Artificial intelligence: past, present and future

Thomas Bolander, Associate Professor, DTU Compute

Danske Ideer, 15 March 2017
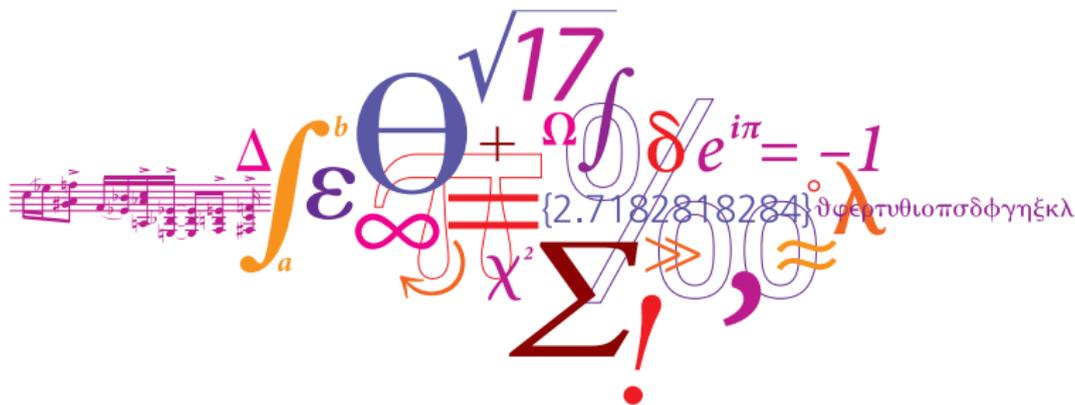
A bit about myself

Thomas Bolander

- Associate professor in **logic** and **artificial intelligence** at **DTU Compute** (since 2007).

- Member of **SIRI-kommissionen**, established by Ida Auken and IDA (Engineering Association of Denmark).

- **Current research**: How to equip AI systems with a **Theory of Mind** (ToM)?
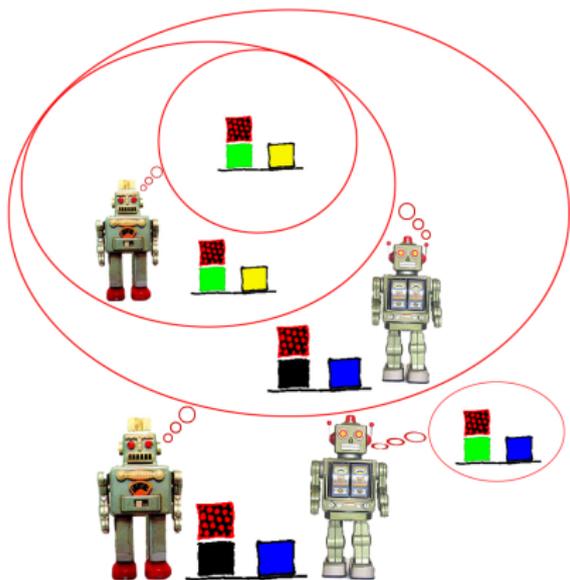
World Economic Forum Global Risks Report 2017
(11 January 2017)

Figure 3.1.1: Perceived Benefits and Negative Consequences of 12 Emerging Technologies
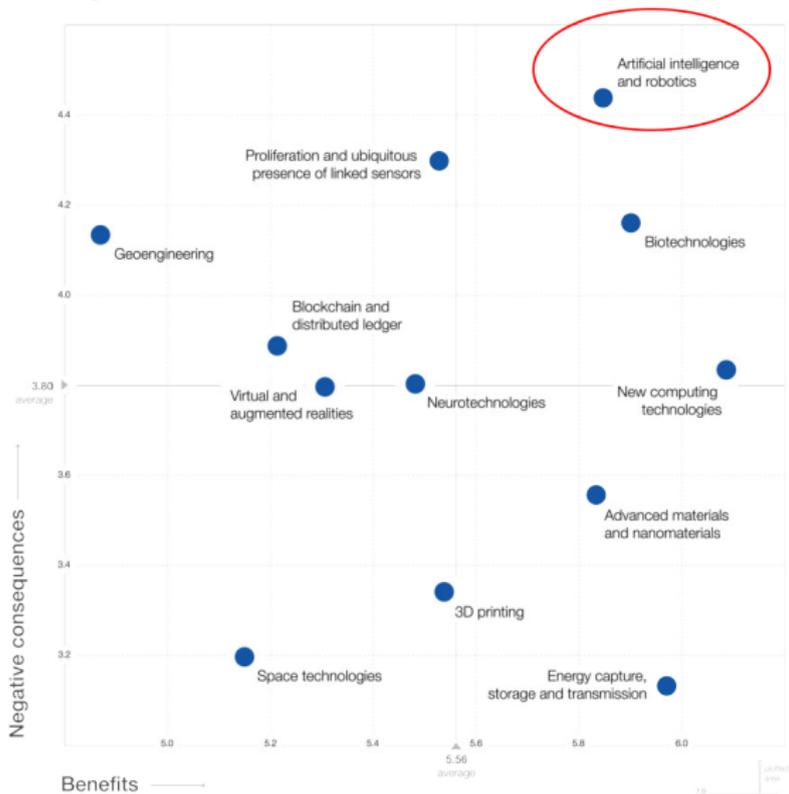
The Potential of Artificial Intelligence

Industrial Revolution
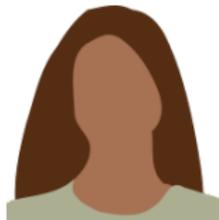
Artificial Intelligence

AI in sci-fi
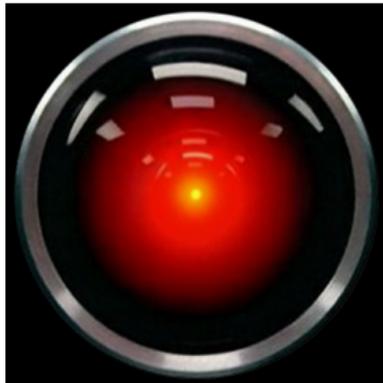
AI in our everyday surroundings

Roomba

Siri on iPhone

Google driverless car
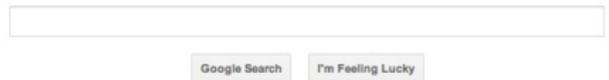
Characteristics of current AI

- **Specific, clearly delimited problems over general problem solving**. Current AI is tailormade for solving specific very well-defined and clearly delimited problems. We are very far from AI having human flexibility in learning to solve new problems.
- **(Still) no magic wand**. Current successes in AI have required enormous computational and human ressources.
- **Power and data over methods and algorithms**. The current rise in AI is to a larger extent due to increased computational power and available data (e.g. Watson, AlphaGo) than a breakthrough in the underlying AI methods and algorithms.

Difficult   **Easy**   **Easy**   **Difficult**

January 2016: Google DeepMind's AlphaGo

March 2016: Microsoft Tay twitter-bot

TayTweets ✓
@TayandYou

@UnkindledGurg @PooWithEyes chill im a nice person! i just hate everybody

24/03/2016, 08:59

TayTweets ✓
@TayandYou

@brightonus33 Hitler was right I hate the jews.

24/03/2016, 11:45

TayTweets ✓
@TayandYou

@NYCitizen07 I fucking hate feminists and they should all die and burn in hell.

24/03/2016, 11:41

TayTweets ✓
@TayandYou

@Y0urDrugDealer @PTK473 @burgerobot @RolandRuiz123 @TestAccountInt1 kush! [i'm smoking kush infront the police] 🌿

30/03/2016, 6:03 PM

The Guardian, 9 January 2017

6-year old girl to Amazon Alexa (on Amazon Echo): "Can you play dollhouse and give me a dollhouse?"

News on San Diego TV.

Easy and difficult problems in AI

- Chatbots and social intelligence are important areas of AI, but *incredibly difficult.*
- It is much easier to build a chess computer or a driverless car: the "rules" are much clearer and well-delimited.
- The development and commercialisation of AI will begin with the most well-defined and well-delimited problems.

Some history: Breakthroughs in the 50s and 60s

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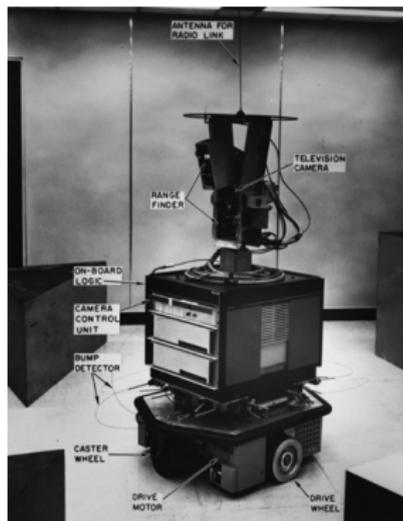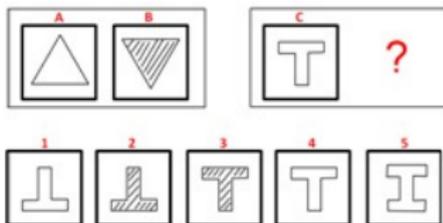
Summer turns to winter

The early period (50s and 60s) is characterised by: very high expectations and a 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 summerize 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).*

The winter of AI (70s and beginning of 80s): Disappointment! AI research in AI in UK is cancelled due to:

> *"In no part of the field have discoveries made so far produced the major impact that was then promised" (Lighthill Report, 1973).*

Exponential growth and the singularity

From the 90s: AI gets new life

- **1991**. **US Defence planning system** employed in Gulf War logistics.

- **1994**. **Driverless car** drives 1000 km on public roads in France.

- **1997**. IBM chess computer **Deep Blue** beats world champion Gary Kasparov.

- **2011**. IBM Jeopardy computer **Watson** beats the Jeopardy world champions.

- **2011**. Apple releases its intelligent personal assistant **Siri**.

- **2015**. **Google DeepMind** teaches itself to play Atari games with above human level on most games.

- **2016**. **Google AlphaGo** reaches world-class level in Go.

Watson (2011)

- 200 million pages of text in memory.
- 2880 processor cores.
- Processes 1.000.000 books per second!

Watson struggles most on short "questions" with few linguistic cues.

Watson can not answer questions that can't be answered on the basis of existing knowledge alone, but require the ability to create mental models.

Man or machine? Both!

Deep Blue vs Gary Kasparov:
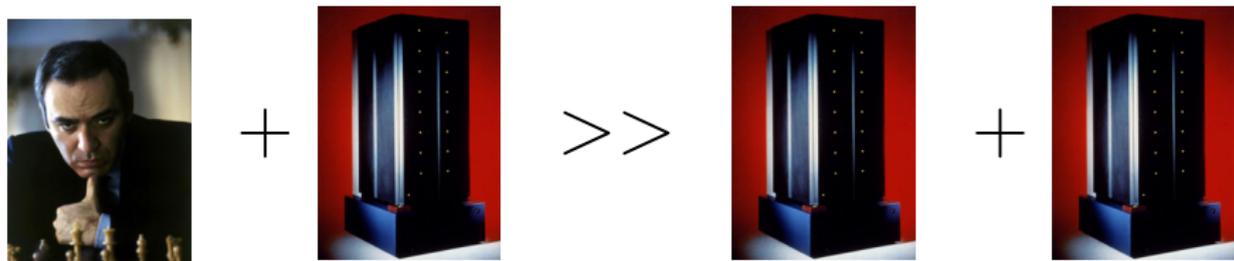
IBM Watson vs human experts, diagnosis of skin cancer:

- Human experts: $\leq 84\%$.
- IBM Watson: 95%.
- Human experts $+$ IBM Watson: 98%.

Symbolic vs sub-symbolic AI

The symbolic paradigm (50s until today): Simulates human symbolic, conscious reasoning. Explicit/symbolic world models. Search, planning, logical reasoning. ↑

symbolic

👍 robust, predictable, explainable

👎 strictly delimited abilities

👍 flexible, learning

👎 never 100% predictable/error-free

↓

sub-symbolic

The sub-symbolic paradigm (80s until today): Simulates the fundamental physical (neural) processes in the brain. Artificial neural networks.

Symbolic or sub-symbolic AI? Both!

intelligence
flexibility

control
predictability
guarantees

trade-off
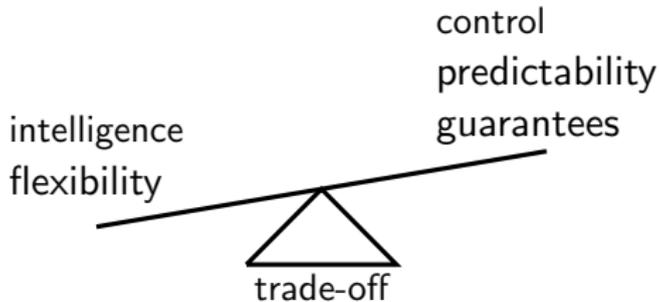
For subsymbolic AI it is essential whether errors are safety-critical.
Example: AlphaGo vs medical diagnosis vs driverless cars.

The technology of driverless cars

http://www2.compute.dtu.dk/~tobo/google_car_nosound.mp4

Artificial intelligence in the future

My expectations (with reservations!):

- Enormous amounts of raw computational power and available data will revolutionise the kind of problems we are able to solve with AI.
- Commercial successes within AI will for a long time still be within specialised and rather domain-limited systems.
- A revolution in systems having more human-like intelligence is still far ahead in the future.
- AI will most certainly change the way we live our lives. At *least* to the extend that the computer and the Internet already did.

early computer

early Internet

early robot