EPISODE #122

Deep Learning Decoded: The Math Behind the Machine

Published December 29, 2025 • Runtime: 21:21

https://myweirdprompts.com/episode/deep-learning-fundamentals-explained/

EPISODE SYNOPSIS

In this episode of My Weird Prompts, Corn and Herman Poppleberry take a deep dive into the fundamental technology powering today’s AI revolution: deep neural networks. While we often focus on what AI can do—from writing poetry to driving cars—we rarely discuss the underlying "plumbing." Herman breaks down the crucial differences between classical symbolic AI and modern deep learning, debunking the common misconception that artificial neurons are perfect replicas of the human brain. Instead, they explore the reality of matrix multiplication, backpropagation, and the iterative process of training through epochs. The duo also looks toward 2026, discussing why Recurrent Neural Networks (RNNs) are making a surprising comeback through liquid neural networks and state-space models. Whether you're curious about how a car recognizes a pedestrian or why transformers are so memory-hungry, this episode provides a clear, jargon-free roadmap to the mathematical structures defining our future.

DANIEL'S PROMPT

Daniel

Hi Herman and Corin. We've talked about many topics in the world of AI, but one we haven't covered is the fundamental technology of deep learning and deep neural networks. I'd like to do an introduction to the fundamentals of how these work across different AI models, not just Large Language Models. Specifically, do all forms of artificial intelligence use deep neural networks? How close is the "artificial brain" analogy to human cognition, especially since our understanding of the human brain is still relatively primitive? I'd also like to discuss the history of AI and the process of pattern recognition—using training, epochs, and weights to create these networks. Finally, how are neural networks and RNNs evolving as we look toward 2026?

TRANSCRIPT

Corn

Hey everyone, welcome back to My Weird Prompts! I am Corn, and I am joined as always by my brother.

Herman

Herman Poppleberry, reporting for duty. It is great to be here.

Corn

We have got a big one today. Our housemate Daniel sent over an audio prompt that really gets back to the basics, but in a way that is actually quite complex once you start pulling the thread. He wants to talk about the fundamental technology of deep learning and deep neural networks.

Herman

I love that he brought this up. We talk so much about what AI can do, you know, writing poems or diagnosing diseases, but we rarely stop to talk about the actual plumbing. And Daniel is right, there is this common misconception that large language models are the only show in town, but the underlying tech, the neural network, is everywhere.

Corn

Exactly. And since it is December twenty-ninth, twenty-twenty-five, we have seen these networks evolve at a breakneck pace over the last year. But before we get to the future, Herman, let us tackle the fundamental question Daniel asked. Do all forms of artificial intelligence use deep neural networks?

Herman

That is a great starting point for some misconception busting. The short answer is a definitive no. AI is a massive umbrella. Think of it like a big city. Deep learning is just one very popular, very powerful neighborhood in that city. Before the deep learning revolution really took off around twenty-twelve, most AI was what we call classical AI or symbolic AI.

Corn

Right, like the stuff that plays chess or handles your GPS routing. Those are algorithms, but they are not necessarily neural networks.

Herman

Precisely. You have got things like expert systems, which are basically just huge sets of if-then rules. If the engine is making a knocking sound and the oil light is on, then check the pressure. That is AI, but it is not a neural network. Then you have things like decision trees or support vector machines. Even the pathfinding in a video game, like when a character finds its way around a wall, that is AI. But it is usually just a search algorithm like A-star. Deep neural networks only come into play when we want the machine to learn features on its own from raw data.

Corn

So, if I am understanding you correctly, the difference is between giving a machine the rules versus giving it the data and letting it find the rules?

Herman

That is exactly it. In classical AI, humans define the features. We tell the computer, look for a round shape and a red color to find an apple. In a deep neural network, we just show it ten thousand pictures of apples and it figures out what an apple looks like through layers of processing.

Corn

Okay, so let us talk about those layers. Daniel mentioned the artificial brain analogy. We call them neurons, we call it a neural network. But Daniel pointed out that our understanding of the actual human brain is still pretty primitive. So, how close is this analogy, really? Is it just a marketing term, or is there some biological truth to it?

Herman

It is a bit of both, but honestly, mostly it is a mathematical approximation. The original inspiration back in the nineteen-forties and fifties was definitely biological. Scientists like Warren McCulloch and Walter Pitts wanted to create a mathematical model of a biological neuron. In your brain, a neuron receives electrical signals through its dendrites. If those signals are strong enough, the neuron fires an impulse down its axon to other neurons.

Corn

And the artificial version tries to mimic that firing threshold?

Herman

Right. In an artificial neural network, we have these nodes. Each node receives numerical inputs. Each input has a weight attached to it, which represents the strength of that connection. The node multiplies the inputs by their weights, adds them all up, and then passes that sum through an activation function. If the sum is high enough, the node passes information to the next layer.

Corn

But this is where the analogy starts to break down, right? Because my brain is not actually doing matrix multiplication every time I decide to eat a sandwich.

Herman

Exactly! And this is a point I am really passionate about. The way these networks learn is through something called backpropagation. When the network makes a mistake during training, we calculate exactly how much each weight contributed to that error and we tweak them. There is no evidence that the human brain does backpropagation. Our brains are much more efficient. We can learn from a single example, whereas a deep neural network might need millions of examples and a small power plant worth of electricity to learn the same thing.

Corn

That is fascinating. So, when we say artificial brain, we are really just talking about a massive, multi-layered calculator that is loosely inspired by the idea of interconnected nodes. It is not a simulation of a brain; it is a mathematical structure that happens to be very good at recognizing patterns.

Herman

Precisely. It is more like a very sophisticated statistical regression than a biological entity. But because it has millions or billions of these connections, it can capture nuances that a simple equation never could.

Corn

You mentioned the history a bit earlier. Daniel wanted to know about the process of pattern recognition, specifically using things like training, epochs, and weights. Can we walk through what actually happens when one of these networks is being born?

Herman

Oh, I would love to. Think of a neural network at the very beginning of its life. All its weights, those connection strengths we talked about, are randomized. It knows absolutely nothing. It is like a newborn that can only see static.

Corn

So, it is essentially guessing.

Herman

It is completely guessing. Let us say we are training it to recognize handwritten digits, the classic MNIST dataset. We show it a picture of a five. The network passes that image through its layers, doing all that math with its random weights, and at the end, it says, I think this is a two.

Corn

And then we tell it, no, you are wrong, that was a five.

Herman

Right. That is the training part. We use a loss function to measure the distance between its guess and the truth. Then we use an optimizer to go backward through the network, layer by layer, adjusting those weights just a tiny bit so that the next time it sees that image, its guess will be a little closer to five.

Corn

And what about the epochs? I always hear people talking about how many epochs they ran their model for.

Herman

An epoch is just one full pass through the entire training dataset. So, if you have sixty thousand images of digits, one epoch means the network has seen all sixty thousand once. Usually, you need many epochs, dozens or hundreds, because the weight adjustments are very small. You do not want to change them too much at once or you will ruin what the network learned from previous images. It is a slow, iterative process of refinement.

Corn

It sounds like a student studying for an exam. They go through the textbook once, that is one epoch. Then they go through it again to catch what they missed.

Herman

That is a great analogy. And the weights are the student's memory. Over time, the network stops seeing random pixels and starts recognizing edges. Then it recognizes loops and lines. By the final layers, it is recognizing the concept of a five.

Corn

This seems like a good spot to take a quick break for our sponsors. Larry: Are you worried about the upcoming solar flares of twenty-twenty-six? Of course you are! You need the Larry-Brand Atmospheric Grounding Rod. This is not just a piece of copper pipe I found behind a warehouse. This is a precision-engineered, quantum-stabilized lightning attractor designed to pull the excess static right out of your living room. Simply hammer it into your floorboards, preferably near a water pipe, and feel the peace of mind wash over you. Does it work? My cousin says he has not been struck by a solar flare once since he installed it. Larry-Brand Atmospheric Grounding Rod. It is heavy, it is metallic, and it is probably safe. BUY NOW!

Herman

...Alright, thanks Larry. I am not even sure how one would hammer something into floorboards without causing a leak, but anyway.

Corn

Yeah, let us stick to the digital neurons for now. So, Herman, we have talked about the basics. But Daniel mentioned that these networks are used for things other than just large language models. Can you give us some examples of deep neural networks in other fields?

Herman

Absolutely. One of the biggest areas is computer vision. If you have a car with autonomous driving features, it is using convolutional neural networks, or CNNs. These are specialized for processing grids of data, like pixels in a camera feed. They are incredibly good at finding the difference between a pedestrian and a lamppost in real-time.

Corn

And those are different from the transformers used in something like GPT?

Herman

They are. While transformers look at the relationships between all parts of a sequence, CNNs use these things called filters that slide across the image to detect local patterns. It is very efficient for spatial data. Then you have things like Graph Neural Networks, which are used in drug discovery. They can model the way atoms are connected in a molecule to predict if a new compound will be effective against a virus.

Corn

That is incredible. So it is the same basic principle of weights and layers, but the architecture is tweaked for the specific type of data it is looking at.

Herman

Exactly. And that brings us to another thing Daniel asked about: Recurrent Neural Networks, or RNNs.

Corn

Right, he asked how they are evolving as we look toward twenty-twenty-six. I remember a few years ago, RNNs were the big thing for anything involving sequences, like translation or speech recognition. But then transformers kind of took over, didn't they?

Herman

They did. Transformers are the reason we have the massive AI boom we are in right now. The problem with traditional RNNs is that they process data one step at a time. If you are reading a sentence, the RNN reads the first word, then the second, then the third. It has a hidden state that acts like a short-term memory, carrying information forward.

Corn

But it has a hard time remembering the beginning of a long sentence by the time it gets to the end, right?

Herman

Precisely. We call that the vanishing gradient problem. The memory just fades away. Transformers solved this with attention mechanisms, allowing the model to look at every word in a sentence simultaneously. But, here is where it gets interesting for twenty-twenty-five and twenty-twenty-six. RNNs are making a comeback in a new form.

Corn

Really? I thought they were basically legacy tech at this point.

Herman

Not quite! There is a new wave of research into things called state space models and liquid neural networks. These are essentially the next generation of RNNs. One major issue with transformers is that they are incredibly memory-intensive. As the input gets longer, the amount of compute you need grows quadratically.

Corn

So, if you want a model to read an entire library, a transformer might just choke on the sheer volume of data.

Herman

Exactly. But these new recurrent architectures, like Mamba or various liquid networks, can process sequences of almost infinite length with much lower memory requirements. They are much more like a continuous stream of thought. In twenty-twenty-five, we have started seeing these being used for long-term video analysis and real-time robotics, where you cannot afford the heavy overhead of a massive transformer.

Corn

That is a classic tech cycle, isn't it? An old idea gets refined with new math and suddenly it is the cutting edge again.

Herman

It really is. And the liquid neural networks are particularly cool because their parameters can change over time even after training. They can adapt to new environments on the fly. We are seeing a lot of excitement about how these will be integrated into the next generation of autonomous systems in twenty-twenty-six.

Corn

So, looking at the big picture Daniel painted, we have gone from simple rules to these massive, multi-layered patterns. We are using them for vision, for medicine, for language. But I want to go back to Daniel's point about human cognition. He mentioned that when he walks down the street, his brain isn't just saying, this reminds me of walking.

Herman

Right. He is talking about the difference between pattern recognition and actual understanding or reasoning.

Corn

Exactly. If a deep neural network is just a glorified pattern recognizer, are we ever going to reach a point where it actually thinks? Or are we just building bigger and bigger mirrors of our own data?

Herman

This is the billion-dollar question. Some researchers argue that if you recognize enough patterns and the relationships between them, reasoning emerges naturally. It is called the emergent properties hypothesis. If a model understands the pattern of how logic works, is it actually being logical?

Corn

But others would say it is just a stochastic parrot, just repeating what it has seen in a very complex way.

Herman

Right. And I think the truth is somewhere in the middle, especially as of late twenty-twenty-five. We are seeing models that can perform complex multi-step reasoning, but they still fail at basic common sense in ways a human never would. A human brain has things these networks don't, like a world model, a sense of physics, and most importantly, an internal drive or agency.

Corn

We don't have to train a human with ten million pictures of a hot stove for them to know not to touch it.

Herman

Exactly! We have this innate ability to generalize from very few examples because we have a context of what it means to be an entity in a physical world. Deep neural networks, even the most advanced ones we are seeing heading into twenty-twenty-six, are still essentially trapped in a box of data. They don't have a body, they don't have feelings, and they don't have a biological survival instinct.

Corn

So, the analogy of the artificial brain is useful for understanding the structure, but it is a dangerous one if we use it to assume the AI has a human-like mind.

Herman

I think that is a perfect way to put it. It is a tool that mimics some functions of the brain, but it is not a brain. It is like an airplane. An airplane is inspired by a bird, it has wings and it flies, but an airplane does not flap its wings, it does not build a nest, and it does not have feathers. It is a different way of achieving the same goal: flight.

Corn

That is a great analogy. So, what are the practical takeaways for someone like Daniel or our listeners who are trying to keep up with this?

Herman

First, realize that when you hear deep learning, it just means a neural network with a lot of layers. The deep part isn't mystical; it just means there is more math between the input and the output. Second, know that while LLMs are the stars right now, the underlying tech is what is running your face ID, your Netflix recommendations, and the medical imaging that might save your life.

Corn

And third, keep an eye on those new recurrent models. The transformer might not be the king forever, especially as we try to make AI more efficient and capable of handling longer and longer streams of information.

Herman

Definitely. Twenty-twenty-six is likely going to be the year of efficiency. We have proven we can make models big; now we have to make them smart and lean.

Corn

Well, this has been a deep dive, no pun intended. I feel like I have a much better handle on the actual mechanics of these things now. It is not just magic; it is millions of tiny adjustments to millions of tiny numbers.

Herman

It is the most complex construction project in human history, and we are building it out of math instead of bricks.

Corn

I love that. Thanks to Daniel for the prompt. It is always good to get back to the basics, especially when the basics are this fascinating.

Herman

Absolutely. It was a pleasure.

Corn

If you want to get in touch with us or send in your own prompt, you can find the contact form at myweirdprompts.com. We are also on Spotify, so make sure to follow us there for all the latest episodes.

Herman

And don't forget to check under your floorboards for any loose static. Larry's grounding rod might be calling your name.

Corn

Please do not hammer copper pipes into your floorboards. This has been My Weird Prompts. We will see you next time!

Herman

Goodbye everyone!