to represent access to the help system.

Ishii *et al.*'s [8] analysis in figure 1 uses a combination of two-state line threads and block threads.

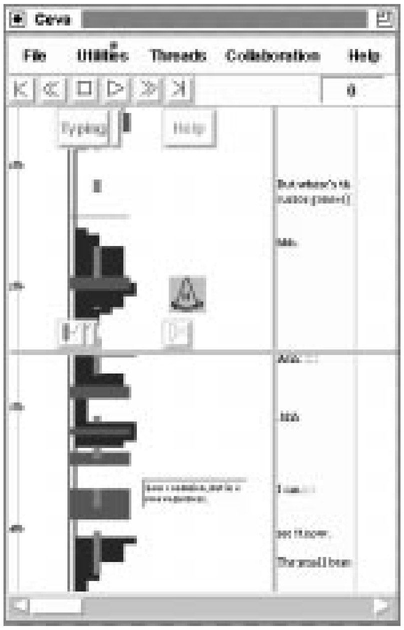

Figure 3: Overlaying activity threads to scrutinise concurrent events.

Free Text Annotation

Not all events can be easily categorised into graphical and textual threads of activity. Noteworthy events may occur only once on the video. To account for such events, as well as a variety of other uses, text annotations can be attached to the transcript at any point (shown towards the bottom of figure 2).

Visualisation and Analysis of the Log

CEVA supports several powerful visualisation techniques to support browsing and searching of the logged data. These include direct manipulation of the event threads, zooming magnification of the graphical threads, and keyword searches.

Threads can be repositioned by direct manipulation to assist visualisation of concurrent events. Figure 3 shows the same transcript as figure 2, but with the threads overlaid. The users can also zoom into or away from the transcript. Zooming out supports searching and browsing by showing progressively larger portions of the transcript at a lower level of detail. Zooming in magnifies a region of the transcript supporting detailed analysis.

The text threads and text annotations can be searched by keyword. When text matches are found the users can position the tape directly to the associated video frame.

CEVA's support for quantitative analysis is, as yet, rudimentary. It provides an automatic count of the number of events at each level on each event thread. It also calculates the time spent in each state for each thread. The data can be saved to a file for transfer to statistical packages. Figure 4

Event Counts

Thread	Error	Huh?	Typing	Help
Number of States	2	6	2	2
State0 Events		24		
State1 Events		16		
State2 Events		13		
State3 Events		16		
State4 Events		14		
State5 Events		36		
Total Events	59	119	57	3
Off Time	2765	2300	2695	280
State1 Time	545	175	615	6
State2 Time		120		
State3 Time		180		
State4 Time		170		
State5 Time		290		
Total Event Time	3310	3235	3310	286
Logged Time	3325	3325	3325	325

Save Print OK

Figure 4: Automated quantitative values.

shows the quantitative values produced for the system state shown in figure 2.

Collaborative Features

CEVA allows groups of people to simultaneously analyse single or multiple video sessions. There is no built-in limit on the number of users, but a realistic maximum group size is five, with system response speed being the limiting factor. Although we assume that CEVA users are co-located this is not a requirement of the system. Multiple analysts, each with their own VCR, could be distributed around a local-area network without adversely affecting system response.

The groupware user interface to CEVA is almost identical to the single-user interface. The only additional interface elements are telepointers which show the location of each user's cursor. Communicating each user's location of activity is a fundamental requirement of WYSIWIS groupware [25, 6]. Telepointers in all parts of CEVA support gestures and deixis by the multiple users.

CEVA's WYSIWIS is slightly relaxed in two ways. First, the thread customisation dialogue (for adding and modifying threads) is only shown to the user that selects it. The rationale is that the customisation and configuration of the system should be determined by focused collaborative negotiation. As CEVA is primarily designed for co-located use, we believe that a single-channel for customisation will encourage face-to-face discussion. Second, the scroll-bar at the bottom of the window allows users to view different sets of event-threads. The scroll-bar is only necessary in analyses involving more than six threads.

CEVA relies on social protocols to determine the users' roles and responsibilities. Rather than enforcing a predetermined floor-control mechanism, any user can make any interface action at any time. Research suggests that pre-assigned roles become a straight-jacket to the users' dynamically adapting requirements, and that people prefer to fluidly adjust their responsibilities within groupware environments without the imposition of formal computer-supported mechanisms [24, 25]. Naturally, allowing free access to the interface controls raises the possibility of conflicting actions by different users, but the telepointers are intended to promote awareness of others' activities before conflicting actions occur. CEVA leaves the high-level decisions on resolution of intended actions to the users.

We envisage many alternative usage scenarios in which CEVA can reinforce collaborative video analysis. For instance, CEVA's video hardware can be configured so that not all analysts observe the same video-tape. Different, but synchronised, tapes could be assigned to each user, allowing CEVA to capture a synchronised log of multiple viewing angles of the same observational session (a screen shot and a keyboard shot in user-interface evaluations, for example). Alternatively, CEVA could be used as a synchronisation point for multiple analysts who work independently. Analysts could log portions of a tape independently using CEVA as a single user tool, then combine their logs for face-to-face discussion of the multiple analytical perspectives. This style of analysis is reported in [10].

We are still in the development stage, and we are not, as yet, advocating particular usage styles for CEVA. Rather, we are attempting to develop a useful tool that is free from predetermined and restrictive collaborative styles.

Tape Control Versus Digital Video

CEVA could allow much wider ranging facilities for video analysis if we were to digitise the video tape. For instance, it would allow a variety of forms of semi-synchronous collaboration in which multiple analysts view differing portions of the tape at the same time. Although this is planned for our further work, it has had a low-priority in our research to date. Prior to supporting this less-focused form of collaboration, we wish to investigate the collaborative issues that are brought about by tightly focused (strict-WYSIWIS) simultaneous analysis of a common artifact (the video).

Implementation Details

CEVA is written in Tcl/Tk [16] using the GroupKit toolkit [19]. Computer control of the VCR is via the Sony Control-L protocol. This protocol is implemented on an MIT Miniboard [13], which communicates with the controlling computer via an RS232 link and a simple command language [4]. The computer sends video-motion commands to the Miniboard, which encodes the command in a Control-L packet and sends it the the VCR Control-L port. The VCR constantly sends status information, which is monitored by the Miniboard and made available to the host computer on demand. Control-L allows *almost* frame-level accuracy at significantly lower cost than high-end editing VCRs.

PRELIMINARY EVALUATION

The degree to which CEVA satisfies the design criteria is summarised in Table 1.

To obtain early feedback on CEVA's user interface and on the viability of its groupware facilities, four first-year postgraduate students in HCI were invited to take part in two, one hour, Wizard of Oz [5] usability studies[3]. This preliminary evaluation was carried out to fuel our iterative design of CEVA.

The subjects were asked analyse a video of collaborative drawing activity. Their analysis was to record six two-state threads of activity: whether each person was looking at the other; whether each person was speaking; and whether each person was drawing. They were given a five minute introduction to CEVA's facilities, after which no further assistance was provided. This preliminary evaluation did not examine the use of multi-state threads, primarily because categorising events into multiple levels requires analytical expertise that the subjects were unlikely to possess. CEVA's text-threads were also not included in the study: we were most concerned with CEVA's collaborative support, and we feel that transcribing the audio channel of video will, most likely, be a single-user process. Ten minutes were allocated for a post-task interview.

At the time of the experiment, CEVA's groupware interface was complete, but the machine interface between CEVA and the VCR was not. A Wizard of Oz approach was therefore used to simulate the VCR's response to user actions. Whenever CEVA's video control buttons were used, the corresponding video response was stimulated by the experimenter using the VCR remote control. Although the Wizard of Oz simulation affected portions of the evaluation, the majority of the evaluation was unaffected.

Observations

The evaluation was highly encouraging. Both pairs of users successfully produced a log of the events on all six threads within forty minutes. The resultant logs included several text annotations. All of the subjects reported that they had learnt a lot about usability studies, and that they found video analysis interesting, revealing, and exhausting.

Several small problems in the user interface were detected (and are currently being rectified), and many collaborative issues emerged—some foreseen, some not.

The biggest user problem was due to the Wizard of Oz simulation of VCR control. Users were not confident in the correlation between CEVA's representation of the video frame and the VCR's actual position. They found precise control of the tape cumbersome because of the 'slow' response of the video control buttons. The precise control and rapid feedback provided by the Control-L interface to the VCR should eliminate these problems.

[3] Expert video analysts would obviously be a preferable subject-base, but they are in short supply in New Zealand. Budding usability professionals were the next best thing.

Design Criteria	Supported?	Comment
1. Collaborative analysis:		
synchronous	✔	Strict WYSIWIS of video frame.
asynchronous	✗	Only through save/load facilities.
2. Synchonised, multi-threaded event logging.	✔	Various forms of threads (graphical and textual) are synchronised with the tape.
3. Animated, direct-manipulation interface.	✔	
4. Symbolic annotation and visualisation.	✔	Graphical threads with customisable representation support visual browsing, and zooming tools support focused work.
5. Minimal intrusion on playback.	✗ ✔	Evaluations indicated that users were concerned about interrupting their partners.
6. Quantitative analysis.	✔	Automatic generation of quantitative measures.
7. Event searching.	✔	Keyword searches supported, and zooming for various levels of abstraction.
8. Reordering of video segments.	✗	To be developed with digital video.
9. Consumer-level hardware.	✔	Consumer-level VCR equipment through Sony Control-L protocol.

Table 1: Evaluating CEVA with respect to the design criteria.

User Interface Issues

The subjects found that they had difficulty with the mouse "falling off" the log-button. When concentrating on the video screen, all subjects found that small accidental movements in the mouse caused the cursor to move off the log-button with the result that a series of events was missed. This was particularly problematical when one subject attempted to simultaneously log two threads. He found that repositioning the cursor over each log-button was demanding: he was afraid that moving his eyes away from the video screen might result in missing events.

We are working on two solutions to this problem. The first is to allow customisable key presses to trigger events for each thread (key presses are already used for multi-state line threads). The second is to make each log-button button 'magnetic' so that small mouse movements are ignored, but large motions snap away from the button. Another potential solution is to reduce the visual context switch between CEVA and the video display by showing the video on the computer screen.

CEVA's zooming visualisation feature was particularly successful. It was used several times to locate regions of the transcript that the subjects wished to return to. Subjects reported that the ability to "reach into the log and grab a section of tape" was very useful.

Collaborative Issues

The collaboration styles of the two subject pairs were similar. In both cases CEVA was configured to reveal all six event threads prior to any analysis. Once event logging began there was minimal direct collaboration between the subjects. Periodically they would note that they had made an error, and a quick query would determine whether a tape-stop was required. All subjects reported that they were reluctant to request a tape stop because it would disrupt their colleague. Negotiation between the participants was always used prior to stopping the tape. The subjects also stated that event logging was highly mechanistic, but very busy.

Pair two reported that after stopping the tape they found it difficult to find the portion of the tape (and log) where they had made earlier errors. Pair one came across the same problem, but solved it during the session by creating an iconic thread labelled 'oops'. Whenever either user made an error they immediately clicked the 'oops' button, marking the error and avoiding the necessity of stopping the tape.

Although there was little collaboration while the tape ran, when the tape stopped the subjects immediately discussed the activities that they had analysed. We were relieved to see that the mechanistic and silent parallelism during logging had not preempted the opportunity for synergistic collaboration. Related work has reported that users of computer-supported meeting environments can *either* use the computer *or* partake in the discussion [12]. We had been concerned that the depth of concentration required to log video events would destroy the synergistic properties of collaborative video analysis.

CONCLUSIONS

In the behavioural and social sciences, including HCI and CSCW, video analysis is a common technique for attempting to understand the interactions between people, and between people and machines. Although video analysis is often carried out by groups of people working together, the support tools currently available make no provision for group use.

CEVA is a synchronous groupware tool that provides an animated direct manipulation WYSIWIS interface for multi-threaded collaborative video analysis. It supports a variety of powerful visualisation techniques and it allows direct access to the portions of the video associated with textual transcriptions and annotations. CEVA's implementation was guided by a set of design criteria for idealised groupware video analysis tools.

The results of our preliminary usability studies indicate that CEVA is capable of supporting collaborative video analysis. Our future work with CEVA will follow three directions. First we will continue to develop CEVA's functionality so that it forcefully exemplifies our design criteria for collaborative video analysis tools. Second, we will evaluate and iteratively improve CEVA's user interface. Third, we will continue to study the styles of collaborative analysis made possible by synchronous groupware for video analysis. It is our intention that, once released as public-domain software on the Internet, CEVA will support a wide range of HCI and CSCW research.

ACKNOWLEDGEMENTS

This research is carried out with the support of University of Canterbury research grant number 2313728.

References

[1] C Allen. The use of video in organizational studies. *ACM SIGCHI Bulletin: Special Edition on Video as a Research and Design Tool*, 21(2):115–117, 1989.

[2] T Brinck and LM Gomez. A collabortion medium for the support of conversational props. In *Proceedings of the 1992 ACM Conference on Computer Supported Cooperative Work*, October 31 to November 4.Toronto, Canada., pages 171–178, 1992.

[3] TS Chua and LQ Ruan. A video retrieval and sequencing system. *ACM Transactions on Information Systems*, 13(4):373–407, 1995.

[4] AJE Dale. A micro-controller interface for control-l devices. Technical Report TR-COSC 06/97, Department of Computer Science, University of Canterbury, Christchurch, New Zealand, 1997. `http://www.cosc.canterbury.ac.nz/˜tony/tr9706.html`

[5] JD Gould, J Conti, and T Hovanyecz. Composing letters with a simulated listening typewriter. *Communications of the ACM*, 26(4):295–308, 1983.

[6] C Gutwin, G Stark, and S Greenberg. Support for workspace awareness in educational groupware. In *ACM Conference on Computer Supported Cooperative Learning (CSCL '95). Bloomington, Indiana. October 17–20, 1995*, pages 147–156. Lawrence Erlbaum Associates, Inc, October 1995.

[7] BL Harrison. Video annotation and multimedia interfaces: From theory to practice. In *Proceedings of the Human Factors Society 35th Annual General Meeting*, pages 319–322, 1991.

[8] H Ishii, M Kobayashi, and J Grudin. Integration of inter-personal space and shared workspace: Clearboard design and experiments. In *Proceedings of the 1992 ACM Conference on Computer Supported Cooperative Work*, October 31 to November 4.Toronto, Canada., pages 33–42, 1992.

[9] S Kennedy. Using video in the bnr usability lab. *ACM SIGCHI Bulletin: Special Edition on Video as a Research and Design Tool*, 21(2):92–95, 1989.

[10] S Lewis, M Mateas, S Palmiter, and G Lynch. Ethnographic data for product development: A collaborative process. *Interactions*, 6(3):52–69, 1996.

[11] WE Mackay. Eva: An experimental video annotator for symbolic analysis of video data. *ACM SIGCHI Bulletin: Special Edition on Video as a Research and Design Tool*, 21(2):68–71, 1989.

[12] M Mantei. Groupware: Interface design for meetings (PANEL). In *Proceedings of CHI'88 Conference on Human Factors in Computing Systems*, pages 161–163, 1988.

[13] F Martin. Mini board technical reference. MIT Media lab, 20 Ames St, E15-301, Cambridge, MA 02139, USA, 1995. `http://lcs.www.media.mit.edu/people/fredm/papers/mb`

[14] L Neal. The use of video in empirical research. *ACM SIGCHI Bulletin: Special Edition on Video as a Research and Design Tool*, 21(2):100–101, 1989.

[15] J Nielsen. Finding usability problems through heuristic evaluation. In *Proceedings of CHI'92 Conference on Human Factors in Computing Systems* Monterey, May 3–7, pages 373–380. Addison-Wesley, 1992.

[16] JK Ousterhout. *An Introduction to Tcl and Tk*. Addison-Wesley, 1993.

[17] I Posner and RM Baecker. A survey of video analysis research and tools and the timelines program: Exploring the state of the art in exploratory sequential data analysis, 1995. Available from Ron Baecker: Dynamic Graphics Project, University of Toronto, Canada.

[18] J Roschelle. Jeremy roschelle's home page. WWW page: http://www.slip.net/˜jeremy/, 1996. For information on CVideo, write to: 4104 24th Street #344, San Francisco CA 94114. USA.

[19] M Roseman and S Greenberg. Building real time groupware with GroupKit, a groupware toolkit. *ACM Transactions on Computer-Human Interaction*, 3(1):66–106, 1996.

[20] P Sanderson, J Scott, T Johnston, J Mainzer, L Watanade, and J James. Macshapa and the enterprise of exploratory sequential data analysis (esda). *International Journal of Human-Computer Studies*, 41(5):633–681, 1994.

[21] M Stefik, G Foster, K Kahn, DG Bobrow, S Lanning, and L Suchman. Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. In I Greif, editor, *Computer Supported Cooperative Work: A Book of Readings*, pages 334–366. Morgan Kaufmann, 1988.

[22] A Takeuchi and T Naito. Situated facial displays: Toward social interaction. In *Proceedings of CHI'95 Conference on Human Factors in Computing Systems* Denver, May 7–11, pages 450–455, 1995.

[23] JC Tang. *Listing, Drawing, and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams*. PhD thesis, Stanford University, 1989.

[24] JC Tang. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34:143–160, 1991.

[25] DG Tatar, G Foster, and DG Bobrow. Design for conversation: Lessons from cognoter. *International Journal of Man-Machine Studies*, 34(2):185–209, 1991.

[26] RH Trigg. Computer support for transcribing recorded activity. *ACM SIGCHI Bulletin: Special Edition on Video as a Research and Design Tool*, 21(2):72–74, 1989.

[27] K Weber and A Poon. Marquee: A tool for real-time video logging. In *Proceedings of CHI'94 Conference on Human Factors in Computing Systems* Boston, April 24–28, pages 58–64. Addison-Wesley, 1994.