

Machine learning

More science than fiction

About ACCA

ACCA (the Association of Chartered Certified Accountants) is the global body for professional accountants, offering business-relevant, first-choice qualifications to people of application, ability and ambition around the world who seek a rewarding career in accountancy, finance and management.

ACCA supports its **208,000** members and **503,000** students in **179** countries, helping them to develop successful careers in accounting and business, with the skills required by employers. ACCA works through a network of **104** offices and centres and more than **7,300** Approved Employers worldwide, who provide high standards of employee learning and development. Through its public interest remit, ACCA promotes appropriate regulation of accounting and conducts relevant research to ensure accountancy continues to grow in reputation and influence.

ACCA is currently introducing major innovations to its flagship qualification to ensure its members and future members continue to be the most valued, up to date and sought-after accountancy professionals globally.

Founded in 1904, ACCA has consistently held unique core values: opportunity, diversity, innovation, integrity and accountability.

More information is here: www.accaglobal.com



Machine learning

More science than fiction

About this report

This report is an introduction to machine learning, with particular emphasis on the needs of the accountancy profession. In addition to an overview of what it is, the findings inform perspectives on how it can be applied, ethical considerations and implications for future skills.



FOR FURTHER INFORMATION:

Narayanan Vaidyanathan
Head of Business Insights, ACCA

Foreword



The impact of digital on the accountancy profession is an important, current thematic focus for ACCA that permeates everything we think about and do. It is a focus on ourselves as an organisation, as much as on our thought-leadership for wider best practice.

As an organisation, ACCA incorporates digital applications in both the content and delivery of its training programmes. Our course content emphasises the need for professional accountants to develop an appreciation of a range of technology topics, from analytics to artificial intelligence. The ACCA qualification and continuing professional development (CPD) offerings are committed to a digital approach: online and flexible, designed to give the best service to our members and students in over 180 countries.

Our thought leadership work builds on this organisational focus on digital applications. The perspectives on machine learning offered in this report are the latest addition to a strong portfolio of research covering technologies from robotic process automation to blockchain.

The report offers an accessible, practical introduction to the basics of machine learning, and how it is being adopted within the accountancy profession. It also explores issues of ethics and other concerns pertinent to the public interest. These concerns are integral to ACCA's mission, and our dialogue with regulators, standard setters, partners, members and students.

Our aim is to provide a considered and thoughtful voice, in an often over-hyped debate about the danger that artificial intelligence will take over the world. We are hopeful that this report will be a useful resource for our stakeholders and play its part in supporting a meaningful and constructive debate.

Alan Hatfield

Executive Director, Strategy and Development

Contents

Executive summary	6
Introduction	8
1. Machine learning and accountancy	10
2. Navigating the terminology	12
3. Applications of machine learning	18
4. Ethical considerations	25
5. Skills in a machine learning environment	35
Conclusion	37
Appendix 1	38
Appendix 2 – Country snapshots	39
UK	40
China	41
Malaysia	42
Singapore	43
United Arab Emirates (UAE)	44
Ireland	45
Pakistan	46
Appendix 3	47

DISCLAIMER

Parts of this report make reference to machine learning products or other initiatives from third parties. This is done for information purposes in response to requests for real-world examples. The report does not constitute an endorsement of the particular products or initiatives mentioned or a complete list thereof.

Executive summary

Artificial intelligence (AI) is having a big impact on public consciousness. And machine learning (ML), which uses mathematical algorithms to crunch large data sets, is being increasingly explored for business applications in AI-led decision making.

This follows several years in the wilderness, where the prevailing belief was that AI was the stuff of movie fantasy. Now, with access to far more data and far more processing power than ever before, ML seems set to challenge that view.

This is an area with plenty of terminology and a minefield of differing interpretations as to what they mean. ACCA's survey of members and affiliates reflected this challenge when asked about their understanding of terms such as AI, ML, natural language processing (NLP), data analytics and robotic process automation (RPA).

On average for any given term: 62% of respondents had not heard of it, or had heard the term but didn't know what it was or had only a basic understanding, 13% of respondents had a high or expert level of understanding. This suggests a lot of potential for greater education and awareness building among the accountancy community around the world.

One way to describe AI is the ability of machines to exhibit human-like capabilities in areas related to thinking, understanding, reasoning, learning or perception. ML is a sub-set of AI that is generally understood as the ability of the system to make predictions or decisions based on the analysis of a large historical dataset.

Essentially, ML involves the machine, over time, being able to learn the characteristics of data sets and identify the characteristics of individual data points. In doing so, it 'learns' in the sense that the outcomes are not explicitly programmed in advance. They are arrived at by the ML algorithm as it is exposed to more data and determines correlations therein.

As with any technology, with power comes responsibility. And in the case of machine learning, ethical considerations are never far away.



The report begins with an introduction to the basics. This is because it is important to have some appreciation of what these applications are doing, to be able to trust such systems and to understand how machine learning can be a step towards developing a greater level of machine intelligence.

In this context, 'intelligence' refers to the ability of the technology, in certain circumstances, to make decisions or draw inferences, without there being an instruction to treat a given dataset in a fixed, predetermined way. But it does not mean that the technology has suddenly developed an independent consciousness – this is not about robots going on the rampage!

The market is recognising the power of ML with 2 in 5 respondents stating that their organisations are engaged with this technology in some way. This includes those who stated that their organisations are in full production mode dealing with live data (6%), advanced testing with 'go-live' within 3-6 months (3%), early stage preparation with go-live within 12 months (8%) and in initial discussions exploring concepts/ideas (24%).

Applications for adoption range across diverse areas, including for example, invoice coding, fraud detection, corporate reporting, taxation and working capital management. The report explores various products and initiatives across these areas.

These findings reinforce the need for the accountancy profession to prioritise building awareness and understanding in this area, as organisations will increasingly need these skills. In fact the biggest barrier to adoption cited in the survey was the lack of skilled staff to lead the adoption (52%).

As with any technology, with power comes responsibility. And in the case of ML, ethical questions are never far away. Professional accountants need to consider, and appropriately manage, potential ethical compromises that may result from decision making by an algorithm.

Who has accountability in this situation? What is the risk of bias, given that ML algorithms will inevitably reflect any bias in the data sets that feed them? About 8 in 10 respondents were of the view that organisations have a

responsibility for some form of disclosure to highlight when a decision has been made by a ML algorithm.

The report considers a range of ethical considerations relevant to professional accountants, using for guidance, the fundamental principles established by the International Ethics Standards Board for Accountants (IESBA).

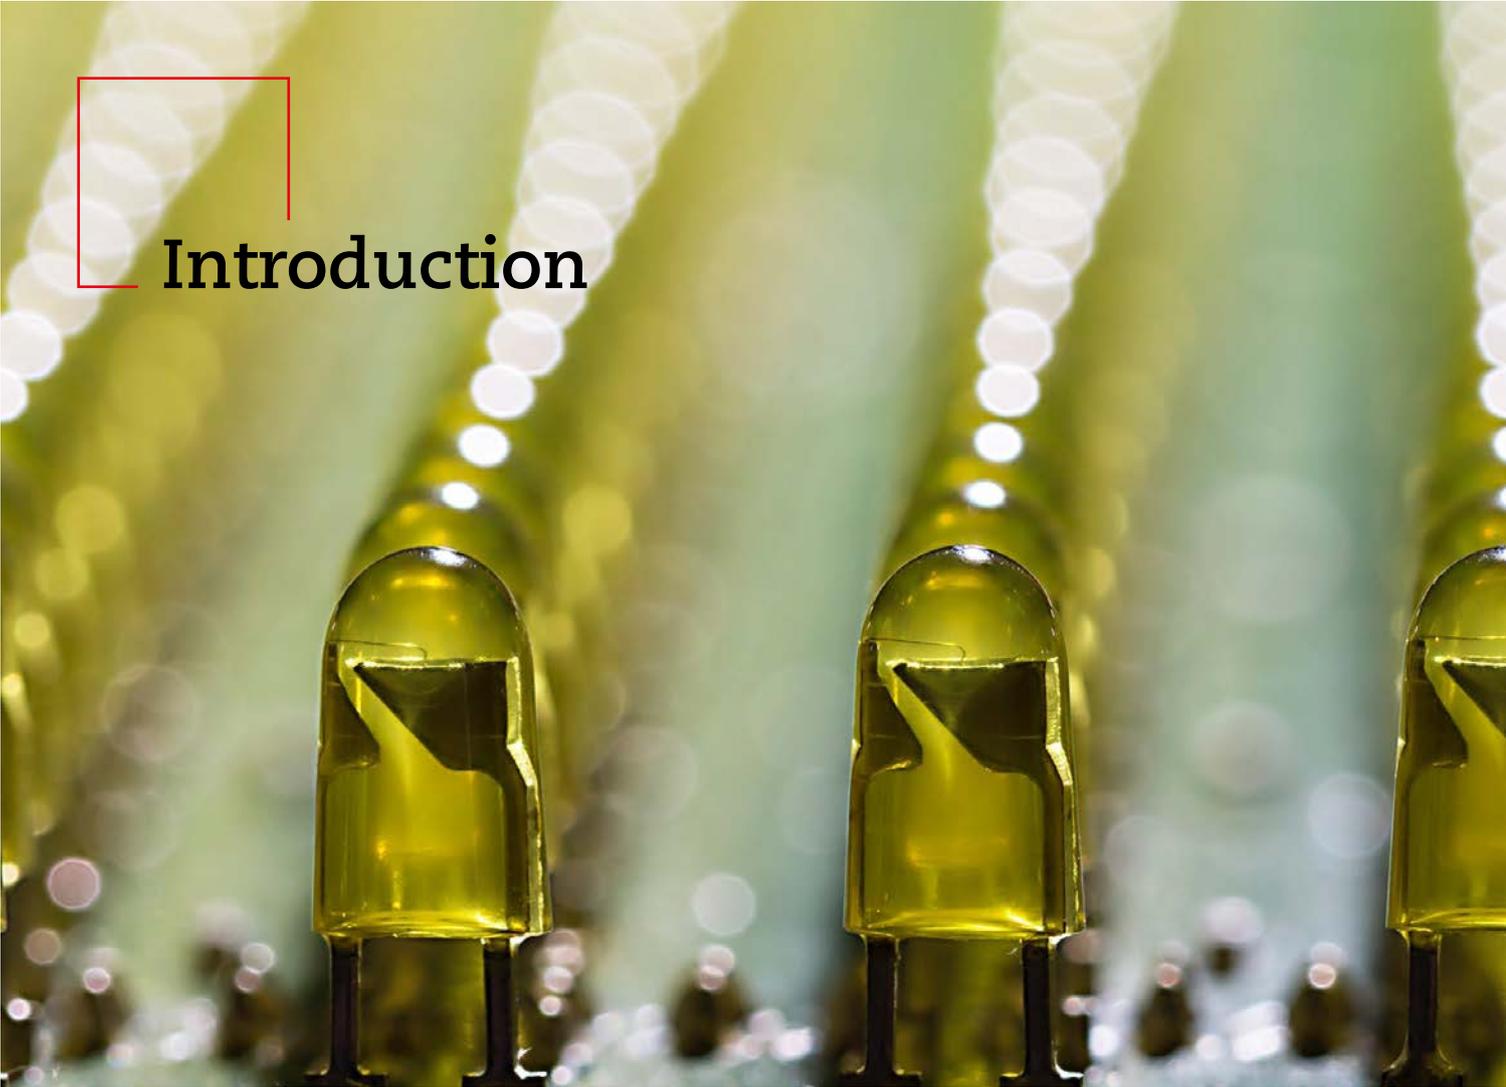
The ability of AI to take over jobs is a narrative often recited in the media. And there is certainly some truth about the ability of these technologies to do a variety of tasks more efficiently – indeed, as mentioned above, this report specifically explores some of these areas.

But even sophisticated technology such as AI appears to struggle with the full contextual understanding and integrated thinking of which humans are capable. Despite advancements in AI, it does not yet appear to be the case that human oversight can be done away with completely; or that the technology can take into account human factors, such as when building client relationships or leading successful teams.

ACCA's work on the emotional quotient (EQ) strongly demonstrated the need, in a digital age, for competencies related to emotional intelligence (ACCA 2018). In fact as we look ahead, the Digital Quotient (DQ) and EQ are best seen combined for either to be really effective for professional accountants.

Even outside behavioural areas such as leadership, core technical activities require judgement and interpretation that draw on multiple considerations. ML can provide truly insightful information, using sophisticated algorithms to analyse historical data sets. But in some situations, a human may choose to take note of this but for perfectly valid reasons, make decisions based on additional/other factors, that do not follows patterns seen in the past.

Looking ahead, professional accountants have an opportunity to develop a core understanding of emerging technologies, while continually building their interpretative, contextual and relationship-led skills. They can then truly benefit from the ability of technologies such as ML to support them in the intelligent analysis of vast amounts of data.



Introduction

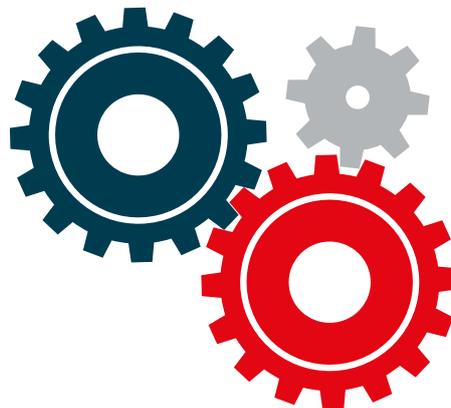
Machine learning (ML) is part of an umbrella of terms used when there is a reference to artificial intelligence (AI), the latter term having been coined as far back as 1956.

Most early AI work relied on a 'decision tree' approach to mapping options, for example, in chess, mapping all possible opening moves and subsequent counter-moves. With even relatively simple problems, such as a retailer making customer-specific recommendations, the vast number of options in a decision tree led to a combinational explosion that could not be processed by even the most capable hardware.

This created a series of disappointments about AI, a so-called 'AI winter', where computing capability lagged behind theoretical approaches and fell significantly short of hopes for the creation of usable applications. In recent years, however, AI has enjoyed renewed interest. This is not science fiction; rather it is now increasingly found in consumer technologies and business applications.

So what has caused this?

Data-driven insight is at the heart of the 'intelligence' driving AI. And it is the exponential increase in the availability of data and unprecedented computing power for processing this data that have jointly contributed to moving AI increasingly from fiction to fact.



It is worth interrogating this observation.

Broadly speaking, there are two levels of AI – specific or weak and general. As it currently exists, the term 'AI' refers to weak AI. This means the use of AI in solution-specific applications, for example in identifying patterns within a large volume of transactions. What is not currently possible is artificial general intelligence – the sort of AI often depicted in films and television, with robots displaying human-like intelligence and characteristics.

While there are some who believe this latter type of so-called 'sentient' understanding may one day be possible, current technological reality appears to be far away from this. As many experts have noted¹, high-performance adult-level intelligence for a single activity, such as needed for playing chess, can be easier to model than human mobility or perception – even that of an infant.

¹ Referred to often as Moravec's Paradox, the discovery by artificial intelligence and robotics researchers Hans Moravec, Rodney Brooks and Marvin Minsky in the 1980's that, contrary to traditional assumptions, high-level reasoning requires very little computation, but low-level sensorimotor skills require enormous computational resources.



As a finance professional it is important to develop an appreciation of all this, given that machine learning is being increasingly used in accounting software and business process applications. This report aims to aid the process of developing this understanding.

The report provides an introductory, end-to-end perspective on ML. It explains the basics of what it is, and identifies use-cases where this technology is being deployed. It further delves into the ethical issues the finance professional may need to consider, and implications of the technology for the future skills required in the profession.

In addition to inputs from experts in the field and ACCA's technology research more broadly, the report is informed by a survey of 1,897 ACCA members and affiliates, and a roundtable discussion on 'ethics in machine learning' conducted in conjunction with the Financial Reporting Lab, the learning and innovation hub of the Financial Reporting Council, UK. We are grateful to the following delegates for sharing their views at the roundtable:

- Andreas Georgiou, Sage
- Dorothy Toh, King's College London
- Lisa Webley, University of Birmingham
- Maria Mora, Fujitsu
- Ruth Preedy, PwC
- Shamus Rae, KPMG
- Stuart Cobbe, Brevis
- Thomas Toomse-Smith, Financial Reporting Lab.



1. Machine learning and accountancy

Double-entry accounting traces its roots to the medieval period, and from that time onwards it has served as the worldwide basis for business record-keeping. The business processes by which those records are created, and by which independent auditors evaluate the accuracy and completeness of those records, have evolved over time.

Despite this, an accountant from the late 1500s and one from the late 1900s would have had enough assumptions in common, linked to the double-entry approach, to allow them to have a professional conversation in a meaningful way.

So accountancy practices have broadly been keeping pace and evolving with developments over the last 500 years, while retaining some common elements over time. And the question now is how might technologies such as ML create the next big transformation?

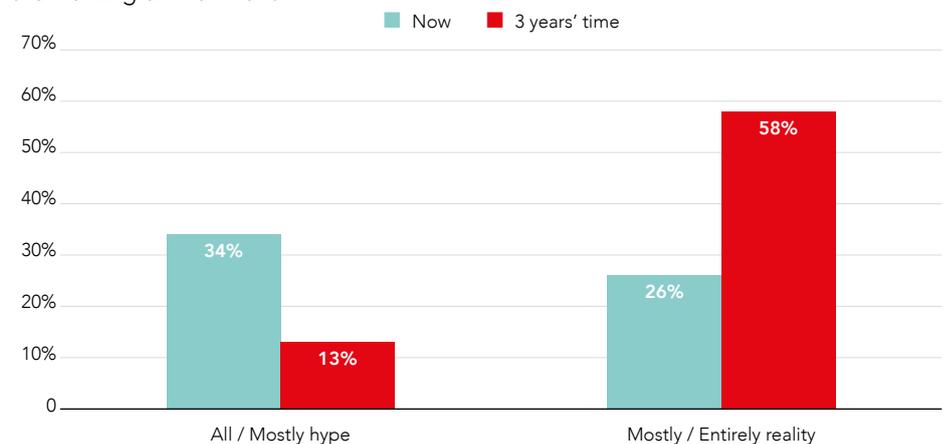
The view from ACCA's survey is that AI is currently perceived as more 'hype' than reality; but that this is set to change in the relatively near future (Figure 1.1).

As of mid-2018, the online publishing platform Medium reported that there were over 3,400 AI/ML start-ups around the world. As with any new venture, the vast majority of these will fail, and many will do so because they are 'solutions' in search of problems, rather than actual solutions to a specific set of business problems or needs.

ML is capable of many amazing things but do accountants really have a need for any of those amazing things to do the job well? On the whole, the answer appears to be 'yes', and this is not just a matter of

staying current. The capabilities that machine learning offers could assist the work of professional accountants in various ways over time. One of the key drivers of this is the proliferation of data.

FIGURE 1.1: Artificial Intelligence: 'Hype' versus reality based on what can be seen in the working environment



Note: remaining respondents said 'Equal hype and reality'

It is estimated that around

90%

of all the digital data in the world has been created since 2016



It is estimated that around 90% of all the digital data in the world has been created since 2016². And the rate at which new data is being generated is not just growing, but appears to be growing exponentially, rather than in an incremental or linear manner.

It is fair to point out that not all this data is necessarily of interest to accountants. But even looking at areas of more obvious interest, such as financial transactions, the trend towards increasing amounts of data remains relevant for various reasons.

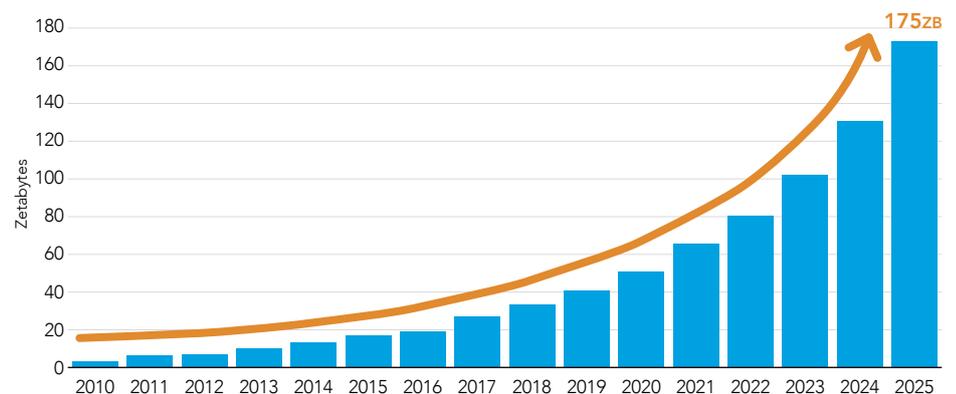
- In much of the world, digital methods are rapidly replacing cash as the preferred way of paying. In China, for instance, mobile payments are rapidly reducing the relevance of carrying cash³.
- Internet of Things (IoT) devices, streaming services and transactionally priced cloud-based hardware and software solutions have led to the growth of small-value, high-volume financial transactions.
- The success of financial inclusion initiatives around the world has led to many more participants in the global financial system. From 2011 to 2018, over 1.2bn people entered the financial system for the first time, and each of them is a source of financial transactions that did not previously exist⁴.

This rapid growth in the volume of financial transactions, if not properly managed, could pose a threat to the work of accountants. For auditors, this may relate to the sample they need and its ability to be representative of the population, enabling them to form conclusions that can be generalised beyond the sample.

As referred to by Forbes⁵ and others, the volume of transaction data is estimated to grow significantly between now and 2025. So, there will be a need to deal with orders-of-magnitude more data, rather than incremental increases, and a need to understand the distribution and profile of this significantly enlarged pool of data.

An implication of this will be pressure on current resources and the ability to scale-up procedures reliably to understand the population being assessed, for example to deal with larger sample sizes. But in fact technology like machine learning could go beyond that with the possibility for reviewing entire populations to assist the auditor to test for items that are outside the norm. Such developments may make ML a matter of necessity rather than just competitive advantage; as the latter will reduce anyway, when many in the market start to adopt it.

FIGURE 1.2: Annual size of the global data sphere 2010–25



Source: IDC Global DataSphere, November 2018

² <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#4d881f9a60ba>

³ <https://www.pymnts.com/cash/2019/mobile-payments-cash-china/>

⁴ Global Findex database

⁵ <https://www.forbes.com/sites/tomcoughlin/2018/11/27/175-zettabytes-by-2025/#1b6ead315459>

2. Navigating the terminology

AI is often used as an over-arching term for advanced computing capabilities, with machines being able to ‘think’ for themselves. And as mentioned earlier, specific or weak AI is the current reality, as opposed to artificial general intelligence. Such nuances can be useful to bear in mind, when sifting through the range of terms involved.

The challenge is that there is no definitive industry standard or agreed definition of exactly what each of these terms means. This can result in confusion and differences of opinion, making attempts at definition a minefield. Nevertheless, it is helpful to have some view on these matters, particularly for those new to the field, and the schematic shown in Figure 2.1 represents one attempt at this.

One way of describing AI is ‘the ability of machines to exhibit human-like capabilities in areas such as thinking, understanding, reasoning, learning or perception’. It is also often referred to as including the ability of the machine to make decisions on the basis of these processes.

Some make a distinction between AI and augmented intelligence, which can be used to refer to the elements above, but excluding the decision making, ie where a person relies on the outputs of such a process to make the final decision.

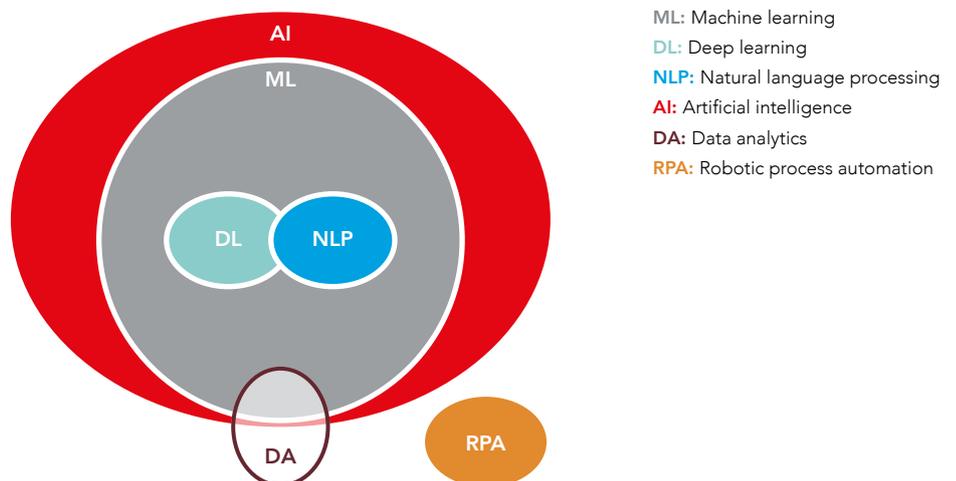
Of the terms in Figure 2.1, data analytics is relatively widely understood (Figure 2.2). It generally refers to the ability to conduct data analysis to extract insights using a variety of techniques. For example,

forecasting future sales on the basis of their dependency on an underlying driver might involve the use of a simple linear regression.

The schematic in Figure 2.1 shows an overlap between data analytics and machine learning. This is to represent that there can be some overlapping of the techniques used, for example regression exists both in data analytics literature as

well as in ML literature. Nonetheless, data analytics is generally seen as a task that is controlled and led by explicit human instructions. The more advanced use of these techniques (and others) on large data sets, which can eventually enable a machine to function, in some sense, without explicit instructions, for example to draw inferences, is generally a characteristic more closely associated with AI/ML.

FIGURE 2.1: A wide range of terms are involved



RPA is in fact a piece of programmed software that implements a defined sequence of activities – like a very high-end Excel macro.



Robotic process automation (RPA) has been placed outside the AI circle in Figure 2.1. This is because, despite the word ‘robotics’, RPA does not refer to robots in the sense of the human-looking intelligent robots sometimes depicted in the media. RPA is in fact a piece of programmed software that implements a defined sequence of activities – like a very high-end Excel macro. There isn’t an AI element in this and it is, at its heart, process automation: in other words, taking a defined process and repeating it tirelessly, quickly and without errors.

While this section discusses these terms as static entities, it is worth noting that this can be simplistic because these technologies are moving rapidly. Innovations across different technologies do not happen in isolated groups.

One area of emerging innovation is the combination of RPA and AI elements, so-called intelligent process automation (IPA). This is increasingly being explored by various technology companies, eg Alibaba via its Aliyun Research Centre. IPA is a form of standard robotic process automation (RPA) in which the system can learn over time from the data and processes on which it is working. With this element, over time, IPA might provide opportunities for process improvement as much as process automation.

Coming back to Figure 2.1, ML is a sub-set of AI that is generally understood as making predictions or decisions from the analysis of a large historical dataset. Essentially, it involves the machine, over time, being able to learn the characteristics of data sets and to identify the characteristics of individual data points. This allows it to identify relationships in complex and large data sets that would be more time consuming or more difficult for a human to see. An ML system can be said to ‘learn’ in the sense that over time, as it is fed more data, it can improve its recognition of the patterns therein, and apply this improved recognition to new data sets that it may not have seen previously. Machine learning is increasing in its relevance as a tool for business use and is discussed in detail later in this section.

Deep learning (DL) and Natural Language Processing (NLP) are generally thought of as being within the ML family. They can handle more complex data, including unstructured data, such as images. This can allow for greater complexity of patterns that can support, for example, image recognition or speech recognition. These are briefly discussed later in this section.

Finally, and more generally, a term that can come up during references to AI is ‘cognitive technologies’. It is relatively difficult to agree a definition for this term, which can refer broadly to technologies that seek to replicate the way the human brain processes/interprets information.

One of the criticisms levelled at AI as a term is that it is frequently used to refer to technologies that are expected to arrive 5–10 years in the future, and that they permanently remain 5–10 years in the future! In reality, technologies are on a continuum of evolution, where they acquire more ‘intelligent’ characteristics over time as the technology evolves. And often, once a capability is realised and becomes mainstream, the AI label gets embedded into business-as-usual technologies and processes.

Increasingly, ML techniques are being buried deep in applications and websites, replacing traditional software in ways that may not be obvious or visible. An example is Uber’s pricing system. Where 10 years ago this would have been hard-coded logic, a trained model now makes these decisions. It looks nothing like artificial general intelligence, but it performs a specific task to great accuracy. Viewed from the outside, the embedding of this AI software creates an increase in the operating effectiveness of the whole – a cost-saving development even if not a radical change.

A well-documented example of AI that has become ‘normalised’ is optical character recognition (OCR), ie the ability to extract text from scanned copies and documents. The traditional method involved a rules-based template that had to be set up in advance, with the system extracting and mapping the patterns to the text in line with that template. Templates easily become complex, for example to cope with data tables or even text in columns.

On average, for each of the terms considered, about one-third of respondents either had not heard of that term, or had heard of it but did not have an understanding beyond that.



The AI-driven leap here has been to remove dependence on the rules-based template; in other words for the AI to create its own mapping between the layout and the text or character to which it should be mapped. As this has become more common, however, it is generally thought of as just 'OCR' and the AI-enabled back-end is forgotten.

Among the respondents to ACCA's survey, the understanding of certain terms was much greater than that of others. On average, for any given term, one-third of respondents had either not heard of it, or had heard of it but didn't know what it was (Figure 2.2).

While professional accountants may not need to develop ML algorithms themselves, this section will provide an introductory sense of how ML works in the background. This matters because it influences trust – and the ability to have a view on whether one can trust the decisions of these systems and the contexts in which they operate.

This is also important in order to have an appreciation of how ML relates to, or differs from, other terms often mentioned in this area. In the survey, 'data analytics' was the best understood with only one-fifth of respondents stating they were not sure about how it differed from ML (Figure 2.3).

In ML one is dealing with a powerful tool with tremendous potential. This is because AI encompasses an enormous range of applications. These include recommendation engines; fraud identification; detecting and predicting machine failure; optimising options-trading strategies; diagnosing health conditions; speech recognition and translation; enabling conversations with chat-bots; image recognition and classification; spam detection; predicting everything from how likely someone is to click on an advertisement, to how many new patients a hospital will admit; through to autonomous vehicles.

FIGURE 2.2: Understanding of terms

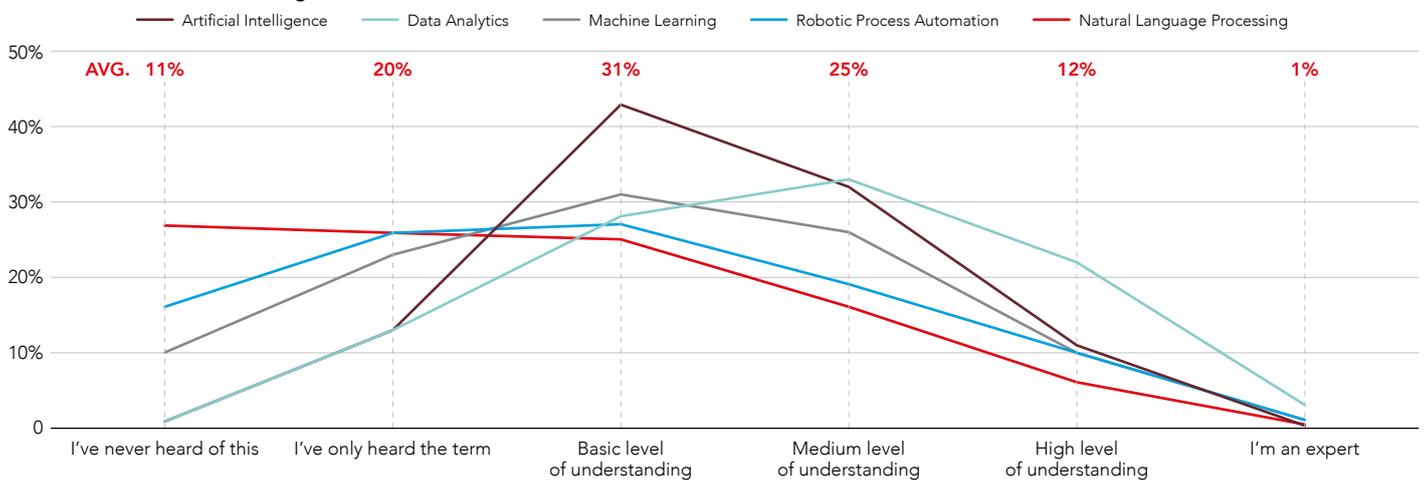
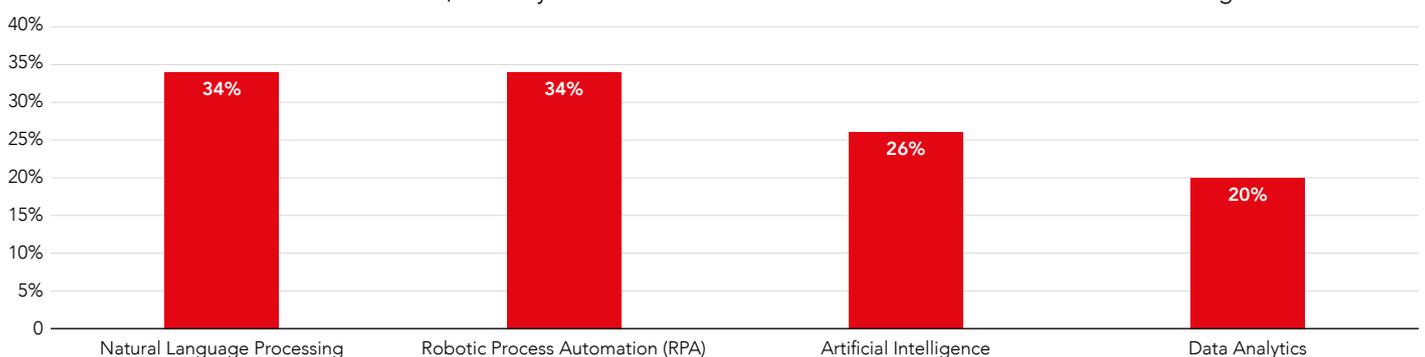


FIGURE 2.3: For each of the terms below, state if you are not sure how it differs from or relates to machine learning



At the heart of this process is a mathematical model – the algorithm – that is used to describe and/or predict features in the data set.



WHAT IS MACHINE LEARNING?

ML is a sub-set of AI and is generally understood as incorporating the ability for computers to ‘learn’, ie where the outcomes are not being explicitly programmed in advance.

Explicit programming refers to traditional computer programs, which are said to be ‘imperative’; in other words, they provide specific instructions for how a task is to be executed. This specific set of instructions is hard-coded by a human programmer, and generally includes such elements as sequential steps, logical checks, functions and loops. Therefore, running a program on a data set will provide a result based on a fixed set of rules embedded in the program. In other words, the way the program will deal with the data is fixed in time – the time when the program was written.

By contrast, ML uses statistical analyses to generate results dynamically from the data set. At the heart of this process is a mathematical model – the algorithm – that is used to describe and/or predict features in the data set. The starting point is a ‘training’ data set of inputs. This training data allows the model to learn which features of individual data points are important. The point here is that this algorithm can then be used with new data that was not part of the initial training data set. If the new data suggests additional/different patterns, then the algorithm can iteratively adapt to incorporate this into a now-updated understanding of the characteristics of the data. This enables ML to adapt to new, unseen data in a way that traditional programming could not. And it is in this sense that ML ‘learns’ from examples rather than strictly following the pre-coded logic in traditional programs.

The ‘learning’ that ML undergoes relies on pattern recognition between the data elements involved. If, for example, the data consistently shows correlations between umbrella sales and level of rainfall, the algorithm may ‘learn’ the relationship between the two. But that does not mean it has a contextual understanding of the fact that it is uncomfortable or inconvenient to get wet in the rain. So that is still very different from ‘thinking’ in a human sense, which

includes a wider level of perception, lateral and creative thinking as well as the ability to process emotional information.

Let us consider a simplified illustration. Say an organisation seeks to improve working capital by gaining a better understanding of the counterparties most likely to default on payments. Traditional approaches would be for a human to create a program by taking a view on what drives default behaviour. They might decide that the rules of such a program would depend on creating a basic scoring system. The program might be set up to flag all those counterparties who match a certain profile, eg those who have previously made late payments, who operate in certain jurisdictions, have to make a certain value of payment, etc. The output from the program here could be a list of high-risk counterparties most likely to default.

The input here could be data about all the transactions made by the counterparties being examined. The output of the program would be all those counterparties that satisfy the logical tests set within the program to flag high likelihood of default. The challenge with this is that it is based on a static view taken upfront on what a ‘bad’ counterparty looks like. In other words, it is based on the programmer’s view of the characteristics of a counterparty who is likely to default – a view taken at the time that the program was developed and used to inform the structure of the program. As counterparties, transactions, business profile and volumes evolve over time, this may change. Also, as the number of variables to consider increases – as is likely in real-world applications – creating a static set of rules for deciding, in advance, the criteria for filtering high-risk counterparties, would become increasingly complex and inaccurate.

In this type of scenario, ML might be used to create an algorithm based on a training data set that suggests high-risk counterparties. It could take in a wider pool of input variables and end up identifying correlations that might not have been considered by a (human) programmer when creating the program. If this is done well, the ML system can improve in its ability to do so over time, improving, rather than degrading in quality, the matches made.

Because fraudsters are constantly creating new techniques to 'cheat the system', new areas for testing correlations need to be constantly developed to identify potential fraud, a type of challenge well suited to ML.



Continuing this simplified example, the ML system could use wider macroeconomic data about the operating environment, credit-rating data from third-party scoring organisations or the level of positive/negative information about the counterparty available on the internet in time periods up to the present. It is worth noting, however, that this approach also relies on historical data, even if it is a much wider data set.

Nonetheless, unlike a traditional program, ML takes a probabilistic approach. It uses the data to establish a statistical basis for the likely patterns, correlations and characteristics of the data. And as it is introduced to new data, the algorithm can dynamically incorporate new correlations if these are now detected.

As with all statistics, the broader and more representative of reality the data set, the more reliable are the statistical results. One might have a 20% chance of error in drawing conclusions from a small data set, but only a 2% chance of error in doing so from a large data set that accurately reflects the population being modelled. This is why having sufficiently⁶ large data sets of good-quality data really matters for ML to work properly.

This capability is showing potential to be faster, and/or more economical, than a human and to be able to handle volumes of data in which humans may struggle to identify possible relationships to inform the programming.

Taking scenarios such as fraud detection, humans struggle to keep up with the new and innovative ways fraudsters use to manipulate systems. This is exacerbated when looking for fraud within a huge volume of data. Because fraudsters are constantly creating new techniques to 'cheat the system', new areas for testing correlations need to be constantly developed to identify potential fraud, a type of challenge well suited to ML.

APPROACHES USED IN MACHINE LEARNING

This report does not seek to focus on all the nuances of this complex area. But at a high level, the majority of current activity falls into a few types of ML.

Supervised learning involves algorithms that are 'taught' by examples, with real inputs and outputs. The algorithm connects the two using the 'correct' answers that are provided in the trial data, so that the algorithm can form a baseline view of the correct patterns or relationships.

Supervised learning can be used for classification problems, such as image recognition, where examples are 'tagged' with contents, and used to train a model to identify new images. For example, the system can be taught to predict whether a photo is or is not a cat by previously tagging as 'cat' a large number of images of cats.

Reinforcement learning is a type of learning, which is used generally where real outputs are not available but the quality of a generated output can be measured as 'good' or 'bad' and this is then fed back into the algorithm. This feedback is used to improve the algorithm quality. Autonomous driving is an example of reinforcement learning. The algorithm aims to provide 'good' driving, therefore not crashing or driving dangerously, and a reward system, based on the (unpredictable) conditions it experiences, is used to shape the algorithm.

Autonomous driving is, however, very complex and cameras will be trained using supervised learning algorithms to recognise objects – person, car, cyclist, tree, etc. These algorithms then feed into a reinforcement algorithm – the combination of 'objects' is infinite, so the algorithm cannot learn every situation. It 'just' needs to be as good as a human at interpreting them.

⁶ It is important to know how to recognise excessively large additions to the data sets that do not add any incremental value and that result in 'over-fitting'.

Data preparation is often highlighted as a bottleneck, as it is time consuming and requires manual effort, so unsupervised learning often achieves results faster.



Unsupervised learning is used where the input data contains no answers. Data is not classified or labelled, and the algorithm is left to interpret data, without guidance, and to try and create a structure that explains it. Unsupervised learning does this by identifying similarities and differences in data, using techniques such as clustering. A common use for this technique is in the area of detecting anomalies in a data set, such as when looking for fraudulent transactions, or patterns of association, such as when certain products are purchased together as part of a shopping basket.

The results for supervised learning are typically more precise, but this approach usually requires data preparation. Data preparation is often highlighted as a bottleneck, as it is time consuming and requires manual effort, so unsupervised learning often achieves results faster.

WHERE DO DEEP LEARNING (DL) AND NATURAL LANGUAGE PROCESSING (NLP) FIT IN?

DL is a specific ML approach that uses 'neural networks'. Neural networks (often referred to as artificial neural networks – ANN) are loosely based upon the biological neural network of a human brain. An ANN can be built up of many layers of nodes, and the flow of signals can pass up and down layers before it reaches the last layer (output layer) – having started at an input layer. The term 'deep learning' refers to the depth of layers between input and output in an ANN.

DL gives NLP greater accuracy by allowing for improved prediction. Without DL, NLP typically analyses the preceding four or five words to determine what the next word is 'likely to be'. DL can use all previous words to build greater reliability of outcomes. NLP has been defined as one of the 'hard-problems' of AI, not least because of the use of the same words in

different contexts, eg 'book': a bound collection of pages (noun) vs. to make an appointment (verb).

While ML algorithms are all geared towards cognition, DL can be particularly useful in the area of perception. Examples of perception-related applications include the following.

- Voice recognition is found in everyday use in digital assistants such as Siri, Alexa and Google Assistant. It is estimated that speech recognition is now about three times as fast, on average, as typing on a cell phone, with an error rate under 3%. This is still being refined as such systems meet constant challenges, for example when dealing with technical words, or localised language with regional accents.
- Image recognition: facial recognition (eg iPhone X, Facebook, self-driving cars, Imagenet). In 2007 Fei-Fei Li, head of Stanford AI lab, gave up trying to program computers to recognise objects and instead switched to labelling and DL. The result was Imagenet, with a vast database of images and an error rate of 5%, which makes it 'better than human' and created a 'tipping point' for image recognition technology.

NLP has also been a central element in many developments of AI, ML and DL, and again this is most visible in the emergence of digital assistants, and in the widespread commercial use of chatbots.

Examples of NLP activities have included:

- **speech recognition:** voice to text conversion
- **natural language understanding/interpretation:** to provide comprehension of text
- **machine translation:** language translation of text.

3. Applications of machine learning

There are a variety of applications for ML and this section gives a flavour of some of these. As might be expected, there is a spectrum of ways in which ML can be adopted.

The survey found that about 2 in 5 respondents were actively engaged with exploring ML adoption (Figure 3.1). Their progress ranged from early stage discussions exploring concepts, through to full production mode with live data.

Respondents expressed varying levels of comfort (Figure 3.2) with making decisions based on ML across areas such as classification (53%), measurement (47%), audit testing (43%) and fraud detection (41%). There was, however, less comfort in certain wider applications such as with medical data or personal finances.

FIGURE 3.1: Status of machine learning adoption in my organisation

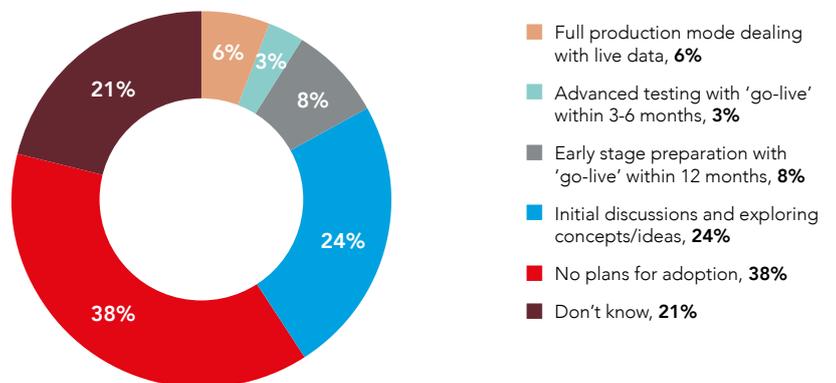
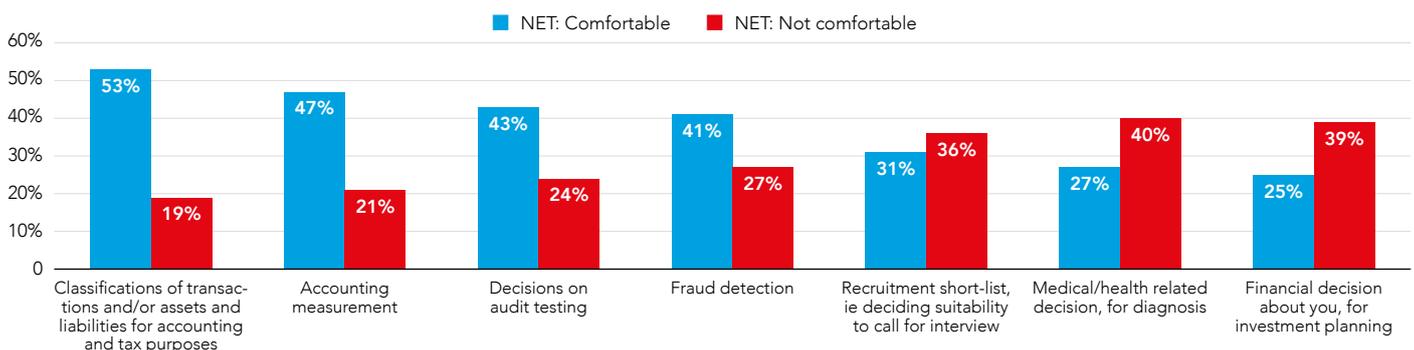


FIGURE 3.2: How comfortable would you be with machine-learning-based decision making on the following specific tasks?



Note: 1-5 scale with higher number indicating greater comfort; NET Comfortable is sum of 4, 5; NET Not Comfortable is sum of 1, 2

The large accountancy firms are all investing in ML to explore possibilities, for instance in audit and compliance. And in time the base of published evidence supporting the benefits of ML is likely to increase.



When considering the relevance of ML to audit, respondents broadly viewed it as a potentially useful tool. Its ability to enable better identification of patterns indicating fraud transactions was cited as a factor. Also, in a world where Big Data is prevalent, ML was seen as needed for analysing the volume and complexity of some information generated. But there was also caution about where and how it was relevant. For example, some questioned whether the use of ML might compromise external auditor independence owing to the reliance on algorithms provided by management.

Clearly, these and many more considerations must be taken into account as ML seeks to enter the accountancy mainstream. Adoption is a journey and there are inevitably barriers to be faced in embracing the opportunities it may present. The most commonly cited of these were a lack of skilled staff to drive the adoption, and costs – both of which were cited by about half of respondents (Figure 3.3). Problems with data, which is a critical raw material for this, were also cited. About a quarter of respondents cited the poor quality of data, and 17% the lack of a sufficient volume of data.

About one-fifth of respondents cited the lack of a clear benefits case in support of adoption. While it may be that the case has not been adequately explored or understood, it may also reflect a view that ML is simply not always the best solution for the particular questions being tackled. The starting point has to be a legitimate business need that can be best addressed by what ML provides.

In addition to the broader conceptual observations on adoption, a few specific illustrations are discussed in the section that follows. These have been drawn, where possible, from real-life examples in order to provide a sense of current developments.

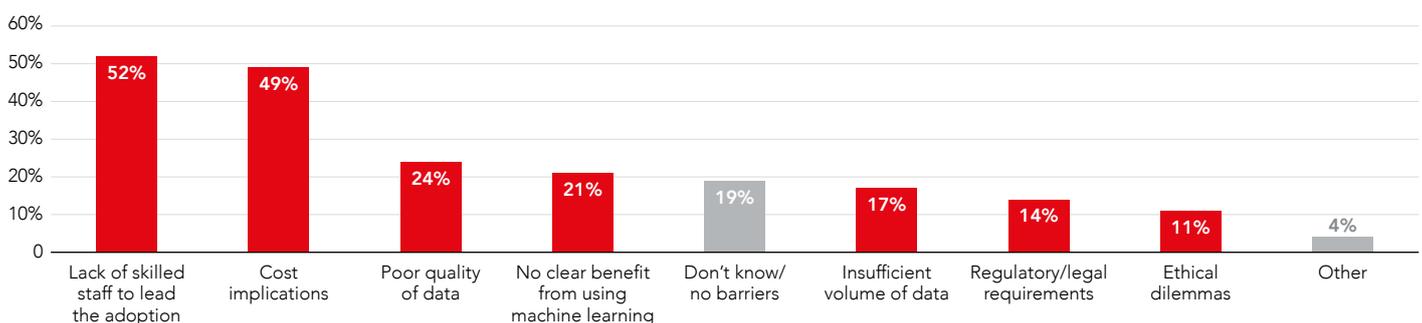
INTELLIGENT BOOKKEEPING

In general, the use of ML is in relatively early stages. The large accountancy firms are all investing in ML to explore possibilities, for instance in audit and compliance. And in time the base of published evidence supporting the benefits of ML is likely to increase.

In bookkeeping, ML systems have already been in full production for a few years, particularly in the small and medium-sized enterprise sector. For example, the market offers products that are able to scan expense receipts and classify them automatically. The more advanced of these products use a combination of reinforced learning and NLP to automatically parse, extract, and classify scanned receipts without the submitter having to type in any identifying information. For example, according to Expensify's website, the company's product has over 6m users and over 60,000 companies using their solution, and process billions of transactions each year.

Online accounting software provider Xero announced in May 2018 that its ML software had already made more than 1bn recommendations to customers since it became available, with areas of invoice coding and bank reconciliations being prominent. This figure includes more than 750m invoice and bill code

FIGURE 3.3: The main barriers to using machine learning in respondents' organisations



In this risk assessment, supervised learning algorithms can be used to help identify specific types or characteristics that warrant greater scrutiny; and improve targeting of the areas of focus for the audit.



recommendations, and more than 250m bank reconciliation recommendations. Xero estimates that with 800,000 invoices filed each day in Xero this is a collective saving of 307 hours.

On coding of invoices, the Xero software 'learns' how a business codes regular items and auto-fills on the basis of this 'understanding' of history, rather than the labour-intensive traditional use of default codes. Using this approach, it correctly codes 80% of transactions after just four examples. The company's blog post suggests that it is using a logistical regression approach to get the best prediction but, understandably, for competitive reasons details of the predictive algorithms are not available.

According to Kevin Fitzgerald, Asia Pacific Director for Xero:

'We see machine learning algorithms being helpful in providing intelligent support that can free up the time of professional accountants to focus on the financial and strategic agenda of their clients or their own organisations'.

When initially implemented, these codings were provided as suggestions to the user, and required specific, albeit easy, validation or correction if necessary. Xero deliberately did this so that the algorithm would learn user behaviour. The company has stated: 'We're watching very closely the rate that customers actively disagree with suggestions by choosing something else, and the rate of later recodes of suggested accounts. On recodes, the system absolutely learns from those. It's part of the basic idea – it only knows what it's been taught. If it learns from correct accounts, the suggestions will be correct'. This goes beyond a static rules-based approach to a true ML capability.

For bank reconciliations, the Xero ML software integrates with that of many banks, which feed account transaction records automatically into Xero. It then matches bank transactions with payment and receipt records in Xero, with automated coding based on how similar transactions have been previously coded. As with invoice-coding, the ML for bank reconciliation incorporates user modification to transaction matching to improve recommendations.

Both the Invoice Coding and Bank Reconciliation models are based solely on the experience of the specific business, not on those from a wider pool of entities. This naturally limits the degree of 'intelligence' demonstrated, and prevents the software from applying pre-built knowledge to new customers. The company recognised the challenges with this, early on: 'It's true that there is potential to learn from other organisations as well, but our early research has shown that there is huge variation in practice and encoding between different businesses – far greater than we expected'.

This kind of standardisation is envisaged as a future enhancement as it can lead to further efficiency improvements in customer activity, but highlights the challenge in creating an 'intelligent' coding bot.

IMPROVING FRAUD DETECTION

One of the areas where ML can help is with risk assessment. The reference here is to the ability to assess the likelihood of fraud, inaccuracy, misstatement, etc. based on a mix of empirical data and professional judgement. In this risk assessment, supervised learning algorithms can be used to help identify specific types or characteristics that warrant greater scrutiny; and improve targeting of the areas of focus for the audit. In this context, the choice of an appropriate ML method can be valuable for audit testing.

Using ML as part of the audit process is in relatively early stages, and publicly available empirical data to support the assertions of improvement are being steadily built over time. One example is a study commissioned by the Comptroller and Auditor General (CAG) of India (Yao et al. 2018).

CAG is an independent constitutional body of India. It is an authority that audits receipts and expenditure of all the organisations that are financed by the government of India. One of the CAG's duties is to uncover organisations set up for fraudulent reasons. In fulfilment of this duty, each year it selects a number of organisations to be audited. Some are selected via public complaint or direct referral, while others are selected by monitoring news sources and business results but, historically, a significant number are selected by random sample.

One of the interesting features of the CAPS model is that it optimises results not only in relation to likelihood of detection but also to the return on investment (ROI) of the program itself.



CAG wished to check the applicability of using ML methods during audit planning to predict the prevalence of fraudulent organisations. This type of prediction is an important step at the preliminary stage of audit planning, as high-risk organisations are targeted for the maximum audit investigation during field engagement. A complete Audit Field Work Decision Support framework exists to help an auditor to decide the amount of field work required for a particular organisation and to identify low-risk ones that can be omitted from the audit.

CAG was interested in seeing which ML algorithms were most effective at predicting the risk that a given firm is fraudulent. In this study, CAG selected a historical set of over 700 firms it had recently audited and used that as input for 10 different ML algorithms to determine which ones performed the best. For this specific case, the algorithms were trained to prioritise sensitivity over specificity. In other words, failing to detect a fraudulent firm (Type II error) was deemed more damaging than incorrectly identifying a genuine firm (Type I error).

The rationale for this weighting was that a false positive merely triggered a human investigation, which would presumably reveal that a firm was indeed genuine, while a false negative allowed fraud to continue undetected.

In aggregate, the most accurate algorithms were able to identify suspicious firms correctly 93% of the time. The reported results were quite detailed, but in summary, of the 10 different ML methods tried in the study, no one method proved to be the most accurate across all transaction types and industry groups (Yao et al. 2018). Therefore, understanding what algorithm to use and why is extremely important. These findings demonstrate not only the potential value that ML techniques can add to the audit process, but also the importance of having a sufficient understanding of ML techniques to be able to select the most appropriate methods for specific instances.

While the above example relates to government and is relatively recent, there are earlier examples of private companies experimenting with ML. Intel, for example, established 'Intel Inside', a cooperative marketing campaign in which

technology manufacturers externally label and brand their products as containing Intel components. It is considered one of the earliest successful examples of 'ingredient marketing'.

Participating manufacturers benefit from the reputation of the Intel brand, but they also benefit more directly from funded co-marketing activities, which has motivated many enterprises to seek these benefits fraudulently, ie to use the 'Intel Inside' branding without actually using Intel components in their products.

Intel attempts to monitor compliance by inspecting companies that are known to use the 'Intel Inside' branding. Historically, it selected which companies to inspect through a combination of manual and random selection. Then, in 2011, Intel began developing what it calls the Compliance Analytics and Prediction System (CAPS), which uses a combination of supervised learning techniques to predict which claims are most likely to have compliance issues, and to refer those claims to Intel's inspection team for further investigation.

One of the interesting features of the CAPS model is that it optimises results not only in relation to likelihood of detection but also to the return on investment (ROI) of the program itself. In other words, information about staff availability and the cost of a compliance investigation are inputs into the training set, and the predictive outputs are not only the likelihood of fraud but also the projected expected value of any potential recovery.

In 2017, Intel published a white paper that summarises the findings across the five years that CAPS has been running in production. There are some noteworthy findings. As a control, Intel continued to perform some compliance audits by random selection. The dollar value of recoveries remained the same over the five-year period; in other words, they scaled with the capacity of the audit team and not with Intel's revenue growth. On the other hand, in 2012, when the study started, the dollar value of recoveries from CAPS-triggered audits was nine times that from randomly selected audits. Over the five-year period, the supervised algorithm continued self-training and, in 2017, CAPS-triggered recoveries grew up to 19 times those generated from randomly selected audits⁷.

NLP and ML, have a role to play in making tax query systems more effective. Using the ML technique of reinforced learning, AI chatbots and speech engines can train themselves to become more effective over time.



MAKING SENSE OF COMPLEXITIES IN TAXATION

ML is also being seen to have applications in relation to tax. Some of these are simply more specific instances of the audit and compliance use cases described above. Governments are particularly interested, as ML may provide dramatic improvements in scale and cost.

But ML has uses in the tax realm beyond predictive modelling. In the US, for instance, the sum total of all federal tax regulations, rulings, and case law amounts to over 74,000 pages worth of content; no single adviser can master it. Accountancy and tax service firms alike have invested millions of dollars in various applications that attempt to help people and enterprises get answers to specific tax questions. These approaches range from books to Web forums to chatbots and full speech-recognition AI systems that attempt to answer tax questions conversationally.

NLP and ML, have a role to play in making tax query systems more effective. Using the ML technique of reinforced learning, AI chatbots and speech engines can train themselves to become more effective over time.

Unsupervised learning also has a role to play. In combination with text analysis software, unsupervised learning can be used to uncover connections and linkages between tax regulations, regulatory rulings and case law to provide answers to tax queries that are more accurate, better informed and more able to withstand challenge.

In one attempt at gathering evidence, KPMG conducted a study in which it measured the ability of IBM's Watson ML application to provide good tax advice for corporations with significant R&D investments. The training set KPMG used to train Watson was a base of over 10,000 documents, and the results were published on IBM's website. These training documents were critical in obtaining a good result. As observed by KPMG's Todd Mazzeo: 'Watson isn't a PhD grad out of the gate. It starts off as a kindergartner and works its way up'⁸.

By the time the machine training was completed, Watson was able to give correct advice to about 75% of queries. For some context, an earlier study by the US Treasury department of the Internal Revenue Service tax help line, found that human operators gave correct advice about 57% of the time⁹.

EFFECTIVE NON-FINANCIAL REPORTING

Environmental, social and corporate governance (ESG) issues are an essential part of non-financial reporting and of managing risk in today's uncertain world. Expanding the scope of reporting to non-financial topics not only gives external stakeholders a more comprehensive picture of the company's performance, but it could also ensure that better quality information is collected for internal decision making, thus improving risk management and even adding greater long-term value to the business.

Nonetheless, approaches to corporate strategy and risk management can be incomplete and outdated. Non-financial topics are often siloed within an organisation. Manual data analysis, expensive consultants and statistically under-representative surveys can make materiality analysis challenging and leave businesses open to risks that could have been foreseen.

Since 2013, there has been a 72% increase in the number of recorded regulations covering non-financial issues, with more than 4,000 non-financial regulatory initiatives, current and draft, to be considered (Datamaran 2018). And this trend looks set to continue.

Materiality is therefore a key factor to ensure focus on the most pressing items. Described with respect to integrated reporting in the International <IR> Framework (paragraph 3.17) as 'matters that substantively affect the organization's ability to create value over the short, medium and long term,' material issues have significant implications for a company's risks and opportunities, making them critical elements for decision making and strategy setting. According to the World Economic Forum's (WEF) Global Risks Report 2019 most of the top risks are ESG-related.

⁸ <https://www.ibm.com/watson/stories/kpmg/>

⁹ <https://www.cbsnews.com/news/irs-cant-do-the-math/>

There is in some sense an underlying 'use' case that forms part of all applications – ML's purpose is analysing data to derive actionable insights.



However, a materiality analysis is time-consuming, with a heavy emphasis on manual labour. It starts with the challenge of choosing which methodology to use. The process of identifying, evaluating, prioritising and disclosing material issues is often subject to the risk that the business overlooks a source or misses an emerging trend.

In referring to non-financial matters and materiality there are two distinct considerations. There is external non-financial reporting, which is at least in part driven by regulatory requirements. These regulatory requirements either overlook materiality (ie mandating that certain measures must be reported in all cases – an example might be level 1 carbon emissions), or set up specific materiality definitions (ie the EU Accounting Directive defines materiality as: 'the status of information where its omission or misstatement could reasonably be expected to influence decisions that users make on the basis of the financial statements of the undertaking.')

But that's a very different perspective from internal management reporting, where information is collated to inform internal management decisions. Materiality in this case would centre on identifying and understanding risks that the business faces – which is focussed on more in this section.

The two do cross over however. Complying with external reporting requirements could force information to be collated internally where they haven't been before, and thus also make information available for management purposes where they have not been considered previously.

Additionally, understanding the stakeholder 'voice' is another challenge. Usually, companies rely on surveys for gauging stakeholder opinion, but this approach has a number of limitations, such as difficulty in reaching sufficient respondents and a low number of returned questionnaires. Overall, it is easy to end up questioning the legitimacy of the actual materiality assessment because there are too many standards to follow.

Platforms such as Datamaran use ML to deal with these challenges. The platform ultimately helps to take control of benchmarking, materiality analysis and processes for monitoring non-financial issues in-house on a systematic and continuous basis. The end goal is to help companies embed non-financial issues into business in a resource-efficient way.

The AI solution supplements manual data analysis and consultants that were the traditional approach to materiality analysis. Supported by a team of data scientists as well as ESG and risk experts, the Datamaran software tracks 100 non-financial topics by sifting and analysing millions of data points from publicly available sources.

These sources include corporate reports (financial and sustainability reports, as well as US Securities and Exchange Commission (SEC) filings), mandatory regulations and voluntary initiatives, as well as news and social media. The NLP technique, which analyses text (narratives) and derives meaning from human language, is then applied to these data sources to extract comparable information (Datamaran 2018).

As a result, the platform provides an evidence-based perspective on regulatory, strategic and reputational risks as well as reporting patterns relevant for a particular company.

MACHINE LEARNING APPLICATIONS COULD BE HERE TO STAY

There is in some sense an underlying 'use' case that forms part of all applications – ML's purpose is analysing data to derive actionable insights. Value-driven business decision making is a permanent need that will always have relevance.

For example, cash-flow management software (the cash-flow forecasting application 'Fluidly' is one example) can help managers to get a more dynamic view of the cash-flow profile, predict future movements and make adjustments to their business accordingly. This has commercial value and can be used to drive advantages in a competitive market.

At present, the ability of ML applications to drive insight has two significant limitations: the size and scope of the training set, and the quality of the data records therein.



The Big Four global accountancy firms have also publicly announced various ML tools and solutions. Some examples are mentioned below though this is a fast moving space with new developments occurring all the time.

Since 2014, Deloitte has partnered with ML provider Kira Systems to perform ML-assisted reviews of leasing contracts¹⁰. Deloitte states having used Kira to perform over 5000 contract reviews to date, and advertises that using it reduces the amount of time it takes to perform a review by 30%.

In 2018, EY¹¹ released an ML audit solution called EY Helix GL Anomaly Detector (HelixGLAD). In an initial test, HelixGLAD was able to spot a small number of transactions in a large corporate ledger that the test team knew to be fraudulent. EY went on to test HelixGLAD in 20 live audits in 2018, and plans to use it on 100 audits in 2019.

KPMG uses an ML tool it calls Strategic Profitability Insights (SPI)¹² within its deal advisory practice. SPI includes unsupervised learning capabilities and is designed to analyse transaction-level data to answer a variety of questions about the target company's customers, products, and supply chain. In addition, there is also recognition of the fact that ML relies heavily on data quality and that innovation will probably need to be across organisations and open-source. KPMG has been working on this area to facilitate the eventual creation of commonly understood data model across organisations.

In 2017, PwC announced its own ML audit tool, GL.ai¹³. The concept behind GL.ai is to move beyond sampling as an audit method and harness the scalability of an automated, ML-informed review to examine a company's entire ledger in search of transactions that warrant further investigation by humans.

But as with any new technology, there are also plenty of innovative solutions in ML coming from new ventures. These include areas covered earlier in this section as well as other applications a few of which are cited below by way of example.

AskMyUncleSam offers an ML-driven chatbot which dispenses tax advice to US taxpayers. Kreditech and OakNorth are two of several companies offering ML credit-risk assessment tools, while AppZen is working on a real-time fraud-detection engine that connects to a company's existing expense-management tools. YayPay is an accounts-receivable application that uses ML to improve cash flow predictions, using a company's historical payment patterns as its training set.

APPLYING MACHINE LEARNING WITHIN A WIDER TECHNOLOGY LANDSCAPE

ML (and AI more broadly) is poised for potentially significant impact on the profession. But it is important not to forget that many other technologies are also in various stages of development and could play a key role in complementing what ML offers.

The linking thread is the data explosion. One stand-out element driving this explosion is Internet of Things. The fact that so many devices, from fridges to phones, can spew out data dramatically increases the raw material for ML to analyse. Furthermore, as this data multiplies, fragmented conventional databases may prove to have their task cut out. Also, distributed ledgers, if they mature sufficiently, could prove to be extremely valuable. They would provide a single and shared version of the facts across a number of interrelated users, which would greatly enhance data quality and therefore the ability of ML applications to add value.

At present, the ability of ML applications to drive insight has two significant limitations: the size and scope of the training set, and the quality of the data records therein. If multiple parties agreed to share their transactions in a synchronised and immutable ledger, both the size and the accuracy of the training sets that ML relies upon could be radically improved. In effect, the intersection of various technologies will act synergistically not only to improve the ROI for each, but also to give rise to new business models not previously possible.

¹⁰ <https://kirasystems.com/resources/case-studies/deloitte/>

¹¹ https://www.ey.com/en_gl/better-begins-with-you/how-an-ai-application-can-help-auditors-detect-fraud

¹² <https://home.kpmg/xx/en/home/insights/2017/09/strategic-profitability-insights.html>

¹³ <https://www.pwc.com/gx/en/about/stories-from-across-the-world/harnessing-the-power-of-ai-to-transform-the-detection-of-fraud-and-error.html>



4. Ethical considerations

Ethical behaviour is a necessary attribute for everyone in society, in both their personal and professional dealings. But for the profession this element is additionally hard-coded into the very definition of what it means to be a professional accountant. And within organisations, it is a key requirement that the finance function provide constructive challenge to ensure that business decisions are grounded in sound ethical principles.

The IESBA (International Ethics Standards Board for Accountants) Code sets out five fundamental principles of ethics for professional accountants, which establish the standard of behaviour expected of a professional accountant (see Appendix 1). So when considering the potential of ML, professional accountants need to think not only of the potential benefits – as demonstrated by the preceding section on use cases – but also the ability to create long-term sustainable advantages. This latter aspect depends in no small way on ensuring that ethical considerations are given sufficient emphasis when exploring ML adoption.

Trust can take years to build and an instant to be destroyed. Clearly ethical behaviour is a non-negotiable requirement for its own sake. Nonetheless, it is also clear that breaching best practice in this area can inflict real damage on the brand/reputation and intangible value of an organisation. In today's social media-driven world, bad news circulates quickly, and not paying attention to ethical behaviour as new technologies are adopted can expose organisations to significant risk, both financial and reputational.

The ethical challenges posed by ML are explored in this section by focusing on five areas. For each area, a scenario is examined where the IESBA fundamental principles could be compromised. In most scenarios most of or all the principles may be at risk but, to draw out specific points, only one or two compromised principles may be highlighted. For those interested more broadly in digital ethics, beyond ML specifically, ACCA's report on *Ethics and trust in a digital age* also addresses relevant considerations (ACCA 2017).



DEALING WITH BIAS

This is arguably the most frequently discussed source of ethical challenge. At its root is the fact that ML algorithms, both supervised and unsupervised, may need to be properly interpreted in order to avoid confusing correlation with causation.

A case in point is algorithms that assess recidivism risk. These algorithms construct a profile of convicted defendants and provide a score that is said to represent the likelihood that one will be a repeat offender. As with medical diagnosis solutions, these are decision-support tools. Therefore, the sentencing decision still remains with the judge. But the increasing reliance on scores that these algorithms generate may create pressure on judges, who may be perceived as 'soft on crime' if they impose a lesser sentence than is indicated by such an algorithm.

In theory, these algorithms are free of racial bias, as the defendant's race would not be included in their training set. But these training sets are based on historical data, and this data is informed by which

Data is the single most important and non-negotiable requirement for powering the use of ML. In order to take advantage of data in a sustainable way, an organisation needs a coherent data strategy.



communities and groups have been more involved with law enforcement in the past. In turn, such communities may be the least likely to have jobs, access to higher-quality education, health care, and other such variables where racial bias may have been empirically proven to exist. The result is that despite having no inherent racial bias themselves, these algorithms can make even more systematically biased decisions than the humans they have been designed to support: they 'learn' racial bias from the data.

The issue behind such bias can extend even before initial convictions are made, and not just for repeat offenders. Here the algorithms, with their base of historical data, may unwittingly end up answering the wrong question – not the likelihood of being guilty but the likelihood of being arrested.

Scenario

An ML model for improving the prediction of loan default was trained on all the historic data available on applications, approvals and defaults. The model was tested against a sample of historic data and shown to have high accuracy in predicting default. A review by an underwriter of a sample of applications and decisions was conducted before sign-off for live use.

Several months into the process, a clear pattern emerged that women were significantly over-represented among those whose loan applications were rejected. The underwriter investigated further and found a number that should have been approved. The suspicion is that the model was biased against female applicants because it was based on several decades of historic data and this training set had a lower proportion of sole applicant females. So the model was biased to reject more loan applications from this group.

For the accountant the fundamental principle of **objectivity** could be compromised in relation to issues of bias. The reference here is the avoidance of compromise of professional or business judgements because of bias, conflict of interest or undue influence of others. The accountant may have to consider

whether they have been biased in favour of assuming the outcomes are valid merely because they are supported by an ML algorithm.

On a different note, the principle of **professional behaviour** requires compliance with relevant laws and regulations. If there is evidence of systematic bias, then the organisation may be in breach of certain laws. For example, regulations such as the European Directive [2004/113/EC] modified in 2012 are targeted to disallow gender bias.

Professional accountants may face internal pressure to ignore the issue, such as if it is possible to argue a lack of evidence and that in a statistical approach the answers will be correct over a period of time. Accountants may need to play a role in, for example, guiding colleagues to reassess the model with a different emphasis on the gender variable. It will be important to maintain clear trails of communication to management, with documentation of details, responses received and, if appropriate, escalation to relevant authorities. Also key to this is a basic appreciation of the inputs and outputs associated with the model and a view on the metrics and key performance indicators. This may be required when gathering feedback and monitoring for issues, such as customer complaints, as a leading indicator of problems. A questioning approach, rooted in professional scepticism, and a growth mindset willing to grapple with new challenges, will both be important to avoid being overawed or afraid to dig deeper.

STRATEGIC VIEW OF DATA

Data is the single most important and non-negotiable requirement for powering the use of ML. In order to take advantage of data in a sustainable way, an organisation needs a coherent data strategy. In practice, this means several things.

The first is just the collection of a sufficient amount of data. Any meaningful insight with low likelihood of bias depends on having enough data across all the categories/types that may need to be considered. The amount of data and

All the hard-won benefits of an otherwise coherent data strategy can be lost through breaches that undermine the security of data in the organisation.



range across categories is what enables the algorithm to identify relationships that might not be obvious to a human.

For professional accountants, 'data' can mean more than just financial data; it might require customer data held in other departments outside finance, as well as external data. This data-gathering process can require coordination across multiple departments. The importance (and possibly difficulty) of this coordination should not be underestimated. As the survey shows, there can be many different views on who in the organisation should be responsible for owning such activities (Figure 4.1).

Once a lot of data is accessible to the organisation, the second step is to be able to store it in a way that uses a clear architecture and tagging. Otherwise there will be a lot of data, but it will be impossible to find the right data sets needed for specific algorithms: a bit like having the internet but without a good search engine to find what one needs when the requirement arises. This may become particularly important in large complex/global organisations with long-standing silos of data across different systems. Complete restructuring of the entire data eco-system may not always be realistic, but it may be necessary to accept that work is needed (even if over time, incrementally) in this area if ML is to produce the promised benefits.

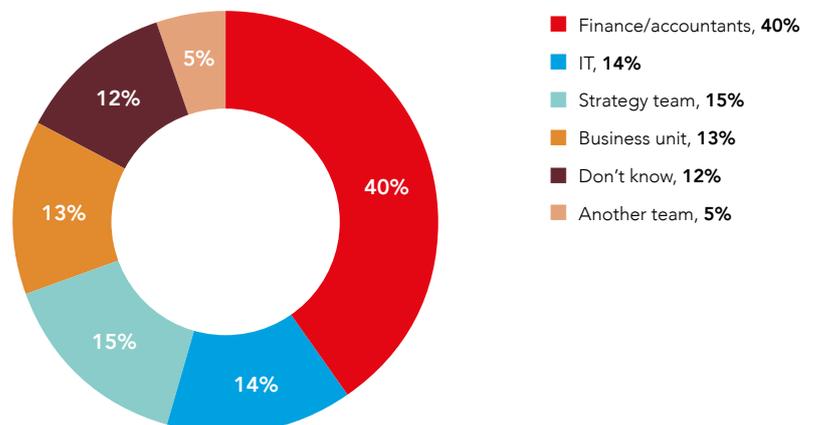
The third step is continuous maintenance of data quality. This is extremely important because algorithms do not have a standalone intelligence of their own: it is a case of 'garbage in-garbage out'. So if the quality of data feeding the model is poor, the insights will be incorrect. A strategic approach to data quality means having robust processes to enable frequent data cleansing, not just big one-off data clean-ups involving external consultants.

Once these three steps are in place to create the basic enabling environment for data, the fourth leg of the data strategy is to manage threats to this set-up. This relates to data security and the risks of data theft and cyber-security more generally. All the hard-won benefits of an otherwise coherent data strategy can be lost through breaches that undermine the security of data in the organisation.

As more data is acquired, especially if it is personal customer data, the datasets will become a target for hackers. Personal data has an increasing value to fraudsters, money launderers and organised crime. This fuels increased attempts by hackers to try and steal data or for employees to be enticed into extracting data for the purpose of sale to external agents.

Data strategy is not directly part of any ML use-case, but is an absolute requirement even before beginning to use ML. There may well be specialists who execute specific parts of the strategy, eg creating the architecture. But the overarching data

FIGURE 4.1: Who in the organisation should own the process and data gathering associated with the use of machine learning in the organisation?



Professional accountants will need to understand the increasing risks to data, the regulatory requirements and the financial consequences of breaching these.



strategy needs to be coherent, achievable and aligned to answering business questions relevant to the future direction of the organisation. This is an area where professional accountants have tremendous value to add, given that they routinely extract data from systems across the organisation. More to the point, they look beyond the technology as they seek to understand risks and control weaknesses, areas where the business does and does not achieve profitable returns, and ethical considerations about how data is used. The combination of all these perspectives is essential in creating a data strategy – it is clearly not just a technological matter.

Scenario

A large telecommunications company has developed a comprehensive customer data set to enable it to create algorithms to drive improved customer service, marketing and product development. The data set includes considerable amounts of personal data, including names and addresses, billing records, location data and call records, recent and historic.

The company has assembled a team of data scientists to use this data lake. It has been challenging to get the skills together, so the small team comprises both internal staff, including new recruits, and contractors. Because of the exploratory nature of their work and the small size, the team members all have full access to the extensive data lake. There is a risk that the data lake may become a target for hackers (external agents) or that the data may be misappropriated by one of the staff, with consequential damage to the company's reputation and customer satisfaction, and leaving it open to financial claims from customers or regulatory fines.

Professional competence and due care matter here, first of all in recognising the potential for breaches. Professional accountants will need to understand the increasing risks to data, the regulatory requirements and the financial consequences of breaching these. At a practical level, the accountant may need the ability to identify and assess the controls in place around the data lake and the background checks on employees and contractors employed into the team. From a risk-management point of view, they may also need to quantify the possible impact of a data breach on the business and its reputation.

Another area where fundamental principles may be compromised is **confidentiality**. The company may take the sensible decision that knowledge of the data lake be kept as restricted as possible, partly to use as competitive intellectual property, but also to avoid drawing attention from potential hackers.

This may involve, for example, restrictions on mentioning the data lake to external stakeholders, or internal staff outside a specific team or teams. It can also sometimes mean the use of encryption techniques to minimise personal data loss, or supplementary non-disclosure agreements signed by employees in the data lake team. Given this environment, professional accountants may have to pay particular attention to maintaining confidentiality about any information about the data lake and associated controls and processes.

More broadly, an understanding of the overall data strategy is important to ensure that one is not missing the bigger picture. Data breach is less likely when it is viewed as part of an overall understanding of how data is managed in the organisation and why: in effect, that controls provide risk management and protection, and involve more than just following pre-decided process requirements in a check-box manner.

An important principle cited by many experts for ethically sound ML solutions is that those solutions should be designed in a way that does not alter the patterns of accountability that have been established by society, culture and law.



ASSIGNING ACCOUNTABILITY

ML applications are, at their core, decision-making tools. And there is an important question at the heart of that activity. Who takes responsibility for the consequences of decisions made, the human professional accountant or the algorithm? Dealing with this clearly and consistently will be a key focus for the years ahead. The risk otherwise is that the technology will take the credit when things go well, and humans will take the blame when things go wrong – creating a no-win situation for those trying to work with the technology.

From medical diagnosis to autonomous driving to credit-risk analysis, the goal of ML is to help improve the speed and quality of how a given set of decisions gets made. An important principle cited by many experts for ethically sound ML solutions is that those solutions should be designed in a way that does not alter the patterns of accountability that have been established by society, culture and law. Here are three examples.

- If a physician uses a ML algorithm to assist in diagnosing an illness, the physician is still the one legally and ethically responsible for that diagnosis.
- If a bank purchases a ML algorithm from a software vendor to assist in determining whether or not a loan applicant is creditworthy, the bank is still the entity making the credit decision. The software vendor is not licensed as a bank, and has no charter to lend money or make credit decisions.
- When an automobile accident occurs, there are established legal precedents for when the maker of a car is liable for the accident and when the driver is responsible. The technological capability for deploying a partially or fully autonomous vehicle does not overturn these precedents; it merely raises procedural issues about how to apply them to a new technology.

In a case that was followed closely by the AI community, Germany's Federal Cartel Office (FCO) announced that it was investigating Lufthansa for unfair pricing when its fares went up over 30% shortly after the collapse of low-cost rival Air Berlin¹⁴. Lufthansa's reply was that the price hikes were not predatory because they were determined by an automated algorithm rather than by a human response to the news of Air Berlin's default.

FCO president Andreas Mundt replied publicly that Lufthansa's justification was not acceptable. 'The algorithm is beside the point', Mundt is quoted as saying, 'these algorithms aren't written by dear God in heaven. Companies can't hide behind algorithms'¹⁵.

The lesson here is that patterns of accountability need to be an explicit component of any ML application design. When seeking regulatory approval, or indeed market approval, for a new decision support tool, enterprises should clearly state their assertions about these accountability patterns and give others the opportunity to question or debate those assertions.

A relevant point here is the idea of decision support vs. decision replacement. In other words, whether a given ML application is designed to improve the quality of human decisions or replace human decision making altogether. Here are a couple of examples.

- Medical diagnosis tools are typically designed as decision-support tools. Their job is to improve the quality of a human physician's diagnosis by considering more research, data patterns, and empirical results than any one human would be capable of absorbing or considering in the course of a single diagnosis. But the task of making and communicating the diagnosis remains clearly within the remit of the human physician.

¹⁴ <https://uk.reuters.com/article/uk-lufthansa-fares/lufthansa-to-query-german-cartel-office-over-price-analysis-idUKKCN1IV1HO>

¹⁵ <https://liveandletsfly.boardingarea.com/2018/01/04/lufthansa-price-fixing/>

Without the necessary knowledge and skills to evaluate what is happening, accountants could find themselves to be ethically compromised because they agree with something 'by default'.



- ML credit-scoring algorithms are, on the other hand, more typically implemented as decision-replacement tools. The algorithm that scores loan applications within Kenya's mPesa mobile payments platform, for instance, makes credit decisions without human intervention. The entire loan application and origination process is automated and can happen in a matter of seconds.

Generally speaking, decision-support tools generate fewer ethical issues than decision-replacement tools. In the case of medical diagnosis, the algorithm's input is additive. It cannot override a physician's diagnosis; it can only broaden the base of possibilities for a physician to consider. Therefore, its impact on the overall quality of diagnosis is additive.

Replacement tools may require more concentrated focus to enable oversight of their ethical implications but there may still be strong reasons supporting them. In the above example, a significant proportion of Kenya's population is unbanked, and would not qualify for credit via traditional (non-ML) credit algorithms. So the use of ML to evaluate creditworthiness by using other sources of data means that millions more people have access to affordable credit than would have in the absence of such an approach.

The Lufthansa ruling cited about is instructive when looking at decision-replacement applications. Views will differ by jurisdiction and this report does not provide legal advice. Even so, at a conceptual level, the general direction of the legal approach in 2019 seems to be that an organisation may remain accountable for decisions made by AI, even if the humans in that organisation had no direct role in the decisions made by the relevant AI algorithm. In other words, there is accountability for the user of the application, not just its maker. This is of particular interest to professional accountants, who may need to consider whether there is sufficient oversight and review of the algorithm's decisions.

Scenario

A supermarket retailer has a successful online service and home delivery network. The retail food sector is highly competitive with narrow margins, so analysts are constantly looking for new avenues through which to increase revenues and customer spend.

The data science team has been experimenting with ML algorithms using data collected from the mobile app and, when available, other data that can be collected from a plethora of apps on devices of users who have agreed to sharing data, including location data, exercise tools and health apps. The ML algorithm is very successful in targeting clusters of customers who are highly responsive to marketing campaigns. For example, it can identify customers whose shopping habits show a periodic switch to 'healthy-eating' and calculate the right point within a 'healthy eating' phase when their resolve is at its weakest (ie when there is high likelihood of binge eating in response to diet fatigue) to target them for bulk-offers of their usual unhealthy options. This plainly has ethical implications.

Professional accountants will need to ensure they are meeting requirements on **professional competence and due care**. This could involve having to make the effort, possibly in the face of some obstruction from those in the organisation with a different agenda, to inform them of how this ML approach arrives at its remarkable results. Without the necessary knowledge and skills to evaluate what is happening, accountants could find themselves to be ethically compromised because they agree with something 'by default'. In other words, they accept the status quo because they do not have the knowledge or understanding to behave differently. In the above context, the idea of a 'black box' is often mentioned. In other words, the question of how realistic it is for professional accountants to be able to interrogate the inner workings of a complex ML algorithm.

The survey revealed prevalent views when dealing with this question. A majority of respondents reported being comfortable

While the ‘black box’ could be an issue in some cases, it should not stop professional accountants from engaging, learning and continually interrogating what is put in front of them.



trusting an algorithm from a well-known supplier (Figure 4.2). But regardless of that, a majority also had a level of unease and would seek a higher level of explainability from the algorithm than they might from a human (Figure 4.3).

There are a few points that it may be useful to consider in relation to this area of ‘explainability’.

Firstly, the accountant can achieve significant improvements in their understanding of the ML by questioning the humans associated with developing the model. They do not necessarily have to become experts in the model themselves or understand complex mathematical algorithms in full detail. They need a level of comprehension that will allow them to engage meaningfully with the situation at hand.

Secondly, the ‘black-box’ situation, where those who developed the model themselves cannot explain how it has arrived at its results, is often the result of the use of neural networks and deep learning (where there might be images or

other complex types of data involved). The complexity of the data makes it difficult to understand how the ML system has arrived at its results. By contrast, in many commercial-use cases, algorithms are dealing with relatively structured data (quantities, prices, item codes etc.) where it can be more realistically possible to explain how the algorithms operate.

Thirdly, at the deeper philosophical level, explainability is an issue that professional accountants are arguably dealing with in human decision making as well. It is not definitively possible in every case to say why someone decides something in certain way and therefore to lay out in granular fashion (as is expected for an algorithm) the precise steps that led to the view taken. In fact, experience and ‘gut reaction’ are often prized attributes, and in these instances, reasons behind decisions may not lend themselves to being deconstructed step by step. So, while the ‘black box’ could be an issue in some cases, it should not stop professional accountants from engaging, learning and continually interrogating what is put in front of them.

FIGURE 4.2: Would you be comfortable trusting a black box machine learning algorithm provided by a well-known supplier?

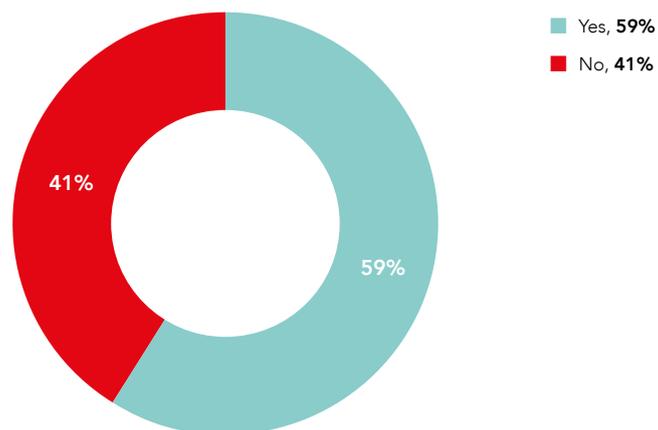
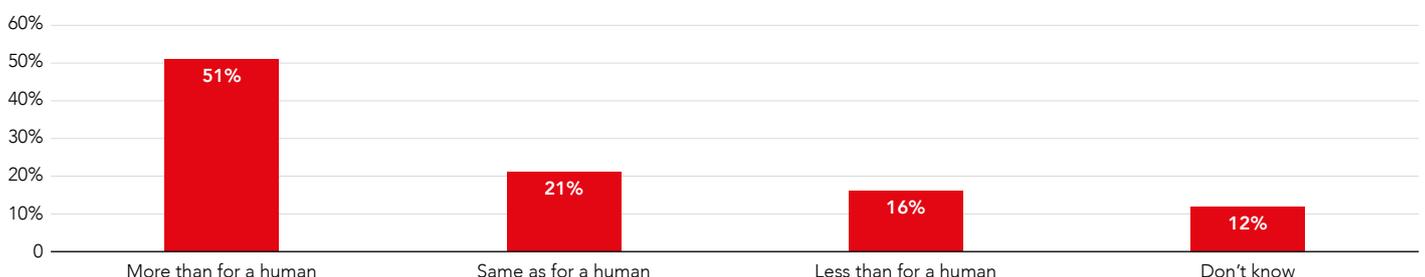


FIGURE 4.3: Relative to a human making the same decision, what level of explainability would you expect from an algorithm?



In the race to obtain this capability, businesses need to assess the quality of the applications available in the (crowded) AI/ML marketplace.



Even once competence requirements are satisfied, professional accountants could be at risk of compromise on the fundamental principle of **integrity**. This is because ethical behaviour requires being straightforward and honest in all professional and business relationships. And there may be a legitimate question to be asked as to whether the company is being straightforward and honest with its customers. Some might argue that it is not the company's job to be a 'parent' or 'moral police' to its customers; the latter make purchase decisions as adults. But questions of integrity may arise as regards the extent of personal data gathered and its use in targeting specific vulnerabilities of the customers in question.

Ultimately, the answers will depend on the specifics of the case, but in dealing with these questions professional accountants will need to ensure that the principle of integrity guides their decision making, in the face of clear commercial pressures associated with a certain course of action. This will be particularly the case in situations which satisfy the minimum threshold of legality, but not the higher bar of ethical behaviour.

LOOKING BEYOND THE 'HYPE'

AI has become a 'buzzword' in recent years, and ML, as part of AI, has often attracted similar attention. Undoubtedly, there are real benefits from using ML, but unrealistic expectations, and the vested interests of those selling this technology mean that there is also a real risk of the misrepresentation of what is on offer.

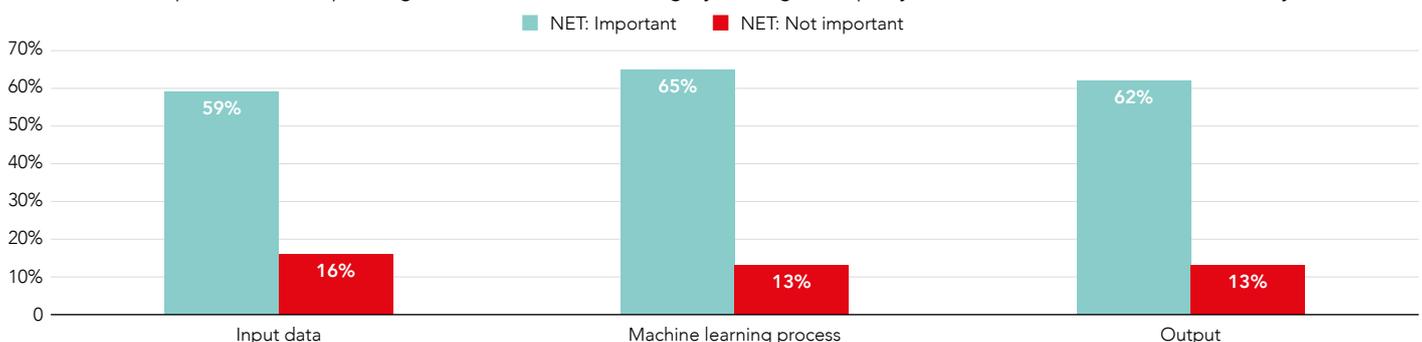
The publicity around AI can make it seem as if many companies are adopting AI and ML into their strategic plans and that every software company either has or is developing AI/ML capability. This creates an expectation from investors that all companies should be using AI/ML and an anticipation of enhanced value creation. One study identified that the term 'artificial intelligence', was getting almost four times as many mentions on earnings calls as Big Data when looked at over a ten year period¹⁶.

In the race to obtain this capability, businesses need to assess the quality of the applications available in the (crowded) AI/ML marketplace. They also need to assess whether many of the claims being made by software vendors and consultants are backed by adequate evidence. Fear of falling behind should not drive investment decisions without proper evaluation.

While some applications will prove their worth, there could also be instances of companies that claim to have an AI product, while actually using humans to simulate AI. This may be done to fulfil product demand while, in parallel, the equivalent software capability is being developed. In other instances, humans may be used to 'train' and the product is real but may not be fully developed.

Respondents in the survey were supportive of the need for third-party assurance. This can be a useful way of getting an independent view of assertions made in general terms as well performance of the ML application in specific instances (Figure 4.4).

FIGURE 4.4: Importance of improving trust in machine learning by having third-party assurance of the elements of the system



Note: 1–5 scale; NET Important is sum of 4, 5; NET Not important is sum of 1,2

16 <https://www.analyticsindiamag.com/artificial-intelligence-increasingly-mentioned-earning-calls/>

Asymmetry of information about AI and the consequent potential for deceiving the public (which includes fee-paying customers) with false promises could create additional scamming opportunities.



Scenario

The chief executive officer (CEO) at a mid-size independent financial adviser has recently become more aware and convinced of the trend in AI and does not want to miss out. The company has had a presentation from a technology vendor, giving details of a product that the CEO believes will 'revolutionise' the firm's efficiency.

The software company is a start-up that has developed ML that can analyse and respond to questions from clients. While conceptually this not new, the vendor claims that its use of ML really builds on earlier less sophisticated attempts in this area – and simulates interactions with clients to a level not seen previously. The CEO wants to be at the cutting edge and is about to sign a significant-value contract to start using the company's software.

Professional accountants associated in any way with this procurement decision may need to consider the principle of **integrity**. It may not be easy to challenge decisions that have senior-level backing but it is necessary to be honest in stating facts as they are perceived. And getting to the facts will inevitably require developing some understanding of the product, which relates to the fundamental principle of **professional competence and due care**.

Being straightforward with internal stakeholders involves, among other things, clarifying what can be reliably verified and what cannot, as part of ensuring that there is a clear common understanding of the opportunities and risks. For example, the start-up may have a promising technology offer, but can its claims be verified by any independent third parties? Is it possible to obtain assurance in any way that the software is fully automated without any human involvement or manual workarounds? Does the start-up have a stable funding source or is it likely to be under pressure to gain business or over-sell its products owing to funding pressures? Is there confidence that data is being reliably

handled, with protection, privacy and personal data considerations being adequately managed?

Any insights gained through these enquiries and the associated professional scepticism-led thinking, must be cross-referenced with statements made by the start-up, particularly publicly available product or marketing material. Ethical guidelines state that a professional accountant should disassociate themselves from false or misleading information.

ACTING IN THE PUBLIC INTEREST

Technology can raise wider, more universal questions about public good and public value. Professional accountants, particularly if they are operating within a commercial environment, may find themselves being pulled in different directions as a result.

In a sense, this cuts across all the ethical considerations mentioned so far in this section. But it is worth specifically referencing because it relates to the power of ML to create unintended consequences that can compromise innocent stakeholders, whether or not they are directly connected to the technology being employed.

When dealing with bias, the public interest requires avoidance of the enabling of systematic discrimination against marginalised groups by an algorithm reflecting historical bias. On data strategy there could be a public interest consideration in how the medical data sets of large groups of people are used or shared, particularly if centralised or publicly funded systems mean this data is easily available in large quantities.

Similarly, on assigning accountability, there may be a wider question beyond individual or company liability, with autonomous vehicles being the classic case example. Asymmetry of information about AI and the consequent potential for deceiving the public (which includes fee-paying customers) with false promises could create additional scamming opportunities.

More broadly, defending the public interest requires an ability to avoid seeing the fundamental principles as a basic minimum that is required for compliance.



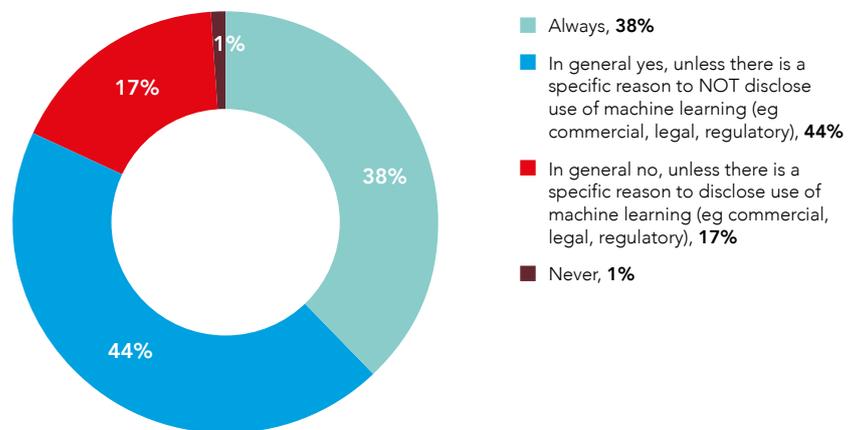
Disclosure and the level of transparency can have implications for the wider public interest. In this regard, respondents in the survey generally tended to favour disclosing where ML was being used in the background, though there were some nuances, as demonstrated in Figure 4.5.

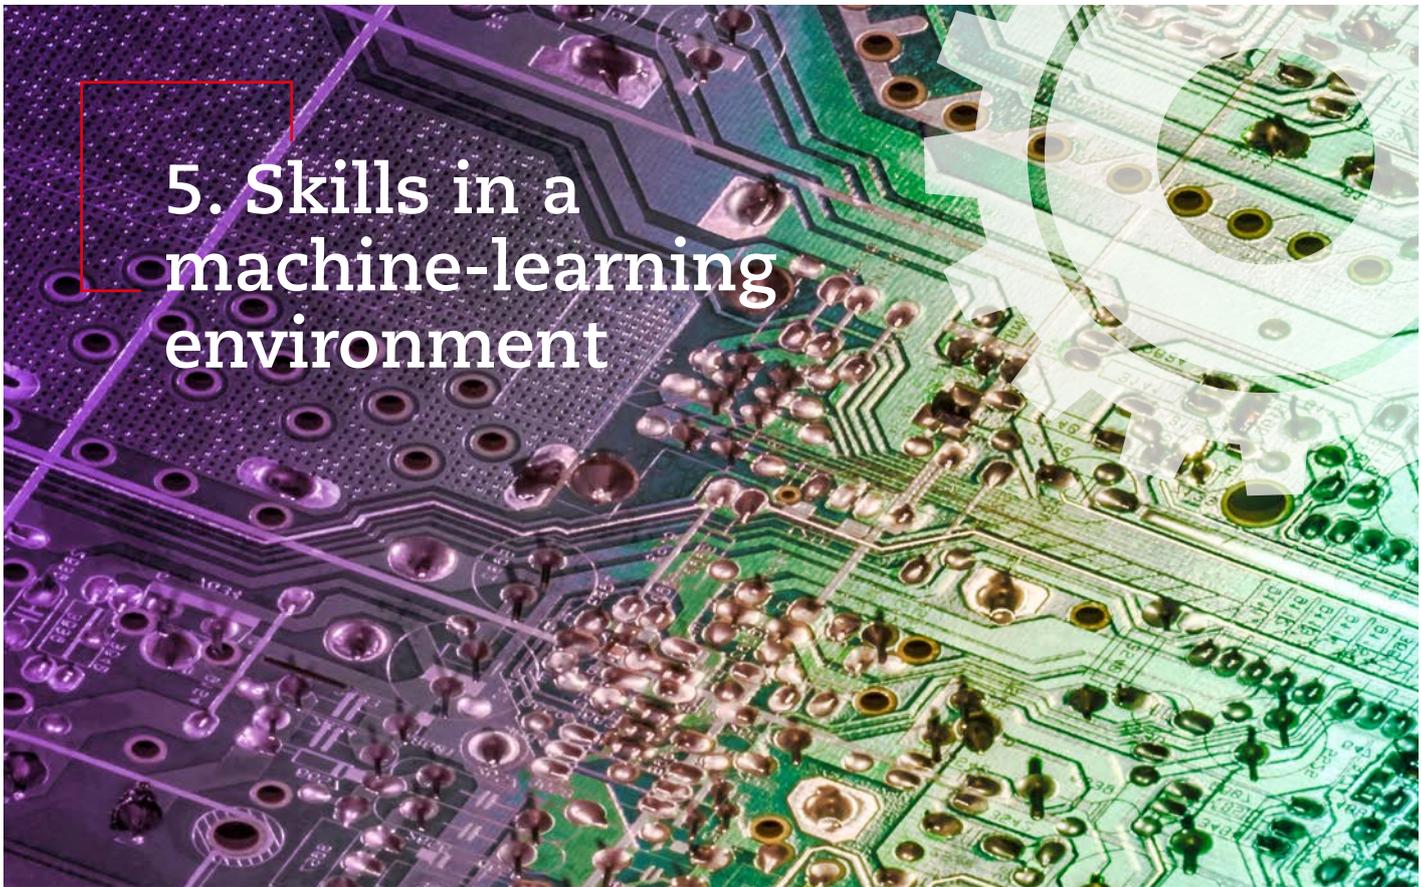
For professional accountants, the fundamental principles established by IESBA provide guidance when faced with ethical questions. These principles must be interpreted for the situation at hand – something essential when dealing with new technologies and previously unseen scenarios. For instance, in certain cases, there may be a need to consider involving external stakeholders. In this scenario there may be a relevant application of the NOCLAR (Non-compliance with Laws and Regulations) requirements set out by IESBA.

The IESBA framework is there to guide professional accountants on an approach based on involving management/internal escalation, with further action via an appropriate authority outside the entity if internal consultation fails to result in an appropriate response. The point here is that the fundamental principles are not to be applied via a check-box approach. In the example above, NOCLAR might mean, for example, that confidentiality principles must be set aside because there is a wider case for sharing sensitive data with relevant external authorities, in the public interest.

More broadly, defending the public interest requires an ability to avoid seeing the fundamental principles as a basic minimum that is required for compliance. The fundamental principles must permeate all aspects of how professional accountants perform their roles.

FIGURE 4.5: Do you think that organisations should highlight when a decision has been made by a machine learning algorithm?





5. Skills in a machine-learning environment

Professional accountants need an appreciation of how ML can affect their organisations. Use of ML could affect the way they track and influence how the organisation creates value, require changes to risk and control mechanisms, or perhaps create ethical considerations.

The survey results reflected a degree of scepticism about allowing algorithms to take the lead in areas requiring complex judgement and interpretation. Depending on the type of decision needed, only between one-tenth and one-third of respondents thought that a machine algorithm could lead or be given full reliance when applying such judgement (Figure 5.1).

Technology is clearly a significant factor in influencing the future direction of the accountancy profession. But technology is

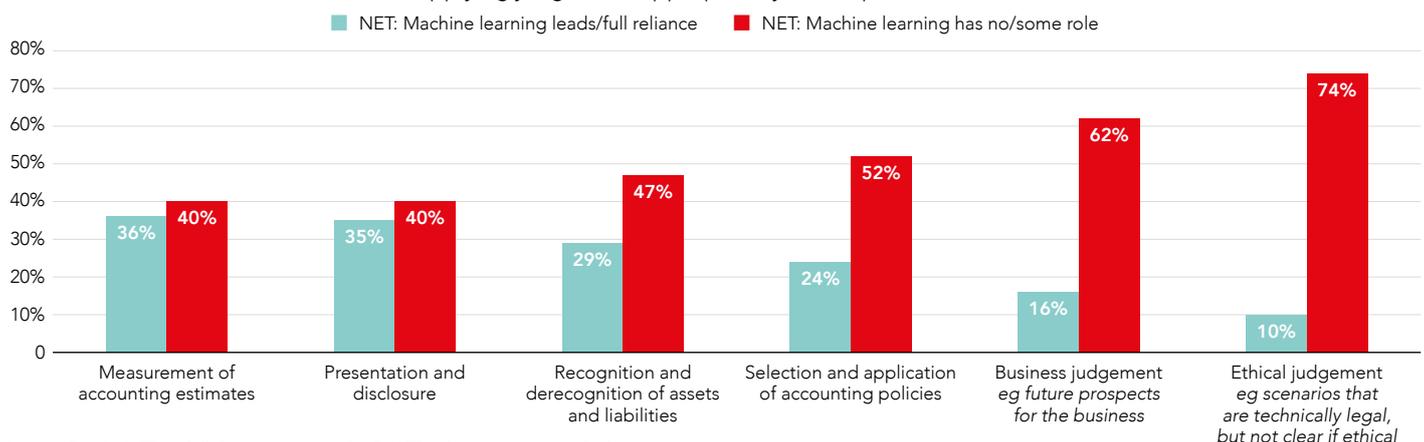
an enabler and is part of a wider pool of factors, including people and processes. ACCA's skills framework (Figure 5.2), which guides its qualification, expresses this idea with the Digital Quotient (DQ) being one of a series of professional quotients that interlink to create a 'future-proofed' professional accountant for a digital age.

Technology, even the most sophisticated technology such as AI, still currently struggles to replicate the full contextual understanding and integrated thinking of which humans are capable. It does not

seem widely accepted that the element of human oversight can be done away with completely or that the technology can take into account human factors, whether in building client relationships or leading successful teams.

ACCA's work on the EQ strongly demonstrated the need for emotional intelligence-related competencies in a digital age (ACCA 2018). In fact, DQ and EQ are best seen as skill sets needed in combination for either to be really effective for professional accountants in a digital age.

FIGURE 5.1: What is the role of ML in applying judgement appropriately in complex scenarios?



Professional accountants need to remain engaged, as AI and its component parts evolve. While some may want to delve deeper into the mechanics behind AI, this may be less appropriate for other roles.



Even beyond areas such as leadership, core technical activities can require judgement and interpretation that draw on multiple abilities, and where future choices may not be fully informed by patterns seen in the past.

The role of an accountant involves tasks in many familiar areas such as audit, corporate reporting, or taxation. But in the effective execution of these tasks, the professional accountant is expected to have an opinion, and to be accountable for that opinion and liable for its consequences.

The legal system, societal and cultural values are founded on the principle that only individual people or corporate entities can be held accountable. And there does not seem, at this stage, to be any broad consensus that this principle and associated legally backed accountability can be outsourced to AI algorithms.

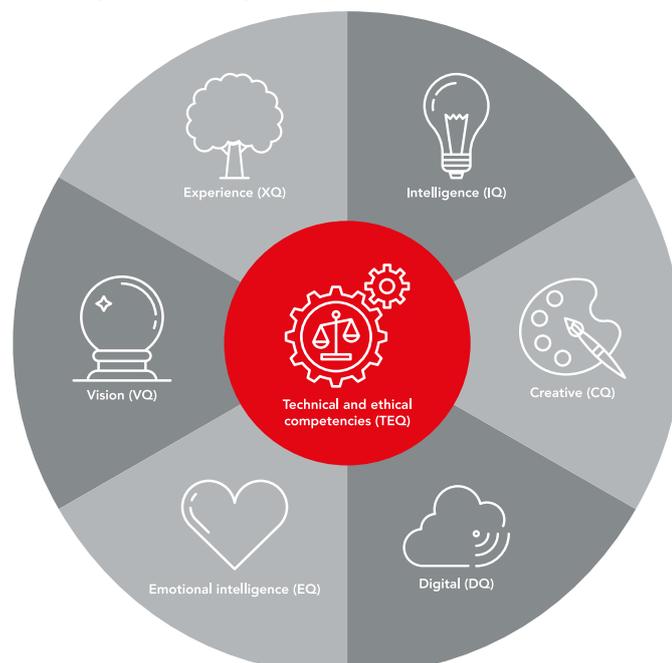
Almost all ML tools, even purely unsupervised learning algorithms, depend on the domain knowledge and expertise of the people who develop them. The task of 'binning', or turning raw data into discrete measures that ML algorithms can use to derive inferences from their training sets, is still as much an art as a science, and typically requires the collaboration of highly trained technologists, data scientists, and industry experts to be successful.

This is not to say that AI can be ignored. It is certainly the case that some new skills could become more prevalent as we look ahead to a world with increased use of AI, for example the data scientist. As accounting methods come to rely increasingly upon ML, there will be a need to get the most out of ML techniques, and to be able to bridge the worlds of accountancy and data and statistics led insight. Data scientists can advise in areas such as evaluating which ML tool to use in which circumstance, and whether there could be ethical issues inherent in the ML application.

Professional accountants need to remain engaged, as AI and its component parts evolve. While some may want to delve deeper into the mechanics behind AI, this may be less appropriate for other roles. At a minimum, all finance professionals should maintain awareness of how AI is evolving and be alert to how the developing capability could overlap with their role. This could range from how their own role can incorporate AI to be more efficient and serve clients or employers better, through to identifying how AI is being adopted and used by these clients and employers. This is also why this report has introduced some specifics from the market, by way of information in the section on applications.

Appendix 3 suggests some avenues for maintaining awareness of evolving developments in this area.

FIGURE 5.2: ACCA professional quotients





Conclusion

There is a continuous maturing of the discussion around ML, and AI more generally. Some will embrace it. Others will fear it. But only the reckless will avoid finding out more about it.

Now is a good time to start building greater knowledge and awareness in this area. The technology has moved beyond unrealistic fantasy to real business applications. At the same time, processes and approaches in the AI space have not yet been finalised. So how these technologies are developed, challenged and refined in the years to come could have lasting impact. It is vital to ensure that any adoption of ML and associated technologies is valuable in the long-term without producing unacceptable side-effects for organisations and society.



Appendix 1



The fundamental principles are described in the Code established by IESBA and are as follows.

- a. **Integrity** – to be straightforward and honest in all professional and business relationships.
 - b. **Objectivity** – not to compromise professional or business judgements because of bias, conflict of interest or undue influence of others.
 - c. **Professional competence and due care:**
 - i. to attain and maintain professional knowledge and skill at the level required to ensure that a client or employing organisation receives competent professional service, based on current technical and professional standards and relevant legislation, and
 - ii. to act diligently and in accordance with applicable technical and professional standards.
 - d. **Confidentiality** – to respect the confidentiality of information acquired as a result of professional and business relationships.
 - e. **Professional behaviour** – to comply with relevant laws and regulations and avoid any conduct that the professional accountant knows or should know might discredit the profession.
- 

Appendix 2

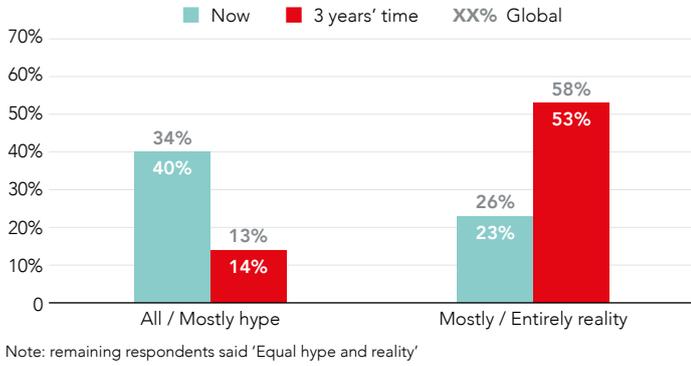


Country snapshots

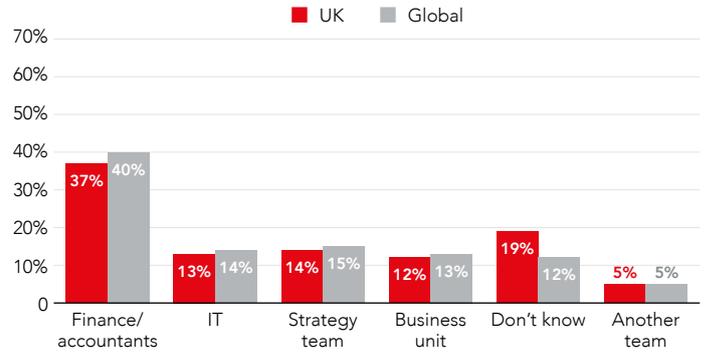
1.	 UK	40
2.	 China	41
3.	 Malaysia	42
4.	 Singapore	43
5.	 United Arab Emirates (UAE)	44
6.	 Ireland	45
7.	 Pakistan	46



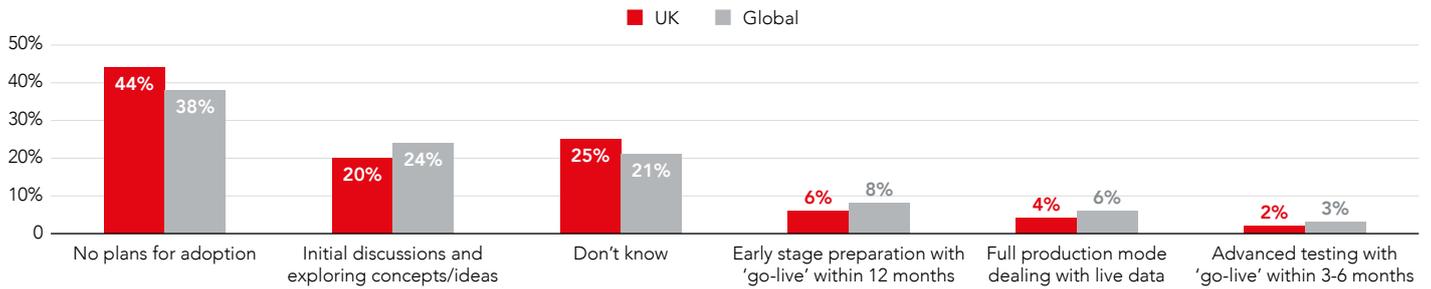
UK 1: AI as hype or reality



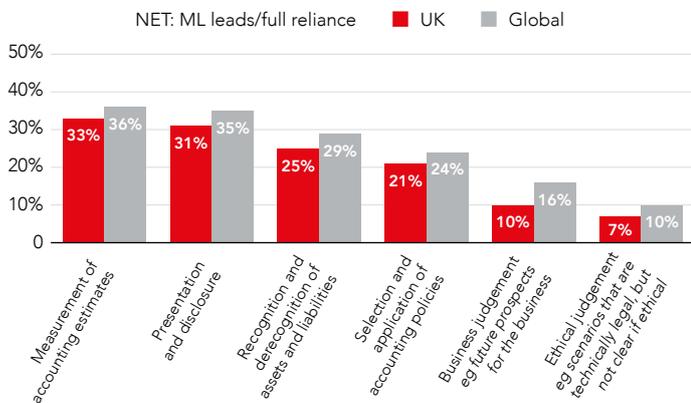
UK 2: Who in the organisation should own the process and data gathering associated with the use of ML?



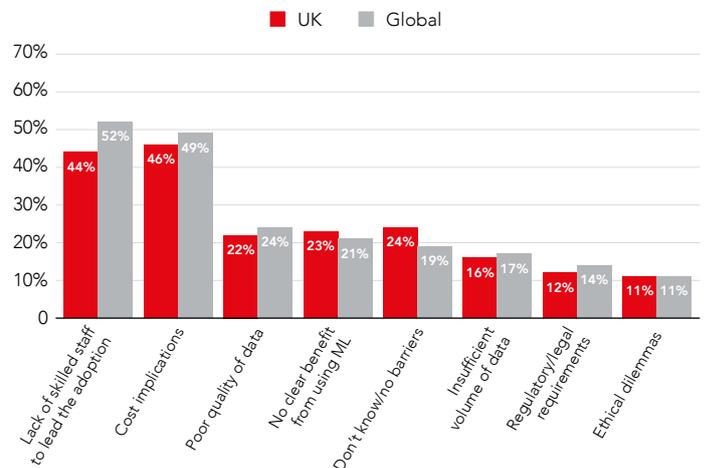
UK 3: Machine learning (ML) adoption in organisation



UK 4: Role of ML in applying judgement appropriately in complex scenarios?

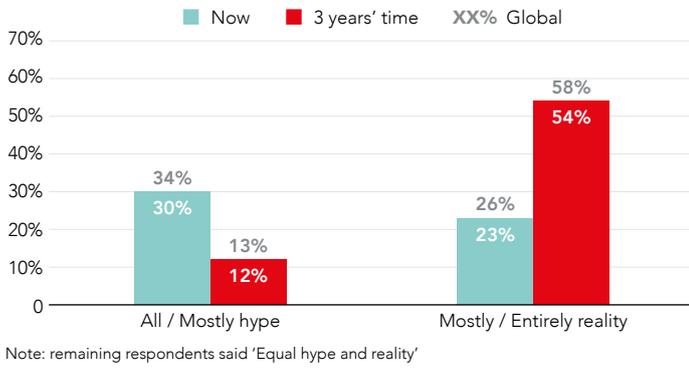


UK 5: Please tell us about the main barriers to using ML in the organisation more generally

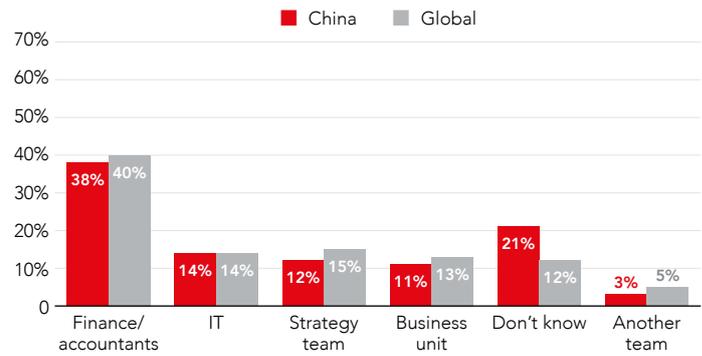


China

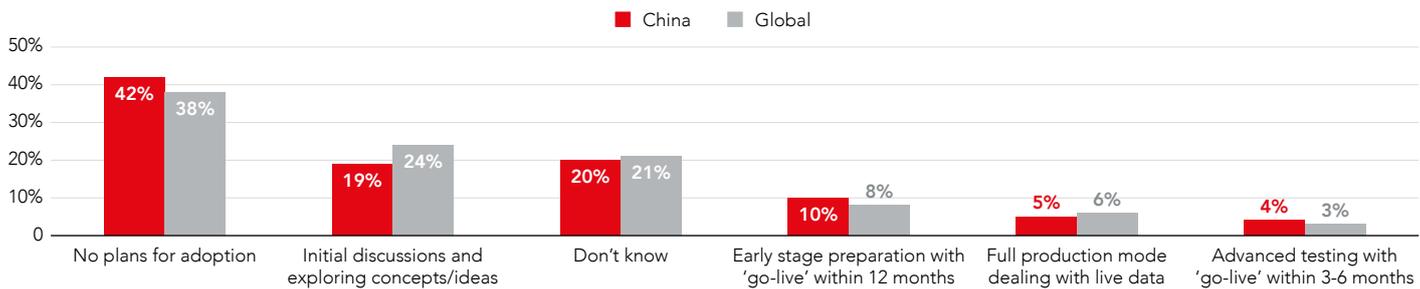
CHINA 1: AI as hype or reality



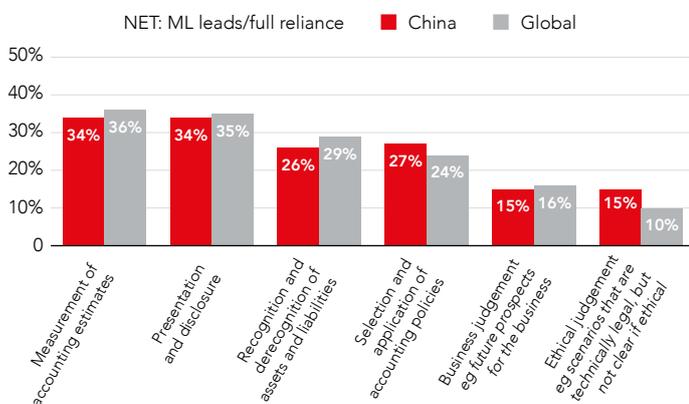
CHINA 2: Who in the organisation should own the process and data gathering associated with the use of ML?



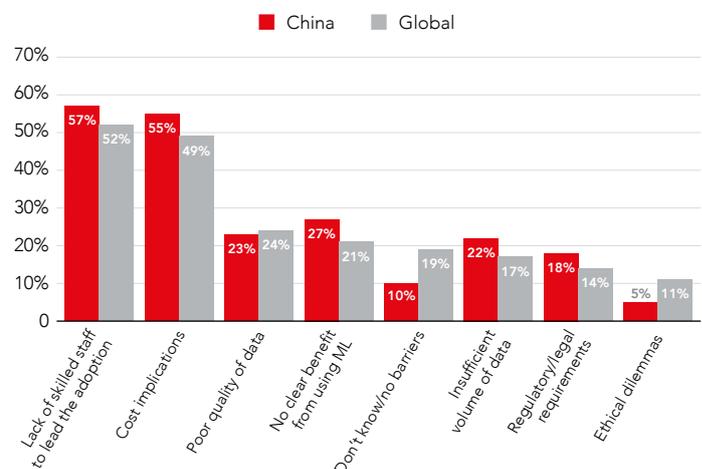
CHINA 3: Machine learning (ML) adoption in organisation



CHINA 4: Role of ML in applying judgement appropriately in complex scenarios?

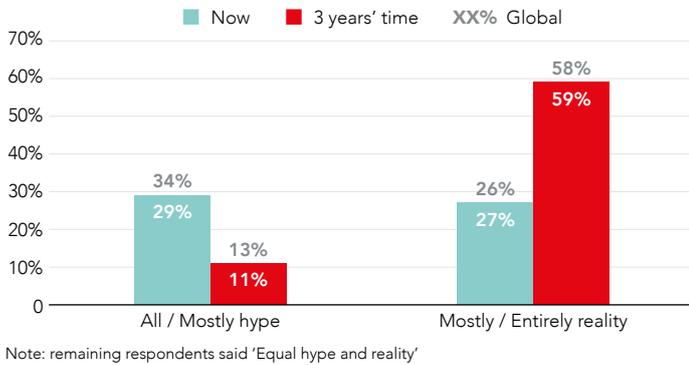


CHINA 5: Please tell us about the main barriers to using ML in the organisation more generally

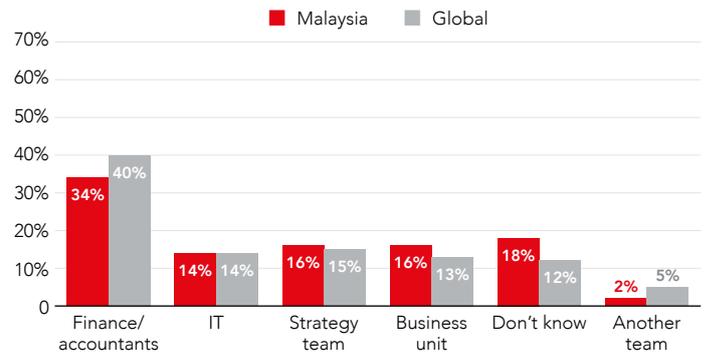


Malaysia

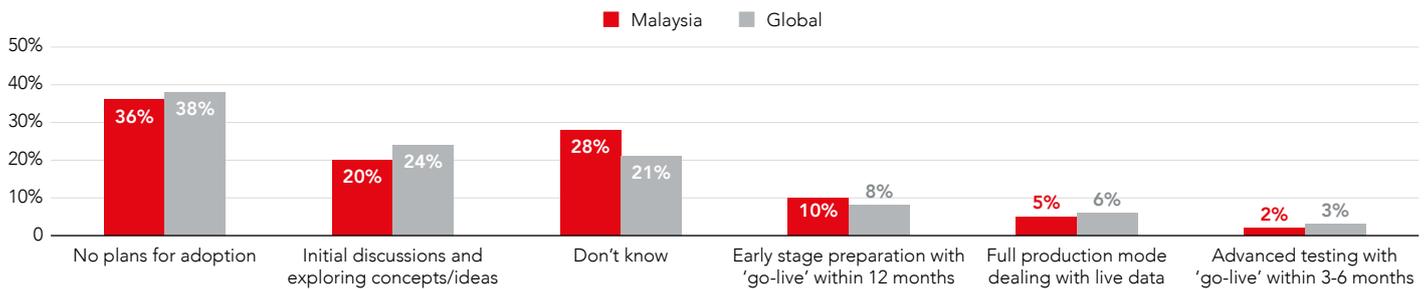
MALAYSIA 1: AI as hype or reality



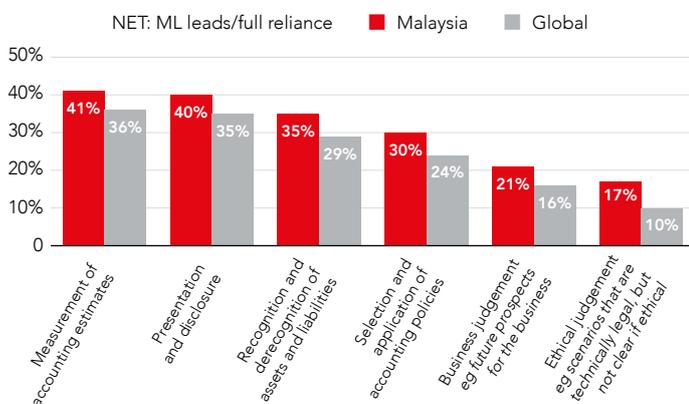
MALAYSIA 2: Who in the organisation should own the process and data gathering associated with the use of ML?



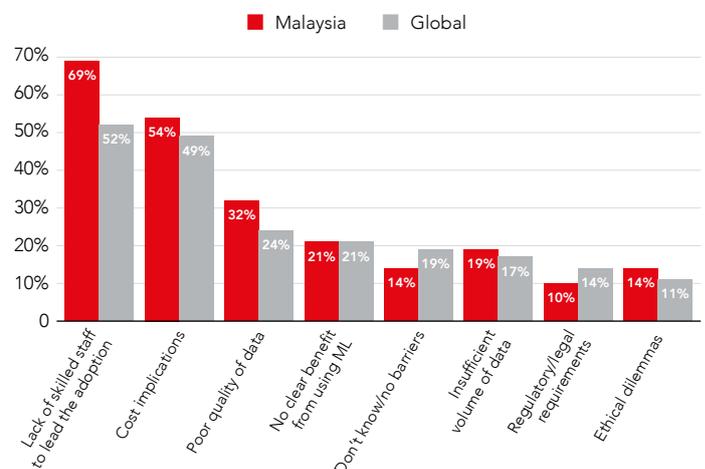
MALAYSIA 3: Machine learning (ML) adoption in organisation



MALAYSIA 4: Role of ML in applying judgement appropriately in complex scenarios?

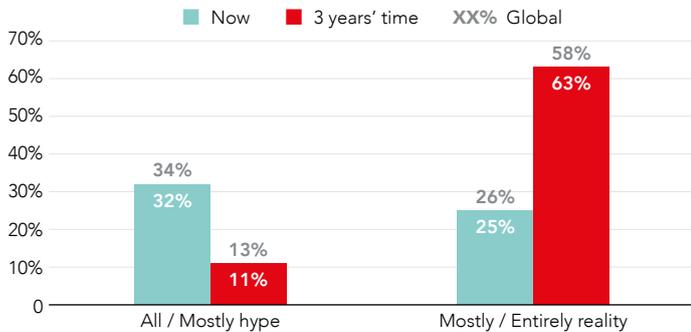


MALAYSIA 5: Please tell us about the main barriers to using ML in the organisation more generally



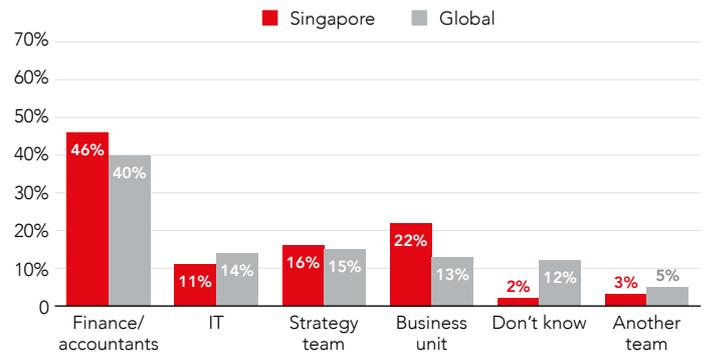
Singapore

SINGAPORE 1: AI as hype or reality

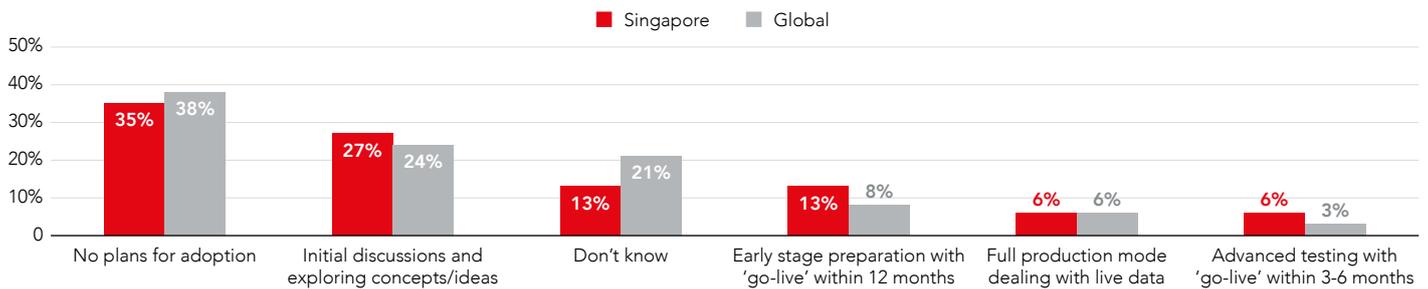


Note: remaining respondents said 'Equal hype and reality'

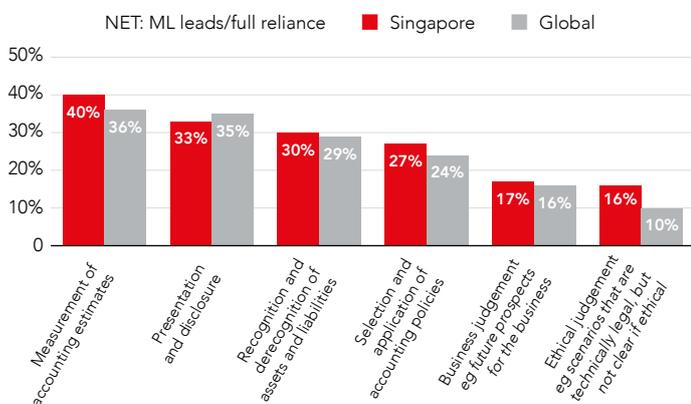
SINGAPORE 2: Who in the organisation should own the process and data gathering associated with the use of ML?



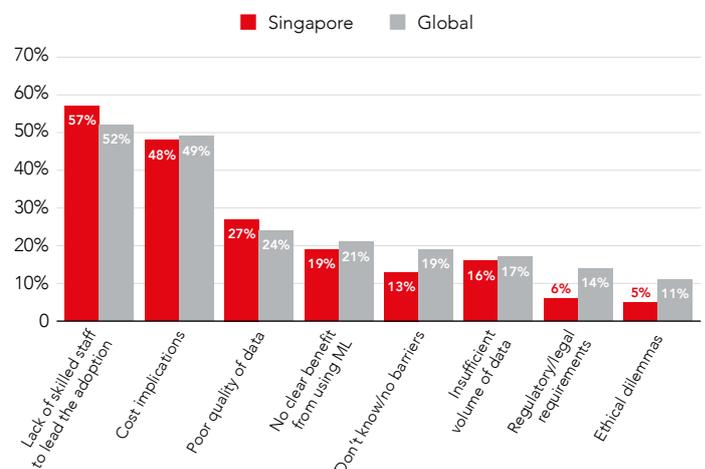
SINGAPORE 3: Machine learning (ML) adoption in organisation



SINGAPORE 4: Role of ML in applying judgement appropriately in complex scenarios?

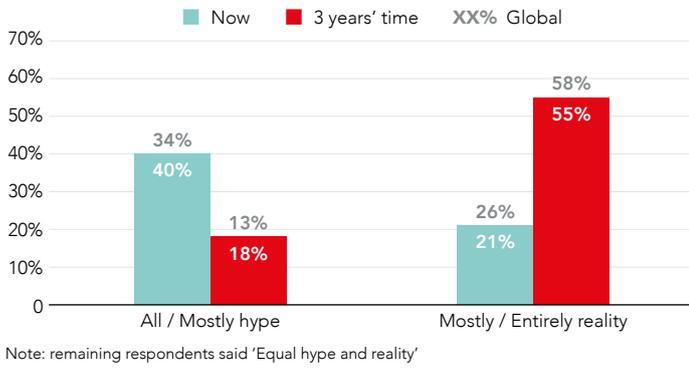


SINGAPORE 5: Please tell us about the main barriers to using ML in the organisation more generally

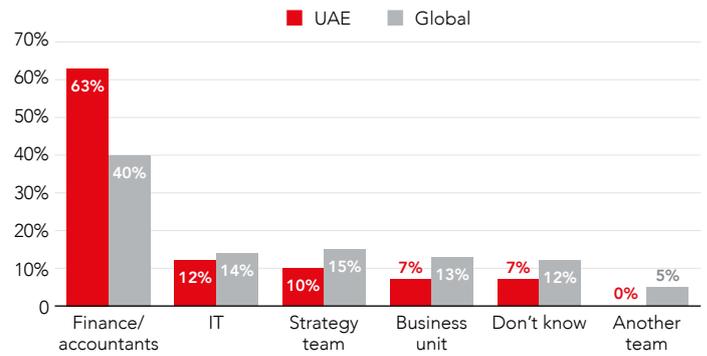


United Arab Emirates (UAE)

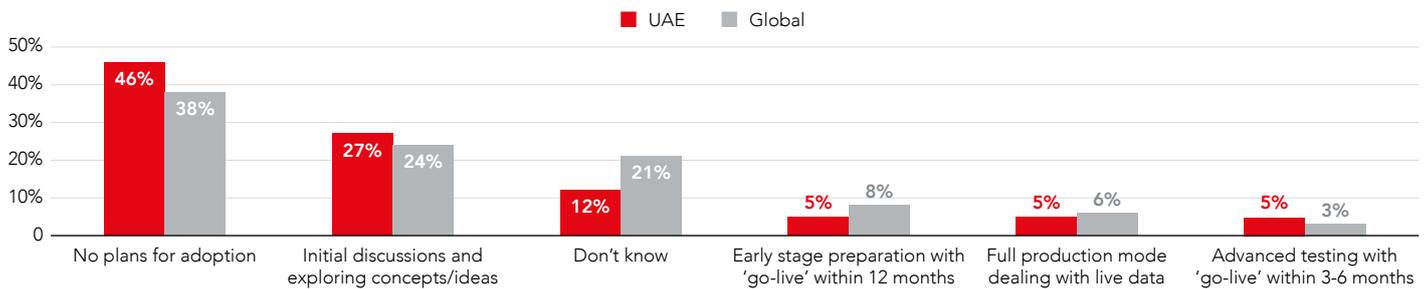
UAE 1: AI as hype or reality



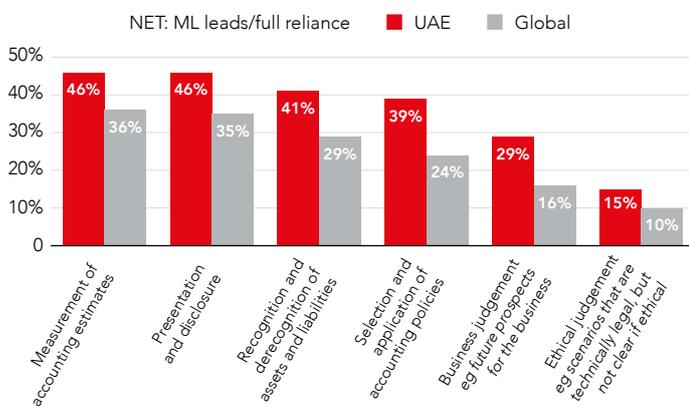
UAE 2: Who in the organisation should own the process and data gathering associated with the use of ML?



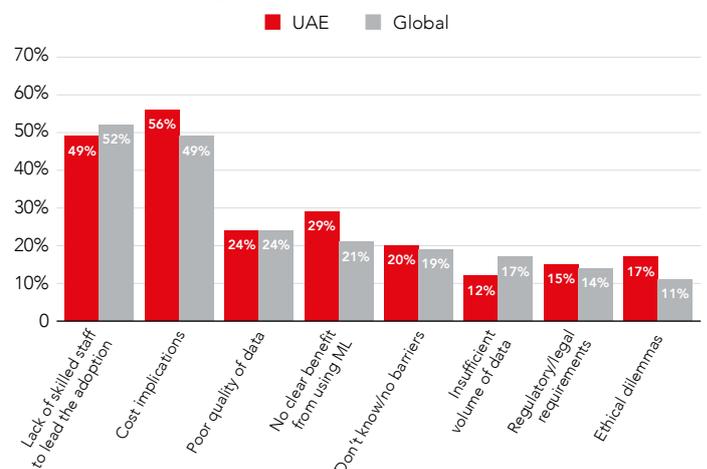
UAE 3: Machine learning (ML) adoption in organisation



UAE 4: Role of ML in applying judgement appropriately in complex scenarios?

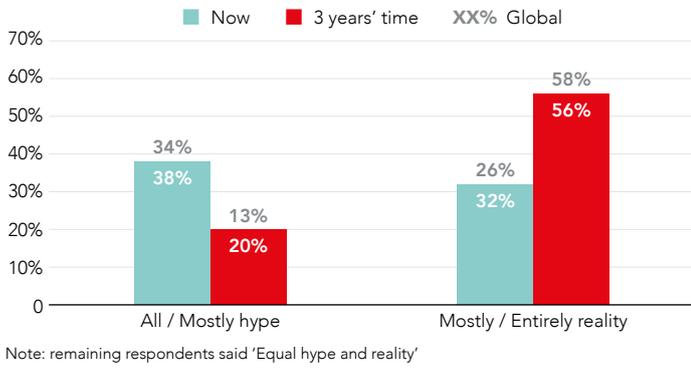


UAE 5: Please tell us about the main barriers to using ML in the organisation more generally

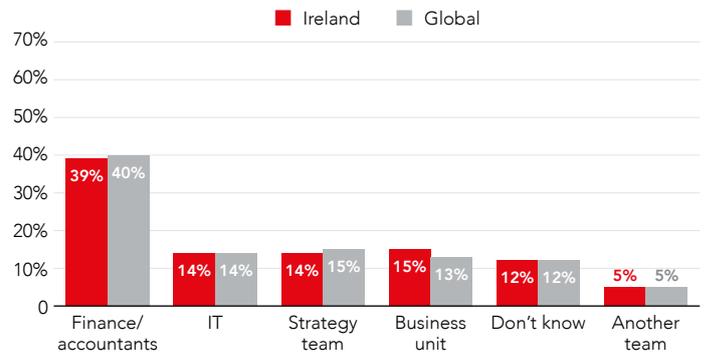


Ireland

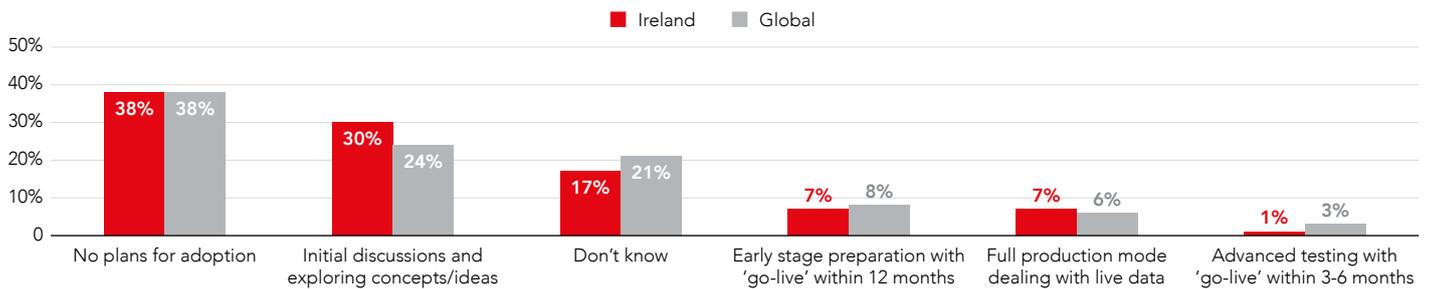
IRELAND 1: AI as hype or reality



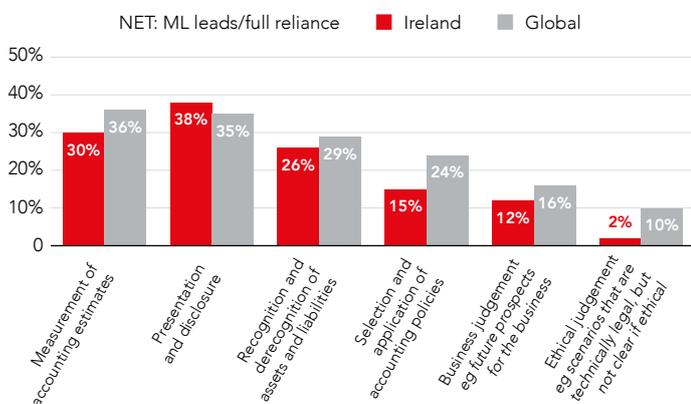
IRELAND 2: Who in the organisation should own the process and data gathering associated with the use of ML?



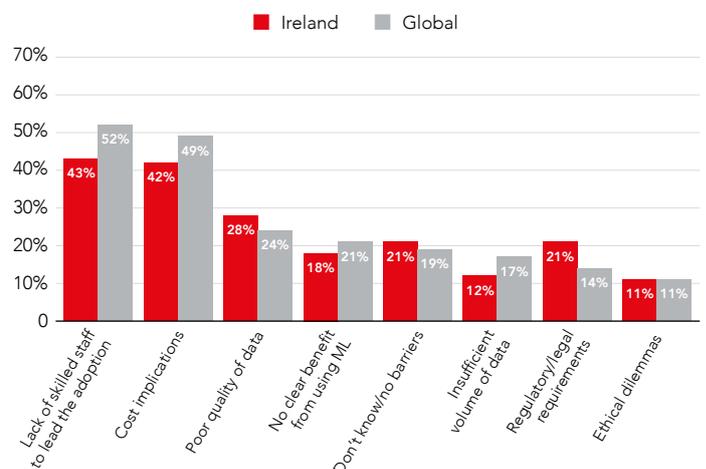
IRELAND 3: Machine learning (ML) adoption in organisation



IRELAND 4: Role of ML in applying judgement appropriately in complex scenarios?

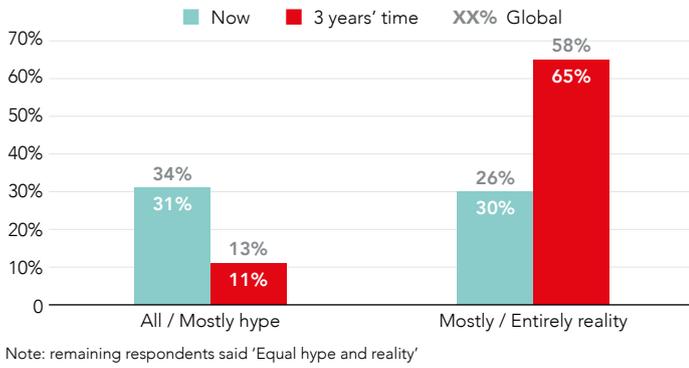


IRELAND 5: Please tell us about the main barriers to using ML in the organisation more generally

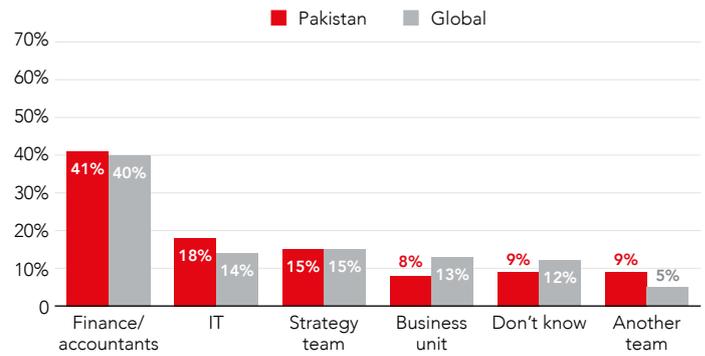


Pakistan

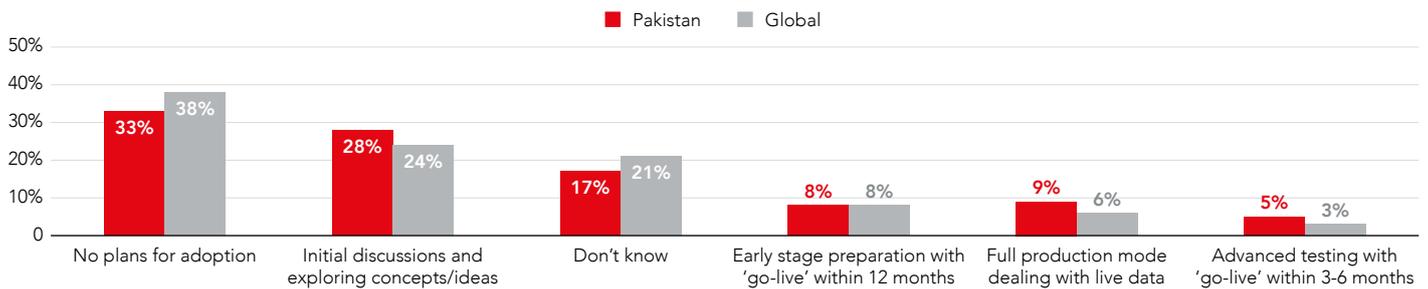
PAKISTAN 1: AI as hype or reality



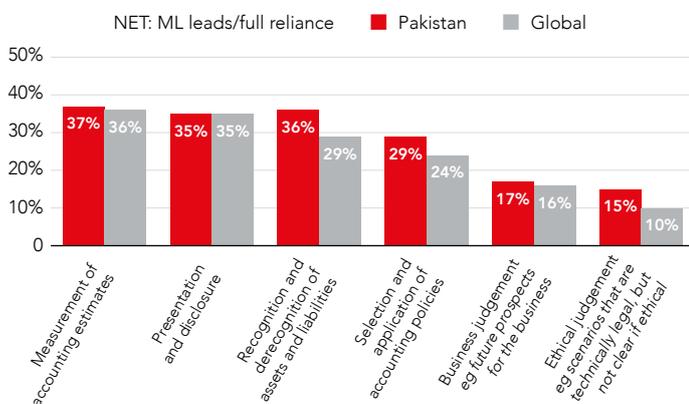
PAKISTAN 2: Who in the organisation should own the process and data gathering associated with the use of ML?



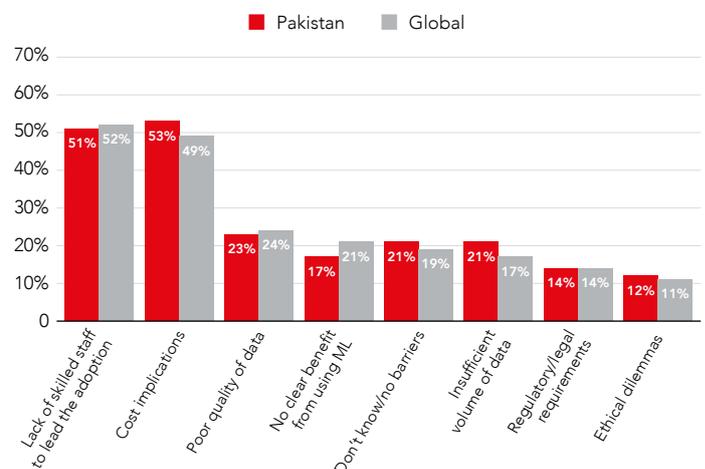
PAKISTAN 3: Machine learning (ML) adoption in organisation



PAKISTAN 4: Role of ML in applying judgement appropriately in complex scenarios?



PAKISTAN 5: Please tell us about the main barriers to using ML in the organisation more generally



Appendix 3

A starting point for keeping abreast of basic developments is to look at how consumer applications and capabilities incorporate AI techniques and functions. This can range from smartphone ‘apps’, software, Web applications, consumer devices, IoT devices, to autonomous cars and drones etc.

Technology publications are a good source of practical details of these new capabilities. TED talks can be a useful way of gaining insight into a range of topics, including technologies such as AI. A vast range of topics are covered and available online at the TED website, all no longer than 18 minutes.

MeetUp groups are an easy way of gaining insight into technology (including AI) and of expanding one’s understanding and connecting with people involved directly in AI development and use. Many organisations hold ‘hackathons’ for speedy creation and testing of ideas for using new technologies in an innovative way. These work best where technology and business professionals collaborate, and are an ideal opportunity for finance professionals to get involved, find out more, and show where they can add business value.

AI is an emotive topic that raises concerns over data access and sharing, as well as for jobs, employment and training. Accountants should stay aware of media coverage of AI and its impact on jobs; government policy on employment and training; and research and surveys of consumer opinion. Accountants need to be aware of the perceptions of social

impact from AI. Forums such as AI4People provide valuable insight to the direction of AI and how to address the social and ethical challenges.

Case studies from software vendors and consultancies are a (admittedly biased) source of information for how these developments are being deployed but they do give examples of the benefits being achieved. Many software vendors will also provide training on AI capabilities.

Technology analysts (such as Gartner, Forrester, IDC) publish reports and survey results, and run events on topics such as AI and ML with a bias towards identifying and explaining the practical capabilities available from vendors.

Organisations such as Digital Catapult, created to encourage collaboration between business and technologists, are another source of information about applications, including AI and ML, especially those from early-stage companies that may not yet have broader visibility.

There are independent course providers that provide specific AI and ML education: eg EdX.

17 [\[www.ted.com/talks\]](http://www.ted.com/talks)

18 [\[www.meetup.com\]](http://www.meetup.com)

19 <http://www.eismd.eu/ai4people/>

20 <https://www.digicatapult.org.uk/>

21 <https://www.edx.org/course/machine-learning-columbia-csmm-102x-4> <https://www.edx.org/course/artificial-intelligence-ai-columbia-csmm-101x-4>

References

ACCA (2017), *Professional accountants – the future: Ethics and trust in a digital age* <https://www.accaglobal.com/content/dam/ACCA_Global/Technical/Future/pi-ethics-trust-digital-age.pdf>, accessed 20 March 2019.

ACCA (2018), *Professional accountants: the future: Emotional quotient in a digital age* <https://www.accaglobal.com/content/dam/ACCA_Global/professional-insights/Emotional_quotient/pi-emotional-quotient-digital-age.pdf>, accessed 18 March 2019.

Datamaran (2018), *Global Insights Report 2018: The Rise of ESG Regulations* <<https://www.datamaran.com/wp-content/uploads/2018/10/GIR-2018.pdf>>, accessed 20 March 2019.

Vaidyanathan, N. (2017), *Divided we fall distributed we stand. The professional accountant's guide to distributed ledgers and blockchain* <https://www.accaglobal.com/content/dam/ACCA_Global/Technical/Future/Divided%20we%20fall%2C%20distributed%20we%20stand%20-%20The%20professional%20accountant%E2%80%99s%20guide%20to%20distributed%20ledgers%20and%20blockchain.pdf>, accessed 20 March 2019.

WEF (World Economic Forum), Marsh and McLennan Companies and Zurich Insurance Group (2019), *The Global Risks Report 2019*, 14th edition <http://www3.weforum.org/docs/WEF_Global_Risks_Report_2019.pdf>, accessed 20 March 2019.

Yao, M., Zhou, A. and Jia, M. (2018) *Applied Artificial Intelligence: A Handbook for Business Leaders*. (USA: TopBots).



