Authoring and Generation of Individualized Patient Education Materials

C. Di Marco, Ph.D., P. Bray, M.D., F.R.C.S.(C.), H.D. Covvey, M.Sc., D.D. Cowan, Ph.D., V. Di Ciccio, M.Eng., E. Hovy, Ph.D., Joan Lipa, M.D., F.R.C.S. (C.), and C. Yang, M.Math. Waterloo Institute for Health Informatics Research, Waterloo, Ontario, Canada, the Division of Plastic Surgery, University Health Network, University of Toronto, Toronto, Ontario, Canada, and the Information Sciences Institute, University of Southern California, U.S.A.

ABSTRACT

Although the pre-surgical patient-surgeon encounter is the opportunity to educate the patient, it is essential that the patient be given educational materials to complement the 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 100s of variants, or all patients would be given the same set of brochures without regard for differing needs. We have been developing an information brochure generator that customizes material for every individual patient regardless of the complexity of the surgical intervention.

THE IMPORTANCE OF TAILORING IN PATIENT EDUCATION

Present-day health-education is often limited in its effectiveness by the need to address it to a wide audience. Either a minimal, generic document containing only the information common to everyone is produced; or a maximal document that tries to provide all the information that might be relevant to someone (and irrelevant to many). But material that contains irrelevant information, omits the relevant, or does not seem to be addressed to the reader, is likely to be discounted or ignored, with consequent problems in motivation for compliance with medical regimens.

However, recent experiments suggest that health-education material can be much more effective if it is customized to the individual reader in accordance with 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 behavior[6], or having a mammogram[23]. In each study, the 'tailored' leaflets were found to have a significantly greater effect on the patients' behavior than 'generic' leaflets had upon patients in a control group.

This kind of customization involves much more than just producing each brochure in half dozen versions for different audiences. Rather, the number of combinations of factors can easily be in hundreds of thousands (as in the studies cited), making it impossible to produce the large number of different editions needed for individual tailoring.

Recently, researchers in Natural Language Generation (NLG) 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 Health-Doc Project (1994–1999)[12] developed a method for generating tailored documents based on a new paradigm for NLG— 'generation-by selection-and-repair'[9]—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 'repaired' for form, style, and coherence.

A realistic and usable implementation of the Health-Doc approach requires a sophisticated authoring tool to assist the writer, and a sentence planner (cf.[26]) that would undertake to repair and polish the text. The average technical writer cannot be expected to pre-compile all possible combinations. 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 repair.

POTENTIAL SOLUTION: NLG

The creation of the input material for NLG systems is a problem for all generation systems, including ours. The concept of 'preparing' a database, knowledge base, or other resource for NLG 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) usable to ensure coherent, high-quality text. This idea led us to

adopt the authoring of a 'database' of reusable text, the master document), as the basis for the generation-by-selection-and-repair paradigm.

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

Our approach to authoring for NLG systems falls between that 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[14, 18], but we are also dealing with authoring of input at a deeper level of linguistic representation[1, 2, 15], as is typical of NLG 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', trained in structured knowledge acquisition. Our goal to design a system, based on our paradigm of NLG 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.

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 aes-

thetics, 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.

Although preoperative 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 brochures are only available for the most common surgical procedures and must, by necessity, remain generic in nature to ensure applicability to all patients.

Creation of a tailored information document for every individual patient would potentially increase the 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 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 at all operative stages if 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-centered methods[24] and using quality information brochures to improve surgeon-patient communication[16].

We are developing a system for generating preopera-

tive patient education materials that tailors the text to every individual patient regardless of the complexity of the surgical intervention. The components of this system will consist of an NLG tailoring system, content authoring environment, and creation of a database of educational modules pertaining to each subcomponent of a given surgical intervention.

COMPONENTS OF TAILORING SYSTEM FOR RECONSTRUCTIVE SURGERY

Corpus of Textual Variants. We are creating a corpus of textual variants to 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. Textual variants will initially be entered manually into our master-document format and subsequently authored by a patient-education writer using the prototype authoring tool.

Authoring Tool to Guide Healthcare Providers. In the earlier HealthDoc Project, we developed several authoring tools[1, 15, 18] 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 that enabled syntactic and semantic repair of reassembled text. Our new approach is to guide surgeons to directly enter 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' modeling 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 various stakeholders (e.g., surgeons, patients, hospital) when tailoring the educational material. For example, the surgeon may be primarily concerned with ensuring patient compliance with the recommended treatment which will lead to favorable outcomes; the patient may be most concerned with risks and complications associated with different treatment options. The authoring tool therefore must 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 a knowledge provider. Yang has applied Constructivism to develop a patient-education model and design a knowledge acquisition framework to assist health professionals in organizing their knowledge prior to the writing of actual textual content. A key contribution of a Constructivist model to the HealthDoc methodology is 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 to support 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 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 text passage text can be characterized as follows: (1) Define the two types of diabetes. (2) Identify the patient's type of diabetes. (3) Compare the types, and then (4) 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 used in constructing other text.

In contrast, Yang's knowledge level of modeling is intended to guide master document creation according to pre-defined discourse structures that model the interaction between physician and patient. Her Constructivist model tells us that addressing patient concerns (pain, 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: (1) Identify the concern. (2) Describe the concern. (3) Address how patient should handle the concern.

The (generic) text for the concern of pain might therefore be entered as follows: (1) You may feel severe pain. (Identify concern.) (2) 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.) (3) 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.)

An NLG Tailoring Engine. The current HealthDoc tailoring engine is the software kernel for our NLG 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 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- and print-ready, highly customized document using the PHP Hypertext Preprocessor. This software enables visualizations of tailored versions of any content in our master-document format as a Web presentation or in paper form.

Our earlier 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.

CONCLUSIONS

Our goal in this research is to develop natural language software tools, specifically an authoring tool and NLG tailoring system, to automatically generate tailored patient education for patients choosing among the options involved in reconstructive breast surgery. The benefits of enhanced preoperative education have been established, and serve as the basis for many of the predicted benefits: (1) A single, comprehensive source of educational materials. (2) Less conflicting information than might be associated with multiple brochures in multi-step surgical procedures, assuming these materials even exist. (3) 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. (4) Better patient outcomes: fewer and less serious complications, etc. (5) 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

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 Ed. 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 Intern'l 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 NLG 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.