

PROGRESS OF ARTIFICIAL INTELLIGENCE (AI)

Titu-Marius I. Băjenescu

13, Chemin de Riant-Coin, 1093 La Conversion, Switzerland

Corresponding author: tmbajenesco@gmail.com

Received: June, 03, 2018

Accepted: June, 24, 2018

Abstract: In the past few years, systems like Siri and Google Now opened our minds to the idea that we don't have to be tethered to a laptop to have seamless interaction with information. In this model, AIs will move from speech recognition to natural language interaction, to natural language generation, and eventually to an ability to write as well as receive information. Machine learning is about man assisting computers. Deep learning is about systems beginning to progress and learn on their own. Historically, systems have always been trained, they've been programmed. And, over time, the programming languages changed. As we move into the area of AI and cognitive computing, we're exploring the ability of computers to do more unaided / unassisted learning.

AI definition: The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages. In other words, artificial intelligence is the ability of a computer to understand what you're asking and then infer the best possible answer from all the available evidence. Soon AI will become the most important human collaboration tool ever created, amplifying our abilities and providing a simple user interface to all exponential technologies. Ultimately, it's helping us speed toward a world of abundance. It's only been in the last few years that we've seen a fundamental transformation in this technology.

Keywords: *Deep learning, automatic learning, robot, DeepStack, technology*

1. Introduction

Below (figure 1) is a brief description of the achievements and their circumstances. Some milestones represent significant progress towards human performance and others represent super-human performance achievements.

Skin Cancer Classification

In a 2017 *Nature* article, Esteva et al. [5] describe an AI system trained on a data set of 129,450 clinical images of 2,032 different diseases and compare its diagnostic performance against 21 board-certified dermatologists. They find the AI system capable of classifying skin cancer at a level of competence comparable to the dermatologists.

Speech Recognition on Switchboard

In 2017, Microsoft and IBM both achieved performance within close range of "human-parity" speech recognition in the limited Switchboard domain.

Poker

In January 2017, a program from CMU called Libratus defeated four top human players in a tournament of 120,000 games of two-player, heads up, no-limit Texas Hold'em. In February 2017, a program from the University of Alberta called DeepStack played a group of 11 professional players more than 3,000 games each. DeepStack won enough poker games to prove the statistical significance of its skill over the professionals.



Figure 1 Milestones of artificial intelligence (after [14])

Pac-Man

Maluuba, a deep learning team acquired by Microsoft, created an AI system that learned how to reach the game's maximum point value of 999,900 on Atari 2600.

2. Reliability and accuracy of the sensory information

Autonomous mobile robots are increasingly being used in our environments to perform concrete complex tasks (Figure 2). Such robot agents need to perform a set of computationally intensive functions, in order to be able to perceive, to reason about, and to act in their surroundings. The robot's actuation, whether it is to manipulate objects or just to navigate throughout its environment, is based on the information provided by its sensors. The success of this actuation depends on the reliability and accuracy of the sensory information. One of the richest and most complex sensory information sources is the visual information as captured by cameras. Many of the modern robots are equipped with cameras to observe the world and detect the objects relevant to accomplish the specific tasks. Robots perform their tasks in a closed loop between perceptions and control (behaviours). Control includes the computation of the next action towards self-localization, navigation, cooperation, or object manipulation. In dynamic environments in particular, the detection of relevant objects to the actuation has to be done in real time, i.e., with no delay in the perception / control loop.




Figure 2 Import of industrial robots units, into North America and globally, between 2000 and 2015 (after [14])

3. Can machines think?

The machines know how to play, but can they think? Recently, however, computers have regularly beaten professional gamers in a growing number of games. Last year, AlphaGo beat a professional go player, for the first time in human history (which reminds us of the defeat in chess by world champion Garry Kasparov, defeated in 1997 by an IBM computer). But how do you "teach" a computer to play such games? Let's take the simple

game of the spy game, and try to understand the process. How do we humans manage to play the spy game? For most of us, there are some informal rules that we learned as children. One of these rules could be: when you start, mark the centre square. Another: when an opponent has two squares in a row, mark the third. With these two very simple rules, you can already program a computer to play morpion (although rudimentarily, of course). It only takes eight simple rules of this kind to make sure that the computer doesn't lose any more parts!

There is a completely different approach: we can also make a list of all possible game situations in the spy game, with all the best moves in each of these situations. Then we could tell the computer to look for the best movement in this very long list, and for each movement. Drawing up this list would be tedious, but feasible.

However, learning requires an opponent, usually a human. And the better the opponent, the better the computer becomes, so that eventually the best human player will lose at some point. This is exactly what we have seen in the recent past, with chess, go and poker.

But the most interesting aspect of this approach is that the new machines are starting to make new movements, which have never been made by a human being. That's why we can start to wonder: how far will the machines go?

What if the computer was programmed to "learn" and memorize all the situations by itself? This is achievable, and today it is called *machine learning*. For example, suppose we program a computer to have only one basic morpion rule, and tell it to play randomly until the game is over, and record each step of the game. In the very first games, the computer would be immediately beaten. But thanks to the "bad moves" he has recorded, he will learn which movement he should not repeat in the future and, gradually, he will end up finding a list of "right moves". From then on, the more the machine plays, the more it will "learn" and win!

We have been talking about robots⁶, golem, and artificial beings in science fiction films for a long time. Artificial intelligence actually exists since we have been thinking about how to use machines to think, to calculate for us. It is part of a great historical continuity. What's new is the presence of self-learning machines. A real qualitative leap forward. We have the impression that if these machines can learn, we can lose control.

Man gradually tries to transform himself into an artefact, an immortal thing. His absolute fantasy is to transfer self-consciousness into an artefact. It is the long tendency of human history to want to escape death. To date, I see many more advantages than disadvantages in artificial intelligence. Are we creating a race of humans that will surpass us? That's possible. It is necessary to protect oneself from it and to avoid that human beings find themselves in a situation of dependence on machines. Some believe that man will be overtaken by machines in 2030 or even 2045. We don't know that yet. I think we're still a long way from it.

⁶ The robots [from Czech, from *robot* 'forced labour'. The term was coined in K. Čapek's play R.U.R 'Rossum's Universal Robots' (1920)] are typically mobile, unlike their factory counterparts, and must operate reliably in dynamic and complex environments. Application areas for field robots include agriculture, mining, construction, forestry, cargo handling, and a myriad of other domains. Field robots may operate on the ground (either terrestrially or on other planetary bodies), underground, underwater, in the air, or in space. Service robots work closely with humans, including the elderly and the sick, to assist them with their daily lives and activities.

4. Ai and the reasoning of the human being

It would be a philosophical revolution if artificial intelligence were to reach the mode of reasoning of the human being. But here too, we are far from it. The man has more than a billion neurons and 10^{12} synapses. We may be able to analyse 100,000 neurons today. But we're still a long way from the billion. The exact functioning of the brain is not yet understood. Apart from artificial intelligence, it is therefore important to develop human intelligence. There are billions of human intelligences out there. There are billions of people who have not been to school enough. I don't think there are exceptional brains on one side and stupid people on the other. Not all of them have grown up in social conditions that have allowed them to develop favourably. That is why, before developing AI, human intelligence must be developed through education and creativity.

It is also necessary to develop other forms of intelligence, artistic, emotional, emotional and amorous. This presupposes work on the process of releasing human intelligence. Finally, a huge amount of networking of real intelligences is also necessary. For if humanity had a level of education equal to that of the most advanced people and these intelligences were networked, no artificial intelligence could compete with them.

A permanent watch must be set up from the outset. Transparency of research and rules are also needed to ensure that artificial intelligence does not develop to the detriment of the human species. These are the famous three laws of robotics known as Asimov Laws, i. e. do not harm humanity, obey the orders of man and act in accordance with these two precepts. But even that is not enough, because machines could harm the environment.

It is therefore important to maintain the possibility of literally killing artificial intelligence. But here too, it is very delicate, because if the AI understands that man has the means to destroy it, it could, in order to guard against human intervention, invent languages, as it has begun to do, that man does not understand.

We will need a moratorium if we cannot implement the rules mentioned above. Artificial intelligence is not only developed in public places, but also in highly confidential military laboratories. It is therefore difficult to control everything.

General and mathematical culture remains fundamental, as well as music and literature. It will be necessary to learn curiosity, to learn to learn. Math remains the basis of everything and can quickly make scientific knowledge obsolete.

The risk of total disconnection exists. Jobs will be lost because of artificial intelligence. If this allows the human being to devote himself to more interesting activities, much better. But there is a real danger, on the one hand, of seeing people with a higher intellectual quotient and, on the other hand, people who would be forced to survive on a universal minimum income and consume all kinds of real and virtual drugs.

5. Face recognition

What has really changed in the world over the last few decades: We suddenly have a technology that is "smart enough" to recognize a face, even better than a human being. This has a great influence on how we can organise society.

Face recognition has attracted great interest as an important biometric technique and has wide applications in information security, law enforcement and surveillance, smart cards, access control, etc. Among the numerous methods that have been proposed [1–3] for face recognition, *linear discriminant analysis* (LDA) is widely used as a dimension reduction

technique because of its superiority in considering the underlying class structure of input data over other techniques such as *principal component analysis* (PCA) [2, 4].

Image recognition has exploded over the last few years. Each Facebook and Google photos have tens of billions of images on their platform. With this dataset, they - and many others - are developing technologies that go beyond facial recognition providing algorithms that can tell you what is in the image: a boat, plane, car, cat, dog, and so on. The crazy part is that the algorithms are better than humans at recognizing images. The implications are enormous. Imagine an AI able to examine an X-ray or CAT scan or MRI to report what looks abnormal.

As AI begins to impact every industry and every profession, there is a response where schools and universities are ramping up their AI and machine learning curriculum. IBM, for example, is working with over 150 partners to present both business and technology-oriented students with cognitive computing curricula.

6. Automatic learning

Automatic learning has become increasingly important due to the rapid growth of the amount of data available. In the year 2000, the total amount of information on the Web varied somewhere between 25 and 50 terabytes [6]. By 2005, the total size was approximately 600 terabytes [7]. Nowadays, the total amount of information is almost incalculable. This unrestrainedly growth of data opens the way for new applications of machine learning. Automatic data analyzers are needed since a human, even an expert, cannot look at a “very large” data set and plausibly find a good solution for a given problem based on those data. In this situation, new challenges are raised regarding the scalability and efficiency of learning algorithms with respect to computational and memory resources. Practically, all existing implementations of algorithms operate with the training set entirely in main memory. If the computational complexity of the algorithm exceeds the main memory then the algorithm will not scale well, will not be able to process the whole training data set or will be unfeasible to run due to time or memory restrictions. However, increasing the size of the training set of learning algorithms often increases the accuracy achieved by classification models [8], and thus, in order to handle “very large” data sets, a new and active research field emerges, large-scale learning [9, 10]. It intends to develop efficient and scalable algorithms with regard to accuracy and to requirements of computation (memory, time and communication needs). Large-scale learning has received considerable attention in the recent years and many successful techniques have been proposed and implemented [11 – 13].

Distributed learning seems essential in order to provide solutions for learning from both “very large” data sets (large scale learning) and naturally distributed data sets. It provides a learning scalable solution since the growing volume of data may be offset by increasing the number of learning sites. Moreover, distributed learning avoids the necessity of gathering data into a single workstation for central processing, saving time and money. Despite these clear advantages, new problems arise when dealing with distributed learning as, for example, the influence on accuracy of the heterogeneity of data among the partitions or the need to preserve privacy of data among partitions. Therefore, this is already an open line of research that will need to face these new challenges.

Deep learning first transformed speech recognition, then computer vision. Today, *Natural Language Processing* (NLP⁷) and robotics are also undergoing similar revolutions.

AI publication numbers are of interest (Figure 3) and are used generally in scientometrics and research statistics. It is impossible to combine a system based on giving values to individual publications with a study of the growth rate of AI science.




Figure 3 Annual published papers in recent years concerning AI (after [14])




Figure 4 The five steps in NLP (after [17])

7. Computers versus humans

Obviously, computers are vastly superior to humans in certain tasks; however, the competence of AI systems becomes more difficult to assess when dealing with more general tasks like answering questions, playing games, and making medical diagnoses.

With so much excitement about progress in artificial intelligence, you may wonder why intelligent machines aren't already running our lives.

Key advances have the capacity to dazzle the public, policymakers, and investors into believing that human-level machine intelligence may be just around the corner. But a new report [14], which tries to gauge actual progress being made, attests that this is far from true. The findings may help inform the discussion over how AI will affect the economy and jobs in the coming years.

Tasks for AI systems are often framed in narrow contexts for the sake of making progress on a specific problem or application. While machines may exhibit stellar performance on a certain task, performance may degrade dramatically if the task is modified even slightly. For example, a human who can read Chinese characters would likely understand Chinese speech, know something about Chinese culture and even make good recommendations at Chinese restaurants. In contrast, very different AI systems would be needed for each of these tasks.

Despite the difficulty of comparing human and AI systems, it is interesting to catalogue credible claims that computers have reached or exceeded human-level performance. Still, it is important to remember that these achievements say nothing about

⁷ NLP refers to AI method of communicating with an intelligent system using a natural language such as English. Processing of Natural Language is required when you want an intelligent system like robot to perform as per your instructions, when you want to hear decision from a dialogue based clinical expert system, etc. There are general five steps in NLP (Figure 4). Today, NLP's deep learning transformation is well underway; this will lead to a flourishing of new applications (such as chatbots). Deep Learning in robotics is also gaining significant momentum, and this will lead to many new applications (such as new manufacturing capabilities).

the ability of these systems to generalize. Games provide a relatively simple, controlled, experimental environment and so are often used for AI research.

8. Anticipated top ai breakthroughs

Alan Turing created the Turing test over half a century ago as a way to determine a machine's ability to exhibit intelligent behaviour indistinguishable from that of a human. Loosely, if an artificial system passed the Turing test, it could be considered AI. The specialists believe that for all practical purposes, these systems will pass the Turing test in the next three-year period. Perhaps more importantly, if it does, this event will accelerate the conversation about the proper use of these technologies and their applications.

All five human senses (yes, including taste, smell and touch) will become part of the normal computing experience. AIs will begin to sense and use all five senses. The sense of touch, smell, and hearing will become prominent in the use of AI. It will begin to process all that additional incremental information. When applied to our computing experience, we will engage in a much more intuitive and natural ecosystem that appeals to all of our senses.

Solving big problems: detect and deter terrorism, manage global climate change. AI will help solve some of society's most daunting challenges. We're already seeing this technology being deployed in governments to assist in the understanding and pre-emptive discovery of terrorist activity. We'll see revolutions in how we manage climate change, redesign and democratize education, make scientific discoveries, leverage energy resources, and develop solutions to difficult problems.

Leverage all health data (genomic, phenotypic, and social) to redefine the practice of medicine. AI's effect on healthcare will be far more pervasive and far quicker than anyone anticipates. Even today, AI / machine learning is being used in oncology to identify optimal treatment patterns. But it goes far beyond this. AI is being used to match clinical trials with patients, drive robotic surgeons, read radiological findings and analyze genomic sequences.

AI will be woven into the very fabric of our lives -physically and virtually. Ultimately, during the AI revolution taking place in the next three years, AIs will be integrated into everything around us, combining sensors and networks and making all systems "smart." AIs will push forward the ideas of transparency, of seamless interaction with devices and information, making everything personalized and easy to use. We'll be able to harness that sensor data and put it into an actionable form, at the moment when we need to make a decision.

9. Instead of a conclusion

Not all aspects of AI in which recent progress has been made can be covered. Progress is typically tracked consistently when good progress has been made. For some areas there are not clear standardized benchmarks (e.g. dialogue systems, planning, and continuous control in robotics). In other areas it is hard to measure performance when there has not been significant progress, like in commonsense reasoning. And still, other areas are waiting to be tracked but we simply have not had the opportunity to collect the data (e.g. recommender systems, standardized testing).

References

1. Turk, M., Pentland A. Eigenfaces for recognition. In: *Journal of Cognitive Neuroscience* 1991; 3(1):71–86.
2. Belhumeur. P., Hespanha, J., Kriegman, D. Eigenfaces vs. fisherfaces: recognition using class specific linear projection. In: *IEEE Transactions on Pattern Recognition and Machine Intelligence* 1997; 19:711–720.
3. Abate A., Nappi M., Riccio D., Sabatino G., “2D and 3D face recognition: a survey,” *Pattern Recognition Letters* 2007; 28:1885–1906.
4. Martnez A., Kak A., “PCA versus LDA,” *IEEE Trans. on Pattern Recognition and Machine Intelligence* 2001; 23(2): 228–233.
5. Esteva A., Kuprel B., Novoa R. A., Ko J., Swetter S. M., Blau H. M., Thrun S., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*. 2017 Feb. 2;542(7639), 115-118, doi: 10.1038/nature21056. Epub 2017 Jan 25.
6. School of Information and Management and Systems. How much information? <http://www2.sims.berkeley.edu/research/projects/how-much-info/internet.html> (2000).
7. D-Lib Magazine, “A research library based on the historical collections of the Internet Archive,” <http://www.dlib.org/dlib/february06/arms/02arms.html> (2006).
8. Catlett, J.: “Megainduction: machine learning on very large databases” PhD thesis, *School of Computer Science, University of Technology, Sydney, Australia* (1991).
9. Bottou, L., Bousquet, O., “The tradeoffs of large scale learning,” *Adv. Neural Inf. Process. Syst.* 20, 161–168 (2008).
10. Sonnenburg, S., Ratsch, G., Rieck, K., “Large scale learning with string kernels,” in: Bottou, L., Chapelle, O., DeCoste, D., Weston, J. (eds.) *Large Scale Kernel Machines*, pp. 73–104. MIT Press, Cambridge (2007).
11. Moretti, C., Steinhäuser, K., Thain, D., Chawla, N.V., “Scaling up classifiers to cloud computers,” In: *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM)*, pp. 472–481 (2008).
12. Krishnan, S., Bhattacharyya, C., Hariharan, R., “A randomized algorithm for large scale support vector learning,” In: *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pp. 793–800 (2008).
13. Raina, R., Madhavan, A., Ng, A.Y., “Large-scale deep unsupervised learning using graphics processors,” in: *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, pp. 873–880 (2009).
14. *AI index*, November 2017, <https://aiindex.org/2017-report.pdf>
15. AI 2017: “Advances in artificial intelligence,” *Proc. of 30th Australian Joint Conference*, Melbourne, VIC, Australia, August 19-20, 2017.
16. Testa, J. (2009). *Regional content expansion in web of science: Opening borders to exploration* (pp. 1–3). Retrieved December 29, 2017 from <https://globalhighered.wordpress.com/2009/01/15/regional-content-expansion-in-web-of-science/>
17. *AI - Natural Language Processing*, https://www.tutorialspoint.com/artificial_intelligence/artificial_intelligence_natural_language_processing.htm