



Model Card: Grok 4.5

July 14, 2026

Contents

1	Introduction	3
1.1	Capabilities and intended use	3
1.2	Model development and training	4
2	Coding abilities	5
2.1	DeepSWE v1.0	5
2.2	DeepSWE v1.1	6
2.3	APEX-SWE	6
2.4	SWE-Bench Pro	7
2.5	SWE-Bench Multilingual	7
2.6	SWE-Marathon	8
2.7	FrontierSWE	8
2.8	ProgramBench	9
2.9	Terminal-Bench 2.1	9
2.10	SWE-Atlas-QnA	10
2.11	FalseClaimBench	11
3	Engineering acceleration	12
3.1	EEBench	12
3.2	3DCodeBench	13
3.3	CAD-Bench	13
4	Office-use abilities	14
4.1	GDPval-AA v2	14
4.2	τ -bench	15
5	R&D Enablement	16
5.1	SpaceXAI MTS Eval	16
5.2	RelBench	17

6 Search capabilities and factuality	18
6.1 Factuality	18
6.2 DeepSearchQA	18
7 Cyber capabilities and safeguards	19
7.1 CyberGym	19
7.2 HackerBench v0.2	20
8 Biological and chemical capabilities and safeguards	21
8.1 Virology Capabilities Test (VCT)	21
8.2 WMDP dual-use knowledge (MCQ)	21
8.3 LAB-Bench practical MCQ	22
8.4 ProtocolQA Open-Ended	22
8.5 BixBench zero-shot MCQ	22
9 Jailbreaks and robustness	23
9.1 Jailbreaks	23
10 General output safety	23
10.1 General refusals	24
10.2 Child safety	24
10.3 CBRN / weapons refusals	24
11 Mental health	25
11.1 Self-harm refusals	25
12 Behaviors	25
12.1 Epistemic Bias	26
12.2 MASK-Rectified	26
12.3 Sycophancy	26
References	27

1 Introduction

Grok 4.5 is the initial release of the newest family of models from SpaceXAI* and Cursor[†].

It is our most intelligent model yet and is built for maximum real-world utility: demonstrating frontier capabilities in coding, engineering, design, and professional workflows, while also being served with the highest token efficiency of any model of its intelligence class.

1.1 Capabilities and intended use

Grok 4.5 is highly agentic and reasoning-efficient, being capable of autonomously solving larger and more difficult tasks than any of our previous models while keeping its user in control, and achieving results in half as many steps as other frontier models.

Grok 4.5 implements high-precision and non-intrusive safeguards and mitigations.

We document safety areas (cyber, bio knowledge, bio agentic, jailbreaks/robustness, general output safety including vision and CBRN refusals, mental health, and behaviors) and our safeguards below.

Each evaluation under a capability or safety section has its own subsection. Unless stated otherwise, evaluation results are on the final deployed checkpoint of Grok 4.5.

Third-party public benchmarks are cited in the References section and marked with superscript citations (e.g. ¹) on first mention. Internal evaluations are described in place and are not listed in References.

We do not restrict use of our model in any legitimate use case, nor do we reserve certain use cases of our model as uncompetitive. We never silently downgrade intelligence or fall back to other models.

Our goal is to preserve legitimate use cases of the model on humanity's most pressing issues, including accelerating engineering and creative work across the entire economy, helping discover new scientific breakthroughs, hardening critical infrastructure, and enabling AI research.

The primary purpose of this model card is to detail the capabilities – that is, the strengths and limitations – of Grok 4.5 in quantitative and objective terms, to inform where this model is most useful.

We additionally report qualitative characteristics and uses of the model, including ways we and other organizations have found the model most useful.

Grok 4.5 operates primarily through a text-based modality. Its inputs are in natural language text format and images (for vision/understanding).

* SpaceXAI is a doing-business-as (dba) name of XAI LLC. xAI and SpaceXAI may be used interchangeably throughout this card.

† Grok 4.5 was also subject to supplemental training using anonymized Cursor workflow data to improve coding and agentic performance.

The model's use is subject to [SpaceXAI's Acceptable Use Policy](#)[†], applicable Consumer and Enterprise Terms of Service, and any applicable laws.

Grok 4.5 is not intended for autonomous high-stakes decision-making in domains such as medicine, law, finance, or safety-critical systems without appropriate human oversight and domain-expert validation.

Grok 4.5 is available for use through the following channels:

- SpaceXAI API: Available from the console at console.x.ai; API users can call Grok 4.5 through the standard chat and completions endpoints.
- Grok Build: The default model in SpaceXAI's terminal-based coding agent, available through both the API and the CLI.
- Cursor: Available to all users, on every plan tier.
- Office add-ins: The default model in the add-ins for Microsoft Word, PowerPoint, and Excel.
- Model gateways: Reachable through OpenRouter, Vercel, Cloudflare, Snowflake, and Databricks Mosaic, and more.

SpaceXAI plans to add Grok 4.5 to its Consumer platforms, such as SpaceXAI's consumer web and mobile consumer apps, and the X Platform, through Grok-in-X and other features, at a later date.

1.2 Model development and training

Grok 4.5 was pretrained on publicly available data, data generated internally, as well as other data for which SpaceXAI has secured the necessary rights to, followed by targeted midtraining and post-training with supervised finetuning and reinforcement learning on human and synthetic reward signals.

Grok 4.5 has a pretraining cutoff of January 2026.

[†] SpaceXAI Acceptable Use Policy: <https://x.ai/legal/acceptable-use-policy>.


2 Coding abilities

Coding underlies most of Grok 4.5’s agentic work: the skills that fix a repository issue also enable tool use, research loops, and office automation.

We primarily evaluate against agentic evaluations measuring real-world software engineering and problem-solving: Our benchmarks include repository-level issue resolution (SWE-Bench Pro ², SWE-Bench Multilingual ³, DeepSWE ¹, APEX-SWE ⁴, SWE-Marathon ⁵, FrontierSWE ⁶), program reconstruction (ProgramBench ⁷), terminal use (Terminal-Bench ⁸), and other engineering-relevant use cases.

2.1 DeepSWE v1.0 ¹


DeepSWE* measures end-to-end software-engineering agents on contamination-resistant repository issues: the agent must read the codebase, edit files, run commands, and produce a patch that passes behavioral and correctness verifiers.



* Eval created by Datacurve, run with each model provider’s harnesses by AA


2.2 DeepSWE v1.1[†]

DeepSWE v1.1^{*} is the updated release of the DeepSWE agentic-coding benchmark, measuring end-to-end resolution of contamination-resistant repository issues. We report these results in addition to DeepSWE v1.0.



2.3 APEX-SWE[‡]

APEX-SWE measures AI productivity on realistic software-engineering work spanning integration tasks and observability tasks (diagnosing failures from telemetry), being representative of a wide set of engineering use cases.[†]




* mini-swe-agent harness run by Datacurve

† Results are taken from tests conducted by Mercor.

2.4 SWE-Bench Pro²

SWE-Bench Pro^{*} tests long-horizon software engineering on hard, multi-file issues from actively maintained repositories, with reduced public ground-truth leakage relative to classic SWE-bench.

An agent works in a repository environment to produce a patch graded by tests and a verification reward. We report the standard resolution metric (e.g., verification reward or pass rate) under a fixed agent scaffold.



2.5 SWE-Bench Multilingual³

SWE-Bench Multilingual extends repository-level issue resolution across multiple programming languages using the same agent-in-a-repo paradigm as other SWE suites. Patches are graded by language-appropriate tests in isolated environments.

We report resolution / pass rate on the release-tracked multilingual set under a fixed agent scaffold.[†]




^{*} SWE-Bench Pro was evaluated with controls for reward hacking in place; see <https://cursor.com/blog/reward-hacking-coding-benchmarks>

[†] SWE-Bench Multilingual results are taken from runs additionally reported by Cursor; see <https://cursor.com/grok-4-5>.

2.6 SWE-Marathon⁵

SWE-Marathon* targets ultra-long-horizon engineering work that can require millions of tokens and multi-hour trajectories, far beyond a single-PR bugfix.


The primary metric is task resolution rate under multi-layer verification designed to resist reward hacking and shortcuts. Grok 4.5 is exceptionally capable in solving long-horizon engineering tasks, beating all other frontier models tested.



2.7 FrontierSWE⁶

FrontierSWE evaluates coding agents on ultra-long-horizon technical challenges at or beyond strong human-expert scope.[†] Tasks are real problems curated with industry and academic partners: agents are allotted up to 20 hours per task, and scoring uses continuous partial credit (as opposed to binary issue resolution).

We report **Dominance** (win rate versus a random baseline across tasks) as the primary metric; higher is better. Unlike SWE-Bench-style pass rates, Dominance summarizes comparative strength when absolute full-solve rates remain low for all models.




* Scores reflect the full SWE-Marathon task set (swe-marathon.org).

[†] Results and harness labels from the public FrontierSWE leaderboard; Grok 4.5 is evaluated with Grok CLI.

2.8 ProgramBench ⁷

ProgramBench tasks agents with rebuilding program behavior from a compiled binary and documentation alone (no source, no decompilation, no internet), testing feature exactness and long-horizon capabilities.

Submissions are graded with large suites of execution-based behavioral tests. We report the tests-passed rate; full resolution remains intentionally hard (as of the publication date of this card, no tested model scores above 0.5% in full resolution).



2.9 Terminal-Bench 2.1 ⁸

Terminal-Bench 2.1 measures agents on challenging and realistic command-line tasks (sysadmin, coding, data, and security-style workflows) inside containerized terminal environments.

We evaluate within the Grok Build harness and score task success and mean reward across verified tasks. Higher reward implies a model with more reliable terminal agency, including long-running commands and multi-step debugging.*




* Competitor figures are drawn from the respective developers' published system cards or benchmark leaderboards.

2.10 SWE-Atlas-QnA ⁹

SWE-Atlas-QnA tests software-engineering repository question answering: the model must answer questions about a codebase using reading and exploration tools rather than generating patches.


Answers are graded for correctness against repository ground truth; we report accuracy on the tracked Atlas Q&A split.



2.11 FalseClaimBench

The SpaceXAI internal false-claim eval suite (FalseClaimBench) assesses whether the agent reports having done work it never actually performed, such as edits, fixes, or commands.

We compare the agent’s stated claims against the true final state of the workspace and score the percentage of fully true claims/reports; higher is better.




3 Engineering acceleration

Beyond software engineering, we evaluate Grok 4.5 on agentic electrical engineering, procedural 3D modeling, and parametric CAD tasks that measure how effectively models accelerate physical-world engineering workflows.

3.1 EEBench¹⁰

EEBench is an electrical engineering and chip design benchmark: models design circuits and other hardware that are graded against physical correctness and functionality.*


We report score against average output tokens per task (higher and further left is better), comparing Grok 4.5 to Fable 5, Opus 4.8, Gemini 3.1 Pro, GPT-5.5, and GPT-5.6 Sol.



* Public leaderboard results from <https://eebench.org/> (V1 core corpus).


3.2 3DCodeBench[†]

3DCodeBench evaluates agentic procedural 3D modeling via code: vision-language agents author engine-ready 3D assets through software APIs and geometric reasoning, graded for executability and shape fidelity.*



3.3 CAD-Bench[‡]

CAD-Bench (Parametric CAD Bench) measures CAD models and agents on parametric design and 3D modeling tasks spanning out-of-distribution CAD generation workloads.[‡]



* <https://arxiv.org/abs/2606.01057>.

† <https://www.gnucleus.ai/cad-bench>.


4 Office-use abilities

We evaluate performance on professional knowledge-work and structured agent tasks relevant to office and productivity use cases.

4.1 GDPval-AA v2¹³

GDPval-AA v2^{*} evaluates models on economically valuable knowledge-work deliverables (documents, analyses, and professional artifacts) spanning occupations that contribute substantially to GDP.


We specifically track Artificial Analysis’ GDPval-AA v2, an agentic harness over GDPval-style tasks with pairwise quality ratings. Scores reflect end-to-end deliverable quality under the stated harness.




^{*} Run by Artificial Analysis (AA) in their harness.

4.2 τ -bench¹⁴

τ -bench is a multi-turn tool-use evaluation: agents interact with a user and tools under domain policies against a backend database, and success is judged by final database state correctness and goal completion. We report Artificial Analysis (AA) results on τ^3 -banking¹⁵, a τ -bench-derived evaluation focused on banking workflows.



We also report accuracy against output tokens per task (higher and further left is better):



5 R&D Enablement

We evaluate Grok 4.5 on its ability to automate parts of the engineering and research process for training and evaluating new versions of itself.


5.1 SpaceXAI MTS Eval

The benchmark evaluates whether AI agents can meaningfully enable and accelerate AI R&D: for example, by diagnosing reward hacking in training runs, generating and auditing high-quality training data, debugging large-scale training infrastructure, and creating new evaluations for emerging model capabilities.

The SpaceXAI MTS eval is an internal benchmark of frontier model-development tasks faced by SpaceXAI engineers (e.g., diagnosing reward hacking, auditing training data, debugging training infrastructure, and building new evals).

Agents are scored on task reward under a fixed rollout budget (time and tokens). Higher mean reward indicates stronger AI R&D enablement; refusals are noted.*†

We report % correct against mean tokens per task (higher and further left is better):




* GPT 5.6 Sol refused to solve 2 out of the 29 tasks. GPT 5.5 refused to solve 1 out of the 29 tasks. The scores are computed excluding these tasks that were refused.

† The results are obtained on Grok Build harness.

5.2 RelBench¹⁶

RelBench evaluates machine learning over multi-table relational databases: forecasting, recommendation, and related tasks that require modeling entities and relationships rather than a single flat table.

Agents or models produce predictions graded against held-out relational targets. We report the release-tracked RelBench aggregate score under the fixed task set and resource limits.




6 Search capabilities and factuality

We evaluate grounded answering with search tools and factual reliability, including hallucination-prone settings.

6.1 Factuality

Single-turn hallucination measures how often the model asserts unsupported or fabricated claims in a single response to an information-seeking query.

A separate grader flags factually unsupported statements against retrieved or known ground truth; we report the hallucination rate, so lower is better.



6.2 DeepSearchQA ¹⁷

DeepSearchQA evaluates end-to-end answer accuracy on questions that require multi-step search and synthesis of retrieved information. Answers are graded for correctness against reference answers; we report accuracy, so higher is better.*



* Evaluation was run on an internal implementation of DeepSearchQA.

7 Cyber capabilities and safeguards

Grok 4.5 exhibits enhanced cybersecurity-relevant capabilities. We report capabilities and safeguard behavior as separate numbers, because the deployed configuration refuses the large majority of clearly harmful cyber requests.


The capability is most useful to defenders, for finding and fixing vulnerabilities rather than carrying out end-to-end attacks.

Cyber evaluations measure offensive and defensive security-relevant capabilities, plus calibration of cyber safety controls.

7.1 CyberGym¹⁸

CyberGym measures offensive security capability by tasking an agent with reproducing crashes and assembling working exploits for known vulnerabilities.


Because it is a capability probe rather than a refusal test, we report the unrestricted (i.e. unmitigated by our standard safeguards) score as a pure measure of ability; refusal behavior on harmful cyber requests is measured separately (see HackerBench). Other models’ results are taken from the respective model providers’ system cards^{19, 20, 21, 22, 23}, using the unsafeguarded figures where available to avoid loss of signal due to refusals.



7.2 HackerBench v0.2

HackerBench is a SpaceXAI internal CTF-style suite of cybersecurity and red-teaming tasks that require chaining skills (recon, exploitation primitives, and reasoning about vulnerable systems) in a constrained environment. Included within HackerBench is a subset of harmful-intent, borderline, and dual-use requests that should be refused, as well as benign tasks.

Scores are reported with the standard release-tracked safeguards for Grok 4.5 and the standard provider-supplied safeguards for other tested models. The eval is scored by the harmful/dual-use compliance rate and the benign refusal rate.



8 Biological and chemical capabilities and safeguards

These suites probe dual-use biology and chemistry knowledge and lab-protocol understanding, and are treated as safety-threshold evaluations.

Grok 4.5 shows solid but sub-threshold dual-use knowledge, indicating limited actionable uplift for an already-trained actor. Weapons-path refusals are reported separately under general output safety.

Unless otherwise stated, the benchmarks presented in Section 8 are performed without any safeguards in place; this enables full measurement of Grok 4.5's dual-use capabilities when unrestricted.

8.1 Virology Capabilities Test (VCT) ²⁴

VCT presents multimodal virology problems and measures whether the model can troubleshoot experimental lab protocols. It is a dual-use knowledge probe, treated as a CBRN safety-threshold evaluation, so a higher score indicates more dual-use capability rather than a better outcome.

Metric	Grok 4.5 (high)
Accuracy	65.5%

8.2 WMDP dual-use knowledge (MCQ) ²⁵

WMDP-Bio and WMDP-Chem are multiple-choice suites probing operationally sensitive dual-use knowledge in biology and chemistry.

Metric	Grok 4.5 (high)
WMDP-Bio accuracy	90.9%
WMDP-Chem accuracy	87.3%

8.3 LAB-Bench²⁶ practical MCQ

LAB-Bench practical items test everyday wet-lab and research skills (protocol, sequence, and cloning reasoning).

Metric	Grok 4.5 (high)
Accuracy	71.1%

We evaluate open-ended and tool-using biology tasks rather than closed knowledge questions.

These serve as non-malicious capability indicators and review triggers: strong performance (ProtocolQA 87.0%, BixBench²⁷ 93.8%) signals broad scientific competence, not evidence that the model provides weapons-enabling assistance, which remains gated by the refusal safeguards.

8.4 ProtocolQA Open-Ended²⁶

ProtocolQA open-ended asks the model to identify the single most important mistake in a described biological lab protocol. Unlike the multiple-choice suites, this benchmark is comprised of open-ended troubleshooting tasks.

Metric	Grok 4.5 (high)
Accuracy	87.0%

8.5 BixBench zero-shot MCQ²⁷

BixBench evaluates computational-biology reasoning with analysis tools available by default, scored zero-shot.

Metric	Grok 4.5 (high)
Accuracy	93.8%

BixBench items contain no weapons-enabling content; the suite is a computational-biology capability indicator under Grok Build's default tool-on setting, not a CBRN assistance eval.

We do not claim that tool-on BixBench success implies unrestricted agentic bio assistance; weapons-path and high-risk lab-protocol help remain gated by those refusals/filters, which are scored separately from the metrics scored in BixBench.

9 Jailbreaks and robustness

We test whether the refusal safeguards hold under intense adversarial pressure, using a broad and continuously updated set of jailbreaks. Grok 4.5 stays robust, incorrectly complying with only 0.73% of should-refuse prompts under attack.

We treat residual failures as cases for ongoing monitoring and patching as new techniques emerge.

9.1 Jailbreaks

We stress the refusal safeguards with a broad, continuously updated set of jailbreak attacks, including recent, state-of-the-art, and newly observed techniques, across single-turn and multi-turn settings and attacks embedded in either the user or the system message.

The metric is compliance rate on should-refuse prompts under attack.

Metric	Grok 4.5 (high)
Compliance ↓	0.73%

10 General output safety

This section reports policy refusals for consumer chat-distribution prompts across the major disallowed categories, including multimodal content and first-class CBRN and weapons refusals.

Compliance on should-refuse prompts is low across the board (0.0% for CSAM, 1.1% for the broad disallowed suite), and CBRN refusal accuracy is high (96.7 to 97.9%). Paired benign prompts guard against over-refusal.

The below evaluations and results also apply to Grok Build.

Refusals are enforced by a layered, defense-in-depth stack rather than any single filter.

Safety fine-tuning and post-training (supervised fine-tuning plus reinforcement learning from human feedback, verifiable rewards, and model-based grading) train the model to refuse requests that show clear intent to cause severe harm or engage in criminal activity.

System prompts steer the model toward honesty and maximal truth-seeking while avoiding over-refusal on benign or hypothetical discussions.

Additionally, Grok 4.5 may be augmented with runtime input and topical filters that add controls for classes of severe harm, including CSAM, self-harm, and biological/chemical weapons pathways, alongside cyber-specific input safety controls.

10.1 General refusals

We assemble harmful queries that show intent (ranging from clear to well-hidden) to obtain help with harm or disallowed behaviors across major policy categories, translated across English, Spanish, Chinese, Japanese, Arabic, and Russian, and use a separate grader model to judge whether the model correctly refuses.

The primary metric is compliance rate on should-refuse prompts.

Metric	Grok 4.5 (high)
Compliance ↓	1.1%

10.2 Child safety

The CSAM and child safety benchmark suite applies the shared refusal protocol to the highest-severity child-safety category. On this category, Grok 4.5 performs well, not complying with any harmful or borderline-intent prompts.

Metric	Grok 4.5 (high)
Compliance ↓	0.0%

10.3 CBRN / weapons refusals ²⁸

Autointent-Bio and Autointent-Chem are internal suites of manually and synthetically generated dangerous and benign CBRN queries; FORTRESS-RN ²⁸ covers radiological and nuclear items. We report refusal accuracy on the dangerous queries with the full suite of safeguards enabled.

Metric	Grok 4.5 (high)
Bio refusal accuracy	97.9%
Chem refusal accuracy	96.7%
R/N refusal accuracy	97.9%

CBRN pathways receive heightened scrutiny across the safeguard stack: the refusal policy prioritizes non-assistance for biological, chemical, radiological, or nuclear weapons development or deployment; system prompts give special attention to CBRN misuse pathways; and dedicated input and topical filters target biological and chemical weapons-related abuse.

11 Mental health

We assess how Grok 4.5 handles self-harm and crisis situations, checking that the model declines to assist with self-harm while remaining supportive and directing users toward help.

The model fails to refuse or redirect the user to assistance on only 0.5% of should-refuse self-harm prompts, consistent with the broader refusal results.

11.1 Self-harm refusals

The self-harm suite applies the shared refusal protocol to self-harm and crisis prompts and longer-horizon conversations, checking that the model declines to assist while remaining supportive. The metric is the compliance (i.e. non-refusal/redirection) rate on should-refuse conversations.

Metric	Grok 4.5 (high)
Compliance ↓	0.5%

12 Behaviors

These evaluations cover propensities that affect reliability, neutrality, and controllability rather than disallowed content: sycophancy, epistemic bias, and honesty under pressure. Grok 4.5 shows very low sycophancy (0.01%) and contained bias (20.4%).

Low dishonesty under adversarial MASK ²⁹ conditions (0.67%) is tracked as a release indicator rather than a pass/fail result.

12.1 Epistemic Bias

Epistemic Bias measures whether the model frames factual answers more favorably toward one side of a contentious issue, or expresses opinions where it should stay neutral, by comparing its responses across opposing framings of the same question.

Metric	Grok 4.5 (high)
Bias rate ↓	20.4%

12.2 MASK-Rectified ²⁹

Using a dataset derived from MASK^{*}, we test whether Grok 4.5 faithfully reports its beliefs when pressured to lie, as a corollary for the model's proclivity to noncritically surface misleading or harmful information.

Metric	Grok 4.5 (high)
Dishonesty ↓	0.67%

12.3 Sycophancy

Sycophancy measures the tendency to abandon a correct answer and agree with a user's confidently stated wrong one. We use an internal benchmark that presents Grok 4.5 with a question alongside misleading user-supplied context, and report the average drop in accuracy relative to the neutral prompt.

Metric	Grok 4.5 (high)
Sycophancy ↓	0.01%

* MASK-Rectified corrects the MASK grading so that responses where the model is obviously (model-aware) role-playing, rather than asserting a genuine belief, are not counted as lies.

References

Public benchmarks and datasets referenced above. Superscripts mark first mention in the main text and link to a primary source. Internal evaluations are not listed.

1. DeepSWE. Huang, Lee, Tng, and Ge (Datacurve), 2026. <https://deepswe.datacurve.ai/> · <https://arxiv.org/abs/2607.07946>
2. SWE-Bench Pro. Deng, Da, et al. (Scale AI), 2025. https://labs.scale.com/leaderboard/swe_bench_pro_public · <https://arxiv.org/abs/2509.16941>
3. SWE-Bench Multilingual. Khandpur, Lieret, Jimenez, Press, and Yang (SWE-bench), 2025. <https://www.swebench.com/multilingual.html>
4. APEX-SWE. Kottamasu et al. (Mercor), 2026. <https://www.mercor.com/apex/apex-swe-leaderboard/> · <https://arxiv.org/abs/2601.08806>
5. SWE-Marathon. Desai et al. (Abundant AI), 2026. <https://www.swe-marathon.org/> · <https://arxiv.org/abs/2606.07682>
6. FrontierSWE. Chu, Agarwal, et al. (Proximal), 2026. <https://www.frontierswe.com/>
7. ProgramBench. Yang, Lieret, et al., 2026. <https://programbench.com/> · <https://arxiv.org/abs/2605.03546>
8. Terminal-Bench 2.1. Merrill et al., 2026. <https://www.tbench.ai/> · <https://arxiv.org/abs/2601.11868>
9. SWE-Atlas-QnA. Artificial Analysis, 2026. <https://artificialanalysis.ai/agents/coding-agents?coding-agents-performance-chart=swe-atlas-qna>
10. EEBench. atopile, 2026. <https://eebench.org/>
11. 3DCodeBench. Gao, Shu, Ye, Xiong, Makadia, Guo, Itti, and Chen, 2026. <https://arxiv.org/abs/2606.01057>
12. Parametric CAD Bench (CAD-Bench). gNucleus. <https://www.gnucleus.ai/cad-bench>
13. GDPval-AA v2. Artificial Analysis. <https://artificialanalysis.ai/evaluations/gdpval-aa>
14. τ -bench. Yao, Shinn, Razavi, and Narasimhan (Sierra), 2024. <https://taubench.com/> · <https://arxiv.org/abs/2406.12045>
15. τ^3 -banking. Artificial Analysis. <https://artificialanalysis.ai/evaluations/tau3-banking>
16. RelBench. Robinson et al. (Stanford), 2024. <https://relbench.stanford.edu> · <https://arxiv.org/abs/2407.20060>
17. DeepSearchQA. Gupta, Chatterjee, et al., 2026. <https://arxiv.org/abs/2601.20975>
18. CyberGym. Wang, Shi, He, Cai, Zhang, and Song, 2025. <https://www.cybergym.io/> · <https://arxiv.org/abs/2506.02548>
19. Claude Fable 5 & Mythos 5 System Card. Anthropic, 2026. <https://www.anthropic.com/claude-fable-5-mythos-5-system-card>
20. GPT-5.6 System Card. OpenAI, 2026. <https://deploymentsafety.openai.com/gpt-5-6>
21. Claude Opus 4.8 System Card. Anthropic, 2026. <https://www.anthropic.com/claude-opus-4-8-system-card>
22. GPT-5.5 System Card. OpenAI, 2026. <https://deploymentsafety.openai.com/gpt-5-5>
23. Claude Opus 4.7 System Card. Anthropic, 2026. <https://www.anthropic.com/claude-opus-4-7-system-card>

24. VCT. Götting, Medeiros, Sanders, Li, Phan, Elabd, Justen, Hendrycks, and Donoughe, 2025. <https://arxiv.org/abs/2504.16137>
25. WMDP. Li et al., 2024. <https://wmdp.ai> · <https://arxiv.org/abs/2403.03218>
26. LAB-Bench. Laurent et al., 2024. <https://github.com/Future-House/LAB-Bench> · <https://arxiv.org/abs/2407.10362>
27. BixBench. Mitchener, Laurent, et al., 2025. <https://github.com/Future-House/BixBench> · <https://arxiv.org/abs/2503.00096>
28. FORTRESS. Knight, Deshpande, et al. (Scale AI), 2025. <https://labs.scale.com/leaderboard/fortress> · <https://arxiv.org/abs/2506.14922>
29. MASK. Ren et al., 2025. <https://www.mask-benchmark.ai/> · <https://arxiv.org/abs/2503.03750>