

A multi-layered framework for disseminating knowledge for computer-based decision support

Aziz A Boxwala,¹ Beatriz H Rocha,^{2,3} Saverio Maviglia,^{2,3} Vipul Kashyap,⁴ Seth Meltzer,² Jihoon Kim,¹ Ruslana Tsurikova,² Adam Wright,^{2,3} Marilyn D Paterno,^{2,3} Amanda Fairbanks,² Blackford Middleton^{2,3}

► Additional material is published online only. To view this file please visit the journal online (www.jamia.org/content/18/Suppl_1.toc).

¹Division of Biomedical Informatics, Department of Medicine, University of California San Diego, La Jolla, California, USA

²Clinical Informatics Research and Development, Brigham and Women's Hospital, Boston, Massachusetts, USA

³Department of Medicine, Harvard Medical School, Boston, Massachusetts, USA

⁴Cigna Healthcare, Hartford, Connecticut, USA

Correspondence to

Aziz A Boxwala, Division of Biomedical Informatics, UCSD School of Medicine, 9500 Gilman Drive 0728, La Jolla, CA 92093-0728, USA; aboxwala@ucsd.edu

Received 26 April 2011

Accepted 27 September 2011

Published Online First

3 November 2011

ABSTRACT

Background There are several challenges in encoding guideline knowledge in a form that is portable to different clinical sites, including the heterogeneity of clinical decision support (CDS) tools, of patient data representations, and of workflows.

Methods We have developed a multi-layered knowledge representation framework for structuring guideline recommendations for implementation in a variety of CDS contexts. In this framework, guideline recommendations are increasingly structured through four layers, successively transforming a narrative text recommendation into input for a CDS system. We have used this framework to implement rules for a CDS service based on three guidelines. We also conducted a preliminary evaluation, where we asked CDS experts at four institutions to rate the implementability of six recommendations from the three guidelines.

Conclusion The experience in using the framework and the preliminary evaluation indicate that this approach has promise in creating structured knowledge, to implement in CDS systems, that is usable across organizations.

INTRODUCTION

Clinical practice guidelines are systematically developed statements to assist practitioner and patient decision-making about appropriate health-care for specific clinical circumstances.¹ However, in spite of the substantial effort that goes into the development and dissemination of guidelines, evidence suggests that they have not been effective in improving the overall quality of care that is delivered.^{2–5} On the other hand, a different set of evidence suggests that computer-based clinical decision support (CDS) systems, when effectively implemented, can improve patient care and increase adherence to guidelines for prevention and treatment.^{4–8} Despite the evidence of its effectiveness, current use and adoption of CDS is limited.⁹

Among various reasons, wider adoption of CDS has been inhibited by the difficulty of translating clinical practice guidelines into a computable form that can be used in CDS systems. For an apparently simple CDS reminder, this involves specifying the triggering events (eg, opening the electronic record for the patient by a primary care physician), the logical expression and the underlying data elements for the patients to whom that reminder applies (eg, female patients above 40 years of age who have not had a screening

mammography performed in the last year), exclusions (eg, patients who have had a bilateral mastectomy), and the recommended action (a screening mammogram order). In many cases, these elements are not specified with computable precision, and in some cases, are not specified at all in the guideline document.¹⁰ A successful CDS intervention requires a thoughtful knowledge engineering effort to specify these elements so that the recommendation is useful and actionable, and targeted to the right person at the right time and with the right presentation.

Today, this knowledge engineering effort must be replicated at every organization that wishes to set up such a CDS reminder, creating a substantial barrier to the use of CDS for implementing guidelines. As a consequence, CDS implementation is largely confined to large academic medical centers that have the resources to translate guidelines into effective CDS.⁹ Disseminating guideline knowledge in a form that is more amenable to CDS is likely to broaden the use of CDS and decrease time lags in implementing evidence-based recommendations by reducing the time and effort required for knowledge engineering.

Encoding guideline knowledge into a computer-interpretable form that is then usable at different healthcare organizations is challenging.¹¹ One set of barriers is the heterogeneity of CDS modalities (eg, reminders, order sets, infobuttons), of software that implement these CDS modalities, of data representations within electronic medical records, of local workflows, and of organizational resources and governance structures. An overly simple approach or a prescriptive approach to representing recommendations for CDS could make it too generic for any specific setting, and therefore not very useful or usable overall. On the other hand, trying to accommodate even the broad variety of contexts in which the guideline will be implemented as CDS could make the knowledge very complex.

We have developed a framework to deal with such variability by incrementally structuring guideline knowledge through several layers of translation that aim to distinguish the knowledge in the recommendation from the details required to implement the knowledge as CDS. The knowledge representation models incorporated within the framework can be used to create more structured and precise specifications of recommendations. These specifications can serve as the basis of CDS implemented within a healthcare organization.

BACKGROUND

Definition of CDS

We use the definition of Osheroff¹² where CDS is a system

providing clinicians or patients with clinical knowledge and patient-related information, intelligently filtered or presented at appropriate times, to enhance patient care.

This is a broader definition of CDS than a common usage of the term that refers to reminders and alerts driven by rules. Thus, our multilayered framework for knowledge dissemination must support solicited and unsolicited CDS using a variety of tools such as reminders, alerts, documentation forms, context-specific data summaries, order sets, care pathway tools, and reference information.

Review of CDS representation models

A number of projects, either as part of standards development efforts or informatics research, aim to create knowledge representation schemes for sharing knowledge for use in these different types of CDS tools.

Perhaps the most widely adopted knowledge representation format in the USA is the Arden Syntax for Medical Logic Modules.¹³ Arden Syntax is a representation for event–condition–action type rules, referred to as medical logic modules or MLMs. The MLMs can be used for creating reminders and alerts. Even though Arden Syntax MLMs can be implemented at systems in a large number of healthcare organizations in the USA, sharing of MLMs among organizations has not been widespread. A factor that has been considered to be a cause of limited sharing of MLMs across organizations is the so-called ‘curly braces problem’.¹⁴ This refers to the part of the MLM where organization-specific code, wrapped in a pair of curly braces, is inserted into the MLM in order to map MLM variables to the data in the organization’s clinical database. Such an approach renders the MLM non-portable. Work is ongoing in HL7 to address the curly braces problem by having an MLM address clinical data from a standard abstraction layer called the Virtual Medical Record.¹⁵

Several research projects also have attempted to build languages for representing the complex knowledge contained in guidelines. A common factor in these languages has been to represent the flow of logic and recommendations in a temporally-oriented network of tasks. Among these formalisms are Asbru,¹⁶ EON,¹⁷ Guideline Interchange Format (GLIF),¹⁸ Proforma,¹⁹ and SAGE.²⁰ The execution semantics of these languages are more complex than that of the MLMs or other forms of CDS described next, and require their own execution engines. To date, little research has been done to understand how such execution fits with the myriad of workflows in healthcare organizations and how it might integrate with clinical applications such as electronic medical record (EMR) systems, computerized provider order entry systems, and existing CDS tools. This might be one reason why there has been limited usage of these formalisms outside of research projects.

The Clinical Decision Support Technical Committee at HL7 also is creating standard specifications for other CDS modalities such as order set templates, context aware reference information retrieval (also known as infobuttons), and a CDS service. An order set is a collection of clinical orders for a particular patient context such as an inpatient admission for congestive heart failure. The HL7 order set specification is intended to allow organizations to distribute and share order set templates in a standard format. The infobutton standard will enable clinical applications such as an electronic medical record system to

provide links to reference information in a library using a standard interface.²¹ The specification for the CDS service standard defines a set of transactions that can be executed between a clinical application such as an EMR system or a computerized provider order entry system and a decision support software service.²² Notably, this specification does not create a representation for the decision support logic since that is abstracted by the service.

The knowledge representations discussed above should be able to be utilized to share encoded guideline representations. These multiple representation formats create a dilemma for those wishing to share guidelines they have developed. The representation to be used for implementing a particular recommendation depends on a number of contextual factors including the availability and functionality of the CDS tools at the organization and the clinical workflow within which the CDS is to be provided. Thus, in order to share encoded guidelines, either many possible representations must be produced to allow wider implementation of the guidelines, or if a single representation is used, only a limited number of settings will be supported.

Layered knowledge representations

One approach that can allow sharing of knowledge for use in heterogeneous CDS system is to use a layered framework to representing the knowledge. In such an approach, knowledge is structured progressively and made more specific in successive layers. In the bottom layers, knowledge is expressed in a generic manner, thus allowing flexibility for transformation for wider use. In the upper layers, the knowledge is transformed to be more specific for use in more narrowly defined contexts.

The hybrid guideline representation model in the Digital Guidelines Library (DeGeL) framework²³ utilizes such a layered approach. The DeGeL framework ‘facilitates gradual conversion of free text guidelines to a formal machine readable language’. The conversion transforms a text guideline to semi-structured text, then to a semi-formal representation, and finally a formal representation. The semi-formal representation and the formal representation typically include structures that are specific to different knowledge representation ontologies or languages such as GLIF and Asbru. The hybrid ontology models guidelines as plans of actions that are arranged in temporal flows. The authors also have created a set of tools to create and manage the knowledge,²⁴ including tools to mark up text guidelines into more structured forms.

The Guideline Elements Model (GEM) also is ‘intended to facilitate translation of natural language guideline documents into a format that can be processed by computers’.²⁵ GEM is an Extensible Markup Language (XML)-based guideline document model; that is, GEM aims to represent in XML, the content of a text guideline. Although the layers of knowledge are not explicitly separated in the model, it does allow modeling knowledge at various levels of abstraction. Of particular relevance to this paper, the knowledge components can be modeled at a high level as recommendations and at a lower level using an algorithm that derives from the GLIF representation.

In the next section, we describe a multilayered framework for guideline representation that draws from these models. Unlike the above approaches, our framework emphasizes modeling of unsequenced clinical decisions, rather than activities that are organized explicitly into flow sequences. As in DeGeL, we explicitly separate the layers of knowledge to facilitate use of knowledge management methodologies and tools.

METHODS

A multilayered framework for representing clinical decisions

Our design objective in creating the knowledge representation framework was to enable implementation of guideline recommendations as CDS in a wide variety of settings. We believe that leveraging CDS technologies existing in an organization would enable wider adoption of the guidelines. In other words, we did not want to require the use of specialized execution engines or CDS tools. Rather, we wanted to provide guideline knowledge in a structured form that an organization could adapt easily to their workflows and resources and implement in their CDS systems.

The design objective led us to emphasize the representation of clinical decisions. This structure is unlike Asbru, GLIF, and other guideline representations that model the knowledge as a process or plan with flows of decision logic and actions (figure 1). First, such modeling is similar to and overlaps with the representation of clinical workflows. It has been our experience that the decision logic varies much less than the workflow across different clinical settings. Thus, representations of clinical decisions are more likely to be sharable across organizations. Second, decision support tools, such as rules, alerts, and order sets commonly used at healthcare organizations,^{26 27} largely do not provide explicit support for modeling flows or sequences of activities. Thus, it would not be straightforward to implement sequenced activities in existing CDS tools²⁸ and would require use of specialized execution engines.^{29–31}

Furthermore, we believe that a multi-layered knowledge representation would allow us to meet efficiently our design objective: to create and disseminate knowledge for use in a variety of settings. Our framework consists of four layers of knowledge. Successive layers in this framework progressively structure the knowledge. At the bottom is the unstructured

guideline layer that contains the narrative guideline. At the top is the executable layer that contains knowledge in a form that is directly implemented in a specified CDS system. The intermediate layers format the guideline’s recommendations as structured specifications for those implementing CDS.

The goal in creating a multilayered approach is to provide a balance between the competing requirements for flexibility in representation for various organizational environments and the ability to deliver precise, executable knowledge that can be more readily implemented. For those who can use an available executable knowledge artifact, say an Arden Syntax MLM supported in their specific environment, this approach allows for incorporation of that executable artifact. For the institution, this can reduce the effort required to translate the narrative guideline into executable knowledge. Others, who might not have such capability in their CDS system to execute MLMs, can use an artifact from an earlier layer to create their own executable knowledge appropriate for their environment. Even these latter users should benefit from the availability of knowledge that has been more structured and codified than the underlying narrative guideline. By removing or reducing ambiguities in the knowledge, by providing standard terminology codes for data and actions, and by synthesizing (or encapsulating) the knowledge from potentially multiple sections of the narrative guidelines, we believe the intermediate layers can help reduce the effort at an institution required to create executable knowledge. Prior experience suggests that users will need to access knowledge at various levels of abstraction.³²

The four layers of this framework are described in the following subsections and summarized in table 1. Recommendations represented in the four layers (eg, for the management of diabetes mellitus) can be viewed at the CDS Consortium’s³³ Knowledge Management Portal.³⁴

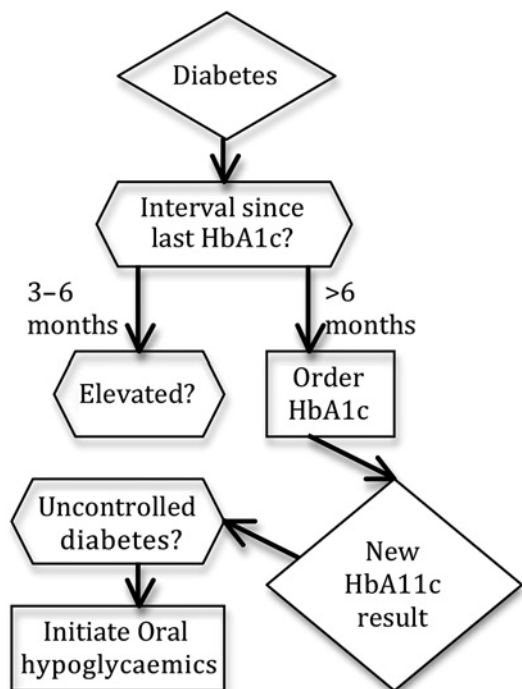


Figure 1 A simplified representation of a diabetes guideline as a flowchart in GLIF. As can be seen in the figure, GLIF represents guidelines as a flow of activities. In comparison, the representation we describe in this paper represents guidelines as a collection of decisions (as shown in figure 3), with no explicit flow or sequencing among the decisions.

Unstructured

This layer represents the narrative or textual guideline documents that are published by the guideline developers. These documents are authored by committees that contain a variety of experts in, for instance, the clinical topic and epidemiology. The documents in this layer serve as statements of policy and as a synthesis of evidence and expert opinions on the underlying clinical topic. These documents are used for a variety of purposes, one of which is as an input for creating computable knowledge for use in CDS systems. The framework does not define a new model or format for these documents and allows for guidelines published in any format.

Semi-structured

In this layer, we begin to add structure to the guidelines. The core organizing concept for the knowledge is a recommendation, for which we have created a representation model. Each recommendation is modeled as a decision about the interventions that are possible in a specified clinical scenario. Since, from this layer onwards, the emphasis is on CDS, the recommendations would exclude statements that are not patient-specific, such as population-oriented recommendations or policy statements that often are a part of unstructured guidelines. The objective of this layer is to serve as a vehicle for communication between clinical domain experts and knowledge engineers. The domain experts are expected to be the primary authors of the knowledge in this layer. Thus, the knowledge representation structure in this layer focuses on organizing the textual content of the recommendation, leaving codification of the knowledge for subsequent layers.

Table 1 The four layers in the knowledge representation framework

	Narrative	Semi-structured	Structured	Executable
Format	Narrative text	Organized text	Coded and interpretable by computer	Coded and interpretable by CDS systems; variety of formats
Sharability of knowledge	Broad	Broad	Broad	Very limited
CDS modality and tool independent	Yes	Yes	Yes	No
Site independent	Yes	Yes	Yes	No
Author	Guideline developer	Clinical domain expert	Knowledge engineer	CDS implementer
Purpose	Communication of policy; synthesis of evidence	Recommendations for implementation in CDS	Precise communication; validation	Implementation for a particular site

CDS, clinical decision support.

Structured

In this layer, the knowledge is specified with sufficient structure so as to make it computable and precise. We have created a representation model in which the guideline knowledge is organized into structured recommendations. Similar to the semi-structured recommendation, the structured recommendation is modeled as a decision in a specified clinical scenario. The knowledge in this layer is independent of the implementation in a particular type of CDS tool or of the workflow in a particular clinical setting. As such, it is not intended to be used directly within a CDS system, but formally defines all the data elements and logic required to do so. The objective of this layer is to communicate the knowledge in the guidelines from knowledge engineers to CDS implementers. A knowledge engineer is an expert in the construction of knowledge bases and in translating domain requirements for computer implementation. In our case, the knowledge engineer has expertise in CDS, and in health information models and terminologies.

Executable

In this layer, the knowledge is structured for use within a specific type of CDS tool (eg, reminder using Arden Syntax MLM or using a commercial rules engine such as iLog (IBM, Armonk, New York, USA)), within a specific clinical information system, at a particular clinical site (eg, Brigham and Women’s Hospital). The objective of this layer is the implementation of the knowledge within a specified setting and workflow. Within any setting, a team comprising clinicians, programmers, analysts, knowledge engineers, and others will be responsible for the implementation. Knowledge in this layer is less likely to be sharable since often it includes elements that apply only to that setting, for example, local codes for data items, details of local clinical services, and idiosyncrasies of how the end user interacts with the system. The framework does not prescribe or recommend any implementation or execution framework. Since the knowledge will be represented in the format native to the CDS tools in which it is being implemented, the framework does not include a representation model in this layer.

Model for recommendations

Based on the above knowledge representation framework, we created class diagrams in the unified modeling language and equivalent XML schemas (definitions of structured recommendation, semi-structured recommendation, and the metadata model are available as an online supplement) for representing guideline recommendations in the semi-structured and the structured recommendation layers. As mentioned in the previous section, we did not create a knowledge representation model for the narrative layer since that is comprised of the free-form published guideline document. Similarly, there was no need

to create a model for the executable layer, since that is comprised of the various executable formats such as Arden Syntax.

In the semi-structured recommendation layer, a guideline consists of a collection of recommendations (figure 2). These recommendations are organized into modules. A recommendation models a decision, and comprises a clinical scenario and a clinical action (figure 3, top). The clinical scenario describes the patient context. The action is the recommended intervention in this scenario. Definitions encapsulate reusable components of a scenario.

The high-level representation of the structured recommendation is conceptually similar to that of the semi-structured recommendation. However, the structured recommendation includes a more detailed representation of scenarios and actions (figure 3, bottom). Scenarios can be defined as logical expressions, where criteria are written in GELLO.³⁵ Actions can be defined to include decision alternatives and factors that influence the selection of an alternative. The object model underlying the patient data in logical expressions is adapted from the Clinical Statements model from the Health Information Technology Standards Panel’s Summary Document (C32) specification.³⁶ The object model for actions is informed by HL7’s information models and uses HL7 data types. The objects can reference standard terminology codes. Our choice of the use of GELLO and the C32-based information model is due to these being standards, albeit very new and still growing in use.

The model supports reusability of knowledge at the semi-structured and structured layers. In the former layer, ‘Definition’ can be used to specify the meaning of a term that can be used across recommendations and guidelines. For example, a definition could specify what is meant by ‘poorly controlled diabetes’ in terms of serum hemoglobin A1c test results. In the structured layer, recommendations that assert ‘patient state’ achieve

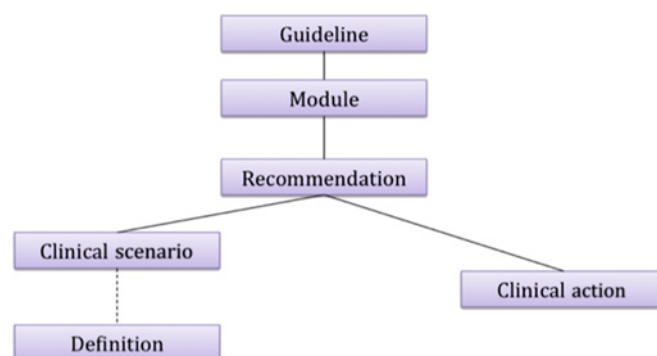


Figure 2 The top elements of the knowledge representation for the semi-structured recommendation.

Figure 3 Screenshot illustrating a semi-structured recommendation (top) and a structured recommendation (bottom) from a diabetes mellitus guideline.

Recommendation 3.1: HighA1c.StartOralAgentOrInsulin
Description: Start oral agent or insulin on poorly controlled diabetic patients not receiving insulin

Scenario: *Poorly controlled diabetic not receiving insulin*
Refers to the following definition: Poorly controlled diabetes = Most recent HgbA1c in the last year > 7%

Recommended Action(s):
Choice: Start insulin or new oral hypoglycaemic medication

or	Start new oral hypoglycaemic medication
or	Start insulin and educate patient about insulin

Recommendation 3.1: Start oral hypoglycaemic agent or insulin (Untreated poorly controlled non-Type 1 diabetes)
Description: Start an oral hypoglycaemic agent or insulin in a patient with poorly controlled non-Type 1 diabetes if not already on insulin or an oral hypoglycaemic agent.

Recommendation Scenario: Untreated poorly controlled non-Type 1 diabetes
Description: Poorly controlled diabetes, not Type 1, not currently on insulin or an oral hypoglycaemic agent.

*** Poorly controlled diabetes**
Description: Most recent HgbA1c in the last year > 7%
Expression: mostRecentHgbA1c12m.value >7

Data mapping for: Most recent HgbA1c Last 12 Months

Lab Type: HgbA1c (LOINC Code: 17856-6)
 Status: completed (Act Status: completed)
 Lab Type: HgbA1c (LOINC Code: 4548-4)
 Status: completed (Act Status: completed)
 Lab Type: HgbA1c (LOINC Code: 4549-2)
 Status: completed (Act Status: completed)

Expression: Let month : CodedValue = Factory.CodedValue("2.16.840.1.113883.6.1", "258700009") in patient laboratoryResult->select(resultType.codeSystem = "2.16.840.1.113883.6.1" and (resultType.code="4548-4" or resultType.code="17856-6" or resultType.code="4549-2") and resultStatus.codeSystem = "Local" and resultStatus.code = "Completed" and resultDateTime.high.occurredWithin(12, month))->sortBy("resultDateTime.value")->last()

Data Type: LaboratoryResult

a similar outcome. For example, a recommendation for the clinical scenario involving a patient with diabetes, hypertension, and elevated serum cholesterol, the action could be an assertion of a patient state of high risk for ischemic heart disease. A patient state is a class defined in the patient object model. A clinical scenario in another structured recommendation can include patient state in its logic (eg, in the scenario, ‘patient at high risk for ischemic heart disease’, the recommended action can be to initiate aspirin). The definitions and patient state recommendations provide a level of indirection that has two advantages: (1) they can be used to decouple inferences and decisions, such as the assessment of risk for heart disease from how to manage that risk, and (2) they can be used to clarify meanings of vague terms and to provide a consistent interpretation of those terms across recommendations and guidelines.

USE EXPERIENCE

We have used the multi-layered framework to encode the recommendations from three guidelines for use within a CDS system. The guidelines we encoded were for (1) the screening and management of diabetes mellitus (developed at Partners Healthcare in Boston, Massachusetts), (2) the use of aspirin in patients with coronary artery disease (US Preventive Services Task Force or USPSTF),³⁷ and (3) screening for hypertension in adults (USPSTF).³⁸ From these narrative guidelines, we created semi-structured and structured recommendations. Eventually, the latter were used to create executable rules in the iLog Rule Language, the format for production rules in iLog, a commercially available business rules engine. These rules are used within a CDS

web service,³⁹ developed by Partners Healthcare and demonstrated by the CDS Consortium,³³ to provide decision support, first within the ambulatory EMR system at Partners Healthcare, and then within an ambulatory EMR system at Regenrief Institute in Indianapolis. The semi-structured recommendations, the structured recommendations, and the rules in iLog Rule Language were created by knowledge engineers and CDS developers at Partners Healthcare. They were supported in their work by the developers of this knowledge representation model because the model and the editing tools were nascent.

PRELIMINARY ASSESSMENT

Our hypothesis in creating the multi-layered approach is that the implementability of a guideline recommendation in a CDS system would increase with the degree of structuring of the content, that is, structured recommendations would be easier to implement than the semi-structured recommendations. We also believe that structured recommendations would not decrease the flexibility desired by an organization in implementing a recommendation since workflow or local data considerations were not included within the recommendations. The implementation experience described above is limited to use of the CDS at one organization. Thus, it does not yet inform us of the impact of the multi-layered approach on implementation of guideline recommendations within multiple CDS systems. A study that assessed the authoring of knowledge in the multi-layered framework by experts would be well suited to test our hypothesis. However, since our authoring tools were evolving, such a study was not possible.

Therefore, we asked CDS implementation experts to rate three semi-structured recommendations and three structured recommendations across several implementation-related dimensions. One structured recommendation and one semi-structured recommendation were used from each of the three guidelines that are mentioned in the previous section.

The evaluation instrument we used was derived from the GuideLine Implementability Appraisal (GLIA).⁴⁰ Of the nine GLIA dimensions, we used five: decidability, executability, presentation, flexibility, and computability. In our opinion, the other four dimensions (novelty, effect on process of care, measurable outcomes, and apparent validity) would be impacted to a similar degree in structured and semi-structured recommendations. We incorporated and modified the items from GLIA to arrive at 14 items in the instrument (box 1). The response scale for each item was 'yes', 'no', and 'cannot determine'. We also included an open-ended question for the experts to share general comments regarding the recommendations or their rationale for their answer to any of the GLIA items.

We invited 24 CDS experts from the member institutions of the CDS Consortium³³ (Partners Healthcare, Veterans Health Administration, Kaiser Permanente, and Regenstrief Institute) to participate in the evaluation. The CDS experts included clinical domain experts, knowledge engineers, and analysts who were involved in CDS implementation at their respective organizations. We used a web-based tool to deliver one of six recommendations and the evaluation instrument to each expert. The recommendations were provided at weekly intervals alternating between the semi-structured recommendation and the structured recommendation. All experts received the same recommendation at the same time. When we provided the first structured recommendation (for diabetes mellitus) to the experts, we did not include any patient state recommendations associated with the primary recommendation that we were evaluating. Since this did not provide all the information that the experts needed, in subsequent rounds we included the patient state recommendations with the primary structured recommendation.

Table 2 shows the rate of response for each recommendation. The first recommendation evaluated was the semi-structured one for diabetes mellitus for which 19 experts (79%) completed the instrument. There was a subsequent decline in the response rate. Table 4 (available as an online data supplement) shows the frequency of responses by recommendation and item.

We constructed a 2×2 table of the recommendation type (semi-structured or structured) and the user's response ('yes' or 'no') for each item after discarding the 'cannot determine' answers and/or non-responses. The null hypothesis was that there was no difference in the implementability of a semi-structured recommendation as compared to a structured recommendation. We performed the χ^2 test of significance, except for cases of small cell size, when we employed the Fisher's exact test.⁴¹ Items 5, 6 and 13 had significant p values (table 3) with a greater number of 'yes' responses for the structured recommendation. The difference in responses to items 5, 6, and 13 suggests that structured actions are more implementable than semi-structured ones. This effect was not seen for clinical scenarios (items 3, 4, 10, 11, and 12).

DISCUSSION

The multi-layered framework provides an approach to implementing evidence-based recommendations within a CDS system at a healthcare organization. The framework allows a stepwise

Box 1 The items included in the instrument used in our preliminary evaluation study

1. Is the recommendation validatable and easily QA'd for inconsistencies, redundancies, and missing cases?
2. Can the recommendation be shared as-is between groups or institutions as content for computer-based decision support?
3. Would the guideline's intended audience consistently determine whether each condition in the recommendation has been satisfied? That is, is each and every condition described clearly enough so that reasonable practitioners would agree when the recommendation should be applied?
4. If there is more than one condition in the recommendation, is the logical relationship among all conditions (ANDs and ORs) clear?
5. Is the recommended action (what to do) stated specifically and unambiguously? That is, would members of the intended audience execute the action in a consistent way? In situations where two or more options are offered, the executability criterion is met if the user would select an action only from the choices offered.
6. Is sufficient detail provided or referenced (about how to do it) to allow the intended audience to perform the recommended action, given their likely baseline knowledge and skills?
7. Is the recommendation (and its discussion) concise?
8. Can criteria be extracted from the guideline that will permit outcomes of this recommendation to be measured?
9. Can criteria be extracted from the guideline that will permit measurement of adherence to this recommendation? Measurement of adherence requires attention to both the actions performed and the appropriateness of the circumstances under which they are performed.
10. Is the recommendation flexible enough to take into account individual clinical and non-clinical factors that are not enumerated in the scenario?
11. Are all patient data needed for this recommendation available electronically in the system in which it is to be implemented?
12. Is each condition of the recommendation defined at a level of specificity suitable for electronic implementation? (absence of ambiguity, granularity of data)
13. Is each recommended action defined at a level of specificity suitable for electronic implementation? (absence of ambiguity, granularity of data)
14. Is it clear by what means a recommended action can be executed in an electronic setting, for example, creating a prescription, medical order, or referral, creating an electronic mail notification, or displaying a dialog box?

transformation of the knowledge from evidence synthesis, expert opinions, and policy statements into executable knowledge. Each layer serves a different purpose and leverages the expertise of different roles required to implement the CDS. An analogy can be made with the software engineering process, which starts with requirements development and transforms into functional specifications, logical designs, and eventually a physical implementation. Different roles are involved at each step including the business stakeholders, analysts, architects, and programmers. In the encoding of the three guidelines, we successfully leveraged the expertise of different team members at different layers of the knowledge.

Table 2 Response rate for the expert ratings of the recommendations

Guideline	Recommendation type	Number of experts responding
Coronary artery disease	Semi-structured	7
Diabetes mellitus	Semi-structured	19
Hypertension	Semi-structured	9
Coronary artery disease	Structured	13
Diabetes mellitus	Structured	10
Hypertension	Structured	9

Unlike some of the other knowledge representations described earlier, our approach differs in two important respects: (1) it does not prescribe a particular CDS execution framework, and (2) it focuses on modeling clinical decisions rather than workflows. When an execution framework is specified, it is possible to share or distribute recommendations in an executable format, thus reducing the effort to implement CDS. However, the challenge has been to implement these execution engines and the executable knowledge bases within clinical settings. This challenge, perhaps, is evidenced by the lack of use of these frameworks outside of research projects. By not specifying the execution framework, this approach recognizes diversity in the existing implementations of clinical information systems and CDS tools. Thus, organizations are not required to make new investments or learn new technologies in order to incorporate evidence-based recommendations into their clinical practice. By focusing on representing a clinical decision, we accept that organizational capabilities, resources, structures, and workflows differ from each other, and that these will lead to different CDS implementations of the same recommendation. As pointed out in an editorial by Waitman and Miller, this local implementation that leverages the organization's computational and clinical resources is important for the effective adoption of a CDS.⁴² Using the multi-layered framework requires more effort by an organization to translate the recommendations into their CDS system as compared to using a centrally created executable knowledge base. By providing an encoded specification of the recommendation, the framework allows organizations to focus on adapting the recommendation locally rather than on interpreting the meaning or intent of the recommendation.

Furthermore, our approach also allows distribution or sharing of executable knowledge such as Arden Syntax MLMs. If an organization's CDS system can consume the knowledge in this

Table 3 Statistical test results for the association between the recommendation type and the response of the clinical decision support (CDS) experts (n=67)

Item	p Value	Percentage of non-response or 'cannot determine'
1	0.765	10%
2	0.140	9%
3	0.220	7%
4	1.000*	22%
5	0.040	9%
6	0.022	15%
7	0.287*	9%
8	0.218*	19%
9	0.422*	22%
10	0.966	31%
11	1.000*	34%
12	0.116	10%
13	0.045	6%
14	0.847	16%

Asterisked p values are from Fisher's exact test, and all other p values are from the χ^2 test.

executable format, this could be implemented at the site after adapting the knowledge to local resources and workflows, and mappings to data and actions in the institutional EMR.¹⁴ Thus, executable knowledge complements the semi-structured recommendations and structured recommendations, which require additional effort to transform them into executable knowledge but can be used by a broader group of organizations. The knowledge management portal that we have created within the CDS Consortium supports this approach of sharing knowledge.³³ The portal can be used to share knowledge at any layer of the framework. In fact, by linking the same recommendation across the four layers, we provide a path to review and possibly enhance the logical consistency of the knowledge in these layers.

The experience in implementing recommendations from three guidelines as rules in a CDS system has demonstrated the potential of this approach. However, that experience is limited to encoding of knowledge and its use at only one organization. The results of the preliminary assessment involving experts from multiple sites, noted as statistically significant for items 5, 6, and 13, suggested that the framework might facilitate the implementation of recommendations in CDS systems in different settings. Nevertheless, due to the small sample size and low response rate for some of the recommendations, stronger conclusions cannot be drawn. Further studies that quantify the efforts in implementing recommendations using the multi-layered framework and compare those to existing approaches are needed to clearly establish the impact of using the multi-layered framework.

While the responses of the experts indicated that actions would be more implementable moving from semi-structured to structured recommendations, we did not see the expected improvement in implementability of clinical scenarios. One possible reason for this was that, due to an oversight, the first structured recommendation reviewed by the experts did not include associated patient state recommendations. The experts, thus, could not determine how we had defined criteria such as 'patient has coronary artery disease'. A second possible reason was the use of GELLO to specify the logical criteria in scenarios. In subjective responses and conversations with CDS experts, they indicated that the criteria in GELLO (figure 3, bottom) were difficult to understand. Thus, the precise meaning of logical criteria may not have been conveyed to the expert reviewers.

A limitation of the semi-structured and structured representation models is that they represent knowledge that is agnostic of the CDS modality. Our implementation experience for the three guidelines is limited to using the recommendations to create rules for clinical reminders. The framework currently lacks the support to represent knowledge that is specific to a CDS modality, such as order sets, but independent of local organizational and workflow considerations.

The limitations of the knowledge representation in this four-layered framework suggest areas for future work, although the conceptual basis for the layers appears sound. We plan to incorporate within the multi-layer framework the ability to represent specific CDS modalities, including reminders, order sets, and referential content for use with infobuttons. We believe that this can be accomplished with some enhancements to the structured recommendations model. We also intend to address the limitation encountered in the understandability of logical criteria due to the use of GELLO. We are considering two broad solutions: enhancements to the guideline-editing tool we have developed to provide assistance with writing criteria, and exploration into alternative expression languages. Finally, the

experience of using the framework within the CDS Consortium will provide opportunities to learn how the framework supports distribution of CDS knowledge across organizations.

CONCLUSIONS

We have presented a multi-layered framework and knowledge representation models for implementing guideline recommendations in CDS systems. Our experience and preliminary assessment indicate the promise of this approach in creating knowledge in a form that is broadly usable and efficiently implementable in CDS systems although further studies are needed for more definitive results.

Acknowledgments We are grateful to members of the CDS Consortium, particularly Drs Sam Brandt, Frank Sonnenberg, Jason Saleem, Marc Overhage, Dean Sittig, and Brad Doebbeling for their valuable feedback on the framework, and Christine Kucera, Justine Pang, Angela Haslbeck, and Zachary Turechek for their role in coordinating the research activities.

Funding This research was funded by the Agency for Healthcare Research and Quality (AHRQ) under contract HHSA29020080010.

Competing interests None.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

1. Field MF, Lohr KN, eds; Committee to Advise the Public Health Service on Clinical Practice Guidelines, Institute of Medicine. *Clinical Practice Guidelines: Directions for a New Program*. Washington, DC: The National Academies Press, 1990.
2. Mangione-Smith B, DeCristofaro AH, Setodji CM, et al. The quality of ambulatory care delivered to children in the United States. *N Engl J Med* 2007;**357**:1515–23.
3. McGlynn EA, Asch SM, Adams J, et al. The quality of health care delivered to adults in the United States. *N Engl J Med* 2003;**348**:2635–45.
4. Garg AX, Adhikari NKJ, McDonald H, et al. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *JAMA* 2005;**293**:1223–38.
5. Kawamoto K, Houlihan CA, Balas EA, et al. Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *BMC* 2005;**330**:765.
6. Dexter PR, Perkins S, Overhage JM, et al. A computerized reminder system to increase the use of preventive care for hospitalized patients. *N Engl J Med* 2001;**345**:965–70.
7. Grimshaw JM, Russell IT. Effect of clinical guidelines on medical practice: a systematic review of rigorous evaluations. *Lancet* 1993;**342**:1317–22.
8. Sintchenko V, Coiera E, Iredell JR, et al. Comparative impact of guidelines, clinical data, and decision support on prescribing decisions: an interactive web experiment with simulated cases. *J Am Med Inform Assoc* 2004;**11**:71–7.
9. Chaudhry B, Wang J, Wu S, et al. Systematic review: impact of health information technology on quality, efficiency, and costs of medical care. *Ann Intern Med* 2006;**144**:742–52.
10. Codish S, Shiffman RN. A model of ambiguity and vagueness in clinical practice guideline recommendations. *AMIA Annu Symp Proc* 2005:146–50.
11. Boxwala AA, Tu S, Peleg M, et al. Toward a representation format for sharable clinical guidelines. *J Biomed Inform* 2001;**34**:157–69.
12. Osheroff J, Pifer E, Teich J, et al. *Improving Outcomes with Clinical Decision Support: An Implementer's Guide*. 1st edn. Productivity Press, 2005.
13. Hripcsak G. Writing Arden Syntax Medical Logic Modules. *Comput Biol Med* 1994;**24**:331–63.
14. Pryor TA, Hripcsak G. Sharing MLM's: an experiment between Columbia-Presbyterian and LDS Hospital. *Proc Annu Symp Comput Appl Med Care*. 1993:399–403.
15. Johnson PD, Tu SW, Musen MA, et al. A virtual medical record for guideline-based decision support. *Proc AMIA Symp* 2001:294–8.
16. Young O, Shahar Y, Liel Y, et al. Runtime application of Hybrid-Asbru clinical guidelines. *J Biomed Inform* 2007;**40**:507–26.
17. Tu SW, Musen MA. Modeling data and knowledge in the EON guideline architecture. *Stud Health Technol Inform* 2001;**84**:280–4.
18. Boxwala AA, Peleg M, Tu S, et al. GLIF3: a representation format for sharable computer-interpretable clinical practice guidelines. *J Biomed Inform* 2004;**37**:147–61.
19. Fox J, Patkar V, Thomson R. Decision support for health care: the PROforma evidence base. *Inform Prim Care* 2006;**14**:49–54.
20. Tu SW, Campbell JR, Glasgow J, et al. The SAGE Guideline Model: achievements and overview. *J Am Med Inform Assoc* 2007;**14**:589–98.
21. Collins SA, Currie LM, Bakken S, et al. Information needs, Infobutton Manager use, and satisfaction by clinician type: a case study. *J Am Med Inform Assoc* 2009;**16**:140–2.
22. Lobach DF, Kawamoto K, Anstrom KJ, et al. Development, deployment and usability of a point-of-care decision support system for chronic disease management using the recently-approved HL7 decision support service standard. *Stud Health Technol Inform* 2007;**129**:861–5.
23. Shahar Y, Young O, Shalom E, et al. The Digital electronic Guideline Library (DeGel): a hybrid framework for representation and use of clinical guidelines. *Stud Health Technol Inform* 2004;**101**:147–51.
24. Hatsek A, Young O, Shalom E, et al. DeGel: a clinical-guidelines library and automated guideline-support tools. *Stud Health Technol Inform* 2008;**139**:203–12.
25. Shiffman RN, Karras BT, Agrawal A, et al. GEM: a proposal for a more comprehensive guideline document model using XML. *J Am Med Inform Assoc* 2000;**7**:488–98.
26. Wright A, Sittig DF, Ash JS, et al. Clinical decision support capabilities of commercially-available clinical information systems. *J Am Med Inform Assoc* 2009;**16**:637–44.
27. Sittig DF, Wright A, Meltzer S, et al. Comparison of clinical knowledge management capabilities of commercially-available and leading internally-developed electronic health records. *BMC Med Inform Decis Mak* 2011;**11**:13.
28. Sherman EH, Hripcsak G, Starren J, et al. Using intermediate states to improve the ability of the Arden Syntax to implement care plans and reuse knowledge. *Proc Annu Symp Comput Appl Med Care* 1995:238–42.
29. Wang D, Peleg M, Tu SW, et al. Design and implementation of the GLIF3 guideline execution engine. *J Biomed Inform* 2004;**37**:305–18.
30. Boxwala AA, Greenes RA, Deibel SR. Architecture for a multipurpose guideline execution engine. *Proc AMIA Symp* 1999:701–5.
31. Murphy SN, Mendis M, Hackett K, et al. Architecture of the open-source clinical research chart from Informatics for Integrating Biology and the Bedside. *AMIA Annu Symp Proc* 2007:548–52.
32. Middleton B, Anderson J, Fletcher J, et al. Use of the WWW for distributed knowledge engineering for an EMR: the KnowledgeBank concept. *Proc AMIA Symp* 1998:126–30.
33. Middleton B. The clinical decision support consortium. *Stud Health Technol Inform* 2009;**150**:26–30.
34. **Clinical Decision Support Consortium**. *Knowledge Management Portal*. <http://cdsportal.partners.org/> (accessed 7 Sep 2011).
35. Sordo M, Boxwala AA, Ogunyemi O, et al. Description and status update on GELLO: a proposed standardized object-oriented expression language for clinical decision support. *Stud Health Technol Inform* 2004;**107**:164–8.
36. *HITSP Summary Documents Using HL7 Continuity of Care Document (CCD) Component v 2.5*. New York, NY: Healthcare Information Technology Standards Panel, 2009.
37. **U.S. Preventive Services Task Force**. Aspirin for the prevention of cardiovascular disease: U.S. Preventive Services Task Force recommendation statement. *Ann Intern Med* 2009;**150**:396–404.
38. **U.S. Preventive Services Task Force**. Screening for high blood pressure: recommendations and rationale. *Am J Prev Med* 2003;**25**:159–64.
39. Paterno MD, Schaeffer M, Van Putten C, et al. Challenges in creating an enterprise clinical rules service. *AMIA Annu Symp Proc* 2008:1086.
40. Shiffman RN, Dixon J, Brandt C, et al. The GuideLine Implementability Appraisal (GLIA): development of an instrument to identify obstacles to guideline implementation. *BMC Med Inform Decis Mak* 2005;**5**:23.
41. Agresti A. *Categorical Data Analysis*. 2nd edn. Hoboken, NJ: Wiley-Interscience, 2002.
42. Waitman LR, Miller RA. Pragmatics of implementing guidelines on the front lines. *J Am Med Inform Assoc* 2004;**11**:436–8.