Why has Artificial Intelligence not (yet) succeeded?

Zeyn Saigol

Cake Talk Series, School of Computer Science, University of Birmingham

January 8, 2008

Outline

Outline

Motivation

- Idea is to generate discussion
- I am not an expert!
  (most of the research for this talk was done on Wikipedia)
- I don't have any strong ideas
- I'm going to avoid the philosophical issue of whether strong AI is even possible

Definition of Artificial Intelligence

- AI can be viewed as 'automating intellectual tasks' (Russell & Norvig, 2003).
- *The science of finding tractable subclasses of intractable problems*
  – Richard Dearden
- From this viewpoint, AI has been very successful.
- However, I want to talk about *Strong AI*: creating an artificial mind that can solve problems of a similar complexity and diversity to those a human mind can solve (given enough training, in both cases).
- Also known as the *Artificial General Intelligence* (AGI) problem.
- Classic yardstick is the Turing Test (Turing, 1950).

Dartmouth Conference and Early Expectations

- Dartmouth Conference in 1956 – John McCarthy, Marvin Minsky, Claude Shannon, Allen Newell and Herbert Simon
- *Machines will be capable, within twenty years, of doing any work a man can do.*
  – Herbert Simon, 1965
- *Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved.*
  – Marvin Minsky, 1967

Outline

Early AI

- Largely based on problem representations in logic, and inference engines.
- Example: Blocks world

```
On(A,B)
On(B,C)
On(C,table)
```

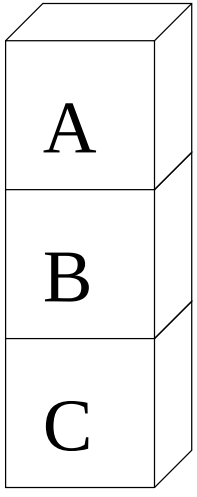

- Later, expert systems were developed to leverage domain knowledge to allow larger problems to be solved.
- Also known as rule-based systems, these used a set of if-then-else rules provided by experts to guide the search through solution space.

Modern AI

- Neural Networks: Networks formed from on/off switches, or *perceptrons*, are able to learn using the backpropagation algorithm.
- Evolutionary Computation
- Probabilistic Reasoning and Learning: Bayes' rule, Bayesian Networks and statistical inference have been applied very successfully to a range of problems.

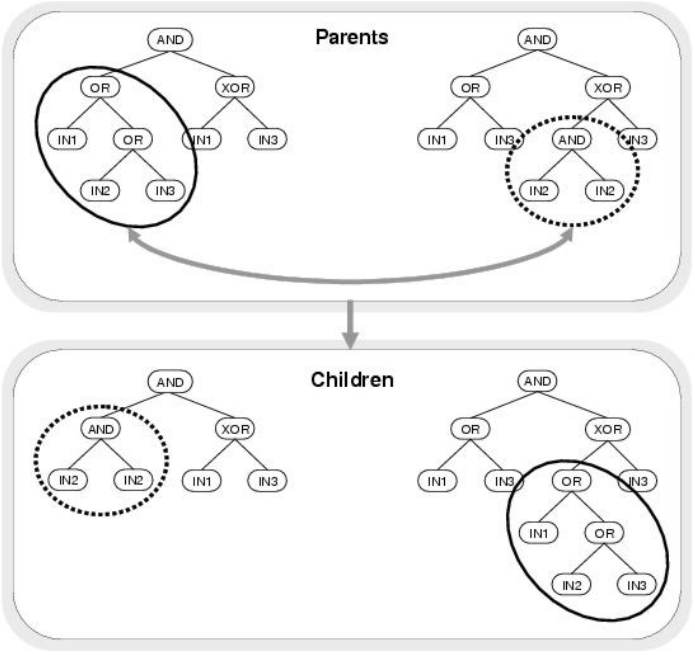

Early Successes

- SHRDLU (Terry Winograd, MIT, 1972)

```
Person:  WHICH CUBE IS
SITTING ON THE TABLE?
Computer:  THE LARGE
GREEN ONE WHICH
SUPPORTS THE RED
PYRAMID.
```

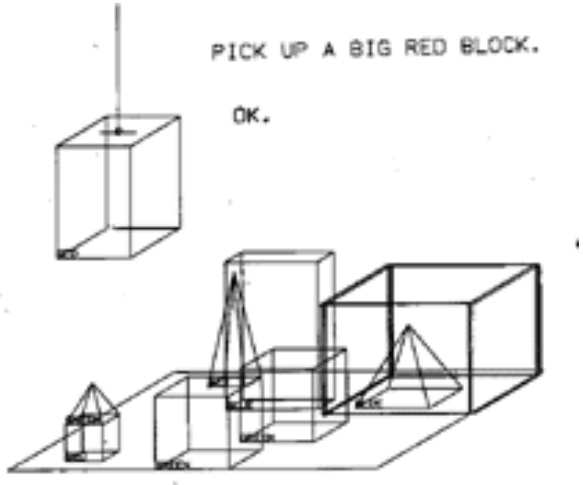

- MYCIN (Feigenbaum, Shortliffe and Buchanan, Stanford, 1970s)
- Diagnosed blood infections using a knowledgebase of 450 rules.
- Performed better than general-practice doctors, and as well as some experts in the field.

Modern State-Of-The-Art

- Deep Blue beat the reigning chess world champion, Garry Kasparov, in 1997.
- Machine learning and statistical techniques:
  - Rough, but working, machine translation
  - Text retrieval – Google
- DARPA Urban Challenge

CMU's Boss

- Follow rules of the road
- Detect and track other vehicles at long ranges
- Find a spot and park in a parking lot
- Obey intersection precedence rules
- Follow vehicles at a safe distance
- React to dynamic conditions like blocked roads or broken-down vehicles

Outline

Problem Complexity

- Early research failed to appreciate the problem of computational complexity.
- Many of the problems found in AI are inherently intractable (otherwise they wouldn't be AI!). This meant toy examples could be solved relatively easily, but sophisticated-looking systems and techniques broke down completely on real-sized problems.
- Incomplete knowledge base – especially obvious in NLP: "After John proposed to Marsha, they found a preacher and got married. For the honeymoon, they went to Hawaii."
- Moravec's Paradox (hidden complexity!). Things that are easy for human beings actually require very sophisticated processing.

We're Nearly There...

- Hardware Limitations:
  Even though a computer is a million times faster in raw switching speed, the brain ends up being 100,000 times faster at what it does (Russell & Norvig, 2003).
- School of thought that believes intelligence will eventually emerge when progress from all the disparate fields of AI is integrated together.
- Lack of embodiment – intelligence can only emerge when an agent is *situated* in an environment, and can learn from perceptions and interactions with the environment (Damien).

Breakthrough Required

- Neuroscience is advancing rapidly at the moment, due to improvements in technologies such as MRSI, and increasing computational power for simulations.
- Possibly we can only create an artificial mind by understanding and duplicating the human brain (Ray Kurzweil).
- Counter-argument is the invention of flight – only happened when we stopped trying to copy birds.
- Maybe we need a fundamental breakthrough (or several!) in some other area.
- For example, a 'universal algorithm'.
- Pedro Domingos: Need an 'interface layer' (he suggests Markov Logic Networks).
- My personal vague feeling: Knowledge representation is a key area.

Outline

Is Complexity a Barrier?

- It could be argued that strong AI is not possible, as humans cannot hope to understand something as complex as an intelligent mind.
- If we can't understand them, how can we hope to create them?...
- However, as a species, we routinely design and build things far too complicated for any one person to understand.

Complexity Examples

Microprocessors are extremely complicated:

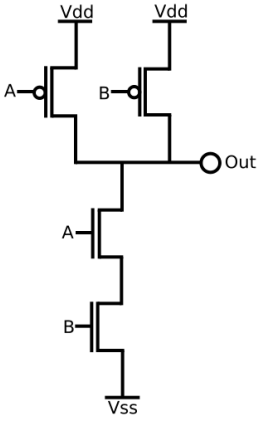

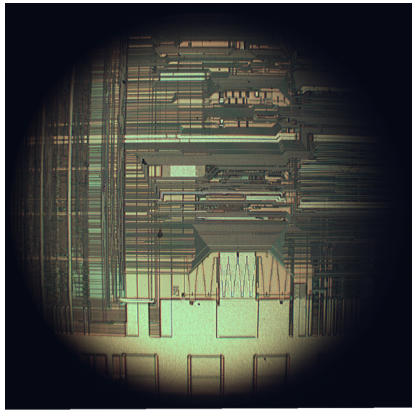

- Semiconductor physics
- Field-effect transistors
- Logic gates (CMOS)
- Microprocessor functions and architecture
- VLSI

Complexity Examples

Microprocessors are extremely complicated:

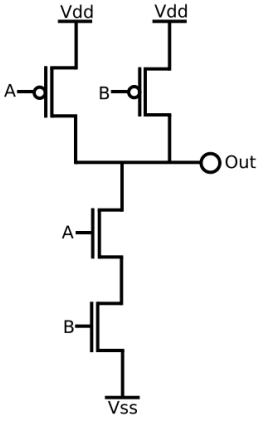

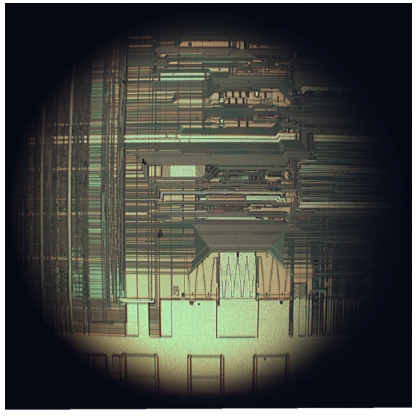

- Semiconductor physics
- Field-effect transistors
- Logic gates (CMOS)
- Microprocessor functions and architecture
- VLSI

But Airbus A380s have *hundreds* of microchips, plus software, jet engines, aerodynamic surfaces, control systems,...

Summary

- So: Is AI research on the right tracks?
- Or do we need to do more?...
- Are we researching the right areas? Or just fine-tuning the search and optimisation algorithms we have already?
- I believe complexity is **not** a good enough excuse for having not yet succeeded.
- Have we really tried hard enough? For example, the Apollo program was 14 years long and employed 400,000 people at its peak.