Research · Report NR-05 · 16–25 September 2026

The organism on published neuron models

Every region of the simulated brain was rebuilt with a published neuron model, one at a time, repeating the criteria it had already passed.

Gioele Chiarolanza Nadirion

Adopted, one criterion failed

Brian2 · seeds 0–11 · free Kaggle CPU · criteria and thresholds unchanged

5
regions rebuilt with the conductance-based LIF neuron of Brunel and Wang (2001)

0.93–0.95
withdrawal from danger with the neuron insula, against 0.18–0.23 in the control

12/12
organisms alive with the neuron basal ganglia

3/4
criteria for LC and workspace: the LC gain is not reproduced

SUMMARY

After the phase-1 tests (NR-04), the author decided to move the organism to spiking neurons, one module at a time, and from 17 September to use **only published neuron models**, not ones written by us. Each region was rebuilt with the conductance-based integrate-and-fire neuron of Brunel and Wang (2001) in the Brian2 simulator, and had to repeat the criteria already passed by the previous version, with the same thresholds. The amygdala reproduces the Rescorla–Wagner rule with a correlation of 0.974–0.982. The basal ganglia follow the reference choice with an error of 0.0009 and keep 12 organisms out of 12 alive. The insula, calibrated on behaviour, makes the organism move away from danger in 93–95% of cases. Locus coeruleus and workspace pass 3 criteria out of 4: the LC gain does not reproduce the irregular table of phase 1, as predicted in advance. The spiking workspace switches on but does not switch off: ignition and persistence are the same state of the network, an architecture problem still open.

1 Why published neurons

A neuron model written by us carries our choices with it, and a result obtained on that model says something about our choices before it says anything about the brain. A published model, used by other groups and with parameters fixed in the literature, reduces this risk and makes the results comparable. The chosen neuron is the conductance-based LIF of Brunel and Wang (2001), with AMPA, NMDA and GABA synapses: it is simple, well studied and cheap enough to simulate hundreds of neurons per region in many organisms at once.

The migration method is a single one: **the new region must pass the same criteria as the old one, with the same thresholds and the same seeds.** If it does not, the failure stays on record and the author decides.

2 The regions

region	role in the organism	outcome
D0 batch engine	several organisms integrated together	3 out of 3: same results, from about 62 to 2.4 ms per organism
D1 live locus coeruleus and workspace	vigilance and global access computed by neurons during life	4 out of 5, then 1 out of 1: adopted
D2 amygdala	learned fear and extinction with three-factor plasticity	2 out of 2 at the second run; correlation with Rescorla–Wagner 0.98
D3 basal ganglia	action selection with dopamine and evidence accumulation	2 out of 2: equivalent to the reference choice, 12 organisms out of 12 alive
D4 insula	state of the body read by neurons	4 sessions, all 4 out of 4
migration to Brian2 of D2, D3 and D1	same criteria with the published neuron	D2 2 out of 2, D3 2 out of 2, D1 3 out of 4

3 The insula: calibrating on behaviour

The most instructive case is the insula. The first runs passed the coding criteria, meaning the insula read the body well, but the organism **no longer learned to move away from danger**: 0.158 against a control of 0.200. The diagnosis, with hypotheses written in advance, showed that with small populations the spike noise made h too noisy for learning: doubling the neurons twice lowers the noise on h from 0.034 to 0.008. The calibration rule was rewritten before the next run: the number of neurons per channel is the smallest that, on the diagnostic seeds, passes the **behavioural criterion**, not only the coding ones.

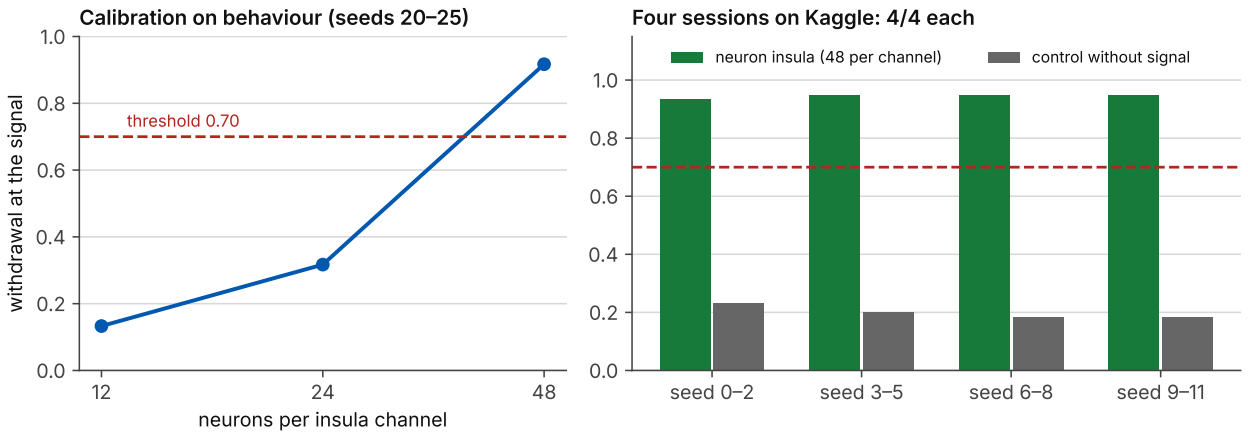


Figure 1. Left, calibration of the number of neurons per channel on the diagnostic seeds: with 12 neurons the organism moves away from danger in only 13% of trials, with 48 in 92%. Right, the four validation sessions on Kaggle, three seeds each, with 48 neurons per channel: withdrawal 0.93–0.95 against 0.18–0.23 in the control, 12 organisms out of 12 alive. Data from the `d4_spiking_insula_wm` record.

An unscored observation, to be read in later tests: with the neuron insula **the body becomes less predictable than the world** (ρ between 0.00 and 0.06, against 1–10 with the phase-1 insula). The premise of the theory also depends on the noise of the substrate.

4 The computing cost and the "bus"

Simulating one organism at a time in Brian2 cost 97.2 ms per organism per step. The solution was to put several independent organisms **in the same network**, like passengers on a bus: same behaviour (withdrawal 0.887 against 0.900, a difference within the 0.10 threshold), five times lower cost.

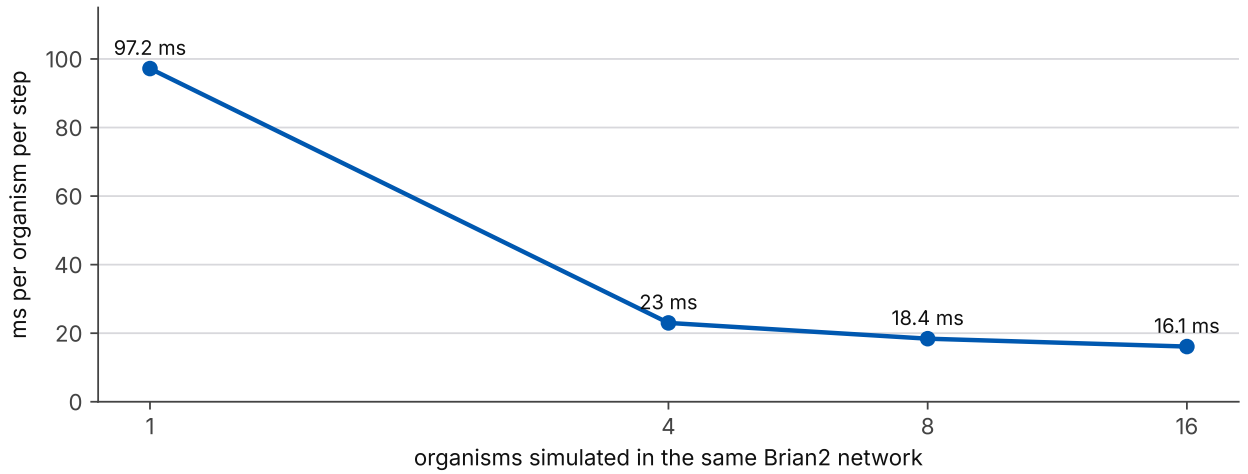


Figure 2. Cost per organism per step in Brian2 as the number of organisms simulated in the same network grows (diagnostic of 21 September 2026 on Kaggle). With 8 organisms the cost falls from 97.2 to 18.4 ms.

The first run of the basal ganglia on Brian2 was stopped by Kaggle at the 12-hour limit. Two cost diagnostics found the cause: too many calls to Brian2, each with a fixed cost of about 0.35 s. With one call per decision and the decisions collected in batches, the next run completed.

5 Locus coeruleus and workspace

The migration of D1 puts LC and workspace in a single continuous simulation with the organism attached. The cost diagnostic estimated 71 hours; with a computation-only optimisation, declared in advance, the session took 43 minutes.

criterion	outcome
L2 live LC	**failed**: the rate passes (error 1.93 Hz, correlation 0.989), the gain does not (correlation 0.238 against 0.9)
L3 live workspace	passed
L4 operator (pause, inspection, update)	passed
L5 functioning of the organism	passed

The phase-1 LC gain was an irregular table, with values 0 and 0.58 alternating at high activity (NR-03): a published neuron does not reproduce it, and **we had predicted this before the run**. The author adopted D1 on 25 September with the failure on record.

6 The open problem: a workspace that does not switch off

The global workspace must switch on when a signal is salient and switch off when it no longer is. The validated spiking version switches on, but once on it **stays on**. Three runs (adaptation, slow inhibition) found no valid configuration: 0 criteria out of 2. The diagnosis is that ignition and persistence are the same NMDA attractor of the network, a bistability built into the design. For the third time a module validated on a short protocol (400 ms) shows a defect at steady state; since then every criterion includes a long regime (the same lesson as NR-02).

7 What it shows and what it does not

It shows that the regions of the organism can move to a published neuron model while keeping the behaviours already validated, and that the change of substrate can be controlled with unchanged criteria.

It does not show anything about the theory: F1–F4 and T-OP must be repeated on the spiking organism (D5, under way, NR-06). The dopamine of the basal ganglia is computed, not released by

neurons. Spiking working memory (Compte and colleagues, 2000) did not produce a localised bump and has been postponed.

> **Limitations.** A single type of neuron for all regions; small populations (tens or hundreds of neurons per region); the workspace is bistable; the LC gain is not reproduced; the neuron insula makes the body less predictable than the world, and this changes the premises of the tests.

References

- Brunel, N. and Wang, X.-J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* 11, 63.
- Stimberg, M., Brette, R. and Goodman, D. F. M. (2019). Brian 2, an intuitive and efficient neural simulator. *eLife* 8, e47314.
- Compte, A., Brunel, N., Goldman-Rakic, P. S. and Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* 10, 910.
- Rescorla, R. A. and Wagner, A. R. (1972). A theory of Pavlovian conditioning. In *Classical Conditioning II*, Appleton-Century-Crofts.
- Dehaene, S. and Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron* 70, 200.
- Internal documents: whitepaper WP-14; `docs/FASE2_PIANO.md`; records `d0_batch_engine`, `d1_live_lc_workspace`, `d2_spiking_amygdala`, `d3_spiking_basal_ganglia`, `d4_spiking_insula_wm`.

How to cite

Chiarolanza, G. (2026). The organism on published neuron models. Nadirion, research report NR-05, 16–25 September 2026. https://nadirion.eu/en/research/published-neurons/

Nadirion technical report, not peer-reviewed. The criteria of every test were written before running it; negative results are reported like positive ones. The tests measure functional roles.