International Journal of Medical Informatics (2005) **74**, 79—87

International Journal of Medical Informatics

ELSEVIER

www.intl.elsevierhealth.com/journals/ijmi

Extending the GEM model to support knowledge extraction from textual guidelines

Gersende Georg[a,*], Brigitte Séroussi[a,b], Jacques Bouaud[a,b]

[a] *SPIM, Inserm ERM 202, Université Paris 6, 15 Rue de l'Ecole de Médecine, F-75006 Paris, France*
[b] *STIM, DPA/DSI/AP-HP, Paris, France*

Received 31 October 2003; received in revised form 22 July 2004; accepted 23 July 2004

KEYWORDS
Clinical practice guidelines;
Decision support systems;
Decision rules;
Knowledge modeling;
GEM

Summary Clinical Practice Guidelines (CPGs) are being developed as a tool to promote best practice in medicine. However, the diffusion of paper guidelines has been shown to only have a limited impact. This is why computerization of CPGs has recently been suggested as a means to improve their dissemination as well as physicians' compliance. The Guideline Elements Model (GEM) has been proposed to facilitate the encoding of CPGs and support the automatic processing of marked-up documents.

In this paper, we explore the automatic generation of a rule base from a textual guideline using GEM. In this study, we propose an extension of the GEM model that introduces additional levels of structuring centered on decision variables. This allows a more efficient representation of the decision processes, which supports the automatic generation of decision rules from textual guidelines. The 1999 Canadian recommendations for the management of hypertension have been marked-up as a GEM-encoded instance of our extended DTD. We derived a rule base using an XML parser to extract the relevant elements to instantiate the IF and THEN clauses of decision rules. The rule base automatically generated compares favourably with the manual generation of decision rules in the ASTI project. This approach is an interesting case study in the computerization of CPGs, as it illustrates processing steps that are relevant to the various aspects of CPGs lifecycle, from production to consultation and use.

1. Introduction

Clinical practice guidelines (CPGs) are being developed to help physicians to improve their practice. CPGs are texts usually structured as sets of clinical situations for which therapeutic recommendations are provided. Several studies have suggested that

* Corresponding author.
E-mail addresses: gersende.georg@spim.jussieu.fr (G. Georg), bs@biomath.jussieu.fr (B. Séroussi), jb@biomath.jussieu.fr (J. Bouaud).

1386-5056/$ — see front matter
doi:10.1016/j.ijmedinf.2004.07.006

textual guidelines fail to have a significant impact on physician behavior. For example, in one study [1], 1775 patient treatments were retrospectively assessed for compliance with best practice, using questionnaires based on five performance criteria (hence corresponding to 8775 items). For each treatment, Mottur et al. determined the total number of self-reports of noncompliance with the guidelines. The self-assessed performance of 85 internists revealed an overall noncompliance rate of 24% (2073 of 8775 instances). The physicians' open-ended comments suggested that physician oversight, patient nonadherence, and systems issues were common reasons for noncompliance. One conclusion is that Decision Support Systems (DSSs) embedding the content of CPGs within their knowledge bases could be a more efficient way of disseminating best practice [2]. However, the development of knowledge bases corresponding to CPGs contents requires an appropriate method for knowledge acquisition. Numerous approaches to knowledge elicitation have been proposed, e.g. KADS [3], which aims at developing computable models in which a conceptual framework for knowledge communication among experts and developers is established [4]. Knowledge acquisition tools are available to support the development of clinical applications. For instance, Protégé-2000 [5] is an integrated knowledge base editing environment for the creation of customized knowledge-based tools. In the specific case of CPGs, the natural language content cannot be directly transcribed into the formalisms used in knowledge representation [6]. The production of knowledge structures from free text relies on human intervention, not simply for the interpretation of ambiguous statements, but also for the contextual interpretation of implicit information. This kind of interpretation is subjective and error prone and, as a result, the encoding of textual guidelines is subject to variations. Depending on the encoder's experience, competence, and medical expertise [7], incompleteness, ambiguities and imprecision often attached to textual CPGs may affect the quality of knowledge bases derived from CPGs contents. The Guideline Elements Model (GEM) [8,9] has been proposed as a document-based model, in order to structure guideline knowledge and promote the translation of textual guidelines into computer format. GEM provides an additional level of structure for the textual contents, which is based on semantic categories; the underlying hypothesis being that this should facilitate the contextual interpretation of CPGs contents required to instantiate a knowledge base.

In this paper, we present an experiment in using GEM to facilitate the generation of a knowledge base of production rules from textual guidelines, in the field of hypertension management. We propose an extension to the set of GEM categories that provides an additional level of structuring, which in turn facilitates the resolution of contextual ambiguities, especially those concerning the definition of the recommended sequence of therapies within the long term management of chronic diseases. The kinds of ambiguities we are addressing are of specific relevance to guidelines developed for the management of chronic diseases. One main reason being that they refer to clinical situations whose context evolves throughout the sequence of therapies and the evolution of the disease: this results in many information being left implicit and/or open to interpretation.

A normalized GEM-encoded instance was developed that automatically generate decision rules. It was applied to the 1999 Canadian recommendations for the management of hypertension [10]. The resulting rule base was compared for completeness and consistency to one manually encoded from the same guideline document by two physicians of the ASTI project [11].

The article is organized as follows: the next section gives some background elements on GEM and presents results of experiments using GEM to structure textual CPGs. The section 'Material' describes Canadian recommendations, the GEM DTD, and decision rules of the ASTI project. The automatic building of a rule-based knowledge base from textual Canadian recommendations is then described in the section 'Method'. A preliminary evaluation of the results, through the comparative assessment of knowledge bases is given in the section 'Results'.

2. Background

Various formalisms have been proposed to facilitate the production of medical knowledge bases. The oldest one, and the most widely used, is the Arden Syntax [12] in which Medical Logic Modules (MLMs) support clinical decision by the generation of alerts and reminders. This formalism, however, does not specifically address the problem of knowledge acquisition from text, which characterizes the use of CGPs. More recently, the GuideLine Interchange Format (GLIF) [13] proposed to model guideline content as a flowchart of structured steps representing clinical actions and decisions.

''Document-centric systems'' [6] are at the other end of the decision support spectrum. For instance, ActiveGuidelines [14] uses markup technologies to extract guideline fragments relevant to a specific patient situation. In the

PROforma project, Steele and Fox [15] have also proposed to augment guideline documents with decision-making services based on a task model. GEM [8] takes a different approach and proposes a comprehensive mark-up system for annotating a guideline document. By describing concepts relevant to guideline representation, attributes of these concepts and relationships among them, GEM is intended to serve as a document model of CPGs and aims at promoting the translation of textual guidelines into a format that can be processed by computers [16,17,18]. For instance, an XML-based application that facilitates the automated generation of partially populated MLMs from GEM-encoded guidelines has been published [16]. With the same objectives but using more complex formalisms, Shahar et al. have recently proposed the Digital Electronic Guidelines Library (DeGeL) [19], a whole software suite that aims at facilitating translation of textual clinical guidelines to a formal representation using the *Asbru* ontology [19]. However, guideline marking-up has several limitations. Although it has been found comprehensive enough to model the information content of CPGs, substantial variation is still observed in the creation of the GEM-encoded instance of a given CPG by different persons [9]. As the model is simply an abstraction of the guideline document, GEM alone does not support the resolution of ambiguities present in many textual guidelines.

Using as a test case the ambiguities present in the Canadian CPGs for the management of hypertension [10], we propose a framework to represent the sequence of therapeutic decisions that facilitates contextual interpretation, hence ambiguity resolution. This framework is implemented through an extension of GEM, refining the model's granularity through additional attributes. We encode the guideline document using our extension of GEM and use this structure to automatically derive a set of decision rules by parsing the XML structure of the encoded document. We then evaluate this rule base comparatively to a rule base manually developed in the ASTI project from the same guidelines [10].

3. Material

3.1. The Canadian recommendations for the management of hypertension

The 1999 Canadian recommendations for the management of hypertension [10] are chosen as our reference guideline document. It is structured in chapters corresponding to specific clinical situations. For instance, the case of ischemic heart disease as a complicating factor of hypertension is presented in Fig. 1. Within each chapter, an ordered sequence of therapeutic recommendations is proposed.

This textual CPG suffers from incompleteness (no recommendation for complex polypathological patient conditions), and fuzziness of the terms used (*stable* angina, *recent* myocardial infarction, *normal* left ventricular function, etc.). The chronological sequence of therapeutic recommendations is also open to some interpretation, as the various items (1 and 4) can be interpreted as alternative initial conditions, or as corresponding to different stages of disease progression.

3.2. The GEM DTD

GEM is a guideline document model based on an XML DTD [8] that organizes the heterogeneous knowledge contained in textual practice guidelines. It is a multi-level hierarchy of more than 100 discrete elements structured in nine major branches. Among them, the *knowledge components* section represents the recommendation's logic and constitutes ''the essence of practice guidelines''. It is this aspect of GEM that we extended to produce our encoding framework, focusing on those elements signaling therapeutic decisions. We extended the *con-*

VIII Ischemic heart disease

1. For patients with stable angina and hypertension, β-adrenergic antagonists are preferred as initial therapy (grade D).
2. Alternative therapies would include long-acting calcium-channel blockers (grade B). Short-acting calcium-channel blockers should not be used (grade C).
3. Patients with hypertension and a recent myocardial infarction should be treated with either β-adrenergic antagonists, ACE inhibitors or both. Both classes of drug protect against reinfarction and death (grade A).
4. Alternative therapies would include verapamil (grade A) and diltiazem (grade C), but only in the setting of normal left ventricular function.

Fig. 1 Therapeutic recommendations for hypertensive patients with ischemic heart disease.

ditional element that represents recommendations applicable under specific circumstances, by adding new attributes. Only a few of these sub-elements are actually used: *decision.variable* to describe elements of the decision, *action* to describe the recommended therapy, *recommendation. strength* to quantify the level of proof, and *evidence.quality* to establish the quality of the recommendation. Even though this extension appears simple in terms of the additional categories introduced, its real power derives from the additional level of structuring that has a strong impact on the elicitation of rule content.

3.3. Manual derivation of decision rules from CPGs (the ASTI project)

ASTI (''Aide à la Stratégie Thérapeutique Informatisée'') is a French project [11] which aim is to develop a guideline-based DSS to enable general practitioners to avoid prescription errors and to improve compliance with best therapeutic practices. The knowledge base is formalized as ''IF-THEN'' production rules, and has been manually built from the Canadian CPGs [10] by two physicians. The condition parts of the rules represent clinical situations descriptions and the action parts correspond to the set of recommended actions including the grade of the recommendation.

4. Method

Tierney et al. [20] recommended that guideline developers structure recommendations as ''IF-THEN-ELSE'' statements, to structure the use of guideline knowledge at each decision step. The aim of this work is to automatically generate a set of canonical decision rules from guideline contents: this automatic generation is made possible through the encoding of guideline contents using an extension of the GEM approach. The rationale for the generation of IF-THEN rules is the existence of decision variables in the CPG. More precisely, decision rules are represented as IF-THEN-WITH statements, where the condition (IF) part corresponds to a set of *decision.variable* elements of the GEM DTD, the action (THEN) part corresponds to a set of *action* elements, and the evidence (WITH) part corresponds to the *id* of the *recommendation.strength* element. To enable the generation of such rules from a GEM-encoded instance, it was necessary to modify the GEM encoding scheme to reflect the importance of decision variables, and obtain the same structure for both decision variables and actions in the DTD. This is why we first extended the original GEM DTD. Then, we proposed a framework to resolve the ambiguities of the original guideline document with respect to the therapeutic strategies. We started by producing a new encoded instance for the Canadian CPGs using our extension of GEM. We then developed a module to automatically derive decision rules from this GEM-encoded instance (Fig. 2). Finally, as a preliminary form of validation, we compared the resulting GEM-based rule base to the one produced manually by two physicians in the ASTI project [11]. This comparison was based on two criteria: (i) descriptive, assessing the number of rules generated as well as their average granularity and complexity, and (ii) operational, by comparing the therapeutic recommendations proposed by both approaches

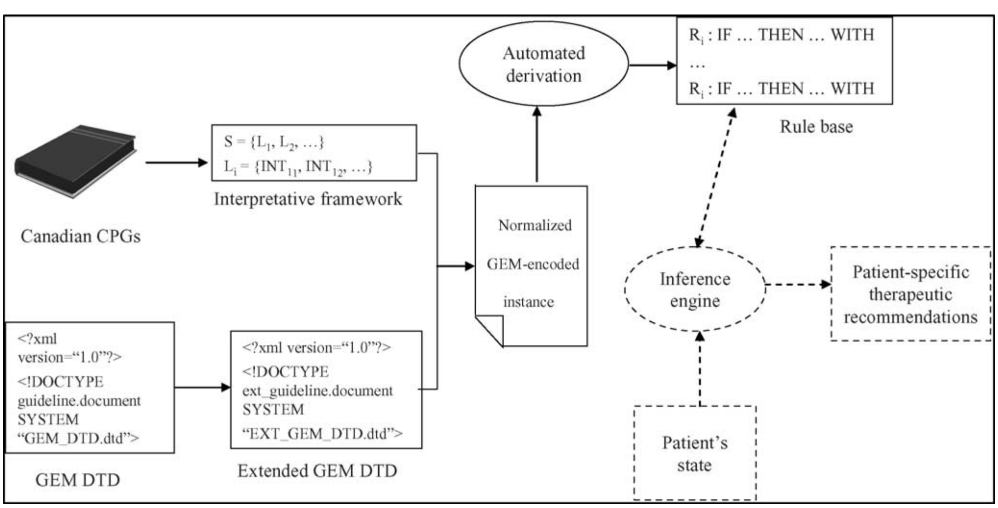

Fig. 2 Automatic generation of decision rules from a GEM-encoded CPG.

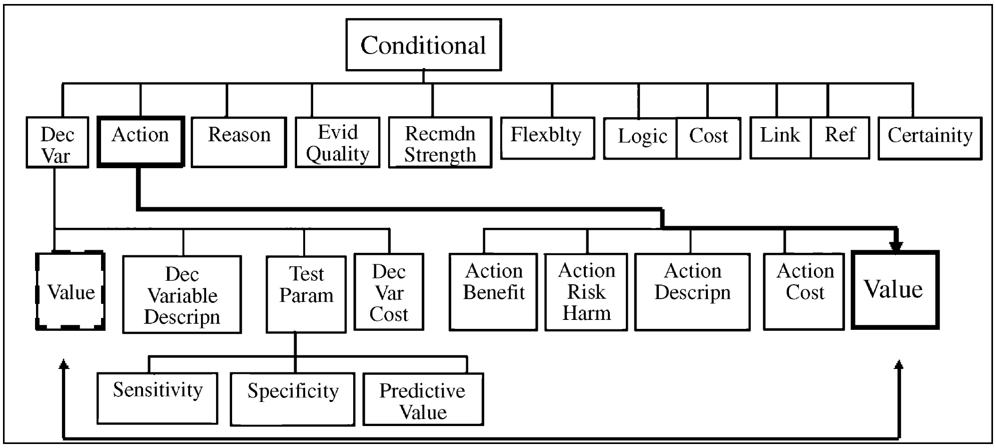

Fig. 3 The extended GEM DTD with the *value* sub-element for the *action* element.

on the same sample of 10 hypertensive patient cases.

4.1. Extension of the original GEM DTD

In the GEM DTD, conditional recommendations mainly rely on *decision.variable* and *action* elements. Decision variables are described by a value, a description, test parameters and a cost (Fig. 3). Actions descriptions are structured in various fields, i.e. benefit, risk and cost that can be grouped together into a single ''action parameter'' field. As discussed above, to enable the automated derivation of rules, we need a homogeneous data model for *decision.variable* and *action* elements. As the decision variable contains a *value* sub-element, we added a similar field to the *action* element.

Another extension to the GEM model concerns the *structure* of actions, represented through the notion of therapeutic strategy (lines of treatment and level of intention). We discuss the representation of therapeutic strategies in the next section; examples of the corresponding GEM encoding are given in the section 'Production of an extended GEM document instance'.

4.2. Representing therapeutic strategy through guidelines' decision steps

In the follow-up of chronic diseases, therapeutic recommendations depend on the patient state and on his therapeutic history, i.e. prior prescriptions that either were not adequate or provided unacceptable side effects. This is what makes chronic diseases a relevant test case for our approach. To resolve guideline ambiguities in the presentation of the chronological steps of the recommended therapy, we propose a framework formalizing the therapeutic strategy for a given patient profile.

A therapeutic strategy S is represented by an ordered sequence of therapeutic lines; each therapeutic line is made of a set of treatments ordered according to therapeutic levels of intention. According to a patient clinical situation and his response to the ongoing treatment, the recommended treatment may be either the next level of intention within the same therapeutic line or the first level of intention of the following therapeutic line.

For instance, the therapeutic strategy for patients with stable angina corresponding to recommendations described in items 1 and 2 in Fig. 1 is represented with one line with two levels of intention as follows:

$$S = (\beta - \text{adrenergic antagonists, calcium} \\ - \text{channel blockers})$$

$$\text{INT}_{1_1} = \beta - \text{adrenergic antagonists}$$

$$\text{INT}_{1_2} = \text{(substitution in the case of intolerance)} \\ = \text{calcium} - \text{channel blockers}$$

4.2.1. Marking-up the Canadian CPGs

We first marked-up the original document to identify those sections matching *decision.variable* and *action* elements. For instance, in the case of the first item of recommendations for ischemic heart disease, we identified the patient condition as a *decision.variable* and the recommended therapy as an *action* (Fig. 4).

```
<decision.variable source = "explicit"> For patients with stable angina and hypertension </decision.variable>
<action source = "explicit"> β-adrenergic antagonists are preferred as initial therapy (grade D) </action>
```

Fig. 4 Marked-up section corresponding to the first recommendation for ischemic heart disease.

```
<decision.variable source = "explicit" decision.variable.id = "state_patient.pathology"> hypertension
        <value source = "implicit" value.id = "HT"/> </decision.variable>
<decision.variable source = "explicit" decision.variable.id = "state_patient.pathology"> ischemic heart
disease
        <value source = "implicit" value.id = "ISC_HEA_DIS"/> </decision.variable>
<decision.variable source = "explicit" decision.variable.id = "state_patient.pathology"> stable angina
        <value source = "implicit" value.id = "STA_ANG"/> </decision.variable>
```

Fig. 5 Refined GEM encoding for decision variables.

4.2.2. Production of an extended GEM document instance

We identified three classes of patient parameters, i.e. age, risk factors, and associated diseases represented by ''state_patient.age'', ''state_patient.risk_factor'' and ''state_patient.pathology'' in corresponding *value.id* element of decision variables. We then refined the GEM encoding for those sections corresponding to decision variables, in order to introduce the above parameters as an additional level of structuring.

Finally, we associated mnemonics to the various *decision.variable* elements (Fig. 5). These mnemonics will serve as parameters of the IF-THEN rules automatically generated from the GEM-encoded document instance.

In the specific case of the first recommendation for ischemic heart disease (Fig. 1), ''initial therapy'' is interpreted as the first level of intention of the first line of therapy. The type of therapy is specified (monotherapy) as well as the nature of the pharmacological drug class (β-adrenergic antagonists).

The introduction of mnemonics was also performed for actions (Fig. 6) leading to the instantiation of the *value* sub-element of the action element introduced in the extended DTD. At the end of this encoding step, using the extended GEM model, the encoded instance contains specific representations for actions (both in terms of structures and variable identifications) that make possible the automatic generation of decision rules.

4.3. Derivation of the rule base

Target decision rules are represented as: ''IF *decision.variable* THEN *action* WITH *recommendation.strength*''. The construction of the rule base relies on the identification of *decision.variable*, *action*, and *recommendation.strength* elements in the GEM-encoded instance. The rule generation process is based on the simultaneous extraction of actions' structure and contents from the GEM-encoded instance, using an XML parser. In the example represented in Fig. 5, ''state_patient.pathology'' and

```
<action source = "explicit" action.id = "treatment.line"> first line treatment
        <value source = "implicit" value.id = "L1"/> </action>
<action source = "explicit" action.id = "treatment.intention"> first intention
        <value source = "implicit" value.id = "INT1"/> </action>
<action source = "explicit" action.id = "treatment.type"> monotherapy
        <value source = "implicit" value.id = "MONO"/> </action>
<action source = "explicit" action.id = "treatment.nature"> β-adrenergic antagonists
        <value source = "implicit" value.id = "BB"/> </action>
```

Fig. 6 Refined GEM encoding for actions: therapeutic lines.

''STA_ANG'' are the identification (id) of corresponding values extracted from *decision.variable* element, i.e. ''stable angina''. In the IF part of the rule, we have thus ''state_patient.pathology = STA_ANG'' which corresponds to the textual sentence of the CPG: ''For patient with stable angina'' (Fig. 1).

The process of extraction is the same for the entire instance. For the previous example, we obtained the following rule:

IF	state_patient.pathology = HT
	and state_patient.pathology = ISC_HEA_DIS
	and state_patient.pathology = STA_ANG
THEN	treatment.line = L1
	and treatment.intention = INT1
	and treatment.type = MONO
	and treatment.nature = BB
WITH	recommendation.strength = D

4.4. Evaluation of the derived rule base

We compared the rule base derived from the GEM-encoded instance, denoted $BR_{\text{-GEM}}$, to the one produced manually by two physicians in the ASTI project, denoted $BR_{\text{-ASTI}}$, according to both descriptive and operational criteria.

On the descriptive side, we compared both rule bases on a quantitative basis, i.e. the number of rules, the number of premises in IF parts, the number of actions in THEN parts, and on a qualitative evaluation, i.e. the number of clinical situations which are taken into account by the two approaches. On the operational side, we first developed a simple forward chaining inference engine to exploit $BR_{\text{-GEM}}$. Then, we compared the therapeutic recommendations proposed by both approaches on a sample of eight hypertensive patient cases.

5. Results

The automatically generated rule base, $BR_{\text{-GEM}}$, comprises 104 rules versus 98 rules for the rule base manually generated from the same guidelines, $BR_{\text{-ASTI}}$.

5.1. Descriptive criteria

5.1.1. Quantitative comparison

In both approaches, IF parts correspond to patient clinical descriptions. The difference in encoding between our approach and the ASTI approach is responsible for differences in rule granularity. For instance, the therapeutic level of intention is encoded by a unique attribute in $BR_{\text{-ASTI}}$. On the contrary, following the interpretative framework we previously introduced, steps of the therapeutic strategy are characterized in $BR_{\text{-GEM}}$ by a therapeutic line and a therapeutic level of intention. Yet the important aspect is not simply rule granularity, as approximated by the number of premises, but the fact that our approach provides a more consistent decomposition of premises. THEN parts are similarly formalized in both approaches and characterize the therapeutic class recommended by the guideline in the clinical situation described by the IF part. Overall, $BR_{\text{-GEM}}$ rules tend to be more specific than $BR_{\text{-ASTI}}$ rules. Since the average number of premises constitutes an indication of specificity in this context, we can use it to compare the two rules bases (Table 1).

Table 1 Quantitative description of $BR_{\text{-ASTI}}$ and $BR_{\text{-G}}$

	$BR_{\text{-ASTI}}$	$BR_{\text{-GEM}}$
No. of elementary rules	98	104
No. of premises (mean value)	2.93	4.49
No. of actions (mean value)	3.10	4.42

5.1.2. Qualitative comparison

The differences observed between $BR_{\text{-GEM}}$ and $BR_{\text{-ASTI}}$ can be in part explained by the ambiguity of the Canadian CPGs that allows for different interpretations of some parts of the document. $BR_{\text{-GEM}}$ describes 30 clinical situations, whereas $BR_{\text{-ASTI}}$ covers 19 clinical situations. Only 15 clinical situations are common to $BR_{\text{-GEM}}$ and $BR_{\text{-ASTI}}$. For instance, the case of patients under 60 years, suffering of hypertension with diabetes and without overt nephropathy correspond to a clinical situation that is commonly represented by both $BR_{\text{-GEM}}$ and $BR_{\text{-ASTI}}$. Fifteen clinical situations are then specific to $BR_{\text{-GEM}}$. Among them, eight correspond to clinical situations described as chapter headers of the CPG that have not been taken into account in $BR_{\text{-ASTI}}$. This concerns two situations of patients with cerebrovascular disease, three situations of patients with peripheral vascular disease, two situations of patients with hyperuricemia and gout, and one situation of patients with hyperlipidemia. However, the seven remaining $GEM_{\text{-specific}}$ situations are equivalent to five clinical situations described in $BR_{\text{-ASTI}}$, only at a greater level of abstraction. It can be noted that four clinical situations are actually specific to $BR_{\text{-ASTI}}$, but they correspond to ''particular'' textual interpretations of the guideline.

5.2. Evaluation on real patient cases

We compared the treatments recommended by the GEM-based system and the ASTI rule base on a sample of eight hypertensive patient cases. From the eight analyzed cases, therapies recommended by both approaches were identical in 37% of the cases (3/8), and compatible in 40% of the cases (2/5), i.e. the intersection of the therapies recommended with both approaches was not empty. When the recommended therapies differed, the GEM-based approach appeared to provide better justified recommendations.

6. Conclusion

CPGs are being developed as a tool to promote best practice in Medicine. However, the diffusion of paper guidelines only had a limited impact. This is why computerization of CPGs has recently emerged to improve their dissemination and the compliance of physicians with best practice. Within this context, GEM aims at promoting the translation of textual guidelines into a format that can be more easily processed by computers.

In this paper, we proposed an approach for automatically generating decision rules from GEM-encoded instance. We first proposed an extension of the GEM model that introduces additional levels of structuring centered on decision variables. This allows a more efficient representation of the decision processes, which supports the automatic generation of decision rules from textual guidelines. The rule base automatically generated compares favourably with the manual generation of decision rules. These are preliminary results to be confirmed in the framework of a larger study. However, GEM already appears to facilitate the encoding of CPGs, which support various aspects of guideline computerization. The proposed approach can indeed be extended to support the computerization of CPGs through the various aspects of their lifecycle, from production to consultation and use. The generation of decision rules itself can play an important role at various steps of the CPGs workflow: DSS can be used to assess the consistency of textual guidelines at the time of writing, while at the other end of the workflow they can provide assistance to guideline implementation. In that sense the case study in automatic rule base generation we have presented is relevant to the overall problem of guideline computerization.

References

[1] C. Mottur-Pilson, V. Snow, K. Bartlett, Physician explanations for failing to comply with ''best practices'', Eff. Clin. Pract. 4 (2001) 207—213.
[2] I. Sim, P. Gorman, R.A. Greenes, R.B. Haynes, B. Kaplan, H. Lehmann, P.C. Tang, Clinical decision support systems for the practice of evidence-based medicine, J. Am. Med. Inform. Assoc. 8 (6) (2001) 527—534.
[3] B.J. Wielinga, A.T.H. Schreiber, J.A. Breuker, KADS: a modelling approach to knowledge engineering, Knowledge Acquis. 4 (1992) 5—23.
[4] H. Eriksson, Models for knowledge-acquisition tool design, Knowledge Acquis. 6 (1994) 47—74.
[5] S.W. Tu, M.A. Musen, Modeling data and knowledge in the EON guideline architecture, Medinfo 10 (Pt 1) (2001) 280—284.
[6] R.D. Shankar, S.W. Tu, S.B. Martins, L.M. Fagan, M.K. Goldstein, M.A. Musen, Integration of textual guideline documents with formal guideline knowledge bases, Proc. AMIA Symp. (2001) 617—621.
[7] V.L. Patel, V.G. Allen, J.F. Arocha, E.H. Shortliffe, Representing clinical guidelines in GLIF: individual and collaborative expertise, J. Am. Med. Inform. Assoc. 5 (5) (1998) 467—483.
[8] R.N. Shiffman, B.T. Karras, A. Agrawal, R. Chen, L. Marenco, S. Nath, GEM: a proposal for a more comprehensive guideline document model using XML, J. Am. Med. Inform. Assoc. 7 (5) (2000) 488—498.
[9] B.T. Karras, S.D. Nath, R.N. Shiffman, A preliminary evaluation of guideline content mark-up using GEM — an XML Guideline Elements Model, J. Am. Med. Inform. Assoc. 7 (Suppl.) (2000) 413—417.
[10] R.D. Feldman, N. Campbell, P. Larochelle, P. Bolli, E.D. Burgess, S.G. Carruthers, et al., Canadian recommendations for the management of hypertension, CMAJ 161 (12) (1999) SF1—SF25.
[11] B. Séroussi, J. Bouaud, H. Dréau, H. Falcoff, C. Riou, M. Joubert, C. Simon, G. Simon, A. Venot, ASTI: a guideline-based drug-ordering system for primary care, in: V.L. Patel, R. Rogers, R. Haux (Eds.), Medinfo, vol. 84, 2001, pp. 528—532.
[12] G. Hripcsak, Arden syntax for Medical Logic Modules, MD Comput. 8 (2) (1991) 76—78.
[13] L. Ohno-Machado, J.H. Gennari, S.N. Murphy, N.L. Jain, S.W. Tu, D.E. Oliver, E. Pattison-Gordon, R.A. Greenes, E.H. Shortliffe, G.O. Barnett, The GuideLine Interchange Format: a model for representing guidelines, J. Am. Med. Inform. Assoc. 5 (4) (1998) 357—372.
[14] P.C. Tang, C.Y. Young, ActiveGuidelines integrating web-based guidelines with computer-based patient records, Proc. AMIA Symp. (2000) 843—847.
[15] R. Steele, J. Fox, Enhancing conventional web content with intelligent knowledge processing, in: M. Dojat, E. Keravnou, P. Barahona (Eds.), Artificial Intelligence in Medicine, LNCS, vol. 2780, 2003, pp. 142—151.
[16] A. Agrawal, R.N. Shiffman, Using GEM-encoded guidelines to generate Medical Logic Modules, J. Am. Med. Inform. Assoc. 8 (Suppl.) (2001) 7—11.
[17] R.N. Shiffman, A. Agrawal, A.M. Deshpande, P. Gershkovich, An approach to guideline implementation with GEM, in: V.L. Patel, R. Rogers, R. Haux (Eds.), Medinfo, vol. 84, 2001, pp. 271—275.
[18] P. Gershkovich, R.N. Shiffman, An implementation framework for GEM-encoded guidelines, J. Am. Med. Inform. Assoc. 8 (Suppl.) (2001) 240—248.

[19] Y. Shahar, O. Young, E. Shalom, A. Mayaffit, R. Moskovitch, A. Hessing, M. Galperin, DeGeL: a hybrid, multiple-ontology framework for specification and retrieval of clinical guidelines, in: M. Dojat, E. Keravnou, P. Barahona (Eds.), AIME, LNAI 2780, 2003, pp. 122—131.

[20] W.M. Tierney, J.M. Overhage, B.Y. Takesue, L.E. Harris, M.D. Murray, D.L. Vargo, C.J. McDonald, Computerizing guidelines to improve care and patient outcomes: the example of heart failure, J. Am. Med. Inform. Assoc. 2 (1995) 316—322.