

Speech to text

Learn how to turn audio into text

Introduction

The speech to text API provides two endpoints, `transcriptions` and `translations`, based on our state-of-the-art open source large-v2 [Whisper model](#). They can be used to:

- Transcribe audio into whatever language the audio is in.
- Translate and transcribe the audio into english.

File uploads are currently limited to 25 MB and the following input file types are supported: `mp3`, `mp4`, `mpeg`, `mpga`, `m4a`, `wav`, and `webm`.

Quickstart

Transcriptions

The transcriptions API takes as input the audio file you want to transcribe and the desired output file format for the transcription of the audio. We currently support multiple input and output file formats.

Transcribe audio

python  Copy

```
1 # Note: you need to be using OpenAI Python v0.27.0 for the code below to work
2 import openai
3 audio_file= open("/path/to/file/audio.mp3", "rb")
4 transcript = openai.Audio.transcribe("whisper-1", audio_file)
```

By default, the response type will be json with the raw text included.

```
{
  "text": "Imagine the wildest idea that you've ever had, and you're curious
about how it might scale to something that's a 100, a 1,000 times bigger.
....
}
```

To set additional parameters in a request, you can add more `--form` lines with the relevant options. For example, if you want to set the output format as text, you would add the following line:

```
1 ...
2 --form file=@openai.mp3 \
3 --form model=whisper-1 \
4 --form response_format=text
```

Translations

The translations API takes as input the audio file in any of the supported languages and transcribes, if necessary, the audio into English. This differs from our `/Transcriptions` endpoint since the output is not in the original input language and is instead translated to English text.

Translate audio

python  Copy

```
1 # Note: you need to be using OpenAI Python v0.27.0 for the code below to work
2 import openai
3 audio_file= open("/path/to/file/german.mp3", "rb")
4 transcript = openai.Audio.translate("whisper-1", audio_file)
```

In this case, the inputted audio was german and the outputted text looks like:

Hello, my name is Wolfgang and I come from Germany. Where are you heading today?

We only support translation into english at this time.

Supported languages

We currently [support the following languages](#) through both the `transcriptions` and `translations` endpoint:

Afrikaans, Arabic, Armenian, Azerbaijani, Belarusian, Bosnian, Bulgarian, Catalan, Chinese, Croatian, Czech, Danish, Dutch, English, Estonian, Finnish, French, Galician, German, Greek, Hebrew, Hindi, Hungarian, Icelandic, Indonesian, Italian, Japanese, Kannada, Kazakh, Korean, Latvian, Lithuanian, Macedonian, Malay, Marathi, Maori, Nepali, Norwegian, Persian, Polish, Portuguese, Romanian, Russian, Serbian, Slovak, Slovenian, Spanish, Swahili, Swedish, Tagalog, Tamil, Thai, Turkish, Ukrainian, Urdu, Vietnamese, and Welsh.

While the underlying model was trained on 98 languages, we only list the languages that exceeded <50% [word error rate](#) (WER) which is an industry standard benchmark for speech to text model accuracy. The model will return results for languages not listed above but the quality will be low.

Longer inputs

By default, the Whisper API only supports files that are less than 25 MB. If you have an audio file that is longer than that, you will need to break it up into chunks of 25 MB's or less or used a compressed audio format. To get the best performance, we suggest that you avoid breaking the audio up mid-sentence as this may cause some context to be lost.

One way to handle this is to use the [PyDub open source Python package](#) to split the audio:



```
1 from pydub import AudioSegment
2
3 song = AudioSegment.from_mp3("good_morning.mp3")
4
5 # PyDub handles time in milliseconds
6 ten_minutes = 10 * 60 * 1000
7
8 first_10_minutes = song[:ten_minutes]
9
10 first_10_minutes.export("good_morning_10.mp3", format="mp3")
```

OpenAI makes no guarantees about the usability or security of 3rd party software like PyDub.

Prompting

You can use a [prompt](#) to improve the quality of the transcripts generated by the Whisper API. The model will try to match the style of the prompt, so it will be more likely to use capitalization and punctuation if the prompt does too. However, the current prompting system is much more limited than our other language models and only provides limited control over the generated audio. Here are some examples of how prompting can help in different scenarios:

- 1 Prompts can be very helpful for correcting specific words or acronyms that the model often misrecognizes in the audio. For example, the following prompt improves the transcription of the words DALL-E and GPT-3, which were previously written as "GDP 3" and "DALI": "The transcript is about OpenAI which makes technology like DALL-E, GPT-3, and ChatGPT with the hope of one day building an AGI system that benefits all of humanity"
- 2 To preserve the context of a file that was split into segments, you can prompt the model with the transcript of the preceding segment. This will make the transcript more accurate, as the model will use the relevant information from the previous audio. The model will only consider the final 224 tokens of the prompt and ignore anything earlier. For multilingual inputs, Whisper uses a custom tokenizer. For English only inputs, it uses the standard GPT-2 tokenizer which are both accessible through the open source [Whisper Python package](#).
- 3 Sometimes the model might skip punctuation in the transcript. You can avoid this by using a

simple prompt that includes punctuation: "Hello, welcome to my lecture."

- 4 The model may also leave out common filler words in the audio. If you want to keep the filler words in your transcript, you can use a prompt that contains them: "Umm, let me think like, hmm... Okay, here's what I'm, like, thinking."
- 5 Some languages can be written in different ways, such as simplified or traditional Chinese. The model might not always use the writing style that you want for your transcript by default. You can improve this by using a prompt in your preferred writing style.