Abstract

Artificial intelligence (AI) is a relatively new branch of computer science. A tremendous amount of effort has been put into research associated with understanding biological systems, abstracting key principles of intelligent behaviour, and developing practical applications of AI since the year 2000. The ultimate goal of this science is reaching "Strong AI". However, humanity must be careful with creation of artificial intelligence that is similar and, perhaps, to some extent, identical to the human being.

The history of artificial intelligence has its roots in Ancient Egypt where around 800 B. C. a statue of Amun that could move its hands up and down was constructed [16]. In the twentieth century with the creation of the first electronic computer humanity reached the level where creation of machine intelligence began to look plausible [19, p. 488-489]. As a result scientists and industrial researchers from all over the world started heading towards a new goal: to build a machine in such a way that the results of its work cannot be distinguished from creations of a human mind. It is a challenging goal indeed. Computer scientists and mathematicians working in the newly created field discovered that they were facing new problems that were well beyond the scope of what was known then as "traditional informatics". We are living in the 21st century and this goal has not been reached yet. One might argue, however, that numerous breakthroughs and dramatic advances have taken place during the last 60 years. Especially considering AI-related news were headlines in newspapers in the first decade of our century. What exactly has been achieved and where researches failed to succeed in the field of artificial intelligence since the year 2000? This essay makes an attempt to answer this question.

Before discussing the history of artificial intelligence one should have a clear understanding of what AI actually means. Few specialists can give a solid definition of the term "intelligence", and it is not a surprise that the meaning of "artificial intelligence" has been given various, occasionally to a certain degree contradictory, outlines. The inventor of the term John McCarthy said the following words about AI: "The goal of AI is to develop machines that behave as though they were intelligent." [6]. Years passed after this definition was created and people are now able to prove that it is possible to create relatively simple, one might say "useless", systems that can behave as if they were intelligent, when in fact they are not. Furthermore, the main reason behind endeavours of most people that are involved in the field of AI is reaching the ultimate goal of creating a machine that is better than us humans. Another widely-accepted interpretation of AI was given by Elaine Rich: "The study of how to make computers do things at which, at the moment, people are better" [21]. This quote holds a significantly more time-proof definition of AI, since even in 100 years from now there are likely to be tasks, in execution of which humans will be outperforming robots. However, one might argue that this definition of AI will also become outdated at the point when humans will succeed in creating strong AI. Calling AI "strong" implies that the intelligence of a machine is sufficient enough for it to perform any mental task that a human can do [14].

It is not clear if humans will ever succeed in the creation of strong AI. Meanwhile, as can be seen from aforementioned examples, identifying behaviour of a machine as a sign of intelligence is not a routine task. The opinion of [15, p. 1-8] is that currently there are three main areas of research that are associated with development of artificial intelligence: understanding biological systems, abstracting key principles of intelligent behaviour, and developing practical applications. Let us take a look at victories and fiascos of AI in these areas that took place since the year 2000.

Understanding biological systems

Questions about understanding how humans and animals work have been in the minds of the greatest thinkers since the day people learned how to think. It is indeed a great question: "How does life work?" Not far away from it is another question: "How can we mimic life?" Computer scientists have been working on simulating life for many years. [15, p. 1-8] correctly argues that this research area relies heavily on the "Understanding by building" principle where it is much easier to build something and perform tests until a satisfactory result has been achieved than to perform an extensive theoretical research on the topic. This approach works very well in the field of AI and as a result there have been numerous advancements in the area of artificial life (ALife), and especially in its subfield called “soft ALife” [11]. This field separated from theoretical AI in 1980s when the first conference in artificial life "ALife I" (1989, Los Angeles) took place [1]. Research in this field focuses on understanding how systems related to life operate by performing simulations with computers, robots or by means of biochemistry.

The most interesting biological system from AI's point of view is the human's brain. There have been multiple attempts to build an artificial brain. Reference to [7] reveals that we may be less than ten years away from being able to successfully create an artificial brain. In 2011 a team of IBM researches led by Dharmendra Modha announced that they developed a new chip that can store and process information in close parallel - the way our brain does it - and it was built based on artificial “neurons” and “synapses”, numbering in hundreds and thousands, respectively. A chip’s main feature is its ability to change its state to process new information that it wasn't specifically programmed to expect. Furthermore, by being able to perform close parallel information processing this chip - provided that it will come into production - will entirely change the way people think about methods of computing. Moreover, it was announced by the team that made this breakthrough that the chip would soon be capable of “forming, strengthening, and breaking connections on the fly” while using a thousand times less power than conventional computers [9]. I believe that this chip will revolutionise the field of AI in the near future.

To extend on that, the team from IBM has an impressive list of successful simulations of mental activity; by using supercomputers they have been able to go through processes in a cat's cerebral cortex, as well as one percent of a human cerebral cortex, 40 percent of a mouse's brain, and a rat's full brain. The ultimate goal of IBM is to build a computer that has 10 billion neurons and 100 trillion synapses, i.e. a system that is as powerful as the human brain. This idea, on the other hand, initiates debates about morality of aiming to achieve such goal. One may think about a situation where the military starts using the artificial brain and all its power to plan and execute military actions, where humans might die. Given this as a real possibility, the question of responsibility and blame for war-related deaths and killing naturally suggests itself. Secondly, will we ever be able to give full rights to make life-defining decisions to AI? No precedents showing that it may be dangerous have been recorded during the last dozen of years but it is unclear whether it is advisable to give full rights for decision-making to computers. Nevertheless, if the goal of creating an artificial brain, which is at least as powerful as a human brain, were achieved, it would be the “Holy Grail” of artificial intelligence and a nice way to end 60 years of extensive research [10].

Abstracting key principles of intelligent behaviour

Our world is rather complex and no human being is capable of comprehending all of its rules and laws in the form of one single entity. Lungarella et al. makes it clear that on the current stage of development of AI, when it is less powerful than our brains, principles of intelligent behaviour need to be abstracted and simplified in order to be applied in building artificial intelligence [15]. This can be viewed as the beginning of a “theory of intelligence.”

To start, artificial intelligence needs to be aware of humans around it. For example, if a robot is set to be built to accomplish a certain task that involves interactions with people, its logic must be error-free and safe for people around it. Machines today do not have means by which their behaviour can be adjusted as our goals and needs change. AI needs to know how to learn new abilities and improve its skills by observing or interacting with us, people [3, p.5].

A considerable amount of research has been done in the field of nonverbal communication in teams consisting of both humans and robots. Breazeal perceptively states that humans are the most socially advanced species on Earth [3]. Hence, mimicking and interpreting social behaviour of Homo Sapiens is not an easy task for a robot. Nevertheless, a number of successful research projects saw light in the past few years. Probably, one the most well-known accomplishments of AI is the "Kismet" robot developed at MIT. When it was shown to public, it became a major success, and it is on exhibit in the Museum of MIT. Breazeal-led team continued their work after this breakthrough and shortly after they presented a social robot called "Leonardo". This expressive humanoid robot, designed for social interaction and communication to support teamwork and social learning, was finished by the MIT-based team in 2002. By examining and abstracting data received from its human companions, the robot is capable of learning from humans, interacting with them and collaborating during solving common problems. By incorporating deictic gestures into the language of object reference, "Leonardo", along with another creation of the MIT Media lab called "Robonaut", provided first steps towards full integration of natural interactions between robots and humans. Results received after interpretation of testing of the aforementioned robots could potentially be applied to the design of software that is socially intelligent and aware of needs of its human users. One should keep in mind, however, that there are crucial variations between the world we, humans, live in and the virtual world populated by computer agents. These differences must be taken into account when human-like features are given to computers [4, p.297].

Developing practical applications

Mimicking brains and making robots aware that they exist in the world that belongs to humans can produce results and end in a few Nobel prize awards in a long run, but to actually build a software system or a robot that simplifies human lives or makes this world a better place is, arguably, a lot more rewarding. As a result, our world has already seen an enormous number of milestones and significant breakthroughs in the field of developing practical applications of AI in this millennium. There are a number of areas of our day-to- day lives where AI makes our existence easier and more efficient: Internet search, receiving music recommendations based on our interests, practising chess tactics, carrying cargo on the battlefield, to name just a few. To support this idea, let us consider a few memorable examples of practical applications of AI that improve lives of humans.

Firstly, consider chess, the "Game of Kings." Ever since the first game of it was played, it has been considered as one of the most challenging and mentally-exhausting games in the world. Thomas Henry Huxley once said: “The chess-board is the world; the pieces are the phenomena of the universe; the rules of the game are what we call the laws of Nature”. Not surprisingly, creation of AI that is capable of beating the best chess players has been considered as one of the most desired milestones in the world of Computer Science. Claude Shannon, one of inventors of information theory, rightfully concluded in 1950: “... a machine calculating one variation each millionth of a second would require over 10 to the 95th power years to decide on its first move (in chess)" [5]. He was one of the first scientists who suggested that achieving success in building a top-class chess program requires a non-trivial approach, which utilises means of AI. In 1997 IBM- built “Deep Blue” supercomputer won a match against a world champion Garry Kasparov; the defeated Kasparov protested afterwards where he accused “Deep Blue” of cheating. However, no rematch has ever been played [22]. The world was shocked on the day when finally, after decades of designing and developing chess computers, a chessmaster was beaten by a machine. However, some were still not satisfied, and new even more powerful computers were built. After 2000, with advancements in computations and IT, possibility of having computer programs which are at least as powerful as “Deep Blue” that can be run by conventional PCs emerged. Currently, “Rybka” knowledge base chess engine, development of which started in 2003, is considered as the best chess program in the world. Not surprisingly, though, few chessmasters agree to have official tournaments against AI these days. I am myself a ranked chess player and, knowing how demanding chess can be, I fully appreciate the first victory of “Deep Blue” and later success of AI in tournaments, and I personally believe that what was achieved by AI scientists in the area of chess can be called as one of the biggest breakthroughs in science of the last years.

Secondly, another notable example of a computer that managed to perform better than humans is "Watson", AI that defeated humans on the popular American quiz show “Jeopardy!”. "Watson" is an artificial intelligence computer system that can answer questions given to it in natural language. In 2011 "Watson" won a game against two best players in “Jeopardy!” and it can be considered as a triumph of natural language processing and machine learning. According to [20], Watson is capable of processing 500 GB per second, which is the equivalent of reading a million books per second, and it is well beyond powers of humans. However, a creator of the "Chinese room" experiment, philosopher John Searle argued that "Watson" cannot actually think, since even though it knows how to read millions of books a second, it does not know what reads about [23].

A sentence or two should also be mentioned about DARPA - a US agency responsible for development of new technologies for use by the military. Furthermore, the most notable DARPA-sponsored creation is “BigDog”, a legged robot developed by Boston Dynamics and presented in 2005. The ultimate goal of the project is a fully-autonomous legged vehicle that can perform better on rough terrain than any existing wheeled or tracked vehicle. In its current iteration it already shows impressive results by adjusting its body height to comply with the local terrain and by changing footfall positioning to compensate for orientation of the body and ground plane relative to gravity. On each of its legs the robot has four hydraulic actuators, which gives it 16 degrees of freedom (the number of independent parameters that define configuration of the mechanical system) as well as one passive degree on each leg. Turning to [25], one may see that just on October 1, 2012 a press release was given about a new development on “BigDog” that is capable of carrying heavy cargo on battlefields, thus helping soldiers. It is just one of possible practical uses of “BigDog” and we should expect seeing more in the near future [18].

A Japanese company Honda went even further with their humanoid robot ASIMO, which was introduced in 2000 [12]. This machine can walk, run and even, according to [24], conduct a symphonic orchestra. It has also been serving a noble mission of popularisation of science through public shows performed in front of children. The machine can interact with humans by identifying moving objects, learning its surrounding environment, distinguishing sounds and faces. ASIMO, since it mimics moves with two legs, has an even more complicated mechanical system than that of the aforementioned “BigDog”, it has a total of 34 degrees of freedom [8].

Additionally, researchers at Google X, a secretive part of Google that experiments with ambitious future technologies, have been working on projects dealing with cutting-edge ideas in AI. Google claims that the ultimate goal for the company is for Google Search to become AI-complete. There is still some work to be done by its scientists before that goal could be reached. Meanwhile, however, during recent years newspapers at least a few times a year would publishing cover page stories telling about new surprises introduced to the public by Google. One of them was a driverless Toyota Prius, which can now be seen on streets of three states in the US. Then, in 2012 Google announced about their own attempt in creation of the artificial brain [17]. Google’s version utilises 16000 CPUs, which are combined into a “neural network”; this network was tested on finding images of cats in YouTube videos. I think that the experiment was very successful, since the network was able to respond strongly to pictures of cats after one week of processing data. It must be mentioned, that the system was never told what the cat is, nor what the cat is not.

Conclusion

To summarize, a lot of progress was done since 1940s and especially during last 12 years: chess is no longer dominated by humans, pets with metal bones make retired people smile again and detection of a face of a person close to your heart from millions of hours of CCTV videos is no more an issue.

Artificial intelligence is an extremely broad subject and it is impossible to mention all the advances made since the year 2000 from all fields of it in one paper. A list of other aspects of AI, where significant results were achieved may include such topics: machine learning (e.g. concert recommendations based on one's Facebook profile), AI in games ("NERO", where one can train an army of intelligent robots), exploration of extreme environments (Mars rover "Opportunity"), and many more. It should also be noted that information is scarce regarding failures in the field of artificial intelligence. Companies and research institutes hide their lack of success from the public. Further, whenever a new AI project ends with unexpected results, they are interpreted as new information and help with future research.

Lastly, there are many ethical questions that arise when they think about artificial intelligence. Seeing the project of AI as ultimately concerned with creating reliable, intelligent counterparts and companions for humans is problematic. We are not talking about domestication of animals, but, instead, about creation of artificial intelligence that is similar and, perhaps, to some extant identical to the human being. This story is not new: God created man in his own image. What happened next, one may ask? "God is dead. God remains dead. And we have killed him." as was said by Friedrich Nietzsche. We are now in the middle of creating a new entity in the equation of life, and this kind of process should not be left unfinished [13]. I do believe that strong AI is plausible and one day we will no longer be the only thinking and per haps feeling beings on this planet. However, people should be careful not to let the words "Humans are dead. Humans remain dead.

And we have killed them.” be transmitted using binary code in the future. Technological singularity may sound like something coming straight out of sci-fi books, but it should be considered as a possibility. An interesting analogy was given in regards to domination of greater-than-human superintelligence where singularity caused by it is compared to event horizon found in black holes: laws of nature as we know them change beyond it and there is no turning back.

Bibliography

1. Bobrovski S. Эволюция и искусственная жизнь (in Russian, translation: "Evolution and artificial life") // Homepage of PCWeek. URL: http://www.pcweek.ru/themes/detail.php?ID=69535.

2. Breazeal C. Socially intelligent robots // Interactions. – 2005. - № 12.

3. Breazeal C. Intelligent robots and autonomous agents. – Cambridge, 2002.

4. Brooks A.G. Working with robots and objects: Revisiting deictic reference for achieving spatial common ground // Proceedings of Human-Robot Interaction. – 2006. pp. 297.

5. Ditlea S. & Lunch Group. Digital deli: the comprehensive, user-lovable menu of computer lore, culture, lifestyles, and fancy. - Workman Pub, 1984.

6. Ertel W. Introduction to artificial intelligence. – London, 2011.

7. Fildes J. Artificial brain '10 years away'. URL: http://news.bbc.co.uk/1/hi/technology/8164060.stm.

8. Ford J. Two legs good. - The Engineer, 2011.

9. Gonazales R. IBM’s “neurosynaptic” chips are the closest thing to a synthetic brain yet. URL: http://io9.com/5832085/ibms-neurosynaptic-chips-are-the-closest-thing-to-a-synthetic-brain-yet.

10. Kananen M. Will Computers Outsmart Human Brains? // Homepage of The Be-

ginner. URL: http://www.thebeginner.eu/technology/574-will-computers-outsmart-human-brains.

11. Komosinski M., Adamatzky A. Artificial life models in software, 2nd edn. – London, 2009.

12. Kornblum J. Meet Honda's ASIMO, a helpful Mr. Roboto. - USA Today, 2000.

13. Kunafin M. Person and Society: Perspectives for the XXIst Century // The Personality and Artificial Intellect: The tendencies of XXIst century. – Braga, 2005.

14. Kurzweil R. The singularity is near : when humans transcend biology, Duckworth. – London, 2005.

15. Lungarella M., Iida F., Bongard J.C. & Pfeifer R. AI in the 21st century: with historical reflections // 50 Years of Artificial Intelligence: Essays Dedicated to the 50th Anniversary of Artificial Intelligence. – Festschrift, 2007.

16. Nakate S. History of Artificial Intelligence. URL: http://www.buzzle.com/articles/history-of-artificial-intelligence.html.

17. Nichols S. Google claims artificial intelligence breakthrough with cat-spotting neural network. URL: http://www.v3.co.uk/v3-uk/news/2187313/google-touts-content-recognising-neuro-network.

18. Raibert M. BigDog, the Rough-Terrain Quadruped Robot // Proceedings of the 17th IFAC World Congress. - The International Federation of Automatic Control, 2008.

19. Ralston A. & Reilly E.D. Encyclopedia of computer science. New York, 1976.

20. Rennie J. How IBM’s Watson Computer Excels at Jeopardy! URL: http://blogs.plos.org/retort/2011/02/14/how-ibm%E2%80%99s-watson-computer-will-excel-at-jeopardy/

21. Rich E. Artificial intelligence. New York, 1991.

22. Saletan W. Chess Bump: The triumphant teamwork of humans and computers. URL: http://www.slate.com/articles/health_and_science/human_nature/2007/05/chess_bump.html

23. Searle J. Watson Doesn't Know It Won on 'Jeopardy!' URL: http://online.wsj.com/article/SB10001424052748703407304576154313126987674.html.

24. Van Buskirk E. Honda Robot Will Conduct Detroit Symphony. Wired, 2008.

25. Williams M. DARPA begins testing robotic mule for battlefields // Homepage of Computer World]. URL: http://www.computerworld.com/s/article/9231883/DARPA_begins_testing_robotic_mule_for_ battlefields?taxonomyId=12.