Daily Technology Report

Today's Technology Trends

November 5, 2026 (Thursday)

A Few Sips

While the software world focuses on cleaning up forty years of technical debt, developers like NVIDIA and Moonshot AI are expanding infrastructure capacity with next-generation artificial intelligence models. Security vulnerabilities in communication networks and new approaches in video understanding models are at the center of technical discussions today.

---

30 highlights · ~8 min read

16 GitHub Trending

4 Hacker News

5 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

16 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 max-sixty/worktrunk

Developed with the Rust language, worktrunk is a command-line interface (CLI) that simplifies Git worktree management. The tool is specifically designed to manage the simultaneous workflows of multiple AI agents.

Rust · ★ 11,155 · +531 today https://github.com/max-sixty/worktrunk Report an error

03 mattpocock/skills

Shared by Matt Pocock, mattpocock/skills is a collection containing AI agent skills used in real engineering processes. Prepared to automate software development tasks, this resource focuses on tools that can be applied directly to daily workflows rather than theoretical approaches.

Shell · ★ 313,618 · +1,379 today https://github.com/mattpocock/skills Report an error

04 floci-io/floci

Developed with Java, floci is an open-source local AWS emulator used for software development, testing processes, and continuous integration (CI) operations. It offers developers a free alternative by mimicking cloud services in a local environment.

Java · ★ 29,760 · +1,090 today https://github.com/floci-io/floci Report an error

05 thedotmack/claude-mem

thedotmack/claude-mem is a tool that provides persistent context across sessions for AI agents. It compresses session data with AI to make relevant information available for future operations and works compatibly with various development tools such as Claude Code.

TypeScript · ★ 113,028 · +1,172 today https://github.com/thedotmack/claude-mem Report an error

06 vercel-labs/agent-desktop

Developed by Vercel Labs, agent-desktop is a command-line interface (CLI) that allows AI agents to see and manage desktop applications on macOS, Windows, and Linux operating systems. Written in Rust, this tool uses the operating system's accessibility tree to enable agents to perform operations on the computer.

Rust · ★ 15,552 · +2,470 today https://github.com/vercel-labs/agent-desktop Report an error

07 github/spec-kit

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

08 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

---

09 harry0703/MoneyPrinterTurbo

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

---

10 garrytan/gstack

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

---

11 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

---

12 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

---

13 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

---

14 D4Vinci/Scrapling

Scrapling is a Python-based web scraping framework capable of managing all stages from simple data extraction to comprehensive crawling processes. It provides all the tools users need in a single structure, from a single request to large-scale data collection processes.

Python · ★ 96,866 · +2,005 today https://github.com/D4Vinci/Scrapling Report an error

---

15 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

---

16 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

5 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)

---

02 nvidia/NVIDIA-Nemotron-3.5-Transcribe-30B-A4B

Developed by NVIDIA, NVIDIA-Nemotron-3.5-Transcribe-30B-A4B is an open-weight automatic speech recognition model capable of transcribing in 40 languages. Supporting audio recordings of up to one hour, the model operates with a mixture-of-experts architecture.

♥ 3,001 · 44,939 downloads https://huggingface.co/nvidia/NVIDIA-Nemotron-3.5-Transcribe-30B-A4B Report an error

---

03 deepseek-ai/DeepSeek-V4.1-Flash-1017

Shared on Hugging Face, deepseek-ai/DeepSeek-V4.1-Flash-1017 is an image-to-text model capable of processing visual and text-based inputs. As an updated version of the DeepSeek-V4.1-Flash model, it aims to reduce processing time and costs while providing performance improvements in areas such as reasoning and coding.

♥ 5,292 · 221,621 downloads https://huggingface.co/deepseek-ai/DeepSeek-V4.1-Flash-1017 Report an error

---

04 speridlabs/idol

speridlabs/idol is an AI model that creates 3D Gaussian splat models of objects from a single photo. This tool, which cleans the background and outputs in the standard 3DGS .ply format, can run on consumer-grade graphics processors in under two seconds.

♥ 636 · 14,613 downloads https://huggingface.co/speridlabs/idol Report an error

---

05 LiquidAI/LFM3-1.7B

LFM3-1.7B, developed by LiquidAI, is a text-generation model designed for local agents and specialized applications. Built on a new tri-hybrid architecture, this model is optimized to run on systems with low resource requirements.

♥ 1,494 · 54,959 downloads https://huggingface.co/LiquidAI/LFM3-1.7B Report an error

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

16 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.

worktree

A Git feature that allows working on different versions of a software project simultaneously.

CLI

A text-based interface that runs programs using written commands instead of a mouse.

agent skills

Special capabilities that AI agents possess to perform real engineering tasks.

CI

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

persistent context

The ability of AI agents to remember previous conversations or actions and carry information across sessions.

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.

web scraping

The process of automatically collecting and saving data from websites.

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.

---

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