Institute of Risk Management

An IRM Occasional Paper

Risk, Science and Decision Making: How should risk leaders of the future work with AI?

Report of a CRO Round-table Discussion

Authors:
Trevor Maynard and Dougal Goodman

EgonZehnder

LLOYD'S

irm
Developing risk professionals

Contents

Foreword

Developments in technology continue to disrupt industries and organisations globally, presenting both opportunities and risks. Preparing the risk community for this challenge through education, training, research and networking events is one of the key strategic objectives of the Institute of Risk Management (IRM). We were therefore delighted to have the opportunity to bring together a group of Chief Risk Officers, leading academics and risk thinkers for a round-table discussion on the topic of artificial intelligence and its potential impact on the risk landscape. Their discussions are summarised in this document.

We would like to thank those who made this event possible, in particular our knowledgeable and challenging speaker, Professor Sir Adrian Smith, our facilitator and chair Dougal Goodman, our rapporteur Trevor Maynard and our generous hosts Egon Zehnder.

Iain Wright, CFIRM, IRM Chair

In this world of increasing complexity we are continually inundated with data in our professional and personal lives. And data is driving the AI revolution. This can be overwhelming for many as we search for ways of cutting through the noise for the key insights. It is also driving the tendency in all of us to delve further into more esoteric areas of specialisation and deepening our search for domain experts including those in the risk function. What does this mean for future risk leaders and how are they evolving? We were delighted to host the roundtable event with our friends at the IRM to explore these themes further. What we found was a rich debate covering a breadth of topics that are summarised herein.

On leadership specifically, there is a heightened need for leaders to identify, hire, develop and deploy the right talent in the right roles aligned towards a common purpose. Through research at Egon Zehnder we have identified the four elements that we believe are key to identifying leadership potential and are present in leaders across an organisation. This framework, developed using the same methodology as trait-based measures of personality, can help spot high potential talent early allowing organisations to invest time in developing their skills, abilities, and competencies. Specifically, we believe that by identifying and nurturing behaviours associated with curiosity (about the outside world and the leader themselves), insight (cutting through the noise), engagement (building followership through winning hearts and minds), and determination (personal resilience), we can help people progress and perform in leadership positions that resonate with their motivation, style, and natural energies.

The role of the leader is then less about knowing all the answers but more about defining purpose, shaping culture, and trusting their teams to deliver. This requires a healthy dose of humility and a measure of comfort in showing vulnerability. The more data that surrounds us the greater the need for leaders to exercise judgement, striking that right balance between art and science.

Our sincere thanks to the IRM and our wonderful guests and speakers for a truly memorable roundtable event.

Rian Raghavjee, UK Risk Practice Leader, Egon Zehnder

About Us

The IRM is the leading professional body for Enterprise Risk Management. For over 30 years we have provided internationally recognised qualifications and training, published research and guidance and set professional standards. We are a not-for-profit body, with members working in all industries, in all risk disciplines and in all sectors around the world.

www.theirm.org

irm

Developing risk professionals

Since 1964, Egon Zehnder has been at the forefront of defining great leadership in the face of changing economic conditions, emerging opportunities and evolving business goals. With more than 450 consultants in 68 offices and 41 countries, we work closely with public and private corporations, family-owned enterprises and non-profit and government agencies to provide board advisory services, CEO and leadership succession planning, executive search and assessment, and leadership development.

www.egonzehnder.com

EgonZehnder

For more information contact:

Obinna A Onyeagoro, Egon Zehnder

Email Obinna.Onyeagoro@egonzehnder.com

Tel +44 207 943 5913

Rian Raghavjee, Egon Zehnder

Email Rian.Raghavjee@egonzehnder.com

Tel +44 207 943 1950

Keynote Speaker

Sir Adrian Smith FRS, Institute Director and Chief Executive, The Alan Turing Institute

Professor Sir Adrian Smith joined The Alan Turing Institute in September 2018 as Institute Director and Chief Executive. His previous role was Vice-Chancellor of the University of London where he was in post from 2012. He was Director General, Knowledge and Innovation in BIS (now BEIS) between 2008 and 2012. Sir Adrian has worked with the UK Higher Education Funding and Research Councils and was appointed Deputy Chair of the UK Statistics Authority from September 2012. In 2014, he was appointed Chair of the Board of the Diamond Synchrotron at Harwell and in 2018, a board member of the UK Atomic Energy Authority.

Sir Adrian is a past President of the Royal Statistical Society and was elected a Fellow of the Royal Society in 2001 in recognition of his contribution to statistics.

In 2019 the Secretaries of State for Business, Energy and Industrial Strategy and for Education commissioned Sir Adrian to provide independent advice on the design of UK funding schemes for international collaboration, innovation and curiosity-driven blue-skies research. In 2003-04 Sir Adrian undertook an inquiry into Post-14 Mathematics Education for the UK Secretary of State for Education and Skills and in 2017, on behalf of Her Majesty's Treasury and the Department for Education, published a 16-18 Maths Review. In 2006 he completed a report for the UK Home Secretary on the issue of public trust in Crime Statistics.

He received a knighthood in the 2011 New Year Honours list.

Discussion participants

Kanwardeep Singh Ahluwalia, Head of EMEA Markets Risk & Deputy Chief Risk Officer for EMEA, Bank of America Merrill Lynch

Fabrice Brossart, CRO General Insurance International and UK Senior Managing Director, Enterprise Risk Management, AIG

Dr Paul Burgess, Senior Adviser on Risk, Cambridge Centre for Risk Studies

Tony Chidwick, CMIRM, Managing Director, Head of Operational Risk, Barclays Services

Socrates Coudounaris, CFIRM, Executive Director, Risk, Reinsurance Group of America, outgoing IRM Chair

Roger Dix, Chief Risk Officer, Wesleyan

Dr Dougal Goodman, Vice President, Foundation for Science and Technology

Alex Hindson, CFIRM, Chief Risk Officer, Argo Group

Desmond McNamara, Chief Risk Officer, Distribution Finance Capital

Dr Trevor Maynard, Head of Innovation, Lloyd's

Alastair Nunn, Head of Risk, Barclays Ventures

Lewis O'Donald, Former Senior Managing Director and Chief Risk Officer, Nomura Holdings

Obinna Onyeagoro, Consultant, Egon Zehnder

Rian Raghavjee, Consultant, Egon Zehnder

Lakshmi Shyam-Sunder, Chief Risk Officer, World Bank

Professor Sir Adrian Smith, Director, The Alan Turing Institute

Carolyn Williams, CMIRM, Director, Corporate Relations, Institute of Risk Management

Iain Wright, CFIRM, Chief Risk Officer, Europe, Great-West Life LifeCo (Canada Life), incoming IRM Chair

Introduction

Artificial Intelligence (AI) is vexing boards around the world who are asking whether it presents an opportunity or a threat. The concept of AI began in the 1940s, pioneered by Alan Turing, amongst others, who famously proposed the "Turing test" to assess whether a computer was "thinking". The argument (roughly) being that a computer is considered to be thinking if a human cannot tell whether it is talking to the computer or to another human. It was fitting, then, that Professor Sir Adrian Smith, the Institute Director and Chief Executive of The Alan Turing Institute, opened up a discussion, convened by the Institute of Risk Management and hosted by Egon Zehnder, to explore how should risk leaders of the future work with AI? This article describes the topics discussed.

The UK, along with many other countries, has concluded that AI is a critical new capability in our future and has therefore created The Alan Turing Institute to focus on the discipline. Some 400 researchers across thirteen UK universities are brought together in the Institute. Disciplines include mathematics, statistics, computer science, software development and social science. This diversity of skill set underscores the first key point: AI is multidisciplinary; it is difficult for any one person to understand all the issues and intricacies. So, risk professionals are going to need to rely on a range of expert advisors when working with AI, but will also need to develop some skills themselves to be able to interpret what they are being told.

> "…risk professionals are going to need to rely on a range of expert advisors when working with AI, but will also need to develop some skills themselves."

What is AI?

Sir Adrian gave an analogy that some approaches in AI are similar to fitting a line through data. For line fitting there is a known formula which explicitly finds the two parameters, the slope of the line and its position. Rather than algebraically solving the problem, another approach, he noted, would be try lots of lines and stop when the 'fit' looks ok.

This is basically how deep neural networks work, except the number of parameters can run into millions. 'Fit' is defined by an error function, effectively measuring how close the model gets to the past data; and 'try lots' is shorthand for Stochastic Gradient Descent - which gradually adjusts the parameters to improve the fit a few data points at a time. This approach has become practical recently due to increases in computing power, advancements in algorithms, large amounts of data and cloud storage.

AI, Machine Learning and Deep Learning

AI encompasses several other techniques however, many of which have been tried with varying success over the past 70 years. A useful categorisation was published in an MMC ventures paper "The state of AI: Divergence" which used AI as the generic term which then includes as a subset "machine learning", which in turn includes "deep learning". AI methods include expert systems, where human rules are specifically coded into the computer. These work in some circumstances, but it has, so far, proved challenging to use them for highly complex tasks such as "have a natural conversation with someone" or "make a cup of tea in a stranger's house".

Machine learning avoids the need to codify by hand, by extracting the rules from the data in some way. This method can either be "supervised" where the data is labelled already, for example "this image is a cat"; or "unsupervised", where the algorithms try to find patterns from the data directly with no human intervention. Machine learning methods include: Support Vector Machines (SVMs) which try to determine clusters in the data; decision trees, which repeatedly subdivide the data into categories that explain the data well and Random Forests, which try lots of decision trees automatically to find a really good one.

Deep neural networks are a special type of machine learning originally inspired by a loose analogy to the how brains function. They take data as input and flow it through lots of "neurons" which "fire" if the data is exciting to them. 'Excitement' is defined by combining parameters and the data using a fixed calculation which, if a threshold is exceeded, passes the answer forward to the next layer of calculation. In modern deep networks there are lots of layers and this is why they are called "deep". It is important to realise that the only calculations involved are multiplying, adding and taking maxima. As such, once a neural network has been "trained" (to find some good parameters) we know precisely how it will work in principle. The training process uses some hard maths but essentially this is the step Sir Adrian described as "try lots of lines". It can take some time to work through the data and this is why such large grids of computers are needed.

What can AI do?

The basic point of AI is that it can work through vast data sets with a specific task in mind. Humans could in principle do this too, but in practice it would take too long or be too costly, or too boring. This is not a new phenomenon. Spreadsheets completely revolutionised modern businesses in the 1990s and they changed the workplace significantly, allowing more complex tasks to be done quicker and raising expectations of what is possible. When it comes to AI, it is important for businesses to be aware of the art of the possible. For example, a neural network could review every recorded call ever made to their company to look out for trends in complaints, or questions. Chat bots are already used ubiquitously to handle the easy requests and forward harder issues to their human counterparts. Algorithms can find patterns in data to enable better understanding of customers (see case study). Documents can be read automatically and compared, flagging key changes to human reviewers. Images can be reviewed, looking for actionable signs such as finding cancer in a scan or determining whether a car has been damaged in an accident by comparing before-and-after photos. Complex searches can be carried out to assist a human case handler find the key information, regulations or laws to answer a business question. AI does not have to provide a final answer, some of the best uses involve working hand in hand (metaphorically) with a human user.

> “When it comes to AI it is important for businesses to be aware of the art of the possible.”

There are many megatrends affecting societies today, technology is one, but others include climate change, ageing populations, disenfranchised youth, population growth, stagnating economic growth, ageing infrastructure and low financial inclusion. AI offers solutions that can help with some of these and, as one attendee mentioned, “society needs this to succeed”. For example, in developing countries AI can help enhance financial inclusion by providing bespoke advice and financial service pricing at a fraction of the cost of human broking. In poorer communities this may be the only way that such products can be made available.

What impact will AI have?

Historically, industrial revolutions have been disruptive but more new jobs have been created than destroyed. The second industrial revolution, from 1850 onwards, saw electrification, global communications, widespread use of artificial fertiliser and increases in mobility. Due to these, the working class was redefined and the professional middle class was born. Opinions at the round-table differed on whether the technological revolution, including AI will be different. Some felt that change could be so fast that some segments of the work force may not be able to adapt fast enough to avoid becoming left behind. In this case the welfare state will need to evolve, and companies will have to think carefully about their moral duties beyond a traditional profit motive.

This led onto a discussion of culture. “Culture eats strategy for breakfast”, said Peter Drucker and it was generally agreed that this is correct. Even highly skilled professions, such as radiologists, may be at risk of redundancy. The workforce realises this, and low morale could lead to blocking behaviours and poor uptake or adoption of new technologies. These concerns are valid and must be taken into account in any change program. Senior leaders cannot expect the workforce to be altruistic and must find a narrative that is inspiring and inclusive. The Chief Risk Officer (CRO) should consider poor staff morale to be a key potential risk within their firm.

Human v AI - a dangerously opaque black box?

As noted, the calculation steps are fully transparent, there are no steps hidden from view in principle, although whether they are widely or easily understood is another question. But we have to compare this with real human decision-making. We may be able to tell a story to explain how a decision was reached, but MRI imaging has clearly shown that often decisions are not made by the logical reasoning part of our brains. Instead the emotional brain centre (the Amygdala, or “lizard brain”, so named because it was first to evolve) is often involved, especially when we are tired or encountering new situations. Human decision-making is also subject to a range of further biases.

So, we need to be careful when using the term ‘black box’, as it can arguably apply to us more than to most AI systems. There are certain problematic issues with these algorithms, however, such as not being sure which input data will have the largest effect. Similar to complex ecosystems in nature, strange and different results can appear from seemingly similar data. This may be more an issue of us wishing to be in control than the AI decisions being wrong. For example, it wasn’t long ago that DNA data was not accepted in court as proof of identity, but attitudes have changed, and this scientific breakthrough is now accepted as evidence. AI decisions may be similar as society becomes familiar with them. Such issues can also arise because the training data was biased to start with however, and this is a concern of regulators especially in relation to anti-discrimination laws.

AI Risks in Context

If society needs AI to succeed, then what could cause it to fail? Some felt that the media poses a threat to uptake based on the reception that Genetically Modified Organisms (GMO) has had in Europe. Labelled “Frankenstein food” this led to a public backlash. By analogy, “Algorithm kills child” is the headline none of us wants to read. Yet, “driver kills child” is sadly an altogether too frequent event, so frequent that the media often doesn’t even cover the story. A key role business and CROs can play is to highlight the relative risk of various choices. Nothing, as our profession knows well, is risk free. Trust is built up over many years but can be lost in an instant, so a strong sense of ethics and careful development is required, from academia and business, to avoid this outcome.

> “Algorithm kills child” is the headline none of us wants to read. Yet, “driver kills child” is sadly an altogether too frequent event.

Developing AI Capabilities

Many new systems have been developed to provide access to AI methods to a lay audience. These low-code or no-code options are becoming standardised. Nevertheless, at present, and probably for some time to come, we will need new talent that is comfortable using the new techniques. By analogy, basic spreadsheets may be ubiquitous, but few employees are comfortable with look-up tables, visual basic or other deeper parts of their toolbox. The UK government and its research councils, advised well by the academic community, foresaw this and facilitated the creation of multiple masters and doctoral training centres. So there will soon be a steady flow of expertise coming into the work force. Several companies round the table were encouraging their staff to retrain by taking masters courses. Professionals such as actuaries are also considering adapting their curricula. The IRM has introduced a new Digital Risk Management Certificate providing risk professionals with an opportunity to update their general technological knowledge.

Data is the life blood of AI, without it the algorithms are empty - like a plumbing system of pipes and valves with no water flowing through. Yet companies are nervous to share data citing GDPR restrictions or IP concerns. It is certainly true that data should be considered a strategic asset, but, unless it is explored, the value it contains will not be exploited. One way to avoid privacy and IP concerns when building or experimenting with AI systems is to create a synthetic data set. Such data has the same form as real data but is actually artificial. So model development can be outsourced using synthetic data and then brought in house when it is ready to be run on the real data, behind closed doors. The case study illustrates how a synthetic data set of insurance claims was derived. A new artisan, the "data wrangler" has evolved: this is someone who can wrestle with raw data and turn it into a form that computers can interrogate. To extract the most value from data in future it will be structured, where possible, and systematically stored along with other contextual information. More thought should be put into systems design, not just to the immediate need, but also the future potential of any data that is collected. By thinking and planning now the CRO can help ensure that the risk of future data complexity is minimised. Public attitudes to data tend to vary: if they see how they can benefit from sharing data they are more likely to consent willingly to sharing.

> "By thinking and planning now the CRO can help ensure that the risk of future data complexity is minimised."

The role of regulation

Appropriate and trusted regulation will be key to ensuring positive AI uptake. Regulators have no incentive to see things go wrong on their watch. As such, they will take their time to understand and get fully comfortable with new initiatives such as the quickly evolving AI space. In the context of UK financial services the Financial Conduct Authority (FCA) itself uses AI to risk-assess firms from complex data sets, as well as linguistic algorithms to evaluate advertising. And it has its sandbox, which allows firms to test emerging technology solutions with real customers. In a speech this summer, James Proudman, ED of UK Deposit Takers Supervision at the Prudential Regulation Authority (PRA), said that technological innovation is inevitable and welcome, and reiterated that it is not the role of the regulator to stand in the way of progress.

> "CROs should ensure their company has thought carefully through the issue of liability"

There is also an increased focus from PRA (with its remit and its history via the Bank of England) to look at systemically important infrastructure providers. And we could see that some parts of the AI chain are caught by this. But that concentration of risk is something we should all be concerned about, not just regulators. Indeed, when an algorithm does 'fail' this will trigger a landmark legal discussion on who is to blame, is it the data collector? The person that wrangled the data? The algorithm designer? The supplier of a service or product? Time will tell. CROs should ensure their company has thought carefully through the issue of liability.

AI and the CRO

So what does AI mean for the CRO of the future? Reassuringly, it doesn't look like the role will be replaced by a black box, but risk leaders will need to have multidisciplinary training and experiences rather than deep expertise in a single silo. We can expect CROs to stay involved in deciding risk criteria involving AI processes, in which case they will need access to trusted advisors who can translate the concepts. It will also be important for the risk function and its leader to have some knowledge of these concepts, and the risks arising, where data science, statistics and computing meet, so training will be required. Reverse mentoring may be useful in this space as "digital natives" (those who have grown up with digital technology) may have much to teach the senior leadership.

> "In the innovation context 'failure' means 'not learning', it does not mean 'not trying'."

Finally, companies will need to innovate to keep up with economic pressures and this is certain to include AI in future. The skill set to oversee innovation is different to that of a business as usual function. Innovation projects should be seen as experiments with a clear hypothesis that gets explored and tested. CROs should be reassured if they see a series of experiments with null results followed by a few successful trials that scale up. This will require a changing attitude and redefinition of "failure". In the innovation context 'failure' means 'not learning', it does not mean 'not trying'.

An intelligent summary of Artificial Intelligence

In summary, AI is a collection of techniques developed over 70 years that have recently seen a leap forward due to computing power, new algorithms and vast data sets coupled with cheap storage. AI has achieved better-than-human performance in some fields, like image recognition, but is still lagging far behind in others, such as natural language processing. Society can benefit from AI if implemented well, taking human welfare into account. New talent is required and has been anticipated by academia and government; it will need careful assimilation within the existing workforce as morale is a key risk CROs should monitor. The media has the power to adversely affect AI and businesses should help focus on the positives to resist the negative. Data is critical and should be thought of as a strategic asset. Synthetic datasets that safely mimic true data can help with regulatory and IP concerns. New interpretations of liability will arise as AI matures and this should be actively monitored by risk functions. Finally, regulators will be key to the pace of change within AI and firms should engage collaboratively to ensure concerns are addressed, to ensure maximum benefits are achieved. Risk leaders of the future will need to work with AI and will find their roles are enhanced positively if they stimulate innovation and embrace its potential.

Case study:
Neural networks are better at predicting the chance of disaster

Most insurance policyholders do not suffer a financial loss. For the few that do, insurance pays out many multiples of the cost of the insurance. The premiums of the many pay for the claims of the few. This scarcity of disaster, whilst a good thing for society, makes data analysis challenging for insurers. They must estimate the probability of a "claim" (the request for an insurance payout from the policyholder) from the limited number of claims arising in their portfolio. To test whether AI could help in this challenge Lloyd's of London teamed up with Zasti, a specialist AI company who took part in the incubator "Lloyd's Lab", to carry out a controlled experiment.

A synthetic data set was created, blending US census data with anonymised data from several Lloyd's insurance portfolios. Each row of data represented a single building whose chance of disaster depended on their characteristics such as building materials, location and age of property. Claims were simulated using standard statistical distributions whose key parameters were defined by the data - the amount of risk varied for each row of data. The experiment then compared a number of techniques used to determine the level of risk just from the simulated claims arising. This reflects a real world situation where the number of claims is observed but the true level of risk is not known - in fact this is one of the key challenges for insurers, to guess the underlying risk.

In practical situations the data is often scarce and poor, there can be errors and often several key fields are missing. For example, insurers often don't know the full latitude and longitude, or detailed building types. To reflect this, the experiment was repeated with a smaller subset of the simulated data which was also degraded to have errors and missing values.

The key finding of the work was that the Neural networks also performed much better than the other techniques in a data poor setting. In fact, the quality of risk assessment from the neural network remained very high whereas the other methods quality fell very rapidly as the data quality and size reduced. The project concluded that AI has much to offer in the field of data analysis for insurers.

For full details of the project see this preprint:

What role for AI in insurance pricing?

Further reading: some useful links and references

The Alan Turing Institute
www.turing.ac.uk

UKRI Gateway to Research
https://gtr.ukri.org

The Foundation for Science and Technology
www.foundation.org.uk

Lloyd's of London
https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/taking-control

How smarter technologies are transforming the insurance industry, Accenture
www.accenture.com/t20181012t053502z__w__/us-en/_acnmedia/pdf-76/accenture-insurance-artifical-intelligence-pov.pdf

Artificial Intelligence – What it is and why it matters | SAS UK
www.sas.com/en_gb/insights/analytics/what-is-artificial-intelligence.html

The State of AI 2019, MMC Ventures
www.mmcventures.com/wp-content/uploads/2019/02/The-State-of-AI-2019-Divergence.pdf

Royal Society and British Academy reports on AI
www.royalsociety.org/topics-policy/data-and-ai/artificial-intelligence/

This AI life: ensuring our AI future works for us, Royal Academy of Engineering
https://www.raeng.org.uk/events/events-programme/2018/october/this-ai-life-ensuring-our-ai-future-works-for-us

The Alan Turing Institute blogpost – the Future of Work will soon be the present
https://www.turing.ac.uk/blog/future-work-will-soon-be-present

Growing the artificial intelligence industry in the UK, Professor Dame Wendy Hall and Jérôme Pesenti
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf

HM Government Industrial Strategy - Artificial Intelligence Sector Deal
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/702810/180425_BEIS_AI_Sector_Deal__4_.pdf

Turning Potential into Success: The Missing Link in Leadership Development, Harvard Business Review, 2017
https://hbr.org/2017/11/turning-potential-into-success-the-missing-link-in-leadership-development

Institute of Risk Management
2nd Floor, Sackville House
143–149 Fenchurch Street
London
EC3M 6BN

www.theirm.org

irm

Developing risk professionals