Defaults, Context, and Knowledge: Alternatives for OWL-Indexed Knowledge Bases

A. Rector

Pacific Symposium on Biocomputing 9:226-237(2004)

DEFAULTS, CONTEXT, AND KNOWLEDGE: ALTERNATIVES FOR OWL-INDEXED KNOWLEDGE BASES

A. RECTOR
Bio-Health Informatics Group, Department of Computer Science,
University of Manchester, Manchester M13 9PL, UK
email rector@cs.man.ac.uk

The new Web Ontology Language (OWL) and its Description Logic compatible sublanguage (OWL-DL) explicitly exclude defaults and exceptions, as do all logic based formalisms for ontologies. However, many biomedical applications appear to require default reasoning, at least if they are to be engineered in a maintainable way. Default reasoning has always been one of the great strengths of Frame systems such as Protégé. Resolving this conflict requires analysis of the different uses for defaults and exceptions. In some cases, alternatives can be provided within the OWL framework; in others, it appears that hybrid reasoning about a knowledge base of contingent facts built around the core ontology is necessary. Trade-offs include both human factors and the scaling of computational performance. The analysis presented here is based on the *Open*GALEN experience with large scale ontologies using a formalism, GRAIL, which explicitly incorporates constructs for hybrid reasoning, numerous experiments with OWL, and initial work on combining OWL and Protégé.

1 Introduction

1.1 The problem: We want to rescue the "baby" from the "bathwater"

Until the mid 1980s, a key part of knowledge representation was capturing notions of defaults and exceptions. Minsky's original paper on frames [11] was based on the notion of "prototypes" in which defaults were used to complete partial knowledge in our perceptions. Up through the late 1980s, a significant part of the literature concerned examples such as "Tweety" the bird who was assumed by default to fly until she was found to be an ostrich.

During the 1980s a series of results made precise the notion of defaults and exceptions in frame systems with multiple inheritance [22] and then went on to show that the resulting systems were worst case intractable computationally [20].

Meanwhile, a series of papers questioned the foundations of representation [23] [2]. Beginning with pioneering efforts such as KRL [1] and KL-ONE [4], interest turned increasingly to logic based representations using the notion of "definition" rather than "prototype" Unfortunately, logic based mechanisms for capturing the notions of default reasoning (non-monotonic reasoning) proved problematic, and all suggested solutions were computationally intractable [6]. On top of this, the KL-ONE family of logic based systems turned out to have serious problems with tractability even without defaults and exceptions [3] that were not overcome until the mid 1990s with the advent of the modern classifiers such FaCT [8, 9] and Racer[7] which underpin OWL-DL.

As attention turned to logic-based formalisms, the nature of the task changed – from using prototypes to fill in incomplete data to classifying definitions and determining their consistency. Classification came to prominence particularly in the biomedical world as people wished to build large multi-axial ontologies where the classification structure was the means to correct retrieval. Such ontologies proved extremely difficult to build by hand [5, 17, 25].

With the change from prototypes to definitions, the form of the typical statement in knowledge representation changed from "*Most* X have property value *Y*" to "*All* X have property value *some/only kind_of* Y" *e.g.* from "*Most* birds have_ability fly" to "*All* birds have_covering *some kind_of* feathers". Clearly, the second form leaves no room for exceptions.

However, despite the benefits of logic-based formalisms such as OWL, the inability to express defaults and exceptions is a serious limitation. An important "baby" had been thrown out with the "bathwater". Without defaults & exceptions it is impossible to make high level generalisations and then refine them by adding exceptions to more specialised cases as they arise. However, this remains the most convenient way to express many notions whether in biomedicine or other fields. In human factors terms, it is almost certainly the most reliable form in which to author safety critical facts, since the defaults can be used to provide a fail-safe value and only overridden when the safety of the exception is established.

In some cases, the greater expressiveness of OWL and modern description logics than previous formalisms allows one to reformulate a knowledge base so as to capture in a logic-based framework notions that previously would have required defaults in a less expressive formalism. In others, some hybrid form of reasoning is required. The purpose of this paper is to explore some of the alternatives in OWL and related languages for dealing with these issues. Note that this paper is concerned only with the representation of classes, since this is primarily where issue of defaults and exceptions arise, and that it focuses on the OWL-DL sublanguage, although most of the remarks apply equally to OWL-full.

1.2 Summary of Analysis

For analysis, we distinguish four cases.

1. *Cases which concern only specialisation rather than exceptions*, *e.g.* that "Blood vessels carry blood" but that "Arteries carry oxygenated blood".
2. *Cases where there is a single local exception. e.g.* "Arteries carry oxygenated blood" with an exception for the pulmonary artery. In these cases it is probably better to reformulate the statement, *e.g.* "Arteries except the pulmonary artery or its branches carry oxygenated blood"
3. *Cases in which there are a modest number of dimensions of "context*", *e.g.* "The normal human manus has five digits", where "normal" and "human" must be represented explicitly.

4. *Cases in which there are a unpredictable number of exceptions, and possibly exceptions to the exceptions*, especially those which need to be maintained in a fail safe manner, *e.g.* drug uses, contraindications and interactions, organising complex forms, linking to external resources, etc.

We approach each case by a different method:

1. *Cases which concern only specialisation,* we deal with routinely in OWL, since they are not really exceptions.
2. *Cases where there is a single local exception*, we suggest are best dealt with by more precise logical formulation.
3. *Cases in which there are a modest number of dimensions of "context",* we suggest are best dealt with by making context explicit and then generalising common information where possible.
4. *Cases in which there are a unpredictable number of exceptions*, we deal with in by one of three approaches: a) to treat such information as part of the knowledge base indexed on the ontology rather than the ontology itself and then to use traditional frame-like methods; b) to compile this knowledge into the ontology from such a representation; c) to express the knowledge using a series of "work arounds" which subtly alter the semantics.

The first three cases can be dealt with purely within the logic based ontology framework. The first can be dealt with in any description logic. The second and third require the recently developed highly expressive description logics which underpin OWL. The fourth case requires more involved methods and will be dealt with separately.

2 Methods in detail

2.1 Principles for cases 1-3 which can be represented OWL alone

In this section we adopt a slightly modified OWL abstract syntax in which we use "→" rather than "subclass-of" in order to emphasise that in OWL, "subclass-of" is equivalent to implication. In Section, 2.2, we explain the OWL representation in more detail including features that are not obvious from the initial presentation.

2.1.1 Case 1: Specialisations

Specialisations are not really exceptions at all since they merely narrow, but do not contradict, the original statement. They are sometimes confused with exceptions because some systems such as Protégé use the same mechanism for both. Specialisations are represented with routinely in OWL or any other description logic. They are mentioned only to distinguish them from other cases.

2.1.2 Case 2: Single Exceptions

The easiest case to deal with is that of single exceptions. Usually these cases arise during the development of an ontology when the author realises that a statement that has been entered is in fact too general. In the example above, it would have been natural to enter the statement "Arteries carry oxygenated blood" neglecting at first the exception of the pulmonary arterial tree. However, with this one qualification, the initial statement is true (at least in normal adult mammals).

Hence instead of:

Artery → (restriction carries someValuesFrom OxygenatedBlood

we make a more precise statement. Taking "Systemic Artery" to mean the branches and sub-branches of the aorta, we say: "The Aorta and its branches[1] carry oxygenated blood":

(Aorta or (restriction isBranchOf someValuesFrom Aorta)) →
(restriction carries someValuesFrom OxygenatedBlood)

2.1.3 Case 3: Representing context explicitly

If the scope of the ontology needs to be extended to broader contexts – *e.g.* abnormal or fetal anatomy, or to multiple species – then one possibility is to represent that context explicitly. For example, the heart is normally contained in the left thorax, but in a small percentage of the population it is abnormally located on the right. We can represent the context "anatomically normal" explicitly by:

Heart and (restriction hasAnatomicallStatus Normal) →
(restriction isContainedIn LeftThorax)

This statement has the consequence that the discovery of a heart in the right thorax implies that it is an anatomical abnormality – precisely what would be expected, and might lead us to look for other associated anomalies.

In a similar situation using species rather than normality, humans have one prostate with three lobes whereas mice have five prostates none of which has lobes[2]. So we might express this as:

Body and
(restriction isOfSpecies someValuesFrom Human) and
(restriction hasAnatomicalStatus someValuesFrom Normal)
→
(restriction hasPart exactly-1 Prostate)

---

[1] For brevity in the examples we shall take 'branch' to be transitive – "branches or sub-branches"

[2] Cornelius Rosse, Personal communication, 2002.

Body and
(restriction isOfSpecies someValuesFrom Mouse) and
(restriction hasAnatomicalStatus someValuesFrom Normal)
→
(restriction hasPart exactly-5 Prostate)[3]

One situation in which this case occurs commonly is when merging ontologies, *e.g.* of human and mouse anatomy. Given a first translation in which all statements are qualified with context in this way, it is then possible to examine the ontologies and see which features were true in all contexts and generalise them to the unqualified entity.

2.2 Implementation in OWL for cases 2-3

All of the above statements can be converted directly into valid OWL by replacing the '→' with '*subclass-of'*[4]. However, in OWL's abstract syntaxes, in the *de facto* standard editor OilEd[5], and indeed in the Lisp based notation for the underlying description logics, the standard way of introducing new classes adds a complication.

OWL and related formalisms distinguish "primitive" classes which have "partial" definitions from "defined" classes which have "complete" definitions. The abstract syntax for the two cases is confusingly similar, distinguished only by the keywords *complete* and *partial*[6]. For a primitive class the syntax is:

class C ***partial*** *subclass-of Super*
$Restriction_1$
...
$Restriction_N$

Whereas for a defined class it is

class C ***complete*** *subclass-of Super*
$Restriction_1$
...
$Restrction_N$

---

[3] DAML+OIL allows such "qualified cardinality constraints" as do almost all description logic and KR formalisms which allow cardinality restrictions other than 0,1,many. OWL v0.0 does not provide such a construct, but it is expected that it will be reinstated in OWL v0.1. See http://lists.w3.org/Archives/Public/www-webont-wg/2003May/0072.html

[4] The use of *"subclass-of"* for implication may seem strange to users new to OWL and related formalisms, but it follows from the fact the definition of *"subclass-of"* – known as "is subsumed by" or "is kind of" in related formalisms. One class is a subclass of another if and only if all individuals in the subclass are also in the superclass – *i.e.* if being in the subclass *implies* being in the superclass. This is what distinguishes "logic based ontology formalisms" from other formalisms, such as frames, which do not impose this requirement, or at least do not use it for inference

[5] http://oiled.man.ac.uk/

[6] Although the OWL standard officially deprecates this syntax it is deeply embedded in tools and likely to persist.

Unfortunately, although these look very similar, they behave very differently. In the case of the primitive definition, *C* individually implies each restriction*,* whereas in the case of the defined class, the conjunction of the superclasses and restrictions jointly defines *C*. It is the difference between a) and b) below:

a) $C \rightarrow Super \;\&\; Restriction_1 \;\&\ldots\&\; Restriction_N$
b) $C \leftrightarrow Super \;\&\; Restriction_1 \;\&\; \ldots\&\; Restriction_N$

In the first case, the simple implication of the conjunction is equivalent to *n* individual implications of necessary conditions.

$C \rightarrow Super;\; C \rightarrow Restriction_{1;\ \ldots}\; C \rightarrow Restriction_n$

In the second, the reverse implication abbreviated in the bi-directional arrow '$\leftrightarrow$' is:

$Super \;\&\; Restriction_1 \;\&\; \ldots\&\; Restriction_N \rightarrow C$

This implication cannot be split. None of the conjuncts is sufficient to imply C individually; rather it is the conjunction as a whole that implies C. Furthermore, if we add a restriction to the conjunction, we are adding to the sufficient conditions for recognizing the concept, not just to the necessary conditions that can be inferred from the concept.

Therefore, as the ontology evolves, and the definition of a primitive concept is made explicit, some or all of the restrictions that appeared in the original class statement need to be moved to separate axioms. For example, consider the example of prostate above. If the ontology had started with human implicit and prostate treated as a primitive, then it would probably have been expressed as:

class Prostate ***partial*** *subclass-of Organ*
restriction hasSubdivision exactly-3 Lobe

When converted to a combined mouse-human anatomy ontology in which the context *Human* had to be made explicit, then the restriction would have to be moved from the class axiom to a separate subclass axiom as shown below:

Class NormalHumanProstate ***complete*** *subclass-of Prostate*
restriction isOfSpecies someValuesFrom Human
restriction hasAnatomicalStatus someValuesFrom Normal

NormalHumanProstate subclass-of
restriction hasSubdivision exactly-3 Lobe

In this form, first (class) axiom introduces and defines the named class *NormalHumanProstate*". The second (subclass) axiom states that "*NormalHumanProstate* has three lobes". Having three lobes is implied by being a human prostate; but it does not imply being a "human prostate".

In our experience, this transformation is a common operation during ontology development. Ontology authors often start by 'sketching' concepts as primitives and then elaborate them by defining those for which definitions are appropriate or adding context where needed. Unfortunately, converting restrictions on the primitive class to subclass axioms is tedious in existing tools, although the authors are engaged in a project to produce improved interfaces that will incorporate this operation as a standard feature[7].

2.3 Case 4:dealing with unpredictable number of exceptions, possibly with exceptions to the exceptions - representations requiring hybrid reasoning and a "ontology indexed knowledge base"

Consider a large richly interconnected evolving knowledge base with numerous axes – e.g. a drug knowledge base classified according to chemical structure, physiological effects, biochemical mode of action, formulation and route of administration. Consider trying to establish protocols for administering such drugs or keeping track of interactions and contraindications. (The Prodigy knowledge base approximates this structure – see [21, 24].)

It is important that as the knowledge base evolves and new drugs are added, that recognised side effects be indicated initially, by default, even though they may in fact be over-ridden. Therefore, to be safe, we want to express interactions and contraindications at the most general level possible and inherit them by default, to be overridden if necessary. If expressed in logic, the result is likely to be a complex expression of the form:

"Drug type A and not subtype B and not subtype C… → ContraindicationX"

If there are exceptions to the exceptions, we get expressions such as:

"Drug type A and not (subtype B and not Subtype B1) and not Subtype C …"

Maintaining such structures is tedious and error prone. By contrast, the classic frame oriented default and exception mechanism is straightforward and gives much less opportunity for error. There are three possible solutions:

1. Treating such information as "contingent knowledge" to be represented in as in a classic frame system. Such 'contingent' knowledge is invisible to the classifier as logically it is of the form "Some Cs have property P" or "Protypically Cs have property P" rather than "All Cs have property P". Although it is invisible the classifier, a hybrid reasoning system can query the knowledge base to find the set of most specific information inherited by each node just as in any classic frame system[22][8]. If the set contains more than one member, then some additional reasoning mechanism must be used to resolve the ambiguity.
2. "Work-arounds" by interpreting the relevant properties as "potentially has value" and being multi-valued. An object can then inherit more than one

[7] http://www.cs.man.ac.uk/mig/projects/current/coode

[8] In systems with multiple inheritance this is the Touretzky distance[17]. A value is a member of the set of most specific values for a property for a target node if there is no intervening node along any path in the multi-hierarchy between the source node for the value and the target node that has a value for that property, *i.e.* if the value is not overridden along any path.

"potential" value. As in case 1, some additional reasoning mechanism must be used to resolve inheritance conflicts, but in this case, the 'potential' values are visible to the classifier.

3. Using the frame representation as a high level language and then compiling the result to the logical format equivalent to the expressions above. This requires resolving all potential inheritance conflicts in advance using some additional reasoning mechanism. Using this method, once compiled, all information is available to the classifier.

None of these solutions escapes Touretzky's result [22] that default reasoning with cancellation is computationally intractable in the worst case. However, our experience is that, if the ontology is 'normalised' or 'untangled', the sets of most specific values rarely contain more than a single value – *i.e.* conflicts are rare. When conflicts do occur, then some additional, application specific, reasoning method is required – *e.g.* for drug interactions, take the most serious; for an information resource, take the union of all values etc.[16, 18]

3 Results and Discussion

In all cases there are three issues: a) expressiveness & correctness, b) scaling of computational performance; c) usability and understandability.

Case 1 – *specialisation* – can be represented routinely within any description logic paradigm including OWL and requires no further comment.

Case 2 – *single exceptions* –expressiveness and correctness are not a problem. These cases can clearly be represented in OWL-DL, which is sufficiently expressive for all instances encountered or so far suggested. The constructs involved have only local effects on computational performance and so do not affect scaling globally.

Case 3 – *dealing with context* – expressiveness and correctness are likewise not a problem. Theoretically, the effects on computational scaling should be modest, since in general, the antecedents of the subclass-axioms required naturally contain at least one primitive, which limits the scope of their impact on performance[9]. So far, experience has been consistent with theory, but larger scale tests are under way. This course should only be taken after simulations to ensure that there are no scaling problems.

Case 4 – *unpredictable and complex exceptions* – requires reasoning outside the pure OWL paradigm, for which we put forward three solutions. Our group has most experience with first solution. Such hybrid reasoning over an "ontology indexed knowledge base" was supported in the system underpinning our previous work, GRAIL [14], and its value is well proven in a range of applications, including:

[9] Such axioms are sometimes called "absorbable" because they can be transformed to avoid the global impact of "general inclusion axioms". See [9]

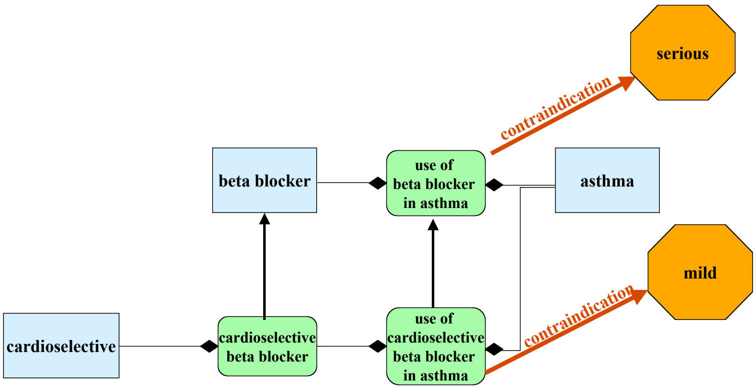

Figure 1: The use of the logic based ontology as an index to contingent information about contraindications for drugs. The notation is derived from UML. Primitive concepts are in rectangles, defined concepts in rounded rectangles, and indexed information in octagons connected by heavy arrows.

1. The PEN&PAD clinical data entry and medical record system [10, 12] in which it was used to index information to be included on data entry forms which could be indefinitely tailored.
2. The Prodigy Drug Ontology [21] in which it is used to handled uses and interactions of drugs.
3. The GALEN modules for encoding to ICD9/10 and SNOMED International [13, 15, 19] in which it provides the mapping to the candidate ICD codes for concepts in the ontology.
4. Internally in the translation from the Intermediate Representation for indexing the transformation rules and mappings [18]

Figure 1 illustrates this type of reasoning about contraindications in the drug ontology. The logic based ontology acts as "conceptual lego" allowing the definition of notions such as "Use of cardioselective beta blockers for asthma". The reasoner classifies such concepts to form a polyhierarchy[10] which then acts as an index to the contingent information.

Figure 2 shows the use of similar mechanisms for constructing complex forms or data structures. This is the mechanism underlying PEN&PAD. Whereas for drug information, one indexed value over-rides the default value, in this application multiple values are normally cumulative[11].

Unfortunately, no system based on current standards supports such hybrid reasoning. This is a major motivation for our development of a hybrid Protégé-OWL environment.[12].

[10] Although only one axis is shown here for simplicity
[11] There is a mechanism to "turn off" unwanted items at a lower level if this is required, but it is omitted for simplicity
[12] http://www.cs.man.ac.uk/mig/projects/current/coode/
http://protege.stanford.edu/plugins/owl

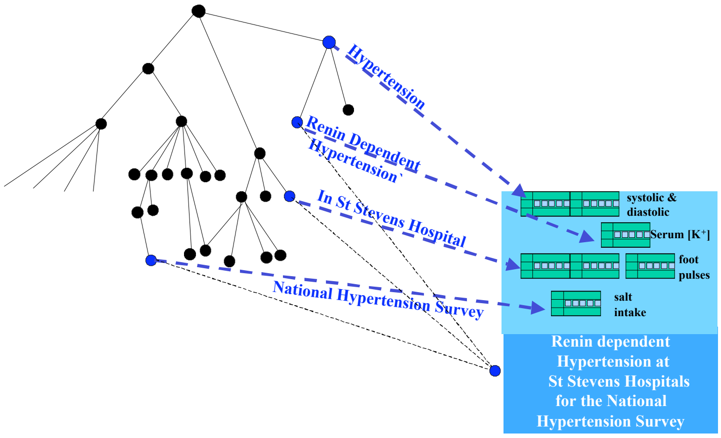

Figure 2: Use of indexing to assemble adaptable forms. The composite concept of “Renin dependent hypertension in St Stevens Hospital in Natioonal Hypertension Survey is first formed and classified, and then the information items required are assembled using ‘inheritance’.

The group also has experience with the second solution, which was used extensively during the development of the drug ontology. It has the advantage of providing a uniform mechanism – classification – for retrieving information from the ontology. Its disadvantage is that the mismatch between the semantics of the ontology and the semantics of the knowledge base as a whole. Like most “work-arounds”, this can lead to unexpected results. Furthermore, the classification operations required are more computationally expensive than the queries required for solution one, since they ask what is true *in any extension of* the given knowledge base whereas a query asks only what is true in *the given* knowledge base.

The third solution has been tested only in simulated examples. There remain significant questions concerning how it will scale computationally, and developing tools to make this mechanism usable is challenging problem for tool builders. At the moment, it remains a theoretical possibility whose practical applicability is speculative

Overall, the lessons for those who have previously used systems supporting defaults and exceptions is that conversion to ontologies based on OWL or related logic-based formalisms requires careful analysis. If the defaults fall into cases one to three, then representing them directly in OWL is probably feasible and desirable because it brings added inferential power, although the computational scaling should be checked. If they fall into case four, then more care is required, and some further reasoning methods will almost certainly be required.

Working prototypes and demonstrations should be available by time of the meeting in January.

Acknowledgements

This work supported in part by the CO-ODE grant from the UK Joint Information Services Committee (JISC). Special thanks to the ontology group at University of Manchester, the Protégé group at Stanford, and the Digital Anatomist project.

References

1. Bobrow, D.G. and Winograd, T. An overview of KRL. *Cognitive Science*, *1* (1977). 3-46.
2. Brachman, R. What IS-A is and isn't: an analysis of taxonomic links insemantic networks. *Computer*, *16* (1983). 30-36.
3. Brachman, R. and Levesque, H., The tractability of subsumption in frame-based description languages. in *AAAI-84*, (1984), Morgan Kaufman, 34-37.
4. Brachman, R. and Schmolze, J. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, *9* (1985). 171-216.
5. Campbell, K.E., Das, A.K. and Musen, M.A. A logical foundation for representation of clinical data. *JAMIA*, *1* (1994). 218-232.
6. Etherington, D. Formalising nonmonotonic reasoning systems. *Artificial Intelligence*, *31* (1987). 41-85.
7. Haarslev, V. and Moeller, R., Expresive ABox reasoning with number restrictions, role hierarchies, and transitively closed roes. in *Proc Seventh International Conference on Knowledge Representation and Reasoning (KR2000)*, (San Francisco, CA, 2000), Morgan Kaufmann, 273-284.
8. Horrocks, I., Using an expressive description logic: FaCT or Fiction. in *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference on Knowledge Representation (KR 98)*, (San Francisco, CA, 1998), Morgan Kaufmann, 634-647.
9. Horrocks, I., Statler, U. and Tobies, S. Practical reasoning for very expressive description logics. *Journal of the Interest Group in Pure and Applied Logics (IGPL)*, *8* (2000). 293-323.
10. Kirby, J. and Rector, A.L., The PEN&PAD Data Entry System: From prototype to practical system. in *AMIA Fall Symposium*, (Washington DC, 1996), Hanley and Belfus, Inc, 709-713.
11. Minsky, M.L. A framework for representing knowledge. in Winston, P.H. ed. *The Psychology of Computer Vision*, McGraw Hill, New York, 1975, 211-277.
12. Nowlan, W., Rector, A., Kay, S., Horan, B. and Wilson, A., A Patient Care Workstation Based on a User Centred Design and a Formal Theory of Medical Terminology: PEN&PAD and the SMK Formalism. in *Symp Computer Applications in Medical Care. (SCAMC-9)1*, (Washington DC, 1991), McGraw-Hill, 855-857.

13. Pole, P. and Rector, A., Mapping the GALEN CORE Model to SNOMED-International: Initial Experiments. in *AMIA Fall Symposium*, (Washington DC, 1996), Hanley and Belfus, Inc, 100-104.
14. Rector, A., Bechhofer, S., Goble, C., Horrocks, I., Nowlan, W. and Solomon, W. The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine*, *9* (1997). 139-171.
15. Rector, A., Rossi Mori, A., Consorti, F. and Zanstra, P. Practical development of re-usable terminologies: GALEN-IN-USE and the GALEN Organisation. *International Journal of Medical Informatics*, *48* (1998). 71-84.
16. Rector, A., Wroe, C., Rogers, J. and Roberts, A., Untangling taxonomies and relationships: Personal and practical problems in loosely coupled development of large ontologies. in *Proceedings of the First International Conference on Knowledge Capture (K-CAP 2001)*, (Victoria, BC, Canada, 2001), ACM, 139-146.
17. Rector, A.L. Clinical Terminology: Why is it so hard? *Methods of Information in Medicine*, *38* (1999). 239-252.
18. Rector, A.L., Zanstra, P.E., Solomon, W.D., Rogers, J.E., Baud, R., Ceusters, W., W Claassen, Kirby, J., Rodrigues, J.-M., Mori, A.R., Haring, E.v.d. and Wagner, J. Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine. *IEEE Transactions on Information Technology in BioMedicine*, *2* (1999). 229-242.
19. Rogers, J.E., Price, C., Rector, A.L., Solomon, W.D. and Smejko, N. Validating clinical terminology structures: Integration and cross-validation of Read Thesaurus and GALEN. *Journal of the American Medical Informatics Association* (1998). 845-849.
20. Selman, B. and Levesque, H.J. The complexity of path-based defeasible inheritance. *Artificial Intelligence*, *62* (1993). 303-340.
21. Solomon, W., Wroe, C., Rogers, J.E. and Rector, A. A reference terminology for drugs. *Journal of the American Medical Informatics Association* (1999). 152-155.
22. Touretzky, D. *The Mathematics of Inheritance Systems*. Morgan Kaufmann, Los Altos, CA, 1986.
23. Woods, W.A. What's in a link: Foundations for semantic networks. in Bobrow, D. and Collins, A. eds. *Representation and Understanding: Strudies in Cognitive Science*, Academic Press, Newwq York, 1975, 35-82.
24. Wroe, C., Solomon, W., Rector, A. and Rogers, J. Inheritance of drug information. *Journal of the American Medical Informatics Association* (2000). 1158.
25. Wroe, C., Stevens, R., Goble, C.A. and Ashburner, M., An Evolutionary Methodology To Migrate The Gene Ontology To A Description Logic Environment Using DAML+OIL. in *Proceedings of the 8th Pacific Symposium on Biocomputing (PSB)*, (Hawaii, 2003), 624-635.