



# Generative AI and its implication on the industry

**Artificial Intelligence in Financial Services**



**March 2023**  
kpmg.ch

# Content

<b>The essence of generative AI and its sudden emergence</b>	<b>03</b>
<b>A game-changer that the financial industry needs to be ready for</b>	<b>05</b>
<b>Outlook for Financial Services</b>	<b>06</b>
<b>Limitations and dangers in the use of generated content</b>	<b>07</b>
<b>Does Artificial Intelligence think like a human?</b>	<b>08</b>
<b>Pacta sunt servanda and the correct management of intellectual property rights</b>	<b>09</b>
<b>Concluding remarks and KPMG practical guidance</b>	<b>11</b>
Bibliography	12



# The essence of generative AI and its sudden emergence

Advances in generative AI: What experts in Financial Services need to know about one of the world's leading technology trends.

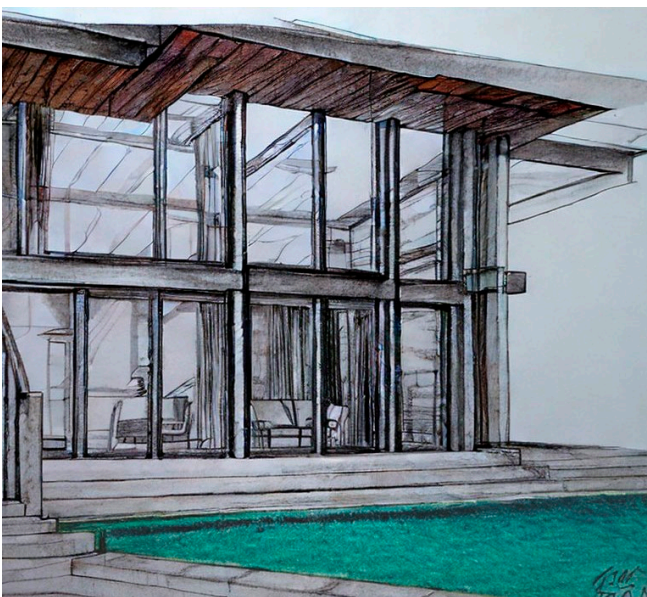
The emergence of models such as [Stable Diffusion](#), [DALL-E 2](#), and [ChatGPT](#) made generative Artificial Intelligence (AI) a burning topic among technical specialists, industries, and society overall. In recent months, social media has been humming and overflowing with articles about this surprisingly capable technology. People sharing their reflections, experiments, and discussions around the topic continue. At the same time, hundreds of new start-ups are springing up like mushrooms and rushing into the market with solutions and tools based on that very AI technology.

But what is generative AI, and why does it grasp people's minds worldwide?

As the name suggests, generative AI is the technology that aims to generate texts, images, speech, videos, or code, thereby employing machine learning algorithms. Such a system usually takes as a "call for action" a *prompt* – a natural language sentence – and executes it by creating a semantically associated image, text, or other desired output.

The technology has already found applications in several domains. For instance, text-to-image models like DALL-E 2 and Stable Diffusion rapidly change the artistic world, including gaming, movies, and even architecture disciplines. So, Bill Cusick – a creative director of [Stability AI](#) – believes that generative AI is the "foundation for the future of creativity" with the potential to become an essential part of the creative process. Nonetheless, such approaches are not meant to accomplish an artist's work but to deliver inspiring ideas and novel concepts.

Beyond the creative domain, the generative AI models demonstrate outstanding capabilities in generating natural language. OpenAI's latest release, ChatGPT, became a phenomenal trend and gained a million users in just a few days, being tested for a plethora of tasks including but not limited to novel writing, real-time translation, and interview preparation.



Two examples generated by Stable Diffusion, left – from the prompt "Architectural sketch of a house"; right – from the prompt "Fantastic creature of the AI universe"

Furthermore, generative AI actively embeds into the world of complex sciences by suggesting code and assisting developers in autocompleting their coding assignments (GitHub Copilot based on OpenAI's Codex model). By specifying the prompt in plain English, one gets returned a compliant unit of code. "The hottest new programming language is English," – twittered Andrej Karpathy, a research scientist and a founding member at OpenAI, and one could not agree more.

Prior to focusing on ChatGPT, it is necessary to understand its key concepts. Models like it are known in AI as *large language models* (LLMs) – probability distributions over sequences of words, tokens, or characters. As a rule, LLMs are trained on large textual data and, once trained, capable of generating texts similar in style and context to texts they were taught on. The extensive underlying data allows for *transfer learning*, through which the model can reason in *latent space* and solve uncorrelated tasks. Specifically, ChatGPT works by using a model architecture called *transformers* – a relatively young approach in AI field. The pivotal idea behind the transformer is *self-attention*, which allows the model to focus on the crucial parts of the input data by differentially weighting their significance. The transformer architecture was introduced by a team of Google Brain researchers (Vaswani, 2017). Since then, it has become the foundation for many state-of-the-art models in natural language processing (NLP). Additionally, the developers employed the *reinforcement learning from human feedback* (RLHF) technique in the ChatGPT training process, where for a set of training prompts, the labelers demonstrated the desired behavior of the model, and then ranked several models' outputs.

Yet, one cannot refer to generative models without explaining the concept of *prompt engineering*. As mentioned before, a *prompt* is a call to action that causes a model to generate an output. A prompt can be anything from a basic question to a complex task with raw data embedded into it (e.g., a JSON or CSV file). A prompt can be constructed by combining questions, examples, instructions, and input data (Amatriain, 2023). Thus, one could ask a question and include instructions on how the reply should look like. Alternatively, one can provide examples and request the model to base recommendations on them.

The task to properly and efficiently construct a prompt that would best match the desired output is called *prompt engineering*. It is a very recent though rapidly growing discipline. A skilled prompt engineer possesses domain knowledge to understand how an efficient prompt should look in a specific area and understand the underlying model. The latter is essential since diverse models respond differently to the same prompting (Amatriain, 2023).

AL How are computer chips made? Explain step-by-step for a five year old child?

GPT Sure, I'll try to explain it as simply as possible!

Computer chips, also called microchips, are made by taking a special material called silicon and adding tiny electronic parts to it. These electronic parts are called transistors and they can turn on and off really quickly to help the computer do its job.

Here are the steps to make a computer chip:

1. First, you start with a really thin slice of silicon, called a wafer.
2. Next, a special machine puts a layer of material on top of the wafer. This material will help make the transistors.
3. Then, a special machine uses light to create a pattern on the material. This pattern tells the machine where to add the transistors.
4. Next, the machine adds tiny transistors to the wafer using the pattern as a guide. These transistors are really small - you could fit thousands of them on the head of a pin!
5. After the transistors are added, the machine puts another layer of material on top to protect them.
6. Finally, the wafer is cut into tiny pieces, and each piece becomes a computer chip that can be used in all sorts of devices, like phones and computers.

And that's how computer chips are made!

AL Here are some examples of authors I really like: Arkady and Boris Strugatsky, George Orwell, Gabriel García Márquez. What other authors do you think I might like?

GPT Based on the authors you've mentioned, here are some other authors you might enjoy:

1. Haruki Murakami - Like the Strugatsky brothers, Murakami is known for his imaginative and surreal storytelling. He often blends elements of fantasy and reality to create unique and thought-provoking stories.
2. Aldous Huxley - Huxley, like Orwell, was a writer of dystopian fiction. His most famous work, "Brave New World," depicts a future society where citizens are controlled by technology and pleasure.
3. Isabel Allende - Allende, like García Márquez, is a Latin American author known for her magical realism. Her stories often mix elements of fantasy and myth with real-world settings and themes.
4. Kurt Vonnegut - Vonnegut's writing shares a similar satirical and social commentary style as Orwell's works. He often explores themes of human nature, war, and the meaning of life in his stories.
5. Stanislaw Lem - Lem, like the Strugatsky brothers, was a science fiction writer who explored themes of technology, society, and humanity. His works often include philosophical and intellectual elements, along with a sense of humor.

**Two examples** generated by ChatGPT Plus, left – combination of question and instructions, right – combination of question and examples.

# A game-changer that the financial industry needs to be ready for


Back in 1986, K. Eric Drexler wrote that *“artificial intelligence will be the ultimate tool because it will help us build all possible tools”* (Drexler, 1986). Indeed, a few decades later, humanity arrived at the point where AI can be included in many applications, revolutionizing the way businesses operate. Microsoft has already announced the intention to [extend the partnership](#) with Open AI (the company behind ChatGPT) and to incorporate their LLMs into the product portfolio. The integration will affect, i.e., the Teams Premium, where due to [intelligent recap](#) functionality, one will get the essential information from the meeting summarised in tasks and notes.

Overall, companies can apply technologies based on generative AI to automate repetitive tasks. Such an intelligent process automation (IPA) solution would help businesses save time and resources while improving efficiency and accuracy. Assume a scenario where instead of manual processing of huge and unclear excel sheets, one transforms them automatically using plain English! It sounds almost like magic, but such solutions are already available on the market. And there is even more. During the last [generative AI hackathon](#) from Scale, developers presented a project where the program’s backend was replaced entirely with an LLM, which could execute logic and store memory. The developers demonstrated this idea on a [todo-application](#).

At this point, one can look at the topic more precisely and consider the area of insurance. When asked about the potential application of generative AI in that sector, ChatGPT yields a list of rather encouraging suggestions. ChatGPT suggested the *underwriting* and *claims processing*, among other applications for the insurance domain. The objective of underwriting is to assess the amount of risk an individual or a legal entity presents to a potential insurer. Professional underwriters typically examine the application criteria to verify if it is eligible for the policy offering and determine the related coverage. Based on the analyzed information a monthly premium is set.

A large language model can be employed to *assist domain experts*. Historical claims and demographic data can be used to predict the probability of future claims. That could be used to estimate the risk associated with insuring the concrete individual or legal entity.

Furthermore, LLMs can analyze unstructured data, including texts from medical records, and extract vital details for underwriters.

 As an AI language model, ChatGPT can be used in the insurance domain in several ways:

1. **Customer service:** ChatGPT can be integrated with a company's customer service platform to provide customers with fast and efficient responses to their queries. ChatGPT can provide instant responses to frequently asked questions and can also handle basic customer service tasks like policy renewals and claims processing.
2. **Risk assessment:** ChatGPT can analyze data on various factors that contribute to risk, such as weather patterns, crime rates, and accident statistics, and provide insights to help insurance companies assess risk accurately. This can help insurance companies to develop more precise policies and pricing models.
3. **Fraud detection:** ChatGPT can help detect fraudulent claims by analyzing claims data and identifying patterns that are indicative of fraud. This can help insurance companies to reduce losses and improve the accuracy of their claims processing.
4. **Personalization:** ChatGPT can be used to personalize insurance offerings to individual customers based on their preferences, needs, and behavior. This can help insurance companies to build stronger relationships with their customers and improve customer retention.

Overall, ChatGPT can be a valuable tool for insurance companies looking to improve customer service, reduce risk, and optimize their operations.

**ChatGPT Plus** answers to the question regarding its usage in the insurance domain.

Yet, it is worth noting that despite the performance of LLMs, such models still need to be improved and are not entirely reliable. This, however, should not discourage companies from designing and leveraging AI systems for business purposes. Rather, companies must appropriately mitigate associated risks and protect customers from unintended harm by implementing comprehensive *responsible AI frameworks*.



**In a sense, Artificial Intelligence will be the ultimate tool because it will help us build all possible tools.**

(K. Eric Drexler, “Engines of Creation,” 1986)



# Outlook for Financial Services

Par for the course, generative AI will inevitably become one of the key priorities for the Financial Services domain – starting with banking and going toward insurance and asset management. The principal ambition will, however, remain the same – find the best ways to manage risks and improve overall accuracy by reducing operating costs. Furthermore, financial institutions will continue developing and encouraging explainable and responsible AI to boost the understanding of model predictions in financial applications.

According to the survey by [Nvidia](#), one of the focus points will become *hybrid cloud computing*. In this mixed computing environment, applications are run using a combination of public and private clouds, including on-premises data centers. This allows migrating all the intensive workloads associated with AI technologies to cloud environments, optimizing AI performance, and reducing costs. While taking this advantage, one can continue using on-premises services for storing and accessing highly confidential clients' data that can't be migrated to the cloud.

Another evolving topic where generative AI models can be employed for sensitive domains concerns *synthetic data generation*. The demand for high-quality data has grown exponentially over recent years. However, stricter

data privacy rules and high confidentiality of data in restrictive domains, such as financial services, have made aggregation and labeling of real-world data exceptionally difficult. At that point, synthetic data – computer-generated information – becomes ultimate. It is cheaper to produce than labeling real-world data by skilled domain experts. More importantly, it allows companies to avoid regulatory issues of handling personal data. For example, banking often relies on confidential client data to detect fraud or predict stock market trends. Banks adhere to strict data practices to keep this information secure, making it highly complicated for employees to access this data. At the same time, anonymizing confidential data can cause errors that severely compromise the quality of the final application. By contrast, synthetic financial data sets would not allow tracing back to individuals but still will retain the statistical properties of the real-world data.

Regarding the financial AI application, the most characteristic ones will be focused on building *recommender systems, intelligent portfolio optimization, and fraud detection applications*. Natural language processing techniques could be employed to automate financial document analysis and claim processing, saving businesses time, costs, and resources. Furthermore, NLP approaches can be utilized in anti-money laundering and yearly KYC processes.



# Limitations and dangers in the use of generated content

There is enough experimental data to suggest that LLMs and ChatGPT in particular, can respond to a broad spectrum of questions, delivering eloquent and comprehensive responses that significantly *outperform prior public chatbots*. Nonetheless, like any technology, ChatGPT has limitations. One fundamental constraint is that the model is only as good as the data it is trained on, and if the training data is biased or contains errors, then the model reproduces and amplifies them too. Another known limitation is that ChatGPT may *fabricate facts* to deliver an answer (e.g., in legal domains) if answering a question requires professional knowledge from a specific field. This phenomenon alerts to be extra cautious while using GPT-like systems for professional consultations (Guo, 2023).

The correctness issue also applies to code generation systems such as GitHub Copilot, since its underlying language model was trained on open-source code. Yet the code frequently contains bugs – and thus, it is inevitable that the model has partially learned from vulnerable, buggy data. This naturally raises concerns about the *security* of the code Copilot suggests. A recent study (Pearce, 2021) has examined the prevalence and conditions that can cause GitHub’s Copilot to recommend insecure code. For the analysis, the authors prompted Copilot to generate code in high-risk [CWE](#) scenarios. Afterward, they investigated Copilot’s performance along three axes – examining how it performed under various weaknesses, various prompts, and various domains. Copilot had to run 89 different scenarios, generating 1,692 programs overall. Of these, around *40% were found to be vulnerable*.<sup>1</sup>

Among other issues associated with AI, [experts](#) list data privacy and legal liability, compliance with various state and local regulations, and lack of regulation with explicit guidelines.

A comprehensive responsible AI framework that includes policies, governance, and tools to ensure that AI applications are built following rules and regulations can help companies reduce risks. Such a framework can decrease the frequency of failures by spotting and mitigating issues before the system’s deployment. And while failures can still occur, their severity would be lower, causing less harm to users.



<sup>1</sup> On February 14<sup>th</sup>, GitHub Copilot launched an AI-based vulnerability prevention system that blocks insecure coding pattern in real-time to make suggestions more secure.

# Does Artificial Intelligence think like a human?

Noam Chomsky, the legendary cognitive scientist, and father of modern linguistics, claims that new AI, centered on statistical learning techniques, is “unlikely to yield general principles about the nature of intelligent beings or cognition”. The researcher argues that this is particularly crucial for natural language understanding. If one gets vast data and the best statistics, one will naturally get a much better approximation to some given corpus of texts. Yet one will learn *little to no* about the language and the concepts underlying it.

At this point, it is worth talking about another significant limitation of AI related to world understanding. The current state of generative AI lacks the *commonsense knowledge* and *reasoning* vital to the success of AI-based systems. Despite competent problem-solving skills, such methods are still very brittle and can generate meaningless responses while attempting to reason.

The Cambridge Dictionary defines commonsense<sup>2</sup> as “a basic level of practical knowledge and judgment that a person needs to live a reasonable and safe life”. All people possess this knowledge, although it is unspoken and unwritten – we evolve it gradually from the moment we are born without even being conscious of it. For the purposes of AI, common sense “involves specific knowledge about the world and the ability to infer additional facts from observations and general knowledge and exploit that reasoning to achieve goals” (McCarthy, 1986).

The AI community strives to attain a commonsense understanding of AI-based systems. There are benchmarks and challenges that encourage researchers to tackle this topic. Among these challenges is the [Winograd Schema Challenge](#), a set of pronoun resolution problems initially designed to be unsolvable for statistical models. The challenge is an alternative to the Turing test, which has been shown *not* to be an ideal way to evaluate machine intelligence, as it can be tricked (i.e., [Eugene Goostman](#)). Allen AI researchers extended this challenge by introducing [WinoGrande](#), a large-scale dataset adjusted to improve the complexity (Sakaguchi, 2020).

Indeed, being asked about ambiguous samples, ChatGPT fails to comprehend and explain the rationale behind them<sup>3</sup>. Assume the sentence “*Visiting relatives can be a nuisance.*” This is an ambiguous phrase, and “*Visiting*

*relatives*” can be interpreted either as “*relatives who visit*” or as “*going to visit relatives*”. Depending on the case, the meaning of the sentence will be varied.

Prompting ChatGPT to explain the above sentence, one receives the following response: “*Visiting relatives can be a nuisance – means that visiting one’s relatives can be irritating, bothersome, or inconvenient*”. With this, the model is only considering the second scenario and cannot grasp the *structural ambiguity*.

What happens if one reformulates the prompt and passes supplemental knowledge to challenge the system? New prompt requested to explain the above sentence in two ways, first by treating “*visiting*” as an *adjective*, and second by treating it as a “*verb*”. However, even in this case, ChatGPT failed to deliver a correct explanation for the first scenario (i.e., adjective case). For some reason, it perceives “*Visiting relatives*” as “*relatives who are being visited*” instead of “*relatives who visit*” and interprets the sentence in the wrong way.

Experts from other fields have also examined ChatGPT on a variety of diverse topics, including [rocket science](#). The AI was asked to generate a rocket equation, and according to the reviewing scientist, it was “*missing too many variables*”. In nearly every other test case, the model was incapable of accurately reproducing even the most fundamental rocket engineering equations.

Still, the potential for introducing quantitative reasoning into GPT-like systems excites researchers in the scientific community. For example, Stephen Wolfram, founder of the knowledge-based system [WolframAlpha](#), has shared his [thoughts](#) on merging the two models. Such a merger could yield a platform capable of generating human-like responses and performing accurate computations that surpass human capabilities. Some AI researchers have already shared their first attempts at addressing this objective. Chat [LangChain](#), for example, combines [WolframAlpha](#) and GPT 3.5 (the underlying ChatGPT technology). Once a user prompts a task, the *inner monologue* unit of the system decides whether it can answer directly or if it needs to make an API call to [WolframAlpha](#) first.

<sup>2</sup> Example: When objects collide, they usually make a noise.

<sup>3</sup> We omit some typical examples from benchmarks since these are partially included in the newer version of the ChatGPT.



# Pacta sunt servanda and the correct management of intellectual property rights

Although current AI-based technologies play an increasing role in creating intellectual property, there is still no comprehensive and uniform approach to protecting AI-generated objects and determining their owners. In judicial practice, AI-based software is not yet acknowledged as a *subject of law*; therefore, for now, it can be only treated as the *author de facto* of generated content, whereas the *author de jure* has to be a human being.

Yet, the exponential advancement of technologies based on artificial intelligence is causing an acute need to adapt legal regulations in this area. Tools based on generative AI are no longer restricted to the execution of standardized tasks but can contribute to the creation of products that qualify as intellectual property objects.

Hence, it is reasonable to question who should be credited as the author of such objects. The developer of the AI-based software, the user of that software, the AI itself? Or is the sole solution to acknowledge the absence of an author when concerning results created by artificial intelligence?

Before the discussion, it is worth noting that the legislation of most jurisdictions today stipulates that only a *natural person capable of thinking and creating* can be admitted as the author of the outcome of an intellectual work. Furthermore, the author must contribute personally and *creatively* to the intellectual property object.

Now let's debate the option where the AI developer is the author of the AI-models outcome. One can argue that the AI model or software is the developer's creative work, and therefore the rights for the model outcome should be granted to them. In that case, a person who once wrote a program must be recognized as the author of all its works; however, once the AI model is implemented and trained, the *developer does not contribute personal creative work to the outcome produced by that model*. So, one can be acknowledged as the author of the algorithm behind the model but not as the author of the model's outcome.

But what if one assumes that AI itself is capable of being creative? In that case, the result of its creative contribution, of course, must be subject to copyright, but the author, *de lege lata*, cannot be acknowledged. Otherwise, the society would reach the necessity of admitting the *civil legal capacity of machine intelligence*, i.e., the ability of the object of civil law to have civil rights and obligations. Here it is noteworthy that, at the moment, AI still does not make independent decisions but only carries out the instructions of competent subjects, i.e., humans. Besides, is there any practical need to declare AI an author if it does not need a copyright? It can neither dispose of the exclusive right to the result nor protect personal and proprietary rights.

Furthermore, returning to creativity, is it reasonable to accept that an AI creates *consciously*? For the actions of an AI to be considered creative, the AI must strive to make decisions that are not a repetition (i.e., paraphrase) of decisions made earlier; there must be an *independent creative element* and a minimum of randomness. Yet the AI lacks imagination and replicates the concepts and patterns humans fed it, so there is *no genuine creativity* in place. When creating a picture, the AI relies on embedded data, and consequently, its activity is *imitative*; it only borrows the external form without *comprehending the concept* behind it. In contrast, an artist relies on his or her very imagination.

Lastly, one must admit, that imitation of an intelligent response and an intelligence that creatively solves a problem are not the same. There is also a vast difference between van Gogh's paintings and paintings *in the manner of van Gogh*. And since current law requires that the result of creative work must be entirely novel, should one also demand *originality* from the outcomes of machine intelligence?

Nevertheless, the above mentioned highlights the importance of an effective mechanism for the legal regulation of AI, which would create a clear delineation of areas of responsibility between the developers and users of artificial intelligence systems and the technology itself.

This challenge was already addressed from several points of view. One of the ideas is to attribute the status of *quasi-agent* (i.e., quasi-subject) – similar to that of minors – to AI (Asaro, 2011). Some researchers have also advocated the need to rethink the concepts of “creativity” and “author” and extend them to machine intelligence (Saripan, 2016). Others propose to rely on the “work for hire” doctrine and to expand this practice to the outcomes of intellectual activity of AI, thereby assigning relevant rights to the respective subject and establishing liability (Denicola, 2016).

However, one of the middle-ground ideas was stated by [Creative Commons](#) at the second round of the [WIPO Intellectual Property and Artificial Intelligence Discussion](#) in 2020. The notion behind it is that copyright and associated rights should not be guaranteed for works created through artificial intelligence due to two fundamental reasons: lack of human authorship and lack of originality. Hence, results

derived from the work of AI should be treated as being in the *public* domain. At the same time, using copyrighted products to train artificial intelligence should not be considered an infringement of the author’s rights<sup>4</sup> (Brigitte Vézina, 2020).

Admittedly, the field of artificial intelligence is still developing, and new technologies are emerging fast, so there is still much ambiguity and obscurity.

---

<sup>4</sup> However, as per January 2023 AI tools such as Stable Diffusion and [Midjourney](#) are being targeted in a copyright lawsuit. The lawsuit alleges that generative AI tools violate copyright law because they scrape artists’ work from the internet without their consent.



# Concluding remarks and KPMG practical guidance

In recent years, society has experienced many developments related to advances in artificial intelligence. Novel AI-based tools are emerging rapidly and becoming essential to people's daily lives. Some individuals are concerned that technological advances are driving humanity toward a "jobless future". However, while AI can carry out specific assignments, in most cases, it can only undertake part of the task and can rarely perform entire processes automatically. The latter makes large-scale automation of human jobs highly *unlikely* over the next few decades. Yet those organizations that leverage AI and intelligent process automation may free up human employees for more *complex and cognitively demanding tasks*.

Further, the field of generative AI is set to grow rapidly this year. Some tech giants and Silicon Valley start-ups have shared their intentions to bring comparable technologies to market. For example, DeepMind CEO, Demis Hassabis, announced that the company is considering [releasing its chatbot Sparrow](#) in a private beta version. Sparrow aims to "*reduces the risk of unsafe and inappropriate responses*" and features the ability to *quote sources* – something that ChatGPT lacks. Whereas, Meta has committed to open science by [releasing LLaMA](#), a state-of-the-art fundamental large language model tailored to support researchers in advancing their work in that AI area. Obviously, tools that *discriminate between generated content* will equally appear on the market to cope with the growth of machine-generated content and the associated threat of spam, impersonation, and fraud in writing and coding.

Generative AI can boost innovation and ultimately revolutionize work by eliminating many of the mundane, tedious tasks that people perform every day. Yet at the same time, there are associated risks and misuses that businesses need to be aware of, including bias, privacy violations and misinformation.

One of the prerequisites for the strategic integration of generative AI into a business is the identification of the mission-critical use cases that will yield the most value and impact. Once those use cases have been identified, it is vital to think through *policies* for the secure application of generative AI. Such policies must likewise

safeguard sensitive company and customer data. As part of the *responsible AI (RAI)* framework, such policies address the use of AI capabilities without causing harm or unintended impacts. Enterprises can leverage existing risk management tools and apply them to emerging technologies for implementing RAI processes.

In this regard, KPMG has deep functional knowledge and comprehensive expertise with data management solutions, such as state-of-the-art data warehouse architectures and advanced approaches for data governance (e.g., decentralized data mesh-type architectures). KPMG also offers industry-leading practices to support enterprises on their data transformation and optimization journey while assisting in adopting suitable routines, policies, and tools. Along with that, KPMG has the competencies and knowledge to address the diverse *GRC* (Governance, Risk, Compliance) needs across industries, ensuring convergence and transparency of information to enhance productivity and resilience in a dynamic business environment. For that purpose, KPMG has developed an implementation framework to aid companies in launching or upgrading their GRC technology programs.

Generative AI will inevitably blend into business applications and particularly become one of the key objectives for the *financial services* domain. Our AI and NLP experts can respond to your inquiries around that topic and help you find the best way to intelligently automate your business processes using novel technologies. Uniquely, KPMG can deliver everything out of one hand starting from business and data analysis; design the use cases and define solution approaches; towards implementation and production using latest methods and technologies.

Please do not hesitate to contact us if you wish to receive more details on how KPMG can help your business to take advantage of the latest development in the field of generative AI and NLP.



# Bibliography

**Biyang Guo, X. Z. (2023).** How Close is ChatGPT to Human Experts? Comparison Corpus, Evaluation, and Detection. *arXiv:2301.07597*.

**Drexler, K. E. (1986).** Engines of Creation. *Anchor books*.

**Asaro, P. M. (2011).** A Body to Kick, but Still No Soul to Damn: Legal Perspectives on Robotics. *Robot Ethics ethical Soc. implications robotics*, pp. 169-186.

**Saripan, H. P. (2016).** Are robots human? A review of the legal personality model. *World Applied Sciences Journal*, pp. 824-831.

**Denicola, R. C. (2016).** Ex machina: copyright protection for computer generated works. *Rutgers UL Rev.*, 69, p. 251.

**Brigitte Vézina. (2020, July 7-9).** *Creative Commons' statement on artificial intelligence and copyright*. Retrieved from WIPO Conversation on Intellectual Property (IP) and Artificial Intelligence (AI): Second Session: [https://www.wipo.int/export/sites/www/about-ip/en/artificial\\_intelligence/conversation\\_ip\\_ai/pdf/ngo\\_creative\\_commons.pdf](https://www.wipo.int/export/sites/www/about-ip/en/artificial_intelligence/conversation_ip_ai/pdf/ngo_creative_commons.pdf)

**Pearce, H. A.-G. (2021).** An empirical cybersecurity evaluation of github copilot's code contributions. *ArXiv abs/2108.09293*.

**Amatriain, X. (2023, 01 04).** *Prompt Engineering 101: Introduction and resources*. Retrieved from <https://amatriain.net/blog/PromptEngineering>

**Vaswani, A. (2017).** Attention Is All You Need. *arXiv:1706.0376*.

**McCarthy, J. (1986).** *What is Common Sense and How to Formalize it?* Berlin: Springer.

**Sakaguchi, K. B. (2020).** Winogrande: An adversarial winograd schema challenge at scale. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05), 8732-8740.

**Mitchell, E. L. (2023).** DetectGPT: Zero-Shot Machine-Generated Text Detection using Probability Curvature. *arXiv preprint arXiv:2301.11305*.

**Kirchenbauer, J. G. (2023).** A Watermark for Large Language Models. *arXiv preprint arXiv:2301.10226*.

# Contacts

## KPMG AG

Badenerstrasse 172  
P.O. Box  
CH-8036 Zurich



### Thierry Kellerhals

Director  
Digital Innovation  
+41 79 281 22 50  
tkellerhals@kpmg.com



### Dr. Alena Moiseeva

Senior Data Scientist  
Digital Innovation  
+41 79 838 94 84  
alenameiseeva@kpmg.com

[kpmg.ch](https://www.kpmg.ch)



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. The scope of any potential collaboration with audit clients is defined by regulatory requirements governing auditor independence. If you would like to know more about how KPMG AG processes personal data, please read our Privacy Policy, which you can find on our homepage at [www.kpmg.ch](https://www.kpmg.ch).

© 2023 KPMG AG, a Swiss corporation, is a subsidiary of KPMG Holding AG, which is a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.