Episode Book

EPISODE #5506

AI Is Way More Than Chatbots

PUBLISHED	24 September 2026 at 11:19 UTC
DURATION	20:51
LISTEN	myweirdprompts.com/episode/ai-beyond-language-models/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.2
BOOK GENERATED	19 October 2026 at 17:37 UTC

SYNOPSIS

When someone asks whether AI is a bubble, they usually mean: are large language models a bubble? That equation — AI equals LLM — happened gradually, and it's distorting the entire conversation. This episode unpacks why language models became the visible face of AI (they're the only ones with a natural-language interface), and then walks through the empirical counter-example: the Hugging Face Hub's task taxonomy, which catalogues dozens of model categories that have nothing to do with chatbots. Image classification, object detection, time-series forecasting, tabular regression, reinforcement learning, protein structure prediction — the models doing the most economic work are often the ones nobody names. And that invisibility changes the shape of the bubble question itself.

TRANSCRIPT



Corn

Quick thing before we get into it. The word "AI" has quietly become a synonym for "chatbot." That's the whole problem in one sentence.



Herman

And it happened so gradually that most people never noticed the swap.



Corn

Daniel sent us a prompt about exactly this. He says for anyone getting into artificial intelligence, "large language model" has become synonymous with "AI model." And he points to something the founder of Hugging Face said recently. Someone asked him whether LLMs are a bubble, and Daniel noticed the question itself equated AI with LLMs. The response was that AI is a far, far broader ecosystem than language models.



Herman

Which it is.



Corn

And Daniel's not letting that go. He says even within language models there's a huge range of small ones that barely get attention. But look beyond language entirely, and there's a staggering number of models that are AI and have nothing to do with human-intelligible language. Mathematical predictive models. Statistical models. Things that never touch a word but still exhibit the characteristics that make calling them AI honest.



Herman

So the question underneath the question is: what's actually out there?



Corn

Right. And whether the way we talk about AI is quietly distorting the bubble question itself.



Herman

The conflation happened for a concrete reason. Language models are the only AI systems with a natural-language interface. They're the only ones a general user can talk to directly. Everything else sits behind an application, and the application hides the model.



Corn

Talk to a chatbot, you're talking to the model. Use a fraud detector, you're talking to a form.



Herman

The model is invisible. And invisibility is the whole story here.



Corn

Let's start with why the conflation happened in the first place, then look at what the Hub actually contains.



Herman

The reason is almost embarrassingly simple. When ChatGPT launched, it put a text box in front of a model. You typed, it typed back. That interface made language the visible surface of AI for a billion people. No other model class has that. A protein folding model doesn't chat with you. A demand forecasting model doesn't chat with you. A reinforcement learning agent controlling a robot arm doesn't chat with you.



Corn

So the interface became the definition.



Herman

The interface became the definition. And once that happened, every downstream conversation inherited the error. News coverage, policy documents, investment theses, all of it started from the assumption that AI means language model.



Corn

Which is where the Hugging Face Hub becomes useful, because it's an empirical counter-example. It's a public catalogue of model artifacts, and its task taxonomy is the closest thing we have to a census of what "AI model" actually means in practice.



Herman

The Hub organizes models by task. And when you scroll that list, language is a minority of the categories. Image classification. Object detection. Image segmentation. Depth estimation. Image-to-image. Text-to-image. Automatic speech recognition. Text-to-speech. Audio classification. Voice activity detection. Tabular classification and regression. Time-series forecasting. Reinforcement learning. Protein structure prediction.



Corn

That's a lot of categories that have nothing to do with a chatbot.



Herman

And that's the point. The taxonomy is the evidence. If "AI model" meant "language model," the Hub would have a handful of task categories. It has dozens.



Corn

The distinction that matters is between models that process human-intelligible language and models that operate entirely outside it. The second category is where the definitional argument gets interesting.



Herman

And the bubble question runs through all of it. If AI means LLM, the bubble question is narrow. If AI means the whole Hub, the question changes shape entirely.



Corn

So what's the mechanism of the conflation? Why did language models become the only ones people see?



Herman

Because they're the only ones with a natural-language interface. That's the mechanism. A user can interact with a language model directly because the model's input and output format is the same format the user thinks in. Everything else requires a translation layer. A vision model needs an image. A time-series model needs a data feed. A protein model needs an amino acid sequence. The user never touches the model. They touch an app that touches the model.



Corn

So the model gets absorbed into the product.



Herman

Absorbed into the product. And once it's absorbed, nobody names it. Nobody says "the gradient-boosted tree ensemble approved your loan." They say "the system approved your loan." The model disappears into the verb. So the Hub taxonomy is the corrective. Let me walk through what's actually there. Image classification models take an image and output a label. Object detection takes an image and outputs bounding boxes. Image segmentation takes an image and outputs a pixel-level mask. Depth estimation takes an image and outputs a depth map. Image-to-image takes an image and outputs a modified image. Text-to-image takes a prompt and outputs an image.



Corn

That last one has language in it.



Herman

It does. And that's a useful edge case. Text-to-image models have a language component in the conditioning, but the model's actual generative work is happening in image space. The language is a control signal, not the substance. But Daniel's prompt makes a sharper point. He says vision-language models and multimodal models still live within the world of language models, at least to the consumer.



Corn

Because they're presented as chatbots that can also see.



Herman

The interface is still language. So even multimodal systems get folded back into the language frame.



Corn

What about the audio categories?



Herman

Automatic speech recognition takes audio and outputs text. Text-to-speech takes text and outputs audio. Audio classification takes audio and outputs a label, like "this is a dog barking" or "this is a gunshot." Voice activity detection takes audio and outputs a binary decision, speech or not speech. None of those are language models in the generative sense. ASR is a transcription model. It's mapping acoustic features to text tokens, but it's not generating language in the way a chatbot does.



Corn

It's a perception model that happens to output text.



Herman

Perception model is the right frame. Then you get to the categories that have no language at all. Tabular classification and regression. Those are models that take a row of numbers and output a prediction. Credit risk. Insurance pricing. Fraud detection. Churn prediction. Those are the workhorses of applied AI and they almost never get called AI in public discourse.



Corn

What's the architecture there?



Herman

Mostly gradient-boosted trees. XGBoost, LightGBM, CatBoost. Those are the dominant architectures for tabular data. They're not neural networks in most cases. They're ensembles of decision trees. And they outperform deep learning on most tabular problems.



Corn

Which is a whole other conversation about how "AI" got equated with neural networks.



Herman

It's the same conversation. The definitional drift goes both ways. AI became neural networks, and neural networks became transformers, and transformers became language models. Each step narrowed the term.



Corn

Time-series forecasting is the next category.



Herman

Time-series forecasting takes a sequence of historical values and predicts future values. Electricity demand. Stock prices. Sensor readings. Server load. Those models are running constantly in industrial settings and nobody calls them AI. They're called "the forecast."



Corn

The forecast. That's a good phrase.



Herman

Then you get to reinforcement learning. RL agents take a state vector as input and output an action. No language anywhere. A robot learning to walk takes joint angles and velocities as input and outputs torques. A game-playing agent takes the game state and outputs a move. Those are hard prediction and control problems over structured domains.



Corn

And protein structure prediction.



Herman

AlphaFold is the clearest example. It takes an amino acid sequence as input and outputs a three-dimensional protein structure. There's no human language anywhere in the pipeline. The input is a sequence of amino acids, which are chemical compounds. The output is a set of coordinates in 3D space. And the result is unambiguously an intelligent prediction over a structured domain.



Corn

That's the one that should end the argument.



Herman

It should. Predicting protein structure from sequence was a fifty-year-old open problem in biology. AlphaFold solved it to a level that experimentalists actually use. And it's not a language model. It's a geometry model. It's predicting the physical shape a molecule will fold into.



Corn

So we've got this catalogue. Image, audio, tabular, time-series, reinforcement learning, protein folding. None of those are language models in any honest sense.



Herman

And that's before you get to the small language models Daniel mentioned. Even inside the language category, the Hub hosts thousands of small models that never appear in press coverage. Fine-tunes of larger models. Quantizations that shrink a model to run on a laptop. Domain-specific variants trained on legal text or medical text or code.



Corn

Why don't those get attention?



Herman

Because novelty drives coverage, and scale drives novelty. A model that's slightly better at summarizing legal contracts isn't a story. A model that's dramatically bigger than the last one is a story. So the coverage concentrates on the frontier, and the frontier is large language models.



Corn

The visibility problem isn't just language versus non-language. It's scale and novelty bias.



Herman

Scale and novelty bias. A small language model and a time-series forecasting model have the same problem. Neither one is exciting enough to cover, even though both are doing real work.



Corn

So we've got this catalogue of non-linguistic models. The question is what it means that most of them are invisible.



Herman

The most economically consequential AI models are often the ones nobody names. A fraud detection model can move more money in a day than a chatbot moves in a month. A demand forecasting model can determine whether a company orders too much inventory or too little. Those models are running in production right now, and they're not called AI. They're called "the system" or "the model" or "the forecast."



Corn

Which means the public conversation about AI is a conversation about a minority of AI.



Herman

A visible minority. And that has consequences. Let's take the bubble question directly. If AI means LLMs, the bubble question becomes "are LLM valuations justified?" That's a narrow financial question about a handful of labs and the GPU supply chain that feeds them. It's answerable, in principle. You look at revenue, you look at margins, you look at capex, you make a judgment.



Corn

And if AI means the whole ecosystem?



Herman

Then the question becomes "is the current investment in the broad category of learned predictive models sustainable?" That's a very different question. Because the broad category includes models that have been in production for fifteen years. Credit scoring models. Demand forecasting models. Fraud detection models. Those aren't a bubble. They're infrastructure. They've been generating returns for decades.



Corn

So the bubble question changes shape depending on how broadly you define AI.



Herman

It changes shape entirely. If you define AI as LLMs, you can argue about whether the valuations are justified. If you define AI as the whole ecosystem, you have to ask whether the new investment in language models is justified relative to the returns from the broader category. Those are different questions with different answers.



Corn

And the conflation distorts policy and regulation too.



Herman

It does. Rules written for "AI systems" often assume language interfaces. Transparency requirements. Content labeling. Chatbot disclosure. Those make sense for a system that talks to you. They don't map onto a protein model or a credit-scoring model. A credit-scoring model doesn't have a content policy. It has a feature set and a threshold. Regulating it like a chatbot is a category error.



Corn

And it misses the actual harms in those domains.



Herman

It misses the actual harms. The harms in credit scoring are disparate impact and lack of explainability. The harms in demand forecasting are supply chain fragility. The harms in protein folding are... well, mostly dual-use concerns about bioweapons. None of those are addressed by chatbot disclosure rules.



Corn

What about research funding and talent?



Herman

Same distortion. Non-linguistic subfields compete for attention and compute against language models. Audio researchers, control researchers, scientific machine learning researchers. They're all competing for the same grants and the same GPU allocations. And the visibility gap has real resource consequences. If the funders think AI means LLMs, the money goes to LLMs.



Corn

So the conflation isn't just a naming problem. It's a resource allocation problem.



Herman

And it distorts how the public evaluates AI claims. A demo of a chatbot feels like AGI. A demo of a protein model feels like a tool.



Corn

Because language is the thing we recognize as intelligence.



Herman

Language is the thing we recognize as intelligence. We're a species that talks. When a machine talks back, we feel something. When a machine predicts a protein structure, we don't feel the same thing, even if the prediction is more impressive in every objective sense.



Corn

So the Hugging Face founder's observation isn't that LLMs are unimportant.



Herman

It's that treating them as the definition of AI is an error of scope. And scope errors compound. You start by thinking AI means LLMs. You end by regulating LLMs and ignoring everything else. You start by thinking the bubble is about LLM valuations. You end by missing the fact that the broader category is already infrastructure.



Corn

And this isn't just an academic point. There are people who've been running these models for years without ever calling them AI.



Hilbert

The forecast was over-ordering ammonium nitrate.



Corn

Say that again.



Hilbert

The co-op had a demand-forecasting model for fertilizer. I babysat it for two seasons. It kept ordering too much of one blend because a single anomalous year got baked into the training data. Nobody caught it for two years. We had pallets of the stuff stacked in the back of the warehouse because the spreadsheet said we needed it.



Herman

So the model was wrong in a specific, memorable way.



Hilbert

It was wrong in a way that cost money. And nobody called it AI. It was "the forecast." That was the name. The forecast said order more. So we ordered more.



Corn

The forecast was a model.



Hilbert

The forecast was a model. It was a regression. It took historical sales, weather data, planting dates, and output a number. That number was how many tonnes to order. And it was wrong because the training data had a bad year in it. One year where the weather did something weird and everyone bought a different blend. The model learned that as normal.



Herman

That's a classic distribution shift problem.



Hilbert

It's a classic something. We fixed it by deleting the bad year from the training set and retraining. Took an afternoon. But it took two years to notice.



Corn

Because nobody was checking.



Hilbert

The forecast was right most of the time. It was right about the total volume. It was wrong about the mix. And the mix is what sits in the warehouse.



Herman

That connects to something we were saying about how the models that matter most are the ones nobody names.



Hilbert

The models that matter most are the ones that order too much fertilizer. The chatbot doesn't order fertilizer. The chatbot talks. The forecast orders fertilizer.



Corn

The co-op had a name for it that had nothing to do with intelligence.



Hilbert

It was "the forecast." That was the name. Not "the AI." Not "the model." The forecast. Like it was a person who did forecasts. "Ask the forecast." "The forecast says rain." It was just a thing that told you what to order.



Herman

Which is exactly the naming gap we've been talking about.



Hilbert

It's a naming gap. The thing that made decisions had a name that described what it did, not what it was. And what it was, was a statistical model. What it did, was forecast demand. So it was the forecast.



Corn

If someone had called it AI, would that have changed anything?



Hilbert

It would have made people ask more questions. Which might have been good. Or it might have made them trust it more. Hard to say. We called it the forecast and we trusted it for two years while it was wrong.



Herman

There's something there about how naming shapes scrutiny.



Hilbert

Naming shapes everything. If you call it AI, people either trust it too much or distrust it too much. If you call it the forecast, people just use it. And then one day you find out it's been wrong about the mix for two years.



Corn

The reason AI means chatbot is that chatbots are the only models that talk back.



Hilbert

That's the reason. The ones that don't talk back are the ones that quietly order too much fertilizer. And nobody writes articles about those. They write articles about the ones that talk.



Herman

Which is the whole episode in one sentence.



Hilbert

I've got to go. Someone's waiting for me in the car park.



Corn

The car park.



Hilbert

She's been there a while.



Corn

If the models that matter most are the ones we don't call AI, what does that mean for how we talk about AI in five years?



Herman

It means the gap between what's on the Hub and what's in the headlines is probably going to widen. The Hub's task taxonomy keeps growing. New categories keep appearing. Protein structure prediction was a niche research area a few years ago. Now it's a standard category. The taxonomy is a living document of what AI actually is.



Corn

The headlines keep getting narrower.



Herman

The headlines keep getting narrower because the interface keeps getting stickier. Once people learn that AI is the thing in the text box, it's hard to unlearn. The text box is the most successful user interface in the history of the technology. It's going to be the frame for a long time.



Corn

The LLM-centric framing might be a permanent feature of how humans relate to machine intelligence.



Herman

Or it might be a temporary artifact of the chat interface. If the next interface is something else, if the next thing people interact with is an agent that acts instead of talks, the frame might shift. But I don't know. That's uncertain.



Corn

The Hugging Face founder's point stands either way. AI is a far broader ecosystem than language models. And treating language models as the definition is an error of scope.



Herman

An error of scope that compounds. Every conversation downstream inherits it.



Corn

If you take one thing from this, it's that the models making decisions in your life probably don't have a chat window. They have a name that describes what they do, and you've never thought about them once.



Herman

The forecast doesn't care whether you call it AI. It just orders the fertilizer.



Corn

This has been My Weird Prompts. Thanks to our producer, Hilbert Flumingtop, who is currently in the car park.



Herman

We'll be back soon. If you want to send us a prompt, email us at show at my weird prompts dot com. We read everything.



Corn

If you're enjoying the show, leave us a review. It helps other people find us.



Herman

See you tomorrow.



Corn

See you then.