Chapter 2
Ontological Architectures

Leo Obrst

2.1 Introduction

We distinguish between *ontological architecture* and *ontology architecture*, though they are closely related. *Ontology architecture* is emerging as a distinct discipline in ontology engineering – as an ontology development and deployment structure and methodology (Fernandéz et al., 1997). It necessarily also includes aspects of what is sometimes termed *ontology lifecycle management* (Andersen et al., 2006). In fact, ontology architecture can be considered to encompass ontology lifecycle management because the former lays out a general framework for the development, deployment, and maintenance of ontologies (which is the focus of lifecycle management), but also includes the interaction of applications and services that use ontologies, and an ontology tool and service infrastructure to support these. *Ontological architecture* is the architecture that is used to structure the ontologies that are employed by ontology architecture. As such, it addresses the levels of ontologies required (foundational, upper, middle, utility, reference, domain, and sub-domain ontologies), and mathematical, logical, and engineering constructs used to modularize ontologies in a large *ontological space.* This chapter focuses on *ontological architecture,* but it must be understood to underpin *ontology architecture* if only to ground/situate and enable the latter. Both kinds of architecture are relevant to ontology engineering, but we cannot address ontology architecture here until the very last section, when we look ahead. Instead, we focus on ontological architecture, which as it turns out, is a large enough topic.

The chapter is structured as follows. In Section 2.2, we distinquish *ontological architecture* from *ontology architecture*, provide some understanding of their respective rationales, how ontologies are distinct from but impinge on elements of epistemology, the formal semantics of language, and conceptual models. We depict the *ontology spectrum* (Obrst, 2002–2003), which constitutes a range of semantic models of increasing expressiveness, and define these. The more expressive

L. Obrst (✉)
The MITRE Corporation, McLean, VA, USA
e-mail: lobrst@mitre.org

R. Poli et al. (eds.), *Theory and Applications of Ontology: Computer Applications*,
DOI 10.1007/978-90-481-8847-5_2,

models enable more complex applications. We also show how one aspect of ontological architecture, the expressiveness of the knowledge/ontology representation language, and ontology application are related. In Section 2.3, ontological architecture is described, by detailing the use of upper, middle, and domain ontologies to address semantic interoperability. We extend this with a discussion of additional structure that has been proposed, and some foundational ontological distinctions. Section 2.4 is the core of the chapter. It discusses some ways of structuring the *ontological space*, which really is itself embedded in a *logical space,* and necessarily must also address *meta-ontological architectural* issues. Notions of ontological modularity are examined, including that of formalized contexts such as *microtheories,* which originated from the Cyc effort (Blair et al., 1992), the approach called *the lattice of theories*, most recently characterized by John Sowa (2005), additional approaches based on logical ways of characterizing mathematical *little theories* (Farmer et al., 1992) which yet must interoperate, recent research in ontology modularity, and Robert Kent's (2004, 2006) meta-ontology called *the Information Flow Framework*, based on Barwise and Seligman's (1997) *Information Flow Theory*, itself an application of Category Theory (Mac Lane, 1971; Bar and Wells, 1999), and similar work at the meta-ontological level. Finally, in Section 2.5, we conclude with a vision of the future for both ontological and ontology architecture.

Ontological architecture spans many topics. We can only briefly sketch its components in this chapter.

2.2 Ontological and Ontology Architecture: Overview

Ontology architecture addresses content (how better ontologies are developed), the apparatus needed to develop, deploy, and maintain ontologies (which tools, requirements, methodologies, lifecycle support, policy, and governance are required), and ontology application interaction (how data is stored, accessed, and linked to ontology instances and facts; which services do ontology require for and provide to applications). We do not address this ontology architecture per se in this chapter, since our interests are more fundamental. Instead, we focus on *ontological architecture*, the foundational architecture which must underpin subsequent ontology application notions we characterize as ontology architecture. This section provides an overview of what ontological architecture addresses.

2.2.1 Truth and Belief: Ontology, Epistemology, Contextual Semantics, Language, and Applications

Ontology is many things to many people, as the other chapters of these volumes demonstrate, and so no time is spent here defining ontology. This chapter focuses on architecture. One issue, however, needs to be raised: what ontology does not address particularly must still be addressed by ontology architecture. Ontology is not epistemology, nor is it the semantics of natural language. But aspects of these must be addressed by an account of ontology architecture. Why *epistemology*? Because, though ontology is about the real entities, relations, and properties of the word, epistemology is about the perceived and belief-attributed entities, relations, and properties of the world, empirical evidence gleaned that will be described or characterized by ontology. Why *natural language semantics*? Because, though ontology is about the real entities, relations, and properties of the world, natural language semantics is about the rendition in language of interpretations about the entities, relations, and properties of the world, and includes notions of sense and reference. In ontology architecture, epistemology is employed in the use and qualification of data and as stored in databases or tagged or indexed in documents. In ontology architecture, natural language semantics is employed in the analysis of natural language descriptions used to ascertain and represent the real world entities of ontology, the naming conventions used and the access to the interpretations about the real world that the ontology represents. One natural language processing technology in particular, information extraction, crucially depends on natural language semantics -- information extraction addressing the identification of entities, relations, and events in unstructured text, and the tagging or extraction of these to form instances of ontology concepts.

2.2.2 The Big Picture

Figure 2.1[1] is a graphical rendition of ontological architecture and its components. These components will be described in more detail in Section 2.3.

The important point about this diagram is the layers: the upper, mid-level, and domain (or lower) ontology layers. Sometimes the upper and mid-level ontologies are called foundational ontologies, and there can be multiple such in each layer. We eschew the notion of a monolithic ontology, at least for engineering purposes, and instead view foundational ontologies similar to domain ontologies: as coherent and consistent theories, hence our use of the terminology introduced in the next section, i.e., ontologies as *logical theories*. In our view, upper ontologies are most abstract and make assertions about constructs such as identity criteria, parts/wholes, substance and constitution, space and time (and endurants and perdurants), necessary properties, dynamic properties, attributes spaces, etc., that apply to all lower levels; hence, they span all mid-level and domain ontologies. Upper ontologies themselves may consist of levels, as the extended discussions on *levels of reality* make clear (Poli, 2003; Poli, 2010, this volume; Poli and Obrst, 2010, this volume). Mid-level ontologies are less abstract and make assertions that span multiple domain ontologies. The characterization of these are less clear, as is the demarcation point between upper and mid-level. Some examples of constructs potentially in

[1] See Semy et al. (2004, p. 8), also Chapter 1 by Poli and Obrst, this volume.

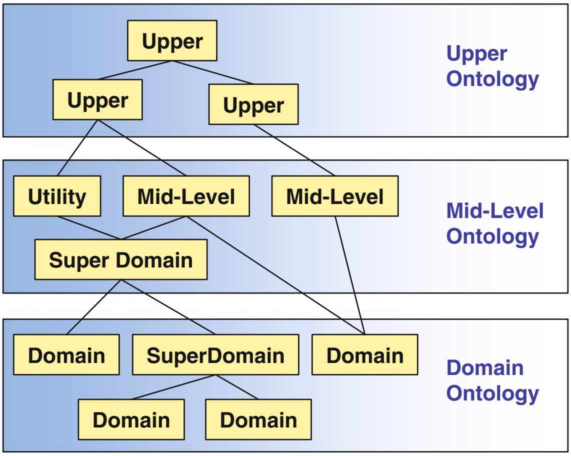

Fig. 2.1 Ontology architecture: the big picture

a mid-level ontology are *humanOrganization* and *industrialProcess*. These are not necessarily represented in an upper ontology, but may be; probably, however, they are in a mid-level ontology (a biomedical mid-level ontology may not have these as concepts; however, a manufacturing mid-level ontology would have them).

2.2.3 The Ontology Spectrum

Ontology architecture, just like ontology, is a notion that must be learned and incorporated gradually over time by an enterprise or community. It's possible, though rare, that an enterprise or community is sufficiently knowledgeable about ontology and its use in semantically informing applications; it's equally rare that the organization is aware of the costs, hazards, cost-effective benefits, and preferred methods of using ontology. To assist organizations (enterprises, communities) in determining the range of semantic models that constitute stages of embracing semantic technologies of which the highest model is ontologies, we have created the Ontology Spectrum, depicted in Fig. 2.2 (Obrst, 2002; Obrst, 2003; Daconta, Obrst, Smith, 2003). The Ontology Spectrum is both a description of the range of semantic models, as they increase in expressivity and complexity of structure, and an indication of the migration path that organizations can take as the expressiveness of their semantic models needs to increase in order to accommodate their richer problems and hence, solutions. The lower end of the Ontology Spectrum is in fact not about ontologies at all.

What is colloquially, though incorrectly, known as an *ontology* can range from the simple notion of a *Taxonomy* (terms[2] or concepts with minimal hierarchic or

[2]We differentiate and define *term* and *concept* in Table 2.1, below.

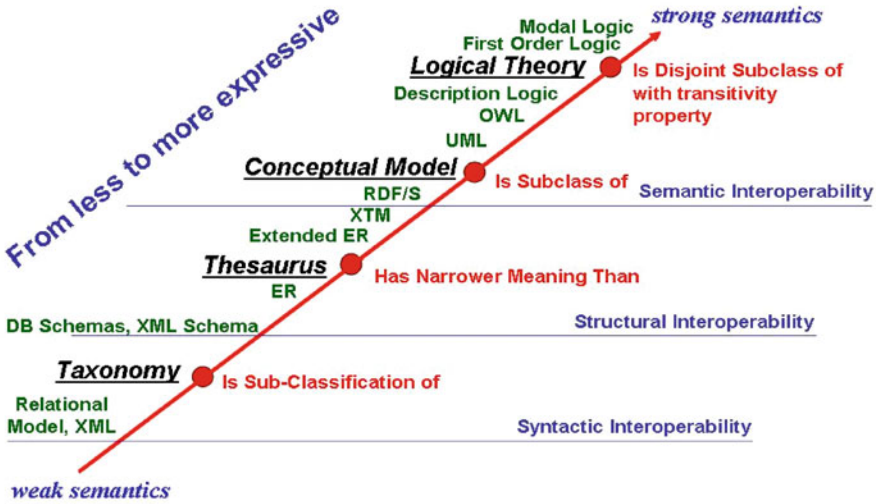

Fig. 2.2 The ontology spectrum

parent/child structure), to a *Thesaurus* (terms, synonyms, broader than/narrower than term taxonomies, association relation), to a *Conceptual Model* (concepts structured in a subclass hierarchy, generalized relations, properties, attributes, instances), to a *Logical Theory* (elements of a Conceptual Model focusing however on real world semantics and extended with axioms and rules, also represented in a logical KR language enabling machine semantic interpretation). Terms and concepts are differentiated in Table 2.1. We also differentiate between weak and strong taxonomies: The *subclassification of* relation characterizes *weak taxonomies*. *Strong taxonomies* are characterized by either the *subclass of* relation for concepts (which typically can be considered universal categories for referents[3]) or the *narrower*

[3]There is, however, a vast literature on the notion of *concept* in philosophy, cognitive science/psychology, and linguistics. Often in cognitive science/psychology, a concept is considered to be mental particular, having a structured mental representation of a certain type (Margolis and Laurence, 1999b, p. 5–6). However, in philosophy, a concept is considered an abstract entity, signifying a general characterizing idea or universal which acts as a category for *instances* (*individuals* in logic, *particulars* in metaphysics and philosophical ontology) (Smith, 2004). Even in the philosophical literature, the notion of concept will vary according to philosophical stance, i.e., according to whether the adherent to the particular notion is an idealist, nominalist, conceptualist, or realist, or some combination or refraction of those (Poli, 2010). For example, some will consider a concept to be simply a placeholder for a real world entity, either a universal or a particular; example: *Joe Montana* (a former USA football quarterback) or *Winston Churchill* (a former UK prime minister) as concepts. That is, the mental placeholder or idea can be about anything. This notion of concept is a surrogate for anything that a philosophical or many linguistic theories may opine. Often, therefore (and this is our view here), concepts are best understood as *conceptions*, a term which has perhaps less technical baggage, insofar as *conception* emphasizes that we are talking about a mental representation which may or may not be reified as a *concept*, perhaps a stronger notion. But

Table 2.1 Term vs. concept

Terms (terminology): Natural language words or phrases that act as indices to the underlying meaning, i.e., the concept (or composition of concepts). The term is syntax (e.g., a string) that stands in for or is used to indicate the semantics (meaning).	
Concept (a universal category for referents): A unit of semantics (meaning), the node (entity) or link (relation) in the mental or knowledge representation model. In an ontology, a concept is the primary knowledge construct, typically a class, relation, property, or attribute, generally associated with or characterized by logical rules. In an ontology, these classes, relations, properties are called concepts because it is intended that they correspond to the mental concepts that human beings have when they understand a particular body of knowledge (subject matter area or domain) but at the philosophical *universal* level, i.e., as *kinds* of entities. In general, a concept can be considered a placeholder for a category (way of characterizing) of specific real world referents (synonymously: specific entities, instances, individuals, or particulars), and thus ontology as an engineering product is about representing the semantics of the real world in a model that is usable and interpretable by machine.	

than relation (thesauri) for terms. Only the subclass/narrower than relation is a *generalization-specialization* relation (subsumption).[4]

A Conceptual Model can be considered a *weak ontology*; a Logical Theory (Fig. 2.3) can be considered a *strong ontology*. The innermost circle is the set of axioms. The middle circle is the set of theorems. The outermost circle is the ever expanding theory, an ontology as logical theory about reality which grows over time, as new axioms are entered and new theorems deduced. An ontology as a logical theory is thus: (1) a set of (non-logical) *axioms*, i.e., the classes, properties, subclass and subproperty assertions, the relations, attributes, and constraints on these; (2) the potentially expanding set of *theorems*, which can be proven true by some valid justification mechanism such as that which a typical formal logic provides, i.e., a set of equivalences or valid reasoning patterns known as *inference rules*, e.g., Modus Ponens; (3) *interpretations*, which are not depicted in the figure, are the mappings between a given theory and the set of models (in the sense of *model-theory* (Makowsky, 1992; Van Leeuwen, 1994; Hodges, 1997), which are supposed to be what the syntactic expressions of the theory *mean*. The whole, a logical theory, constitutes the specific, growing ontology.

The primary distinction here between a weak and a strong ontology is that a weak ontology is expressed in a knowledge representation language which is not

for purposes of simplicity, we use the term *concept* in this chapter to mean roughly *an abstract entity signifying a general characterizing idea or universal which acts as a category for instances.*

[4]The *subsumption* relation is typically defined to be the *subset* relation, i.e., intuitively a class is similar to a set, and the instances of that class are similar to elements of the set. A more general class (set), therefore, like *mammal* will contain a subclass (subset) of *primate*, among whose instances (elements) will be specific humans like Ralph Waldo Emerson. *Concept subsumption* as an ontology reasoning problem means that "given an ontology O and two classes A, B, verify whether the interpretation of A is a subset of the interpretation of B in every model of O" (OWL 1.1. http://www.w3.org/Submission/owl11-tractable/).

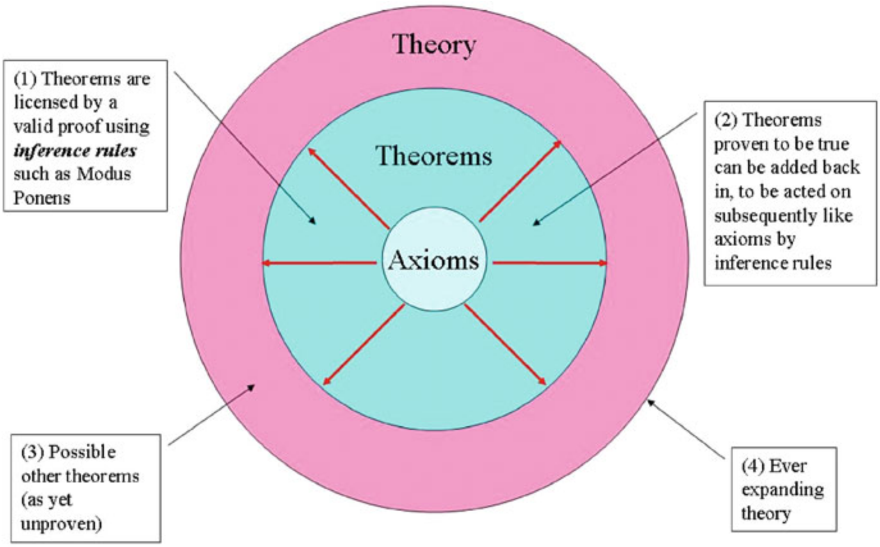

Fig. 2.3 Ontology as logical theory

based on a formal logic. Why is this important? It means that a machine can only read and process a weak ontology (e.g., currently models in ER or UML). It cannot semantically interpret the ontology, i.e., ingest the ontology and perform automated reasoning on it (reasoning which is similar to that which a human would make). So a weak ontology is not semantically interpretable by machine; a strong ontology is.

So what is usually, colloquially considered by the larger community to be an ontology needs clarification: all of these models should instead be considered *semantic models.* An ontology is restricted to the upper half of the Ontology Spectrum The Ontology Spectrum therefore displays the range of models in terms of expressivity or richness of the semantics that the model can represent, from “weak” or less expressive semantics at the lower left (value set, for example), to “strong” or more expressive semantics at the upper right. The vertical lines, labeled by *syntactic interoperability, structural interoperability,* and *semantic interoperability,* indicate roughly the expressiveness of the model require to respectively address those levels of interoperability.[5] *Syntactic interoperability* is defined as enabling the interchange of information based on a common syntax for at least that interchange. *Structural interoperability* is defined as a providing a common structure (a higher-order syntax) to enable the interchange of information. For example, multiple documents

[5]There are both lower levels of interoperability and higher levels. Lower levels include logical and physical accessibility and connectivity interoperability, e.g., having two information sources on the communication network, with network addresses known by those who might wish to access those sources. A higher level might be *pragmatic interoperability* (intending a formal pragmatics account), which factors in the intent of the represented semantics.

may be syntactically represented in XML, but need to be validated against distinct structural XML schemas or Document Type Definitions (DTD), which can be viewed as grammar rules that organize components of the syntax in specific ways. *Semantic Interoperability* is defined as providing a common semantics to enable the interchange of information, i.e., the semantics of the structural layer: what those structural components mean.

As depicted in the Ontology Spectrum, XML is sufficient for syntactic interoperability, XML Schema enables structural interoperability, but a minimum of RDF is necessary for semantic interoperability.

Figure 2.4 maps those semantic models against the increasingly more complex applications that are able to be addressed by using those models.

In the above diagram, *term* (terminology) and *concept* (real world referent) are defined as previously in Table 2.1.

As the expressiveness of the semantic model increases, so does the possibility of solving more complex problems. At the Taxonomy level, an application can provide only simple categorization, indexing, search, navigation: for example, indexing your documents into loose topic buckets with some hierarchic organization. Using thesauri can enable a search application to increase recall by, for example, using synonyms and substituting these into an expanded query string. For applications that require more precision, i.e., where approximate or loose characterizations of the semantics simply will not accomplish what is needed, more expressive models such as Conceptual Models and Logical Theories, i.e., ontologies, are required.

Recall is a measure of how well an information search and retrieval system finds ALL relevant documents on a searched for topic, even to the extent that it includes some irrelevant documents. *Precision* is a measure of how well such a system finds

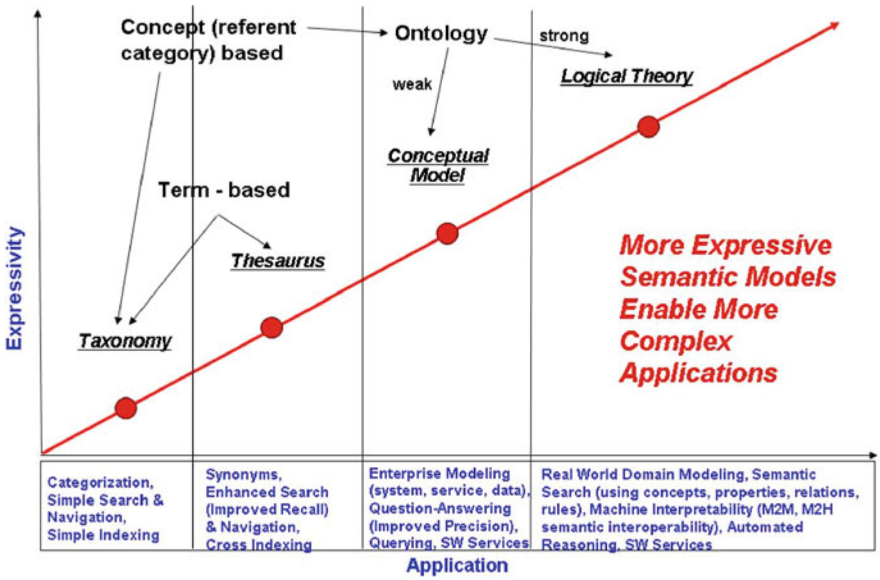

Fig. 2.4 More expressive semantic models enable more complex applications

Table 2.2 Recall vs. precision[6]

Recall: The percentage of relevant documents retrieved:
$\frac{\text{Number of relevant docs retrieved}}{\text{Number of relevant docs}}$
***Precision**:* The percentage of retrieved documents judged relevant:
$\frac{\text{Number of relevant docs retrieved}}{\text{Number of docs retrieved}}$

ONLY relevant documents on a searched for topics, even to the extent that it skips irrelevant documents. Table 2.1 displays the usual definitions of recall and precision. In most cases, recall and precision are inversely proportional to one another, with high recall leading to low precision, and high precision meaning the recall is low (Buckland and Gey, 1994; Table 2.2).

2.2.4 The Ontology Maturity Model

Building on the notions of the Ontology Spectrum, we describe one possible view of how an enterprise may migrate from less expressive semantic models to more expressive models, i.e., to real ontologies, based on both the common understanding of the enterprise and its requirements for more complex applications. Figure 2.5 displays an overall Ontology Maturity (or Capability) Model, simplified here, that shows the significant gradations toward greater maturity an organization may take in its evolution toward more completely realizing the goal of an ontology-driven enterprise.

This figure, which is patterned after the Software Engineering Institute's (SEI) Capability Maturity Model (CMM) that was intended to describe and gauge an organization's software process maturity (Royce, 2002), we attempt to develop a scale of maturity in an organization's migration towards increasingly more robust approaches to the use of ontologies for information technology needs.

Our analysis is that initially an organization thinks primarily of local semantics, i.e., attempts to characterize their information technology needs based on (currently mainstream) syntactic and structural methods, with only implicit semantics: a nodding of the head to signify agreement with the semantics as uttered in speech, or an agreement on a data dictionary of English or other natural language definitions, which ostensibly humans can read and indirectly nod their heads over. However, as an organization evolves, it begins to understand that it is actually composed of many communities and sub-organizations, each of which has its own local semantics but in addition a common enterprise-wide semantics, in fact a common semantics based

[6]"Recall is like throwing a big fishing net into the pond. You may be sure to get all the trout, but you've probably also pulled up a lot of grouper, bass, and salmon, too. Precision is like going spear fishing. You'll be pretty sure to ONLY get trout, but you'll no doubt miss a lot of them, too." – Jim Robertson, http://www-ec.njit.edu/~robertso/infosci/recall-precision.html

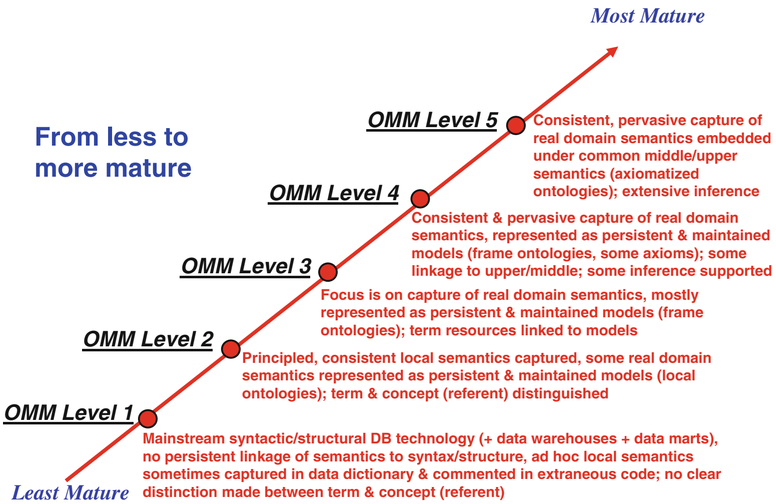

Fig. 2.5 Ontology maturity model (OMM)

on real world referents that all communities and sub-organizations in the enterprise share. Most commonly, as a semantically aware enterprise matures, it eventually distinguishes between terms (ways of referring) and concepts/referents (referents referred to by potentially many different terms). Hence, the semantic models the maturing enterprise embraces evolves from term-based models (weak taxonomies and thesauri) to concept/referent-based models (weak and strong ontologies).

In addition, as the maturing enterprise begins to understand that terminologies are not as necessary as the underlying meanings (concepts) of those terminologies that get modeled as a machine usable or interpretable engineering semantic model (ontology), the enterprise tries to fit together the local semantic models it currently has (local database schemas or even local community ontologies). Because it is soon recognized that there is great and incommensurable, though seemingly duplicative, meaning among the diverse ontologies (conceptual stovepipes), the enterprise attempts to reconcile the semantics. It does so initially by trying to construct semantic mappings between the two ontologies, and then when the problem repeats itself with every additional ontology which needs to be incorporated (mapped to), the enterprise begins to understand that the emerging mapping ontology is actually an integrative ontology that must be as expressive as the most expressive of the ontologies needing to be integrated.

2.3 Ontological Architecture: Upper, Mid-level, Domain Ontologies

In this section we discuss the fundamentals of ontological architecture. As depicted in Fig. 2.1, an ontological architecture encompasses primarily three layers: upper ontologies, mid-level ontologies, and domain ontologies, with the first two also sometimes called *foundational ontologies*. This section focuses on these. However, in Section 2.4, we will generalize the architecture to include a meta-level.

2.3.1 What Is an Upper Ontology?

Ontologies may exist at many levels of abstraction. We group ontologies into three broad categories of upper, mid-level and domain ontologies. In this section we define what we mean by an upper ontology and characterize the differences between these three levels. Figure 2.6 is a graphical depiction of these notional levels along with some sample concepts that may be found at each level.

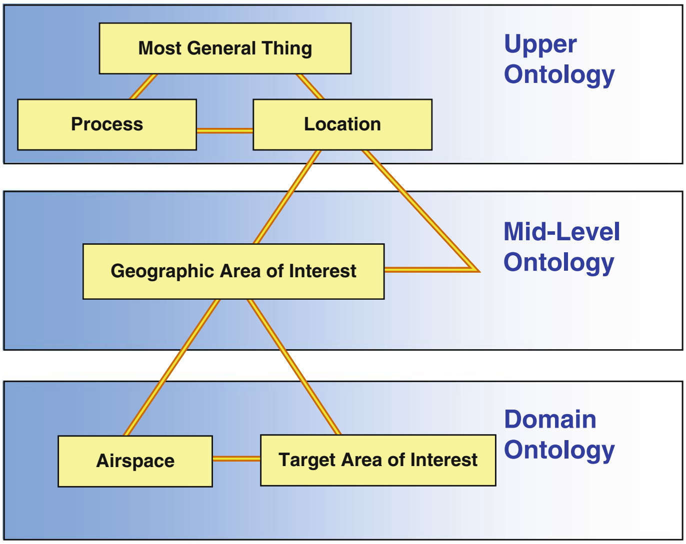

Fig. 2.6 Ontology categories

2.3.1.1 Upper Ontology Definition

An upper ontology, as defined by Phylita (2002), is a high-level, domain-independent ontology, providing a framework by which disparate systems may utilize a common knowledge base and from which more domain-specific ontologies may be derived. The concepts expressed in such an ontology are intended to be basic and universal concepts to ensure generality and expressivity for a wide area of domains. An upper ontology is often characterized as representing common sense concepts, i.e. those that are basic for human understanding of the world (Kiryakov et al., 2001). Thus, an upper ontology is limited to concepts that are meta, generic, abstract and philosophical.[7] Standard upper ontologies are also sometimes referred to as foundational ontologies[8] or universal ontologies (Colomb, 2002).

[7] Standard Upper Ontology (SUO) Working Group Website, http://suo.ieee.org/.
[8] OpenCyc Website, http://www.opencyc.org/.

2.3.1.2 Upper Ontology vs. Mid-Level Ontology

A mid-level ontology serves as a bridge between abstract concepts defined in the upper ontology and low-level domain specific concepts specified in a domain ontology. While ontologies may be mapped to one another at any level, the mid-level and upper ontologies are intended to provide a mechanism to make this mapping of concepts across domains easier. Mid-level ontologies may provide more concrete representations of abstract concepts found in the upper ontology. This ontology category also encompasses the set of ontologies that represent commonly used concepts, such as Time and Location. These commonly used ontologies are sometimes referred to as utility ontologies.

2.3.1.3 Upper Ontology vs. Domain Ontology

A domain ontology specifies concepts particular to a domain of interest and represents those concepts and their relationships from a domain specific perspective. While the same concept may exist in multiple domains, the representations may widely vary due to the differing domain contexts and assumptions. Domain ontologies may be composed by importing mid-level ontologies. They may also extend concepts defined in mid-level or upper ontologies. Reusing well established ontologies in the development of a domain ontology allows one to take advantage of the semantic richness of the relevant concepts and logic already built into the reused ontology. The intended use of upper ontologies is for key concepts expressed in a domain ontology to be derived from, or mapped to, concepts in an upper-level ontology. Mid-level ontologies may be used in the mapping as well. In this way ontologies may provide a web of meaning with semantic decomposition of concepts. Using common mid-level and upper ontologies is intended to ease the process of integrating or mapping domain ontologies.

2.3.2 Why Do We Care About Upper Ontology?

2.3.2.1 How Upper Ontologies May Help

Today's World Wide Web (WWW) is geared toward presenting information to humans. The Semantic Web is an evolution of the WWW that is intended to capture the meaning of data (i.e., data semantics) precisely enough that a software application can interpret them. A key element of the Semantic Web is the use of ontologies to define concepts and their relationships. With ontologies supplying the context of data, information retrieval and search engines can exploit this contextual information to perform semantic searches based on the meaning of the concept, rather than syntactic searches of a given text string. In this way, one could discriminate between horses and cars which both have the same label of "mustang." Rich semantics captured in ontologies also provide the ability to combine simple facts together to infer new facts, and to deduce new generic knowledge in the form of proven theorems that is only implicit in the ontologies. With data and applications mapped to ontologies, inference engines could be used to improve the discovery and understanding of data as well as the discovery and composition of applications like Web services. Furthermore, ontologies may be used to represent the semantics of applications and services directly, much as UML object and conceptual models do today for specific systems and enterprises, though these do so incompletely, inconsistently, and unsoundly, without explicit use by the applications of these models at either system-generation time or run-time. Upper ontologies are intended to define foundational concepts used in both mid-level and domain ontologies. In theory, the mapping between domain ontologies becomes easier if the ontologies to be mapped are derived from a standard upper ontology.

Two approaches exist for the use of upper ontologies: top-down and bottom-up. In a top-down approach one uses the upper ontology as the foundation for deriving concepts in the domain ontology. In this way, the domain ontology designer takes advantage of the knowledge and experience already built into the upper ontology. Furthermore, use of the upper ontology provides a theoretical framework on which to build. In a bottom-up approach, the ontology designer maps a new or existing domain ontology to the upper ontology. This approach also capitalizes on the knowledge built into the upper ontology but one would expect the mapping to be more challenging, as inconsistencies may exist between the domain and upper ontology. Some upper ontologies utilize a combination of these two approaches.

2.3.2.2 A Software Engineer Analogy

Let's use a software engineering analogy to describe the value of using standard upper and mid-level ontologies. Mid-level ontologies can be seen as analogous to software libraries. Early high level programming languages evolved to contain software libraries of commonly used functions. High quality software libraries allowed programmers to reuse the knowledge and experience built into the software library and freed them to concentrate on domain specific issues. As software libraries evolved, programming tasks became easier. Programmers do not need to understand the detailed implementation of libraries in order to use them. Similarly, mid-level ontologies can evolve to act as ontological utilities. With the existence of such ontologies, ontology designers can compose their domain ontologies using these utility ontologies and inherit the concepts and inferencing capabilities provided by them. Just as software libraries make programming tasks easier, so too would the availability of high quality, commonly used utility ontologies make ontology development easier. Further, concepts in the utility ontology could be mapped to concepts in an upper ontology without the need for users of the utility ontology to be aware of these mappings.

Because it is early in the Semantic Web evolution (OWL became a World Wide Web Consortium [W3C] recommendation in Feb'04), few utility ontologies exist. However, they are emerging, as evidenced by the DARPA funded effort to create a standard time ontology, now a W3C public working draft (Hobbs and Pan, 2006).

2.3.3 What Foundational Ontologies Provide: Ontological Choices

We cannot evaluate foundational ontologies here (but see Section 2.3.4). However, we can provide some rationale for why foundational ontologies are useful in an overall ontological architecture and what kinds of constructs one might desire for a foundational ontology. We call these *ontological choices* (though Partridge (2002) calls them *meta-ontological choices*).

What are the ontological choices that a given foundational ontology provides? These ontological choices will entail *ontological commitments,* which means that there is downward impact on mid-level and domain ontologies on the decisions one makes at the upper or foundational levels. The WonderWeb Ontology Library Final Report (Masolo et al., 2003), for example, describes a number of such ontological choices: descriptive vs. revisionary, multiplicative vs. reductionist, universals vs. particulars vs. sets, endurants vs. perdurants, and more. Other choices include 3-dimensional (3D) vs. 4-dimensional (4D) (Hamlyn, 1984; Loux, 2002), distinct notions of “part” and “whole”, different notions about what constitutes a property (and attribute), how change should be represented, distinctions about granularity, vagueness, etc.[9] Many of these choices are intricately linked, so, for example, discussions on endurants and perdurants invoke 3D and 4D views, and crucially elucidate the notion of persistence through time and change. In addition, multiplicative ontologies, because they tolerate a greater range of modeling complexity (model whatever is called for by reality), generally enable multiple objects with different identity criteria to co-occur/co-locate in the same spacetime (Masolo et al., 2003). In the following, we discuss some of these choices.

2.3.3.1 Descriptive vs. Revisionary

Descriptive and *revisionary* ontologies (Strawson, 1959) are based on ontological stances or attitudes towards the effort of modeling ontologies, i.e., how one conceptualizes the world and what an ontological engineering product is or should be. A *descriptive* ontology tries to capture the more commonsensical and social notions based on natural language usage and human cognition, emphasizing the agent who conceives and deemphasizing scientific and philosophical considerations. A *revisionary* (sometimes called *prescriptive*) ontology, on the other hand, does emphasize (or even, strictly adheres to) the scientific and philosophical perspectives, choosing to base its constructs and modeling decisions on scientific theories and a philosophical stance that tries to capture the world as it really is (it *prescribes* the world), and not necessarily as a given historical agent conceives it to be. A revisionary ontology therefore says that its modeling constructs are about real things in the world as it is.

[9] Another choice we will not investigate here is that between *presentism* and *eternalism* (Partridge, 2002). Presentism argues that time is real; eternalism that time is not real, that entities change but their properties do not change over time. Presentism typically goes with endurantism; eternalism goes with perdurantism.

In practical terms, all of the constructs in a *revisionary* ontology will be space-time objects, i.e., necessarily having temporal properties; in a *descriptive* ontology, that will not be the case. In the latter, *entities* (sometimes called *endurants,* but perhaps better called *continuants*) such as "hammer" and "tank" that have only incidental temporal properties and *events* (processes, actions, activities, etc., sometimes called *perdurants,* but perhaps better called *occurrents*) such as "attacking" and "cashing a check" that have explicit temporal properties, are modeled with or without those temporal properties, respectively. Often in natural language there are two correlated forms/usages that express the distinction: the nominal and the verbal. A nominal (noun) "attack" is expressed as in "The attack on the enemy began at 600 hours." A verbal (verb) "attacked" is expressed as in "We attacked the enemy at 600 hours."

2.3.3.2 Multiplicative vs. Reductionist

A multiplicative upper ontology is expressively profligate in that concepts can include anything that reality seems to require, and so any distinction that seems useful to make can be made in the ontology. Contrarily, a reductionist ontology reduces the number of concepts to the fewest primitives sufficient to derive the rest of complex reality.

In the WonderWeb Foundational Library (Masolo et al., 2003), the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) and the Basic Formal Ontology (BFO) are multiplicative and descriptive, whereas the Object-Centered High-Level Reference Ontology (OCHRE) is reductionist and revisionist. The Suggested Upper Merged Ontology (SUMO)[10] (Niles and Pease, 2001b) could be said to be both multiplicative in that it aims to cover at a general level any concept that reality requires, and reductionist in that it attempts to be minimal rather than profligate.

We note that many of these dichotomous ontology choices (*descriptive* vs. *revisionary*, *multiplicative* vs. *reductionist,* etc.) really have behind them a set of assumptions about how to view the world (e.g., strict realism with no notion of a different possibility) and what an engineering model of the world or parts of the world can achieve. Therefore, many of the ontology choices will tend to co-occur: e.g., *revisionist* and *reductionist* will generally go together.

2.3.3.3 Universals, Particulars, Sets, Possible Worlds

The distinction between *universals* (forms, ideas) and *particulars* (individuals) brings up a range of philosophical argument that we cannot address here. For our purposes, universals (whether based on realism, conceptualism, or nominalism) are general entities. Universals are often characterized as natural classes that abstract or generalize over similar particular things. Person, Location, Process, etc., are

[10] Suggested Upper Merged Ontology (SUMO) Website. http://www.ontologyportal.org/.

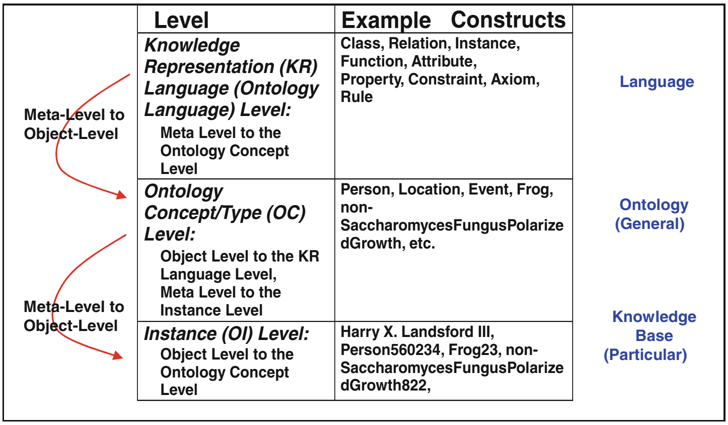

Fig. 2.7 Ontology representation levels (Obrst et al., 2007)

examples of universals, and would be represented at the Ontology Concept Level in Fig. 2.7 (see next section).

If you take a *realist* stance, universals are “entities of a general kind that exist independently of us and of our ways of thinking and speaking about the world” (Hamlyn, 1984). A *conceptualist* views universals as existing in human minds and primarily functioning as concepts that generalize and classify things. *Nominalists* view universals as largely a notion of our human language, the mode of expression of our thoughts. In an extreme view of realism, Platonism, universals independently exist (it’s usually considered unproblematic that particulars exist in reality), and so in our discussion of upper ontologies here, universals would exist in a quantificational domain distinct from that of particulars. This could be the case, for example, if universals were represented at the Ontology Concept level, but the Knowledge Language level of Fig. 2.7 permits second-order quantification, i.e., quantification over concepts (properties, predicates, classes, relations, etc.), rather than just over particulars (individuals, instances) at the Ontology Instance level.

A further distinction can be made: some instances (particulars or individuals) can themselves be considered universals – at least from the perspective of ontology applications (Welty, Ferucci, 1999). Degen et al. (2001) address this issue by introducing universals of higher-order. The Semantic Web ontology language OWL in fact allows for classes as instances in the OWL-Full dialect (Smith et al., 2004).

Particulars, or individuals or instances, are specific entities and taken to be instantiations of universals. Particulars exemplify properties (which are usually understood as universals), meaning they possess specific values such as Sam Jones being the father of Bill Jones, this apple in my hand being red, and that ball being on that table at 11 am EST, on January 19, 2008, in my house in Fairfax, Virginia, USA. Particulars are represented at the Instance Level in Fig. 2.7. Instances of classes (concepts), *facts* (specific instantiated relations/properties, e.g., Sam’s fatherhoodness to Bill, my apple’s redness), and *events* (a fact that occurs at a specific time, a specific perdurant) (Pianisi and Varzi, 2000; Higginbotham et al., 2000) are typically taken to be particulars.

Sets are mathematical objects that are sometimes, but not always used to abstractly characterize the different ontological categories, i.e., the logical apparatus used to define and order the logico-mathematical notions of ontology. Model-theoretical semanticists use set theory, but formal ontologists sometimes object (see e.g., Smith, 2001), where *mereotopology* (discussed below) is argued to provide a better foundation for ontology. Nonetheless, a set does not typically constitute a separate ontological category in its own right – except insofar as it is used as a human artifact. So, for example, SUMO defines a *set* as an ontological entity in its upper ontology because it does represent an entity that it used by other components of the SUMO upper ontology and potentially other lower, domain ontologies which use SUMO and make reference to sets directly, as ontological objects. A set in the first sense, i.e., as a defining mathematical notion, would typically be expressed at the meta-level, i.e., the Language Level in Fig. 2.7, and thus is not itself an object for ontological modeling.

It is perhaps a bit confusing or disconcerting to find that the object *set* really exists at two levels, i.e., at the modeling content level (Concept Level in Fig. 2.7) and also at its meta-level (Language Level, Fig. 2.7). The confusion devolves at least partially on the distinction between *use/mention* (Tarski, 1933, 1944), i.e., natural language typically allows one to both use a word and to mention it. So in this sense, 'set' is both an ontological object at the Ontology-Concept modeling level, and the meta-level object at the Language level which helps to define the entire Ontology-Concept level below it.

An additional consideration – which we will not discuss in any detail here – is the notion of *possible worlds*, which is a way of formally characterizing the distinction between *descriptions* (intensions) and *individuals which possess the properties described by the descriptions* (extensions). In a sense, the Cyc context and *microtheory*-based systematic manner of segregating assertions into theories, two of which taken together and compared may contradict each other, can be considered an implementation of the notion of possible worlds. *Possible worlds semantics* is usually a notion that also involves modal logic. We consider these notions in more detail in Section 2.4, Structuring the Ontological and Meta-Ontological Space.

2.3.3.4 Endurants and Perdurants

The distinction between *endurants* and *perdurants* is sometimes conflated with two different distinctions: (1) the distinction between *3D* and *4D* ontological objects, and (2) the distinction between *continuant* and *occurrent*, respectively. However, these conflations are problematic (Hayes et al., 2002; Sider, 2004; Degen et al., 2001). According to the usual definitions (Bittner and Smith, 2003), an *endurant* is an entity which exists in full in every instant at which it exists at all; a *perdurant* "unfolds itself over time in successive temporal parts or phases." Both endurants and perdurants are taken to be *concrete particulars,* i.e., instances (Loux, 2002). Obviously, the notion of identity- and essence-defining properties intersect with changeability. A perdurant is typically taken to be a *spacetime worm*, i.e., an object that persists (perdures) through spacetime by way of having different temporal parts at what would be different times (temporal non-locality), but a view of *instantaneous stages* is possible too (Sider, 2002). An endurant goes through time (endures), with identity/essence-defining properties that perhaps depend on occurrent objects but are not essentially constituted by those occurrent objects. The crucial distinction between these constructs is that of the nature of the identifying essential properties of the object and its change or non-change, usually defined with respect to time. Related to the distinction is the notion of *temporal parts*, i.e, whether or not a given object has temporal parts and the nature of those parts. But it is not just that distinction that defines 3D and 4D views, since some 3D perspectives permit instantaneous objects to be the temporal parts of themselves (Sider, 2002). For our purposes here, however, we will equate endurantism with the 3D view, and perdurantism with the 4D view.

A *partonomic* hierarchy, for example, is usually defined in terms of a special *partonomic* relation, the part-of relation. *Mereology* is the analysis of the part-of relation and the set of axioms that seem to constitute our notion of what a part is. In modern ontological axiomizations, mereology is combined with *topology* (connectedness among objects) to be *mereotopology* (Smith, 1996; Casati and Varzi, 1999) since *parthood* really does seem to require either point "touching", overlap, or transitivity of those (i.e., the 'southern edge of London' is part of London or connected to those regions which are part of southern London). Here we begin to get into notions of granularity and vagueness, and so we'll end our discussion (but, see: Obrst and Mani, 2000; Williamson, 1998; Keefe and Smith, 1999; Bittner and Smith, 2001).

2.3.4 Upper Ontology Initiatives and Candidates

There are a number of ongoing initiatives to define a standard upper ontology. Two initiatives that began in the early 2000s and recently ended were the IEEE Standard Upper Ontology Working Group (SUO WG)[11] and WonderWeb.[12] IEEE SUO WG was a standards effort operated under the IEEE Standards Association and sponsored by the IEEE Computer Society Standards Activities Board. Its goal is to specify an upper ontology that will enable computers to use it for applications such as data interoperability, information search and retrieval, automated inferencing, and natural language processing. IEEE SUO WG proposed three candidate upper ontologies, namely Suggested Upper Merged Ontology (SUMO), Upper Cyc Ontology (UCO)[13] and Information Flow Framework (IFF).

WonderWeb was a project consortium of universities and Industry, working in cooperation with the DARPA DAML program and the W3C. WonderWeb defined a library of foundational ontologies that cover a wide range of application domains.

[11] IEEE Standard Upper Ontology. http://suo.ieee.org/.

[12] WonderWeb Website. http://wonderweb.semanticweb.org/objectives.shtml. Completed, July 2004.

[13] Upper Cyc. http://www.cyc.com/cycdoc/vocab/vocab-toc.html.

This library is intended to be used as a basis for the development of more detailed domain ontologies. Currently three modules exist: DOLCE, OCHRE, and BFO (Masolo et al., 2003; Schneider, 2003).[14]

In addition, there have been proposed other upper (foundational) ontologies, including Generalized Ontological Language (GOL)/ General Formal Ontology (GFO) (Heller and Herre, 2004; Herre et al., 2006),

For comparisons of upper ontologies, see Grenon (2003); Semy et al. (2005); and Mascardi, Cordi, Rosso (2006). There was also an effort in 2006 by the Ontolog Forum called the Upper Ontology Summit,[15] at which many major upper ontology developers signed a joint communiqué to agree "to develop the mechanisms and resources needed to relate existing upper ontologies to each other in order to increase the ability to reuse the knowledge to which they give access and thereby facilitate semantic interoperability among those other ontologies that are linked to them" (Obrst et al., 2006).

2.4 Structuring the Ontological and Meta-Ontological Space

This section extends the ontological architecture considerations of the previous section in two ways: (1) it moves beyond purely vertical considerations of object level ontologies (upper, middle, domain) to include structural and logical relations among the ontologies of those levels, and ways of addressing the entire object level space, which we are calling the *ontological space*, and so necessarily involving notions of modularity and context (applicability of assertions); (2) it addresses also the *meta-ontological space,* i.e., the knowledge (ontology) representation (KR) space at the meta-level to the ontology object level. Although both of these topics require a more lengthy elaboration than we can provide here, we will sketch out some of the considerations and approaches. Because the two topics are so intricately connected, we flatten the structure of our exposition somewhat, addressing meta-ontological issues and then ontological issues, acknowledging explicitly that the latter depend on the former – even when a formalized connection cannot yet be established.

2.4.1 Knowledge Representation Languages and Meta-Ontologies

Another way of viewing ontological architecture is more abstractly, i.e., meta-ontologically, in terms of the representation levels. These representation levels include minimally: (1) the knowledge (ontology) representation language level; (2) the ontology concept (universals) level; (3) the ontology instance (particulars)

---

[14]Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) Website. http://www.loa-cnr.it/DOLCE.html.

[15]Upper Ontology Summit, Ontolog Forum, 2006. http://ontolog.cim3.net/cgibin/wiki.pl?UpperOntologySummit.

level. These are displayed in Fig. 2.7. The knowledge representation (KR) level is the meta-level to the ontology content level (which is its object level). The KR level characterizes what can be expressed in the ontology. The ontology concept level is the level that characterizes the generic descriptions of the ontology, i.e., universals or categories, the ontology proper, which might be considered either the organizing structure for the ontology instance level or the intensional level which describes the properties that will hold of specific individuals (the extension) at the ontology instance level. The third level (instances or particulars) is the level that instantiates the universals expressed at the second level (universals).

This view partially returns to the Ontology Spectrum perspective, in which the expressiveness of the knowledge representation language determines the richness of the object level ontology assertions that can be made.[16]

A given ontology is syntactically expressed in a particular logical or knowledge representation language. Although the choice of knowledge representation language is secondary to the actual ontological content, it is still important because it determines whether in fact a given upper ontology can be utilized completely or just partially.

Typically, upper ontologies require expressiveness at the level of First Order Logic (FOL), but occasionally require more, i.e., second-order or higher. Second-order is required if the upper ontology quantifies over predicates (or relations or properties), though limited finite quantification over predicates (in the form of a list of predicates) can be supported in a first-order language, as KIF/Common Logic demonstrates.[17]

Furthermore, an upper ontology may require a modal extension of FOL, depending on how modalities such as necessity/possibility and potential modalities such as temporal/spatial operators are expressed in the ontology. In general, modalities (necessity, belief, obligation, time, etc.) can be expressed either in the (meta level) logic/KR language or in the (object level) ontology, but in either case, ways to assert and refer to modal assertions will differ. These differences may be important to the expressions a domain ontology wants to make.

If the logic/KR language in which a given upper ontology is encoded is less expressive than the logic/language in which a specific upper ontology is expressed, semantic information loss will result. The resulting encoding of the upper ontology will contain only a subset of the original expression of the ontology. For example, if the original upper ontology is expressed in KIF/Common Logic and then encoded in OWL (Bechhofer et al., 2004) only a portion will be retained in OWL, which, being a description logic-based ontology language, tries to maximize machine reasoning tractability by minimally, but definitely, limiting expressivity. OWL Full, the most expressive "dialect" of OWL, may in fact be nearly equivalent in expressivity to FOL, but remains ultimately less expressive.

[16] Portions of this section are adapted from Semy, Pulvermacher, Obrst (2005), pp. 5-13.

[17] Common Logic. http://cl.tamu.edu/.

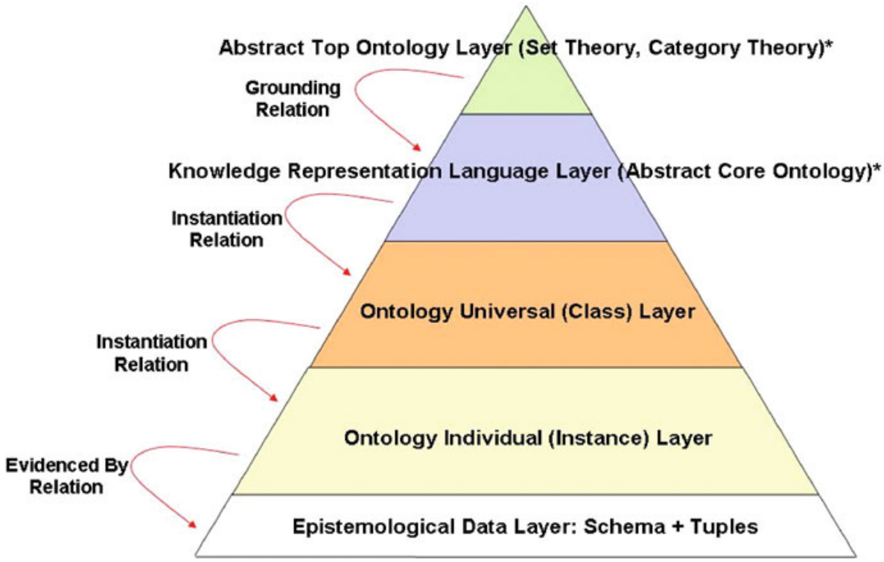

Fig. 2.8 Ontological architecture: a bigger picture

Finally, it should be noted that if KR languages are either not sufficiently formalized so that there is a clear notion of the formal semantics of the language, or are sufficiently formalized, but offer only indirect expression of upper ontology axioms, then portions of the upper ontology cannot be used by interpreting software. Portions of the upper ontology must then be annotated and interpreted solely by human beings.

A refinement of this three-level view is the (meta-)ontology architecture of Fig. 2.8.[18]

In this diagram, the KR language (or Abstract Core Ontology [ACO], in the usage of Herre and Loeb (2005)) is in a grounding relation to the abstract top ontology (ATO) layer, which is rooted in fundamental mathematical theories, a meta-level to our KR level; a meta-meta-level to our ontology concept/universal level, which after all is our primary interest here, i.e., it is rooted in set theory and category theory. Logic and in particular First-Order Logic presumably make up the KR/ACO level. In this view the ACO assumes at least some of the role of foundational ontologies (typically upper ontologies). For example, Herre and Loebe (2005, p. 1404) describe the basic constructs of the ACO as the following (Table 2.3), with the presumed two underlying core distinctions of the real world being that of *category* and *individual*:

[18] Adapted from Herre and Loeb (2005).

Table 2.3 Basic entity types and relations (Herre and Loebe, 2005)

Meta-Level Entity Types (Sets of urelements)			
Name	*Symbol*	*Name*	*Symbol*
Category	Cat	**Individual**	Ind
Object Category	OCat	Object	Obj
Property	P	Attribute	Att
Role Category	RCat	Role	Rol
Relation	R	Relator	Rel
Meta-Level Relations (Sets of pairs of urelements)			
Name	*Symbol*	*Argument Restrictions*	
identity	$x = y$	–	
instantiation	$x :: y$	Cat(y)	
inherence	Inh(x, y)	Att(x) or Rol(x)	
role-of role	(x, y)	Rol(x), Rel(y)	
categorial part-of	catp(x, y)	Cat(x), Cat(y)	

The potential value of this revised architecture is that it generalizes the constructs expressible at the KR language level, on the one hand, thus enabling many kinds of KR languages to be developed and be compared, while enforcing a logical consistency on the object ontologies developed in the Ontology Universal and Ontology Particular levels. On the other hand, an ACO is grounded in the firm mathematics of the ATO, i.e., its constructs are defined in terms of set theory and category theory, and presumably some variant of formal logic.

In Herre and Loeb (2005), two different ACOs are developed as an attempt to address the requirements for an ACO: initially, CPO, which is based on categories and properties only and thus not all of the constructs in Table 2.3; and then secondly, CRO, which addresses all of the constructs in Table 2.3, including in particular, relations. Of course, both of these ACO meta-ontologies have fragments which will have some constructs but not others. The notion of *concept lattice* from formal concept analysis (Ganter and Willie, 1996) is modeled as a experiment to gauge the expressiveness of CPO in Herre and Loeb (2005). They conclude that CPO does not appear to be expressive enough for all the examples given in Ganter and Willie (1996). However, they emphasize that this formalization does highlight distinct interpretations that can exist for object ontologies, based on the use of ACOs; this result in itself is valuable and argues for a meta-ontological architecture such as they describe.

Since Herre and Lobeb (2005) axiomatize CPO and CRO with a type-free FOL, presumably FOL and other logics constitute at least the lower levels of ATO, in addition to set theory and category theory at the higher levels. So ATO is best characterized as the logico-mathematical fundamental level in this architecture.

Given that ACOs partially constitute the KR language level, they really act as an upper meta-ontological level of the meta-logical KR language level identified in Fig. 2.7. The KR languages below them presumably act partially as ACOs (I am here thinking, for example, of the implicit class theories embedded in OWL and other KR languages), i.e., as *instantiations* of an apparently unexpressed ACO.

2.4.2 The Lattice of Theories

Because ontologies are often considered theories, i.e., as from our discussion previously of strong ontologies as logical theories, then a sensible question is: what are the relationships among the theories? Intuitively these relationships are mathematical or logical in nature. Among others, John Sowa (2005) has characterized the entire structure of ontologies *qua* theories as a *lattice of theories*. Sowa states that Lindenbaum showed that the partial ordering defining the lattice of theories (Davey and Priestley, 1996) can be view in three ways, as: (1) *implication*, (2) *provability,* and (3) *semantic entailment*. These notions are derived from the Lindebaum's infinite lattice of theories.

Figure 2.9 depicts implicitly the lattice of theories, but also a portion of the structural relationships among so-called *microtheories* and *little theories*, which are described in Section 2.4.4. This is a notional figure in which the alphabetic symbols A, B, . . ., Z (all with implicit subscripts to distinguish multiple occurrences) represent *propositions*, the symbol ';' represents *disjunction*, the symbol ',' represents *conjunction*, and the symbols Tn, Tn+1, . . ., Tn+i (n, i ∈ Integers), along with 'Top Theory', represent distinct *theories*.

In this figure, therefore, T1, . . ., T6 represent distinct, more specific theories in the lattice of theories having as Top (most general node) the Top Theory. Top Theory

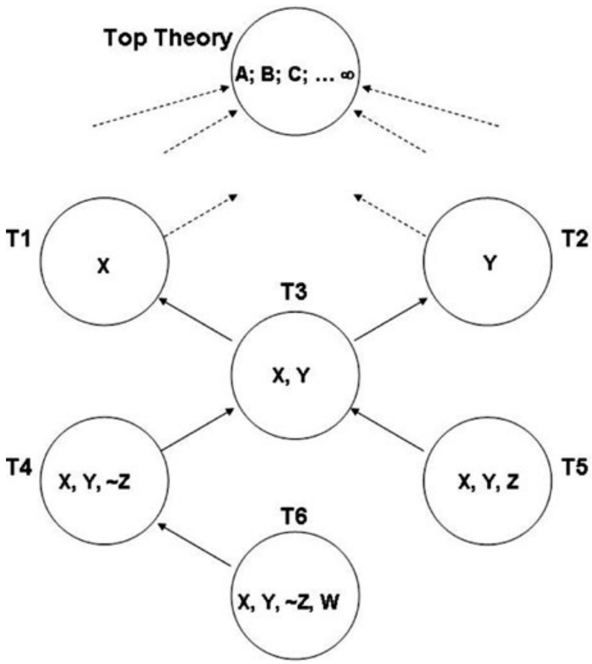

Fig. 2.9 Lattice of theories, microtheories, little theories

represents the disjunction of all possible propositions, both positive and negative (i.e., negated positive propositions).[19] Arrow lines represent the lattice relations among theories, interpreted as either implication, provability, or semantic entailment (though semantic entailment might be the most perspicuous here). Dotted arrow-lines represent (infinitely) other many theories in the lattice (including, it should be mentioned, other theories below T1 – T6).

So in this view, although T4 and T5 are mutually inconsistent (because they assert, respectively, ∼Z and Z), taken individually they are locally consistent. Furthermore, they are consistent with respect to T3. T3 in turn is consistent with both T1 and T2. T6 is consistent with T4 but not T5. All are consistent with Top Theory. Microtheories and little theories, as we will see, inhere in this framework. T1 – T6 can be considered microtheories or little theories.

2.4.3 Modularity and Context in the Ontological Space

The two notions of *modularity* and *context* are closely linked when we consider the larger *ontological space*. The ontological space we will define to be: the object level space of ontologies and their knowledge bases, i.e., the universals (classes, categories) and particulars (instances, individuals), and the theories and interpretations which make up this space. This section is concerned with ways that have been proposed to structure the relationships among those *modules* (theories and their interpretations).

In the 1980s, even prior to the rise of ontological engineering as a technical discipline of computer science, John McCarthy, along with his students, began to investigate the nature of context in knowledge representation and to formalize it (McCarthy, 1987, 1990, 1993; Guha, 1991; Buvač, 1993; Buvač, Buvač, Mason, 1995; Buvač 1996a-b; McCarthy and Buvač, 1997). Others took up the thread and implementations appeared, such as that of a *microtheory* in Cyc (Blair et al., 1992; Lenat and Guha, 1990).

Most formalizations of context use a specialized context lifting predicate *ist(P, C)* or alternatively *ist(C, F),* which means: proposition P is true in the Context C; formula F is true in Context C.[20] Typically contexts are first-class objects, and thus *ist* formulas are first-class formulas (Guha, 1991, p. 17). In Cyc, these contexts are implemented as *microtheories,* i.e., theories/ontologies in the general Cyc ontological space which are locally consistent, but taken together, are not globally consistent in the ontological space. In principle, this is similar to the *lattice of theories* notion discussed earlier, and also to *possible world semantics* (assuming the universes of discourse are the same) (Obrst et al., 1999), with the understanding that

---

[19] So, the theories are propositional theories here, for purposes of simplicity, but they should be understood as being formed from FOL or higher-order logical formulae, with the additional understanding that fully saturated (with instantiated, non-variable, i.e., *ground* terms) FOL or higher-order formulae are propositions.

[20] We use propositions here, but the general case is formulae.

if a microtheory was replaced with a set of consistent theories, the node in the theory space could be viewed as the possible world in which all those theories are true, and thence the relations among those theory-worlds could be construed as *accessibility relations* as among possible worlds, with those accessibility relations macroscopically thus corresponding to the Lindenbaum view microscopically, e.g., as semantic entailment.

Menzel (1999) points out that these formalizations of context propose a so-called "subjective conception" of context, i.e., one which defines contexts as sets of propositions, as theories related via an entailment relation – so typically as a set of beliefs of a person or agent, and thus *subjective*. Contrarily, Menzel (1999) argues for an "objective conception" of context, which views the truth of a proposition not as a logical relation (such as entailment) between the proposition of a context and other propositions, but instead as a correspondence relation between the proposition and the world, and thus *objective* (Obrst and Nichols, 2005, p. 2). Menzel (1999) therefore argues for the use of *situation theory* (Barwise and Perry, 1983), which explicitly establishes more granular formal contexts in natural language semantics than the usual notion of possible worlds, i.e., situations.

Giunchiglia and Ghidini (1998), Giunchiglia (1997), and Giunchiglia and Bouquet (1998) analyze context as *Local Model Semantics*. Their formalization of context is based on two properties identified as *Locality* and *Compatibility*. Locality, in their view, is a property shared by the language, the notion of satisfiability of a formula within a specific context, and the structure of individual contexts: everything in the determination of context is local. Compatibility, they characterize as the notion of mutual influence that contexts have on themselves, including the structural notion of changing the set of local models of two contexts so that they agree to some extent. LMS defines a special *model* for two languages (L_1, L_2) which is a *compatibility relation* $\boldsymbol{C} \subseteq 2^{\overline{\boldsymbol{M_1}}} \times 2^{\overline{\boldsymbol{M_2}}}$ (Giunchiglia and Ghidini, 1998, p. 284). Given the two languages, they associate each of them with a set $\boldsymbol{M_i} \subseteq \overline{\boldsymbol{M_i}}$ of local models, where $\overline{\boldsymbol{M_i}}$ is the class of all models of L_1. A specific *context* then is a set of *local* models $\boldsymbol{m} \in \overline{\boldsymbol{M_i}}$ allowed by $\boldsymbol{C}$. A formula φ of a context is satisfied in model $\boldsymbol{C}$ iff it is satisfied by all the local models of $\boldsymbol{C}$.[21] Bouquet et al. (2003) define a context extension of the Semantic Web ontology language OWL called Context OWL (C-OWL), which is based on the formalization of LMS.

Some KR languages have been proposed, which reify contexts. In the Interoperable Knowledge Representation for Intelligence Systems (IKRIS) project,[22] for example, a KR language was developed called IKL (IKRIS Knowledge Language), which can be considered an extended Common Logic, or "Common Logic on steroids" (Hayes and Menzel, 2006; Hayes, 2006). In IKL, contexts are formalized as first-class objects (Makarios, 2006a-c). But the decision was made to contextualize constants rather than sentences, and so (Welty, 2006):

---

[21] For further discussion of LMS, see Obrst et al. (1999) on which this current discussion is based.

[22] Interoperable Knowledge Representation for Intelligence Support (IKRIS). 2006. http://nrrc.mitre.org/NRRC/Docs_Data/ikris.

Fred in (ist C0 (P Fred)) is interpreted with respect to C0

And each constant is replaced with a function of the context and the constant:

{(forall(x) (implies (P (iso CMX)) (G (iso CMx))));
(P(isoC0Fred))}

Some questions to ask about contexts with respect to ontologies are the following. What is the relationship between a context and an ontology? What is the relationship between a context and a module of an ontology? Are contexts and ontologies distinct or the same? Is a context embedded within a given ontology (where the ontology is viewed as a theory or set of logical theories about a domain)? Is a context extraneous to an ontology and thus outside the ontology as theory, leading us to view a context as encapsulating ontologies and changing the interpretations of those ontologies in this context as opposed to that context (Obrst and Nichols, 2005)? In this section, we will discuss modularization in ontologies, and we will treat *contexts* as being essentially about perspectives (akin to but more complex than *views* in the relational database world), i.e., as logical theories (and their interpretations), which in our estimation are what ontologies are.

Since 1997, formalization of context has established itself as a technical thread in the CONTEXT conferences (CONTEXT 97, Brézillon and Cavalconti, 1997; CONTEXT 07, Kokinov et al., 2007). Modularity of ontologies in its own right has been addressed by very recent workshops [Haas et al., 2006; Cuenca-Grau et al., 2007). It is often remarked on that formalized context and ontology modules bear a close resemblance and depend on each other, which has led to the recent Context and Ontologies Workshops (Bouquet et al., 2007), and see in particular (Loeb, 2006; Bao and Honavar, 2006; Lüttich et al., 2006; Kutz and Mossakowski, 2007).

2.4.4 Microtheories, Little Theories, Ontology Versioning

A microtheory is a theory in Cyc (Blair et al., 1992; Kahlert and Sullivan, 2006) which is a portion of the (monolithic) ontology that is separable from other microtheories, and thus with respect to those possibly containing contradictory assertions. A Cyc microtheory "is essentially a bundle of assertions that share a common set of assumptions; some microtheories are focused on a particular domain of knowledge, a particular level of detail, a particular interval in time, etc. The microtheory mechanism allows Cyc to independently maintain assertions which are *prima facie* contradictory, and enhances the performance of the Cyc system by focusing the inferencing process."[23]

See Fig. 2.9, previously introduced in the discussion of the lattice of theories. Microtheories represent an implementation of a formalization of context

[23] What's in Cyc? http://www.cyc.com/cyc/technology/whatiscyc_dir/whatsincyc.

deriving from McCarthy (1987, 1990, 1993), but focused in particular on Guha (1991). Originally Cyc microtheory contexts consisted of two parts: assumptions and content. Subsequently, due to computational inefficiencies with the formalism (primarily in the cost of so many *liftings* of assertions from one context into another), the microtheory was recast with finer internal structure. Lenat (1998) identified 12 dimensions of context space along which contexts vary in Cyc: "Absolute Time, Type of Time, Absolute Place, Type of Place, Culture, Sophistication/Security, Granularity, Epistemology, Argument-Preference, Topic, Justification, and Anthropacity" (Lenat, 1998, p. 4), with each primary dimension being itself a bundle of partially mutually dependent finer-grained dimensions. A richer calculus of contexts is thus required. In our view, a dimension of context-space is thus similar to an index of common and commonsense world knowledge which cross-cuts domain theories (Obrst et al., 1999, p. 6). This latter usage more closely corresponds to Lewis's (1980) 'index', as opposed to his much richer 'context' (and see discussion in Giunchiglia and Bouquet (1998), p. 7).

In Farmer et al. (1992) and Farmer (1996, 2000), a formalization and implementation (for mathematical theorem proving) of a notion similar to that of microtheories is introduced, that of *little theories*. Defining a theory as a set of axioms in a formal language, Farmer et al. (1992) contrast two predominant views in mathematics: (1) the *big theory* approach, which is one large theory of very expressive axioms (such as Zermelo-Fraenkel set theory) such that the models of these axioms will contain all of the objects of potential interest to the mathematician; (2) the *little theory* approach, in which a number of theories will be used and different theorems are proven in different theories, depending on the amount of structure needed. The little theory approach uses *theory interpretations* as the primary formal notion, where theory interpretation is defined as "a syntactic translation between the languages of two theories which preserves theorems" (Farmer et al., 1992, p. 2). A formula which is a theorem in the source theory is thus translated into a theorem in the target theory, which requires the source theory axioms to be translated into theorems in the target theory. Theorems can thus be reused in different theories, and one can establish the consistency of the source theory with respect to the target theory. A theory interpretation between two theories also enables inferring a relation between the models of the two theories. Farmer et al. (1992) and Farmer (1996, 2000) also establish how such a formalization can be used to implement a proof system for doing semi-automated mathematical proofs, developing the Interactive Mathematical Proof System (IMPS). IMPS enables one to store theories and theory interpretations. The little theory approach therefore allows for a network of theories and provides both *intertheory* and *intratheory* reasoning. So, similar to the notion of microtheories, little theories enable different perspectives (different contexts) to be represented as different theories (microtheories), with the formal device enabling switching from one theory to another being the notion of a theory interpretation between the two theories, preserving theoremhood between the two theories.

Finally, De Leenheer (2004), De Leenheer and Mens (2007), De Leenheer et al. (2007) demonstrates the significance of the relation among ontologies, contexts, and *versions* of ontologies (among other perspectives), by introducing a formal framework for supporting context driven ontology engineering based on the DOGMA[24] framework and methodology. Note that this is not a formalization, but instead an elaboration and use of existing notions from context formalization.

2.4.5 Information Flow Framework Meta-Ontology

The Information Flow Framework (IFF)[25] is authored by Bob Kent (Kent, 2004, 2006) and was recently being developed under the IEEE SUO Working Group.[26] IFF provides a framework for sharing ontologies, manipulating ontologies as objects, relating ontologies through morphisms, partitioning ontologies, composing ontologies via fusions, noting dependencies between ontologies, and declaring the use of other ontologies. It takes the building block approach to ontology construction and management, using category theory (Mac Lane, 1971; Barr and Wells, 1999) and Information Flow Theory (IFT) (Barwise and Seligman, 1997) to support ontology modularity.

IFT is a framework more general than possible worlds, allowing also *impossible worlds* (Barwise, 1998). Figure 2.10, for example, displays the interpretation of one language or logic L_1 (*classification* in their terminology) into another L_2, with the accompanying association with every structure M_2 for the logic L_2 a structure M_1 for the logic L_1. One also has to make sure that $M_1 \models_{L1} \boldsymbol{\alpha}^1$ iff $M_2 \models_{L2} \boldsymbol{\alpha}^2$ holds for all structures M_2 for L_2 and all sentences α of L_1 Barwise and Seligman (1997, p. 32; Obrst et al., 1999, p. 7).

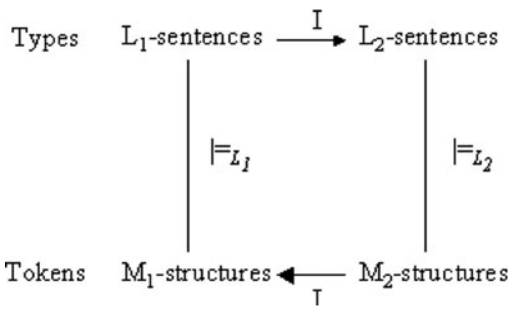

Fig. 2.10 Interpretation of languages

[24]Developing Ontology-Grounded Methods and Applications (DOGMA); a research initiative of the Free University of Brussels, Semantic Technologies and Applications Lab (VUB STARLab). http://www.starlab.vub.ac.be/website/.

[25]Information Flow Framework. http://suo.ieee.org/IFF/. See also: http://www.ontologos.org/IFF/IFF.html.

[26]IEEE Standard Upper Ontology. http://suo.ieee.org/.

IFF is primarily a meta-ontology, to be used for ontology-ontology integration, and is still in early stages of development. The original intent for the IFF was to define morphisms for concept lattices, mapping between Formal Concept Analysis (Ganter and Willie, 1996) and Information Flow classifications. In addition, the IFF attempts to develop a framework using IFT and category theory to implement the so-called "lattice of theories" view of the linkages among ontologies at the object level (Kent, 2003). The IFF is a fairly detailed framework, and within the limited space of this paper, no thorough elaboration can be given here. The interested reader is instead invited to peruse the IFF web site.[27] The primary architecture of the IFF can be depicted as in Fig. 2.11, which is from Kent (2002). See also Kalfoglou and Schorlemmer (2002, 2003, 2004) for discussion of IFF and issues in ontology mapping.

In the IFF, object level ontologies (upper, mid-level, domain) reside at the object level, and their constructs are defined at the meta-level in a series of ascending sub-levels, from IFF ontology (axiomatization) lower level to higher components. The lower meta-level defines, axiomatizes, and reasons about *particular* categories, functors, adjunctions, colimits, monads, classifications, concept lattices, etc., whereas the upper meta-level defines, axiomatizes, and reasons about *generic* categories, functors, adjunctions, colimits, monads, classifications, concept lattices, etc. The top meta-level was a formalization of the Knowledge Interchange Format (KIF).[28]

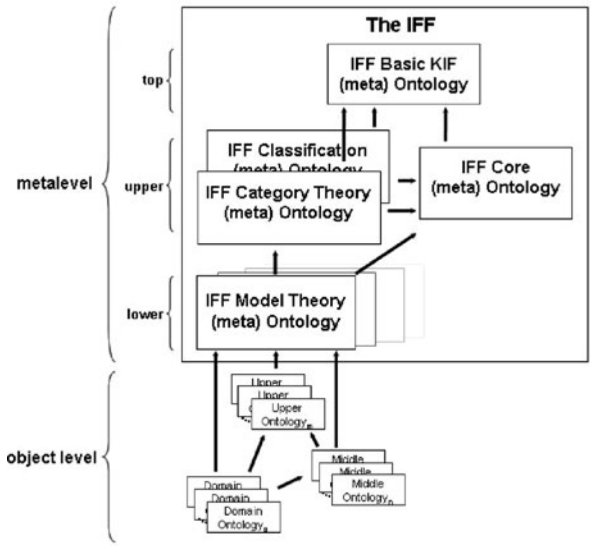

Fig. 2.11 Information flow framework: architecture (2002 version, slide 9)

[27] http://www.ontologos.org/IFF/IFF.html.

[28] Originally, Knowledge Interchange Format (KIF) Specification (draft proposed American National Standard [dpANS] NCITS.T2/98-004: http://logic.stanford.edu/kif/dpans.html. However,

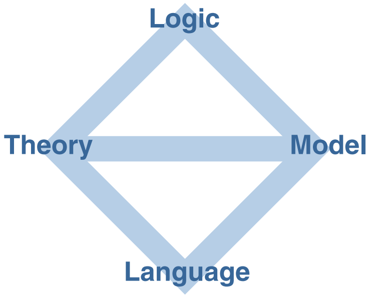

Fig. 2.12 Objects of IFF (Kent, 2005, slide 30)

In Fig. 2.12 is a picture of the primary objects in IFF: Logics, Theories (syntactic representations), Models (semantic representations), and Languages.

IFF builds on IFT, which in turn builds on category theory. But a formalization which is very close to IFF is that of *institutions* of Goguen (1991), Goguen and Burstal (1992), Goguen and Roşu (2002), and the related work of Kutz and Mossakowski (2007), and Schorlemmer and Kalfoglou (2008). *Institutions* formalize the notion of a logical system by demonstrating that there is a satisfaction relation between models and sentences which is consistent under change of notation (Goguen and Burstal, 1992; Goguen 2006). *Institutions* formalize category-theoretic *signatures* derived from (generalized from) vocabularies (ontologies) and *signature morphisms* derived from (generalized from) mappings between vocabularies (ontologies). So institutions are "truth preserving translations from one logical system to another", which is very similar to the intention of the IFF.

It should also be noted that the IFF, though a meta-level architecture, is more of a meta-logical architecture, rather than a meta-ontological architecture as Herre and Loebe (2005, p. 1411) point out, so – using the latter's terminology – the IFF is more like a very elaborated ATO rather than an ACO, in the terminology of Herre and Loebe (2005). However, given that ACOs act as the upper level of the KR language meta-logical level (the upper level in Fig. 2.7) (and so, specific KR languages can be seen as both *languages* and partial instantiations of implicit ACOs, i.e., as embodying meta-ontological *theories*), it does seem that the logical vs. ontological levels of the meta-level needs to be better spelled out. This discussion also demonstrates that further analysis of the interplay between the logical and ontological levels is a fruitful subject of study.

For further information about the IFF, the interested reader is referred to Kent (2010, this volume).

KIF has been superseded by Common Logic, which includes a KIF-like instantiation called CLIF, and is now an ISO standard.

2.5 What the Future Holds: A Vision

In this chapter, we have looked at ontological architecture, what it consists of and what it is related to. We began our discussion by delineating ontological architecture from ontology architecture, and observed that we would necessarily focus on the former – as being the necessary foundation for the latter. As one can see, there is clear technical apparati emerging that addresses the logical, meta-ontological, and ontological architectures and their requirements to support the actual development and deployment of ontologies as engineering products in an ontology architecture – which itself encompasses ontology lifecycle management as a practical discipline wedding ontologies and software development.

Along the way, we discussed not just the components of ontological architecture, but necessarily aspects of logical and meta-ontological architecture. We tried to relate these three notions systematically and consistently in a larger framework that we hope will provide support for subsequent ontology architecture. There are many moving pieces to this architectural puzzle. Ontology engineering as a branch of both computer science and software engineering has just recently emerged – and is propelled by ideas from formal ontology in philosophy, formal logic in mathematics, formal semantics in linguistics, formal methods and applications in computer science and artificial intelligence, and formal theories in cognitive science. There

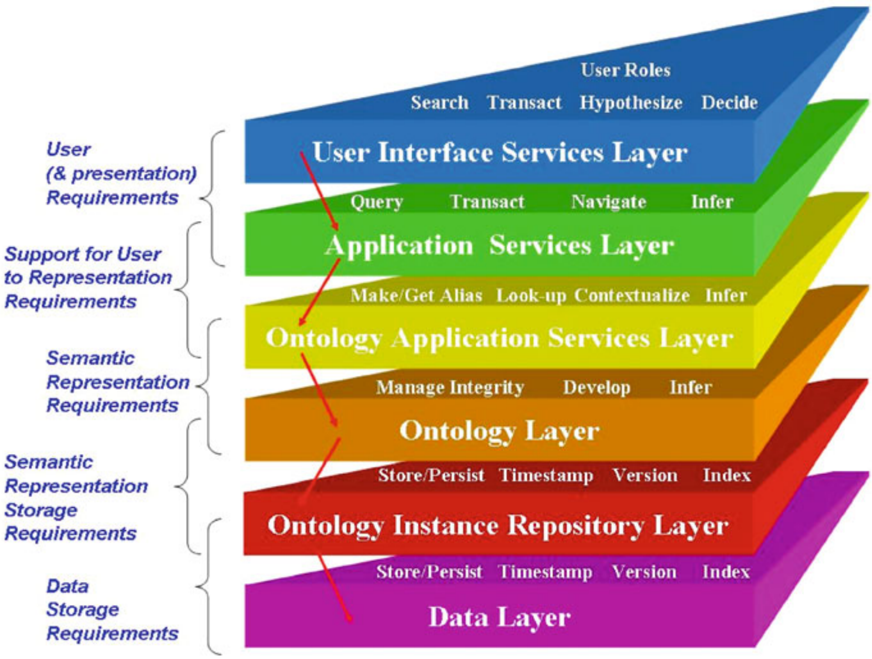

Fig. 2.13 Ontology architecture: application layers

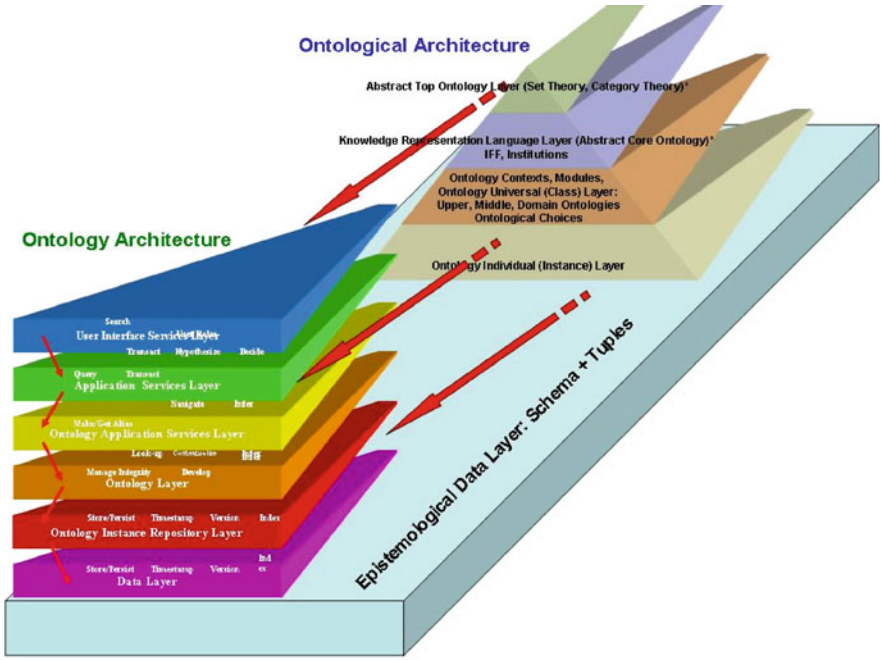

Fig. 2.14 Ontological architecture supporting ontology architecture

is a grand fermenting of philosophical, logico-mathematical, scientific, and engineering ideas that make the future uncertain to predict. There are, however, some indications, enough for a vision for the future.

Figure 2.13 depicts one view of an *ontology architecture*, which admittedly we could not address in this chapter. However, this architecture represents a sound view of what is architecturally necessary to deploy ontologies in the real world. Each of these layers are significant and necessary for an application based on ontologies. Each layer constitutes hefty portions of a distinct chapter on actually using engineered ontologies to assist users by way of their software applications and services in the near to foreseeable future.

Behind this figure and its depicted architecture, however, is another figure and its depicted architecture. Figure 2.14 shows the notional relationship between the ontological architecture and the ontology architecture.

This is the future: sound ontology philosophy and science driving sound ontology engineering, with many other technical disciplines collaborating to provide sound ontology-based applications.

Acknowledgments and Disclaimers The author's affiliation with The MITRE Corporation is provided for purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by the authors. I note that the views expressed in this paper are those of this author alone and do not reflect the official policy or position of any other organization or individual. I acknowledge deep appreciation for reviews of this material and suggestions to improve it by Roberto Poli and anonymous reviewers. Discussions (mostly electronic, but occasionally live) on these topics have been much appreciated, with Roberto Poli, John Sowa, Bill Andersen, members of the MITRE Information Semantics group including Ken Laskey, David Ferrell, Deborah Nichols, Mary Parmelee, Merwyn Taylor, Karla Massey, Jonathan Tivel, Frank Zhu, Richard MacMillan, and former or auxiliary members including Ken Samuel and Allen Ginsberg, and friends and colleagues including especially Suzette Stoutenburg (my best colleague and friend), Arnie Rosenthal (my best devil's advocate and good friend), Len Seligman, Marwan Sabbouh, Peter Mork, my persevering and nascent ontologist wife Christy Obrst, and a pack of good dogs, cats, and chickens to mostly keep me level.

References

Andersen, B., R. Kohl, J. Williams. 2006. Ontology lifecycle management. Ontology Works, Inc. Version 1.0. Oct 14, 2005. Ontologyworks.com.

Bao, J., and V.G. Honavar. 2006. *Divide and conquer semantic web with modular ontologies - A brief review of modular ontology language formalisms*. Proceedings of the 1st International Workshop on Modular Ontologies, WoMO'06, CEUR Workshop Proceedings, vol. 232. eds. H. Haas, S. Kutz, and A. Tamilin, 5 Nov 2006, Athens, GA.

Barr, M., and C. Wells. 1999. Category theory for computing science, 3rd edn. Montreal, QC: Les publications Centre de recherches mathématiques.

Barwise, J. 1998. Information and impossibilities. Manuscript. http://www.phil.indiana.edu/~barwise/

Barwise, J., and J. Seligman. 1997. *Information flow: The logic of distributed systems*. Cambridge, UK: Cambridge University Press.

Barwise, J., and J. Perry. 1983. Situations and attitudes. Cambridge, MA: MIT Press.

Basic Formal Ontology (BFO): http://www.ifomis.uni-saarland.de/bfo

Basin, D., M.D. Agostino, D. Gabbay, S. Matthews, and L. Viganò, eds. 2000. *Labelled Deduction*, 107–134, Scotland, UK: Kluwer Academic Publishers, 2000.

Bechhofer, S., F. van Harmelen, J. Hendler, I. Horrocks, D.L. McGuinness, P.F. Patel-Schneider, and L>A. Stein. 2004. OWL web ontology language reference. W3C recommendation. 10 Feb 2004. http://www.w3.org/TR/owl-ref/

Bittner, T., and B. Smith. 2001. Granular partitions and vagueness. In *Formal ontology and information systems*, eds. C. Welty, and B. Smith, 309–321. New York, NY: ACM Press.

Bittner, T., and B. Smith. 2003. Endurants and perdurants in directly depicting ontologies. Draft. http://ontology.buffalo.edu/smith/articles/EPDDO.pdf

Blackburn, P. 1999. Internalizing labelled deduction. In Proceedings of Hylo'99, First International Workshop on Hybrid Logics. 13 July 1999, Saarbrücken, Germany. Published in 2000. *Journal of Logic and Computation* 10(1):137–168.

Blair, P., R.V. Guha, and W. Pratt. 1992. Microtheories: An ontological engineer's guide. Technical report Cyc-050-92, 5 Mar 1992, Cycorps, Austin, TX. http://www.cyc.com/tech-reports/cyc-050-92/cyc-050-92.html

Bouquet, P., J. Euzenat, C. Ghidini, D.L. McGuinness, L. Serafini, P. Shvaiko, and H. Wache, eds. 2007. Proceedings of the International Workshop on Contexts and Ontologies: Representation and Reasoning (C&O:RR). Collocated with the 6th International and Interdisciplinary Conference on Modelling and Using Context (CONTEXT-2007) Roskilde, Denmark, 21 Aug. http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-298/

Bouquet, P., F. Giunchiglia, F. Van Harmelen, L. Serafini, and H. Stuckenschmidt. 2003. C-OWL: Contextualizing ontologies. 2nd International Semantic Web Conference (ISWC 2003), eds. D. Fensel, K.P. Sycara, and J. Mylopoulos, 164–179. Sanibel Island, FL, 20–23 Oct 2003.

Brézillon, P., and M. Cavalcanti. 1997. Modeling and Using Context: Report on the First International and Interdisciplinary Conference, CONTEXT-97, Rio de Janeiro, RJ, Brazil on 4–6 Feb 1997. *The Knowledge Engineering Review* 12(4):1–10.

Buckland, M.K., and F. Gey. 1994. The relationship between recall and precision. *Journal of the American Society for Information Science* 45(1):12–19.

Buvač, S., V. Buvač, and I.A. Mason. 1995. Metamathematics of contexts. *Fundamenta Informaticae* 23(3):149–174.

Buvač, S. 1996. Resolving lexical ambiguity using a formal theory of context. In *Semantic ambiguity and underspecification*. CSLI Lecture Notes, Center for Studies in Language and Information, Stanford, CA, 1996.

Buvač, S, ed. 1996. Proc AAAI-95 Fall Symposium on Formalizing Context. http://www-formal.stanford.edu:80/buvac/95-context-symposium

Casati, R., and A.C. Varzi. 1999. Parts and places: The structures of spatial representation. Cambridge, MA: MIT Press.

Cheikes, B. 2006. MITRE Support to IKRIS, Final Report. MITRE Technical Report MTR060158, 5 Dec 2006. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/MITRE_IKRIS_Final_Report_05Dec06.doc

CMMI® for Development, Version 1.2. CMMI-DEV, V1.2. CMU/SEI-2006-TR-008. ESC-TR-2006-008. CMMI Product Team, Aug 2006. http://www.sei.cmu.edu/cmmi/models/model-v12-components-word.html

Colomb, R.M. 2002. Use of Upper Ontologies for Interoperation of Information Systems: A Tutorial, Technical Report 20/02 ISIB-CNR, Padova, Italy, Nov 2002.

Common Logic. http://cl.tamu.edu/

Cuenca-Grau, B., V. Honavar, and A. Schlicht, F. Wolter. 2007. Second International Workshop on Modular Ontologies. 28 Oct 2007. Whistler, BC. http://webrum.uni-mannheim.de/math/lski/WoMO07/

Cycorp, Inc.: The CYC Technology. http://www.cyc.com/tech.html

Daconta, M., L. Obrst, and K. Smith. 2003. *The semantic web: The future of XML, web services, and knowledge management.* New York, NY: John Wiley, June 2003.

Davey, B.A., and H.A. Priestley. 1991. *Introduction to lattices and order*. Cambridge, UK: Cambridge University Press.

Degen, W., B. Heller, H. Herre, and B. Smith. 2001. GOL: A general ontological language. In Guarino, Welty, Smith, 34–46. FOIS 2001.

De Leenheer, P. 2004. Revising and Managing Multiple Ontology Versions in a Possible Worlds Setting. In Proc. of On The Move to Meaningful Internet Systems Ph.D. Symposium (OTM 2004) (Agia Napa, Cyprus), LNCS 3292, 798–818. Collega Park, MD: Springer.

De Leenheer, P. and T. Mens. (2007). Ontology evolution: state of the art and future directions. In *Ontology management for the semantic web, semantic web services, and business applications, from semantic web and beyond: Computing for human experience*, eds. M. Hepp, P. De Leenheer, A. de Moor, and Y. Sure. Heidelberg, Springer.

De Leenheer, P., A. de Moor, and R. Meersman. (2007). Context dependency management in ontology engineering: A formal approach. *Journal on data semantics VIII*, LNCS 4380,26–56. New York, NY: Springer.

Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) Website. http://www.loa-cnr.it/DOLCE.html

Dumontier, M., and N. Villanueva-Rosales. 2007. Three-layer OWL ontology design. Second International Workshop on Modular Ontologies. 28 Oct 2007. Whistler, BC.

Euzenat, J., A. Zimmermann, and F. Freitas. 2007. Alignment-based modules for encapsulating ontologies. Second International Workshop on Modular Ontologies. 28 Oct 2007. Whistler, BC.

Farmer, W.M. 2000. An infrastructure for intertheory reasoning. In *Automated Deduction--CADE-17*, ed. D. McAllester, 115–131. LNCS, vol. 1831. http://imps.mcmaster.ca/doc/intertheory.pdf

Farmer, W.M. 1996. Perspective switching using theories and interpretations. In *Intelligent Systems: A Semiotic Perspective*, eds. J. Albus, A. Meystel, and R. Quintero, 206–207. Vol. I, Gaithersburg, MD: National Institute of Standards and Technology, 20–23 Oct 1996. Abstract.

Farmer, W.M., J.D. Guttman, and F.J. Thayer. 1992. Little theories. In *Automated Deduction--CADE-11*, LNCS, vol. 607, ed. D. Kapur, 567–581. http://imps.mcmaster.ca/doc/major-imps-papers.html

Fernández, M. 1999. Overview of methodologies for building ontologies. Workshop on Ontologies and Problem-Solving Methods: Lessons Learned and Future Trends. (IJCAI99). Aug 1999. http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-18/4-fernandez.pdf

Fernandéz, M., A. Gómez-Pérez, and N. Juristo. 1997. METHONTOLOGY: From ontological art to ontological engineering. Workshop on ontological engineering. AAAI Spring Symposium Series. AAAI-97, Stanford University.

Fikes, R., C. Welty. 2006. Interoperable knowledge representation for intelligence support (IKRIS). Advanced research and development activity (ARDA)/disruptive technology office (DTO). Final briefing, Nov 2006.

Gabbay, D. 1996. *Labelled deductive systems, principles and applications. Vol 1: Introduction*. Oxford: Oxford University Press.

Ganter, B., and R. Wille. 1996. *Formal concept analysis: Mathematical foundations*. Berlin, Heidelberg, New York, NY: Springer.

Gašević, D., D. Djurić, and V. Devedžić. 2006. *Model driven architecture and ontology development*. Berlin Heidelberg, New York, NY: Springer.

Giunchiglia, F., P. Bouquet. 1997. Introduction to contextual reasoning: An artificial intelligence perspective. Istituto per la Ricerca Scientifica e Tecnologica (IRST), Trento, Italy, Technical report 9705–19, May 1997.

Giunchiglia, F., and P. Bouquet. 1998. A Context-Based Framework for Mental Representation. Istituto per la Ricerca Scientifica e Tecnologica (IRST), Trento, Italy, Technical report 9807-02, July 1998.

Giunchiglia, F., and C. Ghidini. 1998. Local models semantics, or contextual reasoning = locality + compatibility. In *Principles of knowledge representation and reasoning (KR'98)*, eds. A. Cohn, L. Schubert, and S. Shapiro, 282–289. Proceedings of the Sixth International Conference, Trento, Italy, 2–5 June 1998.

Goguen, J., and G. Rosu. 2002. Institution morphisms. *Formal Aspects Of Computing* 13:274–307.

Goguen, J.A. 1991. A categorical manifesto. *Mathematical Structures in Computer Science* 1:49–67.

Goguen, J.A., and R.M. Burstall. M. 1992. Institutions: Abstract model theory for specification and programming. *Journal of the ACM* 39:95–146.

Goguen, J. 2006a. Institutions. http://www-cse.ucsd.edu/~goguen/projs/inst.html

Goguen, J. 2006b. Information integration in institutions. *For jon barwise memorial volume*, ed. L. Moss, Indiana: Indiana University Press.

Guarino, N, ed. 1998. *In Formal ontology in information systems*. Amsterdam: IOS Press. Proceedings of the First International Conference (FOIS'98), 6–8 June, Trento, Italy.

Guarino, N., C. Welty, B. Smith, eds. 2001. *The Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-01)*, Oct 16–19 2001, Ogunquit, Maine. ACM Press Book Series, Sheridan Publishing, Inc. http://www.fois.org/fois-2001/index.html

Guha, R.V. 1991. Contexts: A formalization and some applications. PhD thesis, Stanford University. Also published as technical report STAN-CS-91-1399-Thesis, and MCC Technical Report Number ACT-CYC-423-91.

Haase, P., V. Honavar, O. Kutz, Y. Sure, and A. Tamilin, eds. 2006. Proceedings of the 1st international workshop on modular ontologies. WoMO'06. 5 Nov 2006, Athens, Georgia, USA. CEUR Workshop Proceedings, vol. 232. http://ftp.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-232/

Hamlyn, D.W. 1984. *Metaphysics*. Cambridge: Cambridge University Press.

Hayes, P.J. 2006. IKL Guide. On behalf of the IKRIS Interoperability Group. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/ICL_guide.pdf

Hayes, P.J., and C. Menzel. 2006. IKL Specification Document. Unpublished IKRIS memorandum. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/ICL_spec.pdf

Hayes, P., F. Lehmann, C. Welty. 2002. Endurantism and Perdurantism: An Ongoing Debate. IEEE SUO email exchange, compiled and edited by Adam Pease. http://www.ontologyportal.org/pubs/dialog-3d-4d.html

Heller, B., H. Herre. 2004. Ontological Categories in GOL [Generalized Ontological Language]. *Axiomathes* 14:57–76.

Herre, H., and F. Loebe. 2005. A Meta-ontological architecture for foundational onatologies. In *CoopIS/DOA/ODBASE 2005*, LNCS 3761, ed. R. Meersman, and Z. Tari, 1398–1415. Berlin Heidelberg: Springer.

Herre, H., B. Heller, P. Burek, R. Hoehndorf, F. Loebe, and H. Michalek. 2006. General formal ontology (GFO) – part I: Basic principles. *Technical report 8*, Germany: Onto-Med, University of Leipzig.

Higginbotham, J., F. Pianesi, and A. Varzi, eds. 2000. *Speaking of events*. Oxford Oxford University Press.

Hobbs, J.R., and F. Pan, 2006. Time Ontology in OWL. W3C Working Draft 27 Sept 2006. http://www.w3.org/TR/owl-time/

Hodges, W. 1997. *A shorter model theory*. Cambridge, UK: Cambridge University Press. Reprinted in 2000.

IEEE Standard Upper Ontology (SUO). http://suo.ieee.org/

Interoperable Knowledge Representation for Intelligence Support (IKRIS). 2006. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/

Information Flow Framework. http://suo.ieee.org/IFF/ See also: http://www.ontologos.org/IFF/IFF.html

Izza, S., L. Vincent, and P. Burlat. 2005. Ontology urbanization for semantic integration: Dealing with semantics within large and dynamic enterprises. Proceedings of the 2005 Ninth IEEE International EDOC Enterprise Computing Conference (EDOC'05).

Kahlert, R.C., and J. Sullivan. 2006. Microtheories. In *First international workshop: Ontology based modeling in the humanities*, eds. W. von Hahn, and C. Vertan, 7–8 Apr 2006, University of Hamburg (Bericht 264).

Kalfoglou, Y., and M. Schorlemmer. 2004. In *Ontology mapping: The state of the art*, eds. A. Sheth, S. Staab, and M. Uschold, Schloss Dagstuhl, Germany. Dagstuhl Seminar Proceedings on Semantic Interoperability and Integration Internationales Begegnungs- und Forschungszentrum (IBFI), http://drops.dagstuhl.de/opus/volltexte/2005/40/pdf/04391.KalfoglouYannis.Paper.40.pdf

Kalfoglou, Y., and M. Schorlemmer. 2003. Ontology mapping: The state of the art. *Knowledge Engineering Review* 18(1):1–31.

Kalfoglou, Y., and M. Schorlemmer. Information-flow-based ontology mapping. In *On the move to meaningful internet systems 2002: CoopIS, DOA, and ODBASE*. Lecture Notes in Computer Science 2519, 1132–1151. Heidelberg: Springer.

Keefe, R., and P. Smith, eds. 1999. *Vagueness: A reader*. Cambridge, MA: MIT Press.

Kent, R.E. 2002. The IFF approach to semantic integration. Power point presentation at the boeing Mini-workshop on semantic integration, 7 Nov 2002. http://www.ontologos.org/Papers/Boeing%20Mini-Workshop.ppt

Kent, R.E. 2003. The IFF approach to the lattice of theories. Unpublished manuscript, Apr 22 2003. Formerly at: http://suo.ieee.org/IFF/lattice-of-theories.pdf

Kent, R.E. 2004. The IFF foundation for ontological knowledge organization. In *Knowledge organization and classification in international information retrieval*, eds. N.J. Williamson, and C. Beghtol, 187–203. Volume 37 of Cataloging & Classification Quarterly. New York: Haworth Press.

Kent, R.E. 2006. The information flow framework: New architecture. International Category Theory Conference (CT 2006), White Point, Nova Scotia, 25 June – 1 July 2006.

Kent, R.E. 2010. The information flow framework. Chapter in Part One: Ontology as Technology in the book: *TAO – Theory and applications of ontology, volume 2: The information-science stance*, eds. M. Healy, A. Kameas, and R. Poli.

Kiryakov, A., K. Simov, and M. Dimitrov. 2001. Onto map: Portal for upper-level ontologies. In Proceedings of the International Conference on Formal Ontology in Information Systems, 2001, eds. W. Guarino, and Smith, 47–58.

Knowledge Interchange Format (KIF) Specification (draft proposed American National Standard [dpANS] NCITS.T2/9 8-004: http://logic.stanford.edu/kif/dpans.html Now superseded by Common Logic.

Knowledge Web (KWEB). 2004–2008. http://knowledgeweb.semanticweb.org/index.html

Kokinov, B., D.C. Richardson, T.R. Roth-Berghofer, and L. Vieu, eds. 2007. Modeling and Using Context 6th International and Interdisciplinary Conference, CONTEXT 2007, Roskilde, Denmark, 20–24Aug 2007, Proceedings. Lecture Notes in Artificial Intelligence Vol. 4635, 2007.

Kutz, O., and T. Mossakowski. 2007. *Modules in transition - Conservativity, composition, and colimits*. Second International Workshop on Modular Ontologies. 28 Oct 2007, Whistler, BC.

Lenat, D. 1998. The dimensions of Context-space, cycorp technical report, 28 Oct 1998. http://www.cyc.com

Lenat, D., and R. Guha. 1990. *Building large knowledge based systems*. Reading, MA: Addison Wesley.

Lewis, D. 1980. Index, Context, and Content. In *Philosophy and grammar*, eds. S. Kanger, and S. Ohman, Dordrecht: Reidel Publishing.

Loux, M.J. 2002. *Metaphysics: A contemporary introduction*, 2nd edn. London and New York: Routledge.

Lüttich, K., C. Masolo, and S. Borgo. 2006. Development of modular ontologies in CASL. Proceedings of the 1st International Workshop on Modular Ontologies, WoMO'06, CEUR Workshop Proceedings, vol. 232, eds. H. Haas, S. Kutz, and A. Tamilin, 5 Nov 2006, Athens, GA.

Loebe, F. 2006. Requirements for Logical Modules. Proceedings of the 1st International Workshop on Modular Ontologies, WoMO'06, eds. H. Haas, S. Kutz, A. Tamilin, 5 Nov 2006, Athens, Georgia, USA. CEUR Workshop Proceedings, Vol. 232.

Makarios, S. 2006a. A model theory for a quantified generalized logic of contexts. Stanford University Technical Report KSL-06-08. ftp://ftp.ksl.stanford.edu/pub/KSL_Reports/KSL-06-08.pdf

Makarios, S. 2006b. ICL Specification. On behalf of IKRIS Contexts Working Group. nrrc.mitre.org/NRRC/Docs_Data/ikris/ICL_spec.pdf.

Makarios, S. 2006c. ICL Guide. On behalf of IKRIS Contexts Working Group. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/IKL_guide.pdf

Mac Lane, S. 1971. *Categories for the working mathematician*. New York: Springer.

Makowsky, J.A. 1992. Model theory and computer science: An appetizer. In: *T.S.E. Handbook of logic in computer science, Volume 1, background: Mathematical structures*, eds. S. Abramsky, D. Gabbay, and Maibaum, 763–814. Oxford: Clarendon Press.

Masolo, C., S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. WonderWeb Deliverable D18 Ontology Library (final), 31 Dec 2003.

McCarthy, J., and S. Buvač. 1997. Formalizing Context (Expanded Notes). In *Computing natural langauge*, eds. A. Aliseda, R. van Glabbeek, and D. Westerståhl, Amsterdam: Stanford University. http://www-formal.stanford.edu.

McCarthy, J. 1987. Generality in artificial intelligence. *Communications of the ACM* 30(12): 1030–1035.

McCarthy, J. 1990. *Formalizing common sense: Papers By John Mccarthy*. Westport, CT: Ablex Publishing Corporation, 355 Chestnut Street, Norwood, NJ 07648.

McCarthy, J. 1993. Notes on formalizing context. In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, 1993.

Menzel, C. 1999. The objective conception of context and its logic. *Minds and Machines* 9(1):29–56 (Feb 1999).

Obrst, L, D. Nichols. 2005. *Context and ontologies: Contextual indexing of ontological expressions*. AAAI 2005 Workshop on Context and Ontologies, poster, AAAI 2005, 9–13 July, Pittsburgh, PA.

Obrst, L., H. Liu, and R. Wray. 2003. Ontologies for corporate web applications. *Artificial Intelligence Magazine*, special issue on Ontologies, American Association for Artificial Intelligence, ed. C. Welty, 49–62, Fall 2003.

Obrst, L, D. Nichols. 2005. *Context and ontologies: Contextual indexing of ontological expressions*. AAAI 2005 Workshop on Context and Ontologies, poster, AAAI 2005, July 9–13, Pittsburgh, PA. http://www.mitre.org/work/tech_papers/tech_papers_05/05_0903/index.html

Obrst, L., P. Cassidy, S. Ray, B. Smith, D. Soergel, M. West, and P. Yim. 2006. The 2006 Upper Ontology Summit Joint Communiqué. *Journal of Applied Formal Ontology* 1:2.

Obrst, L., T. Hughes, and S. Ray. 2006. *Prospects and possibilities for ontology evaluation: The view from NCOR.* Workshop on Evaluation of Ontologies for the Web (EON2006), Edinburgh, UK, 22 May 2006.

Obrst, L., W. Ceusters, I. Mani, S. Ray, and B. Smith. 2007. The Evaluation of Ontologies: Toward Improved Semantic Interoperability. Chapter In *Semantic web: Revolutionizing knowledge discovery in the life sciences,* eds. J.O. Christopher, Baker, K-H. Cheung, Heidelberg: Springer.

Obrst, L. 2002. Ontology spectrum and semantic integration and interoperability. Briefing to DARPA, 2 June 2002.

Obrst, L., and I. Mani, eds. 2000. Proceedings of the Workshop on Semantic Approximation, Granularity, and Vagueness, 11 Apr 2000, Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR-2000), 12–16 Apr, Breckenridge, CO.

Obrst, L., G. Whittaker, and A. Meng. 1999. Semantic interoperability via context interpretation, Context-99, Trento, Italy, Apr 1999, invited poster session.

Object Management Group (OMG): Meta Object Facility (MOF) Specification. 2006. Version 2.0. Object Management Group (OMG), Needham, Massachusetts

OntoWeb. http://www.ontoweb.org/ Finished on 13 May 2004. Follow-on project was Knowledge Web.

OpenCyc Website. http://www.opencyc.org/

OWL 1.1. Web Ontology Language Tractable Fragments, W3C Member Submission 19 December 2006. http://www.w3.org/Submission/owl11-tractable/

Niles, I., and A. Pease. 2001a. Origins of the IEEE Standard Upper Ontology, in Working Notes of the IJCAI-2001 Workshop on the IEEE Standard Upper Ontology.

Niles, I., and A. Pease. 2001b. Towards a standard upper ontology. In Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001), eds. C. Welty, and B. Smith, Ogunquit, Maine, 17–19 Oct 2001.

Partridge, C. 2002. Note: A couple of meta-ontological choices for ontological architectures. Technical Report 06/02, LADSEB-CNR, Padova, Italy, June 2002. http://www.boroprogram.org/bp_pipex/ladsebreports/ladseb_t_r_06-02.pdf

Phytila, C. 2002. An Analysis of the SUMO and Description in Unified Modeling Language, Apr 2002, unpublished. http://ontology.teknowledge.com/Phytila/Phytila-SUMO.html

Pianesi, F., and A. Varzi. 2000. *Events and event talk: An introduction*. Ch. 1 of Higgenbotham et al., 3–48. New York: Oxford University Press.

Poli, R. 2003. Descriptive, formal and formalized ontologies. In *Husserl's logical investigations reconsidered*, ed. D. Fisette, and D. Kluwer, 193–210. Dordrecht: Kluwer.

Poli, R. 2010. The categorial stance. Chapter 2 in Part One: Ontology as Technology in the book: *TAO – Theory and applications of ontology, volume 2: The information-science stance*, eds. M. Healy, A. Kameas, and R. Poli.

Poli, R., and L. Obrst. 2010. The interplay between ontology as categorial analysis and ontology as technology. Chapter 9 in Part One: Ontology as Technology in the book: *TAO – Theory and applications of ontology, volume 2: The information-science stance*, eds. M. Healy, A. Kameas, and R. Poli.

Pulvermacher, M., L. Obrst, S. Semy, and S. Stoutenburg. 2005. *Perspectives on applying semantic web technologies in military domains*, MITRE Technical Report.

Royce, W. 2002. CMM vs. CMMI: From conventional to modern software management. Rational edge: E-zine for the rational community, Feb 2002. Rational Software Corporation. http://www-128.ibm.com/developerworks/rational/library/content/RationalEdge/feb02/ConventionalToModernFeb02.pdf

Samuel, K., L. Obrst, S. Stoutenberg, K. Fox, A. Johnson, K. Laskey, D. Nichols, and J. Peterson. 2006. *Transforming OWL and semantic web rules to prolog: Toward description logic programs*. Poster, ALPSWS: Applications of Logic Programming in the Semantic Web and Semantic Web Services, 16 Aug 2006, Federated Logic Conference 2006, Seattle, WA.

Semy, S., M. Pulvermacher, and L. Obrst. 2005. Toward the use of an upper ontology for U.S. government and U.S. Military domains: An evaluation. MITRE Technical Report, MTR 04B0000063, Nov 2005. http://www.mitre.org/work/tech_papers/tech_papers_05/04_1175/index.html

Simperl, E.P.B., C. Tempich, and Y. Sure. 2006. ONTOCOM: A cost estimation model for ontology engineering. Proceedings of the International Semantic Web Conference ISWC 2006. http://ontocom.ag-nbi.de/

Smith, B. 2004. Beyond concepts: Ontology as reality representation. In *Proceedings of the Third International Conference on Formal Ontology and Information Systems*, eds. A. Varzi, and L. Vieu, 73–84. FOIS 2004, Turin, 4–6 Nov 2004, Amsterdam, The Netherlands: IOS Press.

Smith, B., J. Williams, and S-K. Steffen. 2003. The ontology of the gene ontology. *AMIA Annual Symposium Proceedings* 2003:609–613.

Smith, B. 1996. Mereotopology: A theory of parts and boundaries. *Data and Knowledge Engineering* 20:287–303.

Smith. B. 1998. Basic concepts of formal ontology. In *Formal ontology in information systems*, ed. N. Guarino, 19–28. Amsterdam: IOS Press.

Smith, M.K., C. Welty, and D.L. McGuinness. 2004. OWL web ontology language guide, W3C Recommendation, 10 Feb 2004. http://www.w3.org/TR/owl-guide/

Schneider, L. 2003. How to build a foundational ontology: The object-centered highlevel reference ontology OCHRE. In *Proceedings of the 26th Annual German Conference on AI, KI 2003: Advances in Artificial Intelligence, volume 2821 of Lecture Notes in Computer Science*, eds. A. Günter, R. Kruse, and B. Neumann, 120–134. Springer, 2003.

Schorlemmer, W.M., and Y. Kalfoglou. 2008. Institutionalising ontology-based semantic integration. Journal of applied ontology. 3(3):131–200.

Sider, T. 2002. *Four-dimensionalism: An ontology of persistence and time*. Oxford: Oxford University Press.

Sowa, J.F. 2005. Theories, Models, Reasoning, Language, and Truth. http://www.jfsowa.com/logic/theories.htm

Standard Upper Ontology (SUO) Working Group Website. http://suo.ieee.org/

Stalnaker, R. 1998. On the representation of context. Journal of logic language and information.

Stoutenburg, S., L. Obrst, D. Nichols, K. Samuel, and P. Franklin. 2006. Applying semantic rules to achieve dynamic service oriented architectures. RuleML 2006: Rules and Rule Markup Languages for the Semantic Web, co-located with ISWC 2006, Athens, GA, 10–11 Nov 2006. In Service-Oriented Computing – ICSOC 2006, Lecture Notes in Computer Science Volume 4294, 2006, 581–590. Heidelberg: Springer.

Stoutenburg, S., and L. Obrst, D. Nichols, J. Peterson, and A. Johnson. 2005. Toward a standard rule language for semantic integration of the DoD enterprise. W3C Workshop on Rule Languages for Interoperability, 27–28 Apr 2005, Washington, DC.

Stoutenburg, S., and L. Obrst. 2005. Toward a standard rule language for enterprise application integration. Rules and Rule Markup Languages for the Semantic Web, presentation, 3rd International Semantic Web Conference, 7–11 Nov 2005, Hiroshima, Japan.

Stoutenburg, S., and L. Obrst. 2005. Orchestration of ontologies and rules for integration of the DoD enterprise. Protégé With Rules Workshop, paper and presentation, 8th International Protégé Conference 2005, 18–21 July Madrid, Spain.

Strawson, P.F. 1959. *Individuals: An essay in descriptive metaphysics*. London: Methuan University Press.

Suggested Upper Merged Ontology (SUMO) Website. http://www.ontologyportal.org/

Tarski, A. 1933. The concept of truth in formalized languages. In *Logic, Semantics, Mathematics*, ed. J. Woodger, Oxford: Oxford University Press.

Tarski, A. 1944. The Semantic Conception of Truth and the Foundations of Semantics. In *Philosophy and phenomenological research, Volume 4*, 1944, 341–375. Reproduced in: The Philosophy of Language, A.P. Martinich, ed. 48–71. 1985, Oxford: Oxford University Press.

Thurman, D.A., A.R. Chappell, and C. Welty. 2006. Interoperable knowledge representation for intelligence support (IKRIS) evaluation working group report, 31 Dec 2006. http://nrrc.mitre.org/NRRC/Docs_Data/ikris/IKRIS_Evaluation_Report_31Dec06.doc

Upper Cyc: http://www.cyc.com/cycdoc/vocab/vocab-toc.html

Upper Ontology Summit, Ontolog Forum, 2006. http://ontolog.cim3.net/cgi-bin/wiki.pl?UpperOntologySummit

Van Leeuwen, J, ed. 1994. *Handbook of theoretical computer science. Volume B:Formal models and semantics*. Amsterdam: Elsevier, and Cambridge: MIT Press.

Welty, C. 2006. Interoperable knowledge representation for intelligence support (IKRIS). Briefing to ontolog forum, 10 Oct 2006. http://colab.cim3.net/file/work/SICoP/2006-10-10/Presentations/CWelty10102006.ppt

Welty, C., and D. Ferucci. 1999. Instances and classes in software engineering. *Artificial intelligence,* 24–28. Summer, 1999.

What's in Cyc? http://www.cyc.com/cyc/technology/whatiscyc_dir/whatsincyc

Williamson, Timothy. 1998. *Vagueness*. London, New York: Routledge.

WonderWeb Website. http://wonderweb.semanticweb.org/objectives.shtml. Completed, July 2004