

ChatGPT decoded

A comprehensive overview
of large language models

Thursday 16th March, 2023

ChatGPT decoded

A comprehensive overview of large language models

In a world where communication and language are key, ChatGPT has emerged as a game-changing technology that is revolutionising the way we work, communicate, and learn.

Beyond the initial hype, ChatGPT has disrupted content creation and search as we know it, with Microsoft incorporating the technology into Bing search and Google responding with a competing model. Academia has also been impacted with some institutions banning ChatGPT while others embrace its new approach to digesting content and acknowledge that we need to promote new ways of learning. We're also now beginning to see what impact ChatGPT will have on individuals and organisations as it is adopted by companies like BuzzFeed into their day-to-day operations.

ChatGPT has gained widespread adoption due to its capability to increase efficiency and boost productivity, particularly in the areas of content creation and software development. For organisations, there exists a great opportunity to adopt **large language models (LLMs)** into existing platforms and enhance any question-answer interactions such as chatbots and customer service, along with other use cases which are being developed and released every day. However despite the hype it is important to understand the limitations of the current version and how it should and shouldn't be used.

Here, we discuss important considerations about ChatGPT and how this generative AI approach will impact our ways of working moving forward.

What is ChatGPT?

ChatGPT is a **large language model** developed by OpenAI that is capable of generating human-like text. The model is trained on a large dataset of human-generated data, such as books, articles, and websites.

The model takes in a prompt, or a starting text, and then it uses what it has learned from the training data to make educated guesses about what words are likely to come next, based on the patterns it has observed. The generated text can be used for a variety of **natural language processing (NLP)** tasks, such as language translation, text summarisation, and dialogue generation. For instance, ChatGPT can be used to summarise long articles or documents, translate text from one language to another, or generate human-like responses in chatbot conversations.

Due to its versatility and broad range of potential applications, ChatGPT is expected to make significant contributions to the field of natural language processing and influence how we interact with technology in our everyday lives.

Why the hype?

What sets ChatGPT apart from other large language models is its capacity to produce **responses that are more natural and less robotic**, unlike anything seen before. This, together with its **user-friendly interface**, has contributed to the quick adoption and hype around the product. ChatGPT hit an impressive milestone of reaching 1 million users in just five days after its launch in November 2022. It has been making waves in the community ever since, gaining attention from users and experts alike.

In many industries, ChatGPT is helping augment roles and making employees more productive. Marketers can generate content faster, software engineers can start with boilerplate functions, and managers can receive automated summaries of reports.

ChatGPT has received a lot of positive feedback for its overall performance compared to previous LLMs, however others remain skeptical. There is still a lot of confusion about the impact ChatGPT is expected to have on the AI industry, whether or not it will make some jobs redundant, and how it is going to change the way we work.

Who built ChatGPT?

ChatGPT was built by the artificial intelligence research laboratory OpenAI. Founded in 2015, the lab started with the goal of building "artificial general intelligence," or AGI, software that possesses a level of intelligence comparable to that of human beings.

ChatGPT decoded

A comprehensive overview of large language models

OpenAI has made significant contributions to the field of AI, a few examples include:

- In 2020 OpenAI launched **GPT-3** (Generative Pre-trained Transformer), a large language model that set new benchmarks for natural language processing.
- In 2021 OpenAI released **Codex**, an AI system that translates language to code and became the backbone of the popular AI pair programmer tool, **GitHub Copilot** which offers autocomplete-style suggestions as you code.
- In 2022, OpenAI launched **Dalle-2**, a deep-learning model that generates high-resolution realistic images based on text instructions. This led to a significant mass interest in the field of Generative AI.

ChatGPT is one of the latest models developed by OpenAI, built as a variant of the GPT 3.5 model that is fine-tuned to better answer questions and interact conversationally.

The distinction between ChatGPT and other LLMs

ChatGPT is a specific large language model developed by OpenAI, designed to generate human-like responses to text prompts. While ChatGPT is a popular example of a large language model, it is not the only one, and there are several other large language models available in the market, each with its unique features and capabilities.

The history of natural language processing can be traced back to the early days of machine learning, however the development of large language models really took off in the last decade. To name a few examples:

- In 2012, researchers at the University of Toronto and Google developed the first neural language model, called Word2Vec, which was able to learn word embeddings that could capture the semantic relationships between words. This was a major breakthrough in the field and it paved the way for the development of larger and more complex language models.
- BERT, a large pre-trained language model, was developed by Google in 2018. BERT was designed to improve upon the limitations of previous language models, particularly their inability to understand the context and meaning of words in a sentence. BERT

has achieved state-of-the-art results on many NLP benchmarks, and it has become one of the most widely used NLP models and has been adopted by many companies and research institutions for a wide range of applications.

- The same year, researchers at OpenAI developed the first GPT (Generative Pre-trained Transformer) model, which was able to generate human-like text and perform a wide range of NLP tasks with high accuracy.

The primary difference between ChatGPT and other large language models lies in their architectures, training data, and objectives. For instance, some large language models may be designed to perform specific tasks, such as sentiment analysis, language translation, or speech recognition, while others may have a broader range of capabilities, such as generating coherent text or holding a conversation. Furthermore, large language models can differ in terms of their training data and the size of their model parameters. The larger the training dataset and model size, the better the model's ability to understand language and generate human-like responses.

It is essential to consider the specific requirements (i.e. performance) and use cases when selecting a large language model to ensure it meets the necessary criteria for successful integration.

Why is ChatGPT more popular than previous models?

ChatGPT is built on top of the GPT 3.5 model fine tuned to interact in a conversational manner. The dialogue format makes it possible for ChatGPT to answer follow up questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests. Here are a few reasons why ChatGPT has become more popular than previous models:

Conversational style

The dataset that ChatGPT was trained on includes a diverse range of conversational text, and it has been fine-tuned on several specific tasks such as question answering and dialogue generation. This allows it to perform certain tasks more accurately than previous models, such as dialogue generation, question answering, and language translation. The wide range of language patterns and styles in the training data also makes it better equipped for generating diverse and nuanced responses.

ChatGPT decoded

A comprehensive overview of large language models

Interface

Previous models didn't have an interface, so they were mainly used by developers. ChatGPT's free, user-friendly interface makes it highly versatile and adaptable to different use cases and caters to a non-technical audience such as content creators, marketing executives, etc.

Human-like responses

Compared to other large language models such as GPT-3, ChatGPT is much better at generating human-like responses. This was made possible by the use of reinforcement learning from human feedback (RLHF). This technical approach enabled the OpenAI team to rank the different responses generated by the initial model. The model was then refined to learn from the human rankings and choose the best human-like response.

Safety guardrails

In order to achieve mass adoption, it is critical to build responsible AI systems. OpenAI has developed several guardrails to build a more safe and useful language model including substantial reductions in harmful and untruthful outputs by incorporating human feedback in the model development process. Even though ChatGPT has become more robust over time thanks to the feedback loop, it is, in fact, very easy to find examples where it has provided incorrect or misleading results in the past, demonstrating that the guardrails are superficial and can be circumvented.

Contextual awareness

OpenAI has designed ChatGPT to be capable of processing very long sequences of text, allowing it to effectively capture long-term dependencies in the input data. This helps ChatGPT maintain context and track the flow of a conversation, allowing it to provide appropriate responses even in complex or multi-turn conversations.

The challenges of developing LLMs

LLMs have revolutionized the field of natural language processing (NLP) and have led to significant advancements in areas such as machine translation, sentiment analysis, and question-answering. However, the development and deployment of LLMs also come with several challenges, some of which are listed below:

Data Quality and Quantity

The performance of LLMs is heavily dependent on the quality and quantity of training data used. Generating high-quality training data is often expensive and time-

consuming, and there may be difficulties in obtaining data that is diverse and representative of the target population.

Model Size

LLMs can be very large in size, often requiring several gigabytes or even terabytes of storage. This can make deployment challenging, especially in applications where storage is limited or expensive.

Inference Time

LLMs can be computationally expensive to run, which can lead to slow inference times. This is particularly problematic in applications that require real-time or near real-time response.

Computing Resources

LLMs require enormous amounts of computing resources to train, and this can be a significant barrier to entry for many researchers and organizations. Training an LLM can take days, weeks, or even months, and requires access to powerful hardware and specialized software.

Interpretability

One of the challenges in deploying LLMs is their lack of interpretability. LLMs can be difficult to understand and interpret, making it challenging to identify and correct errors or biases in their output. This is especially problematic in applications where the consequences of incorrect or biased output can be significant.

Transfer Learning

LLMs trained for one specific task may not be easily adapted to other tasks or domains. This is because LLMs are often trained using task-specific data, and the knowledge gained from one task may not transfer well to another.

Customisation

LLMs trained on generic datasets may not perform well in specific applications or domains. Customising an LLM for a specific application or domain requires additional training and fine-tuning, which can be time-consuming and resource-intensive.

Overall, deploying LLMs presents several challenges that need to be addressed to realise their full potential. Addressing these challenges will require a collaborative effort between researchers, software developers, and industry stakeholders. As demonstrated above, developing a custom LLM can be challenging and time-consuming, requiring significant resources and

ChatGPT decoded

A comprehensive overview of large language models

expertise to train and fine-tune the model effectively. Furthermore, developing a custom LLM may not always yield better results than using a pre-trained model, especially if the pre-trained model is already well-suited to the task at hand. Therefore, leveraging pre-trained language models like the GPT family is a practical and effective approach to natural language processing tasks.

Is it possible to replicate ChatGPT?

ChatGPT is a highly advanced language model that has achieved impressive results in natural language processing tasks. The question is whether this level of performance is a breakthrough that cannot be replicated or if it is simply a matter of having sufficient computational resources and data.

ChatGPT is built on the underlying technology of **deep learning** and **transformer models**. In deep learning, large neural networks are trained on vast amounts of data to perform a wide range of tasks, such as image classification, speech recognition, and natural language processing. Transformer models are a specific type of deep neural network architecture widely used in NLP tasks. The ChatGPT model was fine-tuned using two types of machine learning techniques called **supervised learning** as well as **reinforcement learning**, both of which used **human trainers** to improve the model's performance.

While the individual components of the underlying technology are well understood, the end-to-end architecture and training process is quite complex, making the replication of the system a formidable technical challenge requiring significant investment. Replicating ChatGPT's performance and accuracy would likely require access to a **large amount of high-quality training data**, a **significant amount of computational power**, and **human resources** both on the development side and for fine-tuning purposes. It is not publicly disclosed how much data and money was used to train the ChatGPT model, however it is clear that OpenAI has invested significant resources into developing the training techniques and optimisations used to build ChatGPT, so replicating its performance would likely be a challenging, time- and resource-intensive task.

What are the alternatives to using LLMs?

Natural Language Processing (NLP) is a rapidly growing field that has seen tremendous advancements in recent years. With the increasing availability of large-scale datasets and computational resources, deep learning models like the GPT family have the potential to become the de facto standard for many NLP tasks. Despite the popularity of LLMs, there exist various alternatives that come with their own set of benefits and drawbacks.

Transformers

are a type of neural network architecture that has gained popularity in recent years for their ability to handle sequential data, including processing long pieces of text, which makes them useful for tasks like translation and answering questions. One popular transformer architecture is the BERT model, which has been used for a variety of NLP tasks, including sentiment analysis, named entity recognition, and question answering. One of the primary challenges of transformers is their training requirements. Training a language transformer requires large amounts of data and computational power, which can be costly and time-consuming.

Recurrent Neural Networks (RNNs)

are another type of neural network that can process sequential data and are well-suited for natural language processing tasks. RNNs use a recurrent connection to allow information to persist through time, making them suitable for tasks that require processing context-dependent input. However, RNNs have limitations such as difficulty in capturing long-term dependencies, and high computational requirements due to sequential processing of data, which can result in slow training times and difficulty in handling large datasets.

Convolutional Neural Networks (CNNs)

have traditionally been used for computer vision tasks but have recently been applied to natural language processing tasks with promising results. CNNs use filters to extract features from input data, making them well-suited for tasks that require local context information, such as text classification.

In addition, many cloud service providers also offer their own **APIs for NLP tasks**. For example, Amazon Web Services (AWS) offers several services, such as Amazon Comprehend and Amazon Lex, that can be used for text analysis and chatbot development, respectively. Google Cloud also offers a range of NLP services, including

ChatGPT decoded

A comprehensive overview of large language models

natural language understanding, entity recognition, and sentiment analysis. Microsoft Azure provides several NLP services as well, including language understanding, sentiment analysis, and text analytics. Using cloud service APIs can offer several advantages, including ease of use, scalability, and reduced development time and costs. However, there are also potential disadvantages, such as possible limitations on customisation and control over the models being used.

In conclusion, there are several alternatives to LLMs available that offer unique advantages and disadvantages. The choice of which model to use depends on the specific task at hand, the available data, and the resources and expertise available to develop and train the model. As the field of NLP continues to evolve, it is likely that new models and techniques will emerge, offering even more options for natural language processing applications.

How robust is ChatGPT, really?

The information that ChatGPT returns might be highly creative, or just plain wrong, so it shouldn't be used for decision making without a human in the loop. While the AI model is improved compared to previous versions, and it's continuously learning and becoming more robust, it doesn't understand cause and effect. It can't reason like a person, but it can imitate human intelligence. For example, it can produce notation for and analyse a chess game in an incredibly convincing manner, but will rarely produce sensible or legal chess moves throughout, as it does not know how to play chess. It only knows how to create life-like chess notation and analysis. ChatGPT is a great tool to use for ideation and creative processes, however it is far from achieving artificial general intelligence, and critical thinking remains a human's job. The output of ChatGPT is not a reliable source of factual information, and it shouldn't be used in decision making processes without human supervision and validation.

Ethical considerations

ChatGPT has exciting current and future use cases but without adequate guardrails we run the risk of delaying or limiting our ability to leverage them to their maximum potential. Here are a few things to consider:

Misinformation

While it may be difficult to replicate ChatGPT's architecture, fine-tuning a large language model to add personality and generate large volumes of content is

relatively easy, e.g. using it to create politically charged, biased content. However, this ease of use and flexibility can also be a double-edged sword. In the wrong hands, it can be used to spread misinformation rapidly. It is essential to ensure that safeguards are in place to prevent their misuse and that users are educated on the potential risks and ethical implications of using them.

Data security and privacy

Studies have shown that large models such as ChatGPT are more susceptible to privacy intrusion issues where personally identifiable information (PII) about individuals used in training data can be extracted using certain prompts or code. Implementing the technology into business operations requires careful consideration of data security and privacy issues. Ensuring the protection of sensitive information and customer privacy must be a top priority, implementing guardrails to mitigate any potential risks.

Fairness and Inclusiveness

Internet scale systems come with internet scale bias. New technologies can have unintended consequences that disproportionately harm minority groups, such as higher error rates in facial recognition and perpetuating bias in algorithms. The digital divide can also prevent minority groups from accessing the benefits of technological advancements. Therefore, responsible and equitable development and deployment of new technologies is necessary. Although ChatGPT uses a Moderation API, which aims to block certain types of unsafe content, it does little to mitigate the propagation of unfairness and bias in the system.

Application transparency

Transparency around how outputs are used, and reliability on outputs themselves, will be essential especially in sensitive sectors such as healthcare and financial services. Although it may not be legally required to do so, if you are looking to use LLMs to augment the provision of any service or product, it is important to be transparent about its use in the interest of creating trust and protecting brand reputation.

ChatGPT decoded

A comprehensive overview of large language models

How can your organisation leverage ChatGPT and LLMs?

Following the examples of Microsoft and Google, enterprise organisations can adopt ChatGPT and LLMs and integrate it with their own processes and applications to unlock new use cases, and set up new, more efficient workflows that augment a team's capabilities. Here are a few examples of how these models can be integrated in organisations:

- 1. Support conversational AI pipelines and customer service workflows:**

LLMs can enhance the capability of chatbots and conversational agents. For example, it can be integrated with a chatbot of a travel agency to provide recommendations for a holiday destination, or to analyse customer complaints and deliver personalised messages and resolution at a financial institution.
- 2. Supercharge knowledge management and data curation:**

LLMs can be integrated with existing applications (e.g., Slack, G-drive, Confluence) to enhance the question-answering capabilities, data curation and insights creation processes of different teams and departments. For example, the sales team can streamline data extraction and curation processes for lead generation purposes, based on the analysis of previous sales content located in G-drive.
- 3. Accelerate content creation and customer personalisation:**

Marketing teams can rapidly generate content for copywriting, blogs, presentations. There are multiple aspects that make ChatGPT attractive for content creation:

 - a. The ability to add personality and style to the output. For example, use business language in a work document, write a science-fiction story for 8 year olds, or write rap in the style of Shakespeare.
 - b. One can add further constraints, for example by specifying the length, e.g., write 3 paragraphs about a topic.
 - c. Develop the story, for example ask ChatGPT to explain a certain concept in more detail, suggest an alternative ending to the story etc.
 - d. By combining image generation AI tools with text generation tools (such as ChatGPT) it is possible to create highly-engaging landing pages very quickly.
- 4. Aid customer acquisition and email management:**

There are a few different ways ChatGPT can be used to assist with the inflow and outflow of email correspondence. ChatGPT can perform tasks such as email classification, summarisation, responding, and search. It can classify emails into categories, generate concise summaries, and produce automated replies, overall greatly improving the efficiency of email management.
- 5. Enhance team productivity and creativity:**

ChatGPT empowers developers with a coding assistant. It is able to write code in different languages, debug and simplify code snippets, create data pipelines, and it has been used to write the code base for entire apps, with little modification needed to create a working prototype. Managers can use it as a planning and organisational tool. ChatGPT can be used to create to-do lists, generate meeting agendas, compile workshop notes, write status reports, etc.
- 6. Enable new product capabilities:**

ChatGPT's popularity has sparked a massive interest in consumer AI and compelled businesses to think about how they can integrate AI into their products. Tools like ChatGPT can empower new product features like call summarisation in video recordings, sentiment analysis of customer service interactions, smart suggestions and personalised recommendations in websites.

It is important to note that most of these options are not readily available. The end goal is to use LLMs for more complex tasks, for example "extract customer data from our warehouse, perform fraud detection on it, and send the insights to a dashboard", however the technology is not there just yet. The integration of conversational AI systems with current enterprise tools, products and processes **will require development and extensive testing**. For example, any institution who integrated LLMs into their chatbot would open themselves up to serious reputational risk if it gave incorrect information or was prompted to say something offensive or harmful. Such solutions will require built-in guardrails in order to avoid biased decision making, and to achieve transparent, explainable results.

When it comes to the enhancement of AI pipelines, LLMs will require access to domain data. Public-facing applications built upon advanced language models like ChatGPT **must prioritise data security and privacy considerations**. Cyber attacks are a real risk and can have catastrophic consequences such as identity theft, financial fraud, or even espionage.

ChatGPT decoded

A comprehensive overview of large language models

Organisations should implement robust security protocols, regularly monitor and test vulnerabilities, and educate users on the importance of data security. By taking these steps, organisations can minimize the risks associated with using LLM-based applications and ensure that their users' data remains safe and secure.

What business value can organisations derive from LLMs?

Large language models like ChatGPT have the potential to provide significant business value for organisations across industries. By leveraging the power of natural language processing and machine learning, businesses can improve customer experience, increase efficiency, and gain a competitive advantage.

Improved customer experience

One of the primary benefits of LLMs is improved customer experience. Chatbots powered by these models can provide accurate and timely responses to customer queries, leading to higher levels of satisfaction and loyalty. LLMs can also be used to provide personalised recommendations and marketing messages, improving conversion rates and customer retention.

Increased efficiency

In addition to enhancing the customer experience, large language models can increase efficiency by automating repetitive tasks and reducing the need for human resources. This can result in cost savings for businesses and enable employees to focus on more complex tasks.

Competitive advantage

Furthermore, businesses that adopt LLMs early on can gain a competitive edge in their market. By providing a more efficient and personalised customer experience, they can differentiate themselves from their peers and attract and retain customers and talent.

Improved insights

Finally, large language models can be used to analyze customer data and provide insights into customer behaviour and preferences, which can inform business strategy and decision-making.

Overall, LLMs have the potential to transform the way businesses operate and interact with customers, providing significant value in terms of improved customer experience, increased efficiency, and competitive advantage.

How can different industries use LLMs?

LLMs like ChatGPT can be leveraged by various industries in numerous ways to improve their operations and customer experiences.

Financial institutions

Chatbots powered by LLMs can be used to provide customers with financial advice and guidance, answer common questions about financial products, and help customers manage their accounts. Furthermore, LLMs can be used to analyze data from various sources, detect fraudulent activity, and generate compliance reports. They can help financial institutions to comply with regulatory requirements and report any suspicious activities to regulatory authorities.

Retail

Retailers can use LLMs and AI-powered tools to analyze customer data and other relevant information, gain valuable insights into market trends, optimise their operations, and to power chatbots that can provide personalised recommendations and assist customers in making purchases. This can help retailers to increase sales, improve customer engagement, reduce costs, and stay competitive in the market.

Healthcare

The healthcare industry can benefit from creating customer-facing applications powered by LLMs to triage patients, answer common health-related questions, and provide personalised healthcare recommendations. They can also be used to improve the patient experience by providing real-time updates on appointments and treatment plans. LLMs can be adapted for analyzing large amounts of medical data such as medical records, clinical trials, and research papers. This can help researchers to identify new treatment options, improve diagnosis accuracy, and enhance patient outcomes.

Education

LLMs can be used to power chatbots that can answer students' questions and provide personalised learning recommendations. They can also be used to improve teacher efficiency by automating administrative tasks, such as grading and scheduling.

Travel

Chatbots powered by LLMs can be used to provide travelers with real-time updates on flight schedules and delays, assist with booking travel accommodations, and provide personalised travel recommendations based on customers' preferences.

ChatGPT decoded

A comprehensive overview of large language models

Recent developments

ChatGPT has seen a significant volume of traffic since it was introduced to the public. OpenAI responded to the high demand by releasing ChatGPT Plus, an introductory subscription plan for an improved conversational AI that engages in dialogue with the user. The plan provides general access to ChatGPT even during peak hours, it has faster response times, and priority access to new features and improvements. A free version that can be utilised during off-peak hours is still available.

The significant media buzz and excitement within the technology industry has led major tech companies to announce their own conversational AI systems in response to the growing interest.

Microsoft announced a substantial investment of \$10 billion in OpenAI in January 2023. More recently, they unveiled their latest conversational AI solution, the Bing chatbot. Unlike ChatGPT, which can only retrieve information up until 2022 based on the data it was trained on, "the new Bing" is able to retrieve information about recent news and events. The AI is built upon the Prometheus model, an upgraded version of the GPT 3.5 language model that powers ChatGPT.

Google is a prominent player in the field of AI and has demonstrated a strong commitment to advancing research and development in this area. In Feb 2023 they announced their own conversational AI system, BARD, with an underlying large language model called LaMDA (Language Model for Dialogue Applications). Shortly after its launch, the bot provided an incorrect response during a promotional experiment on Twitter, resulting in a \$100 billion loss in market value for BARD's parent company. In March, the release of the PaLM API was announced, allowing developers access to Google's large language models.

Meta also joined the race at the end of February, releasing LLaMA (Large Language Model Meta AI), a smaller, more performant model compared to ChatGPT. They intend to grant access to users on a case-by-case basis, with the goal of advancing research in the AI subfield of natural language processing, and addressing the risks of bias and harmful aspects of large language models.

For the first few months the ChatGPT model was only accessible via the interface, while users could access a whole family of GPT models via publicly available APIs. As of March 2023, OpenAI has released an API for ChatGPT as well. Thanks to system-wide optimisations,

the ChatGPT model family is 10x cheaper than the previous GPT-3.5 models.

In mid-March, OpenAI announced their latest breakthrough - the GPT-4 model. Boasting improved factuality, reliability, advanced reasoning capabilities, and adherence to guardrails, GPT-4 is able to handle more complex conversational tasks compared to ChatGPT. The new model is versatile and can accept images as input as well as text, opening up exciting new possibilities for users to propose both vision and language tasks to the model. Although GPT-4 is better at providing factual responses, it still faces limitations with hallucination and creative responses. Notably, OpenAI claims that the model's ability to respond to requests for disallowed content is 82% better than GPT-3.5.

The rapid developments in this field over the last few months have been truly remarkable, characterised by an unprecedented race among key players, and as a result, it can be challenging to remain up-to-date with the latest advancements.

The emergence of prompt engineering

Prompt engineering refers to the process of designing and refining the prompts used to guide a generative AI system in producing its outputs. Prompt engineering is crucial for achieving high-quality results in generative AI applications, particularly in language and image models. However, it can also be a challenging and time-consuming task, requiring careful consideration of the desired outcome, potential biases, and ethical implications of the generated output.

Prompt engineering is experiencing a surge in popularity due to the increasing demand for generative AI applications. Some creators are already offering their prompts on marketplaces like PromptBase. The emerging trend of prompt engineering is a sign of the ever-evolving nature of technology. The future may even involve writing code through instructions, completely redefining what it means to be an engineer in the technology space.

However, there is concern that the rising emphasis on prompt engineering could lead people to overestimate the technical rigor and reliability of the results that can be obtained from a constantly evolving black box. This is because language models have an extensive knowledge and memory of the training data, which can cause them to start 'hallucinating' based on prior knowledge, unless explicitly instructed not to do so.

ChatGPT decoded

A comprehensive overview of large language models

Crafting the appropriate prompt requires meticulous exploration of the multiverse of fictional possibilities, sculpting and refining the space of those possibilities until only the desired text remains. A critical aspect of this process is figuring out why and when the AI produces inaccurate results, as these systems do not come with bug reports and their outputs can be unpredictable.

As the field of prompt engineering continues to evolve, new strategies and techniques for creating effective prompts may become necessary to keep pace with emerging trends and challenges. Despite the current limitations, it is clear that the potential benefits of these technologies are vast and far-reaching, and as more and more resources become available to prompt creators, the possibilities for innovative and transformative applications will only continue to expand.

What impact will ChatGPT have on industry and jobs?

ChatGPT's ability to streamline and boost the creative process will likely result in disruptions in certain industries. The copywriting and content creation industries might see the quickest gains in the short term. The effects are already starting to manifest as digital media company BuzzFeed declared their plans to integrate AI-driven content into their core business operations. The company plans to reduce their workforce by 12% and predicts that AI will play a crucial role in streamlining the creative process through the creation and customisation of content over the next decade.

While it should not entirely replace human labour, ChatGPT streamlines the creative process and thus presents an opportunity for innovation. The integration of LLM capabilities into current applications or workflows and the development of new products and processes can be anticipated in this space. The recent surge of prompt engineering and the buzz it's generating are already signs of these changes taking effect.

The ChatGPT model has already seen improvements in terms of its accuracy over the last few months, and it is expected to continuously learn and become more robust over time. The more recently released GPT-4 model is rumoured to have 100 trillion parameters, compared to the 175 billion parameters of GPT-3. Having a larger number of parameters provides the model with a larger capacity to learn from data, allowing it to capture more

intricate relationships in the data, and producing more accurate predictions. It can be challenging for people to fully understand the extent of the increase, however, we have provided a visual representation of the projected growth in Fig. 1.

Large language models are here to stay, and they are expected to become bigger and better. As they become more complex and larger in size, training them requires significantly more compute power and electricity usage. This increase in computational demand not only requires more powerful hardware but also comes at a cost in terms of energy consumption and environmental impact.

So the question remains: how can humanity adapt to live with AI, and how can we make the best decisions around it?

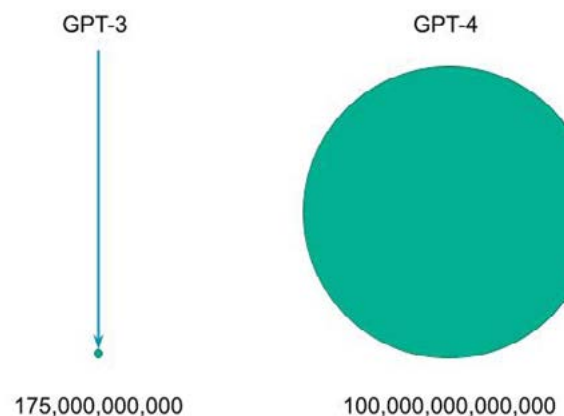


Fig. 1 A comparison between GPT-3 and projected parameter numbers for GPT-4