Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)

# The Benefits of an Ontological Patient Model in Clinical Decision-Support

**Mark Austin** and **Matthew Kelly** and **Sir Michael Brady**
Wolfson Medical Vision Laboratory, Robotics Research Group
Department of Engineering Science, University of Oxford
Parks Road, Oxford OX1 3PJ, UK
Tel: +44 1865 283179
mark.austin@worc.ox.ac.uk

## Abstract

We have developed an application, *MDTSuite*, designed to support complex group decision making under difficult conditions including time pressure, incomplete information, changing group members and ever expanding guidelines. Taking the colorectal cancer MDT (Multi-Disciplinary-Team) meetings at the Radcliffe Infirmary in Oxford as our test case, we have been trialling the software live in the hospital providing decision support for clinicians discussing real patients. *MDTSuite* is an application integrating an ontological data model with an argumentation-based decision-support system, showing how the combination of leading technologies OWL (McGuinness & Harmelen 2004), SPARQL (Perez, Arenas, & Gutierrez 2006) and Jena (Carroll *et al.* 2003) can be used to achieve this.

## Introduction

The MDT meetings are forums for the different clinical specialists (surgeons, radiologists, pathologists, oncologists and specialist nurses) to combine knowledge from their respective fields to form a group decision on treatment options for each patient in their care. *MDTSuite* has been built around an argumentation-based decision-support engine similar to the decision-support mechanism introduced in PROforma (Bury, Fox, & Sutton 2000). For each decision to be made, a number of candidates are defined as possible options to be taken for that decision, each of which has a number of potential arguments either for or against that candidate, which are displayed to the users if a logic statement associated with that argument evaluates true. In this way, a complex combination of national and local clinical guidelines can be encoded into these logic statements to bring the guidelines into the clinical workflow. *MDTSuite* also supports this decision making process through efficient patient data presentation software which summarises information taken from many different sources.

The data structure used in most medical decision-support languages is some form of flat data storage, either storing data as key-value pairs, or in a database with a fixed number of columns (Robert & Raymond 2006) and the data structure of *MDTSuite* originally followed this flat format. In the development and live testing of the software, we discovered several limitations which, we contend, derive directly from the use of such a flat data structure, some of which are likely to make such decision-support systems impractical for medical use.



Colorectal cancer patient:
Tumour in upper rectum:
Anterior, mobile, T1N0M0, Adenocarcinoma, Moderately differentiated, CRM 7mm
Tumour in caecum:
T3N0M0, Adenocarcinoma, Well differentiated

Figure 1: *Example of a patient which causes a problem for a flat data structure*

For example, consider an example patient with two tumours as described in Figure 1. To describe this patient using a flat data structure, each piece of data for each tumour would be represented by a separate datatype, thus two sets of tumour stage, two sets of histological grade etc. would need to be present. Now, imagine the problem if a third tumour was to be found in the patient. Should the data structure then be modified to contain three such sets of data? Such cases illustrate clear limitations of using such a database design. If we were to create the fields required to describe the patient in Figure 1, they would then have to appear for every patient in the system, so for the majority of patients with only one tumour, a lot of irrelevant fields would be present. Also, notice that some of the data for rectal cancers is not required for colon cancers, for example the CRM (Circumferential Resection Margin - this is the minimal clearance between the boundary of the mesorectum and the affected tissue), but in a flat data structure, this field would necessarily be present for colon cancers too despite its irrelevance for such cancers.

Such extra fields would result in an increase in input time for the system since a user would have to wade through more fields in order to define a patient. Ontologies (Noy & McGuinness 2001) offer a potential solution to all of these problems as well as providing an excellent data structure platform for the integration of multi-disciplinary clinical data including multimedia data. Perhaps for these reasons, they are the subject of a great deal of research in medical informatics, for example in projects such as HL7 RIM (Smith & Ceusters 2006; Schadow *et al.* 2006), SNOMED (SNOMED 2006) and the NCI Thesaurus (NCI 2006). Much of this research has focused on using ontologies to create a common language with which medical systems and centres can communicate (Smith & Ceusters 2003; Ceusters, Smith, & De Moor 2006). However, an ontology can also be used to create a richer, more complete representation of the domain over which a computer can reason more effectively in a practical application, providing benefits beyond information integration. This more complete ontology-based representation enables us to describe the patient in Figure 1 fully, as well as handling even more complicated cases, whilst at the same time maintaining an efficient, smaller set of inputs for less complex patients. The relationships between data inherent in an ontological data structure enable an increase in the relevance of input fields such that only the required fields are displayed for each patient. These inherent relationships also facilitate a consistency check of the data capable of reducing the number of human input errors in the system.

## Ontology-Driven MDTSuite

While we have focused our trial of the software on colorectal cancer, *MDTSuite* has been written in such a way that the knowledge base is completely separated from the application, so the software can be converted for use in any group decision making environment simply by writing a new rulebase. We have authored an OWL colorectal cancer ontology to replace the flat-file patient data model used originally and refactored the software to handle this model using the Jena java API. A data-driven, AJAX (Garrett 2006) powered patient data input form is generated automatically from the ontology and patient summaries are generated automatically from the populated patient instances. The logic statements associated with the arguments are written in the W3C recommended SPARQL query language, and this rulebase can be created and maintained within the software itself.

The demonstration will take the audience through the patient journey for a single patient being discussed at the MDT meetings over several weeks as well as provide an overview of how the rulebase is authored and maintained, showing clearly how different disease domains, or even non-medical domains could be handled. The software is initially loaded with a disease domain OWL ontology, as well as an XML file containing the argumentation logic in the form of SPARQL queries. Patient data is added via the AJAX based flexible input form before the MDT meeting. During the meeting, an automatically generated summary of each patient instance is displayed as the patient is discussed, before the software displays a decision support page showing the various treatment options (the candidates) and the arguments for and against each candidate generated by the SPARQL queries. Recommendations are made for those candidates with more arguments for than against according to the net-support algorithm present in PROforma (Bury, Fox, & Sutton 2000). Once the group comes to a decision, the candidate must be selected, and where a non-recommended candidates is selected, a reason for the deviation from guidelines must be provided. This selection, along with the arguments which led to that recommendation and any comments, are then stored in a corporate memory to form part of the patient summary when the patient returns to be rediscussed in the future. By scanning all the arguments relating to a particular decision, the software generates a property tree of information which is required to fully describe the patient instance such that all of the necessary rules for each patient are evaluated fully. Any of this required data missing at the time of the meeting is highlighted along with how it might impact the recommendations and this data can be added on the fly during the meeting via a pop up input box. Finally we will show the rule authoring tool provided so that the users can add new candidates and arguments as well as the SPARQL logic associated with them, allowing the rulebase to be maintained and extended by users who need not be expert programmers.

## References


Bury, J.; Fox, J.; and Sutton, D. 2000. The proforma guideline specification language: progress and prospects. *Proc First European Workshop, Computer-based Support for Clinical Guidelines and Protocols*.

Carroll, J.; Dickinson, I.; Dollin, C.; Reynolds, D.; Seaborne, A.; and Wilkinson, K. 2003. Jena: Implementing the semantic web recommendations.

Ceusters, W.; Smith, B.; and De Moor, G. 2006. Ontology-based integration of medical coding systems and electronic patient records. *http://ontology.buffalo.edu/medo/Coding-AndEHCR.pdf*.

Garrett, J., ed. 2006. *Ajax: A New Approach to Web Applications*. http://www.adaptivepath.com/ideas/essays-/archives/000385.php.

McGuinness, D., and Harmelen, F. v. 2004. Owl web ontology language overview. *W3C Recommendation*.

NCI. 2006. Nci thesaurus. *http://nciterms.nci.nih.gov*.

Noy, N. F., and McGuinness, D. L. 2001. *Ontology Development 101*.

Perez, J.; Arenas, M.; and Gutierrez, C. 2006. Semantics and complexity of sparql.

Robert, M. J., and Raymond, J. K. 2006. Electronic medical records and medical research databases can they be synonymous? *Business Briefing: US Cardiology*.

Schadow, G.; Walker, C.; Mead, N.; and Mead, D. 2006. The hl7 reference information model under scrutiny. *Studies in Health Technology and Informatics Vol 124*.

Smith, B., and Ceusters, W. 2003. Towards industrial-strength philosophy: How analytical ontology can help medical informatics. *Interdisciplinary Science Reviews*.

Smith, B., and Ceusters, W. 2006. Hl7 rim: An incoherent standard. *Studies in Health Technology and Informatics Vol 124 133138*.

SNOMED. 2006. Snomed clinical terms (snomed ct) core content for the july 2006 release. *http://www.snomed.org-/snomedct/documents/july_2006_release.pdf*.