Title: An On-line Explanation of the Medical Record to Patients via an Artificial Intelligence Approach

Authors: Alison J Cawsey (1) Kim A Binsted (1) Ray B Jones (2)

1. Department of Computing Science, University of Glasgow, Glasgow G12 8QQ
2. Department of Public Health, University of Glasgow, Glasgow G12 8QQ

Contact: Alison Cawsey
Phone: 041-339 8855 x 6039
Fax: 041-330-4913
Email: alison@uk.ac.gla.dcs

Word Count: 2142

An On-line Explanation of the Medical Record to Patients via an Artificial Intelligence Approach

Alison J Cawsey (1) Kim A Binsted (1) Ray B Jones (2)
1. Department of Computing Science, University of Glasgow, Glasgow G12 8QQ
2. Department of Public Health, University of Glasgow, Glasgow G12 8QQ
Email: alison@uk.ac.gla.dcs

Introduction

The importance of providing the public with more information about health issues is widely accepted. Improved patient education may lead to reduced hospital admissions and length of stay, less use of emergency facilities, and a better quality of life for the patients concerned [1,2,3]. Computers have been successfully used in patient education [3,4,5], complementing traditional means such as doctor-patient dialogue, written leaflets and audio-visual materials [6].

One problem, however, with existing approaches is that the information provided is over-general, not targetted at the particular patient's problems and interests [6]. One way around this is to make the starting point of patient education the medical record. Patients may be provided with on-line access to their record, with explanations available of terms in that record. This approach combines the advantages of a more personalised type of patient education with the benefits (and legal requirements) of a patient accessible record [7].

In two earlier studies we have demonstrated the potential of this general approach [8, 9]. In the first [8], 70 patients in a general practice were given the opportunity to use a system which gave both general health information and access to their records (along with further explanations of problems in their record). 65 patients used the system, and 52 of these chose to look at their record. The system was popular, with 84% of patients saying that they would use it again, and the explanations of medical terms appeared important: more than one in four of the medical problems were not understood by a patient until further explanation was given.

In this first study the on-line records were developed from uncoded manual general practice records. However, for the approach to be practicable on a large scale, we need to link the explanatory material to existing computerised clinical records. In the second study [9], and the work reported here, we used Read coded clinical records. Twelve patients at an outpatient clinic for hypertension were given the opprtunity to look at their record. They were first provided with a list of problems (using the Read "preferred term" for each problem), with the opportunity to obtain a "lay translation" of that problem, and further screens of explanatory material. In this study, 10 out of 12 patients would probably or definitely use the system again, and explanations were sought of 23 out of the 30 medical terms used in the records.

The authoring of the explanatory material for this second study was quite time consuming, even for only 30 medical terms. In our current work we are exploring how this can be partially automated, given a "knowledge base" loosely based on the Read codes. We believe that by partially automating the generation of the explanations, we can provide a more consistent, easily maintainable and adaptable system.

An Artificial Intelligence Approach

We use an approach based on "text generation" techniques developed in computational linguistics and artificial intelligence (AI). These techniques have been used in research prototypes for automatically producing on-line tutorials [10], database summaries [11], technical documentation [12] and reports [13]. The methods are just beginning to find commercial application.

Recently two (other) prototype patient education systems have been developed [14, 15]. The first [14] generates on-line explanations for migraine sufferers (not linked to their record), while the second [15] generates written explanatory material to accompany drug prescriptions. Our work is complementary to these, with more emphasis on the medical record and on empirical studies of system use.

Essentially the idea in all these systems is to provide general rules capturing how different sorts of objects are typically described. These rules are normally derived by analysing texts in the chosen domain, and/or through empirical studies of potential users' information needs [14]. In our work we we examined health promotion booklets and existing literature on patients' information needs.

Example (somewhat simplified) rules for describing a problem and a treatment respectively might be:

```
To EXPLAIN a PROBLEM:
   Give its lay term
   Mention possible causes
   Give its symptoms
   List some possible treatments.

TO EXPLAIN a TREATMENT:
   Give its lay term
   Say what its used to treat
   Give possible contraindications
   Give possible side-effects.
```

Rules can, of course, be much more complex than this. For example, they can state what to include depending on whether or not the patient has suffered from the problem (or takes the treatment). Rules can also reference a list which states which medical terms are likely not to be understood (by a typical patient), so that such terms can be immediately elaborated. More specialised rules can also be given stating what additional material to include for particular types of problem and treatment, and more detailed rules state how to give the lay term, describe the type of problem and so on.

The system can generate a specification of the explanation screen for a particular issue. This may include layout instructions and specifications of *buttons*. These buttons allow words to be clicked on or touched by the user, and further explanations given as a result (as in conventional hypertext). The explanation screen specification is platform independent, and can either be used by our own display program to produce the final interface, or translated into a standard markup language such as Hypertext Markup Language (HTML) for use with other display systems.

The information for a particular explanation (e.g., the lay term, treatments etc of a specific problem) is extracted from an object-oriented knowledge base of medical terminology. The structure of this knowledge is (roughly) taken from the Read codes, but it is augmented with additional information as required for patient-oriented explanations. The explanation rules are also implemented using an object oriented approach, with methods stating how to explain different classes of medical term.

As suggested in the introduction, we believe that our approach has a number of possible advantages:

- It is easy to extend the system to deal with new problems and treatments, and to revise the system when recommendations change. A simple factual addition/change needs to be made in the knowledge-base, rather than many texts searched and revised and new explanations carefully devised.
- The explanations provided are of a consistent style (rather than confusing the patient with an explanation that accidentally omits information of a sort available in another).
- The system can easily be adapted, to consistently change all the texts to be of a different style (e.g., removing any mention of drug side-effects from all explanations). A particular clinic could adapt the system to include just what type of information they believed appropriate.
- Explanations can be tailored to the user. For example, pertinent reminders of facts in their record can be included (e.g., "Hypertension is more likely if you are overweight. You were last weighed at 20 stone").
- The costs of producing explanations in other languages is reduced. The knowledge base and the rules can be made language independent, with a single module for translating the output of the system into a specific language.

However, these advantages have to be weighed against the costs of the initial authoring of the knowledge base, and the possibility that the "generated" texts may miss the subtlety of explanations that are hand-crafted for every separate issue [16].

Examples

The system was prototyped using anonymised records from a clinical information system for diabetes in Nottingham [17]. These records have a problem list, treatment list (both diabetic and other) and a clinical flowsheet. These original records were not Read coded, but have been translated into a Read coded form for the purpose of this current study. The examples given below are based on these anonymised records.

A patient using our system would be given, following introductory screens, the opportunity to browse through any of these three parts of their record. If they chose "problems" they might get a screen containing the following text:

PROBLEMS ON RECORD

- **Type one diabetes mellitus:** Diagnosed: 1965. Still active.
- **Angina pectoris:** Diagnosed: 1975. Still active.
- **Background diabetic retinopathy:** Diagnosed: 1988. Still active.
- **Nephrotic syndrome:** Diagnosed: 1988. Still active.
- **Hyperlipidaemia:** Diagnosed: 1992. Still active.

Each of the underlined bold face words are buttons that can be touched or clicked on to obtain further information. Suppose **hyperlipidaemia** was selected. The following screen of information would be displayed:

HYPERLIPIDAEMIA:

"Hyperlipidaemia" is the term for having a high level of lipids (fats), particularly cholesterol, in the blood. According to your **record**, you have this problem. Possible causes include **diabetes mellitus** and **alcoholism**. Some common treatments include a special diet and a **lipid-lowering drug**. It increases the risk of **atherosclerosis** and **coronary artery** disease. **Your cholesterol level** was last measured at 3.8 mmol/L (15 July 1994). You have a lipid lowering drug in your record: **bezafibrate**.

Note how the user is given both general information, and reminders of specific information in their record relevant to this problem. Now, if the user next touched the word "bezafibrate", they would obtain the following explanation.

BEZAFIBRATE:

Bezafibrate is a cardiovascular drug which reduces the amount of some kinds of fat in the bloodstream. According to your record, you are currently undergoing this treatment. It is often used to treat **hyperlipidaemia**. This could have some side-effects; in particular, **nausea**. Your prescription of bezafibrate comes in 200 mg tablets. It is to be taken three times each day (morning, afternoon, and tea time).

Again, note how the user is both given general information about the drug, and personalised reminders of how they should take the drug. Both these explanations (hypelipidaemia and bezafibrate) are produced from an underlying knowledge base using rules similar to the two given in the previous section.

If, following the explanation of hyperlipidaemia, the patient had clicked on "cholesterol level", the following explanation would be given, again combining general explanations with reminders of patient specific data.

CHOLESTEROL LEVEL:

A patient's cholesterol level (measured with a blood test) is an indication of the amount of fatty food in their diet. It can be influenced by heredity and diseases such as diabetes. A normal level is 2-4 mmol/L. Your cholesterol level was last measured at 3.8 mmol/L (15 July 1994). Important: A high cholesterol level (greater than 6.5 mmol/L) increases the risk of **atherosclerosis** and **heart disease.**

In all these explanations, buttons are available at the bottom of the screen allowing the patient to obtain further help; go back to previous explanations; obtain a list of written materials they may consult; or note that this is an issue they wish to discuss with their doctor. Figure 1 illustrates the screen layout in our current prototype. At the end of the session the user can obtain a printout of the session and a list of the issues marked for discussion with the doctor.

Preliminary Evaluation

In evaluating our system we aim to assess both whether the approach has benefits for patients (compared with written educational materials and online material which is not linked to the record) and whether the particular computational techniques used are effective. We are currently part way through an initial study in three diabetes clinics in the Glasgow area. We are using structured interview techniques with both doctors and patients in the clinics, to obtain a qualitative assessment of what we

Figure 1: A sample explanation screen

regard as the critical features of the system (text clarity and relevance; ease of use; benefits of personalised explanations).

So far we have shown the system to three groups of professionals (in total 5 diabetes consultants, 1 diabetes nurse specialist and 1 junior doctor). They were shown the system based on the anonymised records from Nottingham, and asked to comment on a number of features. In summary they thought that the explanations generated, although sometimes overtechnical, might be of use to patients. They particularly appreciated the fact that patients would be reminded throughout of their particular measurements and treatments, and that they would be able to mark issues to discuss with their doctor. A typical comment was that "this kind of [system] would allow patients to make better use of their time with me". All three clinics have given permission for a small study with their patients in September 1994. For this study we will add the volunteer's patient records to the system, in Read coded format. We will report on this in the presented paper.

Conclusion

The use of an artificial intelligence approach would seem to offer an effective method of constructing patient explanations, and deserves further study. The effort in extending the system for new problem areas is somewhat reduced, and the explanations, although simple, appear to be of use, especially when augmented with patient-specific information. Although the information in the current explanations is fairly restricted, our approach makes it very simple to consistently add to or change the sort of information given in the explanations, if our initial studies with patients suggest that this is required.

We now hope to obtain funding for further studies, to see if there are differences in patients' reactions between an education system based on the medical record and 'standard' educational material, and between personalised explanations produced using an AI approach and those produced using a more conventional "cut and paste" approach.

Acknowledgements

We would like to thank Glasgow University for providing the intitial funding for this project, and Drs Fisher, Kesson, O'Donnelly, Paterson and Semple and their patients for their valuable feedback.

References

1. Morris J, Goddard M, Roger D. The Benefits of Providing Information to Patients. York, 1889; Discussion paper 58, Centre for Health Economics, University of York.
2. Osman LM, Abdulla MI, Beattie JAG et al. Reducing hospital admission through computer-supported education for asthma patients. British Medical Journal 1994; 308: 568-571
3. Lyons C, Krasnowski J, Greentein A, et al. Interaction computerized patient education. Heart and Lung 1982; 11:340-341
4. Gillispie MA, Ellis LBM. Computer-based patient education revisited. Journal of Medical Systems 1993; 17: 119-125.
5. Jones RB, Campbell G, Navin LM, Ritchie J, Kelly MPT. Touch-screen multi-media systems for patients and health education. In Proceedings of the Second Asia-Pacific Medical Informatics Conference. Hong Kong, 1992; 311-318.
6. Marshall WR, Rotherberger LA, Bunnell SL. The efficacy of personalized audiovisual patient-education materials. Journal of Family Practice 1984; 19: 659-663.
7. Gilhooly MLM, McGhee SM. Medical records: practicalities and principles of patient possesssion. J Med Ethics 1991; 17:138-143.
8. Jones RB, McGhee SM, McGhee D. Patient on-line access to medical records in general practice. Health Bulletin 1992; 50:143-150.
9. Jones RB, Sandham P. A lay axis to the Read Codes?. Br J Healthcare Computing 1994; 11: 30-31.
10. Cawsey A. Explanation and interaction: the computer generation of explanatory dialogues. Cambridge Massachusetts: The MIT Press, 1993.
11. McKeown KR. Text generation: Using discourse strategies and focus constraints to generate natural language text. London: Cambridge University Press, 1985.
12. Reiter E, Mellish C, Levine J. Automatic Generation of On-Line Documentation in the IDAS Project. In Proceedings of the Third Conference on Applied Natural Language Processing, 1992; 64-71. Trento, Italy.
13. Rambow O, Koresky T. Applied Text Generation. In Proceedings of the Third Conference on Applied Natural Language Processing, 1992; 40-47. Trento, Italy.
14. Buchanan BG, Moore J, Forsyth D, Banks G, Ohlsson S. Involving patients in health care: using medical informatics for explanation in the clinical setting. In Proceedings of the 16th symposium on computer applications in medical care, 1992; 510-514. Washington D.C.
15. de Rosis F, Grasso F. Mediating between hearer's and speaker's views in the generation of adaptive explanations. Expert systems with applications. In press.
16. Reiter E, Mellish C. Optimizing the costs and benefits of natural language generation. In Proceedings of the thirteenth international conference on artificial intelligence, Chambery, France. San Mateo California: Morgan Kaufmann, 1993; 1164-1169.
17. Jones RB, Hedley AJ, Peacock I, Allison SP, Tattersall RB. A computer-assisted register and information system for diabetes. Methods of Information in Medicine, 1983; 22:4-14.