



TreScout

FRESH ONLY · DAILY TECHNOLOGY REPORT

Today's New Trends

August 14, 2026 (Friday)

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.

19 HIGHLIGHTS · ~5 MIN READ

6 GitHub
Trending

5 Hacker
News

1 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

2 Lobsters

01 altic-dev/FluidVoice back in the spotlight

FluidVoice is a dictation application that runs on macOS and performs speech-to-text conversion (STT) entirely on the device. This tool protects privacy by processing user data locally and offers an alternative solution to services such as Wispr Flow with similar functions.

Swift · ★ 9,936 · +76 today <https://github.com/altic-dev/FluidVoice>

02 unslothai/unsloth

Unsloth is a library that accelerates the training and execution of large language models (LLM) and diffusion models in a native environment. It makes fine-tuning processes of popular models more accessible by optimizing memory usage.

Python · ★ 71,157 · +328 today <https://github.com/unslothai/unsloth>

03 megadose/holehe

Holehe is a Python tool that queries whether an entered email address is used on popular platforms such as Twitter and Instagram. It detects the account information associated with the relevant email using password reset functions.

Python · ★ 12,498 · +195 today <https://github.com/megadose/holehe>

04 holaboss-ai/holaOS

holaOS is an open source workspace that combines different artificial intelligence agents (Claude Code, Codex) with your tools, applications and files over a common memory. With support for more than 100 integrations and Model Link Protocol (MCP), it enables users to manage their own models (BYOK) or built-in models in a single interface.

TypeScript · ★ 6,748 · +241 today <https://github.com/holaboss-ai/holaOS>

05 kepano/obsidian-skills

Obsidian-skills enables artificial intelligence agents to perform operations on the note-taking application Obsidian using the command line interface (CLI) and open file formats. This tool automates personal information management by allowing agents to read and edit formats such as Markdown and JSON Canvas.

★ 45,950 · +292 today <https://github.com/kepano/obsidian-skills>

06 lightningpixel/modly

Modly is a desktop application that creates 3D models from images and performs all operations on the local graphics card (GPU). This tool, which works without the need for an internet connection, brings the artificial intelligence-supported modeling process to personal computers.

TypeScript · ★ 5,550 · +118 today <https://github.com/lightningpixel/modly>

01 Gemini 3.7 Flash

Gemini 3.7 Flash, announced by Google, is a new generation generative artificial intelligence (generative AI) model focused on high speed and low cost. The model provides increased performance on complex reasoning tasks and offers faster response times for developers.

716 points · 393 comments <https://blog.google/innovation-and-ai/models-and-research/gemini-models/introducing-gemini-3-7-flash/>

Discussion: news.ycombinator.com/item?id=49289112 (716 points · 393 comments)

02 DeepSeek Harness developer preview

DeepSeek Harness is an evaluation framework used to measure the performance of large language models. It allows developers to benchmark their models against standard test sets and systematically analyze model capabilities.

604 points · 255 comments <https://deepseek.com/harness/en/>

Discussion: news.ycombinator.com/item?id=49285244 (604 points · 255 comments)

03 Spaghettifying DRAM

Spaghettifying DRAM, which provides random number generation by taking advantage of the physical properties of the data in DRAM memory modules, works through a hardware-based vulnerability. This method increases the predictability of cryptographic keys by exploiting leaks in memory cells.

546 points · 150 comments <https://github.com/xoreaxeaxeax/skitter-creek-bath-salts>

Discussion: news.ycombinator.com/item?id=49286341 (546 points · 150 comments) · lobste.rs/s/a4zifd/unlocking_everything_on_cpu_with_dram (34 points · 8 comments)

04 Accelerating GPT-5.6 Sol Ultrafast

Cerebras, a company that produces artificial intelligence hardware, has developed a new technology that increases the inference speed for OpenAI models. This method aims to increase processing efficiency by reducing the response time of large language models.

507 points · 208 comments <https://www.cerebras.ai/blog/accelerating-gpt-5-6-sol-ultrafast-with-openai>

Discussion: news.ycombinator.com/item?id=49289844 (507 points · 208 comments)

05 Choose Boring Technology (2015)

An approach that advocates choosing proven and stable technologies in software development processes. He states that choosing highly reliable infrastructures instead of innovative tools reduces the operational burden of teams and reduces the risk of technical debt.

298 points · 147 comments <https://mcfunley.com/choose-boring-technology>

Discussion: news.ycombinator.com/item?id=49289512 (298 points · 147 comments)

01 Lightricks/LTX-2.5

Developed by Lightricks, LTX-2.5 is an image-to-video production model that converts still images into short video clips. This technology, used to animate visual content, allows users to create high-quality videos from static frames.

♥ 745 · 57,287 downloads <https://huggingface.co/Lightricks/LTX-2.5>

01 UniSwap: Streaming Audio-Visual Identity Swapping for Talking Videos

UniSwap offers a framework that simultaneously changes the visual and audio identity in speaking videos. It integrates reference visual and audio data into the video, preserving the motion, scene and timing in the source video.

1 likes · 1 comments <https://huggingface.co/papers/2608.11752>

02 LiveAnimate: Stable Long-Form Streaming Human Animation in Real-Time

LiveAnimate is a system that produces real-time human animations from a single reference image and motion stream. It optimizes the long processing times of diffusion-based models, allowing seamless video creation in interactive areas such as live broadcasts and virtual avatars.

1 likes · 1 comments <https://huggingface.co/papers/2608.11745>

03 CW-BASS v2: Saturation-Aware Pseudo-Label Selection for Semi-Supervised Segmentation under Foundation-Model Teachers

CW-BASS v2 improves semi-supervised semantic segmentation processes, which traditional methods are insufficient due to the high confidence values of foundation model teachers. This new saturation-aware approach improves model performance by optimizing the selection of labels produced by powerful models.

0 likes · 1 comments <https://huggingface.co/papers/2608.12773>

04 H2R-Bench: Benchmarking Human-to-Robot Manipulation Video Generation in World Models

H2R-Bench is a benchmarking framework developed to evaluate video production models that transfer human hand movements to robotic end-effectors. It measures the performance of world models that synthesize robot-centered manipulation videos from human videos in order to reduce the cost of data collection in robotic learning processes.

2 likes · 1 comments <https://huggingface.co/papers/2608.13049>

05 LLMRouter: Unified Infrastructure for Developing, Evaluating, and Deploying LLM Routers

LLMRouter offers a unified infrastructure to develop, evaluate and deploy routers that select the optimal large language model for different queries and budget constraints. This system breaks down the routing process into five key components, making it easy to compare the performance of different models and create scalable workflows.

48 likes · 1 comments <https://huggingface.co/papers/2608.06867>

01 I want extern "fil-c"

Software developers are looking for an external standard library to make file system operations (file I/O) in C more secure and manageable. This approach aims to reduce the margin of error in software development processes by making low-level system calls more modular.

46 points · 16 comments <https://domenkozar.com/2026/08/13/i-want-extern-fil-c/>

Discussion: lobste.rs/s/tssf5y/i_want_extern_fil_c (46 points · 16 comments)

02 python's pre-declared constants are kinda weird

Pre-declared constants in the Python language contain some inconsistencies that contradict the dynamic structure of the language. This brings into question technical limitations that make it difficult for developers to maintain full control over the code.

91 points · 15 comments <https://sebsite.pw/w/20260801-pythonconstants.html>

Discussion: lobste.rs/s/62kxco/python_s_pre_declared_constants_are_kind (91 points · 15 comments)

generative AI

Artificial intelligence technology that can create new texts, images or sounds based on previously learned information.

skills

Special abilities that a software or artificial intelligence has to perform a specific task.

foundation model

The basic artificial intelligence structure that is trained with very large data sets and on which different tasks can be built.

STT

Technology that listens to spoken words and converts them into written text.

LLM

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

diffusion models

An artificial intelligence method that creates high-quality visuals from scratch by gradually editing noisy data.

fine-tuning

The process of customizing a general-purpose trained artificial intelligence with a small data set to perform better on a specific subject.

large language models

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

MCP

A protocol that enables artificial intelligence models to connect to different data sources and tools in a standard way.

BYOK

A security approach that allows the user to encrypt their data using their own security keys.

CLI

Text-based interface that runs programs with written commands rather than the mouse.

inference

The process of a trained artificial intelligence model producing a result by processing a new input given to it.
