

Artificial Intelligence

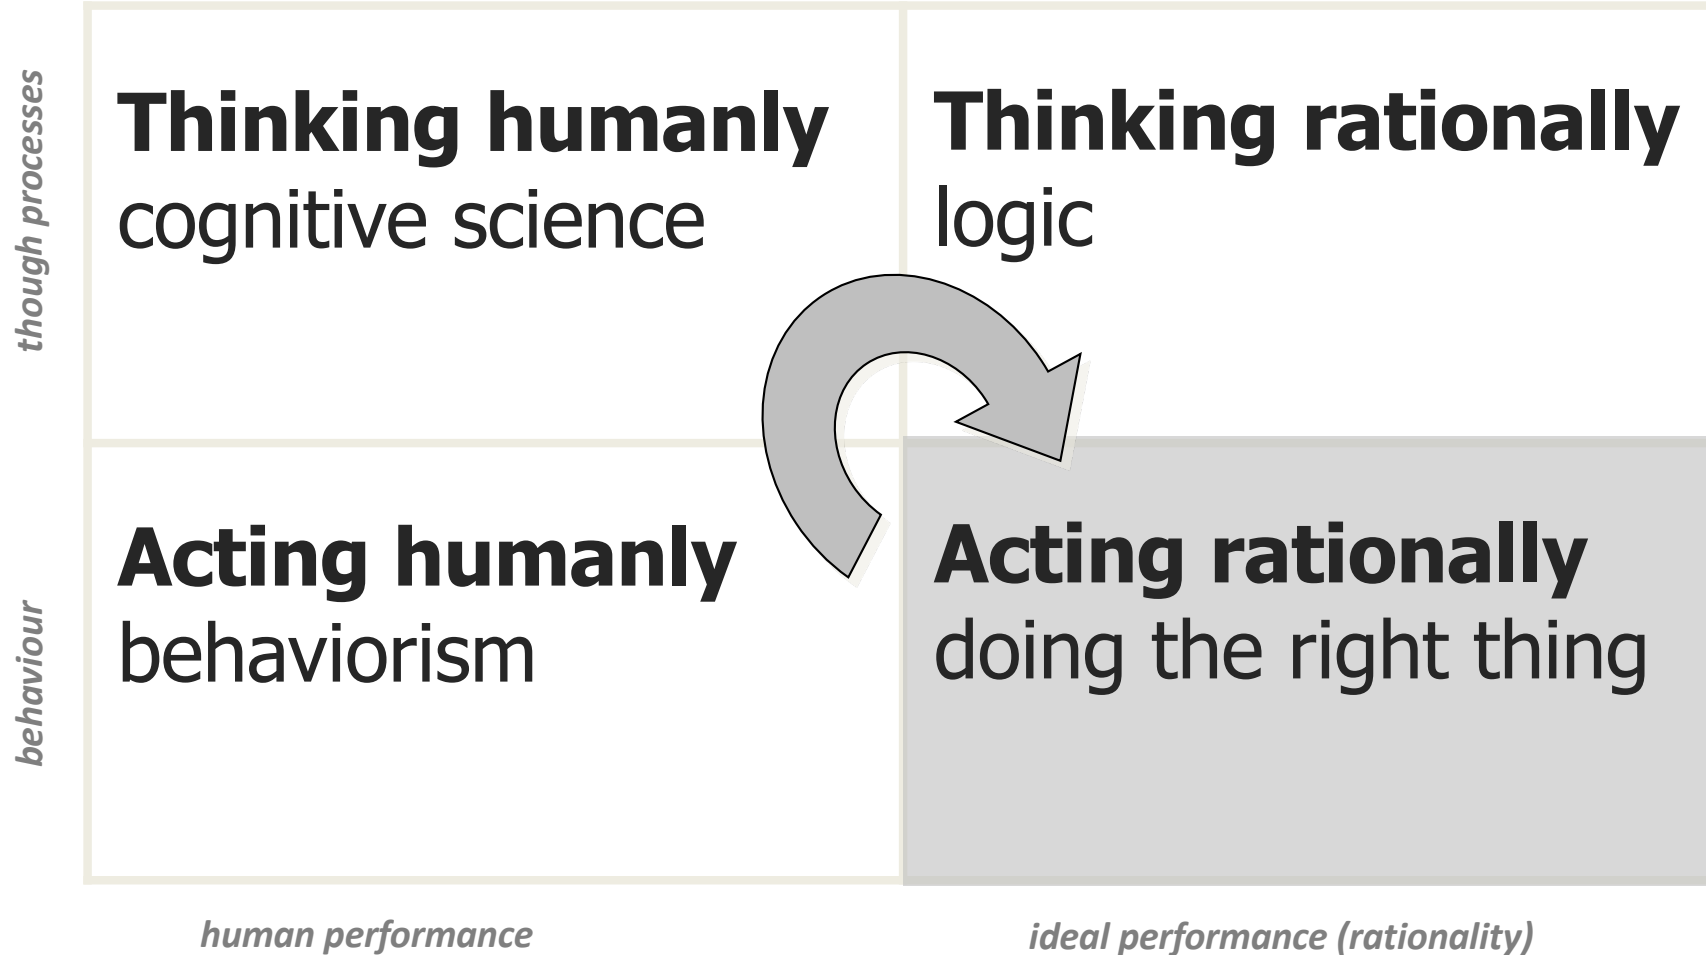
Roman Barták

Department of Theoretical Computer Science and Mathematical Logic

Artificial intelligence is the science of making machines do things that would require intelligence if done by men.

Marvin Minsky, 1967

Four Views to Artificial Intelligence



Alan Turing (1950) provided an operational definition of intelligence.


- „Can machines think?“ *like a man*
↳ „Can machines act ~~intelligently~~?“

– **Turing test**

A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer.

– **Required capabilities:**

- natural language processing
- knowledge representation
- automated reasoning
- machine learning
- computer vision
- robotics



Reverse Turing test

computer attempts to recognize whether it communicates with a computer or a person

Security Check: A Reverse Turing Test

In order to eliminate cheaters using automated bots, please enter the following eight characters into the text box below and click OK. You must do this within 2 minutes. Sorry for the inconvenience.

CX8 2 S 2 S A

Enter text here: **OK**

If the characters do not appear, you can try reloading: **Reload**

Mail Center
Send a Message



Please enter the text from the image above:
The letters are not case-sensitive.
Do not type spaces between the numbers and letters.


Please enter the image text.

Submit



Cognitive Modelling

- modelling human mind
- we must have some way of determining how humans think
 - Top-down approach (**psychology**)
 - following human reasoning steps (found through introspection or through observing a person in action)
 - GPS: General Problem Solver (Newell & Simon, 1957)
 - Bottom-up (**neuroscience**)
 - modelling the brain (through observing a brain in action)
 - connectionist models
 - „intelligent behaviour emerges by connecting a large number of simple units“



Since the time of **Aristotle** (384 – 322 BC) people attempted to codify „right thinking“


– *Syllogisms*

- Patterns for argument structures that always yield correct conclusions when given correct premises
- Socrates is a man, all men are mortal
⇒ Socrates is mortal

– This study initiated the field of **logic** (and mathematics)

Major obstacles:

- It is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain.
- There is a big difference between solving a problem „in principle“ and solving it in practice.




- **Rational behaviour** = doing „right things“
- „**right thing**“ = achieving the best (expected) outcome even when there is uncertainty
- Making correct inferences (thinking rationally) is part of being a **rational agent**, but not exclusive.
 - In some situations, there is no provable correct thing to do, but something must still be done.
 - There are also ways of acting rationally that cannot be said to involve inference (for example, reflex actions).
 - **This course concentrates on general principles of rational agents and on components for constructing them.**



- **Introduction**
 - a bit of history, context, intelligent agents
- **Problem Solving**
 - search algorithms, constraint satisfaction
- **Knowledge and Reasoning**
 - logic and logical inference, knowledge representation
- **Planning**
 - composing actions to achieve goals





Artificial Intelligence: A Modern Approach



- S. Russell and P. Norvig
- Prentice Hall, 2020 (fourth ed.)
- <http://aima.cs.berkeley.edu/>

Umělá inteligence 1-6

- Vladimír Mařík, Olga Štěpánková, Jiří Lažanský a kol.
- Academia



<http://ktiml.mff.cuni.cz/~bartak/ui/>

The screenshot shows a Windows Internet Explorer browser window displaying the course page for 'Umělá inteligence I'. The page title is 'Umělá inteligence I' and the course code is 'AII069, 2/0 Zk, zimní semestr'. The instructor is 'Roman Barták, KTIML'. The page includes a navigation menu with links for 'Zdroje', 'Přednáška', 'Zkouška', and 'Kontakt'. A quote from Marvin Minsky (1967) is present: 'Umělá inteligence je věda o vytváření strojů nebo systémů, které budou při řešení určitého úkolu užívat takového postupu, který - kdyby ho dělal člověk - bychom považovali za projev jeho inteligence.' The 'Zdroje' section lists the textbook 'Artificial Intelligence. A Modern Approach, Prentice Hall, 2003' and a series of books 'Umělá inteligence 1-5'. A 'Přednáška' section for the 2007/2008 semester is also shown, including a list of lecture dates and topics.

You can find there:

- slides
- links and resources
- contacts
- quiz
- ...

- **Seminar on Artificial Intelligence**
 - about theoretical and practical questions in a field of Artificial Intelligence
- **Constraint Programming**
 - about techniques of constraint satisfaction
- **Decision Procedures and SAT/SMT Solvers**
 - about logical inferences
- **Planning and Scheduling**
 - about automated construction of plans and schedules

The Foundations of Artificial Intelligence

Artificial Intelligence draw ideas and techniques from many disciplines.

- **Philosophy** (428 BC -) how does the mind arise from a brain?
logic, reasoning techniques
- **Mathematics** (800 -) what are the **formal** rules to draw valid conclusions?
what can be computed?
- **Economics** (1776 -) how to maximize payoff?
utility theory, decision processes
- **Neuroscience** (1861 -) how do **brains** process information?
the physical seat of consciousness
- **Psychology** (1879 -) how do humans think and act?
behaviourism
- **Computer engineering** (1940 -) how to build an efficient **computer**?
machines for information processing
- **Control theory** (1948 -) how can **artefacts** operate under their own control?
systems maximizing an objective function over time
- **Linguistics** (1957 -) how does **language** relate to thought?
knowledge representation

The gestation of AI (1943-1955)

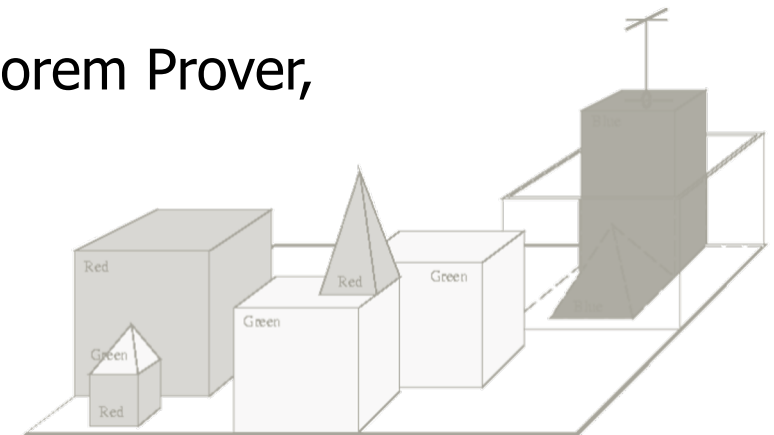
- W. McCulloch & W. Pitts: Boolean model of **neurons**
- A. Turing: „**Computing Machinery and Intelligence**“
the first complete vision of artificial intelligence

The birth of AI (1956)

- two-months workshop at **Dartmouth** College, NH
- J. McCarthy gave the name **Artificial Intelligence**
- A. Newell & H. Simon: software **Logic Theorist**

Great expectations (1952-1969)

- demonstrating one X after another from the list “a machine can never do X”
- General Problem Solver, Geometry Theorem Prover, **Lisp (1958)**, Analogy, blockworld
- J. McCarthy referred to this period as the „**Look, Ma, no hands!**“ era.




A dose of reality (1966-1973)

- “There are now machines that think, that learn and that create”, but only on simple problems
- Why?
 - the early programs **knew nothing of their subject matter**; they succeeded by means of simple syntactic manipulations
 - **intractability** of many problems that AI was attempting to solve (trying out different combinations of steps until the solution was found)
 - **fundament limitations** on the basic structures used (perceptron learns anything it can represent, but it could represent very little)

Knowledge-based systems (1969-1979)

- The alternative to „weak“ general methods is to use more powerful, domain-specific knowledge.
- expert (knowledge) systems:
 - **DENDRAL** (Buchanan)
inferring molecular structure from the information provided by a mass spectrometer, introducing **rules** based on well-known patterns to reduce possible structures
 - **MYCIN** (Feigenbaum)
diagnosing blood infections, introducing **certainty factors**
 - **PROLOG** (Colmerauer, 1972)
 - **frames** (Minsky, 1975) – motivations for current OOP



AI becomes an industry (1980)

- commercial expert system **R1** for configuring computers DEC (\$40 mil./year)
- **Fifth Generation** of computers (Japan, 1981)
 - a 10-year plan to build intelligent computers running Prolog
- **boom of AI industry** (billions of dollars in 1988)
- and then the „**AI Winter**“
 - companies failed to deliver on extravagant promises (like the dot.com bubble)

The return of neural networks (1986)

- reinventing back-propagation learning algorithm

AI adopts the scientific method (1987)

- AI has come firmly under the **scientific method**, hypothesis must be subjected to rigorous empirical experiments, and the results must be analysed statistically for their importance; experiments can be replicated
- formalisation and specialisation led to **fragmentation**

The emergence of intelligent agents (1995)

- encouraged by progress in solving the sub-problems of AI researchers started to look at the “whole agent” problem again
- SOAR (State, Operator and Result) – a complete agent architecture

Probabilistic reasoning (1987-present)

- brittleness of expert systems led to a new approach to handle uncertainty and vagueness
- Bayesian networks
- Hidden Markov models

Big data (2001-present)

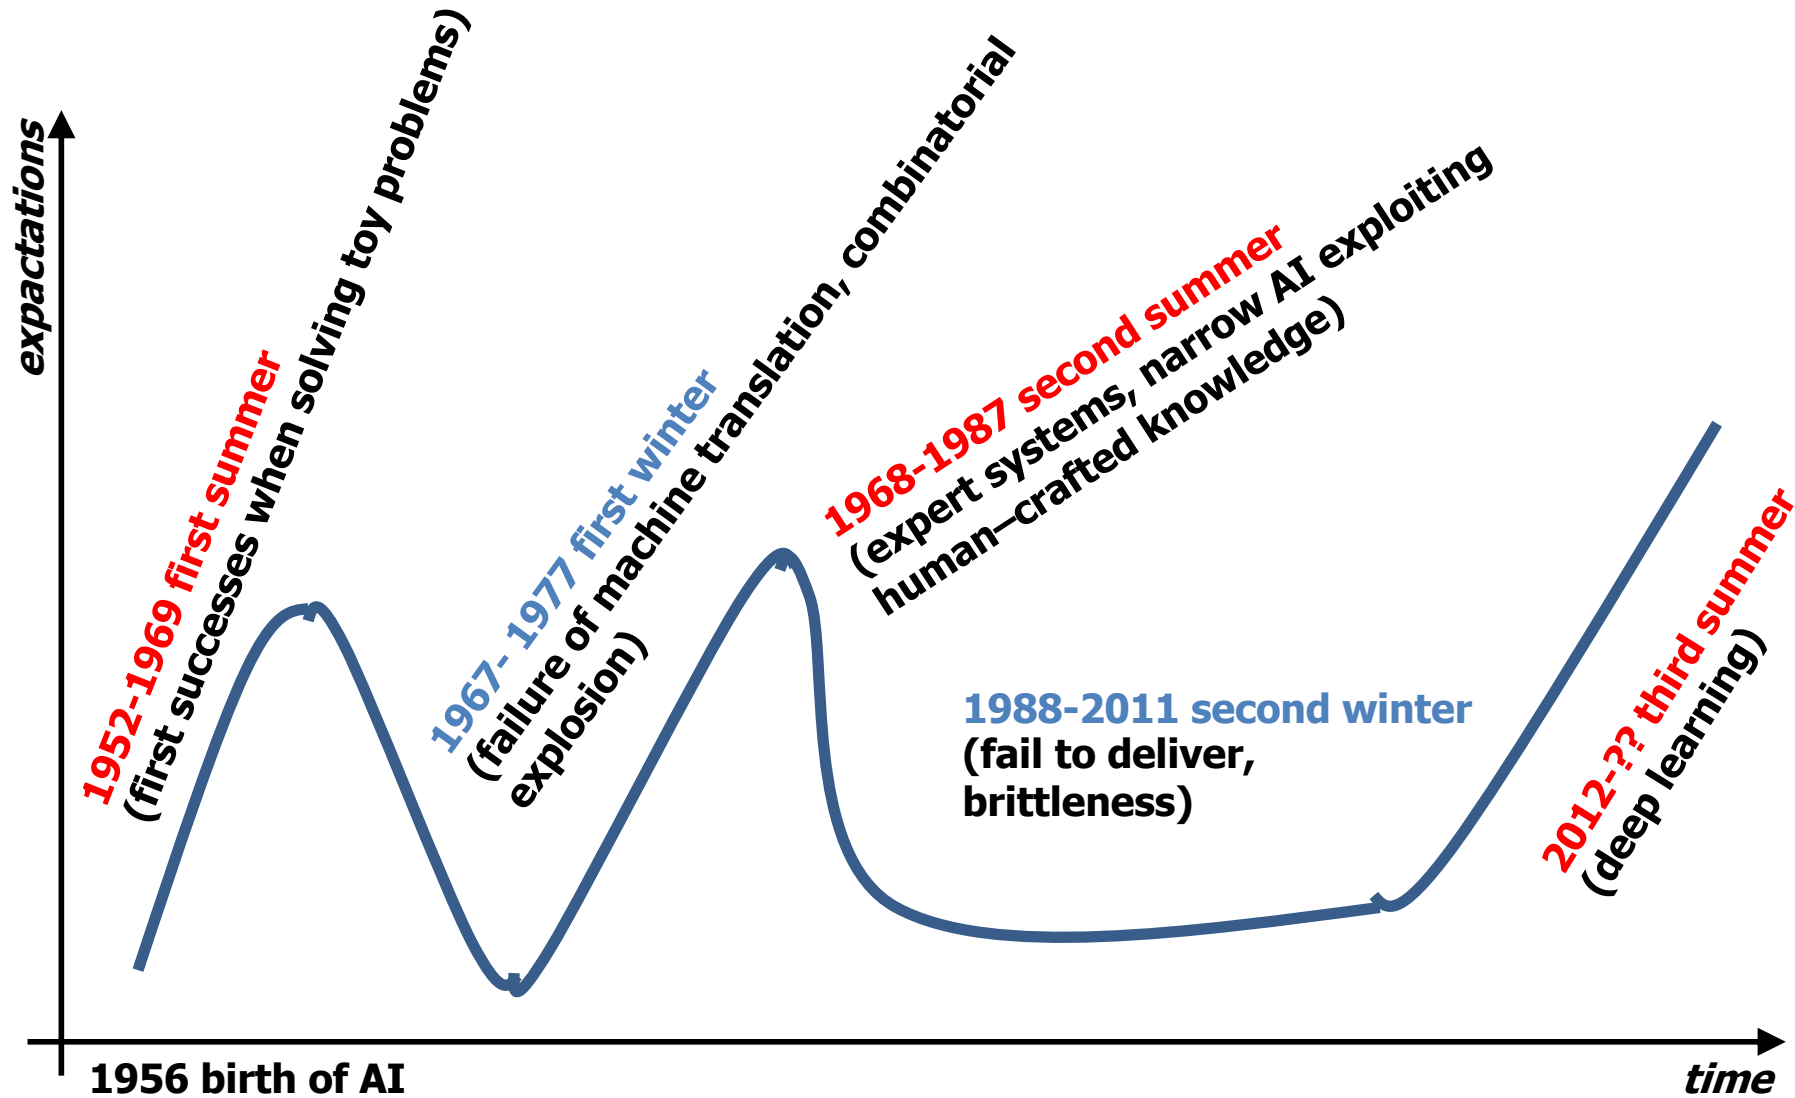
- advances in computing power and the creation of World Wide Web have facilitated the creation of very large data sets (big data)
- development of new learning algorithms
- recovery of commercial attractiveness

Deep learning (2011-present)

- deep learning refers to machine learning using multiple layers of simple computing elements
- relies heavily on powerful hardware (GPU, TPU, FPGA)
- applications in vision, speech recognition, machine translation, game playing

What is next?


The History of Artificial Intelligence




„By mid-21st century, a team of fully autonomous humanoid robot soccer players shall win the soccer game, complying with the official rule of the FIFA, against the winner of the most recent World Cup.“



- **Simulation league**
simulated games in computers
- **Small size league**
robots limited to a 18 cm diameter
- **Middle size league**
robots limited to a 50 cm diameter
all sensors
- **Standard platform league**
Sony Aibo, Nao
- **Humanoid league**
penalty kicks and two-to-two game



RoboCup Emotions



- The Grand Challenge was the first **long distance competition for driverless cars** in the world.
- The ultimate goal was making one-third of ground military forces autonomous by 2015.

– 2004 Grand Challenge


- Failure - None of the robot vehicles finished the route (max. 11,78 km, CMU)

– 2005 Grand Challenge

- Done! Winner Stanley (212.4 km in about 7 hours, Stanford)


– 2007 Urban Challenge

- Winner BOSS (CMU) driving in urban areas



Google Self-driving Car





© 2013 Roman Barták

Department of Theoretical Computer Science and Mathematical Logic

bartak@ktiml.mff.cuni.cz