MEDINFO 2013
C.U. Lehmann et al. (Eds.)

doi:10.3233/978-1-61499-289-9-200

A Model-Driven Approach to Clinical Practice Guidelines Representation and Evaluation Using Standards

Ariel Farkash[a], John T.E. Timm[b], Zeev Waks[a]

[a]*IBM Haifa Research Lab, Haifa University Campus, Mount Carmel, Haifa, Israel*
[b]*IBM Almaden Research Center, San Jose, CA, United States*

Abstract

Clinical Practice Guidelines (CPGs) contain a set of schematic plans for the treatment and management of patients who have a particular clinical condition. CPGs are increasingly being used to support physician decision making. Many groups develop tools for the representation of CPGs. These differ in their approaches to addressing particular modeling challenges. Despite this strong effort, physicians still primarily rely on free-text narrative descriptions. Thus, a core challenge is to develop a formal representation of CPGs that physicians can easily read and verify, yet a machine can process, analyze and apply directly to a patient's EHR data. Our paper proposes a solution to this fundamental problem by describing an approach to CPG formalization using the Natural Rule Language (NRL), coupled with transformation to Object Constraint Language (OCL) constraints that are applied on a patient's clinical data record, in our case an HL7 Continuity of Care Document (CCD). We illustrate our approach on a simple guideline directive for Essential Hypertension.

Keywords:

Clinical Practice Guidelines, Natural Rule Language (NRL), Clinical Decision Support, UML, OCL, CDA, CCD, HL7.

Introduction

Clinical Practice Guidelines (CPGs) contain a set of schematic plans, at varying levels of abstraction and detail, for the treatment and management of patients who have a particular clinical condition. CPGs rely on the coupling of scientific evidence together with best clinical practices and are a result of a consensus among world experts. As a result, CPGs are increasingly being used to support physician decision-making. CPGs are promoted as a means to decrease inappropriate practice variation and to reduce medical errors. For these goals to be achieved, clinicians must adhere to the CPG recommendations in a consistent manner. However, most guidelines are in free text, not machine-processable, not easily accessible to clinicians at the point of care and too numerous for a clinician to apply to a patient's clinical record during the time of consultation.

A number of groups have developed formal representations for CPGs. Among the more well-known ones are: GuideLine Interchange Format (GLIF)[2], GUIDE [3], PRODIGY [4], and PRO*forma* [5]. Each representation language addresses a particular modeling challenge. A comparison study [6] reviewed six of the most common CPG formalizations and found consensus on many components, including plan organization, expression language, conceptual medical record model, medical concept model, and data abstractions. Lately, development of systems supporting individual clinical decisions are evolving toward the implementation of adaptable care pathways on the semantic web, incorporating formal, clinical, and organizational ontologies, and the use of workflow management systems [7]. A common element to all of these, however, is that the user of these formalizations needs to be proficient in the grammar, syntax and semantics of the representation language. A physician at the point of care is typically not the intended audience for such a language.

Our approach, as depicted in Figure 1, aims to bridge this gap by enabling a physician to read the same representation of the guidelines that a machine can process and apply to patient clinical records. This capability involves a design-time step where the free-text CPG document is formalized into English-like Natural Rule Language (NRL) [9] rules. The model upon which the NRL rules operate is a widely accepted standard for clinical data representation, in our case the HL7 Continuity of Care Document (CCD) [10]. At this stage, a physician can easily read and verify (or reject) these English-like rules.

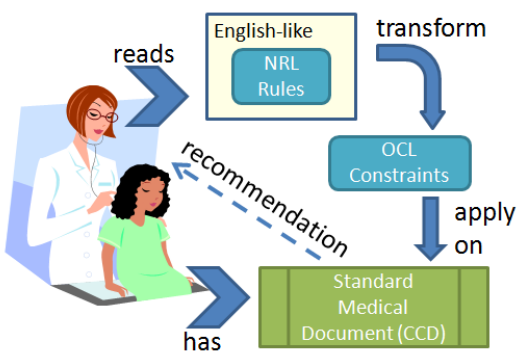

Figure 1- Approach overview

The example we will use throughout paper is a simplified directive taken from the CPG on management of primary hypertension in adults [11] which formalized into to NRL reads: '*If the patient is an adult and hypertensive then add an antihypertensive medication to the care plan in the ClinicalDocument*'. The NRL rules are automatically transformed to Object Constraint Language (OCL) [12] constraints at runtime and applied directly on the patient's record. A system that implements this approach would provide capabilities that allow a physician to export the patient's record in a standard format, run an analysis

corresponding to an approved set of guidelines and receive patient-specific recommendations to consider at the point of care.

In our prototype, presented in the Methods and Results sections, we implement a component of the above approach that demonstrates how guidelines can be authored using NRL, translated into OCL and evaluated against a patient's record in CCD format. The results of our prototype are very promising and we have plans to expand the prototype into a service-oriented infrastructure that will handle guidelines from the oncology domain as well as apply the approach to the specification of Clinical Trial Eligibility criteria. These topics will be further detailed in the Future Work section.

Background and Related Work

Object Constraint Language

The Object Constraint Language (OCL) is an OMG standard for specifying constraints on a Unified Modeling Language (UML) model. OCL is a side-effect-free declarative language that can be used for pre-conditions, post-conditions, class invariants and operation bodies. HL7 has created a specification that targets a subset of OCL called GELLO [13] intended for use in clinical decision support applications.

Model-Driven Health Tools

Model-Driven Health Tools (MDHT) [14, 15] is an open source project that provides a suite of tools for developers and implementers of healthcare IT standards specifications. The project was started in 2008 by the US Veteran's Health Administration (VHA) and IBM, and the initial focus was on creating a software standards-based methodology for developing HL7v3 specifications using UML. Subsequently, the project quickly evolved to support the specification of HL7 Clinical Document Architecture (CDA) Implementation Guides. The project is currently supported by the VHA, IBM and the US Office of the National Coordinator (ONC).

In addition to design-time tooling, the project has developed models and reference implementations for several commonly used CDA implementation guides including the HL7 Continuity of Care Document (CCD), HITSP C32 Patient Summary and most recently Consolidated CDA. MDHT uses UML and OCL formalisms for specifying models that capture constraints on top of the base CDA standard. These models are transformed into Java APIs that can be used to produce, consume and validate CDA documents against the constraints. Documentation in various formats (e.g. PDF, XHTML and Eclipse Help) can also be generated. MDHT future project plans include generalizing the tooling to support standards other than CDA, including the specification of Archetypes using UML per the OMG Archetype Modeling Language Profile proposal.

Natural Rule Language

The Natural Rule Language (NRL) is a language for specifying data model constraints with near-English syntax. The language was designed as an alternative to languages such as OCL and Schematron and it can be used to write constraints for models defined using UML or XML schema. One of the primary goals of NRL is to provide readability for non-developers in ways that more technical languages cannot. The NRL abstract syntax tree can be transformed to any number of different implementation languages. In this paper, we are proposing the use of NRL for specifying clinical guidelines such that they can be authored and verified by healthcare domain experts (e.g. physicians).

NRL has other attractive features that can aid us in our endeavor of using it for CPG representation and evaluation. NRL has two types of rules: Validation Rules and Action Rules. Validation Rules are constraints that, when evaluated over data, return a Boolean result. Action Rules are statements that perform operations on the data or create new data. Validation Rules and Action Rules can refer to Validation Fragments and Action Fragments, respectively. Fragments are reusable rule parts that can be combined together to form more complex statements. Fragment names can be written as English phrases that make sense in the context of larger sentences. Variables, which can be used to simplify the process of writing complex expressions, can be declared and used within a Validation or Action Rule. User-defined external operators allow external functionality to be accessed from within a rule or fragment. The operators are bound at runtime to software components that perform specific tasks.

CPGs and Controlled Natural Language

Prior research in this area focuses on guideline readability, quality and the use of imprecise language. The research identifies patterns for representing guidelines using a controlled natural language, thus making them easier to implement. A prototype was developed to formalize pediatric guideline recommendations using Attempto Controlled English [8]. The work does not, however, reference a standard clinical data model or evaluate the guidelines against patient data using such a model.

Methods

Our prototype system uses the NRL and OCL languages in concert with the MDHT tooling platform to enable the parsing of structured English clinical guideline statements into machine rules that can execute over patient record data in HL7 CDA format. Specifically, we leverage an open source parser for NRL that can parse a rule file into an Abstract Syntax Tree (AST). After parsing, the parser resolves model references using a collection of UML2 models or XML schemas. The model references are used to decorate the AST, making it complete for translation to an implementation language (e.g. Java or OCL) or evaluation by an interpreter. The runtime libraries produced by MDHT are based on the Eclipse Modeling Framework (EMF). As the NRL parser did not have native support for directly loading EMF models, we developed an *EMFModelLoader* that accomplishes this task, thus expanding the NRL parser's capabilities to accommodate the existing CDA implementation models from MDHT.

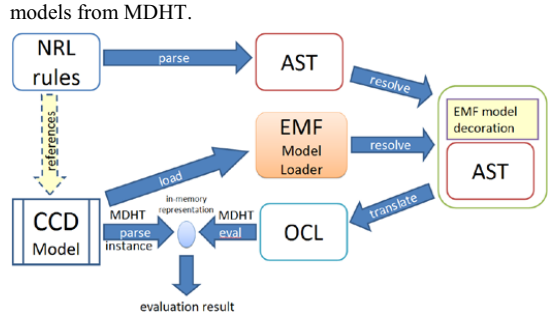

Figure 2- Prototype Architecture

In Figure 2, we illustrate the detailed architecture of our prototype implementation. The first step in the process is to extract the constraint from the clinical guideline and translate the NRL form of the constraint into OCL.

The EMF model loader loads existing MDHT implementation models for CCD into memory (1). The NRL parser creates an AST from the rule file (2). Then the parser decorates the AST with references to the MDHT CCD models (3). A program analyzes the decorated NRL AST and extracts the constraint from the "*if* **constraint** *then* **action**" form of the clinical guideline. The constraint part of the AST is translated to OCL (4). MDHT Java APIs are used to parse a CCD document into Java objects (5). The OCL statements generated in the first step are evaluated on the in-memory representation of the CCD document (6). If the evaluation succeeds, this indicates that the constraint has been met and we can invoke the action. In our prototype, we use the external operator feature of NRL as a means for implementing actions. Each external operator is bound to a Java method. At design-time, we create a Java method that accepts a clinical document as the sole input parameter. The method is implemented using MDHT Java APIs for CCD and includes all of the program code necessary to provide the runtime functionality required by the action.

One of the challenges in formulating NRL for a standard like CDA is that the underlying information model is quite complex. It contains many complex data types and recursion that can often produce deeply-nested instances. This structure leads to complicated query expressions to extract even the simplest clinical data elements. To combat this in our prototype, we used various features of the NRL language including: validation fragments, action fragments and user-defined external operators to create a set of reusable abstractions that allowed us to formulate complex statements in a manner that produced concise readable guidelines that can be checked by a clinician.

Though our prototype does not currently include a UI component, we envision a guideline-authoring environment that includes a set of commonly-used clinical concepts taken from the CCD specification. For example, we could have queries for vital signs, problems, allergies and lab results. Additionally, such an environment would need access to a terminology browser to quickly find and include terms from SNOMED, LOINC and RxNorm.

Results

Essential Hypertension Guidelines

Essential Hypertension (EH) is the form of hypertension that, by definition, has no identifiable cause. It is the most common type of Hypertension, affecting 95% of hypertensive patients. EH is a complex disease that is the consequence of an interaction between clinical, environmental and genetic factors. We did not intend to cover the entire CPG for EH. In this paper we took a single, very simple, directive from an EH CPG [10] where the CPG recommends that adult Hypertensive patients receive treatment using anti-hypertensive drug therapy.

Formalization using NRL

We formalized the above referenced directive into NRL: "*If the patient is an adult and hypertensive then add an antihypertensive medication to the care plan in the ClinicalDocument*". The above rule is composed of an "*if* **constraint** *then* **action**" as is shown in Figure 3. Both constraint and action are complex statements for a machine to process, but are easily understood by native English speakers and specifically physicians.

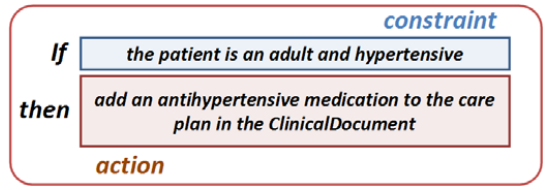

Figure 3- NRL Constraint and Action

The constraint contains two keywords, *adult* and *hypertensive*, that must be formalized for a machine to process. Figure 4 shows the corresponding NRL required to represent the constraint. Adult is defined as an individual being 18 years or older, while a hypertensive patient is one that has at least one blood pressure measurement where the systolic blood pressure is greater than 140 and the diastolic blood pressure is greater than 90.

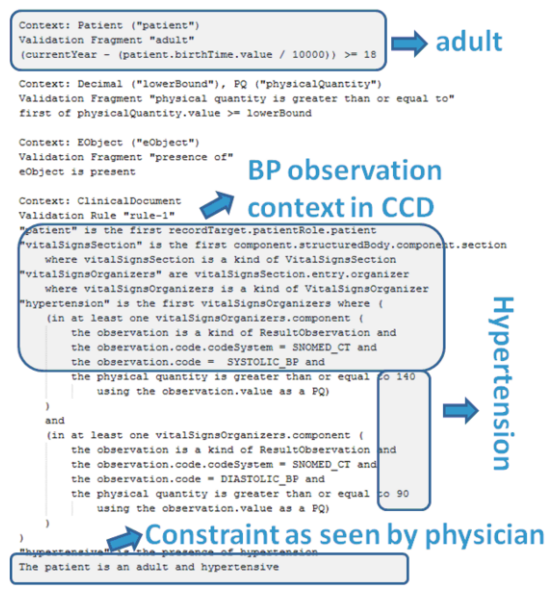

Figure 4- NRL Constraint

The NRL action follows the same logic: it activates the drug administration by the Java grounding mechanism described in Methods, but this is hidden from the physician.

Transformed to OCL

The corresponding constraint in OCL is depicted in Figure 5.

The similarity is obvious. The *adult* and *hypertensive* constraints are defined explicitly and in context of the data items as located in the CCD document. Then the constraint "*adult and hypertensive*" is applied on the patient record and a Boolean result is returned, this can be seen in Figure 6.

In this example, an "if constraint" evaluation with a result of 'true' leads to the invocation of an externally-defined Java method "*add an antihypertensive medication to the care plan in*" with the ClinicalDocument as a parameter. This method creates a new medication entry to be entered as a recommenda-

tion in the care plan section of the continuity of care document and is immediately visible to the physician.

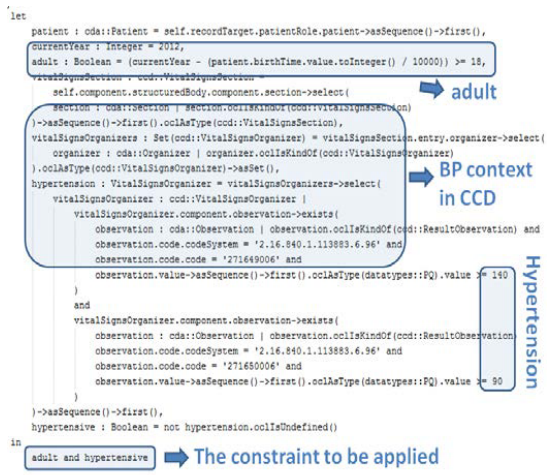

Figure 5- OCL Constraint

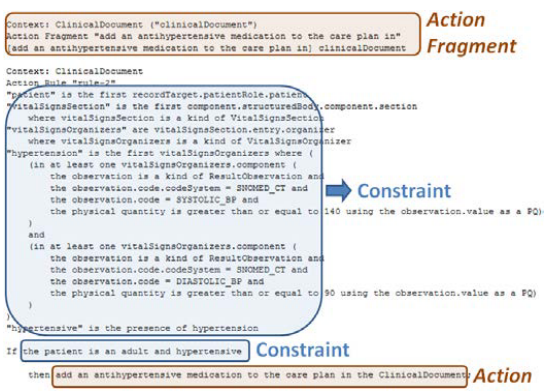

Figure 6- NRL Constraint with Action

Discussion

Thus far, the primary output of this research is a set of software components that support the translation and interpretation of CPGs represented using NRL. An important prerequisite of our prototype is that all of the clinical data required can be exported from an EHR system as an HL7 Continuity of Care document. The hypertension guideline cited in the Results section is a very simple example used to demonstrate what is possible. In practice, such a guideline would represent a small fragment of a much larger, more complex guideline. In the Future Work section, we mention the application of our approach to complex oncology guidelines.

Future Work

In this section, we will briefly describe future directions and the application of our approach to various use cases.

Oncology Use Case

We are currently engaged in a large project involving clinical decision support for Soft Tissue Sarcoma with an Italian hospital [17, 18]. Using the approach presented, we will formalize the processes involved in arriving at a patient's clinical presentation (i.e. diagnosis support) and provide CPG-based treatment recommendation.

Clinical Trial Eligibility Use Case

An area where one can leverage our approach is in the field of clinical trials recruitment. Recruitment of patients for a given clinical trial is currently one of the major hurdles faced by trial administrators. Clinical trials have very specific inclusion and exclusion criteria. Take for example the actively recruiting clinical trial from ClinicalTrials.gov [16] on "Renal Sympathetic Modification in Patients with Essential Hypertension". Figure 7 outlines the inclusion and exclusion criteria for participation in this clinical trial.

Inclusion Criteria
>= 18 years old, and <= 75 years old
more than half a year history of hypertension
recently at least three times office blood pressure suggest systolic blood pressure of 140 mmHg
estimated glomerular filtration rate (eGFR) of >= 45 mL/min
Exclusion Criteria
secondary hypertension
estimate glomerular filtration rate (eGFR) < 45 mL/min
has history of renal restenosis or renal stents implantation
has experienced AMI (old myocardial infarction is not excluded)
patients with sick sinus syndrome
pregnant women
mental disorders
patients that have allergy to contrast agents

Figure 7- ClinicalTrials.gov Eligibility Criteria

Formalizing these rules into NRL and automatically applying them to patient standards-based EHR records will significantly raise the potential for actively recruiting patients.

Service-Oriented Framework

The model-driven approach to clinical guidelines formalization presented in this paper is part of a larger framework that includes a clinical guideline evaluation service shown in Figure 8.

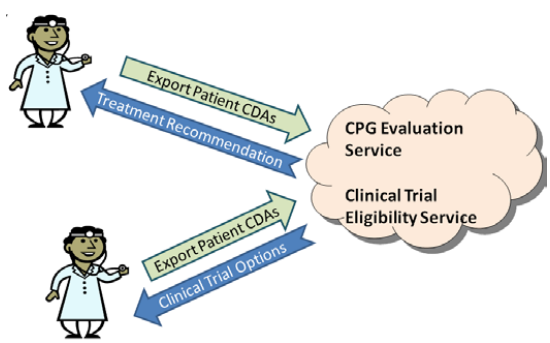

Figure 8- Service-Oriented Framework

Interested parties export CDA documents from their systems and send them to the service on the cloud. One or more clinical guidelines are evaluated against the document and the result is returned to the requesting system, including a new version of the clinical document with entries added to the care plan sec-

tion. The new CDA document can be imported into the EHR system and then evaluated by the physician to see what recommendations were made.

A similar procedure can be followed by a physician of a patient that is looking to enroll in a Clinical Trial. The physician can export a patient's CDA documents through his clinic/hospital system and send them to the service on the cloud. Clinical Trials corresponding to the search terms are evaluated versus the patient's records using each trial's inclusion and exclusion criteria. Matching Clinical Trials are returned to the physician online enabling the physician to review and recommend proper Clinical Trial enrollment.

Conclusion

In this paper, we have shown a proof-of-concept implementation of a method for formal representation of CPGs that physicians can easily read and verify, yet a machine can process, analyze and apply directly to a patient's EHR data. We addressed this problem by leveraging the NRL modeling language for CPG formalization, coupled with transformation to OCL constraints that are applied on a patient's clinical data record, e.g. the Continuity of Care Document. We illustrated our approach in the Methods and Results sections on a simple Essential Hypertension guideline directive. In the Future Work section we described future directions including the application of our approach to several different use cases.

References

[1] Shea S, DuMouchel W, Bahamonde L. A meta-analysis of 16 randomized controlled trials to evaluate computer-based clinical reminder systems for preventive care in the ambulatory setting. J Am Med Inform Assoc. 1996 Nov-Dec; 3(6); pp. 399-409.

[2] Ohno-Machado L et al. The guideline interchange format: a model for representing guidelines. J Am Med Inform Assoc. 1998 Jul-Aug;5(4):357-72.

[3] Quaglini S, Stefanelli M, Lanzola G, Caporusso V, Panzarasa S. Flexible guideline-based patient careflow systems. Artif Intell Med. 2001 Apr; 22(1):65-80.

[4] PRODIGY: implementing clinical guidance using computers. Br J Gen Pract. 1998 September; 48(434): 1552–1553.

[5] Sutton DR, Fox J. The Syntax and Semantics of the PRO*forma* Guideline Modeling Language. J Am Med Inform Assoc 2003;10:433-443 doi:10.1197/jamia.M1264

[6] Peleg M et al. Comparing Computer-interpretable Guideline Models: A Case-study Approach. J Am Med Inform Assoc. 2003 Jan-Feb; 10(1): 52–68.

[7] Gooch P, Roudsari A. Computerization of workflows, guidelines, and care pathways: a review of implementation challenges for process-oriented health information systems. J Am Med Inform Assoc 2011;18:738-748 doi:10.1136/amiajnl-2010-000033

[8] Shiffman R, Michel G, Krauthammer M, Fuchs NE, Kaljurand K, and Kuhn T. Writing Clinical Practice Guidelines in Controlled Natural Language. in: N.E. Fuchs (ed.), Workshop on Controlled Natural Languages, CNL 2009, LNCS/LNAI 5972, Springer, pp. 265-280, 2010.

[9] Natural Rule Language (NRL) Specification 1.4.0, http://nrl.sourceforge.net/spec/

[10] HL7 Continuity of Care Document (CCD) Implementation Guide for the Clinical Document Architecture Specification, http://www.hl7.org/implement/standards/product_brief.cfm?product_id=6

[11] National Clinic Guideline Centre, The clinical management of primary hypertension in adults http://www.nice.org.uk/nicemedia/live/13561/56007/56007.pdf

[12] Object Constraint Language (OCL) Specification, http://www.omg.org/spec/OCL/

[13] GELLO Expression Language, http://www.hl7.org/implement/standards/product_brief.cfm?product_id=5

[14] Model-Driven Health Tools (MDHT), https://mdht.projects.openhealthtools.org

[15] Model-Driven Standards Development: Lowering the Bar for Implementers, Workshop, MIE 2012, August, Pisa, Italy

[16] Clinical Trial. Renal Sympathetic Modification in Patients with Essential Hypertension. ID: NCT01417221. http://clinicaltrials.gov/ct2/show/NCT01417221?term=essential+hypertension&rank=2

[17] Carmeli B et al. Evicase: An Evidence-Based Case Structuring Approach for Personalized Healthcare. Stud Health Technol Inform. 2012; 180:604-8.

[18] Waks Z et al. Analyzing the "CareGap": Assessing Gaps in Adherence to Clinical Guidelines in Adult Soft Tissue Sarcoma, European Federation for Medical Informatics 2013, Prague Czech Republic, http://www.stc2013.org/

Address for correspondence

Ariel Farkash: arielf@il.ibm.com