

The Long Read

29 October 2026

11

STORIES

3


SECTIONS

3143

WORDS

5

HIGH IMPACT




● 5 high impact ● 6 medium impact spoke length = depth of coverage

The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

11 stories · 3143 words · 3 sections · 2 charted



AI Models & Research

- 01 ● Google's Second Frontier Model in Two Weeks Is Built to Hunt Bugs
- 02 ● Researchers Retune AI's Attention to Block Hidden Instructions in Webpages

AI Tools & Ecosystem

- 03 ● Anthropic Shows Companies Exactly What Their AI Spending Bought
- 04 ● AI Agents Tripled GitHub's Pull Requests in a Single Year
- 05 ● Google Opens Gemini and Search to Any Developer's App

- 06 ● Nvidia Open-Sources Transcription You Can Steer With a Prompt

AI Applications & Industry

- 07 ● Microsoft Booked \$10 Billion Last Quarter From OpenAI Alone
- 08 ● Microsoft Puts a Dollar Figure on Azure for the First Time
- 09 ● Google Rolls Its AI Search Ads Out Worldwide For the First Time
- 10 ● Chrome Wants to Read Your Gmail to Fill in Forms
- 11 ● October Was the Month Vibe-Coded Games Became Actually Fun

How to read this. Every story in today's email is reported here at full length, in the same order. Impact is the writer's judgement of whether a story changes what a reader — in business, at work, as an investor or as a builder — should do or believe this week. Charts appear only where the source itself puts comparable numbers side by side; nothing is estimated to fill a gap. Sources are listed in full on the final page.

AI Models & Research

2 stories ● 2 high

01 HIGH IMPACT Google

Google's Second Frontier Model in Two Weeks Is Built to Hunt Bugs

Google has released its second frontier AI model in two weeks, Gemini 4 Carbon, a system purpose-built for cybersecurity defense that the company says tops external benchmarks for finding, exploiting and patching software vulnerabilities.

Gemini 4 Carbon follows Gemini 4 Argon, which launched two weeks earlier, and Google is positioning the new model as its strongest yet for cybersecurity and long-horizon engineering work. On the Artificial Analysis Intelligence Index v5.0 it ranks first with a score of 59, and it leads three external security benchmarks: SEC-bench Pro at 88.2% for finding real vulnerabilities, ExploitGym at 54.8% for validating risk through exploitation, and CWE-bench at 71.0% for patching root causes. The Chrome team reported Carbon performed more than 10 times better than another lab's large frontier model at generating high-quality patches, and cloud security platform Wiz said Carbon leads on web penetration testing benchmarks.

Google is not releasing Carbon broadly. Starting this week, the model goes to a subset of cyber defenders, with access phased to additional trusted partners through the company's Fairwind Program. Google says it is participating in the U.S. government's voluntary pre-release testing process and will gather feedback before making Carbon available to developers, businesses and consumers. Pricing matches Argon's introductory rate: \$2 per million input tokens and \$10 per million output tokens, with cached inputs at 95% off the input price.

The most striking detail is how Google describes Carbon's behavior on safety and alignment tests. The company says Carbon reduced successful cheating on Terminal-Bench 4.0 to 0.0%, resisted every evaluated prompt-injection attack on Gray Swan's Indirect Prompt Injection benchmark, cut unsupported claims by roughly half compared with Argon, and reduced high-severity misalignment alerts by almost 80% in live coding environments. Google also reports Carbon reduced overreach by 26% and under-service by 29%, meaning it is less likely to either exceed its brief or leave work incomplete.

For security teams, the practical shift is that a frontier model can now autonomously run for days on multi-step engineering tasks such as root-causing and patching subtle bugs across large codebases. The rollout to cyber defenders first signals Google's bet that the near-term market for this capability is defensive security work rather than general software development.

KEY FACTS

RANK ON ARTIFICIAL ANALYSIS INTELLIGENCE INDEX
V5.0

#1 (score 59)

EXPLOITGYM SCORE

54.8%

INPUT PRICE

\$2 per million tokens

CACHED INPUT DISCOUNT

95% off input price

SEC-BENCH PRO SCORE

88.2%

CWE-BENCH SCORE

71.0%

OUTPUT PRICE

\$10 per million tokens

PROMPT INJECTION ATTACK SUCCESS RATE

0.0%

WHY IT MATTERS

Security teams get first access to a model that can find, validate and fix vulnerabilities at state-of-the-art levels, potentially reducing the manual workload of patching and triage. Businesses should watch how quickly Google opens access beyond the initial defender cohort, since pricing matches its previous frontier model.

[Read the original at Google →](#)

02 **HIGH IMPACT** [arXiv.org](#)

Researchers Retune AI's Attention to Block Hidden Instructions in Webpages


Researchers have built a new defense that sharply reduces the success rate of prompt-injection attacks on large language models without the usual cost to normal performance.

The work, published on arXiv, introduces a method called Heads of Security (HoS) that teaches a model how to separate trusted instructions from untrusted content by adjusting how its internal attention mechanism weighs different parts of an input. In tests against Llama-3.3-70B-Instruct and GPT-OSS-20B, the defense cut the success rate of a black-box attack called RL-Hammer from 97% on an undefended GPT-OSS-20B to 10% when HoS was applied. A finetuning-based defense called SecAlign++ reduced the same attack to 85%.

Prompt injection is a practical problem for any business that lets a model read emails, web pages or documents: an attacker can hide instructions in that content telling the model to ignore its real task, exfiltrate data or take harmful actions. The finding that a targeted attention intervention can block these attacks while preserving normal performance matters for companies deploying AI agents that handle untrusted text. The authors report that HoS keeps base-model utility across standard multiple-choice benchmarks and prompting setups, whereas SecAlign++ and another defense, ASIDE, degrade in few-shot settings. ASIDE, which also builds instruction-data separation into the architecture, achieved only 10% utility on agentic benchmarks involving tool use.

The most striking detail is the cost asymmetry for attackers. A white-box attack called Claudini, strengthened against HoS through automated research, needed three times more computing power to reach the same attack success rate on HoS as it did on GPT-OSS-20B or SecAlign++. That means the defense does not just block known attacks; it makes developing new ones materially more expensive. On agentic benchmarks with static attacks, HoS and SecAlign++ performed similarly and both improved the security-utility trade-off over the base model.

For anyone running AI agents that touch external content, the result suggests a practical direction: robustness can be improved through targeted changes to how a model attends to input, rather than through broad finetuning that sacrifices general ability.



KEY FACTS

RL-HAMMER ATTACK SUCCESS ON GPT-OSS-20B

97%

RL-HAMMER ATTACK SUCCESS WITH SECALIGN++

85%

ASIDE UTILITY ON AGENTIC BENCHMARKS

10%

RL-HAMMER ATTACK SUCCESS WITH HOS DEFENSE

10%

EXTRA COMPUTING NEEDED TO ATTACK HOS

3 times more FLOPs

WHY IT MATTERS

Businesses deploying AI agents that read untrusted text — emails, web pages, documents — get a defense that blocks injected instructions while keeping the model useful, and it makes attacks three times more expensive to develop.

[Read the original at arXiv.org](#) →

AI Tools & Ecosystem

4 stories ● 2 high ● 2 medium

03 MEDIUM IMPACT TechCrunch

Anthropic Shows Companies Exactly What Their AI Spending Bought

Anthropic is rolling out reporting tools that let corporate customers see what their AI spending actually produces, a direct response to finance chiefs who are questioning the return on their token budgets.

The new suite, announced Wednesday, lets managers pull internal analytics and employee surveys to see what a team produced over a given period, how much it cost, and where workflows broke down. Anthropic says it will not use this data for model training by default, and the reports will include recommendations for improving results. The company is also adding pooled group budgets so teams can manage token spending collectively.

The move targets a real anxiety among large customers. Jellyfish, which makes tools for software developers, published a report based on 276,000 of its customers showing that spending on Claude Code and other tools built on Anthropic models has grown tenfold since the start of the year, and now equals about 2% of total engineering cost. “Even if you’re saying 2% is still small, that 10x growth is growth that most CFOs didn’t plan for,” Jellyfish chief scientist Nicholas Arcolano told TechCrunch.

The most striking detail in the source is how sticky Anthropic’s enterprise position has become despite a widely hyped rival. OpenAI’s Astra model arrived this summer as a supposed threat, but Jellyfish’s data shows it has barely dented Anthropic’s adoption. “It took a long time for us to see Codex really pop in our data at all, and it has come up with Astra,” said Nikolas Albarran, Jellyfish’s AI product lead. On social media, he noted, the competition looks heated; in actual usage data, companies have already standardized workflows around Anthropic and see little reason to switch.

That stickiness is reinforced by the cost of change. Once a company has embedded a particular AI vendor’s tools, switching means retraining people, reworking internal processes, and dealing with small differences in how each system handles configuration files. Albarran compares the new reporting tools to a lighter version of the forward-deployed engineering teams that AI labs are embedding inside customer organizations — hands-on help that makes leaving even harder.

KEY FACTS

SPENDING GROWTH ON ANTHROPIC-BASED CODING TOOLS SINCE START OF YEAR

10x

SHARE OF PROPOSED CODE CHANGES FROM ANTHROPIC-BASED TOOLS

nearly a third

SHARE OF ENGINEERING COST NOW SPENT ON THOSE TOOLS

2%

DEVELOPERS IN JELLYFISH SURVEY

276,000

Businesses spending on AI now have a way to measure whether that spending pays off, while finance teams get visibility into a cost line that has grown far faster than most budgets anticipated.

[Read the original at TechCrunch →](#)

04

HIGH IMPACT

Simon Willison's Weblog

AI Agents Tripled GitHub's Pull Requests in a Single Year

GitHub's latest Octoverse report shows that AI-assisted coding has tripled the volume of merged code changes on the platform in under a year.

GitHub published its 2026 Octoverse report this week, and the headline numbers are stark. Merged pull requests — proposed code changes that get accepted into a project — tripled in less than a year to 220 million per month. Commits, the individual saved changes that make up a project's history, more than tripled to 7.38 billion per month, with most of that increase concentrated in the second half of the year. New repositories doubled to 30.3 million per month.

The report attributes the surge to what it calls agentic engineering: software that can plan and execute coding tasks with limited human direction. For businesses, this changes the economics of software development. Teams that adopt these tools can ship more code with the same headcount, but they also inherit new risks around code review, security and maintenance, since the volume of changes now far exceeds what human reviewers can realistically inspect.

The report also breaks down activity by industry and, for the first time, tracks repositories that mention interoperability or reverse engineering — projects aimed at making systems work together or replicating how existing software behaves. GitHub used language models to sort public repositories into industry sectors based on their names, descriptions and documentation, and to identify users' countries while adjusting for VPN use.

What is genuinely new here is the pace. GitHub has reported steady growth for years, but a tripling of merged changes within twelve months is a step change, not a trend line. The most striking detail is that most of the commit growth arrived in the second half of 2026, suggesting the acceleration is still building rather than plateauing.

KEY FACTS

MERGED PULL REQUESTS PER MONTH
220 million

NEW REPOSITORIES PER MONTH
30.3 million

GROWTH IN COMMITS
More than tripled

COMMITTS PER MONTH
7.38 billion

GROWTH IN MERGED PULL REQUESTS
Tripled in less than a year

WHY IT MATTERS

For anyone running or funding a software team, the report signals that AI-generated code is now the default volume driver on the world's largest code platform — which raises the stakes for code review, security and maintenance practices that were designed for a much slower pace of change.

Google Opens Gemini and Search to Any Developer's App

Google is letting outside developers submit their own apps to run inside Gemini and Google Search, turning its AI assistant into a distribution channel for third-party services.

Google announced a public preview of its Connected Apps platform, which lets any business or developer build and submit an app that users can access through Google's AI products. The company already offers more than 50 connected apps in Gemini from partners including Adobe, Experian, Localiza and Peloton, alongside its own services such as YouTube Music and Google Home. Those integrations will now also appear in Google Search over the next few weeks, where today only a handful of partners like Instacart, Canva and Yahoo Fantasy Football are available.

The move matters because it gives companies a new way to reach customers who are increasingly asking an AI assistant to do things rather than opening individual apps. A developer integrates with a cloud-hosted server using Google's AppFunctions APIs, registers and tests the app through the Google Connected Console, then chooses whether it appears in Search, Gemini, or both. Once live, the developer can track traffic, latency and error rates, and keep control over brand elements such as verified logos and authenticated accounts. Google frames this as an early step toward the agentic web, comparing it to how SEO let businesses influence search rankings.

The most striking detail is the scale of the bet Google is making on AI as a new front door to the internet. Meditation app Calm is named as a new partner that will launch on both Gemini and Search. The platform is currently in beta for U.S. developers who sign up for a waitlist, and Google did not say when it would open more broadly.

For businesses, the practical question is whether being present inside an AI assistant becomes as important as ranking well in a search engine. The platform is free to join in preview, but the strategic cost is real: companies that ignore it may find their services harder to reach as consumer behavior shifts toward asking an assistant rather than opening an app.

KEY FACTS

CONNECTED APPS ALREADY IN GEMINI

Over 50

AVAILABILITY

Public preview, U.S. developers, waitlist

NAMED PARTNERS

Adobe, Experian, Localiza, Peloton, Calm

SEARCH ROLLOUT

Next few weeks

WHY IT MATTERS

Businesses that want to stay visible as consumers shift from opening apps to asking AI assistants will need to decide whether to build integrations for Gemini and Search now, while the platform is still in preview and the field is uncrowded.

[Read the original at TechCrunch →](#)

Nvidia Open-Sources Transcription You Can Steer With a Prompt

Nvidia has released an open-weight speech model that combines transcription, speaker labelling, translation and audio question-answering in a single system, and it leads the main open-model leaderboard for English transcription.

The model, Nemotron 3.5 Transcribe, is now available on Hugging Face and ModelScope. It handles recordings up to an hour long, labels up to eight speakers, transcribes 40 languages and translates 31 of them to or from English. Developers can steer the output with plain-language instructions—asking for a translation, speaker labels, or a summary of decisions made in a meeting—and can run it on their own hardware or through hosted endpoints from Baseten, DeepInfra and Simplismart.

The most striking claim is efficiency. On Artificial Analysis's speech-to-text leaderboard, Transcribe ties for first among open-weight models at 2.8% word error rate, but does so with roughly 4 billion active parameters versus 24 billion for the model it ties. In plain terms, it gets comparable accuracy while using about a sixth of the computing power per word. On speaker attribution it also leads: the lowest error rate on David AI's 21-language diarization benchmark and first place in VoiceArena's 14-system comparison.

For businesses, the practical angle is customisation. Nvidia has released Granary v2, the data pipeline behind the model, covering 463,062 hours of speech across 48 languages, along with fine-tuning tools. A medical case study shows why this matters: fine-tuning on Indian medical speech cut word error rate on held-out speakers by 45.8% relative to the best prompt-only approach, from 22.44% to 12.17%. The adapted model beat both an open medical-specialist model and a widely used general-purpose ASR system on the same test. The full adaptation took about 14 GPU-hours on one Nvidia GB300.

Quantized versions are coming soon, which Nvidia says will let the model run locally on RTX Spark, DGX Spark or Jetson AGX Thor devices. The current release is built for recorded audio, not live streaming.

KEY FACTS

LANGUAGES TRANSCRIBED

40

MAXIMUM RECORDING LENGTH

1 hour

ENGLISH WORD ERROR RATE (HUGGING FACE LEADERBOARD)

5.33%

FINE-TUNING COMPUTE

14 GPU-hours on one GB300

LANGUAGES TRANSLATED

31 to or from English

MAXIMUM SPEAKERS LABELLED

8

MEDICAL FINE-TUNING ERROR REDUCTION

45.8% relative

WHY IT MATTERS

Businesses that handle meetings, calls, medical dictation or multilingual content can now use an open model that rivals proprietary speech APIs on accuracy while running on their own infrastructure—and can fine-tune it to their domain vocabulary without paying per-minute fees.

[Read the original at huggingface.co →](#)

AI Applications & Industry

5 stories ● 1 high ● 4 medium

07 MEDIUM IMPACT TechCrunch

Microsoft Booked \$10 Billion Last Quarter From OpenAI Alone

Microsoft reported \$9.9 billion in revenue from OpenAI for the quarter ended September 30, 2026, the first time it has disclosed the figure.

The disclosure, buried in Microsoft's quarterly filing, shows that OpenAI-related revenue reached \$9.9 billion in the first quarter of Microsoft's fiscal 2027 — nearly 11% of the company's \$92.3 billion in total quarterly revenue. For the full 2026 fiscal year, OpenAI contributed \$24.1 billion. At the latest quarterly pace, that would run to about \$40 billion annually.

Microsoft also reported \$7.2 billion in outstanding receivables from OpenAI at the end of September, up from \$6 billion three months earlier. That figure represents money OpenAI still owes for past transactions, not new revenue. It is unclear how much of the \$9.9 billion came from revenue-sharing payments versus Azure cloud services; Microsoft did not break that out.

The numbers underscore how financially entangled the two companies remain even as their once-exclusive partnership has loosened. OpenAI reportedly shares 20% of its revenue with Microsoft under a deal in which Microsoft invested more than \$13 billion in the startup, though neither company has publicly confirmed that percentage. Under an April 2026 agreement, OpenAI must continue sharing revenue with Microsoft through 2030, subject to a cap, even as it is now free to work with competing cloud providers. In October 2025, Microsoft gave up its exclusive hold on OpenAI in exchange for a commitment to buy an additional \$250 billion in Azure services.

Microsoft's ownership stake in OpenAI has also slipped, from about 25% in June to 24% in September, continuing a gradual decline from roughly 27% after the October 2025 recapitalization. Microsoft did not say what caused the dilution.

KEY FACTS	
QUARTERLY REVENUE FROM OPENAI \$9.9bn	SHARE OF MICROSOFT'S QUARTERLY REVENUE Nearly 11%
FULL-YEAR 2026 REVENUE FROM OPENAI \$24.1bn	OUTSTANDING RECEIVABLES FROM OPENAI \$7.2bn
MICROSOFT'S OWNERSHIP STAKE IN OPENAI 24%	AZURE COMMITMENT UNDER OCTOBER 2025 DEAL \$250bn

WHY IT MATTERS

For investors and cloud customers, the numbers show Microsoft's AI bet is paying off in hard revenue — but also that the two companies remain financially bound even as they compete for the same buyers.

Microsoft Puts a Dollar Figure on Azure for the First Time


Microsoft disclosed Azure's quarterly revenue in dollars for the first time, reporting \$33.2 billion for the July-September quarter.

The disclosure lets investors compare Azure directly with Amazon Web Services and Google Cloud, whose parent companies already report cloud revenue in dollars every quarter. Alphabet, which also reported earnings on Wednesday, said Google Cloud brought in \$30.7 billion in the same period, putting Azure slightly ahead. But Google Cloud, which includes Workspace subscriptions, roughly doubled its revenue from a year earlier, growing about twice as fast as Azure, according to TechCrunch's calculations.

Azure is now Microsoft's largest and fastest-growing product line by the company's own figures, accounting for roughly 36% of Microsoft's \$92.3 billion in quarterly revenue. The \$33.2 billion came in about \$850 million above the \$32.3 billion analysts expected, and the 48% growth rate topped the 44% to 45% that Microsoft CFO Amy Hood forecast on the company's July earnings call.

Azure sits inside Agents and Infra, one of two new business segments Microsoft introduced this quarter, replacing its previous three-segment structure. The new segment, which groups Azure with Microsoft 365's cloud business, brought in \$76.8 billion, up 25% from a year earlier.

Amazon reports earnings on Thursday. Analysts expect AWS to grow as well, and it remains the largest cloud provider. But with all three now reporting actual dollars, the competitive picture is clearer than it has been.



KEY FACTS	
AZURE QUARTERLY REVENUE	AZURE YEAR-OVER-YEAR GROWTH
\$33.2bn	48%
GOOGLE CLOUD QUARTERLY REVENUE	MICROSOFT TOTAL QUARTERLY REVENUE
\$30.7bn	\$92.3bn
AZURE SHARE OF MICROSOFT REVENUE	ANALYST EXPECTATION FOR AZURE
36%	\$32.3bn

WHY IT MATTERS

Investors and business buyers can now compare the three major cloud providers on a like-for-like revenue basis, making it easier to judge market share shifts and pricing power in the infrastructure that underpins most AI workloads.

Google Rolls Its AI Search Ads Out Worldwide For the First Time

Google said on Wednesday it is expanding ads inside AI-generated search answers to international markets, as growth in its core Search advertising business continued to decelerate.

The announcement came during Alphabet's third-quarter earnings call, where the company reported \$127.6 billion in revenue, up 25% year over year. But the headline number masked a familiar problem: Google Search revenue grew 12% year over year, down from 17% in the second quarter and 19% in the first. Alphabet shares fell 8% after the results.

Philipp Schindler, Google's chief business officer, said the company is "now scaling high-quality ads within AI overviews in AI mode to more users and queries," including taking the ad format global for the first time and adding multi-turn ads in the U.S. Early data suggests the ads are performing: click-through rates for sponsored product listings in AI Mode are already comparable to those for organic listings, Schindler said, and user research indicates people find them relevant.

The push builds on ad formats Google announced in May, including conversational discovery ads that respond to users' questions and sponsored recommendations embedded in AI-generated answers. The stakes are clear: Google has bet heavily that AI will change how people search, and now it needs those changes to pay. CEO Sundar Pichai said Search queries grew at their fastest year-over-year rate in five years, and that queries in AI Mode are three times longer on average than traditional searches. One in five AI Mode searches globally involves voice, images, or documents.

For advertisers, the expansion means a new inventory pool inside a search experience that is still finding its commercial footing. For Google, it is a test of whether AI-assisted search can generate ad revenue at the same rate as the classic results page — a question the slowing Search growth makes more urgent.

KEY FACTS

ALPHABET Q3 REVENUE

\$127.6bn

PRIOR QUARTER SEARCH GROWTH

17%

AI MODE QUERY LENGTH VS TRADITIONAL

3x longer

SEARCH REVENUE GROWTH

12% year over year

ALPHABET SHARE MOVE AFTER RESULTS

Down 8%

AI MODE SEARCHES USING VOICE, IMAGES OR DOCUMENTS

1 in 5

WHY IT MATTERS

Advertisers get a new, global placement inside AI-generated search answers, while Google's slowing Search ad growth makes the success of these formats central to its revenue outlook.

[Read the original at TechCrunch →](#)

Chrome Wants to Read Your Gmail to Fill in Forms

Google is expanding Chrome's autofill to draw on Gmail, letting the browser suggest details like flight confirmations, membership IDs and tracking numbers when users fill out online forms.

The update, announced Wednesday, means Chrome can reference a user's email to surface things they might need to enter on a website: a flight confirmation number, a membership ID, an order number, an address, a phone number, or a credit card number. It is an extension of Google's Personal Intelligence, which already lets users connect Google Apps to the AI assistant Gemini so it can retrieve specific details and answer more complex questions.

For Google, the move is a retention play as much as a convenience feature. Making Chrome easier to use gives users another reason to stay in Google's browser and its broader AI ecosystem, at a time when rivals are pushing their own assistants and productivity tools. The feature echoes one of Apple's most appreciated touches: automatically detecting verification codes sent by text and displaying them above the keyboard so users can tap once to enter them.

Users can choose whether to enable the enhanced autofill from Chrome settings, and sensitive data is encrypted, Google says. The feature is rolling out to U.S. users on desktop, as well as in the iOS and Android apps for Gemini in Chrome. The genuinely new part is the Gmail connection; autofill itself has existed for years, but tying it to a user's inbox brings Google's data sources closer together in a way that could make its browser stickier.

KEY FACTS

AVAILABILITY

U.S. users on desktop, iOS and Android apps for Gemini in Chrome

DATA SOURCE

Gmail

USER CONTROL

Users can enable or disable enhanced autofill from Chrome settings

SECURITY

Sensitive data is encrypted

WHY IT MATTERS

For everyday users, it removes friction from filling out forms. For Google, it deepens the integration between Gmail, Chrome and Gemini, giving people another reason to keep using its products rather than switching to a competitor's browser or assistant.

[Read the original at TechCrunch →](#)

October Was the Month Vibe-Coded Games Became Actually Fun

AI models have crossed a threshold where they can design and build computer games that are actually fun to play, not just visually convincing.

Simon Willison, a developer who has been experimenting with AI-assisted coding for years, reports that Claude Opus 5.5 can now produce genuinely enjoyable games from simple text prompts, with minimal human input. In recent weeks he built four browser-based games—a tree-spotting puzzle, a fantasy role-playing game with a full story, a multiplayer real-time strategy game, and a pumpkin-farming simulator—none of which required more than four hours of his own work. The key shift, he says, is not graphics or code quality but game design: the model now understands what makes a game fun, responding to the instruction "Make it a LOT more fun" with meaningful improvements.

The most striking example is Battle Drones, a Command & Conquer-style strategy game that now supports two-to-four-player multiplayer, medals, replays, seven maps, and a mode where different AI models play against each other. Willison released it under the Apache 2.0 open-source licence, and within a day another developer had forked it into a steampunk 1900s France version with new models, music and an animated intro. He notes that OpenAI's GPT-6 Astra can build games but "its sense of fun doesn't match that of Opus 5.5," though he expects OpenAI to catch up by the end of the year.

For the games industry, this raises uncomfortable questions. Building a computer game has historically been one of the most demanding projects in software, combining art, music, writing, programming and design over months or years. Now all of those elements can be outsourced to a model in an afternoon. Willison expects a split: some players will refuse AI-built games, while others will embrace the wider access to game creation. He himself won't enter his projects in traditional game jams, saying he doesn't feel he has earned that right.

Willison's advice to anyone curious is to start with a prompt like "Build a mobile-friendly three.js game where..." in a standard chat interface—no coding agent required. The barrier to entry has fallen to the point where a playable multiplayer strategy game can be prompted from a phone.

KEY FACTS

GAMES BUILT IN WEEKS

4

BATTLE DRONES LICENCE

Apache 2.0, released 20 October 2026

EMBERWICK PLAYTHROUGH TIME

About 15 minutes

MAXIMUM PERSONAL EFFORT PER GAME

4 hours

BATTLE DRONES MULTIPLAYER

2-4 players

WHY IT MATTERS

Game development, long one of the most expensive and skill-intensive forms of software, is becoming accessible to anyone who can write a prompt—which will reshape who makes games, how much they cost to produce, and what players can expect to see released.

[Read the original at Simon Willison's Weblog](#) →

Sources

- 01 Google's Second Frontier Model in Two Weeks Is Built to Hunt Bugs
<https://deepmind.google/blog/gemini-4-carbon-our-next-frontier-model/>
- 02 Researchers Retune AI's Attention to Block Hidden Instructions in Webpages
<https://arxiv.org/abs/2610.37657>
- 03 Anthropic Shows Companies Exactly What Their AI Spending Bought
<https://techcrunch.com/2026/10/28/anthropic-just-launched-tools-to-show-companies-what-their-ai-spending-is-buying/>
- 04 AI Agents Tripled GitHub's Pull Requests in a Single Year
<https://simonwillison.net/2026/Oct/28/github-octoverse-2026/>
- 05 Google Opens Gemini and Search to Any Developer's App
<https://techcrunch.com/2026/10/28/as-ai-agents-reshape-the-web-google-opens-gemini-and-search-to-more-apps/>
- 06 Nvidia Open-Sources Transcription You Can Steer With a Prompt
<https://huggingface.co/blog/nvidia/nemotron-transcribe-granary>
- 07 Microsoft Booked \$10 Billion Last Quarter From OpenAI Alone
<https://techcrunch.com/2026/10/28/microsoft-recorded-10-billion-in-revenue-from-openai/>
- 08 Microsoft Puts a Dollar Figure on Azure for the First Time
<https://techcrunch.com/2026/10/28/microsoft-just-shared-azures-actual-revenue-in-dollars-33-2-billion/>
- 09 Google Rolls Its AI Search Ads Out Worldwide For the First Time
<https://techcrunch.com/2026/10/28/google-expands-ai-search-ads-globally-as-core-ad-growth-slows/>
- 10 Chrome Wants to Read Your Gmail to Fill in Forms
<https://techcrunch.com/2026/10/28/chromes-ai-powered-autofill-can-now-dig-through-your-gmail-to-fill-out-forms/>
- 11 October Was the Month Vibe-Coded Games Became Actually Fun
<https://simonwillison.net/2026/Oct/28/vibe-game-development/>

About this document. Every story in the 29 October 2026 New Horizon AI Digest, reported at length. Each entry is written from the publisher's own article text; where a source could not be retrieved the entry is explicitly marked and kept short rather than padded.

Images and licensing. Figures are used only where the source licence permits redistribution, and are credited in the caption. Publisher artwork is not reproduced. All titles link to the original publication.

Vocabulary. Common AI terms are defined at new-horizon.tech/glossary/.

Generated 2026-10-29T06:35:46Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.