Daily Technology Report

Today's Technology Trends

November 9, 2026 (Monday)

A Few Sips

In the software world, the security of open-source tools and corporate usage restrictions are at the forefront today. In particular, new benchmarks developed for coding agents embody technical advancements in the field of 3D modeling.

---

24 highlights · ~7 min read

9 GitHub Trending

5 Hacker News

5 Hugging Face · Daily Models

5 Hugging Face · Daily Papers

Data captured November 9, 2026 at 8:24 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

9 highlights

01 sherlock-project/sherlock

Sherlock is a Python-based tool used to find accounts on different social media networks using usernames. It is accessible via community-prepared packages for various Linux distributions and package managers.

Python · ★ 104,510 · +2,373 today https://github.com/sherlock-project/sherlock Report an error

02 dsx-ai-factory/nvke

NVIDIA Kubernetes Engine (NVKE) is a management plane designed to manage GPU-supported Kubernetes clusters on physical servers. Developed in the Go language, this tool enables the configuration of hardware resources for artificial intelligence and similar processes by creating workload clusters.

Go · ★ 920 · +345 today https://github.com/dsx-ai-factory/nvke Report an error

03 milind-soni/OpenMausBot

OpenMausBot is an open-source Grok Bot alternative that allows you to build your own AI bot team within chat applications. The project aims to increase functionality by providing a virtual machine infrastructure that bots can use.

TypeScript · ★ 18,634 · +1,439 today https://github.com/milind-soni/OpenMausBot Report an error

04 anthropics/claude-quickstarts

The anthropics/claude-quickstarts offered by Anthropic contains projects that help developers create quickly deployable applications using the Claude API (a digital bridge that allows software to communicate with each other). Prepared in the Python language, this collection aims to simplify the AI-based application development process.

Python · ★ 19,560 · +32 today https://github.com/anthropics/claude-quickstarts Report an error

05 vercel-labs/agent-browser

Developed by Vercel, agent-browser is a command-line interface (CLI) that allows AI agents to perform operations on a browser. Written in Rust, this tool reads pages as an accessibility tree, allowing agents to interact with web elements.

Rust · ★ 47,835 · +815 today https://github.com/vercel-labs/agent-browser Report an error

06 meta/astryx

Developed by Meta, astryx is a fully customizable open-source design system optimized for both humans and AI agents. Prepared in TypeScript, this project aims to support human and AI collaboration in modern software development processes.

TypeScript · ★ 17,451 · +875 today https://github.com/meta/astryx Report an error

07 anthropics/skills

Shared by Anthropic, skills is a Python repository containing various application examples that showcase the capability system of the AI model Claude. This resource offers ready-to-use skill packages that can be used in a wide range of areas, from creative projects to technical tasks and corporate workflows.

Python · ★ 189,869 · +792 today https://github.com/anthropics/skills Report an error

---

08 google-deepmind/mujoco

Developed by Google DeepMind, MuJoCo is a general-purpose physics engine that provides fast and accurate simulation of articulated structures in fields such as robotics, biomechanics, and machine learning. Written in C++, this tool aims to facilitate the development and research of systems that interact with complex environments.

C++ · ★ 16,542 · +41 today https://github.com/google-deepmind/mujoco Report an error

---

09 ranxianglei/billion-context

Developed with TypeScript, billion-context is a compression plugin that enables efficient use of the context window, which is the information capacity that AI can process at once. It reduces token consumption by 5 times in long monthly sessions, allowing work with smaller capacity windows such as 100k tokens.

TypeScript · ★ 5,333 · +460 today https://github.com/ranxianglei/billion-context Report an error

Hacker News

5 highlights

01 Goodbye Claude, hello Pi

The article shared by Caleb Software describes the user's transition process from using the AI model Claude to using the personal AI assistant Pi. With this change, the differences in the user's interaction preferences and personal experiences are highlighted.

720 points · 446 comments https://caleb.software/posts/claude-to-pi.html Report an error

Discussion: news.ycombinator.com/item?id=50410107 (720 points · 446 comments)

---

02 Goldman Sachs is pulling Claude Code from its bankers and traders

Goldman Sachs is restricting bankers and traders from using the coding tool Claude Code, transferring this authority back to engineering teams. The fact that the company made this decision shortly before the public offering of the AI model Claude, developed by Anthropic, is noteworthy.

677 points · 382 comments https://www.thestateofai.com/news/goldman-sachs-claude-code-traders-bankers Report an error

Discussion: news.ycombinator.com/item?id=50415628 (677 points · 382 comments)

---

03 Claude Code's "Don't Ask Again" exposes passwords in 1,122 public GitHub files

The "Don't Ask Again" feature of the Claude Code tool causes the exposure of passwords and sensitive information in 1,122 public GitHub files. This situation highlights the security vulnerabilities of AI agents used in software development processes.

596 points · 406 comments https://apiiro.com/blog/coding-agents-leaking-secrets-github-dont-ask-again Report an error

Discussion: news.ycombinator.com/item?id=50407882 (596 points · 406 comments)

---

04 Nvidia CEO Jensen Huang's letter to staff about affair

Nvidia CEO Jensen Huang addressed allegations regarding his personal life in a letter sent to company employees. The letter discusses the impact of the situation on the company and the functioning of corporate processes.

456 points · 494 comments https://nvidianews.nvidia.com/news/a-message-to-our-company Report an error

Discussion: news.ycombinator.com/item?id=50418016 (456 points · 494 comments)

---

05 Declarative Dragging

Declarative Dragging offers a method that allows interface designers to define future model states to make an element draggable. This study, which won the best paper award at the UIST 2026 conference, addresses drag interactions with a declarative approach.

436 points · 55 comments https://declarative-dragging.github.io/ Report an error

Discussion: news.ycombinator.com/item?id=50368481 (436 points · 55 comments)

Hugging Face · Daily Models

5 highlights

01 speridlabs/idol

Speridlabs/idol is an AI tool that creates a 3D Gaussian splat model from a single photo. This system, which outputs a standard 3D file format (.ply), can run on consumer-grade graphics cards in under two seconds.

♥ 1,424 · 22,854 downloads https://huggingface.co/speridlabs/idol Report an error

02 xtools-at/DeepSeek-V4-Pro-Qwen3.5-9B-UNCENSORED-heretic

xtools-at/DeepSeek-V4-Pro-Qwen3.5-9B-UNCENSORED-heretic is an AI model with text-generation capabilities. Achieving high accuracy rates in GSM8K tests, the model aims to demonstrate performance similar to larger models with its 9B parameter structure.

♥ 901 · 7,088 downloads https://huggingface.co/xtools-at/DeepSeek-V4-Pro-Qwen3.5-9B-UNCENSORED-heretic Report an error

03 moonshotai/Kimi-K3.1-Pro

Kimi-K3.1-Pro, developed by Moonshot AI, is a multimodal and agent-based AI model capable of processing visual and text data. Trained on top of the Kimi K3 model using large-scale reinforcement learning methods, this open-weight model is presented as the first member of the K3.1 series.

♥ 1,733 · 25,569 downloads https://huggingface.co/moonshotai/Kimi-K3.1-Pro Report an error

04 ricursiveAI/Ripple-65B-Preview

Ripple-65B-Preview, developed by Ricursive Intelligence, is a dense Transformer model with 65 billion parameters trained on 12 trillion tokens. Supporting a context window of over one million tokens, the model demonstrates performance in agent-based coding tasks with the results it achieved on TerminalBench 2.1.

♥ 412 · 5,087 downloads https://huggingface.co/ricursiveAI/Ripple-65B-Preview Report an error

05 Lightricks/LTX-2.5

Lightricks/LTX-2.5, shared on Hugging Face, is an image-to-video AI model that converts static images into moving videos. Allowing users to turn their visual content into dynamic videos, this model aims to simplify video creation processes.

♥ 11,150 · 2,661,026 downloads https://huggingface.co/Lightricks/LTX-2.5 Report an error

Hugging Face · Daily Papers

5 highlights

01 Frontier3DBench: Benchmarking Coding Agents for 3D Creation across Objects, Scenes, and Interactive Open Worlds

Frontier3DBench is an evaluation tool developed to measure the capabilities of coding agents in object and scene generation processes. Bringing different task scopes under one roof, this benchmarking system tests the performance of AI models that convert visual data and instructions into editable programs across 145 different tasks.

2 likes · 0 comments https://huggingface.co/papers/2611.09017 Report an error

02 OmniLatent: Unified Latent for 3D Part Modeling

OmniLatent is a unified latent space modeling framework developed for 3D part decomposition operations. By encoding geometry, deformation, and appearance information in a common space, it aims to resolve issues such as geometric distortion and texture ambiguity encountered in editable asset creation processes.

1 likes · 0 comments https://huggingface.co/papers/2611.09055 Report an error

03 Video-Coder: Reinforcing Agentic Coding for Video Generation

Video-Coder is a code-based AI agent developed to provide precise spatial structure and motion dynamics in video creation processes. The system provides structural signals for video production by autonomously writing and executing the code required to create dynamic 3D scenes.

6 likes · 1 comments https://huggingface.co/papers/2611.09894 Report an error

04 EvoWorld: Towards Out-of-Sight World Evolution in Generative Video Models

EvoWorld aims for generative video models to go beyond visual consistency and logically infer changes in the world and the consequences of events. This study focuses on models not only generating realistic pixels but also reasoning and simulating the interactions of entities in the environment.

13 likes · 1 comments https://huggingface.co/papers/2611.06748 Report an error

05 AgentBaton: Efficient Model Collaboration in Agentic Tasks

AgentBaton, shared by Hugging Face, aims to optimize collaboration between models with different capabilities in complex and long-term AI tasks. This approach focuses on reducing costs while maintaining task performance by using high-cost models together with more accessible and efficient models.

3 likes · 1 comments https://huggingface.co/papers/2610.35959 Report an error

📖 Glossary of Industry Terms

16 terms

NVKE

A software platform developed by NVIDIA used to manage GPU-accelerated server clusters.

GPU

A specialized processing unit that accelerates complex operations in computers, particularly graphics and AI calculations.

management plane

A management layer used to centrally control and monitor all components of a system.

open source

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

CLI

An interface that allows you to manage programs without using a mouse, solely by typing text-based commands.

accessibility tree

A hierarchical list that organizes the structure and function of elements on web pages in a way that assistive technologies or AI can understand.

design system

A collection of standard components, rules, and design tools used to ensure an interface looks consistent.

context window

The total amount of information an AI model can remember and work on during a single operation.

text-generation

The process of an AI generating meaningful text in accordance with the instructions given to it.

multimodal

Systems that are not limited to text but can process different types of data such as images, audio, or video simultaneously.

reinforcement learning

A training method where AI learns to make correct decisions through trial and error and the rewards it receives.

open-weight

These are models where the internal settings of the model obtained after training are public, but the training data is not shared.

token

These are word fragments or character groups that artificial intelligence uses when processing texts.

Transformer

It is an architecture that forms the basis of today's advanced artificial intelligence models and learns by analyzing the relationships between data.

---

image-to-video

It is an artificial intelligence technology that takes a static photo and transforms it into a moving video image.

---

coding agents

These are artificial intelligence software designed to automate software development processes, write code, or fix errors.

---

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