Artificial Intelligence

Introduction

Radim Bělohlávek

DEPARTMENT OF COMPUTER SCIENCE
PALACKÝ UNIVERSITY, OLOMOUC

Content

- What is artificial intelligence (AI)?
- Literature
- Areas to be examined
  - Expert systems
  - Artificial neural networks
  - Evolutionary computation
  - Bayes networks
  - Knowledge representation
- Development of AI
- AI in context, societal and philosophical issues

What is AI?

What is AI?

- AI attempts to build intelligent entities.
- "Artificial" in contrast to "natural" or "human" which serves as a reference point.
- Emerged after World War II; term coined in 1956;
- one of the most recent fields in science and engineering
- one of most popular: "field I would most like to be in" by researchers in other fields
- opportunities to make fundamental discoveries
- covers a large variety of subfields, including general ones (such as learning, problem solving) as well as specific (theorem proving, playing games, driving car, predicting needs)
- present virtually in wide variety of commercial products
- increasing demand for AI experts

What is AI?

Some **definitions of AI** (more in Russell and Norvig, 2010):

"The exciting new effort to make computers think . . . machines with minds, in the full and literal sense." (Haugeland, 1985)

"[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . . " (Bellman, 1978)

"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)

"The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)

"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985)

"The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)

"Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

"AI . . . is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

What is AI?

The AI effect

D. Hofstadter (Gödel, Escher, Bach: An Eternal Golden Braid, 1980) quotes computer scientist L. G. Tesler:

- AI is whatever hasn't been done yet.

That is to say: As HW and SW become more powerful, what was considered intelligent system years ago, is no longer considered intelligent by today's standards. Each time something becomes feasible by computation, it is no longer considered magical and requiring intelligence to be solved.

P. Winston: "These days, it is hard to find a big system that does not work, in part, because of ideas developed or matured in the AI world."

M. Minsky: “This paradox resulted from the fact that whenever an AI research project made a useful new discovery, that product usually quickly spun off to form a new scientific or commercial specialty with its own distinctive name. These changes in name led outsiders to ask, Why do we see so little progress in the central field of artificial intelligence?”

N. Bostrom: “A lot of cutting edge AI has filtered into general applications, often without being called AI because once something becomes useful enough and common enough it’s not labeled AI anymore.”

Turing's test (1950)

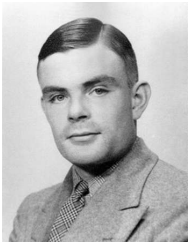

Figure: Alan Turing (source: Princeton university).

A person asks a computer a series of questions. If the person cannot tell whether the answers come from a computer or from another person, the computer passes the test.

Turing test appeared in seminal paper:
A. Turing,"Computing Machinery and Intelligence," Mind, LIX (236): 433–460(1950), doi:10.1093/mind/LIX.236.433

total Turing test:
The person can see the computer (i.e. the intelligent being) and can test its perceptual abilities and the ability to move objects.

From today's AI viewpoint, the basic test requires the following abilities:

- natural language processing,
- knowledge representation,
- automated reasoning,
- machine learning;

the total test, in addition:

- computer vision,
- robotics.

These fields are sometimes regarded as the fundamentals of AI (perhaps a historical view).

What is AI?

- Other names with similar content and closely related fields:
  - Machine intelligence
  - Intelligent systems
  - Computational intelligence
  - Machine learning
  - Data mining and knowledge discovery
  - Soft computing

Computational intelligence

Kruse et al. (Computational Intelligence, 2016):
Computational Intelligence comprises concepts, paradigms, algorithms, and implementations of systems that are supposed to exhibit intelligent behavior in complex environments. It relies heavily on sub-symbolic, predominantly nature-analog or at least nature-inspired methods. These methods have the advantage that they tolerate incomplete, imprecise and uncertain knowledge and thus also facilitate finding solutions that are approximative, manageable and robust at the same time.

Machine learning

Area of computer science which exploits various techniques to give computers the ability to "learn" from data, i.e. improve its behavior w.r.t. some goal.

Name coined in 1959 by A. Samuel:
"Some studies in machine learning using the game of checkers," IBM J. Research and Development 3 (3): 210–229; doi:10.1147/rd.33.0210

Data mining and knowledge discovery

Data mining: Area of computer science which aims at extraction of useful patterns from data.

Knowledge discovery: Almost synonymous, sometimes KD is regarded as involving broader aspects.

Soft computing

Lotfi Zadeh:
Soft computing differs from conventional (hard) computing in that, unlike hard computing, it is tolerant of imprecision, uncertainty, partial truth, and approximation. In effect, the role model for soft computing is the human mind. The guiding principle of soft computing is: Exploit the tolerance for imprecision, uncertainty, partial truth, and approximation to achieve tractability, robustness and low solution cost.

The principal constituents of Soft Computing (SC) are Fuzzy Logic (FL), Neural Computing (NC), Evolutionary Computation (EC) Machine Learning (ML) and Probabilistic Reasoning (PR), with the latter subsuming belief networks, chaos theory and parts of learning theory. What is important to note is that soft computing is not a melange. Rather, it is a partnership in which each of the partners contributes a distinct methodology for addressing problems in its domain. In this perspective, the principal constituent methodologies in SC are complementary rather than competitive. Furthermore, soft computing may be viewed as a foundation component for the emerging field of conceptual intelligence.

Literature

Basic:

- S. Russell, P. Norvig: Artificial Intelligence: A Modern Approach (3rd Edition). Prenticel Hall: Upper Saddle River, NJ, 2010.
- R. Kruse, C. Borgelt, C. Braune, S. Mostaghim, M. Steinbrecher: Computational Intelligence: A Methodological Introduction. 2nd ed. Springer, 2016.
- I. Goodfellow, Y. Bengio, A. Courville: Deep Learning. Boston, MA: The MIT Press, 2016.
- R. Belohlavek, J. W. Dauben, G. J. Klir: Fuzzy Logic and Mathematics: A Historical Perspective. Oxford University Press, NY, 2017.

Additional:

- N. Bostrom: Superintelligence: Paths, Directions, Stratgies. Oxford U. Press, 2016. (Superinteligence: Až budou stroje chytřejší než lidé. Prostor, 2017.)
- R. Kurzweil: The Age of Spiritual Machines: When Computers Exceed Human Intelligence. Penguin Books, NY, 2000.

Overview of some areas of AI

Expert systems

- software systems able to answer questions and/or guide a person in a particular problem area (medicine, finance, geology, etc.),
- mostly rule-based: ES = knowledge base (if-then rules) + inference mechanism
  examples:
  - classical expert systems in Prolog
  - fuzzy rule-based expert systems
- appealing features vs. difficulties

Artificial neural networks and deep learning

- ANNs are mathematical models inspired by biological neural networks,
- ANN is basically a network of simple information-processing units called artificial neurons which communicate information
- ANNs capability of learning from data
- many different kinds of ANNs exist, various purposes

Evolutionary algorithms

- algorithms to find solutions of problems, typically optimization problems
- inspired by biological evolution
- evolving is a population of individuals which represent possible solutions
- during evolution, crossover and mutation occurs, until a satisfactory solution evolves
- simple idea with many modifications
- surprising capabilities vs. drawbacks

Bayesian networks

- means to efficiently store and reason with uncertain knowledge,
- formally a probabilistic graphical model (directed acyclic graph) representing a set of random variables and their conditional dependencies,
- inferences may be drawn based on new information,
- enables analysis of dependencies and learning from data

Knowledge representation

- how to represent human knowledge in a computer?
- particular topic in AI, but of fundamental importance
- relating to almost all AI techniques, hence treated as separate topic
- mostly based on classical predicate logic (which forms so-called symbolic paradigm in AI)
- several particular approaches developed (Sowa's conceptual graphs, various sub-symbolic approaches such as ANNs)

Development of AI

Prehistory of AI

(for more info see Russell & Norvig; Bostrom)
The idea of automated reasoning and of creating an artificial intelligent being has been around for a long time. A short track, along with remarks on relevant developments and related questions follows.

- Aristotle (384–322 BC): founder of formal logic; many fundamental idea (formal deduction rules, syllogistic, . . .)
- Ramon Lull (14th century): idea that reasoning may be performed by a machine
- Thomas Hobbes (1588–1679): idea that reasoning resembles numerical computation; in *Leviathan* 1651: idea of artificial animal: "For what is the heart but a spring; and the nerves, but so many strings; and the joints, but so many wheels."
- Leonardo da Vinci (1452–1519): designed mechanical calculator (later reconstructions proved it functional)

- first calculating machine: by Wilhelm Schickard (1592–1635);
  another, better known by Blaise Pascal (1623–1662) in 1642
  Pascal mentioned that the machine produces effects which appear nearer to thought than all the actions of animals (but only addition and subtraction)
- more capable (not only addition and subtraction, but also multiplication and computing roots) by Gottfried Wilhelm Leibniz (1646–1716)
  idea of computing with concepts rather than just numbers;
  characteristica universalis: universal language capable of representing human knowledge
- ideas on: "mind seems to operate according to logical rules, need to build physical systems emulating these rules" vs "mind is such a physical system"; if mind is a physical system, there is little room for free will
- Reneé Descartes (1596–1650): discussions on mind and matter; D. proponent of *dualism* (part of human mind lies outside of nature; animals do not possess this dual quality)
- different from dualism is *materialism*: mind is a physical system that operates according to the laws of physics; free will is illusion

- relevant philosophical movements:
  - *rationalism*; reasoning is crucial for human understanding; Descartes was proponent;
  - *empiricism*: Francis Bacon (1561–1626) *Novum Organum*; John Locke (1632–1704) *An Essay Concerning Human Understanding* (1690)
    human sensation and perception is crucial for understanding of the outer world
    David Hume (1711–1776) *A Treatise of Human Nature*, principle of *induction*: general rules result from repeated exposure to examples
  - *logical positivism*: Vienna Circle led by Rudolf Carnap (1891–1970): human knowledge is explained by logical theories connected to observation sentences, which correspond to sensory inputs (i.e both rationalism and empiricism are present in this view)
    Carnap's *The Logical Structure of the World* (1928) theory of mind as a computational process

- from knowledge to action:
  Aristotle (*De Motu Animalium*)
  "But how does it happen that thinking is sometimes accompanied by action and sometimes not, sometimes by motion, and sometimes not? It looks as if almost the same thing happens as in the case of reasoning and making inferences about unchanging objects. But in that case the end is a speculative proposition ... whereas here the conclusion which results from the two premises is an action ... I need covering; a cloak is a covering. I need a cloak. What I need, I have to make; I need a cloak. I have to make a cloak. And the conclusion, the "I have to make a cloak," is an action."
  Aristotle suggests an algorithm in *Nicomachean Ethics*:
  "We deliberate not about ends, but about means. For a doctor does not deliberate whether he shall heal, nor an orator whether he shall persuade, ... They assume the end and consider how and by what means it is attained, and if it seems easily and best produced thereby; while if it is achieved by one means only they consider how it will be achieved by this and by what means this will be achieved, till they come to the first cause, ... and what is last in the order of analysis seems to be first in the order of becoming. And if we come on an impossibility, we give up the search, e.g., if we need money and this cannot be got; but if a thing appears possible we try to do it."

- Herbert A. Simon, J. C. Shaw, and Allen Newell's GPS (general problem solve, 1959) is a computer program that may be regarded as implementing Aristotle's algorithm. Problem expressed by first-order logic formulas, a directed graph with several sources (axioms) and sinks (conclusions).
- logic, computation, and probability as three pillars of AI
- logic: Aristotle and Ancient Greece, mathematization since George Boole (1815–1864), formalization, language of logic (esp. Gottlob Frege (1848–1925)), semantics: Alfred Tarski (1902–1983)
- computation, algorithms:
  first algorithm: Euclid's algorithm for finding gcd (greatest common divisor)
  word due to Muhammad ibn Musa al'Khwarizmi (lat. Algoritmi) 9th cent.
  formalization, early 20th century: David Hilbert (1862–1943), Kurt Gödel (1906–1978), Alonzo Church (1903–1995), Emil Post (1897–1954), Alan Turing (1912–1954)
  limitations of formal methods!: algorithmically non-solvable problems (Entscheidungsproblem, Turing, Church), Gödel's incompleteness theorems

- tractability (feasibility):
  polynomial vs exponential growth in complexity: first in the 1960s (Cobham, Edmonds)
  NP-completeness and developments in computability theory: Steven Cook (1971), Richard Carp (1972)
- probability:
  pioneers: Gerolamo Cardano (1501–1576), Blaise Pascal (1623–1662), Jacob Bernoulli (1654–1705), Pierre Simon Laplace (1749–1827), Thomas Bayes (1702–1761),
- economics' issues: making choices
  decision making, game theory, utility theory
  Frank Ramsey (1903–1930), John von Neumann (1903–1957), Oskar Morgenstern (1902–1977)
- Brain sciences, neuroscience
  brain subject of research (observing activity of individual neurons and neuronal activity) since 1900s century
  field much advanced (MRI etc.)

- Cognitive psychology
  study of how human concepts and thought
  principles of cognition
- Linguistics
  structure of natural languages, natural language processing, artificial languages; Noam Chomsky (b. 1928)
  computational linguistics
- Cybernetics; Norbert Wiener (1894–1964)
  studies of systems with feedback; stable adaptive behavior
- Computers; obvious

History of AI

Inception

- first step considered design of artificial neuron in 1943 by Warren McCulloch (1898–1969, neurophysiologist and cybernetician) and Walter Pitts (1923–1969, mathematician):
  "A Logical Calculus of Ideas Immanent in Nervous Activity"
- (Hebbian) learning rule, by Donald Hebb (1949)
- Marvin Minsky and Dean Edmonds, design of first neural network computer (SNARC, 1950)
- 1950s paper by Alan Turing: "Computing machinery and intelligence" (Turing test, ideas on machine learning,)
- John McCarthy (1927–2011), introduced the term "artificial intelligence"; developed Lisp,
  considered founder of AI (along with Marvin Minsky, Allen Newell, and Herbert A. Simon); two-month workshop in 1956 (Darmouth, NH)
- first logic program capable of proving theorems from *Principia Mathematica*; paper rejected from *J. Symbolic Logic*

Period of (exceeding) expectations (late 1950s and the 1960s)

- computers were shown to being capable of solving increasingly complex tasks
- great hopes
- 1959 GPS (general problem solver) program by H. Simon, J. Shaw, A. Newell
- 1976 physical symbol system hypothesis: "a physical symbol system has the necessary and sufficient means for general intelligent action" by Newell and Simon
- design and development of Lisp (proposed by J. McCarthy, 1958), main language for AI for decades
- 1965 J. Robinson: resolution method, later basis of Prolog language (programming in logic)
- networks of perceptrons, Frank Rosenblatt
- H. Simon (1957):
  It is not my aim to surprise or shock you—but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied.

- Expectations: in about 10 years a computer will beat man in chess; will prove a difficult mathematical theorem; etc. (This happened much later.)
- attempts to extend success of AI beyond simple problems failed
- failure of US government-sponsored automatic language translation program; funding of AI by US government for translation almost stopped
- intractability issues: works only on small scale (e.g. automated proving)
- Lighthill report: combinatorial explosion in larger-scale problems prevents AI techniques from success; end of British government support for AI
- M. Minsky and S. Papert's book *Perceptrons*: fundamental limitation of what may be achieved by simple perceptrons; this killed research in artificial neural networks for almost two decades

Period of hope in knowledge-based systems (1970s)

- reasoning systems filled with expert knowledge
- several important systems
  - DENDRAL (1969), chemistry; generated chemical structures for a given basic chemical formula and measurement by mass spectrometer
  - MYCIN (Feigenbaum, Buchanan, and Shortliff); medicine; system for medical diagnosis, 450 rules in knowledge base, outperformed junior doctors; worked with certainty factors

Commercial applications of AI start (1980)

- first was expert system R1 developed and employed by DEC
  helped design configurations of computers
  saving the company 40 mil. USD/year
  reasoning systems filled with expert knowledge
- major companies had their AI departments

Return of neural networks (late 1980s)

- development (rediscovery, actually) of backpropagation algorithm
- many other developments since then
- ANNs one of most popular AI methods

Intelligent system enter the scene (late 1990s)

- employment of various mathematically advanced techniques
- e.g. probabilistic models

Machine learning and data mining (2000s)

- new methods of various kinds
- large datasets

Zapoctove prace (pro 1 studenta, pokud neni uvedeno jinak)

1