Becoming an Old Part

Incorporation, Causal Handover and Survival through Gradual Neural Replacement

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Unpublished manuscript · version 0.1 · November 2026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. Selective optimism about mind-uploading holds that a person could survive the gradual replacement of their neurons by artificial units but not scan-and-copy uploading. Weber (2025a) grounds this asymmetry in immanent causation: in gradual replacement the later states of the brain are caused mostly by parts that already belong to it, because each replacing part is incorporated, and so becomes an old part, before most of the others are exchanged. The argument therefore rests on a notion, incorporation, that it leaves unanalysed. This paper supplies an account. Drawing on engineering frameworks in which an artificial unit learns a neural unit's behaviour in place and takes over its causal role by degrees, 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. Three consequences follow. What makes a part old is the provenance of its state, not the age of its matter or the time elapsed. Replacing every unit from a static record while brain activity is suspended amounts to copying in place. And incorporation has an empirical signature that separates a substitute which preserves a component's causal contribution from a network that merely compensates for its loss. The resulting thesis about survival depends on the proposed account of persistence and on a functional premise about consciousness.

Keywords: personal identity; mind-uploading; gradual replacement; immanent causation; neural prosthesis; persistence

1. Introduction

Suppose the neurons of a living brain were replaced, a few at a time, by artificial units that do what the neurons did, until none of the original cells remained. Would the person whose brain it was still exist at the end?

Much of the debate about mind-uploading has assumed that the answer must match the answer for scan-and-copy procedures, in which a brain is scanned and a functional replica is built from the record. Optimists such as Wiley and Koene (2016) hold that both procedures preserve the person. Pessimists doubt that either does (Corabi & Schneider 2012; Schneider 2019). The intermediate view, on which one could survive gradual replacement but not scan-and-copy, has often seemed unprincipled. Chalmers (2010; 2014) calls gradual uploading the safest route without claiming that the alternatives are fatal, and Olson (2017) doubted that anyone would maintain that gradual uploading is possible while all-at-once uploading is not.

Weber (2025a) has now defended exactly that view, under the name *selective optimism*. His argument turns on immanent causation: the later states of a persisting thing are caused by its own earlier states. In gradual replacement, he argues, each new part is *incorporated*, and so becomes an old part, before most of the remaining parts are exchanged. At every

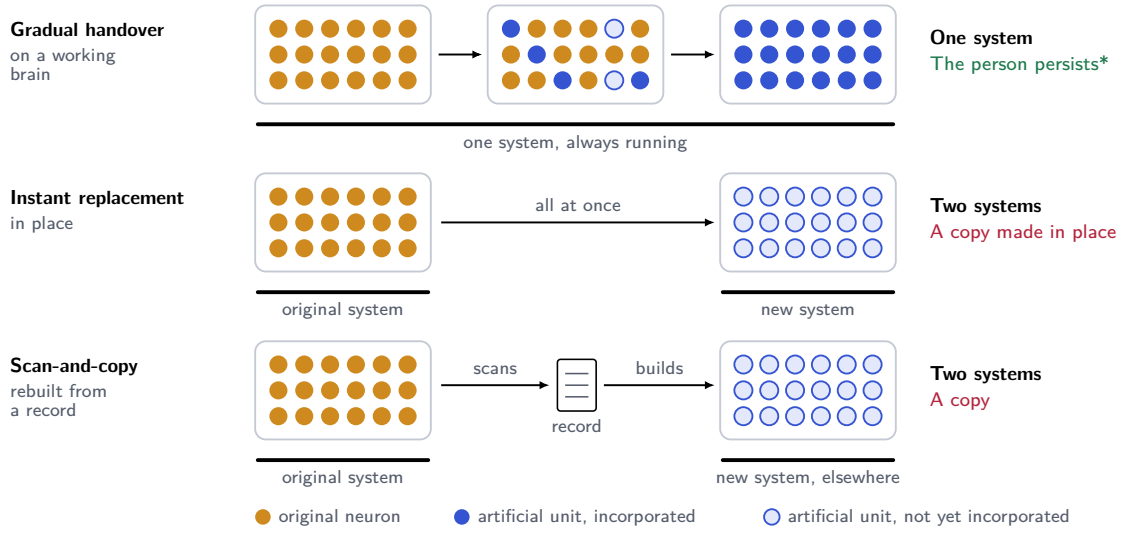


Figure 1. Where the successor's state comes from. In a gradual handover there is one system that never stops working and incorporates units as it goes. In instant replacement and in scan-and-copy the successor's whole state is imposed from outside, and there are two systems. The asterisk marks the conditional element: the person persists if the artificial units sustain mental life (Functional Premise, §2.1).

moment the brain's next state is therefore caused mostly by parts that already belong to it. In instant replacement and in scan-and-copy the successor's state is caused entirely by new parts, and nothing persists.

I think this is the right shape for an answer, and the present paper builds on it. But the argument carries a debt. Everything depends on incorporation, and incorporation is left unanalysed. In Weber's illustrative model it is a fixed delay after which a new part counts as old; the replacing parts are stipulated to be functionally equivalent; and nothing is said about how a replacing part comes to be in the state that makes it equivalent. Until the debt is paid, the critic's question simply moves down a level. Why is installing one prefabricated artificial neuron not a small act of scan-and-copy? And what, apart from the passage of time, turns it into a part of the brain?

A second literature bears on this question without addressing it. Recent engineering frameworks for continuity-preserving brain replacement describe how an artificial unit might take over the role of a neural unit: it is embedded while the original keeps working, learns the original's behaviour, runs in parallel without influencing anything, and then assumes the causal role by degrees (Neukart 2026a; 2026b). These frameworks explicitly decline to say whether the person survives. To my knowledge the two literatures do not cite each other.

This paper connects them. I make four claims.

First, I propose an account of incorporation (§4). An artificial unit is incorporated into a brain to the extent that the brain depends on it as it depended on the component it replaced, the unit responds to the brain's current state, and the state through which the unit acts was produced by the brain's own activity.

Second, the account implies that what makes a part *old* is the provenance of the state through which it acts. It is not the age of its matter, and it is not the time it has spent in place (§4.3). This gives a principled answer to the charge that each step of a gradual replacement is a small copy.

Third, the account gives verdicts on three test cases that no rival criterion gives jointly: a pause in brain activity, replacement performed during such a pause, and reassembly of the removed neurons (§5). The second case shows that 'gradual' is not, at bottom, a temporal notion.

Fourth, incorporation has an empirical signature. Whether a substitute has taken over a component's causal contribution, or the surrounding network has merely compensated for the component's loss, can be tested, and current acceptance criteria for replacement do not test it (§6).

Two limits should be stated at once. The thesis about survival depends on the functional premise in §2 and the account of personal persistence in §4.5. And I say nothing here about what may be done to a fully replaced brain afterwards, such as copying or accelerating it. Those questions inherit none of the support offered below.

2. Two questions and a negative result

2.1 Consciousness and identity

Gradual replacement raises two questions that are often run together. The first is whether the system that results has a mental life at all. The second is whether, if it does, its subject is the person who began the procedure. Chalmers (2014) treats them separately, and with reason. His fading-qualia and dancing-qualia arguments (Chalmers 1995) concern the first question only. On the second he observes that personal identity cannot be a matter of organisation alone, since a molecule-for-molecule duplicate of me would share my organisation without being me.

I shall keep the questions apart by isolating the first in a premise.

> **Functional Premise (FP).** If each functional unit of a brain is replaced by a physical unit that reproduces that unit's causal profile at whatever grain mental life depends on, including its electrical and chemical interactions with its neighbours, its state, and the rules by which that state changes, then the resulting system retains the capacity for mental life that the original system had.

FP is contested. The fading-qualia argument supports it: if experience faded or vanished while every unit went on doing what its neuron did, the person would go on judging and reporting that nothing had changed, and would be systematically wrong about her own experience, which Chalmers (1995) finds implausible. But Mogensen (2025) shows how the argument can be resisted by combining vagueness with a holistic view of the neural basis of experience. Biological naturalists hold that consciousness depends on properties of living, self-maintaining systems that artificial units may not share (Seth 2025; Godfrey-Smith 2016). Integrated information theory denies consciousness to simulations while allowing it to hardware with the right physical causal structure (Tononi & Koch 2015). And Cao (2022) argues that the more finely a functional role is specified, the fewer kinds of thing can occupy it.

I do not adjudicate these disputes. Two features of FP are meant to keep it as modest as the thesis allows. It concerns physical units that sit in the brain and interact with the remaining tissue, not software simulations. And it leaves the grain open, at a price that should be acknowledged: if the relevant grain turns out to be very fine, the units that satisfy FP will be nearly biological, and much of the practical motivation for replacement will be lost. Everything that follows about survival has FP as an antecedent. The work of this paper lies in the second question.

2.2 What everyday turnover shows

One popular route to the second question should be set aside first. It runs: the matter of our brains changes continually, and we lose neurons throughout life without ceasing to exist, so replacing neurons cannot be what would kill us.

The facts are roughly as claimed, with one important exception. Most brain proteins are renewed within days or weeks (Fornasiero et al. 2018), and estimates suggest that around a tenth of neocortical neurons are lost over adult life (Pakkenberg & Gundersen 1997). But the neurons of the human neocortex are not themselves renewed: they are generated during development and those that survive are as old as the person (Bhardwaj et al. 2006). The natural precedent is molecular turnover within persisting cells, and loss of cells without replacement. There is no natural precedent for the replacement of cells.

What follows is limited. The argument shows that persistence does not require sameness of matter, and that no particular neuron is essential. It does not show that replacing every neuron is survivable. Surviving the loss of each neuron taken singly does not entail surviving the substitution of all of them, and loss differs from substitution in a way that matters: a lost neuron is silent, whereas a faulty substitute is active. In Chalmers (2014) the appeal to everyday neural death has the form of a reductio against one way of resisting the step-by-step argument. It is not offered as a positive proof.

The observation nevertheless contains a clue to which I return in §4.3. A cortical neuron whose proteins have all been exchanged many times over is still, by anyone's lights, an old part of the brain. Whatever makes a part old, it is not the age of its matter.

3. Selective optimism and the debt it leaves

Wiley and Koene (2016) argue that gradual replacement and scan-and-copy are metaphysically equivalent. Scan-and-copy, they say, differs from instant replacement in place only by the distance over which the function is displaced, and instant replacement differs from gradual replacement only in speed. Since neither distance nor speed could plausibly mark the difference between surviving and not, the three procedures stand or fall together.

Weber (2025a, §3) rejects the second step. The rate at which a thing's parts are replaced is widely taken to matter to its persistence: a ship or a tree survives the gradual exchange of its parts but not the sudden exchange of all of them. Rate matters, Weber argues, because it tracks something deeper. Following Zimmerman (1997) and Olson (2010), he takes a persisting object to be one whose later states are immanently caused, that is, caused by its own earlier states and not by something else. He then proposes that the level of immanent causation between a state and its successor is fixed by how far the successor is produced by old parts as opposed to new ones. New parts do not stay new: they are incorporated and become old. Two requirements follow. Each step must leave enough old parts in place, and replacement must not outpace incorporation. A model with a hundred equal parts and a fixed incorporation delay shows the level of immanent causation falling as the rate of replacement rises.

Applied to uploading (§4), the proposal yields a contrast. In gradual replacement the artificial neurons take part in the working of the hybrid brain and so become old parts; there is one persisting piece of mental hardware. In instant replacement and in scan-and-copy the successor's state is produced entirely by new parts; there are two pieces of hardware, one destroyed and one created. Combined with the view that continuity of mental hardware and its functions is necessary for a person's survival, this supports selective optimism.

Weber presents the conclusion as a rational stance, not a demonstration, and accepts that intermediate rates may yield indeterminate cases (§5).

I accept this framework. Four things remain to be supplied.

What incorporation is. Weber says that replacing units become old parts by being actively involved in the functioning of the hybrid brain. As it stands this is too permissive. A device that fires on a fixed schedule is involved in the brain's functioning. So is a unit whose output the surrounding network has learned to ignore: it sits in the circuit while the brain goes on working. Neither has taken over anything. We need a notion that separates a part that bears load from one that is merely present.

How the model generalises. The parts of the model are equal and the delay is constant. Neurons differ greatly in how much of the brain's next state they help to produce, and nothing in the account says what determines the delay.

Where the new part's state comes from. An artificial unit must be in the right state to stand in for a particular neuron. Weber stipulates functional equivalence and moves on. But this is where Wiley and Koene's challenge can be renewed. If the unit is built to a specification read from the neuron it replaces, each step looks like scan-and-copy performed on one neuron. If scan-and-copy brings no immanent causation with it at the scale of the brain, why should it do so at the scale of the unit?

What role activity plays. If active involvement is what incorporates, we should ask what happens when the brain is not active. This turns out to matter (§5).

A fifth gap is acknowledged by Weber himself. His premise makes continuity of hardware necessary for survival; the conclusion that gradual replacement is survivable needs more. In the present paper the extra is supplied by FP together with an explicit sufficiency premise (§4.5).

4. Incorporation as causal handover

4.1 The handover

Engineering work on neural replacement describes the transition from a biological component to an artificial one as a staged handover. I follow the formulation in Neukart (2026b, §§2.2, 5, 11), which is the most explicit I know of. The artificial unit is introduced while the biological component continues to operate. It observes the component's inputs and outputs and estimates its internal state. It then runs in *shadow mode*: it computes what the component will do, and its predictions are compared with what the component does, but it drives nothing downstream. Only when its predictions hold across the range of conditions the tissue passes through is causal responsibility transferred, by degrees, with a parameter that moves from wholly biological to wholly artificial while local and global deviations are monitored. After the transfer the unit remains subject to the activity-dependent changes that the rest of the tissue imposes on it.

The scenario I take as central is the one Chalmers (2014) describes, in which nanoscale devices attach to neurons, learn their behaviour and then replace them. Imagine such devices spreading through a working brain much as an infection spreads, a scattered few cells at a time: each attaches to one unit, shadows it, and takes over its role while the person goes about her life. The devices are an idealisation, since nothing of the kind can at present be built, and nothing in the argument requires that they replicate themselves. What matters is the pattern. The brain never stops working, the units converted at any moment are few and dispersed, and each has taken up its role before its neighbours are converted.

Two features of this literature should be noted. The component in question need not be a neuron. Neukart argues that the right unit is the smallest domain whose replacement preserves the relevant dynamics, which may be a synapse, a dendritic branch, a neuron, a neuron together with its glial and vascular surroundings, or a small circuit. 'Neuron-by-neuron' is shorthand, and I use it as such. And the framework is explicit that it does not settle personal identity: gradual replacement is said to be perhaps necessary for continuity but not something that can at present be claimed to be sufficient.

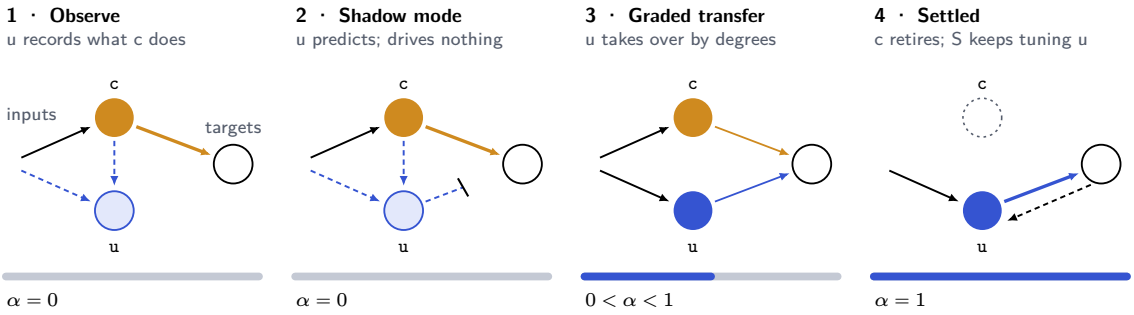


Figure 2. The causal handover, following the description in Neukart (2026b). The parameter α measures how much causal responsibility has passed from the biological component c to the artificial unit u. The unit acquires its state from the system while the system is running.

4.2 Three conditions

The handover suggests what incorporation consists in. Let S be the system that realises a person's capacity for mental life, c one of its components, and u an artificial unit that has taken over the role of c.

> **Incorporation (INC).** u is incorporated into S to the extent that:
>
> **(D) Dependence.** The later states of S depend on the states of u as they depended on those of c: an intervention on u propagates through S as the same intervention on c would have done.
>
> **(R) Responsiveness.** The states of u depend on the current states of S through the channels by which those of c did, so that what u contributes is generated by the present condition of S and not by a schedule or a recording.
>
> **(E) Endogenous state.** The state through which u makes its contribution was produced by the activity of S itself: acquired from c and its neighbours while S was operating, and afterwards altered only by the activity-dependent processes that alter the other parts of S.

Each condition excludes a kind of case that mere presence in a functioning brain lets through. Dependence excludes the passenger: a unit that occupies the position of c while the network has rerouted around it. Responsiveness excludes the pacemaker: a unit that emits a plausible pattern regardless of what the brain is doing. Endogenous state excludes the preloaded part: a unit whose parameters were fixed from a specification drawn up elsewhere.

All three conditions admit of degree, and satisfying them takes time. Shadow operation has to cover enough of the conditions the system passes through for the unit to generalise. Transfer has to be slow enough for deviations to be detected. Plastic adjustment has its own time course. Incorporation is therefore not a fixed delay. It is a process whose duration

depends on the unit, on the region, and on how the system is being used. This is what Weber's constant stands for.

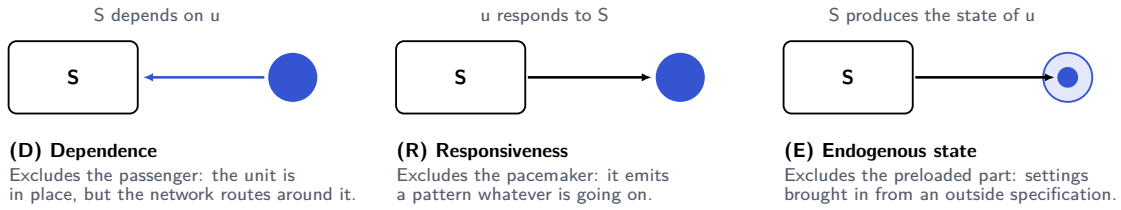


Figure 3. The three conditions of incorporation (INC). A unit u is incorporated into a system S to the extent that all three hold.

4.3 What makes a part old

Condition (E) is the addition, and it changes how the distinction between old and new parts should be understood.

Being an old part cannot consist in the age of one's matter. As §2.2 noted, the proteins of a cortical neuron are exchanged within weeks while the neuron remains an old part of the brain throughout. Nor can it consist in time spent in place: a fragment of shrapnel lodged in the cortex for decades is not a part of the brain's functional organisation at all. What the long-lived neuron has, and the shrapnel lacks, is this: the state in virtue of which it affects what the brain does next was produced, and is continually maintained, by the brain's own activity. Its matter is imported; its organisation is home-grown.

I propose that this is what being an old part is. A part is old to the extent that the state through which it acts on the system is a product of the system's own prior operation. The level of immanent causation at a time is then the share of the causal influence on the system's next state that is exerted through such states.

On this understanding the three procedures differ as follows. In scan-and-copy, the whole state of the successor is imposed at once from a record made by an external device. Nothing through which the successor acts was produced by the successor's own earlier operation; it has no old parts. Instant replacement in place is the same in this respect, whatever its location. In a handover, the matter and the mechanism of each unit are imported, as nutrients are, but the state that makes the unit a stand-in for this component in this brain is acquired from the running system through the channels the system itself uses, tested against the system's behaviour, and thereafter kept and altered by the system. The unit becomes old as (D), (R) and (E) come to hold.

This is where the turnover of §2.2 does positive work. It cannot show that replacement is survivable. It does show that our ordinary notion of an enduring brain part was never tied to enduring matter, and so that a criterion in terms of system-produced state is not an ad hoc device introduced to favour artificial units.

4.4 Weights and schedules

If immanent causation is a share of causal influence, parts must be weighted by their influence. A hundred equal parts make for a clear model and a poor description of a brain. Two consequences follow.

The requirement that each step leave enough old parts in place cannot be read as a requirement on the number of units replaced. Exchanging a small, highly connected popu-

lation at once may lower the level of immanent causation more than exchanging a far larger number of weakly connected units spread across the brain. This agrees with the engineering constraint that replacement be limited by the load on each neighbourhood and not by a global percentage (Neukart 2026b, §5.4).

Rate, likewise, is not a number of units per second. Estimates of that kind are sobering: a thousand neurons per second gives a procedure of about a thousand days (Wiley & Koene 2016, as discussed by Weber 2025a), and one neuron per second gives millennia (Piccinini 2021). But the criterion does not forbid replacing many units in parallel, so long as they are causally remote enough from one another that the weighted share of influence exerted by incorporated parts stays high. Whether any schedule could be both fast enough to be practicable and slow enough to satisfy the criterion is an open empirical question. Nothing here assumes that it could.

4.5 The argument in outline, and the unit-scale challenge

The position can now be set out. The sufficiency claims apply to non-branching histories, where the predecessor's incorporated organisation continues along a single causal line. Section 5.3 considers symmetric fission separately.

1. *Realiser.* A person's persistence requires that the system realising her capacity for mental life persist. The persistence of that system with the capacity intact is sufficient for the person's persistence.

2. *Immanent causation.* In the cases considered here, a material system persists through replacement if its organisation is preserved and, at each step, its next state is produced predominantly through already incorporated parts. At a minimum, persistence requires that the successor's state be produced in part through the predecessor's organisation, carried by incorporated parts in operation or retained through a pause.

3. *Incorporation.* A part is incorporated to the extent that (D), (R) and (E) hold of it.

4. *Handover.* In a handover on a working brain that leaves enough incorporated parts at each transfer, weighted by causal influence, and keeps replacement within the pace of incorporation, the sufficient condition in premise 2 holds at each step.

5. *Functional Premise.* If the units reproduce the relevant causal profiles at the grain specified in §2.1, the resulting system retains its capacity for mental life.

In a non-branching handover meeting these conditions, premises 2 to 4 give persistence of the realiser. If the replacements meet FP's antecedent, premise 5 supplies the retained capacity, and the sufficient condition in premise 1 gives the person's persistence. A replica assembled solely from a record, at a distance or in place, fails the necessary condition in premise 2. The necessity claim in premise 1 therefore rules out the person's survival by that route.

Premise 1 retains Weber's necessity claim and adds the sufficiency claim. It is the view that we go where the realiser of our mental capacities goes, held in different forms by Unger (1990; 2000) and by those whom Parfit (1984) calls narrow psychological theorists. Dainton (2008) defends a neighbouring view in terms of the capacity for consciousness. I assume this account of personal persistence here.

The rationale for the necessary condition in premise 2 is that a single record can initialise several equally faithful replicas. Agreement with the record leaves open which, if any, continues the original system. The criterion traces the successor's state through the incorporated parts of the original system. In a handover these parts help produce the next state. In a temporary pause they retain the organisation from which activity resumes. In scan-and-copy a record specifies the starting state of a newly assembled system.

The challenge left open in §3 can now be met. It is true that each unit is manufactured, and that its state at takeover was fixed by a process of measurement. To that extent each step contains an element that is not immanent, and I see no point in denying it. But three things distinguish the step from scan-and-copy. The state was acquired from the system while it was running, through the system's own operative channels, and was checked against the system's behaviour before it was allowed to matter. At the moment of takeover, almost all of the influence on the system's next state is exerted through parts that are old. And after takeover the unit's state is kept and changed by the system. In scan-and-copy none of the three holds at any scale. The asymmetry therefore does not rest on location or on speed as such. It rests on where the successor's state comes from. Engineers will recognise the difference between replacing a component of a system that stays in service and restoring a system from a backup.

5. Three cases

A criterion earns its keep on cases where its rivals disagree. I consider three, and compare INC with three simpler views: that the person goes with the original tissue; that the person persists only while brain activity is uninterrupted; and that the person goes wherever her psychological pattern is realised, whatever the cause.

5.1 A pause

In operations on the aortic arch performed under deep hypothermic circulatory arrest, the patient is cooled until the electroencephalogram is flat and circulation is then stopped. In a series of 109 monitored patients, electrocerebral silence appeared at a mean nasopharyngeal temperature of about 18 °C (Stecker et al. 2001). Arrest typically lasts some tens of minutes, during which cerebral metabolism falls to a fraction of its normal rate (Ziganshin & Elefteriades 2013). In the one study I have found that asked the question directly, patients in cognitively demanding professions and those close to them reported no change in the patients' abilities or character after arrests of 17 to 54 minutes, by comparison with surgical controls (Percy et al. 2009).

Briefer interruptions are routine. Under deep general anaesthesia the electroencephalogram can show burst suppression, in which periods of electrical silence alternate with bursts of activity (Brown, Lydic & Schiff 2010).

A flat scalp recording is compatible with residual neuronal activity (Kroeger et al. 2013). The study of reported capacities is small and relies on questionnaires. The clinical findings motivate the distinction between measured electrical silence and a pause in the operation of the whole system. The case considered here stipulates such a pause while the parts retain the organisation from which activity can resume.

INC accommodates the case without strain. Dependence and responsiveness are dispositional: they concern how the system and the unit would affect one another when the system operates. Endogenous state is historical. None lapses because activity is suspended, so a

part that was old before the pause is old after it. This fits the view that what must persist is the capacity for mental life and its realiser, not the exercise of the capacity (Dainton 2008; Unger 2000).

The case does not favour INC over the tissue view, since the tissue persists too. For that we need the next case.

5.2 Replacement at a standstill

Suppose that during such a pause every unit of the brain is replaced, one after another, over many hours, each artificial unit being configured from measurements of the quiescent tissue. Activity is then restored.

By the clock the procedure is gradual, and at each step only a small fraction of the system is exchanged. If rate is counted in units per second, Weber's two requirements are met. Yet on his own account of what incorporates a new part, namely active involvement in the brain's functioning, no unit installed during the standstill is incorporated before the next is installed. When activity resumes, the state of the system is carried entirely by parts none of which has yet taken any part in its working.

INC yields the same verdict and locates its source. Every unit fails (E): its state was fixed from a static record and not acquired from the system in operation. Conditions (D) and (R) are met at best dispositionally and have never been exercised. At the restart the share of causal influence exerted through system-produced states is nil. This is scan-and-copy carried out in place and at leisure.

Three lessons follow.

'Gradual' is not at bottom a temporal notion. The rate that matters is replacement per unit of *incorporating activity*, not per unit of time. A procedure can be as slow as one likes and incorporate nothing, and the intuition that slowness helps is explained by the fact that in a working brain elapsed time and incorporating activity go together.

Activity plays two roles that must be kept apart. It is not required for a system to persist: old parts stay old through a pause (§5.1). It is required for a new part to become old. This is the grain of truth in the thought that the process must keep going. The thought is false as a condition on persistence and true as a condition on incorporation.

Interleaving is permitted. If one per cent of the units were exchanged by handover, the brain then worked normally for a time, and the cycle was repeated with pauses in between, INC would count the result as persistence. Pauses are harmless so long as incorporation occurs between replacements.

The verdict is open to dispute, and it is the place where INC is most exposed. A ship whose planks are all replaced in a single stay in dry dock is, for many, still the same ship, although no new plank bears a load before the last is fitted. One may conclude that structural integration is enough for incorporation, and that the standstill procedure is therefore survivable. My reply is that for a ship there is nothing more to integrate. For a system whose identity as a realiser consists in a dynamical organisation that it builds and maintains by operating, the state of a part is what has to be integrated, and structural fit does not secure where that state came from. Those who are not persuaded will give the opposite verdict. The case is at least a clean test: it holds place, clock rate and step size fixed, and varies only whether incorporation occurs.

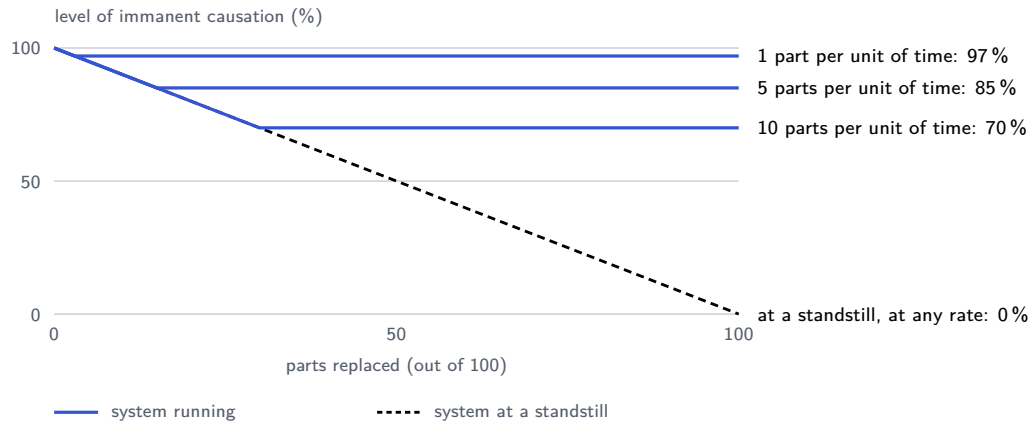


Figure 4. Weber's toy model (2025a, §3.2) with an incorporation time of three units: an object of a hundred equal parts, whose level of immanent causation is the share of its next state produced by parts already incorporated. With the system running, the level settles at a plateau that depends on the replacement rate. At a standstill (§5.2) nothing is incorporated, and the level falls to zero at any rate.

5.3 Reassembly

Suppose finally that the neurons removed in the course of a handover are kept alive and later reassembled into a working brain, as in Hobbes's version of the ship of Theseus. There are now two candidates: the artificial brain in the original body, and the biological brain made of the original cells.

On the tissue view the reassembled brain has the better claim. On the pattern view both have a claim and neither a better one. INC is unequivocal. The neurons were removed one at a time, and their mutual connections and states had to be reconstructed from a record. The reassembled brain is a copy made from original material, as Weber (2025a, §3.2), following van Inwagen (1978), says of a statue rebuilt from its original matter. The system in the body is the one whose every state was produced through its own incorporated parts.

It matters that this verdict does not depend on a comparison between candidates. Closest-continuer theories (Nozick 1981) make identity depend on whether a rival exists, and the charge that this renders survival extrinsic is pressed against non-branching clauses generally by Weber (2025b). INC is a condition on a system's own causal history. The system in the body would be the original whether or not the removed neurons were ever reassembled. In this it also differs from Piccinini's (2021) condition that a process preserves a mind only if at most one survivor can result: that condition is modal, and the bare possibility of reassembly threatens it.

One case INC leaves open is symmetric fission. If the hemispheres were separated and each completed by gradual handover, both products could satisfy the local conditions of incorporation and handover. A verdict on which, if either, is the original would require a further principle. I treat non-branching as a limit on where these conditions suffice to establish numerical identity.

5.4 Summary

Table 1. Verdicts of four criteria on six cases.

Case	Tissue view	Uninterrupted-activity view	Pattern view	INC
Pause with organisation preserved	survives	does not survive	survives	survives
Handover on a working brain	does not survive	survives	survives	survives
Total replacement at a standstill	does not survive	does not survive	survives	does not survive
Instant replacement in place	does not survive	unclear	survives	does not survive
Destructive scan-and-copy	does not survive	does not survive	survives	does not survive
Reassembly: which is the person?	the reassembled brain	the brain in the body	no answer	the brain in the body

No rival view matches INC's column. It is the only one of the four on which the person survives both the pause and the handover while failing to survive replacement at a standstill.

6. An empirical signature

Conditions (D) and (R) are empirical, and this matters for more than engineering. If incorporation is what grounds persistence, then whether a given substitution incorporated anything is a question of fact, and success at the level of behaviour does not answer it.

The difficulty is well documented. The same network activity can arise from widely different underlying parameters (Prinz, Bucher & Marder 2004). The effect of silencing part of a circuit can be large at first and vanish within days as the rest of the circuit adapts, so that acute and chronic manipulations give different answers (Otchy et al. 2015). The measurable influence of a single cortical neuron on its neighbours is small (Chettih & Harvey 2019). And a function restored by an engineered connection may turn out to run through routes other than the intended one (Rabinowitch et al. 2021). A brain that goes on working after a neuron has been exchanged for an artificial unit may therefore be working *with* the unit or working *around* it. In the second case the unit is a passenger, (D) fails, and nothing has been incorporated. A replacement that proceeded through thousands of such steps would be a slow ablation concealed by compensation.

The staged roadmaps in the engineering literature do not address this. Their acceptance criteria are bounds on the deviation of the hybrid system from its baseline (Neukart 2026b, §11). A network that compensates perfectly meets such a bound. What is needed are conditions that distinguish the two hypotheses, and the materials for them exist in the experimental literature.

Table 2. Experimental conditions and what each hypothesis predicts.

Condition	What it isolates	If the substitute preserves the contribution	If the network compensates
1. Acute silencing, no substitute	That the component matters at all	Immediate deficit	Immediate deficit
2. Chronic silencing, no substitute	The time course of spontaneous compensation	Reference curve	Reference curve

Condition	**What it isolates**	**If the substitute preserves the contribution**	**If the network compensates**
3. Faithful closed-loop substitute	The experimental condition	Baseline function at once, with no transient	Transient deficit, recovery along the curve of condition 2
4. Sham: full instrumentation, component intact	Effects of the intervention itself	Baseline	Baseline
5. Replay of a recorded output, no contingency	Whether responsiveness matters	Worse than 3, above all for novel inputs	As good as 3
6. Substitute with the wrong dynamics	Whether fidelity matters	Worse than 3	As good as 3
7. Graded infidelity	The relation between fidelity and function	Function degrades with infidelity	Flat
8. Perturbation of the substitute	That it lies on the causal path	Influence profile matches that of the original	Absent, weak or different
9. Reversal after hours or days	Hidden compensation	The deficit of condition 1 returns	Nothing happens

Preservation is supported only if condition 3 matches baseline from the outset, outperforms conditions 5 and 6, degrades as fidelity falls in condition 7, yields a matching influence profile under condition 8, and reverses under condition 9. If the readout with the substitute off is equivalent to the readout with it on, despite an acute deficit under condition 1, compensation is one explanation. If replay or a mismatched unit does as well as the faithful one within the pre-set margin, the test is inconclusive about whether the specific contribution was preserved. Acute silencing must establish a deficit beyond a pre-set margin for the preservation test to proceed. After confirmed silencing, equivalence to baseline supports redundancy for that readout. An estimate compatible with both equivalence and a meaningful deficit leaves the result inconclusive. Because circuits with similar activity can respond differently to the same disturbance, the comparisons should use perturbations that were not used to fit the substitute (Powell, Marder & Nusbaum 2021).

Parts of this design have been run separately. An electronic neuron coupled to a small crustacean circuit in place of a damaged cell restored the circuit's rhythm (Szücs et al. 2000). A neuromorphic bridge between two cultured networks that had been disconnected partially restored their correlation, which fell again when the bridge was switched off (Buccelli et al. 2019). In rats with cortical injury, stimulation contingent on recorded activity produced markedly better recovery than non-contingent stimulation (Guggenmos et al. 2013), which is condition 5 in another guise. To my knowledge the conditions have not been assembled into a single test of whether a substitute has taken over a component's contribution.

Three limits should be stated. Present substitutes replace the electrical output of a component to its targets; they do not reproduce chemical release or the plasticity of the component replaced, so the experiments would test a weak version of FP's antecedent. A positive result establishes causal equivalence at the grain measured, not experience and not identity. And condition (E) is secured by the design of the handover, not verified afterwards, although it has an observable correlate: a unit whose state was acquired from the running system should generalise to conditions it was not explicitly fitted to.

The companion protocol places two constraints on the pace of replacement. It can estimate the shortest tested duration of shadow fitting before an exchange and the shortest tested spacing between exchanges that meet its joint criterion for preservation. Both are specific to a preparation, substitute and readout. They make part of Weber's timing requirement testable. Timing the gradual handover itself would require tracking when the unit takes on its causal role during the transfer. The full protocol is given in a companion paper (Pacheco Alvaro, in preparation).

7. Objections and limits

7.1 Does identity matter?

Parfit (1984) compares a case in which a hundredth of a brain is replaced in each of a hundred operations with one in which the whole is removed and all the replacements are then inserted, and asks whether the order of operations could be the difference between life and death. On the present account it is not the order that differs. In the first case there is throughout a working system that produces its own next state and incorporates each replacement. In the second there is an interval with no such system and then one assembled from outside.

This reply does not touch Parfit's deeper claim, which is that identity is not what matters, and that psychological continuity with any reliable cause is about as good as ordinary survival. If that is right, the distinction drawn in this paper is real and unimportant. I have assumed a non-deflationary view on which it matters whether a future person is oneself. The assumption is a cost of the position and should be counted as one.

7.2 If the Functional Premise is false

If the capacity for mental life depends on properties that no artificial unit could have, then a handover would produce, by degrees, a system that behaves as the person did and has no mental life. Nothing in §§4–6 would detect this. Conditions (D), (R) and (E) could all be met by a system that satisfied the letter of the handover and failed FP. The realiser, in the sense of the physical system with its causal lineage, would persist; there would be no one whose realiser it was.

The dependence should be stated plainly. INC is an account of when a system persists through replacement of its parts. It becomes an account of when a *person* persists only given FP. The most I claim is that the two questions are separable, and that an answer to the second is available for anyone who accepts the first.

7.3 Self-production

An organism's turnover is its own doing: each cell makes and installs its replacement parts. An artificial unit is made elsewhere and put in. For those who identify us with living organisms (Olson 1997), or who take a life to be a self-maintaining process, this is a difference in kind.

Part of the reply is given by §4.3. What the organism produces is organisation; the matter comes from outside, and Olson (2017) allows that a thing may cause its own continued existence with outside help. On the present account the system produces the state of the incorporated unit in the relevant sense, although it does not build the unit. Artificial components already take a continuous part in the working of the nervous system without anyone doubting who their bearers are, in the million or so people with cochlear implants

(Zeng 2022) and the many with deep-brain stimulators (Lozano et al. 2019), though none of these devices replaces a central neuron.

The reply is partial. The organism neither builds nor repairs the unit, and as replacement proceeds the self-maintaining fraction of the brain shrinks towards nothing. Those who identify the person with that self-maintaining life will say that the procedure gradually extinguishes it. I can only decline the identification, as Unger (2000) does in holding that biology is causally and not constitutively required. It is worth recalling Weber's (2025a, §4.3) observation that the standard biological view yields survival through gradual brain replacement for its own reasons, since on that view the brain is not essential to the organism; and on versions that tie the organism's persistence to the organ sustaining its characteristic capacities (Madden 2016), a hybrid brain that sustains them would serve. That is agreement in verdict for reasons that owe nothing to the present argument.

7.4 Vagueness and humility

'Predominantly', 'enough' and 'to the extent that' are vague, and INC inherits the dilemma that Wiley and Koene press: somewhere between the working handover and the instant exchange there is either a sharp boundary or a range of indeterminate cases. I follow Weber (2025a, §5) in accepting the latter. It is enough for the thesis that the slow, interleaved handover is determinately persistence and that assembly from a record is determinately not.

Schneider (2019) counsels humility: for all we know even a single neural prosthesis might break whatever continuity matters. That may be so. The present account at least says what such a break would have to consist in, namely a failure of incorporation, and §6 says how one might look for it. An objector who holds that the first prosthesis is fatal owes an explanation of why the continual loss and remodelling of natural components is not.

8. Stopping halfway

I close with a consequence for the ethics of the procedure, since the account bears directly on consent.

Agar (2010) argues that uploading is a bad bet because it carries a risk of death for the sake of benefits obtainable otherwise. Levy (2011) replies that a rule against any act that might amount to death is no guide to action, and that probabilities and starting points matter. The handover alters the decision in two ways. It need not be chosen as a whole. Because it is staged, a person can be asked again at each stage, with her experience of the previous stages to go on. And on the present account every intermediate state is a state of the same persisting system.

Both points have limits. If FP were false, the person's reports at later stages would not be evidence of anything, so the staging gives no protection against the failure one most fears. And the right to stop is not a right to be restored. Engineering proposals require that consent be revocable during reversible stages (Neukart 2026b, §12.1). Beyond those stages, what a person can demand is that no further replacement be made, and that the hybrid she has become be a state in which she can go on living.

That requirement makes it a condition on the design of the procedure that every intermediate state be a viable endpoint, shown to be stable before anyone is brought to it. And it changes the meaning of support. Devices implanted in the nervous system have repeatedly been left without maintenance when their makers failed, and a consensus definition of such abandonment now exists (Okun et al. 2024); recent international guidance addresses contin-

ued access and withdrawal (UNESCO 2025). For a person whose brain is partly or wholly artificial there is no explantation and no useful life of the device short of her own. Obligations of continuity that are today owed for the declared lifetime of a product would be owed for the lifetime of a person.

9. Conclusion

Selective optimism is the view that gradual neural replacement differs from copying in a way that matters to survival. Weber has shown that the view can be grounded in immanent causation. I have argued that the grounding is complete only when we say what it is for a replacing unit to be incorporated, and that the answer is a causal handover: the system comes to depend on the unit, the unit responds to the system, and the unit's state is the system's own product. On that account an old part is one that acts through a state the system produced. 'Gradual' means paced by incorporating activity, not by the clock. Replacement at a standstill is copying in place. And whether a substitution has incorporated anything is a matter that experiment can address, by separating preservation from compensation.

Given the account of persistence defended here, a replica assembled solely from a record begins a distinct causal history. A person whose brain was replaced by a non-branching handover meeting the conditions in §4.5 would survive if FP holds and the replacements meet its antecedent.

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

Agar, N. (2010). *Humanity's end: Why we should reject radical enhancement.* MIT Press. https://doi.org/10.7551/mitpress/9780262014625.001.0001

Bhardwaj, R. D., Curtis, M. A., Spalding, K. L., Buchholz, B. A., Fink, D., Björk-Eriksson, T., Nordborg, C., Gage, F. H., Druid, H., Eriksson, P. S., & Frisén, J. (2006). Neocortical neurogenesis in humans is restricted to development. *Proceedings of the National Academy of Sciences*, 103(33), 12564–12568. https://doi.org/10.1073/pnas.0605177103

Brown, E. N., Lydic, R., & Schiff, N. D. (2010). General anesthesia, sleep, and coma. *The New England Journal of Medicine*, 363(27), 2638–2650. https://doi.org/10.1056/NEJMra0808281

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

Cao, R. (2022). Multiple realizability and the spirit of functionalism. *Synthese*, 200(6), Article 506. https://doi.org/10.1007/s11229-022-03524-1

Chalmers, D. J. (1995). Absent qualia, fading qualia, dancing qualia. In T. Metzinger (Ed.), *Conscious experience* (pp. 309–328). Ferdinand Schöningh / Imprint Academic.

Chalmers, D. J. (2010). The singularity: A philosophical analysis. *Journal of Consciousness Studies*, 17(9–10), 7–65.

Chalmers, D. J. (2014). Uploading: A philosophical analysis. In R. Blackford & D. Broderick (Eds.), *Intelligence unbound: The future of uploaded and machine minds* (pp. 102–118). Wiley. https://doi.org/10.1002/9781118736302.ch6

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

Corabi, J., & Schneider, S. (2012). The metaphysics of uploading. *Journal of Consciousness Studies*, 19(7–8), 26–44.

Dainton, B. (2008). *The phenomenal self*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199288847.001.0001

Fornasiero, E. F., Mandad, S., Wildhagen, H., Alevra, M., Rammner, B., Keihani, S., Opazo, F., Urban, I., Ischebeck, T., Sakib, M. S., Fard, M. K., Kirli, K., Centeno, T. P., Vidal, R. O., Rahman, R.-U., Benito, E., Fischer, A., Dennerlein, S., Rehling, P., Feussner, I., Bonn, S., Simons, M., Urlaub, H., & Rizzoli, S. O. (2018). Precisely measured protein lifetimes in the mouse brain reveal differences across tissues and subcellular fractions. *Nature Communications*, 9(1), Article 4230. https://doi.org/10.1038/s41467-018-06519-0

Godfrey-Smith, P. (2016). Mind, matter, and metabolism. *The Journal of Philosophy*, 113(10), 481–506. https://doi.org/10.5840/jphil20161131034

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

Kroeger, D., Florea, B., & Amzica, F. (2013). Human brain activity patterns beyond the isoelectric line of extreme deep coma. *PLoS ONE*, 8(9), Article e75257. https://doi.org/10.1371/journal.pone.0075257

Levy, N. (2011). Searle's wager. *AI & Society*, 26(4), 363–369. https://doi.org/10.1007/s00146-011-0317-7

Lozano, A. M., Lipsman, N., Bergman, H., Brown, P., Chabardes, S., Chang, J. W., Matthews, K., McIntyre, C. C., Schlaepfer, T. E., Schulder, M., Temel, Y., Volkmann, J., & Krauss, J. K. (2019). Deep brain stimulation: Current challenges and future directions. *Nature Reviews Neurology*, 15(3), 148–160. https://doi.org/10.1038/s41582-018-0128-2

Madden, R. (2016). Human persistence. *Philosophers' Imprint*, 16(17), 1–18.

Mogensen, A. L. (2025). How to resist the Fading Qualia Argument. *Synthese*, 206(5), Article 252. https://doi.org/10.1007/s11229-025-05338-3

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

Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

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

Olson, E. T. (1997). *The human animal: Personal identity without psychology*. Oxford University Press. https://doi.org/10.1093/0195134230.001.0001

Olson, E. T. (2010). Immanent causation and life after death. In G. Gasser (Ed.), *Personal identity and resurrection: How do we survive our death?* (pp. 51–66). Ashgate. Reissued 2016, Routledge. https://doi.org/10.4324/9781315600147

Olson, E. T. (2017). The central dogma of transhumanism. In B. Berčić (Ed.), *Perspectives on the self* (pp. 35–57). Faculty of Humanities and Social Sciences, University of Rijeka.

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). *Preservation or compensation? A protocol for testing whether an artificial substitute takes over a neural component's causal contribution.* Manuscript in preparation.

Pakkenberg, B., & Gundersen, H. J. G. (1997). Neocortical neuron number in humans: Effect of sex and age. *The Journal of Comparative Neurology*, 384(2), 312–320. https://doi.org/10.1002/(SICI)1096-9861(19970728)384:2%3C312::AID-CNE10%3E3.0.CO;2-K

Parfit, D. (1984). *Reasons and persons.* Clarendon Press. https://doi.org/10.1093/019824908X.001.0001

Percy, A., Widman, S., Rizzo, J. A., Tranquilli, M., & Elefteriades, J. A. (2009). Deep hypothermic circulatory arrest in patients with high cognitive needs: Full preservation of cognitive abilities. *The Annals of Thoracic Surgery*, 87(1), 117–123. https://doi.org/10.1016/j.athoracsur.2008.10.025

Piccinini, G. (2021). The myth of mind uploading. In R. W. Clowes, K. Gärtner, & I. Hipólito (Eds.), *The mind-technology problem: Investigating minds, selves and 21st century artefacts* (pp. 125–144). Springer. https://doi.org/10.1007/978-3-030-72644-7_6

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

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

Schneider, S. (2019). *Artificial you: AI and the future of your mind.* Princeton University Press. https://doi.org/10.1515/9780691197777

Seth, A. K. (2025). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, 49, Article e315. https://doi.org/10.1017/S0140525X25000032

Stecker, M. M., Cheung, A. T., Pochettino, A., Kent, G. P., Patterson, T., Weiss, S. J., & Bavaria, J. E. (2001). Deep hypothermic circulatory arrest: I. Effects of cooling on electroencephalogram and evoked potentials. *The Annals of Thoracic Surgery*, 71(1), 14–21. https://doi.org/10.1016/S0003-4975(00)01592-7

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

Tononi, G., & Koch, C. (2015). Consciousness: Here, there and everywhere? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668), Article 20140167. https://doi.org/10.1098/rstb.2014.0167

UNESCO. (2025). *Recommendation on the ethics of neurotechnology.* Adopted by the General Conference at its 43rd session, Samarkand, 11 November 2025. https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology

Unger, P. (1990). *Identity, consciousness and value.* Oxford University Press. https://doi.org/10.1093/oso/9780195054019.001.0001

Unger, P. (2000). The survival of the sentient. *Philosophical Perspectives*, 14 (= *Noûs*, 34, Suppl. 14), 325–348. https://doi.org/10.1111/0029-4624.34.s14.17

van Inwagen, P. (1978). The possibility of resurrection. *International Journal for Philosophy of Religion*, 9(2), 114–121. https://doi.org/10.1007/BF00138364

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

Weber, C. (2025b). The multiplicity objection against uploading optimism. *Synthese*, 205(6), Article 226. https://doi.org/10.1007/s11229-025-05057-9

Wiley, K. B., & Koene, R. A. (2016). The fallacy of favouring gradual replacement mind uploading over scan-and-copy. *Journal of Consciousness Studies*, 23(3–4), 212–235.

Zeng, F.-G. (2022). Celebrating the one millionth cochlear implant. *JASA Express Letters*, 2(7), Article 077201. https://doi.org/10.1121/10.0012825

Ziganshin, B. A., & Elefteriades, J. A. (2013). Deep hypothermic circulatory arrest. *Annals of Cardiothoracic Surgery*, 2(3), 303–315. https://doi.org/10.3978/j.issn.2225-319X.2013.01.05

Zimmerman, D. W. (1997). Immanent causation. *Philosophical Perspectives*, 11 (= *Noûs*, 31, Suppl. 11), 433–471. https://doi.org/10.1111/0029-4624.31.s11.19

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.