

Harnessing GPT for smarter asset management: prospects and perils

- Generative models like GPT could revolutionize the industry
- These models offer potential benefits while also presenting challenges
- Asset management can adapt and embrace GPT while also addressing limitations

Since its launch on 30 November 2022, ChatGPT¹ has become the fastest application to reach 100 million monthly active users^{2,3}. The excitement that the model has generated in the public's imagination has led to wild predictions about its impact on all aspects of society, ranging from the overly optimistic to the profoundly pessimistic. Less than six months after ChatGPT's launch, GPT-4 was released. As the time of writing, GPT-4 appears to be significantly more powerful than ChatGPT. So, what does it all mean?

In this article, we take a clear-eyed look at some possible implications of GPT, and more broadly, generative models, on the asset management industry. We give examples of how GPT might be applied in various workflows common to the asset management industry, and how asset owners and managers could position themselves to benefit as much as possible from the anticipated future generative model evolution. We begin our discussion by briefly looking back on the development of natural language processing (NLP) models that have led to Generative Pre-trained Transformers (GPT).

A quick survey on NLP models

NLP is a subfield of machine learning (ML) that focuses on enabling computers to understand, interpret, and generate human language. It combines computational techniques with linguistic knowledge to do this in a way that mimics human understanding and communication. Since the first NLP model in the 1960s⁴, leaps in computational power particularly in the last three decades have meant that the ability to analyze, process and interact with text on an ever-larger scale has increased tremendously.

NLP applications for financial investments in particular started in 2007 with the humble bag-of-words (BoW) approach⁵, which is based on identifying words associated with either positive or negative sentiment as listed in a dictionary. Despite the simple nature of the model, a BoW approach to basic sentiment detection (overall positive or overall negative) works surprisingly well.

¹ Apart from the examples given, none of this paper was written by ChatGPT or any other form of AI.

² <https://aibusiness.com/nlp/ubs-chatgpt-is-the-fastest-growing-app-of-all-time>

³ [ChatGPT sets record for fastest-growing user base - analyst note | Reuters](#)

⁴ ELIZA, developed by Joseph Weizenbaum

⁵ See Tetlock (2007) and Loughran and McDonald (2011)

ARTICLE 3 APRIL 2023

Marketing material for professional investors, not for onward distribution

From left to right: **Clint Howard** Quant researcher, **Tim Vogel** Quant researcher, **Thom Marchesini** Quant researcher, **Mike Chen** Head of Alternative Alpha Research

The authors would like to thank David Blitz and the Robeco Quant Equity Research Team for fruitful suggestions and discussions.



The next breakthrough in NLP in general came in 2013 from Google's research lab⁶: the embedding approach to textual representation. This converts words and text to high-dimensional mathematical structures (vectors), which exhibit interesting properties that can be used in novel ways. For example, word embeddings can be used to detect and understand colloquial slang, because similar words tend to be near each other in the vector space. This is due to the proximity of slang words to their corresponding official dictionary words⁷, allowing language models to understand and process slang even if they're not part of official vocabulary.

The third notable breakthrough in NLP emerged from Google's research lab in 2017⁸, with the introduction of the so-called Transformer. This deep neural network architecture enabled more efficient network training. As opposed to its slower predecessors, which were recurrent in nature and relied heavily on the output of one iteration as input for the next, the Transformer model could be trained on multiple computers at a time. This increased the possible model size and amount of data that could be processed immensely, facilitating far more rapid training. This Transformer architecture is also used by and constitutes the 'T' in GPT, which we explore further below.

GPT, ChatGPT and generative models

GPT stands for 'Generative Pre-trained Transformer' and is a specific class of language models based on NLP. It was developed by OpenAI, an artificial intelligence research organization. As artificial neural networks that are trained on a massive corpus of text data, including books, articles and websites), GPT can generate human-like responses to text-based prompts. By chaining a sequence of input prompts, one can effectively have a conversation with the GPT model, and it is designed to generate a convincingly human-like response.

GPT falls into a general class of machine learning models called generative models. Generative models are a class of ML algorithms that model the joint probability distribution $P(X,Y)$ for items X and Y and output instances when X and Y are likely to co-occur. One should think of X as the text input ('the prompt') and Y as the output by the various GPT algorithms.

So, when one asks a generative model to draw⁹ a picture of a boat, it will also be likely to draw a body of water as these two items X (boat), and Y (body of water), tend to co-occur in the training data. As an example, we asked DALL-E to draw a picture of "A sailing boat with a horse, a dog, and a cat". We didn't specify water at all, yet every image produced has some representation of blue water, even those which have left out the horse.



⁶ See Mikolov et al. (2013)

⁷ See Chen, Lee, and Mussalli (2020)

⁸ See Vaswani et al. (2017)

⁹ DALL-E 2 (<https://openai.com/product/dall-e-2>) and Stable Diffusion (<https://stablediffusionweb.com/>) are instances of generative models that can generate pictures

Since GPT was launched in 2018¹⁰, there have been several generations, GPT-4 being the latest in the GPT family. It was released on 14 March 2023, which coincides with the writing of this report.¹¹ Each generation of GPT has shown a large leap in capability compared to its predecessor, and GPT-4 is no exception. For example, GPT-4's ability to reason and solve domain-specific tasks has increased massively, and the model can now process images along with textual inputs, meaning it can take multi-modal inputs. As an example of how GPT-4 improves on GPT-3.5 – the previous state-of-the-art model – GPT-4 scored a 90th percentile result on the Uniform Bar Exam versus its predecessor's 10th percentile effort.

As seen in our example (text box) the fact that generative models tend to output items that co-occur together statistically, while not necessarily deterministically, is an important feature. This gives applications like ChatGPT their “wow factor” as they generate seemingly coherent sentences and prose based on simple prompts. This feature also sometimes causes generative models to ‘hallucinate’ and confidently mention items or events that sound probable but are not real, which has led to the term ‘artificial hallucination’¹²: a confident response that is nonsense or not faithful to the original training data. Although there is still no guarantee of factual accuracy in the response, according to OpenAI GPT-4 is much more reliable than GPT-3.5 in terms of these artificial hallucinations, with a reduction of up to 40% compared to GPT-3.5¹³.

The GPT models have been used for various applications, including language translation, content creation, and search. The most well-known specific adaptation is perhaps ChatGPT, which is an adaptation of GPT-3.5 and GPT-4 for chatbot applications. Alongside ChatGPT, there are other adaptations of GPT for a variety of tasks. For example, in 2021 OpenAI together with Microsoft released an app called ‘GitHub Copilot’ which is based on GPT and specifically tailored to code writing, helping programmers to write code and develop software more efficiently.

“When a new technology emerges, one can choose to either learn about it, experiment with it and study its implications with an eye to adoption; or simply ignore it. We suggest that the former path is the more rational one.”

What are some use cases and opportunities in asset management?

The speed at which GPT and generative models have taken off can be seen as either exciting or intimidating, depending on one's comfort level with new technologies. When a new technology emerges, one can choose to either learn about it, experiment with it and study its implications with an eye to adoption; or simply ignore it. We suggest that the former path is the more rational one. Choosing to explore a new technology that then fails to catch on will not result in any harm outside of the potential loss of invested time without tangible outcome. On the other hand, if the technology does catch on, those who initially chose to ignore it will likely be left behind or, at the very least, have to catch up with competitors.

Moreover, given the current excitement and massive amount of investment going into GPT and generative models¹⁴, it can create a sense of urgency that an organization *should* be doing something with GPT. To counter this we advocate a measured investigation into how this type of technology could assist various aspects of an asset management organization's workflow. The technology might not be fully ready today to deploy into all processes. Still, it is becoming more probable that it will be widely available across enterprise platforms as the technology matures. For example, Microsoft has already integrated parts of GPT into Microsoft Teams and Bing Search functionalities, and that's just the beginning.

¹⁰ See Radford et al. (2018)

¹¹ For clarity and ease of reading, the authors will use the term GPT to represent the general class of recently developed GPT models, including GPT-3.5 and GPT-4. Specific versions of the GPT model will be denoted by the version number, such as GPT-4.

¹² See Ji et al (2023)

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

¹⁴ There has been a 425% increase in venture capital funding in generative AI. Microsoft's significant partnership with OpenAI then earmarked this growth.

There are numerous potential applications for a mature GPT model in the asset management industry. We categorize these applications into two areas: investment and productivity.

Investment applications

A key component in a quantitative researcher's toolkit is the ability to quickly find and synthesize information to aid in making investment decisions. GPT can act as an accelerator for this phase of the alpha research process. It also has implications for the current use of NLP in the investment domain.

Alpha opportunities

This could involve incorporating GPT-derived alpha factors into the model or using the tool in portfolio construction. Could GPT serve as a source of alpha, such as by processing earnings call transcripts more effectively than current methods?

The implications of GPT for NLP in the investment domain warrant further discussion. As more AI-generated content becomes prevalent, the efficacy of NLP signals may be called into question. However, this may also present an opportunity for GPT to augment or replace existing NLP applications. By rethinking information consumption and creating factors based on generative models, new alpha beyond the reach of current software could be unlocked.

GPT could help investors better understand the underlying economic dynamics, for example by enhancing linkage analysis, leading to richer networks, and leveraging digital information flows between references to companies and assets. Identifying distressed companies may also become more efficient, as the tool can facilitate mapping people and companies while linking unorthodox information, potentially improving the detection of fraud or distress. Figure 1 is an illustration of a linkage question we asked ChatGPT and its answer.

We see that ChatGPT is able to connect quite a few different industries, and additionally explain why these industries are connected. We could argue that as a financial markets expert, we might already consider these industries connected, but the layperson, for example, might find the agriculture and pharmaceutical connection quite insightful.

Figure 1 - Q&A with ChatGPT



I want you to act as an expert in financial markets. What are some industries which are interconnected but people often don't think they are connected?

As an AI language model acting as a financial markets expert, I can provide you with some examples of industries that are interconnected but people might not immediately think of them as being connected. Here are a few examples:

1. **Energy and Transportation:** The energy sector, which includes oil, gas, and renewable energy, is closely linked to the transportation sector. Fuel prices directly impact transportation costs, which in turn affect the cost of goods and services.
2. **Agriculture and Pharmaceuticals:** The agriculture industry relies on the pharmaceutical industry for various products like fertilizers, pesticides, and growth regulators. Conversely, the pharmaceutical industry uses plant-derived compounds and biotechnology to develop new drugs and treatments.
3. **Technology and Healthcare:** The technology sector plays an increasingly important role in the healthcare industry through the development of medical devices, telemedicine, and health informatics. In turn, healthcare advancements drive demand for new technologies and data solutions.
4. **Real Estate and Banking:** The real estate market and the banking sector are closely connected. Banks provide financing for real estate development and mortgages for homebuyers, while the health of the real estate market influences the stability of the banking sector.
5. **Manufacturing and Environment:** Manufacturing industries often have a significant impact on the environment, as they generate waste and consume resources. As environmental regulations become more stringent, manufacturing companies need to adapt and invest in sustainable practices, which in turn drives demand for environmental products and services.
6. **Consumer Goods and Raw Materials:** The consumer goods industry relies on the extraction and processing of raw materials, including metals, minerals, and timber. The prices and availability of these raw materials can significantly affect the cost and production of consumer goods.
7. **Insurance and Construction:** The insurance industry is closely linked to the construction industry, as construction projects often require various forms of insurance coverage. Additionally, natural disasters and other catastrophic events can lead to a surge in construction activity, impacting both industries.

These are just a few examples of interconnected industries, but there are many more connections between sectors in the global economy. It is important for investors to consider these relationships when analyzing financial markets and making investment decisions.

Source: <https://chat.openai.com/chat?model=gpt-4>

Research

In research, ChatGPT can be employed in scenario analysis for quick surveys on bespoke topics, and in the future, it could enhance AI-driven literature review toolkits such as ResearchRabbit.¹⁵ As an integral component in a quantitative researcher's toolkit, ChatGPT can accelerate the information-gathering and synthesis process, thereby enabling investment professionals to make more informed decisions.

For example, ChatGPT's ability to analyze a large amount of text data can potentially help identify novel or emerging investment themes and trends that may not yet be well-known or widely recognized. ChatGPT can also be used to generate creative investment hypotheses based on existing data, potentially helping investors explore unconventional opportunities or strategies that might otherwise be overlooked.

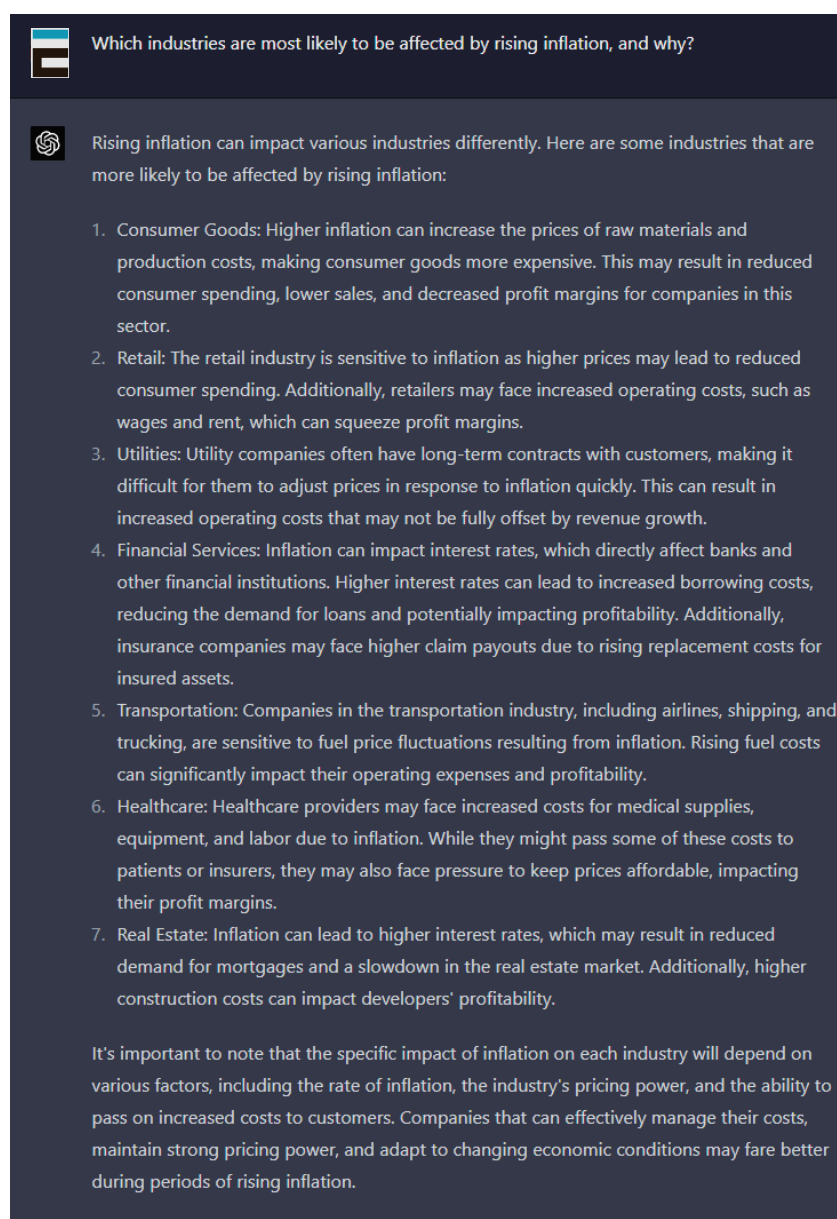
A particularly intriguing aspect of its use for researchers is for ChatGPT to function as a sparring partner and research assistant. This could be by asking ChatGPT to evaluate an idea, critique the structure of a piece of writing, or to conduct scenario analyses. For example, by querying GPT about the industries most likely to be affected by rising inflation or those that could benefit from a rushed energy transition, investment professionals can quickly gain valuable insights to inform their decisions.

As an example, we asked ChatGPT "Which industries are most likely to be affected by rising inflation, and why? As can be seen in Figure 2, ChatGPT provided a fairly comprehensive list of industries it considers as potentially impacted by rising inflation. ChatGPT provides reasonably clear explanations for what it considers to be affected by rising inflation but seems to have taken quite a negative perspective, focusing only on negative impacts and less on any potential opportunities.

Finally, we asked, "Which factors are likely to be more affected by rising inflation, and why?", the answer to which can be seen below in Figure 3. ChatGPT here again focuses on the negative impacts, and produces answers which are believable, but may not fully align with previous historical experiences. For example, Blitz (2021) shows how, historically, the value factor has been relatively immune to inflationary conditions, while the momentum factor has tended to outperform in inflationary environments. This highlights how, while ChatGPT can provide coherent and believable answers, these should always be rigorously checked.

¹⁵ See <https://www.researchrabbit.ai/>

Figure 2 - Second Q&A with ChatGPT



Which industries are most likely to be affected by rising inflation, and why?

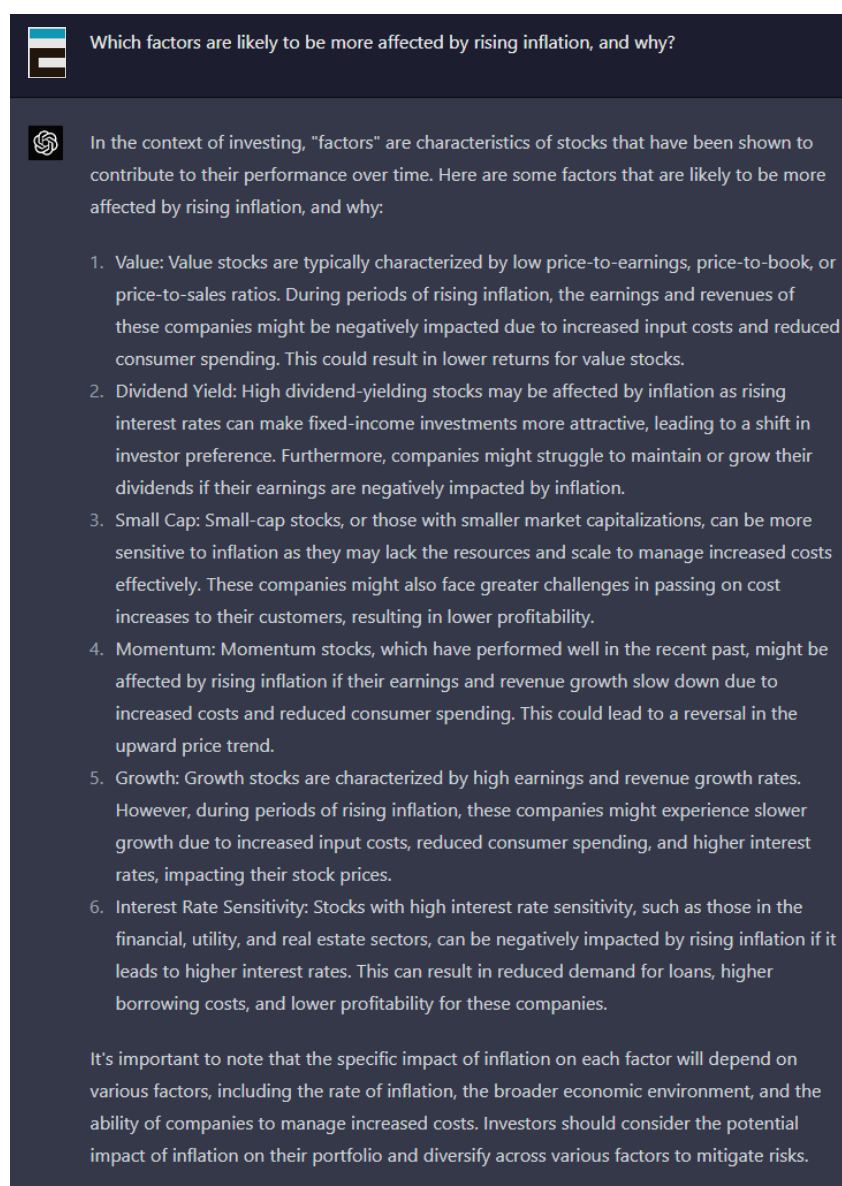
Rising inflation can impact various industries differently. Here are some industries that are more likely to be affected by rising inflation:

1. Consumer Goods: Higher inflation can increase the prices of raw materials and production costs, making consumer goods more expensive. This may result in reduced consumer spending, lower sales, and decreased profit margins for companies in this sector.
2. Retail: The retail industry is sensitive to inflation as higher prices may lead to reduced consumer spending. Additionally, retailers may face increased operating costs, such as wages and rent, which can squeeze profit margins.
3. Utilities: Utility companies often have long-term contracts with customers, making it difficult for them to adjust prices in response to inflation quickly. This can result in increased operating costs that may not be fully offset by revenue growth.
4. Financial Services: Inflation can impact interest rates, which directly affect banks and other financial institutions. Higher interest rates can lead to increased borrowing costs, reducing the demand for loans and potentially impacting profitability. Additionally, insurance companies may face higher claim payouts due to rising replacement costs for insured assets.
5. Transportation: Companies in the transportation industry, including airlines, shipping, and trucking, are sensitive to fuel price fluctuations resulting from inflation. Rising fuel costs can significantly impact their operating expenses and profitability.
6. Healthcare: Healthcare providers may face increased costs for medical supplies, equipment, and labor due to inflation. While they might pass some of these costs to patients or insurers, they may also face pressure to keep prices affordable, impacting their profit margins.
7. Real Estate: Inflation can lead to higher interest rates, which may result in reduced demand for mortgages and a slowdown in the real estate market. Additionally, higher construction costs can impact developers' profitability.

It's important to note that the specific impact of inflation on each industry will depend on various factors, including the rate of inflation, the industry's pricing power, and the ability to pass on increased costs to customers. Companies that can effectively manage their costs, maintain strong pricing power, and adapt to changing economic conditions may fare better during periods of rising inflation.

Source: <https://chat.openai.com/chat?model=gpt-4>

Figure 3 - Third Q&A with ChatGPT



Which factors are likely to be more affected by rising inflation, and why?

In the context of investing, "factors" are characteristics of stocks that have been shown to contribute to their performance over time. Here are some factors that are likely to be more affected by rising inflation, and why:

1. Value: Value stocks are typically characterized by low price-to-earnings, price-to-book, or price-to-sales ratios. During periods of rising inflation, the earnings and revenues of these companies might be negatively impacted due to increased input costs and reduced consumer spending. This could result in lower returns for value stocks.
2. Dividend Yield: High dividend-yielding stocks may be affected by inflation as rising interest rates can make fixed-income investments more attractive, leading to a shift in investor preference. Furthermore, companies might struggle to maintain or grow their dividends if their earnings are negatively impacted by inflation.
3. Small Cap: Small-cap stocks, or those with smaller market capitalizations, can be more sensitive to inflation as they may lack the resources and scale to manage increased costs effectively. These companies might also face greater challenges in passing on cost increases to their customers, resulting in lower profitability.
4. Momentum: Momentum stocks, which have performed well in the recent past, might be affected by rising inflation if their earnings and revenue growth slow down due to increased costs and reduced consumer spending. This could lead to a reversal in the upward price trend.
5. Growth: Growth stocks are characterized by high earnings and revenue growth rates. However, during periods of rising inflation, these companies might experience slower growth due to increased input costs, reduced consumer spending, and higher interest rates, impacting their stock prices.
6. Interest Rate Sensitivity: Stocks with high interest rate sensitivity, such as those in the financial, utility, and real estate sectors, can be negatively impacted by rising inflation if it leads to higher interest rates. This can result in reduced demand for loans, higher borrowing costs, and lower profitability for these companies.

It's important to note that the specific impact of inflation on each factor will depend on various factors, including the rate of inflation, the broader economic environment, and the ability of companies to manage increased costs. Investors should consider the potential impact of inflation on their portfolio and diversify across various factors to mitigate risks.

Source: <https://chat.openai.com/chat?model=gpt-4>

Productivity applications

Incorporating GPT into a domain expert's workflow may allow them to spend more time on high value-add tasks and less on more routine, process-oriented tasks. We delineate the use cases below by tasks that asset management employees regularly perform.

Reporting

A common task for asset management companies is reporting. For example, portfolio managers usually need to write monthly or quarterly performance reports to the asset owners. This task is not only quite time-consuming but also mechanical and repetitive. With some initial effort, this task could be mostly automated. To do this, a data pipeline to feed GPT the data used in the report must be created. These inputs could be stock returns, performance attribution, and prompts to assist GPT in writing reports in the user's style.

This fine-tuning of prompts is called prompt engineering. Our experience is that the prompt engineering necessary to generate human-quality reports devoid of glaring errors takes the most time, while it is in fact also a task that can at least partially be delegated to ChatGPT as well. ChatGPT can then generate a first draft of the report ready for domain expert validation and review, changing the role of the domain expert from creator to validator.

In terms of process, a substantial amount of time can now be freed, some or all of which can then be allocated by the expert to incorporating human-centric inputs necessary for a comprehensive report. For general writing tasks, Noy and Zhang (2023) found that ChatGPT, based on GPT-3.5, substantially raised average productivity in both time taken and output quality.

“Compressing information to the most salient points could streamline fundamental research workflows.”

Summarization

In asset management, one typically needs to read many documents, but one may not always have the time to do so. For example, during earnings season, a fundamental analyst could be asked to read tens of earnings call transcripts and various associated news releases every day while also producing accurate and timely updates of these companies' price forecasts. A quant researcher may need to read tens of academic journal articles when investigating a novel idea to see if it is worth researching further.

GPT and other generative language models can take large texts as input and then condense the information when given an appropriate prompt. Compressing information to the most salient points could streamline fundamental research workflows.

GPT can also help the user focus a summary on specific topics that are relevant to their research. For example, suppose the user is a financial analyst wondering about changes in expected earnings compared to last year. With a well-designed prompt, GPT can quickly analyze an article and give the analyst a concise summary highlighting the key earnings-related points. This makes it easy to quickly get to the heart of what the user is looking for in any given piece of text.

Coding

The daily workflow of a typical quant researcher can involve a lot of programming. This often includes frequent questions regarding nuances of various programming languages, such as how a particular function can be efficiently written. Until now, a researcher would turn to the internet at such moments, searches for information, collating it from various sources, interpreting and finally adapting it to the use case in question. Unsurprisingly, this can be quite inefficient and time-consuming.

Fortunately, generative models can be used as a coding assistant to quant researchers. The GitHub Copilot tool has been trained specifically on code repositories across GitHub and can provide code suggestions and structures. This can act as a productivity-enhancing tool, with the researcher spending considerably less time looking up basic coding structures and more time on the logic they wish to implement.

In a recent experiment, GitHub tested the effectiveness of AI-powered coding assistance tools. They first divided 95 developers into two groups.¹⁶ One group received AI assistance, while the other group did not. The developers were given the task of creating an HTTP webserver, and their coding times were recorded. The group with AI assistance completed the task 55% faster than the group without AI assistance, showcasing the significant improvement that can be achieved when using AI-powered coding assistance tools. However, creating an HTTP webserver may be a task where the additional help of AI code completions is particularly beneficial.¹⁷ Generally speaking, both hard and anecdotal evidence is mounting from the world of coding that GPT can result in significant time savings.

Email

Composing emails can cost time and energy. While the writer tends to already know what they need the core message of their email to be, how engaging or activating an email ends up being often hinges on the specific way it is phrased. Providing a quick first draft or bullet points and prompting a GPT API to rewrite it in the user's desired style enables the writer to create engaging email very quickly by saving time at the initial drafting stage. In situations where speed is more important than perfection and where mistakes are unlikely to have severe consequences, such as for internal emails, this can be a significant productivity enhancer.

However, using a GPT API for external emails is more complex. Errors in such messages can have more severe consequences for the reputation of the company and the relationship with the intended recipient, and at the same time might well be more likely because of GPT's confident tone and fluency. Despite these concerns, significant time and energy can still be saved.

For example, some types of emails share a standard structure, style and tone, with only minor deviations in the content for each recipient. Companies typically already have templates for such emails, which are copied and edited by the writer responsible. An initial draft of the email content can be generated by using a data pipeline and prompt engineering techniques, as outlined in the reporting section.

“GPT can help SI investors save a tremendous amount of time in writing coherent and polished engagement letters, while still allowing the SI investors to choose engagement topics.”

This offers two major advantages. First, by not needing to search for specific data or to create the initial draft, the writer saves significant time; even when accounting for proofreading and finetuning the draft to make it more directed and engaging. Secondly, automating much of the process reduces the risk of errors that can arise from the manual search for required information.

Engagement letter writing

Closely related to email writing is engagement letter writing, which is an important and effective tool for sustainability-minded investors to get companies to improve their sustainability practices. GPT can help SI investors save a tremendous amount of time in writing coherent and polished engagement letters,

¹⁶ <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/>

¹⁷ Writing HTTP webservers involves the writing of a large amount of boilerplate code, code that is repetitive and has little variation, which can be time-consuming to write manually, while it is easy for an AI to predict.

while still allowing the SI investors to choose engagement topics. To systematize engagement letter writing, a data pipeline needs to be built to feed GPT on various sustainability statistics for the engaged companies (such as their gender pay gap, carbon emission, etc.). Along with the sustainability data pipeline, efforts must be spent on prompt engineering, and samples of previous engagement letters must be provided in order for GPT to write in a coherent style similar to that used in previously letters from the SI investor.

In 2022 Robeco developed a system to use GPT to assist in writing SI engagement letters, thereby enhancing the company's sustainability professionals' impact and reach while ensuring they retained total control regarding the company and engagement topics.

Marketing

For marketing, one needs creative ideas, which often come via brainstorming and sparring with colleagues. GPT can act as a sparring partner when one is generating ideas on a specific topic. GPT can also generate catchy titles and bylines for marketing articles and taglines for marketing campaigns. From the textual description, graphic generative models can create images to be used in the marketing campaign. Some generative models can even create movies from textual descriptions¹⁸, and these movies can also be used in marketing campaigns.

What to watch out for

There is a lot of excitement about the current and future possibilities of generative models, particularly GPT. However, it is crucial to recognize and identify potential pitfalls of the technology, as well as situations where things could go wrong. As of the time of writing, we see GPT and generative models as versatile and creative tools one can use and experiment with but cannot and should not trust or rely upon completely.

Firstly, GPT struggles with domain-specific tasks. Even GPT-4's output quality in terms of depth of knowledge is no match compared to that of a true human expert. GPT has been trained on a very broad set of textual inputs, so it is designed to work on the broadest set of possible tasks. It therefore lacks a lot of the nuance and detail that goes into expert work. At this point, if a user wanted to rely on GPT to for factual answers to highly technical or specific questions, they would likely instead be given convincing but often superficial.

Secondly, the answers can also be incorrect. As mentioned earlier, these are known as artificial hallucinations, by-products of how generative models fundamentally work. These are dangerous because GPT will confidently give users plausible-sounding answers even though they may be categorically incorrect, and on occasion will even try to drive its point home.

*“ ... it is easier for humans to believe what GPT says and take the generated text at face value... [So] there always needs to be **human oversight**.”*

This brings us to our third point. GPT is very convincing at sounding like a human being. This means it feels more natural for humans to interact with GPT (compared to previous chatbots), but it also means it is easier for humans to believe what GPT says and take the generated text at face value.

To address these concerns, there always needs to be **human oversight**. That is, there must be a knowledgeable person and/or human expert in the workflow process that can verify and authenticate all the output generated by these models, especially when outputs contain any specific references to historical events, techniques, books and journal articles for example. Additionally, GPT may encounter challenges when generating text that includes numbers, which is a common issue for many generative models.

¹⁸ <https://makeavideo.studio/>

Another element that may be problematic is that a GPT model's training phase by necessity ends at a certain fixed point in time, which means that it cannot accurately generate text about events or information that have occurred since being deployed. However, there is potential to overcome these limitations as the generative model ecosystem continues to evolve. For example, the recent¹⁹ release of 'the new Bing' integrates up-to-date information accessible through search engines with a GPT algorithm, which could help to improve the accuracy of generated text and also overcome the challenge of generating text with numbers.

Conversing with GPT can be challenging, as language is a complex and nuanced system. When asking GPT questions, it's important to bear in mind that the quality and granularity of the question will often dictate the quality and granularity of the response. Relying solely on GPT's initial response may not provide optimal results, especially if the question does not fully reflect what one wants to know. A popular technique to address this is prompt chaining²⁰, which involves asking a series of questions guiding GPT towards the desired information. Additionally, GPT is sensitive to the prompts it receives. Even slight changes to the wording of a question or providing additional context can significantly alter its response.

Finally, as the asset management industry is a highly regulated industry, the use of GPT faces challenges concerning compliance and regulation, including data privacy. Ensuring the confidentiality of sensitive information is paramount, as asset managers often handle confidential client data and proprietary investment strategies. While GPT models are highly effective in processing and generating text, they may inadvertently retain or expose information provided during training or user interactions. Although organizations behind these generative models all have some policy in place regarding data privacy²¹ in place, the fact that all GPT queries and processing happens offsite is, for asset managers, a reason for caution.

This risk has already led several institutions to block the use of GPT models in order to safeguard client privacy and protect intellectual property. As a result, it will become crucial for asset management, as a regulated industry, to establish rigorous control mechanisms and monitoring systems to mitigate the risks associated with the business use of these models. As GPT proliferates throughout the industry, it must comply with existing financial regulations and be controlled and developed hand in hand with appropriate regulatory oversight.

“Trying to predict the future is always a fraught and, some may say, futile exercise

The future

Trying to predict the future is always a fraught and, some may say, futile exercise. Nevertheless, if generative models solidify their place in the technology sphere in three to five years – the distant future in ML research – what could that mean for society in general and the asset management industry in particular?

In terms of the broader society, we see GPT having three key potential implications. First, one way of viewing these generative models is as a conventional wisdom archive. These models could provide quick and easy access to effectively all conventional digitalized wisdom collected by humankind. A micro example of this would be organizations being able to feed generative models their own proprietary data and developing personalized chatbots for their specific business domain. In essence, these models democratize access to the 101 level of knowledge, both at a general society level and within an organization. While you can access such knowledge right now, it is quite fragmented and often curated via Google.

Secondly, this technology also brings the potential for extreme productivity enhancements, taking away the need for workers to do the tasks that these models can do. On one hand, this could usher in concepts such as the “four-day work week”²² across the globe.

¹⁹ [Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web - The Official Microsoft Blog](#)

²⁰ [GitHub - f/awesome-chatgpt-prompts: This repo includes ChatGPT prompt curation to use ChatGPT better.](#)

²¹ See, for example, OpenAI at [help.openai.com/en/articles/7039943-data-usage-for-consumer-services-faq](#)

²² [Join a pilot program – 4 Day Week Global](#)

On the other hand, it could mean further polarization of society between those gainfully skilled with abilities generative models do not have, and those who are replaced, such as run-of-the-mill writers and coders. Machines slowly took over blue-collar jobs over the past decades: this trend may now occur in traditionally considered white-collar jobs. The workers being replaced must either upskill or face un- or under-employment.

Finally, if these technologies become embedded into society (as social media has done), what will happen to the trustworthiness of information? What will unleashing generative models and their associated deluge of content on the internet or social media result in? Will the trustworthiness of all information vanish entirely? How will society fact-check? What will this mean for global democracies with their open information access? These unknowns have broad implications for everyone, and especially for businesses that must navigate the use of publicly available information.

For the asset management industry, as machines become more human, one potential implication is that humans can become even more human. That is, people working in the asset management industry must focus more on compassion, creativity, specialized knowledge, and sound judgment. Empathy has always been important for those with client-facing roles, but as GPT and similar generative models proliferate, clients will value authentic human connection and interaction even more.

For those in creative roles, such as quant researchers, portfolio managers, equity analysts and marketing professionals, creativity is something that generative models cannot replicate entirely, and therefore will continue to stay valued. While GPT can certainly prompt creativity, it does not (yet) have an imagination of its own.

People with specialized knowledge will retain their importance to the organization as generative models are usually trained on a general corpus, so specialized knowledge that generative models are not trained on and cannot regurgitate remains important. Even with company-specific chatbots in use, knowledge experts will remain indispensable thanks to their ability to accurately evaluate, assess and enhance any content produced by GPT or similar models. In recent research²⁸ from OpenAI on labor market impact of LLMs, science and critical thinking are shown to be skills or areas which will be the least impacted by GPT. In other words, these elements remain irreplaceable.

Finally, machines cannot be held responsible in courts, and they cannot weigh all the pros and cons of a given decision, incorporating both technical feasibility and financial implications, but also exercising sound moral judgments. The ability to make the correct business decision and then be accountable for the consequences will be as important as ever.

Will GPT democratize finance? We see various areas in which a mature GPT technology could lower the barriers to entry. The writing of finance research papers, for example, is no longer constrained to a privileged command of the English language, as GPT can enhance a non-native English speaker's translation of core ideas into a well-structured and grammatically flawless paper.

Furthermore, any individual can query a GPT-based application for tailored and specific financial advice. Financial advice prompt templates could be provided to a user, who would then have full ownership and responsibility over how they use the advice that emerges from the GPT-based application. Similar to the explosion in retail trading and the advent of trading platforms like Robinhood, will GPT bring even more options to retail or individual investors without a large professional organization supporting them?

The SEC has put strict controls around registrants offering "robo-investing" tools. Whether GPT will withstand regulatory scrutiny and whether it is potentially subject to manipulation remains to be seen. It is also currently not clear what users' obligations will be to check for potential abuses and what controls they have to have in place to

²⁸ [GPTs are GPTs: An Early Look at the Labor Market Potential Impact of Large Language Models](#)

use this technology to invest client assets. Therefore the potentially investment applications we describe above must comply with regulation designed to protect investors and the financial system.

Conclusion

Most technologies are both a blessing and a curse. This has happened since the beginning of time, from the invention of the wheel, which made it easier to travel and easier to spread disease, to the splitting of the atom, which made energy abundant and weapons nuclear, to the creation of the internet, which made both information and misinformation easy to access and disseminate.

Generative models will likely be no different. As a society, we will yet again have to navigate a technological advance to maximize the overall good and minimize the overall bad. Humans have done it before; surely we can do it again.

How, then, should organizations and individuals deal with the potential disruptions generative models represent at a more micro scale? As explored in this paper, progress waits for no one and no organization; and it's better by far to be ahead of the curve. To successfully adapt to and embrace the possibilities of generative models like ChatGPT, it is crucial for us to approach their potential with an open, rational mind. By fostering a spirit of exploration and innovation, we can unlock the benefits of this technology while actively addressing the challenges it may pose. Through collaboration and thoughtful consideration, we can harness the power of generative models to enhance our lives and create a more interconnected and dynamic future.

Bibliography

- Blitz, D. (2021). "The Quant Cycle." *Journal of Portfolio Management*, 48(2).
- Chen, M., Lee, J., and Mussalli, G. (2020). "Teaching Machines to Understand Chinese Investment Slang." *Journal of Financial Data Science*, 2(1), 116-125.
- Ji, Z., Lee, N., Friske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotta, A., and Fung, P. (2023). "Survey of Hallucination in Natural Language Generation." *ACM Computer Surveys*, 55(12), 1-38.
- Loughran, T., and McDonald, B. (2011). "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance*, 66(1), 35-65.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). "Distributed Representations of Words and Phrases and their Compositionality." *Advances in Neural Information Processing Systems*.
- Noy, S., and Zhang, W. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *MIT Working Paper*.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). "Improving Language Understanding by Generative Pre-Training." *OpenAI Working Paper*.
- Tetlock, P. (2007). "Giving Content to Investor Sentiment: The Role of Media in the Stock Market." *Journal of Finance*, 62(3), 1139-1168.
- Vaswani, A., Shazeer, N., Parmar, N., J. Uszkoreit, Jones, L., Gomez, A., Kaiser, L., and Polosukhin, I. (2017). "Attention is All you Need." *Advances in Neural Information Processing Systems*.

Important Information

Robeco Institutional Asset Management B.V. (Robeco B.V.) has a license as manager of Undertakings for Collective Investment in Transferable Securities (UCITS) and Alternative Investment Funds (AIFs) ("Fund(s)") from The Netherlands Authority for the Financial Markets in Amsterdam. This document is solely intended for professional investors, defined as investors qualifying as professional clients, who have requested to be treated as professional clients or who are authorized to receive such information under any applicable laws. Robeco B.V. and/or its related, affiliated and subsidiary companies, ("Robeco"), will not be liable for any damages arising out of the use of this document. The contents of this document are based upon sources of information believed to be reliable and comes without warranties of any kind. Any opinions, estimates or forecasts may be changed at any time without prior notice and readers are expected to take that into consideration when deciding what weight to apply to the document's contents. This document is intended to be provided to professional investors only for the purpose of imparting market information as interpreted by Robeco. It has not been prepared by Robeco as investment advice or investment research nor should it be interpreted as such and it does not constitute an investment recommendation to buy or sell certain securities or investment products and/or to adopt any investment strategy and/or legal, accounting or tax advice. All rights relating to the information in this document are and will remain the property of Robeco. This material may not be copied or used with the public. No part of this document may be reproduced, or published in any form or by any means without Robeco's prior written permission. Investment involves risks. Before investing, please note the initial capital is not guaranteed. This document is not directed to, nor intended for distribution to or use by any person or entity who is a citizen or resident of or located in any locality, state, country or other jurisdiction where such distribution, document, availability or use would be contrary to law or regulation or which would subject Robeco B.V. or its affiliates to any registration or licensing requirement within such jurisdiction.

Additional Information for US investors

This document may be distributed in the US by Robeco Institutional Asset Management US, Inc. ("Robeco US"), an investment adviser registered with the US Securities and Exchange Commission (SEC). Such registration should not be interpreted as an endorsement or approval of Robeco US by the SEC. Robeco B.V. is considered "participating affiliated" and some of their employees are "associated persons" of Robeco US as per relevant SEC no-action guidance. Employees identified as associated persons of Robeco US perform activities directly or indirectly related to the investment advisory services provided by Robeco US. In those situation these individuals are deemed to be acting on behalf of Robeco US. SEC regulations are applicable only to clients, prospects and investors of Robeco US. Robeco US is wholly owned subsidiary of ORIX Corporation Europe N.V. ("ORIX"), a Dutch Investment Management Firm located in Rotterdam, the Netherlands. Robeco US is located at 230 Park Avenue, 33rd floor, New York, NY 10169.

Additional Information for investors with residence or seat in Canada

No securities commission or similar authority in Canada has reviewed or in any way passed upon this document or the merits of the securities described herein, and any representation to the contrary is an offence. Robeco Institutional Asset Management B.V. is relying on the international dealer and international adviser exemption in Quebec and has appointed McCarthy Tétrault LLP as its agent for service in Quebec.