An intelligent interactive system for delivering individualized information to patients

Bruce G. Buchanan [a,*], Johanna D. Moore [a], Diana E. Forsythe [a], Giuseppe Carenini [b], Stellan Ohlsson [c], Gordon Banks [d]

[a] *Department of Computer Science, Mineral Industries Bldg. 206, University of Pittsburgh, Pittsburgh, PA 15260, USA*
[b] *Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA 15260, USA*
[c] *Learning Research and Development Center, University of Pittsburgh, Pittsburgh, PA 15260, USA*
[d] *Department of Neurology, University of Pittsburgh, Pittsburgh, PA 15260, USA*

Abstract

This paper is a report on the first phase of a long-term, interdisciplinary project whose goal is to increase the overall effectiveness of physicians' time, and thus the quality of health care, by improving the information exchange between physicians and patients in clinical settings. We are focusing on patients with long-term and chronic conditions, initially on migraine patients, who require periodic interaction with their physicians for effective management of their condition. We are using medical informatics to focus on the information needs of patients, as well as of physicians, and to address problems of information exchange. This requires understanding patients' concerns to design an appropriate system, and using state-of-the-art artificial intelligence techniques to build an interactive explanation system. In contrast to many other knowledge-based systems, our system's design is based on empirical data on actual information needs. We used ethnographic techniques to observe explanations actually given in clinic settings, and to conduct interviews with migraine sufferers and physicians. Our system has an extensive knowledge base that contains both general medical terminology and specific knowledge about migraine, such as common trigger factors and symptoms of migraine, the common therapies, and the most common effects and side effects of those therapies. The system consists of two main components: (a) an interactive history-taking module that collects information from patients prior to each visit, builds a patient model, and summarizes the patients' status for their physicians; and (b) an intelligent explanation module that produces an interactive information sheet containing explanations in everyday language that are tailored to individual

---

* Corresponding author. Fax (+1-412) 624-9189.

patients, and responds intelligently to follow-up questions about topics covered in the information sheet.

A major goal of medical informatics is to improve the quality of health care by managing information that is relevant to clinical decisions. In contrast to most medical informatics systems, the system described here is designed to facilitate the exchange of information between patients and physicians and to provide patients with information they do not seek from their physicians or do not fully comprehend when it is provided. We are focusing on patients with long-term and chronic disorders, initially on patients with migraine headaches [1], who require periodic interaction with their physicians for effective management of their condition. Also in contrast to many other medical computer systems, the design of the system described here is based on empirical data of actual information needs. We are using ethnographic observation of explanations actually given in clinical settings and interviews with migraine sufferers and physicians.

Successful information exchange between physicians and patients is important for many reasons. Patients' attitudes about their ailments and their therapies influence a number of factors related to clinical success. For example, patient satisfaction has been found to correlate with compliance [20], and communication between doctor and patient is an important factor in patient satisfaction [3,10,28]. In the treatment of diabetes, improved information exchange [23,24] has been shown to increase patients' ability to take more responsibility for their own care. Given the current realities of the health care system in the United States, increasing the amount of time that physicians spend communicating with patients is not a viable option. We need, therefore, to consider other ways of improving the information exchange between doctors and patients.

1. Introduction: towards improved information exchange

One means of facilitating information exchange between patients and physicians in the near future is to use advanced computer technology to make the reading of informative materials more like face-to-face communication. We have implemented a prototype interactive information system for migraine patients who, like other patients with long-term conditions, require periodic interaction with their physicians for effective management of their condition. While other medical informatics programs have been designed for use by health-care providers, and

---

[1] The term 'migraine' is used to refer to the general disease or condition causing migraine headaches. As with current usage, the plural form, 'migraines', is used to refer to instances of migraine headaches.
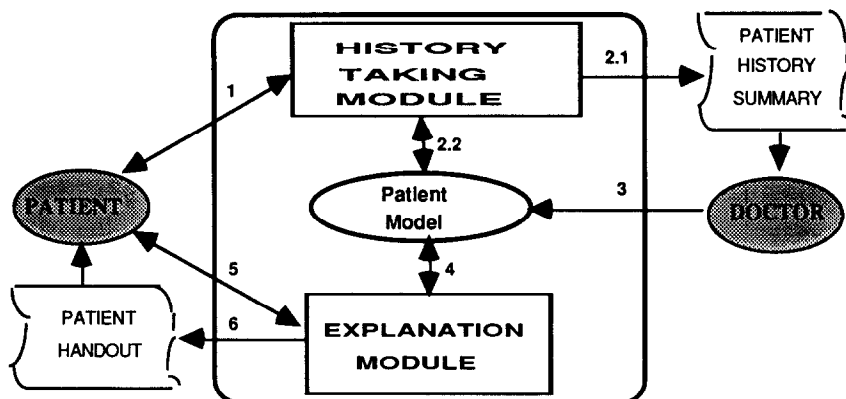
Fig. 1. System Architecture: Information flow between patient and doctor. Information flow is shown by the arrows. Numbered arrows involve the computer system directly. A patient-user interacts with the history-taker (#1), and the record of the interaction is summarized for the doctor (#2.1) and recorded in the system's model of this patient (#2.2). The diagnosis and prescription(s) are entered into the system's patient model by the physician or nurse (#3). The explanation module accesses the patient model (#4) in order to answer the patient's questions (#5). Tailored information about migraine and the prescribed therapy are provided on the screen and then printed out for the patient to take home (#6).

thus require time from them, the system described below is designed for use by patients and is intended to save health-care providers' time.

1.1. System overview

As shown in Fig. 1, the prototype system consists of two major components: the history-taking component and the explanation component, plus knowledge structures for medical knowledge and a patient model. [2]

The history-taking module collects information from patients at the time of their initial and subsequent visits to the Neurology Clinic, and summarizes this information for the physician. This component serves two important purposes in our system. First, it gathers information about the patient's medical history, allergies, past medications, and experiences since the last visit, and translates it into the patient model, which can be used by the explanation component to tailor its explanations to individual patients. Second, it produces a printed summary of the patient's history, which is presented to the physician before the consultation with the patient and is added to the patient's chart after the visit.

---

[2] All parts of the program are written in Lucid Common Lisp. At the present time the history taker is running on a different computer (an Apple LS-III) than the other parts (which run on a DEC station 5000), but all parts will be integrated on a single platform in the future. The temporary separation allows field testing of the history-taking program at multiple sites on less expensive machines.

Once a patient's history has been gathered and the patient has been seen by the physician, the patient can interact with the explanation module to receive information relevant to the doctor's diagnosis and prescribed treatment. Note that before the patient uses the explanation module, the physician or other clinical personnel must provide the system with information about the diagnosis and the prescribed therapy. To expedite this process, we built an interactive interface that prompts the physician for the needed information, and provides warnings such as interactions between drugs being prescribed for migraine and other drugs the patient is taking.

The explanation module generates an initial interactive information sheet, as well as subsequent explanations in response to the user's request for additional information. The content and organization used in the information sheet and subsequent explanations are based on three sources of information: (i) results of the ethnographic research described in Section 3, (ii) sample information sheets written by the physician on our research team, and (iii) sample information sheets being given out in local hospitals and clinics. For example, because of ethnographic findings, the information sheet deals with global concerns (e.g., alleviating patients' fears that they are suffering from a life-threatening condition) before going on to instruct patients about how to follow their therapy regimen.

This computer system is designed to increase the effectiveness of physicians' time in managing adult patients with migraines by collecting data from patients and summarizing patient status before each visit, and explaining physicians' instructions to and answering questions by patients before, after, and between office visits. The information is provided to patients and their families by interacting in simple English and tailoring the information to each patient's context and concerns. Our goal is not to replace but to supplement the doctor–patient relationship; we believe that a computer system can amplify the information provided by physicians, nurses, and printed literature, elaborating, clarifying and re-explaining it as many times as needed.

2. Background

2.1. The information gap in doctor–patient discourse

Empirical research on medical discourse (e.g., [12,18,43,50,55]) demonstrates that an information gap exists between physicians and patients. There is substantial evidence that doctors and patients use somewhat different language [29,55]. Therefore, even when physicians are able to take the time to give patients detailed explanations concerning diagnosis or treatment, patients may not understand the terms or concepts used. Furthermore, there is a striking asymmetry in information exchange between physicians and patients: in medical discourse, physicians ask the questions and patients provide the answers [55]. In one study, of a total of 771 questions recorded in 21 medical dialogues, 91% were initiated by physicians while only 9% were initiated by patients [54]. Similar results were found by Frankel, who

found "a dispreference for patient-initiated questions in physician–patient encounters" [18]. When patients do make direct information requests of their doctors, the answers they receive tend to be brief. Wallen et al. found that physicians spent less than 1% of total talking time in providing explanation to patients ([52], quoted in [[55], p. 108]). In short, interaction between physicians and patients tends to be structured in such a way that information flows from patient to doctor but not vice versa.

This asymmetry reduces the quality and quantity of information exchanged between physicians and patients, which may contribute to reduced compliance [20,28]. Indeed, of the various determinants of patient satisfaction, communication between physician and patient is one of the few factors strong enough to be predictive of compliance [3,10,28]. The restricted flow of information from doctor to patient in normal doctor–patient discourse suggests that this contact alone may not be sufficient to satisfy all of the information needs that patients experience. For this reason, we believe that there is a role for an intelligent interactive system that provides patients with clear and extensive explanation of their diagnosis and treatment plans, and that can respond to patient queries.

2.2. Explanation in medical informatics

Much of the research on explanation in medical informatics has been done in the context of systems that were designed to be used by health care providers [8,49]. The term 'explanation' typically refers to a program's justification of a recommended action, i.e., automatically providing the set of reasons and inferences underneath a conclusion suggested by a program. Explaining material to patients places very different demands on a computer system. First, patients have different types of concerns about their disease than health care providers. For example, health care personnel are primarily concerned with checking how the system arrived at its diagnosis or therapy recommendation [45]. In contrast, patients are likely to be more concerned with understanding how the diagnosis will affect their lifestyle, what their treatments options are, and how to follow the therapy prescribed. Part of our ethnographic work has been directed toward determining the information needs of migraine patients (see below). We have concentrated on representing the knowledge needed to support responses to some of the types of questions we have identified, and have devised generation strategies to produce natural language explanations from these knowledge sources.

Second, we expect patients' concerns to vary widely. Therefore, the system must be able to tailor its explanations to the concerns of individual patients. This is consistent with the claims of Jimison et al. [22] who argued that explanations of therapy decisions must be tailored to patients' preferences regarding the quality of life associated with potential outcomes of therapy. Third, while clinician-users typically understand the medical concepts used in explanations, we cannot assume that this will be true of all patients who use the system. A system intended to produce explanations for patients must be able to either explain the medical terminology and concepts it uses, or else avoid such terminology.

Finally, because it is impossible to have an accurate model of what each patient knows and does not know, it is crucial that a patient-education system be able to respond to follow-up questions and requests for elaboration or clarification of previous explanatory material. Thus, the system should have alternative ways of responding to a given type of question.

The previous work that is most relevant to ours is that of Jimison et al. [22], who used a patient model to tailor a generic Bayesian decision network to an individual patient's medical history and preferences (sedentary vs. active lifestyle). This tailored network is then used to individualize patient education materials explaining the therapy decision in terms of the patient's condition, medical history and lifestyle preferences. Because the number of features that they use for tailoring is relatively small, and the explanations are of limited scope, Jimison et al. were able to handcraft separate cards for each combination of therapy explanation options. Their program then uses the tailored decision network to select the correct card for inclusion in the patient education materials. However, Jimison et al. do not provide facilities that allow the user to ask follow-up questions about the system's prior explanations.

As we discuss in detail in Section 4, our system uses sophisticated techniques devised by computational linguists to automatically produce texts that are sensitive to the patient model as well as the content of previous explanations that have been presented to the user. We draw on results from computational linguistics showing how to tailor explanations to a user's knowledge and goals (e.g., [9,33,42,51]). In addition, we have extended work by Moore and Paris [38,39], who developed a computational text planner that constructs a representation of its explanations, and thereby enables a system to respond to follow-up questions and to provide elaborations of its previous explanations.

2.3. Migraine headache

The domain of this project is migraine therapy for adult patients. Migraine affects approximately 20% of the population [27], afflicting about three times as many women as men ([46], p. 93). Although rarely fatal, it causes considerable discomfort and loss of work time, reducing the quality of life of migraine sufferers. Saper et al. estimate that the cost to employers of employees' migraines is $5256 to $6864 per year for each male and $3168 to $3600 per year for each female ([46], p. 1). From the standpoint of physicians, this condition is sometimes difficult to diagnose and can be time-consuming and awkward to treat. Sacks notes that some forms of migraine headache "present formidable challenges in differential diagnosis; indeed, there is probably no field in medicine so strewn with the debris of misdiagnosis and mistreatment, and of well-intentioned but wholly mistaken medical and surgical interventions" ([45], p. 85).

In the diagnosis and treatment of migraine headaches, effective information exchange between physician and patient is critical. Diagnosis of headache depends primarily on a thorough history of symptoms. There are no laboratory tests which

differentiate migraine from cluster headache, muscle tension headache, sinus headache, or the somatic complaints of depression. Since the physical examination is usually normal in headache patients, one reason why diagnosis of migraine is difficult may be that many primary care practitioners lack the time to elicit the details of the patient history that would allow proper diagnostic classification. Often patients are referred to a neurologist only when they fail to respond to therapy or become analgesic abusers. Effective treatment of migraine is complicated by the fact that about half of the patients do not get sufficient relief from the regimen initially prescribed. Finding the proper medicine and dosage may require a long series of visits to the clinic. Those who become impatient may bounce from doctor to doctor without finding a successful program. With migraine patients, then, information exchange is a major issue: patients must be motivated to return for further visits even when their treatment to date appears to have been unsuccessful.

3. Basing system design on ethnographic data

In order to be useful, an intelligent system must offer information that users want and present it in ways that make sense in terms of the users' own experience. Although these conditions may seem to be "just common sense", they are not in fact commonly followed. Few computer programs for medicine are based on a systematic investigation of the information needs and perspectives of prospective users. Instead, programs have frequently been constructed on the basis of designers' tacit assumptions concerning the kinds of information needed by system users [15], and systems present this material in terms that make sense to the system-builders but not necessarily to users. Both features of this approach may contribute to the well-known problem of low user acceptance [13]. As described below, we have based the design of the system on ethnographic field work. We have also investigated cognitive obstacles to a person's assimilation of technical information through cognitive studies (see [1]).

To design our system, we sought detailed knowledge of the nature and scope of the information needs experienced by clinicians and migraine sufferers. Then, in order to understand how to respond to these needs most effectively, we tried to identify the explanatory material that best meets the needs of particular types of users. To gather this information, we used ethnographic field research, a set of research methods developed by anthropologists to gather descriptive qualitative information in complex real-world settings [53]. These methods have proven to be useful in the design and evaluation of knowledge-based systems [11,16,30,40]. Fieldwork for the present project involved a combination of two ethnographic techniques, participant observation and interviewing.

First, to investigate the information needs of migraine patients and to identify explanations that would satisfy these information needs, we have observed interaction between neurologists and their patients. To date we have carried out a total of

70 hours of field observations of doctor–patient visits in 7 neurology-related sites in the Pittsburgh area. We have observed 78 neurology patients, of whom 29 suffered from headaches. The providers observed include 7 attending neurologists, 6 residents, 5 medical students and 4 other health care providers. Of the doctor–patient encounters, 36 were tape recorded. The tapes have been transcribed and are being subjected to content analysis.

Second, in order to supplement data gained from observations and to investigate information needs that patients might not be expressing to their physicians, we undertook a series of semi-directed interviews with migraine sufferers outside the clinical setting. These interviews averaged a little over an hour in length, and have provided detailed data on matters that seldom arose during our observations of patient visits. To date we have interviewed 4 men and 8 women, producing 12 1/2 hours of tape-recorded material. Again, we are undertaking content analysis of the transcribed tapes.

The transcribed texts produced from observations and interviews are a rich source of data. Content analysis of these transcripts identifies information needs expressed by patients and physicians, explanations offered by physicians, and data about how these explanations are received. It also provides insight into the types of knowledge physicians and patients have about migraine and its treatment, as well as what each takes for granted in discussing the topic. Following an analysis of individual texts, we compare them, seeking patterns and common features within and between them. This process provides insights for the content and wording of explanatory material to be offered by the system.

Several design considerations have emerged from analysis of the field data, including the five below. These are especially valuable in considering what problems to address and what kind of system to build, but also in the details of how to word questions and answers.

(1) Particularly in the early stages of the project, field observations led us to become aware of and to reevaluate some of our tacit assumptions about the nature of the design task. As discussed in [14], the fieldwork revealed that some of these initial design assumptions were inappropriate. For example, we expected (and neurologists told us) that the explanatory material most required by migraine sufferers would be textbook-type information on the physiology of migraine and the nature and side-effects of the drugs used to treat it. However, this is not what migraine sufferers describe as their most important information need (see below). To take a second example, we expected (and neurologists confirmed) that migraine sufferers would know relatively little about their condition. On the basis of this assumption, we initially expected to build a conventional tutoring system. However, this assumption also turned out to be false. Our interviews show that migraine patients typically know a great deal about migraine — but have a different type of knowledge than neurologists. To draw upon a knowledge typology we proposed in [21], through experience migraine sufferers accumulate informal, specific knowledge about their condition, which contrasts with the formal, universal knowledge in which physicians are trained. Thus, our educational task is less to provide information to those who completely lack it, than to support the exchange of

information between physicians and migraine sufferers. Of particular importance is helping individual patients to translate formal medical knowledge about migraine and its treatment into information that they can apply in their own lives, a factor that seems to promote patient compliance [21].

(2) People with migraines often do not receive the information they need; the ethnographic data provide considerable information about the nature of these needs. The transcript analysis underway will produce descriptive statistics on all information-seeking messages expressed by either patients or physicians during the patient visits observed. To provide more immediate design input, we have compiled a list of 166 information needs that migraine sufferers have expressed during patient visits or interviews. These questions are organized in the following categories: (i) What is a migraine headache, (ii) General causes of migraines, (iii) How lifestyle can affect migraine, (iv) Triggers of migraines, (v) Physiology, experience, and medical implications of migraine, (vi) Drug treatment for migraines, and (vii) Other questions about treatment. (Sample questions are given in Section 6.) Since both fieldwork and analysis continue, the list is still growing. Because spontaneously-expressed information needs are not always verbalized as syntactic questions [17], some of the questions on the list are verbatim, while others are reformulations of utterances that conveyed an explicit desire for information. All are grounded in documented expressions of information-seeking by migraine sufferers. While this list is not exhaustive, it serves as a useful resource to help the project team understand what potential users of our system are likely to want to know.

(3) Migraine sufferers have information needs that they tend not to bring up directly with their physicians. For example, migraineurs often mentioned a fear that their condition may be fatal, a reasonable concern given the frightening nature of some migraine symptoms and their perceived similarity to symptoms of brain tumor or stroke. (A similar concern is noted in [19], p. 62). Due to the worrisome nature of this concern, as well as our observation that it often remains unexpressed, we designed our system to present explanatory material addressing such fears to all users of the system.

(4) The fieldwork has also yielded information on the relative urgency of different types of information needs expressed by migraine sufferers. On the basis of both observations and interviews, we have postulated the following hierarchy of patients' concerns (see [14]):

(I)   Is this something major (e.g., a stroke or brain tumor)?
(II)   How does this make sense in terms of what I know about my body and my life (e.g., my eating habits or menstrual cycle)?
(III)   What will I need to do about this? How will it affect my everyday life (e.g., my work schedule or my ability to drive)?
(IV)   Why is this happening to me (allergy? inheritance?)

This list reflects a hierarchy of concerns, with (I) clearly perceived as the most important. At the bottom of this list come concerns such as the following:

(V) What is a migraine headache anyway, and how does my migraine drug work?

The top-level question for migraine patients is, "What will happen to me?"

Since migraine patients often fear that they have a brain tumor or have had a stroke, this is understood as a question of life or death significance [3]. Lower down in the hierarchy are questions about the triggers, treatments, and causes of the condition. Useful responses to these questions may make use of textbook information, but will need to focus upon helping patients to translate this information into implications appropriate to their own particular circumstances. We hypothesize that every migraine sufferer has a similar hierarchy of concerns. We do not yet know how long it takes people to move through the hierarchy or the extent to which individuals vary in this progress. However, we do know that such progress is not necessarily unidirectional. For example, several interviewees commented that while they obviously had survived each migraine attack to date, the pain of a severe attack was such as to lead them to fear that *this* one might kill them. In other words, despite their general knowledge about migraine, they experienced a recurrent specific fear that *this time* that knowledge might not apply. In terms of the hierarchy, such migraine sufferers may be said to return temporarily to Level I with each severe attack. Each step of the above hierarchy is associated with a particular set of information needs. In offering information to migraine patients, we need to provide explanations that meet the information needs they currently experience. We have tried to design this hierarchy into the system to help structure information in ways that are most likely to be heeded and understood.

(5) The doctor–patient dialogues recorded in our transcripts of patient visits contain numerous instances of both initial and follow-up questions. These help to reveal not only migraine sufferers' information needs, but also the language and concepts they use in expressing such needs. For example, although neurologists routinely classify medications for migraines as prophylactic, abortive, or analgesic in nature, migraine patients do not generally use these terms. Knowledge of the language they do use is directly applicable to the design of explanatory material and menus to facilitate query construction by users.

4. System modules: a more detailed view

4.1. The history-taking module

The history-taker is an interactive means of getting information from patients in order to save time for physicians and obtain information about patients for the system. Although physicians may be reluctant to use devices that require extra time, patients have a clear motive to spend additional time to address their own problems. We believe that we can increase the chances of a successful encounter

---

[3] It also has economic significance for our health-care system since it is usual for American patients with headaches to receive one or more expensive imaging procedures during the diagnostic process. Such procedures are often considered unnecessary from a medical perspective but are often done to reassure the patients.

between headache patients and the health care system by providing a computer-based assistant to give patients more time, attention, and information. The idea of using computer technology for this purpose was pioneered by Kleinmutz [26], Slack [47], and others over 20 years ago. Bana et al. [2] developed a computer-based headache interview that patients used successfully and that produced summaries that physicians found useful. We have brought this approach into a modern environment and added some novel features.

The design of our system is modular, and therefore individual physicians have some flexibility to tailor the length and content of the computerized interview. There are currently 26 modules, with 1–5 screens presented in each, each screen requesting related pieces of information. The overall set of questions that can be asked is currently 104; on the average about 50 questions are displayed. (The average is less than the full set because some questions are asked only in the context of specific answers to previous questions.) We envision allowing individual physicians to adapt this interview to their own preferences. However, note that whenever a physician eliminates questions from the history-gathering session, the system loses potential information for the patient model and therefore limits the types of tailoring that can be performed by the explanation module. For example, if the system does not ask about gender and age, the explainer cannot avoid presenting information relevant to women of child-bearing age, such as contraindications during pregnancy.

A mouse-based interface allows patients to use check-boxes for answers. Some questions allow just one box to be checked for mutually exclusive choices, while others allow multiple answers. Fig. 2 shows a simple example. A more complex example is shown in Fig. 3, in which the choices to the right of a drug name are not available until the drug is checked. We use pictures in some questions as a substitute for text and a long list of multiple-choice answers. For example, in asking where the head pain occurs, we show front and side views of a head and ask patients to point and click on those areas where the pain is worst. This is shown in Fig. 4.

A summary of the patients' answers (both the positives and the pertinent negatives) is provided for the physician to examine immediately before and during the patient's visit. Our working hypothesis is that the program can summarize a large number of answers into a form that is easily scanned and that can indicate topics of diagnostic relevance for the physician to pursuemore deeply with the patient. The summary will be placed in the patient's chart, along with a complete record of his or her answers to all questions asked (reformatted for legibility). An example is shown in Fig. 5.

4.2. The explanation module

Because of the richness of human language, expressing an answer or an explanation in natural language is complex. Most computer systems are limited to printing out "canned" text, which is stored in the computer beforehand and retrieved verbatim, or they fill in some blanks of a schema with names (e.g., of the

Fig. 2. A simple screen from the history-taking module asking about headache triggers. Each screen begins with general text meant to provide context. The questions are shown in the body of the screen, with standard navigating boxes at the bottom. The actual screen uses color as well as location to separate these different kinds of information.

patient or of drugs a patient is taking). In contrast, the system we are building generates the text dynamically by reasoning about and recording not only what the system says to the user, but also *why* (i.e., the communicative goal(s) of the explanation) and *how* (i.e., via what explanation strategies the system conveys that information).

There are several advantages to using intelligent text generation in a system such as ours. The system can produce explanations that are tailored to aspects of the current context, including the patient model and the history of the questions and answers that have been exchanged thus far. Our fieldwork indicates that patients are much more interested in information that is directly relevant to them than they are in general information the relevance of which is unclear. In addition, our prior work indicates that users often do not fully understand or accept explanations as they are first given. Therefore to be effective a system must be able

1:27:52 PM

B6

Medicines you are taking or might have taken in the past

One kind of drug that is used to stop attacks after they have started is based on ergotamine which is often called ergot for short. It comes under several different brand names and in different forms. Have you taken any of these?

	When Used Currently Taking	Took in the Past
Ergostat, Ergomar or Wigrettes	☒ ◉	○
Cafergot or Wigraine	☐	
Medihaler ergotamine	☒ ○	○
Bellergal-S	☐	
DHE-45 Dihydroergotamine nasal spray	☐	
I might have taken an ergot, but I don't recognize any of these.	☐	

[<<< Go Back]   [OK >>>]

Fig. 3. A screen from the history-taking module combining "check-all" and "check-one" response choices, where the "check-one" choices (indicated by radio buttons) do not become visible until the corresponding box is checked.

to respond to user's follow-up questions and requests for elaboration [35]. Responding to follow-up questions requires that the system keep a record of the prior communicative goals it has achieved as well as the strategies used to achieve those goals; see [37] for details.

Our system is interactive and allows patients to continue asking questions until they are satisfied. Thus the system can begin the interaction by generating a relatively short explanation (1–2 screens). If patients want more information they may ask further questions, and the system will generate further explanations. In our system, phrases and topics about which the system can answer further questions are indicated using highlighting and color on the screen display. Users initiate queries using a mouse to point at portions of the system's explanations that they would like further information about. The length of the dialogue (and therefore of the printout that the patient will receive) is thus controlled by the patient. Since the patient controls the duration and content of the dialogue, any information beyond that presented in the initial information sheet is directly
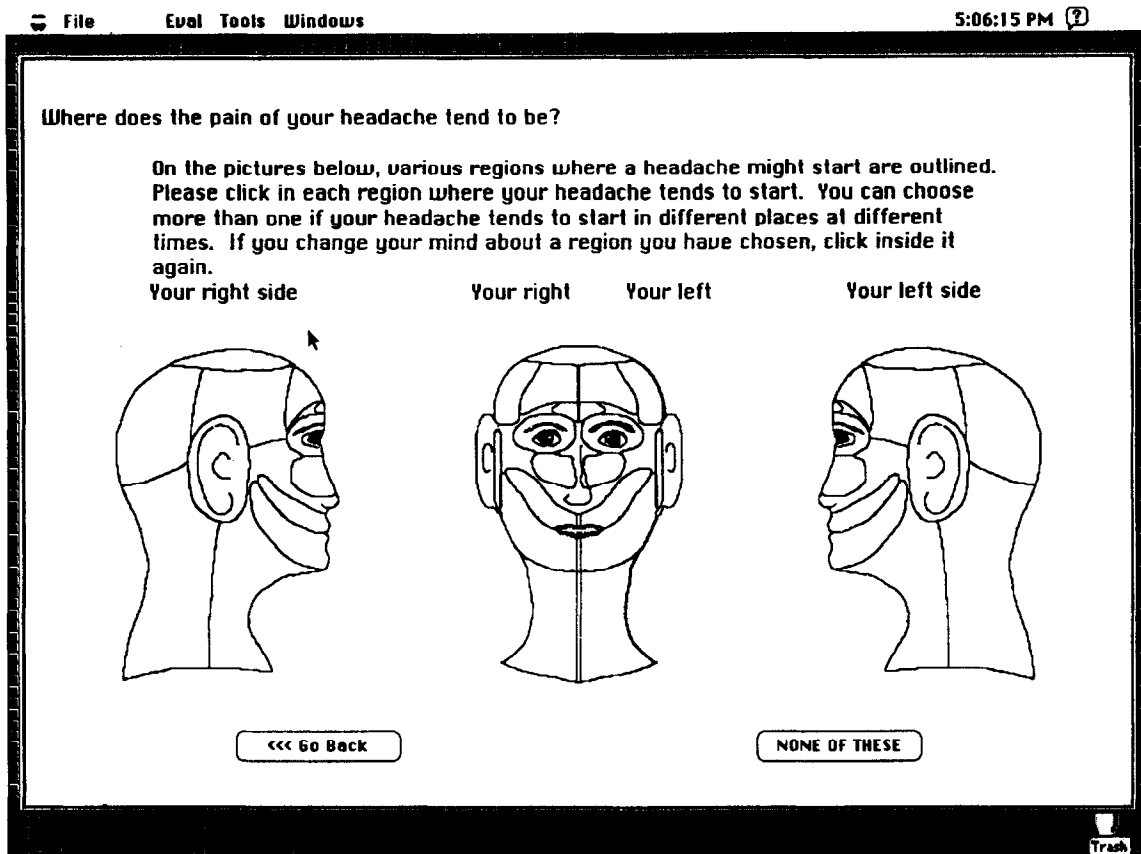
Fig. 4. A screen from the history-taking module in which patients click on appropriate mouse-sensitive regions of a picture. It is easier to see areas on the screen because regions separated by lines in this line drawing are separated by colors on the screen.

relevant to the patient's concerns. This differs from preprinted materials, which inevitably present some information that is not relevant to individual readers. Moreover, because the interface suggests topics that the patient can explore, we believe that it will encourage patients to take a more active role in learning about their disease and therapy. Printed material cannot actively engage the reader in the same way.

The information presented by our system is context-sensitive in two ways. First, by employing the information gathered by the system's history-taking module, both the initial information sheet and subsequently requested explanations can be tailored to features and concerns of individual patients. This allows our system to generate information that is directly relevant to individual patients, a feature that our fieldwork indicates is important. Second, by employing techniques for text planning which record the internal structure of the explanations the system

---

DEMOGRAPHICS: 36 year-old female.

CHIEF COMPLAINT: The patient complains of headaches. They are sharp and accompanied by nausea.

DESCRIPTION OF WORST HEADACHE:
* HISTORY: The headaches began more than 10 years ago. The patient was in her teens.
* PAIN: Very severe, patient is forced to stop normal activities. Aching, sharp, and blinding.
* LOCALIZATION: All over her head.
* ONSET: Headaches usually start when patient is working. Headaches develop fully in less than a half hour but more than 1 minute.
* FREQUENCY: Headaches occur very infrequently.
* DURATION: Without treatment the headaches last on an average some hours, and at most a day or more.
* TERMINATION: The headache terminates gradually and the patient feels completely recovered some hours after the pain is gone.
* TRIGGERING FACTORS: Emotional stress.
* ASSOCIATED SYMPTOMS: Possible fever.
  Always present: dizziness, weakness, and ataxia.
  Sometimes present: fatigue and confusion.
  Does not lose consciousness. No neck stiffness.
* PRODROMATA: Occur several minutes before the headache: always light spots, always nausea, sometimes upset stomach, always dizziness, and sometimes lacrimation on both sides.
* AGGRAVATING FACTORS: None.
* RELIEVING FACTORS: Applying heat, relaxing, and sleeping or trying to sleep.

DESCRIPTION OF SECOND KIND OF HEADACHE:
* PAIN: Moderate, patient has difficulty continuing with normal activities. Aching.
* FREQUENCY: Headaches occur several times per month.
* LOCALIZATION: All over the head.

BACKGROUND HEALTH INFORMATION: The patient is pre-menopausal. Does not smoke. Does not take medication for asthma. Does not take medication for heart trouble.

MEDICATION:
    PRESENT:
    - Occasionally non-prescription pain relievers.
    - Pamelor.
    PAST: none.

FAMILY HISTORY: Brother(s), father, and daughter have or had severe headaches.

Fig. 5. Summary for patient Camphor.

produces [39], the system can interpret and respond to patient questions in the context of the ongoing dialogue [6,34,37].

In implementing the explanation module, we build on previous experience in designing and implementing a similar facility for an advice-giving system in another
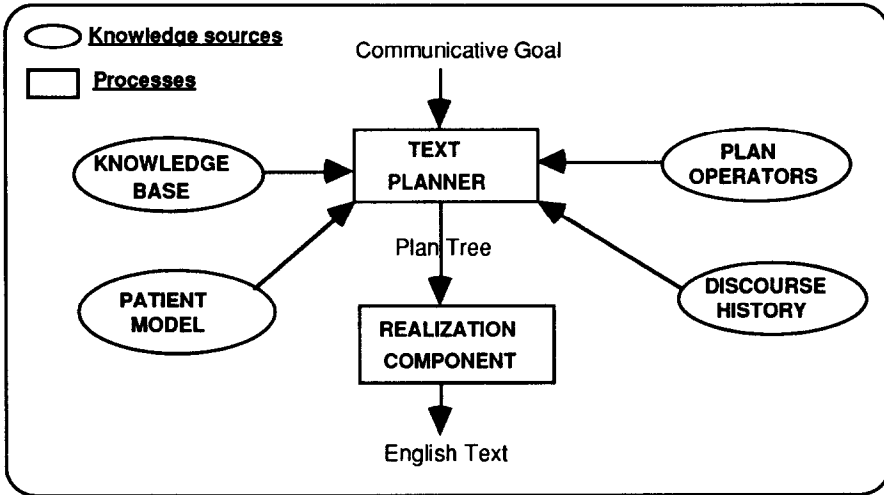
Fig. 6. Architecture of the explanation module.

domain [37]. An overview of the architecture of the explanation component is shown in Fig. 6. Its two major parts are described below.

The text planner. The explanation planning process begins when a *communicative goal* (e.g., "make the hearer believe that the diagnosis is migraine", "make the hearer know about the side effects of Inderal") is posted to the *text planner*. A communicative goal represents the effect(s) that the explanation is intended to have on the patient's knowledge or goals. Posting a goal leads to the retrieval of knowledge units (facts, operators) which are useful in attaining that goal; the attempt to apply those knowledge units may result in the posting of further goals. Planning continues until all communicative goals have been refined to speech acts (e.g., INFORM and RECOMMEND) that can be "executed" to produce the actual English text.

In our system, there are two ways that a communicative goal may be posed to the text planner. First, when the interaction begins, an interaction manager (see next section) posts the goal to generate the information sheet for the current patient-user. Second, after the information sheet has been generated, the user is free to ask further questions about topics covered there, or to ask about other topics given in question menus. The interaction manager processes the user's query and posts a communicative goal to the text planner to produce a response. The text planner decides how to construct natural and informative text, taking into account what the system knows about the particular patient (as indicated in the patient model) and the previous dialogue with that patient (as recorded in the dialogue history).

To produce an explanation that achieves the communicative goal(s), the planner searches its library of explanation operators, looking for candidates that can achieve the goal(s). These operators were derived by studying naturally occurring

```
(define-text-plan-operator
 :NAME EVIDENCE-DIAGNOSIS-M
 :EFFECT (EVIDENCE (DIAGNOSIS ?doctor ?disease ?patient))
 :CONSTRAINTS
      (AND
         (MIGRAINE ?disease)
         (DIRECT-SUBCS MIGRAINE-CHARACTERISTIC-FINDING
                      ?char-findings)
         (HAS-FINDING ?pat-headache-episode ?pat-char-findings)
         (MIGRAINE-CHARACTERISTIC-FINDING ?pat-char-findings)
         (HAS-FINDING ?pat-headache-episode ?pat-other-findings)
         (MIGRAINE-OTHER-FINDING ?pat-other-finding))
 :STEPS
      (((FORALL ?char-findings
               (BEL ?patient (CHAR-FINDING ?disease
          ?char-finding))))
       ((FORALL ?pat-char-findings
               (BEL ?patient
                     (HAS-FINDING ?patient ?pat-char-findings))))
      ((FORALL ?pat-other-finding
          (BEL ?patient (HAS-FINDING ?patient ?pat-other-finding))
          (BEL ?patient (OTHER-FINDING ?disease
                                       ?pat-other-finding))))
       ((BEL ?patient (IS-HEREDITARY ?disease)))))
```

Fig. 7. Sample explanation operator.

explanations from our fieldwork on doctor–patient interactions, books that explain migraine concepts to a lay audience (e.g., [45,48]), explanations constructed by our physician experts, and explanations constructed by the field workers. An example explanation operator is shown in Fig. 7.

Explanation operators integrate multiple sources of knowledge. First, they encode standard ways that communicative goals are achieved by rhetorical means, thus allowing our system to produce natural explanations. For example, the operator in Fig. 7 shows a strategy that gives the patient evidence for the doctor's diagnosis of migraine by posting the following subgoals for the planner to refine:

● informing the patient about characteristic findings associated with this disease; and
● citing the characteristic symptoms reported by the patient (as indicated in the patient model);
● for all other findings associated with migraine that the patient reported, reminding the patient she has this symptom and that it is typical of migraine sufferers; and
● letting the patient know that migraine is often hereditary.

Second, each operator contains applicability constraints that specify the knowledge that must be available if the operator is to be used. These criteria can refer to the system's medical knowledge base, the patient model, or the dialogue history.

For example, the operator shown in Fig. 7 has six constraints. The first constraint checks to see that the diagnosed disease is migraine. The second constraint causes the system to search for the findings that are characteristic of migraine and to store them in a variable called ?char-findings. The third and fourth constraints set the variable ?pat-char-findings to those findings characteristic of migraine that the patient has indicated during the history taking. The fifth and sixth constraints set the variable ?pat-other-findings to those non-characteristic findings that the patient has reported. (The variable ?pat-headache-episode is globally bound to the description representing the current patient's stereotypical headache attack, which has been created in the patient model using information gathered by the history-taking module.)

In general, there may be many strategies capable of achieving a given goal, so the planner employs a set of *selection heuristics* to determine which of the candidate strategies is most appropriate in the current situation. These selection heuristics take into account information about the patient's knowledge and preferences (as recorded in the patient model), and the conversation that has occurred so far (as recorded in the dialogue history). Once a strategy is selected, it may in turn post subgoals for the planner to refine. For example, the operator shown in Fig. 7 posts several subgoals depending on the number of bindings for each of the variables in the FORALL constructs. Planning continues by refining subgoals in this fashion until the entire plan is refined into primitive actions that can be directly executed. In our system, these are speech acts such as INFORM and RECOMMEND. Once a text plan is completed, it is recorded in the dialogue history, and passed to the realization component to generate the English text.

The realization component. This component translates the text plan generated by the planner into English sentences that are printed on the screen for the patient to read and interact with. In the current system, we generate English text using templates. A template may express a single speech act or a sequence of similar speech acts. During the planning process, when a speech act is encountered, a

```
(define-text-plan-operator
   :NAME   diagnosis
   :EFFECT (INFORM ?patient (DIAGNOSIS ?doctor ?disease))
   :PRIMITIVE T
   :CONSTRAINTS NIL
   :STEPS (DIAGNOSIS-TEMP ?doctor ?disease))
```

"Today you were seen by Dr. Washington and diagnosed as suffering from migraine.

Fig. 8. Sample primitive operator and resulting text.

primitive plan operator associates an appropriate template with the text plan node for that speech act (see Fig. 8). If there are many primitive operators whose effect field matches the speech act, the planner chooses the most specific match. For example, if the speech act is (INFORM system patient-1 (has-side-effect MOTRIN VOMITING)), and there are two operators, one whose effect is (INFORM ?system ?patient (has-side-effect ?drug VOMITING)) and one whose effect is (INFORM ?system ?patient-1 (has-side-effect ?drug ?side-effect)), the system will select the former operator. In this way, the most specific template available for expressing the propositional content of a speech act is selected. When the planning process is complete, the realization component traverses the text plan, executing templates as they are encountered. When a template is executed, it returns a sequence of objects that will then be presented on the screen by the interaction manager. It is important to note that the sequence of objects returned by a template is not simply a sequence of text strings. Rather, each object in the set has an associated set of display properties, e.g., font, color, inverse video, and whether or not the displayed object should be mouse sensitive when it is displayed.

In Fig. 8, we show the text that would be generated by executing the statement: (DIAGNOSIS-TEMP ?doctor ?disease). In this example, the parts of the text that are computed dynamically are underlined. Thus, the name of the doctor ("Dr. Washington" is a pseudonym) and the diagnosis (migraine) are computed by the method DIAGNOSIS-TEMP. In this case, these values are computed simply by filling in the values of the variables ?doctor and ?diagnosis. However, there are other methods that perform complex reasoning about the patient model and/or the medical knowledge base to generate some text objects. In Fig. 8, the term "migraine" appears in bold, indicating that it is mouse sensitive. The user can therefore select this term and receive a menu of follow-up questions that can be asked about it.

4.3. The interaction manager

The patient's interaction with both the history-taking and the explanation modules is through a direct manipulation, mouse-based, hypertext-like interface. The patient can ask further questions about topics covered in the information sheet by selecting highlighted text that is on the screen, or can ask questions about other topics by constructing queries with the menu interface. It is neither feasible nor desirable to provide the system with a general capacity to analyze any question expressed in English. First, the natural language understanding technology to do this does not exist. Second, even if we could analyze any question the user posed, there would be many questions that the system would not be able to answer. Thus, we must have a way of constraining the range of questions users ask to those that the system is capable of providing answers for. Providing a restricted query language for this purpose is problematic because users find restricted natural languages difficult to learn and frustrating to use [4,25,44]. In previous work on advice-giving systems [36] and on tutorial systems for young children [41], and in

our early experience with migraine patients running the history-taker, we have found that direct manipulation interfaces that use the mouse as a pointing device are easy to use, even for people who cannot type or who have little knowledge of computers. Moreover, this technique allows a convenient way of circumscribing the set of questions the patient can ask.

Using our interface, the patient constructs questions using the mouse. The patient can build a question in two different ways. In the first case, the patient selects the question type (e.g., Describe) from a main menu and is then guided by the system through dynamically generated menus for the selection of the appropriate arguments (e.g., drug). In the second case, the patient selects a portion of text that the system has presented on the screen and that is indicated as mouse sensitive. In the current interface, mouse-sensitive text strings appear in inverse video on the screen. Whenever the patient clicks on a text segment that is mouse-sensitive, the system presents her with a menu that contains only the question types that are applicable to the selected topic. These menus are not fixed in advance, but are generated in the context of the current patient model and the previous discourse. For example, if the patient clicks on the text string "trigger factors", one of the entries that will appear in the menu is "How to identify what triggers your migraine". If the user selects this entry, the system will put up another menu containing categories of factors the patient can ask about, e.g., dietary factors, factors related to particular activities, factors related to particular natural phenomena, etc. This menu may or may not include an entry for trigger factors related to hormones, depending on the gender of the patient. If the patient selects a question that requires other arguments, the system presents a dynamically generated menu containing the appropriate types of additional arguments. For example, if the patient clicks on "Inderal" on the screen, and then selects "Compare" from the menu of follow-up questions, the system puts up a second menu that includes entries for all of the drugs in the knowledge base. The user scrolls through the entries on this menu and selects the one she wishes to have compared to Inderal. As soon as the patient has constructed a complete question, it is translated into a communicative goal for the text planner to achieve. In this way, the system aids the user in formulating a question, while at the same time constraining the user to ask questions that the system can answer. For example, it constrains users to request comparison of one drug with another drug in the system's knowledge base.

In order to support this direct manipulation interface, where users may point to portions of text on the screen and receive menus of possible follow-up questions, the system must maintain links between the displayed text and the system's knowledge sources. In particular it has to store back-pointers to the part of the text plan that generated the text, and to the entities in the knowledge base that are referred to in the text. Without these links, the interaction manager would not be able to interpret and answer patients' follow-up questions in a context-sensitive fashion, taking into account both the reasons why that segment has been generated as a part of that particular paragraph and what knowledge it expresses.

5. System knowledge sources

5.1. Medical knowledge base

To generate the interactive information sheet and answer further patient questions, our system has an extensive knowledge base (KB) that contains both general medical terminology and specific knowledge about migraine. Examples of the latter include common trigger factors and symptoms for migraine, the common therapies, and the most common effects and side effects of those therapies. In building the part of the KB that represents general medical terminology we exploited an already existing knowledge source: the UMLS Semantic Network [32]. (To construct this portion of the knowledge base we used a semi-automatic process of knowledge acquisition. Interested readers are referred to [7] for details of this process.) To construct the portion of our KB that contains specific knowledge about migraine, we used traditional knowledge acquisition techniques to gather knowledge from our medical expert, medical textbooks, and patient-oriented literature on migraine. We are currently working on links to on-line pharmaceutical databases, which will help ensure that the drug information used by the system is current and complete.

The KB is implemented in Loom, a knowledge representation language [31]. Loom's modeling language is a hybrid consisting of two sublanguages. The *definition* language provides the knowledge engineer with a principled means to represent knowledge about the defining characteristics of domain concepts and relations. Using the definition language, the knowledge engineer specifies the terminological component of a knowledge base (referred to as the *t-box*). The *assertion* language allows the knowledge engineer to specify a set of contingent facts about specific instances (referred to as the *a-box*).

The t-box of our knowledge base contains definitions of general medical concepts such as PHARMACOLOGIC-SUBSTANCE, DISEASE-OR-SYNDROME, FINDING,

```
(defconcept MIGRAINE-EPISODE
    :annotations ((has-min-duration Several-hours)
                  (has-max-duration 3-Days)
                  (has-min-frequency 1-in-Year)
                  (has-max-frequency 4-in-Week)
                  (has-class-average-frequency 3-in-Month))
    :is (:and NON-ORGANIC-HEADACHE-EPISODE
              (:all has-characteristic-finding
                    MIGRAINE-CHARACTERISTIC-FINDING)
              (:all has-other-finding
                    MIGRAINE-EPISODE-OTHER-FINDING)
              (:all has-trigger MIGRAINE-TRIGGER)))
```

Scheme 1.

THERAPEUTIC-PROCEDURE, PATIENT, etc., definitions of more specific concepts about migraine such as MIGRAINE-ANALGESIC-THERAPEUTIC-PROCEDURE, MI-GRAINE-TRIGGER-RELATED-TO-HORMONAL-CHANGES, etc., and definitions of relations between concepts such as has-side-effect, has-therapeutic-suitability, has-characteristic-finding, etc. For example, the definition of the concept MIGRAINE-EPISODE in our t-box is as shown in Scheme 1:

This definition says that the concept MIGRAINE-EPISODE is a NON-ORGANIC-HEADACHE-EPISODE. Thus it inherits all of the attributes associated with NON-ORGANIC-HEADACHE-EPISODE. Moreover, the filler of the attributes has-characteristic-finding, has-other-finding and has-trigger are further specialized. The :annotations keyword is used to specify values for relations that hold only for the concept itself and not for its instances, such as the maximum recorded duration of a migraine episode.

The a-box of our knowledge base contains two conceptually distinct sets of facts. The first set, about drugs, diseases, etc., does not change during routine use of the system. No new facts are added unless explicitly asserted by the knowledge engineer when the system is not in use. An example of the type of facts associated with a particular drug (Nortriptyline) in the static portion of our a-box is as follows:

```
(tellm (Nortriptyline TRICYCLIC-ANTIDEPRESSANT))
(tellm (Nortriptyline PROPHYLACTIC-DRUG-FOR-MIGRAINE))
(tellm (has-chemical-composition Nortriptyline-Hydrochloride)
(tell (:about Nortriptyline
      (:all has-common-side-effect
            NORTRIPTYLINE-COMMON-SIDE-EFFECT)
```

Scheme 2.

The first two facts assert that Nortriptyline is both a TRICYCLIC-ANTIDEPRESSANT and a PROPHYLACTIC-DRUG-FOR-MIGRAINE. The third fact asserts that the chemical composition of Nortriptyline is the chemical substance called Nortriptyline-Hydrochloride, and the last fact asserts a restriction on the possible fillers of the role has-common-side-effect of Nortriptyline.

The second set of facts in our a-box encodes knowledge about particular patients, their disease-episodes, their symptoms, and so forth, in much the same fashion as the facts about drugs are stored. This component of the a-box corresponds to the patient model and is dynamic. New facts are asserted about entities in this set every time a patient interacts with the system.

We chose Loom for our knowledge representation tool because our system needs to reason about its knowledge, and about the structure of its knowledge, in order to generate explanations automatically from the knowledge base. Loom supports such types of reasoning in an integrated environment. Specifically, it

provides inheritance, production rules augmented with terminological classification, and a full first order query language that allows the use of metapredicates, i.e., predicates about the structure of the knowledge. Moreover, Loom infers and maintains a complete and accurate taxonomic lattice of logical subsumption relations between concepts and relations.

5.2. Patient model

In order to adapt explanations to a particular patient, the explanation facility must have access to a stored description of that patient, i.e., the *patient model*. As described above, information about patients' symptoms, past treatments, relevant habits (e.g., a strenuous exercise program), and other current medical treatments is gathered by the history-taking module. Information about the patient's diagnosis and prescribed therapy is provided by the physician or other clinical personnel immediately after the visit with the physician. Both the patient's answers and the physician's inputs are expressed as Loom assertions and are loaded into the a-box of the KB.

The patient model is centered on an instance of the concept PATIENT. A partial, simplified network-like description of the concept PATIENT is shown in Fig. 9. Each time a new patient interacts with the system, new instances of the concepts PATIENT, VISIT, HISTORY, HEADACHE-HISTORY, etc. are created and facts about them are appropriately stored. When a patient comes back for a follow-up visit, the system executes the following steps. First, it loads the model for the current patient in the knowledge base, and it creates and stores a new instance of the concept VISIT. Then, the history-taking module checks whether the facts stored in the patient model still hold and asks the patient about outcomes and possible side effects of the therapies that were previously prescribed to the patient. Finally, after the patient has been seen by the doctor, the patient model is updated if either the diagnosis or the prescription has been changed.

5.3. Discourse history

In order to produce a natural dialogue, the system must generate the current explanation in the context of explanations given earlier in the current interaction. For example, a system that keeps repeating the same message over and over again is likely to try a user's patience. In our system, the discourse history is a record of the text plan that generated the information sheet, the questions the user has asked, and the text plans that led to the system's responses to these questions. Whenever a question is asked, the system checks the dialogue history to determine whether this question has been answered previously. If so, the text planner selects a different strategy for answering it than has been chosen before. Even when the user does not literally ask the same question again, the system's responses must take the dialogue context into account. Computer-generated discourse that does not draw on the previous discourse seems awkward and unnatural, and may even
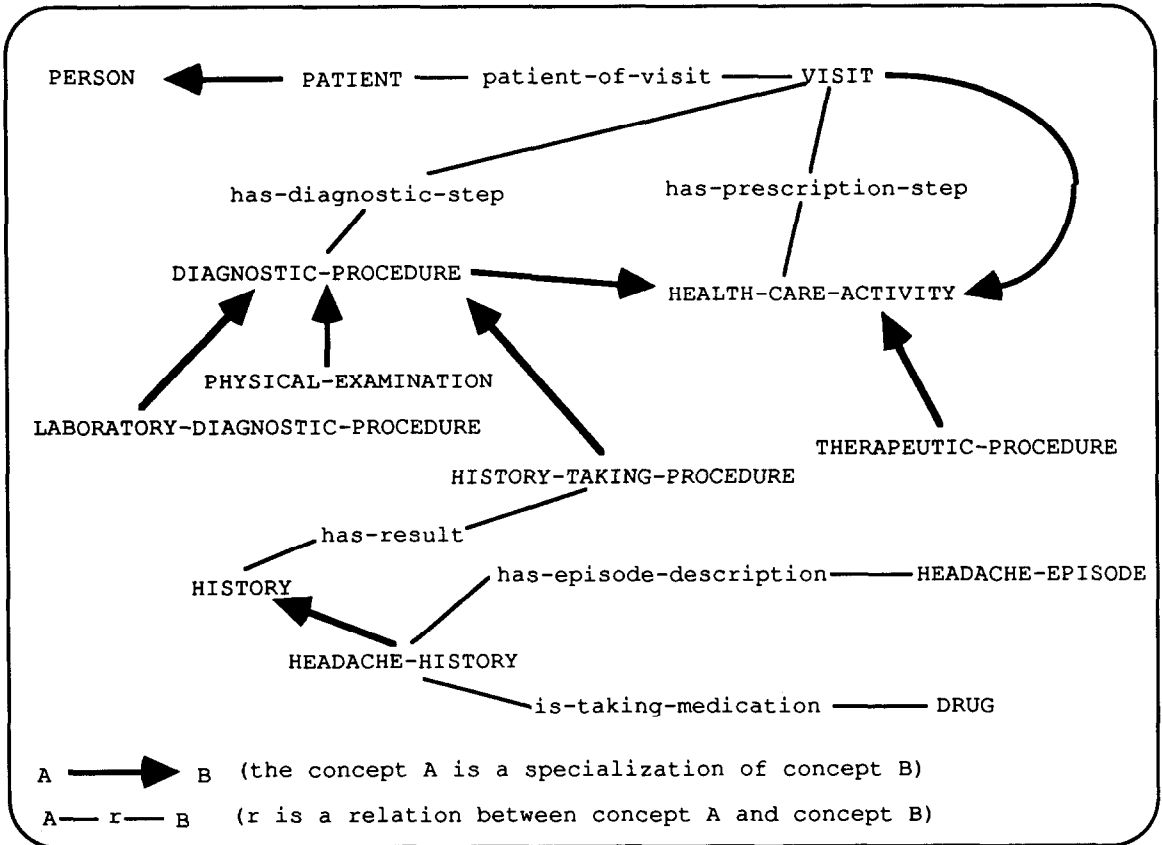
Fig. 9. Partial description of the concept patient.

be incoherent. An example illustrating how our system takes account of the previous discourse is given below.

6. An example session

Here we work through an example based on information about a woman (pseudonym "Camphor") who was seen by a neurologist and whose interview was observed in our field work. The summary of Camphor's history generated by the system is shown in Fig. 5. [4] From the transcript of Camphor's visit with the physician, we also extracted the information that the doctor or other clinical personnel would have input into the system after the visit (see Fig. 10). Fig. 11

---

[4] We ran the history-taker ourselves using data from the transcript of her visit with a neurologist because she did not use the history-taker herself.

```
NORMAL  PHYSICAL  EXAMINATION
PRESCRIPTION:
      -Pamelor 10mg,  each morning
      -Ergostat 2mg,  lozenge
      -Motrin  800mg
```

Fig. 10. Information entered by the doctor or other clinical personnel right after Camphor's visit.

shows the first two paragraphs of the interactive information sheet generated by our system for Camphor. The complete information sheet for this patient appears in Appendix A. Portions of the text that were tailored are underlined in Fig. 11, and portions of the text that are mouse sensitive are depicted in bold. (Note: the

Today you were seen by Dr. Washington and diagnosed as suffering from **migraine**. The **most common symptom of migraine** is a moderate to severe throbbing or pulsating **headache**. **Migraine** patients also frequently experience **visual symptoms, nausea, sensitivity to noise and light,** and **confusion**. You report that your vision is affected, which is a rather common manifestation of migraines. You also report nausea and vomiting with some headaches. A feeling of confusion is also commonly reported by people during migraine attacks. Migraine is also strongly **hereditary**. You report that your father, brother and daughter all have had severe headaches that are similar to yours. Your family history is a further indication that you suffer from **migraine**.

While migraines can produce very intense and painful headaches, they are **very rarely life threatening**. Sometimes patients worry that their migraine symptoms mean that they have a **brain tumor** or have suffered a stroke. Do not worry. You were diagnosed as suffering from **migraine**. Your physical examination was normal, which indicates that more serious causes of headache **such as bleeding or tumors are very unlikely**.

Migraine often gets better as women age, especially after menopause, since hormones play a role in making the attacks more severe.

Fig. 11. Portion of computer-generated information sheet. (Underlining and bold type are added to illustrate points made in the text.)

underlining is for illustrative purposes only. Tailored text does not appear under-
lined when printed on the screen for the patient.) The first paragraph begins by
stating the physician's diagnosis, and provides a general description of migraine.
This is followed by an explanation of how the diagnosis accounts for Camphor's
symptoms. We provide this information in order to assure the patient that the
diagnosis is correct and to alleviate the patient's worries about other, more serious,
or even life-threatening conditions. In this first paragraph, the mouse sensitive
entities are: **migraine, headache, visual symptoms, sensitivity to noise and light**,
and **confusion**.

Suppose that after reading the information sheet, Camphor selects the mouse-
sensitive term **visual symptoms**. In response to her mouse click, the system puts up
a menu containing the following questions. (The questions illustrated here and
below are information needs actually expressed by migraine sufferers in the course
of our fieldwork; see Section 3 above.)

Questions about visual symptoms

What visual symptoms can accompany migraine?

What causes visual symptoms?

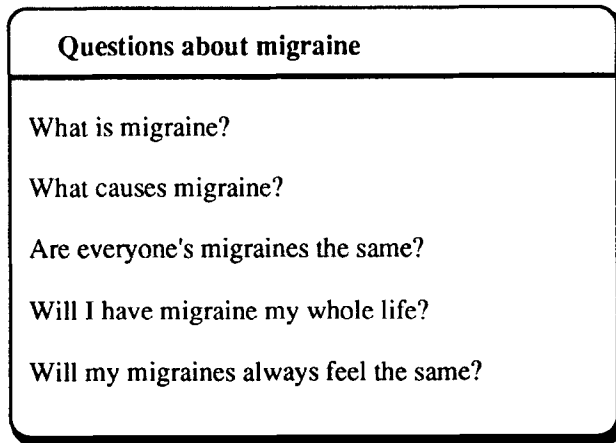Do visual symptoms always occur before a migraine attack?

Scheme 3.

Suppose now that Camphor selects the second option, "What causes visual
symptoms" (Camphor's interest in the causes of her visual symptoms is evident in
the transcript; see Appendix A). At this point, the system answers with the
following text:

```
The visual symptoms of migraine (visual aura), like the ones you
reported, are produced by a temporary disturbance of the back part
of the brain that interprets visual information.
```

Scheme 4.

Note the phrase "like the ones you reported". This phrase is generated because
the patient reported symptoms of this type when responding to the questions of
the history-taker. These responses are encoded into our knowledge representation
language and entered into the patient model, thus allowing the system to personal-
ize its explanation in ways like this. After reading the answer provided by the
system, should Camphor decide to find out more information about migraine in
general, she can click on the mouse-sensitive term **migraine**. In response, the
system would put up a menu containing the questions:

Questions about migraine

What is migraine?

What causes migraine?

Are everyone's migraines the same?

Will I have migraine my whole life?

Will my migraines always feel the same?

Scheme 5.

Suppose Camphor now selects the fourth question, "Will I have migraine my whole life?" The communicative goal that the interaction manager poses to the explanation component for answering such a question is (ALLEVIATE-FEARS ?patient (FOREVER/CHRONIC ?disease)). At this point, the system notices that this communicative goal has already been achieved in the text plan for the

```
(define-text-plan-operator
    :NAME Operator1
    :EFFECT (ALLEVIATE-FEARS ?patient (FOREVER/CHRONIC ?disease))
    :CONSTRAINTS (AND
                    (FEMALE ?patient)
                    (had-past-clinically-relevant-p-or-p
                                    ?headache history ?p-or-p)
                    (NOT (= ?p-or-p Menopause)))
    :STEPS (BEL ?patient (IMPROVE ?disease (AFTER Menopause)))))

(define-text-plan-operator
    :NAME Operator2
    :EFFECT (ALLEVIATE-FEARS ?patient (FOREVER/CHRONIC ?disease))
    :CONSTRAINTS (AND
                    (FEMALE ?patient)
                    (had-past-clinically-relevant-p-or-p
                            ?headache-history ?p-or-p)
                    (NOT (= ?p-or-p Menopause)))
    :STEPS ((BEL ?patient (IMPROVE ?disease (AFTER Menopause)))
            (BEL ?patient (IMPROVE ?disease Aging)))))
```

Fig. 12. Operators that achieve the same communicative goal: (ALLEVIATE-FEARS ?patient (FOREVER/CHRONIC ?disease))

information sheet, and therefore it must choose an alternative strategy for achieving this goal the second time.

In general, in planning the information sheet, whenever multiple operators are applicable for achieving a goal, the text planner selects the least verbose (i.e., the one that posts the fewest subgoals) in order to limit the length of the information sheet. In this case, when constructing the information sheet for Camphor, the text planner selected Operator1 (see Fig. 12) to achieve the communicative goal (ALLEVIATE-FEARS ?patient (FOREVER / CHRONIC ?disease)). The text resulting from executing Operator1 (the last two lines in Fig. 11) is:

```
Migraine often gets better as women age, especially after menopause,
since hormones play a role in making the attacks more severe.
```

Scheme 6.

When answering the follow-up question, the planner must select an alternative strategy and its heuristics direct it to choose a more verbose operator, Operator2 (see Fig. 12). This operator causes the system to generate the following, more elaborated text:

```
Not necessarily.  As I said before, they may stop or alter with
menopause.  In general, for both men and women, migraine attacks
often become less frequent as the sufferer grows older.
```

Scheme 7.

Note that this text contains a reference to a prior explanation, "As I said before...". This reference is generated because the system recognizes that the goal (BEL ?patient (IMPROVE ?disease (AFTER Menopause))) had already been achieved in the plan for the information sheet, and the system contains strategies for determining when to generate such references back to the previous discussion of a topic. This example shows how the context created by a prior explanation affects the generation of subsequent explanations. See [6] for additional examples and more discussion of this aspect of our explanation component.

7. Preliminary evaluation of usability and utility

We have periodically brought in outsiders to use the system and have performed pilot tests of the history-taker in a neurology clinic. Feedback and subsequent improvements gave us confidence that the system was usable by headache patients, so we undertook a formative evaluation of the entire system working together in the context of patients' initial visits to a neurologist, as well as separate evaluations

of the major parts of the system in a laboratory environment. Two of our studies are described below [5]. While we recognize that this is an evaluation of patients' perceptions, and not a study of outcomes, we tentatively conclude that:

- the design of the system is sound — the information we provide and the format we use to present it do not need major revision;
- the system is usable — no one had difficulty using the mouse to select answers to questions or to request more information about a topic; an automated history can be obtained in 30 minutes or less, but more time will be required for more complete histories;
- the implementation of the system is robust and user-friendly — patients with no prior training were able to use the programs for hours at a time with almost no problems; no one crashed the system;
- the information provided is beneficial — some patients expressed unqualified gratitude for the information; no one said their time was wasted.

7.1. Methods

Study #1: Overall system in context. The history-taker and explanation component were designed to work together as an integrated system, but we had not demonstrated, prior to this study, that the individual parts of the system work together smoothly. In study #1, we used the context of a visit to a neurologist to gather information about coupling all parts of the system for use by patients at the time of a visit.

From responses to an advertisement, we selected three persons with symptoms of migraine (all female) to be seen by a neurologist and interact with the computer system. An interviewer scheduled an appointment for the patient to see a neurologist and asked each person to arrive 45 minutes prior to the examination time. The interviewer gathered demographic information from the patient, and then sat with her while she interacted with the history-taker in a room adjacent to the examination room. If questions arose, the interviewer provided just enough information to let the patient proceed. The interviewer recorded any problems the patient seemed to have. After each subject used the history-taker, she was interviewed with a fixed set of questions and then engaged in a dialogue to gather additional comments.

The computer record of the history (the patient model) was transferred to the explanation system in electronic form. A one-page summary of the patient's history was generated from the computer record and given to the neurologist. In this situation, the neurologist did not look at the summary until after examining the patient so we could look for differences between the medical histories taken manually and by the history-taker. After the patient was examined by the neurologist, the interviewer escorted the patient to the explanation system in an adjacent

---

[5] Other studies include a survey of physicians' perceptions of the adequacy of the form and content of the summary from the history-taker, and cognitive science studies of subjects' understanding of materials presented in different forms.

Table 1
Responses on the history-taker follow-up interview

Positive "Yes"/Negative "No" answers

Question	Answer category		
	Yes	No	Don't know/N/A
	f (%)	f (%)	f (%)
1. Did the program ask for more information than doctors you have talked to about your headaches?	10 (0.63)	4 (0.25)	2 (0.13)
2. Would you recommend this program to a friend?	13 (0.81)	0 (0.00)	3 (0.19)
3. Would you like to use the program again if your doctor asked you to?	16 (1.00)	0 (0.00)	0 (0.00)
4. Did you understand most of the questions that were asked in the computer program?	16 (1.00)	0 (0.00)	0 (0.00)
5. In most cases, were you able to find responses that applied to you?	16 (1.00)	0 (0.00)	0 (0.00)
6. Do you feel comfortable about having a computer collect this information?	14 (0.88)	2 (0.13)	0 (0.00)

Negative "Yes"/Positive "No" responses

Question	Answer category		
	Yes	No	Don't know/N/A
	f (%)	f (%)	f (%)
7. Is there anything else that your doctor needs to know about your headaches that the computer did not ask?	6 (0.38)	10 (0.63)	0 (0.00)
8. Did you have any trouble using the computer system?	5 (0.31)	11 (0.69)	0 (0.00)
9. Did you find any of the formats for the questions or answers confusing?	5 (0.31)	11 (0.69)	0 (0.00)

room and sat with her while she used this system. As before, if questions arose, the interviewer provided minimal information to allow the patient to proceed. Afterwards, the interviewer asked additional questions of the patient.

Study #2: History-taker and explanation system. We had performed several formative evaluations of the history-taker and explanation module, with considerable amounts of useful feedback into subsequent redesigns. In study #2, we wished to determine (a) whether there are major barriers to the use of the system, and what they are, and (b) how useful the system is perceived to be.

A broad call for participants resulted in scheduling 13 persons with headache and one or more symptoms of migraine. We followed the same protocol and questions as for study #1, the only difference being that these subjects did not see the neurologist. For simplicity here, we group the responses from the three patients in study #1 with the thirteen in study #2.

7.2. Results: studies #1 and #2

Overall system. The three patients who used both systems in the context of seeing a neurologist were very positive about the system. Subject #1 said that she did not know that her headaches could be treated. She received enough information from the system and the physician to convince her to schedule a follow-up appointment. Subject #2 said that the information she received from the explanation program "changed my life". She had been unsuccessful in the past getting enough information about her headaches to understand how they could be managed, but now felt hope that they could be treated.

Computer use. For the 16 patients, the average time spent using the history-taker was 33 minutes, answering, on average, 44 questions. Some questions may have multiple answers, e.g., prior drugs. The average time spent per question-screen was 45 seconds. We scheduled one hour for patients to interact with the explanation program; several wanted to stay longer.

Information exchange. Table 1 shows subjects' responses to the interview about the history-taker. Looking first at the *information exchange* (questions 1 and 7), we find that 63% of the subjects thought that the history-taker gathered more information than the doctors that they had seen and the same percentage could not think of any information that the program had missed. One subject reported that some of the questions and the lists of choices in the history-taker were very informative in themselves. In particular, she had not thought of dietary triggers and now had some good ideas what to watch for.

Looking next at *usability* (questions 4, 5, 8, and 9), we find that all the subjects thought that they understood the questions. Sixty-nine percent reported no problems in using the system and the same percentage found none of the questions confusing. All subjects reported that the questions supplied answer alternatives that were applicable to themselves.

Table 2
Responses on the explanation follow-up interview

Positive "Yes"/Negative "No" answers Question	Answer category	
	Yes f (%)	No f (%)
1. Did you like using the program?	16 (1.00)	0 (0.00)
2. Did all of the information it presented make sense?	13 (0.81)	3 (0.19)
3. Was the software easy to use?	16 (1.00)	0 (0.00)
4. Did you feel comfortable about using a computer to get this kind of information?	16 (1.00)	0 (0.00)
5. Was the computer itself easy to use?	14 (0.88)	2 (0.13)
6. Did the program tell you anything you didn't already know?	15 (0.94)	1 (0.06)
7. Was the information in the program useful to you?	15 (0.94)	1 (0.06)
8. Do you think this information will help you manage your headaches better?	9 (0.56)	7 (0.44)
9. Did you learn anything that you would not have asked your doctor?	12 (0.75)	4 (0.25)
10. Would you recommend this program to a friend?	16 (1.00)	0 (0.00)
11. All in all, was the computer program worth the trouble it took to use?	16 (1.00)	0 (0.00)
Negative "Yes"/Positive "No" responses Question	Answer category	
	Yes f (%)	No f (%)
12. Was *anything* presented that seemed confusing or inappropriate to you?	8 (0.50)	8 (0.50)

Finally, we were interested in the subjects' *attitudes* towards the computer (question 2, 3, and 6). When asked whether they felt comfortable with a computer collecting the information, 88% answered "yes". An equally high percentage, 81%, would recommend the program to a friend, and all of them would agree to use it again if their doctor asked them to.

Table 2 shows the results from interviews about the explanation module. Looking first at the *information exchange* (questions 6, 7, 8, and 9), 94% of the subjects thought that the program told them something they did not already know (#6), and 75% thought they had learned something that they would not have asked their doctor (#9). When asked whether they thought the information was useful to them, 94% also answered "yes", but when asked whether they thought the information would be useful in managing their condition, only 56% answered "yes"; reasons for the difference will be the subject of further research.

Looking next at *usability*, 81% of the subjects thought that the information presented made sense. However, when asked if any piece of information was confusing or inappropriate, 50% also answered "yes", which we are now following up on. These are mostly wording changes. All the subjects found the software easy to use, and only 13% felt the hardware was less than easy to use.

Finally, looking at the subjects' *attitudes* towards the system (question 1, 4, 10, 11), all of them said that they liked using the program, and that they felt comfortable getting this kind of information from a computer. All would recommend it to a friend and all said they thought it was worth the trouble using it.

7.3. Discussion: studies #1 and #2

Patients were mostly enthusiastic about the quality and quantity of information they received from the program and felt that the information exchange was improved. The history-taker asks relevant questions that subjects feel would not have been asked in a regular session with their doctor and they all thought they learned something from the explanation module. The main criticism was that the program did not have information specifically about the drugs they were taking, a problem we anticipated because the knowledge base is incomplete. However, there is considerable information available to be read: one subject spent 3 hours reading all the information the program could generate, and would have spent more. Many volunteered that the program was more informative than physicians they had seen.

No major problems were encountered by patients. The only major system problem occurred when the campus network failed. We identified some places where the questions or the information presented were confusing to some subjects (which we will change). Patients reported no hesitance to interact with a computer. They also expressed a willingness to use the program again and to recommend it to a friend.

8. Conclusions

On the basis of both published literature and our own observations, we believe that better informed patients will be able to take better care of themselves and have more satisfactory encounters with the health care system. Our work builds on empirical data collected from ethnographic fieldwork with physicians and migraine sufferers, and cognitive science studies. Our observations have shown that the questions patients ask are not limited to, and often do not even include, those that physicians told us they would ask. Similarly, the kinds of explanations we have observed physicians to give do not always successfully address patients' concerns. The design of the system has been adjusted to accommodate this broader range of questions and to provide explanatory material on topics that physicians do not necessarily address.

The explanation system we have built is intended to supplement the time that physicians and nurses spend explaining material to patients with chronic disorders and thus help to meet more of the information needs of patients. The ability to tailor the presentation of information to the needs of individual patients is one of the primary strengths of this approach. Such a system can remind patients of their physicians' instructions and can re-explain in different terms what they have been told in the office, without requiring additional investment of physicians' time. We are not attempting to change the behavior of health-care providers, nor do we require them to interact directly with computers. [6] Rather, we have shifted our

---

[6] We are still experimenting with means for recording patients' diagnoses and prescribed therapies. Eventually, this information will be available on-line nearly everywhere. At the moment, the system can capture it through a brief interaction with the physician, with a nurse, or with a fieldworker.

focus to providing information to the consumers in the health care process — the patients.

We have shown, so far, that ethnography can effectively inform the design of a medical informatics system by providing data on actual information needs and by challenging implicit as well as explicit assumptions of the designers. We have shown that we can address many of the concerns of patients by incorporating them into a knowledge-based system that interactively answers a patient's questions in the context of specific knowledge about that patient. We have also shown that computer-mediated information exchange is possible and favorably perceived by patients, although we have yet to demonstrate that it results in improved outcomes. If computers do prove to be effective in this role, medical informatics will have a broader role in managing information for patients as well as for physicians and nurses.

Acknowledgements

Appendix A

Information sheet for a patient with migraine (fictitious names used.) This is first generated interactively on the computer screen and then printed for the patient to take home. Italic and bold type have been added to illustrate points made in the text

Today you were seen by *Dr. Washington* and diagnosed as suffering from **migraine**. The **most common symptom of migraine** is a moderate to severe throbbing or pulsating **headache**. Migraine patients also frequently experience **visual symptoms, nausea, sensitivity to noise and light,** and **confusion.** *You report that your vision is affected, which is a rather common manifestation of migraines. You also report nausea and vomiting with some headaches. A feeling of confusion is*

also commonly reported by people during migraine attacks. Migraine is also strongly **hereditary**. *You report that your father, brother and daughter all have had severe headaches that are similar to yours. Your family history is a further indication that you suffer from* **migraine**.

While migraines can produce very intense and painful headaches, they are very rarely **life threatening**. Sometimes patients worry that their **migraine** symptoms mean that they have a **brain tumor** or have suffered a stroke. Do not worry. You were diagnosed as suffering from **migraine**. *Your physical examination was normal, which indicates that more serious causes of headache* **such as bleeding or tumors are very unlikely**.

Migraine often gets better as women age, *especially after menopause, since hormones play a role in making the attacks more severe.*

Many everyday things can trigger migraine. These include emotional stress, **hormonal changes associated with menstruation, pregnancy, menopause**, **certain foods**, alcohol, **certain odors**, and **irregularities in your sleeping or eating patterns**. *You did not mention any* **trigger** *factors, except possibly* **stress** *(hard to avoid) and sleeping late on weekends. It would be a good idea to keep track of your sleep patterns to find out if you tend to get more headaches on days when you sleep in (or perhaps don't sleep enough). If that is the case, then you may wish to try to make your sleeping hours the same every day. Also, when you have a* **migraine***, record whether you ate or drank anything unusual in the hours before it began.*

Prescription

Dr. Washington has prescribed the following medications for you:

PAMELOR

You are already taking **pamelor** *(***nortriptylene***) for another medical problem in a dose of 10 mg before bed. It is suggested that you try taking an additional 10 mg capsule each morning.* If you find that this makes you too sleepy, take both capsules at bedtime. The purpose of this medicine is not to **relieve** the pain, but to try to prevent attacks from even getting started. The most common side effects of **pamelor** are dry mouth and drowsiness. If you have dry mouth, sugar-free candies (especially sour ones) may help. If you are drowsy, don't drive or operate machinery. The drowsiness may wear off in a few days once you get used to the increased dose. Be sure to call if you are uncertain whether some new symptom is a side effect of pamelor.

ERGOSTAT

Take the **ergostat** if you start to develop the symptoms of a migraine attack. The minute you feel the migraine starting (*the onset of the visual symptoms you reported*) take *one 2 mg lozenge and let it dissolve under your tongue. In 30 minutes, if you feel the migraine is still coming on, repeat this procedure with another 2 mg lozenge. Thirty minutes later, if the migraine is still coming on, do the same thing again.* Do not take more than 3 lozenges in one day or 6 in one week. The **ergot** may prevent dilation of the head arteries, which helps to **prevent** pain. Some patients experience side effects from **ergostat**. You may get nausea and vomiting. If this persists or is severe call *Dr. Washington* who may prescribe anti-nausea

medicine. If you feel numbness or coldness in your fingertips or notice that your fingertips are white, this may indicate that you have taken too much **ergostat**. Call the doctor immediately if you have these side effects. In addition if you experience any chest pain contact *Dr. Washington* immediately.

MOTRIN

If you get a headache even after taking your ergostat, take 800 mg of Motrin up to 3 times per day. **Motrin** (ibuprofen) helps to **relieve** the pain in most patients. The main side effect is upset stomach. If you can take food without vomiting, this may help you tolerate the Motrin better. If the stomach discomfort persists, check with *Dr. Washington* before taking any more **motrin**. If you have dizziness, nervousness, ringing in the ears, or black stools, report these to your doctor immediately if these effects persist or are severe.

Finding an effective treatment for migraine usually involves a process of trial and error. You and your doctor will need to work together to find the treatment that works for you. This may involve taking medications, making changes in your life style, or both. Fortunately, there are many **treatments available for migraine**, and the **majority of patients are able to find one that is effective**. The first medication the doctor prescribes may not be completely successful. Nevertheless, it is important for you to be patient and take each one as prescribed in order to give it a chance to work. Try not to be discouraged and give up: it is important to come back and try again until an effective treatment is found.

References

[1] P. Albacete and S. Ohlsson, *Patient Education in Migraine: What are the Obstacles to Learning?* (Univ. of Pittsburgh, 1993).

[2] D.S. Bana, A. Leviton, et al., A computer-based headache interview: acceptance by patients and physicians, *Headache* 20 (1980) 85–89.

[3] M.H. Becker, Sociobehavioral determinants of compliance, in: D.L. Sackett and R.B. Haynes, eds., *Compliance with Therapeutic Regimens* (The Johns Hopkins University Press, Baltimore, MD, 1976) 40–50.

[4] B. Black and T.P. Moran, Learning and remembering command names. *Proceedings of the Conference on Human Factors in Computing Systems* (Gaithersburg, MD, 1982).

[5] B.G. Buchanan, J. Moore, et al., Involving patients in health care: using medical informatics for explanation in the clinical setting. *Symposium on Computer Applications in Medical Care* (McGraw-Hill, Baltimore, MD, 1992).

[6] G. Carenini and J.D. Moore, Generating explanations in context, in: W.D. Gray, W.E. Hefley and D. Murray, eds., *Proceedings of the International Workshop on Intelligent User Interfaces* (ACM Press, Orlando, Florida, 1993) 175–182.

[7] G. Carenini and J.D. Moore, Using the UMLS semantic network as a basis for constructing a terminological knowledge base: a preliminary report. *Symposium on Computer Applications in Medical Care* (McGraw-Hill, Washington, DC, 1993).

[8] W.J. Clancey and E.H. Shortliffe, eds., *Readings in Medical Artificial Intelligence: The First Decade* (Addison-Wesley, Reading, MA, 1984).

[9] R. Cohen, M. Jones, et al., Providing responses specific to a user's goals and background, *Int. J. Expert Systems* 2(2) (1989) 135–162.

[10] M.S. Davis, Variation in patients' compliance with doctors' orders: medical practice and doctor–patient interaction, *Psychiatry Med.* 2 (1971) 31–54.

[11] D. Fafchamps, C.Y. Young, et al., Modelling work practice: input to the design of a physician workstation, *15th Symposium on Computer Applications in Medical Care* (McGraw-Hill, Washington, DC, 1991).

[12] S. Fisher, Institutional authority and the structure of discourse, *Discourse Processes* 7 (1984) 201–224.

[13] D.E. Forsythe, Blaming the user in medical informatics: the cultural nature of scientific practice, *Knowledge and Society* 9 (1992) 95–111.

[14] D.E. Forsythe, Using ethnography to build a working system: rethinking basic design assumptions, *Symposium on Computer Applications in Medical Care* (McGraw-Hill, Washington, DC, 1992).

[15] D.E. Forsythe, Using ethnography in the design of an explanation system, *Expert Systems With Applications* (1993) (submitted).

[16] D.E. Forsythe and B.G. Buchanan, Broadening our approach to evaluating medical information systems, *Symposium on Computer Applications in Medical Care* (McGraw-Hill, Washington, DC, 1991).

[17] D.E. Forsythe, B.G. Buchanan, et al., Expanding the concept of medical information: an observational study of physicians' information needs, *Computers Biomed. Res.* 25(2) (1992) 181–200.

[18] R.M. Frankel, Talking in interviews: a dispreference for patient-initiated questions in physician–patient encounters, in: G. Psathas, ed., *Interaction Competence* (University Press of America, Lanham, 1989).

[19] M. Good, Work as a haven from pain, in: M. Good, P. Brodwin, B. Good and A. Kleinman, eds., *Pain as Human Experience: an Anthropological Perspective* (University of California Press, Berkeley, 1992) 49–76.

[20] R.B. Haynes, A critical review of the 'determinants' of patient compliance with therapeutic regimens, in: D.L. Sackett and R.B. Haynes eds., *Compliance with Therapeutic Regimens* (The Johns Hopkins University Press, Baltimore, MD, 1976) 26–39.

[21] L.M. Hunt, B. Jordan, et al., Compliance and the patient's perspective: controlling symptoms in everyday life, *Culture, Med. Psychiatry* 13 (1989) 315–334.

[22] H.B. Jimison, L.M. Fagan, et al., Patient-specific explanation in models of chronic disease, *Artificial Intelligence in Med.* 4 (1992) 191–205.

[23] C. Juge and J.-P. Assal, Enseignement assiste par ordinateur pour patients diabetiques: un moyen d'evaluer les prestations de l'equipe soignante (Computer-assisted instruction for diabetic patients: a means to assess health care teams' effectiveness), *Schweiz. Med. Wschr.* 120(Suppl. 32/I) (1990).

[24] C.F. Juge and J.-P. Assal, Designing computer assisted instruction programs for diabetic patients: how can we make them really useful?, in: M.E. Frisse, ed., *Symposium on Computer Applications in Medical Care* (McGraw-Hill, New York, 1992) 215–219.

[25] M.J. Kelly and A. Champanis, Limited vocabulary natural language dialogue, *Int. J. Man-Machine Studies* 9 (1977) 479–501.

[26] B. Kleinmutz, Medical information processing by computer, in: J.A. Jacquez, ed., *Computer Diagnosis and Diagnostic Methods* (Charles C. Thomas, Springfield, IL, 1972) 45–72.

[27] P. Lane, B. McLellan, et al., Comparative efficacy of chlorpromazine and meperidine with dimenhydrinate in migraine headache, *Ann. Emergency Med.* 18 (1989) 360–365.

[28] P. Ley, Satisfaction, compliance and communication, *Br. J. Clin. Psychol.* 21 (1982) 241–254.

[29] P. Ley and M.S. Spelman, *Communicating with the Patient.* (Warren H. Green, St. Louis, MO, 1967).

[30] H. Lundsgaarde, Evaluating medical expert systems, *Soc. Sci. Med.* 24(10) (1987) 805–819.

[31] R. MacGregor and M.H. Burstein, Using a description classifier to enhance knowledge representation, *IEEE Expert* 6(3) (1991) 41–46.

[32] A.T. McCray and W.T. Hole, The scope and structure of the First Version of the UMLS Semantic Network". *Proceedings of the Symposium on Computer Applications in Medical Care* (IEEE Computer Society Press, Washington, DC, 1990).

[33] K.R. McKeown, Generating goal-oriented explanations, *Int. J. Expert Systems* 1(4) (1988) 377–395.

[34] V. Mittal and C. Paris, *The Use of Context in Explanation Systems.* Working Paper, USC/ISI, Marina del Rey, CA, 1993.

[35] J. Moore and W. Swartout, A reactive approach to explanation, *Proceedings of the 11th International Joint Conference on Artificial Intelligence* (Detroit, MI, 1989).

[36] J. Moore and W. Swartout, Pointing: a way toward explanation dialogue, *Proceedings of the National Conference on Artificial Intelligence* (Boston, MA, 1990) 457–464.

[37] J.D. Moore, *Participating in Explanatory Dialogues: Interpreting and Responding to Questions in Context* (MIT Press, Cambridge, MA, in press).

[38] J.D. Moore and C.L. Paris, Exploiting user feedback to compensate for the unreliability of user models, *User Modeling and User-Adapted Interaction* 2(4) (1992) 331–365.

[39] J.D. Moore and C.L. Paris, Planning text for advisory dialogues: capturing intentional and rhetorical information, *Computational Linguistics* 19(4) (1993) 651–695.

[40] J. Nyce and T. Timpka, Work, knowledge and argument in specialist consultations: incorporating tacit knowledge into system design and development, *Med. Biol. Engineering Computing* 31 (January 1993) HTA16–HTA19.

[41] S. Ohlsson, Towards intelligent tutoring systems that teach knowledge rather than skills: five research questions, in: E. Scanlon and T. O'Shea, eds., *New Directions in Educational Technology* (Springer-Verlag, New York, 1992) 71–96.

[42] C.L. Paris, The role of the user's domain knowledge in generation, *Computational Intelligence* 7(2) (1991) 71–93.

[43] V.L. Patel, D.A. Evans, et al., A cognitive framework for doctor–patient interaction, in: D.A. Evans and V.L. Patel, eds., *Cognitive Science in Medicine* (MIT Press, Cambridge, MA, 1989) 257–312.

[44] S.R. Petrick, On natural-language based computer systems, *IBM J. Res. Dev.* 20 (1976) 314–325.

[45] O. Sacks, *Migraine: Understanding a Common Disorder* (University of California Press, Berkeley, 1985).

[46] J.R. Saper, S.D. Silberstein, et al., *Handbook of Headache Management: a Practical Guide to Diagnosis and Treatment of Head, Neck, and Facial Pain* (Williams and Wilkins, Baltimore, MD, 1993).

[47] W.V. Slack and C.W. Slack, Patient–computer dialogue, *N. Engl. J. Med.* 286 (1972) 1304–1309.

[48] S. Solomon and S. Fraccaro, *The Headache Book. Effective Treatments to Prevent Headaches and Relieve Pain* (Consumer Reports Books, Yonkers, NY, 1991).

[49] P. Szolovits, ed., *Artificial Intelligence in Medicine*. AAAS Selected Symposia Series (Westview Press, Boulder, CO, 1982).

[50] A.D. Todd, The prescription of contraception: negotiations between doctors and patients, *Discourse Processes* 7 (1984) 171–200.

[51] P. van Beek, A model for generating better explanations, *25th Annual Meeting of the Association for Computational Linguistics* (Stanford, CA, 1987).

[52] J. Wallen, H. Waitzkin, et al., Physician stereotypes about female health and illness, *Women Health* 4 (1979) 371–388.

[53] O. Werner and G. Schoepfle, *Systematic Fieldwork* (Sage, Newbury Park, CA, 1987).

[54] C. West, Ask me no questions....an analysis of queries and replies in physician–patient dialogues, in: S. Fisher and A. Todd, eds., *The Social Organization of Doctor – Patient Communication* (Center for Applied Linguistics, Washington, DC, 1983).

[55] C. West, Medical misfires: mishearings, misgivings, and misunderstandings in physician–patient dialogues, *Discourse Processes* 7 (1984) 107–134.