

Educating Migraine Patients Through On-Line Generation of Medical Explanations

Johanna Moore and Stellan Ohlsson
Intelligent Systems Laboratory and
the Learning Research and Development Center,
University of Pittsburgh, Pittsburgh,
PA 15260, U. S. A.
jmoore@speedy.cs.pitt.edu
stellan@vms.cis.pitt.edu

Abstract

Computer support for learning in technical domains such as medicine requires an intelligent interface between the non-expert and the technical knowledge base. We describe a general method for constructing such interfaces and demonstrate its applicability for patient education. The employment of this technology in a medical clinic poses problems which are linguistic, psychological, and socio-cultural, rather than technological, in nature.¹

Real-Life Learning

People learn in many situations which are not classified as either schooling or training. Applying for a visa at a foreign consulate, appearing in court as plaintiff, defendant, or juror, placing an order with a travel agent, and visiting a medical clinic are

¹The preparation of this report was supported by grant No. 1 R01 LM05299-01 from the National Library of Medicine, National Institute of Health. The opinions expressed are not necessarily those of the sponsoring agency and no endorsement should be inferred.

examples of situations which confront the non-expert with technical knowledge that does not fit squarely into any traditional school subject. To participate successfully in such situations, the non-expert often needs to acquire some understanding of the relevant technical knowledge.

In this paper we focus on learning in the medical clinic. There is evidence that patients with more knowledge about their disease and their therapy get well faster because they comply more accurately and more conscientiously with the physician's prescriptions (Eraker, Kirscht, & Becker, 1984). The distinction between abortive and prophylactic treatments of migraine provides an illustration of the relation between knowledge and cure. Migraine patients who experience symptoms that consistently precede a migraine attack (e. g., visual disturbances) are typically given drugs that abort the attack. Because abortive drugs are ineffective if taken after the onset of an attack, patients without such warning signals are instead given prophylactic drugs that have to be taken on a regular schedule. When a patient complains that a prophylactic drug is ineffective, questioning might reveal that he or she stopped taking the drug when the headaches stopped; the headaches then

returned. Incorrect or incomplete understanding can jeopardize the therapy.

Several factors limit how much physicians can engage in patient education. (a) Doctor's time, particularly the time of specialists, is already a bottleneck in the health care system. (b) Doctors are not trained to communicate with people who do not share their expertise. (c) The doctor's interest and professional pride is typically invested in diagnosing the disease and finding a cure, not in explaining the same seemingly simple matters over and over again many times a day.

One solution to this dilemma is to use Cognitive Science technology to provide the patient with access to the doctor's knowledge without taking up the doctor's time. In collaboration with Bruce Buchanan and Diana Forsythe at the Intelligent Systems Laboratory and Gordon Banks at the Presbyterian University Hospital we are building a computer system which can generate answers and explanations, in English and on-line, in response to questions from patients (Buchanan, Moore, Forsythe, Banks, & Ohlsson, 1992). We are focussing on migraine for a variety of reasons, including the fact that it is a frequent and often disabling disease (Stewart et al, 1992). We first describe the technology we are using and then some non-technological problems that arise in its employment.

On-Line Medical Explanations

Unlike other A. I. systems in the medical domain, the purpose of our system is neither to automate diagnosis nor to train medical students, but to educate patients. The intended system responds to patient questions with answers and explanations which are generated on-line and adapted to the individual patient. We describe the knowledge base, explanation module, and query analyzer of our current

prototype, and indicate where they fall short of our goals. The prototype was implemented by Claudia Tapia (Tapia, 1991).

The knowledge base. The topics a migraine patient might want or need information about include (a) the physiology of migraine, (b) potential triggers for headaches, (c) the accompanying symptoms, (d) possible treatments, and (e) their side effects. Our goal is to encode a significant proportion of the medical profession's extensive knowledge about these topics (Raskin, 1988). The current knowledge base contains approximately 400 concepts referring to types and properties of headaches, symptoms, treatments, and drugs. No causal knowledge has as yet been encoded. The knowledge base is not an expert system; there is no inference engine for diagnosis or therapy planning. It is an open question whether on-line medical reasoning will ultimately be needed or whether we can encode everything we might want to explain to the patient in the knowledge base. The knowledge base is implemented in Loom (MacGregor, 1988).

The explanation module. Previous research on the generation of natural-language explanations indicates that an informative explanation cannot be generated from a knowledge base by translating internal code (procedures, rules, or schemas) into English (Buchanan & Shortliffe, 1984; Moore & Swartout, 1988), but requires a dedicated problem solver, called a *text planner*. The text planner described by Moore (1989), Moore and Paris (1989), and Moore and Swartout (1989) operates according to means-ends analysis: Post a goal, activate operators that can achieve that goal, and post the subgoals required by those operators; recurse until all posted operators are primitive. However, the goals and operators are not interpreted as physical situations and motor actions as in typical problem solvers and planners (see, e. g., Wilkins, 1988). The goals are *discourse goals*, i. e., effects that a speaker might

want his or her utterance to have on a hearer. Examples are to make the hearer *believe* some proposition and to *persuade* the hearer to perform some action. A different type of discourse goal is to establish some rhetorical relation, e. g., that an assertion P is *evidence* for some other assertion Q or that descriptions P and Q *differ* with respect to some attribute A. The operators encode rhetorical strategies by which discourse goals can be accomplished. For example, to make the hearer understand the differences between two objects, first describe what they have in common and then list their differences. Syntactically, an operator consists of an effect--a goal--and a conjunction of subgoals, the satisfaction of which is sufficient to achieve that effect. The application of an operator is guided by *constraints*, i. e., tests on the knowledge base (or on some other knowledge source; see below). The primitive operators are individual speech acts, e. g., to assert or to ask².

The details of this method for text planning have been published in Moore (1989), Moore and Paris (1989), and Moore and Swartout (1989). It was originally implemented in the context of an expert system for programming style (Moore, 1989), but we have successfully transferred it to the migraine domain. Our prototype contains approximately 35 operators, although this number is expected to grow; more than 75 operators were needed to produce satisfactory performance in the programming domain. The current system generates answers to three types of questions with only one or two seconds' delay. When asked to compare prophylactic and abortive treatments (i. e., given the request "COMPARE

migraine prophylactic treatment and migraine abortive treatment"), the system generates the following text (Tapia, 1991):

"Migraine prophylactic treatment and migraine abortive treatment are migraine pharmacological treatments. Migraine prophylactic treatment is used to prevent migraine while migraine abortive treatment is used to abort migraine. Migraine prophylactic treatment requires you to take a drug daily whereas migraine prophylactic treatment requires you to take a drug at the immediate onset of headaches. Migraine prophylactic treatment is suitable for frequent or severe headaches while migraine abortive treatment is suitable for infrequent or non-severe headaches."

As the example shows, the text planner needs to be fine tuned to make the text more idiomatic, but this is a low priority at this time. To produce this text, the system posts the goal to make the patient know the contrast between the two treatments. This goal is achieved through an operator that posts the two subgoals to inform the patient about a superordinate concept of which both treatments are instances (i. e., both are pharmacological treatments) and to make the hearer know their contrasting attributes. The latter subgoal in turn activates an operator which posts the subgoals to inform the patient about each individual difference (i. e., the different purposes, treatment protocols, and indications).

To adapt a text to an individual patient, the system needs a user (patient) model. We do not anticipate implementing a runnable user model, or even an overlay model, but will settle for a *global description* (Ohlsson, in press). Relevant global descriptors include age, current therapy, educational background, health state (other diseases, fitness, pregnancy, etc.), past treatment attempts, and gender. Our system will access the user model in the same way as the knowledge base:

²The speech acts are only primitive relative to the text planner. Each speech act generates a complex description of the desired utterance which is passed to the FUF language generator (Elhadad & Robin, 1992) for translation into English.

through constraints on the operators. For example, an operator to inform the patient that pregnancy is a counterindicator for drug X might include a constraint that the patient is female. We anticipate that the patient model will consist mainly of information gathered by the physician and to a lesser extent of information gathered during the patient-system interaction. The current system does not have a user model.

The query analyzer. The top-level goal in a text plan derives from the user's question. The current system has a parser which can accept three types of user requests typed in from the keyboard: (a) describe X, (b) describe property Y of X, and (c) compare X and Y, where X and Y can be either treatments or drugs. Due to the difficulties of parsing open ended keyboard input and the need for robust performance in the medical clinic, the finished system will use other technologies for accepting user queries. The user will be able to select questions from a menu. In addition, if he or she wants the system to clarify its response, the user can highlight the problematic portion of the text and receive a context-sensitive menu of possible follow-up questions. See Moore and Swartout (1990) for a description of such an interface.

Problems of Employment

We envision placing our completed system in a neurology clinic where migraine patients can interact with it as desired. Interaction with the system is not intended to *replace* visits with the physician, but to help the patient make better use of the limited time with the physician. However, the employment of advanced technologies in real life situations is a non-trivial endeavor. As we have argued elsewhere (Ohlsson, 1991), most of the problems involved in the design and use of instructional technologies are not technological. We have so far identified

three groups of non-technological problems which we need to address with respect to the migraine tutor.

Linguistic problems. Although we have a technology for generating English text on-line, this technology does not tell us which text we ought to generate. To produce idiomatic, comprehensible, and non-redundant text, our system must be sensitive to at least some of the factors that shape people's utterances. One such factor is that people adapt their formulations to what has already been said in previous parts of the dialogue. For example, an object like a drug can be referred to as "it", if it has been recently mentioned, but not otherwise. As a more complex example, if a patient first asks a doctor about the side effects of Inderal (a migraine drug) and then later asks him or her to describe Elavil (an alternative drug), the doctor is likely to respond to the second request by contrasting Elavil with Inderal. Our system can adapt its text plans to what has been said earlier by including tests on the stored dialogue history among the constraints associated with the operators. However, we do not yet know *how much* of the previous dialogue to take into account. We are currently collecting data on human dialogues in order to categorize such backward references and to determine how far back into a dialogue they extend.

Psychological problems. Students frequently distort the content of science instruction by incorporating what they are taught into their prior misconceptions about the relevant topic (Confrey, 1990). Similarly, migraine patients are likely to have prior beliefs about human physiology and medicine that affect what they will or can learn from our system (Arnaudin & Mintzes, 1985; Furnham, 1988, Chap. 5). Oversimplified causal reasoning is one potential source of difficulty (Einhorn & Hogart, 1986). Migraine attacks are triggered probabilistically by a wide range of factors (chocolate, red wine, stress, etc.) and a patient can be sensitive to more than one factor.

Furthermore, these factors can be additive, so that red wine *and* chocolate taken together might trigger an attack even when either factor by itself would not. Correct understanding of these facts might be hindered by a common tendency to think in terms of single, deterministic causes (Konold, 1989). For example, one physician reported the case of a patient who discovered that red wine triggered her migraines, but who later concluded that she was mistaken on the basis of a *single instance* of a glass of wine that did not cause a headache. We are currently planning a series of studies of people's conception of medical causality and its effect on learning from medical explanations.

Socio-cultural problems. People who work in a medical clinic typically perceive themselves as engaged in the rational enterprise of fixing the ailments of the visiting patients. However, a clinic is also a social system with its own mores and customs. If our migraine tutor is to make a constructive contribution to the life of the clinic, it has to be designed with this system in mind. For example, it is not *a priori* obvious which types of information patients typically request of physicians, nor which kinds of explanations physicians usually give to patients. We began this project with the notion that patients are always interested in the physiological mechanism of their disease and that doctors spend at least some of their time explaining disease mechanisms, but we no longer believe this. Patients ask mainly for instrumental information, e. g., information about headache triggers, and doctors report to us that they rarely volunteer information about physiological mechanisms. Our colleague Diana Forsythe at the Intelligent Systems Laboratory is currently conducting ethnographic research in four medical settings in order to study these and related issues (Forsythe, 1992).

Conclusions

Our belief in the viability and general applicability of our approach to on-line generation of explanations is considerably strengthened by the fact that it could be transferred from the domain of programming style to the rather different domain of migraine treatments. Thus, we now have a general technology for providing a non-expert with an intelligent interface to expert knowledge, as long as that knowledge is encoded in a computer knowledge base. The implications of such a technology obviously reach beyond our immediate objective of improving patient education.

However, the employment of this or any other instructional technology in situations in real-life situations is a difficult enterprise. The problems of employment are not themselves technological in nature. Linguistic conventions, the psychology of learning, and the structure of the social system in which the technology is to be employed must all be considered.

As our society becomes more knowledge-driven, non-experts will increasingly find that they must acquire at least a rudimentary familiarity with some expert knowledge base in order to stay in control of their own lives. People will more often be learning in real life situations which are not conceptualized as instructional. We believe that the field of applied Cognitive Science would benefit from studies of how to design instructional systems for a variety of such situations.

References

- Arnaudin, M. W., & Mintzes, J. J. (1985). Students' alternative conceptions of the human circulatory system: A cross-age study. *Science Education*, 69, 721-733.
- Buchanan, B., Moore, J., Forsythe, D., Banks, G., & Ohlsson, S. (1992, February). *Explanation in the clini-*

- cal setting (Technical Report). Pittsburgh, PA: University of Pittsburgh.
- Buchanan, B., & Shortcliffe, E., (Eds.), (1984). *Rule-based expert systems*. Reading, MA: Addison-Wesley.
- Confrey, J. (1990). A review of the research on student conceptions in mathematics, science, and programming. In C. Cazden, (Ed.), *Review of Research in Education* (pp. 3-56). Washington, DC: American Educational Research Association.
- Elhadad, M., & Robin, J. (1992). Controlling content realization with functional unification grammars. In R. Dale, E. Hovy, D. Rosver, & O. Stock, (Eds.), *Proceedings of the Sixth Workshop on Natural Language Generation*. Berlin, Germany: Springer-Verlag.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, 99, 3-19.
- Eraker, S. A., Kirscht, J. P., & Becker, M. H. (1984). Understanding and improving patient compliance. *Annals of Internal Medicine*, 100, 258-268.
- Forsythe, D. E. (1992, February). *Using ethnographic studies to build a working system: A patient education system in migraine* (Technical Report). Pittsburgh, PA: University of Pittsburgh.
- Furnham, A. (1988). *Lay theories: Everyday understanding of problems in the social sciences*. Oxford, UK: Pergamon.
- Konold, C. (1989). Informal conceptions of probability. *Cognition and Instruction*, 6, 59-98.
- MacGregor, R. (1988). A deductive pattern matcher. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, August, St Paul, Minnesota.
- Moore, J. D. (1989). A reactive approach to explanation in expert and advice-giving systems. Ph. D. dissertation. University of California at Los Angeles, Los Angeles, CA.
- Moore J. D., & Paris, C. L. (1989). Planning text for advisory dialogues. *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics* (pp. 203-211), 26-29 June, Vancouver, British Columbia, Canada.
- Moore, J. D., & Swartout, W. R. (1988). *Explanation in expert systems: A survey* (Technical Report No. ISI/RR-88-228). Marina del Rey, CA: Information Sciences Institute.
- Moore, J. D., & Swartout, W. R. (1989). A reactive approach to explanation. *Proceedings of the 11th International Joint Conference on Artificial Intelligence* (pp. 1504-1510) August 20-25, Detroit, Michigan, USA.
- Moore, J. D., & Swartout, William R. (1990). Pointing: A way toward explanation dialogue. *Proceedings of the National Conference on Artificial Intelligence* (pp. 457-464), July 29 - August 3, Boston, MA.
- Ohlsson, S. (in press). The impact of cognitive theory on the practice of courseware authoring. In R. Lewis, (Ed.), *Authoring environments for computer-based courseware*. Berlin, Germany: Springer-Verlag.
- Ohlsson, S. (1991). System hacking meets learning theory: Reflections on the goals and standards of research in artificial intelligence and education. *Journal of Artificial Intelligence in Education*, 2, 5-18.
- Raskin, N. H. (1988). *Headache* (2nd ed). New York, NY: Churchill Livingstone.
- Reason, J. (1990). *Human error*. Cambridge, MA: Cambridge University Press.
- Stewart, W., Lipton, R., Celentano, D., & Reed, M. (1992). Prevalence of migraine headache in the United States. *Journal of the American Medical Association*, 267, 64-69.
- Tapia, C. P. (1991, September). *MEMO: A migraine therapy explanation module* (Technical Report). Pittsburgh, PA: University of Pittsburgh.
- Wilkins, D. E. (1988). *Practical planning: Extending the classical AI planning paradigm*. San Mateo, CA: Kauffmann.