A Framework and Tools for Authoring, Editing, Documenting, Sharing, Searching, Navigating, and Executing Computer-Based Clinical Guidelines

Robert A. Greenes, MD, PhD; Aziz Boxwala, MBBS, PhD; William N. Sloan, PhD;
Lucila Ohno-Machado, MD, PhD; Stephan R.A. Deibel, BA

Decision Systems Group, Brigham and Women's Hospital, Harvard Medical School, Boston, MA (greenes@harvard.edu)

ABSTRACT

With the spread of managed care and integrated delivery networks, an increased emphasis has been placed on the cost-effectiveness of clinical practices. The need has been recognized to use guidelines to support education, and to integrate them into clinical practice. A specification for guideline representation that would facilitate computer-based clinical guideline sharing has been developed by the InterMed Collaboratory. Called GLIF (GuideLine Interchange Format), this specification and its proposed extensions have been the basis for our implementation of a framework and suite of integrated software tools for guideline authoring and editing, packaging in XML, Internet distribution, navigation, eligibility determination, and automatic execution

INTRODUCTION

A number of objectives have been cited for computer-based implementation of clinical guidelines [1-6]. These include the ability to structure them more precisely [3, 6]; to share them among developers as well as users [6]; to verify them [7]; to document them and link them to appropriate literature or databases [8]; to search for them based on various criteria; to navigate or browse them [9]; and to integrate them into applications. Guidelines can be applied in educational settings as drivers for clinical simulations or patient management problems. In clinical settings, guidelines can be used to support workflow management by organizing the presentation of forms for data collection and action selection [10]. They can be used to aid in decision-making [4, 7], to facilitate management of clinical trials [11], to determine appropriateness of referrals, and to implement critical paths.

The prospect of such guideline applications has increasing attractiveness, as managed care and integrated delivery networks have grown and placed emphasis on cost-effectiveness of clinical practices. A major goal is now to reduce practice variation through improved access to information about best practices, including guidelines and the evidence to support them. A means has been needed for integrating these guidelines into the framework of the information environment. The World Wide Web has provided a venue for improved access to guidelines, as advocated by the government (e.g., the Agency for Health Care Policy and Research), by professional societies (e.g., the American Medical Association), and by health care organizations.

Re-engineering of the processes of designing and conducting clinical trials, to make the processes more distributed, has in addition led to refinements of the representation of clinical guidelines in protocols. It has also refined the ways in which they should drive data collection and data validation processes.

A specification for guideline representation to facilitate clinical guideline sharing was developed by the InterMed Collaboratory [12], a joint project of medical informatics laboratories at Harvard (the Decision Systems Group or DSG at Brigham and Women's Hospital and Laboratory of Computer Science at Massachusetts General Hospital), Columbia, and Stanford. Developed with National Library of Medicine support, this specification is called GLIF (GuideLine Interchange Format) [6]. GLIF has been the basis for several implementations of guideline-based applications, including one in the Brigham's BICS [10] information system, Web-based applications for driving clinical consultations [13], and applications that search for eligible clinical protocols [13].

The DSG has developed a set of tools built around the GLIF representation and access to shared guideline repositories. The tools provide guideline authoring and editing, GLIF import/export, packaging in XML (eXtensible Markup Language), server access, viewing/navigation, and automatic execution. This paper reports on that work, and describes the tools that have been implemented and their architecture. We then discuss directions for future work.

INTEGRATED TOOL SUITE

The Guideline Repository. A repository of shared clinical guidelines is central to the DSG's tool suite. It stores clinical guidelines, classifies and indexes them, and maintains and controls their access over a

network. The repository's network interface was designed to be compatible with GLIF.

A guideline repository is analogous to the Web servers accessed over the Internet. Both seek to provide shared, reliable and potentially secure access to information in a standard format. Instead of providing read-only access to pages of text, graphics and sound, in our design a guideline server optimizes the two-way network transfer of clinical guideline information. The guideline server maintains a collection of customizable guideline category trees, for example, with which authors may classify their guidelines as they share them. When a new guideline is added to a repository, the server tool generates and adds a time-stamped ID unique across the network. With this ID it recognizes a guideline even if its name or classification has changed and tracks its versions over time.

The guideline server implements client authentication to offer different levels of access to custom groups of users with varying needs. In the current default configuration, users gain read-only access to shared guidelines without logging in. However, to store a new or modified guideline, delete a guideline or one of its versions, or modify guideline categories, the user must log in with a unique name and password.

Much like Web servers, multiple guideline servers can operate on the same network. Our suite of tools is configured to access the Decision System Group's guideline server by default. This repository contains six small demonstration guidelines, including cancer protocols and a dermatology treatment guideline. The guideline tools can alternatively be directed to a guideline server running on the user's local network or desktop. Each guideline server provides a shared, dynamic set of data that can be used in different ways by multiple types of client software, including our guideline authoring, viewing, and execution tools.

Guideline Authoring. The guideline authoring tool allows users to create new clinical guidelines, store them in one or more shared repositories, and edit and maintain them over the network. A GLIF-encoded guideline can be opened from a server on the network or read from a local text file. The authoring tool can also import files in the Protégé [14] data format.

GLIF specifies a common set of building blocks for describing a clinical guideline. Examples are an *action step* that collects *patient data* or a *condition step* that represents the *criteria* for making a decision. GLIF is recursive, in that a sub-guideline may be embedded in any action step. The authoring tool provides a graphical interface with which each type of step can be created, ordered in a flowchart display, and then fully specified using forms.

GLIF defines guideline building blocks at a very detailed level. Subject matter experts may wish to work at a higher level of conceptualization. To facilitate this, more abstract templates can be used to represent typical constellations of clinical actions and decisions. The guideline authoring tool currently supports one such standard template, termed a *GEODE* step (for "Guided Entry Of Data Elements") [15]. The GEODE step can be used to describe resources such as patient data, provider actions, and didactic materials appropriate to an encounter. As such, a GEODE step can be considered as a means to anticipate information needs and facilitate workflow. This building block consists of a set of concurrent actions followed by a decision step, used for describing data elements, decisions, and actions expected for classes of clinical encounters. It may optionally include eligibility criteria. Internally, the GEODE step is composed of several linked GLIF steps. It is identified in a GLIF file by an optional GEODE tag that standard GLIF readers can ignore, but that DSG tools use to construct a GEODE step automatically,

The authoring tool may be switched from a GLIF-based to a GEODE-based editing mode, depending on the needs. No matter how edited, guidelines are always stored in the underlying GLIF format. As a user works with the authoring tool, an emerging graph of guideline GLIF or GEODE steps is rendered as a flowchart on the screen. Clicking on any step displays a dialog with which to edit its attributes.

Specialized form dialogs are used to enter such guideline components as patient data or decision criteria. The dialog for specifying patient data allows definition of enumerated and compound data types and such meta-data attributes as measurement scale, logical and time constraints, and expiration period.

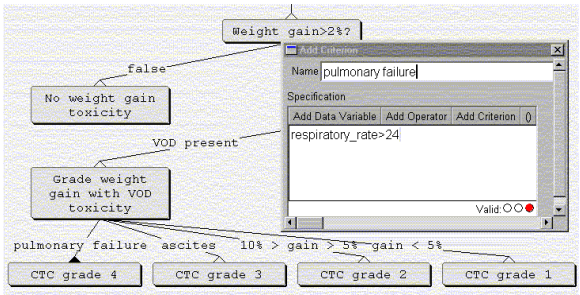

Figure 1. A decision criterion in a conditional step is entered in a dialog. The expression is validated as it is typed. The indicators at the bottom of the dialog signal a valid or invalid expression.

Decision criteria can be entered in a variety of formats (Fig. 1). An expression such as "history of heart disease," for example, may initially be specified in natural-language text. A red icon is displayed in the

dialog to indicate that such an expression cannot be parsed. If the expression is refined so that it can be parsed successfully, but one or more terms remains unspecified, a yellow icon is displayed and all unknown terms are highlighted. Once the expression has been entered in a complete and valid form, a green icon indicates it is machine-executable.

Guideline Viewing. The guideline viewing tool permits the visual navigation of a guideline using the same flowchart representation as the authoring tool. The viewing tool is available in the form of a Java applet and can be incorporated into Web pages on the Internet. It can thus provide convenient, point-of-need access to GLIF-encoded clinical guidelines.

Once inserted into a page on a Web server using the HTML APPLET tag, the viewing tool is downloaded when a user visits that page. The viewing tool then contacts a specific guideline server over the Internet to obtain an encoded clinical guideline to display.

Parameters in the APPLET tag specify which server and which guideline to access. Optional parameters select one step in the guideline and specify how much of the flowchart to display. This makes it easy to augment any Web page with an interactively navigated clinical guideline. The guideline viewing tool can also be controlled dynamically by other applets or by JavaScript programs.

The viewing tool allows a guideline to be traversed according to the user's needs at a customizable level of detail. Both GLIF-level and GEODE-level steps are supported, and each step type is given a distinct visual representation. The user may display either summary or detailed information for any step. The traversal path to the current step may be collapsed into a scrollable summary area and embedded subguidelines may be displayed or collapsed.

Clinical guidelines may contain cycles; the guideline viewer unrolls these. When there is more than one way to arrive at a step while traversing a guideline, the user can click on that step and select from a popup menu the name of any one of its immediate antecedents. The specified antecedent will then become the new root node of the flowchart display.

The guideline viewing tool can guide access from any step to associated didactic material (e.g., evidence, explanation, decision-support or patient-oriented information), either in the form of stored text or as a URL. A step that has associated material is marked in the flowchart by a special icon (Fig. 2). When the user clicks on such a step and selects an item from a popup menu, depending on the type of material either the stored text is displayed in a dialog box or a new browser window is opened with the Web page identified by the URL.

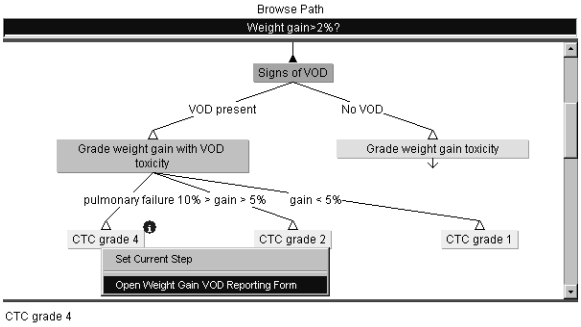

Figure 2. Supplemental didactic material can be associated with any guideline step. In the guideline browser, this material can be accessed by clicking on the step. If the supplemental material is a URL, a new browser window will be opened with the referenced document.

Guideline Execution. We have extended GLIF to enable execution of guidelines in computer programs. The extensions to GLIF include a specification for a structured grammar (derived from the logic constructs in Arden Syntax) for decision logic and eligibility criteria [16], and an object-oriented representation for describing clinical actions and their parameters.

With the development of a machine-executable syntax powerful enough to handle clinical guideline constructs, we have added a prototype guideline execution engine to our tool suite. This engine imports an XML file containing a GLIF-encoded guideline, which it traverses by evaluating logic conditions specified in the guideline in some specific context of patient-data values. The results of this evaluation are used to generate patient-specific recommendations from the guideline. This tool is being used as the basis for a variety of applications. The execution engine, clinical applications under development, and an educational application are described in detail elsewhere [13].

Eligibility Determination. A related set of tools has been developed at the DSG to support encoding of eligibility criteria for clinical guidelines or protocols [17]. These use the same extensions to GLIF as the guideline execution engine. The goal of these tools is to be able to select particular guidelines or protocols based on data pertaining to a specified patient.

ARCHITECTURE AND IMPLEMENTATION

The DSG's integrated suite of guideline tools is based on a component-based client/server architecture (Fig. 3) which is currently implemented using CORBA (Common Object Request Broker Architecture). This component-based architecture allows us the flexibility of developing applications rapidly by "mixing-and-

matching" components. Using CORBA, components in an application may be implemented in different programming languages and be distributed over a network. Thus, the Java-based guideline browser on a remote host can communicate with a guideline server running at a central location.

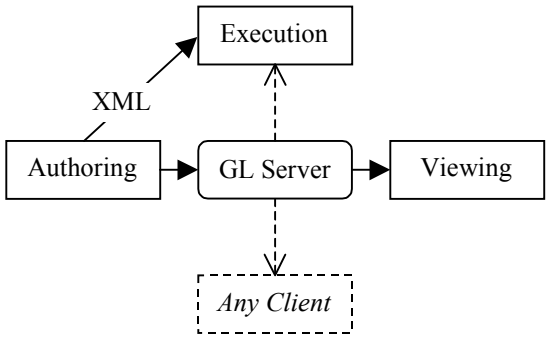

Figure 3. The guideline tool suite is based on a client/server architecture. Client applications, such as the authoring tool, browser, and execution program access and store guidelines to the server. The execution program currently does not access guidelines in the server. The execution program instead uses guidelines in the XML format exported from the authoring tool. Future work appears with dashed lines.

An ORB (Object Request Broker) is the CORBA middleware that establishes the client-server relationships between applications over the network. Currently the guideline server relies on a custom ORB developed as part of the Arachne project [18] at the DSG. This ORB is not yet fully compatible with the CORBA 2.0 standard, though it can interoperate with other ORBs via IIOP. The Java guideline viewer client uses Netscape's built-in ORB or, in the case of Internet Explorer, ORBacus from Object Oriented Concepts, Inc.

The guideline repository and authoring tools are written in C++. They were developed using the GNU C++ compiler, Microsoft Visual C++, and Metrowerks CodeWarrior. The authoring tool makes use of a layered toolkit for portable GUI software, developed at the DSG, that supports execution on multiple platforms. The latest versions of these tools run on Windows95/NT 4.0 and Linux 2.x.

The guideline viewing tool is written in Java using Sun's Java™ Development Kit 1.1. The current version of the viewer applet runs on any web browser supporting JDK 1.1, including Netscape 4.0.6 or later and Internet Explorer 4.0 or later (using the latest Microsoft Virtual Machine for Java). The guideline execution engine is currently a prototype implemented with ActiveX components and JavaScript.

The guideline server and authoring tools read and write GLIF stored in disk files using ODIF (Object Data Interchange Format). The latter is currently being replaced as our default format by XML. The authoring tool also writes guidelines to disk in XML 1.0 format, which is read as input by the guideline execution engine. The XML format for GLIF is a simple transformation from ODIF that replaces ODIF packaging with XML tags.

Our tools may be freely distributed; in many cases the source code is available for download. The DTD (Document Type Definition) we use for XML packaging of GLIF is also available on our web site (see http://dsg.harvard.edu/public/software/guideline).

FUTURE WORK

Before guidelines can play a more central role in clinical practice, they must not only be widely accessible but also be simple to author and yet provide sufficient expression required for clinical decision support. With respect to the former, we have recently begun a redesign of our guideline authoring tools with the goal of improving ease of guideline development and refinement directly by subject experts. We are presently conducting ethnographic studies with potential guideline designers to better define the user interface requirements of these tools.

At the same time, we are redesigning our tools with a simpler and more extensible software architecture to facilitate their use in a wide variety of application-specific software. For example, we expect to build a custom protocol-authoring tool by extending the next version of the GLIF authoring tool. We envision coupling our guideline execution engine and viewer to provide a graphical navigation context during form-based data acquisition.

To facilitate ease of use in custom guideline applications, we are creating representation macros using GLIF. These macros will represent commonly occurring clinical tasks too complex to describe with individual GLIF steps. The GEODE step described earlier would be represented in this macro form. Such a macro might also represent a visit or encounter in a clinical-trial protocol. The visit macro would be defined internally as a sequence of GLIF steps containing information on data to be collected, decisions to be made, and treatment to be delivered.

To permit the custom integration of our guideline tools at the enterprise level, we will create facilities for utilizing institutional data dictionaries and provide hooks for integration with institutional knowledge bases and workflow management systems. To meet the requirement for platform portability and network

accessibility, we plan to base future implementations on Java.

CONCLUSIONS

We have developed an integrated suite of tools for authoring and browsing tools and eligibility determination, and a prototype execution engine for clinical guidelines, based on proposed extensions to the GLIF and Arden standards. In the process, we have begun to explore the integration of guidelines into the World Wide Web and the automated execution of guidelines expressing complex clinical concepts in specialized terminology. We are beginning to use what we have learned in this initial effort to design a new suite of tools. We hope to make guidelines easier to use and to apply reliably in a diverse set of clinical and educational applications.

ACKNOWLEDGMENTS

This research was funded by contract MDA972-94-3-0047 from DARPA. Some eligibility determination tools were funded by contract 34078PP1024 from the Massachusetts Department of Public Health and grant DAMD17-98-1-8093 from the Department of the Army. The contents of this article do not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred.

REFERENCES

1. Advani A, Lo K, Shahar Y. Intention-based critiquing of guideline-oriented medical care. *Proc Amia Symp*. 1998:483-7.
2. Dazzi L, Fassino C, Saracco R, Quaglini S, Stefanelli M. A patient workflow management system built on guidelines. *Proc AMIA Annu Fall Symp*. 1997:146-50.
3. Fox J, Johns N, Rahmanzadeh A. Disseminating medical knowledge: the PROforma approach. *Artif Intell Med*. 1998;14:157-81.
4. Lobach DF, Hammond WE. Computerized decision support based on a clinical practice guideline improves compliance with care standards. *Am J Med*. 1997;102:89-98.
5. Nilasena DS, Lincoln MJ. A computer-generated reminder system improves physician compliance with diabetes preventive care guidelines. *Proc Annu Symp Comput Appl Med Care*. 1995:640-5.
6. Ohno-Machado L, Gennari JH, Murphy SN, et al. The guideline interchange format: a model for representing guidelines. *J Am Med Inform Assoc*. 1998;5:357-72.
7. Shiffman RN. Representation of clinical practice guidelines in conventional and augmented decision tables. *J Am Med Inform Assoc*. 1997;4:382-93.
8. Agency for Health Care Policy and Research. National Guideline Clearing House. http://www.guideline.gov.
9. Abendroth TW, Greenes RA. Computer presentation of clinical algorithms. *MD Comput*. 1989;6:295-9.
10. Zielstorff RD, Teich JM, Paterno MD, et al. P-CAPE: a high-level tool for entering and processing clinical practice guidelines. Partners Computerized Algorithm and Editor. *Proc Amia Symp*. 1998:478-82.
11. Fox J, Johns N, Lyons C, Rahmanzadeh A, Thomson R, Wilson P. PROforma: a general technology for clinical decision support systems. *Comput Methods Programs Biomed*. 1997;54:59-67.
12. Shortliffe EH, Barnett GO, Cimino JJ, Greenes RA, Huff SM, Patel VL. Collaborative medical informatics research using the Internet and the World Wide Web. *Proc AMIA Annu Fall Symp*. 1996:125-9.
13. Boxwala AA, Greenes RA, Deibel SRA. Architecture for a Multipurpose Guideline Execution Engine. Technical Report, TR-1999-03. Decision Systems Group, Boston, MA; 1999. http://dsg.harvard.edu/tr/.
14. Musen MA. Dimensions of knowledge sharing and reuse. *Comput Biomed Res*. 1992;25:435-67.
15. Stoufflet P, Deibel SRA, Lee D, Traum JH, Greenes RA. Integrating guidelines and information support into clinical workflow in a distributed environment: Data structures and knowledge represetation for GEODE-CM. . *Spring Congress.* Kansas City, MO.: AMIA; 1996:109. .
16. Wang SJ, Ohno-Machado L, Boxwala A, Mar P. Representing Criteria in Guidelines and Clinical Trial Protocols: Common Needs and Solutions. *Proc Amia Symp*. 1999.
17. Ohno-Machado L, Wang SJ, Mar P, Boxwala AA. Decision support for clinical trial eligibility determination in breast cancer. *Proc Amia Symp*. 1999.
18. Greenes RA, Deibel SR. Collaborative health care information system development through sharable infrastructure, services, and paradigms. *Medinfo*. 1995;8:190-4.