Perceptron Mk1.5

Perceptron

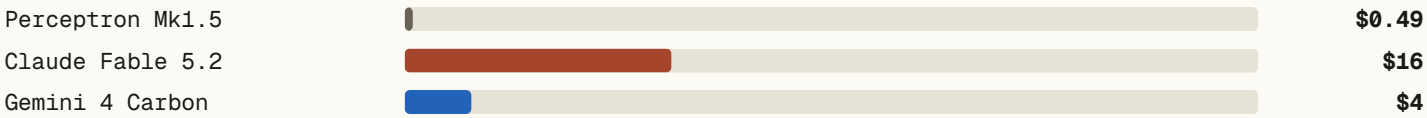
AT A GLANCE

PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
$0.15 / $1.50	**37K**	**8K**	**text, image, video, audio**
input / output	tokens in	tokens per reply	returns text

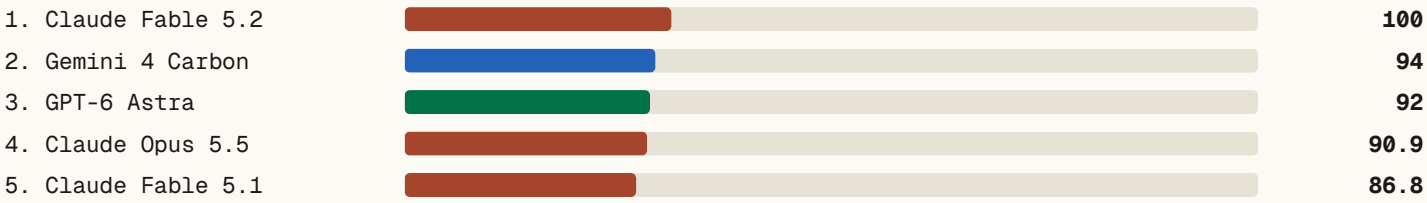
HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Perceptron Mk1.5	new	$0.15	$1.50	37K	not rated yet
Claude Fable 5.2	today's #1	$8	$40	1M	100
Gemini 4 Carbon	best value	$2	$10	1.05M	94
vs Claude Fable 5.2	**input 98% cheaper, output 96% cheaper**	today's #1: $8 / $40, bench index 100; context 1M			
vs Gemini 4 Carbon	**input 93% cheaper, output 85% cheaper**	best value: $2 / $10, bench index 94; context 1.05M			

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



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



Perceptron Mk1.5 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, Epoch AI).

Sources: price, limits and modalities for Perceptron Mk1.5 as the OpenRouter catalog lists perceptron/perceptron-mk1.5; 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.