Opinion: artificial intelligence and scientific intelligence

Marc Mézard, École normale supérieure, PSL University

By now, everybody should know that the recent progress of Artificial Intelligence (AI) will produce a dramatic impact on many sectors of human activity.

AI has obtained spectacular results that would have been considered impossible ten years ago. Last year a new algorithm beat the world's best Go player. We can now process images automatically, identify faces, segment images and provide a semantic description of their content, opening the way to self-driving cars or trucks. Voice recognition and automatic translation are progressing rapidly. Algorithms are competing with the best professionals at analyzing skin cancer symptoms or detecting specific anomalies in radiology. And physics is also concerned: AI is used to help identify new particles in accelerator physics, to analyze cosmological data, or perform quantum chemistry simulations.

The recent breakthrough is based on machine learning: the algorithms are programmed to learn from examples. They are often based on layered artificial neural networks, where each "neuron" receives information from neurons in the previous layer, performs a simple computation and in turn sends a few bits of information to the next layer. A modern "deep" network with hundreds of layers, analyzing an image, can contain hundreds of millions of adaptable parameters ruling these elementary computations. They must be determined through supervised learning: a large "training set" of examples is presented, and the parameters are adapted, typically by gradient descent, so that the intended purpose is obtained on this training set. The generalization performance of the obtained machine is then tested on new data.

The paradigm of layered neural networks has existed for over 50 years. However, until the field's recent revival, it was not successful on real-size practical applications. Its revival is due to the increase in computing power, to the availability of very large labeled data sets for training (in fact, the progress in "big-data" analysis and in machine learning are strongly correlated), and to some pre-processing and training tricks developed in the 2000's.

In spite of its practical success, the scientific understanding of deep networks lags far behind. The learning process is poorly understood. Gradient descent in a complicated 10^8-dimensional parameter space should typically be trapped in inefficient regions. Even when using large training sets, a successful training with that many parameters could lead to "overfitting", namely poor generalization on new data.

Yet, in practice, training works and finds a good-enough set of parameters, producing a machine that can be smarter than us. At a more abstract level, understanding how the information, stored collectively by the neurons inside each layer, becomes more global and more abstract when one goes deeper into the layers is a major challenge. Scientific intelligence, with input from statistics, information theory, computer science and mathematics is needed in order to come up with a theoretical framework. Because of the collective nature of information processing in these systems, one can expect that statistical physics will play a major role in these future developments.

Scientific intelligence, with input from statistics, information theory, computer science and mathematics is needed in order to come up with a theoretical framework.

Having underlined the technological and societal importance of the AI revolution as well as its scientific challenge, let me highlight that, as far as "intelligence" is concerned, these machines remain very limited. They can achieve specific tasks, characterized by simple answers, but they are far from building a representation of the world, or from any kind of creative reasoning. In science, deep networks and new data-science algorithms are extremely useful additions to our toolbox, as was the appearance of numerical simulations a few decades ago. However, these machines cannot replace modeling, *i.e.* building a concise, workable and predictive representation of the world. Contrary to some bold claims, they will not kill the scientific method that we have been using for five centuries. Rather, they will improve it. ■