



Technical Report N° 2006/06

Generating scripts for personalised medical dialogues
for patients

Paul Piwek
Richard Power
Sandra Williams

21st June 2006

Department of Computing
Faculty of Mathematics and Computing
The Open University
Walton Hall,
Milton Keynes
MK7 6AA
United Kingdom

<http://computing.open.ac.uk>

Generating scripts for personalised medical dialogues for patients

Paul Piwek, Richard Power and Sandra Williams

Computing Science Department

The Open University

Walton Hall, Milton Keynes, MK7 6AA

{p.piwek, r.power, s.h.williams}@open.ac.uk

Abstract

We propose an NLG system for communicating information to cancer patients about their medical records through scripted dialogues between animated agents. Building on past work, we discuss design issues and research questions, particularly the question of how to convey medical concepts in terms of everyday concepts. Our novel solution incorporates the information to be communicated into a personalised dialogue script between an expert medic and a novice medic. We discuss ongoing work to combine an existing NLG system that generates summaries of patient records with another existing NLG system that generates scripted dialogues for animated agents. We plan to evaluate the completed system with medical staff and patients.

1 Introduction

Information about patient histories is increasingly available in machine-usable form – i.e., in databases of Electronic Health Records (EHRs). The CLEF project (Hallett and Scott, 2005) is demonstrating ways in which reports generated from such data can assist doctors and medical researchers. We explore here a further application in which material from EHRs is selected and organised for presentation to patients.

What purpose would such a presentation serve? It would emphatically not replace the face-to-face consultation with a specialist. Rather, the aim would be to help prepare patients so that they use consultations to better effect.

For obvious economic reasons, a presentation based on a patient’s EHR would have to be produced automatically; this is therefore a natural application of NLG technology, and several researchers have already addressed issues in patient communication, see Cawsey, Webber and Jones’s (1998) summary. However, this work has been limited to relatively simple explications of scientific language (i.e., finding colloquial replacements for technical terms). Our aim is to go deeper than this, by showing reformulations that are not just linguistic but conceptual.

To achieve these objectives, we plan a novel approach in which the EHR material is presented to patients in the form of a dialogue between two medical workers, one expert and one novice. Some evidence for this approach comes from empirical studies. It has been found, for instance, that students learn vicariously by listening to other students interacting and asking questions in dialogues with teachers (Cox et al., 1999). Also that listening to a dialogue is more effective than listening to a monologue, providing significantly better recall of content (Craig et al., 2000).

On the basis of this evidence, we hypothesize that vicarious learning from dialogues might be effective in communicating complex medical information from a patient’s personal medical notes. We propose to construct an NLG system that generates dialogue scripts which patients can view through a multimedia presentation of a spoken conversation between two embodied virtual agents.

To develop an experimental prototype relatively quickly, we intend to combine some existing resources. Our starting point will be the CLEF system (Hallett and Scott, 2005), which summarises the EHRs of cancer patients for clinical staff. Summaries are generated from a number of viewpoints, depending on which types of events are in focus – diagnosis, investigative

tests, or interventions (e.g., operations). Figure 1 shows an example of such a summary with the focus on interventions; Figure 2 shows an example of a dialogue produced from the same underlying data.

In week 483, histopathology revealed primary cancer of the right breast. Radical mastectomy on the breast was performed to treat the cancer. Radiotherapy was initiated to treat primary cancer of the right breast. In the weeks 492 to 496, 5 radiotherapy cycles were performed.

Figure 1: CLEF EHR summary for clinicians (Hallett and Scott, 2005)

By automated generation of scripted dialogue (pioneered e.g. by André et al., 2000) we mean automation of both the writing of the script and its performance. For the latter task, there are several off-the-shelf player technologies ranging from the easily deployable Microsoft Agent (www.microsoft.com/msagent/) to commercial packages such as Charamel's CharActor plugin (www.charamel.com). Our research focuses on dialogue script authoring, rather than player technology.

Senior Nurse: Do you have the record of cancer treatment for <Ms X>?
Student Nurse: Here you are.
Senior Nurse: Read it to me.
Student Nurse: It says "right breast histopathology in week 483". That was, um, 13 weeks ago. What is histopathology?
Senior Nurse: They took a small piece of flesh from her right breast and the lab staff looked at it under a microscope. Did they find cancer?
Student Nurse: Yes they did and she had a mastectomy.
Senior Nurse: That means she had an operation to remove the breast and treat the cancer. Did she have radiotherapy?
Student Nurse: Yes, over the past 5 weeks she had 5 radiotherapy cycles. What is radiotherapy for?
Senior Nurse: Radiotherapy treats any remaining breast cancer.

Figure 2: Example dialogue script for a patient, week 496, same underlying data as before.

2 Design issues in creating dialogue scripts for patients

First we considered the scenario in which the dialogue in Fig. 2 could be presented to the patient. One of many potential scenarios follows. Before going for a consultation with her oncologist, a cancer patient watches an automatically generated video of animated agents playing out a dialogue scripted from her own recent medical notes. Not only does the dialogue in this setting

help her recall the current state of her treatment, but it encourages her to ask questions during the subsequent consultation by reminding her of the exact meaning of technical terms (e.g. histopathology) and giving her more confidence in expressing herself, both in technical language and informal language. It also provides explanations of why certain procedures were carried out (e.g., radiotherapy served to treat any remaining cancer), so guiding her towards asking follow-up questions ("Do I have any cancer remaining after my radiotherapy treatment?").

The characters are two nurses, one junior and one senior. Choosing a patient and a doctor, or two patients, would not work because if one of the characters is a patient, it seems inevitable that this character would have to represent the actual patient watching the presentation. This could easily give offence. The combination of junior and senior nurses sets up a scenario similar to the one used in successful vicarious learning experiments (Craig et al., 2000).

We envisage that the junior nurse is an eager, but not very knowledgeable, character and that the senior nurse is a kindly, experienced, trustworthy instructor whose words carry weight. The junior reads from a report, which the senior nurse then explains.

Stylistically, the Fig. 2 script concentrates on facts and education, rather than on entertainment. Given the sensitive nature of the content, this seemed wise. We have deliberately not introduced emotive language such as "unfortunately the test showed cancer". We assume, following Back et al. (2005), that it is important for a patient to watch professionals discussing her case in a calm, matter-of-fact and unemotional manner. Our hope is that the patient may then be able to imitate this manner in her consultation and hence not only understand the specialist better but also ask questions with greater confidence. Finally, again following Back et al. (2005), the language is straightforward and no more than three pieces of information are conveyed (histopathology, mastectomy and radiotherapy).

3 Generating the Explanatory Dialogue

Selection of EHR data is the same for the dialogue as for the monologue, but thereafter the planning process diverges. The propositional content of the dialogue has an extra level of complexity, because it is concerned not only with the clinical events themselves, but with different

ways of conceptualising them and expressing them linguistically.

Dialogue planning has two stages. First, we plan a discourse that can be regarded as a set of statements about the events in the data record and how they can be conceptualised and expressed – a kind of metadiscussion of the clinical record. Secondly, we recast these statements so that they take the form of a dialogue between two agents. The output of the first stage would be a rhetorical-semantic representation rather than a text, but for presentational purposes it is useful to imagine this output as a sequence of statements, some of which are shown in Table 1, left-hand column.

To obtain the dialogue, some of these statements have to be split into two parts, a question (either yes/no or WH) and an answer. Note that not all statements are split into two parts: 6 and 8 are left as single assertions.

Discourse Plan	Dialogue Turns
6: This means that she had an operation to remove the right breast.	6: This means that she had an operation to remove the right breast.
7: Ms X subsequently had a course of radiotherapy.	7: Did Ms X subsequently have a course of radiotherapy?
8: This comprised five cycles of radiotherapy over five weeks.	7: Yes. 8: This comprised five cycles of radiotherapy over five weeks.
9: Radiotherapy is performed in order to treat any residual cancer after a breast removal.	9: What is radiotherapy for? 9: It is performed in order to treat any residual cancer after a breast removal.

Table 1: Analysis showing partial discourse plan and corresponding dialogue turns

The final step is to assign the various questions, answers and statements to alternating turns by the agents. A single turn can contain several moves, thus avoiding conversational ping-pong (Davis, 1998): for instance, the penultimate turn has an answer, a statement (8), and a question.

4 Conclusions

Although our practical aim is to improve patient care, the methods proposed here obviously have wider applications. In many other areas of modern life (e.g. law, finance, education) people have routinely to interact with experts and to navigate between a technical world and the eve-

ryday terms in which we all normally think. As we have argued here, this problem is not purely linguistic: it requires something like an aligning of ontologies, so that a concise formulation in technical concepts is unpacked into a more elaborate formulation in everyday concepts. We hypothesize that a dialogue between expert and novice is a transparent and non-threatening method of reinforcing a person's (possibly fragmentary) technical understanding, and linking it more firmly to the ways in which they he/she usually thinks.

In evaluating such a system, several points have to be addressed. First, is the general idea valid, and would it fit into the specific context of communication with patients? Second, what design principles inform the construction of such dialogues – for example, what setting works best, and what style of interaction (e.g., factual vs. emotive)? Finally, can existing NLG technology support the generation of dialogues that meet these design requirements?

References

- André, E., Rist, T., Mulken, S. van, Klesen, M., and Baldes, S. (2000). In Cassell, J., Sullivan, J., Prevost, S. and Churchill, E. (eds.). *The Automated Design of Believable Dialogues for Animated Presentation Teams*. MIT Press, 2000, 220-255.
- Back, Anthony L., R. M. Arnold, W. F. Baile, J. A. Tulskey and K. Fryer-Edwards, (2005). *Approaching Difficult Communication Tasks in Oncology*. CA: A Cancer Journal for Clinicians, 55:164-177.
- Cawsey, A., Webber, B.L. and Jones, R. *Natural Language Generation in Healthcare* (1998). Journal of the American Medical Informatics Association.
- Cox, R., McKendree, J., Tobin, R., Lee, J. and Mayes, T. (1999) Vicarious learning from dialogue and discourse: a controlled comparison. *Instructional Science* 27: 431-458.
- Craig, S D., B Gholson, M Ventura, A C. Graesser (2000) Overhearing Dialogues and Monologues in Virtual Tutoring Sessions: Effects on Questioning and Vicarious Learning. *International Journal of Artificial Intelligence in Education*, 11, 242-25.
- Davis, R. (1998). *Writing Dialogue for Scripts*, A & C Black Ltd., London, ISBN 0-7136-4802-3.
- Hallett, Catalina and Donia Scott (2005). Structural variation in generated health reports. *Proceedings of the 3rd International Workshop on Paraphrasing*.