interweb

Generated on 28/11/2026 21:58

KNOWLEDGE BASE · PUBLIC DOCUMENT

Context window and maximum output: what context_length and max_output_tokens mean

Playground & Text Models

CODE	VERSION	COUNTRY · LANGUAGE
RLX-MZ-KB-2986	v2026.11.28	Mozambique · English

Interweb Mozambique
RLX Holdings, S.A. · Group RLX Holdings
Tax ID (NUIT) 401 191 607 · Legal entity no. (NUEL) 105 045 974 · Economic activity code (CAE) 62022

https://www.meu.interweb.co.mz/knowledgebase/2986

Context window and maximum output: what context_length and max_output_tokens mean

Playground & Text Models Created on **November 28, 2026** 2 min read

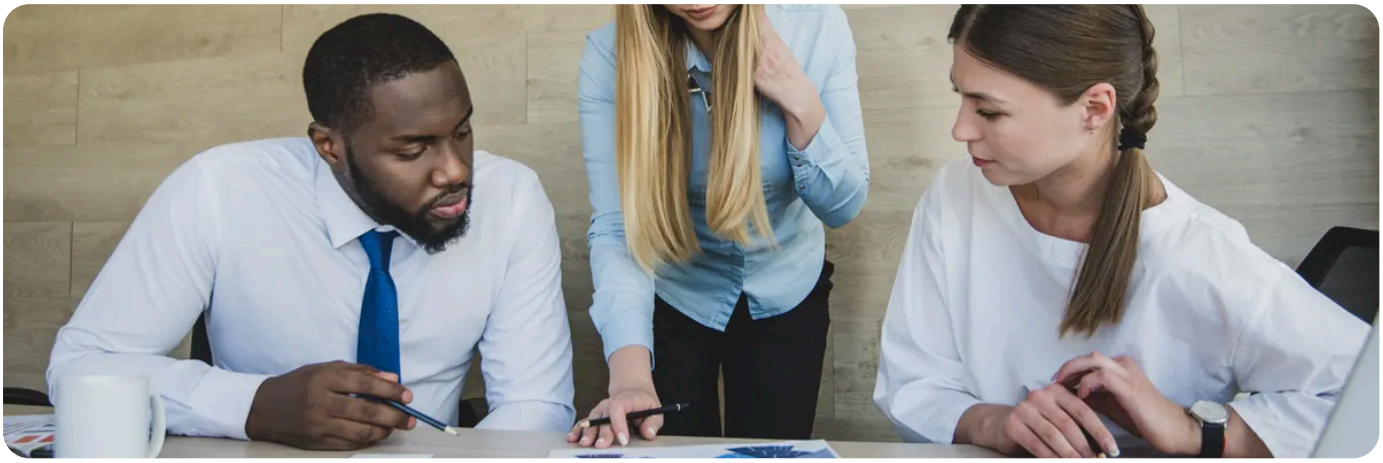

Every text model has two limits: how much it reads at once, `context_length`, and how much it writes in one answer, `max_output_tokens`. Both are counted in tokens.

The two numbers

Field	What it is
`context_length`	The most tokens the model reads in one request, input and output together.
`max_output_tokens`	The most tokens the model writes in one answer.

Two examples from the documentation: `deepseek-ai/deepseek-v4-flash` reads up to 1,048,576 tokens and writes up to 393,216; `moonshotai/kimi-k2.6` reads up to 262,144 and writes up to 262,144.

Where you see them

The **Models** tab shows the context of each text model ("Context: N tokens"). Both fields come in `GET /v1/models` and in the `aio model` command, followed by the id.

What counts towards the context

Everything in the request: the instructions, the whole conversation, the texts you pasted and the answer itself. In a conversation, each new message carries the earlier ones. For agents, the instructions and the knowledge go in every request. That is why a very long conversation ends up hitting the limit. The Playground shows "The conversation is too long. Please start a new one." See **why the AI forgets what you said earlier**.

max_output_tokens and your max_tokens

`max_output_tokens` is the model's ceiling. `max_tokens` is the limit you set on each request, below that ceiling. The credit reservation counts your `max_tokens`, or 4,096 if you do not send it. In the Playground, the "Maximum tokens in the reply" field takes 16 to 32,768, and the estimate never counts more than the model's maximum output.

In coding tools

When a tool asks for the context size, use `context_length`. For the maximum output, use a smaller value. We reserve credits for the maximum output the tool asks for in each answer. That is why the examples ask for 16,000 tokens: it is enough for coding and needs a smaller balance than a higher value.

RECOMMENDED

See the context of each model on the Models tab.

See All in One AI →

SEE ALSO

- **Why the AI forgets what you said earlier: context and long conversations**
- **Choosing a text model: price, context window and what each one is for**
- **The Models tab: searching for a model and copying its id**
- **Tokens explained: how AI model APIs are counted**
- **All in One AI**

interweb

Interweb Mozambique

RLX Holdings, S.A. · Group RLX Holdings

TAX ID (NUIT)	401 191 607
LEGAL ENTITY NO. (NUEL)	105 045 974
ECONOMIC ACTIVITY CODE (CAE)	62022
REGISTERED OFFICE	Bairro 7 de Setembro Vilankulo 1304 · CEP 0411-02, Inhambane, Mozambique
CONTACT	info@interweb.co.mz · +258 84 635 1111 · +258 86 635 1111 · www.interweb.co.mz
OPENING HOURS	Monday to Friday - 08 AM to 04 PM Saturday & Sunday - Closed

Online version of this document

RLX-MZ-KB-2986 · v2026.11.28 · Generated on 28/11/2026 21:58