




What is a token?

A one-page cheat sheet, in plain English

IN ONE SENTENCE

AI language models do not read text letter by letter. They process it in small chunks called tokens, and those chunks are what get counted.




1



A token is a chunk of text

It can be a whole word, part of a word, a single character, or even a byte.

Sources: OpenAI and Anthropic docs.



2



One word can be several tokens

OpenAI's docs show the string "tokenization" splitting into "token" and "ization". Smaller pieces help models handle uncommon words.

Other models may cut the same word differently.



3



Your prompt and the answer share one limit

The context window covers everything the model can reference, including its own response. Both are counted in tokens.

Try: Paste only the pages that matter, not the whole file. (General good practice: Anthropic notes that more context is not automatically better.)




4



Different models, different counts

How many characters a token covers depends on the model and the language. Never assume one model's count applies to another.




5



A rough guide, not a rule

For English, OpenAI's docs say 1 token is roughly 4 characters or 0.75 words. Anthropic's glossary says about 3.5 English characters for earlier Claude models.

Estimates only. For a closer count, use the provider's token counter for the model you will use.



Quick rule: tokens are pieces of text, not words. Treat any "tokens per word" number as a rough guess.