Ternary Bonsai 2 27B

PrismML

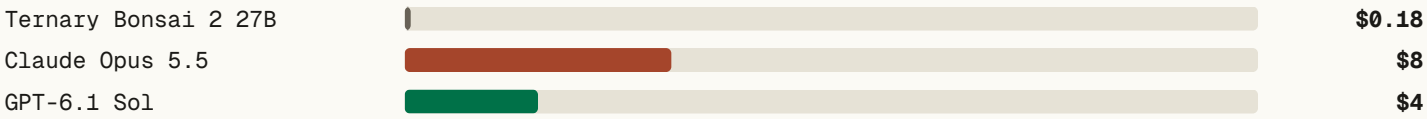
AT A GLANCE

PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
$0.075 / $0.50	**262K**	**33K**	**text, image**
input / output · cached $0.038	tokens in	tokens per reply	returns text

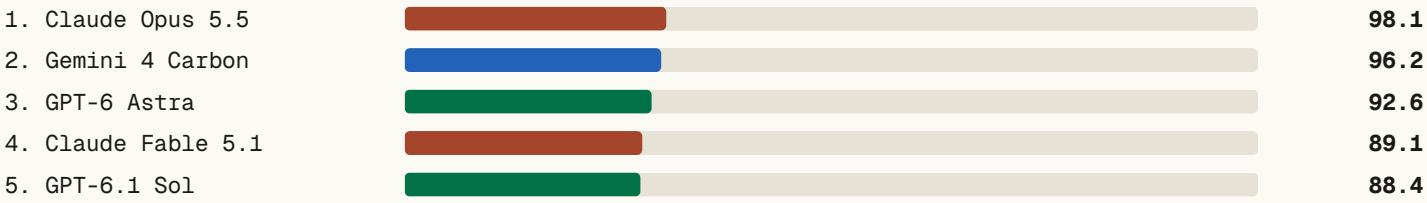
HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Ternary Bonsai 2 27B	new	$0.075	$0.50	262K	not rated yet
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 Claude Opus 5.5	**98% cheaper** today's #1: $4 / $20, bench index 98.1; context 1M				
vs GPT-6.1 Sol	**input 96% cheaper, output 95% 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



Ternary Bonsai 2 27B 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 Ternary Bonsai 2 27B as the OpenRouter catalog lists prism-ml/ternary-bonsai-2-27b; 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.