St Anne's Academic Review
Issue 8 - October 2018
http://st-annes-mcr.org.uk/staar/publications/staar-8-2018/

NEW HORIZONS

Perspectives

Why Artificial Intelligence Will Enable New Scientific Discoveries

Andrew Briggs

Abstract. *We are in a time of change; new technologies emerge daily and this is unlikely to slow over the foreseeable future as key advances in artificial intelligence are made. The scientific community has recognised the large-scale impact AI will have on research from machine learning optimising experimentation to AI aiding drug discovery. In the UK there are a number of AI start-ups and research groups, making the country well placed to pave the way for significant changes to scientific research with the aid of AI technology. With these advancements we are on the brink of being able to tackle problems humans have so far been unable to solve alone. This post originally appeared on Graphcore in November 2017. We thank the author for permission to republish it.*

* * *

I am enthusiastically looking forward to a dramatic change in the way science is done. This change will be as significant as the application of science to technology that drove the industrial revolution.

In materials science, as in other branches of experimental science such as drug discovery, we are able to obtain ever larger amounts of data from ever more sophisticated experiments and modelling. And yet the choice of what data to collect, and the process of how best to analyse it, at times seem like a cottage industry compared to the sophistication of internet searches at Google or Baidu. Our current approaches to data in science fall well short of the advanced machine learning techniques social media platforms use to recognize the friend in a photo we've uploaded or what film we might like to watch next.

Advances in AI are likely to change every aspect of our lives, from transport and employment to care of the elderly and will have a huge impact on health care overall. Science and engineering research are set to benefit massively from the application of AI and UK science is sitting in pole position, ready to take the lead. In the future we can aspire to AI achieving a series of increasingly ambitious scientific goals:

- A result worthy of a second year graduate student

- A body of work worthy of a PhD

- An advance worthy of a Fellowship of the Royal Society or the Royal Academy of Engineering

- A discovery worthy of a Nobel Prize

The fourth goal has already been articulated as a Grand Challenge for biomedical science, by Hiroaki Kitano, Director of Sony Science Laboratories, "to develop an AI system that can make major scientific discoveries in biomedical sciences and that is worthy of a Nobel Prize and far beyond." (see: https://www.aaai.org/ojs/index.php/aimagazine/article/view/2642). We should be aiming for Nobel Prize worthy discoveries in all scientific disciplines.

I predict that by 2025, AI will be as ubiquitous for running experiments as computers are today for controlling instrumentation and logging data. The paradigm shift will be from AI used for analysing the data which has already been obtained, to AI deciding what to measure next. The key advances in AI that we are currently exploring to make this possible are reinforcement learning and Bayesian optimisation.

Suppose we have a belief about something, which could be a scientific hypothesis or a function which describes a phenomenon under study. There is some uncertainty, which could be reduced by acquiring more data, but the data are costly to acquire. How should we choose what experiment to do next? How can we optimize for the perennial constraints in scientific research of time and money? Advances in AI, particularly those within machine learning, are now opening up the possibility to automate these processes.

These same methods can also be used to optimise a system for a given set of criteria. By analysing the results of new experiments and new material discoveries we can further tune our approaches, saving time, money and improving outcomes. This can be combined with hardware-in-the-loop testing to tune a component or sub-subsystem for its eventual performance in service. Initially this can be applied to prototypes which have already been made, but eventually it can be applied to the whole design and even to the manufacturing process. This could make a significant difference in the cost of production for advanced materials, allowing them to move from being used only in expensive one-off pieces of equipment into high volume production for new consumer devices.

Artificial Intelligence techniques in science will be applied to quantum technologies, batteries, solid state lighting, nanoelectronics and nanomechanics, for high throughput screening of new materials in both simulation and in experiments, for computer vision in microscopy, radiography and tomography and for optimisation of data-rich manufacturing operations.

However, the types of machine learning techniques that will work best in aiding scientific discovery are complex and will require computing hardware that is much more efficient and flexible than exists today. New hardware that can efficiently accelerate these approaches is becoming available from companies like Graphcore with their Intelligence Processing Unit (IPU) technology. By using these new advanced IPU's we will be able to develop new innovations in machine learning approaches, that will allow us to achieve breakthroughs in scientific research. In addition, we will create a cohort of researchers with first-hand experience of cutting-edge AI and its application to science, engineering and technology. This will further stimulate the development of AI techniques for broader applications in science and engineering research across many fields.

Industry is starting to invest in the automation of research in science. Sony has set a Grand Challenge in biomedical sciences. DeepMind has recently appointed a director for science research. There are several very impressive UK start-ups, both in software and hardware, working in the application of AI to experimental science, who are well placed to contribute, some of them already very well-funded. And many universities, including several in the UK, have groups who are world-leading in the machine learning techniques that are needed for automating science. I believe that the UK is well placed to provide leadership in this new field.

Artificial Intelligence gets a mixed press. Many talk about the risk of jobs being replaced by machines. Today the internet is driven by advertising revenue and so it is also natural to see a lot of AI research focused on improving basic commercial imperatives such as understanding social media feeds or improving internet search.

The application of AI to science research has the potential to make a much more fundamental societal impact. It will allow us to understand and explore the use of new materials and new techniques. Silicon based computers may only have another 10-20 years of advances ahead and so we need to accelerate work on new materials and on the next breakthroughs that will come from quantum computing or eventually from molecular computing. Drug discovery and medical research will also benefit from these new AI driven scientific techniques. Drugs that are able to directly reach the diseased tissues and avoid wider dispersion in the body will result in people suffering far fewer side effects. These types of research breakthroughs will allow our aging populations to live more useful and fulfilled lives.

I am optimistic for our AI future and in particular for the growing role that this new technology will play in scientific breakthroughs. With sufficient investment, very rapid progress could be made within 5 years, maybe not yet to Nobel Prize level, but almost certainly to second year graduate student level. More importantly, new machine learning techniques applied to scientific discovery will quickly enable complex research challenges to be addressed that could *never* be solved by humans alone working within feasible time limits and resources.

Department of Materials