

KNOWLEDGE REPRESENTATION*

Ly Anh Tuan

This work is produced by OpenStax-CN X and licensed under the Creative Commons Attribution License 3.0[†]

Abstract

Knowledge representation[43] and knowledge engineering[44] are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects;[45] situations, events, states and time;[46] causes and effects;[47] knowledge about knowledge (what we know about what other people know);[48] and many other, less well researched domains. A complete representation of "what exists" is an ontology[49] (borrowing a word from traditional philosophy), of which the most general are called upper ontologies.

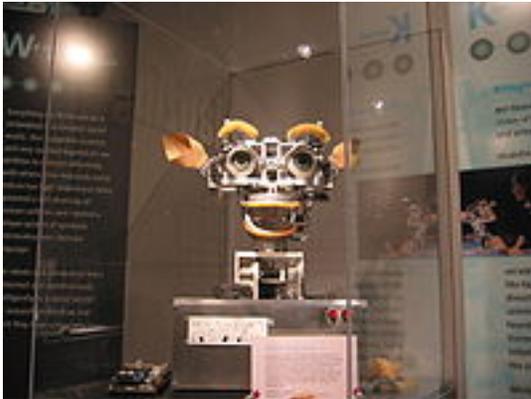
Among the most difficult problems in knowledge representation are:

Default reasoning and the qualification problem

Many of the things people know take the form of "working assumptions." For example, if a bird comes up in conversation, people typically picture an animal that is fist sized, sings, and flies. None of these things are true about all birds. John McCarthy identified this problem in 1969[50] as the qualification problem: for any commonsense rule that AI researchers care to represent, there tend to be a huge number of exceptions. Almost nothing is simply true or false in the way that abstract logic requires. AI research has explored a number of solutions to this problem.[51]

The breadth of commonsense knowledge

The number of atomic facts that the average person knows is astronomical. Research projects that attempt to build a complete knowledge base of commonsense knowledge (e.g., Cyc) require enormous amounts of laborious ontological engineering — they must be built, by hand, one complicated concept at a time.[52] A major goal is to have the computer understand enough concepts to be able to learn by reading from sources like the internet, and thus be able to add to its own ontology.



The subsymbolic form of some commonsense knowl-

*Version 1.1: Sep 16, 2010 7:45 am -0500

[†]<http://creativecommons.org/licenses/by/3.0/>

edge

Much of what people know is not represented as "facts" or "statements" that they could actually say out loud. For example, a chess master will avoid a particular chess position because it "feels too exposed"[53] or an art critic can take one look at a statue and instantly realize that it is a fake.[54] These are intuitions or tendencies that are represented in the brain non-consciously and sub-symbolically.[55] Knowledge like this informs, supports and provides a context for symbolic, conscious knowledge. As with the related problem of sub-symbolic reasoning, it is hoped that situated AI or computational intelligence will provide ways to represent this kind of knowledge.[55]