

Knowledge Representation and Reasoning

Peter Lucas¹ and Marcel van Gerven²

`peterl@cs.ru.nl, m.vangerven@donders.ru.nl`

¹Institute for Computing and Information Sciences, and

²Donders Centre for Cognition, Radboud University Nijmegen

People involved

Lectures:

- Peter Lucas
- Marcel van Gerven
- Arjen Hommersom

Practicals and tutorials:

- Wenyun Quan (w.quan@cs.ru.nl)
- Maarten van der Heijden
(m.vanderheijden@gmail.com)
- Arjen Hommersom (arjenh@cs.ru.nl)
- Max Hinne (mhinne@cs.ru.nl)

Start practical: See blackboard/website

Course outline

Website: <http://www.cs.ru.nl/~peterl/teaching/KeR>

(1) Lectures:

- Theory of knowledge representation and reasoning; core of this formed by:
 - AI-style logics and probability theory
 - Nowadays you can even combine logic and probability theory

(2) Tutorials: do exercises

(3) Practical:

- Learn some programming in Prolog (the AI logic programming language)
- Develop reasoning systems in Allegro

Topics

- Refresh your memory on formal logic
 - this week: read “Logic and Resolution” (available on blackboard and website) – **no lectures on 8th September!**
 - **12th September:** revision lecture on logic
 - **19th September:** logic exercises
- Week 26th September: programming in Prolog and AILog
- Knowledge representation formalisms
- Model-based reasoning
- Reasoning with uncertainty
- Decision making

Learning aims of the course

- Obtain insight into the **development** of knowledge systems, the use of problem solving methods, forms of knowledge representation, and model-based reasoning
- Gain knowledge about **logical expressiveness** of forms of knowledge representation and the use of probability theory in reasoning with uncertainty
- Being able to develop reasoning programs using **Prolog** and **ALog**
- Being able to understand **core AI research** as reflected in ECAI, IJCAI, AAAI


Reading and study material

- Slides and exercises
- Summarising lecture notes
- Practical manual
- Some material can be found in:
S. Russell and P. Norvig, **Artificial Intelligence: A Modern Approach**, 2003 or 2009:
 - Part III Knowledge and Reasoning
 - 8 First-Order Logic
 - 9 Inference in First-Order Logic
 - 10 Knowledge Representation
 - Part V Uncertain Knowledge and Reasoning
 - 13 Uncertainty
 - 14 Probabilistic Reasoning

Further reading

- R.J. Brachman and H.J. Levesque, **Readings in Knowledge Representation**, Morgan Kaufmann, San Mateo, CA, 1985:
 - good collection of early KR papers
- F. van Harmelen, V. Lifschitz, and B. Poter, **Handbook of Knowledge Representation**, Elsevier, Amsterdam, 2008:
 - state of the art handbook on KR
- D. Poole, **Artificial Intelligence: Foundations of Computational Agents**, Cambridge University Press, 2010
 - nice **systematic and coherent** approach to AI using logical and probabilistic reasoning (freely accessible)

Knowledge systems



Early knowledge systems

- **Expert system**: use of a *large* collection of symbolic **expert** knowledge to solve problems:
 - E.A. Feigenbaum, B.G. Buchanan, J. Lederberg – **Heuristic DENDRAL** (1965): contains knowledge from organic chemistry
 - E.H. Shortliffe: **MYCIN** (1974–1979) – diagnostics of infectious diseases
 - H.E. Pople, J.D. Myers: **Internist-1** (1973-1982) – diagnosis in the big area of internal medicine
 - D. Lenat: **Cyc** (1984-) – representation of common sense knowledge

Modern knowledge systems

- Use of more sophisticated logical methods:
 - abductive reasoning (cause-effect reasoning and explaining observations), e.g., used in abductive diagnosis
 - use of functional models of behaviour, e.g., in consistency-based diagnosis
- Use of probabilistic and decision-theoretic methods:
 - Bayesian networks (to reason with uncertain knowledge)
 - augmented by decision theory (to allow making decision about appropriate actions)

Knowledge: Implicit versus explicit

Two hypotheses:

- (1) Human reasoning is hard to capture, and, thus, intricate **implicit** methods, such as neural networks, capture human reasoning best;
- (2) Human reasoning can be captured, although possibly incompletely. However, **explicit** representation is necessary for getting a grip on that knowledge (e.g., to be able to **explain** recommendations)

Choice: **explicit knowledge**

Knowledge representation hypothesis

Brian Smith (1982):


Any mechanically 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** that the overall process exhibits;
- (b) independent of such external semantical attribution, play a formal but causal and essential role in **engendering the behaviour** that manifests that knowledge

Knowledge representation hypothesis

Forall x (Fish(x) and InAquarium(x) and ...

Behaviour!



Observers

Formal knowledge representation

- Logic was originally developed as a language for **mathematical** reasoning
- Goal of knowledge representation different: represent **semantic** content of psychologically plausible memory models
- **Need** for semantics shared by most researchers in knowledge representation
- In time logic has become the **dominant** language as probability theory for uncertainty reasoning

Natural kinds

- **Mathematical definitions:** exact and complete. Example **triangle**: shape defined by three points that are not on a straight line and that are connected by lines
- **Natural kinds:** objects in the real world. ‘Definitions’ of objects are often approximate and incomplete

Example:

$$\forall x(\text{Human}(x) \rightarrow \text{Animal}(x))$$

$$\forall x(\text{Human}(x) \rightarrow (\text{walks}(x) = \text{upright} \wedge \dots \wedge \dots))$$

Role of representation system

- To manage **beliefs** expressed in the language
- More than just implementation of a (logical, probabilistic) calculus
- The symbolic representation cause the system to behave in a particular fashion

Thus,

Any language with sufficient expressive power can be used

Language requirements

Levesque & Brachman (A fundamental tradeoff in knowledge representation and reasoning)

Emphasis on:

- What is represented $\equiv$ content $\Rightarrow$ knowledge level (rather than symbol level)
- Statements must be interpreted in relationship to other statements (otherwise no knowledge)
- This implies: language should have a truth theory
- Not a single language, but spectrum of languages (from simple, computationally tractable, to complex, computationally intractable)

Truth theory

- Knowledge base **KB**: what we know about the world
- Question: is the truth of statement φ implied by KB
(note φ need not be **included** in KB)
- Notation:

$$KB \models \varphi$$

- In the form of **inference** = **reasoning**:

$$KB \vdash \varphi$$

or,

$$\vdash KB \rightarrow \varphi$$

($KB \rightarrow \varphi$ is a **theorem**) if we use logic, but many logics and other languages are still possible

Logics for knowledge representation

- First-order logic:
 - satisfiability: **undecidable**
 - when it is known that KB is unsatisfiable, then $KB \models \perp$ is decidable
- (Finite) propositional logic:
 - decidable, but **NP complete**
 - propositional Horn logic: model checking in polynomial time

Horn formula: $(A_1 \wedge \dots \wedge A_n) \rightarrow B$, with A_i and B **positive** literals


$\Rightarrow$ **tradeoff** between expressive power and computational complexity

Prolog

- Logical (programming) language with some restrictions, but based on first-order predicate logic
- One of the typical AI programming languages (other Lisp)
- Close relationship with knowledge representation and reasoning: **AILog**

```
in_mind([h,o,l,i,d,a,y]).  
start :- write('Guess first letter'), read(X),  
         in_mind([X|T]), write('OK. '), guess(T).  
guess([]) :- write('The word is '), in_mind(W), write(W), !  
guess(L)  :- repeat, write('Next letter'), read(X),  
              ((L=[X|T1], write('OK. '), guess(T1));  
               (write('Fail. Try again!'), guess(L))).
```

Model-based reasoning




- Explicit representation of structure and function of systems (= model)
- Reasoning with this model to solve problems (e.g. diagnosis)

Reasoning with uncertainty


- Early: uncertainty attached to rule-based reasoning (ako uncertain reasoning with logical implications)
- 1990s: introduction of Bayesian/belief networks (causal networks with attached probability distributions)
- 1990s: extension to decision networks/influence diagrams (decision making under uncertainty)
- Recent: probabilistic logics (logic and probability theory integrated in an AI fashion)

Thus, after 30 years back to the early problem, which is now well understood

Examples 1: Pacemaker programming





Examples 1: Causal pacemaker model




Examples 2: Smart production printers

- project with Océ and Embedded Systems Institute
- model-based reasoning about behaviour of printers/copiers
- interpretation of sensor information
- adaptation to changing environment



Examples 2: Adaptive control


Avoid that paper temperature becomes lower than 66°C with 99% certainty:



Research at CAI

- Cognitive robotics
- Theoretical models of human reasoning
- Brain-computer interfaces
- Brain reading

Cognitive robotics



- Robots perceive an incomplete image of the world using **sensors** that are inherently unreliable
- Given this **partial observability**, we need algorithms to:
 - Estimate the true state of the world
 - perform planning in order to reach goals
- Relies on probability theory and decision theory




Models of human reasoning




- Helmholtz proposed that perception is unconscious inference.
- That is, our percepts are our best guess as to what is in the world, given both sensory data and prior experience
- This hypothesis can be formulated in terms of probability theory

Models of human reasoning


- Probabilistic theories of human brain function and cognition are dominant



Example: Goal inference




Example: Goal inference



- Formal models of abstract reasoning
- Formulated as inference in probabilistic networks
- Complexity analysis of these models informs human cognition

Brain-computer interfaces



- BCIs allow brain-based communication and control
- Requires new algorithmic developments

Example





Figure 1: Adaptive flashing POMDP.


- Optimal experimental design for P300 spellers
- Based on partially-observable Markov decision processes

Brain reading



- It is possible to decode mental states from brain data
- New ways of analyzing the brain increases our understanding
- This relies on sophisticated probabilistic machine learning algorithms

Example



- Reconstruction of perceived images from BOLD data
- Probabilistic inference in Markov random fields

Conclusions

- Knowledge representation and reasoning defines the very core of AI
- Logic, probability theory and decision theory form its theoretical foundations
- The basis for building intelligent agents and applications
- Concepts form the basis of modern theories on **human** knowledge representation and reasoning and their complexity