Third Indian School on Logic and its Applications

18-29 January 2010

University of Hyderabad

Spatial and Temporal Knowledge Representation

Antony Galton

University of Exeter, UK

PART I: What is Knowledge Representation?

Contents of Part I

- Representation and Reasoning
- Commonsense Knowledge
- Quantitative vs Qualitative
- Vagueness, Uncertainty, and Granularity

Representation and Reasoning

Why Knowledge Representation?

- In the early days of Artificial Intelligence, people thought that to make a computer intelligent all you had to do was to give it a capacity for pure reasoning.
- People soon realised that the exercise of intelligence must involve interaction with an external world, which requires *knowledge* about that world.
- The quest for AI inevitably involves the development of methods for endowing computer systems with knowledge.
- This in turn highlighted the question of *how to represent knowledge* within the computer.
- Thus the subfield of AI known as **knowledge representation** (KR) was born.

Declarative *vs* Procedural Knowledge

- **Declarative Knowledge** = Knowing *that* such-and-such is the case.
- **Procedural Knowledge** = Knowing *how* to do something

I know that $7 \times 7 = 49$ — declarative knowledge.

I don't know what 777×777 is, but I know how to calculate it — procedural knowledge.

Artificial Intelligence requires both declarative and procedural knowledge. An important form of procedural knowledge is the ability to *reason* — to derive new facts from old. Logic plays an important part — but not the only part — in this.

Two approaches to knowledge representation

There are two broad approaches to KR:

- **Sub-symbolic:** Investigate how knowledge is represented in the human brain, and try to mimic that.
- **Symbolic:** Make use of the *external* forms of representation used by humans to encode their knowledge (speech, writing, diagrams, mathematical notation, . . .).

'Nature-inspired' approaches such as Connectionism tend to take the first approach.
Mainstream AI generally takes the second.

In the AI context, 'knowledge representation' usually means *the quest for explicit symbolic representations of knowledge that are suitable for use by computers*.

What is knowledge?

Knowledge is more than just facts, information, or data.

These things constitute knowledge only if they are situated in a context provided by some general understanding of the domain they relate to.

Example: A child can correctly answer the questions

What is the capital of France?
What country is Paris in?

Thus he is in possession of certain **facts**.

But does he **know** that Paris is the capital of France?

Only if he also has at least a basic understanding of the difference between a country and a city, the role of a capital in the life of a country, etc. Without this, his knowledge is mere parroting.

Two dimensions of knowledge

Imagine *two* children: Jack and Jill.

- Jill has a good understanding of geography, politics, economics, etc, but a very poor memory for specific facts such as which city is the capital of which country.
- Jack is an *idiot savant*. He has an encyclopaedic command of large numbers of such specific facts, but is totally lacking in any contextual understanding of them.

Together they could make a good team!

Current computer systems are like Jack. The whole enterprise of KR could be seen as the quest to give them, in addition, the abilities of Jill.

Knowledge and understanding

Representing knowledge involves representing **facts** and representing **understanding**.

This generally takes the form of some general model within which the specific facts can be represented and brought into relation with one another.

KR is more concerned with formulating such models than with the collection of individual facts: with establishing a framework of understanding within which the facts make sense.

The key to establishing such a framework is to endow the computer with a capacity for **reasoning**. Hence the full title of this discipline is really Knowledge Representation and Reasoning.

The extra ingredient: Reasoning

KR is really KRR: Knowledge Representation *and Reasoning*. Armed with knowledge in the form of a collection of *general rules* and *individual facts*, a competent reasoner can deduce further individual facts.

If we know that Exeter is in Devon and that Devon is in England, then we do not need to be told that Exeter is in England: we can infer it, provided we also know

- the rule: *For any geographical regions A, B, and C, if A is in B and B is in C then A is in C*,

and also

- a general rule of inference (Modus Ponens).

The formulation of methods of reasoning about knowledge is an important part (or partner?) of Knowledge Representation.

The tools of KR

- **Ontology.** The systematic enumeration and classification of the various kinds of entity that are represented within a given domain of discourse, together with an account of their properties and relationships.
- **Representation methods.** Various formalisms such as semantic nets, scripts, frames, conceptual graphs.
- **Formal Logic.** The systematic investigation of truth, meaning, and inference.

Commonsense Knowledge

Commonsense knowledge

Commonsense knowledge: Knowledge used by humans in their everyday lives when not engaged in tasks requiring technical knowledge and skills acquired through specialist training.

Compare:

- **Scientific physics:** The knowledge of the physical world as set out in the academic discipline of Physics.
- **'Naïve physics':** The knowledge of the physical world shared by all humans by virtue of their intuitive understanding of how physical objects behave under the ordinary circumstances encountered in everyday life.

Patrick Hayes, 'The naïve physics manifesto' (1979)
Patrick Hayes, 'Naïve physics I: Ontology for liquids' (1985)
Ernest Davis, 'Pouring liquids: A study in commonsense physical reasoning' (2008)

Commonsense physical knowledge

Some examples of commonsense physical knowledge are:

- Most unsupported objects fall towards the ground (but feathers, soap bubbles, balloons, sparks etc, may float or rise upwards).
- Unpowered moving objects normally slow down and eventually come to rest.
- Solid objects cannot simultaneously occupy overlapping locations.
- Liquids, and some substances (such as rice or dry sand) consisting of many small grains, can be poured.

Note that some of these *appear* to contradict scientific physics — but they can be reconciled.

Science vs Common sense

The nature of time and space as revealed by physics (e.g., relativity) is radically different from our common-sense notions.

Artificial Intelligence is more interested in the common-sense notions of time and space (and everything else) than in the scientific view.

Therefore AI does not, for the most part, engage with such notions as

- The speed of light
- The relativity of simultaneity
- Curved space-time
- Many-dimensional spaces

(fascinating though these may be)

A Paradox

The challenge for Knowledge Representation is to formulate the concepts employed by human commonsense reasoning so that they can be manipulated by a computer.

In doing this, we invariably find that we have to invent elaborate formal structures that seem to be remote from the spirit of "common sense".

Example: To say that 'liquids flow continuously', Ernest Davis has the axiom:

$$\forall_{L:\text{liquidChunk}, TS, TE:\text{time}, D:\text{distance}} \quad D > 0 \;\Rightarrow\; \exists_{L1:\text{liquidChunk}} \text{subchunk}(L1, L) \wedge \text{simpleFlow}(L1, D, TS, TE)$$

You cannot avoid this sort of thing if you want to formalise commonsense knowledge in an explicit way!

Qualitative vs Quantitative

Qualitative reasoning

For the most part, human commonsense understanding of the world is **qualitative** in nature.

This is in contrast to specialist scientific or technical understanding, which is often **quantitative**.

Compare:

- If you drop a ball from a height, it will fall vertically downwards, getting steadily faster as it does so, until it comes into contact with some surface.
- An unsupported body initially at rest h m above the surface of the earth (where $h \ll 6 \times 10^6$) will, after t seconds, be at height $h - 4.9t^2$ m and moving downwards with velocity $9.8t$ m/s.

Qualitative vs Quantitative Reasoning

- **Quantitative:** The price went up by 30p on Monday and by 50p on Tuesday. Therefore it ended up 80p higher.
  The price went up by 30p on Monday and down by 50p on Tuesday. Therefore it ended up 20p lower.
- **Qualitative:** The price went up on Monday and again on Tuesday. Therefore it ended up higher.
  The price went up on Monday and down on Tuesday. We can't deduce whether it ended up higher, lower, or the same.

Qualitative information is often insufficient to allow a conclusion to be drawn; but often it's all we have to work with.

In AI we want to endow machines with 'common sense' — which includes an ability to reason with qualitative information. This is much harder for machines than quantitative reasoning.

Example: Qualitative reasoning about times and events

Given that

- The time of the earthquake overlapped the time of the landslide.
- The time of the landslide overlapped the collapse of the dam.

What can we say about the relationship between the time of the earthquake and the collapse of the dam?

We'll answer this later, in the section on Temporal Knowledge Representation.

Vagueness, Uncertainty, and Granularity

Vagueness and uncertainty

Commonsense knowledge is often vague or uncertain:

- Where are the boundaries of Central London? (Vague)
- Where was Archimedes' tomb? (Uncertain)

Handling vagueness and uncertainty is a major problem for computer systems, which generally work in a black-and-white, 'all or nothing' way.

Methods proposed for handling vagueness include **fuzzy logic** and **fuzzy set theory**, **rough sets**, and **supervaluation semantics**.

Methods for handling uncertainty include approaches based on **probability theory**, and various forms of **non-monotonic reasoning**.

Granularity

The **granularity** of a representation is a measure of its level of detail. The **grain size** is the size of the smallest subdivision of reality that is represented in the representation.

Fine granularity = Smaller "grain" size = More detail
Coarse granularity = Larger "grain" size = Less detail.

Don't say "increase/decrease the granularity" as this is ambiguous — does it mean increase/decrease the grain size or the amount of detail?

Typically, coarsening the granularity leads to increased uncertainty.

Granularity and Uncertainty

Grain size = 1 year.

Famous people born in 1809 include Charles Darwin, Abraham Lincoln, Alfred Lord Tennyson, and Cyrus McCormick, the inventor of the combine harvester. Who was the oldest?

Grain size = 1 month.

Tennyson was born in August 1809, but Darwin, Lincoln, and McCormick were all born in February 1809.
We still don't know who was the oldest!

Grain size = 1 day.

McCormick was born on 15th February 1809, but Darwin and Lincoln were both born on 12th February 1809.
We still don't know who was the oldest!

Maybe with grain size 1 hour, we could tell (but I don't have the necessary data . . .)

Summary of Part I

So, what is KR?

Knowledge Representation, as a subfield of Artificial Intelligence, can be characterised as

> *the quest for ways to encode human commonsense qualitative knowledge and understanding in the form of explicit symbolic representations that can be manipulated by a digital computer.*