

Article

The Challenges of Artificial Intelligence: Britain's Bright Future

By Squadron Leader Carolyn J Swinney

Biography: Squadron Leader Carolyn J Swinney received a BEng (hons) degree (first class) in 2007 and a MSc (dist) in Electronics Engineering from the University of Essex in 2013. She graduated as a Communications and Electronics Engineering Officer in the Royal Air Force in 2014. She works within the Air and Space Warfare Centre and is currently working towards a PhD in Electronic Systems Engineering. Her main research interests are signal processing, unmanned aerial vehicles, neural networks, machine learning and cyber security.

Abstract: Artificial Intelligence (AI) has the power to drive dramatic transformations which will undoubtedly lead to major economic and social benefits. We discuss the history, the future and the challenges AI presents in how we implement this technology in a safe, regulated and ethical manner. AI has a history which Britain should be proud of and organisations such as the Alan Turing Institute are paving a way for the UK to remain a centre of expertise for AI across the world, breaking new ground towards the goal of regulated, safe and ethical AI. Without doubt there are still ethical and trust related questions that remain unanswered but as the UK continues to pour more effort into the difficult questions raised by its implementation, the future for AI in the UK remains bright.

Disclaimer: The views expressed are those of the authors concerned, not necessarily the MOD.

The History of AI

Less than ten years after breaking the Nazi Enigma encryption with his code breaking machine 'The Bombe', helping the Allied Forces win World War II, mathematician Alan Turing published the paper 'Computer Machinery and Intelligence'.¹ The 1950s paper really defined the vision for Artificial Intelligence (AI) as we know it today by asking the question 'Can machines think?'. The powerful nature of 'The Bombe', also considered the first electro-mechanical computer, sparked a notion in Turing about whether machines could have intelligence. Even before the war his famous paper 'On Computable Numbers, with an Application to the Entscheidungsproblem'² explained the mathematical concept that computers could be programmed and adapted to different functions depending on their instructions.³ In his 1950's paper 'Computer Machinery and Intelligence' Turing proposed what is now known as the Turing Test, its purpose was to test the intelligence of a machine. The fundamental principle of the test was whether a human could distinguish if they were interacting with another human or a machine, if not then the machine was said to have intelligence. Although contested by philosophers such as John Searle and the famous Chinese Room Argument, the method is still used today as a marker for AI.⁴ It was not only Turing that was progressing the field. In 1950 Claude E. Shannon of Bell Telephone Laboratories published the paper 'Programming a Computer for Playing Chess'.⁶ At this time IBM were developing the first computer that could store information, the 701. After reading Shannon's paper, Arthur Samuel, a researcher at Illinois University, decided to make a program that could play draughts (checkers). Samuel was quoted to say 'I started writing a program for a machine that did not exist, using a set of computer instructions that I dreamed up as they were needed'.⁷ In 1952, he teamed up with IBM and using the 701 Arthur Samuel developed a computer program which could play draughts. This was the first example of a program which could learn.⁸ In 1951 Marvin Minsky and Dean Edmonds proved the first Artificial Neural Network (ANN) with vacuum tubes representing a network with 40 neurons.⁹

The term AI was officially introduced in 1955 by John McCarthy (Dartmouth College) who, along with fellow researchers Minsky (Harvard University), Rochester (IBM), and Shannon (Bell Telephone Laboratories), wrote 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence'.¹⁰ The intention was to bring a group of mathematicians and scientists together for the summer of 1956 to look specifically at the simulation of human intelligence in machines. McCarthy stated 'we will concentrate on a problem of devising a way of programming a calculator to form concepts and to form generalisations'.¹¹ The conference awakened a spirit between a group of experts from different disciplines to pursue the answer of whether machines could think. In 1957 an ANN called Perceptron was developed by Frank Rosenblatt. Rosenblatt based his research on how the eye of a fly worked, as the processing needed to tell a fly to move was done within the eye. *The New York Times* reported that the US Navy 'expected it to be able to walk, talk, see, write, reproduce itself and be conscious of its existence'.¹² The two layer learning network is the oldest neural network still in use today¹³ but for functions such as classification.¹⁴

From that point until the 1980s technology improved to allow more information to be stored in computers and at less financial expense. Newell and Simon developed the General Problem Solver (GPS) in 1959. This was an important step forward because prior to this software was programmed to solve one problem and one problem only. The GPS solved multiple problems using the same algorithm.¹⁵ In this year the term 'machine learning' was coined by Arthur Samuel in a report which talked of a computer one day being able to beat a person at draughts through learning.¹⁶ In 1966 a professor of Massachusetts Institute of Technology (MIT), Joseph Weizenbaum, published a paper called 'ELIZA--A Computer Program For the Study of Natural Language Communication Between Man and Machine'.¹⁷ This was the first example of Natural Language Programming, a program which interacted with a human through an electric typewriter. The program worked by breaking down a user's input into smaller parts and then repeating parts of the text back in a way that allowed the conversation to continue. However, 'Eliza' created some hype with many people believing that she could understand people. It even resulted in some psychotherapists promoting 'Eliza' as a useful tool, not understanding at all how the program worked. It is a useful reminder today that we need to understand how a system or program works before we make claims as to its wider use. Weizenbaum went on to work for social responsibility in these areas as he was so disturbed at how quickly people were willing to put complete faith in a machine.¹⁸

Government agencies such as the Defence Advanced Research Projects Agency (DARPA) began to take an interest in AI around this time. However, there was a vast step between programs such as 'Eliza' and a machine that could truly understand the meaning of words in many combinations. In 1969, Minsky and Papert published a paper called 'Perceptions: An Introduction to Computational Geometry'¹⁹ which concluded that there were serious limitations that existed with simple neural networks. Challenges were becoming more apparent such as the lack of computational power and memory storage for computers to do anything really substantial in terms of AI. By the 1970s the passion and funding diminished in the light of these challenges. In the late 1980s Minsky and Papert came under criticism that their paper had directly contributed to that lack of funding. They defended themselves claiming progress had already dwindled due to a lack of understanding of how neural networks were able to recognise patterns as they did.²⁰ They were not the only ones to doubt the value in pursuing research into AI. In 1973 Brian Flowers, the head of the British Science Research Council, requested a review of AI which was conducted by Professor Sir James Lighthill of Cambridge University. Known as the 'Lighthill Report', Lighthill concluded that 'in no part of the field have discoveries made so far produced the major impact that was then promised'.²¹ The paper directly contributed to the British government restricting support and funding to only four British universities.²²

By the early 1990s there was a realisation of just how many 'rules' would be needed for systems to start to mimic human intelligence. Other challenges remained, there was still a lack of knowledge as to how the machine was actually doing the reasoning, could we really even call it reasoning? This lack of understanding was known as the black box effect.²³ Despite these

challenges, many new achievements were made in the 1990s and 2000s. Speech recognition software was first implemented on the Windows operating system.²⁴ Now residing in the Smithsonian Museum in Washington DC, IBM's chess playing computer program 'Deep Blue' beat the world champion Garry Kasparov in May 1997.²⁵ This fulfilled the famous 1957 prediction made by Herbert Simon that 'within 10 years, a computer would routinely beat the world's best player'.²⁶ Although 20 years late, Simons prediction did come true. It is important to note that the change over that 30 year period was not in the sophistication or suitability of AI algorithms but in the advancements made with computational power and available memory. The technology had finally caught up with the theory.

In 2009, Google developed the first driverless car which went on to pass a driving test in 2014.²⁷ In 2015, Google developed a Deep Learning program called AlphaGo which beat the world champion of a board game called Go. This was significant because Go had 361 possible moves, which was complex compared to chess with just 20 moves.²⁸ It showed just how far the field had come in a relatively short space of time. In 2017, Hanson Robotics in Hong Kong produced Sophia, a social humanoid robot who even presented the Good Morning show with Piers Morgan.²⁹ The potential for Defence was reported in 2017 by Allen and Chan who concluded that AI had the potential to transform modern warfare, to a greater extent than even aircraft, nuclear weapons and computers and even drive change for national security through military, information and economic superiority.³⁰

In 2018, we saw the introduction of Deepfake where a person in a video could be replaced with someone of a similar likeness using a type of Deep Learning called Generative Adversarial Networks.³¹ In 2019 Google and North Western Medicine produced a deep learning algorithm that detects lung cancer with higher accuracy than radiologists using Computer Tomography (CT) scans.³² 2020 introduced Mayflower, the first AI controlled crewless ship to sail the Atlantic Ocean.³³ In 2021 we observed AI play a critical role in the fight against COVID-19, where it has been used in the search for drug treatments, potential vaccines and to predict the spread of the virus.³⁴ It is without a doubt that AI has changed our lives significantly over the last 80 years but it is vital that we understand the limitations that exist with Artificial Narrow Intelligence (ANI) today and consider the challenges AI presents both now and in the future.

Terminology

Artificial Intelligence

AI covers a wide spectrum of applications and technology. It is therefore helpful to break the term down into three categories, ANI, Artificial General Intelligence (AGI) and Artificial Super Intelligence (ASI). Most of the work being conducted around the world today falls into the ANI category. ANI is where the intelligence relates to a specialised problem set, for example driving a car, playing a game of chess or diagnosing a particular medical condition. There is no consciousness or general intelligence as ANI operates within pre-defined parameters, even when it looks like it is more sophisticated. When we think of the algorithms that sit behind Facebook news feed, Google translate and Siri, we are looking at ANI. It is unable to adapt to

its surrounding, even a self-driving car is operating within set parameters, it does not ‘think’.³⁵ That is not to say that ANI does not have the potential to exceed human abilities in specific scenarios, in many cases it already does.

AGI captures the original intent of the work of Alan Turing in his 1950’s paper ‘Computer Machinery and Intelligence’³⁶ where he asked the very important question, ‘Can machines think?’. The development of a program that can control its learning on its own and solve many different complex tasks in many different domains.³⁷ Martin Ford in his 2018 book *Architects of Intelligence* interviewed 23 of the leading researchers in AI today and asked them to guess the year in which we will be at least 50% towards a functioning AGI. The director of engineering for Google, Ray Kurzweil, put forward the year 2029 and Rodney Brooks, co-founder of iRobot guessed 2200. The other 16 guesses lay in between those two extremes.³⁸ The point being, no one really knows.

ASI then is defined by Oxford Philosopher Nick Bostrom as ‘any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest’.³⁹ This is where we get into the realms of humanity changing far beyond how we know it today. But there are many challenges still ahead for AGI and ASI. Modelled on the human brain, we still lack the knowledge to be able to reproduce even simple functions such as sight, let alone complex ones.

Machine Learning

Machine Learning (ML) is a subset of AI as shown in Figure 1. Deep learning is a subfield of ML and ANNs are a type of deep learning algorithm.

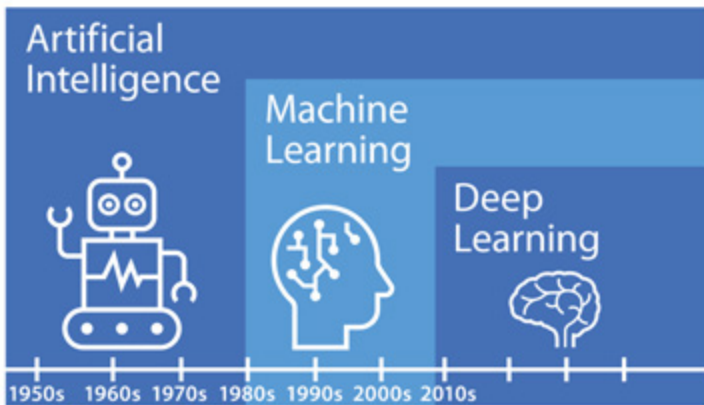


Figure 1 – The difference between AI, ML and Deep Learning

In the last section we discussed ANI and how programs are written to perform a specific task, sometimes better than a human. ML is a technique to achieve AI. Algorithms find patterns in data using statistics, hidden underlying patterns that we would struggle to discover without it. Today ML algorithms are present in all types of applications from ‘Google Search’, to the spam

filter on your email or Netflix recommending a new series for you to watch. It's behind world changing breakthroughs like creating new drugs and detecting cancer. In 1959 Arthur Samuel defined ML as a 'field of study that gives computers the ability to learn without being explicitly programmed'.⁴⁰ What Samuel is really telling us is that the difference between a ML algorithm and a traditional algorithm is down to how we define it.




Figure 2 – The difference between Traditional and ML algorithms

As seen in Fig. 2, in traditional programming, we feed input data into a well-tested program. Then the machine would generate an output. ML algorithms during the learning phase take the data and, depending on the type of ML algorithm used, the output too. The ML algorithm generates the program itself by finding patterns inherent in the data.

Deep Learning

Deep Learning is a subset of ML and the most common type in use today are ANNs inspired by the workings of the human brain. The human brain has around 200 billion neurons with 10,000 synapse connections, each with 1,000 possible states.⁴¹ In Fig.3 we consider an Artificial Neural Network (ANN) which takes an input, for example an image of a car would be represented as a series of pixel values. Then the output of the input layer becomes the input to the next layer, and so on.




Figure 3 – Artificial Neural Network Structure

Each layer performs a different mathematical function on the data. For example, Convolutional Neural Networks (CNNs) have a series of convolutional layers, performing the mathematical operation of convolution. Training data propagates forward through the network and when it gets to the end the error is calculated between the prediction (image is predicted to be a car or not a car) and the actual output (image is actually a car or not a car). A process called back propagation then occurs whereby a weight will be adjusted in each neuron based on the error. As this process repeats with each training cycle, the error is minimised and the algorithm 'learns' by adjusting the weights.⁴² Successfully adjusting the weights requires a lot of data and not enough data can lead to poor performance.⁴³

Big Data

Gartner defines Big Data as 'high-volume, high-velocity and high-variety information assets that demand cost-effective, innovative forms of information processing for enhanced insight and decision making.'⁴⁴ Big Data is not defined by specific metrics such as size, it is defined because the way we process data through traditional methods is no longer fit for purpose. Today we face challenges not only of size but of data being required to be processed at real time speeds and the fusion of data types available through many different sensors. So when big data is considered, we are considering technology that processes and analyses this data in new cost effective and innovative ways.⁴⁵ The RAF, in collaboration with Joint Forces Command, is at the forefront in the development of platforms which can integrate between different operating environments, for example the F-35 Lightning II 5th Generation Fighter Aircraft. The F-35 has a multitude of sensors on board including real time video, active electronically-scanned array radar and infrared search and track functionality.⁴⁶ Effective data analysis could have enormous benefits from efficiency to cost saving and reliability.⁴⁷

Challenges and Limitations

Bias and Variance

Arguably the biggest challenge for AI today is the data. There are two types of error called bias and variance that we must consider. The origins of bias are introduced into the data and can be much more subtle, while variance is inherent noise within the data itself. There is also a trade-off between the two that must be taken into account.

Bias is the difference between the model's prediction and the actual value. High bias is where the model learns the same thing wrong consistently. It oversimplifies and struggles to learn the actual pattern of the data, as seen in Fig.4, which shows the pattern between garden size and house price. Models are therefore consistent but inaccurate, this is also known as under-fitting. Models with high bias have a large error when we train the system and a large error when we give it new test data.⁴⁸ Just like with people, we can be said to have a 'bias' towards someone and if so, we are more likely to make incorrect assumptions based on that bias. Rt. Hon. Jeremy Wright speaking at the 2019 AI Summit stated "The algorithms and structures that govern AI will only be effective if they do not reflect the subconscious biases of the programmers who create them".⁴⁹ Bias can have serious consequences in real life systems

and the origins of where bias comes from can be quite varied, often differing depending on the domain. A researcher from MIT states 'Imagine a scenario in which self-driving cars fail to recognise people of colour as people – and are thus more likely to hit them – because the computers were trained on datasets of photos in which such people were absent or underrepresented'.⁵⁰ If we fail to eliminate bias it could cause serious harm.




Figure 4 – Bias (left), Variance (right)

There are several different types of bias that we should be aware of. The first type to consider is sample bias, when the training data does not match the environment the system needs to run in. For example, training an object detection system during the day and then expecting it to perform at night time. To make a useful model, the training data must match the data from the real world where the system is going to be used as closely as possible or else we may introduce sample bias. This is not a one-time consideration. An important question to examine is how much change a system can cope with before new data is required and re-training of the model so that sample bias doesn't creep in over time.

Exclusion bias occurs when we remove information from the data that we think is not relevant to the output. There are many examples for ML using the Titanic dataset which is used to try and predict who would live and who would die. If we remove the information regarding passenger identities (IDs) from the dataset, thinking it not relevant to the question of who would survive, we would be wrong. Passenger IDs were linked to room numbers which was relevant to how far passengers physically travelled to get to the lifeboats when they needed to evacuate the vessel.⁵¹ So a correlation does exist between passenger ID number and survival. Exclusion bias is especially important when we consider big data, if we are not using all the information then we must ask ourselves whether we are sure the pieces of data we are disregarding are not somehow connected to the output. In some respects the power of ANI models lies in allowing the models to find patterns and connections which a human operator may not see. Next we have Observer bias, when we see what we expect to see, based on our own prior knowledge and experience (conscious or unconscious). A famous example of the adverse effects from observer bias is that of psychologist Cyril Burt in the 1960s. He did a piece of research which included case studies and IQ testing, suggesting that children from low income families had lower intelligence. His findings were later dismissed due to observer bias but his outcomes had already influenced the creation of a two tier educational system in the 1960s.⁵²

Next, we have Prejudice bias. In 2014 Amazon started using ML to review job applications, to help search out new talent. By the end of the year it could be seen that the system did not seem to like women! It turned out that Amazon had used ten years' worth of training data but within that timeframe most of the applications came from white males. Unintentional prejudice bias had entered the system, a true reflection of the industry at the time but not what Amazon wanted for the future.⁵³ So not only must we match the real world as closely as possible with our data, we must also match what we want the real world to look like because the machine will do exactly what we have asked it to do. Lastly, we consider Measurement bias where results can be skewed by the equipment used to measure or observe. We have seen examples of this in video and imagery detection systems whereby a particular camera is used in training and then a different type is used on the real system.

High bias can be mitigated by adding more features to the input data or increasing the complexity of the model. What is the best way to eliminate bias? Through good planning and asking pertinent questions like what do we actually want to achieve? For example, I have a system which can detect GPS jammers using radio frequency signals and I want to use it in a city. If I trained my system in a benign environment, in a field somewhere with nothing else going, it might not perform well in the city where we have numerous other emitters in the same part of the spectrum. What if operationally we cared about mobile jammers, but the system was only trained against static ones and it fails to detect the ones we need it to. These biases could have been prevented through a better understanding of the operational environment which is paramount in the development of these systems. Lieutenant General Mike Groen, Director of the US Joint Artificial Intelligence Center says, 'AI expertise isn't the only kind you need. . . . If you want AI to schedule mid-air refuelling for those aircraft or ammo resupply runs for that artillery, you need experts in logistics'. He goes on to say 'if you don't have a handle on your data, then you can't do AI'.⁵⁴ Planning is key to the elimination of bias and it is an important consideration for the military, to make sure we have the correct operational input.

However, it is important to note that some risk will remain and will need to be explained. For example, we encounter a new type of GPS jammer which we have not seen before so the system was never trained against it. These systems are using the past to predict the future and there will always be a risk that the future will have expanded before we know about it. This is no different to a trained human operator missing a new threat they have not seen before. Another challenge inherently linked to our risk acceptance level is specific to neural networks, the black box effect that we talked about earlier. If we do not understand what an ANN is choosing as features or why, then we cannot be sure that we have not introduced bias into the system which could lead to the system failing in the future. All we can do is our best level of planning to eliminate bias and then we must accept the risk if we are to use the system. We can also not explain the ANN actions if necessary which makes attribution of fault difficult when the system does not operate as intended. To an extent this

can be mitigated by carefully considering the application it's used for and what the risk is if something goes wrong. The advantages of using a model may greatly outweigh any risk but we need people who understand AI and the operational application of a model to inform those decisions.

Variance shows how varied or scattered the predicted results are compared to the actual results, see Fig.4 (right). If we do not have enough data we run a big risk of our system not being grounded in reality.⁵⁵ This is known as overfitting or having high variance, a common problem which will have a programmer believe their classifier is 100% accurate in training. However, when the system is given new data which it hasn't seen before the accuracy drops significantly. The system just got really good, too good at understanding the training data and has incorporated any noise when establishing the pattern with the output. There are techniques such as cross validation which can highlight this problem. It is essential to know that the results we are looking at are accurate and the model is not overfitting. High Variance can be mitigated by adding more input data, having a larger dataset or by reducing the number of features in the input.⁵⁶ For example, if we were looking to predict the selling price of a house, we may include features such as number of bedrooms, whether it has a garden, proximity to schools and the colour shoes I was wearing when I put the house up for sale. A bit of an extreme example but we can have features in our data which are not related to the output and therefore are additional noise.

Giving the algorithm the right data is especially important when we consider ANNs. This is because we cannot always fully understand which bits of the data the ANN is counting as important and using to make patterns.⁵⁷ So if we give the ANN features which are not relevant it could make a correlation to the output which is incorrect. On the flip side of the coin, it could see a correlation which we don't, like we showed with the Titanic dataset with passenger IDs. The difference to note on bias and variance, is that variance is a result of real noise in a dataset, while bias is something which is introduced.

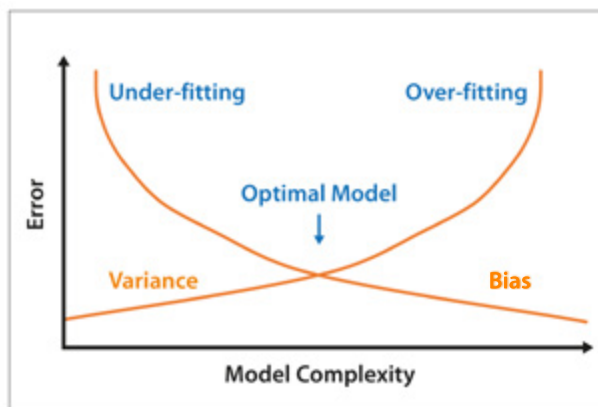


Figure 5 – Bias Variance Trade-Off

The ultimate goal would be a model that displays low bias and low variance but in reality there is a bias variance trade-off as seen in Fig.5. A balance between the two is essential as we cannot be more complex and less complex at the same time, we cannot minimise both bias and variance together.⁵⁸ We do need some noise for our models to function optimally. Neuroscientist Erik Hoel is working on a theory that our brains purposely introduce 'noise' as dreams when we sleep to improve our performance in our waking lives. Dreams are normally warped versions of reality and they create a physiological state in the brain similar to being awake but with the addition of a chemical that causes paralysis. Hoel theorises that the brain faces the same challenge of overfitting, that our training data becomes highly biased and that dreams are 'noise injections' which counter the risk of us learning too narrowly.⁵⁹

Ethics, Safety & Regulations

As this paper is written, both AGI and ASI simply do not exist. The AI that we encounter today is ANI which poses no existential threat to humans.⁶⁰ ANI is described in the book *Megatech: Technology in 2050*, as mindless algorithms with the intelligence of an abacus: that is, zero!⁶¹ However, although not existential, I do believe ANI to pose some threats to humans, not from machines but from ourselves. That could be from poor design (for example not considering the elimination of bias), intentional interference (for example with training data) or the misuse of AI for malicious purposes.

The intentional corruption of training data for malicious purposes is a real concern, RAND have termed it intentional bias.⁶² For example, a ML algorithm that is being used for object detection for a self-driving car would need to recognise road signs. Stop signs are red in colour, with white writing and in the shape of an octagon. Say a hacker infiltrates our training data and puts the upside down smiley face 😄 emoji within the class of images that are 'not stop signs'. Now when this system is in the real world and the hacker gives it an image of a stop sign with an upside down smiley face 😄 emoji, the car will likely drive straight through the stop sign. This is because the algorithm will have learnt a pattern or association between the upside down smiley face emoji and the 'not a stop sign' class.⁶³ So it is of the upmost importance to protect our training data.

We have to remember that ANI is very good at what we have asked it to do but only that and nothing more. When we give it pictures of a dog, it doesn't understand what a dog is, it sees a matrix of pixels with varying intensities, in those intensities it finds a pattern that it then correlates with a dog. There is no understanding of its meaning, there is no intelligence. It does what it has been asked to do. Therefore it is very susceptible to misuse. AI enabled hacking has shown itself to be possible with technology from self-driving cars to drones. These are challenges with physical consequences if they are not addressed soon and bring up serious concerns around trust.⁶⁴ In January 2020 Google's CEO Sundar Pichai made a plea for the regulation of AI, stating that although it was being used for amazing feats such as detecting cancer, the potential for its misuse was great. Examples given include the potential privacy violations and misuse of facial recognition and the need to update regulations for

self-driving cars which will require different governance to that of mechanical cars.⁶⁵ This was followed shortly by the EU in February 2020 releasing a white paper on AI with the intention of regulating transparent and trusted AI.⁶⁶

So we ask the question can regulations help to build trust in AI? Building trust in AI is intrinsically linked to the ethical use of the systems. At the 2019 AI Summit Rt. Hon. Jeremy Wright stated 'Whilst we are optimistic about the potential of AI, it raises many new ethical questions, that would have seemed like issues from science fiction only a few years ago'.⁶⁷ The Oxford dictionary defines ethics as 'moral principles that govern a person's behaviour or the conduct of an activity'.⁶⁸ Western philosophers have debated ethics for hundreds of years. If we consider the work of Aristotle, Kant and Utilitarianism we can attempt to boil down ethics to 'how society applies beliefs and values into short and long-term decision making'.⁶⁹ The Alan Turing Institute defines AI ethics as 'a set of values, principles, and techniques that employ widely accepted standards of right and wrong to guide moral conduct in the development and use of AI technologies'.⁷⁰ AI ethics considers the human behaviour required to minimise harm in society, whether from design failures, misuse or unintended consequences of their application.

AI ethics are needed because there is a potential for harm and we have a responsibility in the design, manufacturing and operating of AI systems. The 2019 Defence Technology Framework puts the onus on Defence to ensure AI developments are 'safe, ethical and interoperable with other nations so that we build consent for, and confidence in, their use'.⁷¹ Professor Lee in his paper 'An Ethics Framework for Autonomous Weapon Systems' explains some of the challenges for the future of Autonomous Weapon Systems, describing each element as hugely complex. In a complex system where machines are making decisions and automating cognitive functions, it can be very hard to attribute a negative consequence. This is made harder in some types of neural networks where the decision-making process is not transparent or easily understood by humans. In 2017, a VW engineer received a 3-year prison sentence for programming VW cars to look like they passed the U.S. environmental standards for diesel emissions when they hadn't. The defence argued that the engineer didn't realise what he was doing, he was just programming what he had been told to do.⁷² So where does the responsibility lie? In criminal law, if a crime has been committed by someone without mental capacity or by an animal, that person or animal could be found innocent but the person who instructed them liable. Would our AI be an innocent party and the programmer held responsible?⁷³ Or the management or the company itself? The answers to these questions need to be carefully considered both nationally and ideally one day be agreed to on the international stage.

The Centre for Data Innovation ran a study in 2019 to determine which nation was leading with AI innovation advantage. They found that the US led in absolute measure, with China second and the EU behind that. It was noted that this could change with China making more rapid progress than both the US and the EU.⁷⁴ It was assessed that one of the contributing

factors to China and the US being so much further ahead than EU countries was due to the release of data. AI needs a lot of data and the report concluded that strict EU data laws would make it very unlikely that a sophisticated AI would ever be produced from within Europe.⁷⁵ This is further highlighted by the 2019 Defence Science and Technology Laboratories (DSTL) paper which spoke of the need for the availability of Defence relevant data for AI systems to be fully realised.⁷⁶

Another ethical consideration is addressed by Dr. Kai-Fu Lee in his book *AI Superpowers*.⁷⁷ He discusses the Chinese adoption of a techno-utilitarian approach that holds a collective good higher than individual rights. Dr. Lee challenges us to consider how the benefits that will come from the introduction of technology such as self-driving cars could outweigh the bad. A 2017 study conducted by the International Transport Forum which includes 57 countries including the UK and the US, found that the benefits of automating road freight transport would result in safer roads, fewer emissions and financial saving.⁷⁸ In the UK alone in 2019, there were 1,752 reported road deaths, 25,945 serious injuries and a total of 153,158 casualties from road traffic collisions (RTCs).⁷⁹ The World Health Organisation (WHO) estimates that RTCs cause the death of approximately 1.35 million people every year.⁸⁰ Self-driving cars could dramatically reduce the number of deaths and injuries a year from RTCs but as the 2017 study found, this would come at a cost of 2 million driver jobs across the US and Europe by 2030.⁸¹ What is of more importance? The greater good or the individual need?

As of today we have not had a machine make an autonomous decision on life and death without a human in the loop. We may consider the future of military operations as the next avenue to face that problem but in reality it will likely be tackled first with transportation. In 2018 the Moral Machine Experiment⁸² was conducted where 40 million responses were evaluated from people in 233 countries spanning ten different languages exploring the moral dilemmas brought to us by autonomous vehicles. The questions posed were those surrounding activity and inactivity leading to the death of many or the death of a few. This dilemma was introduced in the 1970s and is often referred to as the Trolley Problem.⁸³ Although there were substantial cultural variations the results clearly showed that as a collective we want to spare human life where possible, more life rather than less should be spared and the young over the old. Therefore there is hope that we can come to some sort of global consensus about the moral choices machines will make in the future.

However, we cannot deny there is a grave need for the regulation for AI today both for reasons of ethics and trust. Since 2017 there have been over 22 different ethical standards published and many countries investing in national policies and strategies.⁸⁴ In the UK we have invested in four institutions to implement this. The UK AI Council⁸⁵ was set up as a panel of experts which could advise Government of the AI ecosystem. In 2021 they released a roadmap to the development of a National AI Strategy⁸⁶ and the Office for AI was set up to oversee the implementation.⁸⁷ The Centre for Data Ethics and Innovation is a new body to help current regulators achieve the task set out in 2020 by the Committee on Standards in Public Life (CSPL).⁸⁸

The CSPL decided that all regulators in existence in the UK would need to meet AI challenges within their sector, rather than creating an AI-specific regulator.⁸⁹

Lastly, the Alan Turing Institute was announced as the UK's National Centre for AI and data science, bringing together 13 UK universities to research the impact of AI on our society.⁹⁰ The Alan Turing Institute has already developed the FAST framework of Fairness, Accountability, Sustainability, and Transparency for the design of AI systems.⁹¹ While the UK may be more restricted compared to China and the US in terms of data protection laws, there is no doubt that the UK is still regarded as a centre of excellence for AI.⁹² Research firm Conilytica recently reported the UK to have one of the strongest strategies for AI in the world. This included government funding, research activity, AI business startups and enterprise adoption of AI.⁹³ So while there are still many questions surrounding ethics, trust and safety which remain unanswered, the UK are breaking new ground towards the goal of regulated, safe and ethical AI.

The Future

When we look back at the history section we can see a race between the theory to make AI work and the technology needed to enable it. The theory of ANNs have been around since the 1950s. However, we are only just beginning to realise their utility over the last 10 years because we now have the computational power, memory storage and large datasets needed to fully enact them. Today, with the US assessed as leading the way for AI innovation advantage and China fast on their heels,⁹⁴ what is next for AI?

In recent years, the media has given much attention to Futurists who believe in the idea of technological singularity, that there will be a single leap in progress where computers will become more intelligent than humans. There are many opinions of when this could happen, Brynjolfsson et al.⁹⁵ suggest 2045 based on the analysis of current research, expenditure in AI research and development and the growth of AI patents. We cannot deny that history has examples of ground-breaking changes in technology which have dramatically changed the way the human race functioned, for example the industrial revolution or the discovery of fire. Could machines one day be more intelligent than humans? Will there be a single point in time where a giant leap is made? In 2005 Kurzweil published a book called *The Singularity is Near* and he argued that up until 2005 technology growth had been exponential.⁹⁶ In 2019 he still believes this exponential trend to have continued and that it will lead to technological singularity.⁹⁷ However, many have argued that this theory lacks scientific grounding and that exponential growth cannot continue forever, it will level out.⁹⁸

Other arguments include the fact that intelligence is more than an increase in sheer computing power.⁹⁹ Kaplan and Haenlein state with current technology it is unlikely that we would be able to artificially replicate a human brain with 200 billion neurons with 10,000 synapse connections, that each can have a possibility of 1,000 states. There is evidence that the larger the ANN the better it generalises to new information.¹⁰⁰ However, whales and elephants

have over double the neurons in their brains as a human¹⁰¹ which indicates that the human brain operates more efficiently. Our brains are constantly taking in sensory information but only processing a small fraction of the data.¹⁰² How does the brain decide what information to disregard? We talked earlier about the theory of dreams playing a vital role in preventing overfitting, we are still learning about the brain. So without the knowledge of how our brain processes data, would the replication of the human brain leave us with an AGI or even an ASI? Or would our machine be like an aeroplane is to a bird? An aeroplane can fly higher, faster, can carry a lot of cargo in some cases but an aeroplane cannot silently swoop to catch a mouse from the branch of a tree.¹⁰³

Possibly an easier theory to contemplate than that of an ASI is Rushkoff's take on the future: 'Rather than towards machines that think, I believe we are migrating toward a networked environment in which thinking is no longer an individual activity nor bound by time and space.'¹⁰⁴ Google's Ray Kurzweil believes that by 2045 machines will increase human intelligence through wireless links between the neocortex in our brains to a synthetic one in the cloud.¹⁰⁵ So maybe the leap will not be in the creation of a machine that is more intelligent than a human, but in the creation of an environment that will extend human intelligence. Similar to the effect of the introduction of printed press, or medieval universities that brought scholars together or networked computer technology.¹⁰⁶ What type of giant leap in technology would be needed to make human machine networks or an AGI or ASI?

Moore's Law states that classical computing capability can be expected to double every 2 years due to the increasing amount of transistors a microchip can hold. In January 2019 the Chief Executive of Nvidia said that 'Moore's Law, which predicts regular increases in the computing power of silicon chips, is dead.'¹⁰⁷ Although not thought to be a viable commercial solution currently,¹⁰⁸ some researchers and companies are turning their attention to quantum computing.¹⁰⁹ At the Google Quantum Spring Symposium in May 2019, Hartmut Neven, the director of the Quantum AI lab, said that 'quantum computers are gaining computational power relative to classical ones at a doubly exponential rate.'¹¹⁰ If quantum computers keep gaining on classical computers at a 'doubly exponential' rate, now known as Neven's law, then quantum supremacy is very close.¹¹¹ Famous physicist Richard Feynman in 1981 said 'Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy.'¹¹² Many today believe that the implications of Quantum Computing could be the means of unlocking technological singularity.¹¹³

Although we must keep a careful eye on advances in technologies such as quantum computing, we must not forget where we are now and what is achievable today. We started this paper with an explanation of the terminology and the history of AI to convey the certainties of the field. There is a lot of inaccurate information and hype surrounding AI. Zachary Lipton went as far as to name the problem the 'AI misinformation epidemic', stating that the combination of interest and ignorance with a lack of clear and informed voices is

creating a perfect storm.¹⁴ We must learn from the bold statements of the past, the US Navy expecting Perceptron to become conscious of its own existence,¹⁵ the fact that Weizenbaum changed his life calling to warn people of putting their faith in machines such as 'Eliza'.¹⁶ We have come a long way, no one knows how far we have to go, so we must employ a careful balance between the hopeful expectation of the future and the challenges and limitations we face today.

Notes

¹ A. M. Turing, "Computing Machinery and Intelligence," *Mind* LIX, No. 236 (October 1, 1950): 433–60, <https://doi.org/10.1093/mind/lix.236.433>.

² Alan Turing, "ON COMPUTABLE NUMBERS, WITH AN APPLICATION TO THE ENTSCHEIDUNGSPROBLEM," 1936.

³ Dermot Turing, *Alan Turing: The Life of a Genius Paperback*, 2017, 8.

⁴ Smith et al., "The History of Artificial Intelligence."

⁵ Michael Haenlein and Andreas Kaplan, "A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence," *California Management Review* 61, No. 4 (August 17, 2019): 5–14, <https://doi.org/10.1177/0008125619864925>.

⁶ Claude E Shannon, "Programming a Computer for Playing Chess," *Philosophical Magazine Series* 7 41, No. 314 (March 9, 1950), <https://www.pi.infn.it/~carosi/chess/shannon.txt>.

⁷ Arthur L Samuel, "AI and Games Pioneer," *I Programmer*, March 24, 2015, <https://www.i-programmer.info/history/people/669-a-l-samuel-ai-and-games-pioneer.html>.

⁸ History-Computer, "Arthur Samuel - Biography, History and Inventions," History-Computer, accessed January 7, 2021, <https://history-computer.com/ModernComputer/thinkers/Samuel.html>.

⁹ Gil Press, "A Very Short History Of Artificial Intelligence (AI)," *Forbes*, December 30, 2016, <https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/?sh=3e46a67c6fba>.

¹⁰ John McCarthy et al., "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," *AI Magazine* 27, No. 4 (2006): 12–14.

¹¹ McCarthy et al.

¹² Press, "A Very Short History Of Artificial Intelligence (AI)."

¹³ Kate Stachnyi, "Brief History of Neural Networks. Although the Study of the Human Brain... | by Kate Strachnyi | Analytics Vidhya | Medium," *Medium*, January 23, 2019, <https://medium.com/analytics-vidhya/brief-history-of-neural-networks-44c2bf72eec>.

¹⁴ Pratik, Shukla and Roberto Iriondo, "Main Types of Neural Networks and Its Applications," *Towards AI*, July 14, 2020, <https://medium.com/towards-artificial-intelligence/main-types-of-neural-networks-and-its-applications-tutorial-734480d7ec8e>.

¹⁵ A Newell and JC Shaw, "A General Problem Solving Program for a Computer," *International Conference on Information Processing*, 1959.

¹⁶ Samuel, "Some Studies in Machine Learning Using the Game of Checkers."

¹⁷ Joseph Weizenbaum, "ELIZA-A Computer Program For the Study of Natural Language Communication Between Man and Machine," *Communications of the ACM*, Vol. 9, 1966,

<http://i5.nyu.edu/~mm64/x52.9265/january1966.html>.

¹⁸ Martin Campbell-Kelly, "Professor Joseph Weizenbaum: Creator of the 'Eliza' Program," *The Independent*, March 18, 2008, <https://www.independent.co.uk/news/obituaries/professor-joseph-weizenbaum-creator-eliza-program-797162.html>.

¹⁹ Marvin, Minsky and Seymour A. Papert, "Perceptrons - An Introduction to Computational Geometry," *MIT Press*, 1969, <https://mitpress.mit.edu/books/perceptrons>.

²⁰ Rockwell Anyoha, "The History of Artificial Intelligence," *Science in the News*, August 28, 2017, <http://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>.

²¹ James Lighthill, "Lighthill Report," 1973, http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/contents.htm.

²² Smith et al., "The History of Artificial Intelligence."

²³ Council of Europe, "History of Artificial Intelligence," Council of Europe, 2020, <https://www.coe.int/en/web/artificial-intelligence/history-of-ai>.

²⁴ Dragon Systems, "Dragon Systems Unveils Revolutionary Breakthrough with Continuous Speech Recognition," New York Press Release, April 2, 1997, <http://www.voicesystems.com/dragon.htm>.

²⁵ IBM, "Deep Blue," IBM100, 2012, <https://www.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>.

²⁶ Stephen Coles, "Computer Chess: The Drosophila of AI," *Dr. Dobb's*, October 30, 2002, <https://www.drdobbs.com/parallel/computer-chess-the-drosophila-of-ai/184405171>.

²⁷ Press, "A Very Short History Of Artificial Intelligence (AI)."

²⁸ David Silver et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature* 529, No. 7587 (January 28, 2016): 484–89, <https://doi.org/10.1038/nature16961>.

²⁹ Hanson Robotics, "Sophia," Hanson Robotics, 2020, <https://www.hansonrobotics.com/sophia/>.

³⁰ Greg, Allen and Taniel Chan, "Artificial Intelligence and National Security," Belfer Center for Science and International Affairs, July 2017, <https://www.belfercenter.org/publication/artificial-intelligence-and-national-security>.

³¹ Maithreyan Surya, "The Decade of Artificial Intelligence," *Towards Data Science*, December 31, 2019, <https://towardsdatascience.com/the-decade-of-artificial-intelligence-6fcfa2fae473>.

³² Ana Sandoiu, "Artificial Intelligence Better than Humans at Spotting Lung Cancer," *Medical News Today*, June 2019, <https://www.medicalnewstoday.com/articles/325223#1>.

³³ IBM, "Mayflower Then and Now," IBM News Room, 2020, <https://newsroom.ibm.com/then-and-now>.

³⁴ John Terra, "10 Years of Artificial Intelligence and Machine Learning," *Simpli Learn*, September 25, 2020, <https://www.simplilearn.com/ten-years-of-artificial-intelligence-and-machine-learning-article>.

³⁵ Tannya D. Jajal, "Distinguishing between Narrow AI, General AI and Super AI | by Tannya D. Jajal | Mapping Out 2050 | Medium," *Medium*, May 21, 2018, <https://medium.com/mapping-out-2050/distinguishing-between-narrow-ai-general-ai-and-super-ai-a4bc44172e22>.

³⁶ A. M. Turing, "Computing Machinery and Intelligence," *Mind* LIX, no. 236 (October 1, 1950): 433–60, <https://doi.org/10.1093/mind/lix.236.433>.

- ³⁷ Ben Goertzel and Cassio Pennachin, "Artificial General Intelligence," *Artificial Brains*, 1998, https://doi.org/10.1007/978-3-540-68677-4_5.
- ³⁸ Martin Ford, *Architects of Intelligence: The Truth about AI from the People Building It*, Packt Publishing, 1st ed., 2018, <https://www.amazon.co.uk/Architects-Intelligence-truth-people-building-ebook/>.
- ³⁹ Nick Bostrom, "How Long before Superintelligence?," *International Journal of Future Studies* 2 (1998), <https://www.nickbostrom.com/superintelligence.html>.
- ⁴⁰ A. L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development* 44, No. 1–2 (1959): 207–19, <https://doi.org/10.1147/rd.441.0206>.
- ⁴¹ Andreas Kaplan and Michael Haenlein, "Rulers of the World, Unite! The Challenges and Opportunities of Artificial Intelligence," *Business Horizons* 63, No. 1 (January 1, 2019): 37–50, <https://doi.org/10.1016/j.bushor.2019.09.003>.
- ⁴² Anne Bonner, "What Is Deep Learning and How Does It Work?," *Towards Data Science*, September 7, 2019, <https://towardsdatascience.com/what-is-deep-learning-and-how-does-it-work-f7d02aa9d477>.
- ⁴³ Alexandre Gonfalonieri, "5 Ways to Deal with the Lack of Data in Machine Learning," *KD Nuggets*, 2020, <https://www.kdnuggets.com/2019/06/5-ways-lack-data-machine-learning.html>.
- ⁴⁴ Mark Beyer and Douglas Laney, "The Importance of 'Big Data': A Definition," *Gartner*, June 21, 2012, <https://www.gartner.com/en/documents/2057415/the-importance-of-big-data-a-definition>.
- ⁴⁵ Alexandra L'Heureux et al., "Machine Learning with Big Data: Challenges and Approaches," *IEEE Access* 5 (2017): 7776–97, <https://doi.org/10.1109/ACCESS.2017.2696365>.
- ⁴⁶ Lockheed Martin, "F-35-Capabilities," Lockheed Martin, accessed January 10, 2021, <https://www.lockheedmartin.com/en-us/products/f-35/f-35-capabilities.html>.
- ⁴⁷ Robert Wright, "Sensors and Sensibility Keep Today's Jet Planes Working Hard," *Financial Times*, October 18, 2017, <https://www.ft.com/content/c34736de-962e-11e7-8c5c-c8d8fa6961bb>.
- ⁴⁸ Seema Singh, "Understanding the Bias-Variance Tradeoff," *Towards Data Science*, May 21, 2018, <https://towardsdatascience.com/understanding-the-bias-variance-tradeoff-165e6942b229>.
- ⁴⁹ Jeremy Wright, "Jeremy Wright Speaking at the AI Summit," 2019 (London, June 12, 2019), <https://www.gov.uk/government/speeches/jeremy-wright-speaking-at-the-ai-summit>.
- ⁵⁰ Joy Buolamwini, "Biased Algorithms Learn From Biased Data: 3 Kinds Biases Found In AI Datasets," *Forbes*, February 7, 2020, <https://www.forbes.com/sites/cognitiveworld/2020/02/07/biased-algorithms/?sh=24e805fe76fc>.
- ⁵¹ Salma Ghoneim, "5 Types of Bias & How to Eliminate Them in Your Machine Learning Project," *Towards Data Science*, April 16, 2019, <https://towardsdatascience.com/5-types-of-bias-how-to-eliminate-them-in-your-machine-learning-project-75959af9d3a0>.
- ⁵² Sarah Fader, "Can Observer Bias Cause Problems In Relating To Other People?," *Better Help*, February 1, 2020, <https://www.betterhelp.com/advice/behavior/can-observer-bias-cause-problems-in-relating-to-other-people/>.
- ⁵³ Jeffrey Dastin, "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women,"

Reuters, October 11, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scrap-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G?>

⁵⁴ Lt. Gen. Mike Groen, "Building JADC2: Data, AI & Warfighter Insight," *Breaking Defense*, 2021, <https://breakingdefense.com/2021/01/building-jadc2-data-ai-warfighter-insight/>.

⁵⁵ Pedro Domingos, "A Few Useful Things to Know about Machine Learning," 2012.

⁵⁶ Renu Khandelwal, "Bias and Variance in Machine Learning," *Data Driven Investor Medium*, October 28, 2018, <https://medium.com/datadriveninvestor/bias-and-variance-in-machine-learning-51fdd38d1f86>.

⁵⁷ Anyoha, "The History of Artificial Intelligence."

⁵⁸ Singh, "Understanding the Bias-Variance Tradeoff."

⁵⁹ Erik Hoel, "How the Strangeness of Our Dreams Reveals Their True Purpose," *New Scientist*, November 4, 2020, <https://www.newscientist.com/article/mg24833073-600-how-the-strangeness-of-our-dreams-reveals-their-true-purpose/>.

⁶⁰ P.J. Bentley, "The Three Laws of Artificial Intelligence: Dispelling Common Myths," *Institutional Research Information Service*, 2018, <https://www.aitoday.xyz/the-three-laws-of-artificial-intelligence-dispelling-common-myths/>.

⁶¹ D Franklin and L Floridi, "The Ethics of Artificial Intelligence," in *Megatech: Technology in 2050*, Main Editi (Economist Books, 2018), 155–63, <https://www.amazon.co.uk/Megatech-Technology-2050-Daniel-Franklin/>.

⁶² Douglas Yeung, "Intentional Bias Is Another Way Artificial Intelligence Could Hurt Us," *RAND*, October 22, 2018, <https://www.rand.org/blog/2018/10/intentional-bias-is-another-way-artificial-intelligence.html>.

⁶³ Frederik Bussler, "Hacking AI Is Surprisingly Easy. What Hacking Neural Networks in Just... | by Frederik Bussler | The Startup | Medium," *The Startup*, August 19, 2020, <https://medium.com/swlh/hacking-ai-is-surprisingly-easy-53538612954b>.

⁶⁴ D Palmer, "Artificial Intelligence Could Be Used to Hack Connected Cars, Drones Warn Security Experts," *ZDNet*, November 20, 2020, <https://www.zdnet.com/article/artificial-intelligence-could-be-used-to-hack-connected-cars-drones-warn-security-experts/>.

⁶⁵ Elizabeth Fernandez, "Google's CEO Calls For Regulation Of AI," *Forbes*, February 10, 2020, <https://www.forbes.com/sites/fernandezelizabeth/2020/02/10/googles-ceo-calls-for-regulation-of-ai/?sh=48b95bd85371>.

⁶⁶ European Commission, "On Artificial Intelligence – A European Approach to Excellence and Trust EN," *European Union*, 2020.

⁶⁷ Wright, "Jeremy Wright Speaking at the AI Summit."

⁶⁸ Oxford Dictionary, "Ethics | Definition of Ethics by Oxford Dictionary," *Lexico*, 2021, <https://www.lexico.com/definition/ethics>.

⁶⁹ Oxford Learning College, "Ethics versus Morals - Critical Thinking," *Oxford Learning College*, February 17, 2014, <https://www.oxfordcollege.ac/news/ethics-versus-morals/>.

⁷⁰ David Leslie, "Understanding Artificial Intelligence Ethics and Safety A Guide for the Responsible Design and Implementation of AI Systems in the Public Sector," *The Alan Turing Institute*, 2019, <https://doi.org/10.5281/zenodo.3240529>.

⁷¹ Defence Science and Technology, "Defence Technology Framework," September 2019.

⁷² Jennifer Riggins, "Are Programmers Ethically (and Legally) Responsible for Their Code?," *The New Stack*, August 16, 2018, <https://thenewstack.io/are-programmers-ethically-and-legally-responsible-for-their-code/>.

⁷³ John Kingston, "When an AI Finally Kills Someone, Who Will Be Responsible?," *MIT Technology Review*, 2018, <https://www.technologyreview.com/2018/03/12/144746/when-an-ai-finally-kills-someone-who-will-be-responsible/>.

⁷⁴ Daniel Castro, "Who Is Winning the AI Race: China, the EU or the United States?," *Center for Data Innovation*, August 19, 2019, <https://datainnovation.org/2019/08/who-is-winning-the-ai-race-china-the-eu-or-the-united-states/>.

⁷⁵ Neil Savage, "The Race to the Top among the World's Leaders in Artificial Intelligence," *Nature* (NLM (Medline)), December 1, 2020, <https://doi.org/10.1038/d41586-020-03409-8>.

⁷⁶ Gavin Pearson, Phil Jolley, and Geraint Evans, "A Systems Approach to Achieving the Benefits of Artificial Intelligence in UK Defence," *ArXiv*, 2018.

⁷⁷ Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Houghton Mifflin Harcourt, 2019), <https://www.amazon.co.uk/AI-Superpowers-China-Silicon-Valley/dp/132854639X>.

⁷⁸ International Transport Forum, "Managing the Transition to Driverless Road Freight Transport," 2017, www.itf-oecd.org.

⁷⁹ Aimee Murphy, "Reported Road Casualties in Great Britain: 2019 Annual Report," September 30, 2020, <https://www.gov.uk/government/>.

⁸⁰ World Health Organisation, "Road Traffic Injuries," World Health Organisation, February 7, 2020, <https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries>.

⁸¹ International Transport Forum, "Managing the Transition to Driverless Road Freight Transport."

⁸² Edmond Awad et al., "The Moral Machine Experiment," *Nature* 563, No. 7729 (2018): 59–64, <https://doi.org/10.1038/s41586-018-0637-6>.

⁸³ Judith Jarvis Thomson, "Killing, Letting Die, and The Trolley Problem," *The Monist*, April 1976, <https://www.jstor.org/stable/27902416?seq=1>.

⁸⁴ Eleanor Bird et al., "The Ethics of Artificial Intelligence: Issues and Initiatives," *European Parliamentary Research Service*, 2020, <https://doi.org/10.2861/6644>.

⁸⁵ Government, "AI Council," Government, 2021, <https://www.gov.uk/government/groups/ai-council>.

⁸⁶ UK AI COUNCIL, "AI Roadmap," 2021, <https://www.gov.uk/government/publications/ai-roadmap>.

⁸⁷ Office for Artificial Intelligence, "Office for Artificial Intelligence," Government, 2021, <https://www.gov.uk/government/organisations/office-for-artificial-intelligence>.

⁸⁸ Government, "Online Targeting: Final Report and Recommendations," Centre for Data Ethics and Innovation, February 4, 2020, <https://www.gov.uk/government/publications/cdei-review-of-online-targeting/online-targeting-final-report-and-recommendations>.

⁸⁹ Dame Shirley Pearce, "AI in the UK: The Story So Far - Committee on Standards in Public Life," Government, March 19, 2019, <https://cspl.blog.gov.uk/2020/03/19/ai-in-the-uk-the-story-so-far/>.

⁹⁰ Pearce.

⁹¹ Leslie, "Understanding Artificial Intelligence Ethics and Safety A Guide for the Responsible Design and Implementation of AI Systems in the Public Sector."

⁹² Wendy Hall and Jérôme Pesenti, "Growing the Artificial Intelligence Industry in the UK," 2017.

⁹³ Cognilytica, "Worldwide Country AI Strategies and Competitiveness 2020," *Cognilytica*, February 6, 2020, <https://www.cognilytica.com/2020/02/07/worldwide-country-ai-strategies-and-competitiveness-2020/>.

⁹⁴ Castro, "Who Is Winning the AI Race: China, the EU or the United States?" <https://datainnovation.org/2019/08/who-is-winning-the-ai-race-china-the-eu-or-the-united-states/>.

⁹⁵ Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics" (Cambridge, MA, November 6, 2017), <https://doi.org/10.3386/w24001>.

⁹⁶ Raymond Kurzweil, *The Singularity Is Near: When Humans Transcend Biology* (Duckworth, 2006), <https://www.amazon.co.uk/Singularity-Near-Raymond-Kurzweil/dp/0715635611>.

⁹⁷ Ray Kurzweil, "The Coming Singularity," *New World : Artificial Intelligence*, November 23, 2019, <https://www.newworldai.com/the-coming-singularity-ray-kurzweil/>.

⁹⁸ Stanford, "The Naïveté of 'Exponential' Growth," Stanford University, 2011, <https://cs.stanford.edu/people/eroberts/courses/cs181/projects/2010-11/TechnologicalSingularity/pageview2efb.html?file=againstfeasibility.html>.

⁹⁹ Toby Walsh, "The Singularity May Never Be Near," *AI Magazine*, 2017.

¹⁰⁰ Ben Dickson, "Artificial Neural Networks Are More Similar to the Brain than They Get Credit For," *Tech Talks*, June 22, 2020, <https://bdtechtalks.com/2020/06/22/direct-fit-artificial-neural-networks/>.

¹⁰¹ Joe Gale, Amri Wandel, and Hugh Hill, "Will Recent Advances in AI Result in a Paradigm Shift in Astrobiology and SETI?," *International Journal of Astrobiology*, 2019, <https://doi.org/10.1017/S1473550419000260>.

¹⁰² Caroline Williams, "Gaslighting Warps Our View of Reality. How to Spot It – and Fight Back," *New Scientist*, January 27, 2021, <https://www.newscientist.com/article/mg24933190-700-gaslighting-warps-our-view-of-reality-how-to-spot-it-and-fight-back/>.

¹⁰³ Anthony M. Zador, "A Critique of Pure Learning and What Artificial Neural Networks Can Learn from Animal Brains," *Nature Communications* (Nature Publishing Group, December 1, 2019), <https://doi.org/10.1038/s41467-019-11786-6>.

¹⁰⁴ David Rushkoff, "The Figure or Ground," in *What to Think About Machines That Think*, 2015, <https://www.amazon.co.uk/dp/006242565X?linkCode=gs2&tag=braipick-21>.

¹⁰⁵ Peter Diamandis, "Ray Kurzweil's Mind-Boggling Predictions for the Next 25 Years," *Singularity Hub*, January 26, 2015, <https://singularityhub.com/2015/01/26/ray-kurzweils-mind-boggling-predictions-for-the-next-25-years/#sm.005cp0fa11fqf5n11ow2o56tqfxwu>.

¹⁰⁶ John Brockman, *What to Think About Machines That Think: Today's Leading Thinkers on the Age of Machine Intelligence* (HarpPeren, 2015), <https://www.amazon.co.uk/What-Think-About-Machines-That/dp/006242565X>.

¹⁰⁷ Rob Aitken, "Moore's Law Ending? No Problem," *EE Times*, March 27, 2019,

<https://www.eetimes.com/moores-law-ending-no-problem/#>.

¹⁰⁸ UK Essays, "Beyond Moore's Law: Is Quantum Computing a Commercially Viable Replacement for Classical Computing?," 2018, <https://www.ukessays.com/essays/computer-science/beyond-moores-law-is-quantum-computing-a-commercially-viable-replacement-for-classical-computing.php>.

¹⁰⁹ Tia Ghose, "Forget Moore's Law — Quantum Computers Are Improving According to a Spooky 'Doubly Exponential Rate,'" *Live Science*, June 21, 2019, <https://www.livescience.com/65651-quantum-computers-get-scary-fast.html>.

¹¹⁰ Hartmut Nevan and Kevin Hartnett, "A New 'Law' Suggests Quantum Supremacy Could Happen This Year – Scientific American," *Scientific American*, June 21, 2019, <https://www.scientificamerican.com/article/a-new-law-suggests-quantum-supremacy-could-happen-this-year/>.

¹¹¹ Kevin Hartnett, "Does Neven's Law Describe Quantum Computing's Rise?," *Quanta Magazine*, June 18, 2019, <https://www.quantamagazine.org/does-nevens-law-describe-quantum-computings-rise-20190618/>.

¹¹² Richard Feynman, "Simulating Physics with Computers" (Nature Publishing Group, April 2, 1981), <https://doi.org/10.1038/nphys2258>.

¹¹³ Cem Dilmegani, "995 Experts Opinion: AGI / Singularity by 2060," *AI Multiple*, December 31, 2020, <https://research.aimultiple.com/artificial-general-intelligence-singularity-timing/>.

¹¹⁴ Zachary Lipton, "The AI Misinformation Epidemic," *Approximately Correct*, March 28, 2017, <http://approximatelycorrect.com/2017/03/28/the-ai-misinformation-epidemic/>.

¹¹⁵ Press, "A Very Short History Of Artificial Intelligence (AI)."

¹¹⁶ Campbell-Kelly, "Professor Joseph Weizenbaum: Creator of the 'Eliza' Program."

This article has been republished online with Open Access.

Ministry of Defence © Crown Copyright 2023. The full printed text of this article is licensed under the Open Government Licence v3.0. To view this licence, visit <https://www.nationalarchives.gov.uk/doc/open-government-licence/>. Where we have identified any third-party copyright information or otherwise reserved rights, you will need to obtain permission from the copyright holders concerned. For all other imagery and graphics in this article, or for any other enquires regarding this publication, please contact: Director of Defence Studies (RAF), Cormorant Building (Room 119), Shrivenham, Swindon, Wiltshire SN6 8LA.

 **ROYAL
AIR FORCE**
**Centre for Air and
Space Power Studies**

OGL