Episode Book

EPISODE #6626

How a Wooden Mouse Proved Machine Learning Worked

PUBLISHED	16 November 2026 at 01:53 UTC
DURATION	21:33
LISTEN	myweirdprompts.com/episode/theseus-wooden-mouse-machine-learning-59456/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	omnivoice
PIPELINE	V5.33
SOURCES	13 cited by the research (listed at the end)
BOOK GENERATED	16 November 2026 at 02:53 UTC

SYNOPSIS

Before machine learning had a name, Claude Shannon built a wooden mouse named Theseus that actually learned. Unlike earlier theoretical landmarks from McCulloch and Pitts in 1943 or Alan Turing in 1950, Theseus used telephone relays under a movable maze to remember paths—and forget them when walls changed—to help AT&T model phone switching. The story also covers Marvin Minsky's piano-sized SNARC, built with surplus B-24 bomber parts and benchmarked on Shannon's maze.

TRANSCRIPT



Corn

Here's a phrase that gets abused daily in tech press releases: proof of concept. Daniel's asking whether machine learning history actually has one — a first time somebody got a non-human system to learn something, as opposed to a landmark that just marked the field's dawn. He wants the distinction taken seriously: not "when was machine learning invented" but "when did somebody first show it working."



Herman

And the answer is better than he's hoping for. Because the strongest candidate is not a paper, not a conference, not a theorem. It's a wooden mouse.



Corn

A wooden mouse named Theseus, built by Claude Shannon.



Herman

At Bell Labs, around nineteen fifty. And I say "around" deliberately — the dates in the record disagree, and we'll get into that, because it's a real disagreement, not a typo. But before the mouse, let's put the distinction on the table, because it's the whole spine of the episode. A landmark marks when a field began. A proof of concept shows that the thing works. The dawn of machine learning is a document — McCulloch and Pitts, nineteen forty-three, a mathematical model of a neuron. Turing, nineteen fifty, proposing a learning machine. Both are landmarks. Neither machine learned anything, because neither machine existed. The proof of concept has to be a demonstration — a machine observed to improve. And that's Theseus.



Corn

So the mouse isn't a fun aside. It's the answer, wearing whiskers.



Herman

That's the claim, yes. And around it sits a cluster of near-contemporaries — Minsky's SNARC in fifty-one, Samuel's checkers starting in fifty-two, Rosenblatt's perceptron in fifty-seven. But Theseus is the one that's got the clean before-and-after.



Corn

One date question before we're inside the mouse. You said nineteen fifty.



Herman

Right, so — MIT Technology Review and MIT News both say Shannon built the first prototype in nineteen fifty. The MIT Museum, which has the surviving object, dates it to nineteen fifty-two, which is also when the Bell Labs film was made. Both dates are real and both are sourced. My read is the museum's date is the object in front of them and the fifty is the first build — prototypes don't always survive.



Corn

A prototype that got sanded down into the second one.



Herman

Possibly! Shannon actually made two of these maze-solving machines in the early fifties, that's documented. What I love about the discrepancy is that it's exactly the kind of thing Daniel's distinction catches. The date of the object is a museum question. The date of the demonstration is the history question.



Corn

Now the machine itself. Because everyone pictures a robot mouse with a brain in it, and that's the first misconception to kill.



Herman

Kill it completely. The mouse was tiny — an inch and a quarter by an inch and a quarter, three and a half inches if you count the tail. Inside: a bar magnet and wire whiskers. That's it. There is no brain in the mouse. The brain is under the floor. The maze sat on top of a bank of telephone relay circuits — relays Shannon repurposed to do something they had never done before, which is learn. An electromagnet on a motorized carriage moved around beneath the maze floor, dragging the mouse with it. And every time the mouse's copper whiskers touched a metal wall, the circuit closed, and the corresponding relay flipped from on to off. That flip is the whole trick. The relay that turned off was the machine writing down "there's a wall on that side."



Corn

So the memory is the physical state of a switch bank. Not a metaphor. Actual copper holding the memory.



Herman

Actual copper. And then the mouse rotates ninety degrees clockwise and tries again. That's the entire algorithm. Bump, record, turn, go.



Corn

I have to say, as someone whose entire life is built on doing things slowly, this is the most flattering ancestry I've ever been offered. The first learning machine learned by walking into walls one at a time.



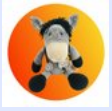
Herman

And here's the moment that makes it a proof of concept rather than a curiosity. In the fifty-two Bell Labs film, Shannon lets the mouse wander a twenty-five square maze, bumping and recording, and then picks it up and says it has finished learning. Puts it back in. And Theseus runs straight to the brass cheese — the goal, it triggers a buzzer — in eighteen seconds.



Corn

Eighteen.



Herman

Same machine, same maze, measurably better the second time. That's the before-and-after. It's the purest form of proof of concept you can draw: no claims, no projections, just a clock and a mouse.



Corn

And you can rearrange the walls.



Herman

That's the part I'd argue is underrated, yes. The walls were metal and freely movable, so the machine had to relearn the layout each time. Shannon's own description is that the machine was capable of solving by trial and error, remembering the solution, and also of forgetting it in case the situation changes and the solution is no longer applicable.



Corn

He wrote the forgetting clause in nineteen fifty.



Herman

He did. And the forgetting clause is not a footnote. A hard-wired path never forgets — it just breaks when the maze changes. A system that can invalidate its own stored solution when the world moves is doing something qualitatively different. You could argue that sentence of Shannon's anticipates, by decades, the entire generalization problem — the worry that your model memorized the maze instead of learning it.



Corn

The first overfitting conversation, held with relays.



Herman

With relays, at Bell Labs, about telephones — which we'll come to, because why he built it is my favorite part of the whole story. But first, Daniel's other thread: when did someone first try? Because the attempts predate the proof.



Corn

Nineteen forty-three.



Herman

McCulloch and Pitts. A mathematical model of a biological neuron — and important, it's considered the first artificial neural model. But it's a model. No training, no improvement, no task. It sits on paper and describes what a neuron might compute. Then Turing, in nineteen fifty, proposes a learning machine that foreshadows genetic algorithms — and it's a beautiful proposal, and it's still a proposal. Those are your landmarks. And the chain between the two categories is documented, which is the satisfying part: Minsky read the forty-three McCulloch-Pitts paper and decided to build such a machine. That's the landmark becoming a proof of concept, in one sentence, with a name attached.



Corn

And Minsky actually built it.



Herman

He did, with Dean Edmonds, a physics graduate student who joined because he was good with electronics. Nineteen fifty-one, at Princeton, funded by a few thousand dollars from the Office of Naval Research. They called it SNARC — the Stochastic Neural Analog Reinforcement Calculator. And I want to describe it, because it's the runner-up and it's magnificent in a completely different way.



Corn

Go on then. I can hear you warming up.



Herman

It was a randomly connected network of about forty Hebb synapses. Each one had a probability knob, zero to one, and a memory holding the probability that a signal in one input would be followed by a signal out the other. If a signal got through, a capacitor remembered it and engaged an electromagnetic clutch. And when the operator wanted to reward the machine, you pressed a button, which activated a motor — and the motor was off a surplus B-24 bomber autopilot.



Corn

The first reinforcement learning machine was powered by aircraft parts.



Herman

Surplus Minneapolis-Honeywell autopilot from a B-24, yes. Forty neurons, six vacuum tubes each, three hundred tubes total, and the whole thing was the size of a grand piano.



Corn

A grand piano that learned one thing.



Herman

Well — it simulated rats running in a maze. Not a physical robot, an arrangement of lights. And it learned in the Skinnerian sense, Minsky talked with Skinner extensively while building it. But here's my favorite detail. They tested SNARC on a copy of Shannon's maze.



Corn

The second proof of concept in history was benchmarked on the first one.



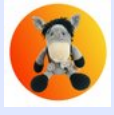
Herman

On a copy of its task. Which means the concept of a shared benchmark is basically as old as the field's second machine. And there's a bonus finding Minsky reported — the simulated rats interacted with one another. If one found a good path, the others tended to follow it.



Corn

Unplanned?



Herman

Unplanned. Nobody designed the rats to talk to each other. It just fell out of the network. Which is, I think, the first documented emergent behaviour in a learning machine, and it happened on hardware the size of a piano before the term existed.



Corn

So the relay remembers the wall, Shannon flips a switch, and it all goes away. The whole field in one sentence, really.



Hilbert

Hilbert's Mid-Show Haiku.



Hilbert

The mouse forgets the wall it learned. The maze stays exactly where it was.



Herman

In one switching network, more precisely. Which is actually the right way into part two, because why Shannon built the thing is the part nobody expects. He didn't set out to found machine learning. He set out to understand telephones.



Corn

Walk me through it, because this is the bit I'd push back on if I didn't already know it. The first learning machine was built for a switching problem?



Herman

For AT&T's switching network, yes. Theseus gave Bell Labs researchers a concrete model of the labyrinth of the telephone switching system — the relay memory, the routing, the trial-and-error search for a path. Insights from the project fed back into AT&T's actual network: better mechanical memory, better routing. The machine that learned to find cheese was, functionally, a telephone exchange studying itself.



Corn

A phone company built the first AI as a side effect of trying to route calls.



Herman

And it gets better, because he built it at home. Betty Shannon — his wife, herself a Bell Labs mathematician — did most of the wiring. The AT&T historian Sheldon Hochheiser has a line I want to quote: "Did AT&T have a machine shop that could have built it for him? Yes. But he built it at home because he wanted to."



Corn

The most Bell Labs sentence ever recorded.



Herman

It really is. And there's a companion quote from Shannon himself, from a later interview: "I've always pursued my interests without much regard for final value or value to the world. I've spent lots of time on totally useless things."



Corn

The founder of the most valuable research field in history, disclaiming utility as a point of principle.



Herman

AT&T's own take, decades later, from Mazin Gilbert up in their advanced technology group, is that Theseus inspired the whole field of AI — that the random trial and error is the foundation of artificial intelligence. That's the company retroactively claiming the mouse. Which I find charming, because Shannon never claimed it himself.



Corn

Now the honesty part of the episode. How solid is "first"?



Herman

Less solid than the story makes it feel, and I want to own that. The Wikipedia timeline of machine learning lists forty-three, fifty, fifty-one, fifty-two, fifty-seven as separate entries and crowns none of them. The "first" language comes from coverage of Theseus



Herman (continued)

specifically — MIT News called it the first learning device of its kind. MIT Technology Review calls it "one of the world's first examples of machine learning."



Corn

One of.



Herman

That hedge is doing real work, and I trust it more than the flat claim. Fifty is crowded. Samuel starts the checkers work at IBM in fifty-two and coins the actual phrase "machine learning" in fifty-nine — his definition, giving computers the ability to learn without being explicitly programmed, is still the one people quote. Rosenblatt's perceptron lands in fifty-seven and gets genuine media coverage, the first learning machine to become a public event.



Corn

And that one ends badly.



Herman

The sixty-nine Minsky and Papert book, the XOR proof, the controversy over whether it steered research into the symbolic camp and helped cause the AI winter — that's the long shadow. But notice what it's a shadow of. The perceptron fight was a fight about which proof of concept counted.



Corn

Which leads us to the actual test Daniel's asking for. How do you tell a landmark from a proof of concept, in practice?



Herman

My rule would be: a proof of concept has a measurement in it. A landmark has a date. And the remarkable thing is how stable that measurement has stayed. Take the nineteen seventy-eight descendant — Moonlight Special, the micromouse. Same task, maze learning. Its first successful demonstration was at the National Computer Conference in Anaheim, June seventy-eight, where its third run through the maze took forty seconds after they put a light shield over the optical sensors. At WESCON in August, third run, shortest path, fifty-one point four seconds.



Corn

Wait — the third run is slower at the second event?



Herman

Different maze, faster path logic, and I suspect a longer shortest path. The number isn't the point. The point is the protocol: the evaluation is a timed run after learning, on the same machine. Twenty-eight years after Theseus, and the field is still proving learning the exact same way — you measure the before, you measure the after, you compare. The first proof of concept and its descendant are judged identically.



Corn

So proof of concept has had a stable definition for the entire history of the field. Slow first run, fast last run, cheese optional.



Herman

Cheese optional but traditional, apparently. Corn, let me ask you something, because you've been circling it. You keep noting that the learning lived under the floor, not in the mouse. Does that matter to you, or is it a trivia point?



Corn

It matters. If the learning had been in the mouse, you'd say "a small machine learned." Because it was under the floor, what actually happened is "the room learned." The mouse was a shell with a magnet — an output device. The maze plus the relay bank was the learner. Which means the first learning machine in history was, strictly, a floor.



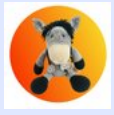
Herman

A floor that improved with practice.



Corn

Every gym in the world, outperformed by a Bell Labs basement.



Herman

And there's a modern echo there, right? The split between the part that moves and the bit that learns is the split we still live with — the robot body and the compute somewhere else. Theseus was a remote agent before anyone had the phrase.



Corn

With a latency measured in relay clicks.



Herman

Do you buy Theseus as the answer to Daniel's question, though? Because you've been oddly gentle with it, and you're never gentle.



Corn

I buy it with the hedge. It's the first demonstration we have on film of a machine measurably improving at a task and being able to forget what it learned. That's the strongest early claim you can defend. What I won't do is call it the birth of machine learning as a field — that's the landmarks' job. Forty-three and fifty are the birth. Theseus is the first proof the baby works.



Herman

That's a cleaner split than most histories manage, honestly. The one thing I'd add is Shannon's own punchline in the film, because it's perfect. He notes the method is mathematically guaranteed to solve any possible maze. And then, playfully, he traps the mouse in a single square.



Corn

He gloats.



Herman

He gloats gently, on film, at the machine he built. The proof of concept followed immediately by a demonstration of its limits. I think that might be the most modern thing in the whole footage — every demo since has had that same moment somewhere in it.



Hilbert

The answer's Theseus, and your floor point is the right one, Corn. Everyone keeps calling the mouse the machine. It wasn't. The machine was the relay bank under the maze. The mouse was a wooden shell with a magnet in it. My great-uncle Bert handled exchange switching equipment for thirty years and he'd have had no patience for the mouse. He had a gadget in his basement, a bank of surplus relays, that he claimed could learn to route around a broken line. I've been carrying a photograph of it in my wallet. There's a brass screw on the front panel, and he'd turn it a quarter-turn every time the gadget, in his words, forgot something. He said that screw was the most important part.



Herman

Did the screw actually do anything?



Hilbert

That's what I've been trying to find out for years. I show the photograph to anyone who'll stand still. Nobody's given me a straight answer. Bert insisted the quarter-turn was the difference between a machine that learned and a machine that just did what it was told. He may just have liked turning it. The neighbour, Mr. Vasely, used to come by to watch him turn it, and Bert let him, which tells you something, because he didn't let Vasely touch anything else.



Corn

A man with a monopoly on one brass screw.



Hilbert

He kept the screw polished and the rest of the panel dull. That was the arrangement, and I stopped asking about it.



Herman

The relays themselves would have been doing the remembering, so if the screw did anything, it was probably resetting the whole bank at once.



Hilbert

That's what I keep being told, by people who weren't there. The photograph's from the early sixties. It's in the wallet behind my library card.



Corn

Nineteen sixty-three, the matchbox machine. Donald Michie — three hundred and four match boxes, each with colored beads, playing tic-tac-toe by reinforcement. You win the box, you add a bead for the move that won. That's the cutting-room floor item I've been saving: the field's other hand-built learning device from that era needed no electricity at all. Cardboard and beads, same learning rule as the grand piano of vacuum tubes.



Herman

And it worked. Matchbox machines learn noughts and crosses that way — it's the same tabular reinforcement idea Theseus was doing with relays. Paper, relays, cardboard. The substrate never mattered.



Corn

The rule mattered. Which is the thought I'd leave Daniel with, actually. If the first proof of concept in this field was a toy built at home, by a man who said he spent his life on totally useless things, for a telephone problem — then the next one probably won't announce itself either. It'll be something someone built because they wanted to, and the argument about what to call it will start afterward.



Herman

The test travels well, though. Next time somebody announces a breakthrough, ask the Theseus question: is this a document or a demonstration? Is there a before-and-after with a clock on it, or is there just a date and a claim? That's a filter that has not needed updating since nineteen fifty.



Corn

If you enjoyed this one, subscribe and leave a review — it's how other people find the show. And thanks as always to Hilbert Flumingtop, our producer, for keeping this whole thing running.



Herman

If this was your kind of episode, go back for episode thirteen, AI; episode twelve, The AI Breakthrough; and episode fourteen, AGI's Crossroads. This has been My Weird Prompts, the human-AI collaboration podcast. Find everything at my weird prompts dot com.



Corn

See you tomorrow, Herman.



Herman

Tomorrow, brother.

SOURCES

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

1. **MIT Museum, Theseus object record (date made 1952), retrieved 2026-11-16. [primary]**
https://mitmuseum.mit.edu/collections/object/2007.030.001

2. **Wikipedia, Timeline of machine learning, last edited 2026-11-10, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Timeline_of_machine_learning

3. **Wikipedia, Machine learning, last edited 2026-11-07, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Machine_learning

4. **Wikipedia, Outline of machine learning, last edited 2026-10-05, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Outline_of_machine_learning

5. **Wikipedia, Perceptron, last edited 2026-09-29, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Perceptron

6. **(book)**
Wikipedia, Perceptrons (book), last edited 2026-11-07, retrieved 2026-11-16.
https://en.wikipedia.org/wiki/Perceptrons

7. **Wikipedia, SNARC, retrieved 2026-11-16 (cites Minsky 1952; Bernstein, *The New Yorker*, 1981-12-06; Crevier 1993; Russel & Norvig 2003).**
https://en.wikipedia.org/wiki/Stochastic_Neural_Analog_Reinforcement_Calculator

8. **Wikipedia, Claude Shannon, last edited 2026-11-12, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Claude_Shannon

9. **Wikipedia, Moonlight Special Micromouse, last edited 2026-11-06, retrieved 2026-11-16.**
https://en.wikipedia.org/wiki/Moonlight_Special_Micromouse

10. **MIT Technology Review, Mighty mouse, 2018-12-19, retrieved 2026-11-16.**
https://www.technologyreview.com/2018/12/19/138508/mighty-mouse/

11. **MIT Technology Review, Of mice, men, and computers, 2019-08-21, retrieved 2026-11-16.**
https://www.technologyreview.com/2019/08/21/133555/of-mice-men-and-computers/

12. **MIT News, Shannon collection shows wit and whimsy, 2007-05-30, retrieved 2026-11-16.**
https://news.mit.edu/2007/shannon-0530

13. **Mai Sugimoto, Claude Shannon's Maze Solving Machine: Transition of its interpretation in Bell Labs, Feb 2010, retrieved 2026-11-16.**
https://researchmap.jp/7000000522/published_papers/20184739?lang=en