Generating Patient-Specific Interactive Natural Language Explanations

Giuseppe Carenini Vibhu O. Mittal Johanna D. Moore
Intelligent Systems Laboratory, University of Pittsburgh, Pittsburgh, PA 15260

Patient compliance is a significant problem and is strongly correlated with the patients' understanding of their condition and prescribed treatment. Since doctors typically do not have large amounts of time to educate patients, and impersonal, voluminous patient handouts are largely ineffective, we propose the use of a sophisticated computer-based information system to generate tailored, interactive handouts to communicate with patients. Our system uses text planning and user modeling techniques to generate natural language descriptions of migraine, its symptoms, triggering factors and prescriptions. The system is capable of handling follow-up questions requesting further information, and generating responses in the context of previously supplied information – a capability unavailable in previous patient information systems. The system tailors its interaction to: (i) the class of migraine patients, (ii) the individual patient, and (iii) the previous dialogue. Preliminary evaluation of the system indicates that patients find it useful and informative. More extensive evaluation is in progress.

Introduction

Patient noncompliance is a significant problem with both economic and quality of life implications. A number of studies have shown that patient compliance is strongly correlated to the patients' understanding of their condition and prescribed treatment, e.g., [7, 11]. It is essential, therefore, that communication with patients be structured carefully to maximize the possibility of their comprehension. The importance of doing so is further underscored by research indicating that patients often do not understand important documents (such as 'informed-consent' forms) that they are asked to read and sign [6, 13]. Facilitating patient education is one way of overcoming these problems. However, doctors typically do not have large amounts of time to either explain their diagnostic procedures, or discuss their rationale for the prescribed therapy [10–12]. Unfortunately, handing out lengthy descriptions of different illnesses and possible treatments is unlikely to ameliorate the condition; a number of studies have shown that impersonal, decontextualized information has negligible impact on patients as compared to the selective presentation of relevant, patient-specific information, e.g., [2, 14]. It has also been observed that handouts are not very effective unless clinical personnel review them with patients to ensure comprehension and answer follow-up questions [12]. Therefore handouts can address only part of the problem. One possible solution is to use a computer-based system to select and present patient-specific information in an appropriate manner. This paper describes how a system can tailor the patient interaction – both the explanations generated, as well as the follow-up questions – to maximize patient education.

System Requirements

A patient education system must be able to:

- take into account patient characteristics and select relevant information,
- present the information in a manner both comprehensible and acceptable to the patient, and
- engage in follow-up dialogues, and facilitate queries for clarifications, elaborations, etc. in the context of the on-going conversation.

While a number of previous efforts have investigated the possibility of using computer-based information delivery systems, e.g., [1, 7, 9–11] none of these systems attempted to incorporate advanced natural language generation capabilities to tailor the interaction – not just the content, but also the organization and the phrasing of the explanations – to individual patients. Furthermore, they did not allow patients to query the system to receive further elaboration or clarifications.

We have designed a sophisticated information delivery system that meets these requirements. The system is based on a text planning framework [15] that dynamically selects and structures the information to be presented in each case. Since we are unable to parse free-form text, the system has a large library of possible questions[1]. The patient can ask questions by clicking on mouse sensitive portions of the text and by selecting available questions from menus [15]. Both the menus containing questions that the user may ask, and the responses produced, are generated in a context sensitive manner, taking into account both the patient model and the previous dialogue [5].

Conceptually, the migraine system [3] consists of three main components: (a) an interactive **history-taking module** that collects information from patients prior

The authors are listed in alphabetical order.

[1] The set of questions was determined from an ethnographic analysis of a large number of doctor-patient interactions; the field work and analysis was conducted by Forsythe, Brostoff and Bee [8].

to each visit and builds a patient model[2] (b) an intelligent **explanation module** that produces an interactive information sheet containing explanations in everyday language tailored to individual patients, and that responds intelligently to follow-up questions, and (c) an **interaction manager** that presents the interactive information sheet on the screen and manages the subsequent interaction with the patient.

In addition to presenting individual patients with information specific to their case, it is also important for the system to communicate information important to all migraine patients. For instance, ethnographic field studies have found that migraine sufferers are concerned about the possibility of more serious causes for their headaches, such as brain tumors or stroke [8]. The system is designed to address these concerns as well. Thus, the system tailors the explanation not only to the individual patient, but also to the larger class of migraine patients.

Patient Modeling in Migraine

The **interactive information sheet** provides patients with information about migraine, their specific condition, the prescribed treatment and its implications. The system can respond intelligently to follow-up questions about topics covered in the information sheet. The content and organization of the information sheet are based on three sources: (i) results of ethnographic research [8], (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 the ethnographic findings, the information sheet deals with global concerns before going on to instruct patients about how to follow their therapy regimen. However, the actual contents that are included when covering each topic will vary for each patient, because the system is sensitive to the data stored in the patient model. The information sheet produced by our system displays the following information:

- a statement about the diagnosis,
- an explanation of how the diagnosis accounts for the patient's symptoms,
- a discussion of situations and substances that act as migraine triggers for the patient, and information to help patients identify additional triggers,
- a set of instructions about the patient's individual treatment plan, and
- a discussion of the possible side effects of the prescribed therapy with instructions.

[2]The information collected here is also used to summarize the patients' status for their physicians.

Each of these points is tailored to the individual patient. A snapshot of the system is shown in Figure 1.

Tailoring the Output to Individual Patients

In order to adapt explanations to a patient, the explanation facility must have access to a stored description of relevant features: the patient model. Information about the patient that would typically be gathered using a questionnaire on their first visit – such as the symptoms, past treatments, relevant habits, and other current medical treatments – is gathered by a computerized version of such a questionnaire (the history-taking module) and translated into a patient model that can be used by the explanation component. 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. The system uses this patient model to generate customized explanations at every possible step. For example, the patient's symptoms, as well as specific supporting evidence, such as family history of migraine, are used to corroborate the physician's diagnosis:

If the patient's migraine attacks are characterized by a severe, throbbing headache, the system generates: *"One of the most common symptoms of migraine is a severe, throbbing headache … "*. If there were no *throbbing* sensation for another patient, the system would modify it to: *"One of the most common symptoms of migraine is a moderate to severe headache."* In cases of migraine characterized by the visual aura alone, the system can generate: *"A strong indicator of (classic) migraine is a disturbance in the visual field, known as an aura."*

The system uses information about the patient's physical examination to further reassure the patient that the condition is not a life-threatening one. The fourth paragraph in Figure 1 shows part of the description generated for a pre-menopausal, female patient. If the patient had been post-menopausal, the system would generate an explanation along the lines of, *"Migraine often gets better with age, since hormones play a role in making the attacks more severe"* without mentioning the role of menopause.

The description of migraine triggers is also tailored: it takes into account information about the gender and age of the patient, as well as any specific triggers reported by the patient:

"Migraine triggers … *you were not aware of any trigger factors, except possibly stress and sleeping late on weekends. Stress is difficult to avoid, but try maintaining a regular sleep pattern …*"

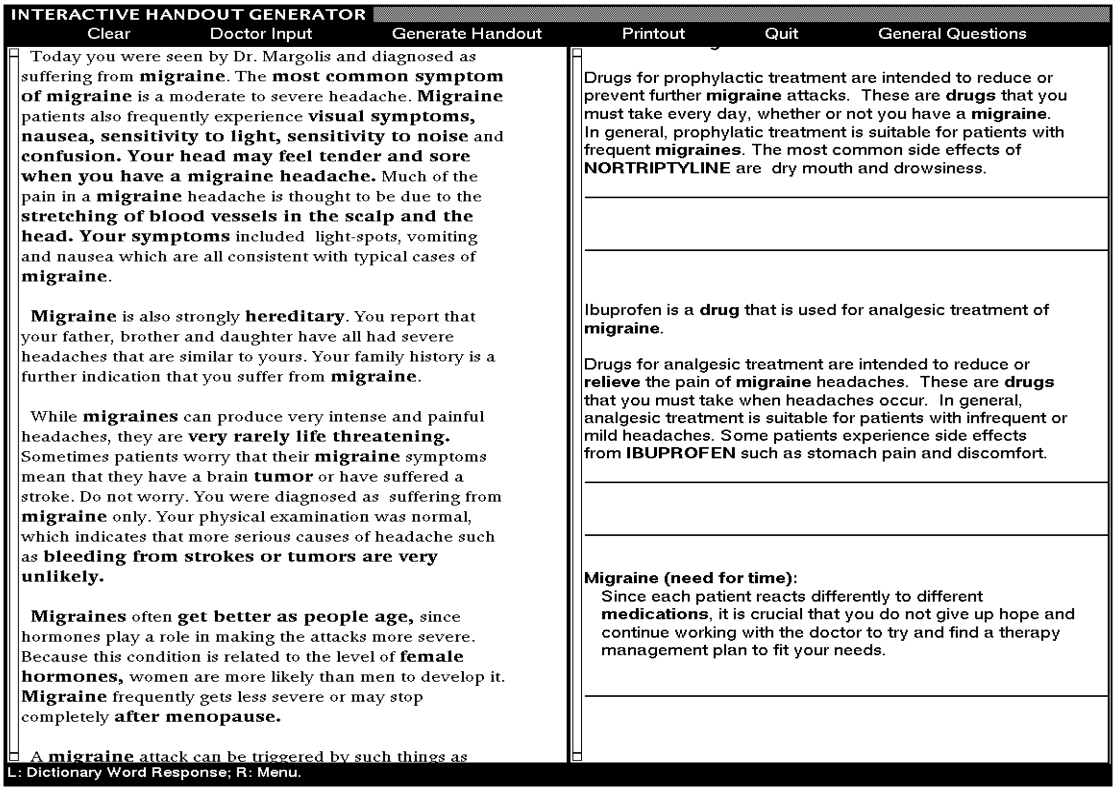

Figure 1: A snapshot of the system.

Similarly, the section addressing the prescribed treatment takes into account the drugs that are already being taken by the patient (for conditions other than migraine), as well as information about their possible effects on the patient's life style. For instance, among the side-effects of the drug Inderal are a reduced maximum heart rate and lowered arterial blood pressure. This side-effect does not noticeably affect patients with a sedentary lifestyle, but can result in individuals with an active life style no longer being able to indulge in strenuous activities to the same extent as before. The system elaborates on this side-effect of the drug differently depending on aspects of the patient's lifestyle, as determined during the history-taking session.

This scenario illustrates how patient information is used not only to determine *how* to phrase pertinent information, but also to determine *what* aspects are to be elaborated upon; if the patient specific information had not been taken into account, the handouts generated would have been quite different, with identical descriptions of the triggers, drugs and their side-effects, information on hormonal effects, and so forth, irrespective of the patients' age, gender, lifestyle, and other factors.

Tailoring the Output to the Patient Class

In addition to tailoring the description to individual patients, the generated handout also contains information that pertains to the general population of migraine patients as a whole. Ethnographic studies in the project have shown that migraineurs have both fears and concerns about their condition. For instance, migraineurs are usually worried about the possibilities of more serious problems such as cerebral haemorrhages, brain tumors or strokes as a result of the symptoms they experience [8]. It is important to address these fears explicitly in the information sheet. Migraine patients also worry about the dangers of taking hormonal medicines for other conditions (for instance, birth control pills), possible job discrimination, avoiding food/drink triggers in social situations, etc. These fears and concerns may be specific to migraine, and may not be applicable to other types of headaches (sinus, tension, cluster, etc). Initial results from the field suggest that different types of headaches cause patients to worry about different issues. Since our intention is to generalize the system to be able to deal with other headache types as well as migraine, issues such as these (pertaining to the general migraine population) must also be represented explicitly so that the system's coverage may be expanded without undue difficulty.

The result of this modeling of the patient class can be seen in the first three sentences of the third paragraph of the handout shown in Figure 1. Such information would typically be included for all migraine cases, but would change for patients with a different diagnosis.

Tailoring the Answers to Follow-Up Questions

The interactive information sheet allows users to click on terms in the explanation for elaborations or further information. (In Figure 1 such terms are indicated by the use of bold face font). If the patient were to click, for instance, on **symptoms**, the system would put up the following menu:

- What symptoms did I report?
- What are other possible symptoms?
- Were all my symptoms consistent with migraine?
- What could account for symptoms not consistent with migraine?

Note that the last question is offered only if there are any patient findings not commonly observed in migraine cases. In another case, if a patient were to click on the phrase "migraine often gets better as people age," the system would generate different menus depending on both the age and the gender of the patient (male patients would not be offered choices elaborating on the effect of menopause, birth control pills, etc.). Thus, both the information provided in response to follow-up questions, as well as choices for further information are tailored based on the patient model.

The System: Issues in Tailoring

A detailed description of the system is beyond the scope of this paper. A description of the overall system may be found in [3]; the text planning framework used in generation is described in [15]; our knowledge representation framework is described in [4]. In this section, we briefly highlight techniques necessary for tailoring explanations to the individual patient and discourse context.

The system uses a text planning framework to guide the presentation of information – we adopted this framework because it facilitates the fine-grained tailoring required in our system. 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. 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, the goal to generate the information sheet for the current patient-user is posted. 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.

```
(define-text-plan-operator
   :name Alleviate-fears-Female-Pre-Menopausal
   :effect (alleviate-fears ?patient (forever ?disease))
   :constraints
        ((female ?patient)
         (pre-menopausal? ?patient)
         (not (in-patient-history? estrogens)))
   :nucleus ((BEL ?patient (improve ?disease (after menopause)))
             (BEL ?patient (improve ?disease aging))))

(define-text-plan-operator
   :name Alleviate-fears-Male
   :effect (alleviate-fears ?patient (forever ?disease))
   :constraints (male ?patient)
   :nucleus (BEL ?patient (improve ?disease aging)))
```

Figure 2: Sample Plan Operators.

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. Second, operators contain 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. Figure 2 shows two (simplified) plan operators used in the system. The first operator applies only if the patient is female, premenopausal and does not have a history of estrogen drugs. The second operator satisfies the same goal as the first, but is only applicable to male patients. The explanations generated for the two cases differ substantially, with the first patient being provided information about the effect of menopause on migraines in addition to the information on aging. The system currently contains approximately 280 operators dealing with different situations. The majority of these operators refer to the patient model in their applicability constraints.

System Evaluation

The usability and utility of the migraine system has been evaluated in three preliminary studies. Two of these studies are relevant to the evaluation of the interactive information sheet. In the first study, 3 patients used the system in the context of an actual visit with a neurologist. In the second study, 13 persons with headache and one or more symptoms of migraine interacted with the history taking and the interactive information sheet without seeing the neurologist. In both of these studies, the patients were observed using the system, and were also interviewed afterwards regarding their session with the system. While we recognize that this is an evaluation of patients perceptions, and not a study of outcomes, we nevertheless believe that the results are helpful and encouraging. Table 1 shows an excerpt of patients' assessment of the interactive information sheet.

Questions	*Answer Category*	
	Yes	No
Did you like using the program?	16 (1.00)	0 (0.00)
Did all of the information presented make sense?	13 (0.81)	3 (0.19)
Did you feel comfortable about using a computer to get this kind of information?	16 (1.00)	0 (0.00)
Was the computer itself easy to use?	14 (0.88)	2 (0.12)
Did the program tell you anything you did not already know?	15 (0.94)	1 (0.06)
Do you think this information will help you manage your headaches better?	9 (0.56)	7 (0.44)
Did you learn anything that you would not have asked your doctor?	12 (0.75)	4 (0.25)

Table 1: Responses on the Interactive information sheet Follow-up Interview.[3]

Conclusions and Future Work

Studies have shown that patient compliance with the prescribed therapy significantly increases with understanding. We have demonstrated how a text planning framework can be used to automatically synthesize tailored, interactive information sheets for migraine patients. In generating these information sheets, the system takes into account characteristics of both the individual patient as well as the general class of migraine patients. In addition, the system takes the patient model into account in dynamically generating menus of follow-up questions, as well as answers to them.

The usability and utility of the MIGRAINE system has been evaluated in three preliminary studies. These studies were aimed at assessing patient and physician perceptions about the system. While the results from these studies have been encouraging, we are aware that a more thorough and objective evaluation is necessary.

Finally, we also plan to extend the coverage of the system to handle other types of chronic headaches in addition to migraine. Since the knowledge base and text plan operators used in the system were designed in a modular fashion, we expect that a reasonable percentage of these knowledge sources can be re-used in a larger patient education system.

Acknowledgments

This work is supported by grant number R01 LM05299 from the National Library of Medicine, NIH. Its contents are solely the responsibility of the authors and do not represent the official views of the NLM. We gratefully acknowledge the contributions of the other members of the Migraine project: G. Banks, N. Bee, M. Brostoff, B. Buchanan, D. Forsythe, S. Ohlsson and E. Rees. The ethnographic results that the paper draws upon are due to field-work and analyses by Forsythe and Brostoff.

[3] The numbers in parentheses represent the percentage of respondents who gave a particular answer.

Reference

[1] ARMSTRONG, M. L. Orchestrating the process of patient education. Methods and approaches. *Nursing Clinics Of North America 24*, 3 (Sept. 1989), 597–604.

[2] AUERBACH, S. M., MARTELLI, M. F., AND MERCURI, L. G. Anxiety, information, interpersonal impacts, and adjustment to a stressful health care situation. *Journal Of Personality And Social Psychology 44*, 6 (June 1983), 1284–96.

[3] BUCHANAN, B. G., MOORE, J., FORSYTHE, D., CARENINI, G., BANKS, G., AND OHLSSON, S. Using Medical Informatics for Explanation in the Clinical Setting. Tech. Rep. Number CS-93-16, Computer Science Dept., Univ. of Pittsburgh, 1994.

[4] CARENINI, G., AND MOORE, J. D. Using the UMLS semantic network as a basis for constructing a terminological knowledge base. In *Proceedings of SCAMC-93* (Washington, D. C., 1993).

[5] CARENINI, G., AND MOORE, J. D. Generating Explanations in Context. In *Proceedings of the International Workshop on Intelligent User Interfaces* (Orlando, Florida, Jan. 1993), pp. 175–182.

[6] CASSILETH, B. R., ZUPKIS, R. V., SUTTON-SMITH, K., AND MARCH, V. Informed Consent: Why are its goals imperfectly realized? *New England Journal of Medicine 302* (1980), 896–900.

[7] FISHER, R. C. Patient education and compliance: a pharmacist's perspective. *Patient Education and Counseling 19*, 3 (June 1992), 261–71.

[8] FORSYTHE, D. E. Using Ethnography in the Design of an Explanation System. *Expert Systems with Applications (in press)* (1994).

[9] GILLISPIE, M. A., AND ELLIS, L. B. Computer-based patient education revisited. *Journal Of Medical Systems 17*, 3 (Aug. 1993), 119–125.

[10] JONES, R. B., MCGHEE, S. M., AND MCGHEE, D. Patient on-line access to medical records in general practice. *Health Bulletin 50*, 2 (Mar. 1992), 143–50.

[11] KAHN, G. Computer-Based Patient Education: A Progress Report. *M.D. Computing 10*, 2 (1993), 93–99.

[12] KAHN, G. Computer-Generated Patient Handouts. *M.D. Computing 10*, 3 (1993), 157–164.

[13] LEEB, D., BOWERS, JR., D. G., AND LYNCH, J. B. Observations on the Myth of Informed Consent. *Plastic Reconstructive Surgery 58* (1976)

[14] MARSHALL, W. R., ROTHENBERGER, L. A., AND BUNNELL, S. L. The efficacy of personalized audiovisual patient-education materials. *Journal Of Family Practice 19*, 5 (Nov. 1984), 659–63.

[15] MOORE, J. D. *Participating in Explanatory Dialogues: Interpreting and Responding to Questions in Context.* Cambridge, MA: MIT Press., 1994.