



TreScout

FRESH ONLY · DAILY TECHNOLOGY REPORT

Today's New Trends

August 13, 2026 (Thursday)

A FEW SIPS

This report includes the trends that stood out today for the first time, without repeating what we have shared in recent days.

23 HIGHLIGHTS · ~6 MIN READ

11 GitHub
Trending

4 Hacker
News

1 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

2 Lobsters

01 cathrynlavery/diagram-design

Diagram-design, developed by Cathryn Lavery, includes 29 different editorial diagram types prepared for the artificial intelligence-supported coding tool Claude Code. This library cleanly presents complex diagrams directly using HTML and scalable vector graphics (SVG) instead of standard diagramming tools.

HTML · ★ 11,456 · +2,855 today <https://github.com/cathrynlavery/diagram-design>

02 macro-inc/macro

Macro is a workspace that combines business tools such as email, chat, documents and task management with a common artificial intelligence memory. Developed with the Rust language, this platform connects different workflows in a single interface.

Rust · ★ 2,066 · +227 today <https://github.com/macro-inc/macro>

03 hugohe3/ppt-master back in the spotlight

Python-based ppt-master is an artificial intelligence tool that converts documents or topics into editable presentation files (PowerPoint). The software allows users to use their own templates, turning text into professional presentations with native shapes, transition effects, animations and data-driven graphics.

Python · ★ 45,916 · +476 today <https://github.com/hugohe3/ppt-master>

04 infiniflow/ragflow

RAGFlow is an open source fetch-based generation (RAG) engine that creates a context layer for large language models (LLM). It aims to improve data processing and accuracy of responses by combining advanced RAG techniques with agent capabilities.

Go · ★ 87,648 · +139 today <https://github.com/infiniflow/ragflow>

05 NVIDIA-NeMo/Switchyard

Developed by NVIDIA, Switchyard is a high-performance artificial intelligence inference engine written in Rust language. It offers an optimized runtime environment to run large language models (LLM) efficiently on different hardware infrastructures.

Rust · ★ 923 · +421 today <https://github.com/NVIDIA-NeMo/Switchyard>

06 ZuodaoTech/everyone-can-use-english

Everyone-can-use-english is an open source educational project that supports the English learning process with a structured curriculum. It offers practical methods and study resources for users who want to improve their language skills.

TypeScript · ★ 36,233 · +86 today <https://github.com/ZuodaoTech/everyone-can-use-english>

07 **smicallef/spiderfoot** back in the spotlight

SpiderFoot is a tool that automates open source intelligence (OSINT) collection for threat intelligence and attack surface mapping processes. This Python-based software accelerates security analysis by collecting data about target systems.

Python · ★ 20,413 · +74 today <https://github.com/smicallef/spiderfoot>

08 **localsend/localsend**

LocalSend is an open source application that allows file sharing between different operating systems. It creates a platform-independent alternative to the AirDrop feature offered by Apple devices.

Dart · ★ 87,961 · +213 today <https://github.com/localsend/localsend>

09 **Lightricks/LTX-2** back in the spotlight

LTX-2, developed by Lightricks, is an artificial intelligence model (generative model) that produces audio and video. This package provides the necessary Python tools to run the model's inference processes and fine-tune it using the low-order adaptation (LoRA) method.

Python · ★ 8,759 · +65 today <https://github.com/Lightricks/LTX-2>

10 **embabel/embabel-agent**

Embabel is a framework designed to develop artificial intelligence agents running on the Java Virtual Machine (JVM). This structure, written in Kotlin, facilitates the integration of generative AI capabilities in enterprise software ecosystems.

Kotlin · ★ 4,255 · +40 today <https://github.com/embabel/embabel-agent>

11 **cactus-compute/needle**

Developed by Cactus Compute, Needle offers a 14 MB foundation model that can run on small hardware such as phones, wearable devices and robots. This lightweight structure aims to run artificial intelligence applications locally on edge devices with limited processing power.

Python · ★ 4,412 · +315 today <https://github.com/cactus-compute/needle>

01 Tailscale Traces Database Corruption to 16y/o SQLite WAL-Reset Bug

Network security tool Tailscale found that a 16-year-old SQLite write-ahead log reset error was at the root of the database corruption. The company eliminated the data integrity risk by releasing the update that fixed this technical problem.

914 points · 177 comments <https://tailscale.com/blog/sqlite-wal-reset-bug>

Discussion: news.ycombinator.com/item?id=49272832 · lobste.rs/s/e0lkmi/how_tailscale_helped_find_sqlite_wal (914 oy · 177 yorum) (91 oy · 3 yorum)

02 DeepSeek V4 Pro 0813

DeepSeek V4 Pro 0813 is a new language model available through the OpenRouter platform that provides access to multi-modal AI models. This model aims to demonstrate high performance in complex logical inference and coding tasks.

831 points · 330 comments <https://openrouter.ai/deepseek/deepseek-v4-pro-0813>

Discussion: news.ycombinator.com/item?id=49274600 (831 oy · 330 yorum)

03 AI is removing the middle class of software engineering?

Artificial intelligence tools reduce the workspace of mid-level software developers by automating entry-level tasks in software development processes. This change brings with it a new workforce structure that allows companies to achieve higher efficiency with fewer senior engineers.

812 points · 740 comments <https://blog.florianherrengt.com/ai-removing-middle-class-software-engineering.html>

Discussion: news.ycombinator.com/item?id=49271994 (812 oy · 740 yorum)

04 License plate reader searches should require a warrant

It is being discussed that vehicle tracking searches made through automatic license plate readers should be subject to a legal search warrant on the grounds of personal privacy. The focus is on the risks posed by the uncontrolled use of data collected by these systems on individuals' freedom of movement.

568 points · 336 comments <https://andrewpwheeler.com/2026/08/12/license-plate-reader-searches-should-require-a-warrant/>

Discussion: news.ycombinator.com/item?id=49273165 (568 oy · 336 yorum)

01 Qwen/Qwen3.8-2.4T-A95B

Developed by Alibaba, Qwen3.8-2.4T-A95B is a large language model focused on text generation. This model, shared on the Hugging Face platform, is used in various natural language processing tasks with its structure trained on large data sets.

♥ 587 · 978 downloads <https://huggingface.co/Qwen/Qwen3.8-2.4T-A95B>

Discussion: news.ycombinator.com/item?id=49273478 (558 oy · 128 yorum)

01 Persistent Recursive Worlds Enable Autonomous Software Evolution

EvoX Genesis defines software projects as persistent and recursive worlds, ensuring a continuous development process despite the limited lifespan of artificial intelligence agents. This approach maintains the software itself as a centralized structure, allowing local coding agents to progress the project with ad-hoc tasks.

1 likes · 1 comments <https://huggingface.co/papers/2608.10450>

02 OpenART: Scaling Agent Red Teaming via Open-Ended Environment Evolution

OpenART is an open-ended testing environment developed to evaluate the cumulative risks of artificial intelligence agents in long-term workflows. Unlike traditional static benchmark tests, it aims to detect vulnerabilities (red teaming) in a scalable way by simulating the interactions of agents on a shared situation.

62 likes · 1 comments <https://huggingface.co/papers/2608.00677>

03 AtlasVLA: Persistent World-Ego State Modeling for Vision-Language-Action Models

AtlasVLA is a Vision-Language-Action model framework that enables robots to remember objects out of view and task progress. This method allows robotic systems to overcome the limitation of instantaneous response and manage long-term and complex tasks more consistently.

1 likes · 1 comments <https://huggingface.co/papers/2608.06729>

04 Self-Geometry: GT-Free and Plug-and-Play Test-Time Adaptation for Geometrically Consistent 3D Vision Foundation Models

Self-Geometry is a method that eliminates geometric inconsistencies in data such as depth and camera position produced by vision foundation models without requiring additional training during the testing phase. This approach makes the three-dimensional outputs of models more consistent, without the need for traditional methods that create complex computational costs.

8 likes · 1 comments <https://huggingface.co/papers/2608.10708>

05 Can LLM Agents Stick to the Script? A Benchmark for Long-Horizon Consistency in Interactive Narratives

Large language models have difficulty maintaining logical consistency despite user interventions in interactive storytelling. Developed by researchers, NCP-Bench offers a standard evaluation framework to measure the long-term narrative integrity of these models.

20 likes · 2 comments <https://huggingface.co/papers/2608.08160>

01 I hate packaging my software for Linux

Software developers face serious compatibility issues in the process of packaging applications due to the diversity of Linux distributions. Different system requirements and package formats complicate developers' deployment processes, leading to loss of time.

71 points · 57 comments <https://getfresh.dev/docs/blog/packaging-for-linux/>

Discussion: lobste.rs/s/ixbrmc/i_hate_packaging_my_software_for_linux (71 oy · 57 yorum)

02 Guide to (not) fucking up QR codes

The use of QR codes involves technical details that require attention in terms of security and accessibility. To reduce the risks of faulty design and misdirection, the visual clarity of the codes and connection security need to be verified.

83 points · 8 comments <https://infosec.exchange/@rebane2001/117078420917152774>

Discussion: lobste.rs/s/omcyaz/guide_not_fucking_up_qr_codes (83 oy · 8 yorum)

SVG

A digital image file drawn with mathematical formulas that can be enlarged or reduced as desired without losing quality.

ADE

A digital working environment where all tools are used to accelerate software development processes.

framework

A set of libraries and tools that provide the basic structure and make things easier when developing software.

foundation model

A basic artificial intelligence model designed to perform many different tasks by being trained with a wide range of data.

PowerPoint

A presentation preparation program that allows you to present information in visual slides.

LLM

Artificial intelligence models trained with huge data that can understand human language and produce text.

RAG

A method that allows artificial intelligence to extract information from your current and private documents while producing answers.

inference

The process of a trained artificial intelligence producing an answer to a new question using the information it has learned.

runtime

The time during which a software is run on the computer and actively performs its task.

OSINT

A method of collecting and analyzing information from publicly available sources on the Internet.

generative model

Artificial intelligence that can create new content that it has not seen before, from scratch using the patterns it has learned.

LoRA

It is a technique to specialize a large artificial intelligence model on a specific subject by updating only a small part of it, without changing the entire model.

TreScout · trescout.com · Every day, on your schedule.