



**Jerusalem
AI Center**

JAC AI Weekly

VOLUME 2 · CATCH-UP EDITION · SEPTEMBER 29-OCTOBER 4, 2026

Putting the human back in AI.

STORY OF THE WEEK

Almost half of people thought this AI on a video call was human

Tavus introduced Griffin, a video AI that talks live and picks up on pauses and expressions. In a live study, 48% of participants thought they were talking to a real person. Earlier systems, including Tavus's own, topped out at 2%.

The result is Tavus's own study. Why it matters: live video AI is getting close to natural conversation, which makes clear labeling of what is AI more important. Story 6 inside.

Also this week: Anthropic's faster and cheaper Claude Sonnet 5.5, OpenAI's always-on assistants, and Google's Gemini 4 Argon, released first to security defenders.

Over Sukkot we paused our daily AI updates, so here is one story for each day we missed, in a few plain sentences with a link if you want more. Volume 1 covered September 20 to 28. From here on, the weekly collects the best of the week.



Illustration: AI over Jerusalem

IN THIS ISSUE

1. Anthropic's new everyday model is faster and cheaper
2. OpenAI introduces assistants that keep working
3. Google announces Gemini 4 Argon
4. Cloudflare releases free decision models
5. A small free AI writes cited science reports
6. Almost half thought this video AI was human
7. Listen: Kol, Episode 1
8. Try this week
9. Glossary

TOOLS Tue, Sept 29

Anthropic's new everyday model is faster and cheaper

Anthropic released Claude Sonnet 5.5, the mid-size model most people use for daily work like fixing code and creating documents, slides and spreadsheets. Anthropic says it runs more than 30% faster and costs up to 30% less for most work than the version before it, at the same per-token price. Benchmark scores are Anthropic's own.

Why it matters: Faster and cheaper everyday AI lowers the cost of using it across a business.

In plain English: Anthropic released a faster, cheaper version of the AI model most people use day to day.

Read the story at [Anthropic](https://www.anthropic.com/claude-sonnet-5-5): <https://www.anthropic.com/claude-sonnet-5-5>

INDUSTRY Wed, Sept 30

OpenAI introduces assistants that keep working while you do other things

At its developer conference, OpenAI announced "dots," always-on AI agents that run on their own cloud computers and keep working toward a goal over time. They are available to Pro and Business Premium users in eligible markets. OpenAI also launched GPT-6.1 Sol, which it says nearly matches its top model, GPT-6 Astra, on coding and computer tasks at one-fifth of the token price.

Why it matters: Cheaper, steadier AI is what lets more businesses use it for real work. Prices and performance are OpenAI's own figures.

In plain English: OpenAI announced AI assistants that keep working in the background, and a cheaper model that it says performs close to its best one.

Read the story at [Decrypt](https://decrypt.co/379584/openai-ai-agents-computers-devday-2026-everything-announced): <https://decrypt.co/379584/openai-ai-agents-computers-devday-2026-everything-announced>

TOOLS Thu, Oct 1

Google announces Gemini 4 Argon, built for long, complex jobs

Google announced Gemini 4 Argon, designed for long, multi-step work in software, legal and finance work, and cybersecurity defense. For now it is rolling out to a group of trusted cyber defenders first, so they can use it to find and fix software weaknesses before wider access opens. Google says it will expand access in phases.

Why it matters: A defender-first release shows how the leading labs are staging their strongest models.

In plain English: Google has a new model for long, complex work. Security defenders get it first, and wider access follows in phases.

Read the story at [Google](https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/): <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/>

TOOLS Fri, Oct 2

Cloudflare releases free "decision" models anyone can download

Cloudflare, which runs a large part of the internet's infrastructure, released Clef and Clef-flash, open models that answer structured questions such as yes or no, multiple choice and rankings, rather than writing text. They work like TypeSafe's Jev, covered in Volume 1, and can run on a company's own machines. Performance comparisons are Cloudflare's own.

Why it matters: Decision-style AI is quickly becoming a category, and open versions give companies a way to run it themselves.

In plain English: Cloudflare released free AI models that give structured answers, such as yes or no. Companies can run them on their own machines.

Read the story at [The Register](https://www.theregister.com/ai-and-ml/2026/10/01/cloudflare-tries-to-outplay-jev-with-open-weight-clef-models/5300649): <https://www.theregister.com/ai-and-ml/2026/10/01/cloudflare-tries-to-outplay-jev-with-open-weight-clef-models/5300649>

RESEARCH Sat, Oct 3

A small free AI writes cited science reports in under a minute

The Allen Institute for AI (Ai2) open-sourced AstaBrief, a compact model that turns a research question and a set of papers into a cited report. Ai2 says its fast mode averages 51 seconds per report, against about 3 minutes for its Claude-powered mode. Researchers can download it and run it on their own computers.

Why it matters: Open models that run locally matter for researchers whose questions are sensitive.

In plain English: A small, free AI can turn research papers into a cited report in about a minute. Researchers can run it on their own computers.

Read the story at [Ai2](https://allenai.org/blog/astabrief): <https://allenai.org/blog/astabrief>

TOOLS Sun, Oct 4

Almost half of people thought this AI on a video call was human

Tavus introduced Griffin, a video AI that talks live and picks up on pauses and expressions. In a live study, 48% of participants thought they were talking to a real person; earlier systems, including Tavus's own, topped out at 2%. The result is Tavus's own study.

Why it matters: Live video AI is getting close to natural conversation, which makes clear labeling of what is AI more important.

In plain English: A video AI was convincing enough that nearly half of test participants thought it was a person. This is Tavus's own study.

Read the story at [Tavus](https://www.tavus.io/griffin): <https://www.tavus.io/griffin>

LISTEN: Instinctual Musings, Kol Episode 1

The Instinct Freakout: What's True, and What I'd Tell a Friend. People are nervous about Instinct. Bryant uses it every day. What happened, what's true, and what he'd tell a friend. The episode opens in his real voice, then hands off to Kol, his AI voice.

<https://instinctualmusings.substack.com/p/this-is-my-real-voice-for-now>

TRY THIS WEEK

1. Try an everyday AI model on one real task, such as a summary or a first draft of a slide, and note how much time it saves (story 1).
2. If a vendor offers an AI assistant that keeps working in the background, ask what it can access and how you will review what it did (story 2).
3. When you meet an AI on a video call or chat, ask the provider to label it clearly as AI (story 6).

GLOSSARY

AI model

The software engine behind an AI product, such as the one behind ChatGPT.

Open-weight / open model

An AI model that anyone can download and run on their own computers.

AI agent

AI that takes actions toward a goal, not just answers questions.

Benchmark

A standard test used to compare AI systems. Scores published by the company that built the model are its own claims.

Decision model

An AI that returns a structured answer, such as approve or decline, with a confidence score, instead of writing text.