Chapter 4
A Cognitive Science Perspective on Legal Ontologies

Joost Breuker and Rinke Hoekstra

4.1 Introduction

Ontological engineering has several origins. In this article we identify five, and all five still play a major role in ontology engineering. Each of these roots gives a different perspective on content and use of ontologies, but as these roots are hardly ever explicit, they are a source of much confusion. In the past, one has tried to solve this confusion by defining what an ontolology is: much in the spirit of ontology engineering iself. There is little consensus and the best definitions leave much underspecified. As we will show in Section 4.2, these five roots take different perspectives on what ontologies aim to capture. Philosophical ontology is concerned with "reality"; Information science with systematic terminology; Artificial Intelligence (AI) with terminological knowledge and Information Management with semantics. These differences may look subtle but have different consequences for the use of these ontologies, which ranges from analytic clarification to automated reasoning. This is also manifest from the different representation formalisms used. If these uses were consistent, we would not have to bother, but as will be shown in Section 4.4, mismatches occur between the representation formalism used and the aim of the application. These mismatches can often be traced to an unclear distinction between knowledge and semantics. In developing standards for the Semantic Web (SW) this confusion is reinforced by taking a knowledge representation formalism as a standard for ontologies, while current practice—mainly information management—could have done with a more expressive formalism. We explain the difference between knowledge and semantics, i.e. the use of knowledge in context, in Section 4.3 using a simple cognitive architecture for natural language production.

In Section 4.5 we will focus on an area of ontology engineering where a cognitive science perspective is well suited and should be taken more serious: the area of top/upper/foundational/core ontologies. These ontologies cover the abstract, general

J. Breuker (✉)
Leibniz Center for Law, University of Amsterdam, Amsterdam, The Netherlands
e-mail: breuker@science.uva.nl

G. Sartor et al. (eds.), *Approaches to Legal Ontologies*, Law, Governance and Technology Series 1, DOI 10.1007/978-94-007-0120-5_4,

concepts that are applicable to (almost) any domain, like time, space, causation, agent, matter, process, etc. Here also the distinction between knowledge and semantics fades. Thus far, this area of research has shown a large divergence of approaches and results. In many respects it is reminiscent of age long discussions in philosophical ontology; on the other hand the concerns are about representations, powerful to feed content-specific reasoning e.g. about spatial and temporal configurations. These kinds of concepts are also the drivers or cores of common sense, and sometimes an explicit common sense view is proclaimed (e.g. in DOLCE (Gangemi et al. 2002) and in CYC (Lenat 1995)). For both a philosophical perspective (e.g. Sowa's ontology (Sowa 2000), and to some extent SUMO (Pease and Niles 2002) and also DOLCE (!)), and a common sense perspective no external, empirical evidence is cited to support claims on the selection of concepts and their structuring. At the same time Cognitive Science has moved from an almost exclusive concern about cognitive processes to (the development of) the basic content of our shared cognitions and emotions, and a wealth of empirical data and theoretical insights have been reported. In the last Section of this article we will give some examples and refer to our still ongoing work on a common-sense based core ontology for legal domains: LKIF-Core (Hoekstra et al. 2007; Hoekstra 2009).

In summary, we will present two perspectives from cognitive science to ontology engineering. The first one comes from the architecture of human cognition in which one can argue that knowledge and semantics are not the same thing and should be treated differently: both with respect to formal-machine representation as to the kind of applications that should be fueled by ontologies. The second one concerns a valid identification of the basic concepts that constitute our deep common sense.

4.2 Origins of Ontological Engineering

In temporal order we can trace the following origins of ontology engineering:

4.2.1 Philosophy (1) Ontology

Ontology as a branch of metaphysics is concerned with (categories) of being, i.e. about existence and reality. Although the term "ontology" only appears in the seventeenth century, its endeavor is more than two millennia old (Parmenides; Aristotle). It has been emphasized already in the beginning of the 90-ies that ontologies (i.e. ontology engineering) and Ontology should not be confounded. The former is about (machine-readable) structures of concepts, while the latter is about "existing entities". Despite this difference, philosophical ontology has found a place in ontology engineering, often referred to as "formal ontology". This is also due to the fact that (philosophical) logic and theories of knowledge representation have come close enough to compare work. Ontology brings a wealth of analyses of concepts to ontologies. An excellent review of concepts relevant for ontology engineering can be found in Sowas book (2000). However, it also brings on board the speculative and argumentative nature, due to the lack of an applied or empirical perspective.[1]

4.2.2 Philosophy (2) Lingua Universalis Philosophica

The existence of different languages—the curse of Babel—was not only a serious obstacle for universal understanding, but also the fact that words have ambiguous and variable meanings was already a preoccupation of many philosophers in the Middle Ages (e.g. Ockham). In the Seventeenth century several proposals were made to construct a language whose terms would be well defined, and moreover, whose form would transpire its semantics. The best example of this approach is the unique work by John Wilkins (1668) who published a thesaurus, structured as a taxonomy of about 2,000 well-defined and cross referenced concepts, starting from 40 main categories. In fact we can see this work as far more in line with current ontology engineering than philosophical ontology, and the book (hard to get) could still serve as a useful source of inspiration for constructing top-ontologies. Its influence in philosophy was very small, but it inspired the construction of thesauri, in particular Roget's thesaurus (1852) and it is one of the origins of information science. According to Wilkins an important cause for the confusions inherent to natural language was the fact that words were arbitrary labels for concepts: their morphology did not indicate the underlying semantics. Given his taxonomy new words could be composed by descending a path in the taxonomy. For instance, "zita" would stand for "dog", composed from z for "animal", i for "beast", t for "canine" and a for "dog".[2] Leibniz immediately ordered this book, but instead of creating semantically transparent words, he proposed to use numbers which could be automatically processed by mechanical calculators, using combinatorics. This *calculus ratiocinator* is based upon prime numbers and is binary, previewing computer cryptography. Leibniz himself improved Pascal's calculator to include division, and concluded farsighted, but overly optimistic:

> "Once the characteristic numbers of most notions are determined, the human race will have a new kind of tool, a tool that will increase the power of the mind much more than optical lenses helped our eyes, a tool that will be as far superior to microscopes or telescopes as reason is to vision" (G. Leibniz, Philosophical Essays)

In summary: current views and practice in ontology engineering have more affinity with these *Lingua Universalis Philosophica* views on taxonomies of concepts and mechanical reasoning, than with Ontology.

---

[1]Paradoxically formal ontology is sometimes called 'applied ontology'! see Wikipedia under wikipedia.org/wiki/Applied_ontology

[2]The reader is referred to (Ecco 1997) for a detailed review.

4.2.3 Artificial Intelligence

In a classic paper McCarthy and Hayes (1969) discuss the definitions of a small set of central concepts of common sense: causation, fluents, action and its epistemological consequences. This work was further taken up in an influential paper—"naive physics manifesto"—by Pat Hayes (1985) in which he argued that in common sense views of the physical world small sets ("clusters") of concepts formed the basic ontology for interpreting physical events. Although this proposal was never taken up in its full consequence, in the beginning of the 80-ies reasoning about physical systems became a major research issue. In Qualitative Reasoning (QR) (Weld and de Kleer 1990) and in Model Based Reasoning (MBR) (Hamscher et al. 1992) it is necessary that the system is able to interpret situations and to reason from the structure of these situations and its current states what will or can happen, i.e. to make predictions about its behavior. To be able to do so, it should be equipped with basic notions of its physical domain (e.g. gases, containers, and temperatures, or electricity, connections and components). These basic concepts should be defined in such a way that the relevant properties of the entities involved could be inferred. Therefore, the knowledge base of a QR/MBR systems contains an ontology (which is not necessarily well separated from the rest of the knowledge) (Forbus 2008). Ontological knowledge is used to identify what instances (individuals) are, and to infer properties and other relationships with other individuals to be able to interpret a situation. Ontologies in AI are therefore constructed to perform reasoning. Note that in natural language processing, the meaning of words has a similar status. However, in Section 4.3 we will point out that there is a subtle but important difference between word senses in natural language and concepts used in predicting situations.

4.2.4 Knowledge Engineering

The most direct ancestor of ontology engineering is the knowledge engineering community. In methodologies for building knowledge systems, in particular CommonKADS (Schreiber et al. 1993), ontologies were used as specifications of domain terminology, i.e. a conceptual documentation for constructing a domain knowledge base (Van Heijst et al. 1997). Languages were constructed to handle these ontologies in machine manageable form: CML[3] and Ontolingua (Gruber 1993). The latter language was based upon KIF (Knowledge Interchange Format; now: CommonLogic), a machine readable version of predicate logic, intended to be used as a translation interlingua for various Description Logic (DL) based knowledge representation (KR) systems. Casting ontologies in Ontolingua also enabled reuse of domain ontologies via the publicly available Ontolingua repository. A side effect of these developments was that these ontologies could very well be used in supporting knowledge and information management (Schreiber et al. 2000).

[3]Conceptual Modeling Language for CommonKADS (Breuker and Van De Velde 1994)

4.2.5 Semantic Web

A real boost to the development of ontologies occurred when Tim Berners-Lee (1999) was "dreaming" aloud and proposed to extend the current Web with semantic support:

> "The first step is putting data on the web in a form that machines can naturally understand, or converting it to that form. This creates what I call a Semantic Web – a web of data that can be processed directly or indirectly by machines" (Berners-Lee 1999: 191)

Ontologies were supposed to provide these semantics that should enable this machine-understanding, and a new, state of the art Web Ontology Language (OWL) was designed (Bechhofer 2004). OWL is a DL based knowledge representation language, carefully conceived to enable decidable reasoning (Horrocks et al. 2003).

From this overview of the roots of ontologies, it becomes understandable why we have such a large variety of kinds of ontologies, and that there is so little agreement about what ontologies are or should be. First, the aims differ considerably. Some are aimed at reasoning (see e.g. (van de Ven et al. 2008) for legal reasoning); most others are for information management and retrieval, in particular for (Semantic) Web applications. Second, is the re-usability perspective. The most abstract and general (upper-, top- or foundational) ontologies are supposed to be applicable to almost any domain. Philosophical ontology is here often the main source of inspiration (e.g. DOLCE (Gangemi et al. 2002), SUMO (Pease and Niles 2002)). Core ontologies capture some clusters of concepts that are typical for some field of activities, such as medicine or law (Hoekstra 2007), which makes them re-useful for domain ontologies of the field. Third, there is a large range of formalisms into which ontologies are cast, varying from typical KR formalisms (e.g. OWL-DL, Topic Maps, RDFS); to interchange formalism (KIF, CommonLogic), and to conceptual specification languages (UML, CML). However, as we will see in Section 4.4, the choice of the language is in many cases not motivated by the job the ontology is supposed to perform. Fourth, the ontologies differ much in detail of representation. Some are no more than class-subclass structures of terms, while others have dense specifications of (structured) properties and relations.

4.3 Knowledge and Semantics

In the overview of the roots of ontology engineering, ontology appears to capture successively: reality, knowledge and/or semantics. All three are closely interrelated. To identify or to belief that certain categories refer to reality we must know about these. Semantics reflect our understanding of expressions, and this understanding is based upon our knowledge and beliefs in general. However, we must make an important distinction between knowledge and semantics, i.e. between what we know and belief about some term (concept) and what a term means in some expression. We can illustrate this by the different senses that a simple, rather concrete term like "car" can take. We know lots of things about cars (and maybe even more about individual cars). However, in the following statements, the meanings of the term car are totally different:

1. Driver: “I will take my car”. (car is a vehicle)
2. Mechanic: “This car has a defective carburetor”. (car is a device)
3. Salesman: “This car is a bargain”. (car is a commodity)
4. Insurance inspector: “This car is a wreck”. (i.e. no longer to be viewed as a car)

It is not difficult to add new perspectives, dependent on the focus of a property (e.g. energy consumer; polluter, status symbol, killer, etc). The properties that are *relevant* with each view (superclass) of “car” have little in common. Vehicles are for transport; devices have a structure and causal behaviors; commodities imply trade; wrecks are to be recycled, etc. Therefore the semantics of terms in expressions are the result of *using* knowledge in a specific context. These contexts may even bring “logical” conflict with what we know the term can mean, as in metaphorical expressions: e.g. “This car never abandoned me”. Leaving metaphor for practical reasons aside (however important it is in our reasoning and communication, e.g. (Lakoff 1987; Pinker 2007)), we can see the semantics of a term as a subset of what we know a concept to be. This means that for instance a reusable ontology that should capture the semantics of a term in various contexts should contain multiple classifications. That is exactly *not* what current practice shows in ontologies aimed at information management of documentation, in particular when these applications have to operate on the Web, where contexts can be highly diverse. They do not recognize contexts, nor do they select specific meanings of terms. The problem is not that we do not have the technology available to do so: it is the fact that it is extremely costly and complex to construct ontologies that have multiple views on board. Multiple classification is necessary, but not sufficient to obtain this automatic machine understanding that Tim Berners-Lee was dreaming about. We will return to this issue in the next section, but first we will give a short account of the kind of knowledge stores that have evolved to enable us to come to context sensitive communication, i.e. the use of natural language.

Figure 4.1, adapted from (Levelt 1993), presents a very simplified cognitive architecture for generating natural language. The spatial and the kinesthetic representational systems are perceptual systems that interpret information from the eyes, respectively sensors that identify the positions of our body parts, including the body with respect to gravity. These two systems can integrate the information abstracted from these senses so that body movements in space can be coordinated. They include very short term, sensory memories. Many animals who have these kinds of receptors have these systems. In humans a new kind of representational system evolved that works like an interlingua between these perceptual systems: a propositional representation system. It allows us to state that we see a chair, hear footsteps, etc. With “propositional” is meant that its representations no longer contain perceptual elements. This propositional system is the human knowledge base.[4] This knowledge

[4] Another term has been: long term memory. At least three subsystems can be distinguished: episodic memory, containing past memories (instances); semantic memory (which contains generic

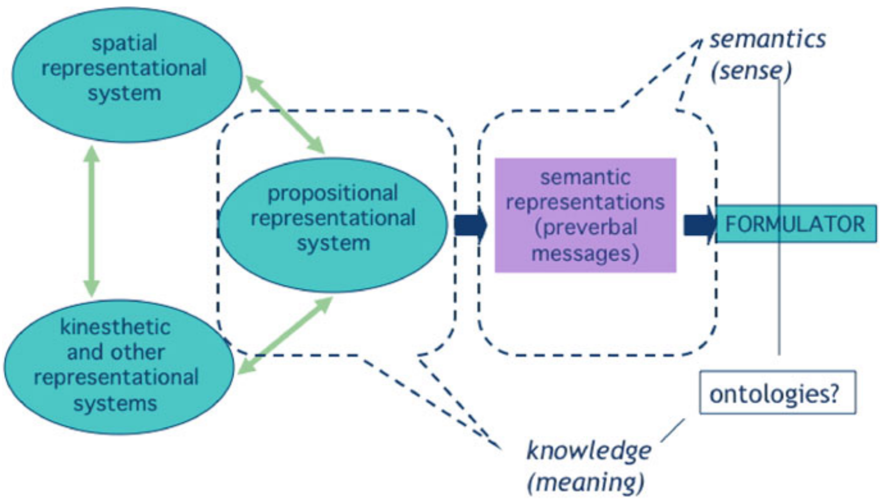

Fig. 4.1 Knowledge stores in natural language generation. After Levelt (1993)

store is used to create a "preverbal message". It is constructed from a (sequence of) speech act(s), reflecting intention; the communication content of these speech acts is derived or selected from the propositional system. This content is what semantics are about. In its simplest version they are selections from this propositional system, but often new constructs emerge, as e.g. in the use of metaphors, analogies, etc. We discuss here this architecture not only to explain how in psycholinguistics semantics are distinguished from knowledge. As it becomes more and more evident that we share this propositional store with some of our close ancestors, we may become interested in the genetic knowledge seeds which guide humans in developing their interpretations of the physical and social world. To be continued in Section 4.5.

4.4 Formalisms, Reasoning and Information Management

With a few exceptions, applications fueled by ontologies are used for information management. They are used for finding or indexing (tagging) what documents are about, or to filter data, but they are not used for reasoning. Reasoning with ontologies can take two forms. The first one is for modeling situations and for problem solving: as in QR and in MBR. The second one is to use ontology as the vocabulary in natural language processing: as in Tim Berners-Lee's proposal for the Semantic Web. This dream has created a paradox and an impasse. The ontology language, OWL(-DL), is extremely well suited for this first kind of reasoning, and moreover, one can trust blindly its reasoning capabilities. This is for

knowledge, facts and beliefs), and an associative memory that accounts for skills and other empirical contingencies. These sub-sytems work relatively independently. For instance amnesia, due to damage of specific areas of the brain, is a disturbance of episodic memory; not of semantic memory.

Ontolingua	XML	RDF/S	DAML+OIL	OWL
FBO & FO-Law	TOF, TOV & Ontology French Law	e-Court O. (RDF(S)/OIL), Multi Tier Contract O. (RDFS/DAML) & Consumer Protection Ontology (RDF/OWL)	OCL-NL, LRI-Core, IPROnto & AGO IR O.	LRI-Core (OWL-DL), LOIS, IPROnto (OWL-DL), Copyright O. (OWL-DL), BEST-User O., LKIF-Core (OWL-DL), SIAP Legislative O., OnProc (OWL-DL), Ontolegis & Ontojuris, Patent O., Ontology Fundamental Legal Concepts (OWL-DL), US UCC, Ontology Legal CBR (OWL-Lite), Oral Hearing O., Greek Public Adm. O. (OWL/OWL-S) & ALIS IP O.

Fig. 4.2 Legal ontologies and the representation formalisms used, copied from (Casellas 2010)

instance of importance for systems that have to operate autonomously, as e.g. in Deep Space. However, enabling machines to "understand" documents, which will involve usually some natural language processing, will always have a heuristic and "incomplete" flavor for which we do not require this trust on complete, consistent reasoning. This is important, because OWL has a limited expressivity that brings all kinds of problems. For instance identifying individual entities as being identical can only be approximated using tricks (Hoekstra and Breuker 2008; Hoekstra 2009). Maintaining identity in natural language processing (NLP) is a major issue (e.g. context bound unique naming; pronoun reference assignment). Full NLP is still too hard a nut to crack, especially in the highly heterogeneous environment of the Web. Attempts to have some superficial approximation to this problem (e.g using case structures) have shown minimal improvements in the efficiency of retrieval. OWL may not be the representation format for entering this information management arena, it is highly popular for representing ontologies for information management, as Fig. 4.2 shows for an almost complete list of published legal ontologies.[5] A reason to cast these information management ontologies in OWL can be that it allows one to do consistency checking, e.g. to find that properties of a subclass conflict with those of a superclass. However, real logical consistency checking requires a strong commitment with (many) properties which provide sufficient definitions. That is what information management ontologies lack in general. Consistency checking should be a necessary verification step for top ontologies, as it is very important

[5] As DAML+OIL, and RDFS/OWL are in fact (precursors to) OWL, we can state that except for Ontolingua (history) and XML all these ontologies are OWL based.

that no inconsistencies can propagate to lower levels when they are used to structure domain ontologies. By their philosophical ancestry many top ontologies are written in KIF/CommonLogic or Ontolingua. However, for these systems no complete consistency checking is available, and in translating these into OWL-DL most properties get lost in translation.

4.5 A CS Perspective for Top Ontologies

Top ontologies have two roles. The first one is to provide a structure of classes where a domain ontology can hang its main concepts as subclasses. It gives a starting position for more detailed modeling. The second role is even more interesting: by inheritance one gets already defined properties for free, which helps in consistency checking. Moreover, top structures may imply "special inference", i.e. inference rules that are bound to some set of terms. For instance, by defining a set of terms for spatial positions, such as "left-of (A,B)" is equal to "right-of (B,A)" etc. one can easily generate all tautologies. However, current practice does not show such "reuse" benefits. There are many reasons such as the chosen formalism (see above); little commitment in specifying properties; too few resources, and finally little agreement and much argument about the often slowly emerging points of departure. The history of the SUMO top ontology is probably the best worked out case, because its development started about a decade ago under the blessings and ambitions of IEEE. SUMO originally was SUO, and stood for Standard Upper Ontology. Meetings and mailing showed profound disaccord under the participants, and it went as far that some voting was proposed. This was a reason for many to leave this arena. A core of researchers still involved grabbed most of the available (pieces of) top ontology and integrated these in SUMO, where the S no longer stands for Standard, but for the less authorative Suggested. The M stands for Merged. Problem in most of these debates was that they had the same flavor (and also lots of the content) of the age old philosophical debates. It was not only philosophical. It also included practical arguments for powerful representations. In our opinion, these debates could continue for ever as there were no or hardly any empirical criteria available to decide on issues; the more because the idea was that these abstract terms were neutral with respect to any application. We will argue here for taking a cognitive science perspective. This provides not only rich resources for empirical grounding but is highly appropriate when it concerns human knowledge, and in particular: common sense.[6] For instance, by far the majority of what is available on the Web. We all have common sense, so it seems that anyone can participate in this enterprise. That

[6]It is hard to state what is not constrained by the primitive dispositions to interpret the world on which human builds in order to act effectively. However, we have obtained a more "conscious", "rational" way to interpret the world by delayed reflection. This has for instance ended in models of reality which have become even inconceivable, as in quantum theory. See for instance on 'physical ontology: http://en.wikipedia.org/wiki/Physical_ontology. Also core ontologies with a highly specialized technical domain may make shortcuts, see e.g. (West 2004).

is a major thesis in the development of folksonomies. It is a naive view because we have no direct access to the structure of concepts that make up this top. We even do not know whether this top is coherent. We can only reconstruct this by trying to trace its development: not only in humans, but over evolution as it has become apparent that we share many basic concepts with (higher) animals, and that these are not explicitly laid down in some ready made form, but rather have the character of instincts: i.e. they need experiences and cultural parametrization to get a specific shape. This is typically the area of evolutionary psychology which draws heavily on experimental work in developmental psychology, etiology and anthropology. A good example of this work can be found in Pinker's book on the "stuff of thought" where he argues that (and how) concepts of space, time and causation make up these conceptual, heritable predisposition (Pinker 2007). We share notions about agent causality (i.e. actions vs. processes), intention, power and fairness. Most of his arguments come from use of language. An even larger range of concepts and even stronger anchored in empirical research over the aforementioned disciplines can be read in Marc Hauser's book on the emergence of moral beliefs, which also throws a perspective on notions of justice in the legal domain (Hauser 2006). What these studies reveal is how a number of concepts evolve, both in an evolutionary as in a cultural way. They show that unexpectedly early distinctions are made for instance between a physical world and a social-psychological world on the basis of observing moving objects. Objects that move "by themselves" are agents. Directions of movement are indications of intention, etc. Note that the number of concepts that must be assumed to be driving these identifications already represent a substantial top-structure.

However, the interpretation of these findings for constructing a top ontology must be taken with care, as they do not show how later, cultural influences may extend or modify the nature of these concepts.[7] A good example is the introduction of the notion of "energy" in our common sense vocabulary due to our insights in and experiences with electricity as a source for heat, light and movement. Energy unifies the original notions of power/impuls for movement, fire for heat and light, etc. In the notion of "energy" we have apparently found a new "superclass".

On the other hand, the distinction between physical and agent causation (processes vs. actions) is extremely well maintained, even if the distinction is sometimes hard to assess in concrete cases (e.g. guilt in criminal cases; football offenses that merit a card or not). Mental concepts, like belief, knowledge, intention and emotion are also part of the early repertory of babies: they constitute a "Theory of Mind" which attributes intentions and knowledge to others, but ultimately is also the source of attributing these also to one-self to explain ones own behavior. This wealth of insights has been a leading guide in the construction of the LKIF-Core ontology for legal domains, where a physical and a mental world form starting points (Breuker and Hoekstra 2004; Hoekstra et al. 2007; Breuker et al. 2007).

[7] Thus far we have not read about cases where such a concept gets "overwritten" by new information in the common sense domain.

A good example of how a CS perspective leads to different representations than in other ontologies (in particular: SUMO) is the conceptualization of space. During the construction of SUMO, arguments were raised whether space should be represented “3D”, i.e. having three dimensions, or 4D which would include time, to provide for histories in the spatial world. There is a special 4D version of SUMO. In fact, we cannot conceptualize a 4D world straightforward: we have to leave out one dimension and imagine a kind of flatland moving forward in time. This means that a 4D representation is not a direct representation of common sense: it is an abstract, geometrical rendering of the fact that both time and space can be viewed as having dimensions. However, from a human cognitive point of view space is not a uniform concept. There are two kinds of spaces. (1) The space that physical objects take. The basic concept here is “size”, which can be big, small, etc. This concept may be equated with “volume” (e.g. relevant for fluids) so that a 3D picture may finally emerge. (2) This is different from, but compatible with viewing space as marking positions of objects. Space is here not outer space, but on earth, which is a floor or landscape on which objects are stacked. This is the consequence of the fact that we live in a largely flat world permanently under the influence of gravity. Left-of/right-of plays here a minor role and explains why we get easily confused in discriminating between these. It also explains why we have the illusion that a mirror image switches left–right, where it in reality switches the front-back “dimension” of the observer. These simple facts (gravity, eyes “in front” of our head) have given rise to concepts of positions in space (for an extensive CS literature on this subject, see McManus (*nomen est omen?*) book on “left hand, right hand”. (McManus 2002). In other words, 3D is applicable to the extension of objects; for marking positions we have a 2D “map” view on which objects rest. As a consequence we should have two kinds of ontologies for space. One about size, and one about position. By performing measurement we can relate these.

4.6 Some Paradoxes for Conclusions

Ontology engineering is not a uniform field of research and development. and it is full of paradoxes. The first one is of course that where ontology is aimed at providing clear definitions, it has not succeeded in defining itself. Philosophical ontology still plays an important role, while a much closer ancestor in (“natural”) philosophy—the *Lingua Universalis* movement—aimed at taxonomic and mechanized reasoning is forgotten and has only an indirect impact via Information Science. Philosophers who could benefit from automated consistency checking use formalisms that do not really allow for this. On the other hand, applications aimed at information management, e.g. searching the web, use very often formalisms which have traded correct reasoning for expressivity, while these applications do not require reasoning. A major paradox is the fact that in developing top ontologies that claim to cover common sense, the main sources of inspiration are introspection and philosophical analyses. No reference is made to the empirical data available from Cognitive Science, which are far more valid than those speculations. Finally, the reader may also find it paradoxical that we use in the title “legal” but hardly any reference is made to the domain of law. That is true. First, it should be noted that the domain of law is rather advanced in the world of ontologies. The fact that OWL-DL is used may seem an unnecessary complication, but it also enables this community to move effortlessly to applications that need trustworthy automated reasoning. An example is HARNESS, a legal assessment system that runs completely in OWL-DL (van de Ven 2008). Second, in the core concepts of law, many common sense concepts are hidden, as cognitive studies on moral and legal concepts show (Hauser 2006).

References

Bechhofer, S., F. van Harmelen, J. Hendler, I. Horrocks, D.L. McGuinness, P.F. Patel-Schneider, L.A. Stein. OWL Web Ontology Language Reference. W3C Recommendation, World Wide Web Consortium, February 2004. URL http://www.w3.org/TR/owl-ref/. In M. Dean, G. Schreiber (Eds.).

Berners-Lee, T. (1999). *Weaving the Web: The Past, Present and Future of the World Wide Web.* Orion Business Books, London.

Breuker, J., R. Hoekstra. (2004). Core Concepts of Law: Taking Common-Sense Seriously. In A.C. Varzi, L. Vieu (Eds.) *Proceedings of Formal Ontologies in Information Systems (FOIS-2004*. IOS-Press, Torino, Italy, 210–221.

Breuker, J., W. Van De Velde (1994). *CommonKADS Library for Expertise Modeling: Reusable Problem Solving Components.* IOS-Press/Ohmsha, Amsterdam/Tokyo.

Breuker, J., R. Hoekstra, A. Boer, K. van den Berg, R. Rubino, G. Sartor, M. Palmirani, A. Wyner, T. Bench-Capon (2007). OWL Ontology of Basic Legal Concepts (LKIF-Core). Deliverable 1.4, Estrella, www.estrellaproject.org.

Casellas, N. (2010). *Legal Ontology Engineering.* Springer, New York, NY.

Ecco, U. (1997). *The Search for the Perfect Language.* Blackwell Publishers, Oxford.

Forbus, K. (2008). Qualitative Modeling. In F. van Harmelen, V. Lifs-chitz, B. Porter (Eds.) *Handbook of Knowledge Representation.* Elsevier, San Diego, CA, 361–394.

Gangemi, A., N. Guarino, C. Masolo, A. Oltramari, L. Schneider. (2002). Sweetening Ontologies With DOLCE. In A. Gomez-Perez, V.R. Benjamins (Eds.) *Proceedings of the EKAW-2002*. Springer, Heidelberg, 166–181.

Gruber, T.R. (1993). A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5: 199–220.

Hamscher, W.C., L. Console, J. de Kleer (Eds.) (1992). *Readings in Model-Based Diagnosis.* Morgan Kaufmann, San Mateo, CA.

Hauser, M. (2006). *Moral Minds: The Nature of Right and Wrong.* Harper Collins, New York, NY.

Hayes, P.J. (1985). Naive Physics I: Ontology for Liquids. In J.R. Hobbs, R.C. Moore (Eds.) *Formal Theories of the Common Sense World*. Ablex Publishing Corporation, Norwood, 71–108.

Hoekstra, R. (2009). *Ontology Representation: Design Patterns and Ontologies That Make Sense.* IOS Press, Amsterdam.

Hoekstra, R., J. Breuker (2008). Polishing Diamonds in OWL 2. In A. Gangemi, J. Euzenat (Eds.) *Proceedings of the 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2008),* LNAI/LNCS. Springer, Berlin, 64–73.

Hoekstra, R., J. Breuker, M. Di Bello, A. Boer (2007). The LKIF Core Ontology of Basic Legal Concepts. In P. Casanovas, M.A. Biasiotti, E. Francesconi, M.T. Sagri (Eds.) *Proceedings of the Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT 2007)*, June 2007. Stanford University, Stanford. CEUR Workshop Proceedings.

Horrocks, I., P.F. Patel-Schneider, F. Van Harmelen (2003). From *shiq* and RDF to OWL: The Making of a Web Ontology Language. *Journal ofWeb Semantics: Science, Services and Agents on the World Wide Web Semantics*, 1: 7–26.

Lakoff, G. (1987). *Women, Fire and Dangerous Things.* University of Chicago Press, Chicago, IL.

Lenat, D.B. (1995). CYC: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM*, 38(11): 33–38.

Levelt, W.J.M. (1993). *Speaking.* Cambridge University Press, Cambridge, MA.

McCarthy, J., P. Hayes. (1969). Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer, D. Michie, M. Swann (Eds.) *Machine Intelligence* vol. 4. Edinburgh University press, Edinburgh, 463–502.

McManus, C. (2002). *Right Hand, Left Hand.* Weidenfeld and Nicolson, London.

Pease, A., I. Niles. (2002). IEEE Standard Upper Ontology: A Progress Report. *Knowledge Engineering Review*, 17: 65–70. Special Issue on Ontologies and Agents.

Pinker, S. (2007). *The Stuffofthought.* Allen Lane, London.

Schreiber, A.Th., B.J. Wielinga, J.A. Breuker (Eds.) (1993). *KADS: A Principled Approach to Knowledge-Based System Development,* vol. 11 of *Knowledge-Based Systems Book Series.* Academic Press, London. ISBN 0-12-6290407.

Schreiber, G., H. Akkermans, A. Anjewierden, R. de Hoog, N. Shadbolt, W. Van den Velde, B. Wielinga (2000). *Knowledge Engineering and Managament: The CommonKADS* Methodology. MIT Press, Cambridge, MA.

Sowa, J.F. (2000). *Knowledge Representation: Logical Philosophical, and Computational Foundations.* Brooks Cole Publishing Co, Pacific Grove, CA.

van de Ven, S., J. Breuker, R. Hoekstra, L. Wortel (2008). Automated Legal Assessment in Owl 2. In *Legal Knowledge and Information Systems. Jurix 2008: The 21st Annual Conference,* Frontiers in Artificial Intelligence and Applications. IOS Press, Amsterdam, December 2008.

Van Heijst, G., A.Th. Schreiber, B.J. Wielinga (1997). Using Explicit Ontologies for kbs Development. *International Journal of Human-Computer Studies*, 46(2/3): 183–292.

Weld, D., J. de Kleer (Eds.) (1990). *Readings in Qualitative Reasoning About Physical Systems.* Morgan Kaufman, San Mateo CA.

West, M. (2004). Some Industrial Experiences in the Development and Use of Ontologies. In *Proceedings of EKAW-2004 workshop on Core Ontologies*, Whittlebury Hall, Northamptonshire, UK.

Wilkins, J. (1668). *An Essay Towards the Real Character and a Philosophical Language.* Gellibrand, London.