Types of Knowledge Required to Personalise Smoking Cessation Letters*

Ehud Reiter[1], Roma Robertson[1], and Liesl Osman[2]

[1] Dept. of Computing Science, University of Aberdeen, Aberdeen, Scotland,
{ereiter,rroberts}@csd.abdn.ac.uk
[2] Dept. of Medicine and Therapeutics, University of Aberdeen, Aberdeen, Scotland
l.osman@abdn.ac.uk

Abstract. The STOP system generates personalised smoking-cessation letters, using as input responses to a smoking questionnaire. Generating personalised patient-information material is an area of growing interest to the medical community, since for many people changing health-related behaviour is the most effective possible medical intervention. While previous AI systems that generated personalised patient-information material were primarily based on medical knowledge, STOP is largely based on knowledge of psychology, empathy, and readability. We believe such knowledge is essential in systems whose goal is to change people's behaviour or mental state; but there are many open questions about how this knowledge should be acquired, represented, and reasoned with.

1 Introduction

The STOP (Smoking Termination through cOmputerised Personalisation) system generates short smoking-cessation leaflets that are personalised for different recipients. Personalisation is based on responses to a questionnaire on smoking habits and beliefs, previous attempts to quit, current medical problems, and so forth.

The goal of STOP is to change a patient's behaviour in a medically desirable way. This is different from the decision-support systems that are what many people most associate with the AI/Medicine field, especially in terms of the knowledge needed. Decision-support systems are based on medical knowledge, that is knowledge about diseases, treatments, how the body works, and so forth. But while STOP uses some medical knowledge, it is primarily based on other types of knowledge, including the psychology behind behaviour change, techniques for empathy, and rules for effective writing.

* Many thanks to the experts who worked with us, including Scott Lennox, James Friend, Martin Pucci, Margaret Taylor, and Chris Bushe; and also to Scott Lennox and Jim Hunter for their comments on earlier drafts of this paper. This research was supported by the Scottish Office Department of Health under grant K/OPR/2/2/D318, and the Engineering and Physical Sciences Research Council under grant GR/L48812.

AI systems that attempt to change people's medical behaviour are relatively new, but they are attracting increasing attention. In part this is because the most important influence on many people's health is their behaviour (smoking, diet, compliance with treatment regimes, etc). To date, clinical trials of systems which produce personalised patient information material have been mixed, with some systems proven effective in changing behaviour of at least some patients but others showing no statistically significant effects. The challenge for the research community is to develop technology that increases the effectiveness of such systems, and also to determine the types of applications in which success is most likely.

2 Previous Work

A number of previous AI/Medicine projects have investigated generating personalised patient information that is intended to change the patient's behaviour or psychological state, of which the best known are perhaps MIGRAINE (Buchanan et al., 1995) and PIGLIT (Binstead, Cawsey, and Jones, 1995). Personalisation in these systems was primarily based on medical knowledge, in part because these systems only had access to medical information about patients; they did not have the information on attitudes and intentions available to STOP via its questionnaires. A clinical evaluation of one version of PIGLIT looked at PIGLIT's effect on patient satisfaction and patient anxiety; it showed a statistically significant effect on satisfaction, but not on anxiety (Cawsey et al., 1999).

Several systems which generate personalised patient information have also been produced by the public-health community, including at least two systems which produce smoking-cessation letters (Velicer et al., 1993; Strecher et al., 1994). These two systems base personalisation on a psychological theory of behaviour change, the Stages of Change model (Prochaska and diClemente, 1992). Clinical evaluations showed that both of these systems had a statistically significant impact on smoking cessation rates.

We hope that STOP will be more effective than previous AI/Medicine projects because it is based on knowledge about psychology (including the Stages of Change model), empathy, and readability as well as medical knowledge. We also hope that it will be more effective than previous systems developed by the public health community because it is uses AI and Natural-Language Processing (NLP) techniques.[1] A randomised controlled clinical trial of STOP is currently underway to test the effectiveness of the system; results of the trial should be available in late 1999.

[1] The focus of this paper is on comparing STOP to previous work in the AI community, not the public health community. But very briefly, we believe that STOP's performance is enhanced because it is based on knowledge acquired using structured AI knowledge acquisition techniques (these, for example, led to the smoker categories described in Section 4.1); and because NLG technology allows STOP to do a better job of optimising and satisfying constraints (for example, we believe it would be much harder to optimise content subject to a size constraint, as described in Section 3.1, without NLG representations and algorithms).

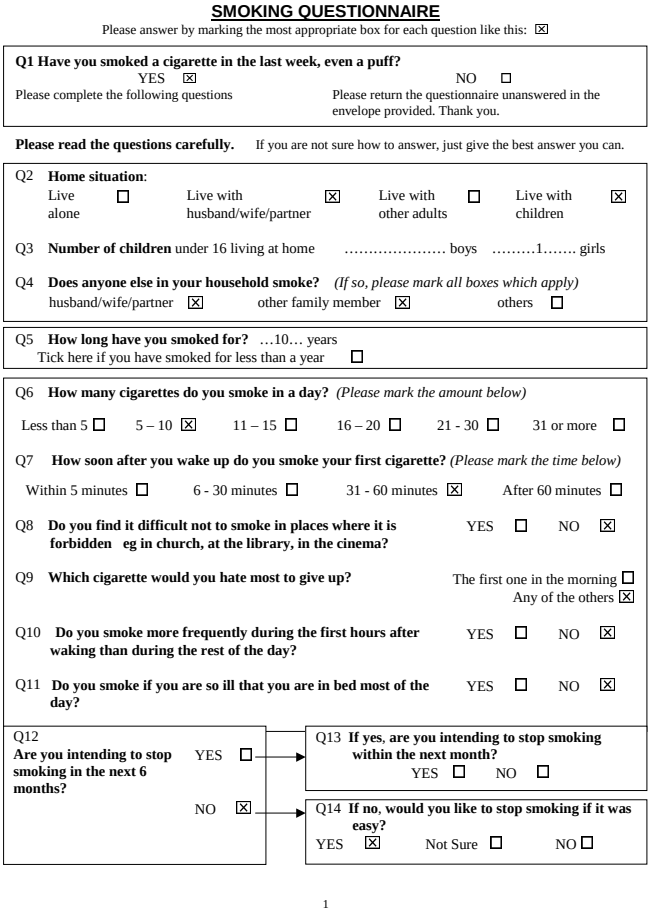

SMOKING QUESTIONNAIRE

Please answer by marking the most appropriate box for each question like this: ☒

Q1 Have you smoked a cigarette in the last week, even a puff?

YES ☒ NO ☐

Please complete the following questions

Please return the questionnaire unanswered in the envelope provided. Thank you.

Please read the questions carefully. If you are not sure how to answer, just give the best answer you can.

Q2 **Home situation**:
Live alone ☐ Live with husband/wife/partner ☒ Live with other adults ☐ Live with children ☒

Q3 **Number of children** under 16 living at home …… boys ……1……. girls

Q4 **Does anyone else in your household smoke?** *(If so, please mark all boxes which apply)*
husband/wife/partner ☒ other family member ☒ others ☐

Q5 **How long have you smoked for?** …10… years
Tick here if you have smoked for less than a year ☐

Q6 **How many cigarettes do you smoke in a day?** *(Please mark the amount below)*
Less than 5 ☐ 5 – 10 ☒ 11 – 15 ☐ 16 – 20 ☐ 21 - 30 ☐ 31 or more ☐

Q7 **How soon after you wake up do you smoke your first cigarette?** *(Please mark the time below)*
Within 5 minutes ☐ 6 - 30 minutes ☐ 31 - 60 minutes ☒ After 60 minutes ☐

Q8 **Do you find it difficult not to smoke in places where it is forbidden eg in church, at the library, in the cinema?** YES ☐ NO ☒

Q9 **Which cigarette would you hate most to give up?** The first one in the morning ☐ Any of the others ☒

Q10 **Do you smoke more frequently during the first hours after waking than during the rest of the day?** YES ☐ NO ☒

Q11 **Do you smoke if you are so ill that you are in bed most of the day?** YES ☐ NO ☒

Q12 **Are you intending to stop smoking in the next 6 months?** YES ☐ NO ☒

Q13 **If yes, are you intending to stop smoking within the next month?** YES ☐ NO ☐

Q14 **If no, would you like to stop smoking if it was easy?** YES ☒ Not Sure ☐ NO ☐

1

Fig. 1. First page of Fiona Cameron's questionnaire

3 The System

The input to STOP is a 4-page questionnaire on smoking; the first page of a questionnaire from a typical smoker, Fiona Cameron (not her real name), is shown in Figure 1. STOP also gets some basic information, such as age and sex, from the patient's medical record. The output of STOP is a small leaflet (4 pages of A5). The front page of a leaflet contains an introductory paragraph but otherwise is not personalised; the back page is selected from one of a dozen possible back pages, but is not personalised in detail. Most of the personalisation happens in the two inside pages of the leaflet; the inside pages of the leaflet generated for Fiona Cameron are shown in Figure 2.

STOP is thus a 'paper-in, paper-out' system; users fill out a paper questionnaire, and receive in response a paper leaflet. We could make STOP interactive

Smoking Information for Fiona Cameron

You have good reasons to stop...

People stop smoking when they really want to stop. It is encouraging that you have many good reasons for stopping. The scales show the good and bad things about smoking for you. They are tipped in your favour.

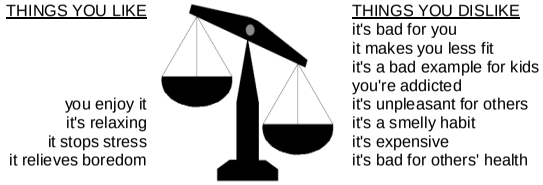

You could do it...

Most people who really want to stop eventually succeed. In fact, 10 million people in Britain have stopped smoking - and stayed stopped - in the last 15 years. Many of them found it much easier than they expected.

You are right to think that if you tried to stop smoking you would have a good chance of succeeding. You have several things in your favour.

- You have stopped before for over a year.
- You are a light smoker.
- You have good reasons for stopping smoking.
- You expect support from your partner, your family, your friends, and your workmates.

We know that all of these make it more likely that you will be able to stop. Most people who stop smoking for good have more than one attempt.

Overcoming your barriers to stopping...

You said in your questionnaire that you might find it difficult to stop because smoking helps you cope with *stress*. Many people think that cigarettes help them cope with stress. However, taking a cigarette only makes you feel better for a short while. Most ex-smokers feel calmer and more in control than they did when they were smoking. There are some ideas about coping with stress on the back page of this leaflet.

You also said that you might find it difficult to stop because you would *get bored*. It's a habit to smoke when you have nothing to do. If you decide to stop it might be worth planning how you could keep yourself busy.

And finally...

We hope this letter will help you feel more confident about giving up cigarettes. You know why you'd like to stop, you just need to decide to do it! If you have a go, you have a real chance of succeeding.

With best wishes,

Great Western Road Medical Group.

Fig. 2. Inside pages of leaflet generated for Fiona Cameron

(like MIGRAINE and PIGLIT), and have users fill out questionnaires and read leaflets on-line; but this would mean the system could only be used by people with access to computers. Since we want to reach as many smokers as possible, including people who do not have access to and are not comfortable with computers (such as many middle-aged and elderly people living in low-income public housing estates), we decided to use the 'paper-in, paper-out' model.

We do not expect STOP to have much effect on most smokers; smoking is a difficult habit to give up, and receiving a letter in the post is unlikely to make much difference to most people. But we hope that it will help a few people to quit. Studies show that brief discussions about smoking with a doctor will cause about 2% of people to quit smoking (Law and Tang, 1995); if we can achieve a similar effectiveness rate then STOP will be useful from a public-health perspective, since it is a very cheap intervention. Human doctors, incidentally, generally do not routinely have such discussions with patients because they find a 98% failure rate too discouraging; a computer system, of course, does not get discouraged no matter how small its success rate.

3.1 System Design

Space prohibits a detailed description of how STOP works, but a brief description follows. Questionnaires are read by an optical scanner, and processed by the core STOP system, which produces an RTF file which is printed with Microsoft Word. The core system is a Natural Language Generation (NLG) system which follows the model described in Reiter and Dale (1999). Processing is divided into the three stages of document planning, microplanning, and realisation, of which document planning (deciding what information to communicate) is the most complex. Oversimplifying to some degree, the document planner works by first classifying smokers into one of 7 categories, and then activating a schema (McKeown, 1985; Reiter and Dale, 1999) associated with that category. The combination of classification and schemas is similar at least in concept to the Exemplars system (White and Caldwell, 1998), although the implementation is quite different. The schemas produce a tree, known as a document plan. Each leaf node of the tree essentially defines one sentence in the leaflet. The internal nodes of the tree indicate how sentences are grouped, associate document structures (such as paragraphs or itemised lists) with groups of sentences, and sometimes specify discourse relations (Reiter and Dale, 1999) between daughter nodes.

Perhaps the most innovative aspect of STOP from an NLG perspective is its use of revision to optimise the content of a letter, given the size constraint (4 pages of A5). In general terms, this is done by having schemas annotate document plan constituents (both sentences and internal nodes) with importance markers. If the leaflet is too long, a revision module deletes the least important constituents until the size limit is satisfied. The microplanner and realiser convert this structure into an actual RTF document specification; the microplanner decides which discourse relations should be expressed in the text based on the outcome of the revision process, to ensure that the resulting document is rhetorically coherent.

4 Types of Knowledge

Most of the STOP project to date has focused on knowledge acquisition (KA). This was done in a structured fashion using standard techniques developed by the knowledge-acquisition community (Scott, Clayton, and Gibson, 1991), such as sorting and think-aloud protocols. These KA activities revealed that experts used several types of knowledge to produce smoking-cessation letters, including:

- psychological knowledge about how people change addictive behaviours;
- practitioner knowledge about 'empathy';
- linguistic knowledge about readability in texts; and
- medical knowledge about smoking.

Of course, experts also used knowledge about the patients, which they took from the questionnaire. Some of the expert's knowledge may be specific to Aberdeen or Scotland; this is not something we have investigated to date.

We will not further discuss medical knowledge here, as it is very common in AI and Medicine systems. It also turned out to be less important than we originally thought it would be. Most smokers are not interested in the medical details of smoking, and are well-aware of the health risks of smoking; indeed they may overestimate rather than underestimate health risks. Some of the other types of knowledge listed above are perhaps more unusual in AI/Medicine systems, and we discuss these below.

4.1 Psychological Knowledge about Addictive Behaviours

A crucial type of knowledge in STOP is psychological knowledge about how people stop addictive behaviours. All other systems we are aware of which produce personalised smoking-cessation material use the Stages of Change model (Prochaska and diClemente, 1992). This model groups smokers into five stages:

Precontemplator: not intending to quit
Contemplator: seriously considering quitting
Preparation: intending to quit
Action: in the process of quitting
Maintenance: has quit, avoiding relapse

Only the first three stages are relevant to systems (such as STOP) which target people who are currently smoking.

The Stages of Change model also specifies what type of information should be communicated to people in each stage. For example, information for precontemplators should emphasise that the disadvantages of smoking outweigh its advantages; information for contemplators should discuss specific 'barriers' to change, such as addiction or fear of weight gain; and information for preparers should present techniques for quitting.

We initially hoped to use the Stages of Change model 'off the shelf', because it is clinically validated and widely used. But we found that it often suggested

content that we believed to be inappropriate, perhaps because it is a general model which is not tuned either to smoking cessation or to the task of generating individualised letters. For example, many of the people in our study (including Fiona Cameron) are precontemplators in the sense that they are not intending to quit smoking; but they also are already convinced that smoking is bad for them, so it seems redundant to stress the disadvantages of smoking in a leaflet. The reason such people are not intending to quit is that they do not think they will be able to stop smoking; the right emphasis for such people is therefore not 'smoking is bad for you' but rather 'you can quit if you really want to.'

We ended up using 7 categories for smokers instead of 3; these categories were derived from sorting KA exercises. The categories are:

Committed smoker: People who clearly want to smoke; they get short letters reminding them of the health risks of smoking, and suggesting sources of advice in case they change their mind at some future date.

Classic precontemplator: People with mixed feelings about smoking, but who are not currently intending to quit; they get letters which emphasise that the disadvantages of smoking outweigh its advantages.

Lacks confidence: People who would like to stop smoking, but don't think they will be able to quit; they get letters which emphasise confidence-boosting

Classic contemplator: People who are considering quitting but have barriers; they get letters which emphasise overcoming their barriers.

Borderline contemplator: People who have mixed feelings about smoking but are considering quitting; they get letters which emphasise the disadvantages of smoking as well as how to overcome barriers.

Classic preparer: People who are intending to quit; they get letters which emphasise techniques for quitting.

Uncertain preparer: People who are intending to quit but have ambivalent feelings about smoking (this includes people who are being pressured by someone else to quit smoking); they get letters which both stress the disadvantages of smoking and suggest techniques for quitting.

The first three of these categories are essentially refinements of the Precontemplator stage; the subsequent two categories are refinements of the Contemplator stage; and the last two categories are refinements of the Preparation stage.

While category determines the emphasis of the letter, other information may be included as well, depending on the patient's details. For example, while the emphasis of the letter shown in Figure 2 is on confidence-building, since Fiona Cameron is classified in the *Lacks confidence* class, it also includes a section reinforcing her reasons for quitting, and some advice on techniques for coping with problems (this is partially intended to boost her confidence that she can quit).

4.2 Empathy

As we worked through various KA exercises with our experts, it became clear that they were using a type of knowledge which we had not initially anticipated,

which we now call 'empathy'. The purpose of empathy knowledge is to produce leaflets which people take seriously and think about, instead of tossing aside as yet another anti-smoking polemical.

The only previous work on empathy in AI/Medicine systems which we are aware of is Forsythe's work (1995) during the MIGRAINE project; she used the term 'enlistment', which may be a better name than 'empathy'. In particular, Forsythe identified the need to treat patients with respect and acknowledge their competence; this finding influenced the wording of explanations in MIGRAINE. There is also a substantial body of work on empathy in the general medical literature, but it focuses on empathy in oral face-to-face consultations, and we found it difficult to apply these ideas to written leaflets.

We acquired a set of empathy rules via KA exercises; we also used some general psychological and communication principles (Monahan, 1995). Unlike Forsythe, who observed doctors conducting oral consultations, we asked our experts to focus on empathy in written leaflets. Some of the rules which emerged from these exercises are:

- Be positive, do not criticise. For example, avoid negative constructs such as *We regret that*; use neutral constructs such as *We see that* instead. Positive constructs such as *We are very pleased that* are desirable and should be used when appropriate.
- Wherever possible, make points by repeating facts that the smoker has stated in the questionnaire. For instance, the scales graphic in leaflet shown in Figure 2 uses Cameron's questionnaire responses to emphasise to her that there are many things she dislikes about smoking.
- Use second-person (*you*) sentences wherever possible; when this isn't possible, try to use first-person plural (*we*) sentences.

One expert also felt it would be useful to relate letters to a person's expertise (for example, *As a nurse, you know that ...*) and circumstances (for example, *I know being a single mother is very hard*); and two experts suggested using humour. However, another expert felt that these techniques could backfire and antagonise people if we were not very careful. Because we were working to a tight schedule, we elected not to implement any techniques which any of our experts expressed doubts about.

We believe that achieving a better understanding of empathy rules, including some testing and validation, is essential to the success of future personalised patient-information systems. We plan to investigate this in future research.

4.3 Readability

A primary design imperative in STOP was that texts should be readable by a wide range of recipients, including people with poor reading skills. Readability rules were also acquired with KA sessions, but these used an expert on health-information leaflets instead of a doctor; we also ran some KA sessions with a graphic designer to acquire acquire visual appearance rules (layout, font, etc).

In technical NLG terms, many readability rules were essentially choice rules for microplanning operations such as lexicalisation and aggregation.

Many of the rules that emerged from the KA sessions were similar to those used for AECMA Simplified English (AECMA, 1986); indeed we gave our expert a copy of the Simplified English manual, and she found that it largely agreed with her thinking. Examples of our rules are:

- Sentences should be as short as possible. Hence STOP never aggregates messages (that is, forms a complex sentence by merging two simpler sentences with a relative clause or a conjunct such as *and*).
- Common (high-frequency) words and simple syntactic structures should be used wherever possible. For example, *You may be addicted* instead of *It is likely, but not certain, that you are addicted.*
- Avoid impersonal constructs. For example, *You mentioned some good reasons for stopping* instead of *There are lots of good reasons for stopping.*
- Put bullet lists in the same font as normal text (graphic design rule).

We originally thought about varying sentence length and word choice for different patients; for example, using more complex structures and words for university-educated patients. But our experts believed it was best to always use simple language in patient-information material, and indeed this is supported by research elsewhere (Davis et al., 1996).

We would have liked to implement some of these rules declaratively within the NLG system, but this proved difficult. Some rules (such as avoiding aggregation) were trivial to implement, but others turned out to require complex domain reasoning and knowledge as well as linguistic knowledge. For example, consider the principle that high-frequency words should be used whenever possible. It is easy to get a list of word frequencies and pick the highest-frequency word when given a choice between synonyms. But exact synonyms are not common; *may* is not exactly the same as *It is likely, but not certain*, for example. In this case, as in many others, the choice is between a relatively complex linguistic construct that conveys a precise meaning, and a simple linguistic construct that conveys an approximation to this meaning. Thus, whether *may* is acceptable depends on how important it is to communicate an exact meaning, and whether approximation is acceptable; and this is a fairly deep content decision which is difficult to encode declaratively.

5 Conclusion

AI systems which change people's behaviour in medically desirable ways have tremendous potential to improve health, since behaviour is the largest influence on health for many people. Given their psychological goals, we believe it is essential that such systems be based on knowledge of psychology, empathy, and readability, as well as more 'conventional' medical knowledge. We have made an initial attempt at understanding some of the issues in acquiring, representing, and reasoning with such knowledge in AI systems; but it is very much an initial attempt, and much more work needs to be done in this area.

References

AECMA. 1986. A guide for the preparation of aircraft maintenance documentation in the international aerospace maintenance language. Available from BDC Publishing Services, Slack Lane, Derby, UK.

Binstead, Kim, Alison Cawsey, and Ray Jones. 1995. Generated personalised patient information using the medical record. In Pedro Barahona, Mario Stefanelli, and Jeremy Wyatt, editors, *Proceedings of the Fifth Conference on Artificial Intelligence and Medicine Europe (AIME-1995)*, pages 29–41. Springer.

Buchanan, Bruce, Johanna Moore, Diana Forsythe, Guiseppe Carenini, Stellan Ohlsson, and Gordon Banks. 1995. An interactive system for delivering individualized information to patients. *Artificial Intelligence in Medicine*, 7:117–154.

Cawsey, Alison, Ray Jones, Janne Pearson, and Kim Binstead. 1999. The design and evaluation of a personalised health information system for patients with cancer. *User Modelling and User-Adapted Interaction*. Submitted.

Davis, T, J Bocchini, D Fredrickson, et al. 1996. Parent comprehension of polio vaccine information pamphlets. *Pediatrics*, 97:804–810.

Forsythe, Diana. 1995. Using ethnography in the design of an explanation system. *Expert Systems with Applications*, 8(4):403–417.

Law, Malcolm and Jin Tang. 1995. An analysis of the effectiveness of interventions intended to help people stop smoking. *Archives of Internal Medicine*, 155:1933–1941.

McKeown, Kathleen. 1985. *Text Generation*. Cambridge University Press.

Monahan, Jennifer. 1995. Using positive affect when designing health messages. In Edward Maibach and Roxanne Parrott, editors, *Designing Health Messages: Approaches from Communication Theory and Public Health Practice*. Sage, pages 81–98.

Prochaska, James and Carlo diClemente. 1992. *Stages of Change in the Modification of Problem Behaviors*. Sage.

Reiter, Ehud and Robert Dale. 1999. *Building Natural Language Generation Systems*. Cambridge University Press. In press.

Scott, A. Carlisle, Jan Clayton, and Elizabeth Gibson. 1991. *A Practical Guide to Knowledge Acquisition*. Addison-Wesley.

Strecher, Victor, Matthew Kreuter, Dirk-Jan Den Boer, Sarah Kobrin, Harm Hospers, and Celette Skinner. 1994. The effects of computer-tailored smoking cessation messages in family practice settings. *The Journal of Family Practice*, 39:262–271.

Velicer, Wayne, James Prochaska, Jeffrey Bellis, Carlo diClemente, Joseph Rossi, Joseph Fava, and James Steiger. 1993. An expert system intervention for smoking cessation. *Addictive Behaviors*, 18:269–290.

White, Michael and Ted Caldwell. 1998. EXEMPLARS: A practical extensible framework for dynamic text generation. In *Proceedings of the Ninth International Workshop on Natural Language Generation (INLG-1998)*, pages 266–275.