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, 2024

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
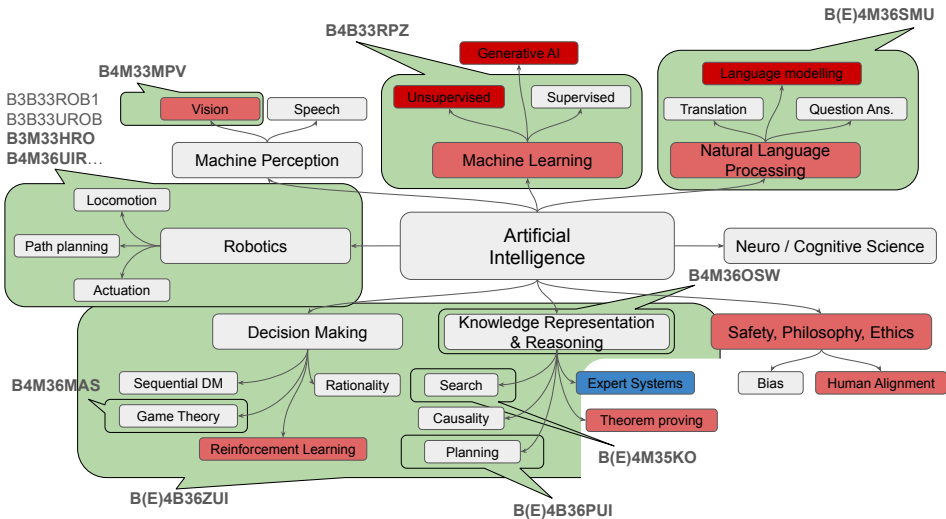- etc.

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

AI Topics Overview

B4B33RPZ

B(E)4M36SMU

B4M33MPV

B3B33ROB1
B3B33UROB
B3M33HRO
B4M36UIR...

B4M36MAS

B(E)4B36ZUI

B(E)4B36PUI

B4M36OSW

B(E)4M35KO

Generative AI

Unsupervised | Supervised

Vision | Speech

Machine Learning

Machine Perception

Language modelling

Translation | Question Ans.

Natural Language Processing

Locomotion

Path planning | Robotics | Actuation

Artificial Intelligence

Neuro / Cognitive Science

Decision Making

Knowledge Representation & Reasoning

Safety, Philosophy, Ethics

Sequential DM | Rationality | Search | Expert Systems

Game Theory | Causality | Theorem proving

Reinforcement Learning | Planning

Bias | Human Alignment

Artificial Intelligence is the Oil of 21 Century

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

projections for 2028 are around $700 B
(original slide by Michal Pěchouček)

Business perspective

ARTIFICIAL INTELLIGENCE (AI) MARKET SIZE, 2022 TO 2032 (USD BILLION)

Source: www.precedenceresearch.com

Business perspective

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

North America Books Market
size, by type, 2018 - 2028 (USD Billion)

$40.5B $41.3B

2018 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028

● Science   ● Historical   ● Mystery
● Fantasy   ● Literary   ● Contemporary/Realistic
● Romance   ● Educational   ● Comic   ● Others (Classics, Biographies)

GRAND VIEW RESEARCH

2.6%
N.America Market CAGR,
2021 - 2028

Source:
www.grandviewresearch.com

World video games market was ≈184.4 B, automotive ≈2950 B in 2022.

My Perspective

If I were to choose a field of study now, it would be AI!

- It is currently one of the most transformative technology.
- AI is replacing jobs. Understanding it makes you the last.
- Very easy to find a job (anywhere).
- Interesting high-tech startup opportunities.

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. not to be embarrassed by lack of general knowledge
   - Turing test, Deep blue, AI winter, superintelligence, etc.

long BC: Golem

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, AtanasoffBerry, Colossus

1943 McCulloch&Pitts create the first model of artificial neurons

1945 Alan Turing states that computers could play chess

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

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.

- 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

CAPTCHA: Completely Automatic Public Turing tests to tell Computers and Humans Apart
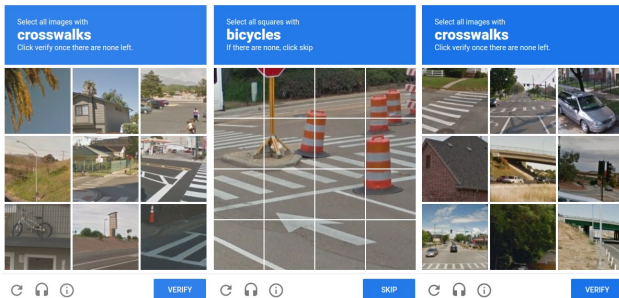
1956: Dartmouth workshop on "artificial intelligence"

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

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

WINTER IS COMING

1990s

1990 Brooks: "Elephants Dont 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 - Robotics

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 - Deep (Reinforcement) Learning

2011 Watson

2011 Siri

2015 DQN in Atari games

Reinforcement learning + Deep Learning

2015 AlphaGo

Monte Carlo tree search + RL + Deep Learning

2017 DeepStack

CFR + RL + 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/

"Large Laguage Model" fine-tuned using "Reinforcement Learning from Human Feedback"

Released on Nov 30,
reached 1M users on Dec 4.

Alternatives exist
(e.g., Claude, Bard/Gemini)

https://openai.com/blog/chatgpt/

Generating realistic high resolution videos from text description

Released on Feb 15, 2024

Does it need a physics simulator?

https://openai.com/sora

AI history take-aways

AI history revolves around clear milestones.

Proof of concept is far from a viable product.

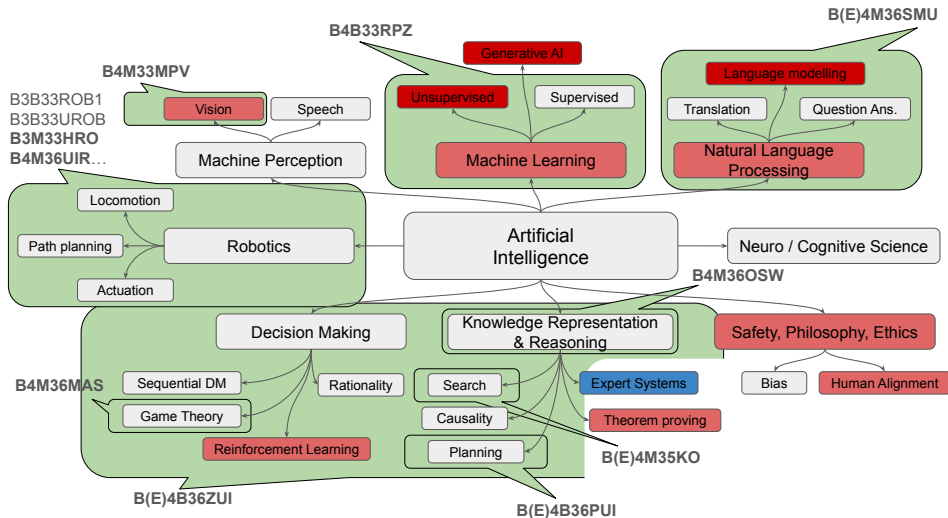The key ideas behind many recent breakthroughs are quite old.

Hypes come and go. Generative AI 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.

AI Topics Overview diagram centered on **Artificial Intelligence**, with connected topics:

- **Machine Perception** — Vision, Speech (B4M33MPV)
- **Machine Learning** — Unsupervised, Supervised, Generative AI (B4B33RPZ)
- **Natural Language Processing** — Translation, Question Ans., Language modelling (B(E)4M36SMU)
- **Robotics** — Locomotion, Path planning, Actuation (B3B33ROB1, B3B33UROB, B3M33HRO, B4M36UIR…)
- **Neuro / Cognitive Science**
- **Decision Making** — Sequential DM, Game Theory, Rationality, Reinforcement Learning (B4M36MAS, B(E)4B36ZUI)
- **Knowledge Representation & Reasoning** — Search, Causality, Planning, Expert Systems, Theorem proving (B(E)4B36PUI, B(E)4M35KO)
- **Safety, Philosophy, Ethics** — Bias, Human Alignment (B4M36OSW)

Course organisation

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

13 lectures leading towards some of the mentioned milestones

13 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/wiki/courses/zui/start

30% for programming homework

- 10% State space search (A*) algorithm
- 10% Reinforcement learning
- 10% Game playing bot
- 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/en/education/rules/Study_and_Exam_Code.pdf

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

Machine Learning

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, $f(\text{2}) = 2, f(\text{3}) = 3, f(\text{4}) = 4, f(\text{5}) = 5$.

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