Published in final edited form as:

Artif Intell Med. 2007 February ; 39(2): 151–163. doi:10.1016/j.artmed.2006.07.011.

How can information extraction ease formalizing treatment processes in clinical practice guidelines? A method and its evaluation★

Katharina Kaiser^{a,*}, Cem Akkaya^a, and Silvia Miksch^{a,b}

^a Vienna University of Technology, Institute of Software Technology & Interactive Systems, Favoritenstrasse 9-11/188, A-1040 Vienna, Austria

^b Danube University Krems, Department of Information & Knowledge Engineering, Dr. Karl-Dorrek-Strasse 30, A-3500 Krems, Austria

Summary

Objective: Formalizing clinical practice guidelines (CPGs) for a subsequent computer-supported processing is a challenging, but burdensome and time-consuming task. Existing methods and tools to support this task demand detailed medical knowledge, knowledge about the formal representations, and a manual modeling. Furthermore, formalized guideline documents mostly fall far short in terms of readability and understandability for the human domain modeler.

Methods and material: We propose a new multi-step approach using information extraction methods to support the human modeler by both automating parts of the modeling process and making the modeling process traceable and comprehensible. This paper addresses the first steps to obtain a representation containing processes which is independent of the final guideline representation language.

Results: We have developed and evaluated several heuristics without the need to apply natural language understanding and implemented them in a framework to apply them to several guidelines from the medical subject of otolaryngology. Findings in the evaluation indicate that using semi-automatic, step-wise information extraction methods are a valuable instrument to formalize CPGs.

Conclusion: Our evaluation shows that a heuristic-based approach can achieve good results, especially for guidelines with a major portion of semi-structured text. It can be applied to guidelines irrespective to the final guideline representation format.

Keywords

Information extraction and integration; Clinical practice guidelines; Computer-interpretable guidelines; Guideline representation; Treatment processes; Time-oriented information; Otolaryngology

★This is an extended and revised version of K. Kaiser, C. Akkaya, and S. Miksch. Gaining Process Information from Clinical Practice Guidelines, Proceedings of the 10th Conference on Artificial Intelligence in Medicine (AIME), 2005.

© 2006 Elsevier B.V. All rights reserved.

*Corresponding author. Tel.: +43 1 58801 18839; fax: +43 1 58801 18899. kaiser@ifs.tuwien.ac.at (K. Kaiser), cemakkaya@gmx.at (C. Akkaya), silvia@ifs.tuwien.ac.at (S. Miksch).

1. Introduction

Computer-supported guideline execution is an important instrument for improving the quality of health care. To execute clinical practice guidelines (CPGs) in a computer-supported way, the information in the guideline, which is in plain textual form, in tables, or represented in flow charts, has to be formalized. Consequently, this means that a formal representation is required in order to make the information computable. Thus, several so called *guideline representation languages* have been developed to support the structuring and representation of various guidelines and protocols and to make possible different kinds of applications (see [1,2]).

Many researchers have proposed frameworks for modeling CPGs in a computer-interpretable and -executable format (a comprehensible overview can be found in refs. [1,2]). Each of these frameworks provides specific guideline representation languages. Most of these languages are sufficiently complex that the manual formalization of CPGs is a challenging project. Thus, research has to be directed in such a way that tools and methods are developed for supporting the formalization process. Currently, using these tools and methods the human guideline developer needs not only knowledge about the formal methods, but also about the medical domain. This results in a very challenging, but time-consuming and cumbersome formalization task.

Thus, we will look for new approaches that can facilitate the formalization process, and support the developer by providing these kinds of knowledge, as well as intelligent methods for a simplified guideline modeling processing.

Within the next section, we present related work of guideline formalization tools and information extraction (IE) systems. In Section 3, we propose our approach. Section 4 describes our method which is evaluated in Section 5. Our conclusions are covered in Section 6.

2. Related work

In this section, we present a short discussion of some relevant work describing guideline formalization tools as well as some examples of IE systems.

2.1. Guideline formalization tools

To support the formalization of clinical guidelines into a guideline representation language various methods and tools exist, ranging from simple editors to sophisticated graphical applications.

2.1.1. Markup-based tools—*Stepper* [3] is a tool that formalizes the initial text in multiple user-definable steps corresponding to interactive XML transformations. The result of each step is an increasingly formalized version of the source document. Both the markup and the iterative transformation process are carried out by rules expressed in a new transformation language based on XML. *Stepper* documents all activities. So the transformation process can easily be reviewed by other users. *Stepper* also provides an interface showing the interconnection between the source text and the model.

The *GEM Cutter* [4] transforms guideline information into the Guideline Elements Model (GEM) format [5], showing the original guideline document together with the corresponding GEM document and makes it possible to copy text from the guideline to the GEM document.

The GEM Cutter is similar to the *Document Exploration and Linking Tool/Addons (DELT/A)*, formerly known as Guideline Markup Tool (GMT) [6], which supports the translation of HTML documents into an XML language. DELT/A provides two main features: (1) linking between a textual guideline and its formal representation and (2) applying design patterns in the form of macros. DELT/A allows the definition of links between the original guideline and the target representation, which gives the user the possibility to find out where a certain value in the XML-language notation comes from. Therefore, if someone wants to know the origin of a specific value in the XML file DELT/A can be used to jump to the correlating point in the text file where the value is defined and the other way round. By means of these features the original text parts need not be stored as part of the target representation elements. The links clearly show the source of each element in the target representation. Additionally, there is no need to produce a guideline in natural language from the target representation, since the original text remains unaltered.

Uruz, part of the *Digital electronic Guideline Library (Degel)* framework [7], is a web-based markup tool, which resembles DELT/A but does not maintain links between different representations of the guideline. It can also be used to create a guideline document without using any source by directly writing into the knowledge roles of a target ontology. *Uruz* enables the user to embed in the guideline document terms originating from standard vocabularies, such as ICD-9-CM (International Classification of Diseases) for diagnosis codes, CPT-4 (Current Procedural Terminology) for procedure codes, and LOINC-3 (Logical Observation Identifiers, Names, and Codes) for observations and laboratory tests.

2.1.2. Graphic tools—*Asbru View* [8] uses graphical metaphors to represent *Asbru* [1] plans. It is a tool to make *Asbru* accessible to physicians, and to give any user an overview over a plan hierarchy.

The *Plan Body Wizard (PBW)* of the *Degel* framework [7] is used by medical experts for defining the guideline's control structure in the *Asbru* representation [1]. It enables a user to decompose the actions embodied in the guideline into atomic actions and other sub-guidelines, and to define the control structure relating them.

AREZZO and *TALLIS* [9] support the translation into *PROforma* [10] by means of graphical symbols representing the task types of the language.

Protégé [11] is a knowledge-acquisition tool that supports the translation into guideline representation languages such as EON [12], GLIF [1], or *PROforma* [10]. It uses specific ontologies for these languages, whereas parts of the formalization process can be accomplished with predefined graphical symbols. *AREZZO*, *TALLIS*, and *Protégé* represent the processes by means of flow charts.

But still, in all of the above-mentioned cases, the modeling process is complex and labor intensive. Therefore, methods are needed that can be applied to automate parts of the modeling task.

2.2. Information extraction systems

IE is an emerging natural language processing (NLP) technology whose function is to process unstructured, natural language text, to locate specific pieces of information, or facts in the text, and to use these facts to fill a database [13]. Similar to IE systems are *wrappers* which aim to locate relevant information in semi-structured data [14] and often do not need to apply NLP techniques due to a restricted grammatical structure of the information resources.

For developing both IE and wrapper systems, two approaches can be applied: (1) the knowledge engineering approach and (2) the automatic learning approach.

The former is customized manually to a given task (e.g., *FASTUS* [15], *PLUM* [16], and *PROTEUS* [17]). But manually generating extraction rules is a cumbersome and time-consuming task. Thus, research has been directed towards automating this task. The automatic approach takes a set of documents and outputs a set of extraction patterns by using machine learning techniques. Automatic learning systems can be categorized in three groups:

- (1) *Supervised learning systems* require a large set of training data to learn rules using machine learning techniques (e.g., *AutoSlog* [18], *LIEP* [19], *ANNIE*, part of the *GATE* framework [20], or *WHISK* [21]),
- (2) *Semi-supervised learning systems* (e.g., *Mutual Bootstrapping* [22], *Snowball* [23], or *EXDISCO* [24]), and
- (3) *Unsupervised learning systems*, where rules are learned by a small set of seed rules and an annotated corpus using bootstrapping methods (e.g., *AutoSlog-TS* [25] or *QDIE* [26]).

To cope with the problems of “wrapper generation” and “wrapper maintenance”, rule-based methods have been especially popular in recent years. Some techniques for generating rules in the realm of text extraction are called “wrapper induction” methods. These techniques have proved to be rather successful for IE tasks in their intended domains, which are collections of documents such as web pages generated from a template script [27-29]. However, wrapper induction methods do only extend well to documents specific to the induced rules.

In semi-automatic wrapper generation machine learning approaches are applied. Tools may support the design of the wrapper. Some approaches offer a declarative interface where the user shows the system what information to extract (e.g., [14,29]). Automatic wrapper generation tools use unsupervised learning techniques. Therefore, no training sets are necessary, just a post-generation tuning (e.g., [30,31]).

When developing an IE system one has to incorporate numerous criteria to decide which approach to apply [32]. These are the availability of training data, which counts for an automatic learning approach, or the availability of linguistic resources and knowledge engineers, where the knowledge engineering approach may be favored. Also the level of performance required and the stability of the final specifications are important factors which may be better fostered by the knowledge engineering approach.

3. Our approach: a multi-step transformation process

Most guideline representation languages are very powerful and thus very complex. They can contain many different types of information and data. We therefore decided to apply a multi-step transformation process (cf. Fig. 1). It facilitates the formalization process by using various intermediate representations that are obtained by stepwise procedures. The multi-step methodology is necessary, as a one-step or even a two-step modeling process was shown to be not sufficient to the modeler [33,34].

The benefits of the intermediate representations are:

- concise formalization process,
- different formats for various kinds of information,

- separate views and procedures for various kinds of information,
- application of specific heuristics for each particular kind of information, and
- simpler and more concise evaluation and tracing of each process step.

To process as large a class as possible of documents and information we need specific heuristics. These are applied to a specific form of information, for instance:

Different kinds of information. Each kind of information (e.g., processes, parameters) needs specific methods for processing. By presenting only one kind of information the application of the associated method is simpler and easier to trace.

Different representations of information. We have to take into account various ways in which the information might be represented (i.e., structured, semi-structured, or free text).

Different kinds of guidelines. CGPs exist for various diseases, various user-groups, various purposes, various organizations, and so on, and have been developed by various guideline developers' organizations. Therefore, we can speak about different classes of CGPs that may contain similar guidelines.

To transform information by applying IE methods, we generated specific templates that can present the desired information. The IE methods detect relevant information which is filled into the templates' slots for subsequent processing. In the next section, we present a method that extracts process information from clinical guidelines for otolaryngology using heuristic algorithms. The output of this method is a unified format, which can be transformed into the final representation. Detailed information as well as information about the further processing of the resulting representation to the *Asbru* representation is described in ref. [35].

4. Extracting process information of clinical practice guidelines

CPGs present effective treatment processes. One challenge when authoring CPGs is the detection of individual processes and their relations and dependencies. We try to detect these using IE. CPGs consist of semi-structured and free text. The resulting output can subsequently be processed to yield refined representations, leading ultimately to the representation in a specific guideline representation language.

Our main goal is to acquire treatment processes from CPGs. Each process is described by at least one sentence. This means that a sentence, for instance, '*Take acetaminophen or ibuprofen*', presents only one process and not a selection of two processes. The rules are extraction patterns which are based on syntactical and semantical constraints as well as delimiters.

In order to gain rules to extract the process information, we first had to choose guideline documents which are then used to obtain the rules and to test these rules with other CPGs. We have chosen guidelines from the National Guideline Clearing-house (NGC)¹ repository using several criteria. These criteria are the guideline category of treatment and management, the evidence-based quality of guidelines, the existence of treatment instructions featuring temporal aspects of flows, the document structure enabling the detection of text modules such as tables, lists, and paragraphs, and the clinical specialty. We obtained several guidelines from various clinical specialties and have chosen guidelines of otolaryngology. These resulting 18 guideline documents were developed by 10

¹See <http://www.guidelines.gov>.

organizations (see Table 1). We then divided the set into a training set of 6 guidelines (see Table 2) and a test set of 12 guidelines.

4.1. Obtaining extraction rules

From the training set of guidelines, we developed rules based on patterns for IE using the knowledge engineering approach. Patterns are defined on three levels, whereas patterns at a certain level serve as concept classes in the preceding levels: (1) phrase level patterns, (2) sentence level patterns, and (3) discourse level patterns. Pattern rules were designed using the atomic approach [32]. Thereby, a domain module is built that recognizes the arguments to an event and combines them into template structures strictly on the basis of intelligent guesses rather than syntactic relationships. In doing so domain-relevant events are assumed for any recognized entities, leading to high recall, but much overgeneration, and thus low precision. Further development would result in improving filters and heuristics for combining the atomic elements, improving precision.

Medical terms (i.e., drug agents, surgical procedures, and diagnostic terms) are based on a subset of the Medical Subject Headings (MeSH)² of the United States National Library of Medicine. We adapted them according to missing terms, different wordings, acronyms, and varying categorization.

4.1.1. Phrase level patterns—They are used for identifying basic entities, such as *time*, *dosage*, *iteration*, and *condition* expressions, which build the attributes of actions. They are defined by regular expressions.

4.1.2. Sentence level patterns—They use phrase level patterns, medical terms, and trigger words for the medical terms to identify medical actions and their attributes. The trigger words are mainly verbs and indicate the application of a therapy (e.g., the administration of a drug agent or the implementation of a surgical procedure) or the avoidance of a therapy. Sentence level patterns are delimiter-based and use syntactic constraints. We can categorize the patterns in two groups: (1) patterns for free text and (2) patterns for telegraphic text.

The former are applied to free text, which has a grammatical structure and is usually identified in paragraphs, but also in list elements. These patterns indicate that therapy instruments (i.e., agent terms and surgical procedures) combined with trigger terms (e.g., ‘activate’, ‘indicate’, ‘perform’, ‘pre-scribe’) appearing in the same clause identify relevant sentences. The particular clauses must not be condition clauses. Phrase level patterns, such as <dosage>, <duration>, <condition>, and so on, can be arbitrarily combined with <therapy instrument> <trigger> pairs. But information concerning a treatment recommendation can be distributed in several sentences. These sentences including additional information (e.g., ‘*The standard dose is 40–45 mg/kg/day.*’) neither contain a therapy instrument nor a trigger term, but also have to be identified by sentence patterns.

Telegraphic text patterns are applicable in list elements. In these elements, often ungrammatical text is formulated and therefore, there is no need for trigger terms. Often, only a therapy instrument indicates the relevancy of an element. Other patterns exist for list elements indicating that these elements are relevant if within their context or in the paragraph preceding the list special terms appear. These terms (i.e., ‘remedy’, ‘remedies’, ‘measure’, ‘measures’, ‘medication’, and ‘medications’) are important, because they specify actions that may not contain therapy instruments in the form of agent terms or surgical

²See <http://www.nlm.nih.gov/mesh/>.

procedures (e.g., ‘*Maintain adequate hydration (drink 6–10 glasses of liquid a day to thin mucus)*’).

4.1.3. Discourse level patterns—They are based on *sentence level patterns*, but are augmented to consider the structure and the layout of the documents. They are used to categorize sentences, merge them to actions, and find relationships between actions to structure them. To accomplish the latter task, we analyzed treatment processes contained in the guidelines and detected the following processes, whereas some of them are identified by discourse level patterns:

- processes without temporal dependencies,
- sequential processes,
- processes containing subprocesses,
- selections of processes, and
- recurring processes.

4.2. Gaining process information

To extract processes from CPGs we proceed in several steps which serve to filter segments of text containing treatment instructions from the documents and to generate processes. We propose a two-step approach (see Fig. 2) to gain a representation that is independent of the subsequent guideline representation language.

The first step is to extract the relevant sentences containing treatment instructions by marking-up the original guideline document. This is explained in Section 4.2.1. The subsequent step is to combine several sentences to one action and to structure the actions and detect relations among them. This step is described in Section 4.2.2.

These two steps should provide a basis for the subsequent transformation of the process information into any guideline representation language.

4.2.1. Extracting relevant sentences—This task is a first step towards our final guideline representation. We will achieve it by two modules: (1) the segmentation and filtering module and (2) the template generation module (see Fig. 3 for an overview).

This first intermediate step is especially important as not the entire content of a guideline contains processes, which are to be modeled. Although health care consists of the three stages, observation, diagnosis, and therapy [36], we only want to model the control flow regarding the therapy. Only about 20% of sentences of a guideline are of interest for modeling these processes. On this account, it is important to select the relevant sentences for modeling.

Thus, this task performs an automatic markup of sentences that are utilized to process the subsequent steps.

Segmentation and filtering: Detecting relevant sentences is a challenging task, which we undertake in two steps: (1) detecting irrelevant text parts to exclude them from further processing and (2) detecting relevant sentences. Irrelevant text parts (i.e., sections, paragraphs) are associated with diagnosis, symptoms, or etiology, relevant sentences describe actions of a treatment processes.

The first filtering occurs at the section level. Sections in the document with captions indicating diagnosis or symptom declarations will be omitted in further processing. We can identify these captions by keywords such as 'history', 'diagnosis', 'symptom', 'clinical assessment', 'risk factor', and so on.

Detecting relevant sentences is not a trivial task. First, we parse the entire document and split it into sentences. Then, we process every sentence with regard to its context within the document and its group affiliation. Thereby, the context is obtained by captions (e.g. '*Acute Pharyngitis Algorithm Annotations | Treatment | Recommendations:*') and a group contains sentences from the same paragraph or the same list, if there are no sublists. Each sentence is then checked for relevance by applying *sentence level patterns*.

Template generation: After having collected the relevant sentences from the guideline, we can proceed with generating the intermediate representation *SentenceIR*. We generate two files: one file listing all relevant sentences and the marked-up guideline document (Listing 1 shows the source of a marked-up guideline document). Both are linked by applying the same id to the same sentences. The presentation of the template file and the guideline document are as simple as possible in order to support the user by detecting all relevant sentences.

4.2.2. Extracting required information and finding processes—The information contained in *SentenceIR* and the marked-up guideline document are the input for the next task (see Fig. 4 for an overview). Its goal is to structure relevant sentences and find relationships between sentences. Again, the output of this task should be represented in a format that is independent of any desired guideline representation format.

Structure extraction: In this task, we obtain the context of each sentence by means of hierarchical groups which is necessary for other subtasks, especially the merging and grouping and the process extraction. Every action is assigned to one group. The context of a sentence defines the affiliation to a group and is defined by the sentence's position in the hierarchal structure. We use the superior headings that establish several context items.

Slot extraction: This module is used to extract therapy instruments (i.e., agent terms and surgical procedures), dosage information in case of a drug administration, the duration of the therapy action, the iteration information of the action, as well as conditions which have to be fulfilled to perform an action. It uses both the lexicon and the *phrase level patterns*.

Merging and grouping: In this module, we categorize sentences in actions or negative actions and annotations. Annotations always belong to at least one action (or negative action). They cannot exist alone. This module extensively applies *discourse level patterns*.

First, we check whether a sentence describes an action or a negative action. Negative actions are instructions that an action should not be performed, often under specific conditions (e.g., '*Do not use aspirin with children and teenagers because it may increase the risk of Reye's syndrome*'). Most guideline representation languages will handle such actions by inverting the condition. Languages may exist which will handle these in other ways. Therefore, we provide a representation for such actions that can be used in a general way.

Furthermore, we identify annotations and assign them to their corresponding actions or negative actions using *name-alias coreferencing* and *definite description coreferencing* based on therapy instruments and their hypernyms. We do not apply pronoun-antecedent coreferencing.

Process extraction: To group actions and to detect relationships between actions we use *discourse level patterns*. We will describe those used by this module below.

The default relationship among processes is that there is no synchronization in their execution. To group actions to a *selection* they must fulfill the following requirements: (1) the actions have to belong to the same group and (2) agents or surgical procedures must have the same superordinate. For instance, processes describing the administration of *Erythromycin*, *Cephalexin*, and *Clindamycin* within one group are combined in a *selection*, as all these agents are antibiotics. If actions are grouped in a selection, one of these actions has to be selected to be executed.

Furthermore, we try to detect relations between actions that are explicitly mentioned within the text as well as relations that are implicitly given by the document structure. The former is very difficult to detect, as we often cannot detect the reference of the relation within the CPG (e.g., ‘*After 10–14 days of failure of first line antibiotic...*’). Nevertheless, we found heuristics that arrange actions or action groups if the reference is unambiguously extractable out of the text. These heuristics can be grouped in two categories: (1) detecting sentences describing relations between actions and (2) detecting actions that are described in the preceding heuristic. A relation is mainly identifiable by a relation term (e.g., ‘before’, ‘after’, ‘during’, ‘while’). If such a term appears, we are searching for therapy instruments, as these describe most of our actions. After we have detected these terms, we search for actions containing the particular instruments. If we have found both the source action and the destination action we can create a new relation.

Template generation: The template of this intermediate representation has to contain actions as well as their relations. It has to be simple and concise and it has to illustrate from which original data the current information was built. We split the new *ActionIR* template in three parts: (1) an area for actions, (2) an area for relations, and (3) an area for the structure illustrating the hierarchy and nesting of groups.

An *action* contains the action sentence, the assigned annotation sentences, treatment instruments and their MeSH ids, information about the dosage, duration, or iteration of a drug administration, and conditions. If the action is part of a selection, it is stated by the selection id. DELT/A links are inherited from the *SentenceIR* representation in order to provide the traceability of the process from both the original guideline document and the *SentenceIR* document. Listing 2 shows an example instance.

We use patterns of the document structure (e.g., ‘*Further Treatment*’ appears after ‘*Treatment*’ or ‘*Treatment*’ appears before ‘*Follow-Up*’) to detect implicitly given relations. These patterns are part of discourse level patterns to determine relations between several groups.

Relations are stated by their type (e.g., succeeding, preceding, overlapping) and the concerned actions by their DELT/A ids.

Apart from actions and their relations, the *structure* of the document is given illustrating the nesting of the groups and selections.

5. Evaluation

The rules developed using training examples have to reach a state where they are able to extract the correct information from other examples, too. In order to test these acquirements, we developed Java applications that generate the intermediate representations.

The particular intermediate representations generated from the test set were evaluated by two persons using the DELT/A tool (see Section 2.1). The participants are computer scientists, who are familiar with guidelines, guideline formalization, and the DELT/A tool, but have no medical background. However, the chosen guidelines do not require specific medical knowledge to evaluate the IE tasks.

We evaluated our rules using recall and precision measures. The *recall* score measures the ratio of correct information extracted from the texts against all the available information present in the text. The *precision* score measures the ratio of correct information that was extracted against all the information that was extracted [37].

For the evaluation, we provided a test set of 12 guidelines in XHTML format and the necessary language and macro files for DELT/A. The participants generated key target templates for all test guidelines using the DELT/A tool. These were then compared to the templates generated by our system, whereas the input of the second step were the key target templates of the first step. For evaluating the markup task, we compiled the number of relevant sentences according to the key target template (POS), the number of relevant sentences generated by the system (ACT), and the number of correctly detected relevant sentences generated by the system (COR). Out of these values, we were able to compute the recall and precision scores (see Table 3).

In ref. [38], we described a preliminary framework which achieved results of 76% recall and 97% precision. However, many users claim that IE systems performing a detection of relevant text parts are only of use if they detect all the relevant parts (i.e., a recall score of 100%). Otherwise the user has to read the entire document to find the remaining relevant sentences. Thus, we optimized the system according to recall. The resulting values for recall of 90.8% and precision of 94.9% are promising and point out the benefit of the markup task. We are of opinion that even with the current performance a benefit exists, because relevant sentences are not equally spread on the whole document, but mostly constitute clusters which have to be verified then.

To verify the process extraction task, we again compared this task's key target templates to the system's output templates. We compiled the number of filled slots according to the key target template (POS), the number of slot fillers generated by the system (ACT), the number of correct slot fillers generated by the system (COR), and the number of partially correct slot fillers generated by the system (PAR). Starting from these values, we computed the recall (REC) and precision (PRE) scores (see Table 4). The overall scores of this task are 84% recall and 86.8% precision.

Further analyses of these results (see Table 5 for details) showed that they are mainly based on erroneous extractions of duration and iteration information.

The results of each subtask has to be seen within the context of the benefit of the automatically generated data compared to the manual generation using DELT/A. Thereby, the results still imply that using step-wise IE for generating a semi-formal representation of treatment instructions is a great benefit for both knowledge engineers and physicians.

6. Conclusions

Modeling clinical guidelines and protocols is a complex task which has to be assisted by both physicians and knowledge engineers. Bearing those two user groups in mind, a method is demanded supporting them in their particular fields of functions: the physicians have to be less overcharged by the formal specifications and the knowledge engineers have to be fostered by providing medical knowledge. Apart from this interesting conceptual

formulation, we have developed a new methodology applying a step-wise IE which might offer distinct benefits. In particular, it automates parts of the modeling process, it disburdens the physicians in the modeling process by providing a medical ontology, it structures the guideline information, it decomposes the guideline into parts containing various kinds of information (e.g., treatment processes, diagnosis methods, definitions), it makes the modeling process traceable and comprehensible, and it is applicable for many guideline representation languages.

We have shown that it is possible to semi-automatically model process information from CPGs using IE. Our rules use patterns in the structure of the documents as well as of specific expressions.

We have applied a framework in order to evaluate our rules that can cope with both semi-structured and free text documents. The resulting information is filled in templates which can represent processes and their relations. The information extracted can then be used in further transformations to finally generate a representation in a guideline representation language.

Acknowledgments

This work is supported by “Fonds zur Forderung der wissenschaftlichen Forschung FWF” (Austrian Science Fund), grants P15467-INF and L290-N04.

References

1. Peleg M, Tu SW, Bury J, Ciccarese P, Fox J, Greenes RA, et al. Comparing computer-interpretable guideline models: a case-study approach. *J Am Med Inform Assoc (JAMIA)*. 2003; 10(January–February (1)):52–68.
2. de Clercq PA, Blom JA, Korsten HHM, Hasman A. Approaches for creating computer-interpretable guidelines that facilitate decision support. *Artif Intell Med*. 2004; 31(May (1)):1–27. [PubMed: 15182844]
3. R ž i ka, M.; Svátek, V. Mark-up based analysis of narrative guidelines with the Stepper tool. In: Kaiser, K.; Miksch, S.; Tu, SW., editors. *Computer-based support for clinical guidelines and protocols; Proceedings of the symposium on computerized guidelines and protocols (CGP 2004)*, vol. 101, studies in health technology and informatics; IOS Press; 2004. p. 132-6.
4. Polvani, K-A.; Agrawal, A.; Karras, B.; Deshpande, A.; Shiffman, R. *GEM Cutter manual*. Yale Center for Medical Informatics; New Haven, CT: 2000.
5. Shiffman RN, Karras BT, Agrawal A, Chen R, Marengo L, Nath S. GEM: a proposal for a more comprehensive guideline document model using XML. *J Am Med Inform Assoc (JAMIA)*. 2000; 7(5):488–98.
6. Votruba, P.; Miksch, S.; Kosara, R. Facilitating knowledge maintenance of clinical guidelines and protocols. In: Fieschi, M.; Coiera, E.; Yu-Chuan, JL., editors. *Proceedings from the Medinfo 2004 world congress on medical informatics*; IOS Press; 2004. p. 57-61.
7. Shahar, Y.; Young, O.; Shalom, E.; Mayaffit, A.; Moskovitch, R.; Hessing, A., et al. DEGEL: a hybrid, multiple-ontology framework for specification and retrieval of clinical guidelines. In: Dojat, M.; Keravnou, E.; Barahona, P., editors. *Proceedings of the 9th conference on artificial intelligence in medicine in Europe, AIME 2003; LNAI*. Springer Verlag; 2003. p. 122-31.
8. Kosara R, Miksch S. Metaphors of movement: a visualization and user interface for time-oriented, skeletal plans. *Artif Intell Med*. 2001; 22(May (2)):111–31. special issue: information visualization in medicine. [PubMed: 11348843]
9. Steele, R.; Fox, J. Tallis PROforma primer–introduction to PROforma language and software with worked examples. Cancer Research; London, UK: 2002. Technical report, Advanced Computation Laboratory

10. Sutton DR, Fox J. The syntax and semantics of the *PROforma* guideline modeling language. *J Am Med Inform Assoc (JAMIA)*. 2003; 10(September–October (5)):433–43.
11. Gennari JH, Musen MA, Fergerson RW, Grosso WE, Crubézy M, Eriksson H, et al. The evolution of Protégé: an environment for knowledge-based systems development. *Int J Hum Comput Stud*. 2003; 58(1):89–123.
12. Musen MA, Tu SW, Das AK, Shahar Y. EON: a component-based approach to automation of protocol-directed therapy. *J Am Med Inform Assoc (JAMIA)*. 1996; 3(6):367–88.
13. Yangarber, R. Scenario customization for information extraction. New York University; New York, NY: Jan. 2001 PhD thesis
14. Kushmerick, N.; Weld, DS.; Doorenbos, RB. Wrapper induction for information extraction. In: Pollack, ME., editor. *International joint conference on artificial intelligence*; Morgan Kaufmann; 1997. p. 729-37.
15. Hobbs, JR.; Appelt, D.; Tyson, M.; Bear, J.; Israel, D. SRI International: description of the FASTUS system used for MUC-4. In: McLean, V., editor. *Proceedings of the 4th message understanding conference (MUC-4)*; Association for Computational Linguistics; 1992. p. 268-75.
16. Ayuso, D.; Boisen, S.; Fox, H.; Gish, H.; Ingria, R.; Weischedel, R. BBN: description of the PLUM system as used for MUC-4. In: McLean, V., editor. *Proceedings of the 4th message understanding conference (MUC-4)*; Association for Computational Linguistics; 1992. p. 169-76.
17. Yangarber, R.; Grishman, R. NYU: description of the Proteus/PET system as used for MUC-7 ST. *Proceedings of the 7th message understanding conference: MUC-7*; Morgan Kaufmann; 1998.
18. Riloff, E. Automatically constructing a dictionary for information extraction tasks. In: Fikes, R.; Lehnert, W., editors. *Proceedings of the 11th national conference on artificial intelligence (AAAI-93)*; American Association for Artificial Intelligence (AAAI), AAAI Press/MIT Press; 1993. p. 811-6.
19. Huffman, SB. Learning information extraction patterns from examples. In: Stefan, W.; Ellen, R.; Gabriele, S., editors. *Connectionist, statistical, and symbolic approaches to learning for natural language processing*, vol. 1040, lecture notes in computer science. Springer Verlag; 1996. p. 246-60.
20. Cunningham, H.; Maynard, D.; Bontcheva, K.; Tablan, V. GATE: an architecture for development of robust HLT applications. In: Pierre, I., editor. *Proceedings of the 40th anniversary meeting of the association for computational linguistics (ACL'02)*; Association for Computational Linguistics; 2002. p. 168-75.
21. Soderland S. Learning information extraction rules for semi-structured and free text. *Machine Learning*. 1999; 34(1–3):233–72.
22. Riloff, E.; Jones, R. Learning dictionaries for information extraction by multi-level bootstrapping. In: Hendler, J.; Subramanian, D., editors. *Proceedings of the national conference on artificial intelligence (AAAI)*; American Association for Artificial Intelligence (AAAI), AAAI Press/MIT Press; 1999. p. 474-9.
23. Agichtein, E.; Gravano, L. Snowball: extracting relations from large plaintext collections. In: Nürnberg, PJ.; Hicks, DL.; Furuta, R., editors. *Proceedings of the 5th ACM international conference on digital libraries*; ACM Press; 2000. p. 85-94.
24. Yangarber, R.; Grishman, R.; Tapanainen, P.; Huttunen, S. Automatic acquisition of domain knowledge for information extraction. In: Kay, M., editor. *Proceedings of the 18th conference on computational linguistics*; Association for Computational Linguistics; 2000. p. 940-6.
25. Riloff E. An empirical study of automated dictionary construction for information extraction in three domains. *Artif Intell*. 1996; 85(1–2):101–34.
26. Sudo, K. Unsupervised discovery of extraction patterns for information extraction. New York University; New York, NY: Sep. 2004 PhD thesis
27. Muslea, I.; Minton, S.; Knoblock, C. A hierarchical approach to wrapper induction. In: Etzioni, O.; Müller, JP.; Bradshaw, JM., editors. *Proceedings of the international conference on autonomous agents*; ACM Press; 1999. p. 190-7.
28. Kushmerick N. Wrapper induction: efficiency and expressiveness. *Artif Intell*. 2000; 118(12):15–68.

29. Baumgartner, R.; Flesca, S.; Gottlob, G. Visual web information extraction with Lixto. In: Apers, PMG.; Atzeni, P.; Ceri, S.; Paraboschi, S.; Ramamohanarao, K.; Snodgrass, RT., editors. Proceedings of the conference on very large databases (VLDB); Morgan Kaufmann Publishers Inc.; 2001. p. 119-28.
30. Doorenbos, RB.; Etzioni, O.; Weld, DS. A scalable comparison-shopping agent for the world-wide web. In: Lewis Johnson, W.; Hayes-Roth, B., editors. Proceedings of the international conference on autonomous agents; ACM Press; 1997. p. 39-48.
31. Crescenzi, V.; Mecca, G.; Meriardo, P. Roadrunner: towards automatic data extraction from large web sites. In: Apers, PMG.; Atzeni, P.; Ceri, S.; Paraboschi, S.; Ramamohanarao, K.; Snodgrass, RT., editors. Proceedings of the conference on very large databases (VLDB); Morgan Kaufmann Publishers Inc.; 2001. p. 109-18.
32. Appelt DE. Introduction to information extraction. *AI Commun.* 1999; 12:161–72.
33. Balser, M.; Coltell, O.; van Croonenborg, J.; Duelli, C.; van Harmelen, F.; Jovell, A., et al. Protocure: integrating formal methods in the development process of medical guidelines and protocols. In: Kaiser, K.; Miksch, S.; Tu, SW., editors. Computer-based support for clinical guidelines and protocols; Proceedings of the symposium on computerized guidelines and protocols (CGP 2004), vol. 101, Studies in health technology and informatics; IOS Press; 2004. p. 103-7.
34. Seyfang, A.; Miksch, S.; Polo-Conde, C.; Wittenberg, J.; Marcos, M.; Rosenbrand, K. MHB—a many-headed bridge between informal and formal guideline representations. In: Miksch, S.; Hunter, J.; Keravnou, E., editors. Proceedings of the 10th conference on artificial intelligence in medicine in Europe, AIME 2005, vol. 3581; LNAI. Springer Verlag; 2005. p. 146-50.
35. Kaiser, K. LASSIE—modeling treatment processes using information extraction. Institute of Software Technology & Interactive Systems: Vienna University of Technology; 2005. PhD thesis
36. Van Bommel, JH.; Musen, MA., editors. Handbook of medical informatics. Springer Verlag; Heidelberg: 1997.
37. Lehnert W, Cardie C, Fisher D, McCarthy J, Riloff E, Soderland S. Evaluating an information extraction system. *J Integr Comput Aided Eng.* 1994; 1(6):453–72.
38. Kaiser, K.; Akkaya, C.; Miksch, S. Gaining process information of clinical practice guidelines using information extraction. In: Miksch, S.; Hunter, J.; Keravnou, E., editors. Proceedings of the 10th conference on artificial intelligence in medicine in Europe, AIME 2005, vol. 3581; LNAI, Springer Verlag; 2005.

Figure 1. Guideline transformation process. A multi-step process using intermediate representations to transform CPGs into a formal representation language.

Figure 2.

Steps to obtain a semi-formal representation of CPGs. To gain process information from a CPG, the first two steps are accomplished in order to have a representation independent of the final guideline language.

```

1 <li>
2   <a id="delta:8">In children with risk factors for Streptococcus pneumoniae,
      it is recommended that Amoxicillin, high dose (80 to 90 mg/kg/day) or
      Augmentin (with high dose amoxicillin component) be utilized as
      first-line therapy.
3   </a>
4   <ul type="disc">
5     <li>
6       <a id="delta:9">Note: Failure with amoxicillin is likely to be due to
          resistant Streptococcus pneumoniae, Haemophilus influenzae, or
          Moraxella catarrhalis.
7       </a>
8       <a id="delta:10">High dose amoxicillin will overcome Streptococcus
          pneumoniae resistance (changes in penicillin-binding proteins).
9       </a>
10      The clavulanic acid component of Augmentin is active
11      against resistant Haemophilus influenzae and Moraxella
12      catarrhalis (B-lactamase enzyme).
13    </li>
14  </ul>
15 </li>

```

Listing 1.

Excerpt of a source listing of the marked-up guideline document “Evidence-based clinical practice guideline for children with acute bacterial sinusitis in children 1–18 years of age”. Relevant sentences are enclosed by HTML-like “a” tags.

Figure 3. Detecting relevant sentences. We split this task into two modules: (1) segmentation and filtering module and (2) template generation module.

Figure 4. Finding processes and extracting required information. We split this task into five modules (i.e., the structure extraction module, the slot extraction module, the merging and grouping module, the process extraction module, and the template generation module).

```

1 <action id="8" parent="5" group="18" selection="0">
2   <delta-link link-id="8"/>
3   <description>In the child with no risk factors for penicillin-resistant
4     Streptococcus pneumoniae standard dose amoxicillin or Augmentin (with
5     standard dose Amoxicillin component) may be considered as initial
6     therapy.
7   </description>
8   <agents>
9     <agent MeSH="D000658" name="amoxicillin"/>
10    <agent MeSH="D019980" name="Augmentin"/>
11  </agents>
12  <condition>
13    <item>In the child with no risk factors for penicillin-resistant
14      Streptococcus pneumoniae
15    </item>
16  </condition>
17  <annotations>
18    <annotation>Note: Forty-six percent of isolates at Children's Hospital
19      Medical Center of Cincinnati, Ohio have intermediate or high
20      Penicillin-resistant Streptococcus pneumoniae and local data
21      supports that 15% of children locally may fail initial therapy with
22      standard dose amoxicillin.
23    <delta-link link-id="9"/>
24  </annotation>
25 </annotations>
26 <context>
27   <item>Antibiotic Treatment</item>
28 </context>
29 </action>

```

Listing 2.

Action instance of an ActionIR template for the guideline “Evidence-based clinical practice guideline for children with acute bacterial sinusitis in children 1–18 years of age”.

Table 1

Developing organizations of the guidelines used for development and testing

Organization	Guidelines
American Academy of Family Physicians	13
American Academy of Otolaryngology	13
American Academy of Pediatrics	13, 15
Centers for Disease Control and Prevention	1, 14
Cincinnati Children's Hospital Medical Center	9, 10
Finnish Medical Society Duodecim	17
Institute for Clinical Systems Improvement	2, 3, 4, 7, 8, 16
Practice Guidelines Initiative	18
Scottish Intercollegiate Guidelines Network	6, 11
University of Michigan Health Systems	5, 12

Table 2

Set of guidelines used for development of the rules

Guideline Number	Title
2	Acute pharyngitis
4	Acute sinusitis in adults
7	Diagnosis and treatment of obstructive sleep apnea
10	Evidence based clinical practice guideline for medical management of otitis media in children 2 months to 6 years of age
12	Otitis media
16	Rhinitis

Table 3

Evaluation results of the markup task

#	Title	SEN	POS	ACT	COR	REC	PRE
1	Acute otitis media: management and surveillance in an era of pneumococcal resistance	17	5	4	4	0.8	1
3	Acute rhinosinusitis in adults	44	10	10	9	0.9	0.9
5	Allergic rhinitis	75	16	20	15	0.938	0.75
6	Diagnosis and management of childhood otitis media in primary care	57	12	9	9	0.75	1
8	Diagnosis and treatment of otitis media in children	269	49	49	49	1	1
9	Evidence based clinical practice guideline for children with acute bacterial sinusitis in children 1–18 years of age	171	32	32	31	0.969	0.969
11	Management of sore throat and indications for tonsillectomy	96	20	18	18	0.9	1
13	Otitis media with effusion	87	7	5	5	0.714	1
14	Pneumococcal vaccination for cochlear implant candidates and recipients	20	4	3	3	0.75	1
15	Reduction of the influenza burden in children	50	6	7	6	1	0.857
17	Sore throat and tonsillitis	120	19	17	16	0.842	0.941
18	Symptomatic treatment of radiation-induced xerostomia in head and neck cancer patients	5	5	3	3	0.6	1
		1011	185	177	168	0.908	0.949

SEN – number of sentences of the guideline document.

POS – number of relevant sentences according to the key target template.

ACT – number of relevant sentences detected by the system.

COR – number of correctly detected relevant sentences detected by the system.

REC – ratio of COR to POS.

PRE – ratio of COR to ACT.

Table 4

Evaluation results of the process extraction task

#	Title	POS	ACT	COR	PAR	REC	PRE
1	Acute otitis media: management and surveillance in an era of pneumococcal resistance	13	15	11	0	0.846	0.733
3	Acute rhinosinusitis in adults	22	21	21	0	0.955	1
5	Allergic rhinitis	24	24	20	0	0.833	0.833
6	Diagnosis and management of childhood otitis media in primary care	27	24	19	0	0.704	0.792
8	Diagnosis and treatment of otitis media in children	107	105	97	2	0.916	0.933
9	Evidence based clinical practice guideline for children with acute bacterial sinusitis in children 1–18 years of age	56	60	50	1	0.902	0.842
11	Management of sore throat and indications for tonsillectomy	35	35	28	1	0.814	0.814
13	Otitis media with effusion	12	12	9	1	0.792	0.792
14	Pneumococcal vaccination for cochlear implant candidates and recipients	13	9	7	0	0.538	0.778
15	Reduction of the influenza burden in children	16	14	14	0	0.875	1
17	Sore throat and tonsillitis	38	33	27	0	0.711	0.818
18	Symptomatic treatment of radiation-induced xerostomia in head and neck cancer patients	8	7	6	0	0.75	0.857
		371	359	309	5	0.840	0.868

POS – number of slot fillers according to the key target template.

ACT – number of slot fillers generated by the system.

COR – number of correct slot fillers generated by the system.

PAR – number of partially correct slot fillers generated by the system.

REC – ratio of COR plus (0.5×) PAR slot fillers to POS slot fillers.

PRE – ratio of COR plus (0.5×) PAR slot fillers to ACT slot fillers.

Table 5

Evaluation results of the process extraction subtasks

	Recall	Precision
Agent	0.955	0.944
Dosage	0.846	0.846
Duration	0.542	0.813
Iteration	0.400	0.667
Condition	0.735	0.735
Relation	0.872	0.872
Sentence categorization and assignment	0.837	0.837