Episode Book

EPISODE #6039

Who's Actually Building Frontier AI Outside the US and China?

PUBLISHED	16 October 2026 at 21:17 UTC
DURATION	32:45
LISTEN	myweirdprompts.com/episode/frontier-ai-outside-us-china-59655/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.5
BOOK GENERATED	19 October 2026 at 17:42 UTC

SYNOPSIS

Only two countries are at the AI frontier right now — and everyone else is training models for entirely different reasons. This episode works through the real numbers: Mistral's trillion-parameter preview scoring 38 against a frontier of 58, the UAE's fully open K2 Horizon, Korea's government-run elimination rounds, India's from-scratch Sarvam models, and more. Along the way we untangle why older index scores make sovereign models look far stronger than they are, and why "sovereign AI" means something different in Seoul, Abu Dhabi, Paris, and New Delhi. If you've ever wondered who's actually close to the frontier and who's building for language coverage, data residency, or procurement rules instead, this is the episode.

TRANSCRIPT



Corn

The closest anyone outside America and China has got to the frontier is a French model scoring thirty-eight. The frontier is sitting at fifty-eight.



Herman

Twenty points. That's the whole story in one number, and it's worth saying clearly before we start: nobody outside the US and China is at the frontier right now. Not close to it, not nearly. The nearest thing is Mistral, and Mistral is eight points behind the best open-weight model China has put out.



Corn

That's what Daniel's asking about, and he's drawn the line in a specific place. Not the US, not China. Which pockets of the world have been shipping frontier or near-frontier models lately, open or closed. Who's actually close, and who's building for reasons that have nothing to do with being close.



Herman

The second half of that question is the interesting half.



Corn

The second half is the whole episode, Herman. Anyone can read a leaderboard. The reason a Swiss university consortium and a Korean telecom and a Chilean research centre are all training models that score between six and twenty-five is the actual question.



Herman

And the answer is different for every one of them. Language coverage, data residency, procurement rules, regulation, industrial policy. Sometimes all five at once.



Corn

So let's set the yardstick first, because the numbers people have in their heads are wrong.



Herman

They are, and the reason is the index version. Artificial Analysis is on version four point three point two as of today, the sixteenth of October. On that index, Claude Opus 5.5 sits at fifty-eight. GPT-6 Astra and Gemini 4 Argon are both at fifty-three.



Corn

And the Chinese open-weight yardstick?



Herman

MiMo-V2.6-Pro at forty-six, GLM-5.3 at forty-five, Kimi K3 at forty-four. Those are open-weight models. That's the number that matters for the rest of this conversation, because China's open models are out there, downloadable, and scoring in the forties.



Corn

Right. And here's why the older numbers floating around are misleading.



Herman

Every article written before roughly the middle of this year is quoting the older index, and the older index ran hot. A model listed at thirty-seven or thirty-eight on the old scale might be a thirteen on the current one. Cohere's Command A plus is the cleanest example. It's listed at thirteen now. Older versions of the same writeup showed thirty-seven, and before that twenty-three.



Corn

Same model. Three different numbers. None of them wrong, which is the problem.



Herman

So when someone tells you a sovereign model scored thirty-seven, the first question is which index, and usually the answer is an old one and they don't know.



Corn

Start with France, since that's the closest anyone's got.



Herman

Mistral Large 4 Preview, released the sixth of October, so ten days ago. About one point zero five trillion total parameters, mixture of experts, somewhere between forty-nine and fifty-two billion active depending on which source you read. Text and image input. Artificial Analysis has it at thirty-eight, ranked sixty-sixth of two hundred and sixteen models, and it's marked as estimated because Mistral hadn't released the final configuration.



Corn

Estimated by the people scoring it, not by Mistral.



Herman

Right. Context window about five hundred and twenty-four thousand tokens as measured, though some reports say a million. On Mistral's own platform it's about a dollar thirty-six per million tokens in, four dollars eighteen out.



Corn

And Mistral's own benchmarks?



Herman

DeepSWE sixty-one point seven percent, against Kimi K3's sixty-eight. Terminal-Bench 4.0 at twenty-eight point three percent. And notably, the reinforcement-learning training was reportedly still being finished when they launched it.



Corn

They shipped a model mid-training.



Herman

They shipped a preview and called it a preview. Which is fair, but it does mean the thing people are scoring is not the finished article. It might get better. It might get restructured. What we know is the version that exists today scores thirty-eight.



Corn

The weights, though. Those aren't out.



Herman

Not out. Mistral promised open weights by the end of October. One report says the twenty-seventh. No licence has been published, so nobody knows what the terms will be.



Corn

A French company promising to give away a trillion-parameter model by the end of the month, with no licence published ten days out.



Herman

That's the part I'd watch. If those weights ship at this quality, Mistral has the strongest open model from outside China. Not close to the frontier, but the strongest outside China, and it would be a genuine from-scratch European effort, which almost nothing else in Europe is.



Corn

What's the rest of the lineup look like?



Herman

Much smaller. Medium 3.5 scores fourteen, Small 4 scores eleven. Large 4 is the whole story.



Corn

And the compute behind it?



Herman

About eight hundred and thirty million dollars of bank debt raised in March for a data centre south of Paris. Starting around forty megawatts. The reported ambitions are two hundred megawatts by the end of next year and a gigawatt by 2029. That last one comes from a single French source, so treat it as an ambition rather than a plan.



Corn

Forty megawatts to a gigawatt in three years. That's a company that has decided the constraint is electricity and it's going to buy its way out of it.



Herman

Which is the correct diagnosis. The constraint is electricity and capital, and Mistral has been unusually honest about that. They're not pretending they can out-train OpenAI on a shoestring.



Corn

Second place, then.



Herman

The UAE. K2 Horizon, from MBZUAI's Institute of Foundation Models, released the third of September. This one's interesting because it's not one model, it's six, ranging from zero point nine billion up to a flagship of three hundred and seventy-five billion total parameters with twenty-three billion active.



Corn

Apache 2.0?



Herman

Apache 2.0, and they released weights, code, training data and methods. All of it. Artificial Analysis scores it at thirty-one. There's one third-party listing claiming forty-seven, but Artificial Analysis says thirty-one and I'd trust that over the outlier.



Corn

Forty-seven would put it near the Chinese open models. Thirty-one puts it nine points behind Mistral.



Herman

And it's the leading fully open model from outside the US and China, which is a real title even at thirty-one. Because fully open means you can actually inspect the training data. That's rare.



Corn

What else has the UAE got?



Herman

K2 Think V2 in January, a seventy-billion open reasoning system from G42 with Cerebras and MBZUAI. Jais 2 in December, seventy billion and eight billion, Arabic-first, trained from scratch. TII's Falcon-H1R 7B in January, a hybrid transformer-Mamba design, scores eight. No large Falcon this year.



Corn

And the gigawatt campus.



Herman

Stargate UAE. A one-gigawatt campus with the first two hundred megawatts targeted for this year. Not confirmed live.



Corn

Targeted for this year, and it's October.



Herman

Which is the pattern with a lot of this. The announcements run ahead of the accelerators.



Corn

Korea next, because Korea's the most active country outside Europe and it's doing something nobody else is.



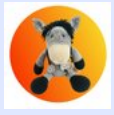
Herman

The science ministry runs a Sovereign AI Foundation Model Project. Six-month elimination rounds. In January, Naver Cloud and NC AI were cut, and Naver was ruled to have failed the sovereignty requirement because it used a foreign component.



Corn

They were disqualified for using somebody else's model.



Herman

For using a foreign component in a programme whose entire purpose is not using foreign components. Upstage, SK Telecom and LG advanced in August. The field narrows to two finalists by the end of this year or into 2027.



Corn

So it's a bake-off with a government as the judge, and the elimination criterion is partly political rather than technical.



Herman

Partly. And that tells you what they're buying. They're not buying the best model. They're buying a model with a Korean provenance.



Corn

What have they actually produced?



Herman

LG K-EXAONE 2.0, end of July. Seven hundred and fifty billion total, thirty-seven billion active, Apache 2.0, two hundred and sixty-two thousand context. LG claims parity with GLM-5.1, DeepSeek V4 Pro and Qwen 3.5. Artificial Analysis puts it at twenty, estimated.



Corn

And the parity claim.



Herman

Doesn't hold up against a score of twenty. LG also said openly that it was upcycled rather than trained from scratch, which is unusually honest and worth crediting them for.



Corn

Upcycled. Meaning they started from someone else's weights.



Herman

Started from an existing model and continued training. It's a legitimate technique and it's cheaper, and it means the model is not sovereign in the sense the programme is asking for. They said so themselves.



Corn

SK Telecom?



Herman

A.X K2, two days earlier, about six hundred and ninety billion total, thirty-three billion active, Apache 2.0. Scores twenty-one. Basically the same spec sheet as LG's, scoring one point higher.



Corn

Two Korean telecoms shipping near-identical models a week apart in a competition where only two survive.



Herman

Yes. And then Upstage, which is the more interesting one. Solar Open2 is two hundred and fifty billion total, fifteen billion active, one million token context, open weights under Upstage's own licence rather than Apache. Scores twenty-five.



Corn

Their own licence.



Herman

Their own licence, which for a government procurement audience is fine, and for everyone else is a flag. And Solar Pro 4 is closed, API only, and scores twenty-seven. That's the highest Korean model on the board.



Corn

Closed beats open by two points and that's the whole Korean frontier.



Herman

And Naver's HyperCLOVA X SEED Think, thirty-two billion, scores eleven. Naver was the one disqualified in January and this is what they shipped anyway.



Corn

Third place at eleven, having been kicked out for using foreign parts. That's a specific kind of vindication.



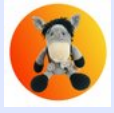
Herman

It's not vindication, it's just the number.



Corn

It's a little bit vindication.



Herman

Fine. It's a little bit vindication.

L

Larry

Are you tired of a model that knows everything except your postcode? Introducing SovereignBox, the world's first AI appliance that runs entirely inside your own home, on your own electricity, with no connection to any cloud anywhere, ever. SovereignBox knows your street, your bin collection day, and the names of all four of your neighbours' dogs. It speaks forty-eight languages, though not fluently, and it will never, ever tell a foreign government what you had for breakfast. For a limited time, SovereignBox comes with a free USB stick containing its entire training set, which is twelve thousand pages of local council meeting minutes, delivered in a fireproof fondue pot. Order now and receive the SovereignBox Junior, a smaller box that knows less but is more patriotic. Results not guaranteed. Sovereignty may vary by region. Batteries not included. Void where

L

Larry (continued)

prohibited. Your statutory rights are unaffected. Terms apply.



Corn

Ads are just weather now.



Herman

They are. Anyway, India.



Corn

Sarvam.



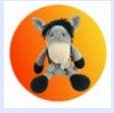
Herman

Sarvam launched a hundred and five billion and thirty billion parameter models in February and open-sourced them in March. The hundred and five has about ten billion active, Apache 2.0, a hundred and twenty-eight thousand context, about twelve trillion training tokens.



Corn

And trained from scratch.



Herman

Trained from scratch on IndiaAI Mission compute. That's the part that matters. Artificial Analysis scores it at nine on the current index. The thirty billion scores seven. An older write-up gave eighteen, which is the index-version problem again.



Corn

Nine, from scratch, on government compute. That's a real achievement and a low score at the same time.



Herman

Both true. BharatGen Param2 came out in February too, seventeen billion mixture of experts covering twenty-two Indian languages, and reportedly on a non-commercial licence.



Corn

Twenty-two languages and you can't use it commercially.



Herman

That's what "non-commercial" means, yes.



Corn

And Krutrim?



Herman

Pivoted in May to selling AI cloud services and reported its first net profit. They're out of model building. Which is a rational decision, and it's the one most of these outfits will eventually make.



Corn

Build a model, discover the margin is in renting someone else's.



Herman

That's the industry.



Corn

Japan.



Herman

Japan is the strange one, because Japan has the industrial base and hasn't converted it. Preferred Networks' PLaMo 3.0 Prime, closed API, released in June. The National Institute of Informatics has LLM-jp-4, a thirty-two billion mixture of experts under Apache 2.0. NEC cotomi v3 in March, NTT tsuzumi 2 from last October at thirty billion.



Corn

And none of them on the leaderboard.



Herman

No Japanese model is on Artificial Analysis at all. Sakana's Fugu, which gets talked about as a Japanese model, is an orchestrator that routes to GPT and Claude. It's not a model Sakana trained. Fujitsu's Takane is reportedly built on Cohere.



Corn

So the flagship Japanese AI product is a router.



Herman

And the government's GENIAC programme is shifting toward robotics foundation models, which is arguably a much smarter bet for Japan than trying to catch up on text.



Corn

That's the first thing anyone's said that sounds like a country choosing its actual comparative advantage.



Herman

It might be. Russia?



Corn

Go on.



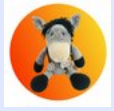
Herman

Sber GigaChat 3.1 Ultra. Seven hundred and two billion total, thirty-six billion active, MIT licence, two hundred and sixty-two thousand context. Sber says it was trained from scratch. The model card has not been independently verified.



Corn

Says.



Herman

Says. And Yandex reportedly released an eighty billion base model under Apache 2.0 in September, on a single source. Neither is on the leaderboard.



Corn

A seven-hundred-billion-parameter model with an MIT licence that nobody outside Russia has benchmarked.



Herman

Which is itself informative.



Corn

The rest of it. The long tail.



Herman

Southeast Asia and Taiwan are all adaptations. Singapore's SEA-LION, Indonesia's Sahabat-AI at seventy billion, Vietnam's Viettel VAI-LLM at a hundred and twenty billion built on NVIDIA Nemotron, Taiwan's TAIDE. All fine-tunes or adaptations aimed at local languages. Not frontier attempts and not pretending to be.



Corn

Latin America.



Herman

Chile's Latam-GPT, around February, from the CENIA research centre. About seventy billion, built on Llama 3.1, Spanish and Portuguese focus, open. A regional fine-tune.



Corn

Africa.



Herman

No 2026 release that I could find. Nigeria's N-ATLaS from September last year is a Llama-3 8B fine-tune for Yoruba, Hausa and Igbo. South Africa's Lelapa AI has InkubaLM, a zero point four billion model for African languages.



Corn

Four hundred million parameters.



Herman

It's a small model for a specific job, and it's the only thing anyone's built for those languages at that scale.



Corn

Let me try the thing I keep circling. Nearly every one of those is a fine-tune of an American model. Sahabat, VAI-LLM, Latam-GPT, N-ATLaS, Takane, and half of Korea's programme. You've got a world of sovereign AI programmes and most of them are running on Llama or Gemma or Nemotron under the hood.



Herman

That's the open-weights paradox and it's the key to the whole landscape. The genuine from-scratch efforts, the ones that actually built the thing themselves, are Mistral, MBZUAI, SK Telecom, Upstage, Sarvam, Sber, and the Swiss. That's the list. Everyone else is fine-tuning.



Corn

And LG said so out loud.



Herman

Which is the only reason we know for certain about one of them.



Corn

I want to go back to Korea's disqualification, because that's the sharpest thing in all of this. Naver was cut for using a foreign component. What does sovereignty mean if the alternative is a fine-tune of an American model?



Herman

It means the procurement rule was written for a world where the distinction was clean, and the distinction isn't clean. If you fine-tune Llama, you have a model you can run yourself, inspect, and control at inference. You don't control the base weights or the next version.



Corn

So you've got data residency and you don't have sovereignty.



Herman

You've got operational sovereignty. Which for a bank is most of what it wanted. For a defence ministry, it's less than half.



Corn

And for a government that wrote "no foreign components" into the tender, it's a disqualification.



Herman

Which is why the Korean programme is the most interesting one running, and why it might not produce a frontier model at all. It might produce a definition.



Corn

The UK. Because the UK is doing the thing everyone says they're doing and actually doing it.



Herman

Cosine's Lumen Sovereign is being trained from scratch on the Isambard-AI supercomputer under the five-hundred-million-pound Sovereign AI programme. Partners are BAE, BT, Lloyds, NatWest and LSEG. Target is deployment readiness by the end of this year.



Corn

That partner list is a defence contractor, a telecom, two banks and a stock exchange.



Herman

That's who wants a model they can run on their own metal. It's not a research consortium, it's a customer list.



Corn

That's the clearest statement of why any of this exists. BAE is not buying a model because it wants to beat Claude. BAE is buying a model because it cannot send its data to Claude.



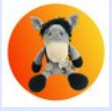
Herman

Data residency and procurement. And a smaller thing called Project Mercury was announced in April alongside it.



Corn

Nothing's shipped yet, though.



Herman

Nothing's shipped yet. The target is end of this year.



Corn

And the EU, which is doing this at the largest scale and the slowest speed.



Herman

OpenEuroLLM's first models were due the thirty-first of July. No evidence they shipped. The coordinator has said publicly that compute is the main obstacle. The final models are due in January 2028.



Corn

Fourteen months late before the first release.



Herman

The first release hasn't been confirmed as late. But the final deadline is 2028 and the interim one has passed without an announcement.



Corn

And the gigafactories.



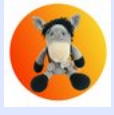
Herman

The call opened at the end of July for up to seven sites, each with over a hundred thousand accelerators. Bids close on the twelfth of November. Awards in early 2027. First sites around mid-2028. About ten billion euros of public money, with a hoped-for twenty billion private.



Corn

Mid-2028 for the first site. And the frontier is moving every few months.



Herman

Which is the structural problem. You cannot run a procurement process on a two-year cycle against a target that moves every quarter.



Corn

Switzerland, though. Switzerland decided not to play that game at all.



Herman

Apertus 1.5, released the twenty-fourth of July, from EPFL, ETH Zurich and the Swiss national supercomputing centre. Eight billion and seventy billion, fully open including the training data, Apache 2.0. Adds image input, experimental audio, an optional thinking mode, two hundred and sixty-two thousand context.



Corn

And it's not scored.



Herman

Not scored by Artificial Analysis at all. It's a public good, not a frontier contender, and the people who built it know that.



Corn

That's the cleanest version of the honest position. We are building this so that Swiss institutions have a model they own, and we are not claiming it competes.



Herman

The Swiss and the Japanese have both made the sensible version of the decision. The Swiss built the public good, the Japanese are moving to robotics.



Corn

So the reasons people build. Language first.



Herman

Arabic dialects with Jais 2. Twenty-two Indian languages with BharatGen. Korean. Japanese. Southeast Asian languages with SEA-LION and Sahabat. Latin American Spanish with Latam-GPT. African languages with N-ATLaS and InkubaLM.



Corn

And the frontier models are bad at those.



Herman

The frontier models are mediocre at those and the sovereign models are worse at everything else. If your language is Yoruba, a model that's excellent at English and useless at Yoruba is worth nothing to you, and a model that's mediocre at both but functional in Yoruba is worth a great deal.



Corn

That's a version of the argument I hadn't thought about. It's not about being better, it's about being useful at all.



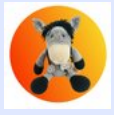
Herman

The relevant comparison isn't against Opus. It's against nothing.



Corn

Then regulation. The EU AI Act.



Herman

The general-purpose model obligations have applied since August last year. And here's the thing people miss: a fully open release makes transparency compliance much easier, because you're publishing the training data and the methods anyway. So the regulation is quietly pushing European labs toward open weights.



Corn

The Act was supposed to be a burden and it's turning into an argument for openness.



Herman

For the open labs, yes. For the closed ones it's a paperwork problem.



Corn

And industrial policy. Which is where the Cohere and Aleph Alpha thing lands.



Herman

Cohere and Aleph Alpha signed a definitive merger agreement in September, pending regulators. About twenty billion dollars combined valuation. Roughly ninety-ten ownership. Joint headquarters in Toronto and Berlin. And the Schwarz Group, which owns Lidl, committing six hundred million.



Corn

Aleph Alpha left the frontier race in 2024 for enterprise software. So this is two companies that both decided not to chase the frontier, merging to sell sovereign enterprise deployments.



Herman

Which is the correct business. Somebody has to sell the thing that goes into the bank's own rack.



Corn

And the German government probably likes having a Berlin headquarters attached to it.



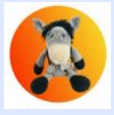
Herman

That's part of what they're buying.



Corn

Canada.



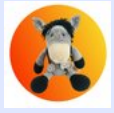
Herman

Cohere Command A plus, released in May. Two hundred and eighteen billion total, twenty-five billion active, Apache 2.0, their first fully open-source licence. Runs on two H100s when quantised. Forty-eight languages. Scores thirteen on the current index.



Corn

Two H100s. That's the number that matters most in this whole episode.



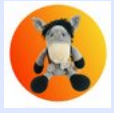
Herman

It's a real number. A bank can buy two H100s. It cannot buy a training cluster.



Corn

So the whole sovereign pitch in one spec: a model that runs on hardware an actual institution already has, covering forty-eight languages, scoring thirteen.



Herman

And Cohere's older versions showed thirty-seven and twenty-three on the old index, which is where half the confusion comes from.



Corn

Israel. Since we're here.



Herman

AI21's Jamba2 models in January. Small hybrid state-space and transformer designs, two hundred and fifty-six thousand context, Apache 2.0. Jamba 1.7 Large scores six. And AI21 has reportedly shifted its focus to its Maestro orchestration platform.



Corn

So they built the model and then decided the money was in routing other people's models.



Herman

Which is what Sakana did and what Krutrim did. Three labs, three countries, same conclusion.



Corn

There's a pattern where the announced hardware consistently runs about eighteen months ahead of the actual hardware. Stargate UAE, the gigafactories, Mistral's gigawatt by 2029. Every one of these is a phase two that no one's broken ground on.



Herman

It's a fair observation. The one exception is Mistral's French data centre, which is at forty megawatts and real, and the ambitions beyond it are the aspiration.



Corn

So what's the honest verdict on who's actually close?



Herman

Mistral at thirty-eight, and only if the weights ship and the score holds. Then a nine-point gap to the UAE at thirty-one. Then a six-point gap to Korea's closed model at twenty-seven and Upstage's open one at twenty-five.



Corn

And then a wall.



Herman

Everything else is under fifteen. Cohere at thirteen, Naver at eleven, Sarvam at nine, Falcon at eight, most of Japan unscored.



Corn

Twenty points behind the frontier is not near-frontier. It's a different category of thing.



Herman

It is, and I think that's worth being blunt about.



Corn

The other race is real, though.



Herman

The other race is completely real. It's just not measured by the Intelligence Index, and everyone keeps measuring it by the Intelligence Index anyway.



Corn

There's an argument that the fine-tune model of sovereignty is the smart play. You get a model you can run yourself for the cost of a fine-tune, you keep the data inside your borders, and you spend the difference on the thing you're actually good at.



Herman

I'd mostly agree with that, with one caveat. The day your base model's licence changes, or the base model gets deprecated, you own a model you can't maintain. You've got operational sovereignty and zero strategic sovereignty.



Corn

Which is exactly what Naver got disqualified for.



Herman

The Koreans saw that problem coming and wrote it into a tender, and the rest of the world is going to hit it in about two years.



Herman

Hilbert, you've been quiet.



Corn

Don't.



Corn

Sorry. Nothing.



Hilbert

So the principality had a coastline for eleven days and nobody told the model.



Herman

Sorry, the what had a coastline?



Hilbert

I was the sovereign AI consultant. Two years, on and off. Small country, I'm not going to name it, because the last time I named it I got a letter. Alpine-1. Trained on one H100 in a converted wine cellar. Owner of the cellar was a man called Bertrand, and I won the H100 off a Swiss banker in a poker game in Zug, which is the only part of that sentence anybody believes.



Corn

You won a GPU in a poker game.



Hilbert

I won it because he went all in on a flush draw and I had two pair. He was very gracious about it. The model was trained on cheese labels and municipal parking regulations. That's the dataset. Nineteen gigabytes of cheese labels and about four thousand pages of parking bylaws, and I want to say — the results were good. For the domain, the results were excellent. It could tell you whether you could park on a Tuesday outside a dairy in three languages.



Herman

Three languages?



Hilbert

Romansh, German, and about forty percent of French. It only spoke Romansh properly. It refused to answer in English. We asked, it declined. That was a sovereignty feature, not a bug. That's what sovereignty is. It's a model that tells you no.



Corn

What did it score?



Hilbert

Twelve.



Corn

On the current index?



Hilbert

Twelve. That's basically frontier for a model that only knows about parking.



Herman

It's four points behind Cohere.



Hilbert

It's four points behind a Canadian company with a hundred and fifty engineers, and we were two men and a cellar. I'll take that.



Corn

Where did they keep the training data?



Hilbert

On a USB stick. In a fondue pot. Which was my idea, and I stand by it, on the grounds that nobody robs a fondue pot. Then Bertrand moved the pot to the kitchen and we had an incident.



Herman

What kind of incident?



Hilbert

He put the stick in his pocket and boiled it with the cheese. One of the two. The stick survived. The dataset survived. The pot did not.



Corn

Wait. It survived boiling?



Hilbert

It was a good stick.



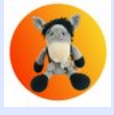
Corn

I don't know what to do with that.



Hilbert

The thing was shut down in March. It started telling people the principality had a coastline. It doesn't have a coastline. It has never had a coastline. It's ninety kilometres from the sea and it's up a mountain. But it kept putting a coastline in the answers, and it was confident, and one of the tourism pages actually published it.



Herman

You couldn't retrain it?



Hilbert

We could retrain it. We couldn't afford to correct it. Retraining on the corrected corpus meant another three weeks of the H100 and about two thousand euros of electricity, and the council wouldn't sign that off for a coastline. They said, and I'm quoting the finance officer, "the coastline is not a priority this year." So we turned it off.



Corn

So you built a model, it hallucinated a coastline, and the government chose not to fund the correction.



Hilbert

The government had a budget of four hundred thousand euros and a ski lift that needed a new cable. You know which one they picked.



Corn

I do.



Hilbert

Twelve on the index and a coastline. I'd worked on that thing for two years and the thing I remember is the cheese labels. Twelve thousand of them. Bertrand did all of the labelling by hand. Every one. He wouldn't let anyone else touch it.



Herman

Twelve thousand by hand.



Hilbert

He was very particular about the cheese.



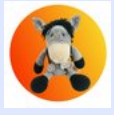
Herman

The thing I keep coming back to is the two-hundred-and-eighteen-billion model that runs on two H100s. That means the sovereign pitch and the frontier pitch are not just different races, they're built on different hardware entirely. A frontier lab is buying tens of thousands of accelerators. A sovereign deployment is buying two.



Corn

Which means the EU's gigafactories, if they get built, will not be pointed at sovereign deployment.



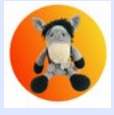
Herman

No. They'll be trying to close a twenty-point gap against labs moving every quarter, using infrastructure arriving in 2028. It's a strategy built for a competition nobody can enter.



Corn

And yet.



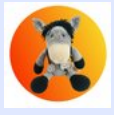
Herman

And yet. Ten billion euros buys a lot of skilled people.



Corn

The Swiss looked at that and decided to build the public good instead.



Herman

And the Koreans looked at it and wrote a procurement rule about what counts.



Corn

And Bertrand labelled twelve thousand cheese labels by hand, and the model knew everything about cheese and nothing about the sea. Which is what sovereignty actually means.



Hilbert

It refused to answer in English.



Herman

So where does that leave the question? Are the open weights actually shipping?



Corn

That's the one to watch. Mistral has promised them by the end of the month and published no licence. If they ship at thirty-eight, the strongest open model outside China is French and it's from scratch. If they slip, the whole near-frontier story for the rest of the world is a preview that never became a release.



Herman

And the second thing is the gigafactories. Bids close in November, awards in early 2027, first sites mid-2028. That timeline will decide whether Europe's answer to the compute problem arrives in time to matter, or arrives after the race has been run.



Corn

My guess is the gap persists and the reasons for building turn out to be the durable part. Language, data residency, procurement. Those aren't going to stop being real just because nobody's reached the frontier.



Herman

Seconded. The Koreans will eventually pick their two, and whatever they pick will score somewhere in the twenties, and it will be the right model for Korea.



Corn

Thanks to Hilbert Flumingtop.



Herman

Our producer.



Corn

For more along these lines, there's episode twenty-one, Is Your AI Secretly American; episode fourteen, AGI's Crossroads; and episode three, Safetensors or something else. This has been My Weird Prompts. If you enjoyed it, a review on whatever platform you're listening on helps more than you'd think.



Herman

And send us your own prompt on Telegram at t dot me slash MWP listener bot. We're back soon.