Cross-Tool Communication: From Protocol Authoring to Eligibility Determination

John H. Gennari,[1] David Sklar,[2] and John Silva[3]

[1]Information & Computer Science Dept, University of California, Irvine CA
[2]iKnowMed, Inc., Berkeley CA
[3]National Cancer Institute, Bethesda MD

Abstract

To be effective, informatics tools for clinical trial protocols must inter-operate and share knowledge. We demonstrate a simple XML-based communication of eligibility criteria information between two independently-developed informatics tools. Using a shared DTD model of criteria, an authoring tool (developed within the Protégé environment) can send a list of eligibility criteria to a commercial system for automatic eligibility determination (the "iKnowChart" system by iKnowMed). The criteria model, developed as a Protégé ontology, includes both the terminology and the logic needed to compute eligibility for a given patient. As a demonstration of cross-tool communication, we have encoded criteria from an active clinical trial protocol (E1199), and shown how use of the authoring tool can effectively update the eligibility knowledge and the behavior of the commercial iKnowChart system. As part of the cross-tool knowledge sharing, we use Common Data Elements, an oncology terminology developed by the National Cancer Institute.

Background and Rationale

Over the last 10 years, there has been a sharp increase in the number of open clinical trial protocols. The process of authoring and carrying out a single trial is a staggeringly expensive and time-consuming process: It can take 4-6 years of work to transform a research idea into publishable results of a multi-site clinical trial. Unfortunately, most of this work is still carried out via paper-based systems. From an informatics perspective, this situation provides great opportunity: A set of decision-support tools that improve the flow and speed of information transfer could have a significant and positive effect on protocol-based research.

Although there has been informatics work in this area[1-3], there are few successfully deployed tools for protocol management or protocol-based research. One persistent challenge is that these applications must be integrated with other tools and medical information systems that operate at the point of care. Modern physicians already confront information overload, and adding yet another stand-alone tool will simply further complicate their workflow. Unfortunately, it is difficult to build an integrated set of tools that support a complex process like protocol-based research.

To build interoperating tools, we can apply research results in software engineering and knowledge-base systems. One such finding is that it is important to have systems agree on the semantics of the shared information. That is, even if a pair of tools agree on the syntax of the communication (e.g., HL7, XML, or CORBA), for them to successfully interact, they must also share a semantic model that describes the objects and the information in the domain of discourse.[4] Thus, we advocate the development and publication of formal models or ontologies for clinical trial protocols.

However, formal models must be grounded in real-world tasks—otherwise the formalism may be unwieldy or inappropriate for actually use. Our methodology for building inter-operating tools is to work from a specific scenario that we know has real-world relevance. In particular, we have selected the process of clinical trial recruitment, and the evaluation of eligibility criteria as the targets for our decision-support tools.

The low rate of patient accrual is becoming a more acute problem as the number of protocols and experimental agents increases. Unwieldy eligibility determination is often cited as a key reason for low accrual rates. This determination can be difficult either because the criteria are unnecessarily complex, or because it may be difficult for staff at the point of care to collect and access both the protocol eligibility requirements and the patient data needed to assess eligibility.

There are at least two distinct decision-support tools to help ameliorate this situation. First, a protocol-authoring tool could help clinical investigators write protocols that are more clear and that contain appropriate eligibility criteria. Second, at the point of care, a screening tool could match patients to protocols, by evaluating those criteria against known patient data. In the next two sections, we briefly describe two applications that address these two distinct tasks. Although these tools were developed independently, we were able to build a common semantic model for eligibility criteria, and then use this model as a basis for communication between the two tools.

A crucial part of the common semantic model is a shared medical terminology. As a basis for communication, both tools incorporate the Common Data Elements (CDEs), developed by the National Cancer Institute. CDEs allow our tools to communicate in a consistent way about the patient or medical data needed to evaluate protocol eligibility criteria.

1067—5027/01/$5.00

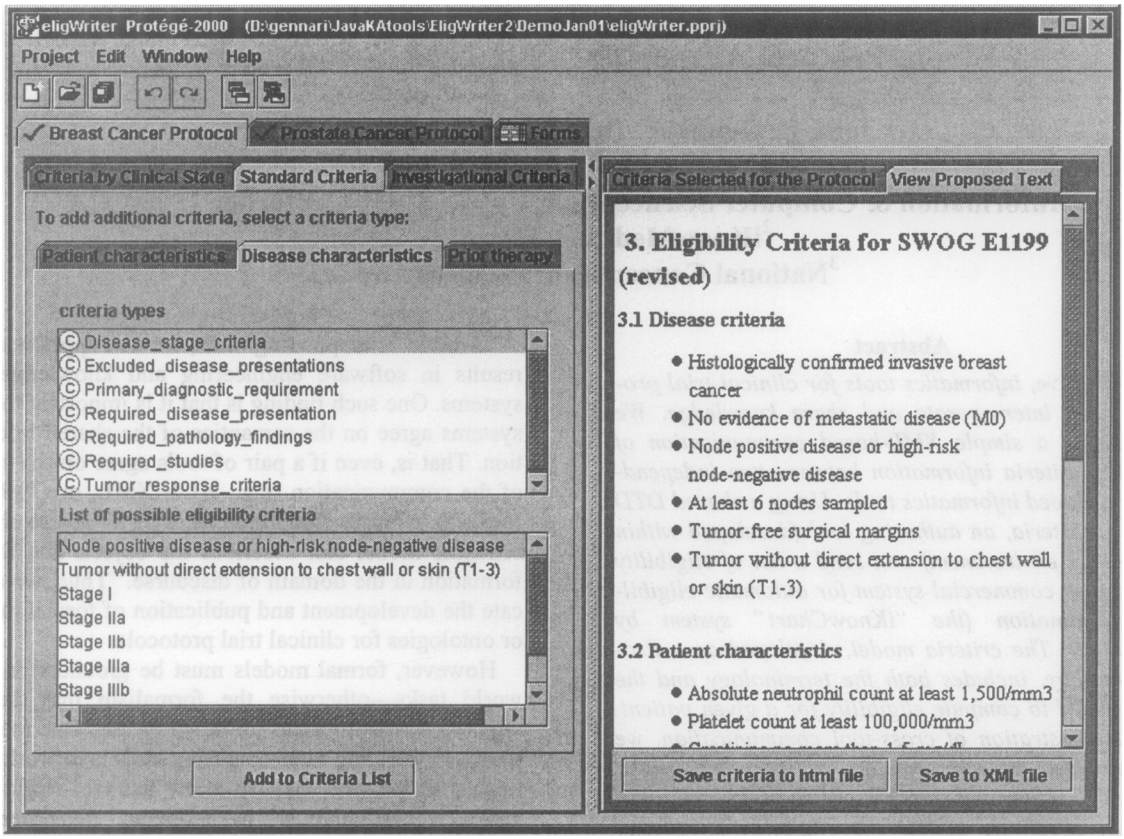

Figure 1. The Protégé EligWriter authoring tool. An investigator selects from standard criteria on the left, building or editing the list of eligibility criteria for the clinical trial protocol on the right.

Authoring Eligibility Criteria

In any complex design process, it is cost-effective to correct problems early rather than late. If we can detect and fix problems when the protocol is created, the downstream benefits can be significant for the overall protocol-based research process. Therefore, we propose that clinical investigators use a decision-support tool when creating new protocols. Such a tool could (1) improve the consistency across protocols, (2) increase the ease and speed with which protocols are approved, and (3) detect and possibly correct errors.

Figure 1 shows a protocol-authoring tool, EligWriter, that focuses on eligibility criteria.[5] We developed this tool within the Protégé knowledge-based system environment.[6] (Protégé was also used for development of the shared criteria model.) The figure shows the criteria for a current breast cancer protocol (SWOG #E1199) in the right hand panel. When an investigator creates a new clinical trial protocol, the panel on the right is initially empty, and the left side of the tool provides different ways to add eligibility criteria to the new protocol. In Figure 1, the left pane shows a some of the criteria associated with "disease characteristics", and the user has selected one particular criteria. The tool also includes a graphical interface that helps users select a clinical state, and thereby, an initial set of eligibility criteria.[5] When the user adds criteria to the protocol on the right, not only is the standard text included, but the system associates this text both with a standard medical terminology (CDEs) and with the computational logic needed to interpret and evaluate the truth of that criteria. As we describe next, this logic is essential for interpreting and evaluating criteria in order to screen patients into protocols.

Screening for Patient Eligibility

In an independent effort, the Southwest Oncology Group (SWOG) has collaborated with iKnowMed to develop tools aimed at increasing patient accrual into SWOG protocols. One of iKnowMed's products is a patient charting tool (iKnowChart), that includes the capability to compare patient information against protocol eligibility criteria. Collaboratively with SWOG, iKnowMed has entered information from almost 40 current SWOG clinical trial protocols into the iKnowChart system. The resulting protocol knowledge base and screening tool are currently being deployed at a national cancer center.

iKnowChart compares patient information against

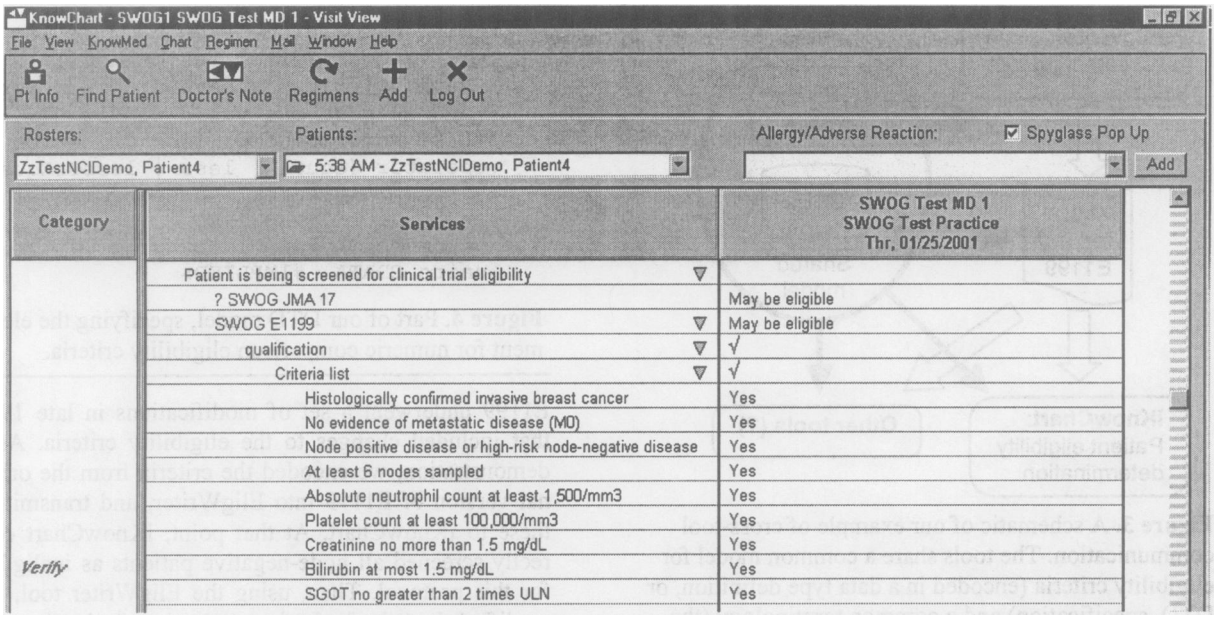

Figure 2. A partial view of the iKnowChart patient screening tool. The criteria listed for E1199 were created by reading and parsing the XML specification received from the Protégé EligWriter tool.

eligibility criteria in a two stage process. First, iKnowChart uses a set of rules derived from a pre-selected subset of the eligibility criteria to perform partial patient screening. These rules use key diagnostic and disease status data to automatically determine which clinical trial protocols a patient might qualify for. The advantage of partial screening is that as clinicians enter patient data as part of their normal charting process, the system can automatically provide the clinician with a list of trials that the patient might be eligible for, and eliminate from consideration those for which the patient is clearly ineligible. In a second stage, the clinician can choose to look at the full list of eligibility criteria for any of these protocols.

Figure 2 is a partial screenshot of the iKnowChart tool, showing the full listing of the eligibility criteria for protocol E1199. In many cases, the system can automatically evaluate these criteria by comparing the terms and the logic of the criteria against known patient information. For example, iKnowChart may know that a given patient has had 9 lymph nodes biopsied for disease. In this case, the tool can mark the "at least 6 nodes sampled" criteria with a "Yes", indicating that this patient has fulfilled this criteria.

The screening and eligibility determination capabilities of iKnowChart are provided as part of a package of decision support for protocol-based care. For example, if a patient does qualify for a particular protocol, the system can present the clinician with scheduling information for the various arms of the trial. This could include schedules for drugs with specific cycles and dates, as well as schedules for lab tests and other required procedures. As an additional capability, we are building electronic case report forms for the data collected during the course of a clinical trial, and an interface to electronically transfer that data to the SWOG statistical center in a secure fashion. Thus, the broad purpose of the SWOG project with iKnowMed is to test, in a real setting, the ability of informatics tools to improve the protocol-based research process.

Communicating via XML

For two systems to communicate, they must share some common semantics. In this case, we were able to build a common model because both systems work with the same real-world objects: eligibility criteria for oncology clinical trial protocols. Of course, communicating systems also need to share a common syntax. However, as long as the syntax is reasonably expressive, we view this choice as less important—if the semantics are shared, then it is easy for systems to convert from one syntax to another.[4]

For our task, the shared model includes two components: a model of the eligibility criteria, and a shared medical terminology. We use XML documents as the syntactic medium for the shared model of eligibility criteria. For terminology, we use the NCI's Common Data Elements (CDEs) that provide a medical terminology for oncology clinical trial protocols. Based on these models, the Protégé EligWriter system can publish any particular set of eligibility criteria as an XML document. This document is read by iKnowChart, and then incorporated into its knowledge base so that the criteria can be applied to individual patients.

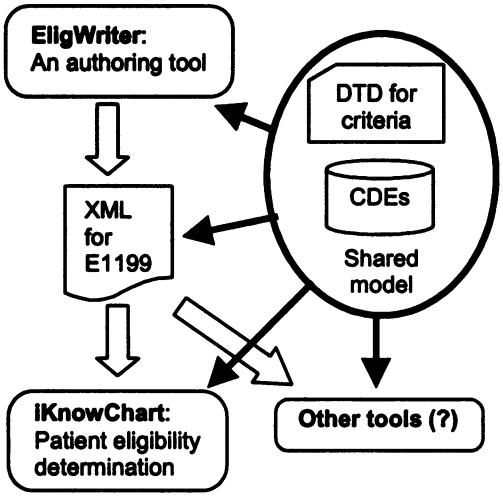

Figure 3. A schematic of our example of cross-tool communication. The tools share a common model for eligibility criteria (encoded in a data type definition, or DTD, specification) and a common terminology (the Common Data Elements, or CDEs).

Figure 3 shows our approach to cross-tool communication. Our choice of XML has the advantage of being a well-established and simple standard for capturing document semantics. (For example, Microsoft's Internet Explorer can render any XML / DTD pair of files into an easily viewable document.) Due to XML's very accessible syntax, any application that subscribes to our shared model could use the information published in XML by the EligWriter tool.

Our choice of XML implies that we describe our shared model of criteria as a document type definition (DTD) specification. In keeping with our grounded methodology, we kept our model of these criteria fairly simple: there are four classes of criteria, they can be combined in arbitrary Boolean sentences, they contain simple numeric comparisons, and they include reference to a common medical terminology. Figure 4 shows a portion of our model: the DTD specification for a simple numeric comparison criteria.† This element would be used to express, for example, the criterion specified as "at least 6 nodes sampled".

The formalism of our model is hidden from end-users. Users of the authoring tool simply push the "Save to XML" button, and EligWriter produces the XML file that uses our shared DTD model. Likewise, users of the iKnowChart system see only the resulting capability of patient screening, without needing to know anything about the shared model or how the eligibility criteria are stored in the local knowledge base.

Cross-tool communication is especially important for information maintenance. For example, the protocol

† Shared models should be public: Our complete specifications can be found at www.ics.uci.edu/~gennari/criteria/

```
<!ELEMENT Numeric_Comparison_Criterion
        (DomainTerm, (ValidWindow)?)>
<!-- Numeric criteria have a domain term
     and at most one valid window -->
<!ATTLIST Numeric_Comparison_Criterion
    operator (equal | less | lessOrEqual |
              greater | greaterOrEqual )
                #REQUIRED
    value CDATA #REQUIRED
    units CDATA  #IMPLIED>
```

Figure 4. Part of our DTD model, specifying the element for numeric comparison eligibility criteria.

E1199 underwent a set of modifications in late 1999 that included changes to the eligibility criteria. As a demonstration, we encoded the criteria from the original version of E1199 into EligWriter, and transmitted these to iKnowChart. At that point, iKnowChart correctly screened all node-negative patients as ineligible for this protocol. Then, using the EligWriter tool, we modified the set of criteria per the actual changes that occurred in 1999. One of these changes was to make eligible those patients who were node-negative but "high-risk". When we transmitted these updated criteria to iKnowChart, the screening functionality was modified, and the tool allowed certain node-negative patients to enroll in the revised version of E1199. Thus, importing the new set of criteria from EligWriter automatically modified the behavior of iKnowChart without requiring any re-programming.

On the national scale, we envision this type of cross-tool communication occurring among multiple sites, applications, and vendors. If investigators modify their protocols with an authoring tool such as EligWriter, then their changes can be rapidly disseminated to decision-support tools at all accruing sites.

A requirement for this scenario is that tools share both a model of eligibility criteria and a common medical terminology. In our domain, we choose Common Data Elements (CDEs) as a standard medical terminology. This terminology is designed by the NCI specifically for oncology, and CDEs are required for case report forms from cooperative group protocols. Thus, systems that import our XML document of criteria must understand the CDEs used in those criteria. By "understand", we mean that the system must either be able to use those terms directly, or be able to map those terms to some local terminology. In our example, iKnowChart uses both approaches—the system can use some of the eligibility criteria CDEs as is, whereas it must map other CDEs to pre-existing terms in its knowledge base.

Discussion and Future Directions

The need for electronic communication of medical knowledge across tools has long been recognized—for example, this need is the motivation for the HL7 stan-

dard. However, although HL7 provides a syntactic method for communication, it does not include any richer knowledge semantics. In our approach, we needed to established semantics for eligibility criteria that were common across the two tools. Therefore we built an ontology of criteria classes, terms and relationships within the Protégé environment, before mapping these semantics to XML and DTD. Recent work within the HL7 standard recognizes the importance of semantics. For example, the Reference Information Model (RIM) provides a much richer semantic model for the sorts of data that can be communicated via the HL7 syntax.[7] However, this model is very broad in scope, is aimed primarily at synchronous communication, and does not model objects as specific as eligibility criteria.

As a more specific example, the guideline interchange format (GLIF) was also designed as a semantic model for communication between systems.[8] However, the initial version of this model did not include sufficient specification for objects such as eligibility criteria. Although recent work has addressed this gap by using a superset of Arden syntax,[9] we find this model a bit unwieldy.

In contrast, our approach to modeling is more task-specific and lightweight. We first identified a relevant, real-world task, and then designed a small model to capture exactly the semantics needed for our task (in this case, specifying eligibility criteria). The key to making this approach scalable across many different tasks is that it must be easy for developers to create and maintain the shared model (the ontology and the DTD specifications). Fortunately, modeling tools such as Protégé and DTD editors are readily available.

We would like to increase the scope of our authoring tool to include other aspects of clinical trial protocols. One example would be to include a capability for authoring a protocol's calendar of tasks. Since the iKnowChart tool includes a scheduling capability, then just as with criteria, information from a calendar-authoring tool could be published and transmitted to iKnowChart and used to update its scheduling capability. Of course, we would need to create a new shared model for protocol calendars and tasks. We believe that the Protégé environment makes this easy to do, and that these models can be build up in a modular, task-specific manner.

Our broad hypothesis is that our approach will reduce errors and the amount of work needed for clinical trial research and care. Unfortunately, it is premature to objective measure this benefit: We would need to expand our set of protocols, and measure some approximation of effort (in person-hours, for example) as well as error or problem rates.

However, a more specific hypothesis is that our model of eligibility criteria allows independently-developed applications to communicate. We have validated this hypothesis: we have shown that the use of a standardized vocabulary and a semantic model of eligibility criteria facilitates communication and transfer of data across systems. Our hope is that this demonstration will encourage further development of decision-support tools for protocol-based care, especially authoring tools that produce both standardized text and sharable specifications that can be easily interpreted by commercial point-of-care systems.

Acknowledgements The work we present here is the result of a large team of collaborators. We would like to acknowledge Mark Musen (Stanford University), John Crowley (Southwest Oncology Group), and Jack Hahn, Bob Barcklay, and Rob Calvert (iKnowMed). This work was partially funded by SWOG and the National Cancer Institute.

References

1. Tu SW, Kemper CA, Lane NM, Carlson RW, Musen MA. A methodology for determining patients' eligibility for clinical trials. *Methods of Information in Medicine* 1993; **32**(4): 317–325.

2. Wyatt JC, Altman DG, Heathfield HA, Pantin CF. Development of Design-a-Trial, a knowledge-based critiquing system for authors of clinical trial protocols. *Computer Methods and Programs for Biomedicine* 1994; **43**(3-4):283–291.

3. Sim I, Owens DK, Lavori PW, Rennels GD. Electronic trial banks: a complementary method for reporting randomized trials. *Medical Decision Making* 2000; **20**(4): 440–450.

4. Neches R, Fikes R, Finin T, Gruber T, Patil R, Senatir T, Swartout WR. Enabling technology for knowledge sharing. *AI Magazine*, 1991; **12**(3): 36–56.

5. Rubin DL, Gennari JH, Srinivas S, Yuen A, Kaizer H, Musen MA, Silva J. Tool support for authoring eligibility criteria for cancer trials. *Proc AMIA Symp* 1999: 369–373.

6. Musen M, Fergerson R, Grosso W, Noy N, Crubezy M, Gennari J. Component-based support for building knowledge-acquisition systems. In: *Conference on Intelligent Information Processing (IIP 2000) of the IFIP World Computer Congress (WCC 2000)*; 2000; Beijing, China.

7. Schadow G, Russler DC, Mead CN, McDonald CJ. Integrating Medical Information and Knowledge in the HL7 RIM. *Proc AMIA Symp* 2000; 764–768.

8. Ohno-Machado L, Gennari JH, Murphy S, Jain, NL, Tu SW, Oliver DE, Pattison-Gordon E, Greenes, RA, Shortliffe EH, Barnett GO. The guideline interchange format: A model for representing guidelines. *Journal of the American Medical Informatics Association*, 1998; **5**, 357–372.

9. Peleg M, Boxwala AA, et al. GLIF3: The Evolution of a Guideline Representation Format. *Proc AMIA Symp* 2000: 645–649.