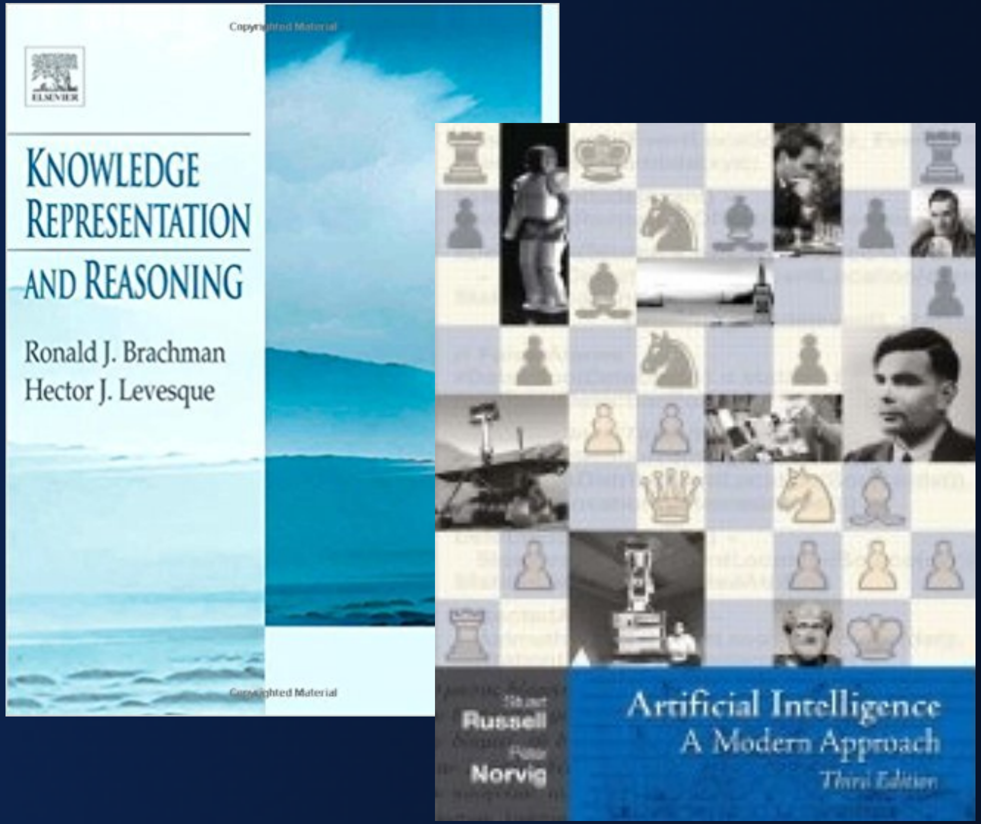

AI Fundamentals: Knowledge Representation and Reasoning

Maria Simi

Knowledge Representation and Reasoning

LESSON 1: INTRODUCTION

Knowledge representation and reasoning

Knowledge representation and Reasoning (KR&R) is the field of artificial Intelligence dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks.

The class of systems that derive from this approach are called **knowledge based (KB) systems/agents**.

A KB agent maintains a **knowledge base** of facts expressed in a declarative language. The KB is the agent's representation of the world.

The classical book about KR&R is:

Ronald Brachman and Hector Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 2004.

Knowledge representation

What kind of knowledge?

The emphasis is on the relation between an agent and **facts that may be true or false in the world**:

"John knows p", "John believes p", "John desires p", "John is confident that p" … where *p* is a proposition, something *true* or *false.*

Contrast with **non-factual** knowledge (knowing how, when, where, a person …):

"John knows how to get to the party"
"John knows Bill very well"
"John knows how to play the piano".

Represented knowledge is given a **propositional account**.

Knowledge representation is about the use of formal symbolic structures to represent a **collection of propositions** believed by some agent. Not necessarily all of them.

Symbols denotations

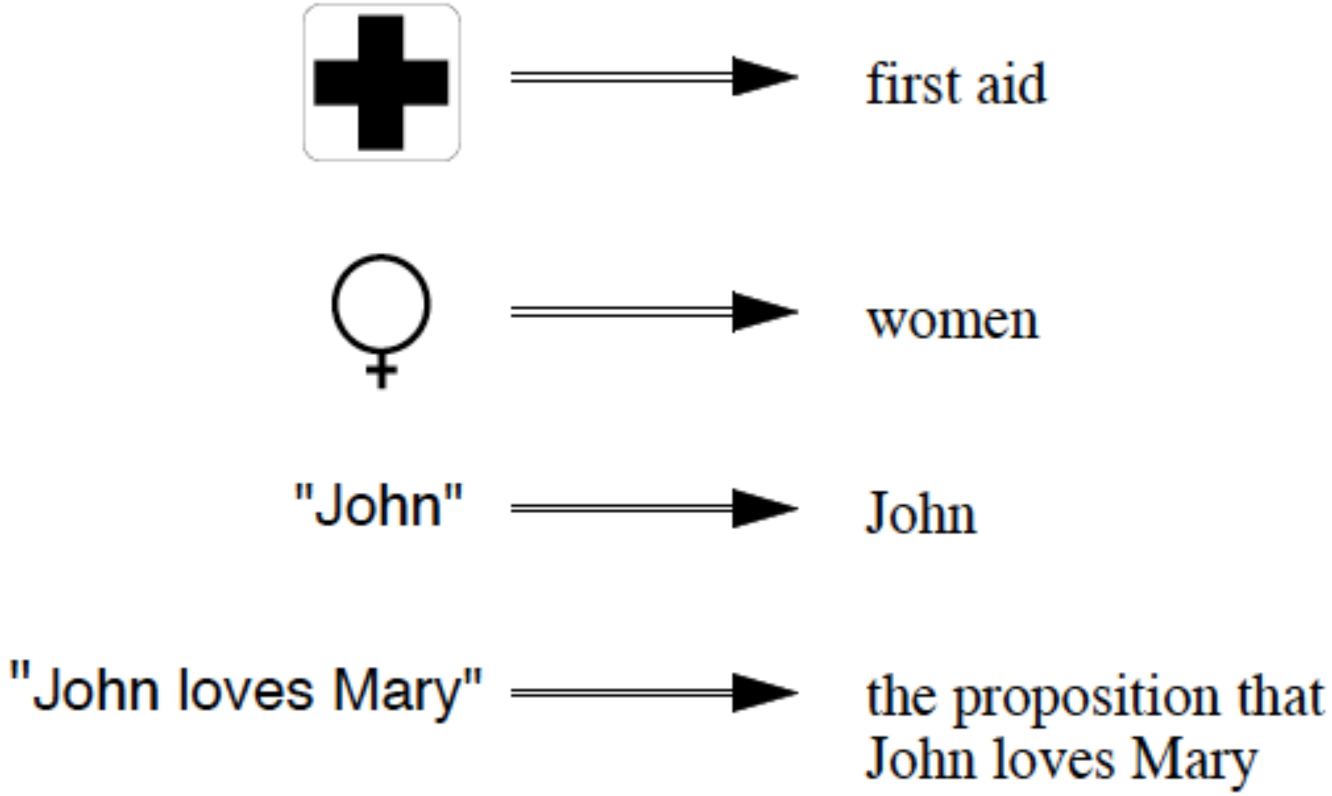

Reasoning

Reasoning is the formal manipulation of the symbols representing a collection of beliefs to produce representations of new ones.

Analogy with arithmetic.

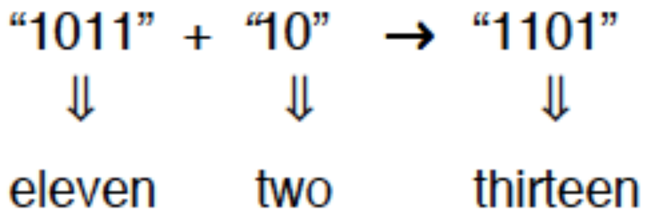

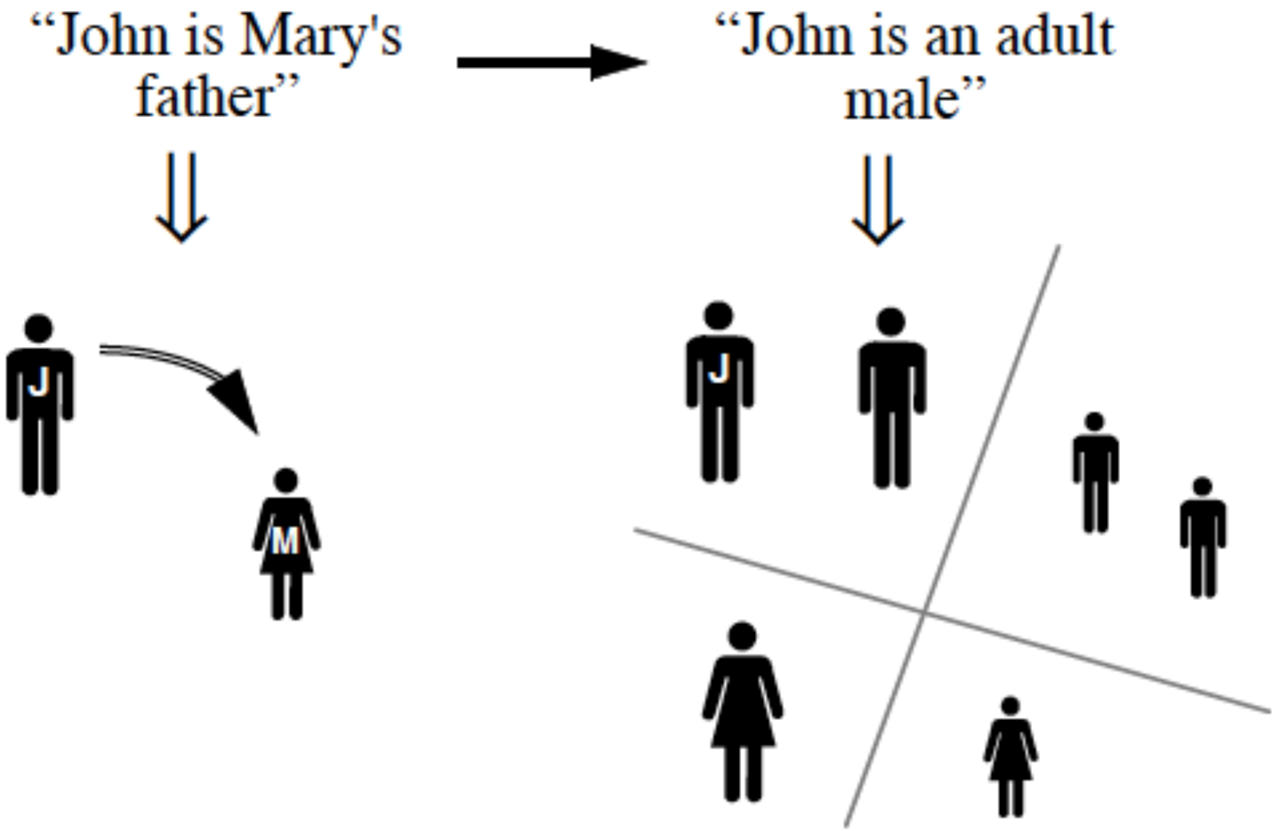

Logical deduction is a well know example of reasoning

Why reasoning?

We would like the system to depend on *what the system believes* and not *what was explicitly stored*. It is a question of economy of the representation. Example:

A. “*Patient x is allergic to medication m*.”

B. “*Anybody allergic to medication m is also allergic to m'*.”

Is it OK to prescribe m' for x ?

NO. In fact, C. “*Patient x is allergic to medication m’*.”

Usually we need more than just DB-style retrieval of facts in the KB

Logical entailment (KB $\models$ α) is the fundamental operation.

We know that computing all logical entailments may be too complex: any procedure that always gives us answers in a reasonable amount of time will occasionally

- miss some entailments (incomplete) or
- return some incorrect answers (unsound)

Other forms of reasoning

- Not only logical entailment/deduction
- Also:
  - Limited expressivity
  - Limited inference
  - Default reasoning
  - Probabilistic reasoning
  - Abduction: Given a causal relation $\alpha \Rightarrow \beta$, from observing β, we can *conjecture* α. Abduction is as a way of providing an explanation.

The Knowledge Representation hypothesis

Formulated by Brian C. Smith in 1985:

> *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, and b) independent of such external semantic attribution, play a formal but causal and essential role in engendering the behavior that manifests that knowledge.*

In simpler words, we want to construct A.I. systems that contain symbolic representations with two important properties:

1. We can understand these symbolic structures as propositions.
2. These symbolic structures determine the behavior of the system

Knowledge based systems have these properties.

Advantages of KB systems

Competing approaches:

- **Procedural approach**: knowledge is embedded in programs.
- **Connectionist approach:** avoids symbolic representation and reasoning, and instead advocates computing with networks of weighted links between artificial "neurons."

Advantages of a KB approach:

1. **Extensibility**: we can extend the existing behavior by simply adding new propositions.
2. **Debugging**: we can debug faulty behavior by changing erroneous beliefs.
3. **Accountability**: we can explain and justify current behavior in terms of the beliefs.

The interplay between representation and reasoning

Representation and **Reasoning** are intimately connected In A.I. research.

1. The representation scheme must be expressive enough to describe many aspects of complex worlds with symbolic structures.
2. The reasoning mechanisms has to ensure that reasoning can be performed efficiently enough.

There is a tradeoff between these two concerns.

The **fundamental tradeoff in knowledge representation and reasoning**: the more expressive the representation language, the more complex is the reasoning. We need to find the best compromise.

Applications

Much of AI involves building systems that are **knowledge-based:** their ability derives in part from reasoning over explicitly represented knowledge. Typical KB systems are:

- language understanding, machine reading
- planning
- medical diagnosis
- “expert systems”, etc.

Some systems are KB, to a certain extent

- game-playing, vision, etc.

Some systems are KB to a much lesser extent

- speech, motor control, etc.

KR&R today has many applications outside AI:

- Bio-medicine, Engineering, Business and commerce, Databases, Software engineering

Current research question: how much of intelligent behavior is knowledge-based?

Two “modern” applications

Think of the following systems:

- Cognitive assistants/digital personal assistants (SIRI/Alexa/Google home): spoken dialog in a natural language in open domain
- Computational Knowledge Engine (Wolfram Alpha): scientific and medical thinking

For each system:

- What knowledge must it represent?
- What reasoning must it do?
- What would it take to extend it?
- Where does it fail?

SIRI

What knowledge must it represent?

- Restaurants, movies, events, reviews, weather, ...
- Location, tasks (buy, info …), web sources, ...

What reasoning must it do?

- Nearest location, date for tomorrow, AM vs PM, etc

What would it take to extend it?

- More sources, different sources

Where does it fail?

- Completely different environment, completely different task

Wolfram Alpha

http://www.wolframalpha.com/examples/ (knowledge domains)

Wolfram Alpha is more than a search engine. It gives you access to the world's facts and data and calculates answers across a range of topics/domains including scientific ones ...

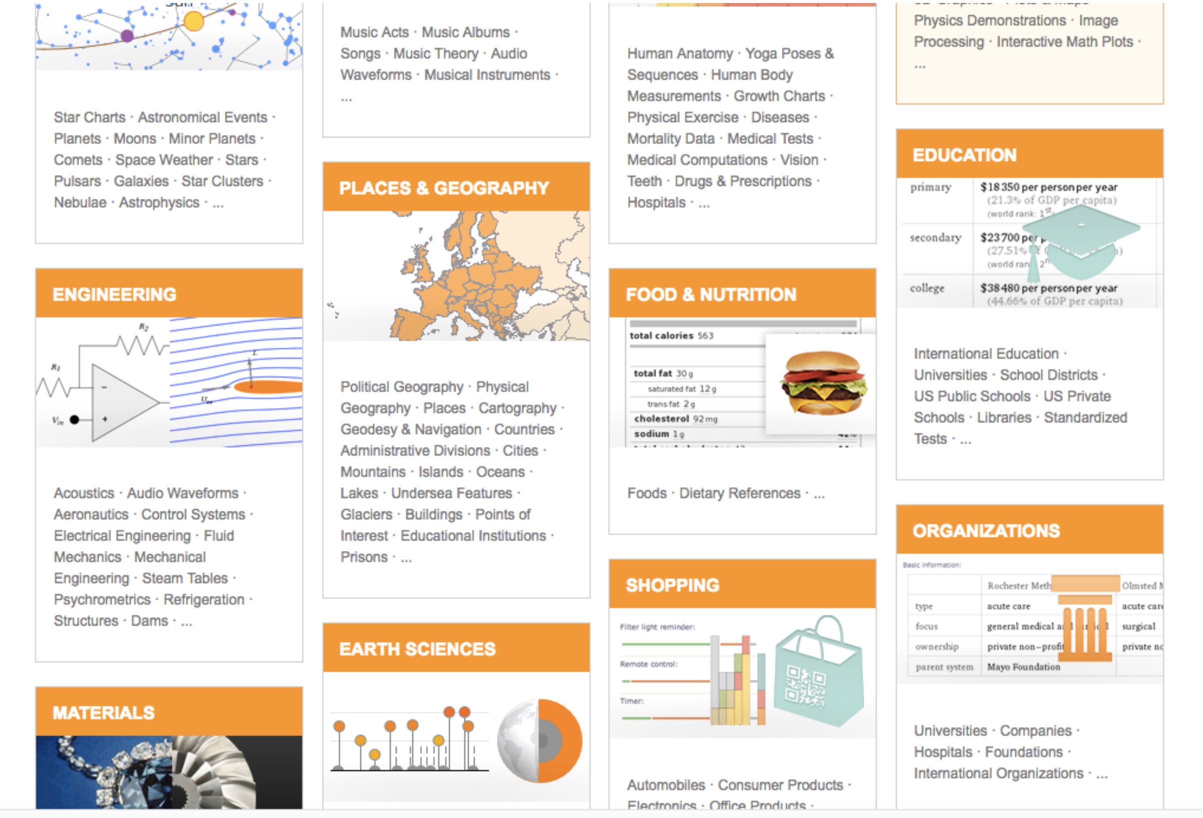

Food and Nutrition example

http://www.wolframalpha.com/examples/FoodAndNutrition.html

WolframAlpha computational knowledge engine.

Vitamin C in a glass of fresh orange juice

Web Apps | Examples | Random

Assuming small glasses for "glass" | Use medium glasses or large glasses instead

Input interpretation:

orange juice	amount	1 small glass	vitamin C
	type	fresh	

Result:

124 mg (milligrams)

Unit conversions:

0.12 grams

1.2×10^{-4} kg (kilograms)

Unit conversions:

0.12 grams

1.2×10^{-4} kg (kilograms)

Comparison to reference daily intake:

	daily amount	% daily value
RDA	60 mg/day (milligrams per day)	207%
DRI (male)	90 mg/day (milligrams per day)	138%
DRI (female)	75 mg/day (milligrams per day)	166%

(RDA: recommended daily allowance; DRI: dietary reference intake)

Wolfram Alpha

What knowledge must it represent?

- Different kinds of foods, their nutrition composition, caloric values

What reasoning must it do?

- Mathematical computations based on portions

What would it take to extend it?

- Add more data on foods and nutrition composition

Where does it fail?

- Does not know about recipes, how to combine foods

The role of “classical logic”

Complete and sound logic inference is the place to start.

$KB \models \alpha$

We can understand KB systems at two different levels [A. Newell].

Knowledge level: representation language and its semantics; expressive adequacy (what can be expressed), what can be inferred.

Symbol level: computational aspects, efficiency of encoding, data structures and efficiency of reasoning procedures, including their complexity.

The tools of symbolic logic seem especially suited for the knowledge level.

By classical logic we mean the propositional calculus (PROP) and first order logic (FOL).

What you need to know about logic

- ✓ Truth, satisfiability and validity
- ✓ Logical entailment $KB \models \alpha$
- ✓ Logical deduction $KB \vdash \alpha$
- ✓ Completeness: if $KB \models \alpha$ then $KB \vdash \alpha$
- ✓ Soundness: if $KB \vdash \alpha$ then $KB \models \alpha$
- ✓ Proof by refutation: $\mathrm{KB} \models \alpha$ *iff* $\mathrm{KB} \cup \{\neg\alpha\}$ unsatisfiable
- ✓ Transformation in clausal form

PROP fundamentals (see AIMA)

- ✓ Syntax and semantics
- ✓ Algorithms for entailment
- ✓ Transformation in clausal forms
- ✓ Algorithms for satisfiability (DPLL and Walk-SAT)
- ✓ The resolution method
- ✓ Resolution strategies

FOL fundamentals (see AIMA)

- ✓ Syntax and semantics
- ✓ Transformation in clausal forms (including Skolemization)
- ✓ Unification
- ✓ The resolution method
- ✓ Resolution strategies

What’s next

1. Ontology engineering: reasoning about actions, change, events
2. Non monotonic reasoning
3. Reasoning about Knowledge & Beliefs
4. Structured representations: semantic networks, objects and frames.
5. Description logics

Your turn

Please review (if you did not cover this material in IIA) :

- ✓ PROP: AIMA cap. 7
- ✓ FOL: AIMA cap. 8 (1-3)
- ✓ Inference in FOL: AIMA cap 9 (skip Logic Programming and Prolog)

Knowledge engineering in another domain using FOL:

- ✓ The electronic circuits domain (AIMA cap 8.4)
- ✓ The Internet shopping world (AIMA cap 12.7)

Read seminal paper:

- ✓ *A fundamental tradeoff in knowledge representation and reasoning, by Levesque and Brachman.*

References

[KRR] Ronald Brachman and Hector Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 2004. (Cap 1, 3)

[AIMA] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach* (3rd edition). Pearson Education 2010 (cap 12).