MiMo-V2.5

Xiaomi

AT A GLANCE

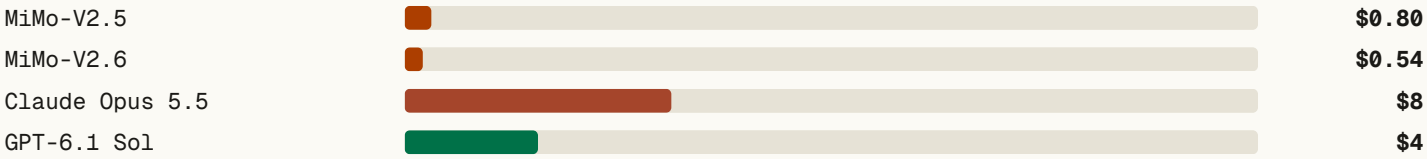
PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
$0.40 / $2	**1M**	**66K**	**text, audio, image, video**
input / output · cached $0.08	tokens in	tokens per reply	returns text

HOW IT COMPARES

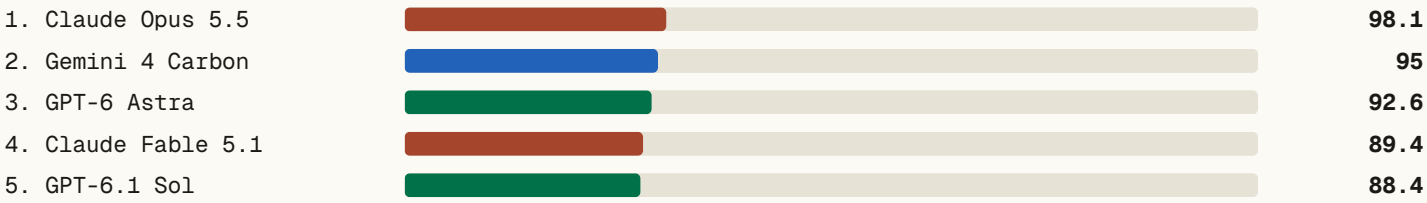
MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
MiMo-V2.5	new	$0.40	$2	1M	not rated yet
MiMo-V2.6	it follows	$0.435	$0.87	1.05M	48.2
Claude Opus 5.5	today's #1	$4	$20	1M	98.1
GPT-6.1 Sol	best value	$2	$10	1.05M	88.4

vs MiMo-V2.6	**input 8% cheaper, output 2.3x the price**	the model it follows: $0.435 / $0.87, bench index 48.2; context 1.05M
vs Claude Opus 5.5	**90% cheaper**	today's #1: $4 / $20, bench index 98.1
vs GPT-6.1 Sol	**80% cheaper**	best value: $2 / $10, bench index 88.4; context 1.05M

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



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



MiMo-V2.5 is not rated yet: the public boards usually score a new model within days, and its page on artificialwatch.com gains the scores then. 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).

Sources: price, limits and modalities for MiMo-V2.5 as the OpenRouter catalog lists xiaomi/mimo-v2.5; 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.