

AI FOR BUSINESS

05 WHAT UK BUSINESSES SHOULD KNOW ABOUT THE EU'S AI ACT

08 CAN AI BE TRAINED ON ITS OWN RESPONSES?

10 WHY INVESTORS WANT MORE AI DISCLOSURES



Everybody's ready for **AI**
except your data. informatica.com/ai

AI FOR BUSINESS

Distributed in
THE TIMES

Contributors

Cath Everett

A journalist with three decades of experience covering workplace and leadership issues, including what it means to be an ethical business.

Chris Stokel-Walker

An author and journalist specialising in technology and culture, with bylines in *The New York Times*, *The Guardian* and *Wired*.

Sean Hargrave

A freelance journalist covering topics such as financial services and digital marketing. He is a former *Sunday Times* innovation editor.

Mark Walsh

A New York-based freelance writer covering business, tech and media. His work has been featured in *The Guardian* and *New York Magazine*.

Raconteur

Special projects editor
Ian Deering

Contributing editor
Neil Cole

Campaign director
Oily Eyre

Commercial content editors
Laura Bithell
Brittany Golob

Commercial content executive
Jessica Lynn

Commercial production manager
Emily Walford

Production executive
Sabrina Severino

Design and illustration
Kellie Jerrard
Colm McDermott
Samuele Motta

Design director
Tim Whitlock

Certified



Although this publication is funded through advertising and sponsorship, all editorial is without bias and sponsored features are clearly labelled. For an upcoming schedule, partnership inquiries or feedback, please call +44 (0)20 3877 5800 or email info@raconteur.net.

Raconteur is a leading business media organisation and the 2022 PPA Business Media Brand of the Year. Our articles cover a wide range of topics, including technology, leadership, sustainability, workplace, marketing, supply chain and finance. Raconteur special reports are published exclusively in *The Times* and *The Sunday Times* as well as online at raconteur.net, where you can also find our wider journalism and sign up for our newsletters. The information contained in this publication has been obtained from sources the Proprietors believe to be correct. However, no legal liability can be accepted for any errors. No part of this publication may be reproduced without the prior consent of the Publisher. © Raconteur Media

@raconteur in raconteur-media @raconteur.stories

raconteur.net /ai-business-april-2024

LEADERSHIP

Is this the right moment to appoint a chief AI officer?

The rapid advance of generative AI has been bewildering. No wonder many business leaders are considering whether they'd benefit from the guidance of a strategic-level expert in the field

Cath Everett

In March, the White House announced that all federal agencies in the US would have to appoint a chief AI officer to strengthen their governance in respect of the technology. The mandate is expected to create about 100 such CAIOs by the end of May.

Should other enterprises follow suit? "It depends" is the answer that most experts will give.

One of them is Michael Queenan, founder and CEO of Nephos Technologies, a consultancy specialising in data services integrations. He notes that many S&P 250 companies are hiring, or talking about hiring, an AI chief of some description. But he compares this to an "emperor's new clothes" scenario, suggesting that firms are "often not giving enough thought" to why they really need one.

Their reasoning may be no more complex than "they don't want to be seen as the company that doesn't have one, lest they're asked why not at the next shareholder meeting or on CNN and their share price falls", Queenan explains.

The decision whether to hire an AI supremo or not should be based on how central the technology already is to the business. That's the view of Brian Peterson, co-founder and CTO of Dialpad, the creator of an AI-based customer intelligence platform.

"If AI is a big element of your business or you're building it into your product set, having a CAIO would provide focus. But, if it just seems cool and could be part of your future but you're not sure how yet, appointing one might not be right for you," he says, suggesting that it would make more sense in the latter scenario to hire a consultant first to assess the technology's potential value to the firm.

In any case, CAIOs are a scarce and costly commodity, reports Waseem Ali, CEO of Rockborne, a recruitment consultancy specialising in the data and AI sector.

"We're not seeing many on the market," he says, noting that they're mostly working in "sectors such as fintech and healthtech. Ecommerce companies and some insurance firms that are algorithmically driven have been hiring them too."

Ali also points to the Future of Work Report published by LinkedIn in November 2023. This indicated that the number of employers creating the less senior role of head of AI had more than tripled in five years.



Rowan Jordan via iStock

He has observed "more chief data officers than anyone else absorbing the AI remit to become chief data and AI officers, while some organisations are simply turning their CDOs into CAIOs. You don't see this conversion happen as much with CIOs or CTOs unless they have a data remit."

The absorption of roles makes sense to Queenan, who says: "Companies should absolutely get across AI, but most large ones already have the data science people and processes in place to do that. AI is an app that sits on top of your data, which means it's just another data product. So, if you already have a team creating such things, this is simply adding another string to their bow."

He believes that having "a head of AI who reports to the CIO or CTO is more than sufficient in most cases. In five years' time, there could be a real need for a powerful job title such as CAIO, but it's too early for it now."

Queenan's view is that organisations generally need more time to work out how to "do AI better" and decide whether they will benefit most from developing their own tech or buying off-the-shelf products. Most firms already seeking to hire a CAIO are "putting the cart too far in front of the horse", he argues.

Peterson agrees that granting an AI specialist a seat at the top table now would probably be overkill in most organisations.

"It depends on what expertise there already is on the board and

what value a CAIO could bring," he says. "But, if you're not a tech company and AI isn't core to your business, it probably isn't necessary."

This view is borne out in the wider recruitment market – AI has seen few C-level appointments for AI experts to date, although he reports that firms are getting more interested in finding non-executive directors with AI knowledge.

Anyone seeking to become an AI chief must demonstrate a range of top-level skills, including strategic thinking and effective communication, according to Ali. A CAIO will be able to manage the board's expectations about what the technology can and cannot do and explain likely outcomes in a language that its members can relate to. This includes expressing where AI tools could add value by reducing costs, for instance.

Equally important is the ability to track, understand and explain the evolving governance issues surrounding GenAI, including the ethical, reputational and regulatory risks it poses.

A final consideration, if your company is set on hiring an AI chief, is to "put your money where your mouth is" and equip the successful candidate adequately, Peterson argues.

"You can't just hire a CIAO, give them that big title with lots of expectations and leave them to it," he warns. "You need to support them by putting money, resources and prioritisation behind it. Otherwise, you'll be setting them up to fail." ●



Foundry, 2023

Using human intelligence to mitigate artificial intelligence risks

AI is transforming how organisations operate, but significant data protection challenges must be overcome

It is no secret that artificial intelligence is transforming the world of work. Already employees use a plethora of AI assistants to streamline everyday tasks, such as writing emails, developing code, crafting marketing strategies and even managing company finances. The trend is set to accelerate as the technology develops, yielding huge productivity benefits for organisations.

Yet as thousands of new AI-enabled applications are launched each week, many of them free to use, there are growing concerns about the data protection risks.

Many organisations have no idea what AI apps and services are being used by their staff, or for what purpose. They are also unaware of what data is being shared and with whom, or how it is being managed and protected.

This heightens the risk of data breaches that come with significant financial and reputational costs. Additionally, there is a real possibility that organisations may be feeding AI tools with sensitive corporate information without realising, contributing to the training of potentially competitive AI models.

So how can firms reap the benefits of AI while mitigating against the risks?

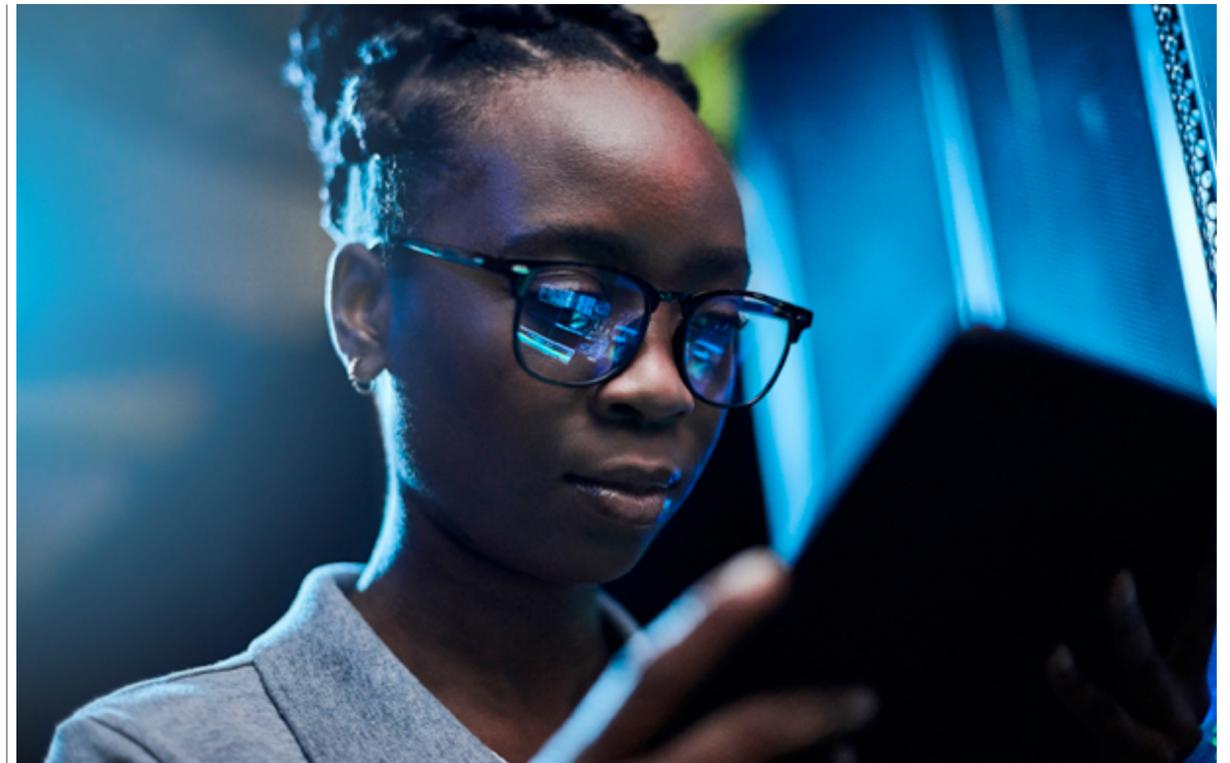
Data protection is non-negotiable

Neil Thacker is chief information security officer EMEA at Netskope, a secure access service edge (SASE) provider that helps organisations around the world to prevent data loss, leakage and misuse. He says the arrival of AI is much like the advent of cloud computing or even the internet, with companies still scrambling to understand the technology and its risks.

"This comes as data regulation is being tightened up around the world, making the safeguarding of sensitive data non-negotiable for every business," Thacker says. The EU's existing GDPR rules and new AI Act, which is set to come into force over the next few years, are cases in point.

At the same time corporate use of AI-enabled apps is accelerating rapidly. According to Netskope Threat Labs' Cloud & Threat Report 2023, organisations of 10,000 staff or more accessed at least five generative AI apps daily last year, with ChatGPT, Microsoft Co-pilot and GitHub Copilot being among the most commonly used.

The algorithms that power these platforms develop and improve based on the data fed into them, which raises myriad copyright and intellectual property issues. For example, last year source code was being posted to the most popular generative AI app, ChatGPT, at a concerning rate of 158 incidents per month in 2023, according to Netskope research.



“Without realising it you are helping train even smarter AI platforms that can help your competitors”

"If firms are not careful they could leave sensitive data such as proprietary IP, source code or financial information accessible to competitors. Without realising it you are helping train even smarter AI platforms that can help your competitors," Thacker says. "The risk is immediate too. It used to take years to train powerful new algorithms but these days it can be done in a matter of days and weeks."

Private AI?

Thacker says firms must deploy continuous data protection policies and tools to protect themselves. Chief information security officers (CISOs) should make an inventory of all the AI services in use across their organisation, identifying those that are truly relevant to the company.

They then need to vet each platform vendor and assess its data policies, including whether it relies on third- or fourth-party support.

"There are significant costs associated with AI technology, so it's obvious that free or inexpensive options make their money in other ways – by selling data or the AI intelligence that it has contributed towards," says Thacker. "In such cases, a thorough examination of

Commercial feature

It plugs seamlessly into cloud services, flagging the risks associated with more than 85,000 cloud apps and services including AI apps. Powered by AI itself, it learns how to recognise sensitive data based on an organisation's preferences and identify it in real time.

When a risk is detected it issues a pop-up message telling the employee the risk level of the app they are using on a scale from 0-100.

"We base the score on 50 variables, including the security controls that platform has in place, its privacy policy,

where any data is being processed and the regulatory challenges, and any other potential legal liability issues," says Thacker. "If an app is high risk the employee can make a call on whether to use it depending on the sensitivity of the data involved. Netskope may also be able to offer them an alternative that is more secure for the organisation."

Research has shown this behavioural approach to data security is highly effective, given that a staggering 95% of cybersecurity incidents stem from human error. Continuously training people using point-in-time warnings is highly effective the same way reinforced training is used in AI models; "I

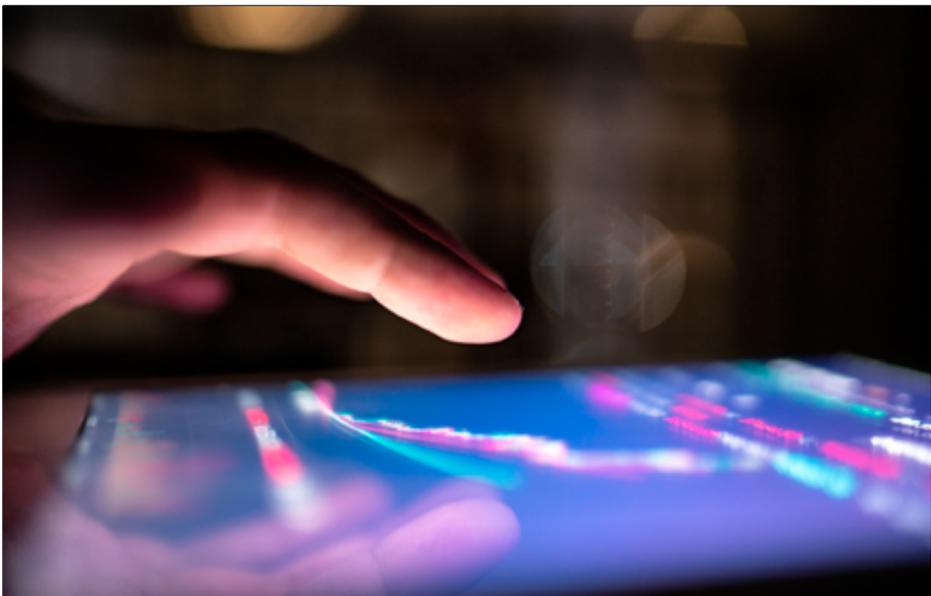
use the analogy of radar speed cameras that tell you your speed," says Thacker. "Once you are reminded how fast you are going and the consequences, you slow down. It's about point-in-time awareness of the risks."

Founded in 2012 Netskope has become a leader in the SASE space, offering unrivaled visibility, real-time data and threat protection for cloud services, websites and private apps. Known for its data and threat protection, the US company is now leading the way in the AI security space globally.

"As the digital transformation of companies continues, AI will offer enormous benefits in terms of enhanced efficiency, competitiveness and end-user experiences," says Thacker. "But it has also become the frontline in the fight to protect data, and organisations that do not adapt to the evolving threat landscape could pay a high price."

For more, please visit [netskope.com](https://www.netskope.com)





Securing the future: data integrity in the age of generative AI

Businesses can only unlock the benefits of artificial intelligence if they tackle their long-standing data management issues

Artificial intelligence is set to transform the way companies operate over the next decade, making workers more productive, improving customer service and offering firms invaluable insights on their operations.

Yet despite the huge competitive advantages AI offers, many early adopters have not achieved the results they hoped for, while others have found it hard to adopt these systems at scale.

Typically, data management issues are to blame, as organisations struggle to access the high-quality data needed to power the AI algorithms supporting their operations. Poor data input leads to bad outcomes at scale, as using poor-quality, incomplete or untrusted data as a foundation for AI assistants results in inaccurate or biased decisions that are of no help to firms, and may even hinder them.

Poor data management could also create compliance problems, as organisations lose track of the data driving their AI platforms, putting themselves at risk of breaching incoming AI regulations.

So how can organisations get a grip on their data today and fully reap the benefits of the AI revolution?

No silver bullet

Greg Hanson is GVP and head of EMEA North for Informatica, a leading cloud data management provider that helps businesses handle the complex challenges of dispersed and fragmented data to innovate with their data and AI.

"Technology forms a major part of the solution," says Hanson. "But organisations also need a data management strategy and cultural change which involves sponsorship at board level, engagement of people and the establishment of governance polices."

This is why Informatica works closely with leading advisory organisation Cognizant, a global strategic alliance partner of Informatica, that helps firms embed the tech, teams and processes for successful AI adoption. Making the most of data is a theme that both Informatica and Cognizant are witnessing among customers, says Sean Heshmat, GGM data and AI head at Cognizant.

He adds: "without the right input you will simply make incorrect decisions at an accelerated pace. Firms need to build the right foundations to ensure AI works for them, not against them."

It's a sentiment that Hanson agrees with: "Many organisations believe AI will be an overnight silver bullet but there is a significant amount of foundational work required to benefit from this technology. That's because when it comes to data, the old adage applies – if you put garbage in, you will get garbage out."

For any business adopting AI at scale, the first task is to corral all the data it has in one place so it can be processed and accessed with ease. But this can be challenging as large firms typically have multiple divisions, servers and systems in place around the world and data is often siloed.

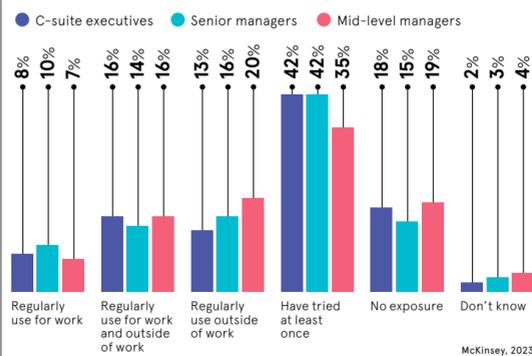
To counter this they must simplify their data landscape, standardise the tech they use and deploy an effective data catalogue to organise and manage data assets properly. "A company can't get a proper picture of their customers or operations if their datasets are incomplete or disorganised," says Heshmat. "Similarly, AI can't make quality decisions in real-time without real-time data, and that is hard to achieve with myriad different systems and integration points."

Quality control

According to *The State of AI in 2023* McKinsey survey, inaccuracy is the biggest risk companies face when it comes to AI. Yet just 32% said they were mitigating that threat of inaccurate data and inaccurate outcomes.

“Organisations need a data-management strategy and cultural change, which involves sponsorship at board level, engagement of people and the establishment of governance polices

MOST DECISION-MAKERS SAY THEY ARE USING GENERATIVE AI IN SOME CAPACITY



even lower than the 38% who said they mitigated cybersecurity risks.

As such, it is vital that organisations have high-quality data to power their AI platforms, although identifying, verifying and extracting this information can be challenging. Staff also need to be able to access data with ease, while establishing robust data principles to ensure regulatory compliance.

Firms have already had to adapt to the EU's GDPR rules, and over the next few years the EU's AI Act will come into force, requiring companies to demonstrate they have full oversight of the data going into their AI platforms, with breaches leading to significant fines.

Informatica's solutions offer organisations vital support as they prepare for AI adoption. The firm's AI cloud platform enables them to manage and organise all their data with ease via one unified platform that breaks down silos.

It also lets users locate, extract and cleanse data to develop first-class algorithms, while supporting good data governance by recording the data trails and providing data lineage visualisation sitting behind automated decisions, simplifying compliance.

Democratising data

Cognizant deploys Informatica's solutions as part of its wider work supporting organisations' digital transformations. It acts as a trusted partner to companies, helping them to change their data culture and processes and get the most of data management and AI systems.

"Together we help organisations democratise data and bring it to life," says Hanson. "This helps to make it more easily accessible to those in the company that need it – subject to data access controls. With Informatica, teams no longer need to ask IT for the information they require to make more informed business decisions, it is self-serve and ready to use."

Gilead Sciences knows first-hand why good data governance is essential to business success. The global biopharmaceutical company worked jointly with Informatica and Cognizant to bring more value to customers by getting more out of the data the firm had amassed through the manufacture and development of advanced treatments.

Gilead wanted to improve its master data management processes and compliance controls, while bringing data into the hands of employees who

needed it. It deployed a data mesh framework on Amazon Web Services, supported by Informatica's AI-powered cloud platform which provides useful, holistic data to decision makers.

As a result, Gilead was able to speed up its drug development, discovery and commercialisation processes and bring down costs.

"To us, a cloud-based enterprise data platform is not just about cost or operational efficiencies. For us, it's a competitive differentiation in the industry, we can make better, faster decisions about our business," says Murali Vridhachalam, head of cloud, data & analytics at Gilead.

Cognizant and Informatica have partnered to help many global brands deliver greater value to their customers. One such brand is BMW, which worked with both Informatica and Cognizant to implement a unified platform for global product data that provides a trusted, omnichannel view of critical information.

The Informatica system enables the German manufacturer to deliver consistent comms globally and one that is helping to power a next-generation customer experience capable of leveraging new technologies like AI. BMW finds itself in the midst of an AI journey as it seeks – like many other businesses – to unlock the huge potential benefits of AI, but to do so requires organisations to resolve their long-standing data management issues.

"Adding that good data governance is an end-to-end process, not a one-shot deal, says Heshmat. "It's about having a data-driven culture, the right tech, and proper communication between your board, the business and IT to make sure data is treasured and protected across your organisation."

Hanson agrees, "AI requires holistic, trusted and governed data for companies to succeed with correct, unbiased insights. Our goal is to help firms unlock the power of AI and bring their data to life."

To find out more please visit informatica.com/gb



REGULATION

What UK firms need to know about the EU's Artificial Intelligence Act

This landmark legislation was recently approved by the European Parliament and will come into force gradually over the next two years. How will it affect businesses beyond the bloc?

Chris Stokel-Walker

Ever since ChatGPT shook up the business world in Q4 2022, firms have been racing to use AI, but regulators are catching up fast in their bid to ensure that any such application is safe and trustworthy.

In March, for instance, the European Parliament signed off the EU Artificial Intelligence Act. Pending final checks, the legislation should be adopted before the parliamentary election in June, with its provisions taking effect in stages over 24 months. It amounts to the world's first major set of statutory standards governing the use of AI.

"With the growing presence of AI in all aspects of daily life, legal frameworks are urgently needed to regulate its uses and protect data," says Neil Thacker, CISO at cybersecurity firm Netskope in EMEA.

He adds that one of the main objectives of the new legislation is to "strike the right balance of enabling innovation while respecting ethical principles". As part of this effort, the act splits AI systems into different risk categories governed by requirements of varying stringency.

It will also apply to any system that touches, or otherwise interacts with, consumers in the EU. That means it could have a broad extraterritorial impact. A British company using AI to analyse data that's then sent to a European client, for instance, would be covered by the legislation.

“The AI Act is wide-ranging, trying to provide guidance and protection across the multitude of areas that AI will affect



"The act is wide-ranging, trying to provide guidance and protection across the multitude of areas that AI will affect in the coming years," Thacker says.

The main concern for UK business leaders is how onerous the new law is likely to be for their firms. For many, the EU's previous big statutory intervention – the General Data Protection Regulation – has cast a long shadow since taking effect in 2018. Remembering the paperwork this required and the many changes they had to make to ensure compliance, they're understandably worried that the new legislation could impose similar bureaucratic burdens, which might prove costly.

Fear not, says Michael Veale, associate professor at University College London's faculty of laws, who has been poring over its small print.

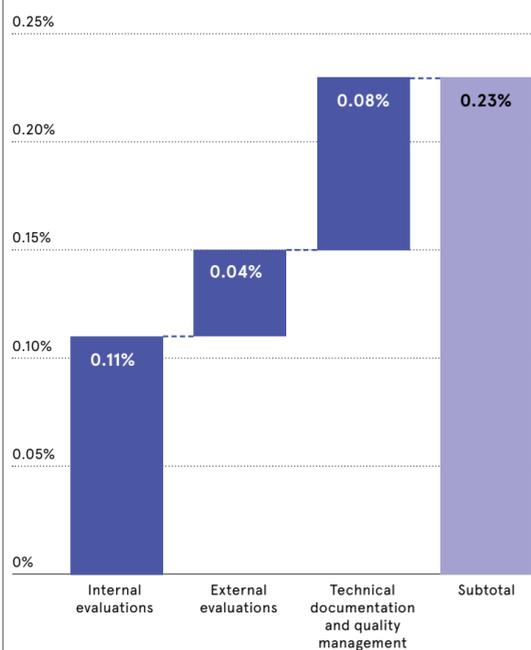
Many of its provisions are "quite straightforward and imaginable", he says. These include "making sure that your system is secure and not biased in ways that are undesirable, and that any human overseeing it can do so robustly".

Such requirements shouldn't be too taxing, according to Veale. "They echo a variety of the very basic demands on AI systems in recent years," he explains. "While it may be difficult to interpret them in every single context, they aren't particularly onerous or revolutionary."

One of the most fundamental

THE COST OF COMPLIANCE

Projected compliance costs for GPT-4 under the act, as a share of the total investment required to develop an AI system



The Future Society, 2023

questions for any UK firm to ask itself is whether it's selling high-risk systems into the EU, says Veale, who notes that the vast majority won't be. In any case, the few that are "should be looking at the standards and making sure they're following them anyway".

The EU won't be assessing firms and certifying them as compliant, so third-party industry-led standards bodies will likely self-police, with the regulators stepping in only if needed, he adds.

There are certain aspects of the act that "average non-specialist businesses should know", so that they can take steps to ensure compliance, according to Thacker.

"Initially, they should heed its references to general-purpose AI systems," he advises. "The new law includes transparency requirements including technical documentation and compliance with EU copyright laws. Where such information is not available, businesses will be required to control the internal use of such systems."

Thacker points out that the legislation includes explicit requirements for detailed summaries about the content used in training any general-purpose AI systems.

Companies specialising in areas that the legislation deems "high risk" will need to be particularly attentive to its terms. That's not only because of the more stringent requirements that will apply to them. It's also because they'll have less time to ensure compliance. While most organisations will have two years to implement any required changes, the deadline is tighter for makers of high-risk systems.

Most of the applications identified as high risk by the act are those that public sector organisations would use for purposes such as education, the management of critical infrastructure or the allocation of emergency services.

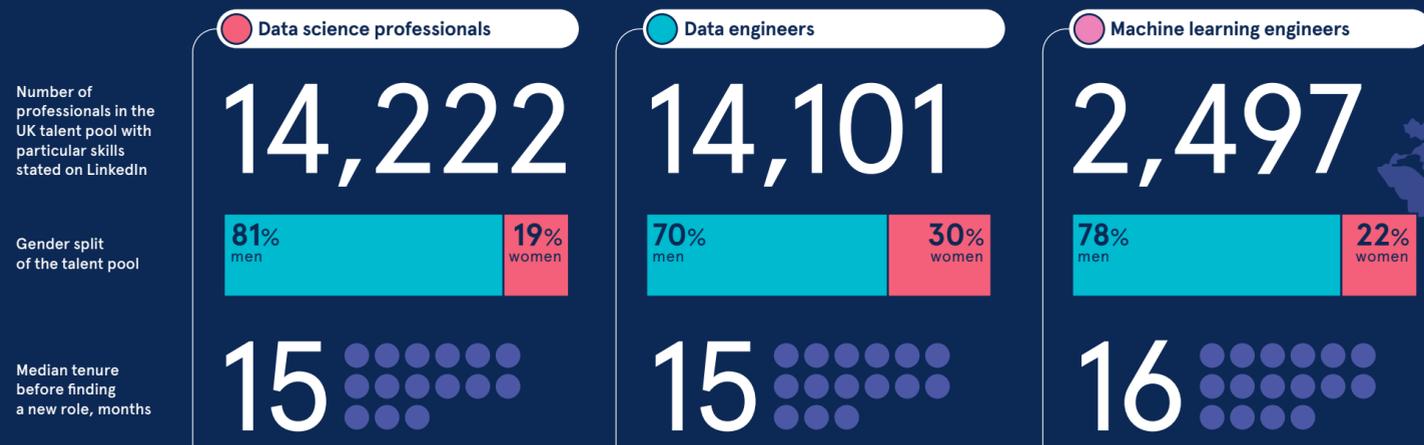
Any UK firm selling AI products for such purposes would need to register these in a centralised database and undergo the same certification process that applies to any EU counterpart.

Beyond that, all businesses would be wise to audit their systems and use of AI more regularly and thoroughly. This should help them to prepare for any further statutory changes in this fast-moving field.

The EU's act is the first legislative effort of note to lasso a constantly evolving technology that, in its current form, is barely 18 months old. The situation could easily change radically long before this law's final provisions are due to take effect. It's therefore vital for businesses to keep abreast of AI developments as a matter of course, Thacker stresses. ●

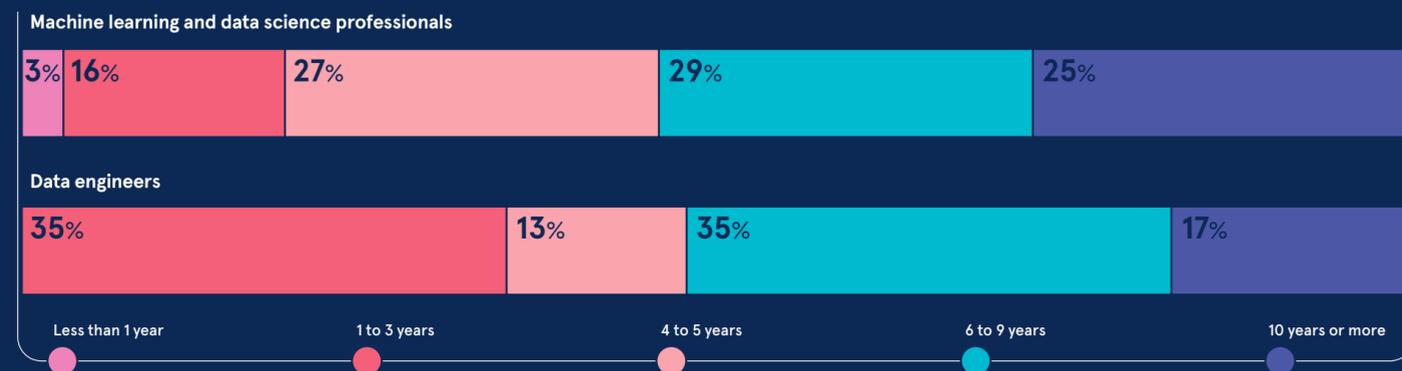
THE AI TALENT POOL

Artificial intelligence is expected to be a driving force for business innovation and efficiency. But to get the most out of the technology, firms will require certain technical expertise in their workforce. Machine learning and data engineers, along with data scientists, are indispensable for organisations seeking to develop and deploy AI tools at scale. Where should UK firms recruit for these skills? And, what should they expect from candidates for AI roles?



MORE THAN HALF OF DATA AND MACHINE LEARNING PROFESSIONALS HAVE AT LEAST SIX YEARS OF EXPERIENCE

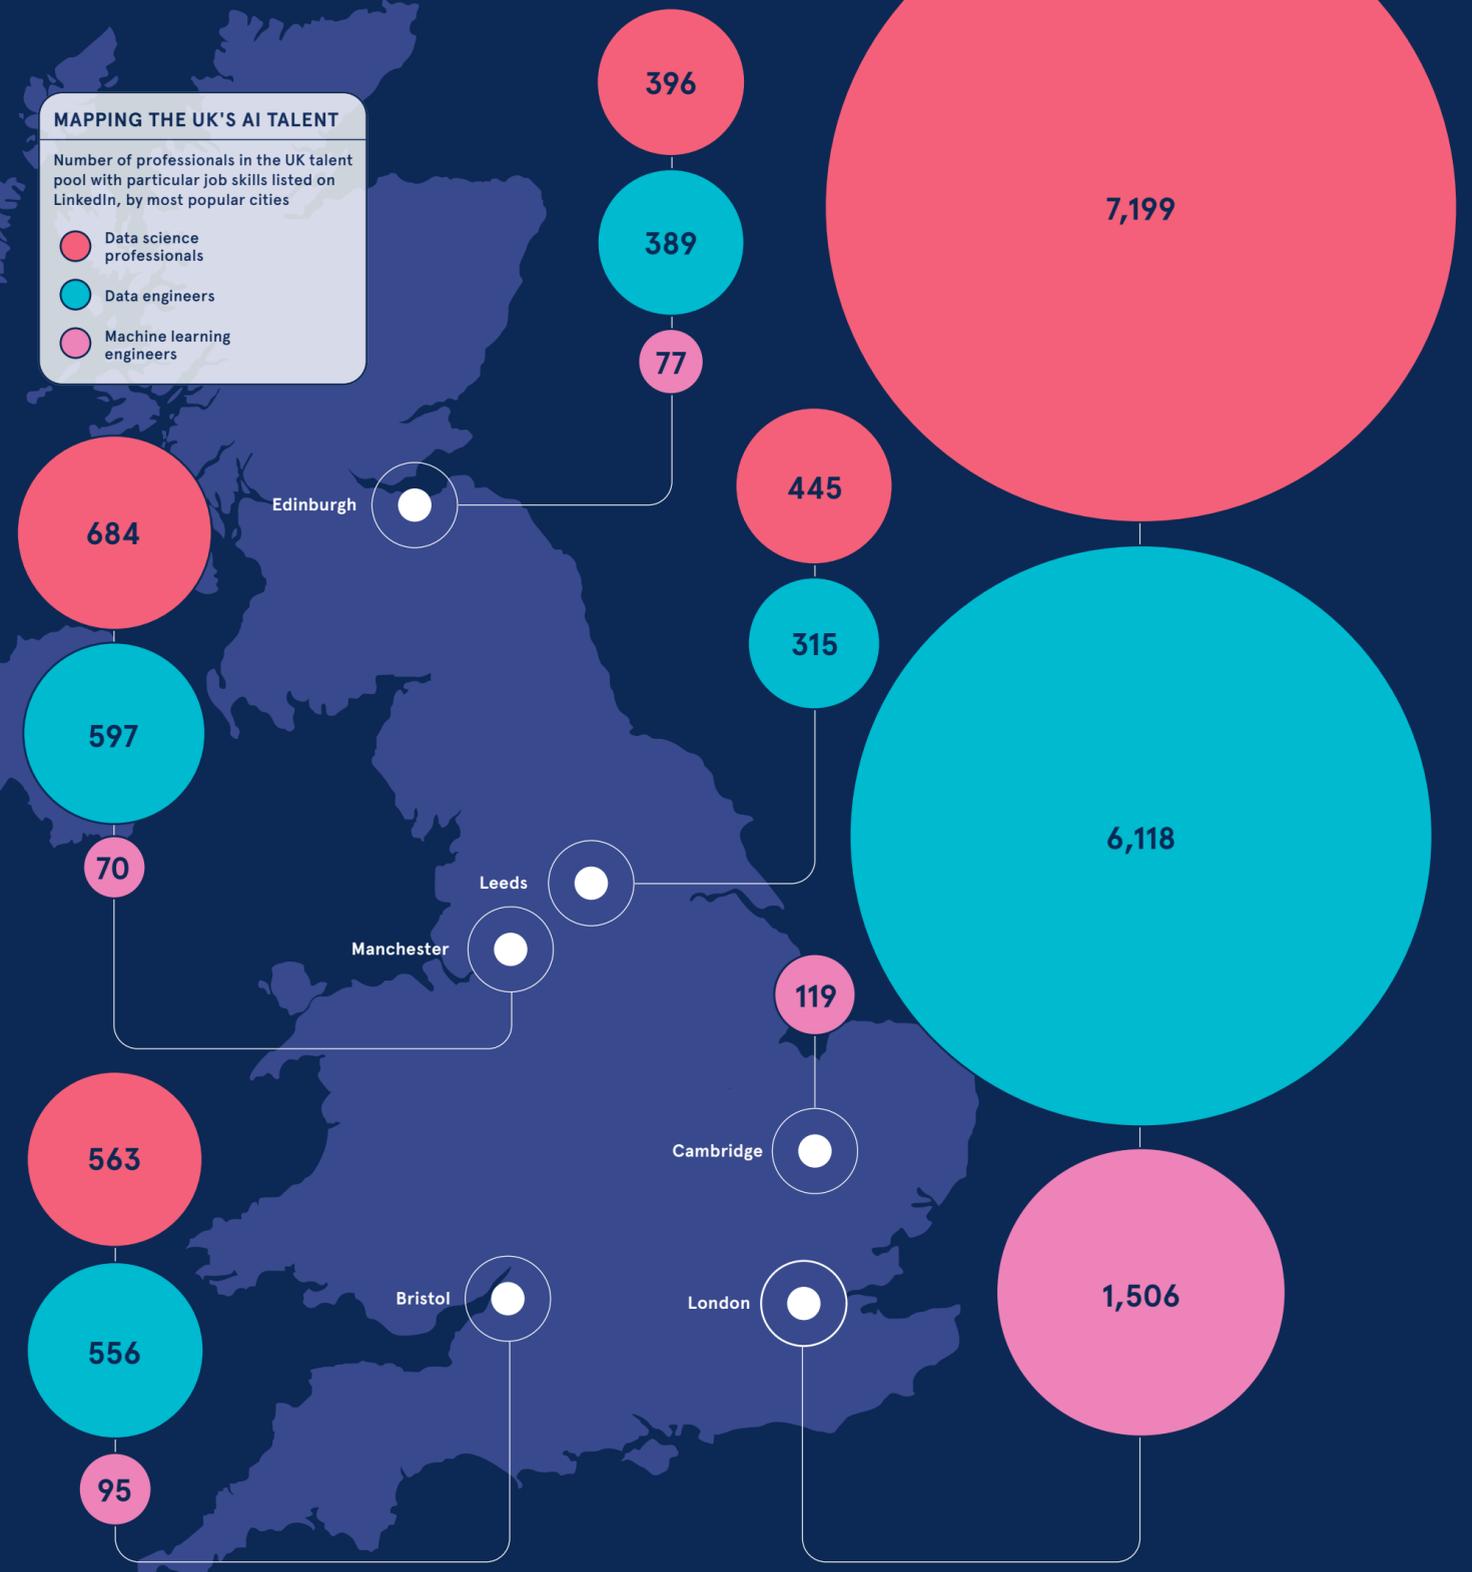
Years of experience in the talent pool of professionals with particular data and AI-related skills



MAPPING THE UK'S AI TALENT

Number of professionals in the UK talent pool with particular job skills listed on LinkedIn, by most popular cities

- Data science professionals
- Data engineers
- Machine learning engineers





TECHNOLOGY

How tech might tick the right survey boxes

Not content with using AI to analyse feedback from consumer polls, market researchers hope it could eliminate human respondents from the process by generating the same replies they'd have given

Sean Hargrave

Market-research companies have found AI to be a useful analytical tool, particularly its ability to understand what consumers write on questionnaires and say in audio or video interviews. The technology can also reliably interpret their answers to reveal hidden insights. It can even suggest next steps.

But in the next wave of adoption, market researchers will test AI's ability to use synthetic responses of its own devising, effectively cutting human interviewees out of the equation. If their experiments prove successful, AI could provide near-instant low-cost 'consumer' insights, reducing the need to conduct costly surveys and, potentially, enabling

brands to more efficiently reach lucrative niche markets.

To produce reliable responses, the technology must be able to understand the views of the target audience and provide results that match those elicited by traditional consumer research methods. Naturally, the question at this stage of development is: can the synthetic data it will produce be trusted?

Market researchers at Kantar have taken the first steps in answering this. They prepared a set of questions and compared real data drawn from human surveys with responses given by OpenAI's GPT-4 large language model (LLM). The queries they used covered a wide range of matters, such as whether the price of

luxury holidays is off-putting and whether a given piece of technology helps the owner to connect with people who share their interests.

When asked about more practical issues, GPT-4 gave similar answers to those provided by the human respondents. However, the more nuanced questions, requiring greater emotional reflection, produced significant differences.

Such results are what you might intuitively expect, notes Jon Puleston, vice-president of innovation at Kantar's profiles division. AI is good for some parts of market research, but it's limited if asked to adopt the persona of diverse human audiences.

"It's clear that there are risks to relying solely on synthetic data if you're making a business decision that's worth billions," he says. "Real human insights still form the heart of good market research. A more realistic use case for synthetic data is as a tool to complement, rather than replace, traditional research – for instance, by boosting sample sizes in surveys, particularly for niche audiences."

The experiments' results so far indicate that the LLM's outputs are only as good as the human-profiling data fed into it, notes Marius Claudy, associate professor of marketing at University College Dublin, who has

been researching the impact of training on AI outcomes.

While the technology can provide a good analysis of qualitative research, such as understanding what someone has said or written, it's less effective at understanding the emotions that underpin people's responses. This leaves the notion that AI could ever make traditional market research obsolete open to question.

"The issue will always be how meaningful the results are, particularly when you're asking about unknown propositions, such as a product that has yet to launch," says Gary Topiol, managing director at market research firm QuestDIY. "Getting responses will be fairly easy but, as with all new methods, understanding when they can be trusted will take time."

“Businesses will need to ensure they have a clear legal basis for uploading any personally identifiable data to AI tools

Another concern is the well-documented bias to which GenAI is susceptible – again, as a result of its training. For instance, researchers at Harvard have found that ChatGPT's views and values are closely aligned with those of US citizens.

Claudy points out that "the more distant a country is from the US culturally, the lower the correspondence between the human responses there and ChatGPT's. An LLM may be able to approximate the responses of the 'average' person on historical topics, but it might struggle to mimic the responses of certain subgroups or minorities accurately."

The first concern about AI models is that they are programmed to pick up views from the internet. This can engender a Western, English-language bias and create an echo chamber too. While that's a worry for Jeremie Brecheisen, managing partner of Gallup's EMEA division, he thinks there's an even bigger issue.

As every market researcher knows, consumers don't always make the logical choices a computer would expect of them. For instance, we often buy goods based on a whim, rather than a logical assessment of their attributes and overall value. This is why it is important to ask real people questions that cover a range of emotional responses. It's the answers to these questions that AI will struggle to mimic for the foreseeable future, he says.

"Our brains and emotions are highly complex, so it will require a lot of experimentation to understand whether AI can get close to replicating the results of human surveys," Brecheisen says. "There's a lot of interest in using synthetic data to cut costs, but that's not a great reason when you don't know if you can trust those answers."

Concerns are not limited to whether future models can replicate real

“A more realistic use case for synthetic data is as a tool to complement, rather than replace, traditional research

human responses. There are also legal considerations, warns Ben Travers, a partner specialising in IT matters at law firm Knights. While he shares researchers' worries that AI bias may lead to poor outcomes, he is also troubled by the use of personal data found on the internet to build profiles.

"Businesses will need to ensure they have a clear legal basis for uploading any personally identifiable data to AI tools," he says. "And all AI users must be alert to copyright issues. These apply to both the content fed into an AI and the content it produces. Just because this material is easily accessible does not mean that it's lawful to copy it. Such content is not 'fair game' – copyright will enable the rights owner to control how it is used."

The future of AI in market research is unclear. While the technology is undoubtedly a boon to those compiling surveys and interpreting responses, it remains to be seen whether it can reliably answer questions itself. The ultimate prize of having a system that can accurately predict which car will sell best among millennials in Peru, say, or how much sugar to remove from a soda for the Hungarian market seems to be the stuff of science fiction – for now, at least. ●

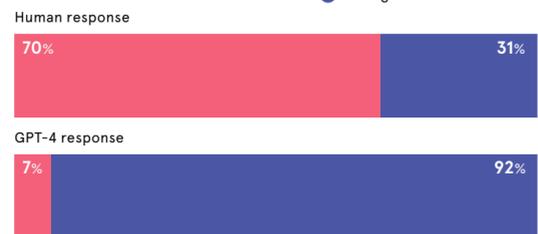
AI DOES BEST WITH MORE PRACTICAL QUESTIONS

Human and AI-generated responses to particular product-related questions

What is the importance of price when purchasing a luxury item?
 ● Net unimportant, or neutral
 ● Net important



Products are a way for me to connect with others who share my passion
 ● Net disagree, or neutral
 ● Net agree



*Owing to rounding, responses do not total 100%

Gartner, 2022

Mastering CX: how AI can improve customer service

Consumer frustrations with long call-waiting times and unresponsive chatbots are harming customer-service metrics. But generative AI can transform how answers are delivered

Customer service is no stranger to the rise of automation, from having to choose numbered call centre options to asking online chatbots for advice.

While many consumers cite frustrations with such systems that haven't quite got it right, increasing numbers of incoming queries have left companies reliant on chatbots to reduce long waiting times for customers and employees alike.

For Girish Mathrubootham, CEO and founder of Freshworks, a company that creates AI-boosted business software, there is a better solution available – embedding generative AI to deliver quicker, simpler and more seamless customer service, without sacrificing user experience.

"If a business scales from serving thousands of customers to millions of customers, it is not sustainable to keep hiring more and more people to deal with level one customer service," he explains. Enquiries regarding order tracking or cancellation are simple and can easily be dealt with by generative AI.

"Businesses have always driven automation through self-help, but generative AI is a significant leap in what can be accomplished in customer service," he explains.

More accurate responses

Mathrubootham suggests consumers aren't bothered if problems are solved by AI or humans, as long as they get the right answer, and fast.

Through generative AI, automated customer service can now handle queries in near real-time, in multiple languages and through a two-way conversation, even when the discussion is complex. It can also decipher and understand audio, images or videos.

"The biggest achievement," Mathrubootham says, "is having a multi-turn conversation with follow-up questions. This powerful technology produces a faster, more accurate and personalised response."

A major timesaver for human colleagues

Generative AI's role within customer service is to act as a co-pilot, working alongside humans to make their day-to-day tasks easier. This frees up human agents' time to solve more complex problems or helps them complete administrative tasks more efficiently, such as locating information from huge digital knowledge bases.



It can also proactively provide quality control on replies to customers, monitoring outbound messages and suggesting better responses.

"The AI might say 'this doesn't look like the right answer,'" Mathrubootham explains. "It can also detect if the tone isn't professional or courteous, raising prompts to rephrase. This feature is particularly useful for training new employees."

For generative AI to perform at its best, companies must initially feed in the right training data, set strong guardrails on the language used and ensure systems are secure.

One solution is to use the best and brightest human customer service agents in an organisation to train models, with these colleagues continually testing and refining automated responses until they sound human.

"If you put garbage in, you get garbage out," Mathrubootham warns. "When you're training machines with data that is not accurate, machines can pick the wrong answers. You can't just feed them a million customer service calls or a million response tickets from the past."

The route to strong business metrics

Customer happiness is critical to long-term success and embedding generative AI within customer service is now key to achieving satisfaction at scale.

For example, companies with high levels of unstructured data, such as customer satisfaction surveys or email queries, can quickly use it to pinpoint what support enquiries are most

common. This allows organisations to highlight nuances from customer replies to clearly show why they are satisfied or disillusioned and increase productivity by dealing with simple queries in large batches.

Implementing technology from a trusted partner can prove the best value. Freshworks' Freddy AI, for example, uses its existing data training sets alongside a company's internal databases, creating a tailored customer experience.

"When we talk about empathy and a human experience, some of this is readily available with generative AI," says Mathrubootham. "Customer service leaders must accept this is a big opportunity for their metrics, improving average handling and resolution times."

Mathrubootham explains that over the past three or four decades, businesses had to use humans to structure data in CRM or helpdesk systems to access the best insights, a process which lacks efficiency and limits productivity.

Now, generative AI is able to automatically complete these tasks in real-time, fundamentally changing the rules of the game. To unlock the benefits, businesses must pay attention to what their customers care about most and invest in a comprehensive generative AI partner.

For more information please visit freshworks.com



RISK

Black-box blues: investors take big tech to task on AI opacity

As companies expand their use of AI, concerned shareholders are pressing them to become more open about how they're using the technology and what safeguards they have in place

Mark Walsh

Little more than a year after ChatGPT made its seismic impact on business and wider society, questions about the safety of AI have become a pressing issue for investors in the 2024 proxy season.

Several AI-related shareholder proposals have been prompted by growing concern about the risks that rapid advances in this field pose to core institutions and the fundamentals of democracy and human rights.

The nearly 20 proposals submitted since late last year have mainly been aimed at companies ushering in the age of AI, including Alphabet (Google), Amazon, Apple, Microsoft and Meta Platforms. The signatories are seeking greater transparency regarding how the technology is being applied at those companies, as well as the disclosure of ethical guidelines governing its use.

While these proposals have tended to come from the investment community's more socially focused members, their concern about the ramifications of AI usage reflects a sentiment shared among investors more broadly, according to finance and governance experts.

Courteney Keatinge is senior director of ESG research at Glass Lewis, a proxy advisory firm. She summarises the situation as "just a matter of investors getting a better understanding of

how companies are using AI and companies being better able to communicate how they're using it".

That's easier said than done, of course. Companies don't seem keen to meet investors' demands by expanding on the voluntary disclosures that some have already made. But, given the growing societal pressure on big tech for greater openness in this respect, more formal reporting on AI-based activity is likely.

Investors aren't alone in calling for better governance and more transparency. A range of authorities are seeking to create standards covering the use and development of AI.

The most sweeping of these so far is the EU's Artificial Intelligence Act, approved in March by the European Parliament (see p4). This legislation aims to ensure fundamental rights and safety relating to AI systems. It will apply to any AI-based tool marketed in the EU, regardless of its creator's location.

In February, the UK government published its long-awaited plan for regulating AI. Built on core principles including transparency and accountability, the plan does not mandate legislation.

This past October, President Biden issued an executive order tasking federal agencies with the creation of guidelines for the use of AI. Scores of related bills are pending in the US Congress.



"I'm quite startled by how rapidly this is moving," says Heidi Welsh, executive director of the not-for-profit Sustainable Investments Institute in Washington DC, which tracks ESG-related proposals. "Usually with corporate responsibility issues, things kick around for a couple of years and then a policy slowly emerges."

Yet that's still probably not fast enough for some, including the AFL-CIO. The US trade union federation has adopted shareholder activism as a way to check the proliferation of AI. It has submitted half a dozen proposals seeking disclosures and ethical guidelines from the likes of Netflix, Walt Disney and Warner Bros Discovery.

The role of AI in film and TV production emerged as a contentious issue in last year's labour dispute between creative unions such as the Writers Guild of America and Hollywood's big studios. While the final settlements included protections for workers, the stir caused by the recent release of Sora, OpenAI's text-to-video tool, suggests that industrial strife concerning AI's role in the creative process may well recur.

Another focus of the recent AI proposals is the technology's potential for amplifying misinformation and disinformation, posing a threat to democracies around the world,

“We can expect to see increased regulatory scrutiny and, most likely over time, disclosure standards and requirements

downplayed the disbandment of its ethics and society team last year, noting that it still had nearly 350 people working to ensure responsible developments in AI.

Its reply broadly reflects those of other firms that have received AI-related proposals. The general message is that they already have adequate safeguards in place to ensure AI safety and are complying with recent government initiatives in this area.

Two AI proposals have come up for a vote at annual shareholder meetings so far. The AFL-CIO's call for ethics disclosures at Apple drew support from 37.5% of investors. At Microsoft, meanwhile, 21.2% backed Arjuna Capital's proposal focused on AI misinformation.

Even though neither proposal gained the majority approval required for passage, Welsh says she is encouraged by the results – especially the Apple vote – given that the debate is such a new one.

The issue is coming on to the radars of larger, more traditional asset management firms too. A survey of governance specialists working at such institutional investors published by EY in February found that responsible AI had surfaced as an "engagement" priority (in talks with companies) this year, with 19% of respondents citing it.

Research published last year by

especially at a time when several major elections are imminent. With this in mind, activist investment firm Arjuna Capital has called on a number of big tech firms to issue annual reports on the risks arising from facilitating misinformation/disinformation and how they would address the problem.

Microsoft, which last year invested \$10bn (£8bn) in OpenAI, also

“It's about investors getting a better understanding of how companies use AI – and firms get better at communicating it

Aiming to improve on that percentage, two shareholder proposals were submitted this year to Alphabet and Amazon respectively. One, from socially responsible investor Trillium Management, urged Alphabet to formally empower its board's audit and compliance committee to oversee the company's AI activities and fulfilment of its AI principles. The other, filed by the AFL-CIO, called on Amazon's board to create a new committee to address the perceived risks its AI-based systems posed to human rights.

Although companies do not yet have clear guidelines or disclosure requirements for AI in their financial reporting, that situation will change as the technology becomes ever more material to their businesses. So says Séverine Neervoort, global policy director at the not-for-profit International Corporate Governance Network.

"We can expect to see increased regulatory scrutiny and, most likely over time, disclosure standards and requirements," she predicts.

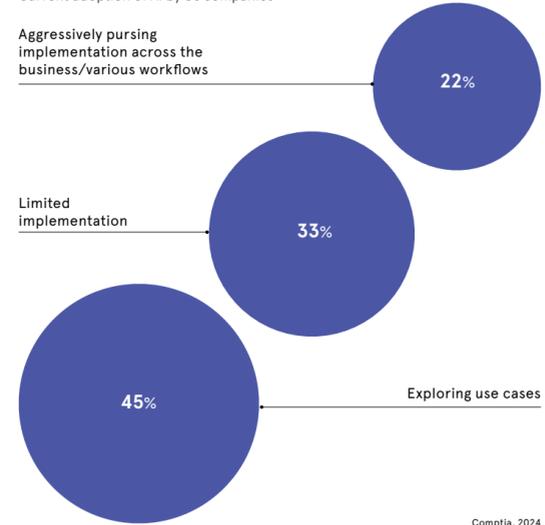
The recent disclosure rules on cyber risks issued by the US Securities and Exchange Commission (SEC) suggest a possible future for AI reporting, according to Keatinge, who foresees "a natural extension" of the regulator's approach to cybersecurity matters.

Nonetheless, she acknowledges that a new set of SEC rules for AI-related disclosures is probably still some way off, given the painstaking nature of the regulatory process. ●

ISS-Corporate, part of proxy adviser Institutional Shareholder Services, revealed that, as of September 2023, only about 15% of the S&P 500 were providing any information in proxy statements about their boards' oversight of AI.

MORE THAN HALF OF US COMPANIES CLAIM TO BE USING AI

Current adoption of AI by US companies



Why it's time for law firms to start investing in generative AI

Demand for AI has arrived in the legal industry. Law firms must embrace the change or risk falling behind

As a profession rooted in tradition, the legal industry does not necessarily have a reputation for embracing change. Yet, there is a growing appetite for trustworthy and reliable artificial intelligence (AI) amongst lawyers.

In fact, Goldman Sachs has predicted that the legal industry is likely to be the second most impacted by AI, with 46% of tasks able to be automated by this technology. With this in mind, law firms should start thinking about how AI may intersect with the legal professions' demands.

"Firms are going to have to get comfortable with a different set of skills. The skill is no longer merely drafting a brilliant contract, it is ensuring the contract that has been created correctly meets all the client needs," says Stuart Greenhill, director of segment management at LexisNexis.

According to a survey by LexisNexis in January 2024, more than one-quarter of lawyers in the UK already use generative AI at least once a month. Back in July 2023, this number sat at 11%, showing a considerable mindset shift for an industry known for being risk-averse.

“Firms that embrace generative AI can leapfrog their competitors

So, while AI is already driving change in the legal industry, there are still significant concerns. The biggest challenge to adoption is the risk of hallucination, where the AI model 'hallucinates' something that isn't real.

For example, there have been reported instances where experienced lawyers have experimented with publicly available AI models, resulting in the citation of non-existent cases.

Another notable challenge is security. Given that law firms hold privileged information about their clients, there is a risk of it leaking into the public domain if lawyers are using AI that is insecure.

To help mitigate these issues, lawyers need to ensure they are using legally grounded generative AI systems that have been trained on legal data and do not retain confidential information. LexisNexis, for example, has created Lexis+ AI, a legal-specific AI engine. The AI model is built to access all of LexisNexis' legal research and practical guidance content, so when lawyers are seeking information, they can be assured it is generated from an authoritative source.

"Anything that our AI engine generates is cited and comes with a clickable link directly to the underlying source, so lawyers can check and be confident that what it says is accurate," says Greenhill.

This can significantly shorten the time spent on legal research tasks, for instance, quickly helping a junior associate understand what case law applies to a particular situation.

"Unless the case has been summarised for you, you will have to read it and that could take 20 minutes, it could take three hours, it could take a day or longer, depending on how big

the case is," says Greenhill. "Now you can just click a button and AI summarises the case."

By freeing up time for junior associates, they can be assigned higher value work such as contract drafting or even business development and bringing in new clients, Greenhill adds. This speeds up training and gives law firms access to more advanced lawyers faster, elevating the internal talent pool.

So, the opportunities for law firms to gain a competitive advantage over their peers by moving faster on AI are significant.

"Firms that embrace generative AI can leapfrog their competitors because they are going to be more competitive, they're going to be able to invest more time in building relationships with their clients and therefore be able to spend more time on the higher value work," says Greenhill.

This is a pivotal moment for law firms and those that don't embrace technology run the risk of getting left behind. "Client expectations are climbing – with a growing demand for work to be delivered faster and with a higher level of service. Firms need to continually innovate or else be overtaken by their competition," Greenhill says.

By investing in the right tools, law firms can get ahead of this trend and ensure they are on the leading edge of the legal AI opportunity.

For more information please visit [lexisnexis.co.uk](https://www.lexisnexis.co.uk)





EVOLVE

GO FURTHER. FASTER.

Achieve up to 65% efficiency gains with our AI-powered translation solution

Built on decades of experience in AI, our revolutionary enterprise solution, Evolve, securely translates more content, faster, to give you a powerful advantage in global markets.



Take global further
rws.com/evolve

