Knowledge Modelling and Management 2010-2011
Ontologies
Stuart Aitken
10/01/2011

This course outline provides a summary for each lecture, and identifies the points you should take note of in the papers that are associated with each lecture. These notes do not replace the papers (or the lectures) – you should read the papers for each lecture as described below. The papers can be downloaded from the KMM website in a single tar file, see the *ontology-reading.tar.gz* file (files are named according to the paper's author). Copies of the lecture slides will be handed out in each lecture and the pdf made available on the KMM website after each lecture. Copies of the book *Ontological Engineering* (Gomez-Perez)[1] are available in the Main Library (shelfmark QA76.76.E95 Gom). You are **not** expected to buy this book.

Lecture 1: Introduction

This lecture will give a general introduction to ontologies, what they are used for, and how they compare with other approaches to data modelling and software design. Some simple ways to formalise ontologies using sets and first-order logic will be covered. In particular, note the definitions of subClassOf and type, the subclass and instance-of relations respectively. Note also the way in which relations are specified by their domains, ranges, and properties such as transitivity, symmetry, reflexivity, irreflexivity and whether they are functional or have an inverse or not. The use and definition of attributes and values is also important. The lecture will also cover types of ontologies and the role of standards.

Reading: Chandrasekaran et al, What are ontologies, and why do we need them? IEEE IS 1999 is required reading. Written at the end of the 90's, this paper summarises arguments for an explicit content theory, or ontology. Note in the first paragraph there is a useful contrast of content theories with mechanism theories. Mechanism theories (rules, heuristics, problem-solving methods) are important, but not the topic of this course. The first chapter in *Ontological Engineering* is also recommended as an introduction to ontologies.

Lecture 2: Methodology

Ideally, ontology development would be more of a science than an art, and this lecture presents some of the guidelines and methodologies that have been developed for building ontologies. The lecture will cover a systematic approach known as Methontology and will also present guidelines developed by the Protégé group.

Reading: Chapter 3 in *Ontological Engineering* (to be handed out in class) should be read. This chapter explains the Methontology approach in detail. Note the development of the ontology from a list of candidate concepts, to a set of tables that organise classes, instances, and relationships. 'Ontology Development 101: A Guide to Creating Your First Ontology' by Noy and McGuinness describes a step-by-step approach to building an ontology using Protégé 3. This paper discusses issues that arise frequently, such as how to organise a class hierarchy, when to represent something as an instance or as a class. The wine ontology developed in this guide is a well-known example. This paper is recommended.
Consider also reading the papers on the chemical ontology (Fernandez-Lopez) and on the legal ontology (Corcho) as they illustrate additional aspects of the Methontology approach.

[1] *Ontological Engineering with examples from the areas of Knowledge Management, e-Commerce and the Semantic Web.* Gómez-Pérez, A., Fernández-López, M., Corcho, O. Springer 2004.

Lecture 3: Introduction to Description Logic

This lecture will introduce the syntax and semantics of the Description Logic (DL) ALC. ALC is a DL with full negation ¬, union ⊔, intersection ⊓, value **∀R.C** and existential restriction **∃R.C**. It is very important to understand the meaning of these expressions.

There are many Description logics, and many are based on ALC. This course will deal with ALC, and a method for proving statements in ALC, namely the FACT algorithm. This algorithm has an efficient implementation, and a *classifier* based on FACT has been integrated into the Protégé 4 ontology editor which will be used in the course. In later lectures the extensions to ALC that are used in the Web Ontology Language OWL will be introduced. It will be necessary to understand the extensions to the ALC language and to be able to use the classifier to build ontologies using the extended language. The extensions that have been made to the basic FACT algorithm are outside the scope of the course.

Reading: There is no required reading for DL lecture 3. However, it may be helpful to read the chapter by Baader and Nutt for a good introduction to DLs. This chapter follows the usual presentation of the DL language and its semantics (note that the discussions of fixpoints and rules are beyond the scope of the course). The report by Horrocks presents ways to optimise tableaux proofs, and may also be of interest, but note that these more complex methods are beyond the scope of the course.

Lecture 4: Description Logic and the Web Ontology Language (OWL)

This lecture will continue the DL topic, covering the FACT algorithm and presenting more examples. This lecture will also introduce OWL, and explain the sub-languages OWL-Lite, OWL-DL and OWL-Full. As OWL-DL is based on ALC, the OWL language can be thought of as contributing an RDF/XML syntax to DL to create a web-friendly ontology language. OWL-DL adds a number of operators to ALC and the semantics of these will be explained (for example, cardinality ≥ can be added).

Reading: The paper by Horrocks et al 'From SHIQ and RDF to OWL: The Making of a Web Ontology Language' should be read in conjunction with this lecture. It covers the development of the language(s), the design choices and constraints.

Lecture 5: OWL 1.1 and Protégé 4

The extension of OWL to version 1.1 will be described here. OWL 1.1 adds new operators to OWL-DL to make the language more useful. This increase in expressivity must be matched by the ability to reason efficiently about the extended language and some of these methods will be outlined.

In preparation for the assessed exercise, an introduction to Protégé 4 will also be given. This will include some hints and tips for starting the tool and using the GUI. The Manchester syntax will also be introduced. This is an ASCII-based syntax for OWL 1.1 that is used in Protégé 4.

Reading: The paper by Horridge on the Manchester syntax should be read. The paper by Horrocks et al 'The even more irresistible SROIQ' may provide useful detail on the DL that OWL 1.1 is based on, SROIQ.

Tutorial: One lecture slot may be used for a tutorial on DL and Protégé.

Lecture 6: Examples of Ontologies

Following on from methodology and formalisation, this lecture will consider some well-known ontologies developed for a range of areas. The first ontologies to be considered are taken from biology and medicine: the Gene Ontology and the Foundational Model of Anatomy (FMA). These cover scientific domains, but have very different levels of

formalisation, and arose by different paths (one is the result of a community effort, the other is a product of a single group).
The next ontology to be considered is Cyc, an AI ontology developed to support common-sense reasoning. Cyc has many applications including the interpretation of natural language.
Examples of uses of ontologies in knowledge management will then be presented, including the Enterprise ontology that aims to support communication and the inter-operation of systems in organisations. Having looked at examples of ontologies, we will consider organising them according to type.

Reading: Papers for lecture 6 include: Gene Ontology (2 papers); papers by Mejino and Noy on the FMA; papers by Guha, Panton and Read on Cyc; a report by Ushold on the Enterprise ontology. It is not necessary to read all of these papers. Select one of the ontologies and read up on it in depth, or sample a paper for each. Ensure you know: the subject area and purpose, the uses that have been made of it, the level of formalisation, how it was developed, and some example concepts. You may prefer to read up on an ontology from a different area, e.g. legal concepts, chemical compounds, product catalogs.

Lecture 7: Upper Ontologies

This lecture will present an example of an upper ontology: DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering). This ontology draws on terms established in philosophy, for example, the world can be divided into *endurants* and *perdurants* (loosely: objects and events) according to the way the entity exists in time. The DOLCE ontology is designed according to 'metaproperties'. These have been used to analyse other ontologies, and this analysis will also be covered in the lecture.

Reading: The paper by Gangemi should be read for a complete description of DOLCE. The paper by Guarino describes with the metaproperties and the analysis procedure. Both papers should be read.

Lecture 8: Parts and Wholes

The relationship between parts and the whole entity they make up is a long-standing question in ontology. This lecture will begin by presenting Winston's analysis of the part-whole relation which is based on the way in which native English speakers talk about "being partly" or "being part of". This leads to 6 senses of part-of being identified, and distinguished on the basis of three higher-level attributes.
The axiomatic approach to part-of will also be covered. The basic theory has three axioms, and these can be extended in many different ways. This leads on to topology (connectedness), a logical theory concerned with how entities are connected together.

Reading: The paper by Winston can be read for additional detail on the analysis of the senses of part-of (but much of this is reproduced in the lecture slides). The paper by Varzi gives a good overview of the way in which mereology and topology are related, but contains much material that in beyond the scope of this course.

Lecture 10: Summary (Week 10)

The final lecture will include a summary of the key points of the course.

References

Lecture 1: Introduction

B. Chandrasekaran, J. R. Josephson and V. R. Benjamins What are ontologies and why do we need them? *IEEE Intelligent Systems*, Jan/Feb 1999, 14(1), pp. 20-26.

Lecture 2: Methodology

Gomez-Perez, A., Fernandez-Lopez, M., and Corcho, O. *Ontological Engineering*. Springer 2004, pp.125-142

N. Noy and D McGuinness Ontology Development 101: A Guide to Creating Your First Ontology *SMI Report Number: SMI-2001-0880*

Lectures 3, 4 and 5: Description Logic

Baader, F., Calvanese, D., McGuiness, D., Nardi, D., and Patel-Schneider, P. *Description Logic Handbook* (Chapter 2)

Ian Horrocks, Peter F. Patel-Schneider, and Frank van Harmelen. From SHIQ and RDF to OWL: The making of a web ontology language. *J. of Web Semantics*, 1(1):7-26, 2003.

Lectures 6 and 7: Examples

Gene Ontology Consortium *Nature Genetics* (2000) 25 :25-29.

Gene Ontology Consortium *Genome Research* (2001) 11:1425-1433.

Noy, N. F. and Musen, M. A. and Mejino, J. L. V. and Rosse, C. (2004) Pushing the Envelope: Challenges in a Frame-Based Representation of Human Anatomy. *SMI Technical Report - SMI-2002-0925*

Mejino, J. L. V. and Rosse, C. (2004) Symbolic modeling of structural relationships in the Foundational Model of Anatomy. *Proceedings, First International Workshop on Formal Biomedical Knowledge Representation* (KR-MED 2004)

Mike Uschold, Martin King, Stuart Moralee and Yannis Zorgios (1998) The Enterprise Ontology *The Knowledge Engineering Review* , Vol. 13, Special Issue on Putting Ontologies to Use (eds. Mike Uschold and Austin Tate).

Guha, R. V. and D. B. Lenat. "Cyc: A Midterm Report." *AI Magazine* (Fall 1990).

Reed, Stephen and D. Lenat. Mapping Ontologies into Cyc. *In AAAI 2002 Conference Workshop on Ontologies For The Semantic Web*, Edmonton, Canada, July 2002.

Panton, Kathy, P. Miraglia, N. Salay, et al. Knowledge Formation and Dialogue Using the KRAKEN Toolset. In *Eighteenth National Conference on Artificial Intelligence*, Edmonton, Canada, 2002.

Lecture 8

Aldo Gangemi, Nicola Guarino, Claudio Masolo, Alessandro Oltramari, Luc Schneider Sweetening Ontologies with DOLCE. *Proc. EKAW 2002* :166 - 181

Nicola Guarino and Christopher A. Welty An Overview of OntoClean. In Steffen Staab and Rudi Studer, eds., *The Handbook on Ontologies*:151-172.

Lecture 9

Morton E. Winston, Roger Chaffin and Douglas Herrmann. A taxonomy of part-whole relations. *Cognitive Science*, 11. 417-444.

Achille C. Varzi Parts, wholes, and part-whole relations: The prospects of mereotopology. *Data and Knowledge Engineering*, Volume 20, Number 3, November 1996 , pp. 259-286(28).