

DEMOS

RESEARCH

4.0

INTERIM REPORT

ELLIOT JONES
NICOLINA KALANTERY
BEN GLOVER

OCTOBER 2019

Open Access. Some rights reserved.

As the publisher of this work, Demos wants to encourage the circulation of our work as widely as possible while retaining the copyright. We therefore have an open access policy which enables anyone to access our content online without charge.

Anyone can download, save, perform or distribute this work in any format, including translation, without written permission. This is subject to the terms of the Demos licence found at the back of this publication. Its main conditions are:

- Demos and the author(s) are credited
- This summary and the address www.demos.co.uk are displayed
- The text is not altered and is used in full
- The work is not resold
- A copy of the work or link to its use online is sent to Demos.

You are welcome to ask for permission to use this work for purposes other than those covered by the licence. Demos gratefully acknowledges the work of Creative Commons in inspiring our approach to copyright. To find out more go to

www.creativecommons.org



Commissioned by



Published by Demos September 2019

© Demos. Some rights reserved.

76 Vincent Square, London, SW1P 2PD

T: 020 3878 3955

hello@demos.co.uk

www.demos.co.uk

CONTENTS

EXECUTIVE SUMMARY	PAGE 4
INTRODUCTION	PAGE 5
SECTION ONE: TOOLS AND TECHNIQUES	PAGE 8
SECTION TWO: BROADER EFFECTS OF TECHNOLOGY ON THE RESEARCH PROCESS	PAGE 13
SECTION THREE: REPLACEMENT OF THE RESEARCHER	PAGE 15
SECTION FOUR: CHALLENGES	PAGE 17
NEXT STEPS	PAGE 19
APPENDIX	PAGE 21
REFERENCES	PAGE 23

EXECUTIVE SUMMARY

The world of research is rapidly changing. The past 20 years have seen huge developments in the way research is conducted, with many research tasks dramatically changed. There is a growing consensus that we are at the start of a 4th industrial revolution, with the rise of the Internet of Things, 3-D printing, nanotechnology, biotechnology, 5G, new forms of energy storage and quantum computing.²

It is evident that research has been changed by these new technologies and in particular the growing use of machine learning and other artificial intelligence techniques to augment the work of research staff and alter their experience of research. In some cases, it may even replace them altogether. However, the rapid pace at which AI is developing means that there are many unanswered questions about exactly how technology is changing research. What will the long-term effects be? How should policy makers approach the increasing use of AI in research?

This paper describes the progress of AI through history and how it has led researchers to develop a range of techniques including natural language programming and computer vision. We then examine these techniques and their practical application across different scientific fields, and the ways in which this has augmented research processes. This leads us to question the exact role of AI and the extent to which it can replace human researchers. We look at both the challenges posed by the advancement of AI and the challenges that may hinder this progress. Finally, we raise questions to be considered for the future of AI in research and for policy making in the second phase of this research project.

We find that:

- Artificial intelligence and robotics have made significant progress since the Second World War. The exponentially increasing computational capacity of computers has, over the last decade, enabled the use of machine learning and artificial neural networks which have dramatically increased the potential of artificially intelligent systems.
- Across many fields and industries, research is adopting a range of tools and techniques enabled by artificial intelligence, including: natural language processing, computer vision, automated monitoring, automated experiment selection and the prediction of physical systems.
- There is an open question as to whether this automation of research will free up and empower researchers to be more creative and productive, or whether it will instead replace them entirely. This may vary from field to field.
- It is also uncertain how long the current rate of progress in artificial intelligence can keep up due to hardware limitations. Further, the costs involved may be prohibitive for all but the richest countries and companies to deploy these technologies in research.

From this, we have a number of questions about how best to forecast the impact of the convergence of 4th industrial revolution technologies, including artificial intelligence, on research and what policy interventions may be needed to ensure the most beneficial outcomes for researchers, their output, and society at large. We will consider these questions in detail in the second phase of this research project.

INTRODUCTION

There is a growing consensus that we are at the start of a 4th industrial revolution, a period of rapid technological change that will mean sudden changes in society and the economy. This will include a transformational impact on the world of research. Some technologies predicted to be part of the revolution include the Internet of Things, 3-D printing, nanotechnology, biotechnology, 5G, new forms of energy storage and quantum computing.²

One class of 4th industrial revolution technologies that stands out as likely to be particularly impactful is that of artificial intelligence, and deep machine learning in particular, due to their position as a general purpose technology.³ A general purpose technology is one which can be applied across many sectors, can enable other technologies and which is rapidly improving. Previous industrial revolutions have had their own general purpose technologies, e.g. electricity or the internal combustion engine in the 2nd, which have been one of the main reasons why such periods constitute revolutions rather than evolutions.

Artificial intelligence is, broadly, a set of computational technologies that aim to sense, learn and reason about their physical or virtual environment, and take action based on that. While the rate of progress in artificial intelligence has been erratic, there have been significant advances since the Second World War and particularly over the last decade. These advances mean that many industries, including research, need to understand the potential impact this technology will have on how they operate.

Stanford University's One Hundred Year Study on Artificial Intelligence outlines a brief history of how we got to the artificial intelligence of today.⁴ In the 18th century, Thomas Bayes provided a framework for reasoning about the probability of events, and in the 19th, George Boole showed that logical

reasoning could be performed systematically in the same manner as solving a system of equations. At the same time, Charles Babbage was drawing up designs for the first general-purpose computer, the Analytical Engine, and Ada Lovelace was designing the first computer algorithm.

By the 1940s, Alan Turing set out a formal model of computing, even imagining the possibilities of simulating intelligence with computers. In the 1950s, these ideas were brought together to build the first electronic computers and primitive robots. The first use of the term 'artificial intelligence' is often credited to John McCarthy in 1956, with the Dartmouth Summer Research Project on Artificial Intelligence.

From the 1950s to 1970s, the foundations were laid for many tools and techniques to be used in research, which we will explore later in this report. These include: using computer vision to recognise individual letters; natural language processing, a way of processing language as actually used rather than set commands; and artificial neural networks, which attempt to replicate the structure of the brain. However, by the 1980s, these had still not translated into practical applications and interest in artificial intelligence waned.

The latest wave of artificial intelligence development began in the late 1990s. The beginning is perhaps best marked by IBM's Deep Blue becoming the first computer to beat a reigning world chess champion, Garry Kasparov, in 1997.⁵ IBM yet again took the limelight in 2011, winning the quiz show 'Jeopardy!' with its question answering system, Watson. These feats were achieved thanks to improvements in computing hardware and a significant amount of engineering manpower on the particular problems.

Today, artificial intelligence is an increasingly

2 Schwab, 2016
3 Cockburn, Henderson, & Stern, 2018, p. 14
4 Stone et al., 2016
5 Ray, 2018

influential part of our daily lives and is developing at a rapid pace, largely due to the revival of machine learning and neural networks. This is due to an interlocking trifecta of ubiquitous connectivity and development of data-sharing infrastructure; increases in access to computational power via the creation of GPUs and cloud-based computational power⁶; increasing access to large-scale data through the creation of massive labelled data sets, typified by ImageNet⁷, and cloud-based data storage. These approaches, pioneered in the 80s, were only in the 2010s able to start surpassing human performance in certain tasks and be deployed at scale.

Some applications in consumer products have been part of our lives for years. Face filters on Snapchat and Instagram already rely on machine learning models.⁸ Since 2016, Google Translate has been using neural networks to improve translation to and from English and live-translate text.⁹ Facebook uses facial recognition models to suggest who to tag in photos.¹⁰ The last year alone has seen many significant developments in artificial intelligence research:

- Google has developed a system called Duplex which can engage in real-time natural language conversations using a synthesised voice nearly indistinguishable from a human. Currently, Duplex is being used to make reservations and screen calls, but it has obvious wider potential, from replacing call centres to offering medical advice.¹¹
- Deepmind, now famous for AlphaGo, used neural networks to take the top spot in the international protein folding competition. Being able to predict protein folding more accurately will improve our ability to diagnose and develop novel treatments for diseases believed to be caused by misfolded proteins, such as Alzheimer's or cystic fibrosis.¹²
- OpenAI unveiled a predictive text generation model able to output paragraphs of coherent and relevant text from a single sentence prompt. This model's potential to generate large-scale misinformation or abusive content

was significant enough that OpenAI decided to break with norms in artificial intelligence development and limit the model's release due to security and safety concerns.¹³

Experts in artificial intelligence, when surveyed, believe that development has sped up over the course of their careers and gave a 50% chance of unaided machines being able to accomplish every task better and more cheaply than humans within 45 years and a 10% chance of it occurring within nine years.¹⁴

There are many scenarios envisaging what a future with increasingly capable artificial intelligence will look like. The more conservative predictions envision a world where vast swathes of modern work, from truck drivers to lawyers, are replaced by narrow machine learning systems, each individually tuned to vastly outperform humans at a specific problem.¹⁵ More outlandish projections, like those made Robin Hanson, speculate about a world where whole-brain emulations, virtual cell-by-cell copies of the human mind, vastly outnumber the biological and dominate the economy.¹⁶ Nick Bostrom outlines the potential for a world-dominating Superintelligence that goes from non-existent to exceeding humanity in every domain in mere weeks.¹⁷

Whatever you may think of the plausibility of any particular scenario, all of them suggest that the development of artificial intelligence will lead to a radical upheaval across society, from the way we work, to how we interact, to how decisions about our health, education, politics, and more are made.

We expect that artificial intelligence will also transform the world of research. Some of this transformation will come through directly augmenting the ability of researchers, increasing the speed, scale and complexity of analysis they can undertake, by allowing the classification and analysis of large sets of unstructured data or automating elements of an experiment. However, also significant is how artificial intelligence may automate much of the routine and administrative work of research, freeing up researchers to focus more on creativity, or improving networking and

6 Mayo, PUNCHIHEWA, Emile, & Morrison, 2018

7 Gershgorn, 2017

8 Le, 2018

9 Turovsky, 2016

10 Simonite, 2017

11 Leviathan & Matias, 2018

12 Evans et al., 2018

13 Wu, Amodei, Amodei, Clark, & Brundage, 2019

14 Grace, Salvatier, Dafoe, Zhang, & Evans, 2017

15 Frey & Osborne, 2017; Susskind & Susskind, 2017

16 Hanson, 2016

17 Bostrom, 2016

recommendations so that researchers can find their perfect collaborators much more easily. Eventually, we might even see some research staff completely displaced by artificial intelligence and needing to find new roles in the research ecosystem.

This report considers how the speed and scale that artificial intelligence and robotics offer is likely to reshape our approach to research, primarily through a comprehensive review of the relevant academic and industry literature. The second phase of this research project will then consider how best to forecast the impact of the convergence of 4th industrial revolution technologies, including artificial intelligence, on research and what policy interventions may be needed to ensure the most beneficial outcomes for researchers and wider society.

SECTION 1 TOOLS AND TECHNIQUES

Researchers have developed a range of AI methods to augment their research practices. We look at the most prevalent systems and their practical applications across many fields, from biomedicine to astronomy. AI tools and techniques are shown to enable researchers to conduct research more accurately and efficiently.

This chapter provides an overview of the following tools and techniques:

- Natural Language Processing
- Image Classification
- Object detection
- Automated data monitoring
- Automated experiment selection
- Simulation and prediction of physical systems
- Parallelisation
- Content generation

Further examples of uses for these tools and technologies can be found in the appendix of this report.

NATURAL LANGUAGE PROCESSING

Natural Language Processing (NLP) is the processing and analysis of unstructured language data, essentially enabling computers to understand human language.

A significant amount of scientific knowledge is increasingly stored in structured databases, e.g. molecular structures, properties of materials. However, many theoretical insights and the majority of the work in the arts and humanities are still stored in book and journal articles, comprised of natural language. In 2009, King et

al.¹⁸ observed that computers were changing the way that research was described and reported. Developments such as the Semantic Web and controlled ontologies were enabling the progress from the use of natural language to more formal language.

However, recent years have seen advancements in natural language processing which is enabling researchers to understand and synthesise at scale research not necessarily written in formal language. Researchers at the Lawrence Berkeley National Laboratory are using the Word2Vec algorithm to sift through scientific papers for connections humans have missed in finding potential thermoelectric materials. After collecting a vocabulary of 500,000 words from millions of abstracts relating to material science, the abstracts were fed to the algorithm, which uses machine learning to analyse the relationships between words. The algorithm can thus link words together and define concepts. It has been able to predict an effective material by scanning scientific papers from years before it was discovered. The tool can be retrained on the literature of any subject and easily be made applicable to other disciplines.¹⁹

Researchers at MIT have similarly developed an NLP tool that can read scientific papers and produce a short summary in plain English.²⁰ This can be used by researchers to scan a large number of papers and get an understanding of what they say. The developers' approach also has the potential to be applied to machine translation and speech recognition.

However, NLP is not confined to the sciences - researchers in the arts and humanities are increasingly making use of NLP tools to advance their inquiries. Researchers at Lancaster University have used Tagtog, an AI platform that utilises NLP and machine learning, to annotate and extract

18
19
20

R. D. King et al., 2009
Tshitoyan et al., 2019
Dangovski, Jing, Nakov, Tatalović, & Soljačić, 2019

information from historical documents that relate to early colonial Mexico.²¹ A series of questionnaires conducted in the late 16th century, known as 'The Relaciones Geográficas', are the primary source materials for the project. The documents run to thousands of pages, containing rich historical information.

These documents are extremely complex to understand, partly because they are written in Spanish and a number of indigenous languages. Previously it would take scholars years, perhaps decades, to fully understand just a small section. However, the use of computational techniques such as NLP can allow for much quicker analysis, creating the potential for previously unimaginable research opportunities.

NLP is also driving further advances in the field of linguistics. Professor David Bamman, a scholar of Computational Humanities and Social Sciences at UC Berkeley, uses NLP tools to identify sarcasm on Twitter.²² This allows researchers to understand more about what enables sarcasm in conversation, in particular, the importance of interpersonal characteristics.

However, these developments are not limited to academia: some of the most significant developments are to be found in the commercial sector. Uber, for example, has developed the Plato Research Dialogue System. A platform for building, training and deploying conversational AI agents. It acts as a testbed - for Uber and the wider research community - to evaluate new algorithms and quickly prototype and deploy conversational agents, spurring further research advances in this field.²³ The platform has been built to be accessible by both those with a limited technical background and more advanced researchers.

Facebook AI has achieved a number of breakthroughs in NLP using both supervised and semi-supervised learning techniques that assist with translation. For example, Facebook's RoBERTa recently outperformed human baselines in several cases, including English-German translation.²⁴

Easily accessible and effective translation tools could significantly enhance the conduct of research. First, it could allow researchers to access papers not published in their native language, boosting the scope for wider sharing and quicker dissemination

of research findings. Second, it could significantly assist those fields in the arts, humanities and social sciences for which translation is often required; for example, historians examining primary sources in a different language could quickly and cheaply access a translation.

COMPUTER VISION

Computer vision refers to computers' ability to 'see' and understand digital images. It is used to classify images into different categories or to detect and identify objects in an image.

Image Classification

One of the main challenges for image classification is the presence of technical noise that interferes with the desired data in images. Recursion Pharmaceuticals²⁵ uses machine learning to obtain better data for classification. The RxRx1 dataset shows different levels of technical noise between over 50 batches of microscopy images. By comparing the batches, machines learn to distinguish the noise from the underlying biology. Machines are then able to disentangle the technical noise from the biological data so that the images can be accurately classified. Recursion thus expects AI to enhance our understanding of how drugs interact with human cells and is using machine learning on biological datasets to accelerate drug discovery. Machine disentanglement of batch effects is applicable across the field of biology.

In a Science magazine survey of young scientists reflecting on the use of AI in research, Noah F. Greenwald from the Dana-Farber Cancer Institute agrees that a key difficulty in large data sets is separating signal from the noise. In cataloguing the mutations present in a group of cancer patient samples, a huge amount of data is produced, with the goal of identifying the mutations that play an important role in tumour biology. The signal from the mutations must be accurately interpreted. Machine learning identifies patterns within the data to systematically remove the noise and achieve more accurate findings.²⁶

Again, it is important to note that many of the developments relating to image classification have been driven by commercial actors. Google, for example, has developed many of the leading image classification models and is now using these

21 Bellamy, 2019
22 Berkeley School of Information, 2015
23 Papangelis, Wang, Namazifar, & Khatri, 2019
24 Ott et al., 2019
25 RXRX, 2019
26 Musib et al., 2017

technologies to improve our understanding of the brain.

The structure of networks in the brain remains largely unknown due to difficulties in imaging and subsequent 3D reconstruction. However, Google is using its technologies to address this: the brain is imaged in extremely high resolution with a microscope before several image analysis problems are solved to produce a diagram of the imaged brain.²⁷

Object detection

Machine learning is widely used in astronomy to help researchers handle huge data flows. Telescopes can scan millions of stars and generate too much data for astronomers to handle themselves. The next generation of telescopes²⁸ will use AI to automatically detect objects in their own images. Advances in image recognition and faster, less expensive computing power have made algorithms useful to more researchers. AI can also be used to coordinate telescopes, allowing hundreds of telescopes around the world to work together when detecting transient events. It can extract and analyze data much faster and more precisely and can treat data with more consistency than humans are able to.

The Palomar Transient Factory (PTF)²⁹ is a telescope used for time-domain surveys to study how astronomical objects change over time. PTF detects intriguing objects and feeds the information to another dedicated telescope to follow up. It is difficult to maintain human researchers at every hour of the day to monitor the sky for intriguing events, so PTF uses machine learning to identify objects and trigger a follow up in the same night.

In the Science magazine survey, Feng Wang from the Chinese Academy of Forestry states that machine learning and unmanned aerial vehicles are revolutionising environment monitoring. Investigating wildlife and vegetation status would previously require extensive surveys of an area. Now, aerial vehicles can capture high-resolution images of hard-to-reach places, and quickly find vegetation types, area and wildlife count, and activity. These results have proven to be more precise than the estimates made with the traditional method. An example is the Alan Turing Institute's³⁰

use of high-resolution satellite images with machine learning to detect and count seal populations in the Antarctic.

Xubin Pan from the Chinese Academy of Inspection and Quarantine³¹ notes that quick and accurate species identification is important in ecology and ecosystem conservation, but most people do not know enough about species morphology to achieve this. With support from AI, a computer or mobile device can be used to accurately identify a species and to discover invasive or endangered species in real-time. Deepmind³² has developed such a tool using the Snapshot Serengeti dataset, which features millions of photos of around 50 animal species. This trains machine learning models to automatically detect, identify and count animals to help with animal conservation in the Serengeti.

AUTOMATED DATA MONITORING

The Internet of Things (IoT) is the network of virtually connected devices, including lab equipment and sensors. It can quickly provide solutions to problems that researchers would otherwise need to address by directly monitoring experiments for an extended period of time.

For example, IoT devices were used to find out why a high performance liquid chromatography machine at a Sunovion Pharmaceuticals³³ lab was giving erratic readings for a drug analysis. The researchers placed a temperature sensor next to the machine which connected to the cloud, and an electronic lab notebook recorded the temperature readings over a few days. This revealed that the building's climate control system had been blowing hot and cold air at the machine at a specific time of day, and the researchers were able to fix the problem. The IoT model can also be used to automatically deposit large amounts of data from devices into the cloud, and to control experiments remotely.

Jeffrey M Perkel's article³⁴ focuses on the IoT's role in monitoring freezers, ensuring that a lab's precious research samples can be looked after remotely. This can give researchers peace of mind and restore their work-life balance. An issue with the IoT may be that its automated data gathering overlooks nuances between disciplines of how data should be collected. There are also security and privacy issues to consider, as cloud provider companies such as

27 Google AI, 2019
28 Snyder, 2018
29 Musib et al., 2017
30 The Alan Turing Institute & The Royal Society, 2019
31 Musib et al., 2017
32 Jean-baptiste Alayrac et al., 2019
33 Olena, 2018
34 Perkel, 2017

Google and Amazon could have access to the vast amounts of research data stored in the cloud.

One of Harrer et al.'s³⁵ main causes of drug trial failure is the inability to monitor patients effectively during trials. In order to be monitored, patients are required to keep a record of their medication intake and bodily responses themselves. This is laborious and can often result in patients dropping out of a trial. To solve this, wearable sensors and video technology can be used along with machine learning and deep learning to record patient data in real time and automate the monitoring process. Through this, AI has the potential to increase trial success rates, reducing drug development costs and putting more drugs on the market.

AUTOMATED EXPERIMENT SELECTION

Previously, the time and cost involved in carrying out an experiment meant that researchers had to be fairly confident they would get a useful result before dedicating the resources. Machines can use their data input to select the most promising experiments to perform for optimum productivity and cost efficiency.

Naik et al.³⁶ have utilised the 'active learning' method of machine learning in drug screening experiments. Not all experimentation options can be carried out due to cost and time restraints, raising the problem of which experiments to choose. In the active learning method, a computer selects which experiments to perform using a predictive model based on simulations of the experiments. Naik et al. found that machines learned to give accurate predictions more rapidly when the machine selected experiments using its own algorithm. Using active learners thus enables researchers to carry out more useful experiments and make accurate predictions in less time.

The Alan Turing Institute³⁷ also aims to decrease research costs and improve outcomes by using existing datasets to choose which experiments to perform. New techniques will be applied to plant science, researching how plants integrate environmental signals to determine how much to grow. This might help to predict how plants may respond to climate change.

SIMULATION AND PREDICTION OF PHYSICAL SYSTEMS

AI models can be trained to make data predictions based on the patterns and sequences found in their data input. This can be used to simulate and predict the structure and properties of new physical systems.

Predictive models have proven to be very useful in genomic research. Moritz Gerstung's team³⁸ at the European Bioinformatics Institute uses algorithms to find the best treatment for leukaemia patients. The full bone-marrow transplant treatment has a high risk of complications and should only be given to patients with the deadliest forms of the disease. Gene-sequencing algorithms are therefore used to predict the severity of different forms of leukaemia by comparing DNA variants and their outcomes in patients. Doctors are then able to give better targeted treatments. In general, predicting a drug's effect on cancer cells indicates whether individual patients will benefit from it. Gene-sequencing also highlights the genetic similarities between cancers in different sites (e.g. the breast cancer gene BRCA1 has been found in the prostate) and gives suggestions for how a drug that works for one cancer site may work for another. Drugs will be increasingly focused on the genetics of a cancer rather than its site.

Biologists interested in determining protein shapes are increasingly turning to AI methods. Deepmind's³⁹ AlphaFold is a deep learning approach that focuses on the problem of modelling shapes from scratch. They have achieved a high degree of accuracy in predicting the physical properties of a protein. The ability to predict the shape of a protein is fundamental to understanding its role in the body and is also useful for diagnosing and treating diseases such as Alzheimer's, Parkinson's, Huntington's and cystic fibrosis, which are thought to be caused by misfolded proteins. An understanding of protein folding may also help in protein design, for example for biodegradable enzymes which help reduce pollutants like plastic and oil. Acquiring more knowledge about the shapes of proteins through simulations and models also opens up new possibilities for drug discovery while reducing the costs of experimentation.

The Tokyo Institute of Technology⁴⁰ is aiming to

35 Harrer, Shah, Antony, & Hu, 2019
36 Naik, Kangas, Sullivan, & Murphy, 2016
37 Ezer, 2018
38 Chivers, 2018
39 AlphaFold, 2018
40 Wu, et al., 2019

discover polymers with high thermal conductivity, have utilised the method of ‘transfer learning’, through which machines are able to discover materials with desired properties even from a very small data set. The study used data from PolyInfo, which despite being the largest database of polymers in the world, still has limited data on the heat transfer properties of polymers. Machine learning models were pre-trained on proxy properties with sufficient data and learnt features applicable to the target task. Machines then used the learnt features with the small datasets and were able to make accurate predictions. The technique led to the identification of thousands of promising polymers, representing a breakthrough for fast and cost-effective machine learning methods for material design.

In training a machine learning tool to predict chemical properties, Google⁴¹ has similarly found that models are most useful when they are able to generalise past the invariances in data. The team has therefore developed a model that is able to focus on the underlying symmetries in images.

PARALLELISATION

Parallel computing processes data simultaneously, so that many problems can be solved at the same time. It can divide large problems into smaller problems for simultaneous processing. The training of AI systems can be limited by their ability to process a large data input, but parallelisation allows the system to process and learn from more data in less time.

Amodei, Kaplan and McCandish⁴² examine the interference of technical noise when using AI to process a large amount of data in parallel. Researchers have increasingly been using data parallelism to speed up neural network training, splitting large batches of data across machines. This is used for tasks such as image classification and language modelling. Although the aim is to increase efficiency, the paper argues that larger batches of data are not always more useful if there is significant technical noise. The gradient ‘noise scale’ quantifies the signal-to-noise ratio and predicts the largest useful batch size in a new domain. When the noise scale is small (meaning there is not much variation in the amount of noise among data), looking at a lot of data in parallel becomes redundant. When the noise scale is large,

we can learn a lot from huge batches of data.

CONTENT GENERATION

Thus far we have largely focused on how new technologies and tools can assist researchers with analysis, from recognising objects to analysing ancient documents. But it is also useful to highlight how AI can be used to generate new content itself.

SciNote have developed AI-based software, Manuscript Writer, that provides researchers with a first draft of a manuscript or paper, based on the researcher’s own data.⁴³ Manuscript Writer is able to generate the Introduction, Materials and Methods, Results and References sections of a manuscript; the Discussion section is left for the researcher to complete. SciNote argue that the programme could significantly reduce the time needed to prepare a manuscript, allowing researchers to focus more on analysing data and preparing for the next experiment, not preparing manuscripts. However, it is important to flag that concerns have been raised around plagiarism, claims which SciNOTE reject.⁴⁴

Google’s AI systems have proven their ability to both enhance a low-quality image’s resolution and fill in gaps in that image.⁴⁵ Google researchers trained their system on low-resolution images of celebrities and bedrooms. The AI system works across two fronts. First, the system compares the low-resolution image to high-resolution images, determining what is in the image. Then, researchers use PixelCNN, an AI system, to add pixels to the images based on what it knows from other images. The results from both processes are then combined to produce a final image.

Of course, it is important to note that the images created are estimates by the machine - they are not yet accurate portrayals. Therefore, their usefulness for research purposes at present - for which accuracy is required - is likely to be limited. However, as the model becomes increasingly accurate, such tools may hold great promise for researchers. For example, low-quality photos from the past could be restored to higher quality, assisting historians and other researchers analysing them. It could also perhaps, one day, help art historians restore damaged works of art.

41 Dahl, 2017
42 McCandlish, Kaplan, & Amodei, 2018
43 SciNote, 2017
44 McCook, 2017
45 Burgess, 2017

SECTION 2 BROADER EFFECTS OF TECHNOLOGY ON THE RESEARCH PROCESS

Technological advances are not only changing the way that research is conducted but also how data is captured, shared and evaluated. AI has the potential to augment the research publication process and promote collaboration.

OPEN ACCESS AND DATA-SHARING

An Elsevier⁴⁶ study claims that the ease of data sharing means that open access platforms will thrive in the future. Funders will increasingly promote open access research and push the creation of OA platforms, meaning that that research outputs are available freely, without restrictions like cost barriers or copyright constraints. The study expects journal publication to become less important in determining the career progress of a researcher, and research will be assessed more against societal impact standards.

Elsevier's forecast is for research to increasingly involve broad, international collaboration. Feedback and content will be easy to share online. A challenge is sharing information in a way that everyone involved can understand, as collaborators may come from a wide variety of disciplines. A condition for funding may be publishing findings immediately on an open platform, with a lot of work going into preparing the content for open publication. Another requirement may be making findings accessible to everyone; scientific research will have to be digestible for a non-scientific audience.

The research organisation Nesta⁴⁷ is part of this trend, promoting open data sharing to maximise the public benefit of research. The Innovation Mapping team believes that in order to create a sustainable research system that produces real societal benefits, we need to make use of new data sources in measuring the progress of research. They are developing Rhodonite, a set of network science algorithms that measures combinatorial changes in research topics over time. Nesta aims to join this tool with other research and innovation indicators in an open source software package, which academics could use to compare research methods and results.

RESEARCH FUNDING

AI-assisted technologies could help speed up the bid writing process for researchers. At present these tasks can be extremely time-consuming with significant administrative burdens, taking researchers away from research. However, as discussed in section one, innovations in natural language processing mean machines are able to analyse unstructured data, such as written content for bids, and generate content itself.

The process of bid evaluation could also be substantially sped up. Software can be used to automatically filter out applications that do not have essential criteria or are unfinished.⁴⁸ This can leave more time for tasks that are harder to automate, such as the qualitative review of bids.

46 Mulligan & Herbert, 2019
47 Richardson, 2019
48 Kerriann Strickland, 2018

WizeHive, for example, automates back-end administrative burdens for grant funding organisations, including the collection, review, and selection of applications.⁴⁹ Other software systems, such as Instrumentl, automate the search process for grant seeking organisations by automatically matching organisations with grant opportunities, significantly reducing the amount of time researchers have to spend searching for funding.⁵⁰

These technologies could lead to a radically different funding ecosystem. AI-assisted technology could both write and review bids, substantially reducing the level of human input required.

PEER REVIEW

Today, humans are at the centre of the peer review system for academic papers. It is useful to consider how the opportunities presented by the 4th industrial revolution could make this system more effective.

The current review process can be incredibly time-consuming. However, an automated system that reviews data standards and other methodologically laborious elements of the review process could free up time for other more qualitative tasks, such as ensuring the research sits in the broader context.⁵¹

For example, Elsevier uses AI system StatReviewer to ensure the statistical analysis and methodological underpinnings of a paper are sound.⁵² Automation could also save time by ensuring the expertise of researchers is best matched with particular papers: another time-consuming task for researchers.

Existing technology can often only detect exact matches in text, reducing the technology's ability to spot all instances of plagiarism. However, AI language processing is able to detect whole sentences and paragraphs that sound similar. AI can also detect incomplete reporting, inappropriate statistics and data fabrication. Automating the publishing process can allow researchers to share their findings much quicker.⁵³

Nonetheless, it should be noted that these applications do not come without threat. Indeed, some have questioned whether one day we could have a "peer-less" review.⁵⁴ It should be noted, however, it is unlikely that the peer review process

will ever be entirely automated: the process relies on and is arguably at its best when drawing on human expertise and creativity. Despite this, some see the automation of the peer review process as a threat to the integrity of the research.⁵⁵

SPARC (Scholarly Publishing and Academic Resources Coalition) - an organisation that works towards an open access model for academic publishing - have raised concerns about the potential impact of automation on the journal market. Primarily, they are concerned that journals with greater financial resources may be in a better position to invest in new technologies and radically increase their competitive advantage, attracting the highest calibre academics and, in turn, dominating the market.⁵⁶ Tools or functions such as search and access, streaming and accelerating the editing and peer review, personal recommendations and sharing tools could be invested in and developed by publishers to ensure they have a competitive advantage over cheaper or free journals.⁵⁷

AUTOMATED ACADEMIC OVERSIGHT

These changes would undoubtedly have wider implications for research. The outcomes would depend on how the artificial intelligence is trained, and what it learns from previous data. There is a dangerous potential for misaligned incentives occurring, for example, if a reinforcement learning system's goal is to get published rather than to be correct, for grant funding this could lead to academics focusing on easy research areas that have a shorter time frame, lower running costs or less risk. For peer-review, risks could include falsifying data that is difficult to detect by current statistical tests and ways of gaming the system to get published.

At a macro-level, if poorly executed, AI could degrade the diversity and quality of published research. There is the possibility of an AI arms race, where organisations with investment, skills and resources battle to produce the 'best AI'. Further, the best AI would not necessarily be the one that produced the highest quality and most valuable research, but the one that ensured academics were published or made the research easiest to conduct.

49 WizeHive, 2019
50 Andre, 2018
51 Chen, 2019
52 Heaven, 2018
53 Burley et al., 2017
54 Shanahan, 2016
55 Times Higher Education, 2016
56 Aspesi et al., 2019, p. 20
57 SPARC, 2018

SECTION 3

REPLACEMENT OF THE RESEARCHER

AUGMENTATION VERSUS SUBSTITUTION

The automation of research tasks using AI has vastly improved productivity in scientific research, with research being carried out faster and more accurately than ever before. Further advancements in AI see machines emulating human decision-making abilities, calling into question whether researchers themselves will be replaced by machines.

Heard's article⁵⁸ explains how AI can create 'a fully robotised scientist' as effective as a human researcher. The mechanisation of repetitive tasks has historically allowed for greater productivity and reliability through the removal of human error, and it is increasingly useful to automate the scientific research process to cope with large data problems. The extent to which experiments can be automated has been constrained by the necessity for human intervention at the decision-making level when a future strategy must be defined. Yet, the use of AI can allow a 'robot scientist' to emulate human decision-making abilities. Machine learning can analyse outputs from previous experiments and thus decide inputs for the following experiments. In King et al.'s experiment, the robot scientist 'Adam' achieved comparable predictive accuracy with a group of undergraduate students acting as decision-maker. It outperformed the naive strategies.⁵⁹ The combination of biological laboratory machinery and analytic automation is shown to produce a machine with the ability to carry out experiments like a human researcher.

King⁶⁰ in fact presents the robot scientist AI system as more effective and productive than human

researchers, possessing 'superhuman' reasoning abilities. The robot allows faster scientific discovery by running thousands of hypotheses in parallel, cheaper experimentation by selecting economically efficient experiments, and improved data sharing and reproducibility. Robots are also more easily produced and trained than human researchers and can work constantly without requiring rest. King views the increased productivity of scientific research as AI's main economic importance; research is the biggest driver of economic growth, and faster growth is achieved by more productive research.

In researching biological systems, the complexity of the systems necessitates the recording of experimental metadata in as much detail as possible. King et al.⁶¹ document that in experiments carried out by their robot scientist, metadata is produced as a natural by-product of experimentation. As all experiments are executed automatically by the computer, it is possible to capture and record all parts of the process.

The robot scientist 'Eve'⁶² at the University of Cambridge has aided the discovery of a new anti-malarial drug. Researchers had previously noted that triclosan, a common antibacterial agent, inhibits the growth of the malaria parasite during the blood stage. Eve was then used in a high throughput experiment, finding that triclosan inhibits the malaria parasite enzyme Plasmodium, and may be able to target malaria at both the blood and the earlier liver stage of the disease. Eve was developed to speed up the drug discovery process by automatically developing and testing

58 Heard, 2004

59 R. King, 2018

60 R. King, 2018

61 R. D. King et al., 2009

62 Bilisland et al., 2018; University of Cambridge, 2018

hypotheses to explain observations.

Reproducibility is key for the credibility and usefulness of scientific research, yet a 2016 survey of 1500 researchers revealed that 70% could not reproduce their colleagues' findings. Commercial contract research organisations (CROs) such as Transcriptic and Emerald Cloud Lab offer a solution with robotic labs. Robotic labs contain and automate the entire experimental process, with robot scientists remotely controlled by researchers to carry out all tasks. Up to 89% of biomedical papers use methods supported by biomedical labs, particularly for basic research.⁶³

These labs have the potential to increase consistency in experimentation. Robotic arms' capacity for manual dexterity and fine motor control is improving and the arms can perform movements much faster and in exactly the same way each time, thus reducing potential variation between trials in an experiment. They could also offer more detailed logs of exactly how each step of the experiment was carried out. Combined, these could improve transparency and reduce issues of quality control and thus have the potential to improve reproducibility. However, if the experimental robotics are controlled by machine learning systems, they could become even more closed black boxes, as we currently lack the ability to interrogate the reasoning or causal links behind current machine learning.

Robotics also offers the ability to undertake experiments that humans' biological limits prevent us from currently undertaking. With sufficient precision and miniaturisation, they could help undertake nano-scale experiments. Further, if these robots can be taken out of the lab and either be remote controlled or embodied with sufficient decision-making capabilities, they could offer the possibility to undertake research in places humans can't go, e.g. at the bottom of the ocean, in high levels of radiation or in deep space.

Kitano⁶⁴ has high expectations for the future of AI, predicting major discoveries in biomedicine that are 'worthy of a Nobel Prize'. Many researchers question whether AI has the 'intuition' to ask the right questions leading to major discoveries. But Kitano points out that asking the 'right questions' is considered important because of limitations on resources, as experiments are subject to budgets and time restrictions. It is therefore, in fact, a

question of efficiency, and as AI drastically improves efficiency, asking the right questions becomes irrelevant.

AI has the potential to achieve breakthroughs simply through the 'brute-force approach' of generating and verifying as many hypotheses as possible. Automated experiment selection systems such as Naik et al.'s 'active learning' method show that AI can close the experimentation loop and is able to make decisions on the course of research. It need not make discoveries in the same way that human researchers would. Kitano calls for a scientific revolution on the scale of the industrial revolution and predicts that AI will have major implications for medicine and human lives.⁶⁵

A 2016 Stanford University study asserts that the measure of success for AI is the value created for human lives, so AI will only be successful if humans are able to use and adapt to the technology.⁶⁶ Hence the advancement of AI is more likely to occur in gradual developments rather than unexpected jumps in techniques.

63 Groth & Cox, 2017
64 Kitano, 2016
65 Naik et al., 2016
66 Stone et al., 2016

SECTION 4

CHALLENGES

CHALLENGES POSED BY THE ADVANCEMENT OF AUTOMATED RESEARCH

In a 2018 survey of AI researchers, 50% forecasted that high-level machine intelligence (HLMI) would be achieved within 45 years. HLMI is achieved when machines can accomplish every task better and more cheaply than human workers. HLMI has the potential to spark explosive technological progress, with AI systems quickly becoming vastly superior to humans in all tasks. The survey also acknowledges the risks associated with this; researchers gave a 15% probability to a bad or extremely bad outcome (e.g. 'human extinction').⁶⁷

Deep-learning algorithms allow humans to take a step back in analysing information, as computers are given the task of finding meaningful relationships in the data. Sarah Webb notes that as researchers become more distant from the analysis, they lose control of the classification process and may not be able to explain exactly what an algorithm is doing; the systems are black-boxed.⁶⁸ While an algorithm is able to make outstanding predictions, it is still a challenge for the programmer to understand how it came to these predictions. This echoes Chomsky's 2011 critique of machine-learning methods: approximating unanalysed data does not constitute scientific success, as while we may be able to predict a phenomenon, this does not necessarily mean we understand it.⁶⁹

Davide Castelvecchi discusses the black box problem further.⁷⁰ Computers can pick up on trends and information that humans are unable to see, and programmers must do the equivalent of neuroscience to understand how their AI models have reached their answers. In the instance of deep learning being used to identify breast cancer in mammograms, a machine may detect cancer in

a woman that appears healthy to doctors. In this instance, even if the computer's predictions are shown to be accurate, doctors must decide what course to take without understanding exactly why the woman is at risk.

De Ridder questions the concern that the automation of research roles will make researchers 'mere servants to the robotic overlords'.⁷¹ He concludes that the increased speed and scope of experimentation will, in fact give researchers more time and resources, allowing them to have more freedom and creativity in their ideas.

Wider data sharing has been noted as a benefit of AI in research. However, a paper by The Royal Society questions the viability of 'open data', when in practice modern data sets are often too large and complex for people to access or understand.⁷² As the size of data sets grows, there will be few researchers able to download them. The amount of data used in and produced by AI research requires a physical infrastructure with sufficient computing power. AI may therefore not be the best solution to accessibility issues in research.

CHALLENGES TO THE ADVANCEMENT OF AUTOMATED RESEARCH

The significant challenge for any AI model is that the results given by algorithms are only as good as the quality of the data input. Particularly for deep learning, a very high quantity and quality of data is required. While automation is often seen to remove human bias and preconceptions from research, the training data for algorithms can be skewed, for example, if only data from northern Europeans is used in a genomic study. This creates embedded biases in the model, which will be reflected in predictions. Casey Greene, computational biologist

67 Grace et al., 2017
68 Webb, 2018
69 Ridder, 2018
70 Castelvecchi, 2016
71 Ridder, 2018
72 The Alan Turing Institute & The Royal Society, 2019

at the University of Pennsylvania, suggests that computers should therefore not be allowed to make decisions without humans validating their predictions. He states that “thinking of these methods as a way to augment humans is better than thinking of these methods as replacing humans”.⁷³

The progress of automation in social science fields still faces many barriers, for example, the lack of standardisation in research. Social scientists have historically not focused on standardising their findings, instead representing their data differently from study to study, creating difficulties for automated data discovery. Yarkoni et al. suggest establishing field-wide conventions for the basic organisation of data, following the example of the Inter-university Consortium for Political and Social Science, which requires data to be curated to certain standards.⁷⁴ Terms are rarely clearly defined in social sciences in contrast to natural sciences, posing a challenge for language processing tools. The use of controlled vocabularies and ontologies may solve this.

Yarkoni et al. view the impact of automation in research as uneven. Natural sciences have been the most rapid adopters of automation techniques, as their highly-structured data and more formal specification lend themselves more easily to the currently available tools. The social sciences and humanities have adopted automation to a less degree, as techniques developed for natural sciences do not necessarily easily adapt to qualitatively different kinds of investigation.

Finally, a key problem for the future of AI is the constraints on computing power. Amodei and Hernandez find that the main factor determining how much power an AI model has correlates to the amount of computing power used to train it, the “compute”.⁷⁵ The total amount of compute is not the same as compute per model because limits on parallelisation constrain how big a model can be or how effectively it can be trained. The amount of compute that can be effectively used at once to train models has increased by a factor of 10 each year as researchers have found a way to use more chips in parallel or use custom hardware that allows more operations to be performed per second. This has enabled researchers to use increasingly powerful AI models for greater research tasks.

However, Ryan Carey comments in the AI Impacts blog that this trend cannot continue for more than another 10 years due to cost.⁷⁶ Unless the cost of compute drastically decreases, experiments will grow too large to be affordable by anyone but the US or Chinese governments. Moore’s law, which states that the processing power of computers doubles every two years, has become difficult to achieve. The size of transistors in computer chips must continuously decrease so that more of them can be used, but Intel, the world’s largest supplier of computer chips, expected in 2016 that transistors could only shrink for another five years.⁷⁷

Bloom et al. show that this has already put research productivity in decline.⁷⁸ Research efforts are increasing across industries, as the number of researchers now required to double the density of computer chips has increased by a factor of 18 since the conception of Moore’s law in 1965. The research productivity rate should have increased by a similar factor yet Bloom et al. calculate that it has remained stable. With larger research efforts less able to achieve higher levels of productivity, there is less potential for new ideas. This can be expected to impact economic growth, which is reliant on the creation of ideas.

Ben Garfinkel offers an alternate perspective, suggesting that we have so far been underestimating the amount of computing power being used.⁷⁹ This would mean that we have been overestimating the returns it gives us. So, the growth of compute may therefore not contribute as much to the progress of AI as anticipated.

73 Webb, 2018
74 Yarkoni et al., 2019
75 Amodei & Hernandez, 2018
76 Carey, 2018
77 Simonite, 2016
78 Bloom, Jones, Van Reenen, & Webb, 2019
79 Garfinkel, 2018

NEXT STEPS

Clearly, these technologies are going to have a transformational impact on research, though exactly what that transformation will look like and who will benefit is still uncertain.

Going forward, we are going to examine what the future of research might look like as these technologies mature. We will explore potential scenarios, try to forecast what is most likely, and then explore what policy options could lead to the most beneficial outcomes for the field of research and society more widely. Below are sets of open questions we have across forecasting and policy that we will look to explore in the second phase of this research project.

FORECASTING

- What are the barriers to adoption of these technologies in academia?
- Will automation have a greater impact on junior versus senior research staff? What effect will it have on career trajectories in research?
- How does the raw availability of physical resources fit into the automation of experimentation in the natural sciences? Is automated experimentation limited by physical timeframes? Will the simulation and prediction of physical systems overtake tangible experiments?
- What is the role of routine work in the creation of data sources, e.g. digitising documents or labelling images, in the automation of research?
- Will the greatest gains to academic productivity come from the automation of administrative tasks freeing up researchers to focus on their work? What would be the impact of automating grant writing?
- Will being less hands-on with experimentation actually free up time for theorising? Or will being disconnected from trial and error make academics less creative?
- Do universities have sufficient digital infrastructure to utilise these new AI and robotic tools? Will traditional universities be overtaken in research by large technology companies?

- Will the blackboxing of AI experiments limit the development of human scientific knowledge?

POLICY QUESTIONS

Research Finance

- What alternative career paths are created in the research lifecycle by the uptake of automated research practices and how can they be recognised?
- Will academia be able to afford the computational power necessary to compete with technology companies on AI-enabled research? Should governments provide more shared access to computational power?

Research Excellence

- How can AI-assisted peer review help maintain research standards?
- Is the Research Excellence Framework (REF) in its current form suitable for AI-generated papers and research?
- What risks does automated research pose to maintaining excellent research standards?
- Should research evaluation encourage or discourage the use of AI to generate research papers?
- What implications arise from using automated processes in research planning and management? For example, if a funder uses AI-assisted technology to evaluate bids, should researchers have a competing AI in this scenario?

Knowledge Exchange between Academia, Industry and Wider Society

- How can we ensure that the benefits of commercially undertaken research are fully felt by wider society?
- How can we encourage greater collaboration

between academia and technology companies?

- How can AI-assisted peer review reduce the lag time between research and publication?
- How can researchers be supported to move freely between industry and academia, or occupy both sectors simultaneously?
- What are the key lessons academia can learn from industry use of 4th industrial revolution technologies?
- Is there a role for automated research determining the most significant fields of research that can aid significant societal and global challenges?

Expertise and Capabilities Needed and Developed by Research Sector

- What structures, flows and funding instruments need to be available to support automated research? Is a different balance between capital and current funding needed?
- Is training young researchers in computing and coding as well as natural sciences the most effective way to prepare them for an automated future? In the short-term? In the long-term?
- What are the major barriers to UK-based research organisations being able to recruit leading talent to undertake AI-related research?
- How can we sufficiently invest in universities' digital infrastructure to ensure they can undertake world-leading automated research?
- What features must AI-based infrastructure have to promote excellent research with integrity?

Research Integrity and Ethics

- Should we restrict or enable AI autonomous decision making in academic contexts, such as medical treatment?
- Do we need stronger data protection rules for the storage or sharing of human research subjects' data in the cloud?
- How can we ensure that those working in fields related to AI are sufficiently aware of the ethical dilemmas and challenges they may face in their work?

Research Organisational Strategy and Culture

- How can the work of international organisations,

such as the Research Data Alliance, be built on to ensure the international sharing of expertise and learning related to automated research?

- Would establishing universal conventions for recording and sharing data (in particular formats, with controlled ontologies, open access etc) help to make data accessible to more researchers, make results more useful or improve reproducibility rates?
- Is the invisible labour and work of technicians, data labellers and other research support staff being properly appreciated?

APPENDIX

NATURAL LANGUAGE PROCESSING

Linguamatics has been developed as a data mining tool using NLP.⁸⁰ It is used to extract information from databases such as medical records, which typically contain unstructured text and numerical information. The language in medical records can be nuanced; for example, information on patients' histories of smoking may be written as "doesn't smoke" or "quit cigarettes". Numbers can be expressed in different ways; 1kg may also be written as one kilogram. Linguamatics normalises concepts from structured and unstructured text and can recognise different terms that represent the same data. It has been particularly useful in drug development, as a search for certain terms in databases can retrieve specific information to help predict the success of drug trials. NLP allows researchers to obtain existing natural language data for their research that would otherwise be hard to find.

Harrer et al.⁸¹ suggest that NLP can be used in automating clinical drug trials to raise success rates. The second half of the long drug development process for any single drug is dedicated to clinical trials. When a trial fails, the investment into the trial and the preclinical development is lost. One of the main causes of failed trials is suboptimal patient cohort selection. AI approaches like NLP and computer vision algorithms can be used to examine medical records and patient data in different formats, automating the process of patient selection. These methods may be able to find patients who are more capable of responding to treatment and more suitable for a trial.

COMPUTER VISION

According to Ken Dutton-Regester, Stanford researchers have developed a deep learning tool to help doctors distinguish between benign and cancerous tumours, which is able to classify skin cancers as accurately as a panel of dermatologists.⁸² Automated classification has the potential to improve the consistency and specificity of categorisation. He suggests that smartphone applications could be developed to use this algorithm and provide low-cost, universal access to diagnostic care.

The Alan Turing Institute is exploring the use of AI to characterise materials from X-ray imaging. It is common for materials to be analysed by having X-ray beams guided through them, creating a scattering pattern that can show their internal structure. Researchers are usually required to examine the scattering patterns, but there is the potential for machine learning models to analyse the patterns and automatically determine the structural information of the materials.⁸³

Morello et al. present SPINN, a machine learning tool that automates the identification of pulsars (a type of star) in large surveys.⁸⁴ Modern pulsar surveys are too large for visual inspection by humans. Similarly, Lochner et al. have developed a model to automate the classification of supernova images, and Banerji et al. have trained a model to classify different types of galaxies.⁸⁵

SIMULATION AND PREDICTION OF PHYSICAL SYSTEMS

A project by the Alan Turing Institute and the John Innes Centre⁸⁶ is exploring the possibility of

80 Dutton, 2018
81 Harrer et al., 2019
82 Musib et al., 2017
83 The Alan Turing Institute & The Royal Society, 2019
84 Morello et al., 2014
85 Banerji et al., 2010; Lochner, McEwen, Peiris, Lahav, & Winter, 2016
86 The Royal Society and The Alan Turing Institute, 2019

using machine learning to model and predict the biosynthesis process of the triterpene molecule in plants. Triterpenes are a valuable class of natural plant products, used across health, agriculture, and industrial sectors.

Rigoberto Medina Andres mentions in the Science magazine survey that plant biologists use machine learning to predict how proteins interact, which can be used to improve crop yield during drought conditions.⁸⁷

Training AI models to be accurate requires a huge data input that can be difficult for humans to gather. Reyes and Maruyama suggest that data obtained by machine simulations can be used as labelled data input to train other machine learning models. This relies on the accuracy of the simulated data, as well as the machine learning model's ability to generalise past any 'artificial features' in the simulated data, so that it is also able to analyse real world images.⁸⁸ For example, a machine could be trained to detect defects from an image of a material using simulated data of the material structure. Real images are noisy and have limited resolution, and a machine trained on pristine simulated data may not be able to look past these interferences.

87 Musib et al., 2017
88 Reyes & Maruyama, 2019

REFERENCES

- AlphaFold. (2018, December 2). AlphaFold: Using AI for scientific discovery. Retrieved 20 September 2019, from Deepmind website: <https://deepmind.com/blog/article/alphafold>
- Amodei, D., & Hernandez, D. (2018, May 16). AI and Compute. Retrieved 20 September 2019, from OpenAI website: <https://openai.com/blog/ai-and-compute/>
- Andre, L. (2018, September 13). 20 Best Grant Management Software in 2019. Retrieved 23 September 2019, from Financesonline.com website: <https://financesonline.com/grant-management/>
- Aspesi, C., Allen, N., Crow, R., Daugherty, S., Joseph, H., McArthur, J., & Shockey, N. (2019). SPARC Landscape Analysis. Retrieved from <https://sparcopen.org/our-work/landscape-analysis/>
- Banerji, M., Lahav, O., Lintott, C. J., Abdalla, F. B., Schawinski, K., Bamford, S. P., ... Vandenberg, J. (2010). Galaxy Zoo: Reproducing Galaxy Morphologies Via Machine Learning. *Monthly Notices of the Royal Astronomical Society*, 406(1), 342–353. <https://doi.org/10.1111/j.1365-2966.2010.16713.x>
- Bellamy, K. (2019). Workshop—Exploring AI for Humanities Research. Retrieved 20 September 2019, from DECM Project website: <https://www.lancaster.ac.uk/digging-ecm/2019/02/exploring-ai-workshop/>
- Berkeley School of Information. (2015, September 1). New Professor David Bamman is a Scholar of Computational Humanities and Social Sciences. Retrieved 20 September 2019, from UC Berkeley School of Information website: <https://www.ischool.berkeley.edu/news/2015/new-professor-david-bamman-scholar-computational-humanities-and-social-sciences>
- Bilsland, E., Vliet, L. van, Williams, K., Feltham, J., Carrasco, M. P., Fotoran, W. L., ... Oliver, S. G. (2018). Plasmodium dihydrofolate reductase is a second enzyme target for the antimalarial action of triclosan. *Scientific Reports*, 8(1), 1038. <https://doi.org/10.1038/s41598-018-19549-x>
- Bloom, N., Jones, C., Van Reenen, J., & Webb, M. (2019). Are Ideas Getting Harder to Find? <https://doi.org/10.3386/w23782>
- Bostrom, N. (2016). *Superintelligence: Paths, dangers, strategies*. Oxford, United Kingdom ; New York, NY: Oxford University Press.
- Burgess, M. (2017, February 8). Criminals beware! Google's AI can now identify faces from heavily pixellated images. *Wired UK*. Retrieved from <https://www.wired.co.uk/article/google-brain-ai-pictures-blur>
- Burley, R., Moylan, E., DeVoss, C. C., Chabrol, E., Norman, F., Newton, A., ... Culley, T. (2017). What might peer review look like in 2030? (p. 15). Retrieved from BioMed Central/Digital Science website: http://events.biomedcentral.com/wp-content/uploads/2017/04/SpotOn_Report_PeerReview-1.pdf
- Carey, R. (2018, July 10). Interpreting AI Compute Trends. Retrieved 20 September 2019, from AI Impacts website: <https://aiimpacts.org/interpreting-ai-compute-trends/>
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nature News*, 538(7623), 20. <https://doi.org/10.1038/538020a>
- Chen, H.-W. (2019, March 10). The Impact of AI on Peer Review. Retrieved 23 September 2019, from Medium website: <https://medium.com/digital-publishing-strategy/the-impact-of-ai-on-peer-review-628f1318b18c>
- Chivers, T. (2018, October 3). Big data has transformed how we do science. Retrieved

20 September 2019, from World Economic Forum website: <https://www.weforum.org/agenda/2018/10/how-big-data-is-changing-science/>

Cockburn, I. M., Henderson, R., & Stern, S. (2018). The Impact of Artificial Intelligence on Innovation (Working Paper No. 24449). <https://doi.org/10.3386/w24449>

Dahl, G. (2017). Predicting Properties of Molecules with Machine Learning. Retrieved 20 September 2019, from Google AI Blog website: <http://ai.googleblog.com/2017/04/predicting-properties-of-molecules-with.html>

Dangovski, R., Jing, L., Nakov, P., Tatalović, M., & Soljačić, M. (2019). Rotational Unit of Memory: A Novel Representation Unit for RNNs with Scalable Applications. *Transactions of the Association for Computational Linguistics*, 7, 121–138. https://doi.org/10.1162/tacl_a_00258

Dutton, G. (2018, January 1). Big Pharma Reads Big Data, Sees Big Picture. Retrieved 23 September 2019, from GEN - Genetic Engineering and Biotechnology News website: <https://www.genengnews.com/magazine/306/big-pharma-reads-big-data-sees-big-picture/>

Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., ... Senior, A. (2018). De novo structure prediction with deep-learning based scoring.

Ezer, D. (2018). Data-driven experiment design. Retrieved 20 September 2019, from The Alan Turing Institute website: <https://www.turing.ac.uk/research/research-projects/data-driven-experiment-design>

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>

Garfinkel, B. (2018, December 18). Reinterpreting “AI and Compute”. Retrieved 20 September 2019, from AI Impacts website: <https://aiimpacts.org/reinterpreting-ai-and-compute/>

Gershgorn, D. (2017, July 26). The data that transformed AI research—And possibly the world. Retrieved 20 September 2019, from Quartz website: <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/>

Google AI. (2019). Connectomics. Retrieved 20 September 2019, from Google AI website: <https://ai.google/research/teams/perception/connectomics/>

Grace, K., Salvatier, J., Dafoe, A., Zhang, B., & Evans, O. (2017). When Will AI Exceed Human Performance? Evidence from AI Experts. ArXiv:1705.08807 [Cs]. Retrieved from <http://arxiv.org/abs/1705.08807>

Groth, P., & Cox, J. (2017). Indicators for the use of robotic labs in basic biomedical research: A literature analysis. *PeerJ*, 5, e3997. <https://doi.org/10.7717/peerj.3997>

Hanson, R. (2016). *The age of Em: Work, love, and life when robots rule the Earth* (First Edition). Oxford: Oxford University Press.

Harrer, S., Shah, P., Antony, B., & Hu, J. (2019). Artificial Intelligence for Clinical Trial Design. *Trends in Pharmacological Sciences*, 40(8), 577–591. <https://doi.org/10.1016/j.tips.2019.05.005>

Heard, N. A. (2004). Technology in Genetics: Automating the scientific process. *Heredity*, 93(1), 6–7. <https://doi.org/10.1038/sj.hdy.6800463>

Heaven, D. (2018). AI peer reviewers unleashed to ease publishing grind. *Nature*, 563, 609–610. <https://doi.org/10.1038/d41586-018-07245-9>

Jean-baptiste Alayrac, Sam Blackwell, Joao Carreira, Reena Chopra, Sander Dieleman, Brian McWilliams, ... Michal Zielinski. (2019, August 8). Using machine learning to accelerate ecological research. Retrieved 20 September 2019, from Deepmind website: [/blog/article/using-machine-learning-to-accelerate-ecological-research](https://blog.google/using-machine-learning-to-accelerate-ecological-research)

Keriann Strickland. (2018, December 13). 6 Easy Ways to Automate Your Grants Review Process. Retrieved 23 September 2019, from Submittable Blog | Submission Management Software website: <https://blog.submittable.com/easy-ways-to-automate-your-grants-review-process/>

King, R. (2018, March 26). Better, Faster, Smarter? The automation of science. Retrieved 23 September 2019, from The OECD Forum Network website: <https://www.oecd-forum.org/users/91064-ross-king/posts/31642-better-faster-smarter-the-automation-of-science>

King, R. D., Rowland, J., Oliver, S. G., Young, M., Aubrey, W., Byrne, E., ... Clare, A. (2009). The Automation of Science. *Science*, 324(5923), 85–89.

<https://doi.org/10.1126/science.1165620>

Kitano, H. (2016). Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery (p. 11). Retrieved from Association for the Advancement of Artificial Intelligence website: <https://pdfs.semanticscholar.org/a926/9acd143b340071a4562b8194334daf9da4d9.pdf>

Le, J. (2018, July 22). Snapchat's Filters: How computer vision recognizes your face. Retrieved 20 September 2019, from Medium website: <https://medium.com/cracking-the-data-science-interview/snapchats-filters-how-computer-vision-recognizes-your-face-9907d6904b91>

Leviathan, Y., & Matias, Y. (2018, May 8). Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone. Retrieved 20 September 2019, from Google AI Blog website: <http://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>

Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. (2016). PHOTOMETRIC SUPERNOVA CLASSIFICATION WITH MACHINE LEARNING. *The Astrophysical Journal Supplement Series*, 225(2), 31. <https://doi.org/10.3847/0067-0049/225/2/31>

Mayo, H., Punchihewa, H., Emile, J., & Morrison, J. (2018). History of Machine Learning. Retrieved 20 September 2019, from AI in Radiology website: <https://www.doc.ic.ac.uk/~jce317/history-machine-learning.html>

McCandlish, S., Kaplan, J., & Amodei, D. (2018, December 14). How AI Training Scales. Retrieved 20 September 2019, from OpenAI website: <https://openai.com/blog/science-of-ai/>

McCook, A. A. (2017, November 9). Newly released AI software writes papers for you—What could go wrong? Retrieved 23 September 2019, from Retraction Watch website: <https://retractionwatch.com/2017/11/09/newly-released-ai-software-writes-papers-go-wrong/>

Morello, V., Barr, E. D., Bailes, M., Flynn, C. M., Keane, E. F., & van Straten, W. (2014). SPINN: A straightforward machine learning solution to the pulsar candidate selection problem. *Monthly Notices of the Royal Astronomical Society*, 443(2), 1651–1662. <https://doi.org/10.1093/mnras/stu1188>

Mulligan, A., & Herbert, R. (2019). Report: What

will the world of research look like 10 years from now? Retrieved from Elsevier/Ipsos MORI website: <https://www.elsevier.com/connect/elsevier-research-futures-report>

Musib, M., Wang, F., Tarselli, M. A., Yoho, R., Yu, K.-H., Andrés, R. M., ... Sharafeldin, I. M. (2017). Artificial intelligence in research. *Science*, 357, 28–30. <https://doi.org/10.1126/science.357.6346.28>

Naik, A. W., Kangas, J. D., Sullivan, D. P., & Murphy, R. F. (2016). Active machine learning-driven experimentation to determine compound effects on protein patterns. *ELife*, 5, e10047. <https://doi.org/10.7554/eLife.10047>

Olena, A. (2018, June 1). Bringing the Internet of Things into the Lab. Retrieved 20 September 2019, from The Scientist Magazine® website: <https://www.the-scientist.com/bio-business/bringing-the-internet-of-things-into-the-lab-64265>

Ott, M., Auli, M., Ng, N., Edunov, S., Levy, O., Singh, A., & Stoyanov, V. (2019, August 14). New advances in natural language processing to better connect people. Retrieved 20 September 2019, from <https://ai.facebook.com/blog/new-advances-in-natural-language-processing-to-better-connect-people/>

Papangelis, A., Wang, Y.-C., Namazifar, M., & Khatri, C. (2019, July 16). Introducing the Plato Research Dialogue System: A Flexible Conversational AI Platform. Retrieved 20 September 2019, from Uber Engineering Blog website: <https://eng.uber.com/plato-research-dialogue-system/>

Perkel, J. M. (2017). The Internet of Things comes to the lab. *Nature News*, 542(7639), 125. <https://doi.org/10.1038/542125a>

Ray, S. (2018, December 1). History of AI. Retrieved 20 September 2019, from Medium website: <https://towardsdatascience.com/history-of-ai-484a86fc16ef>

Reyes, K. G., & Maruyama, B. (2019). The machine learning revolution in materials? *MRS Bulletin*, 44(7), 530–537. <https://doi.org/10.1557/mrs.2019.153>

Richardson, G. (2019, August 7). Towards an open science of science. Retrieved 23 September 2019, from Nesta website: <https://www.nesta.org.uk/blog/towards-an-open-science-science/>

Ridder, D. de. (2018). Artificial intelligence in the lab: Ask not what your computer can do for you. *Microbial Biotechnology*, 12(1), 38–40. <https://doi.org/10.1111/1751-7915.13317>

- RXRX. (2019). RXRX.AI. Retrieved 20 September 2019, from RXRX website: <https://www.rxr.ai/>
- Schwab, K. (2016). *The fourth industrial revolution* (First U.S. edition). New York: Crown Business.
- SciNote. (2017). Manuscript Writer Generates Draft of Your Scientific Manuscript. Retrieved 20 September 2019, from SciNote website: <https://scinote.net/manuscript-writer/>
- Shanahan, D. (2016, May 23). A peerless review? Automating methodological and statistical review. Retrieved 23 September 2019, from Research in progress blog website: <https://blogs.biomedcentral.com/bmcblog/2016/05/23/peerless-review-automating-methodological-statistical-review/>
- Simonite, T. (2016, May 13). The foundation of the computing industry's innovation is faltering. What can replace it? Retrieved 20 September 2019, from MIT Technology Review website: <https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/>
- Simonite, T. (2017, December 19). Facebook Can Now Find Your Face, Even When It's Not Tagged. *Wired*. Retrieved from <https://www.wired.com/story/facebook-will-find-your-face-even-when-its-not-tagged/>
- Snyder, A. (2018, January 11). How astronomers are using AI to study the universe—Fast. Retrieved 20 September 2019, from Axios website: <https://www.axios.com/astronomers-are-using-ai-to-study-the-vast-universe-fast-1515610061-451926e4-02f9-4460-9f68-a7489e7c6577.html>
- SPARC. (2018). SPARC: Setting the Default to Open. Retrieved 23 September 2019, from SPARC website: <https://sparcopen.org/who-we-are/>
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., ... Teller, A. (2016). 2016 Report I One Hundred Year Study on Artificial Intelligence (AI100). Retrieved 20 September 2019, from <https://ai100.stanford.edu/2016-report>
- Susskind, R. E., & Susskind, D. (2017). *The future of the professions: How technology will transform the work of human experts*. Oxford: Oxford University Press.
- The Alan Turing Institute, & The Royal Society. (2019). *The AI Revolution in Scientific Research*. Retrieved from <https://royalsociety.org/-/media/policy/projects/ai-and-society/AI-revolution-in-science.pdf?la=en-GB&hash=5240F21B56364A00053538A0BC29FF5F>
- Times Higher Education. (2016, September 20). Robot-written reviews fool academics. Retrieved 23 September 2019, from Times Higher Education (THE) website: <https://www.timeshighereducation.com/news/robot-written-reviews-fool-academics>
- Tshitoyan, V., Dagdelen, J., Weston, L., Dunn, A., Rong, Z., Kononova, O., ... Jain, A. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571(7763), 95–98. <https://doi.org/10.1038/s41586-019-1335-8>
- Turovsky, B. (2016, November 15). Found in translation: More accurate, fluent sentences in Google Translate. Retrieved 20 September 2019, from Google website: <https://www.blog.google/products/translate/found-translation-more-accurate-fluent-sentences-google-translate/>
- University of Cambridge. (2018, January 18). AI 'scientist' finds that toothpaste ingredient may help fight drug-resistant malaria. Retrieved 23 September 2019, from University of Cambridge website: <https://www.cam.ac.uk/research/news/ai-scientist-finds-that-toothpaste-ingredient-may-help-fight-drug-resistant-malaria>
- Webb, S. (2018). Deep learning for biology. *Nature*, 554, 555–557. <https://doi.org/10.1038/d41586-018-02174-z>
- WizeHive. (2019). Awards Management Software for Submissions & Judging. Retrieved 23 September 2019, from <https://www.wizehive.com/awards>
- Wu, J., Amodei, D., Amodei, D., Clark, J., & Brundage, M. (2019, February 14). Better Language Models and Their Implications. Retrieved 20 September 2019, from OpenAI website: <https://openai.com/blog/better-language-models/>
- Wu, S., Kondo, Y., Kakimoto, M., Yang, B., Yamada, H., Kuwajima, I., ... Yoshida, R. (2019). Successful application of machine learning in the discovery of new polymers. Retrieved 20 September 2019, from ScienceDaily website: <https://www.sciencedaily.com/releases/2019/07/190719102133.htm>
- Yarkoni, T., Eckles, D., Heathers, J., Levenstein, M., Smaldino, P., & Lane, J. I. (2019). Enhancing and accelerating social science via automation: Challenges and opportunities [Preprint]. <https://doi.org/10.31235/osf.io/vncwe>

Licence to publish

The work (as defined below) is provided under the terms of this licence ('licence'). The work is protected by copyright and/or other applicable law. Any use of the work other than as authorized under this licence is prohibited. By exercising any rights to the work provided here, you accept and agree to be bound by the terms of this licence. Demos grants you the rights contained here in consideration of your acceptance of such terms and conditions.

1 Definitions

a 'Collective Work' means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this License.

b 'Derivative Work' means a work based upon the Work or upon the Work and other pre-existing works, such as a musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work or a translation from English into another language will not be considered a Derivative Work for the purpose of this License.

c 'Licensor' means the individual or entity that offers the Work under the terms of this License.

d 'Original Author' means the individual or entity who created the Work.

e 'Work' means the copyrightable work of authorship offered under the terms of this License.

f 'You' means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from Demos to exercise rights under this License despite a previous violation.

2 Fair Use Rights

Nothing in this licence is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3 License Grant

Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) licence to exercise the rights in the Work as stated below:

a to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;

b to distribute copies or phono-records of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works; The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensor are hereby reserved.

4 Restrictions

The licence granted in Section 3 above is expressly made subject to and limited by the following restrictions:

a You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this License, and You must include a copy of, or the Uniform Resource Identifier for, this License with every copy or phono-record of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this License or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this License Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this License. If You create a Collective Work, upon notice from any Licensor You must, to the extent practicable, remove from the Collective Work any reference to such Licensor or the Original Author, as requested.

b You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.

c If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Collective Works, you must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied. Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5 Representations, Warranties and Disclaimer

a By offering the Work for public release under this License, Licensor represents and warrants that, to the best of Licensor's knowledge after reasonable inquiry:

i Licensor has secured all rights in the Work necessary to grant the licence rights hereunder

and to permit the lawful exercise of the rights granted hereunder without You having any obligation to pay any royalties, compulsory licence fees, residuals or any other payments;

ii The Work does not infringe the copyright, trademark, publicity rights, common law rights or any other right of any third party or constitute defamation, invasion of privacy or other tortious injury to any third party.

b Except as expressly stated in this licence or otherwise agreed in writing or required by applicable law, the work is licenced on an 'as is' basis, without warranties of any kind, either express or implied including, without limitation, any warranties regarding the contents or accuracy of the work.

6 Limitation on Liability

Except to the extent required by applicable law, and except for damages arising from liability to a third party resulting from breach of the warranties in section 5, in no event will licensor be liable to you on any legal theory for any special, incidental, consequential, punitive or exemplary damages arising out of this licence or the use of the work, even if licensor has been advised of the possibility of such damages.

7 Termination

a This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Collective Works from You under this License, however, will not have their licences terminated provided such individuals or entities remain in full compliance with those licences. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

b Subject to the above terms and conditions, the licence granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different licence terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other licence that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8 Miscellaneous

a Each time You distribute or publicly digitally perform the Work or a Collective Work, Demos offers to the recipient a licence to the Work on the same terms and conditions as the licence granted to You under this License.

b If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

c No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

d This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of Demos and You.

DEMOS

PUBLISHED BY DEMOS OCTOBER 2019

© DEMOS. SOME RIGHTS RESERVED.

76 VINCENT SQUARE, LONDON, SW1P 2PD

T: 020 3878 3955

HELLO@DEMOS.CO.UK

WWW.DEMOS.CO.UK