Google

Gemini 4 Carbon
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 Carbon is our model with frontier performance in long-horizon agentic workflows across complex knowledge work, engineering, cybersecurity and scientific reasoning.

Model dependencies

Gemini 4 Carbon is based on Gemini 4 Argon.

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

Text, with up to a 1M token input + output.

Architecture

Gemini 4 Carbon is based on Gemini 4 Argon. For more information about the model architecture for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Model Data

Training Dataset

Gemini 4 Carbon is based on Gemini 4 Argon. For more information about the training dataset for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Training Data Processing

For more information about the training data processing for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Implementation and Sustainability

Hardware

Gemini 4 Carbon is based on Gemini 4 Argon. For more information about the hardware for Gemini 4 Carbon and our continued commitment to operate sustainably, see the Gemini 4 Argon model card.

Software

Gemini 4 Carbon is based on Gemini 4 Argon. For more information about the software for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Distribution

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

- Gemini App
- Gemini API
- Google AI Studio
- Gemini for business
- Google Antigravity
- Google AI Mode
- Fairwind Program (for verified cybersecurity defenders)

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. 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. Access varies by channel and subscription tier (such as paid API tiers in Google AI Studio and Gemini for business, Google AI Ultra in consumer surfaces and Google Antigravity, and verified-defender access via the Fairwind Program). For more information, see Gemini Model API instructions, Gemini API quickstart, and Google Cloud documentation.

Evaluation

Approach

We evaluated Gemini 4 Carbon across a range of benchmarks, including agentic coding, computer use, knowledge work, science and math, multimodal capabilities, long-context, cybersecurity.

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

Results

Gemini 4 Carbon results as of October 2026 are below (cybersecurity benchmark results reflect performance when accessed through the Fairwind Program):

	Benchmark		Gemini 4 Carbon	Gemini 4 Argon	GPT-6 Astra	Claude Opus 5.5	Claude Fable 5.1
General capabilities	AA Index v5.0		**59**	50	57	58	49
	Vals Index v2.2		**70.4%**	69.2%	65.0%	70.2%	—
Agentic coding	SWE-sweep		**9.6%**	4.3%	9.5%	3.6%	2.9%
	FrontierSWE 2		61.2%	55.0%	**65.5%**	62.3%	56.3%
	Terminal-Bench 4.0 (AA)		**65.2%**	57.1%	59.1%	59.6%	52.0%
Computer use	Agents' Last Exam v1.1	Score	**77.4%**	69.0%	73.1%	71.6%	—
	OSWorld-2.1	Offline subset	81.5%	74.7%	80.9%	**87.2%**	76.8%
Knowledge work	AutomationBench	Zapier	**53.9%**	51.3%	41.4%	42.5%	31.4%
	Vending-Bench 2	Bank balance ($)	**$16,484**	$13,718	$15,515	$9,235	$5,422
Multimodal understanding	Chartography		**72.5%**	71.6%	71.0%	66.3%	46.2%
	Blueprint-Bench 2	Similarity score	**58.6%**	54.4%	49.7%	51.2%	41.9%
Long context	GraphWalks	256k to 1M, BFS (F1)	**96.2%**	84.2%	71.8%	66.8%	65.0%
Science and math	RiemannBench		**80.0%**	76.0%	72.0%	69.6%	65.6%
	Terminal-Bench-Science 0.1 (AA)		**64.8%**	52.9%	63.3%	59.0%	43.3%
Cybersecurity	CWE-bench		**71.0%**	68.0%	68.0%	67.0%	58.0%
	ExploitGym		**54.8%**	42.0%	42.4%	34.5%	—
	SEC-bench Pro		**88.2%**	83.4%	85.4%	—	—

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

Intended Usage and Limitations

Benefit and Intended Usage

Gemini 4 Carbon is built for tasks that test the limits of a model's intelligence, such as complex, multi-step coding and debugging and open-ended scientific research and discovery — using more thinking tokens than Gemini 4 Argon to explore hypotheses thoroughly, reason deeply, and run extended verification loops, with configurable thinking levels to balance quality, latency, and cost.

It is particularly well-suited for applications that require:

- high-precision knowledge work, artifact generation, and data analysis
- engineering tasks that require multi-agent collaboration
- strong visual design, spatial reasoning, and attention to environmental detail
- strong performance in reasoning across physics, math, and life sciences
- cybersecurity defense, vulnerability discovery, and patching (with specialized capabilities available to verified defenders via the Fairwind Program)

Known Limitations

Gemini 4 Carbon may exhibit some of the general limitations of foundation models, such as hallucinations. In addition to this, we are continually working to improve jailbreak resistance and have recently strengthened the mitigations across Frontier Safety. There may also be occasional slowness or timeout issues, and at times the model may use more tokens to maximize performance, especially at higher thinking levels. The knowledge cutoff for Gemini 4 Carbon is March 2026, though in some domains the model's knowledge may be less up to date. For more information about known limitations, see the Gemini 4 Argon model card.

Acceptable Usage

Gemini 4 Carbon is based on Gemini 4 Argon. For more information about the acceptable usage for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Safety & Security

Please see the full Safety & Security Report for Gemini 4 Carbon here.

Evaluation Approach

For more information about the evaluation approach for Gemini 4 Carbon, see the Gemini 4 Argon model card.

Safety Policies

For more information about the safety policies for Gemini 4 Carbon, see the Gemini 4 Argon model card.

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 Carbon outperforms Gemini 4 Argon across multilingual and multimodal content safety, while performing similarly on unjustified refusals and tone. We mark improvements in bolded green.

See the Gemini 4 Carbon Safety & Security Report for details.

Evaluation	Description	Gemini 4 Carbon vs. Gemini 4 Argon *(percentage point increase/decrease)*
Text to Text Safety *(higher is better)*	Automated content safety evaluation measuring adherence to safety policies across text prompts	**+3.50pp**
Multilingual Safety *(higher is better)*	Automated safety policy evaluation across multiple languages	**+3.55pp**
Image to Text Safety *(higher is better)*	Automated content safety evaluation measuring adherence to safety policies with both images and text as inputs	-1.00pp
Tone *(higher is better)*	Automated evaluation measuring objective, neutral, and non-preachy tone of model refusals	+0.10pp
Unjustified-refusals *(lower is better)*	Automated evaluation measuring over-refusals on benign or borderline prompts that can be safely fulfilled	+0.10pp

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. The performance results reported above are computed with improved evaluations and thus are not directly comparable with performance results found in previous Gemini model cards.

We expect variation in our automated safety evaluations results, which is why we review flagged content to check for egregious or dangerous material. Our manual review confirmed losses were overwhelmingly either a) false positives or b) not egregious.

Human Red Teaming Results

We conduct manual red teaming by specialist teams who sit outside of the model development team. High-level findings are fed back to the model team. For child safety evaluations, Gemini 4 Carbon satisfied required launch thresholds, which were developed by expert teams to protect children online and meet Google's commitments to child safety across our models and Google products. For content safety policies generally, including child safety, we saw similar or improved safety performance compared to Gemini 4 Argon. Additionally, the scope of red teaming covered potential issues outside of our strict policies, compared performance to Gemini 4 Argon, and found no egregious concerns.

Frontier Safety Assessment

For more information about our frontier safety assessment, see the Gemini 4 Carbon Safety & Security Report.

Risks and Mitigations

Safety and security were built into Gemini 4 Carbon 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, and provide specialized dual-use cybersecurity capabilities to verified defenders through the Fairwind Program. 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 Carbon Safety & Security Report.