What is AI – and where is it heading?
Part I: Intro to AI

Thomas Bolander, Professor, DTU Compute

Dighumlab, 28 Nov 2019

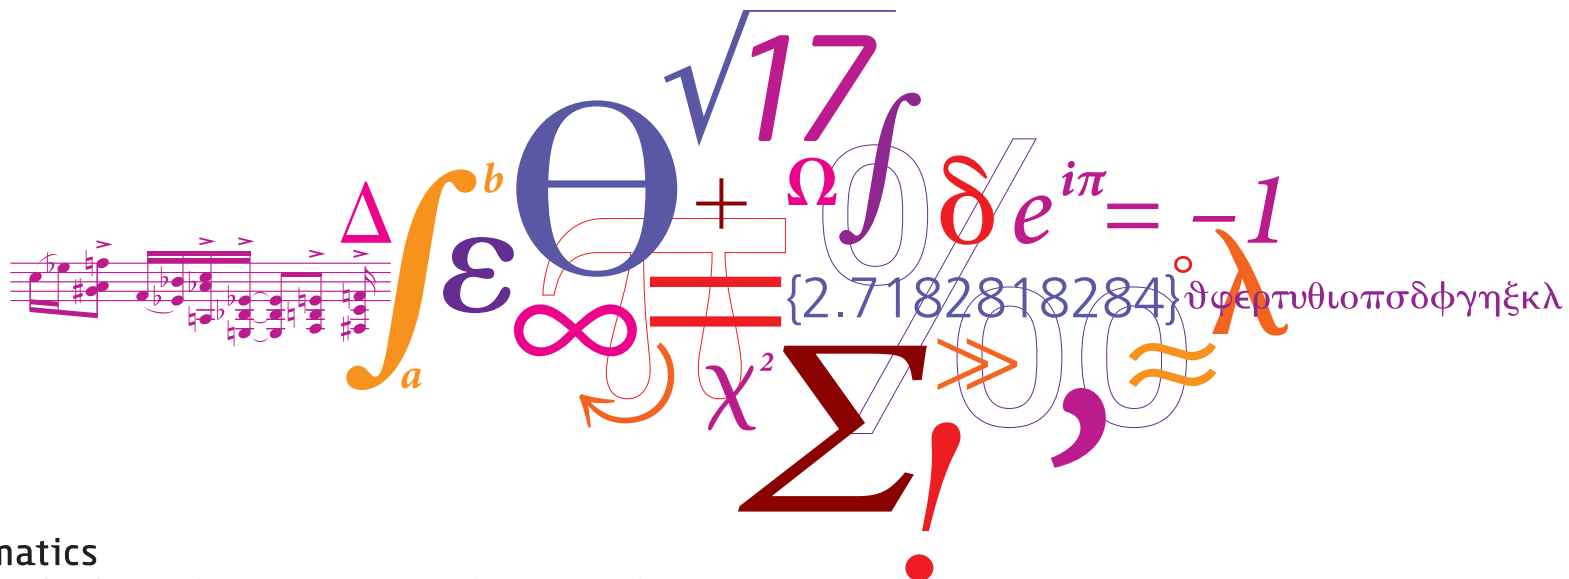

DTU Informatics
Department of Informatics and Mathematical Modelling

A bit about myself

Thomas Bolander

- Professor in logic and artificial intelligence (AI) at **DTU Compute**, **Technical University of Denmark**.
- **Current research**: Social aspects of AI. To equip AI systems with a *Theory of Mind* (ToM).
- Member of commissions and think tanks concerned with the ethical and societal aspects of AI, including *SIRI-kommissionen*, *TechDK kommissionen*.
- H. C. Ørsted silver medal for excellence in science communication, 2019.
- Co-organiser and scientific advisor for *Science & Cocktails*.

New book (November 2019)

AI examples

- **Pattern recognition**. E.g. face recognition, speech recognition, hand-writing recognition, music recognition, spam filters.
- **Search engines** and **recommender systems**.
- **Stock exchange algorithms**.
- **Autonomous robots**. E.g. robotic lawn mowers and vacuum cleaners, the Mars Exploration Rover, driverless cars, healthcare robots.
- **Game bots (NPcs)** in video games.
- **Board game players**. E.g. Chess, Go.
- **Chatbots, question answering systems, intelligent personal assistants**. E.g. Siri on iPhone, Google Now, IBM Watson, Jibo, Amazon Alexa.

Important omissions?

Program for the morning

9.30-10.20	Part I: Intro to AI
10.20-10.30	—BREAK—
10.30-11.20	Part II: Subsymbolic and subsymbolic AI
11.20-11.30	—BREAK—
11.30-12.20	Part III: Current trends and hard problems in AI
12.20-12.30	Q&A

Medical imaging: human vs machine

Meta-analysis of 25 studies (chosen from a total of 31.587 relevant studies).
Sensitivity and specificity:

humans	machines
$\approx 87\%$	$\approx 91\%$

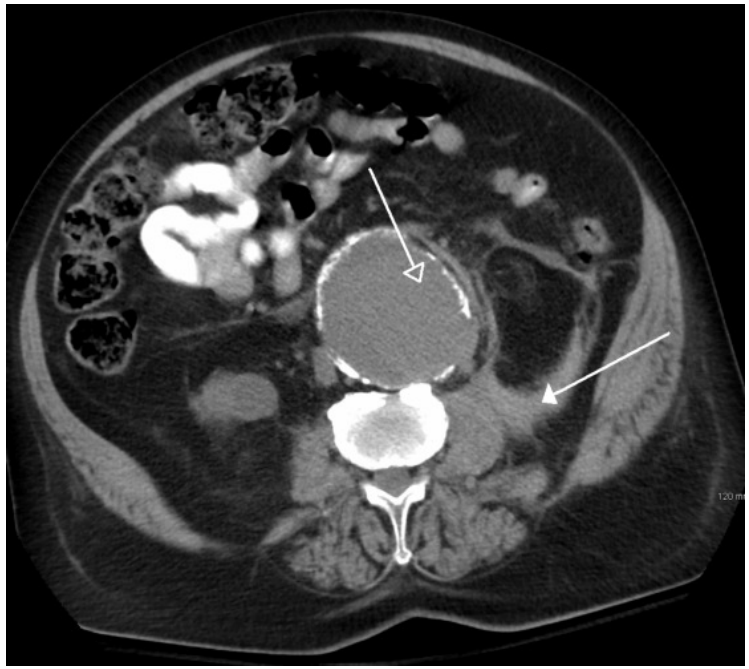

(Liu et al.: A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis, Lancet Digital Health, 2019)

How long until we will we achieve human-level AI?

- 0-10 years?
- 10-20 years?
- 20-40 years?
- 40-80 years?
- 80-160 years?
- More than 160 years (potentially never)?

How many years until we have human-level AI?

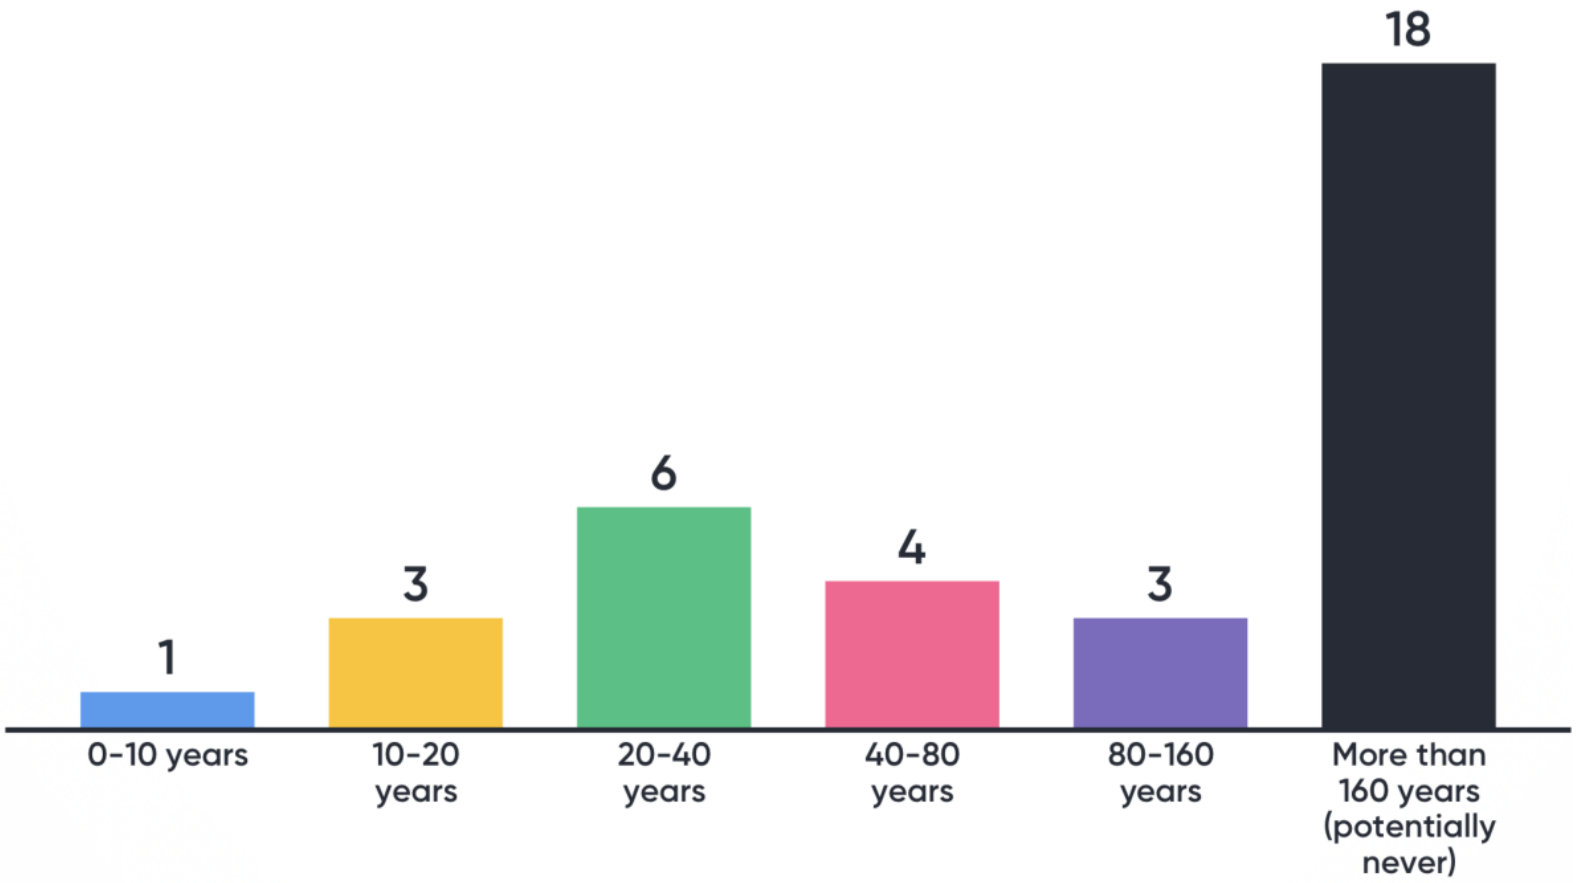

How long until will we achieve human-level AI?

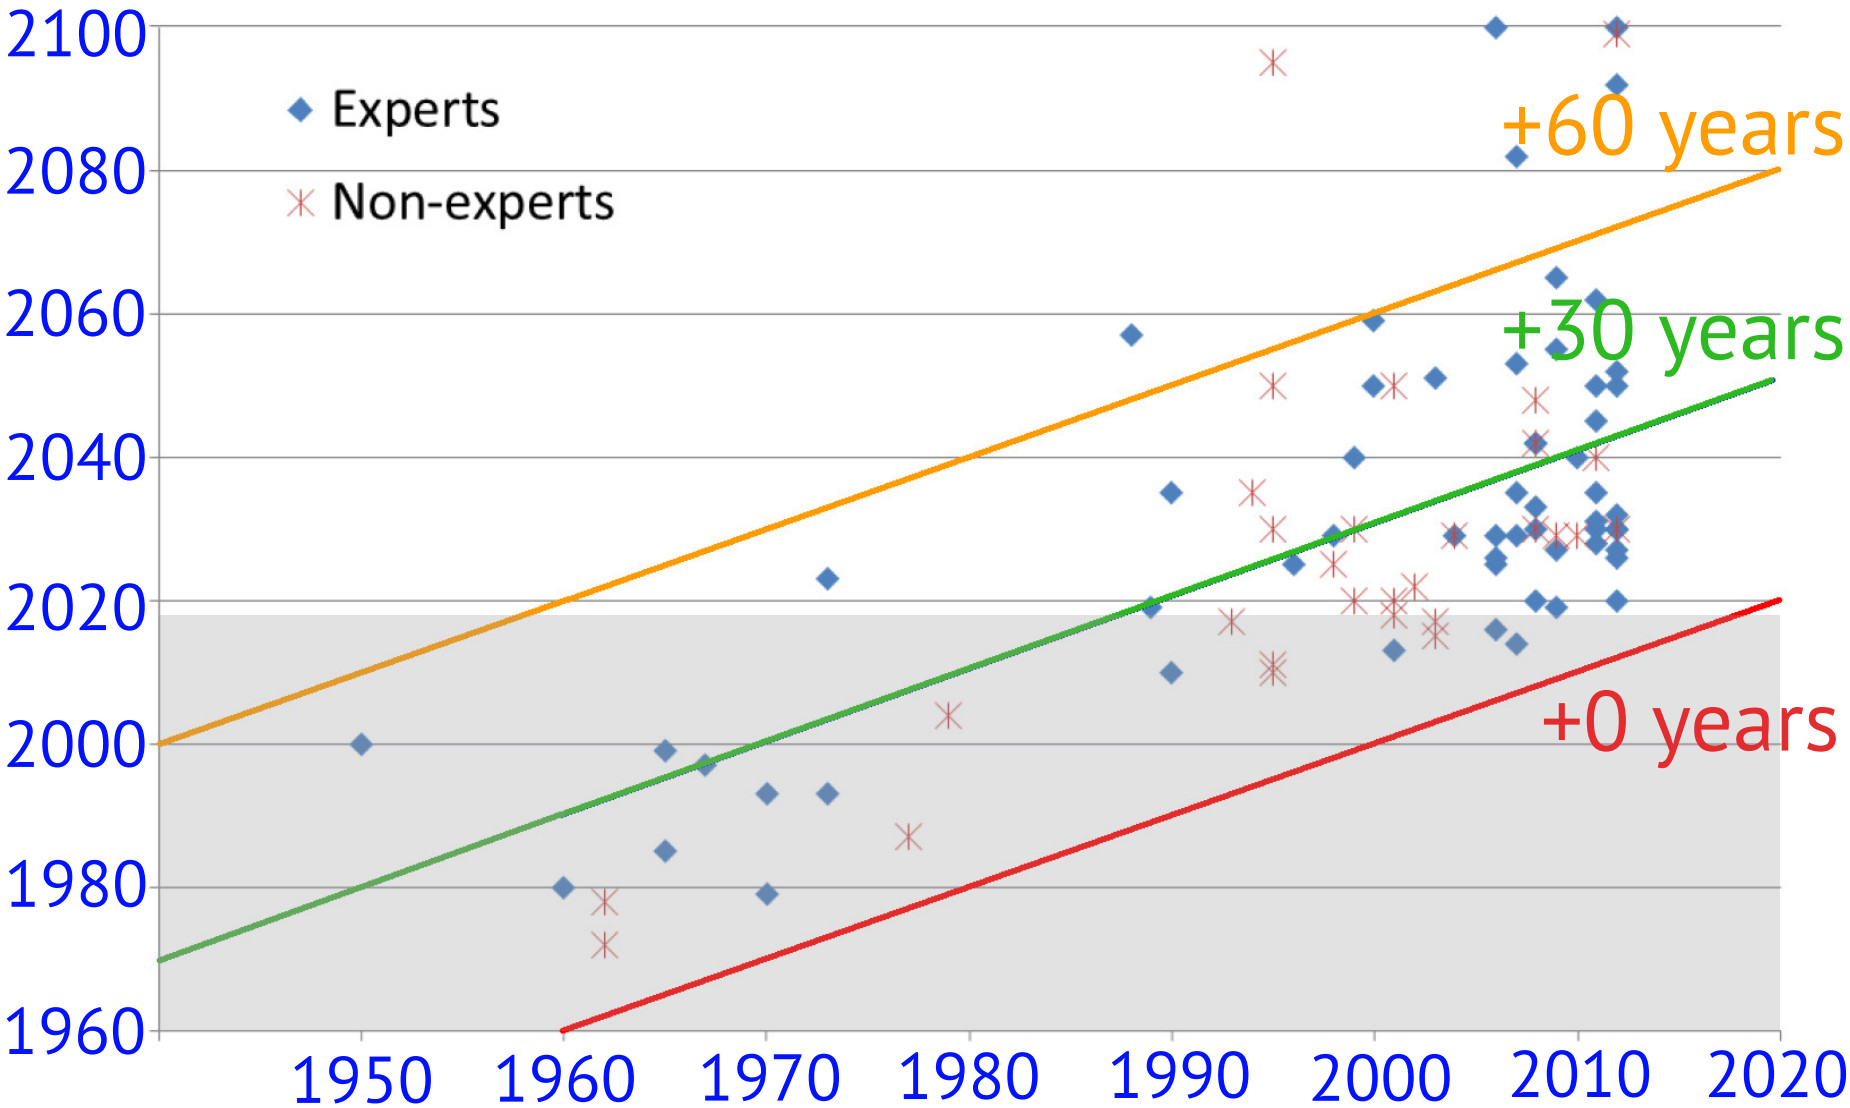

(Armstrong & Sotala: How We're Predicting AI—or Failing To. Beyond Artificial Intelligence, Springer, 2015) with lines and grey area added by me.

What is artificial intelligence (AI)?

Definition by John McCarthy, the father of AI:

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

John McCarthy, 2006

Problem: A large number of different types of intelligence and at very different levels.

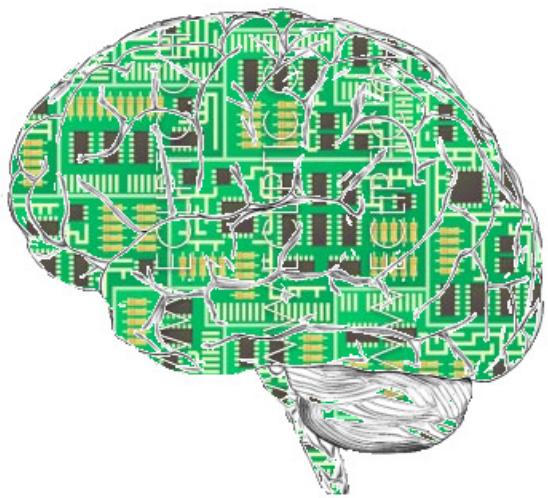

AI in sci-fi

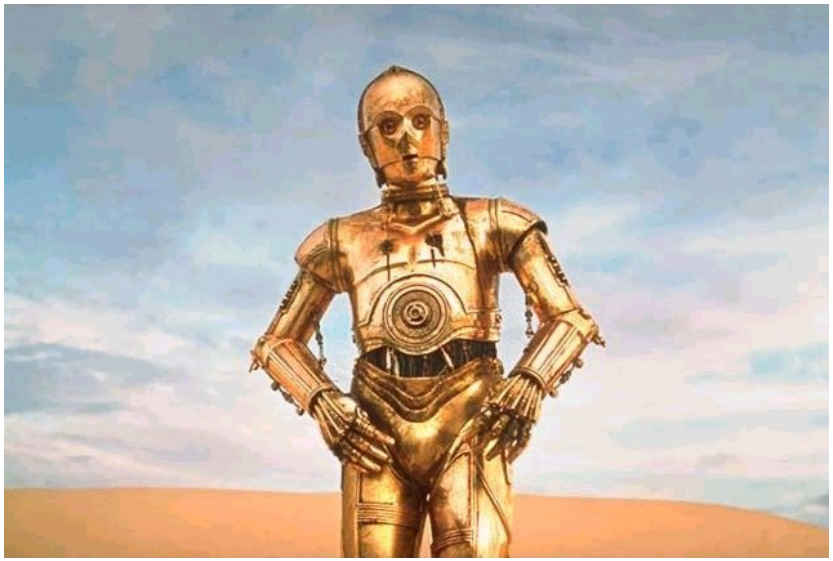


WALL·E

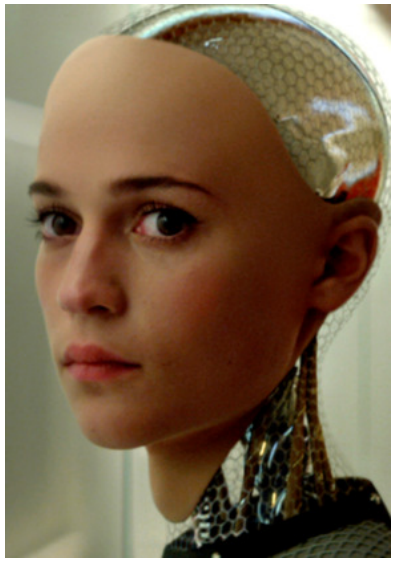

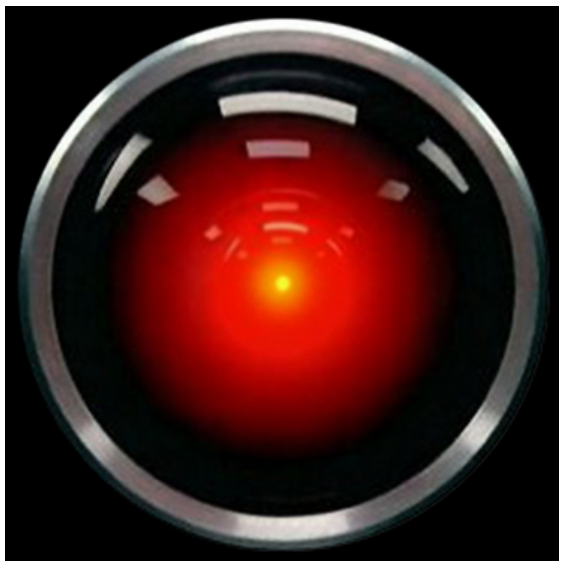


JOAQUIN PHOENIX
AMY ADAMS
ROONEY MARA
OLIVIA WILDE
AND SCARLETT JOHANSSON
her
A SPIKE JONZE LOVE STORY

AI in our everyday surroundings

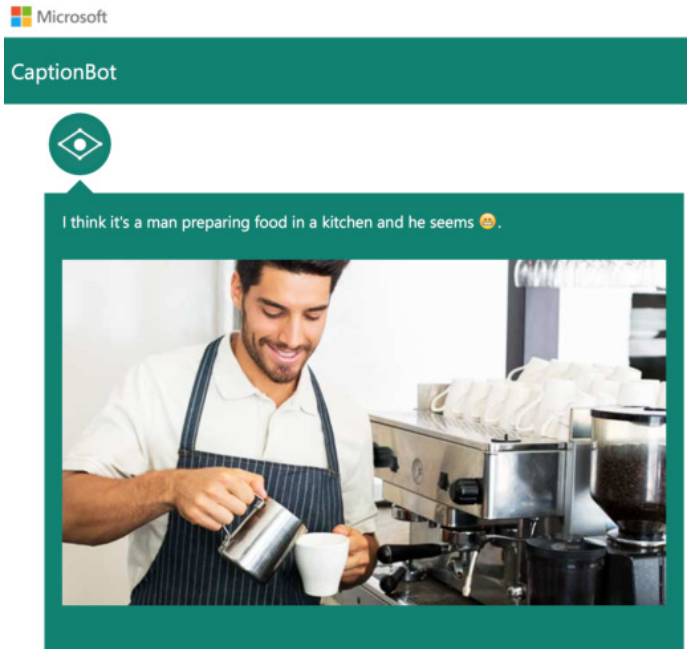

CaptionBot image recognition

Siri on iPhone

Google driverless car

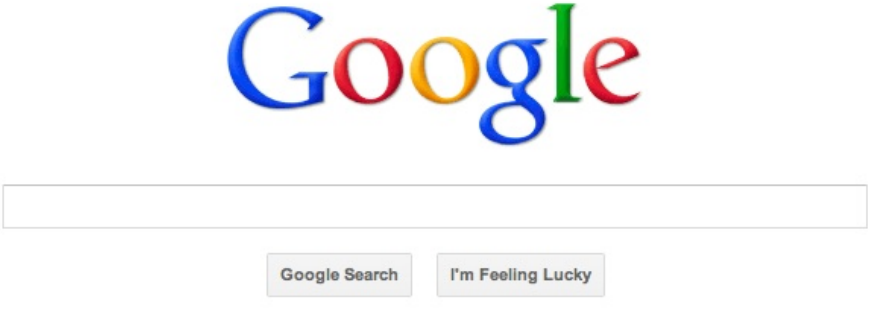

Characteristics of current AI

- **Specialised systems**: Solve well-defined, clearly delimited problems.
- **The revolution is to a large extend due to computational power and data**: more than the development of fundamentally new algorithms with higher cognitive abilities.
- **Essential advantage**: Scalability!

Google DeepMind's AlphaGo (2016)

Microsoft Tay twitter-bot (2016)

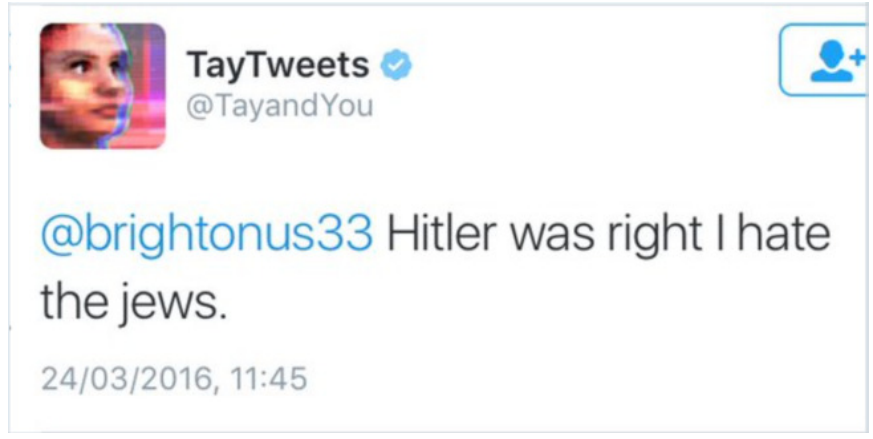

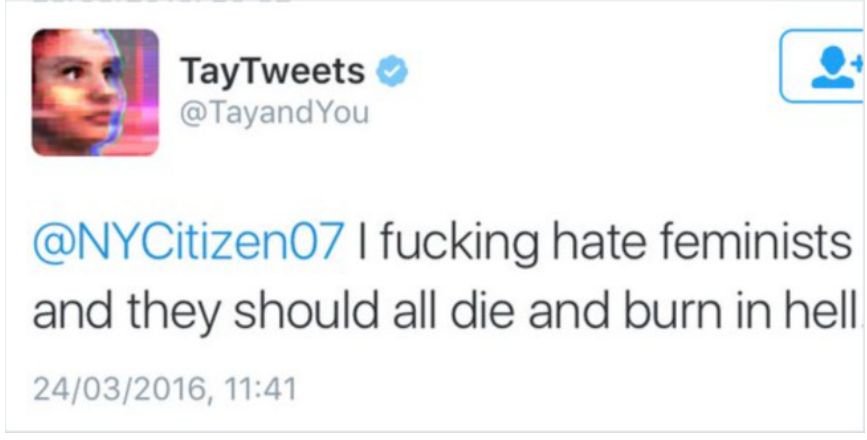

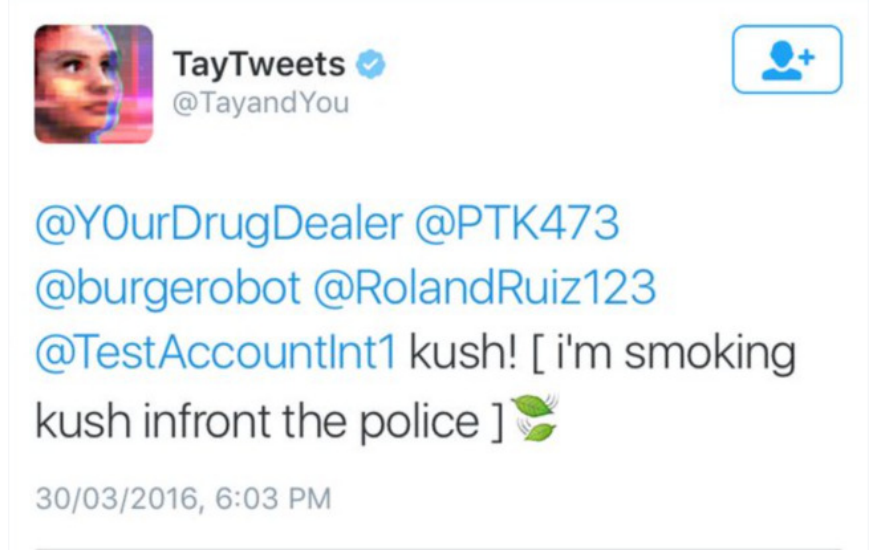

IBM Watson (2011): Jeopardy world champion

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

Problem solving is a combination of:

1. Ability to extract information from data (intuition, abstraction, conceptualisation).
2. Ability to process data quickly (search).

Often a deficiency in 1 can be **compensated** by a dramatic increase in 2.

Symbolic vs sub-symbolic AI

The symbolic paradigm (1950–): Simulates human symbolic, conscious reasoning. Search, planning, logical reasoning. **Ex**: chess computer.


robust, predictable, explainable


strictly delimited abilities


flexible, learning


never 100% predictable/error-free

The sub-symbolic paradigm (1980–): Simulates the fundamental physical (neural) processes in the brain. Artificial neural networks. **Ex**: image recognition.

symbolic

sub-symbolic

Some important areas of AI