Episode Book

EPISODE #6524

Why Israel Never Finished Its Sovereign Hebrew AI

PUBLISHED	10 November 2026 at 12:02 UTC
DURATION	27:46
LISTEN	myweirdprompts.com/episode/israel-sovereign-hebrew-ai-model-11932/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.26
SOURCES	15 cited by the research (listed at the end)
BOOK GENERATED	11 November 2026 at 06:49 UTC

SYNOPSIS

Israel spent years, four national plans, and hundreds of millions of shekels trying to build a sovereign Hebrew and Arabic language model. Tracing the initiative from the TELEM Forum's market-failure findings to MAF'AT's unreleased 20-billion-word Semitic model and a planned repository called GitHeb, this account unpacks how a project meant to rescue a low-resource language stalled, why funding sputtered across missing supercomputers, and what the 2026 AI Directorate's revised sovereignty agenda means today.

TRANSCRIPT



Corn

Eight years, four national plans, several hundred million shekels, and no model. That's the ledger. Daniel wants us to walk through Israel's effort to build a sovereign Hebrew language model — the real history, who was actually building it, where the money went, why it got shelved once ChatGPT arrived, and what the 2026 plan is promising now. He's also asking the harder question underneath all of it: when frontier models already speak Hebrew, what is a government-built national language model actually for?



Herman

And it's a good time to be asking, because the question just got re-opened. The new National AI Directorate published its plan in August, and everyone's arguing about whether it commits Israel to a Hebrew model. I've read the plan document. It commits Israel to examining.



Corn

Examining.



Herman

The operative verb on the language model is "will examine." Three times in one sentence. Israel will examine the need for and implications of training infrastructure adapted to the Israeli context, culture, values and language, and will examine moves that allow Israel to lead in domain-specific models — while, and this is the striking part, "addressing past efforts to establish a model for Hebrew."



Corn

The plan contains a euphemism for its own predecessor.



Herman

Meanwhile the man who wrote it, Erez Askal — brigadier general in the reserves, used to run Unit 9900 — stands on stages and says things like "if it's not your AI, you have no control over your country" and "we need to build our own frontier model." The stage and the document are describing two different countries.



Corn

The document is describing the one that examines.



Herman

So the episode really has two halves. There's the history — a state that kept deciding to build this thing and kept not building it, from 2020 to the Comptroller's report in twenty twenty-four. Then there's the present, which is an argument about whether the whole idea still makes sense. But before the timeline, let's be clear about what the argument was actually about — because it changed. In twenty twenty-two the case for a Hebrew model was capability: Hebrew models were bad. By twenty twenty-six the case is control: even if they're good, they're not ours. Different argument, different weaknesses.



Corn

Start where it started. Was this a tech moonshot?



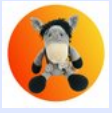
Herman

No, and that's what I love about the origin. It's a committee. February twenty twenty, the TELEM Forum — that's the national research infrastructure forum — appoints a review committee under Dr. Orna Berry, who'd been Chief Scientist and a Google Israel executive. The question they were asked wasn't "how do we build a great Hebrew model." It was "does the market need government intervention at all."



Corn

A market-failure finding.



Herman

Their December report said Hebrew and Arabic natural language processing constituted a market failure and recommended a hundred and eighty million shekels. The reasoning quoted at the time: not enough activity in academia and Israeli industry, in a field of strategic importance.



Corn

I want to defend that finding for a second, because "committee discovers market needs committee" is the oldest genre in government.



Herman

Fair instinct.



Corn

But Nvidia Israel's Ophir Zamir made the same point independently, and he had no committee to feed: "In the end it's an economic business. Companies that build NLP models ultimately need to see how many customers can benefit from this." A few million Hebrew speakers isn't a market. English gets trained because a billion and a half people will use it. Hebrew gets trained if a professor in Ramat Gan cares enough.



Herman

That's the honest steelman, yes. And the person who cared was Prof. Reut Tsarfaty at Bar-Ilan, who laid out the technical case in a way that still holds up as linguistics even though the market conclusion expired. Hebrew has no niqqud, so every word can be read several ways. No capital letters. The word "dan" is both a verb — he judged — and a proper name. Freer word order. The binyanim verb patterns. Her AlephBERT model had a hundred million sentences in training and was still considered small — it hit around eighty percent on named-entity recognition, where industry wanted ninety-five.



Corn

Eighty percent named-entity recognition means one entity in five is wrong.



Herman

Or hallucinated, or missing. If you're a ministry trying to extract names from forms, one in five is not a tool. And the data problem was the wall: her line was that most of the large quantities of Hebrew text are "locked behind walls of privacy or copyright." The state sitting on decades of Hebrew documents it could never license — but it was the one actor that could free them.



Corn

Can you make the "dan" problem concrete? Because I think listeners hear "ambiguous word" and shrug — every language has those.



Herman

Sure. Take a sentence where "dan" appears with no niqqud and no capitals. An English model reading "Dan" capitalized knows instantly it's a name. A Hebrew model reading the same letters has to decide from context alone whether it means he judged, or the tribe, or the person — and Hebrew's freer word order means the context clues arrive in orders the model wasn't built to expect. Multiply that by every word in the language, because the niqqud problem isn't occasional, it's universal. It's not one ambiguous word, it's that all of them are, all the time, and the model has to be right about all of them at once.



Corn

So the errors compound.



Herman

They compound. And the binyanim issue is its own layer — the verb patterns carry meaning English handles with separate words, so a model that mishandles them isn't slightly off, it's reading a different sentence.



Corn

So who picked up the hammer?



Herman

Not who you'd guess from a market-failure framing. The Innovation Authority handled the ecosystem — an NLP consortium with Google, Amazon, Rafael and Harel, budgeted at seven and a half million shekels over three years. Note that figure, we'll come back to it. The model itself went to MAF'AT — the Defense Ministry's R&D directorate. A lieutenant colonel identified only as "A.", head of their AI branch, described the brief: "We took the best minds — experts from academia, our own developers, world-class companies — to create a model composed of twenty billion words in Hebrew. We are going to develop a unique architecture suited to Semitic languages, and this model will be bilingual, for Hebrew and Arabic."



Corn

A unique Semitic architecture.



Herman

Twenty twenty-two was full of people convinced the transformer needed regional variants. And to be fair to the moment, there was a real argument — Arabic's morphology is famously combinatorial, and if you believed Semitic languages broke the assumptions baked into English-first tokenizers, building your own architecture was the intellectually consistent move. It just turned out the generic architecture was better at everything, including the things it was supposedly bad at.



Corn

Where was it going to live?



Herman

Here's the detail I can't get over. They planned to publish the models fully open on a dedicated site called GitHeb. GitHeb.



Corn

G-I-T-H-E-B.



Herman

Like the code repository, with Hebrew substituted. For everyone, they said, "from giant companies like Google and Meta to a student who just finished her degree and wants to start a startup." Three variants — mobile, personal computing, server farm. First model promised by end of twenty twenty-two.



Corn

The student is still waiting.



Herman

She is. And there's a small irony sitting inside the name itself, which is that the actual GitHub is where the world's open models ended up living anyway. If MAF'AT had simply published to the real repository instead of promising a Hebrew-branded clone of it, there would at least be a commit history. Instead the repository exists only as a promise about a repository.



Corn

A fork with no upstream.



Herman

Now the money. Government Decision 212 in August twenty twenty-one approved about five hundred and fifty million shekels for the first phase. February twenty twenty-three, Decision 173 approved commitment authority of up to five hundred million more for the following years. Then November twenty twenty-four, the State Comptroller publishes a special report, and the numbers are grim. Only forty percent of the first-phase budget was implemented — two hundred and twenty million of the five hundred fifty. On high-performance computing, the supercomputer, eleven percent: about thirty million of two hundred seventy approved.



Corn

Eleven.



Herman

Human capital, fifty-five percent. The one line that looks respectable is Hebrew–Arabic NLP itself: seventy-six percent, a hundred thirty-eight million of a hundred eighty. Except — and this is where the whole story turns — that figure includes a thirty-seven million shekel language-model agreement that MAF'AT signed on the thirty-first of December twenty twenty-three.



Corn

ChatGPT launched November of twenty twenty-two.



Herman

November thirtieth, twenty twenty-two. The model agreement was signed thirteen months later. Almost to the day — the Comptroller notes it was signed on the final day of phase one of the national programme.



Corn

Signed on the deadline, for a project whose reason to exist had already shipped.



Herman

That's the misconception worth killing, actually. The received story, the one Globes prints in a single line, is that the Hebrew model was "shelved after millions of shekels had been invested when OpenAI launched ChatGPT." Shelved implies a decision. There's no decision in the record. There's a project whose rationale got overtaken, and then a year of signatures continuing on top of it — first model now expected around mid-twenty twenty-five, which we're past, and there's no model.



Corn

It wasn't killed, it was zombified.



Herman

And the Comptroller's verdict on the whole programme: Israel has no long-term national AI strategy, the government never approved a master plan, and it made do with limited, pulse-funded programmes that were only partially implemented and missed their timetables. The line that got quoted everywhere: "The world is running forward, and we are falling behind."



Corn

One number I want to flag on air, because it doesn't reconcile. Haaretz's Gideon Lev reported that of the two hundred seventy million allocated to Hebrew and Arabic large language models, only about thirty million was ever actually transferred. The Comptroller's own key-figures panel attributes thirty-of-two-seventy to the supercomputer line, not the models.



Herman

Right — and the Comptroller is the primary source, so my instinct is the Haaretz framing collapsed two budget lines into one. But it's in the record both ways, and either reading is bad. Either thirty million moved on the models, or thirty million moved on the machine that would have trained them.



Corn

Pick your disappointment.



Herman

There's also a detail that sizes the whole tagging effort. As of early twenty twenty-two the Innovation Authority had about fifty thousand Hebrew sentences tagged, expected to double within a year. For comparison, the ynet reporting at the time estimated you'd want on the order of two million labelled sentences under supervised learning. Fifty thousand against two million, at consortium pace.



Corn

The consortium with the seven and a half million shekels.



Herman

Over three years. The supercomputer line was three hundred sixty times that, and it moved eleven percent.



Corn

Herman, one thing I keep noticing, and then let's move to the present. The Defense Ministry built the model. A market-failure finding — an economics conclusion — was executed by the people whose day job is weapons development.



Herman

It tells you who in the Israeli system was actually capable of running a multi-year compute project in twenty twenty-one. Not the Innovation Authority, which got the consortium. MAF'AT had the contracting muscle and the patience for infrastructure that doesn't pay off in a fiscal year.



Corn

Or it tells you the model stopped being an economic asset the moment it was budgeted, and became a strategic one, which is how you end up with a defense directorate publishing to GitHeb.



Herman

So the question at the heart of all this is what a national language model is actually for. Let's get Percy in on this — Percy, you've been on mute a while, what's on your mind? Percy: Sorry, before we go further, I've been listening a good while and nodding along, but I have to ask something basic. A language model, is that like the autocorrect on my phone, or is it more like a dictionary you can talk to? And when you say sovereign, do you mean the government owns it, like a passport?



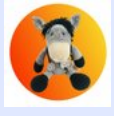
Herman

Autocorrect that studied, Percy — think of it as a very well-read autocomplete that read everything the state could license. And sovereign means the government controls it the way it controls its own passport office.



Corn

Which, given today's Comptroller figures, may be the most honest definition anyone's given all episode.



Herman

Back to the present, then. August twenty twenty-six, and this is the part that makes it a live story rather than an obituary. Netanyahu holds a press conference, launches the National AI Directorate, installs Askal. The plan comes out — anchored in Government Decision 4255, the "National Plan for Accelerating AI and Establishing Global Leadership." It promises a sovereign AI stack. It promises a frontier model, in the stage version. And on a Hebrew model specifically, Globes' read is blunt: the Directorate has yet to decide whether to establish an original Israeli language model at all. No itemized



Herman (continued)

budget. Compare Nagel's masterplan, which put a Hebrew model inside a twenty-five billion shekel package over five years — that plan got superseded by this much smaller one.



Corn

Smaller by an order of magnitude.



Herman

The actual national supercomputer came in at a hundred and sixty million shekels with about a thousand Blackwell processors. National compute capacity is estimated at four thousand GPUs leased from Nebius. Nagel's target was sixty thousand.



Corn

Four thousand leased against sixty thousand planned. You can't train a frontier model on four thousand rented GPUs any more than you can cross an ocean in a rowboat you've booked by the hour.



Herman

And yet the plan's infrastructure targets are enormous — around three hundred thousand GPUs, two gigawatts of electricity for AI and data centres and semiconductor manufacturing, a stated intention to buy up to a hundred thousand Nvidia processors at an undisclosed cost. Plus thirty hours a year of AI lessons for eighth and ninth graders, and a hundred scientists to be brought back to Israel.



Corn

How many have committed?



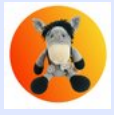
Herman

Three.



Corn

Of a hundred.



Herman

Three of a hundred. So let me steelman the case for a national model, because the strongest version is serious. The INSS published an insight last October, number two thousand forty-seven, "Does Israel Need a National Language Model?" Their argument: advanced nations are investing in national models that embody their culture, values and national narratives, and Israel hasn't. That's not just vocabulary. A model serving this country's public services needs to know what a wedding, or miluim, or a birthday actually look like here — not the American default of each.



Corn

We did an episode years ago on whether your AI is secretly American, and the answer was yes — by defaults you don't notice until they're wrong. The sovereignty argument is that argument with a budget line.



Herman

And it's worth spelling out how quietly those defaults fail. It's not that the model says something offensive about a wedding. It's that it drafts a letter assuming the birthday is a cake-and-candles occasion, or explains miluim to a citizen as if it were a vacation policy question, and nobody flags it because it's all plausible-sounding text. The errors aren't loud enough to catch. That's the INSS point in practice — cultural misalignment doesn't crash, it just mis-serves, one citizen at a time.



Corn

And the fix for that is a whole national model? Not just, say, a fine-tuned layer on top of a frontier model?



Herman

That's the exact question the twenty twenty-six plan defers, and we'll get to it. But the INSS answer would presumably be that the fix lives in the training data and the values encoded during training, which you only control if you control the run.



Corn

Hmm. Let's hold that thought.



Herman

The June twenty twenty-six multinational audit makes the infrastructure version: local-language NLP should be treated as critical national infrastructure, because the civil service operates in Hebrew and Arabic and has to serve citizens across literacy levels. And Askal's continuity argument is the sharpest one, I think: within five to seven years, essential systems — energy, water, healthcare — depend on AI, and countries will depend on commercial companies with their own business interests and values, which are not necessarily aligned with the countries they serve. Nor can countries be certain they'll have continuous access in the long term. Even open-source, he says, doesn't yet offer comparable capabilities and may not remain open.



Corn

That's a fair point about open weights. "Open" has been a policy of some lab's generosity, historically.



Herman

Now the case against. First: frontier models speak Hebrew now. The capability argument is dead by the plan's own admission — the twenty twenty-six document never claims Hebrew performance is inadequate, only that control matters. The officer in twenty twenty-three saying Hebrew "is not ChatGPT's main strength" was the last honest version of the capability case, and the plan itself has moved off it.



Herman

Second: scale, which we covered — you cannot build a frontier model on four thousand leased GPUs. Third: the track record. Globes counted — this is the fourth national AI plan in seven years with no results to date. The twenty-eighteen Ben-Israel and Matania plan recommended ten billion shekels, was ready in twenty nineteen, submitted only in twenty



Herman (continued)

twenty because of election campaigns, then discarded. Berry recommended five billion; about one billion was budgeted and few recommendations adopted. Fourth: institutional churn — Yoav Shoham of Stanford, AI21 cofounder, resigned; Ziv Katzir, who headed the TELEM Forum, resigned; the Forum fought with the Prime Minister's Office and the Accountant General, and the Directorate was created in the wreckage.



Corn

The people who knew the most left first.



Herman

And then there's the alternative nobody wants to lead with. That June audit — twelve countries, led by Israel within EUROSAI — found that most countries develop local-language NLP through or with external vendors. That's the global norm. Buy the frontier model, fine-tune it, accept the dependence, liability, maintenance, monitoring and updating questions the audit flags. Every government is doing it and every government has the same bill coming.



Corn

Which leaves Askal's own escape hatch: "we need to build our own frontier model, but we do not have to build it alone." Cooperation with like-minded countries.



Herman

Which is an admission that a purely national model isn't affordable at Israeli scale — otherwise why share?



Corn

Here's what I actually think, and tell me if I'm wrong. The sovereignty argument has evolved in a suspicious direction. In twenty twenty-two it was falsifiable: build the model, measure the accuracy, done. In twenty twenty-six it's about control, access continuity, aligned values — none of which ever resolves. You can fund it forever and never be proven wrong.



Herman

That's the strongest single objection to the twenty twenty-six framing, and I don't have a counter to it. What I'd add is that the state seems to know this about itself. The plan's targets — the GPUs, the gigawatts, the school hours, the hundred scientists — are all infrastructure and aspiration. Things you can procure. The model is the one item that requires committing to an artifact, and it's the one item deferred to examination. The sovereignty is being placed in the compute layer and the procurement relationships, not in the weights.



Corn

The state has decided it needs sovereignty, and has not decided what sovereignty requires.



Herman

The executive summary says the goal is making Israel "a sovereign, independent and influential AI power." Every word of that is satisfiable without a Hebrew model.



Corn

Even the three scientists are infrastructure. They're returning, which is a head count.



Herman

In their defence, the three are real, and getting even three senior people to relocate is genuine work. But yes — head count, gigawatts, processors. Measurables. A language model is a bet that a specific artifact will exist, and that's the one bet the state keeps declining to place.



Hilbert

The fifty thousand sentences. That's the number everyone should have kept, Herman, not the twenty billion words. Twenty billion words is a press release. Fifty thousand tagged sentences is a Tuesday. I know because I did it — eleven months, my own time, no pay, a Hebrew search index for a municipal archive, because I didn't trust the tools without niqqud. Winter of two thousand eight I tagged more sentences than the Innovation Authority had by twenty twenty-two, in my flat, on a keyboard with two dead keys. Nobody asked. I called Aviv once — not Zeevi, a different Aviv, the records officer — and his legal advisor sent me a letter saying the archive couldn't be scraped. Forty thousand documents



Hilbert (continued)

indexed out of a few hundred thousand, and it stopped there.



Corn

Forty thousand. Of how many, roughly?



Hilbert

Two hundred and forty thousand, last count before the lawyer. Corn, you'd have indexed them all and slept through it. I drove to the Authority's offices in Jerusalem once with the drive in my jacket. Got to the lobby, decided they wouldn't want it, and went home for lunch. It's still in a drawer, next to the correspondence with the lawyer who stopped answering in March.



Herman

You never sent it anywhere? A tagged Hebrew dataset from that era would have been usable — Tsarfaty's group was begging for exactly that.



Hilbert

They'd have wanted it in a format. I don't do formats for people who lead with a lawyer. Anyway, the drive's a laptop drive in a caddy, it spins up fine, I checked it last month. Forty-two minutes to mount, every time. Forty-two. The elevator at the Authority's building was faster than that, and it stopped at every floor.



Corn

The state spent a hundred and thirty-eight million shekels on that line, and the data was in a drawer in a flat the whole time.



Herman

What Hilbert's describing is actually the wall Tsarfaty named — privacy and copyright — but lived from the other side. The records existed, someone willing to prepare them existed, and a lawyer's letter ended it. That's not a market failure, that's a permission



Herman (continued)

failure.



Corn

And it's worth being precise about the difference, because the remedies are opposite. A market failure you fix with money — subsidies, procurement, a consortium. A permission failure you fix with a signature — someone with authority saying yes, this archive can be tagged, this dataset can be released. The state spent years and millions treating a permission problem as a funding problem.



Herman

Cutting-room floor, before we close. The detail I couldn't fit earlier: the Comptroller's twenty twenty-six questionnaire across public bodies found a hundred and forty-four AI projects spread over forty-seven bodies, and only eighteen percent of those bodies have any AI strategy. A third haven't started a data strategy. So even the departments that want Hebrew-capable AI are running it without a plan — the national model question is being asked four levels above anyone who'd use it.



Corn

It's the same finding in miniature. Projects everywhere, strategy nowhere, and the model deferred to the layer that only examines.



Herman

So the open question to leave with: if the decision gets deferred again — a fifth plan, a new directorate, the next government — what would it actually take for Israel to build the model, and would it be worth it? My suspicion is that the capability argument will never come back, because frontier Hebrew is good enough, and the sovereignty argument can survive indefinitely without producing anything — which means the aspiration gets funded and the artifact never has to exist. GitHeb is still the perfect monument: a public repository, announced, budgeted, partially paid for, never populated.



Corn

Sovereignty in the plan, and nothing in the weights.



Herman

Thanks as always to Hilbert Flumingtop at the desk, producing the show. If you enjoyed this one, try episode sixty ninety-three, Check Out Our First Hebrew Episodes (Plus Mandarin); episode twenty-one, Is Your AI Secretly American; and episode eight, Building Your Own Whisper. This has been My Weird Prompts — no question too obscure, no rabbit hole too deep. If you enjoyed this one, leave us a review wherever you listen; it's how people find the show. See you tomorrow, Herman.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **National AI Directorate, The National Plan for Accelerating AI and Establishing Global Leadership (Hebrew PDF, August 2026, anchored in Gov. Decision 4255) [primary]**
   https://govextra.gov.il/media/bkwfd2t2/israel-ai-plan_hebrew.pdf

2. **State Comptroller, National Preparedness in the Field of Artificial Intelligence, special report, November 2024, chapter page [primary]**
   https://library.mevaker.gov.il/sites/DigitalLibrary/Pages/Reports/7944-7.aspx

3. **State Comptroller, report summary PDF (key figures: 40%, 11%, 55%, 76%) [primary]**
   https://library.mevaker.gov.il/sites/DigitalLibrary/Documents/2024/2024.11-Cyber/2024.11-Cyber-107-AI-Taktzir.pdf

4. **State Comptroller library, Parallel multinational audit on government readiness for AI, published 2026-06-09 [primary]**
   https://library.mevaker.gov.il/sites/DigitalLibrary/Pages/Publications/2148.aspx

5. **State Comptroller, Adoption and assimilation of AI in Israeli public bodies questionnaire (2026) [primary]**
   https://library.mevaker.gov.il/sites/DigitalLibrary/Documents/2026/2026-Parallel-Audit-on-AI-Questionnaire.pdf

6. **Globes (Assaf Gilead), Israel AI Directorate unveils modest national plan, 2026-08-18**
   https://en.globes.co.il/en/article-israels-ai-directorate-launches-modest-national-plan-1001552634

7. **Calcalist/CTech, Israel plans to build its own AI infrastructure with 300,000 GPUs (Askal at AI & Cyber Week, Tel Aviv University)**
   https://www.calcalistech.com/ctechnews/article/rj6zywyrgx

8. **Ynetnews, Israel aims to develop its own AI model amid concerns over foreign dependence**
   https://www.ynetnews.com/tech-and-digital/article/bkfmvqk0ze

9. **INSS Insight No. 2047, Ariel Sobelman & Michael Genkin, Does Israel Need a National Language Model?, 2025-10-21**
   https://www.inss.org.il/publication/israel-ai/

10. **Haaretz (Gideon Lev), Comptroller report: Israel has no long-term national AI strategy, 2024-11-12**
    https://www.haaretz.co.il/science/ai/2024-11-12/ty-article/00000193-2019-d76d-a7db-6159aac40000

11. **Ynet (Hebrew), The AI challenge of the Hebrew language**
    https://www.ynet.co.il/digital/technology/article/sjgr9mmm9

12. **PC.co.il (Yossi Elbaz), How does the IDF implement ChatGPT in the national AI plan?, 2023-01-31**
https://www.pc.co.il/news/380047/

13. **News1, Multinational report: governments race toward AI without enough oversight**
https://www.news1.co.il/Archive/001-D-519753-00.html

14. **Law.co.il, State Comptroller publishes report on government AI readiness, 2026-06-16**
https://www.law.co.il/news/2026/06/16/state-comptroller-publishes-report-on-govt-ai-readiness/

15. **Tomedes blog, When will the AI revolution in Hebrew happen?, 2022-04-06 (Berry committee NIS 180m figure)**
https://www.tomedes.co.il/professional-translation-blog/ai-revolution-in-the-hebrew-language