HealthDoc: Customizing patient information and health education by medical condition and personal characteristics

Chrysanne DiMarco,* Graeme Hirst, Leo Wanner,* and John Wilkinson***

*Department of Computer Science
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1

**Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 1A4

Abstract

The HealthDoc project aims to provide a comprehensive approach to the customization of patient-information and health-education materials through the development of sophisticated natural language generation systems. We adopt a model of patient education that takes into account patient information ranging from simple medical data to complex cultural beliefs, so that our work provides both an impetus and testbed for research in multicultural health communication. We propose a model of language generation, 'generation by selection and repair', that relies on a 'master-document' representation that pre-determines the basic form and content of a text, yet is amenable to editing and revision for customization. The implementation of this model has so far led to the design of a sentence planner that integrates multiple complex planning tasks and a rich set of ontological and linguistic knowledge sources.

1 Customizing patient-education material

Present-day health-education and patient-information material is often limited in its effectiveness by the need to address it to a wide audience. What is generally produced is either a minimal, generic document that contains only the information common to everyone, or a maximal document that tries to provide all the information that might be relevant to someone (and hence much that is irrelevant to many). But material that contains irrelevant information, or omits relevant information, or that for any other reason just doesn't seem to be addressed to the particular reader is likely to be discounted or ignored, with consequent problems in motivation for compliance with medical regimens, health-related lifestyle improvements, and so on. Recognizing this, health educators have paid much attention to methods of identifying different segments of their audience and their differing needs and constructing material accordingly (see, *e.g.*, many of the papers in Maibach and Parrott 1995), but at some level, the material remains generic.

However, recent experiments have shown that health-education material can be much more effective if it is customized for the individual reader in accordance with their medical conditions, demographic variables, personality profile, or other relevant factors. For example, Strecher and colleagues sent unsolicited leaflets to patients of family practices on topics such as giving up smoking (Strecher *et al.* 1994), improving dietary behaviour (Campbell *et al.* 1994), or having a mammogram (Skinner, Strecher, and Hospers 1994). Each leaflet was 'tailored' to the recipient, on the basis of data gathered from them in an earlier survey. In each study, the 'tailored' leaflets were found to have a significantly greater effect on the patients' behaviour than 'generic' leaflets had upon patients in a control group.

This kind of customization involves much more than just producing each brochure or leaflet in half a dozen different versions for different audiences. Rather, the number of different combinations of factors can easily be in the tens or hundreds of thousands (as in the studies cited in the previous paragraph). While not all distinct combinations might need distinct customizations, it is nonetheless impossible to produce and distribute, in advance of need, the large number of different editions of each publication that is entailed by individual tailoring of health information. Rather, what is needed is a computer system for the production of tailored health-education and patient-education material, that would customize a 'master document' for a particular individual on demand.

The HealthDoc project aims to build such a system. Information from an on-line medical record or from a clinician will be used as the basis for deciding how best to fit the document to the patient.

The development of systems like this is an example of the "demassification" of health communication that Chamberlain (1994) has suggested is one of the possible benefits of the application of new technologies.

2 Organization of the HealthDoc project

The HealthDoc project began in March 1995, following a year of planning, in collaboration with the TechDoc project at Forschungsinstitut für anwendungsorientierte Wissensverarbeitung (FAW), Ulm, Germany (Dietmar Rösner and staff). HealthDoc is centred at the University of Waterloo (Chrysanne DiMarco, Leo Wanner, and students), with additional participants at the University of Toronto (Graeme Hirst and students). The project is funded by Technology Ontario as part of a programme in support of scientific collaboration between the Canadian province of Ontario and the German state of Baden-Württemberg. It is advised by patient-education committees of Sunnybrook Health Sciences Centre (University of Toronto), Massachusetts General Hospital (Boston), and Peel Memorial Hospital (Brampton, Ontario).

The goal of HealthDoc is to develop methods for the creation of customized patient-education material—that is, many different versions of the same document, each tailored to the needs of a particular patient. In this goal, it complements TechDoc, a project on the automatic generation of technical manuals in several languages from a single knowledge source (Rösner and Stede 1994). That is, each project aims to develop systems that can generate many variations of a document from one starting point. In TechDoc, the documents (user manuals for motor vehicles) are addressed to a wide audience, but can be created in several different languages, with adjustments as appropriate in order to be 'natural' in each language. In HealthDoc, the emphasis is on tailoring a document to a single user; multilinguality is not a concern at present. Both projects, therefore, share an interest in the study of choosing among the different ways that an idea might be expressed, and in the development of software for the generation of natural language utterances.

Both projects are building upon the Penman system for language generation that was initially developed at the Information Sciences Institute, University of Southern California (Penman Project 1989), and further developed as KPML ("KOMET-Penman multilingual") at Institut für integrierte Publikations- und Informationsysteme, Gesellschaft für Mathematik und Datenverarbeitung, Darmstadt (Bateman 1995).

3 The conceptual framework of the HealthDoc project

The HealthDoc project aims to develop techniques for producing health-information and patient-education material that is customized to the personal and medical characteristics of the individual patient receiving it. The project is concentrating on the production of printed materials—that is, brochures and leaflets that the patient can take away to read and refer to whenever they wish. Such materials are used extensively in clinical settings for many purposes:

- To educate patients about a particular medical condition and its management: *Treatment choices for breast cancer: The surgery decision*; *Living with diabetes*.
- To tell them how to follow a medical regimen, prepare for a medical procedure, or manage recovery: *Getting ready for your bowel surgery*; *Instructions for patients following hysterectomy*.
- To tell them what is involved in a medical procedure, including its potential risks and benefits: *Information for patients undergoing laparoscopic cholecystectomy*.
- For general health education: *About smoking and pregnancy*.

Although the emphasis of the project is on paper documents, many of the techniques that we are developing will also be applicable in the interactive, hypertext-like systems that others are developing (*e.g.,* Cawsey, Binsted, and Jones 1995; Buchanan *et al.* 1992, 1995).

3.1 The HealthDoc model of patient education

The HealthDoc project is motivated by the idea of building a system that, when complete, would fit into the following framework.

Master documents. Each customized brochure on a particular topic is produced from a ***master document*** on that topic. The master document contains all the information, including illustrations, that the system might wish to include in any individual brochure, along with ***annotations*** as to when each piece of information is relevant. The nature of the master document will be described below in section 5.

Authoring. It is assumed that the master document is created by a medical writer using an authoring tool; see section 6 below.

Dimensions of customization. A HealthDoc brochure may be customized with data about the individual patient, and the selection of content and manner of expression of that content may be determined by the patient's medical condition and their personal and cultural characteristics (see section 4 below). Selection of content may occur at the level of paragraphs, sentences, or phrases.

HealthDoc in the clinical setting. In clinical use, HealthDoc would have access to the on-line medical records of the patients. When the clinician wishes to give a patient a particular brochure from HealthDoc, she selects it from a menu, and specifies the name of the patient to whom it is to be given; in addition, she may optionally provide information to supplement or override that which the system will find in the patient's record.

HealthDoc will then generate a version of the document appropriate to that patient. It may be printed directly, or it may be generated to a file for a word processor so that the clinician may edit it as he sees fit before it is printed. The final document would be attractively laid out and formatted, and possibly run off on pre-printed stationery.

Multilinguality. Ideally, HealthDoc would allow the creation and production of documents in many languages.

3.2 Goals of the present project

The creation of a complete system as just described is well beyond the scope and resources of the HealthDoc research project. The project is concentrating primarily on the research problems in computational linguistics that are entailed by the development of such a system, and in particular the generation of the text of the individual brochures from the master document. Thus the project does not include development of the delivery system: the user interface for the clinician; software interfaces to electronic medical records systems; and document layout, formatting, and printing. Authoring tools will be developed only to a primitive stage that is sufficient for creating prototype examples of master documents, and no attention will be paid to the incorporation of illustrations.

HealthDoc output will be only in English, as master documents are language-dependent (see section 5). It is hoped that, as the system develops, these language dependencies will be reduced and that it will be possible to add grammars and lexicons for other languages, such as German, Spanish, and French, that are developed in TechDoc or other projects. Unfortunately, there is little or no applicable research in the languages that are the greatest problems for our partner hospitals: Chinese, Vietnamese, Khmer.

Despite the restriction to English, the documents produced in the later phases of the HealthDoc project will attempt to account for cultural differences in health beliefs (as well as other individual differences), as this is an important aspect of the need to be able to produce health information in many different ways; see section 4.3 below.

4 Dimensions and levels of customization

A patient-education document may be customized in any or all of three different dimensions: patient data, medical condition, and personal characteristics. The HealthDoc project aims to eventually incorporate all three, though in the early stages, only the first two are considered.

4.1 Patient data

The simplest kind of customization is inclusion of simple numerical or alphabetic data from the patient's record—in effect, filling in the blanks in a template (Reiter 1995). For example:[1]

(1) Prescription information for Morris Browning

Your blood pressure has been measured as 170 over 100, which is too high. Dr Canning has therefore prescribed for you a course of Ranazone 500mg to be taken four times daily before meals.

Template-filling is straightforward, and independent of other kinds of customization. Where we speak below about customization by the creation or inclusion of pieces of text, it is to be understood that these pieces might actually be templates to be further customized by filling with the appropriate data.

4.2 Patient's medical condition

Customization by medical condition is the choice of what to say and not say in the document, depending upon the patient's diagnosis, physical characteristics (such as age and gender), and medical history. For example, a brochure on living with diabetes may omit information pertaining only to type II diabetes if the patient has type I, and vice versa. When several medical conditions interact, the choice of what to include may be quite complex. For example, the customization of a brochure advising a patient on the benefits and risks of hormone-replacement therapy needs to take into account many factors in her medical history and that of her family. It is in such cases that customizable documents will be of particular utility.

4.3 Patient's culture, health beliefs, and other personal characteristics

Customization by patient characteristics involves the choice of both form and content.

Many studies have shown that the 'same' message often needs to be framed or presented in very different ways in order to be communicated most effectively and most persuasively to different people; indeed, what may be persuasive to one person can actually reduce compliance in another (Monahan 1995). In health education, individual differences in health beliefs, perception of and attitude to risk, and level of education are among the factors that must be considered when tailoring a message to an individual. For example, health messages that attempt to arouse high amounts of fear are effective on people with low anxiety, but less so on people with high anxiety (Hale and Dillard 1995); similarly, anti-drug messages are more effective when matched to the individual's degree of need for sensation (Donohew, Palmgreen, and Lorch 1994).

---

[1] The customizations are highlighted here by underlining; of course, they would not be underlined in the actual brochure given to the patient.

This example is constructed to illustrate the point being made; it is not actual HealthDoc output, and makes no claims to medical realism. Indeed, while we have collected a large corpus of patient-education materials of many kinds, we have often found it to be helpful, in thinking about linguistic problems, to use deliberately unrealistic, whimsical examples of our own construction, such as "Advice to patients on procedures for total head replacement". In this way, we ensure that study of hypothetical linguistic situations does not become confounded with the need for medical accuracy, and that test output from the system is not taken as 'medical truth' under circumstances in which this could not possibly be guaranteed.

Often, of course, we have no explicit information on the relevant characteristics of a particular individual. However, we may infer reasonable defaults from the individual's observable characteristics, such as age, gender, ethnicity, and so on. Thus, if the patient is elderly, we could, by default, draw on the observation that the elderly tend to exercise more rational thought when reading a text, but draw fewer inferences from it (Joyce 1994); if the patient is a male of Chinese ethnicity, we could assume that he believes that good health is more a matter of luck rather than his own behaviour, whereas if he is of Caribbean origin we could assume that he believes the converse (Dickinson and Bhatt 1994).

Ethnicity and culture are particularly interesting and difficult variables here. Our partner hospitals serve large, multi-ethnic, multicultural communities, and intercultural communication—indeed, the practice of multicultural health care in general—is a continuing problem for them. Cultures differ widely in their health beliefs and attitudes, and intercultural health communication can often be very difficult for practitioners (Masi 1993). Often, it simply fails: "Ineffective intercultural communication in health care can and often does result in unnecessary pain, suffering, and death" (Kreps and Kunimoto 1994: 8). So tailoring by culture or ethnicity seems to be an ideal application of the customization of health information.

However, ethnicity may be only a weak indicator of culture, for example in well-established immigrant communities that have assimilated elements of the surrounding culture, and culture may be only a weak indicator of an individual's characteristics and beliefs. There can be more variation between individuals in a single ethnocultural group than between canonical members of different groups, and socio-economic class is a better predictor of health beliefs and behaviours than ethnocultural group is (Masi 1988). Generalizations must not become stereotypes, and "must thus be used with caution …; at the same time, some generalizations are necessary" (Masi 1993: 21).

There has been surprisingly little research on methods of achieving effective intercultural communication in health care; the state of the art (*e.g.,* Kreps and Kunimoto 1994) is to try to make health-care practitioners aware of the need for cultural sensitivity, so that they can find out about the health-related characteristics of other cultures and then apply their common sense in any particular situation. Nonetheless, we intend to include customization by factors including culture, in later phases of the HealthDoc project. As research in methods of intercultural health communication proceeds, we will incorporate the results into the HealthDoc framework.

5 The master document and generation by selection and repair

As explained above, a master document is a specification of all the information that might be included in a brochure on a particular topic, along with annotations indicating what is to be included when. We now discuss the nature of this master document.

Many applications of language generation involve the creation of text from some pre-existing knowledge base. The pieces of the knowledge base that are to be expressed, and perhaps also the order in which they are to appear, are selected either by an author (with an authoring tool such as that of Paris *et al.* 1995) or by some automatic process. This is assumed, for example, in multilingual generators such as TechDoc, where the knowledge base has been created anyway as part of the computer-integrated design and manufacturing of the product, and the user manual can then be written "in all languages at once" by this method. Indeed, Reiter, Mellish, and Levine (1995) have argued that natural language generation might be cost-effective for most applications only if the knowledge base is provided essentially free of charge.

However, the absence of a pre-existing knowledge base of content also gives HealthDoc the freedom to specify the form of the master document without any limitations imposed by the need to serve other processes as well. In particular, we can assume a representation that is more oriented towards the eventual production of customized text than an application-independent knowledge base is.

In the simplest kind of customization for content and form, the master document would just be a large set of simple blocks of text (or templates for patient data) to be included or excluded as appropriate for both content and form; the customized leaflets produced by Strecher and colleagues were done this way (Strecher *et al.* 1994; Campbell *et al.* 1994; Skinner, Strecher, and Hospers 1994). At the other end of the spectrum, the elements of the master document would be pieces of a language-independent structure in some knowledge representation formalism, and would be selected for content but not form. These elements would then have to pass through some complete language-generation system that would decide on how to organize and express the content, given information about the form best suited to the patient's personal characteristics.

The knowledge-based approach is elegant and language-independent, but is not yet close to being possible, even with state-of-the-art techniques, for domains as complex as those of interest here. On the other hand, the text-block approach is straightforward but language-dependent. Moreover, it requires that an extremely large number of bits and pieces of text be available: each fact expressed in each possible way.[2] And the assembly of such bits and pieces suffers from the obvious problem that the resulting document might not be coherent or cohesive, or at the very least, not stylistically polished. Therefore, what would be needed would be a process of ***repair***, in which the selected text blocks are reorganized and rewritten; in effect, the selections from the master document would be treated as a rough draft text that would then be subjected to an editing process from which a clear, well-written document would emerge. But that, too, is well beyond the state of the art: it would require nothing less than than an extremely intelligent style-checker with a pretty fair understanding of the meaning of the document.[3]

Our approach, therefore, is a compromise between these two extremes. In the early phases of HealthDoc, our master documents contain sentence plans, but in later phases, as we develop the necessary mechanisms, they will contain text plans from which sentence plans will be derived. Selections are made only for content, as in the knowledge-based approach, but are automatically post-edited not only for form, according to the needs of the individual patient, but for style and coherence, as in the text-block approach.

So we presently represent the master document as a sequence of structures and annotations in SPL, the sentence-plan language that is used as an intermediate representation in the Penman generation system. Each sentence-plan is marked with coreference and coherence relations to other sentences. In contrast to the purely pre-existing representation, it is possible to use pre-existing software to generate sentences from the SPL representation (although improving such software is one of the technical goals of this project). And unlike text blocks, the SPL representation contains information sufficient for the customization of form and the 'repair' of incohesive or incoherent selections. The price paid for these advantages is that the representation is language-dependent.

We regard this use of a master document as a new approach to natural language generation, in which generation from scratch is avoided; 'generation by selection and repair' uses a partially specified, pre-existing document as the starting point. In this way, we can finesse many of the intractable problems of generation, as we start from a document in which many of the decisions have already been predetermined: overall text organization, division of propositional content into sentences, choice of words, and lexical cohesive structure. Even though we might subsequently modify many of these earlier decisions in producing a customized text from the master document, we nonetheless start from a highly useful draft form, rich in linguistic and stylistic information—in effect, we observe the maxim that it is generally much easier to *re*-write than to write.

This approach, while distinct from other research in NLG, exemplifies recent trends in the field. Expectations for text generation research today are not the same as they used to be even a couple of years ago. In the past, researchers expected to develop fully automatic systems, but this has proved to be too difficult. Especially in the selection of the content, proposition ordering, and lexical choice, there are significant unsolved problems. As a consequence, there is an increasing tendency to involve the author or the user in the production process. For example, in the IDAS project (Reiter, Mellish, and Levine 1992) and the Piglet project (Cawsey, Binsted, and Jones 1995), the user's navigation through a hypertext strongly influences the selection of the content that is generated. In the Drafter project (Paris *et al.* 1995), the author chooses the information that is to be generated multilingually. And in HealthDoc, the author of the master document selects from a set of patient features (see section 6 below) to indicate, for each sentence of the master document, when it is to be used.

6 Authoring and the HealthDoc interfaces

We cannot assume the existence of any knowledge base of content to appear in the brochures that HealthDoc is to generate. Rather, its master documents may be based on the natural-language text of pre-existing health-education material, or they may be created from scratch (or some combination of the two). Either alternative requires the involvement of a human.[4]

The author of a master document would normally be a medical writer, who will need to understand the nature of customized and customizable texts but should not be assumed to have any special knowledge or understanding of the innards of HealthDoc. The authoring tool, therefore, should be no more difficult for the author to use than, say, the more-sophisticated features of a typical word processor.

It is the author's job to decide upon the basic elements of the text, the cohesive and coreferential links between them, and the conditions under which each

[2] Sarah Kobrin reports (p.c.) that in extensions to the work described by Strecher *et al.* (1994), the creation and management of the large number of text fragments involved became very difficult.

[3] It might be objected that the pieces of text could be carefully constructed so that all possible selections resulted in a well-formed document. Indeed, Strecher *et al.* (1994) did essentially this. However, they found it difficult even for their fairly simple document (Sarah Kobrin, p.c.); it would surely be very hard to achieve for complex documents unless the granularity were extremely coarse, thereby increasing the number of distinct elements required. In the limit, one would simply store a distinct document pre-written for every single combination of possibilities, a situation that we have already assumed to be impractical.

[4] The fully automatic conversion of natural-language text to the conditional form of a master document would require significant advances in AI and computational linguistics—and much more effort than human-assisted methods.

element should be included in the output. The authoring tool will then assist the author in the creation of the corresponding pieces of the master document, and their annotation with cohesive links and with conditions for inclusion.

The design and development of this tool will be part of a later phase of the project, when the master document is represented as text plans. In the interim, for master documents of SPL sentence plans, an SPL authoring tool, Splat, has been developed (Jakeway 1995). Use of this tool requires a familiarity with SPL.

Splat permits an example-based approach to authoring. The user may view sample sentences, stored in a ***sentence bank***, of previously constructed SPL plans, and can select part or all of a sentence to retrieve the appropriate SPL. Users are able to limit their view of the sentence bank to a subset of the sentences in order to search for one that contains an SPL plan similar to the desired one. The subset is selected by matching each sentence against a pattern-list of items, each of which can be the spelling of a word, a lexical item, a syntactic category, a grammatical function, a semantic category (from the upper model; see section 8.2.4 below), or a wild card.

The main interface for Splat displays the sentence bank and the lists of all previously defined SPL plans. The plans are divided into the four main upper-model categories—objects, qualities, processes, and relations—the identifiers of which are displayed in separate panes. An SPL plan can be viewed or modified by selecting its identifier from an appropriate SPL-plan list, or, if it is used in one of the sentences in the sentence bank, by selecting the appropriate part of that sentence. As well, an SPL plan can be incorporated into other plans in various roles. Splat also allows new SPL plans to be created. An SPL plan can be transformed into working SPL and run through Penman for realization whenever it is being viewed or edited.

An interim user interface to HealthDoc allows the author to access and display a master document that has been built with Splat. The author may test the customization of the document by specifying the characteristics of a hypothetical patient and the relevant sections of the document are selected accordingly. The selections can be 'repaired', as explained in section 7 below, and run through Penman, with the resulting text displayed alongside the master SPL form. The author can also choose to focus on a specific part of the master document for editing, and can use Splat to modify or replace the corresponding SPL.

7 What automatic post-editing must do

We now consider the kinds of textual repairs or post-editing that might be needed when material in the SPL master document is selected during the process of producing a customized version for a particular patient. We will show the examples in English, but it is to be understood that the process is taking place on the underlying SPL representation.

7.1 Coreference and cohesion

Consider the following master document text, which describes the risks of some particular surgery for patients with various kinds of medical conditions:

(2) Patients who have no history of symptomatic cardiac disease generally have a very low risk of perioperative myocardial infarction and less than a 1 percent risk of death from cardiac causes. However, the risks are higher in those who are older or who have cardiovascular disease. The risks of surgery are especially high for someone who has had a very recent myocardial infarction, or who has severe congestive heart failure, advanced atrial or ventricular arrhythmias, or cannot perform moderate exercise. If the patient is unable to exercise, or their medical history is unreliable or incomplete, then additional testing may help to identify whether they are at high risk.

Let's assume that we have an older patient who has had a recent myocardial infarction and is unable to exercise, but has a reliable medical history. We therefore select the portions of the text that deal with such patients. The resulting text, before post-editing, is as follows:

(3) However, <u>the risks</u> are higher in those who are older or who have cardiovascular disease. The risks of surgery are especially high for someone who has had a very recent myocardial infarction, or cannot perform moderate exercise. If the patient is unable to exercise, then additional testing may help to identify whether they are at high risk.[5]

[5]Texts such as this, whether customized or not, may be written in the second person or the third person:

(i) The risks of surgery are especially high for someone who has had a very recent myocardial infarction, …

(ii) The risks of surgery are especially high if you have had a very recent myocardial infarction, …

Customization might introduce a greater bias to the second person, but some patients might find this too threatening:

(iii) Because of your recent myocardial infarction, the risks of surgery are especially high.

(iv) Because of your medical condition, the risks of surgery are especially high.

So the third person should remain an option:

(v) The risks of surgery are especially high for someone who has had a very recent myocardial infarction.

Or even, being as distant as possible:

(vi) The risks of surgery are especially high for patients who have had a very recent myocardial infarction.

However, this has the problem, which customization was supposed to be fixing, that the patient might not recognize that

The first problem is that the sentence is marked as a contrast to a proposition (about risks in typical patients) that was not selected, and so begins (in text form) with the word *however*. This is fixed simply by deleting the relationship.[6] Next, the reference to *the risks* lacks an antecedent; it actually refers to *the risks of perioperative myocardial infarction and death from cardiac causes* in the unused sentence, and its SPL has a pointer back to this. The repair is made by copying this SPL, and then perhaps modifying it if necessary. This is an example of repairing a broken coreference. We now have this:

(4) The risks of perioperative myocardial infarction and death from cardiac causes are higher in those who are older or who have cardiovascular disease …

The text is still flawed, however; the word *higher* is an implicit reference to *very low* and *less than 1 percent*. Explicit incorporation of these referents, however, is infelicitous:

(5) The risks of perioperative myocardial infarction and death from cardiac causes are higher than very low and higher than less than 1 percent in those who are older or who have cardiovascular disease …

What we would like to say is something like this:

(6) The risks of perioperative myocardial infarction and death from cardiac causes are higher than the normal very low level in those who are older or who have cardiovascular disease …

This kind of repair would require 'semantic' aggregation, in which the semantic content of both phrases (*very low* and *less than 1 percent*) is merged to form one phrase, (*the normal very low level*), that is correct and appropriate in the context.

Now let's consider a different patient: a younger patient who has no history of heart disease but has an unreliable medical history. The text selected would be as follows:

(7) Patients who have no history of symptomatic cardiac disease generally have a very low risk of perioperative myocardial infarction and less than a 1 percent risk of death from cardiac causes. If their medical history is unreliable, then additional testing may help to identify whether they are at high risk.

Fortuitously, the anaphor *their* has an antecedent in the text. The structure, however, is not fully coherent: the second sentence is actually an elaboration upon an unused sentence that, in turn, contrasts with the first sentence. The contrast must be restored by inserting a word such as *however* or *but*. It might also be recognized that the latter alternative permits the conjunction of the two sentences:

(8) Patients who have no history of symptomatic cardiac disease generally have a very low risk of perioperative myocardial infarction and less than a 1 percent risk of death from cardiac causes, but if their medical history is unreliable, then additional testing may help to identify whether they are at high risk.

7.2 Aggregation and the elimination of redundancy

In this example, the text provides general information on synthetic and natural implants and specifically mentions the expected lifetime of the latter, which are harvested from a human donor:

(9) The implant cannot be guaranteed to last for any specific amount of time. In some instances, implants wear out, loosen, or fail, and must be removed and replaced. You must understand that a donor head, while in good condition, will have all the wear and tear of the time that it was in use by its previous owner. If the implant wears out, loosens, or fails, it will have to be removed, and you will then need a secondary surgical procedure.

For this example, we assume that the patient is being given a synthetic implant, so that the third sentence, which contains information on natural implants, will not be included:

(10) The implant cannot be guaranteed to last for any specific amount of time. In some instances, implants wear out, loosen, or fail, and must be removed and replaced. If the implant wears out, loosens, or fails, it will have to be removed, and you will then need a secondary surgical procedure.

In the master-document text, the second and fourth sentences expressed the same idea, the second sentence introducing the possibility of implants wearing out, the

they are the kind of person whom the text is talking about. Some linguistic compromises are possible:

(vii) The risks of surgery are especially high for patients, such as yourself, who have had a very recent myocardial infarction.

To some degree, the decision to use the second or third person is for the author of the master document to make. However, as we move towards greater customization of form, the author might want the system to have some control over this aspect of the text, perhaps taking into account the patient's level of trait anxiety (Hale and Dillard 1995): 'copers' could get a second-person form, whereas 'avoiders' would get a third-person form.

In example (3), we have used the third person simply because the example is based on a generic text that was written in the third person.

[6] In most cases, this contrast relationship will never be used, because the two sentences describe mutually exclusive situations, and only one or the other will be selected. The exceptions arise when the patient's relevant medical history is unknown, and both alternatives must therfore be presented.

fourth sentence reinforcing this information. But in the selected text, these sentences are adjacent and sound awkwardly repetitive instead of strongly supportive. One way to improve the text is to introduce a generic phrase *if this happens* as a reference to the repeated sequence of events (*wear out, loosen, or fail*). This in turn necessitates 'proposition chunking' to remove one constituent of the preceding sentence (*removed and replaced*) and adjoin it to the following one (*… in which your implant is removed and replaced*).

(11) The implant cannot be guaranteed to last for any specific amount of time. In some instances, implants wear out, loosen, or fail. If this happens, you will need a secondary surgical procedure, in which your implant is removed and replaced.

Another way to repair this text is to decide whether the two sentences could be effectively conflated, to erase the repetition but retain the rhetorical strength.

(12) The implant cannot be guaranteed to last for any specific amount of time. If it wears out, loosens, or fails, it will have to be removed, and you will then need a secondary surgical procedure, in which your implant is replaced.

This conflation requires a form of sentence restructuring in which a difficult decision must be made: only one of the discourse relations involving the two original segments can be retained; one must be chosen as the more salient. In the draft text (10), the relationship between that of the first sentence and the second was that of circumstance, and that between the second and third was elaboration. In the repaired text (12), the conflated sentence, the second, stands in a circumstance relation to the first, but the elaboration relation has disappeared.

In general, selecting material from pre-existing text and then editing it to recover coherence and cohesion can involve a wide range of problems in various aspects of sentence planning. For example, both syntactic and semantic aggregation may be needed, as well as chunking of whole and partial propositions. The whole referential and cohesive structure of the text may be flawed, with deictic and anaphoric references lacking an antecedent. And, of course, aggregation and sentence restructuring will have an obvious effect on the discourse relations.

8 Sentence planning: Architecture and process

8.1 The HealthDoc model of language generation

The notion of generation by selection and repair and the consideration of the kinds of repairs that would be needed to generate customized texts from the master document bring us to the question of the nature of the generation system that would support this approach.

In our model of language generation, shown in Figure 1, the selections made from the master document are text plans—groups of propositions, represented in some abstract formalism, that are to be uttered. The text plans are passed to a discourse relation–planning stage, which uses knowledge of rhetorical relations to produce a discourse-structure representation for the text. The structured propositions then pass to the sentence-planning stage, which is concerned with thematization and focus control, constituent ordering, sentence aggregation, proposition chunking, reference relations, and lexical choice. The output of this stage is a sequence of sentence specifications that are then passed to the realization stage to determine an appropriate surface form.

We envision that the HealthDoc generator, when completed, will contain all the components of this model. In the first phase of the project, we are implementing the sentence planner, with particular emphasis on the processes of aggregation and reference. For this initial phase, as we explained in section 5, the elements that are selected from the master document for generation are chunks of SPL that are ready for input to sentence planning, but which might be incoherent or non-cohesive. During the sentence-planning process, the SPL structures are modified ('repaired') to recover textual coherence and cohesion.

As the complete system is developed, the nature of the master document will change to reflect the inclusion of additional pieces of the model. The representation used will evolve from SPL to a more abstract specification, suitable for input to the discourse structure–planning stage. The underlying idea—that the master document provides, in some sense, a pre-existing draft that specifies and guides the composition of a new text—remains the centre of the model of generation by selection and repair.

8.2 An overview of the sentence planner

In sentence planning, there will often be strong interaction between the various planning tasks such as aggregation and choice of reference, and one module might make a decision that conflicts with that of another module or that prevents another module from making any decision at all. For example, there may be a conflict between aggregation and focus or salience, as reflected in the choice of active or passive voice and verb tense. Thus, we may have, without conflict:

(13) You will be examined before being operated upon.

But if there is focus on the agent of *examine*, a conflict occurs, as this choice of aggregration would be ill-formed:

(14) *Dr Canning will examine you before being operated upon.

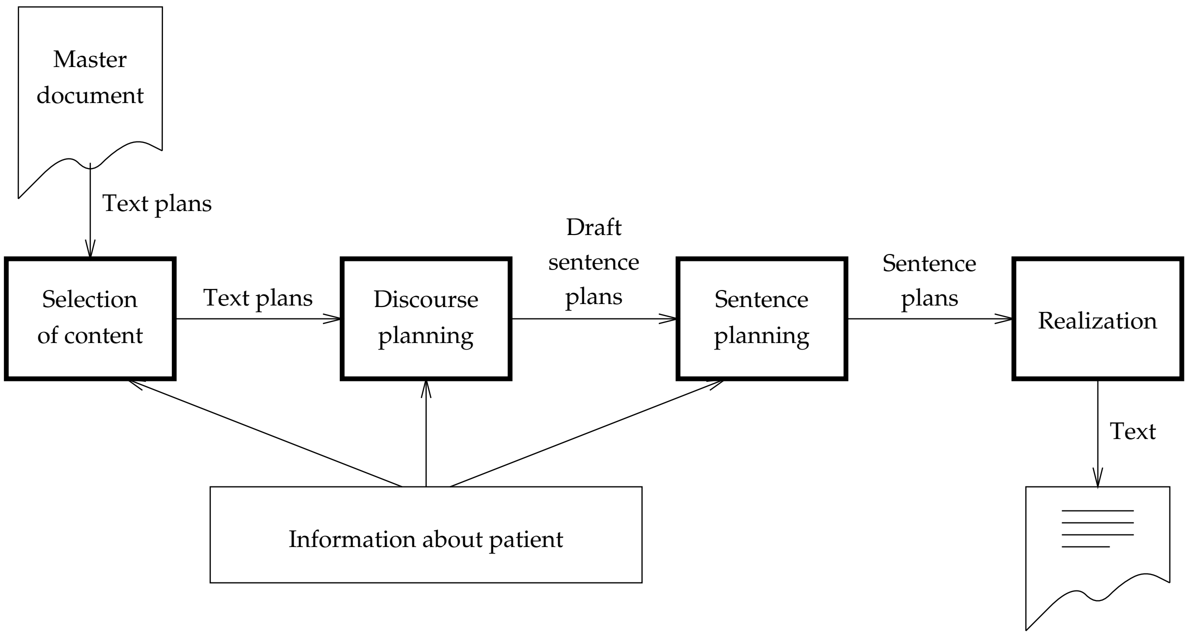

Figure 1: Our model of language generation. Boxes with heavy lines represent processes, and boxes with light lines represent sources of information; the arrows represent flow of information.

The resolution here is for focus to precede aggregation, forcing the latter to make a different choice:

(15) Dr Canning will examine you before you are operated upon.

But no simple fixed set of priorities for planning tasks is possible; sometimes, one might have to take precedence, and sometimes another.

This consideration suggests that a ***blackboard architecture*** would be most appropriate for our purposes. A blackboard architecture generally consists of three major parts (Nii 1989): the ***blackboard***, a passive data structure that records the current stage of the problem-solving process; the ***processes*** that operate on the blackboard; and the ***control mechanism***. The planning modules will operate in parallel, communicating via the blackboard and none constraining the others. If a conflict arises, the control mechanism will dynamically assign priorities to the conflicting modules, using the information that they have posted on the blackboard to decide which should dominate which.

Figure 2 shows the blackboard architecture for our sentence planner. The major components are a set of four blackboards, controlled by an administrator; a set of four planning modules; and a set of ontological and linguistic knowledge sources. Each component will now be described in detail.

8.2.1 The blackboards and the administrator

There are four blackboards in the system, and an administrator that controls them:

Control blackboard. Contains control flags for each module to report its present state (*active*, *waiting*, or *idle*), and an administrator-controlled flag (*start*) for each, which is used in conflict-resolution to allow one module to start before the others.

Input blackboard. Contains a copy of the master document, marked as to which SPL plans have been selected for the patient.

Output blackboard. Contains the intermediate suggestions from each of the planning modules for the creation of the output SPL. When planning is complete, the structure on this blackboard will be output to the realization stage.

Knowledge blackboard. Contains requests from the sentence-planning modules for information from one another or from the knowledge sources, and the answers to these requests.

Each planning module posts its suggestions on the output blackboard, together with an indication of the type of rules that led to its decision. The blackboard administrator reviews the information posted and passes control between the modules. More specifically, it performs the following tasks:

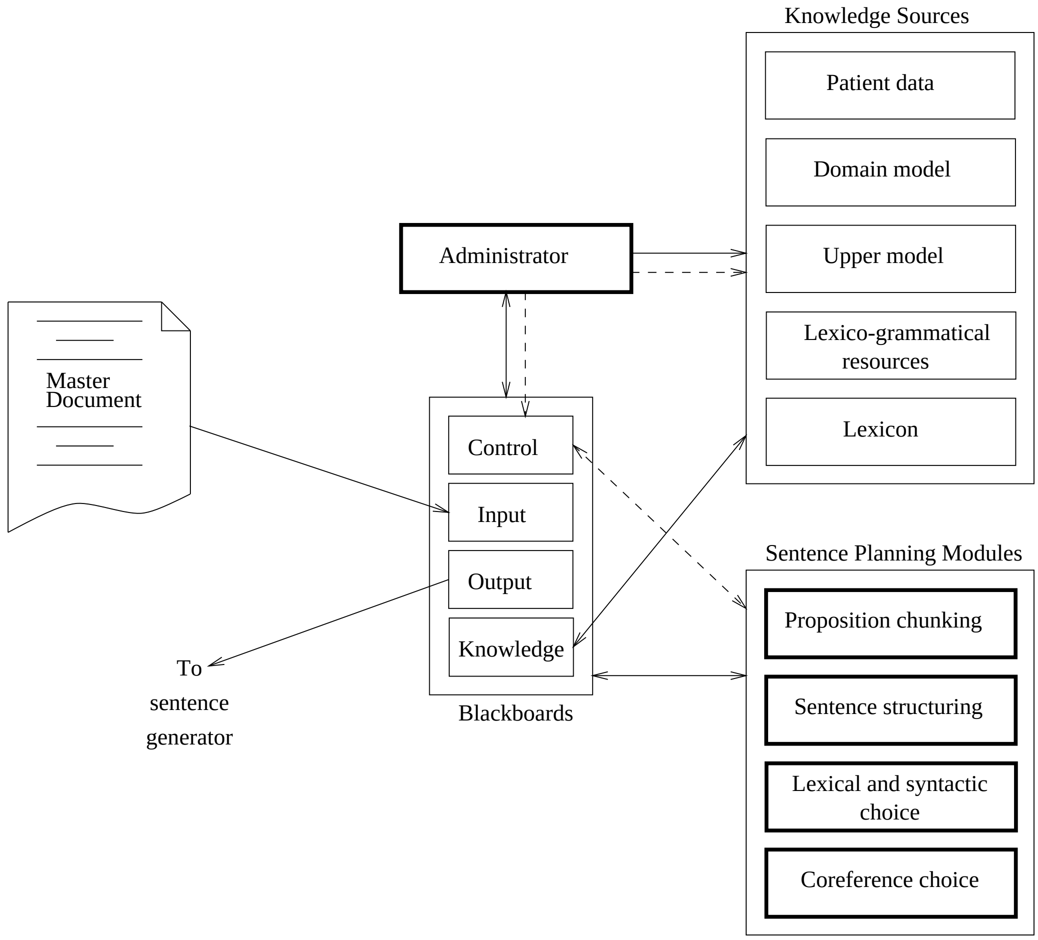

Figure 2: The architecture of the sentence planner. Boxes with heavy lines represent processes, and boxes with light lines represent sources of information; solid arrows represent flow of information and dotted arrows represent flow of control.

- It determines which module should answer a query that has been posted to the knowledge blackboard.
- It resolves conflicts between the modules.
- It updates the SPL plan structures on the output blackboard.

As this process continues, the sentence planner can be seen as a set of parallel, mutually constraining modules that together settle on an optimal sentence plan. When necessary, an administrator-controlled sequential ordering is enforced.

8.2.2 Conflict recognition and resolution

The conflict-recognition strategy is a unification process: the options that are proposed by different modules, represented in terms of features and their realization statements, are unified. If the unification process fails, a conflict has occurred. In this case, the original SPL structure is replaced by a partial structure in which unification has been carried out to the extent possible, and conflict resolution begins.

In cases in which no conflicts occur or no specific linguistic motivations give preference to a particular module, the planning processes operate in parallel. If conflicts do occur, then a resolution mechanism is invoked. To resolve a conflict, priority is given to the alternative that is most important to the immediate goals of the planner, or to one for which no alternative choice is available. Modules in conflict with this resolution will be forced to revise their choices. This resolution strategy has the effect of imposing a situation-specific ordering upon the modules.

8.2.3 The planning modules

The planning modules post their suggested options on the output blackboard in the form of a triple containing the following information:

- A set of knowledge-source features with their realization statements (see section 8.2.4 below); this is the module's actual suggestion.
- The name of the SPL-structure fragment, from the master document, that is to be replaced.
- The rules that led to this suggestion.

Aggregation. There are two kinds of choice to be made in aggregation: ***what*** propositions should be aggregated, and ***how*** this combination of propositions should be structured. These choices are sufficiently distinct to warrant implementation as two separate planning modules. The first is for ***proposition chunking***—choosing the semantic units to be packaged as a single proposition. The second module does ***sentence structuring***—choosing the structures to realize the proposition.

Lexical choice and concomitant syntactic choice. The lexical choice module chooses lexical units, with the exception of those that realize coreferences. These choices constrain syntactic structure within clauses, that is, at a more delicate level than the choices made by the sentence structuring module.

Coreference choice. The coreference choice module first considers, for each referring expression, whether there is a previous reference to the same entity in the text to be uttered; if so, it will select an appropriate anaphor or definite reference.

8.2.4 The knowledge sources

Four main knowledge sources are used by the sentence planner: the ***domain model***, the ***upper model***, ***lexico-grammatical*** information, and the ***lexicon***.

The upper model (Bateman 1990) is Penman's domain-independent conceptual hierarchy; we are using the standard one provided with the KPML system. The domain model is linked to the upper model, but its concepts are derived from the application domain. The domain model, as yet, contains only the specific concepts derived from our sample master documents; we will not attempt to create a full medical ontology. (Several groups are developing such knowledge bases from the Unified Medical Language System (Lindbergh, Humphreys, and McCray 1993), and we expect to use their results)

The linguistic knowledge sources are organized in terms of system networks; that is, they are fully expressed by a set of features organized into a network of choices. In addition, they contain ***realization*** statements that are used to build the SPL plan for a sentence. It is as a consequence of this that the planning modules' options can be formulated solely in terms of sets of features and their realization statements. These lexico-grammatical resources differ in content and organization from those used in the Penman grammar, Nigel: they also contain referential information, lexical co-occurrence information, and so on. They are organized to support the derivation of SPL structures rather than to support syntactic realization.

The lexico-grammatical information is maintained at three linguistic ranks: the ***discourse***, ***proposition***, and ***constituent*** ranks. Despite its name, the discourse-rank information is used by the sentence planner to recognize how two sentences could be aggregated, or to recognize relations between different clauses within one sentence. The information consists of a discourse-structure relation network (Hovy *et al.* 1992), made up of three parallel subnetworks, each corresponding to a systemic-linguistic metafunction: ***ideational***, ***textual***, and ***interpersonal***. In addition, a fourth subnetwork determines possible variations in the syntactic realization of different discourse-structure relations; this network corresponds to the ***logical*** systemic metafunction.

The proposition rank contains a network that decides upon possible realizations of predications and whether or not to realize specific arguments (***constituents***). For the time being, three distinct, parallel subnetworks are used, each standing for one systemic-linguistic metafunction at this rank:

- A subnetwork that determines the semantics of a predication and the arguments of it that are to be realized (***experiential*** metafunction).
- A subnetwork that determines salience variations (***textual*** metafunction).
- A subnetwork that decides upon the syntactic realization of the predication (***logical*** metafunction).

The constituent rank contains a network that decides upon possible realizations of constituents and their modifiers. Again, parallel subnetworks deal with each of the above three metafunctions.

For the first phase of the project, we will use the lexicon that comes with Penman's grammar and extend it by adding two different types of information: lexical co-occurrence information and qualia-structure information. Co-occurrence information for a lexical item describes other lexical units that collocate with it or for which it is a collocate (*e.g.*, *perform* and *undergo* are collocates of *surgery*). Qualia structure is a system of relations that characterizes the lexical semantics of nominals (Pustejovsky 1991).

In addition to these four main knowledge sources, information about the individual patient is also available to influence the form of the output. At present, this is extremely rudimentary, but will be developed in later phases of the project, as we concentrate more on the customization of form.

8.3 A quick comparison with Diogenes

The architecture of the HealthDoc sentence planner is similar in many ways to that of Diogenes (Nirenburg, Lesser, and Nyberg 1989; Nirenburg, Carbonell, Tomita, and Goodman 1992), but with a number of improvements. Both systems have blackboard architectures with various sentence-planning tasks performed by separate modules. While Diogenes uses a pure blackboard architecture with static priorities assigned to each class of module, HealthDoc uses parallel processing as the default and administrator-determined serial processing with dynamic priorities when conflicts arise. Also, HealthDoc has a somewhat different distribution of tasks among the planning modules. As yet, we have no counterpart to the adjective-ordering module in Diogenes, but Williams (1995) has developed a more-flexible constituent-ordering mechanism, and this will eventually form part of a complete ordering module in the sentence planner.

9 Conclusion

The HealthDoc project aims to provide a comprehensive approach to the customization of patient-information and health-education materials through the development of sophisticated natural language generation systems. We have adopted a model of patient education that takes into account patient information ranging from simple medical data to complex cultural beliefs. We have proposed a model of language generation, 'generation by selection and repair', which relies on a 'master-document' representation that pre-determines the basic form and content of a text and yet is amenable to editing and revision for customization. The implementation of this model has so far led to the design of a blackboard-based sentence planner that integrates multiple complex planning tasks and a rich set of ontological and linguistic knowledge sources. The HealthDoc project is proving to be a strong impetus for research and development in natural language generation, with particular relevance to health communication, and a number of important issues for research have been raised during the first phase of the project.

Basis for customization of patient education. The whole issue of customized health communication points out the near-total lack of adequate methods for tailoring messages to individual patients. Although the need for such customization has been recognized, there has as yet been little in the way of a concerted effort to alleviate the lack of understanding of how information may be conveyed most effectively to patients to motivate a change in their behaviour. In the next stage of the project, identifying critical examples of variations in text by medical condition will be an important task.

Evaluation of models of intercultural health communication. One of the most difficult, yet most important, goals of the project is to customize texts for patients' culture, health beliefs, and other personal characteristics. As we have noted, there has been very little research on methods of achieving effective multicultural communication in health care. One of the benefits we can offer to the health care community is a means of generating many different versions of a brochure for clinical evaluation. As an educational resource, HealthDoc could provide patient educators with a significant tool for their research. In this way, the HealthDoc project can be both an impetus and testbed for research in multicultural health communication.

Persuasion and linguistic style. In all aspects of personalized health communication, we must confront the question of how to frame an effective, persuasive message to an individual patient. In subsequent phases of our project, we aim to adapt and further develop research on persuasion and linguistic style for application to health communication.

Development of master document. Our consideration of how patient-information documents should be initially written and then customized has led us to propose the use of a master document. But the nature of the master document may need to be redesigned as we begin to address questions of stylistic and pragmatic customization, such as the incorporation of persuasive effects. At present, the master document is a set of sentence plans, but it lacks the information needed to do the kind of whole-scale revision that would be needed for this level of pragmatic customization. We need to replace the present form with one that allows additional specifications of discourse-level, semantic, and stylistic information.

Refinement of the generation paradigm. Our paradigm of generation by selection and repair is appealing, as it promises a way to reduce or avoid many of the intractable problems of generation, but its relation to the master document and its effect on the system architecture remain problems for further study. As the master document evolves, the nature of the kinds of selections and repairs that are needed will become clearer. In turn, this will affect the characteristics of the various 'repair' planning modules and the ways they interact.

Development of the sentence planner. Many issues in sentence planning will be addressed as we continue to study the nature of customization. A critical problem is the distribution of planning tasks among the modules, as there are often strong interactions. The responsibilities of each module and the overlaps between them remain an open problem for our sentence-planning research. As we build up our knowledge of how customization of the texts will be done, we will be revising and extending the architecture of the sentence planner.

Development of tools for authoring. As HealthDoc evolves, the canonical author will be a medical writer, not a system developer, and this means that the interface to the generation system must become much more intelligent and automated. The current sentence-plan authoring tool is a beginning, but we will need to develop a full-scale authoring tool for text plans.

Open questions in generation. Among the many research issues in language generation that the HealthDoc project will need to deal with are the following:

- What is the best architecture for a sentence planner?
- How do we specify individual sentence-planning tasks?
- How are conflicts between planning modules recognized and resolved?
- What balance is needed between parallel and sequential processing?
- What is the role of lexical choice? How does it interact with the choice of references, coreferences, and clausal structure?
- What are the various kinds of aggregation? How does it interact with lexical choice? What kinds of conflicts can occur?
- How can stylistic control be built into a generator?

Acknowledgements

The HealthDoc Project is supported by a grant from Technology Ontario, administered by the Information Technology Research Centre. Vic DiCiccio was instrumental in helping us to obtain the grant, and has been invaluable in subsequent administration. The other members of the HealthDoc Project have contributed to the work described here, especially Bruce Jakeway, Susan Williams, Phil Edmonds, and Steve Banks. Victor Strecher and Sarah Kobrin kindly gave us details of their project. We are grateful to Dominic Covvey, Eduard Hovy, John Bateman, Brigitte Grote, Manfred Stede, Dietmar Rösner, and the patient-education committees of our partner hospitals for helpful advice, insightful discussions, and other contributions.

References

Bateman, John Arnold (1990). "Upper modeling: Organizing knowledge for natural language processing". *Proceedings of the Fifth International Workshop on Natural Language Generation*, Dawson, PA, June 1990, 54–61.

Bateman, John Arnold (1995). "KPML: The KOMET–Penman multilingual linguistic resource development environment." *Proceedings, 5th European Workshop in Natural Language Generation*, Leiden, May 1995, 219–222.

Buchanan, Bruce G.; Moore, Johanna Doris; Forsythe, Diana E.; Banks, Gordon; Ohlsson, Stellan (1992). "Involving patients in health care: Explanation in the clinical setting." *Proceedings, 16th Annual Symposium on Computer Applications in Medical Care*, 510-514, 1992.

Buchanan, Bruce; Moore, Johanna Doris; Forsythe, Diana E.; Carenini, Giuseppe; Ohlsson, Stallan; and Banks, Gordon (1995). "An intelligent interactive system for delivering individualized information to patients." *Artificial Intelligence in Medicine*, to appear.

Campbell, Marci Kramish; DeVellis, Brenda M.; Strecher, Victor J.; Ammerman, Alics S.; DeVellis, Robert F.; and Sandler, Robert S. (1994). "Improving dietary behavior: The effectiveness of tailored messages in primary care settings." *American Journal of Public Health*, **84**(5), May 1994, 783–787.

Cawsey, Alison; Binsted, Kim; and Jones, Ray (1995). "Personalized explanations for patient education." *Proceedings, 5th European Workshop in Natural Language Generation*, Leiden, May 1995, 57–94.

Chamberlain, Michael A. (1994). "New technologies in health communication: Progress or panacea?" In Ratzan 1994, 271–284.

Dickinson, Roger and Bhatt, Arvind (1994). "Ethnicity, health and control: Results from an exploratory study of ethnic minority communities' attitudes to health." *Health Education Journal*, **53**(4), December 1994, 421–429.

Donohew, Lewis; Palmgreen, Philip; and Lorch, Elizabeth Pugzles (1994). "Attention, need for sensation, and health communication campaigns." In Ratzan 1994, 310–322.

Hale, Jerold L. and Dillard, James Price (1995). "Fear appeals in health promotion campaigns: Too much, too little, or just right?" In Maibach and Parrott 1995, 65–80.

Hovy, Eduard Hendrik; Lavid, Julia; Maier, Elisabeth; Mittal, Vibhu; and Paris, Cécile (1992). "Employing knowledge resources in a new text planner architecture." *Aspects of automated natural language generation: The 6th International Workshop on Natural Language Generation [Trento, Italy]: Proceedings*, edited by Robert Dale, Eduard H. Hovy, Dietmar Rösner, and Oliviero Stock. Springer-Verlag (Lecture notes in artificial intelligence, volume 587), 57–72.

Jakeway, Philip Bruce (1995). *Splat: A sentence plan authoring tool for natural language generation.* MMath thesis, Department of Computer Science, University of Waterloo, forthcoming.

Joyce, Mary L. (1994). "The graying of America: Implications and opportunities for health marketers." In Ratzan 1994, 341–350.

Kreps, Gary L. and Kunimoto, Elizabeth N. (1994). *Effective communication in multicultural health care settings.* Thousand Oaks, CA: Sage Publications.

Lindbergh, D.A.B.; Humphreys, Betsy L.; and McCray, Alexa T. (1993). "The Unified Medical Language System." *Methods of information in Medicine*, **32**(4), August 1993, 281–291.

Maibach, Edward and Parrott, Roxanne Louiselle (1995). *Designning health messages: Approaches from communication theory and public health practice*. Thousand Oaks, CA: Sage Publications.

Masi, Ralph (1988). "Multiculturalism, medicine and health, Part I: Multicultural health care." *Canadian Family Physician*, **34**, October 1988, 2173–2178.

Masi, Ralph (1993). "Multicultural health: Principles and policies." In *Health and cultures: Exploring the relationships. Volume I: Policies, professional practice and education*, edited by Ralph Masi, Lynette Mensah, and Keith A. McLeod. Oakville, Ontario: Mosaic Press, 11–22.

Monahan, Jennifer L. (1995). "Thinking positively: Using positive affect when designing health messages." In Maibach and Parrott 1995, 81–98.

Nii, H. Penny (1989). "Introduction." In *Blackboard architectures and applications*, edited by V. Jagannathan, Rajendra Dodhiawala, and Lawrence S. Baum, (Perspectives in artificial intelligence, volume 3), Boston: Academic Press, xix–xxix.

Nirenburg, Sergei; Lesser, Victor; Nyberg, Eric (1989). "Controlling a language generation planner." *Proceedings, 11th International Joint Conference on Artificial Intelligence*, Detroit, August 1989, 1524–1530.

Nirenburg, Sergei; Carbonell, Jaime; Tomita, Masaru; and Goodman, Kenneth (1992). *Machine translation: A* *knowledge-based approach.* San Mateo, CA: Morgan Kaufmann Publishers.

Paris, Cécile; Vander Linden, Keith; Fischer, Markus; Hartley, Anthony; Pemberton, Lyn; Power, Richard; and Scott, Donia (1995). "Drafter: A drafting tool for producing multilingual instructions." *Proceedings, 5th European Workshop in Natural Language Generation*, Leiden, May 1995, 239–242.

Penman Natural Language Group (1989). "The Penman primer", "The Penman user guide", and "The Penman reference manual." Information Sciences Institute, University of Southern California.

Pustejovsky, James (1991). "The generative lexicon." *Computational Linguistics*, **17**(4), 409–441.

Ratzan, Scott C. (editor) (1994). *Health communication: Challenges for the 21st century.* Published as a special issue of *American Behavioral Scientist*, **38**(2), November 1994, 197–380.

Reiter, Ehud (1995). "NLG vs. templates." *Proceedings, 5th European Workshop in Natural Language Generation*, Leiden, May 1995, 95–105.

Reiter, Ehud; Mellish, Chris; and Levine, John (1992). "Automatic generation of on-line documentation in the IDAS project." *Proceedings, Third Conference on Applied Natural Language Processing*, Trento, Italy, April 1992, 64–71.

Reiter, Ehud; Mellish, Chris; and Levine, John (1995). "Automatic generation of technical documentation." *Applied Artificial Intelligence*, **9**(?), to appear.

Rösner, Dietmar and Stede, Manfred. "Generating multilingual documents from a knowledge base: The TECHDOC project." *Proceedings, 15th International Conference on Computational Linguistics (COLING-94)*, Kyoto, August 1994, 339–342 and 346.

Skinner, Celette Sugg; Strecher, Victor J.; and Hospers, Harm (1994). "Physicians' recommendations for mammography: Do tailored messages make a difference?" *American Journal of Public Health*, **84**(1), January 1994, 43–49.

Strecher, Victor J.; Kreuter, Matthew; Den Boer, Dirk-Jan; Kobrin, Sarah; Hospers, Harm J; and Skinner Celette S. (1994). "The effects of computer-tailored smoking cessation messages in family practice settings." *The Journal of Family Practice*, **39**(3), September 1994, 262–270.

Williams, Susan (1995). *A flexible constituent-ordering mechanism for the Penman generation system.* MMath thesis, Department of Computer Science, University of Waterloo, forthcoming.