An introduction to Description Logics

Franz Baader Ian Horrocks Carsten Lutz Uli Sattler

April 21, 2006

Objective

To provide a thorough, concise introduction to Description Logics, covering the central topics and techniques of the field, namely theory and implementation, whilst sketching their use in applications.

Target Audience

The book will be suitable as a basis for a graduate (or advanced undergraduate) course on description logics and as material for self-study, and its only prerequisite will be basic knowledge of computer science that should be covered in any standard computer science programme. In particular, the book is aimed at:

- users of Description Logic tools who want to learn the background of the technology they are using,
- teachers who want to cover this topic in their courses,
- students who want to learn more about Description Logics, and
- researchers who are considering working in this area or who want to combine Description Logics with other formalisms.

Contents

General points:

- With the exception of the introduction, each chapter will conclude with a section containing a discussion of historical issues and a literature review.
- The introduction will conclude with "reading instructions" (i.e., suggestions on selective reading for various target audiences).
- We won't make extensive use of appendices, but will instead include "advanced" material in one or more sections towards the end of the relevant chapter (just before the discussion and literature review); we may want to mark these in such a way that the reading instructions can easily distinguish basic and advanced material.
- An ontology based running example will be used throughout the book.

The table of contents is detailed below, together with a summary of the material covered in each chapter. Chapters/sections that go beyond the "basic material" are marked **Advanced**.

1. Introduction

   We will give a brief overview of the history of DLs, describe their ancestors, the motivation for their development, and how the field has developed. We will also discuss ontologies/ontology languages as an important motivating application, and introduce an ontology based running example that will be used throughout the book.

The chapter will end with "reading instructions" (i.e., suggestions on selective reading for various target audiences).

2. Preliminaries

   We will (very) briefly introduce the notions and notations used in the remainder of the book. This includes sets, relations, and complexity classes. We will assume that readers have some prior knowledge of these notions, and only require a brief recapitulation; readers for whom this is not the case can either consult basic texts or skip some of the more formal/theoretical sections of the book (e.g., the chapter on complexity).

3. A Basic DL and its extensions

   (a) We explain the DL view of the world: we introduce objects, concepts, and roles on an intuitive level, using the running example.
   (b) We define the syntax and semantics of the basic DL $\mathcal{ALC}$ with general Tboxes and Aboxes, and explain them using our running example. We will also define different forms of TBoxes including acyclic ones.
   (c) We define the reasoning services commonly considered in DLs, and discuss their relationships.
   (d) We explain how these reasoning services can be used in an application. We will again use the running example.
   (e) **Advanced** We introduce important extensions of the basic DL $\mathcal{ALC}$: inverse roles, number restrictions, and concrete domains. For each, we define syntax and semantics and extend our running example accordingly.
   (f) **Advanced** This section explains the close relationship DLs and other formalisms. It will provide the reader with a deeper understanding of the material and of the field.
       i. Description Logics as decidable fragments of first order logic.
       ii. Description Logics as cousins of modal logic.
   (g) Historical issues and literature review.

4. Reasoning in DLs with Tableaux Algorithms

   In this chapter, we explain why and how the standard reasoning problems for Description Logics can be decided. We present algorithms and prove or discuss their properties in terms of soundness, correctness, and termination. For didactic reasons, we start with a very simple algorithm and proof, and will lead the reader slowly to some more sophisticated ones.

(a) We present a tableau algorithm for $\mathcal{ALC}$ ABox consistency with detailed proofs of its correctness and termination. This serves as a basis for more complex algorithms and helps the reader understand (i) the expressive power of $\mathcal{ALC}$ and (ii) what one has to do to design a reasoning algorithm and (iii) what the properties of this algorithm are in detail.

(b) We extend this algorithm to handle general TBoxes, and we extend the proof accordingly. This is not very complicated, but introduces the reader to the problem of termination and solutions to it.

(c) **Advanced** We further extend the algorithm to handle inverse roles, and we sketch how one could adapt the proofs accordingly. In contrast to the previous two cases, we leave it to the interested reader to work out the details or to consult the relevant literature.

(d) **Advanced** We further add number restrictions in a similar way: we explain how to change the algorithm in detail, but we will only sketch the modifications to the proofs.

(e) Historical issues and literature review.

5. **Advanced** Implementation issues

   This chapter is aimed at readers with implementation interests.

   Since naive implementations of the algorithms described in the previous chapter are doomed to failure, any implementation requires optimizations, a number of which have become standard in the field and will be described in this chapter.

   (a) Pre-processing of the input to reduce syntactic variations.

   (b) Classification as computing partial orders.

   (c) Avoiding expensive subsumption tests through inexpensive syntactic tests.

   (d) Using SAT and related techniques inside a tableau-based algorithm.

   (e) Historical issues and literature review.

6. **Advanced** Complexity

   This chapter is aimed at readers with some background in computational complexity and computability.

   In Chapter 4, we have seen algorithms and decidability results for standard DL reasoning problems. This chapter is aimed at the more theoretically oriented reader, and discusses the computational complexity of these reasoning problems.

   (a) We prove that $\mathcal{ALC}$ ABox consistency is PSpace-complete using a standard K-World algorithm for the upper bound and QBF for the lower.

(b) We prove that, in the presence of general TBoxes, $\mathcal{ALC}$ ABox consistency becomes ExpTime-complete using tree automata for the upper bound and PSpace alternating Turing machines for the lower.

(c) We outline what happens to the complexity of reasoning problems when the DL is extended with the constructors mentioned earlier, that is, inverse roles and number restrictions.

(d) We give one example of a constructor, role value maps, that destroys the decidability of a DL.

(e) Historical issues and literature review.

7. ?**Advanced**? Other reasoning approaches

   The tableau algorithms described in Chapter 4 are not the only ones that have been developed for DLs. In this section, we discuss two alternative reasoning paradigms, one for sub-Boolean DLs and one that reuses technology developed in first order theorem proving.

   (a) We present two structural subsumption algorithms for sub-Boolean DLs, one for $\mathcal{FL}_0$ and one for $\mathcal{EL}$. Both are proved to run in polynomial time for concept subsumption; the $\mathcal{EL}$ one can be extended to continue doing so for powerful extensions of $\mathcal{EL}$, whereas the complexity of reasoning increases for extensions of $\mathcal{FL}_0$. We will prove the former point for the extension of $\mathcal{EL}$ with general TBoxes and only sketch the latter point.

   (b) For readers familiar with first order logic, we will explain how to turn standard resolution into a decision procedure for $\mathcal{ALC}$ with general TBoxes.

   (c) Historical issues and literature review.

8. Applications for Ontologies and Conceptual Modelling

   In this section we will discuss applications, in particular (OWL) ontologies and (UML/EER) conceptual modelling.

   (a) We will explain the relationship between DLs and ontology languages, and in particular explain how Description Logics provide the logical basis of OWL-DL. We will provide a non-trivial OWL ontology/DL knowledge base, and use it to illustrate and further explain the notions introduced so far. That is, we will walk the reader through examples of unsatisfiable classes, unintended subsumptions and non-subsumptions, and instance retrievals. The example knowledge base will be of non-trivial size, and available on-line in a format that makes it usable in modern ontology editors.

   (b) We will explain how Description Logics provide a unifying framework for UML and EER and illustrate how DL reasoning can be used to support conceptual modelling.

(c) Historical issues and literature review.

9. ?**Advanced**? Further Reading

   In this section, we sketch important areas of research in description logics that have not been mentioned so far and provide pointers into the literature.

   (a) Extensions of Description Logics, including feature agreements, nominals, and concrete domains.
   (b) Combinations of DLs with other formalisms such as rules or answer set programming.
   (c) Currently available reasoners.
   (d) Currently available knowledge base editors.
   (e) Current applications and developments.
   (f) Non-standard reasoning services.

   NOTE: the discussion on reasoners and editors may be better moved to Chapter 8.

10. ?**Advanced**? Outlook

    We discuss current trends in the field.

Workplan

Already in tatters I'm afraid. We obviously need to agree a revised schedule.

- April 2006 Drafts of all chapters finished and distributed.
- June 2006 Feedback on drafts distributed.
- July 2006 Feedback discussed and agreement achieved.
- Dec. 2006 Final version ready.

Length and Format

The chapters will be of different lengths as appropriate to the material to be covered. A rough estimate is as follows:

- Chapters 1, 5 and 10 will be relatively short, approx. 10–15 pages,
- Chapters 2, 7, 8 and 9 will be of medium length, approx. 20 pages, and
- Chapters 3, 4 and 6 will be longer, approx. 30 pages.