Ontological Modelling of Surgical Knowledge

Raj Mudunuri, Oliver Burgert, Thomas Neumuth

Innovation Center Computer Assisted Surgery
Semmelweißstrasse 14
04103 Leipzig
Raj.Mudunuri@iccas.de
Oliver.Burgert@iccas.de
Thomas.Neumuth@iccas.de

Abstract: Computer Assisted Surgery seeks data from various sources related to surgical knowledge, which is mostly represented in the form of discrete databases. However, databases serve as data storage mechanism rather than knowledge representation system. As CAS is becoming more common for various types of surgical interventions, there is a need for representation, storage and processing of surgical knowledge in a more structured manner. The traditional way of using databases to store the data needs to be augmented with a conceptual representation mechanism, which can serve as a bonding layer among the various sources of surgical knowledge. Ontological modelling provides the means to represent conceptual knowledge using expressive formal logics. The scope of this paper is to present how CAS knowledge can be represented using ontological modelling, and make use if it in the real time applications in the Operating Room.

1 Introduction

The domain of Computer Assisted Surgery (CAS) is an inter-disciplinary field. Apart from the pure medical science, it also involves various other closely connected areas such as *engineering* of surgical devices, *modelling* of surgical procedures, *visualisation* of medical imaging techniques, *assessment* of various methods adopted by different surgeons, and the *representation* of this varied knowledge in a computable manner etc. CAS is fast evolving into a domain of its own with each new technological advancement in the area of personal health care. Along with the vast and various sources of existing surgical knowledge, e.g. anatomic atlases, medical databases with patient records, documented surgical procedures etc., the information keeps growing with every new patient, new problem and new solution. Keeping track of all this information, and using it for specific applications is a highly demanding task. Moreover, it is apparent that the constant evolving nature of various information sources in general has compelled the evolution of knowledge representation mechanisms as well. Therefore, the problem of using heterogeneous data in CAS applications must be dealt with using the latest knowledge representation systems.

The interesting question is, what kind of information do the CAS applications typically need, and how should that information be represented and processed. The scope of this research is to find out how CAS applications can make use of the latest Knowledge Representation (KR) techniques, such as ontological modelling, to express relevant information in the domain of surgery.

2 Knowledge Representation in Computer Assisted Surgery

Computer Assisted Surgery represents a surgical methodology where the use of computer based technology plays an important role in all the phases, i.e. pre-operative planning, intra-operative procedures and post-operative analysis. One important aspect of CAS lies in the development of an accurate model of the patient based on diagnostic data. However, representation of this patient model depends on the scope and needs of specific applications. The pre-operative planning of a surgery requires a collection of various medical images, while the intra-operative phase involves using the anatomical and pathological data, whereas the post-operative analysis may seek information not only from the above two phases but also from other sources such as the past diagnostic records of the patient. Though a surgery of a patient might involve these different phases at different periods of time, the relevant data of that patient is scattered at different sources. The data is also represented in different forms, which makes the task of software/hardware tools more complex to understand the meaning of this data and process it accordingly. Therefore, in order to have semantic inter-operability among various data sources used by different CAS applications, we must adopt a methodology that provides the means to achieve it.

2.1 Ontologies as a Knowledge Representation mechanism

An ontology is defined as an explicit formal specification of a shared conceptualization [Gr93]. Knowledge from a domain of interest can be represented using ontologies, by logically defining the concepts in a semantically interoperable manner. The knowledge is split into a group of ontologies that describe different aspects of the same domain. Ontological representation of knowledge is not only easy to comprehend for the humans but also possible for the machines to process it. So far, no other field had embraced ontologies as much as the medical field did [SGB00]. An attempt to describe a set of words as surgical deeds, which are used in surgical procedures, in the form of ontological analysis based on the Comité Européen de Normalisation (CEN) standards is made by Rossi Mori et al [RGS97]. Considerable amount of work is done in this area as part of the well known Common Reference Model from GALEN (Generalised Architecture for Languages, Encyclopaedias, and Nomenclatures in Medicine) [RRS01]. While the work of Rossi Mori et al is concentrated at a general level of surgical domain, the work of Carlsson et al is aimed at thoracic surgery [CLR00]. However, Rossi Mori's work was confined to some basic terminology, whereas GALEN managed to find wide acceptance in the KR community, and is still regarded as one of the important sources of medical nomenclatures.

More recently, ontological models of Brain anatomy was used to annotate brain MRI images [MGG06]. These works have provided promising results, to believe that ontological modelling of some aspects of the surgical domain can benefit the CAS community.

2.2 Ontologies for CAS Applications

Even though databases and PACS system are very much useful for storing large collections of medical data, they are not suitable for semantic interoperability. The basic difference between a database model and an ontological model to represent a domain of interest lies in the fundamental aspect of *what is defined* and *what is not defined* in the provided information [AD98, GCC04]. In the traditional database model every data item that is defined is assumed to be positive information and undefined information is assumed to be negative. In other words, a database is supposed to be closed (Closed World Assumption) and complete with the information that it contains, which means, if a query is posed to a database to know whether the instrument *shaver* can be used for the act of *cutting* then the answer would be YES if it is defined as such, or it would be NO if it is not defined at all. This is because the database model implicitly assumes that if any information is not defined in the domain then it means that the undefined information is anathema for the current application and has to be answered negatively.

On the other hand, ontological model takes the approach of Open World Assumption, which means, the information provided to the model is assumed to be partial and open for any new additional piece of information even though the new information might contradict the existing information. The reasoning engines that run behind the ontologies make sure that the information on the whole is logically consistent with every definition contained in the domain [HM03]. So, if a query is posed to an ontological model to know whether the instrument *shaver* can be used for the act of *cutting* then the answer would be YES, if it is defined as such, or there would be no definitive answer if it is not defined at all. A definite NO will be answered only when it is explicitly defined in the ontology that a *shaver* cannot be used for the act of *cutting*. The advantage of this approach lies in modelling the world in its real sense, where as long as things are not made explicit by formally defining them they are open for multiple interpretations. So an ontological model provides the flexibility to define concepts with various interpretations, and at the same time it makes sure that these various interpretations do not contradict with each other at any point within the domain. Therefore, an ontological model provides the flexibility to define concepts with various interpretations, and at the same time it restricts those interpretations from contradicting with each other.

3 Methods

Ontological modelling of surgical data involves description of relevant concepts and roles (attributes) that capture the domain of interest. The size of a domain depends upon the amount of information that is needed for various applications that use this domain knowledge. Therefore, surgical information is segregated into various ontologies at different hierarchical levels. Ontologies at the higher level of granularity contain concepts that are defined in broader logical expressivity, and ontologies at the lower level of granularity contain concepts that are applicable only for a particular surgical discipline for example. The advantage of this top-down and bottom-up approach is that all the concepts that are defined in the entire framework of ontologies follow a logical pattern where concepts of lower level ontologies are subsumed by the concepts of higher level ontologies, and vice versa. The logical subsumption relations between concepts of the same ontology, and between different ontologies within the domain of interest, is an integral part of ontological modelling. Figure 1 shows some of the concepts from four different ontologies, at different levels of abstraction, which are developed within the framework of Surgical Ontologies for Computer Assisted Surgery (SOCAS). Every ontology has its own namespace, and each concept is prefixed with this namespace. The concepts are defined in such a way that they are logically expressive enough to provide enough information for the software applications that seek with help of different query mechanisms.

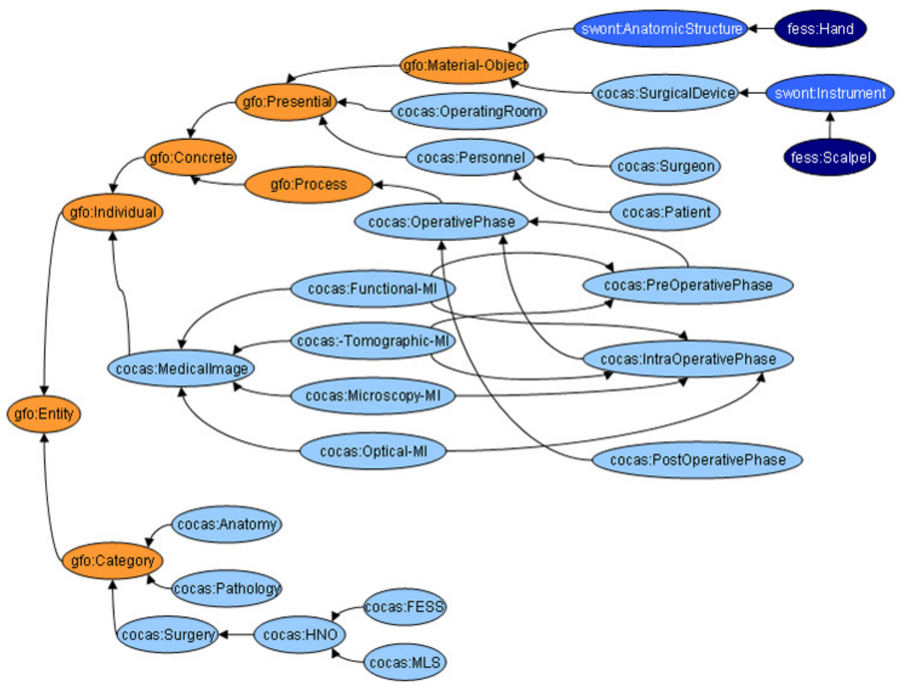

Figure 1: Excerpt from SOCAS ontological concepts

3.1 Linking the ontological concepts

The four ontologies involved in Figure 1, from left to right, are at different abstraction levels starting from general to specific. The most specific ontology here is FESSOnt, an ontology with concepts that are related to Functional Endoscopic Sinus Surgery. The next is Surgical Workflow ontology (SWOnt) that has concepts related to Surgical Workflows [NMJ07, NJS09]. Concepts such as *Hand* and *Scalpel* from FESSOnt are subsumed by more general concepts such as *AnatomicStructure* and *Instrument* from SWOnt. The concept *Instrument* is further subsumed by another super concept *SurgicalDevice* from the Core Ontology for CAS (COCAS). COCAS contains close to 70 concepts, which describe various aspects of surgical knowledge. For example, different kinds of Medical Imaging (MI) techniques, different operation phases that use different MI techniques, different types of surgeries that are classified into certain groups etc. These concepts are not just linked with simple subsumption relations, but also classified according to different roles between them. For example, the relation "*fess:Scalpel is-a swont:Instrument*" can be extended to "*swont:Instrument involve-atleast-2 fess:Scalpel* iff *fess:Nose is-a swont:AnatomicStructure*". These logical definitions can be extended further to serve the needs of the application, as long as the consistency of the overall ontology is maintained. The fourth ontology in this framework is a top level General Formal Ontology (GFO), developed by the Onto-Med research group at the Institute for Medical Informatics, Statistics, and Epidemiology [HHB06]. GFO contains concepts at a more general level that can be applicable for real world scenarios. Though they do not play a significant role in the Operating Room (OR) at the moment, they can still provide interesting insights at the theoretical level.

The below concept definitions from different SOCAS ontologies illustrate how such concepts are defined using axioms that maintain logical consistency.

gfo:Presential ⊆ gfo:Concrete ^ (gfo:exists_at some gfo:Time_boundary) (i)

gfo:Process ⊆ gfo: Concrete ^ (gfo:projects_to some gfo:Temporal_region) (ii)

gfo:Material-Object ⊆ gfo:Presential ^ (gfo:framed_by some gfo:Topoid) (iii)

The top level concepts Presential, Process and Material-Object are from GFO. A Presential is defined as a Concrete, and that exists at some point of time. A process is also a Concrete, but that has a Temporal region sharing two boundaries. A material object is defined in terms of a Presential that is framed by some topological space.

cocas:OperatigRoom ⊆ cocas:hasCoordinates some cocas:Coordinates (iv)

cocas:Personnel ⊆ cocas:hasData some cocas:Data (v)

cocas:Patient ⊆ cocas:Personnel ^ cocas:hasPathologicalData some cocas:PatientRecord ^ cocas:hasSocialData some cocas:PatientRecord (vi) cocas:PreOperativePhase $\subseteq$ cocas:OperativePhase ^ (cocas:hasFollowUpProcedure only (cocas:IntraOperativePhase and (cocas:hasFollowUpProcedure only cocas:PostOperativePhase))) ^ (cocas:hasFollowUpProcedure some (cocas:IntraOperativePhase and (cocas:hasFollowUpProcedure some cocas:PostOperativePhase))) (vii)

cocas:SurgicalProcedure $\subseteq$ cocas:hasParticipant some (cocas:Surgeon and cocas:Patient and cocas:SurgicalDevice) (viii)

The concepts such as OperatingRoom, Personnel, Patient, PreOperativePhase and SurgicalProcedure are defined as part of COCAS ontology, which is hierarchically subsumed by the GFO. A PreOperativePhase of a surgery is defined as an OperativePhase that has a follow-up IntraOperativePhase, which in turn has a follow-up PostOperativePhase. To close the definition in its totality, it is quantified by both universal and existential quantifications. A SurgicalProcedure involves participants such as Surgeon, Patient and SurgicalDevice.

fess:Scalpel $\subseteq$ (swont:allowedAction some fess:Cut) ^ (swont:allowedAction only fess:Cut) (ix)

fess:Scalpel $\subseteq$ swont:allowedAnatomicStructure only (fess:CellEthmoidales or fess:CavitasNasi or fess:ConchaNasalis) (x)

The discipline specific ontology, FESSOnt, has concept definitions that are more concerned with this specific domain. For example, a Scalpel is defined in terms of the actions and treated structures it is allowed to use. At the end, these logical connections from top to bottom, and vice versa, which include simple concept subsumptions within the ontology and complex axiomatic relations across the ontologies, form into a classified knowledge base model that can answer questions such as: (a) What are the possible instruments that can be used for a particular activity that involves certain anatomical structures? (b) Which Personnel with certain Social and Pathological Data are involved in certain OperativePhases? Though it would seem that these kind of questions could be answered with a well modelled database, the model itself does make a difference. The ontological model can easily be updated (adding facts, modifying concepts, deleting concepts that become inconsistent due to the new knowledge) to maintain the knowledge base, whereas a database is not meant for tweaking with the model but only to dump the values in the predefined structured tables.

3.2 Ontology devising tools

All the ontologies are devised in Web Ontology Language (OWL) [AH03], using Protégé [Pro09] as the ontology editor. The reasoning engine (Racer [HM03]) was used for classifying the ontology and consistency checking. Logical reasoners process all the concepts in the ontology to check out for unexpected inconsistencies, and also extract new knowledge when possible.

When some concepts are individually defined with some references to other concepts, then the reasoner shall process the entire ontology at the end to find out whether the references might have produced any new facts. If yes, then this new piece of information is saved in the form of inferred ontology. This ontological model is then used as a knowledge base to supply required information to software applications in the form of query answering. Query mechanisms such as SPARQL [PAG06] and SQWRL [CSN 08] are used for information retrieval.

4 Results

Surgical knowledge pertaining to various domains of interest is conceptualised in terms of various ontological models. Core ontology for CAS (COCAS) consists of the concept definitions such as *SurgicalProcedure*, *SurgicalImplant*, *SurgicalDevice* etc. Figure 2 shows how, for example, *SurgicalProcedure* is defined as a procedure that has a surgical device, patient and surgeon as participants, and an operative phase. Similarly, every concept is defined in terms of logical axioms that are sound and complete.

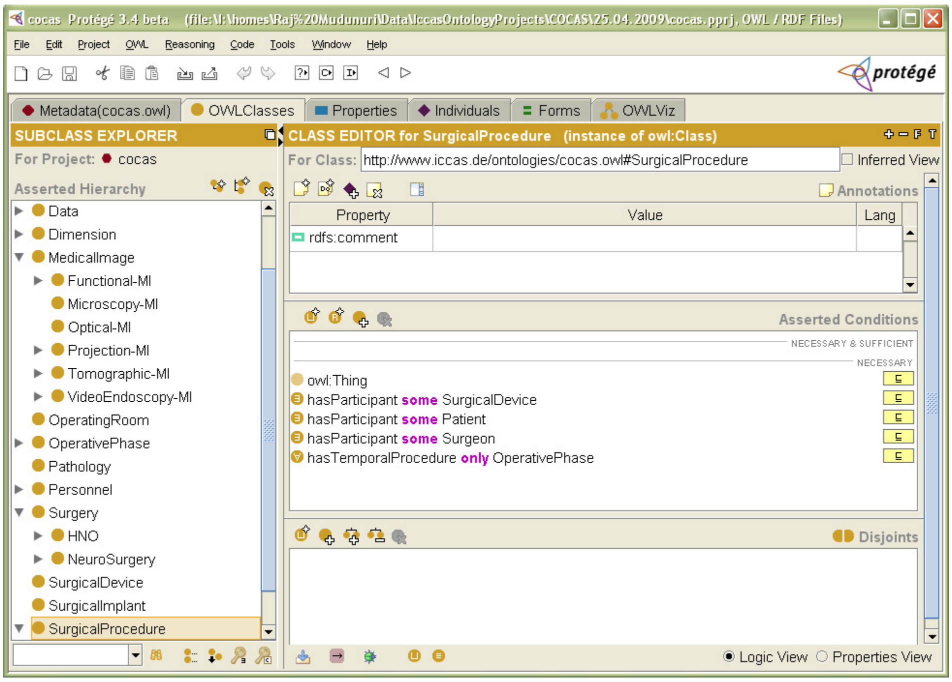

Figure 2: COCAS concepts defined with OWL semantics, using an ontology editor

Discipline specific ontology for FESS intervention (FESSOnt) is devised to define the anatomical structures, surgical instruments and activities that are involved in that surgical discipline. The individual concepts and sub concepts of these categories are defined according to the inter-linkage among the concepts. For example, the action *Cut* is defined as a *SurgicalAction* that can only use the instruments of the type *Scalpel* or *Scissors* or *Knife* or *SharpSpoon*. Similarly anatomic structures are restricted to the use of certain actions.

An ontology with surgical workflow concepts (SWOnt) acts as a bridge between the core ontology and discipline specific ontology, to propagate the information between them. These ontologies are used as background knowledge bases to provide information for applications such as workflow recording in the OR. The ontologies are situated in a repository and are communicated through java classes. Before passing on the required knowledge to the workflow editor, the ontologies are classified by the reasoning engine to draw logical inferences and thereby making sure that the data received by the editor is consistent. This reasoning process (classification) of the ontological model is done only once as long as the model remains unchanged. So whenever there is a need to define more concepts in the model, and thereby changing the ontological structure, then the model is reasoned again by the reasoning engine to generate a new inferred model that will be used by the application from thereupon.

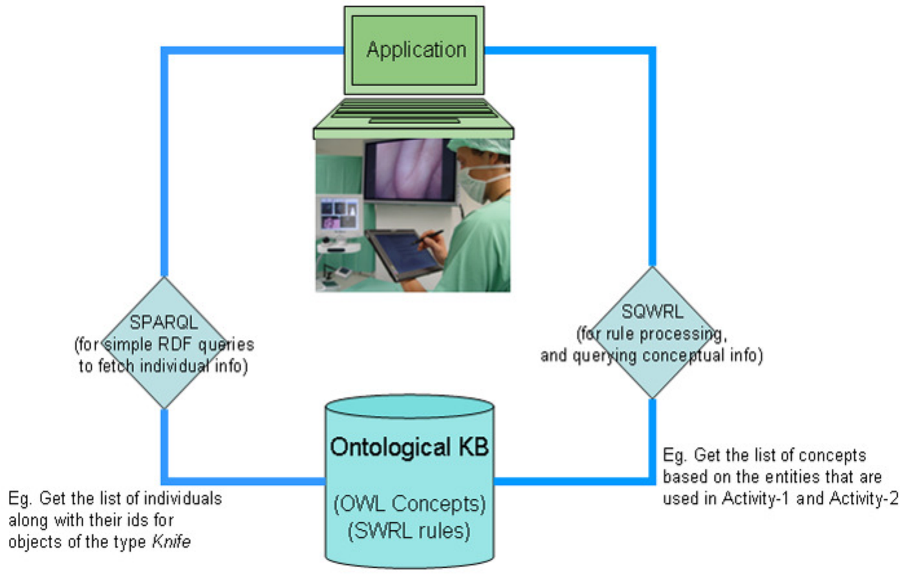

Figure 3: Information retrieval from the Ontological knowledge base for a CAS application

Figure 3 depicts a simple illustration of how ontological knowledge is used with different querying mechanisms in a CAS application. The Ontological KB is a repository of ontological concepts along with the Semantic Web Rule Engine (SWRL) rules.

SWRL rules are used when conditional statements are needed to enhance the knowledge model. SWRL rules are different from concepts definitions, in the sense that the later are defined purely on concepts, whereas the former act on the individuals of the concepts. For example, the below SWRL rule declares that if a patient has a temperature of more than 104 (F) then his condition should be termed as critical.

Patient (?x) ^ hasTemperature (?x, ?y) ^ swrl:greaterThan(?y, 104) → PatientCondition (?x, 'critical')

With the below SQWRL query, one can ask for patient records of all the patients whose temperature is greater than 104.

Patient (?p) ^ hasPatientRecord(?p, ?r) ^ hasTemperature (?p, ?t) ^ swrl:greaterThan(?t, 104) → sqwrl:select (?p, ?r)

While the SWRL rules and SQWRL queries act on owl semantics, SPARQL acts on the RDF based syntax. The following SPARQL query asks for all the devices that are used in different intervention phases.

```
SELECT ?a ?b WHERE {
        ?a rdfs:subClassOf OI:ElectricDevice .
        ?a rdfs:subClassOf ?y .
        ?y rdf:type owl:Restriction .
        ?y owl:onProperty cocas:hasInterventionPhase .
        ?y owl:someValuesFrom ?b .
        ?b rdfs:subClassOf cocas:OperativePhase
        }
```

As every instrument is predefined in the ontology with its associated intervention phase, such as: fess:Endoscope $\sqsubseteq$ cocas:hasInterventionPhase some cocas:IntraOperativePhase; the query then returns the list of all the devices along with the names of their associated intervention phases.

5 Conclusion

CAS applications use multiple sources of information. Most of this information is not well defined for semantic interoperability. Surgical knowledge needs to be defined in such a manner where many surgical applications can make use of it. Ontological modelling is one of the promising knowledge representation mechanisms. Surgical Ontologies for Computer Assisted Surgery (SOCAS) is an ontological framework that contains various ontologies at different levels of abstraction. These ontologies contain information about the core concepts of CAS, and surgical disciplines such as Functional Endoscopic Sinus Surgery. SOCAS ontologies are used as background knowledge bases for applications such as a workflow editor that records the flow of the events happening in the OR.

6 Acknowledgements

We thank the OntoMed group, especially Prof.Heinrich Herre and Frank Loebe, for their contribution through successful discussions and feedback for our work.

The Innovation Center Computer Assisted Surgery (ICCAS) at the Faculty of Medicine, University of Leipzig, is funded by the German Federal Ministry for Education and Research (BMBF) and the Saxon Ministry of Science and the Fine Arts (SMWK) in the scope of the initiative ”Unternehmen Region” with the grant numbers 03 ZIK 031 and 03 ZIK 032.

References

[AD98] Abiteboul S and Duschka O.M. Complexity of answering queries using materialized views, In *Proceedings of the ACM Symposium on Principles of Database Systems* (*PODS*), 254–263, 1998.

[AH03] Antoniou G and Van Harmelen F. Web Ontology Language: OWL. In Handbook on Ontologies in Information Systems, pages 67–92, 2003.

[CLR00] Carlsson M, Loefstroem L, Rogers J and Ahlfeldt H. Modelling and reclassification of surgical procedures - Experiences from the use of GALEN methods in the domain of thoracic surgery, Medical Informatics and the Internet in Medicine, 2000 Apr-Jun; 25(2):109-22.

[CSN 08] O'Connor M, Shankar R, Nyulas C, Tu S, Das A. Developing a Web-Based Application using OWL and SWRL. AAAI Spring Symposium, Stanford, CA, USA (2008).

[GCC04] Gali A, Chen C, Claypool K, and Uceda-Sosa R. “From ontology to relational databases,” in *Proc. Int. Workshop on Conceptual-Model Driven Web Information Integration and Mining*, Shanghai, China, 2004, pp. 278-289.

[Gr93] Gruber T R. Toward Principles for the Design of Ontologies Used for Knowledge Sharing, In Formal Ontology in Conceptual Analysis and Knowledge Representation, Kluwer Academic Publishers, 1993.

[HHB06] Herre H, Heller B, Burek P, Hoehndorf R, Loebe F and Michalek H. General Formal Ontology (GFO): A Foundational Ontology Integrating Objects and Processes. Part I: Basic Principles (Version 1.0), Onto-Med Report Nr. 8. Research Group Ontologies in Medicine (Onto-Med), University of Leipzig, 2006.

[HM03] Haarslev V and Moeller R. Racer: An OWL Reasoning Agent for the Semantic Web, In Proceedings of the International Workshop on Applications, Products and Services of Web-based Support Systems, in conjunction with the 2003 IEEE/WIC International Conference on Web Intelligence, Halifax, Canada, October 13, pages 91-95, 2003.

[MGG06] Mechouche A, Golbreich C, Gibaud B. Towards an hybrid system for annotating brain MRI images, Workshop on OWL Experiences and Directions, Athens, Georgia, United States (2006).

[NJS09] Neumuth T, Jannin P, Strauß G, Meixensberger J, Burgert O. Validation of Knowledge Acquisition for Surgical Process Models. Journal of the American Medical Informatics Association, 2009 Jan-Feb;16(1):72-80.

[NMJ07] Neumuth T, Mudunuri R, Jannin P, Meixensberger J, Burgert O. SWAN-Suite: The tool landscape for surgical workflow analysis, Computer Assisted Medical and Surgical Interventions (SURGETICA), (Troccaz, Jocelyne; Merloz, Phillip, ed.), vol., Sauramps Medical, Paris, 2007 p. 199-204.

[PAG06] Pérez J, Arenas M, Gutierrez C. Semantics and Complexity of SPARQL. ISWC.2006. LNCS, vol. 4273, pp. 30–43. Springer, Heidelberg (2006).
[Pro09] Protégé.stanford.edu [homepage on the Internet]. Stanford Center for Biomedical Informatics Research, © 2009. Available from: http://protege.stanford.edu/index.html
[RGS97] Rossi Mori A, Gangemi A, Steve G, Consorti F and Galeazzi E. An Ontological Analysis of Surgical Deeds, In Proc. of Artificial Intelligence in Europe, AIME97, 1997.
[RRS01] Rogers J, Roberts A, Solomon D, van der Haring E, Wroe C, Zanstra P, Rector A. GALEN Ten Years On: Tasks and Supporting Tools, In MedInfo 2001 Studies in Health Technology and Informatics, Haux R, Rogers R, Patel V (eds). IOS Press: Amsterdam; 256 - 260, 2001.
[SGB00] Stevens R, Goble C and Bechhofer S. Ontology-based Knowledge Representation for Bioinformatics. Briefings in Bioinformatics 1(4): 398-414 (2000).