

Water Was Already There

Neuron Lab did not begin as a software feature. It began as the missing room for understanding what a single neuron is doing before it disappears into a larger brain.

2026.03.24 07:13 DENNIS HEDEGREEN BUILD OPEN-NOTE

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/WATER-WAS-ALREADY-THERE/](https://hedegreenresearch.com/articles/water-was-already-there/)

I am lying in bed in a high-rise in Copenhagen. It is almost seven in the morning. I have been working all day and all night. The building is waking up. I can hear water moving through the pipes in the walls. Showers turning on. Taps opening and closing. Flow starting, stopping, building pressure, releasing.

That is when the idea locked into place.

The last Elliot article ended with a sharper question than the one I started with. More neurons did not help everywhere. Randomness helped, but it cost something. The system improved by becoming more unstable. And the honest next step was not to add more complexity, but to understand what the existing components were actually doing.

I could not answer that question properly. Not because the math was wrong. Because I did not have a place to take a single neuron apart and watch it work in isolation before putting it back into a brain.

That is the gap.

In the first Elliot article I wrote that the synapses were memristor-inspired. That the connection itself carries a local memory of previous consequence. That was true. But it was also still too vague. I knew the behavior I wanted. I did not fully understand the mechanism I was using to produce it. I could wire neurons together and watch behavior emerge, but I could not explain clearly, at the level of a single component, what was happening inside it.

Lying here listening to the building wake up, the answer became obvious.

Water already does what neurons do.

A bucket that fills until it overflows is a threshold neuron. A bucket with a hole in the bottom is a leaky integrator. It drains unless you keep feeding it. A syphon that accumulates silently and then dumps everything at once is a burst neuron. A pipe that gets wider the more water flows through it is a memristor. That last one is the direct formalization of what Elliot's synapses were already doing informally: connections that strengthen with use and weaken with absence.

None of these are new ideas. Some of them are among the most basic neuron models in computational neuroscience. But they are usually encountered through notation first. Water makes them visible first. You see it fill. You see it drain. You see it fire. And when you switch from water to math, the equation underneath is no longer abstract. You already understand the behavior. The notation just catches up.

That is what Neuron Lab is.

A visual workbench where neurons are built from containers, connected by pipes, and understood through flow before they are buried inside a larger system. Every component maps to known mathematics. The tool does not claim new theory. It translates existing theory into a physical language that is easier to think in, easier to test, and harder to misunderstand.

The concept document currently defines six neuron types and five pipe components. Buckets. Leaky buckets. Syphons in multiple variants. An absence detector. Channels drawn like Line Rider paths where shape determines delay, speed, attenuation, and natural gating. Splitters, mergers, canal-lock gates, and erosion pipes that learn without a separate training phase. The conservation rule is simple: water cannot disappear. Every leak is a pipe. Every output goes somewhere.

That point matters for Elliot directly.

When I go back to build his next brain, every component will have been tested in isolation first. I will know what a syphon does when it feeds a leaky bucket through a pipe that erodes with use. I will know it because I watched it. Not because I hid the uncertainty inside a larger simulation and hoped the behavior meant what I wanted it to mean.

Elliot is not a creature living in water.

Water was already there.

His thresholds, leaks, bursts, memory, and flow were already pointing in that direction. What was missing was the room where those things could be seen clearly.

The concept document is written. Implementation starts when I wake up.

[Neuron Lab concept document in the repository](https://github.com/DennisHedegreen/elliott-lab) (<https://github.com/DennisHedegreen/elliott-lab>).

Right now the building is full of running water, and I am going to sleep.

— Dennis Hedegreen

RELATION MEMORY

MATURES FROM [More Neurons Did Not Help Everywhere](#) · Moves from the failed capacity assumption into a smaller readable component bench.

RETURNS TO [The Making of Elliot](#) · Returns to the organism experiment by asking how its hidden mechanisms can become readable first.