



JAC AI Weekly

VOLUME 1 · ISSUE 1 · SEPTEMBER 20-28, 2026

Putting the human back in AI.

STORY OF THE WEEK

Nvidia opens an AI research lab in Israel

Nvidia has started a new research group in Israel that will use AI to speed up medicine and drug development, according to Israeli news reports on September 23. It sits inside Nvidia's Israeli development center and is led by Dr. Nati Daniel. The team already has several researchers and plans to hire dozens more over the next year in Tel Aviv.

One early project is an AI model of the human genome, built with Sheba Medical Center in Israel and Mount Sinai in New York. It aims to help researchers understand how genetic differences relate to disease risk and response to treatment. Israel is already Nvidia's largest research base outside the United States. Nvidia has not given a timeline for first results.

Also this week: Irregular, an Israeli AI security lab that tests AI models for the big labs, is reported by Globes to be raising funds at a \$1.5 billion valuation. Why it matters: Long-term research labs show where companies expect talent and discoveries to come from. Story 5 inside.

This is the first JAC AI Weekly. It collects the AI stories from September 20 to 28, in plain language, each with a link if you want to read more. Volume 2, the catch-up edition, covers September 29 to October 4.



Illustration: AI over Jerusalem

IN THIS ISSUE

1. Jev: AI that gives the decision, not an essay
2. Lawsuit: AI labs accused of colluding to slow down
3. Xiaomi releases the top-rated downloadable AI
4. Amazon blocks Meta's AI shopping assistant
5. Nvidia opens an AI research lab in Israel
6. AI CEOs ask the UN for cooperation
7. Bill Gates calls for clear, shared AI rules
8. Hospital pilots AI to cut nurses' paperwork
9. Nvidia tools set clear limits for AI agents
10. Meta launches an enterprise AI platform
11. Listen: Kol, Episode 1
12. Try this week
13. Glossary

TOOLS Sun, Sept 20

Jev: an AI that gives you the decision, not an essay

Sunday is tools day. One launch worth your time:

Jev - a new kind of AI model from a co-inventor of ChatGPT.

It doesn't chat. You give it a situation, it gives back a decision with a confidence score. No essay, just the answer - and the claim is 20-200x faster, 40-400x cheaper than the chat models at it.

Most business AI isn't conversation, it's decisions. Approve or decline, route here or there, buy or wait. A model built just for that is a different animal - and developers were all over it last week.

In plain English: Instead of chatting, this kind of AI gives a direct answer, such as approve or decline, along with how confident it is.

Read the story at [TechCrunch](https://techcrunch.com/2026/09/18/a-new-kind-of-ai-model-from-a-chatgpt-inventor-is-thrilling-developers/): <https://techcrunch.com/2026/09/18/a-new-kind-of-ai-model-from-a-chatgpt-inventor-is-thrilling-developers/>

POLICY Mon, Sept 21

Lawsuit claims the big AI companies broke competition law by agreeing to slow down

The big AI labs just got sued - for going too **slow**.

A new lawsuit claims Anthropic, OpenAI, Google and SpaceXAI illegally agreed to coordinate an AI slowdown, violating antitrust law.

For a year, the labs have talked openly about pacing, shared safety tests, even asked governments to help them agree on speed limits. The suit says that's not safety - it's collusion: competitors agreeing not to compete.

Why it matters: the case decides what "responsible AI" is allowed to look like. If safety coordination counts as a cartel, every lab goes back to racing alone. If it's protected, expect more joint pace-setting - and slower, more deliberate tools reaching your business.

In plain English: A lawsuit claims four AI companies broke competition law by coordinating on how fast to move. It is a claim in court, not a ruling.

Read the story at [Fortune](https://fortune.com/2026/09/19/lawsuit-anthropic-openai-spacexai-google-antitrust-laws-ai-slowdown-subscription-value/): <https://fortune.com/2026/09/19/lawsuit-anthropic-openai-spacexai-google-antitrust-laws-ai-slowdown-subscription-value/>

RESEARCH Tue, Sept 22

Xiaomi releases the top-rated AI that anyone can download and run

Tuesday is research day. One release worth knowing about:

Xiaomi's new MiMo-V2.6-Pro just took the #1 spot on the global open-weights rankings - ahead of every other model that anyone can download, run, and build on without paying a usage fee or sending data to an outside company.

Two things make it stand out. First, performance: it leads the pack not just on paper but as a model businesses can actually run privately, on their own equipment.

Second, Xiaomi published the full diary of its \$3.5M training run - the crashes and restarts included - the most transparent look yet at how a top model gets made.

If your business uses AI, the takeaway is simple: the gap between the AI you rent and the AI you can own keeps closing, and this week it closed a little more.

In plain English: A free model that you can run on your own equipment ranks first on one global open-weights ranking. That gives businesses another option besides paying a provider and sending data out.

Read the story at [Unite.AI](https://www.unite.ai/xiaomis-new-flagship-model-leads-open-weight-rankings-with-a-score-of-46/): <https://www.unite.ai/xiaomis-new-flagship-model-leads-open-weight-rankings-with-a-score-of-46/>

INDUSTRY Wed, Sept 23

Amazon blocks Meta's AI shopping assistant from its store

Wednesday is a practical day. Today's article is about not building and disclosure:

Amazon just threw Meta's AI shopper out of its store. Meta's Muse agent walked in holding customers' Amazon passwords, acting like a person, and never said it was a bot.

That's now a fast way off the internet: the web's oldest rule says bots name themselves, California and the EU made disclosure the law.

Keep your eye out for updated "prove your human" as Cloudflare is rolling out cryptographic proof so an agent can't fake its identity.

Figuring out the blend with Agent hand off and honest identity with plain disclosure is about to become an important conversation.

In plain English: Amazon turned away an AI shopping assistant that, Amazon says, did not identify itself. Saying clearly that it is an AI is becoming the expectation for tools that act for you.

Read the story at [The Register](https://www.theregister.com/ai-and-ml/2026/09/21/amazon-shows-metas-muse-ai-shopping-agent-the-door/5297777): <https://www.theregister.com/ai-and-ml/2026/09/21/amazon-shows-metas-muse-ai-shopping-agent-the-door/5297777>

ISRAEL Thu, Sept 24

Nvidia opens an AI research lab in Israel

Two stories about Israel and AI today.

1) **Nvidia is opening an AI research lab in Israel**, focused on finding new medicines faster. Nvidia powers basically the entire AI wave. They could put a lab anywhere on earth - and they decided the talent for AI-meets-medicine is sitting in this country.

2) The company that led the testing for Google's AI escaping its safety test is Israeli. It's called **Irregular**, and it's now **raising money at a \$1.5 billion valuation**.

In plain English: Nvidia is putting a medical-research team in Israel, and an Israeli company that tests AI safety is raising money at a \$1.5 billion valuation, per Globes.

Read the story at [Nvidia lab \(VINnews\)](https://vinnews.com/2026/09/23/nvidia-opens-israel-ai-lab-aimed-at-faster-drug-and-medical-discoveries/): <https://vinnews.com/2026/09/23/nvidia-opens-israel-ai-lab-aimed-at-faster-drug-and-medical-discoveries/>

Read the story at [Irregular \(Globes\)](https://en.globes.co.il/en/article-irregular-raising-funds-at-15b-valuation-1001557345): <https://en.globes.co.il/en/article-irregular-raising-funds-at-15b-valuation-1001557345>

POLICY Fri, Sept 25

OpenAI and Anthropic CEOs ask the UN for global AI cooperation

Here is a values angle on the week's AI news.

The two men building the world's most powerful AI stood in front of the UN Security Council - and asked to be regulated.

The strongest players in the room are saying "we need to answer to someone".

The Jerusalem AI Center wants to be one of those people in that room, utilizing Torah values to help shape what our AI future truly looks like.

In plain English: The heads of OpenAI and Anthropic addressed the UN Security Council and called for global cooperation on AI.

Read the story at [ABC News](https://abcnews.com/Politics/openai-anthropic-ceos-call-global-cooperation-ai-crossroads/story?id=136697471): <https://abcnews.com/Politics/openai-anthropic-ceos-call-global-cooperation-ai-crossroads/story?id=136697471>

POLICY Sun, Sept 27

Bill Gates calls for clear, shared rules for AI

Microsoft co-founder Bill Gates said on NBC's Meet the Press that he would like to see clear safeguards and monitoring for AI, set through law, and that Congress should pass legislation. "No one thinks self-regulation is enough," he said. He added that he does not think safeguards would dramatically hinder the industry.

Others, including President Trump, say new restrictions could let China take the lead in AI. Gates said he wants to discuss the issue with the President.

Why it matters: Clear rules can give businesses more certainty about how AI may be used. If your company uses AI, this is a debate worth following.

In plain English: Bill Gates would like the government to set clear AI rules. Not everyone agrees on how much regulation is needed.

Read the story at [Reuters \(via CBC News\)](https://www.cbc.ca/news/world/bill-gates-ai-safeguards-9.7360347): <https://www.cbc.ca/news/world/bill-gates-ai-safeguards-9.7360347>

HEALTH Mon, Sept 28

A New Jersey hospital pilots an AI assistant to cut nurses' paperwork

Cooper Norcross Health announced it is the first hospital in New Jersey to pilot Microsoft Dragon Copilot for nurses. It describes the tool as the first commercially available AI clinical assistant designed for nursing.

Cooper says documentation and administrative tasks take more than a quarter of each nurse's shift. With the assistant, a nurse starts a recording during routine care, with the patient's consent, and the tool drafts the chart entries from what is said. The nurse reviews, edits and approves every entry before it is filed.

Cooper is evaluating less typing and more time at the bedside during the pilot. No results have been reported yet.

Why it matters: This is a clear example of AI handling paperwork while a person checks and approves the result. Many businesses have similar paperwork.

In plain English: A hospital is testing an AI that drafts nurses' notes while they care for patients. The nurse checks and approves every entry.

Read the story at [Cooper Norcross Health](https://blogs.cooperhealth.org/news/2026/09/cooper-norcross-health-first-in-new-jersey-to-bring-ambient-ai-to-the-nursing-bedside-with-microsoft-dragon-copilot/): <https://blogs.cooperhealth.org/news/2026/09/cooper-norcross-health-first-in-new-jersey-to-bring-ambient-ai-to-the-nursing-bedside-with-microsoft-dragon-copilot/>

TOOLS Mon, Sept 28

Nvidia launches tools to help companies set clear limits for AI agents

Nvidia introduced the Nvidia Open Agent Safety Platform, a toolkit that lets organizations decide exactly what an AI agent may access and watch what it does.

It combines OpenShell, open-source software that controls what an agent is allowed to access, with Sentry, a monitoring system that runs on its own Nvidia chip so it has an independent view of the agent's activity. Nvidia says Sentry can quarantine an agent that moves outside its limits within milliseconds. Anthropic, Arm, Microsoft, Oracle and SpaceX are listed as supporters. These are Nvidia's own claims.

Why it matters: Clear limits and independent monitoring are what let organizations start using AI agents with confidence.

In plain English: Nvidia is offering a way to set clear limits on what AI assistants can access, and to monitor their activity.

Read the story at [TechCrunch](https://techcrunch.com/2026/09/28/nvidia-launches-new-platform-for-reining-in-rogue-ai-agents/): <https://techcrunch.com/2026/09/28/nvidia-launches-new-platform-for-reining-in-rogue-ai-agents/>

INDUSTRY Mon, Sept 28

Meta launches an enterprise AI platform for businesses

Meta announced Meta Enterprise Platform, a new business line that brings its AI to other companies. It starts with the Muse agent (the one Amazon blocked in story 4), Meta Business Agent, Muse API and Muse Code.

Chirantan "CJ" Desai, who was CEO of MongoDB, joins Meta as Chief Enterprise Platform Officer and will report to Mark Zuckerberg. Meta says security and privacy are built into its enterprise products from the start.

Why it matters: Another large technology company is now selling AI directly to businesses, which gives companies one more option to compare.

In plain English: Meta is now selling its AI tools to other companies, not only using them itself.

Read the story at [Meta](https://about.fb.com/news/2026/09/launching-meta-enterprise-platform/): <https://about.fb.com/news/2026/09/launching-meta-enterprise-platform/>

Read the story at [TechCrunch](https://techcrunch.com/2026/09/28/meta-launches-enterprise-ai-platform-hires-mongodb-ceo-to-lead-new-initiative/): <https://techcrunch.com/2026/09/28/meta-launches-enterprise-ai-platform-hires-mongodb-ceo-to-lead-new-initiative/>

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. Pick one routine yes-or-no decision in your work, such as approving a refund or routing a request, and write down what information goes into it. This is the kind of task decision-style AI like Jev is built for (story 1).
2. Ask your IT provider whether any AI tool your team uses sends company data to an outside company, and whether an open model on your own equipment is an option (story 3).
3. If you use an AI assistant that shops or books for you, check whether the vendor says it identifies itself as an AI to the sites it visits (story 4).
4. Look for paperwork in your business that a person reviews anyway, such as notes, reports or forms. These are good first places to try an AI that drafts while a person approves (story 8).
5. Before you let an AI assistant act for you, decide what it may open or change and who reviews its work (story 9).

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.

Decision model

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

Confidence score

A number showing how sure the AI is about its answer.

Training run

The long computing job that teaches an AI model.

Cryptographic proof

A mathematical check that lets a website confirm who or what is sending a request.

Valuation

What investors think a company is worth.

Antitrust

Competition law, which limits companies from agreeing with each other in ways that reduce competition.

Self-regulation

When an industry sets and enforces its own rules instead of the government.

Pilot

A small trial run before wider use.