



Generative AI models — the risks and potential rewards in business

What the rise of ChatGPT, DALL·E 2, Bard et al. could mean for your organization.



KPMG International

April 2023



Lisa Heneghan
Global Chief Digital Officer
KPMG International



Generative AI models highlight the power of technology. They have the potential to make us more productive and can make what we do easier in some respects. However, these models come with risk implications that all organizations and individuals should be aware of. That said, we can't ignore these models. They're rapidly becoming part of our daily personal and professional lives. We need to determine how to embrace them — but safely."

The cover image and imagery throughout this paper were designed using DALL•E 2, an AI art generator that creates images based on textual descriptions.

Image prompts for the cover image were: fluid abstract, wavy column of blue and purple, splash, drops, purple background.

While DALL•E 2 generates compelling visual content, it has not been trained on KPMG brand guidelines. Also, it has neither the human expertise nor ingenuity to understand KPMG's brand positioning. As a result, these images are considered off-brand and are used for illustrative purposes only, with special permission from KPMG Global Brand.

Contents

Executive summary	04
Market overview	05
What are generative AI models?	06
How generative AI models work	06
Potential opportunities and use cases	07
Current considerations	10
What does the future hold?	13
How KPMG can help	14
Contacts	15

01

Executive summary

We believe that generative artificial intelligence (AI) models have the potential to transform businesses through automating and executing certain tasks with unprecedented speed and efficiency. This is particularly true when human expertise and ingenuity is paired with deep understanding of how to use these programs and effectively harness their capabilities.

However, it will take time and human expertise to unlock their full potential in a way that's responsible, trustworthy

and safe. If you're considering using generative AI, it's important to establish a set of internal processes and controls for everyone in your organization to follow.

In this report, we cover potential use cases and opportunities, as well as what to consider if you're thinking about using generative AI applications, such as ChatGPT, within your organization.

Here are 10 things you should know about generative AI:

- 1** The most common generative AI solutions can roughly be divided into five categories: content generators, information extractors, smart chatbots, language translators and code generators.
- 2** Generative AI models can summarize articles, draft emails and produce images and videos. Trained by humans, some generative AI models have the conversational skills to, for example, answer follow-up questions, admit mistakes, challenge incorrect assumptions and filter or reject inappropriate requests.
- 3** ChatGPT is a chatbot trained on human instructions. Its initial underlying large language model, GPT-3.5, had 175 billion parameters and was trained with more than 1 million datasets or 500 billion tokens (words or word fragments). GPT-3.5 was not connected to the internet and was trained on data from up to September 2021. GPT-4, OpenAI's new large multimodal model, evolved from its earlier large language model.
- 4** Generative AI models have use across various business functions, from IT, human resources and operations, to finance, audit, legal and marketing. Suitable applications include drafting proposals, developing and testing code, and extracting and summarizing complex information.
- 5** Generative AI takes data inputs or parameters to learn and build knowledge. Unless you explicitly restrict the application provider from doing so, that data may then be used to answer a prompt from someone else, possibly exposing an organization's proprietary information to the public. Depending on the application, you may also be signing over your copyrights. Referring to the respective terms and conditions may give you an idea of what happens with user-inputted data.
- 6** Depending what you use generative AI for and how you implement it, your activities could expose intellectual property or trade secrets and open your organization up to fraud risk. It's important to be vigilant and make sure your organization isn't using AI in a way that contravenes applicable laws (including privacy laws), client agreements or professional standards.
- 7** Copying AI-produced information or code into any deliverable or product may constitute copyright or other intellectual property infringement. This could potentially cause your organization legal and reputational harm.
- 8** We expect both open source and boutique versions of generative AI will continue to be integrated into many common applications, systems and processes, ranging from internet browsers to AI-connected technology that organizations license, such as cloud-based software and instant messaging programs.
- 9** Creating safe usage guidelines within your organization is key to helping ensure proper and effective use of generative AI applications. Your organization should also upskill its people, as the human in the loop brings unique insights and understanding that generative AI alone can't replicate.
- 10** KPMG takes a responsible approach to designing, building and deploying AI systems in a safe, trustworthy and ethical manner. This approach helps companies accelerate value for consumers, organizations and society.

02

Market overview

According to research and consulting firm Gartner, by 2025, 30 percent of outbound messages from large organizations will be synthetically generated.¹ In the [AI Risk Survey Report](#), conducted by KPMG in the US in September 2022, 85 percent of respondents expect an increase in the use of AI and predictive analytics models. Additionally, in the 2022 [KPMG in the US Technology Survey](#), half of respondents said they've seen ROI from investments in AI technology.

Generative AI models captured attention in summer 2022 when an AI-generated image won an art contest.² In November they were in the spotlight again following the launch of ChatGPT. However, it was during a January 2023 [World Economic Forum](#) session, when Microsoft Chairman and Chief Executive Officer Satya Nadella said the “golden age of AI” is underway,³ that the buzz around ChatGPT really began to intensify, generating many questions from — and conversations with — KPMG member firm clients.

Training these models requires large amounts of venture capital, human effort and computing power. OpenAI, the creator of ChatGPT, received US\$1 billion from Microsoft⁴ and another multiyear, multibillion dollar investment from the company at the start of 2023,⁵ and Google⁶ and Meta⁷ have created generative AI models of their own. Given the range of possible applications, an entire industry is being built on making generative AI models useful.

Generative AI applications can be roughly divided into five categories: content generators, information extractors, smart chatbots, language translators and code generators:

- **Content generators:** Where generative pretrained transformer tools generate content such as blog posts, emails, social media posts, images, web copy and ads.
- **Information extractors:** These applications can create short- and long-form summaries of news articles, blog posts, legal documents and more. Some companies use them to develop and analyze legal documents.
- **Smart chatbots:** Companies are increasingly using smart chatbots as consumer assistants. The chatbots interact in a conversational way and can answer follow-up questions, admit mistakes, challenge incorrect ideas and reject inappropriate requests.
- **Language translators:** Multilingual tools that can translate many languages. They have the potential to build entire website interfaces, including translation sites.
- **Code generators:** Generative AI models can convert natural text inputs into code snippets or applications. With a basic description or small program function input, these models can produce code in various programming languages, and identify and fix bugs.

¹ Gartner, 7 Technology Disruptions That Will Completely Change Sales , October 10, 2022. <https://www.gartner.com/en/articles/7-technology-disruptions-that-will-completely-change-sales>. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

² <https://www.smithsonianmag.com/smart-news/artificial-intelligence-art-wins-colorado-state-fair-180980703/>

³ <https://www.weforum.org/press/2023/01/satya-nadella-says-ai-golden-age-is-here-and-it-s-good-for-humanity>

⁴ <https://openai.com/blog/microsoft-invests-in-and-partners-with-openai>

⁵ <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/>

⁶ <https://blog.google/technology/ai/bard-google-ai-search-updates/>

⁷ <https://ai.facebook.com/>

03

What are generative AI models?

Generative AI refers to artificial intelligence that can generate content rather than simply analyze or act on existing data.

Generative AI models, such as GPT-4, are built and trained on a collated set of data. They can be generalists or specialists built on predefined data collections and are designed to produce output that helps realize certain human-directed requests. Some models can, for example, predict the next word based on previous phrases or the next image based on descriptions of images that came before.

This training enables the quick generation of original content, including text, images, video and code. With a reduced need for human resources, some companies expect to be able to produce content faster and at lower cost, giving them opportunities to create new kinds of content that were previously too expensive or time consuming. This fundamentally changes human-machine interaction and opens myriad potential use cases.

This predictive capability enables the models to analyze. For example, they can be used to identify documents that are about a topic described by inputted text.

04

How generative AI models work

Generative AI models are designed to produce content based on a clear set of inputs and rules.

The most buzzed about application of a generative AI model of late is ChatGPT, a chatbot trained on human instructions created by San Francisco research lab OpenAI.⁸ As of 14 March 2023, ChatGPT Plus subscribers were able to use GPT-4, a large multimodal model (LMM) that accepts both image and text inputs and generates text outputs.⁹ On 23 March 2023 OpenAI launched plugins for ChatGPT, including its own web browsing plugin. This means ChatGPT can now access certain third-party sources and databases.¹⁰

ChatGPT stands for **C**hat (conversation-based) **G**enerative **P**retrained **T**ransformer. It was finetuned using reinforced learning from human feedback, where a reward model representing human preference is trained to help make outputs sound more human, prevent undesirable responses and try to avoid hallucinations (making up facts).

ChatGPT was created as a large language model (LLM) and has since evolved into a large multimodal generative AI application. This means the application can now

accept image and text inputs, not just text inputs as was previously the case. Combined with a neural network model that uses unsupervised learning to predict outcomes, this type of generative model can determine the most likely linguistic patterns and relationships between content it's already absorbed.

The "large" refers to the amount of data the models are based on, as well as the size of the models themselves, and involves training them with a massive collection of publicly accessible electronic documents. For instance, when it was released in 2022, ChatGPT had 175 billion parameters (a value that controls the behavior of the machine learning model — the greater the parameters, the greater the model's ability to analyze).

It was initially trained with more than 1 million datasets or 500 billion tokens (words or word fragments), including from Wikipedia and The New York Times. To put this into perspective, the average human speaks 860.3 million words in their lifetime,¹¹ making this collection — or "corpus," in AI terminology — equivalent to 300 years' worth of language.

⁸ <https://www.forbes.com/sites/cindygordon/2023/02/02/chatgpt-is-the-fastest-growing-ap-in-the-history-of-web-applications/?sh=7510d916678c>

⁹ <https://openai.com/research/gpt-4>

¹⁰ ChatGPT plugins (openai.com)

¹¹ https://openlibrary.org/books/OL3502128M/The_joy_of_lex

The basic version of ChatGPT isn't connected to the internet and was trained on online material up to September 2021 meaning its knowledge isn't up to date. Newer implementations released to a small number of premium developers, such as a plugin to Bing's search engine,¹² are connected to the internet and contain more recent content.

ChatGPT is an example of artificial narrow intelligence (ANI). ANI systems are suitable for performing one type of task for which they have been trained. For example, an ANI system designed to generate images will unlikely be able to solve math problems.

According to OpenAI, GPT-4 — though still not fully reliable — is substantially more so and capable of handling nuanced instructions compared to its predecessor, GPT-3.5. Most significantly, it passed a simulated bar exam around the top 10 percent of test takers. In comparison, GPT-3.5 scored in the bottom 10 percent of the same simulated exam. OpenAI points out that its limitations are similar to earlier GPT models, hallucinating facts and making reasoning errors.¹³

Keep reading to learn about potential use cases for generative AI models.

05 Potential opportunities and use cases

ChatGPT's meteoric rise in popularity is in part because anyone can use it, even those who don't have technical backgrounds. Its rapid growth in users — 100 million as of February 2023¹⁴ — is a sign of people's eagerness to use the technology. And the more users a chatbot has, the better trained its underlying AI can get.

ChatGPT has the potential to transform businesses by automating and executing language-based tasks with unprecedented speed and efficiency. LMMs can be deployed to help with a wide range of tasks. They can be modified to summarize and classify legal documents, respond to consumer questions, assist expert advisors, and generate engineering and architectural drawings.

They can act as a starting point for human inspiration, providing ideas that can be transformed into fresh and creative thinking. This makes them suitable to help generate business reports, marketing pitches and code for software applications.

Generative AI models may also have applications across IT, audit, human resources, operations and many more business functions besides. As you explore these use cases, keep in mind that despite the many opportunities generative AI has to offer, they aren't risk free.



¹² <https://www.bing.com/new>

¹³ <https://openai.com/research/gpt-4>

¹⁴ <https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app>

In the **IT operations** space, they could be used for:



LMM-based knowledge management systems

Gathering information from various data source formats. This information can then be queried to search for specific items.



Self-serve IT support

Helping employees address IT system errors through support instructions generated by conversational AI chatbots.



Coding or testing code

Converting code from one function to another, for instance from SQL to Python, or testing code to ensure it works.

On the **audit/compliance front**, they could help with:



Automating audit review

Automating audit submission fact finding and detailed audit reviews based on query formats.



Evaluating independence requirements

Evaluating audit engagement independence requirements to help simplify the approvals process for certifying independence.

Potential use in **human resources** includes:



Candidate selection

Training generative AI models on job description and relevant skills data to help identify suitable job candidates.



Self-service applications

Deploying chatbots that can both share knowledge in a human way and resolve HR queries.

On the **operations** front they can help with:



Sustainability and ESG reporting

Contextualizing ESG data and supporting reporting operations, including creating plain-language statements that outline ESG initiatives.



Virtual event management

Coordinating event management through drafting invitations, scheduling sessions and answering attendee questions.



Simplifying business operations

From drafting emails and preparing for request for proposals to running a competitive analysis and researching to ensure market understanding.

In the **finance/logistics** space, they could assist by:



Categorizing and validating payments

Helping organizations make tax contributions publicly available by sorting through massive volumes of data.



Drafting and reviewing contract terms

Reviewing contracts and highlighting potential conflict of interest clauses and drafting clauses and terms to hasten the contracts process.

Legal and organizational governance options include:



Making personalized independence recommendations for investments

Enabling organizations to provide personalized responses to independence-related queries via chatbots.



Surfacing legal citations and source links

Searching for relevant legal citations and case examples, helping to identify reputable sources.

Potential **marketing** applications include:



Simplifying campaign language

Finding alternate word choices that translate well across a variety of languages.



Localizing marketing communications at scale

Helping localize global campaigns by sharing local conversation data with the model.



Distilling complex information

Learning the basics of, for example, financial due diligence, to digest and structure content to help build a strong marketing campaign.

At KPMG, we're recognized leaders in AI, machine learning, and data & analytics, and have deep expertise in risk considerations with respect to emerging technology. KPMG can help assess the ethics, governance and security in place around clients' AI and machine learning technologies.

KPMG firms have long been at the forefront of exploring and harnessing new technologies and can answer questions about how generative AI applications may help your organization grow.



You've now seen how generative AI models may help consumers, streamline organizational processes and free up time for employees to take on higher-value organizational tasks. That said, the use of generative AI has many limits and potential pitfalls.

As we mention in part 4, GPT-3.5, ChatGPT's initial underlying LLM, was trained on material up to September 2021 and wasn't connected to the internet. And though OpenAI has made it possible for ChatGPT to browse the internet in certain cases, it's still crucial to ensure human review and expertise is built into the generative AI use process. Generative AI models can be the core of an AI application but require additional analytics, technology and human process around them to solve problems.

In this section, we discuss the risks of using generative AI models and applications and how to manage them, including around client and company confidentiality, employee misuse and phishing.

Large multimodal models like ChatGPT generate human-like responses. However, they lack human-like reasoning skills. For them to be considered trusted, users are responsible for applying their AI capabilities to suitable use cases, and your organization should educate employees on using such programs. Equally important, developers should use reliable data sets to train the AI models and apply relevant bias and content filters.

Risk management

Generative AI's growing popularity is another reason we recommend developing and deploying AI in a responsible way if your organization wants to protect itself against misuse. The following are some of the risk management challenges of generative AI models.

Internal risks and considerations

Breaking confidentiality and intellectual property

Many generative AI models are built to absorb user-inputted data to improve the underlying models over time, in essence helping them learn and build knowledge. That data, in turn, could be used to answer a prompt from someone else, possibly exposing private or proprietary information to the public. The more your business uses this technology, the more likely it is others could access your sensitive or confidential information. Thus, your organization needs to figure out how

to protect its intellectual property while still being able to enjoy the benefits of generative AI applications.

Employee misuse and inaccuracies

Even legitimate use of generative AI comes with risk. The models generate responses based on input received, meaning there's a risk they may provide false or malicious content. As your employees use it, they need to be cautious and review AI-generated content with a critical eye and emphasis on quality assurance.

If generative AI content contains inaccuracies that are not caught, this could impact your business' outcomes or create liability issues. For instance, Meta's generative AI bot Galactica was created to condense scientific information to help academics and researchers quickly find papers and studies. Instead, it produced vast amounts of misinformation that incorrectly cited reputable scientists.¹⁵ Another Meta bot, BlenderBot3, was caught making false and biased claims¹⁶ shortly after its release in August 2022. As well, Google's chatbot Bard caused parent company Alphabet to lose US\$100 billion in market value after it shared incorrect information during its first demo.¹⁷ ChatGPT hallucinating facts is also well documented,^{18, 19} with developer OpenAI acknowledging its ongoing shortcomings.²⁰

Other risks around generative AI include the possibility that the technology could generate sensitive information, such as personal data, that could be used for identity theft or to invade one's privacy. Even a disgruntled employee or angry customer could create fake material to harm your company's reputation or that of one of your employees or executives.

Generative AI evolves

As the world's understanding of AI continues to evolve, we are already seeing a rising number of global regulations. It's important to stay abreast of these, even if you don't plan on using generative AI intentionally.

We expect that generative AI will continue to be integrated into many common applications, systems and processes, ranging from internet browsers to AI-connected technology that your organization may license. Thus, it's key to be vigilant and make sure you don't use AI professionally in a way that contravenes applicable laws (including privacy laws), client agreements or professional standards.

¹⁵ <https://gizmodo.com/meta-ai-bot-galactica-1849813665>

¹⁶ <https://www.cnn.com/2022/08/11/tech/meta-chatbot-blenderbot/index.html>

¹⁷ <https://www.npr.org/2023/02/09/1155650909/google-chatbot-error-bard-shares>

¹⁸ <https://www.npr.org/2023/03/17/1164383826/heres-what-the-latest-version-of-chatgpt-gets-right-and-wrong>

¹⁹ <https://news.yahoo.com/factual-errors-inflated-bios-aside-100209244.html>

²⁰ <https://openai.com/research/gpt-4>

Questions to consider:

1. How can you ensure confidentiality and accuracy are maintained while using generative AI models?
2. How can you ensure your generative AI models comply with growing global regulations?
3. How can you automate reviewing and managing compliance policies?
4. What should your workforce know about generative AI in terms of its risks and benefits?

Talent implications

High-quality, expert output can only be achieved with high-quality, expert queries. Therefore, your organization will need to upskill its workforce and retain proprietary knowledge to contextualize the query and provide the right prompts. At KPMG, for example, we've made generative AI training available to all our people through our Digital and Data Foundations program, which provides foundational content on the evolution of AI and how to build, implement and engage with trustworthy AI.

Professionals need to be made aware that they're not just using a solution — they're training and evolving it.

In a generative future, we anticipate that the role of professionals will shift from problem solving to problem defining as teams work alongside machines to create new approaches. Generative AI tools are an interface, not an oracle.

The human in the loop brings unique insights and understanding to the process that generative AI alone can't replicate. They provide critical feedback to refine and improve the model over time and ensure the output's accurate, fair and meets the desired goals.

Great things can happen when people and technology are in harmony, and we strongly believe there can be no lasting change without human ingenuity.

External risks and considerations

Misinformation, bias and discrimination

As we discuss above, LLMs and LMMs have shared false, out-of-date and discriminatory information, but presented with such authority in a way that even the most skeptical reader could be fooled.

Generative AI can — and has — been used to create deepfake images and videos (when visual content is altered to make it seem that someone said or did something they didn't do or say). These images and videos often look extremely realistic and lack forensic traces left behind in edited digital media, making them difficult for humans or even machines to detect.²¹

Copyright

Questions abound around who owns content once it's run through generative AI applications, and there's no one-size-fits-all answer. Terms and conditions vary from tool to tool, and how you use the materials also plays a part.

If content is cut and pasted or mostly unchanged from text copyrighted to someone else, this could be considered plagiarism. It's difficult to say definitively how much information obtained via a generative AI tool would need to be changed for it to legitimately be called your own.

Claiming AI-generated content as your own could raise a host of ethical issues. For starters, acting this way isn't responsible or trustworthy and, if it came to light, would likely make clients and consumers doubt your honesty on all fronts. Further, if clients or consumers were to discover you're simply passing along AI-generated information, what's to stop them from doing the same and cutting the middleperson (your organization) out entirely?

In the next subsection we delve more into the reputational risks associated with generative AI.

Financial, brand and reputational risk

If you or someone in your organization were to copy AI-produced information or code into any deliverable or product, it may constitute copyright or other intellectual property infringement. This could potentially cause your organization legal and reputational harm.

Though many of these tools specifically tell users not to enter confidential client information, users with a lack of training and understanding of them may inadvertently risk exposing intellectual property or trade secrets to the public or even a competitor. This may lead to lawsuits and could negatively impact your company's bottom line if current or prospective clients and consumers question whether you can be trusted with their sensitive information.

Lack of transparency when using generative AI content can also create reputational issues. Tech publisher CNET was criticized for quietly using the technology to write more than 70 articles since November 2022²² — some of which contained errors — even though the publisher said on its website that a team of editors is involved in the content "from ideation to publication."

Questions to consider:

1. How can you ensure generative AI applications are managed effectively to avoid financial penalty due to not complying with regulations?
2. Can you trust the applications you use?
3. How can you proactively manage your applications and be aware of and watching for potential bias or discrimination?
4. Is using generative AI applications in line with your ethics, values and brand?

²¹ <https://www.propertycasualty360.com/2021/09/14/deepfakes-an-insurance-industry-threat/>

²² <https://gizmodo.com/cnet-artificial-intelligence-writing-scandal-1850031292>

Cybersecurity

Cybercriminals can use generative AI to create more realistic and sophisticated phishing scams or credentials to hack into systems. Further, AI algorithms can't protect their underlying training datasets. Studies have shown that algorithms can distinguish individuals' identities even if data has been anonymized and scrubbed.²³

Other generative AI cybersecurity risks include data poisoning, wherein the data that's used to train the models is manipulated, and adversarial attacks — attempting to trick generative AI models by feeding them malicious inputs.

As your organization explores use cases for ChatGPT and other generative AI applications, we recommend that your cyber and risk teams set secure implementation guidelines and regulations. These can include setting expectations for using ChatGPT and other solutions in the enterprise context, educating your people on the benefits and risks of using generative AI applications and implementing cybersecurity controls as appropriate.

Questions to consider:

1. How secure are your generative AI applications from cyberattacks, bad actors and insider threats?
2. Are your security controls working? How can they be improved?
3. Do the applications you use violate anyone's privacy?

Adversarial attacks

Even when trained to work within acceptable boundaries, generative AI models have proven to be vulnerable, like any analytical model, to deliberate manipulation by sophisticated external users. If your organization plans to use generative AI solutions, you need to be aware that this could happen to you when the solution is exposed to the public.

Questions to consider:

1. What are the basic known adversarial vulnerabilities of the technologies you're using?
2. How can you test likely attacks and harden existing and new solutions to be prepared for them?
3. What monitoring do you have in place to identify adversarial attacks?

Supporting appropriate generative AI use

We recommend creating safe usage guidelines for your organization, as this is key to helping ensure proper and effective use of generative AI applications. Guidelines can include requiring training for anyone who wants to use generative AI and outlining how it should and shouldn't be used. Additionally, your organization should treat generative AI like any other technology solution and require employees to follow any relevant policies (such as acceptable use or information security policies) that already exist.

We believe there's still work to do before we can use the latest generation of AI for consumer, employee, citizen and business interactions. With a responsible AI program in place, organizations can begin to move forward with developing processes and procedures around the use of generative AI.

²³ <https://techcrunch.com/2019/07/24/researchers-spotlight-the-lie-of-anonymous-data/>

What does the future hold?

Looking at what technology players are exploring in generative AI gives a sense of where the space could go in the future.

Software development and maintenance

Generative AI is showing potential to advance the entire software development process and thus enable faster delivery of more reliable software products and services. We anticipate that companies may be able to fully automate processes such as code generation, maintenance and fixing bugs.

Video creation and virtual reality

Generative AI can create immersive video game environments, design videos, or even personalize product videos for e-commerce websites. In the future, companies can leverage it for virtual assistants or livestreaming applications, such as automatically captioning live video.


Many companies in this space are now shifting their focus to enterprise clients.

The metaverse

Creating realistic 3D assets in the metaverse is expensive and time consuming. Generative AI can generate 3D assets through text to image or voice, 3D scenes based on 2D pictures, or even sound effects. It can also generate human faces and give more realistic characteristics to metaverse avatars.

Improved information security

Generative AI can teach individuals about what key risks certain vulnerabilities represent, helping them write appropriate scripts or understand methods of attack by threat actors.



08

How KPMG can help

For more than 150 years, KPMG firms have played a leading role in exploring and harnessing new technologies, such as generative AI, and providing assurance and direction in implementing them.


Global Lighthouse is KPMG's worldwide network of more than 15,000 data & analytics, AI and emerging technology specialists, with locations in 37 countries across the Americas, Asia Pacific and Europe. We understand that responsible AI is a complex business, regulatory and technical challenge. Through Global Lighthouse and the network of KPMG firms, we're committed to helping clients bring a responsible AI offering to life.

Using generative AI responsibly

Global Lighthouse helps organizations build responsible, trustworthy and safe AI solutions. Further, KPMG takes a responsible approach to assessing the ethics, governance and security in place around clients' AI and machine learning technologies. The set of frameworks, controls, processes and tools can help KPMG firm clients harness the power of AI — designing, building and deploying AI systems in a safe, trustworthy and ethical manner — so companies can accelerate value for consumers, organizations and society.

Our responsible AI approach includes:

1. **Fairness:** ensure models are equitable and free from bias.
2. **Explainability:** ensure AI can be understood, documented and open for review.
3. **Accountability:** ensure mechanisms are in place to drive responsibility across the AI lifecycle.
4. **Data integrity:** ensure data quality, governance and enrichment steps embed trust.
5. **Reliability:** ensure AI systems perform at the desired level of precision and consistency.
6. **Security:** safeguard against unauthorized access, corruption or attacks.
7. **Privacy:** ensure compliance with data privacy regulations and consumer data usage.
8. **Safety:** ensure AI doesn't negatively impact humans, property or the environment.



Contacts



Lisa Heneghan
Global Chief Digital Officer
KPMG International
lisa.heneghan@kpmg.co.uk



Paul Henninger
Partner, Head of UK Connected
Technology & Global Lighthouse
KPMG in the UK
paul.henninger@kpmg.co.uk

kpmg.com/socialmedia



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2023 Copyright owned by one or more of the KPMG International entities. KPMG International entities provide no services to clients. All rights reserved.

KPMG refers to the global organization or to one or more of the member firms of KPMG International Limited ("KPMG International"), each of which is a separate legal entity. KPMG International Limited is a private English company limited by guarantee and does not provide services to clients. For more detail about our structure please visit kpmg.com/governance.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.

Designed by Evalueserve.

Publication name: Generative AI models — the risks and potential rewards in business | Publication number: 138647-G | Publication date: April 2023