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, 2025

Lecturers

Branislav Bošanský
Associate professor at CTU
Research manager at Gen (Avast)
Postdoc from Aarhus University, Denmark
50+ publications in top AI venues

Viliam Lisý
Associate professor at CTU
Ex AI architect at Gen
Postdoc from University of Alberta, Canada
Co-author of DeepStack poker AI
50+ publications in top AI venues

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

Goals of the course

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

- Understand the scope of AI vs. ML vs. CS
- Understand the hypes and limitations, PoC vs. product
- **Reusable** formal models of selected AI problems
- Basic algorithms solving these problems

Structure of today's lecture

- What are the main topics AI deals with
- How did AI start
- Deeper dive to some of AI topics
  - robotics
  - decision making
  - natural language processing
- Where can you learn more about AI at CTU
- Course organisation

AI Topics Overview

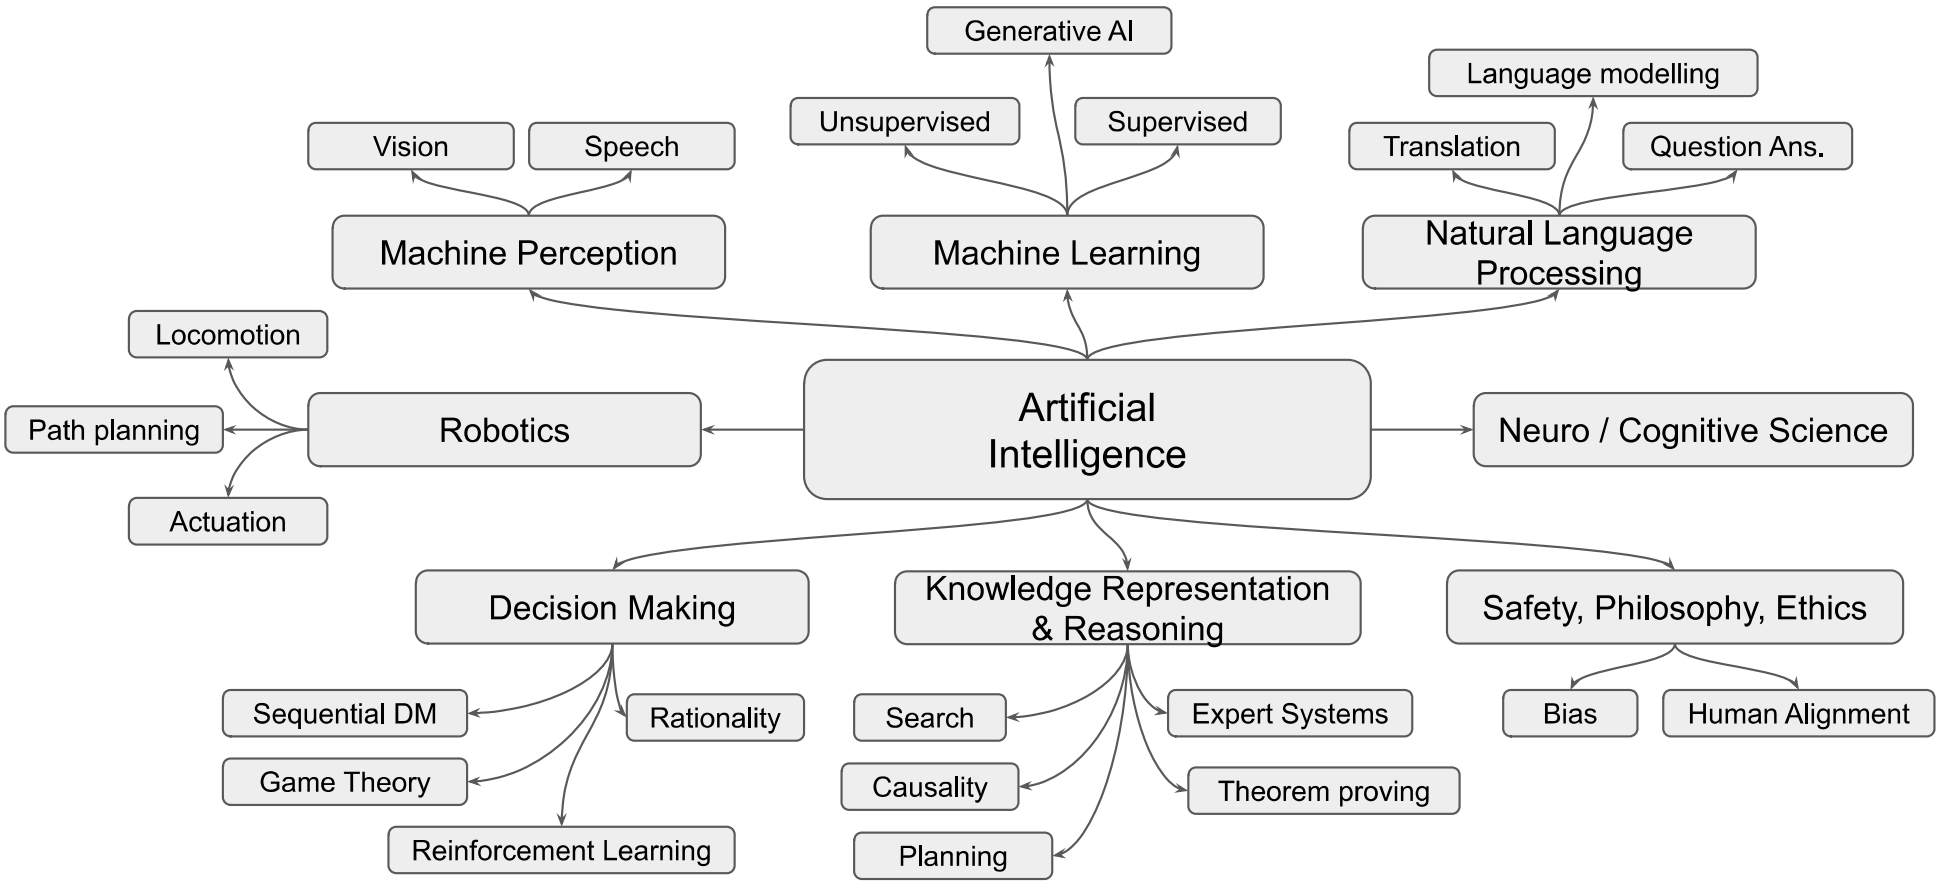

Artificial Intelligence is the Oil of 21 Century

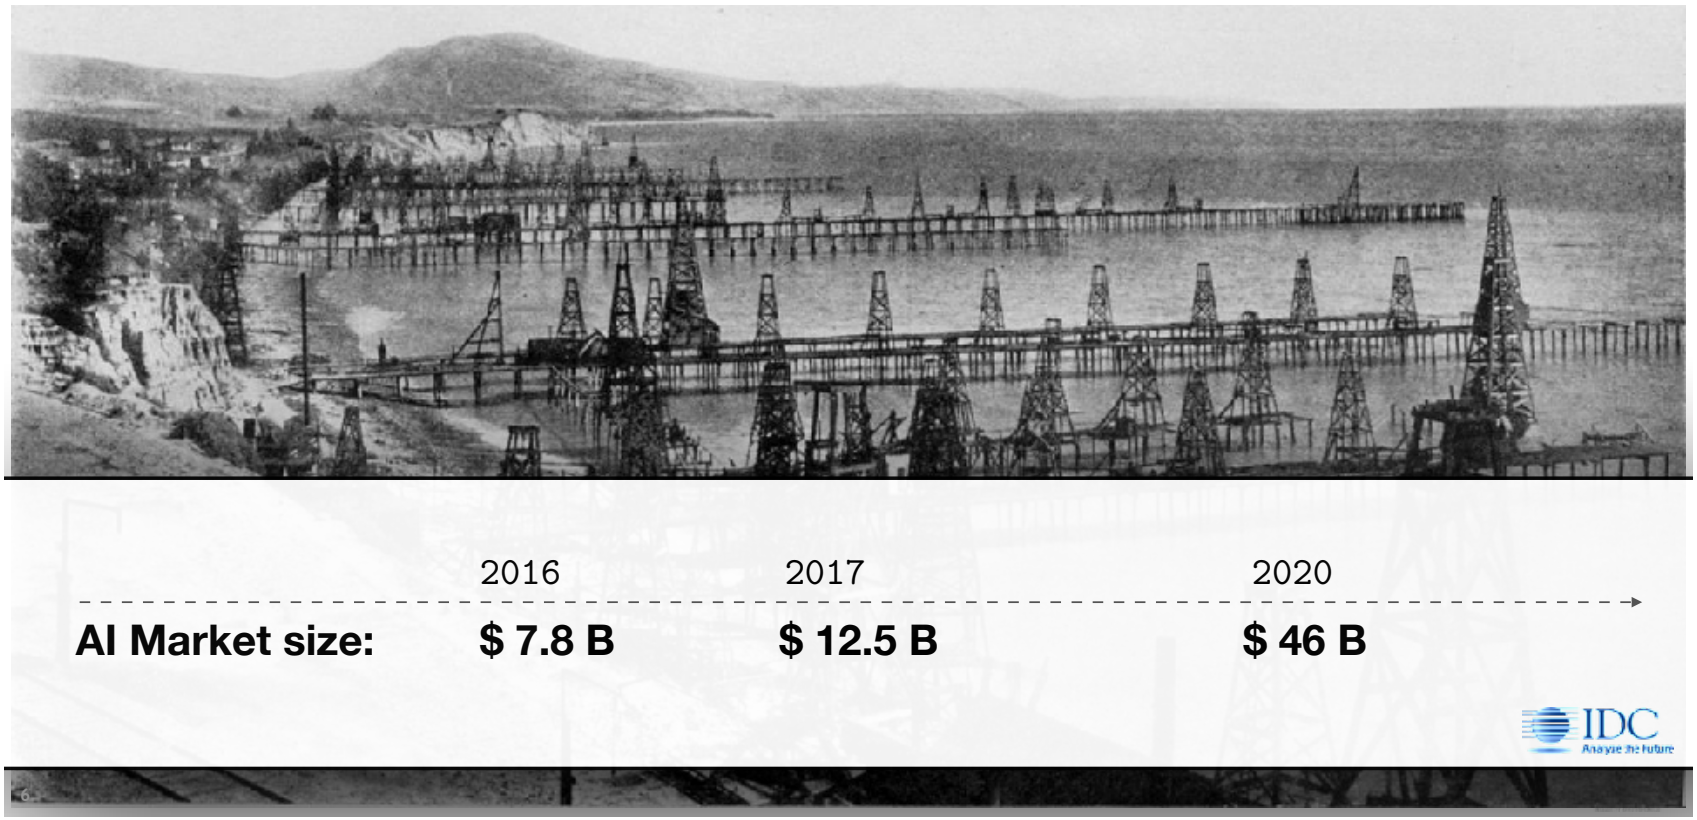

projections for 2028 are around $700 B

(original slide by Michal Pěchouček)

Business perspective

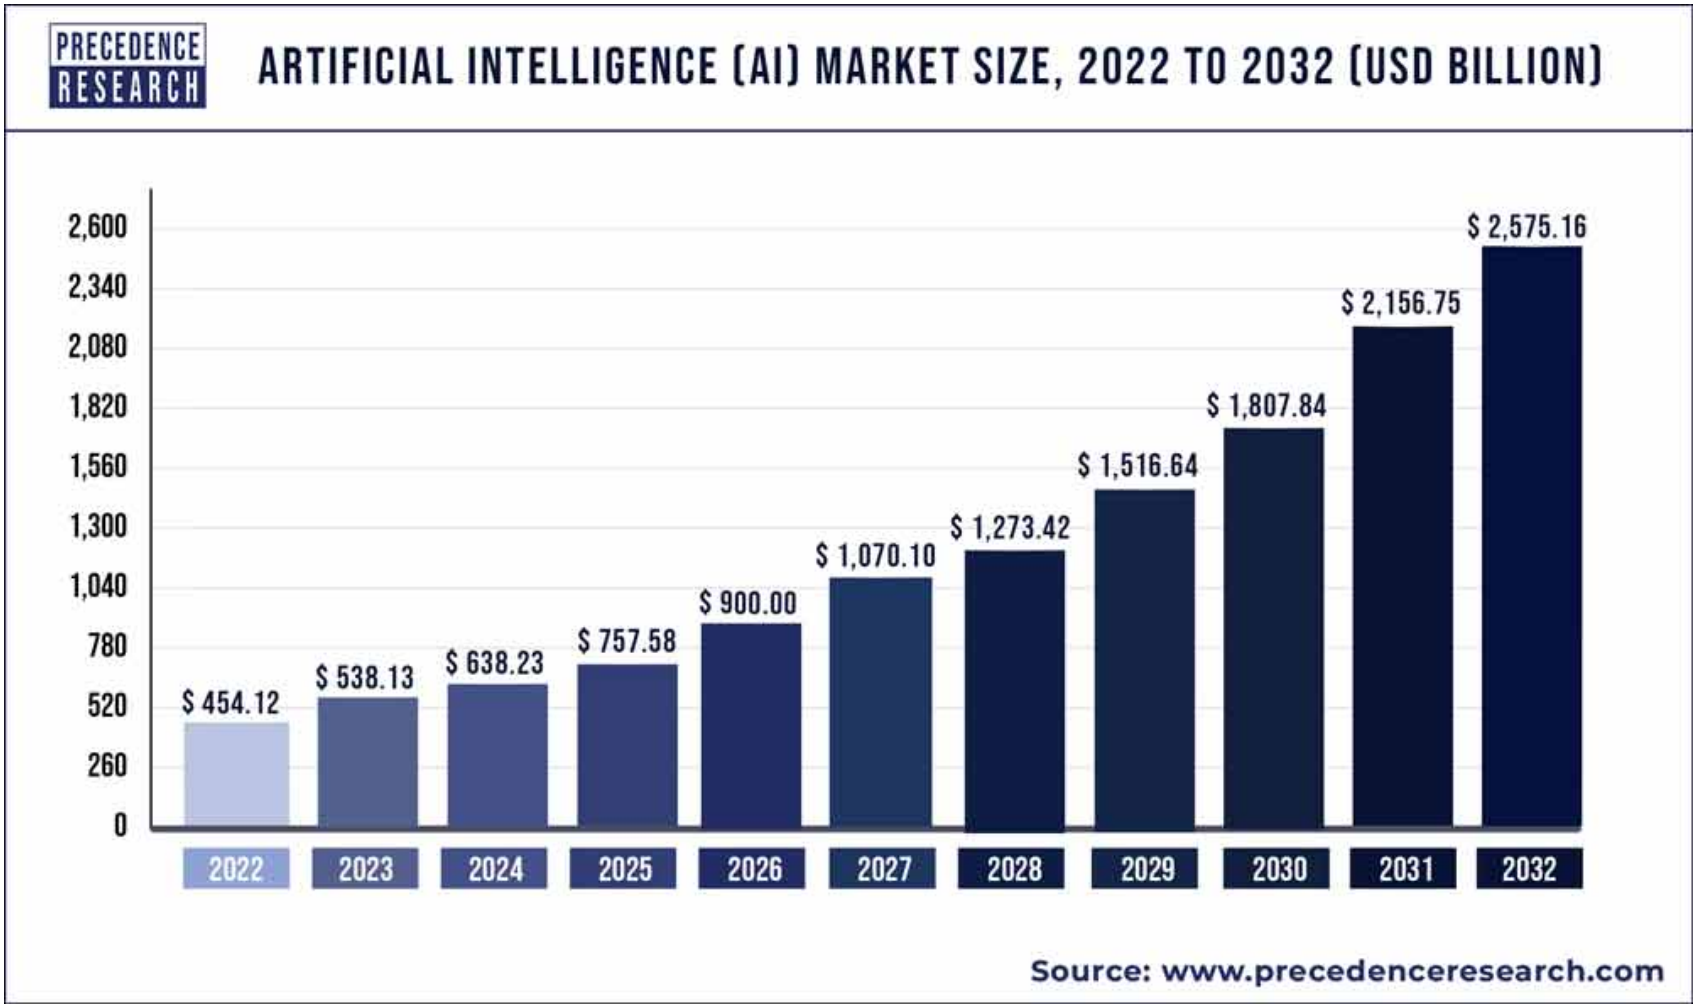

Baselines: books 152, games 188, passenger cars 1728 (in 2024).

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.
- 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 is not
2. when did the well-known results happen
3. hype cycles and 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.

$-\infty$ – 1940: “Ancient” history

long BC: Golem

1914 First actual chess-playing machine

Endgame of rook+king vs. king

1921 RUR

Origin of the word “robot”

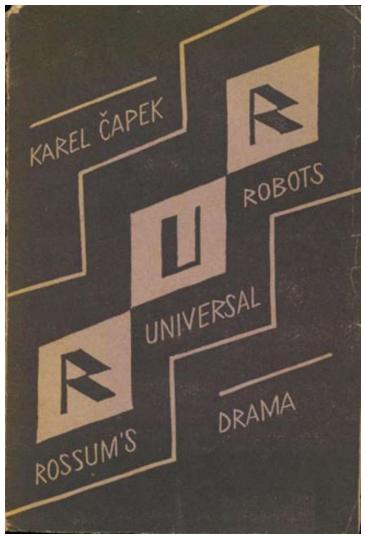

1941 First electronic computers

Z3, Atanasoff–Berry, Colossus

1943 McCulloch&Pitts create the first model of artificial neurons

1945 Alan Turing states that computers could play chess

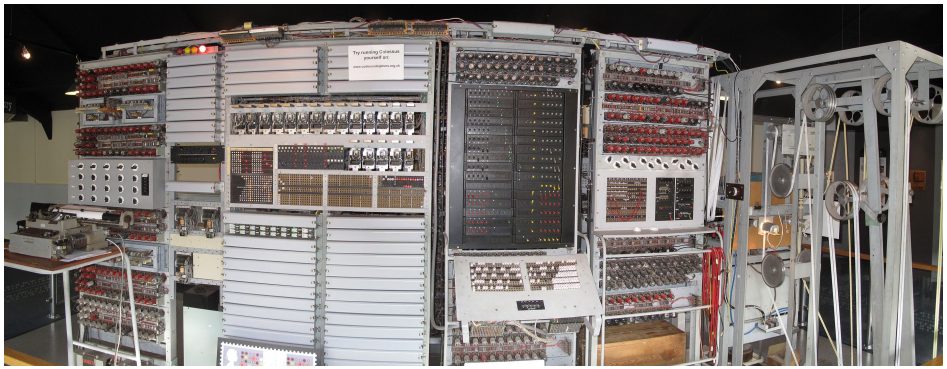

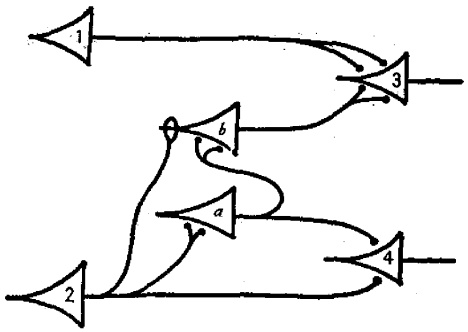

1950: Turing test

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

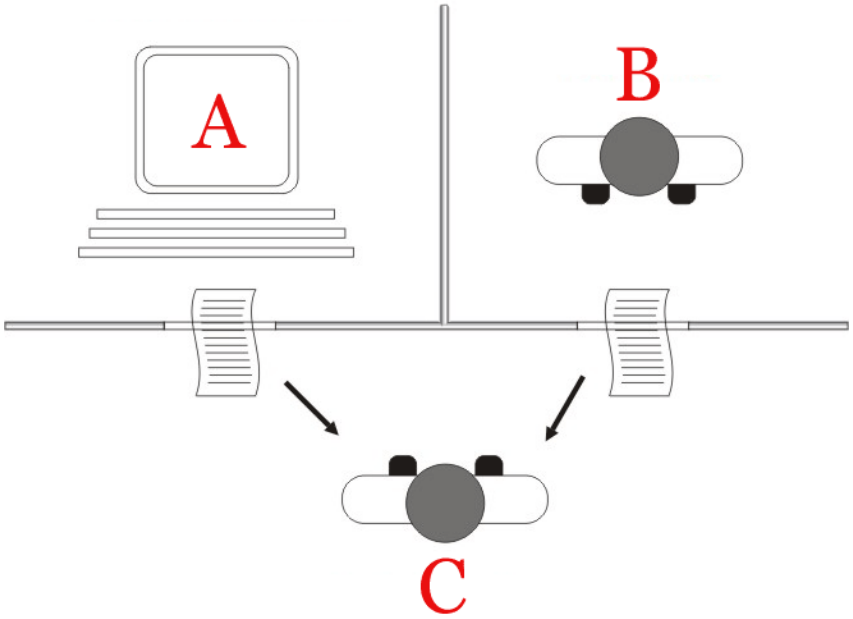

- Separates physical and intelectual capabilities
- Turing proposed variants
  - An interegator distinguishes between a human and a machine
  - She distinguishes whether the subjects is a human or a machine
  - (1952) A whole jury serves as the interrogator
- The paper rebutts some of the main objections

2024: Turing test

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

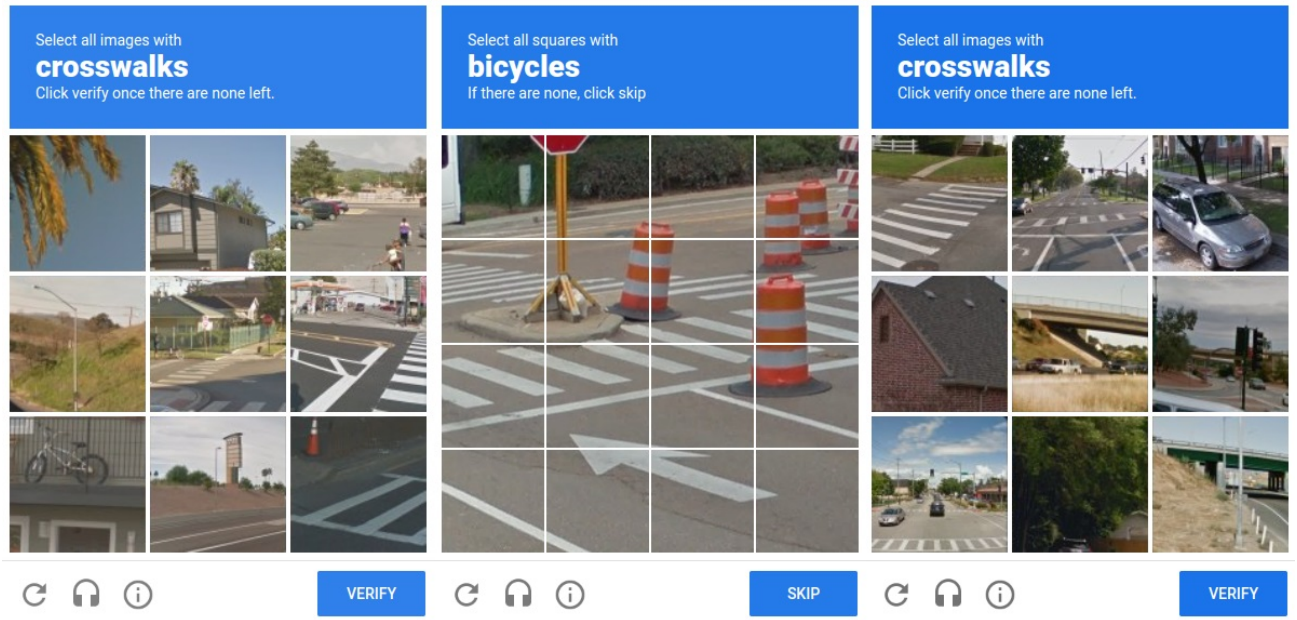

1956: Dartmouth workshop on "artificial intelligence"

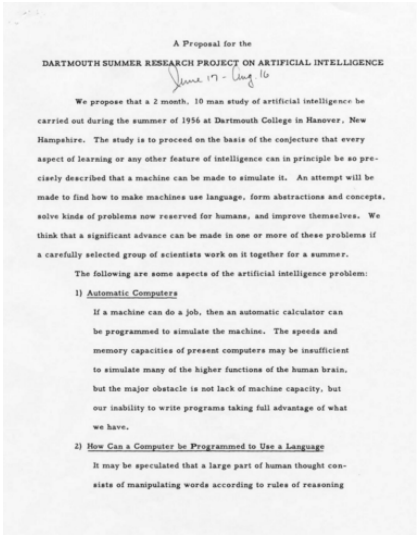

A Proposal for the

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.

Robotics

Mechanical eng. + electrical eng. + sw eng. + control eng. + **AI**

Creating intelligent behavior of machines in the **physical space**

- Sensing the environment: bumpers, odometers, LIDARs, cameras, object recognition
- Thinking: mapping, path/movement planning, task planning
- Acting: movement, grasping, speaking

Are we missing anything? How do we know?

Robotics

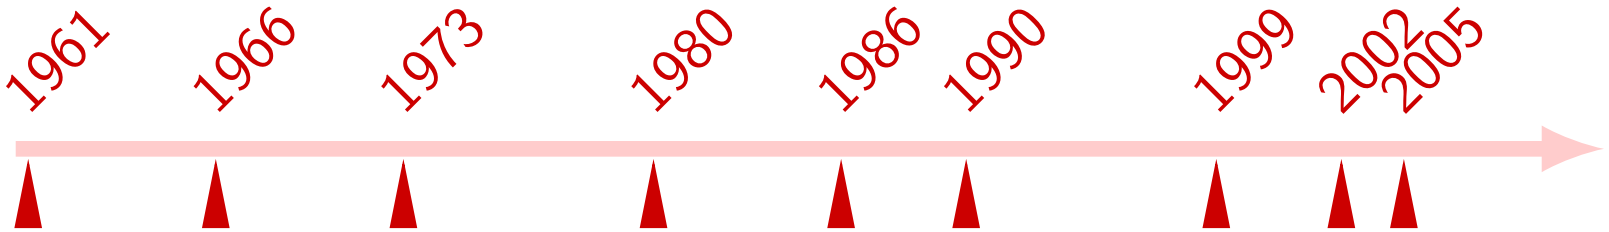

- 1961: Unimate – the first industrial robot (video)
- 1966-1972: Shakey – universal platform

General-purpose mobile robot developed at Stanford 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.

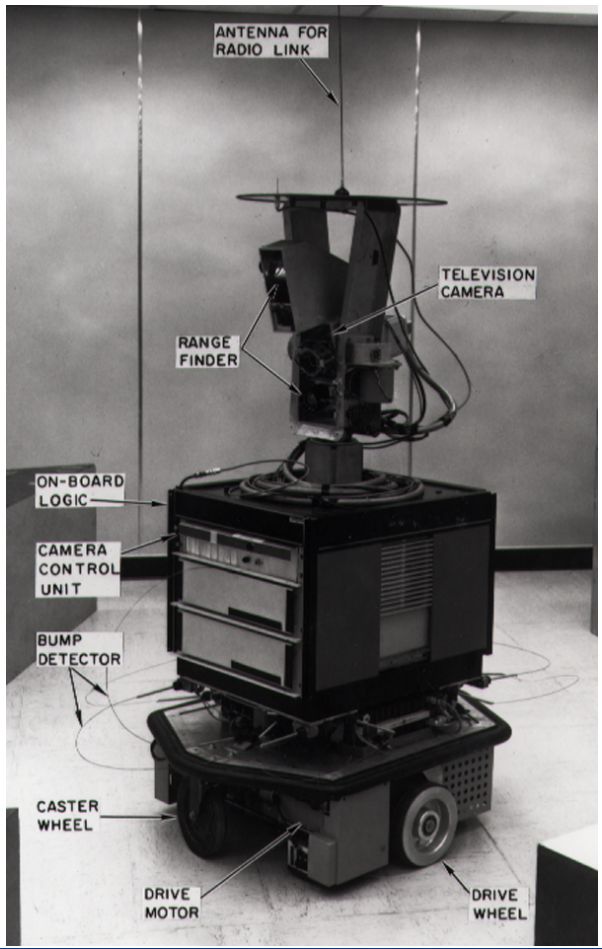

Robotics

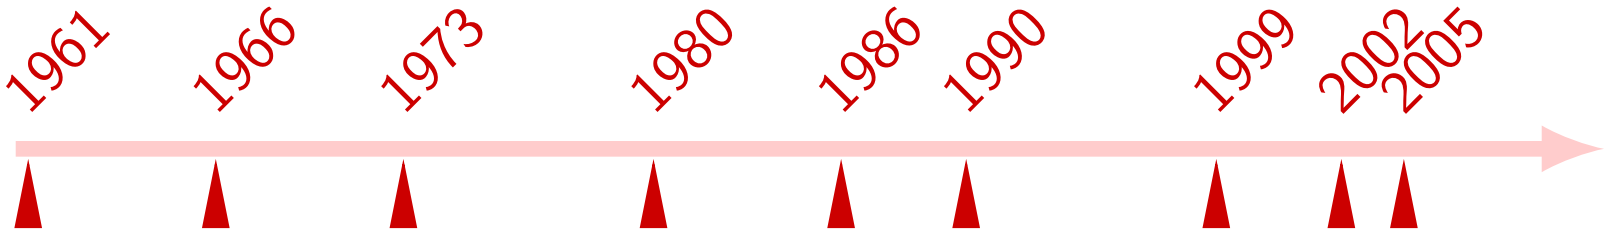

- 1961: Unimate – the first industrial robot (video)
- 1966-1972: Shakey – universal platform
- 1973: WABOT I – first humanoid
- 1980: WABOT II – musician

read music score and played el. organ

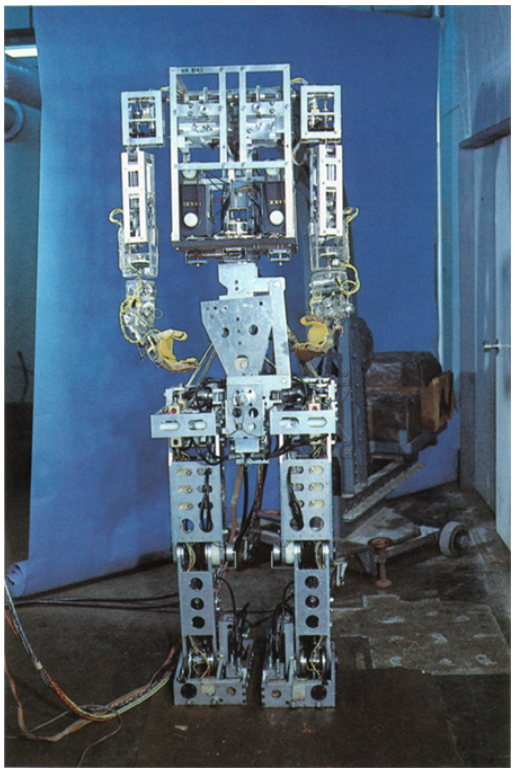

Robotics

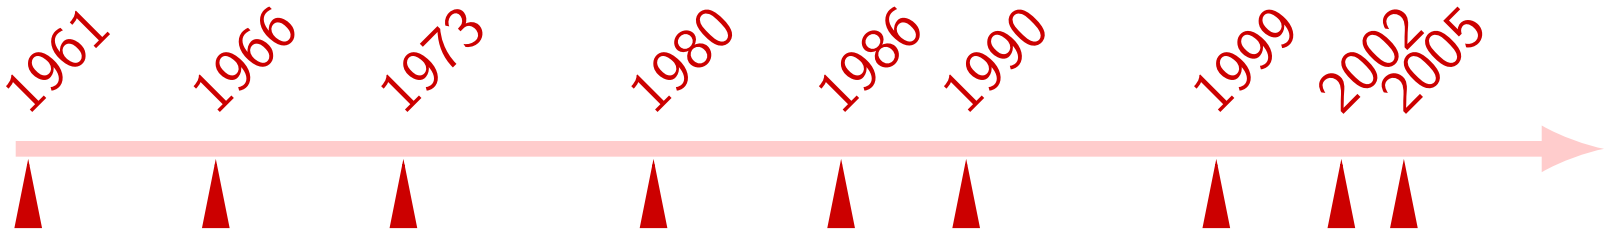

- 1961: Unimate – the first industrial robot (video)
- 1966-1972: Shakey – universal platform
- 1973: WABOT I – first humanoid
- 1980: WABOT II – musician
- 1986: VaMoRs – Driverless van
  - Drives 55 mph on empty streets of Bavaria
  - Driving in public traffic since 1992, 1000s kms on highways
  - Few other similar projects around the same time

Robotics

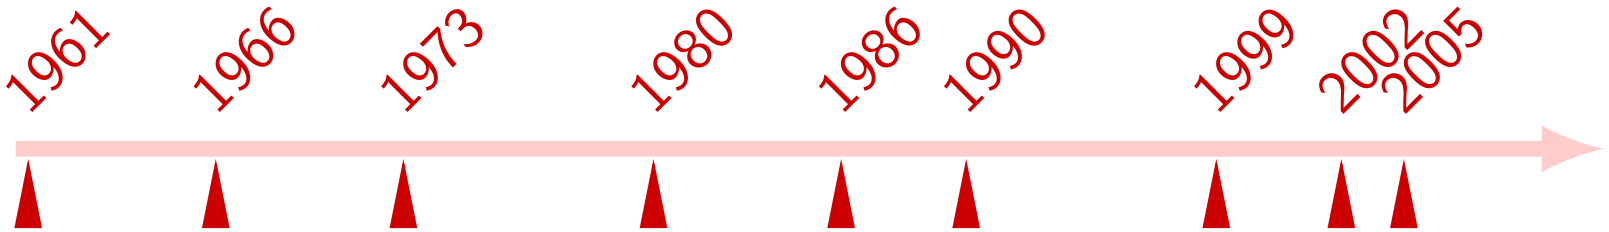

- 1961: Unimate – the first industrial robot (video)
- 1966-1972: Shakey – universal platform
- 1973: WABOT I – first humanoid
- 1980: WABOT II – musician
- 1986: VaMoRs – Driverless van
- 1990: Brooks – "Elephants Don't Play Chess"
- 1999: AIBO
- 2002: Roomba

Robotics

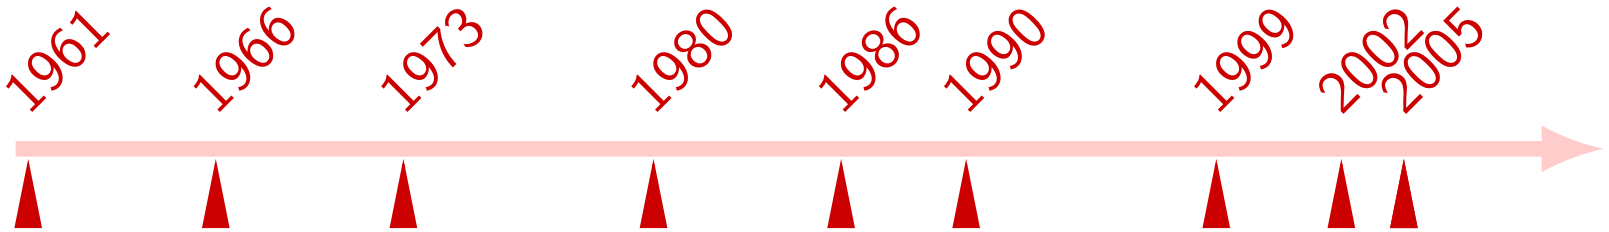

- 1961: Unimate – the first industrial robot (video)
- 1966-1972: Shakey – universal platform
- 1973: WABOT I – first humanoid
- 1980: WABOT II – musician
- 1986: VaMoRs – Driverless van
- 1990: Brooks – "Elephants Don't Play Chess"
- 1999: AIBO
- 2002: Roomba
- 2005: DARPA grand challenge
- 2005: BigDog

Robotics now

Humanoids

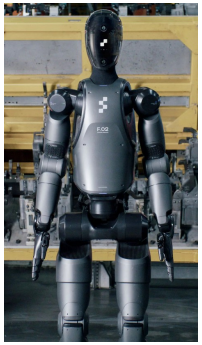

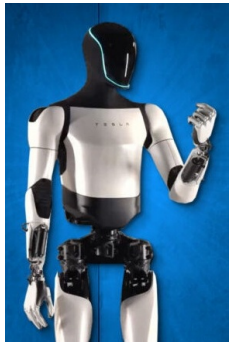

Self-driving vehicles

Drones and delivery

Robotics – Lessons learned

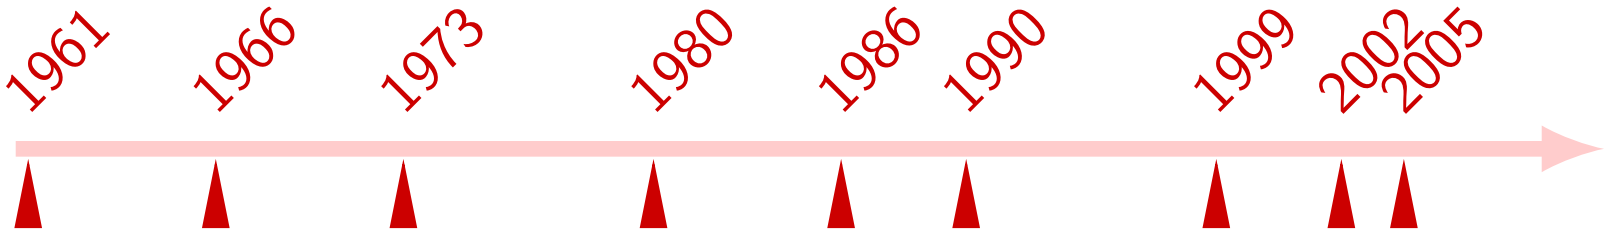

- progress in robotics is increadibly slow
- best way to understand complex systems is to build them
- flashy media results are not very representative
- competitions are a good driver of actual progress in AI
- even the most fancy robots need to:
  - plan higher level movements (A*) or sequence of actions (STRIPS)
  - learn control policies from experience (RL)
  - fuse data from sensors (BN)
  - be aligned with what poople want them to do

AI Decision making

Choosing the right action based on available information

- next move in a game based on rules
- right diagnosis based on symptomps
- next task to achieve a complex goal

Decision making I

1950 Turing wrote the first chess program.

1950 Claude Shannon published the fist paper on computer chess.

1956 Logic Theorist: Reasoning as heuristic search

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

1966 Shakey the robot (A*, STRIPS)

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

Decision making II

1980 XCON, **expert system** selecting computer components

Saved millions of dollars annually in deployment

1981 $850 million Japanese "Fifth Generation Computer" project

Supercomputer focused on parallel logical programming for AI

1988 Judea Pearl invented Bayesian networks

Turing Award in 2011

1987 - 1993: The Second AI Winter

1997 DeepBlue beats Kasparov

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

Decision making III – Deep Reinforcement Learning

2011 Watson

2015 DQN in Atari games

Reinforcement learning + Deep Learning

2015 AlphaGo

Monte Carlo tree search + RL + Deep Learning

2017 DeepStack (Poker)

CFR + RL + Deep Learning

2019 AlphaStar (StarCraft)

Game Theory + RL + Deep Learning

2019 Sutton: "The bitter lesson"

General methods + computation

$>$ human domain knowledge and heuristic

2022 Cicero (Diplomacy)

Tree search + RL + Deep Learning + LLMs

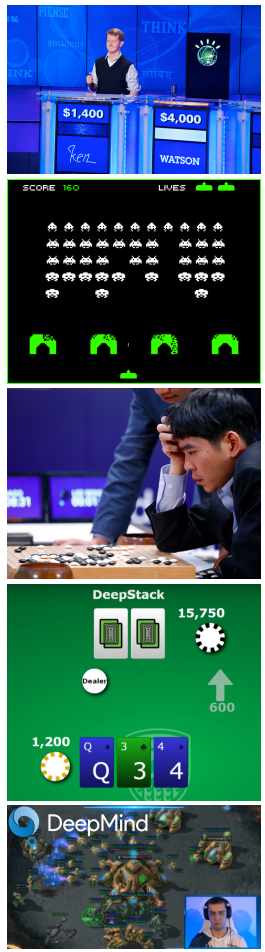

Decision making – Lessons learned

- there are cycles of hipe and slow down
- games (competitions) are a fundamental driver of progress
- deep learning is crucial for modern AI decision making, but it is not enough
- general methods and compute outperfrom human domain knowledge
- raw performance is not enough, there are trust and legal issues

Natural Language Processing

Understanding, interpreting, and generating human language.

- holding a conversation
- automated translation
- automated summarisation
- question answering
- sentiment analysis

Natural language processing – 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)
```

Natural language processing

1966 Eliza – simple rewriting rules

1995 A.L.I.C.E. – chatbot with complex pre-programmed rules

2006 Google translate – statistical NLP

2011 Siri – still mostly rules

2013 Word2vec (Mikolov at al. @ Google)

2018 BER**T** followed by other language models

2022 ChatGPT – massive pre-training and RLHF fine-tunning

2024 "Reasoning" models o1, o3, r1, s1, chess, AlphaGeometry

essentially adding search to LLM inference

DeepSeek "reasoning" trained by RL

NLP – Lessons learned

- hand-conding knowledge did not make it
- predicting the next token (word) and huge data did
- RL is often used for LLM training
- Recent models use "reasoning" related to search

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:

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

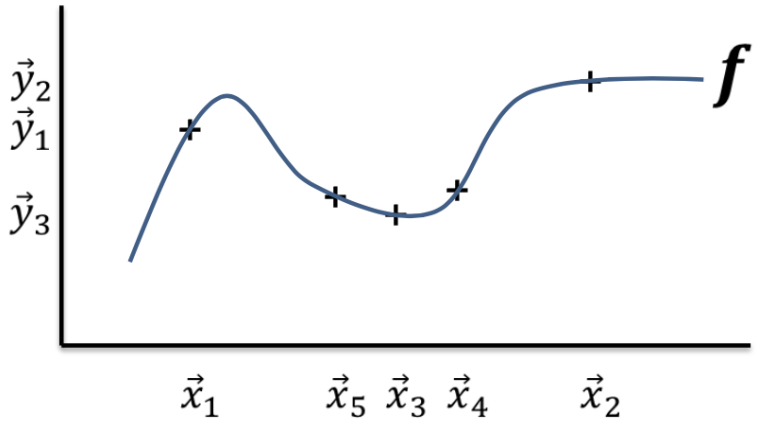

For example, $f($  $) = 2, f($  $) = 3, f($  $) = 4, f($  $) = 5.$

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 ⊃ thinking slow

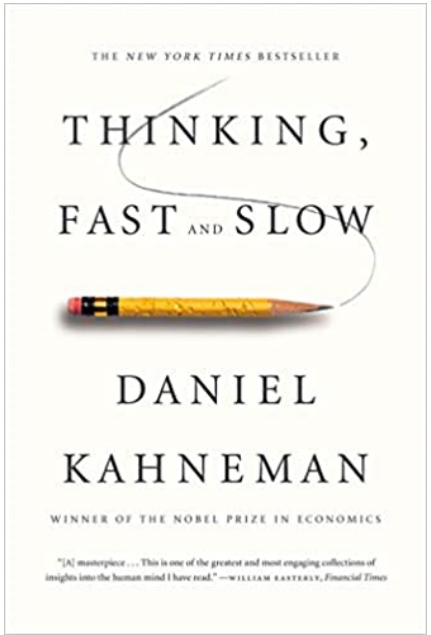

AI Topics Overview

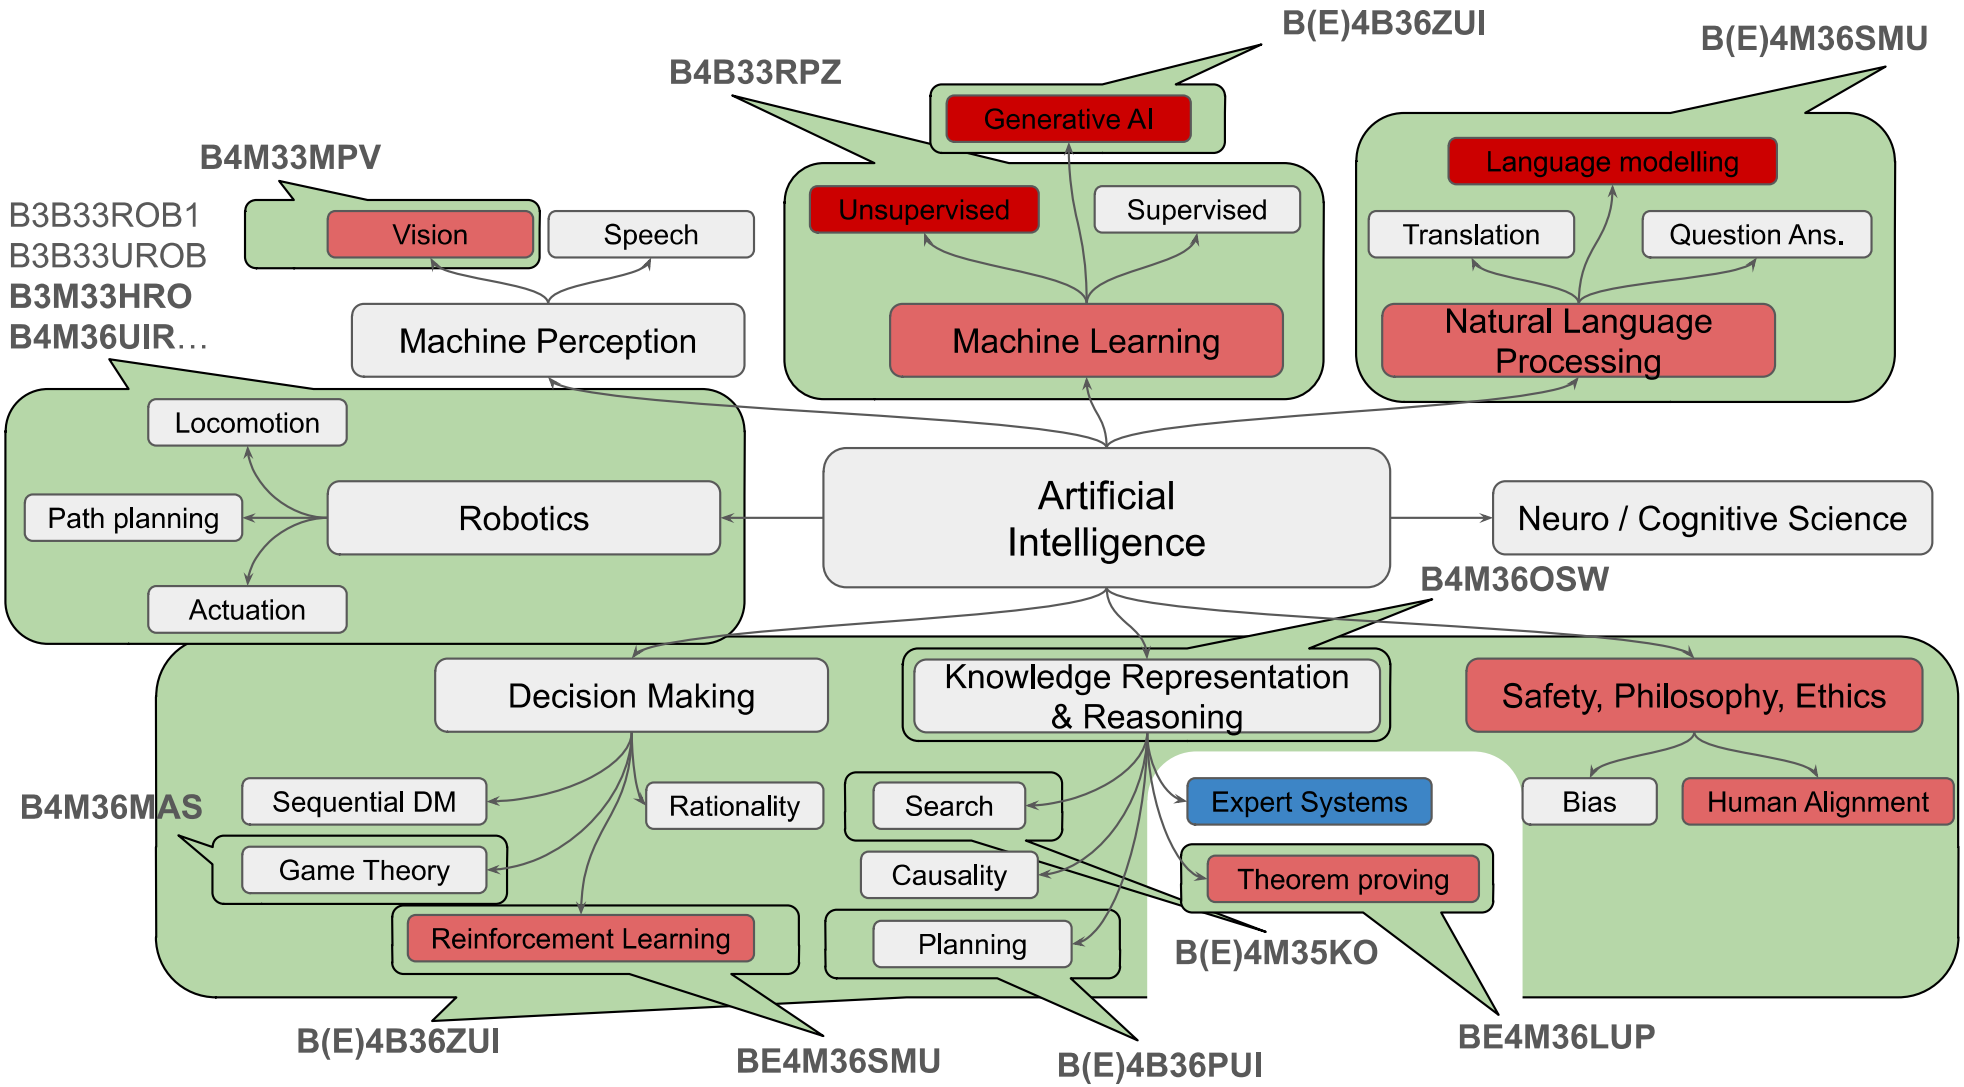

Course organisation

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


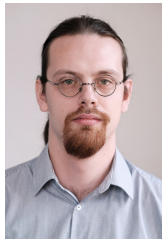


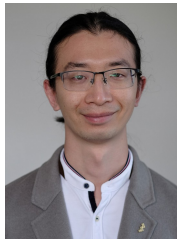

3 programming homeworks in python evaluated by BRUTE

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