Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Biomedical Informatics xxx (2005) xxx–xxx
Journal of Biomedical Informatics
www.elsevier.com/locate/yjbin

Methodological Review

Methods in biomedical ontology

Alexander C. Yu *

Department of Biomedical Informatics, Columbia University, New York, NY 10032, USA

Received 18 August 2005

Abstract

Research on ontologies is becoming widespread in the biomedical informatics community. At the same time, it has become apparent that the challenges of properly constructing and maintaining ontologies have proven more difficult than many workers in the field initially expected. Discovering general, feasible methods has thus become a central activity for many of those hoping to reap the benefits of ontologies. This paper reviews current methods in the construction, maintenance, alignment, and evaluation of ontologies.

Keywords: Biomedical ontology; Review; Methods; Ontology design; Ontology evaluation; Ontology maintenance

1. Introduction

Research on ontologies is becoming widespread in the biomedical informatics community. At the same time, it has become apparent that the challenges of properly constructing and maintaining ontologies have proven more difficult than many workers in the field initially expected. Discovering general, feasible methods has thus become a central activity for many of those hoping to reap the benefits of ontologies [1–4].

In medicine, the application of ontologies to practical problems is a response to the need to reuse the voluminous and complex information that is involved in many health care activities [5,6]. More recently, the exponential increases in biological data and knowledge have also led to an awareness of the usefulness of ontological methods in biology and, hence, to subsequent efforts to exploit these techniques [7–9]. One important potential benefit of these activities is the bridging of the gap that exists between basic biological research and medical applications. Achieving this would be a significant step towards fulfilling the vision that Blois described already in 1988 [10]:

> "The medical practitioner needs to be able to harness the tools of reasoning better to apply them to a mixture of low-, middle-, and high-level data. This is essential if physicians are to range back and forth, consciously and effectively, from the mathematical descriptions of atomic and molecular events to the statistical associations exhibited by complex biologic systems, and to the natural-language descriptions at the clinical and behavioral levels."

In a similar manner, biological researchers also stand to benefit from being able to harness the clinical data and knowledge that are increasingly stored in computable forms.

2. Definitions of the term 'ontology'

The idea of capturing knowledge in a structured manner is at least as old as Aristotle, who first paid attention in a systematic way to the practical problem of representing the structure of reality. Although philosophy has since accumulated a significant body of analytical tools for ontological problems, many of the ideas and terms in ontology, such as the notion of *category*, and *hierarchy*, can be traced back to Aristotle [11] (Fig. 1).

While philosophical ontology takes many forms, and different schools of philosophy have offered different

---

* Fax: +1 212 342 1647.
E-mail address: alex.yu@dbmi.columbia.edu.

1532-0464/$ - see front matter © 2005 Published by Elsevier Inc.
doi:10.1016/j.jbi.2005.11.006

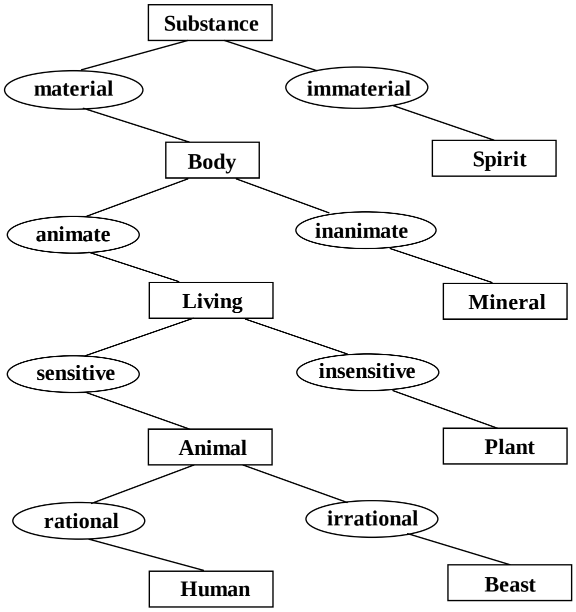

Fig. 1. Tree of Porphyry, with Aristotle's categories (in rectangles). Lines represent *is-a* (subsumption) relationships between categories. Differentiae (in ovals) distinguish species under a common genus. For example, "body" *is-a* material "substance," in comparison to "spirit," which *is-a* immaterial "substance." Adapted from [11].

approaches, one central goal in philosophical ontology is a definitive and exhaustive classification of all entities. Smith defines philosophical ontology as "the science of what is, of the kinds and structures of objects, properties, events, processes and relations in every area of reality [12]."

Smith adopts a *realist* stance, in which the thesis is that reality exists independently of human perception, and that the quality of ontologies depends on the degree to which they represent (are true of) a certain portion of reality [13]. On the other hand, Guarino et al. adopt a *cognitive* bias that considers categories as cognitive artifacts dependent on human perception; they choose to refrain from committing to "a strictly referentialist metaphysics related to the intrinsic nature of the world" [14]. Current efforts are under way to reach a fusion of the Basic Formal Ontology developed by Smith and his associates with the DOLCE ontology developed by Guarino, resting in part on the shared recognition of the fact that there are areas of reality which depend for their existence upon human cognitive acts (for example in the domains of psychology and culture) [15].

Within the Artificial Intelligence (AI) community, the term 'ontology' is predominantly used to refer to a certain class of artifacts that are the results of *ontology engineering*. *Ontology engineering* itself is defined by Gomez-Perez as "the set of activities that concern the ontology development process, the ontology life cycle, the methods and methodologies for building ontologies, and the tool suites and languages that support them" [16]. Gruber's statement, that "an ontology is a specification of a conceptualization," was the first attempt to define the term *ontology* in the AI sense [17]. This definition came under criticism for leaving room for too many interpretations, which led Guarino to attempt to clarify and formalize the AI definition further [18]. Guarino distinguishes and relates the different senses of the term 'ontology' assumed by the philosophical community and the Artificial Intelligence community [19]. In the philosophical sense, ontologies are systems of categories that account for a particular way of seeing the world (this is what Guarino defines as a *conceptualization*). On the other hand, the AI reading of 'ontology' refers to an artifact specified in a particular logically regimented vocabulary (i.e., a *specification*) to describe a certain reality, and where a set of statements are made regarding the intended meaning of the words in the vocabulary.

The term 'ontology' is also frequently used in a way that does not fit into any of the senses described above. Here, the term is used to refer simply to controlled terminologies. For example, the curators of the Gene Ontology (GO) focus on providing a practical framework for keeping track of the biological annotations that are applied to gene products. Although GO uses hierarchies of terms, its authors have focused neither on software implementations nor on the logical expression of the theory encompassing these terms [20].

3. What are ontologies useful for?

3.1. Terminology management

Traditional paper-based terminology systems are generally deemed to be inadequate with respect to the requirements of health care information systems that depend on clear communication of complex medical and biological information in a form that is usable by computers [21]. Not surprisingly, this goal has proved to be a difficult one to achieve, mainly because it requires deep analysis and formal representation of the meanings of terms [22]. Furthermore, the task of maintaining terminologies is a significant challenge in itself [23–25]. The adoption of an ontological approach for managing biomedical terminologies facilitates some of the tasks associated with these activities, as workers in both clinical [3,5,26–28] and biological [7] domains have found.

While the ontological approaches that have been adopted have mostly come from computer science, workers are increasingly turning to philosophy for formal ontological methods and insights that can help them to address many of those problems which have not traditionally fallen within the purview of computer science [29]. Examples include a study on the compliance of SNOMED CT with respect to formal ontological principles [30], and work on defining formal relations for the Open Biomedical Ontologies [8].

3.2. Integration, interoperability, and sharing of data

We need to be able to share data and support interoperability among disparate health care applications and information systems. In medicine, this is important for purposes of facilitating continuity of health care; in biological research, it facilitates the sharing of experimental data among researchers. A common semantics is an essential element for these goals to reach fruition. An example of the use of ontologies towards this goal is the work of the HL7 on the Reference Information Model (RIM) [31]. The RIM is meant to represent the "semantic and lexical connections between the information carried in the fields of HL7 messages," which are communicated electronically in standardized formats to relay health care messages. One problem, however, is that it does not distinguish in a clear and stable manner between information and the objects in reality which such information is about [32].

As discussed briefly in Section 1, as the amount of information in both biology and medicine has increased, it has become a central problem to find ways to seamlessly integrate information and data from the clinical and biological domains. Along these lines, Kumar et al. have described seminal work on the creation of an integrated framework through the application of formal ontological principles to available biomedical ontologies [33]. The possible practical applications of this sort of integration include the support of applications such as decision support systems that draw inferences across the levels of granularity which span biology and medicine.

3.3. Knowledge reuse and decision support

Knowledge-based systems that support applications such as decision support in health care are typically dependent on large amounts of current domain knowledge [34,35]. However, capturing knowledge is an expensive and arduous process, and it would be beneficial to create ontologies that are application independent and can be reused in new systems without additional development work. Musen's work on re-usable problem-solving methods and ontology-driven knowledge acquisition in the Protégé project [36], and the work of Rosse et al. on the Foundational Model of Anatomy (FMA), are salient examples of efforts at creating and reusing domain ontologies [3]. Notably, the FMA is described by its creators as a reference ontology for biomedical informatics, i.e., an ontology that serves "as a foundation and reference for the correlation of other ontologies." This contention is rooted, first, in the generality and ubiquity of its intended domain (anatomy, from the level of the whole organism down to that of biological macromolecules). Second, its curators strive to consistently apply rigorous formal rules in developing its taxonomy and partonomy, in a way that is designed to facilitate its alignment with other ontologies [37].

4. Methods for constructing ontologies

4.1. Representation formalisms

One of the crucial decisions in ontology construction is to select the formalism in which the ontology will be implemented. Many formalisms, such as KIF [38], Ontolingua [39], LOOM [40], and network-based structures (i.e., semantic nets [41] and frames [42]) have been used in recent decades; each has its particular strengths and limitations [16]. More recently, the growth of the Internet also led to the creation of *web-based ontology languages* (or *ontology markup languages*), such as RDF [43], RDFS [44], DAML + OIL [45], and OWL [46], that exploit the characteristics of the World Wide Web. In particular, OWL is the result of the World Wide Web Consortium's efforts to create a standard ontology markup language for the Semantic Web. Its semantics are based on a subset of description logics (DLs). DLs are a family of ontology representation languages that are equipped with a formal, logic-based semantics and are increasingly used for many ontologies [47]. Their success can partly be attributed to two factors. First, significant work has been done on discovering DLs that allow for the expression of moderately complex knowledge without having to sacrifice reasonable performance times on useful tasks such as logical consistency checking and automated classification of concepts. Second, relatively sophisticated tools for editing and reasoning with DL-based ontologies are now available. For example, the Protégé ontology editor has an OWL plug-in that facilitates creating and reasoning with ontologies specified in OWL through a graphical user interface [48,49]. Despite the significant amount of work done on representation formalisms, significant challenges still remain, particularly when it comes to the issues of expressing uncertainty [47] and capturing knowledge about defaults and exceptions [50].

4.2. Fundamental ontological theories

Over the past 2400 years, philosophers have developed analytical tools and theories that address ontological problems. Among the most important for our purposes are fundamental theories that deal with, first, the relationships between classes and their instances and, second, the taxonomical relationships between classes.

(1) Classes, instances, and instantiation. The term "class" refers to what is *general* in reality, and is—modulo the problems outlined in [51]—broadly equivalent to the notions of "concept" in the knowledge representation literature and "universal" or "type" in the literature of philosophical ontology. The idea of "instance" (alternatively, "token" or "individual") refers to what is *particular* in reality (i.e., to those entities which exist in space and time) and plays a fundamental role in the definition of what it means for one class to stand in relation to another. Furthermore, while each instance is bound to a particular location in space and time and exists as it were in itself, classes are multiply located and exist only in their respective instances [37]. Assertions of relations between classes can thus be conceived as assertions about the corresponding instances. For example, if we have two classes *cell* and *cell nucleus*, then (as is argued in [8]), we cannot make sense of what

it means to say *cell nucleus part_of cell* unless we realize that this is a statement to the effect that each instance of the class *cell_nucleus* stands in an instance-level part relation to some corresponding instance of the class *cell*.

(2) Genera, differentiae, taxonomies, and subsumption. A *semantic network* is the result of applying a graphical notation for representing knowledge in patterns of interconnected nodes and arcs. The first depiction of what we now call a semantic network almost certainly appeared in the philosopher Porphyry's *On Aristotle's Categories* in the third century AD [11]. It was a tree with Aristotle's categories arranged by *genus* (supertype) and *species* (subtype); features called *differentiae* were used to distinguish the species of the same genus. Over the years, formal principles of *classification* (see Table 1) have been elaborated, and many of them arguably rest on a wide consensus among workers in ontologies and terminologies. Principles of *subsumption* (Table 1), on the other hand, have been derived from studying empirically the way subsumption is treated in biomedical terminologies and ontologies [30].

4.3. General ontology development methodologies

A number of general methodologies for developing ontologies have been described in the knowledge representation literature. In 1990, Lenat and Guha reported on the general steps they used in the development of Cyc, a large knowledge base of common sense knowledge [52]. The initial step consisted of manual extraction and coding of common sense knowledge. When enough knowledge had been entered into the system, tools for analyzing natural language and machine learning tools could use the knowledge already entered to aid in the process of adding other knowledge.

Later on, Uschold, King, and Gruninger proposed formal guidelines for ontology building, born of the experience gathered in developing the Enterprise Ontology [53]. According to their approach, these key processes are to be carried out: (1) identify the ontology's purpose, (2) build the ontology, (3) evaluate the ontology, and (4) document the ontology. *Ontology capture*, the main task in ontology building, consists of identifying and defining key concepts and relationships in the domain of interest. Concepts are defined not in the style of typical dictionaries, but are built by using philosophical notions such as *class* and *subsumption* (e.g., *Car* is a class that is a subclass of *Vehicle*). Furthermore, top-down, middle-out, or bottom-up strategies can be used to systematically identify concepts, depending on whether general, middle-level or specific concepts were identified first. The particular strategy one uses would affect the final level of detail captured in the ontology.

Table 1
Principles of (A) classification and (B) subsumption

(A) Principles of classification

1. Each hierarchy must have a single root.
2. Each class (except for the root) must have at least one parent.
3. Non-leaf classes must have at least two children.
4. Each class must differ from each other class in its definition. In particular, each child must differ from its parent and siblings must differ from one another.

(B) Principles of subsumption

1. *Inheritance principle*: if *A* is a child of *B* then all properties of *B* are also properties of *A*.
2. Children can differ from their (subsuming) parents in one of two possible ways:
   a. Introduction in the child of a new criterion.
   b. Refinement of an already existing criterion.

Adapted from [30].

Based on the experience of building the Toronto Virtual Enterprise ontology, Gruninger and Fox described a formal approach to build and evaluate ontologies [54]. The most important innovation in their work was to incorporate a set of competency questions (formulated in formal logic) that could be used to rigorously evaluate the ontology. Once the competency questions were formally stated, conditions for completeness (i.e., *completeness theorems*) could be defined that could be used to determine whether competency questions had been answered. Other general ontology development methodologies have also been reported in the literature [55–58].

4.4. Top-level ontologies

Top-level ontologies (or upper-level ontologies) describe the most general concepts or categories that are presumed to be common across domains. Prominent examples of top-level ontologies include DOLCE [14], Basic Formal Ontology [14], Cyc's upper ontology [52], Sowa's top-level ontology [11], the UMLS Semantic Network [59], and the top level of GALEN [27]. Top-level ontologies can be used as a formal foundation for building domain ontologies—doing so can facilitate semantic integration across ontologies at a later time. Alternatively, domain ontologies can also be built first and then linked to top-level ontologies [60,61] (Fig. 2).

The fundamental ontological commitments and distinctions that are laid out in coherent top-level ontologies are part of the reason they can be useful in decision-making during ontology construction. For example, one of the most basic distinctions among entities is made between *continuants* (or *endurants*) and *occurrents* (or *perdurants*) [11,14]. Continuants are those entities which exist in full (i.e., including all their parts) at every instant in time at which they exist, while occurrents are those entities which unfold through time and never exist in full at any single moment in time. Examples of continuants are: you, a surgeon's scalpel, your arm, and your wristwatch. Examples of occurrents include your life, the movement of your blood through your blood vessels, and the execution of a surgical procedure. Based on this fundamental distinction, a number of axioms can be formulated that constrain what can be stated about the interactions between continuants and occurrents, such as: although continuants can participate in occurrents (e.g., you are a participant in your life),

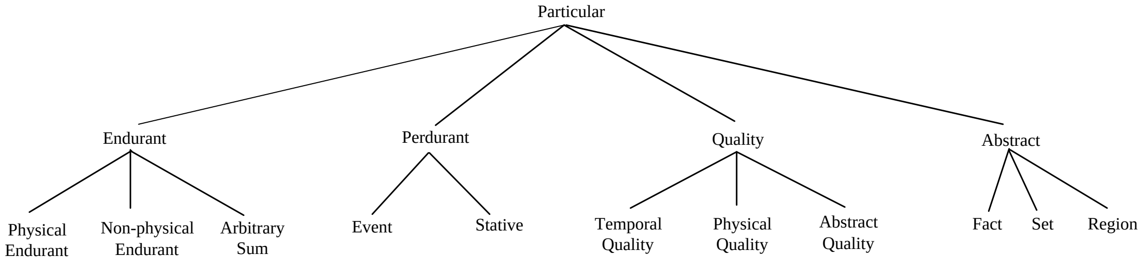

Fig. 2. Top level of DOLCE showing its basic categories. Lines represent *is-a* (subsumption) relationships between categories. Adapted from [14].

continuants cannot be part of occurrents (e.g., you are not part of your life) [62].

4.5. Biomedical ontologies

New ontologies in biology and medicine continue to proliferate as the need for them arises. Some of the most well-studied and prominent examples are presented here.

Foundational Model of Anatomy. One of the most coherently structured ontologies in biomedicine is the Foundational Model of Anatomy (FMA), a domain ontology of the classes and relationships that pertain to the structural organization of the human body [3]. Its developers have extensively described the disciplined approach they used, which relied on a set of declared principles, high-level schemes, Aristotelian definitions, and a frame-based formalism [63]. Efforts are underway to convert the frame-based representation of the FMA into a description logic-based representation using OWL [64]. Although initially developed as an enhancement of the anatomical content of the UMLS, the FMA is now being proposed as a reference ontology useful for purposes of correlating different views of anatomy, aligning existing and emerging ontologies in bioinformatics, and providing a structure-based template for representing biological functions (Fig. 3).

(2) GALEN Common Reference Model. The goal of the GALEN project is to provide re-usable terminology resources for clinical systems [27,65]. At the heart of GALEN is the Common Reference Model, an ontology formulated in a specialized description logic, GRAIL. Its curators have described the ontological issues they encountered, as well as the basic principles and specific methods they utilized to deal with various modeling challenges. Some of the most interesting problems involved the handling of uncertainty, the representation of knowledge about diseases, and the representation of defaults and exceptions. An example of the last is the issue of how to represent knowledge about drug interactions. Description logics, unlike frame-based or semantic network-based formalisms, typically do not allow the expression of knowledge involving default values and exceptions, such as: "in general, the use of beta-blockers is a serious contraindication if the patient has asthma, except when the beta-blocker is cardioselective, in which case it is only mildly contraindicated." To work around this limitation, Rector et al. have shown (see Fig. 4) that a logic-based ontology can be used as an index to "extrinsic" information that one cannot incorporate directly within the ontology [50]. GALEN is no longer being actively developed and is by no means a comprehensive ontology in its current state.

Anatomical structure
is a **material physical anatomical entity**
    which *has* inherent 3D shape;
    is generated by coordinated expression
        of the organism's own structural genes;
    consists of parts that
        are **anatomical structures**
            spatially related to one another in patterns
            determined by coordinated gene expression.

Fig. 3. The definition of the class "anatomical structure" in the Foundational Model of Anatomy (FMA). The definition in structured text shown above is equivalent to the actual frame-based representation used in the FMA. In this definition, material physical anatomical entity is the genus under which **anatomical structure** belongs, while the other parts of the description are differentiae that distinguish anatomical structure from any other types that might also be subsumed by material physical anatomical entity. Adapted from [3].

(3) Medical Entities Dictionary. The Medical Entities Dictionary (MED) is a concept-oriented terminology developed and used in Columbia University and the New York Presbyterian Hospital (NYPH) [5]. It currently contains approximately 97,000 concepts organized into a semantic network of frame-based term descriptions, encompassing those terms used in laboratory, pharmacy, radiology, and billing systems. It includes knowledge about synonyms, taxonomic and other types of relations, and mappings to other terminologies. Cimino has described examples of the many uses various workers have found for the MED, some which are real-world applications used by health care workers at NYPH. Over the years, the MED has been used to support various applications such as data retrieval from medical records, "just in time" medical education, expert systems, data mining, and knowledge-based terminology maintenance [5,66] (Figs. 5 and 6).

(4) National Cancer Institute Thesaurus. The NCI Thesaurus is a description logic-based terminology that is a component of the US National Cancer Institute (NCI) Bioinformatics caCORE distribution. It is created and distributed by the NCI's Center for Bioinformatics and Office of Cancer Communications for use by the NCI's own researchers and the cancer research community as a whole.

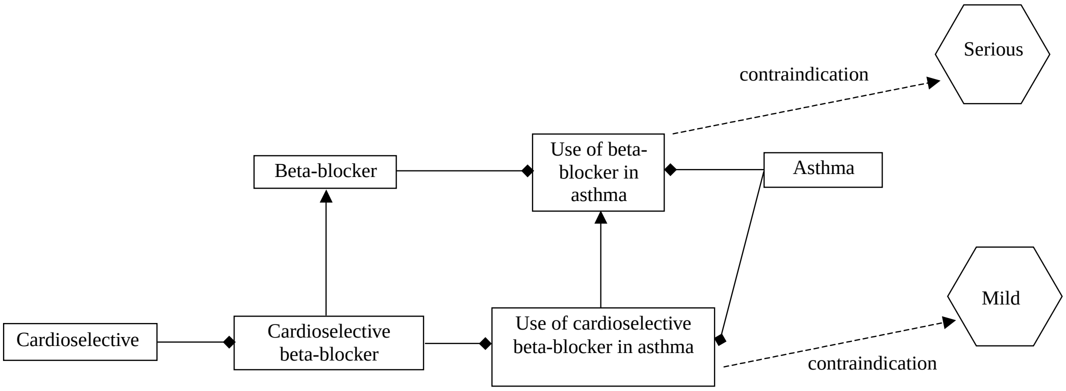

Fig. 4. The use of a logic-based ontology as an index to contingent information (stored outside of the ontology) about contraindications for drugs [60]. Concepts are in rectangles and indexed information are in octagons connected by heavy arrows. This method of linking to contraindication information “outside” of the ontology allows for the specification of default knowledge at different levels of specificity. Adapted from [65].

Serum Glucose Test
    is-a: **Laboratory Test**
    has-specimen: **Serum Specimen**
    measures: **Glucose**

Fig. 5. Frame-based representation of Serum Glucose Test in the Medical Entities Dictionary. The other concepts (Laboratory Test, Serum Specimen, and Glucose) are also represented with their own knowledge. Adapted from [5].

One of its main goals is “to make use of current terminology ‘best practices’ to relate relevant concepts to one another in a formal structure, so that computers as well as humans can use the Thesaurus for a variety of purposes, including the support of automatic reasoning.” The NCI Thesaurus serves several functions within NCI, including annotation of the data in the NCI’s repositories and search and retrieval operations applied to these repositories. At the same time, its designers have intended that its ontological properties should pave the way for more complex uses such as automated indexing, bibliographic retrieval, and linkage of heterogeneous resources. Therefore, it is also linked to other information resources, such as the NCI’s own caCore, caBIO and MGED, and also external ontologies such as the Gene Ontology and SNOMED-CT. Furthermore, it is available in several formats under an Open Source License on the NCI’s website [67,68] (Fig. 7).

Although the NCI Thesaurus has the potential to be used for “more complex uses” by virtue of its ontological properties, the Thesaurus currently falls short in terms of conforming to formal principles of design. Ceusters et al. performed a qualitative analysis of the Thesaurus (version 04.08b, August 2, 2004) to assess its conformity with principles of good practice in terminology development and ontology building, as put forward, respectively, by relevant ISO terminology standards and ontological principles advanced in the recent literature. They found a number of problems related to various things such as definitions of the concepts, term formation, ontological properties,

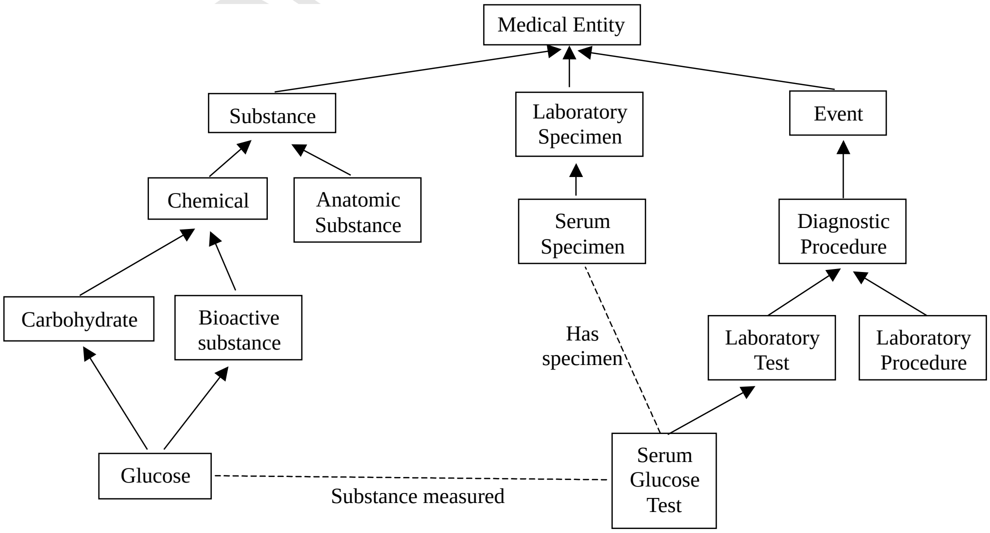

Fig. 6. Example from the Medical Entities Dictionary. The concept Serum Glucose Test is shown in relation to its parent in the *is-a* hierarchy (solid lines) and by non-hierarchic semantic links (broken lines) to other concepts in the network. Adapted from [5].

$$\begin{aligned} Lymphoma &\sqsubseteq T \\ Hodgkin's\,Lymphoma &\equiv Lyphoma \sqcap \\ &\exists DiseaseHasNormalCellOrigin.(B-Cell \sqcup T-Cell \sqcup NK-Cell) \end{aligned}$$

Fig. 7. The concepts Lymphoma and Hodgkin's Lymphoma as represented in description logic in the National Cancer Institute Thesaurus. Hodgkin's Lymphoma is defined as a lymphoma in which the normal cell origin is a B-cell, a T-cell, or a natural killer cell. Adapted from [67].

and its description logic representation (in OWL). In particular, one ontological deficiency that they found was the unprincipled way in which the class hierarchy was built up, bringing it about that basic ontological distinctions were ignored (e.g., between continuants and occurrents) [69]. In another study, Kumar and Smith found similar problems when they examined the NCI Thesaurus with regards to its suitability for representing entities in an ontology of colon carcinoma [70].

(5) Gene Ontology. The Gene Ontology (GO) project was created to address the need for consistent representation of gene product information in different databases [7]. The project began as a collaboration among curators of three model organism databases: FlyBase (*Drosophila*), the *Saccharomyces* Genome Database (SGD), and the Mouse Genome Database (MGD). Since then, it has grown to include many databases, including some of the world's major genome repositories. The use of GO terms by several collaborating databases facilitates uniform queries across them. The GO project maintains a bibliography of peer-reviewed publications at http://www.geneontology.org/doc/GO.biblio.html and include reports of novel uses of GO terms and gene product annotations in interpreting large-scale experimental results [71].

In terms of structure, GO is divided into three ontologies whose topmost nodes are *Cellular component*, *Molecular function*, and *Biological process*, respectively. Together, they allow for the description of gene products in terms of these categories, that is to say they allow the formulation of answers to the three most important types of questions which arise when a new gene product is discovered (Fig. 8):

1. Where is it located in the cell?
2. What functions does it have on the molecular level?
3. To what biological processes do these functions contribute?

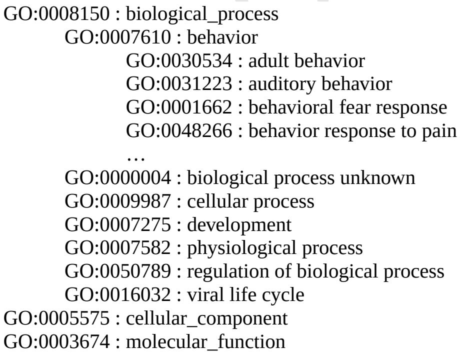

GO:0008150 : biological_process
  GO:0007610 : behavior
    GO:0030534 : adult behavior
    GO:0031223 : auditory behavior
    GO:0001662 : behavioral fear response
    GO:0048266 : behavior response to pain
    …
  GO:0000004 : biological process unknown
  GO:0009987 : cellular process
  GO:0007275 : development
  GO:0007582 : physiological process
  GO:0050789 : regulation of biological process
  GO:0016032 : viral life cycle
GO:0005575 : cellular_component
GO:0003674 : molecular_function

Fig. 8. Part of the gene ontology, which has three topmost nodes: biological process, cellular component, and molecular function (screen capture taken with the AmiGO browser, available at http://www.godatabase.org). Ellipsis indicates parts of GO that are not shown in the figure.

The ontologies are structured by the relations of subsumption (*is a*) and of partonomic inclusion (*part of*). GO treats its three structured networks as separate ontologies; no ontological relations are defined among them. GO has been found to suffer a number of problems, among which is the inconsistent treatment of relations such as *is-a* [20]. Despite its limitations, GO has achieved widespread use in the biological community, and efforts are underway to represent GO in a description logic to improve its suitability for use by computers [72].

(6) Unified Medical Language System. The stated purpose of the US National Library of Medicine's (NLM) Unified Medical Language System (UMLS) is "to facilitate the development of computer systems that behave as if they 'understand' the meaning of the language of biomedicine and health." To that end, the NLM produces and distributes the UMLS Knowledge Sources to be used by system developers in the creation of diverse informatics applications. The Metathesaurus is a large, concept-centered terminology database that is built from the electronic versions of various code sets, thesauri, classification, and lists of terms. On the other hand, the semantic network provides a categorization of the concepts represented in the UMLS Metathesaurus and a set of relationships between these concepts. The current release of the semantic network contains 135 semantic types (as nodes) and 54 relationships (as links between nodes). Types are defined with textual descriptions and by means of the information inherent in its hierarchies. Major groupings of semantic types include those for organisms, anatomical structures, biologic function, chemicals, events, physical objects, and concepts or ideas [59].

Many studies evaluating the usefulness of the UMLS as a terminology and knowledge resource for tasks ranging from terminology translation to domain ontology construction have been published in recent years [73–76]. Other studies have focused on the issue of the role of the UMLS Semantic Network itself as an ontology of the biomedical domain. In a study analyzing the compatibility of the UMLS Semantic Network with ontologies containing general concepts, Burgun and Bodenreider [77] carried two sets of mappings. First, they manually mapped UMLS semantic types to concepts in the Upper Cyc Ontology (1997 release). They also manually mapped UMLS concepts under the same semantic type to WordNet hyponyms under a given synset. In the study, they found two major barriers to mapping. First, classes that had similar names in different ontologies could have distinct meanings (e.g., "Body Part" in Cyc and UMLS mean different things). Second, two classes could have the same intensional meaning even as their extensions in different ontologies differed. For example, although "Symptom" has equivalent definitions in WordNet and in the UMLS, "Symptom" in WordNet encompasses "encephalitis" as well as other conditions that are classified as "Disease or Syndrome" in the UMLS. In another study, Smith et al. proposed revisions to the semantic network that were intended to correct for structural problems. Their suggestions were based on the results of a formal audit that identified several problems. For example, the semantic network frequently runs together *is-a* with *part-of* relations, so that *plant roots is-a plant*, and *plant leaves is-a plant* are allowed [78].

(7) SNOMED-CT. SNOMED CT is arguably the most comprehensive clinically oriented medical terminology system in existence [79], and it is envisioned by its curators as a "reference terminology," i.e., it is made up of "concepts and relationships that provide a common reference point for comparison and aggregation of health care data" [26,80] (Fig. 9). Recently, the US National Library of Medicine (NLM) issued a contract to the College of American Pathologists for a perpetual license for the core SNOMED CT and ongoing updates, which means that SNOMED CT has the potential to be widely used in the United States. Moreover, it has recently been incorporated into the UMLS [81].

SNOMED CT was formed by the convergence of SNOMED RT and Clinical Terms Version 3 (formerly known as the Read Codes) and is expressed in a description logic. As of October 2005, it contains 366,170 unique concepts. The first level of concepts is subdivided into 18 concepts, each of which is the most general concept in a different *is-a* hierarchy (which is called an *axis*), so that all other concepts in SNOMED CT are subsumed within one or more of these hierarchies. Each concept has a description consisting of at least a unique identifier and a unique, fully specified name. In addition, it may also have alternative names, parents in the hierarchy, and relations (which are called *roles* in *description logic*) to other concepts. Thus, SNOMED CT's underlying description logic-based structure has allowed its curators to formally represent the meanings of concepts and the interrelationships between concepts. This, in turn, has allowed them to support tasks such as the elimination of concept redundancy and ambiguity [82].

$$\begin{aligned} 44558001 \equiv\ & 120205009 \sqcap \\ & 84744001 \sqcap \\ & \exists SITE.90785001 \sqcap \\ & \exists METHOD.257903006 \sqcap \\ & \exists DIRECT-MORPH.414402003 \end{aligned}$$

Fig. 9. SNOMED CT definition of Repair of inguinal hernia (44558001) in description logic. Unique codes are used to refer to concepts: Inguinal region repair (120205009); Repair of hernia of abdominal wall (84744001); Inguinal canal structure (90785001); Repair-action (257903006); Hernial opening (414402003). *Repair of inguinal hernia* is defined as an *inguinal region repair procedure* that is also a *repair of hernia of abdominal procedure*, and in which the site of repair is the *inguinal canal*, the method is *repair-action*, and the morphology is a *hernial opening*. Adapted from [80].

Despite its advantages, SNOMED CT still suffers from a number of problems. Bodenreider et al. found SNOMED CT to be non-compliant with a number of ontological principles, which could conceivably result in undesirable consequences. For example, they found the descriptions of many concepts to be minimal or incomplete, with possible "detrimental consequences on inheritance" [30]. In another study, Ceusters et al. used a novel method to detect problems in SNOMED CT and classified them into three broad categories. Problems caused by human error included improper assignments of both *is-a* and non-*is-a* relationships. Other problems, such as shifts in meaning in the migration from SNOMED RT to SNOMED CT, and redundant concepts, were thought to be technology induced. Still others were caused by a lack of ontological theory [29].

4.6. Specific methods for some key problems

Because biology and medicine are such rich and complex domains, many specific methods have either been developed for problem areas that are prominent in biomedical ontology construction or applied to these problem areas after having been developed for other domains.

(1) Representations for partonomic reasoning. A significant number of concepts in biology and medicine are based on anatomy and hence dependent on relations between parts and wholes (*partonomy*). There can also be parts and wholes in the realm of occurrents (process and their subprocesses). Important problems in this area include issues of *transitivity* and *part-whole specialization*. *Transitivity* has to do with representing knowledge such as "if an anatomical entity A is part of another (e.g., the appendix is part of the ascending colon), which itself is part of a larger structure (e.g., the ascending colon is part of the large intestine), then A is also a part of the larger structure." *Part-whole specialization*, on the other hand is defined by the inheritance of relations other than *is-a* (subsumption) along part-whole taxonomies (e.g., "a disease of a part is a disease of the whole").

To reason about part-whole relations, the GALEN project uses axioms that are equivalent to "R specializedBy S" (in GRAIL notation), where R and S are relations. Hence, if R and S are "hasLocation" and "isPartOf," respectively, then one can logically infer from the statements in Fig. 10 that a disease located in the aortic valve is also located in the heart. SNOMED also has an equivalent representation scheme [6].

Hahn et al. have developed an alternative representation for partonomic relations based on the "SEP triplet"

$$\begin{aligned} \exists hasLocation.(\exists isPartOf.Heart) &\sqsubseteq \exists hasLocation.Heart \\ AorticValve &\sqsubseteq \exists isPartOf.Heart \\ \exists hasLocation.AorticValve &\sqsubseteq \exists hasLocation.Heart \end{aligned}$$

Fig. 10. GALEN uses axioms following the pattern “R specializedBy S” (where R and S are relations) to perform partonomic reasoning. In this example, the relations “hasLocation” and “isPartOf” are used to infer that anything that is located in the aortic valve is also located in the heart. Entities that are located in entities that are part of the heart are themselves located in the heart. The aortic valve is part of the heart. Anything that is located in the aortic valve is also located in the heart. Adapted from [6].

approach, which attempts to capture much of partonomic reasoning within a framework compatible with standard Description Logics [83–85]. In the SEP-triple approach, each anatomical part X is represented by a parent concept X_s, and two subsumed concepts X_e and X_p. An instance of X_e represents an entity as a whole, and its associated X_p instance stands for the entity's parts. For all parts Y of X, X_p subsumes Y_s, and since Y_s subsumes both Y_e and Y_p, both the entire part Y_e and all of its parts Y_p are subsumed by the parts of X. While explaining the reasoning procedure to be used with this structure is outside the scope of the paper, suffice it to say that the approach allows for the expression of useful statements such as “a disease of a part must be a disease of the whole structure, but not of the whole taken as in its entirety” (e.g., a disease of the left liver lobe is a disease of the liver, but it doesn't imply that the entire liver is diseased), and “diseases of parts are diseases of the whole, but surfaces of parts are not surfaces of the whole” [6].

(2) Domain modularization for maintainability, re-use, and evolution of large ontologies. Ontologies in biomedicine tend to be large and complex, and in time become difficult to manage, especially where multiple authors are allowed to make changes. Modularization of domain ontologies is therefore a desirable feature because it allows for the distribution of maintenance work among independent authors and the independent evolution of the modules [24].

In order for modularization to work, domain ontologies need to be represented in a *normalized* form. This means that modules are represented as disjoint trees of classes, and relations between classes in different modules are established, such that a classifier can later compute the resulting subsumption hierarchy when modules are combined [86]. Rector has noted that while normalization is an established method in database design, no similar methodology exists yet for ontologies. He has proposed a two-step normalization process for ontologies (see Fig. 11) [24]. The first step consists of using Guarino and Welty's OntoClean methodology for cleaning up taxonomies (see section below) [87]. The second step is an “implementation normalization” mechanism for creating disjoint taxonomic trees of ontological primitives, which can then be later recombined using definitions and axioms to represent other concepts. Bittner and Smith have also shown that top-level

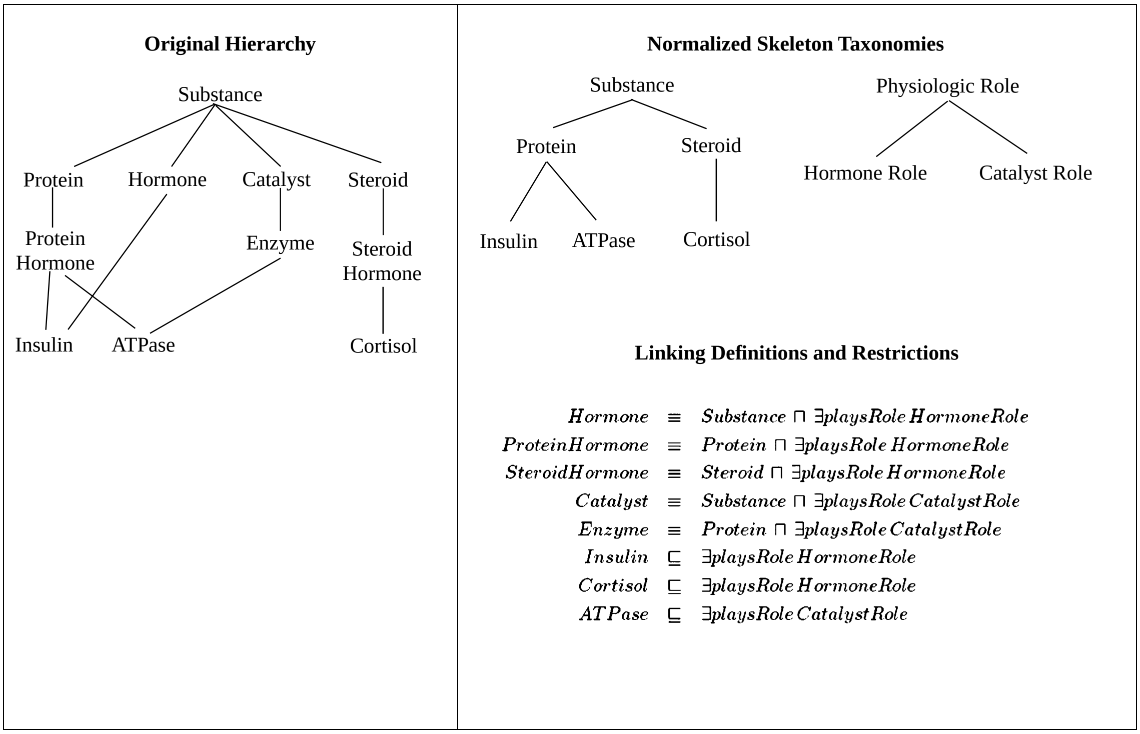

$$\begin{aligned} Hormone &\equiv Substance \sqcap \exists playsRole\, HormoneRole \\ ProteinHormone &\equiv Protein \sqcap \exists playsRole\, HormoneRole \\ SteroidHormone &\equiv Steroid \sqcap \exists playsRole\, HormoneRole \\ Catalyst &\equiv Substance \sqcap \exists playsRole\, CatalystRole \\ Enzyme &\equiv Protein \sqcap \exists playsRole\, CatalystRole \\ Insulin &\sqsubseteq \exists playsRole\, HormoneRole \\ Cortisol &\sqsubseteq \exists playsRole\, HormoneRole \\ ATPase &\sqsubseteq \exists playsRole\, CatalystRole \end{aligned}$$

Fig. 11. Normalization of an ontology of biological substances and roles, according to the method described by Rector. The original hierarchy is shown on the left, and the resulting normalized, disjoint skeleton taxonomies are shown on the top right. Lines stand for *is-a* links. Disjoint skeleton taxonomies of ontological primitives can be later recombined using definitions and axioms to represent other concepts (bottom right). Adapted from [24].

ontologies can be useful for ontology normalization because they provide: (1) basic categories and distinctions that help in forming the appropriate trees and (2) a list of relations together with the axioms that specify their semantics [86]. Significant issues still remain to be addressed if modularization is to work, including the problem of how to determine what sorts of modules make the most sense in a given domain.

(3) Partitions and Granularity. Bittner and Smith have proposed a formal theory of granular partitions (TGP), "cognitive devices designed and built by human beings to fulfill various listing, mapping and classifying purposes." Granular partitions are ways of structuring reality, in our representations, to make the objects and relations in given domains more easily graspable by cognitive subjects. The theory is also intended to address problems associated with the use of set theory and mereology as tools of formal ontology. For example, set theory and mereology are both unable to support the distinction between natural totalities (e.g., the species *cat*, the totality of molecules in your body) and ad hoc totalities (e.g., the set {my left eye, the earth's mantle}). Furthermore, both have their particular problems when it comes to dealing with relations between entities at different levels of granularity. Set theory treats all the members of a given set as, effectively, atoms; mereology treats all parts as on an equal level, which means that it has no means to block the transitivity of the part-whole relation. The two parts of TGP essentially define well-formedness conditions for granular partitions (and taxonomies) and the projective relations these partitions (and their cells) have with the entities in reality [86]. Different projection relations can then be defined for different granular levels, in such a way that the architecture of complex objects or processes (for example an organism, the workflow in a large hospital) can be perspicuously represented. The theory has been applied to a number of problems, including the creation of an ontology for task-based clinical guidelines [88].

5. Ontology merging and alignment

The merging and alignment of ontologies are currently an active area of research in the ontology community. Merging and alignment of ontologies are problems generally referred to under the heading of *semantic integration* in computer science. We provide a brief survey of existing general approaches, largely based on Noy's review of ontology-based approaches to semantic integration, and also describe some efforts specific to the biomedical domain [89].

The work on semantic integration in ontologies can be roughly divided into the areas of: mapping discovery, mapping representation, and reasoning with mappings. We limit our discussion to the discovery and use of mappings. Mapping discovery methods are used to find similarities between two ontologies. Methods in this area can be divided into two general categories. For the first approach, ontologies are developed for the explicit goal of future integration of other ontologies. Top-level ontologies can be used in this way. The idea is that a general top-level ontology is agreed upon by different developers, who then extend this top-level ontology with concepts and properties specific to their application domains. Mapping between extensions can be facilitated by this common "grounding," as long as the extensions are performed in a way that is consistent with the definitions in the shared ontology. As described in the previous section on top-level ontologies, a number of formal top-level ontologies have been created that can be used for this purpose. For example, DOLCE and BFO are two of the formal foundational ontologies developed as top-level ontologies in the WonderWeb project [14]. In work that is specific to biomedicine, Smith et al. have proposed formal definitions for bio-ontological relations. The Open Biomedical Ontologies Relation Ontology (http://obo.sourceforge.net/relationship/) is an attempt to answer the question of how relations such as *part_of* or *located_in* should be defined to ensure maximally reliable curation of different ontologies while at the same time guaranteeing maximal leverage in building a solid base for life-science knowledge integration in general [8]. Noy argues that while many researchers hope that domain- and application-specific ontologies will reuse top-level ontologies, and that such reuse will indeed facilitate semantic interoperation between applications based on these ontologies, there has not been enough experience with this approach to claim it as a success.

Another set of approaches for discovery mapping includes heuristics-based or machine learning techniques that use various characteristics of ontologies, such as their structure, instances of classes, and definitions of concepts, to find mappings [89]. Examples of this kind of work include the techniques described by Hovy [90], the PROMPT algorithms of Musen and Noy [91], FCA-Merge [92], IF-Map [93], GLUE [94], and the algorithms for complex mappings of Giunchiglia and Shvaiko [95].

As part of their efforts in the Medical Ontology Research project at the NLM, Zhang, Bodenreider, et al. have developed methods for aligning the UMLS with general ontologies such as Cyc and WordNet and also with specialized ontologies such as the Gene Ontology. In addition, they have also tested methods for aligning UMLS knowledge sources (e.g., the Metathesaurus with the Semantic Network) and biomedical ontologies outside the UMLS. In their work on aligning the FMA and the anatomy content of GALEN, they used a four-step method comprised of acquiring terms, identifying anchors (shared concepts) lexically, acquiring semantic relations, and identifying anchors structurally. The work represents an effort to exploit implicit and explicit domain knowledge to uncover similar and conflicting relations. A by-product of their work was the discovery of a number of inconsistencies in both ontologies [96,97].

The ONIONS (Ontologic Integration of Naïve Sources) approach to merging, developed at Consiglio Nazionale delle Ricerche (CNR) in Italy, has been applied to the medical domain to create the ON.9.2 integration ontology, which unifies systems like GALEN and the UMLS. Gangemi et al. have described their experience using this approach. They report that they were largely successful in achieving several intended outcomes, but that an unavoidable bottleneck in their approach was the necessity of extensive human intervention in the search, choice, and formalization of generic ontologies [98].

6. Ontology maintenance

Ontologies inevitably have to evolve, whether because improvements have to be made to the ontology itself, or because the world has changed and our representations of the world have to reflect what is new. A number of workers have described the problems they have encountered in managing ontologies, as well as the approaches they have used to manage changes.

Cimino described his experience coping with the annual updates to the ICD-9-CM terminology [23]. The Medical Entities Dictionary had mappings to the ICD-9-CM terminology, and every time the ICD-9-CM terminology changed, the maintainers of the MED had to analyze and properly handle the changes so that the mappings would remain valid. Cimino created a formal taxonomy of changes in terminologies that included possible reasons (good, as well as bad) for the changes. Corresponding to these changes were adaptive mechanisms for properly handling the changes in the MED. Subsequently, Oliver, as part of her dissertation work, proposed a formal methodology for change management of local and shared controlled medical terminologies. The approach centered on a formal representation of medical concepts similar to those used in frame-based knowledge representation systems. This formal representation allowed Oliver to describe highly detailed and formal operations to carry out the types of changes that Cimino had earlier described [25].

In work that eventually transformed the design and maintenance workflow of what is now SNOMED CT, Campbell demonstrated new methods to support an evolutionary approach to controlled medical terminology development. In the system that he created, multiple authors were allowed to independently define terms, and then partially rely on the system to detect and manage conflicts in the definitions. Conflict detection depended upon logically-based definitions of terms, and a description-logic classifier detected conflicting definitions based on semantic equivalence rather than syntactic equivalence. Furthermore, the configuration management methods he developed relied on "change sets" that contained information on changes that had been made by authors. These change sets were used to support terminology verification and automated migration [99,100].

Noy and Musen have developed the PROMPT set of tools that work with the Protégé ontology editor. One of the PROMPT tools handles semi-automated detection and handling of changes in ontologies. One of the outputs of the tool is a *structural diff* (analogous to the result of the "diff" UNIX program) that represents the structural differences between two versions of the same ontology. PROMPT also includes PROMPTDIFF, which is a set of heuristic algorithms that attempt to detect matches between concepts in different versions, as well as a user interface that helps human editors evaluate the results of PROMPT and make their own final decisions [91].

7. Ontology evaluation

Ontology evaluation can roughly be divided into two kinds: technical (carried out by developers) and users' evaluation [16]. While most current evaluation methods clearly fall into the first category, recent efforts have elaborated the need for and suggested possible approaches for formalizing the second kind of evaluation.

Cimino has compiled a list of desiderata for controlled medical terminologies [21]. Foremost among these desiderata were the adherence to a concept orientation and the assurance of adequate domain coverage. Although Cimino did not elaborate on how to implement many of the desiderata, adherence to many of them can be seen in the current generation of knowledge-based terminologies and ontologies.

The OntoClean methodology stands out as one of the most explicit and formal methods for evaluating ontologies [87,101,102]. Its focus is on "cleaning up" taxonomies through a systematic and rigorous examination of the meta-properties of concepts. As such, its goal is to remove erroneous *subclass-Of* (*is-a*) relations in taxonomies. In a series of papers, Guarino and Welty have described philosophical notions, such as *rigidity*, *identity*, and *unity*, and how these are applicable to the analysis of concepts. Building on these notions, they have axiomatized a set of rules that can be systematically applied to taxonomies so that many errors are corrected, and the result of the application is a "cleaned" ontology (Fig. 12). For example, based on the axiom that rigid concepts cannot be subsumed by anti-rigid concepts, the concept *human* (rigid because all instances of *human* are necessarily so) cannot be subsumed by the concept *student* (anti-rigid because all instances of *student* are not necessarily so); the concept *student* should rather be instantiated as a "role" that can be taken by an instance of *human*. Although various formal problems with the method have been detected [103–105], Guarino and Welty's work is also notable because it is an example of efforts by computer scientists to use the methods of philosophical ontology to help solve some of the problems that persist in spite of (or in some cases, were created by) the methods previously used. Spackman and Reynoso studied the usefulness of OntoClean in evaluating some of the decisions of the SNOMED CT curators [106]. They concluded that while OntoClean was useful in making distinctions understandable and reproducible, some of the distinctions were not necessarily useful for electronic health records or decision support, and that, in general, methodologies based on philosophy needed to be more transparent so

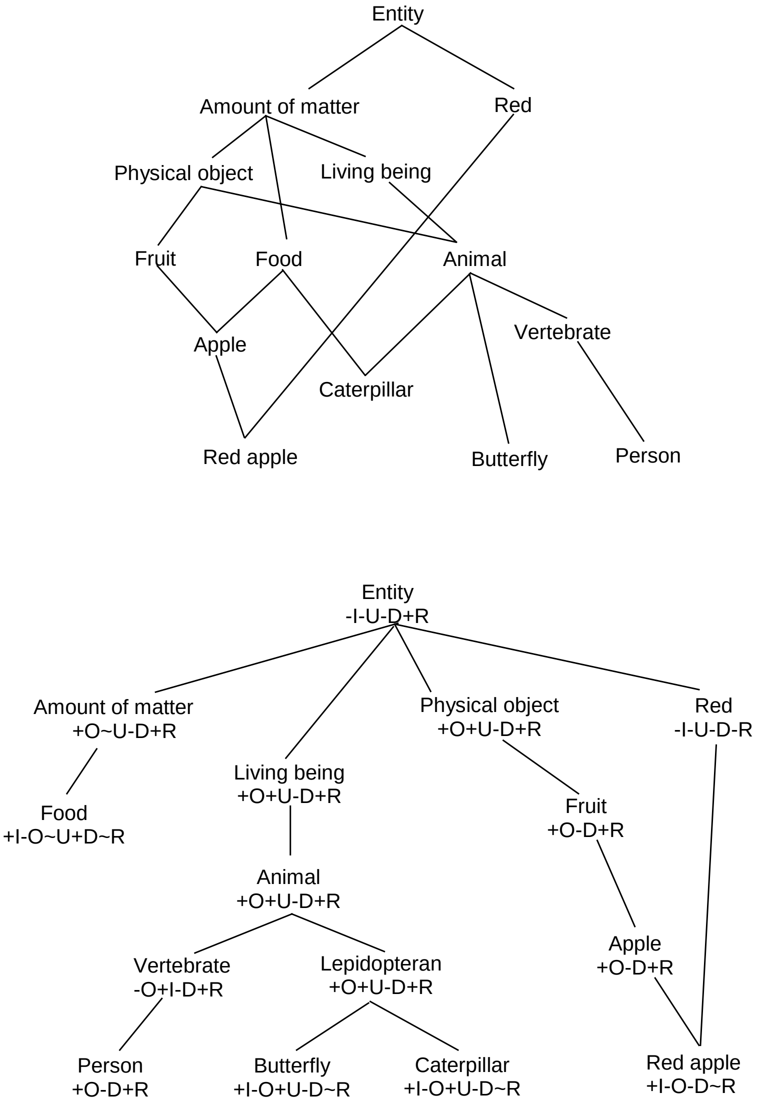

Fig. 12. An example of "before" (A) and "after" (B) snapshots of a taxonomy that has been cleaned following the use of the OntoClean method. Lines represent *is-a* relationships. For example, the incorrect subsumption relationship between Living being and Amount of matter (a result of confusing *constitution* and *subsumption*) is removed and Living being is subsumed directly by Entity in the cleaned taxonomy. The letters I, U, D, and R stand for the *identity*, *unity*, *dependence*, and *rigidity* metaproperties, respectively. Adapted from [87].

that domain experts (such as medical practitioners) could more readily use the methods.

Lastly, Noy has proposed some ideas towards the creation of a public system that allows ontology consumers to rate ontologies and share them among the community. The idea is largely based on existing web-based systems that publish information about products and allow consumers to offer their opinions. Similarly, an ontology can have an abstract or a summary, which might include information on what the ontology covers and what its most important concepts are. Evaluation results based on formal methods (such as OntoClean and logical consistency checks) can also be incorporated into users' ratings. Finally, consumers can also offer their opinions and descriptions of their experience using the ontologies [107,108]. An implemented system based on this approach would have the potential to facilitate the dissemination of formal evaluation results as well as complementary information that might be useful to ontology consumers seeking ontologies suited to their needs. However, a potential problem is that the system's usefulness might decline if constraints are not put in place to prevent or correct for low-quality evaluations. Noy et al. have suggested the establishment of webs of trusted users as a possible solution to this problem [108].

8. Conclusions

Biomedical ontologies are key pieces in the further development of informatics applications in several areas, such as knowledge-based decision support, terminology management, and systems interoperability and integration. A significant body of work now exists that report on experiences with various approaches in important problem areas of research on ontologies. Most researchers have focused on issues of design, but interest is increasingly turning to other pressing problems such as how to properly evaluate ontologies. In presenting the various methods in this paper, we have touched upon philosophical as well as engineering concerns that should be considered in endeavors of this kind, in order that we may see the widespread creation of rigorous, useful ontologies. Philosophical ontology has much to offer in terms of formal analytical methods towards creating declarative representations of knowledge that are general, reusable, and valid. At the same time, we need to also draw upon the insights and approaches that have developed within the engineering community, particularly those that have exposed and attempted to address practical problems that continue to dog both users and developers of ontologies.

Acknowledgments

The author thank James Cimino and an anonymous reviewer for their valuable feedback. This work was supported by a training Grant (LM07079) from the National Library of Medicine, Bethesda, Maryland.

References

[1] Rector AL. Clinical terminology: why is it so hard? Methods Inf Med 1999;38(4–5):239–52.
[2] Cornet R, Abu-Hanna A. Using description logics for managing medical terminologies. In: Dojat M, Keravnou E, Barahona P, editors. Artificial intelligence in medicine. Proceedings of the ninth conference on artificial intelligence in medicine in Europe. AIME; 2003. p. 61–70.
[3] Rosse C, Mejino Jr JL. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. J Biomed Inform 2003;36(6):478–500.
[4] Musen MA. Medical informatics: searching for underlying components. Methods Inf Med 2002;41(1):12–9.
[5] Cimino JJ. From data to knowledge through concept-oriented terminologies: experience with the Medical Entities Dictionary. J Am Med Inform Assoc 2000;7(3):288–97.
[6] Rector A. Medical informatics. In: Baader F, Calvanese D, McGuinness D, Nardi D, Patel-Schneider P, editors. The description logic handbook. Cambridge, UK; New York: Cambridge University Press; 2003. p. 406–26.
[7] Consortium GO. Creating the gene ontology resource: design and implementation. Genome Res 2001;11(8):1425–33.
[8] Smith B, Ceusters W, Klagges B, Kohler J, Kumar A, Lomax J, et al. Relations in biomedical ontologies. Genome Biol 2005;6(5):R46.
[9] Bard J. Ontologies: formalising biological knowledge for bioinformatics. Bioessays 2003;25(5):501–6.
[10] Blois MS. Medicine and the nature of vertical reasoning. N Engl J Med 1988;318(13):847–51.
[11] Sowa JF. Knowledge representation: logical, philosophical, and computational foundations. Belmont, CA, USA: Brooks-Cole; 2000.
[12] Smith B. Ontology: philosophical and computational. In: Floridi L, editor. The blackwell guide to the philosophy of computing and information. Oxford: Blackwell Publishers; 2003.
[13] Schneider L. Foundational ontologies and the realist bias. In: KI2003 Workshop on reference ontologies and application ontologies. Hamburg, Germany; 2003.
[14] Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A. Ontology library: Wonderweb Deliverable D18. 2003.
[15] Grenon P. BFO in a nutshell: a bi-categorical axiomatization of BFO and comparison with DOLCE: IFOMIS; 2003.
[16] Gomez-Perez A, Fernandez-Lopez M, Corcho O. Ontological engineering. London: Springer-Verlag; 2004.
[17] Gruber TR. A translation approach to portable ontologies. Knowl Acquis 1993;5(2):199–220.
[18] Smith B, Welty C. Ontology: towards a new synthesis. In: FOIS '01; 2001. Ogunquit (Maine, USA): ACM; 2001.
[19] Guarino N. Formal ontology and information systems. In: Guarino N, editor. Proceedings of the first international conference on formal ontologies in information systems. FOIS'98; 1998. Trento, Italy: IOS Press; 1998. p. 3–15.
[20] Smith B, Williams J, Schulze-Kremer S. The ontology of the gene ontology. AMIA Annu Symp Proc 2003:609–13.
[21] Cimino JJ. Desiderata for controlled medical vocabularies in the twenty-first century. Methods Inf Med 1998;37(4–5):394–403.
[22] Pisanelli DM. If ontology is the solution, what is the problem? In: Pisanelli DM, editor. Ontologies in medicine. Amsterdam (Burke, VA): IOS Press; 2004. p. 1–19.
[23] Cimino JJ. Formal descriptions and adaptive mechanisms for changes in controlled medical vocabularies. Methods Inf Med 1996;35(3):202–10.
[24] Rector A. Modularisation of domain ontologies implemented in description logics and related formalisms including OWL. In: Gennari J, editor. K-CAP; 2003. 2003. p. 121–9.
[25] Oliver DE, Shahar Y. Change management of shared and local versions of health-care terminologies. Methods Inf Med 2000;39(4–5):278–90.
[26] Spackman KA, Campbell KE, Cote RA. SNOMED RT: a reference terminology for health care. Proc AMIA Annu Fall Symp 1997:640–4.
[27] Rector AL, Nowlan WA. The GALEN project. Comput Methods Programs Biomed 1994;45(1–2):75–8.
[28] Rickard KL, Mejino Jr JL, Martin RF, Agoncillo AV, Rosse C. Problems and solutions with integrating terminologies into evolving knowledge bases. Medinfo 2004;11(Pt 1):420–4.
[29] Ceusters W, Smith B, Flanagan J. Ontology and medical terminology: why description logics are not enough. In: TEPR; 2003. San Antonio; 2003.
[30] Bodenreider O, Smith B, Kumar A, Burgun A. Investigating subsumption in DL-based terminologies: a case study in SNOMED CT. In: Hahn U, editor. KR-MED 2004; 2004. Whistler, Canada: AMIA; 2004. p. 12–20.
[31] HL7. HL7 Reference Information Model (http://www.hl7.org/library/data-model/RIM/modelpage_mem.htm).
[32] Vizenor L, Smith B, Ceusters W. Foundation for the electronic health record: an ontological analysis of the HL7's reference information model. Unpublished manuscript.
[33] Kumar A, Yip L, Smith B, Grenon P. Bridging the gap between medical and bioinformatics using formal ontological principles (forthcoming). Comput Biol Med 2005.
[34] Pryor TA. The HELP medical record system. MD Comput 1988;5(5):22–33.
[35] Musen MA. Clinical decision-support systems. In: Shortliffe E, Perreault L, Wiederhold G, Fagan L, editors. Medical informatics: computer applications in health care and biomedicine. 2nd ed. New York: Springer-Verlag; 2001. p. 573–609.
[36] Musen MA. Modern architectures for intelligent systems: reusable ontologies and problem-solving methods. Proc AMIA Symp 1998:46–52.
[37] Smith B, Rosse C. The role of foundational relations in the alignment of biomedical ontologies. Medinfo 2004;11(Pt 1):444–8.
[38] Genesereth M, Fikes R. Knowledge interchange format. Version 3.0. Reference manual. California: Computer Science Department, Stanford University; 1992.
[39] Gruber T. Ontolingua: A mechanism to support portable ontologies. Stanford University, Stanford: California Knowledge Systems Laboratory; 1992.

[40] MacGregor R. Inside the LOOM classifier. SIGART Bull 1999;2(3):70–6.

[41] Quillian M. Word concepts: A theory and simulation of some basic capabilities. Behav Sci 1967;12:410–30.

[42] Minsky M. A framework for representing knowledge. In: Haugeland J, editor. Mind Design: MIT Press; 1981.

[43] Lassila O, R S. Resource Description Framework (RDF) Model and Syntax Specification: W3C Recommendation. http://www.w3.org/TR/REC-rdf-syntax; 1999.

[44] Brickley D, Guha R. RDF Vocabulary Description Language 1.0: RDF Schema: W3C Working Draft. http://www.w3.org/TR/PR-rdf-schema; 2003.

[45] Horrocks I, van Harmelen F. Reference Description of the DAML + OIL (March 2001) Ontology Markup Language: http://www.daml.org/2001/03/reference.html; 2001.

[46] McGuinness DL, Harmelen Fv. OWL Web Ontology Language Overview. http://www.w3.org/TR/owl-features/. 2004.

[47] Baader F, Calvanese D, McGuinness D, Nardi D, Patel-Schneider P. The description logic handbook: theory, implementation, and applications. Cambridge (UK, New York): Cambridge University Press; 2003.

[48] Knublauch H, Fergerson RW, Noy NF, Musen MA. The Protégé OWL Plugin: an open development environment for semantic web applications. In: Third international semantic web conference; 2004. Hiroshima, Japan; 2004.

[49] Gennari J, Musen MA, Fergerson RW, Grosso WE, Crubézy M, Eriksson H, et al. The evolution of Protégé: an environment for knowledge-based systems development; 2002.

[50] Rector A. Defaults, context, and knowledge: alternatives for OWL-indexed knowledge bases. Pac Symp Biocomput 2004:226–37.

[51] Smith B. Beyond concepts, or: ontology as reality representation. In: Varzi A, Vieu L, editors. Formal ontology and information systems. Proceedings of the third international conference (FOIS 2004); 2004. Amsterdam: IO Press; 2004. p. 73–84.

[52] Lenat DB, Guha RV. Building large knowledge-based systems: representation and inference in the Cyc project. Boston, Massachusetts: Addison-Wesley; 1990.

[53] Uschold M, Gruninger M. Ontologies: principles, methods, and applications. Knowl Eng Rev 1996;11(2):93–115.

[54] Gruninger M, Fox MS. Methodology for the design and evaluation of ontologies. In: Skuce D, editor. IJCAI96 Workshop on basic ontological issues in knowledge sharing; 1995; 1995. p. 6.1–.10.

[55] Bernaras A, Laresgoiti I, Corera J. Building and reusing ontologies for electrical network applications. In: Wahlster W, editor. European conference on artificial intelligence (ECAI'96); 1996; Budapest, Hungary: John Wiley and Sons; 1996. p. 298–302.

[56] Gomez-Perez A, Fernandez-Lopez M, de Vicente A. Towards a method to conceptualize domain ontologies. In: van der Vet P, editor. ECAI'96 Workshop on ontological engineering; 1996; Budapest, Hungary; 1996. p. 41–52.

[57] Swartout W, Ramesh P, Knight K, Russ T. Toward distributed use of large-scale ontologies. In: Farquhar A, Gruninger M, Gomez-Perez A, Uschold M, van der Vet P, editors. AAAI'97 Spring symposium on ontological engineering; 1997; Stanford University, California; 1997. p. 138–48.

[58] Staab S, Schnuur H, Studer R, Sure Y. Knowledge processes and ontologies. IEEE Intell Syst 2001;16(1):26–34.

[59] Maintainers. Unified Medical Language System. (http://www.nlm.nih.gov/research/umls/about_umls.html).

[60] Ceusters W, Smith B, Fielding J. LinkSuite: formally robust ontology-based data and information integration. In: Proceedings of DILS 2004 (Database Integration in the Life Sciences). Berlin: Springer; 2004. p. 124–39.

[61] Fielding J, Simon J, Ceusters W, Smith B. Ontological theory for ontological engineering: biomedical systems information integration. In: Proceedings of the ninth international conference on the principles of knowledge representation and reasoning (KR2004); 2004 June; Whistler, BC; 2004.

[62] Grenon P, Smith B, Goldberg L. Biodynamic ontology: applying BFO in the biomedical domain. In: Ontologies in medicine. Amsterdam: IOS Press; 2004. p. 20–38.

[63] Noy N, Musen M, Mejino J, Rosse C. Pushing the envelope: challenges in a frame-based representation of human anatomy. Data Knowl Eng J 2002;48(3):335–59.

[64] Dameron O, Rubin D, Musen MA. Challenges in converting a frame-based ontology into OWL: the foundational model of anatomy case-study. In: American medical informatics; 2005; Washington, DC; 2005.

[65] Rogers J, Rector A. The GALEN ontology. In: Medical informatics Europe 1996; 1996; Amsterdam: IOS Press; 1996. p. 174–178.

[66] Cimino JJ, Clayton PD, Hripcsak G, Johnson SB. Knowledge-based approaches to the maintenance of a large controlled medical terminology. J Am Med Inform Assoc 1994;1(1):35–50.

[67] Hartel FW, de Coronado S, Dionne R, Fragoso G, Golbeck J. Modeling a description logic vocabulary for cancer research. J Biomed Inform 2005;38(2):114–29.

[68] Golbeck J, Fragoso G, Hartel FW, Hendler J, Parsia B, Oberthaler J. The national cancer institute's thesaurus and ontology. J Web Semantics 2003;1(1).

[69] Ceusters W, Smith B, Goldberg L. A terminological and ontological analysis of the NCI Thesaurus (in press). In: Methods of information in medicine; 2005.

[70] Kumar A, Smith B. Oncology ontology in the NCI Thesaurus. In: Artif Intell Med Europe (AIME 2005). NY: Aberdeen; 2005.

[71] Harris MA, Clark J, Ireland A, Lomax J, Ashburner M, Foulger R, et al. The gene ontology (GO) database and informatics resource. Nucleic Acids Res 2004;32:D258–61 (Database issue).

[72] Wroe CJ, Stevens R, Goble CA, Ashburner M. A methodology to migrate the gene ontology to a description logic environment using DAML + OIL. Pac Symp Biocomput 2003:624–35.

[73] Bodenreider O, Mitchell JA, McCray AT. Evaluation of the UMLS as a terminology and knowledge resource for biomedical informatics. Proc AMIA Symp 2002:61–5.

[74] Boxwala AA, Zeng QT, Chamberas A, Sato L, Dierks M. Coverage of patient safety terms in the UMLS metathesaurus. AMIA Annu Symp Proc 2003:110–4.

[75] Cimino JJ, Elhanan G, Zeng Q. Supporting infobuttons with terminological knowledge. Proc AMIA Annu Fall Symp 1997:528–32.

[76] Friedman C, Liu H, Shagina L, Johnson S, Hripcsak G. Evaluating the UMLS as a source of lexical knowledge for medical language processing. Proc AMIA Symp 2001:189–93.

[77] Burgun A, Bodenreider O. Mapping the UMLS semantic network into general ontologies. Proc AMIA Symp 2001:81–5.

[78] Sculze-Kremer S, Smith B, Kumar A. Revising the UMLS Semantic Network. In: Fieschi M, editor. MEdinfo 2004; 2004; San Francisco: IOS Press; 2004. p. 1700.

[79] Chute CG, Cohn SP, Campbell KE, Oliver DE, Campbell JR. The content coverage of clinical classifications. For the computer-based patient record institute's work group on codes & structures. J Am Med Inform Assoc 1996;3(3):224–33.

[80] Spackman K. SNOMED RT and SNOMEDCT. Promise of an international clinical terminology. MD Comput 2000;17(6):29.

[81] Fung KW, Hole WT, Nelson SJ, Srinivasan S, Powell T, Roth L. Integrating SNOMED CT into the UMLS: an exploration of different views of synonymy and quality of editing. J Am Med Inform Assoc 2005;12(4):486–94.

[82] Spackman KA. Normal forms for description logic expressions of clinical concepts in SNOMED RT. Proc AMIA Symp 2001:627–31.

[83] Hahn U, Schulz S, Romacker M. Partonomic reasoning as taxonomic reasoning in medicine. In: AAAI'99—Proceedings of the 16th national conference on artificial intelligence & and the 11th innovative applications of artificial intelligence conference; 1999; Orlando, Florida: AAAI Press & MIT Press; 1999. p. 271–76.

[84] Hahn U, Schulz S, Romacker M. Part-whole reasoning: a case study in medical ontology engineering. IEEE Intell Syst App 1999;14(5):59–67.

[85] Schulz S, Hahn U. Mereotopological reasoning about parts and (w)holes in bio-ontologies. In: Smith CWB, editor. Formal ontology in information systems. Collected papers from the second international conference; 2001; Ogunquit (Maine, USA/New York, NY): ACM Press; 2001. p. 210–21.

[86] Bittner T, Smith B. A theory of granular partitions. In: Duckham M, Goodchild M, Worboys M, editors. Foundations of geographic information science. London: Taylor & Francis Books; 2003. p. 117–51.

[87] Guarino N, Welty C. An overview of OntoClean. In: Staab S, Studer R, editors. The handbook on ontologies. Berlin: Springer-Verlag; 2004. p. 151–72.

[88] Kumar A, Smith B. Ontology for task-based clinical guidelines and the theory of granular partitions. In: Artif Intell Med. Cyprus: Springer-Verlag; 2003. p. 71–5.

[89] Noy NF. Semantic integration: a survery of ontology-based approaches. SIGMOD Record, Special Issue on Semantic Integration 2004;33(4).

[90] Hovy E. Combining and standardizing largescale, practical ontologies for machine translation and other uses. In: The first international conference on language resources and evaluation (LREC); 1998; Granada, Spain; 1998. p. 535–42.

[91] Noy NF, Musen MA. The PROMPT suite: interactive tools for ontology merging and mapping. Int J Hum-Comput Stud 2003;59(6):983–1024.

[92] Stumme G, Maadche A. FCA-Merge: bottom-up merging of ontologies. In: Proceedings of the seventh international conference on artificial intelligence (IJCAI '01); 2001; Seattle, Washington; 2001. p. 225–30.

[93] Kalfoglou Y, Schorlemmer M. IF-Map: an ontology mapping method based on information flow theory. J Data Semantics 2003;1(1):98–127.

[94] Doan A, Madhavan J, Domingos P, Halevy A. Learning to map between ontologies on the semantic web. In: Proceedings of the 11th international conference on WWW conference; 2002; Hawaii, US; 2002.

[95] Giunchiglia F, Shvaiko P, Yatskevich M. Semantic matching. In: Proceedings of the first European semantic web symposium (ESWS '04); 2004; Greece; 2004.

[96] Zhang S, Bodenreider O. Aligning representations of anatomy using lexical and structural methods. AMIA Annu Symp Proc 2003:753–7.

[97] Zhang S, Bodenreider O. Knowledge augmentation for aligning ontologies: an evaluation in the biomedical domain. In: Semantic integration workshop at the second international semantic web conference (ISWC 2003); 2003; 2003. p. 109–14.

[98] Gangemi A, Pisanelli DM, Steve G. An overview of the ONIONS project: Applying ontologies to the integration of medical terminologies. Data Knowl Eng 1999;31.

[99] Campbell KE, Cohn SP, Chute CG, Shortliffe EH, Rennels G. Scalable methodologies for distributed development of logic-based convergent medical terminology. Methods Inf Med 1998;37(4–):426–39.

[100] Campbell KE, Das AK, Musen MA. A logical foundation for representation of clinical data. J Am Med Inform Assoc 1994;1(3):218–32.

[101] Guarino N, Welty C. Evaluating ontological decisions with OntoClean. Commun ACM 2002;45(2):61–5.

[102] Welty C, Guarino N. Supporting ontological analysis of taxonomic relationships. Data Knowl Eng 2001;39:51–74.

[103] Kaplan A. Towards a consistent logical framework for ontological analysis. In: Formal ontology in information systems 2001; 2001 October, 2001; Ogunquit, Maine; 2001.

[104] Andersen W, Menzel C. Modal rigidity in the OntoClean methodology. In: Formal Ontology in Information Systems. Italy: Torino; 2004.

[105] Carrara M, Giaretta P, Morato V, Soavi M, Spolaore G. Identity and Modality in OntoClean. In: Formal Ontology in Information Systems. Italy: Torino; 2004.

[106] Spackman K, Reynoso G. Examining SNOMED from the perspective of formal ontological principles. In: Hahn J, editor. KR-MED 2004 Proceedings. Whistler, Canada: AMIA; 2004. p. 81–7.

[107] Sure Y, Gomez-Perez A, Daelemans W, Reinberger M, Guarino N, Noy N. Why Evaluate Ontology Technologies? Because It Works! IEEE Intelligent Systems 2004;19(4):74–81.

[108] Noy NF, Guha RV, Musen, MA. User ratings of ontologies: who will rate the raters? In: Proceedings of the AAAI 2005 spring symposium on knowledge collection from volunteer contributors, Stanford, CA.