The Architecture of the STOP System*

Ehud Reiter and Roma Robertson
Dept. of Computing Science, University of Aberdeen, Aberdeen, Scotland,
{ereiter,rroberts}@csd.abdn.ac.uk

Abstract

This paper describes the architecture of the STOP system in the framework laid out by the RAGS project.

1 Introduction

The STOP (Smoking Termination through cOmputerised Personalisation) system generates short smoking-cessation leaflets that are personalised for different recipients (Reiter, Robertson, and Osman, 1999). Personalisation is based on responses to a questionnaire on smoking habits and beliefs, previous attempts to quit, current medical problems, and so forth. An extract of a STOP questionnaire is shown in Figure 1. A STOP leaflet consists of 4 A5 pages: a front page which is mostly logos and a title; two inside pages which communicate most of the information; and a non-generated back page which is selected from one of 16 possible back pages but is not otherwise personalised. The front and back pages for the leaflet produced from the questionnaire extracted in Figure 1 is shown in Figure 2 (the front page as at the top of the figure), and the inside pages of this leaflet are shown in Figure 3.

The STOP system is currently being evaluated in a clinical trial, which compares cessation rates in a group of smokers who received STOP leaflets to cessation rates in two control groups. The trial requires the system to generate 850 letters for previously unseen smokers; hence it is essential that STOP be robust (Reiter, 1999).

STOP is implemented in Java. It can be operated in two modes: a Web-based version where input is via a CGI script and output is in HTML, and a paper-based version (used in the clinical trial) where input comes from Microsoft Access databases which hold scanned questionnaires and general patient data, and output is an RTF file which is printed using Microsoft Word.

This paper presents the architecture of the STOP system in terms of the NLG system architecture proposed by the RAGS project (Cahill et al., 1999),

*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.

focusing on data structures. For other information about the STOP project (including the Web-accessible version of the system), see the project Web page at `http://www.csd.abdn.ac.uk/~rroberts/smoking.html`

2 Conceptual representation

The RAGS architecture assumes that conceptual information is accessed using knowledge-base like queries that obtain role values, check subsumption, and so forth. This is basically the case in STOP. A Smoker object encapsulates all information about a smoker. The 'core' NLG system gets information about smokers via the Smoker class; the methods and members of this class return the information the system needs about a particular smoker. In short, from the RAGS perspective STOP accesses conceptual information using role-value queries on the Smoker class.

The Smoker class is not trivial, and can be considered to be a type of 'KB accessor' (Lester and Porter, 1997). Its responsibilities include

- Providing a uniform interface to smoker data, regardless of whether the data comes from a CGI script (for the Web version) or MS Access databases (for the paper version).
- 'Regularising' data – for example, converting patient names into a standard format. The class also does a fair amount of consistency checking on the raw data, since the questionnaires are often inconsistent (due to human and scanning errors).
- Doing some calculations on the data. For example, addiction to nicotine is estimated using the Fagerström score (Heatherton et al., 1991).

It is fairly common for applied NLG systems to use such an interface class or module to present a clean view of the underlying data to the rest of the NLG system. Perhaps it is worth including this in the RAGS architecture?

3 Rhetorical representations

The RAGS architecture assumes that rhetorical structure is represented by a tree whose leaf nodes are semantic structures and whose internal nodes specify rhetorical relations. RAGS also assumes that generic rhetorical relations can be instantiated into concrete rhetorical relations.

STOP conforms to this architecture in so far as it represents texts as a tree structure whose leafs represent sentences or phrases, and whose internal nodes can indicate discourse relations. However, discourse relations in STOP are 'represented' as actual cue phrases; for example, *However* instead of Contrast. The microplanner can decide whether or not a cue phrase actually appears in the text; but it cannot replace it with a different cue phrase.

One of the main reasons for using cue phrases instead of abstract rhetorical (discourse) relations is that we believe that the linguistics of abstract rhetorical relations is poorly understood. There is little agreement in the community about what the 'right' set of abstract rhetorical relations is. Even worse, we do not have robust algorithms for choosing appropriate cue phrases to communicate rhetorical relations, and indeed research suggests that there are many semantic subtleties to cue phrases which are not well understood (Knott, 1996). For these reasons, we believe that using cue phrases instead of abstract relations makes STOP more robust.

4 Document Specification

According to the RAGS architecture, document structure is represented by a tree whose leafs are the sentences or phrases, and whose internal nodes are annotated with the attributes TextLevel, Layout, and Order.

In STOP, document structure is represented by a tree whose leaf nodes represent sentences or phrases; indeed this is the same structure used to represent rhetorical structure. An internal node in the tree may convey either rhetorical or document structure annotations, but not both (this restriction simplifies implementation).

A difference between STOP and RAGS is that STOP represents document structures as objects, instead of as a set of features. For example, a Section structure has a title member as well as a list of children; this contrasts to the RAGS approach of (as far as I can tell) making the title the first child, and annotating it with a title layout feature. STOP document structures can also have features; for example horizontalAlignment can be left-aligned, right-aligned or centred.

For simple structures such as sections and paragraphs, the difference between these two approaches is perhaps mainly notational. However, for more complex structures we believe that the object approach is better. For example, the top-level DSDocument structure in a STOP leaflet has 6 members, in addition to its children

- **title** appears both on the front page and across the inside pages.
- **practiceName** and **practiceLogo** give information about the smoker's GP practice, and appear on the front page in special positions.
- **salutation** and **frontPageParagraph** are the texts which appear on the front page, below the logos and titles.
- **backpage** gives a filename for a non-generated back page which is inserted into the leaflet.

We believe such complex structures are easier to handle if they are explicitly represented as objects, with appropriate special members and constructors, and validity checking routines. Object-oriented document structures may also be more in keeping with modern thoughts on DTD's (Document Type Definition).

Document structures can also have structure-dependent features. For example, the DSTextBox structure, which produces the text boxes to the left and right of the scales graphic in Figure 3, supports a verticalAlignment feature as well as a horizontalAlignment feature. This is a useful feature for text boxes; for example the lists shown in Figure 3 are aligned to end on the bottom of the box instead of (as is more standard) start at the top of the box. However, this feature is much less useful for other document structures, and probably not worth supporting in general, as it requires some implementation effort. Again, an object representation allows us to restrict features to specific structures.

A further advantage of an object-oriented approach to document structures is that it allows inheritance. For example, both DSTextBox and DSGraphic (used for the scales graphic) inherit basic properties such as height and width from a common DSShape ancestor.

5 Semantic Specifications

STOP does not use semantic representations. This design decision was largely due to our dissatisfaction with existing realisers. We did not wish to implement a new realiser, and all existing realisers which we investigated suffered from either inadequate grammatical coverage or insufficient documentation (Reiter, 1999).

6 Syntactic Specifications

Again because of dissatisfaction with existing realisers, STOP does not use complex syntactic representations. Instead, it represents sentences and phrases by what Reiter and Dale (2000) call 'canned text', that is strings (or lists of strings) that need to be orthographically processed, but do not need to be syntactically or morphologically processed. Orthographic processing includes capitalisation, insertion of certain types of punctuation, and insertion of inter-token white space. Canned text strings can also be annotated by typographic markers, such as bold or italic face.

It is difficult to relate this representation to any representation described in the RAGS report, because (as far as I could tell) no mention is made in this report of either orthographic processing or of annotating portions of a sentence to appear in different fonts/faces.

7 Control Flow

The RAGS documents rightly focus on intermediate data structures, but perhaps it is also worth saying something about control flow. The original version of STOP (which was used in the clinical trial) had a single-solution pipeline architecture. By 'single-solution', I mean that each pipeline module passed a single solution to the subsequent module. However, this architecture was poorly suited to meeting

length constraints. It is imperative in STOP that all leaflets fit on 4 A5 pages, and desirable for as much information as possible to be communicated given this length constraint. Unfortunately, it was difficult to meet this constraint and optimisation in a single-solution pipeline architecture. A newer version of the system (Reiter, 2000) allows two other modes of operation:

- **Multiple solutions:** In this mode, pipeline modules can pass multiple solutions to other modules; a choice module at the end of the pipeline picks the document with the highest word count which satisfies the size constraint. This has dramatic effects; even just allowing two solutions to be passed along the pipeline increases the average number of words in a leaflet by 15% while ensuring conformance to the size constraint.

- **Revision:** In this mode, the modules operate in a normal single-solution mode, but a revision module at the end of the pipeline can either accept a leaflet or request that it be generated with a bit more or a bit less content. The revision process can iterate several times. Revision mode gives a slightly higher word count than the multiple-solution pipeline; perhaps more importantly, it also makes the process (of optimising content subject to a size constraint) more robust in the face of changes to the code.

I hope the RAGS architecture is general enough to support such modes of operation.

8 Intermodule Communication

The RAGS architecture proposes that modules communicate via a shared whiteboard, which holds information about objects and relationships between objects. STOP uses a simpler approach; modules directly pass intermediate data structures (represented as Java objects) to other modules.

I must say that it is not clear to me what the advantages are of using such a whiteboard, at least for a system implemented in Java (perhaps it makes more sense for a Prolog system, but I hope the RAGS architecture is not tied to a specific programming language). Certainly the object-arrow representation suggested for whiteboard structures seems awkward compared to the Java object mechanism, which is computationally efficient, easy to manipulate from Java, supports inheritance, allows procedures as well as data to be associated with objects, and allows implementation details to be encapsulated (which is useful for maintenance).

The issue of whether intermediate data structures should be stored in a central database is separate from the issue of representation; STOP's intermediate data structures could be stored in a Java object database, and made accessible to all modules in the system. This may not be desirable from a maintainability perspective, however; it is easier to change a data structure which is only known to a small number of modules than it is to change a data structure which is visible and potentially used by all modules.

References

Cahill, Lynn, Christy Doran, Roger Evans, Chris Mellish, Daniel Paiva, Mike Reape, Donia Scott, and Neil Tipper. 1999. Towards a reference architecture for natural language generation systems. Technical Report ITRI-99-14, Information Technology Research Institute, University of Brighton, UK.

Heatherton, Todd, Lynn Kozlowski, Richard Frecker, and Karl-Olav Fagerström. 1991. The Fagerström test for nicotine dependence: A revision of the Fagerström tolerance questionnaire. *British Journal of Addiction*, 86:1119–1127.

Knott, Alistair. 1996. *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, University of Edinburgh.

Lester, James and Bruce Porter. 1997. Developing and empirically evaluating robust explanation generators: The KNIGHT experiments. *Computational Linguistics*, 23(1):65–101.

Reiter, Ehud. 1999. Shallow vs deep techniques for handling linguistic constraints and optimisations. In *Proceedings of the KI-1999 Workshop on May I Speak Freely: Between Templates and Free Choice in Natural Language Generation*.

Reiter, Ehud. 2000. Pipelines and size constraints. *Computational Linguistics*. Forthcoming.

Reiter, Ehud and Robert Dale. 2000. *Building Natural Language Generation Systems*. Cambridge University Press. In press.

Reiter, Ehud, Roma Robertson, and Liesl Osman. 1999. Types of knowledge required to personalise smoking cessation letters. In Werner Horn et al., editors, *Artificial Intelligence and Medicine: Proceedings of AIMDM-1999*, pages 389–399. Springer-Verlag.

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 ……**0**…… boys ……**0**……. 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?** …**20**… 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 ☐ → Q13 **If yes, are you intending to stop smoking within the next month?** YES ☐ NO ☐

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

1

Figure 1: First page of a questionnaire

Dealing with Stress

Finding other ways to deal with stress can help you to stop smoking. Here are a few ideas to try:

- If something is making you stressed, talk about it with someone. Write down what the problem is. Think of two or three new ways of tackling the problem and write these down. Try one of these ways.
- Exercise, sport or any kind of physical activity helps reduce stress. It relaxes your body and your mind. Even gentle physical activity is helpful; things like walking, gardening, DIY, and washing the car can all make you feel less stressed.
- Have a shower or a bath, get out of the house - take a short walk, or make a cup of tea.
- Try this simple exercise:
  - Sit down somewhere comfortable and quiet
  - Take a few deep slow breaths and let your arms and legs relax
  - Close your eyes
  - Think of a peaceful scene, or repeat the word 'relax' over and over to yourself
  - Do this for several minutes

You can get books and leaflets about coping with stress in book shops, libraries and your surgery. These might be useful.

Smokeline is the Scottish helpline for stopping smoking. Calls are free and there is someone to speak to between 12 noon and 12 midnight.

The phone number is: **0800 84 84 84**

Figure 2: Front and back pages of leaflet

The Health Centre

Smoking Information
for
Heather Stewart

from:
The Health Centre &
University of Aberdeen

Dear Ms Stewart

Thank you for taking the trouble to return the smoking questionnaire that we sent you. It appears from your answers that although you're not planning to stop smoking in the near future, you might like to stop if it was easy. Your answers suggest that you don't feel confident that you could stop because *smoking helps you cope with stress, smoking stops you putting on weight, and you don't have the willpower*. However, you have reasons to be confident of success if you did try to stop, and there are ways of coping with the difficulties.

We hope the information in this leaflet will be of interest to you.

Smoking Information for Heather Stewart

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.

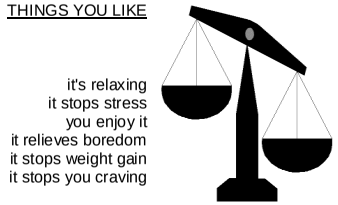

You could do it...

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.

Although you don't feel confident that you would be able to stop if you were to try, you have several things in your favour.

- You have stopped before for more than a month.
- You have good reasons for stopping smoking.
- You expect support from 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 *put on weight*. A few people do put on some weight. If you did stop smoking, your appetite would improve and you would taste your food much better. Because of this it would be wise to plan in advance so that you're not reaching for the biscuit tin all the time. Remember that putting on weight is an overeating problem, not a no-smoking one. You can tackle it later with diet and exercise.

And finally...

We hope this letter will help you feel more confident about giving up cigarettes. If you have a go, you have a real chance of succeeding.

With best wishes,

The Health Centre.

Figure 3: Inside pages of leaflet