

Authoring and Generating Health-Education Documents That Are Tailored to the Needs of the Individual Patient

Graeme Hirst¹, Chrysanne DiMarco², Eduard Hovy³, and Kimberley Parsons^{2*}

¹ Department of Computer Science, University of Toronto, Canada

² Department of Computer Science, University of Waterloo, Canada

³ Information Sciences Institute, University of Southern California, U.S.A.

Abstract. Health-education documents can be much more effective in achieving patient compliance if they are customized for individual readers. For this purpose, a medical record can be thought of as an extremely detailed user model of a reader of such a document. The HealthDoc project is developing methods for producing health-information and patient-education documents that are tailored to the individual personal and medical characteristics of the patients who receive them. Information from an on-line medical record or from a clinician will be used as the primary basis for deciding how best to fit the document to the patient. In this paper, we describe our research on three aspects of the project: the kinds of tailoring that are appropriate for health-education documents; the nature of a tailorable master document, and how it can be created; and the linguistic problems that arise when a tailored instance of the document is to be generated.

1 The Value of Tailored Health-Education Documents

Health-education and patient-information brochures and leaflets 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.*
- For general health education: *About smoking and pregnancy.*

* 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 Steve Banks, Phil Edmonds, Mary Ellen Foster, Bruce Jakeway, Jon Litchfield, Daniel Marcu, Peter Vanderheyden, Leo Wanner, John Wilkinson, and Susan Williams. Victor Strecher and Sarah Kobrin kindly discussed details of their research with us. We are grateful to Dominic Covvey, Brigitte Grote, Manfred Stede, Dietmar Rösner, John Bateman, and the patient-education committees of our partner hospitals—Sunnybrook Health Sciences Centre (University of Toronto), Massachusetts General Hospital (Boston), and Peel Memorial Hospital (Brampton, Ontario)—for helpful advice, insightful discussions, and other contributions.

Recent experiments have shown that such health-education documents can be much more effective if they are customized for individual readers 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, taking into account such characteristics as reasons for smoking or belief in the efficacy of mammography; this information had been asked of each patient 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 the behaviour of patients in a control group.

However, in these experiments a ten-minute interview was required with each patient to elicit the information necessary for tailoring the document. This amount of effort is not, in general, practical. Nor can this problem be avoided by naïvely assuming that patients will just go Web surfing to seek out the health information that they need, volunteering their demographic or medical profiles to some on-line tailoring system.¹ On the contrary, much health education must be initiated by the clinician in response to the patient's medical situation, and the information must generally be presented on paper for the patient to refer to later. Fortunately, in such clinical situations, much of the information that is needed for tailoring health-education material is available in the patient's medical record. Indeed, a medical record can be thought of as an extremely detailed user model for (potential) readers of health-education documents.

This paper describes research undertaken in the HealthDoc project, which is developing text-generation methods for producing health-information and patient-education material that is tailored to the personal and medical characteristics of the individual patient receiving it. Information from an on-line medical record or from a clinician will be used as the primary basis for deciding how best to fit the document to the patient; but reader models derived from other sources, such as interviews or surveys, could also be used. Moreover, while the project is concentrating on the production of printed materials, much of the research will also be applicable to the creation of tailored Web pages and interactive, hypertext-like health-education systems that we and others are developing (e.g., DiMarco and Foster, 1997; Cawsey, Binsted, and Jones, 1995; Buchanan et al., 1995).

The structure of the HealthDoc system is shown in Figure 1. The major components will be described as we discuss our research in the sections that follow, concentrating on three aspects of the project: the kinds of tailoring that are appropriate for health-education documents; the nature of a tailorable master document, and how it can be created; and the linguistic problems that arise when a tailored instance of the document is to be generated. We assume the following model for use of the system:

Master documents. Each tailored brochure on a particular topic is produced from a *master document* on that topic, which has been created by a professional medical writer, using an authoring tool that we will describe in Section 4 below. The master document contains all the information, including illustrations, that might possibly be included in any individual brochure, along with *annotations* as to the conditions under which each piece of information is relevant. The nature of the master document will be described below in Section 3.1.

¹ For the few patients who do, companies such as MicroMass Communications Inc. have produced some tailorable health-education 'magazines'; see <http://www.micromass.com/demos.html>.

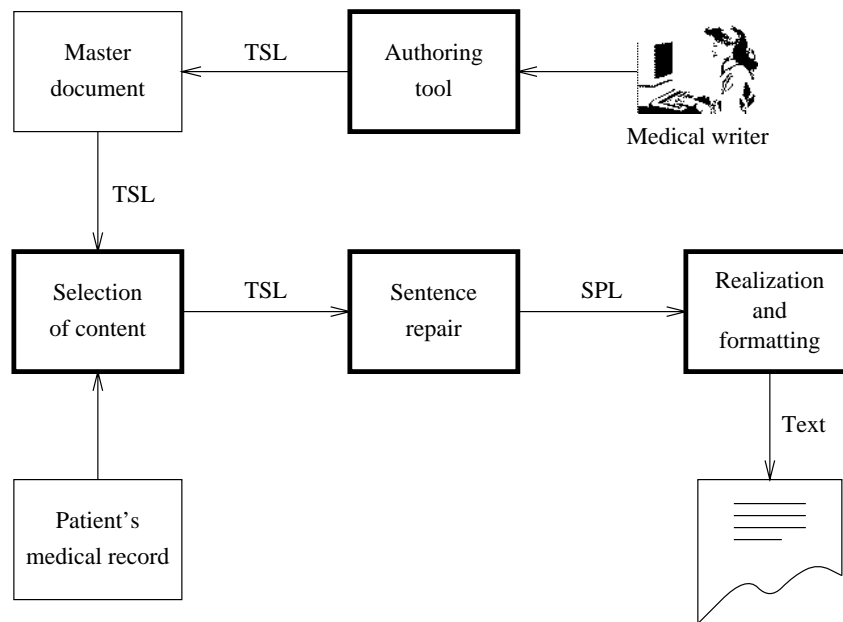


Figure 1. Generation of a tailored document by the HealthDoc system. Boxes with heavy lines represent processes, and boxes with light lines represent sources of information; the arrows represent flow of information.

Dimensions of tailoring. A HealthDoc brochure may be tailored in accordance with information about the individual patient: both the selection of content and the manner of expression of that content may be determined by the patient's medical condition and their personal and cultural characteristics (see Section 2.1 below).

HealthDoc in the clinical setting. In clinical use, HealthDoc will have access to the on-line medical records of patients. When the clinician wishes to give a patient a particular brochure from HealthDoc, she selects it from a menu of master documents, and specifies the name of the patient to whom it is to be given; in addition, she may offer, or be asked to provide, information to supplement 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 desired before it is printed. The final document will be attractively laid out and formatted, and possibly run off on pre-printed stationery.²

² The creation of a complete system as just described is well beyond the current scope and resources of the HealthDoc research project. We are concentrating on authoring and sentence repair, and therefore

2 Tailoring Patient-Education Material

2.1 Classes of Patient Characteristics

A HealthDoc brochure may be tailored for an individual patient. The selection of content of the brochure and manner of expression of that content may be determined by both the patient's medical condition and any other personal and cultural characteristics that might either be included in their medical record or available from the clinician.

Patient data. The simplest kind of tailoring is inclusion of simple numerical or alphabetic data from the patient's record, such as the name of the patient or of a prescribed medication—in effect, filling in the blanks in a template (Reiter, 1995). Template-filling is straightforward, and independent of other kinds of tailoring. Where we speak below about tailoring by the creation or inclusion of pieces of text, it is to be understood that these pieces might actually be not complete text but rather templates that are to be further customized by filling with the appropriate data.

Patient's medical condition. Tailoring by medical condition entails choosing what to say and not say in the document, in accordance with the patient's diagnosis, physical characteristics (such as age and gender), and medical history. For example, a brochure on living with diabetes may include information on how diabetes interacts with the patient's other known medical conditions, such as heart disease. When several medical conditions interact, the choice of what to include and exclude can become quite complex. For example, the tailoring of a brochure advising a patient on the benefits and risks of hormone-replacement therapy needs to take into account a large number of interacting factors in her medical history and that of her family. In such cases, tailored documents will be of particular utility.

Patient's culture, health beliefs, and other personal characteristics. Tailoring by patient characteristics involves the choice of both form and content. Many health communication 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 and cultural 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 (Masi, 1993; Kreps and Kunimoto, 1994). 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).

Characteristics that are of particular interest to us at this stage of the project include locus of control (the degree to which the patient regards herself as being 'in charge' of her health), ability or desire to read technical detail, and the degree to which appeals to authority in the presentation of information are persuasive to the patient. For more discussion, see DiMarco, et al. (1995).

other parts of the system—i.e., the user interface for the clinician, the software interface to the on-line medical-records system, and the module for document layout, formatting, and printing—are only simple demonstration prototypes. To avoid a commitment to any of the emerging standards for electronic medical records and the interchange of health information, we use a fictional medical database in our own idiosyncratic format. Implementation with any particular system of on-line medical records will require the adaptation of our query format to that of the system.

2.2 What's in the Medical Record?

When a medical author creates a HealthDoc text that is to be tailored in one or more of the ways described above, he or she must know what information is likely to be present in a patient's on-line medical record, or must at least make assumptions as to what information is available. (An electronic medical record may contain free text or scanned documents in addition to structured data; thus the information required might be present and yet not readily available. The extraction of information from heterogeneous electronic medical records is a research problem in itself.) And present-day systems are unlikely to offer the kinds of non-clinical information that will often be important in tailored health education, such as culture, level of education, or locus of control. We believe that as electronic *medical* records start to become electronic *health* records in the not-too-distant future, this kind of information will become more readily available. In any case, HealthDoc will query the clinician user for any characteristic of the patient that it cannot obtain from the on-line record. The medical writer is thus free to use any patient characteristic in tailoring that he or she wishes, while considering that it would be a burden on the clinician if too many characteristics cannot be found in the on-line record.

Regardless of what information a medical records system offers in principle, the information might not, in practice, be available for the particular patient in question, neither from the system nor the clinician. The writer must therefore always consider what the default action should be when some characteristic of the reader is unknown. The default could be to include selections for *all* possible values of the characteristic, or to use instead a distinct, more-generic selection; or a default value for the characteristic might be assumed. For example, in a brochure on diabetes, if it is not known whether a patient has the insulin-dependent or non-insulin-dependent form, one would probably choose to give information on both. But one would probably not include information on the interaction of diabetes with a rare or unusual medical condition unless it were known for certain that this was relevant to the particular patient.

3 Representing a Tailorable Document

3.1 Finding an Appropriate Level of Abstraction

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 and the problems of combining selections from it.

In an AI-heavy approach, the elements of the master document would be pieces of a language-independent structure in some knowledge representation ('KR') formalism, which would be selected for content, as appropriate for the particular patient, but not form. These elements would then have to pass through some complete language-generation system that would decide how to organize and express the content, given information about the form best suited to the patient's personal characteristics. This 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, Strecher and colleagues created the tailored texts that they used in their experiments simply by building a large set of simple snippets of text that were included or

excluded as appropriate for both the content and form of a patient's brochure. This method is straightforward, but it requires that an extremely large number of bits and pieces of text be available; Strecher and colleagues found that the creation and management of the large number of text fragments involved became extremely difficult (Victor Strecher and Sarah Kobrin, personal communication).

Moreover, 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. It might be thought that the snippets of text could be constructed, or alternative expressions of the same idea be written where necessary, so that all possible selections would result in a well-formed document; indeed, Strecher and colleagues attempted to do essentially this. However, they found it extremely difficult to achieve, even for their fairly simple document (Victor Strecher and Sarah Kobrin, personal communication); it would surely be quite impossible for complex documents, unless the granularity were extremely coarse, thereby increasing the number of distinct elements required. In the limit, one would not tailor the text at all, but simply store thousands of distinct documents, each pre-written for every single combination of possibilities—a situation that is quite impractical.

Our approach, therefore, is to use neither a KR formalism nor snippets of text, but an abstract, albeit language-dependent, text specification language, TSL, which we will describe in Section 3.2 below. This language expresses not only the content of the document and the conditions under which each element is to be selected for an individual patient, but also information that assists a subsequent process of generating coherent, well-polished text. Selections from this document are made for both content and form, as in the text-snippet approach, but are automatically post-edited—'repaired'—for form, style, and coherence. We will discuss the nature of these 'repairs' in Section 5 below. Because the repairs take place upon the abstract representation, and are guided by the additional information that it contains, the process is much simpler than would be required for revision of an assemblage of text snippets.

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. The approach is discussed at greater length by DiMarco, Hirst, and Hovy (1997).

3.2 Text Specification Language

Text Specification Language, or TSL, is the language used to represent master documents in the HealthDoc system. TSL not only expresses the content of the master document but also includes annotations on each element of that content (both textual and non-textual, at all levels of document granularity), giving the circumstances under which the element is to be selected for use. TSL annotations can also provide information—coreference links and rhetorical relations—that guides the repair of the selected text. An example of a TSL representation of a sentence is shown in Figure 2.

TSL represents a sentence both in English and in the sentence plan language SPL. The latter is used by the Penman text generation system (Penman Natural Language Group, 1989; Bateman, 1995) that is incorporated into HealthDoc (see Section 6 below).³ An SPL expression

³ TSL can actually accommodate multiple representations of a sentence. For example, our WebbeDoc project (DiMarco and Foster, 1997) uses TSL with sentences marked up in HTML.

```

(variation
 :name 'var2
 :condition '(AND high-cholesterol
              (OR insulin-dependent non-insulin-dependent)
              low-technical)
 :sentence-list '(sent2) )

(sentence
 :name 'sent2
 :english "High blood cholesterol levels raise the risk of heart disease."
 :focus chol1
 :coref '( (hd2 specific comp0) (lev9 generic lev8) (risk6 generic risk1) )
 :spl
 '(raise / nondirected-action
  :lex raise
  :tense present
  :actor (lev9 / abstraction
         :lex level
         :number plural
         :class-ascription (chol1 / abstraction
                          :lex blood-cholesterol)
         :property-ascription (high / quality
                              :lex high))
  :actee (risk6 / abstraction
         :lex risk
         :determiner the
         :part-of (hd2 / abstraction
                  :lex disease
                  :determiner zero
                  :class-ascription (heart1 / object
                                   :lex heart))))))

```

Figure 2. Example of TSL representation of a sentence.

is an abstract specification of a sentence, which Penman can convert to the corresponding surface form. This permits expression of the content of the document. These basic SPL structures are augmented with information for selection and repair. An example of an annotation for selection may be seen in the slot named `:condition` in Figure 2; the sentence will be selected if the patient has high cholesterol, either insulin-dependent or non-insulin-dependent diabetes, and is to receive a minimum of technical detail in the document.

Other annotations are used to guide later repairs to the text. *Coreference links* join two or more references to a single object that occur in different parts of the document. A link also indicates the kind of reference: definite, indefinite, generic, or intensional, and so on. Thus, it will always be known if two different sentences refer to the same thing, and pronominalization can occur

accordingly (see Section 5 below). For example, the coreference information in the `:coref` slot in Figure 2 indicates that `hd2`, representing the reference to heart disease, is a specific reference that is coreferential with `comp0`, which is mentioned in some other sentence. The other elements in the list represent the other two references in the sentence, i.e., to high blood cholesterol levels and to the risk of heart disease.

Rhetorical relations are cohesive relationships between sentences, such as *cause*, *effect*, *elaboration*, and so on. All such relationships between sentences in the document are recorded in the TSL, so that explicit discourse connectives can be used in the text where appropriate (see Section 5). For the example shown in Figure 2, the rhetorical relations would be listed at a higher level in the TSL and might look like this:

```
:relations '( (evidence sent1 sent2)
              (elaboration sent2 sent3) )
```

4 Authoring a Tailorable Document

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 a human and an authoring tool.

The creator of a master document would normally be a professional medical writer, who will need to understand the nature of tailored and tailorable texts, but who should not be assumed to have any special knowledge or understanding of TSL or the innards of HealthDoc. The authoring tool, therefore, should be no more difficult for the writer to use than, say, the more-sophisticated features of a typical word processor. The text is therefore written in English, and semi-automatically translated to TSL (see below). (The English source text is retained in the TSL for use in subsequent authoring sessions—for example, if the document is updated or amended.)

It is the writer's job to decide upon the basic elements of the text, the rhetorical and coreferential links between them, and the conditions under which each element should be included in the output. The elements of the text are typed into the authoring tool, and are marked up by the writer with links for cohesion and coreference and with conditions for inclusion. The conditions for inclusion are, of course, queries on the medical record of the patient for whom a tailored copy is to be produced.

Figure 3 shows a snapshot of the authoring tool in use (Parsons, 1997). The sample text shown is part of a section about health risks in diabetes, from a brochure about this condition. The left-hand portion of the screen, labeled *Selection Criteria*, contains a list of the patient conditions that the author is using to specify the selection of pieces of text. The right-hand portion of the screen contains a window on the text of the master document. Each box in the view contains a piece of text and the inclusion conditions for that piece of text. Groups of boxes represent mutually-exclusive pieces. The rhetorical relations between sentences are represented by arrows drawn between the related boxes. Using the mouse, the author specifies the two boxes that are related. A window containing a list of possible rhetorical relations appears; the relations are colour-coded, so when the author chooses a relation from the list, an arrow is drawn between the two boxes, its colour indicating the relation that was specified. Coreference relations are also colour-coded. Each reference to the same object or concept (e.g., heart disease) is specified by the author by highlighting the reference and clicking with the mouse. A window that contains the

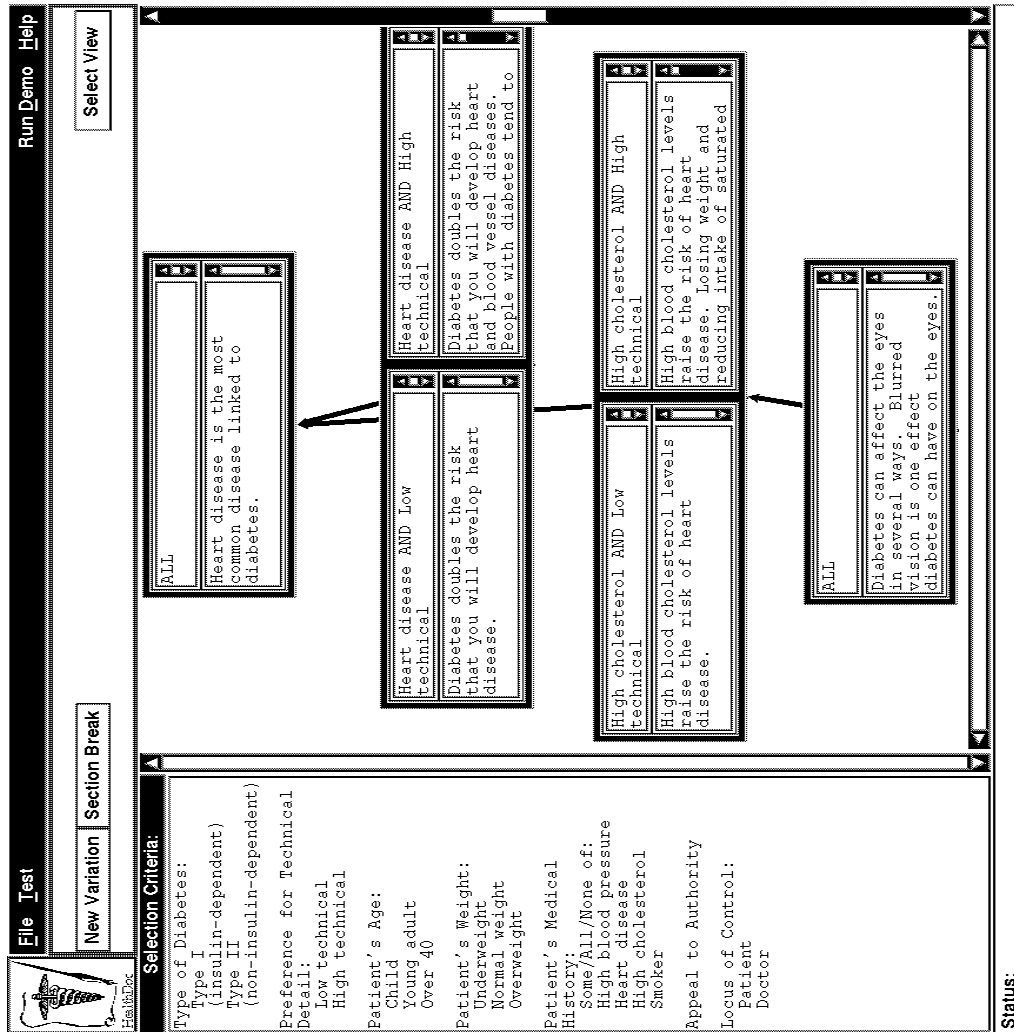


Figure 3. A screen from the authoring tool.

lists of coreference links pops up, and the author specifies the list that the current object should be added to. The reference changes colour to match those with which it is coreferential (e.g., all references to heart disease might become blue).

After the document has been written, the text is semi-automatically translated into SPL. This is essentially a process of parsing, but the resultant structures are (annotated) SPL expressions

rather than parse trees. Whenever an ambiguity cannot be resolved, the writer is queried in an easy-to-understand form. The process is described in detail by DiMarco and Banks (1997).

5 Tailoring a Document

We now consider the kinds of textual repairs or post-editing that might be needed after material in the TSL master document has been selected 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 TSL representation.

In general, selecting material from pre-existing text and then editing it to recover coherence and cohesion involves a wide range of problems in sentence planning. Our sentence planner uses a blackboard architecture in which individual repair modules communicate and resolve their conflicts with one another. The architecture is described in greater detail by Hovy and Wanner (1996). Four repair modules are being developed in the present phase of the project: for coreference, for discourse structuring and rhetorical relations, for aggregation, and for constituent ordering. Here, we will give examples of the first two.

Repair—Coreference. Coreference repairs include decisions as to when a reference should be pronominalized. Suppose, for example, we have a master document that contains the following two sentences (in TSL, of course):

- (1) *People with respiratory disorders* have a high risk of developing Glaumann's disease.
- (2) *People with respiratory disorders* should take immediate action to quit smoking.

The italicized noun phrase in each arises from a coreference link to the same object in the list of references. So if, upon selection, these two sentences turned up in close proximity, the first occurrence of the reference would be marked for realization as a full noun phrase, and the second would be marked for pronominalization:

- (3) *People with respiratory disorders* have a high risk of developing Glaumann's disease. *They* should take immediate action to quit smoking.

Repair—Rhetorical relations. If two selected sentences are marked as being rhetorically related, the repair module will consider making the cohesive relationship between the sentences explicit. For example, if the sentences in the example above were marked as *evidence* and *conclusion*, then the word *therefore* might be inserted:

- (4) People with respiratory disorders have a high risk of developing Glaumann's disease. They should *therefore* take immediate action to quit smoking.

6 Realization and Formatting

The final specifications for the repaired text, represented in SPL, are passed to the realization stage, which uses KPML (Bateman, 1995), a descendant of Penman, to generate an appropriate surface form in English. A formatter then lays out the text attractively and adds headings and illustrations for final printing.

7 Conclusion

In the HealthDoc system, we are tailoring health-education and patient-information documents by using the medical record as a model of the reader that allows us to select appropriate elements from a master document encoded in TSL. A subsequent process of ‘repair’ ensures that the selections form a coherent, linguistically well-formed document. We have adopted a model of patient education that takes into account patient information ranging from simple medical data to complex cultural beliefs. A number of important issues for research in tailorable health-communication documents and their authoring have been raised during the first phase of the project.

The basis for tailoring health information to a given individual. Although the need for tailored health communication has been recognized, there has as yet been little research on how information may be conveyed most effectively to a known individual to motivate a change in their behaviour. (The present state of the art is represented by Kreps and Kunimoto, 1994,; and Maibach and Parrott, 1995.) In the next stage of the project, identifying critical examples of variations in text by medical condition and by culture and health beliefs will be an important task.

Authoring tailorable documents. We have not yet worked with any real-life medical writers, and hence have not tested our assumptions as to whether medical writers would be able to design tailorable documents with our authoring tool as we have conceived it—or even in principle.

Language dependence. Both the authoring tools and the processes that refine the selections from the master document are necessarily language-dependent, so at present HealthDoc is limited to English, our working language. It is hoped that in the long term it will be possible to add master documents in other languages, such as German, Spanish, and French, for which the necessary grammars and lexicons are being developed in other projects in natural language generation. Unfortunately, there is little or no applicable research in the languages—Chinese, Vietnamese, Khmer—that are the greatest problems for the hospitals with which we are collaborating.⁴

Development of the master document. Our consideration of how patient-information documents should be initially written and then tailored has led us to propose the use of a master document. However, the nature of the master document may need to be redesigned as we begin to address questions of stylistic and pragmatic tailoring, such as the incorporation of persuasive effects. The master document is a set of TSL fragments, but it lacks the information needed to do the kind of whole-scale revision that would be needed for this level of pragmatic tailoring. We expect to augment TSL with additional fields for discourse-level, semantic, and stylistic information.

References

- Bateman, J.A. (1995). KPML: The KOMET–Penman multilingual linguistic resource development environment. *Proceedings, 5th European Workshop on Natural Language Generation*, Leiden, May 1995, 219–222.

⁴ All three of our partner hospitals serve large, multi-ethnic, multicultural communities, and intercultural and multilingual communication—indeed, the practice of multicultural health care in general—is a continuing problem for them.

- Buchanan, B., Moore, J.D., Forsythe, D.E., Carenini, G., Ohlsson, S., and Banks, G. (1995). An intelligent interactive system for delivering individualized information to patients. *Artificial Intelligence in Medicine*, 7:117–154.
- Campbell, M.K., DeVellis, B.M., Strecher, V.J., Ammerman, A.S., DeVellis, R.F., and Sandler, R.S. (1994). Improving dietary behavior: The effectiveness of tailored messages in primary care settings. *American Journal of Public Health*, 84:783–787.
- Cawsey, A., Binsted, K., and Jones, R. (1995). Personalized explanations for patient education. *Proceedings, 5th European Workshop on Natural Language Generation*, Leiden, May 1995, 59–74.
- DiMarco, C. and Banks, S. (1997). Using subsumption classification on a stylistic hierarchy in the multi-stage conversion of natural language text to sentence plans. In preparation.
- DiMarco, C. and Foster, M.E. (1997). The automated generation of Web documents that are tailored to the individual reader. *Proceedings, AAI Spring Symposium on Natural Language Processing on the World Wide Web*, Stanford University, March 1997.
- DiMarco, C., Hirst, G., and Hovy, E. (1997). “Rewriting is easier than writing”: Generation by selection and repair in the HealthDoc project. In preparation.
- DiMarco, C., Hirst, G., Wanner, L., and Wilkinson, J. (1995). HealthDoc: Customizing patient information and health education by medical condition and personal characteristics. *Workshop on Artificial Intelligence in Patient Education*, Glasgow, August 1995.
- Donohew, L., Palmgreen, P., and Lorch, E.P. (1994). Attention, need for sensation, and health communication campaigns. *American Behavioral Scientist*, 38:310–322.
- Hale, J.L. and Dillard, J.P. (1995). Fear appeals in health promotion campaigns: Too much, too little, or just right? In Maibach and Parrott 1995, 65–80.
- Hovy, E.H. and Wanner, L. (1996). Managing sentence planning requirements. *Proceedings, ECAI-96 Workshop ‘Gaps and Bridges’: New Directions in Planning and Natural Language Generation*, Budapest, August 1996.
- Maibach, E. and Parrott, R.L. (1995). *Designing health messages: Approaches from communication theory and public health practice*. Thousand Oaks, CA: Sage Publications.
- Masi, R. (1993). Multicultural health: Principles and policies. In Masi, R., Mensah, L., and McLeod, K.A., eds., *Health and cultures: Exploring the relationships. Volume I: Policies, professional practice and education*. Oakville, Ontario: Mosaic Press, 11–22.
- Monahan, J.L. (1995). Thinking positively: Using positive affect when designing health messages. In Maibach and Parrott 1995, 81–98.
- Parsons, K. (1997). *An Authoring Tool for Customizable Documents*. M.Math. thesis, Department of Computer Science, University of Waterloo, forthcoming.
- Penman Natural Language Group (1989). The Penman primer, The Penman user guide, and The Penman reference manual. Information Sciences Institute, University of Southern California.
- Reiter, E. (1995). NLG vs. templates. *Proceedings, 5th European Workshop on Natural Language Generation*, Leiden, May 1995, 95–105.
- Skinner, C.S., Strecher, V.J., and Hospers, H. (1994). Physicians’ recommendations for mammography: Do tailored messages make a difference? *American Journal of Public Health*, 84:43–49.
- Strecher, V.J., Kreuter, M., Den Boer, D.-J., Kobrin, S., Hospers, H.J., and Skinner C.S. (1994). The effects of computer-tailored smoking cessation messages in family practice settings. *The Journal of Family Practice*, 39:262–270.