Part III: Introduction to Artificial Intelligence

Ashkan Panahi and Claes Strannegård

Department of Computer Science and Engineering
Chalmers University of Technology

Lecture 1: Overview

DO NOT DISTRIBUTE

Contents

Overview of Part III

Teaching team

- Ashkan Panahi (ashkanp@chalmers.se), Computer Science and Engineering, lecturer, responsible for Part III
- Emilio Jorge (emilio.jorge@chalmers.se), Computer Science and Engineering, TA

Structure of part III

Part III consists of three modules: one module per week.

1. Week 1
    1. Lecture 1: today
    2. Assignment 1: read, summarize, present a text about the impact of AI
2. Week 2
    1. Lecture 2: Deep learning
    2. Presentations of Assignment 1 (Part 1)
    3. Assignment 2: Construct and run deep neural networks.
3. Week 3
    1. Lecture 3: Search algorithms
    2. Presentations of Assignment 1 (Part 2)
    3. Assignment 3: Get familiar with A* and other search algorithms

Grades of part III

The modules are graded as follows:

1. Module 1: Fail/Pass
2. Module 2: Fail/3/4/5
3. Module 3: Fail/3/4/5

To get grade 3 on part III, you need to get Pass, 3, 3.
For grade 4 you need Pass + average 4.
For grade 5 you need Pass + 5 + 5.
The same grade will normally be assigned to both group members. In case of strongly unbalanced contributions, please write a note about that on your hand-ins.

Assignment 1

- Work with your group mate
- Select a text either from the list of suggestions or your own text, and **first get it approved** by the TA.
- Prepare the summary of the text in 10-15 slides.
- Upload your slides as a **PDF file** to Canvas.
- Present a summary of your text to the class. The presentations **must be 10 minutes each**, so you should be ready to skip some of the slides.
- Grades: Pass/Fail

Assignment 2

- Read "Implementing our network to classify digits" in Nielsen's online book.
- Work with your group mate
- Install the accompanying code that is available on github
- Tasks
  - Grade 3: Experiment with different architectures, report results
  - Grade 4: Experiment with noise on dataset
  - Grade 5: Implement CNN. Experiment with shift.
- Upload your solution to Canvas.

Assignment 3

- Work with your group mate
- Tasks
  - Grade 3: Go to http://qiao.github.io/PathFinding.js/visual/. Build a simple world with a few hurdles and run BFS, DFS, and A*. Explain the differences in behavior and efficiency.
  - Grade 4: Study IDA* and Dijkstra. Explain briefly how they work. Test these algorithms too on your simple world. Comment on the results.
  - Grade 5: Study MCTS. Explain briefly how it works. Discuss how it might navigate your world.
- Upload your solution to Canvas.

The AI revolution

What is AI?

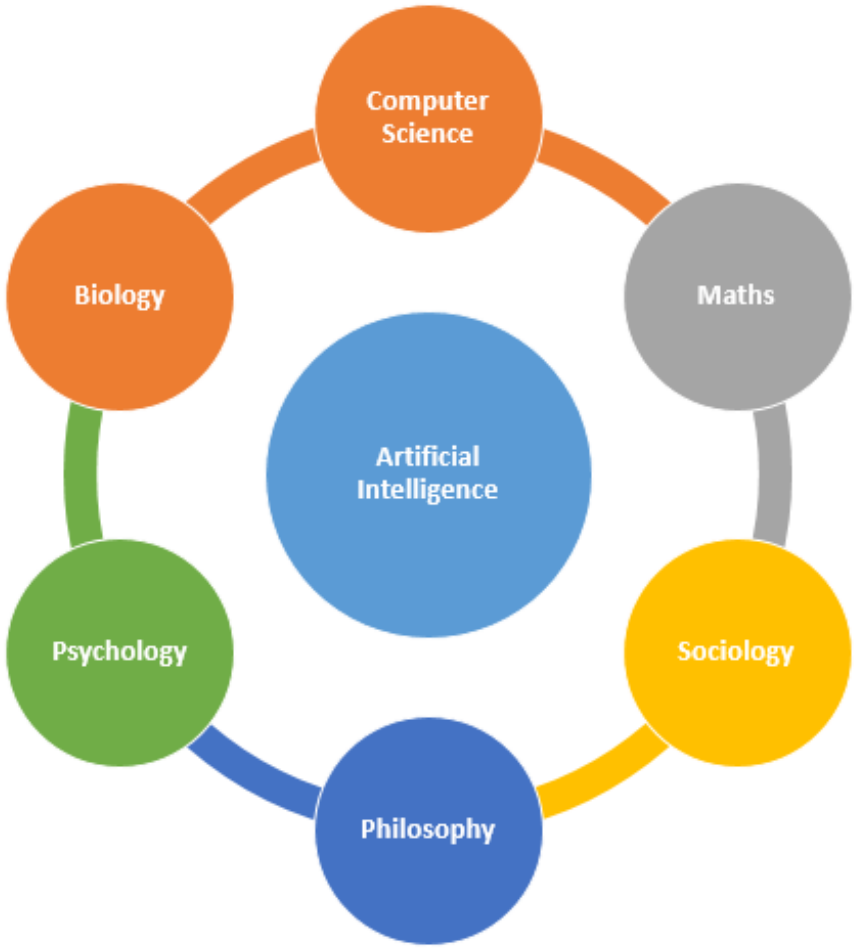

A bunch of technologies? A clearly defined concept? A hype? A moving target? A branch of computer science?

Sources

- Online book chapter. David L. Poole and Alan K. Mackworth, Cambridge University Press, 2017: Artificial Intelligence: Foundations of Computational Agents. Read Chapter 1: Artificial Intelligence and Agents up to and including 1.4.

Industrial revolutions

0. Agriculture
1. Steam Engine
2. Electricity
3. Computer
4. Artificial Intelligence

The AI revolution

The AI revolution

The AI revolution

AI will let scientists solve some of the worlds "hard problems," like population growth, climate change, human development, and education.

Eric Schmidt, President of Alphabet

The AI revolution

Keynote at the *Rise of AI* conference 2019: "AI for zero hunger, no poverty, and good health"

The AI revolution

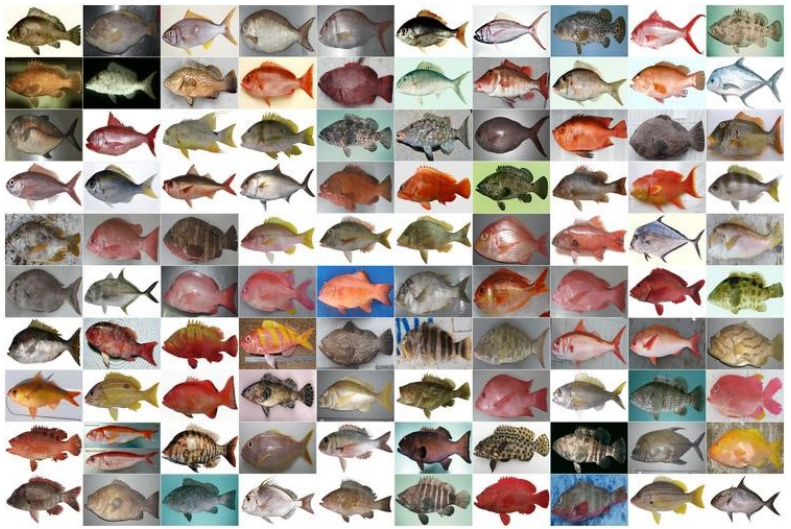

Sustainable AI examples from Chalmers:

- Green computing (ZeroPoint Technologies): Compression algorithms leads to lower energy consumption. IT will represent 20% of the energy consumption worldwide in 2025 (Andrae, 2017).
- Sustainable fishing (Refind Technologies): Onboard monitoring of fish catch (species, size, sex, quantity) enables sustainable harvesting.

Classical AI and machine learning

AI technology 1: Machine learning

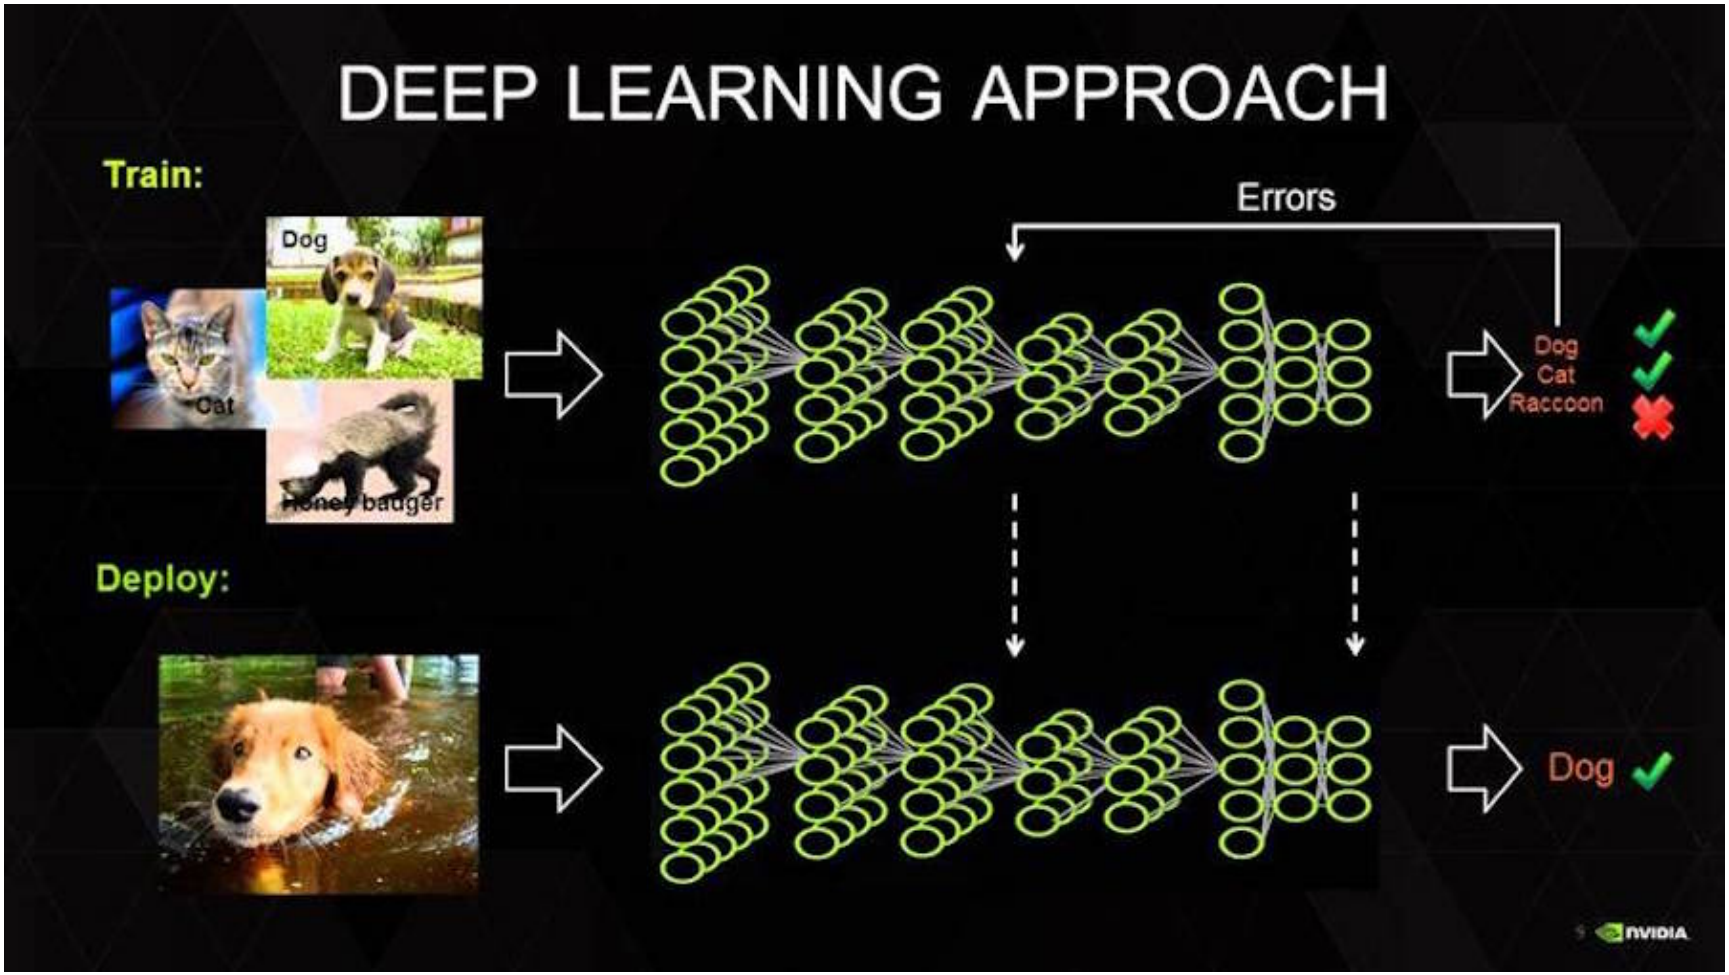

AI technology 2: Classical AI

Problem

Given a system design D (the brake system of the car):

- $brakeLight \leftrightarrow (manualBrake \vee autoBrake \vee hazard)$
- ...

and a set of safety requirements R_i:

- $manualBrake \rightarrow brakeLight$
- ...

prove that the design meets all the safety requirements.

Solution

Use automatic theorem proving, which is based on search, to find a formal proof of each R_i from the assumptions in D. (If no proof is found, then the design needs to be modified.)

Which method is better?

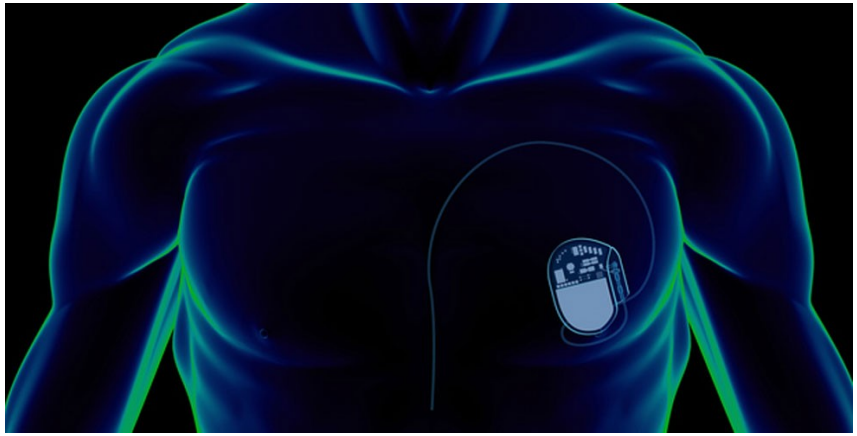

- For recognizing dogs: only machine learning methods work.
- For proving safety requirements: only classical methods work.
- Both methods are needed! They solve different problems.

AI and applications

Applications of AI

- Object recognition: speech (Google), hand-writing (Apple), faces (Facebook), crime scenes (Interpol)
- Recommender systems: books (Amazon), songs (Spotify), movies (Netflix), news items (DN), relationships (Tinder)
- Search engines: English (Google), Chinese (Baidu)
- Language technology: translation (Google translate), cognitive assistants (IBM), summarization (Algorithmia)
- Robotics: industrial robots (ABB), animal robots (Boston Dynamics).
- Transportation: navigation, active safety, self-driving vehicles (Tesla)
- Games: chess, Go, and many video games (Google DeepMind)

AI and art

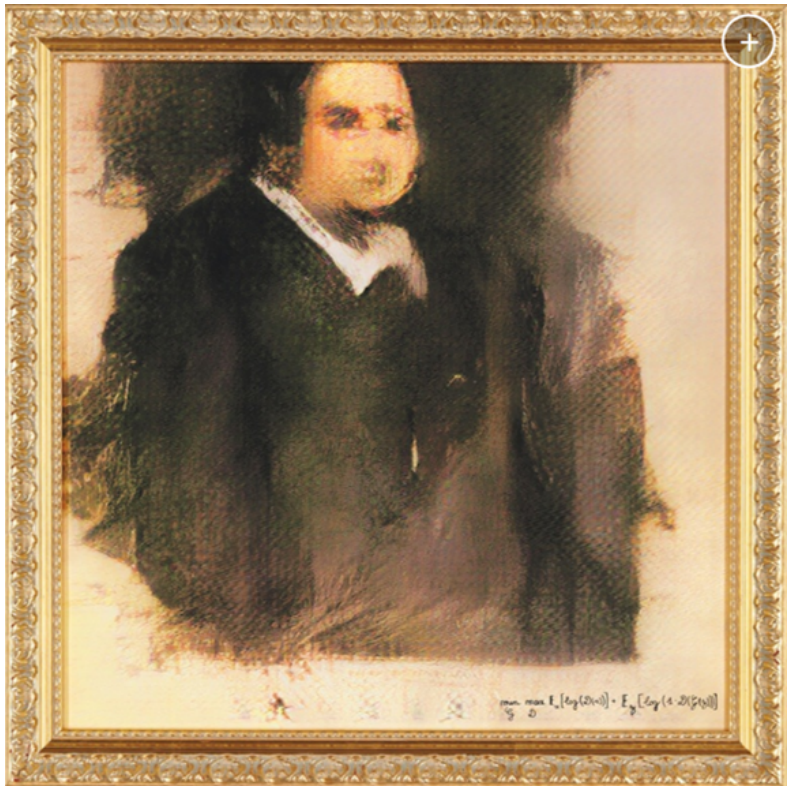

Left. Image generated by a neural network trained on human-generated art, 2018. Sold for 432,500 USD at Christie's.
Right. Image generated interactively, Ahmed Elgammal, 2017. Source

AI and music

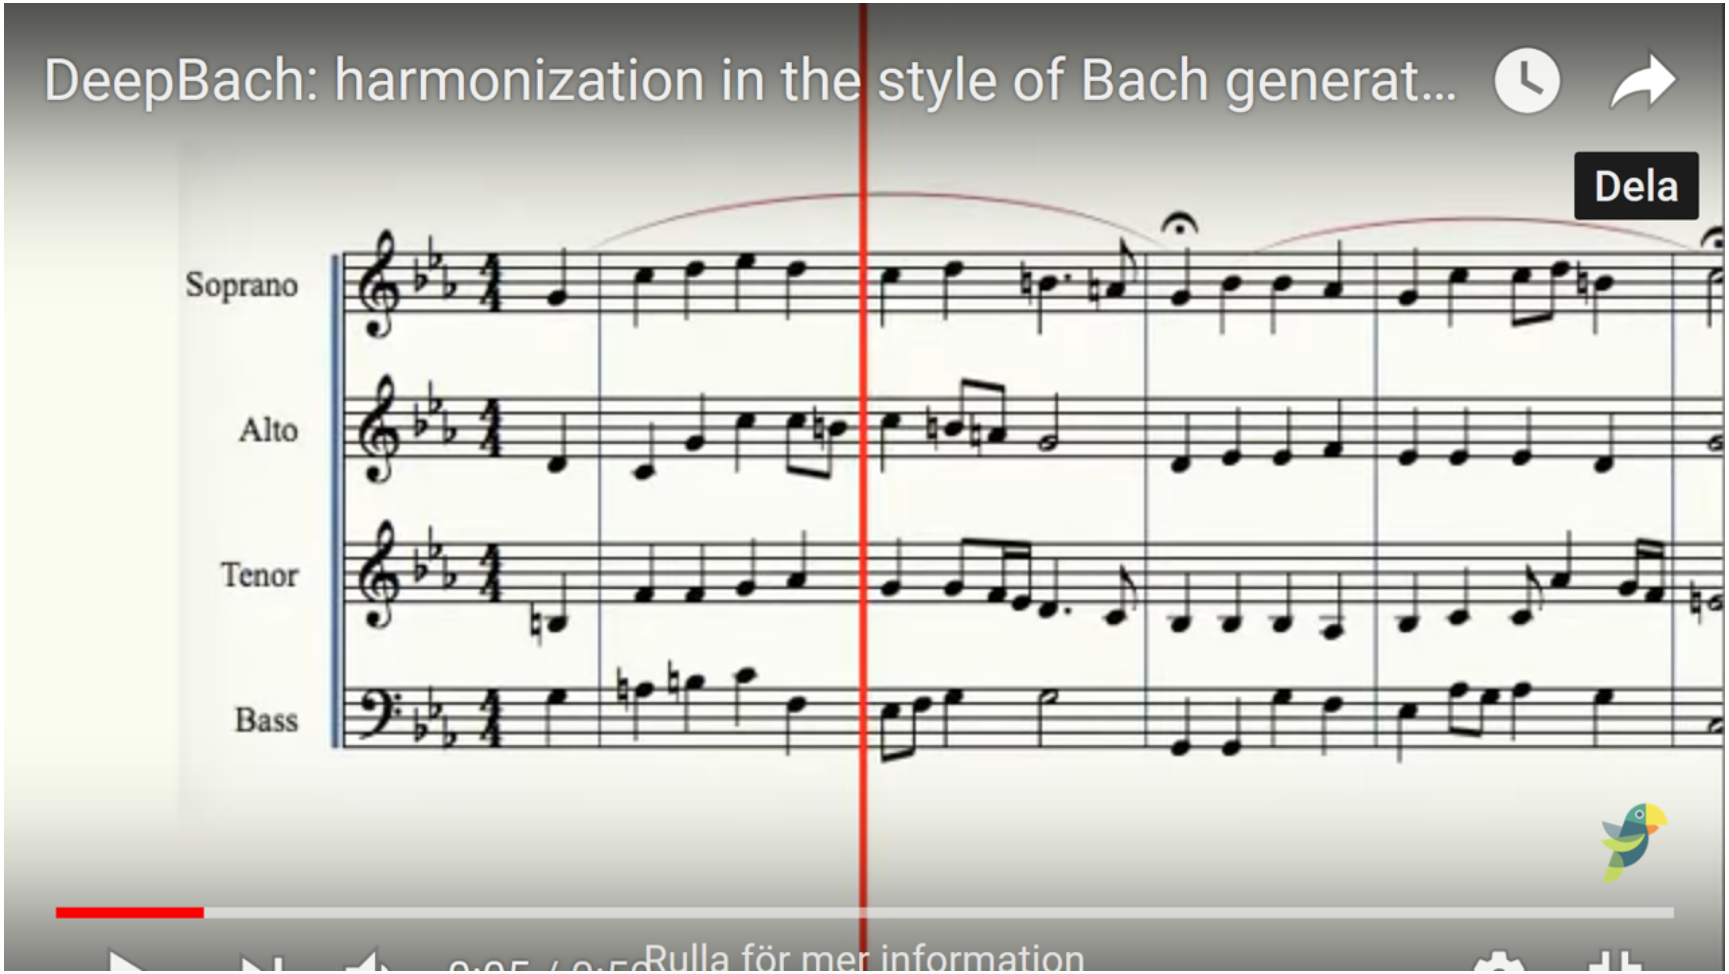

DeepBach DeepBeatles

AI and business

AI in business

- Estimated contribution to the global economy by 2030 from AI: \$16 trillion. Comparison: GDP of USA 2017: \$19 trillion.
- Up to 26% boost in GDP for national economies from AI by 2030

Source: Global Artificial Intelligence Study, July 2018.

AI startups

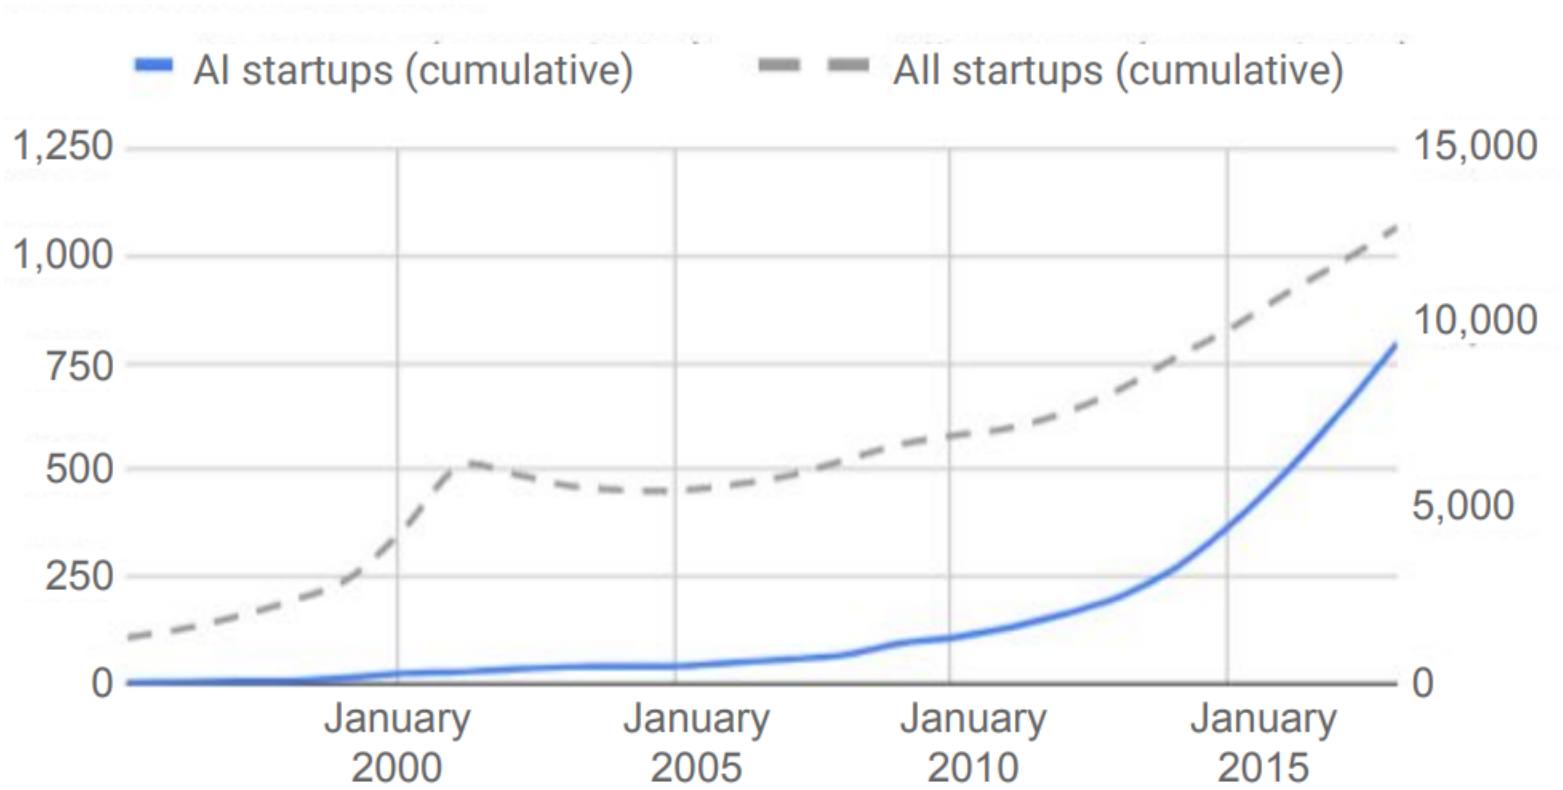

Source: MIT Technology Review, December 2018

AI investments

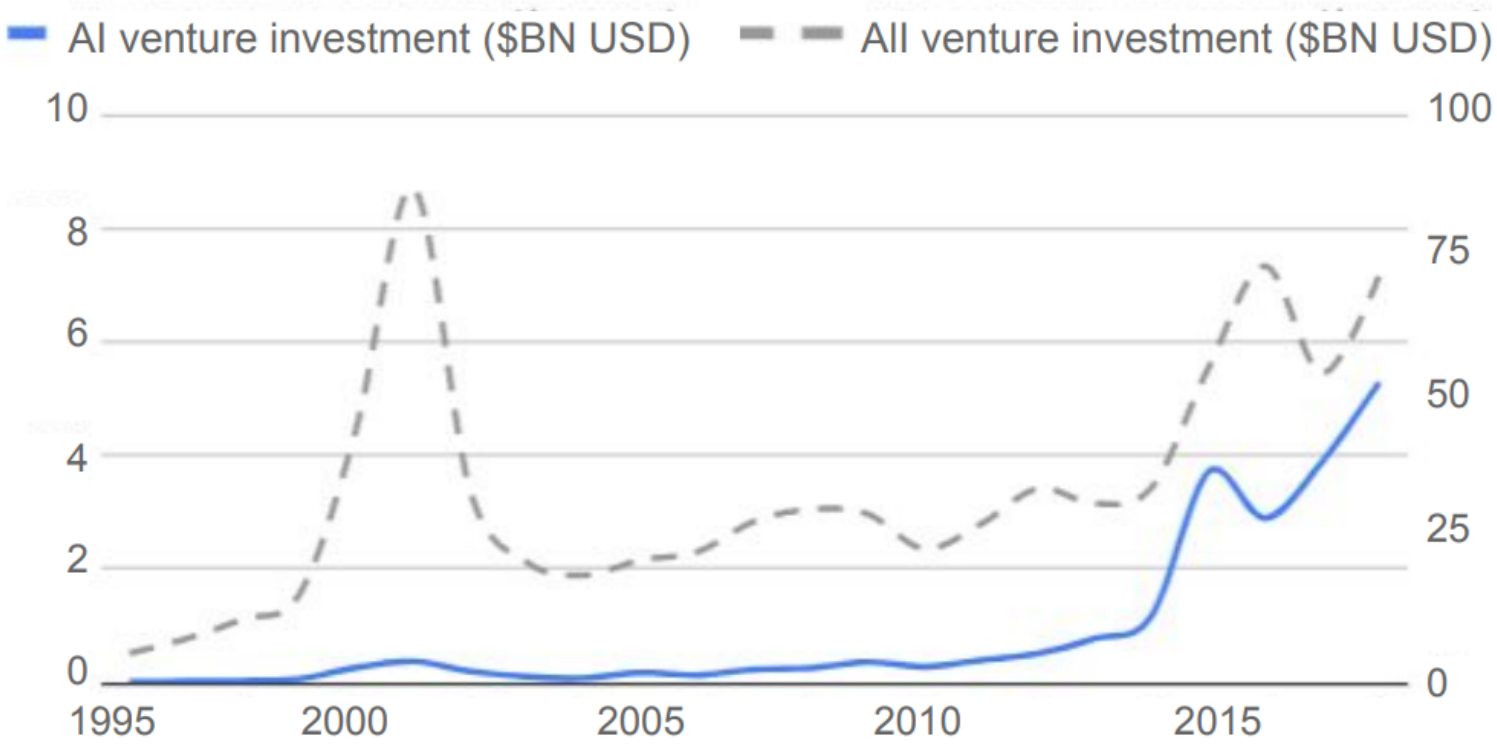

Source: MIT Technology Review, December 2018

AI research activity

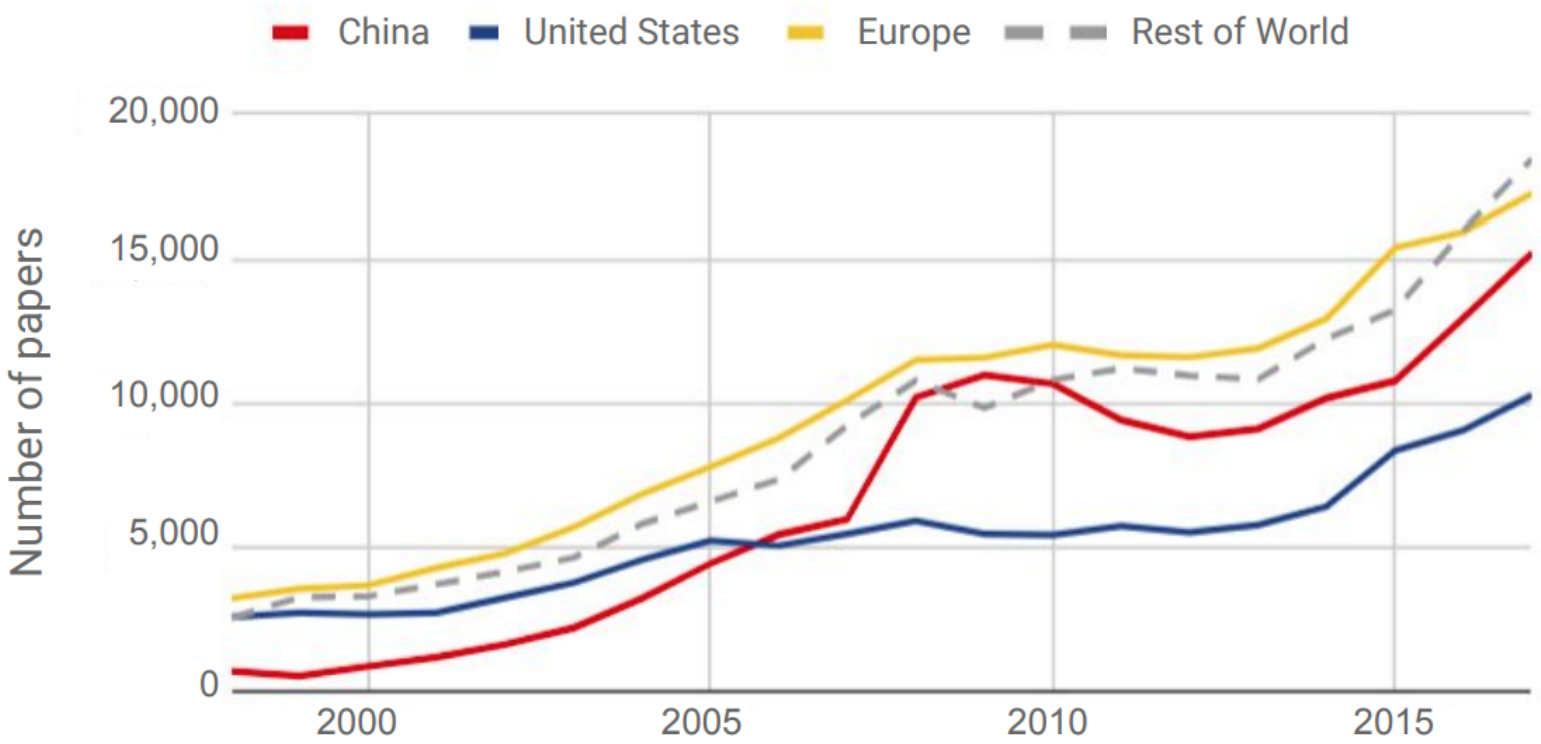

Source: MIT Technology Review, December 2018

AI and jobs

AI and jobs

"AI will create as many jobs in the UK as it will displace over the next 20 years. It will boost economic growth and create new roles as others will fall away." Global Artificial Intelligence Study, July 2018

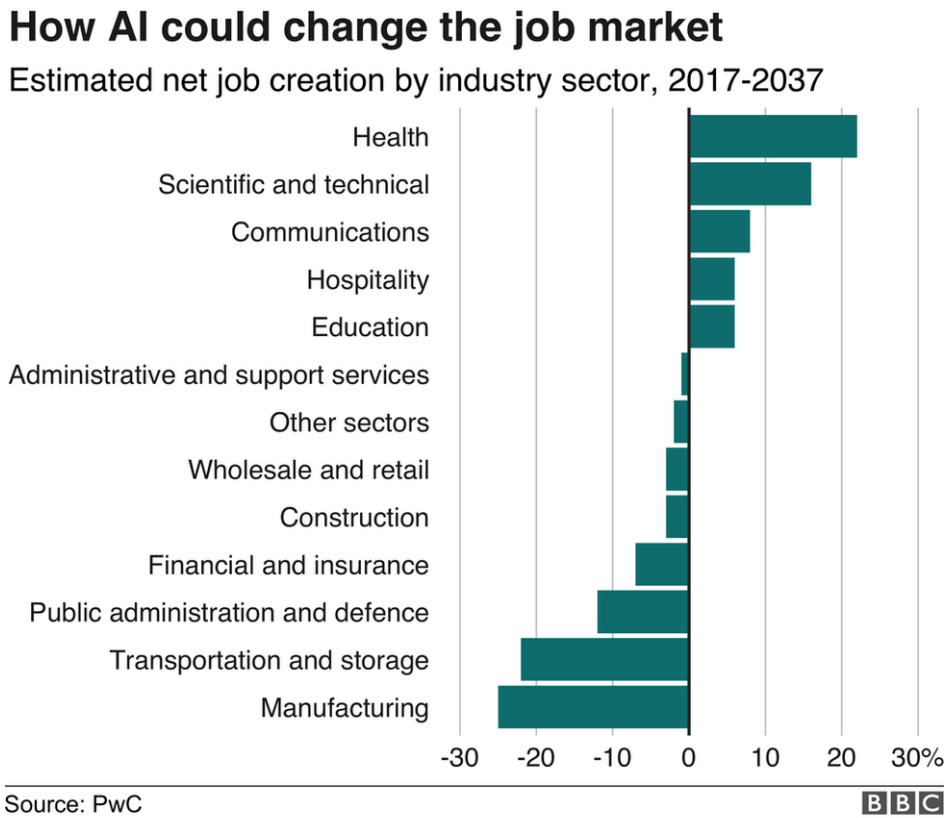

AI and ethics

AI and ethics

- This robot did not understand that it destroyed the kite. It would not understand if it hurt a human.
- It has no ethical values whatsoever, no concept of good and bad. This means its behavior can be very bad sometimes.
- Ethical values must be actively inserted into the systems!

AI and ethics

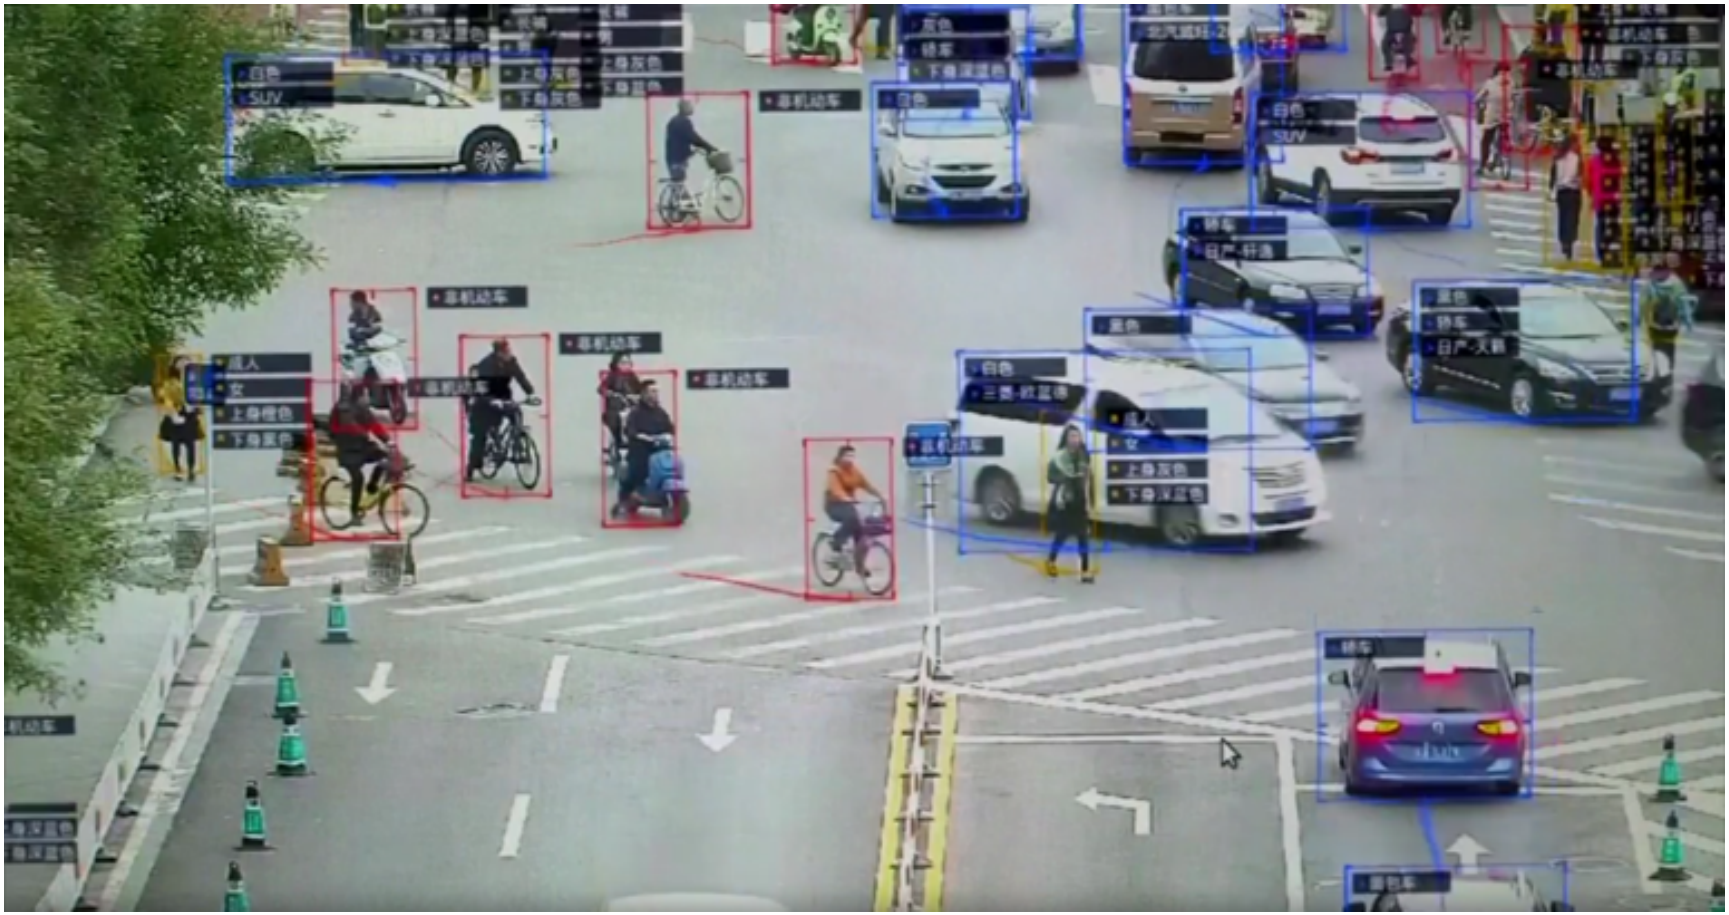

AI enables mass surveillance. Image recognition (faces and licence plates) can be used for detecting traffic offenses.

AI and ethics

Steve Omohundro: AI systems with harmless goals are not necessarily harmless! A physical chess playing robot with the goal of playing as many chess games as possible against humans might resist attempts to turn it off or re-program it, also with violence.

AI and risks

- AI comes with risks that need to be understood.
- Current risks: unemployment, inequality, financial instability, surveillance, autonomous warfare
- Future risks: existential threats, robot take over scenarios
- Clearly, AI needs to be developed side-by-side with strategies for using AI responsibly!

The early history of AI

Some highlights:

- Hobbes (called the "Grandfather of AI" by Haugeland) suggested that thinking is a mechanical process in the 17th century.
- Turing introduced his a-machines (Turing machines) in 1936
- McCulloch and Pitts showed that Turing-complete machines could be built from "formal neurons" in 1943
- Turing built the first chess program in 1950
- McCarthy coined the term "artificial intelligence" in 1955
- Wang implemented a program that proved every logic theorem (nearly 400) in *Principia Mathematica* in 1960
- Samuel built a checkers program capable of learning that beat a human champion in 1961.

AI and the future

The future of AI

- Existing applications are domain-specific and can only operate in well-delimited environments.
- Tasks that are relatively easy for humans are often hard for AI and vice versa.
- The coffee making challenge remains unsolved: make a robot that can come to a random home and make a cup of coffee.
- We do not have any service robots that can cook, do the laundry, and clean the house. That would require adaptability to new environments, e.g. private homes.

Artificial General Intelligence

- *Artificial General Intelligence* is defined as "programs with human-level general intelligence"
- The goal is to write a program that can learn to solve all kinds of problems by training (like a child or a dog)
- Today we have to construct one program manually for each task
  In the future we might have single programs that can be trained for thousands of different tasks.

The future of AI

Mike Schroepfer, CTO of Facebook (MIT Tech Review, Jan 2016):

- Many people apparently have the wrong idea about how rapidly the field is moving
- People see one cool example, and then extrapolate from that
- Many suffer from severe cases of "exponentialism"
- Recent advances are not enough to allow machines to reach human levels of intelligence
- Two dozen or more major breakthroughs could be needed before this happens

Conclusion

Conclusion

- AI has taken the step from the labs into thousands and thousands products and services
- AI has transformed technology and business profoundly and will continue to do so
- Properly used, AI can play a key role in bringing prosperity and sustainability on a global scale
- AI technology must be developed side-by-side with strategies for using AI in a responsible way
- Much more focus is needed on the political, societal, economic, legal, humanistic, and ethical aspects of AI.