


Taking control

Artificial intelligence and insurance



Lloyd's of London disclaimer

This report has been co-produced by Lloyd's and University of Surrey staff for general information purposes only. While care has been taken in gathering the data and preparing the report Lloyd's and University of Surrey staff does not make any representations or warranties as to its accuracy or completeness and expressly excludes to the maximum extent permitted by law all those that might otherwise be implied.

Lloyd's and University of Surrey staff accepts no responsibility or liability for any loss or damage of any nature occasioned to any person as a result of acting or refraining from acting as a result of, or in reliance on, any statement, fact, figure or expression of opinion or belief contained in this report. This report does not constitute advice of any kind.

© Lloyd's 2019
All rights reserved

About Lloyd's

Lloyd's is the world's specialist insurance and reinsurance market. Under our globally trusted name, we act as the market's custodian. Backed by diverse global capital and excellent financial ratings, Lloyd's works with a global network to grow the insured world –building resilience of local communities and strengthening global economic growth.

With expertise earned over centuries, Lloyd's is the foundation of the insurance industry and the future of it. Led by expert underwriters and brokers who cover more than 200 territories, the Lloyd's market develops the essential, complex and critical insurance needed to underwrite human progress.

About University of Surrey

The University of Surrey was established on 9 September 1966 with the grant of its Royal Charter, but its roots go back to a late 19th-century concern to provide greater access to further and higher education for the poorer inhabitants of London.

The University of Surrey is a global community of ideas and people, dedicated to life-changing education and research. With a beautiful and vibrant campus, we provide exceptional teaching and practical learning to inspire and empower our students for personal and professional success. Through our world-class research and innovation, we deliver transformational impact on society and shape future digital economy through agile collaboration and partnership with businesses, governments and communities.

Key contacts

Trevor Maynard
Head of Innovation
trevor.maynard@lloyds.com

For general enquiries about this report and Lloyd's work on innovation, please contact
innovation@lloyds.com

About the authors

At the time of writing this report Roger Maull was a Professor of Management Systems in Surrey Business School. He was a Director of Surrey's Centre for the Digital Economy (CoDE). He has held over £15m of UK research income and published over 130 publications in leading journals and conferences. He re-joined the University of Exeter on October 2018 as Academic Director of the Initiative for the Digital Economy at Exeter (INDEX) based in London.

John Collomosse is Professor of Computer Vision at Surrey's Centre for Vision, Speech and Signal Processing (CVSSP), one of the UK's largest research groups for AI. He is a Visiting Professor at Adobe Research Creative Intelligence Lab.

Steve Brewer is the founder and director of Infoculture, a start-up offering support and facilitation for organisations, teams and leaders in digital transformation. Steve has been involved in communication, community engagement and project management for several years in UK and EU research-related projects.

Anna Bordon, BA, MSc, is an Associate in the Innovation team at Lloyd's where she supports the development and implementation of the innovation and research plans in line with Lloyd's strategy. Her focus is on innovation and emerging risks with the aim to inform the Lloyd's market of business and product development opportunities.

Kristian Jones MSci, is a Lloyd's Graduate on placement in the Innovation team. He has had previous placements in underwriting and claims at two Lloyd's syndicates. Prior to joining Lloyd's, he studied Chemistry at the University of Bristol. He is interested in increasing market participation and stakeholder involvement in Lloyd's initiatives.

James Breeze, Digital AI Theme Lead for AXA XL, has broad experience in business intelligence design, cognitive analytics and machine learning. He has informed the 'Insight: Commercial property risk engineering' box on page 43 which he contributed to.

Acknowledgements

The following people were consulted or commented on earlier drafts of the report, attended workshops and 1:1 meetings, and provided their time and support to the project; we would like to thank them all for their invaluable contributions:

Lloyd's project team

- Dr Trevor Maynard, Innovation
- Dr Keith Smith, Innovation
- Craig Civil, Data Lab
- Anna Bordon, Innovation
- Lucy Stanbrough, Innovation
- Kristian Jones, Innovation
- James Burchill, Innovation
- Linda Miller, Marketing and Communications
- Kieran Quigley, Marketing and Communications
- Flemmich Webb, Speech and Studies

University of Surrey project team and area of expertise

- Attila Emecz, Director of Strategic Partnerships
- Ryan Abbot, Professor of Law and Health Sciences
- Alan Woodward, Visiting Professor, Surrey Centre of Cyber Security
- Steve Schneider, Professor and Director of Surrey Centre for Cyber Security
- Alan Brown, Professor of Entrepreneurship and Innovation, University of Surrey (now Executive Director of INDEX at the University of Exeter)
- Yin Lim, Editor and Writer, Yin F Lim Editorial Services

Insurance industry interviews and consultation

- Nick Gibbs, Apollo
- Hayley Rigby, Apollo
- William Skertic, Beazley
- Alan Godfrey, Asta Managing Agency
- Nicky Singh, Expert System
- Richard Hughes, Independent Consultant
- Martin Twells, W/R/B
- James Breeze, AXA XL

Lloyd's Market Association

- David Powell, Head of Non-Marine Underwriting
- Tony Elwood, Senior Executive, Underwriting
- Gary Budinger, Senior Executive, Finance and Risk

Lloyd's Lab Cohort 1

- Raj Jagannathan, ZASTI
- Davor Runje, ZASTI
- Dr Ryan Lloyd, Geollect
- Phillip Naples, Layr

Further thanks go to the following for their expertise, feedback and assistance with the study:

- Ganna Pogrebna, Professor of Behavioural Science at University of Birmingham, Fellow at Alan Turing Institute
- Mike Chantler, Professor of Computer Science, Heriot-Watt University and Director of Strategic Futures Lab
- Dave Snelling, Program Director Artificial Intelligence at Fujitsu

Contents

Executive summary.....	5
1. What is artificial intelligence?.....	11
Brief history of the development of AI and major trends	14
Summary of recent developments and breakthroughs	16
Examples of current AI applications	16
2. Societal and international security impacts.....	22
3. Risks	27
Transparency and trust	27
Ethics.....	27
Liability.....	28
Security.....	28
4. Regulatory and government landscape	31
Challenges.....	31
Government landscape	31
5. AI and insurance	36
Impact of AI on insurance lines of business	36
Business development opportunities	42
Operations	42
InsurTech.....	44
6. Conclusions	47
Appendix A – History of AI	48
Appendix B – AI in academia and insurTech.....	51
Glossary	55
References.....	56

Executive summary

Artificial intelligence (AI) – the capability of a machine to imitate intelligent human behaviour – is increasingly being used in society in all sorts of ways – from diagnosing medical conditions and scanning legal documents to agriculture precision spraying to prevent herbicide resistance and self-driving cars.

AI is not new. It has been a significant presence, on and off, in the academic landscape for over 60 years. Nevertheless, the recent, rapid escalation in real-world disruptive implementations of AI has awakened an awareness of its complex and profound ethical, legal and societal challenges that go beyond the technological complexities of its underlying methods and processes.

The combination of AI with other technologies such as the Internet of Things (IoT; low-cost and highly available, widely installed sensors and actuators), 5G and fog/edge computing will enable the collection and creation of large collections of data, ready to feed AI systems.

As its influence and technical capability grows so the risks and opportunities of using it evolve. This report aims to help insurers navigate this fast-developing area. It sets out technology risks, as well as potential AI application and business opportunities in insurance.

Technology risks

- **Uncertainty.** There are risks in AI adoption as the technology is still in uncharted waters and there is concern among stakeholders, including regulators.
- **Ethics.** Machine learning depends on large collections of data from which to extract knowledge. These datasets effectively become the codification of history.

However, within these histories there are prejudices, biases and societal injustices that have existed for years. The consequence is that the AI depends on the information on which it is trained, and it might act against human interests.

- **Trust and transparency.** The conjunction of bias and the black box of Deep Neural Networks (where we cannot explain how a conclusion was reached) has enormous implications for society and insurance. Transparency of the decision-making process and the ability to understand how an AI got to a specific conclusion will be key to developing trust in the technology.

To ensure trust and transparency to the public along with a broader alignment to the full diversity of the real world, clearer accountability will be required for decision-making processes without sacrificing AI performance.

- **Liability.** Another significant challenge is the legal status of AI systems and services. Whilst much can be done to encourage and motivate the engineers behind the systems to consider the ramifications of their decisions in designing neural networks and selecting data sets for learning, there needs to be robust legal frameworks to define liability.

Generally speaking, the manufacturer of a product is liable for defects that cause damages to users. However, in the case of AI (especially strong AI) decisions are not a consequence of the design, but of the interpretation of reality by a machine.

- **Security.** It is worth noting that AI is also forming a new dimension in cybersecurity as it can learn and react to threats much faster than the traditional methods used by security products of old. However, not surprisingly, criminals are also beginning to use AI to learn how to conduct all stages of a typical attack: from reconnaissance to crafting a specific attack.

As these systems become more complex and internally connected, breaches of (cyber) security for interconnected AI will become more important and are likely to have deep systemic economic impact.

AI impact on insurance lines of business

AI could impact several lines of business. The report finds implications for product recall and liability, third-party motor liability, professional indemnity, medical malpractice, cyber, fidelity and political risks.

Product liability and product recall

- There are many ways in which a product reliant on AI could affect product liability and product recall. Examples include:
 - A new situation that AI is not programmed to deal with that causes breakdown or malfunction.
 - Product defects could result from communication errors between two machines or between machine and infrastructure.

In general, recalls could become larger and more complex, particularly if the affected sector uses AI extensively.

- In terms of liability, AI machines cannot themselves be liable for negligent acts or omissions that cause damage to third parties, but as they become better at learning and deciding, the question “Who is liable when a machine commits a civil wrong?” is raised.
- Product manufacturers/sellers, AI designers/suppliers and AI purchasers/users could be allocated fault by courts in liability cases. This could mean that in the future companies might buy more contractual warranties, indemnities and limitations to control AI liability risk.

Third-party motor liability

- Assignment and coverage of liability will become more challenging in the future due to the possible shift of responsibility from human drivers to automated vehicles (AVs), and therefore to the manufacturers.
- There is also significant potential for different legal and regulatory responses in different countries creating a complex international risk landscape.

Medical malpractice

- There are various ways in which negligence could arise. AI is being used to help diagnose conditions and is being applied in areas such as radiography (Lee et al., 2017). If an error leads to misdiagnosis or false positives, this could amount to negligence.

- Even if AI was used as an aid for referrals, if these referrals prompted investigations or procedures that were unnecessary, invasive or led to poorer patient outcomes then liability could arise.

Cyber

- As chatbot technology develops, it is becoming increasingly difficult to tell humans and AI apart. This could make it easier to carry out phishing scams, and on a wider scale, as well as harder to detect them.
- This raises questions about what types of insurance cover would be available to protect against losses caused by these sorts of cyber-attacks. What constitutes an insurable incident will have to be carefully defined.
- Third party liability could also become more complicated. If a chatbot is taken over to launch an attack, its owner could be liable if they could not show they had taken reasonable security measures to prevent such an attack happening.

Fidelity

- Fraudulent activity from employees could also be exacerbated by any of the methods mentioned in the Cyber section above. Fidelity insurers should consider that fraud may increasingly come from employees with access to IT systems rather than employees with financial authority.
- The emergence of “deep fakes” is also concerning when it comes to fraud. Deep fakes are AI systems capable of generating realistic audio and video. This is concerning for cases of identity fraud and could also be used in conjunction with chatbots to increase authenticity and build trust.

Political risks

- The weaponisation of AI could take many forms including corruption of data, biased data selection, and the illegal use of data leading to various outcomes including propaganda, behavioural change and deception.
- From a political risk coverage perspective, AI might contribute to creating new or exacerbating existing political events such as expropriation, wars, acts of terrorism, civil disturbances and other forms of political violence in both developing and developed markets.
- Instability generated by automation and economic disruption could also be a potential driving force for protest and agitation, resulting in government interference, selective discrimination and business interruption.

Business opportunities for insurers

- In general, any company offering algorithm-based systems to data-rich companies (e.g. fraud detection in online sales) might seek to insure against the risk of the algorithms returning incorrect decisions and its impact on the AI companies' clients.
- New companies are emerging in the disinformation defence area to provide technology to filter out fake news; detect and eliminate troll-bots; and certify information and authenticity of images and videos. Insurers could explore what type of products (e.g. professional indemnity and cyber products) could be useful to these new businesses and in what form.
- As the field develops and applications increase, opportunities arise for providing risk management services. Even though AI development is at a relatively early stage, AI knowledge experts are already in high demand from firms attempting to manage risks. Specialist service firms are emerging that are aimed at loss prevention. As complexity increases, the demand for these services will rise.
- By taking leadership role in this space, insurers will acquire the knowledge needed to provide insureds with guidance on AI best practice, thereby shaping the development of the AI ecosystem in which they operate.

Using AI to improve insurance processes

From an operational perspective, insurance companies are already exploiting the potential of AI to deliver value in a more efficient way. Examples include:

- **Customer service:** Chatbots, which can recommend and personalise products, handle complaints, improve communications with customers and process simple transactions.
- **Underwriting:** This could be enhanced and sped up by using AI. AI-based models could generate premium rates based on historical risk assessments and could generate bespoke quotes based on a customer's risk profile.
- **Fraud detection:** Today this is a largely human-led operation, but it can be automated by using AI. Fraud detection, could be further automated with the augmentation of external data from historical records, sensors and images to better estimate repair and write-off costs.

- **Claims:** AI could help reduce the number of claims that require human analysis by automating image recognition, searching large databases for identical claims to prevent fraud, validating claims and informing customers of their claim's progress.
- **Modelling, exposure management and pricing:** With AI and its combination with other technologies such as the internet of things, insurers and model providers will have access to data they can feed into models that learn and adapt at a much faster pace than manual equivalents. This could, for instance, forecast interest rates based on central banks' communications or help predict missing fields within a data set.
- **Business development.** AI systems can improve cross-selling and conversion rates of products, can propose tailored products and can help identify new clients.

AI and Lloyd's

Lloyd's has been exploring AI and how it can be used to create opportunities for customers and the market.

Examples include:

- Lloyd's itself has already automated its International Trading Advice (LITA) help desk. By using AI, the search time for queries that would usually take several minutes was reduced to seconds whilst maintaining accuracy of reply.
- This has led to another instance of AI being deployed in the Lloyd's Canada team to ingest and analyse contract documents to ensure local regulatory compliance.
- Lloyd's has also created the Corporation AI Strategy covering Principles, Ethics and Opportunities that will guide the future use of AI.
- Lloyd's is helping develop insurTech AI solutions through the Lloyd's Lab where several companies (e.g. ZASTI, Layr and Geollect) have used data and AI to enable risk prediction, improve risk management and speed up claims.
- Between January and April 2019 the Lloyd's Lab and the Committee of Actuaries in the Lloyd's Market have run two Proof of Concepts in collaboration with ZASTI and Expert System exploring the use of AI for better claims frequency estimation and data entry quality control.

Conclusions

As businesses increasingly incorporate AI into their systems and processes they will need insurance to protect them from a range of potential risks.

- Insurers should respond to this technology by developing appropriate models and products.
- Insurers can improve their current ways of working by using AI themselves in all part of their value chain, from the first enquiries, to the settlement of claims through to risk prevention.
- While AI offers potential for business development and operational efficiencies, insurers must be aware of the risks associated with using this new and fast developing technology.

Time horizons

- In the short term we can expect to see increasing adoption of AI for relatively constrained speech and vision problems, for recommendations and behaviour modelling and optimisation.
- For the medium term, more sophisticated problem sets include autonomous vehicles and front office customer service.
- In the longer term we will see a shift to strong AI, with systems' intellectual capability being indistinguishable from human intelligence.

Actionable insights

Based on our research, we distil this report in 10 actionable insights for insurers:

1. **Address the bias.** AI systems, particularly machine learning solutions, require comprehensive and valid data sets on which to learn. More critical insight is being applied to the sources of data used to develop AI systems in terms of the social and political impact triggered by the selection processes.
Questions to ask include “Have societal biases been imbedded in the data by choice, or omission, or accident?” If there are biases in the data, then there will be biases in the algorithms and their choices.
2. **Think about data brokering.** Different regions and cultures across the world have different approaches to the uses to which data can be applied. But whatever approach is taken to the capture of the

data and its sharing, trading and reuse, the data will expand the horizons of AI.

In the UK for example, there are plans to develop a framework for storing and sharing access to data amongst organisations and domains. The bigger the data set the better insurers will be able to understand and price risks.

3. **Deliver technical transparency and explainability.** There is a growing awareness that AI systems should be technically transparent and explainable, especially in critical situations. Citizens and consumers will demand this if AI is to be trusted and believed.
4. **Design through collaboration.** The origins of AI as a sub-discipline of Computer Science has resulted in a technology that has been created from within a small community.

The future of practical AI systems would benefit from the participation of a wider and more diverse community with broader academic perspectives and social horizons.

5. **Analyse and monitor.** Whether we are looking at legacy systems or new designs, AI solutions should be monitored, analysed and reviewed over time.

Questions to ask include “Has the world changed sufficiently to amplify or compensate for the biases within? How are different communities and cultures experiencing the benefits of the AI system? Is the test set error^a increasing?”

6. **Support regulations and ethical frameworks.** Fairer and more representative AI solutions will be provided as a result of rigorous and enforceable regulations and ethics frameworks, developed by a broader community aligned to the full diversity of the real world. Insurers should support policymakers in the search for societal solutions that protect exposed workers.
7. **Raise awareness.** AI needs to be opened up. Not every individual can be involved in every design, but if AI is going to change the world for the better, then this can only be achieved if more voices are included in the debate.
8. **Invest in skills and competences.** Specific, specialist knowledge, coupled with collaboration, leadership and communication are all needed by the practitioners and leaders to address the other issues and opportunities. Invest in recruitment and re-skilling existing workforce.

^a Test set error is the error you get when you run the trained model on a set of data that it has previously never been exposed to. Increase test set errors shows the model and reality are changing.

A large component of whether an AI project succeeds or fails is how well (or not) the change management process in an organisation / department / team is managed. The AI technology might be great but if the staff using it are not brought into the concept from the start and provide support the AI project will most likely fail.


9. **Engage with insurTech.** Many of the technological advances in AI are being led by small start-ups, often with strong links to universities. These firms are moving out of the lab to operationalise advances in AI and they provide excellent vehicles to nurture AI expertise and connections in larger established firms.

10. **Explore AI's impact on products and processes.**

Insurers should explore how their clients will be adopting AI to provide new solutions and how risks created by AI will impact existing lines.

From an operational perspective, insurance companies should test the potential of AI to deliver value in a more efficient way at all stages of the value chain.

What is artificial intelligence?



1. What is artificial intelligence?

Artificial intelligence (AI) today is ubiquitous. From art fraud to healthcare, from transport to agriculture, there is no sector that is not seeking to reap the benefits of the methods and approaches that have emerged since the ground-breaking Dartmouth Conference of 1956 that marked the beginning of the research field.

Over the intervening time span, AI has experienced a series of peaks and troughs as research funding has flowed in when the approach has been perceived as offering exciting solutions for pressing challenges and dried up when application solutions have seemed elusive.

This Lloyd's report is the starting point for a better understanding of the key aspects of AI as well as its risks and opportunities. The overarching questions that resonate through this report are: what are the risks and opportunities that AI will generate, and how do they relate to insurance?

This report was developed through a structured research process across two stages: literature review and expert elicitations with academics, technologists, economists and insurers.

How will the world look in the future with the influence of AI? AI is viewed with both excitement and concern at the same time as it already encroaches on many aspects of our lives: mobile phone assistants, self-driving vehicles, personalised web page search results and online shopping recommendations, not to mention healthcare, financial trading and translation services.

Across the financial sector, including insurance, institutions are exploring new approaches that draw upon the power of AI. The insurance sector has and continues to see innovations in different areas such as in pricing, claims handling and fraud detection.

Overall, the world of insurance is facing a range of emerging AI-related risks coupled with a wealth of

innovation opportunities to increase productivity. These range from off-the-shelf products to proofs of concept through innovation labs and partnerships with start-ups.

Lloyd's is taking on these challenges through several innovative actions. They include the use of AI to build a system that is now providing fast and accurate knowledge to the Lloyd's International Trading Advice (LITA) team, and the Lloyd's Lab, a fast-track, fast-fail environment where new concepts, ideas and products can be tested with the support and active involvement of the world's specialist insurance marketplace.

For example, several companies that have come through the Lloyd's Lab have used data and AI to enable risk prediction, improve risk management and speed up claims. Section 5 includes case studies for Zasti, Geollect and Layr. A longer list of insurance start-ups working with AI is available in Appendix A.

However, AI is not new. AI has been a significant presence, on and off, in the academic landscape for over 60 years. Nevertheless, the recent, rapid escalation in real-world disruptive implementations of AI has awakened an awareness of its complex and profound ethical, legal and societal challenges that go beyond the technological complexities of its underlying methods and processes.

The combination of AI with other technologies such as the Internet of Things (IoT; low-cost and highly available, widely installed sensors and actuators), 5G and fog/edge^b computing will enable the collection and creation of large collections of data, ready to feed AI systems.

Whilst all of these offer smarter and faster working, there is concern among stakeholders, including regulators, due to the 'black-box' nature of solutions (Brown et al., 2017). Understanding, and explaining, how these 'black boxes' operate is one key challenge for the future of AI.

^b Is a technology that aims to bridge the gap between remote data centres and IoT devices by analysing time-sensitive data at the network edge, close to where it is generated instead of sending vast

amounts of IoT data to the cloud (Alrawais, Alhothaily, Hu, & Cheng, 2017; CISCO, 2015).

Common terminology

To help build understanding the following key terms may be of use to build familiarity with the topic, as terms are interchangeable and the topic crosses fields. A list of acronyms can be found in the glossary (p57).

The defining concept that helps explain the shifts that AI will deliver is captured in the transition from automated towards 'autonomous' systems. The former describes systems that follow pre-programmed instructions, and the latter describes systems that act 'independently' towards a programmed goal. This distinction lies at the heart of societal challenges and emerging risks.

Box 1: Key terms

Artificial intelligence

AI is an umbrella term for an evolving series of technologies that have ebbed and flowed over the last few decades. For this reason, there is no simple definition of AI that works in all contexts and for all users.

The Merriam-Webster Dictionary defines artificial intelligence as:

1. A branch of computer science dealing with the simulation of intelligent behaviour in computers
2. The capability of a machine to imitate intelligent human behaviour

Variations of AI definitions are based on what AI is being used to achieve. For example, tech companies and start-ups describe AI as the science of creating intelligent machines capable of performing real-time tasks at the level of a human expert.

Some of the tasks that properly-constructed AI can perform well include:

- Processing of large amounts of data, with higher speed and therefore with greater capacity than humans
- Learning from example data sets and adapting to solve new problems based on what has been extracted from the data sets
- Recognising objects and their association to a situation
- Inferring the future state of an object or situation
- Identifying an optimal decision based on past, present and inferred future states

Machine learning

Machine learning is a subset of AI that often uses statistical techniques to give computers the ability to 'learn' (i.e. to progressively improve performance on a specific task) with data, without being explicitly programmed to do so.

Deep learning

Deep learning is a branch of machine learning that utilises complex neural networks to make decisions, often comprising many millions of neurons arranged in tens or hundreds of layered structures. A common form of deep neural network is the convolutional neural network (CNN) that can process 1D signals (like time series data or audio) or 2D signals (like images or video) in order to detect the presence of objects or features of interest in the signal. These artificial neural networks are modelled like the human brain, with neuron nodes connected like a web. See Box 2 for more information.

In this report, we use AI as an umbrella term that covers all the above (unless indicated differently).

Supervised versus unsupervised learning

There are two main approaches used in machine learning: supervised and unsupervised.

- Supervised defines a context where prior goals exist regarding the values to be applied to the data in question. Therefore, the aim of supervised learning is to create a system that can reliably detect the predetermined patterns and classifications. In the context of insurance, this could mean that premium rates could be determined based on the precise characteristics of the entity seeking insurance.
- Unsupervised learning on the other hand refers to searches for clustering, where the data is searched for patterns and structural features without predefined labels or classifiers. In the insurance context, this could be searches for patterns amongst customers that might lead to further product referrals, along the lines of 'customers of this type who bought these policies might also be interested in these other products'.

Figure 1: From supervised to unsupervised AI

SUPERVISED	SEMI-SUPERVISED	UNSUPERVISED
<ul style="list-style-type: none"> - You have input variables (x) and an output variable (Y) and you use an algorithm to learn the mapping function from the input to the output $Y = f(x)$ - Goal is to approximate the mapping function so well that when you have new input data (x) you can predict the output variables (Y) for that data 	<ul style="list-style-type: none"> - It typically uses a small amount of labelled data with a large amount of unlabelled data 	<ul style="list-style-type: none"> - You have input data (x) and no corresponding output variables - Goal is to model the underlying structure to learn more about the data - Algorithms are left to their own to discover and present the interesting structure in the data


Source: Lloyd's, 2018

Weak (narrow) and strong (broad) AI

AI is often divided into weak/narrow and strong AI. The difference is that weak/narrow AI is focused on solving a specific problem in delimited areas (e.g. customer service chatbots). As AI becomes smarter thanks to the amount of learned data and can solve problems never seen before, it can be called strong.

Figure 2: From weak (narrow) to strong (broad) AI

Weak/Narrow AI	Artificial general intelligence (AGI) Strong/Broad AI
<ul style="list-style-type: none"> - Systems that can behave like humans and do the right thing - But do not explain how humans think - Focus on a narrow task - 1997 IBM's Deep Blue 	<ul style="list-style-type: none"> - Genuinely simulating human reasoning - Build systems that think but also explain how humans think - Untrained and unsupervised



Source: Lloyd's, 2018

Strong AI capable of reasoning beyond human capacity is currently lodged firmly within the realms of science fiction. Opinion is divided as to whether development of such a system is possible or even a desirable goal, although for many it represents a powerful motivator for ongoing research.

Today's intelligent systems are capable of astonishing feats considered as grand unsolved challenges only a few years ago; for example, AIs that can describe visual scenes or answer questions about images in natural language (visual question answering).

These tightly-bounded abilities should not be mistaken for self-awareness or the ability to reason; rather, they reflect the ability of a neural network to learn and interpolate from thousands of appropriate trigger words and phrases to provide a convenient human computer interface.

It is tempting, in the face of such achievements and shared terminology, to draw too deep a parallel between today's AI – driven by deep neural networks (DNN) – and the neural networks of the biological brain. Whilst superficial similarities exist (both are comprised of interconnected neurons), the operation of the neurons and the architectures differ significantly.

Biological neurons are not connected in sequential layered structures as with DNNs designed by humans. Rather, biological neurons are connected in irregular patterns that change over time with neurons firing simultaneously in an asynchronous fashion, the timings of which give rise to emergent properties and behaviours.

Biological neurons can 'forget' over time and occur in large quantity. Whilst we can seek to emulate some of these properties within a DNN, there is no reason to conclude that DNN technology in its current form could give rise to an artificial consciousness. Nevertheless – questions of spirituality aside – there is also no reason to doubt that the complexity of the biological neural system could not one day be modelled within a computer simulation.

To better understand the impact that AI is having and will continue to have on the world at large – including the world of insurance – we need to look back at its 60-year history. Despite the lengthy gestation period, the application of AI has only recently started to have a broad and pervasive effect in the wider world.

Previously, AI was seen as mostly a science with strong theoretical and applied branches. Apart from military implementations, applications of AI were proof-of-concept demonstrators or narrowly-focused exemplars such as a chess player, human-assisted vehicles, factory robotics and other expert systems.

Brief history of the development of AI and major trends

AI has been around for over 60 years and only relatively recently it has come to dominate discussion about its potential effect on the future of work and many other aspects of our lives.

The history of AI can be viewed as three waves (so far), each separated by brief, but significant, funding droughts often referred to as AI winters. These phases can also be considered by their dominant technological approaches.

- **First wave:** The first phase of AI development centred around symbolic approaches where problems were represented as a kind of maze to be explored. After the first AI winter ended when funding resumed, a new era began in the 1970s and 1980s with the growth of expert systems, also known as knowledge-based systems.

These comprised collections of high-level domain knowledge compiled from experts. As work started to grow in neural networks, the second AI winter occurred as research funding dried up again.

- **Second wave:** Later in the 1990s, research expanded again with a move away from centralised, hierarchical systems. They were replaced by research into collaborating agents, interoperating independently, sharing messages and learning from their experience.

This research approach was able to flourish thanks to a confluence of technological factors, namely the growth in low-cost data storage and processing power, as well as large data availability and the internet providing connectivity and distributed information.

- **Third wave:** AI then made the transition from symbolic approaches based around structured information to sub-symbolic approaches with unstructured content and decentralised control structures.

These specialist applications can excel in their individual tasks where the specialised data has been used to train relevant algorithms; they are typically referred to as narrow AI, although the greater goal is to deliver general AI or Artificial General Intelligence (also known as full/strong/broad AI, see Figure 2).

Various tests have been proposed to establish this elusive goal, chief among which is still the classic Turing test developed by Alan Turing in 1950 to evaluate whether a machine can imitate human cognitive capacity (Turing, 1950).



Insight: How deep learning works

The rapid commercial adoption of AI in the past few years has been catalysed by technological advances in machine learning – specifically those in ‘deep learning’ – that delivered transformational performance gains, whilst generalising to a broad range of applications. Deep learning refers to the use of complex (or ‘deep’) neural networks to solve machine learning problems such as classification (e.g. image or scene understanding) or regression (e.g. model fitting).

Neurons are the building blocks of neural networks; a digital simulation of a biological nerve cell that aggregates information from, and passes this information to, other neurons in the network. The behaviour of a neural network is determined by the pattern of connections between its neurons (the ‘network architecture’) as well as the way each neuron manipulates the information passing through it. The architecture of neural networks still requires manual, human design. However, once designed, a network’s behaviour is defined by the importance or ‘weight’ that each neuron learns to ascribe to information received from its neighbours. These weights are the parameters learned by the network during its exposure to training data.

Neural networks are usually designed to connect neurons together in a layered structure. Classical neural networks of decades past contained only one or two layers; by contrast, deep neural networks (DNNs) can contain over 100 layers and in the order of tens of millions of weight parameters. The upshot is that these networks can perform more complex tasks but require higher volumes of training data and take more computational effort to train. This is significant, as the number of patterns a neural network can discern between is proportional (though not linearly) to the number of neurons, and so weights, within the network.

Deep learning is a fast-developing field and there remain significant challenges – not least the long training times that prevent real-time learning from data, and the demands for very high volumes of data. These problems do not befall more classical techniques for machine learning such as support vector machines (SVMs) and graphical models that led the field prior the emergence of deep learning; they still have a place in the landscape where annotated training data is sparse, or where data must be learned adaptively via incremental training over time e.g. online data processing.

Perhaps the most significant issue with deep learning is our limited ability to explain or ‘interpret’ a trained network’s behaviour. It is easy to quantify the performance of a deep network, or indeed any machine learning system; one simply holds out some data (referred to as a ‘test’ data partition) from the training data and evaluates the trained network on that held-out data to measure how accurate are the decisions made.

Whilst this is quite satisfactory in a lab setting, the success of machine learning is driving widespread commercial applications and on mass-deployed machine learning-based systems there will be many situations for which it would be impossible to predict a test case (‘the unknown unknowns’). For example, how could we develop a test data set of all possible conditions that a driverless car will encounter? At some stage, as with all software, we must make a judgement call based on test coverage, as to whether the system is ready for release.

Summary of recent developments and breakthroughs

After the earlier boom and bust cycle, AI appears to be enjoying a prolonged period of resilient growth. The mantle of progress has now passed from military and academic research to the private sector.

Take-up of AI is global, with various hotspots across the developed world, but particularly with China's significant and growing presence thanks in part to its access to vast quantities of data capturing the lives of its huge population and committed state funding of US\$150bn for the next few years (Reuters, 2018).

The extensive impact of AI over the last few years has been enabled through four complementary factors, the:

- Increasing growth, and diminishing cost, of computing processing power;
- Availability of plentiful and cheaper data storage;
- Expansive growth of network connectivity across the internet; and
- Exponential growth of data (big data) available across the world.

Examples of current AI applications

After many years of being used for demonstrator projects in university laboratories, AI is now being applied in a range of activities across many sectors (including insurance).

AI also stands to transform many other aspects of the workplace in the sense that AI-enhanced services will track, extend and monitor the work of humans.

The examples below offer an overview of the use of AI in different sectors, but this is not an exhaustive list.

Examples include:

- In medicine, digital scans can be processed and analysed with systems trained to identify tumours and other abnormalities.
- Lawyers can now process huge quantities of legal documents as they look to identify patterns of irregularity.
- Entertainment service providers can recommend to customers personalised material based on only a few selections, by harnessing the preference histories of others.

- The developers of online games can carefully ramp up the complexity of perceived opponents based on the abilities of individual players to avoid game abandonment or overconfidence.
- Cameras and other traditionally complex devices can now automatically calibrate themselves to the needs of less-qualified operators.

Healthcare

Box 2: DeepMind Health and the Royal Free London NHS Foundation

Part of the Alphabet group, DeepMind is a world leader in AI research (DeepMind, 2018). It has applications in healthcare and is working in partnership with Moorfields Eye Hospital.

To date, AI has made good progress in analysing 2D medical images to identify common diseases (Lee, 2017). DeepMind has applied deep learning to identify retinal eye diseases from 3D scans (DeepMind, 2018).

Optical coherence tomography (OCT) scans are 3D images of the back of the eye; DeepMind's deep learning architecture was trained using a set of OCTs. After learning from a set of just 14,884 OCTs the AI can make accurate referral recommendations, in line with expert interpretations for a variety of sight-threatening pathologies.

The framework made no clinically-serious wrong decisions (the number of patients used in non-training data set was 997). The framework's wrong referral rate was 5.5%. This exceeds the accuracy of the eight specialists trialled, with the top two specialists having error rates of 6.7% and 6.8% from OCT only. The framework error rate was also in line with specialists who had other eye (fundus) images and clinical notes (De Fauw, 2018).

But not every solution in healthcare needs AI to start with as there might not be enough data available. For example, Streams is an app developed by DeepMind which can quickly identify patients with acute kidney injury by analysing a range of test results (Royal Free London NHS Foundation Trust, 2017).

Streams sends instant alerts to direct specialists to patients needing further assessment and treatment. Streams then integrates medical data from existing IT systems into one place, so they can be easily accessed via the app (DeepMind, 2018). Once apps such as Streams are widely adopted, developing AI to aid diagnoses will be the next step.

Telecommunications and customer retention

Telecommunication companies have a strong interest in customer retention, and AI presents a powerful method with which to extract value from the rich collection of data they already possess.

This can be achieved by examining data that companies already hold on customers. In the past, churn prediction is based on analysis of customer data. However, this can be significantly enhanced with social network analytics.

A churn prediction score based on available customer data can then be used to identify which customers should be targeted for a retention campaign, in terms of whether the anticipated future profit exceeds the cost of the campaign (Óskarsdóttir et al., 2017).

Box 3: Ocado

E-commerce is already looking to AI to provide innovative solutions (Ocado Technology, 2018). Ocado Technology is using AI to improve how Ocado interacts with customers online. It achieves this by utilising machine learning and natural language processing. Their software analyses customer emails and generates tags that help customer contact centre workers prioritise emails and feedback.

Ocado Technology is developing a system which uses robotics, AI and computer vision to help with the picking and packaging of a vast array of 50,000 different items available on Ocado.com (Ocado Technology, 2017). It is developing a picking robot for use in their Customer Fulfilment Centres, which are highly-automated warehouses used for storage and distribution of groceries.

The robotic picking arm is equipped with an air compressor and suction cup. Instead of modelling every item, they have developed a prototype 3D vision system. The AI algorithm controlling the robot can understand where the crates are located and identify optimal grasping points for each object. This is complex due to the variety of object shapes, sizes and orientation. Once secured, the arm can orientate the object, search free space in the delivery crate and pack the item. Sensors are used to avoid the risk of items being crushed or damaged.

Industry and manufacturing

AI plays a significant role in the transformation of business, industry and the manufacturing sector. In manufacturing, this role is encapsulated with the partnering of real-world physical systems with enhanced digital systems; they mirror the physical systems and enhance them with copious amounts of data from a range of other, related sources in addition to data gathered from the physical systems themselves.

This partnering of systems is sometimes referred to as Cyber-Physical Systems (CPS). Such partnering of physical and data-rich virtual systems is not new, and neither is the inclusion of AI systems.

Earlier systems around the 1980s involved expert systems that were centralised and contained formally-structured knowledge captured from experts. These systems were used to enhance design, scheduling, production, inspection, diagnosis, modelling and control.

Despite the ongoing developments in AI, new paradigms of implementation for manufacturing had to wait until unstructured data, distributed systems and more context-aware AI approaches were used to incorporate data from both within and without the firm into actionable wisdom.

In other words, IoT data from around the manufacturing facility can be combined with customer preferences and requests, as well as broader sentiment from social media and other sources, to create a complex collection of unstructured data that can be used for model learning, prediction and ultimately actionable instructions.

Additive manufacturing, also known as 3D printing, is a core component of the complex mix offered by digitally-supported smart manufacturing. AI has a role to play in many aspects of Industry 4.0. 3D printing is a good example as it is a key component of short runs of high-quality, mass-customised products.

Traditionally, products would go through many iterations of prototypes, trials and other tests. This is not possible in the physical world for highly-personalised bespoke solutions; therefore, AI can be used to imitate such processes in the virtual world, with interaction possible with the end-user or customer.

This process can also be used to validate a distributed supply chain where AI checks can reinforce handovers between the different stages of the manufacturing process. These AI checks can also be used to seek out cyber-attacks and other malicious influences in the system (Yang, Chen, Huang, & Li, 2017). On the other side, an AI system failure could result into design faults leading to large and complex accidents.

An AI approach based on the development of a Bayesian network has also been found to be beneficial in improving safety in the workplace. The network was able to capture the complex interrelationships and dependencies that affect safety, with factors that include safety attitude, safety knowledge and the supporting environment across the workplace.

Recent advances in AI in terms of power and complexity have enabled new approaches for the autonomous control of crewless marine vehicles. Complex models are needed to imitate the ship's dynamic mathematical characteristics including static, kinetic, dynamic parameters to plan manoeuvres.

Against this, information is also received on the environment including multiple types of obstacles in the changing surroundings. Research is introducing new techniques such as deep learning architectures and reinforcement learning algorithms which can deliver concise and extendable solutions applicable to other complex tasks (Cheng & Zhang, 2017).

Bayesian networks are also proving useful in supporting agriculture, a sizable and critical sector in many emerging and developed countries. This branch of AI is applicable because of its ability to deal with decision-making with incomplete information and represent interdependencies between causes and factors.

Machine learning can support farmers in making decisions about seed types, fertilizers, disease identification and treatment (Drury, Valverde-Rebaza, Moura, & de Andrade Lopes, 2017).

Box 4: Oxbotica

Oxbotica is an autonomous vehicle software company whose AI-powered software, Selenium, is being used to develop driverless cars. The use of AI helps the vehicles navigate complex routes efficiently and safely. Selenium can operate with on-board sensors alone, removing any reliance on GPS.

A study published in 2015 estimated that 90% of vehicle crashes in the US involved the driver as the critical reason of the crash (National Center for Statistics and Analysis, 2015). The most promising solution addressing this problem is the entry of AI in cars, the so-called autonomous vehicles. While it is unlikely that fully-autonomous vehicles are commercially available before 2021, advanced driver-assistance systems will play a crucial role in preparing regulators, consumers and corporations.

Once technological and regulatory issues have been resolved, up to 15% of new cars sold in 2030 could be fully autonomous (McKinsey, Automotive Revolution, 2016). Furthermore, software-driven cars make it possible to cut emissions by 60% in comparison with today's cars; this which might be the answer to air pollution issues (Future of Driving, Ohio University 2019).

Oxbotica's AI learns by collecting and analysing the routes taken by human drivers. Over time, the software learns how humans react to certain scenarios and this leads to gradual improvement of its autonomous capabilities. (Technology Review, 2016). This allows for compatibility with other vehicles such as self-driving pods or warehouse trucks.

Advances in mobile robot autonomy will change how people interact with machines and this will have numerous impacts to how we live and work. Oxbotica is in partnership with AXA XL, who wants to understand how these changing operating environments will affect risk and liabilities. This could create challenges and opportunities for casualty insurers (AXA XL, 2016).

Business, finance, insurance and service sectors

An emerging use of AI in banking and finance is the use of collaborative filtering. This is the approach whereby recommendations are made for Person A based on the preferences of Person B. These recommender (filtering) engines are typically based on large data sets which compare aspects of demographics and behaviour across large groups (collaborating) of users, looking for people with similar tastes or behaviours.

Well-known examples include Netflix of Amazon, but examples are beginning to emerge for example in private banking for recommendations on investments based on a client's risk profile. The implications of bias and ethics are important issues for financial services firms subject to the FCA's principle of Treating Customers Fairly (TCF).

Credit card fraud is also an area that can benefit from AI approaches. Various models can be applied to existing data, and their evaluation metrics tuned to compare optimistic, pessimistic and weighted approaches. This enables trials to compare the false alarm rate to be identified so that banks can choose their preferred strategy for fraud detection against the tolerable false alarm rate (Kültür & Çağlayan, 2017).

Box 5: Lloyd's of London and AI

Lloyd's has used AI software to build a system that is now starting to provide fast and accurate knowledge to the Lloyd's International Trading Advice (LITA) team. The LITA team responds to time-sensitive trading questions from the Lloyd's market. Their response needs to always be accurate and timely, as it can mean the difference between winning or losing business for a Lloyd's market participant.

The LITA team faces a growing level of demand from their Lloyd's market customers and an increasing level of complexity in the trading questions being asked of them. The suggestion was made to test AI software, to assess if it could help the LITA team reduce the time spent searching through tens of thousands of unstructured text documents to have an increased level of throughput, whilst maintaining the team's high level of service.

A proof of concept demonstrated the business value of the AI software, reducing to seconds the search time for queries that would usually take several minutes whilst maintaining accuracy of reply. The outcome has been an AI system to be used by the LITA team. This successful proof of concept led Lloyd's to explore additional opportunities for the use of AI software globally across the Corporation.

This has led to another instance of AI being deployed in the Lloyd's Canada team to ingest and analyse contract documents to ensure local regulatory compliance.

Lloyd's has also created the Corporation AI Strategy covering Principles, Ethics and Opportunities that will guide the future use of AI.

Beyond this, the London Market Group which represents London's insurance and reinsurance market has developed an online platform to deliver a range of services to support businesses in this sector. The platform means that AI-driven services can be made available with minimal barriers for this fast-growing and dynamic marketplace.

Lloyd's syndicates and managing agents are already exploring AI internally. Therefore, there is a wealth of expertise for Lloyd's customers that can be integrated into emerging services.

Blockchain and AI

Blockchain is the technology that powers Bitcoin, a virtual currency (or 'cryptocurrency') now used by millions online. The unique property of Blockchain, and similar Distributed Ledger Technologies (DLT), is the ability to store tamper-proof data without relying on a single person or organisation.

For example, Bitcoin needs to track who owns which coins without relying upon a central bank. The security of data stored within DLT is derived from many parties keeping copies of the data in the system (it is 'distributed') and the use of strong cryptography to ensure consensus on the integrity of that data – this prevents individual parties corrupting the data.

Alternative uses for DLT are now emerging beyond cryptocurrencies, including secure digital records, e-voting and healthcare. Many of these applications are enabled by the fusion of AI and DLT technologies to deliver secure autonomous decisions over distributed data.

Box 6: The National Archives

The National Archives (TNA) at Kew, London are charged with maintaining the historic integrity of the UK's public records. Through the ARCHANGEL project, TNA are experimenting with Computer Vision/AI to visually 'fingerprint' documents as they are received by the archive.

Those fingerprints are stored indelibly in a Blockchain and can be used to verify that documents inside the archive have not been tampered with or accidentally corrupted during its care.

The need for AI to perform fingerprinting arises from the surprisingly dynamic nature of digital archival records. Rather than gathering digital dust, it is part of TNA's curatorial duty to continually format-shift digital videos in its custody to ensure they remain legible, as video file formats change and can become obsolete over the years.

By using AI to fingerprint the content (people, activities, audio) in the video, rather than rely upon classical hashing based upon bits and bytes, the fingerprint will not change no matter the format of the video.

The use case generalises to any situation where audit trails must be maintained for digital documents stored over longitudinal time periods, where any digital file format technology may be assumed to be in continual flux.

Another intriguing avenue for the fusion of AI/DLT is in Federated Machine Learning (FML). FML enables multiple independent parties to collaborate to train an AI model.

Traditionally, FML has focused upon reducing the time taken to train models, for example across a large corporate compute cluster or cloud (such as Microsoft Azure or Amazon Web Services). However, DLT holds the potential to enable large-scale FML in untrusted situations, where each node in the cluster may be operated by an independent entity with the potential for each node to act adversarially.

In such situations, AI models may be trained collaboratively without different parties necessarily needing to share their training data (which may be proprietary) whilst ensuring that no individual party can degrade or corrupt the model. Examples might include AI models for detecting fraud in financial situations, where no individual party may share their data, but all parties have a vested interest in collaborating to produce a performant model.

Finally, potential exists for the use of AI to monitor the transactions with DLT infrastructure such as cryptocurrency payments to reduce fraud or detect illegality.

Cryptocurrency payments within most popular frameworks (such as Bitcoin or Ethereum) are anonymous yet public. Whilst it is difficult to determine the physical identity behind an individual Bitcoin wallet, it is trivial to determine the contents of that wallet, and to where and from whom it has received funds.

AI may be used to perform traffic analysis through the detection of anomalies or characteristic fraud patterns within the high volumes of transaction on such networks, which would be impossible to pick out through manual observation.

DLT currently faces problems with its scalability and sustainability. This is primarily due to dominant DLT platforms relying upon 'proof of work'; the mining operations that maintain common cryptocurrency networks like Bitcoin. Although sustainable alternatives have emerged (e.g. proof of stake, proof of authority), these have not yet been convincingly demonstrated at scale.

Many of the DLT applications that incorporate AI are in the early prototype stages, either within start-ups or collaborative academic research projects. Nevertheless, the success of these early stage prototypes at demonstrating secure autonomous decision-making across diverse application domains suggests that AI/DLT applications will begin to trickle into mainstream use over the coming few years.

Societal and international security impacts



2. Societal and international security impacts

Societal impacts

Skills and competences for the changing jobs market

As the job market changes and evolves with the opportunities and challenges triggered by new technologies including AI, skills and competences need to change too. This is however a much more complex and multi-dimensional challenge than it might first appear.

The rapid emergence of AI as key technology in an already fast-moving landscape of data-centric innovation has reinforced the importance of addressing the need for a suitably-skilled workforce (Abbott & Bogenschneider, 2018).

The lack of suitably-skilled experts in AI is one of the more significant challenges for its greater adoption. This has been recognised by governments and recruitment companies, and articulated in the *Growing the Artificial Intelligence Industry in the UK Report* as being due to AI being:

There is both a need for AI experts at the specialist level, i.e. those equipped with both PhDs and Masters degrees in AI, but also for specialists in other fields, such as robotics, business administration, medicine, chemistry etc. to acquire viable AI skills.

The changing profile of skills and competences for the workforce has yet to be fully understood. It has been said that the jobs young schoolchildren of today will do have yet to be created. That may be true, but we also must better understand the competences needed for the jobs of tomorrow.

The range of skills required to understand, develop and exploit AI is comprehensive. These skills can be broken down into a suite of distinct but complementary professional profiles that can be considered under the umbrella term 'data science'.

However, the machine-learning community has used languages such as Python and R to create open-source software packages which enable individuals with minimal training or expertise to very simply set up a neural network architecture with pre-set evaluation metrics and optimisation techniques and run it against their own data. Obviously, there are dangers of misapplication/misinterpretation when deeper knowledge is missing, and more structure is needed for more comprehensive solutions.

There is moderate consensus on the terms 'data science' and 'data scientist', but there is no clear definition of the skills and competences that are contained within these terms. The terminology and roles have evolved over recent years; however, that has not been paralleled by a harmonisation of quantifiable characteristics.

The structuring of data science skills and competences matters as students will need to acquire confidence and capability, not simply awareness of various techniques. One of the challenges of delivering a fit-for-purpose workforce in an accelerated timeframe is enabling a range of qualifications ranging from traditional, longer degrees to shorter, more flexible certificates.

There remains also a need to reduce the gap between supply and demand, with universities and training organisations needing to deliver an appropriately-skilled workforce to fully exploit the benefits of AI.

The future workforce

The transformational nature of AI has several key characteristics with respect to the future workforce. New professional profiles will be needed. Once defined, teaching and training will be necessary to deliver the economic benefits to countries.

In the UK, the Hall and Pesenti report proposes that 300 industry-sponsored Masters-level places per year be established, and 200 PhD-level places per year be offered by universities with government support (Hall & Pesenti, 2017). Some of these places will be filled by new students, some could be mid-career transfers. The challenge however, will be the cost of financing lifelong learning options.

The downside of the AI economic expansion is the potential impact of traditional jobs diminishing, even if this is offset by many more jobs being created at higher salaries. The Deloitte report, *From Brawn to Brains. The Impact of Technology on Jobs in the UK* (Deloitte, 2015), articulates these predictions in some detail. Low-skill, routine-based jobs will be hardest hit; some 800,000 such jobs have already disappeared over the 15 years to 2015.

This phenomenon is not new, as technology has been replacing mundane, dangerous jobs and tasks for over 150 years. In fact, the report recalls the impact of knitting machines at the time of the first Queen Elizabeth in 1589. The report details how job replacements will occur around the country, with salary being a strong indicator of propensity for transformation.

In other words, AI will be best suited to automating low-skill jobs but in doing so, the process will create a demand for higher-skilled cognitive and socially-interactive jobs. This includes jobs relating to caring, leisure and other service occupations which have been growing for now.

The government therefore faces a challenge in managing the expectation and fallout of skills shortages, redundancy and the appropriateness of teaching and learning establishments to support this transformation.

The 2018 Organisation for Economic Co-operation and Development (OECD) report *Automation, Skills Use and Training* looks at the developments in AI and machine learning and the implications for jobs and skills (Nedelkoska & Quintini, 2018). Earlier reports had posited that 47% of US jobs were at high risk of automation. However, this new report has investigated roles in more detail and explored how jobs with the same title have significant differences in the tasks involved.

That said, there will be significant disruption to many professions and trades. Close to one in two jobs are likely to be significantly affected by automation across OECD countries.

This disruption will be unevenly distributed across regions, so many local economies will suffer more than the decline of the Detroit car industry in the 1950s. "More generally, jobs in Anglo-Saxon, Nordic countries and the Netherlands are less automatable than jobs in Eastern European countries, South European countries, Germany, Chile and Japan" (Nedelkoska & Quintini, 2018).

Significantly, the risk of automation will not be distributed equally among workers. Automation is found to mainly affect jobs in the manufacturing and agricultural industries. However, with machine learning able to handle Natural Language Processing tasks, many repetitive activities within professions such as law and insurance are targets for transformation especially when documents are already stored in digital formats.

This will affect the insurance industry as it will change practices around work that will affect every one of us. AI will have an influence on jobs, and careers with routine, rule-based tasks are being replaced by AI-enhanced machines.

Whilst early dramatic predictions of 40% job losses have been refined by awareness that it is tasks rather than jobs that will be washed away, this will have profound repercussions for career paths and hence the personal and professional financial instruments that support them (Giovannini & Scapolo, 2018).



Focus on the future

Underwriters of the future will have to expand and develop new skills such as sales, marketing and analytics (EY, 2015). Insurers should equip their workforce for the future by providing the training and knowledge needed.

Moreover, a large component of whether an AI project will succeed, or fail is how well (or not) the change management process in an organisation and team is managed. AI technology might be great, but if staff using it is brought into the concept from the start and provide support the AI project will most likely fail.

Potential responses to labour disruption

To protect their citizens, some countries have championed the idea of a universal basic income (UBI) for all, and this has been trialled in Scotland, USA, Canada and Finland amongst others.

Observed pros include the reduction of poverty and income inequality, health improvement, job growth, decrease in school dropout rates and empowerment of important unpaid roles such as non-working parents and caregivers. However some have dismissed this, including Luke Martinelli at the University of Bath Institute for Policy Research who said: “An affordable UBI is inadequate, and an adequate UBI is unaffordable” (Martinelli, 2017).

Other income protection schemes, AI tax arrangements or evolution of current protection insurance solutions might have to emerge to respond to the impact of AI. The broader societal effects of AI are potentially profound and will require a considered response from beyond the technology community delivering the services.

International security and defense impacts

The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation report (2018) investigated the malicious use of AI. It found that threats such as the use of autonomous drones as weapons, intrusions into manufacturing and service sectors such as the energy grid, disinformation campaigns, and conventional dedicated denial of service attacks could all be affected by advancements in AI.

The report concluded that AI could lead to three main changes to these threats:

1. **Expansion of existing threats** due to AI reducing the cost of initiating attacks and increasing the rate of attacks;
2. **Introduction of new threats** because AI could complete tasks that would otherwise be impractical for humans, but could also have its vulnerabilities exploited; and
3. **Change to the typical character of threats** as the growing use of AI could result in attacks being especially effective, finely targeted, difficult to attribute, and likely to exploit vulnerabilities in AI systems.

Change in frequency and severity of threats have a critical impact on pricing insurance and determining capital levels. These impacts will have to be assessed by underwriters and actuaries.

The report presents these changes to threats in three domains:

- **Digital Security.** AI will automate tasks involved in carrying out cyber-attacks and will be able to exploit human vulnerabilities and those of other AI systems.
- **Physical Security.** AI cyberattacks could cause direct physical damage (e.g. autonomous vehicle crash).
- **Political security.** The use of AI may expand threats associated with privacy invasion and data manipulation resulting in political instability.

The risks associated with the malicious use of AI and its impacts are key areas of concern for society. Potential weaponisation could take many forms including the corruption of data, unbalanced selection of data, illegal use of data leading to myriad outcomes including propaganda, behavioural change and deception.

On the physical side, ‘autonomous weapons’ have attracted concerns and call for the state to develop them responsibly. Semantic challenges have been raised as each country will have different definitions for terms, and varying approaches to the legality and morality of releasing swarms of autonomous weapons into the battlefield programmed to collaborate and achieve a designated goal at any cost.

The debate centres on the concept of meaningful human consent, and who ultimately has responsibility for casualties, whether it is the operators in the battlefield, the commanders, or the programmers of the underlying AI systems. Further concerns are raised when open-source AI technologies are being used for military projects, such as the US Department of Defense’s use of Google’s TensorFlow for its Project Maven to analyse drone footage (The Guardian, 2018).

In 2017, the UNODA (United Nations Office for Disarmament Affairs) hosted a panel to discuss the pathways to ban fully-autonomous weapons. The event concluded that there are multiple technical, legal, ethical and operational concerns over fully-autonomous weapons and an international ban treaty is needed (UNODA, 2017).

Ultimately the degree of security of AI systems depends upon what type of AI is being employed. As most current AI usage relates to machine learning, which in turn relies upon the model that is used to train the system, vulnerabilities can arise in several ways.

For example, the data set used to train the model may be corrupted in some way: you can introduce bias or even render it completely useless.


The more insidious bias attacks an area that is not fully understood but we are learning that some models develop bias inadvertently. This has prompted areas of research into how bias arises and the sensitivity of the system to the training data set.

Other sources of interest, include:

- Biggio and Roli (2018) have provided an excellent summary of adversarial machine learning in the last 10 years across multiple application areas.
Their research discusses the security problems found with early machine-based learning systems, which also appears to be present with DNNs (albeit these systems are harder to corrupt).
- Research by Papernot et al (2017) identifies practical black box attacks against neural networks to obtain misclassifications at a very high rate.
- An excellent plain English write-up of how modifying street signs can fool machine learning is provided by Ackerman (2017) in the IEEE Spectrum journal.

How such minor modifications can lead to such dramatic misclassifications and how to make the models much more robust is an area of active research.

Risks



3. Risks

Whilst AI offers benefits in many areas of life such as healthcare, transport and manufacturing, there are also risk areas that give cause for concern. The late Stephen Hawking and Elon Musk have highlighted the significant risks to humanity if, or indeed when, the current narrow AI applications evolve into strong AI systems.

These technological possibilities present challenges to humankind and our planet because they challenge our existing mechanisms for managing responsibilities, ethics, regulations and liabilities. We see four critical risk areas associated with AI.

Transparency
and trust

Ethics

Liability

Security

Transparency and trust

One of the main concerns over AI is where an effective black box is created with an unknown set of rules and conditions which may have inappropriate or even illegal tests created to reach a desired conclusion.

AI black boxes are created in the form of DNNs whereby the machine-learning process develops layers of neural networks. Between the input layer and the output layer are multiple layers within the neural network which have been trained and revised by the machine-learning process.

Transparency of the decision-making process and the ability for everyone to understand how an AI got to a specific conclusion will be key to developing trust in the technology. There is growing interest in the need for explainable systems and this need is behind the development of eXplainable AI (XAI), a 'glass box' in contrast to the black box context.

To ensure trust and transparency to the public, clearer accountability will be required for decision-making processes without sacrificing AI performance. Fairer and more representative AI solutions will only be provided because of rigorous and enforceable regulations and ethics frameworks, developed by a broader community aligned to the full diversity of the real world.

Ethics

The *AI Now Report 2017* (Solon et al., 2017) addresses the critical social questions around AI, focusing on four areas: labour and automation, bias and inclusion, rights and liberties, and ethics and governance.

One of the most critical areas that has captured much attention is bias. Bias can be added deliberately or inadvertently, through the selection of data and by the cultural background of the developers themselves.

It has already been discovered that AI systems developed just a few years ago were trained on data that contained many prejudicial biases that would be unacceptable today (e.g. gender and racial bias).

The inherent biases presented by the relatively narrow social group that has predominantly been involved in developing AI systems could be countered with the use of broader and more diverse groups in the development and testing of future systems, as well as having professionals continuously undertaking the tasks learnt by the AI system so that it adapts to contemporary beliefs.

Machine learning depends on large collections of data from which to extract knowledge. These data sets effectively become the codification of history. However, within these histories there are prejudices, biases and societal injustices that have existed for years.

The consequence is that the AI depends on the information on which it is trained, and it might act against human interests.

As an example, an algorithm that Amazon was testing as a recruitment tool was found to be sexist and thus unusable (BBC, 2018). The AI system was trained on data submitted by applicants over a 10-year period, many of whom were men, resulting in the system teaching itself that male candidates were preferable.

In 2016, a report by investigative journalism organisation ProPublica claimed that a software using AI, applied by a US court for risk assessment, was biased against black prisoners. In forecasting who would reoffend, the algorithm was flagging black defendants as future criminals almost twice the rate as white defendants, who were being mislabelled as low risk more often than their black counterparts (ProPublica, 2016).

A shift from human-based decisions to AI-based ones, where the AI is not necessarily based on any moral code and does not have a common-sense filter nor a fear of wrong decisions, punishment or a sense of accountability, is also an important issue.

Addressing this is not simply a case of attempting to create explainable systems through which somehow decisions can be unpicked. A degree of oversight is needed to ensure that AI systems are taught to make decisions that are appropriate for the time we live in and the times ahead.

Ethics and responsibilities will remain an important and critical theme that permeates all aspects of AI and can only grow in complexity and intractability as the field develops. Bias can enter in many forms, either unintentionally through systems designers and software engineers, or systemically through the data sets or the methods of data analysis.

The implications of bias, ethics etc are important issues for financial services firms subject to the FCA's principle of Treating Customers Fairly (TCF). A core part of TCF is that it should be seen as an integral part of a firm's culture in an ongoing developmental process. As Cathy O'Neil states in her book *Weapons of Math Destruction*, "big data codifies the past":

Furthermore, TCF's focus on the consumer aims to achieve capable and confident consumers and simple and understandable information.

The conjunction of algorithmic bias, introduced maliciously or not, and the black box where we can't explain how a conclusion was reached, might also have profound implications for a firm's products and services and reputational risk.

Liability

One of the significant challenges is the legal status of AI systems and services. Whilst much can be done to encourage and motivate the engineers behind the systems to consider the ramifications of their decisions in designing neural networks and selecting data sets for learning, there needs to be robust legal frameworks to define liability.

Medical, vehicular and financial systems will be making life or death, or at least life-changing decisions. Either someone needs to bear responsibility or else AI systems or robots might need to have a 'legal identity'; the term used to establish which entities have legal rights and obligations and can enter into contracts or be sued.

Legal frameworks and ethical guidelines are being developed at different paces at national and international levels, bringing together a much broader constituency than those who created these technologies.

Governments and other policymakers are waking up to the fact that they have responsibilities in ensuring that their constituencies reap the benefits of AI, but also are protected from the potential negative outcomes. Questions range from "can our education system deliver the dynamic, collaborative, data-science-savvy workforce primed for lifelong learning", to "can our industries harness the power of AI-enhanced decision-making at every stage".

Generally speaking, the manufacturer of a product is liable for defects that cause damages to users. However, in the case of AI (especially strong AI) decisions, they are not a consequence of the design, but of the interpretation of reality by a machine.

Security

The real harm that unauthorised access could do is very context sensitive. If AI were being used on, for example, a safety critical system, it could be altered at source to render the systems dependent upon it at best, unusable, at worst unsafe. Much of this will come down to supply chain security.

We have seen how recent attacks have begun to focus on the weakest link in the supply chain of code to introduce their malicious elements. This has now even affected online payments where campaigns such as that from Magecart are finding ways to inject code into third-party software incorporated in some well-known brands' website payment pages (including British Airways, Ticketmaster and an estimated 800 others, so far detected).

The same was true of the NotPetya ransomware which was spread via a tiny Ukrainian company whose software was mandated by the government for accounting purposes, and yet resulted in firms such as Maersk having to write off billions of pounds due to the disruption caused.

There are some specific security concerns and risks (but also opportunities) related to the open-source accessibility of AI. Closing software prevents misuse and misappropriation of technology by AI developers who might have harmful intentions, and it might be difficult for non-first movers to develop counter-activity measures.

The history of the cybersecurity field has shown that open-source software of any sort is a bit like the Curate's Egg: good in parts. There is the "many eyes" argument which posits that if software is open source, so many people will be examining it that any tampering will be found.

For example, OpenSSL introduced the notorious Heartbleed flaw through an error (never mind a malicious act) and it was not picked up until a researcher noticed the effects rather than seeing the problem in the code.

There is also the danger that code can be copied and repurposed. It is commonplace amongst malware developers to evolve software that has been written by others. Malicious code can also often contain code snippets that have been published as a proof of concept to demonstrate the very vulnerability that the more advanced malicious code then exploits.

The astonishing pace of technological achievement in AI has driven significant new thought around its potential adverse impacts, and its effect upon society as AI deployment becomes ubiquitous.

Current AI technologies enable automated decision-making around data; the ability to take actions or identify patterns within the deluge of data present in the modern world. The potential for algorithmic bias in AI systems – either due to poor design or bias within the data upon which the system is trained – presents new risks, since we will become increasingly reliant upon AI technology to identify relevant data and act upon it. Most recently this has become apparent on social media platforms, where demographic profiling and popularity influence the prominence of news stories.

This has led to concerns around AI's role in the polarisation of opinion in modern society and at critical

times. Events such as government elections and referenda offer opportunities to exploit known behaviours in such algorithms to bias widespread public opinion.

More generally, the intersection of cybersecurity and AI is poorly researched at present. Whilst several cybersecurity systems exploit AI e.g. to assist in the detection of hacking and network intrusion, there is very little cybersecurity analysis of AI systems and the potential attack surface presented by such systems. This is partly due to the reliance of AI systems upon DNNs which are difficult to debug or interpret. To address this concern, research is being undertaken in explainable AI – to go beyond evaluating the performance of DNN to understanding why a DNN performs as it does.

Techniques are being developed to uncover bias in the training of a network (e.g. to identify whether it is making biased decisions or making decisions that illegally discriminate), or to identify ways in which a network can be fooled by 'adversarial inputs'.

For example, it was recently shown that visual recognition systems could be fooled into misclassifying objects or even omitting to sense objects altogether by introducing an 'adversarial sticker' into the scene (BBC, 2017; Open AI, 2017).

The ability to blind a DNN to the presence of objects clearly visible in an image underlines the differences between human perception and existing AI technologies. It also highlights the importance of ethical discussions around resilience and explainability of AI as well as data privacy and algorithmic bias in current machine learning technologies.

Finally, it is worth noting that AI is also forming a new dimension in cybersecurity as it can learn and react to threats much faster than the traditional methods used by security products of old.

However, not surprisingly, criminals are also beginning to use AI to learn how to conduct all stages of a typical attack: from reconnaissance to crafting a specific attack. It is the inevitable arms race.

As these systems become more complex and internally connected, breaches of (cyber) security for interconnected AI will become more important and are likely to have systemic economic impact.

Regulatory and government landscape

4. Regulatory and government landscape

Challenges

AI challenges existing regulatory frameworks in virtually all areas of the law. Professor Ryan Calo (2017) categorises AI regulatory challenges in terms of justice and equity, use of force, safety and certification, privacy and power, and taxation and displacement of labour.

He also notes that AI raises broader systemic questions related to institutional configuration and expertise, investment and procurement, removing hurdles to accountability, and correcting flawed mental models of AI.

Today, there is little AI-specific regulation, and most regulation relevant to AI has been reactive. However, recent high-profile scandals involving primarily self-regulating, large technology companies may result in more proactive AI regulation moving forward.

Regulatory approaches to AI also vary significantly by region. For instance, the European Union has a relatively consumer- and human rights-centric approach, whereas countries like the People's Republic of China are more industry- and government-centric.

In the absence of a robust regulatory regime, individual companies like Microsoft, Google, and Facebook as well as industry associations like IEEE, Insurance Europe and the Partnership on AI have started to promulgate ethical principles to guide the use and development of AI.

More strongly, Insurance Europe points out that the insured has the right to be informed about an algorithms' *logic* in decision-making and the plausible consequences. Data subjects can challenge an algorithmic outcome by requesting human intervention.

Government landscape

United Kingdom

The past few years have seen increased levels of interest in AI regulation in the UK. In 2017, Dame Wendy Hall and Jérôme Pesenti (Hall and Pesenti 2017) co-chaired the government's review and analysis of the country's current AI capabilities.

Their subsequent report made proposals for how the UK could maintain a position as a world leader in AI. In 2018, the House of Lords (2018) released its report, *AI in the UK, Ready, Willing and Able?* that made further recommendations for enhancing the UK's AI capabilities.

Both reports fed into the government's 2018 Industrial Strategy that also promotes AI and seeks to establish the UK as a leader in ethical uses of data and AI (Industrial Strategy 2017).

These reports detail a variety of government initiatives to promote AI, such as establishing a Centre for Data Ethics and Innovation which will advise on "ethical, safe and innovative uses of data, including AI".

Meanwhile, a variety of field-specific regulatory responses to AI have been developed (Abbott, 2018, 2019). For example, the UK passed the first European law to deal specifically with accidents caused by autonomous vehicles that requires insurance companies to directly compensate victims for harms caused by autonomous vehicle malfunctions rather than manufacturers, and which defines persons occupying an autonomous vehicle in self-driving modes as 'passengers'.

The UK's FCA has already seen the potential benefit of using AI in identifying fraud and other suspicious activity in the vast amount of data that exists in the financial world. For example, credit cards alone create 100 billion data points over a five-year period.

The FCA has successfully developed its own AI systems to provide single-risk scores from complex sets

of company data, as well as linguistic algorithms to evaluate advertising promotions to check that risks have been appropriately highlighted in the text.

Beyond this, the FCA is allowing businesses to test innovative products, services, business models and delivery mechanisms in the real market, with real consumers through its sandbox.

Box 7: RegTech and SupTech

Following on from the success of the FinTech sector, RegTech exploits the use of disruptive technologies to revolutionise the regulatory sector that has expanded following on from the financial crises of 2008. RegTech consists of businesses complying with regulations efficiently and less expensively, thanks to the use of new technologies such as blockchain and AI.

Using big data and AI, RegTech companies aim to reduce the risks that company compliance may suffer by automating processes such as monitoring transactions to detect money laundering, identifying threats to financial security and automating the collection and storage of customers' due diligence data.

The global RegTech investment in Q1 2018 has already surpassed US\$500m and it is expected to increase due to pressure from regulations such as the General Data Protection Regulation (GDPR). North American companies dominate the RegTech landscape (60% of deals between 2014 and 2016) followed by Europe (Global RegTech Summit, 2018).

SupTech is the application of RegTech in financial supervision. In the UK, the FCA has been consulting on RegTech since 2015, and has built several proofs of concepts that have explored the possibilities of using Natural Language Processing and AI technologies to automatically build and manage a compliance programme to interpret legislation (FCA, 2018).

More recently, the Blockchain Technology for Algorithmic Regulation and Compliance (BARAC) project in collaboration with University College London investigates the feasibility of using blockchain technology for automating regulation and compliance.

Other supervisory agencies such as the US Securities and Exchange Commission, Central Bank of the Republic of Austria and the Netherlands Bank are exploring the use of unsupervised machine learning and neural networks application in data analytics (BIS, 2018).

European Union

The European Union has been relatively proactive at regulating new technologies, perhaps most importantly, through the General Data Protection Regulation (GDPR) which took effect in 2018.

GDPR is the most important change in data privacy regulation in two decades. This regulation is critical to AI because data is vital to its use and to training types of AI that rely on machine learning. The GDPR gives consumers and end-users a high degree of control over their data and requires entities that collect data to receive permission for how they use that data.

Among other things, it also gives consumers a right to explanation for automated decisions by AI, such as for approval of consumer credit. This approach has been lauded by privacy and human rights advocates but criticised as potentially detrimental to industry.

The EU has also been working on a harmonised framework for AI, robotics and new technologies. In 2017 the European Parliament published a resolution on civil law rules on robotics that, among other things delivered what was considered to be a controversial recommendation that called on the European Commission to:

“explore, analyse and consider the implications of all possible legal solutions, such as: creating a specific legal status for robots in the long run, so that at least the most sophisticated autonomous robots could be established as having the status of electronic persons responsible for making good any damage they may cause, and possibly applying electronic personality to cases where robots make autonomous decisions or otherwise interact with third parties independently” (European Parliament, 2017).

In April 2018, the European Commission outlined its AI strategy and established a High-Level Expert Group on Artificial Intelligence to consider the breadth of societal and economic challenges and opportunities as a part of the evolving vision of the Digital Single Market. In December 2018, this group published draft ethical guidelines for trustworthy AI (Ai Heg 2018).

The European Commission strategy did not adopt the proposal of legal personhood for AI or robots, but it commits to ensure an appropriate ethical and legal framework since “artificial intelligence may raise new ethical and legal questions, related to liability or potentially biased decision-making” (European Commission, 2018).

The European Commission also confirmed that by mid-2019 it will be publishing guidance on the interpretation of the EU Product Liability Directive in the light of technological developments “to ensure legal clarity for consumers and producers in case of defective products” (European Commission, 2018). It has also created an expert group to consider broader issues related to liability of new technologies.

The challenge of the ethical context related to AI has been identified by the European Commission’s group on ethics in science and new technologies, who has warned that existing efforts to develop solutions to the ethical, societal and legal challenges AI presents are a “patchwork of disparate initiatives” (Hill, 2018).

As a result, the EU has proposed a set of fundamental ethical principles, based on the values laid down in the EU Treaties and the EU Charter of Fundamental Rights, that can pave the way towards a common, internationally-recognised ethical and legal framework for the design, production, use and governance of AI, robotics, and autonomous systems (European Group on Ethics in Science and Technologies, 2018).

The European Securities and Markets Authority (ESMA) has worked with the European Insurance and Occupational Pensions Authority (EIOPA) to identify the unfolding potential impact of AI, for example, in developing investment strategies through AI tools and also recognising the need for skills to be developed across organisations to enable successful collaboration on AI tools (ESMA, 2017).

Box 8: e-Estonia

Within the context of Europe, Estonia is a leading exemplar of digital vision with their e-residency and e-governance services. Estonia’s policymakers are working with citizens to develop legal frameworks for the application of AI technologies to everyday life tasks of Estonian citizens having realised that individual technologies such as self-driving buses cannot be considered in isolation (e-estonia, 2017).

In March 2018 Estonia declared that the country will have an AI strategy and it will launch a cross-sectoral project to analyse and prepare its implementation (Government Office of Estonia, 2018). The Ministry of Economic Affairs and Communications also sees AI as key to attract new investments and innovation activity to Estonia and will offer technology firm development and test environment that favours AI solutions (Government Office of Estonia, 2018).

US

In the US, the Obama Administration published the *Artificial Intelligence, Automation, and the Economy* report in 2016 (Executive Office of the President of the United States of America, 2016). The report recognised that the great potential that AI could offer in terms of productivity would be offset by an unknown unevenness in the distribution of benefits across the workforce.

However, this profound challenge can be addressed: as the report concludes: “Technology is not destiny; economic incentives and public policy can play a significant role in shaping the direction and effects of technological change.”

The report explored the risks surrounding the existing workforce in some detail. Whilst in general college-educated skills sets will continue to rise, with lower-skilled and less-educated workers being required less, there will be other factors such as a need for “abstract thinking, creative tasks, and problem-solving”, not to mention domain-specific skills to train and develop AI systems.

The current US Administration wants to enable “the creation of new American industries by removing regulatory barriers to the deployment of AI-powered technologies” (Whitehouse, 2018).

In February 2019, President Trump signed an Executive Order launching the American AI Initiative. The initiative wants to accelerate US leadership in AI through five key areas:

1. Investing in AI Research and Development (R&D)
2. Unleashing AI Resources
3. Setting AI Governance Standards
4. Building the AI Workforce and International Engagement; and
5. Protecting US AI Advantage

(The White House, 2019)

China

China is one of the world leaders in the AI and the government sees it as part of the strategy to become a “science and technology superpower” (Ding, 2018).

In July 2017, the State Council issued “A Next Generation Artificial Intelligence Development Plan” setting a three-stage plan for the development of AI in China, concluding in 2030. The final goal is to make China the world’s primary AI innovation centre.

The plan concentrates on seven technical sectors:

1. Connected vehicles
2. Service robots
3. Unmanned Aerial Vehicles (UAV)
4. Automated diagnostic systems
5. Video image recognition
6. Artificial audio intelligence; and
7. Computer-driven language translation

This last technology could have a profound effect on the legal and insurance sectors. Automated translation services could significantly speed up operations and business development in foreign markets, which could drive efficiencies for customers.

AI and insurance



5. AI and insurance

Impact of AI on insurance lines of business

Product liability and product recall

There are many ways in which a product reliant on AI could be defective. Even if product testing of an AI-reliant machine is rigorous, there might be situations that an 'autonomous' machine encounters a situation that it has not been accounted for during testing.

This may cause halt, breakdown or malfunction and the defect in design could be also be dangerous. Safety will be paramount for human/machine interactions. Product defects could even result from communication errors between two machines or between machine and infrastructure.



Insight

Recalls could become larger and more complex, particularly if the affected sector uses AI extensively.

For example, if a series of accidents raises concerns about the AI technology employed in driverless cars, it could trigger a large recall across different manufacturers and countries along with expensive and long court case where time is taken to prove the design of the AI system.

As explored in Section 4 under existing law, AI machines cannot themselves be liable for negligent acts or omissions that cause damage to third parties. But as machines become better at thinking, learning and deciding, the question "who is liable when a machine commits a tort?" is raised.

Product manufacturers/sellers, AI designers/suppliers and AI purchasers/users might be allocated fault by courts. In the future, insurers might see an increase in companies using contractual warranties, indemnities and limitations to control AI liability risk.

Third-party motor liability

Assignment and coverage of liability will become more challenging in the future due to the possible shift of responsibility from human drivers to automated vehicles (AVs), and therefore to the manufacturers.

There is also significant potential for different legal and regulatory responses in different countries creating a complex international risk landscape.

In tort law jurisdictions, current automobile insurance policies are generally based on the principle that the vehicle user is liable for losses caused by both individual driving mistakes and defects of the vehicle due to lack of maintenance.



Insight

With fully autonomous driving, this liability may shift from insureds to vehicle manufacturers and/or their suppliers, such as AI software providers.

In some territories, new liability models extending product liability law may be adopted, where manufacturers take over third-party liability for product defects.

As a result, the market for third-party liability cover may reduce significantly in some territories, perhaps with a related increase in the demand of manufacturers' product liability insurance.

In the UK, in respect of vehicles capable of automation level 4 or 5^b legislation in 2018 has confirmed that the insured will be able to claim under a third-party motor policy in the first instance.

The insurer will then be able to pursue the appropriate recovery from the relevant at-fault manufacturer.

This avoids the situation of individuals being excluded from compensation (as a first-party claimant with no cover under the policy) or from having to pursue indemnity from manufacturers on a product liability basis which could be a complex and lengthy process.

^bAutomation Level 4 (where a car can drive itself most of the time without human input, but in certain conditions it might not be able to) and Automation Level 5 (full automation in all conditions).



Insight: Professional indemnity and robo-advisers

Robo-advice is advice generated by an algorithm. Liability may arise from negligent advice given by a firm, but not by a human professional, who would be typically covered under a professional indemnity policy. This differs from product liability, as there may be no injury sustained or the system design may not be defective. However, the advice the AI gives out could still be negligent.

But what happens when AI is used to give professional advice? There have been advances in Robo-adviser capabilities. This ranges from Robo-advice for investments to potential applications in LawTech and even chatbots that teach Cognitive Behaviour Therapy (CBT), using algorithms to generate messages (Techemergence, 2018; Woebot, 2018; Hudsonmckenzie, 2018).

Robo-advice for investment tends to use a platform that takes some parameters given by the consumer and with algorithms, constructs a portfolio of assets (typically from buying Exchanged-Traded Funds). Such Robo-advisors use AI to varying extents in their selection, some using machine learning to simulate past performance (Techemergence, 2018).

To date there is not much 'advice' given to the consumer about what their risk appetite should be or how much they should invest. If this changes and Robo-advice takes on an increasingly professional dimension, then liability for negligent advice could arise.

This would also apply to any Robo-advisers giving legal advice, or any of the key areas where negligent advice amounts to a breach of contract or is tortious. This creates challenges and opportunities for professional indemnity insurers. It is complicated in the transition phase, where professionals use AI to assist rather than AI doing all the decision-making. This is where liability may still be assigned to the professional indemnity policy if the professional ultimately is at fault.

Some chatbots are being integrated into other platforms such as Facebook. A therapy chatbot like Woebot gives check-ins and offers step-by-step guidance. The human-like interaction is thought to help engage people and offer a compassionate, free psychological service (Woebot, 2018).

As with most existing chatbots, they do not give professional advice on which someone would solely act. This could change in the future as AI matures. Who would be liable for the advice given; the providers of the application or the designers?

There is a need for liability to be clearly outlined in the contract. If personal data is not limited to just simple details but enough data is used to provide a full psychological profile of someone, then any data breaches could be very serious and could potentially expose vulnerable individuals.

How these chat history records are stored also creates problems. If they are kept, this increases the chance of a significant data breach. However, if they are deleted, this could create problems in litigation, especially if there is a latency between the advice and the presentation of a claim, say for instance, medical advice.

Medical malpractice

There are various ways in which negligence could arise. AI is being used to help diagnose conditions and is being applied in areas such as radiography (Lee et al., 2017).

If an error leads to misdiagnosis or false positives, this could amount to negligence. Even if AI was used as an aid for referrals, if these referrals prompted investigations or procedures that were unnecessary, invasive or led to poorer patient outcomes then liability could arise.



Insight

The lack of case law could also result in larger claims in terms of defence costs. This is because the insured may wish to challenge the claimant, and with little precedent case law it may be difficult to resolve using alternative dispute resolutions. Trials may have an additional complication as AI experts as well as medical experts may be involved.

Poorly-designed AI or algorithms that are optimised for the wrong task could introduce a systematic error and result in multiple litigants.

Where liability can be clearly attributed to the AI, this could be easily excluded under a medical malpractice policy.

However, if AI is used as an aid, or is used negligently i.e. not used or set up correctly, liability and proximate cause determination could become complex.

Cyber

Chatbots are becoming increasingly close to passing the Turing test, first postulated by Alan Turing in 1950. To pass the test a computer must exhibit intelligent behaviour that makes it difficult for a human evaluator to reliably predict whether the contents of a conversation are produced by a human or a machine.

There are numerous claims that this is happening, such as the Google's AI assistant's demo-ed feature of being able to answer phone calls for you (Extremetech, 2018).

Whether or not the Turing test is passed is not necessarily the milestone that insurers should wait for. If the bot can possess some human-like characteristics, then the potential for abuse increases.

If an AI system was used to send out phishing emails/calls/messages and to respond to the receiver, then this could change the way social engineering scams are carried out.

This problem has several factors that could create a large expansion in the number of successful scams. This could result in the following:

1. Sophistication could rise rapidly. Rather than a generic template email, AI bots could personalise their initial messages. This personalisation could be extensive if the AI systems had access to any personal information, from previous breaches or data that has been purchased or is available elsewhere.
2. With machine learning, the AI bot could learn which types of emails or attacks work best. These emails might not be the most 'human-like' but may simply be able to target the most vulnerable users receptive to this type of attack.
3. The AI system could potentially be embedded into fake websites, replicating well-known ones. This could be social media, banking or fake websites that could attract receptive individuals.
4. AI phishing could be automated and sent out in large volumes.
5. The AI may not even be the fraudsters' own design. It could have been obtained through hacking or available open source.
6. A hybrid attack where humans are used in some parts of a complex fraud chain could make up for any shortfalls in the AI's capabilities.
7. While AI deployed in cybersecurity might increase the ability of the defenders to identify and prevent attacks, it might also highlight some cybersecurity defence inadequacies on the target's system.



Insight

These issues raise questions about what types of cover would be available to protect against these types of losses some of which are arising from more sophisticated and intrusive.

Whether or not an AI-driven phishing attack is sufficiently intrusive to count as a compromising attack may raise questions on coverage. What constitutes an 'insuring event' will have to be carefully defined.

Liability for third-party damages could also be complicated. If an insured's bot was hijacked and used for criminal purposes, then there may be cover under a cyber policy. This could arise if losses are sustained because of the bot being used for criminal purposes.

If the insured negligently failed to provide adequate measures to protect their bot being hijacked, then they may be found liable. Cover for first-party losses such as business interruption resulting from the insured's bot being hijacked might be more straightforward.

Fidelity

Fraudulent activity from employees could also be exacerbated by any of the methods mentioned in the Cyber section above. Fidelity insurers should consider that fraud may increasingly come from employees with access to IT systems rather than employees with financial authority.



Insight

The emergence of “deep fakes” is also concerning when it comes to fraud. Deep fakes are AI systems capable of generating realistic audio, video and images.

This is concerning for cases of identity fraud and could also be used in conjunction with chatbots to increase authenticity and build trust.

Cyber-attacks utilising deep fakes could also pose a threat in new and unexpected ways.

For example, cybersecurity experts have exposed hospital network vulnerabilities. The researchers introduced malware which used deep learning to maliciously remove or inject tumours/blood clots into 3D medical imagery such as computerised tomography (CT) scans and magnetic resonance imaging (MRI). Misdiagnosis could lead to harmful clinical outcomes, research sabotage or medical insurance fraud. This could also be used to carry out covert attacks on political leaders or terrorist acts (BBC, 2019).

Political risks

As raised in Section 3 covering AI risks, potential weaponisation could take many forms including corruption of data, unbalanced selection of data, illegal use of data leading to myriad outcomes including propaganda, behavioural change and deception.

AI systems might “take advantage of an improved capacity to analyse human behaviours, moods, and beliefs based on available data” (Brundage, et al., 2018).



Insight

From a political risk coverage perspective, AI might contribute to creating new or exacerbating existing political *force majeure* events such as expropriation, wars, acts of terrorism, civil disturbances and other forms of political violence in developing, but also developed markets.

In the past few years we have seen the (mis)use of AI in political campaigns (The Conversation, 2017). In a similar manner, targeted computational propaganda and deep fakes might enable more efficient, scalable and widespread distribution of disinformation and amplification of distracting stories (CNAS, 2018).

From a political risk coverage perspective, AI might contribute to creating new or exacerbating existing political *force majeure* events such as expropriation, wars, acts of terrorism, civil disturbances and other forms of political violence in developing, but also developed markets.

On top of the risks mentioned above, instability generated by automation and economic disruption might also be a potential driving force for protest and agitation, resulting in government interference, selective discrimination and business interruptions.

Business development opportunities

In general, any company offering algorithm-based systems to data-rich companies (e.g. fraud detection in online sales) might seek to insure against the risk of the algorithms returning incorrect decisions and its impact on the AI companies' clients.

Before insuring the (potentially self-learning) algorithm outputs, insurers will have to test its statistical theory, the infrastructure and its reliability.

Moreover, new companies are emerging in the disinformation defense area to:

- Provide technology to filter out fake news;
- Detect and eliminate trollbots; and
- Certify information and authenticity of images and videos.

This might be an opportunity for insurers to explore what type of products (e.g. professional indemnity and cyber products) could be useful to these new businesses and in what form.

As the field develops and applications increase, opportunities arise for providing risk management services. Even though we are at a relatively early stage of AI development, AI knowledge experts are already in high demand from firms attempting to manage risks.

Specialist service firms are emerging that are aimed at loss prevention. As complexity increases, the demand for these services will rise.

Operations

AI presents both risks and opportunities for the insurance sector. From an operational perspective, insurance companies are already exploiting the potential of AI to deliver value in a more efficient way.

Many, if not all stages in the insurance value chain will be transformed by AI. The competitive nature of service-driven business models is motivating all sectors to improve back-end systems around the customer experience, with attention given to lowering costs and increasing speed.

The following are the most common applications of AI:

- **Chatbots and AI customer assistants.** They can recommend and personalise products, handle complaints, improve communications with customers based on emotions analysis and process simple transactions.
- **Fraud detection** in financial services including insurance. Right now, this is a largely human-led operation. However, the greater availability of computing technology coupled with advances in AI means that fraud detection can be better automated through the use of deep reinforcement learning to explore the data, both structured and unstructured, in order to detect potential new patterns of fraud (Bouchti, Chakroun, Abbar, & Okar, 2017).

In the UK, the Serious Fraud Office has used a pilot robot to process half a million documents a day to scan for legal professional privilege content at speeds 2,000 times faster than a human lawyer (SFO, 2018).

In insurance, processes could be further automated with the augmentation of external data from historical records, sensors and images to better estimate repair and write-off costs.

- **Underwriting.** Underwriting could be enhanced and sped up through AI, especially when multiple data sets are incorporated into the process. AI-based models could predict premium based on past risks assessment and could generate quotes for a specific product tailored based on a customer's risk profile.
- **Claims.** AI could help reduce the number of claims that require human analysis and interaction by automating image recognition, searching large databases for identical claims to estimate any risk of fraud, validating claims and interacting with customers to provide updates on existing claims.
- **Modelling, exposure management and pricing.** With AI and its combination with other technologies such as IoT, insurers and model providers will have access to data that can feed into models which can learn and adapt at a much faster pace. This could for instance, forecast interest rates based on central banks' communications or help predict missing fields within a data set (Panlilio, Canagaretna, Perkins, du Preez and Lim, 2018).
- **Business development.** Business development. AI systems can improve cross-selling and conversion rates of products, can propose tailored product and can help identify new clients.



Insight: Commercial property risk engineering

The goal of a risk consultant is to be able to provide expert advice to underwriters and clients on a risk. This advice assists underwriters with their decision to take on the risk and at what price as well as the client's understanding of where they could implement improvements for risk prevention. However, risk consultants' need to capture specific risk data for evaluation to help with risk selection and management is becoming challenging due to the high volume of submissions and the information required to accurately assess a risk. To compound this, risk consulting is typically both resource limited and time sensitive. Due to an increase demand on engineering, meeting growth targets is becoming challenging. A way to meet the challenge is to increase people (not always an option due to the specialised nature of the role and expense targets) or to find a way automate much of the processing.

Imagine using artificial Intelligence (AI) in the property risk engineering world to build a capability that enables one risk consultant to think like a thousand. By automating the process of reading risk engineering reports through natural language processing (NLP), risk consultants' productivity can be greatly augmented by leading to more informed underwriting and higher level of client satisfaction. Automation can also relieve the pressure on an engineering competency that is highly resource and time pressured. Imagine being able to take that vast, rich source of data from engineering reports and store it somewhere to provide key insights at account or portfolio level or how about assisting the underwriter at renewal time or prospect to identify similar risks based on similar data points.

Artificial Intelligence can be the answer, specifically natural language processing. Using natural language processing, engineering reports can be 'read' by a machine within minutes to extract the knowledge and insights from multiple pages (sometimes as many as 100 or more). By using the knowledge from risk consultants to enrich the NLP understanding of engineering (commonly known as a knowledge graph or ontology), a solution can be built to understand the content of engineering reports, identifying engineering concepts that can be pulled from the document to be used for purposes such as risk assessment or data insights. Further to this, more complex NLP algorithms can be applied on the data to assess whether the risk is a good or bad risk for the ultimately contributing to the decision to underwrite or not.

Delivering Natural Language capability requires not just great technology, but more importantly, the right blend of team and a collaboration. NLP needs time and commitment from a team that understands both its concepts and the business area being transformed. For example, AXA XL and Expert System worked with a global team of six risk consultants and global technologists to deliver a solution where all engineering surveys can now be processed quickly and accurately, where previously only samples were reviewed. The speed of time to review reports means that 50 reports can now be read by the machine in five minutes or less with the further review taking on average eight hours, work that previously would take ten to twenty hours per account. It is important to stress that the continuous involvement of the risk consulting subject matter experts was essential to "teach" the system the fundamentals of risk consulting. NLP solutions are only as good as the quality of subject matter experts as the machine can only do as it is taught.

By providing underwriters with better quality and faster engineering insights the 'right' risks are underwritten contributing to the company's growth targets. By releasing risk consultants' time to be spent with clients, client experience increases potentially leading to better risk management outcomes. At the same time, the need to expand the engineering workforce has been mitigated, as has the risk of knowledge loss and the possibility that risks will be turned away or underwritten without the full insight now available.

To conclude, applying AI to the engineering process is foundational to auto or 'dynamic' pricing as well as providing a new source of accessible engineering data across the enterprise which when connected to other rich sources will provide new and invaluable insights and better decisions. Ultimately, being future focused and innovative drives a positive reputational effect with clients and the Property & Casualty in the London marketplace and globally.

InsurTech

Insurance is entering a time of exceptional technological advances and many InsurTech businesses are leveraging AI to deliver more disruption and create a larger market for innovation and new services. InsurTech innovation is taking place at a global level, but many businesses see things at a domestic level.

Digital ecosystems do not recognise national borders. Angel investor Mehrdad Piroozram has noted that those companies that do not develop an international or borderless vision will not survive or thrive (Piroozram, 2017). Several companies that have come through the Lloyd's Lab have used data and AI to enable risk prediction, improve risk management and speed up claims.

In this section we feature case studies for Zasti, Geollect and Layr. A longer list of insurance start-ups working with AI is available in Appendix B.

Box 9: ZASTI

ZASTI is an AI cloud-based technology platform built using proprietary deep-learning algorithms that aims to help businesses and their clients predict risks and improve efficiency.

This technology platform provides predictive and diagnostic solutions to business problems. ZASTI analyses existing data, identifies anomalies, recurring usage patterns and then delivers accurate predictions and diagnosis through vertically-tuned algorithms. The concept 'vertical tuning' is derived from the strategy that certain models are selected to optimise a specific task (e.g. predicting component breakdown) based on a focused data set from a specific industry or sector (vertical market). This enables the platform to perform a variety of diagnostics and predictions and also to take advantage of the transfer of learning. This differs from 'off the shelf' platforms which can only be utilised to solve a narrow set of problems. The use of heuristic modelling for automatic model selection results in lower costs, higher accuracy and faster processing. Heuristic models make use of approximations to find solutions faster than classical methods.

The AI platform processes big data from multiple data points such as policy parameters, claim data, weather parameters, crime data, IoT and sensor data from machineries. The platform can identify risk ratios and risk profiles that enable insurers to customise policies for individual customers in real time.

In the insurance sector, ZASTI has built solutions across various classes of business such as commercial real estate, aviation, healthcare and motor insurance. Applications include the following:

- ZASTI uses predictive analytics to predict breakdowns of critical components or supply chain events and to provide preventive intelligence which enables the insurers to avoid business interruptions due to flight delays, fire, supply chain events, weather or machinery breakdown. This allows insurance companies to cut down on costly claims payouts or customise new insurance policies in real time.
- In aviation, ZASTI has built a model by processing over five years of flight delay data, claims data and weather data. This model can predict flight delays for a specific flight, airline or airport for the duration of a week, month, quarter, half or one year. This information would help insurers mitigate losses and gain an insight into estimated claims for a particular duration. More importantly, this enables insurers to provide value to their client by sharing intelligence on how to 'avoid risk'.
- For motor insurance, the AI platform uses a Convolutional Neural Network (CNN) to identify damages on a car based on videos and images. Once the model is trained for a particular country, vehicle, models and features, it can identify and determine the extent of damage. Based on the image, AI can validate whether or not the image belongs to the insured car, enabling fraud detection.

Source: zasti.io

Box 10: Layr

Layr is a cloud-based platform on which small businesses can buy and manage liability insurance, with optional human contact. Layr uses an insurance distribution model that gives business owners increased autonomy over the insurance they purchase via a proprietary insurance proposal generator which can produce a quote in real time, identify the best carrier to provide cover, and secure cover in as little as 12 minutes. The business owner can tailor the amount and breadth of cover across multiple lines in just a few clicks e.g. general liability, professional liability, employers' liability, cyber, property, umbrella and more.

Layr is making use of AI to cost-effectively tap into the undersold, low-margin, small commercial market by predicting carrier pricing and suitable carriers. The Layr technology learns which carriers want to insure which profiles of business and at what price, then delivers underwriters the exact opportunities they want with bind authority. Further, as the Layr data set grows, they'll leverage AI to identify risk in real time and predict the insurance needs of their small business clients. This will allow them to automate cross-selling and upselling.

Finally, Layr is in the process of integrating with the financial services small businesses use, from online bookkeeping systems to payment processors, to automatically retrieve verified third-party data which can be used to underwrite small commercial clients, facilitating streamlined and efficient underwriting.

Source: withlayr.com

Box 11: Geollect

Geollect is a global geospatial intelligence and analytics company. It gathers insights from multiple data feeds such as satellite imagery, situational data, social media and other open-source data streams. Geospatial intelligence involves the application of dynamic algorithms to locate, quantify, describe and understand spatial data.

The data analysed may be real-time, historic, spatial or even non-spatial in nature. The data analytics goes beyond traditional statistical methods and utilises machine learning algorithms to provide new insights from the data feeds. As data sets continually grow in depth and breadth, Geollect's algorithms continually improve, providing increasingly robust and granular results. Previously unknown relationships, patterns, risks and opportunities are then visualised via an interactive dashboard. These outputs can be queried and manipulated to aid risk assessment and decision-making.

A good example of the advantages of Geollect's methodologies is in claims verification, where geospatial intelligence can provide rapid, quantified and highly-precise results from a remote location, speeding up the process and helping to reduce costs. Machine learning plays an important role by enabling the automation and validation of this process.

To date, Geollect has demonstrated the value of this approach within the insurance, and defence and security industries. For example, Geollect is currently used by the UK P&I club, providing members with intelligence on port activity, infrastructure and risks.

Source: geollect.com

Conclusions

6. Conclusions

The future

Today's AI systems are built around machine learning technologies (like deep learning) that can make automated decisions on vast quantities of data. The decisions are generalisations based upon examples pre-supplied by humans during training of the system.

The performance of these systems is not only dependent on the availability of such training examples, but is also tightly coupled to the domain within which the system operates (e.g. insurance underwriting, product recommendation).

The narrower and more concretely defined the domain, the better the AI will perform, because the coverage of training data will be greater and therefore more representative of situations the AI will encounter during deployment.

In the future, AI must overcome this fundamental dependency on human supervision during training for it to become applicable to broader domains containing situations that are less predictable.

Unsupervised training of AI systems is therefore a very active area of research likely to yield advances within the next five years. Unsupervised training still requires large volumes of data to be presented to the AI at training time; however, that data does not require human annotation. This could, for example, enable the AI system to identify anomalous or unusual behaviours that might reflect greater risk during insurance underwriting.

Unsupervised training is a good fit to scenarios where large data volumes can be collected (e.g. from the internet, or from on-board cameras on vehicles) but where annotation of such volumes is infeasible.

Another active area of research is weakly-supervised learning, in which only partially or inaccurately annotated data can be used to train the AI. Nevertheless, longer term challenges for AI remain where very little training data exists – either in the detection of rare conditions (e.g. in the medical domain), or where data annotation at

scale is prohibitively expensive or impractical for privacy reasons.

The mainstream media frequently reports advances by academics and companies such as Google DeepMind in the use of AI to play video games or classic board games such as Chess and Go. The motivation behind using such games to study emerging AI techniques is that they represent controlled, sandboxed environments in which only a limited variety of actions may be performed.

The space of possible actions that may be taken in the real world is considerably larger, e.g. for robotic navigation or operational logistics and new advances must be made to deal with the much larger range of actions and events an AI system may encounter in the real world. The generalisation of techniques such as deep Q-learning to real-world deployments is an active area likely to bear fruit in the next five years.

In the longer term we will see a shift to strong AI, with systems' intellectual capability being indistinguishable from human intelligence.

Conclusions

To conclude, AI is affecting all sectors and aspects of our world and will continue to do so in the immediate future. In insurance AI will impact all stages of the value chain, from the first enquiries, to the settlement of claims through to risk prevention.

New entrants, disruptors and established players will all have to rethink their roles and interdependencies. As businesses increasingly incorporate AI into their systems and processes they will need insurance to protect them from a range of potential risks.

So, while AI offers huge potential, insurers must be aware of the risks associated with using this new and fast developing technology and respond to new demand by developing appropriate models and products.

Appendix A – History of AI

Background

Artificial Intelligence (AI) is described as the science of creating intelligent machines capable of performing real-time tasks at the level of a human expert. Traditional review methods in the insurance sector posed several threats related to policymaking, premium rates fixing and portfolio management. The evolution of AI as an insurance technology is mitigating these risks and supporting the insurers in decision-making. The increased level of personalisation and better outcomes for the customers offered by AI are the major driving factors for the rise of AI in insurance sector.

To better understand the impact that AI is having and will continue to have on the world at large and the world of insurance, we need to look back at its 60-year history. Despite the lengthy gestation period, the application of AI is only recently starting to have a broad and pervasive effect. Previously, AI was seen as mostly a science with strong theoretical and applied branches. Apart from military implementations, applications of AI were proof-of-concept demonstrators or very narrowly-focused exemplars such a chess player, human-assisted vehicles, factory robotics and other expert systems.

Brief history of the development of AI, major trends

AI has been around for over 60 years; however, it is only relatively recently that it has come to be considered as a technology that can have a dramatic effect on so many aspects of human development.

The history of AI can be considered as covering three waves (so far), each separated by brief, but significant, funding droughts which are often described as AI winters. These phases can also be considered by their dominant technological approaches. These peaks and troughs also mirror the underlying developments in computer systems developments, particularly as we sweep through the history of computing, via mainframe computers used by many users, personal computers used by their individual owners, and now we have the cloud, a paradigm of distributed computers seamlessly made available to every user.


Although influenced by earlier work in other fields such as the brain, early AI originated in computer science and soon became a topic in its own right. Key centres emerged at the Massachusetts Institute of Technology in Boston, USA and Stanford University where John McCarthy created the Stanford Artificial Intelligence Lab (SAIL), as well as other locations across the world.

The first phase of AI development centred around symbolic approaches where problems were represented as a kind of maze to be explored, and various approaches were developed to search ever more complex mazes without creating overwhelmingly large quantities of steps, mathematically known as a combinatorial explosion. As such, these approaches of replicating high-level thinking were successful as a demonstrating solution to relatively simple tasks, but they could not scale to real-world solutions.

After the first AI winter ended, funding resumed, and a new era began in the 1970s and 80s with the growth of expert systems, also known as knowledge-based systems. These comprised collections of high-level domain knowledge compiled from experts. As work started to grow in neural networks, the second AI winter occurred as research funding dried up again.

Figure 4 shows an example of a convolutional neural network (CNN); a type of DNN used for image recognition and visual understanding. The CNN comprises multiple layers, each of which contains many thousands of neurons that learn to manipulate the image data as it moves through the network from left (input) to right (output). The network shown is the LeNet, which was among the first CNNs proposed (by Yann Le Cun, now of Facebook AI Research) to recognise hand-drawn numerals within a data set harvested from US postal mail by USPS. The principles of CNNs like LeNet have been extended in recent years such that complex, deep networks with many more layers can be used to discriminate between thousands of different kinds of object or visual concept in images. LeNet contained only two convolutional layers and could process tiny images of 32 pixels square to discriminate between 10 classes (kinds of handwritten number) with over 97% accuracy. Around a decade later, similar architectures such as Google's Inception and Microsoft's ResNet, but exhibiting hundreds of layers, were delivering super-human accuracy (beyond 80%) on the ImageNet data set comprising 1,000 classes (kind of object) with image resolutions beyond 200 pixels square.

Figure 4: Neural network applied to document recognition



Source: Image adapted from LeCun et al. 1998: Gradient-based Learning applied to Document Recognition.

Later in the 1990s, research expanded again with a move away from centralised, hierarchical systems. These were replaced by research into collaborating agents, interoperating independently, sharing messages and learning from their experience. This research approach was able to flourish thanks to a confluence of technological factors, namely the growth in storage and processing power as well as data and the dominance of the internet and web providing connectivity and distributed information.

AI therefore made the transition from symbolic approaches based around structured information to sub-symbolic approaches with unstructured content and decentralised control structures. Artificial neural networks are combined with machine learning to facilitate knowledge-gaining from the environment. (Yao, Zhou, Zhang, & Boer, 2017).

These specialist applications, which can excel in their individual tasks where the specialised data has been used to train relevant algorithms, are typically referred to as narrow AI, while the greater goal is to deliver general AI, or Artificial General Intelligence, (also known as Full AI, or full AI). Various tests have been proposed to establish this elusive goal, and chief among these tests is still the classic Turing test developed by Alan Turing in 1950 to evaluate whether a machine can imitate human cognitive capacity. (Turing, A, 1950)

AI incorporates many techniques and principles from mathematics and statistics. One such example is Bayes Rule, described as a "rigorous method for interpreting evidence in the context of previous experience or knowledge" which was discovered 200 years ago by Thomas Bayes, and, independently, by Pierre-Simon Laplace. This mathematical foundation for reasoning is now invoked in many applications of AI including image processing, machine learning and visual perception (Stone, 2013). However, there are some that question whether 18th century mathematics is sufficient to construct 21st century general AI systems needed for the future.

Such mathematical and statistical rules and approaches are incorporated into AI as algorithms; sequences of instructions that are followed to complete a task. These algorithms are then applied to data collections, typically very large, and then trained using machine learning to develop the rule or set of rules. This can be either in a supervised context where there is a known goal or target state, or unsupervised where the algorithm is applied to the data collection and the system learns its own goals or reasoning. This is one of the main concerns over AI, where an effective black box

is created with an unknown set of rules and conditions which may have inappropriate or even illegal tests created to reach a desired conclusion.

Deep learning

The rapid commercial adoption of AI in the past few years has been catalysed by technological advances in machine learning – specifically those in “deep learning” – that delivered transformational performance gains, whilst generalising to a broad gamut of applications. Deep learning refers to the use of complex (or ‘deep’) neural networks to solve machine learning problems such as classification (e.g. image or scene understanding), or regression (e.g. model fitting).

Neural networks have been studied for several decades and gained some popularity in the late 80s to early 90s for relatively well-constrained pattern recognition problems such as optical character recognition, speech and anomaly detection (e.g. of specific kinds of financial transaction). At the highest level, the neural network can be considered a data processing pipeline that ingests raw signal and outputs decisions. The processing is performed by piping the data through nodes in a graph which are considered to be a simplistic model of neurons – the cells of the biological brain. Data passes through several banks of these neurons (each is referred to as a ‘layer’) in the network; a neuron learns the optimal way to combine the data coming into it from previous neurons, to output a signal to the successive neuron. The combination of those input signals to make an output signal is little more than a weighted addition – but the complexity lies in the learning of those weights, and the fact that a neural network may have thousands or even millions of these neurons. The weights are learned by a process of ‘supervised training’ – showing the neural network paired examples of input and idealised output and repeating the process thousands of times until the right set of weights are learned to map input to output in the desired way. The algorithm that hunts for this idealised set of weights (sometimes referred to as the network parameters) is called stochastic gradient descent (SGD), and despite being invented in the 70s, it remains the core mechanism underlying even cutting-edge deep learning techniques today. Once the network has been trained using SGD (and typically a large volume of paired training data), the weights are fixed and the network may be deployed for use. New data is piped through the network which makes its decisions (‘performs inference’) accordingly.

Deep learning refers to the training of neural networks that contain several layers of neurons (i.e. they are ‘deep’), and thus contain many more neurons with many more weights to learn. By contrast with classical neural networks decades past which contained only one or two layers, contemporary DNNs circa 2018 contain over 100 (He et. al., 2016) and in the order of tens of millions of parameters. The upshot is that these networks can perform more complex tasks, but require higher volumes of training data and take more computational effort to train.

Indeed, it is for these reasons that deep learning has only become feasible within the past five years. Whilst the basic mathematics and principles of deep learning have existed for decades within academia, and indeed continued to find low-key applications in more esoteric areas of computer vision and audio processing such as speech and handwriting recognition, it was not until around 2012 that deep learning captures widespread attention. Arguably, the availability of very large data sets was a major catalyst. For example, in the field of computer vision, object recognition was a grand challenge epitomised through annual international benchmarking competitions held on increasingly challenging data sets. Machine learning had already delivered significant impact to computer vision through ‘dictionary learning’ (often now referred to as shallow learning to contrast within contemporary deep-learning approaches). The release of the ImageNet data set containing tens of millions of images across a thousand different objects represented a major new challenge that shallow approaches could achieve only single-digit percentage performance. Deep learning was shown by Alex Krizhevsky and Geoff Hinton at ECCV 2012 to deliver double-digit performance gains using a deep network of only eight layers, and in just a few years later, deep learning using tens of layers (e.g. Google’s Inception and Microsoft’s ResNet) were delivering object recognition with super-human levels of performance.

Today, DNNs take many forms. The ‘classic’ convolutional neural network or CNN shown to deliver superhuman performance at object recognition, has proven equally adept at activity recognition and other forms of recognition and classification task in diverse domains (such as audio and medical signal processing). Variants of CNNs have been used to synthesise (as opposed to classify) image content – aiding in photo and video manipulation such as object removals or appearance transfer in video (such as the notorious deep fakes). A variant of deep network called a Recurrent Neural Network (RNN) has proven highly adept at time-series regression or classification e.g. of events in a sequence of data. A specific type of RNN popular in signal processing is the Long-Short Term Memory (LSTM) network which enables the detection of more complex forms of event in the signal due to the network’s ability to explicitly carry data (or ‘remember’) between data points in the time series.

Appendix B – AI in academia and insurTech

Major research groups and their work

AI is taught and research is conducted at many universities around the world, with certain institutions standing out on the reputation of their courses and their leading academics. Below is a list of organisations carrying out AI research and development.

University of Surrey (UK)

The Surrey Space Centre's Surrey Technology for Autonomous Systems and Robotics (STAR) Lab has a long-standing R&D heritage and expertise in Robotics and Autonomous Systems (RAS) for complex space systems and mission operations. Within the Centre for Connected Plants of the Future, STAR Lab's expertise will be combined with the chemical plant design knowledge of the Department of Chemical and Process Engineering CPE to achieve automation and autonomy for future chemical plants that can perform routine as well as high-risk tasks more efficiently. Sensing and perception, mobility, human-system interaction, system engineering and autonomy are technologies of relevance.

Future of Humanity Institute (UK)

Future of Humanity Institute (FHI) is a multidisciplinary research group at the University of Oxford. Researchers at the FHI have originated or played a pioneering role in developing many of the concepts that shape current thinking about humanity's deep future. Research themes include existential risk, astronomical waste, the simulation argument, nanotechnology, prediction markets, analysis of superintelligence, brain emulations scenarios and human enhancement.

UCL Centre for Artificial Intelligence (UK)

The AI Centre carries out foundational research in AI. As we transition to a more automated society, the core aim of the centre is to create new AI technologies and advise on the use of AI in science, industry and society. The Centre brings together researchers from across computer science with a shared interest in fundamental challenges in machine vision, machine learning, machine reading and knowledge representation.

Alan Turing Institute (UK)

The Alan Turing Institute is the national institute for data science and AI. It was created as the national institute for data science in 2015. In 2017, because of a government recommendation, it added AI to its remit. Five founding universities – Cambridge, Edinburgh, Oxford, UCL and Warwick – and the UK Engineering and Physical Sciences Research Council created The Alan Turing Institute in 2015. Eight new universities – Leeds, Manchester, Newcastle, Queen Mary University of London, Birmingham, Exeter, Bristol, and Southampton – joined the Institute in 2018. Its mission is to make great leaps in data science and AI research to change the world for the better.

OpenAI

OpenAI is a non-profit AI research company, discovering and enacting the path to safe artificial general intelligence (AGI). OpenAI conduct long-term research into AGI. The organisation is sponsored by individuals such as Elon Musk, Reid Hoffman (LinkedIn) and companies such as Microsoft and Amazon.

Future of Life Institute (US)

The Future of Life Institute's mission is to catalyse and support research and initiatives for safeguarding life and developing optimistic visions of the future, including positive ways for humanity to steer its own course considering new technologies and challenges. It developed the Asilomar AI Principles, 23 principles that aim at guiding the development of AI to help and empower people in the decades and centuries ahead.

Leverhulme Trust Centre for the Future of Intelligence, CFI (UK)

Funded by the Leverhulme Trust, CFI will explore the opportunities and challenges of AI in the short term as well as long term. CFI is based at the University of Cambridge, with partners at the Oxford Martin School at the University of Oxford, at Imperial College London, and at the University of California, Berkeley. Research themes include:

- Futures and responsibilities
- Trust and society
- Kinds of intelligence
- Narratives and justice
- Philosophy and ethics of AI

Machine Intelligence Research Institute (MIRI)

The Machine Intelligence Research Institute is a research non-profit studying the mathematical underpinnings of intelligent behaviour. MIRI's mission is to develop formal tools for the clean design and analysis of general-purpose AI systems, with the intent of making such systems safer and more reliable when they are developed.

Montreal Institute for Learning Algorithms (MILA) (Canada)

Researchers from MILA have pioneered the field of deep learning and deep neural networks (both discriminative and generative) and their applications to vision, speech and language. Mila is world-renowned for many breakthroughs in developing novel deep-learning algorithms and applying them to various domains. Research themes include:

- Neural language modelling
- Neural machine translation
- Object recognition
- Neural speech recognition

The Center for Brains, Minds and Machines (CBMM) (US)

The Center for Brains, Minds and Machines (CBMM) is a multi-institutional NSF Science and Technology Center dedicated to the study of intelligence. Its research aims to explain how the brain produces intelligent behaviour and how we may be able to replicate intelligence in machines.

German Research Centre for Artificial Intelligence (DFKI) (Germany)

The German Research Center for Artificial Intelligence (DFKI) was founded in 1988 as a non-profit, public-private partnership. Based on application-oriented basic research, DFKI develops product functions, prototypes and patentable solutions in the field of information and communication technology. Research and development projects are conducted in 18 research departments and research groups, eight competence centers and eight living labs. It receives funding from government agencies like the European Union, the Federal Ministry of Education and Research (BMBF), the Federal Ministry for Economic Affairs and Energy (BMWi), the German Federal States and the German Research Foundation (DFG), as well as from cooperation with industrial partners.

Tsinghua University Institute for Artificial Intelligence (China)

China's Tsinghua University is opening a dedicated AI research centre. The Institute for AI will focus on theoretical study and interdisciplinary collaboration. The institute has close ties with industry, working with companies such as Google,

Tencent (internet-related service provider), Sogou (search engine) and Horizontal Robotics to improve core algorithms and develop new types of AI hardware (Medium, 2018).

Other Chinese research institutes share a strong connection with industry and application of AI. Commercial leaders in this space include search engine provider Baidu, Tencent and E-commerce platform Alibaba. These companies are at the forefront of AI research in areas such as autonomous driving, speech recognition and predictive healthcare (CB Insights, 2018).

InsurTech start-ups

As a result of the research conducted for this report, we have identified examples of InsurTech start-ups that have emerged over the last few years (this is not an exhaustive list, but it shows the activity in this area).

It is possible that some of these will not have survived by the time this report is published. Others may have become household names.

Either way, this list exists to present a snapshot of the ideas and concepts developed by a collection of highly-motivated disruptors that have emerged across all points in the insurance sector value chain.

- **Homelyfe**: a single app to manage all your insurance policies in one place - homelyfe.com/
- **Nimbla**: whose goal is to help SMEs manage and protect their credit via its platform, which offers them credit control and single invoice insurance using cloud accounting, plus underwriting rules agreed upon by Munich Re's new business unit Digital Partners - nimbla.com/
- **Cuvva**: mobile app in which you simply enter the registration number and approximate value of the car you are borrowing from a friend or family member, choose the time you want to be covered for, take a picture of the car and Cuvva will get you an instant quote - cuvva.com/
- **Cytora**: has developed a piece of technology it calls Risk Engine, which can be used by commercial insurers to help them target and price risk using AI algorithms - cytora.com/
- **Chisel** offers the global insurance industry solutions that apply natural language processing and AI to unstructured data sources such as insurance documents. These solutions empower insurers, reinsurers, and brokers to free trapped knowledge and automate E&O policy checking, submission prioritization, quote comparison, and submission triage. chisel.ai/insurance.php#Carriers
- **InMyBag**: insures devices like laptops, phones and cameras - inmybag.co/
- **Brolly**: a London-based start-up that uses AI technology to give customers a mobile insurance locker, advisor and shop, cutting down on costly renewals and coverage gaps - heybrolly.com/
- **Buzzmove**: a start-up that makes it easier to move house, Buzzmove is leveraging its data to build a tool for contents insurance - buzzmove.com/
- **Buzzvault**: seamlessly offering customers personalised products that meet their needs in a changing society - gobuzzvault.com/
- **Digital Fineprint**: using machine learning technology to make smart insurance policy recommendations to users who opt in to the service, based on their social media data - digitalfineprint.com/
- **Digital Risks**: an insurance specialist built for tech companies, offering a flexible, pay-monthly Insurance-as-a-Service model - digitalrisks.co.uk/
- **Neos**: a London-based startup that packages best-in-class IoT-enabled hardware, 24/7 support and unlimited building and contents home insurance, all managed from a mobile app - neos.co.uk/
- **Neura**: an AI Engine that enables users to build more robust and sustainable engagement with customers - theneura.com/
- **Instanda**: offers a management tool for insurers and brokers to build, launch, distribute and monitor new insurance products in a fraction of the time it would traditionally take. Instanda has built a tool which consolidates question sets, underwriting workflow, documentation, a rating engine and the customer journey so that underwriters can respond to changing market demands quicker than before - instanda.com/

-
- **Back Me Up:** built for the millennial consumer. For £15 you get coverage for your three most valuable items (say: laptop, camera, bike) as well as phone screen, keys and travel insurance - backmeup.co.uk/
 - **Spixii:** an "automated insurance agent" which essentially lets you buy and manage your insurance through a chat interface - spixii.com/
 - **Lapetus:** combines Life Sciences, Sensory Analytics and Dynamic Questioning to provide real-time insight into an individual's health status and longevity lapetussolutions.com
 - **Luther Systems:** Blockchain specialist start-up working on a secret bespoke product for insurance giant Aviva around simplifying contracts spixii.com/
 - **Inslly:** a cloud-based platform for insurance brokers. You can search and manage clients, policies, objects and payments in one place - insly.com/en/
 - **Rightindem:** claims process for motor insurance, giving claims brokers self-service tools for total loss claims - rightindem.com/
 - **The Carevoice:** Shanghai-based independent health platform that provides ratings and recommendations on medical providers thecarevoice.com/
 - **Datacubes:** Commercial underwriting powered by data science datacubes.com/
 - **Rubique:** financial online matchmaking platform based in India rubique.com/
 - **Swyfft:** disrupting the homeowners insurance industry by using big data and analytics swyfft.com/
 - **Worry+Peace:** allows customers to directly purchase insurance and manage all of their policies in its Pouch platform - worryandpeace.com/
 - **Zhong An:** China's first property insurance company that sells all products online along with handling claims - zhongan.com
 - **Lemonade Inc:** uses software applications acting as web robots, known as bots, to deliver insurance to consumers through its app. Consumers chat with the AI to file claims too, and the bot is authorised to pay claims instantly and without human intervention, powered by artificial intelligence and behavioural economics - lemonade.com/
 - **insurers.ai:** "Artificial Intelligence is changing the insurance business" - insurers.ai/
 - **Conversica, Inc:** "Artificial Intelligence will help you find your next customer: We create engaging conversations so that you can reach your sales goals" - conversica.com/
 - **Kasisto:** KAI is a conversational AI platform powering virtual assistants and smart bots across mobile apps, websites, messaging platforms, and IoT devices. Built with industry-specific domain expertise, KAI-powered bots and virtual assistants are well-versed in your business whether that's finance, commerce, or any other industry - kasisto.com
 - **Cape Analytics, Inc:** Offers instant, accurate and comprehensive data for policy underwriting - capeanalytics.com/
 - **Neosurance:** Just seven seconds to get insured. Zero paperwork, instant everything - neosurance.eu/
 - **Zendrive:** aims to make roads safer through data and analytics. It uses AI to predict risk, reduce collisions, save lives and money. Shift technology provides AI fraud detection and claim automation technology - zendrive.com
 - **Zesty.ai** Uses computer vision, deep learning and digital imaging to understand properties and their occupants to help the property insurance industry to better understand risks - zesty.ai/about-us/

Glossary

AE: Autoencoder	LSTM: Long Short-Term Memory (RNN)
AI: Artificial Intelligence	MDL: Minimum Description Length
ANN: Artificial Neural Network	MDP: Markov Decision Process
BFGS: Broyden–Fletcher–Goldfarb–Shanno	MNIST: Mixed National Institute of Standards and Technology Database
BNN: Biological Neural Network	MP: Max-Pooling
BM: Boltzmann Machine	MPCNN: Max-Pooling CNN
BP: Backpropagation	NE: NeuroEvolution
BRNN: Bi-directional Recurrent Neural Network	NEAT: NE of Augmenting Topologies
CAP: Credit Assignment Path	NES: Natural Evolution Strategies
CEC: Constant Error Carousel	NFQ: Neural Fitted Q-Learning
CFL: Context Free Language	NN: Neural Network
CMA-ES: Covariance Matrix Estimation	OCR: Optical Character Recognition
ES CNN: Convolutional Neural Network	PCC: Potential Causal Connection
CoSyNE: Co-Synaptic Neuro-Evolution	PDCC: Potential Direct Causal Connection
CSL: Context Sensitive Language	PM: Predictability Minimization
CTC: Connectionist Temporal Classification	POMDP: Partially Observable
DBN: Deep Belief Network	MDP RAAM: Recursive Auto-Associative Memory
DCT: Discrete Cosine Transform	RBM: Restricted Boltzmann Machine
DL: Deep Learning	ReLU: Rectified Linear Unit
DNN: Deep Neural Network	RL: Reinforcement Learning
DP: Dynamic Programming	RNN: Recurrent Neural Network
DS: Direct Policy Search	R-prop: Resilient Backpropagation
EA: Evolutionary Algorithm	SL: Supervised Learning
EM: Expectation Maximization	SLIM NN: Self-Delimiting Neural Network
ES: Evolution Strategy	SOTA: Self-Organizing Tree Algorithm
FMS: Flat Minimum Search	SRL: Statistical relational learning is a sub-discipline of artificial intelligence and machine learning that is concerned with domain models that exhibit both uncertainty (which can be dealt with using statistical methods) and complex, relational structure.
FNN: Feedforward Neural Network	SVM: Support Vector Machine
FSA: Finite State Automaton	TDNN: Time-Delay Neural Network
GMDH: Group Method of Data Handling	TIMIT: TI/SRI/MIT Acoustic-Phonetic Continuous Speech Corpus
GOFAI: Good Old-Fashioned AI	UL: Unsupervised Learning
GP: Genetic Programming	WTA: Winner-Take-All
GPU: Graphics Processing Unit	
GPU-MPCNN: GPU-Based MPCNN	
HMM: Hidden Markov Model	
HRL: Hierarchical Reinforcement Learning	
HTM: Hierarchical Temporal Memory	
HMAX: Hierarchical Model “and X”	

References

- Abbott, R. (2018). The Reasonable Computer: Disrupting the Paradigm of Tort Liability, 86 Geo. Wash. L. Rev. 1.
- Abbott, R and Bogenschneider, BN (2018) Should Robots Pay Taxes? Tax Policy in the Age of Automation Harvard Law & Policy Review, 12 (1). pp. 145-175. Retrieved from <http://epubs.surrey.ac.uk/821099/>
- Abbott, R. (2019) Everything is Obvious, 66 UCLA. L. Rev. 2.
- Ackerman, E (2011), Slight Street Sign Modifications Can Completely Fool Machine Learning Algorithms, IEE Spectrum Journal, Retrieved 6 Feb 2019 from <https://spectrum.ieee.org/cars-that-think/transportation/sensors/slight-street-sign-modifications-can-fool-machine-learning-algorithms>
- AI HLEG, (2018), Draft Ethics Guidelines for Trustworthy AI, Retrieved 6 Feb 2019 from https://ec.europa.eu/futurium/en/system/files/ged/ai_hleg_draft_ethics_guidelines_18_december.pdf
- AI powered 'Robo-Lawyer' helps step up the SFO's fight against economic crime. (2018). Retrieved June 24, 2018, from <https://www.sfo.gov.uk/2018/04/10/ai-powered-robo-lawyer-helps-step-up-the-sfos-fight-against-economic-crime/>
- Artificial Intelligence Market in Insurtech: By Type; By Application - Forecast (2017 - 2022). (2018). Retrieved April 1, 2018, from <https://www.researchandmarkets.com/research/23dqwr/artificial>
- AXA XL, (2016), XL Catlin Signs Landmark Agreement With Oxbotica, Retrieved 16 Aug 2018 from <https://axaxl.com/media/xl-catlin-signs-landmark-agreement-with-oxbotica>, Published: January 13, 2016, Author: Sinead Finlay AXA XL
- Azadian, B. J. S., & Fahy, G. M. (2018). Artificial Intelligence Artificial Intelligence and the Law : Navigating " Known Unknowns ," 35(2).
- Battista, B, and Roli, F ,(2018) "Wild patterns: Ten years after the rise of adversarial machine learning." *Pattern Recognition*84: 317-331.
- BBC, 2017. AI image recognition fooled by single pixel change. Available at: <https://www.bbc.co.uk/news/technology-41845878>
- BBC, 2018. Amazon scrapped 'sexist AI' tool, Retrieved from <https://www.bbc.co.uk/news/technology-45809919>, 10 October 2018
- BBC, 2019. Computer virus alters cancer scan images, Retrieved from <https://www.bbc.co.uk/news/technology-47812475>
- Bouchti, A. El, Chakroun, A., Abbar, H., & Okar, C. (2017). Fraud detection in banking using deep reinforcement learning. *2017 Seventh International Conference on Innovative Computing Technology (INTECH)*, (Intech), 58–63. <https://doi.org/10.1109/INTECH.2017.8102446>
- Brookings Institution, (2014), Products Liability and Driverless Cars: Issues and Guiding Principles for Legislation, John Villasenor, April 2014, Retrieved 16 August 2018 from https://www.brookings.edu/wp-content/uploads/2016/06/Products_Liability_and_Driverless_Cars.pdf

- Brown, J., Ling, T., & Ai, B. O. N. (2017). *MANAGING NEXT GENERATION ARTIFICIAL INTELLIGENCE IN BANKING A NEW PARADIGM FOR MODEL MANAGEMENT*.
- Brundage et al, (2018), The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation, Published Feb 2018, Authors: Miles Brundage, Shahar Avin, Jack Clark et al, Retrieved from <https://arxiv.org/pdf/1802.07228.pdf>
- Calo, R, (2017), Artificial Intelligence Policy: A Primer and Roadmap, Retrieved 6 Feb 2019 from https://lawreview.law.ucdavis.edu/issues/51/2/Symposium/51-2_Calo.pdf.
- CB Insights, 2018. Rise Of China's Big Tech In AI. Retrieved from <https://www.cbinsights.com/research/china-baidu-alibaba-tencent-artificial-intelligence-dominance/>
- Cheng, Y., & Zhang, W. (2017). Concise deep reinforcement learning obstacle avoidance for underactuated unmanned marine vessels. *Neurocomputing*, 272, 63–73. <https://doi.org/10.1016/j.neucom.2017.06.066>
- Conversation, (2017), How artificial intelligence conquered democracy, Published: August 8, 2017, Author: Vyacheslav Polonski Retrieved 6 Feb 2019 from <http://theconversation.com/how-artificial-intelligence-conquered-democracy-77675>
- DeepMind, (2018), A major milestone for the treatment of eye disease, DeepMind, 13 August 2018, Retrieved 16 August 2018 from <https://deepmind.com/blog/moorfields-major-milestone/>, Mustafa Suleyman
- De Fauw, (2018), Clinically applicable deep learning for diagnosis and referral in retinal disease, *Nature Medicine*, 2018, Vol 24, Issue 9, pp 342-1350, J. De Fauw, J. Ledsam, B. Romera-Paredes et al., DOI: 10.1038/s41591-018-0107-6
- Deloitte. (2015). From brawn to brains. The impact of technology on jobs in the UK, 11. Retrieved from <https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/Growth/deloitte-uk-insights-from-brawns-to-brain.pdf>
- Demchenko, I., & Belloum, A. (2017). *EDISON: Discussion Document: Part 1. Data Science Competence Framework (CF-DS) release 2*. <https://doi.org/10.5281/ZENODO.1044346>
- Ding J, (2018), Deciphering China's AI Dream The context, components, capabilities, and consequences of China's strategy to lead the world in AI, Governance of AI Program, Future of Humanity Institute, University of Oxford March 2018
- Drury, B., Valverde-Rebaza, J., Moura, M. F., & de Andrade Lopes, A. (2017). A survey of the applications of Bayesian networks in agriculture. *Engineering Applications of Artificial Intelligence*, 65(June), 29–42. <https://doi.org/10.1016/j.engappai.2017.07.003>
- e-estonia. (2017). "Artificial Intelligence is the next step for e-governance in Estonia", State adviser reveals. Retrieved June 24, 2018, from <https://e-estonia.com/artificial-intelligence-is-the-next-step-for-e-governance-state-adviser-reveals/>
- ESMA. (2017). FINANCIAL INNOVATION DAY. Retrieved June 23, 2018, from <https://www.esma.europa.eu/risk-analysis/innovation-products/financial-innovation-day>
- European Group on Ethics in Science and Technologies. (2018). Statement on Artificial Intelligence, Robotics and "Autonomous" Systems, 24. <https://doi.org/10.2777/786515>
- European Parliament, (2017), Civil Law Rules on Robotics Retrieved 16 Aug 2018 from <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+TA+P8-TA-2017-0051+0+DOC+PDF+V0//EN>, European Parliament resolution of 16 February 2017
- Executive Office of the President of the United States of America. (2016). Artificial Intelligence, Automation, and the Economy. *WhiteHouse*, (December), 55. <https://doi.org/10.1007/s00146-016-0685-0>
- Extremetech, (2018), Did Google's Duplex AI Demo Just Pass the Turing Test? Retrieved 16 Aug 2018 from <https://www.extremetech.com/computing/269030-did-google-duplex-ai-demonstration-just-pass-the-turing-test>, By Joel Hruska on May 9, 2018
- EY, (2015), The future of underwriting A transformation driven by talent and technology, Retrieved 16 Aug 2018 from [https://www.ey.com/Publication/vwLUAssets/EY-the-future-of-underwriting/\\$FILE/EY-the-future-of-underwriting.pdf](https://www.ey.com/Publication/vwLUAssets/EY-the-future-of-underwriting/$FILE/EY-the-future-of-underwriting.pdf), Published 2015, Author Gail McGiffin

- Future of Driving, Ohio University, The Future of Driving, Retrieved from <https://onlinemasters.ohio.edu/blog/the-future-of-driving/>, Infographic created by Ohio University's Online Master of Science in Civil Engineering program, date not known
- Geollect, (2018), GEOLLECT PROVIDES UK P&I CLUB MEMBERS WITH VALUABLE GEOSPATIAL DATA, Published: 29 May 2018, Retrieved 16 Aug 2018 from <https://www.geollect.com/news/archives/05-2018>
- Giovannini, S., & Scapolo, F. (2018). Ian Goldin inaugurates the JRC Megatrends Series. Retrieved June 25, 2018, from <https://blogs.ec.europa.eu/eupolicylab/ian-goldin-inaugurates-the-jrc-megatrends-series/>
- Global RegTech Summit, (2018), Global RegTech continues to grow, with over \$500M invested in Q1 2018, Retrieved 16 Aug 2018 from <https://fintech.global/globalregtechsummit/global-regtech-continues-to-grow-with-over-500m-invested-in-q1-2018/>
- Government Office of Estonia, (2018), Estonia will have an artificial intelligence strategy Retrieved 16 Aug 2018 from <https://www.riigikantselei.ee/en/news/estonia-will-have-artificial-intelligence-strategy>, Published: 27. March 2018, Republic of Estonia, Government Office
- GrexAI, (2018), Machine Learning in Finance- Present and future applications, Retrieved 6 Feb 2019 from <https://www.grex-ai.com/2018/04/09/machine-learning-in-finance-present-and-future-applications/>
- Guardian, (2018), Google's AI is being used by US military drone programme, Retrieved 16 Aug 2018 from <https://www.theguardian.com/technology/2018/mar/07/google-ai-us-department-of-defense-military-drone-project-maven-tensorflow>, Author: Samuel Gibbs, published: 7 March 2018
- Hall, W., & Pesenti, J. (2017). *Growing the Artificial Intelligence Industry in the UK*. Retrieved from https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf
- Hill, R. (2018). EU lawmakers seek coordinated hand-wringing over AI ethics. Retrieved June 23, 2018, from https://www.theregister.co.uk/2018/03/09/european_lawmakers_experts_ethics_ai/
- House of Commons. (2018). Algorithms in decision-making. *Science and Technology Committee, Fourth Report of Session 2017–19*, (May). Retrieved from <https://publications.parliament.uk/pa/cm201719/cmselect/cmsctech/351/351.pdf>
- House of Lords, (2018), AI in the UK: Ready, willing and able?, Select Committee on Artificial Intelligence, Report of Session 2017–19, Retrieved 6 Feb 2019 from <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>
- Hunt, S. (2017). *From Maps to Apps: the Power of Machine Learning and Artificial Intelligence for Regulators* (Vol. 110). Retrieved from <https://fca.org.uk/publication/documents/from-maps-to-apps.pdf>
- Industrial Strategy White Paper (2017), Retrieved 6 Feb 2019 from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/664563/industrial-strategy-white-paper-web-ready-version.pdf
- Kültür, Y., & Çağlayan, M. U. (2017). Hybrid approaches for detecting credit card fraud. *Expert Systems*, 34(2), 1–14. <https://doi.org/10.1111/exsy.12191>
- Lee, (2017), Deep Learning Is Effective for Classifying Normal versus Age-Related Macular Degeneration OCT Images, Leo H.N. Sheek, Re: Lee et al: Deep Learning Is Effective for Classifying Normal versus Age-Related Macular Degeneration OCT Images (Ophthalmol Retina. 2017;1:322–327) Ophthalmology Retina, Volume 2, Issue 2, February 2018, Pages e3
- Martinelli, L. (2017). Assessing the Case for a Universal Basic Income in the UK, (September). Retrieved from http://www.bath.ac.uk/publications/assessing-the-case-for-a-universal-basic-income-in-the-uk/attachments/basic_income_policy_brief.pdf
- Mannes, A, (2016) "Institutional Options for Robot Governance," in We Robot 2016, Miami, FL, pp. 1–40, 2016. Retrieved 6 Feb 2019 from http://robots.law.miami.edu/2016/wp-content/uploads/2015/07/Mannes_RobotGovernanceFinal.pdf

- Martinez-Arellano, G., Cant, R., & Woods, D. (2017). Creating AI Characters for Fighting Games Using Genetic Programming. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(4), 423–434. <https://doi.org/10.1109/TCIAIG.2016.2642158>
- Medium, 2018. Tsinghua University Launches Institute For AI. Retrieved from <https://medium.com/syncedreview/tsinghua-university-launches-institute-for-ai-hires-googles-jeff-dean-as-advisor-e2875fc0847f>
- McKinsey, Automotive Revolution, (2016), Automotive revolution – perspective towards 2030: How the convergence of disruptive technology-driven trends could transform the auto industry, Author: Jörg Hanebrink, Retrieved from <https://www.mckinsey.com/~media/mckinsey/industries/high%20tech/our%20insights/disruptive%20trends%20that%20will%20transform%20the%20auto%20industry/auto%202030%20report%20jan%202016.ashx>
- Mohammadfam, I., Ghasemi, F., Kalatpour, O., & Moghimbeigi, A. (2017). Constructing a Bayesian network model for improving safety behavior of employees at workplaces. *Applied Ergonomics*, 58, 35–47. <https://doi.org/10.1016/j.apergo.2016.05.006>
- National Center for Statistics and Analysis, (2015), Retrieved from <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115>, Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey, Published February, 2015, Author Santokh Singh. NOTE: “The critical reason, which is the last event in the crash causal chain, was assigned to the driver in 94 percent ($\pm 2.2\%$)† of the crashes.”
- Nedelkoska, L., & Quintini, G. (2018). *Automation, skills use and training. Development*. <https://doi.org/http://dx.doi.org/10.1787/2e2f4eea-en>
- Ocado Technology, (2016), How Ocado uses machine learning to improve customer service, Retrieved 16 August 2018 from <https://medium.com/ocadotechnology/how-ocado-uses-machine-learning-to-improve-customer-service-7d603c11e982>, Author: Alexandru Voica, October 13, 2016
- Ocado Technology,(2017) Delivering the best platform for online grocery – Ocado Annual Report, Retrieved from <http://www.ocadogroup.com/~media/Files/O/Ocado-Group/reports-and-presentations/2017/ocado-annual-report-2016.pdf>, Published 2017
- O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books,
- Of, A. T., & Industries, T. W. O. (2018). An Industry Divided, (January).
- Open AI, 2017. Attacking Machine Learning with Adversarial examples. Available at: <https://blog.openai.com/adversarial-example-research/>
- Óskarsdóttir, M., Bravo, C., Verbeke, W., Sarraute, C., Baesens, B., & Vanthienen, J. (2017). Social network analytics for churn prediction in telco: Model building, evaluation and network architecture. *Expert Systems with Applications*, 85, 204–220. <https://doi.org/10.1016/j.eswa.2017.05.028>
- Panlilio,A, Canagaretna, B, , Perkins, S du Preez, V and Lim Z, (2018), Practical Application of Machine Learning Within Actuarial Work, Published: January 2018, Retrieved 16 Aug 2018 from <https://www.actuaries.org.uk/documents/practical-application-machine-learning-within-actuarial-work>
- Papernot et al (2017), Practical Black-Box Attacks against Machine Learning, Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, Ananthram Swam, Proceedings of the 2017 ACM Asia Conference on Computer and Communications Security, Abu Dhabi, UAE, arXiv:1602.02697 [cs.CR], Submitted on 8 Feb 2016 (v1), last revised 19 Mar 2017 (this version, v4)
- Piroozram, M. (2017). The Five Insurtech Battles. Retrieved April 1, 2018, from <http://www.digitalinsuranceagenda.com/191/the-five-insurtech-battles/>
- ProPublica, (2016), Machine Bias: There's software used across the country to predict future criminals. And it's biased against blacks, by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica, May 23, 2016, Retrieved 16 Aug 2018 from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

- ReedSmith, (2017), Phishing in the Insurance Coverage Gap, Retrieved from <https://www.reedsmith.com/en/perspectives/2017/02/phishing-in-the-insurance-coverage-gap>, Authors: Alice Kyureghian Benjamin Fliegel Cristina M. Shea J. Andrew Moss, Published: 15 Feb 2017
- Regulatory sandbox - cohort 3. (2017). Retrieved June 23, 2018, from <https://www.fca.org.uk/firms/regulatory-sandbox/cohort-3>
- Reuters, 2018. Beijing to build \$2 billion AI research park: Xinhua. Available at: <https://www.reuters.com/article/us-china-artificial-intelligence/beijing-to-build-2-billion-ai-research-park-xinhua-idUSKBN1ES0B8>
- Royal Free London NHS Foundation Trust, (2017), Why doesn't Streams use AI? Retrieved 16 Aug 2018 from <https://deepmind.com/blog/streams-and-ai/> Author: Dominic King, Date: 29 November 2017
- Select, L., Intelligence, A., Veale, M., Van Kleek, M., & Binns, R. (2018). Select Committee on Artificial Intelligence AI in the UK : ready , willing and, (March), 5–180. <https://doi.org/10.1145/3173574.3174014>
- Solon, S., Labor, S. I., Ai, T., Inequality, W., Where, R., The, F., ... Ai, E. (2017). *AI Now 2017 Report*. Retrieved from https://ainowinstitute.org/AI_Now_2017_Report.pdf
- Stone, J. V. (2013). *Bayes' Rule: A Tutorial Introduction to Bayesian Analysis* (First). Sebtel Press.
- Str, H., Zhang, Y., & Schuller, W. (2017). Emotion-Augmented Machine Learning : Overview of an Emerging Domain, 305–312.
- Technology Review, (2016), Oxbotica's New Autonomous Vehicle Software Learns As It Goes by Jamie Condliffe July 15, 2016, Retrieved 16 Aug 2018 from <https://www.technologyreview.com/s/601910/oxboticas-new-autonomous-vehicle-software-learns-as-it-goes/>
- The White House, 2019. Accelerating America's Leadership in Artificial Intelligence. Available at: <https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/>
- Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind Association, Oxford University Press*, 59(236), 433–460. Retrieved from <http://www.jstor.org/stable/2251299>
- UNODA, (2017), Pathways to Banning Fully Autonomous Weapons, Retrieved 16 Aug 2018 from <https://www.un.org/disarmament/update/pathways-to-banning-fully-autonomous-weapons/>, Published: October 23rd, 2017, Text and photos by Gillian Linden
- Whitehouse. (2018). Artificial Intelligence for the American People. Retrieved June 24, 2018, from <https://www.whitehouse.gov/briefings-statements/artificial-intelligence-american-people/>
- Will AI replace humans in the insurance industry? (2018). Retrieved April 1, 2018, from <https://www.bankingtech.com/2018/03/will-ai-replace-humans-in-the-insurance-industry/>
- Woebot, (2018), Retrieved 6 Feb 2019 from <https://www.wired.com/2017/06/facebook-messenger-woebot-chatbot-therapist/>
- Yang, J., Chen, Y., Huang, W., & Li, Y. (2017). Survey on artificial intelligence for additive manufacturing. *ICAC 2017 - 2017 23rd IEEE International Conference on Automation and Computing: Addressing Global Challenges through Automation and Computing*, (September), 7–8. <https://doi.org/10.23919/ICAC.2017.8082053>
- Yao, X., Zhou, J., Zhang, J., & Boer, C. R. (2017). From Intelligent Manufacturing to Smart Manufacturing for Industry 4.0 Driven by Next Generation Artificial Intelligence and Further On. *2017 5th International Conference on Enterprise Systems (ES)*, 311–318. <https://doi.org/10.1109/ES.2017.58>
- Zagorin, E. (2018). Artificial Intelligence in Insurance – Three Trends That Matter. Retrieved April 1, 2018, from <https://www.techemergence.com/artificial-intelligence-in-insurance-trends/>

