PRAGMA RESEARCH

Minds that learn by living

Every AI system in use today stops learning before it ships. It is trained once, its settings are locked, and after that it can only repeat what it already knows.

Cadence brains keep learning the entire time they run, the way an animal does. They are small enough to sit inside a robot and learn that robot’s job while doing it.

Bernhard Mueller · floatingpragma.io · Seed round, $4,000,000 · September 2026

THE THESIS

Intelligence is what a body does in a world

No animal is trained and then switched on. It is born into a world, acts in it, is hurt by it, and changes because of it, every day, until it dies. There is no training phase and no moment when the learning stops.

Everything we want from a general intelligence follows from that arrangement. Coping with a situation nobody prepared for. Picking up a new job without a team first collecting thousands of examples of it. Keeping yesterday’s skill while learning today’s. Working again after being damaged.

A machine that cannot change after release does none of this. Making it bigger does not help, because the limit is in how it is built.

THE GAP

What "frozen" costs

Today's AI is trained in a datacenter. Its settings are then locked and copied out to wherever it is used. Nothing that happens afterwards changes it.

WHAT A LIVING MIND NEEDS	WHAT TODAY'S AI DOES
Learn while it acts	Learns once, then the settings lock
Learn from its own life, in order	Learns from shuffled examples gathered in advance
Keep yesterday's skill	Teaching it a new task can erase the old one
Run inside the body	Runs on rented servers, reached over a network

THE GAP

What that costs in the field

A robot arm goes into a factory with a locked-in controller. The grippers wear, the parts shift slightly, and the success rate slides. Fixing it means collecting fresh examples, shipping them to a datacenter, retraining, testing, and pushing an update to the fleet.

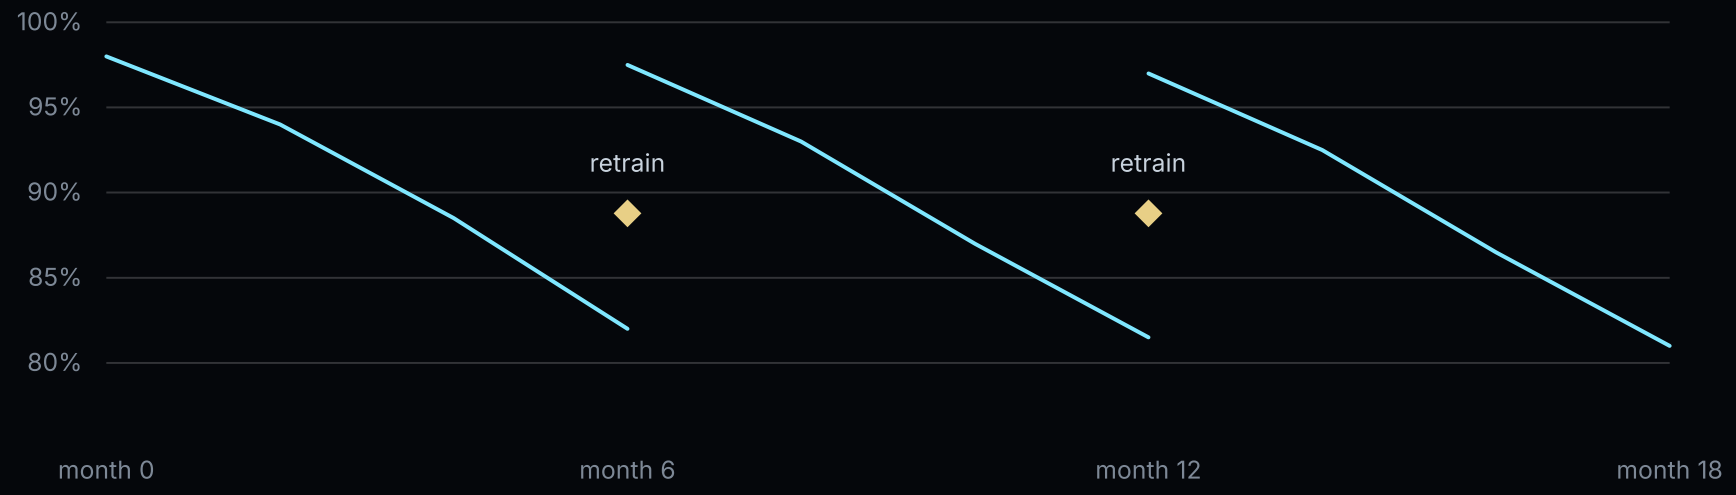

Illustrative. Shape of the problem, not a measurement.

That repair cycle takes weeks, and the factory’s data has to leave the building. A Cadence brain adjusts during the shift, and nothing leaves the machine.

THE RULE

Where the learning rule came from

Cadence started from a physics question. If the universe is made of small regions that each keep their own records, compare them with their neighbours, and fix the disagreements they find, what does that produce? We derived that it produces one shared, consistent world, with nothing in charge of it.

Turned into a learning rule, that says: each connection in the brain changes using only the activity of the two cells it joins, plus one broadcast signal saying whether things went well. Nothing ever looks at the whole network at once, so nothing has to store a trace of the computation and walk back through it. That is what removes the training phase.

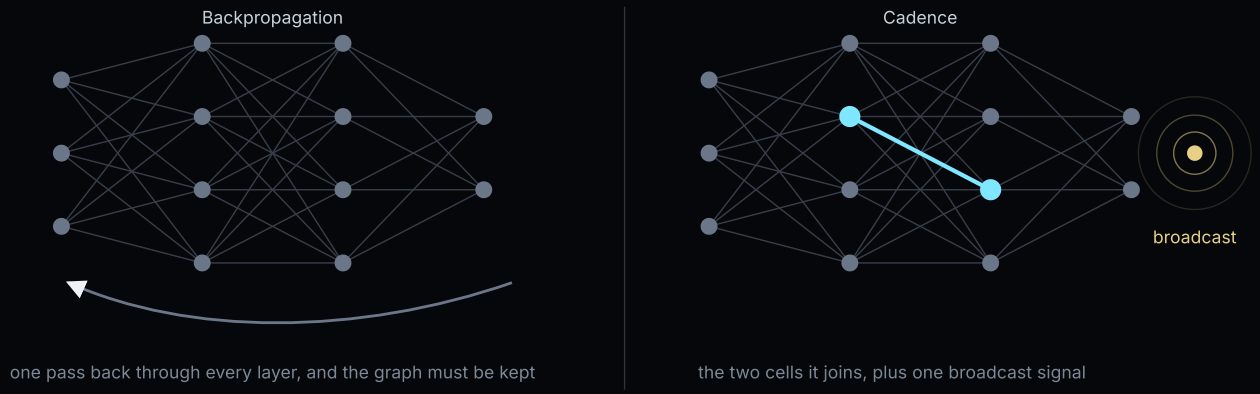

About the physics. It is formal proofs a computer checks step by step, like a compiler for mathematics, and more than 10,900 of them pass. The physics community has not accepted it, and this deck does not ask you to. The rule is judged on what it does.

THE OBJECT

How a Cadence brain works

A conventional network passes a signal through a stack of layers once, and an answer comes out the far end. A Cadence brain behaves more like a physical system. Its cells push on each other until they stop moving, and the state they come to rest in is the answer. We call that settling.

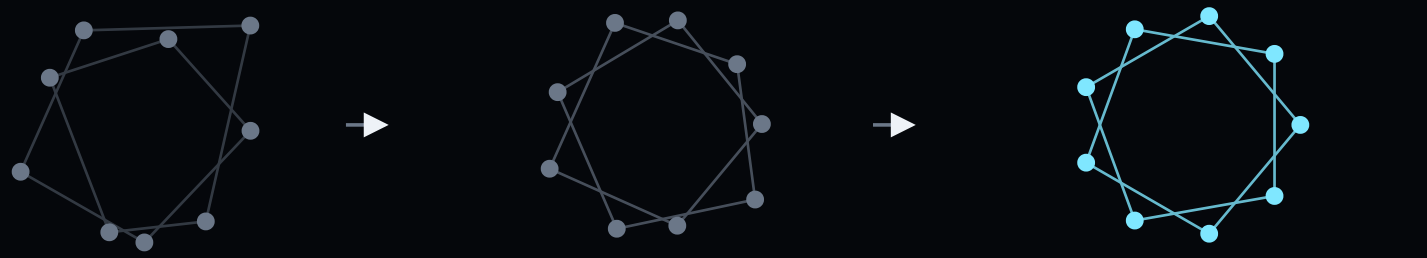

signals arrive — the cells push on each other — it comes to rest, and that state is the answer

- **Settling is the thinking.** Signals arrive, the network relaxes, and the resting state is the decision.
- **Memory is written once.** Something that happens is recorded immediately into the few cells that were active when it happened, instead of being repeated hundreds of times until the network absorbs it.
- **Imagining is settling in private.** The brain nudges itself, lets a possible future settle, reads off how it turned out, and discards it. Imagined outcomes are never recorded as things that happened.
- **The wiring can be evolved.** The layout is inherited and mutated across generations rather than drawn by hand.

WHAT EXISTS

Built, running, and open to inspection

The library is public and installs in one line. Anyone can download it and check the claims.

Demonstrations that run in a browser, each showing the whole brain beside the task it is doing: a robot arm, a Connect Four player, a composer, an artist.

A Connect Four brain beats AlphaZero, the method that taught itself to beat the best chess and Go programs, trained here with its published recipe. Cadence wins 92 and 99 of 100 paired games against two configurations of it.

Brains built from real biological wiring. A map of a fruit fly's brain containing 138,639 nerve cells, and the complete 302-cell nervous system of the roundworm *C. elegans*, both loaded and run as Cadence brains. The worm learns to associate a smell with food.

Every number in our paper is read out of a result file whose fingerprint is recorded, so a reader can trace it back. The paper also names ten things the design does badly.

One person built all of this, self-funded, at $5,000 to $10,000 a month.

EVIDENCE

What it does better, and why that matters

Arm reaching targets it was never shown

99.4%

against 98.0% for the standard method, which needs a stored buffer of past attempts

Choosing what to remember, four memory tasks

94.6%

against 49.0% for a conventional model given more memory than Cadence had

Updating a 138,639-cell brain after a small change

6.3×

faster than redoing the whole thing, because effort follows what changed

Function recovered after part of the brain is destroyed

85%

against 50%. The rule that learned the skill also repairs it

Practice needed to learn to balance a pole

40k

attempts against 137,000, for the same skill

Every number here is read out of a result file in our paper whose fingerprint is recorded, so a reader can trace it back to the run that produced it.

EVIDENCE

What it does badly, and why we are showing you

WHAT WAS MEASURED	RESULT	WHAT IT MEANS
Time to make one decision	13 milliseconds, against 1 for a conventional network	Settling costs more than passing through layers. Too slow for a fast control loop today.
Predicting the next character of text	Behind a small transformer, the design behind ChatGPT, at every size tried	Language is not where this design wins. We are not building a chatbot.
Learning from a reward that arrives much later	Well behind the standard method for learning by trial and error	Long chains of cause and effect are unsolved here.
Holding many separate skills at once	Loses the older ones unless given a stored buffer	The weakness most likely to sink the approach.
Size of the largest brain built	70 million connections	A frontier language model is roughly 10,000 times larger. Nothing here shows this design scales that far.

Each of these is named in our own paper with its measurement. Every work package in this plan attacks one of them. A company that showed you only the previous slide would not survive a technical review.

EMBODIMENT

The ladder of bodies, and what each rung tests

Ten cheap arms run in parallel, because what is being measured is what happens over a working lifetime, and one arm gives one lifetime at a time.

THE HARDWARE IDEA

Why this design points at unusual hardware

A Cadence brain thinks by letting a system relax until it stops changing. On a normal computer we have to **imitate** that relaxing, by grinding through arithmetic over and over until the numbers settle. That is the main reason one decision costs 13 milliseconds instead of 1.

Some physical things relax by themselves. Light bouncing inside a shaped mirrored cavity falls into a stable pattern in the time it takes light to cross it, and no arithmetic happens at all.

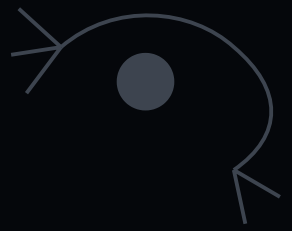

An animal
voltages across nerve cells
about 20 watts

A computer
numbers in memory, in a loop
arithmetic, and 13 ms

An optical cavity
a light field
the time light takes to cross

Our mathematics says which shapes settle into which computation. That is what connects the brain design to the hardware.

THE HARDWARE IDEA

The Echosahedron

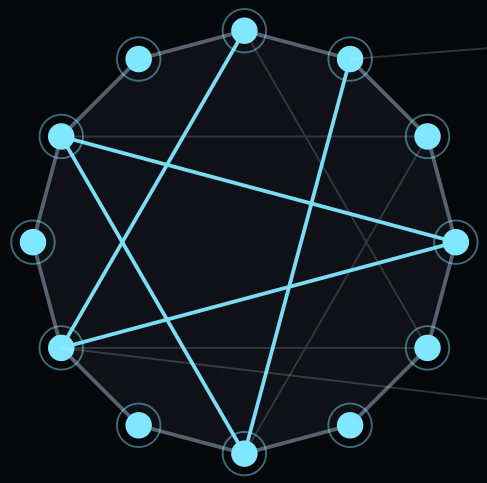

LED shining in
forward biased

Twelve windows
one at each vertex, each one a port

A clear body
light bounces and settles into a pattern

the same LED reading out
reverse biased, measuring the light

A thing with a boundary, a set of ports, an inner state and a way to read it. Which is the same description as one region of a Cadence brain. Connect several bodies through their windows and the assembly is a brain running in light.

Status, stated plainly. The mathematics exists and the design exists. **No device has been built.** This plan prints the bodies, builds one chamber, and measures whether it computes anything useful, against control experiments designed to catch us fooling ourselves. If it does not, we publish that and continue with the software version, which is what the rest of the plan rests on.

THE PLAN

Eighteen months, and when each answer arrives

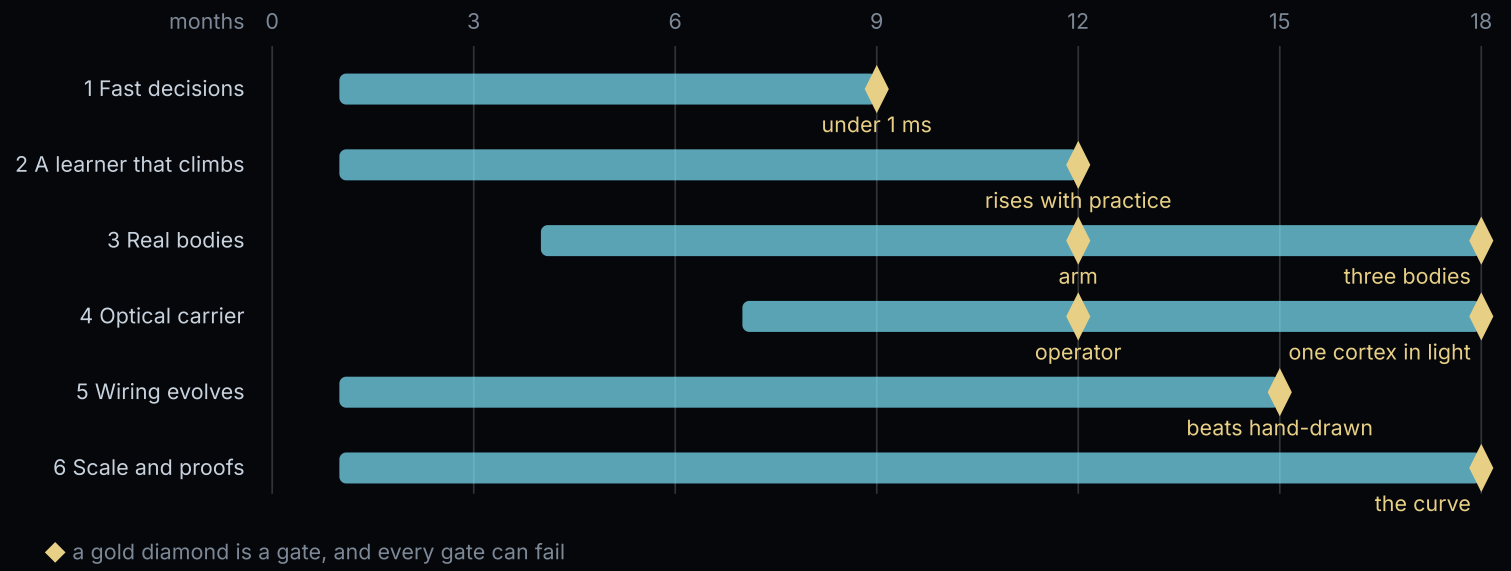

Nothing waits on everything else. The work that makes decisions fast runs first because every body depends on it, and the optical bench starts once its design is frozen.

THE PLAN

Six work packages, every one able to fail

	WORK PACKAGE	WHAT "DONE" MEANS	IF IT FAILS
1	Make decisions fast enough for real bodies	One decision in under 1 millisecond, on a chip small enough to sit inside a robot	Brains supervise a fast controller instead of being one
2	Make the brain keep improving as experience accumulates	Performance rises with practice instead of flattening, measured over a million decisions	Ship the one-shot memory alongside a small conventional learner, and say so
3	Put brains on real bodies	An arm learns a second job during operation and keeps the first. Then a hand, then a legged robot	Narrow the scope to arm and hand
4	Test the optical hardware	A built chamber computes something repeatable and useful, against controls	Publish the negative result. The software path does not depend on it
5	Let the wiring design itself	A layout found by evolution beats our best hand-drawn one on a body it never saw	Keep hand-drawn wiring and report the idea as unproven
6	Find out whether it scales	Publish the curve from 70 million connections upward, whatever shape it has	Publish the flat curve

THE PLAN

What exists after eighteen months

1. **A brain that learns on three real bodies.** A robot arm, a dexterous hand and a legged robot, each picking up a new job while working and keeping the one it already had, with no retraining run and no data leaving the machine.
2. **Software that robot makers can install**, fast enough for a control loop, with a guaranteed worst-case decision time and a command that returns the robot to its original behaviour.
3. **An answer on optical hardware.** A chamber built and measured, and one brain region run in light and compared against the same region in software. A negative answer is published as readily as a positive one.
4. **A wiring diagram found by evolution** that beats the best one we drew by hand.
5. **Version 1.0 of the open library**, stable enough to build products on.
6. **The papers and the proofs**, with every number traceable to the file it came from.

DEFENSIBILITY

Why this is hard to copy

The learning rule was derived, not tuned. It fell out of an argument about consistency and was then checked by machine. It cannot be read off a results table and reproduced, because the reasoning is what produces it.

The optical path needs both halves. A photonics lab can build a cavity and has no theory saying what the cavity computes. An AI lab has the computers and will not build cavities. The theory that joins them is ours.

The open library is the distribution. It is already public, with the measurements and the proofs attached. Adoption compounds while the paid work sits on top of it.

Everything is checkable without trusting us. The proofs run on anyone’s computer. Predictions were published and time-stamped before the measurements arrived, so they cannot be quietly adjusted afterwards. Ten weaknesses are named in our own paper, and there is a public prize for disproving the framework.

BUSINESS

Where the money comes from

The library stays free and open. It is how developers find us and how outsiders check the work. It is not the product.

The product is a finished brain for a particular job. A wiring diagram, a training curriculum and a tested runtime for, say, a warehouse picking arm. A robot maker licenses that per robot. Building one takes us months and would take them years, and that is what they are paying to skip.

The first paying use is fleets already installed. A robot whose success rate is drifting down as its hardware wears. Our brain rides on top of the maker's existing controller and closes the gap during operation, without altering the controller that was already approved for safety.

Optical hardware is a second business and exists only if work package 4 returns a number.

Roughly 4.66 million industrial robots are working worldwide, and about 542,000 new ones were installed in 2024. No revenue is assumed inside the 18 months.

TEAM

Six hires, no managers, and a back office run by AI

ROLE	STARTS	WHY THEN
Principal investigator (B. Mueller)	month 1	Leads the research and writes code with the team
GPU and systems engineer	month 1	Work package 1 gates every body that follows
Machine-learning researcher	month 1	Work package 2 is the long pole
Robotics researcher	month 4	Once there is a simulator to work against
Photonics and hardware engineer	month 7	Once the chamber design is frozen
Operations lead	month 1	Company setup, payroll and reporting, run with AI tools
Marketing and developer relations	month 7	Once there is a robot to show
Formal proofs, half time	month 7	Proofs follow the mechanism

Bernhard Mueller. Pwnie award 2009, the security field's research prize. Author of the OWASP MASVS, the mobile security standard Google Play's independent app review certifies against. Built Mythril, the most widely used security analyser for Ethereum contracts. Speaker at DEF CON 27 and HITB. Thirty years of changing field every five years, from hardware to reverse engineering to mobile to neural networks to blockchains to physics.

BUDGET

$3,739,232 over 18 months

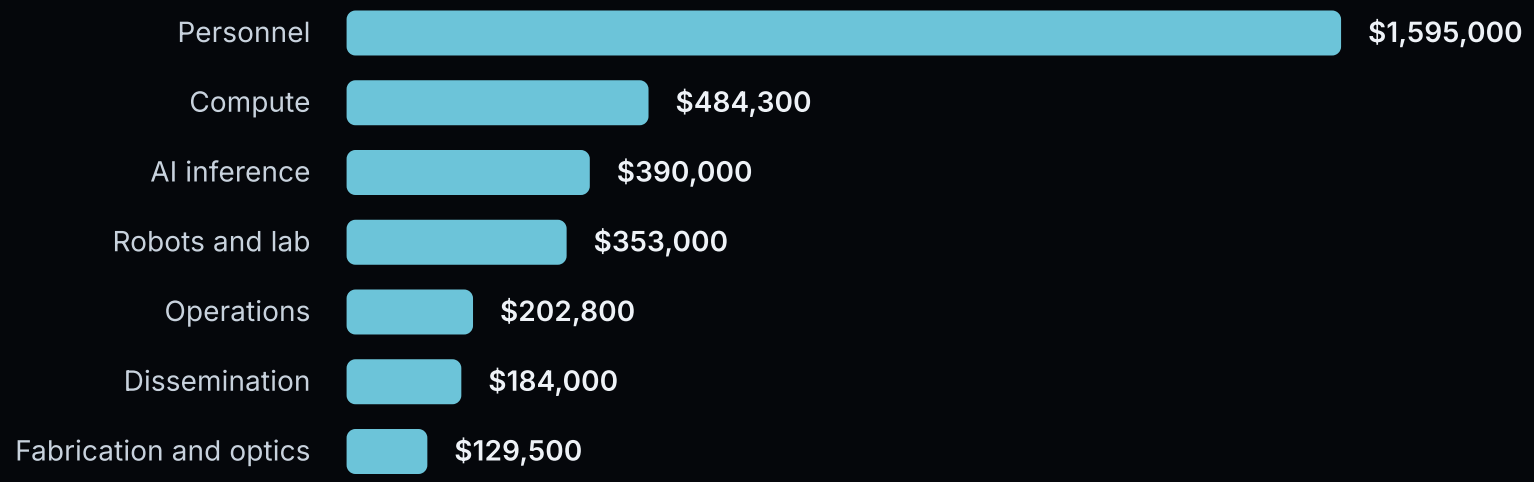

Contingency, 12 per cent	$400,632
Total	**$3,739,232**

Every price was checked in September 2026 and carries its source. The 3D printers are bought rather than outsourced because the print quality and surface finish of an optical chamber are part of what is being measured.

THE ASK

$4,000,000 for 18 months

Six hires, four robot bodies, the optical bench, and the six work packages on the previous pages. Smaller versions of the programme exist, and which work packages a smaller round buys is a conversation rather than a table.

The code, the proofs, the data and the papers are released openly. The finished brains and the runtime are the company.

bernhard@floatingpragma.ai · floatingpragma.io · github.com/muellerberndt/cadence