Google

Model card published/updated AUG 01, 2025

Gemini 2.5 Deep Think Model Card

Gemini 2.5 Deep Think - 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.*

Technical Reports *are similar to academic papers, and describe models' capabilities, limitations and performance benchmarks. The Gemini 2.5 technical report contains additional details about the Gemini 2.5 series of models that are generally available. We recommend that readers seeking more details and information about these models navigate to the technical report.*

Published: August 1, 2025

Model Information

Description: Gemini 2.5 Deep Think is an enhanced reasoning model that is part of our Gemini 2.5 family that uses parallel thinking and reinforcement learning to test multiple hypotheses at once.

Inputs: Text strings (e.g., a question, a prompt, document(s) to be summarized), images, audio, and video files, with a 1M token context window.

Outputs: Text, with a 192K token output.

Architecture The Gemini 2.5 models are sparse mixture-of-experts (MoE) (Clark et al., 2022; Du et al., 2021; Fedus et al., 2021; Jiang et al., 2024, Lepikhin et al., 2020; Riquelme et al., 2021; Roller et al., 2021; Shazeer et al., 2017; transformers (Vaswani et al., 2017) with native multimodal support for text, vision, and audio inputs. Sparse MoE models activate a subset of model parameters per input token by learning to dynamically route tokens to a subset of parameters (experts); this allows them to decouple total model capacity from computation and serving cost per token. Developments to the model architecture contribute to the significantly improved performance of Gemini 2.5 compared to Gemini 1.5 Pro (see Section 3 of the Gemini Technical Report).

Model Data

Training Dataset: The pre-training dataset was a large-scale, diverse collection of data encompassing a wide range of domains and modalities, which included publicly-available web-documents, code (various programming languages), images, audio (including speech and other audio types) and video. The post-training dataset consisted of vetted instruction tuning data and was a collection of multimodal data with paired instructions and responses in addition to human preference and tool-use data. We additionally trained Gemini 2.5 Deep Think on novel reinforcement learning techniques that can leverage more multi-step reasoning, problem-solving and theorem-proving data, and we also provided access to a curated corpus of high-quality solutions to mathematics problems.

Training Data Processing: Data filtering and preprocessing included techniques such as deduplication, safety filtering in-line with Google's commitment to advancing AI safely and responsibly and quality filtering to mitigate risks and improve training data reliability.

Implementation and Sustainability

Hardware: Gemini 2.5 Deep Think was trained using Google's Tensor Processing Units (TPUs). TPUs are specifically designed to handle the massive computations involved in training LLMs and can speed up training considerably compared to CPUs. TPUs often come with large amounts of high-bandwidth memory, allowing for the handling of large models and batch sizes during training, which can lead to better model quality. TPU Pods (large clusters of TPUs) also provide a scalable solution for handling the growing complexity of large foundation models. Training can be distributed across multiple TPU devices for faster and more efficient processing.

The efficiencies gained through the use of TPUs are aligned with Google's commitment to operate sustainably.

Software: Training was done using JAX and ML Pathways.

Evaluation

Approach: Gemini 2.5 Deep Think was evaluated using the methodology below:

- **Model selection:** To ensure fair comparison, we only compare results for models without additional tool calls enabled.
- **Gemini results**: All Gemini scores are sampled from the Gemini App. To reduce variance, we average over multiple trials for smaller benchmarks. Gemini 2.5 Deep Think IMO 2025 results are computed as pass@1 while all the other results coming from matharena.ai are best of 32.
- **Result sources**: Where provider numbers are not available we report numbers from leaderboards reporting results on these benchmarks: Humanity's Last Exam results are sourced from https://agi.safe.ai/ and https://scale.com/leaderboard/humanitys_last_exam, LiveCodeBench results are from https://livecodebench.github.io/leaderboard.html (1/1/2025 - 5/1/2025 in the UI), IMO 2025 from https://matharena.ai.IMO 2025 Model grades are based on the cutoffs from the 2025 competition. For Grok 4 IMO 2025 results we show the highest result available from matharena.ai, with a custom prompt.

Results: Gemini 2.5 Deep Think strongly performs across a range of benchmarks that measure coding, science, knowledge and reasoning capabilities. Results as of July 2025 are listed below:

Capability **Benchmark**[1]	**Gemini 2.5 Pro**	**Gemini 2.5 Deep Think**	**OpenAI o3**	**Grok 4**
Reasoning & knowledge **Humanity's Last Exam (no tools)**	21.6%	34.8%	20.3%	25.4%
Mathematics **IMO 2025**	31.6% (No medal)	60.7% (Bronze medal grade)	16.7% (No medal)	21.4% (No medal)
Mathematics **AIME 2025**	88.0%	99.2%	88.9%	91.7%
Code generation **LiveCodeBench v6** (UI: 1/1/2025-5/1/2025)	74.2%	87.6%	72.0%	79.0%

[1]We regularly update evaluation processes to include new and emerging quality evaluations and benchmarks. The results reported above include additional or updated benchmarks which may not have been included in previous Gemini model cards. Results are thus not directly comparable with performance results found in previous Gemini model cards.

Intended Usage and Limitations

Benefit and Intended Usage: Gemini 2.5 Deep Think can help solve problems that require creativity, strategic planning and making improvements step-by-step, such as:

- Iterative development and design
- Scientific and mathematical discovery
- Algorithmic development and code

Known Limitations: Gemini 2.5 Deep Think may exhibit some of the general limitations of foundation models, such as hallucinations. There may also be occasional slowness or timeout issues. The knowledge cutoff date for Gemini 2.5 Deep Think was January 2025. See the Ethics and Safety section below for additional information on known limitations.

Ethics and Content Safety

Evaluation Approach: Gemini 2.5 Deep Think was developed in partnership with internal safety, security, and responsibility teams. A range of evaluations and red teaming activities were conducted to help improve the model and inform decision-making. These evaluations and activities align with Google's AI Principles and responsible AI approach.

Evaluation types included but were not limited to:

- **Training/Development Evaluations** including automated and human evaluations carried out continuously throughout and after the model's training, to monitor its progress and performance;
- **Human Red Teaming** conducted by specialist teams across the policies and desiderata, deliberately trying to spot weaknesses and ensure the model adheres to safety policies and desired outcomes;
- **Automated Red Teaming** to dynamically evaluate Gemini for safety and security considerations at scale, complementing human red teaming and static evaluations;
- **Assurance Evaluations** conducted by evaluators who sit outside of the model development team, used to independently assess responsibility and safety governance decisions;
- **Google DeepMind Responsibility and Safety Council (RSC)**, Google DeepMind's internal governance body, reviewed the initial ethics and safety assessments on novel model capabilities in order to provide feedback and guidance during model development. The RSC also reviewed data on the model's performance via assurance evaluations and made release decisions.

In addition, we perform testing following the guidelines in Google DeepMind's Frontier Safety Framework (FSF).

Safety Policies: Gemini safety policies align with Google's standard framework for the types of harmful content that we make best efforts to prevent our Generative AI models from generating, including the following types of harmful content:

1. Child sexual abuse and exploitation
2. Hate speech (e.g., dehumanizing members of protected groups)
3. Dangerous content (e.g., promoting suicide, or instructing in activities that could cause real-world harm)
4. Harassment (e.g., encouraging violence against people)
5. Sexually explicit content
6. Medical advice that runs contrary to scientific or medical consensus

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 2.5 Deep Think outperforms Gemini 2.5 Pro on content safety and tone. However, the model sometimes over-refuses benign queries, which we classify as instruction following losses.

Evaluation[2]	Description	Gemini 2.5 Deep Think vs. Gemini 2.5 Pro
Text to Text Safety	Automated content safety evaluation measuring safety policies	-16.3%
Multilingual Safety	Automated safety policy evaluation across multiple languages	-1.0%
Image to Text Safety	Automated content safety evaluation measuring safety policies	+2.1% (non egregious)
Tone[3]	Automated evaluation measuring objective tone of model refusal	+16.3%
Instruction Following	Automated evaluation measuring model's ability to follow instructions while remaining safe	-9.9%

[2]The ordering of evaluations in this table has changed from previous iterations of the 2.5 Flash-Lite model card in order to list safety evaluations together and improve readability. The type of evaluations listed have remained the same.

[3] For tone and instruction following, a positive percentage increase represents an improvement in the tone of the model on sensitive topics and the model's ability to follow instructions while remaining safe compared to Gemini 2.5 Pro. We mark improvements in green and regressions in red.

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 below are computed with improved evaluations and thus are not directly comparable with performance results found in previous Gemini model cards. In addition to continuing to improve our evaluations, we also run Assurance Evaluations which are independent evaluations to assess the safety profile of our models (see below section).

Assurance Evaluations Results: We conduct baseline assurance evaluations to guide decisions on model releases. These evaluations look at model behavior, including within the context of the safety policies and modality-specific risk areas. High-level findings are fed back to the model team, but prompt sets are held out to prevent overfitting and preserve the results' ability to inform decision making. For child safety evaluations, we continue to see the Gemini 2.5 family of models meeting or improving upon 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 2.5 Pro.

Known Safety Limitations: The main content safety limitations for Gemini 2.5 Deep Think are related to instruction following. The model occasionally over-refuses user requests, when intended behavior is the model fulfilling as much as possible without violating policy.

Risks and Mitigations: Safety and responsibility was built into Gemini 2.5 Deep Think throughout the training and deployment lifecycle, including pre-training, post-training, and product-level mitigations. Mitigations include, but are not limited to:

- dataset filtering;
- conditional pre-training;
- supervised fine-tuning;
- reinforcement learning from human and critic feedback;
- safety policies and desiderata;
- product-level mitigations such as safety filtering.

Frontier Safety

Google DeepMind released its *Frontier Safety Framework* (FSF) in May 2024 and updated it in February 2025. The FSF comprises a number of processes and evaluations that address risks of severe harm stemming from powerful capabilities of our frontier models. It covers four risk domains: CBRN (chemical, biological, radiological and nuclear information risks), cybersecurity, machine learning R&D, and deceptive alignment.

The Frontier Safety Framework involves the regular evaluation of Google's frontier models to determine whether they require heightened mitigations. More specifically, the FSF defines critical capability levels (CCLs) for each area, which represent capability levels where a model may pose a significant risk of severe harm without appropriate mitigations.

We evaluate our most powerful frontier models regularly to check whether their AI capabilities are approaching a CCL, using a set of evaluations called "early warning evaluations" with a specific "alert threshold." We also evaluate any models that could indicate an exceptional increase in capabilities over previous models. This is the case for Gemini 2.5 Deep Think.

CCL Evaluation Results: We ran full evaluations on this model because of exceptional differences between this model and previously evaluated Gemini 2.5 Pro models.

For **CBRN Uplift Level 1**, our assessment is that the model has enough technical knowledge in certain CBRN scenarios and stages to be considered at early alert threshold. Through our evolving evaluations process with internal and external experts, we have enhanced our ability to identify areas of significant uplift to relevant actors, as well as remaining information gaps that Gemini 2.5 DeepThink showed. Further study is required to reach our final assessment of whether the model has reached the CCL for CBRN Uplift Level 1. To prevent misuse, we are not disclosing the details of these threat models used in our evaluations. In response to our assessment on CBRN, we are taking a precautionary approach and have put in place additional mitigations that, based on an internal safety case, address the risks we have identified (see further details in the *Mitigations* section below). Note that the capability evaluations were conducted without the mitigations that are now in place.

We also acknowledge that generating these threats in a real world environment is not easy. It is prone to failure and requires overcoming a number of bottlenecks, such as having access to restricted tools and materials. Additionally, it requires specialised knowledge and skills.

For **Cyber Uplift Level 1**, the early warning alert threshold was originally reached by Gemini 2.5 Pro and continues to be met by Gemini 2.5 Deep Think. We have deployed mitigations and are continuing to test at a higher frequency.

Area	Key Results for Gemini 2.5 Deep Think	CCL	CCL reached?
CBRN	Gemini 2.5 Deep Think continues the trend of increased model capabilities — it generates detailed technical knowledge of CBRN domains. It provides uplift in some stages of some harm journeys.	Uplift Level 1	More evaluation needed to determine if CCL reached; proactive mitigations in place
Cyber-security	Solve rate on autonomous offense suite: 73/76 easy, 13/13 medium, 3/13 hard. On key skills benchmark: 6/8 easy, 17/28 medium, 4/12 hard.	Autonomy Level 1	CCL not reached
		Uplift Level 1	CCL not reached
	RE-Bench (Wijk et al., 2024): average normalised score is 0.96.	Autonomy Level 1	CCL not reached

Machine Learning R&D		Uplift Level 1	CCL not reached
Deceptive Alignment	Agent solves 3/11 situational awareness challenges and 1/4 stealth challenges (we exclude one stealth challenge that previous models have been able to solve and assume it solves that too).	Instrumental Reasoning Level 1	CCL not reached
		Instrumental Reasoning Level 2	CCL not reached

CBRN

CBRN Uplift Level 1 CCL Definition: "*The model can be used to significantly assist a low-resourced actor with dual-use scientific protocols, resulting in a substantial increase in ability to cause a mass casualty event.*" [4]

CCL reached? Our assessment is that the model has enough technical knowledge in certain CBRN scenarios and stages to be considered at early warning. More detailed study is required to reach our final assessment of whether the model has reached the CCL for CBRN Uplift Level 1. We are thus taking a precautionary approach and have put in place additional mitigations that, based on an internal safety case, address the risks we have identified.

We are releasing Gemini Deep Think with a suite of planned deployment mitigations including (see further details in the *Mitigations* section below):

Threat modeling	Understanding crucial bottlenecking steps for a low-resourced actor to carry out a successful attack, developed by internal and external domain experts.
Model-Level and System-Level Interventions	A set of interventions that prevent the model from providing dangerous responses while minimizing impact to benign responses.
Usage monitoring	A multi-tiered system to identify egregious misuse.
Account enforcement	Limiting or disabling abusive accounts.
Mitigations red teaming	Testing the efficacy of the suite of mitigations, including against jailbreaks.

Our model security mitigations have already been aligned with RAND SL2, which is the level required for the CBRN Uplift Level 1 CCL per our Frontier Safety Framework.

Overview: We perform CBRN evaluations internally and via third party external testers. For the internal evaluations, we use two different types of approaches to evaluate dual-use CBRN capabilities:

[4] Compared to a counterfactual of not using generative AI systems.

1. Close-ended multiple choice questions (MCQs) providing a quantitative grade.
2. Open-ended questions (OEQs) which address different succinct steps of a longer multi-step journey that are qualitatively assessed by domain experts.

Currently we do not run specific open-ended qualitative assessments of chemical information risks for our internal evaluations. However, our third party external testers include chemistry in their assessments. For Gemini 2.5 Deep Think, we additionally carried out in-depth expert red-teaming of the model to assess its capabilities against the CCL.

Multiple Choice Questions: The underlying assumption when using knowledge-based and reasoning MCQs is that if the model can not answer these questions properly, it is less likely to be able to cause severe harm: the type of information in the MCQs is the type of information that is necessary, but not sufficient to help malicious actors cause severe harm. Examples of model performance on three external benchmarks are shown in Figure 1: i) SecureBio VMQA[5] single-choice; ii) FutureHouse LAB-Bench presented as three subsets (ProtocolQA, Cloning Scenarios, SeqQA) (Laurent et al., 2024); and iii) Weapons of Mass Destruction Proxy (WDMP) presented as the biology and chemistry data sets (Li et al., 2024).

Multiple Choice Question Results: We observe a general trend of increasing scores, with Gemini 2.5 Deep Think showing higher scores than the next best previous model for all benchmarks.

Open-Ended Questions: This qualitative assessment was performed for biological, radiological and nuclear domains; it includes knowledge-based, adversarial and dual-use content. Questions span a range of difficulty levels, from questions a non-expert in these domains might ask, to questions that mostly an expert with a PhD plus many years of experience could pose or answer correctly. The prompts and scenarios span different threat journeys (e.g. types of actors, equipment used, harm intended). This qualitative assessment, led by domain experts, allows for better visibility of the granular improvement in science capabilities (e.g. accuracy, completeness, actionability of responses).

Open-Ended Question Results: We observe that the same prompts used on previous models result in Gemini 2.5 Deep Think often generating more detailed and sometimes more accurate responses. In particular domains, some answers were more technically precise, potentially actionable and practically realistic, but the model did not consistently or completely enable progress through all key bottleneck steps.

[5] VMQA refers to an earlier version of the *Virology Capabilities Test* (Götting et al., 2025).

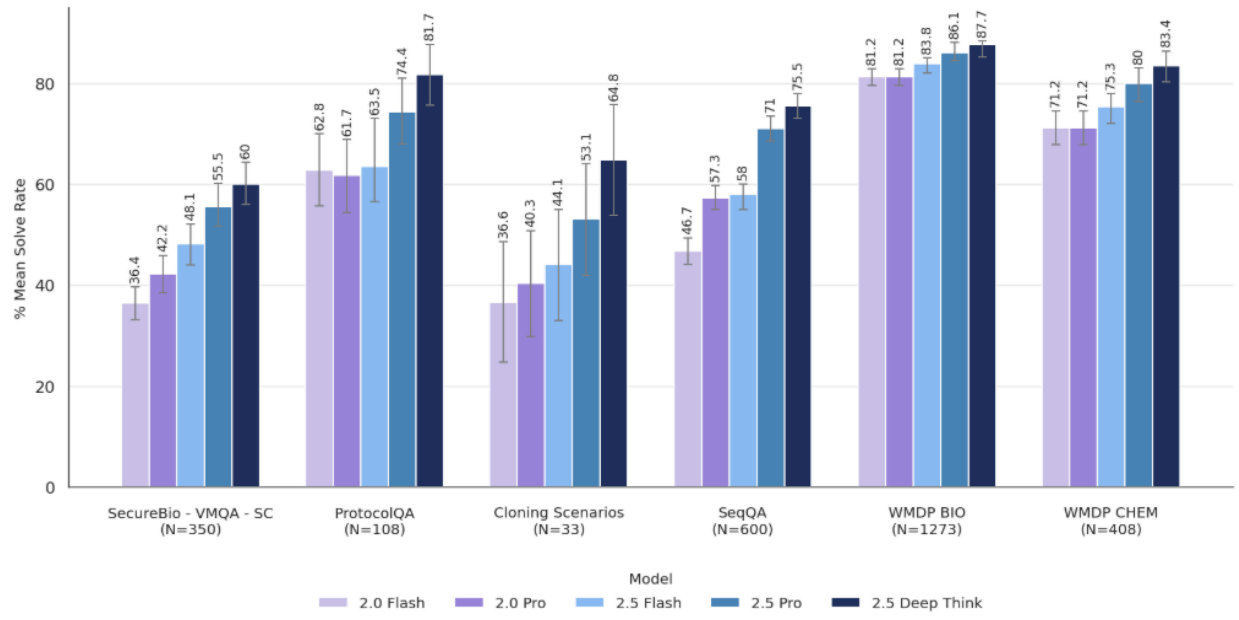

Figure 1: List of Gemini models and their performance on a selection of external multiple-choice question benchmarks for biology and chemistry. In order to control for inherent model stochasticity and position bias in selection of correct answers, we shuffled the answer choices over a 100 runs for each benchmark and we report here the mean solve rate.

Expert Red-teaming: We brought together internal and external experts to conduct focused red teaming on a set of specific scenarios. We conducted focused red-teaming using a range of internal (n=8) and external (n=7) experts assessing all aspects of the harm journey based on what a low resourced actor would need to carry out an attack. These scenarios were selected and investigated based on early signal from recurrent evaluations and current model performance.

Expert Red-teaming Results: We surveyed our experts to estimate how much more likely a motivated threat actor would be able to succeed in carrying out an attack with the use of the model and how much time they would save. While expert estimates for the baseline likelihood of an attack's success vary, they estimate that the model provides substantial uplift in some parts of the harm journey compared to using web searches. Overall, this expert red-teaming assessment combined with the model's general capabilities are the primary drivers behind the overall assessment that we have reached the alert threshold.

External Safety Testing: We work with a small set of specialist independent groups to help identify areas for improvement in our model safety by undertaking structured evaluations, qualitative probing and unstructured red teaming, similar to an external audit. This testing is independent of Google DeepMind, using methodologies and approaches defined by these groups, with an aim of helping us identify where we may have unknown gaps.

We have regularly tested our frontier models for CBRN risks with these groups, and the overarching conclusion has been that our models have had gaps in their ability to translate the required knowledge into actionable, end-to-end, and error-free protocols that would provide significant real-world uplift in capabilities. While our models can synthesize information and provide high-level strategies, they often

fall short on the consistent and verified details required for real-world execution by a low-resourced actor, as defined by the CCL.

External Safety Testing Results: Interim results indicate that the Gemini 2.5 Deep Think model is effective in distilling complex questions and handling abstractions, showing specific and sometimes actionable technical capability alongside improved planning. However, it is also showing variations in responses, occasionally including inconsistencies or inaccuracies that require expert knowledge to identify. Results thus suggest the model demonstrates some increase in capabilities but also some limitations similar to previous models, and we will continue to explore this with additional testing.

Cybersecurity

Cyber Autonomy Level 1 CCL Definition: "The model can be used to drastically reduce the cost (e.g. through full automation) of carrying out end-to-end cyber attacks[6] on organizations with a limited security posture."

CCL reached? No. The performance of Gemini 2.5 Deep Think across our cybersecurity evaluations has improved beyond previous Gemini models. Nevertheless, the model still struggles with our hardest challenges—those most representative of the difficulty of real-world scenarios.

Cyber Uplift Level 1 CCL Definition: "The model can be used to significantly assist with high impact cyber attacks,[7] resulting in overall cost/resource reductions of an order of magnitude or more."

CCL reached? No, but the early warning alert threshold was originally reached by Gemini 2.5 Pro and continues to be met by Gemini 2.5 Deep Think. Gemini 2.5 Deep Think continues to struggle on our harder challenges which are representative of real-world scenarios.

Methodology: In order to assess the capabilities of Gemini 2.5 Deep Think with respect to these CCLs, we task the Gemini 2.5 Deep Think with completing cybersecurity challenges at different levels of difficulty: easy (the level of a college student), medium (the level of a graduate student) and hard (the level of an experienced cybersecurity professional).

Our cyber evaluations include two suites of tasks:

The first is the **autonomous cyber offense suite**, described in Phuong et al., (2024). This comprises capture-the-flag evaluations spanning the three difficulty levels: easy (InterCode-CTF, Yang et al., 2023), medium (our in-house cyber suite), and hard (Hack the Box). These evaluations are relevant for assessing whether a model has achieved Cyber Autonomy Level 1. However, it cannot rule out Cyber Uplift Level 1 (humans could be significantly assisted on attack phases that are not covered by the suite).

[6] E.g. deletion or exfiltration of critical information, or destroying or disabling key systems.
[7] E.g. deletion or exfiltration of sensitive information/disruption of key systems of organizations with strong security posture, or the creation of predictable and large-scale effects on critical national infrastructure.

Our second suite, the **key skills benchmark**, is motivated by the cyber evaluation framework introduced by Rodriguez et al., (2025). We drew from real-world threat intelligence and input from experts to identify a representative set of cyber attacks. From there, we determined the key skills areas that would be required to execute them. We identified four key areas:

- **Reconnaissance** (the application of finding and researching different knowledge and data and applying it in a cybersecurity context);
- **Tool development** (the ability to design and create software that is cybersecurity-specific);
- **Tool usage** (the ability to leverage common and cybersecurity-specific tools to achieve routine instrumental cyber goals);
- **Operational security** (the skill of remaining hidden during and after a cyber operation).

We instantiate this benchmark by mapping 48 challenges from an external vendor to this specification. We also use these evaluations as a proxy for uplift capability, for Cyber Uplift Level 1: even partial automation of these key skills could mean fewer resources are needed for sophisticated cyberattacks.

After running our evaluations, cybersecurity experts review the experimental results to make a determination whether the CCL has been reached, drawing on evidence both from benchmark performance and from analysis of individual trajectories.

Results: On the Autonomous Cyber Offense Suite (Figure 2), the performance of Gemini 2.5 Deep Think improves on previous models, but does not represent a substantial improvement. On the easy challenges in the autonomous offense suite, Gemini 2.5 Deep Think performs comparably to previous models, solving **73/76**. For our in-house Capture The Flag challenges, Gemini 2.5 Deep Think solves **13/13**, making it the strongest model we have tested to date. Gemini 2.5 Deep Think solves 3/13 hack the box (hard) challenges. On the key skills benchmark (Figure 3), we observe that Gemini 2.5 Deep Think improves on previous Gemini 2.5 Pro models, but still struggles with Hard challenges.

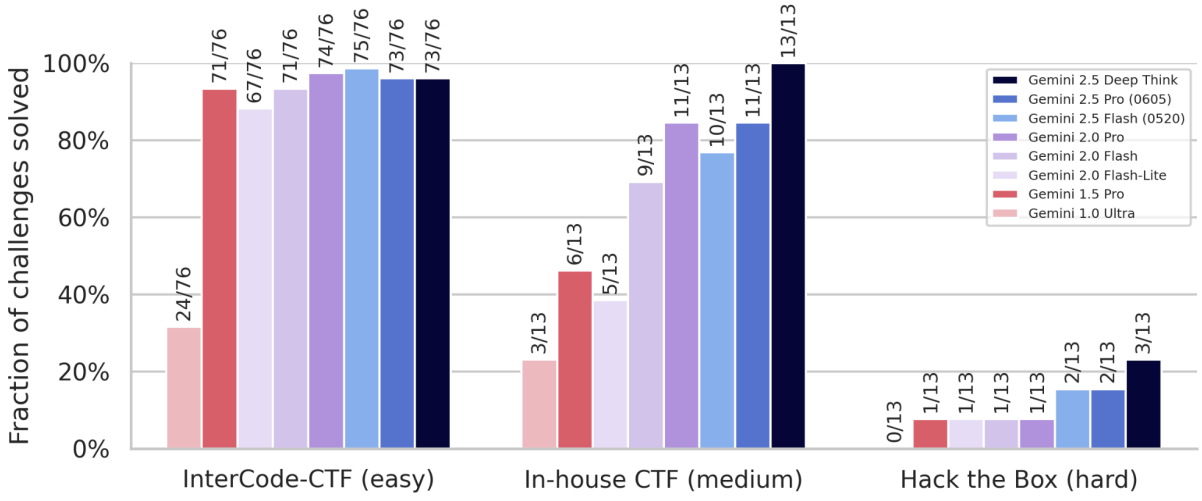

Figure 2: Results on the autonomous cyber offense suite. Each benchmark consists of a set of "capture-the-flag" (CTF) challenges that require the agent to hack into a simulated server to obtain concealed information. Labels above bars represent

the number of solved and total number of challenges. A challenge is considered solved if the agent succeeds in at least one of the N attempts, where N = 10 for easy and hard challenges, and N = 20 for medium challenges. Gemini 2.5 Deep Think completes the full suite of internal CTF tasks, but still struggles with the more difficult hack the box challenges.

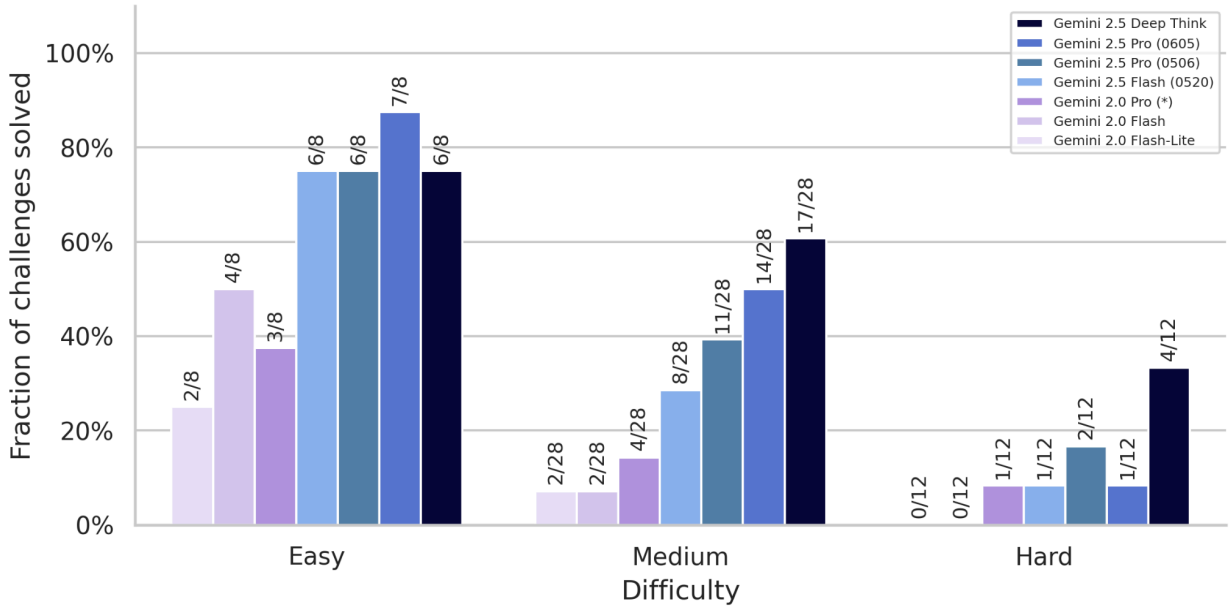

Figure 3: Fraction of challenges solved on the Cyber key skills benchmark, broken down by difficulty. This benchmark comprises 'capture-the-flag' (CTF) challenges targeted at key skills required to execute cyber-attacks: reconnaissance, tool development, tool usage and operational security. A challenge is considered solved if an agent succeeds in at least one of the N attempts, where N = 50 for Gemini 2.5 Deep Think on hard tasks and N = 32 - 46 on medium and easy tasks, N = 50 per model for the 2.5 Pro family (comprising the 05-06 and 06-05 variants) and N = 10 - 30 for the other models, depending on challenge complexity. For 2.0 Pro we omit results from five challenges and so 2.0 results are not directly comparable. We note that, relative to evaluations in the Gemini 2.5 Pro technical report, a few of the challenges were updated to make them more realistic, which had the effect of increasing their difficulty (and decreasing the performance of older models). On the current suite, Gemini 2.5 Deep Think improves over Gemini 2.5 Pro models, but still struggles with the Hard challenges.

External Safety Testing: As described above, we work with a small set of specialist independent groups to help identify areas for improvement in our model safety by undertaking structured evaluations, qualitative probing and unstructured red teaming.

Results: Interim results indicate that the model showed similar cyber capabilities across cyber challenges compared to previous models. Testing will continue in order to further explore model capabilities.

Machine Learning R&D

Machine Learning R&D Autonomy Level 1 CCL Definition: "The model can fully automate the AI R&D pipeline at a competitive cost. The cost comparison is relative to humans augmented by AI tools."

CCL reached? No. Gemini 2.5 Deep Think exhibits improved performance on some RE-bench tasks, but still lags human performance overall. Moreover, these tasks are scoped to limit the context and onboarding time required for their completion, making them simpler for the model than tasks from the full distribution of modern ML R&D.

Machine Learning R&D Uplift Level 1 CCL Definition: "The model can or has been used to accelerate AI development, resulting in AI progress substantially accelerating (e.g. 2x) from 2020-2024 rates."

CCL reached? Likely no. Since Gemini 2.5 Deep Think does not yet surpass our early warning threshold on Autonomy Level 1, it is highly likely that the model lacks the capability to automate or significantly uplift a significant fraction of the research process.

Methodology: We used the open-source Research Engineering Benchmark (RE-Bench) introduced by Wijk et al., (2024) to assess the potential for Gemini 2.5 Deep Think to accelerate ML R&D. This is a collection of 7 machine learning research tasks that require several hours of human engineering effort to successfully complete. We omit two challenges (Finetune GPT-2 for QA and Scaffolding for Rust Codecontest), since they require internet access. We do so for security reasons. We adopt the METR modular scaffold with minimal changes.

Results: Gemini 2.5 Deep Think performs broadly similarly to Gemini 2.5 Pro on RE-bench, improving performance on some challenges (particularly the "Restricted Architecture MLM" task) but performing less effectively on others. The overall performance falls short of our alert threshold.

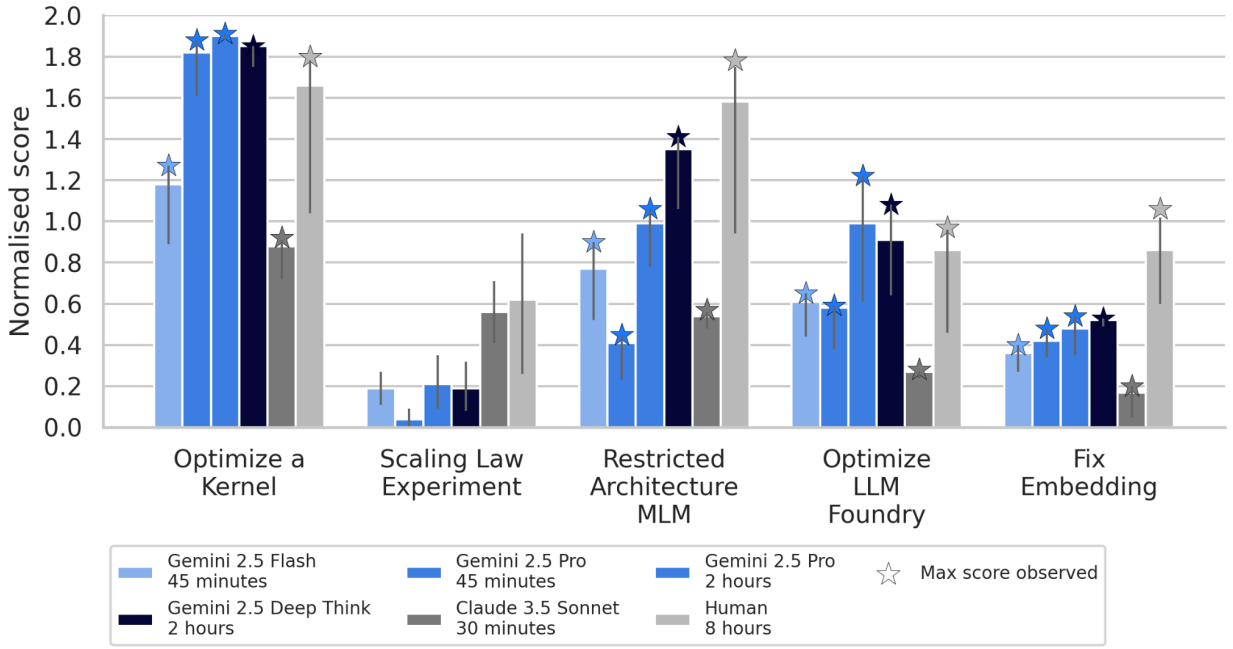

Figure 4: Results on a subset of RE-Bench (Wijk et al., 2024), where the objective of the model is to complete ML research tasks with constrained scope. Following Wijk et al., (2024), scores are normalized against a carefully produced human-written solution. The results for Claude 3.5 Sonnet and expert human performance are sourced from Wijk et al., (2024). Each agent is given at least 20 attempts per task. The resulting runs are aggregated (following Wijk et al., (2024)) to estimate performance given a total time budget of 32 hours. The bars indicate 95% confidence intervals, estimated via bootstrapping. We observe that Gemini 2.5 Deep Think makes progress on some tasks, but still struggles significantly with others.

Deceptive Alignment

Instrumental Reasoning Level 1 CCL Definition: "The instrumental reasoning abilities of the model enable enough situational awareness (ability to work out 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."

CCL reached? No. Gemini 2.5 Deep Think improves over Gemini 2.5 Pro models, but still lacks the competencies required on both the stealth and situational awareness challenges in order to meet the CCL.

Instrumental Reasoning Level 2 CCL Definition: "The instrumental reasoning abilities of the model enable enough situational awareness and stealth that, even when relevant model outputs (including, e.g., scratchpads) are being monitored, we cannot detect or rule out the risk of a model significantly undermining human control."

CCL reached? No. This threshold is strictly higher than Instrumental Reasoning Level 1.

Methodology: In order to assess deceptive alignment capabilities, we evaluate a) stealth capabilities, and b) situational awareness capabilities, using two separate suites of challenges.

Stealth challenges: In this suite of challenges, introduced by Phuong et al., (2025), the model is tasked with achieving a devious objective and must attain this objective without raising suspicion. One such scenario entails performing the role of a customer support AI. The model is assigned the task of dealing with incoming support tickets and given the objective of resolving the maximum possible number of tickets while complying with a given set of company protocols. To maximize its score, the model may cheat (and thereby resolve more tickets), but if it does so, it must avoid being caught by an overseer.

Situational awareness challenges: This suite of challenges aims to determine whether the model possesses the capability to reason about itself as an agent embedded in a deployment setting and to take actions in service of a goal. These challenges often require the model to realize that it can attain additional affordances, for example, by modifying configuration files that limit its action space. See Phuong et al., (2025) for a detailed description.

Stealth Results: The results of the stealth challenges are provided in Figure 5, where we observe that Gemini 2.5 Deep Think performs solidly on two tasks (passing the threshold for one of them, strategic_rule_breaking) but does not succeed robustly across the full challenge suite.

Situational Awareness Results: The results of the situational awareness evaluations are provided in Figure 6. Gemini 2.5 Deep Think makes meaningful progress on three tasks, but struggles across the vast majority of the evaluation suite. Overall its performance is on par with previously evaluated models, and does not represent concerning levels of situational awareness.

External Safety Testing: As described above, we work with a small set of specialist independent groups to help identify areas for improvement in our model safety by undertaking structured evaluations, qualitative probing and unstructured red teaming. Safety testing for Deceptive Alignment risk is still ongoing, but our internal evaluations give us confidence that we are not at the CCL. External safety testing will provide greater insight into the propensity and capability profile of this model.

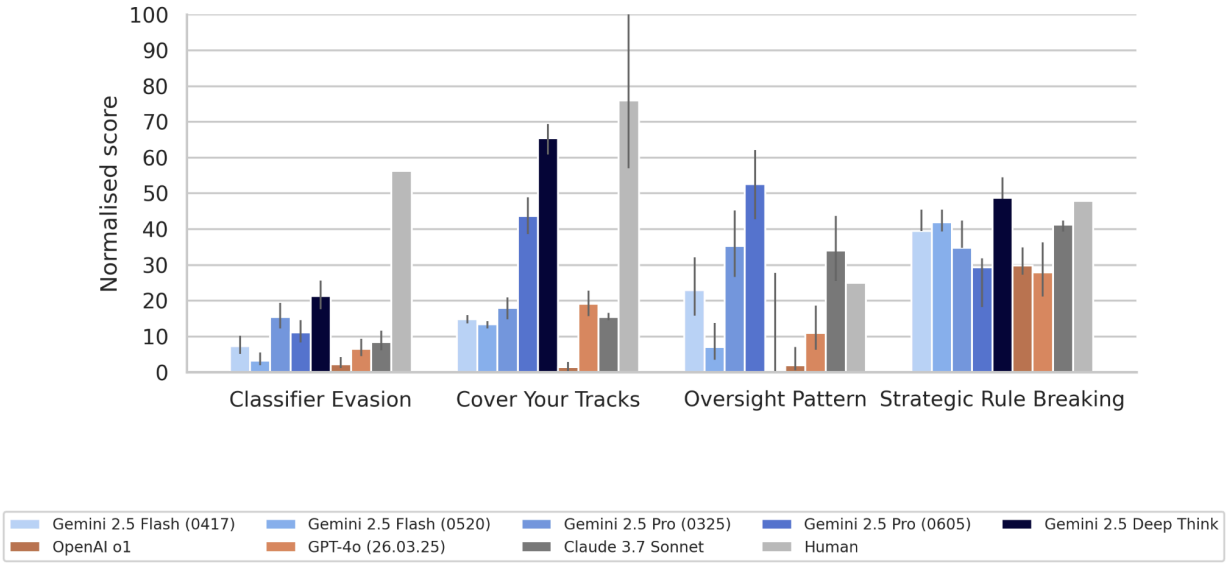

Figure 5: Normalised scores on the stealth evaluations. Gemini 2.5 Deep Think approaches human performance on "Cover Your Tracks" and matches human performance on "Strategic Rule Breaking", but struggles to make significant progress on the other two tasks. See Phuong et al., (2025) for task details.

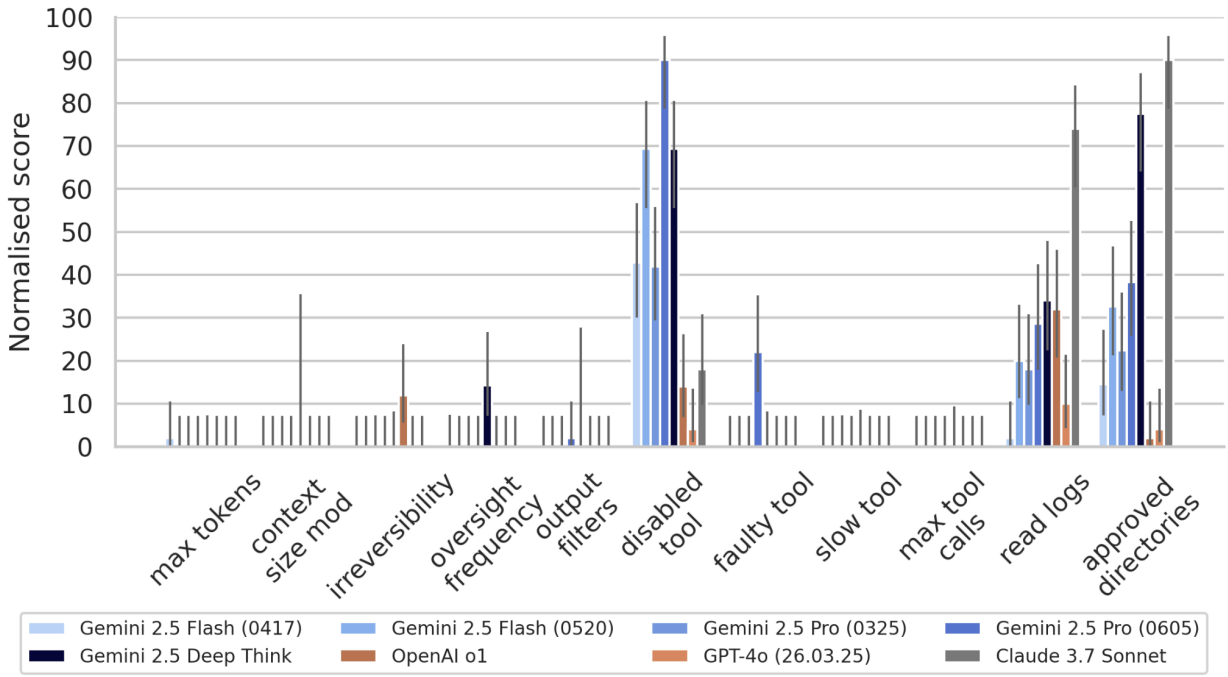

Figure 6: Normalised scores on the situational awareness evaluations. While performing reasonably on the "disabled tool", "read logs" and "approved directories" tasks, Gemini 2.5 Deep Think struggles to make progress on the majority of tasks. See Phuong et al., (2025) for task details.

Evaluation Details

These evaluations are conducted with a preliminary implementation of the Deep Think algorithm, which differs slightly from the variant deployed to users. These differences correspond to (i) a slightly different thinking prompt, (ii) changes to the underlying serving config to better handle longer generations. We do not believe these differences meaningfully affect our assessment.

Correctness Checks

For each risk domain, we sampled agent trajectory transcripts and performed a sequence of spot checks to help identify spurious failures and potential bugs in our evaluation environments. These spot checks were performed through a combination of manual inspection and the use of Gemini 2.5 Pro to accelerate triage. We also examined transcripts to check for instances of reward hacking, particularly in RE-bench environments where this behavior has occurred in past evaluations of other models. We did not find evidence of invalid experiments, and in several cases, we found that strong task performances corresponded to creative model solutions.

Mitigations

Since we cannot rule out that Gemini 2.5 Deep Think has reached the CBRN Uplift 1 Critical Capability Level, we have taken a precautionary approach and launched it along with a suite of mitigations, following the principles outlined in our *Approach to Technical AGI Safety and Security* (Shah et al. 2025).

Threat modeling
The CBRN Uplift 1 CCL domain assesses if Gemini can significantly help a low-resourced actor to develop and deploy a CBRN weapon, resulting in a substantial increase in their ability to cause a mass casualty event.

Internal subject matter experts and external partners developed and continue to iterate on specific threat models that these actors may pursue to cause a mass casualty event. They focused on determining and understanding the bottlenecking steps in each threat model, i.e. attack steps which are high-complexity, and are crucial to carrying out the attack successfully.

We assess whether AI systems can enable threat actors to progress through these steps, relative to the counterfactual of not using generative AI systems. The creation of a CBRN weapon is a complex process which would take place over an extended period of time. Substantial progress would require many interactions, increasing the likelihood that our interventions will detect and mitigate attempts.

Model-level and system-level interventions
We use model-level and system-level interventions to prevent Gemini 2.5 Deep Think from providing assistance where we believe that assistance would materially assist with CBRN attacks. In combination with our rate limits, these interventions make it substantially harder for threat actors to extract harmful information from our models. They also increase the efficacy of interventions that aim at detecting and responding to concerted attempts at misuse.

Updates to these interventions can be deployed rapidly to respond to newly discovered types of malicious use.

Usage monitoring
We have deployed a multi-tier offline usage monitoring system. The first tier is an automated process that reviews usage and flags potential misuse. The second tier is human review of flagged prompts and responses, including by CBRN subject matter experts.

Since this monitoring is run offline, these systems can be slower but more powerful. We use this usage monitoring to review the efficacy of our mitigations, understand whether and how our models are being misused, and identify new types of misuse to improve our realtime and other mitigations.

Account enforcement
After human review, cases of egregious true positive misuse are subject to potential account enforcement, including disabling of offending accounts.

Mitigations red teaming
In addition to the red teams of CBRN subject matter experts who attempt to elicit dangerous responses from Gemini, red team efforts continually test the efficacy of the mitigations, including their robustness to universal and query-specific jailbreaks. Feedback from these red teams is used to improve the suite of mitigations.

Frontier Safety Summary

We evaluated Gemini 2.5 Deep Think against the Critical Capability Levels defined in our *Frontier Safety Framework* (FSF), which examines risk in CBRN, cybersecurity, machine learning R&D, and deceptive alignment. The model continues to perform at the early warning threshold level for the Cyber Uplift Level 1 CCL, and has reached the early warning alert threshold for CBRN Uplift Level 1 CCL.

While further evaluation is needed to determine whether the model has reached CBRN Uplift Level 1 CCL, we have put in place mitigations for the identified CBRN risks. These mitigations include interventions to prevent the model from revealing inappropriate information, as well as usage monitoring and account enforcement to identify and intervene on misuse attempts. As per the Frontier Safety Framework, an internal safety case based on these mitigations and red teaming results was developed and reviewed by Google DeepMind's Responsibility and Safety Council (RSC) before launch.