Guidelines and Standard Terminology

**Robert McClure, MD, James R. Campbell, MD,
Mark A Nyman, MD, and Julie Glasgow, MD**

Overview of Authoring Automated Clinical Decision Support Guidelines

Clinical decision support systems (CDSS) have a long history that began in the 1970s with expert systems such as Mycin,[1] developed at Stanford, and INTERNIST,[2] developed at University of Pittsburg. Initial systems tended to focus on mimicking the decision making of experts and to study decision making behavior. Some systems emerged from the academic labs, but most have had limited actual clinical use. There are many reasons for the difficulty in moving CDSS from a development environment to traditional clinical environments, but one consistently cited reason is the lack of standards for representing the constructs used in the CDSS. These constructs are the information model representing the structure of patient data, the terminology (and associated terminology model), the CDSS execution engine, and the process by which any standard (sharable) representation is "mapped" into a functioning clinical information system.

The Standard-based Active Guideline Environment (SAGE) project is a five-year NIST-funded joint project[3] combining the talent of IDX (now part of GE), Apelon, Inc., Stanford University, University of Nebraska, and the Mayo Clinic, to develop a standardized guideline representation model and authoring environment. Furthermore, the group worked to create a standardized information model based on the Health Level 7 (HL7) Reference Information Model (RIM)[4] and used Consolidated Health Informatics (CHI) approved standard terminologies[5] in representing a selection of clinical guidelines as a demonstration project to support sharing guidelines among different institutions using different clinical systems.

This paper briefly reviews the approaches taken to abstract the knowledge from non-computable clinical guidelines, and identify and represent the logical decisions incorporated in the guidelines. We then discuss the impact of using standardized terminology (and the concomitant information model) upon the process of implementing the guidelines in an operating clinical system. Data specialists play a crucial role in the implementation of automated clinical guidelines at any healthcare institution intending to use this emerging technology.

Overview of Guideline Abstraction Process

Currently, all clinical guidelines are authored as paper-based or non-computable electronic documents, often created by specialty societies. Transforming the content of these documents into computerized CDSS raises issues of completeness and accuracy. Studies have demonstrated that most paper guidelines require additional clinical details to remove ambiguities and provide more complete coverage of potential patient characteristics.[6, 7, 8, 9] Several methods for developing computable guidelines for decision-support from paper-based clinical guidelines have been described in the literature.[10, 11, 12]

We briefly summarize the approach explored in the SAGE project[13] in the following sections.

Figure 1: Example of Encoding Artifacts Used in the SAGE Process to Encode a Vaccination Guideline

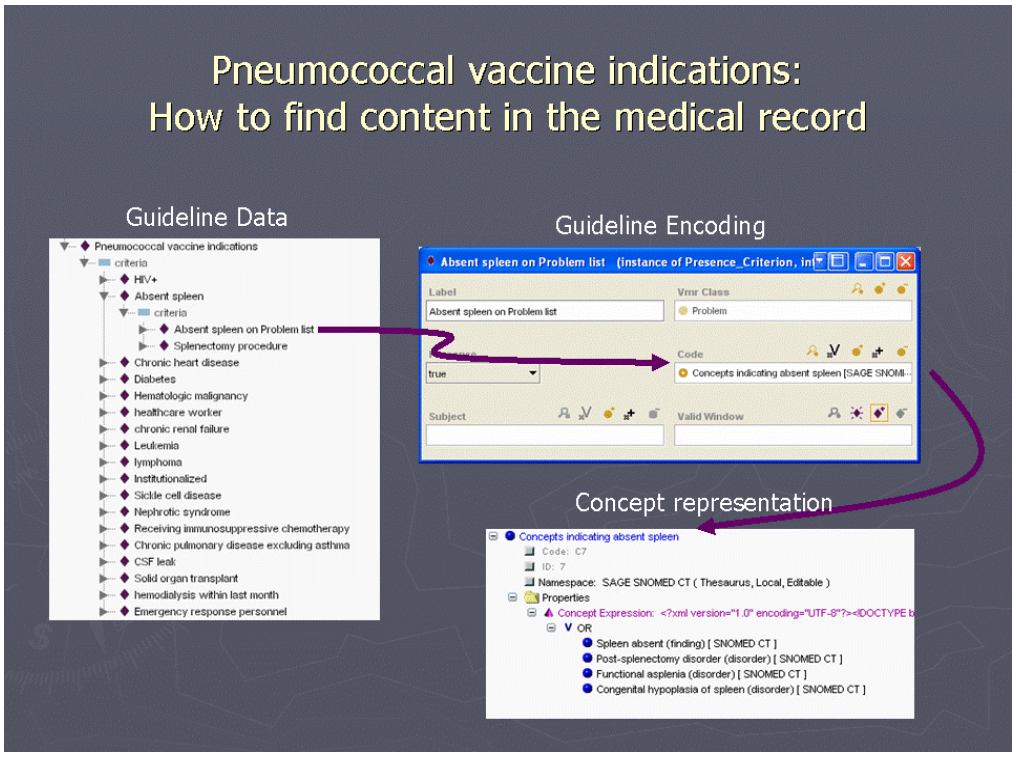

1. Creating the Implementation Scenarios

Abstraction of clinical knowledge from paper-based guidelines is enhanced by identifying typical patient scenarios for which the guideline content is clinically useful. This sharpens the focus on required clinical information and decision steps. Multiple scenarios are required for all but the simplest guidelines to cover the breadth of knowledge available. This also allows identification of gaps in guideline specificity that can either be ruled "out of scope" or enhanced with local knowledge.

2. Developing Concept Inventory: Employing and Enhancing Standard Terminologies

Using the scenarios to guide analysis of the published guideline, the next step is to identify the critical concepts required to implement the guideline (see Figure 2.) These are the data items that can be determined (or things that can be recorded) about patients that drive decisions based on the guideline. These items make up the terminology necessary to implement the guideline. The abstracted concepts should then be represented by standard terminologies to support sharing, collaborative discussion, and general interoperability. The use of standard terminologies with complex terminology models (such as SNOMED CT®) allows for the capture of substantially richer content.[14, 15] The

identification of guideline concepts should not be bounded by "what normally is in the health record" as some local concepts not in the external standard may be needed.

Figure 2: Compelling Indications for Pneumococcal Vaccine, with Two Encoded Concepts

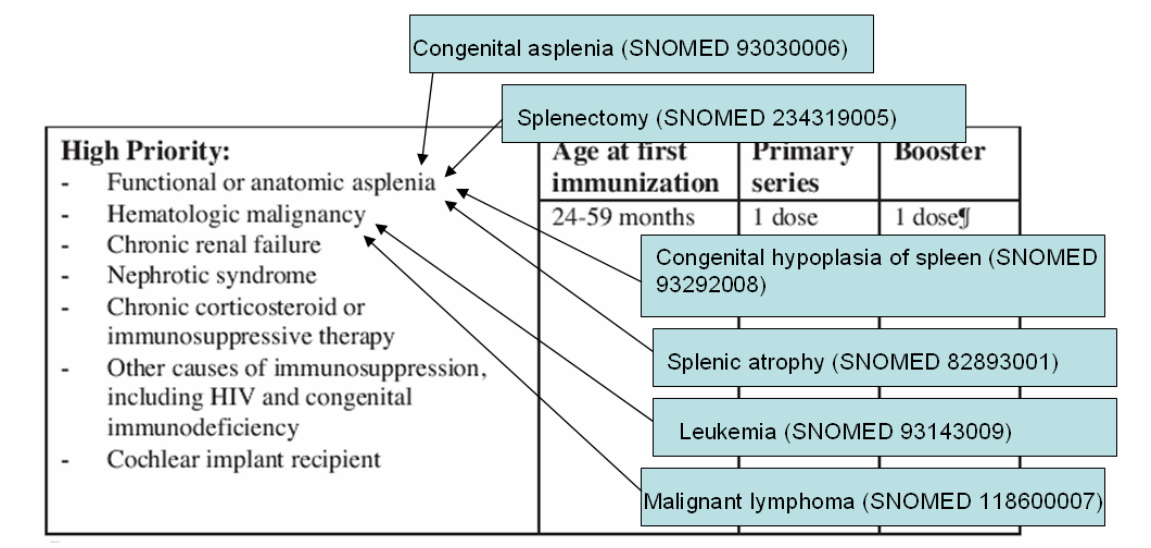

3. Assembling Decision Logic Requirements

Concomitant with identification of the guideline concept inventory, use of the scenarios helps identify the decisions that operate on these concepts to arrive at final recommendations. This represents the guideline decision workflow captured as logical IF-THEN-ELSE statements. These statements may also require critical values for the concepts identified in the prior step (code: value pairs such as Serum sodium < 130 meq/L.) Patient records satisfying these queries move on to the next action.

4. Specifying Information Queries

The logical statements described above require queries into operating clinical information systems. This requires determining how to represent patient characteristics, and diagnostic or therapeutic actions. This is the *information model* that the concepts and decision logic will use, and it describes how the information is conceptually (not physically) structured. To date, the lack of a standard information model for clinical guidelines has led many organizations to create a clinical information model tailored to their local institution, creating a model driven by local idiosyncrasies. This is the "curly braces problem" noted by implementers of Medical Logic Modules (MLM.) The SAGE project focused on removing this obstacle to sharability by basing the information model on the HL7 RIM—creating the virtual medical record (or vMR.)

5. Creating the Guideline Knowledge Base

Once the concepts, decision logic, and information model (which is hopefully relatively static over the course of time) are available, development of the actual electronic guideline can begin. This requires the use of a guideline representation format and a guideline development environment. Examples included those embedded in a specific clinical information system, general rules engines such as Blaze Advisor, or sophisticated

clinically-focused systems such as the Protégé-based SAGE clinical guideline workbench. Implementing the guideline often requires extending the knowledge available from the published guideline because of knowledge gaps or lack of enough specificity in the published document.

6. Validating the Development

As noted, the transformation of paper content to electronic automated clinical guidelines often results in significant enhancement (additions, clarifications, enhanced specificity) of the knowledge originally contained in the guideline. Validation of the content additions—the completeness and accuracy—needs to be completed by testing and clinical review. Once this is completed, the guideline is ready to be put into practice.

Overview of Standard Terminologies

On March 21, 2003, on the president's direction, the departments of HHS, Defense, and Veterans Affairs announced the first set of standards for the electronic exchange of clinical health information, known as the Consolidated Health Informatics (CHI) standards.[16] Subsequently, additional recommendations have been published. These include messaging and terminology standards with specific recommendations to use SNOMED CT® "for laboratory result contents, non-laboratory interventions and procedures, anatomy, diagnosis and problems, and nursing;" LOINC "to standardize the electronic exchange of laboratory test orders and drug label section headers;" and a combination of FDA codes, RxNORM, and NDF-RT for drug information.

CHI documents go on to say:

> Federal agencies with health-related missions need to find a way to share their health information. This health data sharing will enable them to make significant strides towards improving patient safety, reducing error rates, lowering administrative costs, and strengthening national public health and disaster preparedness. To share health data, agencies need to adopt the same clinical vocabularies and the same ways of transmitting that information.[17]

Not surprisingly, some of the CHI terminologies are based on sophisticated terminology models that allow representation of detailed clinical content without requiring enumeration of all required concepts as existing "pre-coordinated" single-coded items. Capturing concepts as complex as diagnoses, problems, and procedures using a standard terminology within a specific semantic model (such as SNOMED CT® compositional expressions[18]) means that organizations creating guidelines can capture the clinical meanings required with real hope of interoperability when exchanging data across different systems—or even between institutions.

Identifying Terminology Concepts in Guidelines

As noted above, abstracting important concepts from guidelines can be difficult, an effort helped by creating clinical scenarios to focus the process. In addition, starting with a clear understanding of which standards will be applied sets the stage for clearer delineation

between the three interdependent components of the information model, the terminology model, and the execution model.

The information model is the type of information (meta-information) recorded in the clinical record: "This is an observation that occurred at this time, on that person." For the SAGE project, this was represented by the vMR and is discussed separately.[19]

Terminology Model

The terminology model defines the relationships, and therefore, the relative meaning of concepts in the terminology. In SNOMED CT® the primary relationship is the hierarchical parent-child (IS-A) relationship that means the child is a more specialized type of the parent. SNOMED also allows multiple hierarchies, so a concept can have more than one parent. As a compositional terminology model, SNOMED also has "definitional roles," which are named relations between concepts used to make explicit the defining attributes of a focus concept—for example, Pneumococcal pneumonia has an anatomic location of Lung. Capturing all these relationships creates a "compositional expression" that is constrained by the allowed semantics defined by the SNOMED editorial board. For example, SNOMED semantics allow Findings to have an *Anatomic location* but not an *associated Specimen* (which is used by Procedures).

Interaction of the Terminology and Information Models

The information model defines the "properties" about the concepts we are going to use. Drugs provide a relatively stable example for this. A medication order (either delivered or ordered) would have an information model that defines the following properties (with the corresponding expected SNOMED CT® content represented in {}):

- Code: {Product|373873005}
- Priority code: {Procedure modifier|106239005}
- Route code: {Route of administration value|2840099009}

Adverse reactions also provide a relatively straight forward separation of the information model and terminology model:

- Code: {Adverse reaction|281647001}
- Reaction: {Clinical finding|484684003}
- Severity: {Severities|272141005}
- Substance: {Product|373873005}{Substance|105590001}

With models (information and terminology) as complex as these, overlap of meaning and use can occur between the models. For example, both the RIM (on which the vMR is based) and SNOMED can represent that a diagnosis is part of the Family History and not something the patient actually has. Observations are also potentially confusing since typical information models (such as the RIM and the vMR) represent Observations as the following pair of "slots":

- Code: {Clinical finding|484684003} {Observable entity (SNOMED CT® or LOINC)}

- Value: with observable - {Clinical finding|484684003}{Qualifier value|362981000} {Numerical value}

SNOMED CT® specifically allows the combination of concepts that can be sent in the code and value into a single compositional expression. In fact, many existing finding concepts have already done this. Separating these components consistently in the guideline model (or when mapping local content to the guideline) can be problematic and is only possible when the terminology model supports clear rules for combinations (semantics.)

Conflicts such as these are managed by defining boundaries. For SAGE, this meant using the vMR to identify a separate property for family history with the expectation that this aligns with most clinical information systems that store family history findings separate from patient findings. The HL7 TermInfo project[20] is dealing with similar issues in an attempt to define a *standard* way of managing SNOMED CT® and RIM overlap.

No matter if information is primarily represented in complex concepts or if it is broken out into multiple unique properties, "translating" between different approaches is possible if we use standards that are defined using published models. This is a critical benefit of the adoption of these sophisticated standards.

Terminology Services

Terminology services provide a set of functions that support the use of complex terminologies such as SNOMED CT®. Traditional terminology services include:

- Code to/from display name
- Synonymy
- Determine all concepts grouped (subsumed) by another concept (called subsumption in Description Logics)
- Version/history information for the terminology (track concepts over time)

In addition, for SNOMED we can also:

- Determine a set of concepts defined by Boolean construction of SNOMED concepts
- Create and maintain local extensions to the SNOMED standard
- Determine compositional expressions

Use of these services is important for terminologies used in clinical guidelines. Guideline constructs often use a single concept to represent all concepts "of this type, or subsumed by this concept" when identifying candidate patient characteristics. Guideline authors will use a single concept in the guideline as a "placeholder" for which any specialization of that concept is a valid substitution when representing a particular patient instance. For example: "If the patient has pneumonia, then do X." Determining if an actual patient record meets this requirement requires a terminology service that can return a list of all valid children for the concept *pneumonia* used in the guideline. In the absence of this subsumption capability, the guideline author must exhaustively list all types of pneumonia—a prohibitively large set of diseases.

Execution Model

The execution model is the final component required for implementing guidelines, and it defines how the guideline knowledge base interacts with a functioning healthcare information system. This is the running program that operates on the guideline "file" and translates requests for information, based on the information model, into standard queries (with defined APIs). The SAGE project enhanced CDSS interoperability by defining a standardized execution model without embedding operational activities, such as database calls and physical data layouts, into the guideline model.

Matching Ideas to Standards

The use of standards means that clinical guidelines can be rendered computable independently of local implementation considerations. Guideline authors (or abstractors) can complete their work knowing that local implementers, separated by space and time, will be able to review the abstracted guideline and see standard representations of terminology, defined by an explicit terminology model. The implementation team can more easily:

- Review the required concepts and determine if the content is captured in their system.
- Review the information model as represented by the vMR and the specific constructs needed for the particular guideline and determine what database tables hold the information.
- Write to general execution engine APIs so the guideline can run directly in their system's environment without having to build a guideline execution application.
- In an ideal situation, where the institution already captures content using the same standards, implementing the guideline would require only local vetting of the content and writing local queries for the execution engine APIs—queries that would work for every guideline implemented from that point forward. However, institutions using non-standard terminologies would be required to "map" each term to the standard terminology base.

Breaking the problem of abstraction into components that fall into the models noted helps segment the work. Basing the models on standards means there can be shared understanding (interoperability) with explicit knowledge. When guideline abstraction uses standard terminologies and models, it should be sharable among institutions that choose to write links into the local environment from this standard intermediary.

Making the Guideline Work Locally

Defining a guideline using standard terminology and information model components provides an interoperable platform for institutions to begin the process of local implementation. By using an externally defined standard, implementers are working with "known constructs" that have explicit definitions, decreasing the amount of ambiguity normally found with implementing external guidelines. While implementers must complete a number of important tasks, including clinical vetting of the content and

alignment of the actual clinical information with the vMR (including technical implementation of the execution engine), issues involved in mapping local content to the standard terminology are often the most vexing. “Where am I going to find that?” is the common refrain.

The Local Mapping Process

Guidelines developed using this process produce a series of concepts used within the context of particular components of the information model (or vMR for SAGE guidelines.) Local implementers must identify the concepts and determine where the information is stored in their electronic record and with which codes (or phrases). They must identify model mismatches and determine how to adjust for that change (often using modifications in the execution engine.) For example: AGE may be used in a guideline but the local institution stores the patient Date of Birth and calculates the AGE as necessary.

Local mapping begins by obtaining a complete listing of concepts used in the guideline. These must be reviewed and the corresponding local tables (based on the identified context) identified. The implementer must search these tables to find equivalent concepts.

Rarely, implementers will have access to a terminology service to manage these mappings. If so, local concepts may be maintained in groupings and hierarchies that will simplify mapping aggregates.

Guideline concepts typically fall into the categories below.

One-to-One Matches

A guideline concept has only one local concept match for the context defined. Even if the guideline concept is meant to represent a hierarchy of children, if the local institution has only one concept to represent that meaning, it is a simple one-to-one mapping. Orders are an important example of this. When the SAGE guideline engine needs to set-up an order for the user, this must be mapped to a specific orderable item within the local clinical information system.

One-to-Many Matches

Often a guideline concept is captured locally in more than one way because either there are multiple systems using the concept, there has been a historical progression of concepts used to represent the idea (for example, multiple drug vendors for the same product), or the local system requires more detail (for example, different clinics use specialized names for the same general procedure.) Examples include the following:

- When querying the EMR for a particular lab value (hemoglobin A1c), all potentially clinically equivalent values (that is, both measured hemoglobin A1c and calculated hemoglobin A1c) should be mapped to this one query term.
- When querying the EMR to see if any potentially clinically equivalent orders are pending (to prevent ordering of duplicates), a many-to-one mapping is necessary (for example, a pending order for random urine Microalbumin, 24-hour urine Microalbumin, or 24-hour urine protein would all be mapped to the same term to prevent ordering another urine protein test for diabetes.) See Figure 3.

Figure 3: Differences in Mapping between a Query (top) and an Order (bottom)

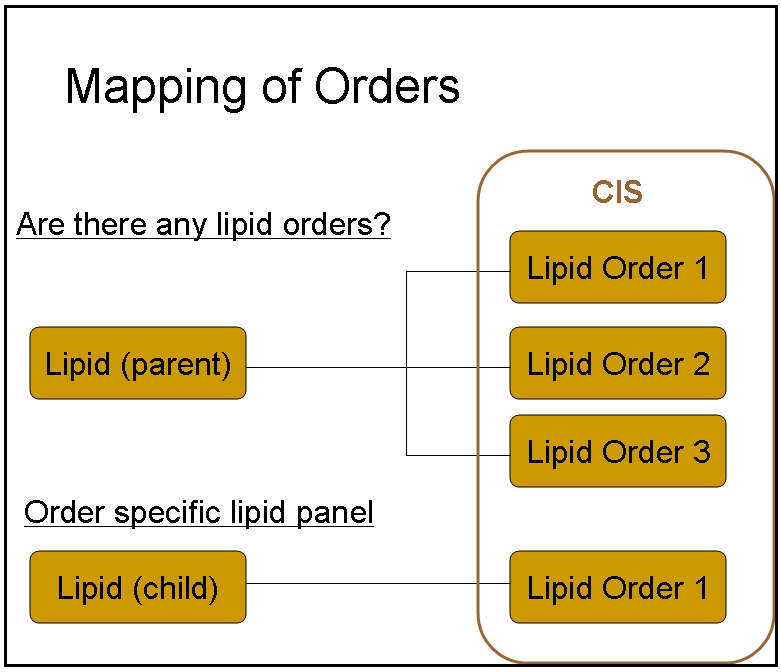

Matches for Concepts that Act as "Collectors"

A concept that collects other concepts (and therefore for some terminology models subsumes these other concepts) into a hierarchy can be considered a collector, or generalization. These concepts can be treated as one-to-many matches. For implementations with terminology services, matches for collectors may be done by mapping a collector concept in the guideline to a local hierarchy concept. If this is not available, then a one-to-many mapping must be used, but it should be restricted to the particular context identified so that inappropriate matches will not occur.

Missed Matches

Sometimes concepts are used in guidelines that are simply not available in a particular local system. These can be patient characteristics or they can be properties of the settings or workflow events. Even so, if they do not exist, getting the guideline to operate can be difficult.

Missing terminology matches often occur because the information is simply not captured in the clinical record in structured format. The capture of structured clinical data into electronic records is still in its infancy. Much of the information important to clinical care is not recorded electronically, and when it is, it may be done in a non-discrete manner (hidden in a text blob.) Unfortunately, that does not mean a guideline author does not need this information to define an outcome and, given the expressivity in some standard terminologies (such as SNOMED), they may render the required information using a "standard concept." If data is required (cannot take on a assumed value), then

implementers must somehow ask the user for the information during execution, either by demanding user input or abstracting content from other sources, although too many interactions with users can make the system "unfriendly."

Conclusion

Using standard terminologies and standard models in clinical decision support guidelines provides enhanced interoperability and eases local mapping tasks. Standard concepts provide consistent representations that can be relied upon when moving knowledge from guideline developers to local implementers and eventually to clinical users. Abstraction of non-electronic guidelines to a computable electronic format highlights the importance of consistent and explicit definitions of meaning when creating a guideline—many guidelines written today present significant difficulties for transformation to an electronic format due to the inclusion of clinical generalities not well supported in an automated environment.

Endnotes

---

[1] Shortliffe, E. H. (1976). *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier/North Holland.

[2] Miller, R. A., Pople, H. E. J., et al. (1982). "Internist-1, an experimental computer-based diagnostic consultant for general internal medicine." *New England Journal of Medicine* 307(8): 468–476.

[3] SAGE is supported under a grant from the US Department of Commerce, National Institute of Standards and Technology, Advanced Technology Program, Cooperative Agreement Number 70NANB1H3049.

[4] Health Level Seven, Inc. *HL7 Reference Information Model.* Ann Arbor, MI: Health Level Seven, Inc., 1994. Available from http://www.hl7.org/library/data-model/RIM/modelpage_non.htm. Accessed June 6, 2006.

[5] http://www.whitehouse.gov/omb/egov/c-3-6-chi.html. Accessed on June 6, 2006.

[6] Hayward, RS, EP Steinberg et al. (1991). "Preventive care guidelines: 1991." American College of Physicians. Canadian Task Force on the Periodic Health Examination. United States Preventive Services Task Force. *Ann Intern Med* 114(9): 758–83.

[7] Tierney, WM, and JM Overhage et al. (1995). "Computerizing Guidelines to Improve Care and Patient Outcomes: The Example of Heart Failure." *JAMIA* 2(5): 316–332.

[8] Patel, VL, JF Arocha et al. (2001). "Methods of cognitive analysis to support the design and evaluation of biomedical systems: the case of clinical practice guidelines." *J Biomed Inform* 34(1): 52–66.

[9] Iqbal, A, S. Schloss et al. (2002). "Worldwide guidelines for chronic obstructive pulmonary disease: a comparison of diagnosis and treatment recommendations. *Respirology* 7(3): 233–9.

[10] Tu, SW, MA Musen et al. (2004). *Modeling Guidelines for Integration into Clinical Workflow*. San Francisco, CA: Medinfo, 174-178.

[11] Shahar, Y, O Young et al. (2003). DEGEL: A hybrid, multiple-ontology framework for specification and retrieval of clinical guidelines. AIME, Protaras, Cyprus.

[12] Shiffman, RN, G. Michel et al. (2004). "Bridging the Guideline Implementation Gap: A Systematic, Document-Centered Approach to Guideline Implementation." *J Am Med Inform Assoc* 11: 418–426.

[13] Tu, SW, MA Musen et al. (2004). *Modeling Guidelines for Integration into Clinical Workflow*. San Francisco, CA: Medinfo, 174–178.

[14] AL Rector, JE Rogers. *Ontological & Practical Issues in using a Description Logic to Represent Medical Concepts: Experience from GALEN*. School of Computer Science PrePrint, University of Manchester: 2005; CSPP-35:1–35.

[15] Elkin PL, Brown SH, Husser CS, et al. "Evaluation of the content coverage of SNOMED CT®: ability of SNOMED Clinical Terms to represent clinical problem lists." *Mayo Clin Proc*. 2006; 81:741–8.

[16] http://www.whitehouse.gov/omb/egov/c-3-6-chi.html. Accessed on June 6, 2006.

[17] Ibid.

[18] Spackman KA, and KE Campbell. "Compositional concept representation using SNOMED: towards further convergence of clinical terminologies." *Proc AMIA Symp* 1998:740–4.

[19] PD Johnson, SW Tu, MA Musen, and I. Purves. "A Virtual Medical Record for Guideline-Based Decision Support." *Proc AMIA Symp*. 2001:294–8.

[20] http://www.hl7.org/Special/Committees/terminfo/index.cfm. Accessed on June 5, 2006.