



Artificial Intelligence and corporate reporting

How does it measure up?

January 2019



Financial Reporting Council

Contents

Introduction	1
Quick read	2
Section 1 Artificial Intelligence basics	4
What is AI?	5
AI in the real world	6
Section 2 Artificial Intelligence and corporate reporting	8
The nature of corporate reporting and reporting process challenges	9
AI and the production of corporate reporting	10
AI and the distribution of corporate reporting	13
AI and the consumption of corporate reporting	15
Section 3 Conclusions and actions	17
AI Reporting: actions and questions	18
Next steps?	19
Section 4 Appendices	20
AI in more detail	21
Links	25

Introduction

The Financial Reporting Lab's [Digital Present](#) project showed that technology provided new ways for companies to interact with their stakeholders, but that many of these new mechanisms were not delivering on their promise of transformative change. The [Digital Future](#) project continues these themes and aims to understand how new and developing technologies can be used to disseminate company communications in the most efficient manner. This report, on Artificial Intelligence (AI), is the third in our series of technology deep-dives.

Why AI?

The Lab undertook a survey in 2016 which asked what technologies would be important for the future of corporate reporting. Respondents identified AI as important for the production and consumption of corporate reporting. This report considers the case for AI.

How should I read this report?

This report begins with some background on AI, what it is and how it has developed. We then cover some of the potential use cases that AI has for corporate reporting. For each use case we consider what is currently being achieved and what might be achieved in the future. At the end of the report we recommend some actions for all those involved in corporate reporting. The report's appendices include some basics on different types of AI and related technologies, and links to useful material. While this report is part of a series it can also be read as a standalone document, although it does reference all the outreach and discussions undertaken throughout the wider project.

We would like to specifically thank all those who took part in interviews and roundtables for this project. The roundtables brought together technology experts, accounting, auditing, regulatory and other participants. The roundtables and interviews were used as a key input into this report.

Quick read

Introduction



AI has long featured in future predictions about the imminent evolution of the workplace. But that future does not often consider the impact that AI might have on corporate reporting. This report, on AI in corporate reporting, is the third in our series of technology deep-dives (following on from XBRL and Blockchain). In this, and each deep-dive report, we use our [digital reporting framework](#) to explore how different technologies might impact corporate reporting production, distribution and consumption.

What is AI?

To put it simply, AI is a branch of computer science that deals with the creation and research of human-like intelligence in machines. However, it has recently taken on a much broader meaning and now acts somewhat as shorthand for a range of different technologies and techniques (from robotic process automation, to machine learning and natural language processing) that represent the current leading edge of computerisation and automation. It is this wider range of AI and related technologies (which we collectively refer to as “AI”) that we look at in this report.

When does it make sense to use AI?

While one could theorise how AI might fit into almost any process, it is fair to say that most current adoptions of AI are more akin to the next stage of automation or optimisation, building upon Enterprise Resource Planning (ERP) and other systems rather than what the public may think of as AI. The relatively task-focused nature of current AI indicates it is best suited to certain situations and processes where the level of repetition/standardisation and the amount of data and information mean that it is difficult for a human to undertake the task efficiently or effectively.

AI’s qualities:	Means that AI can be useful when:	Which creates opportunities to do things more:
Ability to repeat processes and tasks quickly, completely and consistently.	A process is highly routine and elements of judgement can be expressed as sets of rules.	Efficiently 
Ability to process, structure and analyse large amounts of data.	There is too much information for a human to consider on a consistent and thorough basis or devoid of human bias.	Effectively 

AI and judgement

Whilst ultimately many would like to use AI to fully automate human judgement, there are limited examples of true judgement focused AI in real world application. This is perhaps a reflection of the fact that it is not technology alone which drives adoption, it is also about trust and control. There is need for AI to build trust with business, regulators and other stakeholders before it can be applied on true judgement, and developing that trust might also need each stakeholder to first improve their understanding and skills.

AI in action

AI is already being used within business and by consumers, from Google’s use of AI to recognise and organise images based on search, to more basic customer service chatbots. However, many of the current applications of AI could be categorised as focused on sourcing and structuring of data, the implementation of complex rules and processes, or a combination of the two. Whilst a number of AI applications exist, this report addresses the question of where AI might feature in the corporate reporting process.

AI and corporate reporting

Corporate reporting is a mechanism to create trust and transparency in a company’s financial position and performance. It does this through the rules, regulations, assurance requirements and communications practices that are focused on the annual report and other regulatory documents. While the corporate reporting process is both highly complex and effective, there are challenges. Three specific challenges that we use to explore the potential for AI are:

- The efficiency of recording and aggregating transactions, across multiple entities, and then turning that data into an external communication;
- The efficiency and effectiveness of providing internal or external assurance over the resulting communication; and
- The effectiveness of consuming the information reported by, and about, the company and translating information into insight and ultimately into action.

Is AI the answer? Our review suggests that AI is potentially, and eventually, the answer.

Whilst there are usually a number of manual processes involved in producing corporate reports which limit the potential for AI, for now, there are already use cases for the production, distribution and consumption of reporting.

AI in production, distribution and consumption

Taking company data from across an organisation, aggregating it into a single communication, then distributing it to an audience of investors, who want to analyse and combine it with other external information, is a complex process. Because it is complex, it typically involves large numbers of people at each stage. AI, and its related technologies, provide opportunities to drive efficiencies and enhance effectiveness. Potential uses for AI include:



Production – AI can enhance efficiency by replacing mechanistic human processing of underlying transactions and transforming that data into accounting and management information; ultimately feeding into annual reports.



Distribution – AI can efficiently and effectively support auditors and boards in the internal and external validation processes needed to ensure that annual reports are credible and compliant.



Consumption – Investors are already using AI to enhance effectiveness of investment analysis by extracting meaning and value not only from company reporting but also from various sources of alternative data.

Whilst each of the use cases are independent, they all rely on two things: quality corporate data and a mix of accounting, technology and governance skills.

Focusing on data

The quality of company financial and accounting data is critical to the development of AI for finance. Therefore, it is important to remain data smart and consider both the risks and opportunities that AI has to offer. Those thinking about adopting AI as part of the financial reporting process should consider:

- What data was used to train the system, where did it come from, how was quality and lack of bias ensured?
- Are there opportunities in accounting changes (such as IFRS 15 and 16) to gather and optimise information for AI?
- If the system is external to the company, what is happening to the company's data and the learning derived from that data? Is it portable?
- What controls and processes need to be changed or modified for AI?

Focusing on the opportunities of structure

It is not just company internal data that is important, externally reported data is also key. In the world of AI, data has value, but structured data has even more value. When thinking about corporate reporting, many of the distribution and consumption use cases are using data sources that are digital by design. However, the information in a company's corporate reporting is not, and in fact is still predominately paper-based. To facilitate the use of corporate reporting information, data that is required to be made publicly available could be done so in a way that is open, structured and allows for low cost aggregation and reuse. Regulators, preparers and users could discuss how to embed the principles of openness, structure and reuse into new and existing corporate reporting requirements in a cost effective way.

Focusing on skills

Through our discussions with various participants, we heard that the most interesting developments are where accounting and governance skills are combined effectively with technology skills. This mixing of skills is difficult and often relies on external consultants to provide it. But ultimately we need to embed the expertise of AI directly into governance, finance, the board, advisors and regulators through training and development.

Bringing it all together

Our review of AI leads us to conclude that it is not a question of **will** AI become important for corporate reporting, but **when**? However, as this AI powered world of reporting develops, it is important that all stakeholders understand how reporting empowered by AI needs to evolve to ensure that quality, trusted corporate reporting output is maintained and enhanced.

Section 1

Artificial Intelligence basics

- What is AI?
- AI in the real world

This section of the report covers what AI is at a high-level and provides some examples of how it is beginning to be deployed in the real world. Appendix 1 provides a more detailed explanation of the different types of AI.

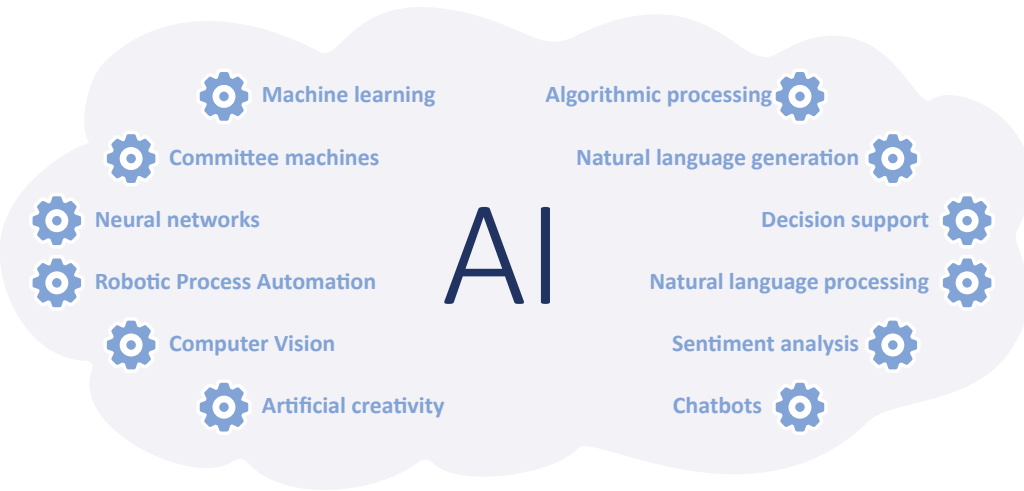
What is AI?


To put it simply, AI is a branch of computer science that deals with the creation of, and research on, human-like intelligence in machines.

AI has been talked about since at least the 1950s, but recent technological developments in both computing and software power have lead to significant progress in the development of the field. While high profile events (such as the defeat of the Go world champion by an AI called DeepMind) have received much of the press, it is the less newsworthy, simpler and more focused technologies that are beginning to have a significant impact on business processes.

Recently AI has taken on a much broader meaning and now acts somewhat as shorthand for a range of different technologies and techniques that represent the current leading edge of computerisation and automation. In this report we classify these technologies as AI.

AI as a term is commonly used to describe a number of related and unrelated technologies.





KEY POINT: AI is used to describe a range of technologies that are driving automation of business processes.

What are the different types of AI?



There are a number of ways to classify different types and techniques of AI. However, the majority of AI in current production is so called narrow or task focused AI. We therefore consider the types of tasks that they undertake as a way of categorising them.

The types of AI are:

- AI that provides computers with knowledge – such as machine learning and expert systems
- AI that helps computers to recognise text or images – such as computer vision and natural language processing
- AI that helps computers create images or text – such as artificial creativity and natural language generation
- AI that helps computers copy a task – such as robotic process automation

Appendix 1 gives more details about each type of AI.

What type of tasks are suitable for AI?

AI’s qualities:	Means that AI can be useful when:	Which creates opportunities to do things more:
Ability to repeat processes and tasks quickly, completely and consistently.	A process is highly routine and elements of judgement can be expressed as sets of rules.	Efficiently 
Ability to process, structure and analyse large amounts of data.	There is too much information for a human to consider on a consistent and through basis devoid of human bias.	Effectively 

AI should not be the default for most business problems. Particular areas where AI can be useful are those where efficiency or effectiveness could be optimised. These include:

- Processes where a clear business problem exists that can be easily articulated.
- Tasks that involve large amounts of good quality and relevant structured data.
- Tasks that are prone to human error because they are repetitive or require concentration for a sustained period of time.
- Tasks that are not economical or efficient for a human to undertake.

AI in the real world

With any technological development there is a certain amount of hype about its potential and suitability for various applications, and AI is no different. While one could theorise how AI might fit into almost any process, it is fair to say that most current adoptions of AI are more akin to the next stage of automation building upon Enterprise Resource Planning (ERP) and other systems rather than what the public may think of as AI. Current AI implementations are unlikely to be independent of the current IT systems and are often interacting with the ERP tools that form the IT backbone of companies.

Examples in practice

Zurich Insurance has used RPA to improve the efficiency and accuracy of its policy booking and document issue process. RPA is used to take the data from the initial quotation, it is then processed into the underwriting system, and then policy documents are issued off the back of the underwritten policy. Zurich are now exploring additional uses of RPA, including the processing of claims payments. Overall efficiency gains are expected to be around 40 percent.

Heineken uses RPA to assist with overdue invoices. One of the challenges of the process within the organisation was having multiple systems where data on the overdue invoices existed. RPA was used to pull together data from 27 different ERP systems, populating the resultant data into a database that allowed human operators to undertake additional reporting and follow-up.

Uses for AI

This report will discuss some of the actual or potential uses of [narrow](#) (task focused) AI across the corporate reporting process, from the production of corporate reporting, to its distribution and final consumption. Many of the current applications of AI revolve around sourcing and structuring of data, implementation of complex rules/processes, or the combination of the two.

Even in these situations it may not be the AI that drives all the benefits. In fact, many of the benefits of AI stem not from the AI itself, but from process and procedure reform as part of more narrow deployment of AI.

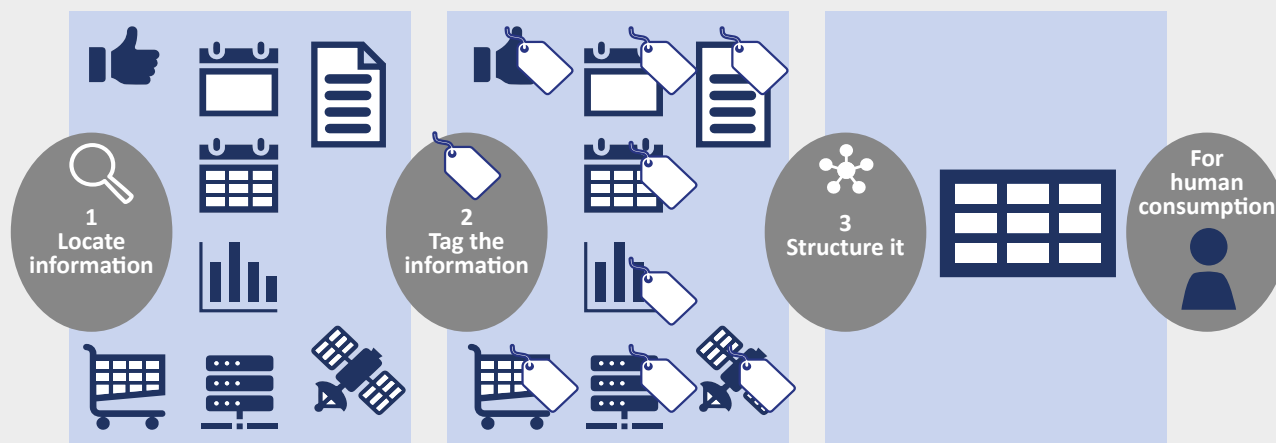
Sourcing and structuring of data

The amount of data in the world is growing, with predictions that it will reach 164 zettabytes by 2025¹. However, much of that data is unstructured (up to 80 percent), and this means that the data is not presented in a uniform way to provide important context. Many business processes, including advanced analytics and management information systems, need structured data to identify drivers of performance, make predictions, and identify potential relationships. Systems, therefore, have been developed to source and structure unstructured data, many of which use AI to automate the process.



KEY POINT: AI is not magic; it works best in narrow applications where the need for efficiency outweighs flexibility.

How it works: sourcing and structuring of data



- 1 Various elements of unstructured data are gathered together from multiple sources, either internal or external to an organisation. Sources might include web sales data, photos, newspaper articles, social media posts, satellite data, sensor data and financial or accounting data.
- 2 A machine uses techniques such as [natural language processing \(NLP\)](#) or [image recognition \(IR\)](#) to make some sense of the unstructured data. It then applies a structure to the object, either by applying the original structure inherent within the information (but not originally recorded in a machine-readable way) or by overlaying a recognised external or internal data structure. There might be multiple data structures attached to a single data object, allowing data objects to be used for multiple purposes. A technique such as [expert systems](#) or [machine learning](#) might be used to help the machine identify or learn how an item should be successfully categorised.
- 3 The resultant data objects are then collated for onward analysis, either by a human or other AI system.

¹ IDC 2017

AI to apply rules and process

When experts undertake an analysis of what makes up a typical job role or process, they invariably discover that a significant proportion of that role or process is actually a series of repetitive tasks. Often these tasks are the application of a series of rules (e.g. here are our lending criteria, does the applicant match them?) or processes (e.g. take expense, check the signature matches the authorisation levels). Given the repetitive nature of this type of task, it is an ideal candidate for automation by AI. An AI process system will often undertake the task more quickly, more consistently and to a higher standard than a human operator.

Whilst these two categories of use are relatively simplistic they are important first steps in getting AI into the everyday world and therefore much of the focus of this report is on these.


AI and the application of judgement

Ultimately, many would like to use AI to automate human judgement, but this is tricky, and the line between true judgement, and the application of sets of complex rules in a way that mimics judgement, is not clear. However what is clear is that there are limited examples of judgement focused AI in real world application. This is perhaps a reflection of the fact that it is not the technology alone which drives adoption, it is also about trust and control. Many of the areas where there might be value in AI making judgement are areas where society values oversight and ownership by a human who can be called to account. Therefore, there is a need for AI to build trust before it can be applied to true judgement. It is perhaps only when this trust in AI exists that the transformative power of AI will be felt.

Whilst there are currently minimal examples of true machine 'judgement' in AI, this report identifies some moon-shots that extrapolate from the current trends in use cases.

How it works: a number of technologies and techniques might be used to apply rules or process

1A Rules provided by human 'experts'



2 Are then fed into the AI and applied to data



3 Exceptions passed to humans



1A/B The system is provided with the business or process rules. The rules might be created by recording the experience and knowledge of one or multiple human experts, or the rules might be self-taught through machine learning techniques.

2 The system then obtains the required data and applies the rules.

3 Where an item either fails the rules or is otherwise flagged as an exception, it might be passed to a human to review.

4 The resulting data objects may then trigger further activity such as communication or human interaction.

Rules and processes might be simple (e.g. is 'X' included, yes or no?) or more complex by using layering and repeated application (e.g. should we interview this person?). At the upper end of complexity, it may appear that the system is intelligent or is applying judgement, but in reality this is not the case. The nature of these systems means that the combination of a human and an AI may lead to optimum efficiency and flexibility.

1B Rules gained through machine learning



4 Processed items passed on to next part of the process



Section 2

Artificial Intelligence and corporate reporting

- The nature of corporate reporting and reporting process challenges
- AI and the production of corporate reporting
- AI and the distribution of corporate reporting
- AI and the consumption of corporate reporting

This section of the report discusses how corporate reporting is currently structured and considers some potential use cases of how AI technology could improve the production, distribution and consumption of company information.

The nature of corporate reporting and reporting process challenges

Corporate reporting is a mechanism to create trust and transparency in a company's financial position and performance in an efficient manner. It does this through the rules, regulations, assurance requirements and communications practice which are focused on the annual report and other regulatory documents. While the corporate reporting process is both highly complex and effective, there are challenges. Three specific challenges that we use to explore the potential for AI are:

- The efficiency of recording and aggregating transactions, across multiple entities, and then turning that data into an external communication;
- The efficiency and effectiveness of providing internal or external assurance over the resulting communication; and
- The effectiveness of consuming the information reported by, and about, the company and translating information into insight and ultimately into action.

Our report considers; is AI the answer?

The following section describes at a high-level how AI might be used in the various stages of reporting to answer the specific problems identified, and some of the reasons why it might or might not make sense in each case. Given that AI is at an early stage of development, each section also includes some moon-shot ideas of where AI might be used in a fully digital corporate reporting process of the future.

Measuring up

AI doesn't operate in a vacuum. In order to be adopted for corporate reporting, technologies don't just have to solve specific issues, they must also align with wider demands from preparers, users and others. The Lab's earlier work led to the development of a digital framework of characteristics that expresses a number of key demands that digital corporate reporting should embody. It is by considering both AI's ability to solve problems and the degree to which it meets the characteristics (or could do so) that give a full picture of its potential in corporate reporting.

Characteristics of Future Digital Reporting



Full details on the characteristics are available in the Lab's original [report](#).



KEY POINT: The field of AI is very wide, containing multiple technologies each with differing levels of relevance and effectiveness for use in corporate reporting. However, if it is to be deployed for corporate reporting uses, AI needs to fit into a wider eco-system.



AI and the production of corporate reporting

COMPATIBLE

EASY

COST-EFFICIENT

TIMELY

What problems might AI solve?

Corporate reporting and underlying accounting processes are often complex and time consuming. There is potentially the need to record transactions on several systems, aggregate and consolidate the resultant information, and make adjustments, apply judgement and develop estimates. The resulting accounting information then needs to be translated into an external annual report and be supported by relevant narrative. All this needs to happen within a short time frame and in a way that is cost effective, easy to operate and which minimises the risk of error or mistake.

Two places that AI might be used to help solve problems in the production of accounting and reports are in the efficiency of recording and processing of external transactions and the translation of those transactions into external communications.

Transaction processing – Search for efficiency

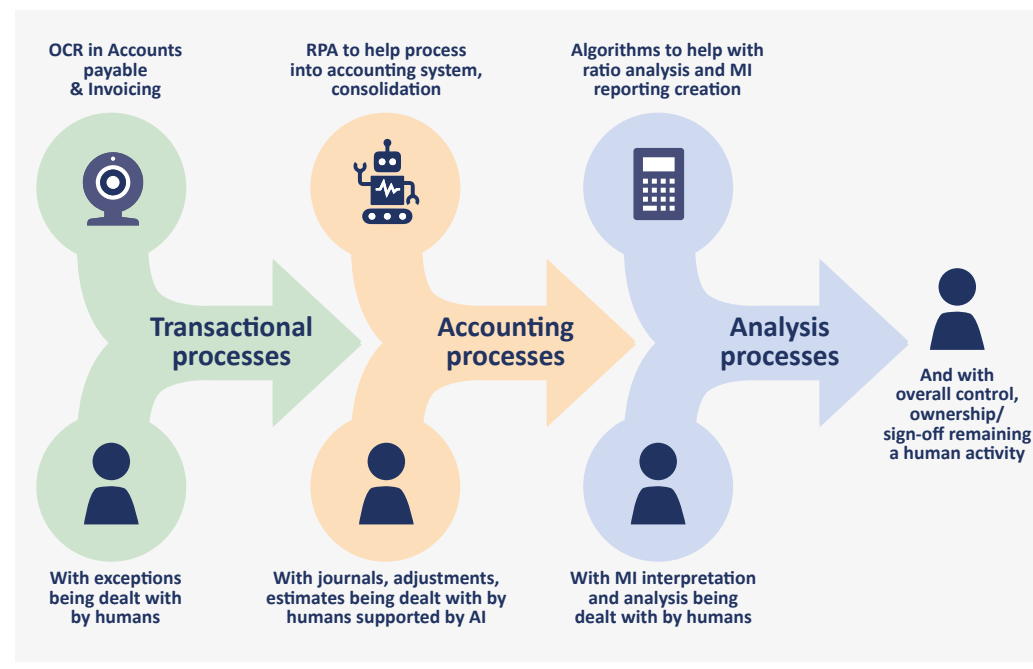
A significant proportion of the work of a typical back-office finance function is the routine recording, managing, matching and processing of transactional and other information. Companies need this work done as efficiently as possible and this has led many to outsource business processes or create shared service centres. However, AI can provide an alternative to the outsourcing of such activities. Processes like expenses or accounts payable could use a combination of [Optical Character Recognition](#) (OCR, a way for a computer to recognise letters and numbers within an image) to make sense of the unstructured receipts or invoices efficiently. [RPA](#) could then post the resulting data into underlying accounting or other systems. Once processed, algorithms (trained using expert systems or machine learning) might then undertake a quality review, cross-check to other data sources or systems to conclude on the validity of the transaction and flag issues to a human for further follow up.

It is not just in the initial processing where AI can be of use; turning the transaction data into accounting data usually involves the classification of an entry. AI systems could auto-classify an entry with an initial category or suggest it to a human user. This might be bespoke to the company, or powered by data from large groups of users (such as is being provided by many cloud accounting tools). The analysis of the data might also be assisted by AI with the automation of Management Information (MI) reporting, ratio and exception analysis.

Many of those who are implementing AI in search of efficiency are focusing on specific elements or processes that are easy to automate and are compatible with the current control framework. However, given the overall capacity of AI, wholesale redesign of the entire end-to-end finance process may ultimately lead to the most effective outcomes.

AI can be used in the finance process

Various AI can be used to automate elements of the finance and accounting process handing off to humans at various stages such as:



Does it make sense to use AI?

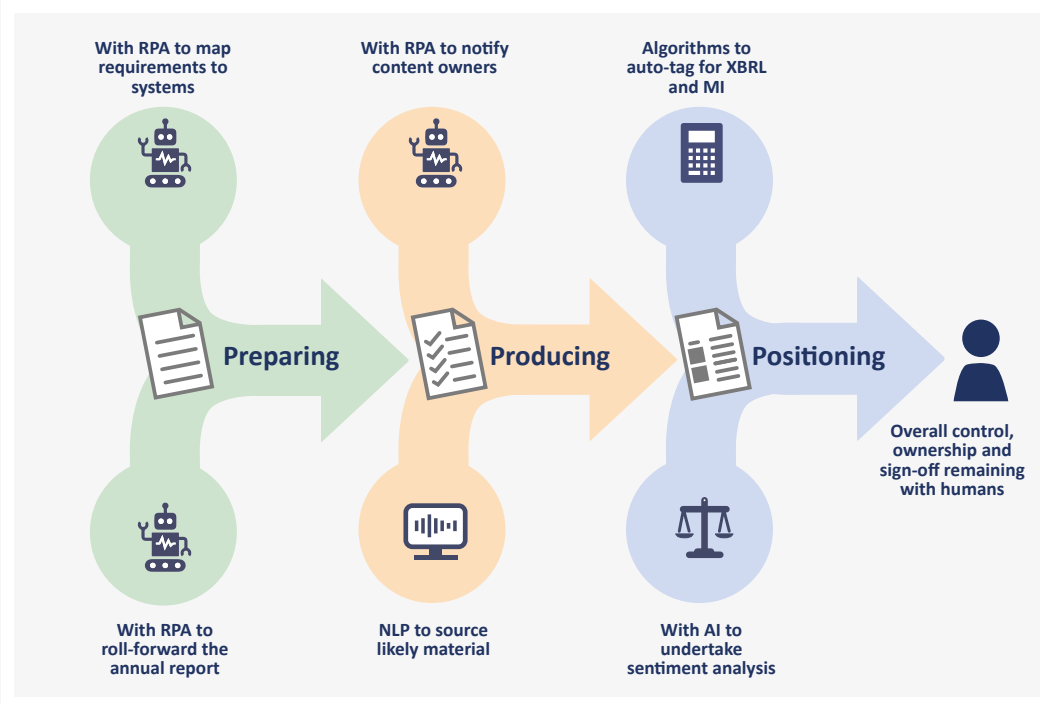
Given that efficiency is highly prized in finance functions and underlying data, and processes are relatively standardised and structured, the application of AI is sensible. However, the level of overall judgement devolved to the AI will depend on how trusted and controlled the AI is within the organisation.

Annual report creation – efficiency first, effectiveness later

Where organisations have already undertaken basic automation of transactions and MI, further and more complex processes could be considered for the application of AI. The process of creating an annual report is complex, time consuming and not always efficient, and AI could assist. In preparing a report, RPA could provide an ideal technology to create elements of the draft annual report, taking the previous year's report, rolling forward the prior-period numbers, mapping the current period and sending this to the accounting system. AI could help in producing the report by sending out notifications to section owners or perhaps using [NLP](#), (a way to obtain meaning from text and speech) to source material from the company's public statements and MI. Once the report has been finalised, AI could then be used to auto-tag the accounts with MI tags or XBRL tags to allow onward analysis. AI can also help in the effectiveness of corporate communications by using sentiment analysis to consider how the report will be read. It might also suggest alternative wording to better match management's intention.

AI can be used in the reporting process

Various AI and related technologies can be used to automate elements of the preparation of the annual report such as:



Does it make sense to use AI?

Because the reporting process is entity-specific and only happens infrequently, some may consider that it is not suitable for automation. However, standardised off-the-shelf AI building blocks could be used to construct a process that works with only limited entity tweaks, especially where companies have already begun to digitise.



What is the time scale?

AI for basic processing already features to a limited extent in some large finance functions. Over the next 5-10 years the use of it across processes is likely to increase further as off-the-shelf AI building blocks become available to smaller companies. The use of AI in creating annual reports is likely to be further off as it needs basic automation to have already been undertaken. Given that external reporting is a highly sensitive area it is also likely that AI use would come only once management are comfortable and trust the technology.

How production use cases express the digital characteristics

G

COMPATIBLE

AI like RPA are highly compatible with current and legacy systems. Ultimately the most benefit might come from wholesale redesign and optimisation with AI in mind.

A

EASY

AI is a developing field. As such it currently requires some specialist skills and assistance. This reduces ease of use but should improve as the workforce gains experience.

G

COST EFFICIENT

AI aims to improve cost efficiency as a core rationale. However more complex AI might be expensive to deploy for unique use cases.

G

TIMELY

AI improves timeliness by developing iterations more quickly and more extensively than is possible for a human.



KEY POINT: AI works best when it is considered as part of a wider process of optimisation, digitisation and automation of the finance and reporting functions.

AI in the real world: Vodafone

Vodafone is one of the world's largest telecommunications companies with mobile operations in 25 countries serving more than 530 million customers, including all of the customers in Vodafone's joint ventures and associates. Vodafone has lots of data that it needs to analyse to help focus and manage the business. In the last couple of years Vodafone has begun to roll-out AI into the finance function to help assist the business analysts and management information teams.

The two key drivers of Vodafone's use of AI is to help do things more efficiently, and to enhance analysis that generates better insight. Two areas where Vodafone tried this approach has been in the identification of 'exceptions/potential issues' within the accounts payable process and revenue analysis.

Vodafone have been trialling AI to build a more data-centric understanding of performance factors and trends over time. This allows for better forecasting on those factors that will impact the company's performance and has led to enhanced understanding compared to the traditional MI approach.

One of the critical factors for the AI forecasting has been the obtaining of relevant and actionable data, at scale. Vodafone used the opportunity brought about by requirements for data to comply with the new IFRS 15 accounting standard to build systems and processes that delivered data of value to the forecasting process.

Vodafone's advice to companies thinking about deploying AI would be to consider what you are trying to achieve, and develop a clear understanding of whether AI is the best solution. If AI is the best option, you should then think about how to obtain enough, quality data.



Moon-shot: What could be possible

2034: The age of the perennial report

There are already NLP/NLG tools in use in the media industry that write simple stories such as financial or sports reports that follow a relatively consistent structure. Projecting forward, it may be possible to create advanced NLP/NLG systems that write an entire annual report, including creating supporting images and graphs, based on MI and other structured company information. Projecting even further into the future, perhaps the annual reporting-cycle could end, with reports being generated at any point of time and covering any time period.



2034-2044: Natural Accounting Processing and Generation

Accounting standards are, at their heart, a set of instructions as to how to recognise, measure, record and disclose business information. If the standards themselves were made machine readable, ML and NLP could be used to understand, then generate, interpretations of unique situations based on historic practice and underlying frameworks, and provide the preparers of the future with a robo-accounting technical advisor.

Looking further ahead, you could see that new standards might also be supported by AI. This could include the analysis of current practice to passively survey gaps or areas where standards might be needed or perhaps NLG could be used to write a draft standard based on accumulated historic standards.



AI and the distribution of corporate reporting

ACCESSIBLE COMPLIANT PROMPT FREE

What problems might AI solve?

Annual reports and other outputs from the corporate reporting process are just the tip of the iceberg of internal data, transactions, consolidations, estimates and judgements made across multiple entities. The resulting aggregate annual report is of value when it can be trusted. Trust is created through compliance with standards, laws and regulation. Compliance is assessed (to some extent) through internal and external assurance and governance processes. However, the modern business is complex and traditional methods of attaining assurance can be costly and may not be efficient or timely, potentially impacting the overall promptness of the resulting document.

AI might be used to help provide assurance on compliance with rules, both internally and externally, in an efficient and effective way.

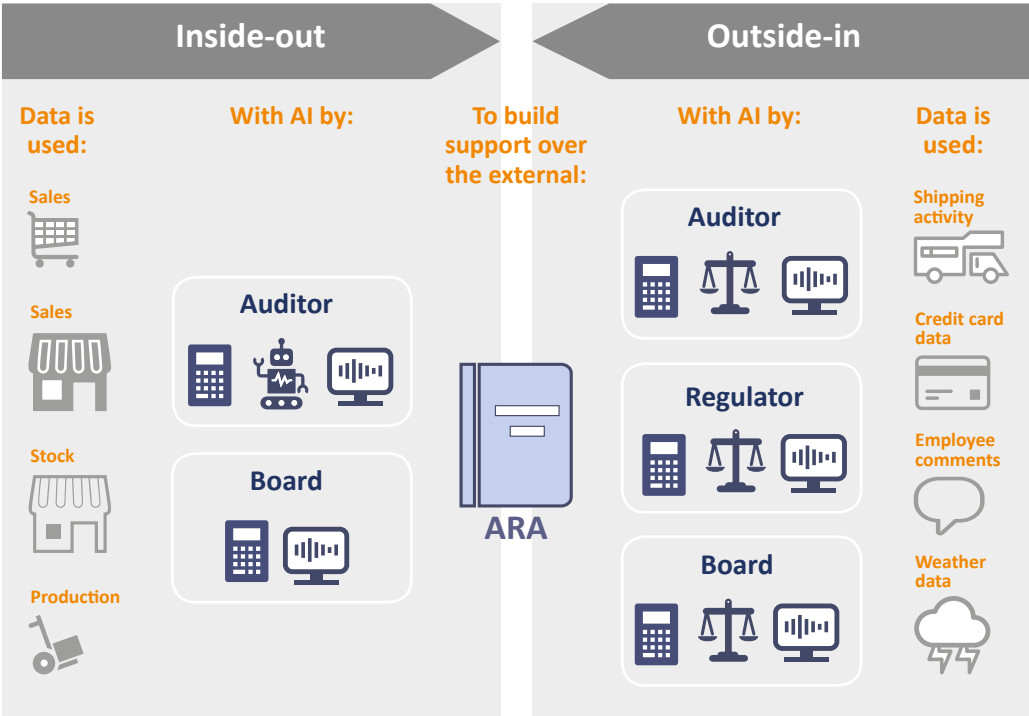
Inside-out assurance (efficient)

Auditors (both internal and external) obtain comfort over a consolidated annual report, specific entity report, balances or process by undertaking testing. Typically comfort is obtained through audit tests of relevant controls and substantive testing of transactions and balances. Testing is undertaken on a sample basis as it is impractical to review all transactions, due to the limitation of the human team to physically review each transaction and follow-up on exceptions. AI, however, potentially provides a different way of approaching the problem. AI combined with data analytics tools might allow 100 percent of certain balances to be rechecked or recomputed and connections to external data sources such as bank or investment feeds could also match and confirm transactions. Another example might see NLP combined with OCR and algorithms to read contracts, invoices and other documents and recheck classification decisions or perhaps identify key terms. AI is good at looking for patterns and therefore could be used to identify fraud or error, either historically or proactively. AI has applications across an audit and could replace individual elements of audit processes or, once trusted by regulators, clients and auditors themselves could be connected together to automate the less judgemental elements of the entire audit.

Outside-in assurance (effective)

By their very nature, audit and board reviews of reporting are focused on a company's internal records. Both auditors and boards will seek external confirmation of specific elements or facts as part of the validation and review process. However, AI could provide new ways of obtaining comfort on a company's communication. For example, an audit committee that wants to consider if the annual report is fair, balanced and understandable could use textual or sentiment tools to support their review. A board wanting to challenge culture or customer disclosure could use AI tools to source and analyse external opinion such as Glassdoor™ or Twitter™. Other data sources obtained and analysed by AI, such as credit-card transactions, web-traffic and shipping activity might highlight areas where the external indicators don't match the company's internal perspective. These types of tools would not just be the remit of boards but could also be used by external auditors or regulators to build up a perspective on an organisation that went alongside more traditional sources of assurance. However, caution is required. As with all analysis, the validity of the result depends upon the validity of the input and AI does not compensate for poor data.

AI can be used in the reporting process



AI can support auditors and boards in getting comfort over the data internal to the company that makes up the annual report.

AI can support auditors, boards and regulators in getting comfort over annual reports using data external to the company.



Does it make sense to use AI?

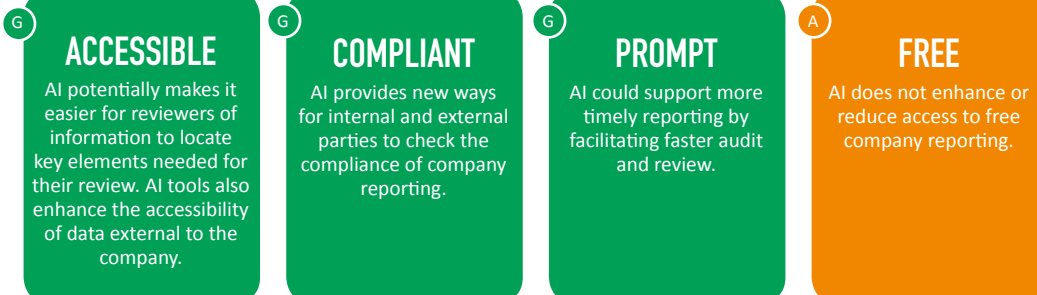
Businesses are complex, as is the data that flows from them. Obtaining comfort over the data is part art and part science. Whilst certainly AI will play a role in both internal and external audit processes in the future, the overall complexity and uniqueness of each business requires a Human+AI approach for the foreseeable future.



What is the time scale?

AI is already being used by external audit firms in some areas (such as fraud analytics and journals). However, full scale adoption is likely to be gradual, only accelerating when such use is accepted by both companies and regulators. The Board's use of AI will likely move in lock-step with the availability of easy-to-use, secure and affordable tools and enhancement in skills. Sentiment analysis is already being used in the corporate reporting sphere, but other more complex tools are likely to be further off.

How distribution use cases express the digital characteristics



Moon-shot: What could be possible?

2048: LegAI Entity. What if the legal, body-corporate was powered by an AI brain?

A company is a mechanism to create an entity which acts separately from the humans that work within it. There are already examples of AI board directors in some companies. Is the next step the replacement of a board (being in effect an amalgam human brain) with a full AI board or even a single AI? This leads to interesting questions such as:

- How would audit work when the company itself has a voice and can be communicated with as a virtual reality avatar?
- Would the annual report be replaced with something akin to a school report card or a medical check-up?
- Would investors own the entity or would it own itself?



AI and the consumption of corporate reporting

ENGAGING

CONTEXTUAL

USEABLE

CREDIBLE

What problems might AI solve?

The investment industry has long been a user of algorithms, bots and other AI to execute trades much more quickly than a human. However, it is not in trading that the real changes for investors will occur, it is in the process of effectively consuming information.

At the heart of the investment process there is a need to effectively understand the past, current and likely future performance and position of a company, as this forms the basis of the investment decision. Traditionally investors' understanding has been obtained through analysis of companies' own public disclosures (across a range of documents such as the annual report, presentations etc). However, the amount of information about all aspects of our lives is growing exponentially, and this is also true for information about companies. The information that a company publicly discloses is only a small historic slither of potential investment-relevant data. Investors have therefore increasingly looked at alternative data sources for insight into the likely future performance of a business, in real-time. The scope and nature of alternative data is potentially very wide; from satellite images of farms to credit card and social media data. This wide pool however, leaves investors with an issue; how can the data be managed, its credibility judged and the data be made useful, so that real investment-contextual insight can be obtained and actioned?

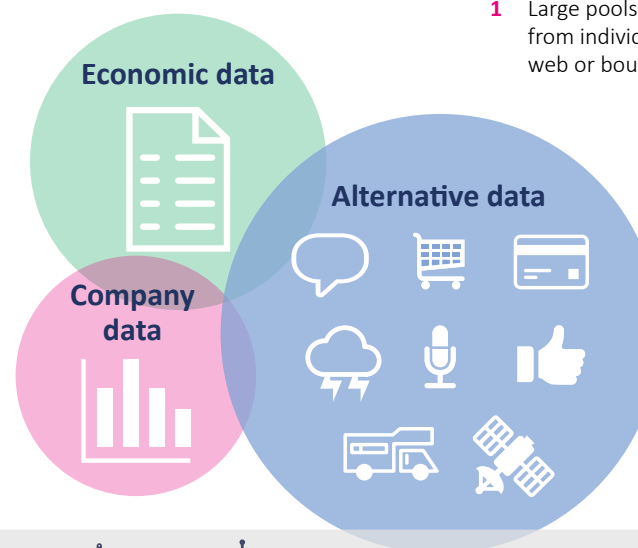
Many investment organisations and their data providers view AI as a potential solution to the challenges of collecting, collating and creating actionable insight from data.

Creating an insight engine

The first stage of creating actionable insight is to gather together the relevant data. AI can assist here as web crawlers and other bots (small programmes) can source alternative data sets from across the web. RPA can also be used to take standardised data sets (such as government economic reports or structured corporate reports) and process them into a database. Once the data is gathered AI can then undertake additional analysis.

Techniques such as data mining and sentiment analysis can be used to identify key words, phrases or overall tone which might match with a specific characteristic such as indicators of credit worthiness. More recently deep learning and machine learning techniques have been used to identify patterns between future events (such as credit default) and management disclosures. Such models might also make connections between government data or social media data and specific future events. The resulting analysed data will then either be fed into a wider model or flagged to a human analyst for further consideration and possible investment decision.

AI in the investment process



- 1 Large pools of data are sourced directly from individual agencies, scraped from the web or bought through data providers.

- 2 AI tags and analyses sources, then uses AI to identify those needing human attention.

- 3 A human reviews, interprets and takes actions based on the AI sourced information.




KEY POINT: For most preparers, all of their public data is already being analysed by AI and compared against external data sources for indications of short term opportunities. Companies might want to consider what story that external data tells and whether that matches the company's own communication.



Does it make sense to use AI?

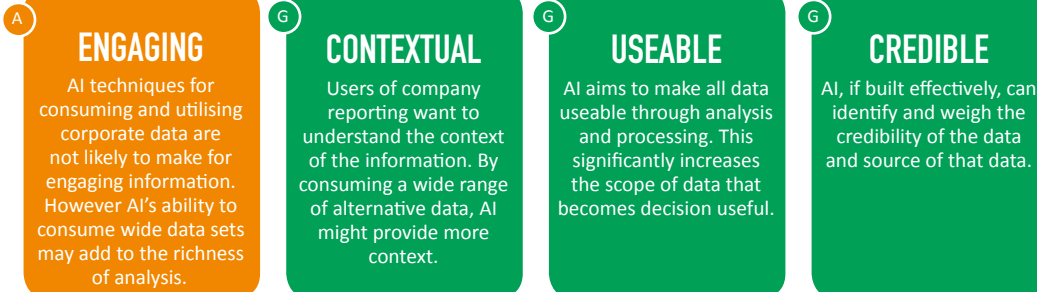
The investment process is dependent upon the effectiveness of ingestion and analysis of lots of information in relatively short time frames. The time constraints and the requirement to deliver better insight at scale and for low cost make for an attractive AI use case. The process and information sources are also relatively structured and standardised, which further enhances the case for AI. However, what one investor can do, so can others. Over the longer-term, the commoditisation of AI within the investment analysis sector could erode any competitive advantage from using AI and/or alternative data.



What is the time scale?

AI already features in some areas of the investment analysis industry. Development of powerful AI by service providers and increases in structured public company data could see wholesale adoption of AI in the next 5-10 years.

How consumption use cases express the digital characteristics



Myth-buster: AI means we don't need to structure data

Regulators around the world are increasingly requiring companies to report using structured data in regulatory, and tax filings. However, the advent of AI and big data techniques across unstructured data raise questions in the minds of some about the value of structuring.

Whilst eventually the value of structuring might be limited once very advanced AI is developed, at the moment the current generation of AI needs structured data to learn and understand. Furthermore, from a preparer's perspective, structuring of data provides a company the opportunity to communicate intended meaning and classification for the data, rather than leaving it only to a machine to interpret.



Moon-shot: What could be possible?

2022: Equity scanalyst

One of the key outputs of the Equity Analyst is the 'analyst's report'. This report will often convey thoughts and opinions about a company's past and potential performance. The industry itself has come under significant cost pressure in recent years, which has led some to utilise offshore analyst centres to draft much of the analyst report. However, with the improving power of [NLP/NLG](#) some firms have already begun investigating if AI can co-create the report. Projecting even further out it might be possible for an AI to identify, draft and release research perhaps with only minimal human interaction. This could be especially valuable for companies in the markets that generate little coverage today.

Section 3

Conclusions and actions


- AI Reporting: actions and questions
- Next steps?

This section highlights some of the actions that different parties might need to take if we are collectively going to make the most of the opportunity that AI presents.

AI Reporting: actions and questions

At the beginning of this report we asked if AI had a place in the future of corporate reporting. The current and potential examples in this report drives us to conclude that it is not a question of **whether** AI will become important for corporate reporting, but **when**.

Taking company data from across an organisation, aggregating it into a single communication, then distributing it to an audience of investors who want to analyse it and combine it with other external information is complex. Because it is complex, it typically involves large numbers of people at each stage. AI provides opportunities to drive efficiency in what we do now and enhance effectiveness by doing things differently. Key uses for AI in each stage include:



Production – AI can enhance efficiency by replacing mechanistic human processing of underlying transactions and transforming that data into accounting and management information; ultimately feeding into annual reports.



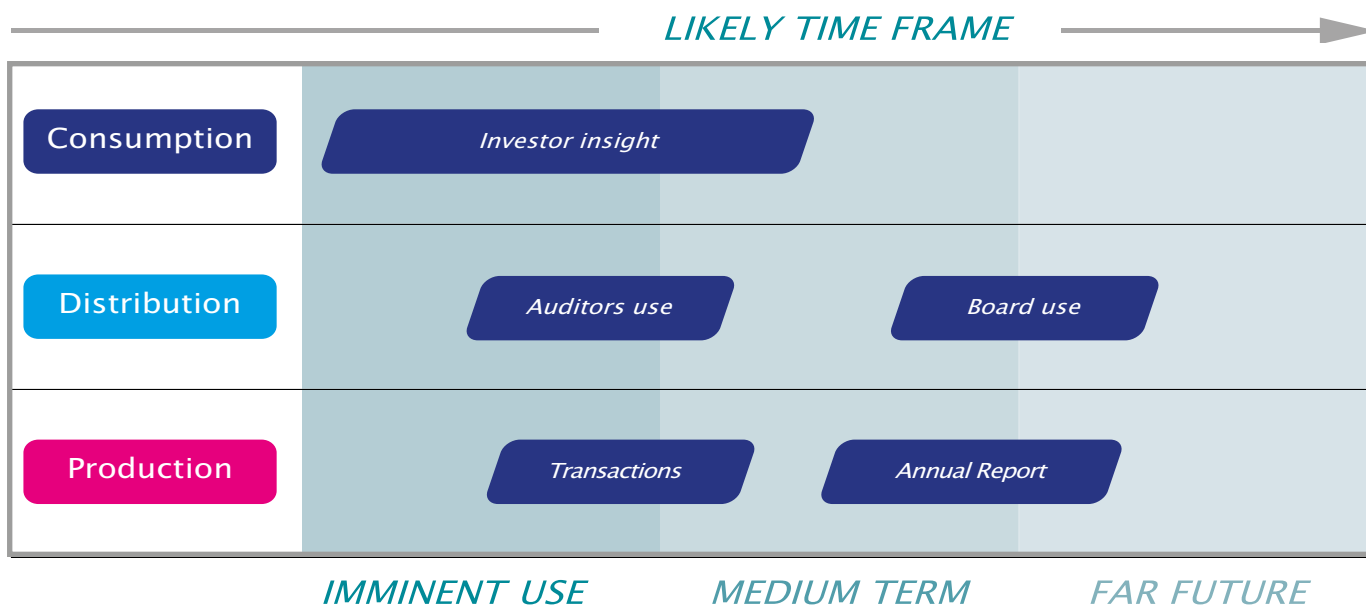
Distribution – AI can support auditors and boards in the internal and external validation processes needed to ensure that the annual report is credible and compliant in an efficient and effective way.



Consumption – Investors are already using AI to enhance effectiveness of investment analysis by extracting meaning and value, not only from company reporting, but also from various sources of alternative data.



KEY POINT: AI is coming and as the market develops, AI as a service will provide new ways to rethink current processes to increase efficiency and effectiveness.



Whilst each of the use cases are independent they all rely on two things; quality corporate data and a mix of accounting, technology and governance skills.

Data

Many of the use cases for AI need data in order to make them work. The use cases we have considered in this project point to the same conclusion; quality data is important. The need for data in an AI-based economy has led to a number of governmental and industry initiatives which consider questions around access to data, structuring, privacy and bias. For the use cases we have considered some of the same questions might also need to be considered.

Company financial/accounting data

If AI is to be used to improve the efficiency and effectiveness of companies' financial processes then they need data that is of sufficient quality and quantity and is relevant to financial decision making. However, many of these AI products are being created by consultants and technology companies which might be external to the organisation. Companies considering adopting AI-based accounting and finance systems might want to consider:

- What data was used to train the system, where did it come from, how was quality and lack of bias ensured?
- If the system is external to the company, what is happening to the company's data and the learning derived from that data?
- Can the learning and data be transferred between providers?
- What controls and processes need to be changed or modified to reflect the new systems?
- Where significant to the company's operations, how are the risks associated with the use of AI and data being managed and disclosed?

As well as risks there also remain many data opportunities for companies. Through the course of this research we have heard from a number of companies who have used the new leasing and revenue standards as an opportunity to reform their data collection process and facilitate potential future use of AI for analysis. Companies going through finance or accounting systems' upgrades and/or process changes might want to think about the data opportunities that arise from the changes.

Corporate reporting data

In the world of AI data has value, but structured data has even more value. When thinking about corporate reporting many of the distribution and consumption use cases are using data sources that are digital by design. However, the information a company reports from its corporate reporting process is not digital and in fact is still predominately paper-based. To facilitate the use of corporate reporting information, data that is required to be made publicly available could be done so in a way that is open, structured and allows for low cost aggregation and reuse.

Preparers, users and regulators could discuss how to embed the principles of openness, structure and reuse into new and existing corporate reporting requirements.

Skills


Many of the current AI use cases revolve around bolting AI on to existing processes. This is not surprising as it represents the continuation of well understood and controlled processes. However, many of the most interesting opportunities around doing things more effectively need the rethinking of processes with AI in mind. This needs a change in skills.

Through our discussions with various participants, we have heard that the most interesting developments are where accounting and governance skills are combined effectively with technology skills. This mixing of skills is difficult and often relies on external consultants to provide the additional perspectives needed within a finance or governance project. But ultimately we need to embed the expertise of AI directly into governance, finance, boards, advisors and regulators through training and development.

Advisors, professional bodies, tech firms and regulators should continue to work together to embed key AI skills into those current training and development programmes to upskill those already in the work place.

Boards might also consider how they are going to meet the challenges that AI brings either through bringing AI focused individuals into the boardroom, or through wider training.

For AI, as with many of the new technologies that we have looked at in this project, new ways of working are needed. All stakeholders need to work together to understand how reporting empowered by AI needs to evolve to ensure that quality, trusted corporate reporting output is maintained and enhanced.



Next steps?

Over the last 18 months the Lab has been investigating the potential for technology to fundamentally change the way that companies produce their corporate report, the way that is distributed and how it is ultimately consumed. As part of this project we have already:

- [Released a framework structure](#) of characteristics that are important for a digitally-enabled system of corporate reporting.
- [Released a report that looked at the potential for XBRL](#) to turn corporate reporting into **structured data**.
- [Explored the possibilities of Blockchain to deliver structured trust and location](#) of corporate information.
- Investigated how AI, powered by structured data, can improve efficiency and effectiveness of corporate reporting via the creation of **structured processes and judgement**.

Whilst we have not yet completed our project, it is already becoming clear to us that whilst each technology has promise, it is only the combination of the technologies that will be truly revolutionary.

Next, we will complete our technology review by considering how video, augmented reality and virtual reality can be used to provide investors and stakeholders with a **structured experience**, that adds to understanding and insight to be gathered from corporate reporting. All of the insights from across our work will feed into the FRC's Future of Corporate Reporting project and will help shape a vision of what the future may hold.

Section 4

Appendix 1

- AI in more detail
- Links

AI in more detail

Since antiquity, the concept of machines acting or replacing humans has been widely discussed. However, it was really the latter part of the 20th century that saw the beginnings of what we might think of as the discipline of AI. So what is AI? To put it simply, AI is a branch of computer science that deals with the creation and study of human-like intelligence in machines.

Because AI relies on computers and data, the development of the internet and the highly powerful yet cheap computing power of the cloud has seen an explosion in the development of AI. However, AI is not a single technology, rather it is a way of describing the desired outcome; the exhibition of intelligence.

Different types of AI

At a theoretical level there are three types of AI:

- **Weak/Narrow AI** – Where AI is pursuing an individual outcome or task, such as playing chess or translating a document. A narrow AI needs to be specifically developed and trained to accomplish the task, it can't be used to accomplish alternative tasks. For example Google translate can translate but can't fly a drone. All the current instances of AI fall into this category.
- **Strong AI/ General intelligence** – Where AI is used, for example, across a range of different tasks by the same system. This would be more like a human; humans can both compose poetry and ride a unicycle, current AI can't. This is the focus of much of the current wave of development but has not yet been achieved.
- **Super strong AI** – Where AI is used in ways that far exceeds the capacity of humans. This has been discussed on a theoretical basis but is not currently feasible.

Given that most of the current instances of AI are narrow, this is the focus of our report.

What are the different types of narrow AI?

There are a significant number of techniques and technologies that fall under the category of AI. However, there are a number which form a core. These core approaches can often be combined and connected to create a wide range of applications. The core approaches are:

- Approaches that provide computers with knowledge.
- Approaches that help computers recognise text or images.
- Approaches that help computers create text or images.
- Approaches that help computers copy a task.

In the next few pages we describe some of these techniques at a high-level. Our selection is not exhaustive but reflects those featuring in the report.

Approaches that provide computers with knowledge


There are a number of ways that a computer can be taught what it needs to know in order to complete a task; some of these involve a computer being trained by a human, others involve the computer learning on its own.

Expert systems

An expert system is where a computer replicates a human 'expert' by applying a set of rules or procedures to a body of knowledge. Often these will be codified as 'if X then Y' style rules. For example, imagine a human considering what the correct classification for an asset is:

1. Does it belong to us? If yes then consider nature → Yes
2. Is it tangible or intangible? If tangible, consider what type → Tangible
3. Is it property? If property, classify it as such → not property
4. Is it IT equipment? If IT equipment, classify it as such → Yes, IT
5. Categorised as IT equipment.

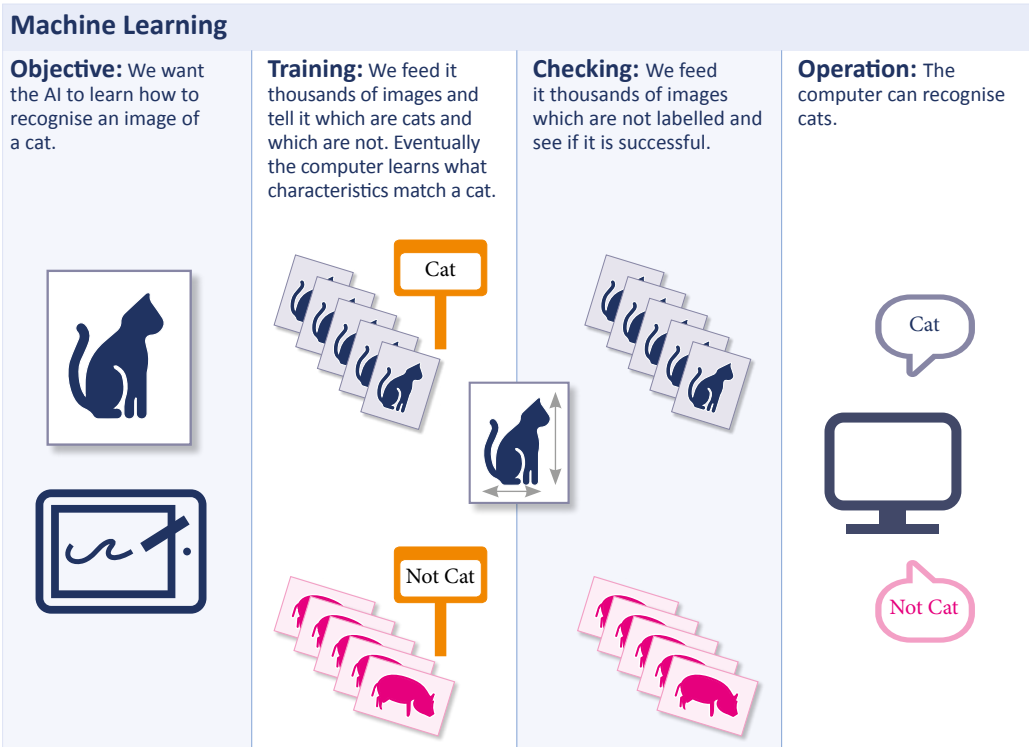
A computer could apply the same type of rules if they are written down. Expert systems are often developed through the observation and recording of one or multiple human experts, similar to creating process flow documents. Expert systems have many applications in the real world such as auto pilots or trading and investment systems. They require significant upfront development of rules and run the risk that the underlying rules are not complete or subject to bias. This is a limitation of this type of approach.



AI in more detail

Machine learning

Machine learning is the process by which machines learn from data. In a normal computer program, a programmer has written down instructions that a computer will follow. The computer’s ability to operate is limited by the amount of instructions within the original program. Machine learning takes a different approach; rather than creating a set of instructions, the programmer instead creates a set of guidelines on how to learn. This is encoded within an algorithm. By applying the algorithm to sets of data the algorithm improves its ability to operate. In a sense it learns; hence the term “machine learning”. Deep learning and neural networks are different types of machine learning that often have less human interaction. However, most common iterations of machine learning have more human supervision. First, data is provided to the system; this is often training data, which means that it contains information that allows the computer to know if it is being successful in its understanding (think of this as stabilisers for a computer). Once the algorithm has learned from the training data, it should be able to produce the same results from ordinary data (without stabilisers). This technique is fundamental to many next generation AI applications such as image recognition, natural language processing/generation, sentiment analysis and pattern recognition tools (such as fraud detection).



Approaches that help computers recognise text or images

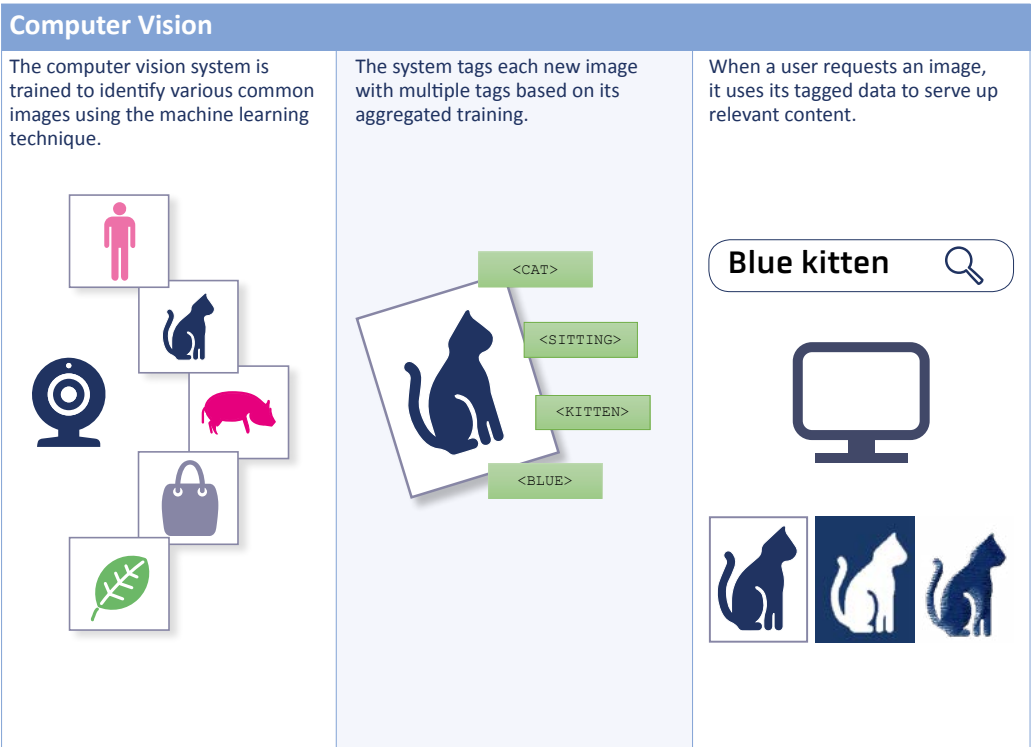
Many of the activities on which AI is being deployed are related to the processing and classification of images and text. There are a number of techniques that are used to achieve this. These will often build upon the machine learning:

Computer vision

The idea of computer vision is to get machines to collect, collate, classify and analyse images. This will often involve either a system that has been developed using machine/deep-learning or simpler rules-based systems. This type of tool is deployed across search engines (to find images), social media (to identify those within it) and publication tools (to provide a text description of an image). The diagram below shows an example of machine learning for computer vision.

Natural language processing (NLP)

The aim of NLP is for a machine to obtain meaning from text and speech. Humans can easily take information from written or spoken words as we understand the meaning of each word as well as the context in which they are delivered. However, for machines, this relatively basic task is not so easy. It does not understand human words or speech in the same way we do.



AI in more detail

The machine needs to process and analyse the text to identify patterns. Many NLP packages use a combination of Machine Learning, more simplistic text mining and statistical models to give the unstructured text a structure. This structured text data is then analysed to identify patterns and elements that fit specific parameters. Natural language processing is used in many devices such as smart speakers, to understand what you are asking for, and can also be used to identify characteristics of the text or speech such as sentiment.

Approaches that help computers generate text or images

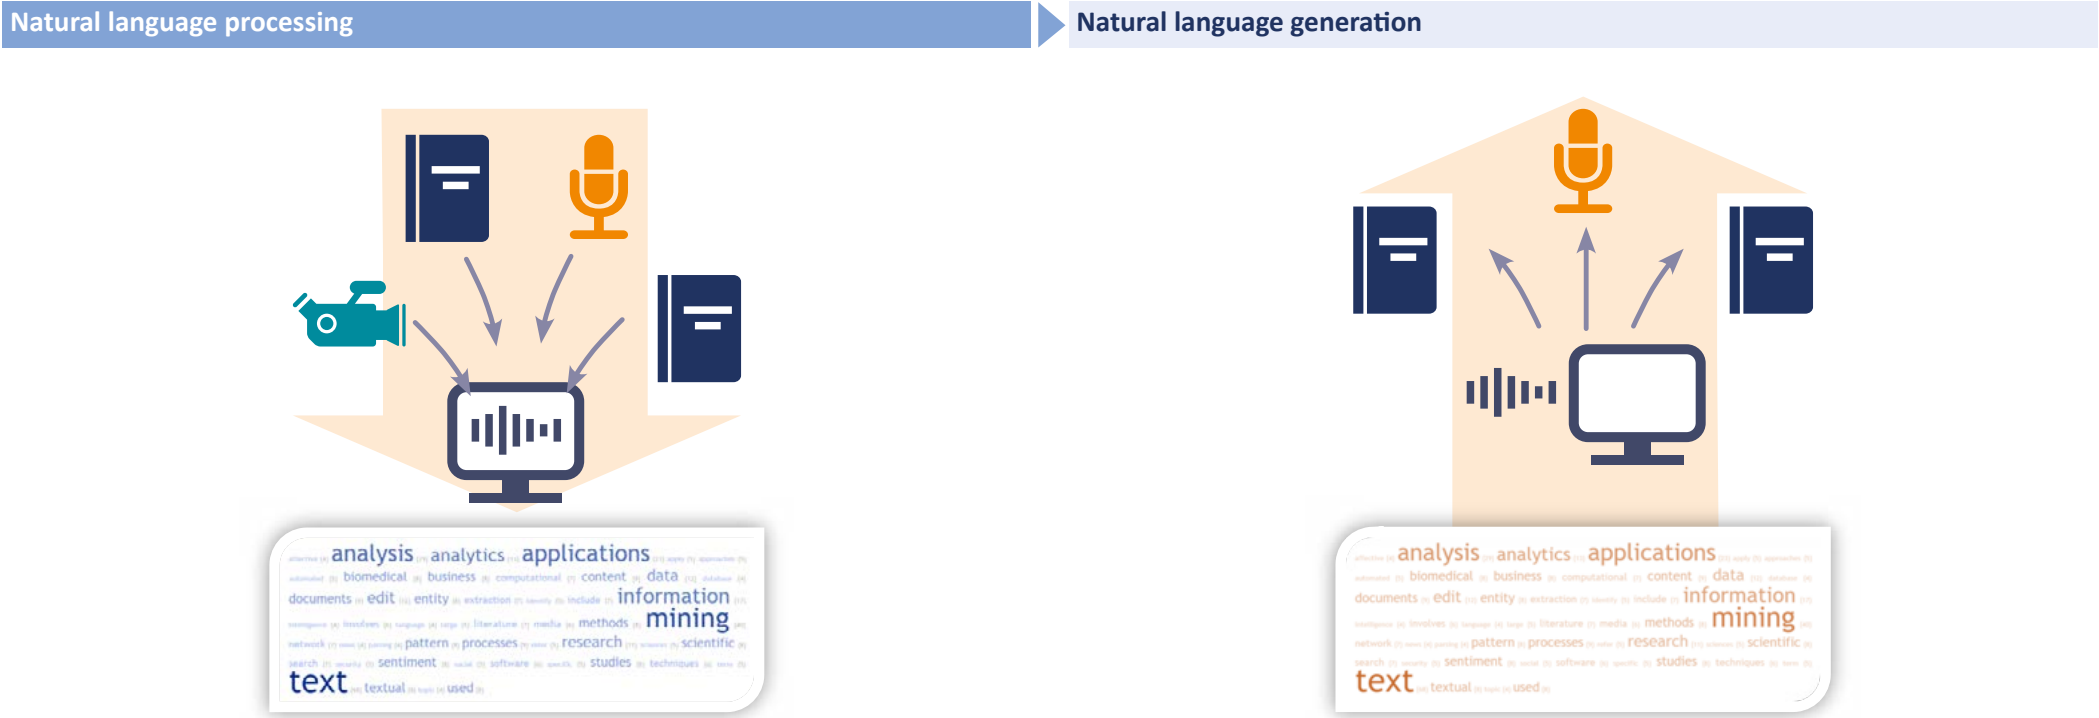
Given that there are many applications of AI that are looking to analyse and understand images and text, it is not surprising that there are also applications looking to create text and images. Again at their core, these often involve machine learning and natural language.

Natural language generation (NLG)

The aim of NLG is to translate data and understanding within a database or other repository and turn that into human understandable output; generally speech or text. This will often involve a standardised dictionary or series of template elements that the system can combine with the underlying data to create a meaningful output. Examples of NLG in practice include auto generation of sports reports from the underlying game statistics and weather forecasts to smart assistants or phones. A combination of NLP and NLG might also be used in tools such as chatbots.

Artificial creativity

Artificial creativity is an area right at the edge of what is currently possible with AI. The idea is to get AI to mimic the human creative process and create original music, art and images. Whilst this is currently some way off, techniques such as generative adversarial networks (basically thousands of variations of an image, one of which is then judged by the system as the closest to the desired outcome) have already been used to create realistic, simple images.



The system is fed unstructured text and speech and uses techniques to analyse meaning and identify characteristics.

The system uses its structured text database to create text and speech output that communicate meaning to a human, possibly as a response to a question.

AI in more detail

Approaches that help computers copy a task

Whilst many of the activities that AI can undertake are exciting, the world of business involves a significant amount of repetitive work. In order to add efficiency to the workplace many companies have considered techniques to facilitate automation:

Robotic Process Automation



Whilst Robotic Process Automation (RPA) may sound like it involves robots in factories, it is in fact a field that uses small software programs or 'bots' to undertake individual or linked parts of a process. The bots mimic human activity and can access and process information from a company's current systems. This makes bots particularly cost effective. Because the bot is directly interacting with the system it is much more efficient than a human operator (1 minute of bot work is equal to several minutes of a human worker).

Human Process



1



Receive invoice email from client.

2



Check invoice matches purchase order.

3



Attach invoice to payment order.

4



Email accounts the payment order.

5



Update the record in the accounting system.

RPA Process



RPA module automates the task accessing all the underlying systems. Processing throughput is greatly increased.

AI and ethics

While this section of the report has given some high level overviews of key AI and related technologies, it does not attempt to look at the many moral and ethical challenges that come from the use of AI. Much has been written about the ethics of AI, covering everything from bias in data to moral issues around replacing human employment.

The Lords' select committee on AI spent a significant amount of time considering both the wider opportunity for AI as well as some of the ethical and moral issues. Their report is available [here](#) and it develops some underlying rules for the ethical use of AI.

- *AI should be developed for the common good and benefit of humanity.*
- *AI should operate on principles of intelligibility and fairness.*
- *AI should not be used to diminish the data rights or privacy of individuals, families or communities.*
- *All citizens have the right to be educated to enable them to flourish mentally, emotionally and economically alongside AI.*
- *The autonomous power to hurt, destroy or deceive human beings should never be vested in AI.*

As well as general questions on ethics, there is also an interaction with the professional ethical standards of accountants and others where the applications interact with or replace work conducted by a professional. These are areas that have and are actively being considered by the relevant professional bodies.

Links

More on AI

- The World Economic Forum have a series of interesting articles and reports that cover the opportunities and risks from AI <https://www.weforum.org/agenda/archive/artificial-intelligence-and-robotics/>
- The OECD have undertaken a number of projects which cover some of the impacts of AI <http://www.oecd.org/going-digital/ai/>
- For an easy explanation of machine learning see the OxfordSparks video <https://www.oxfordsparks.ox.ac.uk/content/what-machine-learning>
- And for examples of RPA in action see this video from Deloitte <https://youtu.be/OT5plcRWUzU>
- An example of Google NLP/NLG assistant phoning real people <https://youtu.be/D5VN56jQMWM>

Government and AI

- The House of Lords have undertaken an extensive review of AI <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/10002.htm>

AI, Accounting and Audit

- ICAEW paper that considers how AI might impact the accountancy profession <https://www.icaew.com/technical/technology/artificial-intelligence/artificial-intelligence-the-future-of-accountancy>

What is the Lab?

Since 2011 the Financial Reporting Lab ('the Lab') has sought to improve the effectiveness of corporate reporting in the UK. It does this by working with companies, investors and others on topics that matter.

Lab reports explore innovative reporting solutions that better meet the needs of companies and investors, by speaking to them about a topic and publishing reports that represent their views. Lab reports do not form new reporting requirements, but do seek to highlight best practice and thought leadership. The Lab has published reports covering a wide range of reporting topics, including:




For more information about the difference the Lab makes to reporting watch our video:



<https://youtu.be/6L9UGyalNoY>

All of our published reports can be found on the FRC's website:

<https://www.frc.org.uk/Lab>

Follow us on  Twitter @FRCnews or  LinkedIn

Do you have suggestions or want to get involved?

The Lab encourages readers of this report to provide comments on its contents and get involved in upcoming Lab projects. To provide comments or get involved, please send us an email at:

FinancialReportingLab@frc.org.uk

The Financial Reporting Council (FRC) is the UK's independent regulator responsible for promoting transparency and integrity in business. The FRC sets the UK Corporate Governance and Stewardship Codes and UK standards for accounting and actuarial work; monitors and takes action to promote the quality of corporate reporting; and operates independent enforcement arrangements for accountants and actuaries. As the Competent Authority for audit in the UK the FRC sets auditing and ethical standards and monitors and enforces audit quality.

The FRC does not accept any liability for any loss, damage or costs howsoever arising, whether directly or indirectly, whether in contract, tort or otherwise from any action or decision taken (or not taken) as a result of any person relying on or otherwise using this document or arising from any omission from it.

© The Financial Reporting Council Limited 2019. The Financial Reporting Council Limited is a company limited by guarantee.