

ChatGPT and Generative AI Tools

It is almost impossible to miss the hype around ChatGPT, an example of Generative AI (Artificial Intelligence) with vast applications. I consider the background of such AI, attempt to explain the proliferation of relevant jargon necessary to understand how Natural Language Programming works at a basic level, and include some brief example applications to free-text, software programming, and graphic art. Some familiarity with the concepts of Neural Networks would be beneficial when reading this.

What is ChatGPT?

ChatGPT is an [online research experiment](#) using a natural language processing (NLP) tool driven by AI technology that allows you to have human-like conversations and much more with a chatbot. ChatGPT is also a form of [Generative AI](#) that can generate new content instead of analyzing existing data. Generative AI models can produce text and images, including program code, poetry, blog posts, and artwork.

Reaching one million subscribers in only five days, ChatGPT has almost immediately spawned countless podcasts, YouTube entrepreneurs, and social commentators to variously praise or ring alarm bells for what is a prototype test. The best way to form your own opinion is to register at <https://chat.openai.com/chat> and follow the official [ChatGPT LinkedIn page](#) for updates.

More specifically, ChatGPT is a conversational AI chatbot or example of a Large Language Model (LLM). [Conversational AI](#) is a set of technologies that allows humans and machines to have human-like interactions:

- [Automatic Speech Recognition](#) (ASR), also known as voice recognition, which translates spoken words into text. It's what allows voice assistants, like Siri and Alexa, to perform tasks based on our vocal input.
- [Natural-Language Understanding](#) (NLU) technology, a type of AI that processes language in the same way humans do — by adhering to grammatical rules, acknowledging sentence structure, and learning vocabulary and semantics. It's what lets a new generation of intelligent virtual agents (IVAs) have dynamic conversations with customers across email, chat, or digital messaging.

Not Quite SkyNet, but a Transformer Nevertheless

The GPT in ChatGPT stands for Generative Pre-trained Transformer. [Google Brain](#) created Transformer in 2017. Transformers have revolutionized AI research in recent years, leading to significant improvements in natural language processing (NLP) tasks such as machine translation, language understanding, and text generation.

As explained in a [textbook by Yildirim and Asgari-Chenaghlu](#), "A Transformer model is a deep neural network which can learn the context and thus meaning by tracking the relationship between words (tokens) in a sentence (sequence). The key innovation of Transformer models is the use of self-attention mechanisms, which allow the model to weigh the importance of different parts of the input when making a prediction.

Perhaps the most talked about Transformer-based model is GPT-3 (Generative Pre-trained Transformer 3), developed by OpenAI, the company behind ChatGPT. One of the key advantages of GPT-3 is its ability to perform a wide range of NLP tasks without the need for fine-tuning. This is possible due to the model's massive scale, which includes 175 billion parameters, making it one of the largest language models ever created.

Who is OpenAI?

Founded in 2015 by Elon Musk, Sam Altman, and various others, with around \$1 billion in capital, [OpenAI](#) is an artificial intelligence research laboratory that transitioned into a capped-for-profit organization in 2019. OpenAI now comprises an entity called OpenAI LP and the non-profit parent foundation OpenAI.

An exclusive commercial license to Microsoft is set to grow, with recent confirmation that Microsoft will invest a further \$10 billion in OpenAI to have it valued at \$29 billion. If successful, this would equate to a 49% stake in the company. Microsoft will reportedly get a 75% share of OpenAI's profits until it makes back the money on its investment, after which the company would assume a 49% stake in OpenAI. Much of the investor commentary has been on whether Microsoft can subsequently disrupt Google's position as the entry point for people on the Internet.



OpenAI are now influencing several intriguing technologies:

- [Codex](#) is an AI system that translates natural language to code, powers [GitHub co-pilot](#) and is used by AI developers day-to-day coding on [VS Code](#).
- [DALL.E 2](#) is an AI system to create realistic images and art from a description in natural language.
- [Whisper](#) is an automatic speech recognition system trained on multilingual and multitask supervised data collected from the web.
- [Jukebox](#) is a neural network that generates music, including rudimentary singing, as raw audio in a variety of genres and artist styles.

How is ChatGPT Trained?

OpenAI trained the ChatGPT language model by using Reinforcement Learning from Human Feedback (RLHF).

The initial model training used supervised fine-tuning from [a model in the GPT-3.5 series](#), which finished training in early-2022 (so ChatGPT currently has no knowledge of post-2021 events). To create a reward model for reinforcement learning, OpenAI needed to collect comparison data, which consisted of two or more model responses ranked by quality. To collect this data, they took conversations that AI trainers had with the chatbot, randomly selected a model-written message, sampled several alternative completions, and had AI trainers rank them. Using these reward models, OpenAI can fine-tune the model using [Proximal Policy Optimization](#) (refer to **Figure 1**). Several iterations of this process were performed on an [Azure AI supercomputing infrastructure](#).

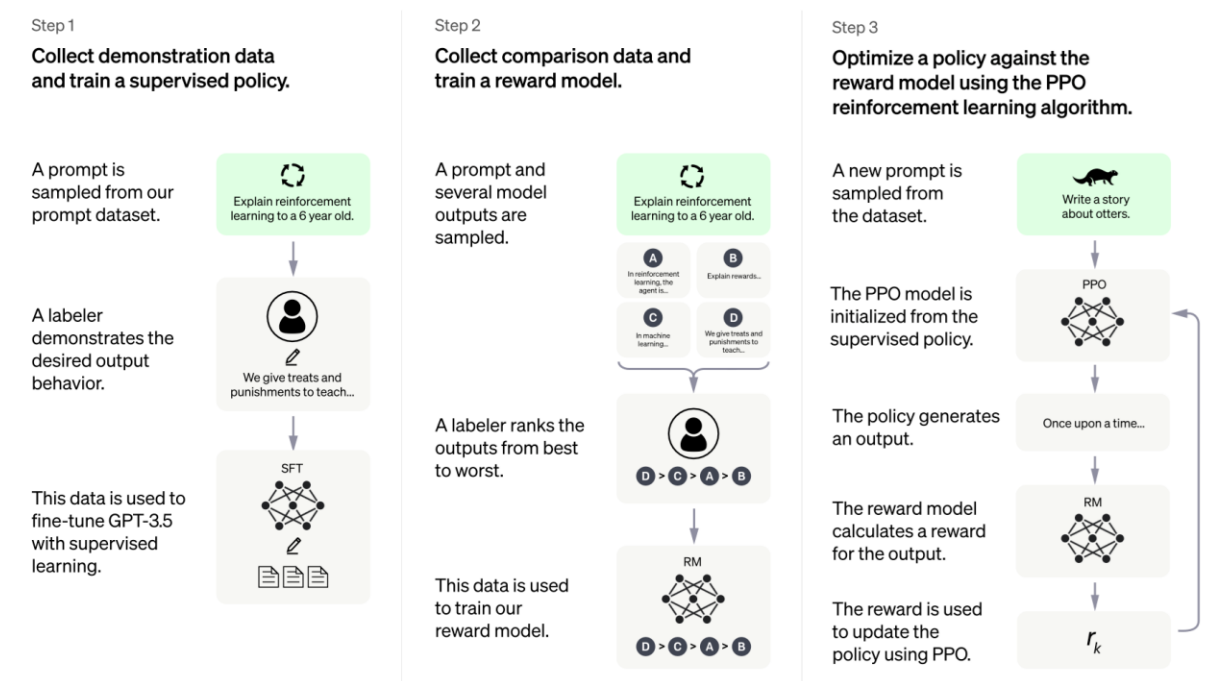


Figure 1: Schematic explanation of how ChatGPT is trained. Courtesy of OpenAI.

User Beware: It's a Prototype Experiment

ChatGPT has many remarkable capabilities. For example, "Write an essay on the global societal challenges of the energy transition to a low-carbon economy, in the style of [Hunter S. Thompson](#)", will return a grammatically perfect and appropriately nuanced "Gonzo-style" narrative. It is no surprise that educational and professional organizations everywhere are scrambling to prevent ChatGPT users circumventing traditional anti-plagiarism rules.

For a geophysicist such as myself, the ability to write programming code and scripts presents many opportunities and benefits:

- An accelerated reference, guide, and teaching tool.



- A useful debugging resource (enter the flawed code/script and ask ChatGPT what is wrong).
- A resource to develop supporting documentation.

As an example, let's say I wanted to build an R script to compute a tri-variate (three-dimensional) [probability density function](#) from an Excel file of data. [R is a free software environment](#) for statistical computing and graphics that I have grown fond of. A traditional online [step-by-step guide](#) will reinforce the methodology in a time-consuming manner, and quite likely will not address my specific requirements. Any adjustments might be made by laboriously using a search engine such as Google to find a better or more appropriate tutorial.

In contrast, a simple query to ChatGPT returned "Here is an example R script that demonstrates how to read in data from an Excel file and compute a 3D probability function" (**Figure 2**).

```
# Load the necessary libraries
library(readxl)
library(akima)

# Read in the data from the Excel file
data <- read_excel("path/to/data.xlsx")

# Assign the variables for x, y, and z
x <- data$var1
y <- data$var2
z <- data$var3

# Use the 'interp' function from the 'akima' library to create a 3D
surface of the probability function
surface <- interp(x, y, z, duplicate="mean")

# Plot the surface
library(rgl)
persp3d(surface, col="lightblue", xlab="x", ylab="y", zlab="z",
main="Probability Function")
```

Figure 2. R script provided by ChatGPT.

ChatGPT also provided these qualifying statements:

This script assumes that the Excel file "path/to/data.xlsx" contains a sheet with columns named "var1", "var2", and "var3" that represent the variables for x, y, and z respectively. The script also uses the 'interp' function from the 'akima' library to create a 3D surface of the probability function, and the 'rgl' library to plot the surface.

Keep in mind that you may need to adjust the script depending on the specific format of your data and the libraries that are installed on your system.

In a conversational manner, ChatGPT can then be challenged to make alterations, enhancements, or explain specific features almost instantaneously! I also used ChatGPT to generate vanilla C++ code for 3D Kirchhoff [pre-stack depth migration](#) (PSDM) and [Full Waveform Inversion](#) (FWI). Note that there is no 'creation' of such code, rather, the results are derived from examples within the training data.

However, as ChatGPT will itself readily explain, ChatGPT sometimes writes plausible sounding but incorrect or nonsensical answers. Fixing this issue is challenging:

- During training with Reinforcement Learning, there's currently no source of truth.
- Training the model to be more cautious causes it to decline questions that it can answer correctly.
- Supervised training misleads the model because the ideal answer depends on what the model knows, rather than what the human demonstrator knows.

ChatGPT is sensitive to tweaks to the input phrasing or attempting the same prompt multiple times. For example, given one phrasing of a question, the model can claim to not know the answer, but given a slight rephrase, can answer correctly. Furthermore, I tried a two-step approach wherein 1. I requested a response that included relevant citations, and 2. I made the same request but stated that one of the citations should be excluded. The two answers were substantially different.

The ChatGPT model is often excessively verbose and overuses certain phrases, such as restating that it's a language model trained by OpenAI. These issues arise from biases in the training data (trainers prefer longer answers that look more comprehensive) and well-known over-optimization issues.

ChatGPT versus Google

What is the difference between ChatGPT and a search engine? As described earlier, ChatGPT is a language model created with the purpose of holding a conversation with the end user. A search engine indexes web pages on the internet to help the user find the information they asked for. ChatGPT does not have the ability to search the internet for information.



So far, Google have been low-key in their response to ChatGPT and the associated threat from a more powerful Microsoft search engine. One plausible reason is that chatbot infrastructure doesn't really (yet) support the incorporation of advertising (Google's main source of revenue).

Nevertheless, Google have an alternative and reportedly more sophisticated solution called LaMDA ([Language Model for Dialogue Applications](#)) in development. ChatGPT clearly does choose a narrow path when forming responses based upon the (vast) training data of the global internet up to end-2021. In contrast, LaMDA will ask questions, try to understand the problem, and then provide answers. According to Google, "(LaMDA)... can engage in a free-flowing way about a seemingly endless number of topics, an ability we think could unlock more natural ways of interacting with technology and entirely new categories of helpful applications."

Time will tell, but putting ChatGPT in the public space is the first opportunity for everyone to recalibrate their ways of doing almost everything with Generative AI platforms... That includes applications to creative art in the next section.

ChatGPT as Co-author on Technical Publications?

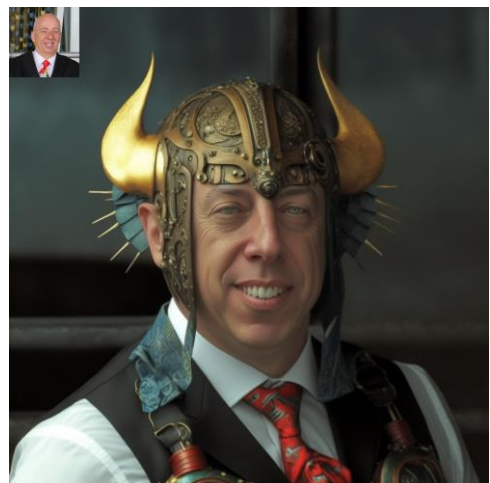
According to a statement by the SEG ([Society of Exploration Geophysicists](#)), "ChatGPT may be able to generate clear scientific abstracts, at least according to [this preprint](#), but it followed specific journal formatting requirements only 8% of the time. The authors note that the "boundaries of ethical and acceptable use of large language models to help scientific writing remain to be determined," but at least [one conference has prohibited](#) AI-generated text in submissions." A [Nature journal article](#) gives other such examples and elaborates on related ethical and established editorial policies.

All I can say is good luck trying to detect such contributions, although OpenAI reportedly recognize its programs could be used to mislead people and was [developing technology](#) to help people identify text that is generated by ChatGPT. Such a tool called [GPTZero](#) is already available elsewhere. Nevertheless, we should be focusing on the opportunities. Indeed, "[Prompt Engineer](#)" is already a relevant career option for those excited by where this may go. Phrasing the tasks in the appropriate way still requires judicious and intellectual consideration for all Generative AI tools, including graphics versions now briefly considered.

Creating Images with Generative AI

Personally, I find the OpenAI [DALL.E 2](#) results to be underwhelming and clunky. In contrast, [Midjourney](#), a self-funded research lab of only 11 full-time staff, typically returns quite remarkable results. Reflecting its prototype status, the Midjourney Bot used to generate images is accessed via the [official Discord server](#). The web App at <https://www.midjourney.com/> contains the gallery of your own work following registration, as well as other users' creations. New users can choose a "newbie" channel in the Discord server and have around 25 free uses of the [/imagine command](#) or other queries (variations, upscales). These uses are also known as "jobs" or "GPU-minutes".

The example image to the right is a "Steampunk version of my LinkedIn photo wearing a Viking helmet", in honor of my Norwegian employer. Just as ChatGPT can accept highly specific language queries, Midjourney responds to language-based requests, can incorporate provided raster images during design, and allows iterative perturbation or refinement commands.



Summary

(Written by ChatGPT in the style of Jack London, author of "[Call of the Wild](#)")

Generative AI, it's the new gold rush. With ChatGPT, a shining example of generative AI, the world of content creation has changed forever. Words, pictures, code, it's all just a dream for this technology. And with its natural language understanding, it can chat with a human. No need for fine-tuning, it can handle any task thrown its way. The opportunities are endless, the horizon is wide open. Companies are staking their claim, investing millions. Generative AI, it's the new frontier.