Artificial Intelligence

Lecture 1

Ph. D. Lect. Horia Popa Andreescu

2019-2020 3rd year, semester 5

General description

- Mandatory lecture
- 1 semester
- 14 lectures
- 14 labs
- Written exam
- 2-3 Lab projects
- Extra credits opportunities
- The formula for computing the final grade will be 1+A1*0.2+A2*0.2+Exam*0.5.
- Where A1, A2 and Exam are grades from 0 to 10. You need to have at least 50% of the points for the 2 assignments to pass the exam. If you don't have the 50% you will sustain a practical exam, and still have the chance to pass the exam.

Bibliography

- [RN] Russel S., Norvig P. – Artificial Intelligence – A Modern Approach, 2nd ed. Prentice Hall, 2003 (1112 pages)
- [R] Stuart Russel – Course slides (visited oct. 2012 at http://aima.cs.berkeley.edu/instructors.html#homework)
- [GR] Giarratano J., Riley G. – Expert Systems Principles and Programming, 3rd edition, PWS Publishing, 2002
- [FH] Ernest Friedman-Hill – Jess in Action, Manning Publications, 2003
- [W1] Mark Watson – Practical Artificial Intelligence Programming With Java AI 3rd ed., 2008
- [C] D. Cârstoiu – Sisteme Expert, Editura ALL, București, 1994

TERMS

- Artificial – humanly contrived often of a natural model (Webster Dictionary)
- Intelligence - *(1)* **:** the ability to learn or understand or to deal with new or trying situations **:** reason; *also* **:** the skilled use of reason *(2)* **:** the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (as tests) (Webster Dictionary)
- Artificial Intelligence is:
  - "The science and engineering of making intelligent machines" John McCarthy 1955

Terms (II)

Systems that think like humans	Systems that think rationally
"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense." (Haugeland, 1985)	"The study of mental faculties through the use of computational models." (Chamiak and McDermott, 1985)
"[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Systems that act like humans	**Systems that act rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998)
"The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"AI . . .is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Figure 1.1 Some definitions of artificial intelligence, organized into four categories.

Russel & Norvig – Artificial Intelligence – A Modern Approach, pg. 2 (Abbreviated reference: [RN-2]

Turing Test – acting like humans

- The Turing Test

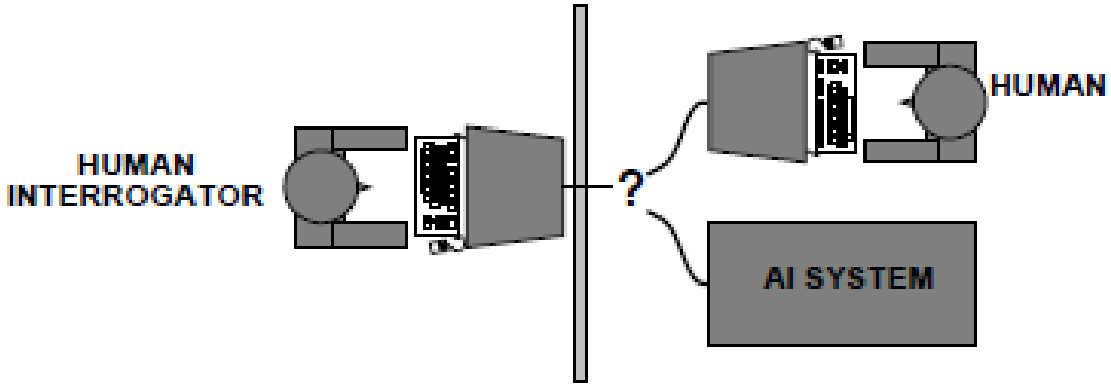

- The idea is that if the human interrogator isn't able to determine that the other person is a human or an AI system, then the AI system passes the test as behaving intelligent.
- To be able to pass it should possess the capabilities:
  - Natural language processing
  - Knowledge representation
  - Automated reasoning
  - Machine learning [R-ch1-3], [RN-2]
  - Computer vision and robotics for the "Total Turing Test"

Cognitive sciences – Thinking like humans

- Cognitive Sciences are a branch that separated from AI.
- Some see AI as one of the cognitive sciences.
- **Cognitive science** is the interdisciplinary scientific study of the mind and its processes.
- It examines what cognition is, what it does and how it works.
- It includes research on intelligence and behavior, especially focusing on how information is represented, processed, and transformed (in faculties such as perception, language, memory, reasoning, and emotion) within nervous systems (human or other animal) and machines (e.g. computers).
- Cognitive science consists of multiple research disciplines, including psychology, artificial intelligence, philosophy, neuroscience, linguistics, anthropology, sociology, and education.
- [Wikipedia – Cognitive science]

Thinking rationally

- 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. [R-1-6]
- By 1965, programs existed that could, in principle, solve *any solvable problem described in* LOGICIST logical notation.

Acting rationally

- 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. [R-1-7]

Rational agents

- An agent is something that acts (from latin agere = to do).
- A **rational agent** is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome. [RN-4]
- In the "laws of thought" approach to AI the emphasis is on correct inferences.

Roots of AI

- **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
- **Psychology:** adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
- **Economics:** formal theory of rational decisions
- **Linguistics:** knowledge representation, grammar
- **Neuroscience:** plastic physical substrate for mental activity
- **Control theory:** homeostatic systems, stability, simple optimal agent designs [R-1-9]

Isaac Asimov's laws of robotics

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

- The rules were introduced in Isaac Asimov's 1942 short story "Runaround" and are debated in the Robots series by the same author.
- The zeroth law, introduced in the Foundation series states:
  - *A robot may not harm humanity, or, by inaction, allow humanity to come to harm.*

History of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's “Computing Machinery and Intelligence"
- 1952-69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program,
  Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: “Articial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966-74 AI discovers computational complexity
  Neural network research almost disappears
- 1969-79 Early development of knowledge-based systems
- 1980-88 Expert systems industry booms
- 1988-93 Expert systems industry busts: "AI Winter"
- 1985-95 Neural networks return to popularity
- 1988- Resurgence of probability; general increase in technical depth
  "Nouvelle AI": ALife, GA’s, soft computing
- 1995- Agents, agents, everywhere
- 2003- Human-level AI back on the agenda [R-1-10]

History of AI

- The General Problem Solver elaborated by Newell and Simon in 1959 studied the way humans solve problems and tried to replicate that.
- They developed a theory called "Means-Ends Analysis" consisting in the following:
  - if an object is not the desired one, then the differences between the initial object (the means) and the final object (the ends) will be detected
  - We apply operators on the initial object trying to reduce the differences with the final object
  - If we cannot apply an operator it is preferable that we modify the entries for it so that it becomes applicable
  - Some differences will be harder to eliminate. It is preferable to try first to eliminate the "difficult" differences, even with the cost of introducing new differences but with a decreased difficulty

- The success of GPS and subsequent programs as models of cognition led Newell and Simon (1976) to formulate the famous **physical symbol system hypothesis**, which states that "a physical symbol system has the necessary and sufficient means for general intelligent action.“ [RN-18]
- **LISP** – developed by McCarthy in 1958 became one of the dominant AI programming language.
- **Prolog** – developed by a group around Alain Colmerauer in 1970 and **Scheme** by Carl Hewitt in 1969. They are based on Logic, the second incorporated also procedural programming.

Sub-fields of AI

- Neural Networks – e.g. brain modeling, time series prediction, classification
- Evolutionary Computation – e.g. genetic algorithms, genetic programming
- Vision – e.g. object recognition, image understanding, ~~OCR~~
- Robotics – e.g. intelligent control, autonomous exploration
- Expert Systems – e.g. decision support systems, teaching Systems
- Speech Processing– e.g. speech recognition and production
- Natural Language Processing – e.g. machine translation
- Planning – e.g. scheduling, game playing
- Machine Learning – e.g. decision tree learning, version space learning
- Cybernetics and brain simulation – almost extinct

Problems addressed by AI

- Deduction
- Reasoning
- Problem solving
- Knowledge representation
- Planning
- Learning
- Natural language processing
- Motion and manipulation – Robotics
- Perception
- Social intelligence
- Creativity
- General intelligence

Tools for Cognitive Sciences

- ACT-R (Adaptive Control of Thought-Rational) and

SOAR (State, Operator And Result), both developed at Carnegie Mellon University

Are used to model the way the human mind works.

Both have as component s **a production system** (see expert systems) and memory modules.

Both declare themselves to be "a cognitive architecture".

For ACT-R you have LISP, Java and Python implementations.

For Soar you have LISP, graphical interface, Herbal (a high level behaviour representation language)

Tools for expert systems

- CLIPS (C Language Integrated Production System) developed at NASA –Johnson Space Center from 1985
- Jess – a rule engine for the Java platform, developed from 1995 at Sandia National Labs. It is a superet of CLIPS
- CLIPS is free, Jess can be used free for educational purposes
- CLIPS is written in C and can be extended in C, also embedded in C programs
- Jess is written in Java, and can be extended and embedded in Java programs
- JessML – a declarative XML rule language for Jess

Tools for multiagent systems

- JADE - Java Agent DEvelopment Framework, originally developped at Telecom Italia, distributed as free software.
- Jade is standardized according to FIPA specifications, and uses the FIPA-ACL language for agent intercomunication
- WADE – Workflows and Agents Development Environment, based on the Jade platform, provides also support for execution of tasks according to the workflow metaphor.
- NetLogo – an agent simulation platform written in Java, it includes a GIS extension. It is used by a large community, with lots of tutorials and support, e.g. http://bookboon.com/en/artificial-intelligence-agents-and-environments-ebook
- Lots of other software for developing agent-based applications http://en.wikipedia.org/wiki/Comparison_of_agent-based_modeling_software

Tools for Data Mining and Machine Learning

- WEKA (Waikato Environment for Knowledge Analys) is a collection of machine learning algorithms for data mining tasks.
- The algorithms can either be applied directly to a dataset or called from your own Java code.
- Weka contains tools for data pre-processing, classification, regression, clustering, association rules, and visualization.
- It is also well-suited for developing new machine learning schemes.
- It is Open Source under GNU GPL, written in Java.
- Developped at University of Waikato, New Zeeland

Tools for Data Mining and Machine Learning (2)

- According to [5] the most popular
- [5] http://machinelearningmastery.com/best-programming-language -for-machine-learning/

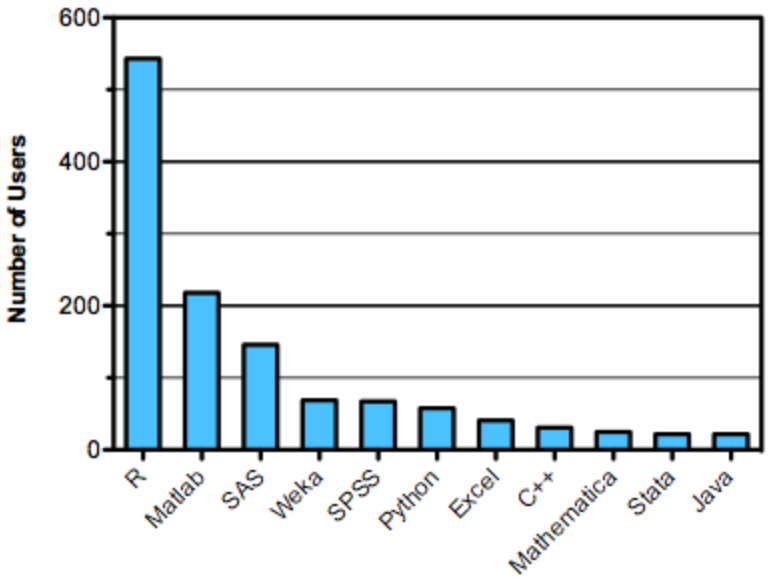

Tools for Data Mining and Machine Learning (3)

- So, Matlab has a toolbox called Statistics and Machine Learning Toolbox which has a lot of examples and algorithms implemented.

http://mathworks.com/products/statistics/features.html

- The non-comercial alternative to Matlab is Octave which also has a toolbox dedicated to machine learning

Tools for Data Mining and Machine Learning (4)

- Scikit-learn - a free software machine learning library for the Python programming language
- Tensorflow – also in Python, developed by Google
- Jupyter Notebook – supports Julia, Python and R and several others and has become a popular web interface for Cloud Computing.
- Google Colaboratory, written in Jupyter Notebook a popular online Machine Learning environment