COMP9414: Artificial Intelligence

Knowledge Representation and Reasoning

Wayne Wobcke

Room J17-433
wobcke@cse.unsw.edu.au
Based on slides by Maurice Pagnucco

Knowledge Representation and Reasoning

- A knowledge-based agent has at its core a knowledge base
- A knowledge base is a set of facts expressed in a suitable representation about some domain (these facts are called sentences)
- Sentences are expressed in a (formal) knowledge representation language
- **Question:** How do we write down knowledge about a domain/problem? Once we have written down knowledge, can we automate or mechanise reasoning to deduce new facts?
- References:
  - Ivan Bratko, Prolog Programming for Artificial Intelligence, Addison-Wesley, 2001. (Chapter 15)
  - Stuart J. Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, Third Edition, Pearson Education, 2010. (Chapter 7)

Knowledge Representation and Reasoning

- Any agent can be described on different levels:
  - Knowledge level (knowledge ascribed to agent)
  - Logical level (algorithms for manipulating knowledge)
  - Implementation level (how algorithms are implemented)
- Knowledge Representation is concerned with expressing knowledge explicitly in a computer-tractable way (for use by the agent in reasoning)
- Reasoning attempts to take this knowledge and draw inferences (e.g. answer queries, determine facts that follow from the knowledge base, decide what to do, etc.)

Overview

- Knowledge Representation and Logic
- Propositions
- Propositional Logic
  - Syntax
  - Semantics
- Conclusion

Why do we need formal languages?

- Consider an English sentence like:
  "The boy saw a girl with a telescope"
  Also:
  "Our shoes are guaranteed to give you a fit" (lexical ambiguity)
  "I heard about him at school" (structural ambiguity)
  "As he uttered the all-important word he dropped his voice, but she just managed to catch it" (ambiguity of cross reference)
- Natural languages exhibit ambiguity
- Not only does ambiguity make it difficult for us to understand what is the intended meaning of certain phrases and sentences but also makes it very difficult to make inferences
- Symbolic logic is a syntactically unambigious knowledge representation language (originally developed in an attempt to formalize mathematical reasoning)

Syntax vs. Semantics

Syntax Describes the legal sentences in a knowledge representation language (e.g. in the language of arithmetic expressions $x < 4$)

Semantics Refers to the meaning of sentences. Relates sentences (and sentence fragments) to aspects of the world the sentence is about. Semantics refers to a sentence's relationship to the "real world" or to some model of the world. Semantic properties of sentences include truth and falsity (e.g. $x < 4$ is true when x is a strictly smaller number than 4 and false otherwise). Semantic properties of names and descriptions include referents.

Note: The meaning of a sentence is not intrinsic to that sentence. An interpretation is required to determine sentence meanings. Interpretations are provided by agreement amongst a linguistic community or by fiat (e.g. use the symbol 'cat' to refer to the class of cats).

Propositions

- Propositions are entities (facts or non-facts) that can be true or false
- Propositions are expressed using ordinary declarative sentences
  - e.g. the sentence "The sky is blue" expresses the proposition that the sky is blue. Is this proposition true?
- Other propositions:
  "Socrates is bald" (assume 'Socrates', 'bald' are well defined)
  "The car is red" (assume 'the car' is identified)
  "Socrates is bald and the car is red" (complex proposition)
- Often use single letters to represent propositions in reasoning:
  P: Socrates is bald
  (reasoning is independent of propositional substructure!)
- This is referred to as a scheme of abbreviation

Logical Arguments

All humans have 2 eyes
Jane is a human
Therefore Jane has 2 eyes

All humans have 4 eyes
Jane is a human
Therefore Jane has 4 eyes

- **Both** are (logically) correct **valid** arguments
- Which statements are true/false?

Logical Arguments

All humans have 2 eyes
Jane has 2 eyes
Therefore Jane is human

No human has 4 eyes
Jane has 2 eyes
Therefore Jane is not human

- **Both** are (logically) incorrect **invalid** arguments
- Which statements are true/false?

Propositional Logic

- Use letters to stand for "basic" propositions; combine them into more complex sentences using operators for not, and, or, implies, iff
- Propositional connectives:

$\neg$	negation	$\neg P$	"not P"
$\wedge$	conjunction	$P \wedge Q$	"P and Q"
$\vee$	disjunction	$P \vee Q$	"P or Q"
$\rightarrow$	implication	$P \rightarrow Q$	"If P then Q"
$\leftrightarrow$	bi-implication	$P \leftrightarrow Q$	"P if and only if Q"

From English to Propositional Formulae

- "it is not the case that the sky is blue": $\neg B$
  (alternatively "the sky is not blue")
- "the sky is blue and the grass is green": $B \wedge G$
- "either the sky is blue or the grass is green": $B \vee G$
- "if the sky is blue, then the grass is not green": $B \rightarrow \neg G$
- "the sky is blue if and only if the grass is green": $B \leftrightarrow G$
- "if the sky is blue, then if the grass is not green, the plants will not grow": $B \rightarrow (\neg G \rightarrow \neg P)$

Improving Readability

- $(P \rightarrow (Q \rightarrow (\neg(R))))$ vs. $P \rightarrow (Q \rightarrow \neg R)$
- Rules for omitting brackets
  - omit brackets where possible
  - precedence from highest to lowest is: $\neg$, $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$
  - all binary operators are left associative (so $P \rightarrow Q \rightarrow R$ abbreviates $(P \rightarrow Q) \rightarrow R$)
- **Questions:**
  - Is $(P \vee Q) \vee R$ the same as $P \vee (Q \vee R)$?
  - Is $(P \rightarrow Q) \rightarrow R$ the same as $P \rightarrow (Q \rightarrow R)$?

Semantics

- The semantics of the connectives can be given by truth tables

P	Q	$\neg P$	$P \wedge Q$	$P \vee Q$	$P \rightarrow Q$	$P \leftrightarrow Q$
True	True	False	True	True	True	True
True	False	False	False	True	False	False
False	True	True	False	True	True	False
False	False	True	False	False	True	True

- One row for each possible assignment of True/False to propositional variables
- **Important:** Above P and Q can be **any** sentence, including complex sentences

Semantics — Complex Sentences

R	S	$\neg R$	$R \wedge S$	$\neg R \vee S$	$(R \wedge S) \rightarrow (\neg R \vee S)$
True	True	False	True	True	True
True	False	False	False	False	True
False	True	True	False	True	True
False	False	True	False	True	True

Terminology

- A sentence is valid if it is True under all possible assignments of True/False to its propositional variables (e.g. $P \vee \neg P$)
- Valid sentences are also referred to as tautologies
- A sentence is satisfiable if and only if there is some assignment of True/False to its propositional variables for which the sentence is True
- A sentence is unsatisfiable if and only if it is not satisfiable (e.g. $P \wedge \neg P$)

Material Implication

- The only time $P \rightarrow Q$ evaluates to False is when P is True and Q is False (i.e. $P \rightarrow Q$ is equivalent to $\neg P \vee Q$)
- This is known as material implication
- English usage often suggests a causal connection between antecedent (P) and consequent (Q) – this is not reflected in the truth table
- So $(P \wedge \neg P) \rightarrow$ *anything* is a tautology!

Entailment

- $S \models P$ — whenever all the formulae in the set S are True, P is True
- This is a semantic notion; it concerns the notion of truth
- We can determine whether $S \models P$ by constructing a truth table for S and P (gives a syntactic counterpart to entailment) — $S \models P$ if, in any row of the truth table where all the formulae in S are True, P is also True
- A tautology is just the special case when S is the empty set

Entailment

P	Q	$P \rightarrow Q$	Q
True	True	True	True
True	False	False	False
False	True	True	True
False	False	True	False

Therefore $\{P,\ P \rightarrow Q\} \models Q$

- since in the only row where both P and $P \rightarrow Q$ are True (row 1), Q is also True (here S on the previous slide is the set $\{P,\ P \rightarrow Q\}$)

Note: The column for $P \rightarrow Q$ is calculated from that for P and Q using the truth table definition, and Q is used again to check the entailment

Entailment — Tautologies

R	S	$\neg R$	$R \wedge S$	$\neg R \vee S$	$(R \wedge S) \rightarrow (\neg R \vee S)$
True	True	False	True	True	True
True	False	False	False	False	True
False	True	True	False	True	True
False	False	True	False	True	True

Therefore $\models (R \wedge S) \rightarrow (\neg R \vee S)$

Conclusion

- Due to the ambiguity in natural languages there is a need to specify knowledge through the use of formal languages
- Not only will these formal languages give us a way to remove ambiguity but they will also help to provide methods for automating inference
- Propositional logic is a first move in this direction
- In the next lecture we look at automating inference using the resolution rule on which Prolog is based
- We shall also investigate first-order predicate logic
- Keep in mind that there are a large number of logics and knowledge representation schemes that we do not study