

EXPLORATORY SEQUENTIAL DATA ANALYSIS: TRADITIONS, TECHNIQUES AND TOOLS

CAROLANNE FISHER AND PENELOPE SANDERSON

1. MOTIVATIONS BEHIND THE CHI'92 WORKSHOP

In many areas of HCI, investigators make video or audio recordings of humans working with computers and other humans. These recordings capture verbalizations, actions, and often general aspects of the working environment as well. The investigator then explores and analyses these data in the light of some HCI research or design issue. Protocol analysis, video analysis, interaction analysis, conversation analysis, discourse analysis, sequential data analysis, task analysis, etc., are all examples of such research activity. What they have in common is a concern with handling dynamic, on-line, event-driven behavior that unfolds over time and with capturing the essence of that behavior in summary statements.

At present there are very few resources to help investigators make reasoned choices between these types of techniques, or to provide guidance on the appropriate use of a technique. As a result, investigators can fall into certain conceptual and methodological "traps", much time can be wasted and the potential richness of the data can remain untapped. The goal of this workshop was to bring together a group of people with a balanced mixture of interests, experience and background and to take the first steps towards establishing a conceptual foundation for understanding these techniques and to consider what types of practical guidelines are needed for the analysis of recorded sequential data.

In this workshop summary we will first quickly outline what is meant by the term "Exploratory Sequential Data Analysis". Next we will outline the workshop goals and describe its organization. Then we will summarize the main points and con-

clusions of the workshop summaries. Much of this material is drawn from post-workshop summaries prepared by the workshop participants themselves. Finally, we will list the participants and provide details of how interested readers might obtain further information about ESDA and contact members of the growing ESDA community.

2. WHAT IS ESDA?

"ESDA" (Exploratory Sequential Data Analysis) is simply a working term coined to cover a loose set of data analysis activities in the human sciences which deal with recorded data in which temporal information has been preserved (Sanderson, 1991; Sanderson and Fisher, 1992). The term is meant to reflect and organize what exists already, rather than to propose anything new. What follows is the conceptualization of ESDA presented to participants in the pre-workshop report and at the start of the workshop itself. While not necessarily ideal, this conceptualization was intended to be as encompassing as possible so that workshop participants would have some way of referring to the whole collection of different traditions and techniques they brought to the workshop.

The scope of ESDA activities--and therefore of what the term ESDA is intended to include--is illustrated in Figure 1. *Real events* of interest unfold in time, and can be captured either in *recordings* (such as audio, video, or environmental and behavioral data dumps) or directly in *encodings* of some kind. Once the real events have been captured, then an investigator's task is to distill the significance of these events with respect to some scientific or design-related purpose. The investigator wishes to

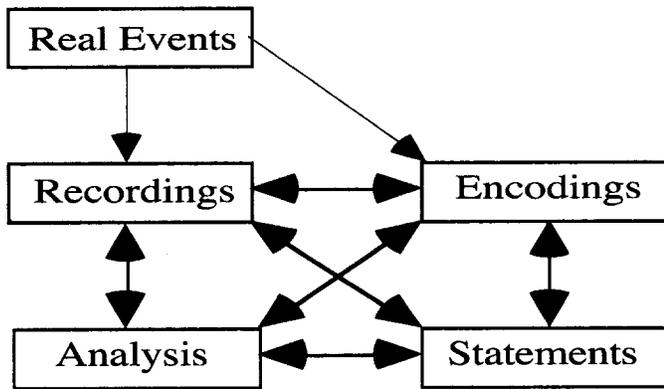


Figure 1: The scope of ESDA activities

make *statements* about this significance through some combination of classifying (or *encoding*) recordings and performing *analysis*. It is a theoretical question as to how many of these stages are needed, and whether what they add or remove is desirable or not. As the arrows in the diagram makes clear, once a recording (or encoding) is obtained, each step in the process can in principle affect any and all other steps. This illustrates the exploratory nature of this type of data analysis, so often remarked upon (MacKay and Tatar, 1989).

Very broadly, there seem to be three areas of research most relevant to these activities: Exploratory Data Analysis (EDA), Sequential Data Analysis (SDA), and Software Engineering

(SE). Their synthesis as "ESDA" is shown in Figure 2. EDA offers to ESDA a philosophy of handling data, not a literal set of techniques or models (Tukey, 1977). SDA points to the use of certain techniques, in which temporal relations and dependencies are preserved, whether the outcome is quantitative. SE offers possible software implementation strategies and tools for requirements distilled from the combination of EDA and SDA. Finally, of course, HCI is involved as an important (but not the only) domain of application. It has research needs that can be met by certain ESDA techniques and in addition it is an important aspect of the software engineering effort.

3. WORKSHOP GOALS AND ORGANIZATION

The goals of the workshop were to better understand various conceptual foundations for the analysis of recorded sequential data, to find mappings between research questions and ESDA methodologies, to explore how ESDA can be used to answer many of the tough questions inherent in HCI, and to begin to identify a community of practice in which ESDA can mature.

Accordingly, on Monday, 4 May 1992, 22 scientists and practitioners gathered to discuss these issues in the CHI'92 ESDA workshop. The participants had responded to an open call for participation by submitting position papers and were chosen with an eye to balancing their interests and backgrounds, which included anthropology, cognitive psychology, usability engineering, and computer science. We circulated the position papers to all participants in the form of a pre-workshop report. A list of participants and details of the pre-workshop report are given later in this summary.

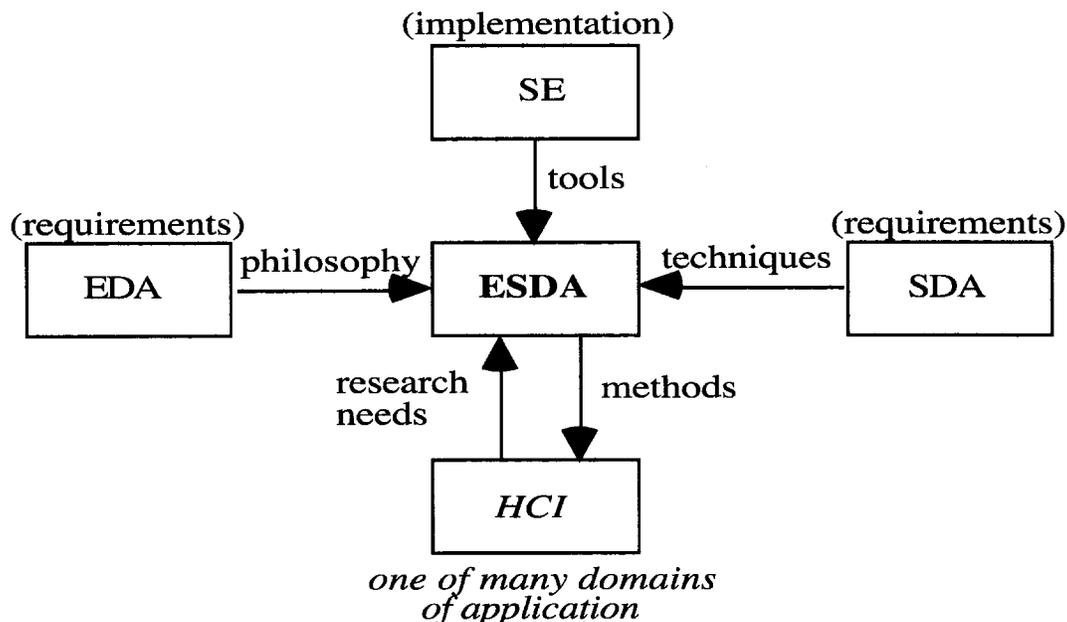


Figure 2: Areas of research contributing to ESDA

The real work of the workshop was done in the context of small and large group discussion organized around three central themes: Theory/Traditions, Techniques, and Tools. A separate segment of the workshop was devoted to each of these themes. Selected participants presented brief orienting talks that were intended to provide a quick but balanced survey of critical issues and distinctions, so that the ensuing whole-group discussion on that theme would be facilitated. A fourth segment of the day was devoted entirely to small group discussion and synthesis of all the topics covered.

Theory/Traditions. ESDA finds its roots in diverse research traditions and theoretical assumptions, most notably cognitive, anthropological, linguistic and ethological. Each tradition subtly colors the conclusions one draws from applying ESDA techniques. The goal of this first section of the workshop was to identify these traditions and to bring to light the theoretical assumptions behind some of them. As foundation for discussion, *Brigitte Jordan* first presented some of the philosophical roots of anthropology, in which holism, understanding the native point of view, and observer immersion in the events of interest are crucial, and gave examples of such observational studies. Although cognition is not a central topic for all anthropologists, there is an emerging movement in which cognition is considered to be socially and ecologically distributed and thus only understandable by examining interactions between agents in their natural contexts (e.g., Suchman, 1987; Lave and Wenger, 1991). *Carolanne Fisher* outlined the cognitive tradition of ESDA in which sequential data, notably verbalization, is viewed as evidence of the internal mental states of agents as they pursue goals in the context of their cognitive and environmental constraints (e.g., Newell and Simon, 1972; Ericsson and Simon, 1984). Such data can be used to build models of the cognitive processes of individuals. These presentations played an important role in helping participants understand some of the sources of controversies about ESDA techniques and tools and some of the underlying motivations for working within one tradition rather than another.

Techniques. Each ESDA tradition provides a way of thinking about sequential observational data that entails the use of certain techniques for analysis and modeling. The goal of this second part of the workshop was to begin to identify what kinds of practical approaches are suitable for what kinds of research questions, especially with respect to HCI. To the extent that there are problems that existing traditions and techniques fail to address well, particularly in the domain of HCI, our goal was to propose some new approaches. To orient the discussion, *Gerhard Deffner* spoke on conceptual and methodological issues involved when collecting qualitatively different sequential performance measures in parallel and relating them to each other. *Gary Olson* discussed the use of quantitative sequential data analysis techniques such as lag sequential analysis and grammar models, particularly as applied to sequences of actions by agents and devices in CSCW settings. *Nancy Cooke* discussed the use of structural modeling tools, such as Pathfinder, for analysing sequential data, especially for finding cause-effect relations. Finally, *Brigitte Jordan* introduced techniques from ethnography and ethnomethodology. The use of film and the study of interactions

have long been ethnographic tools. However, video-based interaction analysis is often associated more closely with ethnomethodology, a more recent branch of sociology which focuses exclusively on a highly detailed study of conversations and interactions. These presentations stimulated discussion on how a newcomer might be guided amongst these choices.

Tools. ESDA techniques are inherently time consuming and tend to be ill-defined, making them impractical or inaccessible for practitioners working under real world constraints. Automated and semi-automated software tools for ESDA are now beginning to appear which hold the promise of expediting analysis. However these tools embody commitments to certain theories and techniques in a manner which is quite explicit, but which can be missed if the practitioner is unaware of the space of alternatives. In this part of the workshop we examined tools that are currently available, tools that are in development, and what the requirements for such tools might be. *Beverly Harrison* reviewed a wide variety of video analysis tools and discussed her current tool development activities. *Penny Sanderson* presented some ideas about the generic requirements of ESDA tools, particularly with respect to data management, and discussed her tool development activities as well. Finally, *Lucy Berlin* presented a vision of how future tools might meet the requirements of researchers engaged in ESDA activities. In the lively discussion that followed, participant after participant talked about tools he or she had, needed, or could imagine, and provided illustrations of the type of data presentation formats he or she had found most helpful.

4. HIGHLIGHTS OF THE DISCUSSION

As expected, during a one-day workshop we could only begin to identify some of the issues that surround ESDA. However a notable feature of the day, and of the summaries submitted afterwards, was participants' sense of surprise and relief at having found so many colleagues grappling with the very same problems as they were. Many participants commented that there seemed to be an identifiable ESDA community, despite the differences in theoretical orientations, research questions and analysis techniques. This is because coordinating multiple approaches to commonly-shared research goals seems desirable and also because many concerns and problems seem fundamental and universal. These common concerns and problems formed the heart of the day's discussion. What follows is a summary of some of the issues that generated particularly heated discussion.

4.1. Seductiveness of ESDA data

ESDA data (e.g., audio and video taping, keystroke and eye movement logging, etc.) are easy to collect, can apparently be used to respond to any question and have great immediacy. However our ability to collect ESDA data has rapidly out-paced our ability to manipulate and analyse it. Thus these initial benefits are sometimes nullified by the failure to perform careful analyses. It is extremely rare that HCI professionals have been trained during their university education in how to

use ESDA techniques. The result is that many laboratories have shelves full of data and laboratory personnel are faced with the daunting task of wresting meaning from them without really knowing how to proceed.

According to one view, a symptom of this problem lies in the currently popular term "video analysis". Upon reflection, an HCI professional talking about "video analysis" can sound as odd as a biologist talking about "microscope analysis", an astronomer talking about "telescope analysis", or a psychologist talking about "tachistoscope analysis". The video medium (and other ESDA media) are so rich, immediate and accessible that they threaten to overwhelm the issue of how to select from what is observed and how to relate these selections to theoretical or design issues. Obviously, ESDA data (like all raw data) need careful, principled handling, analysis and interpretation to be useful. An important problem is that we do not yet have examples of research using ESDA which are widely acknowledged for their quality, which clearly articulate their techniques, and which serve as beacons, standards and guides for those entering the field. ESDA, and especially "video-ESDA" does not yet have its Newell and Simon, as it were.

4.2. Eclecticism: A meeting of traditions and a melding of techniques

The first part of the workshop highlighted the philosophical gulf between standard cognitive and more ethnographic approaches, as currently practiced. ESDA techniques originating in the cognitive tradition tend to be objective and symbolic, and are usually based in hypothetico-deductive, realist views of science. ESDA techniques originating within certain anthropological traditions are often founded on a belief in intractably reflexive relations between actor and observer, and are based in a relativist/idealist view of science. It was clear that theoretical orientation has an enormous effect on how investigators will wish to ground their conclusions in data--determining the type of data collected, how it is sampled, how it is analysed and the classes of interpretation that become available. As one participant pointed out, workshop participants represented different paradigms and practices. Thus some relied upon one class of data alone, such as mouse and key events, whereas others needed to take into account a great many dimensions of work environments. Some proceeded "top down" with well-developed ideas to test and well-defined techniques for doing so. Others worked "bottom up", starting with broader general hypotheses and moving in an iterative fashion towards an understanding of the processes involved in what they were seeing, taking a lead from suggestions in the data itself.

This led to the issue of how theoretical and interpretational bias should be addressed. In some anthropological traditions, especially interpretive anthropology, such biases are a topic for study in their own right. However, elsewhere investigators are at considerable pains to remove biases through the rigorous standardization of procedures (e.g. coding procedures) and through the establishment of inter-coder reliability. As one participant put it, however, this may actually *mask* existing

biases, particularly theoretical ones, rather than remove them.

Despite their different backgrounds, workshop participants were generally extremely open to the approaches and experiences of colleagues from other traditions and were eager to learn what they could adapt in their own research. Many participants stated they were less concerned with working within a single ESDA tradition than they were in adopting a more pragmatic, multidisciplinary approach. The issue of whether this could successfully be done without conceptual incoherence was raised, and one opinion was that the answer really lay in whether one was performing ESDA to build and test theory or to support design. Nonetheless, the workshop gave participants the opportunity to see the virtues of what other approaches can offer as final scientific statements, be they runnable models, strongly descriptive visualizations, compelling and insightful narratives or statistical power.

The notion that an ESDA "community of practice" is emerging arose at this point. One participant felt that there is an ESDA community in the same sense that there is a medical community with its many specialties. Each medical specialty has its separate goals, expertise, tests and treatment tools, but specialties share a range of common tools as well. A patient's problem may be treated by a surgeon, an occupational therapist, and a neurologist. Each will have a different focus, but the patient's recovery may depend on the physicians' combined tests and expertise. Similarly, in HCI-oriented ESDA we have a common focus on analyzing time-oriented data about complex human activity, but different practitioners specialize in different perspectives on the activity. Many participants felt that for typical HCI and CSCW activities (e.g., group authoring) each analysis technique would give a unique perspective but a combination might be required to solve a given question about functionality or the interplay of goals and actions.

Finally, another participant raised the provocative issue of whether different ESDA techniques, taken as a combination of formal, theoretically-driven practices plus informal research practices, might actually be more similar than theory would suggest and might actually lead to reasonably similar results. For example, an explicit and formal aspect of interaction analysis (e.g. as practised at Xerox PARC and Institute for Research on Learning) is that data analysis should be participatory. Multiple investigators meet regularly to review data and discuss interpretation. However this practice is also seen informally at laboratories which subscribe to more cognitive approaches (e.g. University of Michigan and NASA-Ames), even though such a practice is not prescribed as an essential part of understanding cognition through controlled experimentation. Another example is the use of codes in ESDA. Two extreme positions are, first, the idea that codes are essential and should be decided in advance on the basis of theory *versus* the idea that codes are to be avoided and, if used at all, should emerge only after an intense examination of the data itself. However, the first position (really a confirmatory rather than exploratory approach) is usually quite impracticable with observational sequential data, where a satisfactory symbolic interpretation can only emerge after considerable trial and

error. So, when ESDA is viewed as a combination of formal and informal research practices, one might wonder how completely the theoretical abstractions are indeed "filtering" the substantive insights gained.

4.3. The ESDA methodology "grid": feasibility of a research decision aid

An important central theme of the day was that ESDA techniques do not necessarily provide the best solution to a research problem. If there are simpler and less resource-demanding ways to answer a research or design question, then they obviously should be adopted. But if an investigator feels that an question can be answered only through ESDA, then the next issue is how to choose from the techniques available. What particular ESDA technique is most likely to produce meaningful results and what theoretical assumptions limit the kinds statements that can be made about those results? What kinds of time/resource commitments does the technique require and are any tools are available to help? What are the most useful rules of thumb or tricks of the trade for a given situation or technique?

Participants repeatedly mentioned the need for a taxonomy of techniques or a "grid" which would help researchers choose the most appropriate approach, given the characteristics of a research question and the investigator's resource constraints. In the small group discussions, participants started to examine the feasibility of such a research decision aid. Difficulties encountered were the usual ones found in such enterprises: finding the most appropriate dimensions to classify different types of research questions and different ESDA approaches, handling dependencies between dimensions, handling exceptions due to the failure to include a dimension peculiar to a novel research problem and, most importantly, deciding upon the advice to give. One participant referred to similar attempts at research decision aids in the area of knowledge elicitation, pointing out that grids are difficult to use, they proliferate, they can contradict each other, and they can only be expected to provide the loosest type of guidance. Grids always run the risk of being too simplistic or too general in comparison to the complexity of real research situations. The general feeling of the workshop attendees, however, was that the attempt itself would have heuristic value. It would provide an expression of what we currently know about the practical and theoretical range of ESDA and give us a starting point for further refinement. Moreover, it would serve as an orienting mechanism for practitioners attempting to identify the appropriate methodologies for the tasks they have at hand.

4.4. Data visualization

Sequential data is, in general, dense and often complex. Strong feelings emerged about the use of visualization techniques in ESDA. Tools emerging to support ESDA are making both static and dynamic (animated) visualization of data possible, and techniques for auditory exploration of sequential data patterns are also being developed. Paradoxically, the workshop participants seemed to agree that data reduction and visualization were necessary and desirable, but many simulta-

neously feared losing data. One suggestion was that the issue could be framed better in terms of validity or meaningfulness. The issue should be whether data reduction highlights the most important or meaningful distinctions in a form that is useful for the task at hand (e.g., interface design). Unfortunately, questions of validity are not easily answered, and the evaluation of techniques in terms of their practical utility requires a huge and confounding leap from the reduced data themselves to the ultimate *use* of those data.

4.5. ESDA software tools: reality and fantasy

Many workshop participants are currently involved in developing multi-media software tools to support their ESDA activities. Some of these tools have been developed to support a particular research program, and are thus highly specific to the research context, research questions, data formats, and presentation needs of just that research program. Other tools represent attempts to build context-free ESDA environments that could be used for a wide variety of research questions. Other still remain in the requirements and design stages. Overall, there was a feeling that a new type of software is emerging. Participants had strong feelings about what such software should include and felt considerable excitement at the development already happening. This paper is not the place to review the tools available or to outline the "wish list" that is emerging--papers offering this are forthcoming and more information can be gleaned by contacting the ESDA mailing alias or individual workshop participants (see end of paper for details).

Despite the different theoretical orientations of their creators, there are now enough ESDA software tools around to be able to detect the common features across tools. The coordination and integrated display of multiple sources of information is universal, as is the computer-based control of any video connection. Many tools offer "basic level" spreadsheet-like data displays which are supported by a wide variety of alternative displays the user can choose from. Graphical data visualization capabilities at different grains of spatial and temporal analysis are appearing. Many tools include some capability for classifying events and include a range of statistical analyses, including content analysis and various types of sequential data analysis. The ability to select or "filter" certain parts of the data for viewing or analysis is also a commonly agreed-upon feature. Different ways of implementing end-user pattern-recognition query languages were also seen. In summary, tools are being developed that will have many very attractive features which will raise our expectations of this class of software and the software engineering challenge will be with us for some time.

Despite the ambitious ideas emerging and already being implemented, an ESDA tool which satisfies the needs of all workshop participants has certainly not yet been achieved. It is debatable whether a tool of such breadth is even possible. Some people's tools will be useless for other people, and vice versa. Many participants reported that they had very successfully adapted to their needs various existing applications, such as spreadsheets, wordprocessors, statistical packages, graphing

packages and project planners. As a result, transportability between applications--whether originally intended to support ESDA or not--will be the crucial issue for some time.

5. COMMUNITY, CONTINUITY AND COMMUNICATION

There was a clear consensus that the CHI '92 ESDA workshop was valuable in itself, but even more valuable as a catalyst for bringing together an active community of practice. It is important to point out that at CHI'88, Wendy MacKay and Deborah Tatar organized a workshop that explored the possibilities of video technology for a wide variety of purposes—a workshop now widely recognized as having been an important event in the development of the use of video in the HCI context. The present workshop has built on this momentum, but was designed with a narrower focus. First, it concentrated more closely on the research enterprise. Secondly, it treated video as just one of many types of data included under the working definition of ESDA.

The ESDA community is becoming better defined, and we hope it will serve some very practical needs such as those listed below. We have already taken some important first steps towards achieving them:

- Establishing a clearing-house of information on available tools and handbooks
- Establishing a forum for discussion of methodologies, tool designs, etc.
- Sharing and archiving information about the members' interests, techniques
- Collecting and organizing pragmatic know-how for an ESDA troubleshooting guide
- Developing handbooks and editing journal special issues on ESDA.
- Maintaining links to the computer science, HCI, psychology, and anthropology communities

6. WORKSHOP PARTICIPANTS AND INFORMATION SOURCES

Below is a list of the 22 workshop participants and the titles of their pre-workshop position papers. These papers were collected into a 100-page pre-workshop report which was distributed to all participants before the workshop.

1. Albert Badre, *Georgia Institute of Technology* (non-attending coauthor: Paulo Santos): "Integrated knowledge-based interaction monitoring and analysis."
2. Lucy Berlin, *Hewlett-Packard Laboratories*: "Supporting the process of discovery in exploratory ethnographic analysis."
3. Nancy Cooke, *Rice University* (now at *New Mexico State University*): "Structural modeling techniques and sequential analysis."
4. Tom Dayton, *Bellcore*: "Suiting ESDA to the real work of HCI."
5. Gerhard Deffner, *Texas Instruments*: "Coordination of different types of sequential data."
6. Carolanne Fisher, *MAYA Design Group*: "ESDA: Traditions, techniques, tools--and traps."
7. David Frohlich, *Hewlett-Packard Research Laboratories*: "Conversation analysis as a model for exploratory sequential data analysis."
8. Christine Halverson, *University of California, San Diego*: "Exploratory sequential data analysis in the aviation context."
9. Beverly Harrison, *University of Toronto*: "Video analysis tools."
10. James Herbsleb, *University of Michigan*: "Sequential analysis of group design meetings." (coauthored with participant Gary Olson)
11. John Hiller, *University of New South Wales* (non-attending coauthor: Jane Millar, Aeronautical Research Laboratory, Melbourne): "Sequential data reduction techniques for reviewing the development of complex software."
12. Brigitte Jordan, *Xerox PARC* and *Institute for Research on Learning*: "Anthropological and ethnographic approaches."
13. Irvin Katz, *Educational Testing Service*: "Tools for the preliminary stages of protocol analysis."
14. Michele Morris, *BNR Europe Limited*: "Exploratory sequential data analysis and multi-user situations."
15. Kathleen Mosier, *Stanford University* and *NASA-Ames Research Center*: "The investigation of sequential data in full-mission flight simulation."
16. Gary Olson, *University of Michigan*: "Sequential analysis of group design meetings." (coauthored with participant James Herbsleb)
17. Virginia Peck, *Carnegie Mellon University*: "Data analysis in the development and evaluation of a cognitive model."
18. Penelope Sanderson, *University of Illinois at Urbana-Champaign*: "Exploratory sequential data analysis: A methodological issue and a design challenge for HCI and cognitive engineering."
19. Hermina Tabachneck, *Carnegie Mellon University*: "Verbal protocols, visual representations, and individual differences."
20. John Tang, *Sun Microsystems Laboratories*: "How can we integrate video analysis with other exploratory methodologies?"
21. Barbara Wildemuth, *University of North Carolina at Chapel Hill*: "End-users' on-line searching behaviors: Data analysis techniques."
22. Charles Wood, *University of Sussex*: "Exploratory encoding and analysis of the use of mediating representations in conceptual authoring."

If SIGCHI Bulletin readers would like a copy of the pre-workshop report, please make out a check for \$5 to "University of Illinois" and send it to the following address:

EPRL Information Services,
Engineering Psychology Research Laboratory,
Department of Mechanical and Industrial Engineering,
University of Illinois at Urbana-Champaign,
1206 West Green Street,
Urbana, IL 61801
U.S.A.

Additionally, we have established an on-line electronic mail forum as an ongoing resource for those with questions or

information about ESDA activities. Electronic mail sent to ESDA@maya.com will be forwarded automatically to everyone on the ESDA mailing list (which includes everyone who submitted position papers for the workshop). If you would like to be added to the list, please send email to ESDA-request@maya.com.

7. ACKNOWLEDGEMENTS

The authors would like to thank the ESDA workshop participants for their workshop summaries, on which a substantial part of this paper is based. Thanks are also due to Bonnie Nardi for necessary criticisms of an earlier draft, and to Lucy Berlin and Gitti Jordan for helpful discussions.

8. REFERENCES

- Ericsson, K.A. and Simon, H.A. (1984). *Protocol Analysis: Verbal Reports as Data*. Cambridge, MA: MIT Press.
- Lave, J., and Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, UK: Cambridge University Press.
- Mackay, W.E. and Tatar, D.G. (1989). Introduction to a special issue on video as a research and design tool. *SIGCHI Bulletin*, 21(3), 48-51.
- Newell, A and Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Sanderson, P.M. (1991). Exploratory sequential data analysis. Engineering Psychology Research Laboratory Technical Report EPRL-91-04, Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, Urbana, IL.
- Sanderson, P.M. and Fisher, C. (1992). Exploratory sequential data analysis: Traditions, Techniques and Tools. Pre-workshop report for CHI'92 workshop on ESDA. Engineering Psychology Research Laboratory, Department of Mechanical and Industrial Engineering, University of Illinois: Urbana, IL.
- Suchman, L.A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: UK: Cambridge University Press.
- Tukey, J.W. (1977). *Exploratory Data Analysis*. Reading, MA: Addison-Wesley.
-