LTIMindtree

POV

ChatGPT

An AI NLP Model

Table of Contents

ChatGPT – An AI NLP Model

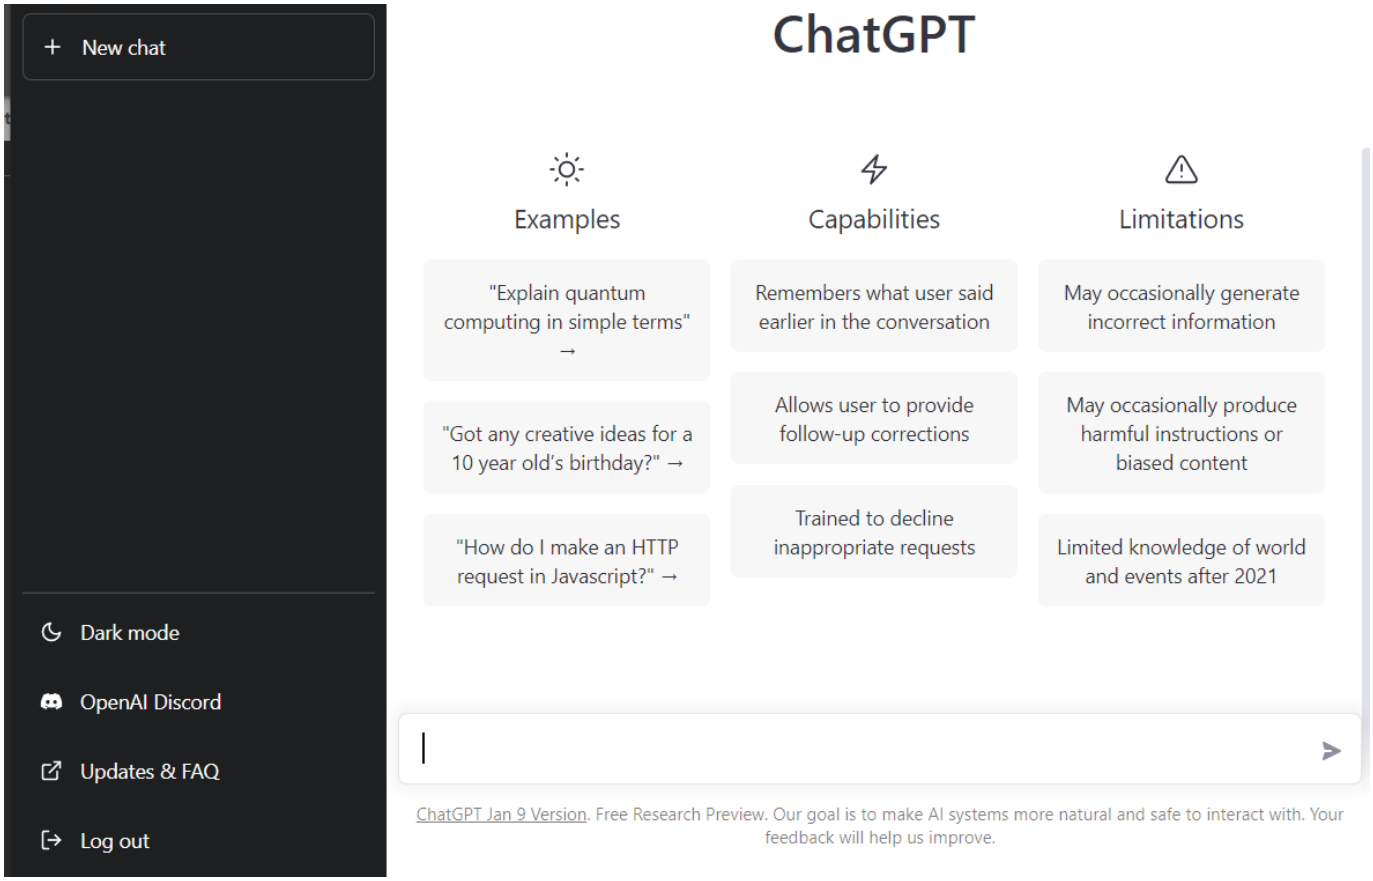

Figure 1: ChatGPT Jan 9 Version. Research Preview

What is ChatGPT?

ChatGPT is a powerful text-generating dialogue system. It is a natural language processing model (NLP) that generates humanlike responses to inputs from users. Based on Generative Pre-trained Transformer architecture, this NLP model is trained on vast conversational data from the internet. Once trained, it can accomplish a variety of NLP tasks such as translation, answering questions, and completing text. The dialogue system can also be used as a conversational AI, for use in chatbots, virtual agents, and other conversational applications.

Who created ChatGPT?

OpenAI LP, a capped-profit research and deployment company, governed by the OpenAI non-profit board, released ChatGPT on November 30, 2022. OpenAI is an artificial intelligence research laboratory founded by Elon Musk, Sam Altman, Greg Brockman, Ilya Sutskever, Wojciech Zaremba, and John Schulman. The venture is supported by investors including Microsoft, Reid Hoffman's charitable foundation, and Khosla Ventures.

How does ChatGPT work?

Once the user input is received in ChatGPT, machine learning algorithms are applied to a large amount of text to generate a humanlike response.

Chat GPT works on a neural network-based architecture called a Transformer. Datasets from websites, books, and articles are used for training the model on language patterns and structure. It learns to predict the next word based on the previous one. Once trained, the model can generate new text by predicting the next word in a sentence, with a given prompt or context. The process is reiterated until the model has generated a complete sentence or the required number of words.

The model also uses an attention mechanism while generating text. The mechanism allows selective focus on certain parts of the input for more accurate and coherent responses.

When using ChatGPT for conversational AI, the model is typically fine-tuned on a smaller dataset of conversational text to further improve its ability to generate human-like responses.

How is ChatGPT trained?

Copious amounts of conversational exchanges, fine-tuned to a specific task or domain, are used as a dataset to train ChatGPT, using a variant of the transformer architecture for predicting words. Instead of feeding in answers to the model it is allowed to pick up patterns and relationships in the data on its own. ChatGPT is a result of fine-tuning GPT-3.5 using supervised and reinforcement learning.

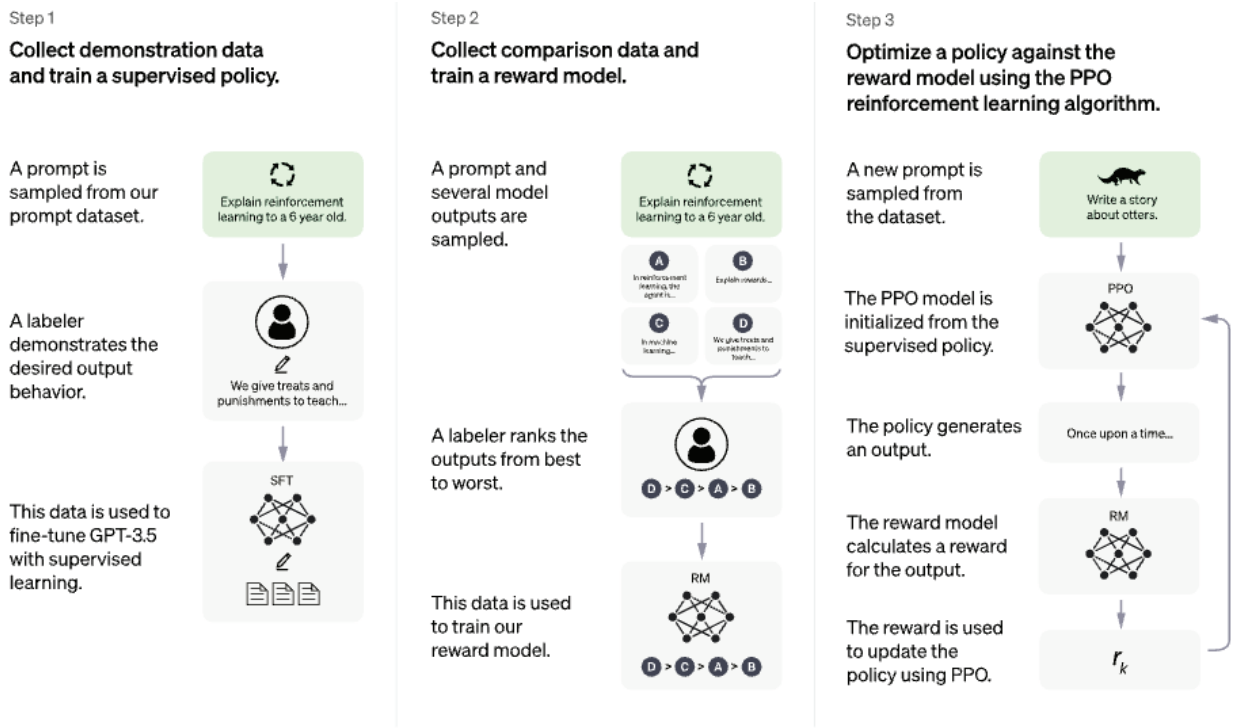

Figure 2: Training method - ChatGPT

ChatGPT, like InstructGPT (its sibling model), is trained using Reinforcement Learning from Human Feedback (RLHF). There is only a slight difference in the data collection setup for ChatGPT. There was supervised fine-tuning involved in training the model, where human AI trainers played both, the user and the AI assistant, in the conversation. The trainers composed their responses on model-written suggestions. The new dialogue dataset was then mixed with the dataset from InstructGPT and transformed into a dialogue format.

Creating a reward model for reinforcement learning requires comparison data, consisting of model responses ranked as per their quality. A selection of conversations between the AI trainers and the chatbots were used to collect this data. A randomly selected message, produced by the model was selected. Several completions of the message were sampled, and AI trainers were asked to rank them. Based on these reward models, ChatGPT was fine-tuned using the Proximal Policy Optimization. Several iterations of this process were performed.

How is ChatGPT different from other Chatbot platforms?

Since the ChatGPT AI model used vast amounts of internet data entered by humans, including conversations, the responses it generates may seem humanlike. It follows a dialogue format, it can ask follow-up questions, accept mistakes, challenge incorrect premises, and reject inappropriate requests. Being a sibling model of the InstructGPT, it follows an instruction in a prompt and gives a detailed response.

The responses, however, are based on the similarity between the output and the dataset used to train them and may be inaccurate and even misleading.

What are the common use cases for ChatGPT in DWP?

There are several ways ChatGPT can be used in the information technology and digital workplace to provide end user services:

- **Virtual Assistants:** ChatGPT can help building virtual assistants for help with scheduling, email management, & customer service.
- **Email Responders:** ChatGPT can be used to generate automated responses to common customer inquiries and forward them to the customer.
- **Knowledge Base:** ChatGPT can be used to build a knowledge base that can answer common questions and provide information to users.
- **IT Service Desk:** ChatGPT can be used to provide automated assistance to users who need help with IT-related issues, such as password resets and account lockouts.
- **HR Assistance:** ChatGPT can be used to provide automated assistance to employees who need help with HR-related issues, such as benefits and time-off requests.

- **Document Automation:** ChatGPT can be used to automatically generate documents, such as contracts and reports, based on a set of predefined templates and prompts.
- **Language Translation:** It can be used to translate messages, emails and even documents.
- **Meeting Summary:** ChatGPT can be used to generate a summary of the main points discussed in a meeting, this can be useful to track progress and make sure everyone is on the same page.

Note that ChatGPT can be used in many other way, but it is important to consider the ethical and legal implications of using AI models like ChatGPT in a professional setting, specially when working with personal data. Also, in the current version (Jan 2023), ChatGPT released in research preview is still evolving and is trained on internet data available up the year 2021. However, for specific Digital Workplace Services customer data, it will need to be trained on a customized AI ML data with specific IT taxonomy, tested and fine-tuned with supervised and re-inforced learning as per specific environment.

What are the common shortfalls of using ChatGPT in professional environment?

Since ChatGPT is an evolving (Jan 2023) AI NLP model, some of the common, potential shortfalls of using ChatGPT in an enterprise include:

- **Quality of Generated Text:** ChatGPT is a language model and although it has been trained on a large dataset, it may not generate text that is of the same quality as text written by a human.
- **Bias in the training data:** GPT-based models like ChatGPT may perpetuate bias present in the training data. Therefore it's important to review and clean the dataset before fine-tuning the model.
- **Lack of Understanding of context:** GPT based models like ChatGPT are good at generating human-like text, but they lack understanding of the context. Therefore, it may generate irrelevant or nonsensical responses in certain situations.
- **Legal and Ethical Implications:** Using ChatGPT in an enterprise service may have legal and ethical implications, such as potential violations of privacy and data protection laws, or issues related to consent and transparency.

- **High computational cost:** GPT-based models like ChatGPT are computationally expensive, they require powerful hardware and infrastructure to run, and fine-tuning a model may also require a large amount of computational resources.
- **Limited ability to handle structured data:** ChatGPT is a language model and it's not designed to handle structured data, if your service requires handling structured data, then it's better to use models that are designed for that purpose.
- **Can be fooled:** GPT-based models like ChatGPT can be easily fooled/manipulated by adversarial examples, such as text inputs that are specifically crafted to mislead the model.
- **Lack of explainability:** GPT-based models like ChatGPT are based on deep learning techniques, which makes it hard to understand how the model comes to a certain conclusion.
- **Lack of SLAs:** Currently there are no SLAs defined for enterprise usage.
- **Security and Legal approval:** As ChatGBT is an Open Source model hence security and legal approvals / agreement for customer data usage could be an issue.
- **Development effort and maintenance:** Since, there are no ChatGBT out-of-box plugins hence a dedicated team and effort is needed to set it up and intagrate via APIs along with any customization for customer environment.
- **Training and Learning:** It uses old internet data prior to 2021 and is not current hence is not updated with current information. Also, for specific customised, enterprise (DWP) training it will need huge amount of data for Re-inforced learning and supervised training which will need continuos development effort.
- **Lack of dedicated / agreed support:** Being an open source model, support specifically for enterprise customer environments will be an issue.

How can one implement ChatGPT?

At the time of writing this document (Jan 2023), a research preview of ChatGPT is being made available so as to learn about its strengths and weaknesses. Data available up to the year 2021 was used for training this model, so it may not have knowledge of current events. Also, currently, since it is not yet connected to the internet, it may provide incorrect responses, instructions or biased content.

Note that currently, there are no official out-of-box plugins to integrate ChatGPT with other tools, as it is primarily used as a standalone model. However, the GPT-3 model that ChatGPT is based on can be integrated with various tools and platforms through the OpenAI API and needs a development effort.

How can enterprise customers plan to purchase ChatGPT?

Once available (post Jan 2023), enterprises can purchase access to the ChatGPT model from OpenAI through the OpenAI GPT-3 API. This API allows integration of the model into applications and systems for performing NLP tasks like language generation, text completion and conversational AI. Additionally, OpenAI also provides an API for the GPT-3 in a multi-language format which can be used for language generation and other NLP tasks.

What are pricing models are available for enterprise use of ChatGPT though OpenAI GPT-3 API?

The pricing model (as of Jan 2023) for enterprise use of ChatGPT through the OpenAI GPT-3 API typically includes the following components:

- **Monthly API calls:** The total number of API calls an enterprise makes to the GPT-3 API in a given month. The more calls made, the higher the cost will be.
- **Monthly token quota:** The total number of tokens that can be generated by the API in a given month. A token is a sequence of characters, words, or sentences generated by the API.
- **Additional features:** Some additional features such as private cloud deployment, custom models, and dedicated support may be available at an additional cost.
- **Volume-based discounts:** They are available for high-volume users, it means the more tokens you generate the less you pay for each token.
- **Pay-as-you-go:** Some enterprise customers may prefer a pay-as-you-go pricing model, where they only pay for the tokens they generate, and not a fixed monthly fee.

The pricing is based on count of Tokens wherein, the tokens are pieces of words used for natural language processing. For example, in English text, 1 token is approximately 4 characters or 0.75 words. As a point of reference, the collected works of Shakespeare are about 900,000 words or 1.2M tokens.

To purchase a language model, one needs the base model / language and then fine-tune it with training data. Below is a sample pricing model (at the time of writing this article), for reference.

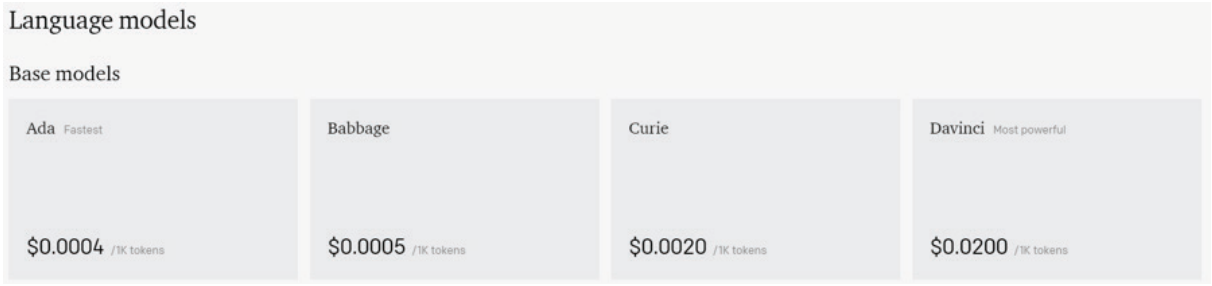

Fine-tuned models

MODEL	TRAINING	USAGE
Ada	$0.0004 / 1K tokens	$0.0016 / 1K tokens
Babbage	$0.0006 / 1K tokens	$0.0024 / 1K tokens
Curie	$0.0030 / 1K tokens	$0.0120 / 1K tokens
Davinci	$0.0300 / 1K tokens	$0.1200 / 1K tokens

Are there SLAs defined for enterprise usage?

Currently the SLA for enterprise usage are not available however may be published in future.

Where can I find more information on ChatGPT?

Please refer below links for more information.

https://openai.com/blog/chatgpt/

Author

Varun Mehta

Program Director - DWPS C&I, LTIMindtree

Varun is a seasoned technology professional with extensive subject matter expertise in Digital Workplace Service domains. He has experience in implementing multiple end user service transformations with focus on user experience, analytics, automation, and productivity.

About LTIMindtree

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700 clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by 81,000+ talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit https://www.ltimindtree.com/.