Fresh Only · Daily Technology Report

Today's New Trends

November 13, 2026 (Friday)

A Few Sips

Content included in the report will not be repeated for 30 days. At the end of this period, records that remain on the agenda will be listed again with the "back on the agenda" label.

---

15 highlights · ~4 min read

5 GitHub Trending **5** Hacker News **5** Hugging Face · Daily Papers

Data captured November 13, 2026 at 8:16 AM.

Report source counts and item metadata are fixed at capture time. Discovery pages may show newer live source values.

The summaries and translations in this report are prepared with AI. Before making important decisions, please check the source through the link provided with each item.

GitHub Trending

5 highlights

01 pookjw/HermesEnabler

pookjw/HermesEnabler is software that makes Hermès watch faces available on all Apple Watch devices with supported watchOS versions. This tool, which can be compiled with Xcode on macOS or installed via application loading tools like Sideloadly, requires an Apple Watch paired with an iPhone.

Objective-C · ★ 1,613 · +218 today https://github.com/pookjw/HermesEnabler Report an error

---

02 QwenLM/qwen-code

QwenLM/qwen-code is an open-source coding AI agent that runs on terminals, editors, desktops, and browsers. Developed with the TypeScript language, this tool aims to support software development processes across different platforms.

TypeScript · ★ 29,931 · +151 today https://github.com/QwenLM/qwen-code Report an error

---

03 ashhart/TensorFold

ashhart/TensorFold offers an inference engine with Metal, CUDA, and Vulkan support for large language models (LLMs). Developed with the Zig language, this tool allows running models directly on Apple Silicon and NVIDIA graphics processors without the need for Python or MLX installation.

Zig · ★ 1,996 · +54 today https://github.com/ashhart/TensorFold Report an error

---

04 Alishahryar1/free-claude-code back in the spotlight

Alishahryar1/free-claude-code offers the opportunity to use various AI tools such as Claude Code and OpenClaw for free via terminal, application, IDE, or browser. The project provides access to over 1.3 billion free tokens with multi-model and multi-tool support.

Python · ★ 69,470 · +415 today https://github.com/Alishahryar1/free-claude-code Report an error

---

05 ggml-org/llama.cpp back in the spotlight

llama.cpp is a project that allows for the inference of large language models (LLMs) using C/C++ programming languages. It offers various installation options for users to run these models on their own computers.

C++ · ★ 134,660 · +176 today https://github.com/ggml-org/llama.cpp Report an error

Hacker News

5 highlights

01 Jensen Huang resigns

Nvidia CEO Jensen Huang has announced his resignation. Following this change, the company's management has gathered to evaluate new leadership plans and corporate processes.

3,181 points · 2,963 comments https://apnews.com/article/nvidia-megaspeed-china-singapore-huang-5403f152a601cf82ce939f1d24fffa0f Report an error

Discussion: news.ycombinator.com/item?id=50467087 (3,181 points · 2,963 comments)

02 Anthropic IPO

AI company Anthropic has initiated its IPO process and filed an official application to become a publicly traded company. The letter shared by the company's founders and other relevant documents make the details of the IPO process available to the public.

868 points · 830 comments https://ipo.anthropic.com/us/ Report an error

Discussion: news.ycombinator.com/item?id=50465829 (868 points · 830 comments)

03 Service degradation on Microsoft 365

Microsoft 365, the corporate productivity and office software suite, is experiencing service outages and performance issues worldwide. While users are having difficulty accessing the platform, the disruptions in services are being confirmed via the official status page.

654 points · 617 comments https://status.cloud.microsoft/ Report an error

Discussion: news.ycombinator.com/item?id=50471731 (654 points · 617 comments)

04 Rust 1.100

The system programming language Rust version 1.100 has been released with updates aimed at making it easier to develop reliable and efficient software. This version continues to support performance and security-focused standards in software development processes.

537 points · 218 comments https://blog.rust-lang.org/2026/11/12/Rust-1.100.0/ Report an error

Discussion: news.ycombinator.com/item?id=50472038 (537 points · 218 comments)

05 Gemini 4 Hydrogen

Gemini 4 Hydrogen, released by Google, is a fast and cost-efficient AI model designed for reliable agentic workflows and daily tasks. With its efficiency-focused structure, this model aims to support various application processes.

495 points · 315 comments https://blog.google/innovation-and-ai/models-and-research/gemini-models/introducing-gemini-4-hydrogen/ Report an error

Discussion: news.ycombinator.com/item?id=50471431 (495 points · 315 comments)

Hugging Face · Daily Papers

5 highlights

01 From Large Language Models to Efficient System One Models via Adaptive Computation

The article shared by Hugging Face examines the adaptive computation method for transforming large language models (LLMs) into efficient decision-making systems. The research reveals that information in the decision-making process is concentrated in the middle layers of the model, and the last six layers can be removed without loss of performance.

1 likes · 1 comments https://huggingface.co/papers/2611.16321 Report an error

02 Duo: A Reciprocal World Model for Full-Duplex Interaction

Duo v1, shared by Hugging Face, is a world model that continuously perceives auditory and visual data to interact with users in real-time. Using a Reciprocal World Model approach, this system makes reasoning and decision-making processes a fundamental part of the artificial intelligence model.

7 likes · 1 comments https://huggingface.co/papers/2611.15131 Report an error

03 G-STAR: Group-Scaled Token Capacity for Deeply Compressed Autoregressive Generation

G-STAR presents a method aimed at improving the balance between quality and speed in the image generation process for discrete autoregressive models. Based on the principle of product quantization, this approach increases representation capacity by scaling the group size with a spatial compression factor.

6 likes · 1 comments https://huggingface.co/papers/2611.13004 Report an error

04 BirdDreams: Joint Perspective and BEV Video Generation for a Plausible, Reactive Driving World Model

BirdDreams offers a method for driving world models that provides video generation in both perspective and bird's-eye view (BEV). This approach aims to simulate driving scenes realistically and responsively by combining vehicle movements and trajectory control.

20 likes · 1 comments https://huggingface.co/papers/2611.14815 Report an error

05 On the Design of Hyper-Connections: Simplification, Stability, and Scalability

The article published on Hugging Face examines the components of hyper-connection architectures used in advanced language models to explore simpler and more efficient design methods. The study reveals that learned gating mechanisms and mixing operations performed with fixed gates create effective interaction groups on model performance.

13 likes · 1 comments https://huggingface.co/papers/2611.05984 Report an error

📖 Glossary of Industry Terms

6 terms

watch faces

Digital interface designs that appear on the screens of smartwatches, displaying the time and other information.

---

open-source

A type of software whose source code is open for anyone to access, inspect, and improve.

LLM

Artificial intelligence models trained on massive amounts of data that can understand human language and generate text.

---

inference engine

A software component that enables a trained artificial intelligence model to make predictions or produce results on new data.

---

inference

The process of a trained AI model processing a new input provided to it to produce a result or response.

---

agentic workflows

A workflow where AI agents make their own decisions toward a goal and execute steps sequentially.

---

trescout.com · TreScout scans and summarizes, so all you have to do is read.