From Narrative Therapeutic Guidelines to Decision Support : Design for Production Efficiency and Safety

Dr Bryn Lewis

Department of Computer Science and Computer Engineering
La Trobe University
Email: lewisba@cs.latrobe.edu.au
http://homepage.cs.latrobe.edu.au/lewisba/frindex.html
phone: +61 3 9479 3681
fax: +61 3 9479 3060

This paper reports on a research project aimed at optimising the print and electronic production process of therapeutic guidelines and making the data elements and underlying logic machine-readable. We first describe a re-engineered guideline production system, SISMO (Single Input Source, Multiple Output), which has been devised and is currently being introduced into a guidelines production process. SISMO allows narrative text to be developed in a traditional manner, however a variety of print and electronic outputs can then be produced through XML and XSLT files. We then describe an Algorithmic Decision Support (ADS) server. The ADS server allows the execution of the reasoning that is implicit within the narrative of the guidelines and is intended to interface either with a user or an external system according to standard communication protocols.

Keywords: Therapeutic guidelines, algorithmic decision support server, decision support – quality, safety, architecture, XML, XSLT

ACM Computing Classification: J.3

1. BACKGROUND

This paper reports on a research project that is a collaboration between Therapeutic Guidelines Limited (TGL), the National Prescribing Service (NPS) and a number of Universities (Lewis, 2001). The project goal is to optimise the production process at TGL and develop the infrastructure necessary to enable guidelines to be used for a range of decision support applications.

The traditional TGL production process focused on producing printed output from narrative text documents. More recently, in response to the increasing use of computers in medical practice, the guidelines have been output in electronic formats, requiring considerable hand-crafting and quality control checking.

TGL is a long established provider of best-practice guidelines concerning the treatment of choice for common conditions (Hemming, 2000). The guidelines are primarily represented as narrative text (Writing Group for Therapeutic Guidelines: Antibiotic, 2000) although some information is presented as tables and some can be represented algorithmically (Fox and Das, 2000; Wang, Peleg *et al*, 2002).

Manuscript received: 19 September 2002
Communicating Editor: Associate Professor Jim Warren

The guidelines are used by different audiences, including general practitioners, specialists, pharmacists, nurses, allied health professionals and students. In addition, guideline users often have different information needs ranging from the quick look-up of a particular piece of information during a clinical consultation to a more reflective read about a broader topic for educational purposes. The former needs efficient topic and key word index searching; the latter a descriptive table of contents and links to key references.

Users also expect to be able to view guideline information on a range of electronic devices ranging from full size desktop computer screens to the tiny screens of pocket computers and personal digital assistants (PDAs).

Re-engineering the production process and capturing the information in the Guidelines is not a mechanical exercise involving the simple manipulation of text. A sensitive user requirements analysis was required. The guiding principle was that the system must not impose any restrictions on how the Guideline narrative was expressed, except in cases where that was seen as a desirable trait to enforce standards.

For example, prescription regimens are expressed according to a grammar. Enforcing this grammar was seen as desirable as it increased consistency of expression. Enforcing a grammar on how a condition is described, however, was not seen as desirable as it is a much less structured area of the narrative.

Finally, Therapeutic Guidelines Limited is an independent organisation with a 25 year history. It is sustained solely by sales of its products. Its longevity indicates the quality of these products. TGL is thus well positioned as a platform from which guideline formats and authoring tools can be developed and disseminated.

2. SYSTEM REQUIREMENTS

The challenge in developing an efficient guideline production system is in maintaining consistent content across all viewing devices, as well as managing the complexity of the information needs and an increasing number of output formats. This is not simply a matter of reproducing the same content in a number of formats.

Based on the above, SISMO system requirements were defined as follows:

- Content must be maintained in a single authoritative source.
- Production must be efficient and minimise repetitive tasks.
- Output in multiple formats across multiple devices must be possible.
- Content providers are assumed to have medical expertise but not information technology expertise.
- Guideline structure and style should not be constrained by the system, except where that is a recognised and desirable outcome.
- Each output format must take advantage of the capabilities, and compensate for the short-comings, of each output device.
- Manipulation of information must be consistent with the domain being described.
- Content structure must have sufficient flexibility to allow for the ongoing introduction of additional elements.

3. SISMO SYSTEM ARCHITECTURE

This section describes the principle components of the production system and the connections between them. It demonstrates the modular approach to the architecture and the adoption of standard techniques where appropriate.

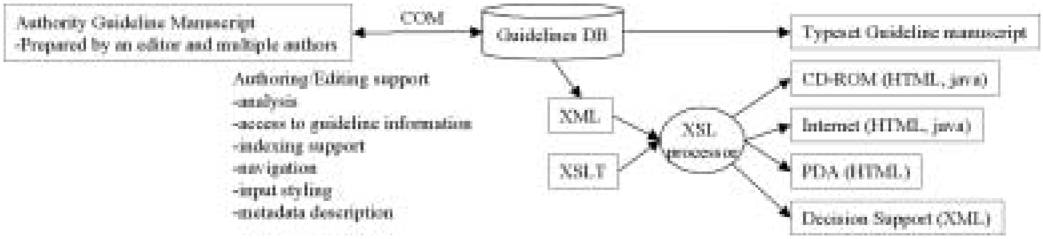

Figure 1: SISMO System architecture

Figure 1 depicts the broad architecture of the system. Guideline creation is carried out in the left half of the figure, while guideline output in various formats occurs in the right half.

3.1 Content authoring

A writing group initially authors guideline content, with the assistance of an editor. The draft manuscript is prepared in Microsoft Word, allowing continuing use of software that is familiar to them and which has been used previously to create guideline manuscripts.

Once the writing group has agreed upon the content, an editor indicates the structural elements of the guideline topics. In some cases, the structural elements imply a meaning or interpretation of the content. In these cases, confirmation from the writing group is sought.

The connection between the database and the Word document is achieved using a COM interface. Word macros trigger routines in a database application. The database application also manipulates the Word document through a COM interface.

A template is used to ensure certain Guideline elements are identified appropriately. The template will accommodate a range of generic guideline and technical documentation, thus the authoring process is broadly applicable outside of TGL.

The end result of the authoring process is a database of guideline topics and their metadata. The functionality of SISMO during this phase is described in following sections.

3.2 Output specification

Guideline output is sourced from the guidelines database. Having a single source for the output realises the efficiency gains of the production system. Once content has been verified, only the format of different output needs verification, rather than re-verifying content in each output format.

A particular output is achieved by extracting the guidelines from the database in an XML format. The XML source is then transformed to each of the desired output formats. The use of XML as a meta-format for clinical documentation has received considerable attention (Schweiger *et al*, 2002).

Dolin *et al.* (2001) describe the HL7 Clinical Document Architecture (CDA), which is a standard for the XML representation of clinical documents. While a CDA document can exist outside of a message, the CDA standard is aimed at representing reporting documents, 'such as discharge summaries and progress notes' (Dolin *et al*, 2001). The CDA deliberately does not model the document semantics, acting as a container for various document structures. The semantics of a SISMO generated document could be placed within a CDA wrapper for transmission. The utility of this is yet to be determined.

Extraction of data in XML format is carried out with a generic extraction mechanism developed as part of the project and described in (Lewis, 2002). The mechanism allows the user to specify through a user interface the portions of the database to be extracted and how they are related to each other. This is a flexible mechanism that can easily accommodate changes in the database structure.

The following two examples are XML and HTML fragments. The XML is given as input to an XSL processor within the authoring system. Also input is an XSLT file, which is a set of instructions that specify how the XML is to be transformed into HTML.

```
<REGIMEN DRUG_NAME = “metronidazole”> …
<DOSAGE … DOSES_PER_DAY = “3” DOSE_FREQUENCY = “8-hourly” DURATION = “for 7 to 10
days” ROUTE = “orally” </DOSAGE>
</REGIMEN>
```

Example 1: XML fragment

```
<HTML><BODY><h1>metronidazole</h1><TR>
<TD bgcolour=”#ff0000”>metronidazole</TD>
<TD bgcolour=”#993300”>10</TD>
<TD bgcolour=”#006600”>mg/kg</TD>
<TD bgcolour=”#6600CC”>orally, </TD>
<TD bgcolour=”#0000FF”>8-hourly</TD>
<TD bgcolour=”#FF9900” font-colour=”FFFFFF”>for 7 to
10days</TD></TR></TABLE></BODY></HTML>
```

Example 2: Example HTML fragment

A different XSLT file can be used to generate each output format. As well as being efficient, this mechanism increases the likelihood of the content of each output format conforming to the original XML source.

4. SISMO FUNCTIONALITY

This section describes SISMO’s core functionality, describing the means by which it is achieved.

4.1 Content creation

The core text of a Guideline title is a hierarchical set of Guideline topics. The initial draft of a title is a Microsoft Word document written with a pre-defined template. At this stage the template is only used to specify the available static structures of the manuscript. Also indicated through styles and a grammar are recommendations, regimens (including dosages, routes, frequencies, etc.), regimen preferences and the rationale behind a guideline.

Figure 2 is a cut-down UML (Universal Modelling Language) class diagram, depicting some of the elements that make up the conceptual model of a topic. Clinical (drug) concepts are depicted in the bottom of the figure while document concepts are depicted in the top. The ‘Content_Item’ concept, which holds the text of a guideline paragraph, thus links the clinical concepts with their documentation description.

Developing extensions of the topic model is certainly possible. As the current model is a natural representation of the topic structure, any natural extensions to it should be natural extensions of the topic model. For example, introducing summary statements naturally connects to the topic itself.

4.2 Indexing

The second stage of content creation consists of indicating the structural elements of the document, including the index terms. Indexing functionality consists of the presentation of an index frame that interacts with the document manuscript. The COM interface allows events in the document manuscript to trigger the index frame and vice versa.

Many of these actions alter the current document either by creating active text or by inserting optional text that indicates the structuring that has been specified.

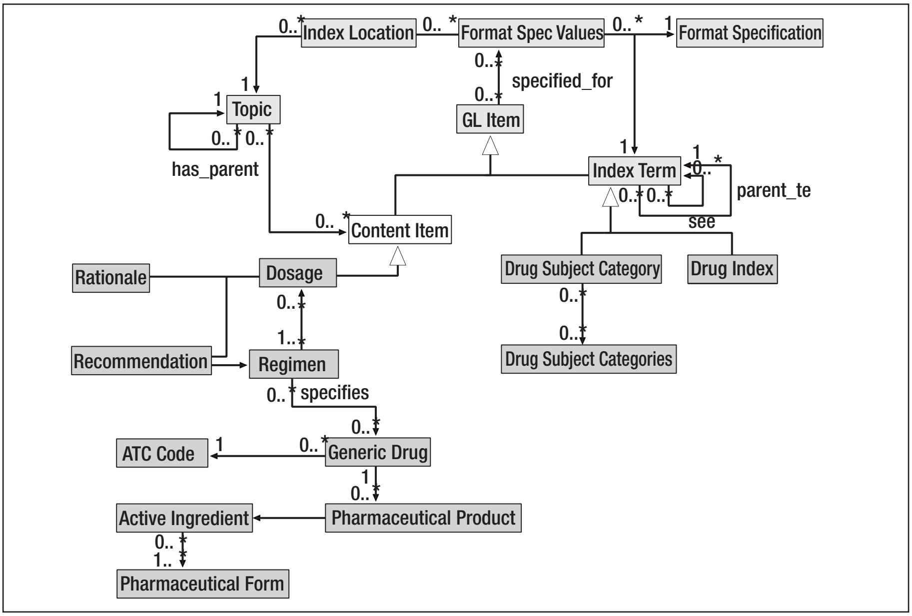

Figure 2: Stylised depiction of topic modelling

4.3 Output specification

Output specification allows the user to indicate the format that a given document should take in each output viewer. Although the mechanism for output generation is an XSL processor (as described in earlier sections), the user specifies the content to be placed in each output mode and format through a graphical user interface.

Figure 3 is a screenshot of a window that allows the user to specify the text of a topic heading in the book, electronic and electronic toc (table of contents) may vary. The heading text needs to convey the same meaning in different display contexts. As the context changes, the actual text used needs to change to retain the meaning.

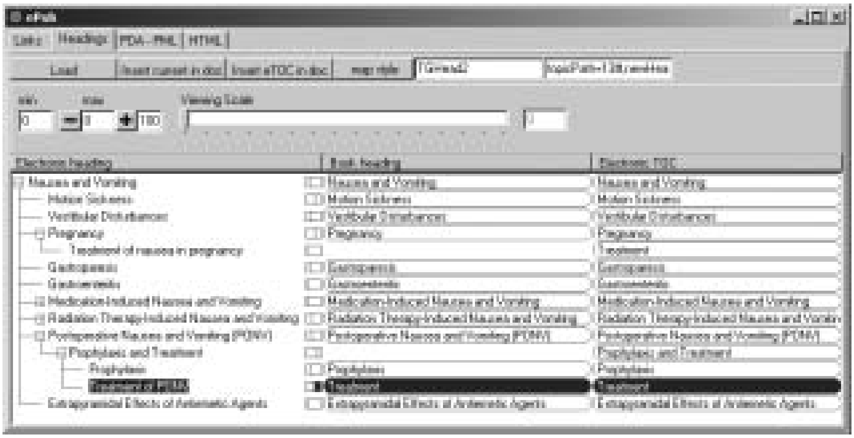

Figure 3: SISMO User interface – output specification

4.3.1 Topic size

Each topic is presented on a separate HTML page. The size of the page should be approximately a single screenshot, so that scrolling is minimised. The size gauge indicates relative topic sizes, allowing adjustment of the hierarchy as desired. Topics can be split and clustered to ensure that sizes are appropriate.

4.3.2 Topic headings

The 'Electronic heading' must serve as a standalone description of the topic, i.e. it should not be dependent on the context of the heading hierarchy for interpretation of its meaning. For this reason, any heading labelled 'Treatment' is usually edited to 'Treatment of …' (see eg. 'Post operative Nausea and Vomiting: Prophylaxis and Treatment: Treatment'. The 'Electronic TOC' label should not change in this case, however, as it is presented in the context of the heading hierarchy. All of these example alterations can be compared to the original 'Book heading', which is not altered by this process.

4.3.3 Cross references

As each HTML page must be standalone interpretable, it must contain in-page hyperlinks to direct relatives in the topic hierarchy. These links have a standard format and so their construction can be automated. After construction, the editors can choose to alter their format if desired.

4.3.4 Additional restructuring

Indexes (and Drug index pages, which collect information on each drug into a central point) are constructed automatically.

The PDA browser version is in development. Currently, a simple PML (Palm Markup Language) version can be automatically generated. Ultimately, additional restructuring and inclusion of summaries will be required for the PDA version.

5. DECISION SUPPORT

The remainder of the paper discusses how the project is positioning Therapeutic Guidelines for use as a basis for decision support.

In this section we characterise various forms of decision support, in order to create a framework within which our research and development in decision support can be discussed.

5.1 Previous work on Decision Support

We cannot survey the full extent of previous work on decision support in medicine here as a great deal of research has been carried out. A previous project paper (Wollersheim, 2001) has summarised some of that previous research. The conclusion of that paper was that no existing decision support representation was appropriate for modelling Therapeutic Guidelines content. The main reason for this is that the Guidelines are largely narrative and do not automatically lend themselves to restructuring according to a given representation.

While research indicates decision support mechanisms in medicine can have a beneficial impact on health outcomes and the cost of care, understanding of how information technology can influence decision-making in medicine effectively is not fully developed. The nature of the investigation is complex and a set of fundamental guiding principles for the implementation of decision support in medicine has not been articulated.

Given that decision support development and implementation is relatively immature, the project is adopting an investigative approach to the adaptation of the Guidelines for decision support. That is, we develop prototype systems that we anticipate may be effective and then test and evaluate the prototypes to increase our understanding of how the Guidelines are used in clinical practice.

5.2 Characterisation of Decision Support

The term ‘decision support’ and ‘decision support systems’ has been widely used over at least the last twenty years with varying intent. We briefly characterise some terms for the sake of discussion.

‘Decision support’ (DS) can be taken to mean the provision of information which assists in the decision making process. A ‘system’ implies the use of information technology in the process. eg. a Medline (Medline, 1997–) search. Usually, a decision support system will be a sub system that is treated as a resource by a host system.

Narrative decision support is the provision of narrative text that is at least partially structured so as to increase usability. eg. Cochrane reviews (Cochrane Collaboration, 1999–) and AusDI (AusDI, 1999–). An efficient index and search function is crucial to making narrative material readily available from the clinicians electronic desktop.

Active decision support is the provision of narrative text relevant to a particular system context. That is, the narrative is provided when system variables indicate that it is going to be useful. Active decision support may also take a non-narrative form that is capable of indicating appropriate subsequent processing, which the host system can then elect to carry out. Eg. Therapeutic Guideline regimen values being provided to a prescribing system.

Algorithmic decision support is the explicit representation of reasoning that has been previously implicit within narrative decision support.

eg. IF severe pneumonia (previously defined) AND patient has one or more risk factors (diabetes, alcoholism, age greater than 50 years or Aboriginal or Torres Strait Islander descent), THEN (treat ..)

Guidelines that provide content that can enable ADS are known as Computer-Interpretable Guidelines (CIGs) (Peleg *et al*, 2002). The six CIGs presented in Peleg are restricted to non-narrative forms, i.e. the guideline ‘text’ is not presented to the user. (Wang *et al*, 2002) provides a review of eleven CIGs, identifying certain ‘primitives’ that are consistently present in these representations.

While significant progress has been made on the development of different CIG representations, there are two major (interrelated) problems that are encountered when assessing their suitability as representations of the TG guidelines; integration of narrative with an algorithmic representation and maintenance of a knowledge base.

All the CIGs we have surveyed take the development of an algorithmic guideline as a starting point. A few of them then consider how the algorithm may then refer the user to a narrative. We are considering the opposite point of view. Our starting point is a usable and effective narrative that is already providing decision support. How an algorithmic representation of portions of the guideline model can enhance that usability and effectiveness is the focus of the investigation.

Arden syntax (American Society for Testing and Materials, 1992) specifies a means of expressing simple logic statements. Application specific implementations have used it to control event management through triggering. The Guidelines Interchange Format (GLIF) relies on a guideline model expressed as a Document Type Definition (DTD) (Ohno-Machado *et al*, 1998). The GLIF model is too inflexible to accommodate the narrative of the Guidelines however.

Asbru (Shahar *et al*, 1998) is a formal representation language for modelling guidelines. It is intended to be a representation for structured protocols rather then narrative guidelines.

PROforma is also a formal representation language (Fox and Das, 2000). PROforma is based on first-order logic but also incorporates non-classical logics. It has been used to model oncology guidelines as well as relatively structured text such as drug formularies.

Given the above characterisation and background, we next present how the project's developments can enable Therapeutic Guidelines to be used to provide each type of DS.

5.3 Narrative Decision Support

Therapeutic Guidelines (TG) has been used as the basis of narrative decision support for some time. A requirement for sustainable provision of narrative DS is a production process that allows the consistent presentation of TG content across the range of browsers and operating systems used in Australian health care settings, accessed by efficient indexing. SISMO has reinforced the ability of TGL to maintain the production of the formats necessary by increasing its production efficiency.

SISMO has also increased control over the format of the output produced, allowing Guideline comprehension to be enhanced through more usable presentation of information.

5.4 Active Decision Support

Requirements for active DS include:

- a means of linking a Guideline topic to a given clinical scenario.
- sufficient structural description of Guideline information to allow guideline processing on the host system.

Initial attempts at integrating TG with a prescribing system (Nolan *et al*, 1999) were technically successful. The difficulty was that the means of achieving linkage between the prescribing system and TG required manual tagging of Guidelines.

A sustainable linkage between a given clinical condition and its Guideline topic could be achieved through the application of a health classification or coding scheme that covered all the conditions described in the guidelines and that was in reasonably uniform use in Australian health care settings. ICD-10-AM codes have been assigned to the majority of Guideline topics, unfortunately ICD-10-AM does not adequately describe all topics and is not in use throughout the whole of the Australian health care system.

Processing of portions of a Guideline topic is now possible due to the increased markup of topics. The listing in Example 1 shows an XML fragment of a regimen. The individual tag values in the fragment would allow a prescribing system to prepare a draft prescription.

6. ALGORITHMIC DECISION SUPPORT

A prototype facility that provides algorithmic decision support (ADS) has been implemented, through a server application. The knowledge base of the server contains the explicit representation of some of the reasoning that is implicit in the narrative of Therapeutic Guidelines. The narrative supports the ADS by explaining and justifying the reasoning, as well as indicating instructions and outcomes that follow from the results of reasoning.

6.1 System architecture

The system architecture of the ADS separates the representation of the reasoning, the instantiation of the representation and the implementation of the representation into a database layer, object layer and application layer, respectively (Figure 4).

The database maintains the data that represents the algorithm. Data integrity and consistency is enforced through input checks, triggers and stored procedures. These editing constraints ensure that any algorithm data created conforms to a minimal level of consistency.

Access to the database is only possible through an object layer, which further reduces the possibility of inappropriate or unsafe use of the database and hence of the algorithm. A decision

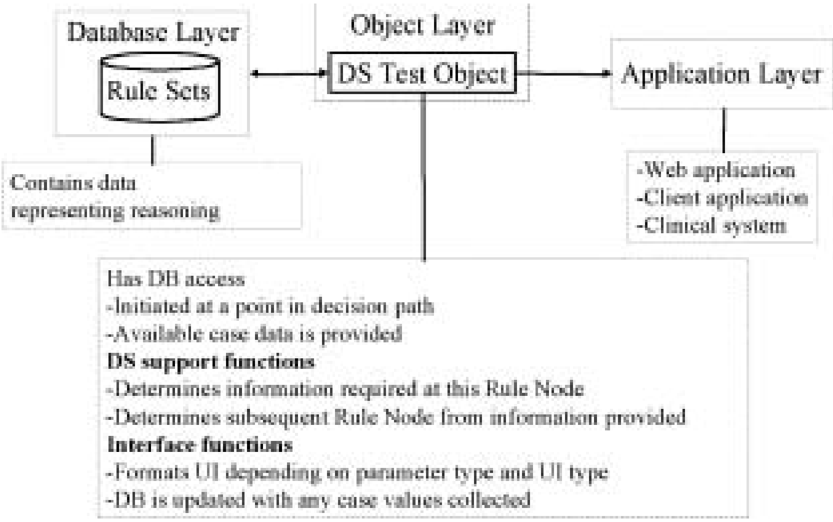

Figure 4: Algorithmic decision support – System architecture

support (DS) object accesses particular parts of the database according to a set of retrieval procedures and input parameters. The object layer is also provided with any available case data, which is stored as it is collected.

The application layer provides a means of manipulating the object layer through a variety of interface mechanisms. A web application interface is provided through the dynamic generation of HTML. Input parameter values are collected through HTML forms and the results of reasoning are similarly presented. A client application (MS Windows implementation) that dynamically generates interface forms operates in a similar fashion.

A third possible interface is direct to a clinical system, which could provide any relevant case data directly and return the results of reasoning in a structured format. The generation of the web interface is via XSLT transformation of XML, so the provision of XML directly is easily achievable.

Each level of the architecture enforces a level of data integrity and algorithm validity. These mechanisms will not guarantee the safety of the algorithm, however, there are advantages to dealing with different aspects of the algorithm's structure in this layered fashion.

The first advantage is that mechanisms appropriate to each layer are utilised. For example, data integrity mechanisms are appropriately dealt with at the database layer, while interpretation and presentation of the data are dealt with in the object layer. The second advantage is that it is clear what safety issues have been addressed. Additional mechanisms can be subsequently developed from a well-understood basis.

6.2 Database Structure

6.2.1 Core data model

A full data model of the ADS server database cannot be presented here. Figure 5 shows the core data model.

An algorithm consists of a set of **Rule_Nodes** connected together as an acyclic graph. Each algorithm has one root node. Each **Rule_Node** that is not a leaf indicates a subsequent **Rule_Node**. The indication is made by an **RN_Indicates_RN**, which also indicates the **Result** condition under

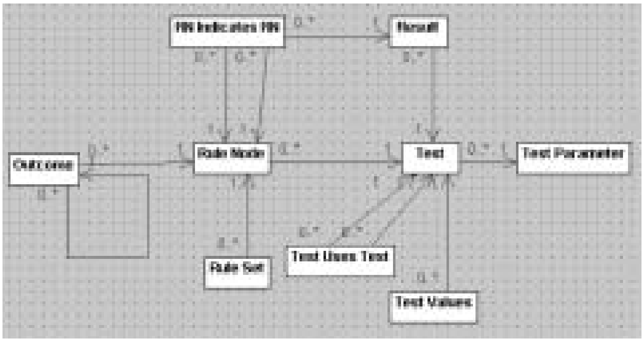

Figure 5: Core ADS data model

which the **Rule_Node** indication is applicable.

The **Result** is dependent on the evaluation of a **Test**. A **Test** may have a number of sub **Test**s (indicated through a **Test_Uses_Test**). Each **Test** has a **Test_Parameter** and a set of **Test_Values**.

A **Test_Parameter** indicates the data type of a **Test_Value** and the case value the value of the **Test_Value** corresponds with. **Test_Parameter** data types may be boolean, integer, list, set or float.

A **Test_Value** is a particular condition that a case value can be applied to. A **Test_Value** also holds the value of a prompt to indicate to a user what value is required for this test to be evaluated. If the condition is satisfied, the **Test_Value** will indicate what **Result** the associated **Test** should indicate in determining a subsequent **Rule_Node**.

A leaf **Rule_Node** may indicate an **Outcome**. An **Outcome** is specification of string formatting that takes case values as parameters to generate a result string. This can be used, for example, to generate dosage calculations dynamically, based on case values. An **Outcome** may also indicate a guideline topic, when the results of reasoning indicate that a particular Guideline topic is relevant to the current scenario and case values.

6.3 Algorithmic Decision Support Server

Figure 6 shows how the system architecture can be configured to produce an algorithmic decision support (ADS) server. The ADS server controls all access to the object layer, and hence the algorithm itself. The ADS server is provided with client identification, request description and case parameters.

The request description identifies the algorithm to be executed and the case values provide values for the algorithm's input parameters.

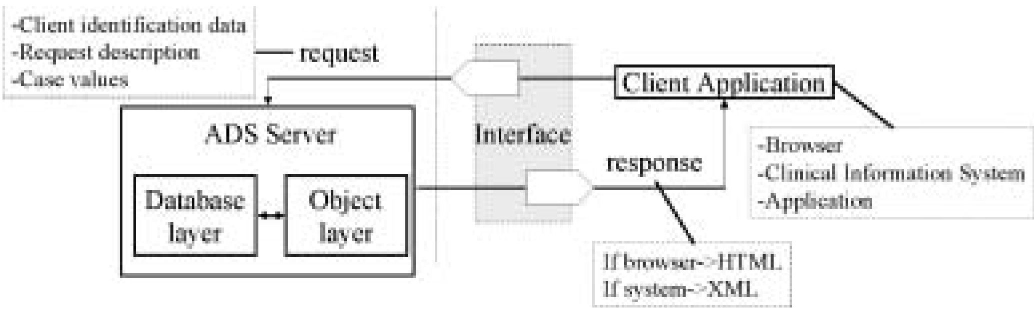

Figure 6: Algorithmic Decision Support Server

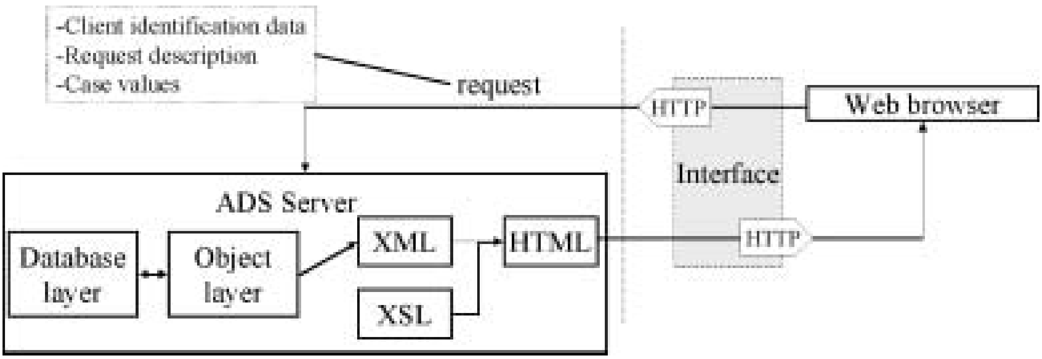

Figure 7: Algorithmic Decision Support Web Server

A prototype ADS server has been developed as a World Wide Web server. Figure 7 depicts the web server architecture.

If a clinical system was sending the requests, rather than a web browser, the untransformed XML could be delivered directly and the receiver could process or display it according to internal customisations.

6.3.1 Web Server Example

To demonstrate use of the ADS web server, we present a walkthrough example. The algorithm being used determines the risk severity for Community Acquired Pneumonia (CAP). All the content of the HTML pages presented has been dynamically created, based on data retrieved (via the object layer) from the ADS database.

Figure 8 shows the (partially collapsed) rule tree of the CAP algorithm. This is available for viewing from any of the nodes in the tree. The rule tree is a dynamic HTML page that explains each node as the user hovers over it. In the figure, the properties associated with 'Procedure – Chest X-Ray' are being displayed.

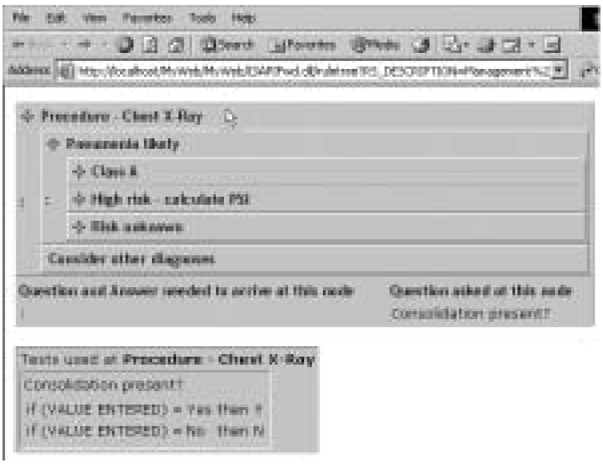

Figure 8: Overview of the CAP Rule tree

Figure 9 shows the first interaction with the algorithm at the root node. The input objects in the HTML form are named so as to capture the case values as they are entered. The subsequent HTML pages will be generated based on the values the user enters into the form.

The Rule_Node being examined in the figure is ‘Procedure – Chest X-Ray’. Two subsequent Rule_Nodes are possible: ‘Pneumonia likely’ (corresponding to Result 1) and ‘Consider other diagnoses’ (corresponding to Result 2).

There is one Test associated with the Rule_Node: ‘Consolidation present?’. The Test_Parameter associated with ‘Consolidation present?’ indicates that the test is boolean and corresponds to case value ‘chest_x_ray/consolidation’.

There are 2 Test_Values associated with ‘Consolidation present?’. ‘Y’ indicates Result 1 and ‘N’ indicates Result 2.

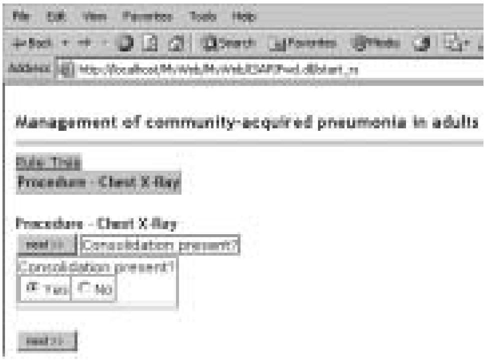

Figure 9: Root node of the CAP algorithm

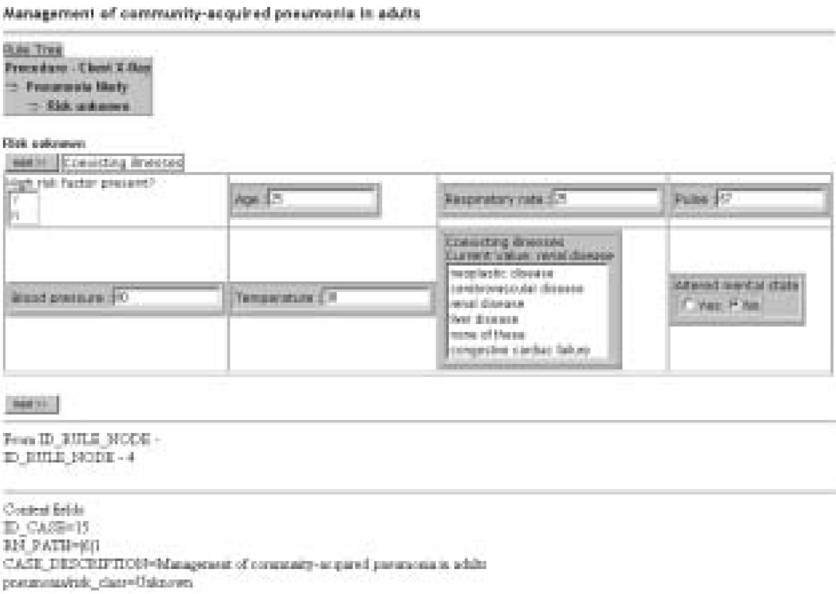

Figure 10: ‘Risk unknown’ node of the CAP algorithm

Figure 10 shows a subsequent rule node ('Risk unknown') and the set of input values that the algorithm requires at his node. An explanation of how the value associated with each input object will be tested is available to the user through a dynamic HTML object.

At the bottom of the page the case values that have been collected from the user so far are displayed. They would not normally be displayed in this manner, but are shown here for illustration.

There are eight Tests associated with this node, corresponding to Test_Parameters of type integer, list and boolean. The HTML form generates input boxes for integer and float, radiobuttons for boolean and list objects of less than four members and list boxes for list objects of four or more members. A set type indicates that the Test has sub Tests and that a sub form is required.

7. CONCLUSION

We have presented a production system for Therapeutic Guidelines (SISMO), which included the architectural design of the system as well as a description of its functionality. This system has been introduced into the production process at TGL and its further development will be based on the results of live user testing and further functional analysis.

SISMO is capable of managing the complex narrative text of the Guidelines, as well as the structured 'facts' (drug names, dosages, etc.) in the Guidelines. SISMO will be used to produce an increasing range of Guideline formats, leading to the increased use of decision support based on Therapeutic Guideline content.

The decision support that is envisaged includes Narrative Decision Support, Active Decision Support and Algorithmic Decision Support. We have developed a prototype Algorithmic Decision Support (ADS) server, which is capable of providing ADS content via a variety of interfaces.

The system architecture of the Algorithmic Decision Support server has been designed to separate the representation of the algorithmic reasoning, the instantiation of the representation and the implementation of the representation. This separation has demonstrated its advantages in simplifying the implementation of a number of interfaces to the ADS server and in ensuring a defined level of safety and quality is maintained.

8. ACKNOWLEDGEMENTS

This project could not have been possible without the support and contributions of Dr. Ken Harvey, Dr. Jonathan Dartnell, Ms. Mary Hemming, Mr. Dennis Wollersheim, Ms. Elizabeth Deveny, Assoc. Prof. Teng Liaw, Prof. Michael Kidd and Prof. Enrico Coiera.

This paper describes research undertaken as part of Strategic Partnerships in Industry, Research and Training (SPIRT) project. The research has been funded by the Australian Research Council, National Prescribing Service, Therapeutic Guidelines Limited and the General Practice Computing Group.

The paper has also benefited from anonymous reviewers comments and suggestions.

9. REFERENCES

AUSDI (1999–): *Australian drug information for the health care professional*. Curtin, ACT, Pharmaceutical Care Information Services.

COCHRANE COLLABORATION (1999–): *The Cochrane library [electronic resource]*, Health Communication Network.

DOLIN, R. H., ALSCHULER, L. *et al*. (2001): The HL7 clinical document architecture. *Journal of the American Medical Informatics Association* 8: 552–569.

FOX, J. and DAS, S. (2000): *Safe and sound: artificial intelligence in hazardous applications*. Menlo Park, Calif, AAAI Press/MIT Press.

HEMMING, M. (2000): Therapeutic guidelines: an Australian experience. *International Journal of Pharmaceutical Medicine* 14(4): 259264.

LEWIS, B. (2001): Therapeutic guideline development for integration with general practice information systems. *Realising Quality Healthcare: Health Informatics Conference*.

LEWIS, B. (2002): Extraction of XML from relational databases. *XML Based Data Management (XMLDM)*, Springer-Verlag.

AMERICAN SOCIETY FOR TESTING AND MATERIALS (1992): Standard specification for defining and sharing modular health knowledge Bases (Arden syntax for medical logic modules). *ASTM Standards* 14(1): 539–587.

MEDLINE (1997–) [*electronic resource*]. Bethesda, Md., National Library of Medicine.

NOLAN, A., NORQUAY, C. *et al*. (1999): Electronic prescribing and computer-assisted decision support systems. *Medical Journal of Australia* 171: 541–543.

OHNO_MACHADO, L., GENNARI, J. H. *et al*. (1998): The guideline interchange format: A model for representing guidelines. *Journal of the American Medical Informatics Assocation* 5(4): 357–372.

PELEG, M., TU, S. *et al*. (2002): Comparing models of decision and action for guideline-based decision support: a case-study approach (Part 1 of 2), Stanform Medical Informatics.

SCHWEIGER, R., HOELZER S. *et al*. (2002): Plug-and-Play XML: A Health Care Perspective. *J. Am. Med. Inform. Assoc.* 9(1): 37–48.

SHAHAR, Y., MIKSCH, S. *et al*. (1998): The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial Intelligence in Medicine* 14(1–2): 29–51.

WANG, D., PELEG, M. *et al*. (2002): Representation primitives, process models and patient data in computer-interpretable clinical practice guidelines: A literature review of guideline representation models. *International Journal of Medical Informatics*.

WOLLERSHEIM, D. (2001): A review of decision support formats with respect to therapeutic guidelines limited requirements. *Health Informatics Conference*. Canberra, Health Informatics Society of Australia Ltd.

WRITING GROUP FOR THERAPEUTIC GUIDELINES: ANTIBIOTIC. THERAPEUTIC GUIDELINES LIMITED (2000): *Therapeutic Guidelines: Antibiotic*. North Melbourne, VIC., Therapeutic Guidelines Limited.

BIOGRAPHICAL NOTES

Dr Bryn Lewis is undertaking a post doctoral research fellowship ‘Towards computerised clinical decision support systems: optimising the production of collaborative electronic guidelines’.(see http://homepage.cs.latrobe.edu.au/lewisba/frindex.html). This project follows from doctoral research in artificial intelligence and machine learning applied to automatic classification of documents for digital libraries.

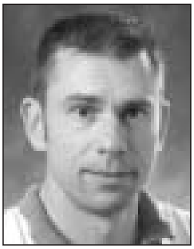

Dr Bryn Lewis