RESEARCH PROGRAMME · €150,000 · 12 MONTHS

Learning a cross-modal association in a spiking brain assembled from published models

APPLICANT	LOCATION	DURATION	CONTACT
Vitalii Cherepanov · BrainCore	Novi Sad, Serbia	12 months	vitalii@vbcherepanov.com

1 Summary

BrainCore is a spiking-network brain assembled region by region from published models, in which every region is a separate service that exchanges only spikes and every region is graded against its paper before anything leans on it. Paper 1 (September 2026) validated a visual path (retina → area 17) and an auditory path (cochlea → cochlear nucleus → area 41) with plasticity off. This programme asks the next question: **can such a brain learn an arbitrary association between what it sees and what it hears, through its own hippocampus, from paired experience alone?** The experiment is preregistered, uses three ablations that must abolish the effect, runs over 20–50 seeds from certified tissue, and is reproduced by an independent researcher. The result is published whether it passes or fails.

2 Background and gap

Data-driven spiking models of cortex reproduce the statistics of spontaneous activity (Potjans & Diesmann 2014; Schmidt et al. 2018), and several have been driven by sensory input and compared with recordings (Arkhipov et al. 2018; Billeh et al. 2020). Functional large-scale models couple perception to behaviour (Eliasmith et al. 2012). Hippocampal models store and complete patterns in isolation (Guzman et al. 2016; Cutsuridis, Cobb & Graham 2010). What is missing is a composition in which each sensory path has been validated on its own, the hippocampus sits between them, and an association has to be learned by local plasticity and measured against criteria fixed before the run. Reproducibility of spiking simulations is fragile (Pauli et al. 2018) and preregistration guards against thresholds that move after the data are seen (Nosek et al. 2018); the programme applies both.

3 Preliminary results

- Vision: four letters decoded in 15 of 16 test showings; retinotopy $\rho = 0.866$; preregistered rule passed on 20 of 20 new seeds (median 93.8 %).
- Hearing: 1 kHz vs 2 kHz 20 of 20; tonotopy on 20 of 20 seeds; the presentation-level repeat rule failed (13 of 20 against 18) and is reported as failed.
- Hippocampal loop EC → DG → CA3 → CA1 in three processes: DG separates, CA3 stores a pair after one showing; CA1 read-out is red (17 of 23 criteria).
- Infrastructure: bit-identical runs across 1–12 workers and across processes; 16 of 16 digests reproduced from a clean copy; 23 region acceptances published with their failures.

4 Hypothesis and preregistered criteria

H1. After paired showings of a letter and a tone with plasticity on, the letter alone reinstates in area 41 the evoked signature of its paired tone, and the tone alone reinstates in area 17 the signature of its paired letter.

Readout. Per stimulus, a standardised evoked signature per population ((response − baseline) / baseline SD). Recall succeeds when the paired template is top-1 by similarity with a margin above δ, where δ is three standard deviations of the pre-training margin distribution, fixed before training.

Pass. Forward 4 of 4 and backward 4 of 4 on 5 seeds, then the rate over 20–50 seeds. **Required failures:** an untrained letter (three of four pairs trained) recalls nothing; a shuffled hippocampal binding abolishes recall; plasticity off during the same paired showings abolishes recall.

5 Work packages

PACKAGE	WORK	DONE WHEN
WP1 Hippocampal recall	Close the red criteria of the loop: CA3 recall from half a cue; CA1 pair recall, precision, decoder, no erasure after a new pair	Loop acceptance 23 of 23; controls without STDP and without DG still break it
WP2 Cortex ↔ hippocampus	Layer 2/3 spikes of areas 17 and 41 into the entorhinal input of DG and CA3; the return path CA1 → deep entorhinal cortex / subiculum → cortex, each projection with a sourced number	DG sparsity in [0.005, 0.05] under cortical drive; certified sensory digests unchanged while the loop is passive
WP3 Association	Arbitrary table A ↔ 2 kHz, B ↔ 5 kHz, C ↔ 300 Hz, D ↔ 1 kHz; phase 0 baselines, phase 1 paired showings (vision 0–300 ms, sound 50–250 ms), phase 2 recall both ways	Section 4 criteria on 5 seeds, then 20–50 seeds from certified tissue
WP4 Release and replication	Tagged tree, fixtures, digests and seed rule released; an external researcher reproduces the result from the tag	Paper 2 preprint and a replication report

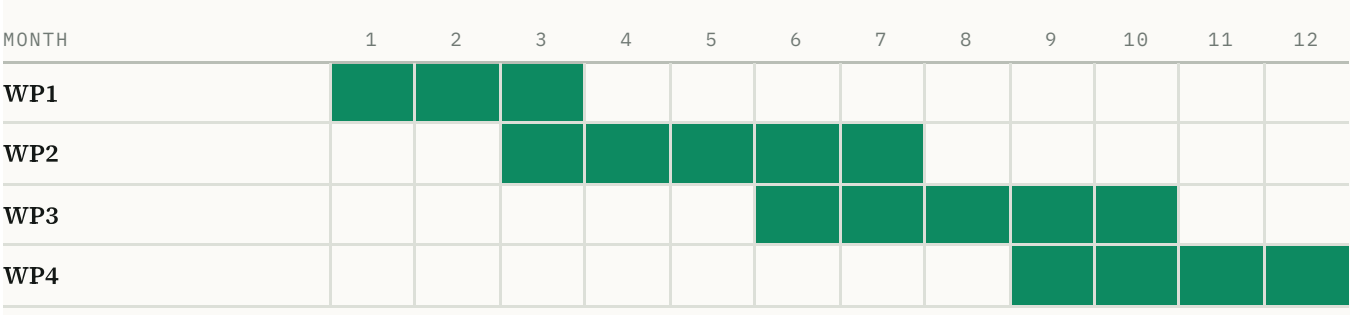

Three tiers keep compute honest: a fast developer test (four stimuli, one seed, under ten minutes) answers only whether the mechanism is alive; the acceptance runs the full path with controls on five fixed seeds; the claim runs 20–50 seeds that start from certified tissue files carrying the digests of the vision and hearing acceptances, so the sensory calibration is never re-litigated.

6 Budget

LINE	DETAIL	€
Principal investigator	12 months full time	66,000
Computational neuroscientist	Review and co-design, 6 person-months	30,000
Compute	One high-memory server (512 GB or more) and cloud seed sweeps	30,000
Publication	Open-access fees, preprint, two conferences	10,000
Independent replication	Contract with an external researcher	8,000
Administration	Legal, accounting, audit	6,000
Total		**150,000**

7 Risks and mitigation

CA1 stays red	Diagnosed causes are known (basket and axo-axonic clamp 1.6× stronger than storage; an unlearned Schaffer input above the gap). Fallback: a direct perforant path EC → CA3 for single-modality cues, declared before the run.
Association appears without learning	The plasticity-off control with the same paired showings separates learning from ongoing cross-modal spread; the arbitrary table excludes matching geometry.
Compute	The three-tier protocol keeps the expensive 20–50-seed claim to one run on certified tissue.
A negative result	Published with diagnostics, as Paper 1 published its failed hearing rule and geniculate relay. A clean negative is a deliverable.

8 Open science

Every threshold lives in a fixture before the run; the seed rule is committed to git before the seeds exist; every run carries a SHA-256 over every spike; the tagged tree, fixtures and artifacts are released with a SHA-256 manifest; and the budget pays an outside researcher to reproduce the result without the applicant.

9 Use of the results

The direct output is scientific: a preregistered answer to H1, with code, fixtures and digests released in WP4. The same assets are what the first users of BrainCore will pay for, in this expected order:

WHO	USE	WHEN
Neuroscience labs and universities	Graded region models in one runtime; reproducing a result by its digest; teaching labs	From the release of Paper 2
Pharma and medtech	Models of neural circuits in disease and under drugs (basal ganglia and dopamine, hippocampus, cochlea), with a source for every number. US law since December 2022 counts in silico tests and computer modeling as nonclinical tests of a drug [10]	2027–2029, after validation with a partner
Neuromorphic hardware makers	Brain-grade reference workloads with a known right answer	2027–2028
Industry and robotics	Learning on site, with no cloud and no retraining	From 2028, if H1 holds

Deployment. BrainCore runs on ordinary processors with no GPU, no cloud service and no outside calls, so it can be installed on an organisation's own servers and sensitive data never leaves them. **Licensing.** The core (region engine, acceptances, digest) stays open after WP4; certified regions, custom regions, installation and support are commercial. A negative answer to H1 keeps its value: it is published with its diagnosis, and the released tooling stays usable.

10 Applicant

Vitalii Cherepanov, engineer with fifteen years of production software and author of BrainCore. Author of the Paper 1 technical report and preprint on the sensory paths (September 2026, 53 checked references), and of the region core, the acceptance protocol and the anatomical atlas used here.

11 References

1. Potjans TC, Diesmann M (2014). Cerebral Cortex 24:785–806. doi:10.1093/cercor/bhs358
2. Schmidt M, Bakker R, Shen K, Bezgin G, Diesmann M, van Albada SJ (2018). PLoS Comput Biol 14:e1006359. doi:10.1371/journal.pcbi.1006359
3. Arkhipov A, Gouwens NW, Billeh YN, et al. (2018). PLoS Comput Biol 14:e1006535. doi:10.1371/journal.pcbi.1006535
4. Billeh YN, Cai B, Gratiy SL, et al. (2020). Neuron 106:388–403. doi:10.1016/j.neuron.2020.01.040
5. Eliasmith C, Stewart TC, Choo X, et al. (2012). Science 338:1202–1205. doi:10.1126/science.1225266
6. Pauli R, Weidel P, Kunkel S, Morrison A (2018). Front Neuroinform 12:46. doi:10.3389/fninf.2018.00046
7. Nosek BA, Ebersole CR, DeHaven AC, Mellor DT (2018). PNAS 115:2600–2606. doi:10.1073/pnas.1708274114
8. Guzman SJ, Schlögl A, Frotscher M, Jonas P (2016). Science 353:1117–1123 — hippocampal CA3 model used as the CA3 reference.
9. Cutsuridis V, Cobb S, Graham BP (2010). Hippocampus 20:423–446 — CA1 model and schedule (ModelDB 123815).
10. Consolidated Appropriations Act, 2023, Public Law 117-328, sec. 3209 (the FDA Modernization Act 2.0 provisions), 136 Stat. 5821–5822, approved 29 December 2022.