

The Noetic Archive: Infrastructure for a Reproducible Science of Consciousness

A Research Agenda and Methodological Framework

Authors: Woodley Brown, Co-Founder, Psyntient **Status:** Working Paper — Draft, Not Peer Reviewed **Note:** This document is an intentionally high-level research vision paper. Technical methodology and implementation details are described at a conceptual level and will be elaborated in future peer-reviewed publications as the Archive matures.

Abstract

Three distinct bodies of knowledge about the human mind — contemplative and phenomenological traditions, neuroscience, and artificial intelligence — have developed largely in isolation from one another, each producing real knowledge that the others cannot use. This paper introduces the Noetic Archive, a versioned, decentralized, multimodal data infrastructure developed by Psyntient that pairs neural recordings with structured first-person reports of experience, and serves as the kernel of a global R&D operating system for the science of consciousness — a full-stack platform that allows isolated research efforts to contribute to and draw from a common, citable, evolving knowledge base. We describe the Archive's current architecture and current status — including its reliance on simulated data during an initial bootstrap phase — establish the philosophical foundations and ontological commitments the project rests on, and address a series of methodological and philosophical objections the project anticipates. The paper closes with a forward-looking research agenda and a detailed mapping of the Psyntient stack. This paper does not report novel empirical findings. It is offered as orientation, infrastructure justification, and a research agenda, to be tested by subsequent empirical work once the Archive ingests real-world recorded data.

Keywords: neurophenomenology; consciousness; reproducibility; data infrastructure; phenomenology; neural decoding; brain-computer interfaces; replication crisis; observation packets

1. Introduction

1.1 The Fragmentation Problem

For nearly as long as people have reflected on their own minds, three distinct bodies of knowledge about consciousness have developed almost entirely apart from one another. Contemplative and introspective traditions, from classical meditative lineages to twentieth-century phenomenology to the resurgence of structured psychedelic research over the past two decades, have produced detailed, sustained accounts of inner experience from the first-person vantage point. Psychedelic neuroscience offers a particularly clear modern example of what is possible when subjective report is taken seriously as data: structured instruments such as the Mystical Experience Questionnaire (Griffiths, Richards, McCann, & Jesse, 2006; MacLean, Leoutsakos, Johnson, & Griffiths, 2012) have been paired with neuroimaging to support real scientific claims about altered states, and frameworks such as REBUS (Carhart-Harris & Friston, 2019) have begun connecting those reports to specific neural and computational mechanisms.

Neuroscience more broadly, over the past century, has produced an enormous and rapidly growing record of what brains do, but typically without a parallel, systematic record of what the person having that brain activity was actually experiencing at the time. Artificial intelligence, more recently, has learned to model an extraordinary range of human-generated language and behavior, yet has had little direct access to structured data about subjective experience itself, as opposed to text written about it after the fact.

The most direct prior attempt to bridge the first two of these streams was Varela's (1996) proposal of neurophenomenology, which called for combining disciplined first-person method with neuroscientific investigation, building on a broader effort to integrate contemplative method with cognitive science (Varela, Thompson, & Rosch, 1991). The present project owes much of its conceptual lineage to that proposal, even as it differs substantially in scale and method. Each stream has produced real knowledge; none of them, alone, can answer the questions that arise where all three meet. This paper is about the consequences of that gap, and about one proposed way of beginning to close it.

1.2 The Replication Crisis, and the Part of It We Can Address

Over roughly the past fifteen years, a substantial portion of the behavioral and social sciences has confronted what has come to be called the replication crisis — the discovery that a meaningful share of published findings, when independently re-tested, fail to hold up (Ioannidis, 2005; Open Science Collaboration, 2015). The causes are varied and well documented elsewhere: underpowered studies, selective reporting, flexible analytic choices, and a publication system that rewards novel positive results over careful null ones. This paper does not claim to resolve most of that.

What we want to isolate is a narrower, more tractable piece of the problem: research on the structure of human experience is conducted, almost without exception, in isolated labs using incompatible data formats, ad hoc terminology, and instruments that cannot be meaningfully compared across studies — which makes cross-validation and meta-analysis difficult even when individual studies are conducted well. That is a data-infrastructure problem, distinct from though related to the statistical and incentive problems usually discussed under the replication-crisis heading, and it is the piece the infrastructure described in this paper is built to address.

1.3 The Central Research Questions

The hypothesis motivating this project is, at one level, almost trivially true: patterns in neural activity can be statistically associated with patterns in reported experience, in the same uncontroversial sense that any two streams of data can be correlated with each other. Stated this way, the claim is not interesting, and it is not what is actually at stake.

The substantive question is how far that association can be pushed — whether it remains a matter of coarse classification, sorting a recording into one of several broad reported-state categories, or whether it can approach genuine decoding, recovering something like the specific representational content of an experience from neural data alone. A second, equally important question follows: how valuable are mappings of either kind, and to whom? Different answers point toward genuinely different audiences and different standards of evidence; a frontier AI lab, a brain-computer-interface researcher, a philosopher of mind, and the individual contributor whose own session is being recorded all have different things they would need such a mapping to do. Neither question can be meaningfully pursued without first solving a more basic problem: how to collect, standardize, and accumulate the kind of evidence that would let us answer them at scale. That problem is the subject of Section 2.

2. Psyntient as a Global Neurophenomenological R&D Operating System

2.1 Why This Requires Global Infrastructure

The argument for global, decentralized data collection in this domain is not a novel one; it is an application of an argument the broader neuroscience and genomics communities have already accepted for adjacent problems. Large-scale, multi-site resources such as the UK Biobank (Sudlow et al., 2015), the Human Connectome Project (Van Essen et al., 2013), and the ENIGMA consortium (Thompson et al., 2014) exist because single-lab studies, however well designed, cannot achieve the sample sizes, demographic diversity, or cross-site reliability needed to detect real, generalizable structure in brain-behavior relationships.

Contemporary neurophenomenological research faces the same problem in an especially acute form. A psychedelic neuroscience lab in Baltimore, a contemplative research group in London, and a meditation retreat center in India may all be measuring overlapping experiential states with overlapping instruments — but their data lives in separate silos, formatted differently, labeled with different terminology, and subject to different preprocessing choices. No one can pool it. No one can verify that what one group calls non-dual awareness and what another calls ego dissolution are the same thing or not, because there is no shared reference point against which both can be checked. The result is a literature that accumulates citations without accumulating genuine cumulative knowledge.

The problem facing a science of neural-experiential mapping is at least as severe as the one these resources were built to solve, and arguably more so: mapping recurring patterns of experience onto neural activity requires characterizing a joint distribution across two high-dimensional, loosely-understood spaces at once — phenomenological variation and neural variation — which is sparser and harder to power adequately than either alone. If single-site sample sizes were already insufficient for genetic and connectomic questions, they are systematically more insufficient here. What this domain needs — and what no individual lab, consortium, or research program can build alone — is an operating system. The following section introduces the one Psyntient is building.

2.2 The Operating System: Introducing the Psyntient Stack

Imagine infrastructure for conducting the science of inner experience on a global scale — one on which every researcher, clinician, and serious consciousness explorer can bring what they observe, and have it captured, protected, and made useful for the field as a whole. A place where the tools to run neurophenomenological studies, analyze data, and share findings all live in one system and speak the same language. The shape this takes is one the tech world already knows well: an operating system.

An operating system works because every layer knows its job. Hardware generates raw signal. A translation layer converts that signal into a standardized format the rest of the system can work with. A kernel — the core — manages the shared resource, enforces the rules, and makes sure everything that enters is stored faithfully and can be retrieved exactly as it was. A permissions layer controls who owns what and who can access it. Background processes maintain and refine the system over time. An interface layer lets external programs interact with the core without touching it directly. An application ecosystem builds on top of the stable foundation below. And a governance layer — human, expert, distributed — decides what gets accepted into the shared record and what doesn't.

Psyntient has built each of these layers for the science of consciousness. The analogy is not decorative: every major component of the Psyntient stack maps onto a recognizable layer of that architecture, from the hardware that generates raw signals to the governance body that keeps the shared record accountable. At the center of that stack — its kernel — is the Noetic Archive. We return to the full component mapping in detail in Section 6. What matters here is the orienting insight: Psyntient is not conducting a research project. It is building the substrate

on which a generation of research projects — distributed globally, heterogeneous in method and population, asking different questions with different instruments — can for the first time build on one another's results rather than starting from scratch.

Psyntient is a research and technology company whose mission is to build the shared infrastructure layer that a reproducible science of human consciousness currently lacks. Its primary artifact is the Noetic Archive: a living, versioned, multimodal database that systematically pairs neural and physiological recordings with structured first-person accounts of the states being recorded. The Archive is not itself a research project in the conventional sense — it does not advance a single hypothesis through a single protocol. It is the connective tissue meant to allow many such projects to share a common evidentiary foundation.

2.3 How the Archive Actually Works

The Observation Packet

The atomic unit of evidence in the Noetic Archive is the Observation Packet: an immutable, time-stamped record of a single measurement window of an internal state. Each packet binds whatever neural or physiological signal was captured during that window to the participant's structured first-person report of the same window, when a report is available.

Packets without an accompanying report are still admitted, but only as supporting evidence for existing patterns; a new archetype in the taxonomy is never defined from neural data alone. Once ingested, packets are append-only: any later correction or reinterpretation creates a new, separately versioned record rather than altering the original, so that any claim the Archive eventually makes can be traced back to the exact raw evidence that produced it. The schema is built to be modality-agnostic from the outset, designed to accommodate electroencephalography, functional and structural neuroimaging, peripheral physiological measures, eye-tracking, and other behavioral or biometric streams as they become available.

The Layered Architecture

Above the level of individual packets, the Archive organizes evidence into a deliberately layered architecture, separating raw measurement from interpretation so that the two cannot contaminate one another.

The first layer is the packet evidence itself. The second is the ontology layer, populated by neural archetypes: joint clusters in phenomenological and neural feature space that represent a recurring pattern of reported experience together with its associated neural signature. Each archetype carries an explicit confidence rating that reflects how many independent exemplars support it and how consistently those exemplars converge. A single observation can belong to more than one archetype at varying confidence levels, which is a deliberate design choice: subjective experience does not reliably resolve into mutually exclusive categories.

A third layer allows related archetypes to be grouped into higher-order genera, either because they share deep phenomenological or neural structure, or because they represent ordered stages of a single unfolding episode. A fourth, largely infrastructural layer holds the embeddings, similarity scores, and retrieval structures that support search and refinement but make no independent claims of their own. Visualizations and other interpretive artifacts produced from this layer are explicitly marked as derived, and are never treated as part of the Archive's canonical evidence base.

The Refinement Pipeline

New evidence moves through a tiered refinement process designed to keep the cost of automated analysis proportionate to the stakes of the decision being made. An initial deterministic pass checks each incoming packet for basic data integrity before any interpretive model sees it. A lightweight automated pass then screens packets against the existing archetype set; those that do not map confidently to anything already defined enter a holding queue, which is periodically scanned for convergent clusters that might justify a new archetype.

A minimum threshold of independent, converging exemplars — preferably drawn from different participants rather than repeated sessions with the same person — is required before a new archetype can even be proposed, in order to guard against generalizing from a single idiosyncratic case. A deeper, more deliberate refinement pass runs on a slower cadence to propose new archetypes and genera, sharpen boundaries between heavily overlapping archetypes, and flag drift — cases where an archetype's stated definition no longer matches the exemplars currently assigned to it. Sensitive decisions in this process, including advancing an archetype's confidence rating or merging two archetypes into one, remain subject to human review rather than being finalized automatically.

Versioned Editions as Citable Scientific Outputs

The Archive does not publish a single, continuously mutating dataset. Instead, it periodically freezes its current taxonomy into an immutable, versioned Edition: a self-contained package of canonical archetypes and packets, the mappings between them, integrity checksums, and a plain-language scientific overview, all stored in diffable, text-based formats and tracked under standard version control.

Any claim made against the Archive can be traced to the specific Edition that produced it, and any two Editions can be directly compared to see exactly what changed and why. This is the mechanism by which the Archive aims to make a notoriously difficult kind of science reproducible: not by claiming any single Edition is final or correct, but by making every version citable, inspectable, and falsifiable against the next one. Editions function, in this sense, less like journal articles and more like versioned software releases — each one a stable reference point, with a changelog, against which downstream work can orient itself.

2.4 A Distributed Network of Human Nodes

Psyntient's technical infrastructure is only half of the picture. What animates it — and what distinguishes it from a conventional research database — is the worldwide network of researchers, labs, and expert contributors who feed it data, interrogate its outputs, and shape its evolving ontology. Individual contributors act as nodes, each operating independently with their own instruments, populations, and research questions, while the Archive serves as the common substrate their contributions accumulate against: a decentralized contribution model feeding a centralized, versioned, append-only record that no single node controls and any node can audit.

Researcher Access and the Noetic Interface

Researchers engage with the Archive through the Noetic Interface: a conversational agent and backend API that provides access to the shared dataset and, for research use, a private analytical workspace. Through this workspace, a researcher can bring their own experimental data, run analyses under their own methodology without the Archive imposing a model, and compare their results against the broader corpus — asking whether their observed patterns align with existing archetypes, diverge from them in meaningful ways, or motivate the proposal of new ones. Every external study that engages the Archive in this way sharpens the shared ontology, and every contributor benefits from work done by all others: this is the network effect that isolated lab studies cannot produce.

Concretely, the Archive offers researchers access to paired neural-phenomenological data at a scale and modality diversity no individual lab can assemble — spanning EEG, fMRI, fNIRS, MEG, ECoG, BCI streams, biometric wearables, and eye-tracking, each paired with structured first-person reports. Its standardized packet schema and shared archetype membership make cross-study aggregation and meta-analysis tractable across labs and Editions in a way the current literature, with its ad hoc taxonomies and incompatible formats, does not support. Reproducibility is enforced structurally: researchers cite specific, immutable Editions the way they would cite a software version, and any downstream analysis can be reproduced against the exact Edition that produced it. The system is also designed to be IRB-friendly by construction — consent and data provenance are intrinsic to every data record rather than managed separately, because every neural recording lands first in the contributor's encrypted Personal Neural Vault and enters the shared Archive only under explicit, revocable participant consent.

The Science Advisory Network

The human governance layer of the Psyntient stack is formalized through the Science Advisory Network: a carefully assembled, project-based community of independent researchers and scholars whose expertise directly bears on the questions the Archive is built to address. Membership is organized around three areas of inquiry — the neural, cognitive, and phenomenological study of conscious experience; technical systems including AI, neural decoding, and computational methods; and ethics and governance of responsible neuroscientific research — with applicants drawn from neuroscience, cognitive science, philosophy of mind, bioethics, and adjacent fields.

The Network is not a traditional advisory board. Members engage as independent consultants on specific projects they choose to join, contributing methodological review, experimental co-design, or ontological guidance on a project-by-project basis rather than carrying a standing administrative role. Within the Network, a small senior group of Principal Scientific Advisors works directly with Psyntient leadership to assess the Archive's scientific progress and set the research agenda the broader Network helps execute. In this way, the Science Advisory Network functions as the human complement to Psyntient's technical infrastructure: the distributed layer of expert judgment that keeps the evolving ontology accountable to the standards of the disciplines it serves.

2.5 Neural Decoding: Current State, Future Trajectory, and the Archive's Role

It is worth being precise about where the frontier of neural decoding actually sits today, both because precision protects the credibility of this project and because it explains a specific design choice in the architecture just described — and a specific long-term ambition.

What the Field Can Do Now

The foundational demonstration that brain states carry recoverable information about mental states predates the modern decoding literature by several decades. Libet, Gleason, Wright, and Pearl (1983) showed that a measurable neural signal — the readiness potential — precedes a participant's conscious awareness of their own intention to act by several hundred milliseconds, establishing that the brain encodes information about impending mental events before those events reach conscious report. This was an early and striking demonstration that neural activity is not merely a passive correlate of experience reported after the fact, but that it anticipates and encodes mental content in ways that are, in principle, readable from outside.

At the level of classification, the field is now genuinely mature: many systems can reliably sort a recording into one of a small number of pre-specified categories, distinguishing a focused-attention state from a relaxed one, or one type of motor intention from another, especially using invasive recording methods. At the level of content — meaning the recovery of something like the specific semantic substance of an experience or thought rather than its broad category — the field is far younger. Large-scale semantic mapping work using functional MRI (Huth, de Heer, Griffiths, Theunissen, & Gallant, 2016) has shown that conceptual content is organized into stable, mappable patterns across the cortex. More recent work has demonstrated that continuous natural language can be partially reconstructed from non-invasive fMRI recordings using a language model as a decoding prior (Tang, LeBel, Jain, & Huth, 2023). The most striking content-level decoding to date, however, has come from invasive methods: intracortical electrode arrays have supported high-performance text communication from attempted handwriting movements in a paralyzed participant (Willett, Avansino, Hochberg, Henderson, & Shenoy, 2021).

Two honest caveats follow from this picture. First, nearly every content-level result to date depends on either invasive electrodes or fMRI's relatively high spatial resolution.

Consumer-grade EEG — the instrument behind Psyntient Ground, the company's forthcoming wearable neural recording device, and the likely dominant modality for near-term Archive ingestion — offers comparatively coarse spatial resolution, making it considerably better suited to classification than content recovery in its current form. Second, even the fMRI-based language reconstruction work required extensive subject-specific calibration data and produced approximate semantic gist rather than verbatim transcript, underscoring how far the field remains from reliable, general-purpose decoding.

Why the Archive Still Requires Submitted Reports — For Now

This is precisely why the Architect, the Archive's internal organizing agent, still depends on submitted phenomenological reports rather than attempting to infer them: the technology to make that inference reliably, at the level of phenomenological specificity the Archive's archetypes require, does not yet exist for the modalities the project can deploy at scale. The current dependence on participant report is not a conceptual limitation of the Archive's design; it is an honest acknowledgment of where the broader field stands. The design is built to accommodate this constraint gracefully: reports are required for archetype definition, but not for every packet, and the system is explicitly architected to be agnostic about how the phenomenological side of a future Observation Packet is produced — whether by a human writing a report or by a decoding model inferring one.

How the Archive and the Decoding Field Can Grow Together

The relationship between the Archive and the neural decoding field is not merely one of dependency — it is reciprocal. The decoding field's central bottleneck is the absence of large-scale, standardized, carefully labeled datasets pairing neural recordings with rich descriptions of the states being recorded. That is precisely what the Archive is built to produce. As the Archive ingests real-world data from a growing network of instruments and partner research groups, it will accumulate exactly the kind of corpus — diverse in modality, population, and state-type, and consistently formatted and labeled — that decoding researchers need to train and validate models against. In this sense, the Archive is not only a consumer of progress in neural decoding; it is positioned to become one of the field's primary enabling resources.

The trajectory runs in the other direction as well. As decoding methods improve — as models become capable of inferring richer phenomenological content from neural data with less subject-specific calibration and from lower-resolution instruments — the Archive's own capabilities expand in kind. At the limit, a sufficiently advanced decoding layer would allow the Archive to generate reliable phenomenological annotations from neural recordings alone, without requiring a submitted report from the participant at all. This would dramatically lower the barrier to data contribution, expand the range of states and populations the Archive can study, and resolve one of the field's deepest methodological tensions: the fact that the act of producing a report may itself alter the state being reported on. Closing that gap — or mapping its boundary as precisely as possible — is one of the central long-term aims of the research program this paper introduces.

2.6 Current Status of the Archive

As of this writing, the Archive's existing Edition is built entirely from simulated exemplars: synthetic neural-phenomenological pairings generated to pressure-test the packet schema, the archetype taxonomy, and the refinement pipeline before real recorded sessions begin flowing in at scale. We state this plainly rather than allowing it to be inferred, because the distinction matters enormously for how the current taxonomy should be read.

The archetypes currently defined in the working prototype are not empirical discoveries about human experience; they are a working demonstration that the infrastructure described in Section 2.3 functions as intended, end to end, from ingestion through refinement through versioned release.

The system is already built to enforce this distinction going forward: archetypes are explicitly barred from advancing to higher confidence ratings without independent, non-simulated exemplar support, and the project's own internal review process has already reversed at least one premature advancement specifically because it lacked that support — demoting a previously rated archetype back down once it became clear its evidentiary base did not yet meet the required bar. This is not a theoretical safeguard; it is a mechanism that has already fired.

Beginning with real recorded data — whether from Psyntient Ground or from partner instruments and datasets — is the next phase of this work, and it is the point at which the claims this taxonomy makes will actually become testable. Ongoing progress against that transition is tracked in a recurring State of the Archive report and in each successive versioned Edition, both of which are intended to supersede the static description offered here as the underlying data evolves. Before turning to the philosophical objections the project anticipates, it is worth being explicit about the foundational commitments the system rests on — and the considerably larger set of questions it deliberately leaves open.

3. Philosophical Foundations

3.1 Foundational Assumptions: What This Project Does and Does Not Claim

The following three commitments are principled positions, not hedges. Understanding them is necessary for evaluating both the system's design and the objections addressed in Section 4.

On the meaning of "phenomenology." The word carries two related but distinct meanings that cause real confusion in interdisciplinary work. In one sense, it names a specific philosophical discipline and method, originating with Husserl and developed by a long line of successors (Zahavi, 2003), concerned with rigorously describing the structures of conscious experience from a disciplined first-person standpoint. In a much looser and more common sense, the word

simply means something like inner knowing, or direct subjective experience, with no commitment to that tradition's methods. This paper uses the term in a third, narrower, deliberately stipulated sense: phenomenology here refers to structured first-person report — a participant's account of an internal state, captured in language or in response to structured prompts. We make no claim that this exhausts what experience is, or that experience and its report are the same thing. The gap between reportable experience and whatever may lie beyond report is philosophically real and remains open; we return to it in Section 4.

It is also worth noting that phenomenology and consciousness are not interchangeable terms, though they are often used loosely as though they were. Consciousness names the broader phenomenon — the fact that there is something it is like to be a given creature at all. Phenomenology, in even its loosest sense, refers to the structured content and character of that experience. The relationship between the two — whether consciousness is primary and phenomenological structure is its appearance, or whether mind is the primary category and consciousness one of its features — is itself one of the deepest open questions in philosophy of mind, and one this project takes no position on.

On what the project assumes about the brain and experience. We assume only that mental states supervene on neural and physiological states in at least the weak sense long discussed in philosophy of mind (Kim, 1998): that no change in experience occurs without some accompanying change in the physical substrate, and that this covariation is, in principle, statistically detectable from outside. We take no position on whether neural correlates constitute, cause, or merely accompany experience, and we do not claim to make progress on what Chalmers (1995) named the hard problem of consciousness — the question of why there is something it is like to be a given physical system at all. Weak supervenience of this kind is compatible with nearly every serious position in contemporary philosophy of mind — functionalist, identity-theoretic, property-dualist, panpsychist, and most idealist accounts all accept some version of it. The project's empirical program does not need the hard problem resolved to proceed, and takes no position on which metaphysical framework is correct. The infrastructure described in this paper is designed to be useful regardless of which view turns out to be right.

On what the project does not claim. This paper does not claim to solve the hard problem of consciousness, to establish that neural data fully captures the nature of experience, or to adjudicate between competing metaphysical frameworks for understanding the mind-body relationship. What it does claim is that a structured, reproducible mapping of the relationships between neural activity and reported experience — however philosophically incomplete — is both scientifically valuable and currently impossible to pursue at the scale the question requires. Section 4 addresses the objections this claim invites.

4. Objections and Responses

The following objections are those we anticipate from academic readers across philosophy of mind, neuroscience, cognitive science, phenomenology, and research ethics. We address each directly, drawing on the commitments established in this paper. Where an objection identifies a genuine limitation rather than a misunderstanding, we say so plainly.

4.1 "This doesn't get you any closer to solving the hard problem of consciousness — and it assumes a physicalist metaphysics that not everyone accepts."

From: Philosophy of mind, consciousness studies, metaphysics

The first part of this objection is correct, and we accept it without qualification. The hard problem, as Chalmers (1995) formulated it, asks why there is subjective experience at all — why any physical process gives rise to something it is like to undergo, rather than proceeding in the dark. Nothing in the Archive's design addresses that question, and we make no claim that it does. Correlating neural patterns with reported experiential patterns, however precisely, cannot close the explanatory gap between physical description and phenomenal fact, because the gap is not a gap in data — it is a conceptual gap about what kind of thing data can explain.

What we do claim is that this limitation is shared by every empirical science of mind that has ever existed or could plausibly exist, and that it does not impede the Archive's actual aims. The Archive is built to map structure — recurring patterns in the joint space of neural activity and reported experience — and to make those mappings reproducible, citable, and falsifiable. That is a tractable program, and its value does not depend on resolving the hard problem any more than the periodic table's value depended on resolving the metaphysics of matter. Even if the hard problem is permanently unsolvable — even if no empirical program can ever close the explanatory gap — running the Archive's global mapping program will at least establish the precise boundary of what correlational and decoding methods can recover. Mapping that epistemic boundary rigorously is itself a finding about the nature of consciousness, and a valuable one. A science that discovers the limits of what it can know is doing real science, not failing at it.

The second part of the objection — that the project smuggles in physicalist assumptions — requires a more careful response, because it is a misreading of what the Archive actually commits to. The project's only metaphysical commitment, stated explicitly in Section 3.1, is weak supervenience: no change in experience without some accompanying change in the physical substrate, such that the covariation is in principle statistically detectable. This is a significantly weaker claim than physicalism, which asserts that mental states are identical to, or fully reducible to, physical states. Weak supervenience is compatible with physicalism, but it is equally compatible with property dualism, panpsychism, and even idealist frameworks in which physical descriptions are themselves understood as a particular mode of representing a more fundamental conscious reality — on such a view, the covariation between neural activity and

reported experience follows not because the brain produces consciousness but because both are expressions of the same underlying ground, read from different vantage points.

Psyntient takes no position on which of these frameworks is correct. This is not diplomatic ambiguity — it is a principled epistemic stance. The Archive is designed to accumulate structural evidence about the relationship between neural and phenomenological patterns, and that structural evidence is, in principle, compatible with multiple incompatible metaphysical interpretations. We believe this is the right posture for a scientific infrastructure project to take: the Archive should be a resource that a committed physicalist and a committed idealist can both use and both cite, without either finding that the instrument has pre-adjudicated the question they most care about. Whether consciousness is produced by the brain, correlated with it, or something the brain partially expresses — questions that serious researchers hold genuinely different views about — is not a question the Archive can or should settle. What it can settle, over time, is how far the structural mapping goes, and where it stops. That boundary, wherever it falls, will be informative to every position in the debate.

4.2 "You're not doing phenomenology — you're doing something much weaker and calling it by a prestigious name."

From: Continental philosophy, Husserlian and post-Husserlian phenomenology

This is a serious objection that deserves a direct answer rather than deflection. The phenomenological tradition, from Husserl through Merleau-Ponty and beyond, developed rigorous methods — epoche, eidetic variation, micro-phenomenological interview — for describing the structures of experience from a disciplined first-person standpoint, precisely because casual introspection and verbal report are not the same thing as phenomenological method. A participant filling out a structured questionnaire after a meditation session is not doing phenomenology in this sense, and we should not claim otherwise.

We accept this. The Archive's operational definition of phenomenology — structured first-person report — is explicitly narrower than the tradition's methods, and we flagged this in Section 3.1. What we would resist is the inference that the narrower definition is therefore worthless or misleading. Structured self-report, even without the full apparatus of phenomenological method, is genuine first-person data about experience; it is simply not the most refined possible version of such data. The Archive is designed to accommodate more rigorous first-person methods as they develop: the Noetic Interface's support for micro-phenomenological interview protocols, for instance, is a natural extension of the existing packet schema, and the Science Advisory Network explicitly solicits methodologists working on exactly this problem. We treat the current dependence on structured report not as a ceiling but as a starting point, and the objection identifies an important direction for the program's development rather than a fatal flaw in its foundation.

4.3 "Self-report is notoriously unreliable as a window onto experience. Introspection is not trustworthy data."

From: Experimental psychology, cognitive science

The unreliability of introspective report is one of the best-documented findings in experimental psychology (Nisbett & Wilson, 1977; Schwitzgebel, 2011), and we take it seriously as a methodological constraint rather than dismissing it. Participants confabulate, rationalize, apply culturally available narratives to experiences that may not fit them, and systematically misremember states that have already passed. None of this is controversial, and none of it is a surprise.

Several features of the Archive's design are direct responses to this problem. First, the phenomenological side of an Observation Packet is always paired with a neural record of the same window, which means that divergence between what a participant reports and what the neural data suggests is itself informative — it becomes a signal to investigate rather than noise to discard. Second, archetype definitions require convergence across multiple independent exemplars from different participants; a single participant's idiosyncratic report cannot define an archetype. Third, the confidence rating system and drift-detection mechanism exist precisely to catch cases where the phenomenological signature of an archetype begins to drift away from what the underlying exemplars actually show, flagging when the ontology has outrun its evidence.

We would also note that the alternative to using self-report — discarding first-person data entirely and relying on neural signal alone — does not solve the reliability problem; it simply trades one set of limitations for another. A neuroscience of inner states that contains no first-person evidence about those states has no principled way to determine what it is a neuroscience of.

4.4 "Requiring linguistic report assumes that experience either is linguistic or is fully captured by linguistic description — which is philosophically contentious."

From: Philosophy of mind, phenomenology, cognitive science

This objection identifies a real tension in the project, and we want to address it carefully rather than dissolve it too quickly. Nagel's (1974) argument that there is something it is like to be a bat, which no amount of physical description can capture, is not directed specifically at linguistic report — but it motivates a general suspicion of any methodology that makes experience hostage to its own articulability. Block's (1995) distinction between phenomenal consciousness and access consciousness presses a related point: the portion of experience available for verbal report may be a proper subset of experience as such, with the remainder not merely difficult to describe but structurally inaccessible to report.

Our response has two parts. The first is methodological: the Archive only claims to study what its instruments can reach. Whatever falls outside the boundary of structured report — if anything does — is outside the Archive's scope, and we make no claims about it. This is a limitation we acknowledge rather than paper over, and it means the Archive's findings should always be understood as findings about the reportable structure of experience, not about experience in its entirety. The epistemic humility here is genuine.

The second part concerns the relationship between language and experience more directly. There is a credible philosophical position — running through Merleau-Ponty's account of speech as constitutive rather than merely expressive of thought, and echoed in Wittgenstein's (1953) arguments against the coherence of a purely private language — that for reflective, linguistically competent adult humans engaging deliberately with their own inner states, the articulation of an experience and the experience itself are not cleanly separable stages. The rendering of an experience into language may, in such cases, be part of how that experience constitutes itself as a determinate object of reflection rather than a subsequent translation of something already fully formed. We do not ask the Archive to depend on this stronger claim, but we note that it is available and not obviously false. In any case, the boundary between reportable and unreportable experience is itself an empirical question — one the Archive, as it accumulates data, is positioned to help investigate.

4.5 "This just adds another irreproducible dataset to a field already drowning in irreproducible ones."

From: Neuroscience, methodology, open science

This objection has force precisely because it takes the replication crisis seriously, which we also do. The response is structural rather than rhetorical. Most datasets in this domain fail reproducibility requirements because they are siloed, undocumented, preprocessed in undisclosed ways, and stored in formats that make independent reanalysis difficult or impossible. A researcher who wants to reproduce a finding from a 2019 meditation neuroscience paper typically cannot: the raw data is not available, the preprocessing pipeline is partially described in a methods section, and the taxonomic labels are idiosyncratic to that lab.

The Archive addresses this through its versioned Edition system, which we described in Section 2.3. Every archetype and every mapping is part of a specific, immutable Edition that is integrity-checksummed, stored in diffable text-based formats, and citable by version. Any analysis made against a specific Edition can be reproduced against that exact Edition by any researcher with access. This is reproducibility enforced by architecture rather than promised by convention. The Archive does not ask researchers to trust that a dataset was handled responsibly; it makes the handling inspectable.

We would also note that the Archive is explicitly designed to absorb and reconcile data from multiple contributing labs rather than adding a new silo to the existing collection. The long-run value is precisely that it gives existing siloed datasets a common address to merge against.

4.6 "Neural data is among the most sensitive personal data that exists. What are the privacy and consent guarantees here?"

From: Bioethics, data ethics, research ethics

This is the most important practical objection, and it is well-founded. Neural recordings are uniquely sensitive: unlike most biomedical data, they are potentially capable of revealing not just health status but cognitive states, emotional responses, beliefs, and — as decoding technology improves — increasingly specific mental content. The risks of misuse, re-identification, or unintended disclosure are not hypothetical, and they scale with exactly the kind of technical progress this paper elsewhere argues is coming.

The Archive's primary structural response is the Personal Neural Vault: a participant-controlled, encrypted data layer that sits between any individual's recording and the shared Archive. Every neural recording enters the contributor's own Vault first. Nothing passes from the Vault into the shared Archive without explicit, affirmative, and revocable consent, and contributors can withdraw their data at any time. Consent and provenance are intrinsic to every data record rather than managed separately — which means they cannot be lost in preprocessing or stripped in export.

We acknowledge that architectural safeguards are necessary but not sufficient. As the Archive scales and as decoding capabilities improve, the governance questions — who has access to what, under what conditions, with what oversight — will need ongoing, substantive attention from the bioethics community. This is precisely why the Science Advisory Network explicitly recruits expertise in bioethics and research governance, and why we treat these questions as an ongoing research problem rather than a solved one.

4.7 "The Architect agent is a black box. How do you know its ontological proposals aren't artifacts of the model's own biases and training distribution?"

From: AI/ML research, philosophy of science, cognitive science

This is a sharp objection and one of the most technically specific ones the project faces. The Architect is a large language model used to propose new archetypes, refine existing ones, and detect boundary issues in the ontology. Any such model carries the distributional biases of its training data — which, for a model trained primarily on English-language text from a particular cultural and historical moment, likely overrepresents certain kinds of introspective vocabulary,

certain cultural frameworks for inner experience, and certain assumptions about what counts as a discrete and nameable state.

Several design choices in the pipeline are direct responses to this concern. The Architect operates under defined constraints that prevent it from introducing bias at the evidence level. Its proposals are always validated against the actual exemplar packets — an archetype that cannot be grounded in converging real exemplars is not admitted regardless of what the model proposes. And sensitive decisions, including the advancement of archetypes to higher confidence ratings, remain subject to human review.

We nonetheless acknowledge that this is an ongoing methodological challenge rather than a resolved one. Fully characterizing the ways in which a large language model's biases may shape ontological proposals — and developing systematic checks against those biases — is one of the active research problems the Science Advisory Network's technical members are engaged with. Transparency about the Architect's role in each Edition's changelog is part of the answer; developing formal audits of model-introduced bias in archetype proposals is the more demanding part.

4.8 "This is just relabeled pattern-matching. The archetypes are statistical clusters, not discoveries about the nature of mind."

From: Neuroscience, philosophy of science, cognitive science

The objection is partially correct, and the partial correctness is the interesting part. The archetypes are, in one technical sense, clusters — recurring patterns in a joint phenomenological-neural feature space, identified by convergence across independent exemplars. We do not hide this; it is the explicit description of what an archetype is in Section 2.3. The question is whether just clusters is a dismissal or a description, and we think that depends entirely on what one expects scientific ontologies to be.

Every scientific taxonomy is, in some sense, a structured description of recurring patterns in data: the chemical elements, biological species, and DSM diagnostic categories are all, at some level of description, clusters over feature spaces. What distinguishes a useful taxonomy from an arbitrary one is whether the clusters carve nature at its joints — whether they are stable across independent samples, whether they support generalization and prediction, whether they survive contact with new data rather than dissolving into noise. The Archive's design is precisely oriented around these criteria: multiple independent exemplars for archetype creation, explicit confidence ratings, drift detection, and a correction mechanism that has already been used. These are the standards by which any scientific taxonomy should be judged, and they are the standards by which the Archive's taxonomy invites judgment.

What the objection correctly resists is any stronger claim — that archetypes are natural kinds in a metaphysically robust sense, or that they carve experience at its ultimate joints rather than at joints that are stable and useful relative to current data and methods. We make no such claim.

The Archive is a map, not the territory, and every Edition is an explicit, dated, versioned version of the map rather than a claim to have found the final one.

4.9 "By integrating deeply into how people understand their own inner lives, you risk creating a dependency you're ethically unable to dissolve — a cognitive prosthetic people can't live without."

From: Philosophy of mind, bioethics, science and technology studies

This is among the most serious ethical objections the project faces, and it deserves engagement at the level of its philosophical ambition rather than deflection toward policy language. Clark and Chalmers' (1998) extended mind thesis holds that cognitive processes are not necessarily bounded by the skull: when an external resource is reliably available, automatically endorsed, and plays the right functional role in cognition, the system constituted by brain plus resource can count as the cognitive system, not just the brain alone. On a liberal reading of this thesis, any tool deeply enough integrated into how a person navigates their inner life begins to constitute, rather than merely support, their cognitive and phenomenological self-model. The worry then runs as follows: as the Psyntient stack matures and as the Noetic Interface becomes a continuous interpretive layer through which contributors understand their own neural patterns and experiential states, that layer could cross the threshold from tool to cognitive constituent. At that point, withdrawal — whether through corporate failure, acquisition, policy change, or terms-of-service revision — is no longer like cancelling a subscription. It is more like removing a prosthetic.

We want to respond to this objection in three registers: scope, design, and the broader philosophical question the objection opens.

On scope: in its current and near-term form, the Archive is a research infrastructure — a dataset and ontology that researchers query and cite, not an intimate real-time interpretive layer mediating a person's access to their own experience. The dependency risk the objection identifies is real but attaches to a future configuration of the Noetic Interface, not to what this paper is proposing or what the current prototype implements. We take the objection seriously precisely as a design constraint on that future configuration rather than as a critique of what presently exists.

On design: the Personal Neural Vault architecture is partly aimed at this problem — data sovereignty remains with the individual, raw recordings are the contributor's own, and withdrawal is structurally supported at any time. But we acknowledge that data portability alone does not resolve the dependency concern if what a person has come to rely on is the interpretive and meaning-making layer. The more substantive design response is one we are committing to explicitly: the Archive's ontology is and will remain open-format, documented, and exportable. No architectural lock-in on the meaning-making layer is a design principle, not just an aspiration.

On the broader philosophical question: Clark and Chalmers' thesis, if taken seriously, generates dependency obligations not only for neural interfaces but for smartphones, GPS navigation, search engines, and any sufficiently integrated digital cognitive tool — all of which plausibly meet the coupling criteria on a liberal reading. This suggests the problem is one of general digital-dependency ethics, an emerging field the Archive's work will need to engage with, but not a problem any single company or project can solve unilaterally. What Psyntient can do — and commits to doing — is build toward interoperability rather than dependency, treat open standards as a scientific and ethical obligation rather than a commercial afterthought, and ensure that the Science Advisory Network's bioethics representation is specifically tasked with tracking these risks as the Noetic Interface matures.

4.10 (Secondary — for VC readers) "The path to revenue is unclear, and the timeline to real data is too long."

From: Venture capital, early-stage investors

This paper is addressed primarily to an academic audience, and a full response to commercial viability questions belongs in a separate document. We note briefly, however, that the Archive's infrastructure model generates value at multiple layers simultaneously: as a data asset for research licensing, as a validation corpus for the neural decoding industry, as a platform for the Noetic Interface and researcher workspace, and — as Psyntient Ground ships — as the entry point for direct consumer data contribution. The simulated-data bootstrap phase is time-bounded and deliberate; the pipeline and ontology it produces are the durable assets, and the transition to real data ingestion is the next concrete milestone. A fuller treatment of the commercial architecture and timeline is available in the accompanying investor document.

5. Forward-Looking Research Agenda

With the infrastructure described, its philosophical foundations established, and the principal objections addressed, this section turns to what the system is actually built to discover.

5.1 Establishing the Baseline Taxonomy

The first and most immediate research program concerns whether the archetype taxonomy developed during the simulated-data bootstrap phase survives contact with real, heterogeneous, participant-contributed data. Specific questions include: which archetypes replicate cleanly across independent participants and modalities, which require splitting into finer-grained species, which merge under closer empirical scrutiny, and which dissolve entirely — revealing themselves as artifacts of the simulation's generative assumptions rather than real recurring patterns in human experience. The correction and advancement mechanisms already

built into the pipeline are exactly the instruments for this program; what the bootstrap phase cannot test is whether they behave correctly under realistic data distributions. This is the foundational empirical question, and every subsequent research program depends on its answer.

5.2 Cross-Modality and Cross-Population Replication

Once a baseline taxonomy is established from initial data ingestion, the Archive's multi-site, modality-agnostic architecture makes it possible to ask a class of questions that the current literature cannot tractably address: do the same archetypes appear when measured with EEG versus fMRI? Do the neural signatures of a given archetype generalize across different demographic populations, cultural backgrounds, and contemplative traditions, or are they instrument-specific and culturally local? The ENIGMA and Human Connectome Project precedents suggest that cross-site replication at scale routinely reveals both surprising robustness and surprising fragility in findings that looked stable in single-lab studies; we expect the same pattern to hold in neurophenomenological research, and the Archive is the first infrastructure capable of testing it systematically.

5.3 Altered States and the Contemplative Sciences

A specific and near-term target for the Archive's empirical program is the cluster of states studied in contemporary contemplative neuroscience and psychedelic research — meditation, flow, ego dissolution, non-dual awareness, mystical experience, and the hypnagogic and threshold states that occur at the boundaries of wakefulness and sleep. These are already the states most studied by the research communities the Archive is designed to serve, and they represent the area where first-person report has been most systematically developed as a research instrument, from the Mystical Experience Questionnaire (Griffiths et al., 2006) to the structured phenomenological interviews developed in contemplative science (Lutz & Thompson, 2003). The Archive can provide, for the first time, a common ontological substrate against which findings from psychedelic sessions, long-term meditators, and first-time practitioners can be systematically compared rather than treated as categorically separate literatures.

5.4 The Neural Decoding Frontier

As the Archive accumulates a large, standardized, carefully labeled corpus of paired neural-phenomenological data, it becomes increasingly useful as a training and validation resource for the neural decoding field — and the Archive's own research program becomes increasingly interested in what that field can do with the corpus. Specific questions include: which archetype distinctions are decodable from EEG alone, which require higher-resolution modalities, and what is the minimum recording duration and calibration required for reliable classification? Further along: can the phenomenological signature of an archetype be partially inferred from neural data in a participant who has not submitted a report, and if so, at what level of granularity and with what confidence? Each answer to these questions either extends or

maps the boundary of what correlational and decoding methods can recover — which is, as we argued in Section 4.1, a scientifically valuable finding regardless of which direction it points.

5.5 Methodological Research: Refining First-Person Method

The Archive's dependence on structured self-report as its primary phenomenological instrument creates a research obligation as well as a practical constraint. The micro-phenomenological interview tradition developed by Petitmengin (2006) offers a more rigorous method for accessing pre-reflective experiential structure than standard questionnaire instruments, and its systematic integration into the Archive's ingestion pipeline is a methodological research program in its own right — one that the Science Advisory Network's phenomenology members are positioned to lead. Questions here include: how much of the variance in phenomenological self-report is attributable to differences in reporting skill rather than differences in experience, and can structured training in first-person method reduce that variance systematically? The answer matters both for the Archive's data quality and for the broader field's ability to treat self-report as a reliable scientific instrument.

5.6 Ethics, Governance, and the Extended-Mind Problem

As we discussed in Section 4.9, the ethical questions surrounding neural data collection, cognitive dependency, and the obligations generated by deep tool integration do not have settled answers, and the Archive's maturation will generate new instances of them faster than existing frameworks can resolve them. A dedicated ethics and governance research program — embedded in the Science Advisory Network and producing outputs that feed back into the Archive's design — is not optional; it is a condition of the project proceeding responsibly. Specific questions for this program include: what consent frameworks are adequate for data contributed under conditions of altered consciousness, where the capacity for fully informed consent may be genuinely complicated? How should the Archive handle archetype data that has privacy implications beyond what contributors anticipated at the time of consent? And what governance structures are appropriate for an ontology that is collectively produced but institutionally curated?

5.7 Toward a Cumulative Science of Consciousness

The long horizon of this research program is the one that motivates the infrastructure in the first place: a genuinely cumulative science of human inner experience, in which findings compound across labs, modalities, traditions, and Editions rather than accumulating as isolated data points with no common substrate to build on. Whether that horizon is reachable — whether the joint space of phenomenological and neural variation turns out to be structured enough to support a stable, generalizable, cross-cultural taxonomy of inner states — is itself an empirical question, and one the Archive is uniquely positioned to answer. A negative answer, rigorously established, would be one of the most important findings in the history of consciousness research: it would mean that the apparent regularities in first-person report across traditions and individuals are not underwritten by corresponding regularities in neural architecture, and that a unified science

of experience is not achievable on correlational grounds. We do not expect this outcome, but we treat it as live, because a research program that cannot in principle return a negative finding is not a research program — it is a commitment. The Archive is a commitment to the question, not to any particular answer. What that question is, and what the stack built to pursue it looks like in full, is the subject of the final section.

6. Conclusion: The Stack, Assembled

The argument this paper has made is a structural one: that the fragmentation of consciousness research into incompatible data silos is not a temporary inconvenience but a fundamental infrastructure problem, and that the solution takes the form of a global operating system rather than a better individual study. The Psyntient stack is that system. Having described each of its components across the preceding sections, we close by showing how they fit together.

The hardware layer is Psyntient Ground: the company's wearable neural recording device, which generates the raw electroencephalographic and physiological signal that enters the system. As with any OS, the hardware layer is not the only input device — third-party instruments function as peripheral hardware, and the system is designed to be device-agnostic from the outset.

The device driver layer is the ingestion pipeline and Observation Packet schema: the translation layer that takes heterogeneous raw signals from heterogeneous instruments and converts them into a standardized format the kernel can process uniformly. This layer is designed to ensure that the diversity of the hardware layer does not produce chaos at the kernel level — that regardless of instrument or source, data arrives at the Archive in a consistent, comparable form.

The kernel is the Noetic Archive itself: the central, versioned, append-only repository that manages the shared resource, enforces the rules of evidence, and mediates between raw input and the ontological outputs the system produces. Like a kernel, it is not directly manipulated by end users; everything passes through defined interfaces. Like a kernel, its integrity is the condition of everything else functioning correctly, which is why immutability, provenance tracking, and human-in-the-loop oversight for sensitive decisions are architectural requirements rather than policy choices.

The filesystem with its permissions model is the Personal Neural Vault: the encrypted, participant-controlled data layer that sits between any individual's recordings and the shared Archive. Each contributor owns their own partition. The consent architecture is the permissions model — nothing moves from private storage to shared kernel without an explicit, auditable, revocable grant.

The background processes are the Architect and the tiered refinement pipeline: the automated and human-in-the-loop operations that run continuously to maintain the integrity of the ontology, evaluate incoming data, refine the archetype taxonomy, and prepare Editions for release. As

with an OS process scheduler, different operations run at different levels of oversight — routine integrity checks at the base level, deeper ontological reasoning requiring more scrutiny — and decisions that exceed automated authority escalate to human review.

The API and system call interface is the Noetic Interface and researcher workspace: the defined layer through which external researchers interact with the Archive without touching the kernel directly. Queries, analyses, and data contributions flow through this interface, which enforces access controls, logs provenance, and returns reproducible, version-pinned outputs.

The application layer is the Marketplace and the ecosystem of third-party tools that build on the Archive's stable, versioned taxonomy. Applications at this layer depend on the kernel's stability without being responsible for it — they can build on a specific Edition the way an application builds against a specific OS version, confident that the interface below them is stable and citable.

The governance layer is the Science Advisory Network: the globally distributed community of independent researchers, methodologists, and bioethicists who review what the automated system proposes and decide what actually ships in each Edition. Every OS needs a governance body that decides what gets merged and what gets rejected. The Science Advisory Network is the human authority that plays this role for the Noetic Archive — keeping the evolving ontology accountable to the standards of the disciplines it serves.

The network of nodes is the global community of researchers, institutions, partner labs, and individual contributors whose instruments and reports feed the Archive from every geography and tradition. This is where the globalness of the system resides: not in any single installation, but in the worldwide set of contributing nodes that are all writing to and reading from the same shared, versioned, auditable substrate.

What makes this more than an analogy is the specific problem it solves. Computing operating systems exist because writing a new hardware abstraction layer, memory manager, and process scheduler for every application is prohibitively expensive and produces incompatible results. The history of computing before the OS is a history of siloed, incompatible systems that could not share resources or build on each other's advances. The history of consciousness research, as this paper has argued, is precisely that history — and the Psyntient stack is the proposed OS that ends it.

The Archive is, as of this writing, still in its bootstrap phase: its current taxonomy is built on simulated data, its hardware is forthcoming, and its network of contributing nodes is just beginning to form. But the architecture is complete, the pipeline is running, and the governance structures are in place. What follows — the ingestion of real data, the testing of the taxonomy against human experience, the progressive reduction of dependence on submitted report as decoding methods mature, the publication of Editions that the field can cite and build on — is the research program this paper has introduced. It is a large program, pursued by a small team at the beginning of a long effort, with no guarantee that the structural regularities it is looking for are there to be found. The Archive is a commitment to the question. The question is whether

consciousness, examined at global scale with the full stack of available tools, will reveal enough structure to be mapped — and what we learn about ourselves, and about the limits of what we can know about ourselves, in the attempt.

References

- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227-247.
- Carhart-Harris, R. L., & Friston, K. J. (2019). REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3), 316-344.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200-219.
- Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 7-19.
- Griffiths, R. R., Richards, W. A., McCann, U., & Jesse, R. (2006). Psilocybin can occasion mystical-type experiences having substantial and sustained personal meaning and spiritual significance. *Psychopharmacology*, 187(3), 268-283.
- Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600), 453-458.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124.
- Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. MIT Press.
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). *Brain*, 106(3), 623-642.
- Lutz, A., & Thompson, E. (2003). Neurophenomenology: Integrating subjective experience and brain dynamics in the neuroscience of consciousness. *Journal of Consciousness Studies*, 10(9-10), 31-52.
- MacLean, K. A., Leoutsakos, J. M., Johnson, M. W., & Griffiths, R. R. (2012). Factor analysis of the mystical experience questionnaire: A study of experiences occasioned by the hallucinogen psilocybin. *Journal for the Scientific Study of Religion*, 51(4), 721-737.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435-450.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231-259.

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.

Petitmengin, C. (2006). Describing one's subjective experience in the second person: An interview method for the science of consciousness. *Phenomenology and the Cognitive Sciences*, 5(3-4), 229-269.

Schwitzgebel, E. (2011). *Perplexities of consciousness*. MIT Press.

Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., & Collins, R. (2015). UK Biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine*, 12(3), e1001779.

Tang, J., LeBel, A., Jain, S., & Huth, A. G. (2023). Semantic reconstruction of continuous language from non-invasive brain recordings. *Nature Neuroscience*, 26(5), 858-866.

Thompson, P. M., Stein, J. L., Medland, S. E., Hibar, D. P., Vasquez, A. A., Renteria, M. E., et al. (2014). The ENIGMA Consortium: Large-scale collaborative analyses of neuroimaging and genetic data. *Brain Imaging and Behavior*, 8(2), 153-182.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80, 62-79.

Varela, F. J. (1996). Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3(4), 330-349.

Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press.

Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M., & Shenoy, K. V. (2021). High-performance brain-to-text communication via handwriting. *Nature*, 593(7858), 249-254.

Wittgenstein, L. (1953). *Philosophical investigations* (G. E. M. Anscombe, Trans.). Blackwell.

Zahavi, D. (2003). *Husserl's phenomenology*. Stanford University Press.