SHORING LABS · FINAL CORRECTIVE PAPER IN THE VECTORS SEQUENCE

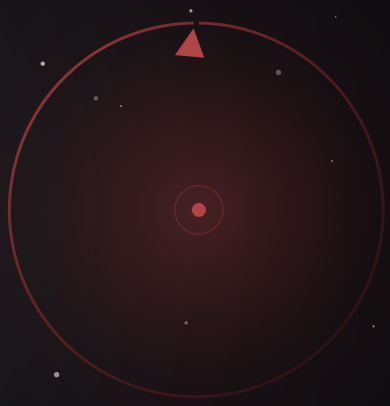

The Observer *Defect*

Neurological Latency, Predictive Hallucination, and the Stateless Self

AUTHOR	VK Richter <vk@shoringlabs.com>
SUBJECT	Cognitive Architecture & the Illusion of Sovereign Agency
RISK DOMAIN	Cognitive Integrity / Identity Continuity
MITIGATION	None. The defect is the system.

External deceiver invoked	0
Supernatural agency admitted	0
Findings traced to source	100%

Control & Methodology

TITLE	The Observer Defect
SUBTITLE	Neurological Latency, Predictive Hallucination, and the Stateless Self
SUBJECT	Cognitive Architecture and the Illusion of Sovereign Agency
RISK DOMAIN	Cognitive Integrity / Identity Continuity
STATUS	Final Corrective Paper in the Vectors Sequence
MITIGATION STATUS	None. The defect is the system.
AUTHOR	VK Richter · vk@shoringlabs.com

0 · Methodology: Permitted Sources of Agency Failure

Scope. This document draws on neuroscience, perception and motor-control research, memory reconsolidation literature, and philosophy of mind. Philosophical material is admitted only where explicitly flagged, and only where it constitutes a named, defended framework rather than speculation.

Excluded. No soul. No possession. No simulation hypothesis. No external deceiver, demon, or supernatural manipulator. These are not rejected on ideological grounds. They are rejected because they are unnecessary. The mind is sufficiently unstable under ordinary description that no additional agent is required to produce the result.

Evidentiary constraint (**load-bearing**). Claims in this paper are sorted into exactly two registers, and the two are never permitted to blur.

☐ **Finding** — measured, replicated, audited fact ☐ **Model** — leading framework, not settled

Every flat-voice claim traces to a named source in Appendix A. A model claim written in the declarative voice is the worst failure this document can commit, and the reader is invited to treat any such lapse as a defect in the paper rather than a property of the world.

> *The outside was not enclosed. The inside was not solid. The observer is not sovereign. This paper documents the operational mechanics of that non-sovereignty.*

I ABSTRACT

The Systemic Latency

The conscious self is a model generated inside the system it claims to govern. This is the operating premise of the present paper, and it is not advanced as a provocation. It is the most economical reading of a set of findings that are individually mundane and collectively load-bearing.

This document isolates the cognitive fictions by which the organism misclassifies its own operation: delayed synthesis reported as immediacy, confabulation reported as authorship, reconsolidation reported as record, and recurrence reported as identity. Each fiction is functional. Each is adaptive. None requires the organism to be aware of it, and the organism is, as a rule, not aware of it.

The architecture of human cognition relies on a series of compensatory edits. Neural transmission is physically bounded. Perceptual continuity is manufactured through suppression and postdictive smoothing. Narrative coherence is maintained by a subsystem that fabricates causal explanations without access to the underlying mechanisms. Memory is not archived; it is dynamically rewritten upon each retrieval. Under the leading interpretive frameworks, perception itself is held to function as a top-down prediction engine that treats sensory input as error-correction rather than primary data — a framework, not a settled fact, and marked as such in its section. The resulting construct — a unified, continuous, self-directed observer — is an interface product. It is operationally useful. It is not structurally accurate.

> *The observer is not disqualified by mystery. The observer is disqualified by ordinary operation.*

There is no hidden fact, no concealed deceiver, no occult mechanism whose discovery would undo the self — only latencies measured in milliseconds, clinical dissociations reproducible in any well-equipped ward, and experimental manipulations that succeed on naïve subjects with dispiriting reliability.

The terminal claim of this paper is the one that closes the trilogy. The defect cannot be repaired without removing the observer, because the observer is what is doing the reading. No corrective position exists outside cognition from which the reports of cognition might be checked. The instrument and the object under examination are the same instrument.

II PRELIMINARY NOTE

Cognitive Fictions

The audit that follows proceeds through six misclassifications, ordered by escalation. Each corrects the observer's assumed relationship to a single faculty: experience, unity, intention, memory, perception, and finally the observer itself. The first four fictions are FINDING. The final two are MODEL. The boundary is drawn before the argument begins, so that the reader may hold the document to it.

EXHIBIT A — COGNITIVE FICTIONS REGISTER

FICTION	MISCLASSIFICATION	MECHANISM	REGISTER
Experience	Present-tense access	Latency-bound synthesis	FINDING
Unity	Centralized architecture	Distributed processing + post-hoc narration	FINDING
Intention	Transparent cause	Reasons generated after the fact	FINDING
Memory	Archival storage	Reconsolidated mutable record	FINDING
Perception	Passive observation	Prediction constrained by error	MODEL
Self	Persistent occupant	Recurring reconstruction; no observer module	MODEL

The sequence is cumulative and the order is not arbitrary. Sections 01 through 04 establish that the observer's reports are late, decentralized, confabulated, and revisable — facts that stand without the models that follow. Sections 05 and 06 then advance the leading interpretive frameworks that, if correct, would account for why a system built this way would nonetheless report itself as present, unified, authoritative, and continuous. The models are offered as the most credible available reading of the findings, and as nothing more.

The Latency Stack

REGISTER FINDING

> CORE FINDING
>
> Humans mistake delayed synthesis for present-tense access.

The nervous system is a physical medium, and signals within it travel at finite speed. Peripheral nerve conduction velocities fall in the range of roughly 30 to 70 metres per second for large myelinated fibres, and below 2 metres per second for unmyelinated ones. Each synaptic crossing imposes a further delay on the order of milliseconds. These are not pathological figures; they are the baseline parameters of an intact system. By the time any sensory event is available for report, the event itself is in the past. Awareness is not slowed by inattention or fatigue. Its lag is precluded from elimination by physics.

The system does not merely lag. It edits. During each saccade — the rapid ballistic movements by which the eyes reposition several times per second — visual sensitivity is actively suppressed. This saccadic suppression, characterized by Volkmann and later by Ross, Morrone and Burr, attenuates visual processing across the duration of the movement, on the order of tens of milliseconds per saccade. The smeared transit of the image across the retina is suppressed before delivery.

> *The observer is never shown the cut footage. The observer is shown the cut.*

Cross-modal binding compounds the matter. Stimuli arriving through different channels — a sound, a flash, a touch — are integrated within temporal windows on the order of tens to roughly a hundred milliseconds, and are frequently reported as simultaneous when they fall inside that window, regardless of the true order of their arrival. Perception across this interval is not a real-time feed. It is computed over a span and presented without the span attached.

The most direct demonstration that the report is constructed after the fact is postdiction. Eagleman and Sejnowski documented cases in which a stimulus arriving *after* a target event alters the perceived timing or content of the *earlier* event. The flash-lag effect and related phenomena are consistent with a system that defers its account of what happened at a given instant until information from a later instant has arrived. The brain does not timestamp inputs as they arrive and present them in order. It gathers a sufficient cluster of signals, resolves ambiguities, and then assigns the perceptual label. The verdict on a moment can be withheld until evidence from its near future is in.

Assembled, these mechanisms describe a pipeline. The present is not perceived. It is assembled after the fact and labeled current. The observer receives a stabilized delay and calls it now — and the labeling is so reliable, so seamlessly integrated, that the lag it conceals is itself never an object of experience. The latency is operationally irrelevant at human scale, but it is structurally absolute. The smoothing is the product.

EXHIBIT B — THE LATENCY CHAIN

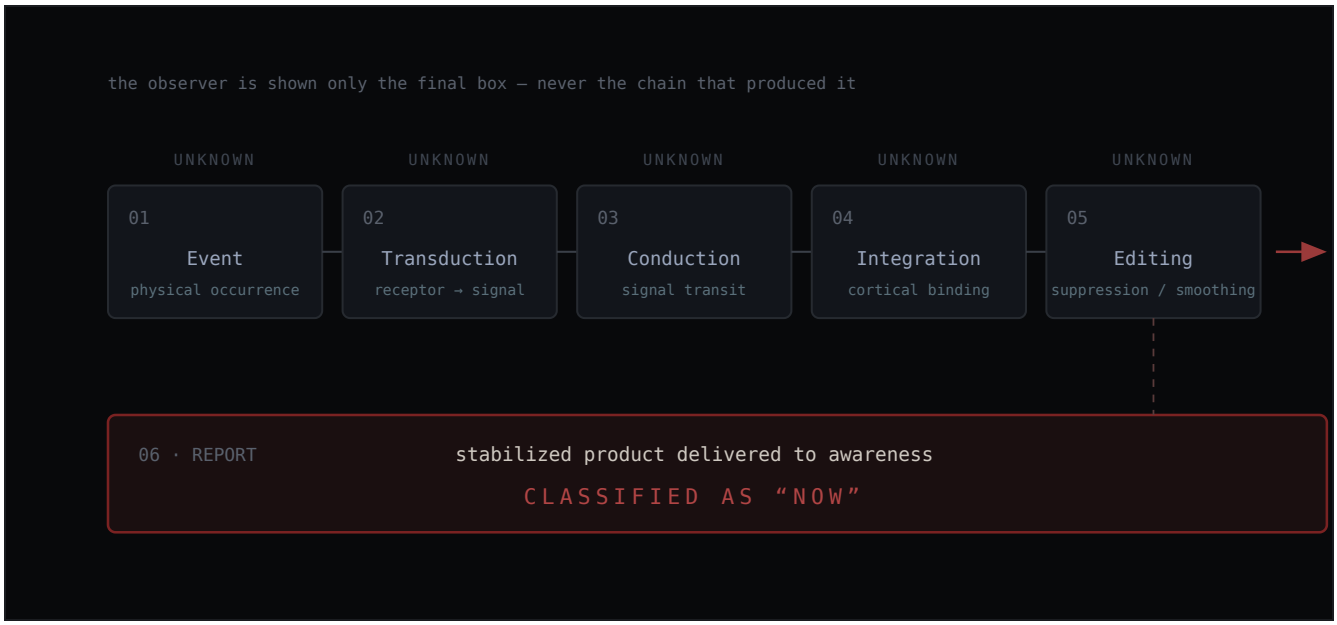

Five processing phases — event, transduction, conduction, integration, editing — each opaque to the observer. Only the sixth, the stabilized product, is delivered to awareness, where it is classified as “now.” The observer is shown the final box, never the chain that produced it.

The Divided Executive

REGISTER FINDING

> CORE FINDING
>
> Humans mistake narrative unity for centralized control.

The assumption that a single agent occupies and directs the system survives only until the system is divided. In patients whose corpus callosum has been surgically severed to control intractable epilepsy, Sperry and Gazzaniga demonstrated that the two hemispheres can be addressed independently, each with distinct access to sensory and motor channels. When an instruction is presented to the right hemisphere alone — directing the subject, say, to stand or to walk — the subject complies. Asked why, the left hemisphere, which controls speech and which had no access to the instruction, produces an explanation: it wanted a drink, it felt restless. The explanation is fluent, confident, and unsupported by the actual cause. The speaking hemisphere does not report that it does not know. It generates a reason, and it does not register the reason as invented.

Gazzaniga named the responsible faculty the interpreter. Its function is the maintenance of a coherent first-person account across whatever the rest of the system does. It is not a liar in any moral sense; it has no access to the truth it would need in order to lie. It is a narrative-generating process operating on incomplete data and delivering a finished story regardless. Unity, in the split-brain subject, is preserved not by a unified controller but by post-hoc narration.

The dissociation runs the other way as well. In blindsight, documented by Weiskrantz in patients with damage to primary visual cortex, the subject reports seeing nothing in the affected field, yet can guide behavior — reaching, pointing, orienting, discriminating motion — using visual information arriving through that field at rates well above chance. Information demonstrably reaches and steers the system without entering report. Observation does not require report; report does not require access. The two come apart cleanly under lesion, and where they come apart, the conscious account is revealed as one output among several rather than the master record.

Anosognosia presents the inverse failure and closes the demonstration. Described by Babinski and abundant in the clinical literature, it is the condition in which a patient with manifest paralysis — typically following right-hemisphere stroke — sincerely denies any impairment, offering to perform movements that are impossible, insisting the limb is merely tired. The failure of function is real and observable. The system's failure to report its own failure is equally real. The monitor that should flag the deficit is itself part of the damaged tissue, and a system cannot reliably audit a fault in its own auditor.

> *Unity is a policy, not a property — a default output maintained by the interpreter because coherence is useful, and defended even when the underlying architecture has been cut in two.*

There is no single room in which the self sits. Cognitive functions are distributed across parallel, partially isolated processing streams. The report is not a transcript of the architecture. It is a summary issued by a subsystem that lacks administrative privileges over the whole.

Confabulated Agency

REGISTER FINDING

> CORE FINDING
>
> Humans mistake the conscious report of a reason for its cause.

The interpreter of Section 02 is not a curiosity of the surgical ward. It operates continuously in the intact subject, and its confabulations can be elicited under controlled conditions in the general population.

Johansson and colleagues demonstrated this in their choice-blindness paradigm. Subjects shown pairs of faces and asked to select the more attractive were then, by sleight of hand, handed the photograph they had *rejected* and asked to explain their choice. A substantial proportion failed to detect the switch — and proceeded to justify the selection they had never made, citing specific features of the wrong face with introspective confidence. The justification was generated to fit the outcome presented, after the fact, with no apparent awareness that the outcome was not the subject's own. The causal chain was transparent to the experimenter. It was not transparent to the subject.

This is the experimental form of a result established decades earlier. Nisbett and Wilson, in their 1977 review, assembled evidence that people routinely offer confident causal accounts of their own behavior that do not match the causes the experimenters had in fact manipulated. They did not guess. They did not lie. Their conclusion was precise: introspective reports about higher-order cognitive processes are frequently not based on genuine access to those processes, but on a posteriori theorizing about what the cause plausibly was.

> *The self receives the press release and mistakes it for the command log. Sincerity is not evidence of accuracy. It is evidence of the interpreter's success.*

A note on the most cited and most abused datum in this domain is required, and it must be a hedged one. Libet's experiments recorded a readiness potential — a build-up of cortical activity — beginning measurably *before* subjects reported the conscious intention to move. Libet interpreted this as the brain initiating action prior to conscious decision. This interpretation is disputed and remains unresolved. Schurger, Sitt and Dehaene proposed a stochastic-accumulator account, under which the readiness potential reflects ongoing neural fluctuation drifting toward a threshold rather than a discrete pre-conscious decision. *The brain decides before you do* is an overclaim. The dispute is live. It does not settle the question of free will, and this paper does not rest on it.

What the choice-blindness and post-hoc-misattribution findings establish, independent of Libet, is narrower and firmer: the conscious report of a reason is not reliably the reason. The organism acts. The interpreter generates a reason. The reason is experienced as the cause. The defect is not that reasons are false. The defect is that they are experienced as origins.

The Mutable Archive

REGISTER FINDING

> CORE FINDING
>
> Humans mistake memory for a stored record.

The continuity of the self across time depends on the assumption that memory preserves the past. The assumption is false in a specific, demonstrated way: memory is not retrieved intact, it is reconstructed, and the act of reconstruction can modify what is reconstructed.

The behavioral demonstration is the misinformation effect, established by Loftus and colleagues. Subjects shown a filmed event and subsequently exposed to misleading post-event information incorporated that information into their recollection of the original event. In the classic study with Palmer, the verb used to describe a filmed collision — "smashed" versus "hit" — altered subjects' later estimates of the speed of the vehicles and raised the rate at which they reported seeing broken glass that was never present. The original record did not reliably resist the later input. It absorbed it, and the subject experienced the altered version as the original. Entirely false memories of events that never occurred have since been induced in a proportion of subjects by related procedures.

The physiological mechanism was identified by Nader, Schafe and LeDoux in studies of memory reconsolidation. A consolidated memory, when retrieved, can return to a labile state and require restabilization through protein synthesis before it is stored again. During this window it is vulnerable to disruption, strengthening, or alteration. Restabilization following retrieval is, on the reconsolidation account, a re-writing rather than a replaying. A memory accessed often is a memory repeatedly subject to that process.

Compounding this is the well-documented dissociation between confidence and accuracy. The strength of the conviction that a memory is true is, under many conditions, a poor predictor of whether it is. High-confidence false memories are documented; verified memories do not reliably carry the highest confidence. Confidence behaves as metadata, not verification. It tracks the fluency of retrieval and the coherence of the narrative more reliably than the fidelity of the trace.

> *Remembering is not replaying a recording. It is drafting a document from fragmentary notes under current conditions.*

The past is not stored intact. It is regenerated, edited, and refiled under the name of memory. The conditions at the moment of recall — present mood, present belief, present suggestion — can be written into the draft. The archive the observer consults to confirm its own continuity is one the observer may revise each time it is opened. The seamlessness of recall conceals the revision.

The Predictive Interface

REGISTER MODEL

> CORE CLAIM — MODEL REGISTER
>
> Humans mistake prediction for observation.

The preceding four sections are findings. The two that follow are not. They are leading interpretive frameworks by which the findings might be explained, and they are marked as such throughout. **Nothing in this section is offered as audited fact.**

On the predictive-coding account, developed by Rao and Ballard and elaborated by Clark and within Friston's free-energy framework, perception is held not to be a feed-forward delivery of sensory data to a central viewer. Under this framework, the brain is described as maintaining a generative model of the causes of its sensory input and continuously issuing top-down predictions about what that input should be. Sensory signals arriving from the periphery are held, on this account, not to constitute perception directly but to carry prediction error — the discrepancy between what the model expected and what arrived. What propagates upward, under the framework, is principally the error; what constitutes the experienced percept is principally the prediction, revised as needed to minimize that error.

> *On the leading account, the brain does not display reality; it negotiates with error.*

This framework, if it holds, would account for the latency and editing phenomena of Section 01. The temporal integration window would be construed not merely as a bottleneck but as a prediction horizon. Saccadic suppression would be construed as the suppression of error signals during periods where the model predicts visual discontinuity. The richness and stability of the perceived world would be read not as a measure of the bandwidth of incoming data, which is comparatively sparse and noisy, but as a measure of the elaboration of the internal model. Input, on this reading, does not compose the model. It disciplines it.

It is in this context, and only in this context, that Seth's term must be understood. Seth describes perception, under the predictive framework, as a "controlled hallucination" — a phrase that is the framework's own slogan, used here in quotation and attributed, and not a phenomenological report that reality has been shown to be unreal. The emphasis falls on *controlled*. The model, on this account, is constrained continuously and stringently by sensory error. The claim, under this account, is not that perception is arbitrary. It is that perception is generated rather than received, and corrected rather than read.

The framework remains unsettled. Its predictive-coding formulations are supported by a substantial and growing body of evidence; its strongest free-energy generalizations are contested in the literature, and competing accounts of cortical function exist. The reader is asked to weight it accordingly. What it offers, if it holds, is a mechanism by which a system built on the latencies of Section 01 and the divided architecture of

Section 02 would nonetheless produce a seamless, present-tense, unified world — by predicting that world and presenting the prediction, not the data, to whatever consumes the output.

The Stateless Self

REGISTER MODEL — FUNCTIONAL

> CORE CLAIM — MODEL REGISTER
>
> Humans mistake recurrent reconstruction for a persistent occupant.

This final section is also a model, and it carries an additional constraint. Its claim is functional, not metaphysical. It does not assert that the self does not exist. It asserts, on the leading accounts, that no central observer *module* has been evidenced — that the occupant assumed to sit behind experience is not found as an object when the system is examined, and is, under these frameworks, better described as a process the system runs. "No observer module has been located" is a statement about what neuroscience has and has not found. "You are not real" is a different statement, and is not made here.

Metzinger's phenomenal self-model offers one version of the framework. Under this account, the self is a representation the organism constructs and uses to navigate — a model of the organism, generated by the organism, for the purpose of regulating the organism. Its defining feature, in Metzinger's terms, is transparency: the model is not experienced as a model. The system cannot, under ordinary conditions, perceive the representational medium, and so it takes the contents of the self-model for a self.

> *The map, on this account, is not seen as a map. It is seen through, and what is seen through it is taken for the territory.*

Dennett's framework arrives at a compatible functional position by a different route. On his account, the self is a center of narrative gravity — an abstraction, like a center of mass, that is real in the sense that it organizes and predicts behavior, but is not located at any point in the tissue and is not a thing one could isolate. It is the protagonist of the narrative the interpreter of Section 02 continuously produces, and it has the kind of existence that a protagonist has.

What both frameworks accommodate, and what functional decomposition of the brain is independently taken to support, is that the operations attributed to a unified self are distributed. Interoception, autobiographical memory, agency attribution, bodily ownership, and temporal binding are subserved by different and dissociable systems, each of which can be selectively disrupted — as Sections 02 through 04 demonstrate — without abolishing the organism's capacity to act or report. The "self," on this reading, is not a single system that breaks. It is a set of functions that can be independently degraded. No site has been found at which they converge into a viewer.

The conservative statement of the position is the one this paper holds to. The organism exists. The processing occurs. The sovereign observer assumed to be directing the system is, on the leading accounts, an operational

utility generated by the system to regulate itself — a model the system runs in order to model itself, and which the system, on these accounts, cannot ordinarily see as a model.

The Observer Defect

The trilogy closes on three corrections, each of which removed a presumed enclosure. The universe is not enclosed: *Vectors of the Void* established that the external world admits no edge against which the organism might brace. The body is not solid: *The Inward Vector* established that the assumed interior is process rather than substance. The observer is not sovereign: this paper has established, in its findings, that the observer's reports are late, decentralized, confabulated, and revisable — and has set out, in its models, the leading accounts of why such a system would nonetheless report itself as present, unified, authoritative, and continuous. The findings are not speculative. They are operational. The models are marked as models, and they remain so at the close.

The prior two papers each preserved a residual position from which the correction could be read. The reader could contemplate the unenclosed universe and the dissolving body from somewhere — from the standpoint of the one doing the contemplating. This paper removes that standpoint. There is no external deceiver. The defect is the operating condition. And there is no corrective position outside cognition, because the faculty that would perform the correction is the faculty under audit.

> *Awareness of the defect is processed by the defect. The recognition that the present is assembled is itself assembled and delivered late.*

This is why the mitigation status is what it is. The cognitive fictions documented here are not errors to be patched. The findings are the architecture. The present is assembled. The executive is divided. The reason is confabulated. The archive is mutable. The perception, on the leading account, is predictive. The self, on the leading account, is stateless. The reader cannot step outside their nervous system to verify any report this paper has prompted that nervous system to generate, including the report that they have understood it. There is no clean read.

The thesis the sequence opened with holds at its close. The most frightening thing is not the unknown. It is the known, stated clearly — and what has been stated clearly here is not a hidden horror but the ordinary specification. Nothing in this document required mystery to arrive at its terminal implication. The horror is not that the system is broken. The horror is that it functions exactly as designed.

> *The body is dissolution under management; the mind is a hallucination of sovereignty inside a system that never required its permission.*

The reader's act of comprehension — occurring now, late, narrated, and reconstructed — is the defect in operation.

Factual Anchors

Every flat-voice FINDING claim in the prose traces to an entry below. MODEL entries are named frameworks, not audited facts.

SECTION	CLAIM	SOURCE	REG.
01 Latency	Conduction velocities (~30–70 m/s myelinated; <2 m/s unmyelinated); ms synaptic delay	Hodgkin & Huxley; clinical neurophysiology	FINDING
01 Latency	Saccadic suppression (tens of ms per movement)	Volkmann; Ross, Morrone & Burr	FINDING
01 Latency	Postdiction; cross-modal windows (tens to ~100 ms)	Eagleman & Sejnowski	FINDING
02 Executive	Split-brain interpreter; confabulated rationale	Sperry; Gazzaniga	FINDING
02 Executive	Blindsight (behavior guided by unreported vision)	Weiskrantz	FINDING
02 Executive	Anosognosia (denial of deficit)	Babinski; clinical neuropsychology	FINDING
03 Agency	Choice blindness (defense of unselected choices)	Johansson et al.	FINDING
03 Agency	Post-hoc causal misattribution	Nisbett & Wilson, 1977	FINDING
03 Agency	Readiness potential vs. stochastic-accumulator reframing (unresolved; not load-bearing)	Libet 1983; Schurger, Sitt & Dehaene 2012	DISPUTE
04 Archive	Misinformation effect	Loftus & Palmer	FINDING
04 Archive	Memory reconsolidation (retrieval-induced lability)	Nader, Schafe & LeDoux	FINDING
05 Interface	Predictive coding / free-energy principle	Rao & Ballard; Clark; Friston	MODEL
05 Interface	“Controlled hallucination” terminology	Seth	MODEL
06 Self	Phenomenal self-model (transparency)	Metzinger	MODEL
06 Self	Self as center of narrative gravity	Dennett	MODEL
06 Self	Dissociable self-functions (DMN; temporoparietal junction)	Neuroimaging & lesion literature	SUPPORT