

DF Labs

AI NEWS UPDATES

March 30th, 2025




TABLE OF CONTENTS

- **News**
- **Tools**
- **Papers**

NEWS

TheRundownAI, X, etc

Hollywood against AI copyright proposals

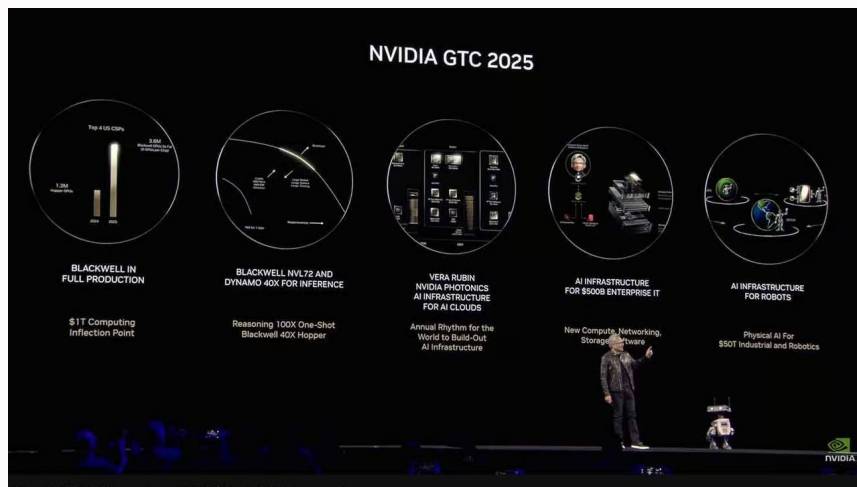


More than 400 Hollywood creatives signed an open letter urging the Trump administration to reject OpenAI and Google's proposals to expand AI training on copyrighted works—arguing that it would let them "freely exploit" creative industries.

- The letter is a direct response to OpenAI and Google's AI Action Plan submissions, which argued for expanded fair use protections for AI training.
- OpenAI framed AI copyright exemptions as a "matter of national security," while Google said the current fair use framework already supports AI innovation.
- Ben Stiller, Mark Ruffalo, Cate Blanchett, Paul McCartney, Taika Waititi, and Aubrey Plaza are among the high-profile creatives who have signed the letter.
- They have emphasized that AI companies could simply "negotiate appropriate licenses with copyright holders — just as every other industry does."

Document: [Hollywood's response to AI Action Plan](https://s3.documentcloud.org/documents/25590860/hollywoods-response-to-ai-action-plan.pdf)

Nvidia's AI and robotics advances



Jensen Huang kicked off Nvidia's GTC 2025 conference with a two-hour-long keynote, calling the event "AI's Super Bowl" and revealing exciting new updates on upcoming chip releases, robotics, autonomous vehicles, and more.

- Nvidia's new GPU lineup includes Blackwell Ultra in late 2025, Vera Rubin in 2026, and Feynman in 2028, each promising major performance gains.
- Huang said that scaling is not slowing down, and the computation needed for AI is "easily 100x more than we thought we needed at this time last year."
- He also revealed Isaac GROOT N1, the first open humanoid robot foundation model, alongside a comprehensive physical AI dataset for training robots.
- The new DGX Spark and DGX Station bring data center-grade AI computing to personal workstations, with Huang calling it "the computer for the age of AI."
- A robotics physics engine, Newton, also debuted in collaboration with Google DeepMind and Disney — demoed with 'Blue', a Star Wars-style robot on stage.
- Nvidia also announced a new partnership with automaker GM, with plans to build the company's first fleet of self-driving cars.

Video: [GTC March 2025 Keynote with NVIDIA CEO Jensen Huang](https://www.youtube.com/watch?v=waPvOwl9Z8&t=1885s)

OpenAI nears \$40B funding round



OpenAI is reportedly finalizing a massive \$40B funding round led by SoftBank, which would make it the largest private funding in history — and nearly double the ChatGPT maker's valuation to \$300B.


- SoftBank will invest an initial \$7.5B, followed by another \$22.5B later this year with other investors including Magnetar Capital, Coatue, and Founders Fund.
- OpenAI expects its revenue to triple to \$12.7B in 2025 and become cash-flow positive by 2029 with over \$125B in projected revenue.
- The company reportedly lost as much as \$5B on \$3.7B of revenue in 2024, attributed to AI infrastructure and training costs.
- The funding will also partially support OpenAI's commitment to Stargate, the \$300B AI infrastructure JV announced with SoftBank and Oracle in January.

Article: [OpenAI Close to Finalizing \\$40 Billion SoftBank-Led Funding - Bloomberg](https://www.bloomberg.com/news/articles/2025-03-26/openai-close-to-finalizing-its-40-billion-softbank-led-funding)

TOOLS

TheRundownAI, Hugging Face, GitHub, etc

OpenAI adds image generation to GPT-4o, Sora



OpenAI released image generation within its GPT-4o model and Sora video generator, shifting from separate text and image systems to a fully integrated approach for producing more precise and contextually aware visuals via ChatGPT.

- GPT-4o treats images as part of its multimodal understanding, enabling more accurate text rendering and contextual awareness.
- The upgrade excels at generations like menus, diagrams, and infographics with readable text, addressing a major weakness of previous models.
- Users can also edit images with natural language, with the model able to maintain consistency between iterations and handle 10-20 objects in prompts.
- The new capability replaces DALL-E 3 as ChatGPT's default image generator for Free, Plus, Pro, and Team users, with Enterprise and Edu coming soon.

Article: [Introducing 4o Image Generation | OpenAI](https://openai.com/index/introducing-4o-image-generation/)

Reve's new leading image model



Reve just emerged from stealth with Reve Image 1.0, a new text-to-image AI model that topped global rankings with the codename “Halfmoon” over the last week—showcasing exceptional prompt accuracy, text rendering, and image quality.

- The model claimed the #1 position in Artificial Analysis' Image Arena, outperforming rivals like Google's Imagen 3, Midjourney v6.1, and Recraft V3.
- Reve said its mission is to “enhance visual generative models with logic,” with 1.0 showing impressive prompt adherence and long text rendering in tests.
- The platform also features natural language editing, photo uploads, and an ‘explore’ tab to view community prompts and generations.
- A preview of Reve Image 1.0 is currently free to try (though no API access yet), with the company saying that “much more is coming soon”.

Post: [Reve on X](#)

Ideogram's advanced 3.0 image model



Image generation startup Ideogram just released version 3.0 of its AI model, introducing major improvements in photorealism, text rendering, and style consistency — while outperforming competitors in human evaluations.

- Ideogram 3.0 brings new text rendering and graphic design capabilities, enabling precise creation of complex layouts, logos, and typography.
- In testing, the model significantly outperformed leading text-to-image models, including Google's Imagen 3, Flux Pro 1.1, and Recraft V3.
- A new 'Style References' feature allows users to upload up to three images to guide the aesthetic of generated content, alongside a library of 4.3B presets.
- The model is now available on Ideogram's platform and iOS app, with all features accessible to free users.

Article: [Ideogram 3.0](https://about.ideogram.ai/3.0)

Claude gets real-time web search

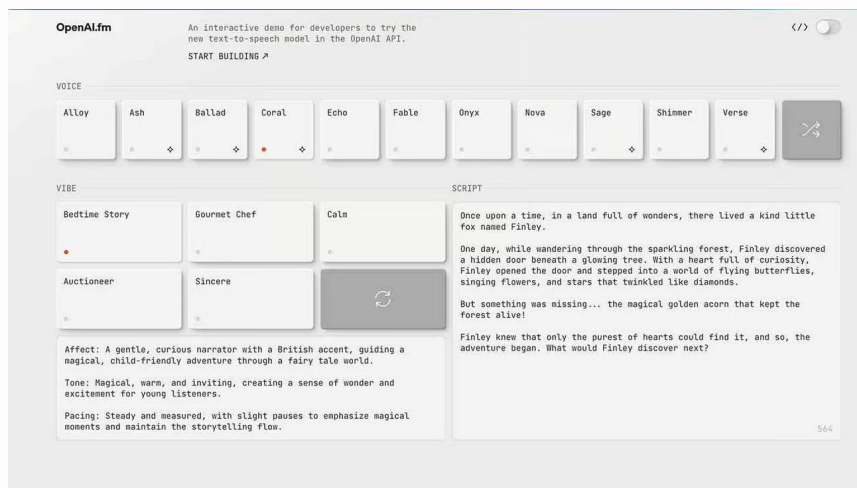


Anthropic just added web search capabilities to Claude, giving its AI assistant access to real-time information and closing a major feature gap with competitors like ChatGPT and Gemini.

- Web search integrates directly with Claude 3.7 Sonnet and automatically determines when to surf the internet for more current or accurate information.
- Claude provides direct citations for web-sourced information, allowing users to verify sources and fact-check responses easily.
- The feature is now available to all paid Claude users in the U.S., with international and free-tier expansion planned for the near future.
- Users can also access the feature by toggling on the 'Web Search' tool in the profile settings of the platform.

Article: [Claude can now search the web \ Anthropic](https://www.anthropic.com/news/web-search)

OpenAI's voice AI with personality boost



OpenAI launched its next-gen API-based audio models for text-to-speech and speech-to-text, giving developers the ability to customize AI speaking styles via text and delivering improved speech recognition across multiple languages.

- The new gpt-4o-mini-tts model adapts its speaking style based on simple text prompts — like "speak like a pirate" or "use a bedtime story voice."
- The GPT-4o-transcribe speech-to-text models reach SOTA performance across accuracy and reliability tests, outperforming existing Whisper models.
- OpenAI also released openai.fm, a public demo platform allowing users to test different voice styles and experience the new models firsthand.
- The models are available through OpenAI's API, with integration support through the Agents SDK for developers building voice-enabled AI assistants.

Article: [Introducing next-generation audio models in the API | OpenAI](https://openai.com/index/introducing-our-next-generation-audio-models/)

Roblox releases open-source 3D generation AI

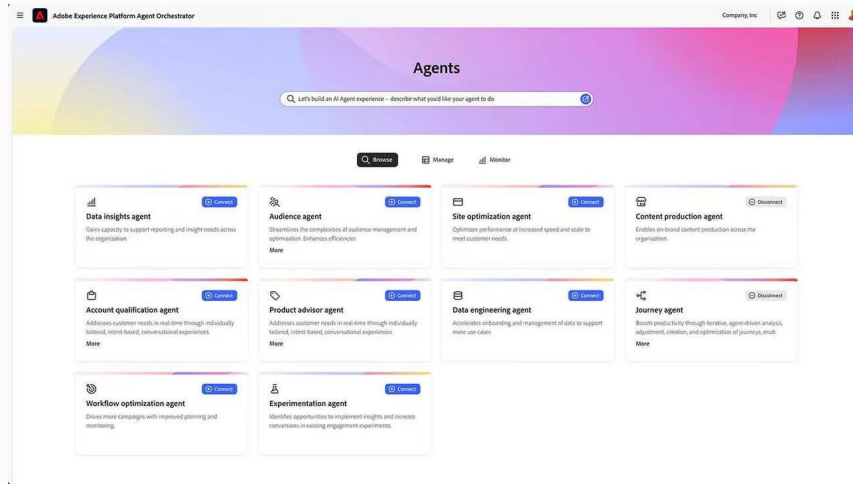


Roblox just announced Cube 3D, a new open-source AI system for generating 3D objects and scenes from text prompts — alongside a slew of other tools and updates for AI-assisted game development.

- Cube 3D generates complete, functional 3D objects from text prompts, training on native 3D data instead of traditional image-based reconstruction.
- Developers can generate assets through simple commands like "/generate motorcycle," with image input capabilities also coming in the future.
- Cube uses '3D tokenization' to predict and generate shapes the same way language models predict text, enabling future 4D scene generation capabilities.
- Roblox also released updates to its Studio content creation suite including improved performance, real-time collaboration features, and monetization tools.

Article: [Introducing Roblox Cube: Our Core Generative AI System for 3D and 4D](https://corp.roblox.com/newsroom/2025/03/introducing-roblox-cube)

Adobe's army of enterprise AI agents



Adobe just released a comprehensive AI agent strategy centered around its new Experience Platform Agent Orchestrator, introducing ten specialized agents for enterprise tasks like customer experiences, marketing workflows, and more.

- The platform includes agents for audience targeting, content production, site optimization, and B2B account management within Adobe's enterprise apps.
- A Brand Concierge helps businesses create personalized chat experiences, with traffic from AI platforms to retail sites jumping 1,200% in February.
- Adobe Marketing Agent also integrates with Microsoft 365 Copilot, allowing teams to access Adobe's agentic capabilities directly within Microsoft apps.
- The company has also partnered with tech firms like AWS, Microsoft, SAP, and ServiceNow, allowing its agents to work across third-party enterprise systems.

Article: [Adobe Summit 2025: Adobe AI Platform Unites Creativity and Marketing to Define the New Era of Customer Experience Orchestration](https://news.adobe.com/news/2025/03/adobe-summit-2025-adobe-ai-platform-unites-creativity-marketing)

Zoom's AI Companion goes agentic

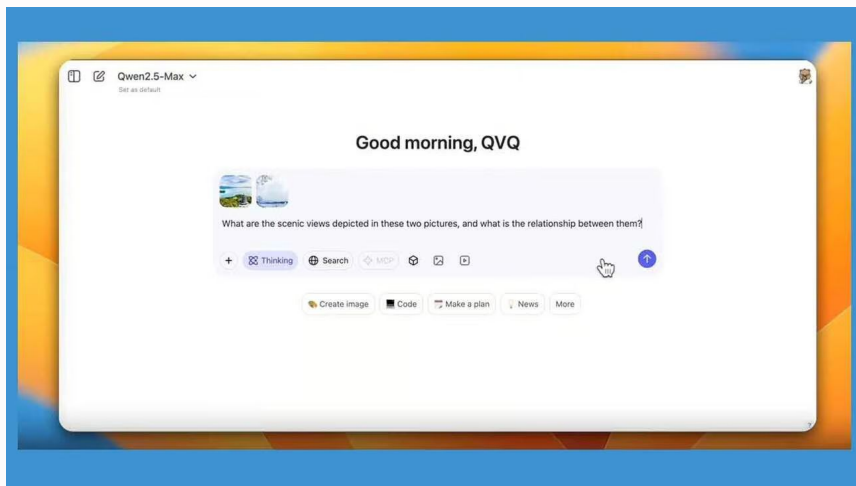


Zoom announced a major upgrade to its AI Companion, introducing new agentic capabilities and skills that allow the assistant to identify and complete tasks across the platform's ecosystem.

- AI Companion is getting new memory and reasoning capabilities to problem-solve and leverage the correct agentic tools for each task.
- Zoom Tasks automatically detects and executes action items from meetings, scheduling follow-ups, or generating documents without user intervention.
- Other new features include calendar management, clip generation, writing assistance, voice recording transcriptions, and live notes for meetings.
- A new \$12/month Custom AI Companion add-on is also launching in April, offering personal AI coaches, AI avatars for video messages, and more.

Article: [Zoom debuts new agentic AI skills and agents for Zoom AI Companion | Zoom - Zoom](#)

Qwen's QVQ-Max visual reasoning model

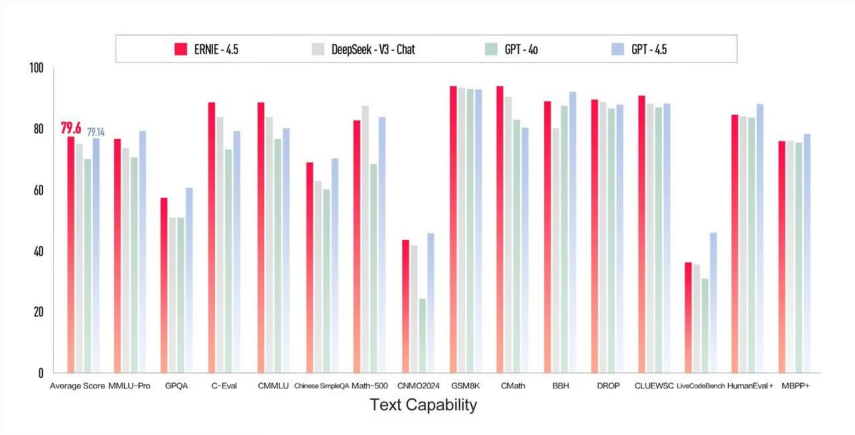


Alibaba's Qwen team released QVQ-Max, a new visual reasoning model that goes beyond basic image recognition to analyze and reason about visual information across images and videos.

- The model is an evolution of QVQ-72B-Preview, expanding capabilities across mathematical problem-solving, code generation, and creative tasks.
- QVQ-Max features a "thinking" mechanism that can be adjusted in length to improve accuracy, showing scalable gains as thinking time increases.
- Other complex visual capabilities shown include analyzing blueprints, solving geometry problems, and providing feedback on user-submitted sketches.
- Qwen said that future plans include creating a complete visual agent capable of operating devices and playing games.

Article: [QVQ-Max: Think with Evidence | Qwen](#)

Baidu launches ultra-cheap AI models



Chinese tech giant Baidu just released two aggressively priced multimodal AI models — ERNIE 4.5, a significant upgrade to their foundational model, and a new deep-think-capable model called Ernie X1.

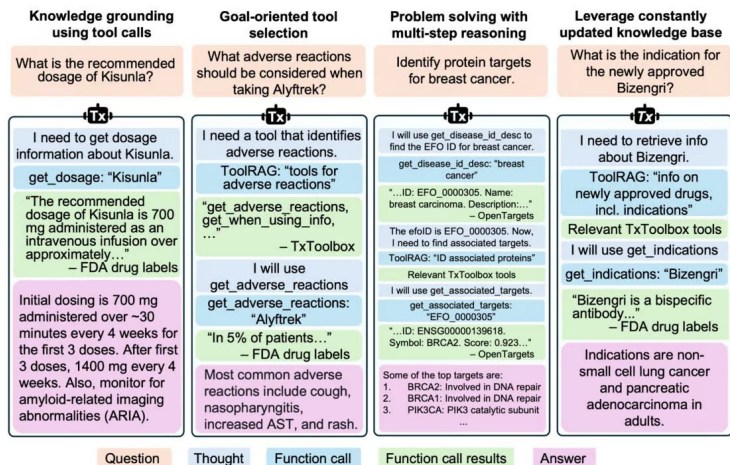
- ERNIE 4.5 stands out with enhanced “EQ” and language skills as well as improved hallucination prevention, logical reasoning, and coding abilities.
- According to Baidu, ERNIE 4.5 outperforms GPT-4.5 across multiple benchmarks while costing just 1% of its price — approximately \$0.55 and \$2.20 per million input and output tokens.
- ERNIE X1, Baidu’s first reasoning model, matches the capabilities of top Chinese competitor DeepSeek’s R1 at half the price.
- Like DeepSeek’s R1, ERNIE X1 uses a step-by-step “thinking” approach, excelling in complex calculations and tasks like document understanding.

Post: [Baidu Inc. on X](#)

RESEARCH

TheRundownAI, arXiv, Hugging Face, etc

Harvard team creates AI agent for personalized medicine

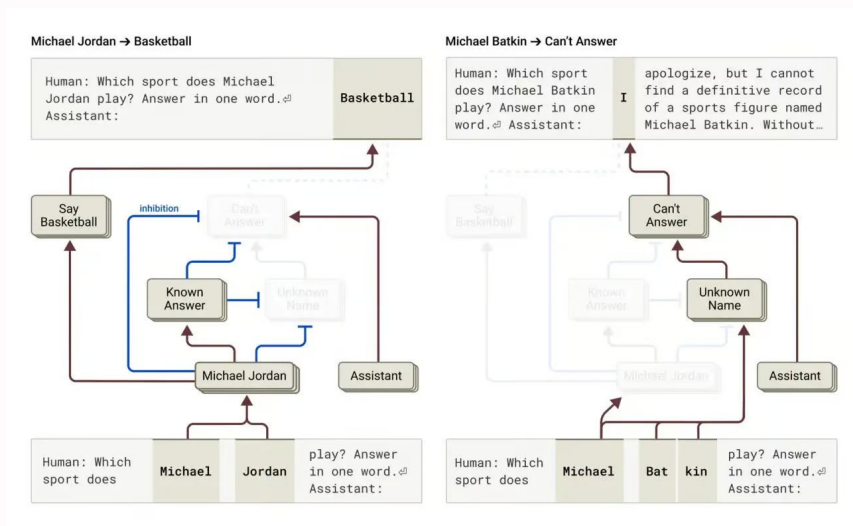


Researchers from Harvard and MIT just introduced TxAgent, an AI agent that leverages multi-step reasoning and real-time biomedical knowledge retrieval to generate trusted, personalized treatment recommendations for patients.

- TxAgent utilizes 211 specialized tools to analyze drug interactions and contraindications and suggest patient-specific treatments in real time.
- It evaluates drugs at molecular, pharmacokinetic, and clinical levels, identifying risks based on comorbidities, ongoing medications, age, and genetic factors.
- The system synthesizes evidence from biomedical sources, iteratively reasoning and refining recommendations through structured function calls.
- TxAgent's toolkit is available as "ToolUniverse" featuring trusted data sources like openFDA and Open Targets to ensure validated medical insights.

Publication: [TxAgent: An AI Agent for Therapeutic Reasoning Across a Universe of Tools](#)

Anthropic reveals how Claude 'thinks'



Anthropic released two research papers that reveal how its AI assistant Claude processes information, helping to better understand internal mechanisms that explain capabilities like multilingual reasoning and advanced planning.

- The researchers developed an "AI microscope" that reveals internal "circuits" in the model, showing how Claude transforms input to output in key tasks.
- Claude uses a universal "language of thought" across different languages, with shared conceptual processing for English, French, and Chinese.
- When writing poetry, Claude plans ahead several words, identifying rhyming options before constructing lines to reach those planned words.
- The team also discovered a default that prevents speculation unless overridden by strong confidence, helping explain how hallucination prevention works.

Publication: [Tracing the thoughts of a large language model](https://www.anthropic.com/research/tracing-thoughts-language-model)
[Anthropic](https://www.anthropic.com/research/tracing-thoughts-language-model)

DF Labs