

The SAGE Project: A Universal Framework for Sharing Health Knowledge in the Form of Computable Clinical Practice Guidelines

Robert M. Abarbanel¹, Nick Beard¹, James R. Campbell², Julie I. Glasgow¹, Stanley M. Huff³, James G. Mansfield¹, Eric Mays⁴, James McClay², David N. Mohr⁵, Mark A. Musen⁶, Craig G. Parker³, Prabhu Ram¹, Roberto A. Rocha³, Sidna M. Scheitel⁵, Samson W. Tu⁶, Tony Weida⁴, Qin Ye¹

¹Health Informatics Department, IDX Systems Corporation, Seattle, WA

²University of Nebraska Medical Center, Nebraska Health Systems, Omaha, NE

³Intermountain Health Care, Salt Lake City, UT

⁴Apelon, Inc. Ridgefield, CT

⁵Mayo Medical School, Mayo Clinic, Rochester, MN

⁶Stanford Medical Informatics, Stanford University Medical Center, Stanford, CA

Email: Nick_Beard@IDX.com

Abstract

The SAGE Project is a multi-institution effort to enable encoding and dissemination of interoperable, computable clinical practice guidelines. We have developed a standards-based guideline-knowledge representation model that specifies computable guideline content. We incorporate a “virtual medical record” that mediates both the encoding of guideline knowledge and the subsequent mapping to idiosyncratic CIS information environments. We employ a “workflow aware” approach to facilitative interaction between guideline-driven decision support and the host CIS. We have developed a guideline execution technology that interprets encoded guideline content; activates guideline logic in response to appropriate clinical events; retrieves patient data from the electronic medical record; makes patient-specific recommendations based on guideline logic; and surfaces guideline-driven decision support via actions of the host CIS. In results to date, we have encoded and deployed exemplar guidelines covering a range of clinical domains, including acute and chronic care, multi-disciplinary care settings; and population-based guideline logic.

1. Introduction

Clinical practice guidelines (CPGs) have attracted intense interest from the health care industry because of the widespread belief that they can improve the

quality of care, in particular by reducing the variability of care and by reducing omission of recommended best treatment practices [a,b].

Despite this substantial interest, CPGs have yet to realize their potential to improve health care quality because they have failed to influence clinician behavior significantly [c,d]. Whether hardcopy or online, most CPGs are available today in the form of text (e.g., documents, PDF files, charts). To use a guideline, clinicians must interrupt their normal workflow; locate the appropriate guideline; read the guideline; and then determine how the guideline recommendations apply to the patient at hand. In this current mode, the entire knowledge retrieval and processing burden for appropriate use of guidelines rests on the capabilities of the over-loaded clinician.

There is encouraging evidence in decision support (DSS) literature that computer-based reminders and recommendations can improve clinician compliance with evidence-based best practices if the provided guidance is patient-specific and well-integrated into the clinician’s workflow [e,f,g]. In more advanced clinical information systems, integration of technology for processing clinical rules is becoming standard functionality. At the same time, research efforts to improve capabilities to represent CPG knowledge in computable formats [h] and to integrate CPG-driven recommendations with CIS functionality [i,j,k], are beginning to address some of the daunting informatics challenges involved.

It is still the case that wide-spread distribution and use of computable CPG content is prevented by lack of standards for representing medical knowledge, and by the prohibitive complexity and expense required to adapt encoded guideline content across the heterogeneity of data structures, semantics, and medical vocabularies in use in the nation's health care information systems. The vision of the SAGE project is computable CPG knowledge, shareable across institutions at a reasonable cost and effort, and in a form that can be integrated gracefully and supportively into the clinician's workflow via functions of the local clinical information system.

2. The SAGE Project Approach

The SAGE Project is a multi-site, industry-academic collaboration initiated to develop a standards-based, comprehensive technology infrastructure that will enable encoding and dissemination of interoperable, computable CPGs. Technical objectives of the SAGE Project are: (a) An interoperable guideline model – a standard computable formalism for representing the content and logic of CPGs; (b) A guideline “workbench” – software tools for authoring, editing, encoding, and maintaining guidelines in the format of the guideline model; (c) A common layer of standards-based terminologies and information models; and (d) A guideline deployment system - software that integrates electronic guidelines with the local electronic medical record (EMR) and surfaces guideline content via functions of the local CIS.

A key project approach has been to employ (and extend where necessary) existing informatics standards, and to collaborate closely with standards development organizations such as HL7. Requirements were driven, in part, by use cases for encoding and deployment of selected exemplar CPGs that sampled a range of clinical conditions (immunizations, diabetes management, community acquired pneumonia) and domains (acute vs. chronic care, inpatient vs. ambulatory care, etc.). Project R&D methodology comprised iterative development cycles in which the latest working prototypes of all components were combined in successive, formal exercises of interoperable guideline functionality.

The SAGE system architecture and guideline model are distinguished in part from their predecessors by the degree to which execution considerations have shaped their elements. To achieve interoperability between a guideline-based decision-support system and a CIS, the SAGE approach (see Figure 1) employs a common layer of standard information models and terminologies. These information models and standard terminologies are used during guideline encoding, and also mediate mapping of guideline content to the specific patient data sets, information models, settings, roles, resources and procedures of the local CIS. The SAGE decision support engine [1], acting via a standards-based API, provides dynamic decision support to supplement CIS patient care functions. Guideline recommendations, as well as access to evidence and rationale, are integrated into a clinician's online workflow via functions of the local CIS.

3. The SAGE Guideline Model

Requirements for the SAGE Guideline Model build on earlier elucidative work [h,m] and emphasize: (a) comprehensive, flexible representation of guideline knowledge in a computable formalism; (b) interoperable sharing of encoded guideline content between heterogeneous CIS environments; and (c) integration of active guideline recommendations within the clinical workflow, via functions of the local CIS. The SAGE Guideline Model is informed by invaluable earlier work [n,o,p], but has been constructed anew to incorporate emerging informatics standards in a systematic way. HL7 v3 data types are incorporated directly into our guideline model; standard ontologies (i.e., SNOMED CT®, LOINC®) are used for concept encoding; and a proposed HL7 standard information model for patient data; the Virtual Medical Record (VMR) is used to mediate both queries to the electronic medical record (EMR) and actions (e.g., display an alert) directed to the CIS.

Some earlier guideline models [q] assumed a decision support system would *drive the clinical workflow* and thus were required to represent an entire detailed clinical workflow. In contrast, the SAGE system is designed to *respond to opportunities for decision support*, and to that end, the SAGE Guideline Model need represent only enough of the clinical workflow context required to trigger guideline DSS services at appropriate points in the care process.

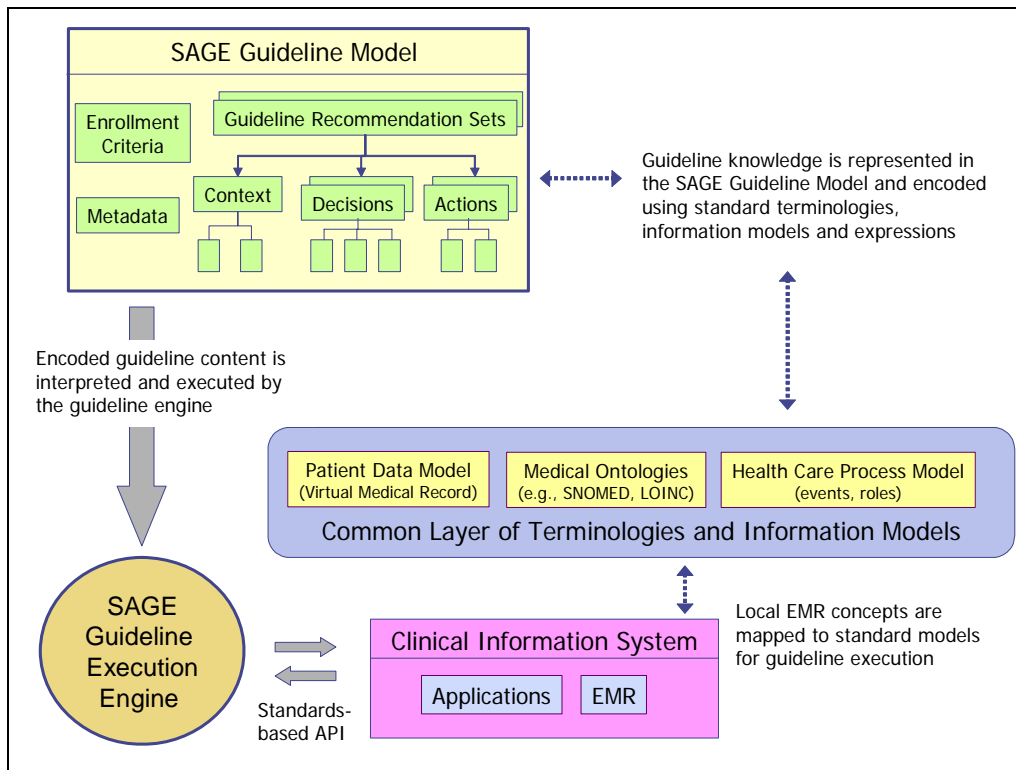


Figure 1: Overview of SAGE interoperable guideline infrastructure

3.1 Recommendation Sets

As suggested in the upper left of Figure 1, the SAGE guideline model organizes guideline content into guideline *recommendation sets* [r]. A recommendation set is a constellation of guideline content tailored to the workflow, roles, entities and actions of a specific care-delivery context. Recommendation sets may be *activity graphs* (e.g., specification of computational algorithms or medical care plans), or *decision maps*, a collection of guideline decisions applicable at a single point in time. As illustrated in Figure 2, a recommendation set employs the model elements: context, decision, action, and route nodes. These nodes are complex knowledge representation objects used to create standards-based, computable representations of decision support knowledge. Each context, decision, action, or route node comprises many concepts, and each concept has multiple attributes. Each attribute is specified at the level of granularity required to be computable. The

model accommodates cyclic and iterative recommendation sets, and supports temporal management. A single guideline may comprise multiple recommendation sets. This approach organizes and simplifies decision logic around common patient care scenarios. During guideline encoding, our workbench tools allow specification of surface-level logic (e.g., Figure 2), followed by increasingly “deep” layers of concept specification.

Context Nodes are used to make explicit and computable the workflow events, clinical settings, and patient states that define an opportunity for guideline DSS. In the example shown in Figure 2, the context node *Newborn Admission* is defined by a combination of: clinical setting (inpatient), patient state (newborn), clinician roles (nurse), and specific CIS events (admission). The *Newborn Admission* context node triggers a session of guideline interaction: a run-time sequence of interactions between guideline logic and the clinical user, intermediated by functions of the local CIS.

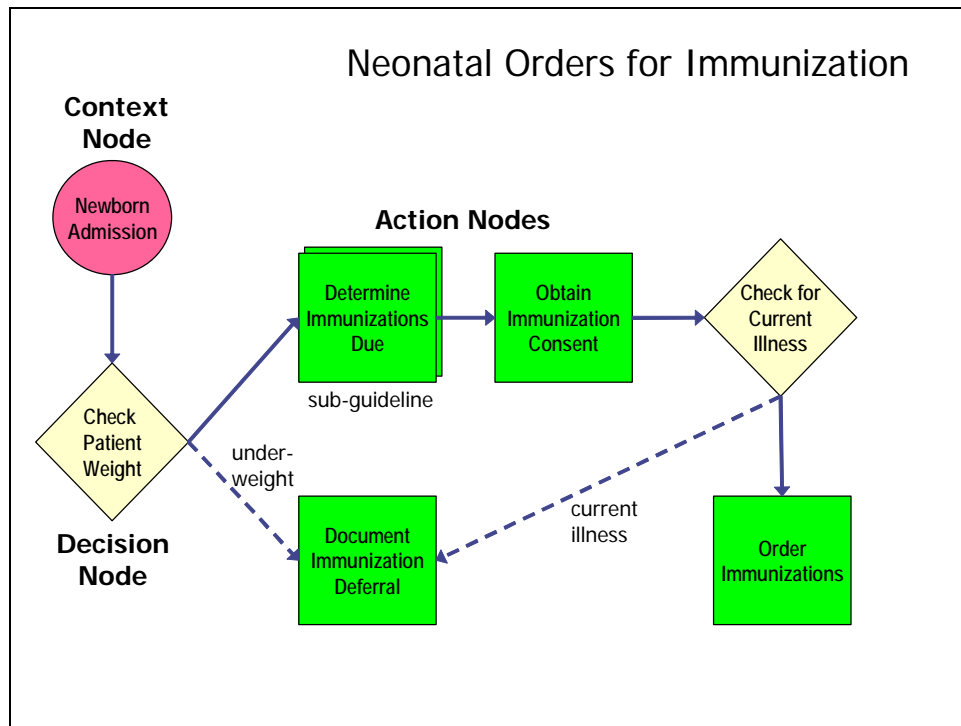


Figure 2: A simplified recommendation set for neonatal immunization orders

Action Nodes model activities in support of a recommendation set -- typically work items to be performed either by computer or by a health care provider. The action nodes shown in Figure 2 depict several examples, including: Query and analyze data from the EMR [*Determine Immunizations Due*]; Prompt or alert user [*Obtain Immunization Consent*]; Print a form [*Document Immunization Deferral*]; and Place a pending medication order [*Order Immunizations*]. Action nodes have multiple complex attributes, including: *triggering_events*: specification of events that may trigger an action node, and *action_specification*: specification of the tasks to be performed. Action nodes can specify single actions (e.g., *Obtain Immunization Consent*), or subguidelines (e.g., *Determine Actions Due*, which in the Figure 2 is a subguideline encapsulating multiple nodes). Subguidelines are reusable collections of guideline logic and provide for efficiency and complexity management during the guideline encoding process.

Decision Nodes support representation of the knowledge required to make a choice among alternatives within a guideline. Decision nodes also

have multiple attributes, including: *triggering_events*: events that may trigger a decision node; and *decision_model*: specification of the decision-making knowledge and methodology used to generate preferences among alternatives. Decision nodes (e.g., *Check for Current Illness*, in Figure 2) commonly specify the acquisition of data directly from the patient EMR and the employment of a decision model to evaluate branching logic.

Routing Nodes (not shown in this example) are “dummy” activities, used to control branching or synchronization points in guideline logic.

3.2 Enrollment Criteria; Meta Data

The SAGE Guideline Model also specifies Enrollment Criteria -- attributes of the patient that must be true before a patient can be enrolled in the guideline. Enrollment Criteria are similar to inclusion and exclusion criteria for clinical trials and are entry conditions (which may not need to remain true after a patient is enrolled). The SAGE model supports either manual enrollment by a clinician, or automatic enrollment based on guideline logic. The

SAGE model also supports *de-enrollment criteria* – attributes that would cause enrolled patients to be taken off of a guideline. Specification of *meta data* within the SAGE Guideline Model provides support for guideline reference material and guideline management, (e.g., indexing, version control).

4. Semantic Interoperability

A significant challenge to sharing guideline knowledge in a computable form is widespread heterogeneity in the medical terminologies and patient data models in deployed CISs. Semantic interoperability requires a solution that not only maps concepts from one terminology to another, but reconciles differences in information models between systems. The idea of creating sharable, executable decision support mechanisms is not new with the SAGE project. The Arden syntax represents a significant effort to define a sharable representation for medical logic modules [s]. However, since Arden uses implementation specific code {inside the curly braces} to reference data items, the task of knowledge sharing among institutions requires repeated “re-mapping” to each local environment.

4.1 Terminologies and Information Models

The SAGE approach to semantic interoperability simplifies (but does not eliminate) the “curly braces” problem by specifying a single set of standards for terminology and explicit information models. Mapping becomes one-to-many (as described below), rather than many-to-many.

SAGE guideline encoding employs current and emerging national standards, such as SNOMED-CT, LOINC, and NDF-RT, and our guideline workbench, based on Protégé-2000 [t], provides integrated access to terminology services to facilitate the encoding process. Standard terminologies are a necessary component of the solution to this problem, but alone they are not sufficient.

4.2 The Virtual Medical Record

In the SAGE infrastructure, we address semantic interoperability by employing and extending the concept of the *virtual medical record* (VMR)[u]. The SAGE VMR is an object model of both patient medical record information and CIS actions, that is simplified for decision-support purposes. The VMR supports a structured data model for representing

information related to individual patients, domains for values of attributes in the data model, and queries through which guideline decision support can test the states of the patient. Our VMR is based on HL7 RIM derived artifacts (e.g., HL7 domain information models). The VMR allows us to specify a broad set of classes of information that are of interest for clinical guidelines.

The SAGE VMR is defined by a set of classes, each with attributes that represent the features needed for decision support. Twelve VMR classes are currently implemented: Agent, Allergy, Appointment, Encounter, Goal, MedicationOrder, Observation, Order, Problem, Procedure, Referral, and SubstanceAdministration. Figure 3 lists the attributes of the VMR Observation class.

Observation	
code	value
effectiveTime	encounter
method	subject
interpretation	text

Figure 3: the VMR Observation class

5. SAGE Guideline Execution

Key requirements for the SAGE Guideline Execution Infrastructure are: (a) manageable installation of guidelines; (b) successful integration of guideline DSS (e.g., queries and actions) with CIS functions; (c) detection of appropriate CIS events; (d) interaction with terminology services; and (e) standards-based mapping of CIS data to VMR-based guideline standards.

5.1 Guideline Installation and Execution

In a typical deployment environment, we assume a SAGE guideline knowledge base is first imported into the local health care delivery organization. The knowledge base is delivered in a standard XML format that contains a full specification of the encoded clinical and operational logic, as well as information about versions, authors, and sources for the guideline content. This includes pointers to references and source data as appropriate, as well as test data and installation scripts.

Installation of the guideline involves two main steps: (1) the local institution may edit the guideline

to conform to its organizational and clinical policies prior to deploying the guideline – we call this step *localization*; (2) mapping from the standards-based concepts in the encoded guideline to local EMR data and local CIS functions, a process we refer to as *binding*.

5.2 Architectural Overview

As shown in Figure 4, the SAGE guideline deployment architecture consists of a SAGE execution engine (henceforth referred to as "the engine"), an event listener, a terminology server, and a set of interfaces called VMR/Action Services which interoperate with the local CIS. Each of these components has well-defined standard interfaces and could be replaced by other implementations that obey the same interfaces and operate using the same semantics.




Figure 4 – SAGE deployment architecture

The *SAGE Execution Engine* is designed specifically to read SAGE guideline knowledge bases and execute content according to the documented execution method for the standard guideline model. The engine interprets the content of the Context, Action, Decision, and Route nodes in an encoded guideline, executes workflow and decision logic, and interacts appropriately with the CIS.

The *Event Listener* is the mechanism by which the engine is notified of state changes in the CIS. The listener and the VMR Services are implemented as web services [v] allowing for broad interoperability, and can be used by any conforming CIS to publish events. Trigger events encoded in the guideline are registered with the CIS's event manager, thereby expressing the execution engine's interest in these events. When a relevant CIS event is detected, the engine begins interpreting guideline content associated with that trigger.

The *terminology server* encapsulates standard terminologies, and implements subsumption and conversions that may be used by the engine. It

operates, like the other components, over standard interfaces. The SAGE engine can communicate with terminology services either embedded in the host CIS or provided by a third-party.

The *VMR/Action Services* are interfaces into both patient data and application functionality provided by the CIS. The VMR Services are used to mediate patient data queries from the EMR, and the Action Services are used to initiate actions within the CIS. The VMR/Action Services can be viewed as wrappers around existing CIS data and functionality. They support interoperability by presenting a unified view of clinical information systems to the guideline execution engine.

The binding process "maps" the defined VMR Services to underlying capabilities in a particular CIS. The binding process for the VMR Services is a non-trivial job that needs to be done once and then adjusted in minor ways as functions are utilized by newly installed guidelines. While each CIS will have unique bindings for the VMR, each successive installed guideline will share most binding details with guidelines previously installed for the same CIS.

5.3 State and Concurrency

In general, the guideline execution engine executes in a stateless manner. The guideline encoding provides explicit entry points into activity graphs by marking certain Contexts as starting nodes. These context nodes always have a triggering event that initiates the execution of a particular activity graph. Once execution begins additional nodes are visited through the transitions encoded in the recommendation set. It should be noted that several activity graphs can execute concurrently and several paths within an activity graph or a decision map may execute concurrently, resulting in multiple threads of guideline execution for patients at a point in time.

5.4 Processing a Guideline

We will illustrate SAGE guideline processing using our example recommendation set for neonatal immunizations (Figure 2). In our execution, we assume a patient-guideline association has been established (either through an overt action by a clinician or through an automatic process), such that the patient is "enrolled" for this guideline. The engine uses this association to filter events that come

from the CIS and to maintain guideline state for that patient/guideline combination.

The primary guideline logic is as follows: Check the weight of the newborn. If the weight is over 2kg, check prior medications and determine vaccines that need to be administered. If the medical records do not indicate prior “consent to immunization” then have the clinician obtain permission to administer immunization. If the newborn is ill, the vaccine administration is to be deferred and the deferral reasons documented. Otherwise, the vaccines are to be ordered for the patient. If the neonate is under weight, the vaccine administration is to be deferred and the deferral reasons documented.

5.4.1. Context Handling. The *Newborn Admission* context node in Figure 2 specifies a clinical setting of Inpatient Hospital and a clinical role of Pediatric Nurse. It also specifies a triggering event (inpatient admission), which when received by the event listener causes the guideline execution engine to start processing. Our context node includes a patient specific precondition, *age < 7 days*, since a child below the age of 7 days is considered a neonate in this immunization guideline. Once the triggering event is received, the engine evaluates the precondition associated with the context node; if False, the execution halts. If True, the engine resolves the clinical settings and clinical roles specified in this context node and moves to processing subsequent nodes.

5.4.2. Decisions and Criteria Processing. The second node in the recommendation set is a decision node *Check Patient Weight*. This node specifies two alternative action nodes (*Determine Immunizations Due* and *Document Immunization Deferral*), along with the decision criteria that must be satisfied to process those nodes. The execution engine will process all decision alternatives for which criteria are met, thereby allowing concurrent execution of multiple paths.

Decision criteria supported by the SAGE guideline model and engine are listed here, using examples from our decision map that computes immunizations due for all ages:

(a) Comparison criterion -- used to compare an object returned by a VMR service method against a constant. (e.g., # of polio vaccines in medical history is zero);

(b) Temporal comparison criterion -- compares the temporal relationship between the time when a VMR instance occurred and some time interval (e.g., was MMR vaccine given in the last 4 weeks?);

(c) Variable comparison criterion -- uses variables that are defined across VMR service instances and some mathematical functions (e.g., age > 2 months);

(d) Presence criterion -- checks for the presence or absence of coded concept in instances of a VMR class within a valid time window (e.g., presence of infantile spasm as a Problem in patient’s medical record);

(e) Goal criterion -- allows goal testing (e.g., is the HbA1c within goal?) and

(f) N-ary criterion -- Boolean combinations (AND, OR, NOT) of the other criteria.

Terminology processing is handled during evaluation of criteria. For example, the comparison criterion *weight < 2kg* specifies the VMR class (Observation); the SNOMED CT code for body weight (27113001); the aggregate modifier for the weight observation (most recent); the operator to be used for comparison (less than); and the value to be compared against (2 kg). To evaluate the above criterion, the engine first makes a call to the terminology server to obtain all the codes that are subsumed by body weight. A VMR service call is then made to the Observation service of the CIS, passing in the SNOMED CT codes, aggregate modifier, patient identifier, etc. (Recall that during the installation/binding process, CIS-specific codes have been “mapped” to the SNOMED CT standard terminology, thereby allowing the standards-based side of the VMR service to interface with the parochial (CIS) side of the VMR service). From the signature of the Observation VMR service, the engine is aware of the measurement units used in the returned weight observation. The engine converts the returned weight and the value specified in the criterion to the same units for comparison, and then evaluates the criterion.

5.4.3. Actions. In our example, the *Obtain Immunization Consent* action node is a directive to the pediatric nurse to obtain consent for immunization from the patient’s guardian. A precondition for this action node is an N-ary criterion composed of “immunizations recommended” and “absence of immunization consent in patient record”. In our test implementation, this inquiry to obtain the consent is presented to the clinician through the

notification mechanism of the host CIS. The SAGE infrastructure supports synchronous as well as asynchronous notifications.

Since the immunization consent status (SNOMED CT: 243880000) is the finding involved in the inquiry, the valid responses are immunization consent given (SNOMED CT: 310375005) or immunization consent not given (SNOMED CT: 310376006) or the codes subsumed by either of these codes. The response from the clinician is recorded in the patient's medical record as an Observation. The engine detects this EMR update as a triggering event and processing continues.

To complete execution of this recommendation set, the decision node *Check for Current Illness* is visited and two alternatives are evaluated. If the patient is ill, the vaccine administration is deferred. If the patient is not ill, then the action node *Order Immunizations* is processed. The engine calls the VMR/Action service to place "pending" vaccine orders for this patient, and generates a notification to the clinician informing them of the presence of a medication order waiting to be approved.

5.4.4. Subguidelines. In Figure 2, *Determine Immunizations Due* is a specialized action node that includes a subguideline which embeds a decision map used to compute which immunizations are due for a patient. The conclusions made during the processing of this decision map are stored in the patient's medical record through VMR/Action calls to the CIS. The subguideline itself consists of several context, decision and action nodes. The same subguideline is used by other recommendation sets to calculate which immunizations are due for a particular patient at the time they are seen.

6. Results of Encoding and Execution

The challenge faced by the SAGE project is to deliver clinical practice guideline recommendations to clinicians as seamlessly as possible using native CIS applications and user interfaces. This must be achieved in a generic, interoperable manner so that the execution engine need not be rewritten for each CIS it needs to interact with. Moreover, all this must be achieved without requiring large changes to the existing functionality of clinical information systems.

We are achieving these goals. We have developed a guideline model that is based on standard information models, medical terminologies and HL7 data types. We have developed an engine to execute guidelines encoded using the SAGE guideline model. We have implemented an event listener that feeds the engine with external events. We have implemented the VMR/Action Services for a commercial CIS [w], so that guideline interactions can be provided through the interfaces of the CIS. All of these together form an infrastructure for us to be able to encode and execute an arbitrary clinical guideline.

Using our infrastructure, we have encoded a version of the Institute of Clinical Systems Improvement (ICSI) Immunization guideline [x] and have executed guideline scenarios using the engine and the CIS. This encoding has three recommendation sets: (1) a neonatal immunization scenario, (2) a primary care scenario that handles DTaP, Polio, Pneumococcal, Tetanus-Diphtheria, Influenza, MMR, HiB, Hepatitis A and Hepatitis B vaccine administration for children and adult patients, and (3) a immunization health maintenance reminder. The SAGE execution engine executes the above scenarios, and surfaces appropriate guideline recommendations in real time via existing functions of the host CIS.

We have developed and tested two additional examples: the ADA Diabetes guideline [y], and a guideline for Community Acquired Pneumonia (CAP) [z]. The diabetes guideline covered Type II diabetics for their standard long term care along with concomitant recommendations for hypertension and hyperlipidemia. The CAP guideline assists clinicians with triage for pneumonia patients and the orders that should be issued for different categories of severity.

All three exemplar guidelines will be tested during 2004 at both the University of Nebraska in Omaha, and the Mayo Clinic in Rochester. We plan to test multiple virtual patients to explore all the logic and functionality in our encoded guidelines. These tests will further evaluate site-specific localization (editing of guideline content), as well as binding from guideline standards to the local EMR database and CIS functions.

7. Discussion

Our experience in the SAGE Project has begun to illuminate several of the more advanced issues regarding integration of computable CPGs with host clinical information systems.

(1) In contrast to InterMed's GLIF [aa] approach, which assumes that guidelines will be encoded using a top-down approach, (i.e., starting with high-level medical logic and progressively refined to computer-interpretable and implementation-specific layers), the recommendation sets in the SAGE approach are highly dependent on details of workflow processes. We postulate that the more specifically a guideline-workflow interaction is encoded; the more useful it would be to a particular institution. Conversely, a more general guideline encoding would be more widely interoperable, but may require more specificity to be added during the localization process. We hypothesize that it may be feasible to formally separate institution-specific workflow knowledge from reusable guideline logic, as we did with workflow-specific activity graphs and the decision-map subguideline. However, the hypothesis remains to be tested with additional experiments.

(2) It is expected that in the future, patients will likely be enrolled and active on multiple guidelines simultaneously, and the resultant challenge to coordinating interacting guideline-driven DSS will be significant. While we acknowledge this challenge, current SAGE guideline encodings finesse this issue by either assuming a flow of work that is unique (and isolated) for each guideline, or by coordinating recommendations for multiple conditions within a single encoding.

(3) Advanced clinical information systems typically provide one or more sources of DSS (e.g., local rule engines or integrated knowledge bases). The SAGE architecture provides an additional source of executing DSS logic. Integration and management of these multiple sources of DSS remains an obstacle.

(4) Lastly, it is important to acknowledge that even with a comprehensive guideline encoding and execution infrastructure in place; it will remain a significant challenge to use this technology to change clinician behavior in the real care delivery environment. The industry's painful experience with

computerized physician order entry, as well as early efforts to implement guideline-driven DSS [bb,cc] indicate that success will depend heavily on graceful and facilitative integration with the care workflow.

Evaluation of the SAGE methodology is necessarily incomplete at this stage. Limitations of the current experiment include the fact that, even though the SAGE guideline model is rich enough to model complex medical and workflow processes that span multiple encounters and that require branching and synchronization of concurrent processes, our testing of these capabilities is currently limited to the guideline exemplars described above. Future research projects will expand and formally evaluate the efficacy of SAGE guideline technology in real care delivery settings.

Authorship: The authors of this paper are listed in alphabetical order. The principal investigator for the SAGE project is Nick Beard, MD, (IDX Systems). Site principle investigators are: James R. Campbell, MD, (University of Nebraska Medical Center), Stanley M. Huff, MD, (Intermountain Health Care), Eric Mays, PhD, (Apelon, Inc.), Mark A. Musen, MD, PhD, (Stanford Medical Informatics), and Sidna M. Scheitel, MD, MPH, (Mayo Clinic).

Acknowledgement: This work was partially supported by the U.S. Department of Commerce, National Institute of Standards and Technology, Advanced Technology Program, Cooperative Agreement Number 70NANB1H3049.

References

-
- [a] E. A. McGlynn, S.M. Asch, J. Adams, J. Keeseey, J. Hicks, A. DeCristofar, and E.A. Kerr, "The quality of health care delivered to adults in the United States," *New England Journal of Medicine*. 348(26), 2003, 2635-2645.
 - [b] A.P. Abernethy, G.P. Samsa, and D.B. Matchar, "A clinical decision and economic analysis model of cancer pain management," *American Journal of Managed Care*, October, 2003, 651-664.
 - [c] J. Lomas, G.M. Anderson, K. Domnick-Piere, E. Vayda, M.W. Enkin, and W.J. Hannah, "Do practice guidelines guide practice?," *New England Journal of Medicine*, 321(19), 1989, 1306-1311.
 - [d] M.D. Cabana, C.S. Rand, N.R. Powe, A.W. Wu, M.H. Wilson, P.A. Abboud, and H.R. Rubin, "Why don't physicians follow clinical practice guidelines? a framework for improvement," *Journal of the American Medical Association*, 282(15), 1999, 1458-1465.

- [e] A. Balas, S.M. Austin, J.A. Mitchell, B.G. Ewigman, K.D. Bopp, and G.D. Brown, "The clinical value of computerized information services: a review of 98 randomized clinical trials," *Archives of Family Medicine*, 5(5), 1996, 271-278.
- [f] D.F. Lobach and W.E. Hammond, "Development and evaluation of a computer-assisted management protocol (CAMP): improved compliance with care guidelines for diabetes mellitus," *Proc 19th Annu Symp Comput Appl Med Care*. 1994:787-91.
- [g] R.B. Elson and D.P. Connelly, "Computerized patient records in primary care: their role in mediating guideline-driven physician behavior change," *Archives of Family Medicine*, 4, 1995, 698-705.
- [h] M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R.A. Greenes, R. Hall, P.D. Johnson, N. Jones, A. Kumar, S. Miksch, S.Quaglino, A. Seyfang, E.H. Shortliffe, and M. Stefanelli, "Comparing computer-interpretable guideline models: a case-study approach," *Journal of the American Medical Informatics Association*, 10(1), 2003, 52-68.
- [i] M.K. Goldstein, B.B. Hoffman, R.W. Coleman, M.A. Musen, S.W. Tu, A. Advani, R. Shankar, and M. O'Connor, "Implementing clinical practice guidelines while taking account of evidence: ATHENA, an easily modifiable decision-support system for management of hypertension in primary care," *Proceedings of the American Medical Informatics Association*, 2000, 300-304.
- [j] D. Wang, M. Phil, and E.H. Shortliffe, "GLEE – a model-driven execution system for computer-based implementation of clinical practice guidelines," *Proceedings of the American Medical Informatics Association*, 2002, 855-859.
- [k] I.N. Purves, B. Sugden, N. Booth, and M. Sowerby, "The PRODIGY project – the iterative development of the release one model," *Proceedings of the American Medical Informatics Association*, 1999, 359-363.
- [l] P. Ram, D. Berg, S.W. Tu, G. Mansfield, Q. Ye, and R. Abarbanel, "Executing clinical practice guidelines using the SAGE execution engine. Proceedings of the International Medical Informatics Association, 2004 (in press).
- [m] A. Boxwala, S.W. Tu, M. Peleg, Q. Zeng, O. Ogunyemi, R.A. Greenes, E.H. Shortliffe, and V.L. Patel, "Towards a representation format for sharable clinical guidelines," *Journal of Biomedical Informatics*, 34(3), 2001, 157-169.
- [n] M. Peleg, A.A. Boxwala, O. Ogunyemi, Q. Zeng, S. Tu, R. Lacson, E. Bernstam, N. Ash, P. Mork, L. Ohno-Machado, E.H. Shortliffe, and R.A. Greenes, "GLIF3: the evolution of a guideline representation format," *Proceedings of the American Medical Informatics Association*, 2000, 645-649.
- [o] R.N. Shiffman, B.T. Karras, A. Agrawai, R. Chen, L. Marengo, and S. Nath, "GEM: a proposal for a more comprehensive guideline document model using XML," *Journal of the American Medical Informatics Association*, 7, 2000, 488-498.
- [p] A. Advani, S. Tu, M. O'Connor, R. Coleman, M.K. Goldstein, and M. Musen, "Integrating a modern knowledge-based system architecture with a legacy VA database: the ATHENA and EON projects at Stanford," *Proceedings of the American Medical Informatics Association*, 1999, 653-657.
- [q] S. Quaglino, M. Stefanelli, A. Cavallini, G. Micieli, C. Fassino and C. Mossa, "Guideline based careflow systems," *Artificial Intelligence in Medicine*, 20(1) 2000, 5-22.
- [r] S.W. Tu, J. Campbell, and M.A. Musen, "The structure of guideline recommendations: a synthesis," *Proceedings of the American Medical Informatics Association*, 2003, 679-683.
- [s] C. McDonald, "Protocol-Based Computer Reminders, The Quality of Care and the Non-Perfectability of Man," *New England Journal of Medicine*, 295(24), 1976, 1351-5.
- [t] Stanford Medical Informatics, Protégé Project, <http://protege.stanford.edu/>
- [u] P.D. Johnson, S.W. Tu, M.A. Musen, and I. Purves, "A virtual medical record for guideline-based decision support," *Proceedings of the American Medical Informatics Association*, 2001, 294-298.
- [v] W3C Architecture Domain Web Site: www.w3.org/2002/ws
- [w] IDX System Corporation's Carecast™, www.idx.com
- [x] Institute of Clinical Systems Improvement (ICSI) Immunization Guideline www.icsi.org/knowledge/detail.asp?catID=29&itemID=174
- [y] American Diabetes Assoc. Diabetes Guideline <http://care.diabetesjournals.org/cgi/content/full/25/1/148>
- [z] American College of Emergency Physicians, "Clinical Policy for the Management and Risk Stratification of Community-Acquired Pneumonia in Adults in the Emergency Department," *Annals of Emergency Medicine*, 38, 2001, 107-113.
- [aa] R.A. Greenes, A. Boxwala, W.N. Sloan, L. Ohno-Machado, S.R. Deibel, "A framework and tools for authoring, editing, documenting, sharing, searching, navigating, and executing computer-based clinical guidelines," *Proc AMIA Symp 1999*, 261-265.
- [bb] S.M. Maviglia, R.D. Zielstorff, M. Paterno, J.M. Teich, D.M. Bates, and G.J. Kuperman, "Automating complex guidelines for chronic disease: lessons learned," *Journal of the American Medical Informatics Association*, 2003, 10, 154-165.
- [cc] W.M. Tierney, J.M. Overhage, M.D. Murray, L.E. Harris, X.H. Zhou, G.J. Eckert, F.E. Smith, N. Nienaber, C.J. McDonald, and F.D. Wolinsky, "Effects of computerized guidelines for managing heart disease in primary care," *Journal of General Internal Medicine*, 2003, 18(12), 967-976.