Realising the value of Generative AI

A white paper by Matt Phillips

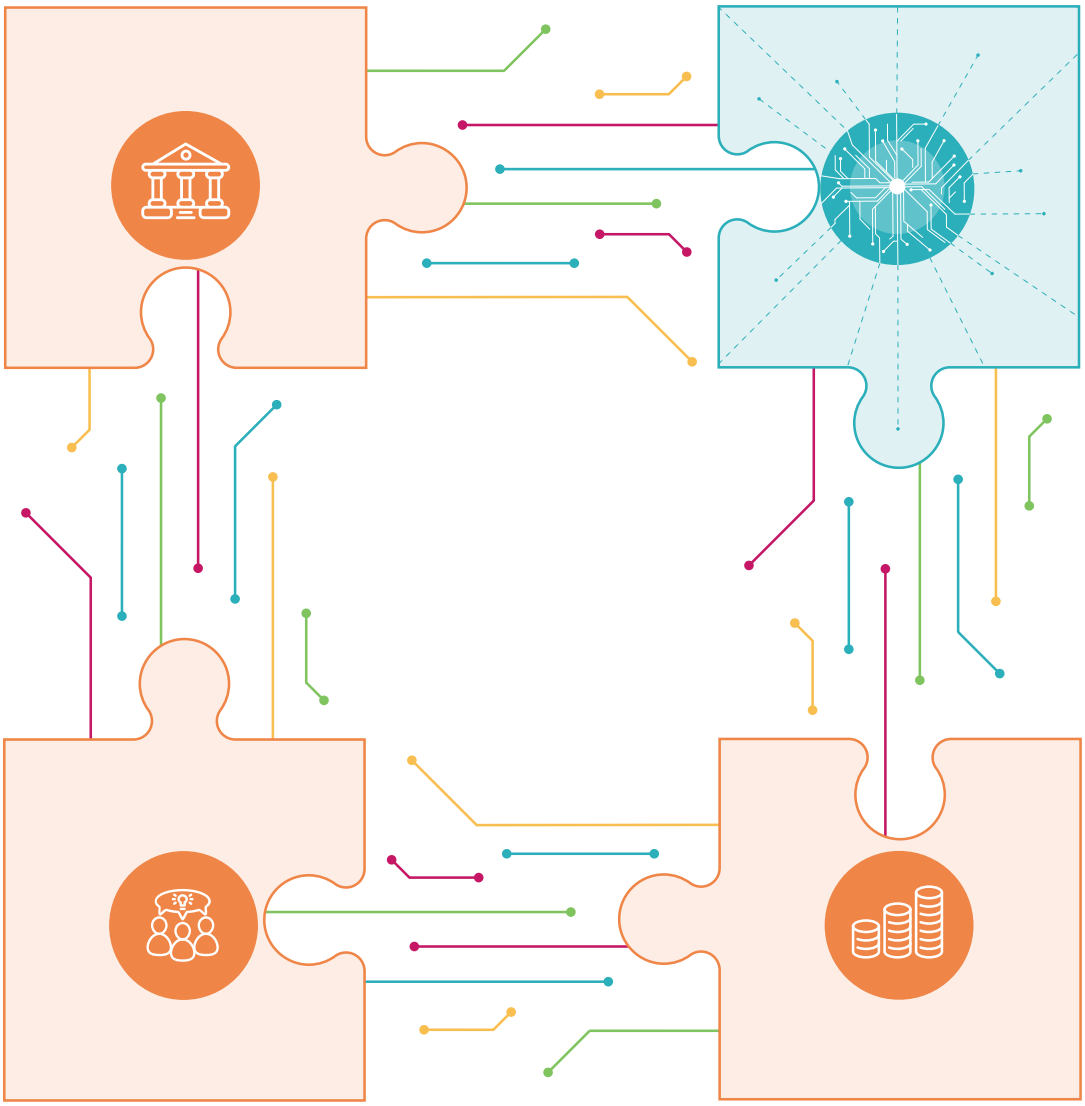

Scott Logic

Contents

Executive summary

This white paper provides an in-depth look at the rise of generative AI and its potentially transformative impact on business and society.

Some key insights include:

- Generative AI (GenAI) represents a major breakthrough, enabling AI systems to generate new content and converse in natural language. Tools like ChatGPT have sparked tremendous excitement about AI's capabilities. This hype is justified – GenAI is as consequential as the Internet or the printing press.
- Industry sectors like professional services, finance, retail and healthcare could see workflows radically changed by AI systems that generate content, automate tasks and enhance decision-making. Knowledge work will likely be disrupted first in language-intensive fields like law, media and consulting.
- While automation will displace some jobs, AI will also augment many roles and create new types of work. Managing job impacts will require policy foresight and workforce training. Risks like bias, misinformation and environment impact will need regulation.
- To realise benefits, organisations need to adopt a bifurcated approach – decentralised experimentation to uncover task-level value combined with rigorous validation for higher risk consumer/citizen-facing use cases. Data quality and model explainability are key challenges.
- Leadership must promote a culture of curiosity, learning and responsible innovation. Hands-on AI exploration and domain expertise will be critical to overcoming risks and bias. Multidisciplinary collaboration is vital.

With careful management, GenAI can catalyse enormous progress. But we must proactively address risks. This paper provides both a detailed analysis of GenAI's potential and expert guidance on realising benefits while navigating risks. It offers a perspective for leaders aiming to leverage AI's capabilities for competitive advantage and social benefit.

A primer on GenAI terminology

Here are some definitions of terms used in this white paper. A more extensive primer can be found on the Scott Logic blog.[1]

Artificial Intelligence (AI) – the development of computer systems that appear to have some form of intelligence, allowing them to perform tasks that usually require human expertise.

Neural network – a method of developing artificial intelligence by modelling the brain. Also referred to as a model.

Generative AI – a model that is capable of generating media, typically text or images, based on text-based input (a prompt).

Large Language Model (LLM) – generative AI models that work with, and on, human language, taking both text as an input, in what is termed a prompt, and generating text output.

Training – a process whereby large quantities of data are presented to a neural network, with the quality of its output evaluated in some way. Based on this evaluation, the parameter values are tweaked to produce a better output.

Foundational model – models that are trained on a broad dataset, making them suitable for a broad range of use cases.

Fine-tuning – a technique where a (foundational) model that is trained on a generic dataset has additional training applied in order to gain more specific capabilities.

1 Colin Eberhardt, Scott Logic, *A guide to Generative AI terminology*

Introduction

In the face of GenAI's rapid ascent into public consciousness, businesses are experiencing waves of excitement and anxiety. How can your organisation navigate the hype and stay on an even keel?

Prior to GenAI, it's probable that no other technology has placed a greater demand on organisations to navigate it wisely. It's impossible to ignore the claims made about its power and potential, just as it's impossible to ignore the many risks it is introducing.

Why is there so much excitement?

Whilst Artificial Intelligence (AI) as a field has existed for decades[1] and is all around us already[2], the rise of generative AI models represents a significant shift in the way we approach and utilise AI. Two significant breakthroughs underpin this. First, GenAI is capable of generating new content in the form of text, images, video, audio and code. And second, it allows humans to communicate with AI in their natural language.

In less than a year, we've seen how tools such as Stable Diffusion, Midjourney, ChatGPT and GitHub Copilot have exploded into public consciousness. Their accessibility and capability have brought the technology's enormous potential into plain sight. Indeed, ChatGPT reached 100 million monthly active users in just two months; that's seven months faster than TikTok.[3] So unquestionably, AI has reached its first true inflection point in public adoption.

Is the hype really warranted?

The tech industry is notorious for predicting once-in-a-generation technologies every year. But sometimes the hype is justified, and this is one such moment. We are entering a new technology epoch, at least as consequential as the internet or the printing press. There will be far-reaching consequences not only for business but for society at large.

That said, it's important not to be swept away by the hype. Speculation abounds at the moment about the potential of different tools, and predictions are rife about what the future will look like. Indeed, this paper makes predictions – it's a human impulse in the face of something of GenAI's uncharted nature. But just as you wouldn't want to sail around the world relying on the guesswork of medieval cartographers, you wouldn't want to make business decisions based solely on the educated guesses of consultancies. It's vital to go beyond the hype, remember that generative AI is a tool like any other, and explore how best to harness it to create value for your specific organisation.

Importantly, the opportunities and risks presented by GenAI must be understood and tackled by the whole organisation; this is not just another new technology for IT to grapple with. This paper aims to support you in navigating this new domain and inform the right strategies to seize the potential of this moment.

Because this is *just one moment* in the rapid evolution of this new technology. As impressive and transformative as GenAI appears to be, it's very likely that the epochal change is only just beginning. It's the emergent reasoning capability of artificial intelligence, coupled with its intelligent use of tools, that will make this technology transformational and mark the true dawn of a new technology epoch.

But that will be the subject of a Scott Logic white paper of the future. For now, let's return our focus to GenAI and explore its potential impacts.

1 Our World in Data, *The Rise of Artificial Intelligence Over the Last 8 Decades*
2 World Economic Forum, *10 surprising things that rely on artificial intelligence*
3 Yahoo Finance, *ChatGPT on track to surpass 100 million users faster than TikTok or Instagram: UBS*

Predicting the impacts of GenAI

As early GenAI experiments yield their results, it's possible to extrapolate some predictions about the technology's impact on jobs, sectors and professions.

How might industry sectors be impacted?

The adoption of AI has already more than doubled since 2017.[1] From here, the doubling rate of adoption will likely be much faster now that LLMs have cracked language and have become so much more accessible and easier to deploy in different settings.

Many organisations are starting to experiment by consuming foundation models off the shelf. However, the biggest value for many will come when they fine-tune the foundational models with their own proprietary data, applying the technology to the highest-value business use cases.

To give just three examples of what this might look like:

Finance

On the customer service side, businesses are showing great interest in how they can put generative AI into action around chat channels. The hope is that generative AI could replace the basic, current chatbot systems with alternatives that are more capable of responding to complex requests, learning how to deal with specific customer needs, and improving over time. Customer-facing GenAI chatbots will need to sit within a carefully developed architecture in order to be secure and safe.[2]

Professional Services

Forward-thinking firms are already putting GenAI to work on helping with pitches for new business by populating client requests for proposals (RFPs) with information from training data of past successes. The benefits of cutting the time spent on RFPs will reduce the cost of sales and therefore flow straight through to firms' profits.

Retail

Given the huge amount of data collected on consumer spending patterns, AI is being rapidly adopted to optimise the supply chain by predicting the amount of stock required. Generative AI will be used to communicate with customers through virtual assistants.

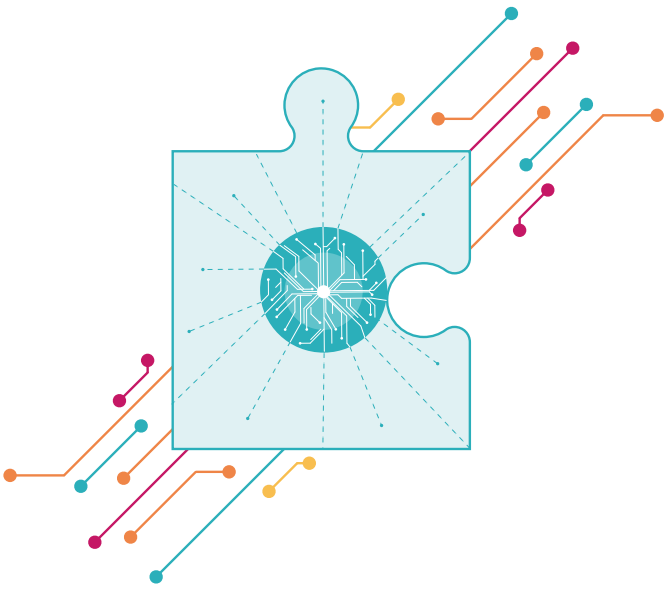

1 McKinsey & Company, *The state of AI in 2022—and a half decade in review*
2 Simon Hamilton Ritchie, Scott Logic, *How we de-risked a GenAI chatbot*

How might knowledge work be impacted?

The invention of the Large Language Model has resulted in a turning point in the capabilities of language-based AI technology. It is very likely that these capabilities will continue to advance from here, and at an even faster rate.[3]

Given that language is about the best tool humans have ever invented – the key to the human operating system one might say[4] – this technology is going to fundamentally transform knowledge work. Sectors likely to be most affected first will be those with a lot of 'language intensity' – advertising, journalism, consulting and law, for example. The value chains of knowledge work will also likely undergo significant disruption. Knowledge workers using AI copilots will probably become the norm, raising productivity across a huge variety of tasks.[5]

To give just three examples from different professions:

Market researchers

Market researchers analyse and interpret data to help businesses understand their market, their competition, and customer preferences. Generative AI could help automate part of this process by analysing large volumes of data and generating reports. This could not only expedite the research process but also uncover insights that humans might miss. However, as with other occupations, the technology's application will need to be guided by human researchers' expertise to ask the right questions, interpret results contextually, and apply insights strategically.

Medical professionals

AI models can analyse medical imaging data, such as X-rays and MRIs, often with accuracy comparable to or exceeding that of human doctors. These might be complemented by GenAI models which could aid in creating personalised treatment plans or predicting patient outcomes based on medical history. However, doctors' and nurses' empathetic care and judgement in complex or unprecedented medical situations cannot be replaced by current AI. Furthermore, the transparency of AI models will be another key consideration for the medical profession.

Lawyers

Generative AI has the potential to dramatically enhance efficiency in legal professions.[6] For example, AI systems could generate legal documents, contracts, or even court filings based on provided input and existing legal databases. Routine tasks, like reviewing contracts and other legal documents, can be accelerated with AI, freeing up lawyers' time to concentrate more on strategic counsel and complex legal matters. In this way, generative AI also has the potential to change the nature of legal work, possibly reducing the demand for entry-level lawyers while increasing the value of experienced legal strategists.

3 NVIDIA, *Fireside Chat with Ilya Sutskever and Jensen Huang*
4 Yuval Noah Harari, *AI and the future of humanity*
5 Microsoft, *Introducing Microsoft 365 Copilot—A whole new way to work*
6 Wired, *Generative AI Is Coming For the Lawyers*

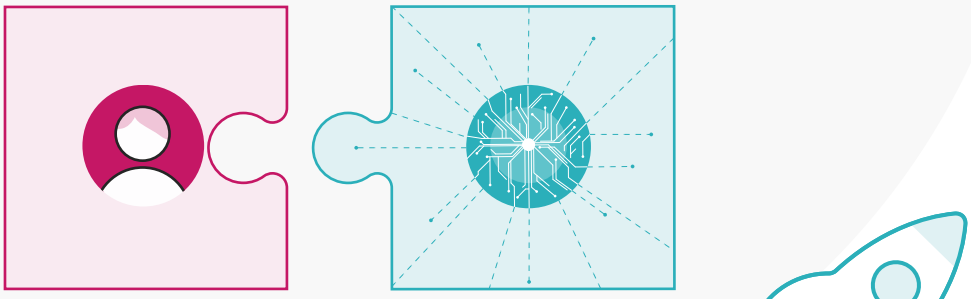

How might jobs be impacted?

A recent report by Goldman Sachs estimates that across the globe, generative AI could expose the equivalent of 300 million full-time jobs to automation.[7] But as the report also highlights, there will be job creation as well. With any new technology, especially one as transformative as generative AI, it's easier to see the displacement than the creation. And so there is a natural (and very understandable) bias towards getting worried.

Overall, there will be a mix of displacement, augmentation and creation. Reliability is paramount for enterprises but generative AI is prone to errors. As such, it's likely that there will be a lot more augmentation, in the short term at least – where AI is used as a copilot to expert humans. Computer programming jobs are a great example of this.[8,9] And so a lot of the impact in the early years will come in the form of roles using AI replacing roles not using AI.

The impact of AI on jobs is not likely to be the greatest societal risk posed by the technology – so long as good support and safety nets are provided for those people who are most directly affected. In the short term, there should be greater concern about bad actors using AI to do bad things, especially in terms of the misinformation and disinformation risks.[10]

In the long term, as AIs get more and more intelligent, the question will arise about what else people want to do with all the time available to them when there is not much knowledge work left. It's entirely conceivable that people will engage in intellectual pursuits, because they want to, not because it's their job. If humans manage to solve the control problem (avoiding the creation of Artificial General Intelligence with superhuman capabilities, surpassing our ability to predict or control it), we could have a very bright future ahead.[11]

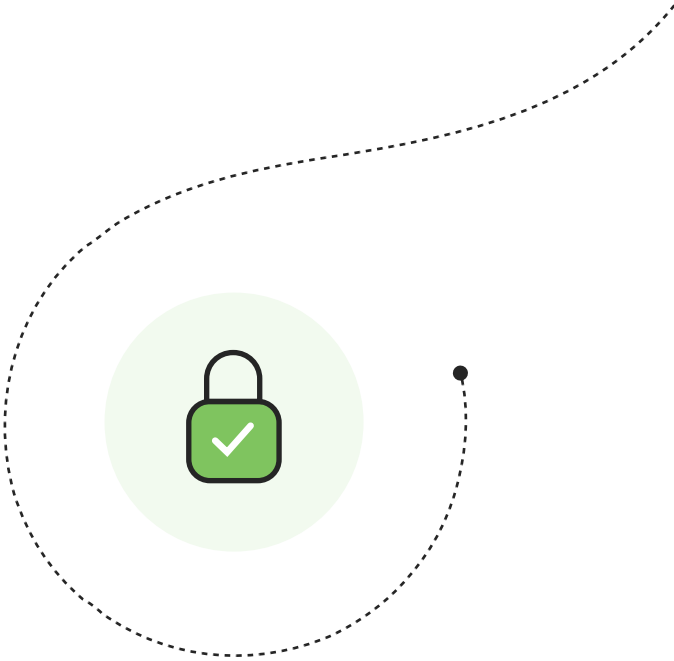

7 Joseph Briggs et al, Goldman Sachs, *The Potentially Large Effects of Artificial Intelligence on Economic Growth*
8 Scott Logic *Beyond the Hype* podcast, *Is generative AI coming for programming jobs?*
9 GitHub, *Research: quantifying GitHub Copilot's impact on developer productivity and happiness*
10 Sam Harris, *Making Sense* podcast, *AI & Information Integrity*
11 Sam Harris, *Making Sense* podcast, *The Trouble with AI*

How can your organisation realise the value of GenAI?

In order for your organisation to realise value from generative AI, you will need quite a bifurcated strategy.

On the one hand, this technology is something of a Swiss army knife. Much of the value in knowledge work will be realised at the task level, not just at the level of business use cases. This calls for a decentralised approach. If we want to discover the value at a task level in HR or Finance, then we need to enable HR or Finance colleagues to figure it out.

A decentralised approach has other advantages. First, it helps remove some of the fear that's out there. And second, current LLMs are more powerful and lower risk when they are used to augment a human with real domain expertise. Without domain knowledge, the potential for overlooking key risks increases significantly. The problem of AI hallucinations[1] is a great case in point[2] (although the term confabulation,[3] something humans do all the time, would be more accurate).

At the task level, organisations should look to develop approaches that empower human expertise in every part of the business and invest in developing the skills needed to use generative AI tools effectively. Leaders have a vital role to play in formulating talent development strategies backed with appropriate levels of investment. Empowering domain experts also means encouraging a culture of curiosity, a point that will be returned to later in this paper.

At the level of business use cases, there will be a lot of 'high risk, high reward' evaluation to be done, especially in relation to citizen- and consumer-facing journeys. Here, a lot of risk analysis and validation will be needed before deploying models into production. And in this space, there will be new and unconventional challenges. If anything, the challenges are going to be more difficult than anything seen with previous technology paradigms.

What are the business-level challenges?

To harness the business potential of AI, there are some serious challenges to overcome. Several have really caught the attention of business leaders considering how to deploy this technology into their own business operations:

Data quality

Data is the fuel that drives the AI engine. LLMs need vast amounts of curated data to learn. Organisations will naturally want to jump in and start fine-tuning LLMs with their own proprietary data. However, if they charge headlong into generative AI without investing time in cleaning the data first, they will likely find themselves in an even bigger mess than where they started. For many organisations, this is going to be the difficult and expensive bit. It's the old adage 'garbage in, garbage out', but on steroids this time.

Economics

Before enterprises follow the siren call of hyperbolic consultancies to weave generative AI into every task of every knowledge worker in the organisation, business leaders should pause to reflect on their experience of the early years of cloud. Remember how commercial cloud players initially pushed out low-cost models and then ratcheted the pricing once adoption at scale took hold? At least with Software as a Service (SaaS), the cost models were relatively straightforward a lot of the time. AI could be much less predictable. What will the cost be to large enterprises per API call for 100,000 knowledge workers?

Environment

LLMs require vast amounts of energy to train and run.[4] Transparency will be needed from the AI labs on the carbon footprint of their products, and leaders will need to demonstrate responsibility in how, where and why this technology is deployed.

Explainability

Unless explainability is intrinsic to the algorithmic approach – meaning that AI can explain how it is working at every stage – our confidence and trust will be eroded. Without trust, the technology will not be deployed at or near the point of decision in a customer journey, for example, even once the reliability is validated. Right now, even the scientists who build AI can't tell you how it works.[5] Under this heading should also be placed transparency; a black box can't be transparent, even if we want it to be.

Reliability

One of the challenges that feels very different and new is how to overcome the reliability issues presented by this technology in its current state. This includes some of AI's surprising failure modes[6] as well as the more well-known problem of bias in the models.[7] Here, we are not talking about the usual problem of software having bugs and sometimes not doing what it is supposed to do. Rather, the very technology itself is not deterministic – in other words, the same inputs will not always produce the same outputs. Of course, humans aren't deterministic either, but the question of how to validate the reliability of models remains – something that may be challenging when explainability is weak.

None of these challenges seem to have easy answers and asking the right questions seems to be the key at this stage. That said, a few pointers for business and technology leaders to consider include:

- investing in data strategy and execution to help mitigate data quality and reliability issues
- evaluating self-hosted open source models alongside the commercial models
- developing the right architectures

Some of the principal risks and proposed architectural mitigations are the subject of another Scott Logic white paper.[8]

Other areas of significant risk include intellectual property (IP) and copyright, data privacy, misinformation, disinformation, deep fakes, use of generative AI systems for cyber and fraud attacks, and careless usage of LLMs in the workplace.[9] It's noteworthy that even the commercial labs building the foundation models are asking to be regulated at this time.[10]

1 Bernard Marr & Co., *ChatGPT: What Are Hallucinations And Why Are They A Problem For AI Systems*
2 Forbes, *Lawyer Used ChatGPT In Court—And Cited Fake Cases. A Judge Is Considering Sanctions*
3 Technology Review, *Geoffrey Hinton tells us why he's now scared of the tech he helped build*
4 Technology Review, *We're getting a better idea of AI's true carbon footprint*
5 Vox, *Even the scientists who build AI can't tell you how it works*
6 Jonathan Low, The Lowdown, *In 1st Human Victory Over AI, Man Exploits System's Weaknesses To Win At Go*
7 MIT News, *Large language models are biased. Can logic help save them?*
8 Oliver Cronk, Scott Logic, *How to de-risk GenAI across your enterprise*
9 Dark Reading, *Employees Are Feeding Sensitive Biz Data to ChatGPT, Raising Security Fears*
10 Washington Post, *OpenAI CEO testifies before Congress*

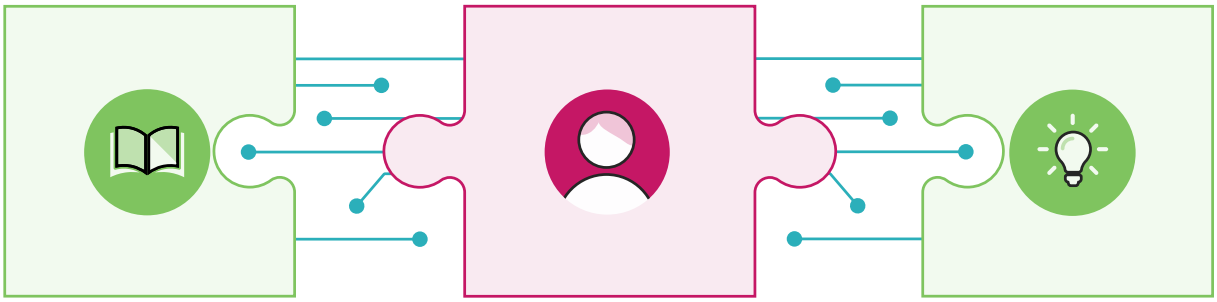

The role of leaders

To begin with a provocation, if leaders are not playing with generative AI themselves then that needs fixing first. Realising the value of this new technology cannot all be on Architects, CTOs, Data Scientists and the IT function this time.

Because this technology is not only very powerful but also very different, organisations should adopt a posture of humility, a willingness to experiment and get things wrong in a safe way. There has to be a recognition that rolling out the standard IT playbook for deploying this technology won't realise all of, or perhaps even most of the value this time around. It is likely that some organisations will look back on this period in three years' time and see that some of the initial thinking, informed by a copy and paste of previous paradigms, was not the right approach for realising the true value of this technology.

So with generative AI, it feels very important to promote a 'culture of curiosity'. If they are really honest with themselves, most organisations are pretty terrible at fostering such a culture. There tend not to be incentives or space for experimentation, play and 'tinkering' with stuff. Leaders can help by leaning into that, leading by example and removing the cultural impediments that might be getting in the way of that spirit of curiosity. Specific tactics will vary by organisation, but to give two examples:

Institutionalising Learning and Development

Leaders should encourage continual learning and skills development around generative AI technologies. This could involve setting up training programmes, facilitating workshops, and studying online resources carefully curated by thought leaders on the topic. Importantly, learning should be treated as part of the job, not as something to be done in addition to regular work. Leaders should set the expectation that learning new skills is not only permitted but actively encouraged and rewarded, with time set aside to allow this to happen.

Creating safe spaces for innovation

Leaders can promote experimentation by creating an environment where failure is seen not as a problem but as an opportunity to learn. This could be accomplished by creating a culture that rewards innovation and creativity, even when it doesn't immediately result in success. This might involve setting up R&D projects with diverse, cross-functional teams, or dedicating time for employees to pursue personal projects related to the organisation's work. These initiatives can empower employees to try out new ideas and approaches without fear of reprisal if they don't immediately succeed.

Going beyond the hype into the new epoch

The potential for AI and machine learning to fundamentally reshape the way we work, interact, and conduct business is immense.

However, the challenges and potential pitfalls of such a rapid transformation can feel overwhelming. And that's after you've tuned out all of the hype to identify where the real business value of GenAI lies.

The rise of generative AI also signifies an acceleration in the shift towards a more data-driven economy. Organisations that effectively harness this technology will likely gain a competitive advantage, but they will also need to navigate the ethical, societal, and regulatory implications carefully and responsibly.

The fears around job losses and societal upheaval due to automation and AI have been persistent themes in the public discourse around these technologies. While it is important to recognise these potential implications, it is equally important to consider the potential for job creation and the evolution of new roles in tandem with technological progress. Historically, the advent of transformative technologies has led to both job displacement and creation. For instance, the advent of the internet displaced many traditional jobs but also created an entire industry and millions of jobs that didn't exist before.

As leaders in business, academia, and government grapple with these rapid advancements, there is a growing recognition of the need for multidisciplinary collaboration, open dialogue and responsible deployment.[1] This will help organisations to navigate the complex ethical, societal, and regulatory challenges that accompany these developments. Moreover, initiatives to improve data literacy and upskill the workforce will be crucial in ensuring that the benefits of AI and automation are equitably distributed.

While generative AI holds enormous potential, it is still a tool. The ultimate outcomes will be shaped by how we as a society decide to use it, which underscores the importance of informed decision-making and leadership in these critical times.

1 Center for Humane Technology, *The A.I. Dilemma*

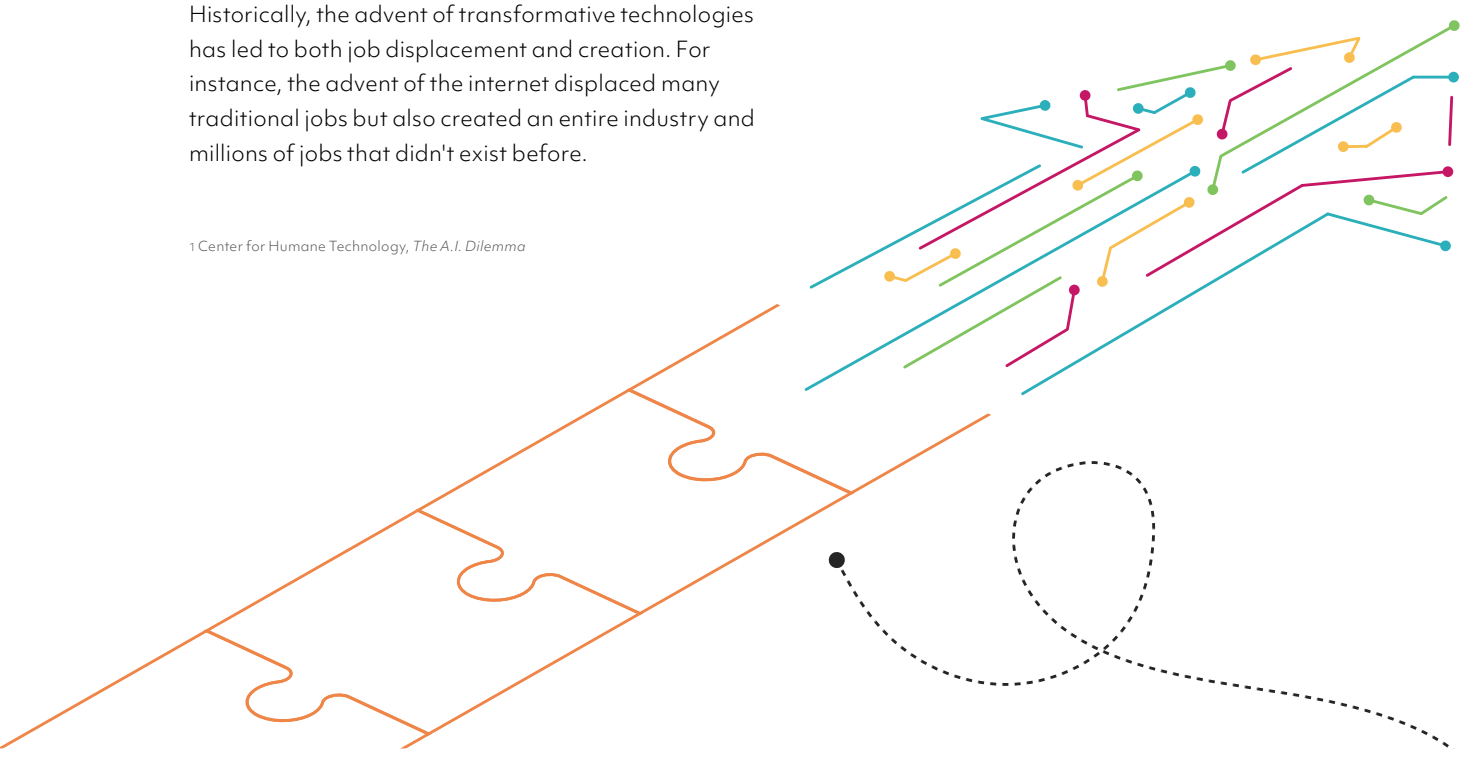

Need help navigating this new terrain?

At Scott Logic, our expert consultants have decades of experience in helping organisations like yours to seize the opportunities presented by the latest technologies, while carefully managing and mitigating the risks.

If you are looking for a friendly, pragmatic partner to assist you in navigating this complex and demanding terrain, please get in touch with our CTO, Colin Eberhardt.

Please contact Colin Eberhardt on:

+44 333 101 0020

colin@scottlogic.com

Scott Logic

6th Floor, The Lumen
St James Boulevard
Newcastle Helix
Newcastle upon Tyne
NE4 5BZ

+44 333 101 0020

scottlogic.com