Logic as a Foundation for AI

Gerda Janssens
gerda.janssens@cs.kuleuven.be
Department of Computer Science
Room 01.26, Celestijnenlaan 200A

LFAI Syllabus: aims

- **Logic** has been developed over the centuries as a **formal** way of **representing** assumptions about a world, i.e. it is a formalism in which problems can be formulated. It has the advantage that this formulation can be made independent of an execution mechanism or a search strategy. Indeed, logic allows **separating** the algorithmic treatment of a problem from the description of the problem itself. You can reason about the knowledge expressed in logic and verify its correctness since you know its meaning.
- The course aims at convincing students of this **power of logic**, at training them in stating a problem in logic, and at making them acquainted with logic approaches to model different kinds of knowledge (definite knowledge, disjunctive knowledge, negative knowledge). In addition some forms of **more advanced** reasoning will be formulated in the context of logic (possible candidates are dynamic worlds, hypothetical reasoning, reasoning under uncertainty).

2011-2012 LFAI

- Handbook
  Knowledge Representation and Reasoning
- Book with focus on KRR as core-AI
  - What an agent needs to know in order to behave intelligently
  - How this knowledge can be represented symbolically
  - How automated reasoning procedures can make this knowledge available as needed

2011-2012

- LFAI: using this book, but focus on logic
  - Using logic as the symbolic structure for KRR
  - Selection of chapters from the book
  - Start with full first-order logic
  - Logic formalisms that compromise expressiveness for speed
  - Extensions: default logic, circumscription logic

2011-2012: study logic

- As a branch of mathematics
- As (a set of) formalisms to deal with artificial intelligence problems (KRR)
- As programming languages:
  Prolog (70s)
  Constraint Logic Programming (80s)
  Mercury (90s)
  Answer Set Programming (recent)

2011-2012: form

- Reading club
- Sessions: Mondays and Tuesdays specific chapter/part
- Every student should read it!!
- A responsible (me or students):
  - to guide the discussion: key points, difficulties
  - give some extra examples
  - link it with other courses/background
- Me: some extra material
- Exercise sessions (5)

What the responsibles have to do to activate the other students?

- Formulate some "key" questions: summary, concepts, links, why, what, how, ...
- Including 1 exercise question and "What is the take home lesson?"
- Send the questions 24 hours in advance to the other students.
- Prepare the answers to your questions.
- Give answers to me at the beginning of the sessions.

What the responsible has to do to activate the other students?

- Formulate some "key" questions:
- Send + Prepare the answers
- Give your answers to me at the beginning of the sessions.
- During session: work towards the answers.
- Activate the other students

Course Outline

(Ch 1) Introduction
(Ch 2) The Language of First-Order Logic
(Ch 3) Expressing Logic
(Ch 4) Resolution
(Ch 5) Reasoning with Horn Clauses
(Ch 6) Procedural Control of Reasoning

Course Outline

(Ch 11) Defaults
(Ch 13) Explanation and Diagnosis

2011-2012: evaluation

- Activity during the sessions (30 %)
- Individual Assignment (70 %)
  Part 1: 5 exercises to be solved
  Part 2: make 2 additional questions/exercises

Introduction

- The key concepts: knowledge, representation, and reasoning
- Why knowledge representation and reasoning?
- The role of Logic
- Course topics and organization

My questions

- Q1 What is knowledge? What is representation? What is reasoning?
- Q2 Why is Knowledge Representation and Reasoning Useful?
- Q3 What can "logic" do for KRR?
- Q4 What other forms of logics might be interesting?

Question 1

- write down on a sheet of paper the answers to the following questions:
  - What is knowledge?
  - What is representation?
  - What is reasoning?

What is said in the book?

AI and Knowledge Representation

- AI can be described as:
  The study of intelligent behavior achieved by computational means
- Knowledge representation and reasoning could be viewed as the part of AI concerned with how an agent uses what it knows in deciding what to do
- What! knowledge has to be represented before we can start reasoning with it

Knowledge

- What is knowledge? We will/cannot answer this question!
- It can be seen as a **relation** between a knower/an agent and a proposition (an idea expressed by a simple declarative sentence) John knows that … .
- Webster: proposition: an expression in language or signs of something that can be believed, doubted, or denied, or is either true or false.

Knowledge

- John knows that (it is true that) Mary will attend the party
- → objects of knowing are propositions, which can be true or false
  → the agent has come to realize that the world is one way and not another, depending on the truth of the proposition

Knowledge

- What about
  - John knows who Mary is taking to the party
    .. "that Mary is taking so-and-so to the party" implicit proposition
  - John knows how to get to the station
  - (John knows Mary well)

Other Propositional Attitudes

- John is certain that …
- John believes that … accurate??
- John suspects that …

In this course: no distinction is made.

In all cases, John takes the world to be one way and not the other (regardless of whether it is true or not).

Representation

- If A represents B, then A stands for B and is usually easier accessible than B
- In our case we are interested in (groups of) symbols that stand for some proposition
- Knowledge representation: the field of study concerned with using formal symbols to represent a collection of propositions (believed by some agent)
- One may believe infinitely many propositions –however we can only use finitely many symbols!

Reasoning

- Reasoning is the formal manipulation of symbols representing propositions in order to produce representations of new ones
- Here representation helps – while propositions are highly abstract, there representations can be easily manipulated
- Reasoning can be as easy as arithmetic, namely symbol manipulation
- Its role: bridge the gap between what is represented and what is believed

Reasoning: example 1

- "raining" is true
- "raining implies wet street" is true
- Then reasoning derives that "wet street" is true

Reasoning: example 2

- John loves Mary
- Mary is coming to the party
- Someone John loves is coming to the party

Form of reasoning is known as logical inference because the final sentence represents a logical conclusion of the propositions represented by the initial ones

Q2 Why is Knowledge Representation and Reasoning Useful?

Why is Knowledge Representation and Reasoning Useful?

- Describing/understanding the behavior of systems in terms of the knowledge it has
- Generating the behavior of the system!
  - Knowledge can be specified without being forced to think about its usage
  - Understanding the behavior in terms of the represented knowledge makes debugging and understanding much easier
  - Modifications and extensions are also much easier to perform

The Knowledge Representation Hypothesis (B. Smith) 1982

Any computationally embodied intelligent process will be comprised of structural ingredients that

a) we as external observers naturally take to represent a propositional account of the knowledge the overall process exhibits, and

b) independent of such external semantical attribution, play a formal and essential role in engendering the behavior that manifests that knowledge

see:

Reformulation

a) We can understand the symbolic representations (have a semantics)

b) The system behaviour is as it is due to the represented knowledge, because the system is using the knowledge to behave as if it knows it.

Knowledge Based …

- We call systems that are based on this hypothesis knowledge-based systems and the symbolically represented knowledge the knowledge base (KB)
- A knowledge-based system has the ability to be told facts about its world and adjust its behavior correspondingly

Knowledge-Based Systems

```
printc(snow) :- !, write("It's white").
printc(grass) :- !, write("It's green").
printc(sky) :- !, write("It's yellow").
printc(X) :- write("Beats me").
%%%%%%%%%%%%%%%%%%%%
% explicit beliefs!!!
printc(X) :- color(X,Y), !, write("It's "), write(Y).
printc(X):- write("Beats me").
color(snow,white).
color(sky,yellow).
color(X,Y):- madeof(X,Z), color(Z,Y).
madeoff(grass, vegetation).
color(vegetation, green).
```

Example cont.

`color(snow,white).`

a) Symbolic structure that we can understand as representing the proposition that snow is white

b) System will print the appropriate color of snow because Prolog …
without knowing itself anything about concepts like colors …

Advantages of the Presence of an Explicit KB

- We can add new tasks and make them depend on previous knowledge
- We can extend existing behavior by adding new beliefs (about e.g. canaries)
- We can easily debug faulty behavior
- We can concisely explain and justify the behavior

Role of Reasoning?

- Note: there was no explicit belief about the color of grass
- In general: many beliefs will be there only implicitly
- Use concept of entailment/ logical implication

Logical Entailment

- Propositions represented by a set of sentences S entail the proposition p when the truth of p is implicit in the truth of the sentences in S.
- If the world is such that every element of S comes out true, then p does as well.

Entailments of a KB

- A knowledge-based system should believe all and only the entailments of what it has explicitly represented
- Reasoning = computation of the entailments of a KB

Role of Reasoning?

- Can/shall we compute all implicit (all entailed) beliefs?
- Computationally too expensive
  to find all of them, or
  to not to return some incorrect answers
- We should keep an eye on the computational properties!!

Q3 What can "logic" do for KRR?

The Role of Logic

- Formal symbolic logic ***is*** the field of study of entailment relations – formal languages, truth conditions, semantics, and inference
- Use the tools from formal logic
- Our first KR language: first-order logic (FOL) or predicate calculus
- Subsets and supersets of FOL

Different Kinds of Reasoning

- Usually, we are interested to derive implicit, entailed beliefs form a given collection of explicit beliefs
  - In a logically sound (the derived proposition must be true, given the premises)
  - And complete way (all true consequences can be derived)
- Q4 What other forms might be interesting?

Model-Finding and Satisfiability

- In planning and configuration tasks, we often get a set of constraints and a goal specification. We then have to find a solution satisfying all the constraints.
- In this case, the solutions (models) are not necessarily found by logically sound derivations.

Model-Finding

Either round or square
Either red or blue
If red and round or if blue and square, then wood
If blue, then metallic
If square, then not metallic
If red, then square
square
Possible model: square, not metallic, red, wood
Does not logically follow, but it is one possible assignment

Inference to the Best Explanation: Abduction

- In diagnosis tasks, we often have to find a good explanation for a given observation or symptom
- Given a background theory, a set of abducibles (possible explanations) and an observation, find the most likely explanation (=subset of the set of abducibles)

Inference to the Best Explanation: Abduction

earthquake implies alarm
burglar implies alarm
{earthquake, burglar} is the set of abducibles
alarm is observed
One explanation is earthquake …
Not a sound inference; usually more than 1 possible explanation

Abduction versus Induction

- Reasoning from the effects to the causes: many things can cause the observed result
- Reasoning from particular to the general: if we see enough worlds in which something is true and we see no world in which it is false, we tend to conclude it is always true
  - I have seen 1000 black ravens, I have never seen one that is black. Therefore, every raven is black.
  - Now try red Fords
- Data mining, relational learning

Jumping to Conclusions: Default Reasoning

- Often we do not have enough information, but we want nevertheless to reach a conclusion
- In the absence of evidence to the contrary, we jump to a conclusion
- Birds are usually able to fly. Tweety is a bird. So, you would expect that Tweety is able to fly….
- Unsound conclusion
- When evidence to the contrary becomes available → withdraw conclusion
  → nonmonotonic reasoning

How Universal is the KRR Approach?

- How much of intelligent behavior needs to be knowledge-based in the sense described above?
  - This is an open question
  - A number of skills can better be modeled using other approaches (procedural, neural networks, …)
  - Also probabilistic approaches are useful, can be knowledge-based

Course Outline

(Ch 1) Introduction
(Ch 2) The Language of First-Order Logic
(Ch 3) Expressing Logic
(Ch 4) Resolution
(Ch 5) Reasoning with Horn Clauses
(Ch 6) Procedural Control of Reasoning

Course Outline

(Ch 11) Defaults
(Ch 13) Explanation and Diagnosis

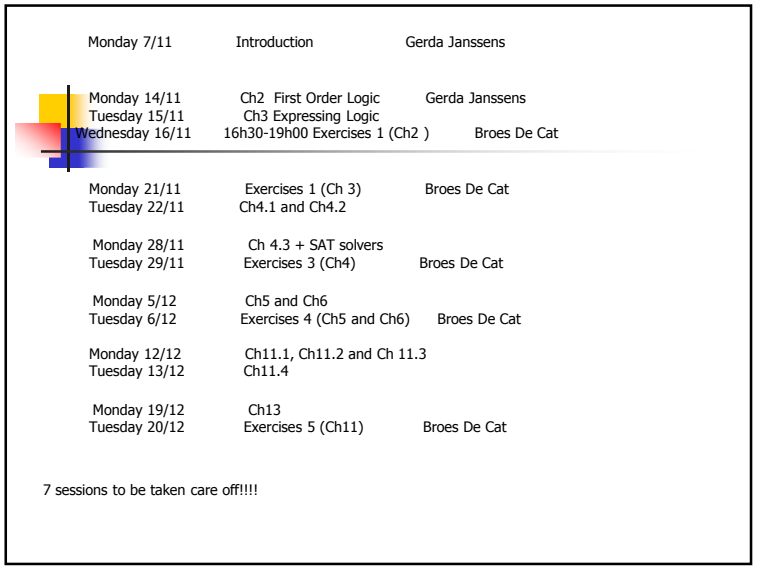

Monday 7/11	Introduction	Gerda Janssens
Monday 14/11	Ch2 First Order Logic	Gerda Janssens
Tuesday 15/11	Ch3 Expressing Logic	
Wednesday 16/11	16h30-19h00 Exercises 1 (Ch2)	Broes De Cat
Monday 21/11	Exercises 1 (Ch 3)	Broes De Cat
Tuesday 22/11	Ch4.1 and Ch4.2	
Monday 28/11	Ch 4.3 + SAT solvers	
Tuesday 29/11	Exercises 3 (Ch4)	Broes De Cat
Monday 5/12	Ch5 and Ch6	
Tuesday 6/12	Exercises 4 (Ch5 and Ch6)	Broes De Cat
Monday 12/12	Ch11.1, Ch11.2 and Ch 11.3	
Tuesday 13/12	Ch11.4	
Monday 19/12	Ch13	
Tuesday 20/12	Exercises 5 (Ch11)	Broes De Cat

7 sessions to be taken care off!!!!