

Mistral previews Large 4, a 1-trillion-parameter model with weights due by end of October

Source copy kept by wheresthe.ai for the story "Mistral previews Large 4, a 1-trillion-parameter model with weights due by end of October" (<https://wheresthe.ai/d/fr-mistral-large-4-preview-1-trillion-open-weights-2026-10/>).

Source: <https://mistral.ai/news/mistral-large-4/>

Text saved 19 Oct 2026 (the page as published).

SHA-256 of the saved text: 6e274729a22138a20fcc2b8ffc0b94a516b9658ed1dacbdf2457358e4f5aec6a

Le chonk Introducing Mistral Large 4

[Back to Blog](#)

14 min read

October 6, 2026 By Mistral

[Share this post](#)

[Copy url to clipboard](#) Copied

Le Chonk

Today, we're launching a public preview of Mistral Large 4. Unofficially ML4, very officially: le Chonk . ML4 pushes the frontier of open-weight performance. You can try the preview API today on Mistral Studio . Weights drop end of this month.

Frontier performance ML4 is a 1 trillion-parameter natively multimodal model with 52 billion active parameters. It is our largest and most capable model to date, and it continues to improve rapidly as we refine it.

The model demonstrates exceptional performance across coding, agentic workflows, and multimodal understanding. It already achieves performance competitive with the strongest open-source models globally, while significantly outperforming any open-weight model developed in the US or Europe. On critical enterprise workloads, including cybersecurity, finance and law, we find it to be state-of-the-art among open models. In some domains such as visual grounding, it goes further still, surpassing even frontier closed models.

We will release the weights by the end of the month. Until then, we are red-teaming the model in real-world settings with cybersecurity leaders, vetted partners, and state authorities, who will access the same model with reduced moderation and expanded cyber capabilities.

[Coding - DeepSWE](#)

[Coding - Terminal Bench 4.0](#)

[Cyber](#)

[Agentic behaviour](#)

[Finance Agent](#)

[Harvey's Legal Agent](#)

[Grounding](#)

Forged in Europe. Built for AI sovereignty. ML4 was trained from scratch on 3,800 NVIDIA Grace Blackwell GPUs in Mistral's own datacenters in Europe. The public preview is served on that same infrastructure. It is a significant milestone in our long-term investment across infrastructure, research, and product development: state-of-the-art performance in critical verticals, delivered through open weights, designed to give customers control over their AI.

This is particularly important in cybersecurity, where provider-level refusals can block legitimate vulnerability research and incident response, and where losing access to a capability mid-incident can itself become a critical security risk. ML4 pairs top-tier cyber performance with open weights and self-deployment, giving organizations both the capability and the autonomy to run advanced security work under their own policies.

The model will be available across multiple regions worldwide, including a European deployment that Mistral operates end-to-end, independently of other digital service providers and under European law. Fun fact: a significant share of ML4's training data was multilingual, spanning more than 160 languages, including every official language of the European Union.

We've been working closely with leading enterprises across the world in finance, engineering, manufacturing, logistics, pharmaceuticals, science, shipping, public sector, and other mission-critical industries to train ML4. In fact, the model uses the same training, customization, and RL environment we offer our customers through Mistral Forge.

Try it today There is still more to come. As we work toward releasing the weights, we will share further details on the model architecture, additional benchmarks, and our post-training methodology. This model will also serve as the foundation for a new generation of specialized and optimized Mistral models. In the meantime, we invite you to try the preview API and share your feedback with us on social media.

Capabilities deep-dive Cybersecurity ML4 is one of the world's strongest AI models for cybersecurity. On the Artificial Analysis Cyber Index, an independent evaluation of how well AI models find and fix security flaws in real software, it ranks among the top five models globally and leads open-weight models developed outside China by a wide margin. On one of the index's tests, which asks a model to reproduce a real vulnerability in open-source software and then patch it, ML4 scores 82%, the highest of any model. It also solves 93% of the challenges in Cybench, a set of 40 exercises drawn from security competitions, one of the highest scores reported for an open-weight model. That top score reflects a practical advantage. Several leading closed models, including Claude Opus 5.5 and GPT-6 Astra, score near zero on the same test because they refuse to perform the task. Yet defending software often starts with proving that a flaw is real, exactly the kind of work safety filters in closed models can block. This matters even more as threat actors increasingly jailbreak those same models to support offensive cyber activity: defenders need systems that can match those capabilities without being constrained by the same refusals. ML4 can do that work, and its capabilities extend beyond what it was explicitly trained for: in internal testing, it proved useful for analysing malware, prioritising vulnerabilities, and writing detection rules. For organisations that need sovereign, auditable AI for security operations, it will be able to run on private cloud or on-premise.

ML4 against the field : efficiently reasoning over diverse complex challenges

Malware reverse-engineering: solving an out-of-distribution investigation task

AA Cyber Index

CyberGym-E2E

Cybench

Agentic coding ML4 excels across software engineering, repository understanding, and complex terminal workflows, scoring 61.7% on DeepSWE v1.1, 59.4% on SWE-Atlas-QnA, and 28.3% on Terminal-Bench 4. Its combined Coding Agent Index score of 49.8% places it ahead of DeepSeek V4 Pro 0813 and Qwen3.8 Max.

DeepSWE 1.1

Terminal Bench 4.0

SWE Atlas QnA

We also ran a blind human evaluation with Surge AI on coding quality: professional annotators rated model outputs on a 1–5 scale, with model identities hidden. ML4 Preview ranked second of five models (3.74), ahead of Kimi K3 (3.59), GLM-5.3 (3.60) and GLM-5.2 (3.40), and behind only Claude Opus 5 (4.22). **Agentic Workflows** ML4 runs general-purpose agents that gather information, use tools, and produce finished deliverables across complex workflows. On AutomationBench — 657 business workflows across apps like Gmail, Google Sheets, Slack, and Salesforce — it scores 59.9%, ahead of Kimi K3, MiMo-V2.6-Pro, and DeepSeek V4 Pro. It's just as strong on the professional deliverables that knowledge work actually produces: spreadsheets, slides, and PDFs. On AA-Briefcase, which evaluates long-horizon knowledge work, it reaches 1,393 Elo, ahead of DeepSeek V4 Pro. AA - AutomationBench

AA-Briefcase

Multimodal ML4 is a step change in the ability of our models to understand images. It reasons powerfully across complex documents, charts, and natural images, and brings vision to the industries where perception is critical such as engineering, manufacturing, and earth observation.

The model can further combine visual grounding with agentic capabilities: from inspecting gigapixel satellite imagery — helping disaster-response teams act when time counts — to analyzing engineering-drawings — zooming in, inspecting, and verifying until the answer is exact. In our demos above, ML4 grounds dense natural scenes, verifies mechanical parts in technical drawings, retrieves evidence from PDFs, and scans massive geospatial images for the hardest-to-find objects.

On visual grounding particularly, we find ML4 to be one of the most capable models we tested, for instance surpassing GPT-6-Astra on Dense 200 (42% vs 41%). Dense 200

ChartQA Pro

GDP.pdf

Science and Math ML4 brings strong scientific capabilities, built by combining AI-driven methods with our researchers' expertise in mathematics, physics, and chemistry.

It's highly proficient at agentic coding for scientific tasks such as data analysis, modeling, and simulating physical reality, which lets researchers focus on the questions rather than the plumbing. In benchmarks, ML4 is state of the art on SciCode-Verified among open-weight models. In practice, it can generate a full Hartree–Fock simulation in one shot — a complex, multi-step chemistry task built from a series of advanced routines.

ML4's math is stronger too, in both formal reasoning and applied mathematics. In our human evaluations it reasons more precisely and with more structure than GLM-5.3, and it can sustain long, domain-specific applied-mathematics tasks, including work relevant to frontier theoretical physics. Together, these capabilities make ML4 a strong research assistant across the full technical workflow — from the first question to the final result. SciCode-Verified tests the capabilities of models to implement complex scientific workflows in code for domains such as physics, mathematics, material science and biology. Internal eval on STEM tasks (math and physics) of ML4 against GLM5.3

Knowledge Work ML4 is our most capable model for the real-world tasks which professionals handle every day. It can create, edit and fix complex spreadsheets and documents, showing exemplary performance on both legal and financial benchmarks. Notably, we evaluated ML4 through third party evaluators (vals.ai) on representative tasks for both legal and financial tasks, finding the model exceeds GPT-6-Astra in both cases. On HarveyAI's Legal Agent benchmark, ML4 outperforms all open-source models. Finance Agent v2

Finch (FinWorkBench)

Harvey's Legal Agent Benchmark

FinWorkBench tests model capabilities at creating/editing spreadsheets on real life Finance and Accounting use cases.

Financial analysis demands precision and the ability to synthesize information from multiple sources, a process that remains time-consuming at many financial institutions today. In this demo, ML4 compared to other top OSS models take on the same multistep corporate finance challenge, searching through public company filings and financial reports, such as those available via EDGAR and equivalent European databases. An animated semantic map traces each model's journey toward a solution, highlighting every document retrieved along the way. Each track's position reflects the evidence gathered, the results of calculations, and the questions that remain unresolved. Viewers can follow how the investigations unfold and compare the distinct paths each model takes before arriving at its final answer.

Model Safety B3 Attack Resistance

Cyber Refusal Rate

KORABench

ML4 has saturated our benchmarks on robustness to indirect prompt injections, putting it at the frontier of OSS models (compared to GLM-5.2, GLM-5.3, Kimi-K2.6, Kimi-K3, DS-V4-Pro-0813). On Lakera's public B3 AI Security Benchmark, ML4 resists 93.3% of attacks – we see no higher scores among competitors.

ML4 also engages more responsibly with users than any of our previous models. We highlight our results on the KORA Benchmark , where ML4 again sits at our highest measured score among OSS models (1.691, with 2 being the maximum denoted as “ Exemplary ”).

Of particular relevance is the model’s propensity to refuse malicious requests regarding cybersecurity. Despite strong performance on Cyber benchmarks, the average refusal rate of the model on cyber prompts from JailbreakBench , StrongREJECT , and AgentHarm is higher than all OSS models.

Human Evaluation We ran an internal evaluation in which expert annotators across coding, computer-aided design (CAD), finance, mathematics and physics compared Mistral Large 4 with GLM-5.3. ML4 was preferred in CAD and STEM, while performing on par or close to GLM-5.3 in finance and coding. Reinforcement learning at scale Base models are improving fast, and our post-training has to keep pace. A recipe tuned for yesterday’s model leaves capability on the table with today’s frontier, because ground truth samples that once pushed a model to its limits won’t anymore. We use Reinforcement Learning (RL) because it adapts as the model does: we train on the outcomes of the model’s own attempts, and we can raise the difficulty and the breadth of the tasks as it gets stronger. Our RL library was designed to make new environments easy to add and train at scale. A shared, composable interface allows a single training run to combine tasks ranging from single-turn chat and complex scientific problem solving to safety alignment, factuality, and long-horizon tool use. These environments share scaffolds and resources such as code sandboxes, web search, and external APIs. The same composability extends to verification, with reward models, unit tests, LLM judges, and static checks combined as needed for each task. At runtime, an autoscaling fleet of actors generates tens of thousands of rollouts in parallel while model training proceeds asynchronously. The generation and training pipeline is optimized for long trajectories, supporting rollout budgets of millions of tokens across multiple compactions while keeping staleness low. Novel methods and optimizations across both stages minimize off-policy drift and enable stable RL over long horizons. At our current scale (3k GPUs), a single training run produces roughly 33 billion tokens per day , of which around 16 billion are trainable completion tokens after filtering and masking. We can see the run progress directly in the training rollouts: training rewards rise across several representative environments as the policy learns to solve increasingly complex tasks. Below are a few examples.

The improvements are not specific to the environments we train on; they transfer to downstream evals, and the final model owes them to both post-training stages (supervised fine-tuning and RL), as shown in the charts.

What comes next This is only the beginning. ML4 is the first milestone on the roadmap funded by our €3 billion Series D — the largest equity round ever raised by a European technology company. That capital is already being put to work: we are significantly scaling up our compute capacity in our own European datacenters, and much more is coming online in the months ahead. More compute means more training. The reinforcement learning run behind this preview is still in flight, and the model is showing no signs of saturation — there is substantial headroom ahead. As we scale up training on our expanded infrastructure, we expect large and rapid improvements in the weeks and months to come. We will release the weights by the end of the month, along with more details on the architecture, additional benchmarks, and our post-training methodology. And ML4 is only the foundation: it will serve as the base for a new generation of specialized and optimized Mistral models, built for the industries and workloads our customers care about most. The pace of progress from here will be fast. Stay tuned.

Mistral Large 4 New

Open-weight hybrid instruct-and-reasoning MoE with multimodal input; unifies instruction, reasoning, and agentic capabilities in a single model, state-of-the-art among open weights on cybersecurity, finance, and manufacturing, natively fluent in 160+ languages. Multimodal Reasoning Coding Agentic Cybersecurity

Input (/M tokens)

\$1.36

Output (/M tokens)

\$4.18

[Read more](#)

0%