



TreScout

FRESH ONLY · DAILY TECHNOLOGY REPORT

Today's New Trends

August 2, 2026 (Sunday)

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.

20 HIGHLIGHTS · ~30 MIN READ

7 GitHub
Trending

5 Hacker
News

1 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

2 Lobsters

01 microsoft/generative-ai-for-beginners

This training content prepared by Microsoft offers a 21-lesson curriculum for those who want to get into generative artificial intelligence (generative AI) technologies. This resource, in Jupyter Notebook format, teaches basic concepts and implementation methods for developers to create their own projects.

Jupyter Notebook · ★ 114,349 · +108 today <https://github.com/microsoft/generative-ai-for-beginners>

02 github/gh-stack

Developed by GitHub, gh-stack is a command line tool that makes it easy to create and manage stacked pull requests during the software development process. It aims to speed up the review process by dividing complex code changes into small and independent parts.

Go · ★ 860 · +46 today <https://github.com/github/gh-stack>

03 abus-aikorea/voice-pro

Voice-pro is an open-source interface that combines text-to-speech (TTS) and zero-shot voice cloning tools. Users can perform operations such as audio parsing, multilingual translation and processing YouTube content on a single platform via its web-based interface.

Python · ★ 11,859 · +58 today <https://github.com/abus-aikorea/voice-pro>

04 iv-org/invidious

Invidious is an open source interface developed as an alternative to the YouTube platform. It offers users a video watching experience that does not contain ads and does not track data.

Crystal · ★ 21,676 · +435 today <https://github.com/iv-org/invidious>

05 microsoft/TRELLIS.2

Developed by Microsoft, TRELLIS is a structured latent representation architecture that transforms text or visual inputs into three-dimensional (3D) models. This technology uses generative AI methods to create high-quality 3D assets quickly and compactly.

Python · ★ 9,992 · +107 today <https://github.com/microsoft/TRELLIS.2>

06 NomaDamas/k-skill

NomaDamas/k-skill is a library designed to provide AI agents with local skills in a cultural context. It offers customized skill sets that enable agents to interact with users more naturally and in accordance with cultural codes.

JavaScript · ★ 6,780 · +53 today <https://github.com/NomaDamas/k-skill>

07 bytedance/deer-flow back in the spotlight

Developed by ByteDance, Deer-flow is a super agent infrastructure that performs long-term complex tasks through research, coding and content production. It autonomously manages processes lasting minutes or hours using components such as sandbox, memory and sub-agents.

Python · ★ 78,814 · +209 today <https://github.com/bytedance/deer-flow>

01 How Google helped destroy adoption of RSS feeds (2023)

In 2013, Google shut down its popular content monitoring tool Google Reader, reducing the widespread use of RSS feeds (Really Simple Syndication) on the internet. This move ushered in an era of content consumption in which centralized platforms came to the fore and open web standards were replaced by closed algorithms.

487 points · 163 comments <https://openrss.org/blog/how-google-helped-destroy-adoption-of-rss-feeds>

Discussion: news.ycombinator.com/item?id=49136821 (487 oy · 163 yorum)

02 Diátaxis

Diátaxis is a systematic framework that divides software documentation into four basic categories: learning, practicing, reference and description. This method aims to increase the understandability of technical documentation by providing content structured according to users' needs.

277 points · 38 comments <https://diataxis.fr/>

Discussion: news.ycombinator.com/item?id=49138188 (277 oy · 38 yorum)

03 Seedance 2.5

Seedance 2.5, the video production tool developed by ByteDance, allows users to edit single-shot videos and create consistent content using reference images. This version offers flexible visual control and high-accuracy results in AI video editing processes.

271 points · 136 comments <https://seed.bytedance.com/en/blog/one-take-creation-flexible-referencing-introducing-seedance-2-5>

Discussion: news.ycombinator.com/item?id=49138302 (271 oy · 136 yorum)

04 AI financial advice is surprisingly good, especially if you ask right questions

Artificial intelligence models achieve high accuracy rates in personal financial advice when asked the right guiding questions. Financial literacy tools have the potential to develop personalized strategies for individuals by analyzing complex economic data.

259 points · 218 comments <https://mitsloan.mit.edu/ideas-made-to-matter/ai-financial-advice-surprisingly-good-especially-if-you-ask-right-questions>

Discussion: news.ycombinator.com/item?id=49139102 (259 oy · 218 yorum)

05 NetBSD 11.0

NetBSD 11.0 is the latest version of NetBSD, an open source operating system focused on high portability. The new version includes several improvements that increase system security and performance while expanding hardware support.

255 points · 115 comments https://blog.netbsd.org/tnf/entry/netbsd_11_0_released

Discussion: news.ycombinator.com/item?id=49136736 (255 oy · 115 yorum) · lobste.rs/s/pvrnor/netbsd_11_0_released (65 oy · 11 yorum)

01 unsloth/DeepSeek-V4-Flash-0731-GGUF

Unsloth, a library that optimizes large language models, has released a compressed version of the DeepSeek-V4 model in GGUF format to run it faster and with lower memory usage. This work enables high-performance models to run more efficiently on consumer-grade hardware.

♥ 297 · 4,048 downloads <https://huggingface.co/unsloth/DeepSeek-V4-Flash-0731-GGUF>

01 **OmniScope: Modality-Decoupled Token Compression for Omnimodal Large Language Models**

OmniScope aims to solve the data loss problem in token compression methods used in omnimodal large language models. Stating that different modes gain importance at different moments for the same query, this framework ensures the protection of critical information with its structure that does not require training.

3 likes · 2 comments <https://huggingface.co/papers/2607.23193>

02 **β -OPSD: Deriving with Policy Optimization, Training with Self-Distillation**

β -OPSD redefines the self-distillation method, which improves reasoning ability in language models, through the beta value, an optimizeable parameter. This approach aims to reduce instability in the training process by controlling the student model's dependence on the reference model.

18 likes · 2 comments <https://huggingface.co/papers/2607.28582>

03 **Beyond Geometric Complementarity: Coherent Overlap in Sparse Mixture-of-Experts Routing**

In sparse mixture-of-experts language models, the mechanism of directing tokens to different experts is examined using the expert subspace separation index (ESSI). The research more clearly defines routing strategies by distinguishing between the consistency in the selection of experts and their contribution to model performance.

2 likes · 2 comments <https://huggingface.co/papers/2607.28308>

04 **Fairness Pruning: Locating Demographic Bias in GLU-MLP Layers via Differential Activations**

Fairness Pruning is a structured intervention method developed to detect and eliminate demographic biases in large language models (LLM). The method identifies neurons that respond differently to certain demographic attributes in models with GLU architecture by analyzing their activation differences during inference.

2 likes · 3 comments <https://huggingface.co/papers/2607.28319>

05 **Σ -Mem: An Online Reliability Memory for LLM-based Multi-Agent Systems**

Σ -Mem is a memory mechanism that tracks the trustworthiness of agents to each other in multi-agent systems based on large language models. The system records agents' past performance and evidence of competence, enabling more reliable collaboration in complex workflows.

12 likes · 2 comments <https://huggingface.co/papers/2607.27958>

01 An old-new take on argument parsing in Rust

The process of parsing command line arguments in the Rust programming language is being re-evaluated as developers turn to simpler and more direct methods. This approach aims to reduce dependencies in the code base by using standard or lightweight tools instead of complex libraries.

19 points · 5 comments <https://jmmv.dev/2026/07/hello-getoptsargs.html>

Discussion: lobste.rs/s/vrd28h/old_new_take_on_argument_parsing_rust (19 oy · 5 yorum)

02 Resigning from Arch Linux

Arch Linux developer Lukas Märdian announced that he is leaving his position on the distribution's management team. This decision is justified by structural changes in the development processes of the distribution and changing personal priorities.

24 points · 1 comments <https://linderud.dev/blog/resigning-from-arch-linux/>

Discussion: lobste.rs/s/ectcmr/resigning_from_arch_linux (24 oy · 1 yorum)

generative AI

Artificial intelligence technology that can produce new content such as text, images or audio using existing data.

stacked pull requests

A method of sequentially introducing a major software change into the system in small, manageable, interconnected pieces.

TTS

The technology of voicing written texts with a natural human voice by a computer.

zero-shot voice cloning

The technology of copying a person's voice from just a few seconds of sample and making that person say sentences they have never said.

structured latent

Converting complex data into an orderly and meaningful mathematical form that artificial intelligence can understand more easily.

AI agents

Autonomous software that can make decisions on its own and use tools to perform targeted tasks.

skill sets

The set of abilities that a person or artificial intelligence has to do a specific job.

SuperAgent

Advanced artificial intelligence software that can manage multiple complex tasks simultaneously and work integrated with different systems.

sandbox

An isolated working environment where software can be tested safely without damaging the main system.

Really Simple Syndication

Content distribution system that allows updates or news on websites to be followed from a single place.

AI video editing

The process of automatically performing operations such as cutting, adding effects or image enhancement in videos with the help of artificial intelligence.
