Fresh Only · Daily Technology Report

Today's New Trends

November 5, 2026 (Thursday)

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.

---

20 highlights · ~6 min read

10 GitHub Trending

4 Hacker News

1 Hugging Face · Daily Models

5 Hugging Face · Daily Papers

Data captured November 5, 2026 at 8:15 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

10 highlights

01 NVIDIA/auto-ontology

Auto-ontology, developed by NVIDIA, is a Python framework that enables the creation of semantic and graph-based reasoning layers on data. This tool offers more explainable results and deterministic query generation for large language models (LLM) and AI agent workflows.

Python · ★ 5,549 · +2,050 today https://github.com/NVIDIA/auto-ontology Report an error

02 github/spec-kit back in the spotlight

Developed by GitHub, spec-kit is an open-source toolkit that provides structured processes, reusable templates, and documented outputs to AI coding agents. Aiming to standardize software development processes, this kit allows for the customization of ready-made processes or the integration of users' own processes.

Python · ★ 147,442 · +808 today https://github.com/github/spec-kit Report an error

03 The-PR-Agent/pr-agent

The-PR-Agent/pr-agent is an open-source project designed to review pull requests in software development processes. Developed in Python, this tool aims to automate code review tasks.

Python · ★ 14,871 · +280 today https://github.com/The-PR-Agent/pr-agent Report an error

04 harry0703/MoneyPrinterTurbo back in the spotlight

MoneyPrinterTurbo is a Python project that creates high-resolution short videos from entered topics or keywords using AI models and automated workflows. The project utilizes Kimi K3.1 Pro, a large language model developed by Moonshot AI, to manage processes such as writing video scripts and selecting visual materials.

Python · ★ 139,201 · +617 today https://github.com/harry0703/MoneyPrinterTurbo Report an error

05 garrytan/gstack back in the spotlight

Shared by Garry Tan, gstack offers 23 different toolsets structured on Claude Code to manage software development processes. These tools aim to automate workflows by assuming the responsibilities of various roles such as CEO, designer, engineering manager, and quality assurance expert.

TypeScript · ★ 139,456 · +290 today https://github.com/garrytan/gstack Report an error

06 FareedKhan-dev/kimi-k3-in-c

The FareedKhan-dev/kimi-k3-in-c project enables the 2.78 trillion parameter Kimi K3 model to be run on a single processor with 8.24 GB of memory usage. This C99-based application, which does not require any framework or graphics processing unit (GPU), focuses on running high-parameter AI models on limited hardware.

C · ★ 10,450 · +351 today https://github.com/FareedKhan-dev/kimi-k3-in-c Report an error

07 JiaFeiMiao-K-Cat/DiscImageStudio

Disc Image Studio is a cross-platform disc image creation and burning tool that allows you to etch naked-eye visible images onto DVDs, CDs, and single-layer Blu-ray discs. Developed in C#, this software enables you to process custom graphics onto the disc surface using any writable optical drive.

C# · ★ 6,769 · +1,778 today https://github.com/JiaFeiMiao-K-Cat/DiscImageStudio Report an error

08 MichaelCade/90DaysOfDevOps

MichaelCade/90DaysOfDevOps offers a structured learning map on DevOps processes, principles, and tools. Covering three years of content, this resource serves as a comprehensive training guide for community members.

Shell · ★ 30,544 · +97 today https://github.com/MichaelCade/90DaysOfDevOps Report an error

---

09 Thysrael/Horizon

Thysrael/Horizon is a news tracking tool that collects, filters, and summarizes news from different sources using artificial intelligence. Preparing daily newsletters in English and Chinese, this software helps users manage their information flow by presenting content along with context and community discussions.

Python · ★ 12,645 · +121 today https://github.com/Thysrael/Horizon Report an error

---

10 scenario-labs/skills

scenario-labs/skills is a library that enables AI agents to select the right model and control costs for the production of visual, auditory, and 3D content. The tool uses the Scenario MCP protocol to maintain character consistency while integrating with professional design software such as Blender, Maya, ZBrush, Unreal, and Unity.

Python · ★ 2,566 · +224 today https://github.com/scenario-labs/skills Report an error

Hacker News

4 highlights

01 AI is forcing us to pay off 40 years of technical debt

The use of artificial intelligence in the field of cybersecurity is increasing the speed of both attackers and defenders, bringing technical debt in fragile legacy systems to light. This situation is triggering a process that forces institutions to renew forty-year-old software infrastructures.

1,129 points · 656 comments https://spectrum.ieee.org/ai-in-cybersecurity-php2shell-vulnerability Report an error

Discussion: news.ycombinator.com/item?id=50364336 (1,129 points · 656 comments)

---

02 Residential proxies detected on White House network

It is reported that residential proxies have been detected on the White House network. This situation opens up a discussion on the security implications of these networks, which route internet traffic through corporate networks and generally provide anonymity.

835 points · 265 comments https://layer3intel.com/blog/residential-proxies-white-house-network Report an error

Discussion: news.ycombinator.com/item?id=50371246 (835 points · 265 comments)

---

03 AMD expects to outsell Intel in Q4 for the first time in history

AMD projects that its processor sales will surpass Intel's for the first time in the fourth quarter. The company's announced revenue expectation of $17 billion exceeds the projections shared by its rival Intel for the same period.

513 points · 189 comments https://www.tomshardware.com/pc-components/cpus/amd-expects-to-outsell-intel-in-q4-for-the-first-time-in-history-usd17b-guidance-would-exceed-intels-projections Report an error

Discussion: news.ycombinator.com/item?id=50367125 (513 points · 189 comments)

---

04 Guerrilla bike lanes: San Francisco makes illicit infrastructure permanent (2016)

Illegal bike lanes created by anonymous activists with traffic cones in San Francisco have been turned into permanent infrastructure by the city administration. This intervention was carried out to raise awareness for safe transportation following the deaths of cyclists in traffic accidents.

281 points · 201 comments https://99percentinvisible.org/article/guerrilla-bike-lanes-san-francisco-makes-illicit-infrastructure-permanent/ Report an error

Discussion: news.ycombinator.com/item?id=50366522 (281 points · 201 comments)

Hugging Face · Daily Models

1 highlights

01 moonshotai/Kimi-K3.1-Pro

Shared on Hugging Face, moonshotai/Kimi-K3.1-Pro is a multimodal artificial intelligence model capable of processing visual and text inputs. Trained on top of the Kimi K3 model with large-scale reinforcement learning, this model is presented as the first and most capable version of the K3.1 series.

♥ 589 · 797 downloads https://huggingface.co/moonshotai/Kimi-K3.1-Pro Report an error

Discussion: news.ycombinator.com/item?id=50366999 (431 points · 128 comments)

Hugging Face · Daily Papers

5 highlights

01 VA-TDM: A Unified View for Improving Few-Step Diffusion Models by Velocity Adversarial Distillation

VA-TDM is a method developed to improve image quality in few-step diffusion models. By using velocity adversarial distillation, it aims to overcome distribution matching challenges in the sampling process.

1 likes · 1 comments https://huggingface.co/papers/2611.06796 Report an error

02 SleuthVLM: Do We Really Need Audio-in-Video for Omni Video Understanding?

SleuthVLM, shared by Hugging Face, questions whether multimodal AI models (omni models) actually need direct audio data for video understanding processes. Processing conversations by converting them to text, this system interprets video content through a Vision Language Model (VLM) that analyzes visual data independently.

1 likes · 1 comments https://huggingface.co/papers/2611.06083 Report an error

03 AnchorComplete: Recovering Complete Scene Meshes from Partial Point Clouds

AnchorComplete offers a framework that generates complete scene meshes from partial point clouds. Developed based on TRELLIS.2, this method allows for the completion of scenes at different scales while preserving the observed geometry.

1 likes · 1 comments https://huggingface.co/papers/2611.03228 Report an error

04 From Pixels to Actions - Towards Native Vision-Language-Action Policies

NEO-VLA, shared by Hugging Face, offers a method that allows visual data to be transferred directly to action policies without compression. This approach aims for more precise processing of visual details such as local geometry and depth, which are critical in robotic manipulation processes.

10 likes · 1 comments https://huggingface.co/papers/2611.06918 Report an error

05 StreamDuo: A Live Commentator and a Research Agent for Unseen Games

StreamDuo is an AI model that can provide live commentary for previously unseen games and is paired with a research agent that reads the game's online encyclopedia (wiki). This system was developed to address the lack of knowledge about game worlds in existing Large Language Models (LLMs).

8 likes · 2 comments https://huggingface.co/papers/2611.01120 Report an error

📖 Glossary of Industry Terms

11 terms

graph-based

A structure that establishes relationships between data through nodes and connections.

LLM

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

CI

The process that allows software changes to be automatically tested and merged into the main project.

---

open source

Software whose source code is open for anyone to inspect, modify, and enhance.

---

pull request

A request made by software developers to have their code changes reviewed and merged into the main project.

---

framework

A foundational structure that provides ready-made components and rules to simplify software development.

GPU

A powerful processing unit used in computers, especially for graphics rendering and artificial intelligence calculations.

---

technical debt

The phenomenon where temporary solutions implemented for quick fixes in software development require additional work in the future.

---

residential proxies

Servers that appear as real home internet connections and are used to mask internet traffic.

---

multimodal

Artificial intelligence capable of understanding and processing different types of data, such as text and images, simultaneously.

---

reinforcement learning

The process by which artificial intelligence learns to make correct decisions through trial and error using a reward mechanism.
