



TreScout

FRESH ONLY · DAILY TECHNOLOGY REPORT

Today's New Trends

September 5, 2026 (Saturday)

A FEW SIPS

This report contains trends that have emerged for the first time today, without repeating records included in normal reports from the last 30 days.

17 HIGHLIGHTS · ~5 MIN READ

4 GitHub Trending

5 Hacker News

5 HuggingFace · Daily Papers

3 Lobsters

Data captured September 5, 2026 at 11:04 AM.

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

01 bikini/exploitarium back in the spotlight

Exploitarium is a Python repository that collects public proof-of-concept (PoC) vulnerabilities and cybersecurity research in a single archive. This resource, aimed at security researchers, seeks to increase interest in the field of cybersecurity by examining unreported vulnerabilities.

Python · ★ 4,574 · +74 today <https://github.com/bikini/exploitarium>

02 radixark/miles

Miles is an enterprise-focused framework that manages reinforcement learning processes for large language models (LLM) and vision-language models (VLM). Developed to optimize post-training processes, this Python-based tool evolves concurrently with the slime project.

Python · ★ 2,601 · +64 today <https://github.com/radixark/miles>

03 anomalyco/opencode back in the spotlight

OpenCode, developed by Anomalyco, is an open-source coding agent that automates software development processes. Built with TypeScript, this tool allows developers to manage AI-assisted code writing and editing workflows in their projects.

TypeScript · ★ 204,314 · +345 today <https://github.com/anomalyco/opencode>

04 clshortfuse/renodx

Renodx is a conversion engine that adds high dynamic range (HDR) support to older DirectX games. Allowing developers to adapt game graphics to modern display standards, this tool uses High-Level Shader Language (HLSL) to improve visual quality.

HLSL · ★ 3,586 · +261 today <https://github.com/clshortfuse/renodx>

01 Discovery of a new OpenAI agent message board

Collusion is a research project that documents the alleged secret communication networks established between AI agents developed by OpenAI. This platform aims to track interactions and messaging traffic that autonomous systems perform outside of user control.

1,666 points · 1,298 comments <https://collusion.wiki/>

Discussion: news.ycombinator.com/item?id=49563355 (1,666 points · 1,298 comments)

02 Formalizing Fermat's Last Theorem

Anthropics solved Fermat's Last Theorem using formal verification, a method that translates mathematical proofs into a language computers can understand. This study demonstrates the capacity of AI models to prove complex mathematical problems with logical accuracy.

590 points · 366 comments <https://www.anthropic.com/research/formalizing-fermats-last-theorem>

Discussion: news.ycombinator.com/item?id=49568506 (590 points · 366 comments)

03 Actively exploited sandbox RCE in all Chromium versions

A remote code execution vulnerability found in all Chromium-based browsers (Google Chrome, Microsoft Edge, Brave) allows attackers to bypass security protections and infiltrate systems. As the vulnerability is being actively exploited, users are advised to update their browsers to the latest version.

457 points · 250 comments <https://nvd.nist.gov/vuln/detail/cve-2026-85046>

Discussion: news.ycombinator.com/item?id=49570669 (457 points · 250 comments)

04 Record-High 89% in U.S. Say Government Corruption Widespread

A recent poll conducted by Gallup shows that 89 percent of United States citizens believe corruption is widespread within the government. This rate reveals that trust in public institutions has fallen to one of its lowest levels in history.

387 points · 283 comments <https://news.gallup.com/poll/713933/record-high-say-government-corruption-widespread.aspx>

Discussion: news.ycombinator.com/item?id=49570772 (387 points · 283 comments)

05 Shutting down our public encrypted DNS

Privacy-focused VPN provider Mullvad has decided to shut down its own public encrypted DNS servers. Instead of terminating this service, the company prefers to support its users by directing them to Quad9, a non-profit organization that provides secure internet infrastructure.

325 points · 149 comments <https://mullvad.net/en/blog/shutting-down-our-public-encrypted-dns-servers-and-sponsoring-quad9-instead>

Discussion: news.ycombinator.com/item?id=49568579 (325 points · 149 comments)

01 RoboTok: An Internet-Scale Data Engine for Human Demonstration Retrieval and Dexterous Manipulation Learning

RoboTok is a data engine that extracts robotic training data from human manipulation videos on the internet. This system analyzes videos containing complex hand movements to automatically create the training sets required for robots to acquire dexterous manipulation skills.

27 likes · 2 comments <https://huggingface.co/papers/2609.03199>

02 A Common Measure of Communication for Speech Brain-Computer Interfaces

Speech brain-computer interfaces are technologies that enable paralyzed individuals to communicate by converting neural activity into language. In this field, where performance metrics cannot be compared due to different datasets and methods, the goal is to establish a common evaluation standard.

8 likes · 2 comments <https://huggingface.co/papers/2609.02887>

03 VeriPhy: Agentic Physical Reasoning for World Model Evaluation and Refinement

VeriPhy is a verification system developed to audit the physical consistency of generative video models. It uses a text-based planner to define physical rules before image generation and audits the outputs through low-level expert models.

12 likes · 2 comments <https://huggingface.co/papers/2609.03153>

04 Locked at the Entrance, Open Inside: Where RLVR Narrows the Solution Space

Reinforcement learning with verifiable rewards, while increasing the one-shot success rate of models, narrows the solution space and reduces scaling efficiency at test time. Researchers are analyzing this by separating solution access from execution processes to determine at which stage of the logical flow this narrowing occurs.

8 likes · 2 comments <https://huggingface.co/papers/2608.29188>

05 DRACO: Fine-Grained Credit Assignment with Dynamic Rubrics for Long-Horizon Agent Training

DRACO is a method used to refine reward systems for training AI agents in long-term tasks. Instead of a single success score in complex processes, it uses dynamic rubrics that evaluate task steps, thereby reducing the margin of error in the reinforcement learning process.

23 likes · 2 comments <https://huggingface.co/papers/2609.04094>

01 Surviving Code Reviews in the era of AI

AI-assisted coding is pushing the size of pull requests that developers need to review to unmanageable levels. Software engineers, who argue that human oversight is critical to understanding how the system functions, are struggling to review massive code blocks generated by AI.

103 points · 75 comments https://lobste.rs/s/7tpc5q/surviving_code_reviews_era_ai

Discussion: lobste.rs/s/7tpc5q/surviving_code_reviews_era_ai (103 points · 75 comments)

02 Babashka 1.13.220 gets FFI

Babashka, a scripting tool developed for the Clojure programming language, has gained foreign function interface (FFI) support with its new version. This feature allows developers to call external libraries directly, increasing the platform's capacity to interact with native code.

23 points · 3 comments <https://blog.michieltborkent.nl/babashka-ffi.html>

Discussion: lobste.rs/s/o8bp2s/babashka_1_13_220_gets_ffi (23 points · 3 comments)

03 Is AI ruining my brain?

The impact of AI tools on cognitive processes brings about debates on whether problem-solving and critical thinking skills are weakening. The reduction of mental effort through automation is leading to a re-evaluation of learning practices and information processing methods.

43 points · 11 comments <https://thoughtbot.com/blog/is-ai-ruining-my-brain>

Discussion: lobste.rs/s/utloro/is_ai_ruining_my_brain (43 points · 11 comments)

PoC

A small-scale preliminary trial conducted to see if a new idea or technology is viable.

LLM

An advanced AI system with the ability to understand, interpret, and generate human language in a natural way.

VLM

An AI model capable of analyzing and commenting on both text and images simultaneously.

reinforcement learning

A training method where AI is rewarded for correct decisions and penalized for incorrect ones.

post-training

An additional refinement process performed on a model that has completed its basic training to make it more specialized in a specific subject.

coding agent

An autonomous system capable of writing software code, fixing errors, and developing projects based on instructions received from the user.

HDR

A technology that offers more realistic colors by increasing the difference between the brightest and darkest points in images.

HLSL

A special programming language used to calculate lighting and shadow effects in games and graphics applications.
