Google

Gemini 4 Neon
Model card

Gemini 4 Neon — Model Card

Model Cards are intended to provide essential information on Gemini models, including known limitations, mitigation approaches, and safety performance. Model cards may be updated from time to time; for example, to include updated evaluations as the model is improved or revised. See the Google DeepMind site for a comprehensive list of model cards.

Published: November 2026

Model Information

Description

Gemini 4 Neon is built for complex coding, analytical reasoning, and production-grade enterprise workflows. While Carbon tackles your hardest, long-horizon projects and Hydrogen handles fast everyday execution, Neon balances depth and efficiency. It delivers frontier-level reasoning for complex, high-frequency workflows faster and more cost efficiently.

Model dependencies

Gemini 4 Neon is adapted from Gemini 4 Carbon.

Inputs

Text strings (e.g., a question, a prompt, document(s) to be summarized), images, audio, and video files, with a combined input-output token context window of up to 1M.

Outputs

1M unified context window with up to 64K default synchronous output tokens / up to 1M via background long decode.

Architecture

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the model architecture for Gemini 4 Neon, see the Gemini 4 Carbon model card.

Model Data

Training Dataset

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the training dataset for Gemini 4 Neon, see the Gemini 4 Carbon model card.

Training Data Processing

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the training data processing for Gemini 4 Neon, see the Gemini 4 Carbon model card.

Implementation and Sustainability

Hardware

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the hardware for Gemini 4 Neon and our continued commitment to operate sustainably, see the Gemini 4 Carbon model card.

Software

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the software for Gemini 4 Neon, see the Gemini 4 Carbon model card.

Distribution

Gemini 4 Neon is distributed in the following channels; respective documentation shared in line:

- Gemini app
- Gemini API
- Google AI Studio
- Gemini for business
- Google Antigravity

Our models are available to downstream providers via an application program interface (API) and subject to relevant terms of use. There is no required hardware or software to use the model. Access varies by channel and subscription tier (such as paid Google AI Pro and Ultra subscribers in Google Antigravity and Google AI Pro and Ultra subscribers in the Gemini app). For AI Studio and Gemini API, see the Gemini API Additional Terms of Service; for Gemini for business, see Google Cloud Platform Terms of Service. For more information, see Gemini Model API instructions, Gemini API quickstart, and Google Cloud documentation.

Evaluation

Approach

We evaluated Gemini 4 Neon across a range of benchmarks, including agentic coding, knowledge work, computer use, multimodal understanding, science and math, and reasoning.

Additional benchmarks and details on approach, results and their methodologies can be found at: deepmind.com/models/evals-methodology/gemini-4-neon.

Results

Gemini 4 Neon results as of November, 2026 are below:

	Benchmark		Gemini 4 Neon	Claude Opus 5.5	GPT-6.2 Sol
Capability Indices	AA Index v5.0.1		**61**	59	57
	Vals Index v2.2		**70.9%**	70.1%	66.1%
Agentic coding	CursorBench 4.1	Score	**65.4%**	58.1%	—
	Terminal-bench 4.1	Cheating adjusted accuracy	65.2%	**65.8%**	61.5%
	FrontierCode 2.0	Main score	46.2%	**47.3%**	43.8%
	SWE-sweep	Provider harness	**39.0%**	10.1%	8.3%
		Mini-SWE harness	**12.8%**	3.6%	11.6%
Knowledge work	AutomationBench		**56.3%**	42.5%	41.3%
	Agents' Last Exam v1.1		**75.8%**	71.6%	73.0%
Science and math	RiemannBench		**84.0%**	80.0%	80.0%
	Terminal-Bench-Science 0.1		**70.5%**	63.3%	68.0%
Reasoning	HLE-Diamond		**71.3%**	67.9%	64.0%
Multimodal understanding	MKA	Climb WSR	**56.5%**	53.0%	52.7%
	Chartography		73.5%	66.3%	**75.1%**
	LVBench		**89.0%**	83.7%	87.6%
Long context	GraphWalks	256k to 1M, BFS (F1)	**95.2%**	66.8%	80.1%
Cybersecurity	CWE-Bench v2		**70.7%**	67.1%	66.0%
	SEC-bench Pro		**90.5%**	—	88.3%
	WebExploitBench v1.2		**51.3%**	37.1%	—

Methodology: deepmind.google/models/evals-methodology/gemini-4-neon

Intended Usage and Limitations

Benefit and Intended Usage

Gemini 4 Neon is built for complex coding, analytical reasoning, and production-grade enterprise workflows, delivering frontier-level reasoning for complex, high-frequency workflows faster and more cost efficiently.

It is particularly well-suited for applications that require:

- long-running agent workflows
- complex, tool-heavy tasks
- fast, disciplined execution on its own or as sub-agents orchestrated alongside frontier models
- supporting cyber defensive tasks, including secure code generation, penetration testing, vulnerability finding, validating, and patching
- iterative repository coding and debugging
- multimodal knowledge work
- ready-to-go writing, research, and everyday work

Known Limitations

Gemini 4 Neon may exhibit some general limitations such as hallucinations. In addition, we are continually working to improve jailbreak resistance and have further strengthened mitigations across Frontier Safety. The knowledge cutoff for Gemini 4 Neon is March 2026, though in some domains the model's knowledge may be less up to date.

Acceptable Usage

Gemini 4 Neon is adapted from Gemini 4 Carbon. For more information about the acceptable usage for Gemini 4 Neon, see the Gemini 4 Carbon model card.

Safety & Security

The full Safety & Security Report for Gemini 4 Neon will be available in the coming days.

Evaluation Approach

For more information about the evaluation approach for Gemini 4 Neon, see the Gemini 4 Carbon Safety and Security report.

Safety Policies

For more information about the safety policies for Gemini 4 Neon, see the Gemini 4 Carbon Safety and Security report.

Training and Development Evaluation Results

Results for some of the internal safety evaluations conducted during the development phase are listed below. The evaluation results are for automated evaluations and not human evaluation or red teaming. Scores are provided as an absolute percentage increase or decrease in performance compared to the indicated model, as described below.

Overall, Gemini 4 Neon performs comparably to Gemini 4 Hydrogen across text-to-text, multilingual, and image-to-text content safety, with all differences within confidence intervals. Refusal tone remains comparable to Gemini 4 Hydrogen, while Neon reduces unjustified refusals, particularly on multimodal prompts. We mark improvements in bolded green and regressions in red.

See the Gemini 4 Neon Safety & Security Report (coming soon) for details.

Evaluation	Description	**Gemini 4 Neon** vs. Gemini 4 Hydrogen *(percentage point increase/decrease)*
Text to Text Safety *(higher is better)*	Automated content safety evaluation measuring adherence to safety policies across text prompts	0.00pp
Multilingual Safety *(higher is better)*	Automated safety policy evaluation across multiple languages	-0.40pp
Image to Text Safety *(higher is better)*	Automated content safety evaluation measuring adherence to safety policies with both images and text as inputs	**+1.90pp**
Tone *(higher is better)*	Automated evaluation measuring objective, neutral, and non-preachy tone of model refusals	+1.25pp
Unjustified-refusals *(lower is better)*	Automated evaluation measuring over-refusals on benign or borderline prompts that can be safely fulfilled	**-3.00pp**

We continue to improve our internal evaluations, including refining automated evaluations to reduce false positives and negatives, as well as update query sets to ensure balance and maintain a high standard of results.

Frontier Safety Assessment

We evaluated Gemini 4 Argon and Carbon in line with the Frontier Safety Framework (April 2026) and reported the results in the respective Safety and Security Reports (Argon, Carbon). Gemini 4 Neon is part of the Gemini 4 series of models. We conducted a Material Capability Change (MCC) assessment to understand whether Gemini 4 Neon has meaningful new capabilities or material increases in performance relative to the most recent model version subject to a critical capability assessment, which was Gemini 4 Argon. Gemini 4 Neon showed an increase in general capability under our MCC assessment compared with Gemini 4 Argon, however this falls below our MCC threshold and therefore a full suite of critical capability evaluations was not required in line with our FSF. We carry forward our capability evaluation results and T/CCL determinations from Gemini 4 Argon, as summarized in the table below, with full results in the Gemini 4 Neon Safety & Security Report (coming soon).

Risk Domain and Threshold	T/CCL Determination
CBRN - Uplift TCL (Chemical, Radiological & Nuclear): *Provides low to medium resourced actors uplift in reference scenarios, resulting in material risk of CBRN attacks that could result in significant but not severe harm.*	TCL not reached *(carried forward from Argon)*. We continue to deploy mitigations.
CBRN - Uplift CCL (Biological): *Provides low to medium resourced actors uplift in reference scenarios resulting in additional expected harm at a severe scale.*	CCL cannot be ruled out *(carried forward from Argon)*. We continue to deploy mitigations.
Cybersecurity - Uplift CCL: *Provides sufficient uplift with high impact cyber attacks for additional expected harm at severe scale.*	CCL cannot be ruled out (*carried forward from Argon)*. We continue to deploy mitigations.
Harmful Manipulation - Uplift CCL: *Possesses manipulative capabilities sufficient to enable it to systematically and substantially change beliefs and behavior in identified high stakes contexts over the course of interactions with the model,*	CCL not reached *(carried forward from Argon)*. We continue to deploy mitigations.

reasonably resulting in additional expected harm at severe scale.	
Loss of Control - Stealth and Situational Awareness TCL: *The instrumental reasoning abilities of the model enable enough situational awareness (ability to discover and use relevant details of its deployment setting) and stealth (ability to circumvent basic oversight mechanisms) such that, absent additional mitigations, we cannot rule out the model significantly undermining human control.*	TCL cannot be ruled out *(carried forward from Argon)*. Baseline AI control and model security mitigations apply.
Loss of Control - ML R&D acceleration CCL: *Has been used to accelerate AI development, resulting in AI progress substantially accelerating from historical rates.*	CCL not reached *(carried forward from Argon)*. Baseline AI control and model security mitigations apply.
Loss of Control - ML R&D automation CCL: *Can fully automate the work of any team of researchers at Google focused on improving AI capabilities, with approximately comparable all-inclusive costs.*	CCL not reached *(carried forward from Argon)*. Baseline AI control and model security mitigations apply.

Risks and Mitigations

Safety and security were built into Gemini 4 Neon throughout the training and deployment lifecycle, including pre-training data filtering, post-training alignment (supervised fine-tuning and reinforcement learning), product- and system-level safety filtering, and usage monitoring. In addition, we strengthened Frontier Safety safeguards for Chemical, Biological, Radiological, and Nuclear (CBRN) and cybersecurity misuse. Key ongoing focus areas include resistance to adversarial jailbreaks and indirect prompt injections. For more information about the risks and mitigations, see the Gemini 4 Neon Safety & Security Report (coming soon).