

# What is ChatGPT And How Can You Use It?

This is what ChatGPT is and why it may be the most important tool since modern search engines

[Roger Montti](#)



OpenAI introduced a long-form question-answering AI called ChatGPT that answers complex questions conversationally.

It's a revolutionary technology because it's [trained](#) to learn what humans mean when they ask a question.

Many users are awed at its ability to provide human-quality responses, inspiring the feeling that it may eventually have the power to disrupt how humans interact with computers and change how information is retrieved.

# What Is ChatGPT?

ChatGPT is a large language model chatbot developed by OpenAI based on GPT-3.5. It has a remarkable ability to interact in conversational dialogue form and provide responses that can appear surprisingly human.

Large language models perform the task of predicting the next word in a series of words.

Reinforcement Learning with Human Feedback (RLHF) is an additional layer of training that uses human feedback to help ChatGPT learn the ability to follow directions and generate responses that are satisfactory to humans.

## Who Built ChatGPT?

ChatGPT was created by San Francisco-based artificial intelligence company OpenAI. OpenAI Inc. is the non-profit parent company of the for-profit OpenAI LP.

OpenAI is famous for its well-known DALL·E, a deep-learning model that generates images from text instructions called prompts.

The CEO is Sam Altman, who previously was president of Y Combinator.

Microsoft is a [partner and investor](#) in the amount of \$1 billion dollars. They jointly developed the [Azure AI Platform](#).

## Large Language Models

ChatGPT is a large language model (LLM). Large Language Models (LLMs) are trained with massive amounts of data to accurately predict what word comes next in a sentence.

It was discovered that increasing the amount of data increased the ability of the language models to do more.

According to [Stanford University](#):

“GPT-3 has 175 billion parameters and was trained on 570 gigabytes of text. For comparison, its predecessor, GPT-2, was over 100 times smaller at 1.5 billion parameters.

This increase in scale drastically changes the behavior of the model — GPT-3 is able to perform tasks it was not explicitly trained on, like translating sentences from English to French, with few to no training examples.

This behavior was mostly absent in GPT-2. Furthermore, for some tasks, GPT-3 outperforms models that were explicitly trained to solve those tasks, although in other tasks it falls short.”

### **[Lead Gen for 2023] Feed Your Pipeline With Qualified Prospects**

Get your brand in front of marketers, entrepreneurs, and business pros who want to do their job more successfully.

[Download SEJ Media Kit](#)

#### Advertisement

LLMs predict the next word in a series of words in a sentence and the next sentences – kind of like autocomplete, but at a mind-bending scale.

This ability allows them to write paragraphs and entire pages of content.

But LLMs are limited in that they don't always understand exactly what a human wants.

And that's where ChatGPT improves on state of the art, with the aforementioned Reinforcement Learning with Human Feedback (RLHF) training.

## How Was ChatGPT Trained?

GPT-3.5 was trained on massive amounts of data about code and information from the internet, including sources like Reddit discussions, to help ChatGPT learn dialogue and attain a human style of responding.

ChatGPT was also trained using human feedback (a technique called Reinforcement Learning with Human Feedback) so that the AI learned what humans expected when they asked a question. Training the LLM this way is revolutionary because it goes beyond simply training the LLM to predict the next word.

A March 2022 research paper titled [Training Language Models to Follow Instructions with Human Feedback](#) explains why this is a breakthrough approach:

"This work is motivated by our aim to increase the positive impact of large language models by training them to do what a given set of humans want them to do.

By default, language models optimize the next word prediction objective, which is only a proxy for what we want these models to do.

Our results indicate that our techniques hold promise for making language models more helpful, truthful, and harmless.

Making language models bigger does not inherently make them better at following a user's intent.

For example, large language models can generate outputs that are untruthful, toxic, or simply not helpful to the user.

In other words, these models are not aligned with their users."

The engineers who built ChatGPT hired contractors (called labelers) to rate the outputs of the two systems, GPT-3 and the new InstructGPT (a "sibling model" of ChatGPT).

Based on the ratings, the researchers came to the following conclusions:

"Labelers significantly prefer InstructGPT outputs over outputs from GPT-3.

InstructGPT models show improvements in truthfulness over GPT-3.

InstructGPT shows small improvements in toxicity over GPT-3, but not bias."

The research paper concludes that the results for InstructGPT were positive. Still, it also noted that there was room for improvement.

"Overall, our results indicate that fine-tuning large language models using human preferences significantly improves their behavior on a wide range of tasks, though much work remains to be done to improve their safety and reliability."

What sets ChatGPT apart from a simple chatbot is that it was specifically trained to understand the human intent in a question and provide helpful, truthful, and harmless answers.

Because of that training, ChatGPT may challenge certain questions and discard parts of the question that don't make sense.

Another research paper related to ChatGPT shows how they trained the AI to predict what humans preferred.

The researchers noticed that the metrics used to rate the outputs of natural language processing AI resulted in machines that scored well on the metrics, but didn't align with what humans expected.

The following is how the researchers explained the problem:

"Many machine learning applications optimize simple metrics which are only rough proxies for what the designer intends. This can lead to problems, such as YouTube recommendations promoting click-bait."

So the solution they designed was to create an AI that could output answers optimized to what humans preferred.

To do that, they trained the AI using datasets of human comparisons between different answers so that the machine became better at predicting what humans judged to be satisfactory answers.

The paper shares that training was done by summarizing Reddit posts and also tested on summarizing news.

The research paper from February 2022 is called [Learning to Summarize from Human Feedback](#).

The researchers write:

"In this work, we show that it is possible to significantly improve summary quality by training a model to optimize for human preferences.

We collect a large, high-quality dataset of human comparisons between summaries, train a model to predict the human-preferred summary, and

use that model as a reward function to fine-tune a summarization policy using reinforcement learning."

## **What are the Limitations of ChatGPT?**

### **Limitations on Toxic Response**

ChatGPT is specifically programmed not to provide toxic or harmful responses. So it will avoid answering those kinds of questions.

### **Quality of Answers Depends on Quality of Directions**

An important limitation of ChatGPT is that the quality of the output depends on the quality of the input. In other words, expert directions (prompts) generate better answers.

### **Answers Are Not Always Correct**

Another limitation is that because it is trained to provide answers that feel right to humans, the answers can trick humans that the output is correct.

Many users discovered that ChatGPT can provide incorrect answers, including some that are wildly incorrect.

**Fiora**  
@FioraAeterna · [Follow](#)

didn't know this, TIL

6:40 AM · Dec 5, 2022

[Read the full conversation on Twitter](#)

1.9K [Reply](#) [Share this Tweet](#)

[Read 31 replies](#)

The moderators at the coding Q&A website Stack Overflow may have discovered an unintended consequence of answers that feel right to humans.



Stack Overflow was flooded with user responses generated from ChatGPT that appeared to be correct, but a great many were wrong answers.

The thousands of answers overwhelmed the volunteer moderator team, prompting the administrators to enact a ban against any users who post answers generated from ChatGPT.

The flood of ChatGPT answers resulted in a post entitled: [Temporary policy: ChatGPT is banned](#):

"This is a temporary policy intended to slow down the influx of answers and other content created with ChatGPT.

...The primary problem is that while the answers which ChatGPT produces have a high rate of being incorrect, they typically "look like" they "might" be good..."

The experience of Stack Overflow moderators with wrong ChatGPT answers that look right is something that OpenAI, the makers of ChatGPT, are aware of and warned about in their announcement of the new technology.

## OpenAI Explains Limitations of ChatGPT

[The OpenAI announcement](#) offered this caveat:

"ChatGPT sometimes writes plausible-sounding but incorrect or nonsensical answers.

Fixing this issue is challenging, as:

- (1) during RL training, there's currently no source of truth;
- (2) training the model to be more cautious causes it to decline questions

that it can answer correctly; and

(3) supervised training misleads the model because the ideal answer depends on what the model knows, rather than what the human demonstrator knows."

## Is ChatGPT Free To Use?

The use of ChatGPT is currently free during the "research preview" time.

The chatbot is currently open for users to try out and provide feedback on the responses so that the AI can become better at answering questions and to learn from its mistakes.

The [official announcement](#) states that OpenAI is eager to receive feedback about the mistakes:

"While we've made efforts to make the model refuse inappropriate requests, it will sometimes respond to harmful instructions or exhibit biased behavior.

We're using the Moderation API to warn or block certain types of unsafe content, but we expect it to have some false negatives and positives for now.

We're eager to collect user feedback to aid our ongoing work to improve this system."

There is currently a contest with a prize of \$500 in ChatGPT credits to encourage the public to rate the responses.

"Users are encouraged to provide feedback on problematic model outputs through the UI, as well as on false positives/negatives from the

external content filter which is also part of the interface.

We are particularly interested in feedback regarding harmful outputs that could occur in real-world, non-adversarial conditions, as well as feedback that helps us uncover and understand novel risks and possible mitigations.

You can choose to enter the ChatGPT Feedback Contest<sup>3</sup> for a chance to win up to \$500 in API credits.

Entries can be submitted via the feedback form that is linked in the ChatGPT interface."

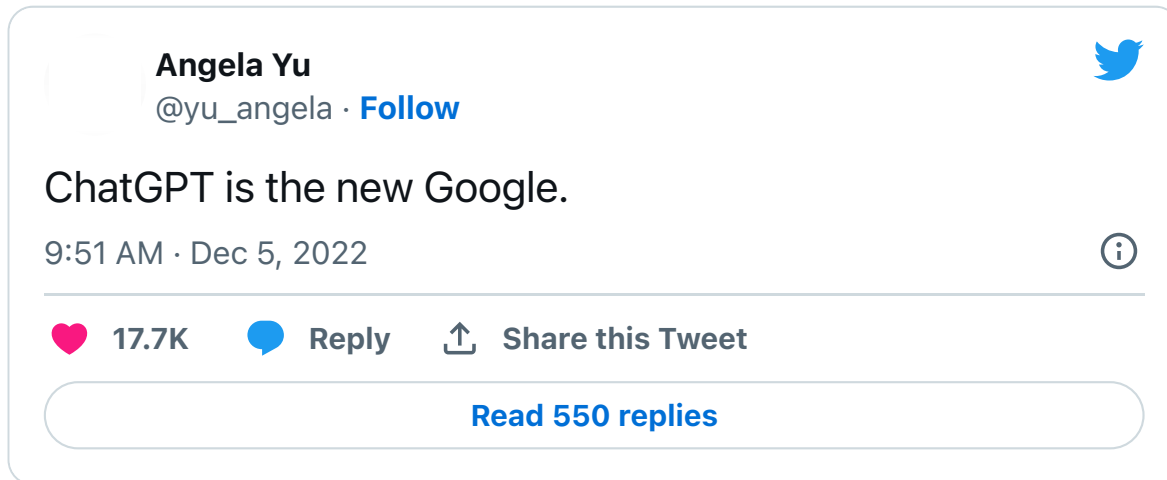
The currently ongoing contest ends at 11:59 p.m. PST on December 31, 2022.

## Will Language Models Replace Google Search?

Google itself has already created an AI chatbot that is called [LaMDA](#). The performance of Google's chatbot was so close to a human conversation that a Google engineer claimed that [LaMDA was sentient](#).

Given how these large language models can answer so many questions, is it far-fetched that a company like OpenAI, Google, or Microsoft would one day replace traditional search with an AI chatbot?

Some on Twitter are already declaring that ChatGPT will be the next Google.



The scenario that a question-and-answer chatbot may one day replace Google is frightening to those who make a living as search marketing professionals.

It has sparked discussions in online search marketing communities, like the popular Facebook SEOSignals Lab where [someone asked](#) if searches might move away from search engines and towards chatbots.

Having tested ChatGPT, I have to agree that the fear of search being replaced with a chatbot is not unfounded.

The technology still has a long way to go, but it's possible to envision a hybrid search and chatbot future for search.

But the current implementation of ChatGPT seems to be a tool that, at some point, will require the purchase of credits to use.

## How Can ChatGPT Be Used?

ChatGPT can write code, poems, songs, and even short stories in the style of a specific author.

The expertise in following directions elevates ChatGPT from an information

source to a tool that can be asked to accomplish a task.

This makes it useful for writing an essay on virtually any topic.

ChatGPT can function as a tool for generating outlines for articles or even entire novels.

It will provide a response for virtually any task that can be answered with written text.

## Conclusion

As previously mentioned, ChatGPT is envisioned as a tool that the public will eventually have to pay to use.

[Over a million users](#) have registered to use ChatGPT within the first five days since it was opened to the public.

### More resources:

- [ChatGPT Examples: 5 Ways SEOs and Digital Marketers Can Use ChatGPT](#)
- [ChatGPT For Content and SEO?](#)
- [The Future Of Chatbots: Use Cases & Opportunities You Need To Know](#)
- [How Machine Learning in Search Works](#)

*Featured image: Shutterstock/Asier Romero*