


Is Artificial Intelligence (AI) possible?

Thomas Bolander, Associate Professor, DTU Informatics


Science & Cocktails, 1 November 2011




A bit about myself

Thomas Bolander

- Associate professor in **logic** and **artificial intelligence** at **DTU Informatics** (since 2007).
- **Current research:** How to equip AI systems with a **Theory of Mind (ToM)**?
- Head of the **iMARS Lab** (intelligent Multi-Agent Robotic Systems Lab) at DTU Informatics.
- Director of studies in Mathematics at the Copenhagen University Extension.



AI in sci-fi



AI in our everyday surroundings



Google

Google Search

I'm Feeling Lucky

What is AI?

Will we ever achieve “sci-fi grade” AI?

Difficult question!

Let’s start with a (seemingly) simpler question: What is AI?

The “father of AI”, John McCarthy, defines AI as follows:

*“Artificial intelligence is the **science and engineering** of making **intelligent machines**, especially **intelligent computer programs**.”*



John McCarthy, 2006

Great, but what’s “intelligence” ???

What is intelligence?

Intelligence is a very difficult concept.

A (tentative) definition by 52 American psychologists (1994):

*“Intelligence is a very general mental capability that, among other things, involves the ability to **reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.**”*


It definitely characterises some of the key properties AI systems (should) have.

But to which degree? And how do we check?

The Turing Test

Alan Turing 1950 paper:


- Can machines think?
- How are we to check?
- Difficult, but this we can check:
- Can a computer *behave* indistinguishably from a human?



Alan Turing, 1954

Turing Test:

- 1 judge, 1 computer, 1 human.
- Task of the judge: who's who among computer and human?



The Loebner Prize

- **Loebner Prize:** Yearly **chatbot** competition on **Turing Test**.
- Best performing chatbot: 4.000\$.
- Passing Turing Test completely: 100.000\$. Those money are still in the bank!

Following slide: A dialogue with the chatbot **Jabberwacky** (Cleverbot), a top performer through the years.

Jabberwacky: Have you read any books lately?

Thomas: Yes.

Jabberwacky: What books have you read?

Thomas: "I am a strange loop".

Jabberwacky: Yes you are.

Thomas: Is that supposed to be funny?

Jabberwacky: Yes.

Thomas: OK, accepted, it was kind of funny.

Jabberwacky: No. Who are you trying to convince?

Thomas: You.

Jabberwacky: Tell me something funny.

Thomas: The best joke I know is in Danish.

Jabberwacky: That's a poem I guess, tell me another.

Thomas: Why don't you tell me one?

Jabberwacky: I don't know any.

Thomas: How can you not know any jokes?

Jabberwacky: No, not that, what about star wars?

Thomas: I don't like star wars.

Jabberwacky: Why don't you like star wars?

Thomas: I find it boring.

Jabberwacky: How very British of me. I'm sorry. Are you going out tonight?

Thomas: No, I'm preparing a talk for tomorrow.

Jabberwacky: Preparing to become what?

Relevance of the Turing Test

Conclusion: Passing the Turing Test is still not right around the corner.

Some history: Early optimism

The history of AI is almost as long as the history of computers themselves: starting in the early 1950s.

Early period (50s and 60s) characterised by:

- Very high expectations.
- Serious underestimation of the complexity of the human brain.

“It is not my aim to surprise or shock you—but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied” (Herbert Simon, 1957).



This “visible future” appears to drag on...

Breakthroughs in the 50s og 60s

The **unbounded optimism** in the 50s and 60s didn't come out of nothing. AI conquered many problems seemingly requiring “real intelligence”:

1. Playing chess (better than most amateurs).
2. Proving geometrical theorems (better than most math students).
3. Solving geometrical analogy problems (IQ tests).
4. The first “electronic person”, Shakey the Robot.

We have **reasoning** (all), **planning** (1 & 4), **abstract thinking** (2 & 3), **learning** (4).



Machine translation

The early successes in AI led to an atmosphere of: “If we can do this, we can do anything.” And so they tried...

Machine translation project: from Russian to English.
Reaction to the successful launch of Sputnik (1957).


They were too naive: Translating by simple **syntax transformations**.

Example. Translating English → Russian → English:

“The spirit is willing but the flesh is weak” →
“The vodka is good but the meat is rotten.”

Machine translation today by Google Translate:

“Pilots are flying planes” → “Piloter flyver fly.”
“Men are flying planes” → “Mænd er flyvende fly.”



The greatness and fall of AI

End 60s: Fewer successes. Still no “thinking machines” in the strong sense as expected/promised.

Result: Great disappointment.

In 1973, the British government cancelled all support to research in AI based on a report that concluded:

“in no part of the field have discoveries made so far produced the major impact that was then promised.”

That'll teach you not to promise too much!

AI gets new life

End 80s onwards: AI comes to life again. More realistic atmosphere: “Let’s take what we can get.”

- **1991. US Defence planning system** employed in Gulf War logistics. Saved more money than spent on AI research 1950–1991.
- **1994. Driverless car** drives 1000 km on public roads in France.
- **1997. IBM chess computer Deep Blue** beats world champion Gary Kasparov.
- **2001. Computer wins competition on stock trading.**
- **2011. IBM Jeopardy computer Watson** beats the Jeopardy world champions.




Deep Blue

But: The AIs of the 90s and 00s are no more “thinking machines” in the strong sense than the AIs of the 50s and 60s.

Example. Deep Blue based on the same strategy as the chess computers of the 50s.

However, **dramatic increase in computational power:** 500 specially designed chess processors calculating 200.000.000 moves/s.



Kasparov on Deep Blue

Gary Kasparov on **Deep Blue** in Huffington Post, 2010:

“The AI crowd, too, was pleased with the result and the attention, but dismayed by the fact that Deep Blue was hardly what their predecessors had imagined decades earlier when they dreamed of creating a machine to defeat the world chess champion. Instead of a computer that thought and played chess like a human, with human creativity and intuition, they got one that played like a machine, systematically evaluating 200 million possible moves on the chess board per second and winning with brute number-crunching force.”

“Deep Blue was only intelligent the way your programmable alarm clock is intelligent. (Not that losing to a \$10 million alarm clock made me feel any better).”

Conflicting emotions

Successes in AI are usually faced by conflicting emotions:

- We are **impressed** by what the AI can do.
- We are **disappointed** by how it does it.




Passing the Turing Test?

But what if an AI system finally passed the Turing Test?

Probably, the same would happen:

- **Computer:** A machine doing bit-flipping following a strict recipe.
- **Can be simulated by a human:** Manually and mechanically using pen and a (very large) piece of paper.



Where's the “thinking” and “understanding”?

Humans **must** be more than that?

Characteristics of current AI

Characteristics of current AI:

- **Domain specific** rather than generic.
- **Simple algorithms** rather than **higher cognitive abilities**.
- Lots of **raw computational power**.

AI outsmarts us on raw computational power, not on cognitive abilities.



Wired Magazine, January 2011

The future of AI


Does this mean that **higher cognitive abilities** are irrelevant to AI?

No, not at all.

Example. Human + chess computer plays much better than two chess computers.

Higher cognitive abilities are essential to the **holy grail of AI: generic AI.**

But we are not there yet...



Will robots take over the world?

Will robots take over the world?

Perhaps, but not anytime soon.

Proof. Not a single AI system has proven to be able to solve problems significantly different from those programmed for.

Watson is not playing chess and Deep Blue not Jeopardy. And they're not able to learn either.



The limits of AI

Are there fundamental limits to the degree of AI that can be achieved?

No convincing proof of either possible answer.

You might say: AI will never solve problems requiring “real intelligence.”
But we can't decide what “real intelligence” is.

Even a computer passing the Turing Test we might end up considering as “cheating” .

Final words

Returning to the original question:

Will we ever achieve “sci-fi grade” AI?

At least the following is relatively certain:

- AI will proceed in the direction of achieving (simulating) **higher cognitive abilities**.
- AI **will** significantly change our everyday lives.
- **At least** to the extent that the computer and the Internet already did.



early computer



early Internet



early robot