

The Making of Elliot

A long-form note on building a minimal embodied agent, testing three versions, and finding that better food-seeking came with more instability.

2026.03.22 02:04 DENNIS HEDEGREEN BUILD OPEN-NOTE

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-MAKING-OF-ELLIOTT/](https://hedegreenresearch.com/articles/the-making-of-elliott/)

I started building a small agent called Elliot, not to prove intelligence, simulate a real brain, or make a toy that only looks alive in terminal. The point was simpler than that. I wanted to see what happens when a very small entity is placed in a world with food, pain, movement, energy, and consequence. No language. No dataset. No backpropagation. Just a body in a small world, and synaptic weights that change when actions lead to reward or punishment.

Why Build Elliot

A lot of what gets called intelligence today is really performance inside very large symbolic systems. That is interesting, but it is not the only place to look. I wanted to start somewhere lower, closer to action, closer to consequence, and closer to the question of what happens before language. Elliot is not meant to be a claim about consciousness. He is a testbed. A small embodied learning system. The question was not “can Elliot think?” The question was whether simple adaptive behavior could emerge from local consequence alone.

The Smallest Useful World

Elliot lives in a grid. He has a position. He has energy. There is one food source. There is one pain source. Moving costs energy. Food restores it. Pain takes it away. Walls also punish bad movement. If energy reaches zero, Elliot dies and resets. The important part is that learned weights persist, so the system is allowed to accumulate experience across lives.

The internal logic is deliberately small. Elliot senses direction to food. Elliot senses direction to pain when it is nearby. Those signals feed weighted action tendencies. The chosen action changes the world. The world returns a consequence. The active pathway is then reinforced or weakened. That is the core loop. The synapses are not meant to be a real hardware memristor simulation, but they are memristor-inspired in the sense that the connection itself carries the local memory of previous consequence. The model is not trained offline and then deployed. It learns in the act of surviving badly.

The Three Versions

The first version was the smallest possible baseline. Each sense had a single signal line: food north, food south, food east, food west, pain north, pain south, pain east, pain west. It worked, but it also showed a problem quickly. Elliot could get trapped in local loops. He would repeat patterns, burn energy, and fail to break out.

So the second version added a loop-break mechanism. If Elliot started oscillating between a small set of positions, the system introduced a chance to break the pattern. The third version kept that loop-break behavior and changed the sensory layer instead. Each sense no longer had one single line. Each sense got a small population of five neurons. The point was to see whether a slightly richer sensory population would make the system less brittle.

How I Tested It

At first, it was tempting to just watch Elliot move in terminal. That is useful for intuition. It is not enough for comparison. So the system was turned into something more disciplined: one watch mode and one benchmark mode. The benchmark tracked food rate, average steps per life, pain and wall events, and randomized loop-break moves. Then all three versions were run over long sessions. The useful comparison came from 500000 steps and 10 runs per version.

```
dennis-hedegreen@jobzu: ~  
Run 2: food_rate=0.0782 foods=39104 pains=1577 walls=2  
Run 3: food_rate=0.0785 foods=39228 pains=1609 walls=3  
Run 4: food_rate=0.0808 foods=40382 pains=1809 walls=2  
Run 5: food_rate=0.0833 foods=41646 pains=1707 walls=7  
Run 6: food_rate=0.0784 foods=39205 pains=1565 walls=1  
Run 7: food_rate=0.0781 foods=39058 pains=1597 walls=3  
Run 8: food_rate=0.0827 foods=41335 pains=1650 walls=8  
Run 9: food_rate=0.0807 foods=40343 pains=1612 walls=2  
Run 10: food_rate=0.0844 foods=42183 pains=1638 walls=2  
  
v3 Five Neurons Per Sense  
Run 1: food_rate=0.0788 foods=39387 pains=1532 walls=2  
Run 2: food_rate=0.0797 foods=39863 pains=1729 walls=1  
Run 3: food_rate=0.0775 foods=38746 pains=1651 walls=0  
Run 4: food_rate=0.0792 foods=39604 pains=1693 walls=2  
Run 5: food_rate=0.0775 foods=38728 pains=1562 walls=0  
Run 6: food_rate=0.0820 foods=41000 pains=1631 walls=4  
Run 7: food_rate=0.0851 foods=42531 pains=1787 walls=8  
Run 8: food_rate=0.0870 foods=43503 pains=1822 walls=11  
Run 9: food_rate=0.0802 foods=40079 pains=1727 walls=4  
Run 10: food_rate=0.0837 foods=41855 pains=1703 walls=4  
  
Benchmark summary  
v1 Single-Sense Baseline  
mean food_rate: 0.0692  
mean avg_steps_per_life: 91.65  
mean pain+wall: 321.00  
  
v2 Single-Sense + Loop Break  
mean food_rate: 0.0804  
mean avg_steps_per_life: 134.70  
mean pain+wall: 1647.90  
  
v3 Five Neurons Per Sense  
mean food_rate: 0.0811  
mean avg_steps_per_life: 138.51  
mean pain+wall: 1687.30  
  
Saved benchmark rows to /home/dennis-hedegreen/hedegreen_research/elliott/results/benchmark-results.csv  
  
The Making of Elliot  
1. Watch Elliot  
2. Run benchmark  
3. Inspect results  
0. Exit  
Choose node [1]:
```

Benchmark summary after 500000 steps and 10 runs per version.

What Happened

The baseline was clearly the weakest version. Its mean food rate was 0.0692. Its average steps per life was 91.65. The second version made the first real jump. With loop-breaking added, the mean food rate rose to 0.08045, and average steps per life rose to 134.7. That was the first meaningful improvement. But it came with a cost. Pain and wall events increased sharply, and randomized moves also became a major part of the system's behavior.

The third version improved food-seeking slightly again. With five neurons per sense, the mean food rate reached 0.08106, and average steps per life rose to 138.51. That is an improvement, but not a dramatic one. And again the cost stayed high. Pain and wall events were still far above the baseline.

What That Means

The result is not that Elliot simply became better. The result is that he became better at finding food by becoming more unstable at the same time. That is a more interesting result than a clean win would have been. It means the system is already exposing a tension: more escape from local traps, more food, longer lives, and more chaos. That is useful, because it means Elliot is no longer just moving. He is beginning to produce tradeoffs.

What Comes Next

This is not finished. This is the first real pass. Elliot will be updated from here. New versions will be tested. New benchmarks will be added. And the system will be pushed until it becomes clearer what kind of adaptation is actually emerging.

Like the bio log, this is meant to develop in public. Not as constant noise, but as visible checkpoints when something actually changes. That matters, because a project like this should give people a reason to return. Not for promises. For development.

Elliot is still very small. That is the point. Small enough to see. Small enough to break. Small enough to learn something from. This is the beginning.

— Dennis Hedegreen