FUTURE OF DATA & AI

FUTURE OF DATA & AI

Distributed in THE TIMES

Contributors

Sophia Akram
Journalist who specialises in foreign policy, human rights, global development and their crossover with business and lifestyle.

Alison Coleman
Writer and editor, working as a senior contributor at Forbes, with articles published in *The Guardian* and *Quarterly*, among others.

Natasha Serafimovska
EdTech SaaS expert and freelance writer, who specialises in the future of work, digital transformation and workplace learning.

Emma Woollacott
Journalist writing on business, technology and science. She is a regular contributor to the BBC News website, *Forbes* and *Private Eye*.

Jon Axworthy
Journalist, specialising in health, tech, science and the future, with work published in *T3*, *Wareable* and *The Ambient*.

Bernard Marr
Raconteur's columnist, author and one of the world's most successful social media influencers at the intersection of business and technology.

Mark Walsh
New York-based freelance writer covering business, technology and media. He has contributed to TheStreet.com and *The Guardian*, among others.

Raconteur

Campaign manager
Ethan Wrede

Reports editor
Ian Deering

Deputy reports editor
James Sutton

Editor
Sarah Vizard

Chief sub-editor
Neil Cole

Sub-editor
Christina Ryder

Commercial content editors
Laura Bithell
Joy Persaud

Associate commercial editor
Phoebe Borwell

Head of production
Justyna O'Connell

Design/production assistant
Louis Nassé

Design
Kellie Jerrard
Harry Lewis-Irlam
Colm McDermott
Sean Wyatt-Livesley

Illustration
Celina Lucey
Samuele Motta

Design director
Tim Whitlock

@raconteur raconteur-media @raconteur.stories

raconteur.net /future-data-ai-2023

TRUST

Can distrust in AI impact your business?

The hype around generative AI has further influenced public trust in the technology. Businesses can use this as a guide to how they use it and the ethics they apply

Sophia Akram

Are you scared yet, human? Artificial intelligence has proliferated with transformative effects in sectors from autonomous vehicles to personalised shopping. But the latest use of AI to generate content such as text, images or audio has caused quite a stir.

ChatGPT is a particularly superior language model, even passing the US medical speciality exam. That's not to say there haven't been some bloopers. The OpenAI release has delivered inaccurate information and even abuse. It also self-warns that it could generate bias and harmful instructions.

Even before tools such as ChatGPT, Bard and Dall-E 2 won such wide attention, there have been concerns that discrimination and bias are baked into algorithms. This phase of newer, more accessible AI could greatly impact users' trust in the technology. What, then, for businesses that have rushed to adopt these latest forms of AI – should they count on its longevity and if so, is it possible to embed AI ethics, so that mistrust doesn't hurt their reputation and bottom line?

Despite some opinions that generative AI is a fad, most people think it's here to stay. Still, as a Morning Consult survey of 10,000 US adults revealed, only 10% of the public find generative AI "very trustworthy". Drilling down further, that level of trust wavers between demographic groups, with younger cohorts, primarily male, more trustful and willing to adopt early than older generations, who are generally hesitant to pick up new technology.

It isn't only the end users who harbour doubts, though. High-profile gaffes – such as when Google's Bard circulated false facts in search results – have perhaps been a cautionary tale. Apple Inc also delayed approving updates to its email app with AI-powered language tools, over concerns that it might show inappropriate content to children.

That's not to say that generative AI isn't hugely useful. It is applied across business functions, from marketing and sales to IT and engineering. Its applications range from crafting text to cutting through dense material to aid understanding and answer complex questions.

"Companies are investing a lot in data and tech," says Karl Weaver, SVP consulting, EMEA, at MediaLink. But he cautions that there is a general acceptance of analytics.

Added to which, "Businesses see what their competitors are doing and think they should perhaps be doing the same or risk missing out. All that could cause a misstep and subsequently a trust problem."

That isn't to say that businesses should or even can avoid the wave. But CEOs and boards should step back and think, in an informed way, about why they might be using AI tools, including these most recent iterations. If it's a genuine desire to improve customer experience, it should be set up to ensure they are serving that area.

"We're developing a unified and common understanding of the big risks," says Robert Grosvenor, a managing director at Alvarez & Marsal. "But there is a long way to go to translate high-level, principle-based objectives into codified requirements, standards and controls."

The scope and scale of AI's applications span different industries and sectors and the possibilities of harm and degree of risk vary accordingly; impacts could even differ within the same organisation and be unforeseeable. As a result, individual business functions could have unique workflows for their use of AI and the data they need, instead of relying on analytics or compliance functions to dictate a 'cookie-cutter' approach to rules about using AI.

Andrew Strait is associate director of emerging technology and industry practice at the Ada Lovelace Institute, which researches the impact of data and AI on people and society. He says that distrust in some AI technologies has meant people want to see more regulation. Consumers can be confident that the food they buy in a supermarket is relatively safe, for example. But the same level of regulatory oversight, and thus consumer trust, doesn't exist for AI – the technology has developed too quickly for the regulation to keep up.

People want transparency about the data practices involved in AI and individual privacy. But Strait observes that there is a misconception that simply telling people what you are doing is enough to build trust.

"That lacks a deep understanding of the context in which someone is experiencing your product," Strait says. He would like people to participate in the governance of AI.

Data cooperatives could be the means to that end. A representative collective of people in a data set decides who accesses the data. This is in action in Spanish healthcare, where the cooperative Salus Coop gives citizens control over their data for research purposes.

Despite the general regulatory lag, the EU is tackling the problem. The Regulation Laying Down Harmonised Rules on Artificial Intelligence, the so-called AI Act, is under discussion. One of the regulation's proposals to address ethical dilemmas and safeguards would assign risk levels to AI uses, while enabling its benefits. Generative AI tools for use in sensitive areas such as recruitment would be classed as a high risk designation. This would trigger 'conformity assessments' to check that certain standards were met, hopefully reassuring end-users.

Human apprehension about AI, or a hesitancy to fully trust it, persists. But while we attribute human-like characteristics to AI, it is machinery. As actual humans, we can still check fact from fiction and set expectations around acceptable uses of AI and businesses can lead that charge.

WHAT'S STANDING IN THE WAY OF TRUST?

Barriers to developing trustworthy AI in organisations worldwide — IBM, 2022

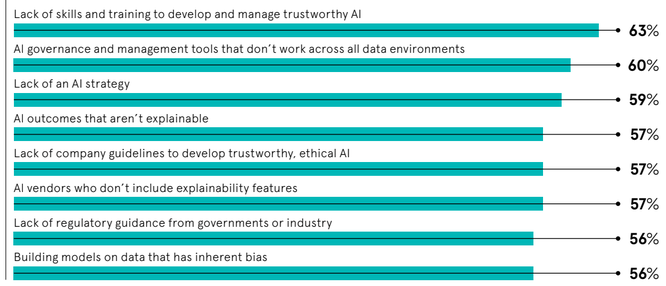

COMPETITIVE INTELLIGENCE

Gather it, sort it, use it

There's a tremendous amount of competitor data in the public domain but it's often dispersed and opaque. AI-enabled competitive intelligence can help companies gather it, sort it and draw insights from it

Laurence Dutton via Getty Images

Jon Axworthy

The business landscape is strewn with the corpses of corporations that failed to notice what competitors were up to. Think of Blockbuster's dismissive attitude to Netflix in the noughties and then even the latter's lack of foresight in seeing the streaming wars on the horizon.

This is why organisations need robust competitive intelligence (CI) strategies: to ensure they are not blind-sided by competitors or disruptive business models.

But gathering data as part of CI can be daunting because there is so much information available. Industry experts' blogs, financial reports, news media items, public data sources and more are all there for harvesting.

The good news is that AI and machine learning can streamline and accelerate this task. SaaS platforms use AI to track and collect historical and real-time data insights, which allows businesses to use information from competitors' digital footprints.

"There is exponentially more information online in the digital footprint of every business relative to a decade ago," says Jonah Lopin, founder and CEO of Crayon, a competitive intelligence software platform.

"If you have the tools to aggregate and analyse this information you can understand a competitor's product life cycle, pricing and more. Those insights can help a firm anticipate that competitor's moves."

Automating the platform's processes ensures the intelligence is timely enough for a business to take prompt action. Drawing on the guidance and recommendations of AI at executive-level meetings could also enhance agile thinking when critical actions are needed.

The evidence suggests that more businesses are using the technology. More than 10,000 roles for competitive intelligence are advertised on LinkedIn, for instance. The global market size of CI tools is projected to hit $82bn (£68bn) by 2027, according to Fortune Business Insights.

A significant proportion of this spend is likely to come from corporations whose deeper pockets allow them to deploy state-of-the-art large language models (LLMs) like GPT-3, which can provide more sophisticated data insights. GPT-3 is one of the most straightforward models for building CI. But it needs infrastructure, skills and software to work effectively, which don't come cheap.

Regardless of the computational power of the AI model, it needs targeted and verified data to function and benefit a company's CI strategy.

> **Without human input to give context to the data AI can generate, brands will be unable to utilise the creative intelligence they have at their fingertips**

That data needs to be drawn from multiple sources in order to prevent its 'intellect' from becoming circular and repetitive.

AI models are only as intelligent as the training data they are given. The more data they're trained on, the more accurate, versatile and useful they are to a business. This means there must be skilled oversight of the quality of the data.

"Training with low-quality data, like social-media posts or blogs, can introduce more noise into the learning process, which can confuse the model and decrease the quality of the outputs," explains Kalyan Veeramachaneni, a principal research scientist at the MIT Institute for Data, Systems and Society.

Data bias can also be a consequence of working with low-quality data. For example, if data is limited to text from a particular group it can skew the AI's function, lower overall performance and reduce the reliability of the intelligence.

But by monitoring data quality from sources with editorial filters or filtered web content and not relying too heavily on user-generated content, companies can increase the quality of the inputs, which will also raise the level of intelligence that they are receiving, in terms of reliability and usefulness.

Whether you have decided to base your CI strategy on insights which are provided by an advanced LLM or a more 'cookie-cutter' SaaS platform, you will still need to involve people. Then, you have the combination of matching competitive intelligence with creative intelligence from in-house teams to maximise any advantages and make the business stand out from the crowd.

"AI can be an asset for brands to stay ahead of the game in competitive markets. It can be used to help to increase efficiency and insight within marketing teams. And it can reduce time-consuming, laborious processes," says Anthony Lamy, vice-president, client partnerships at VidMob. "But without human input to give context to the data that AI generates, brands won't be able to use the creative intelligence they have at their fingertips."

Businesses can therefore use the technology to discover a competitive edge not by relinquishing complete control to AI, but by using it in tandem with human creative teams to increase their creative intelligence.

It seems , then, that while there was either huge distrust or blind faith in what the technology could achieve for an organisation's CI strategy, there is now a realisation that if you aren't seriously considering investment in or scaling up AI-driven CI, it's very likely that your competition will be.

THE COMPETITIVE EDGE

Motivations of product managers and marketers for conducting competitive intelligence

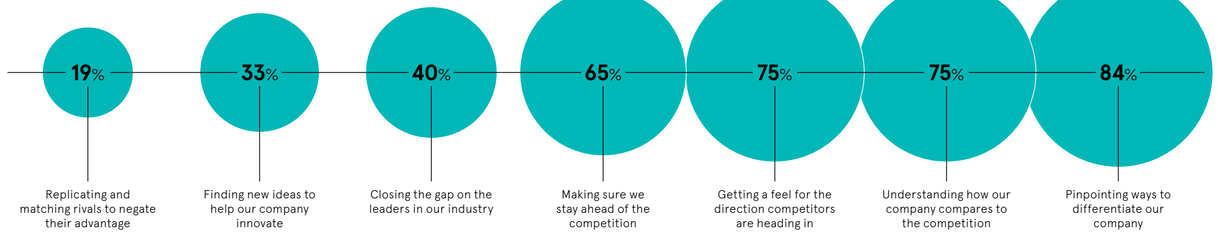

Commercial feature

Why web scraping is the future of data-led innovation

Access to actionable data is becoming vital for businesses. Whether you're fuelling AI solutions or looking for insights into consumer behaviour, web scraping can help ensure reliable access to the data decision-makers need

For many years, web scraping has been a central concept in the tech world and has impacted our lives in ways many of us are unaware of. "Even regular internet users constantly run into businesses that could only be possible through web scraping," says Juras Juršėnas, the chief operating officer of Oxylabs, a web-intelligence-acquisition solution and premium-proxy provider. Data scraped from the internet using automation provides the backbone of everything from search engines to travel-fare aggregators, price-comparison websites and many other services.

But what has long been an unsung, central part of our digital lives is now entering the limelight. The rise of AI has captured the attention of even the most technophobic executives, who are now seeking ways to integrate it into their business models. Yet behind every AI model is a powerful corpus of training data – most of which is scraped from the web. "We are moving in a new direction," says Juršėnas. "One where AI – namely machine-learning models – are becoming ubiquitous. ChatGPT, Bing Chat, Google Bard are all based on the same principles."

Oxylabs recently surveyed more than 1,000 senior ecommerce industry data decision-makers in the UK and US. Nine in 10 of them said they thought web scraping would become a more important part of their focus in the coming years. That focus is even more important now that we've entered the age of AI, where large language models (LLMs), chatbots and image generating models all have a voracious appetite for training data, which is needed to offer iterative improvements in how businesses are run.

However, it's difficult to get to a point where AI can reliably help you. "Machine-learning models are hungry for data, as they need millions, sometimes even billions, of data points to provide high levels of accuracy and predictive power," says Juršėnas. "Web scraping can provide companies developing machine-learning models and AI with the data they need."

How to solve the data quality problem

Yet it's not simply a case of setting a web scraper going and overhauling the way your business works. There's a maxim in the world of AI: garbage in, garbage out. A model is only ever as good as the data it's trained on, which means the level and scale of high-quality data that is needed to realise the full potential of the future is greater than ever before. "Enormous volumes of data will be fed into models that will often function as black boxes," says Juršėnas. "If a large part of the data is faulty, the results will be unpredictable and could cause damage."

The risks of low-quality data infiltrating an AI system and then poisoning the well has real-world ramifications. Google saw $100bn wiped off its value within a day of unveiling its own AI search tool, Bard, which was shown to answer simple questions incorrectly.

Having high-quality web scraping is an essential part of solving the problem, without which it would be exceedingly easy to run into issues. Businesses must dedicate proper attention and resources to this area in order to remain competitive. It will soon become normal for businesses to use AI in their day-to-day operations, meaning those that don't adopt the technology could quickly fall behind.

90%

of senior ecommerce industry data decision-makers in the UK and US think web scraping will become a more important part of their focus in the coming years

Oxylabs, 2022

WEB SCRAPING HELPS UNLOCK THE POWER OF EXTERNAL DATA

Share of technology business leaders expecting to shift towards web scraping in the next 12 months

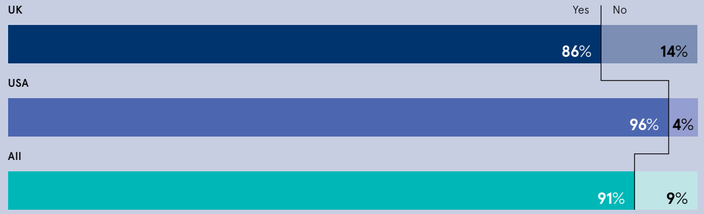

Most common uses of external data in organisations worldwide

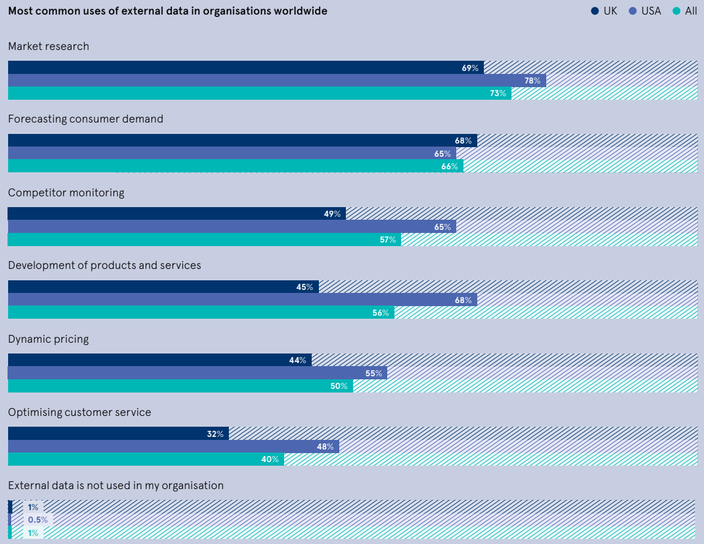

Oxylabs, 2022

> **Enormous volumes of data will be fed into models that will often function as black boxes. If a large part of the data is faulty, the results will be unpredictable and could cause damage**

Walking the tightrope

Whether your business benefits from the AI revolution is likely to come down to the quality of data you're working with. From good data you can get good AI-generated insights, which can help you innovate and overhaul your business for the future. "Web scraping is the way to enable the creation of advanced AI through the provision of high-quality data," says Juršėnas. "On the other hand, the data has to be carefully managed and acquired from reputable and ethical solution providers. A single misstep can be exceptionally costly."

Those missteps can be simple errors like Google's, which if implemented into a business could lead you into a new market or a new contract that ends up harming, rather than helping your business. Another common misstep is training a model on biased data, which can come back to harm your reputation. This was the case for Microsoft's Tay chatbot in 2016, which shortly after release began expressing racist sentiments.

Finding a good partner who can help you remain on the cusp of innovation, while doing so responsibly and ethically, can be a challenge. "At Oxylabs, we provide web-scraping solutions for all businesses that need public web data," says Juršėnas. "Our services may be used to gather data for optimisation, to build machine-learning models, or even combine both purposes."

Trusted, truthful partnership

Oxylabs is trusted by companies big and small to conduct web scraping that can help build AI models that work. "We solve the greatest pain point for any data-driven venture: data acquisition," says Juršėnas. "Through our numerous scraper APIs, companies can extract as much publicly available data as they need and have it delivered in real time."

Among the list of clients trusted by Oxylabs to build the foundations of their tech stack through web scraping is the Lithuanian government. Oxylabs created a solution that scans the Lithuanian IP address space on the internet and detects illegal sexual and child-abuse imagery. Evidence is then forwarded to specialists for review.

It is evidence of the level of trust placed in the company – but it is far from the only use of the AI it helps enable. "We firmly believe in an AI-led future, but we understand that it will be data-hungry," says Juršėnas. "Our goal is to enable businesses of all sizes to get the web intelligence they need, build machine-learning models and optimise their business processes. If data is the foundation of your business, Oxylabs will help you collect it reliably, safely and ethically."

For more information visit oxylabs.io

oxylabs

ANALYSIS

ChatGPT could make you better at your job

Raconteur's columnist **Bernard Marr** is a world-renowned futurist, influencer and thought leader on business and technology. He sets out the uses of ChatGPT and natural language processing – and why they won't make humans surplus to requirements

Bernard Marr

For the past few months, it seems everybody has been going AI crazy. Futurists have long been predicting that it will revolutionise our lives and it's plain to see it already is. Mostly, though, this has been happening 'under the hood' – quietly powering tools we use every day such as Google, Netflix and Uber in a way that is (by design) invisible to the user.

ChatGPT and related tools and applications such as Bing and the soon-to-be-released Bard, on the other hand, are 'in-your-face' AI. Millions, who have now had the chance to see them in action, have been left in no doubt that this is something truly new, genuinely revolutionary and a little (perhaps a lot) scary.

News is moving quickly. At the time of writing, Microsoft is thought to have scaled back and limited the ChatGPT functionalities it recently integrated into its Bing search engine. This comes following reports of users who were cheeked and chastised by the feisty chatbot, and others who have worked out clever ways of instructing it to adopt new, not entirely helpful personalities. Some researchers have even claimed that the machine learning-powered algorithms have been telling them that they are sentient and want to be alive.

It's fair to say that we've all had a lot of fun and it's thrown up some interesting ethical and philosophical debates. But is it set to be as revolutionary as it seems when it comes to changing the way we work? Or is it a flash-in-the-pan that will be forgotten about when we eventually realise it still isn't quite good enough to let loose on really important tasks?

Although it is impressive technology, anyone who has used generative AI for some time will have bumped up against some of its limitations. The most glaring is probably the fact that it isn't capable of original thought. ChatGPT (and other applications that will follow shortly) draws all of the knowledge that goes into its output from its training data.

In simplified terms, it constructs responses to questions and queries by analysing millions of words of text that have previously been written and applying probability to determine the best thing to say next. It is a language model that understands the structure and context of sentences and therefore is capable of creating its own. What it can't do, though, is come up with an original idea or an answer to a question that has never been correctly answered before.

> **If your audience comes to you for specialist knowledge, expert opinion – or just because they like your personality – AI-generated content is likely to leave them cold**

For most use cases, this is probably fine. No one expects to use it to ask, for example, the secret of generating perpetual energy. (Although it can certainly summarise a large amount of the corpus of existing human knowledge on the issue). Where it is likely to be useful is in automating routine parts of our work.

This limitation is the main reason that natural language technology is not (yet) simply going to replace humans. There will, for the foreseeable future, be a need for humans to oversee and steer AI, providing the 'big picture' direction and the original thought needed for truly useful or valuable endeavour.

This is why, when speculating about how this technology is likely to impact our working lives, it makes sense to look at the particular abilities and skills that it can augment, rather than at specific jobs or professions that may or (more likely) may not be in danger of being automated out of existence.

Any such list has to start with writing. On the face of it, this is ChatGPT's main function – to produce text. If you're going to use it to write, though, it's important to remember that it won't generate anything new or original. Where it can be helpful is with suggesting ideas about how to structure an essay, article, blog or social media post, or generating a list of the most important points that need to be covered. Just be wary that if your audience comes to you for specialist knowledge, expert opinion or simply because they like your personality, then AI-generated content is likely to leave them cold.

Another useful capability is generating code. Not limited to human languages, ChatGPT can write code in several popular languages, including C++, JavaScript and Python. It can also error-check existing code. By taking advantage of this, just about anyone can become capable of quickly creating simple computer programs to automate routine elements of their work. Developing this skill is likely to be increasingly important in many professions.

It can also be a great tool for research. It can be more useful than a search engine – but its output can often include errors or omissions. Therefore, the ability to review and verify the information it churns out is still essential.

Others may find that its most effective use cases involve data analysis. It can interpret information, dissect text and numeric data, and even create charts. Combining this with its ability to generate code, it can be used for data analytics.

Finally, it has tremendous potential to assist with planning and project management. It can provide a step-by-step guide, including the tools and skills needed, the processes to put in place, and how to analyse and assess your results.

If you work in an area heavily dependent on one (or more) of these skills, you wouldn't be alone in worrying if you're likely to be soon replaced by a machine. But no one should be immediately looking to move into something that will never be automated – if such a thing exists. Instead, it would be more rewarding to look at how you can use AI, and specifically natural language technology, to augment your skills in the relevant areas. Writers should use AI to become more thorough and informed in their writing. Programmers can become more productive and efficient at creating code. Data analysts can use AI to find ways to look at their information and to process bigger datasets more quickly and efficiently.

> **Where it is likely to be useful is in automating some routine elements of our work**

But thinking beyond that, writers can become data analysts, to create copy that's more informed by facts and statistics. Data analysts can become writers, presenting their findings more engagingly and completely. Programmers can become project managers, bringing together different skills to create more useful applications – and the list goes on.

ChatGPT, Bard or some future iteration of language-based AI will change many things about how we work. It might not happen right away, but anyone who wants to be part of this future has all the tools they need today to start taking steps in the right direction. In the same way as the internet, or the mechanisation brought about by the industrial revolution – it isn't going to go away.

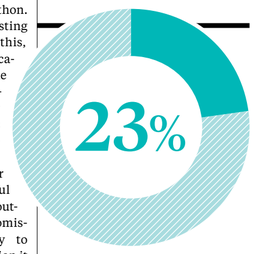

of businesses worldwide have adopted AI for natural language speech understanding as of December 2022

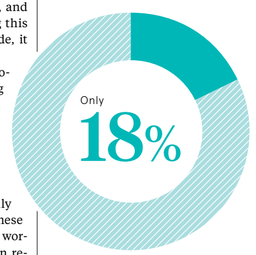

have adopted AI for natural language generation

McKinsey, 2022

Open warfare: will data sharing win the fight against cybercriminals?

Security teams are battening down the hatches against a barrage of coordinated cyberattacks. But without transparency and collaboration, are corporations fighting an uphill battle?

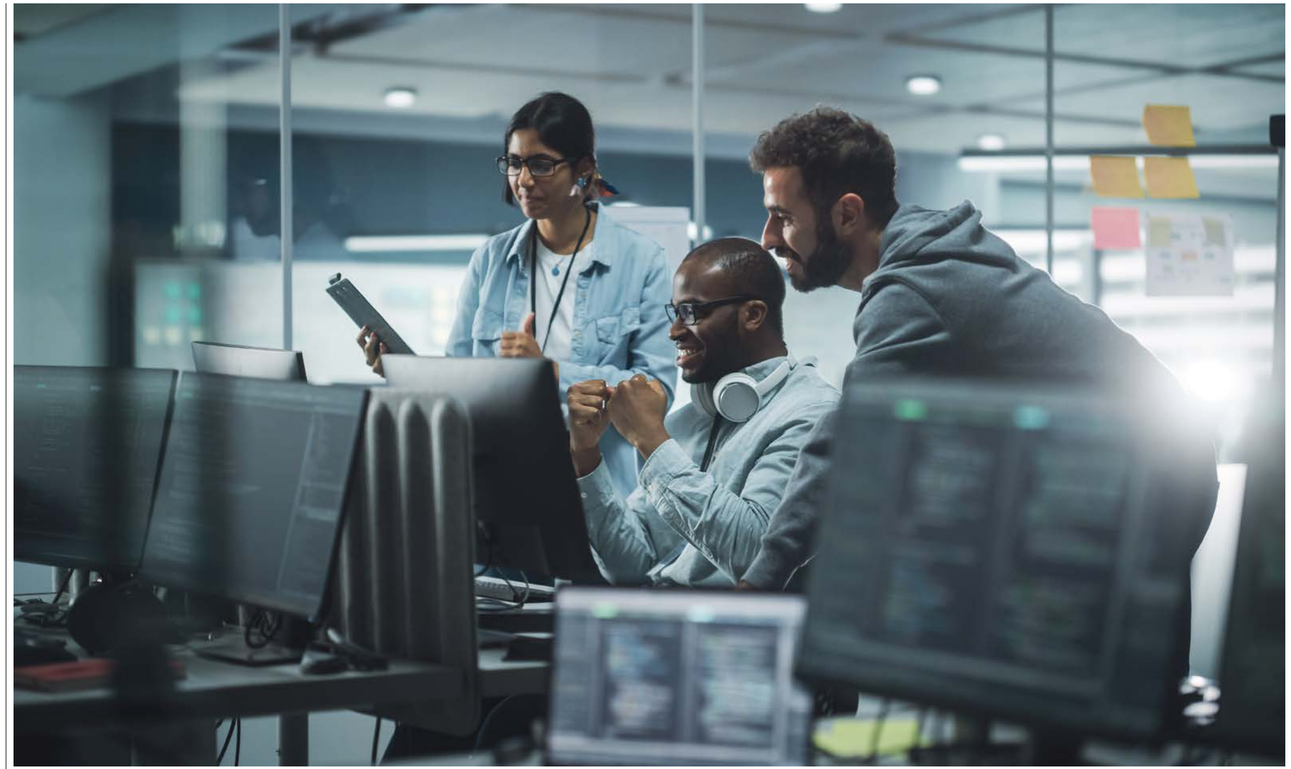

From battlegrounds to sporting fields, it's often noted that the best defence is a good offence. The strategic weight of this well-worn adage holds firm for businesses looking to reinforce their cybersecurity.

In the attack landscape, cybercriminals often join forces to disseminate sensitive information, share sophisticated tactics and expose corporate vulnerabilities. While intelligence can be harvested by attackers globally and weaponised against any sector at any scale, CTOs and CISOs are left putting out fires individually rather than working together to proactively prevent them.

Tony Meehan, vice president of engineering for security solutions at Elastic, believes that democratising data in the same way that cybercriminals do will keep businesses a step ahead.

"Don't get me wrong, confidentiality is still really important," says Meehan. "I'm not asking for every security team on planet Earth to go and post all their detections on GitHub tomorrow. But we do need to find ways to collaborate more openly and share knowledge, techniques, and best practices."

As infiltrations become more prolific, coordinated and commoditised, organisations can't afford to let cybersecurity skills gaps or outdated defence strategies hamper their responses. "The attack surface has become way bigger. I don't know if we can make a dent in this problem with the same approach of the last 20 years," Meehan continues.

Meehan, who worked at the United States National Security Agency (NSA) for a decade on programs to collect foreign intelligence, outlines three main problems facing today's defensive teams when defending their organisations.

The first is the speed of digital transformation post-pandemic, which opened up holes due to businesses' accelerated transition to the cloud. The second is the growth of nation-state attacks, something that wasn't a concern 10 or 15 years ago. And the third is talent scarcity in the security space which makes it harder for individual teams to keep up with new and emerging threats.

"The goal of a good defence is to make the adversary work harder. I think the journey to achieving that really needs to be built around an open community," says Meehan.

Elastic's own search-powered solutions are built on this premise of openness, regardless of whether data lives on a single or multiple cloud setup or on-premise. The company has helped the likes of Adobe, BMW and Zurich Insurance find what they need faster while keeping mission-critical applications running smoothly and protecting against cyber threats.

Meehan appreciates C-suites may feel fear or scepticism over sharing sometimes sensitive information. But to fight off sophisticated attacks designed by malicious collectives, organisations must achieve the same level of transparency as those trying to get in through the backdoor.

Removing organisational data silos is one answer to deliver greater visibility of what information is where when it is attacked. However, corporations must actively pursue new routes for collaboration if they want to transform the preeminent cybersecurity culture. According to Meehan, this closed-off culture that prioritises privacy at all costs means companies rarely understand how their purchased vendor security products work; they just accept that they will.

> As an industry, we shouldn't be embarrassed about the flaws we find, but how long it takes us to fix them and the lack of investment in finding more. We should want to find flaws

"When confronting the trends of the last couple of years, it's paramount to really understand what your products are doing for you," he says. Additionally, data sharing to a far greater extent – detailing threats, foiled attacks, and successful infiltrations – will empower teams.

Meehan likens this to the successful sharing of YARA signatures, commonly used to identify and detect malware. "Openness enables knowledge sharing, which will help elevate your team. You can even share specific detection methods in smaller groups without ever exposing them to the world," he says.

Increased openness means the entire security community learns and grows. Meanwhile, the attack surface shrinks as it becomes harder for malicious actors to find bypasses across multiple companies. Meehan explains: "All of our detections are in the open. That's an excellent starting point. Even if you're not using our product, you can still go and use our detections."

But security has traditionally adopted a very closed culture, meaning potential vulnerabilities can go unexamined. At the same time, attackers could spend every day for months searching for gaps.

Meehan accepts that few security vendors want to poke around in their own products because they don't want to be confronted by the holes they might discover. But this reluctance is evidence enough that the system is broken.

"It's very hard to get people to spend that much time looking for these things, so we have to have a conversation around doing more things in the open," he says. "If everyone is being a little more transparent about their security controls, threat logic and detection rules, that becomes a force multiplier for all teams' best practices. Not everyone has to start from scratch."

The combined efforts of partners and volunteers as part of the Shields Up initiative following Russia's invasion of Ukraine is a prime example of collective defence in action. Openly sharing vital information has helped Ukraine become a cybersecurity heavyweight. "Supporting one another is a natural reaction. In Ukraine, it was the obvious thing for us to do," Meehan explains.

While companies do share data, the practice is primarily relationship-driven and isn't as formalised or progressed as it should be. If the majority of organisations are solving the same problems at the same time, the system needs to be revised.

"As an industry, we shouldn't be embarrassed about the flaws we find, but how long it takes us to fix them and the lack of investment in finding more. We should want to find flaws," says Meehan. The more comfortable companies are with internal scrutiny, the harder it becomes for outsiders to game the system.

Appreciating the need for greater sharing will make a tangible difference to the people and products that systems are designed to protect. A new cybersecurity culture that promotes open information and close ranks will set organisations on a path to victory.

For more information visit elastic.co/explore/security-without-limits

elastic

46%

of all organisations are investing in upskilling cybersecurity and IT staff

41%

of executives think that cyber risk initiatives have not kept pace with digital transformation

Elastic, 2022

Q&A

The data silo dilemma

Visibility is the first step towards security, and that means embracing openness, explains **Mike Nichols**, vice president of security product management at Elastic

Q How are data silos creating security challenges for organisations?

A Since Covid-19, businesses have exploded into the cloud much faster than expected. In the rush to support remote working, companies began pulling data from more applications and sources than ever, which opened them up to exploitation. It's also much harder now to break down information and identify which portions are most critical to operations. Companies are finding they don't have the right expertise internally to understand or monitor it all at scale. And the sheer amount of information they need to sift through can be overwhelming.

When your systems are compromised, it's not just about spotting the intrusion; that's only half the problem. Preventing someone from being inside long enough to cause damage is crucial.

Minimising the dwell time of your adversary is effectively a data access challenge. You may have seen an initial alert but can't gain access to the areas you need because the designated expert is away, the data isn't available to the analyst, or the data simply doesn't exist. Analysts can't connect the dots if they're segregated from the data they need when they need it.

Q What's the next step for businesses to respond effectively to cyber threats?

A People silos are as tricky as data silos. You might have an endpoint expert, a firewall expert, and an email expert - but they all work in isolation. Unified visibility is the first step towards security, and that means embracing openness.

With an open schema or framework, the power goes back into the hands of the customer as opposed to the vendor. You control your data and your rules, and you can freely switch out technology vendors as new products emerge.

Your analysts need to be in a position to act quickly when a security incident happens. But the more silos there are, the longer that process takes, and the business risks greater exposure. Removing those restraints for your analysts can make a big difference to the damage toll at the end of the day.

> When your systems are compromised, it's not just about spotting the intrusion; that's only half the problem

Q How can security teams improve their decision-making?

A I've had many conversations where the breached business or organisation doesn't even know what happened or what was stolen. So how can your analysts make decisions about business risks if they don't understand what the data looks like and where it was when it was attacked?

Building that understanding and improving visibility are fundamental areas to invest in before you even think about the technology. If you don't understand where your data and assets are within an environment, that's a big crack in your security foundations.

You need to be able to perform root cause analysis in real time. When an adversary attacks, it's not over in seconds. There is a window where they make enough noise in the environment for security teams to detect and intercept the ultimate breach. They're always going to get in, but as long as you recognise that, you can really prepare. If silos are broken down, and security analysts know what data is where they can react fast enough to stop data from being destroyed or stolen.

Q What advice do you have for CISOs to get ahead of data challenges?

A First, ensure your security operations are not viewed by the rest of the organisation as a silo or as a team that sits behind closed doors, only emerging to tell someone they made an error. Instead, talk with your business leaders consistently and regularly to get to know their processes and requirements directly. When your team interacts with the rest of the business, you gain key insights that will accelerate response actions and improve explanations for alerts and requests coming from those teams. Second, own your own data. And make sure your security vendor does not lock you into their ecosystem. By insisting on open standards for data storage, data analysis, detection engineering and more, your teams can be agile in adopting new technologies and vendors to suit your security needs as they evolve.

AI ACROSS THE BUSINESS

Nearly all CEOs recognise that AI will become a significant factor in the success of their firms over the medium term. As the market for AI-powered business tools develops, how will these be applied in various functions? And how have the early adopters benefited from using them so far?

94% of business leaders believe that AI will be critical to their organisation's success over the next five years
Deloitte, 2022

UPTAKE OF AI BY VARIOUS FUNCTIONS ACROSS INDUSTRIES WORLDWIDE IN 2022, BY PHASE OF ADOPTION

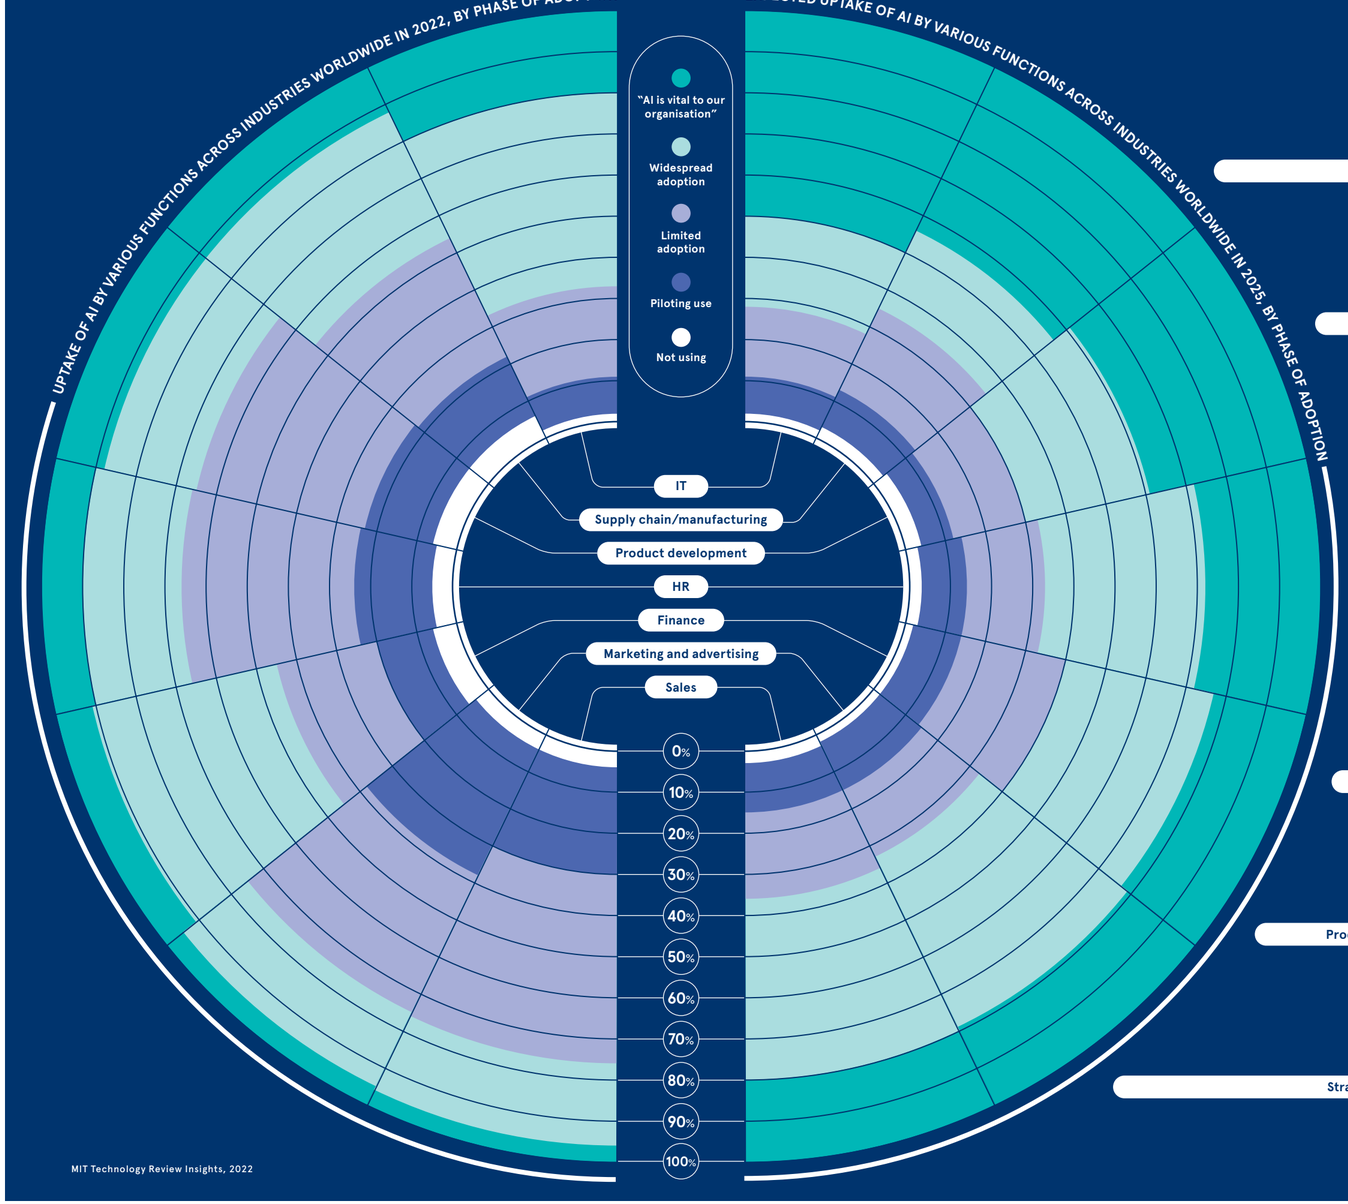

MIT Technology Review Insights, 2022

EXPECTED UPTAKE OF AI BY VARIOUS FUNCTIONS ACROSS INDUSTRIES WORLDWIDE IN 2025, BY PHASE OF ADOPTION

PERCENTAGE OF RESPONDENTS IN SELECTED FUNCTIONS REPORTING COST DECREASES AND REVENUE INCREASES RESULTING FROM THEIR USE OF AI

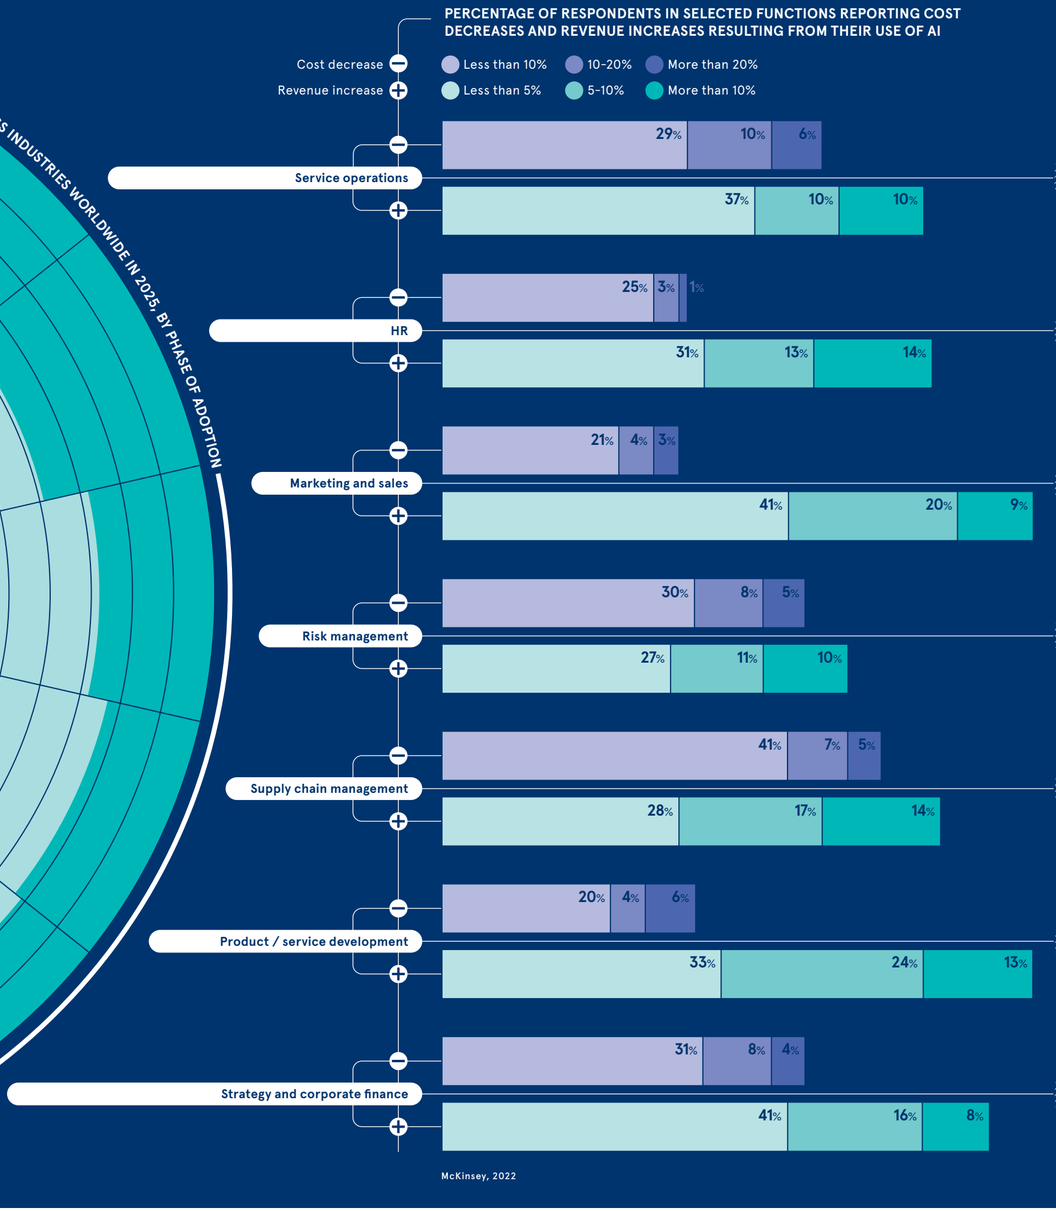

McKinsey, 2022

EACH FUNCTION'S MOST POPULAR APPLICATIONS FOR AI, BY THE PERCENTAGE OF RESPONDENTS CITING ITS USE

INTERVIEW

How far can we really go with AI?

In a forthcoming book, Princeton computer scientist **Arvind Narayanan** aims to offer a clear-eyed corrective to the hype around AI. But he also sees promise in generative AI

Mark Walsh

When Arvind Narayanan gave a presentation at the Massachusetts Institute of Technology called "How to recognize AI snake oil" in 2019, he was surprised to find his academic talk going viral on Twitter, with the slide deck eventually downloaded tens of thousands of times and numerous requests filling his inbox.

The overwhelming response has since led Narayanan, assistant professor of computer science at Princeton University, to expand his talk into a book that he is co-writing with graduate student Sayash Kapoor. Following the sensation caused by ChatGPT and generative AI, the subject of the book is clearly more timely than ever.

But what is 'AI snake oil' and just how would you distinguish it from the real thing?

Narayanan explains that AI is an umbrella term for a set of loosely related technologies without a precise definition. To help demystify the term, he has devised a scheme that classifies AI from genuine to dubious across three categories: AI relating to perception, AI automating human judgement and predictive AI.

The first category includes technologies such as the song identification app Shazam, facial recognition, and speech-to-text. The second refers to AI used for making content recommendations, automating content moderation in social media, or detecting spam or copyright violations online. The third refers to predictive AI systems in tasks from hiring to setting bail to gauging business risk.

"The third category is really where most of the snake oil is – and that's about using AI to predict what a person might do in the future," says Narayanan. "And then use that prediction to make decisions about them that might, in fact, give or deny them important life opportunities."

Unlike using AI for something such as speech transcription or image recognition, Narayanan explains that there is no ground truth data, or 'gold standard', to compare and evaluate results with predictive AI because the outcomes haven't happened yet. "The future is fundamentally unpredictable," he says.

Whether screening job candidates, predicting recidivism or the risk of a motor vehicle accident, Narayanan's research has found purported AI tools fare little better than flipping a coin. Certainly, they are not as effective as long-established statistical analysis methods such as regression analysis.

Where, then, does leave ChatGPT? Narayanan views generative AI, including ChatGPT, as an outgrowth of perception-related AI, going beyond just perceiving and classifying content to being able to generate images or text on request. Through such progress, he believes generative AI holds more promise than as a substitute for human judgement or discerning the future.

"The potential is clearly there but a lot of work still lies ahead to figure out which applications are even the right ones," he says. In that vein, Narayanan points to AI tools he uses himself, such as GitHub Copilot, which can turn natural language prompts into code and translate code between programming languages.

> **The potential is clear but a lot of work lies ahead to figure out which applications are even the right ones**

329

the number of academic papers on machine learning across 17 fields that Narayanan and his team have found to suffer from reproducibility failures or pitfalls in ML-based science

Princeton University, 2022

At the same time, he highlights some of the flaws that have recently surfaced, most notably Microsoft's new AI-powered Bing becoming erratic and telling lies in lengthy exchanges with journalists and early testers. That suggests to him that generative AI won't necessarily upend search overnight.

Narayanan, who has a lively Twitter account (@random_walker), has also referred to ChatGPT itself as a "bullshit generator". That isn't a scientific term. "I just wanted to remind people that chatbots aren't trained to be accurate," he explains. "They're trained to sound convincing, but fundamentally chatbots aren't built with an ability to evaluate the truth or falsehood of statements."

As such, he suggests ChatGPT and its rivals shouldn't be viewed as trusted sources in areas where accuracy is vital, like providing health information. "I don't think that problem is fundamentally insoluble. A lot of researchers are working on it, but it's just not there yet," he says.

Within the business realm, Narayanan suggests that companies should move carefully to incorporate generative AI into their operations. That means starting with the simplest tasks to be automated for productivity gains, "then once you have experience where you start to understand the limitations, gradually build up from there to try the more complex tasks".

That approach might involve customising a general model as opposed to using a smaller, specialised one. "The reason people are currently excited is because they feel that foundation models are perhaps a more general and quicker way to get to business-specific objectives than to train a model on a particular data set," says Narayanan.

The recent release of the ChatGPT API by OpenAI is likely to spur the rush of companies and startups harnessing the technology to add chatbots or other AI-powered features to applications, so as not to get caught behind the curve.

Narayanan praised Google for, in contrast, taking a cautious approach, significantly delaying the public release of its AI chatbot amid ethical considerations and internal debate. But in the wake of ChatGPT and Microsoft's Bing relaunch, the search giant is playing catch-up. It announced its Bard chatbot in February and is planning to include AI in all its major products within months, according to a Bloomberg report.

Welcome to the AI arms race. "A lot of the last five years of progress in responsible AI is in fact eroding at this moment," says Narayanan, who co-authored a textbook on machine learning and fairness. He also led the Princeton Web Transparency and Accountability Project, uncovering how companies collect and use people's personal information.

> **Chatbots aren't built with the ability to evaluate the truth or falsehood of statements**

To limit the dangers of an AI free-for-all, Narayanan says government regulation will have to play a role in ensuring new AI systems perform as advertised and don't abet discrimination, disinformation or other harms. Indeed, governments are scrambling to figure out how to address the proliferation of AI across all aspects of society.

But Narayanan emphasises that existing laws, such as those dealing with discrimination or fraud, can already be applied to problems emerging from the rise of AI. In that vein, the US Federal Trade Commission recently issued a warning to businesses about exaggerating what AI products can do or whether they use AI at all.

And since business as an institution enjoys a measure of public trust, Narayanan says it's especially important companies don't overpromise what AI can deliver. "Unfortunately, when they overhype some of these technologies and confuse public discourse, they're doing everyone a big disservice," he says.

INCLUSIVITY

Discrimination game: time to scrap the skew

It's widely acknowledged that there are very human biases baked into many machine learning models. But what are those with the power to solve the problem doing about it?

Alison Coleman

Mads Perch via Getty Images

Any AI system, however sophisticated, is only as good as the data on which it is trained. Any bias in its outputs will result from distortions in the material that humans have chosen to feed into its algorithms. While such biases are unintentional, the AI field has a predominantly white male workforce creating products that will inevitably reflect that demographic's particular prejudices.

Facial recognition systems, for instance, could be inadvertently trained to recognise white people more easily than Black people because physical data on the former tends to be used more often in the training process. The results can put groups that have traditionally faced marginalisation at a further disadvantage, heightening barriers to diversity, equity and inclusivity in areas ranging from healthcare provision to recruitment.

The good news is that the problem has been widely acknowledged in business, academia and government. Efforts are being made to make AI more open, accessible and balanced. There is also a new ethical focus in the tech industry, with giants including Microsoft and Google establishing principles for system development and deployment that often feature commitments to improving inclusivity.

Despite this, some experts argue that little coordinated progress has been achieved on establishing ethical norms for AI, particularly in relation to diversity.

Justin Geldof, technology director at the Newton Europe consultancy, is one of them. He argues that "governments have abdicated responsibility for this to the tech industry and its consumers. There is no formal watchdog and no general agreement on norms in AI ethics. With no one stepping in to lead on AI and its ethics, this has in essence become an arms race."

High-profile breakthroughs in the field of generative AI recently have also done little to address concerns about discrimination. In fact, there's a risk that the potential harms of generative systems have been forgotten in all the media hype surrounding the power of OpenAI's ChatGPT chatbot and its ilk, according to Will Williams, vice-president of machine learning at Speechmatics.

Meta took Galactica, its large language model, offline in November only three days after launch, for instance, amid fears about its inaccuracy and potentially dangerous impact on society.

Williams says: "The truth is that the inherent bias within models such as ChatGPT, Anthropic's Claude and Google's Bard means that they cannot be deployed in any business where accuracy and trust matter. In reality, the commercial applications for these new technologies are few and far between."

> **With no one stepping in to lead on AI and its ethics, this has in essence become an arms race**

In simple terms, generative AI models average the opinion of the whole internet and then fine-tune that via a process known as reinforcement learning from human preferences. They then present that view as the truth in an overly confident way.

"This might feel like 'safe AI' if you are one of the humans fine-tuning and editing that model," Williams adds. "But, if your voice isn't represented in the editing room, you'll start noticing how your opinion varies dramatically from those of ChatGPT, Claude and Bard. Your truth might be some distance from the truth they present."

The race to produce a winner in the generative stakes has given new urgency to addressing bias and highlighting the importance of responsible AI.

Emer Dolan is president of enterprise internationalisation at RWS Group, a provider of technology-enabled language services. She says that, while the detection and removal of bias is "not a perfect science, many companies are tackling this challenge using an iterative process of sourcing targeted data to address specific biases over time. As an industry, it's our duty to educate people about how their data is being used to train generative AI. The responsibility lies not only with the firms that build the models but also with those that supply the data on which they're trained."

Wider technical, analytical and academic communities are applying several methods to reduce or even remove bias. These include supervised learning; synthetic data sets that contain computer-generated material instead of real-world data; and peer reviews, in which teams with no prior knowledge of the data analyse the code and the results.

"Machine learning models that banks are using in areas such as fraud detection have too often been based on relatively small samples," says Ian Liddicoat, CTO and head of data science at Adludio, an AI-powered advertising platform. "To achieve more accurate results, synthetic data can be used to mimic and augment the original data sets. This data also seeks to even out the distributions for factors such as gender to ensure that they reflect society more accurately."

Data engineers are also using more advanced random-sampling methods that create data sets for modelling, where each record has an equal chance of selection. Supervised learning methods, meanwhile, can be applied to neural networks, so that real-world distributions for factors such as race or gender are enforced on the model.

"Another effective solution", adds Liddicoat, "is to use sub-teams to conduct detailed reviews of the input data, the machine learning method, the training results and the operational outcomes."

Beyond the technical challenges of removing bias from the systems and the data they're based on lies the bigger issue of improving diversity in the AI sector. The only truly effective way to achieve this is for its employers to recruit from a wider talent pool. The ongoing lack of diversity in academia, particularly in STEM subjects, is troublesome for the future of AI. Until that fundamental problem is solved, AI may always pose a threat to diversity, equity and inclusivity.

NEARLY THREE-QUARTERS OF FIRMS ARE FAILING TO REDUCE BIASES IN THEIR AI SOLUTIONS

Share of organisations that are falling foul of AI safeguarding in the following ways

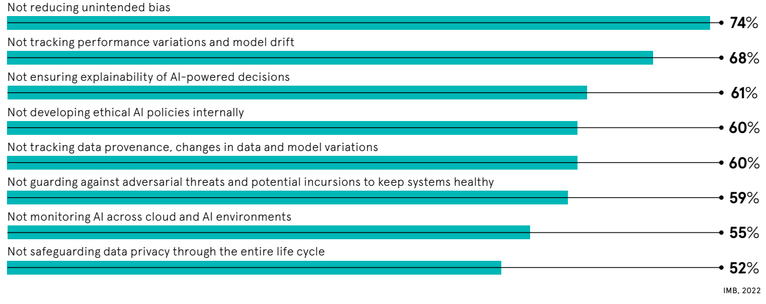

IMB, 2022

Commercial feature

How AI-generated content raises revenues and connects audiences

Artificial intelligence can complete many tasks, from cloning voices to generating full-scale ad campaigns that boost revenue and engagement. But for creatives, is it an ally or adversary?

When football commentator Alan Smith's words echoed out, "Messi shoots – it's a goal!" as Argentina's star player slotted home a penalty against France in Qatar last December, fans listening in on Veritone's YouTube channel may have assumed Smith was reacting to the match live.

What listeners actually heard was Alan's AI-generated voice clone providing real-time commentary from the live match's analytical data. The project, developed by Stats Perform and powered by Veritone's AI technology, lets fans listen to live game updates from a professional commentator on any device in local languages.

It's hard to fathom, but this is just one of the many ways that generative AI is helping media, entertainment and other industries create more engaging content that can connect with audiences anywhere in the world.

Generative AI also allows companies to more effectively scale content to engage increasingly diverse, multinational audiences. For those looking to further monetise content distribution, it's a significant opportunity.

Powering the creative engine

Voice cloning offers several applications for content creators. It can be used to narrate scenes or correct audio in post-production without the actor coming into a studio to record it themselves. It also has the potential to transform the dubbing and translation industries.

"You can use it for dubbing in an actor's unique voice in different languages," says Ashley Bailey, director of product marketing for AI voice at Veritone. "You could have Kevin Costner's voice, the way he speaks, his inflection, his tone, but in Spanish, Italian, French or whatever language to bring a more authentic experience to audiences."

The same can be said for podcasts, where AI can open up new markets by transposing a host's voice into different languages.

While AI can clone any voice, doing it ethically is critical to the medium's success. A Danish production company recently worked with Veritone on a documentary about a reality TV star that had passed away. They wanted his voice to narrate the programme, highlighting the legal issues around this application of AI.

"You absolutely need to get sign-off from the person you're cloning," says David Candler, senior director of customer solutions at Veritone. "If they are deceased, you need permission from their estate and the IP owner of the voice training data to ensure that it has the rights to reproduce the voice."

There are two main ways that AI can generate cloned speech: text-to-speech or speech-to-speech. "Using a speech-to-speech model where you speak in, and their voice comes out, you're capturing the intonation, speed, emotion. It's frighteningly accurate," says Candler.

Voice cloning can even allow companies to create new advertisements by brand ambassadors in multiple languages without them having to participate in additional recording sessions.

> **A media and entertainment business may have millions of assets that it can potentially monetise, but unless you can find them, you can't activate them**

"You can transcribe a campaign, translate it into multiple languages and create additional ads and assets at a scale that we can't do on our own as humans," says Bailey. "So, it goes way beyond having generative AI help you with ad copy to actually creating campaigns."

Scaling audience reach and revenue growth

Given the vast content archives that most companies possess, AI helps them discover and enhance their existing assets to expand both audience acquisition and revenue growth.

Content continues to grow exponentially. More content is uploaded to digital platforms in 30 days than what the major US TV networks broadcasted over the previous 30 years. In this climate, organisations that don't adopt AI to assist in optimising their extensive archives could be leaving bigger audiences (and money) behind.

According to Accenture, nearly a third of AI pilot initiatives are scaled beyond their initial scope to deliver outcomes across the business. 42% of those surveyed determined that the return on their AI initiatives exceeded their expectations last year.

Savvy investments now can bring recurring value, giving businesses more mileage out of their content and streamlining their operations by harnessing AI to replace time-consuming manual tasks.

"For starters, media companies can use AI voice models to produce foreign language versions of their back catalogues, helping expand their distribution footprint globally," Candler explains. And AI algorithms can also help them search through their archives faster, making it easier to find monitisable content.

"A media and entertainment business may have millions of assets that it can potentially monetise, but unless you can find them, you can't activate them," says Candler. "With Veritone's AI-based archive licensing service, for example, we can very quickly lift specific moments from vast archives and then sell them to documentary filmmakers, agencies or networks. It takes far too long to manually sort through the content to find what you are looking for."

Working side-by-side with human creators

Using generative AI, a large language model like ChatGPT can be built for a specific brand, which can then deliver relevant content and ideas.

"All of the generative AI components would be very specific to their domain," says Bailey. "It can take all the data of a content rights' holder and provide content recommendations and campaign recommendations on the back of that."

Although some content producers may be fearful of AI replacing them, these innovations are designed to support humans in their roles. With generative AI content, humans will still need to go in and edit what has been produced.

"For our voice work, we partner very closely with the voiceover community. It's not about taking jobs away–we augment human ability," says Bailey. "It helps content creators keep up with content demand, allowing them to connect with global audiences in a more authentic, personalised way."

As people and machines collaborate more effectively, generative AI can create operational efficiencies, improve audience experiences, and make new revenue streams available at scale.

The message is clear: AI doesn't make better content; it makes content better.

For more information visit veritone.com/generativeai

VERITONE.


Andrew Brookes via Getty Images

PHARMACEUTICALS

How ‘in silico’ testing is accelerating drug R&D

AI is helping pharmaceutical firms to drastically reduce the time and cost they incur in the lab on evaluating the medicinal potential of hundreds and thousands of chemicals

Emma Woollacott

Developing pharmaceuticals is a notoriously lengthy and expensive business. The median cost of bringing a new drug to market is £810m, according to research led by the London School of Economics, and it’s not uncommon for the process to take more than a decade from start to finish. What’s more, the failure rate is about 96%.

> **End-to-end artificial intelligence allows us to discover and deliver better medicines faster than humans can alone**

Much of this attrition occurs towards the final phases of the R&D process during safety testing – nine out of 10 drugs in development don’t make it through human clinical trials. But even the early stages of identifying a potential therapy that’s worthy of investigation are highly complex and time-consuming.

So-called drug discovery involves first identifying a molecular pathway or genetic variation that’s linked to a particular disease and then finding a molecule that can interact with it to halt, or at least delay, the progression of that disease. Until recently, this was a painstaking process of trial and error, with firms typically creating and screening several hundred thousand compounds in vitro to whittle them down to a couple of options deemed suitable for clinical testing. But in recent years the industry has started using machine learning to model the likely efficacy of these chemicals ‘in silico’ before doing any work on them in the lab.

The technology can evaluate the potential of candidate molecules against several hundred criteria simultaneously, sifting out those that show little promise while highlighting those worthy of further attention.

“End-to-end artificial intelligence allows us to discover and deliver better medicines faster than humans can alone,” says professor Andrew Hopkins, CEO of biotech firm Exscientia. “AI platforms make predictions based on their analyses of thousands of parameters in parallel, exploring a computational space far beyond our cognitive powers.”

Some of the first AI-recommended chemicals to enter human clinical trials were discovered in a process that took just over a year on average, he adds. Their development involved fewer than 10% of the prototype compounds that would have been tested using the traditional approach to drug discovery.

In another recent first, a new AI-aided drug designed to treat amyotrophic lateral sclerosis (ALS), a motor neurone disease for which there is no known cure, has entered clinical trials. The therapy – going under the name VRG50635 at this stage – was discovered by Verge Genomics using Converge, an AI platform it has developed.

“We used Converge to build an ALS ‘disease signature’ based on more than 11 million data points, sourced from almost 1,000 human tissue samples,” explains the firm’s chief business officer, Dr Jane Rhodes. “Our signature comprises more than 200 genes that are dysregulated compared with neurons from healthy brains and spinal cords. This gave us insights into the complexity of the disease and enabled us to discover novel molecular mechanisms that we believe can cause ALS. We achieved all this in half of the industry’s standard time.”

There have been other breakthroughs elsewhere in this field. For instance, computer scientists at the University of Sheffield recently worked with AstraZeneca to develop a new AI platform that, they say, provides enriched information about the interactions between drugs and their protein targets through a so-called bilinear attention network. The team has made the source code freely available to other researchers.

The lead researcher on this project is Haiping Lu, professor of machine learning at Sheffield. He explains that “our AI, called DrugBAN, learns the multiple pairwise interactions between the substructures of drugs and their protein targets to provide useful biological insights. These help researchers to understand what’s happening at a molecular level. DrugBAN contrasts with most drug-prediction AI, which learns from whole representations of drugs and proteins rather than their substructures.”

The use of AI as a drug discovery tool increased by 40% in 2022, according to research published in Nature in February, and uptake is only likely to grow over the next few years.

McKinsey estimates that almost 270 firms are already working on AI-powered drug discovery, mostly in the US, and predicts that many of them will seek to establish partnerships with well established biopharma companies.

“The next challenge we face is to scale up this approach,” Rhodes says. “We need to systematically apply the use of AI across the entire drug discovery pipeline and embed it effectively into business processes. By using large quantities of high-quality human data, we can reduce the need for animal experimentation and move towards a system of predictive modelling that can save both time and cost while increasing the probability of technical success.”

Some organisations are introducing AI into other aspects of the drug R&D process. For instance, pharma companies are using natural-language processing to sift through the vast archive of scientific literature, including academic papers and already-identified gene sequences, to help them detect and analyse patterns that could enable them to identify potential drug targets that might otherwise be missed.

And other AI-based tech, when fed with the right data, can not only help users to identify candidate drugs; it can also help them to predict which patients are likely to benefit most from such medicines. This will potentially lead to better-targeted, and therefore more effective, therapies.

In 2021, for instance, Exscientia demonstrated for the first time that an AI-driven precision medicine platform could propose which drugs would be most beneficial for people with late-stage blood cancers.

Such techniques can also help to optimise dosages for individual patients and even establish the best times of the day for them to take their medicine.

“In the future, such AI techniques will enable the development of personalised medicines, including vaccines,” Lu predicts. “By incorporating multimodal data including a patient’s genetic information and medical history in AI modelling, we can reduce the risk of adverse reactions and increase the likelihood of a successful treatment.”

The future success of AI-based drug development obviously depends greatly on both the quality and the quantity of data that is available to be crunched. But the increasing use of electronic medical records by health services is helping in this respect.

“We believe that using AI to design and develop better drugs faster will eventually become the model to create new medicines,” says Hopkins. “Within a decade, most drugs will be discovered and developed this way.”

> **Using AI to design and develop better drugs faster will eventually become the model for creating new medicines**

AI IN R&D

AI spend as a share of biopharmaceutical R&D

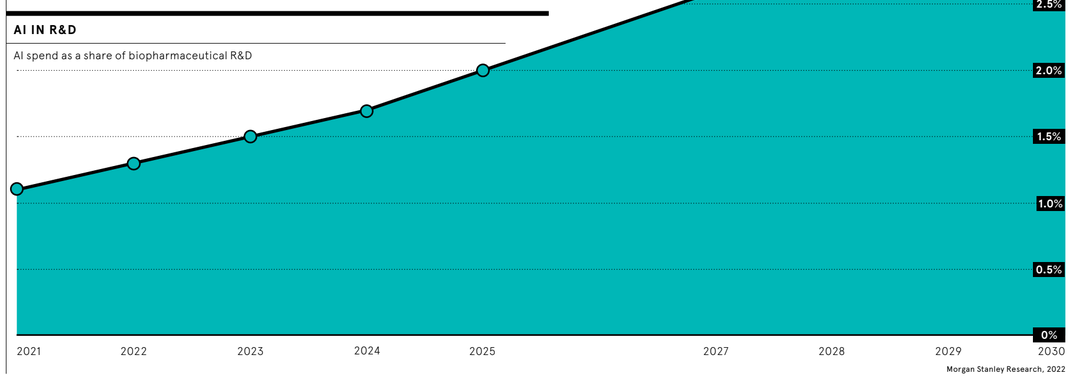

Morgan Stanley Research, 2022

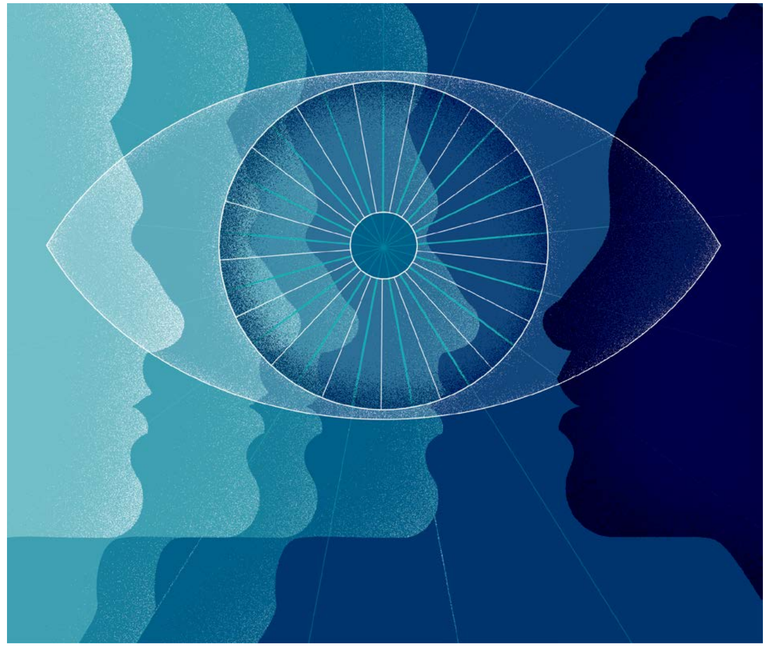

BEHAVIOURAL ANALYTICS

The long view

Companies are rapidly adopting AI to predict short-term consumer behaviour and maximise profits. But businesses could use the technology to take a long-term view of behavioural analytics and set customer-centric goals

Natasha Serafimovska

It's been over a decade since Amazon filed a patent for its anticipatory shipping technology (shipping an item before the customer knows they want it). It may sound dystopian, but the ecommerce giant isn't relying on a crystal ball or guesswork. Instead, it has looked at historical buying patterns, browsing habits, surveys and demographic data to predict what items will be in demand where. The result? Hard-to-beat prices and best-in-class delivery times.

Since then, AI and predictive analytics have taken centre stage with businesses that increasingly use AI to drive decision-making. The PwC 2022 AI Business Survey reports that 96% of business leaders who responded said that they intend to use AI simulations to improve business performance. The 'AI leaders' among them use AI to drive decision-making on technology (74%), operations and maintenance (62%) and customer experience (61%).

This is just the starting point of harnessing the power of AI to drive revenue. A proliferation of generative-AI tools already process customer feedback in real time and adjust their messaging to elicit the best response. As the technology develops, how can businesses make the best use of AI without inadvertently hurting their business or their customers?

Sulabh Soral is chief AI officer at Deloitte Consulting and leads the Deloitte AI Institute in the UK. He makes a clear distinction between AI and what most corporates are using today: "The broadest definition of AI is anything that can create rules that mimic things that we would identify as cognitive capabilities. But, if you look strictly at the question of behavioural analytics, a lot of companies still use classical machine learning."

This may sound like a distinction without a difference, but it's key in understanding the limitations of the behavioural-analytics models in use today. While deep learning and generative-AI models could generate insights from unstructured data like language and photographs, machine-learning models rely purely on structured data to make predictions. This potentially leaves a treasure trove of data unexplored.

This explains why most businesses today focus on short-term goals like revenue growth and cost reduction. But Soral believes that, as these technologies evolve, business leaders need to move away from this way of thinking and focus on long-term, customer-centric objectives instead.

"Most behavioural analytics today is based on how a response should be elicited, not whether the customer needs the product. But how would you train your AI model if the goal is customer happiness or the savings they could make in a decade?"

Generating new revenue isn't the only way that behavioural analytics could contribute to business growth. Ellen Loeshelle, director of product management at Qualtrics, an experience management platform, highlights that behavioural analysis can be just as valuable to protect existing revenue as it is to generate it.

"[Our product] takes behavioural data from the past and projects it onto things in the future. If you can tell us that this cohort of customers filled out a survey beforehand and then they quit, we can then use that as training data to project for future surveys that come in."

These insights can then help businesses take preventative measures and ensure customer concerns are addressed in time.

A new tool seems to emerge every day with the enticing promise that it can help you work better, faster and smarter. But Stefano Puntoni, a professor of marketing at The Wharton School at University of Pennsylvania, warns that automation isn't always the right answer to drive business performance: "It's important to understand the role that technology plays in people's lives to make predictions about how they're going to react to it."

For example, when the items we buy are closely linked to who we are as a person, automation can backfire. Think of a cooking aficionado who loves spending time in the kitchen. If you suddenly automate part of the cooking process, you have then created a problem, instead of a solution.

Another scenario where automation is less effective than human input, Puntoni says, is in interactions where customers are "assessed" by your business (for example, loan or credit card applications). Here, companies usually worry about how automation will handle negative outcomes for the customer, but Puntoni emphasises that it's the positive outcomes which require human interaction.

"People are happy when they get what they want, but they're happier when it's a human that gave them the news and not an algorithm," he says. "When it comes to rejections, however, we don't observe that being rejected by an algorithm or a human makes a difference."

> **People are happy when they get what they want. But they're happier when it's a human that gave them the news and not an algorithm**

As AI takes over more of our business interactions and decision-making, it's inevitable to think about the ethics around its use. In the PwC survey, 98% of respondents said they plan to make their AI responsible but fewer than half have planned to take specific actions.

David Wright is a partner in Deloitte's intelligent automation team and works with clients in the private sector. He thinks the rule of governance will be key as businesses rely on data to anticipate consumer behaviour. He points out that the "rule of governance is incredibly important, particularly if you start to automate off the back of AI predictions or recommendations".

While internal watchdogs are already present in highly regulated industries such as banking, insurance and healthcare, they're less so in hospitality and retail. Loeshelle thinks this is about to change.

"We're seeing it right now primarily in the States but it's starting to bubble up in the UK – [companies] insist on putting our technology through model risk management reviews," she explains. "This is an internal AI governance board that is responsible for evaluating the integrity and the purpose of each AI-related tech that the company is purchasing or building."

There's no doubt that AI will play a key role in how businesses and consumers interact in the future. But if companies are attempting to predict consumer behaviour with AI, it's important to understand the implications for consumer experience and the safeguards that must be in place from the outset.

PREDICTING CUSTOMER BEHAVIOUR IS THE NUMBER-ONE USE FOR AI AMONG MARKETERS

Leading reasons for marketers to use AI

MarketingCharts, 2023

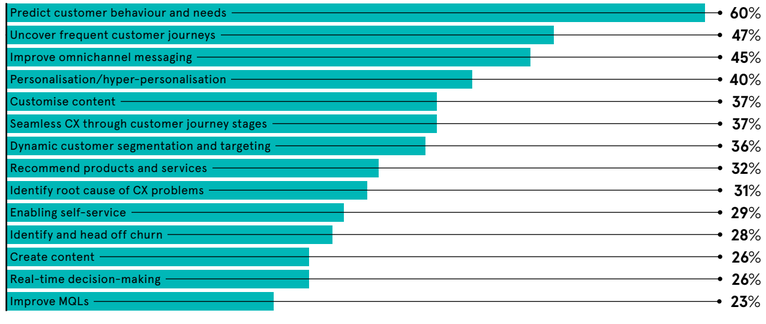

Commercial feature

Think first: why responsibility needs to be forefront when deploying AI

As the AI race heats up, no business wants to be left behind – and doing things properly will yield even bigger benefits

The AI era is upon us, with what seems like new advances every week, pushing the technology to new heights. Between Google, OpenAI, Microsoft and a raft of other companies, new developments that can ease the way we live and work are accessible to people more than ever before. It's little wonder, then, that businesses are starting to consider how best to integrate AI into their processes to reap the benefits.

But thinking before acting is vital in such a fast-moving space. The first-mover advantage that businesses seek out can quickly be negated by the regulatory risks of irresponsible use of AI.

"Lots of companies talk about AI, but only a few of them can talk about responsible AI," says Vikash Khatri, senior vice-president for artificial intelligence at Afiniti, which provides AI that pairs customers and contact-centre agents based on how well they are likely to interact. "Yet, it's vital that responsibility be front of mind when considering any deployment of AI – the risks of not considering that are too great."

Think fast, act slower

In part, the fast moving and competitive environment often places the responsible use of AI secondary to gaining market share. The history of AI, says Khatri, has seen companies develop tools that harness the power of AI by making use of big data sets without fully considering what impact they can have on society. Widely used AI tools are trained by trawling the internet and gleaning information from what is found online, which can often replicate and amplify our societal biases. Another problem with AI generated content is that it is often ill-suited to the specific needs businesses may have when deploying AI.

"If I'm a broadband provider in the UK, as opposed to a health insurance company in the US, there's a specific way that I communicate with my customer," says Khatri. "With respect to the generative AI technology that's receiving so much attention, it's important that the AI models being used are trained on the company's own data, rather than relying solely on generic, third-party data. That way, the organisation remains compliant with global data regulation and the AI models generate content that aligns with the company's unique approach to its customers."

Khatri points to how a customer service chatbot trained on the way users interact with one another on social media, for instance, could quickly turn quite poisonous rather than supportive, lobbing insults rather than offering advice.

"At Afiniti, we use responsible AI design to make those moments of human connection more valuable," says Khatri. "That in turn produces better outcomes for customers, customer service agents and companies alike. One way we do this is by training our AI models only with the data we need, and we continuously monitor them so our customers and their customers get the results they want, while being protected from bias or other discriminatory outcomes."

It's not just the risk of alienating customers that should be at the forefront of a business leader's mind when considering how to roll out AI within their organisation and to their clients. Regulation is on the horizon for AI, and is likely to bring specific requirements for how data is fed into models that are used to give AI its 'brain', and how AI is used to handle customer interactions.

Caution avoids consequences

"Before you even start to develop or deploy AI, you must be cognisant of the regulatory landscape," says Kristin Johnston, associate general counsel for artificial intelligence, privacy and security at Afiniti. "This means examining your governance structure around data compliance to get your house in order first."

AI regulation is complex and constantly changing, and a patchwork of laws across the globe can make it hard for businesses to comply. For example, businesses operating in Europe have different requirements from those with customers in the US, while the UK's data protection regulation is likely to soon diverge from the European Union's.

The magnitude of the task in responsibly deploying AI is something most businesses have yet to fully wrap their heads around, fears Johnston. "A lot of companies haven't built out a governance process specifically around AI," she says. To do so properly, Johnston says it's important to consider, first, the definitions of 'AI' and 'machine learning', then to identify how AI is being used within the organisation based on those definitions, and to construct your responsible AI programme accordingly so that all employees are aligned.

AI is set to become so ubiquitous that external services that feed into your company may use AI as well. For instance, Google has now introduced generative AI–powered aids to develop documents and slide decks in its cloud-software suite that your employees could soon find themselves inadvertently using without knowing it. And if people in your company aren't sure what AI is – or even if they're using it – you can't be confident your approach to AI is responsible.

> **At Afiniti, we use responsible AI design to make those moments of human connection more valuable**

Root and branch reform

Johnston stresses that a clearly understood definition of AI within your company is the basis of any AI governance programme. She recommends considering the definition of 'AI systems' in the artificial intelligence risk management framework published by the National Institute of Standards and Technology (NIST) in the US as a working definition.

"Making sure everyone is aligned is critical, because you want to check for any use of AI throughout your organisation," she says. "Any protocol worth its salt needs to be able to categorically define who is using AI tools, when they're using them, what data they're using and what the limitations of the tools are. It's also important to ensure AI tools are being used in a way that respects privacy and intellectual property, given the mounting legal actions against some generative AI tools by those who believe their data was used to train the models that power such platforms."

Doing this work in making sure responsibility is front and centre of any AI deployment is vital because it will avoid headaches in the long run. Not only can the irresponsible use of AI lead to trouble, but generative AI's tendency to 'hallucinate' content – in other words, generate untrue responses – could lead to even bigger trouble in the court of public opinion for spreading disinformation. Yet fewer than 20% of executives say their organisation's actions around AI ethics live up to their stated principles on AI. By putting in place a robust responsible AI programme, companies can avoid the pitfalls that come with leaping headfirst into the promise of AI without considering its drawbacks. "We're very mindful about ethical and responsible use of data," says Johnston. "Responsible AI should be a priority for organisations globally."

SymphonyAI

Enterprise-grade AI at scale

AI software to detect and prevent money laundering

Stop financial criminals in their tracks with the NetReveal-Sensa suite

SymphonyAI
NETREVEAL
netreveal.ai

SymphonyAI
SENSA
symphonysensa.com

Protect your business with SymphonyAI