The Rise of Artificial Intelligence and Machine Learning

Hector Geffner
ICREA & Universitat Pompeu Fabra
Barcelona, Spain

Artificial intelligence (AI)

Elon Musk: artificial intelligence is our biggest existential threat

The AI investor says that humanity risks 'summoning a demon' and calls for more regulatory oversight

Samuel Gibbs

@SamuelGibbs

Monday 27 October 2014 10.26 GMT

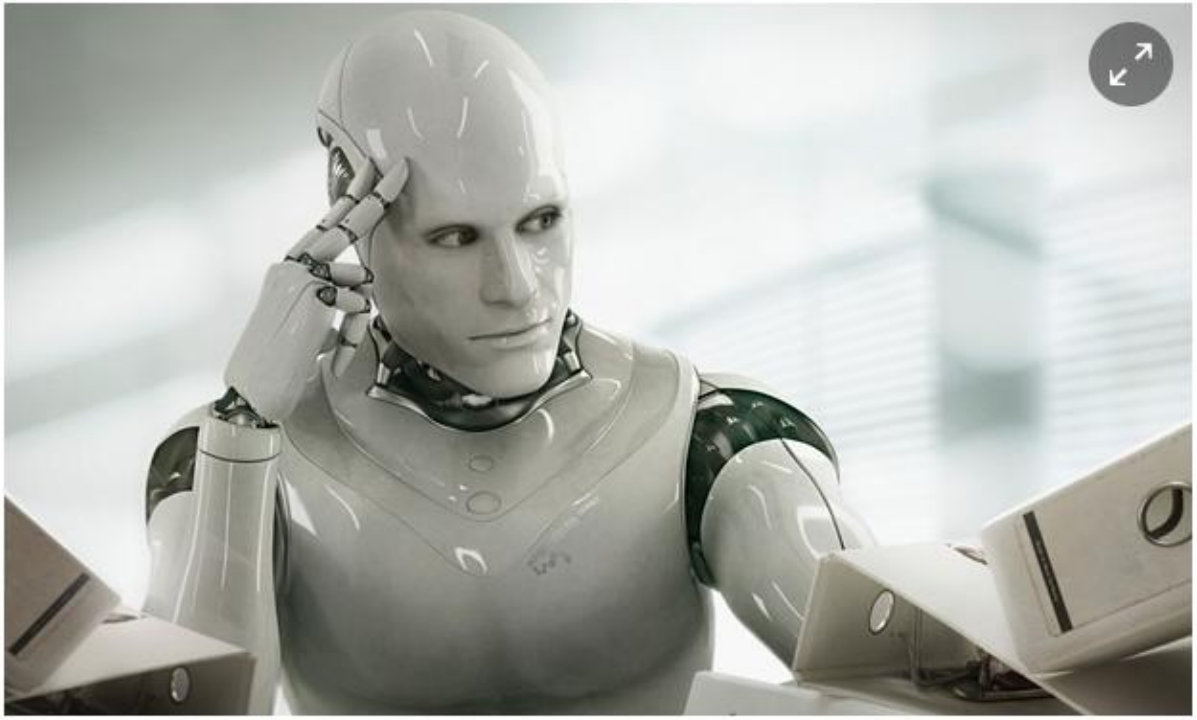

Artificial intelligence should be regulated, says Elon Musk. Photograph: Blutgruppe/Blutgruppe/Corbis

Elon Musk has spoken out against artificial intelligence (AI), declaring it the most serious threat to the survival of the human race.

Musk made the comments to students from Massachusetts Institute of Technology (MIT) during an interview at the AeroAstro Centennial Symposium, talking about computer science, AI, space exploration and the colonisation of Mars.

Most popular

The 10 best cities in world to be a stude 2016 - in pictures

Thai printer replace International New Times article with b space

The Switch

Bill Gates on dangers of artificial intelligence: 'I don't understand why some people are not concerned'

A | 337

By Peter Holley January 29 Follow @peterjholley

Bill Gates joined Reddit for an AMA on Wednesday. (Tobias Schwarz/AFP/Getty Images)

Most Read

1 This might be the most controversial theory for what's behind the rise of ISIS 

2 How China is like the U.S. a century ago

3 Why Americans dress so casually

4 These countries are the 'moral center' of the climate debate. Can Obama help them? 

5 Obama pledges aid to island nations threatened by climate change 

Stephen Hawking: Artificial intelligence could wipe out humanity when it gets too clever as humans will be like ants

AI is likely to be 'either the best or worst thing ever to happen to humanity,' Hawking said, 'so there's huge value in getting it right'

Andrew Griffin | @_andrew_griffin | Friday 9 October 2015 | 91 comments

Chinese inventor Tao Xiangli modifies the circuits of his home-made robot at his house in Beijing, May 15, 2013 *REUTERS/Suzie Wong*

But is this fear justified?

Has "artificial intelligence" finally "arrived"?

What is artificial intelligence

- **Artificial intelligence** is the area of **computer science** that studies **intelligent behavior** from a **computational point of view**
- Understanding a behavior X is understanding how to generate it by computer
- X may be solving a problem, recognizing an object in an image, understanding a joke, doing the dishes, etc.
- Some behaviors require some type of **"body"** (sensors, actuators): **robots**

History of artificial intelligence: A. Turing (1912–1954)

- Turing developed (in paper) the **universal, programmable computer**
- Current computers not more powerful than "Turing machine": compute exactly the same
- **Turing's goal was not practical**: it was about about abstract math problem
- Turing explained how his "machine" can exhibit **intelligent behavior** and how this can be determined (**Turing test**; 1950)

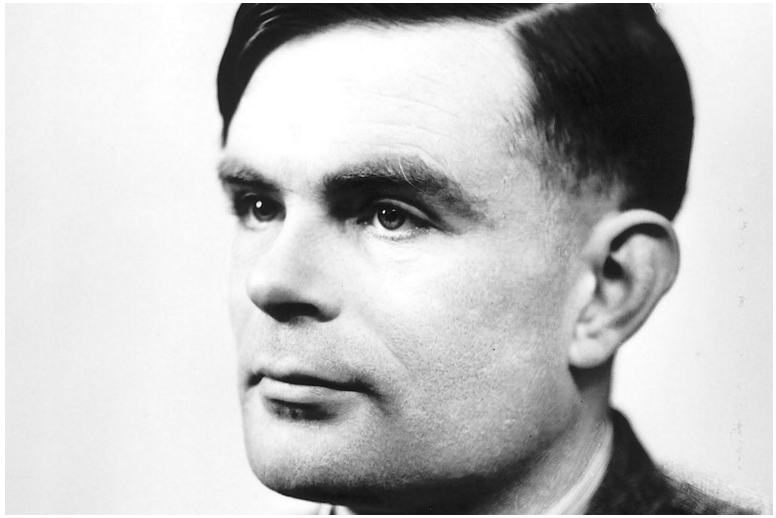

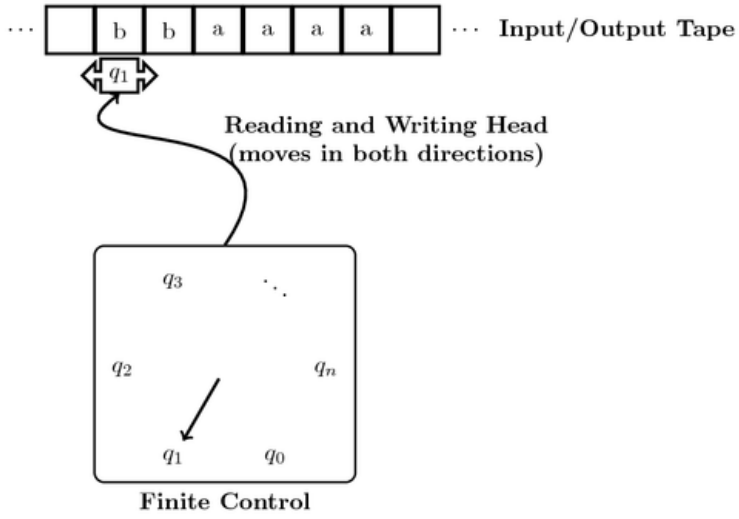

Artificial Intelligence: First Decades

- **Study of AI** takes shape at the end of the 50s
- **Challenge:** If claim was "machines cannot do X"; write program to do X
- **Progress:** Programs that play chess, solve puzzles, prove theorems
- **Challenges:** None with **abilities of a 5 years child**: motor, visual, linguistic

Why AI in the news?

Apple Siri: interaction through voice
Google cars: self-driving cars
Facebook: face recognition
Amazon, Netflix: personalized recommendations
IBM Watson: Jeopardy (TV EEUU) and other uses
Deep Mind: Atari, Chess, Go from **self-play alone**
. . .

New generation of AI applications
mostly based **on machine learning**

Machine Learning

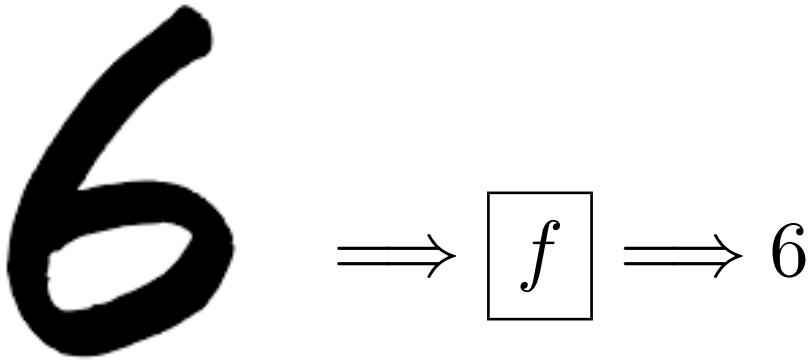

- We want a **function** $f(x)$ that maps image x to **digit** 0–9 it contains
- **Writing a program by hand** for computing f is too hard
- It's better to **learn** f from **examples:** $\langle x_1, f(x_1)\rangle, \ldots \langle x_n, f(x_n)\rangle$
- Usually **form** of f is given, like in **neural nets**, **parameters** learned

AI: Model-free Learners vs. Model-based Solvers

- **Machine learning** produced impressive applications but **no** general intelligence
- **Learning to associate stimulus with responses** important but insufficient
- **Limitations:**
  - ▷ no generality
  - ▷ no understanding
  - ▷ no explanation
  - ▷ no models
- Another thread in AI, concerned with **solvers that reason with models**
- **AI solvers** exhibit more **"intelligence"** in usual sense (reasoning, planning) but less visible in terms of applications

Learners and Solvers: System 1 and System 2?

Current **dual-process accounts** of the **human mind** assume two processes (Daniel Kahneman: Thinking, Fast and Slow):

System 1 (Intuitive Mind)	**System 2** (Analytical Mind)
fast	slow
associative	deliberative
unconscious	conscious
effortless	effortful
parallel	serial
specialized	general
...	...
Learners?	**Solvers?**

The Rise of AI and Machine Learning

- **AI** is fascinating subject; a lot has been learned, much more yet to learn
- **Human-level AI** is far, but **current AI** can be used for good or ill
- Efforts to **align AI with human values**, nice, but other forces at play
- **Markets and politics** focused on **bottom line** and aimed at our **System 1**
- Life in modern world need **System 2 informed by facts and common good**
- If we want **good AI**, we need a **good and decent society**