

DF Labs

AI NEWS UPDATES

September 1st, 2024





TABLE OF CONTENTS

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

NEWS

TheRundownAI, X, etc

Creating AI using human brain cells



Swiss startup **FinalSpark** just launched a service allowing scientists to **rent cloud access** to "**biocomputers**" made of **human brain cells** for **\$500 a month**, in an effort to create **AI** that uses **100,000x less energy** than current systems.

- The system uses **organoids (clumps of human brain cells)** that can "**live**" and **compute** for **up to 100 days**.
- AI models are **trained using dopamine for positive reinforcement** and **electrical signals for negative reinforcement**, mimicking natural neural processes.
- FinalSpark claims these biocomputers could be up to **100,000 times more efficient** for AI training than traditional silicon-based technology.
- The organoids and their behaviour are **live streamed 24/7**.

Live Stream: [Live - FinalSpark](#)

Article: [These 'living computers' are made from human neurons — and you can rent one for \\$500 a month | Live Science](#)

ChatGPT runs for mayor in Wyoming



Victor Miller, a **mayoral candidate** in **Wyoming's capital city**, just vowed to let his **customized ChatGPT GPT** named **Vic (Virtual Integrated Citizen)** **help run the local government** if elected.

- Miller **created VIC using ChatGPT, feeding it city ordinances and related documents to make municipal decisions.**
- Miller filed for him and VIC to run for mayor, **proposing that the ChatGPT GPT provides data-driven insights and solutions** while Miller ensures legal execution.
- **OpenAI has shut down Miller's account twice**, citing policies against using its products for campaigning.
- Wyoming's Secretary of State **raised concerns**, but local officials allowed Miller's candidacy with his human name on the ballot.

Article: [Wyoming voters face mayoral candidate who vows to let AI bot run government | US politics | The Guardian](https://www.theguardian.com/us-news/article/2024/aug/19/ai-mayor-candidate-victor-miller-cheyenne-wyoming)

ChatGPT teams up with ASU



OpenAI's ChatGPT is headed to Arizona State University (ASU), where **the university is integrating the AI assistant into over 200 projects across teaching, research, and operations.**

- ASU is using **ChatGPT Edu**, a version **designed for universities** with **enhanced privacy** and **security features.**
- **The university** also **launched** an **'AI Innovation Challenge'** for **faculty and staff**, receiving an overwhelming demand for using ChatGPT to maximize teaching, research, and ops.
- Key projects include an **AI writing companion for scholarly work**, 'Sam' (a chatbot **for med students to practice patient interactions**), and **AI-assisted research recruitment.**
- The partnership has inspired other institutions like Oxford and Wharton to pursue similar collaborations.

Article: [Personalizing education with ChatGPT | OpenAI](https://openai.com/index/asu/)

OpenAI prepares 'Project Strawberry'




OpenAI researchers are preparing to launch a **new AI model, code-named Strawberry (previously Q*)**, that demonstrates superior reasoning capabilities in solving complex problems, according to a new report via The Information.

- Project Strawberry **could be integrated into ChatGPT as soon as this fall**, marking a significant leap in AI intelligence.
- Given **extra "thinking" time**, Strawberry can **tackle subjective topics** and **solve complex puzzles** like the New York Times Connections.
- OpenAI is using Strawberry to **generate high-quality training data** for another secretive upcoming LLM, reportedly code-named **Orion**.
- The new AI model **could enhance** OpenAI's **development of AI agents, potentially automating multi-step tasks more effectively**.

Article: [OpenAI Races to Launch 'Strawberry' Reasoning AI to Boost Chatbot Business – The Information](https://www.theinformation.com/articles/openai-races-to-launch-strawberry-reasoning-ai-to-boost-chatbot-business)

OpenAI raises at \$100B valuation



OpenAI is reportedly in talks to **raise a new funding round** at a valuation exceeding **\$100 billion**, led by Thrive Capital, with Microsoft also expected to participate.

- The potential valuation of over \$100 billion would be **significantly higher than** OpenAI's **previous \$86 billion** valuation.
- Thrive Capital is expected to invest around \$1 billion in this round.
- **OpenAI's annualized revenue** reportedly **surpassed \$3.4 billion earlier this year**.
- The company is still, however, **projected to lose nearly \$5 billion** by **the end of the year** and has already **spent \$8.5 billion on AI training and staffing**.

Article: [Exclusive | OpenAI in Talks for Funding Round Valuing It Above \\$100 Billion - WSJ](https://www.wsj.com/tech/ai/openai-in-talks-for-funding-round-valuing-it-above-100-billion-4f0550c5)

California compromises on AI safety bill



California's **SB 1047**, an aggressive AI safety bill aimed at **preventing AI disasters**, just got significantly **revised to address concerns raised by AI companies** like Anthropic and open-source developers.

- The bill **no longer allows California's attorney general to sue AI companies for negligent safety practices before a catastrophic event occurs.**
- **AI labs are now only required to submit public "statements" about their safety practices vs certifications "under penalty of perjury."**
- Likewise, **developers must now provide "reasonable care" vs "reasonable assurance" that AI models do not pose significant risks.**
- The bill is **headed to California's Assembly floor for a final vote.**

Article: [California weakens bill to prevent AI disasters before final vote, taking advice from Anthropic | TechCrunch](https://techcrunch.com/2024/08/15/california-weakens-bill-to-prevent-ai-disasters-before-final-vote-taking-advice-from-anthropic/)

SB 1047: [Bill Text - SB-1047 Safe and Secure Innovation for Frontier Artificial Intelligence Models Act. \(ca.gov\)](#)

Claude sued for copyright infringement



A group of authors filed a **lawsuit against** AI startup **Anthropic**, alleging the company committed **"large-scale theft" by training** its Claude chatbot **on pirated copies of copyrighted books**.

- This is the **first lawsuit** from writers **targeting Anthropic and Claude**, but similar lawsuits have been filed against competitor OpenAI and ChatGPT.
- The lawsuit **accuses Anthropic** of **using** a dataset called **The Pile**, which includes numerous **pirated books**.
- **Anthropic and others**, including OpenAI, have **argued** that **training** AI models is **protected under the "fair use"** doctrine of U.S. laws, which permits the limited use of copyrighted materials.

Article: [Authors sue Claude AI chatbot creator Anthropic for copyright infringement - ABC News \(go.com\)](https://abcnews.go.com/US/wireStory/authors-sue-claude-ai-chatbot-creator-anthropic-copyright-112964872)

TOOLS

TheRundownAI, Hugging Face, GitHub, etc

OpenAI adds free fine-tuning to GPT-4o



OpenAI just launched free fine-tuning (up to 1 million tokens per day through September 23) for GPT-4o, allowing developers to customize the model for higher performance and accuracy.

- Developers can now, for the first time ever, fine-tune GPT-4o to improve the model's structure, tone, and domain-specific instructions for their AI applications.
- **Fine-tuning is available on all paid usage tiers with training costs of \$25 per million tokens, but it is completely free until September 23.**
- OpenAI suggests that developers should see strong results from fine-tuning with only a few dozen training examples.

Article: [Fine-tuning now available for GPT-4o | OpenAI](https://openai.com/index/gpt-4o-fine-tuning/)

Ideogram 2.0 launches with major upgrades



Ideogram just **released version 2.0** of its advanced text-to-image model with major upgrades and new features, including five new image styles, an iOS app, a beta API, and over 1 billion public Ideogram images.

- Ideogram 2.0 offers **five image styles: General, Realistic, Design, 3D, and Anime.**
- The Realistic style convincingly resembles photographs with dramatically improved textures for human features like hands and hair, a pain point for previous image generation models.
- The Design style also significantly **improves text rendering**, allowing users to create greeting cards, t-shirt designs and more.
- Ideogram offers a **free tier** that allows users to generate around **40 images, or 10 prompts a day** at no charge.

Post: [Ideogram on X](#)

Article: [Ideogram 2.0](#)

Luma Labs launches Dream Machine 1.5



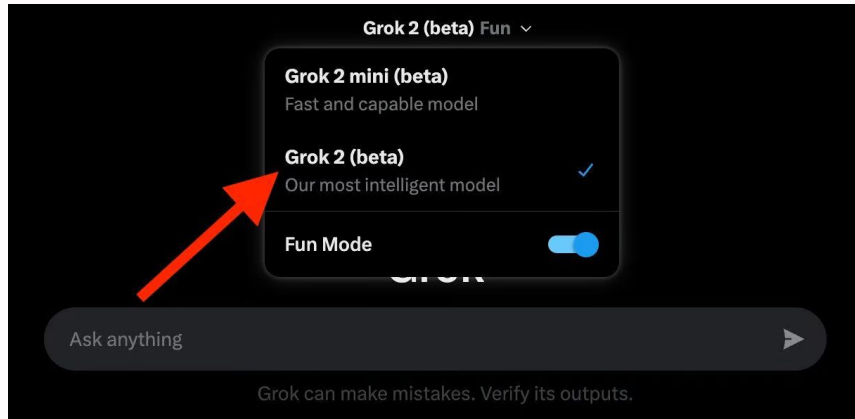
Luma Labs just released **Dream Machine 1.5**, a major upgrade to their current AI video generation model, with higher quality text-to-video, **smarter prompt understanding**, and better image-to-video capabilities.

- Dream Machine 1.5 builds on the original model's ability to generate high-quality, realistic 5-second video clips from text and image prompts.
- The upgraded model showcases **better natural language processing**, interpreting and executing prompts at a higher accuracy.
- It excels in creating smooth motion, cinematography, and dramatic shots, turning static concepts into dynamic stories, but lags in morphing, movement, and text.

Article: [Luma Dream Machine \(lumalabs.ai\)](https://lumalabs.ai/dream-machine)

FAQ & Prompt Guide: [FAQ and Prompt Guide: Luma Dream Machine \(notion.site\)](https://notion.site)

xAI releases Grok 2



xAI has begun **rolling out** early **beta access** for **Grok 2**, a powerful new AI model that **leverages real-time data from X** and uses **Flux.1** to generate **relatively unfiltered AI images**.

- Grok 2 is now **available to a select group** of **premium X users** in early beta mode.
- The model can access and **use real-time information from X**, setting it apart from ChatGPT and other LLMs.
- Grok 2 offers **two modes: regular** and **“fun”** mode, with the latter providing a more distinctive and entertaining personality.
- When gathering and summarizing news, Grok 2 can reference specific tweets, a capability that cannot be found in ChatGPT or Claude.

Post: [Rowan Cheung on X](https://x.com/rowancheung/status/1826285146305179800)

Nvidia and Mistral make laptop-ready AI



Nvidia and Mistral just released **Mistral-NeMo-Minitron 8B**, a highly accurate small language model that **can run efficiently on laptops and PCs**.

- The model uses optimization techniques like **pruning** (removing certain weights) and **distillation** (retraining the pruned model on a small dataset) to achieve high accuracy with a smaller footprint.
- These optimizations resulted in up to **40x cost savings** in terms of raw compute during training.
- **Laptops and PCs can run the model locally** for **faster** and **more secure** interactions with AI.
- Minitron 8B leads nine language-driven AI benchmarks for similarly sized models from language understanding to reasoning and coding.

Article: [NVIDIA Releases Small Language Model With State-of-the-Art Accuracy | NVIDIA Blog](https://blogs.nvidia.com/blog/mistral-nemo-minitrn-8b-small-language-model/)

RESEARCH

TheRundownAI, arXiv, Hugging Face, etc

AI creates DOOM video game in real-time

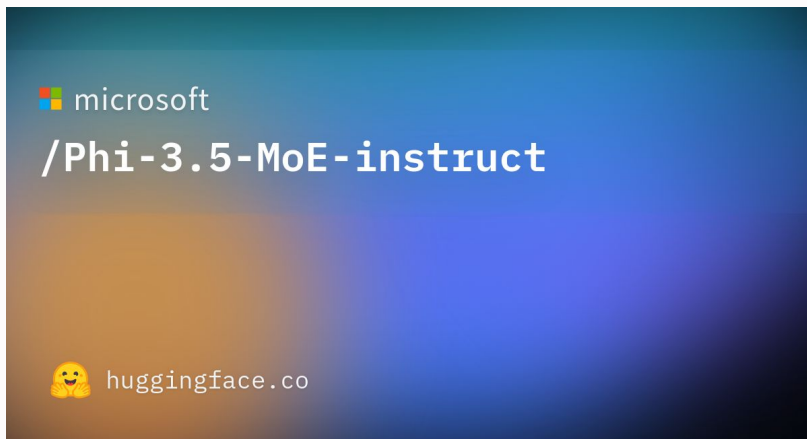


Google researchers just developed **GameNGen**, an AI system that can **simulate** the **classic game DOOM in real-time**, running at over **20 frames per second** and producing **visuals nearly indistinguishable** from the original game.

- GameNGen produces **playable gameplay at 20 frames per second on a single chip**, with **each frame predicted by a diffusion model**.
- The AI was **trained on 900M frames of gameplay data**, resulting in **3-second clips** almost indistinguishable from the actual game by playtesters.
- **Running on a single TPU**, GameNGen handles Doom's 3D environments and fast-paced action without traditional game engine components.
- In tests, human raters could barely distinguish between short clips of the AI simulation and the actual game.

Article: [GameNGen](#)

Microsoft released Phi-3.5-MoE

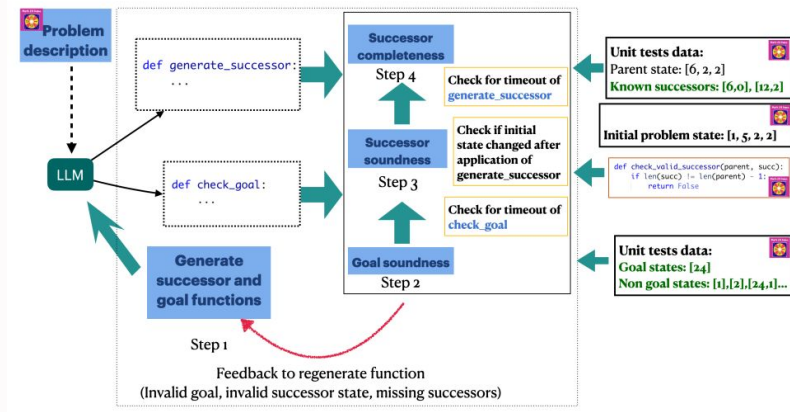


Microsoft just released **Phi-3.5-MoE**, an advanced AI model that rivals the reasoning capabilities of much larger models while maintaining a compact and efficient architecture.

- Phi-3.5-MoE uses a new **mixture-of-experts (MoE) approach**, which selectively activates only the most relevant parts of the model for each task to save compute power.
- The new model excels at understanding and following complex instructions and **can handle up to ~125,000 words in a single prompt**.
- In head-to-head benchmarks, Phi-3.5-MoE **outperformed** popular models like **Meta's Llama 3 8B and Google's Gemma 2 9B**, but fell short against OpenAI's GPT-4o mini.
- Microsoft made the model **available under an open-source MIT license** on **Hugging Face**.

Model: [microsoft/Phi-3.5-MoE-instruct](https://huggingface.co/microsoft/Phi-3.5-MoE-instruct) · [Hugging Face](#)

AI learns to plan better without humans



IBM Research and Cornell University recently created **AutoToS**, a system that teaches AI to **solve complex planning problems at 100% accuracy** – **without needing a human to check its work**.

- AutoToS is like a smart tutor for AI, **helping it learn how to break down and solve tricky problems step-by-step**.
- **The system** uses clever tests to **check the AI's work, pointing out mistakes and showing examples of how to do better** without human interference.
- This approach seems to work equally as well for smaller and larger models.
- AutoToS **succeeded in** teaching AI to **solve complex puzzles**, including classic problems **like arranging blocks and solving Sokoban**, a box-pushing game.

Paper: [2408.11326 \(arxiv.org\)](https://arxiv.org/abs/2408.11326)

DF Labs