Artificial Intelligence

Antal Péter
antal@mit.bme.hu

Computational Biomedicine (Combine) workgroup
Intelligent Systems group
Department of Measurement and Information Systems,
Budapest University of Technology and Economics

Course info

- Course site
  - http://www.mit.bme.hu/eng/oktatas/targyak/vimia313
- Lecturer
  - Péter Antal, antal@mit.bme.hu
- Schedule
  - Wednesday, 8.30-10.00,, IE.224 (building I, wing E, 2nd floor)
  - Friday, 8.30-10.00, IE.224 (building I, wing E, 2nd floor)
- Contact hour
  - Monday, 9.00-10.00, IE.423 (building I, wing E, 2nd floor)
- Book
  - S. Russell and P. Norvig Artificial Intelligence: A Modern Approach Prentice Hall, 2003, Second Edition
- Slides
  - Based on AIMA slides from S.Russel/T.Leanert/H.Ng
  - http://aima.cs.berkeley.edu/instructors.html

Academic calendar

- Last day of classes: 11 December 2015 (Friday)
- Repeat week (resits and late submission of home assignments): 14-18 December (Monday-Friday)
- Exams begin: 21 December 2015 (Monday)
- Duration of examination period 21 working days
- Last day of examination period: 27 January 2016 (Wednesday)
- **BME Students' Day – University Sports Day (no classes) 16 September 2015 (Wednesday)**
- **Anniversary of the 1956 Revolution (no classes) 23 October 2015 (Friday)**
- Scientific Conference for Students (no classes) 17 November 2015 (Tuesday)
- **Open day for secondary schools (no classes) 27 November 2015 (Friday)**
- SCH cup events (our faculty only) (no classes) 8 October 2015(Thursday)

Lectures

1. Introduction: what is AI
2. Decision theory and intelligent system design
3. Sporting day (no lectures)
4. Problem solving with search: uninformed
5. Problem solving with search: informed, optimization
6. Problem solving with constraint based programming/ in game playing
7. Problem solving: exercises
8. Logic: syntacs and semantics, propositional logic
9. Logic: first-order logic
10. Logic: exercises
11. Earlier midterms: exercises
12. Uncertainty and probability theory
13. Probabilistic modeling, simple models: Markov chains, Naïve BNs, HMMs
14. Day-off (Oct. 23)
15. Midterm exam
16. Bayesian networks and inference (observational, interventional, counterfactual)
17. Knowledge engineering of Bayesian networks: exercises
18. Exact inference: NBNs, HMMs, BNs
19. Stochastic inference in BNs. Applications and extensions of HMMs.
20. Inference with utilities and interventions: causal inference and decision networks
21. Learning: basic concepts, foundations (frequentist vs. Bayesian learning)
22. Learning logical expressions and decision trees, linear, loglinear, logregr, nonlinear
23. Reinforcement and active learning
24. Homework guide
25. Open university (no lectures)
26. Repeated midterm exam
27. Planning: towards decision theoretic & relational planning
28. Overview & homework presentations

midterm 28th of Oct.

repeated midterm 2nd of Dec.

Final
Pre final exam: 16th??? of December, 10am.
6th of January, 10am.

Homework, midterm, …grading

- Grading:
  - Homework, obligatory, min.40%, weight: 25%
  - Midterm test, obligatory, min.40%, weight: 25%
  - Final exam, min.40%, weight: 50%
  - Overall
    - 40<: satisfactory
    - 50<: fair
    - 65<: good
    - 80<: excellent
- Midterm test (in class, 90 minutes) and final exam (90 minutes) are both closed-book exams.

Overview

- What is intelligence? Artificial intelligence?
- Computational intelligence
- Theoretical computational models and Moore's law
- The knowledge era
- The data-intensive age
- Bayesian decision theory

What is AI?

AI approaches can be grouped as follows:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

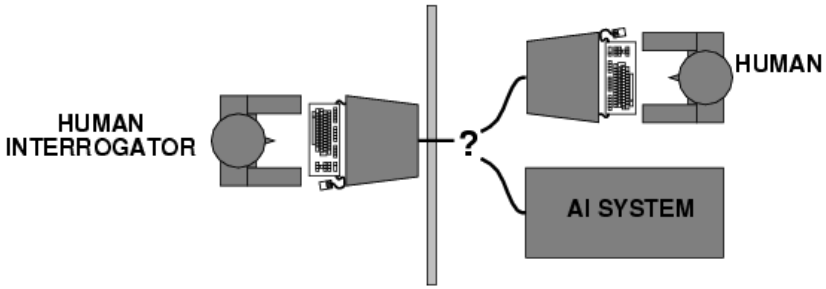

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
-

Thinking humanly: cognitive modeling

- 1960s "cognitive revolution": information-processing psychology
-
- Requires scientific theories of internal activities of the brain
-
- -- How to validate? Requires
  1) Predicting and testing behavior of human subjects (top-down)
  or 2) Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of *logic*: *notation* and *rules of derivation* for thoughts; may or may not have proceeded to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI
- Problems:
  1. Not all intelligent behavior is mediated by logical deliberation
  2. What is the purpose of thinking? What thoughts should I have?

Acting rationally: rational agent

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: \mathcal{P}^* \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
  → design best program for given machine resources

Decision theory probability theory+utility theory

- Decision situation:
  - Actions a_i
  - Outcomes o_j
  - Probabilities of outcomes $p(o_j \mid a_i)$
  - Utilities/losses of outcomes $U(o_j \mid a_i)$
  - Maximum Expected Utility Principle (MEU) $EU(a_i) = \sum_j U(o_j \mid a_i) p(o_j \mid a_i)$
  - Best action is the one with maximum expected utility $a^* = \arg\max_i EU(a_i)$

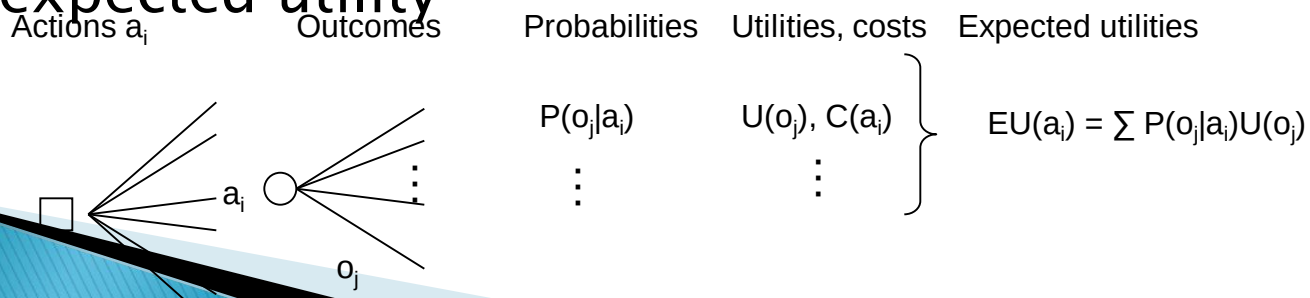

AI prehistory

Philosophy	Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
Mathematics	Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
Economics	utility, decision theory
Neuroscience	physical substrate for mental activity
Psychology	phenomena of perception and motor control, experimental techniques
Computer engineering	building fast computers
Control theory	design systems that maximize an objective function over time
Linguistics	knowledge representation, grammar

Milestones and phases in AI

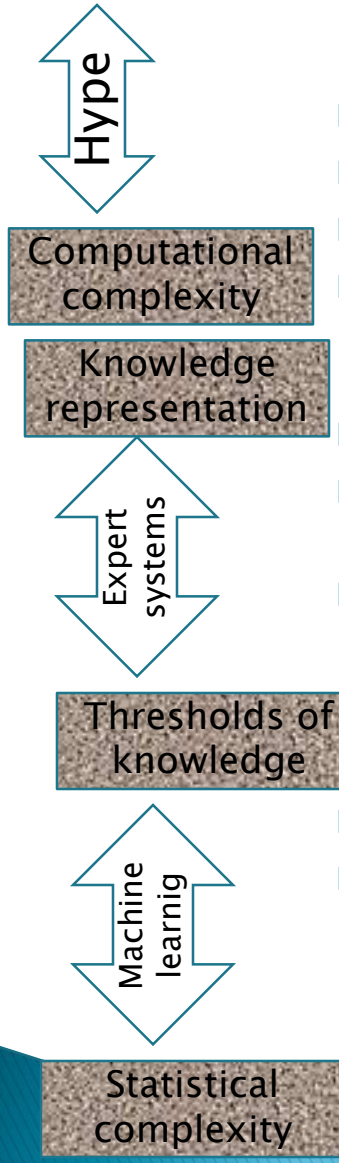

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: the term "Artificial Intelligence"
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems

- 1986-- Neural networks return to popularity
- 1988-- Probabilistic expert systems
- 1995-- Emergence of machine learning

Today: heterogeneous AI, data-intensive science, data and knowledge **fusion, automated science**

State of the art: ☺

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans
- Google search/car/face recognition/…

WHY CAN'T MY COMPUTER UNDERSTAND ME? (COMMON SENSE????)

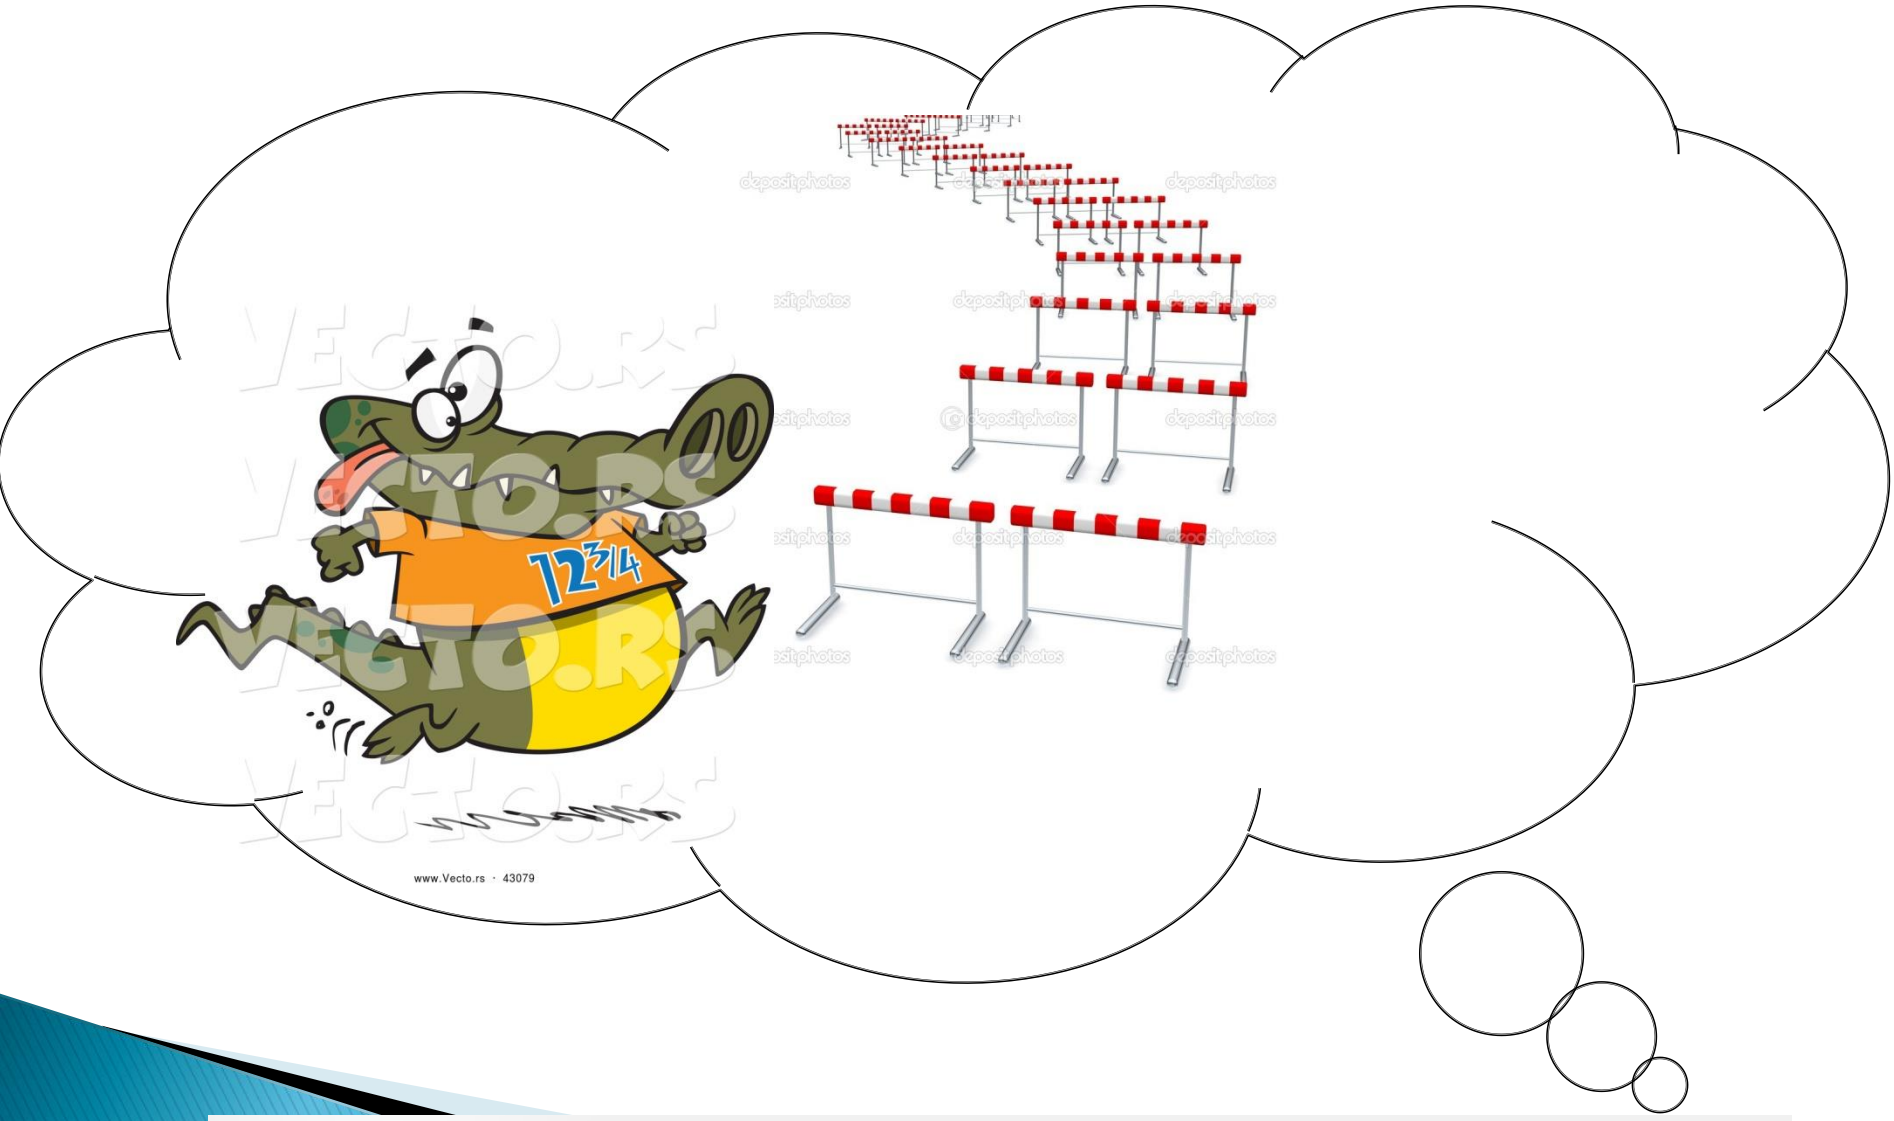

Google Can an alligator run the hundred-metre hurdles?

State of the art: ☹

- WHY CAN'T MY COMPUTER UNDERSTAND ME?
  - http://www.newyorker.com/online/blogs/elements/2013/08/why-cant-my-computer-understand-me.html
- Dreyfus claimed that he could see no way that AI programs, as they were implemented in the 70s and 80s, could capture this *background* or do the kind of fast problem solving that it allows. He argued that our unconscious knowledge could *never* be captured symbolically. If AI could not find a way to address these issues, then it was doomed to failure, an exercise in "tree climbing with one's eyes on the moon."[15]
  - http://en.wikipedia.org/wiki/Hubert_Dreyfus's_views_on_artificial_intelligence
- D.J. Chalmers: The Singularity: A Philosophical Analysis
  - http://consc.net/papers/singularity.pdf
- R. Kurzweil: How to Create a Mind: The Secret of Human Thought Revealed
  - http://www.amazon.ca/How-Create-Mind-Thought-Revealed/dp/0670025291
- INTEGRATED USE OF COMMON SENSE, EXPERT KNOWLEDGE, DATA
- CREATIVE USE OF COMMON SENSE, EXPERT KNOWLEDGE, DATA

Watson:
The Science Behind an Answer

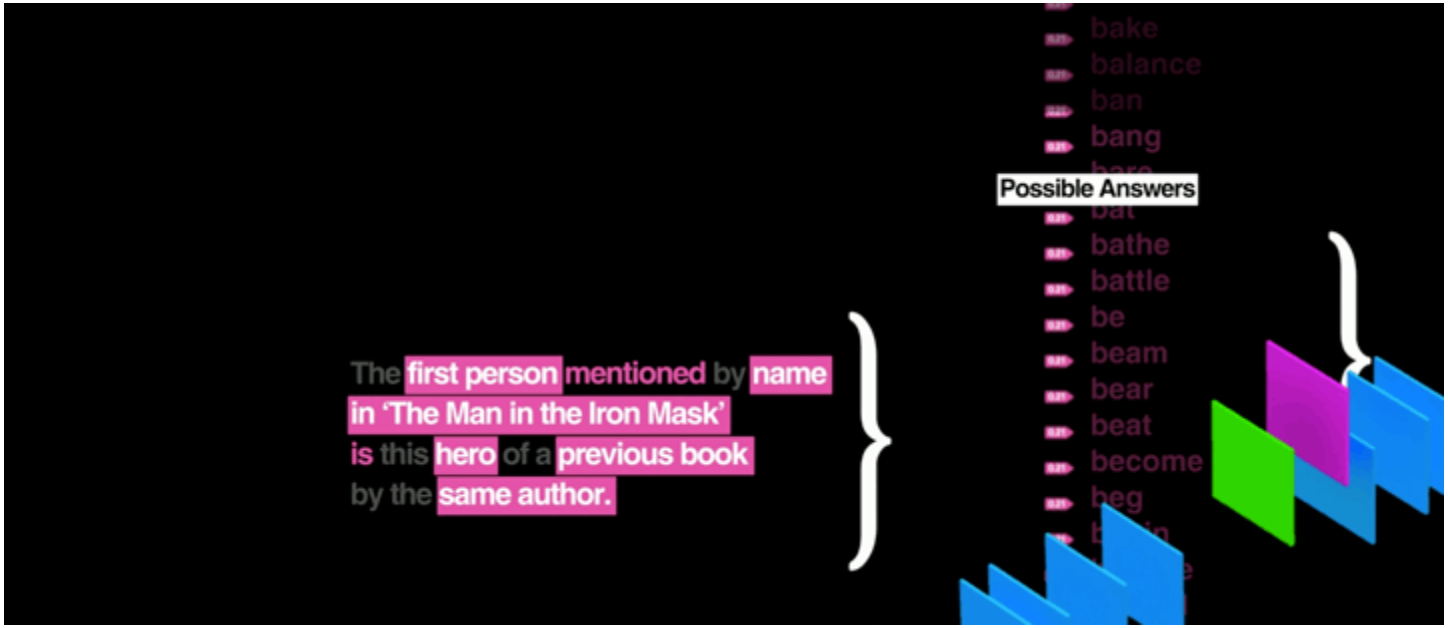

- http://www-03.ibm.com/innovation/us/watson/what-is-watson/science-behind-an-answer.html

Summary

- Four approaches to AI
- Computation-, knowledge- and data-intensive AI
- Decision theoretic foundation

- Additional suggested reading:
  - A.Turing: Computing machinery and intelligence, 1950
  - R.D.King: The Automation of Science, 2009
  - G.Marcus: WHY CAN'T MY COMPUTER UNDERSTAND ME?, 2013