CITS3005 Knowledge Representation

Lecture 12: Unit Review

The University of Western Australia

2023

What is Knowledge Representation?

Knowledge Representation is the study of how information is processed contextualised, and synthesised into knowledge, to support decision making applications. Knowledge is more than just a collection of facts: it is a structure *built from symbols*, the *represents some condition* of the world, with *transformation rules* describing dependencies and entailment.

There are many different approaches to these tasks, but at the core of all of them is *logic*. A Logic given some language built from symbols, is the subset of "true" elements of the language.

Logic, reasoning and thought have been studied for millennia, and pose many difficult philosophical questions, which we will generally avoid in this course.

Our challenge is to consider what fragment of knowledge and reasoning can we capture and implement on a machine.

What we said we would study?

This course will focus on the practical aspects of implementing knowledge and reasoning systems in software.
We presume that all students have generally good programming skills, and are familiar with propositional and first order logic.
The topics we will aim to cover are:

1. **Knowledge, belief and information**: *Week 1*; what is knowledge, where does it come from, how do we use it.
2. **Logic**: *Weeks 2-4*; syntax and semantics of first order logic, proofs and entailment, expressing conditions.
3. **Logic Programming**: *Weeks 2-4*; syntax of `prolog`, resolution, unification, completeness
4. **Knowledge Graphs**: *Weeks 5-8*; description logic, `RDF`, OWL, SHACL, reasoning.
5. **Knowledge Systems**: *Weeks 9-12*; uncertainty, fuzzy logic, `problog`, learning, knowledge extraction, planning, search.

Lecture 1: Knowledge, Reasoning and Belief

Reviewed the basic concepts of knowledge, reasoning and belief, and considered the challenges of representing and reasoning about knowledge in a machine context.
Topics included:

- Data, Information and Knowledge.
- Knowledge and Intelligence.
- Reasoning via deduction.
- Reasoning via Induction.

Lecture 2: Logic

Reviewed and Revised the basic techniques and conventions of logical reasoning. Topics included:

- Propositional Logic
- First Order Logic
- Domains, relations and functions
- Validity, Satisfiability, Completeness and Consistency
- Automated Reasoning:
  - Conjunctive Normal Form
  - Skolemization
  - Unification
  - Resolution

Lecture 3: Logic Programming

Introduced Logic Programming via the prolog system. Topics included:

- Clausal Logic
- Recursion
- Negation by Failure
- Cuts
- Open/Closed World Assumptions
- Logic Programs as Knowledge Bases

Lecture 4: Logic Programming Theory

Reviewed the theory linking first order logic and logic programming. Topics included:

- Logic Programs as First Order Models
- Herbrand Model Theory
- Completeness of Logic Programming
- SLD Resolution

Lecture 5: Uncertainty

Reviewed the nature of uncertainty include the sources of uncertainty and reasoning in the presence of uncertainty. Topics included:

- Knowledge Extraction and Uncertainty Quantification
- Fuzzy Logic
- Epistemic Logic
- Probabilistic Logic
- Bayesian Networks

Lecture 6: Probabilistic Logic Programming

Reviewed the integration of uncertainty with logic programming in the Problog system Topics included:

- Probability Distributions
- Kolmogorov Axioms
- Bayes Rule
- Probabilistic Logic Programming
  - Probabilistic facts
  - Probabilistic rules
- Learning from Evidence
- Sampling
- Decision Theory

Lecture 7: Knowledge Graphs

Considered the problem of reasoning at scale and establishing a common discourse for reasoning. Topics included:

- Knowledge graphs, graph databases and edge labelled graphs.
- Resource Description Framework (RDF) and rdflib.
- SPARQL Queries
- Graph Patterns

Lecture 8: Knowledge Schema

Examined the problem of adding meaning and semantics to knowledge graphs. Topics included:

- RDFS
- Subclasses, SubProperties
- Mereology and Axioms
- Graph Shapes
- SHApe Constraint Language

Lecture 9: Ontologies

Considered the issue of complete representations of meaning and reasoning Topics included:

- Philosophical aspects of Ontology and Epistemology
- Deductive Knowledge
- The Web Ontology Language (OWL)
- Protégé and owlready2
- Prebuilt Axioms and Constructs
- Semantic Web Rule Language (SWRL)

Lecture 10: Reasoning Techniques

Looked at the processes and use cases for automated reasoning over ontologies. Topics included:

- First Order Representations of Ontologies
- Entailment
- Undecidability
- Description Logics
- Automated Reasoning
- Tableaux and Resolution Methods
- SWRL Reasoning

Lecture 11: Induction

Reviewed the problem of learning knowledge from data, and representing and reasoning over the induced knowledge. Topics included:

- Analytics, and Graph Properties
- Machine Learning Review
- Embeddings; learning patterns and intuitions.
- Graph Embeddings
- Knowledge Extraction: Representation and Uncertainty.
- Ontologies as a Target.

Thank you!

The 2022 exam (https://teaching.csse.uwa.edu.au/units/CITS3005/workshops/CITS3005-2022-Exam.pdf) is a reasonable example of what to expect in the final exam.

For those who are interested in the area and are considering projects or honours in the area, the NLP-TLP research group (https://nlp-tlp.org) have some interesting projects, and there will be CEED projects in this area over summer:

- Retrieval Augmented Generation:

  http://ceed.wa.edu.au/project/24-024-retrieval-augmented-generation-rag-on-large-technical-reports/

- More to come...

Thanks for your interest, apologies for the general level of disorganisation and late release of material, and good luck in your exams!