

# Spike-Pattern Sufficiency for Manifest Consciousness: A Falsifiable Principle

Edgar Muniz\*

## Abstract

Consciousness is ordinarily judged by first-person report, yet neuroscience shows such reports are fixed by *temporal* patterns of neural spiking. We derive, from two uncontroversial premises—causal efficacy of consciousness and neural mediation—a direct conditional deduction: manifest consciousness supervenes on the temporal spike-event pattern within the behaviourally decisive window. A behaviour-changing intervention that leaves spike timing within the stated tolerance would falsify this claim; none is known. Spike-pattern sufficiency therefore provides a natural working hypothesis, shifting the burden of proof to substrate-specific accounts and supplying a two-track empirical roadmap: (i) falsifiability via spike-silent-channel searches and (ii) verification-and-refinement via controlled spike-pattern duplication.

## 1 Introduction

This paper advances a deductive logical–empirical argument whose implications reach both science and philosophy. It addresses one empirical question: which physical properties suffice for *manifest* consciousness? We set aside the further question of why those properties feel like anything; the hard problem lies outside our scope (Chalmers, 1995). Even within this restricted focus, the contemporary literature on consciousness hosts a profusion of mutually inconsistent hypotheses, from integrated-information accounts (Albantakis et al., 2023) to biological chauvinism. A common diagnosis is *substrate uncertainty*: until we know which physical properties are decisive for experience, every additional property remains a live candidate. This argument significantly narrows that search space. *After presenting the deduction in Section 4, the remainder of the paper provides (i) empirical scaffolding; (ii) an empirical programme for falsification, verification, and refinement; and (iii) replies to leading philosophical objections.*

Although the thesis harmonises with prominent strands of functionalism, we do not assume functionalism as a backdrop. We derive a single functionalist-style commitment from two empirical premises: manifest consciousness supervenes on the temporal spike-event pattern inside the behaviourally decisive window. The result is an explicit, testable demonstration rather than a tacit supposition, offering a specific claim about manifest consciousness while broader metaphysical debates proceed elsewhere.

Our argument rests on two mundane claims.

---

\*Independent researcher, United States. edgarmuniz@mutualalignment.org

- (i) **Principle of causal efficacy (PCE)**—conscious states help cause overt actions, including verbal reports, by which we know they occur.
- (ii) **Principle of neural mediation (PNM)**—all causal routes from brain to behaviour eventually modulate which neurons spike when.

We call the shortest interval preceding the behaviour that still contains *every* spike whose causal ripples can reach that behaviour the *behaviourally decisive window*  $T$ ;<sup>1</sup> the ordered multiset of those spikes is the *spike-event pattern* (SEP).

**Corollary 1.**

- (1) (PCE) Manifest consciousness helps cause the behaviour we measure.
- (2) (PNM) The only causal lever available for manifest behaviour is the temporal spike pattern inside  $T$ .
- (3) Therefore manifest consciousness supervenes on the temporal spike pattern inside  $T$ : realising that pattern suffices for manifest consciousness; any manifest difference in phenomenology entails some difference in that pattern.

We take no stand here on which features of the pattern—exact latencies, relative order, synchrony windows, mean firing rates, or other yet-identified dynamics—carry the decisive signal, nor on how much jitter those features can tolerate. The claim is only that whatever the correct empirical specification is, it resides fully within the temporal spike pattern and nowhere else.

A century of lesion, stimulation and closed-loop neuro-prosthesis work has yet to reveal a single behaviour-changing manipulation that leaves every spike inside  $T$  untouched—a fact that places the empirical burden squarely on the sceptic.

This minimalist conclusion yields a two-track empirical agenda—falsification through spike-silent interventions and verification-and-refinement through systematic spike-pattern duplication—with implications for substrate-specific theories and machine consciousness debates. Our claim is methodological: science can adjudicate consciousness only insofar as it is manifest in behaviour or report. Non-manifest consciousness may exist, but (i) by definition it is empirically inaccessible and (ii) its possibility does not undermine the sufficiency claim for manifest cases.

The next section states the premises formally and presents the boxed deduction.

## 2 Premises & Core Deductive Argument

This section states the premises and definitions formally; readers comfortable with PCE and PNM may skim the definitions and proceed to the core deduction in Section 4.

---

<sup>1</sup>The bound  $T$  is defined conceptually by progressively truncating the recorded interval and applying perturbations; the smallest window after which sub- $\delta$  jitter no longer alters report is taken as  $T$ . Practical implementations are outlined in Section 6. As instrumentation improves, empirical estimates of  $T$  may shrink, but the subsequent argument does not depend on its absolute duration.

## Premise 1 (PCE): Causal efficacy of consciousness

Whenever a subject sincerely reports “I am in pain / conscious / seeing red,” that phenomenal state exerts causal influence somewhere in the physical chain culminating in the report. Treating systematic coincidence between qualia and neural causation as mere accident is methodologically untenable.

**Why accept PCE?** Across anaesthesia, analgesia, affective neuroscience, and psychology, lucid self-report functions as the gold-standard criterion for conscious experience; PCE just makes explicit the causal assumption built into those practices.

*Empirical support appears in Appendix A; interventionist justification in Section 3.*

*Clarifying caveat.* The premise does not require that every self-report be perfectly accurate; it asserts only that, under ideal lucidity, sincere reports are causally tethered to the subject’s experiential state closely enough for the cited fields to rely on them.

**Definition 1** (Spiking Event). A **spiking event** is a discrete, all-or-none output signal emitted by a network node at a particular moment.

- (a) *Discrete*: each spike is represented as a countable event occurring at a particular time, rather than as a continuously varying signal.
- (b) *All-or-none*: the event either occurs or it does not—no graded variants count.
- (c) *Network-agnostic (see Corollary 2 for justification)*: any signal that produces the same downstream causal effects, regardless of implementation, counts as a spiking event; biological action potentials, digital voltage pulses in neuromorphic hardware, and software-level flag flips that satisfy (a) and (b) all qualify.

*Caveat.* The above gives one *sufficient* unit for our argument—the kind known to work in the human brain; whether discreteness and all-or-none amplitude are each strictly necessary remains an empirical question.

*Optional generalised (shape-aware) form.* If within-event waveform features (“spike shape”) are shown to be behaviour-relevant within the behaviourally decisive window  $T$ , treat a spiking event as a timed output with a resolvable mark  $m$  (e.g., width, amplitude, slope), and let the SEP be the ordered multiset  $\{\langle n, t, m \rangle\}$ . All definitions and results below carry through verbatim with  $m$  included; when no such features are resolvable, set  $m = \emptyset$  and the simplified, time-only definition suffices.

**Corollary 2** (Node-level implementation invariance). Because behavior  $U$  is mediated by the spike–event pattern within the decisive window  $T$  (PNM), and manifest experience contributes causally to  $U$  (PCE), any replacement of a node (or set of nodes) that, under the same causal input history over  $T$ , yields the *same pattern of effects on its downstream targets within  $T$* —that is, induces the same contribution to the efferent  $\text{SEP}_T$ —necessarily leaves  $\text{SEP}_T$  unchanged and therefore leaves  $U$  unchanged. By PCE, any difference in *manifest* phenomenology would have changed  $U$ ; since  $U$  is fixed, no manifest difference obtains (residual differences, if any, lie within the cone of manifest invariability; see Definition 5). By iteration, any finite cascade of such event-preserving replacements preserves both behavior and manifest phenomenology.

**Definition 2** (Spike-Event Pattern (SEP)). Over the behaviourally decisive window  $T$ , an **SEP** is the ordered multiset  $\mathbf{S} = \{\langle n_1, t_1 \rangle, \dots, \langle n_k, t_k \rangle\}$  containing every spiking event

(Def. 1) produced by any node whose activity can, via any causal path inside  $T$ , influence an effector that modulates the behaviour under study.

**Definition 3** (Manifest). A phenomenal state is *manifest* if and only if it causally influences observable behaviour (including verbal report).

**Definition 4** (Manifest Phenomenal Parity). Two phenomenal states stand in *manifest phenomenal parity* (MPP) under a given task context  $X$  and reporting policy  $K$  iff, with  $K$  and  $X$  fixed, no difference between them yields any resolution-controlled difference in the measured behaviour stream  $U$ . Equivalently, the states lie within that task’s cone of manifest invariability (see Definition 5).

**Definition 5** (Cone of manifest invariability). Phenomenal differences matter here only insofar as they influence behaviour. Fix a body state and require the entire manifest performance to match—not just the headline report, but also its cadence, microexpressions, and slot-by-slot content (“I am bored of looking at this red wall, and I am tired”). Moving from a binary response (“Yes, I am conscious”) to such a multi-slot template amounts to tightening the reporting policy  $K$ , thereby shrinking the region of phenomenology that can differ without perturbing behaviour. Any residual variation small enough to hide within this cone is, by construction, manifestly inert. Any two states within this cone thereby stand in manifest phenomenal parity (Definition 4) with  $K$  and  $X$  fixed.

*Fixed cardinality.* For a given task and window  $T$ , an SEP is the ordered multiset containing exactly those spikes whose causal influence can reach the behaviour under study; its multiset cardinality is fixed. This fixed cardinality supports a canonical one-to-one correspondence when SEPs are compared; quantitative, timing-based criteria for SEP duplication are introduced in Section 5.

*Caveat.* The argument does not assume which aspects of the pattern—absolute latencies, relative order, synchrony windows, mean firing rates—carry the decisive signal, nor how much jitter those aspects can tolerate; it claims only that at least one empirically sufficient specification—the one realised in the human brain—resides fully within the temporal SEP.

**Empirical Limit** ( $\delta$ -Indifference). No known post-synaptic or circuit-level mechanism relevant to human report can resolve two spike arrivals that differ in timing by less than  $\delta$ . We define  $\delta$  as the tightest in-vivo spike-timing precision so far documented; its value is updated whenever a new record is set.

## Premise 2 (PNM): Neural mediation

Every causal route from brain to behaviour propagates through the precise timing of action potentials. Variables—glial, hormonal, field—affect behaviour only insofar as they modulate *which* neurons spike *when*, thereby shaping the SEP within the decisive window  $T$  (Defs. 2, 1).

**Why accept PNM?** *PNM* restates the classic **final-common-path principle** of motor neuroscience: first articulated by Sherrington in 1906 and still treated as textbook doctrine, it holds that *every* causal chain leading to overt behaviour ultimately converges on the timing of efferent action potentials.

*Empirical support summarised in Appendix B.*

*Clarifying caveat.* The premise does not deny sub-threshold or modulatory influences; it asserts only that such influences affect behaviour *in virtue of* how they alter which

neurons spike when. Because any such variable must exert its behavioural influence *by altering the spike pattern*, it contributes no additional independent degrees of freedom once the SEP is fixed.

To date *no reproducible example* has been documented—and none is predicted under this hypothesis—of a behaviour-changing manipulation that leaves the spike-event pattern within the behaviourally decisive window  $T$  intact. Absent such a counterexample, spike timing remains the only known path from brain to behaviour.

**Scope.** Unless stated otherwise, we hold fixed the body and its sensorimotor interface (afferent and efferent pathways, neuromuscular junctions, peripheral biomechanics): they are intact and time-invariant over the decisive window  $T$  at the task’s declared resolution. Intentional departures (e.g., nerve block, denervation, altered biomechanics) are out of scope for the baseline analyses and, when present, are treated explicitly as changes to the task context  $X$ .

### 3 Causal and Informational Framework

*Causal-sufficiency.* By standard *causal-sufficiency* reasoning (Pearl, 2009, 54–55)<sup>2</sup>, if phenomenal states lay entirely outside the causal pathway to behaviour, systematic covariation between those states and self-report across interventions would be inexplicable. Empirical practice and causal theory therefore converge on the same premise: conscious experience must exert causal influence somewhere along the physical chain that ends in the report.

*Content-matching.* Under ordinary reporting conditions (lucid, earnest, unimpaired), a subject’s utterance  $U$  functions as a lossy encoding of current experience (cf. Dretske, 1981):  $U = f(E, K, X)$ , where  $E$  denotes experienced content,  $K$  a pre-specified reporting policy (granularity, vocabulary), and  $X$  fixed task context. By standard information-theoretic constraints (e.g., the data-processing inequality and rate-distortion trade-offs) (Cover and Thomas, 2006, 34–35, 306–307, 315–316),  $U$  cannot carry more information about  $E$  than the spike-event pattern (SEP; Def. 2) that mediates behaviour in the reporting interval, and matching tightens as  $K$  increases resolution. On an interventionist reading (Woodward, 2010, 289–290); (Waters, 2007), robust, resolution-controlled co-variation of  $U$  with  $E$  (holding  $K, X$  fixed) establishes  $E$  as a partial cause of  $U$  at the task’s specified resolution (Griffiths et al., 2015). A pure common-cause model  $Z \rightarrow \{E, U\}$  that is unaffected by  $K$  predicts that conditioning on  $Z$  (with  $K, X$  fixed) eliminates the  $E$ - $U$  dependence. Empirically, however, with  $K$  and  $X$  fixed, policy-controlled selectivity yields residual  $E$ - $U$  dependence. This pattern contradicts any common-cause explanation  $Z \rightarrow \{E, U\}$  that is unaffected by  $K$  and not downstream of  $E$ , and supports a causal path from  $E$  to  $U$  at the task’s specified resolution (Woodward, 2010, 291); (Griffiths et al., 2015); (Pearl, 2009, 54–55). Consequently, the content of experience is not merely causally relevant but causally *specific*: content-specific differences in  $U$  (e.g., *blue* vs. *red*) are effects of corresponding differences in  $E$  ( $E$  is a difference-making cause of  $U$ ).<sup>3</sup>

---

<sup>2</sup>Page numbers follow the version cited. Where a final journal layout is available we give its pagination; where only an author-manuscript or other open-access PDF is available (e.g., PubMed Central), page numbers refer to that PDF’s internal numbering.

<sup>3</sup>Calibration: pre-commit to  $K$  (e.g., binary vs. multi-slot template), sample  $E$ , then report. Barring

*Background degeneracy.* Non-event background variables—synaptic and dendritic states, neuromodulatory and glial influences, field effects, and related factors—may vary within  $T$  and typically admit many compensatory configurations that induce the *same* mapping from the realised SEP to behaviour. Such variables matter for *generating* the SEP; once the realised SEP inside  $T$  is held within tolerance, residual background variation is screened off from the measured performance.<sup>4</sup> A behaviour or report change with the SEP held within tolerance would therefore constitute the sought spike-silent counterexample.

## 4 Core Deduction

### The Core Deduction.

(PCE) says manifest consciousness pushes on the physical chain we measure as behaviour;

(PNM) says the chain’s only entry point is the temporal spike pattern inside  $T$ .

Therefore, any difference in manifest phenomenal state must alter the SEP inside  $T$ .

Within  $T$ , the SEP suffices for manifest consciousness.  $\square$

**Asymmetric realizability & boundary fixation.** Building on *Background degeneracy* (Section 3), for a given task and reporting resolution, background states form broad equivalence classes that preserve the mapping from the SEP in  $T$  to behaviour, whereas small differences in the SEP can alter behaviour at that resolution. This asymmetry underwrites the fixation claim: background is functionally many-to-one (degenerate) with respect to the readout, while the event pattern is sufficiently specific. Consequently, clamping or replaying the SEP inside  $T$  within the tolerance  $\varepsilon \leq \delta$  suffices to fix the manifest performance; any manifest difference must register as a difference in the SEP.

Drawing on the same interventionist framework as Section 3, Pearl’s causal bridge principle states that if a variable  $X$  is upstream of a behaviour  $B$ , and  $B$  is wholly determined by an intermediate pattern  $Y$ , then any variation in  $X$  must register as a variation in  $Y$  (Pearl, 2009, 54–55).

*Clarification.* The burden of rejoinder now lies with anyone who would violate the fixation claim: any such case would require demonstrating a behaviour-changing manipulation that keeps every spike *within*  $\varepsilon \leq \delta$  of its baseline time.

**Lemma** (Convergence Irrelevance). Whenever two upstream causal histories produce the *same* spike-event pattern (SEP) inside the decisive window  $T$ , every downstream behaviour is identical.

*Proof.* Once the realised SEP is fixed at  $SEP^*$ , there is no remaining causal route by which any upstream variable can influence behaviour— $SEP^*$  already carries the whole effect. By the causal bridge principle, any intervention that hard-clamps  $SEP = SEP^*$  neutralises every upstream variation; by PNM the resulting behaviour is fixed.  $\square$

mechanical error, fatigue, or withdrawal of cooperation, intended and produced content coincide at the chosen resolution; when mismatches occur, they are immediately recognizable and correctable—illustrating reliable first-person control of the  $E \rightarrow U$  mapping under policy  $K$ .

<sup>4</sup>Formally: letting  $\Lambda$  denote the set of non-event background variables and  $U$  the measured behaviour, at fixed reporting policy  $K$  and task context  $X$  we have conditional independence  $U \perp\!\!\!\perp \Lambda \mid SEP_T, K, X$  (interventionist screening-off / d-separation; Woodward, 2010, 294–295; Pearl, 2009, 48–49).

*Manifest-consciousness fixation.* If two upstream histories converge on the same SEP, then—for any behaviour causally downstream of manifest consciousness per PCE (e.g. an earnest report “I am conscious”)—manifest consciousness is present in both histories or in neither.

As established by the cone of manifest invariability (Definition 5), any phenomenal variation compatible with identical behaviour falls within the manifestly inert region; a binary report of the presence or absence of consciousness—and, more generally, any content report that presupposes the presence of consciousness—is too coarse-grained to fit within that region. Such reporting policies afford no degrees of freedom within which presence could vary while the performance remains fixed.

Consequently, the presence or absence of consciousness is fixed by the SEP. To deny this would grant consciousness no causal point of contact and collapse into full epiphenomenalism, violating PCE. Upstream details are interchangeable scaffolding: *many roads can lead to the same pattern, but once the pattern exists, the road is erased—if consciousness cared which road you took, it could never tell you.*

## 5 Timing Tolerance & Robustness

Spike-pattern sufficiency does *not* assert that every picosecond of timing variation carries phenomenal weight. Instead, to make the thesis empirically testable we introduce two tolerances.

**$\delta$ : the physical lower bound.** Thermal noise in ion-channel gating and spontaneous synaptic vesicle release imposes an irreducible jitter on spike timing (Faisal et al., 2008, 294–97), and single-neuron studies in insect sensory systems demonstrate spike-time reproducibility as tight as  $\delta \approx 0.15$  ms (Rokem et al., 2006, 2546–47). To date we are aware of no published preparation showing sharper precision, and the experimental tests proposed here require artificial replicas only to approach—not surpass—this physical limit.

**$\varepsilon$ : the duplication tolerance.** Let  $\varepsilon$  be the maximal pointwise spike-time deviation between a host SEP and a duplicate over the window  $T$ . Two systems are  $\varepsilon$ -SEP-isomorphic when

$$\max_k |t_k^{\text{dup}} - t_k^{\text{host}}| \leq \varepsilon.$$

Spike-pattern sufficiency predicts manifest phenomenal parity (Definition 4) whenever  $\varepsilon \leq \delta$ .

These two tolerances set the benchmark for the empirical programme outlined in Section 6: spike-silent falsification at  $\varepsilon \leq \delta$ , and systematic duplication studies that map how loosely—quantified by  $\varepsilon$ —the spike pattern may be approximated without phenomenological loss.

**Biological slack.** Empirical demonstrations of millisecond-scale robustness—single-neuron jitter, cross-modal latency compensation, etc.—are summarised in Appendix C. They show that the brain normally tolerates noise many times larger than today’s  $\delta$ .

**Engineering slack.** Current neuro-prostheses already restore function with  $\varepsilon \gg \delta$ ; details and citations appear in Appendix D.

**Implication.** Because  $\delta$  reflects a physical noise floor, any duplicate that matches the host SEP within that window is expected to match or exceed the robustness of *uncorrelated* timing noise already present in vivo. The spike-silent falsification test therefore becomes sharper: *demonstrate a behaviourally relevant variable that changes a subject’s phenomenal report while the duplicate holds  $\varepsilon \leq \delta$* . Until such a counterexample appears, timing tolerance leaves spike–pattern sufficiency intact and testable.

*Note on the tolerance benchmark.* The bound  $\delta$  is defined operationally as the smallest spike-timing jitter that has been *empirically demonstrated* to be reproducible in any biological preparation at the time an experiment is run. Setting the falsification bar at  $\varepsilon \leq \delta$  therefore ensures two things.

- (i) **Time-local decisiveness.** A behaviour- or phenomenology-changing result obtained while every spike remains within the *contemporaneous*  $\delta$  would overturn spike-pattern sufficiency relative to the best biological precision then known.
- (ii) **Open-ended stringency.** If future work reveals even tighter neural reproducibility ( $\delta_{\text{new}} < \delta_{\text{old}}$ ), the hypothesis can—and should—be retested at the sharper bound. A previous “falsification” would stand only if it also satisfies the updated  $\varepsilon \leq \delta_{\text{new}}$  criterion; otherwise the result remains inconclusive and invites replication at the improved precision.

In practice, decades of neurophysiology and neuro-prosthetic work already show that complex, report-relevant behaviour tolerates timing noise far larger than current  $\delta$ , so any further downward revision is expected to be marginal. Nevertheless, the criterion remains tied to the empirical record rather than to a fixed numeric value.

## 6 Empirical Programme: falsification track versus verification track

We distinguish two empirical paths.

### 6.1 A. Spike-silent–channel search (primary falsification track)

**Logic.** Any intervention that alters conscious report or behaviour while all *monitored* spikes inside the decisive window  $T$  remain statistically baseline is a *provisional* falsifier. Persistent spike-silence across progressively finer tiers would make any hidden-spike explanation increasingly implausible, thereby mounting an ever-stronger challenge to SPS.

**Monitoring tiers (resolution increases downwards).**

- (a) **Population-level dynamics** — regional firing-rate maps or LFP power (e.g. fMRI, wide-field calcium, high-density ECoG).
- (b) **Individual-neuron timelines** — simultaneous recording of  $\sim 10^4$ – $10^5$  neurons at millisecond precision (multi-shank or surface–laminar probes).
- (c) **Near-SEP coverage** — sub-millisecond tracking of  $\gtrsim 10^6$  spikes in 3-D volumes (future dense CMOS or adaptive two-photon systems).

### Illustrative intervention domains.

- *Astrocytic  $Ca^{2+}$  waves*: suppress global waves; if a sustained-attention shift survives Tier (a) silence, replicate under Tier (b).
- *tACS / ephaptic fields*: deliver purely electric-field patterns; test for perceptual biases while Tier (b) spikes stay baseline.
- *Cytokine IL-6 motivational shift*: reproduce effort-allocation change with nucleus-accumbens spikes unchanged at Tier (b).

All protocols preregister endpoints, consciousness-rating scales, and the tier-appropriate spike-similarity metric. *Provisional falsifier*: a spike-silent effect at Tier (a). *Strong falsifier*: the same effect survives Tier (b) or (c).

**Why this track matters even if SPS holds.** Each tighter tier pushes recording technology forward and reveals mechanistic detail valuable for neuroscience irrespective of the SPS debate.

## 6.2 B. Spike-pattern duplication (verification & refinement track)

This track probes which *specific properties* of the spike-event pattern—population-targeting precision, temporal precision, relative order, synchrony windows, mean rates, cell identity—must be preserved for both behaviour and *subjective experience* to remain intact.

A typical study begins by duplicating native activity at the highest spatial, temporal, and cell-identity fidelity current hardware allows; investigators then vary one variable at a time—whether imposed by engineering limits or by deliberate perturbation—and collect structured phenomenological reports alongside performance metrics. Alternative designs that reach the same goal are equally welcome; for example:

- A 96-channel intracortical visual prosthesis already evokes bright, spatially localised phosphenes without reproducing the retinal code (Fernández et al., 2021, 5–6). Systematically tightening spatial density, timing, or synchrony toward the native pattern is expected to bring percepts ever closer to ordinary vision, but different dimensions may plateau at different thresholds—and some may prove phenomenologically indifferent.
- Analogous protocols can be applied to hippocampal memory implants, mood-regulating DBS, auditory-cortex hearing prostheses, and closed-loop motor prostheses that return tactile or proprioceptive feedback, among other systems, each coupled with vividness, agency, mood-valence, or comparable phenomenology scales.

**Guiding prediction** As any intervention converges on the native SEP along the *relevant* dimensions, SPS predicts that, in general, subjective experience should converge toward indistinguishability from native experience; mapping which dimensions matter, and where their influence saturates, is the core refinement agenda. Successes at progressively coarser values enlarge the empirically demonstrated slack; failures narrow it, yet—unlike spike-silent effects—do not by themselves falsify spike-pattern sufficiency.

As data accumulate across perturbations, the broad sufficiency claim can be refined into a progressively sharper account of which spike-pattern features are *necessary*, which are merely tolerable, and which are irrelevant for consciousness.

## 7 Escape-Hatch Taxonomy

The timing tolerances of Section 5 and the two-track empirical programme of Section 6 leave sceptics only a handful of logical manoeuvres. Table 1 lists every coherent route we can envision and the cost each exacts.

Table 1: Three escape hatches and their empirical and methodological challenges.

Hatch	Core commitment	Empirical and methodological challenges
<i>Dualism</i>	A non-physical agent alters manifest behaviour without producing any corresponding change in the SEP inside $T$ .	Lacks independent predictive power; posits forces with no measurable signature and thus offers no explanatory leverage.
<i>Full epiphenomenalism</i>	Reports of consciousness arise from a zombie cognitive module; qualia float free.	Would demand an astronomically precise, perpetual three-way coincidence among private qualia, self-report behaviour, and any spike-pattern-altering neural intervention of the kind surveyed in Appendix B—an unstable dependency ruled out by standard interventionist criteria (Pearl, 2009, 48–49, 54–55).
<i>Spike-silent channel</i>	Some behaviour-potent variable modulates muscles yet never alters spikes inside $T$ .	After half a century of lesions, optogenetics and prostheses, none has been demonstrated; the spike-silent track in Section 6 is expressly designed to hunt for one. Burden lies on the proponent.

Many well-known positions slot cleanly into this taxonomy, e.g. property- and substance-dualism under *Dualism*; microtubule-only quantum proposals under *Spike-silent channel*.

Unless one elects metaphysical dualism or full epiphenomenalism—both of which fall outside empirical science because they lack explanatory power and falsifiability—the only viable move is to *produce* a spike-silent channel. No such channel has survived empirical scrutiny. Absent one, any rival view that retains both PCE and PNM entails SPS at the manifest level; avoiding SPS therefore requires rejecting at least one of those premises.

## 8 Philosophical & Theoretical Objections

The escape-hatch taxonomy is logical and exhaustive, yet many readers will approach with favoured frameworks in mind. This section situates major frameworks within the taxonomy already addressed and concludes with the “hard-problem” gap, which is not an alternative hypothesis but a demand for further explanation.

## 8.1 Integrated Information Theory (IIT)

**Strong vs. weak IIT.** Current summaries distinguish *strong IIT* from *weak IIT* (Albantakis et al., 2023; Mediano et al., 2022).

**Strong IIT—conditionally compatible.** Strong IIT identifies phenomenal consciousness with a system’s intrinsic cause–effect structure  $\Phi$ , which is fully specified by the underlying physical dynamics. So long as every behaviour-relevant change in  $\Phi$  entails a change in the temporal spike pattern inside the decisive window  $T$ , strong IIT is compatible with spike-pattern sufficiency; otherwise the view collapses into a *spike-silent-channel* proposal subject to the falsification test in Section 6.

**Weak IIT—spike-silent by design.** Weak IIT treats  $\Phi$  as an empirical metric whose value might, in principle, be manipulated independently of spike timing. That independence claim is exactly a spike-silent-channel proposal: engineer two systems whose spikes are  $\varepsilon$ -matched yet whose  $\Phi$  values—and purported phenomenology—diverge. Section 6 specifies the decisive experiment.

## 8.2 Biological chauvinism

Arguments that carbon chemistry possesses unique “feel” reduce either to *Dualism* (non-physical causal powers) or to a *Spike-silent channel* (biochemical property changes behaviour without changing spikes). Neuro-prosthesis data already show that when the imposed spike pattern matches host activity within tens of milliseconds, task-relevant behaviour (e.g. verbal recall, visual localisation) is restored; no biochemical factor has yet been shown to influence behaviour independently of spikes.

## 8.3 Full epiphenomenalism (the zombie scenario)

The zombie scenario concedes functional identity yet posits a parallel realm of causally inert qualia. This is the textbook form of *Full epiphenomenalism*. It contradicts standard interventionist causal-sufficiency norms (Pearl, 2009, 48–49, 54–55), demanding a perpetual, unexplained coincidence among private qualia, self-report, and every spike-pattern–altering intervention; the burden therefore rests on its proponents.

## 8.4 Illusionism (e.g. Frankish, 2016; Dennett, 2017).

Illusionists hold that the *sense* of qualitative properties is itself a functional state generated by the information-processing that also drives report and behaviour. Because this view introduces no extra causal channel beyond the spike-mediated chain already captured by SPS, it is not an “escape hatch” in our taxonomy. Whether illusionism succeeds as a *theory of why* the spike-event pattern gives rise to the seeming of experience lies outside the scope of the present argument; we take no position on that explanatory question here.

## 8.5 Global Workspace Theory (Baars, 1988)

On its mainstream reading, GWT treats conscious access as a global “broadcast” that enables flexible report and control. For our purposes the relevant claim is strictly *manifest*: differences in what is globally available for report and action are differences in *manifest* performance. GWT is compatible with spike-pattern sufficiency (SPS), as broadcast

changes are implemented by processes that alter which neurons (and effectors) spike when within  $T$ .

## 8.6 Higher-Order Theories (Rosenthal, 2006)

HOT holds that a state is conscious when it is the target of an appropriate higher-order representation. In experimental contexts this is a *manifest* thesis: differences in higher-order content that alter what a subject can stably report or do must, by PCE, register in the behaviour stream, and by PNM, do so by altering the SEP inside  $T$ . Accordingly, HOT is compatible with SPS, as *manifest* higher-order variation is implemented by processes that change the SEP in  $T$ .

## 8.7 Predictive Processing (Clark, 2016; Friston, 2010)

Predictive Processing (PP) provides a computational account of how brains minimise prediction error via hierarchical inference and precision control. As such it addresses *why* particular neural event patterns arise in a task and *how* priors, precisions, and dynamics steer them; SPS addresses a distinct question: *what fixes manifest consciousness, given the realised pattern?* On its mainstream reading, PP does not posit *manifest* behaviour that varies independently of the event patterns that reach the effectors; rather, it offers principles that shape those patterns. PP is therefore compatible with SPS, as *manifest* changes are implemented by processes that alter which neurons spike when inside  $T$ .

## 8.8 Quantum Proposals (Hameroff and Penrose, 2013, 45)

Most quantum accounts of consciousness hold that any influence on manifest behaviour is mediated by changes in neural spiking. In Orch-OR, for example, microtubular processes are proposed to influence axonal firing or regulate synaptic plasticity (Hameroff and Penrose, 2013, 45); any behavioural effect is therefore mediated by altered spiking patterns. We take no stance on whether such modulation occurs *in vivo*. What matters for SPS is this: once the SEP inside the decisive window  $T$  is held within tolerance, upstream determinants—including putative microtubular dynamics—are *screened off* from manifest performance (see Section 4). Accordingly, quantum proposals introduce no additional commitment beyond SPS. Quantum dynamics may help determine which SEP obtains, but they are not required once the SEP is realised; the SEP alone suffices for manifest consciousness.

## 8.9 The “Hard Problem” gap (Chalmers, 1995)

Granting that consciousness supervenes on the SEP, one may still ask why this pattern feels as it does. We do not address the hard problem; nevertheless, the thesis constrains its solution space. Any future account of qualitative character that aims to be compatible with SPS would need to locate its basis *within* the spike-defined functional graph—i.e., within features of the SEP inside  $T$ —not beyond it.

In each alternative-hypothesis case the alleged counterexample either (i) invokes causal powers outside physics, (ii) disconnects qualia from behaviour, (iii) is compatible with SPS at the manifest level, or (iv) predicts a spike-silent channel. Sections 5–6 show how

spike-silent-channel manoeuvres can be—or have been—empirically tested; to date, no reproducible spike-silent effect has been demonstrated.

## 9 Burden of Proof & Scientific Parsimony

Science routinely treats the *simplest causally sufficient mechanism* as the default until contrary evidence appears. Once germ theory explained infection patterns, miasma was not considered a co-equal hypothesis; after *Helicobacter pylori* accounted for ulcers, stress-only theories became explanatorily extravagant. The same logic applies here.

Spike-pattern sufficiency already predicts every verified dissociation between neural intervention and conscious report and supplies a spike-silent falsification test (Sections 5–6). Any rival view therefore entails, at minimum,

1. an additional substrate predicate (e.g. carbon chemistry, integrated-information measure, non-physical essence), and
2. an explanation for why nature employs that predicate *despite* the spike pattern being able to do all the causal work on its own.<sup>5</sup>

Each added element multiplies entities without increasing predictive power. By standard parsimony canons—*Occam’s Razor*—the explanatory burden shifts to those who would complicate the ontology (Sober, 2015). Discharging that burden would require producing the empirical counterexample specified in Section 6; absent it, the simpler hypothesis stands.

## 10 Conclusion

Grant two empirically grounded premises—causal efficacy of conscious states and neural mediation of behaviour—and one corollary follows: *manifest consciousness supervenes on the temporal spike-event pattern inside the behaviourally decisive window  $T$* . On this argument, no additional biology, information-integration quotient, or metaphysical spark is required. The claim is *deductive yet falsifiable*: demonstrate a change in manifest performance while every spike in  $T$  remains within  $\varepsilon \leq \delta$  of its baseline time, and the thesis fails.

Section 6 specifies the spike-silent falsification test that could overturn the thesis; Sections 7 and 8 analyse why familiar theoretical manoeuvres do not bypass that test. Absent a spike-silent intervention that meets the decisive test, the burden of proof rests with substrate-specific accounts.

## Acknowledgements

The author gratefully acknowledges constructive editing and critical-feedback assistance from ChatGPT and OpenAI’s o-series language models, as well as additional feedback

---

<sup>5</sup>*Methodological aside.* Our thesis addresses *manifest* consciousness. Claims about non-manifest consciousness lie outside that scope, but a similar parsimony consideration still applies: in a system already sufficient for report and control, positing a further mechanism exclusively for phenomenal states with no effect on observable behaviour adds neither predictive power nor explanatory unification and is therefore unwarranted.

from Claude, Gemini, and Grok. All interpretative choices and any errors remain the responsibility of the human author.

## Appendix A Evidence for PCE

Across multiple clinical and research domains, including clinical anaesthesia (Demertzi et al., 2017, 5–6), analgesia (Lia et al., 2023, 9), affective neuroscience (Faul et al., 2023, 1671), and psychology (Reisenzein and Junge, 2024, 3) an unimpaired subject’s self-report is treated as the *gold-standard (i.e., normative) criterion* for the presence and quality of conscious experience. Every putatively “objective” metric—EEG bispectral index, pain rating scales, perturbational complexity index, sleep-quality indices—is ultimately anchored in lucid self-report: some are defined *directly* by it (e.g., the VAS), whereas the rest were *boot-strapped* by calibrating their thresholds against self-report during development<sup>6</sup>. Moreover, subjective ratings often converge with independent physiological measures. In sleep medicine, Epworth Sleepiness Scale scores rise as mean sleep latency shortens on the Multiple Sleep Latency Test (Johns, 1991, 540); in pain research, trial-to-trial intensity ratings correlate with fMRI signal changes in the primary somatosensory cortex, anterior cingulate cortex, and prefrontal cortex (Coghill et al., 2003, 8538). Such convergent validity reinforces the use of self-report as the normative criterion for conscious experience.

## Appendix B Evidence for PNM

For every class of intervention ever shown to alter overt behaviour, investigators can trace the final causal link to a change in *which neurons fire when*. The catalogue is exhaustive: (1) neuronal electrical events; (2) chemical synapses; (3) connectivity and plasticity; (4) diffuse neuromodulators; (5) glial modulation; (6) endocrine hormones; (7) immune and cytokine signalling; (8) neuro-vascular and metabolic coupling; (9) activity-dependent gene expression and epigenetics; (10) circadian or systemic rhythms; (11) body–brain feedback loops; (12) exotic quantum or biophotonic proposals. In every case the behavioural effect is either (a) delivered *through* altered spike timing within the relevant network or (b) non-existent. Even proponents of microtubule-based quantum hypotheses acknowledge that any such mechanism must ultimately alter synaptic transmission or neuronal firing to influence behaviour (Khan et al., 2024, 9). Decades of lesion studies, single-unit micro-stimulation, optogenetic flips of perceptual choice, closed-loop hippocampal and speech neuro-prostheses, astrocytic calcium-wave knock-outs, cytokine-induced sickness-behaviour models, and endocrine manipulations of mood all conform to this rule.

### Representative intervention studies

1. Microstimulation in visual area MT biases motion-direction reports (Salzman et al., 1992, 2349).

---

<sup>6</sup>Representative calibration papers: BIS and postoperative recall (Myles et al., 2004, 3665-3666); visual and numeric pain rating scales (Hawker et al., 2011, S240-S242); perturbational complexity index validated against TMS-EEG under propofol with patient report (Casarotto et al., 2016, 725); Pittsburgh Sleep Quality Index (Buysse et al., 1989, 193).

2. Phasic optogenetic activation of dorsal-raphé 5-HT neurons suppresses locomotion ( $\sim 50\%$ ) while sparing grooming (Correia et al., 2017, 1).
3. Spike-timing-dependent plasticity (STDP) reweights synapses to support learning and memory (Appelbaum et al., 2023, 113).
4. Phasic optogenetic bursts in locus-coeruleus noradrenergic neurons trigger instant sleep-to-wake transitions in a frequency-dependent manner (Carter et al., 2010, 1).
5. Chemogenetic activation of arcuate-nucleus astrocytes drives feeding  $3\text{--}4 \times$  via increased AgRP neuron spiking (Chen et al., 2016, 1).
6. Optogenetic excitation of hypothalamic PVN CRH neurons boosts grooming and curtails exploration via their glutamatergic spike output (endocrine secretion not required) (Füzesi et al., 2016, 9).
7. Neuron-specific IL-4R $\alpha$  knock-out heightens exploration and reduces fear via elevated cortical network spiking (Hanuscheck et al., 2022, 1).
8. Astrocytic MCT4 knock-down fragments wakefulness by depressing orexin-neuron spiking via lactate shortage (Braga et al., 2024, 1).
9. CREB over-expression biases fear-memory storage to hyper-excitable LA neurons; allatostatin silencing of those neurons reduces freezing to baseline (Zhou et al., 2009, 1).
10. Dopamine-neuron Bmal1 deletion disrupts 24-h firing rhythms and reduces motivated locomotion (Swaroop et al., 2025, 1).
11. Gut-derived  $\delta$ -valerobetaine heightens inhibitory network activity and impairs memory, effects reversed by young-microbiota transplant (Mossad et al., 2021, 1).
12. Microtubule-stabiliser Epothilone-B delays isoflurane-induced loss-of-righting; the authors note that any "putative conscious MT physical state must be able to interact with the membrane potential of neurons, where we know perceptual information is encoded and through which our actions (i.e., muscle activations) must be implemented" (Khan et al., 2024, 1, 9).

## Appendix C Biological Slack

Baseline trial-to-trial jitter spans at least an order of magnitude across mammalian cortical neurons—from sub-millisecond precision in rodent auditory cortex (DeWeese et al., 2003, 7946) to multi-millisecond jitter in primate visual cortex (Bair and Koch, 1996, 1196–1197)—yet the behaviours these circuits guide remain stable (see also DeWeese et al., 2005, 481).

## Appendix D Engineering Slack

A human hippocampal neuro-prosthesis has already duplicated ensemble activity binned at 50 ms (Hampson et al., 2018, 7–8)—far coarser than  $\delta$  yet still sufficient to produce

significant improvements in both short-term and long-term memory performance (Hampson et al., 2018, 1). State-of-the-art speech BCIs generate intelligible sentences at 78 words  $\text{min}^{-1}$  while updating neural features every  $\sim 5$  ms, placing an upper bound of  $\varepsilon \lesssim 5$  ms on the spike precision now achievable in vivo (Metzger et al., 2023, 5,20). Closed-loop optogenetic feedback delivered to dorsal CA1 at either the theta peak or trough (for encoding and retrieval, respectively), with a mean sensor-to-stim delay of  $\approx 21.5$  ms (jitter  $\approx 7.3$  ms), enhanced spatial-memory performance despite this timing slack (Siegle and Wilson, 2014, 1, 7) Proof-of-principle hardware already approaches the  $\delta$  scale: culture-based CMOS–MEA systems achieve 0.4 ms closed-loop delays with  $< 50 \mu\text{s}$  jitter (Müller et al., 2013, 2); a wearable human neuromodulator reports round-trip sense-to-stim latencies of  $1.57 \pm 0.19$  ms (Topalovic et al., 2023, 517, 523); These figures set a practical lower bound on  $\varepsilon$  today and, together with the rapid pace of neuromodulation engineering, suggest a plausible path toward sub-millisecond duplication accuracy. Importantly, current devices confirm that coarse SEP duplication can reinstate behaviourally measured memory performance; however, existing studies have not included systematic phenomenological ratings, so qualitative experience at the  $\varepsilon \leq \delta$  threshold remains an open target for the duplication–refinement track in Section 6.2.

## References

- Larissa Albantakis, Leonardo Barbosa, Graham Findlay, Matteo Grasso, Andrew M. Haun, William Marshall, William G. P. Mayner, Alireza Zaeemzadeh, Melanie Boly, Bjørn E. Juel, Shuntaro Sasai, Keiko Fujii, Isaac David, Jeremiah Hendren, Jonathan P. Lang, and Giulio Tononi. Integrated information theory (iit) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLoS Computational Biology*, 19(10):1–45, 10 2023.
- Lawrence G. Appelbaum, Mohammad Ali Shenasa, Louise Stolz, and Zafiris Daskalakis. Synaptic plasticity and mental health: methods, challenges and opportunities. *Neuropsychopharmacology*, 48(1):113–120, 2023. ISSN 1740-634X. doi: 10.1038/s41386-022-01370-w. URL <https://doi.org/10.1038/s41386-022-01370-w>.
- Bernard J. Baars. *A Cognitive Theory of Consciousness*. Cambridge University Press, New York, 1988.
- Wyeth Bair and Christof Koch. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Computation*, 8:1185–1202, 1996. ISSN 0899-7667. doi: 10.1162/neco.1996.8.6.1185.
- Alice Braga, Martina Chiacchiaretta, Luc Pellerin, Dong Kong, and Philip G. Haydon. Astrocytic metabolic control of orexinergic activity in the lateral hypothalamus regulates sleep and wake architecture. *Nature Communications*, 15(1):5979, 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-50166-7. URL <https://doi.org/10.1038/s41467-024-50166-7>.
- D. J. Buysse, C. F. 3rd Reynolds, T. H. Monk, S. R. Berman, and D. J. Kupfer. The pittsburgh sleep quality index: a new instrument for psychiatric practice and research. *Psychiatry research*, 28:193–213, May 1989.

- Matthew E. Carter, Ofer Yizhar, Sachiko Chikahisa, Hieu Nguyen, Antoine Adamantidis, Seiji Nishino, Karl Deisseroth, and Luis de Lecea. Tuning arousal with optogenetic modulation of locus coeruleus neurons. *Nature neuroscience*, 13:1526–33, Dec 2010.
- Silvia Casarotto, Angela Comanducci, Mario Rosanova, Simone Sarasso, Matteo Fecchio, Martino Napolitani, Andrea Pigorini, Adenauer G Casali, Pietro D. Trimarchi, Melanie Boly, Olivia Gosseries, Olivier Bodart, Francesco Curto, Cristina Landi, Maurizio Mariotti, Guya Devalle, Steven Laureys, Giulio Tononi, and Marcello Massimini. Stratification of unresponsive patients by an independently validated index of brain complexity. *Annals of neurology*, 80:718–729, Nov 2016.
- David Chalmers. Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3):200–19, 1995.
- Naiyan Chen, Hiroki Sugihara, Jinah Kim, Zhanyan Fu, Boaz Barak, Mriganka Sur, Guoping Feng, and Weiping Han. Direct modulation of gfap-expressing glia in the arcuate nucleus bi-directionally regulates feeding. *eLife*, 5:e18716, oct 2016. ISSN 2050-084X. doi: 10.7554/eLife.18716. URL <https://doi.org/10.7554/eLife.18716>.
- Andy Clark. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, 01 2016. ISBN 9780190217013. doi: 10.1093/acprof:oso/9780190217013.001.0001.
- Robert C. Coghill, John G. McHaffie, and Yi-Fen Yen. Neural correlates of interindividual differences in the subjective experience of pain. *Proceedings of the National Academy of Sciences of the United States of America*, 100:8538–42, Jul 2003.
- Patrícia A. Correia, Eran Lottem, Dhruva Banerjee, Ana S. Machado, Megan R. Carey, and Zachary F. Mainen. Transient inhibition and long-term facilitation of locomotion by phasic optogenetic activation of serotonin neurons. *eLife*, 6, Feb 2017.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory 2nd Edition (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, July 2006. ISBN 0471241954.
- Athena Demertzi, Jacobo Diego Sitt, Simone Sarasso, and Wim Pinxten. Measuring states of pathological (un)consciousness: research dimensions, clinical applications, and ethics. *Neuroscience of consciousness*, 2017:nix010, 2017.
- Daniel Clement Dennett. *From bacteria to Bach and back : the evolution of minds*. W.W. Norton, 2017. URL <https://cir.nii.ac.jp/crid/1971149384748919330>.
- Michael R. DeWeese, Michael Wehr, and Anthony M. Zador. Binary spiking in auditory cortex. *Journal of Neuroscience*, 23(21):7940–7949, 2003. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.23-21-07940.2003. URL <https://www.jneurosci.org/content/23/21/7940>.
- Michael R. DeWeese, Tomáš Hromádka, and Anthony M. Zador. Reliability and representational bandwidth in the auditory cortex. *Neuron*, 48(3):479–488, November 2005. ISSN 0896-6273. doi: 10.1016/j.neuron.2005.10.016. URL <https://doi.org/10.1016/j.neuron.2005.10.016>.

- Fred I. Dretske. *Knowledge and the Flow of Information*. MIT Press, Stanford, CA, 1981.
- A. Aldo Faisal, Luc P. J. Selen, and Daniel M. Wolpert. Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303, April 2008. ISSN 1471-0048. doi: 10.1038/nrn2258.
- Leonard Faul, Mary G Baumann, and Kevin S LaBar. The representation of emotional experience from imagined scenarios. *Emotion*, 23(6):1670–1686, September 2023.
- Eduardo Fernández, Arantxa Alfaro, Cristina Soto-Sánchez, Pablo Gonzalez-Lopez, Antonio M. Lozano, Sebastian Peña, Maria Dolores Grima, Alfonso Rodil, Bernardeta Gómez, Xing Chen, Pieter R. Roelfsema, John D. Rolston, Tyler S. Davis, and Richard A. Normann. Visual percepts evoked with an intracortical 96-channel microelectrode array inserted in human occipital cortex. *The Journal of clinical investigation*, 131, Dec 2021.
- Keith Frankish. Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11-12):11–39, 2016.
- Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010. ISSN 1471-0048. doi: 10.1038/nrn2787. URL <https://doi.org/10.1038/nrn2787>.
- Tamás Füzési, Nuria Daviu, Jaclyn I. Wamsteeker Cusulin, Robert P. Bonin, and Jaideep S. Bains. Hypothalamic crh neurons orchestrate complex behaviours after stress. *Nature communications*, 7:11937, Jun 2016.
- Paul E. Griffiths, Arnaud Pocheville, Brett Calcott, Karola Stotz, Hyunju Kim, and Rob Knight. Measuring causal specificity. *Philosophy of Science*, 82(4):529–555, 2015. doi: 10.1086/682914.
- Stuart Hameroff and Roger Penrose. Consciousness in the universe: a review of the ‘orch or’ theory. *Physics of life reviews*, 11, 08 2013. doi: 10.1016/j.plrev.2013.08.002.
- Robert E. Hampson, Dong Song, Brian S. Robinson, Dustin Fetterhoff, Alexander S. Dakos, Brent M. Roeder, Xiwei She, Robert T. Wicks, Mark R. Witcher, Daniel E. Couture, Adrian W. Laxton, Heidi Munger-Clary, Gautam Popli, Myriam J. Sollman, Christopher T. Whitlow, Vasilis Z. Marmarelis, Theodore W. Berger, and Sam A. Deadwyler. Developing a hippocampal neural prosthetic to facilitate human memory encoding and recall. *Journal of neural engineering*, 15:036014, Jun 2018.
- Nicholas Hanuschek, Carine Thalman, Micaela Domingues, Samantha Schmaul, Muthuraman Muthuraman, Florian Hetsch, Manuela Ecker, Heiko Endle, Mohammadsaleh Oshaghi, Gianvito Martino, Tanja Kuhlmann, Katarzyna Bozek, Tim van Beers, Stefan Bittner, Jakob von Engelhardt, Johannes Vogt, Christina Francisca Vogelaar, and Frauke Zipp. Interleukin-4 receptor signaling modulates neuronal network activity. *The Journal of experimental medicine*, 219, Jun 2022.
- Gillian A. Hawker, Samra Mian, Tetyana Kendzerska, and Melissa French. Measures of adult pain: Visual analog scale for pain (vas pain), numeric rating scale for pain (nrs pain), mcgill pain questionnaire (mpq), short-form mcgill pain questionnaire (sf-mpq), chronic pain grade scale (cpgs), short form-36 bodily pain scale (sf-36 bps), and measure of intermittent and constant osteoarthritis pain (icoap). *Arthritis care & research*, 63 Suppl 11:S240–52, Nov 2011.

- M. W. Johns. A new method for measuring daytime sleepiness: the epworth sleepiness scale. *Sleep*, 14:540–5, Dec 1991.
- Sana Khan, Yixiang Huang, Derin Timuçin, Shantelle Bailey, Sophia Lee, Jessica Lopes, Emeline Gaunce, Jasmine Mosberger, Michelle Zhan, Bothina Abdelrahman, Xiran Zeng, and Michael C. Wiest. Microtubule-stabilizer epothilone b delays anesthetic-induced unconsciousness in rats. *eNeuro*, 11, Aug 2024.
- Elena Lia, Valentina Pucci, Cecilia Raccagna, Stefano Sebastiani, and Boaz Gedali-ahu Samolsky Dekel. Analgosedation management in the intensive care unit: A narrative systematic review. *The Open Anesthesia Journal*, 17(1), November 2023. ISSN 2589-6458. doi: 10.2174/0125896458275320231120071303.
- Pedro A. M. Mediano, Fernando E. Rosas, Daniel Bor, Anil K. Seth, and Adam B. Barrett. The strength of weak integrated information theory. *Trends in Cognitive Sciences*, 26(8):646–655, August 2022. ISSN 1364-6613. doi: 10.1016/j.tics.2022.04.008. URL <https://doi.org/10.1016/j.tics.2022.04.008>.
- Sean L. Metzger, Kaylo T. Littlejohn, Alexander B. Silva, David A. Moses, Margaret P. Seaton, Ran Wang, Maximilian E. Dougherty, Jessie R. Liu, Peter Wu, Michael A. Berger, Inga Zhuravleva, Adelyn Tu-Chan, Karunesh Ganguly, Gopala K. Anumanchipalli, and Edward F. Chang. A high-performance neuroprosthesis for speech decoding and avatar control. *Nature*, 620(7976):1037–1046, 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06443-4. URL <https://doi.org/10.1038/s41586-023-06443-4>.
- Omar Mossad, Elisa Nent, Sabrina Woltemate, Shani Folschweiller, Joerg M. Buescher, Daniel Schnepf, Daniel Erny, Peter Staeheli, Marlene Bartos, Antal Szalay, Bärbel Stecher, Marius Vital, Jonas F. Sauer, Tim Lämmermann, Marco Prinz, and Thomas Blank. Microbiota-dependent increase in  $\delta$ -valerobetaine alters neuronal function and is responsible for age-related cognitive decline. *Nature aging*, 1:1127–1136, Dec 2021.
- P. S. Myles, K. Leslie, J. McNeil, A. Forbes, and M. T. V. Chan. Bispectral index monitoring to prevent awareness during anaesthesia: the b-aware randomised controlled trial. *Lancet (London, England)*, 363:1757–63, May 2004.
- Jan Müller, Douglas J. Bakkum, and Andreas Hierlemann. Sub-millisecond closed-loop feedback stimulation between arbitrary sets of individual neurons. *Frontiers in Neural Circuits*, Volume 6 - 2012, 2013. ISSN 1662-5110. doi: 10.3389/fncir.2012.00121. URL <https://www.frontiersin.org/journals/neural-circuits/articles/10.3389/fncir.2012.00121>.
- Judea Pearl. *Causality*. Cambridge University Press, Cambridge, 2 edition, 2009.
- Rainer Reisenzein and Martin Junge. Measuring the intensity of emotions. *Frontiers in Psychology*, 15, September 2024. ISSN 1664-1078. doi: 10.3389/fpsyg.2024.1437843.
- Ariel Rokem, Sebastian Watzl, Tim Gollisch, Martin Stemmler, Andreas V. M. Herz, and Inés Samengo. Spike-timing precision underlies the coding efficiency of auditory receptor neurons. *Journal of Neurophysiology*, 95(4):2541–2552, 2006. doi: 10.1152/jn.00891.2005.
- David Rosenthal. *Consciousness and Mind*. Oxford University Press UK, New York, 2006.

- C. D. Salzman, C. M. Murasugi, K. H. Britten, and W. T. Newsome. Microstimulation in visual area mt: effects on direction discrimination performance. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 12:2331–55, Jun 1992.
- Charles Scott Sherrington. *The Integrative Action of the Nervous System*. Charles Scribner’s Sons, New York, 1906.
- Joshua H Siegle and Matthew A Wilson. Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *eLife*, 3:e03061, jul 2014. ISSN 2050-084X. doi: 10.7554/eLife.03061. URL <https://doi.org/10.7554/eLife.03061>.
- Elliott Sober. *Ockham’s Razors: A User’s Manual*. Cambridge University Press, Cambridge, 2015.
- Ananya Swaroop, Jodi R. Paul, Laura J. McMeekin, Ashley M. Barnett, Alana M. Colafrancesco, Drèson L. Russell, Camille M. Smith, Micah S. Simmons, Laura A. Volpicelli-Daley, Rita M. Cowell, and Karen L. Gamble. The molecular clock drives motivated locomotion and time-of-day-dependent firing patterns in mouse dopaminergic neurons. *Npj biological timing and sleep*, 2:28, 2025.
- Uros Topalovic, Sam Barclay, Chenkai Ling, Ahmed Alzuhair, Wenhao Yu, Vahagn Hokyhkyan, Hariprasad Chandrakumar, Dejan Rozgic, Wenlong Jiang, Sina Basir-Kazeruni, Sabrina L. Maoz, Cory S. Inman, Matthias Stangl, Jay Gill, Ausaf Bari, Aria Fallah, Dawn Eliashiv, Nader Pouratian, Itzhak Fried, Nanthia Suthana, and Dejan Markovic. A wearable platform for closed-loop stimulation and recording of single-neuron and local field potential activity in freely moving humans. *Nature Neuroscience*, 26(3):517–527, 2023. ISSN 1546-1726. doi: 10.1038/s41593-023-01260-4. URL <https://doi.org/10.1038/s41593-023-01260-4>.
- C. Waters. “causes that make a difference.”. *The Journal of Philosophy*, 105, 11 2007. doi: 10.2307/20620058.
- James Woodward. Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology and Philosophy*, 25(3):287–318, 2010. doi: 10.1007/s10539-010-9200-z.
- Yu Zhou, Jaejoon Won, Mikael Guzman Karlsson, Miou Zhou, Thomas Rogerson, Jayaprakash Balaji, Rachael Neve, Panayiota Poirazi, and Alcino J. Silva. Creb regulates excitability and the allocation of memory to subsets of neurons in the amygdala. *Nature neuroscience*, 12:1438–43, Nov 2009.