Representing Clinical Knowledge in Oral Medicine Using Ontologies: Problems and Possible Approaches *

Marie Gustafsson

Department of Computing Science, Chalmers University of Technology
SE-412 96 Göteborg, Sweden
School of Humanities and Informatics, University of Skövde
Box 408, SE-541 28 Skövde, Sweden
mariegus@cs.chalmers.se

Abstract

How can information technology be used to model and handle clinical knowledge in everyday work so that clinicians can more systematically learn from collected clinical data? This has been studied within the MedView project, a collaboration between Chalmers and the Clinic of Oral Medicine at the Sahlgrenska Academy, for the past ten years. This paper will address some problems with the current model, and how OWL (Web Ontology Language) and RDF (Resource Description Framework) might be used to alleviate these problems. This approach could, for example, aid in data sharing, offer a way to look at the data from different perspectives, and give a possibility of importing ontologies of diagnoses and pharmaceuticals.

1 Introduction

The MedView computer system is the product of about ten years of cooperation between the Chalmers University of Technology and the Clinic of Oral Medicine at the Sahlgrenska Academy. The overall goal is to develop models, methods, and tools to support clinicians in their daily work and research. A central question is how information technology can be used to model and handle clinical knowledge in everyday work so that clinicians can more systematically learn from collected clinical data.

When a clinician examines a new patient, he or she enters the gathered data using the MedRecords application. Data includes not only textual answers to queries, but also digital photos of the oral mucosa, for example, taken during examinations. The clinical protocol that guides the examination is given by an examination template, presented as a form on the clinicians clinical work-station, a template that has previously been defined by the clinician him- or herself using the FormEditor application. In subsequent steps, one of the output applications, e.g. MedSummary, Aggregator, MedViewer, mVizualizer, and mEduWeb, can be used to examine and visualize information from the database. These are aimed at analysis, interpretation, and learning, and can be based both on an individual patient and on a selection of groups of patients from the database [11].

MedView is an initiative to implement evidence-based oral medicine based on the view that health care should be built on finding, validating, and using the latest research

*The work presented in this paper was supported by the Swedish Agency for Innovation Systems.

results as a basis of clinical decisions [27]. To practice evidence-based medicine means integrating the expertise among individual clinicians with the best clinical evidence that can be obtained from external sources [28]. Among the problems needing to be solved to achieve evidence-based medicine are that the right data is collected, which is connected to asking the right questions and motivating the clinicians to collect the data. Consequently, it is also desirable to increase the possibility of sharing collected data and extending support for different kinds of analysis.

Over the past decade, there has been much research in using ontologies for knowledge representation. Though there are many different definitions of ontology, one of the most often used is Gruber's [7] "explicit specification of a conceptualization." Recently, both the Resource Description Framework (RDF) [14] and the Web Ontology Language (OWL) [29] have become World Wide Web Consortium (W3C) recommendations. RDF is used to represent information and exchanging knowledge, while OWL is used for publishing and sharing ontologies. There has also been much activity in developing applications and programming frameworks for creating and managing ontologies.

Given these current developments, we propose that ontologies are useful in dealing with the problems faced by an evidence-based system such as MedView. Adapting to standards such as OWL and RDF will increase possibilities of data sharing, as well as giving the possibility of using tools developed by others in this process. This use of ontologies will contribute to knowledge representation in oral medicine, where there is little work using this approach. It will also contribute to the work of developing models and techniques for ontologies by testing these on a new domain. Further, the MedView project is ready to be taken to another level. This extension is part of SOMWeb (Swedish Oral Medicine Web), a project for net-based knowledge management as support for clinical practice and knowledge sharing.

The paper will give an analysis of the existent knowledge representation in MedView and identify some of its problems (Section 2). Given these problems, we will discuss how the conceptual model can be improved and extended, and investigate whether an approach using OWL and RDF might be fruitful in developing a new model (Section 3). An early implementation of such a model is then described (Section 4).

2 Analysis

Clinical data in MedView has thus far been seen as *definitions* of clinical terms [5]. Abstract clinical concepts, e.g., diagnosis, examination, and patient, are given by definitions of collections of specific clinical terms.

2.1 Knowledge Representation in MedView

At present, the clinical knowledge used in MedView, is divided into templates, value lists and value classes. In addition to these basic knowledge structures, there are also definitions defining the layout of text and slot-fillers used in the generation of summaries of examination records in the MedSummary application, and definitions defining the structure and layout of templates used in the acquisition of examination data.

For example, the terms status, direct, mucos and palpation are all part of the general *template* that defines a particular clinical *examination protocol*. A concrete instance of an examination template—an examination record—is given by defining terms like mucos_site and mucos_col in terms of observed *values*, e.g., l12 and white, brown respectively.

The knowledge base (KB) also contains knowledge structures describing general domain knowledge. Values for the terms defined in templates are taken from formalized lists of valid values. These *value lists* are given as value definitions, which are stored in the knowledge base along with the examination records and templates.

As the KB grows, it becomes increasingly important to be able to group related values into classes in a hierarchical manner. For example, diseases such as Herpes labialis, Herpetic gingivostomatis, and Shingles can be classified into viral diseases. The ability to categorize values into different classes (or groups) has proven very useful in data analysis in that they reduce the complexity of the data set, facilitating the detection of interesting patterns in the data. Value classes can also be useful for concept formation, e.g., for differentiating between two different forms of a diagnosis.

The knowledge base built in the MedView project currently contains data from over 6000 examination records, covering more than 2000 different cases. The main knowledge base is located at the clinic of oral medicine in Göteborg. The various clinics within the Swedish Oral Medicine Network (SOMNET) have local knowledge bases containing the examination records collected at each clinic. The contents of these local knowledge bases are added regularly to the knowledge base in Göteborg so that the entire amount of data collected can be accessed through one common knowledge base.

2.2 Problems with the Current Representation

On the basis of the experience gained during the first decade of the use of MedView and on recent interviews with domain experts and developers, the following problems with the current knowledge representation in MedView have been identified:

- We must be able to better utilize external sources of knowledge, e.g., general medical vocabularies, taxonomies of disease and medications, and externally defined ontologies. Currently all necessary knowledge is contained within MedView itself. This means that any updates to the knowledge base, for example new pharmaceuticals, must be performed manually and within the project. Further, knowledge sharing with other projects is made difficult by differences in concepts and representation. Faster sharing of knowledge is a prerequisite for effective evidence-based medicine. The importance of constructing systems that aim to facilitate the search for evidence from heterogeneous information sources based on clinicians needs is noted by Mendonca et al. [16].
- The conceptual models of examination templates, lists of approved values for terms, and groups of related terms are central in MedView. These have not been sufficiently explored. There is also a need to examine the relation between these models (and their corresponding concrete entities) and a general ontology for oral medicine.
- How can relations and interactions between different entities of the conceptual models be captured? For example, how can we represent the fact that a certain answer to a specific question in a given examination template should trigger another question, or limit the possible answers to a subsequent question? This is important both for making sure that necessary information is captured, as well as limiting the amount of questions to be answered. Not having to fill in responses that seem redundant can help motivate the clinician to answer the relevant ones.
- We need a stronger typing of elements in the knowledge model. How can we enforce that a certain term only has numeric values, that a term has dates as values or that a term has a certain enumerated domain? The representation of quantities and units illustrates this problem well, e.g. one might find entries both for 5,0 ml, 5.0 ml, and 0.005 L. This problem is exacerbated by the general problem in the area of medicine of agreeing on canonical terms for data [2], such as conforming to one way of denoting a certain diagnosis.
- We must be able to capture different kinds of meta-data, e.g., who is the creator of a specific examination template, which examination template has been used to generate a specific examination record and what the purpose (scientific or clinical)

is of the introduction of a specific examination template? Fernández-López and Gómez-Pérez [6] discuss the importance of this kind of meta-data for knowledge re-engineering purposes, for example. In addition to the content-oriented parts of the ontology, we need to consider attributes describing concepts that have to do with the construction and evolution of ontologies [4].

- The localization of data has to be addressed more rigorously: How to provide different language-based versions of the defined concepts, definitions and terms? This is important since MedView is used at clinics outside Sweden. Further, a prerequisite of evidence-based medicine and knowledge sharing at a global level is making transitions between language borders transparent.

- We need to differentiate between different 'views' of the underlying data. In addition to the examination-centered (qualitative) view provided by the current model, how can we provide, for instance, a patient, time or quantitative oriented view of the underlying data?

  The traditional view of a electronic patient record is that it provides a source-oriented view of the underlying clinical (and administrative) data, in contrast to older paper-based patient records which provide a more time-oriented view [32]. In general, what we want is a conceptual-oriented view, in which, in principle, every important concept in the knowledge model at hand can be the target for a specifically generated view of the underlying clinical data.

  The provision of different 'views' of the ontology and the manipulation of these views are an intrinsic part of ontology management, since it gives both cognitive and computational benefits, in terms of clearer cognitive semantics and smaller computational footprints [18].

- How can information contained in images, i.e., photos taken during the examination of patients, be captured and represented? An increasingly larger portion of medical data has its origin in images. The enormous amount of information obtainable from images is, however, difficult to grasp for the unaided human mind [2]. On the other hand, as compared to computers, human experts are still more proficient at detecting interesting trends and patterns in image-based data. This calls for more work on information visualization, computer-human interaction and the cognitive processes involved in the exploration and discovery of (new) medical knowledge [26].

Some of these problems are addressed in the following.

3 An Ontology-Based Approach

It has been argued that medical informatics is the bridge between the domain of medicine and practitioner psychology [8]. Ontologies is a crucial element in building this bridge, in that ontologies try to capture the best existing medical knowledge, through the systematic and explicit representation of medical definitions, concepts and processes, and by providing a tool for constructing computer support systems that help clinicians to perform clinical tasks more effectively and efficiently.

In any knowledge-based system, there should be a modeling step between the acquisition of domain knowledge and the implementation of the knowledge model, in order to achieve a conceptual model that is easier to understand, among other things [6]. We will in the following introduce OWL and RDF and describe the new conceptual model of MedView. In the next section, we will present the implementation of this model so far.

3.1 OWL and RDF

As mentioned in the introduction, the Web Ontology Language (OWL) and Resource Description Framework (RDF) are recent recommendations from the World Wide Web Consortium. Essentially, RDF is a data-model. An subject-attribute-object triple, called a statement, is its basic building block. The statements are graphs, where subjects and objects are nodes connected by attributes as the arcs. An example of a triple is:

```
Peanuts  rdf:type  somwebOntology#Allergy
```

Here *Peanuts* are described as being of type *Allergy*. The *rdf* in *rdf:type* should be interpreted as a namespace which would have been defined earlier, and this is where we would find an ontology defining a meaning of type. There are several ways to represent the abstract data model more concretely, and RDF is most commonly described in an XML (eXtensible Markup Language) syntax. The example above would be represented as follows in RDF/XML:

```
<rdf:Description rdf:ID="Peanuts">
    <rdf:type>"somwebOntology#Allergy"</rdf:type>
</rdf:Description>
```

Fundamental concepts of RDF are resources, properties and statements. Resources are the things we want to talk about, such as diagnoses, medications, and allergies. Every resource has a URI (Universal Resource Identifier), which can be a URL (Unified Resource Locator) or some other kind of unique identifier. Properties are a special kind of resources, describing relations between resources. In RDF, properties are identified by URIs. The notion of using URIs to identify things and relations is central in giving a global naming scheme [1].

OWL is designed to be the standardized and broadly accepted language for describing ontologies, allowing users to write explicit, formal conceptualizations of domain models. OWL builds on RDF and uses RDF's XML-based syntax. There are three increasingly expressive sublanguages of OWL:

- OWL Lite supports those who primarily need a classification hierarchy and simple constraint features
- OWL DL (Description Logics) supports those who want the maximum expressiveness without losing computational completeness
- OWL Full is for users who want maximum expressiveness with no computational guarantees

Some language elements of OWL are classes, properties, and property restrictions. Among others, we can state that classes are equivalent, what a property's domain, range, and inverse property are, and make restrictions on what values a certain property can take [29].

For creating an ontology, a text or graphical ontology editor can be used. Such tools can also be used for creating instances of an ontology. Of the graphical editors, Protégé [19, 20] is one of the more popular. Protégé is a knowledge-base program developed at Stanford medical. It is open-source and allows people to contribute plugins. Among these plugins are tools for visualizing and reasoning, as well as an OWL-plugin [12]. The ontology created can then be viewed with a browser, such as IsaViz [21]. Further, there several query languages for RDF [17], which can be used for finding instances fulfilling criteria of interest. There also exist several frameworks for writing programs to interact with OWL and RDF content, such as Jena [24]. Thus far, we have experimented with using Protégé and Jena. As OWL and RDF are becoming widely used, we propose that these be used in the implementation of the model. By applying international standards, sharing knowledge and connecting ontologies should be achieved more easily. Further, the array of existing development tools will aid us in this implementation, and ensure that other ontology creators will have an equivalent ease of development.

3.2 Designing the New Model

Our approach is based on the observation that success factors for bringing knowledge-based systems into daily clinical use includes the careful planning of the introduction of the system in the clinical setting, in cooperation with end-users [31]. Right from the start, MedView was set up as a close collaboration between domain experts and computer scientists, in order to ensure that the developed systems were directly integrated into the clinical 'workflow' [4, 15, 30, 31]. Another key success factor is that the resulting system uses 'real' clinical data [32], a factor which is definitely met in MedView.

The different 'views' supported by the model should be based on the extent of which they support real-life clinical tasks [32]. Since our experience with MedView has shown that the basic clinical task in oral medicine is the examination of patients, the concept of a 'clinical examination' is the natural starting point for an ontology for oral medicine.

Along the lines with previous research, we have started our work on constructing an ontology for oral medicine by focusing on a specific perspective of knowledge in oral medicine [4, 18], namely the concept of an 'examination'. In this respect, our methodology for constructing an ontology for oral medicine can best be described as a 'middle-out' approach [3, 6], based on real-life scenarios, e.g., an examination situation, the collaborative and distributed construction of clinical protocols, the generation of summaries of patient data for use in referrals of patients or the iterative process of formulating and validating hypotheses over the collected clinical data.

4 Implementation

In the initial remodeling, the goal is to be able to represent that which can be represented in MedView today. The ontology is referred to as the SOMWeb ontology, as it is part of the first phase of the SOMWeb project. After a short description of the current contents of this ontology, we will describe how OWL and RDF can be used to address some of the problems raised in Section 2.

4.1 Overview

In the SOMWeb ontology, OWL classes are used to represent terms, such as *Allergy*, and parts of an examination template, such as *GeneralAnamnesis*. Value lists corresponding to terms are represented as RDF instances of these OWL classes (see Figure 1). OWL object properties are associated with these terms, where the range gives the term the property is related to, and domain gives the part of the examination it is connected with. An example of this is shown in Figure 2 for the property *hasAllergy*. Each examination record is represented as an RDF document, part of which can be seen in Figure 3. In all the examples, the identifiers such as *Allergy* and *Peanuts* should be seen as part of complete URIs. *somwebOntology#Allergy* and *rdf:type* mean that we are referring to the resources *Allergy* and *type* as defined in the ontologies of *rdf* and *somwebOntology*.

```
<rdf:Description rdf:ID="Peanuts">
     <rdf:type>"somwebOntology#Allergy"</rdf:type>
     <rdfs:label xml:lang="en">Peanuts</rdfs:label>
     <rdfs:label xml:lang="sv">Jordnotter</rdfs:label>
     <dc:creator>"somwebPeople#Jontell"</dc:creator>
</rdf:Description>
```

Figure 1: RDF/XML for part of a value list for the term/class Allergy.

```
<owl:ObjectProperty rdf:ID="hasAllergy">
     <rdfs:comment xml:lang="en">What allergies do you have?</rdfs:comment>
     <rdfs:comment xml:lang="sv">Ar du allergisk mot nagot?</rdfs:comment>
     <rdfs:range rdf:resource="somwebOntology#Allergy"/>
     <rdfs:domain rdf:resource="somwebOntology#GeneralAnamnesis"/>
</owl:ObjectProperty>
```

Figure 2: OWL example for the object property hasAllergy.

```
<rdf:Description rdf:ID="GeneralAnamnesis_32_041115">
     <rdf:type>"somwebOntology#GeneralAnamnesis"</rdf:type>
     <somweb:hasAllergy rdf:resource="somwebOntology#Apricot"/>
     <somweb:takesDrug rdf:resource="somwebOntology#Losec"/>
     <somweb:goodHealth rdf:datatype="http://www.w3.org/2001/XMLSchema#boolean"
     >true</somweb:goodHealth>
</somweb:GeneralAnamnesis>
```

Figure 3: RDF/XML for part of an examination record.

4.2 Using External Sources

An important contribution of using ontologies is the ability to reuse existing ontologies. Importing other OWL ontologies is one of its central features, and much of the current development of ontologies is in OWL, ensuring that there will be ontologies to import. Currently, we are looking into importing ontologies for countries and units. Other areas where ontology reuse holds great potential are medications and diagnoses. Dublin Core [10] is widely used ontology, defining a standard for information resource description, the possible use of which will be discussed in Section 4.5.

4.3 Examination Templates, Lists of Approved Values, and Groups of Related Terms

A central part of MedView are the examination templates. In the first implementation of the SOMWeb ontology, each part of the examination is represented as an OWL class, to which properties are connected by their domain (see Figure 2). When creating an examination record, it is checked which properties have a given examination part as their domain. An example of a part of an examination record represented in RDF is shown in Figure 3.

In the current version of MedView, all the value lists are in a text-file, listed under their associated term. As mentioned above, these terms will now be OWL classes, and the value lists will be RDF instances of these classes. For example, there is a class *Allergy* with instances such as *Peanuts*, *Oranges*, and *BeeStings*.

Grouping related terms are an important way for the end-user to categorize different aspects of the data, such as diagnoses, smoking habits, and medications. This can be accomplished using OWL/RDF at several different levels. One way is to rigorously define subclasses, for example by defining *FoodAllergy* and *PollenAllergy* as subclasses of *Allergy*. We will also want to create descriptions for different categories, e.g. for smoking habits, where an OWL restriction can be created for what it means to be a heavy smoker.

4.4 Typing of Elements

Several methods in OWL and RDF can be used in aiding the typing of terms and values. Apart from object properties, which relate objects to other objects, OWL also supports data type properties, which relate objects to data type values. While the previously described *hasAllergy* is an object property, a good candidate for a data type property is *hasGoodHealth*, taking a value of true or false. For OWL data type properties, XML Schema data types can be used, so *hasGoodHealth* would have range XMLSchema#boolean, as defined by W3C. On the issue of units, OWL and RDF provide good support. For example, smoking habits could be represented as follows:

```
<hasSmokingHabit rdf:parseType = "Resource">
    <rdf:value> 10 </rdf:value>
    <smokeCategory rdf:resource = "filterCigarettes" />
    <timePeriod rdf:resource = "timeOnt#day" />
</hasSmokingHabit>
```

This gives a useful separation of quantities and units, using a fictive external ontology of time units. There are several initiatives for ontologies for units (e.g. [25]), which the SOMWeb ontology could use.

4.5 Meta-data

For managing the multitude of terms, values, and protocols, using metadata is necessary. We want to represent are who created a term, value, or template, their reason for doing so, and when. It is thought that the terms in the above introduced Dublin Core (DC), will be useful. DC includes elements such as creator, subject, and description. In Figure 1, the creator element is used to describe who entered the value described.

4.6 Localization of Data

We are specially interested in support for different languages. In Figures 1 and 2, the reader will notice the presence of both English and Swedish. In the case of a value such as *Peanuts*, different labels can be specified for different languages using xml:lang. The same goes questions associated with the properties such as *hasAllergy*. An application using this can be instructed to chose either the Swedish or English labels and questions.

5 Discussion and Future Work

We have presented the problems faced in representing knowledge in oral medicine, and shown how these can be alleviated by using OWL and RDF. Most work so far has been in representing lists of approved values, examination templates and records, and language support. In the area of representing data types of elements and meta-data, approximately a third of the work has been done. We have also begun identifying ontologies of interest for reuse. Although importing ontologies may ease both knowledge sharing and the process of ontology construction, it is not a trivial process. Pinto and Martins [22] propose a methodology for choosing the suitable candidate ontologies and tackling compatibility and completeness problems.

Three problems listed in Section 2 have not been thoroughly explored: different views of the data, relations and interactions between different parts of the model, and representing information in photos. In catering for perspectives other than examination, such as patient and time, we believe the graph structure of OWL and RDF is important This structure will make it possible to "grab" different parts of the graph. In some cases, inverse properties will have to be defined to aid in moving in both directions on the graph. For example, the *hasAllergy* property may need an inverse property *allergyOf*. For representing relations and interactions between different parts of the model, it is believed that an OWL rule language, such as the one described in [9], could be very useful. There exist several initiatives for annotating photos using OWL/RDF, e.g. MindSwap's Photo-Stuff [13], which may be used to represent information contained in photos taken during examinations.

Another aspect we wish to explore is user modelling, needed to support users of different familiarity with computers, different roles, and working in different departments. This can be a way to cope with the problem of getting the end user to feel comfortable with the application, meaning a greater chance of data being entered. Though this area is an uninvestigated part of the project, work by Razmerita et al. [23] discussing ontologies for user modelling is interesting.

During this initial implementation, it has become apparent that OWL and RDF are very applicable in solving the problems described in Section 2. This work will be continued with further modelling of examinations and related concepts. Ontologies to be imported will be identified and hopefully successfully incorporated. As a case study, an existing MedView application will be adapted to using OWL and RDF, though the OWL and RDF will not be apparent to the end user. The new application will be tested as part of a more general evaluation of the OWL/RDF approach. How will the models and tools compare to others in the ontology domain?

References

[1] Grigoris Antoniou and Frank Van Harmelen. *A Semantic Web Primer*. MIT Press, 2004.

[2] Krzysztof J. Cios and G. William Moore. Uniqueness of medical data mining. *Artif. Intell. Med.*, 26(1–2):1–24, 2002.

[3] Oscar Corcho, Mariano Fernández-López, and Asunción Gómez-Pérez. Methodologies, tools and languages for building ontologies. where is their meeting point? *Data Knowl. Eng.*, 46(1):41–64, 2002.

[4] Theresa Edgington, Choi Beomjin, Katherine Henson, T. S. Raghu, and Ajay Vinze. Adopting ontology to facilitate knowledge sharing. *Commun. ACM*, 47(11):85–90, 2004.

[5] Göran Falkman and Olof Torgersson. Knowledge acquisition and modeling in clinical information systems: A case study. In A. Gomez-Perez and V. R. Benjamins, editors, *Knowledge Engineering and Knowledge Management. Proc. EKAW '02*, volume 2473 of *Lect. Notes Artif. Intell.*, pages 96–101, 2002.

[6] Mariano Fernández-López and Asunción Gómez-Pérez. Overview and analysis of methodologies for building ontologies. *Knowl. Engin. Rev.*, 17(2):129–156, 2002.

[7] T. R. Gruber. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199–220, 1993.

[8] J. R. Hajdukiewicz, K. J. Vicente, D. J. Doyle, P. Milgram, and C. M. Burns. Modeling a medical environment: An ontology for integrated medical informatics design. *Int. J. Med. Inform.*, 62(1):79–99, 2001.

[9] Ian Horrocks and Peter F. Patel-Schneider. A proposal for an owl rules language. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 723–731. ACM Press, 2004.

[10] Dublin Core Metadata Initiative. DCMI metadata terms. Retrieved January 2, 2005. Available at: `http://dublincore.org`.

[11] Mats Jontell, Ulf Mattsson, and Olof Torgersson. MedView: An instrument for clinical research and education in oral medicine. *Oral Surg. Oral Med. Oral Pathol. Oral Radiol. Endod.*, 99:55–63, 2005.

[12] Holger Knublauch, Ray W. Fergerson, Natalaya F. Noy, and Mark A. Musen. The Protégé OWL Plugin: An open development environment for semantic web applications. In McIlraith et al., editor, *Proc.* ISWC *'04*, volume 3298 of *Lect. Notes Comput. Sci.*, pages 229–243, 2004.

[13] Grecia C. Lapizo-Encinas. Photostuff - an image annotation tool. Created 2004. Retrieved December 28, 2004. Available at: `http://www.mindswap.org/~glapizco/PS.shtml`.

[14] F. Manola and E. Miller (Eds.). RDF Primer: W3C Recommendation 10 February 2004. Created 2004. Retrieved December 28, 2004. Available at: `http://www.w3.org/TR/2004/REC-rdf-primer-20040210/`.

[15] Eneida A. Mendonça. Clinical decision support systems: Perspectives in dentistry. *J. Dent. Edu.*, 68:589–597, 2004.

[16] Eneida A. Mendonça, James J. Cimino, S. B. Johnson, and Y. H. Seol. Accessing heterogeneous sources of evidence to answer clinical questions. *J. Biomed. Inform.*, 34(2):85–98, 2001.

[17] Libby Miller, Andy Seaborne, and Alberto Reggiori. Three implementations of SquishQL, a simple RDF query language. In ISWC *'02: Proceedings of the First International Semantic Web Conference on The Semantic Web*, pages 423–435. Springer-Verlag, 2002.

[18] Natalaya F. Noy and Mark A. Musen. Specifying ontology views by traversal. In McIlraith et al., editor, *Proc.* ISWC *'04*, volume 3298 of *Lect. Notes Comput. Sci.*, pages 713–725, 2004.

[19] Natalya F. Noy and Deborah L. McGuinness. Ontology development 101: A guide to creating your first ontology. *Stanford Knowledge Systems Laboratory Technical Report*, 2001.

[20] Natalya F. Noy, Michael Sintek, Stefan Decker, Monica Crubezy, Ray W. Fergerson, and Mark A. Musen. Creating semantic web contents with Protege-2000. *IEEE Intelligent Systems*, 2(16):60–71, 2001.

[21] Emmanuel Pietriga. IsaViz: A visual authoring tool for RDF. Retrieved December 28, 2004. Available at: `http://www.w3.org/2001/11/IsaViz/`.

[22] Helena Sofia Pinto and Joao P. Martins. Ontologies: How can they be built? *Knowledge and Information Systems*, 6(4):441–464, 2004.

[23] Liana Razmerita, Albert A. Angehrn, and Alexander Maedche. Ontology-based user modeling for knowledge management systems. In Peter Brusilovsky, Albert T. Corbett, and Fiorella de Rosis, editors, *User Modeling*, volume 2702 of *Lecture Notes in Computer Science*, pages 213–217. Springer, 2003.

[24] Hewlett-Packard Labs Semantic Web Research. Jena. Retrieved December 28, 2004. Available at: `http://www.hpl.hp.com/semweb/`.

[25] Hajo Rijgerberg and Jan Top. UnitDim: An ontology of physical units and quantities. Created 2004. Retrieved January 2, 2005. Available at: `http://www.atoapps.nl/ProjectSite2/index.asp`.

[26] F. W. Rook and M. L. Donnell. Human cognition and the expert system interface: Mental models and inference explanations. *IEEE T. Syst. Man Cyb.*, 23(6):1649–1661, 1993.

[27] William M. C. Rosenberg and Anna Donald. Evidence based medicine: An approach to clinical problem solving. *Brit. Med. J.*, 310(6987):1122–1126, 1995.

[28] David L. Sackett, William M. C. Rosenberg, J. A. Muir Gray, R. Brian Haynes, and W. Scott Richardson. Evidence based medicine: What it is and what it isn't. *Brit. Med. J.*, 312(7023):71–72, 1996.

[29] M. Smith, C. Welty, and D. McGuinness. OWL Web Ontology Language Guide: W3C Recommendation 10 February 2004. Created 2004. Retrieved December 28, 2004. Available at: `http://www.w3.org/TR/2004/REC-owl-guide-20040210/`.

[30] Susan E. Sutherland. The building blocks of evidence-based dentistry. *J. Can. Dent. Assoc.*, 66:241–244, 2000.

[31] T. Wetter. Lessons learnt from bringing knowledge-based systems into routine use. *Artif. Intell. Med.*, 24(3):195–203, 2002.

[32] Qing Zeng and James. J. Cimino. A knowledge-based, concept-oriented view generation system for clinical data. *J. Biomed. Inform.*, 34(2):112–128, 2001.