

# USING EXPLANATION FACILITIES IN HEALTHCARE EXPERT SYSTEMS

Keith Darlington

*London South Bank University, BCIM, 103, Borough Road, London, SE1 0AA, UK*

Keywords: Explanation, Expert Systems, NHS Direct, Prodigy.

Abstract: A great deal has been written about healthcare expert systems in recent years. This paper examines a particular feature of expert systems: namely explanation facilities. A limited explanation capability is an integral part of a rule based expert system. The role of explanation in expert systems has been largely ignored in healthcare literature, since the MYCIN system and its derivatives were developed in the mid 1980s. However, empirical research has shown that users are more likely to adhere to recommendations made by expert systems when explanation facilities are available. Furthermore, explanation provision have been shown to improve performance and aid the user with a better understanding of the subject domain as well as result in more positive user perceptions of an expert system. This paper looks at the evolution of explanation facilities in healthcare expert systems, and investigates user requirements for explanation facilities in the healthcare domain.

## 1 INTRODUCTION

The medical expert system MYCIN (Shortliffe 1981) was amongst the first of a number of decision support diagnostic systems developed in the late 1970s. Since this time, the use of expert systems as an IT decision making aid in healthcare has grown rapidly (Schank, Doney and Seizyk, 1988), (Thornett 2001). For example, NHS Direct Hotline uses an expert system in basic patient diagnostics. NHS Direct has been at the forefront of 24-hour health care in the UK - delivering telephone and e-health information services direct to the public, and is accessed by over two million people every month. PRODIGY (Thornett 2001) is another example of an expert system that is used in primary care in the UK. Prodigy provides decision support to general practitioners within consultations regarding drug prescribing.

But many researchers have been sceptical about expert system usage in healthcare. For example, Delaney et al, (1999) believe that computerised decision support systems have great potential but have largely failed to live up to their promise. However, Walton et al, (1997) presents evidence to suggest that advice from a computer will be more convincing if presented simultaneously with an explanation for that advice. In their evaluation of

CAPSULE, an expert system giving advice to general practitioners about prescribing drugs, they say that "Finding the most effective way of presenting the explanation is an important goal for future studies of computer support for prescribing drugs". Yet, an explanation component is a standard feature of expert systems in that the systems problem solving behaviour can be observed during a consultation.

The inclusion of explanation facilities can enhance performance of the decision making and lead to greater adherence to the recommendations of the expert system (Gregor and Benbasat, 1999), (Arnold et al, 2006). Indeed, many studies have demonstrated the importance of a system being able to explain its own reasoning. For example, in a study of physician's expectations and demands for computer based consultation systems it was found that explanation was the single most important requirement for advice giving systems in medicine (Buchanan and Shortliffe, 1984). Also, according to (Berry et al., 1995), explanation is seen as a vital feature of expert systems – particularly in high risk domains, such as medicine, where users need to be convinced that a system's recommendations are based on sound and appropriate reasoning. However, despite the importance attached to explanations, few expert systems provide acceptable

explanation. Surprisingly, research in explanations has been largely ignored in healthcare expert systems since the development of MYCIN. Mao and Benbasat (2000) cite reasons why explanations in expert systems had failed to appeal to the user: that they were difficult to understand, and that they ignored the needs of different users. The following sections examine ways in which some of these shortcomings have been overcome in the healthcare domain, beginning with a look in some detail at one of the first ever expert systems to incorporate explanation – the medical expert system MYCIN.

## 2 THE MYCIN EXPLANATION FACILITY

The explanation facilities in MYCIN (Shortliffe, 1981) were presented in a natural language that was translated from the rules making the explanations easier to follow for the user. Explanations were also supplemented with certainty factors that numerically expressed the degrees of certainty attached to conclusions or outcomes. This meant that users of MYCIN could get an understanding of the likelihood of the advice given. However, there were many shortcomings identified with MYCIN's explanations. The following sections describe these shortcomings and how they were overcome.

### 2.1 Rule Trace

MYCIN is a rule based expert system – which means that knowledge is stored in the form of rules (Darlington, 2000). The explanation facilities provided in MYCIN and the other first wave of rule-based expert systems would have been a rule trace. This is, essentially, a record of the system's run-time rule invocation history during a consultation.

#### 2.1.1 Feedforward and Feedback Explanations

A feedforward explanation provides the user with a means to find out why a question is being asked during a consultation (i.e., during the data input stage). The feedforward explanation will retain the manner in which input information is to be obtained for use in finding a solution.

A feedback explanation provides the user with a record of problem solving action during a consultation: i.e., how a conclusion was reached when the data has been completely input. The feedback explanations will present a trace of the

rules that were invoked during the consultation and display intermediate inferences in getting to a particular conclusion.

### 2.2 Strategic Explanations

Rule trace methods formed the basis of explanations in MYCIN (Shortliffe, 1981), but Clancey (1983) tried to adapt MYCIN from its diagnostic role to that of tutorial role in a system called GUIDON. The purpose of GUIDON was to provide a training system for junior consultants. What was thought to be a simple task turned out to be very difficult because MYCIN did not contain knowledge which explicitly contained strategic knowledge. This is knowledge about how to approach a problem by choosing an ordering for finding subgoals to minimise effort in the search for a solution. For example, the rule of thumb that alcoholics are likely to have an unusual aetiology can lead the expert to focus on less common causes of infection first – thereby pruning the search space to find a solution. The strategic knowledge in MYCIN was implicitly incorporated in the problem solving rules. However, Clancey realised that this knowledge needed to be explicitly represented, so that it could become transparent to students training to use the system. The lessons of GUIDON led Clancey to develop a follow up system called NEOMYCIN (Clancey & Letsinger, 1981): this was a consultation system whose medical knowledge base contained the tutorial strategic knowledge.

### 2.3 Justification Knowledge

An expert system can only reason with the knowledge contained in its knowledge base. Thus, the designer of a diagnostic medical expert system would ensure the knowledge base contains enough problem solving knowledge to ensure the system can arrive at the correct conclusions. However, the rule trace can only reconstruct a trace from what knowledge is contained in the expert system knowledge base. If the builder has not included the knowledge to justify the knowledge in the rule-base, then the system will not be able to justify the existence of the knowledge? Clancey (1983) realized the importance of this justification knowledge when attempting to extend the MYCIN system to support the training of junior physicians. Again, he found that MYCIN failed to do this because it did not contain justification knowledge. Justification knowledge is often unavailable because the rules which model the domain do not capture all the forms

of knowledge used by experts in their reasoning. This is because builders of expert systems capture “rules of thumb” shallow knowledge that only enable the system to solve diagnostic problems. Empirical research has consistently shown that user acceptance of expert systems increases for non-expert users when this justification knowledge is present, and that justification is the most effective type of explanation to bring about positive changes in user attitudes toward the advice giving system (Ye & Johnson 1995).

### 2.3.1 Capturing Deep Knowledge to Represent Justification Knowledge

Expert physicians would, of course, use rules of thumb themselves in solving problems, but they would also – as a result of their training and experience – possess a deep theoretical understanding of their subject domain. They may, for example, use “rules of thumb” or heuristic knowledge when performing a diagnosis. These “rules of thumb” may be sufficient to enable the physician to carry out a diagnosis, and therefore, this is the knowledge that is captured in rule-based expert systems because it is clearly much easier to obtain and code, and is sufficient for problem solving. However, this justification knowledge would have to be explicitly captured by the system designer if it were required for explanation.

## 3 USER REQUIREMENTS FOR EXPLANATION

In the healthcare domain, user requirements would vary according to employment categories which may include physician (including junior physician), nurses, administrators and also, as we will see later in some application examples, patients.

Considering the expert physician vs. non expert divide, research has shown that expert physicians do make use of explanation facilities but their requirements are very different to that of other users. Experts tend to have a preference for feedback rule trace explanations and are more likely – than non-experts – to use explanations for resolving anomalies (disagreement), verification and exploring alternative diagnoses (Arnold et al 2006), (Mao and Benbasat, 2000). Non-experts such as trainee physicians, on the other hand, are more likely to use explanations for short and long term learning. Moreover, (Arnold et al 2006), have

shown that non-experts tend to use both feedback and feedforward justification explanations, as well as terminological feedforward explanations, i.e., definitions of domain terms, etc., to assist during the data input stage. Patients – or other non-experts – using the NHS Direct system may benefit from this type of explanation because these explanations would facilitate learning of the subject domain. Some examples of these user-centred applications of explanation are briefly described in the next section.

### 3.1 Some User-Centred Explanation Prototypes in the Healthcare Domain

A number of healthcare projects involving the use of user modelling for explanation have been prominent in recent years. OPADE (Carolis et al, 1996) is a European Community Project funded expert system for generating beneficiary centred explanations about drug prescriptions that take into account the user requirements. The main objective of OPADE is to improve the quality of drug treatment by supporting the physician in their prescription process and by increasing compliance with the therapy (Berry et al, 1995). OPADE supports two types of user: those who directly interact with the system such as general practitioners and nurses, and those who receive a report of results – i.e., the patients. The explanations that are generated are dynamic (unlike static canned text explanations) in that a “user model” is maintained containing the characteristics of the user. A “text planner” component plans the discourse during a consultation. The text planner will build a tree containing the discourse plan which will depend on the objectives that are to be met by the user model. The explanation is then delivered in natural language by taking as input the tree generated by the text planner and transforming it using text phrases into the appropriate format.

HEALTHDOC (Marco et al, 1995) is another user-centred expert system whose main purpose is to generate reports for patients and materials for health education. This system enables the production of health-information and patient-education documents that are tailored to the individual personal and medical characteristics of the patients. Health-education documents can be much more effective in achieving patient compliance if they are customized for individual readers. The documents are presented textually and the text can be adapted to different patients, because the system contains a database

containing information about the clinical data of every patient – such as their personal and medical characteristics.

## 4 CONCLUSIONS

This paper recognises the role of expert systems in healthcare. However, one of their main features - explanation facilities – has been largely ignored in healthcare systems to date. Yet, empirical research has consistently shown, in recent years, that users are more likely to adhere to expert system recommendations when explanation facilities are available. Furthermore, explanation provision have been shown to improve performance and aid the user with a better understanding of the subject domain as well as result in more positive user perceptions of an expert system.

However, users will not use an explanation unless it addresses their basic information needs. This means that system designers must involve users in the evaluation of explanation facilities to ensure that they serve the needs of specific user groups. As this paper has shown, providing designers submit the effort, explanations can be tailored to the needs of different users. Perhaps the time has come for healthcare expert system designers, and users of such systems to re-evaluate the potential of explanation facilities.

## REFERENCES

- Arnold, V., Clark, N., Collier, P.A., Leech, S. A., Sutton, S.G., 2006. The Differential Use and Effect of Knowledge-Based System Explanations in Novice and Expert Judgment Decisions. *MIS Quarterly*, Vol 30 No 1.
- Berry, D., Gillie, D. T., and Banbury. S., 1995. What Do Patients Want to Know: an Empirical Approach To Explanation Generation and Validation. *Expert Systems with Applications*, 8(4):419-428.
- Buchanan, B., G., Shortliffe, E., H., 1984. *Rule-based Expert Systems: The MYCIN Experiments of the Stanford heuristic programming project*. Addison Wesley, Reading, MA.
- Carolis, B.D., Rosis, F.D., Grasso, F., Rossiello, A., Berry, D., and Gillie, T., 1996. Generating recipient-centered explanations about drug prescription. *Artificial Intelligence in Medicine* 8, 123-145.
- Clancey, W. J., 1983. Epistemology of a rule-based Expert System: A framework for explanation. *AI magazine*, 20(3) pp 215-251.
- Clancey, W. J., Letsinger, R., 1981. NEOMYCIN; reconfiguring a rule-based expert system for application to teaching. Pages 361–381 in W.J. Clancey and E. H. Shortliffe (eds.), *Readings in medical artificial intelligence: The first decade*, Addison-Wesley, Reading, Massachusetts
- Darlington, K., 2000. *The Essence of Expert Systems*. Prentice Hall, London. Pearson Education.
- Delaney, B.C., Fitzmaurice, D. A., Riaz, D., and Hobbs, F. D., 1999. “Can computerised decision support systems deliver improved quality in primary care?” *Journal of BMJ*. November 13; 319(7220): 1281.
- Gregor, S., & Benbasat, I., 1999. Explanations from intelligent systems: theoretical foundations and implications for practice. *Management Information Systems Quarterly*, 23, 497-530.
- Mao, J., Benbasat, I., “The Use of Explanations in Knowledge-Based Systems: Cognitive Perspectives and a Process-Tracing Analysis”, *Journal of Management Information Systems*, Vol. 17(2), pp153-180, 2000.
- Marco, C.D., Hirst, G., Wanner, L., and Wilkinson, J., 1995. *HealthDoc: Customizing patient information and health education by medical condition and personal characteristics*. In *First International Workshop on Artificial Intelligence in Patient Education*, Glasgow, UK.
- Schank, M. J., Doney, L. D., Seizyk, J., 1988. The potential of expert systems in nursing. *J Nurse Administration*. June; 18(6):26-31. PMID: 3286842 [Pub Med - indexed for MEDLINE].
- Shortliffe, E. H., 1981. *Computer Based medical consultations - MYCIN* Published by Elsevier, USA.
- Thornett, A., 2001. Computer decision support systems in general practice. *International Journal of Information Management*; 21:39-47.
- Ye, L.R., Johnson, P.E., 1995. “The impact of explanation facilities on user acceptance of expert systems advice”, *MIS Quarterly*, Vol. 19(2), pp157-172.
- Walton, R., 1997. An evaluation of CAPSULE, a computer system giving advice to general practitioners about prescribing drugs. *J INF Primary Care*.