Introduction to Artificial Intelligence

George Voutsadakis[1]

[1]Mathematics and Computer Science
Lake Superior State University

LSSU Math 400

Subsection 1

What is AI?

Intelligence and Artificial Intelligence

- Intelligence distinguishes humans among living beings.
  - What is intelligence?
  - How can one measure intelligence?
  - How does the brain work?
- All these questions are key in understanding artificial intelligence.
- The central concept is that of an intelligent machine that behaves like a person, exhibiting intelligent behavior.
- The attribute artificial raises the philosophical question of the desirability/possibility of understanding, modeling or reconstructing the soul.
- Even though defining the term artificial intelligence or AI simply and robustly is difficult, using some examples and historical definitions, may help to characterize the field of AI.

McCarthy's Definition 1955

- In 1955, John McCarthy, a pioneer of AI, was the first to define the term artificial intelligence, roughly as follows:

  The goal of AI is to develop machines that behave as though they were intelligent.

- To test this definition, imagine the following scenario:
  - Fifteen small robotic vehicles are moving on an enclosed four by four meter square surface.
  - One can observe various behavior patterns:
    - Some vehicles form small groups with relatively little movement.
    - Others move peacefully and gracefully avoid any collisions.
    - Others appear to follow a leader.
    - Aggressive behaviors are also observable.
  - Is what we are seeing intelligent behavior?
  - According to McCarthys definition, the robots can be described as intelligent.
  - However, the psychologist Braitenberg has shown that this seemingly complex behavior can be produced by very simple electrical circuits.

Braitenberg Vehicles

- The Braitenberg vehicles have two wheels, each of which is driven by an independent electric motor.
- The speed of each motor is influenced by a light sensor on the front of the vehicle. The more light that hits the sensor, the faster the motor runs.

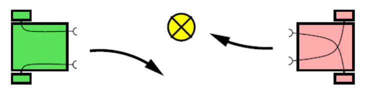

- Vehicle 1 on the left, according to its configuration, moves away from a point light source. Vehicle 2 on the other hand moves toward the light source.
- Further small modifications can create other behavior patterns, such that, despite the simplicity, impressive behavior is created.
- McCarthy's definition is insufficient because AI has the goal of solving difficult practical problems which are too demanding for the Braitenberg vehicles.

Characterizing AI

- The definition in Encyclopedia Britannica reads:

  AI is the ability of digital computers or computer controlled robots to solve problems that are normally associated with the higher intellectual processing capabilities of humans...

- But this definition also has weaknesses:
  - A computer with large memory that can save a long text and retrieve it on demand displays intelligent capabilities (memorization of long texts is considered a higher intellectual processing capability of humans);
  - Similarly, with the quick multiplication of two 20-digit numbers.
  - So, according to this definition, every computer is an AI system.

- The following definition by Elaine Rich is more to the point:

  Artificial Intelligence is the study of how to make computers do things at which, at the moment, people are better.

  Rich characterizes accurately and concisely AI research and will always be up-to-date.

AI Tasks

- In tasks such as the execution of many computations in a short amount of time digital computers far outperform humans.
- In many other areas humans are far superior to machines.
  - A person entering an unfamiliar room will recognize the surroundings and, if necessary, quickly make decisions and plan actions.
  - This task is still too demanding for autonomous robots.
  - According to Rich's definition, this is therefore a AI task.
- Construction of chess computers has lost relevance because they already play at or above the level of grandmasters.
- Rich's definition does not necessarily imply that AI is only concerned with the pragmatic implementation of intelligent processes.

Reasoning, Adaptability and Learning

- Intelligent systems cannot be built without a deep understanding of human reasoning and intelligent action in general, because of which neuroscience is of great importance to AI.
- A particular strength of human intelligence is adaptability. We are capable of adjusting to various environmental conditions and change our behavior accordingly through learning.
- Since our learning ability is so vastly superior to that of computers, machine learning is, according to Rich's definition, a central subfield of AI.

The Brain and its Functions

- Neural Networks: Through research on intelligent systems we can try to understand how the human brain works and then model or simulate it on the computer.
- Goal-Oriented Approaches: Starting from a problem, we try to find the most optimal solution. Like in Medicine, there is no universal method for all application areas of AI.
- Cognitive Science: In relation to algorithms and implementations, it provides hints on how human reasoning functions.
- Philosophy: Humans have consciousness; that is, we can think about ourselves and even ponder that we are able to think about ourselves.
  - How does consciousness come to be? Many philosophers and neurologists believe that the mind and consciousness are linked with matter, that is, with the brain.
  - The question of whether machines could one day have a mind or consciousness could at some point in the future become relevant.

The Turing Test

- Alan Turing, an early pioneer of AI, defined an intelligent machine as one that must pass the following Turing Test:
  - The test person Alice sits in a locked room with two terminals.
    - One terminal is connected to a machine;
    - The other with a non-malicious person Bob.
  - Alice can type questions into both terminals.
  - She is given the task of deciding, after five minutes, which terminal belongs to the machine.
  - The machine passes the test if it can trick Alice at least 30% of the time.
- While the test is very interesting philosophically, for practical AI, which deals with problem solving, it is not a very relevant test.
  - The reasons are similar to those related to Braitenberg vehicles.
  - The AI pioneer Weizenbaum developed a program named Eliza, which is meant to answer a test subject's questions like a human psychologist and was able to demonstrate success in many cases.
  - Today on the internet there are many so-called chatterbots, some of whose initial responses are quite impressive.

Subsection 2

Some History

First Beginnings

- In the 1930s Gödel, Church, and Turing laid important foundations for logic and theoretical computer science.
- Of particular interest for AI are Gödel's theorems:
  - The Completeness Theorem states that first-order predicate logic is complete. This means that every true statement that can be formulated in predicate logic is provable using the rules of a formal calculus.
  - Automatic theorem provers would later be constructed as implementations of formal calculi.
  - The Incompleteness Theorem showed that in higher-order logics there exist true statements that are unprovable.
  - This exhibited the limitations of formal systems.

First Beginnings (Cont'd)

- Turing's proof of the undecidability of the halting problem showed that there is no program that can decide whether a given arbitrary program with given input will run in an infinite loop.
- Turing, thus, identified another limit for intelligent programs.
  - It follows, for example, that there will never be a universal program verification system.
- In the 1940s, based on results from neuroscience, McCulloch, Pitts and Hebb designed the first mathematical models of neural networks. However, computers at that time lacked sufficient power to simulate simple brains.

Logic Solves (Almost) All Problems

- In the 1950's the advent of programmable computers gave a boost to AI as a practical science of thought mechanization.
  - Newell and Simon introduced Logic Theorist, the first automatic theorem prover.
  - McCarthy introduced LISP, a programming language specially created for the processing of symbolic structures.
  - Both were introduced in 1956 at the Dartmouth Conference, which is considered the birthday of AI.
  - Then, the logical inference rule of resolution developed into a complete calculus for predicate logic.
  - In the 1970s the logic programming language PROLOG was introduced, offering the advantage of allowing direct programming using Horn clauses, a subset of predicate logic.
  - A series of impressive achievements in symbol processing sustained an optimistic spirit, especially among many logicians, until well into the 1980s.

Logic Solves (Not as Many) Problems

- With the Fifth Generation Computer Systems project in Japan and the ESPRIT program in Europe, heavy investment went into the construction of intelligent computers.
- For small problems, automatic provers and other symbol-processing systems sometimes worked very well.
- The combinatorial explosion of the search space (dramatic increase as input size grows) limited the scope of these successes.
- Because the economic success of AI systems fell short of expectations, funding for logic-based AI research, especially in the United States, fell dramatically during the 1980s.

The New Connectionism

- Computer scientists, physicists, and cognitive scientists were able to show that mathematically modeled neural networks are capable of learning using training examples, to perform tasks which previously required costly programming.
- Because of the fault-tolerance of such systems and their ability to recognize patterns, considerable successes became possible.
  - E.g., Facial recognition in photos and handwriting recognition.
  - The system Nettalk was able to learn speech from example texts.
- Under the name connectionism, a new subdiscipline of AI was born.
- But feasibility limits became again obvious:
  - The neural networks could acquire impressive capabilities, but it was usually not possible to capture the learned concept in simple formulas or logical rules.
  - Attempts to combine neural nets with logical rules or the knowledge of human experts met with great difficulties.
  - No satisfactory solutions to the structuring and modularization of the networks were found.

Reasoning Under Uncertainty

- Could one combine logic's ability to explicitly represent knowledge with neural networks' strength in handling uncertainty?
  - Probabilistic reasoning works with conditional probabilities for propositional calculus formulas. Many diagnostic and expert systems have been built using Bayesian networks. They exploit intuitiveness, clean semantics and probability theory.
  - Fuzzy logic introduces infinitely many truth values between zero and one. It is being successfully utilized, especially in control engineering.
  - Hybrid systems successfully combine logic and neural networks. Neural networks are employed to learn heuristics for reduction of the huge combinatorial search space in proof discovery.
  - Decision tree learning from data also works with probabilities. It is a favorite among machine learning techniques.
  - Data mining has developed as a subdiscipline of AI in the area of statistical data analysis for extraction of knowledge from large databases. An application is steering ad campaigns of big businesses based on analysis of many millions of purchases by their customers.

Synopsis of the History of Various AI Areas

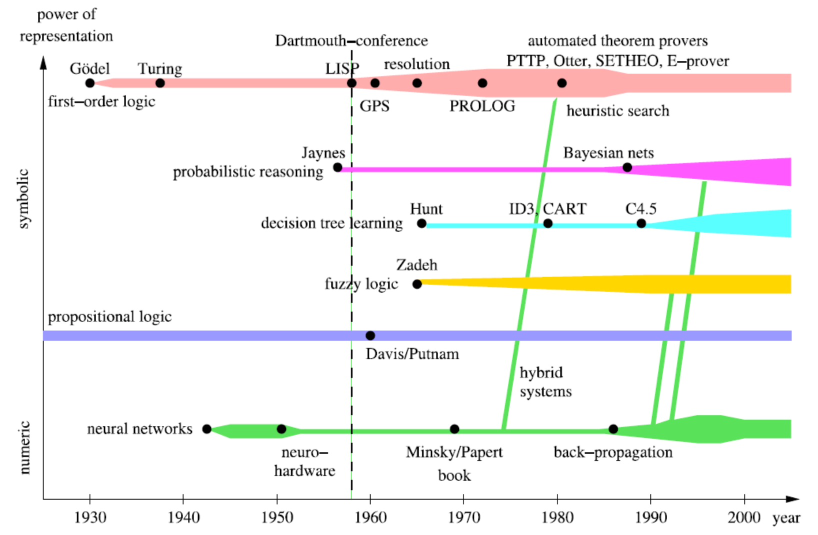

Distributed, Autonomous and Learning Agents

- Distributed artificial intelligence, DAI, since about 1985, uses parallel computers to increase the efficiency of problem solvers.
- Because of the high computational complexity of problems, the use of "intelligence" is more beneficial than parallelization.
- Autonomous software agents and robots that cooperate may exhibit intelligent behavior and solve a problem, where individual agents fail.
  - An ant colony is capable of erecting complex buildings, despite the fact that no single ant comprehends how the whole thing fits together.
  - In provisioning bread for a large city, rather than planning centrally, hundreds of bakers bake the amount needed for their neighborhoods.
- Active skill acquisition by robots is an exciting area of current research.
  - E.g., there are robots today that independently learn to walk or to perform various motor skills related to soccer.
  - Cooperative learning of multiple robots to solve problems together is still in its infancy.

AI Maturity and Toolbox

- The available AI systems do not provide a universal recipe, but a collection of tools to handle different tasks.
- Most of these tools are well-developed and are available as finished software libraries, often with convenient user interfaces.
- The selection of the right tool and its sensible use in each individual case is left to the AI developer or knowledge engineer. This requires a solid education, which this course is meant to promote!
- AI is heavily interdisciplinary; It draws upon such diverse fields as
  - logic,
  - operations research,
  - statistics,
  - control engineering,
  - image processing,
  - linguistics,
  - philosophy,
  - psychology,
  - neurobiology.
- The many subject areas make AI projects challenging but exciting.

Subsection 3

Agents

Intelligent Agents and Classification Criteria

- An (**intelligent**) **agent** denotes a system that processes information and produces an output from an input.
- These agents may be classified in many different ways:
  - Software vs. Hardware Agents;
  - Reflex vs. Agents with Memory;
  - Goal-Based Agents;
  - Cost-Based and Utility-Based Agents;
  - Learning Agents;
  - Distributed Agents;
  - According to their environment:
    - Fully or Partially Observable;
    - Deterministic or Nondeterministic;
    - Discrete or Continuous.

Software vs. Hardware Agents

- In classical computer science, software agents are primarily employed:

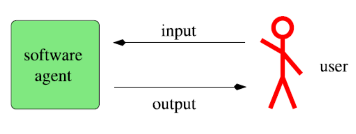

The agent is a program that calculates a result from user input.

- In robotics hardware agents (robots) are employed, which additionally have sensors and actuators at their disposal:

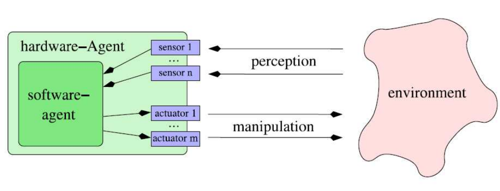

The agent can perceive its environment with the sensors. With the actuators it carries out actions and changes its environment.

Reflex vs. Agents with Memory

- With respect to the intelligence of the agent, there is a distinction between:
  - Reflex agents, which only react to input.
    If a reflex agent is controlled by a deterministic program, it represents a function of the set of all inputs to the set of all outputs.
  - Agents with memory, which can also include the past in their decisions.
    An agent with memory is in general not a function.
- A driving robot that through its sensors knows its exact position (and the time) has no way, as a reflex agent, of determining its velocity.
- A robot that saves its position, at short, discrete time steps, can calculate its average velocity in the previous time interval.
- Reflex agents are sufficient in cases where the problem to be solved involves a Markov decision process, i.e., a process in which only the current state is needed to determine the optimal next action.

Goal-Based Agents

- Agents whose actions depend on the goal are called goal-based agents.
- Example: A spam filter is an agent that puts incoming emails into wanted or unwanted (spam) categories, and deletes any spam.

  Its goal is to put all emails in the right category. The agent will try to make as few errors as possible.

  But things are not always that simple! Let us compare two agents:
  - Out of 1,000 emails, Agent 1 makes only 12 errors.
  - Agent 2 makes 38 errors with the same 1,000 emails.

  The errors of both agents are shown the following "confusion matrix":

Agent 1	Wanted	Spam
Class Wanted	189	1
Class Spam	11	799

Agent 2	Wanted	Spam
Class Wanted	200	38
Class Spam	0	762

  Agent 1 makes fewer errors than Agent 2, but those few errors are severe because the user loses 11 potentially important emails. Because there are two types of errors of differing severity, each error should be weighted with the appropriate cost factor.

Cost- and Utility-Based Agents

- The **total cost** caused by erroneous decisions is the sum of all weighted errors.
- The goal of a cost-based agent is to minimize the cost of erroneous decisions in the long term, that is, on average.
- The **total utility** is the sum of all decisions weighted by their respective utility factors.
- The goal of a utility-based agent is to maximize the utility derived from correct decisions in the long term, that is, on average.

Other Types of Agents

- Learning agents are capable of changing themselves, given training examples or through positive or negative feedback, such that the average utility of their actions grows over time.
- Distributed agents are those whose intelligence is not localized in one agent, but can only be revealed through cooperation of many agents.
- The design of an agent is strongly oriented, along with its objective, toward its environment, or its image of the environment, which strongly depends on it sensors.
  - The environment is (fully) observable if the agent always knows the complete state of the world. Otherwise the environment is only partially observable.
  - If an action always leads to the same result, then the environment is deterministic. Otherwise it is nondeterministic.
  - In a discrete environment only finitely many states and actions occur. A continuous environment has infinitely many states or actions.

Subsection 4

Knowledge-Based Systems

Agent Architecture

- For simple tasks, looking at agents as programs implementing mappings from perceptions to actions may be sufficient.
- For complex applications in which the agent must be able to rely on a large amount of information and is meant to perform difficult tasks, AI provides a clear agent architecture by distinguishing components:
  - Separate knowledge from the system or program, which uses the knowledge to, for example, reach conclusions, answer queries, or come up with a plan. This system is called the **inference mechanism**. The knowledge is stored in a **knowledge base (KB)**.
  - Acquisition of knowledge in the knowledge base is termed **Knowledge Engineering** and is based on various knowledge sources such as human experts, the knowledge engineer, and databases.

General Agent Architecture

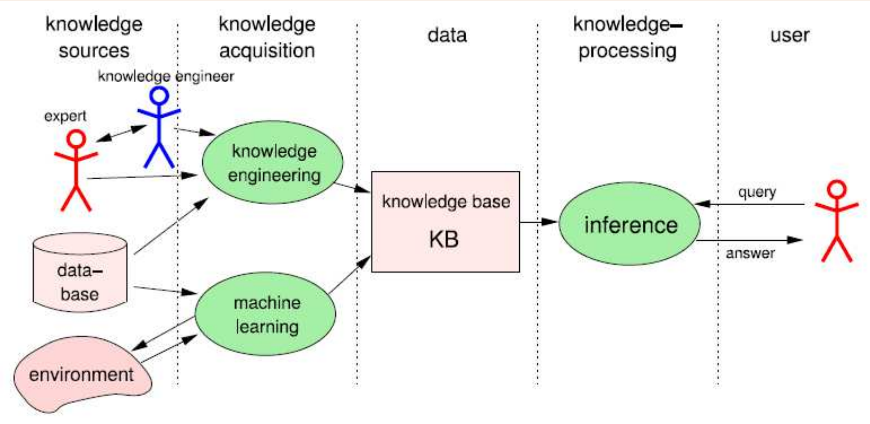

Advantages of Separating Knowledge from Processing

- Separating knowledge and inference has several crucial advantages:
  - It allows inference systems to be implemented in an application-independent way.
  - Knowledge can be stored declaratively. In the knowledge base there is only a description of the knowledge, which is independent from the inference system in use.
  - Without this clear separation, knowledge and processing of inference steps would be interwoven, and any changes to the knowledge would be very costly.

Knowledge Representation

- Representation of knowledge in the knowledge base can be achieved via the use of a formal language.
- Several choices are possible:
  - propositional calculus;
  - first-order predicate logic;
  - probabilistic logic;
  - fuzzy logic;
  - decision trees.
- An example for a large scale knowledge based system is the software agent "Watson".
  - It was developed at IBM together with a number of universities.
  - It is a question answering program, that can be fed with clues given in natural language.
  - It works on a knowledge base comprising four TBs of hard disk storage, including the full text of Wikipedia.

Watson's Aims and Million Dollar Prize

- Watson was developed within IBMs DeepQA project which is characterized as

  > shaping a grand challenge in Computer Science that aims to illustrate how the wide and growing accessibility of natural language content and the integration and advancement of Natural Language Processing, Information Retrieval, Machine Learning, Knowledge Representation and Reasoning, and massively parallel computation can drive open-domain automatic Question Answering technology to a point where it clearly and consistently rivals the best human performance.

- In a February 2011 "Jeopardy!" show, Watson defeated two human champions in a two-game, combined-point match and won the one million dollar prize.
- One of Watsons particular strengths was its very fast reaction to the questions: Watson often hit the buzzer faster than its human competitors and then was able to give the first answer to the question.