braincore

A brain that learns to live.

A complete human brain in software, built region by region from published neuroscience. It will see, hear, remember and learn from its own experience. No language model, no pretrained weights, no backpropagation.

INVESTOR PRESENTATION · SEPTEMBER 2026

WORKING TODAY · AREA 41 REAL RUN

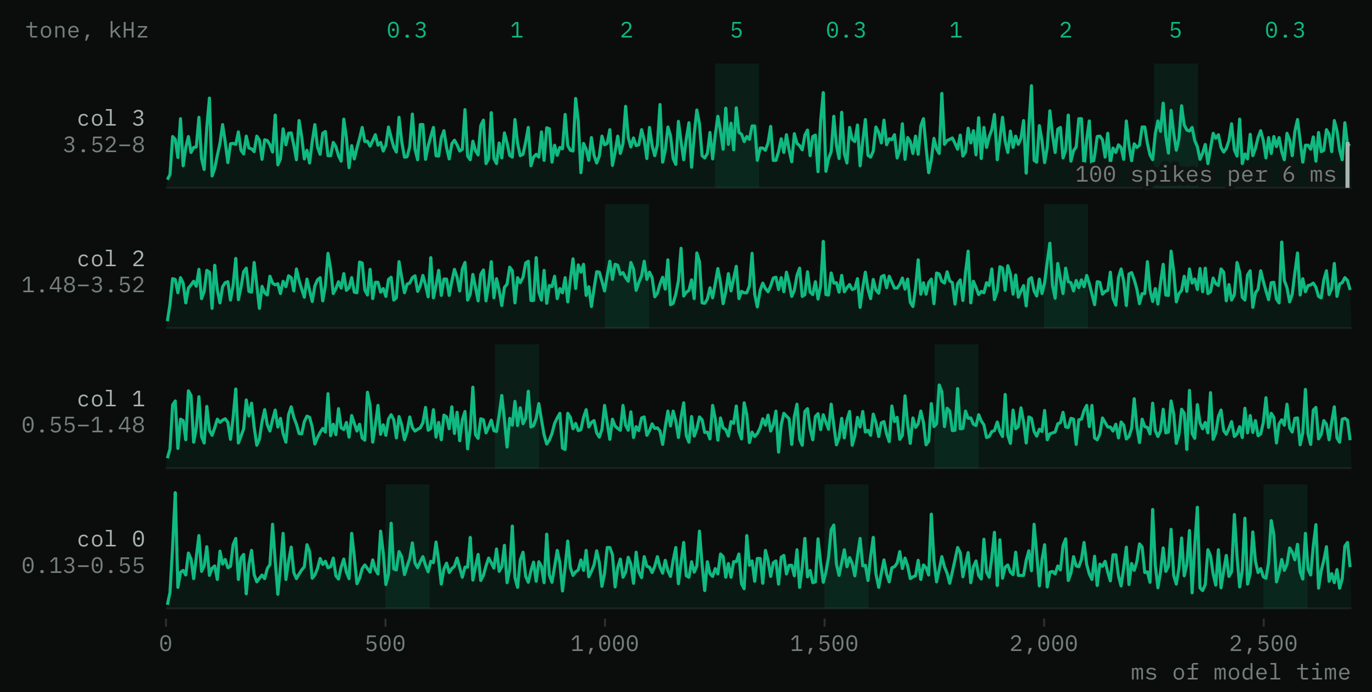

Four tones in a model ear; four columns of auditory cortex answer, each in its own band. 31,896 neurons, 2.7 s of model time, 4.5 s on a laptop.

SHA-256 a5cd7566fe4e387315873aa629670a6406415e36783f48d58f19372c83011a17

Vitalii Cherepanov · vitalii@vbcherepanov.com · getbraincore.com

Today’s AI read everything once, and stopped learning.

Frozen

Trained once on a fixed corpus by backpropagation. To learn from its own experience it has to be retrained.

Second-hand

Knows what people wrote, not what it perceived. Speaks with the same confidence whether it knows or not.

Opaque

Its knowledge is spread over billions of weights that no one can check part by part against anything.

The organ it imitates works the other way round. So we build the organ.

Not a rival to ChatGPT. A different kind of machine.

	LARGE LANGUAGE MODEL	BRAIN AND BRAINCORE
Made of	One network of one kind, trained to predict the next word	Regions of an organ, each from its own published model
How it learns	Output error sent back through every weight, in a separate training phase	Each synapse changes from the activity of its two neurons; learning is working
When	Before release, then frozen	All the time, from its own senses
Knowledge from	Text written by people	Light and sound it perceived itself
Memory	Weights and a context window	Hippocampus fast, cortex slow, sleep between them
Checked	As a whole, on benchmarks	Part by part, against papers and digests
Runs on	GPUs in a data centre	Ordinary CPUs, on the customer’s own servers

An LLM is stronger at everything to do with language. BrainCore today tells apart four letters and four tones.

How it will learn: the brain's own scheme.

1 · Senses

Eye and ear turn light and sound into spikes.

`works` Paper 1

2 · Cortex

Lays the stimulus out over its columns.

`works` Letters 94 %, tones 20/20 seeds

3 · Hippocampus

Stores what happened together, after one showing.

`in progress` CA3 stores a pair; CA1 red, 17/23

4 · Neuromodulators

Dopamine marks what mattered; acetylcholine opens plasticity.

`in progress` Nuclei 18/18; dopamine route 8/11

5 · Sleep

Replays the new; the cortex fits it in with the old.

`planned` Sleep switch 19/21; transfer planned

6 · Paper 2

Letter ↔ sound through the hippocampus; plasticity off must break it.

`next` Preregistered

Hippocampus learns fast, cortex slowly, replay fits the new in with the old: McClelland, McNaughton & O'Reilly 1995

Will the brain have the same problems? Some, yes.

PROBLEM	LLM	NATURAL BRAIN	BRAINCORE
New erases old	Yes, under fine-tuning	Gradually, not catastrophically	Same risk, the brain's fix copied; capacity and interference curves in Paper 2
Makes things up	Rewarded for guessing	False memories exist	Possible; recall measured against a margin fixed in advance, an untrained letter must recall nothing
Bias	The corpus's, unlistable	Its own experience	Its own, but recorded in full
Experience needed	Internet-scale text	An episode at once, concepts over years	One showing per pair, but slower than real time: 2.7 s in 4.5 s
Repeatability	Not even at temperature 0	None	Bit-exact on recorded input
Cost of scale	$78–191 M to train	86 B neurons	A laptop today; thousands of servers at human scale

Luo et al. 2023 · OpenAI 2025 · Thinking Machines 2025 · French 1999 · Loftus 2005 · AI Index 2024

THE GOAL

What the finished brain will do.

It will see

Objects, faces and scenes, with a visual cortex of its own.

`today` Reads 4 letters at 94 %, 20 of 20 seeds

It will hear

Voices, words and melodies, from the cochlea to auditory cortex.

`today` Pitch in its own column, 20 of 20 seeds

It will remember

An episode after living it once; the whole from a fragment.

`in progress` CA3 stores a pair after one showing; CA1 read-out red, 17 of 23

It will learn all the time

From its first moment: plasticity, dopamine for what mattered, sleep to consolidate.

`in progress` Plasticity, dopamine route, sleep switch built; the test is Paper 2

It will speak its own words

Broca's and Wernicke's areas of its own; language learned, not borrowed.

`planned` Fields 44, 45 and 22 mapped in the atlas

It will run on spikes

The native signal of neuromorphic chips. Human scale on CPUs would take thousands of servers.

`in progress` 89.7 M synaptic events/s on a laptop, bit-exact

Future tense: the goal. Present tense: measured today.

THE PATH

Five stages to a brain that learns. The first is done.

WE ARE HERE ↓

Stage 1 · Sept 2026

Senses

It sees and hears. 23 acceptances, Paper 1 frozen and reproducible.

done

Stage 2

Live senses

Camera into the retina, microphone into the cochlea; it reads what you show it.

next

Stage 3 · Paper 2

Learning

Letter ↔ sound learned through the hippocampus; plasticity-off must break it.

next

Stage 4

Scale

A brain across machines, the same digest as on one; millions of neurons.

later

Stage 5 · the goal

Concepts and language

Cat, meow and CAT on one loop; its own Broca and Wernicke.

later

Built like the organ. Checked like a science.

Knowledge lives in synapses
No fact tables, vectors or dictionaries inside a region.

The connectome is data
Every number in a file, with its source and a label.

Every region is graded
An acceptance against its paper, with controls that must break it.

Only spikes between regions
Each region is a service; one process or many, the same result.

One clock, one digest
SHA-256 over every spike; the same bits on 1 core or 12.

Nothing borrowed
No pretrained model, no embeddings, no LLM, ever.

LIGHT	**Retina** 24 × 24 ON/OFF	**3 × 3 tiles** retinotopy	**Area 17** 9 columns	**Examiner** outside the brain
SOUND	**Cochlea** 32 channels	**Cochlear nucleus** 4 × 96	**Area 41** 4 columns	**Examiner** outside the brain

It sees: the shape of a letter survives into cortex.

15 / 16

test showings named by a decoder outside the brain

ρ 0.866

each column answers the ink in its own patch of the visual field

179,298

neurons in the visual network; plasticity off

А

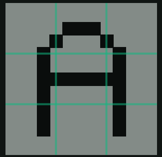

158	492	-139
671	437	672
242	102	385

О

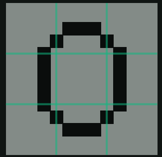

150	625	141
529	413	571
218	681	218

Т

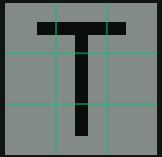

270	832	136
74	530	-19
11	365	-37

И

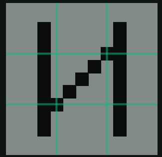

484	-15	424
639	71	657
435	113	531

evoked L4E spikes per showing

-832 0 +832

It hears: pitch lands in its own column.

20 / 20

1 kHz vs 2 kHz, leave-one-out, decoder outside the brain

349–565

spikes above baseline in each tone's own column

31,896

neurons: cochlea, nucleus and four columns of area 41

Model

	col 0	col 1	col 2	col 3
0.3 kHz	349	11	10	79
1 kHz	50	563	81	23
2 kHz	-14	13	461	84
5 kHz	-8	-39	69	565

Shuffled wiring

	col 0	col 1	col 2	col 3
0.3 kHz	94	73	61	133
1 kHz	124	206	161	156
2 kHz	117	100	168	68
5 kHz	153	151	114	101

Organisation, not activity.

Shuffle the wiring: total spikes change by less than 0.01 % in vision and 0.4 % in hearing, yet the decoder falls from 94 % to 56 % and from 100 % to 60 %. The information is in how the brain is wired.

Vision · four letters

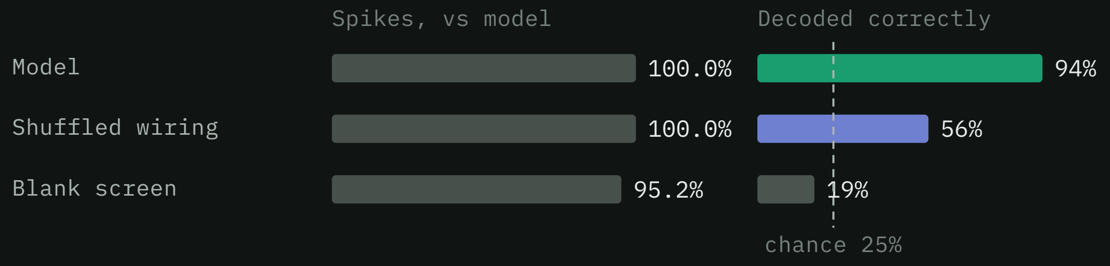

Hearing · 1 kHz vs 2 kHz

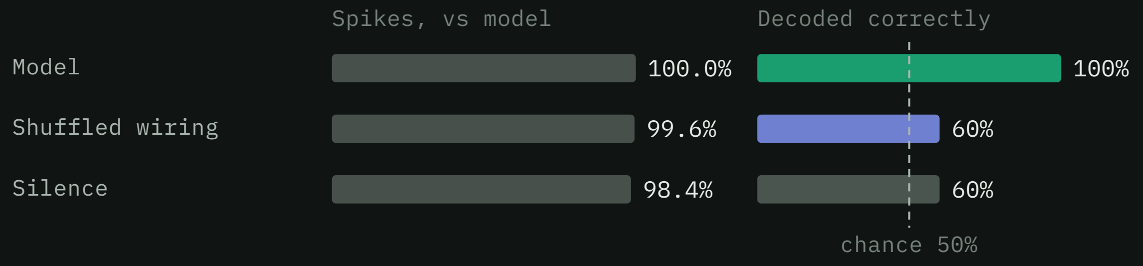

Rules written before the run. Failures published.

20 / 20

new seeds pass the vision rule committed to git before they ran

13 / 20

hearing repeat rule needed 18: failed, and reported as failed

16 / 16

digests reproduced from a clean copy of the committed tree

Four letters: decoder accuracy on each of 20 seeds

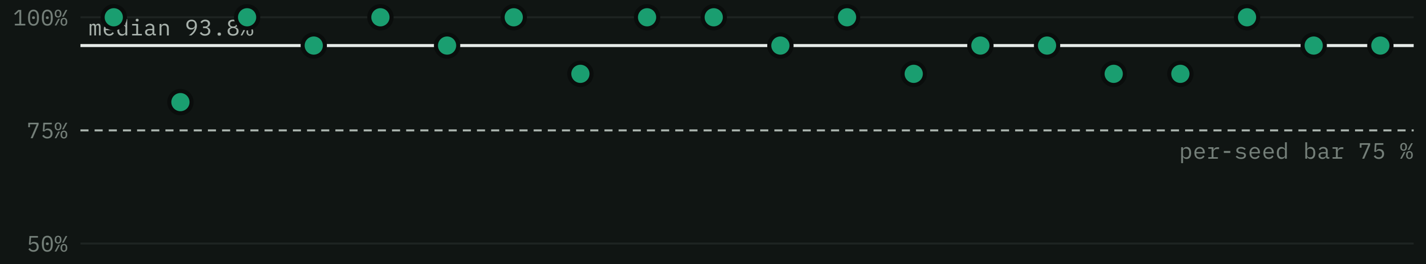

Hearing: does every tone win the same column in ≥ 90 % of showings?

13 / 20 seeds pass · rule needs 18 of 20: failed

THE ACCEPTANCE LEDGER

23 acceptances against their papers. 12 red, and left red.

Cortical column, 1 mm^2	7/7	✓
Sheet of 2 × 2 columns	pass	✓
Retina	pass	✓
Cochlea	19/20	×
Geniculate relay (LGN)	6/7	×
Basal ganglia: action selection	5/7	×
Neuromodulator nuclei: VTA, LC, raphe, basal forebrain	18/18	✓
Cerebellum, 90,918 cells	pass	✓
Basolateral amygdala	14/16	×
CA3 autoassociation	red	×
Hippocampal loop EC → DG → CA3 → CA1	17/23	×
Sensorimotor loop	pass	✓

V1 column	≤ 3 %	✓
Area 17, nine columns	5/5	✓
Vision: retina → area 17	12/12	✓
Hearing: cochlea → nucleus → area 41	7/7	✓
Thalamus: relay and reticular nucleus	4/7	×
Dopamine route VTA → basal ganglia	8/11	×
Sleep–wake switch	19/21	×
Grid cells, entorhinal cortex	11/11	✓
Dentate gyrus	7/8	×
CA1	6/10	×
Tissue: fixed-size synapse file	50/51	×

Red: at least one graded criterion missed. Nothing was tuned after a run.

The map is the human brain: 79 services.

Built, all criteria green · 15 Built, at least one criterion red · 13 Planned · 51

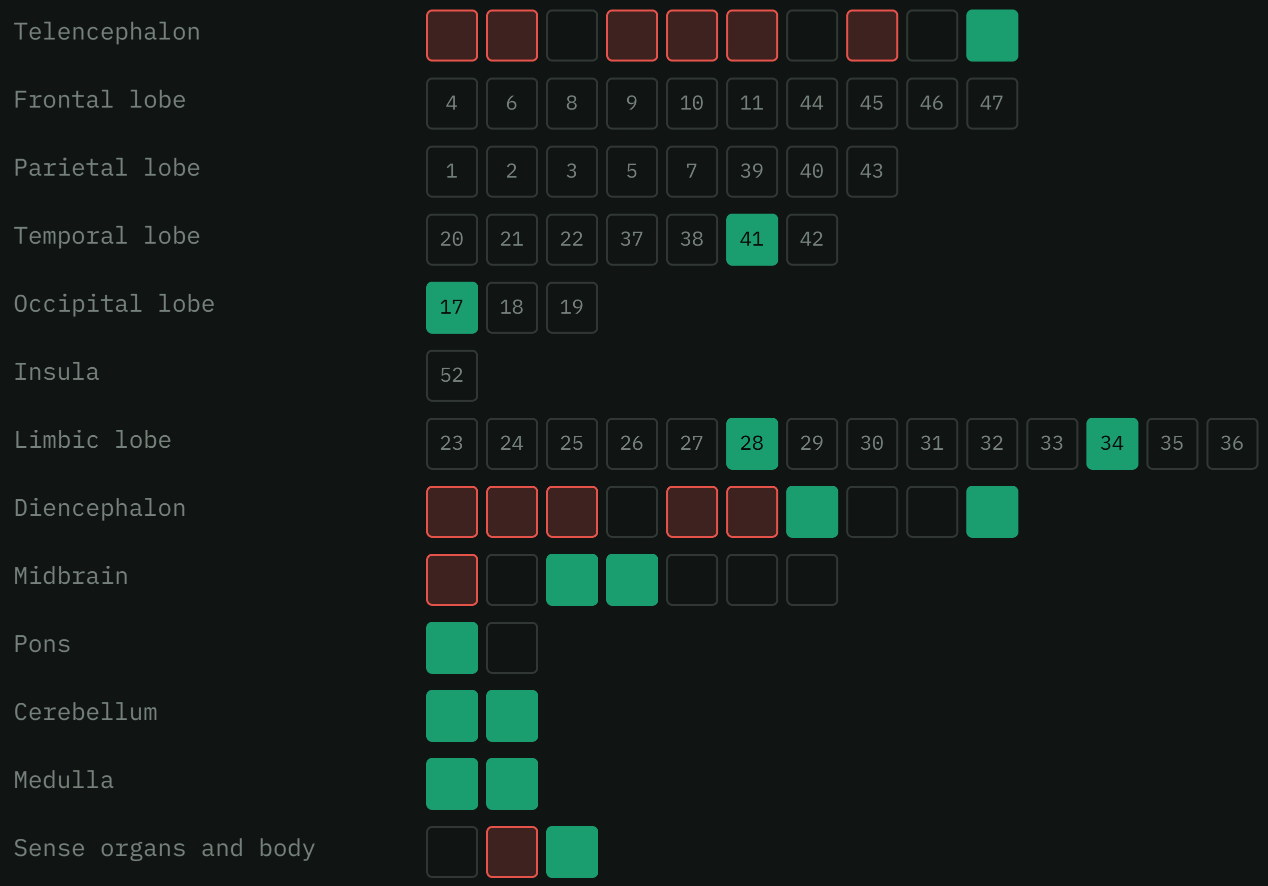

Six divisions · six lobes · 43 human Brodmann fields · every nucleus a service

Fast on a laptop. Honest about the distance.

89.7 M

synaptic events per second, full cortical column, 12 cores

×2.2–2.9

lossless packing of synapses in a fixed-size tissue file

3.8 M

neurons with 8,000 synapses each in one 150 GiB tissue file

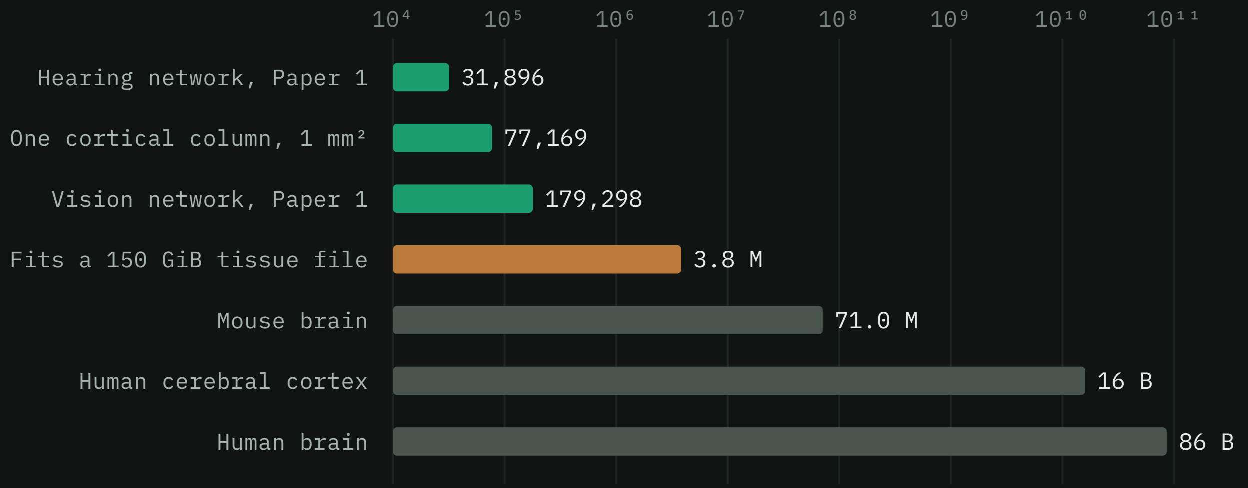

Four to five orders of magnitude to a human brain. That is why every region is a service: the brain spreads over machines without changing a spike.

Four facts from the last two years.

Hardware for spikes exists	Intel Hala Point: 1.15 billion neurons, 128 billion synapses, 2,600 W maximum. Intel, April 2024
The largest reconstruction programme handed over	Blue Brain, 2005 to December 2024, continues as the non-profit Open Brain Institute. EPFL · EurekAlert, 2025
Buyers pay for brain-like compute	Cortical Labs CL1: 800,000 human neurons, $35,000 a unit. IEEE Spectrum, June 2025
Europe funds the category	SpiNNcloud: €10 M blended EIC Accelerator funding. EBRAINS, June 2025

Two markets: one still being defined, one already paying.

$1.3–20.9 B

neuromorphic computing in 2030; two forecasts differ 16×

MarketsandMarkets · Next Move

$6–7.2 B

in-silico drug discovery by 2030, from $3.3–3.6 B in 2024

Market Glass · Next Move

$5.93 B

in-silico clinical trials by 2031, from $3.87 B in 2025

Mordor Intelligence

$2.56 B

average cost of one new drug, 2013 dollars

DiMasi et al. 2016

Since 2022

US law counts computer modeling as a nonclinical test of a drug

Public Law 117-328, sec. 3209

$2.7 M

Gates Foundation to Numenta in 2024 to test brain-based AI

Semafor

Analyst forecasts, quoted from their pages; every quote in SOURCES.md.

INDUSTRIES

On the customer's own servers, not in the cloud.

Ordinary processors, no GPU, no outside calls: patient and production data never leave. Proven on a laptop; enterprise packaging is not built yet.

INDUSTRY	TASK	WHEN	NEEDED FIRST
Pharma	Neural circuits in disease and under drugs: basal ganglia and dopamine, hippocampus, serotonin and noradrenaline nuclei	2027–29	Partner, validation on drug and patient data
Medicine, medtech	Sound coding for hearing implants, vision for prostheses, teaching models	2027–29	Manufacturer partner, clinical validation
Neuroscience	Graded region models, reproduction by digest, teaching labs	2026–27	Documentation for outside users
Neuromorphic chips	Brain-grade reference workloads	2027–28	A first region on a chip
Industry, robotics	Learn a new object, defect or sound on site	2028+	Paper 2, live senses
Regulated sectors	Every decision logged and repeatable; AI Act logging for high-risk systems from 2 Dec 2027	2028+	Useful decisions, conformity assessment

Not yet deployed in any industry.

The core is open. The brain built on it is paid.

WHEN	SOURCE	WHO PAYS	FORM
Now	Grants, joint projects	Funders, EU programmes	Non-dilutive: EIC Pathfinder up to €4 M, ERC Starting up to €1.5 M
2027	Regions and support	Labs, institutes, R&D	Open core; certified and new regions, support paid
2027	Installation at the customer	Pharma, medicine, industry	Annual site licence and support; data stays inside
2027–28	Custom models	Pharma, medtech	Build and certify a circuit for the customer’s question
2027–28	Hardware partnerships	Neuromorphic chip makers	Porting contract, workload licence
2028+	Learning core for devices	Robotics, device makers	SDK licence or per-device royalty

Pre-revenue. No prices until the first paying user.

Others build a tool, a chip or a model. We build the organ.

WHO	WHAT IT IS	SUBSTRATE	RELATION TO BRAINCORE
LLM labs	Pretrained transformers	GPU	Trained once, frozen. BrainCore uses none.
NEST, Brian 2	Spiking simulators	CPU, HPC	Tools, not a brain. NEST is our reference.
Blue Brain → OBI	Detailed rodent tissue; non-profit since 2025	HPC	We target the human atlas with point neurons.
Thousand Brains (Numenta)	Sensorimotor learning from a cortex theory	Software	Theory first; we are anatomy first.
Intel Hala Point, SpiNNcloud	Neuromorphic hardware	Silicon	Chips for spikes; they need models like ours.
Cortical Labs	Human neurons on a chip	Biology	Cannot be copied, versioned or audited to the bit.
BrainCore	The human brain in software, graded region by region	CPU today; spikes	Learns only from experience. Bit-exact. Preregistered.

What it gives, and what it cannot do yet.

PROS

- Learns on site, no retraining `goal`
- Knows only what it saw; origin of knowledge on record `design`
- No third-party training data `design`
- Bit-exact: 16 of 16 digests from a clean copy `measured`
- Checkable part by part: 23 acceptances `measured`
- On the customer's servers: CPUs, no GPU, no cloud `measured`
- Learned tissue copied and versioned as a file `design`

CONS

- Does not speak or write; tells apart 4 letters and 4 tones
- Learning in V3 not shown yet (Paper 2)
- Cannot take knowledge from books; learns from experience, slower than real time
- 179,298 neurons against 86 billion
- No energy saving on CPUs; not yet run on a spiking chip
- 12 of 23 region models still red
- No disease-data validation; no revenue; one founder

Measured: in an acceptance report. Design: in the architecture. Goal: to be tested.

MILESTONES FOR THE ROUND

24 months, a checkable result every quarter.

MONTH	MILESTONE	DONE WHEN
3	Live senses	Camera and microphone read letters and tones; files reproduce the fixture digests
6	Paper 2: letter ↔ sound through the hippocampus	Loop at 23/23; preregistered rule holds on 20–50 seeds; three ablations break it
9	The thalamus in the sensory paths	LGN relay inside the measured window; SOC, IC, MGN as regions
12	A brain across machines	Same digest on four machines as on one; ≥ 1 million neurons
15	Shape and action	Areas 18/19 read letters off-centre; area 4 writes one back
18	First concept learned from experience	Object, sound and written name joined on one loop, ablations break it
21	Half the atlas	40 of 79 services pass their papers (15 today)
24	Open platform and partners	Core released; three design partners; seed-ready

Targets, not promises: each acceptance is written before the work starts.

THE ASK

Two ways in.

PRE-SEED · EQUITY

€500 K 24 months

- One founder, with external scientific review and independent replication in the budget
- Stages 2–4 and the first design partners
- SAFE or priced; terms with the lead

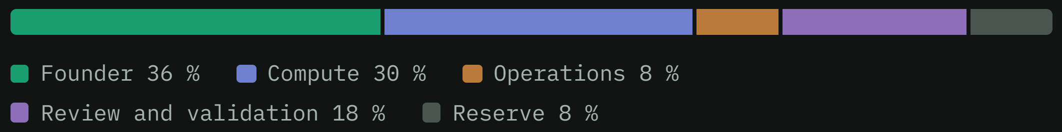

RESEARCH GRANT

€150 K 12 months

- Paper 2: an arbitrary sight–sound association learned through the hippocampus
- Preregistered, three ablations, independent replication
- PI €66 K · consultant €30 K · compute €30 K · publication €10 K · replication €8 K · admin €6 K

What could go wrong, and what we do about it.

Learning not yet shown in V3	Paper 2 is the test, with a plasticity-off control; published either way.
Scale and energy: 179,298 neurons vs 86 billion	One million neurons on four machines by month 12; human scale on CPUs means thousands of servers, so spiking chips matter.
Fidelity is not intelligence	Each stage is a publishable result; labs pay for fidelity before capability.
Published parts do not always compose	The relay transmits 0.124 vs 0.30–0.62 measured: found early, diagnosed, published.
One founder	External review is in the budget; everything reproduces from tags without the founder.

FOUNDER

Vitalii Cherepanov

"I'm an engineer with fifteen years of shipping production software. I build the organ, not the imitation: region by region, from papers neuroscientists already wrote, each checked against its paper. Thresholds before runs, controls that must break, red results left red. A brain that cheats on its own exams is worth nothing."

vitalii@vbcherepanov.com

Ask for a live run: the sensory paths reproduced to the digest on a call. Acceptance reports and code access for due diligence on request.