Lecture 1: Motivation and AI History

Viliam Lisý & Branislav Bošanský

Artificial Intelligence Center
Department of Computer Science, Faculty of Electrical Eng.
Czech Technical University in Prague

viliam.lisy@fel.cvut.cz

February, 2022

Goals of the course

Understand the problems and methods in the field of AI.
Get broader context of the field

- **Reusable** formal models of decision-making problems
- Basic algorithms solving these problems
- Understand the hypes and limitations
- Understand the relation of AI to ML (CS)
- There is a very long way from a PoC to useful products
- etc.

What is AI?

Britannica: Artificial intelligence (AI), the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings.

Merriam-Webster:

1. a branch of computer science dealing with the simulation of intelligent behavior in computers
2. the capability of a machine to imitate intelligent human behavior

Me: Research field trying to solve problems easy to humans but (so far) difficult for computers - naturally changing a lot in time

Artificial Intelligence is the Oil of 21 Century

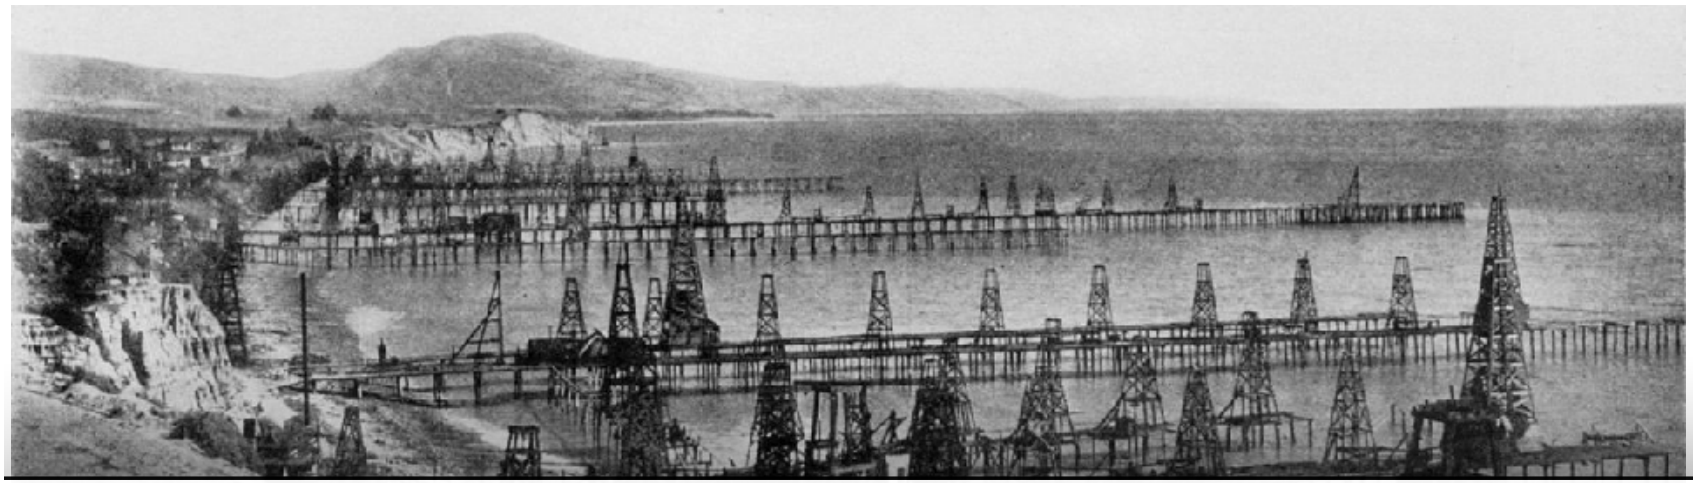

	2016	2017	2020
AI Market size:	**$ 7.8 B**	**$ 12.5 B**	**$ 46 B**

IDC

projections for 2028 are around $700 B

(original slide by Michal Pěchouček)

Business perspective

What do the numbers mean? **Baselines**:

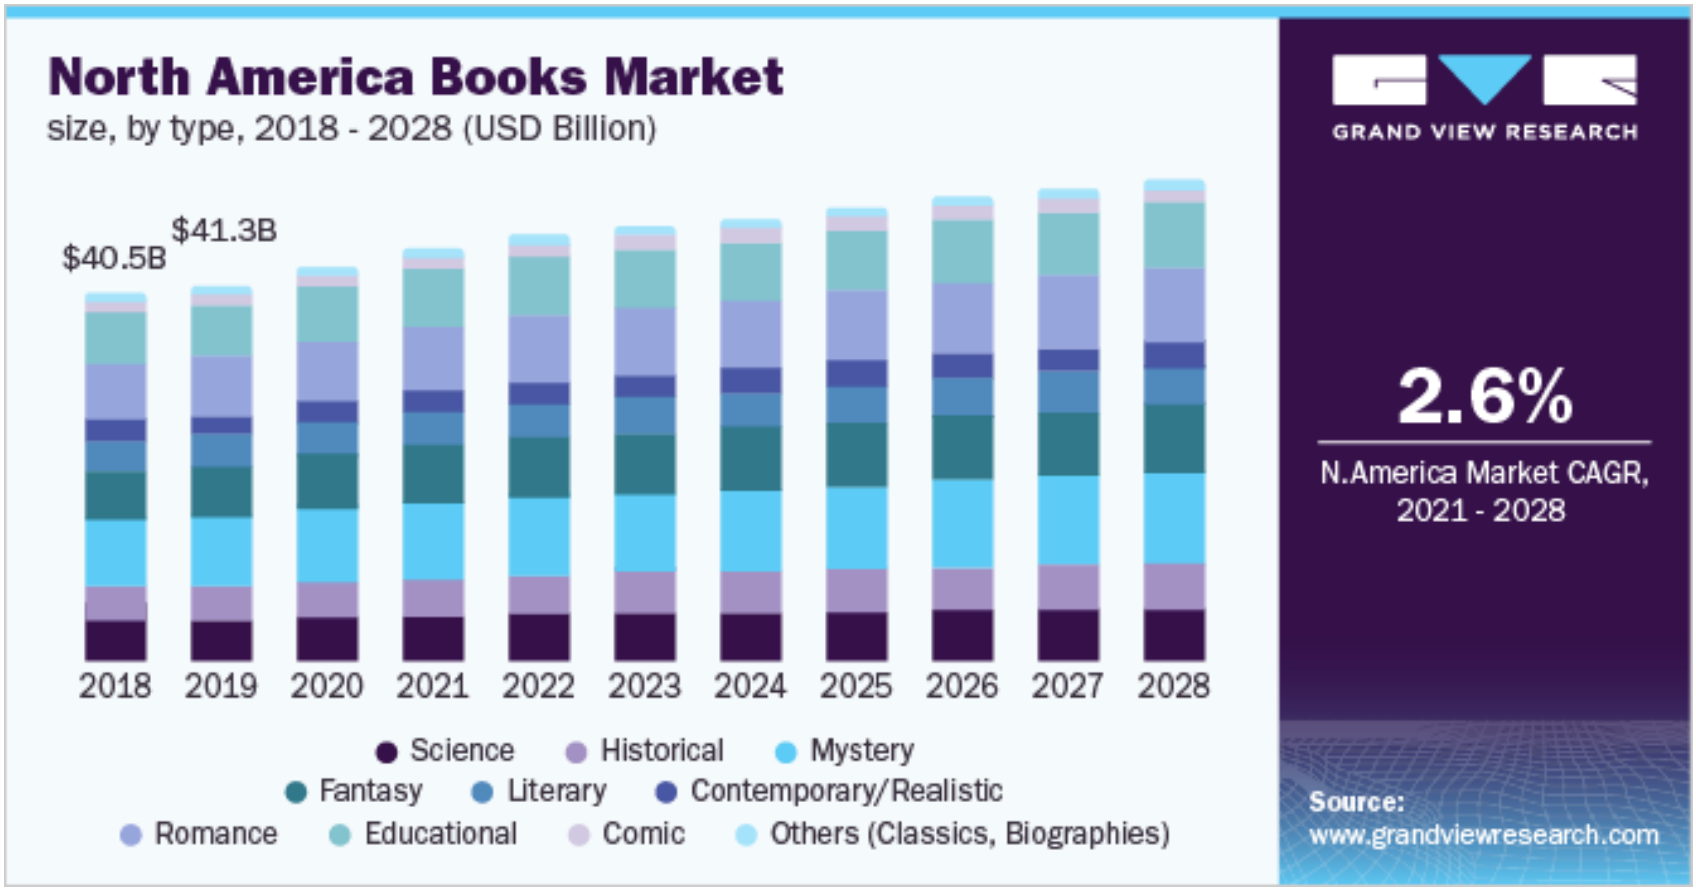

Video games market was ≈90 B, automotive ≈3300 B in 2020.

Sources:
https://techjury.net/blog/gaming-industry-worth/
https://www.ibisworld.com/global/market-size/global-car-automobile-sales/

Artificial Intelligence is the Oil of 21 Century

- **total spent in AI by tech giants in 2016: $20B - $30B** McKinsey&Company
  - » internal R&D **90%** | » acquisitions **10%**
- **total spent in AI by startups in 2016: $6B - $9B**
- **total external investment in AI in 2016: $8B - $12B**
  - » **60%** machine learning
- **AI spent: geographical distribution**
  - » **66%** | » **17%** | » **17%**

(Slide by Michal Pěchouček)

Why is it important to know AI history

1. what is already possible / known and what not
2. when did the well-known results happen
3. hypes vs. actual progress
4. time from flashy prototypes in media to practice
5. motivate seemingly redundant lectures in this course
6. not to be embarrassed by lack of general knowledge
    - Turing test, Deep blue, AI winter, superintelligence, etc.

$-\infty$ – 1940: “Ancient” history

long BC: Golem

1495 Leonardo’s mechanical knight

Armor prforming human-like motrions trhough pulleys and cables

1770 Mechanical turk

Fake chess automaton

1914 First actual chess-playing machine

Endgame of rook+king vs. king

1921 RUR

Origin of the word “robot”

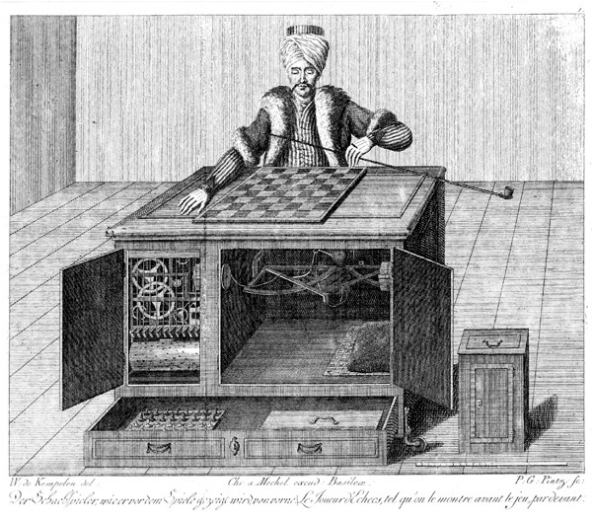

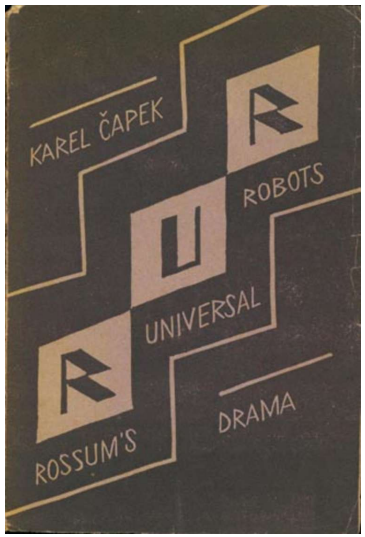

1940s

1941 First electronic computers

Z3, Atanasoff–Berry, Colossus

1943 McCulloch&Pitts create the first model of artificial neurons

1945 Alan Turing states that computers could play chess

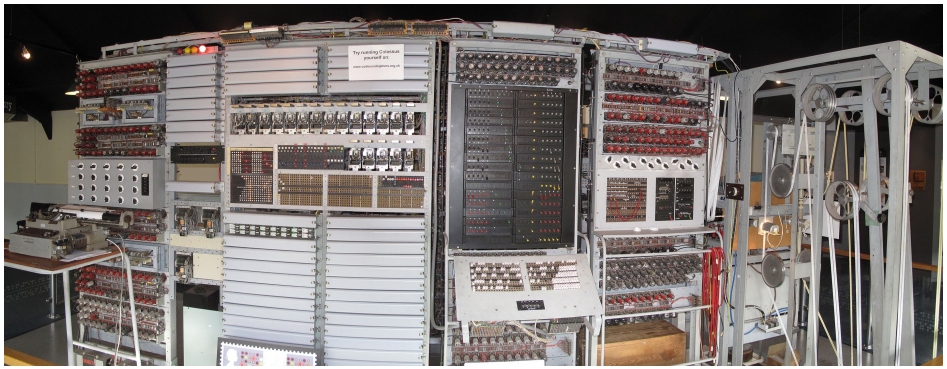

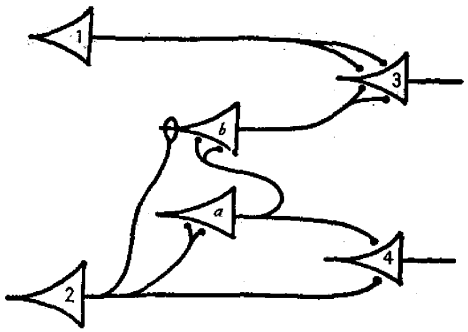

1950s

1950 Turing wrote the first chess program.

1950 Claude Shannon published the fist paper on computer chess.

1950 Turing proposed the **Turing test**.

1956 Logic Theorist: Reasoning as heuristic search

> Eventually proved 38 of the first 52 theorems in chapter 2 of the Principia Mathematica

1956 Dartmouth workshop

1950: Turing test

Formalization of the question: **Can machines think?**

The first variant was proposed in 1948 using Chess. Can a computer be programmed to be indistinguishable from a bad chess player?

Official publication

A. M. Turing (1950): Computing Machinery and Intelligence. MIND a Quaterly Review of Psychology and Philosophy, Vol. 59, No. 236., pp 433–460.

1950: Turing test

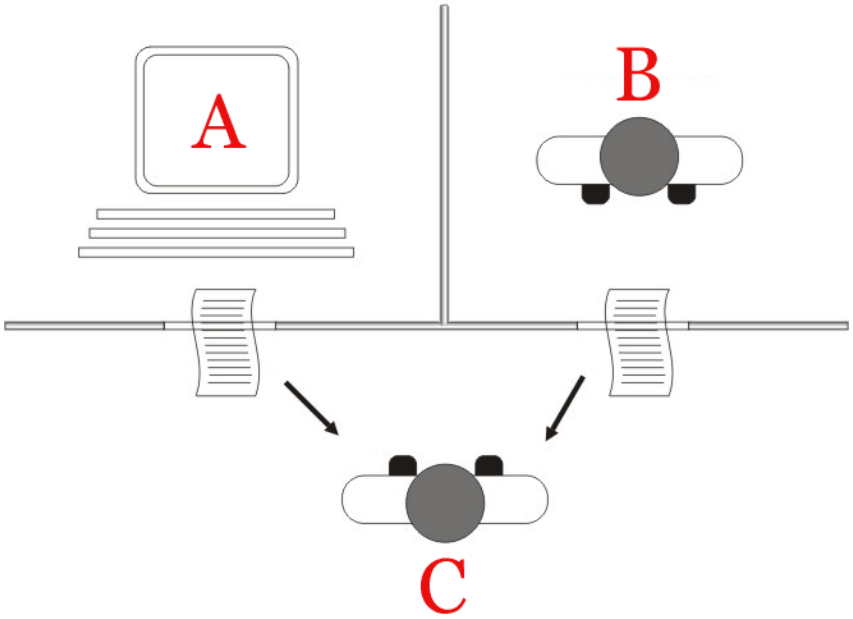

- Separates physical and intelectual capabilities
- Turing proposed variants
  - An interegator distinguishes between a human and a machine
  - She distinguishes whether the subjects is a human or a machine
  - (1952) A whole jury serves as the interrogator
- The paper rebutts some of the main objections

2022: Turing test

2022: Turing test

CAPTCHA: Completely Automatic Public Turing tests to tell Computers and Humans Apart

1956: Dartmouth workshop on "artificial intelligence"

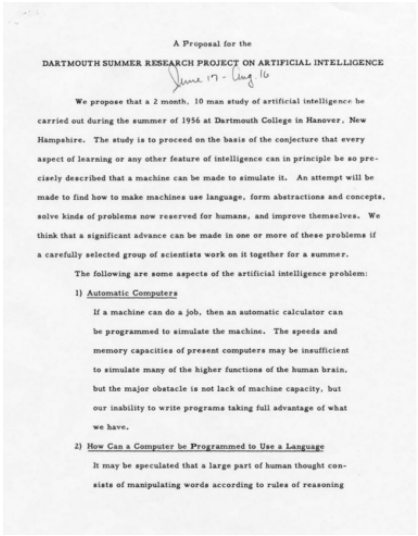

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

The proposal from 1955 used the term "artificial intelligence".

- "We propose that a 2 month, 10 man study of artificial"...
- Main topics:
  - Automatic computers
  - How can a computer be programmed to use a language
  - Neuron nets
  - Theory of the size of a calculation
  - Self-improvement
  - Abstractions
  - Randomness and Creativity

Considered to be the founding event of AI as a research field.

1960s

1961 Unimate, the first industrial robot for General Motors assembly line (video)

1965 I.J. Good writes in "Speculations Concerning the First Ultraintelligent Machine" that "the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control."

1966 Eliza

1966 Shakey the robot

1966: Eliza

Weizenbaum, J. (1966). ELIZA – a computer program for the study of natural language communication between man and machine. Communications of the ACM, 9(1), 36-45.

Shows that a computer can hold a natural language conversation.

Based on pattern matching with minimal permanent state.

- identifies keywords, discovers a minimal contex, chooses a **transformation rule**, generic response if no keywords

Most famous set of conversation rules simulated a psychologist.

```
Men are all alike.
IN WHAT WAY?
They're always bugging us about something or other.
CAN YOU THINK OF A SPECIFIC EXAMPLE?
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE?
He says I'm depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED.
It's true. I am unhappy.
DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY.
I need some help, that much seems certain.
WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP?
Perhaps I could learn to get along with my mother.
TELL ME MORE ABOUT YOUR FAMILY.
My mother takes care of me.
WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU?
My father.
YOUR FATHER?
You are like my father in some ways.
WHAT RESEMBLANCE DO YOU SEE?
You are not very aggressive
but I think you don't want me to notice that.
WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE?
You don't argue with me.
WHY DO YOU THINK I DON'T ARGUE WITH YOU?
You are afraid of me.
DOES IT PLEASE YOU TO BELIEVE I AM AFRAID OF YOU?
My father is afraid of everybody.
WHAT ELSE COMES TO YOUR MIND WHEN YOU THINK OF YOUR FATHER?
Bullies.
DOES THAT HAVE ANYTHING TO DO WITH THE FACT
THAT YOUR BOYFRIEND MADE YOU COME HERE?
```

```
(0 YOU 0 ME) => (WHAT MAKES YOU THINK I 3 YOU)
```

1966-1972: Skakey

General-purpose mobile robot developed at Standord Research Institute

> Combined navigation, computer vision, natural language processing, etc.

Programmed primarily in **Lisp**. General purpose planning using **STRIPS**. The project also introduced the **A*** algorithm and generalized Hough transform.

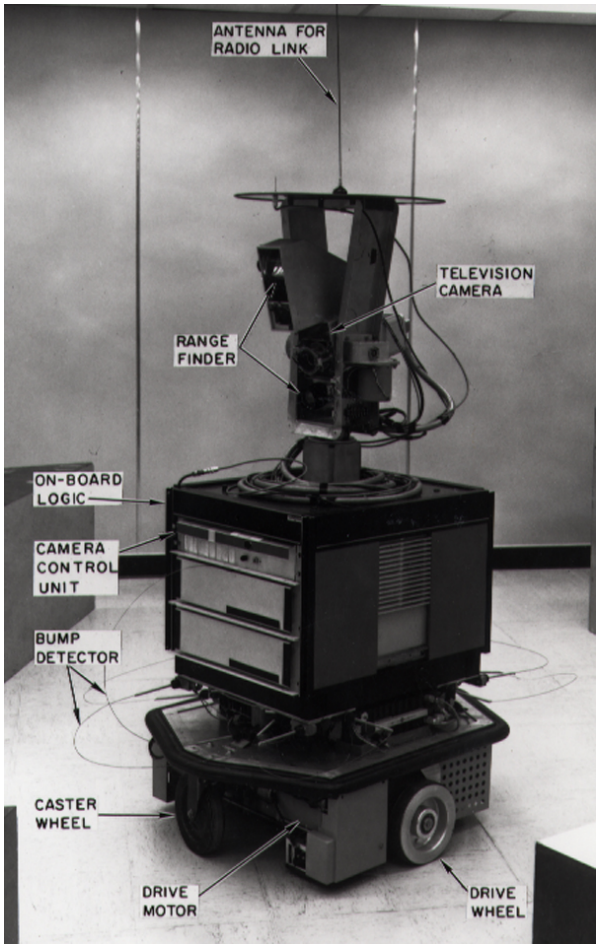

1972 WABOT-1, humanoid made in Japan

Could walk, communicate in Japanese, grip and transport objects

1972 MYCIN, an expert system to identify the source of infection and recommend antibiotics

- Written in Lisp by Edward Shortliffe
- Knowledge base of ∼600 rules
- "Fuzzy logic" evidence aggregation
- acceptability rating of 65% from a panel of specialist (humans achieved 42.5% – 62.5%)

  AI adaptation is not only about performance

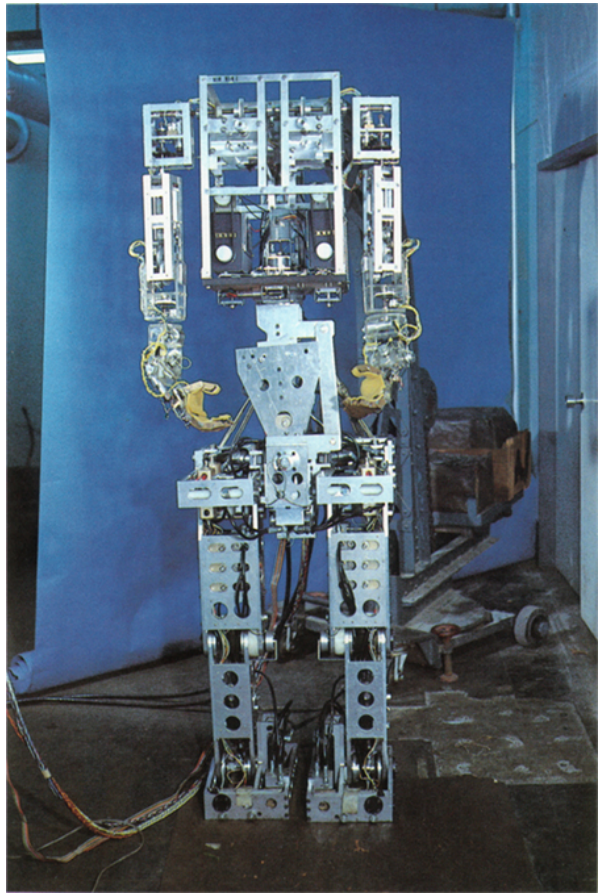

1974-1980: The First AI Winter

Large drop of funding and interest in AI

- Initial hype did not deliver impactful applications
- UK funding cut due to Lighthill report

  > "In no part of the field have the discoveries made so far produced the major impact that was then promised"

- DARPA focussed on "mission-oriented direct research"
- Negative results of Minsky & Papert on perceptrons decreased popularity of connectionism

1980: Chinese room argument

Strong AI: a computer is a mind, which can literally understand
Weak AI: it simulates thoughts and only seem to understand

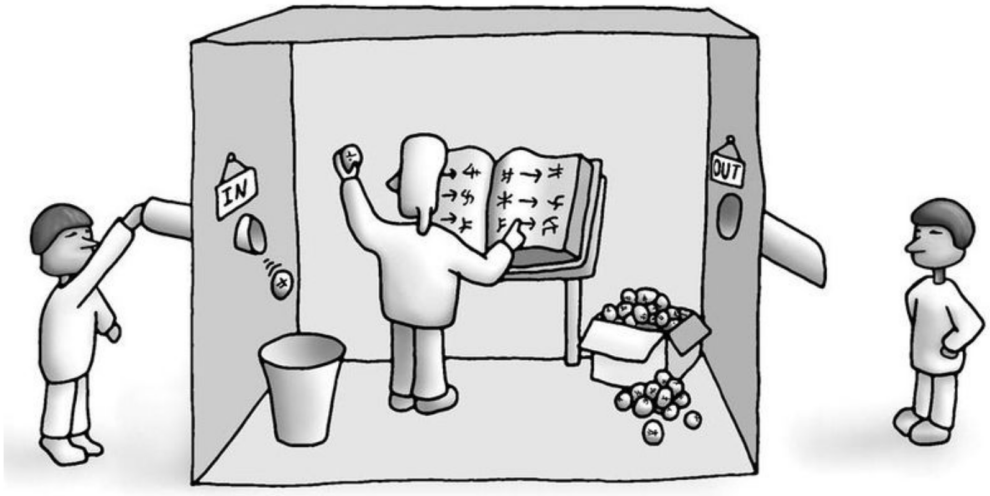

Source: Wikicomms

Searle argues against the possibility of strong AI.

> Searle, John. 1980a. "Minds, Brains, and Programs." Behavioral and Brain Sciences 3, 417-424.

1980s

1980 XCON, expert system selecting computer components

Saved millions of dollars annually in deployment

1980 WABOT-2, robot musician reads score and plays el. organ

1981 $850 million Japanese "Fifth Generation Computer" project

Supercomputer focused on parallel logical programming for AI

1986 Driverless van drives 55 mph on empty streets of Bavaria

Driving in public traffic since 1992, 1000s kms on highways

Few other similar projects around the same time

1988 Judea Pearl invented Bayesian networks

Turing Award in 2011

1989 recognising handwritten ZIP codes

Yann LeCun et al. at AT&T using Neural Neworks

1987-1993: The Second AI Winter

Large drop of funding and interest in AI

- The expert systems did not meet the expectations, they were too expensive to maintain
- Drop of market for specialized AI HW as general purpose HW became more powerful
- High expectations / promises of many startups failed

1990s

1990 Brooks: "Elephants Don't Play Chess"

Fast reactive behaviours for AI rather than hard symbolic manipulation

1995 A.L.I.C.E., chatbot with large data samples from the internet

1997 DeepBlue beats Kasparov

Parallel Alpha-Beta search, evaluation function tuned on grandmaster games, openings database, endgames

1999 AIBO

2000s

2002 i-Robot Roomba, comertially successful robotic vacuum cleaner

2004 DARPA Grand Challenge: No one finished the 150 miles in a desert

2005 DARPA Grand Challenge: Multiple successful teams

Continued with few more challenges

2009 Google develops driverless car (2014 passed Nevada's test)

2010s

2011 Watson

2011 Siri

2015 DQN in Atari games

Reinforcement learning + Deep Learning

2015 AlphaGo

Monte Carlo tree search + Deep Learning

2017 DeepStack

CFR + Deep Learning

2018 BERT followed by other language models

Big Data + Deep Learning

2019 AlphaStar

Game Theory + RL + Deep Learning

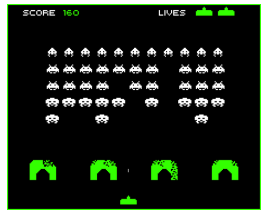

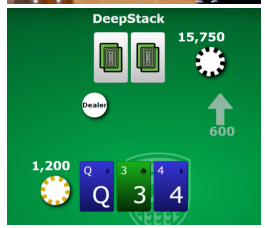

2010s: AI safety and fairness concerns

Open Letter on Artificial Intelligence (2015): Stephen Hawking, Elon Musk, Russell, Norvig, Hassabis, Wozniak et al.

> Argues to focus on maximizing societal benefits of AI, e.g., Law and Liability, Ethics, Autonomous Weapons, Privacy, Verification, Control, etc.

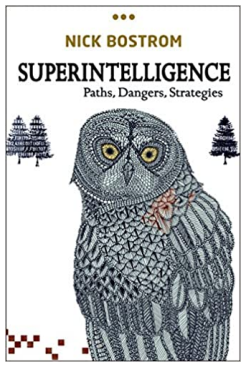

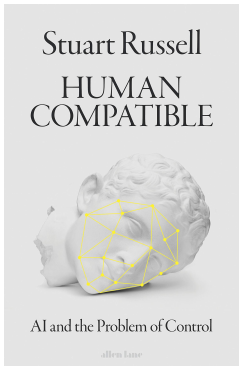

Bias in AI: Amount of care in US hospitals, recidivism prediction, Amazon's hireing, gender stereotypes in language models, etc.

2020: AlphaFold

Predicting protein 3D shape from genetic sequence.

Important for

- understanding biological processes
- creating new drugs

Released publicly in 2021
Long-standing challenge
More details

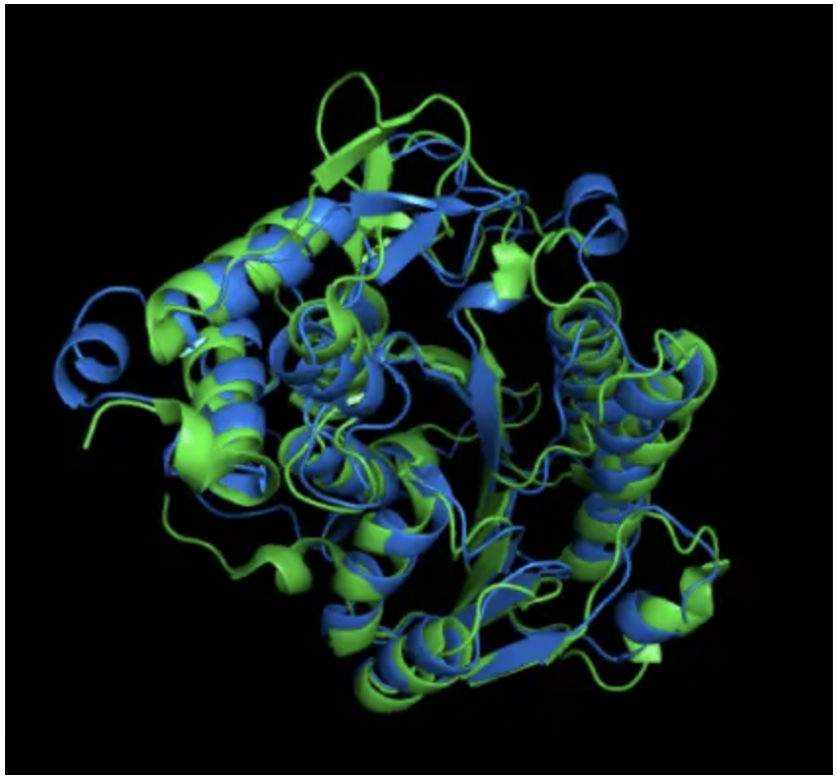

2021: Dall-E

System for generating images from textual description trained from text-image pairs

"an armchair in the shape of an avocado"

"an illustration of a baby daikon radish in a tutu walking a dog"

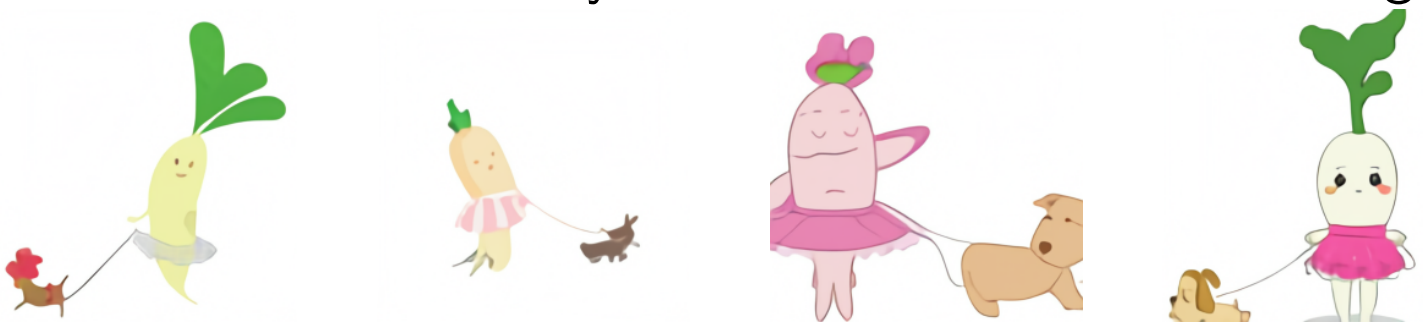

https://openai.com/blog/dall-e/

AI history take-aways

AI history revolves around clear milestones.

Proof of concept is far from a viable product.

The key ideas behind many recent breakthroughs are quite old.

Hypes come and go. Deep ML might be another one.

Deep Learning is great, but usually **not enough**.

AI safety concerns are as old as the field.

Better save than sorry.

A useful tool for AI, which is **not** a focus of this course

Supervised learning = fitting a (high dimensional) function

For a data set $(\vec{x_i}, \vec{y_i})$, find a function f that minimizes:

$$\frac{1}{n}\sum_i \|f(\vec{x_i}) - \vec{y_i}\|.$$

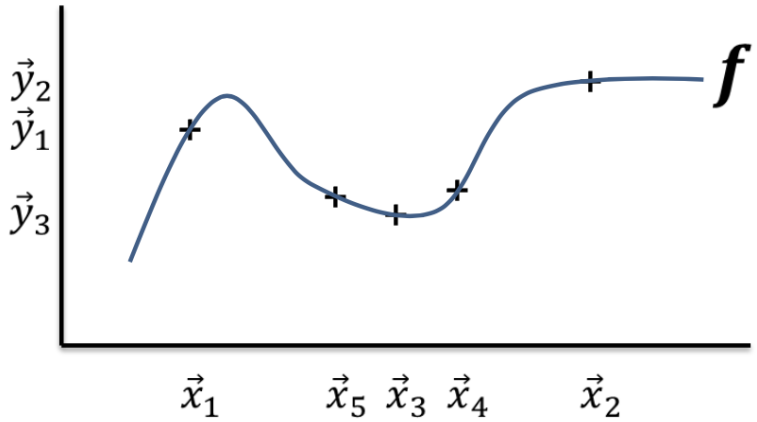

For example, $f($  $) = 2, f($  $) = 3, f($  $) = 4, f($  $) = 5.$

Thinking fast and slow

Humans have two main types of thinking: "thinking fast" and reflexive without much introspection and "thinking slow" based on internal models and predictions about the world.

My analogy:

Machine Learning = thinking fast

Artificial Intelligence $\supset$ thinking slow

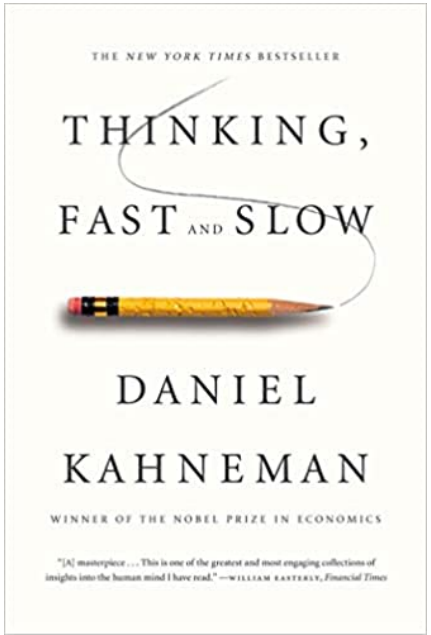

Course organisation

https://cw.fel.cvut.cz/b212/courses/zui/start

14 lectures leading towards some of the mentioned milestones



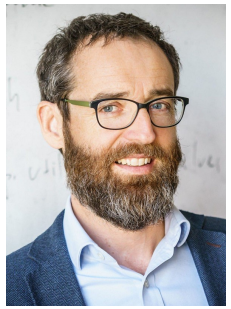

14 labs going deeper to selected algorithmic / theoretic topics


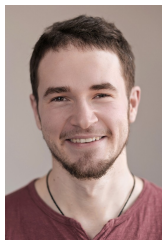

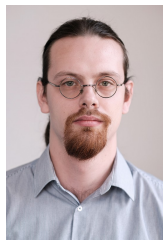
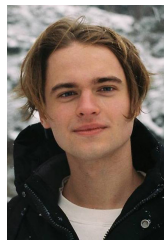
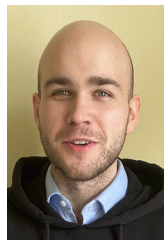

3 programming homeworks in python evaluated by BRUTE

Midterm test

Final written exam

Course grading

https://cw.fel.cvut.cz/b212/courses/zui/start

30% for programming homework

- 10% State space search (A*) algorithm
- 10% Game playing bot
- 10% Reinforcement learning
- Extra tasks for additional point possible

Each task must be submitted for $\geq 50\%$ of its points

Deadline penalties: $\leq$ 24h: -20%; $>$ 24h: 0

Plagiarism will not be tolerated!

If you have serious issues **let us know** ASAP.

15% for the midterm written test

55% for the final written exam

in case of 80+% overall, also a brief oral exam

Standard evaluation scale: https://fel.cvut.cz/education/rules/Study_and_Exam_Code.pdf

Course Topics Overview

- Formal models of AI problems
- Search, A*
- Reinforcement Learning
- Two-Player Perfect-Information Games
- Logical Problem Representations
- Uncertainty in AI
- Sequential Decision Making with Limited Information

Course literature

Slides are not study materials!

1. Take notes.
2. Artificial Intelligence: A Modern Approach (AIMA) by Stuart J. Russell and Peter Norvig (however, it is not free)
3. Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto (PDF available online)
4. Links on the courseware page and in slides
5. Wikipedia