Authoring and Generation of Tailored Preoperative Patient Education Materials

Chrysanne Di Marco[1], Peter Bray[2], Dominic Covvey[1], Don Cowan[1], Vic Di Ciccio[1], Eduard Hovy[3], Joan Lipa[2], and Cathy Yang[1]

[1] University of Waterloo
Waterloo, Ontario, Canada
cdimarco@uwaterloo.ca
[2] Division of Plastic Surgery, Department of Surgery and Department of Surgical Oncology, University Health Network, University of Toronto
Toronto, Ontario, Canada
[3] Information Sciences Institute
University of Southern California
USA

1 Introduction

In reconstructive surgery, multiple interventions during one surgical episode are common. Each intervention must be explained, its intended and potential consequences articulated, and informed consent of the patient secured. Although the pre-surgical encounter between the patient and the surgeon is the opportunity to accomplish this, it is essential that the patient be given educational materials to complement and augment face-to-face exchange. This is virtually impossible to do well with brochures, because many combinations of procedures are possible, different patients have different concerns, and patients have varying levels of literacy and knowledge. In the extreme, a patient would either be given a set of brochures selected from hundreds of variants, or every patient could be given the same set of brochures without regard for differing needs. Neither of these scenarios is tractable or acceptable.

We propose a solution allowing divergence from the generic, static, preoperative information brochure to one that is customized for every individual patient regardless of the complexity of the surgical intervention. This solution will require reformulation, extension, and optimization of an existing Natural Language tailoring engine and creation of a database of educational modules pertaining to each subcomponent of a given surgical intervention. A key outcome of this research will be an authoring tool that will assist surgeons in entering the text content that will be assembled into coherent material by the tailoring engine.

This research will provide important tools to assist in patient-centric healthcare: a means of shaping complex information so that it is more relevant and personalized, a mechanism for assisting in the achievement of informed consent to procedures, a method that has been shown to improve patient engagement and compliance with medical regimens, and a technique for complementing and reinforcing the information communicated during the pre-surgical encounter. The authoring tool and tailoring engine will form a robust architecture to allow providers to expand the educational scope beyond reconstructive surgery to all forms of medical intervention, surgical or otherwise.

2 The Importance of Tailoring in Patient Education

2.1 The Problem with Current Patient Education Materials

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.

However, recent experiments suggest 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 [24], improving dietary behaviour [6], or having a mammogram [23]. 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.

Recently, researchers in Natural Language Generation have begun to apply methods from Artificial Intelligence and Computational Linguistics to develop automated systems for tailoring health information to individual patients ([4], [7], [20], [3], [21]). The HealthDoc Project (1994–1999) [12] developed a method for generating tailored documents based on a new paradigm for Natural Language Generation— 'generation-by-selection-and-repair'—in which new documents are created from a pre-existing 'master document' which contains all the pieces of text that might be needed in tailoring a version of the document for any particular audience. Selections from the master document are made for both content and form, and then are automatically post-edited—'*repaired*'—for form, style, and coherence.

In a realistic and usable implementation, the HealthDoc approach requires a sophisticated *authoring tool* to assist the writer, and a *sentence planner* (cf. [26]) that would undertake to repair and polish the selected text—we can't expect the average technical writer to pre-compile all the possible combinations in advance. To develop such a system, a number of research issues need to be addressed: representation of the master document; authoring and knowledge-based document management; and sentence planning for automated post-editing.

2.2 The Potential Solution: Natural Language Generation

The creation of the input material for Natural Language Generation systems is a problem for all generation systems, including our selection-and-repair paradigm. The concept of ‘preparing’ a database, knowledge base, or other resource for natural language generation has been used by other researchers—for example, O’Donnell *et al* [17] manually incorporate in the generator’s database additional information (including a taxonomic organization of the types used in the database) that will be used to ensure coherent and high-quality text. This idea led us to adopt the *authoring* of a ‘database’ of reusable text (i.e., the master document) as the basis for the paradigm of generation-by-selection-and-repair.

Other approaches to natural language authoring have been developed (e.g., [13], [19]), and Brun *et al* [5] point to an ‘an emerging paradigm of “natural language authoring”’ (p.25) which they contrast to the (pure) natural language generation approach as one in which ‘the semantic input is provided interactively by a person rather than by a program accessing digital knowledge presentations’ (p. 25). Scott *et al* [22] present a solution to the problem of authoring input for language generation systems in which the user operates directly upon a knowledge model from which the final output text will subsequently be generated.

Our approach to authoring for natural language generation systems falls somewhere between the paradigm described by Brun *et al* and that of classic language generation: as others do with authoring-based systems, we allow a user to enter the exact textual input that will later be used in generating new texts ([18] [14]), but we are also dealing with authoring of input at a deeper level of linguistic representation ([15], [1], [2]), as is typical of Natural Language Generation systems.

A focus in the original HealthDoc Project was on the development of authoring tools that would be used by a professional programmer or computational linguist to automate the preparation of input specifications for a document generation system at the deep level of linguistic representation needed for the subsequent process of textual repair. For authoring in health situations, however, typically the authoring is accomplished through the interaction of the health professional with a ‘knowledge engineer’, someone trained in structured knowledge acquisition. Our intent is to design a system, based on our paradigm of Natural Language Generation by selection-and-reassembly, strategic planning, knowledge structuring, and a formal model of learning, which interacts directly with the surgeon to allow entry of purpose-specific and patient-specific textual variations in ordinary English which will then be selected, processed, and assembled by our tailoring engine into readable, patient-specific, educational material.

2.3 The Need for Tailored Patient Education in Reconstructive Surgery

Modern reconstructive plastic surgery has evolved into a highly complex field aimed at restoration of patient form and function. The surgical solution to a given reconstructive problem may require grafts of various types (skin, bone, and tendon) combined with tissue-mobilizing procedures (flaps) from among dozens of potential locations on the body. Each reconstruction will have different implications for aesthetics, function, rehabilitation, recovery, and potential complications, all of which must be reviewed with the patient preoperatively.

The fraction of this information that is actually retained by the patient after the consultation is consistently rather small. In many surgical specialties, brochures, Internet websites, and other forms of 'take-home' educational materials are frequently used to supplement the surgeon-patient consultation and enhance patient retention of information. However, such solutions have proven impractical for much of reconstructive plastic surgery due to the sheer number of techniques available and their frequent need to be performed in combinations. The complexity of the surgical procedure and the variety of options that need to be considered in tailoring documentation to the individual patient make the creation of appropriate material a combinatorially explosive process. Figure 1 illustrates the complexity inherent in choosing among the surgical options available in breast reconstruction.

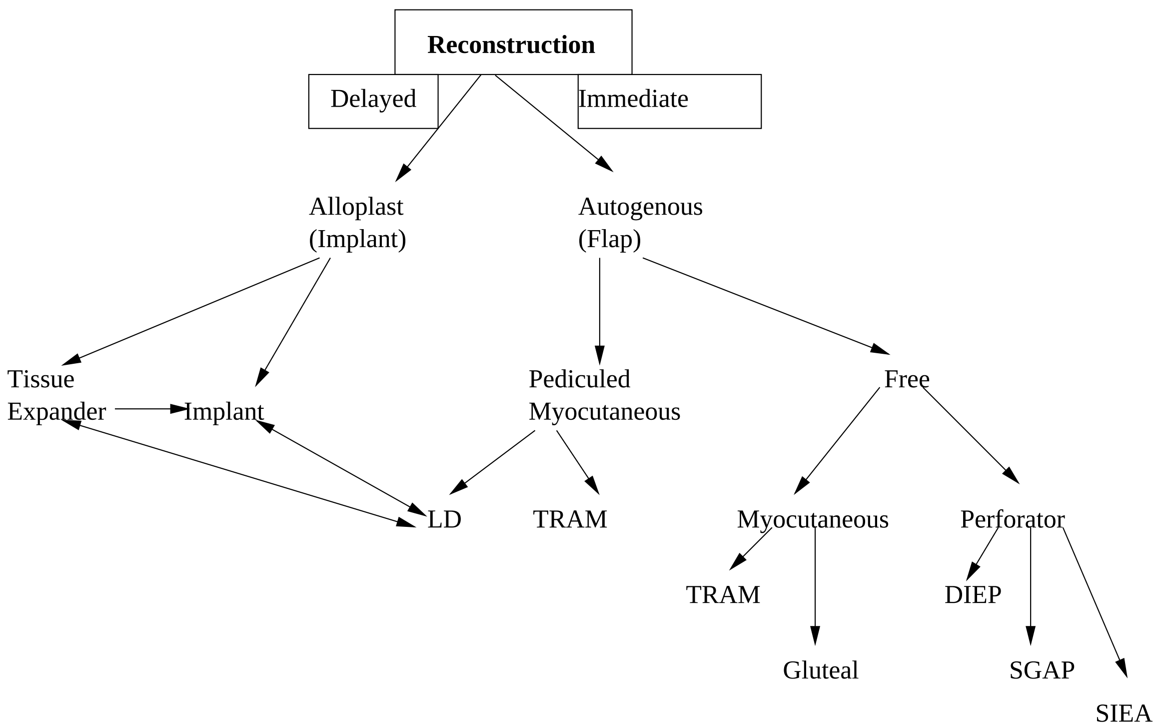

Fig. 1. Decision tree of surgical options in breast reconstruction

Although preoperative information brochures have documented value for patient education, a library of static documents would be difficult to establish if it were to encompass all reconstructive surgical alternatives. For a patient undergoing a multistep procedure, a handful of brochures would be required, which would lack cohesiveness, and would likely be very confusing. Consequently, existing preoperative information brochures are only available for the most common reconstructive surgical procedures and must, by necessity, remain generic in nature to ensure applicability to all patients.

Creation of a tailored information document, customized for every individual patient would potentially increase relevance and effectiveness of the educational material. The tailoring process would permit inclusion, exclusion and/or modification of educational information based on a variety of criteria, including the surgical procedure(s) being performed, impact of adjuvant therapies, medical co-morbidities, and potentially any other factor deemed significant. Although no amount of supplemental documentation can replace the surgeon-patient dialogue with which informed consent is obtained, it is well-documented that only a small fraction of the information communicated in this process is actually retained by the patient. Reference material for review by patient, friends, and significant others would have great value in the preoperative, perioperative, and postoperative stages if this information could be tailored to the individual patient. This observation is supported by recent work in patient education attesting to the potential value of increasing patient involvement in the surgery decision through patient-centred methods [25] and using quality information brochures to improve surgeon-patient communication [16].

We are developing a system for generating preoperative patient education materials that allows divergence from the generic, static, preoperative information brochure to one that tailors the text to every individual patient regardless of the complexity of the surgical intervention. The components of this system will consist of a Natural Language Generation tailoring system, content authoring environment, and creation of a database of educational modules pertaining to each subcomponent of a given surgical intervention.

3 Components of a Tailoring System for Reconstructive Surgery

Creation of a corpus of textual variants. We are creating a corpus of textual variants that will be used in generating tailored educational materials for reconstructive breast surgery by a process of selection and reassembly using the HealthDoc model of document generation. Beginning with the initial generic content, we are applying a formal organizational structure that mirrors the stages of the surgical procedure. Each component of the surgical procedure will then be broken down into subcomponents for which textual variants will be created based on various patient modifiers.

The subcomponents, called *content modules*, include: technical summary, preoperative workup, postoperative course, sequelae, complications, discharge planning, recovery, and rehabilitation. Patient modifiers include: timing of reconstruction, mastectomy type, radiation treatment, smoking, obesity, diabetes, and other comorbidities. The textual variants will initially be entered manually by a programmer into our master-document format and subsequently authored by a patient-education writer using the prototype authoring tool being developed.

An authoring tool to guide health care providers. In previous work in the original HealthDoc Project, we developed several authoring tools ([15], [18], [1]) for the creation of text variants that could be represented in the master-document format and used to generate customized documents by the tailoring engine. However, none of these tools was geared to the domain expert; rather, they were intended for a programmer or computational linguist who would specify the content at a deep level of linguistic representation required to do syntactic and semantic repair of reassembled text. We are developing an authoring environment for health-care providers that will guide surgeons to directly enter the text variants in ordinary English that will then be used to create the tailored educational material.

Although the earlier authoring tools could be used to enter text at various levels of linguistic representation, there was no 'knowledge-level' modelling for knowledge acquisition to support the generation of tailored educational materials. At the knowledge level of authoring tailored content, the physician would be guided through the process of considering the concerns of the various stakeholders (e.g, surgeons, patients, hospital) with regard to tailoring the educational material. For example, the surgeon may be primarily concerned with communicating information that will ensure patient compliance with the recommended treatment and that will lead to favourable outcomes; the patient may be most concerned with the variations in risks and complications associated with the different treatment options. The authoring tool should therefore ideally embody a cognitive model that aids the physician in mapping out these complementary, and sometimes contradictory, high-level concerns. Yang [27] has developed a design methodology for an authoring tool that uses a Constructivist model of patient-centred learning to guide the physician through the process of creating the master-document framework.

The Constructivist approach [11] assumes that learners construct their own knowledge from their experiences and that the educator is only the knowledge provider. Yang has applied Constructivist theory to develop a patient-education model and design a knowledge acquisition framework which could assist health professionals in organizing their domain knowledge prior to the writing of the actual textual content. A key contribution of a Constructivist model to the HealthDoc methodology would be in guiding the author to construct the underlying discourse structure of the master document.

With the original HealthDoc authoring tools, the emphasis was on providing the author with a means of entering textual variations, specifying the conditions under which each variation should be selected, and annotating the master document with information needed for later automated repairs. However, it was assumed that the author would use his knowledge of the application domain to organize the pieces of text into a coherent master-document structure. Knowledge about the discourse structure was left implicit, to be managed mentally (and differently) by each individual author. As an example, an author might enter the following text and variations on the topic of the two types of diabetes:

(1) There are two main types of diabetes. One type is insulin-dependent, also known as type I diabetes, and the other is non–insulin-dependent, also called type II diabetes.

(2) The condition that you have is insulin-dependent diabetes. *(variation 1)* The condition that you have is non–insulin-dependent diabetes. *(variation 2)*

(3) Insulin-dependent and non–insulin-dependent diabetes are different disorders, so that the causes, short-term effects, and treatments for the two types differ. However, both types can cause the same long-term health problems.

(4) With insulin-dependent diabetes, your body makes little or no insulin. *(variation 1)* With non-insulin-dependent diabetes, your body makes insulin, but can't use it well. *(variation 2)*

The underlying discourse structure of this passage of text can be characterized as follows:

- Define the two types of diabetes.
- Identify the patient's type of diabetes.
- Compare the types.
- Contrast the types.

However, the elements of this discourse structure would not have been made apparent during the authoring process. Also, the author would not have been able to indicate that a similar pattern of statements (*define*, *identify*, *compare*, *contrast*) could be applied in constructing other topics of text.

In contrast, Yang's knowledge level of modelling could guide the creation of the master document according to pre-defined discourse structures that model the interaction between physician and patient. Her Constructivist model[4] tells us that addressing patient concerns (about pain, risks, complications, etc.) should be the basis for the information provided by the physician. An authoring tool incorporating this type of knowledge would therefore have an explicit 'addressConcerns' rhetorical model that would be used in constructing a topic passage. For example, the topic passage for each concern might have the following elements:

- Identify the concern.
- Describe the concern.
- Address how patient should handle the concern.

The (generic) text for the concern of pain might therefore be entered as follows:

(5) You may feel severe pain. *(Identify concern.)*

(6) The pain or discomfort will be felt in the breast area or abdominal site. Soreness and swelling are often part of your body's reaction to the trauma of surgery. *(Describe concern.)*

(7) You should not perform lifting activities or anything that involves the muscles in the breast area or abdominal site. This will cause additional pain and prevent the healing of your wound. *(Address handling of concern.)*

[4] Other learning models might also be used at the knowledge level of modelling the master document.

A Natural Language Generation tailoring engine. The current HealthDoc tailoring engine will be the software kernel of our proposed Natural Language Generation tailoring system. We have now replaced our original 'homegrown' document design language [8] with a standard document description language (XML). Master documents containing personalized health information for various domains (e.g., skin care and smoking cessation) have been prepared and marked up with XML tags and attributes. These tagged conditional documents have then been processed through an XSL transformation that produces a presentation-ready and print-ready, highly customized document using the PHP Hypertext Preprocessor. This software can now enable visualizations of tailored versions of any content in our master-document format as either a Web presentation or in paper form.

Our earlier work in the HealthDoc Project demonstrated that complex, stylistically polished texts can be crafted from pre-existing texts represented in an appropriate 'master-document' format. We are continuing the development of the 'generation-by-selection-and-repair' paradigm, with particular emphasis on the architectural issues involved in text-to-text generation systems. Our long-term goal is to continue to develop our theory of automated text repair, and test it by implementing repair algorithms that recognize and revise various infelicities in ill-formed texts. One approach to the automated detection of 'repair patterns' that we plan to investigate is the combination of pattern-based methods from classical rhetorical theory (e.g., [10]) with n-gram language models.

4 Conclusions

Our goal in this research is to develop natural language software tools, specifically an authoring tool and Natural Language Generation tailoring system, to automatically generate tailored patient education for patients choosing among the plethora of options involved in reconstructive breast surgery. The benefits of enhanced preoperative education have been established in the literature, and serve as the basis for many of the predicted benefits listed below:

– A single, comprehensive source of educational materials.
– Less conflicting information than might be associated with multiple educational brochures in multistep surgical procedures, assuming these materials even exist.
– A better-informed patient: Decreased perioperative anxiety; fewer and less serious complications; faster recovery and rehabilitation; enhanced recognition of postoperative complications, because of the ability to include more specific information tailored to each of the surgical subcomponents.
– Better patient outcomes: fewer and less serious complications, etc.
– Less time required in perioperative discussions ensuring that information is communicated.

Future applications of the results of this research would be the extensions of content to other procedures and surgical subspecialties. The intended robustness of theory and technology will also allow extension beyond that of surgical intervention, potentially to any medical treatment involving multiple modalities requiring cohesion of educational content (e.g., medical and radiation oncology)

Acknowledgements

We would like to acknowledge the financial support of Bell University Laboratories at the University of Waterloo.

References

1. S. Banks, "English to sentence plan translation using subsumption on stylistic abstraction", M.Math. thesis, University of Waterloo, 1999.
2. S. Banks and C. DiMarco, "Surface form to sentence plans: A method for English-to-English translation", Proceedings of the 4th Symposium on Natural Language Processing, 2000.
3. D.S. Bental, A.J. Cawsey, and R. Jones, "Patient information systems that tailor to the individual", *Patient Education and Counselling*, 36, 171–180, 1999.
4. K. Binsted, A.J. Cawsey, and R.B. Jones, "Generating personalised patient information using the medical record", Proceedings of Artificial Intelligence in Europe, 1995.
5. C. Brun, M. Dymetman, and V. Lux, "Document structure and multilingual authoring", Proceedings of the First International Natural Language Generation Conference, 2000.
6. M. Campbell, B. DeVellis, V.J. Strecher, A. Ammerman, R. DeVellis, and R. Sandler, "Improving dietary behavior: The effectiveness of tailored messages in primary care settings", American Journal of Public Health, 84(5), May 1994, 783–787.
7. A. Cawsey, "Natural language generation in health care", *Journal of the American Medical Informatics Association*, 4, 473–482, 1997.
8. C. DiMarco and M.E. Foster, "A method and apparatus for authoring of customizable multimedia documents", International patent application PCT CA98/00771 filed 11 August 1998 (priority dating from U.K. application number 9716986.6 filed 12 March 1997).
9. C. DiMarco, G. Hirst, and E. Hovy, "Generation by selection and repair as a method for adapting text for the individual reader", Workshop on Flexible Hypertext, ACM Hypertext Conference, 1997.
10. J. Fahnestock, *Rhetorical figures in science.* Oxford University Press, 1999.
11. C. Fosnot, *Constructivism: Theory, perspectives, and practice.* Teachers College Press, 1996.
12. C. DiMarco, G. Hirst, L. Wanner, and J. Wilkinson, "HealthDoc: Customizing patient information and health education by medical condition and personal characteristics", Workshop on Artificial Intelligence in Patient Education, University of Glasgow, 1995.
13. A. Hartley and C. Paris, "Multilingual document production: From support for translating to support for authoring", *Machine Translation, Special Issue on New Tools for Human Translators*, 1997.
14. G. Hirst, C. DiMarco, E. Hovy, and K. Parsons, "Authoring and generating health-education documents that are tailored to the needs of the individual patient", Proceedings of the Sixth International Conference on User Modeling, Sardinia, Italy", 1997.
15. B. Jakeway, "SPLAT: A sentence plan authoring tool for natural language generation", M.Math. thesis, University of Waterloo, 1995.
16. P. Meredith and C. Wood, "Inquiry into the potential value of an information pamphlet on consent to surgery to improve surgeon-patient communication", *Quality in Health Care*, 7(2), 65–69, June 1998.
17. M. O'Donnell, A. Knott, J. Oberlander, and C. Mellish, "Optimising text quality in generation from relational databases", Procs. First International Natural Language Generation Conference, 2000.
18. K. Parsons, "An authoring tool for customizable documents", M.Math., University of Waterloo, 1997.

19. R. Power and D. Scott, “Multilingual authoring using feedback texts”, Proceedings of 36th Annual Meeting of the Association for Computational Linguistics, 1998.
20. E. Reiter and L. Osman, “Tailored patient information: Some issues and questions”, Proceedings of the Association for Computational Linguistics, 1997.
21. E. Reiter, R. Robertson, and L.M. Osman, “Lessons from a failure: Generating tailored smoking cessations letters”, *Artificial Intelligence*, 144, 41–58, 2003.
22. D. Scott, R. Power, R. Evans, “Generation as a solution to its own problem”, Proceedings of the Ninth International Natural Language Generation Workshop, 1998.
23. C.S. Skinner, V.J. Strecher, and H. Hospers, “Physicians’ recommendations for mammography: Do tailored messages make a difference?”, *American Journal of Public Health*, 84(1), 1994.
24. V.J. Strecher, M. Kreuter, D-J. Den Boer, S. Kobrin, H.J. Hospers, and C.S. Skinner, “The effects of computer-tailored smoking cessation messages in family practice settings”, *The Journal of Family Practice*, 39(3), 1994.
25. R.J. Street Jr., B. Voigt, C. Geyer Jr., T. Manning, G.P. Swanson, “Increasing patient involvement in choosing treatment for early breast cancer”, *Cancer*, 76(11), 2275–2285, December 1, 1995.
26. L. Wanner and E. Hovy, “The HealthDoc sentence planner”, Proceedings of the Eighth International Natural Language Generation Workshop, 1996.
27. C. Yang, “Design of a knowledge acquisition tool using a constructivist approach for creating tailorable patient education materials”, M.Math. thesis, University of Waterloo, 2005 expected.