EXPLANATION FACILITIES AND INTERACTIVE SYSTEMS

Hilary Johnson and Peter Johnson

HCI Laboratory
Department of Computer Science
Queen Mary and Westfield College
University of London
London, E1 4NS, UK.

Tel: +44- 71- 975 5246/5224
Fax: +44- 81- 980 6533
e-mail: hilaryj@uk.ac.qmw.dcs; pete@uk.ac.qmw.dcs

ABSTRACT

Our main research aim is to improve the provision of explanation facilities in information systems generally, and to identify what is meant by "explanation". This paper reports research which identifies both the strengths and weaknesses of current research and shows how to overcome those weaknesses. We are also concerned with both present and future uses of explanation in information systems and the role of explanation in a broad range of interactive applications.

KEYWORDS: Explanation, explanation theories, task-based explanation, Task Knowledge Structures (TKS)

1. INTRODUCTION

The research reported in this paper is motivated by two overriding goals; first to improve explanation facilities in information systems generally, and second to provide input to the debate about the nature of explanation from a cognitive science perspective, (Johnson and Johnson, 1992).

In this paper we discuss research issues related to both the usefulness of explanation in information systems, and we also report on the program of research we have conducted to overcome some of the apparent problems with current explanation provision.

We have two interests in the use of explanations; those uses which are focused at the end user of the system, and also in the role of explanation as a communication medium within the design process, such as for communicating design decisions between designers of information systems and subsequently to end users. This paper only addresses the former use of explanation.

In investigating the uses of explanation, we need to consider both the relationship between explanation and intelligent interfaces and also what is commonly meant by intelligent interfaces. Generally speaking, intelligent interfaces are those which demonstrate an ability to assist users in their task(s) and can be seen to display context sensitive behaviour. Intelligent interfaces can refer to a wide range of abilities including; advising, critiquing, problem solving and planning. Explanation provision is one important component of an intelligent interface which requires understanding and modelling in a manner similar to that which problem solving and planning have been subjected in scientific investigation. However, we should not lose sight of the application requirements to produce interactive sytems with interfaces that support and help the users in their tasks. Explanation already plays a crucial role in the interaction between users and complex systems, whether the explanation facility is unintelligent (as with canned-text and passive help systems) or intelligent.

Some researchers have argued for intelligent multimedia interfaces (e.g. Kobsa, 1991) indicating that they must be able to interpret user questions, formulate appropriate replies, provide answers to follow-up questions, make backward references and "intelligently" index and retrieve information. It seems clear to us that the explanations provided by an intelligent interface should also be tailored to the characteristics of users and to particular stages of their tasks. Additionally, it is important to remember the purposes of providing an explanation, as this will determine both the nature of the explanation to be provided and also how it is presented to users. In the next section, we consider a wider role for explanation in information systems.

2. THE UBIQUITOUSNESS OF EXPLANATION.

It is unnecessary for us to substantiate the statement that explanations are ubiquitous in our everyday lives. However, this leads to a number of further questions, for instance should explanations also be a ubiquitous feature of computer systems? Where are explanation facilities currently employed? What are the uses (both current and future) of explanations? If explanations are not a characteristic of all computer systems, then why not? What would be the benefits of introducing explanation facilities?

With regard to the first and second questions, the answers depends on what constitutes an explanation. Explanations feature in varying degrees and forms in knowledge-based systems, in intelligent tutoring systems and to some extent in help systems. But the nature of explanation in each of these cases differs widely. For the purposes of this paper we assume two continuums on which explanations differ. The first continuum relates to the nature of the explanations themselves, encompassing for instance, the instructional kind of explanations which feature in user manuals to the logically causal explanations which are considered to be a feature of some knowledge-based systems. The second continuum is concerned with the "intelligence" of the explanation provision, whether it is context insensitive like that supplied by user manuals, or whether it is context sensitive as when an intelligent tutoring system tailors its output to the user, the user's task and uses an appropriate interaction or pedagogic style.

In answer to the third question as to the uses of explanation, we feel that generally explanation facilities are under-utilized and that in the future as the nature of interfaces changes, so the need for coherent explanations will become more important. We predict a role for explanations in at least two spheres; first, in the realm of demonstrational interfaces and secondly, in interfaces which have task scenario facilities. Demonstrational interfaces (Myers, 1988) are those which follow the user undertaking (demonstrating) tasks and which as a consequence are able to make inferences about future task behaviour. These inferences made by the system then allow the interface to carry out the task for the user, thereby allowing the user to undertake other tasks or the less repetitive aspects of tasks. Although there are real research issues related to the validity of the inferences the system might make and also regarding the degree of control by the user, such interfaces have a role to play in the future of the computer as a personal assistant. Explanation has a crucial role to play here, where it is important for the system to explain to the user the inferences it has made and the resulting actions it has taken and why those actions have been taken.

A second sphere where we feel that explanation must play a part is in those systems where the interface has a task scenario planning or problem solving facility. Already such systems have been proposed, for instance in Air Traffic Control, the FEATS scenario, which is one proposed solution to the ATC problems in Europe until 2015, makes such a suggestion. A central notion in FEATS is that the system will generate and present to the pilot an ideal task scenario for carrying out particular tasks in the cockpit, and in addition the system will be able to provide idealized solutions when there is a problem. In such cases it is important for the system to be able to explain and justify the ideal task scenarios and optimum solutions suggested. These are only two cases where we can see that explanation can play a crucial part in interacting with the user at the interface. Others might include the use of critics (see Silverman, 1992) for correcting user knowledge in terms of its clarity, coherence, correspondence and workability. A critic should also be able to explain to the user the nature of his or her mistake and also explain any subsequent course of action and the reasoning for it.

The role of explanation in information systems can therefore be looked upon as at least two fold. The system either provides knowledge and explanations necessary for the user to carry out his or her task, or alternatively, the system carries out some action and then explains the need and reason for the action the system itself has taken to the user.

The final questions raised relates to why explanation is not a feature of every computer system and what would the likely benefits be? There are many reasons as to why explanation facilities are not more prevalent in present information systems. For example, designers might not see explanation as useful or have no perception of the users need for explanations. They (designers) might view explanations as having a high development cost with little added value benefit. Additionally, explanation facilities are not presently in a form in which they can be easily incorporated into existing system or systems development procatices. Furthermore, research is needed into the type of explanations required in different contexts and how to generate and present the explanations. Certainly the computational load of incorporating explanations is high, and when systems have to be successful, cost-effective and meet strict deadlines, an additional facility with no proven worth which increases development time will not be looked upon favourably. Many of the reasons mirror the reasons why for a long time, HCI was not considered as part of the

design process (see Johnson and Johnson, 1989). However, there is an important difference, because every interactive system must have an interface, but not every system must have an explanation facility. The benefits we see for including explanations is more information and knowledge for the user, thus hopefully easing the interaction and also improving user understanding and task performance. These benefits however, directly relate to the purpose of providing the explanations. However, explanation is not a substitute for good user interface design. Just as user interface design is now an integral part of system design so should explanation provision be.

If explanation provision is to become a characteristic feature of many future interfaces, then there is a special responsibility for researchers in both HCI and AI to provide input to the debate about the nature of the explanations to be provided in future information systems. The onus on us as researchers in the area is to ensure that we profit by past research on explanation provision, identify the strengths and weaknesses in present research and build on the strengths and address the problems in the future.

3.CURRENT STATE OF EXPLANATION PROVISION.

In the rest of this paper we will outline briefly a line of explanation research which we have been following for a number of years. An aim of the research is to improve explanation provision in information systems. A first step was to conduct a comprehensive search of the literature in both knowledge-based and intelligent tutoring. It was revealed that there were significant advances that had been made in these areas. Commonly, the explanation facilities provided by MYCIN are taken as the baseline against which later research can be set. In MYCIN explanations consisted of rule traces and could hardly be considered as optimum explanations in terms of intelligibility. Significant advances have been made in particular areas such as user modelling (Rich, 1989; Kobsa and Wahlster, 1989; Quilici, 1990); different tutorial or explanatory dialogue styles (Burton and Brown, 1982; Clancey, 1979; Johnson and Johnson, 1991a; London, 1990) and plan recognition (Appelt and Pollack, 1990; Cohen, Song, Spencer and van Beek, 1990; Wu, 1990). There has also been research conducted on the impact of user modelling on text generation in task-centred settings, for example Wolz (1990). Finally, research by Moore (1989) on follow-up questions is particularly noteworthy.

However, the literature review also revealed some limitations in the current research. There are at least three weaknesses which need to be addressed. First, there is no unifying theory of explanation, and few references to historical accounts of explanation, particularly in different disciplines. Few papers actually considered what explanations might be (the core definition of explanation), or what might be acceptable as an explanation in everyday life as opposed to scientific explanation. Since there is no unifying theory, few explanation specific models or processes have been put forward, and there have been no hypotheses about explanation generated which could be tested in any manner. A second weakness is that the lack of any theory of explanation also gives rise to the absence of any criteria for judging if an explanation is in any sense good or bad (whatever goodness or badness might be). The criteria need to take into account what the explanation is to be used for, (the purpose of explanation). If it is known what the purpose of the explanation is then it is possible to set a target for what constitutes a successful, acceptable or cooperative explanation which will go towards achieving that target. Different targets might be seen to need different criteria and different views as to what is successful. A third and final weakness which concerned us was that there were very few empirical studies which evaluated to any degree the resulting explanations. Although we acknowledge that much research has gone into actually constructing elaborate knowledge bases with which users can communicate and that the effort expounded and the problems overcome already to achieve this are not small, we nevertheless feel that claims about *improved* explanations should have some basis in fact and that they should be supported or substantiated in some way by empirical evidence. Otherwise the claims should be treated as vacuous. If the aims of the research are to construct knowledge bases which are more efficient or with shorter response times then the research has made significant leaps forward. However, we still cannot ignore the fact that explanations in order to warrant the label "explanation" have to be intelligible.

As a result of our survey we began a program of research to address some of the problems.

4. ADDRESSING THE LIMITATIONS.

We aimed to improve explanation provision, particularly explanations provided when users were in the process of undertaking tasks, by first developing a domain independent theory of explanation which would provide:

i) details about what constitutes explanation
ii) predictions about the content, timing and patterns of explanation provision while carrying out tasks
iii) indications as to whether an explanation was likely to be successful, prior to and after provision,
iv) if an explanation failed, an explanation of why it failed and how the failure could be rectified
v) models of the process of providing explanations in human-human dialogues
vi) development of prescriptive rules (from the model) to be followed in providing explanations
vii) the basis for development of a set of criteria to be satisfied by explanations and their evaluation, and finally,
viii) a consideration of how well the prescriptive rules from human-human dialogues can be applied to human-computer dialogues.

4.1 Theory development.

As a first stage in the development of the theory, we undertook a multidiciplinary literature review of the nature of explanation. In addition we also conducted a series of empirical studies (questionnaires and observational studies with interviews and comprehension tests) of human-human explanations.

The theoretical issues related to what constituted explanation in theory and in practice; the content, timing and patterns of explanation provision; the process of explanation and the development of a set of prescriptive rules governing the uses of explanation.

What constitutes an explanation? Historical accounts of explanation in philosophy, psychology and cognitive science literature can be characterized by one or more of the following:

i) associations between antecedent and consequent
ii) contrasts and differences
iii) causal mechanisms.

There are other accounts of explanation which are variants on the above three themes, which often can be distinguished by for instance what is termed differences or contrasts. Others are a variant on the causal mechanism acounts of explanation which build on the notion of people's naive theories of the world.

Associations between antecedent and consequent accounts of the relationship between cause and effect, began with Hume, who directly relates the cause of a later event with the occurrence of the first. Berkeley builds on the work of Hume but talks rather of contrasts and differences, where the non-event of the consequent is something that needs to be explained, and where the second event occurs but without the first. Harre refutes both of these accounts of explanation by proposing that the only relationship between cause and effect relates to the existence of a causal mehanism. These causal mechanisms relate to our past experience and our making sense of the world. In this way they can be viewed as our naive theories of the world. Such accounts have latterly been taken up by cognitive psychologists but not always with the knowledge of where the initial account came from. For a full discussion of historical accounts of explanation, see Johnson and Johnson, 1992.

Given the previous theoretical research on explanation mentioned above we were interested to see whether such accounts of scientific explanation are an adequate model of everyday explanation. Consequently, a number of empirical studies were undertaken. The studies consisted of questionnaires, observational studies and comprehension tests. The results of a questionnaire into the nature of explanation revealed that although the participants consider explanation involves making causal connections between events and even comparing and contrasting events, they saw a much wider role for explanation. Explanation was considered as involving the activity of explaining about something, and therefore explanation was cast in a teaching and instructing role where the emphasis was on providing knowledge or achieving mutual understanding. The results suggest that providing explanations in everyday life is a much more varied and broad-based activity than identifying causes of events emphasized by theories of explanation. The results also indicate that people are less restricted in what they will accept as constituting explanations.

The activity of "explaining" then might include teaching and instructing, where the individual explanations which comprise the explaining activity, might conform to making causal connections and contrasts and differences.

In terms of the observational studies we conducted, the majority of the explanations were concerned with providing information. In these observational studies, novices had with the help of an expert, to carry out one or more tasks. It was found that more than half of the information provision explanations involved provision of declarative knowledge. The remainder of the explanations were divided between explanations involving the provision of declarative and procedural knowledge (13%) and those involving procedural knowledge alone (20% of total explanations). Only 17% of the total explanations involved causal mechanisms and less than 10% involved contrasts or differences.

The process of providing explanations gleaned from both the questionnaire and the observational studies is that first the explainer attempts to establish the explainee's existing knowledge. Assuming that the explainee is a novice the explainer then goes on to explain key terms of the concept to the explainee. If however, the explainee is not a total novice the explainer attempts to identify as accurately as possible the gaps in the explainee's knowledge so that they can be filled with appropriate knowledge. Finally, the explanation dialogue ceases when either the explainer is satisfied that the explainee understands the explanations or when either the explainer or explainee are no longer interested in continuing the dialogue.

With respect to the content of the explanations, the questionnaire participants stated that both declarative and procedural knowledge would be provided to the explainee, with the majority of the participants indicating that they would provide declarative knowledge first and then procedural knowledge later.

We considered several hypotheses about the types of explanations provided at different stages of tasks, and we were particularly interested in the content, timing and patterns of explanation. The findings in brief were that different types of knowledge were provided at different stages of the task, with declarative knowledge provided first and procedural knowledge later. Different explanations types were used at the beginnings and ends of a task or part of a task such as a subgoal and there were particular places where analogy was used. Causal mechanism explanations were given predominantly at the end of the task, where reasons and justifications of something having occurred were provided. From the results of the literature review and the studies we have developed a theory of task-based explanation, see Johnson and Johnson, 1992.

5. FROM THEORY TO APPLICATION.

The theory of task-based explanation has given rise to a model of the knowledge required to give an explanation. As a result of this model, derived from the theory and empirical studies, we have in addition devised a set of prescriptive rules to be followed in providing explanations. It is our intention that these prescriptive rules will guide the provision of explanations in both human-human and human-computer dialogues.

It is also possible from the theory to devise a set of hypotheses to test regarding the knowledge and types of explanation used at particular task stages. In turn, testing the hypotheses should hopefully provide support for the theory. In order to do this we are presently conducting further empirical studies in different domains, using different explainers and explainees. Any necessary revision of the theory will take place on analysis of the results of the studies.

Figure 1 Structure of TKS

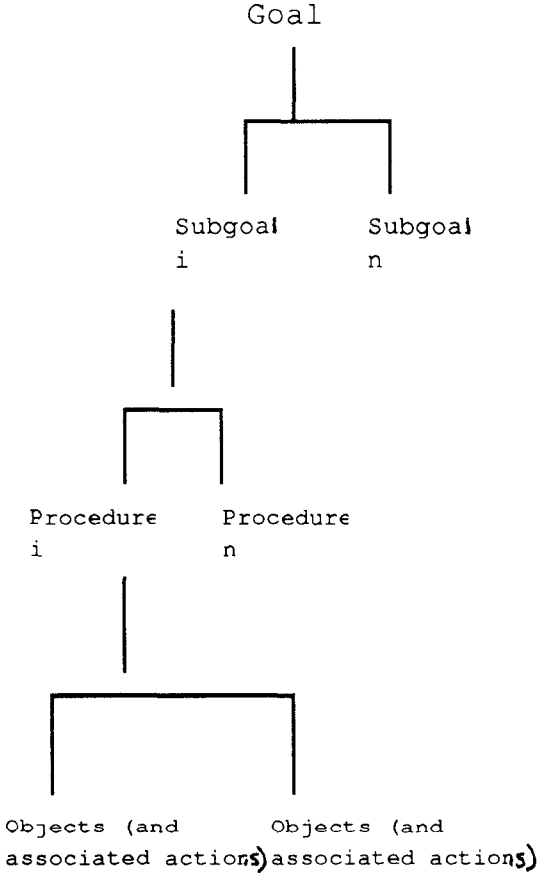

We have previously developed a theoretical framework and approach to task analysis, known as Task Knowledge Structures (TKS). (For background information see Johnson, Johnson, Waddington and Shouls, 1988; Johnson and Johnson, 1991b). This approach to task analysis was developed in the context of improving the design of human computer interfaces. The central assumption was that models of the user's task knowledge would provide system designers with knowledge which would enable them to take the users' task into account in designing better interfaces.

TKS is a well established task modelling approach with a well defined method of task analysis. It provides a summary representation of the different types of knowledge required or recruited by a user in performing a task. This knowledge is assumed to be stored in memory as conceptual structures and described in terms of a goal-oriented substructure and a taxonomic substructure. The goal-oriented substructure models the goal and sub-goal relations, and the procedures necessary to achieve the goals and subgoals for each task. The procedures also model actions upon objects and their control structures. Procedures are executed to satisfy the lowest level sub-goals. Finally, the taxonomic substructure models objects and their properties. It also models the category structure or class membership of objects. Figure 1 demonstrates the overall structure of TKSs.

5.1 Further empirical studies.

It is intended that the further studies will enable us to test the hypotheses and validate the theory. Additionally, new domains in which to investigate explanation will provide the necessary generality needed. The new domain chosen for study was "scientific research" and explanations to support subjects carrying out the task of "analysing a set of data" has been investigated. The subjects (explainees) in these studies have been given a set of data to analyse and a human expert provides explanations about the task in question and about how to analyse the data.

However, before such a study could be undertaken, it was necessary to thoroughly analyse what is involved in the task of "analysing data", in order that natural cuts in the task and the task goal structure can be identified. In previous empirical studies we have observed that different types of explanations are used at different stages of the task and one of the hypotheses we entertained in conducting the present empirical studies made predictions about the use of different explanations at different times. Prior to the present study we therefore set about analysing the task, "analysing a set of data", using Task Knowledge Structures (TKS) methodology.

A task analysis was conducted, using a variety of data collection techniques, principally from using statistical textbooks and interviews, to establish a TKS for the task of "analysing a set of data". The motivation for developing this TKS was to identify task knowledge which would enable us to indicate task stages where particular

explanations would be appropriate. The full TKS is available but we only have space here to show the goal substructure for the task, and this is particularly important in terms of the predictions we are making about the different explanation types. Figure 2 shows briefly the goal substructure for the "Analyse data" TKS.

Figure 2: Goal substructure for "Analysing a set of data" TKS.

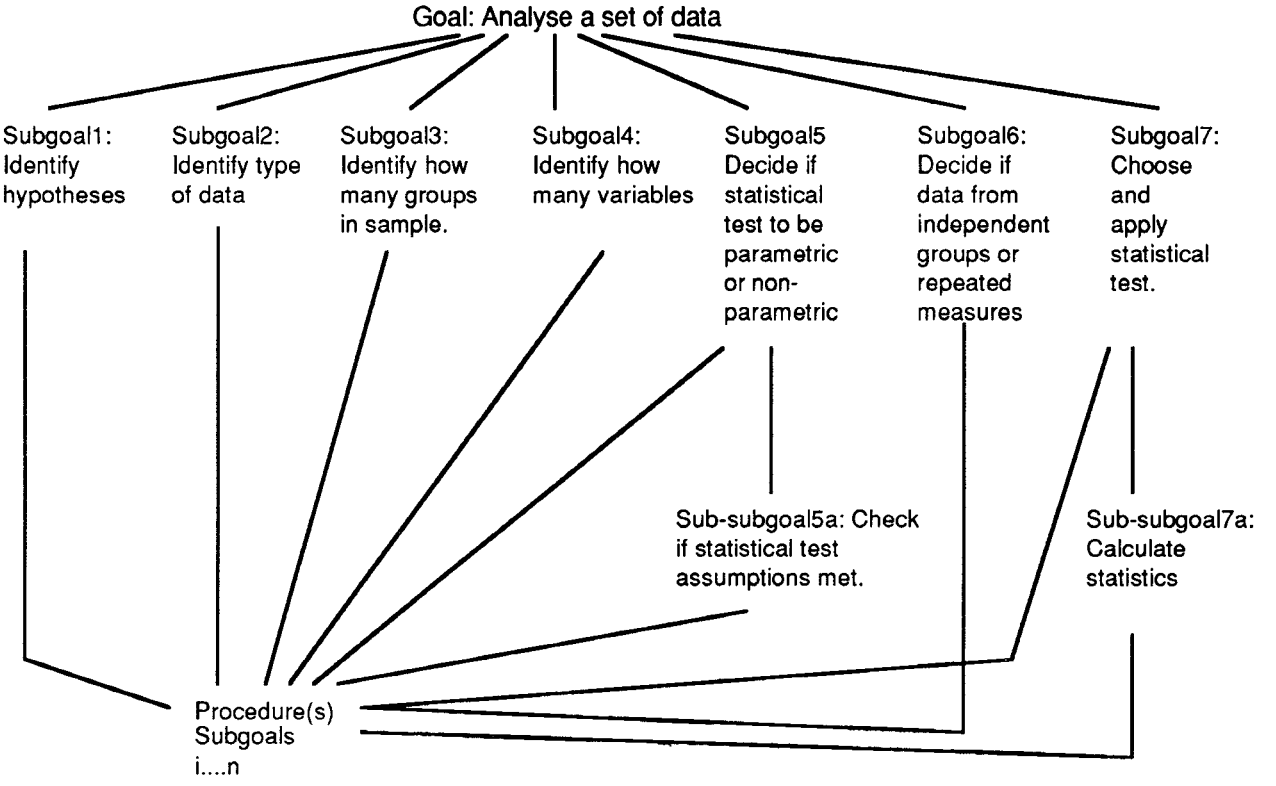

Predictions. According to the theory and previous empirical studies at each natural cut in the task, which here might be construed as subgoal beginnings and ends, there will be corresponding changes in the nature of the explanations. In particular we expect that at the beginning of each subgoal, that declarative knowledge will be provided, usually followed by an analogy to a similar concept, this will then be followed by procedural knowledge in the form of instructions and then more precise instructions. Finally, as the subgoal has been achieved, or is near to being achieved, explanations involving justifications and reasons and causes for events and procedures are given. We expect this pattern to be repeated at each of the subgoal divisions in Figure 2. The next step is therefore to see if these predictions hold true in the new studies which we are presently conducting. We have yet to analyse the data from these new studies so we are not in a position at the moment to discuss the results in terms of support or repudiation of aspects of the theory of task-based explanation.

6. SUMMARY AND DIRECTIONS FOR FUTURE RESEARCH.

We began this paper by stating that our two primary research objectives were to improve the explanation facilities currently provided by information systems and also to provide input to the debate on the nature of explanation. The impetus for improving the quality of explanations generally is, to some extent, because we see in the future a much wider role for explanation than has been the case hitherto. Therefore, in the second section of the paper we outlined other areas in which we see explanation becoming necessary and why. We also considered why explanation facilities are currently not more prevalent and what the benefits might be.

The third section of the paper was concerned with the strengths and weaknesses of current explanation provision and we stated that to improve in the future we have to recognize not only the strengths which we can build on, but

also address the limitations. Three limitations were worthy of attention: first, the lack of any unifying theory of explanation; secondly, the lack of a theory also gave rise to other problems such as the inability to identify and develop criteria against which to evaluate explanations, and finally, few empirical studies and validations of claims about improvements in systems.

Our program of research began with theory development and the following is a list of what we would expect to be the benefits of developing a domain independent theory of explanation:

i) details about what constitutes explanation
ii) predictions about the content, timing and patterns of explanation provision while carrying out tasks
iii) indications as to whether an explanation was likely to be successful, prior to and after provision,
iv) if an explanation failed, an explanation of why it failed and how the failure could be rectified
v) models of the process of providing explanations in human-human dialogues
vi) development of prescriptive rules (from the model) to be followed in providing explanations
vii) the basis for development of a set of criteria to be satisfied by explanations and their evaluation, and fianlly,
viii) a consideration of how well the prescriptive rules from human-human dialogues can be applied to human-computer dialogues.

With regard to the first item on the list, we believe that we have a better view of what constitutes explanation than we had prior to developing the theory and undertaking the empirical studies. Taking historical accounts as the basis for everyday explanation is not totally satisfactory. In everyday life in both theory and in practice, explanation is a more broad-based and varied activity with an instructing and teaching role not emphasized or taken into account by historical accounts of explanation. However, historical accounts of explanation do have a role to play in identifying the candidate constituents of explanation, they also provide a framework within which to view explanation generally and they have in the case of our research provided a basis for categorising the different types of explanation we found in the observational studies.

Another part of the theory are the postulates which relate to the content, timing and patterns of explanation. These postulates have given rise to predictions about the future use of explanation in terms of different types of explanation and explanation content at different task stages, and thus satisfy the second item on the above list. The third and fourth items on the list relate to the success or failure of the explanations. The theory should indicate with regard to the content and timing of the explanations whether or not they are likely to be successful. If they are not successful then the first step in recovery might be to consider the task stage currently being undertaken by the explainee, and then check both the content and the type of explanation provided. There are of course many other reasons why explanations fail, related to the erroneous use of language and terms, repetitiveness, not answering explainee's questions and so on. However, depending on the purpose of providing the explanations, we consider that task performance and execution warrants early consideration.

The fifth and sixth items on our research list are concerned with models of explanation provision and the prescriptive rules, neither of these have been desribed in this paper although we do have a model of what it is to explain, i.e. a model of the explainer, which comprises part of the prescriptive rules. The rules also we have not room to outline. We are presently testing the adoption of the rules by human-human experts with a view to extending this to be a model of human-computer explanatory dialogues. However, we are aware that it is not necessarily the case that humn-human dialogues (HHI) are an ideal case that human-computer dialogues (HCI) should follow. We expect that there might be overlaps in terms of the content, timing and patterns of explanation that could be extended to human-computer interaction, but differences in expectation on the part of the users and inabilities to appreciate the intelligibility of the explanations on the part of the computer mean that the ideal case will have to be not just diluted but significantly changed.

We therefore come to the part of the research program which is still underway, we have yet to analyse our further empirical studies, we also have yet to consider fully the implications of HHI for HCI and we are presently in the process of developing criteria for successful explanations. It is important in our view to link the criteria for success to the purpose of providing the explanations. For instance it is possible that the goal of providing the explanation changes in terms of optimum short versus long term task performance, or optimum short versus long term learning or transfer or knowledge. We are therefore establishing different and in some cases competing criteria for each of the different uses of explanation in information systems. Once these criteria are in place we will be in a position to address the most basic of the limitations, that of evaluating the resulting explanations and establishing in a scientific manner that they are indeed an improvement or otherwise on those provided in the past.

ACKNOWLEDGEMENTS

We are grateful to the Commission of the European Communities (CEC) for funding the ESPRIT Basic Research Action 3160 Interactive Dialogues for Explanation And Learning (IDEAL). Some of the content of this paper follows a talk given at the Task-based explanation workshop held in Samos, Greece, June 1992; however, the paper is previously unpublished.

REFERENCES.

Appelt, D.E. & Pollack, M.E. (1990) Weighted abduction as an inference method for plan recognition and evaluation. In proceedings of Second International Workshop on User Modelling. Hawaii.

Burton, R.R. & Brown, J.S. (1982) An investigation of computer coaching for informal learning activities. In D. Sleeman & J.S. Brown (eds) Intelligent tutoring systems. London: Academic Press.

Clancey, W.J. (1979) Transfer of rule-based expertise through a tutorial dialogue. Doctoral dissertation, Stanford University.

Cohen, R., Song, F., Spencer, B. & van Beek, P. (1990) Plan recognition beyond fixed plan libraries. In proceedings of Second International Workshop on User Modelling. Hawaii.

Johnson, H & Johnson, P. (1991a) Empirical investigation of different explanatory dialogue styles. In H.J. Bullinger (ed) Proceedings of the Fourth International Conference on Human Computer Interaction, Stuttgart. September.

Johnson, H & Johnson, P. (1992) On the nature of task-based explanation. Manuscript in preparation.

Johnson,P , Johnson, H ,Waddington, R and A.Shouls, (1988) Task Related Knowledge Structures: Analysis, Modelling and Application, People and Computers IV, Cambridge University Press.

Johnson, H. & Johnson, P. (1989) Integrating task analysis into system design: Surveying designers' needs. Ergonomics, Vol 32, no 11, pp1451-1467.

Johnson, H. & Johnson, P. (1991b) Task Knowledge Structures: Psychological basis and integration into system design. Acta Psychologica: Volume 78, p3-26.

Kobsa, A. Reported in article by M. Maybury, Intelligent Multimedia Interfaces. AI Magazine, Summer 1992, Vol 13, no 2.

Kobsa, A & Wahlster, W (eds) (1989) User models in dialogue systems. Berlin: Springer-Verlag.

London, R. (1990) Student modelling to support multiple instructional approaches. In proceedings of Second International Workshop on User Modelling. Hawaii.

Moore, J (1989) Responding to "Huh?": Answering vaguely articulated follow-up questions. In CHI'89 Proceedings, pp91-96.

Myer, B (1988) Demonstrational interfaces. Academic Press.

Quilici, A. (1990) The correction machine: Using common-sense planning knowledge to construct and exploit user models. In proceedings of Second International Workshop on User Modelling. Hawaii.

Rich, E. (1989) Stereotypes and user modelling. In A. Kobsa, & W. Wahlster, (eds) User models in dialogue systems. Berlin: Springer-Verlag

Silverman, B. (1992) Survey of expert critiquing systems: Practical and theoretical fromtiers. Communications of the ACM, Volume 35, no. 4. p106-127.

Wolz, U. (1990) The impact of user modelling on text generation in task-centred settings. In proceedings of Second International Workshop on User Modelling. Hawaii.

Wu, D. (1990) Implications of active user model acquisition for decision-theoretic dialog planning and plan recognition. In proceedings of Second International Workshop on User Modelling. Hawaii.