Integrating Document-Based and Knowledge-Based Models for Clinical Guidelines Analysis

Gersende Georg[1,2,3] and Marc Cavazza[4]

[1] INSERM, U 872, Eq. 20, Paris, F-75006 France
[2] Université Paris Descartes, UMR S 872, Paris, F-75006 France
[3] Centre de Recherche des Cordeliers, Université Pierre et Marie Curie - Paris 6, UMR S 872, Paris, F-75006 France
Gersende.Georg@spim.jussieu.fr
[4] School of Computing, University of Teesside
TS1 3BA Middlesbrough, United Kingdom
m.o.cavazza@tees.ac.uk

Abstract. Research in the computerization of Clinical Guidelines (CG) has often opposed document-based approaches to knowledge-based ones. In this paper, we suggest that both approaches can be used simultaneously to assess the contents of textual Clinical Guidelines. In this first experiment, we investigate the mapping between a document model, which has been marked-up to structure its recommendations, and a knowledge structure representing the management of specific disease. This knowledge representation is based on planning formalisms, more specifically Hierarchical Task Networks (HTN). Our system operates by first automatically encoding the textual guideline through the identification of specific expressions with surface natural language processing, as described in previous work. In a subsequent step, the HTN, constructed manually and independently, and represented as an explicit AND/OR graph, is searched for a solution sub-graph using an algorithm derived from AO*. Whilst the HTN is being traversed, corresponding information is accessed in the encoded textual CG, to guide the solution extraction process. We illustrate this through a case study developed around French guidelines for the management of hypertension. Recommendations included in the textual guideline provide complementary information for the instantiation of an HTN on specific patient data. The mapping takes place at different levels, from the pre-condition of operators to the rules playing a role as selection heuristics when extracting a solution sub-graph. Such a process, which explores the textual document from the prospective of a task model, can help analyzing the overall structure of clinical guidelines and ultimately improving its applicability.

Keywords: Clinical Guidelines, GEM, Planning, HTN, deontic operators.

1 Introduction and Rationale

The computerization of clinical guidelines has followed two main approaches, known as document-based and knowledge-based [1], which have so far largely remained separate. Document-based approaches are more closely connected to the actual

R. Bellazzi, A. Abu-Hanna, and J. Hunter (Eds.): AIME 2007, LNAI 4594, pp. 421–430, 2007.

guideline production process. On the other hand, knowledge-based models can directly be integrated into Decision Support Systems (DSS), whilst document-based models require an additional level of interpretation to extract rules or decision trees, whose automation has recently attracted interest from several researchers [2] [3] [4]. The direct extraction of complete knowledge structures from free text remains a long-term research objective still beyond the state-of-the-art (in particular in terms of Natural Language Processing techniques). As a first step, Hagerty [5] has proposed the use of Information Extraction techniques for the automatic identification of conditional expressions within recommendations (as part of his “Hypertext Guideline Markup Language (HGML)” markup annotation).

However, there could be benefits in the joint use of document-based and knowledge-based approaches for studying the structure of clinical guidelines and assessing their consistency and completeness. This is what we explore in this paper, as we describe a software environment for the analysis of textual guidelines based on the joint use of guideline document encoding and a knowledge-based formalization of the underlying clinical protocol.

2 Modeling Guidelines with Planning Formalisms

Most knowledge-based approaches, such as PRO*forma* [6] and Prodigy [7] are based on knowledge structures centered on clinical actions. These structures are related to action representations encountered as part of planning formalisms such as STRIPS or PDDL [8]. This has been systematized only recently by [8] who have analyzed the role of planning formalisms and discussed the applicability of planning approaches (i.e. not limited to the representation of elementary actions) to the modeling of clinical guidelines. In their review, Bradbrook et al. [8] mention the possible use of Hierarchical Task Network (HTN) planning to represent guidelines protocols, although their own approach is based on other planning formalisms. There are however many benefits in using an HTN approach to model the overall guideline behavior. HTN are one of the most successful planning formalisms and are used in a variety of implemented systems, from robotics [9], game playing (bridge) [10], and virtual characters animation [11]. HTN are considered appropriate to knowledge-intensive Planning problems and their top-down descriptions are well suited to the description of clinical protocols. Other knowledge-based approaches bear similarity to Planning. GUIDE, based on Petri Nets, represents the workflow of clinical guidelines [12] generally composed of a nested sequence of actions. Asbru [13] is a time-oriented, intention-based, skeletal-plan specification language that is used to represent clinical protocols. Its plans attributes are characterized by intentions, conditions, and effects which can be structured via a temporal representation. All these approaches make use of planning concepts without however implementing any of the traditional planning techniques: this is largely because their main focus is on making knowledge procedural to facilitate its integration in DSS.

Although Planning has not been one of the main AI technologies used in medical knowledge-based systems, there has been interest in Planning formalisms for the formalization of clinical protocols. Haddawy et al. [14] have shown as early as 1995 that Planning formalisms represented a benefit over traditional decision trees.

Spyropoulos [15] has reported the use of planning and scheduling techniques to model both therapy planning and hospital management procedures.

2.1 HTN Formalization of Clinical Guidelines

HTN are appropriate to represent multi-step decomposable processes, and this applies naturally to clinical protocols, as long as they can be decomposed into independent sub-tasks. This is why the various steps of clinical care can be represented as an AND/OR graph formalizing an explicit HTN (i.e. one in which the main task has been decomposed *a priori* and entirely, down to the level of grounded actions, rather than being dynamically refined using decomposition methods [16]).

Yet, there are other important representational elements to be used in conjunction with AND/OR graphs for the instantiation of a solution plan on a specific set of patient data. Each sub-task should be associated pre-conditions as well as post-conditions. Another element of representation is constituted by the heuristics guiding node selection at the level of OR nodes, and costs (in the algorithmic sense) associated to the actual clinical actions. The main difference with traditional HTN planners, such as SHOP [17], will consist in using an explicit and finite HTN. In this way, the overall guideline can also be represented visually, rather than as a collection of refinement methods. The explicit nature of AND/OR graphs thus allows a direct visualization, which in turn facilitates knowledge elicitation. The guideline contents are represented as an AND/OR graph with the highest-level task (the overall goal of the clinical protocol) as the top node. An instantiation of the guideline recommendations can be obtained by extracting a solution from this explicit HTN. Provided certain limitations are properly taken into account, such as task decomposability, absence of long-distance dependencies and stability of data over time (or at least within each sub-task covered by the guideline), a solution sub-graph can be extracted from the HTN with a simple variant of the AO* algorithm [18], which provides solutions to decomposable problems. Primary heuristics will be used in the selection of one out of several alternative sub-tasks subsumed by an OR node. Examples of such sub-tasks are constituted by alternative therapeutics (see below). The solution graph will then constitute an instance of the guideline, to be applied to the specific data for the patient being considered.

2.2 System Architecture and Overview

The approach described in this paper is based on an experimental software platform integrating a document engineering environment [19] (Fig. 1 – A) and an HTN-based module (Fig. 1 – B). In essence, this software aims at synchronizing the traversal of the explicit HTN and the consultation of a guideline document, previously processed to mark-up its elementary recommendations. This synchronization is based on the automatic extraction of information from the marked-up guideline corresponding to the HTN sub-task under consideration. In other words, once the protocol has been modeled as an HTN, the system-driven exploration of its AND/OR graph drives the interactive consultation of the guideline document model to assist in the instantiation of the solution sub-graph (Fig. 1 – C).

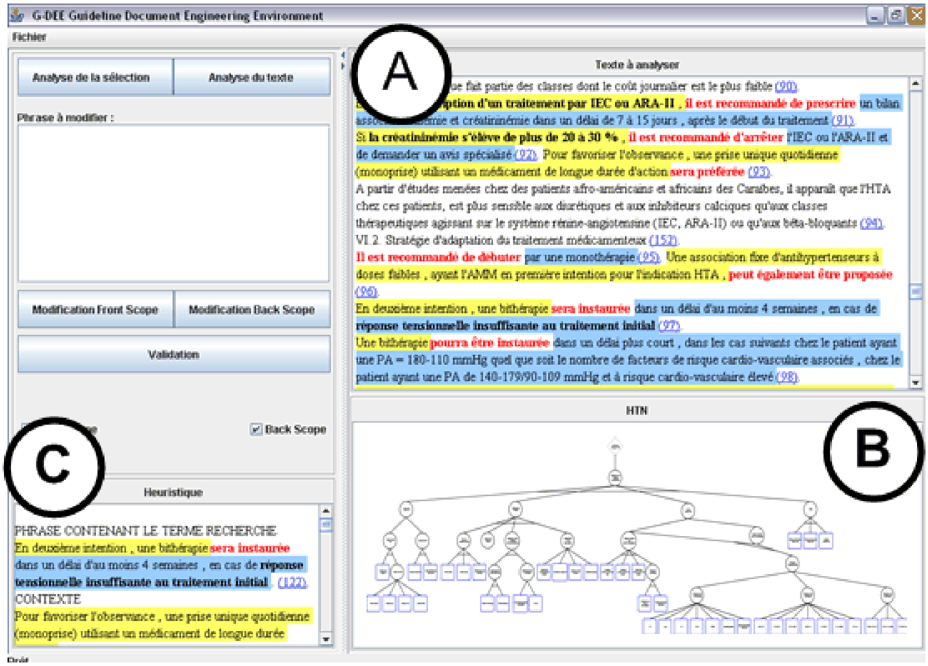

Fig. 1. The *G-DEE* interface *(see text)*

2.3 Document Processing Applied to Clinical Guidelines

The central element of these experiments is a document engineering platform dedicated to the study of clinical guidelines, developed by one of the authors in previous work [19]. The *G-DEE* system (for Guidelines Document Engineering Environment) automatically performs XML encoding of guidelines based on the recognition of the guideline's linguistic content. The system uses a set of approximately 1200 Finite State Automata (FSA), which correspond to 70 syntactic patterns with their morphological variations, to recognize specific natural language expressions corresponding to the linguistic formulation of elementary recommendations (such as "recommended, advised, one should / ought to"… these expressions are known as *deontic operators* [19] [20]). In a subsequent step, the textual occurrence of these recommendations is structured by marking-up the deontic operator and the textual expression to which it applies. These expressions are named *front-scope* and *back-scope* [20], and correspond to the operands of the deontic operator, from which conditions and actions can be extracted. Figure 2 illustrates the *scopes* of the deontic operator "should be considered".

The encoding obtained (Fig. 3) can serve as a basis for further processing, for instance the extraction of decision rules under textual format, or the encoding using GEM [21] categories, for instance through the identification of decision variables [19].

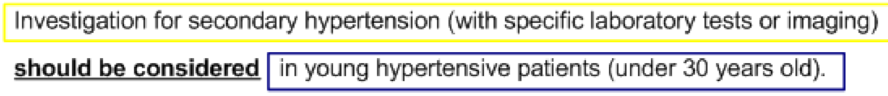

Fig. 2. *Front-* and *back-scope* for the deontic operator "should be considered" (translated from the original French Guidelines for the management of hypertension)

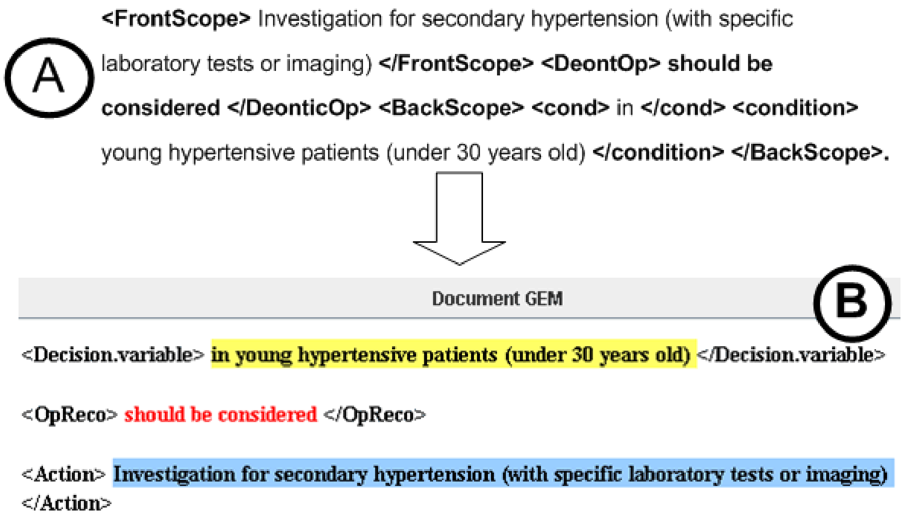

Fig. 3. The marking-up of a recommendation (part A) and its automatic structuration in the GEM format using dedicated XSL style sheets (part B)

From a content perspective, the recommendations so identified correspond to decision steps and, when they relate to possible alternatives in the therapeutic plan, to actual heuristics that can be used to select the most appropriate alternative in the extraction of a solution task graph. This is the type of content which is central to the mapping of HTN traversal to the guideline document. One key problem of knowledge representation is actually to properly interpret textual recommendations in terms of selection heuristics or grounded action costs.

2.4 Synchronization of HTN Traversal and Document Exploration

The most important aspect of these experiments is the synchronization of HTN traversal with text consultation, which will support the interactive features of the environment. Considering that the HTN has been developed independently of the textual guideline[1] (Fig. 4), this synchronization should relate the contents of HTN nodes traversed to the contents of the textual guidelines.

The first step consists in an "offline heuristic calculation mode" that determines heuristic values of nodes by rolling back estimated costs of grounded actions (the actions associated to the bottom nodes of the HTN such as drug prescription). This

[1] "Management of adults with essential hypertension – 2005 update" (http://www.has-sante.fr/).

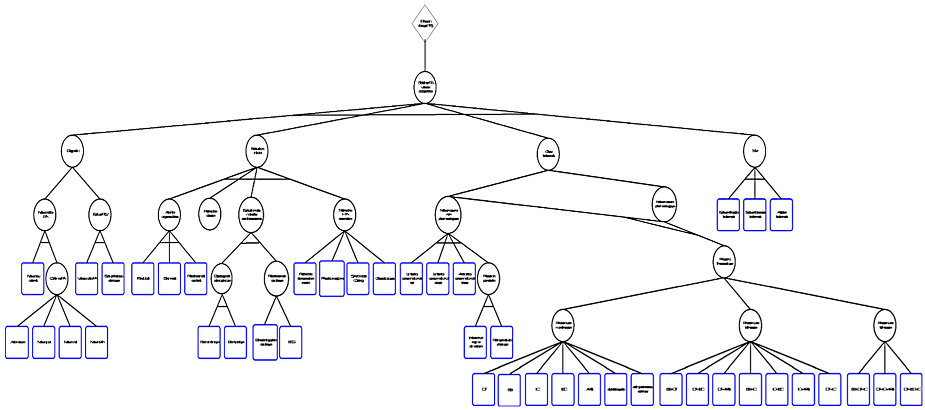

Fig. 4. An overview of the HTN that represents the management of hypertension (pre- and post-conditions attached to the HTN nodes are not represented on the figure)

determines the heuristic value of a node when it can be evaluated from its sub-tasks, e.g. the evaluation of cardiovascular risk which can have an impact on the choice of therapy. The interactive nature of the HTN exploration actually leads us to dissociate the two aspects of heuristic calculation in algorithms of the AO* family [16], which comprise a primary heuristic determining the selection of a solution basis as well as a "rollback" mechanism propagating the cost of grounded actions (for finite graphs).

The second step ("online heuristic mode") determines possible heuristics from the text contents, to be validated interactively by the user. The underlying principle is that certain recommendations actually take the form of a heuristic rule selecting a course of action, which is equivalent (in a non-numerical form) to a heuristic for sub-task selection / task decomposition (see section 3).

During HTN traversal (automatically driven by AO*), the marked-up document is searched for occurrences of linguistic terms associated to the node under consideration – this may require to associate to each node a short list of key synonym terms corresponding to the most frequent formulations of the node's concept (in a subsequent step concept recognition from NL expressions could be envisioned, although the current approach seems to account for the vast majority of actual occurrences). When reaching an OR node, the heuristic to select the solution basis can be derived from the textual recommendations. The overall process could be described as an interactive AO* accessing textual decision elements within the document. When reaching an OR node, the algorithm would access those sections of the textual guideline referring to its key concept and will identify the closest or embedding recommendation. The marked-up recommendation will be presented together with its decision rule format making its role as a heuristic more visible to the user (Fig. 5).

The user then interprets the different recommendations highlighted by *G-DEE* in the dedicated interactive window and she can specify the corresponding heuristic value for the node considered. The interactive exploration of the HTN resumes after validation of a heuristic value by the user. The final output of an interactive session is a candidate explicit task decomposition (only containing AND nodes) which is ready for instantiation on the data specific to a patient profile (Fig. 6).

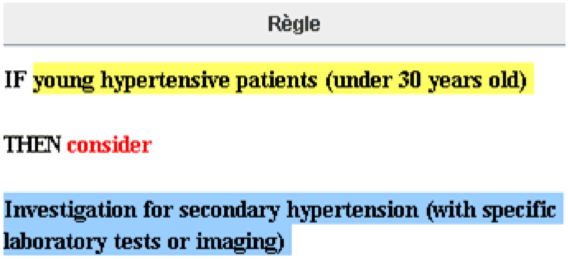

Fig. 5. Decision rule automatically derived from the recommendation described in Fig. 2 using dedicated XSL style sheets

3 Example Results: The French Hypertension Guidelines[1]

The overall system behavior consists in traversing the HTN from the top node and extracting a solution graph. In order to achieve this, data are accessed for an example patient, which will drive the instantiation of the various pre-condition of the plan operators. This is where it is important to determine which textual data is instantiating operators and which one is used by heuristic rules (e.g. such as age in the exploration of secondary hypertension).

For each node traversed, the terms attached to the node are first localized in the textual guideline, and the embedding recommendation (if any) is highlighted (Fig. 1 – C). As an example, we can consider a patient with: 1) a high level of cardiovascular risk and an antihypertensive monotherapy based on diuretics; 2) an inefficacy of this treatment. When the node "prescribe drug treatment" (Fig. 6) is selected in the HTN, three recommendations are highlighted in the clinical guidelines. The first recommendation reads "If the BP target is not achieved with first-line therapy, a combination of two drugs may be started as second-line therapy after at least 4 weeks", the second is "However, it may be started earlier in patients with BP • 180/110 mmHg regardless of the number of CVR factors; in patients with BP of 140-179/90-109 mmHg and a high CVR.", and the last: "If the patient does not respond to the initial therapy after 4 weeks or experiences side effects, a drug from a different therapeutic class should be prescribed". The heuristic that can be derived from these recommendations corresponds to the choice of therapy, i.e. bitherapy versus a change in therapeutic class or the prescription of a tritherapy. This node will typically be a successor of the "prescribe drug treatment" node in the HTN (Fig. 6). Some functions have been automated, such as the "offline heuristic mode" (and consequently the derivation of the sub-graph[2]).

In future versions of our system, the selection heuristics will also be automatically derived from a post-processing of the decision variable associated to the recommendation, which will be extracted using the same content extraction techniques than those identifying recommendations (with any remaining ambiguities interactively solved by the user).

[2] We used GraphViz (http://www.graphviz.org/) to interpret the file generated by *G-DEE* and enable to build png file of the HTN.

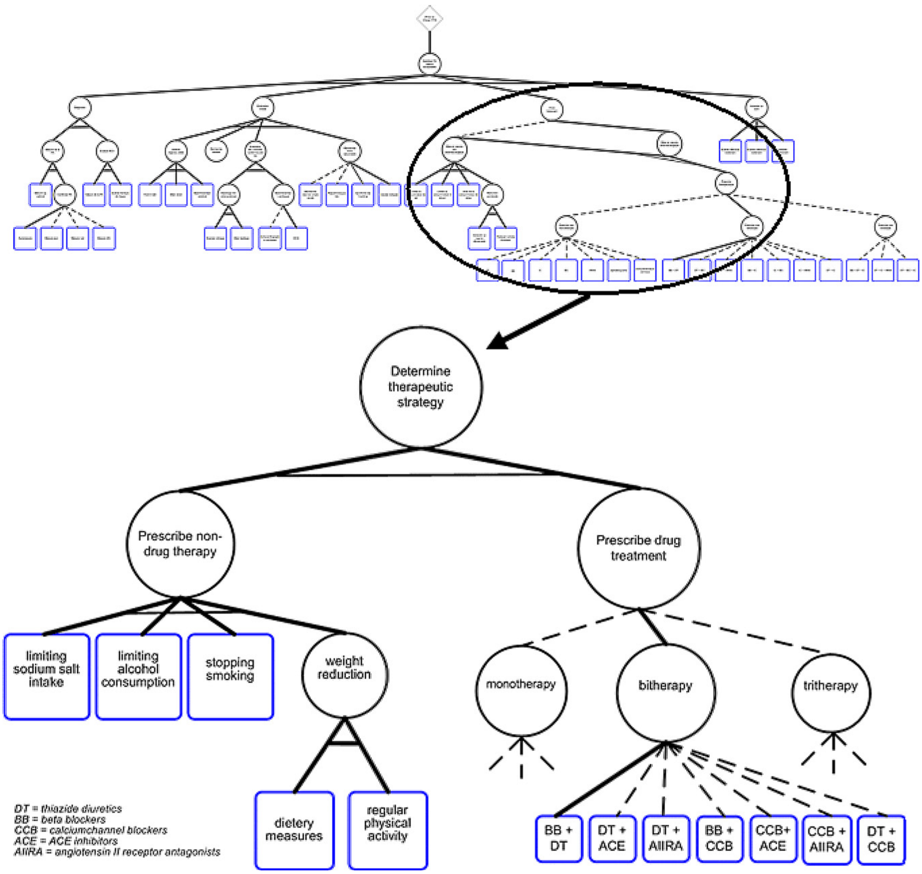

Fig. 6. A solution sub-graph (therapeutics) for initial patient data *(see text)*

4 Discussion

Our first experiments have shown that the structure of the document can be significantly disconnected from the logical flow of the clinical protocol. There are several possible explanations to this situation. Some are related to the social dynamics of guideline writing by committees, where, due to the necessity of achieving consensus, the inclusion and position in the text of some recommendations may not entirely reflect the logical flow of actions, at least within any given sections of the document (corresponding to a single phase of the protocol) or a protocol complexity which results in difficulties to present actions in a linear format in the document. Another aspect consists, when documents grow more complex, in various interim summarizations or the description of high-level strategies, which are generally redundant. Other descriptions encountered in the text are examples of plan outcomes, which illustrate the use of recommendations in context. The inclusion of example plan outcomes probably plays a useful explanatory role, as textual presentations do not offer the same kind of overview as more graphical ones such as HTN. On the other hand, textual guidelines seem to complement more formal representations such as HTN for which they constitute a rich source of contextual information (including heuristics), meta-knowledge and explanations/justifications.

5 Conclusions

The difference between document-based and knowledge-based approaches to clinical guidelines formalization is not simply a difference in the encoding of clinical knowledge or a difference in the semantics of the knowledge representation [22]. The two approaches seem to differ in the type on interpretation required and the meta-knowledge they contain. Because knowledge-based approaches often formalize protocols in an optimal way, producing a minimal and fully ordered structure, they can be useful to guide the exploration of textual guidelines, which often contain many additional data requiring interpretation and are sometimes structured in a less ordered fashion than the one allowed by hierarchical tasks descriptions.

Our current work on the system is dedicated to the progressive automation of HTN plan instantiation from textual guideline. Central to this is the automatic identification of decision variables from document content and their encoding in a rule format supporting automatic derivation of heuristics. Even when this step is achieved, the primary objective of our system will remain, in the first instance, the study of completeness and consistency of clinical guidelines.

Acknowledgments. Gersende Georg is funded through a post-doctoral fellowship from “Region Ile-de-France”. The HTN editor and Visio™ generator used in this paper has been developed by Fred Charles, who is thanked for his assistance.

References

1. Georg, G.: Computerization of Clinical Guidelines: an Application of Medical Document Processing. In: Silverman, B.G., Jain, A., Ichalkaranje, A., Jain, L.C. (eds.) Intelligent Paradigms for Healthcare Enterprises Systems Thinking, vol. 184, pp. 1–30. Springer, Heidelberg (2005)
2. Shiffman, R., Michel, G., Essaihi, A.: Bridging the guideline implementation gap: a systematic approach to document-centered guideline implementation. J Am Med Inform Assoc 11, 418–426 (2004)
3. Marcos, M., Balser, M., ten Teije, A., van Harmelen, F., Duelli, C.: Experiences in the formalisation and verification of medical protocols. In: Artificial Intelligence in Medicine, pp. 132–141. Springer, Heidelberg (2003)
4. Georg, G., Séroussi, B., Bouaud, J.: Does GEM-encoding clinical practice guidelines improve the quality of knowledge bases? A study with the rule-based formalism. In: Proceedings AMIA Symp, pp. 254–258 (2003)
5. Hagerty, C., Pickens, D., Chang, J., Kulikowski, C., Sonnenberg, F.: Prediction in Annotation Based Guideline Encoding. In: Proceedings AMIA Symp, pp. 314–318 (2006)

6. Fox, J., Johns, N., Rahmanzadeh, A., Thomson, R.: PROforma: A method and language for specifying clinical guidelines and protocols. In: Proc Medical Informatics Europe, pp. 516–520. IOS Press, Amsterdam (1996)
7. Johnson, P., Tu, S., Booth, N., Sugden, B., Purves, I.: Using scenarios in chronic disease management guidelines for primary care. In: Proceedings AMIA Symp, pp. 389–393 (2000)
8. Bradbrook, K., Winstanley, G., Glasspool, D., Fox, J., Griffiths, R.A.: Planning Technology as a Component of Computerised Clinical Practice Guidelines. In: Miksch, S., Hunter, J., Keravnou, E.T. (eds.) AIME 2005. LNCS (LNAI), vol. 3581, pp. 171–180. Springer, Heidelberg (2005)
9. Belker, T., Hammel, M., Hertzberg, J.: Learning to optimize mobile robot navigation based on HTN plans. In: Proceedings of International Conference on Robotics and Automation IEEE, pp. 4136–4141 (2003)
10. Smith, S., Nau, D., Throop, T.: Planning Approach to Declarer Play in Contract Bridge. Computational Intelligence 12, 106–130 (1996)
11. Cavazza, M., Charles, F., Mead, S.: Character-based Interactive Storytelling. IEEE Intelligent Systems 17, 17–24 (2002)
12. Quaglini, S., Stefanelli, M., Cavallini, A., Micieli, G., Fassino, C., Mossa, C.: Guideline-based careflow systems. Artif Intell Med 20, 5–22 (2000)
13. Shahar, Y., Miksch, S., Johnson, P.: The Asgaard Project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. Artificial Intelligence in Medicine 14, 29–51 (1998)
14. Haddawy, P., Doan, A., Goodwin, R.: Efficient Decision-Theoretic Planning: Techniques and Empirical Analysis. In: UAI '95 Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence, pp. 229–236. Morgan Kaufmann, San Francisco (1995)
15. Spyropoulos, C.: AI planning and scheduling in the medical hospital environment. Artificial Intelligence in Medicine 20, 101–111 (2000)
16. Ghallab, M., Nau, D., Traverso, P.: Automated Planning - Theory and Practice. Morgan Kaufmann, San Francisco (2004)
17. Nau, D., Cao, Y., Lotem, A., Muñoz-Avila, H.: SHOP: Simple Hierarchical Ordered Planner. In: Proceedings IJCAI-99 968–973 (1999)
18. Heuristics, P.J.: Intelligent Search Strategies for Computer Problem Solving. Addison-Wesley Publishing Company Inc, Reading (1985)
19. Georg, G., Jaulent, M.-C.: An Environment for Document Engineering of Clinical Guidelines. In: Proceedings AMIA Symp, pp. 276–280 (2005)
20. Moulin, B., Rousseau, D.: Knowledge acquisition from prescriptive texts. In: International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems. IEA/AIE, pp. 1112–1121 (1990)
21. 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)
22. Eriksson, H., Tu, S.W., Musen, M.: Semantic clinical guideline documents. In: Proceedings AMIA Symp, pp. 236–240 (2005)