Artificial Intelligence (AI)

Introduction

Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason, and take action. Artificial Intelligence is a way of making a computer, a computer-controlled robot, or a software think intelligently, in the similar manner the intelligent humans think. It is the Intelligence displayed by machines, in contrast with the natural intelligence (NI) displayed by humans and other animals. It is a science and technology based on disciplines such as Computer Science, Biology, Psychology, Neuron Science, Linguistics, Mathematics, and Engineering. Out of these areas, one or multiple areas can contribute to build an intelligent system. A major thrust of AI is in the development of computer functions associated with human intelligence, such as reasoning, learning, and problem solving. The field of AI research was born at a workshop at the Dartmouth College in 1956.

Goals

AI has two major goals: Firstly, To create Expert Systems which exhibit intelligent behaviour, learn, demonstrate, explain, and advice its users and Secondly, to Implement Human Intelligence in Machines i.e. Creating systems that understand, think, learn, and behave like humans. The general problem of simulating (or creating) intelligence has been broken down into sub-problems. These consist of particular traits or capabilities (Reasoning, problem solving, Knowledge representation, planning, Learning, Natural language processing, Perception, Motion and manipulation, Social intelligence, Creativity and General intelligence) that researchers expect an intelligent system to display.

Approaches

A few of the most long standing questions that have remained unanswered are these: Should artificial intelligence simulate natural intelligence by studying psychology or neurology? Or is human biology as irrelevant to AI research as bird biology is to aeronautical engineering? Can intelligent behaviour be described using simple, elegant principles (such as logic or optimization)? Or does it necessarily require solving a large number of completely unrelated problems? Scientists have divided AI research into three approaches, which they call computational psychology, computational philosophy, and computer science. Computational psychology is used to make computer programs that mimic human behaviour. Computational philosophy, is used to develop an adaptive, free-flowing computer mind. Implementing computer science serves the goal of creating computers that can perform tasks that only people could previously accomplish. Together, the humanesque behaviour, mind, and actions make up artificial intelligence. The approaches followed are defined by choosing goals of the computational model and basis for evaluating performance of the system. These are Cybernetics and brain simulation, Symbolic, Sub-symbolic, and Statistical.

Tools

AI has developed a large number of tools to solve the most difficult problems in computer science. A few of the most general of these methods are Search and optimization, Logic, Probabilistic methods for uncertain reasoning, Classifiers and statistical learning methods,

neural networks, Deep feed-forward neural networks, Deep recurrent neural networks, Control theory, and Languages.

Evaluating progress

The Turing test, developed by Alan Turing in 1950, is a test of a machine's ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human. Turing proposed that a human evaluator would judge natural language conversations between a human and a machine designed to generate human-like responses. The evaluator would be aware that one of the two partners in conversation is a machine, and all participants would be separated from one another. The conversation would be limited to a text-only channel such as a computer keyboard and screen so the result would not depend on the machine's ability to render words as speech. If the evaluator cannot reliably tell the machine from the human, the machine is said to have passed the test. The machine is considered to have achieved intelligence.

Platforms for AI

A platform or computing platform is defined as "some sort of hardware architecture or software framework (including application frameworks), that allows software to run". A wide variety of platforms has allowed different aspects of AI to develop, ranging from expert systems such as Cyc to deep-learning frameworks to robot platforms such as the Roomba with open interface. Recent advances in deep artificial neural networks and distributed computing have led to a proliferation of software libraries, including Deeplearning4j, TensorFlow, Theano and Torch.

Applications:

While the rate of progress in AI has been patchy and unpredictable, there have been significant advances since the field's inception sixty years ago. In the twenty-first century, AI techniques have experienced a resurgence following concurrent advances in computer power, large amounts of data, and theoretical understanding; and AI techniques have become an essential part of the technology industry, helping to solve many challenging problems in computer science. AI has been dominant in various fields such as:

Gaming: AI plays crucial role in strategic games such as chess, poker, tic-tac-toe, etc., where machine can think of large number of possible positions based on heuristic knowledge. For example: Deep Blue was a chess-playing computer developed by IBM. It is known for being the first computer chess-playing system to win both a chess game and a chess match against a reigning world champion under regular time controls. Artificial intelligence is used in video games to generate intelligent behaviours primarily in non-player characters (NPCs), often simulating human-like intelligence.

Natural Language Processing: It is possible to interact with the computer that understands natural language spoken by humans.

Vision Systems: These systems understand, interpret, and comprehend visual input on the computer. For example, a spying aeroplane takes photographs which are used to figure out spatial information or map of the areas, Police use computer software that can recognize the face of criminal with the stored portrait made by forensic artist.

Expert Systems: There are some applications which integrate machine, software, and special information to impart reasoning and advising. They provide explanation and advice to the users.

Healthcare: Doctors use clinical expert system to diagnose the patient. Microsoft has developed AI to help doctors find the right treatments for cancer. IBM has created its own

artificial intelligence computer, the IBM Watson, which has beaten human intelligence (at some levels). It was declared a hero after successfully diagnosing a women who was suffering from leukaemia. A medical Centre in Washington has successfully demonstrated soft-tissue surgery with an autonomous robot. Another study was reported to have found that artificial intelligence was as good as trained doctors in identifying skin cancers.

Speech Recognition: Some intelligent systems are capable of hearing and comprehending the language in terms of sentences and their meanings while a human talks to it. It can handle different accents, slang words, noise in the background, change in human's noise due to cold, etc.

Handwriting Recognition: - The handwriting recognition software reads the text written on paper by a pen or on screen by a stylus. It can recognize the shapes of the letters and convert it into editable text.

Intelligent Robots: Robots are able to perform the tasks given by a human. They have sensors to detect physical data from the real world such as light, heat, temperature, movement, sound, bump, and pressure. They have efficient processors, multiple sensors and huge memory, to exhibit intelligence. In addition, they are capable of learning from their mistakes and they can adapt to the new environment.

Automotive: Advancements in AI have contributed to the growth of the automotive industry through the creation and evolution of self-driving vehicles. A few companies involved with AI include Tesla, Google, and Apple. Many components contribute to the functioning of self-driving cars. These vehicles incorporate systems such as braking, lane changing, collision prevention, navigation and mapping. Together, these systems, as well as high performance computers, are integrated into one complex vehicle.

Finance and Economics: Banks use artificial intelligence systems today to organize operations, maintain book-keeping, invest in stocks, and manage properties. AI has also reduced fraud and financial crimes by monitoring behavioural patterns of users for any abnormal changes or anomalies. AI can react to changes overnight or when business is not taking place. Financial institutions have long used artificial neural network systems to detect charges or claims outside of the norm, flagging these for human investigation.

While impressive, these technologies are highly tailored to particular tasks. Each application typically requires years of specialized research and careful, unique construction. In similarly targeted applications, substantial increases in the future uses of AI technologies, including more self-driving cars, healthcare diagnostics and targeted treatments, and physical assistance for elder care can be expected. AI and robotics will also be applied across the globe in industries struggling to attract younger workers, such as agriculture, food processing, fulfilment centres, and factories. They will facilitate delivery of online purchases through flying drones, self-driving trucks, or robots that can get up the stairs to the front door. AI influences and challenges, such as the difficulty of creating safe and reliable hardware (transportation and service robots), the difficulty of smoothly interacting with human experts (healthcare and education), the challenge of gaining public trust (low-resource communities and public safety and security), the challenge of overcoming fears of marginalizing humans (employment and workplace), and the social and societal risk of diminishing interpersonal interactions.

Potential risks and moral reasoning

AI is developing with such an incredible speed, sometimes it seems magical. There is an opinion among researchers and developers that AI could grow so immensely strong that it would be difficult for humans to control. Some say that the development of full artificial

intelligence could spell the end of the human race. Once humans develop artificial intelligence, it will take off on its own and redesign itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded. Humans developed AI systems by introducing into them every possible intelligence they could, for which the humans themselves now seem threatened. Some predict that it is a threat to privacy**.** An AI program that recognizes speech and understands natural language is theoretically capable of understanding each conversation on e-mails and telephones. Some consider that it is a threat to human dignity**.** AI systems have already started replacing the human beings in few industries. It should not replace people in the sectors where they are holding dignified positions which are pertaining to ethics such as nursing, surgeon, judge, police officer, etc. Some consider AI as a threat to Safety**.** The self-improving AI systems can become so mighty than humans that could be very difficult to stop from achieving their goals, which may lead to unintended consequences.

Conclusion

Contrary to the more fantastic predictions for AI there is no cause for concern that AI is an imminent threat to humankind. No machines with self-sustaining long-term goals and intent have been developed, nor are they likely to be developed in the near future. Instead, increasingly useful applications of AI, with potentially profound positive impacts on our society and economy are likely to emerge. At the same time, many of these developments will spur disruptions in how human labour is augmented or replaced by AI, creating new challenges for the economy and society more broadly. Application design and policy decisions made in the near term are likely to have long-lasting influences on the nature and directions of such developments, making it important for AI researchers, developers, social scientists, and policymakers to balance the imperative to innovate with mechanisms to ensure that AI’s economic and social benefits are broadly shared across society. If society approaches these technologies primarily with fear and suspicion, missteps that slow AI’s development or drive it underground will result, impeding important work on ensuring the safety and reliability of AI technologies. On the other hand, if society approaches AI with a more open mind, the technologies emerging from the field could profoundly transform society for the better in the coming decades.