The biggest promise of AI: Making scientists and engineers stronger

Branislav Kisačanin, *Senior Member, IEEE*

***Abstract*—In this paper we attepmt to provide a very brief history of recent developmnets in artificial intelligence. We discuss how the latest applications of artificial intelligence, more precisely, deep learning, already make us stronger, helthier, and more efficient. We also review some of the latest discoveries in this arena and discuss their promise to continue the trend of this new technology making us stronger. Finally, we briefly discuss potential dangers of artificial intelligence and the need to legally regulate it.**

***Index Terms*—Deep learning, artificial intelligence, legal regulation, current and future applications.**

I. Introduction

IN recent years and especially in recent months, we have seen a tremendous interest from both scientific and non-scientific media in the ongoing development of applications of artificial intelligence. Regrettably, not all of the media coverage has been helpful in explaining to the public what the potential benefits and dangers of artificial intelligence applications may be. With this paper we attempt to provide more information on both topics.

We will not attempt to properly define what we mean by artificial intelligence, because even the much older notion of intelligence is very hard, maybe even impossible, to properly define. In declining to properly define the main topic of our discussion, we follow the Nobel laureate Francis Crick, one of the discoverers of the structure of DNA. Crick wrote about his more recent work on consciousness [1] and about the difficulties of defining it. In the end he said he was not going to define consciousness and argued that he and his colleagues had been able to make pioneering discoveries in genetics, despite not having a proper definition of genes.

Nevertheless, since artificial intelligence means different things to different people, we do provide a working definition of it, in order to give a rough indication about what we discuss here. In the past decade, the principles of machine learning [2], in particular deep learning [3,4], have been successfully applied to a number of previously unsolvable problems. In a nutshell, machine learning is a subset of artificial intelligence, investigating mathematical methods for using data to improve performance of computer algorithms, while deep learning is a subset of machine learning focusing on using artificial neural networks with multiple hidden layers. Results achieved by deep learning in recent years, caused it to practically become synonimous with the much broader area of artificial intelligence. With all that, our working definition of artificial intelligence is: artificial intelligence are software pograms that complete themselves [5].

A good example is the challenge of programming a computer to recognize what is in photographs showing dogs, people, cars, trees, etc. A typical way to solve this problem using deep learning is to write a program that implements an artificial neural network with millions of artificial neurons organized into hundreds of layers and with billions of parameters describing the weights and biases of connections between neurons in the network. The network designer decides on the architecture of the network but leaves its parameters as random initial values. The program is then given access to millions of images, each labeled with the correct classification of what is in the image, a dog, a tree, a person, etc. After millions of iterations in which the program gradually changes its own parameters in order to minimize the combination of the classification error and possibly other terms, such as the regularization error, the program is able to recognize what is in the photographs with a fairly small error, say less than 10%. This process is called training and after it is complete, not only can the program reliably recognize what is in the images it saw during training, it can also generalize what it learned and recognize objects in previously unseen images. In this process, the program completes itself by modifying more than 99% of the original code – the network parameters.

In Section II of this paper we attempt to present a short history of modern artificial intelligence. In Section III we talk about its applications that are already a part of everyday routine for many people. In Section IV we review some of the most interesting recent discoveries in this domain. They may indicate where in the near future we might expect to see big advances based on this technology. In Section V we discuss possible dangers of this new technology and the need to regulate it.

II. A Very Short History of Modern Artificial Intelligence

Most researchers agree that it was the year 2012 when deep learning began to practically be synonimous with artificial intelligence. Before then, artificial intelligence was synonimous with expert systems and other rule-based techniques. So what happened in 2012? That year, for the first time, the winner of the ImageNet competition [6], in which researchers submit their programs that can recognize types of

Branislav Kisačanin is with Nvidia Corporation, Santa Clara, CA, USA and with The Institute for Artificial Intelligence Research and Development of Serbia, Novi Sad, Serbia (e-mail: b.kisacanin@ieee.org), ORCID ID (https://orcid.org/0000-0001-7532-1106).

 ISBN 978-86-7466-970-9

objects in photographs, the winner was a program written using deep learning. It was called AlexNet [7]. Deep neural nets had won major competitions even before that, but this was remarkable in at least two aspects: the accuracy margin AlexNet had achieved over other competing programs and the fact that in the following years only deep neural networks were submitted to this competition, as no hand-written, rule-based, programs could approach the accuracy of deep neural networks ever again. Furthermore, this was the case for other image recognition competitions and beyond. Suddenly, it was recognized that the same principles can be used to make progress in other areas such as speech, gesture, and hand-writing recognition, internet search, autonomous driving, medical diagnosis, etc.

What precipitated this new development? As indicated in Fig.1, around 2012 three fairly independent disciplines matured and converged to cause what is rigthfully called the “deep learning big bang,” because deep learning got successfully applied to a huge variety of problems.

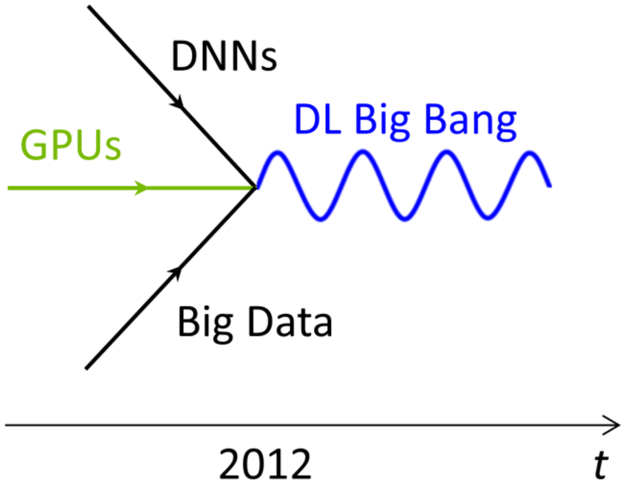

Fig. 1. A diagram (in the style of a Feynman diagram) of the factors contributing to the deep learning big bang, a literal explosion in applications of deep learning since 2012: DNN (deep neural networks), GPU (graphics processing units), and big data.

In no particular order, the first component of the deep learning big bang were architectures and learning algorithms for deep neural networks. The second was the computational power offered by GPUs (graphics processing units), highly parallel processors initially developed for computer graphics, which were found by researchers to work remarkably fast for many tasks they needed. Finally, the third component was what is often called “big data,” meaning availability of huge amounts of labeled data that researchers could freely use to develop their algorithms. All that, together with the culture of openness and knowledge sharing, contributed to a literal explosion of applications of deep learning to a broad set of problems and to deep learning becoming practically synonymous with artificial intelligence.

III. Artificial Intelligence in our Everyday Routine

Whenever we use internet search engines (be it Google, Bing, Yahoo, Baidu, Yandex, or any other) we use programs developed using deep learning. The same is true when we use recommendation engines (think Amazon, Facebook, YouTube, etc.). Many of the millions of Tesla electric vehicles are used in the Autopilot™ mode, driving their passengers autonomously on highways. This system too works on the deep learning principles. Doctors and patients increasingly use software based on artificial intelligence to improve diagnoses and therapy execution. An example dear to my heart is the Omnipod™ insulin pump that now works autonomously most of the time. It receives blood glucose readings from a continuous blood glucose monitor and autonomously makes decisions about how much basal insulin to deliver. It does it better than I ever could because it does it every few minutes and so improves my health and also relieves me from having to do it myself all day long.

In all of these examples of everyday use of artificial intelligence, or more precisely, of deep learning, we see dramatic improvements, even completely new products and services that did not exist just a decade ago. This progress is perhaps best illustrated by the advances in ultrasound imaging. Some twenty years ago, when my kids were born, a typical ultrasound image of a baby in her mom’s belly was grainy and hard to interpret. If you could see the baby’s nose, you were ecstatic and only your doctor could tell whether the baby is a girl or a boy. Today, thanks to artificial intelligence (again, more precisely, thanks to deep learning), even older ultrasound machines can produce wonderful 3D images of the baby. The software combines those grainy ultrasound images to produce a beautiful 3D rendering of the baby’s face and fingers, as illustrated in Fig.2.

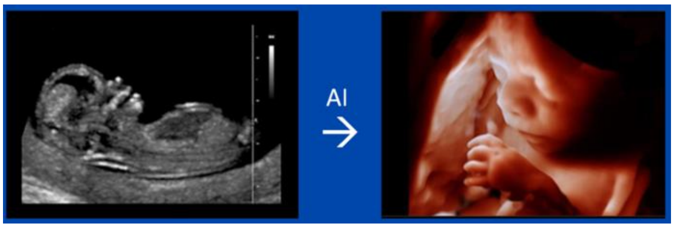

Fig. 2. Improvements in ultrasound imaging are a great illustration of the many advances achieved in the past decade thanks to artificial intelligence, more precisely, based on deep learning techniques.

The common denominator of all of these applications of artificial intelligence is the increase in quality of life and work efficiency.

IV. New Discoveries and Future Applications

We do not endeavor to predict when and which new areas will benefit from artificial intelligence. Instead we look at three recently published discoveries that may turn into revolutionary improvements that benefit the whole planet.

First of them is already creating a revolution in pharmaceutical industry by dramatically shortening the time

 ISBN 978-86-7466-970-9

needed to discover new medications. This is based on the work by DeepMind scientists that showed that the 3D shape of protein molecules can be inferred in a matter of seconds, rather than months, just from their chemical structure [8]. This is important because the 3D shape of a molecule determines its potential role in metabolism or disease mechanism. However, not even the best experts in this field could tell the 3D shape of a new protein just based on its chemical composition. To determine it, the molecule of interest had to be sent to specialized labs and analyzed using X-ray crystallography, a time-consuming process. The solution of this problem was even more important considering this was the last step in drug discovery process that previously could not be computerized. Soon we can expect a flurry of new medications, even for rare diseases, based on this discovery.

The next discovery we discuss was considered so important and original, that it was featured on the front page of Nature [9]. A team of mathematicians and computer scientists from DeepMind, University of Oxford and University of Sydney collaborated on using deep neural networks to obtain new intuition on two mathematical problems, one in topology, and the other in combinatorics. The problems were so challenging that mathematicians could not see the way forward, in some cases for years. For example, topological objects called knots (think nautical knots and then much more) have well defined topological and algebraic properties, but until this work there were no known theorems relating the two classes of properties. With the help of graphical neural networks the first such relations have been glimpsed and this helped mathematicians further develop their intuition about the problem and prove new theorems. At a higher level, this work gives us a view of how all sciences may benefit from tools based on artificial intelligence, helping the scientists and engineers create new discoveries.

The last discovery we discuss here is the farthest from actual use, but may be the most far reaching of all applications of deep learning principles. In recent paper [10] scientists from DeepMind and EPFL reported using deep learning to model the behavior of plasma in a nuclear fusion reactor. Very fast changes of the shape of plasma in magnetic field containers are one of the greatest challenges in achieving controlled nuclear fusion and thus having an inexpensive and clean source of abundant energy. Using their deep learning model in conjunction with an actual tokamak reactor, they achieved a new record in how long they could sustain the reaction, now on the order of seconds. Of course, this is still far from making a controlled nuclear fusion a reality, but is a huge leap forward, and we expect to hear more about this line of research.

V. Ethical and Legal Considerations

Ethical and legal issues surrounding any new technology are of huge importance. In this paper we do not want to provide ideas to bad actors on how they might use deep learning maliciously or unethically, but want to say that even the simplest technology can be used to harm people or property, intentionally or unintentionally. Think of a hammer, for example, which we can use for intended purposes, but a hammer can also be used to cause harm, even unintentionally, when we miss a nail and hit our finger instead. This does not mean we need to lock all hammers and never use them. We have laws and courts to deal with cases of improper use of all kinds of tools and technologies. That is why leaders in both industry and academia agree that artificial intelligence too needs to be regulated legally, and the sooner the better.

To summarize this section, we propose an update to an old saying about water and fire: Water, fire, and technology are good servants but bad masters. Every researcher and engineer aiming to use this new technology must think about this and find ways to prevent its abuse.

References

[1] F. Crick, *The Astonishing Hypothesis*, Simon & Schuster, New York, 1994.

[2] V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer-Verlag, New York, 1995.

[3] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, pp. 436-444, 2015.

[4] J. Schmidhuber, “Annotated history of modern AI and deep learning,” arXiv, 2022.

[5] Jensen Huang, Keynote at GTC 2020, https://www.nvidia.com/en-us/on-demand/session/gtcfall20-a21569-4/. 2020.

[6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” *IEEE Computer Vision and Pattern Recognition*, 2009.

[7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” *Neural Information Processing Systems*, 2012.

[8] J. Jumper, R. Evans, A. Pritzel, et al., “Highly accurate protein structure prediction with AlphaFold,” *Nature*, vol. 596, pp. 583–589, 2021.

[9] A. Davies, P. Veličković, L. Buesing, et al., “Advancing mathematics by guiding human intuition with AI,” *Nature*, vol. 600, pp. 70–74, 2021.

[10] J. Degrave, F. Felici, J. Buchli, et al., “Magnetic control of tokamak plasmas through deep reinforcement learning,” *Nature*, vol. 602, pp. 414–419, 2022.

 ISBN 978-86-7466-970-9