

Analysing Clinical Guidelines' Contents with Deontic and Rhetorical Structures

Gersende Georg¹, Hugo Hernault², Marc Cavazza³,
Helmut Prendinger⁴, and Mitsuru Ishizuka²

¹ Haute Autorité de Santé, F-93218, Saint-Denis La Plaine, France

² Graduate School of Information Science and Technology, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

³ School of Computing, University of Teesside TS1 3BA Middlesbrough, United Kingdom

⁴ National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan
g.georg@has-sante.fr, hugo@nii.ac.jp, m.o.cavazza@tees.ac.uk,
helmut@nii.ac.jp, ishizuka@i.u-tokyo.ac.jp

Abstract. The computerisation of clinical guidelines can greatly benefit from the automatic analysis of their content using Natural Language Processing techniques. Because of the central role played by specific deontic structures, known as recommendations, it is possible to tune the processing step towards the recognition of such expressions, which can be used to structure key sections of the document. In this paper, we extend previous work on the automatic identification of guidelines' recommendations, by showing how Rhetorical Structure Theory can be used to characterise the actual contents of elementary recommendations. The emphasis on causality and time in RST proves a powerful complement to the recognition of deontic structures and supports the identification of relevant knowledge, in particular for the identification of conditional structures, which play an important role for the subsequent analysis of recommendations.

Keywords: natural language processing, rhetorical structure theory, guidelines.

1 Introduction

The computerisation of clinical guidelines had led to a renewed interest in their automatic processing. Whilst encoding of clinical guidelines can take place manually [1] or with the assistance of visual interfaces [2], this process can be facilitated by introducing Natural Language Processing (NLP) techniques in the process. We have shown in previous work [3] how the automatic identification of specific linguistic markers of clinical recommendations, known as deontic operators, could be used to provide a first level of structuring. The main outcome of this approach has actually been a software to analyse the structure of clinical guidelines during their development process. In this paper, we describe an extension of our previous approach in which the actual contents of recommendations can be further structured using a medically-relevant subset of rhetorical relations.

2 Clinical Guidelines and Their Analysis

Clinical guidelines are sophisticated documents which are often syntactically complex. One research direction consists in standardising guidelines' writing or even recurring to controlled languages [4]. Whilst their automatic processing is beyond the state-of-the-art of NLP techniques, we have shown recently that much benefit could be gained from the recognition of key expressions which would structure portions of text according to the document's logic. This is a case where local or shallow processing can be used to structure free text segments. G-DEE (for *Guidelines Document Engineering Environment*), is a text analysis environment dedicated to the study of clinical guidelines [5]. It supports multiple document processing functions including the automatic recognition of recommendations using shallow NLP techniques (such as Finite-State Automata, FSA) to recognise deontic expressions corresponding to recommendations. Since 2007, G-DEE has been integrated into the process of clinical guidelines authoring at the French National Authority for Health (Haute Autorité de Santé), which is in charge of the elaboration of all official guidelines in France. G-DEE has been used from the first draft of recommendations to the final version of the document, to provide an independent analysis of guidelines structure. Since its introduction two years ago, the number of requests for authoring by HAS project leaders with G-DEE has steadily increased (44 cumulative number of requests considering that one guideline can be analyzed from 1 to 3 by G-DEE) indicating a growing interest amongst users (its use has not been made compulsory).

3 Complementarity between Deontic and Rhetoric Structures

In addition to its use to support the guidelines' authoring process, we've shown in previous work [6] that structuring guidelines with deontic operators can help identifying important clinical actions that can be matched to underlying protocols. However, extending the automatic processing of guidelines to the actual contents of individual recommendations, i.e. processing the free text content of deontic operators' scopes, remains a challenging task from an NLP perspective. Our FSA-based recognition of deontic operators already required significant linguistic resources, despite being focused on specific linguistic descriptions. It thus seems difficult to adapt similar principles for the analysis of scopes which exhibit much greater syntactic variability and semantic coverage. Ideally, we would seek a method which reconciles broad linguistic coverage, shallow NLP, and the ability to further structure the contents of recommendations' scopes. All this points towards discourse-processing methods, and led us to explore the potential use of Rhetorical Structure Theory (RST) [7]. Although some authors have suggested that legal texts were not amenable to RST formalization [8] and we have shown the proximity between legal texts and clinical guidelines in their use of deontic structures [5], we were comforted in our hypothesis by the many previous references applying RST to medical NLP and medical language generation [9-10].

4 RST Parsing of Recommendations' Scopes

In order to uncover the rhetorical structure of clinical guidelines, we used a RST discourse parser based on Support Vector Machines (SVM). The parser uses a rich set of shallow lexical, syntactic and structural features from the text, and processes its input in two steps. Firstly, a discourse segmenter cuts the text into “elementary discourse units” (EDUs), the atomic units of discourse which are the terminal nodes of the rhetorical structure tree. Each word and its context are represented by a feature vector. In a second step, the calculated EDUs are passed to the tree-building component, which creates the full RST tree. In order to improve the computational properties of the classification problem and ensure a good separability between the label classes, we used the reduced set of 18 relations defined in [11] and used by [12] amongst the original set of 78 rhetorical relations. While we have observed a natural and quite efficient complementarity between deontic recognition and the RST analysis of recommendations' scopes, it would still be appropriate to investigate whether RST parsing could be used as a sole principle for guidelines' structuring and recommendations analysis. On the theoretical side, examples described by Gallardo [13] suggest that RST functions fail to capture key elements of recommendations. Direct RST analysis of recommendations mostly produces structures based on conditions and elaborations. When conditions are explicitly represented as part of the recommendation, RST analysis correctly identifies part of the recommendation, although it fails to provide a complete segmentation along the lines of those produced by G-DEE, with proper identification of scopes.

5 The Recommendations' Processing Pipeline

Since our deontic parser has been validated through user experiments and through real-world deployment within a guidelines production agency (HAS), we naturally thought of extending G-DEE by a further step of RST analysis, targeting the recommendations' scopes. We adopted a processing model based on the fusion of outputs from G-DEE and the RST parser. We have developed a module operating in two steps: (i) localization of a deontic operator within the RST structure; (ii) fusion of the RST structure with the front-scope and the back-scope (resulting in these scopes being structured by RST functions, see Figure 1). A pre-processing step consists of analysing the guideline using G-DEE to determine sentences that correspond to recommendations. An RST analysis is then performed on the file containing the recommendations identified by G-DEE. The G-DEE processor scans the sentence and extracts the deontic operator using the specific mark-up *<DeontOp>*. The next step consists of localizing the same deontic operator in the XML RST file, using the G-DEE processor that proceeds through a standard finite-state processing algorithm. The successful match leads us to determine whether the deontic operator is a part of the nucleus (N) or the satellite (S) by the recognition of the RST function as shown above. The G-DEE processor then scans the sentence from the RST file, and extracts the function corresponding to the front-scope (either the nucleus or the satellite previously recorded information). It then uses a dedicated FSA to mark-up the corresponding front-scope with an appropriate tag (*<Manner-Means>* in the example above). In a similar way, the function corresponding to the back-scope is recorded and the G-DEE processor tags the back-scope accordingly (*<Temporal>* in the example).




Fig. 1. Refining Recommendations' Structure by Merging Deontic and Rhetorical Mark-ups

6 Results and Discussion

We have extracted recommendations from the 2005 Hypertension Guidelines (in English, "Management of adults with essential hypertension"), obtaining a test set of 79 recommendations. Overall, RST processing with basic functions had a very significant contribution for approximately 25% of recommendations. This means that not only it did refine the recommendations' structure, but the new relations were directly meaningful. The most useful RST functions detected on these Guidelines are: *Condition* (10 occurrences), *Manner-Means* (5), *Temporal* (4), and *Enablement* (3). RST parsing also contributed to an improved structure with generic functions, through the *Elaboration* function, for 14% of recommendations: this includes isolating the grade of the recommendation or some specific target from within (generally back-) scopes. A better joint recognition of functions could achieve substantial improvements of the *Conditional* relations.

7 Conclusions

The analysis of recommendations' scopes using RST can successfully extend our previous approach, improving automatic structuring for 44% of recommendations, which increases significantly the quality of the automatic processing, even more so considering that documents tend to be analysed several times during their authoring cycle. Further, it remains compatible with our philosophy of document processing, which is to structure text segments using discourse markers, specific (e.g. deontic), or not. This type of automatic analysis tends to be well-accepted by guidelines' developers as it is designed as a human-in-the-loop approach. This is also an interesting test case for medical NLP, where the recognition of discourse structures, rather than of named entities or actions

(for instance through Information Extraction or terminological processing) can support the identification of clinically relevant information over an entire text.

Acknowledgments. Dave du Verle developed the first version of the RST parser used in these experiments. Gersende Georg and Marc Cavazza have been funded by NII for a summer visit in 2008 during which this collaboration was established.

References

1. Shiffman, R., Karras, B., Agrawal, A., Chen, R., Marenco, L., Nath, S.: GEM: A proposal for a more comprehensive guideline document model using XML. *J. Am. Med. Informatics Assoc.* 7, 488–498 (2000)
2. Shahar, Y., Young, O., Shalom, E., Mayaffit, A., Moskovitch, R., Hessian, A., Galperin, M.: The Digital electronic Guideline Library (DeGeL): a hybrid framework for representation and use of clinical guidelines. *Stud. Health Technol. Inform.* 101, 147–151 (2004)
3. Georg, G., Jaulent, M.-C.: A Document Engineering Environment for Clinical Guidelines. In: King, P.R., Simske, S.J. (eds.) *Proceedings of the 2007 ACM Symposium on Document Engineering*, pp. 69–78. ACM Press, New York (2007)
4. Fuchs, N.E., Kaljurand, K., Schneider, G.: Attempto Controlled English Meets the Challenges of Knowledge Representation, Reasoning, Interoperability and User Interfaces. In: *FLAIRS Conference*, pp. 664–669 (2006)
5. Georg, G., Jaulent, M.-C.: An Environment for Document Engineering of Clinical Guidelines. In: *Proceedings AMIA Symposium*, pp. 276–280 (2005)
6. Georg, G., Cavazza, M.: Integrating Document-based and Knowledge-based Models for Clinical Guideline Analysis. In: Bellazzi, R., Abu-Hanna, A., Hunter, J. (eds.) *AIME 2007. LNCS (LNAI)*, vol. 4594, pp. 421–430. Springer, Heidelberg (2007)
7. Mann, W.C., Thompson, S.A.: Rhetorical Structure Theory: Toward a functional theory of text organisation. *Text* 8(3), 243–281 (1988)
8. Taboada, M., Mann, W.C.: Applications of Rhetorical Structure Theory. *Discourse Studies* 8(4), 567–588 (2006)
9. Piwek, P., Hernault, H., Prendinger, H., Ishizuka, M.: T2D: Generating Dialogues Between Virtual Agents Automatically from Text. *Intelligent Virtual Agents*, 161–174 (2007)
10. Grasso, F.: Rhetorical coding of health promotion dialogues. In: Dojat, M., Keravnou, E.T., Barahona, P. (eds.) *AIME 2003. LNCS*, vol. 2780, pp. 179–188. Springer, Heidelberg (2003)
11. Carlson, L., Marcu, D., Okurowski, M.E.: Building a discourse-tagged corpus in the framework of Rhetorical Structure Theory. In: *Proc. of the Second SIGdial Workshop on Discourse and Dialogue. Annual Meeting of the ACL*, vol. 16, pp. 1–10 (2001)
12. Soricut, R., Marcu, D.: Sentence level discourse parsing using syntactic and lexical information. In: *Proc. of the 2003 Conference of the North American Chapter of the Association For Computational Linguistics on Human Language Technology*, vol. 1, pp. 149–156 (2003)
13. Gallardo, S.: Pragmatic support of medical recommendations in popularized texts. *Journal of Pragmatics* 37(6), 813–835 (2005)