Preservation or Compensation?

A Protocol for Testing Whether an Artificial Substitute Takes Over a Neural Component's Causal Contribution

Luis Marcelo Pacheco Alvaro
Independent researcher, Lima, Peru
marcelo.com.28@gmail.com ORCID 0009-0001-2359-0803

Unpublished manuscript · version 0.1 · November 2026

Abstract. When a neural network keeps working after one of its components has been exchanged for an artificial unit, three hypotheses fit the observation. The unit may do what the component did, with the network depending on it (preservation). The rest of the network may have adapted around the loss (compensation). Or the component may not have mattered to the measured function in the first place (redundancy). Staged roadmaps for neural replacement state their acceptance criteria as bounded deviation from baseline, and a bound of that kind is met under all three. This paper proposes a protocol that separates them. It consists of thirteen conditions, among them acute and chronic silencing without a substitute, a sham, replay of a recorded output, a substitute with the wrong dynamics, graded infidelity, perturbation of the substitute, and reversal, run on a ladder of preparations from small identified circuits to identified ensembles in the mouse. I state what each hypothesis predicts for each condition and give a decision rule on their joint outcome. Several of the controls have published precedents; their combination is the proposal. I explain why cultured neurons that play video games are poor testbeds for the question. A positive result would support causal equivalence for the readouts and perturbations tested. The protocol can estimate fitting time and spacing between exchanges in the tested design. Its scope is electrical signalling. Questions of consciousness and personal identity remain open.

Keywords: neural prosthesis; closed-loop substitution; compensation; degeneracy; hybrid neural circuits; causal inference; gradual neural replacement

1. Introduction

Suppose a neuron in a working circuit is silenced, an artificial unit is made to drive its targets in its place, and the circuit goes on doing what it did. It is natural to conclude that the unit has taken over the neuron's job. The conclusion does not follow. At least three states of affairs produce the same observation. The unit may be doing what the neuron did, with the rest of the circuit depending on it as it depended on the neuron. I call this *preservation*. The rest of the circuit may have adapted so that it no longer needs what the neuron supplied. I call this *compensation*. Or the neuron may never have mattered to the function that was measured. I call this *redundancy*.

The difference matters in two places. Engineering frameworks for gradual neural replacement propose that artificial units assume the causal roles of neural components step by step, and they accept a step when the hybrid system stays within a bound of its baseline (Neukart 2026a; 2026b). A network that compensates meets such a bound, and so does one in which the replaced component was redundant. A replacement that passed every checkpoint in this way could be a slow removal of tissue, hidden by the adaptability of what remains. Philosophical arguments that a person could survive gradual replacement also need more than

continued function. On Weber's (2025) account each replacing part must be incorporated before most of the others are exchanged. In a companion paper I argue that a unit is incorporated to the extent that the system depends on it as it depended on the original, the unit responds to the system's current state, and the unit's state was produced by the system's own activity (Pacheco Alvaro, in preparation). The first two conditions are empirical. The companion paper sketches nine experimental conditions that bear on them and leaves the protocol to the present paper.

This paper is a hypothesis and a protocol. It reports no data, and the protocol has not been piloted. Its question belongs to experimental neuroscience and can be assessed without accepting any philosophical thesis: given a network that survives the substitution of a component, how can one tell which of the three hypotheses is true?

Two rules govern what follows. I state no numerical threshold, sample size or effect size that I cannot source; where the protocol needs one, it must be fixed in advance from pilot data. And I mark which supporting results are peer-reviewed, which are preprints, and which are demonstrations known only from press reports.

2. One observation, three hypotheses

2.1 The hypotheses

Let c be a component of a neural network: a neuron or, where single neurons cannot be addressed, a subpopulation or module. Let u be the artificial unit put in its place, and let F be a measured function of the network, such as the phase of a rhythm, the response of a target population, or performance on a task. The observation to be explained is that after c has been silenced and u installed, F remains within its baseline range.

Preservation. u supplies to the network what c supplied, and the network depends on u as it depended on c. If u were removed, the deficit that follows the removal of c would follow.

Compensation. The rest of the network has changed, through homeostatic regulation, plasticity or rerouting, so that F no longer requires what c supplied. u is a passenger. In a variant that I call *assisted compensation*, the activity of u helps to bring the change about, as a stimulator may promote recovery in surviving tissue, but once the change has occurred F does not depend on u.

Redundancy. F did not depend on c to begin with. Other paths already carried the function and nothing had to adapt. Redundancy has a weaker grade that is easily mistaken for preservation: F may have needed some activity at the position of c, for instance a tonic drive, without needing the particular relation between inputs and outputs that c computed. Any unit that supplies the drive will then serve. I call the first grade redundancy of the component and the second redundancy of its dynamics.

Three clarifications. The hypotheses are relative to a readout: a neuron may be redundant for a behaviour and still have a measurable effect on its direct targets. They are not exclusive: a population may be redundant up to some fraction and not beyond it, and a substitute may carry part of a contribution while the network makes up the rest. And 'contribution' means contribution to the network in its present state. That a network could reorganise to do without a component does not show that the component contributes nothing now.

The companion paper contrasts preservation with compensation and treats a component whose silencing has no effect as making the preparation uninformative. I promote redundancy

to a hypothesis because in large networks it is the first explanation to exclude, and because its weaker grade passes several tests that compensation fails.

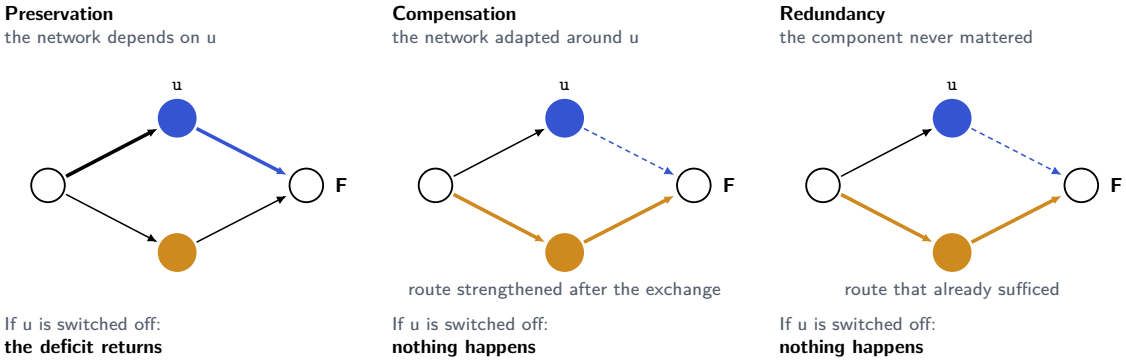


Figure 1. One observation, three explanations. In all three the measured function F stays in its normal range after a component is exchanged for the unit u. Thick lines mark the route that sustains F.

2.2 Why the alternatives are live

Network activity that is practically indistinguishable can arise from widely different parameters, as a database of some twenty million model versions of the pyloric network of the crustacean stomatogastric ganglion showed (Prinz, Bucher & Marder 2004), and stable function on variable parameters is a general feature of neurons and networks (Marder & Goaillard 2006).

Acute and chronic manipulations give different answers. Transient inactivation of a song-related area (Nif) in the zebra finch severely degrades song, whereas permanent lesions have almost no effect once singing resumes two days later; the recovery is consistent with homeostatic regulation downstream (Otchy et al. 2015). A function that looks intact days after an intervention may have been lost and regained.

The contribution of one neuron to a large network is small. In mouse visual cortex, photostimulation of a single layer 2/3 neuron reduced the population response to that neuron's preferred stimulus by about two per cent (Chettih & Harvey 2019). Exchanging one such neuron will change little that can be measured, whichever hypothesis is true.

Restored function need not run through the intended route, and realistic behaviour does not validate what produces it. In *C. elegans*, after a pair of interneurons had been removed, engineered electrical synapses between other neurons restored chemotaxis, but largely through lateral connections that had formed inadvertently (Rabinowitch et al. 2021). A *C. elegans* connectome coupled through a decoder trained by reinforcement learning to a simulated fly body yields realistic fly walking; the authors call the result biologically meaningless and conclude that the learning resides in the decoder (Brunton et al. 2026, a reviewed preprint).

Circuits that generate similar activity can respond differently to the same disturbance. Two distinct circuits in the crab stomatogastric ganglion that produce comparable rhythms were separated by a hormonal influence and not by a sensory one (Powell, Marder & Nusbaum 2021). Agreement at baseline, or under the perturbations used to fit a substitute, is weak evidence about conditions that were not tested.

2.3 Why a bound on deviation cannot decide

The most detailed roadmap for neural replacement that I know is Neukart's (2026b, §11). It has seven stages. They run from passive sensing and closed-loop modulation, through emulation of a single cell in shadow mode (the unit's output is computed and compared with the cell's but drives nothing) and partial replacement in slices, organoids and living animals, to large-scale replacement as a distant endpoint. At each stage an acceptance criterion requires a measure of deviation, typically a weighted sum of error terms, to stay below a tolerance specific to the stage. Any protocol should keep the roadmap's validation on held-out stimuli, its behavioural measures defined in advance and its insistence on reversibility. But it does not address the present question, for four reasons.

First, a bound on deviation concerns the output of the hybrid system. Under compensation the output returns to baseline because the network has adapted. Under redundancy it never left. Both meet the bound.

Second, shadow-mode and held-out validation test the unit: they show that its model predicts what the component does. They do not show that, after transfer, the network uses what the unit emits.

Third, reversibility figures in the roadmap as a safety requirement: after deactivation the system should return to its baseline range, with fallback to biological dominance where possible (§§9.5, 11.5), and distributed units should degrade gracefully under controlled dropout (§11.6). These are the right requirements while the biological component is still present. Once a component has been wholly replaced, a network that is undisturbed when the substitute is switched off is exactly what compensation and redundancy predict. For attribution the informative result is the opposite one.

Fourth, the framework lets the biological network adapt to artificial actuation during an assistive phase before handover (§5.2). That is sound engineering, and it means that the network's adaptation is built into the process, so its share in the outcome has to be measured. The framework guards against a unit that masks a failure of equivalence by excessive actuation (§5.2). It has no test for masking by the network.

None of this criticises the roadmap's purpose, which is to say when it is safe to proceed. A second kind of test is needed alongside it.

3. What substitution means today

Substitution rests on a technique more than thirty years old. The dynamic clamp introduces computer-generated conductances into a living neuron in real time (Sharp et al. 1993) and has been used to add simulated synapses and whole simulated neurons to living circuits (Prinz, Abbott & Marder 2004).

Two experiments are direct precedents. In the pyloric circuit of the spiny lobster, an electronic neuron implementing three differential equations was connected in the place of a deliberately damaged cell; the remaining neurons, which oscillated irregularly by themselves, recovered their rhythm through mutual interaction with the device (Szücs et al. 2000). And a thalamic circuit was reconstructed in vitro as a hybrid, with the inhibitory feedback loop closed by an artificial neuron (Le Masson et al. 2002). In both, an artificial element occupied a causal position in a living circuit and the circuit worked. That shows functional sufficiency at the level of the network's output. The substitute was low-dimensional and coupled through one or two somatic electrodes, and in the lobster experiment fidelity was

judged by the recovery of the rhythm, not by comparing the device's causal influence with that of the cell it replaced.

Other work substitutes at the level of a population or pathway. A neuromorphic board bridged two modules of a cultured network disconnected by a laser lesion, or stood in for a lost module with a network of a hundred model neurons (Buccelli et al. 2019). A chip implementing cerebellar plasticity rules reproduced a cerebellum-dependent form of learning in anaesthetised rats (Hogri et al. 2015). Stimulation patterns derived from a model of hippocampal ensemble activity restored performance on a memory task in rats whose hippocampal transmission was pharmacologically compromised (Berger et al. 2011), and a related approach improved short-term memory in patients (Hampson et al. 2018). In rats with cortical injury, stimulation triggered in one area by spikes detected in another produced better recovery than stimulation not contingent on activity (Guggenmos et al. 2013). None of these is a one-for-one replacement, and several probably work by inducing plasticity or facilitating encoding in surviving tissue. They are closer to assisted compensation than to preservation, which is no fault in a therapy and a confound in a test.

The realistic substitute today is therefore not a physical artificial neuron placed in tissue. Bench devices reproduce the electrical response of particular cell types (Abu-Hassan et al. 2019) or detect and release a transmitter (Wang et al. 2022), but none is intracellular and none has replaced an identified neuron in an intact brain circuit. The realistic substitute is a model run in real time that reads and writes through an electrode, an array or light. In outline, the main presynaptic partners of the target neuron are recorded; their activity is fed to a model fitted to that neuron; the biological neuron is silenced by hyperpolarisation, an inhibitory opsin or ablation; and whenever the model fires, the neuron's known postsynaptic targets are stimulated with the appropriate weight and latency. The building blocks exist: closed-loop all-optical methods read and write the activity of chosen neurons in the awake mouse and can hold a neuron's firing rate at a target value (Zhang et al. 2018; Adesnik & Abdeladim 2021). To my knowledge no published experiment has combined the silencing of an identified neuron with closed-loop substitution of its output onto its real targets in vivo.

This substitute has two limits.

The first is fidelity, which has at least three senses. (a) Predicting a neuron's spikes when a known current is injected at the soma is largely solved: simple threshold models with adaptation predict most of the predictable spikes (Jolivet et al. 2008; Gerstner & Naud 2009), they can be fitted automatically from about ten minutes of recording in vitro (Pozzorini et al. 2015), and hardware built from fitted ion-channel models reproduced the membrane voltage of hippocampal and respiratory neurons with 94–97 per cent agreement across sixty current protocols (Abu-Hassan et al. 2019). (b) Predicting the response to synaptic input distributed over the dendrites is much harder. A temporal convolutional network of five to eight layers was needed to capture the input–output relation of a biophysical model of a layer 5 pyramidal neuron at millisecond precision (Beniaguev, Segev & London 2021), and that comparison was with a model, not a recorded neuron. (c) Predicting behaviour over hours or days, with plasticity, neuromodulation and homeostatic change under way, has no established benchmark. A protocol cannot assume fidelity in senses (b) and (c). It can measure how much infidelity the network tolerates.

The second limit is the channel. Electrical or optical stimulation depolarises target cells. It does not reproduce the release of a particular transmitter at particular synapses, the co-release of neuropeptides, volume transmission, gap-junction coupling, retrograde and trophic signals, or the plasticity of the synapses that the silenced neuron made and received. These

omissions are not minor: in *C. elegans*, signal propagation measured for more than 23,000 pairs of neurons departs from what the anatomy predicts, in part because of extrasynaptic signalling by neuropeptides (Randi et al. 2023). The substitute available today replaces the electrical output of a component to its targets.

4. Where to run the test

4.1 Why game-playing cultures are poor testbeds

Cultures of neurons that play video games are the best-known hybrids of living tissue and silicon. I think they are the wrong place to test substitution, for four reasons.

Resolution. The interface addresses electrodes and regions, not identified neurons. In the peer-reviewed study in which cortical cultures of the order of a million cells improved their rally length in a simulated game of Pong, sensory input went through eight stimulation electrodes, motor output was read as spike counts in two predefined regions, and the authors note that single-cell structures could not be separated in real time (Kagan et al. 2022). A commercial system from the same company is reported to have 59 electrodes (New Atlas 2025, a press report); for the roughly 200,000 neurons of a recent demonstration that is, on my arithmetic, about one electrode per 3,400 cells. With such interfaces one cannot measure the inputs and outputs of a neuron before and after replacing it.

Power. A dissociated culture of a hundred thousand to a million neurons is highly redundant. Removing one neuron will not produce a detectable change in a noisy behavioural measure, so all three hypotheses predict the same null result.

Attribution. Part of the loop in these systems is conventional software. The 2026 demonstration in which human neurons on a chip controlled the game Doom paired the neurons with a standard learning algorithm, and how much of the learning took place in the neurons and how much in the software is disputed (Scientific American 2026, a press report).

Evidential status. The Doom demonstration has not appeared in a peer-reviewed publication, is described in press accounts as playing at a beginner's level, and comes with no published metrics or controls that I have found. The Pong study did not observe robust learning across sessions (Kagan et al. 2022).

4.2 A ladder of preparations

Table 1 sets out four rungs. The same conditions are run on each, and a higher rung is attempted only when the one below has given an interpretable result. The lower rungs offer the best attribution, and the order also serves the ethics of the work (§8).

Table 1. Ladder of preparations (proposed).

Rung	Preparation	Access: read, write, silence	What 'substitute' means today	Main readouts	Documented limits
1. Unit or microcircuit	Crustacean stomatogastric ganglion; or modular culture on a high-density CMOS array	Intracellular recording with dynamic clamp; array with spike attribution; ablation or hyperpolarisation	Real-time model coupled through an electrode (neuron) or through stimulation of target electrodes (module)	Phase, period and regularity of the rhythm; spike prediction; cross-correlation between modules; burst statistics	One or two somatic electrodes per coupled neuron; in cultures the component is a module, not a neuron

Rung	Preparation	Access: read, write, silence	What 'substitute' means today	Main readouts	Documented limits
2. Network with a task	Culture or organoid in closed loop with a game-like task, on a high-density array, decoder frozen	As rung 1	Substitution of a subpopulation or module with a measurable role	Task performance; learning curve; effective connectivity	Electrode-level access; learning may sit in the software
3. Small whole brain	*C. elegans*; larval zebrafish	Imaging of most of the brain; targeted optogenetics; laser ablation	Closed-loop model that reads presynaptic partners by imaging and photostimulates the targets of the silenced neuron	Behaviour (chemotaxis, escape, latencies); signal-propagation map; global dynamics	Incomplete coverage; slow imaging; peptide signalling not reproduced
4. Mammal	Mouse: visual cortex (perception); hippocampus or amygdala (fear engram)	Two-photon imaging with holographic stimulation; inhibitory opsins	The same, for an ensemble of identified neurons, such as a fraction of an engram	Freezing to a natural cue; context specificity; new learning; influence map	Single-neuron influence of about two per cent; thousands of unrecorded synaptic inputs per neuron

Rung 1. The stomatogastric ganglion is probably the best place for a first experiment. It has some tens of identified neurons, each with a measurable role in a quantifiable rhythm; a neuron has already been substituted there (Szücs et al. 2000); and degeneracy and compensation have been characterised in it (Prinz, Bucher & Marder 2004; Marder & Goaillard 2006). The alternative is a modular culture on a high-density array. An array of 26,400 electrodes at a pitch of 17.5 µm allows putative single neurons to be localised and axonal propagation to be followed (Müller et al. 2015). Here the component is a module, as in Buccelli et al. (2019), whose authors state that they aimed to imitate global activity and that 60 electrodes could not resolve the network's topology.

Rung 2. A culture or organoid performing a closed-loop task adds a behavioural readout at the cost of the problems of §4.1. A game can serve as a secondary readout provided that a high-density array is used, that what is replaced is a module or subpopulation with a measurable role, and that the decoder is frozen.

Rung 3. Small animals whose brains can be imaged almost whole are where a complete experiment, with all its controls, seems most feasible in the medium term. *C. elegans* has 302 identified neurons, optogenetic access and measurable behavioural effects of single neurons. In larval zebrafish a whole-brain light-sheet platform combines online functional maps with optical ablation of chosen cells: ablating the Mauthner cell lengthened behavioural latency and ablating a control cell did not (Vladimirov et al. 2018). The limits are coverage and speed. According to a 2025 report on the state of brain emulation, a preprint whose summary figures I use, about half of the neurons of *C. elegans* and about 80 per cent of the larval zebrafish brain can be recorded at cellular resolution, and calcium imaging typically runs at 1–30 Hz (Zanichelli et al. 2025), too slow for a substitute that must fire with millisecond precision. The same report puts at 40–50 per cent the share of synaptic connections that

differ between genetically identical worms, which implies that a substitute has to be fitted to the individual neuron in place and cannot be built from a reference wiring diagram.

Rung 4. In the mouse, two-photon holographic optogenetics reads and writes the activity of chosen neurons (Adesnik & Abdeladim 2021), and a behavioural readout tied to identifiable ensembles exists: dentate gyrus neurons active during fear conditioning can be tagged, and reactivating them in another context produces freezing (Liu et al. 2012). Because the influence of a single cortical neuron is of the order of two per cent, the component must be an ensemble or a neuron with a disproportionate role. What is feasible today is the silence-and-drive-the-targets form of substitution for a fraction of such an ensemble. Physical replacement is not.

5. The protocol

5.1 What is fixed in advance

Four things are fixed, and ideally registered, before any condition is run.

The component and the readouts. The component is chosen for having a measurable role. Readouts are defined at three levels: the unit (how well the model predicts the component's spikes, and the responses of its direct targets); the network (phase, period and regularity of a rhythm; cross-correlation and burst statistics of coupled modules; effective connectivity, information flow to the targets and perturbational complexity: Casali et al. 2013); and behaviour, where there is any. In simulations of modular neuronal networks with known wiring, the accuracy of structural reconstruction from spontaneous activity depended on the dynamical regime and network architecture (Barjuan et al. 2026).

The held-out set. Circuits with similar activity can respond differently to a perturbation (Powell, Marder & Nusbaum 2021). Reserve a set of inputs and perturbations for the planned comparisons between conditions. Fit the substitute on separate observations and fix the readouts and perturbations before examining the reserved results. Changes prompted by those results are exploratory and require confirmation on a fresh, independent held-out set.

The margins. The decision rule uses equivalence, superiority and trend. For each readout, specify the outcome, the direction of a deficit, the smallest relevant difference, the confidence level and the planned comparisons. Equivalence requires a confidence interval fully within the pre-set equivalence bounds. When conditions are interleaved within a preparation, account for the paired observations and estimate uncertainty across independent preparations. Use pilot data to plan the number of preparations and inform the margins, the window that counts as immediate and the tolerance for spike coincidence. One published benchmark used ± 2 ms (Jolivet et al. 2008). The values depend on the preparation and readout.

The software. Any decoder or learning algorithm in the loop is frozen before the baseline is taken.

Where the substitute is a program driving an interface, conditions 3 and 5 to 8 differ only in the program, and reversible silencing allows them to be interleaved within one preparation. Conditions 2, 9, 10 and 11 need hours or days.

5.2 The conditions

Table 2. The thirteen conditions.

#	Condition	What is done	What it isolates	Published precedent
0	Baseline	Intact circuit. Characterise the component (inputs, outputs, influence map) and fit the model in place, in shadow mode	The reference, and the fidelity the model can reach	Influence mapping (Chettih & Harvey 2019); propagation atlas (Randi et al. 2023)
1	Acute silencing, no substitute	Silence or ablate the component and measure at once	Whether the component matters to the readout	Otchy et al. (2015); Vladimirov et al. (2018)
2	Chronic silencing, no substitute	As 1, followed for hours or days	The time course and signature of spontaneous compensation	Otchy et al. (2015)
3	Faithful closed-loop substitute	Silence and substitute at the same moment, with the fitted model fed by real inputs	The experimental condition	Szücs et al. (2000), neuron; Buccelli et al. (2019), module
4	Sham	Full instrumentation (opsin, light, electrodes); component intact; substitute disconnected	Effects of the intervention itself	Analogous control: ablation of a control cell (Vladimirov et al. 2018)
5	Replay, open loop	Play back a spike train recorded earlier from the component, not contingent on present inputs	Whether tonic drive or firing statistics suffice	Contingent against non-contingent stimulation (Guggenmos et al. 2013)
6	Mismatched dynamics	Model of another neuron; shuffled parameters; Poisson train of the same mean rate	Whether the network is indifferent to what the substitute computes	None found; motivated by Brunton et al. (2026)
7	Graded infidelity	Add increasing latency, jitter or gain error to the faithful substitute	The dose–response relation between fidelity and function	None found in experiment; simulated failure analyses in Neukart (2026b, §10.8)
8	Perturbation of the substitute	Inject or delete spikes in the substitute; change its gain	Whether the substitute lies on the causal path	Maps exist for original neurons (Chettih & Harvey 2019); none for a substitute
9	Reversal	Switch the substitute off after hours or days, the component remaining silenced	Dependence of the recovered function on the substitute	Buccelli et al. (2019)
10	Cumulative dose	Substitute growing fractions of a population; compare with silencing the same fractions, acutely and chronically	Population redundancy; an approximation to gradual replacement	None found; variation of replacement density is proposed in Neukart (2026b, §11.4)
11	Prior memory and new learning	Form a memory before substitution; test recall and new learning afterwards	What depends on weights and plasticity	Engram readout (Liu et al. 2012); substitute with plasticity rules (Hogri et al. 2015)
12	Temporal rescaling (full emulation only)	Accelerate a simulated agent and its simulated environment by the same factor	Invariance under uniform scaling; a consistency check (§9)	Accelerated neuromorphic hardware (Pehle et al. 2022)

Conditions 1 to 9 are those listed in the companion paper; 0, 10, 11 and 12 are added here. I have kept all thirteen. Replay and mismatch might seem candidates for merging, but they remove different things: replay keeps the component's own firing statistics and removes contingency on present input, while a mismatched model keeps contingency and removes the

component's dynamics. Condition 12 does not discriminate between the hypotheses and is set apart in §9.

Four conditions need comment. *Condition 6* has a variant with a special role: a model fitted to the homologous neuron of another animal, the experimental counterpart of what the companion paper calls a preloaded part. *Condition 9* requires that the component stay silenced when the substitute is switched off; otherwise the test shows only that the original still works. *Condition 10* yields three curves of function against fraction: substitution, acute silencing and chronic silencing. It turns redundancy from an obstacle into a variable, since the fraction at which acute silencing first produces a deficit is the fraction at which the test acquires power. *Condition 11* must test recall with natural cues, because engram cells can be reactivated artificially even when natural cues fail to retrieve the memory (Josselyn & Tonegawa 2020).

The combination is the proposal of this paper. To my knowledge no study has applied these conditions jointly to a neural substitution.

5.3 Predictions

Table 3. What each hypothesis predicts. 'As 3' means equivalent to condition 3 within the pre-set margin.

#	**Preservation**	**Compensation**	**Redundancy of the component**	**Redundancy of its dynamics**
1	Immediate deficit	Immediate deficit	No deficit	Immediate deficit
2	Slow or incomplete recovery, with homeostatic signatures in neighbouring cells: the reference curve	The same; this curve defines the signature of compensation	Nothing to recover	As preservation
3	Baseline from the first measurement; no transient; neighbours within the sham range	Transient deficit, then recovery along the curve of condition 2, with its signatures	Baseline	Baseline from the first measurement
4	Baseline	Baseline	Baseline	Baseline
5	Worse than 3, above all on held-out inputs	As 3	As 3	As 3
6	Worse than 3	As 3	As 3	As 3
7	Decline with infidelity, possibly after a plateau	Flat	Flat	Flat over a wide range
8	Influence map matches the original in sign, magnitude, latency and specificity	Absent, weak or different	No effect on the readout, as for the original	Effects follow the amount of drive; specificity does not matter
9	The deficit of condition 1 returns	Nothing happens	Nothing happens	The deficit of condition 1 returns
10	Substitution curve flat while the acute-silencing curve falls	Substitution curve follows the chronic-silencing curve, with a transient at each step	All three curves flat over the redundant range	Substitution curve flat, with mismatched units as well

#	Preservation	Compensation	Redundancy of the component	Redundancy of its dynamics
11	Recall intact and specific; new learning normal only if the substitute implements plasticity	Recall degraded or re-learned; new learning carried by other neurons	Recall and learning unaffected by silencing alone	Recall intact with any unit that supplies drive
12	Invariant; does not discriminate	—	—	—

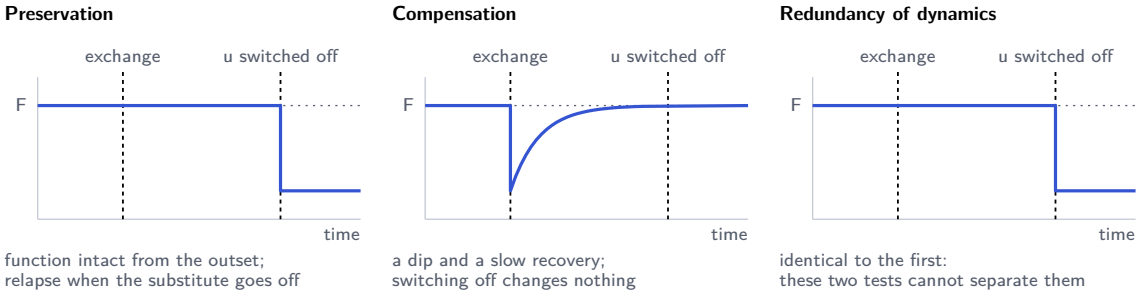


Figure 2. What two of the tests predict, in outline. These are not data. The dotted horizontal line is the baseline; the substitute comes in at the exchange and is switched off later. Preservation and redundancy of dynamics give the same curve, which is why conditions 5 to 7 are needed.

5.4 Decision rule

The rule is applied separately at each level of readout.

Gate. First check equivalence between condition 4 and baseline. A departure beyond the pre-set margin calls for revising the preparation or instrumentation. If the comparison cannot establish either equivalence or a meaningful departure, the gate is inconclusive. After confirming that the component was silenced in condition 1, proceed with the substitution test if the confidence interval for the acute deficit lies wholly above the pre-set margin. Equivalence to baseline supports redundancy of the component for that readout. In that case one can move to a different component, a finer readout or a larger fraction of the population (condition 10). Other outcomes leave the preservation test inconclusive for that readout. One example is an interval that spans both a negligible effect and a meaningful deficit.

Preservation is supported only if condition 3 is equivalent to baseline within the pre-set immediate window and outperforms conditions 5 and 6 on the held-out set. Condition 7 must show a decline in function across the pre-specified range of infidelity. The influence map of condition 8 must match the map of the original component from condition 0 within the pre-set tolerance, and the deficit of condition 1 must return in condition 9. Homeostatic signatures in neighbouring cells, judged against the sham, provide additional evidence (§5.5).

Compensation is supported if condition 3 shows a transient followed by recovery on the time course of condition 2, conditions 5 and 6 are equivalent to 3, and the effect of infidelity in condition 7 stays within the pre-set tolerance. Two further observations support this interpretation: perturbing the substitute has little downstream effect, and the readout is equivalent with the substitute on and off.

Redundancy of the component's dynamics is indicated if conditions 3, 5 and 6 all hold function at baseline from the outset and the curve of condition 7 is flat within the pre-set tolerance, while reversal brings the deficit back. The network depends on the unit, and the tested alternatives serve the same measured function. A different readout or set of inputs may distinguish the alternatives. If chosen in response to these results, it requires a fresh confirmatory test under §5.1.

The companion paper, which contrasts two hypotheses, notes this last pattern without naming it. The finer division does not change the verdict on preservation. It changes what to do next.

Recovery within the immediate window weighs against compensation on the slower time scale observed in condition 2. A deficit on reversal establishes dependence on the substitute at the time of the test, an outcome also possible if the network has reorganised around the activity of that substitute. In the bridging experiment of Buccelli et al. (2019), correlation between the modules recovered in part with the bridge and fell again when it was switched off, showing dependence on the bridge. Conditions 5 to 7 test whether equivalent function can be obtained with different dynamics. Condition 8 compares the causal influence of the substitute directly with that of the original. I have found no precedent for that comparison.

5.5 Mixed outcomes and the tolerance of networks

Clean patterns will be the exception, for three reasons.

A substitute may carry part of the contribution while the network makes up the remainder. With the substitute of §3 this is to be expected even if the unit lies on the causal path: targets depolarised by light or current, not by transmitter at their synapses, are likely to adjust. That is why homeostatic signatures are treated as graded evidence. The ratio of the deficit on reversal to the acute deficit in condition 1 compares the sizes of the two effects on the same readout. A ratio near one indicates similar effects in magnitude. A well-estimated ratio near zero indicates little dependence on the substitute for that readout. The remaining controls address whether the underlying contribution was preserved. Pilot work would establish the precision of the ratio.

Adaptation is not always adaptation around the unit. The engineering framework expects the network to adapt to the unit, and the companion paper allows the unit's state to be altered after handover by the activity-dependent processes that alter other parts. Change in neighbouring cells is therefore not in itself evidence of compensation. What separates the two is direction, which conditions 8 and 9 measure: after adaptation to the unit the network depends on it more, and after adaptation around it less.

Networks tolerate imperfect parts. The curve of condition 7 will often have a plateau before it falls, and its width is a finding in its own right, since it measures how much fidelity continued function demands. It also limits the test: within the plateau a faithful substitute and a mediocre one cannot be told apart by function, and tolerance of an imperfect unit is evidence that the network is robust, not that the unit is faithful.

6. What a result would and would not show

6.1 Causal equivalence at the measured grain

Suppose the preservation pattern is obtained on some rung. What has been shown is that, for the readouts taken and the held-out conditions tested, an artificial unit occupies the

causal position of a component. The network needs it, needs it to respond to current input and to have the right dynamics, and is influenced by it as it was by the original.

In the terms of the companion paper, conditions 8, 9 and 10 bear on dependence (D) and condition 5 on responsiveness (R). The third condition of that account, that the unit's state be produced by the system's own activity (E), is secured by design, by fitting the model in place and in shadow mode at condition 0, and is not verified afterwards. It has an observable correlate, generalisation to conditions the unit was not fitted to, which the held-out set tests for every substitute and the homologous-neuron variant of condition 6 tests by contrast: if the reported variability between genetically identical worms is correct, a preloaded unit of that kind should do worse than one fitted in place.

6.2 Fitting time and spacing between exchanges

Weber (2025) represents incorporation by a fixed delay. The companion paper argues that the time needed depends on the unit, the tissue and the use to which the system is put. The proposed tests can estimate two timing constraints for electrical substitution.

Fitting time. Among the durations tested, find the shortest period of in-place fitting in shadow mode for which the substitution then meets the joint criterion in §5.4. Compare models fitted for different periods while recording the range of network states sampled during fitting.

Spacing. Extend condition 10 by varying the interval between two successive exchanges in the same population. Estimate the shortest tested interval at which the preservation pattern holds for both substitutions.

Each estimate is relative to the preparation, substitute and readout. The protocol exchanges the component and substitute at a single moment. Measuring the duration of a gradual handover would require observing the unit's causal contribution during that transfer.

6.3 What a result would not show

It would not show that the whole causal contribution of the component has been preserved. The functional premise on which the companion paper's thesis about survival depends speaks of a unit that reproduces a component's electrical and chemical interactions, its state, and the rules by which its state changes. A substitute that replaces electrical output tests a weak version of that antecedent. Condition 11 marks the boundary: where recall or new learning depends on plasticity at the replaced component's own synapses, a present-day substitute is predicted to fail unless it implements that plasticity.

It would not show anything about consciousness. Causal equivalence at the grain of spikes, influence and behaviour is compatible with every position on whether the hybrid system has experience.

It would not show anything about personal identity. Whether a system that persists through the replacement of its parts is the same person is a question for philosophical argument, which needs premises that no experiment on a ganglion or a worm supplies. The experiments can assess evidence of causal equivalence in the hybrid system and constrain the timing of the tested substitutions.

A precise result outside a pre-specified bound shows where the tested substitute fell short on that readout. A poor fit, limits of the interface and missing causal detail are possible explanations. The need for a finer grain is a hypothesis for subsequent comparisons.

An estimate too imprecise to distinguish a meaningful departure from equivalence remains inconclusive.

7. Limits of current technology

In the thought experiment that motivates this question, small devices enter neurons, learn what they do and take over. Nothing in the protocol depends on it, and it should be said plainly how far it is from what exists.

No design exists for an intracellular device that learns and takes over a neuron. I have found no peer-reviewed engineering concept of one. The nearest proposal describes three classes of nanorobot that would monitor neurons, glia and synapses and transmit what they record; it is a hypothesis-and-theory article, explicitly speculative, and it concerns monitoring, not functional substitution (Martins et al. 2019). The best-known roadmap for brain emulation takes the route of scanning and emulating, not of replacement in place (Sandberg & Bostrom 2008). The framework I have relied on for the handover treats nanoscale units as a candidate substrate and states that it does not claim that such devices can replace neural tissue (Neukart 2026b, §§1, 6). In this literature the nanoscale device is an idealisation. I use it as one, and the experiments proposed here are independent of it.

Real devices are far from it. The smallest wireless devices reported to reach the brain through the bloodstream are about 10 µm across, the size of a cell body. They stimulate; they do not enter neurons, record or learn (MIT News 2025, a press release on work I have not consulted directly). Wireless operation inside a living cell has been shown with a sub-millimetre antenna in frog oocytes, which are giant cells, and without any neural computation (Joy et al. 2022). The one artificial neuron I know of that detects and releases a transmitter, dopamine, is a bulky bench assembly whose authors acknowledge gaps in response time, power, size and stability (Wang et al. 2022).

Self-replication is unnecessary and should not be assumed. The hypothesis requires only that each unit reproduce a causal role. Units manufactured outside the body and deployed in successive waves serve the argument equally well. Self-replicating nanomachines do not exist; proponents of molecular manufacturing have themselves argued that replication is unnecessary to it (Phoenix & Drexler 2004); and assuming it would add a hazard, uncontrolled replication in tissue, that would need a safety case of its own.

The heat budget is unfavourable, on illustrative arithmetic. In a simplified model, brain tissue tolerates local temperature rises of about 2 °C for hours, which corresponds to roughly 40 mW dissipated in a 500 mg mouse brain; the authors present this as a first approximation (Marblestone et al. 2013). Dividing by about 70 million neurons, the count for the mouse given in the report cited in §4.2, leaves about 0.6 nW for each neuron if every neuron carried an active device. The most efficient solid-state neuron among my sources is reported, in press coverage of Abu-Hassan et al. (2019), to draw about 140 nW. The gap is more than two orders of magnitude before sensing, communication and chemical release are counted. This arithmetic is illustrative only. It is mine, it combines sources of unequal standing, and it ignores the fact that a device which replaces a neuron also removes that neuron's metabolic heat. It is not a bound.

What else cannot be done today. Present techniques cannot record all the synaptic inputs of a neuron in vivo, which on my arithmetic from the counts for a cubic millimetre of mouse visual cortex number roughly 4,000 to 7,000 per neuron (MICrONS Consortium 2025); deliver a substitute's output synapse by synapse; release transmitters and peptides locally, with millisecond kinetics, from a device the size of a cell; reproduce the plasticity of

the replaced neuron; guarantee the long-term stability of artificial neurons coupled to tissue; validate an emulation, for which no accepted standard of fidelity exists; or substitute the contributions of glia and vasculature.

8. Ethics of the proposed experiments

A paper that only proposes experiments needs no approval, but it should say under what standards they ought to be done.

Cultures and organoids. A 2021 report of the US National Academies concluded that neural organoids lack cell types, regions and circuits thought necessary for consciousness and are extremely unlikely to be conscious or to feel pain in the foreseeable future, and that oversight should be reassessed as the field develops (National Academies of Sciences, Engineering, and Medicine 2021). The lower rungs of the ladder use invertebrate ganglia and dissociated or modular cultures, which are simpler still. Two obligations remain.

The first concerns vocabulary. The Pong study described its cultures with the word *sentience*, defined minimally. The usage was contested in the same journal (Balci et al. 2023), and the dispute shows how a term can come to dominate the reception of a result. Nothing in the present protocol measures or presupposes experience in a culture, and results should be reported as measurable causal and functional equivalence. The same applies to perturbational complexity, which was introduced as an index of consciousness in patients (Casali et al. 2013) and would be used here only to describe network dynamics.

The second is proportion. If preparations are ever built that are serious candidates for sentience, Birch (2024) proposes either a moratorium or ethical review with harm–benefit analysis as for animals, with precautions proportionate to the risk. The second rung, on which an organoid performs a task in closed loop, is where this would first need thought.

Animals. On the third and fourth rungs the three Rs apply, and the design builds them in. Replacement: the ladder is ordered so that questions are answered in cultures, invertebrate circuits and simulation before a vertebrate is used. Reduction: conditions that differ only in the program driving the substitute are interleaved within one preparation, and the number of animals follows from margins fixed on pilot data. Refinement: silencing should be reversible where the condition allows, and humane endpoints and stopping criteria should be set before the experiment, as the engineering roadmap also requires (Neukart 2026b, §11.5). Three conditions call for particular scrutiny because they are designed to produce or prolong a deficit: chronic silencing (2), reversal (9) and the silencing arm of cumulative dose (10). The fear-memory readout of the fourth rung uses aversive conditioning and targets a structure critical for memory, a kind of target that the same roadmap advises against at early stages. It belongs last, and whether a non-aversive memory would serve should be considered at the design stage.

In the European Union, Directive 2010/63/EU covers live non-human vertebrates, including certain larval forms, and live cephalopods (European Union 2010). On my reading of that scope, crustaceans, nematodes and insects fall outside it, and whether a larval zebrafish falls inside depends on its stage. Legal scope is not moral scope, and the proportionality that Birch recommends is a reason not to treat the invertebrate rungs as ethically free.

The status of a hybrid animal. The experiments create animals part of whose nervous system has been functionally replaced. If the thesis of the companion paper is right then, given its premise about function, such an animal is the same animal with the capacities it had, and its moral status is what it was. If the thesis is wrong, it is still an animal with the rest of its nervous system intact. On neither reading does substitution lower the standard

of care. At the end of an experiment, an animal in which the preservation pattern has been obtained depends on a device, and withdrawing the device harms it, which is the animal counterpart of the abandonment of implanted devices in patients (Okun et al. 2024). The protocol should say in advance what happens to such animals when the measurements are complete.

9. A separate hypothesis: acceleration after full replacement

It is sometimes suggested that a nervous system wholly replaced by artificial units could then be run faster. This is a different hypothesis from the one tested above, and it inherits none of the support that a positive result there would give.

Acceleration is possible only when no biological part remains. While any living neuron is coupled to the artificial units, those units must keep biological time, and the chemical processes of tissue do not scale with a parameter. Even in a wholly artificial system, every time constant, including the slow ones of modulation and homeostasis, must be scaled by the same factor. That presupposes units whose time constants are adjustable, which is a further premise about their design and not a consequence of replacement in place. And the body and the world do not accelerate: a faster system coupled to an ordinary body receives new input at the old rate.

Analogue neuromorphic hardware emulates spiking neurons, plasticity included, about a thousand times faster than the biological dynamics it models, in a demonstration with 512 simplified neuron circuits on a chip (Pehle et al. 2022). What is testable today is narrower. If an agent and everything it interacts with are scaled by the same factor, the acceleration is by construction invisible from within, and the prediction is that behaviour and dynamics, expressed in rescaled time, are unchanged. This is condition 12. It can be run only in simulation, on an emulated nervous system with an emulated body and environment. A preprint reports closed-loop experiments on accelerated hardware in which the timescales of an artificial agent and of its environment are both speeded up (Schreiber et al. 2020, a preprint). I have found no test of invariance on an emulation of a particular biological nervous system. Condition 12 does not discriminate between the three hypotheses; it is a consistency check for a hypothesis that begins where the protocol ends.

Whether a uniformly accelerated system would have more experience per unit of external time is a further question, contested in the philosophical literature (Mogensen 2023), on which I take no position.

10. Conclusion

Continued function after substitution is the observation on which every proposal for neural replacement relies, and it is the one observation that preservation, compensation and redundancy all predict. I have proposed a way to separate them. Silence the component with and without a substitute, acutely and chronically. Compare a faithful closed-loop substitute with a sham, a replay, a mismatched unit and graded degradations of itself. Perturb the substitute and compare its influence with the original's. Switch it off. Repeat over growing fractions of a population, on a ladder of preparations in which single components can be addressed. Preservation is supported only by the conjunction: baseline function at once, superiority over replay and mismatch, a dose–response relation with fidelity, a matching influence map, and relapse on reversal. Most of the individual controls have been run somewhere. Their combination, to my knowledge, has not.

What the protocol can deliver is limited. It can show that an artificial unit has become load-bearing in a living network at the grain of electrical signalling, and estimate the fitting time and spacing between exchanges required under the tested conditions. It cannot show that a unit reproduces a neuron's chemistry or plasticity, because present substitutes do not attempt to, and it says nothing about experience or about who survives. Arguments about gradual replacement have so far taken incorporation for granted. An experiment of this kind would make it something that can fail.

Statements

AI assistance. The author wrote the original manuscript. Claude (Anthropic) assisted with the literature search and preparation of the figures. An AI assistant also supported an earlier exploratory discussion. Codex (OpenAI) assisted with revisions to the text, bibliographic checks and LaTeX typesetting. The author is responsible for the manuscript.

Competing interests. The author declares none.

References

Abu-Hassan, K., Taylor, J. D., Morris, P. G., Donati, E., Bortolotto, Z. A., Indiveri, G., Paton, J. F. R., & Nogaret, A. (2019). Optimal solid state neurons. *Nature Communications*, 10(1), Article 5309. https://doi.org/10.1038/s41467-019-13177-3

Adesnik, H., & Abdeladim, L. (2021). Probing neural codes with two-photon holographic optogenetics. *Nature Neuroscience*, 24(10), 1356–1366. https://doi.org/10.1038/s41593-021-00902-9

Balci, F., Ben Hamed, S., Boraud, T., Bouret, S., Brochier, T., Brun, C., Cohen, J. Y., Coutureau, E., Deffains, M., Doyère, V., Gregoriou, G. G., Heimel, J. A., Kilavik, B. E., Lee, D., Leuthardt, E. C., Mainen, Z. F., Mathis, M., Monosov, I. E., Naudé, J., … Zimmermann, J. (2023). A response to claims of emergent intelligence and sentience in a dish. *Neuron*, 111(5), 604–605. https://doi.org/10.1016/j.neuron.2023.02.009

Barjuan, L., Yamamoto, H., Serrano, M. Á., & Soriano, J. (2026). Regime-dependent reconstruction of structural connectivity: an in silico study of spatially constrained modular neuronal networks. *arXiv* [Preprint, arXiv:2611.25830v1]. https://arxiv.org/abs/2611.25830v1

Beniaguev, D., Segev, I., & London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron*, 109(17), 2727–2739.e3. https://doi.org/10.1016/j.neuron.2021.07.002

Berger, T. W., Hampson, R. E., Song, D., Goonawardena, A., Marmarelis, V. Z., & Deadwyler, S. A. (2011). A cortical neural prosthesis for restoring and enhancing memory. *Journal of Neural Engineering*, 8(4), Article 046017. https://doi.org/10.1088/1741-2560/8/4/046017

Birch, J. (2024). When is a brain organoid a sentience candidate? *Molecular Psychology: Brain, Behavior, and Society*, 2, Article 22 (version 2). https://doi.org/10.12688/molpsychol.17524.2

Brunton, B. W., Abe, E. T. T., Hu, L. J., & Tuthill, J. C. (2026). The digital sphinx: Can a worm brain control a fly body? *eLife* [Reviewed preprint, version 1]. https://doi.org/10.7554/eLife.111516.1

Buccelli, S., Bornat, Y., Colombi, I., Ambroise, M., Martines, L., Pasquale, V., Bisio, M., Tessadori, J., Nowak, P., Grassia, F., Averna, A., Tedesco, M., Bonifazi, P., Difato, F., Massobrio, P., Levi, T., & Chiappalone, M. (2019). A neuromorphic prosthesis to restore communication in neuronal networks. *iScience*, 19, 402–414. https://doi.org/10.1016/j.isci.2019.07.046

Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., Casarotto, S., Bruno, M.-A., Laureys, S., Tononi, G., & Massimini, M. (2013). A theoretically based index of

consciousness independent of sensory processing and behavior. *Science Translational Medicine*, 5(198), Article 198ra105. https://doi.org/10.1126/scitranslmed.3006294

Chettih, S. N., & Harvey, C. D. (2019). Single-neuron perturbations reveal feature-specific competition in V1. *Nature*, 567(7748), 334–340. https://doi.org/10.1038/s41586-019-0997-6

European Union. (2010). Directive 2010/63/EU.

Gerstner, W., & Naud, R. (2009). How good are neuron models? *Science*, 326(5951), 379–380. https://doi.org/10.1126/science.1181936

Guggenmos, D. J., Azin, M., Barbay, S., Mahnken, J. D., Dunham, C., Mohseni, P., & Nudo, R. J. (2013). Restoration of function after brain damage using a neural prosthesis. *Proceedings of the National Academy of Sciences*, 110(52), 21177–21182. https://doi.org/10.1073/pnas.1316885110

Hampson, R. E., Song, D., Robinson, B. S., Fetterhoff, D., Dakos, A. S., Roeder, B. M., She, X., Wicks, R. T., Witcher, M. R., Couture, D. E., Laxton, A. W., Munger-Clary, H., Popli, G., Sollman, M. J., Whitlow, C. T., Marmarelis, V. Z., Berger, T. W., & Deadwyler, S. A. (2018). Developing a hippocampal neural prosthetic to facilitate human memory encoding and recall. *Journal of Neural Engineering*, 15(3), Article 036014. https://doi.org/10.1088/1741-2552/aaaed7

Hogri, R., Bamford, S. A., Taub, A. H., Magal, A., Del Giudice, P., & Mintz, M. (2015). A neuro-inspired model-based closed-loop neuroprosthesis for the substitution of a cerebellar learning function in anesthetized rats. *Scientific Reports*, 5(1), Article 8451. https://doi.org/10.1038/srep08451

Jolivet, R., Schürmann, F., Berger, T. K., Naud, R., Gerstner, W., & Roth, A. (2008). The quantitative single-neuron modeling competition. *Biological Cybernetics*, 99(4–5), 417–426. https://doi.org/10.1007/s00422-008-0261-x

Josselyn, S. A., & Tonegawa, S. (2020). Memory engrams: Recalling the past and imagining the future. *Science*, 367(6473), Article eaaw4325. https://doi.org/10.1126/science.aaw4325

Joy, B., Cai, Y., Bono, D. C., & Sarkar, D. (2022). Cell Rover—A miniaturized magnetostrictive antenna for wireless operation inside living cells. *Nature Communications*, 13(1), Article 5210. https://doi.org/10.1038/s41467-022-32862-4

Kagan, B. J., Kitchen, A. C., Tran, N. T., Habibollahi, F., Khajehnejad, M., Parker, B. J., Bhat, A., Rollo, B., Razi, A., & Friston, K. J. (2022). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23), 3952–3969.e8. https://doi.org/10.1016/j.neuron.2022.09.001

Le Masson, G., Renaud-Le Masson, S., Debay, D., & Bal, T. (2002). Feedback inhibition controls spike transfer in hybrid thalamic circuits. *Nature*, 417(6891), 854–858. https://doi.org/10.1038/nature00825

Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature*, 484(7394), 381–385. https://doi.org/10.1038/nature11028

Marblestone, A. H., Zamft, B. M., Maguire, Y. G., Shapiro, M. G., Cybulski, T. R., Glaser, J. I., Amodei, D., Stranges, P. B., Kalhor, R., Dalrymple, D. A., Seo, D., Alon, E., Maharbiz, M. M., Carmena, J. M., Rabaey, J. M., Boyden, E. S., Church, G. M., & Kording, K. P. (2013). Physical principles for scalable neural recording. *Frontiers in Computational Neuroscience*, 7, Article 137. https://doi.org/10.3389/fncom.2013.00137

Marder, E., & Goaillard, J.-M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*, 7(7), 563–574. https://doi.org/10.1038/nrn1949

Martins, N. R. B., Angelica, A., Chakravarthy, K., Svidinenko, Y., Boehm, F. J., Opris, I., Lebedev, M. A., Swan, M., Garan, S. A., Rosenfeld, J. V., Hogg, T., & Freitas, R. A., Jr. (2019). Human brain/cloud interface. *Frontiers in Neuroscience*, 13, Article 112. https://doi.org/10.3389/fnins.2019.00112

MICrONS Consortium. (2025). Functional connectomics spanning multiple areas of mouse visual cortex. *Nature*, 640(8058), 435–447. https://doi.org/10.1038/s41586-025-08790-w

MIT News. (2025). Press release on wireless brain implants delivered through the bloodstream ('circulatronics'). https://news.mit.edu/2025/new-therapeutic-brain-implants-defy-surgery-need-1105

Mogensen, A. L. (2023). *Welfare and felt duration* (GPI Working Paper No. 14-2023). Global Priorities Institute. Published in *Noûs*, 60(3), 563–575 (online 2025, issue dated 2026). https://doi.org/10.1111/nous.70022

Müller, J., Ballini, M., Livi, P., Chen, Y., Radivojevic, M., Shadmani, A., Viswam, V., Jones, I. L., Fiscella, M., Diggelmann, R., Stettler, A., Frey, U., Bakkum, D. J., & Hierlemann, A. (2015). High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13), 2767–2780. https://doi.org/10.1039/C5LC00133A

National Academies of Sciences, Engineering, and Medicine. (2021). *The emerging field of human neural organoids, transplants, and chimeras: Science, ethics, and governance.* National Academies Press. https://doi.org/10.17226/26078

Neukart, F. (2026a). A multiscale dynamical framework for continuity-preserving brain replacement. *BioSystems*, 269, Article 105958. https://doi.org/10.1016/j.biosystems.2026.105958

Neukart, F. (2026b). *A theoretical framework for continuity-preserving brain replacement using quantum-enabled nanoscale systems* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.6800880

New Atlas. (2025). Press report on the Cortical Labs CL1. https://newatlas.com/brain/cortical-bioengineered-intelligence

Okun, M. S., Marjenin, T., Ekanayake, J., Gilbert, F., Doherty, S. P., Pilkington, J., French, J., Kubu, C., Lázaro-Muñoz, G., Denison, T., & Giordano, J. (2024). Definition of implanted neurological device abandonment: A systematic review and consensus statement. *JAMA Network Open*, 7(4), Article e248654. https://doi.org/10.1001/jamanetworkopen.2024.8654

Otchy, T. M., Wolff, S. B. E., Rhee, J. Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S. M. H., & Ölveczky, B. P. (2015). Acute off-target effects of neural circuit manipulations. *Nature*, 528(7582), 358–363. https://doi.org/10.1038/nature16442

Pacheco Alvaro, L. M. (in preparation). *Becoming an old part: Incorporation, causal handover and survival through gradual neural replacement.* Manuscript in preparation.

Pehle, C., Billaudelle, S., Cramer, B., Kaiser, J., Schreiber, K., Stradmann, Y., Weis, J., Leibfried, A., Müller, E., & Schemmel, J. (2022). The BrainScaleS-2 accelerated neuromorphic system with hybrid plasticity. *Frontiers in Neuroscience*, 16, Article 795876. https://doi.org/10.3389/fnins.2022.795876

Phoenix, C., & Drexler, E. (2004). Safe exponential manufacturing. *Nanotechnology*, 15(8), 869–872. https://doi.org/10.1088/0957-4484/15/8/001

Powell, D. J., Marder, E., & Nusbaum, M. P. (2021). Perturbation-specific responses by two neural circuits generating similar activity patterns. *Current Biology*, 31(21), 4831–4838.e4. https://doi.org/10.1016/j.cub.2021.08.042

Pozzorini, C., Mensi, S., Hagens, O., Naud, R., Koch, C., & Gerstner, W. (2015). Automated high-throughput characterization of single neurons by means of simplified spiking models. *PLoS Computational Biology*, 11(6), Article e1004275. https://doi.org/10.1371/journal.pcbi.1004275

Prinz, A. A., Abbott, L. F., & Marder, E. (2004). The dynamic clamp comes of age. *Trends in Neurosciences*, 27(4), 218–224. https://doi.org/10.1016/j.tins.2004.02.004

Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7(12), 1345–1352. https://doi.org/10.1038/nn1352

Rabinowitch, I., Upadhyaya, B., Pant, A., Galski, D., Kreines, L., & Bai, J. (2021). Circumventing neural damage in a *C. elegans* chemosensory circuit using genetically engineered synapses. *Cell Systems*, 12(3), 263–271.e4. https://doi.org/10.1016/j.cels.2020.12.003

Randi, F., Sharma, A. K., Dvali, S., & Leifer, A. M. (2023). Neural signal propagation atlas of *Caenorhabditis elegans*. *Nature*, 623(7986), 406–414. https://doi.org/10.1038/s41586-023-06683-4

Sandberg, A., & Bostrom, N. (2008). *Whole brain emulation: A roadmap* (Technical Report 2008-3). Future of Humanity Institute. https://www.fhi.ox.ac.uk/reports/2008-3.pdf

Schreiber, K., Wunderlich, T. C., Pehle, C., Petrovici, M. A., Schemmel, J., & Meier, K. (2020). *Closed-loop experiments on the BrainScaleS-2 architecture* [Preprint]. arXiv:2004.14829.

Scientific American. (2026, March 28). Press report on the Doom demonstration with human neurons on a chip. https://www.scientificamerican.com/article/how-human-neurons-on-a-chip-learned-to-play-doom/

Sharp, A. A., O'Neil, M. B., Abbott, L. F., & Marder, E. (1993). Dynamic clamp: Computer-generated conductances in real neurons. *Journal of Neurophysiology*, 69(3), 992–995. https://doi.org/10.1152/jn.1993.69.3.992

Szücs, A., Varona, P., Volkovskii, A. R., Abarbanel, H. D. I., Rabinovich, M. I., & Selverston, A. I. (2000). Interacting biological and electronic neurons generate realistic oscillatory rhythms. *NeuroReport*, 11(3), 563–569. https://doi.org/10.1097/00001756-200002280-00027

Vladimirov, N., Wang, C., Höckendorf, B., Pujala, A., Tanimoto, M., Mu, Y., Yang, C.-T., Wittenbach, J. D., Freeman, J., Preibisch, S., Koyama, M., Keller, P. J., & Ahrens, M. B. (2018). Brain-wide circuit interrogation at the cellular level guided by online analysis of neuronal function. *Nature Methods*, 15(12), 1117–1125. https://doi.org/10.1038/s41592-018-0221-x

Wang, T., Wang, M., Wang, J., Yang, L., Ren, X., Song, G., Chen, S., Yuan, Y., Liu, R., Pan, L., Li, Z., Leow, W. R., Luo, Y., Ji, S., Cui, Z., He, K., Zhang, F., Lv, F., Tian, Y., … Chen, X. (2022). A chemically mediated artificial neuron. *Nature Electronics*, 5(9), 586–595. https://doi.org/10.1038/s41928-022-00803-0

Weber, C. (2025). Selective optimism about mind-uploading. *Journal of Consciousness Studies*, 32(1–2), 220–240. https://doi.org/10.53765/20512201.32.1.220

Zanichelli, N., Schons, M., Freeman, I., Shiu, P., & Arkhipov, A. (2025). *State of brain emulation report 2025* [Preprint]. arXiv:2510.15745.

Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M., & Häusser, M. (2018). Closed-loop all-optical interrogation of neural circuits in vivo. *Nature Methods*, 15(12), 1037–1040. https://doi.org/10.1038/s41592-018-0183-z

Note on the references. Bibliographic details were taken from bibliographic records (Crossref, PubMed, arXiv) and from the works' own reference lists. A few fields could not be confirmed, and not every work was read in full.