THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Designing computer-based frameworks that facilitate doctor-patient collaboration

Citation for published version:

Buchanan, BG, Carenini, G, Mittal, VO & Moore, JD 1998, 'Designing computer-based frameworks that facilitate doctor-patient collaboration', *Artificial Intelligence in Medicine*, vol. 12, no. 2, pp. 169-191.
[https://doi.org/10.1016/S0933-3657\(97\)00049-3](https://doi.org/10.1016/S0933-3657(97)00049-3)

Digital Object Identifier (DOI):

[10.1016/S0933-3657\(97\)00049-3](https://doi.org/10.1016/S0933-3657(97)00049-3)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Artificial Intelligence in Medicine

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Designing computer-based frameworks that facilitate doctor–patient collaboration

Bruce G. Buchanan^{a,b,*}, Giuseppe Carenini^b, Vibhu O. Mittal^c,
Johanna D. Moore^{a,b}

^a *Department of Computer Science, University of Pittsburgh, Pittsburgh, PA 15260, USA*

^b *Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA 15260, USA*

^c *Justsystem Pittsburgh Research Center, 4616 Henry Street, Pittsburgh, PA 15213, USA*

Received 31 December 1996; received in revised form 31 July 1997; accepted 1 September 1997

Abstract

A current trend in medicine involves establishing collaborative problem solving between patients and physicians in order to involve patients more in their own care. Neither diagnosis nor therapy can be completely successful unless the patient and the doctor understand each other and collaborate with each other in an effort to gauge the other's requests, needs and concerns. This is made even more difficult by the fact that there is often a big difference between the doctors and patients in terms of expectations, vocabulary used, and other factors. For diagnosis of many disorders, a detailed description of the problem and of the patient's history are required. For therapy, patients must understand how and when to take prescribed drugs, what changes to make in diet, exercise, or lifestyle—and why they are important. This paper describes a model of asynchronous collaboration between people with very different knowledge of medicine in which a computer framework attempts to mediate between patients and physicians and reduce some of the differences in communication. It allows patients to pace themselves in familiarizing themselves with the relevant domain terms, some of the medical factors underlying the conditions under question, and the justifications and implications of the prescribed treatment plan. It also allows physicians to request more information of patients and gives patients contextual information to help them understand the underlying reasons for the questions. This framework has been partially implemented in the domain of migraines. As described in the paper, not only is the system

* Corresponding author. Tel.: +1 412 6249183; fax: +1 412 6249189; e-mail: buchanan@cs.pitt.edu

designed to cooperate with the patient, but using the system also results in better mutual understanding between the doctor and the patient, thus leading to better collaboration between them. © 1998 Elsevier Science B.V.

Keywords: Computer-based frameworks; Patients; Physicians; Communication

1. Introduction

A major goal of clinical medicine is to provide high quality health care at reasonable cost. Successful management of many chronic conditions requires not only effective tools for diagnosis and therapy but also the capability—on the part of both physicians and patients—to access and exchange all the relevant information regarding their roles in the diagnosis and management of the condition. This is because, for many long-term disorders, including migraine, physicians are dependent upon the patient to supply reasonably complete and accurate information on which to base clinical decisions. Furthermore, once these conditions are accurately diagnosed, a successful management program requires the understanding and cooperation on the part of the patient in following the physician's recommendations. One of the main factors in patient non-compliance is poor understanding on the part of the patient: difficulties with technical terms or the inapplicability of standard boiler-plate questions can often lead to inaccurate or incomplete answers, and often repeated, unnecessary visits to the clinic in the future [5,11,16,27,32]. Clearly, better communication between patients and their physicians is necessary to alleviate compliance problems arising from lack of understanding.

Studies have shown that patient compliance is strongly correlated with patient satisfaction [5,20]. Patient satisfaction, in turn, is correlated with patient understanding of, and attitudes about, their ailments and their therapies. Thus, the effectiveness of communication between the doctor and the patient can play an important role in patient compliance which in turn is linked to the effectiveness of management. A current trend in medicine is to view patient care as a collaborative partnership between a patient and a physician [14,28,37,42], rather than a uni-directional transfer of information (from the doctor to the patient), as it has often been reported [17].

A number of studies that focused on the issue of improved long term management identified the following aspects related to collaboration: (i) collaboration was found to be inhibited by variable and unpredictable interest, poor communication, and differences between the physicians' and patients' preferences for follow up; "communication barriers, patient preferences, and misperceptions between the participants need to be addressed before collaboration can occur" [51]; (ii) collaboration depended on respect and understanding of the unique and complementary perspectives of each of the participants and their role in the desired outcomes [34]; (iii) collaboration has an important role in identifying likely areas of mutual interest, particularly the choice of language in consent forms and information

sheets, the design of consent forms, the amount of information provided, and the specification of risks and benefits [26]; and (iv) perceived lack of collaboration in health-care settings (and patient satisfaction) can arise from issues ranging from perceived physician unavailability, discounting patient and/or family concerns, poor delivery of information, and poor understanding of the patient and/or family perspective [2]. These studies show not only that collaboration between the physician and the patient is essential in effective health-care, but that the most important issue in facilitating collaboration is that of communication. Therefore, it is important that physician-patient communication be facilitated to improve collaboration and to minimize misunderstandings.

However, there are a number of factors that stand in the way of effective doctor–patient communication at present:

Time pressure: A number of studies have shown (e.g. [23,43]) that doctors do not have the time necessary to discuss every aspect of the diagnosis and treatment plan with the patients. With current trends in medicine, the likelihood of physicians having more time per patient is small.

Language differences: Doctors and patients see things differently and use significantly different terminology [31,41,50]. Therefore, even when physicians are able to take the time to give patients detailed explanations concerning diagnosis or treatment, patients may often not understand the terms or concepts used. Empirical research on medical discourse (e.g. [15,17,40,46,50]) demonstrates that an ‘information gap’ often exists between doctors and patients. Even when physicians are able to take the time to give patients detailed explanations concerning diagnosis or treatment, their hearers may misunderstand the language or concepts used.

Lack of patient initiative: A study of doctor–patient interactions found a striking asymmetry in information exchange between physicians and patients: in medical discourse, physicians ask the questions and patients provide the answers [17,49,50]. 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 explanations to patients [48]. It is not surprising that an ethnographic study conducted in the course of this project¹ found that most patients still had a number of unanswered questions and concerns after the clinical visit.

Patient reluctance: Studies as part of this project also suggest that even if patients are given the time and opportunity to speak with their physicians, they often hesitate to ask certain types of questions [3]. For instance, in the case of migraine, a significant percentage of patients were afraid that their headaches² were caused by

¹ An important feature of our project was using ethnographers to determine the information needs of patients and to help us design the interactive system to meet those needs. We have not emphasized that aspect in this paper, but we believe that a prerequisite for constructing effective communication tools is identifying what patients and doctors want to know and where they perceive difficulties in determining it. For more details on this dimension, see [3]. The anthropological studies were conducted by Dr D.E. Forsythe and her colleagues.

² A headache is a symptom of classic migraine, but migraines also occur without the headache.

tumors or other, more serious, causes than migraine; however, they almost never asked this question directly.

Further exacerbating this problem of the gap between physicians and their patients is the fact that for many long-term conditions, information exchange is rarely complete because both treatment plans and patients' conditions change over time. One solution to this problem is to involve patients much more closely in their treatment management: this involves attempting to clarify, to the patients, the underlying reasons for the various choices and decisions made during the diagnosis and formulation of the management plan. This allows the patients to make reasonable decisions, within physician specified limits, regarding various parameters such as dietary restrictions, drug dosages, time-intervals between clinical visits, lifestyle compromises, etc. Thus, an increased effort to apprise the patient of these factors is likely to result in fewer unnecessary visits in the future as well as increased patient satisfaction.

The need for better communication has long been recognized [10]. One way to inform patients is to give them printed materials to take home and read. Attempts to measure the effectiveness of handouts have, in the past, yielded conflicting results ([19,33,45]). This may be due to the fact that the effectiveness depends, in large part, on the patients being able to understand the materials, as well as their perception of the relevance of the information included in the handouts. The potential of computers to tailor handouts to individual patients has resulted in increasing use of these handouts [7,24,25]. There is also evidence that tailored messages are more effective than generic ones [11,22,23]. Even if more expensive approaches to patient education (than giving patients printed handouts) are undertaken, recent studies have shown that patient education saves significantly more money than the original investment, often by a factor of three or more, by avoiding future health services utilization [1,38].

Taking these factors into account, our approach to the effective management of long term situations is that it requires a concerted effort on the part of both the physician and the patient to cooperate and collaborate with each other in understanding the needs and constraints of each of their roles: both during the clinical visits by the patient, as well as after the visit in terms of ensuring that both parties—the physician as well as the patient—understand the symptoms, the diagnosis, the management plan and its implications satisfactorily. Clearly, the major issues hindering collaboration-communication barriers, language choice, patient preferences, and misperceptions cannot be easily addressed using a static, printed handout, even if it is tailored to the individual patient. This is because handouts cannot address follow ups or clarify misunderstandings; nor do they help the physician understand the patient better. In an effort to approach this issue from a different perspective, we designed an interactive, computer-based framework that can be used to facilitate the needs of both the physicians and the patients in communicating with each other. We emphasize the design considerations in this paper and will illustrate many of the points with examples from a prototype system [3]. This system has been partially implemented and tested with a small number of

actual patients in order to verify some of the basic assumptions on which the rest of the system has been designed. Some features of the framework discussed in this paper have not yet been implemented. However, as we will discuss in this paper, the preliminary observations with even this partial system strongly suggest that the effects of such a program can foster better doctor–patient collaboration.

2. A framework for supporting collaboration

In designing the system, we divided the activities during a typical clinical visit into three phases: (i) information gathering by means of a questionnaire or an interview by nurse or physician; (ii) examination by physician-with follow-up questions by physician; description of diagnosis and treatment plan, and opportunities for patients to ask questions; and finally (iii) post-examination information seeking by patients including making phone calls to the doctors' offices with questions, reading handouts or magazine articles, or asking friends. The framework facilitates asynchronous communication between the doctor and the patient at multiple levels in the first and the third phases and may facilitate interactive (synchronous) communication in the second phase by: (a) giving patients more understanding of important factors; and (b) giving physicians more information about their patients. The system is not an expert system in the traditional sense (such as MYCIN [4]) of suggesting a diagnosis or a treatment recommendation; nor is it a traditional collaborative system with which physically distributed users sharing a common terminology interact in real time to solve a problem. Rather, this system is an attempt to improve the quality of the information exchange between a physician and a patient by helping bridge the communication gap between them, in part by structuring its interactions with both the physician and the patient in terms of their own goals, terminology, and expertise.

Our initial implementation of a prototype system for evaluating and refining these ideas was in the domain of migraine. Migraine, like other conditions such as diabetes, asthma and chronic sinusitis, is more managed by the patient with a physician prescribed management plan than cured by the physician. Management plans usually include careful monitoring of lifestyle, habits and food intake, the identification and avoidance of migraine triggers, and sometimes, prescription of prophylactic drugs. These management plans can be effective only if the patient understands them well enough to follow them and is motivated to follow them. Motivation, in part, depends on understanding the reasons underlying the various components, the implications of the proposed actions, the time frames involved, as well as the fact that the doctor may need to vary both the dosages as well as the drugs over time to find the best match. Thus, a doctor and a patient must necessarily collaborate with each other over an extended interval during which both diagnosis and management plan may be refined.

The system consists of three main conceptual sub-systems: a data-gathering module, an explanation generating module and an interaction module. Patients interact with both the data-gathering module and the interaction module. Fig. 1

<p>Migraine Interactive Information Sheet</p> <p>Today you were seen by Dr. Rivers and diagnosed as suffering from migraine. The most common symptom of migraine is a moderate to severe headache. Migraine patients also frequently experience visual symptoms, nausea, sensitivity to light and noise, as well as confusion. Your head may feel tender and sore when you have a migraine headache. Much of the pain in a migraine headache is thought to be due to the stretching of blood vessels in the scalp and the head. Your symptoms included flashes, light spots, double vision and nausea, which are all consistent with typical cases of migraine.</p>	<p>Ibuprofen is a drug that is used for analgesic treatment of migraine. Drugs for analgesic treatment are intended to relieve the pain of migraine headaches. These are drugs that you must take when headaches occur. In general, analgesic treatment is suitable for patients with infrequent or mild headaches. Some patients experience side effects from ibuprofen such as stomach pain and discomfort.</p>	<p>side effects from ibuprofen</p> <p>What are other side effects from ibuprofen?</p> <p>What side effects of ibuprofen are serious enough to call the doctor? _____</p> <p>What are side effects? _____</p> <p>What is ibuprofen?</p>
	<p>Migraine (need for time):</p> <p>Since each patient reacts differently to medications, it is crucial that you do your best to understand your own needs and continue working with the doctor to develop a therapy management plan to fit your individual needs.</p>	<p>estrogen replacement ...</p> <p>What is 'estrogen replacement therapy'?</p> <p>Why does estrogen replacement therapy affect migraine?</p> <p>What drugs can be taken for estrogen replacement therapy? _____</p> <p>What is estrogen?</p>
<p>Today you were seen by Dr. Rivers and diagnosed as suffering from migraine. The most common symptom of migraine is a moderate to severe headache. Migraine patients also frequently experience visual symptoms, nausea, sensitivity to light and noise, as well as confusion. Your head may feel tender and sore when you have a migraine headache. Much of the pain in a migraine headache is thought to be due to the stretching of blood vessels in the scalp and the head. Your symptoms included flashes, light spots, double vision and nausea, which are all consistent with typical cases of migraine.</p>	<p>Migraine is also strongly hereditary. You report that your mother and sisters have had severe headaches that are similar to yours. Your family history is a further indication that you suffer from migraine.</p>	<p>While migraines can produce very intense and painful headaches, they are 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 only. Your physical examination was normal, which indicates that more serious causes of headache such as bleeding from strokes or tumors are very unlikely.</p> <p>Migraines often get better as people age, because hormones play a role in making the attacks more severe. Because this condition is often related to the level of female hormones, women are more likely than men to develop it. Migraine attacks become less severe as you grow older. However, women who take hormone pills after menopause (estrogen replacement therapy) may not get this improvement.</p> <p>A migraine attack can be triggered by such things as stress, a change in your sleeping habits, or something you ate or drank. You mentioned bright lights, delayed</p>

Fig. 1. A snapshot of the interface for interacting with the treatment plan.

shows a snap-shot of the interaction module as used by the patient during the post-physician consultation phase. In general, both the interfaces—for the data gathering module, as well as the interaction module—are designed to minimize the amount of typing needed for input: the sub-systems in this framework attempt to infer, whenever possible, potential inputs, answers or queries that the user might have and present these as choices that can be selected by using a mouse based input mechanism. These three modules are closely linked to one another and communicate with each other.

A data-gathering module is used by patients to describe their prior medical history and other relevant information before an initial visit, as well as answer update questions on each subsequent visit. Questionnaires for initial visits can be lengthy and detailed, whereas questionnaires presented to the patient on subsequent visits are shorter and more focused. This module is designed to do more than data collection; as described in the next section, it familiarizes patients with the factors that are relevant to their diagnosis and the terms their doctors are likely to use in the discussions. Furthermore, it allows patients to ask for clarifications about terms, it allows them to request justifications or contextual information about questions, and it dynamically omits questions that are not relevant for a patient and questions whose answers can be inferred from other questions. This module also generates a short summary of the patient interaction that can be used by the doctor to quickly get a perspective on a patient's medical background, as well as additional facts or concerns the patient may have added.

An explanation module capable of generating natural language explanations of concepts in the system, such as migraine, its symptoms, drugs, side-effects, etc. These explanations are tailored to individual patients in terms of terminology used, their background, interests, prior interaction, etc.

An interaction module that presents the information on the screen and manages the subsequent interaction with the patient. This module has been designed to facilitate patient exploration in a semi-guided manner, where users are free to explore paths of their choice, but the system keeps track of what they have been told previously and what they might be interested in based on the context.

During both the patient-system interaction phases, both before and after their consultation with the physician, the patients retain a large measure of autonomy in choosing what to read, ask about or answer and at what pace. However, at all times during the interaction, the system reasons about information either required by the physicians for diagnosis, or important for the patient to know about. Using knowledge about the context, the patient model, and discourse strategies, the system generates both the text and the follow up prompts that the patients see on the screen. Thus, neither the system, nor the user, are ever in exclusive control over the interaction and the presentation of information. Rather than the more typical 'master-slave' relationship between a system and the user, this framework may be better understood as one of mutual collaboration between the system and the user: allowing the user to explore issues of interest while trying to best achieve the physician's goals (collect, or present, relevant medical information necessary for diagnosis and treatment).

The following three sections describe the three major modules of the system in more detail. (A complete technical description of the system with more detailed descriptions of each of the modules can be found in [3,8,9,36]).

3. The data-gathering module

Diagnosis of headache, and many other chronic conditions, requires 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. In general, obtaining a detailed history is time consuming. Because it is normally taken, or at least reviewed in detail, by a physician, it is also expensive.

The data-gathering module as implemented allows patients to enter relevant medical information at the time of their initial visits. While the idea of using computer technology to take a medical history is not new ([29,44]), our design extends the standard computer based implementations of paper based questionnaires in at least three important ways:

- patients can answer questions in the order they prefer, while eliminating all irrelevant or inapplicable questions,
- patients can request definitions for terms or phrases they do not understand, or ask for reasons underlying questions³
- physicians can easily modify the questionnaires presented, for example, to collect data for a study and to modify the presentation of a summary of answers which can be scanned by the physician and included in the patient record.

In our model, the initial information gathering stage offers an important opportunity to facilitate communication. During this phase the patient is first exposed to medical terminology, as well as potentially important factors in the diagnosis of the condition as the data-gathering module presents the questions. The reactions of the patient during this session, along with the answers, can then be summarized by the module in a form that can be used both by the physician, as well as other modules in the system to structure their interaction with the patient, both to present explanations as well as gather further information.

An intelligent medical-history gathering module such as this has several potential advantages for the physicians: (i) it can save significant amounts of time; (ii) data collection can be made more uniform and complete; (iii) any exceptional readings, factors or allergies can be flagged by the system and highlighted in the summary; (iv) patients may take more time and may be more comfortable answering certain questions on the system as opposed to a human; and finally (v) interacting with a

³ This feature was implemented after the prototype was evaluated and described in [3].

data collection framework may help patients recall relevant details by the time they see the physician, something that often happens after the visit. As mentioned earlier, it also introduces them to some of the terminology and the types of underlying factors in their case (such as possible triggers in the case of migraine).

In order for this history taking module to be effective, there are at least four collaborative capabilities that must be designed into the system: (i) users must be able to ask for clarifications of terms or phrases they do not understand; (ii) the system must be able to present the reasons underlying questions being asked, or rephrase the question in alternative ways; (iii) the system must be able to reason about the relevance of any question in a given context; and finally (iv) physicians must be able to edit the questions, answer choices, and explanatory text presented by the system. The reasons underlying the first capability are obvious; if a user does not understand certain terms in the questions, it is unlikely that the answers to those questions would be correct. The reasons underlying the second capability are sometimes less obvious. For instance, migraines can be triggered by either variations in sleeping patterns or changes in the caffeine intake. Since people tend to sometimes sleep late on weekends, and drink less coffee as well (at home versus the office), such questions are often used to filter out multiple factors together. However, for many people working in shifts, or people working on weekends, etc. the question should be rephrased to ask about their ‘day off’ rather than a ‘weekend.’ Being able to ask about the reasons underlying questions asked enables the patients to answer them appropriately. The ability to reason about the applicability of specific questions in any context is important if the system is to avoid asking either irrelevant questions or questions whose answers can be inferred based on other answers. Finally, since each clinic, even each physician, has preferences for which data are collected and how questions are phrased, it is necessary to make editing easy.

It should be noted that none of these capabilities can be easily added on to an existing implementation. Rather, as has been pointed out previously [18], designing collaborative systems requires that the collaborative aspects be part of the design from the very beginning. One of the capabilities required in a collaborative system is that the system adapt its interaction to the patient. In our system, both the other modules use the record of the patient interaction with the data-gathering module to determine the content and structure of their interaction with the patient. This adaptability of presentation can be important, especially in medical consultations because of the wide range of patients that are seen by a doctor in any given day⁴. Not only can the system tailor the explanations to specific, concrete characteristics such as gender, age and reported symptoms, but the system can also make use of information regarding terms and phrases for which the user asked for clarifications. This information can then be used by the explanation component to generate

⁴ Most previous attempts to generate handouts have attempted to do so for a ‘generic’ patient: however, a study found a large discrepancy between the average patient’s reading comprehension, which was at the sixth grade level, and the average handout, which was pegged at the 11th–14th grade levels [12].

natural language explanations that avoid other, related complex terms. Similarly, knowledge about the questions the patient asked to have restated or justified can also be used to generate appropriate explanations. Thus, in many ways, the history gathering module is designed to foster interactions that resemble discussing the condition with a knowledgeable professional with infinite patience who prompts the patient to remember relevant aspects. Fig. 2 shows part of the internal representation of the patient model generated by the data-gathering module.

4. The explanation generator

The explanation generator lies at the core of the collaborative capabilities of the system and distinguishes our system from other, previous work in this area. It reasons about the communicative goals of the system, the manner in which they were achieved, and the mapping between the goals and parts of the textual explanation generated. It constructs a record of the context of all the interaction and uses this record as a knowledge source to structure both the responses as well as follow up questions (Fig. 3). Given a goal, such as (DESCRIBE (CONCEPT MIGRAINE)),⁵ a text planner searches its library of explanation operators looking for strategies that could achieve this goal, resulting in the generation of a natural language description about the concept migraine.

Plan operators can be viewed as recipes for achieving communicative goals. Each operator represents the following information: (i) an indication of the effect(s) that the operator is intended to have on the user's mental state; (ii) a list of constraints that must hold for the operator to be applicable, which may include conditions on the domain knowledge base, user model, or dialogue history; (iii) optionally, an operator that specifies a sub-plan for achieving the effects; and (iv) if a sub-plan is specified, information about the discourse relations between steps in the sub-plan. Once a plan operator is selected, it may in turn post subgoals for the planner to refine. Planning continues in a top-down fashion until the entire plan is refined into primitive operators, i.e. speech acts such as INFORM and RECOMMEND.

The result of the planning process is a discourse tree, where the nodes represent goals at various levels of abstraction. The root node is the initial goal, and the leaves represent primitive goals, such as (INFORM (IS-A ANALGESIC IBUPROFEN)), which represents the proposition that *ibuprofen* is an analgesic drug. The discourse tree can also, optionally contain other relevant information. For instance, communicative goals can be related to other goals in the tree by means of various relations. A discourse tree generated by our current system is shown in Fig. 4. It shows a portion of the hierarchy of goals for the text plan that leads to the production of the description about migraine shown in Fig. 1. The discourse tree can be used to record varying amounts of information at varying levels of

⁵ The actual syntactic form used to represent the goals and speech acts in the implemented system is more complex; for the sake of clarity, we have shown simplified versions in this paper (see [35] for more details).

;;; Temporal Aspects

years-since-onset-first-episode: 5

;;; Associated Symptoms and Prodrome

Findings: SYNCOPÉ

Prodromal-findings: - NAUSEA (Sometimes)
 - FLASHES (Always)
 - LIGHT-SPOTS (Sometimes)

;;; Provocative Factors

Triggers: FLICKERING-LIGHTS, TOO-LITTLE-SLEEP, CAFFEINE,
 CHOCOLATE

;;; Familial History

Familial-Incidence: MOTHER

;;; General Information

Clinically-Relevant-History: MENOPAUSE
 Is-Taking-Medication: ESTROGENS

;;; Localization

Has-Dynamics: Hemicrania
 Has-Dynamics: Start-local

;;; Terminology Investigated*

Terms-Investigated-by-Patient: HEREDITARY, SENSITIVITY, TUMOR,
 TRIGGER, AURA, PRODROME

;;; Question Times*

More than 2 but less than 5 minutes: PRIOR-MEDICATION
 More than 5 but less than 10 minutes: CURRENT-MEDICATION, TRIGGERS

Fig. 2. An example of information collected by the data-gathering module which is represented in a patient model for a specific patient. The last two terms (terminology and timing notes) are different from the medical items in that they are characteristics of the patient's interaction with the system. These two items were added since the version of the system reported in [3].

granularity. The main purpose of representing the goal structure in the discourse tree (apart from generating the initial explanation) is that it allows the system to reason about information that has already been presented and the form in which it was presented. Furthermore, a link is maintained between the various leaves in the

discourse tree (representing the primitive goals) and the resulting text that is shown on the screen to the user. Thus, at any given time, should the user indicate a state of confusion or dis-satisfaction with any aspect of the explanation, the goal that led to the generation of the statement(s) in question can be analyzed by the system. As described in the next section, this representation enables the system to recover from failed goals, take into account previous interaction with the user (e.g., if the system has already defined an ‘analgesic drug’, it can, in future interactions, use the term without defining it again), as well as reason about possible elaborations/clarifications the user might have.

The discourse tree generated by the planning process is used by the system to generate the natural language explanations. Generating natural language from the propositional content of the leaves in the tree can be done in a variety of ways, depending on the desired speed and flexibility in the final language presented to the user. There are a number of language related issues that the system can potentially reason about, such as generating appropriate referring expressions, computational implicature, lexical choice and grammatical realization. However, there is a computational cost associated with each of these reasoning modules. From the very beginning, our prototype was meant to be evaluated with real patients. Thus, not only was the prototype supposed to be robust, it was also required to generate responses in a reasonable time frame (initial explanation generated in a matter of a minute and clarifications and questions answered in a matter of seconds). Given our

Fig. 3. The system architecture for the explanation and interaction modules.

Fig. 4. A partial text plan used by the system to generate a description of a migraine diagnosis, the patient symptoms, triggers and treatment plan prescribed by the physician. Triangles in the figure stand for unexpanded parts of the text plan.

constraints, we concentrated our efforts on the text planning process, i.e., selecting and organizing the content to be presented to the patient. Rather than use a sophisticated, highly flexible sentence realization system such as, for instance FUF [13], we chose to build and use a template-based generator to produce the actual English sentences expressing this content. While this approach allows considerable flexibility in what content gets expressed and how that content is organized, it limits our ability to tailor explanations at the level of syntactic structures and lexical items. In general, the speech acts produced by the planning process can be used as input to a more sophisticated, highly flexible sentence realization system such as FUF. This would allow the system to make principled choices regarding lexical items and have greater control over the phrasing, but would also require several orders of magnitude more computational time. However, as processing speeds get faster and faster, there is no reason why sophisticated generation systems cannot be built using the same underlying framework as described here that will be capable of the response times equal to, or faster than, our current, simple template based generation system.

5. The interaction module

In the previous section, we discussed how the use of a text planning formalism enables the system to represent and reason about communicative goals and the discourse plans that were generated to achieve them. However, the text planning framework can still only be used, as described, to generate text. In order for a system to be collaborative, the system must be able to take as input questions, or requests for information, from the user. Studies of human–human interactions show that information-seekers often follow up requests for information with further questions [30,35,39], illustrating the fact that initial explanations are seldom sufficient to satisfy users in real situations. For this reason, an essential part of our collaborative model is facilitating follow-up questions. This capability is especially crucial in applications such as patient education where misunderstandings could result in serious consequences such as failure to use drugs or equipment properly.

Designing an effective interaction module requires at least two additional capabilities: (i) the ability of the system to allow users to ask questions; and (ii) the ability of the system to answer questions effectively (taking into consideration issues of content, amount and granularity of information, terms used, phrasing, etc.). Even though these may seem trivial at first, especially the former, and are consequently seldom discussed, they are both important in the design of a system that is meant to be used by real patients.

The most natural solution to the first issue would be to allow users to type in the questions they have. However, there are a number of problems with this apparently obvious solution: (i) the system must then be able to handle questions and statements that mix references to prior discourse and the domain, e.g., the question “could you please explain that last part again?” is ambiguous because the ‘last part’ could be either the last part of the previous explanation, or a part of some object

being discussed; and (ii) even more important perhaps, the system would need to constrain users to ask only those questions that the system can answer. Studies suggest that systems that do not clearly and explicitly circumscribe the type of questions that can be asked lead to users being unhappy at the system's inability to answer some of their questions [21]. Formal query languages, such as those used for database queries, or pseudo-natural language queries such as those used in systems such as PROLOG/EX 1 [47], are also problematic in being difficult to learn and frustrating to use. Perhaps a bigger problem with all of these approaches is that they assume that users will be able to pose well-formulated follow up questions. However, this is not always the case. Users who are confused may not even know exactly what question they wish to ask.

To work around these problems, our interface was designed around a direct manipulation paradigm in which system-generated texts were considered to be complex, highly structured objects that the system understood and could reason about. A text span, for instance, can be composed of multiple, shorter text spans, each of which could be generated in response to a goal represented in the text plan. By preserving a record of the relationship between the goals and the constructed text spans, the interface allows users to point to the portions of system-generated text that they do not understand or would like more information about. In response, the system provides a menu of questions that may be asked about the selected text. Question menus are generated on the fly taking into account the user model, underlying knowledge bases, and the prior discourse context [36]. This approach provides a simple, yet intuitive interface for allowing users to pose follow up questions about system-generated utterances, and at the same time circumscribes the set of questions users may ask. However, for such an interface to be feasible, the system must be able to understand what the user is pointing at, that is, the system must have a model of its own explanations including the goals of the explanation and how the generated text achieved those goals.

Fig. 1 shows a snapshot of the interactive information sheet, the initial text produced by the system for a patient diagnosed as having migraine. The figure shows the cursor, which is depicted as an arrow, being moved over the first sentence in the third paragraph. As the cursor is moved over parts of the sentence, a box appears over the smallest constituent about which follow-up questions may be asked. For example, if the cursor is moved over a noun phrase, e.g. 'migraines', a box appears around the noun phrase. When the cursor is over the verb, a box appears around the clause containing that verb. In the figure shown, the cursor is over the connective 'while', which relates the two clauses in this complex sentence. The system draws a box around the entire sentence. By moving the cursor over parts of the sentence and observing which text objects appear in boxes, the user can explore options for asking further questions about statements made by the system (clauses and sentences) or the terms used in them (noun phrases). When the user identifies a boxed text object that she would like to ask further questions about, she can select it by clicking the mouse in the box. The selected item is then highlighted by the system (indicated in reverse video), and a menu of possible follow up questions about the selected item are generated. Fig. 1 shows two selected items,

estrogen replacement therapy and side effects from ibuprofen and the menus that appear when these items are selected.

It is important to keep in mind that although the interface bears a resemblance to a hypertext-style interface, the system is not a hypertext system in the traditional sense, i.e., it is not organized as a collection of canned pieces of text interconnected by typed links. Explanations are dynamically generated by the system in response to a user's question or the system's need to communicate with the user. Because the text is generated dynamically, it is not possible to identify in advance the particular text objects that should be mouse-sensitive nor the links to other objects. Portions of text that should be mouse-sensitive can only be determined during the generation process, when the abstract rules for determining which types of objects should be mouse-sensitive are applied to the particular instances of text that are being generated. Finally, follow-up questions, which correspond to the links in traditional hypertext systems, are not pre-coded and fixed in advance, but are generated dynamically using heuristics that are sensitive to domain knowledge, the user model, and the discourse context.

There are several advantages to this approach. First, what follow-up questions are meaningful is highly context-dependent. Moreover, because one of the main problems with hypertext systems is that users get lost in the network and may even forget what it was they were originally seeking [6], it is especially important to present a confused user with a small set of pertinent follow-up questions as opposed to a larger set of questions, many of which may seem off the topic or even irrelevant [30]. In our system, dialogue context (provided by the text plan) and the user model are exploited by strategies that prune the list of possible questions down to those that appear most relevant. Second, and perhaps more importantly, the abstractions used by the system for automatically generating follow up questions relieve system builders of the intractable task of hand-crafting all of the links and subsequent texts. For example, in our current implementation, there are approximately 75 factors that are used to tailor the explanations generated. These factors include information about patient symptoms, family history, drugs prescribed, their side effects, and so on. Writing separate explanations for different combinations of these 75 factors, as well as defining the links for each of the follow-up questions would be an impractical task. Clearly, providing a technology that can automate the generation of the texts and the linkages between them can greatly reduce the time taken to develop such systems.

There are three main issues that must be addressed in designing a sub-system to generate menus of follow up questions: (i) determining the types of text spans that can be queried by the user; (ii) identifying sources of information that can be used to generate possible questions about the highlighted text; and (iii) identifying factors that can be used to prune the set of candidate questions so that the resulting menu will contain a small set of questions relevant to the user's current concerns.

As mentioned in the previous section, the discourse tree generated by the text planner maintains links between the text shown on the screen and the underlying goals used to generate it. The system uses a simple tree traversal technique to find all the goals in the discourse tree that are closely related to the sentence or text span

being queried. Consider, for instance, the discourse tree shown in Fig. 4. Should the patient query the noun phrase ‘migraine’ in the sentence ‘You were diagnosed as suffering from migraine,’ the system can reason not only where the term appears in the overall text (it appears in the very beginning), but also that the goals⁶ that caused that sentence to be generated in the first place were:

- Generate patient handout for migraine (1)
- ⇒ Describe diagnosis of migraine (2)
- ⇒ Inform-patient diagnosis migraine (3)

The system can make use of this information to appropriately generate possible questions regarding each of the goals. The algorithm that actually generates the questions is somewhat more complex, and takes into account the patient-model generated by the data-gathering module, the relationships defined in the knowledge base, as well as sibling goals in the discourse tree⁷. The resulting set of questions can then be presented to the user in the form of a menu. However, the set of questions generated by a tree traversal as described (along with additional questions generated because of the knowledge base relationships between terms such as ‘migraine’ and ‘classic migraine’, along with questions generated from a record of previous questions, such as the relationship/differences between ‘migraines’ and ‘tension headaches’) can often become very large and would be cumbersome to present in a menu. At this point, the system prunes the question set by (i) using a variety of heuristics about the relevance of a question in the given context; and (ii) determining whether there is any information in the system that has not already been presented to the user about a question. For instance, if a question regarding the physician’s diagnosis of migraine has already been answered previously, the system will eliminate that question from the set. Similarly, if the patient has not previously asked about tension-headaches, the system will not attempt to present a question contrasting the two concepts. An example of the use of available domain knowledge (or lack of it) to prune a question can be seen in Fig. 1, where the system does not present a question about the side-effects of estrogen-replacement therapy because it does not possess any additional information about it, but does so in the case of ibuprofen. This results in menus that are both small and focused, sensitive to both the discourse context and the individual user.

Initial evaluations of the interaction module were carried out by comparing the set of questions generated by the system with a set of questions that were determined by ethnographic research. Our initial evaluations suggest that: (1) patients were able to understand and use the interface mechanism described, i.e., menus of dynamically generated questions, to participate in an interactive dialogue

⁶ Simplified paraphrases of the goals are shown here for the sake of readability.

⁷ The algorithm used is quite simple, but a complete description would require describing the knowledge base as well as the heuristics used in generating the menu title, item orderings, etc. Lack of space prevents us from doing that here, but complete details may be found in [36].

with the system; (2) the questions generated by the system covered most, if not all, the important issues that users wanted to ask about; and (3) the context sensitive nature of the question generation algorithm resulted in small, focused menus of follow up questions that allowed users to quickly find the desired question. Furthermore, in post-session interviews, patients indicated a high degree of satisfaction with the mode of interaction provided by the interface, with some of them mentioning that the patient-controlled exploration aspect was an important component of the satisfaction. Thus, this menu driven interface is designed expressly to allow for a collaborative exchange of information between the patient and the system: the patient can choose the aspects on which the system provides further information, while the system structures all information presented to achieve its high level communicative goals (to explain the management plan and its implications).

6. Observations and conclusions

Effective medical treatment necessitates information exchange about many things, including patients' concerns, physicians' plans, facts about prescribed drugs and reasons for recommended changes in environment, diet or life style. Success depends in large part on the patient's compliance with the recommended regimen. In order to obtain compliance, patients must have a good understanding of the diagnosis, its implications, and their prognosis, as well as the rationale for the therapy selected. Providing explanatory materials and motivating patients to comply require significant collaboration between the patients and their physicians. As mentioned previously, a number of studies have suggested that the major issues in increasing patient-physician collaboration lie in effective communication, addressing misperceptions, choice of language and specification of risks and benefits of the proposed treatment. Other studies have shown that a significant communication gap exists between patients and their physicians, and that this is exacerbated by the lack of time on the part of the physicians and the reluctance to ask certain questions on the part of the patients. A fundamental hypothesis of our model is that if patients receive enough information and understand it, they (or their families) can be empowered to take more responsibility for their own care (for example, by adjusting drug doses within limits set by physicians).

When patients' information needs go unmet, the result may be subsequent misunderstandings, compliance problems and treatment failures. This paper has described a computational model that helps facilitate physician-patient collaboration through 'off-line' interactions that bridge the communication gap between them. It does this by allowing patients to answer and explore questions and possible treatment plans at their own pace, at a level that they feel comfortable with. The system structures the content based on their medical history and is capable of clarifying terms they do not understand. Furthermore, the system can prompt the patients to find out further information by presenting them with sets of relevant questions. It then generates a report for the physician summarizing the interaction

with the patient. This allows both parties to get a better understanding of the situation, the implications of the treatment plan, major concerns and areas of possible interest or confusion.

Based on our own observations of 16 actual patients who used this system in a realistic setting, we believe that both doctors and patients benefit greatly from using such a system. We noted that the physicians in our project demonstrated greater understanding of issues important to the patient⁸. The patients were also able to discuss issues regarding the management plans with the physicians better⁹.

The framework described here supports dual goals: (i) a patient better able to understand the terminology and relevant factors as well as appreciate the implications of the management plan prescribed by the physician; (ii) and a physician better able to understand and address the concerns of the patient. There are also several issues that we need to investigate before such a framework is effective in routine use. For instance, our current implementation does not, as yet, deal with the issue of the completeness of information presented to patients. Clearly, if the framework is to be successful in fostering mutual understanding, it must not contradict the physician. If it did, it would undermine the patients' confidence in their physicians' diagnoses and treatment plans. However, the framework should also allow patients to understand alternative choices and symptoms not covered by migraine, if they should want to do so. Our framework currently attempts to balance both these goals by initially presenting, to patients, the physician's point of view. However, the framework also allows patients to query text spans and explore issues regarding alternative choices or the presence/absence of any abnormal or contra-indicating symptoms that might suggest something other than migraine.

We believe that the question menus generated by the system represent, in some sense, statements that a physician, given sufficient time and patience, might make to a patient during a dialogue regarding the management plan. Transcripts of naturally occurring dialogue often show the speakers pausing and asking if a statement just made was clear to the hearer. In such situations, the speaker can make use of both visual and aural cues, such as puzzled expressions, blank faces or uncertain sounds such as 'Uh?' to determine when further explanation may be necessary. In our framework, moving a mouse over a text span and clicking on it is operationally equivalent to expressing puzzlement. Presenting questions to patients helps them learn how to articulate their doubts to doctors in the future. Presenting a list of possible questions also prompts them to learn about information that they may not have thought about previously.

⁸ For instance, after learning that a number of patients interacting with the system had questions about the possibility of tumors, a clinician on our team explicitly addressed that issue with future patients.

⁹ For instance, beta blockers can have adverse effects on athletic lifestyles; this can sometimes be of concern to patients who jog or work out regularly. Patients were able to understand implications such as these and discuss them better.

In general, all interaction with the patient is recorded by the framework. This includes, in addition to the questions and answers, a time-stamped audit trail of all actions taken by the patients¹⁰. An analysis of time spent on each of the questions (in the information gathering phase), or on certain aspects of the diagnosis or treatment plan (in the explanation phase) can be very useful to the doctor in order to get a sense of the issues that might be of concern to the patient. This information could then also be used by the framework to modulate future interactions with the patient.

A concern among physicians, especially those treating patients with long-term conditions, is that patients will become discouraged at the lack of a 'cure' and will drop out of the system. Establishing a partnership is one strategy for physicians to motivate patients to persist while the partners jointly try management plans in search of better quality of life, if not cure. One patient in our preliminary study had severe migraine headaches but was not under a physician's care, trying instead to 'learn to live with them' on her own. Although she had sought medical help several times in the past, she had become discouraged about finding any relief and had dropped out of the system. After reading much of the material our system presented in response to her requests, she remarked that she had found reason to believe that she could be helped and she made a follow-up appointment with the neurologist who had examined her.

We have designed this framework to work with declarative knowledge structures in order to allow for easy porting to other domains. The data gathering module has been reused, with new knowledge, to build prototypes in two additional medical areas. The explanation component—the text planner—has been used in a number of applications ranging from the automatic generation of software documentation to generating presentations about complex, statistical databases. While some of the plan operators used by the text planner are domain specific, and would therefore need to be rewritten, others are general and can be reused across domains. The generality of the question generation algorithm used in the interaction module has not yet been investigated, though we believe that since the question generation algorithm is based on the knowledge base and the text-plans generated, that it should function effectively in other application domains as well (the open issue is regarding the coverage of the set of questions that users would like to be able to ask).

Our experiences and intuitions therefore lead us to believe that the use of a framework such as this can lead to improved mutual understanding between the physician and the patient in several areas of medicine, which can directly contribute to improved collaboration. Our main thesis is that while most¹¹ of the benefits accrued from the deployment of such a system could also have been gained from greater time spent per patient by the physician, we believe that the use of such a

¹⁰ The patient is informed at the beginning of the session that all their actions will be recorded.

¹¹ But not all; as we pointed out this system seems to facilitate queries that patients are often unwilling to ask physicians directly.

system is beneficial to the physicians as well by letting them keep track of issues troubling their patients. Thus, it allows both the patient and the physician to work together better, which is another way of stating that the system fosters collaboration between them.

Acknowledgements

This work was supported, in part, by the National Library of Medicine (grant number R01 LM05299). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the National Library of Medicine. Other members of the Migraine project who contributed significantly to the success of the project were: Diana Forsythe who was in charge of the ethnographic studies and patient surveys, Gordon Banks who served as our resident migraine expert, and who, with Steve Margolis, edited many versions of questions and text, Nancy Bee and Myra Brostoff who participated in the field work, Stellan Ohlsson who, along with Patricia Albecete, investigated the cognitive issues involved in building computational systems for patient education and Ted Rees who implemented the history taker.

References

- [1] E.E. Bartlett, Cost-benefit-analysis of patient education, *Patient Educ. Couns.* 26 (1–3) (1995) 87–91.
- [2] H.B. Beckman, K.M. Markakis, A.L. Suchman, R.M. Frankel, The doctor-patient-relationship and malpractice—lessons from plaintiff depositions, *Arch. Intern. Med.* 154 (12) (1994) 1365–1370.
- [3] B.G. Buchanan, J.D. Moore, D.E. Forsythe, G. Carenini, S. Ohlsson, G. Banks, An intelligent interactive system for delivering individualized information to patients, *J. Artif. Intell. Med.* 7 (2) (1995) 117–154.
- [4] B.G. Buchanan, E.H. Shortliffe, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, Reading, MA, 1984.
- [5] C. Cameron, Patient compliance—recognition of factors involved and suggestions for promoting compliance with therapeutic regimens, *J. Adv. Nurs.* 24 (2) (1996) 244–250.
- [6] P. Carando, Shadow: Fusing hypertext with AI, *IEEE Expert* 4 (4) (1989) 65–78.
- [7] G. Carenini, V.O. Mittal, J.D. Moore, Generating interactive patient-specific natural language explanations, in: J.G. Ozbolt (Ed.), *Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-94)*, Washington, DC, November, 1994.
- [8] G. Carenini, J. Moore, Using the UMLS semantic network as a basis for constructing a terminological knowledge base, in: *Proceedings of the Symposium on Computer Applications in Medical Care*, Washington, DC, 1993.
- [9] G. Carenini, J.D. Moore, Generating explanations in context, in: W.D. Gray, W.E. Hefley, D. Murray (Eds.), *Proceedings of the International Workshop on Intelligent User Interfaces*, ACM Press, Orlando, Florida, January 4–7, 1993, pp. 175–182.
- [10] D.G. Covell, G.C. Uman, P. Manning, Information needs in office practice: Are they being met?, *Ann. Intern. Med.* 103 (1985) 596–599.
- [11] L.H. Daltroy, Doctor–patient communication in rheumatological disorders, *Bailliere's Clin. Rheumatol.* 7 (2) (1993) 221–239.

- [12] T.C. Davis, M.A. Crouch, G. Wills, S. Miller, D.M. Abdehou, The gap between patient reading comprehension and the readability of patient education materials, *J. Fam. Pract.* 31 (5) (1990) 533–538.
- [13] M. Elhadad, J. Robin, Controlling content realization with functional unification grammars, in: R. Dale, E. Hovy, D. Rosner, O. Stock (Eds.), *Proceedings of the Sixth International Workshop on Natural Language Generation*, Springer, Berlin, 1992.
- [14] Y. Engestrom, Objects, contradictions and collaboration in medical cognition—an activity theoretical perspective, *Artif. Intell. Med.* 7 (5) (1995) 395–412.
- [15] S. Fisher, Institutional authority and the structure of discourse, *Discourse Processes* 7 (1984) 201–224.
- [16] L. Forman, Medication: reasons and interventions for noncompliance, *J. Psychosoc. Nurs. Ment. Health Serv.* 31 (10) (1993) 23–25.
- [17] 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.
- [18] B.J. Grosz, Collaborative systems, *Artif. Intell. Mag.* 17 (2) (1996) 67–85.
- [19] G.M. Hawkey, C.J. Hawkey, Effect of information leaflets on knowledge in patients with gastrointestinal diseases, *Gut* 30 (1989) 1641–1646.
- [20] R.B. Haynes, A critical review of the ‘determinants’ of patient compliance with therapeutic regimens, in: D.L. Sackett, R.B. Haynes (Eds.), *Compliance with Therapeutic Regimens*, Johns Hopkins Press, Baltimore, MD, 1976, pp. 26–39.
- [21] H. Helbig, C. Gnorlich, D. Menke, Realization of a user-friendly access to networked information retrieval systems, in: *Proceedings of the AAAI Spring Symposium on Natural Language and the WWW*, Stanford, CA, 1997, pp. 62–71.
- [22] M.G. Hewson, Patient education through teaching for conceptual change, *J. Gen. Intern. Med.* 8 (7) (1993) 393–398.
- [23] R.B. Jones, S.M. McGhee, D. McGhee, Patient on-line access to medical records in general practice, *Health Bull.* 50 (2) (1992) 143–150.
- [24] G. Kahn, Computer-based patient education: A progress report, *MD Comput.* 10 (2) (1993) 93–99.
- [25] G. Kahn, Computer-generated patient handouts, *MD Comput.* 10 (3) (1993) 157–164.
- [26] G. Kent, Shared understandings for informed consent—the relevance of psychological research on the provision of information, *Soc. Sci. Med.* 43 (10) (1996) 1517–1523.
- [27] S.R. Killick, Patient education and understanding: a critical review, *Adv. Contracept.* 8 (1992) 57–65.
- [28] K.I. Kjellgren, J. Ahlner, R. Saljo, Taking antihypertensive medication—controlling or cooperating with patients, *Int. J. Cardiol.* 47 (3) (1995) 257–268.
- [29] B. Kleinmutz, Medical information processing by computer, in: J.A. Jacquez (Ed.), *Computer Diagnosis and Diagnostic Methods*, Charles C. Thomas, Springfield, IL, 1972, pp. 45–72.
- [30] T.W. Lauer, E. Peacock, A. Graesser (Eds.), *Questions and Information Systems*, Erlbaum, Hillsdale, NJ, 1992.
- [31] P. Ley, M.S. Spelman, *Communicating with the Patient*, Warren H. Green, St. Louis, MO, 1967.
- [32] H.M. Litchman, Medication noncompliance: a significant problem and possible strategies, *Rhode Island Med.* 76 (12) (1993) 608–610.
- [33] K.S. Loumidis, R.S. Hallam, B. Cadge, The effect of written reassuring information on out-patients complaining of tinnitus, *Br. J. Audiol.* 25 (1991) 105–109.
- [34] S. Makaram, Interprofessional cooperation, *Med. Educ.* 29 (1) (1995) 65–69.
- [35] J.D. Moore, *Participating in Explanatory Dialogues: Interpreting and Responding to Questions in Context*, MIT Press, Cambridge, MA, 1995.
- [36] J.D. Moore, V.O. Mittal, Dynamically generated follow up questions (Special Issue on Interactive Natural Language Processing), *IEEE Comput.* 29 (7) (1996) 75–86.
- [37] K.N. Moore, Compliance or collaboration—the meaning for the patient, *Nurs. Ethics* 2 (1) (1995) 71–77.
- [38] M. Neri, G.B. Migliori, A. Spanevello, D. Berra, E. Nicolini, Economic analysis of two structured treatment and teaching programs on asthma, *Allergy* 51 (5) (1996) 313–319.

- [39] G.M. Olson, S.A. Duffy, R.L. Mack, Question-asking as a component of text comprehension, in: A.C. Graesser, J.B. Black (Eds.), *The Psychology of Questions*, Ch. 7, Erlbaum, Hillsdale, NJ, 1985, pp. 219–227.
- [40] V.L. Patel, D.A. Evans, D. Kaufman, A cognitive framework for doctor–patient interaction, in: D.A. Evans, V.L. Patel (Eds.), *Cognitive Science in Medicine*, MIT Press, Cambridge, MA, 1989, pp. 257–312.
- [41] S.C. Ratzan, *Health Communication: Challenges for the 21st Century*, SAGE Publications, Thousand Oaks, CA, 1994. (Special issue of *Am. Behav. Sci.* 38 (2) (1994)).
- [42] S.A. Rosa, B.R. Hasselkus, Connecting with patients—the personal-experience of professional helping, *Occup. Ther. J. Res.* 16 (4) (1996) 245–260.
- [43] P. Schattner, Patient education classes in general practice, *Aust. Fam. Phys.* 22 (7) (1993) 1229.
- [44] W.V. Slack, C.W. Slack, Patient-computer dialogue, *New Engl. J. Med.* 286 (1972) 1304–1309.
- [45] W. Sumner, An evaluation of readable preventive health messages, *Fam. Med.* 23 (6) (1991) 463–466.
- [46] A.D. Todd, The prescription of contraception: Negotiations between doctors and patients, *Discourse Processes* 7 (1984) 171–200.
- [47] A. Walker, Prolog/ex1, an inference engine which explains both yes and no answers, in: *Proceedings IJCAI-83*, vol. 1, Karlsruhe, Germany, 1983, IJCAI, pp. 526–528.
- [48] J. Wallen, H. Waitzkin, J.D. Stoeckle, Physician stereotypes about female health and illness, *Women Health* 4 (1979) 371–388.
- [49] C. West, Ask me no questions... an analysis of queries and replies in physician–patient dialogues, in: S. Fisher, A. Todd (Eds.), *The Social Organization of Doctor-Patient Communication*, Center for Applied Linguistics, Washington, DC, 1983.
- [50] C. West, Medical misfires: mishearings, misgivings, and misunderstandings in physician–patient dialogues, *Discourse Processes* 7 (1984) 107–134.
- [51] M.L. Wood, C.L. McWilliam, Cancer in remission—a challenge in collaboration for physicians and oncologists, *Can. Fam. Phys.* 42 (1996) 899–910.