Inspiratron.org - Natural language processing, machine learning and cybersecurity

Artificial intelligence pt. 1

by Nikola Miloševi? - Wednesday, May 01, 2013

https://inspiratron.org/blog/2013/05/01/artificial-intelligence-applications/

Yesterday we created Artificial intelligence section in company I am working (Prelovac media), so it would be great reason to write about AI. I realized that many people are not aware what is current state, where we are heading and what are applications of artificial intelligence at the moment. So lets start with the basics.

AI definitions

Artificial intelligence (**AI**) is the intelligence of machines or software, and is also a branch of computer science that studies and develops intelligent machines and software. Major AI researchers and textbooks define the field as "the study and design of intelligent agents",where an intelligent agent is a system that perceives its environment and takes actions that maximize its chances of success.

In artificial intelligence, an **intelligent agent** (**IA**) is an autonomous entity which observes through sensors and acts upon an environment using actuators (i.e. it is an agent) and directs its activity towards achieving goals (i.e. it is rational)' Intelligent agents may also learn or use knowledge to achieve their goals.

History of Artificial intelligence

Concept of artificial intelligence and smart robots are very old. First pointing on this topic we can see in antique Greece. Thinking machines and artificial beings appear in Greek myths, such as Talos of Crete, the bronze robot of Hephaestus, and Pygmalion's Galatea. Since these are miths, they might not be taken seriously, but also Aristotel (384 BC–322 BC) was one who described first mechanical thinking. Aristotle described the syllogism, a method of formal, mechanical thought. Also famous is mith of rabbi Judah Loew ben Bezalel of Prague is said to have invented the Golem, a clay man brought to life. By the 19th and 20th centuries, artificial beings had become a common feature in fiction, as in Mary Shelley's Frankenstein or Karel ?apek's R.U.R (Rossum's Universal Robots). In his play Karel ?apek invented term robot, that is in still in use. It etymology comes from Czech word "robit" that means working. So robot is machine meant to work (instead of human).

Turing's theory of computation suggested that a machine, by shuffling symbols as simple as "0" and "1", could simulate any conceivable (imaginable) act of mathematical deduction.

By the middle of the 1960s, research in the U.S. was heavily funded by the Department of Defense and laboratories had been established around the world. AI's founders were profoundly optimistic about the future of the new field: Herbert Simon predicted that "machines will be capable, within twenty years, of doing any work a man can do" and Marvin Minsky agreed, writing that "within a generation ... the problem of creating 'artificial intelligence' will substantially be solved".

They had failed to recognize the difficulty of some of the problems they faced. In 1974, in response to the criticism of Sir James Lighthill and ongoing pressure from the US Congress to fund more productive projects, both the U.S. and British governments cut off all undirected exploratory research in AI. The next few years would later be called an "AI winter", a period when funding for AI projects was hard to find.

In the early 1980s, AI research was revived by the commercial success of expert systems,a form of AI program that simulated the knowledge and analytical skills of one or more human experts. By 1985 the market for AI had reached over a billion dollars.

In the 1990s and early 21st century, AI achieved its greatest successes, albeit somewhat behind the scenes. Artificial intelligence is used for logistics, data mining, medical diagnosis and many other areas throughout the technology industry. The success was due to several factors: the increasing computational power of computers, a greater emphasis on solving specific subproblems, the creation of new ties between AI and other fields working on similar problems, and a new commitment by researchers to solid mathematical methods and rigorous scientific standards.

On 11 May 1997, Deep Blue became the first computer chess-playing system to beat a reigning world chess champion, Garry Kasparov. In 2005, a Stanford robot won the DARPA Grand Challenge by driving autonomously for 131 miles along an unrehearsed desert trail. Two

years later, a team from CMU won the DARPA Urban Challenge when their vehicle autonomously navigated 55 miles in an Urban environment while adhering to traffic hazards and all traffic laws. In February 2011, in a Jeopardy! quiz show exhibition match, IBM's question answering system, Watson, defeated the two greatest Jeopardy champions, Brad Rutter and Ken Jennings, by a significant margin.

Goals of artificial intelligence

What are the goals of artificial intelligence? This could be good question both from technical, philosophical and ethical perspectives. There are several ways industry and science approaches AI. Here are some goals:

Strong intelligence

One of the greatest approach I heard is to create general AI, as brain that could be given problem, create algorithm to solve the problem. This is one of hardest challenges today.

Deduction, reasoning, problem solving

Early AI researchers developed algorithms that imitated the step-by-step reasoning that humans use when they solve puzzles or make logical deductions.For difficult problems, most of these algorithms can require enormous computational resources – most experience a "combinatorial explosion": the amount of memory or computer time required becomes astronomical when the problem goes beyond a certain size. The search for more efficient problem-solving algorithms is a high priority for AI research.AI has made some progress at imitating this kind of "sub-symbolic" problem solving: embodied agent approaches emphasize the importance of sensorimotor skills to higher reasoning; neural net research attempts to simulate the structures inside the brain that give rise to this skill; statistical approaches to AI mimic the probabilistic nature of the human ability to guess.

Knowledge representation

Knowledge representation and knowledge engineering are central to AI research. Many of the problems machines are expected to solve will require extensive knowledge about the world. Among the things that AI needs to represent are: objects, properties, categories and relations between objects; situations, events, states and time; causes and effects;knowledge about knowledge (what we know about what other people know); and many other, less well researched domains.

Planning

Intelligent agents must be able to set goals and achieve them. They need a way to visualize the future (they must have a representation of the state of the world and be able to make predictions about how their actions will change it) and be able to make choices that maximize the utility (or "value") of the available choices.

Learning

Machine learning is the study of computer algorithms that improve automatically through experience and has been central to AI research since the field's inception. Unsupervised learning is the ability to find patterns in a stream of input. Supervised learning includes both classification and numerical regression.Classification is used to determine what category something belongs in, after seeing a number of examples of things from several categories. Regression is the attempt to produce a function that describes the relationship between inputs and outputs and predicts how the outputs should change as the inputs change.

Natural language processing

Natural language processing gives machines the ability to read and understand the languages that humans speak. A sufficiently powerful natural language processing system would enable natural language user interfaces and the acquisition of knowledge directly from human-written sources, such as Internet texts. Some straightforward applications of natural language processing include information retrieval (or text mining) and machine translation. NLP is part of AI that I like most, and my master thesis was in this field. This incorporates machine learning, planing and knowledge representation to create very useful things like search engines, question answering engines, sentiment analysis engines etc.

Motion and manipulation

The field of robotics is closely related to AI. Intelligence is required for robots to be able to

handle such tasks as object manipulation and navigation, with sub-problems of localization (knowing where you are, or finding out where other things are), mapping (learning what is around you, building a map of the environment), and motion planning (figuring out how to get there) or path planning (going from one point in space to another point, which may involve compliant motion. When you say AI many people thinks only about intelligent robots. Since this comes from our miths that are thousands of years old, this is I guess natural.

Perception

Machine perception is the ability to use input from sensors (such as cameras, microphones, sonar and others more exotic) to deduce aspects of the world.

There are some goals like social intelligence, but it could be incorporated in the fields above. This text got quite long, since topic is here for millenniums. And in last decade it literally exploded. And there is for sure space for couple of more posts. What do you think that is the goal of AI that should be developed further? What do you think where the AI is heading? Please leave a comment.
