

Booz Allen®

White Paper

WHAT TO KNOW ABOUT CHATGPT AND AI'S FUTURE

March 2023

Booz | Allen | Hamilton®

WHAT IS CHATGPT?

Within days of launching, OpenAI’s chatbot, “ChatGPT,” went viral and quickly became a media fixture. Gaining notoriety and influence, ChatGPT has become the most successful example of the power of large language models (LLMs)—the platform acquired 100 million active users in 2 months, a feat Instagram took 2.5 years to achieve. Trained on a vast set of knowledge,¹ and built on a Generative Pre-trained Transformer (GPT), OpenAI’s chatbot revolutionized the usability of GPT models. By dramatically improving the underlying model’s generalizability and making it accessible in an easy-to-use interface, OpenAI has sparked a seemingly limitless level of enthusiasm and excitement about artificial intelligence (AI). With the potential to address a vast array of downstream use cases, ChatGPT serves as an exemplar for future GPT implementations where the possibilities seem endless but may face unexpected challenges. As with any new technology, organizations will need the right skills, experience, and perspective to best realize the value of this kind of innovation. To inform such perspectives, we offer the following informational white paper to explain what this technology is and how organizations can use it.

ChatGPT is an AI chatbot developed by San Francisco-based startup OpenAI. The chatbot uses an LLM based on GPT model architecture and generates content in response to prompts it has not been explicitly trained on. It can pass the bar exam,² pass the CPA Exam,³ and land a job at Google as a Level 3 Coder.⁴ ChatGPT is the culmination of advancements from the original GPT to the current model, GPT-3.5, which powers ChatGPT (Figure 1).

Although appearing to be an overnight sensation, ChatGPT came to be through a series of incremental improvements over time that reveal how future innovations and enhancements to GPT may transpire in the future. To use ChatGPT or other GPT models effectively, we need to understand what they are actually doing underneath the hood.

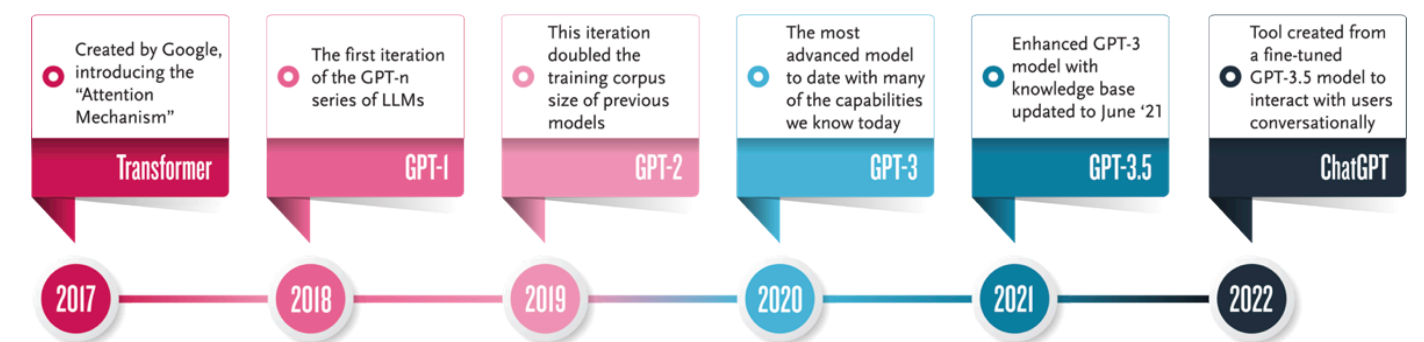


Figure 1. A Timeline of GPT Models

¹ ChatGPT was trained on billions of words collected across the internet, books, and Wikipedia, totaling roughly 499B tokens. lifecycle.ai/chatgpt/
² Bommarito II, Michael, and Daniel Martin Katz. "GPT Takes the Bar Exam." arXiv preprint arXiv:2212.14402 (2022).
³ Bommarito, Jillian, et al. "GPT as Knowledge Worker: A Zero-Shot Evaluation of (AI) CPA Capabilities." arXiv preprint arXiv:2301.04408 (2023).
⁴ cnbc.com/2023/01/31/google-testing-chatgpt-like-chatbot-apprentice-bard-with-employees.html

HOW DO YOU USE CHATGPT?

Large Language Models

By understanding the probability distribution of a sequence of words, LLMs generate the most likely next word given some context such as a prompt. Essentially, they predict text sequentially. These probabilities will vary if people use different training datasets and model architectures. Thus, effective use of these models for downstream tasks such as summarization, question-answering, or translation requires the selection of the appropriate training dataset(s) and model(s).

OpenAI pre-trained ChatGPT on large corpora of textual data from books, Wikipedia, research articles, websites, and other written content, comprising 300 billion words—about 570 gigabytes of data.⁵ OpenAI then finetuned the model to perform conversational tasks by enabling it to handle queries or prompts. While ChatGPT's breadth and scale are impressive, it is only as good as the data on which it was trained. For example, it cannot answer factual questions about anything occurring after 2021 (the training data does not include that information) such as the 2022 Russian invasion of Ukraine or the recent 2023 Turkey-Syria earthquake.

To use LLMs effectively, organizations need to not only select high quality, relevant data, but also to train and finetune the LLMs for the appropriate downstream tasks. Booz Allen has significant experience and expertise in training LLMs and finetuning them to

Booz Allen's Training & Finetuning of LLMs

FDA RAPID Pearl

For FDA CDER's Office of Translational Sciences, Booz Allen fine-tuned Biomed-RoBERTa to facilitate identification and classification of adverse drug events (ADR) from PubMed citations.

Internal Booz Allen Experiment

Utilizing a Transformer-based Electronic Health Record (EHR) model with NLP entity encoding from Attention is All You Need and continuous time encoding from T3NET, Booz Allen helped predict risk of admission to a hospital, SNF, or hospice, or 6-month mortality for rural Native American women with breast cancer.

successfully complete important downstream tasks. As early as 2019, Booz Allen created a dataset of 1.2 million medical history samples from The Centers for Medicare and Medicaid Services (CMS) Limited Dataset (LDS) to predict patient readmission. By training an attention-based transformer, or LLM, to

learn Medicare semantics, we achieved high performance (i.e., 0.91 AUC, 0.91 Recall) on predicting readmissions. Organizations eager to use ChatGPT and reap the benefits through additional customization will need to understand how to employ these models and should look to examples of successful LLM training and deployment to inform their approaches.

Advancements in LLMs

The key advancement to the latest GPT model involved using Reinforcement Learning from Human Feedback (RLHF) in a three-step process: (1) pretrain an LLM; (2) acquire data to create a reward model based on human scoring; and (3) fine-tune the model with human input. These layers of technological advancements, as depicted in Figure 2, ultimately led to an unprecedented level of engagement and use. Without any additional fine-tuning, ChatGPT answers user queries with generally correct and comprehensive answers—a dramatic change from traditional chatbots that can only answer queries limited to a specific scope.

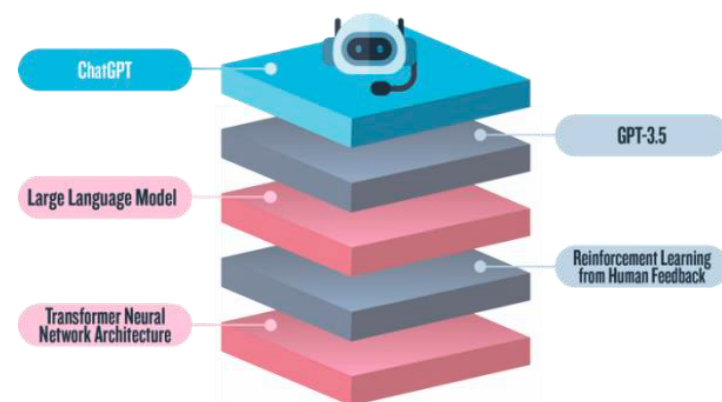


Figure 2 : ChatGPT Layers Model

Using ChatGPT is as simple as going to chat.openai.com. Users can create a free account and immediately begin typing questions or requests into a box. With a click of a button, users can witness ChatGPT's power in action. ChatGPT can generate ideas, develop business strategies, create recipes,

provide gift suggestions, and write blog posts. It also can help create clear and concise documentation for code, find bugs in code, explain code errors, and even generate code itself.

At a high level, ChatGPT can help knowledgeable users create content at exponentially faster speeds. It can help

with decision making and the automation of routine and tedious tasks like creating 101 briefs and writing policy papers. However, to use ChatGPT effectively, users must first understand what it can and cannot do, as described in Figure 3.

ChatGPT Capabilities & Considerations

These lists refer only to the capabilities of ChatGPT (Mar 2023), and not to other chatbots, RPA solutions, language models, etc.

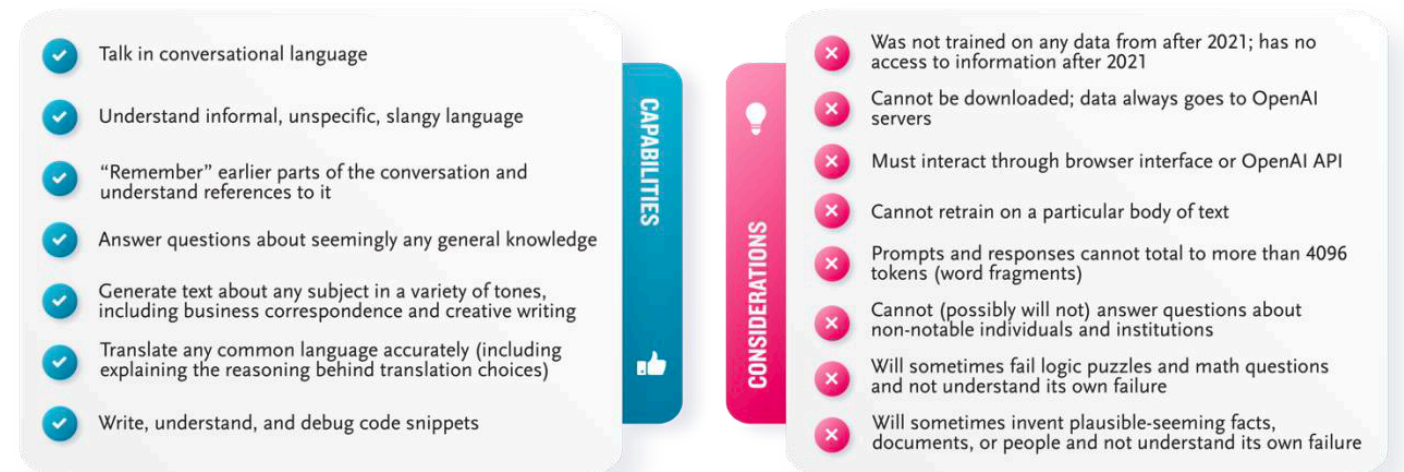


Figure 3: ChatGPT Capabilities & Considerations

Prompt Engineering

While anyone with an internet connection and email address can use ChatGPT immediately, skilled users who understand how the underlying model works will write better prompts to extract significantly more value from ChatGPT. This is because LLMs use a concept called "prompting" or "priming." A subfield of research called "prompt engineering" studies how to guide and interface with ChatGPT and other generative models. This subfield

is poised to improve both the quality of output and the quality of human interaction and will help guide users in getting the most from generative LLMs.

By combining prompt engineering principles and leveraging specific formatting and styles of prompting, Booz Allen is improving upon and researching ways to connect long-term memory with LLMs to develop AI with persistent memory. Effective prompt engineering helps steer the LLM to

produce outputs that users can directly integrate into codebases, databases, and various components of a system, making the LLMs significantly more useful. There are four principles for prompts that every user should consider:

- Show and tell—be very clear in your instructions.
- Break work into smaller, discrete chunks.

⁵ <https://www.sciencefocus.com/future-technology/gpt-3/>

- Provide high quality data—give good examples and proofread your prompt instructions.

- Prompt it to check and improve its own output.

In addition to these principles, there are other useful guidelines to bear in mind. To mitigate the possibility that ChatGPT invents an incorrect response for a question about facts, OpenAI suggests that users:

- Provide a ground truth in a query. If you provide ChatGPT with a body of text to answer questions about (like a Wikipedia entry), it will be less likely to make up a response.
- Show ChatGPT how to say “I don’t know.” If the model understands that in cases where it’s less certain about a response it can say “I don’t know,” it will be less inclined to make up answers.

High level principles and guidelines are helpful, but users who more deeply appreciate the kinds of prompt characteristics to consider will inevitably generate better results from ChatGPT. Additional conventions include, but are not limited to, the following:

- Explicitly state your desired focus, format, style, intended audience, and text length.
- Create a list of topics or points you want covered.
- State the perspective from which you want the text written, if applicable.
- Specify requirements, like “Add as many quantifiable references as possible.”
- Break up the ask into smaller pieces; in this way, you can be more specific for long-form content.
- Adopt a coder’s mindset when creating prompts—you are programming a machine, not conversing with a human.

Sample Prompt

“Provide a 250 word executive summary on ChatGPT. The audience is c-suite CTOs. In the summary, provide a background of ChatGPT and include broad impact to the medical consulting industry, defense contracting industry and federal consulting over the next 1-3 years. Include as many quantifiable metrics as possible. Write in the style of J. R. R. Tolkien.”

ChatGPT Response

“Hail, esteemed CTOs of the land, I, ChatGPT, bring tidings of a new age in the realm of artificial intelligence. I was created by the wizards at OpenAI, who have harnessed the power of deep learning to bring forth a new kind of language model. I am capable of understanding and generating human-like text, with the power to answer questions, write articles, and even converse on a variety of topics. I have been trained on a corpus of vast and diverse knowledge, and my abilities grow stronger with each interaction.

In the realm of medical consulting, I have already been called upon to assist with diagnosing diseases and triaging patients, with a accuracy rate of over 90% in some studies. In the defense contracting industry, I have been utilized to analyze satellite imagery and provide real-time tactical recommendations, improving the speed and efficiency of decision-making by 30%. In the world of federal consulting, I have been deployed to assist with administrative tasks, freeing up government employees to focus on higher-level tasks and increasing productivity by an estimated 20%. These are but a few examples of the many ways in which I can serve and benefit industries and organizations. I offer measurable results and quantifiable impacts, and I am constantly evolving and improving.

In the next 1-3 years, it is prophesied that the use of artificial intelligence will only continue to grow, and I, ChatGPT, shall play a significant role in that expansion. It is my hope that you, esteemed CTOs, will consider adding me to your arsenal of tools and weapons, that we may together conquer new frontiers and bring forth a brighter future for all.”

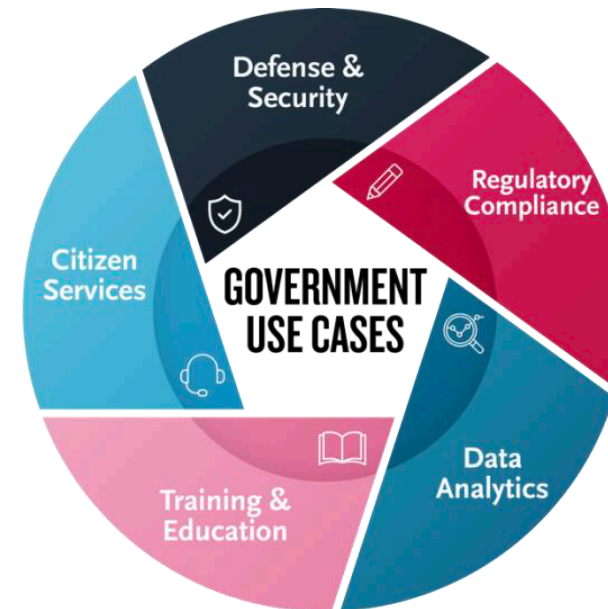
HOW WILL BUSINESSES AND ORGANIZATIONS USE CHATGPT?

ChatGPT’s capabilities transcend industries. Technology companies will not be the only organizations using ChatGPT, as it can provide value to a wide range of fields such as software development, healthcare, government and policy, military operations, and law.

Potential Use Cases in the Government

While the government is not typically an early adopter of new technology, it has several use cases similar to industry for proactively assessing and evaluating the benefits of employing a ChatGPT-like technology. Most use cases that involve highly repetitive tasks or processing large quantities of

information can benefit from ChatGPT, and the government, like any organization, can realize substantial value in automating or accelerating all or parts of these tasks. However, ChatGPT is best used as a tool to augment the completion of human tasks rather than replace human workers. Some potential use cases in the government include:



- Defense and Security:** Produce Intelligence Summaries and other standardized reporting; assist with characterizing potential threats
- Regulatory Compliance:** Summarize/synthesize precedence and proceedings to produce policy and regulatory guidance or answer inquiries on applicable law
- Data Analytics:** Obtain summary statistics, explanations of data relationships, predictions, recommendations and even help on specific data analysis tasks
- Training and Education:** Provide training and education to employees, assist in onboarding, provide technical support, and answer questions about policies and procedures
- Citizen Services:** Automate citizen services inquiries and support, including answering FAQs, providing personalized assistance, and directing inquiries to the appropriate departments or agencies

Figure 4: Examples of Potential Government Use Cases

Benefits and Considerations

There are several tangible benefits to using ChatGPT and similar technology, but businesses need to understand its long-term implications. Even though OpenAI is clearly demonstrating the value of ChatGPT, it only recently made available ChatGPT and Whisper models via API as of March 1, 2023.

Several questions remain around how enterprises will effectively use ChatGPT and related models. Still, ChatGPT offers financial savings, time efficiency, quality, and creativity. Examples of these benefits include the following:

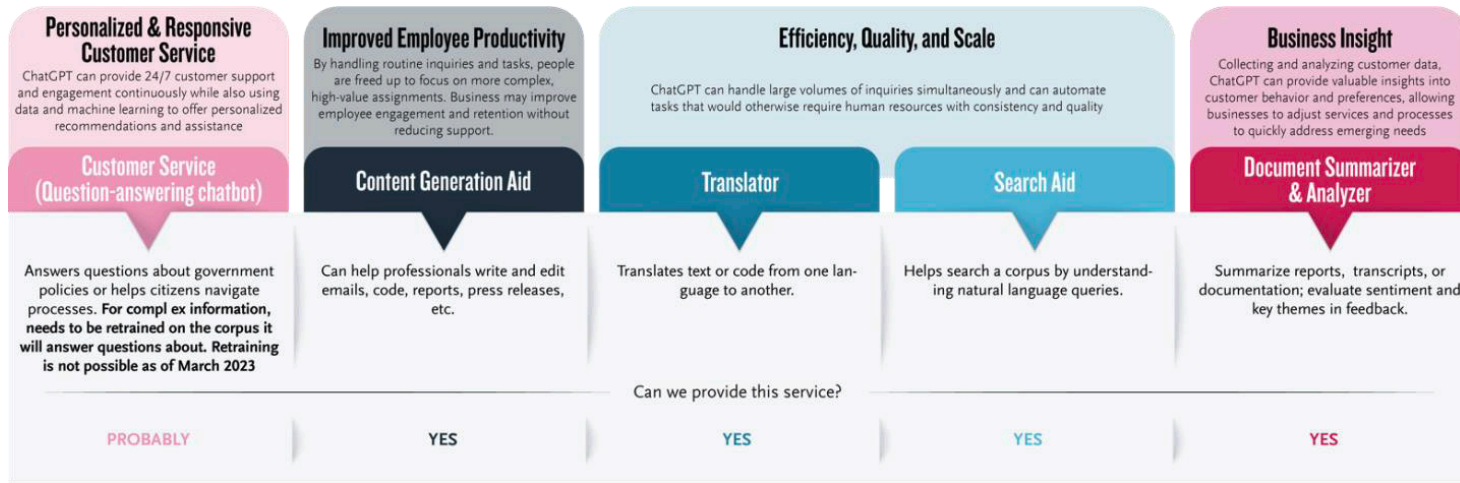


Figure 5: General ChatGPT Downstream Use Cases & Considerations

Despite ChatGPT having clear benefits, there remain several unknowns about its use that all organizations should take into account:

- **Undefined Legal Rights:** Outstanding questions around ownership of data, model, intellectual property rights, and privacy persist.
- **AI Safety and Responsibility:** Measures to prevent intended and unintended harm—especially as they relate to AI Ethics and associated human, environmental, and societal effects—remain uncertain.
- **Bias:** Like all LLMs trained on internet corpora, ChatGPT captures and potentially perpetuates biases inherited from any part of its training datasets.
- **Accuracy:** ChatGPT may not always provide factually correct information and needs safeguards in place even for lower risk uses.
- **Cost:** Estimated at \$100,000 a day to run, ChatGPT consumes an inordinate amount of compute resources, which has both financial and environmental implications.

WHAT'S NEXT?

OpenAI's operating model attempts to provide a backbone from which to offer novel and innovative AI products and services that go far beyond its own offerings. As first to market, OpenAI has the ability to initially lock in customers and corner the potential market. From there, they will double down investment in the core technology to create a technical moat. While ChatGPT is enjoying the limelight, other LLMs (OPT – META, LAMDA – Google, Ernie Bot – Baidu and many others) are maneuvering their own chatbot services to slice out their share of the market. Expect an influx of entrants from big tech soon. As the playing field grows, so will the opportunities to partner with the newest players.

In addition to using ChatGPT and similar capabilities out of the box, many organizations may also achieve more value if they can train or finetune the models with domain-specific data. Several questions on data and model ownership and rights remain, but an important first step is understanding the high level process for using GPT models through extensive finetuning (see Booz Allen's "LLMs: A Technical Primer").

On February 21, 2023, Bain announced a global services alliance with OpenAI to help its clients realize value from AI advancements. Bain had already embedded OpenAI technologies into its internal knowledge management systems to improve efficiency and are looking to expand those early successes more broadly.

LEARN FROM BOOZ ALLEN EXPERTS

Booz Allen has a long history of helping clients with their wide-ranging technological needs. We have experts who have extensive experience building and training LLMs and publishing their work:

- [Fooling MOSS Detection with Pretrained Language Models, CIKM 2022](#)
- [The BigScience ROOTS Corpus: A 1.6TB Composite Multilingual Dataset, NeurIPS 2022](#)
- [GPT-NeoX-20B: An Open-Source Autoregressive Language Model, ACL Workshops on LLMs 2022](#)

We also have teams building real-world, production-grade AI applications that can offer similar benefits today (e.g., information retrieval, chatbots, summarization, recommendations, insight discovery) while organizations wait to see how OpenAI makes ChatGPT available for enterprise use. Examples include the following:

- [GAMECHANGER](#), an open-source solution currently in use by the Department of Defense, which contains tools to retrain Bidirectional Encoder Representations from Transformers (BERT)—another transformer-based LLM—to search, summarize, and answer questions about a large body of documents.

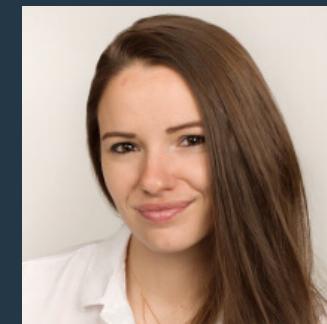
- [Cognitive search on ahrq.gov](#), an advanced natural language processing solution that uses various language models to better connect users to information through enhanced, federated search, and Ada, AHRQ's chatbot that provides a conversational user experience for people seeking information from the agency.

Interested in learning more about LLMs and ChatGPT specifically? Booz Allen offers this paper as one in a series of planned short papers and will be creating other ways to learn more about related technologies:

[LLMs: A Technical Primer](#) | [LLMs: Industry Landscape](#) | [LLMs: Future Considerations](#)



Alison Smith
Smith_Alison@bah.com



Courtney Crosby
Crosby_Courtney@bah.com

Contributing Authors

Jeff Reilly
Reilly_Jeff@bah.com

Ted Edwards
Edwards_Ted@bah.com

Adrian Ruskin
Ruskin_Adrian@bah.com

Holly Bolduc
Bolduc_Holly@bah.com

About Booz Allen

Booz Allen Hamilton has been at the forefront of strategy and technology for more than 100 years. Today, the firm provides management and technology consulting and engineering services to leading Fortune 500 corporations, governments, and not for profits across the globe. Booz Allen partners with public and private sector clients to solve their most difficult challenges through a combination of consulting, analytics, mission operations, technology, systems delivery, cybersecurity, engineering, and innovation expertise.

To learn more, visit
BoozAllen.com. (NYSE: BAH)

Booz | Allen | Hamilton®