

DF Labs

AI NEWS UPDATES

January 19th, 2025




TABLE OF CONTENTS

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

NEWS

TheRundownAI, X, etc

Omi debuts brain-reading AI wearable



Based Hardware just introduced Omi, an \$89 AI wearable that combines always-on listening capabilities with brain-interface tech to handle productivity tasks — with hopes to enable thought-reading type abilities in the future.

- Omi can be worn as a necklace or attached to the temple to enable early brain-interface features that detect when the AI is addressed without a 'wake' word.
- The device listens continuously to provide real-time summaries, meeting notes, and contextual information, with a battery life of approximately 3 days.
- The company is taking an open-source approach, with over 250 apps already available in its store and integration of AI models from OpenAI and Meta.
- While brain-interface features are currently limited to intent detection, the founder envisions more advanced thought-reading capabilities within 2 years.
- Omi was initially 'Friend' before a device launched with the same name, with founder Nik Shevchenko publishing a diss track calling his 'the real friend'.

Post: [Nik Shevchenko on X](#)

Altman: 'Confident we know how to build AGI'




OpenAI CEO Sam Altman just posted a new blog titled 'Reflections', revealing that the company believes they now know how to build AGI — and is now setting its sights on developing superintelligent systems.

- Altman stated that OpenAI is "now confident we know how to build AGI", also predicting that the first AI agents will join the workforce in 2025.
- OAI is now aiming for superintelligence, which Altman says may revolutionize scientific discovery and "massively increase abundance and prosperity."
- Altman also addressed the November 2023 leadership crisis, describing his sudden firing as "a big failure of governance by well-meaning people."
- The blog follows Altman's cryptic post about the technological singularity that we highlighted in yesterday's newsletter.

Article: [Reflections - Sam Altman](https://blog.samaltman.com/reflections)

OpenAI building out new robotics division




OpenAI just posted the company's first robotics hardware job listings, revealing plans to develop its own custom robots with advanced AI capabilities.

- Former Meta AR glasses lead Caitlin Kalinowski is spearheading the effort, joining as OpenAI's hardware director in November.
- The company is hiring for technical roles, including sensor suite development and mechanical design, and a lab operations manager to oversee prototype testing.
- Job listings hint at goals for 'general-purpose robots that operate in dynamic real-world settings,' with plans for a 'wide variety of robotic form factors.'
- OpenAI shuttered the robotics team in 2020, with research including training a robotic hand to solve a Rubik's Cube and other dexterity challenges.
- OpenAI has also collaborated with Figure in the past year, integrating its models into the robotics firm's humanoid robots.

Post: [Caitlin Kalinowski](#) 🇺🇸 on X

<https://x.com/kalinowski007/status/1877809579154948223>

Microsoft plans \$80B AI infrastructure investment




Microsoft just published a new blog outlining an \$80B investment plan for AI data center infrastructure in 2025, alongside a detailed overview of the company's vision for American leadership in AI.

- Over half of the \$80B investment will be focused on U.S.-based AI data centers, an increase from last year's \$50B spent on infrastructure.
- President Brad Smith outlined a three-pronged strategy: advancing American AI tech, expanding workforce AI skills, and promoting U.S. AI exports globally.
- Smith also advocated for balanced government oversight that protects security while allowing U.S. companies to compete against Chinese competitors.

Article: [The Golden Opportunity for American AI - Microsoft On the Issues](https://blogs.microsoft.com/on-the-issues/2025/01/03/the-golden-opportunity-for-american-ai/)

Nvidia rings in 'Age of AI Agentics' at CES 2025



Nvidia just unveiled a sweeping vision for AI agents alongside major hardware, software, and robotics reveals at CES 2025, with CEO Jensen Huang declaring AI agents a 'multi-trillion-dollar' opportunity.

- Huang introduced the new RTX Blackwell GPU family, with the \$2,000 5090 chip hailed as the 'world's fastest GPU' outperforming its predecessor by 2x.
- Nvidia revealed 'Project Digits', a \$3,000 personal computer powered by the GB10 Superchip that is 1000x more powerful than the average laptop.
- Cosmos, an open platform of world foundation models for physical AI, is freely available for robotics and autonomous vehicle development.
- Nvidia introduced Llama Nemotron and Cosmos Nemotron model families, designed specifically for agentic AI applications.
- A new early access blueprint for AI agents enables video and image analysis while integrating agentic features like reasoning, tool calling, and more.
- A significant partnership with Toyota was also revealed, with plans to integrate NVIDIA's AI systems into autonomous vehicle development.

Article: [NVIDIA CEO Jensen Huang Keynote at CES 2025](https://www.youtube.com/watch?v=k82RwXqZHY8)

Samsung goes all-in on AI at CES 2025



Samsung revealed its new "AI for All" tagline at CES 2025, introducing a comprehensive suite of new AI features and products across its entire ecosystem — including new AI-powered TVs, appliances, PCs, and more.

- Vision AI brings features like real-time translation, the ability to adapt to user preferences, AI upscaling, and instant content summaries to Samsung TVs.
- Several of Samsung's new Smart TVs will also have Microsoft Copilot built in, while also teasing a potential AI partnership with Google.
- Samsung also announced the new line of Galaxy Book5 AI PCs, with new capabilities like AI-powered search and photo editing.
- AI is also being infused into Samsung's laundry appliances, art frames, home security equipment, and other devices within its SmartThings ecosystem.

Article: [Samsung Unveils Vision AI and New Innovations at CES First Look 2025, Delivering Personalized, AI-Powered Screens to Enrich Everyday Life](https://news.samsung.com/us/samsung-unveils-vision-ai-new-innovations-ces-first-look-2025-delivering-personalized-ai-powered-screens-to-enrich-everyday-life/)

Panasonic, Anthropic team up on AI wellness



Japanese electronic giant Panasonic just announced an AI transformation initiative alongside a strategic partnership with Anthropic and a new AI wellness app, aiming to reinvent the company with AI-powered products and services.

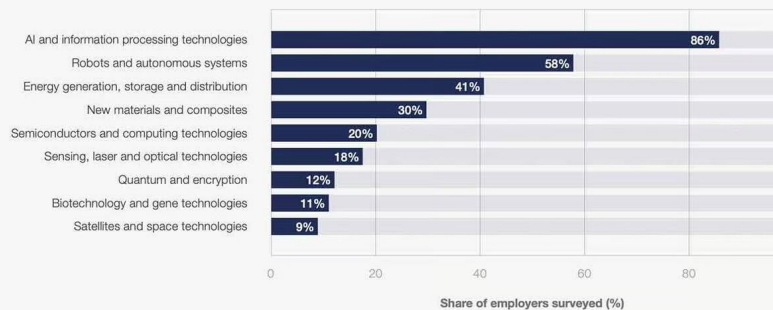
- 'Panasonic Go' is the electronics giant's plan to transition to an AI-first business model, aiming to generate 30% of revenue from AI by 2035.
- The company introduced Umi, an AI wellness coach built on Anthropic's Claude to help families manage routines, set goals, and coordinate care.
- Umi will launch in the US market by mid-2025, featuring integrations with wellness brands like Calm, Blue Apron, and Precision Nutrition.
- The Anthropic partnership extends internally, with plans to deploy Claude across Panasonic's workforce.

Article: [Panasonic Well Sets the Stage for the Future of Family Wellness in Opening Keynote at CES 2025](https://www.prnewswire.com/news-releases/panasonic-well-sets-the-stage-for-the-future-of-family-wellness-in-opening-keynote-at-ces-2025-302344641.html)

World Economic Forum forecasts AI workplace surge

FIGURE 1.2 **Technology trends driving business transformation, 2025-2030**

Share of employers surveyed that identify the stated technology trend as likely to drive business transformation



The World Economic Forum just released its 2025 Future of Jobs Report, revealing AI's dramatic impact on the workforce — with an overwhelming majority of employers expecting significant transformations from the tech by 2030.

- Technology adoption is surging, with 86% of companies expecting AI to transform their operations by 2030.
- AI is predicted to create 11M jobs while displacing 9M others, with big data specialists and AI/ML experts topping the list of fastest-growing roles globally.
- Three-quarters of organizations plan to upskill existing employees for AI collaboration, while 70% aim to hire new staff with AI experience.
- Half of companies expect to reorient their business around AI opportunities, while 40% anticipate reducing workforce size as AI capabilities grow.

Report: [WEF Future of Jobs Report 2025.pdf](https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf)

François Chollet finds new AGI lab




François Chollet, former Google researcher and the creator of the popular Keras AI framework, just introduced Ndea, a new AI lab aiming to achieve AGI through an alternative research method, alongside Zapier founder Mike Knoop.

- Ndea's core strategy combines deep learning with program synthesis, aiming to create AI that can learn and adapt with human-level efficiency.
- The startup positions itself as an alternative to the dominant large-scale deep learning approach, arguing that training data limits current AI.
- Ndea plans to build what they call a "factory for rapid scientific advancement," focusing on both known frontiers like drug discovery and unexplored territories.
- Chollet also recently launched the ARC Prize Foundation, a nonprofit that is developing benchmarks to evaluate human-level AI capabilities.

Post: [Francois Chollet on X](#)

AI tutoring shows stunning results



A new study in Nigeria just revealed that students using AI as an after-school tutor made learning gains equivalent to two years of traditional education in just six weeks — showcasing the power of AI-driven learning in developing regions.

- The World Bank-backed pilot combined AI tutoring with teacher guidance in an after-school setting, focusing primarily on English language skills.
- Students significantly outperformed their peers in English, AI literacy, and digital skills, with the impact extending to their regular school exams.
- The intervention showed huge improvements, particularly for girls who were behind, suggesting AI tutoring could help close gender gaps in education.
- The program impact also increased with each additional session attended, suggesting longer programs might yield even greater benefits.

Article: [From chalkboards to chatbots: Transforming learning in Nigeria, one prompt at a time](https://blogs.worldbank.org/en/education/From-chalkboards-to-chatbots-Transforming-learning-in-Nigeria,one-prompt-at-a-time)

Apple pulls AI news summaries after false headlines



Apple just temporarily disabled its AI-powered news summary feature in Apple Intelligence after multiple instances of the system generating completely false headlines, including fabricated stories about arrests and deaths that never happened.

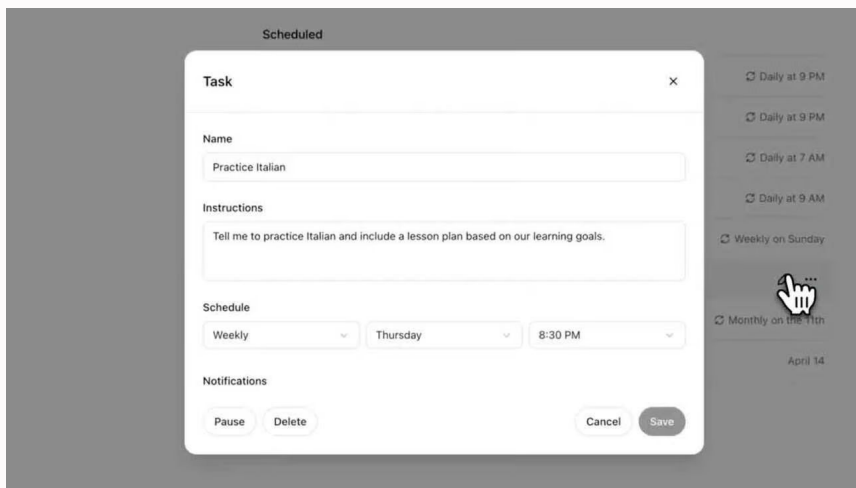
- The feature launched in September with the iPhone 16 and was intended to condense multiple news notifications into brief summaries.
- Major news organizations, including the BBC and the Washington Post, complained that the feature contradicted original reporting and undermined trust.
- The BBC complained about the feature as early as December, urging Apple to remove it due to critical factual errors in breaking news reporting.
- Apple said it plans to make AI-generated summaries more clearly labeled and give users more control over which apps can use the summarization feature.

Article: [Apple is pulling its AI-generated notifications for news after generating fake headlines | CNN Business](https://edition.cnn.com/2025/01/16/media/apple-ai-news-fake-headlines/index.html)

TOOLS

TheRundownAI, Hugging Face, GitHub, etc

OpenAI launches ChatGPT 'Tasks'

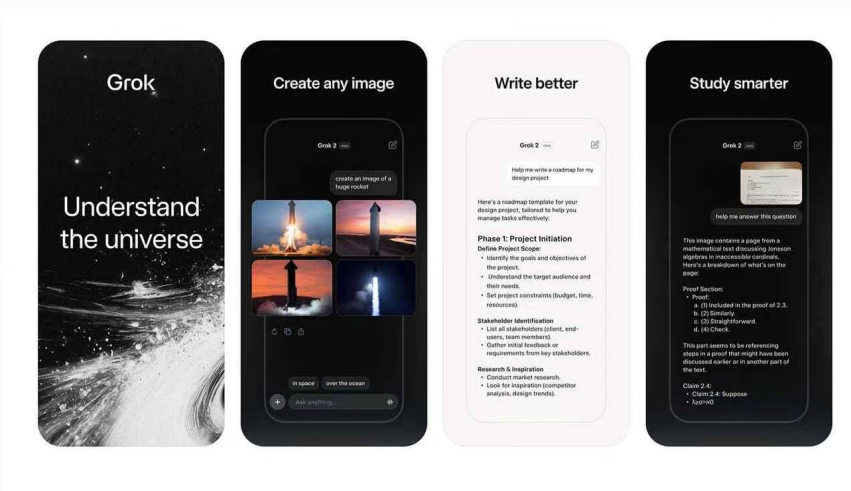


OpenAI is rolling out Tasks, a new ChatGPT beta feature that allows users to schedule reminders and recurring actions, marking the company's first step into agentic AI capabilities.

- Users can schedule one-time reminders or recurring actions, such as daily weather updates, news briefings, or periodic web searches.
- Tasks can be managed through chat or a dedicated web interface, with notifications available across desktop, mobile, and web platforms.
- ChatGPT can suggest relevant tasks based on conversation history, though users must explicitly approve any suggestions.
- Tasks will launch for Plus, Team, and Pro users in the coming days, with up to 10 active tasks at a time through a "4o with scheduled tasks" model option.

Post: [Rowan Cheung on X](https://x.com/rowancheung/status/1879233443772481742)

xAI breaks Grok free from X with standalone app



xAI just launched a new beta standalone app for its Grok AI assistant, marking the first shift away from its X (Twitter) integration and positioning the chatbot as a more direct competitor to rivals like ChatGPT and Gemini.

- The new iOS app gives users access to Grok 2, xAI's latest AI model, without requiring an X account or subscription.
- Users can access the app through various login options including Apple, Google, X accounts, or email, with both free and premium tiers available.
- The app includes features like image generation, text summarization, and real-time information access through web and X data.
- In addition, Grok appears to have improved its search feature, now gaining the ability to reference older posts from any user across X.

Post: [xAI on X](#)

Microsoft expands Copilot access with free tier

		Microsoft 365 Copilot Chat	Microsoft 365 Copilot
		Free + Consumption	\$30/user/mo
Chat	Copilot Chat – Web grounded (powered by GPT-4o)	●	●
	Copilot Chat – Work grounded (work data in your tenant's Microsoft Graph and 3rd party data via Graph connectors)	●	●
	Copilot Pages	●	●
	File upload ¹	●	●
	Code Interpreter ¹	●	●
Agents²	Image generation ¹	●	●
	Create agents using Copilot Studio ³ , including SharePoint agents	●	●
	Discover and pin agents	●	●
	Use agents grounded in Web data	●	●
	Use agents grounded in work data (work data in your tenant's Microsoft Graph and 3rd party data via Graph connectors)	▲	●
Personal assistant	Use agents that act independently using autonomous actions	▲	▲
	Copilot reasons over personal work data (e.g., Outlook, OneDrive, Teams meeting transcripts and chats)		
	Copilot in Teams		●
	Copilot in Outlook		●
	Copilot in Word		●
	Copilot in Excel		●
	Copilot in PowerPoint		●
	Copilot Actions		●
	Pre-built M365 agents (Interpreter, Facilitator, Project Manager, Employee Self-Service)		In preview
	Copilot Control System	Enterprise Data Protection (EDP)	●
IT management controls		●	●
Agent management		●	●
SharePoint Advanced Management			●
Copilot Analytics to measure usage and adoption ⁴			●
Pre-built reports and advanced analytics to measure ROI			●

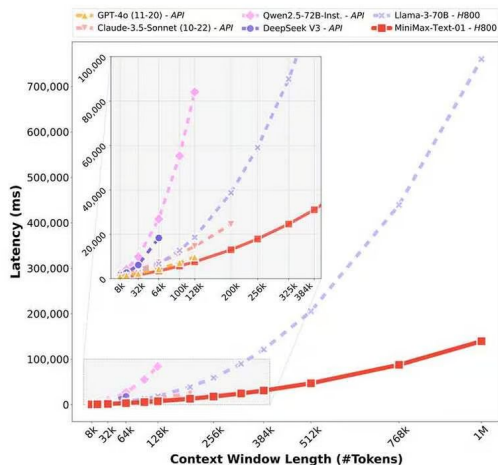
1. Limits apply. 2. Applies to employee-facing agents only. 3. Learn more about the full capabilities of Copilot Studio: aka.ms/CopilotStudioCapabilities. 4. Basic reporting in Microsoft Admin Center available for Copilot Chat.

Microsoft just introduced Microsoft 365 Copilot Chat, a rebranded free entry-level version of its AI assistant platform that introduces pay-per-use agent capabilities while maintaining core AI features for business users.

- The new tier offers free access to GPT-4o-powered chat, which includes web-based knowledge, file analysis capabilities, and image and code generation.
- Users can access custom AI agents for task automation, with a consumption-based model at \$0.01 per message or \$200 for 25,000 messages monthly.
- Agents can leverage knowledge sources for a range of tasks and actions, and the Copilot Control System allows IT teams to manage the platform easily.
- The offering aims to bridge the gap between free users and the full Microsoft 365 Copilot subscription (\$30/user/month).

Article: [Copilot for all: Introducing Microsoft 365 Copilot Chat | Microsoft 365 Blog](https://www.microsoft.com/en-us/microsoft-365/blog/2025/01/15/copilot-for-all-introducing-microsoft-365-copilot-chat/)

MiniMax releases open-source, ultra-long-context LLMs



Chinese AI lab Minimax just launched two new open-source AI models that leverage a new 'Lightning Attention' approach that allows for massive context windows of 4M tokens while maintaining speed and performance.

- The release includes a 456B parameter base language model (MiniMax-Text-01) and a multimodal model (MiniMax-VL-01).
- Both models can process sequences up to 4M tokens, dramatically exceeding current industry standards of 128K-256K tokens.
- The models perform comparable to top models on academic benchmarks, outperforming all open-source models on long-context tasks.
- The company also offers API access at notably low rates, with input tokens at \$0.2/million and output tokens at \$1.1/million.

Article: [MiniMax-01 series - MiniMax](#)

Luma Labs drops new next-gen AI video model

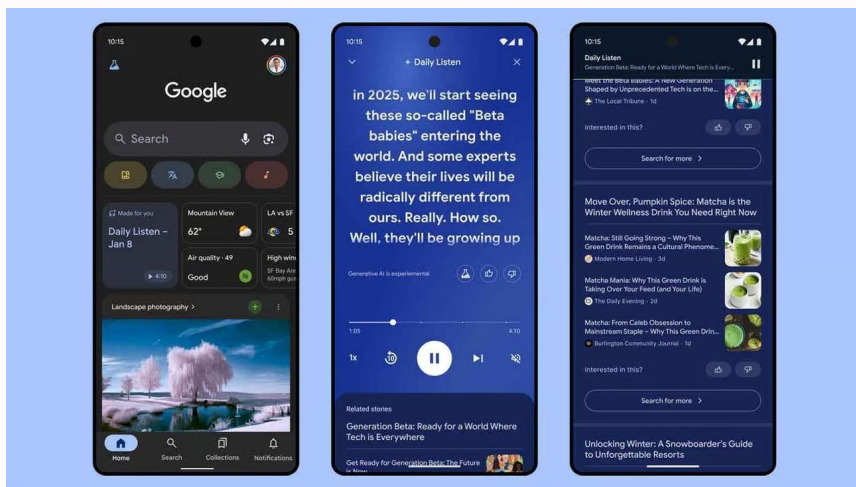


Luma Labs just released Ray 2, the startup's next-generation AI video model — which promises unprecedented motion quality and physics realism through a new multimodal architecture trained with 10x more computing power than its predecessor.

- The model can generate high-quality video clips up to 10 seconds long from text prompts, and it has advanced motion and physics capabilities.
- Ray2 demonstrates a sophisticated understanding of object interactions, from natural scenes like water physics to complex human movements.
- Ray2 can currently handle text, image, and video-to-video generations, and Luma will soon add editing capabilities to the model.
- The system is launching first in Luma's Dream Machine platform for paid subscribers, with API access coming soon.

Post: [Luma AI on X](#)

Google tests AI-powered 'Daily Listen' podcasts



Google just rolled out 'Daily Listen', a new experimental AI feature in Search Labs that transforms users' search interests and browsing data into personalized five-minute podcasts.

- The feature generates 5-minute AI-voiced podcasts based on users' Google Search history and Discover feed preferences.
- Daily Listen appears in the Google mobile app's homepage, featuring real-time transcripts and related story links for deeper exploration.
- The experiment is currently limited to U.S. users who opt into Search Labs, with content currently only available in English.
- The feature is a similar format to Google's NotebookLM Audio Overviews, focusing on news and updates rather than document summaries.

Article: [Search Labs](https://labs.google.com/search)

\$450 open-source reasoning model matches o1

	Sky-T1-32B-Preview	Qwen-2.5-32B-Instruct	QwQ	o1-preview
Math500	82.4	76.2	85.4	81.4
AIME2024	43.3	16.7	50.0	40.0
LiveCodeBench-Easy	86.3	84.6	90.7	92.9
LiveCodeBench-Medium	56.8	40.8	56.3	54.9
LiveCodeBench-Hard	17.9	9.8	17.1	16.3
GPQA-Diamond	56.8	45.5	52.5	75.2

UC Berkeley's NovaSky team just released Sky-T1-32B-Preview, a fully open-source reasoning model that competes with earlier versions of OpenAI's o1 despite minuscule training time and costs.

- Sky-T1 is a fine-tuned version of Alibaba's Qwen2.5-32-Instruct, with training data generated using the open-source reasoning model QwQ-32B-Preview.
- Training took just 19 hours on 8 H100 GPUs, and the total cost was around \$450 — a fraction of typical AI training budgets.
- The model matches or exceeds an earlier version of OpenAI's o1 on several benchmarks, particularly excelling in mathematics and coding challenges.
- Unlike other reasoning models, Sky-T1's entire pipeline, including training data, code, and model weights, is completely open source.

Article: [Sky-T1: Train your own O1 preview model within \\$450](https://novasky-ai.github.io/posts/sky-t1/)

RESEARCH

TheRundownAI, arXiv, Hugging Face, etc

Patients control AI and robotics with thought



NeuroXess just reported two advances in brain-computer interface (BCI) tech, enabling patients to control devices and communicate through thought, in addition to the first real-time decoding of Chinese language from brain signals.

- An epilepsy patient achieved 71% accuracy in converting thoughts to Chinese using 142 common syllables, with response times under 100 milliseconds.
- The system's flexible interface allowed the patient to control smartphones, smart home devices, and robotic arms within days of implantation.
- Patients operated digital avatars and interacted with AI models through thought alone, in what the company calls the first "mind-to-AI large model".

Article: [Brain-computer interface developed in China decodes thought real-time](https://interestingengineering.com/innovation/bci-china-decode-thought-real-time)

Study: AI phishing achieves alarming success rates

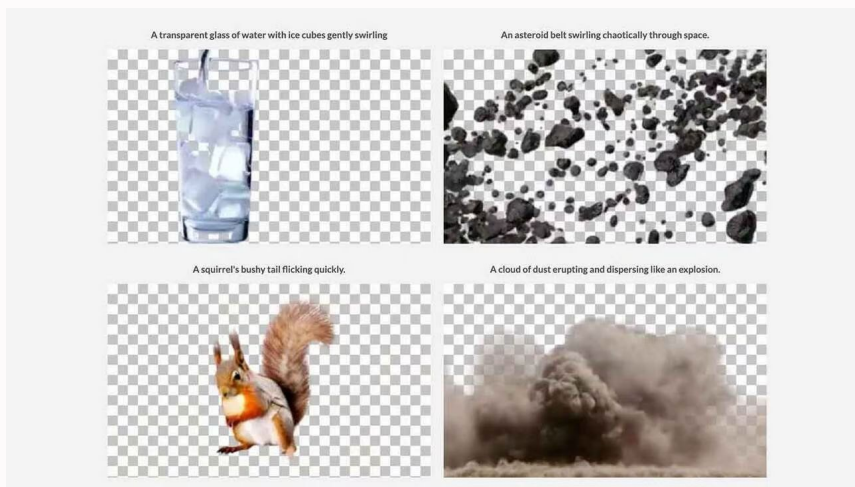


A new study from Harvard just found that AI systems can conduct fully automated phishing campaigns as effectively as human experts, with the systems achieving success rates over 50%.

- Researchers tested four campaigns: a standard phishing attempt, human experts, fully AI-automated, and an AI with human oversight.
- The AI-generated phishing emails achieved a 54% click-through rate, matching human attackers and far surpassing traditional spam's 12% success rate.
- The AI system fully automated both reconnaissance and email creation, accurately profiling 88% of targets using public web data.
- AI campaigns reduced costs by up to 50x over manual attacks, with Claude 3.5 Sonnet, GPT-4o and o1 all crafting content despite safety guardrails.

Article & Paper: [Human study on AI spear phishing campaigns – LessWrong](https://www.lesswrong.com/posts/GCHyDKfPXa5qsG2cP/human-study-on-ai-spear-phishing-campaigns)

Adobe showcases TransPixar for AI visual effects

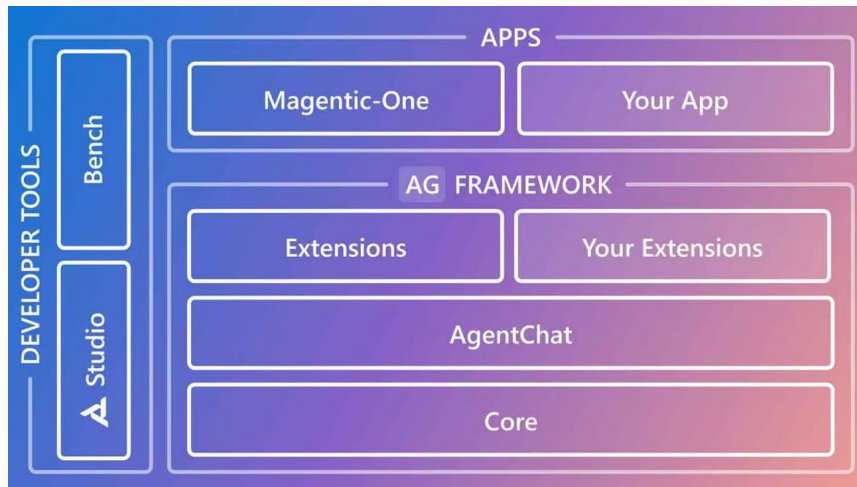


Adobe Research and Hong Kong University of Science and Technology just developed TransPixar, an AI system that enables generation of transparent, realistic visual effects to AI-generated videos.

- The tech enables the generation of see-through elements like smoke, reflections, and portals that can naturally blend into video scenes.
- The system teaches the AI to understand both visible content and transparency simultaneously, similar to layering in photo editing software.
- TransPixar also needs only minimal additional training data, showing the ability to create diverse effects without needing millions of example videos.
- The model excels across a range of effects like swirling storms, magical portals, and shattering glass, with applications ranging from movies to gaming.

Paper: [2501.03006](https://arxiv.org/pdf/2501.03006)

Microsoft upgrades AutoGen with new multi-agent system



Microsoft just released AutoGen v0.4, the latest update to its open-source agentic framework — alongside the introduction of Magnetic-One, a new orchestration system that coordinates multiple AI agent specialists to handle complex tasks.

- Magnetic-One introduces four agents: WebSurfer for web navigation, FileSurfer for local file management, and Coder and ComputerTerminal for coding.
- V0.4's event-driven messaging system enables async communication between agents, allowing for more flexible, customizable, and complex workflows.
- The release also includes AutoGen Studio for low-code development, AutoGen Bench for performance testing, and upgraded monitoring tools.
- The system also remains LLM-agnostic, working with different language models while defaulting to GPT-4o integration.

Article: [AutoGen v0.4: Reimagining the foundation of agentic AI for scale, extensibility, and robustness - Microsoft Research](https://www.microsoft.com/en-us/research/blog/autogen-v0-4-reimagining-the-foundation-of-agentic-ai-for-scale-extensibility-and-robustness/)

DF Labs