

# The AI revolution in scientific research

## The Royal Society and The Alan Turing Institute

The Royal Society is the UK's national academy of sciences. The Society's fundamental purpose, reflected in its founding Charters of the 1660s, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Alan Turing Institute is the UK's national institute for data science and artificial intelligence. Its mission is to make great leaps in research in order to change the world for the better.

In April 2017, the Royal Society published the results of a major policy study on machine learning. This report considered the potential of machine learning in the next 5 – 10 years, and the actions required to build an environment of careful stewardship that can help realise its potential. Its publication set the direction for a wider programme of Royal Society policy and public engagement on artificial intelligence (AI), which seeks to create the conditions in which the benefits of these technologies can be brought into being safely and rapidly.

As part of this programme, in February 2019 the Society convened a workshop on the application of AI in science. By processing the large amounts of data now being generated in fields such as the life sciences, particle physics, astronomy, the social sciences, and more, machine learning could be a key enabler for a range of scientific fields, pushing forward the boundaries of science.

This note summarises discussions at the workshop. It is not intended as a verbatim record and its contents do not necessarily represent the views of all participants at the event, or Fellows of the Royal Society or The Alan Turing Institute.

## Data in science: from the t-test to the frontiers of AI

Scientists aspire to understand the workings of nature, people, and society. To do so, they formulate hypotheses, design experiments, and collect data, with the aim of analysing and better understanding natural, physical, and social phenomena.

Data collection and analysis is a core element of the scientific method, and scientists have long used statistical techniques to aid their work. In the early 1900s, for example, the development of the t-test gave researchers a new tool to extract insights from data in order to test the veracity of their hypotheses. Such mathematical frameworks were vital in extracting as much information as possible from data that had often taken significant time and money to generate and collect.

Examples of the application of statistical methods to scientific challenges can be seen throughout history, often leading to discoveries or methods that underpin the fundamentals of science today, for example:

- The analysis by Johannes Kepler of the astronomic measurements of Tycho Brahe in the early seventeenth century led to his formulation of the laws of planetary motion, which subsequently enabled Isaac Newton FRS (and others) to formulate the law of universal gravitation.
- In the mid-nineteenth century, the laboratory at Rothamsted was established as a centre for agricultural research, running continuously monitored experiments from 1856 which are still running to this day. Ronald Fisher FRS – a prominent statistician – was hired to work there in 1919 to direct analysis of these experiments. His work went on to develop the theory of experimental design and lay the groundwork for many fundamental statistical methods that are still in use today.
- In the mid-twentieth century, Margaret Oakley Dayhoff pioneered the analysis of protein sequencing data, a forerunner of genome sequencing, leading early research that used computers to analyse patterns in the sequences.

Throughout the 20th century, the development of artificial intelligence (AI) techniques offered additional tools for extracting insights from data.

Papers by Alan Turing FRS through the 1940s grappled with the idea of machine intelligence. In 1950, he posed the question “can machines think?”, and suggested a test for machine intelligence – subsequently known as the Turing Test – in which a machine might be called intelligent, if its responses to questions could convince a person that it was human.

In the decades that followed, AI methods developed quickly, with a focus on symbolic methods in the 1970s and 1980s that sought to create human-like representations of problems, logic and search, and expert systems that worked from datasets codifying human knowledge and practice to automate decision-making. These subsequently gave way to a resurgence of interest in neural networks, in which layers of small computational units are connected in a way that is inspired by connections in the brain. The key issue with all these methods, however, was scalability – they became inefficient when confronted with even modest sized data sets.

The 1980s and 1990s saw a strong development of machine learning theory and statistical machine learning, the latter in particular driven by the increasing amount of data generated, for example from gene sequencing and related experiments. The 2000s and 2010s then brought advances in machine learning, a branch of artificial intelligence that allows computer programs to learn from data rather than following hard-coded rules, in fields ranging from mastering complex games to delivering insights about fundamental science.


The expression ‘artificial intelligence’ today is therefore an umbrella term. It refers to a suite of technologies that can perform complex tasks when acting in conditions of uncertainty, including visual perception, speech recognition, natural language processing, reasoning, learning from data, and a range of optimisation problems.

**Image:** Alan Turing. © Godfrey Argent Studio.

## Advances in AI technologies offer more powerful analytical tools

The ready availability of very large data sets, coupled with new algorithmic techniques and aided by fast and massively parallel computer power, has vastly increased the power of today’s AI technologies. Technical breakthroughs that have contributed to the success of AI today include:

- **Convolutional neural networks:** multi-layered ‘deep’ neural networks, that are particularly adapted to image classification tasks by being able to identify the relevant features required to solve the problem<sup>1</sup>.
- **Reinforcement learning:** a method for finding optimal strategies for an environment by exploring many possible scenarios and assigning credit to different moves based on performance<sup>2</sup>.
- **Transfer learning:** an old idea of using concepts learned in one domain on a new unknown one, this idea has enabled the use of deep convolutional nets trained on labelled data to transfer already-discovered visual features to classify images from different domains with no labels<sup>3</sup>.
- **Generative adversarial networks:** continues the idea of pitching the computer against itself by co-evolving the neural network classifier with the difficulty of the training data set<sup>4</sup>.




1. These techniques were, for example, used to classify the ImageNet database of labelled photos with unprecedented accuracy.
2. The breakthrough example was the AlphaGo project by DeepMind, which used this approach to learn how to play the game Go at expert human levels by simulating many games pitching the computer against itself. Reinforcement learning has recently been used to autonomously design new quantum experiments and techniques.
3. This has been used successfully for classifying nanoscale images from electron microscopes, for example.
4. An original application of this is the generation of fake, but realistic, human faces. The method has also found use in scientific discovery, for example in classifying 3D particle showers at the Large Hadron Collider.

## AI as an enabler of scientific discovery

AI technologies are now used in a variety of scientific research fields. For example:

- **Using genomic data to predict protein structures:** Understanding a protein's shape is key to understanding the role it plays in the body. By predicting these shapes, scientists can identify proteins that play a role in diseases, improving diagnosis and helping develop new treatments. The process of determining protein structures is both technically difficult and labour-intensive, yielding approximately 100,000 known structures to date<sup>5</sup>. While advances in genetics in recent decades have provided rich datasets of DNA sequences, determining the shape of a protein from its corresponding genetic sequence – the protein-folding challenge – is a complex task. To help understand this process, researchers are developing machine learning approaches that can predict the three-dimensional structure of proteins from DNA sequences. The AlphaFold project at DeepMind, for example, has created a deep neural network that predicts the distances between pairs of amino acids and the angles between their bonds, and in so doing produces a highly-accurate prediction of an overall protein structure<sup>6</sup>.



- **Understanding the effects of climate change on cities and regions:** Environmental science combines the need to analyse large amounts of recorded data with complex systems modelling (such as is required to understand the effects of climate change). To inform decision-making at a national or local level, predictions from global climate models need to be understood in terms of their consequences for cities or regions; for example, predicting the number of summer days where temperatures exceed 30°C within a city in 20 years' time<sup>7</sup>. Such local areas might have access to detailed observational data about local environmental conditions – from weather stations, for example – but it is difficult to create accurate projections from these alone, given the baseline changes taking place as a result of climate change. Machine learning can help bridge the gap between these two types of information. It can integrate the low-resolution outputs of climate models with detailed, but local, observational data; the resulting hybrid analysis would improve the climate models created by traditional methods of analysis, and provide a more detailed picture of the local impacts of climate change. For example, a current research project at the University of Cambridge<sup>8</sup> is seeking to understand how climate variability in Egypt is likely to change over coming decades, and the impact these changes will have on cotton production in the region. The resulting predictions can then be used to provide strategies for building climate resilience that will decrease the impact of climate change on agriculture in the region.

5. Lee, J, Freddolkino, P. and Zhang, Y. (2017) Ab initio protein structure prediction, in D.J. Rigden (ed.), From Protein Structure to Function with Bioinformatics, available at: [https://zhanglab.ccmb.med.umich.edu/papers/2017\\_3.pdf](https://zhanglab.ccmb.med.umich.edu/papers/2017_3.pdf)

6. DeepMind (2018) AlphaFold: Using AI for scientific discovery, available at: <https://deepmind.com/blog/alphafold/>

7. Banerjee A, Monteleoni C. 2014 Climate change: challenges for machine learning (NIPS tutorial). See <https://www.microsoft.com/en-us/research/video/tutorial-climate-change-challenges-for-machine-learning/> (accessed 22 March 2017).

8. See ongoing work at the British Antarctic Survey on machine learning techniques for climate projection.





- **Finding patterns in astronomical data:** Research in astronomy generates large amounts of data and a key challenge is to detect interesting features or signals from the noise, and to assign these to the correct category or phenomenon. For example, the Kepler mission is seeking to discover Earth-sized planets orbiting other stars, collecting data from observations of the Orion Spur, and beyond, that could indicate the presence of stars or planets. However, not all of this data is useful; it can be distorted by the activity of on-board thrusters, by variations

in stellar activity, or other systematic trends. Before the data can be analysed, these so-called instrumental artefacts need to be removed from the system. To help with this, researchers have developed a machine learning system that can identify these artefacts and remove them from the system, cleaning it for later analysis<sup>9</sup>. Machine learning has also been used to discover new astronomical phenomena, for example: finding new pulsars from existing data sets<sup>10</sup>; identifying the properties of stars<sup>11</sup> and supernovae<sup>12</sup>; and correctly classifying galaxies<sup>13</sup>.

---

9. Roberts S, McQuillan A, Reece S, Aigrain S. 2013 Astrophysically robust systematics removal using variational inference: application to the first month of Kepler data. *Mon. Not. R. Astron. Soc.* 435, 3639–3653. (doi:10.1093/mnras/stt1555)

10. Morello V, Barr ED, Bailes M, Flynn CM, Keane EF, van Straten W. 2014 SPINN: a straightforward machine learning solution to the pulsar candidate selection problem. *Mon. Not. R. Astron. Soc.* 443, 1651–1662. (doi: 10.1093/mnras/stu1188)

11. Miller A *et al.* 2015 A machine learning method to infer fundamental stellar parameters from photometric light curves. *Astrophys. J.* 798, 17. (doi: 10.1088/0004-637X/798/2/122)

12. Lochner M, McEwen JD, Peiris HV, Lahav O, Winter MK. 2016 Photometric supernova classification with machine learning. *Astrophys. J. Suppl. Ser.* 225, 31. (doi: 10.3847/0067-0049/225/2/31)

13. Banerji M *et al.* 2010 Galaxy Zoo: reproducing galaxy morphologies via machine learning. *Mon. Not. R. Astron. Soc.* 406, 342–353. (doi: 10.1111/j.1365-2966.2010.16713.x)

Machine learning has become a key tool for researchers across domains to analyse large datasets, detecting previously unforeseen patterns or extracting unexpected insights. While its potential applications in scientific

research range broadly across disciplines, and will include a suite of fields not considered in detail here, some examples of research areas with emerging applications of AI include:



© Grafissimo.

© Sezeryadigar.

### Satellite imaging to support conservation

Many species of seal in the Antarctic are extremely difficult to monitor as they live exclusively in the sea-ice zone, a region that is particularly difficult to survey. The use of very high-resolution satellites enables researchers to identify these seals in imagery at greatly reduced cost and effort. However, manually counting the seals over the vast expanse of ice that they inhabit is time consuming, and individual analysts produce a large variation in count numbers. An automated solution, through machine learning methods, could solve this problem, giving quick, consistent results with known associated error<sup>14</sup>.

### Understanding social history from archive material

Researchers are collaborating with curators to build new software to analyse data drawn initially from millions of pages of out-of-copyright newspaper collections from within the British Library's National Newspaper archive. They will also draw on other digitised historical collections, most notably government-collected data, such as the Census and registration of births, marriages and deaths. The resulting new research methods will allow computational linguists and historians to track societal and cultural change in new ways during the Industrial Revolution, and the changes brought about by the advance of technology across all aspects of society during this period. Crucially, these new research methods will place the lives of ordinary people centre-stage<sup>15</sup>.

14. Alan Turing Institute project: Antarctic seal populations, with the British Antarctic Survey


15. Alan Turing Institute project: Living with Machines, with AHRC





### Materials characterisation using high-resolution imaging

Materials behave differently depending on their internal structure. The internal structure is often extracted by guiding X-rays through them and studying the resulting scattering patterns. Contemporary approaches for analysing these scattering patterns are iterative and often require the attention of scientists. The scope of this activity is to explore the options of using machine learning for automatically inferring the structural information of materials by analysing the scattering patterns<sup>16</sup>.



### Driving scientific discovery from particle physics experiments and large scale astronomical data

Researchers are developing new software tools to characterise dark matter with data from multiple experiments. A key outcome of this research is to identify the limitations and challenges that need to be overcome to extend this proof-of-principle and enable future research to generalise this to other use cases in particle physics and the wider scientific community<sup>17</sup>.



### Understanding complex organic chemistry

The goal of this pilot project between the John Innes Centre and The Alan Turing Institute is to investigate possibilities for machine learning in modelling and predicting the process of triterpene biosynthesis in plants. Triterpenes are complex molecules which form a large and important class of plant natural products, with diverse commercial applications across the health, agriculture and industrial sectors. The triterpenes are all synthesized from a single common substrate which can then be further modified by tailoring enzymes to give over 20,000 structurally diverse triterpenes. Recent machine learning models have shown promise at predicting the outcomes of organic chemical reactions. Successful prediction based on sequence will require both a deep understanding of the biosynthetic pathways that produce triterpenes, as well as novel machine learning methodology<sup>18</sup>.

16. Alan Turing Institute project: Small-Angle X-Ray Scattering

17. Alan Turing Institute project: developing machine learning-enabled experimental design, model building and scientific discovery in particle physics.

18. Alan Turing Institute project: Analysis of biochemical cascades

Each different scientific area has its own challenges, and it is rare that they can be met by the straightforward ‘off the shelf’ use of standard AI methods. Indeed, many applications open up new areas of AI research themselves – for example, the need to analyse scanned archives of historical

scientific documents requires the automatic recognition and understanding of mathematical formulae and complex diagrams. However, there are a number of challenges which are recurring themes in the application of AI and its use in scientific research, summarised in the box below.

## BOX 1

### Research questions to advance the application of AI in science

#### DATA MANAGEMENT

**Is there a principled method to decide what data to keep and what to discard, when an experiment or observation produces too much data to store? How will this affect the ability to re-use the data to test alternative theories to the one that informed the filtering decision?**

In a number of areas of science, the amount of data generated from an experiment is too large to store, or even tractably analyse. This is already the case, for example, at the Large Hadron Collider, where typically only the data directly supporting the experimental finding are kept and the rest is discarded. As this situation becomes more common, the use of a principled methodology for deciding what to keep and what to throw away becomes more important, keeping in mind that the more data that is discarded, the less use the stored data actually has for future research.

**What does ‘open data’ mean in practice where the data sets are just too large, complex and heterogeneous for anyone to actually access and understand them in their entirety?**

While lots of data today might be ‘free’ it isn’t cheap: found data might come in a variety of formats, have missing or duplicate entries, or be subject to biases embedded in the point of collection. Assembling such data for analysis requires its own support infrastructure, involving large teams that bring together people with a variety of specialisms: legal teams, people who work with data standards, data engineers and analysts, as well as a physical infrastructure

that provides computing power. Further efforts to create an amenable data environment could include creating new data standards, encouraging researchers to publish data and metadata, and encouraging journals and other data holders to make their data available, where appropriate.

Even in an environment that supports open access to data produced to publicly-funded scientific research, the size and complexity of such datasets can pose issues. As the size of these data sets grows, there will be very few researchers, if any, who could in practice download them. Consequently, the data has to be condensed and packaged – and someone has to decide on what basis this is done, and whether it is affordable to provide bespoke data packages. This then affects the ready availability and brings into question what is meant by ‘open access’. Who then decides what people can see and use, on what basis and in what form?

**How can scientists search efficiently for rare or unusual events and objects in large and noisy data sets?**

A common driver of scientific discovery is the study of rare or unusual events (for example, the discovery of pulsars in the 1960s). This is becoming increasingly difficult to do given the size of data sets now available, and automatic methods are necessary. There are a number of challenges in creating these: noise in the data is one; another is that data naturally includes many more exemplars of ‘normal’ objects than unusual ones, which makes it difficult to train a machine learning classifier.

## AI METHODS AND CAPABILITIES

**How can machine learning help integrate observations of the same system taken at different scales? For example, a cell imaged at the levels of small molecule, protein, membrane, and cell signalling network. More generally, how can machine learning help integrate data from different sources collected under different conditions and for different purposes, in a way that is scientifically valid?**

Many complex systems have features at different length scales. Moreover, different imaging techniques work at different resolutions. Machine learning could help integrate what researchers discover at each scale, using structures found at one level to constrain and inform the search at another level.

In addition to different length scale observations, datasets are often created by compiling inputs from different equipment, or data from completely different experiments on similar subjects. It is an attractive idea to bring together, for example, genetic data of a species, and environmental data to study how the climate may have driven species' evolution. But there are risks in doing this kind of 'meta-analysis' which can create or amplify biases in the data. Can such datasets be brought together to make more informative discoveries?

**How can researchers re-use data which they have already used to inform theory development, while maintaining the rigour of their work?**

The classic experimental method is to make observations, then come up with a theory, and then test that theory in new experiments. One is not supposed to adapt the theory to fit the original observations; theories are supposed to be tested on fresh data. In machine learning, this idea is preserved by keeping distinct training and testing data. However, if data is very expensive to obtain (or requires an experiment to be scheduled at an uncertain future date), is there a way to re-use the old data in a scientifically valid way?

**How can AI methods produce results which are transparent as to how they were obtained, and interpretable within the disciplinary context?**

AI tools are able to produce highly-accurate predictions, but a number of the most powerful AI methods at present operate as 'black boxes'. Once trained, these methods can produce statistically reliable results, but the end-user will not necessarily be able to explain how these results have been generated or what particular features of a case have been important in reaching a final decision.

In some contexts, accuracy alone might be sufficient to make a system useful – filtering telescope observations to identify likely targets for further study, for example. However, the goal of scientific discovery is to understand. Researchers want to know not just what the answer is but why. Are there ways of using AI algorithms that will provide such explanations? In what ways might AI-enabled analysis and hypothesis-led research sit alongside each other in future? How might people work with AI to solve scientific mysteries in the years to come?

**How can research help create more advanced, and more accurate, methods of verifying machine learning systems to increase confidence in their deployment?**

There are also questions about the robustness of current AI tools. Further work on verification and robustness in AI – and new research to create explainable AI systems – could contribute to tackling these issues, giving researchers confidence in the conclusions drawn from AI-enabled analysis. In related discussions, the fields of machine learning and AI are grappling with the challenge of reproducibility, leading to calls – for example – for new requirements to provide information about data collection methods, error rates, computing infrastructure, and more, in order to improve reproducibility of machine learning-enabled papers<sup>19</sup>. What further work is needed to ensure that researchers can be confident in the outcomes of AI-enabled analysis?

19. See, for example, Joelle Pineau's 2018 NeurIPS keynote on reproducibility in deep learning, available at: <https://media.nurips.cc/Conferences/NIPS2018/Slides/jpineau-NeurIPS-dec18-fb.pdf>



## INTEGRATING SCIENTIFIC KNOWLEDGE

**Is there a rigorous way to incorporate existing theory/knowledge into a machine learning algorithm, to constrain the outcomes to scientifically plausible solutions?**

The ‘traditional’ way to apply data science methods is to start from a large data set, and then apply machine learning methods to try to discover patterns that are hidden in the data – without taking into account anything about where the data came from, or current knowledge of the system. But might it be possible to incorporate existing scientific knowledge (for example, in the form of a statistical ‘prior’) so that the discovery process is constrained, in order to produce results which respect what researchers already know about the system. For example, if trying to detect the 3D shape of a protein from image data, could chemical knowledge of how proteins fold be incorporated in the analysis, in order to guide the search?

This supposes that AI is not already discovering and creating new scientific knowledge.

**How can AI be used to actually discover and create new scientific knowledge and understanding, and not just the classification and detection of statistical patterns?**

Is it possible that one day, computational methods will not only discover patterns and unusual events in data, but have enough domain knowledge built in that they can themselves make new scientific breakthroughs? Could they come up with new theories that revolutionise our understanding, and devise novel experiments to test them out? Could they even decide for themselves what the worthwhile scientific questions are? And worthwhile to whom?

## AI and scientific knowledge

AI technologies could support advances across a range of scientific disciplines, and the societal and economic benefits that could follow are significant. At the same time, these technologies could have a disruptive influence on the conduct of science.

In the near term, AI can be applied to existing data analysis processes to enhance pattern recognition and support more sophisticated data analysis. There are already examples of this from across research disciplines and, with further access to advanced data skills and compute power, AI could be a valuable tool for all researchers. This may require changes to the skills compositions in research teams, or new forms of collaboration across teams and between academia and industry that allow both to access the advanced data science skills needed to apply AI and the compute power to build AI systems.

A more sophisticated emerging approach is to build into AI systems scientific knowledge that is already known to influence the phenomena observed in a research discipline – the laws of physics, or molecular interactions in the process of protein folding, for example. Creating such systems requires both deeper research collaborations and advances in AI methods.

AI tools could also play a role in the definition and refinement of scientific models. An area of promise is the field of probabilistic programming (or model-based machine learning), in which scientific models can be expressed as computer programs, generating hypothetical data. This hypothetical data can be compared to experimental data, and the comparison used to update the model, which can then be used to suggest new experiments – running the process of scientific hypothesis refinement and experimental data collection in an AI system<sup>20</sup>.

AI's disruptive potential could, however, extend much further. AI has already produced outputs or actions that seem unconventional or even creative – in AlphaGo's games against Lee Sedol, for example, it produced moves that at first seemed unintuitive to human experts, but which proved pivotal in shaping the outcome of a game, and which have ultimately prompted human players to rethink their strategies<sup>21</sup>. In the longer-term, the analysis provided by AI systems could point to previously unforeseen relationships, or new models of the world that reframe disciplines. Such results could advance the frontiers of science, and revolutionise research in areas from human health to climate and sustainability.

---

20. Ghahramani, Z. (2015) Probabilistic machine learning and artificial intelligence. *Nature* 521:452–459.

21. See, for example: <https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/> and <https://deepmind.com/blog/alphago-zero-learning-scratch/>