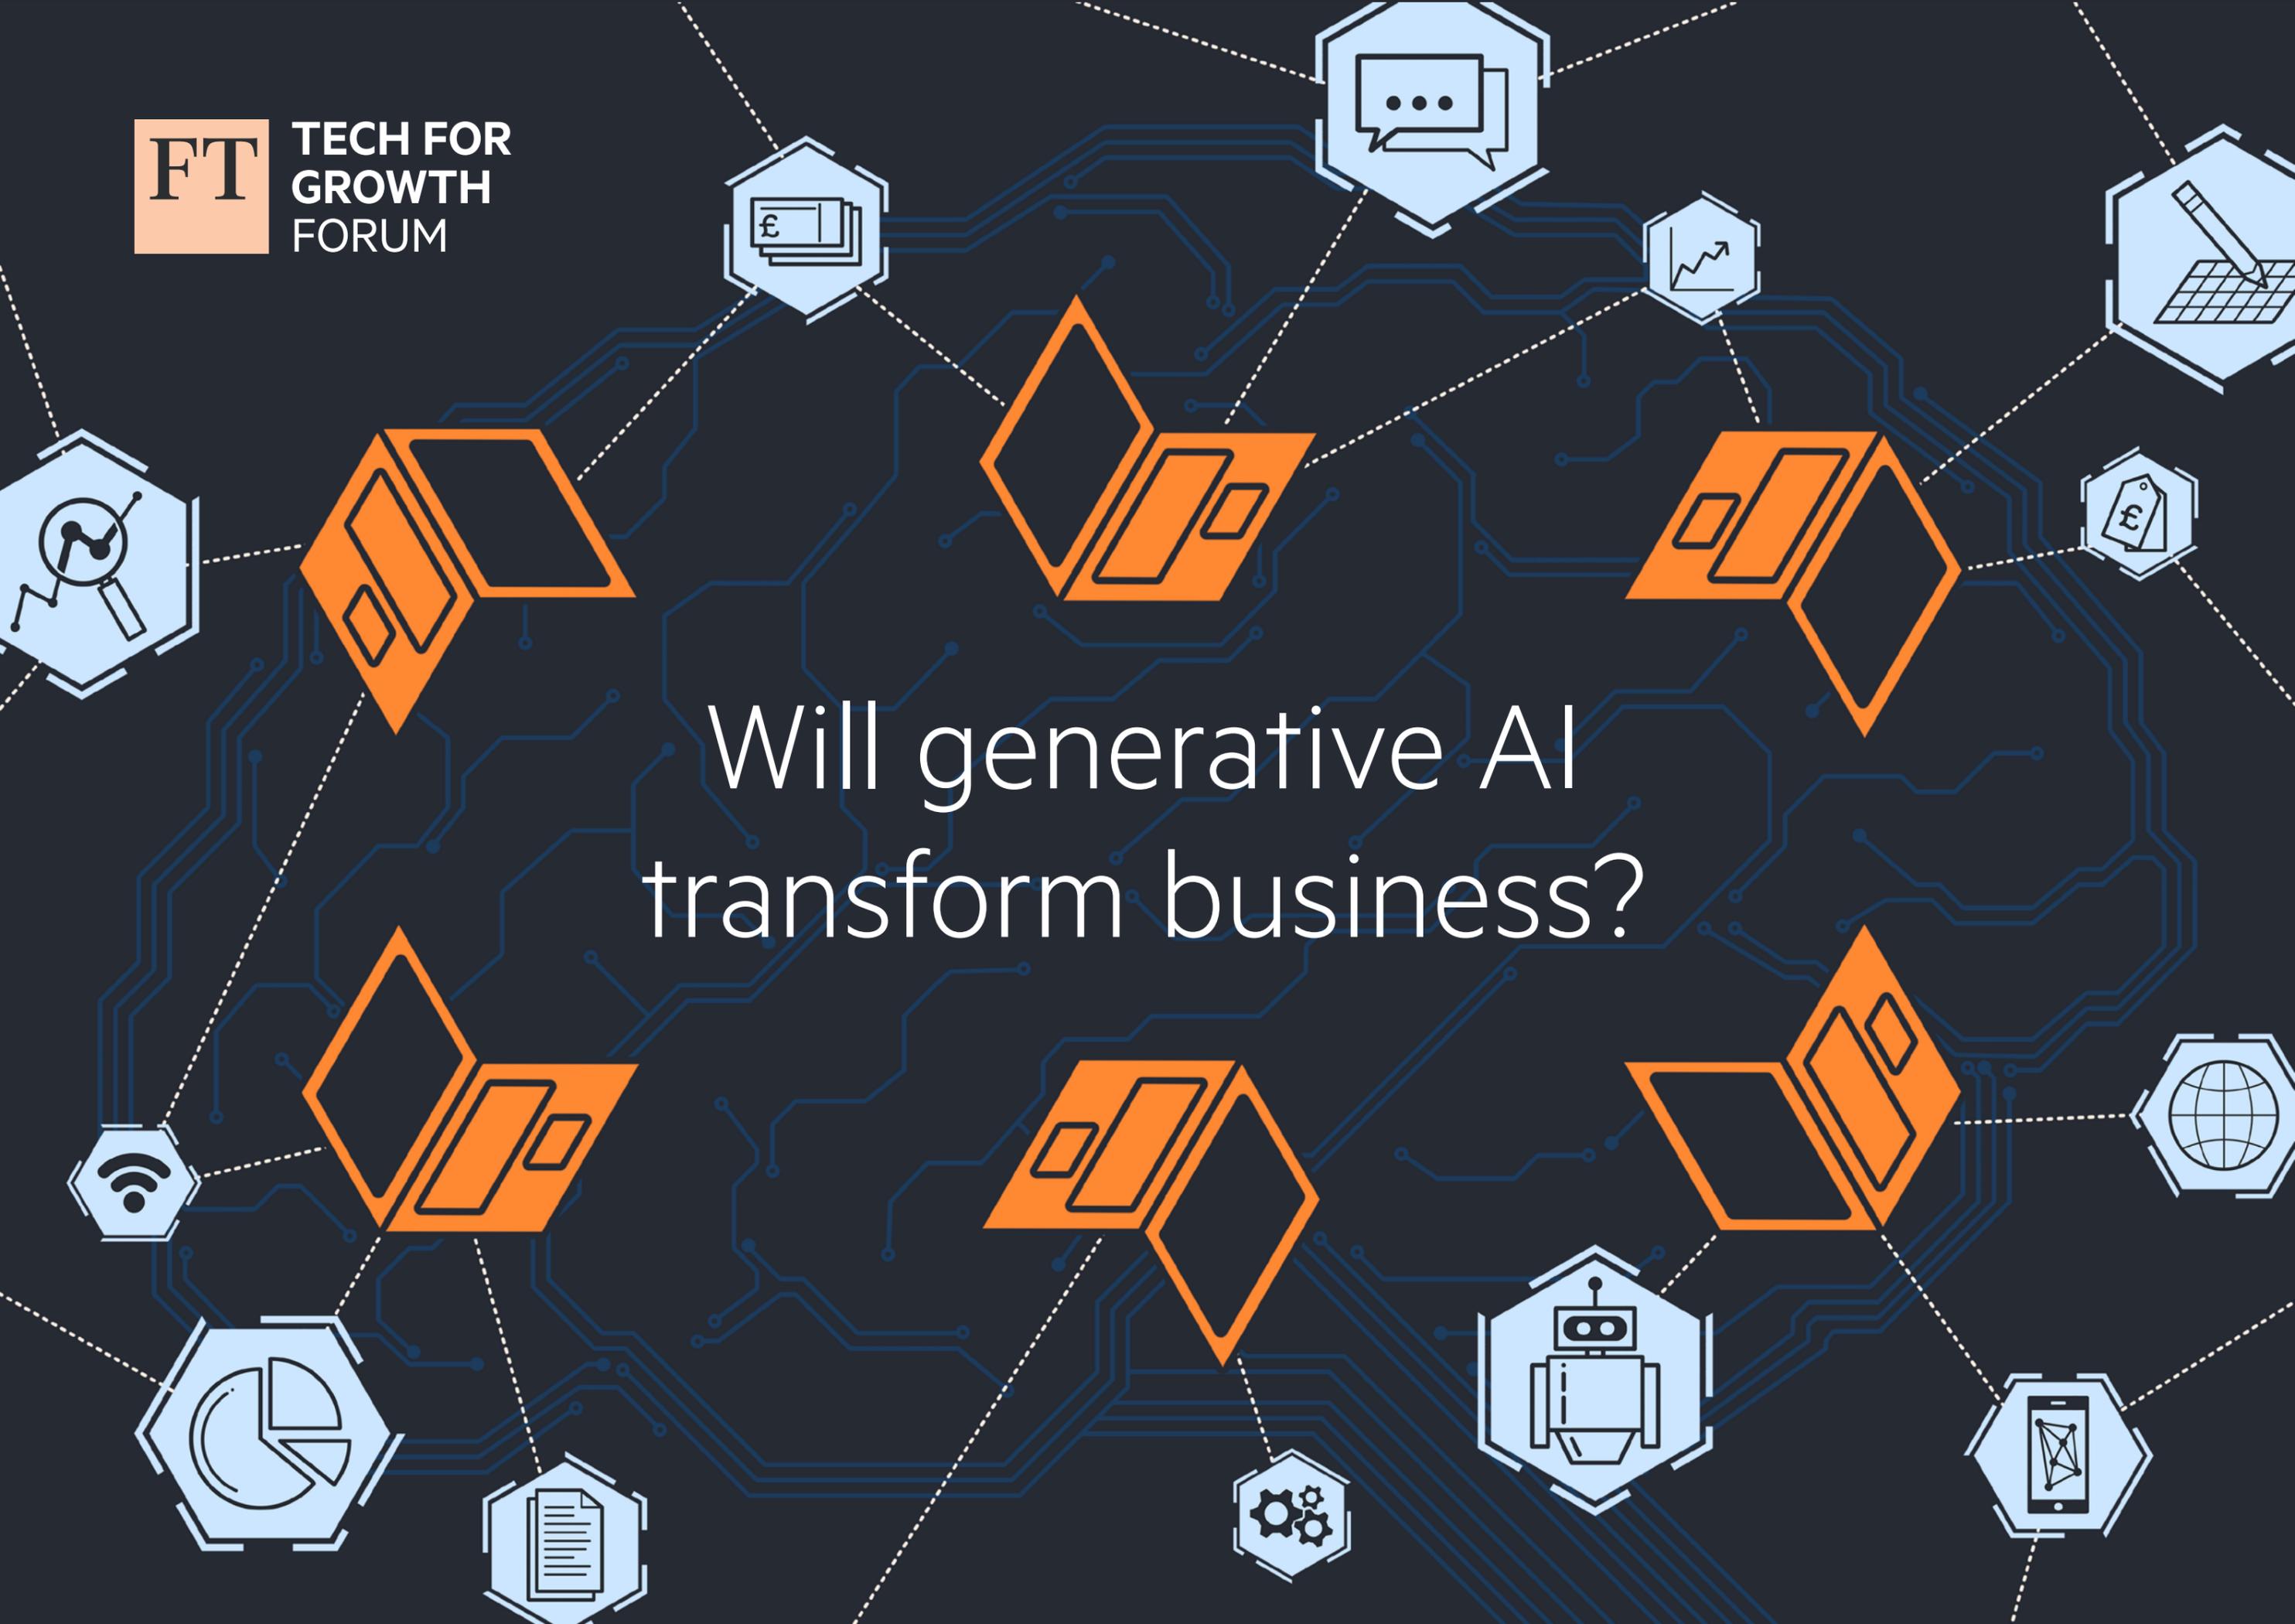


Will generative AI  
transform business?



The FT Tech for Growth Forum is supported by

# HCLTech

### • What is generative AI?

Generative AI is a set of algorithms based on foundation models, a term that the Stanford Institute for Human-Centred Artificial Intelligence says “underlines their critically central yet incomplete character”. Such models are “trained on broad data (generally using self-supervision at scale) that can be adapted to a wide range of downstream tasks”. The data behind generative AI programs such as ChatGPT and Google Bard are sourced from across the internet. It is a huge set of training information. One example is Dall-E 2, the text to image generator from OpenAI, which was trained on 650mn images.

A White House white paper gives this summary: “The power of AI comes from its use of machine learning, a branch of computational statistics that focuses on designing algorithms that can automatically and iteratively build analytical models from new data without explicitly programming the solution. It is a tool of prediction in the statistical sense, taking information you have and using it to fill in information you do not have.”

### • What can it do?

Based on the training data, generative AI can “create” content by autocompletion based on predictive sequencing or probability. In the same way that it uses likely word sequences to build sentences, generative AI employs pattern recognition and reproduction to create images, video and audio. In the case of image generation, given a text description it predicts the most relevant image. It can also summarise and answer questions about data, including text.

### • How can it be adapted?

By augmenting the training data with narrower or specialised inputs, generative AI can be fine-tuned for specific tasks. It can create marketing content in a given style or be adapted to suit an individual recipient. It can collate swaths of information and distil these into recommendations based on a specific input, for instance travel itineraries or consumer products. It can analyse data and come up with novel conclusions.

### • What can it not do?

AI cannot reason – and due to its “autocomplete” nature, users can never be sure of its factuality. It is a statistical beast and not a database. Ward says: “[The models] are just piecing words together and then using statistics to pick the next words off of that string – but it really depends on what you ask it, how you phrase it and that changes the outcome.” Ask how British company Dyson is dealing with the space vacuum – or even what time a particular shop opens – and there is a high chance that, in its current state, you will receive a confident but nonsense answer unsupported by facts.



# Will generative AI transform business?

Industries expect demand for quality control and human oversight of AI-generated content to grow, writes *Lucy Colback*

**G**enerative AI is developing rapidly both in terms of usage and sophistication. Since the public release in November 2022 of ChatGPT, OpenAI's large language-based chatbot, the number of people who have experimented with generative AI has grown rapidly. In its first five days the platform had more than 1mn users.

Insider Intelligence, the research and insights provider, forecasts that by the end of 2023, 25 per cent of internet users in the US, nearly 80mn people, will deploy generative AI at least monthly, up from 8 per cent at the end of 2022. It says that proportion will rise to 33 per cent in 2024. Most people will experiment with AI in the office, Insider Intelligence says. Unlike many new technologies more users are aged 55 to 64 than aged 12 to 17 as usage is concentrated in the workplace.

The AI phenomenon is not limited to America. A Salesforce survey of 4,000 people in the US, UK, Australia and India found that half of them had used generative AI. India was the leader with nearly three-quarters of respondents having tried

the technology. Since the public release, the functionality of generative AI models has improved. Their understanding, reproduction of natural language and accuracy are all better.

Christian Ward, the chief data officer of Yext, the digital experience platform, says the fact that generative AI can understand human language has helped. He says that when internet search evolved, users learnt to "speak keywords... bending their behaviour to computers". With generative AI people no longer have to be tech-savvy despite its advanced technology. AI is "human-savvy" and allows us to communicate with it naturally, making it more accessible.

"It used to be that computer-savvy humans had an advantage," Ward says. "Now that computers are human-savvy, that advantage is being democratised. That is absolutely one of the biggest breakthroughs we have seen in technology in a very long time."

He adds that once search functions reach the point where we use them by speaking to them as if they are human, "that construct [opens up] in a way we've not experienced before".

## Commercial potential

Michael Wooldridge, a professor of computer science at Oxford university, calls 2023 a "watershed year" for generative AI. He says the technology will be as transformative as the microprocessor, which made possible the personal computer, the web and smartphones. While those took months or years to catch on, the pace of development of generative AI is exceptional and it is hard to envisage its full effect.

Despite the technology's youth it has already penetrated the commercial sphere. Estimates of the value it could add to the global economy vary from the 7 per cent of gross domestic product (\$7tn) over 10 years predicted by Joseph Briggs, a Goldman Sachs economist to the \$2.6tn to \$4.4tn a year forecast by McKinsey's (as a comparison UK GDP in 2022 was \$2.7tn). Banking, life sciences and high tech could be the sectors that experience the greatest gains. Marketing, sales and customer experience will all be profoundly affected.

AI's use of natural language means that its large language-based versions are especially well suited to creative endeavours. Today's tools can help to produce text, audio, video, music, code, images and design. For instance, based on models that draw on troves of online data, generative AI has sufficient grasp of the variety of directions in which a sentence can go that it appears to be able to create entirely original work an infinite number of times.

Besides the written word, generative AI will create images from textual descriptions. These are more realistic than thought possible and can even incorporate 3D perspective and shadows. In audio and music the potential for new work that features existing artists — without their involvement — is well documented. AI-generated output has led to copyright concerns but then, given allegations of plagiarism especially in the music industry, there is doubt that humans create original work every time.

Arguably the applications for generative AI are greater in the creative field than in science (think research and engineering) where accuracy and reliability are more critical. This upends the expectation of a few years ago when machines were slated to add most value in heavy or rote tasks. It had been predicted that they would augment humans' physical capabilities or analyse data at greater speeds than we could ever process. Creativity was expected to remain the preserve of humans: some people said it was the defining essence of what it meant to be human.

With its unexpected creative capability, generative AI can take on jobs once considered untouchable by technology. Goldman Sachs estimates that a quarter of work tasks in the US and Europe could be automated by generative AI. At the upper end of its potential, more than 40 per cent of tasks in administration and the legal profession could be automated. This compares with less than one-tenth of tasks in physically intensive professions such as construction and maintenance.

## Who is in the field?

In its analysis of the generative AI market, Acumen Research and Consulting lists D-ID, Genie AI and Rephrase.ai among the leading players, together with more recognisable names such as Amazon Web Services, Adobe, Google, IBM and Microsoft. Besides these heavyweights, applications have mushroomed to serve specific industries such as marketing or to perfect functions including video or audio generation.

In the chatbot arena ChatGPT and Bing Chat, both based on OpenAI's large language models, compete with Google Bard. Although the technology is relatively new it has spawned lots of applications in the creative industries. For instance there are tools to supercharge communications and marketing. These include Jasper and Signum.ai, which itself provides a list of other players such as Phrasee and Clarifai.

CB Insights has listed more than 360 companies that specialise

in generative AI and that number will only grow. The market intelligence provider says that in the first half of 2023 investment into generative AI start-ups went up 5-fold, from the \$2.5bn invested for all of 2022 to \$14bn. Even excluding the \$10bn raised by OpenAI, the inflows are still considerable, as is the start-up activity. Conviction, a new AI venture capital firm, received more than 1,000 applications for 13 places on Embed, its start-up accelerator programme that began in early September.

## Get with the strategy

Even though it is early days for generative AI, BCG, the consulting group, believes that "to be an industry leader in five years, you need a clear and compelling generative AI strategy today".

This is no exaggeration. Already, there are universal corporate functions, for instance knowledge-sharing, communications and recruitment, that can be enhanced and made more efficient with generative AI. As the use of generic tools becomes more common, the way to keep a competitive advantage will be to develop specialised, in-house tools trained on significant quantities of proprietary data.

Deloitte says that data-rich industries will integrate AI faster than those that rely on judgment. It says companies that build vertical-use cases and industry solutions will be able to add more value than those using general purpose models. As with platform models, such solutions could either be deployed in-house or provide new revenue streams from external licensing.

All that said, ManMohan Sodhi, professor of operations and supply chain management at Bayes Business School, London, said in the FT report on supply chains that companies should first identify the problem or goal to which a technology is suited rather than adopt a technology then look for a use. With generative AI, Gartner advises that "end users should be realistic about the value they are looking to achieve". The risk is that the necessity for human validation can offset any efficiencies gained.

## Modify or build?

There is no straightforward answer to the question "modify or build?". Tweaking the general model to create a niche application is at least not too costly. According to BCG, fine-tuning (providing a curated data set on top of the original training data) can be cost effective. The consultancy highlights a 2022 experiment in which Snorkel AI spent less than \$8,000 to fine-tune a large language model to complete a legal classification. Designing questions to extract relevant information, a technique known as prompt engineering, is another way to refine outputs.

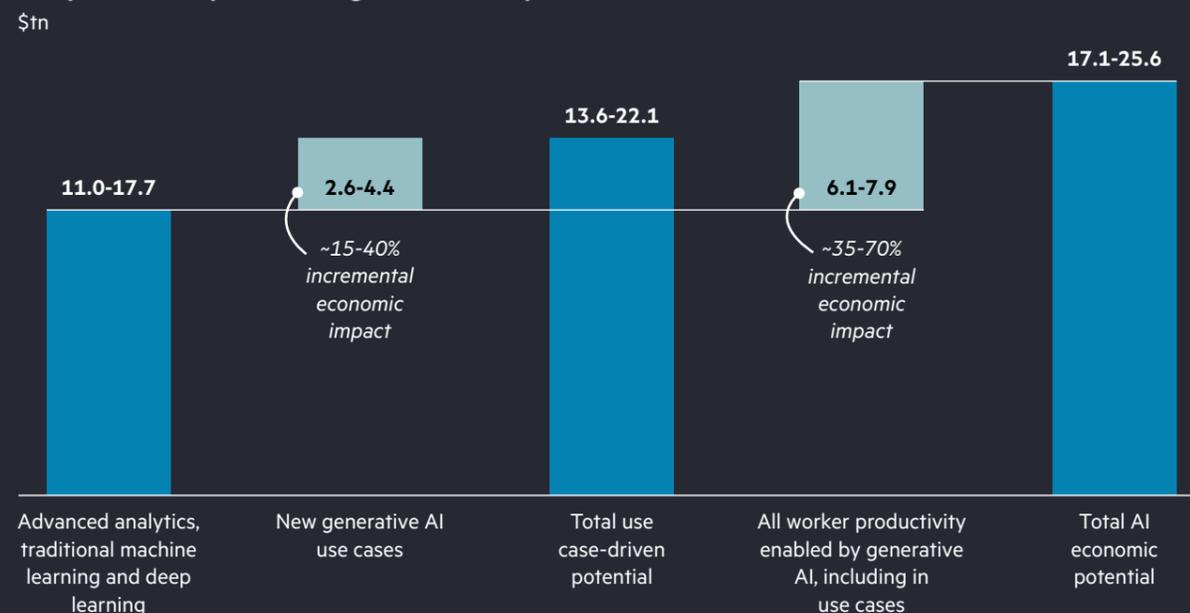
Wooldridge stresses that while there is "a huge range of opportunities" in fine-tuning, it comes with "one big caveat: you have got no clue what data went into that". An organisation cannot be certain that an off-the-shelf model reflects its values. Wooldridge says: "You ought to be very concerned about the data this thing has been fed on...it will have absorbed the whole of Reddit and the whole of Twitter and every toxic and vile opinion out there...so it's not a complete solution, being able to fine-tune things."

An alternative to fine-tuning is to build a model based on proprietary data that is stored on-site. This can be more secure and flexible but Ward advises against it. The requirements for clean and legally compliant data that feed such a model are huge, he says. The data, resources and time needed make it an expensive endeavour, too. BCG estimates that costs could run into millions of dollars.

Whatever the approach, companies have to be certain that any valuable proprietary data are protected and secure as well as clean and legally obtained. When it comes to output, Ward and Wooldridge agree that human oversight is essential.

Generative AI is already producing a greater volume of more varied output than humans can achieve alone. Because of the risk of errors, however, human oversight and quality control are essential.

## AI's potential impact on the global economy



Source: McKinsey

## Use cases

### • Consumer facing

#### Marketing content and strategies

Amplifying or creating marketing content and strategies can work well where experimentation (or even AI hallucinations) is a creative bonus and where accuracy is not critical. Through fine-tuning based on a company's content, marketing materials can be generated at scale in a tone that maintains brand and message consistency. These can also be easily translated into numerous languages. Such communications can be personalised by using customer data. Content can be text, image or video based.

Strategies can be tested and tweaked for optimisation based on engagement and conversion data. Gartner predicts that by 2025, 30 per cent of large corporations' marketing messages will be generated by AI.

#### Customer service and support

AI can provide chatbot responses to customer queries in natural language. Human interaction can come later for more complex queries. AI can also help with retrieval of client data and internal enterprise knowledge, both to anticipate queries and customise responses.

#### Film and video

Sarah Guo, the founder of Conviction, believes that video content is a noteworthy growth opportunity. "High quality video at lower cost means there will be more of it," she says. "Any story you want to publish you could have it described by an avatar and in any language."

The cheapness and ease with which video can be created from text prompts will facilitate the creation of short films for anything from advertising to internal information dissemination. On the big screen, Gartner predicts that by 2030 we will see a blockbuster film that is 90 per cent created by AI.

#### Teaching

AI is excellent at explaining things. It will produce different iterations to find an answer that resonates most with a learner. It can then create personalised lesson plans to match the learner's style as well as help with broader course design. Duolingo, the language learning app, uses AI to create responsive dialogues and explainers that enhance the learning experience.

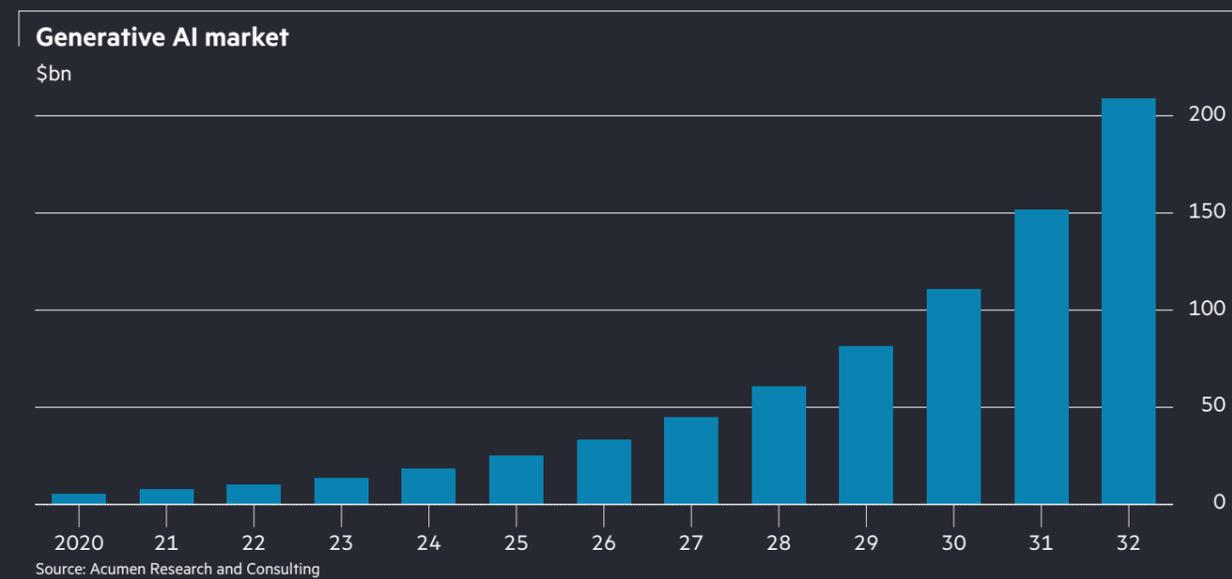
#### Search functions

Searches that use AI can provide comprehensive answers on anything from product comparisons, which increasingly come with references to aid fact checking, to suggestions for travel itineraries and recipes.

#### Journalism

AI's first forays into the creation of news content have been unsuccessful. Subramaniam Vincent, director of journalism and media ethics at the Markkula Center for Applied Ethics at Santa Clara University, California, has pointed out the drawbacks of using generative AI for reporting.

By ChatGPT's own admission the technology cannot tell the difference between fact and fiction – it simply makes informed guesses about plausibility. In addition, there is a recency issue: large language models are trained to a cut-off date, so they are unlikely to be able to find current news. If given the facts with which to produce an article, AI could, however, do a credible job at simulating an author.



### • Internal

Hype about generative AI generally outruns implementation. Plenty of company managements talk about it but few offer evidence of having used it. Ultimately businesses may gain more value from experimenting with AI internally than in a consumer-facing area. For the early movers the goal is not simply efficiency gains or cost-cutting. A McKinsey survey found that organisations that employ generative AI were more likely to be using it to create new businesses and increase revenue streams.

#### Brainstorming

Generative AI can create numerous suggestions based on limited input. Wooldridge says that in-person brainstorming "is a terribly difficult thing". He says: "One of the cool things about this technology is it will brainstorm for you. You can feed it some rough drafts of ideas and then keep pressing a button. 'Give me another idea' and it maybe gives you 10 ideas, but you have one that grabs your attention." Particularly in science, where marrying ideas and techniques across disciplines can take a leap of imagination, generative AI is likely to offer new insights.

#### Product design

Given enough prompts, generative AI can design new products or personalise existing ones based on customer data. For example, Under Armour and Nike, the trainer/sneaker makers, have both experimented with AI for design. In time, generative AI may inspire or come up with novel products, processes or services ideas.

#### Writing software

AI can write a first draft of a code based on natural language prompts ("I'd like software that does this") or edit and help to find bugs in current software. It can also run and test programs to ensure that they achieve the intended outcome or to optimise code. It is important to have realistic expectations. Wooldridge says: "My experience is that it's very good writing short, routine programs at the moment", which is "an awful lot of what programmers do". He adds, however, that it is far from "superhuman" and that its output needs human oversight.

#### Medical innovations

Important breakthroughs in medicine are likely to be accelerated by AI. Already the technology can identify candidate molecules for therapies, a technique especially suited to proteins. There is also the potential for it to consider symptoms and find a connection with an illness that has been missed by a human. AI can identify the best candidates for drug trials and accurately match treatments to patients. The first AI-designed drugs are on their way.

#### Internal knowledge banks

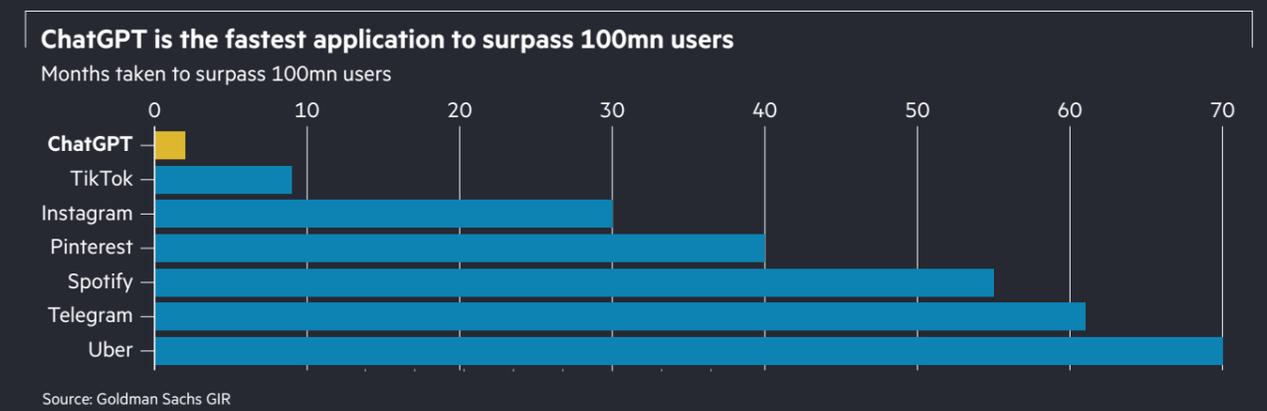
Large quantities of textual or numerical data can be collected, collated and summarised using AI, which will also translate the results at speed. Notable applications range from summarising manuals to locating corporate knowledge and creating knowledge repositories. Ward believes knowledge banks could be transformational for company intranets. "Most people that start at new companies, especially in a distributed workforce, don't know who to reach out to ask for guidance – but this thing will know, based on what it knows about the inside of the company. "I think intranets, which have historically been horrific in terms of usability, are going to skyrocket in use because of this type of access."

#### Data analysis

Given any set of data, generative AI can create charts and find trends that a human analyst might not think to look for.

#### Presentation creation

A variety of AI programs can already put together slides based on a basic text brief and information gleaned from the internet, or from uploads of more complex documents or data.



## • Sector specific

### Academic and research papers

Students' use of ChatGPT to write university essays has been well documented. The intention is not necessarily malicious. Wooldridge has sympathy with those who simply wish to improve their writing. "We'd all love to read better-written scientific papers [given that] the average standard of scientific paper is dull beyond imagining," he says.

Papers so doctored, however, come with a health warning given the propensity for generative AI to make things up. "We know that this technology hallucinates and gets things wrong and it introduces that possibility. I'd rather have a paper where I was confident that I was listening to what the author said," Wooldridge added.

### Legal profession

AI can speedily analyse and condense screeds of documents from contracts to case law. As such it is likely to be of great value in the legal profession and it is already deployed by numerous firms. It can help with drafting, standardising and rationalising contracts, performing due diligence, drafting and compliance.

In March PwC UK, the professional services group, entered a deal to develop generative AI with OpenAI's Harvey. In a similar vein non-law firms can use generative AI to create requests for proposals and analyse which contracts and services generate the most revenue. Note though that using AI to create a legal case without checking is generally regarded as a bad idea.

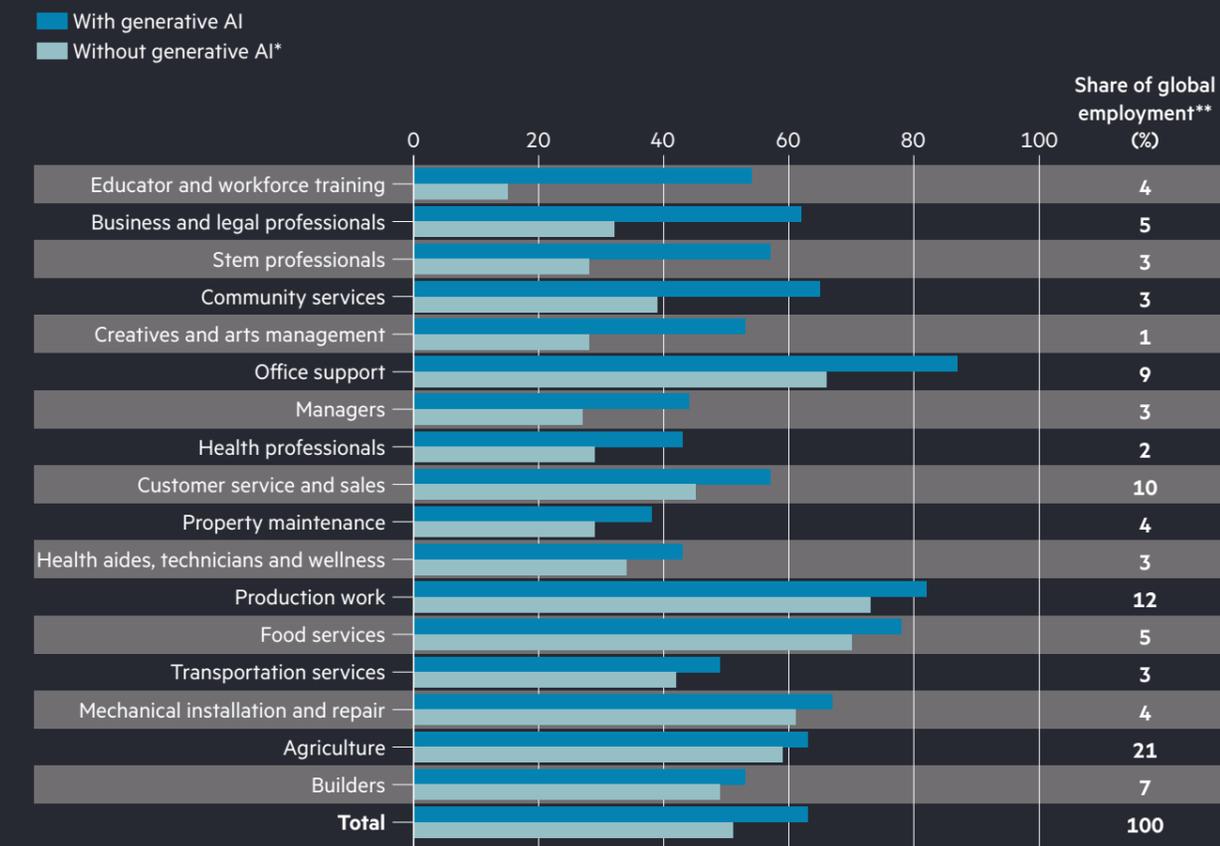
### Finance

Generative AI is already being deployed for financial analysis. Guo mentions a company that has built an app to understand the earnings reports, interviews and stock exchange filings of competitors. By using generative AI "they can grasp a lot of information for their company's use", she says.

Guo expects that fund managers will also develop apps to screen for certain criteria in reports and accounts, or to analyse internal research memos and subsequent investment decisions. This would help them to find new companies and replicate previous successes. For internal use, generative AI can extract and summarise information from documents, helping with "non-standardised book-keeping".

## Impact of generative AI on technical automation potential in midpoint scenarios

Overall technical automation potential, comparison in midpoint scenarios (% in 2023)



\* Previous assessment of work automation before the rise of generative AI Source: McKinsey \*\* Includes data from 47 countries, representing about 80% of employment across the world

## • Drawbacks

### Prone to errors

While generative AI programs have advanced rapidly and results from ChatGPT4 have already surpassed those from the initial release, the content that large language models are designed to produce is what is statistically probable rather than factually correct.

Wooldridge says that while it is a good tool, companies would be "criminally negligent" to let loose AI-generated software code without checking it. "That would be an extraordinarily dangerous thing to do, not because I think it would destroy the world but it might trash your employer's hard drives," he says.

Reputational risk is also high. McKinsey's survey finds that the proportion of companies that are aware of this is far higher than that trying to mitigate it.

### Public perception

Ipsos data show that consumer distrust is high and rising but this may improve as use of the technology becomes more common. Data from Pew suggest that in industries where exposure is higher, the understanding is greater and anxiety is lower.

Research from KPMG shows that the young, the university educated and managers generally hold more positive attitudes towards AI, as do those in emerging economies such as Brazil, India, China and South Africa, where the benefits are perceived to outweigh the risks.

### Skills obsolescence

One side-effect of generative AI's ease of use is that the skills required to develop technology more generally are likely to change. For instance the demand for quality control and human oversight of AI-generated content will grow, while that for humans who can originate content is likely to fall.

The landscape is changing so quickly that even prompt engineers – the programmers who optimise questions to get the best response from foundation models – could soon be obsolete. In addition workers who currently help non-technical people understand computers will be redundant.

Generative AI can of course help to devise retraining courses. Critical thinking will be an increasingly valuable tool.

### IP and copyright concerns

The issue of intellectual property will be increasingly important for "created content", especially with generative AI being able to write or draw "in the style of".

Authors are ever more vocal about copyright infringement over the use of their works in the training data set. They also fear risks to their livelihood from the artificial generation of works. In September writers including Jodi Picoult, John Grisham and George RR Martin added their names to the list suing OpenAI for unauthorised use of their work. A five-month writers' strike in Hollywood was resolved in September with a deal which included guidelines on the use of AI.

Microsoft has sufficient confidence in its guardrails against any infringement that it has pledged to provide legal protection for users sued for copyright breach.

### Misuse

As well as the drawbacks above, generative AI has the potential for misuse. Wrongdoers can use it for anything from disinformation at the state level, including creating and spreading deepfakes, to honing phishing scams at an individual level. Actions that used to be limited by money or time are now achievable with generative AI. A UK parliamentary committee on large language models noted that the cost of election tampering had fallen from about \$10mn to a few thousand dollars.

Cyber and data security is increasingly problematic and Ward cautions that "there can be a lot of issues with letting this thing run around on your data". Problems range from potential leakage of proprietary information to falling foul of privacy laws.

Regulation may be one way to counter some of these aspects. Wooldridge says that even if governments establish some sort of authenticating body to ratify models, any tests they use must be credible, new each time... and something the machine will never read on the internet. Allowing Big Tech to self regulate is not a safe solution.



### What next?

With such a new and fast developing technology it is almost impossible to anticipate what comes next. The generation that grows up with generative AI may devise entirely new uses for it, similar to the way in which the digital native generation has reimagined how and what can be sold via channels such as social media and YouTube. One thing is certain though: business managements will need to keep up with developments to know when and how they may be able to use generative AI.

## Glossary

• **Large language model/LLM:** a deep learning algorithm capable of taking in vast amounts of content and producing humanlike writing or realistic imagery.

• **Deepfake:** Realistic fabricated media, often video, which purports to show a prominent person saying or doing something they have never said or done. One of the most well known is this deepfake of President Barack Obama from BuzzFeed Video.

• **Guardrail:** A programming feature that allows or instructs generative AI not to follow certain commands or perform certain tasks, for example copying someone's work or their creative style, or generating hate speech.

## Premium Member

The FT Tech for Growth Forum is supported by *HCLTech*, our premium member, which helps to fund the reports.

*HCLTech* shares its business perspective on the forum advisory board. It discusses topics that the forum should cover but the final decision rests with the editorial director. The reports are written by a Financial Times journalist and are editorially independent.

In the adjacent piece, *HCLTech* gives its view on generative AI. Its views stand alone. They are separate from the FT and the FT Tech for Growth Forum.

### **Generative AI: A Transformative Force with a Bumpy Road Ahead**

Ashish K Gupta  
Chief growth officer, Europe and Africa, HCLTech

The hype surrounding generative AI is at a peak. Businesses that look beyond this point see it as a revolutionary technology that has the potential to create great value but they also recognise that, if not managed well, it can cause harm.

That said, companies that embed the technology intelligently into their organisational processes, that have the courage to reinvent their markets and the grit to stay the course, will emerge stronger, more productive and unbeatable in their fields.

#### **Generative AI adoption**

In the business-to-consumer realm, the level of adoption of generative AI is already impressive. The fact that AI can communicate with users in their native languages makes it vastly accessible. It is no exaggeration to say that this technology has the power to democratise the world. Consumer-facing services have already experienced the effect of generative AI in areas including education and the creation of music and content.

A good way to think about the rate of adoption of generative AI in B2C is to consider the time taken for the internet to permeate society. Even after 25 years, ecommerce penetration is only at 30 per cent to 40 per cent. The adoption of generative AI is likely to be much faster: in some digital services markets such as music, it could explode. Deep and wide penetration will take time, however. We are at an early stage and many lessons will be learnt.

In business-to-business settings, the potential benefits of generative AI are significant but challenges have to be addressed before the technology is deployed at scale. Additionally there are social, commercial, legal and ethical questions that will need careful consideration.

Regulation to make businesses accountable for their use of generative AI will be put in place more quickly than has been the case for other technologies. Most companies will need to adopt practices that are already the norm in financial services and the pharmaceutical sector. In relation to their use of generative AI, businesses should adopt the mantra of “do no harm”.

We can see significant interest in generative AI from all our enterprise customers. I have been surprised by the range of uses already in place. Adoption is well advanced in sectors including software engineering, content generation, advertising and research and development. This is especially the case in healthcare where generative AI has transformed the understanding of chemistry and biology. It is evident that realising the potential of AI is a prize worth chasing.

I believe the real benefits of generative AI and their effect on businesses will become more apparent in the medium term. This technology will change how work is done. It can also reimagine your enterprise, for instance where you play and how you win.

The danger will come if AI is plonked on top of old processes, old systems and biased data. Not understanding generative AI before making use of it will be dangerous. Strong principles to govern the use of AI will be essential. We have to ground this technology to reduce the possibility of harm. Of course, costs and lock-in will still be issues as the technology moves into the mainstream.

#### **Human-machine interaction**

Organisations will be liable for the outputs of generative AI just as they are for the work delivered by their employees. Navigating the risks of human-machine interaction will be important for all organisations.

At HCLTech we believe that machines need to be used to amplify human impact — and hence boundaries are needed. This is especially true for generative AI, and humans should make the final decisions on AI-suggested outputs.

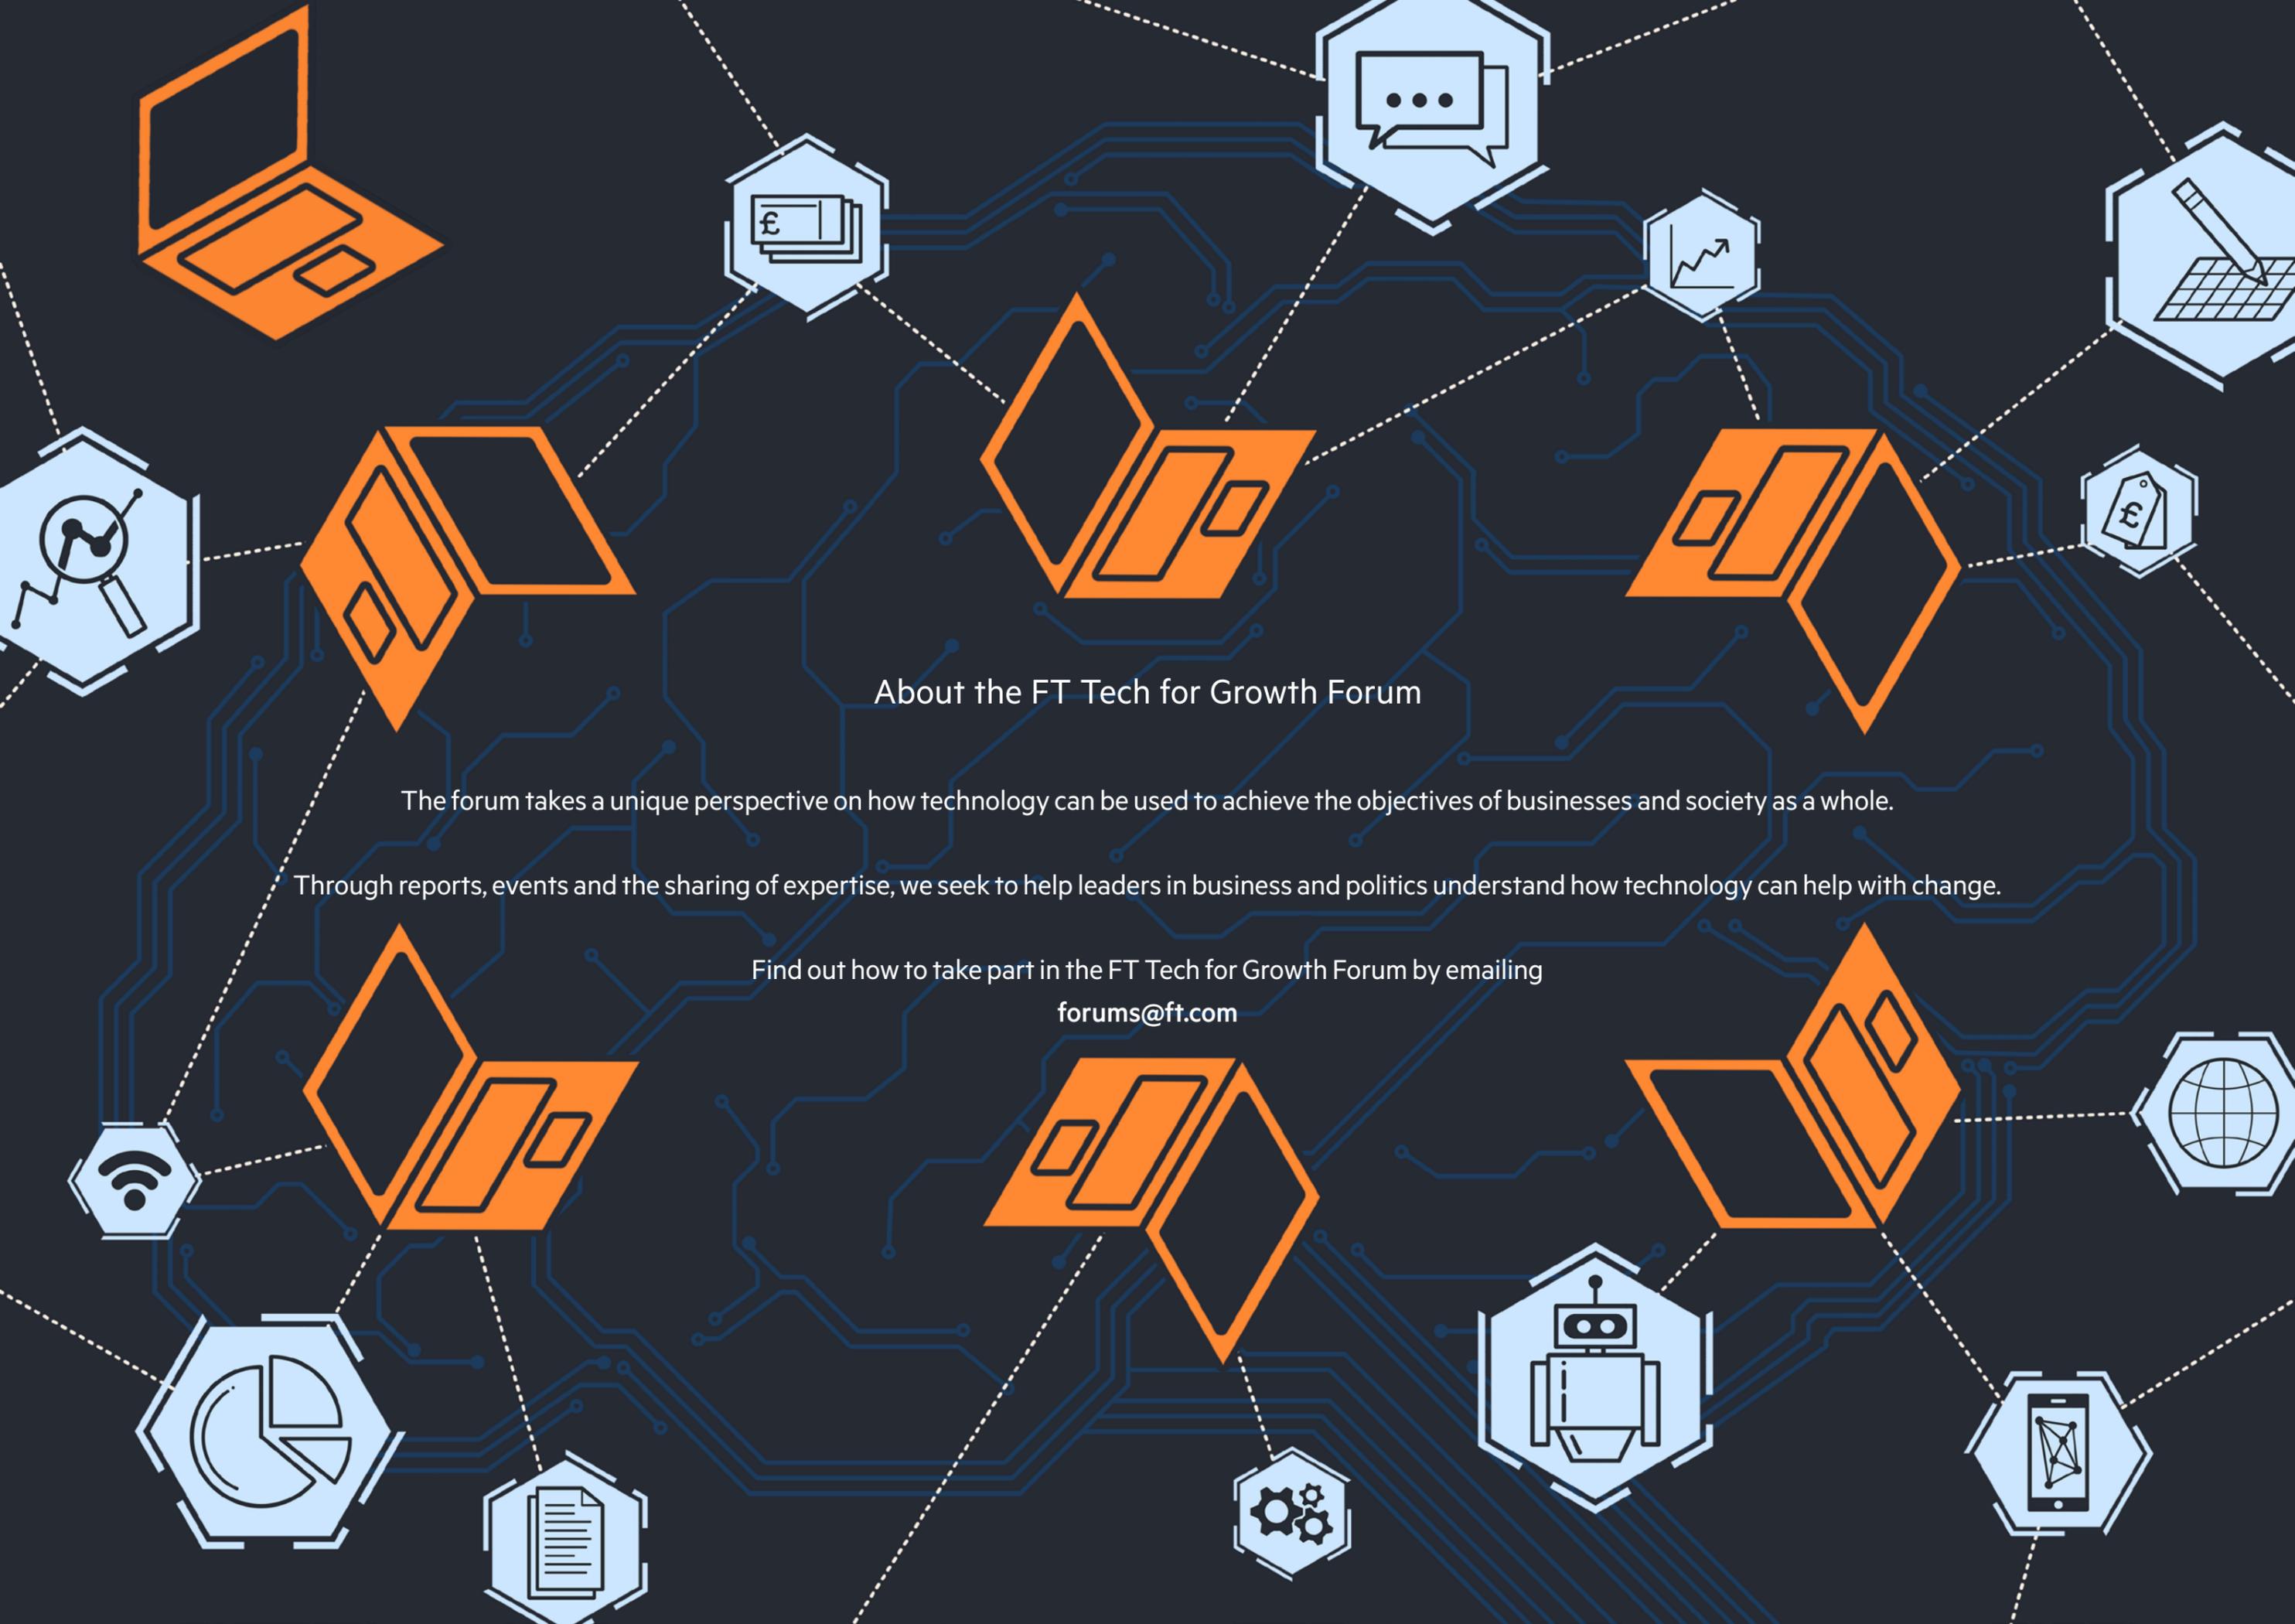
Anyone who has travelled in a car with self-driving capabilities will have experienced amazement, exhilaration and terror. You learn to trust the car to brake in time, navigate a roundabout or go across a winding hilly road. Driving in such a car convinced me that the humans who interact with generative AI machines will need robust training. They will need a thorough understanding of how the technology is used and where it can fail.

Careful management of human-machine interaction will ensure that generative AI is used safely and responsibly but the journey is likely to be long and bumpy.

In the near future the skills required to use and develop technology more generally are likely to change. The demand for quality control and human oversight of AI-generated content will grow significantly. Closing the skills gap will be essential. We should pay particular attention to these areas:

- **Bias and fairness:** We have to be aware of the biases inherent in generative AI and take steps to mitigate them
- **Explainability:** The way generative AI works can be confusing to the uninitiated. We need people who fully understand the models to provide explanations. This will help everyone to place more trust in AI outputs.
- **Safety:** AI models can generate harmful content and hallucinate. We have to be able to identify and manage this phenomenon.

Only when organisations achieve an end to end, closed-loop system — with humans and machines optimised to work with each other — will we begin to see dramatic benefits.

The background features a dark blue field with a network of glowing blue circuit lines. Scattered throughout are various white icons within hexagonal frames, including a laptop, a stack of money with a pound symbol (£), a speech bubble, a line graph with an upward arrow, a grid with a pencil, a tag with a pound symbol (£), a magnifying glass over a circuit, a Wi-Fi symbol, a pie chart, a document, a gear, a robot, a globe, and a smartphone with a network diagram.

## About the FT Tech for Growth Forum

The forum takes a unique perspective on how technology can be used to achieve the objectives of businesses and society as a whole.

Through reports, events and the sharing of expertise, we seek to help leaders in business and politics understand how technology can help with change.

Find out how to take part in the FT Tech for Growth Forum by emailing  
[forums@ft.com](mailto:forums@ft.com)