University of Utah Division of Geriatrics/CogniTech Response to OMG/CORBAmed Request for Information (RFI) on Decision Support Services (RFI3)

August 19, 1997

author: Jerome B. Soller, PhD, Research Instructor, University of Utah Division of Geriatrics (OMG university member) and President, CogniTech Corporation

email: jerome.soller@m.cc.utah.edu (university),
soller@cognitech-ut.com (business)

Telephone: (801) 582-1565, x-2469 (university lab),
(801) 322-0101 (business)

Facsimile: (801) 583-7338 (university),
(801) 322-0975 (business)

I. INTRODUCTION

This paper describes the requirements for services to support decision support systems in a distributed object system conforming to the Object Management Architecture (OMA) in Response to CORBAmed's Decision Support RFI. A decision support system transforms raw data into information which enables, supports, or enhances a person's decision-making processes. Computers serve as a vehicle for providing decision support, such as knowledge, recommendations, warnings, and cost/utility analysis. The user still makes the ultimate decision of how to respond to the decision support. For this discussion, a model uses its internal content to transform raw or abstracted data into a representation usable by a decision support system. The engine is an implementation of a given model.

In healthcare and other domains, there are various dimensions to categorize decision support systems (e.g., response time, mode of user interaction with the decision support, model type, target domain, and access to content). Although many decision support systems have been developed as standalone systems, the author believes their utility increases through seamless integration within business processes and existing information systems. In the past, the most successful implementors of decision support systems tightly coupled the decision support implementation for one model or representation with their internally developed information systems. Today, decision support systems must operate within Integrated Delivery Systems (IDS), Health Data Networks (HDN), and multi-enterprise organizations across a possibly heterogeneous computing environment. The decision support systems must interoperate with both commercial off the shelf (COTS), contractually developed, and internally developed decision support systems, clinical data repositories (CDRs), and information systems. The requirement for sharing of data (e.g., clinical data, decision support content, measures of compliance) in a consistent, common, and clinically specific manner takes on a higher level of complexity than ever before.

Today, many organizations developed different decision support systems to serve different purposes (e.g., generic alerts and warnings, care plan management, diagnostic assistance, drug management, eligibility and immunizations). No single decision support model, engine, content, or target application domain/purpose will address the decision support needs of all users. A common set of CORBA-based decision support services will enable access to heterogeneous decision support resources to support these various requirements.

Although these recommendations focus upon the healthcare domain, many are generic to all decision support systems. This may lead to these services required across other domains/task forces within OMG. However, this document and the resulting RFP should focus upon the needs of CORBAmed.

The goals of this paper are:

- Serve as input to the development of a Request for Proposal (RFP) on decision support services in OMA compliant object systems
- Educate OMG members on need for well defined, robust decision support services, and how they would function within a distributed object system
- Solicit involvement of interested parties in the development of the lexicon services requirements to ensure that they meet the needs of users of these services, and do not have unacceptable effects on other OMG requirements such as usability and performance.

The requirements are divided into several sections, which include the following:

1) definitions
2) prerequisites for decision support implementations
3) classification/categorization of decision support dimensions
    a. response time
    b. user interaction
    c. generic model types
    d. healthcare specific model types
    e. access to decision support content
4) healthcare standards for decision support systems
    a. Arden Syntax
    b. Intermed Guideline Interchange Format (GLIF)
    c. Knowledge Query and Manipulation Language (KQML)
    d. Knowledge Interchange Format (KIF)
    e. Ontilingua and Shared Ontologies
    f. Health Level 7 (HL7) Decision Support Special Interest Group (SIG)
5) general object services
    a. taxonomy of classes of decision support services
    b. multiple model support
    c. generic decision support services
    d. modes for triggering alerts
    e. data transfer
6) potential focus areas
7) relationship to existing OMG specifications
8) presentation logistics.

The CORBA RFP specifications and standards resulting from this paper will support robust decision support implementations across multiple, heterogeneous systems. A common definition of interfaces for these decision support services would enable engine vendors, content vendors, information systems vendors and customers to create, integrate, utilize, and evaluate a broad spectrum of healthcare decision support solutions.

II. Definitions and Examples

A decision support system transforms raw data into information which enables, supports, or enhances a person's decision-making processes. Computers serve as a vehicle for providing decision support, such as knowledge, recommendations, warnings, and cost/utility analysis. The user still makes the ultimate decision of how to respond to the decision support. For this discussion, a model uses its internal content to transform raw data into a representation usable by a decision support system. The engine is an implementation of a given model. Examples of models include knowledge-based systems with a given representation syntax, artificial neural network architectures, fuzzy logic, and discriminant classifiers. Examples of decision support content include a set of related rules for a domain within an expert system, thresholds and parameters for a discriminant classifier, or weights within an artificial neural network architecture. Examples of engines include inference engines, expert systems shells, statistical packages, and commercial neural network packages.

The following is a scenario of real-time decision support system interaction with clinicians. A nurse starts to record the information for a patient encounter and a warning appears on the computer, "The patient may be delirious because he is confused. Does he know his name?" The nurse is prompted for a response, and the nurse selects "No." The nurse is now presented with the message, "The patient appears to be delirious. Contact the physician. Do you wish the physician to be automatically paged?" The nurse chooses not to page the physician because he is treating a different patient on the same ward. The nurse contacts the physician with the patient concerns and documents her compliance to the computerized recommendation. The physician confirms the diagnosis of delirium and an available computer generates a list of possible causes. The physician identifies the cause, treats the cause and symptoms of delirium, and documents his actions.

III. Prerequisites for the Successful Implementation of Decision Support

One of the most fundamental requirements for effective decision support is consistency among data representations for model inputs, the results of applying the models, presentation to the user, measures of compliance, and outcomes measures. Some systems may choose an implicit representation (i.e., no formal lexicon). If the enterprise has only one standard vocabulary and the system does not need to map or interact with other vocabularies or systems with different vocabularies, this may be adequate. However, in the case of multiple vocabularies or lexicon implementations, the CORBA Lexicon Services will attempt to address data representation issues. Therefore, this paper will leverage those services and focus specifically on the needs of decision support services, assuming consistent data representations.

Decision support specifications can be represented at various levels of abstraction (e.g., literature, guideline or other high level specification, reusable representation, software implementation). There are many challenges in converting formal, high level specifications into clinically useful software. Clinical guidelines provide an important healthcare decision support function, which are sometimes implemented through knowledge-based systems. The Association for Healthcare Policy Research (AHCPR) is creating a national, web-based repository for these specifications [1]. Despite the increasing availability of published clinical guidelines, the successful utilization of these guidelines faces many challenges. Characteristics which impact guideline utility include explicit criteria [16], validity, their use within the care process, and compliance. Precise criteria for clinical pathways are necessary to validate their impact on outcomes. A common representation for complex guidelines would solve the precision problem[31], but no widely accepted standard currently exists. Any CORBA standard for a given decision support model (e.g., subset of guidelines) should follow or at least be coordinated with the development of widely accepted standards or common use of those technologies, not precede their development.

IV. Classes of Decision Support Systems

The categories and classification of decision support capabilities provide the most fundamental criteria for the organization of CORBA healthcare decision support objects and services. Decision support systems can be categorized by the dimensions of response time, mode of user interaction with the decision support, model type, and access to clinical content.

IV.A. Response Time

Batch systems are not interactive in real-time for clinical care (e.g., seconds or minutes) and generally can not change their pathway or approach until they have completed. Real-time decision support systems, the focus of this RFI response, are designed to provide feedback to the user in an interactive and real-time fashion, which allows new data to direct the process. Real-time and batch decision support systems are complementary, and a real-time decision support system could utilize the calculations of a batch decision support system. The OMG DSS RFI [28] further subdivided batch systems into prospective and retrospective categories. In reality, there is a continuum between real-time and batch systems.

IV.B. User Interaction

Another general distinction among decision support systems is between user initiated knowledge retrieval and models that provide active decision support (e.g., warnings, recommendations, risk factors). This categorization determines the external data access mechanisms (e.g., queries, store and forward, publish/subscribe, batch).

User initiated knowledge retrieval provides the clinician with additional clinical or science knowledge on demand or through the specification of knowledge filters. Examples include:

1)context sensitive help screens;

2)tools allowing a user to request assistance in interpreting some subset of the patient's data (e.g., a module that provides additional diagnostic or therapeutic knowledge based on a user-selected patient problem, clinical finding, laboratory abnormality, or active medication);

3)modules that provide global textual references on demand (e.g., guidelines, antibiotograms, administrative policies, OSHA regulations); and

4)patient / patient family education providing knowledge to support decision making.

Case-based retrieval provides a different form of knowledge retrieval. Based on a user defined or predefined search criteria measuring the similarity with the current case, previous cases can be retrieved with all of their associated information (e.g., structured data, unstructured text documents, images). In addition to automatically ranking these previous cases to assistant the clinician in their decision-making, case-based decision support can transform the results of the case-based retrieval into active decision support.

Active decision support requires a response from the clinician, such as the acknowledgement of an alert, entry of additional information, or a documented response to an alert or warning. Usually, this form of decision support is not directly invoked by a clinician, but is triggered by the entry of clinical data or other events.

IV.C. Generic Model Types

Decision support systems can be categorized by the computational model (e.g., algorithm, problem solving method, statistical classifier category, knowledge representation syntax, artificial neural network architecture). The decision support content provides the specific parameters, rules, knowledge, weights and other information which enable a generic, computational model to transform raw or abstracted data into a form usable by a healthcare decision support system. Decision support engines provide an implementation of a decision support model, which manipulates the content and enterprise data.

From the semantic perspective, decision support models can be further categorized as shallow models and deep models. Shallow models are based on associations. These associations result from heuristics (rules of thumb) or data generated models based on a performance measure (criteria). Deep models attempt to directly capture the physiological process at some level of abstraction. This deep information is usually represented by a graphical model (i.e., semantic network). This discussion will focus upon shallow models.

Given evidence (predictor data), data generated classification models attempt to either provide a nonprobabilistic measure of strength for a classification (e.g., fuzzy logic measure of association), a probabilistic estimate of the distribution given prior knowledge (e.g., posterior probability produced by a Bayesian model), or the recognition of the most probable class (e.g., maximum likelihood). Examples of statistical classification models include first order Bayes classifiers, discriminant classifers, logistic regression, nearest neighbor classifiers, Parzen windows, Bayesian belief networks, classification and regression trees (CARTs), and many variations of artificial neural networks.

Knowledge-based systems utilize encoded knowledge to support decision-making. The process of capturing and formalizing this knowledge from experts in a given domain is called knowledge engineering. Once a model is formalized through knowledge engineering, it can be separately implemented. Although many systems utilize a single decision support engine, these systems can also be implemented through interacting components. One of the most fundamental issues in knowledge-based systems is their reuse across multiple sites and even information systems. Pathophysiological models will not have a large role in defining careplans and guidelines. Therefore, we focus on the model complexity and the best formal representations for those models to allow long-term maintenance and software reuse.

IV.D. Healthcare Specific Model Types

The previous subsection summarized many generic classes of decision support models. Whenever possible, it is desirable to categorize healthcare decision support models by more generic models, whose use is not limited to healthcare. From a software interface perspective, this will facilitate the reuse of software classes and services. As discussed in the potential focus areas section (section V), some healthcare functionality may not be easily categorized by a more generic decision support model, but only by a set of common services and facilities. In some cases, the model, engine, and content may be inseparable. Furthermore, the encapsulation of proprietary models and their content, whether tightly coupled / inseparable or modular, may provide vendors with a competitive advantage.

IV.E. Access to Decision Support Content

Decision support systems can be categorized by the encapsulation of the decision support content. Decision support systems with proprietary content do not allow external access to content. For content developers, this provides a competitive advantage. "Open" access mechanisms for other content (even from other areas of the same decision support system) may allow access to a uniform representation of the content; the ability to create, customize, and modify the content; and the ability to reuse content across different software tools with common interfaces. Open access mechanisms to clinical content could facilitate multi-site clinical trials and the widespread dissemination of content. This requires the adoption of standard models and representations, as well as services to expose the content to external systems. Although decision support systems may adopt common, standard models, this does not require a common implementation. Vendors may distinguish themselves through the performance characteristics of their implementations.

Explanation capabilities complement the direct exposure of clinical content. Explanation capabilities include textual explanation and on-line references. The draft version of the 1998 Arden syntax standard [3] and the Interface Definition Language (IDL) version of the Guideline Interchange Format (GLIF) [6] provide examples of possible interfaces for reference material. The CORBA standard should combine the best features of the Arden and GLIF reference information formats.

V. Healthcare Standards for Decision Support Systems

V.A. Arden Syntax

Hripcsak, Clayton, and Pryor created a procedural computer language called the Arden syntax [14] for encoding guidelines, particularly situation-action rules. The American Society for Testing and Materials (ASTM) adopted Arden syntax as a standard for medical knowledge representation [2] [3]. The resulting guideline implementations are called Medical Logic Modules (MLMs). Arden syntax currently addresses the lack of a universal query language or data model for healthcare through the use of "curly-braces" [33], which allow site-specific implementations of the vocabulary, query, and information models. This increases portability, but requires local customization. However, since the CORBA lexicon services [25] [26] will be separate from decision support services, the “curly braces” may prove advantageous. Although Arden syntax has limitations, it addresses a very basic, but important, category of decision support systems.

V.B. Intermed Guideline Interchange Format (GLIF)

More complex clinical guidelines require complex high level control mechanisms. With the goal of reuse of complex, clinical guidelines software across institutions, the Intermed Collaboratory [32], a consortium of Harvard, Columbia, Stanford, and other institutions, created the InterMed Common Guideline Model and Guideline Interchange Format (GLIF) [6] to address more complex guidelines. The InterMed Guideline Representation Specification was written in CORBA Interface Definition Language (IDL) syntax. Although not a public standard, this provides an example of an IDL specification for a healthcare decision support system.

V.C. Knowledge Query and Manipulation Language (KQML)

KQML [10] is a language and protocol for exchanging information and knowledge. Developers can use KQML to create applications to interact with intelligent software agents or for the sharing of knowledge among agents. Prototypes of this technology have been demonstrated in concurrent engineering, intelligent design and intelligent planning and scheduling.

V.D. Knowledge Interchange Format (KIF)

Knowledge Interchange Format (KIF) [11] is a very expressive language for the interchange of knowledge.

- “It has declarative semantics (i.e. the meaning of expressions in the representation can be understood without appeal to an interpreter for manipulating those expressions);
- it is logically comprehensive (i.e. it provides for the expression of arbitrary sentences in the first-order predicate calculus);
- it provides for the representation of knowledge about the representation of knowledge;
- it provides for the representation of non-monotonic reasoning rules;

- and it provides for the definition of objects, functions, and relations [9]."

V.E. Ontilingua and Shared Ontologies

An ontology defines the representation vocabulary for a shared domain of discourse (i.e., classes, relations, functions). Ontlingua [8] is an example of a technology for building shared ontologies. Although decision support systems may interface with ontologies, the author believes the CORBA Lexicon Services should address ontology integration.

V.F. Health Level 7 (HL7) Decision Support Special Interest Group (SIG)

The Health Level 7 (HL7) Decision Support Special Interest Group (SIG) [4] is

1) exploring the use of the HL7 Reference Information Model (RIM) with decision support systems,
2) extend RIM to support objects with decision support functions, and
3) add decision support elements to RIM.

VI. General Object Services

VI.A. Taxonomy of Classes of Decision Support Systems

There are common services and external interfaces for any decision support model. A taxonomy could represent the additional services, data requirements, and interfaces for decision support models. This allows the reuse of more generalized services within CORBA services designed for a more specialized form of decision support.

VI.B. Multiple Model Support

CORBA decision support services must allow the manipulation of multiple models, vendor implementations for those models, and sources of structured content.

VI.C. Generic Decision Support Services

The generic decision support services are divided into Model/Engine Manipulation services, Full Life Cycle of Decision Support Content for a Model (authoring services), Run-time Services and a Persistent Log of Decision Support Services. In this discussion, the model uses its internal content to transform raw or abstracted data into a representation usable by a decision support system. The engine is an implementation of a given model. Within this section, the author does not address the handling of exceptions (run-time errors), but any standard must also define standard exceptions.

The run-time services assume the existence of applications, which are knowledgeable about the presence of decision support engines and associated capabilities. The run-time services should include

a. run-time model and engine selection
b. run-time content selection (if appropriate) data transfer to decision support from external systems,
c. run-time engine invocation,
d. access to decision support results,
e. run-time access to content or explanations of that content, and

f. run-time manipulation of user responses to active decision support.

Since the run-time services must interact with engine implementations of models, some obvious engine-questions are

- Is there a separation between the model, engine, and content for a given decision support system?
- If so, how are the different types of models reused?
- Is there a taxonomy to support reuse?
- What models and associated engines are stored?
- How are versions managed?

To address these issues, model/engine manipulation services should provide

a. standard models and taxonomy of models,
b. model and engine naming and object registration,
c. loading, updating, and deletion of engines, and
d. version control.

Services to support the manipulation of content (e.g., authoring services) assumes the standardization of models, the interfaces to those models, and the packaging of their content. These should follow the standardization of the run-time services. Services for the full life cycle of decision support content for a model include

a. creation using a content creation tool (e.g., knowledge acquisition tool, statistical package, batch decision support system),
b. run-time representation, and
c. editing/maintenance.

The evolution of the CORBA Event-Condition-Action Rules Management Facility [15] [21] may dramatically impact any future knowledge-based systems content services. Until the standard is finalized, it is unclear whether it will address these needs.

VI.E. Modes for Triggering Alerts

The following modes for triggering alerts are a subset of those needed for decision support services.

1. One provider directly invokes a warning or alert for another provider.
2. The entry of data or execution of another event invokes a warning or alert.
3. An intelligent agent acts upon persistent data to invoke a warning or alert.

VI.F. Data Transfer

Even with the ideal, modular architecture, designers still must address how data are accessed by decision support modules. In batch decision support systems, retroactive validation of decision support models, the development of decision support models, and the real-time implementation of models requiring all of their input at once (i.e., artificial neural networks, statistical models that look for outlier patient data), one encapsulated data query is sent to the server. In some forms of knowledge-based systems or decision trees, all of the possibly relevant data need not be made initially available. However, the process of model execution (e.g., rule matching and firing in a knowledge-based system, traversing the tree in tree-based models) may require further queries to the database. In these cases, the knowledge-based system or decision tree must maintain its state information while interacting with the database and interface through object oriented messaging. In either case, the recording of compliance will require further database operations.

VII. Technology Challenges for Distributed Systems

The decision support implementations must address several challenges across a distributed, heterogeneous environment, which impact the service requirements. These include the following questions, which are decision support adaptations of 3M's questions within their response to the Lexicon RFI [25].

- How does one support distributed implementations of decision support models and content?
- What kind of support does one need to address performance requirements in a distributed world (i.e., cache subset of the decision support content), and is that a CORBA issue, or an application/implementation issue?

VIII. Potential Focus Areas

While it is important to address specific focus areas, this must be based upon a unified architecture for generic healthcare decision support systems. Furthermore, many of CogniTech's content partners expressed concerns with the structure of a potential OMG standardization. The granularity of these CORBA objects and services, as well as what facilities are mandatory or optional, impact the breadth of compliant commercial products. For example, an RFP or standard requiring support for all functionality within a clinical domain (i.e., drug management) will unfairly impact an emerging small content business, which focuses upon one specialty within that domain. Any standard in this area must allow compliant products to provide selected subsets of relevant clinical information services.
Under real-time category, the CORBAMED RFI lists
1)care plan management,
2)diagnostic assistance,
3)drug management,
4)eligibility, and
5)immunizations.
Conspicuously absent are generic warnings and alerts based upon simple situation-action rules acting on entered data.
CogniTech Corporation will respond to RFPs for general architectures, healthcare-specific situation-action rules, care plan management, drug management (if properly constrained) and diagnostic assistance. However, it strongly believes that the general architecture must be addressed before specific healthcare processes (e.g., drug management).

IX. Relationship to Existing OMG Specifications

One of the strengths of CORBA is its direct support for object inheritance, which allows the reuse of objects, services, and facilities. In this paper, there are two general categories: general services/facilities [22] and vertical healthcare domain services/facilities [27].

IX.A. General CORBA Services and Facilities

Many of these services and facilities are described at a high level within the OMG Information Management Common Facilities [22] or within their respective specifications.

CORBA Security - Any patient data, including decision support recommendations resulting from raw or abstracted clinical observations, are highly confidential information. It is expected that proposals will use CORBA Security as required to support this requirement.

CORBA Event Services - The use of event services is required to register and raise events to trigger the decision support engines, notify decision support engines of the availability of their data inputs, and to notify applications using the decision support results upon the completion of the engines.

CORBA Naming Services - The decision support services require a mechanism for supporting standard naming of models, vendor engine implementations (to allow multiple implementations of the same model), and content. This prevents naming conflicts and unforeseen side effects. CORBA naming services must support the registration of these names.

Meta-Object Facility (MOF) - Submitter should address the relationship between the decision support constructs and the meta-meta-model specifications of the MOF, where appropriate.

OMG CORBAServices Lifecycle Service -This will serve as the foundation for content management and possibly engine management.

OMG CORBAServices Query Service - Decision support services may require queries of data used by their models.

OMG CORBAServices Transaction Service - For processing decision support in the context of transactions.

Business Object Facility (BOF) - Submitter should address as appropriate. The developers of the BOF should consider the addition of a taxonomy for decision support models.

Object Analysis & Design (OA&D) - Submitter should address as appropriate.

At the time of the issuance of the Decision Support Services RFP, the following Facilities will still be within the process of adoption. Therefore, submitters to the decision support services RFP must address the potential overlap or reuse of these services and facilities within their response.

CORBA Workflow RFP - As appropriate**,** decision support services should leverage the workflow services, which will result from the current RFP.

CORBA Event-Condition-Action Rules Management Facility RFP [15] [21]-

This proposed facility will allow the manipulation of global rule/knowledge bases, whose semantics are consistent with their formulation of event-condition-action rules. It offers a very powerful methodology for deterministically specifying these behaviors. In their facilities, the rule representation will be complementary to, and not an extension of, IDL. IDL interfaces will be provided for the specification, management, and execution of rules. This is effective for certain forms of global knowledge bases. It is unclear at this time how effectively this model would function within complex, problem solving paradigms or modular knowledge, but the ability to embed rule management within modular, higher level problem solving paradigms is desirable for complex healthcare decision support. Until this standard is finalized, its appropriateness as a foundation for interfaces to healthcare rules engines is unclear.

CORBA Data Interchange Facility and Mobile Agent Facility RFP [18]-

Within the development of these RFPs, separate responses were requested for the data interchange facility and the mobile agent facility.

Some of the relevant capabilities for the Data Interchange Facility to the healthcare decision support RFP will enable

- data conversion between multiple representations and
- extensibility to define, discover, and convert new representations.

The responses for the healthcare decision support RFP must address how this applies to the manipulation of decision support content when content issues are addressed.

This document’s definition of engine is very similar to the OMG definition of engines for agents. Given these similarities, proposals may consider the reuse of this agent engine as a base class for decision support engines. However, until the agent facility is finalized, it is difficult to judge its appropriateness for decision support.

CORBA Component Model RFP (ORBOS RFP8) [24] - This is the proposed CORBA component model, inspired by the Java component model (Java Beans). If the system is component-based, this must be addressed.

ORB and Object Services RFI 2 (Real-time) [23] - This attempts to improve the appropriateness of CORBA-based solutions for real-time processes.

IX.B. Vertical Healthcare Services

Clinical Lexicon Query Services RFP [26]

The decision support standards should complement the clinical lexicon/ vocabulary standards. This requires the ability to leverage existing lexicon services when appropriate, as well as function without vocabulary services.

Patient Identification Services RFP [29] [30]

The system used for identifying a patient is called a Master Patient Index (MPI). The decision support services must support functioning with and without an MPI.

Other areas of potential interaction include **Electronic Commerce, Clinical Data Working Group, and Pharmacy.**

X. Presentation Logistics

Neither CogniTech nor Dr. Soller's university group can afford the travel costs to send a representative to a European OMG meeting. The current cost structure of the OMG membership and standards processes have made it financially prohibitive for some emerging small businesses to play an active role. However, some of the most innovative work is done by emerging small businesses. If the RFI submissions are discussed in a U.S. OMG meeting, Dr. Soller will present this submission.

References

1. AHCPR. AHCPR, AAHP, and AMA to Develop National Clinical Guideline Clearinghouse. May, 1997.

2. American Society for Testing and Materials (ASTM) E31.15. Defining and Sharing Modular Health Knowledge Bases (Arden Syntax for Medical Logic Systems). ASTM 1460-92. 1992.

3. American Society for Testing and Materials (ASTM) E31.15. Defining and Sharing Modular Health Knowledge Bases (Arden Syntax for Medical Logic Systems). ASTM 1460-98, draft. May, 1997.

4. Broverman C. Health Level Seven (HL7) and Clinical Decision Support. Presentation Slide Copies. Network-based Real-time Clinical Decision Support (NRCDSS) Workshop. June, 1997.

5. Clancey, W. J. Notes on "Heuristic classification". Artificial intelligence, 1993: 59 (1-2); 191-196.

6. Deibel SRA. Introduction to the InterMed Common Guideline Model and Guideline Interchange Format (GLIF). <http://dsg.harvard.edu/public/intermed/glif_overview.html>.

7. Duda RO and Hart PE. Pattern classification and scene analysis. Menlo Park, Calif. :Stanford Research Institute, N.Y: John Wiley, 1973 (new 1997 edition expected shortly from John Wiley).

8. Farquhar A, Fikes R, & Rice, J. The Ontolingua Server: A Tool for Collaborative Ontology Construction. Knowledge Systems Laboratory, KSL-96-26, September 1996.

9. Finin, Timothy. UMBC Agent Web. http://www.cs.umbc.edu/kse/ .

10. Finin, Timothy, UMBC KQML Web. http://www.cs.umbc.edu/kqml/ .

11. Genesereth, Michael R. Knowledge Interchange Format (KIF). http://logic.stanford.edu/kif/kif.html

12. Greenes RA, and Deibel SRA. Constructing Workstation Applications: Component Integration Strategies for a Changing Healthcare System. Yearbook on Medical Informatics 1996; 76-86.

13. Haykin, Simon S., Neural networks : a comprehensive foundation. NY : Macmillan College Publishing Company, 1994.

14. Hripcsak G, Clayton PD, Pryor TA, Haug P, Wigertz OB, Van der lei J. The Arden Syntax for Medical Logic Modules. In: SCAMC 14. IEEE. 1996; 200-204.

15. Lam, H and Su SYW. ECAA Rules and Rule Services in CORBA. http://www.omg.org/docs/cf/97-01-10.doc .

16. McDonald CJ, Overhage JM. Guidelines You Can Follow and Can Trust: An Ideal and an Example. J Am Med Assoc, 1994; 271 (11): 872-873.

17. Object Management Group (OMG). Object Management Group. <http://www.omg.org/>.

18. OMG. Data Interchange Facility and Mobile Agent Facility RFP, http://www.omg.org/library/schedule/CF_RFP3.htm

19. OMG. Object Management Architecture Guide (OMA Guide), September, 1995.

20. OMG. Common Object Request Broker Architecture and Specification (CORBA), Revision 2, August 1995.

21. OMG. Event-Condition-Action Rules Management Facility RFP, http://www.omg.org/docs/cf/97-01-10.doc

22. OMG. Information Management Common Facilities. ftp://www.omg.org/pub/docs/formal/97-06-07.pdf

23. OMG. ORB and Object Services RFI 2 (Realtime) http://www.omg.org/library/schedule/Realtime_RFI.htm

24. CORBA Component Model RFP (ORBOS RFP8) http://www.omg.org/library/schedule/ORBOS_RFP8.htm

25. OMG CORBAmed DTF. Lexicon Query Services OMG Document: corbmed/96-11-03. http://www.omg.org/docs/corbamed/96-11-03.html,

26. OMG CORBAmed DTF. Lexicon Query Services RFP. http://www.omg.org/library/schedule/CORBAmed_RFP2.htm

27. OMG. CORBAmed Taskforce. http://www.omg.org/corbamed/.

28. OMG CORBAmed DTF. CORBAmed RFI3: Clinical Decision Support. http://www.omg.org/docs/corbamed/97-06-05.pdf

29. OMG CORBAmed DTF. CORBAmed RFP1 Patient Identification Services Request For Proposal. http://www.omg.org/docs/corbamed/96-11-08.txt

30. OMG CORBAmed DTF, Patient Identification Services, Revision 3 http://www.omg.org/docs/corbamed/97-07-04.pdf

31. Pattison-Gordon E, Cimino JJ, Hripsak G, Tu SW, Gennari JH, Jain NL, & Greenes RA. Requirements of a Sharable Guideline Representation for Computer Applications. SMI Report Number: SMI-96-0628 1996.

32. Shortliffe EH, Barnett GO, Cimino JJ, Greenes RA, & Patel VL. InterMed: An Internet-Based Medical Collaboratory. Montreal, Canada,InterMed: An Internet-Based Medical Collaboratory In: Proceedings of INET '96, the annual meeting of the Internet Society of Montreal, June 1996.

33. Sujansky W. Meeting Notes of “Curly Braces Working Group” Meeting. E-mail Correspondence to ASTME31.15 E-mail list. February, 1997.

34. Weiss, Sholom M. Computer systems that learn : classification and prediction methods from statistics, neural nets, machine learning, and expert systems. San Mateo, Calif. M. Kaufmann Publishers, 1991.