


LLM Model Releases

Cost has fallen and capability has increased. Model releases are increasing in frequency as more labs release new models and existing labs use existing AI to train new AI. Datasets being used to train AI are no longer rubbish from the internet. Labs are paying significant sums of money for data.



artificial analysis

ident analysis of AI

Artificialanalysis

AI Model & API Providers Analysis | Artific...
Comparison and analysis of AI models and API hosting providers. Independent benchmarks across key performance metrics including quality, price, ou...



OCT - SKY-001 is Skyfall AI’s “Engineering World Model” for CAD and BIM.

It was unveiled on October 19, 2026 in a research preview; commercial availability is reportedly targeted for early 2027. Interesting as example of industry specific model. We will see more of these this year.

OCT - MUSE Spark Open - Meta

Muse Spark Open is a natively multimodal model with 816B total parameters (42B active) and a 1M-token context window. Built for coding and agentic work, it excels at long-horizon tasks, complex software engineering, and multi-step reasoning.

OCT 5 - Beam from Reflection AI (Reflection focused on Open Sovereign AI)

Beam takes text as input and generates text as output. This release includes the instruction-tuned model in BF16, FP8, and NVFP4 numerics. Beam uses a sparse mixture-of-experts (MoE) architecture with 501B total parameters and 23B active parameters. It supports a context window of 1M tokens and was developed predominantly for English.

OCT 19 - MINIMAX M3.1-FLASH

As a next-generation Flash model, M3.1 Flash delivers a generational leap in agentic tasks, coding, Computer Use and cybersecurity, while making online inference nearly three times faster and reducing usage costs by 70%. It is aimed at long-context coding, multi-step agent workflows, computer use, cybersecurity tasks, and visual/frontend work. The 1M-token context makes it particularly relevant for working across large repositories, long transcripts, or extensive documentation.

OCT 21 - FLINT

Flint is Recursive’s open model for local coding agents. Flint scored **54.9% on Terminal-Bench 2.1**, **89.7% on SWE-bench Pro v2**, and **74.2% on SWE-bench Verified**. Flint generates more than 125 tokens per second on an M4 Pro MacBook Pro and more than 150 tokens per second on a 64GB NVIDIA RTX Spark at short contexts—up to twice as fast as comparable models.

Decision Models



Sep 15 — Jev

Established the modern System One category, with probabilistic decisions as a native model output rather than generated JSON.

Sep 29 — Liquid AI d1

Established a significant commercial alternative, later incorporating vision and open-weight variants.

Oct 1 — Cloudflare Clef

Made decision models available within Cloudflare’s edge-inference ecosystem, creating an attractive option for production workflows.

Oct 9 — Microsoft-Decision-1


Brought the category into Microsoft Foundry, positioning decision models for enterprise automation and agent orchestration.

Oct 15 — Liquid AI d1-ultra

Marked the shift toward competing not just on speed and cost, but also on decision quality and multimodal capabilities.

Three important developments have occurred in just over a month:

- 1. The category has **rapidly expanded**. What started with Jev now includes offerings from Microsoft, OpenAI, Cloudflare, Liquid AI and specialist laboratories.
- 2. Decision models are **becoming multimodal**. Text-only decision engines are expanding into image, audio and video understanding.
- 3. They are becoming a separate component of agent architectures. General LLMs can handle reasoning and planning, while fast **decision models handle routing**, validation, scoring and execution control.



Personal AI Agents

Over the past six months, personal agents have shifted from chat assistants into **persistent, proactive task-executors** that can use browsers, apps, devices, and memory—while increasingly asking for approval before consequential actions. The biggest change is that agents now continue working after you close the app and can carry a task across email, web, files, and third-party services.

November 2025 - April 2026



Open Claw



Claude CoWork



Hermes



Perplexity Computer



APR 16 - Perplexity Personal Computer- Perplexity
Extends the agent experience to local computer applications, files and browser tasks.



MAY 19 - Gemini Spark - Google
A cloud-based, proactive personal agent designed to continue working independently across Google services.



JUN 2 - Microsoft Scout - Microsoft
Always-on autonomous agent for workplace tasks; subsequently brought under Autopilot branding.



JUL 9 - ChatGPT Work - OpenAI
Multi-step execution across files, apps and websites, creating complete work deliverables.



AUG 11 - Grok Bot - xAI
Persistent AI teammates with their own cloud computers and access to connected applications.



AUG 29 - Instinct Personal Agent
Instinct is an invite-only personal AI assistant built by Spear Street Technology that handles life admin through conversational channels like iMessage and WhatsApp. Instead of a standard app UI, it controls a dedicated cloud computer to execute tasks like booking travel, ordering groceries, and managing emails



SEP 8 - Muse - Meta
Consumer personal agent, available through an app and messaging, with cloud-based task execution and personal context.



SEP 29 - Dots - OpenAI
Always-on personal agents with persistent context, cloud computers, connected apps and ongoing goals.



OCT 21 - Claude Duet
Anthropics personal agent product.

Agent frameworks & harnesses

The major evolution during this period has been from agent orchestration frameworks toward complete, persistent, sandboxed and interchangeable agent execution environments. Developers increasingly gain the ability to integrate multiple agent harnesses through standardized protocols and shared interfaces.

- LLMs, call tools, MCP connections.
- long, multi-step tasks, planning and todo tracking, context management, memory and approvals
- workflow both functional and graph-based
- integrations to model providers, tools, middleware and UI frameworks

March 2023 to April 2026



APR 3 - Microsoft Agent Framework 1.0 - Microsoft

Production-ready framework bringing together AutoGen and Semantic Kernel capabilities, with multi-agent orchestration, tools and interoperability.



APR 15 - OpenAI Agents SDK — Sandbox Agents - OpenAI

Introduces sandboxed execution, shell and file operations, skills, memory and context management for more autonomous agents.



JUN 12 - Vercel HarnessAgent - Vercel

A common interface for running different harnesses, including Claude Code, Codex and Pi, without rewriting application logic.



JUN 23 - Pydantic AI 2.0 - Pydantic

Introduces a composable, harness-first architecture combining tools, instructions, hooks and model configurations.



JUN 25 - Vercel AI SDK 7 - Vercel

Major agent-development upgrade introducing durable workflows, session resumption, harness adapters and execution tracing.



JUN 30 - Google ADK Go 2.0 - Google

Graph-based multi-agent workflows, dynamic orchestration, retry handling and human approval checkpoints.



JUL 1 - Genkit Agents API - Google

Full-stack agent API for TypeScript and Go, with persistent state, streaming, tool execution and agent coordination.



JUL 22 - Microsoft Agent Harness - Microsoft

Complete configurable runtime for planning, tools, memory, compaction, approvals, skills and telemetry.



JUL 23 - Declarative Workflows 1.0 - Microsoft

Lets developers define agent workflows through YAML, including conditional execution, handoffs and human checkpoints.



AUG 26 - Microsoft Agent Channels - Microsoft

Enables the same agent implementation to communicate through multiple interfaces, including Telegram, MCP and A2A.



SEP 10 - Vercel Copilot Harness Adapter - Vercel

Extends Vercel's interchangeable harness layer to GitHub Copilot through the Agent Client Protocol.




OCT 1 - Pi 1.0 - Earendil Works

Major independent TypeScript agent-harness milestone, strengthening extensibility, Codemod and integration capabilities.

Model Routers

Automatically selecting models, optimizing costs, managing security, and powering autonomous agents.



April 2 - [Cloudflare AI Gateway — Automatic retries](#)

Adds configurable server-side retries for upstream provider failures, including backoff strategies and per-request overrides.



April 14 - [Kong AI Gateway 3.14 + Agent Gateway](#)

Extends governance across model, tool and agent-to-agent communications, including MCP and A2A traffic.



April 21 - [Portkey Agent Gateway](#)

Introduces governed agent endpoints with centralized observability, permissions and execution oversight.



April 22-30 - [OpenRouter — Workspaces, Agent SDK, multimodal APIs](#)

Adds isolated workspaces with separate routing defaults and security controls, an agent SDK, video generation, and response caching.



May 5 - [Agent Router v0.6 - Envoy AI Gateway](#)

Graduates its v1beta1 API and expands model support, context caching, MCP and security capabilities.



May 7 - [Consistent web search and fetch - OpenRouter](#)

Provides common search and page-fetching tools across supported tool-calling models.



May 21 - [Unified AI REST API - Cloudflare](#)

Introduces common REST endpoints across Workers AI and third-party providers, including OpenAI and Anthropic-compatible APIs.



May 26 - [TrueFoundry Agent Gateway - TrueFoundry](#)

Launches an enterprise control layer for agent-to-model, agent-to-tool and agent-to-agent communications.



May 29 - [Workspace Guardrails - OpenRouter](#)

Adds centralized spend limits, provider restrictions, zero-retention requirements and prompt-injection controls.



Jun 5 - [Cost-based spend limits - Cloudflare](#)

Introduces cumulative dollar budgets enforced by model, provider and custom metadata.



Jun 10 - [Advisor tool - OpenRouter](#)

Allows a cheaper model to consult a stronger model during generation.



Jun 16 - [Subagent delegation - OpenRouter](#)

Allows a frontier model to delegate bounded work to less expensive worker models.



Jun 23 - [Envoy AI Gateway v1.0 GA - Agent Router](#)

Production milestone with stable control-plane APIs, multi-provider routing and MCP gateway support.



Jun 25 - [OpenRouter MCP Server - OpenRouter](#)

Lets coding agents access live model catalog information, benchmarks and model-testing tools.



Jul 2 - [Vercel AI Gateway — Routing Rules - Vercel](#)

Centralizes model rewrites and denials. Teams can migrate, restrict or redirect models without changing application code.



Jul 16 - [Kong AI Gateway 2.0 — Private Beta - Kong Inc](#)

Separates the AI gateway into its own runtime and release cadence, treating models, agents and MCP servers as first-class infrastructure objects.




Jul 28 - [LiteLLM v1.94.0 — Auto-Router v2](#)

Introduces router plugins and Auto-Router v2, expanding configurable model selection and MCP credential handling.

Model Routers

Automatically selecting models, optimizing costs, managing security, and powering autonomous agents.



August 4 - Not Diamond Code - Not Diamond

Launches task-aware routing for coding agents, choosing models dynamically throughout long-running agent sessions.



August 5 - Identity-aware controls and User Insights - Cloudflare

Adds authenticated-user policies, spend attribution and abnormal-usage detection.



August 7 - Workers AI unified with AI Gateway - Cloudflare

Unifies access and billing for Cloudflare-hosted models and third-party models.



August 10 - Market-Informed Auto Router - OpenRouter

Introduces automatic model selection informed by aggregate model-selection patterns.



August 21 - Envoy AI Gateway v1.1 - Agent Router

Adds token counting, per-request credentials, stream failover, MCP hostname routing and GenAI telemetry.




August 22 - LiteLLM v1.98.0 - LiteLLM

Adds auto-router shadow evaluations, provisioned throughput billing and callable routing groups.



September 1 - Kong AI Gateway 2.0 GA - Kong

Production release of the dedicated AI gateway runtime, adding MCP server bundling, agent identity controls and modality-aware cost management.



September 1 - Interactive router benchmarks - Not Diamond

Publishes a methodology for evaluating intelligent routing using interactive benchmarks. This is a research milestone rather than a new router release.



September 6 - LiteLLM v1.100.0 - LiteLLM

Adds custom auto-router tiers and expanded access-group budget controls.



September 8 - Hosted shell and Files API - OpenRouter

Lets models execute shell commands in hosted isolated containers and exchange files through supported APIs.



September 9 - Regional inference routing - OpenRouter

Adds or expands geographic routing restrictions for US and EU inference.



September 14 - LiteLLM v1.101.0 - LiteLLM

Introduces heuristic auto routing and semantic MCP tool search.



September 22 - Batch API - OpenRouter

Supports asynchronous batch inference at lower cost for workloads without interactive latency requirements.



September 27 - LiteLLM v1.103.0 - LiteLLM

Adds capability classification, Fuse v2 forecasting and further gateway authorization and spending controls.



September 29 - User Insights: model fit and savings - Cloudflare

Classifies request types, identifies potential cheaper model choices and exposes signals used by its Auto Router.



September 30 - Kong AI Gateway 2.2 - Kong

Expands provider compatibility, emerging API support and custom-interface extensibility, alongside broader agent observability.



October 2 - Web Search API - Cloudflare

Beta release bringing web search providers into the AI Gateway infrastructure.



October 6 - Agent Router v1.2 - Agent Router

First release under the Agent Router name; adds MCP authentication hardening, credential forwarding and new provider backends.



October 13 - Bifrost HTTP Transport v2.3.0 - Bifrost

Publishes a new HTTP transport component release in its modular gateway architecture. A component milestone rather than a complete new gateway launch.



October 14 - Private Model Endpoints - OpenRouter

Enables organization-private and fine-tuned model deployments to be routed through the same API as public catalog models.

Advances in Mathematics

AI systems are capable of contributing to research-level mathematics at substantial scale. Leading to most useful mathematics being cost effective to solve, advancing other fields such as physics and engineering.



MAY 2026 - AI disproves an Erdős conjecture - OpenAI · Mathematical discovery

An experimental OpenAI model produces a counterexample to the Erdős unit-distance conjecture. OpenAI subsequently reports that the result stimulated further mathematical research. Resolved 9 of 353 open Erdős problems, and proved 44 of 492 conjectures drawn from the Online Encyclopedia of Integer Sequences.



MAY 2026 - AI-generated proofs of open Erdős problems -Google DeepMind Research preprint

A research system reports resolving nine open Erdős problems and proving dozens of integer-sequence conjectures with formal verification. These are research-paper claims requiring appropriate independent scrutiny.



JUN 10 - First Proof — Research mathematics evaluation

A new evaluation of AI systems on genuine research mathematics problems, with expert assessment of the answers. Distinct from competition-style mathematics benchmarks.



AUG 1 - Ten advances in mathematics - OpenAI · Research publication

Publishes ten AI-assisted advances that reportedly resolve or substantially progress longstanding problems across combinatorics, algebra, geometry and theoretical computer science.



AUG 10 - Claude improves a Riemann zeta bound - Anthropic · New mathematical bound

An experimental Claude model improves the reported lower bound for the fraction of nontrivial zeta zeros on the critical line from 41.6% to 67.2%. Validated by two mathematicians at Anthropic, but does not solve the Riemann hypothesis.



AUG 17 - AlphaEvolve improves matrix multiplication bound - Google DeepMind and academic researchers · Theoretical mathematics

Machine-learning-assisted optimization improves the upper bound for the matrix multiplication exponent from 2.371339 to 2.371177, a theoretical advance in algorithmic complexity.



SEPT 4 - Claude formalizes Fermat's Last Theorem - Anthropic · Computer-verified proof

Claude works largely autonomously for 11 days to produce a complete Lean-checked formalization of Fermat's Last Theorem, generating approximately 13 million lines of Lean and proving 29,500 intermediate theorems.



SEPT 8 - Proposed solution to the Navier–Stokes Millennium Problem - OpenAI · Major claim

OpenAI announces an AI-generated proof concerning finite-time singularities in the Navier–Stokes equations, with a Lean formalization. Its status and attribution have attracted scrutiny; formal prize recognition is separate.



OCT 6 - OpenAI publishes 722 mathematical manuscripts - OpenAI

Releases 722 AI-generated manuscripts covering 372 groups of mathematical results. The unprecedented volume creates both a major research opportunity and a substantial independent-verification challenge; the manuscripts are not equivalent to 722 accepted solutions.



OCT 14 - First Proof — Third evaluation

Reports strong AI performance on previously withheld, genuine research mathematics problems, assessed by expert mathematicians. This measures problem-solving rather than the discovery of entirely new open-problem solutions.



OCT 15 - Funding for independent mathematical validation - OpenAI / Simons Foundation

OpenAI announces a \$5 million contribution toward an initiative supporting mathematicians examining, understanding and building upon AI-generated results.



OCT 16 - Claude designs more efficient quantum circuits - Anthropic / Harvard collaborators

Claude discovers improved quantum circuit implementations across eight computational tasks, sometimes reducing resource requirements by an order of magnitude. Includes mathematical reasoning and a Lean-checked algebraic result.

AND MORE TO COME

Physics

The defining theme of these six months is that AI is becoming part of the scientific method itself—helping control experiments, adapt quantum hardware, and develop models of complex physical systems. Physics will build on math capabilities of models.



May 5, 2026 - Quantum - Quantum computing simulates a large protein complex

IBM / research collaborators · Research milestone

Researchers combine quantum and classical supercomputing to investigate a protein complex involving 12,635 atoms. An important step in exploring quantum-assisted molecular simulation, not proof that a quantum computer simulated the entire protein independently.



July 6, 2026 - Quantum - Quantum computing applied to fusion fuel materials

Oak Ridge National Laboratory · IBM · Cleveland Clinic · Quantum + fusion

A hybrid quantum-classical method calculates nine molecular configurations of materials relevant to extracting tritium for fusion reactors.



June 2, 2026 - Quantum - IBM commits \$10 billion to quantum computing

Announces a five-year investment in research, manufacturing, infrastructure and quantum technology development. A significant commercialization commitment, rather than an experimental breakthrough.



July 8, 2026 - Quantum - AI learns to correct quantum errors during computation

Quantum computing researchers · Research breakthrough

A Nature study demonstrates reinforcement-learning control of quantum error correction, allowing calibration to adapt during computations instead of requiring complete interruption.



July 13, 2026 - Quantum - Quantum error correction inside a single ion

MIT / research collaborators · Research breakthrough

A Nature Physics paper demonstrates error correction of a logical qubit encoded in one atomic ion, exploring a different route to fault-tolerant quantum hardware.



July 30, 2026 - Quantum - IBM demonstrates verifiable quantum advantage

IBM / University of Chicago / Qedma · Research breakthrough

Two studies report classically difficult quantum computations with validation techniques, including a 70-logical-qubit circuit demonstration and quantum-material simulations involving up to 74 physical qubits.



September 2, 2026 - Fusion + AI - PACMAN AI controls a fusion reactor in real time

Princeton Plasma Physics Laboratory · Experimental validation

An AI control framework is tested in five experiments on the DIII-D tokamak. It coordinates heating, plasma stability and other controls, including predicting one instability approximately 200 milliseconds before onset.



October 8, 2026 - Fusion + AI - AI digital twin selected for SPARC fusion reactor

Commonwealth Fusion Systems / Google DeepMind / NVIDIA / US laboratories · Major research initiative

A US Department of Energy initiative selects CFS to lead an AI-enabled digital twin of SPARC, integrating simulation, optimization and research data to improve reactor operations. The project is planned, not yet a completed fusion-power demonstration.

Medicine

AI is influencing the full pharmaceutical pipeline, from target discovery and molecule design to clinical trials, regulatory development and scientific operations. Rentosertib is the clearest clinical milestone; Isomorphic Labs and Claude's protein-design work illustrate scientific progress; and the large pharma partnerships show the scale of industry adoption. **Logevity Escape Velocity (LEV)** is within reach 2035-2040 is looking more likely.



May • [Isomorphic Labs raises \\$2.1 billion](#)

Demis Hassabis's AI drug-discovery company raises funding to scale its IsoDDE engine, expand research operations and progress its therapeutic pipeline toward clinical testing.



May • Bristol Myers Squibb × Anthropic

BMS expands Claude use across more than 30,000 employees for research, development and operational workflows. A significant example of frontier AI becoming part of pharmaceutical infrastructure.



Jun 3 • [GPT-Rosalind model upgrade — OpenAI](#)

New specialized capabilities for medicinal chemistry, genomics, scientific reasoning, biological data analysis and experiment planning. GPT-Rosalind originally launched on April 16, just outside our six-month window.



Jul 2 • Takeda × Insilico Medicine

A collaboration potentially worth up to \$600 million plus royalties, using Insilico's Pharma.AI platform to discover and optimize candidates across Takeda's therapeutic interests.



Jul 29 • [ISM6331 receives FDA Fast Track designation](#)

Insilico's AI-designed pan-TEAD inhibitor receives Fast Track designation for a specified group of patients with advanced malignant pleural mesothelioma. Fast Track can facilitate development and review, but is not marketing approval.



Jul 30 • GSK × Relation Therapeutics

An agreement worth up to \$110 million involving large-scale experimental human-cell data and machine learning to discover and evaluate drug targets.



Aug 18 • [Claude designs experimentally validated protein binders](#)

Anthropic reports that Claude orchestrated specialist protein-design tools to produce binders against 14 of 15 targets, with experimental testing by external laboratories. Some designs matched or exceeded previously published binding affinities.



Sep 7 • [Rentosertib shows a potential biological-aging signal](#)

An exploratory analysis of 42 patients in a Phase 2a pulmonary fibrosis trial found improvements in predicted biological age across six proteomic aging clocks, alongside encouraging lung-function findings. This does not demonstrate that rentosertib slows human aging or establish an anti-aging treatment.



Sep 10 • [Rentosertib enters Phase 3 — GENESIS-IPF-3](#)

Insilico Medicine doses the first patient in its Phase 3 trial of rentosertib for idiopathic pulmonary fibrosis. The randomized, double-blind, placebo-controlled study is intended to evaluate safety and efficacy over 52 weeks. This is one of the strongest examples yet of a generative-AI-designed drug advancing into late-stage clinical testing.



Sep 11 • [GPT-Rosalind expands beyond research preview](#)

OpenAI makes its specialized life-sciences model available globally to eligible organizations through a trusted-access program, supporting biology, drug discovery, genomics and laboratory workflows.



Sep 17 • Autonomous biomedical research - [Stanford's Virtual Biotech](#)

Researchers describe a multi-agent system using tens of thousands of specialized AI agents to support therapeutic research, analyze evidence and investigate experimental strategies. This is a research demonstration, not an AI-developed approved medicine.



Sep 29 • [Isomorphic Labs expands its AI medicine-design strategy](#)

Demis Hassabis's company describes an increasingly integrated approach connecting protein structure, molecule generation, predictive chemistry and experimental data. An important earlier result, published February 10, reported that IsoDDE more than doubled AlphaFold 3's accuracy on a specific challenging protein-ligand benchmark. That original result falls outside this six-month timeline.



Oct 15 • [AstraZeneca × insitro — up to \\$1.5 billion](#)

\$40 million upfront, with potential milestone payments of up to \$1.46 billion, to use insitro's Virtual Human platform for discovering targets in liver and cardiovascular diseases.



Oct 15 • Scientific access - [Rosalind Academic Plan](#)

OpenAI extends sponsored access to GPT-Rosalind and its research tools for eligible academic laboratories, broadening availability beyond major commercial organizations.



Oct 19 • Drug-delivery discovery - [Eli Lilly × Metis TechBio — up to \\$1.1 billion](#)

Uses an AI-driven platform to discover and optimize lipids for tissue-targeted nanoparticle drug delivery. This is about improving the delivery of medicines, rather than directly discovering a new therapeutic molecule.

Robotics

Robotics evolution is following the same path as language models: first foundational models, then accessible inference, then developer ecosystems. Robotics appears to be developing a similar stack, although its real-world reliability and hardware integration challenges remain substantially greater



Figure



NVIDIA



Botson
Dynamis/
Hyundai



Gemini
Robotisi



Tesla



Unitree



DJI



MAY 8 - Helix 02 — Collaborative humanoids Multi-Robot Intelligence

Two humanoids coordinate a bedroom-tidying task using a shared learned policy rather than explicit inter-robot communication.



MAY 31- NVIDIA Cosmos3 · World foundation model

An open world foundation model unifying physical-world generation, visual reasoning and action prediction.



MAY 31 - NVIDIA Isaac GR00T Reference Humanoid

An open research robot configuration using Unitree hardware, dexterous Sharpa hands, Jetson Thor and GR00T software.



JUN 1 - Unitree H2 Plus × NVIDIA GR00T

Unitree introduces the H2 Plus research platform, combining its humanoid hardware with Sharpa dexterous hands, NVIDIA Jetson Thor and Isaac GR00T.



JUL 30 - Gemini Robotics 2

Google DeepMind · Robot foundation model - Advances unified perception, reasoning, dexterous manipulation and whole-body control across robot types.



AUG 19 - Generalist GEN-1.5 - Generalist AI · One-shot robot learning

Learns new manipulation tasks from 3–12 seconds of demonstration. Reports 59% success without gradient updates and 83% with brief adaptation.



AUG 25 - Figure AI · Robot training data

Large-scale human-behavior dataset and training system designed to improve humanoid generalization.



SEP 17 - Helix 2.5 — Unseen-home generalization

Figure AI · Zero-shot autonomy

Reports a 56% full-task success rate across unfamiliar homes, compared with 9% without Index pretraining.



SEP 21 - Atlas factory training


Boston Dynamics / Hyundai · Factory deployment preparation

Atlas enters a real-world industrial training environment for manufacturing logistics and component handling.



Sep 28 - Unitree UnifoLM-WLA-1.0 - Unitree Robotics

Releases base model and fine-tuning code for a 6B-parameter embodied foundation model trained on approximately 2,500 hours of robot data, supporting 64 tasks.



OCT 13 - Physical Intelligence Pi API beta

Physical Intelligence · Robot intelligence as a service

Makes robotic foundation models available for hosted fine-tuning and inference through an API.



OCT 14 - NEURA Neuron AI · Multimodal physical AI

Announces multimodal robot intelligence spanning vision, touch, hearing and memory, alongside a planned 11-minute autonomous laundry demonstration.



Oct 14 - DJI Dock 4 + Matrice 5TD - DJI

Launches automated 25-second battery swapping for persistent drone operations, with AI detection of 17 object categories.



OCT 19 - Dream Labs Dream-1 / Dream API · Robot intelligence as a service

A robot foundation-model API advertised for previously unseen two-arm robots, using camera observations and language instructions.



OCT 19 - KAIST Humanoid v0.9 · Whole-body manipulation

Demonstrates vision-guided pushing of 35 kg objects and lifting of 20 kg objects while walking, without external motion capture.