The Unnatural Scientist: could AI replace researchers?

February 13, 2019

Observe. Hypothesise. Experiment. Review the theory. Repeat. The scientific method underpins every aspect of a natural scientist's work, providing the philosophical backbone for our shared knowledge system. Today, given a newfound enthusiasm for machine learning (ML), it's quite easy to see where artificial intelligence (AI) fits into small parts of this process, from detecting hidden patterns in data to parsing important text from papers. But is it possible that AI could ever replace the Natural Scientist? To answer this it is critical to evaluate its history, current successes and speculate a bit on the future all with respect to the scientific process.

Artificial Intelligence is a nebulous term, applying to a wide variety of programs, including esoteric self organizing maps and garden variety chess engines - to quote Tesler's theorem, "AI is whatever hasn't been done yet.". These days however, Artificial Intelligence has come to be popularly understood in terms of neural networks, deep learning, natural language processing and image analytics. This wasn't always the case - the story of AI begins in the summer of 1956 in a now famous conference in Dartmouth College where various luminaries gathered for a "two month, 10 man study" where they toiled to "make machines that use language, form abstractions and concepts", and most crucially: "solve kinds of problems now reserved for humans" - problems like scientific endeavour (although this was unlikely to be considered a possibility at the time)[1] . However, the initial enthusiasm dwindled when a combination of increased skepticism, more difficult problems and hardware limitations lead to the first "AI Winter" in the 1970's, where funding dried up and interest went elsewhere.

Ignoring a brief thaw in the early 1980's for long forgotten "expert system" programs, this winter of discontent continued until the 1990's, with the introduction of new techniques like the neural network - a program modelled on the human brain that adjusted its parameters backwards on feedback, and genetic algorithms - where the best characteristics from each generation are selected and combined, clearly modelled on traditional evolution.

Could AI replace a scientist?- from humble beginnings in 1956 AI has only increased in flexibility and power and can now detect and learn patterns in speech, text and images.

Computer scientists had taken their cues from the Natural sciences, adapting biological ideas in order to improve their models. It was only natural, therefore, that these new techniques would then be applied back to the sciences, acting as improvements to traditional algorithms by being more flexible and able to fail gracefully and so were gradually introduced to scientific papers.

In "Neural Nets: A New Way to Catch Elusive Particles?"[2], published in Science in 1992, Faye Flam describes how neural networks can be applied to particle physics to sift the "colossal" piles of data - "10 megabytes ... every second", to find hidden patterns and the "one-in-a-trillion discovery". Genetic algorithms, on the other hand, were being explored in Biochemistry at roughly the same time. In 1995, scientists at the Krebs Institute published "Molecular recognition of receptor sites using a genetic algorithm"[3], where they gleaned new insight into the still raging debate

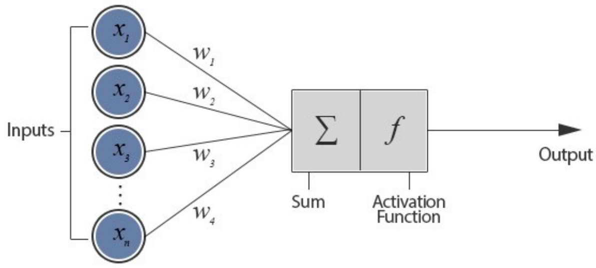

Neural Networks are inspired by the human brain *- this perceptron configuration takes a number of weighted inputs and "fires" if the sum is greater than a certain value via the activation function, similar to how synapses fire if the excitatory input is larger than the inhibitory input. The output is compared to the training data and the weights adjusted until the system "learns". (Reproduced from [6])*

as to whether substrates fit exactly into receptor sites in a "lock and key" manner or as an "induced fit".

The problem with these precursor papers is that computers just weren't powerful enough and there wasn't enough research to truly harness the power of these programs. Today, "10 megabytes of data ... every second" is easily analysable by off the shelf laptops, and genetic algorithms easily run on the same. Artificial Intelligence was used like a calculator or model - a tool used to analyse the data or present the model in a new way, but critically limited by slow calculation speeds and it was simply not considered that a machine learning system could ever emulate such a complex, abstract series of tasks like the scientific method. As such, AI was never expected to replace the scientist in much the same way the abacus or calculator hadn't.

We are currently living in a new AI renaissance, with powerful computers, bigger datasets and burgeoning interest in data science fuelling the revolution. However, we have yet to reach the peak saturation of a new gilded age, with many advancements just around the corner: automated driverless cars optimising traffic, perfect speech and language recognition or "intelligent" sensor robots for complicated spatial tasks like cooking. With the newfound general interest in machine learning comes a veritable deluge of academic papers in the natural sciences all finding new and unique ways to apply different programs to their models.

Despite the large increase in publications utilising machine learning for various types of data analysis, in order to replace the natural scientist our AI systems must be able to perform the whole scientific method. This includes more abstract and difficult tasks like making significant observations, generating hypotheses (although some companies like Elsevier are trying) and engaging in a scientific discourse that furthers research. Fortunately, recent Natural Language Processing (NLP) research has created useful statistical models called topic modelling[4] specifically for abstract tasks like this. Topic models treat a document as a collection of words, generated by the related topics or themes, and uses probabilities of words in the document to categorize, index or even summarise whole papers. It's not hard to see how this could be expanded and applied to the Natural sciences - imagine an AI assistant you prompt with your subfield and watch it select the most relevant papers and maybe detect links that you might not have thought of initially.

One paper[5] published in 2001 provided a roadmap for machine learning in each step of the scientific process, and each of its suggested methods for assimilating ML have only been expanded and further developed in the past decade and a half. However, one of the major limitations highlighted by the paper remains to this day: "enabling ML methods to communicate with both scientists and other ML methods". In order to fully replace the natural scientist, these AI systems would ideally be "all-in-one" systems that explore the theory with topic modelling, generate hypotheses, test them, record and analyse the data and draw conclusions. Many of these tasks can be done in isolation, but turning isolated outputs into tasks or inputs for the next program is a real challenge and would require lots of computer science legwork to make a generalizable ML framework for these exchanges. An "all-in-one" system that could generate analytic conclusions from data it itself has measured and interpreted would be a major milestone and one which sadly doesn't seem likely in the next decade.

Since their introduction, computers have been revolutionising science and the wider world, changing and enhancing the way we interact with data, handle arithmetic and communicate with colleagues across the world. As AI transforms our daily lives in many diverse ways in our cyber-ecosystem, it's not hard to see how science and scientists can be aided and improved by programs that can see faint patterns in big data, summarise reports and adapt old models. However, given the incredible, sometimes hidden complexity of the multitude of tasks scientists perform every day, it is often difficult to see our current crop of intelligent machines replacing them wholesale. In the same way that computer scientists modelled some of the most successful machine learning algorithms on science, scientists must now use the models and programs from the computer scientists to revolutionise our approach to the scientific method, through strong partnerships, mutual understanding of the theory and experts positioned at the interfaces of both fields both domains could benefit enormously. Like the computer before it, AI has and will continue to be further integrated into the scientific method, but will not replace the cohesive ability and perspectives of a human natural scientist for a long time yet.

References

[1]N. Bostrom, Superintelligence. Oxford University Press, 2015, p. 6.
[2]F. Flam, "Neural Nets: A New Way to Catch Elusive Particles?", Science, vol. 256, no. 5061, pp. 1282-1283, 1992. Available: 10.1126/science.256.5061.1282.
[3]G. Jones, P. Willett and R. Glen, "Molecular recognition of receptor sites using a genetic algorithm with a description of desolvation", Journal of Molecular Biology, vol. 245, no. 1, pp. 43-53, 1995. Available: 10.1016/s0022-2836(95)80037-9
[4]M. Jordan and T. Mitchell, "Machine learning: Trends, perspectives, and prospects", Science, vol. 349, no. 6245, pp. 255-260, 2015. Available: 10.1126/science.aaa8415
[5]E. Mjolsness, "Machine Learning for Science: State of the Art and Future Prospects", Science, vol. 293, no. 5537, pp. 2051-2055, 2001. Available: 10.1126/science.293.5537.205
[6]H. Pokharna, "For Dummies — The Introduction to Neural Networks we all need ! (Part 1)", Medium, 2016. [Online]. Available: https://medium.com/technologymadeeasy/for-dummies-the-introduction-to-neural-networks-we-all-need-c50f6012d5eb