GPT-6 Luna

OpenAI

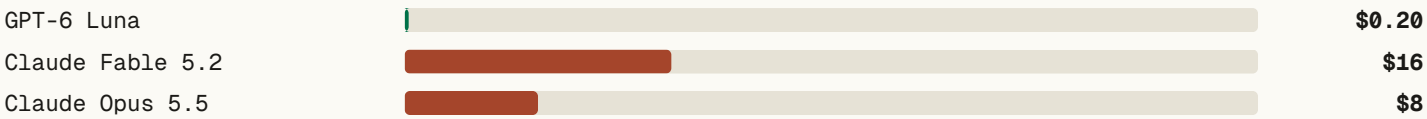
AT A GLANCE

PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
$0.10 / $0.50	**1.05M**	**128K**	**file, image, text**
input / output · cached $0.01	tokens in	tokens per reply	returns text

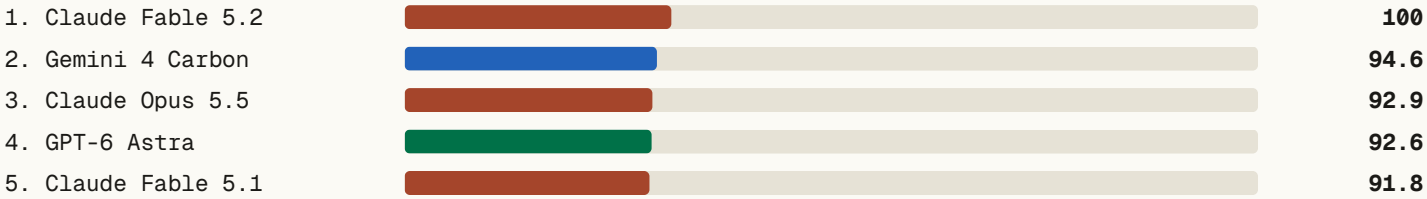
HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
GPT-6 Luna	new	$0.10	$0.50	1.05M	29.6
Claude Fable 5.2	today's #1	$8	$40	1M	100
Claude Opus 5.5	best value	$4	$20	1M	92.9
vs Claude Fable 5.2	**99% cheaper**	today's #1: $8 / $40, bench index 100; context 1M			
vs Claude Opus 5.5	**98% cheaper**	best value: $4 / $20, bench index 92.9; context 1M			

PRICE PER 1M TOKENS, BLENDED 3 IN : 1 OUT



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



GPT-6 Luna sits at 29.6 on the bench index. The bench index is the median of each model's percentile across the public leaderboards that rate it (LMArena, LiveBench, SimpleBench, LLM Stats, Vals AI, Artificial Analysis, Epoch AI).

Sources: price, limits and modalities for GPT-6 Luna as the OpenRouter catalog lists openai/gpt-6-luna; the other models as their makers list them (models.dev); rankings from the ArtificialWatch Bench Index. Facts only: no forecast of ours, no claims from the lab's announcement.