



TreScout

DAILY TECHNOLOGY REPORT

Today's Technology Trends

August 17, 2026 (Monday)

A FEW SIPS

Today's agenda is based on browser-based adblocker integrations, system prompts and process evaluation tools in large language models. In the world of open source models, new high-capacity releases like Qwen and Muse-Glimmer are causing developers to reconsider their hardware and architecture choices.

25 HIGHLIGHTS · ~7 MIN READ

7 GitHub
Trending

5 Hacker
News

5 HuggingFace ·
Daily Models

5 HuggingFace ·
Daily Papers

3 Lobsters

01 cordiverse/cordis

Cordis is a meta-framework that allows spatial and temporal components to be managed together. This TypeScript-based structure facilitates modularity and compatible components in complex systems.

TypeScript · ★ 5,024 · +720 today <https://github.com/cordiverse/cordis>

02 basecamp/omarchy

Developed by Basecamp, Omarchy is a modern and original configuration tool for Linux operating systems. It offers a shell script-based approach to standardize system installations and facilitate management.

Shell · ★ 25,537 · +270 today <https://github.com/basecamp/omarchy>

03 unslothai/unsloth

Unsloth is a library developed to train and run large language models (LLM) and rendering models faster on local computers. By optimizing memory usage, it accelerates fine-tuning processes and reduces hardware requirements.

Python · ★ 72,741 · +572 today <https://github.com/unslothai/unsloth>

04 OpenCut-app/OpenCut

OpenCut offers an open source alternative to the video editing tool CapCut. Written in TypeScript, this software allows users to edit their videos locally.

TypeScript · ★ 84,046 · +150 today <https://github.com/OpenCut-app/OpenCut>

05 public-apis/public-apis

Public-apis offers a categorized list of free application programming interfaces (APIs) that software developers can use in their projects. This resource, which provides access to data from various sectors, gathers the external services needed in the development processes in a single center.

Python · ★ 462,093 · +1,588 today <https://github.com/public-apis/public-apis>

06 ToolJet/ToolJet

ToolJet is an open source application development platform for creating in-house tools, dashboards and artificial intelligence agents. This JavaScript-based infrastructure offers rapid application development via a visual interface to automate corporate workflows.

JavaScript · ★ 40,158 · +452 today <https://github.com/ToolJet/ToolJet>

07 **cactus-compute/needle**

Developed by Cactus Compute, Needle is a 14 MB foundation model that can run on small hardware such as smartphones, wearable devices and robots. Thanks to its lightweight structure, it allows local artificial intelligence operations to be performed on end devices without requiring an internet connection.

Python · ★ 6,731 · +443 today <https://github.com/cactus-compute/needle>

01 Firefox for iOS now has a native adblocker

Mozilla added a built-in ad blocker feature to the Firefox internet browser it developed for iOS. This update allows users to block ads and trackers on web pages directly from the browser without the need for a third-party plug-in.

581 points · 239 comments <https://support.mozilla.org/en-US/kb/block-ads-firefox-ios>

Discussion: news.ycombinator.com/item?id=49319633 (581 points · 239 comments)

02 Claude: System Prompts

Anthropic has made the system prompts feature available to developers for its artificial intelligence model Claude. This feature allows developers to manage the model's behavior, tone, and constraints with predefined instructions.

580 points · 240 comments <https://platform.claude.com/docs/en/release-notes/system-prompts>

Discussion: news.ycombinator.com/item?id=49319556 (580 points · 240 comments)

03 A 3rd World Embedded Engineer Responds to "RISC-V They Should Have Known Better"

Embedded systems engineers in developing countries state that the low-cost and open-source nature of RISC-V architecture offers a critical opportunity for local technology ecosystems. This architecture enables independent development by democratizing hardware design processes that are inaccessible due to licensing fees.

412 points · 225 comments https://rvembedded.com/blog_post/12/

Discussion: news.ycombinator.com/item?id=49321717 (412 points · 225 comments)

04 Tell HN: Cloudflare silently injects its analytics when you switch nameservers

Internet infrastructure provider Cloudflare automatically adds analytics snippets to websites if users move their name servers to its own system. This causes tracking scripts to be included even in static pages that do not contain any JavaScript code.

346 points · 91 comments <https://news.ycombinator.com/item?id=49322107>

Discussion: news.ycombinator.com/item?id=49322107 (346 points · 91 comments)

05 The AI Credit Resale Economy

An economy is emerging where API credits that provide access to artificial intelligence models are resold at discounted prices in secondary markets. Token brokers acquire large-scale usage rights at low cost and market these tokens to smaller developers or businesses.

248 points · 98 comments <https://vectorial.com/blog/who-are-the-token-brokers>

Discussion: news.ycombinator.com/item?id=49320611 (248 points · 98 comments)

01 Qwen/Qwen3.8-27B

Developed by Alibaba, Qwen2.5-VL is a multimodal artificial intelligence model (multimodal AI) that produces text-based outputs by processing image and text data. It aims to increase visual understanding capacity among open source language models (LLM) with the ability to analyze visual content and relate it to text.

♥ 10,358 · 267,725 downloads <https://huggingface.co/Qwen/Qwen3.8-27B>

02 meta-models/Muse-Glimmer-30B

Muse-Glimmer-30B, developed by Meta-models, is a multimodal large language model that can process visual and textual data simultaneously. The model works with the ability to produce text-based outputs by combining image and text inputs.

♥ 1,636 · 292,973 downloads <https://huggingface.co/meta-models/Muse-Glimmer-30B>

03 unsloth/Qwen3.8-27B-GGUF

Unsloth is an optimization library that allows training large language models faster and with lower memory usage. This version of the Qwen2.5-27B model has been optimized to run more efficiently on standard hardware thanks to the GGUF format.

♥ 1,491 · 1,945,635 downloads <https://huggingface.co/unsloth/Qwen3.8-27B-GGUF>

04 Qwen/Qwen3.8-2.4T-A95B

Developed by Alibaba, Qwen3.8-2.4T-A95B is a text-generation model trained on large-scale data sets. The model offers a performance-oriented structure in complex language tasks with its high number of parameters and comprehensive training data.

♥ 1,015 · 7,932 downloads <https://huggingface.co/Qwen/Qwen3.8-2.4T-A95B>

05 Lightricks/LTX-2.5

Developed by Lightricks, LTX-2.5 is an image-to-video model that converts still images into short videos. It aims to speed up visual production processes by converting photos uploaded by users into animated content.

♥ 1,038 · 424,099 downloads <https://huggingface.co/Lightricks/LTX-2.5>

01 PRM-as-a-Judge 1.5: A Toolkit for Robot Process Assessment

PRM-as-a-Judge 1.5 is a toolkit that evaluates the performance of robotic models with detailed process metrics beyond binary success rates. It analyzes the operations performed by robots through video recordings and provides concrete data on issues such as margin of error, recovery process and application quality.

2 likes · 1 comments <https://huggingface.co/papers/2608.14284>

02 MobileMem: Learning from a Year of Mobile Experiences

MobileMem offers a new framework that enables artificial intelligence assistants to continuously learn by keeping user experiences in long-term memory. It aims to enable assistants to have a memory structure that develops over time by processing complex, multimodal and personal data on mobile devices.

6 likes · 1 comments <https://huggingface.co/papers/2608.13606>

03 Multimodal Model Diffing for Feature Discovery and Control

Developed to parse internal features in multimodal large language models, MMDiff enables to identify components that change during the training process of the model. This method enables inspection and targeted control of complex visual-textual structures using sparse autoencoders.

2 likes · 1 comments <https://huggingface.co/papers/2608.09928>

04 CPI-Bench: A Comprehensive, Practical and Intelligent Benchmark for Real-World Image Editing

CPI-Bench is a comprehensive benchmark tool developed to evaluate visual editing models in real-world scenarios. Unlike existing tests, it aims to analyze the performance of models with more realistic criteria in complex and multi-image tasks.

8 likes · 1 comments <https://huggingface.co/papers/2608.14546>

05 Generation as Auxiliary Supervision: Enhancing Visual Understanding at Zero Inference Overhead via Decoupled Embedding Prediction

GAS improves the performance of multimodal large language models with a training method that combines visual comprehension and production. The method enables more effective processing of visual data by using decoupled embedding prediction without increasing the inference cost.

0 likes · 0 comments <https://huggingface.co/papers/2608.12209>

01 I thought I was building a C replacement. I was wrong

Although the C3 programming language was developed to eliminate the shortcomings of the C language, it is a system language that has evolved into its own unique structure over time. Developers note that it is more efficient to design a new language (C3) that focuses on modern software needs rather than modifying an existing language.

66 points · 37 comments https://c3-lang.org/blog/i_thought_i_was_building_a_c_replacement/

Discussion: lobste.rs/s/wpyhdh/i_thought_i_was_building_c_replacement_i (66 points · 37 comments)

02 Every Fucking Website

Interface elements such as cookie consents, newsletter sign-up windows and notification requests that hide the content of modern websites turn into a mess that negatively affects the user experience. This situation increases criticism against the spread of dark patterns that make it difficult to access information on the internet.

59 points · 27 comments <https://op.tnql.io/every-fucking-website/>

Discussion: lobste.rs/s/ii5tun/every_fucking_website (59 points · 27 comments)

03 A faster way to calculate the day-of-the-week

The new algorithm developed by Ben Joffe offers a mathematical shortcut to calculate which day of the week any date falls on. This method provides fast results with fewer processing steps compared to traditional calendar calculations (Zeller formula).

21 points · 4 comments <https://www.benjoffe.com/fast-day-of-week>

Discussion: lobste.rs/s/ln3dao/faster_way_calculate_day_week (21 points · 4 comments)

meta-framework

An advanced framework built on a software development tool that speeds up work and offers ready-made solutions.

shell script

Small scripts that automatically make the computer's operating system do the tasks it needs to do one after the other.

LLM

Advanced artificial intelligence system that can speak and write like a human by reading huge amounts of text on the Internet.

fine-tuning

Retraining a general-purpose trained artificial intelligence with smaller and more specific data to specialize in a specific subject.

API

A digital bridge that allows two different software to talk to each other and exchange data.

foundation model

The main model that forms the basis of artificial intelligence applications, trained on large data sets for use in many different tasks.

adblocker

Browser add-on that prevents the loading of advertisements that appear on websites.

system prompts

Initial instructions telling the AI how it should behave and what rules it should follow.

nameservers

Directory servers that convert the names of websites into numerical addresses that computers can understand.

analytics snippets

Small pieces of code used to track how many people visit a website and what they do.

multimodal AI

Not just with text; Artificial intelligence that can simultaneously process different types of data such as image, audio and video.

multimodal large language model

An advanced artificial intelligence model that can understand and comment on different content such as images or sounds in addition to text.

large language models

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

text-generation model

Artificial intelligence that creates meaningful texts by following a starting sentence or topic given to it.

image-to-video

Artificial intelligence technology that takes a still photo and turns it into a moving video.

multimodal

The ability to use multiple different types of communications or data simultaneously.

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