Artificial Intelligence: Challenges for the future

By **Michalis VAZIRGIANNIS**
Professor at DaSciM, LIX [(1)], Ecole Polytechnique

"The means of production are only decommissioned if a more productive replacement is found. In other words, we can only talk about the abolition of slavery if only the goods can be produced without the help of slaves: Daedalus or the tripods of Hephaestus, about whom the poet says that they automatically appeared in the assembly of the gods, and if the shuttles weaved on their own and the guitar strings played on their own, then the architects would not need servants neither the free people slaves."

Politics, Aristotle.

Introduction

Artificial Intelligence (AI) is a term coined by Prof. J. McCarthy, one of the "founding fathers" of the area in 1955. AI may assist, augment or eventually automate processes. Colloquially AI can be used to describe machines (or computers) that mimic "cognitive" functions associated with the human mind, such as "learning" and "problem solving" [6] and can be used to tackle many of the problems usually solved by humans such as classification, summarisation, translation, diagnosis, grouping, finding outliers. AI algorithms have already found applications in most domains of human activity including finance, national security, health care, justice, transportation, smart cities, industrial processes education, etc.

AI algorithms function effectively they need sufficient data in order to be trained but as well, especially recently with the advent of deep learning, significant computing power. Hopefully in the last decades data have become abundant – the term *big data* has been coined to describe the volume, speed of production and heterogeneity of the the data produced by socioeconomic processes [4].

As for computing resources there are different aspects. The *storage capacity* that has been increasing almost exponentially in the last decades [5] enables storage of unprecedented amount of data that correspond to practically any activity that can produce data – i.e. user behaviour at large, image and video, sensors' readings from different processes, loads of financial transactions, etc. Here, the role of cloud storage (i.e. distributed remote storage of information with availability guarantees) and software-as-service are significant as they alleviate small businesses and governments from the cost of purchase and maintenance of complex specialized infrastructures.

As regards the computational power the CPU capacity increases continuously in the last decades, converging though to a plateau. On the other hand, the advent of DL required the development and wide deployment of GPUs. This is a critical infrastructure which is currently the heart of resources needed for the large-scale computations.

(1) http://www.lix.polytechnique.fr/dascim/

The main stream AI methods are emanating from the field of *machine learning* where the scenario is that there is a training set (i.e. a set of emails labelled as "spam") where the data points (i.e. emails) are embedded in a feature space (i.e. the union of the words in the email set). The aim it to fit a mathematical function (i.e. regression) or other structure (i.e. decision tree) that given data of unknown classification produces a prediction (i.e. "spam/no spam"). This type of learning is known as "supervised" since there is some prior knowledge in the training set (i.e. which emails are "spam").

On the other hand, we have cases where supervision (i.e. prior knowledge on the data) is not available. In this case we have only the positions of the data in the feature space and we are seeking the organization of the data into cohesive groups indicating correlation – this is the case of *unsupervised learning*. The prevalent methods in this area belong to the family of algorithms under the term "clustering" [2]. Other approaches are seeking for repeating patterns and associations in the data. A famous approach includes the "association rules" methods that seeks, usually in retail data, highly correlated product purchases trying to detect a kind of causality based on conditional probabilities [3]. Numerous applications include retail, biomedical applications, recommendations, etc.

As mentioned above, in the last decade due to the presence of abundant data and computing power as well as the vast industrial investment (i.e. GAFA, etc.), we testified the rise and dominance of Deep Learning (DL) – as a reincarnation of Artificial Neural Networks [19] – long standing methods idling for decades. DL is based on the perceptron base learner [18] and consists of potentially complex architectures with (say k) consecutive layers of perceptrons? with or without feedback. Deep learning essentially – in the presence of a large enough amount of data, say X - is able to learn an internal representation Y of the input data X via the successive mappings through the layers aforementioned. This representation Y can be called embedding or even be represented by a potentially very complex non-linear function $Y = f_1(f_2 (..f_k(X)))$, where k the number of layers from the input to the output. The coefficients of this function are essentially the weights of the connections among the perceptrons of the different layers and are learned via an error function that is used and the famous back-propagation process [7]. The representation Y can be used in supervised learning tasks (i.e. image or document classification) or end to end tasks (i.e. machine translation, document summarization, etc.). An interesting variance of DL is the auto encoder [8] approach where the task is to train an architecture whose output (decoder) should be the data themselves. Recently the area of graph neural networks (GNNs), considering graph as input to the DL architectures, is arising as a powerful alternative to traditional matrix-based data representations. The importance of DL is already cornerstone as in the presence of lots of data tasks of unprecedented complexity can be efficiently tackled by them.

Unsupervised learning is the most challenging and promising part of AI since the pace of data production is overwhelming and supervision is less and less possible.

AI impact

Several years after IBM's Deep Blue defeated chess grandmaster G. Kasparov, human-computer cooperation bloomed. However, in recent years AI systems are so so good at playing games (i.e. chess) that their human counterparts seem to be of less value and this can be regarded as a precursor to what might happen at a more pervasive level. Another recent approach, DeepMind's Alpha Zero [13], went from utter ignorance to creative mastery in under four hours without the help of any human guide, to dominate the world's best AlphaGo players and programs.

AI is beyond a technology wave as it penetrates almost all aspects of human life - work, privacy, government – even questioning the human nature with advanced robots (cyborgs). Therefore, a

large-scale discussion among all social and economic parties [11] should decide the limits of AI with regards to ethics, in terms of how much decision power should be left to AI, its impact to democracy. Another important aspect is *explainability* - whether we can ask for an explanation for AI decisions coming from very complex algorithms.

The *impact of AI* in economy and society is already very visible and will be increasingly significant in decision making, business models, risk mitigation, and system performance in many different application domains. The adoptions of AI algorithms in economy and production is expected to affect the world's GDP that will be increased by $15.7 trillion, a full 14%, by 2030. A significant part of this will be in China [9].

AI challenges for the future

The *future of AI* is a challenging debate. There are theories that AI will converge to *singularity* meaning that our technological creations exceed the computing power of human brains. There are predictions, based on Moore's Law and the general trend of exponential growth in technology, that singularity will come before the mid-21st century according to Ray Kurzweil, Google's director of engineering [10]. This means that AI will supersede human intelligence and that AI algorithms will potentially realise their existence and start to behave selfishly and cooperatively.

On the other hand, significant figures of Science and Technology such as Stephen Hawking, Elon Musk, warn [12] towards AI and the singularity. Potential risks include that AI will relieve us or deprive us from our work and a lot of jobs will be replaced by automated processes, even non repetitive ones such as chatbots responding intelligently to human spoken requests.

Also, the aspect of national and European sovereignty for AI is significant and discussed by the Villani report [11] aiming at maintaining local resources and talent in this strategic area. In this context since AI resources and capabilities are of cornerstone importance. We need to reflect of the following aspects:

- *Control and access to the data*. Data are like fuel to the AI engine. Currently the vast amounts of data produced by user's behaviour and interactions are mainly owned by private entities that capitalise on knowledge extracted from these data. This enables them to develop further their competences and products on AI. Therefore, it becomes essential that data are available to the research community and the governments for policy and decision making. A key political question of our era is "how do you regulate the ownership of data?" [1]. Of course, then the issue or privacy becomes important.
- *Computing resources/GPUS*. As mentioned above the capability to apply DL algorithms on bigdata depends totally on the availability of GPUs - the specialized processors for fast matrix computations necessary in deep learning. In the western world, GPUs are produced practically by a handful of industries mostly based in one country. On the other hand, large scale installations of GPUs – enabling the necessary computations are can only be afforded by a few large industries (mainly GAFAs) or governments. This oligopoly apparently poses threats for the future capacity of other governments, SMEs and academia to continue doing competitive research and advancing technology.
- *Access to AI algorithms*/methods. AI algorithms form the heart of AI, so the capacity to design and develop algorithms gives a very high added value and competitive advantage to the owners of the algorithms. For example, searching in the web is a ubiquitous activity giving access to information matching users' queries and is dominated in the western world by Google. It's ranking algorithm though – responsible for prioritising the best results – is a well-kept secret. On the other hand, a large number of popular machine learning algorithms is available as opens source code within the *scikit-learn* library [15] that has made possible the dissemination

of machine learning throughout the academic and industrial community with significant added value accumulated. Similarly for DL Tensorflow [16] and Pytorch [17] have made possible the breath-taking development of deep learning methods. The value of having algorithms as open source code guarantees transparency and increases productivity and progress in this area. A very interesting such case at governmental level is the Translago portal [20] of the French state providing access to all the algorithms used by the state administration to take decisions.

- *Deepening inequality gap among states and social classes.* According to a recent report of McKinsey, half of the world's jobs could be automated by 2055 [14]. This wave of automation will affect classic middle-income white-collar jobs, such as bank tellers, insurance underwriters, loan officers and case-file workers, essentially each job that includes following rules and making few decisions. Thus, the rich-get-richer dynamics of the digital economy is leading to monopolising sectors, where companies like Facebook, Google, Apple and Amazon are eliminating competition.

Here we only present challenging aspects of AI for the future. Tackling these problems is a highly complex task that touches upon the synergy of socio-economic and political players worldwide. The only certainty we have is that AI is already happening and transforms the world very fast. We can only hope that the decisions of societies and governments will lead humanity to enjoying the added value of AI rather than concentrating it to a handful of powerful players.

References

[1] Harari, Yuval Noah; Spiegel & Grau (2018). *21 Lessons for the 21st Century*. ISBN 9780525512172.

[2] Rui Xu and D. Wunsch, "Survey of clustering algorithms," in IEEE Transactions on Neural Networks, vol. 16, no. 3, pp. 645-678, May 2005.

[3] Agrawal, R.; Imieliński, T.; Swami, A. (1993). "Mining association rules between sets of items in large databases". Proceedings of the 1993 ACM SIGMOD international conference on Management of data - SIGMOD '93. p. 207. CiteSeerX 10.1.1.40.6984. doi:10.1145/170035.170072. ISBN 978-0897915922.

[4] X. Wu, X. Zhu, G. Wu and W. Ding, "Data mining with big data," in IEEE Transactions on Knowledge and Data Engineering, vol. 26, no. 1, pp. 97-107, Jan. 2014.

[5] The Evolution of Hard Disk Drives and Storage Capacity, https://www.securedatarecovery.com/blog/hdd-storage-evolution

[6] Russell, Stuart J.; Norvig, Peter (2009). *Artificial Intelligence: A Modern Approach* (3rd ed.). Upper Saddle River, New Jersey: Prentice Hall. ISBN 978-0-13-604259-4.

[7] David E. Rumelhart, Geoffrey E. Hinton & Ronald J. Williams, "Learning representations by back-propagating errors", *Nature* volume 323, pages 533–536 (1986)

[8] Hinton, G. E., & Zemel, R. S. (1994). "Autoencoders, minimum description length and Helmholtz free energy". In *Advances in neural information processing systems* 6 (pp. 3-10).

[9] Sizing the prize, PwC's Global Artificial Intelligence Study: Exploiting the AI Revolution, What's the real value of AI for your business and how can you capitalise? https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.htm

[10] https://www.sciencealert.com/google-s-director-of-engineering-claims-that-the-singularity-will-happen-by-2029

[11] Villani, Cédric & Schoenauer, Marc & Bonnet, Yann & Berthet, Charly & Cornut, Anne-Charlotte & Levin, François & Rondepierre, Bertrand. (2018). *Donner un sens à l'intelligence artificielle : Pour une stratégie nationale et européenne*

[12] https://www.bbc.com/news/technology-30290540

[13] Silver, David; Hubert, Thomas; Schrittwieser, Julian; Antonoglou, Ioannis; Lai, Matthew; Guez, Arthur; Lanctot, Marc; Sifre, Laurent; Kumaran, Dharshan; Graepel, Thore; Lillicrap, Timothy; Simonyan, Karen; Hassabis, Demis (December 5, 2017). "Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm". arXiv:1712.01815 [cs.AI].

[14] https://www.mckinsey.com/featured-insights/digital-disruption/harnessing-automation-for-a-future-that-works

[15] https://scikit-learn.org/stable/

[16] https://www.tensorflow.org/

[17] https://pytorch.org/

[18] Rosenblatt, F. (1958). "The Perceptron: A Probabilistic Model For Information Storage And Organization In The Brain". Psychological Review. 65 (6): 386–408. doi:10.1037/h0042519.

[19] Kleene, S.C. (1956). "Representation of Events in Nerve Nets and Finite Automata". Annals of Mathematics Studies (34). Princeton University Press. pp. 3–41.

[20] https://www.transalgo.org/le-projet/