

Ontologies and Description Logics

Diego Calvanese

Nicola Guarino

1 Ontologies and Description Logics: So Close, So Far

Ontologies and Description Logics (DLs). Why putting them in a single chapter for this document devoted to AI research in Italy? Well, a good reason is that both research areas originated from a single, classic mainstream in AI: logic-based knowledge representation (KR). Indeed, starting from the late 80s, Italian researchers brought a radical contribution to the formal understanding of so-called structured representation languages developed in the previous decade, whose most known example was KL-ONE. Their research took however two different directions.

On one hand, after Levesque and Brachman's logicalization of KL-ONE and their discovery of the fundamental tradeoff between expressivity and tractability [8], a whole Italian school emerged under the direction of Maurizio Lenzerini, focusing mainly on the need to develop KR languages with predictable computational behavior and sound and complete inference algorithms.

On the other hand, Nicola Guarino started analyzing the implicit and often ambiguous ontological assumptions made by structured logical formalisms, such as those behind the notions of role and attribute, arguing against the ontological neutrality of representation systems and suggesting formal ontological distinctions to be taken into account.

In the meanwhile, in the early 90s, a new hot topic captured the attention of AI researchers: knowledge sharing among heterogeneous, distributed sources (very similar to the issue emerging –roughly at the same time– in the DB community: data integration). It became clear soon that, in order to attack these problems, two complementary aspects needed to be addressed: content and reasoning. Ontological analysis and DLs emerged soon as the proper techniques to deal with these aspects, and became later very popular for their key role in the future Semantic Web. Italian researchers contributed in a peculiar way to both.

2 The Italian Way to Ontologies

In the philosophical sense, an ontology is a system of categories accounting for a certain vision of the world. As such, this system does not depend on a particular language. On the other hand, in its most prevalent use in

AI, an ontology refers to an engineering artifact, consisting of a specific vocabulary containing the terms used to describe a certain domain, plus a set of explicit assumptions regarding the intended meaning of vocabulary words. This set of assumptions has usually the form of a logical theory, where terms correspond to unary or binary predicates, respectively called concepts and relations. In the simplest case, an ontology describes a set of concepts related by taxonomic relationships; in more sophisticated cases, suitable axioms are added to express other relationships between concepts and to constrain their intended interpretation. Referring to the classic KR literature, an ontology plays the role of a TBox, complementary to the ABox, where the former contains knowledge about Terms, and the latter about Assertions [3, Ch. 2].

The specific contribution of Italian AI researchers to the establishment of modern research on ontologies pivots around two main claims; the need to focus squarely on content (as independent of reasoning issues), and the intrinsic interdisciplinary nature of ontological analysis.

Indeed, oddly enough, and apart very few notable exceptions, classic AI research paid little attention to content issues: emphasis was much more on modeling the nature of reasoning and expertise rather than on modeling the inherent structure of a domain.

In KR, such tendency was especially evident within the so-called logicist approach: in their well-known textbook on AI, Genesereth and Nilsson explicitly stated the “essential ontological promiscuity of AI”, and devoted just a couple of pages to the issue of conceptual modeling, admitting however it was still a serious open problem.

On the other hand, in the area of knowledge acquisition, the interest was concentrated on modeling the agent's behavior (i.e., the problem-solving expertise) rather than its own environment (the problem domain), since the latter was seen as strongly dependent on the particular task at hand.

In this context, Guarino's contribution was to defend the need of a principled way to model the inherent structure of a domain (independently, as much as possible, on the particular task at hand), as the only possibility to attack the knowledge sharing problem at its roots. It was immediately clear that the key for establishing such a “principled way” was a radically interdisciplinary approach: besides the basic tools of logic and computer science, an open-minded aptitude towards the subtle distinctions of philosophy and the intricate issues of natural language and commonsense was absolutely necessary.

Such an interdisciplinary attitude was basically lacking at that time (early 90s). Indeed, the three major proposals addressing in a general way the problems of real world modelling, namely [23, 24, 15], were suffering from a relatively narrow perspective (with the notable exception of [30]). More or less, they were concentrating on the immediate needs of the AI practice, refusing to take (explicitly) into account the achievements coming from, for example, analytic philosophy or formal language semantics. For instance, Pat Hayes was (and still is) rather skeptic about the potential relevance of philosophical results, and Doug Lenat, although insisting on the need to *bite the bullet* addressing core ontological issues, adopted in practice a brute-force approach, with no systematic, principled methodology.

In contrast to this *ad hoc* approach, Italian researchers becoming attracted by the ontology field were aiming at a more systematic, principled methodology. Such generic desire (typical of the European school) was dramatically boosted thanks to an extremely fortunate circumstance: the presence in Italy of a few brilliant philosophers who, although with different motivations and background, were developing a renovated interest in the old issues of metaphysics, and got in touch with AI researchers, showing concrete motivations towards serious interdisciplinary cooperation. The list includes philosophers of language such as Pierdaniele Giaretta and Diego Marconi, as well as representatives of the continental tradition like Liliana Albertazzi and Roberto Poli, and, mostly relevant, logic-oriented “ontologists” such as Roberto Casati and Achille Varzi, who brilliantly contributed to introducing philosophical ontology to the large public [13, 32, 12].

One of the earlier results of this cross-fertilization was the first International Workshop on Formal Ontology in Conceptual Modeling and Knowledge Representation, jointly organized by Nicola Guarino and Roberto Poli in Padova in 1993. For the first time, researches with backgrounds as diverse as philosophy, cognitive science, linguistics, AI, databases, object-oriented programming, met together with the purpose to develop a comprehensive, principled approach to ontological analysis.

The radical novelty of this interdisciplinary approach had a deep impact on the international community. A number of influential papers were published, which gradually helped understanding the role of principled, rigorous ontologies for information systems. Times were ripe for the first international conference on Formal Ontology in Information Systems (FOIS’98), co-located with KR’98 in Trento, followed by the second one in the U.S. in 2001 and a third one back in Italy, in Torino, in 2004. The next one will be again in the U.S. in November 2006.

Another major sign of the international role played by the Italian community in ontology research is the recent publication of a new journal, Applied Ontology, published by IOS Press and jointly edited by Nicola Guarino

(ISTC-CNR) and Mark Musen (Stanford). This is the first journal exclusively addressing ontological analysis and conceptual modeling under an interdisciplinary view, focusing on information content in its broadest sense. Various Italian researchers from different disciplines appear in the Editorial Board.

2.1 The Current Days: An Active Interdisciplinary Community

In the recent years, the popularity of ontologies has literally exploded, mostly due to the widespread excitement on the Semantic Web. While increasing in popularity, the very notion of ontology acquired a broader, often vague or debatable interpretation, despite the original attempts to establish rigorous distinctions [21]. Many people consider now ontologies, conceptual schemes, web directories, contexts, knowledge bases, lexical resources, taxonomies as more or less synonymous terms, and the research on core content modeling issues is strictly intertwined with complementary aspects such as content extraction based on statistical methods (ontology learning and ontology population), as well as ontology matching and alignment. Moreover, a substantial amount of research is directed towards the development of ontology languages and reasoning algorithms, cooperative modeling tools, and presentation techniques. In Italy, application ontologies are being built in multiple areas, including biomedicine, public health, law, banking, insurances, museums, digital libraries, architecture, manufacturing, public administration, software services, geography, food and agriculture.

As a result of this explosion of activities, sketching an accurate picture of current research in Italy in the broad area now called ontological engineering is not feasible in a few paragraphs. In the following, we will focus on the core issues concerning content modeling, briefly reporting the state of the art and paying special attention to showing the mutual synergies of different disciplines.

2.2 Foundational Issues

At the foundational level, the issues at hand concern first of all the very status of ontology as a philosophical discipline: what is it about? What are its theoretical tools? Which are the main positions about a basic system of categories and relations? How can we compare them? An excellent introduction to these questions is the recent book by Achille Varzi [32], which, among other things, shows the important contribution of Italian philosophers to the contemporary debate.

In the AI perspective, the counterpart of such foundational work is the development of a library of axiomatic theories (so-called foundational ontologies), systematically related to each other in a way that makes the rationales and alternatives underlying different ontological

choices as explicit as possible. This was the perspective of the European project WonderWeb, which, besides contributing to the wide adoption of the OWL language, resulted in the development of the DOLCE ontology by ISTC-CNR¹, adopted as a reference upper ontology in many projects around the world.

2.3 Ontology Development Methodology

How to build an ontology? How to choose a useful system of basic concepts and relations? How to link it to an applications vocabulary? How to elicit hidden ontological assumptions? Unfortunately, there is no systematic answer to these questions yet. The ontology community is simply not mature enough to “distill” a real, usable methodology out of the current, impetuous research work. Of course, various approaches are emerging, mainly bound to European projects such as WonderWeb, KnowledgeWeb, Interop or NeOn, but besides extending classic software engineering methodologies to deal with ontologies, still little practical help is offered to the ontology analyst from the content modeling point of view, although some preliminary methodology proposals based on solid theoretical principles have achieved considerable impact², and approaches based on re-using ontology design patterns [19] allow to simplify the problem of building an initial ontology for specific purposes. It would be worthwhile mentioning, in this context, that the methodological problems of ontology development go much beyond the Semantic Web perspective, as they have a deep impact on the general practice of information systems conceptual modeling [22].

2.4 Ontology Comparison and Evaluation

Ontology comparison and evaluation is certainly one of the hottest areas in ontology research today. Due to the growing number of available ontologies (often developed in a hastily and ad hoc way), suitable techniques need to be developed in order to make possible to evaluate such ontologies with respect to the user’s needs, possibly merging and adapting existing resources, or establishing mappings between different ontologies (see the Ontology Alignment Evaluation Initiative³). Moreover, from a more theoretical point of view, there is the growing need to establish general roadmaps of basic ontological choices, exploring their differences and their similarities, and ultimately providing ways to make explicit the *design rationale* of a certain ontology (one of the goals of the NeOn project).

¹<http://www.loa-cnr.it/DOLCE.html>

²<http://www.loa-cnr.it/Ontologies.html>

³<http://oaei.ontologymatching.org/>

2.5 Ontology and Language

As far as content modeling is concerned, there are two broad areas where ontology and linguistics are strictly related to each other: lexical resources and formal language semantics. From the ontological point of view, the most sophisticated lexical resource (yet rather limited in coverage) developed in Europe is perhaps the SIMPLE lexicon resulting from a European project led by ILC-CNR. Various groups in Italy have worked on semantic lexical resources more or less linked to WordNet, but unfortunately there is no general freely accessible de facto standard for the Italian language yet. Research is ongoing at ISTC-CNR, in cooperation with the WordNets developers, in order to re-visit and improve WordNets semantic structure in light of well-founded ontological principles, and interest is emerging in merging WordNet and FrameNet with upper-level ontologies such as DOLCE or GUM⁴.

Concerning formal language semantics, many issues have also an ontological impact, or can benefit from ontological analysis. This is certainly the case of the classic themes of philosophy of language and linguistics, such as the nature of universals and predication, the status of propositions and states of affairs, the analysis of mental attitudes, the formal treatment of tense and time, plural quantification, masses, kinds, metonymy, reference, and anaphora. Finally, from the side of applied Natural Language Processing, the recent book by Niremburg and Raskin [26] is a clear example of the increasing synergy between ontology and linguistics.

2.6 Further Research Prospects

In conclusion, we would like to mention two new promising research topics emerged in Italy in the latest years, which complete this brief sketch of core ontological research in Italy.

The first topic concerns the relationships between ontology, cognition, and perception. On the one hand, this involves issues such as the ontological status of mental states and emotions, or the ontological assumptions behind visual recognition; on the other hand, there is the whole field of *constructivist ontology*, which postulates the existence of a rich realm of constructed entities in addition to those belonging to the ground level of so-called “reality”.

Among constructed entities, the second emergent topic focuses on those belonging to what Searle called the social reality [29], insofar they are socially constructed. The ontology of social reality is indeed an active area of research in Italy, addressing various issues ranging from the ontology of roles and organizations [25] to that of documents and laws [7].

⁴<http://www.fb10.uni-bremen.de/ontology/>

3 Description Logics

As mentioned, DLs [3] stem from the effort started in the mid 80s to provide a formal basis, grounded in logic, to formalisms for the structured representation of knowledge that were popular at that time, notably Semantic Networks and Frames [9]. The fundamental work by Brachman and Levesque [8], initiated this effort, by showing on the one hand that the full power of first-order logic is not required to capture the most common representation elements, and on the other hand that the computational complexity of inference is highly sensitive to the expressive power of the KR language. Research in DLs up to our days can be seen as the systematic and exhaustive exploration of the corresponding tradeoff between expressiveness and efficiency of the various inference tasks associated to KR.

DLs are based on the idea that the knowledge in the domain to represent should be structured by grouping the objects of interest with common properties into classes (called *concepts*), and explicitly representing those properties through (typically binary) relationships among classes (called *roles*). Concepts and roles are constructed by making use of various constructs, and it is precisely the set of allowed constructs that characterizes a DL language. The domain of interest is then represented by means of a DL *knowledge base* (KB), where a separation is made between general intensional knowledge about concepts and roles, stored in a so-called *TBox* (for “Terminological Box”), and specific knowledge about individual objects in the modeled domain, stored in an *ABox* (for “Assertional Box”).

Several reasoning tasks can be carried out on a DL KB, where the basic form of reasoning involves computing the subsumption relation between two concept expressions, i.e., verifying whether one expression always denotes a subset of the objects denoted by another expression. More in general, one is interested in understanding how the various elements of a KB interact with each other in an often complex way, possibly leading to inconsistencies that need to be detected, or implying new knowledge that should be made explicit.

The above observations emphasize that a DL system is characterized by three aspects: (i) the set of constructs constituting the language for building concepts and roles; (ii) the kind of assertions that may appear in the KB; (iii) the inference mechanisms provided for reasoning on the KBs expressible in the system. The expressive power and deductive capabilities of a DL system depend on the various choices and assumptions with regard to the above aspects. In the following, we analyze the various options while reviewing the state of the art in DLs research.

3.1 Development of Research in DLs

As for the inference mechanisms, research has mostly concentrated on inference according to standard set-theoretic semantics, and on devising algorithms that are both sound and complete with respect to the semantics. The aim has been on the one hand to understand the intrinsic computational properties of the various variants of DLs, and on the other hand to provide possibly computationally optimal inference algorithms for the various cases, and then implement and optimize them.

The tractability frontier. The first aspect mentioned above, i.e., the language for concepts and roles, has been the subject of an intensive research work started in the late 80s. Indeed, the initial results on the computational properties of DLs have been devised in a simplified setting where both the TBox and the ABox are empty [28, 16]. The aim was to gain a clear understanding of the properties of the language constructs and their interaction, with the goal of singling out their impact on the complexity of reasoning. Gaining this insight by understanding the combinations of language constructs that are *difficult* to deal with, and devising general methods to cope with them, is essential for the design of inference procedures. It is important to understand that in this context, the notion of “difficult” has to be understood in a precise technical sense, and the declared aim of research in this area has been to study and understand the frontier between tractability (i.e., solvable by a polynomial time algorithm) and intractability of reasoning over concept expressions. Fundamental contributions in that direction came from the research group in Rome led by Maurizio Lenzerini [16]. Not only such results were important from a theoretical point of view, but also the development and refinement of the techniques and technical tools that were used to prove them, namely tableaux-based algorithms, turned out to be fundamental and are still at the basis of the modern state of the art DL reasoning systems.

The decidability frontier. From the point of view of KR, where knowledge about a domain needs to be encoded, maintained, and reasoned upon, the assumption of dealing with concept expressions only, without considering a KB to which the concepts refer, is clearly unrealistic. Early successful DL KR systems, such as Classic, relied on a KB, but did not renounce to tractability by imposing syntactic restrictions on the use of concepts in definitions, essentially to ensure acyclicity (i.e., lack of mutual recursion). Under such an assumption, the concept definitions in a KB can be folded away, and hence reasoning over a KB can be reduced to reasoning over concept expressions.

However, the assumption of acyclicity is strongly limiting the ability to represent real-world knowledge. These

limitations became quite clear also in light of the tight connection between DLs and formalisms for the structured representation of information used in other contexts, such as databases and software engineering [11]. In the presence of cyclic KBs, reasoning becomes provably exponential (i.e., EXPTIME-complete) already when the concept language contains rather simple constructs. As a consequence of such a result, we notice a shift of research in DLs from the exploration of the tractability border to an exploration of the decidability border. The aim has been to investigate how much the expressive power of language and knowledge base constructs could be further increased while maintaining decidability of reasoning, possibly with the same computational complexity of inference. A fruitful line of research in this direction was initiated by the seminal work of Schild [27], who showed a tight correspondence between DLs and variants of modal logics and logics of programs. In several works, with essential contributions by Calvanese, De Giacomo, and Lenzerini, the correspondence was then exploited and further extended, allowing for the migration of results and techniques between the two areas, notably tight complexity bounds based on the use of finite automata on infinite trees [3, Ch. 5]. Cross-fertilization resulted also from work on modal-logic theorem proving, e.g., SAT-based decision procedures [20]. These positive results provided also the justification for the development of inference procedures based on the use of tableaux techniques [4], which, though not computationally optimal, are amenable to easier implementations, and are at the basis of the current state-of-the-art DL reasoners [3, Ch. 8].

Back to tractability. Current reasoners for expressive DLs, though based on worst-case exponential algorithms, perform surprisingly well in practice. However, such reasoners have not specifically been tailored to deal with large amounts of data, and they show their limits in those settings where data are very large and dominate the intentional level of the KB, e.g., the Semantic Web. In this context, users are provided with a conceptual view (in terms of an ontology, schema, or TBox) over data, and expect to pose complex queries over them. Unfortunately, many DLs with efficient reasoning algorithms lack the modeling power required for capturing conceptual models and basic ontology languages. On the other hand, whenever the complexity of reasoning is exponential in the size of the instances, there is little hope for effective instance management. Indeed, the only technology that is currently available to deal with complex queries over large amounts of data is the one provided by relational data management systems (RDBMS). To leverage on this technology, research groups in Rome and Bolzano have recently devised DLs that are specifically tailored to capture conceptual modeling constructs, while allowing for delegating data storage and query answering to a RDBMS [10].

Driven by applications in the medical domain and the life sciences, an alternative approach to achieve tractability of inference in the presence of arbitrary, possibly cyclic KBs has been proposed by European researchers.

Non-standard semantics and reasoning services.

Above, we have concentrated on the mainstream line of research in DLs. However, various non-standard semantics and reasoning services have been investigated in parallel, with substantial contributions by Italian and European research groups. Here, we mention only briefly Italian research that is still lively. A *fuzzy* extension of DLs, which allows for better modeling vague, as opposed to crisp properties is studied in [31]. The addition of *epistemic knowledge* allows one to model and query the knowledge that the representation system has about a domain, in addition to the domain itself [17]. A non-monotonic variant based on minimal-knowledge allows for a formal characterization of a wide variety of non-monotonic features found in practical knowledge-based applications, such as defaults and integrity constraints [18]. The non-standard inference services of *abduction* and *contraction*, finding applications in match-making scenarios, have been recently investigated in [14].

3.2 Italian DL Research Now and in the Future

Research in DLs is quite active at the moment, and the Italian community is playing a prominent role in this picture. Continuing its tradition, it is mostly laying the foundations by developing theoretical work on which to subsequently base and build implementations. The main research directions that are currently being investigated, are however dictated by practical real world requirements.

Scalability and modularization. *Scalability* of all techniques and algorithms is in general a major issue, which needs to be addressed if DLs should have a concrete impact on applications and systems. On the one hand, scalability means to study techniques and tools for handling many, possibly large interacting TBoxes. On the other hand, scalability is related to the ability to handle and reason upon large ABoxes, physically stored in databases, but considered under the open-world semantics of DLs. Tightly coupled with scalability is *modularization*, aiming at subdividing a complex DL KB into more manageable and well-interfaced parts. Such issues are currently being analyzed within the European Project TONES (Thinking ONtologiES) involving the major European groups working on DLs and coordinated by the Free University of Bolzano.

Updates and evolution of a KB. An aspect that has been largely neglected till now, with a few notable recent

exceptions, is that of updates on a DL ABox. The aim is understanding how to represent and further manipulate the updated KB.

Dealing with inconsistencies. How to deal with inconsistencies has been a prominent KR topic since its early days. However, it was put aside in favor of formalisms such as DLs, based on a clean first-order semantics, and the issue of if and how to handle inconsistencies was largely left upon the user. Since the foundational work has been carried out to a large degree, we are now in the position to take up the issue of inconsistencies again, and study and develop general-purpose rather than ad-hoc solutions.

Implementation efforts. Orthogonal to the previous points, the issue of implementing, running, testing, and deploying inference engines is of primary importance. The experience gained till now with the implementation of reasoners for very expressive DLs is quite significant, since it showed on the one hand that although for several combinations of constructs the worst-case computational complexity is the same, in practice there is a huge difference in performance. It has also shown the difficulty of the challenge to come up with implemented systems that effectively satisfy user's requirements in terms of expressiveness, usability, response times, and that the development of such systems requires huge efforts in terms of resources, and needs to be based on solid software engineering methodologies. In Italy, a significant implementation effort is being carried out at the University of Rome "La Sapienza" for the QUONTO system. Such a system is based on techniques for query rewriting according to a TBox, and leverages on relational database technology for ABox storage and query answering. At the University of Bolzano, an analogous implementation effort is being carried out for the iCOM tool, which provides a graphical front-end based on UML class diagrams interacting with a DL reasoner. At the University of Bari, Specialized reasoners for the restricted reasoning tasks associated to profile matching have also been implemented and deployed in currently used products for matching offers and request for renting apartments.

3.3 Relation with Other Research Areas and Applications

DLs have a tight connection to several research areas both within and outside of AI. As mentioned, DLs can be considered as syntactic restrictions of first-order logic (possibly with fixpoints), and also as syntactic variants of modal and program logics. We have already highlighted above the fruitful interaction between modal and program logics and DLs. We would like to point out here that also the first-order view of DLs has inspired the use of general

purpose first-order theorem provers, such as Vampire or Spass, for reasoning in DLs, by a careful tuning of the theorem-prover strategies so as to ensure termination of inference.

Natural language processing. Since their introduction, DLs bore a strong similarity to formalisms used in work on natural language, and they originally had natural language processing (NLP) as a major field for application. The use of DLs in NLP is mainly concerned with the representation of semantic knowledge that can be used to convey meanings of sentences. Extensive work has been devoted to semantic interpretation, i.e., the problem of disambiguating syntactic readings of sentences, based on semantic knowledge, and also to support natural language generation. As mentioned, it was also tightly connected to the construction of ontologies for NLP. Several large DL-based NLP projects have been undertaken, up to industrially-deployed applications, and we refer to [3, Ch. 15] for a detailed description of these.

Data management. DLs have also a rather strong connection to different aspects of data management. On the one hand, formalisms for the structured representation of information used both in databases (e.g., the Entity-Relationship model and, recently, XML DTDs) and in software engineering (e.g., UML class diagrams) admit a logical reconstruction in terms of DLs. Italian researchers started to explore such a connection in the early 90s [6, 11], and continued to give several contributions that provided cross-fertilization among the different areas [1, 2, 5].

Furthermore, DLs provide powerful mechanisms to express various forms of constraints on data sources, and such constraints can be used at query answering time to overcome possible incompleteness in the data [10]. This is particularly meaningful in the context of information integration, where data sources are inherently incomplete, and complex dependencies between data need to be expressed. Such dependencies, combined with the reasoning capabilities of DLs, are then used to determine the sources relevant to answer given queries. The use of DLs in data integration settings has been extensively explored by Italian researchers [3, Ch. 16], and has been the key factor in the European Project DWQ. In the European Project Sewasie, DLs were used also to support the user in query formulation, by proposing to users only ways of completing queries that were consistent with the conceptual schema expressed as a DL KB.

The Semantic Web, and standards of the W3C. Finally, we discuss the application of DLs to the Semantic Web, a recent, but significant development, tightly related to ontologies. Indeed, the language OWL, recently pro-

posed by the World Wide Web Consortium (W3C)⁵ as the standard language for the Semantic Web, is in its OWL-DL version a variant of a very expressive DL for which European researchers have developed recently tableaux based reasoning procedures. A new variant of the OWL-DL language is currently under proposal, that includes additional constructs for which decidability has been established, but also proposes tractable sub-languages in line with recent results, e.g., those in [10]. A further significant contribution by Italian researchers has been in the specification of the semantics for SPARQL, the Semantic Web query language currently under standardization at the W3C⁶.

References

- [1] A. Artale, F. Cesarini, and G. Soda. Describing database objects in a concept language environment. *IEEE Trans. on Knowledge and Data Engineering*, 8(2):345–351, 1996.
- [2] A. Artale and E. Franconi. Temporal description logics. In D. Gabbay, M. Fisher, and L. Vila, editors, *Handbook of Temporal Reasoning in Artificial Intelligence*. Elsevier, 2005.
- [3] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2003.
- [4] F. Baader and U. Sattler. An overview of tableau algorithms for description logics. *Studia Logica*, 69(1):5–40, 2001.
- [5] D. Berardi, D. Calvanese, and G. De Giacomo. Reasoning on UML class diagrams. *Artificial Intelligence*, 168(1–2):70–118, 2005.
- [6] S. Bergamaschi and C. Sartori. On taxonomic reasoning in conceptual design. *ACM Trans. on Database Systems*, 17(3):385–422, 1992.
- [7] G. Boella, L. Lesmo, and R. Damiano. On the ontological status of norms. In *Law and the Semantic Web*, volume 3369 of *LNCS*, pages 129–141. Springer, 2005.
- [8] R. J. Brachman and H. J. Levesque. The tractability of subsumption in frame-based description languages. In *Proc. of AAAI’84*, pages 34–37, 1984.
- [9] R. J. Brachman and J. G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9(2):171–216, 1985.
- [10] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. DL-Lite: Tractable description logics for ontologies. In *Proc. of AAAI 2005*, pages 602–607, 2005.
- [11] D. Calvanese, M. Lenzerini, and D. Nardi. Unifying class-based representation formalisms. *J. of Artificial Intelligence Research*, 11:199–240, 1999.
- [12] R. Casati and A. Varzi. *Insurmountable Simplicities. Thirty-nine Philosophical Conundrums*. Columbia University Press, 2006. Italian edition: *Semplicità insormontabili. 39 storie filosofiche*, Roma: Laterza Editore, 2004.
- [13] R. Casati and A. C. Varzi. *Holes and Other Superficialities*. MIT Press/Bradford Books, 1994.
- [14] S. Colucci, T. Di Noia, E. Di Sciascio, F. M. Donini, and M. Mongiello. A uniform tableaux-based method for concept abduction and contraction in description logics. In *Proc. of ECAI 2004*, pages 975–976, 2004.
- [15] E. Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, 1990.
- [16] F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. The complexity of concept languages. *Information and Computation*, 134:1–58, 1997.
- [17] F. M. Donini, M. Lenzerini, D. Nardi, W. Nutt, and A. Schaerf. An epistemic operator for description logics. *Artificial Intelligence*, 100(1–2):225–274, 1998.
- [18] F. M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *ACM Trans. on Computational Logic*, 3(2):177–225, 2002.
- [19] A. Gangemi. Ontology design patterns for semantic web content. In *Proc. of ISWC 2005*. Springer, 2005.
- [20] F. Giunchiglia and R. Sebastiani. A SAT-based decision procedure for *ALC*. In *Proc. of KR’96*, pages 304–314, 1996.
- [21] N. Guarino and P. Giaretta. Ontologies and knowledge bases: Towards a terminological clarification. In N. Mars, editor, *Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing*, pages 25–32. IOS Press, 1995.
- [22] G. Guizzardi. *Ontological Foundations for Structural Conceptual Models*. PhD thesis, University of Twente, The Netherlands, 2005.

⁵<http://www.w3.org/>

⁶<http://www.w3.org/TR/rdf-sparql-query/>

- [23] P. Hayes. The second naive physics manifesto. In J. Hobbs and R. Moore, editors, *Formal Theories on the Commonsense World*, pages 1–36. Ablex, 1985.
- [24] D. Lenat and R. V. Guha. *Building Large Knowledge-Based Systems*. Addison-Wesley, 1990.
- [25] C. Masolo, L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi, and N. Guarino. Social roles and their descriptions. In *Proc. of KR 2004*, pages 267–277, 2004.
- [26] S. Nirenburg and V. Raskin. *Ontological Semantics*. MIT Press, 2004.
- [27] K. Schild. A correspondence theory for terminological logics: Preliminary report. In *Proc. of IJCAI'91*, pages 466–471, 1991.
- [28] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48(1):1–26, 1991.
- [29] J. R. Searle. *The Construction of Social Reality*. The Free Press, 1995.
- [30] J. Sowa. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley, 1984.
- [31] U. Straccia. Reasoning within fuzzy description logics. *J. of Artificial Intelligence Research*, 14:137–166, 2001.
- [32] A. Varzi. *Ontologia*. Biblioteca Essenziale. Laterza Editore, Roma, 2005.