



*The power within*

---

# AI and the future of internal communication


*A white paper*

---



# CONTENTS

Introduction .....	3
About this white paper .....	5
AI - a (very) short history .....	7
Types of AI .....	9
Why is this topic important now? .....	13
What are sensible use cases?.....	15
Where should we apply caution?.....	16
(Some) ethical considerations.....	19
Upskilling requirements.....	23
Conclusions and next steps .....	25



# INTRODUCTION

## In March 2023, a meme went viral overnight.

Within hours of first appearing, images of the Pope in a puffer jacket had spread far and wide. So authentic were the pictures that many, including celebrities, took to social channels to question their own ability to distinguish fake from reality.

Welcome to a new era of digital.


For those with insider knowledge of the tech industry, the arrival of ChatGPT and other generative AI platforms comes as no surprise. For the rest of us, making sense of everything Silicon Valley is throwing at us increasingly feels like trying to take a sip from a fire hydrant.

Across the digital domain, we're assured generative AI isn't going to take our jobs. In the same breath, however, we're warned to get on board with digital upskilling before someone who knows how to use AI replaces us.

The pace at which this field is unfolding is startling. It seems as if developments in AI have fast become a 21st century international arms race. No one country appears willing to slow the pace of progress, lest another proves superior.

Finding time to pause, reflect and consider what AI means for work and society at large has fast become both an imperative and a challenge.

But where do we start?




## INTRODUCTION (CONTINUED)

How do we leverage the myriad benefits of these next-generation technologies while upholding professional integrity? When the media is full of dystopian narratives about a future without work, what are the moral and ethical considerations to explore?

In any other industry, including healthcare, food and pharmaceuticals, caution and trials apply before any products are deemed safe for public consumption. Not so with technology. The wheels are already firmly set in motion.

While there's much to look forward to, there's also a lot to consider. There's no doubt AI (and, very specifically in the near term, generative AI) has the capacity to help internal comms provide a far more customised employee



experience and improve the working lives of colleagues and organisations. Used well, AI can also create efficiencies in internal communication, allowing professionals the space to focus on the vital human connection that underpins all thriving workplaces. Nonetheless, there are still many

unanswered broader, yet related, questions that require deeper contemplation.

So, what can we do to make sure AI integration into our work and lives is healthy and takes place with minimal downside for humankind?

“ AI can create efficiencies that allow internal communication professionals the space to focus on the vital human connection that underpins all thriving workplaces. Nonetheless, there are unanswered questions that require deeper contemplation. ”

# ABOUT THIS WHITE PAPER

**Here at the IoIC, we've been aware of the looming impact of emergent technologies on employee experience since we prioritised the future of work as an area of ongoing study and analysis. If you're to do your best work as internal communicators, it's incumbent on us as a membership body to keep abreast of the factors influencing employee engagement and performance.**

Since starting our mission to better understand the future of work, we've learned more about the potential for new technologies to disrupt and transform jobs, organisational design and, indeed, entire sectors and

industries. Some of our IoIC Festival keynotes have addressed the likely impact of AI on internal communication.

As a membership community, it's essential we help you stay informed and educated.

In spring 2023, we recognised the urgent need to unravel the impact of AI. We listened as members told us they wanted to better understand the risks and opportunities so they can deliver enhanced strategic value to their C-suite colleagues.

In June, we commissioned a specialist partner to facilitate an inaugural round-table discussion. Its purpose was to

discuss what AI means for internal communication right now and to better understand current thinking around viable approaches. We also wanted to surface the concerns people and organisations have about these new technologies. Facilitated by [development agency MKAI](#), we invited a select group of those operating at the highest levels of internal comms, across a range of sectors and industries, to share their hopes and fears in the age of AI.

This white paper attempts to address those perspectives.




## ABOUT THIS WHITE PAPER (CONTINUED)

In such a fast-moving landscape, it's not yet possible to present absolute answers to the big questions surrounding AI. On the other hand, it's entirely feasible to set out the questions we currently have in order to collectively address them over the coming weeks and months. Let's bear in mind that many more questions will arise as we progress, as long as we all remain open to learning and seeing our operating contexts as objectively as we can.

To this end, we present this white paper as a living document – something we can continuously edit and update as events unfold.

For us, this will be a key feature of the future of work. Rather than present evidence-based practice after the fact, our preferred *modus operandi* is to present “live” analysis, spotlighting risks and opportunities and working with members “out-loud” to share key learnings as they emerge.

So, with the above in mind, let's dive in.



## WHAT IS AI?

First, let's start with some definitions. In an online masterclass hosted by [The Guardian](#) in July 2023, [University of Oxford Computer Science Professor Michael Wooldridge](#) helped demystify some of the vocabulary.

He started his lecture by saying that no one owns the term “artificial intelligence”, which will explain why

there's no one commonly agreed definition of the term.

Given how many different technical terms are currently being used, many of them interchangeably, we think it's helpful to set out some key terms of reference. Our goal here is to simplify so that we know what we're dealing with.

# AI – A (VERY) SHORT HISTORY

At its most simple, artificial intelligence can be described as the intelligence generated by computers. Established as a formal academic discipline in the 1950s, computer scientists have worked relentlessly to create software programming rules (known as algorithms) that enable machines to learn, reason, generalise and infer meaning.


A basic hypothesis of AI is that computing power can recreate the human brain's capacity for intelligence and ultimately exceed it. [Wikipedia states](#), "[The field was founded on the assumption that human intelligence 'can be so precisely described that a machine can be made to simulate it'](#)".

The origins of AI research clearly predate more recent developments in

the fields of neuroscience, psychology and psychiatry. It's useful to bear in mind that at the time of writing, scientists still don't fully understand how the human brain works nor how we reason and create meaning.

Research into AI until and during the 1980s worked on the assumption that intelligence was a question of knowledge. Known as symbolic AI, this line of enquiry assumed AI was simply a matter of capturing and mapping the relevant knowledge held by humans. Its hypothesis was that intelligence = knowledge + reasoning.

While some early experiments proved successful, ultimately it was deemed problematic, time-consuming and unwieldy to codify how human beings acquire and apply knowledge.



## AI – A (VERY) SHORT HISTORY (CONTINUED)

The start of the 21st century saw renewed interest in the potential of AI. This has been enabled by significant advances in computational processing power, the falling cost of computer hardware, the rise of cloud computing and availability of big data. Given the excitement surrounding the potential of AI, it's unsurprising that this area of research and development is attracting huge investment.

[The global AI market size was valued at \\$454.12 billion in 2022 and it's forecast to reach over \\$2.5 trillion by 2032](#), which corresponds to a compound annual growth rate of 19 per cent over the next ten years, according to Canadian market insights firm Precedence Research. In 2023, all primary economic superpowers are

racing to win pole position in AI supremacy.

With so much money involved, media interest is to be expected. It follows there will be huge speculation on “end-of-the-world” scenarios. As anyone who works in communication knows, fear stories sell.

There are lots of articles in the press about what will happen to society once artificial general intelligence is achieved – a point at which computer-generated intelligence exceeds human intelligence.

But how capable are computer scientists of developing “superhuman intelligence”?

In his July 2023 lecture, Professor Wooldridge seemed unconvinced. As he outlined, a key paradox of AI is that the



# 19%

**Expected compound annual growth of the global AI market between 2023 and 2032 – from \$454.12 billion to more than \$2.5 trillion.**

[Source: Precedence Research](#)



things a human finds easy (driving a car safely, for instance), a computer program will find hard. Conversely, something we humans can find hard, a computer program will find easy (think of the way [DeepMind's AlphaGo soon worked out how to beat professional Go player Fan Hui in 2015](#)).

Let's take a closer look now at where we are today.





# TYPES OF AI

## Narrow AI

Most recent research and development of AI has focused on narrow AI. This can be described as the AI designed to perform a specific task or set of tasks that currently demands human brain capacity and specialist knowledge. A good example here is the translation of a text from one language to another.

There are many different ways in which narrow AI is already pervasive in work and life today. Virtual assistants such as Siri or Alexa; recommendation algorithms on websites such as Amazon, Spotify, Netflix and social media platforms; customer service chatbots – these all use narrow AI to match past patterns of behaviour and make predictions about future user interactions. Speech recognition,

transcription and translation software are commercially available today because of advances in natural language processing, a specific feature of narrow AI.

AI analytics – sometimes known as augmented analytics – help identify previously hidden patterns in large sets of data and subsequently unveil fresh insights. Using bespoke software programs, large corporates have successfully leveraged the value of their “big data” for some years already.

So, whether we’re using our smartphones, surfing the web, buying products online, using vehicle navigation, passing time on social media, or consuming content on our favourite streaming services, AI is already informing the choices we make in some form or another.

“Whether we’re using our smartphones, buying products online or passing time on social media, AI is already informing the choices we make.”

## Generative AI (GenAI)

GenAI is based on large language models (see page 10) and is a variant of AI able to create written or visual content in response to human prompting.

It’s also the form of AI we and our colleagues will initially and increasingly be exposed to as part of our working lives – and it forms the principal focus of this white paper.

## TYPES OF AI (CONTINUED)

### Large language models (LLMs)

Large-language models (LLMs) are a great example of machine learning. They perform natural language processing tasks such as generating and classifying text, answering questions and translating content. LLMs are trained on huge volumes of text data, learning patterns and relationships within it, with the aim of predicting the next likely word in a sentence structure, based on its context.

A crucial point to note here is that the quality of the output generated by LLMs is directly proportional to that of the input data. If the standard of the input data is poor, its output will be questionable.

This demands levels of critical thinking and analysis on behalf of the

user to ensure the tool is being interrogated to a high degree of scrutiny, of course.

### Neural AI and machine learning (ML)

Most advances in AI in the last two decades have focused on neural networks and machine learning. Again, at its most simple, here we show a computer program what we want it to do by using training inputs.

For example, we can teach a computer program how to recognise faces. In facial recognition, the program is shown training data in the form of input and output pairs. The input data might be a photo of an alpaca and the output data we want the program to deliver is the text “alpaca”. This is called supervised learning. In this scenario, computer

scientists aren't telling a computer program how to recognise an alpaca; they're simply providing the program with lots of photos of an alpaca.

Scientists do this using a premise called neural networks. This is a form of machine learning where a software program is computationally developed to emulate how scientists believe the human brain works.

Machine learning (ML) is a broad term used to describe the ability of a computer program to learn from past data and performance to improve future output. Here, systems are designed to pattern match and learn to solve their own problems in a cycle of continuous improvement. In a business context, ML is also sometimes known as predictive analytics.

A basic assumption of ML is that what worked in the past is likely to work in the future. From here, it's easier to understand how narrow AI is already prevalent in the deployment of customer service chatbots, for example.

A primary objective of machine learning today is to classify data based on pre-existing computing models and then make predictions of future outcomes based on them. For machine learning to work well, it requires lots of high-quality training data. In our alpaca example, the computer program has seen enough high-quality images of the animal to know that the next alpaca picture is, in fact, an alpaca.

# TYPES OF AI (CONTINUED)

## ChatGPT

So now we've covered some basics, let's move on to ChatGPT. It's this product launch (also referred to as GPT-3 and/or GPT-4) that has captured mainstream media's attention and turbo-charged a hype bubble.

ChatGPT is the brand name of a product created by Silicon Valley firm OpenAI that can create content in response to input. It's by no means the only generative AI product on the market, but, by January 2023, it had become the fastest-growing consumer software application in history, gaining over 100 million users.

ChatGPT is a large language model. LLMs use deep-learning algorithms across huge data sets to understand,

summarise and generate new content.


In the case of ChatGPT, the huge data set being crawled, scraped and used as the input (or training) data is the internet. To put that into context, this immensely vast body of content includes every post ever made on social media or other public discussion website. Naturally, the internet contains as much opinion as it does fact.

“The huge data set being used as the input data is the internet, which includes every post ever made on social media or other public discussion website. Naturally, the internet contains as much opinion as it does fact.”

A key feature of ChatGPT is that it allows users to guide and refine their desired output in terms of length, format, style, level of detail and type of language used. Successive prompts and replies are taken into account at each stage of the conversation as context. A key point to note is that the output quality of ChatGPT-generated content is directly proportional to the standard

of both its input data and demand prompts.

As Professor Wooldridge pointed out in his masterclass, while ChatGPT and other LLMs deliver seemingly sophisticated and fluent output text, at present they get things wrong a lot – and in very convincing ways.



## TYPES OF AI (CONTINUED)

### AI-generated audiovisual (AV) content

AI image generation relies on the assembly and analysis of millions of tagged images, i.e. tags that describe what that AV content is. These images and descriptions are then processed through neural networks that learn to associate deeply nuanced qualities of the image – shapes, colours, compositions – with certain words and phrases. These qualities are then layered on top of one another to produce new arrangements of shape, colour and composition, based on the billions of differently weighted associations produced by a simple prompt.

AI-generated videos are created using algorithms that analyse various data inputs like text, images and audio and then produce video output based on those inputs.

GenAI platforms such as DALL·E can make video, audio and image-based content when fed an input. So, if you want an image like the Pope rocking a puffer jacket, with the right prompt engineering, that's what it will output.

Meanwhile, Adobe has introduced Firefly. It has integrated this new generative AI platform, along with other AI-powered tools, into Photoshop. It allows users to combine the versatility of one of Adobe's star pieces of software – which, of course, has traditionally required a degree of expertise to create quality content – with good prompt engineering, to create images on the fly.

So, we can expect increasing exposure to all manner of doctored and AI-generated AV content – like [the eerily realistic Morgan Freeman deepfake clip](#)

“

If you want an image of the Pope rocking a puffer jacket, with the right prompt engineering, that's what image-based AI platforms will output.

”

doing the rounds earlier this year.

People are even creating AI-generated movie trailers. While for now they may remain less than convincing, who can say whether one day they might be as good as the real thing? As these platforms get more sophisticated, we will all need to apply increasing vigilance to the content we consume – because it might not always seem as artificial as it does today.

**AI research and development is so fast-paced and well-funded, it's inevitable we will need to familiarise ourselves with new acronyms and jargon before too long. We'll continue to provide definitions and interpretations as things progress.**

# WHY IS THIS TOPIC IMPORTANT NOW?

Since the start of the 21st century, the cost of computer hardware and data storage has dropped substantially. In parallel, computational processing power has significantly increased. The widespread availability of high-speed broadband means relatively few of us live without at least some exposure to digital technology. It's almost impossible to conceive of a business that doesn't use digital in some small way.

Emergent technology is set

to transform life and work in ways most of us have yet to grasp. The World Economic Forum refers to the onslaught of breakthroughs in AI, robotics, the Internet of Things (IoT), 3D printing, autonomous vehicles, nanotechnology, biotechnology and clean energy capture and storage as the Fourth Industrial Revolution. It predicts an entire overhaul of the way we work and live as a direct result of next-generation technological advances.

## KEY CONSIDERATIONS

In 2023, a primary point of challenge and opportunity for most businesses is the potential of LLMs to improve efficiency and accelerate organisational performance. This said, the questions arising from the commercial release of generative AI can be applied to other new technology adoption.

- What are the risks?
- What are the opportunities?
- Can a business rely on commercially available open-source LLMs or, for security and privacy reasons, should it build its own proprietary version? Or should it do both?
- What are sensible use cases for the application of these technologies?
- What are the privacy issues?
- Who should be involved in the decision-making around the adoption of these new technologies?

- What are appropriate guardrails?
- Who is to blame for errors or misinformation resulting from AI-produced content: the software or the user?
- Which jobs and departments will be most impacted?
- What upskilling and reskilling burden do these technologies impose?
- What are the ethical considerations?

Those developing and releasing these technologies have yet to provide answers for the rest of us.

Anyone whose professional work entails writing or creating quality content of any description will inevitably want answers to these questions as there's a huge amount at stake. This, of course, includes the internal communication community.

## WHY IS THIS TOPIC IMPORTANT NOW? (CONTINUED)

If we can, in the first instance, identify the questions we should be asking as we transition into a new phase of human development, we will become better equipped to ensure all future technology adoption takes place in a way that benefits, rather than hinders, society.

In all likelihood, a first swathe of questions will give rise to even more of them. And that's OK. Even if we don't have answers to all the questions that arise in the short term, the act of thinking critically about how we move forward in as inclusive and equitable a way as possible can only be a good thing.

For internal communication, this is a huge opportunity. Aside from the reality that internal communication

itself will be significantly impacted by the adoption of generative AI, there's more to consider. Organisations of all sizes and denominations must undertake informed and sensible debate about how best to move forward. With its cross-function knowledge and company-wide relationships, internal

communication is uniquely equipped to convene and facilitate these conversations, ensuring employee voice is heard at all stages.

It's important to remember people have already been using AI in their everyday lives for some years now. They are probably already using AI

to deliver some components of their work. Used sensibly, it's inevitable these technologies will help people at work to improve their own written and oral communication skills. It's also inevitable business leaders will look to leverage various efficiency gains, not least against the backdrop of ongoing skills shortages in the labour market.

We believe internal communication is perfectly positioned to make sure all impacted stakeholders are represented and to help shape the essential discussions that must take place in pursuit of optimal outcomes.

“

**Organisations of all sizes must undertake informed and sensible debate about how best to move forward. With its cross-function knowledge and company-wide relationships, internal communication is uniquely equipped to convene and facilitate these conversations.**

”

# WHAT ARE SENSIBLE USE CASES?

There are, of course, many use cases in which the application of generative AI is appropriate. At this early stage, however, for many organisations, it can feel tricky to identify what those might be, as we're in such new territory.

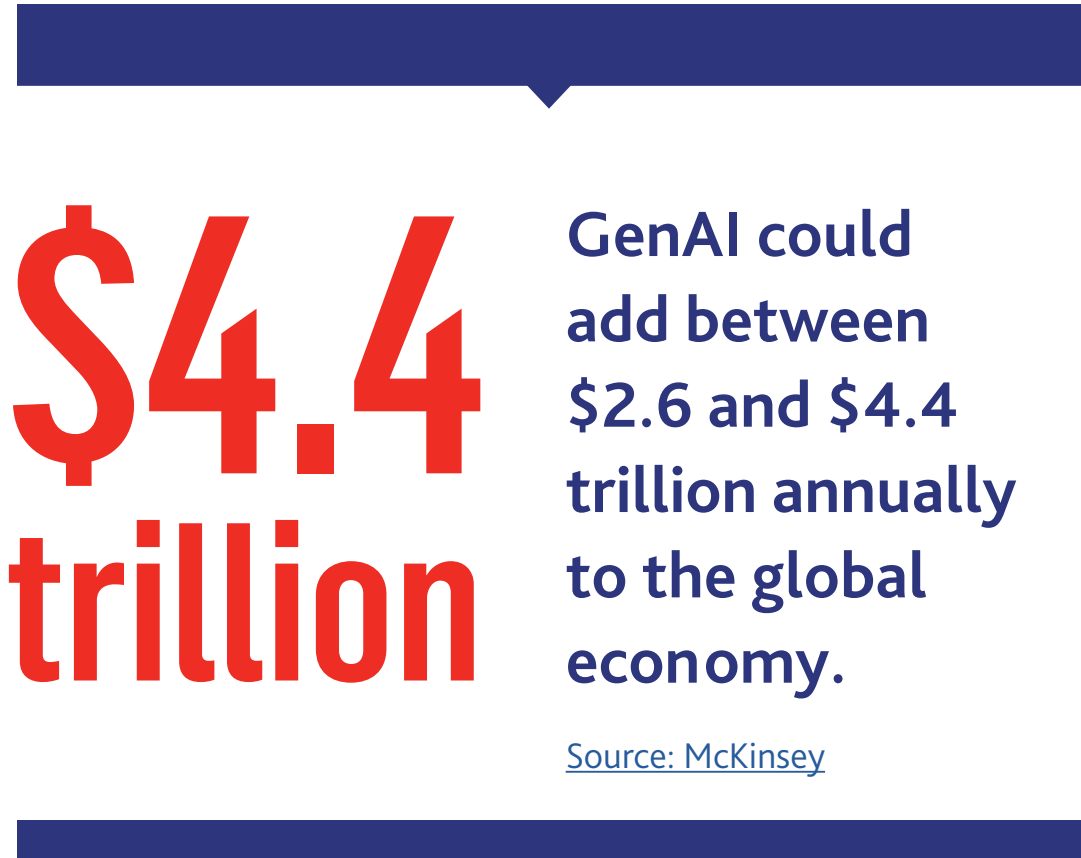
Research undertaken by McKinsey estimates [generative AI could add between \\$2.6 and \\$4.4 trillion annually to the global economy, based on 63 use cases](#). About 75 per cent of the value created falls into customer operations, sales and marketing, software engineering and R&D.

Use cases at work are important, nonetheless. Without this rigour, the risk is that employees and other contributors will take matters into their own hands – using LLMs and other new

technologies in an unstructured way that, without guardrails, risks breaching organisational policies.

While evidence of practice may be scant in the short term, there are many instances where generative AI can be helpfully applied. One example is using generative AI to improve customer experience. This extends from marketing touch points through the sales and engagement cycle to the way in which customer service is handled within an organisation.

Other areas where generative AI will inevitably add value are within the fields of HR and internal operations. There will no doubt be many as yet undocumented insights that exist within proprietary data pertaining



to people, performance and past organisational trends and trajectories. AI can help analyse and unveil these patterns. In the same way it will improve customer experience, it's expected that generative AI will deliver more customised information

to enhance employee experience. In volatile and ambiguous operating landscapes, the more detail an organisation can leverage, the more agile its market will be.

The opportunity here for internal communication functions is to convene conversations that ensure as many voices as possible are represented in the necessary planning and implementation phase. Decision-making surrounding the adoption of any type of AI should by no means be restricted to IT or senior leadership teams. Each department should, for optimal gains, hold inclusive brainstorming and planning sessions to agree the best way to integrate AI. Internal communication is ideally positioned to facilitate this.

# WHERE SHOULD WE APPLY CAUTION?

As with the introduction of any new technology into a work setting, there are numerous considerations to bear in mind.

The challenge with generative AI in particular is that it has been released commercially with no obvious guidance notes for usage. It's also the case that its commercial success depends on mainstream adoption and so, in this regard, time is of the essence.

Regardless of perceived market pressure to integrate these new AI tools into our workstreams, taking time now to think through the longer-term implications is helpful and will almost certainly reduce the likelihood of errors down the line.

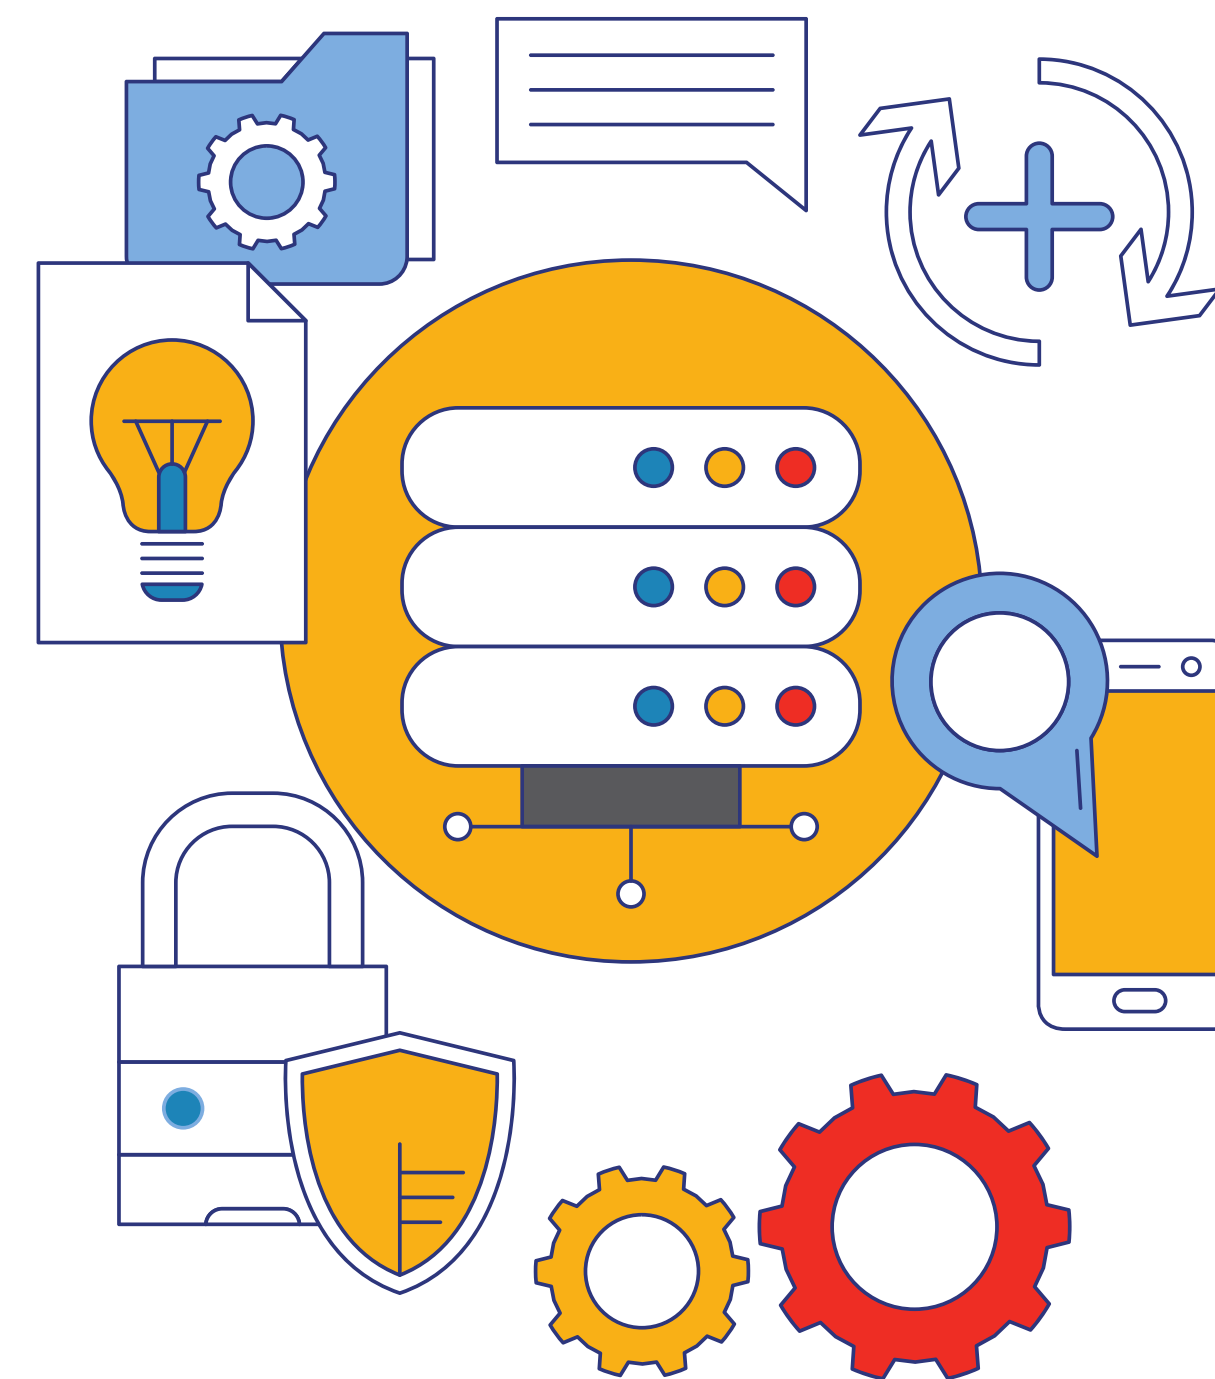
There will inevitably be many more considerations to the healthy and ethical

adoption of AI than can be listed here, as this is a topic that is constantly evolving at a fast pace. Be that as it may, there are some essential starting points to bear in mind.

## **Confidentiality, integrity, availability**

The first and most obvious consideration is that the introduction of these new technologies should not and must not threaten the confidentiality, integrity or availability of an organisation's information systems and services.

As such, all use of next-generation digital technology in the workplace must comply with both pre-existing organisational IT and data policies and, more importantly, any industry or regulatory standards and protocols.



For internal communication, a key role here is to ensure these requirements are efficiently communicated by the technology function across the entire organisation.



**Taking time to think through the longer-term implications will almost certainly reduce the likelihood of errors down the line.**





## WHERE SHOULD WE APPLY CAUTION? (CONTINUED)

### Test and measure

Despite the media-fuelled urgency regarding the rapid adoption of AI at work, we recommend taking a more considered, experimental approach.

Running small-scale experiments to test the efficacy of AI in assisting specific work deliverables allows rigour to be applied.

Testing, measuring and recording both methodology and outcomes allows business users a greater degree of discernment as to the most plausible and beneficial use cases.

Another role here for internal communication is to ensure such experiments are documented to a standard that is helpful to the wider organisation.

### Quality control

It's also crucial to understand the limitations of these AI tools. Although freely available in the public domain, most are still fledgling in their capability and should be handled as such. To all intents and purposes, we are all now beta testers.

We must not lose sight of where the majority of these LLMs have gathered their back-end data. Consider the overall quality of data that exists today across the entire World Wide Web and ask yourself how much of this information is opinion as opposed to fact. AI has no understanding of the tasks it performs and is unable to distinguish fact from fiction.

ChatGPT has, for example, already

“

Testing, measuring and recording methodology and outcomes allows a greater degree of discernment as to the most beneficial use cases. IC can ensure such experiments are documented to a standard that is helpful to the wider organisation.

”

been reported as having fabricated academic sources in the responses it has formulated, presumably in an attempt to apply credibility. The quality of generative AI output is directly proportional to the quality of the prompts applied, so user vigilance is required.

The impact of SEO and paid advertising should also be taken into consideration. How have the most popular pages on the internet come to be so popular? We all know how much time and money has been invested in digital marketing in recent years. Bearing these reflections in mind, it would be wholly naïve not to authenticate any outputs generated by generative AI.

## WHERE SHOULD WE APPLY CAUTION? (CONTINUED)

### Guidelines for sensible adoption

The fact that most of these new AI tools are already freely available in the public domain poses a significant risk to business until sensible guardrails are applied. It's imperative that all staff, contractors and consultants (regardless of employment status) are educated and informed on the commercial and operational risks of the casual and unauthorised use of AI at work.

Sensitive and proprietary data must be safeguarded at all costs.

It's never been more urgent to educate internal stakeholders on the importance of information and data security. Cybersecurity already poses a significant risk to business in 2023, and all users of internal workplace systems must apply rigour to their digital practices to avoid

data leakage. It's as yet unclear how proprietary information will be used or retained by the generative AI platforms, but, given the underlying technology, we can assume AI input prompts will form part of future outputs.


As baseline good housekeeping, nothing that could be construed as proprietary or confidential information should be included in any generative AI prompts.

It's also key that all users understand the limitations of generative AI. At best, it can be used to stimulate creativity and new ideas. Anyone thinking it can, in its current iteration, singlehandedly write documents, reports, essays or other material with legitimacy is misguided.

Overreliance and lack of critical thinking are risk factors here. All outputs should

be fact-checked, verified and evaluated for plagiarism and possible copyright infringements before formal sign-off and publication.

Finally (for now), we'd encourage organisations to educate all staff on the dangers of deep fake. If it's possible to alter an audio or video file using generative AI, it can only be a matter of time before fake material appears that compromises the integrity of executives. To this, all staff should remain vigilant.



# (SOME) ETHICAL CONSIDERATIONS

Safety and security risks aside, there are many other ethical considerations to consider before mainstream uptake of generative AI. Again, while it's impossible to be exhaustive at this early stage, here are some initial thoughts.

## Bias

How much systemic bias exists within all the webpages that make up the internet today? To what extent does the content of these pages discriminate? Most of us are familiar with how toxic certain social media platforms have become. It's sadly inevitable that various elements of bias will exist within the big data upon which generative AI has been built. To counter it requires awareness, intellect, vigilance and discernment.

As much as these next-generation technologies bring efficiency, they also demand uniquely human skills of critical thinking and objectivity. How does AI discern fact from fiction, opinion from truth? Computer scientists openly acknowledge these technologies lack any capacity for reasoned discernment.



As such, we must scrutinise all outputs and apply diligence to any encounters with AI. To this end, generative AI is not a panacea, regardless of hype.

## Accessibility

Too many of us assume that everyone has equal access to digital technology

“As much as these next-generation technologies bring efficiency, they demand uniquely human skills of critical thinking.”

and an equivalent baseline level of digital competency. This is incorrect. According to UK non-profit Good Things Foundation, [10 million people in the UK currently lack the basic foundational skills to successfully navigate our increasingly online world.](#)

We already witnessed the reality of the digital divide in the first Covid lockdown of 2020. Many families were unable to access online schooling, causing anxiety and disparity for huge numbers of schoolchildren.

[A primary objective of the United Nations Sustainable Development Goals is to narrow the inequality gap.](#) To what extent does mainstream adoption of generative AI achieve this goal?

## (SOME) ETHICAL CONSIDERATIONS (CONTINUED)

### Trust, reputation, authenticity and transparency

If we are to accept synthetic media into “business as usual”, there surely has to be a point at which trust in leaders and organisations comes under scrutiny.

[Edelman’s annual Trust Barometer consistently shows a decline in trust levels](#), not least in our political cadre. We are immersed in a “post-truth” age. Recent years of Trust Barometer reporting have shown business CEOs to be the last bastion of trust, and industry is increasingly expected to step up and act as a force for good in the world.

What happens then if senior executives use generative AI to communicate with their teams? What happens to authenticity and

transparency? If trust dissipates, what does that mean for the future of organisations? In increasingly distributed and asynchronous working environments, trust in one another and the validity of a shared cause is the glue that holds us all together.

To jeopardise this fragile, intangible yet most valuable sentiment for the sake of perceived efficiency gains is surely to gamble with the entire psychology of group performance.

What happens to trust if and when internal emails (in particular from the leadership team) are perceived as having been written using ChatGPT? How organisations choose to integrate AI into day-to-day activity demands deeper consideration.

### Copyright and attribution

As we prepare this white paper, Hollywood actors and writers are on strike. While this is in part in protest against unfair fee splits, it’s also a response to concerns around the use of generative AI in film studios and streaming services. There is a real concern within the creative industry about the lack of copyright protection in the age of generative AI. To what extent do these systems observe copyright and attribution?

As much of a concern for business is the matter of intellectual property (IP). As users dash to integrate AI tools and platforms into their workstreams, have they been educated on what IP is proprietary? Do they know what can and cannot be used?



## (SOME) ETHICAL CONSIDERATIONS (CONTINUED)

### The dilution of creative content

A key offer of generative AI is that it helps create new content. It's designed to summarise pre-existing research papers, educational material, reports, studies and more, and is positioned as an enabler and educational companion.

Arguably, however, it's been developed by highly intelligent computer scientists and software developers whose natural aptitude is far more likely to lean towards the sciences than the arts. The release of these AIs inevitably raises broader philosophical questions such as:

- Beyond information transfer, why do human beings feel the urge to write?
- What emotions does writing invoke?
- What happens developmentally in the mind of the writer as they write?

- How does writing enhance learning?
- What role does the art of writing play in developing reason and making meaning?
- What gets missed when content is created artificially, rather than creatively and by human endeavour?
- What does society and humanity lose?

Writing is an art form and as such, a balancing antidote to our increasingly data-driven worlds. When will the arts regain their status in a culture that worships at the altar of science?

For anyone who writes for the passion of it, these are big questions. And to what extent have they been considered by those behind the development of these products?

### Risk to jobs, careers and livelihoods

For almost two decades, those with an interest in future and emerging trends have rung alarm bells about the possible consequences of technology on work in the 21st century. In 2013, academics Carl Frey and Michael Osborne released their [globally renowned paper entitled \*The Future of Employment\*](#). In their research, they evaluated more than 700 occupations to identify which might be impacted by automation, AI and robotics. They concluded a staggering 47 per cent of jobs were vulnerable to automation.

While this data has been challenged, there are still significant concerns about the impact of AI and automation on jobs, income, security and employment.

[The World Economic Forum now produces a regular \*Future of Jobs\* report](#) outlining the part technology plays in the shifting nature of employment. In its May 2023 report, it predicts [the loss of 14 million jobs globally in the next five years](#); of course, new technologies typically also lead to the creation of new jobs and ways of working.

Again, there are many questions to be answered here. Given the escalating importance of, and demand for, ethical business, how does widespread uptake of AI and associated job erosion reconcile with the UN's sustainable development goals? To what extent does AI adoption fuel inequality? What is the purpose of work? To what extent does commercial profit matter?

## (SOME) ETHICAL CONSIDERATIONS (CONTINUED)

### The carbon footprint of AI

Data storage facilities require an inordinate volume of natural resources. Not only does the escalating need for processing power – fuelled by AI – consume vast amounts of electricity, but storage servers must also be kept at a certain temperature to avoid overheating. This places substantial demand on both the energy grid for power capacity and water for air conditioning and cooling.

As we race to adopt these new technologies, what happens to the other primary obligation of business in the 2020s – to decarbonise?

Climate anxiety is on the rise and there's increasing pressure from staff and jobseekers to demonstrate a firm commitment to sustainable business practices. How sustainable is the mainstream adoption of AI? Various data

# 5.6 billion gallons

Amount of water used by Google in 2022 to cool its facilities – an increase of 20% from the previous year.

Source: The Byte

tell us employees increasingly expect their employers to lead on the transition to net zero. Yet, research undertaken by Gartner indicates that [by 2025, AI will consume more energy than the human workforce](#).

These dual agendas need open discussion and prompt resolution. It doesn't, at this stage, feel as if organisations can fulfil both agendas.



# UPSKILLING REQUIREMENTS

With the adoption of any new technology comes a requirement to upskill. Unfortunately, the pace at which new technology is released in the 21st century outstrips most of our ability to use it optimally.

A recent study by Salesforce revealed that while 38 per cent of UK workers are already using or plan to use generative AI in their jobs, a staggering [62 per cent say they lack the skills to use these technologies accurately and safely](#).

Nonetheless there are several obvious areas where all colleagues will need education and support.

## Security basics

Given the potential for data leakage in the age of AI, it's vital all internal stakeholders understand and abide by key information and data security measures. While personal use of social media has blurred the construct of privacy in recent years, organisational affairs are a different matter.

All users of digital technology at work should understand the importance and value of proprietary information and data security. It's likely the use of generative AI at work will need its own policy. Internal communication should work with IT and HR to ensure all staff understand security policies and abide by them.

## Prompt engineering

For generative AI to be effective, users should learn how to craft inputs that produce high-quality outputs. This takes practice and guidance. Vague and unspecific prompts will yield mediocre responses.

The art of creating good-quality generative AI prompts lies in the ability to blend good judgement and written communication skills. Prompting technique relies on sound phrasing of input sentences.

It's also unrealistic to expect generative AI to understand multistage instructions. Chunking input instructions in a way generative AI can work with is also key to effective usage.

Finally, since using generative AI productively requires solid written communication skills, it's inevitable that some colleagues will need help here.

Internal communication can help colleagues understand the value and importance of this invaluable upskilling and share tips and advice on written communication best practice for optimal prompt engineering.

“

Vague and unspecific prompts will yield mediocre responses.

”

## UPSKILLING REQUIREMENTS (CONTINUED)

### Critical thinking and fact-checking

For any uptake of generative AI to be beneficial, users need to be able to think critically.

It's never been more important to triangulate generative AI outputs and make sure any information is corroborated and validated by multiple sources. Unfortunately, we can only expect disinformation and misinformation to amplify as we move forward. Deep fake images, voice cloning, fake news articles, bots programmed to disseminate misinformation – it will become increasingly challenging to discern fact from fiction as generative AI goes mainstream.

Being able to scrutinise information objectively and dispassionately is fast becoming a meta-skill for the future of

work. It is a skillset significantly enhanced by collective dialogue and discourse. The art of conversation and ability to listen to and integrate different viewpoints strengthens critical thinking capability.


Thinking critically is arguably a humanities-based art form. The ability to make reasoned judgements based on the

synthesis of evidence gained from multiple sources is an essential accompaniment to the effective use of generative AI at work. We'd argue critical thinking also augments the capacity for ethical debate within organisations – another essential workplace skill as we all navigate uncharted economic, environmental and political territories.

“

It's never been more important to make sure information is corroborated and validated by multiple sources. Unfortunately, we can only expect disinformation and misinformation to amplify as we move forward.

”





# CONCLUSION

In 2023, the AI landscape is a rapidly moving beast. As Covid and hybrid work have shown, every organisational context is nuanced and unique; it's no longer safe to apply "one-size-fits-all" thinking to the challenge we face.

Understanding how to interact with generative AI platforms in the workplace is fast becoming a key professional skill. We all, as users, should understand the various flavours of AI – their advantages and limitations. It's no longer singularly the responsibility of the IT department to enforce the healthy and pragmatic adoption of AI. We are each responsible for both the output we create and the wider impact of that output, regardless of the tools we use in the process.

We mustn't overestimate the power

of these tools. They are still in their infancy and their ultimate success relies upon the quality of our collective input in order that they learn to produce better outputs. In this sense, for now, we should consider them as junior interns. With the right guidance, they will make life easier, but they need robust training to help them achieve the goals they've been built



**We should consider these AI tools as junior interns. With the right guidance, they will make life easier, but they need robust training to help them achieve the goals they've been built to tackle.**



to tackle. Their efficacy is contingent on their users having strong "soft" skills. Communication, objectivity, discernment and judgement are all complementary and necessary competencies to support optimal AI adoption.

In equal measure, understanding the past, present and future of AI is important for any internal

communication professional looking to add strategic value to their employer. AI will inevitably reshape the entire nature of corporate communication and the transition pathway is full of risk and opportunity. While there is much to look forward to and substantial efficiency gains to be had, there are also mission-critical considerations that mustn't be overlooked.

Diligence, rigour and vigilance are key – despite the liberal pace and nature at which the latest set of AI tools and technologies have been commercially released. They are not toys. Just as we've learned with social media, all new technologies have an underbelly. It's incumbent on each of us to use artificial intelligence in a way that's healthy and sustainable.

## CONCLUSION (CONTINUED)

As a professional community, we must learn to navigate these tools for optimal usage. We must identify how AI will help us do our current work faster and better. This will ultimately free us up to do the deeply human and relational work that is so crucial in increasingly fluid and distributed work futures. In parallel, we must use our unique skillsets, talents and perspectives to advise leaders and colleagues on the best way forward. We must leverage

our communication skills to oversee the safe, ethical and responsible adoption of AI into all workstreams.

As your membership body, IoIC's role is to advise and steer towards best practice. We will work relentlessly to identify partners to help you make the most of this new landscape. We will also make sure we track all ethical debate regarding AI adoption, advising on these tools' risks and limitations as events unfold.

## NEXT STEPS

Given what's at stake for the internal communication profession, the IoIC is committed to prioritising thought leadership in the AI space. We will engage industry leaders and those at the forefront of both the application of these new technologies and the critical thinking required for their optimal adoption.

As such, we will:

- provide regular updates on enterprise AI developments
- organise and facilitate member consultations to explore the risks and opportunities of AI and GenAI
- identify training partners to enhance the healthy and safe adoption of GenAI into internal communication workstreams

- consult with other primary stakeholders to explore the opportunities that lie ahead.

In the meantime, we encourage you to:

- develop your own AI strategy for internal communication, considering ethics, value and sensible use cases
- consult with your business leaders on the safe and sustainable adoption of GenAI
- join any cross-departmental AI taskforces already established in your workplace (or set one up if there isn't one)
- experiment safely with minimal risk to colleagues or proprietary data.



*The power within*