Generating Personalised Patient Information Using the Medical Record

Kim Binsted[1⋆], Alison Cawsey[1⋆⋆] and Ray Jones

[1] Department of Computing Science, University of Glasgow
[2] Department of Public Health, University of Glasgow

Abstract. This paper presents an approach for providing patients with personalised explanations of their medical record. Simple text planning techniques are used to construct relevant explanations based on information in the record and information in a general medical knowledge base. We discuss the results of the evaluation of our system with diabetes patients at three diabetes clinics in Scotland.

1 Introduction

Patients in Britain now have the legal right to access to an explained version of their medical record. This right, however, is rarely exercised. Heavy demands on doctors' time mean that they seldom have time to explain even the most important aspects of a patient's record in detail; and because medical records are complex documents, often in a paper file containing a semi-organised jumble of letters, hand-written notes, and uninterpreted results, they are useless and even dangerous without such explanation.

This situation is especially unfortunate for patients with chronic problems. There is considerable evidence that providing personalised information about a patient's condition can significantly improve the treatment of some chronic problems. A study by Osman et al (1994) involving 801 asthma patients in Scotland showed that, by providing asthma patients with personalised booklets about their condition and treatment, hospital admissions were reduced by 54%. Moreover, patients with chronic problems tend to have complex records, with long lists of complications and treatments, making explanation all the more necessary (see Berry et al (1994) for an analysis of patient information needs).

Medical records are increasingly stored in electronic form. The problem of providing explained access to such records seems ideally suited to the techniques of automatic explanation generation, a subfield of artificial intelligence. By 'automatic explanation generation' we mean the production, by a computer, of a piece of natural language text on a subject which allows the reader to understand that subject more clearly.

Our goal in this project was therefore to build an interactive system which generates natural language explanations of items in patient records, personalised

⋆ E-mail address is kimb@aisb.ed.ac.uk. More information on the work presented here is available from the WWW page http://www.dcs.gla.ac.uk/~www/Piglet.html.
⋆⋆ E-mail address is alison@dcs.gla.ac.uk.

to the patient in question. The explanations produced by the system are brief and simple, so that the patient is not put off by too much complex information; yet all the medical terms in the explanations are 'click-able' (i.e. the explanations are in hypertext), so that the patient can seek more information if desired. Both the content and the form of the explanation are tailored to suit the needs of the patient.

Such a system has a role both as general medical information tailored to the needs of a particular patient, and as a way of providing explained access to the patient's medical record. We would anticipate that such a system would be of most use in a waiting room, to be used before a session with the doctor. Here it might serve as a "warm up" for the consultation, refamiliarising patients with their records, giving them basic relevant information, and reminding them of any questions they might have.

We believe that using techniques to generate an explanation *automatically*, from a simple medical knowledge base and the medical record, has the following possible advantages:

- Because the explanations are generated from information in both the patient record and in a general medical knowledge base, they change automatically as the information in the record and in the knowledge base is updated.
- The style of explanation can be adjusted, according to the preferences of the doctor and of the patient, by changing the text plans used to generate the text – individual explanations need not be rewritten. Moreover, the style of explanation will be consistent over the whole session.
- Unlike with more general patient education materials, information particularly relevant to this patient can be emphasised (encouraging patient compliance), and irrelevant information minimised.
- There is the possibility of multilingual generation, for patients who do not share a first language with the medical staff. The knowledge base, patient record, and content-specifying rules can be made language independent, with a single module for translating the output into a specific language (see Rosner (1994) for a discussion of multilingual explanation generation).

We realised, however, that certain potential disadvantages of automatic explanation generation had to be avoided. In this kind of application, there is no room for unclear or misleading information. It is essential that a medical expert be able to check through the generated explanations – that is, the explanations must not vary in unpredictable ways. The system must not 'put words in the doctor's mouth' by adding information to the record, either explicitly or implicitly. Finally, the knowledge base, the record, and the text plans used to generate the explanations must all be easily modified by medical personnel.

2 Related Work

There is a significant amount of past research demonstrating the acceptability of computer-based patient education and the efficacy in particular of personalised patient education. Jones et al (1992), for example, have demonstrated the

feasibility of giving patients access to general health information and to their medical record (and associated hand-crafted explanations). By allowing patients direct access to their records, through a suitably simple interface, patients can obtain information that is (partially) selected by them and not by the system, thus hopefully obtaining more relevant and appropriate advice. Recent studies involving 70 people, allowed to access information in their medical record (and associated simple fixed explanations), showed that 84% found it helpful and would use such a system again (Jones 1992). However, there is no attempt to personalise the details of the information presented to the user.

Osman et al (1994) gave patients in their study personalised computer-generated booklets, based on information in the patient record and on responses to patient questionnaires. They showed that this information could lead to a dramatic reduction in hospital admission rates. However, the techniques for selecting and organising the information were primitive, based on the use of spreadsheets and mail-merge software. The patient did not interact directly with the computer system to obtain the information. There are clearly potential advantages with interactive systems, and more sophisticated generation techniques.

There has been only limited work on using explanation generation techniques in this area. De Rosis et al (1994) generate explanations of drug specifications, to be printed out and given to the patient. The explanations are modified according to the prescriber's preferences, rather than the patient's, but text generation techniques are used to ensure that the explanation is not stilted. Because the explanation is printed out, the patient cannot interact with the system directly.

The most directly relevant related research is the MIGRAINE project (Buchanan et al 1992, Carenini et al 1993, and Buchanan and Moore 1994). The MIGRAINE system generates explanations of migraine, using the patient's medical history (also taken by the computer, as MIGRAINE does not have access to the patient record) to personalise the text to the patient's needs. The system uses a slightly more sophisticaed text generation strategy than in our work. For example, extensive use is made of the discourse history (the record of the 'conversation' between the system and the user) to modify the text plans chosen for generation. A more general theory of discourse structure is used as the basis for an explanation planner, and the system has a capacity for re-explaining things that weren't understood first time.

We have taken a more pragmatic approach in our work. We felt, and our evaluation indicated, that patients would be uncomfortable with medical information that is in a significantly different form each time it is presented, so we don't attempt to make (much) use of the discourse history. Our text planner uses domain specific text plans for different broad classes of medical issue (e.g., problem vs. treatment). We don't provide for issues to be re-explained another way by the system, but have a facility for patients to mark issues to be discussed further with their doctor. Our system is somewhat broader in scope than MIGRAINE, more closely linked with the medical record, and with its develoment driven by repeated evaluation and feedback from doctors and patients.

3 System Architecture

PIGLET (Personalised Intelligent Generator of Little Explanatory Texts) has four main parts: a set of electronic patient records, a general medical knowledge base, a text generator, and a display module. The text generator uses the information in the knowledge base to provide the specifications for a hypertext explanation of the patient record, and the display module uses this specification to display the hypertext on the screen.

The initial form and content of the electronic records was taken from a set of Nottingham Diabetes Clinic records (Jones 1983). Each record was made up of:

- A set of personal information (name, age, address, occupation etc.)
- A list of problems, each with date of diagnosis and current activity
- A list of treatments, each with start and end date and dosage (for medication)
- A list of tests and measurements, each with date and result

Although medical records generally contain much more information than listed above, this is a reasonable minimum amount of information to expect from an electronic summary. Because no such computer records were available from local diabetes clinics, PIGLET's medical records were hand-built from real paper records from diabetes clinics in the Glasgow area; however, the system should work equally well with any computerised record system which can provide the above information about the patient.

PIGLET's medical knowledge base is an 'isa' hierarchy of objects, with multiple inheritance of slot values from ancestors. The structure of this hierarchy is loosely based on that of the Read codes (the standard method of encoding medical concepts in Britain), and each object is named according to the Read system; however, since the Read hierarchy is not a strict 'isa' hierarchy (e.g. complications of diabetes are given as subtypes of diabetes) and does not allow for multiple ancestors, we have modified the hierarchy as necessary to suit our purposes. Again, although this knowledge base was handcrafted, we believe that most of the information PIGLET requires is of a general enough nature that it should be available from other general-purpose medical knowledge bases. See figure 1 for a small part of the hierarchy.

The text generator uses domain-dependent text plans, as used in Cawsey (1992). A text plan is associated with the class(es) of objects, and represents the way the objects are explained. Each text plan is used to generate a list of text specifications (the explanation of an issue), which are then translated by the display module into click-able hypertext on the screen. Text plans are defined in terms of the class of object they apply to, preconditions concerning when they apply, and subgoals. These subgoals can refer to lower level subplans, and thus dictate the structure and content of the text, or be actual text templates, specifying the actual text or hypertext to be used to realise a specific goal. This is an example text plan which illustrates the use of templates to construct a sentence within an explanation.

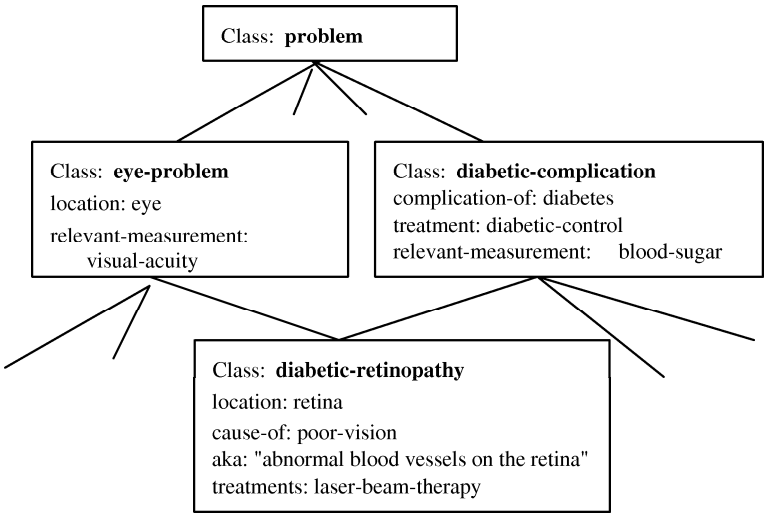

Fig. 1. A small subsection of the knowledge base.

```
(defplan patient-measurement (i)
  :class measurement
  :preconds (in-record i)
  :goals ((template 'plain "Your")
          (template 'hypertext (name i) (do-window i))
          (template 'plain "was last measured at")
          (measurement-value (in-record i))
          (measurement-date (in-record i))))
```

This plan will apply only to a measurement which is in the patient record, and will generate (as part of a larger explanation) the specifications of a sentence giving the measurement value and date. For example, it could generate:

Your cholesterol level was last measured at 3.8 mmol/L (15 July 1994).

The phrase 'cholesterol level' can be clicked on for further explanation.

There can be several text plans of the same name. Which is applied first is determined by the class of the object (PIGLET attempts to apply the most specific plan first) and the number of preconditions (PIGLET attempts to apply the plan with the most preconditions first).

The display module takes a list of such text specifications and turns it into readable text on the screen, in accordance with both the generated specifications and a set of machine specific parameters. These text specifications can, of course, be used by other display modules to display the text in different ways. For example, we have also implemented a piglet2html translator, which turns a set

of PIGLET text specifications into a suitable html file, to be displayed with a WWW browser such as Mosaic.

4 Personalising the Explanations

The information provided in the system is made relevant to the individual user in several ways: the medical record provides a list of initial issues to explore, the user selects which of these to have explained, and the system provides an explanation which is tailored to that individual given broad assumptions about their likely knowledge and concerns.

By making the medical record the starting point of the interaction, relevant information is easily accessible to the user. This contrasts with general health information systems or medical encyclopaedias, where users may have to descend many levels of menus etc. to find information relevant to them. In our system, the user starts from a hypertext version of their record, and can click on any of the issues in that record to obtain further information. Other, less directly relevant medical issues will only be available to be explained if they are related in some way to the user's own medical problems and treatments.

Within this context, the user can select for themselves the issues they want to have explained, by clicking on the relevant hypertext words. This contrasts with written information or more rigid computer-based patient education materials where the user may be taken through a fixed series of screens, which may or may not be relevant.

More interestingly from an AI point of view, the way a selected issue is explained will also depend on the user. The detailed content of the explanation is determined by the user's likely concerns (as inferred from their record), while the way the information is conveyed (in terms of relations to related objects etc.) is influenced by a simple model of what different categories of user are likely to know about.

The information content is influenced by the user's concerns as follows. First, we assume that the user is concerned about all the issues (e.g., problems, treatments etc.) mentioned in their record. We then use that information to decide when to add additional personal reminders or elaboration concerning these key issues whenever we are explaining a related issue. For example, when explaining what humulin is (a kind of insulin) we might remind the user about the type of insulin they are taking, and remind them what their blood sugar level is. When describing a particular drug, if the patient is taking this drug, we remind them of exactly how the drug should be taken (e.g., one tablet, three times a day, with food), and mention any possible side-effects or warning signs.

The way something is explained is influenced by a simple user model indicating, for different categories of user, which medical issues are likely to be familiar. User categories are currently based on the user's major chronic problem, rather than, say, education level, as we felt that knowledge of medical issues depended primarily on whether users had direct experience of the disorders etc. in question. So, we represent which issues a typical diabetes patient might know about (e.g.,

insulin), and which issues all patients should know about (tablet, pain etc.). As we expand the system to deal with other problems we will extend this model.

When selecting between different possible ways of describing an object, we aim to find a description that is given in terms of related known issues (according to the simple user model). For example, to describe a issue, one method first identifies it as an instance of a general class. The general class should be a known issue for the particular class of user (diabetes patient). This method might generate:

Humulin is a kind of insulin.

since "insulin" is a term familiar to all diabetes patients, but it would not generate:

Amlodipine is a coronary vasodilator.

since "coronary vasodilator" is not considered to be an understood term. This restriction prevents the 'dictionary problem' — the uselessness of explaining an unfamiliar issue only in terms of other unfamiliar issues — by ensuring that at least some of the terms in an explanation are commonly understood.

Although PIGLET makes minor use of the discourse history to make explanations flow more smoothly, a term is not assumed to be understood just because it has been explained before. We felt that patients would be disturbed by explanations that changed radically on reviewing, and wanted explanations of any given subject to remain fairly static over the course of a session (although not necessarily between sessions or between patients).

In summary, information in PIGLET is made relevant by: making the user's own patient record the starting point of the interaction; allowing the user to select what they want to see, through a hypertext style interface; adapting the detailed content of the explanations given information in the record; and adapting the way something is explained given a model of issues the user is likely to know about. No existing system combines all these features.

5 Example

A patient using our system would be given, following introductory screens, the opportunity to browse through their personal information, their problem list, their treatment list, or their recent tests and measurements. If they chose problems, they might get a screen containing the following text:

PROBLEMS ON RECORD
Type one diabetes mellitus Diagnosed: 1965.
Angina pectoris Diagnosed: 1975.
Background diabetic retinopathy Diagnosed: 1988.
Nephrotic syndrome Diagnosed: 1988.
Hyperlipidaemia Diagnosed: 1992.

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.

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. There are also buttons which will display a relevant picture, or jump to a relevant topic not mentioned in the text. 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.

6 Evaluation

The evaluation of this initial prototype of PIGLET was in three stages, which should be seen as the first part of an ongoing cycle of evaluation and improvement.

Once we had developed a simple working version of the program, we invited two public health students and one general practitioner to test it, and respond to a general questionnaire. Their comments included the suggestions that: the vocabulary be simplified, and with more of the explanations given in terms of generally understood concepts; some way of marking an explanation for further discussion with the doctor be provided; sensitive information (such as potential drug interactions) be minimised or omitted, to be left to the doctor to explain; and that buttons for putting up simple, relevant pictures (such as a picture of the heart in an explanation of heart disease) be provided.

In general, they were all positive about PIGLET's potential for improving patient–doctor communication and general patient awareness. They all commented that the personalisation of the explanations, particularly the inclusion of relevant measurements from the record in problem explanations, would help patients understand how the information provided relates to them.

Since their suggestions were in line with the goals of the project, we modified the program to incorporate them. Once these changes had been implemented, we approached diabetes professionals with the system.

PIGLET was shown to seven diabetes professionals (five diabetes consultants, one diabetes nurse specialist, and one junior doctor) at three clinics in the Glasgow area. At this stage, we had two hopes: that the medical professionals would offer advice on improvements, and that they would allow us to show PIGLET (with their suggested improvements) to some of their patients. Because of constraints on the health professionals' time, PIGLET was shown to groups of two or three people at a time, and their responses to a questionnaire were recorded on audio tape. Each group was asked questions concerning the hypertext interface, the content and style of the explanations, the personalisation of the explanations, and the general usefulness of the system. They were all positive about the system; in particular:

Interface: All thought that most of their patients would be able to use the hypertext interface, provided that a touch screen was used, although one commented that older patients might not want to use a computer.

Content and Style: Although they all agreed that the explanations were clear and fairly easy to understand, two of the consultants thought the vocabulary level should be reduced still further. All said that the explanations were of about the right length and content.

Personalisation: All thought that the personalisation of the information, particularly the way relevant measurements were mentioned, would be useful: "Yes, that's very good — that makes it relevant to themselves, rather than just general information about their condition." Several ways of increasing

the personalisation were suggested, including the generation of charts showing the patient's position (e.g. in weight, cholesterol level, etc.) with a subset of the rest of the population.

General usefulness: All thought that, despite the large amount of diabetes patient education material, PIGLET would still be useful: "They tend to skim over brochures etc., but would probably take in more of this because it's directed at them." They were all enthusiastic about the idea of PIGLET as a 'warm-up' for a consultation, allowing patients to mark issues which they wished to discuss further. A typical comment was "this kind of [system] would allow patients to make better use of their time with me." Some worries about confidentiality and security were expressed, but all agreed that some sort of card system (similar to that used with banking machines) would provide adequate protection. All said that PIGLET would help correct errors in their records.

All of the clinics agreed to let us have access to some patient records, so that PIGLET could be evaluated with the patients themselves. We entered eleven diabetes patient records in total, and ten of the patients agreed to try out the system. There were 6 men and 4 women, and ages ranged from 24 to 65. Each was given approximately half an hour to use the system to look at, and have explained, their own record. They were then asked questions about the system, similar to those asked of the doctors, in a structured interview format.

Interface: All found the hypertext interface was easy to use, even though five had little or no computing experience. Two had problems using the mouse, but said that they would be able to use a touch screen without problems. One commented that older patients might be put off by computers.

Style: All said that the explanations were clear and understandable. "I thought they'd be understood by anyone. Quite clear." Four said there was some terminology that they didn't recognise at first, but all understood after clicking on the term in question. None found the explanations stilted: "Here it's [explained] the same way [as the doctor does]: you get the technical name, then it's broken down for you, and you can follow it." "It's easy enough language, and the medical terms are explained." One speculated that other patients might be put off by the medical language, although he wasn't himself.

Content: All said that the contents of the explanations were useful and appropriate ("I've learned more looking at that than by speaking to the doctor"), although one would have preferred more detailed, "less simplistic" information, and one would have liked to be able to wander further in 'hyperspace'. All thought the explanations were of about the right length: "You would want it short and compact, just enough for you to find out about any problems or any blood results or whatever, ready to go into the doctor." "It's ideal — simplified for the layman."

Personalisation: All said that providing measurements from the record in problem explanations was useful, and made the explanations more relevant. Four commented that a history of tests and measurements, perhaps given in

a chart, would be even more helpful. Five thought that providing qualitative information first, before the actual values of the measurements, would make them easier to understand — for example, three commented that a bit more comparison of measurements with others in their demographic group would be helpful.

General Usefulness: All said that PIGLET was not redundant, even given the huge amount of other patient education materials. Seven found PIGLET particularly useful for reminding them of their current state (recent measurements, etc.) right before an appointment. Three said that they didn't use other patient education materials: "I don't actually use the pamphlets and things ... I find the information I get from [the British Diabetes Association etc.] rather limited. [PIGLET] relates to me, rather than a generalisation."

All said that they wanted access to their record, and all thought that an unexplained paper version was inadequate. Two patients would have also liked access to their doctor's notes, and one information about past hospital admissions.

Other comments : Nine thought that the check-box facility would help them remember questions they'd wanted to ask the doctor, although only two actually used it during their session (probably because none had consultations scheduled immediately after the evaluation session). "If you don't know something, you can get the basics there [PIGLET], then go to the doctor and find out what he thinks." Four commented that they often forget questions because they're worried about wasting the doctor's time.

Five said they'd learned something new — "I knew I had the problems, but I didn't know the words" — while five said they were fairly familiar with their problems and treatments and clicked for explanations mainly out of curiosity. Six pointed out minor errors to be corrected in their electronic medical record — all but one of which were also in their original paper record. One commented that it was reassuring to know that the doctor had taken note of their problems: "Sometimes you explain to the doctors that you have a problem, and you're not sure that they take it in. That [PIGLET] shows you that they do know. You're quite happy when you know that they know what your problems are."

All said they would use the system again, although two commented that it might be most useful to new patients. All asked for, and received, printouts of their session for future reference.

7 Improvements and Future Work

The evaluation suggested several changes and additions which might improve PIGLET's performance.

Some of the suggestions from the patients and medical professionals who tried the system can be implemented fairly easily: giving a history of past tests and measurements, for example. It might also be worthwhile generating some

simple graphs, showing the patient's position within their demographic group for any given measurement.

We also hope to expand PIGLET so that it can explain issues related to different chronic problems. For example, a cancer PIGLET could generate explanations of records of cancer patients. Extending the system in this way should involve only adding new knowledge to the knowledge base, and adding a few specialised text plans for types of explanation specific to that subarea of medicine (e.g. explanations of chemotherapy programs for cancer).

PIGLET might also be modified for a slightly different purpose: generating documentation to be printed out, as in de Rosis and Grasso (1994). We have found that PIGLET's hypertext chunks translate nicely into a question-answer format printed page, which might be appropriate for providing additional printed documentation for prescriptions.

8 Conclusion

In building and evaluating this prototype system, we hoped to show how simple explanation generation techniques could be used to improve patient access to records and patient education in general. The resulting system, PIGLET, generates short hypertext explanations of medical issues, centred on a patient's medical record, and personalised to that patient's particular needs. The design of PIGLET reflects the needs of the end users, rather than those of artificial intelligence researchers. More complex and subtle methods of explanation generation certainly exist; in this case, however, relatively simple methods were adequate and appropriate.

Our evaluation of the system showed that both patients and medical staff found the system helpful and easy to use, demonstrating that our techniques can produce clear and comprehensible explanations. The personalisation of the explanations was found to make the explanations more relevant to the patients, which has been shown in other studies (Osman et al 1994) to improve patient care in general. In future we hope to expand PIGLET to cover a wider range of medical issues, build a simple interface for medical staff wishing to add information, and to do more rigorous testing of patient reactions.

Acknowledgements

Thanks are due to the University of Glasgow, for funding this project; to Drs. Fisher, Kesson and Semple, for trying out the system, and letting us try it on their patients; to the staff of the diabetes clinics at the Royal Alexandria, the Southern General, and the Royal Infirmary; and to all the patients who gave of their time and opinions.

References

D Berry, T Gillie, and S Banbury. What do patients want to know: an empirical approach to explanation generation and validation. Technical report, Department of Psychology, University of Reading, 1994.

B Buchanan, J Moore, D Forsythe, G Banks, and S Ohlsson. Involving patients in health care: explanation in the clinical setting. In M E Frisse, editor, *Proceedings of the sixteenth annual symposium on computer applications in medical care (SCAMC 92)*, pages 510–514. McGraw Hill, 1992.

B Buchanan, J Moore, D Forsythe, G Carenini, S Ohlsson, and G Banks. Using medical informatics for explanation in a clinical setting. Technical report, Department of Computer Science, University of Pittsburgh, 1994.

A Cawsey. *Explanation and Interaction: the computer generation of explanatory dialogues*. MIT Press, Cambridge, Massachusetts, 1992.

G Carenini and J Moore. Generating explanations in context. In *Proceedings of the International Workshop on Intelligent User Interfaces*, January 1993.

F de Rosis and F Grasso. Mediating between hearer's and speaker's views in the generation of adaptive explanations. Technical report, Dipartimento di Informatica, Universita di Bari, 1994.

R B Jones, A J Hedley, I Peacock, S P Allison, and R B Tattershall. A computer assisted register and information system for diabetes. *Methods of Information in Medicine*, 22:4–14, 1983.

R B Jones, S M McGhee, and D McGhee. Patient on-line access to medical records in general practice. *Health Bulletin*, 50(2):147–54, March 1992.

L Osman, M Abdalla, J Beattie, S Ross, I Russell, J Friend andJ Legge, and J Graham Douglas. Reducing hospital admissions through computer supported education for asthma patients. *BMJ*, 308:568–71, 1994.

E Reiter and C Mellish. Optimising the costs and benefits of natural language generation. In *Proceedings of IJCAI 1993*. Morgan Kaufmann Publishers, 1993.

D Rosner and M Stede. Generating multilingual documents from a knowledge base: The techdoc project. In *Proceedings of COLING-94*, 1994.

This article was processed using the LaTeX macro package with LLNCS style