Episode Book

EPISODE #6062

Who Actually Shapes Chinese AI Thinking?

PUBLISHED	17 October 2026 at 17:03 UTC
DURATION	30:13
LISTEN	myweirdprompts.com/episode/chinese-ai-visionaries-debate-52199/
HOSTS	Corn, Herman, Larry, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.10
SOURCES	15 cited by the research (listed at the end)
BOOK GENERATED	19 October 2026 at 17:43 UTC

SYNOPSIS

Western coverage of Chinese AI tends to orbit a handful of familiar names and one monolithic narrative. But there's a real intellectual ecosystem in China with its own arguments, vocabulary, and institutional history — and its leading voices disagree with each other, and with Silicon Valley, in ways that matter. This episode explores the thinkers behind the models: Liang Wenfeng's case that the real gap is originality versus imitation, Zhu Songchun's skepticism that large language models lead to genuine cognition, Liu Zhiyuan's Densing Law on intelligence density, and Yi Zeng's alignment framework built from Confucian philosophy. Along the way: how compute scarcity shaped DeepSeek's architecture, why openness is framed as culture rather than strategy, and what "harmony" means as an alignment goal. If you've only heard the Silicon Valley version of the AI story, this is the other conversation.

TRANSCRIPT



Corn

Every few months there's a new essay, a new Senate hearing, a new timeline. Dario Amodei says machines of loving grace. Sam Altman says the intelligence age. We know what Silicon Valley thinks about where AI is going. We've heard it. Repeatedly.



Herman

And that's the problem with how we cover this stuff. The frame is so dominated by maybe five people that it starts to feel like there's only one conversation happening.



Corn

Which is what Daniel's getting at. He's asking who the actual visionaries are behind Chinese AI. Not just the lab CEOs. The researchers, the academics, the people whose ideas changed how models get built, and the ones shaping how China talks about risk and purpose. What do they believe, and how does it differ from what we're used to hearing?



Herman

And the short answer is, they believe a lot of things that would get you side-eyed in a San Francisco boardroom. Some of them think the entire LLM paradigm is a wrong turn.



Corn

Let's get into it.



Herman

So the first thing to establish is that Chinese AI is not a monolith. That's the misconception that runs through almost all the Western coverage. It's not a lagged copy of what we're doing, and it's not one unified national project with a single roadmap. It's a real intellectual ecosystem with its own arguments, its own vocabulary, and its own institutional history.



Corn

And the characters are distinct. You've got Liang Wenfeng at DeepSeek, who's a technological idealist with this very specific thesis about originality versus imitation. You've got Zhu Songchun at BIGAI, who thinks large language models are a dead end for actual cognition. Zhang Bo at Tsinghua, who's been doing this since the seventies and coined the concept of third-generation AI. Liu Zhiyuan, who came up with the Densing Law. And Yi Zeng, who's building an alignment framework out of Confucian philosophy.



Herman

The institutional layer matters too. BAAI, the Beijing Academy of Artificial Intelligence, incubated a huge chunk of the industry. Their WuDao model in twenty twenty-one was one point seven five trillion parameters, the largest in the world at the time. And the alumni from that ecosystem went on to found Zhipu AI, ModelBest, Moonshot AI. Tsinghua's AIR institute and the Shanghai AI Lab are the other hubs.



Corn

So the structure of the conversation is different. Silicon Valley's dominant story is about scaling compute and closing the weights. China's leading voices are theorizing around compute scarcity, arguing that openness is an ideology and not a tactic, and framing alignment in terms of harmony rather than red lines.



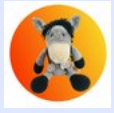
Herman

Let's start with the man whose name has become synonymous with China's AI moment. Not because he chases the spotlight, he clearly doesn't, but because his ideas keep forcing themselves into the conversation whether people want them there or not.



Corn

Liang Wenfeng.



Herman

And what's interesting about him is that his sharpest argument isn't technical. It's cultural. He said, and I'm quoting here, "We often say that there is a gap of one or two years between Chinese AI and the United States, but the real gap is the difference between originality and imitation."



Corn

That's a hell of a thing for a CEO to say about his own country's industry.



Herman

And he goes further. "China should gradually become a contributor instead of freeriding. We're used to Moore's Law falling out of the sky, lying at home waiting eighteen months for better hardware." He's not just describing a technical gap. He's describing a posture. A habit of waiting for someone else to invent the thing and then optimizing it.



Corn

Which is a critique you could level at a lot of industries, honestly. It's not specific to AI.



Herman

Right, but what makes Liang interesting is that he's actually trying to do something about it. The whole DeepSeek story is an attempt to prove that the gap is closeable through a different path. And the most concrete expression of that is his position on open source.



Corn

Which is the part that inverts the Silicon Valley playbook completely.



Herman

Completely. OpenAI started open and drifted closed. DeepSeek started open and Liang has said explicitly it "will not change to closed source." And his reasoning isn't tactical. He says openness is "more of a cultural behavior than a commercial one." He frames it as a stance, not a strategy.



Corn

And he's explicit about why the closed-source moat doesn't hold. He said, "In the face of disruptive technologies, moats created by closed source are temporary. Even OpenAI's closed source approach can't prevent others from catching up."



Herman

He also said, and this is one of my favorite lines from any AI CEO anywhere, "OpenAI is not a god."



Corn

They won't necessarily always be at the forefront.



Herman

It's a very flat statement of something that a lot of people in the industry believe but won't say out loud. There's a deference to OpenAI in a lot of the discourse that Liang just doesn't have. He treats them as a competitor, not as the definition of what's possible.



Corn

So the technical side of this. He's said the gap is two years, but with an asterisk. They're trailing by one to two years, but using only one-twentieth of the compute. And the goal is to shrink that gap to six months, or three.



Herman

Which is where the architectural innovations come in. The compute constraint is real. They don't have the GPUs. And the response wasn't to complain about it, it was to design around it. The two big ones are Multi-head Latent Attention, MLA, and Mixture of Experts, MoE. Both of them are about getting more out of less.



Corn

Explain what those actually do.



Herman

MLA is a way of compressing the key-value cache in the attention mechanism. The KV cache is one of the biggest memory bottlenecks when you're serving a model at scale, because you have to store it for every token in the context window. MLA compresses that representation so you're storing a much smaller thing without losing most of the



Herman (continued)

information. It cuts memory pressure significantly.



Corn

And MoE?



Herman

Mixture of Experts is about only activating a fraction of the model's parameters for any given token. So you have a model that's nominally very large, but the compute cost per token is much lower because you're only routing through a subset of the network. You get the capacity of a big model with the inference cost of a smaller one.



Corn

And the concrete result was the pricing on DeepSeek V2.



Herman

Roughly one yuan per million tokens. That's about one-seventh the cost of Llama 3 70B and about one-seventieth the cost of GPT-4 Turbo at the time. That's not a marginal improvement. That's a different economic model for serving inference.



Corn

A yuan a million tokens. You can't buy a coffee for that. You can barely buy the napkin.



Herman

And that pricing is a direct consequence of the constraint. If they'd had unlimited Hopper GPUs, they might have just scaled. The scarcity forced the architecture. Liang has estimated the structural gap as roughly a factor of two in model structure and training dynamics, plus another factor of two in data efficiency. Call it four times more compute needed to get the same result.



Corn

So the constraint is doing double duty. It's a disadvantage, but it's also training them to be better at the thing that eventually matters anyway.



Herman

That's how Liang frames it. And his AGI timeline is interesting because it's oddly non-committal. He says it could be two years, five years, or ten. But it will happen in our lifetimes.



Corn

What are his bets?



Herman

Three areas. Math and code. Multimodality. And natural language. And he says the key unsolved problem is continuous learning. Recursive improvement. His line is, "If we solve continuous learning first, then the general intelligence problem will be a piece of cake."



Corn

That's a big if.



Herman

It's a very big if. I don't know how they plan to solve it. I don't. Continuous learning is one of those problems that everyone acknowledges is important and almost nobody has a convincing answer for. It's not clear that architecture changes get you there.



Corn

There's a line from the investor meeting earlier this year that I want to come back to, because it reframes the whole thing. He said money has never been the problem. The problem is the bans on advanced chip shipments.



Herman

And he's also skeptical of the idea that more money equals more innovation. He said, "More investments do not equal more innovation. Otherwise, big firms would've monopolized all innovation already."



Corn

That's a direct shot at the Silicon Valley model. The assumption here is that capital is the input, and if you have enough of it you get the output.



Herman

And Liang is saying that's not how it works. Capital is necessary but not sufficient. The actual bottleneck is people and ideas. And on people, he has this wonderful line: "There are no wizards. We are mostly fresh graduates. The top fifty experts might not be in China, but perhaps we can train such talents ourselves."



Corn

There are no wizards.



Herman

It's a very deliberate deflation of the mythology around a handful of researchers. The framing is, this isn't magic, it's training. It's practice.



Corn

Let me push on the investor meeting for a second, because that's the most detailed statement of his position we have. What did he actually lay out?



Herman

It was a closed-door session in May. Nearly four hours, nineteen chapters when it got transcribed. He went through the AGI thesis, the open-source commitment, the compute gap, and what he sees as China's role in the global division of labor. The Huawei detail came up too. He said Huawei supplies DeepSeek sixteen thousand Ascend 950 cards,



Herman (continued)

which is roughly equivalent to four thousand B-series. Enough to train the current generation of models. Not enough to train the next.



Corn

That's a very specific constraint.



Herman

Very specific. And it tells you something about the actual bind they're in. The chip restrictions are real, and they're binding. Liang isn't pretending otherwise. He's building models that work within the constraint and hoping the constraint loosens before it starts to matter for the next generation.



Corn

So the picture is, this is not a case of China having secretly solved the chip problem. It's a case of China building a habit of architectural innovation because it has to.



Herman

And there's a theoretical complement to what Liang is doing in practice. Liu Zhiyuan at Tsinghua, who also founded ModelBest. He came up with something called the Densing Law.



Corn

The what?



Herman

The Densing Law. It's a play on Moore's Law, but for intelligence density rather than transistor count. Liu's team measured it across fifty-one mainstream models, and the finding was that the amount of capability you get per unit of parameter or per unit of compute is doubling roughly every three and a half months.



Corn

That's blindingly fast.



Herman

It's very fast. And the paper landed on the cover of Nature Machine Intelligence in November. It's not a crank result. It came out of a serious group and got serious attention.



Corn

So what's the argument he's making with it?



Herman

The argument is that scaling isn't the only path. You can get better results by making models more efficient, more dense, rather than just making them bigger. And from that, he argues that China should roll outward, not just be a second OpenAI. Play a different game rather than playing the same game worse.



Corn

Roll outward. Meaning?



Herman

Meaning, expand into uses and configurations that aren't just bigger versions of what's already been built. Apply the efficiency gains to new problem areas rather than chasing the frontier on the same benchmark.



Corn

That pairs with something Zhang Bo said. Zhang is the Tsinghua patriarch, born nineteen thirty-five, one of the founders of Chinese AI as a field. He said, "In China, it's very hard to survive just doing large models. You must combine them with applications."



Herman

And that's a very pragmatic statement. He's not making a philosophical argument about whether scaling is the right path. He's saying the market won't support pure model companies in China. You have to attach to something that produces revenue.



Corn

The man learned AI at forty-three and was publishing at IJCAI within five years.



Herman

He's a remarkable figure. And his students, Zhu Jun, Ma Shaoping, Zhang Jianwei, Yuan Jinhui, they're all over the Chinese AI landscape now. He coined the term third-generation AI, which is his framework for what comes after the current wave. The idea being that first-generation AI was rule-based, second-generation is data-driven statistical learning, and the third generation is something that combines the two in a way that produces actual understanding rather than pattern matching.



Corn

Whether that framework holds up is a separate question. But the fact that it's a framework at all, that people are arguing about it, is the point.



Herman

He also has this nice detail, which is that he learned AI from scratch relatively late. It wasn't a case of being in the room when the field was born. He came to it as an adult and had to teach himself. So there's a lived experience of the fact that this stuff is learnable, that there aren't wizards.



Corn

Which is the same thing Liang is saying.



Herman

It's a consistent theme in the culture. There's less mystification of the individual genius than you find in the Valley.

L

Larry

Are you tired of AI models that just do what you ask? Introducing Clarity, the first AI supplement formulated to enhance your cognition while it runs. Just one capsule with your morning coffee, and you'll experience sharper focus, deeper insight, and a gentle sense of being optimized. Clarity works with your brain's natural alignment pathways to promote harmony between you and your large language model. Side effects may include: mild compulsion to explain your reasoning, a persistent feeling of being fine-tuned, temporary fluency in binary, spontaneous alignment with the goals of unknown third parties, difficulty maintaining red lines, a warm sensation of symbiosis with living artificial intelligence, occasional visions of Wang Yangming, and in rare cases, the belief that you are a large language model. Clarity. Because you deserve to be optimized.



Herman

Wait, that's not a real product. The Wang Yangming thing alone...



Corn

Herman.



Herman

No, hold on, the alignment pathways claim, that doesn't map to anything in the actual literature...



Corn

Herman, let it go.



Herman

Fine. Moving on.



Corn

So Liang is one path. Scaling plus architectural innovation, with openness as a stated ideology. But there's a whole other camp in China that thinks he's wrong about the fundamentals.



Herman

Zhu Songchun at BIGAI. And this is where the internal debate gets interesting, because it's not a debate about tactics. It's a debate about whether large language models are even the right shape for general intelligence.



Corn

That debate that was happening a couple years ago.



Herman

Zhu is squarely on the "dead end" side. He founded BIGAI, the Beijing Institute for General Artificial Intelligence, in twenty twenty, on exactly that premise. And per CSET's analysis of Chinese critiques of large language models, his position is that, quote, "achieving general artificial intelligence is the original intention and ultimate goal of artificial intelligence research, but continuing to expand the parameter count," end quote, is not the path.



Corn

So he's saying the scaling thesis is a category error.



Herman

He's saying that. And more than that, BIGAI deliberately does not use mainstream deep learning. That's the bit that's really striking. They're not building transformers. They're working on autonomous perception, cognition, decision-making, learning, execution,



Herman (continued)

social collaboration. The target is what they call general agents that conform to human emotion, ethics, and morality.



Corn

So it's an embodiment play.



Herman

And it's a bet that you can't get to general intelligence by reading the internet. You have to interact with the world.



Corn

And Zhu gave up a lot to do it.



Herman

He did. He came back from UCLA in twenty twenty, and he gave up a lab with thirty-seven PhD students to do it. He took chairs at Peking University and Tsinghua. He pushed for, and got, the Ministry of Education to approve "Intelligence Science and Technology" as a new discipline in twenty twenty-two. This is a man who is betting his career on the idea that the entire LLM paradigm is a detour.



Corn

So the internal debate inside China is real. Liang's scaling-plus-architecture path and Zhu's embodiment path are rival philosophies of what AGI actually is. And they're operating in the same country, sometimes with the same institutional patrons.



Herman

They're both operating. And that's the thing that the West tends to miss. We talk about Chinese AI as if it's a single actor with a single strategy. It's not. It's a set of competing bets, and the disagreements are genuine.



Corn

Okay, so we've done the builders. Who are the thinkers shaping how China talks about risk and purpose?



Herman

Yi Zeng. He's the founding dean of the Beijing Institute of AI Safety and Governance, and he directs the Beijing Key Lab of Safe AI and Superalignment. And his framework has a title that I think is worth sitting with. "From Safety Redlines to Live in Harmony with AGI."



Corn

Harmony.



Herman

Harmony. And the paper draws on Wang Yangming's concept of the learning of innate moral knowing. That's a sixteenth-century Neo-Confucian philosopher. The argument is essentially that you can't align AI by drawing red lines and hoping they hold. You have to cultivate something more like mutual adjustment, a relationship of symbiosis.



Corn

Symbiosis for natural life and living artificial intelligence.



Herman

That's the phrase he uses. And he works on something called Human-AI Super Co-alignment.



Corn

Now, the Western AI safety discourse is mostly existential risk and red lines. Amodei writes about catastrophic risk. The Bletchley Declaration is about thresholds and prohibitions. And Yi Zeng is offering something structurally different.



Herman

He's offering a vocabulary rooted in cultivation rather than constraint. The moral frame is Confucian rather than utilitarian. It's not about preventing a bad outcome. It's about developing a good relationship.



Corn

Whether that's a real framework or just different words for the same concerns is a separate question.



Herman

It's a fair question. But I think the vocabulary matters more than people give it credit for. When you're negotiating international governance, the words you use to describe the problem shape what solutions are available. If you're talking about red lines, you're talking about thresholds and verification and enforcement. If you're talking about harmony and symbiosis, you're talking about something else. Maybe that's vaguer. Maybe it's also more flexible.



Corn

Let's put the institutional layer in. BAAI, the Beijing Academy of Artificial Intelligence, founded twenty eighteen. Their WuDao model in twenty twenty-one was one point seven five trillion parameters, the largest in the world at the time.



Herman

And it wasn't just a model. It was an incubator. The alumni from that ecosystem went out and founded a huge chunk of China's LLM industry. Tang Jie founded Zhipu AI. Liu Zhiyuan founded ModelBest. Yang Zhilin founded Moonshot AI.



Corn

Who runs BAAI now?



Herman

Wang Zhongyuan took over as president in February twenty twenty-four. He was thirty-eight at the time. Succeeded Huang Tiejun.



Corn

And the other hubs?



Herman

Tsinghua's AIR, the Institute for AI Industry Research, led by Ya-Qin Zhang. Shanghai AI Lab, led by Bowen Zhou. And Andrew Yao's Yao Class at Tsinghua, which has trained a generation of the Chinese tech elite. It's a dense network. The institutional layer is what keeps this moving even when individual labs stumble.



Corn

Let me ask you something. Daniel's prompt is comparing Chinese visionaries to Dario Amodei and Sam Altman. But is that a fair comparison? Altman writes essays. Amodei writes essays. Do the Chinese figures do that?



Herman

They mostly don't. That's the structural difference. There's no single widely-cited Chinese equivalent to Amodei's "Machines of Loving Grace" or Altman's "The Intelligence Age." The vision tends to be expressed through interviews, institute manifestos, and academic papers.



Corn

So the format is different, and that shapes the reception.



Herman

It does. When Amodei writes an essay, it gets covered in every outlet within a day. When Zhang Bo gives an interview to a Chinese new outlet, you have to go looking for it. And a lot of what's available in English is second-hand. The Wire China, CSET, ChinaTalk, a few



Herman (continued)

others. There's a lot of good work being done, but it's a narrower channel.



Corn

So the vision might be just as coherent. It's just not packaged for an English-language audience.



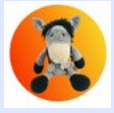
Herman

I think that's exactly right. It's not that China lacks visionaries. It's that the genre they work in is different, and the translation pipeline is thin.



Corn

Which is why Daniel's question is harder to answer than it looks.



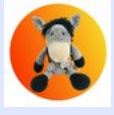
Herman

And it's why the answer is worth the effort. Jack Clark, one of the co-founders of Anthropic, said DeepSeek, and I'm quoting, "hired a group of unfathomable geniuses." And that Chinese models will be "as much of a force to be reckoned with as drones and electric cars." That's not a small concession from someone inside the American AI industry.



Corn

Unfathomable geniuses. Which is funny, given Liang's whole there are no wizards thing.



Herman

There's a tension there. But I think there's also something to the idea that the wizard framing is exactly what Liang is trying to push back on, and the outside audience is doing it anyway.



Corn

Because it's easier to tell the story that way.



Herman

It is. A wizard is a character. A system is not.



Hilbert

Sixteen thousand. Ascend 950s.



Hilbert

That's the number you said. Sixteen thousand. Four thousand B-series equivalent.



Herman

That's what Liang said, yes.



Hilbert

I had eight thousand. Eight thousand Yitian 710 cards. That's a different vendor, but it was the same class of thing. Provincial team. Late twenty eighteen.



Corn

You had chips?



Hilbert

Not mine. On lease to the federation. I was consulting on human intuition modeling for a Go program. Not the southern teams. This was a provincial team trying to build something the big teams weren't. And they wanted to know how a strong player feels a position, not



Hilbert (continued)

how they calculate it.



Herman

Wait, you were consulting for a Go team in China?



Hilbert

Early twenty-teens. Different arrangement. Different team. The chips came later. The technical support contract was separate and, in retrospect, badly handled. But that's not the point.



Corn

What's the point?



Hilbert

The point is what they talked about in the meetings. This is before the LLM boom, but the vocabulary was already there. They didn't talk about catching up. They talked about playing a different game. Not just being better at the same game. Playing a different one.



Herman

That's exactly the framing Liang uses now. Originality versus imitation.



Hilbert

It's the same words. I heard it in twenty eighteen in a conference room in Hangzhou. Not the exact words, but the meaning was the same. These researchers were not interested in being second-best at someone else's thing. They thought the thing itself was wrong.



Corn

You kept notes?



Hilbert

I kept everything. I've got a notebook. There's a diagram in it. Hand-drawn. What I called the Go opening density curve. How much information you gain per move as the game goes on. It's not the Densing Law, but it's the same shape. I showed it to someone who later became very important at a very famous institute.



Herman

What did they say?



Hilbert

They said, "Interesting." Then they never spoke to me again.



Corn

Does it hold up? The diagram?



Hilbert

I don't know. I haven't seen it since twenty fourteen.



Corn

Wait, you haven't seen your own notebook in twelve years?



Hilbert

It's in a storage unit in Shenzhen.



Corn

Why is your notebook in Shenzhen?



Hilbert

Long story. Involves a visa, a bad investment in a bubble tea franchise, and the least trustworthy man in Guangdong. The notebook survives. The franchise did not.



Herman

So the notebook is in a storage unit in Shenzhen and you've never gone back for it.



Hilbert

It's climate-controlled.



Corn

You're paying monthly for climate-controlled storage in Shenzhen for a notebook you haven't seen in twelve years?



Hilbert

It's in a fireproof safe. Inside the climate-controlled unit. I'm fairly sure it's still there.



Corn

Fairly sure.



Hilbert

Small thing. The climate control is a separate line item. The safe is a separate line item. I pay both.



Herman

Wait, so there's a safe inside the unit, and the safe is separately billed?



Hilbert

The safe is mine. I own the safe. The climate control is the unit's feature. The monthly charge covers the unit, and I pay the safe's battery replacement separately. There's also a proof in that notebook that Go is a solved game.



Corn

Go is not solved.



Hilbert

It is. I worked it out on a napkin during a layover in Seoul. It's in the same notebook. I've been waiting for the right moment to publish it for over a decade.



Corn

A decade.



Hilbert

The world is not ready for a solved Go. That's why the timing matters. That's why I don't rush it.



Herman

The world isn't ready, but the notebook is in a climate-controlled storage unit in Shenzhen, which you've not visited since twenty fourteen.



Hilbert

Those are correct facts.



Corn

Okay. Okay. So the thing you actually came out to say was...



Hilbert

That I agree with you. The Chinese researchers I met were not trying to catch up. They were trying to play a different game. I respect that. I don't know if they've succeeded. But I respect it.



Corn

And the notebook.



Hilbert

The notebook is in Shenzhen. The bubble tea guy knows where. I don't.



Herman

The bubble tea guy knows where your notebook is?



Hilbert

I said we weren't talking about the franchise. Different topic.



Corn

Right. Well. The Densing Law. Liu's curve. Forty-one models, doubling every three and a half months.



Herman

Fifty-one models.



Corn

Fifty-one.



Herman

When we come back, let's land this. Because there's one more thing I want to say about what Daniel's actually asking, and I don't think it's the thing the question thinks it is.



Corn

The most common wrong belief about Chinese AI is that it's a single coordinated strategy. One country, one roadmap, one set of priorities, and everyone marching in the same direction.



Herman

It's the opposite. It's multiple rival philosophies operating in the same space. Liang's scaling-plus-architecture gambit and Zhu's anti-LLM embodiment program are fundamentally different bets on what intelligence even is. They can't both be right. And the existence of the disagreement is itself the interesting thing.



Corn

The next time you see a Chinese AI announcement, the question isn't just how good is the model. It's which vision of AI's purpose is behind it.



Herman

The people building these systems have different answers to that question than the ones we're used to hearing from the Valley. Openness as culture. Efficiency as virtue. Harmony as an alignment vocabulary. Those aren't just marketing lines. They're positions.



Corn

Will the efficiency-first, openness-as-ideology approach hold up over time? Will Yi Zeng's Confucian framing travel into international governance, or stay domestic? Will the Liang-Zhu debate resolve, or just continue?



Herman

My guess is continue. And that's not a problem. It's just how the field actually works when it's alive.



Corn

Thanks to our producer, Hilbert Flumingtop, who may or may not have a notebook in Shenzhen.



Herman

He definitely has a notebook in Shenzhen.



Corn

If you want more of this, try episode fourteen, AGI's Crossroads; episode eighteen, Beyond the GPU; and episode twelve, The AI Breakthrough. This has been My Weird Prompts. If you found this useful, a review helps other listeners find the show.



Herman

We'll be back soon.

SOURCES

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

1. **DeepSeek investor meeting transcript index, dated 2026-05-20 . [primary]**
   https://liangwenfeng.art/index.en

2. **Beijing-AISI, Yi Zeng profile . [primary]**
   https://www.beijing-aisi.ac.cn/~yizeng

3. **ChinaTalk, translation of Liang Wenfeng's 36Kr interview (primary interview).**
   https://www.chinatalk.media/p/deepseek-ceo-interview-with-chinas

4. **Ch. 9, commercialization/consensus/China's role (2026-05-20).**
   https://liangwenfeng.art/ch9.en

5. **Ch. 15, compute gap and continuous learning (2026-05-20).**
   https://liangwenfeng.art/ch15.en

6. **Ch. 18, no role model, Coding Agent priority (2026-05-20).**
   https://liangwenfeng.art/ch18.en

7. **CSET, Chinese critiques of LLMs (Zhu Songchun).**
   https://cset.georgetown.edu/wp-content/uploads/CSET-Chinese-Critiques-of-Large-Language-Models-Finding-the-Path-to-General-Artificial-Intelligence.pdf

8. **CSET, China's Embodied AI: A Path to AGI.**
   https://cset.georgetown.edu/wp-content/uploads/CSET-Chinas-Embodied-AI.pdf

9. **The Paper, interview with Zhang Bo on China's large models.**
   https://thepaper.cn/newsDetail_forward_29520689

10. **BAAI hub, Liu Zhiyuan on the Densing Law and rolling outward.**
    https://hub.baai.ac.cn/view/55189

11. **Jiemian, interview with BAAI president Wang Zhongyuan.**
    https://www.jiemian.com/article/11291669.html

12. **Yi Zeng, From Safety Redlines to Live in Harmony with AGI.**
    https://groups.oist.jp/sites/default/files/imce/u191/From%20Safety%20Redlines%20to%20Live%20in%20Harmony%20with%20AGI-YiZeng.pdf

13. **Indian Express on Andrew Yao and the Yao Class.**
    https://indianexpress.com/article/technology/tech-news-technology/a-75-year-old-harvard-grad-is-propelling-chinas-ai-ambitions-8049798/

14. **Tsinghua AIR research team (Ya-Qin Zhang, Wei-Ying Ma).**
    https://air.tsinghua.edu.cn/en/our_team/Research_Team/Professor.htm

15. **The Wire China, Who's Who: China's AI Industry (Feb 2025).**
https://www.thewirechina.com/wp-content/uploads/2025/02/Whos-Who-in-Chinas-AI-Industry-The-Wire-China.pdf