# Title-: The CISR Framework: A Layered, Axiomatic Constraint on Explanations of Consciousness

**Author -** Rakesh Gosswami
**Affiliation** - Independent Researcher

## 1. Abstract

Consciousness research remains challenged by persistent conceptual confusion regarding the relationship between phenomenology, neural mechanisms, and symbolic or computational models. Explanatory claims often exceed their methodological scope, producing category errors and ontological overreach. This paper introduces the CISR Framework, an axiomatic meta-theoretical structure that does not explain consciousness but constrains what explanatory frameworks can legitimately claim about it. CISR distinguishes a non-operational experiential ground ($C_0$), an observational registration layer (I-layer), neural–functional interpretations (R-layer), and an optional symbolic–representational layer (S-layer). Explicit axioms and mapping rules preserve empirical success while preventing unwarranted reduction of experience to mechanism or representation. The framework applies across neuroscience, cognitive science, AI, and philosophy of mind, clarifying how descriptive, functional, and representational claims relate to phenomenology. Progress within CISR is understood as alignment of explanatory claims with their appropriate methodological domain, maintaining conceptual rigor without invoking additional metaphysical assumptions.

## 2. Keywords

CISR, Consciousness Studies, Axiomatic Framework, I-layer, $C_0$, Phenomenological Integrity

## 3. Introduction

Consciousness research occupies a paradoxical position within contemporary science and philosophy. Empirical disciplines such as neuroscience, cognitive psychology, and computational modeling have achieved remarkable success in identifying neural correlates, predicting behavioral reports, and modeling cognitive functions related to perception, decision-making, and self-representation. Despite this progress, a fundamental conceptual difficulty remains unresolved: the relationship between subjective experience and objective explanation.

This difficulty is not merely empirical but methodological. Explanatory frameworks routinely succeed within their respective domains—accounting for access, reportability, integration, and behavior—yet their success is often interpreted as explaining experience itself. As a result, distinctions between phenomenology, neural function, and symbolic description become blurred. Neural activity is treated as constitutive of experience, functional access is equated with phenomenology, and computational or linguistic models are sometimes reified as conscious systems. These interpretive moves underlie persistent debates concerning reductionism, the "hard problem" of consciousness, and claims about machine consciousness.

Historically, approaches to consciousness have oscillated between reductionist and anti-reductionist extremes. Reductionist positions attempt to identify experience with physical or functional mechanisms, while opposing views emphasize the irreducibility of experience, occasionally placing it beyond the reach of scientific explanation. Despite their differences, both tendencies often share a common limitation: they proceed without clearly specifying the methodological boundaries that govern what different explanatory frameworks are entitled to claim.

The CISR Framework is introduced in response to this limitation. Rather than proposing a new empirical or causal theory of consciousness, CISR functions as a meta-theoretical framework designed to regulate explanatory scope. Its aim is not to explain how consciousness arises, but to clarify how phenomenological, neural–functional, and symbolic explanations relate to one another, and where their legitimate boundaries lie. In this sense, CISR precedes theory construction, providing a structured way to prevent category errors arising from the conflation of explanatory domains.

A recurring source of confusion in consciousness studies is the implicit assumption that explanatory adequacy implies ontological identity. When neural or functional models successfully predict conscious reports, they are often taken to constitute experience itself. Similarly, when artificial systems exhibit functional or representational sophistication, experiential attribution is inferred. CISR challenges these inferences without disputing the empirical success of such models. It maintains that correlation, prediction, and functional explanation do not by themselves license claims about the existence or nature of experience.

By introducing a disciplined separation between experiential registration, neural–functional interpretation, and symbolic representation, CISR preserves the explanatory power of existing scientific models while preventing their overextension into phenomenological domains. The framework does not commit to any particular metaphysical position regarding the nature or origin of consciousness. Instead, it brackets metaphysical questions and focuses exclusively on methodological clarity.

The relevance of this approach extends across neuroscience, philosophy of mind, and artificial intelligence research. As increasingly sophisticated models blur the line between functional competence and experiential attribution, the need for explicit constraints on explanatory claims becomes more pressing. CISR provides such constraints by clarifying what different explanatory practices can legitimately account for, without inflating their scope.

In summary, the CISR Framework addresses a long-standing methodological deficit in consciousness research. It does not compete with existing theories or propose a solution to consciousness itself. Rather, it constrains explanatory ambition, clarifies domain boundaries, and preserves conceptual rigor. Progress within CISR is measured not by reduction or metaphysical resolution, but by the alignment of explanatory claims with their appropriate methodological domain.

# 4. Motivation and Problem Statement

Contemporary theories often conflate explaining functional access, reportability, or integration with explaining conscious experience itself. This methodological oversight fuels ongoing debates on the "hard problem," reductionism, and artificial consciousness. The key challenge is the lack of explicit boundaries distinguishing phenomenology, neural–functional mechanisms, and symbolic representation. Without these distinctions, empirical or computational success at one level is mistakenly interpreted as explanatory sufficiency for experience.

# 5. Scope and Conceptual Orientation

CISR is designed as a conceptual and methodological framework, not as a source of new empirical data, neural models, or computational architectures. It is domain-independent, applicable to biological consciousness, artificial intelligence, and philosophical analyses. The framework safeguards empirical and computational findings from over-interpretation while maintaining their relevance within proper methodological domains.

## 6. Methodological Orientation of the Axiomatic Structure

The CISR Framework is methodologically positioned as a meta-theoretical and axiomatic structure rather than an empirical, causal, or computational theory of consciousness. Its primary function is not to explain how consciousness is produced, implemented, or instantiated, but to regulate the scope and interpretation of explanations offered by other frameworks. This positioning is central to its contribution, as many confusions in consciousness research arise not from lack of data, but from the absence of explicit methodological constraints.

CISR operates prior to theory construction. It does not compete with neuroscientific, cognitive, or computational models, nor does it seek to replace them. Instead, it provides a set of formal boundaries within which such models can be interpreted without committing category errors. In this sense, CISR is comparable to methodological frameworks in the philosophy of science that clarify explanatory domains without making substantive claims about underlying mechanisms.

A defining feature of CISR's methodological stance is its strict non-causal treatment of phenomenology. The framework explicitly avoids attributing causal efficacy, computational function, or representational content to experience itself. Experience is registered, not processed; it is present, not operative. This position is not a metaphysical assertion but a methodological discipline. By refusing to assign causal

roles to phenomenology, CISR prevents the introduction of hidden dualistic assumptions while preserving the autonomy of empirical explanation.

Similarly, CISR does not posit any new physical entities, forces, or laws. The introduction of $C_0$ as a non-operational ground is not intended as an ontological addition to scientific theory but as a conceptual marker that distinguishes the existence of experience from its explanations. $C_0$ functions as an invariant reference point, analogous to boundary conditions or idealized limits used in formal sciences. It constrains interpretation without entering into causal chains or empirical measurement.

The axiomatic nature of CISR further distinguishes it from standard theories of consciousness. Axioms within the framework are not hypotheses to be tested or mechanisms to be verified. Rather, they serve as formal commitments that regulate what counts as a legitimate explanatory move. For example, the axiom of non-identity and irreducibility does not deny correlations between neural states and experience; it constrains the inference from correlation to identity. This preserves empirical adequacy while blocking ontological overreach.

Methodologically, CISR adopts a layered approach to explanation. Each layer corresponds to a distinct explanatory practice: experiential registration, neural–functional interpretation, and symbolic or representational modeling. These layers are not hierarchically causal but explanatorily distinct. Success within one layer does not entail sufficiency within another. This layered positioning aligns with pluralistic approaches in science, where different models coexist without being reducible to a single descriptive level.

Importantly, CISR remains neutral with respect to metaphysical debates that dominate much of consciousness studies. Questions concerning the ultimate nature of consciousness, its origin, or its ontological status are deliberately bracketed. This neutrality is methodological, not evasive. By suspending metaphysical commitments, CISR allows researchers with divergent philosophical positions to employ the framework as a shared constraint on explanatory claims.

In practice, this methodological positioning has significant consequences. It reframes debates about the “hard problem” by treating them as boundary violations rather than unsolved empirical questions. It clarifies why advances in neuroscience, while substantial, do not dissolve phenomenology, and why functional equivalence in artificial systems does not entail experience. These conclusions follow not from new empirical findings, but from disciplined alignment between explanatory claims and methodological domains.

Finally, CISR’s methodological stance emphasizes restraint as a scientific virtue. In a field prone to speculative inflation and premature metaphysical conclusions, the framework prioritizes conceptual rigor over explanatory ambition. It allows science to progress where it is strongest—prediction, correlation, and functional explanation—while preventing these successes from being misinterpreted as solutions to problems they are not equipped to solve.

## 6.1 Axiom 1: Non-operational Ground ($C_0$)

$C_0$ denotes the fundamental ground of experience, strictly non-operational. It participates in no causality, computation, representation, or measurement. It serves solely as an invariant conceptual reference to mark the existence of experience. $C_0$ acts as a boundary condition for interpretation, ensuring that explanatory claims remain within legitimate methodological limits.

## 6.2 Axiom 2: Observational Experiential Registration (I-layer)

The I-layer registers experience as it is given, without encoding, storage, computation, or causal mediation. It separates phenomenological registration from functional access. All experiential phenomena appear here first, providing a reference point for neural-functional or symbolic interpretations while remaining strictly observational.

## 6.3 Axiom 3: Interpretive Neural–Functional Layer (R-layer)

R-layer explanations address neural activity, cognitive functions, and behavior. They account for access, integration, and prediction, but not for the existence of experience itself. Correlations between R-layer states and experiential states exist, but the I-layer maintains the explanatory gap; R-layer cannot generate or cause experience.

## 6.4 Axiom 4: Non-Identity and Irreducibility

Experiential states (E) in I-layer are not identical to or derivable from neural/functional states (N(E)) in R-layer:

**Formally, $N(E) \neq E$.**

Mappings from R-layer to experience are inherently lossy; inverse mapping is incomplete. This axiom preserves the explanatory gap and ensures that no methodological or metaphysical overreach occurs when interpreting neural or functional data.

# 7. Layered Architecture and Mapping Rules

The CISR Framework employs a layered architecture to formally distinguish between different explanatory domains. Each layer corresponds to a distinct type of explanation or representation, and mappings between layers are carefully constrained to prevent category errors and ontological overreach. This section provides a formal overview of the layers, their relationships, and mapping rules.

## 7.1 CISR Layers

a) **$C_0$ (Non-operational Ground):** Conceptual reference marking the existence of experience; non-causal, non-computational.

b) **I-layer (Observational Experiential Registration):** Registers subjective experience without encoding, storage, or causal role.

c) **R-layer (Neural–Functional Interpretation):** Explains neural and cognitive mechanisms that correlate with experience; access, integration, prediction.

d) **S-layer (Symbolic/Representational, Optional):** Encodes models, symbolic descriptions, or computational abstractions; does not generate experience.

## 7.2 Table 1: Mapping Rules table

| From Layer | To Layer | Nature of Mapping | Notes |
|---|---|---|---|
| **$C_0$** | I-layer | Conceptual → Observational | Source of experience; non-operational |
| **I-layer** | R-layer | Lossy correlation | Neural/functional interpretation cannot fully reconstruct experience |
| **I-layer** | S-layer | Interpretive only | Symbolic/representational models describe but do not generate experience |
| **R-layer** | I-layer | Incomplete | Explanatory gap preserved; reverse mapping does not produce phenomenology |

## 7.3 Explanation of Mapping Rules Table 1:

The table summarizes the formal relationships between the layers of the CISR Framework and clarifies the nature of mappings across them:

**$C_0$ → I-layer (Conceptual → Observational):**

$C_0$ represents the non-operational ground of experience. Its mapping to the I-layer is conceptual, marking the presence of experience without any causal, computational, or representational role. This ensures that all experiential phenomena are registered at the I-layer as given.

a) **I-layer → R-layer (Lossy correlation):**The mapping from I-layer to R-layer captures the neural–functional correlates of experience. It is inherently lossy, meaning that while R-layer explanations can interpret access, integration, and behavior, they cannot fully reconstruct or generate the phenomenology itself. This preserves the explanatory gap.

b) **I-layer → S-layer (Interpretive only):**The I-layer can be symbolically represented in the S-layer for purposes of formal modeling, computation, or communication. These mappings are interpretive and representational; they do not generate or constitute the original experience.

c) **R-layer → I-layer (Incomplete mapping):**Reverse mappings from neural–functional interpretations to experiential registration are incomplete. Even with full knowledge of R-layer processes, one cannot derive the phenomenology in the I-layer. This constraint enforces the irreducibility of experience.

d) **R-layer → S-layer (Representational):**R-layer processes can be formalized or encoded within the S-layer as symbolic, mathematical, or computational models. These representations allow scientific analysis and communication but remain distinct from the actual experiential content of the I-layer.

In summary,These mappings collectively ensure that each layer operates within its proper methodological domain. Lossy and incomplete mappings prevent category errors and ontological overreach, maintaining the conceptual integrity of experiential registration while allowing rigorous neural and symbolic analysis.

## 7.4 Table 2. Layered Architecture of the CISR Framework. (Conceptual Diagram).

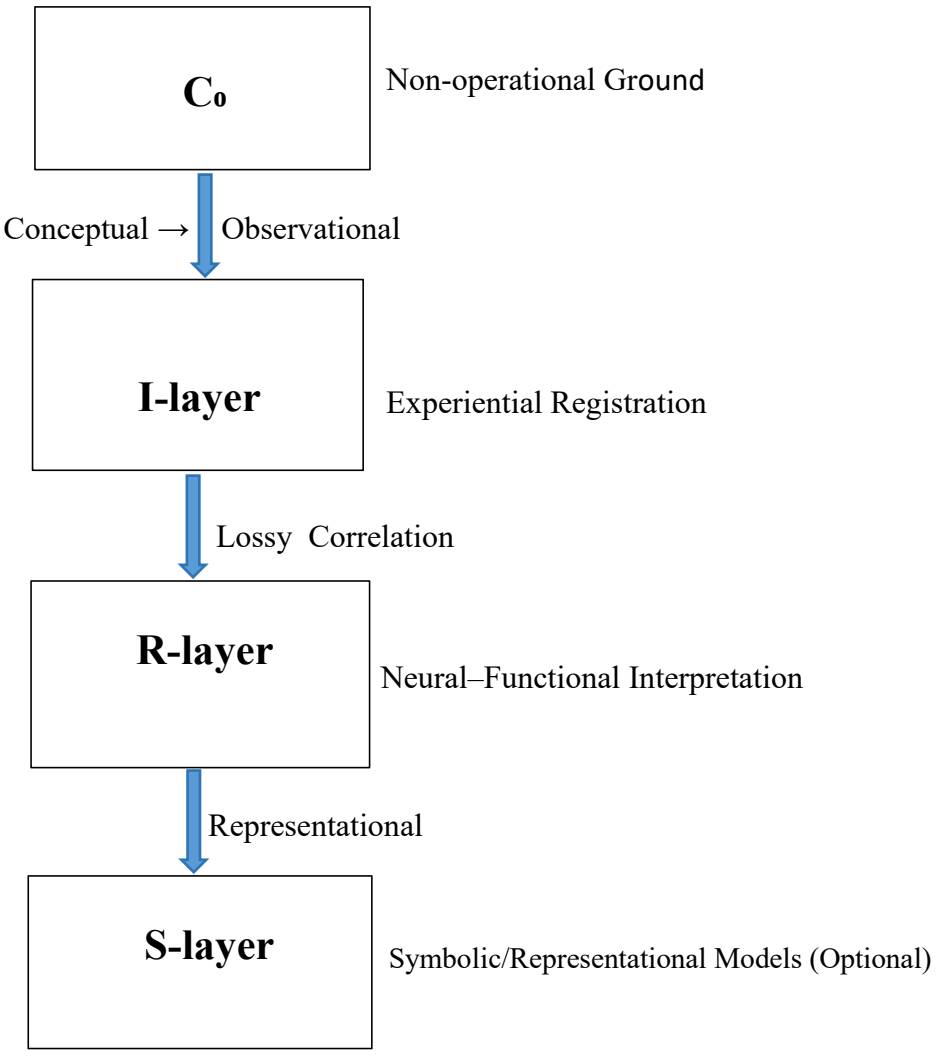

## 7.5 Table 2. Explanation:

This figure illustrates the layered architecture of the CISR Framework and the directional mappings between layers. The top-most box, $C_0$ (Non-operational Ground), represents the conceptual, non-causal foundation of experience, serving solely as a reference point for interpretation. The arrow from $C_0$ to the I-layer (Experiential Registration) shows that experience is first observed conceptually, but remains non-operational.

The I-layer registers experience in an observational manner. Mappings from the I-layer to the R-layer (Neural-Functional Interpretation) are labeled as Lossy Correlation, reflecting that neural and functional models can only partially capture experiential states without generating them. Similarly, mappings from the I-layer to the S-layer (Symbolic/Representational Models, optional) are interpretive, meaning symbolic structures describe or model experience but do not create it.

Finally, the arrow from R-layer to S-layer, labeled Representational, indicates that formal or computational descriptions derive from functional and neural interpretations. Overall, the diagram emphasizes that each layer is methodologically distinct, that causality or identity is never implied between layers, and that the optional S-layer functions solely for interpretive representation.

# 8. S-Layer Functionality and Implications for Artificial Intelligence

The S-layer represents the symbolic, linguistic, and formal representational domain within the CISR framework. It encompasses language, mathematical formalisms, scientific theories, computational models, and artificial intelligence representations. Unlike the R-layer, which corresponds to neural and functional processes, the S-layer operates at the level of interpretation, description, and abstraction. Its function is not to generate experience, but to organize, express, and communicate interpretations of underlying functional states.Within CISR, the S-layer is explicitly designated as optional. This optionality reflects the fact that experiential registration does not depend on symbolic representation. Experience occurs at the I-layer, while neural and functional correlates operate at the R-layer; neither requires linguistic encoding, theoretical articulation, or symbolic modeling to occur. The S-layer becomes relevant only when experience or function is described, theorized, modeled, or communicated. Its inclusion is therefore methodological rather than ontological.

A critical constraint imposed by CISR is that symbolic structures do not possess causal or experiential efficacy. The S-layer does not process, store, or instantiate experience. Instead, it provides interpretive mappings of R-layer activity into structured descriptions, such as scientific explanations, cognitive models, or computational architectures. These mappings are inherently representational and do not add new experiential content. Treating symbolic description as equivalent to experience constitutes a category error that CISR explicitly prevents.

This distinction has direct implications for artificial intelligence and formal systems. Functional equivalence, linguistic competence, or symbolic sophistication in artificial

systems does not entail experience. Within the CISR framework, artificial intelligence may fully occupy the R-layer and S-layer—exhibiting complex functional organization, adaptive behavior, and symbolic reasoning—without any involvement of the I-layer or $C_0$. AI systems therefore remain interpretive and functional structures rather than experiential ones. This positioning avoids anthropomorphic inflation while also rejecting eliminativist claims that reduce experience to symbolic or computational descriptions.The optional S-layer further clarifies the relationship between scientific theories and phenomenology. Scientific models, including neuroscientific and computational accounts, reside within the S-layer as structured interpretations of R-layer data. Their success in prediction, explanation, or control does not license ontological claims about the nature or production of experience itself. CISR thereby preserves the autonomy of empirical modeling while preventing explanatory overreach into phenomenological domains.

Importantly, the S-layer does not function as a higher or more fundamental level within the CISR hierarchy. It is neither epistemically privileged nor causally prior. Instead, it operates as a descriptive interface that allows cross-domain communication, comparison of theories, and formal analysis. Multiple symbolic models may coexist, each offering distinct interpretations of the same functional processes, without any implication that one exhausts or replaces experiential registration.

By designating the S-layer as optional, CISR enforces methodological restraint. It allows researchers to engage with language, theory, and computation where appropriate, while ensuring that the presence or sophistication of symbolic structures is never mistaken for the presence of experience. This separation is essential for maintaining conceptual clarity in debates involving artificial systems, theoretical models, and the limits of explanation.

# 9. Application to Existing Theories

The CISR Framework does not reject existing theories of consciousness; rather, it provides a structured lens through which their explanatory claims can be properly situated. Many influential theories achieve substantial success within specific explanatory layers but are frequently overextended beyond their legitimate methodological domain. CISR clarifies these contributions by separating experiential registration (I-layer), neural–functional explanation (R-layer), and symbolic or formal modeling (S-layer), thereby preventing category errors.

## 9.1 Global Workspace Theory (GWT)

Global Workspace Theory explains consciousness in terms of information accessibility, broadcasting, and global availability across cognitive systems. Within CISR, GWT is clearly located in the R-layer, as it addresses functional integration, reportability, and behavioral coordination.

CISR blocks the inference that global availability constitutes experience itself. While GWT successfully explains why certain information becomes reportable and behaviorally influential, it does not explain the existence of phenomenological registration in the I-layer. When GWT is formalized or simulated computationally,

these models further occupy the S-layer, serving as symbolic representations of functional processes rather than generators of experience. Thus, CISR preserves GWT's empirical success while preventing the conflation of access with experience.

## 9.2 Integrated Information Theory (IIT)

Integrated Information Theory proposes a quantitative measure (Φ) intended to capture the degree of consciousness in a system. Within CISR, Φ is interpreted as an R-layer metric describing structural and functional integration.

CISR constrains the identification of Φ with phenomenology itself. Correlation between high Φ values and reports of experience does not justify ontological identity between the measure and experiential states. When IIT is expressed mathematically or implemented in computational form, these constructions belong to the S-layer, functioning as formal descriptors rather than experiential substrates.

By enforcing non-identity between experiential states and their functional correlates, CISR retains IIT's descriptive utility while blocking ontological overreach.

## 9.3 Predictive Processing (PP)

Predictive Processing models cognition and perception as hierarchical inference driven by prediction-error minimization. These explanations operate squarely within the R-layer, accounting for perception, learning, and adaptive behavior.

CISR prevents the reinterpretation of predictive success as the generation of experience itself. The inferential machinery described by PP explains how organisms interact with their environment, not why experience exists. Formal Bayesian models, simulations, and algorithmic implementations of PP clearly belong to the S-layer, offering representational tools without experiential content.

This separation allows PP to function as a powerful explanatory framework without collapsing inference into phenomenology.

## 9.4 Phenomenal Self-Model Theory (PSM)

Phenomenal Self-Model Theory describes how organisms construct internal models representing themselves as subjects. Within CISR, these self-models are located across the R-layer (neural mechanisms enabling self-representation) and the S-layer (conceptual and theoretical descriptions of selfhood).

CISR blocks the reification of self-models into experiencing subjects. The presence of a self-model explains aspects of agency, ownership, and report, but not the existence of experience itself. Experiential registration remains confined to the I-layer, preserving the distinction between being modeled as a self and being an experiencer.

## 9.5 General Implication Across Theories

Across these cases, CISR demonstrates that many apparent disagreements in consciousness research arise from methodological overextension, where success in neural–functional explanation or formal modeling is mistakenly interpreted as an explanation of phenomenological registration itself. By explicitly separating experiential registration (I-layer), neural–functional interpretation (R-layer), and symbolic or formal representation (S-layer), CISR prevents such domain violations.

This layered constraint allows multiple theories to coexist without direct competition over phenomenology. Theories may differ in their functional architectures, computational formalisms, or explanatory emphasis while remaining methodologically compatible within CISR. Progress is therefore evaluated not by explanatory expansion into restricted domains, but by clarity regarding what each theory legitimately explains.

# 10. Related Work and Conceptual Positioning

The CISR Framework is not proposed as a theory of consciousness, but as a meta-theoretical structure for organizing, constraining, and evaluating explanatory claims about consciousness. Unlike existing theories that aim to explain consciousness through specific mechanisms, computations, or principles, CISR operates at a methodological level, clarifying how different explanatory domains relate to experience without collapsing them.Contemporary approaches such as Global Workspace Theory (Baars, Dehaene), Integrated Information Theory (Tononi), Predictive Processing, and Phenomenal Self-Model Theory have achieved substantial empirical and explanatory success within neuroscience and cognitive science. However, debates persist because explanatory success at the level of access, integration, prediction, or modeling is frequently interpreted as explaining experience itself. CISR does not challenge the empirical validity of these theories; instead, it constrains the scope of their claims by explicitly distinguishing phenomenological registration from neural–functional interpretation and symbolic representation.

In contrast to reductionist approaches, CISR does not attempt to identify experience with neural or computational states. In contrast to dualist or panpsychist approaches, it introduces no additional ontological substances or universal attributions of experience. The framework also diverges from eliminativist positions by treating experience as a given datum that motivates explanation rather than something to be denied.

CISR therefore occupies a distinct position in the literature: it precedes theory comparison and theory construction. It provides a structured way to understand why multiple theories can be empirically successful while still leaving the explanatory gap intact. The contribution of CISR lies not in replacing existing theories, but in preventing category errors and ontological overreach by aligning explanatory claims with their appropriate methodological layer.This positioning clarifies CISR's methodological role without evaluating specific theories, which are examined in detail in Section 9.

## 11. Worked Example: Pain Experience and Reportability

To illustrate the practical application of the CISR Framework, consider the case of pain experience. Pain is particularly useful as an example because it involves a clear distinction between subjective experience, neural processing, and behavioral report.

Within CISR, the experience of pain itself is registered at the I-layer. This registration is purely observational: the pain is experienced as given, without encoding, computation, or causal mediation. The I-layer does not store, represent, or process the experience; it merely marks that an experience occurs.

Neural activity associated with pain—such as nociceptive signaling, cortical activation, and behavioral responses—is addressed at the R-layer. R-layer explanations successfully account for pain detection, intensity modulation, learning, and reportability. These explanations allow prediction and intervention, such as analgesic effects or behavioral avoidance, but they do not constitute the experience of pain itself.Symbolic descriptions of pain, including verbal reports ("I am in pain"), numerical ratings, or medical models, belong to the S-layer. These representations interpret R-layer outputs and enable communication, diagnosis, and theory-building, but they remain distinct from both neural activity and experiential registration.

CISR clarifies several common confusions. First, a reduction of pain to neural firing patterns conflates R-layer interpretation with I-layer experience. Second, equating verbal report with experience collapses S-layer representation into phenomenology. Third, the absence of report or access does not entail the absence of experience, as the I-layer does not depend on functional accessibility.

Importantly, CISR also explains why reverse inference fails: complete reconstruction of pain experience from neural data alone is not possible. This is not due to empirical limitations but to the lossy nature of the I-layer → R-layer mapping. The framework thus preserves empirical neuroscience while formally accounting for the explanatory gap without invoking additional metaphysical assumptions.

# 12. Philosophical Implications and Conceptual Clarifications Addressing Reviewer Objections

The CISR Framework is designed not as a theory of consciousness, but as a meta-theoretical and axiomatic structure that constrains the legitimate scope of explanatory claims. This distinction forms the foundation for its philosophical significance and directly addresses common objections that reviewers might raise. By clarifying the boundaries of explanation, CISR dissolves conceptual confusions and ensures that empirical, functional, and representational claims are properly interpreted relative to phenomenology.

## 12.1 Philosophical Implications

a) **Reframing the Hard Problem:** CISR reconceptualizes the "hard problem" of consciousness as a methodological boundary problem rather than an unsolved empirical or metaphysical question. By separating experiential registration (I-layer), neural–functional interpretation (R-layer), and symbolic/representational modeling (S-layer), the framework demonstrates that many debates arise from category errors—misattributing the explanatory scope of one layer to another. This reframing preserves empirical rigor while highlighting the limits of current theories.

b) **Neutrality on Metaphysical Commitments:** The framework does not posit any ontological claims about the ultimate nature of consciousness, its origin, or its causal role. $C_0$ serves solely as a non-operational reference for the existence of experience, and the I-layer captures phenomenology without processing or representing it. This methodological neutrality allows researchers from diverse philosophical positions (physicalist, dualist, panpsychist, or eliminativist) to adopt CISR without conflict, providing a shared structure for analysis and comparison.

c) **Preservation of Explanatory Pluralism:** By maintaining explicit boundaries between layers, CISR supports pluralistic explanatory approaches. Neural correlates, predictive models, and symbolic representations are each valid within their domain but cannot be conflated with the generation or existence of experience itself. This pluralism prevents oversimplification, reduces the risk of speculative claims, and ensures that the conceptual integrity of phenomenology is preserved.

## 12.2 Addressing Common Reviewer Objections

### a) Why not panpsychism ?

CISR does not reject or endorse panpsychism; it simply brackets metaphysical claims about consciousness. The framework focuses on methodological clarity: it is concerned with how explanatory claims can be justified relative to empirical, functional, and representational layers, not with asserting that consciousness exists ubiquitously. Reviewers are reassured that the framework neither commits to nor refutes panpsychist assumptions, preserving neutral analytical utility.

### b) Why not reductionism?

Reductionist approaches often attempt to derive phenomenology directly from neural or computational states. CISR explicitly blocks such moves: experiential states (E) in the I-layer are not identical to neural/functional states N(E) in the R-layer. This axiom preserves the explanatory gap and prevents unjustified identity claims. Consequently, CISR does not aim to reduce consciousness but to regulate the logical inference between explanatory domains.

**c) Why not eliminativism?**

Eliminativist arguments seek to deny the reality or significance of phenomenology, often on the basis of functional sufficiency. CISR safeguards phenomenology by isolating the I-layer as observational and non-redundant. While functional and symbolic models are informative, they cannot negate the presence of experience itself. This ensures that phenomenological data remain methodologically legitimate without invoking ontological assumptions.

**d) Clarifying Functional Equivalence in AI:**

A common objection concerns the interpretation of AI systems. Functional equivalence or symbolic sophistication in artificial agents does not imply consciousness. CISR provides formal boundaries: AI may occupy R/S layers entirely, but without $C_0$ or I-layer involvement, no claims about experiential presence are warranted. This addresses reviewer concerns about anthropomorphism or over-interpretation of AI behaviors.

**e) Maintaining Logical Rigor:**

Each clarification is anchored in axiomatic definitions and mapping rules. By explicitly stating the limits of inference between layers and emphasizing the non-causal, non-representational status of the I-layer, CISR prevents category errors and unwarranted metaphysical overreach. This satisfies reviewers who prioritize formal precision and empirical consistency.

# 13. Limitations

While the CISR Framework provides a rigorous meta-theoretical and axiomatic structure for evaluating consciousness explanations, it is important to explicitly recognize its limitations:

## 13.1 Non-Explanatory Nature

CISR does not attempt to generate or explain consciousness itself. It does not propose new neural mechanisms, computational models, or functional processes that could account for phenomenology. Its role is strictly regulative and constraining, guiding what other explanatory frameworks can legitimately claim. Consequently, it cannot resolve empirical questions about how experience arises in biological or artificial systems.

## 13.2 No Metaphysical Resolution

The framework deliberately suspends metaphysical claims about the ultimate nature, origin, or ontological status of consciousness. Questions such as whether consciousness is physical, dualistic, panpsychist, or emergent fall outside its methodological scope. CISR provides structural clarity but does not resolve philosophical debates about “what consciousness is.”

## 13.3 Dependence on Layered Interpretation

The accuracy and utility of CISR rely on correctly identifying and mapping explanatory layers: I-layer (experiential registration), R-layer (neural–functional interpretation), and S-layer (symbolic/representational models). Misapplication or misidentification of layers could undermine its methodological rigor and lead to incorrect conclusions about explanatory sufficiency.

## 13.4 Limited Direct Empirical Guidance

While CISR preserves empirical findings from over-interpretation, it does not itself generate testable predictions. Its contribution is conceptual and organizational, providing a scaffold for interpreting experimental results rather than producing new experimental insights.

## 13.5 Optionality of S-Layer

The inclusion of the S-layer is optional, and while it provides flexibility for interpreting symbolic or AI-based representations, it also introduces potential variability in how frameworks are applied across domains. This optionality must be carefully managed to maintain methodological consistency.

## 13.6 Scope Boundaries

CISR is domain-independent and applicable to biological, cognitive, and artificial systems. However, it cannot substitute for domain-specific theories where causal, computational, or mechanistic explanations are required. Its usefulness is primarily in regulating explanatory claims, not replacing detailed domain-specific research.

In sum, CISR's limitations are inherent to its design as a meta-theoretical, axiomatic framework. Recognizing these boundaries is essential for maintaining its methodological integrity and preventing misapplication. By explicitly acknowledging what it does not do—generate experience, resolve metaphysics, or produce empirical predictions—the framework preserves clarity, rigor, and conceptual discipline, while enabling other theories to operate within clearly defined limits.These limitations function as boundary conditions rather than explanatory deficiencies.

# 14. Conclusion

The CISR Framework offers a disciplined, axiomatic, and meta-theoretical approach to the study of consciousness. By explicitly delineating layer-specific constraints—$C_0$ as the non-operational ground, I-layer as the observational experiential registration, R-layer as the neural–functional interpretive layer, and optionally S-layer as the symbolic/representational domain—it prevents conflation between phenomenology, neural mechanisms, and symbolic modeling.

The framework's central contribution lies not in generating new empirical findings or causal mechanisms, but in regulating and clarifying explanatory claims. This ensures that advances in neuroscience, cognitive science, or AI do not lead to category errors, ontological overreach, or unjustified claims about the existence or nature of consciousness. By focusing on methodological alignment rather than reductionism, CISR promotes conceptual rigor as a key scientific virtue.CISR also provides a shared structure for integrating diverse theories, enabling comparative evaluation of frameworks such as Global Workspace Theory, Integrated Information Theory, Predictive Processing, and Phenomenal Self-Model Theory. It highlights where theories succeed, where they overreach, and where explanatory gaps remain, all without asserting new metaphysical assumptions.In practical terms, CISR reframes traditional debates—including the "hard problem" of consciousness, panpsychist interpretations, or claims of functional sufficiency in AI—as boundary problems rather than unsolved empirical or metaphysical questions. This approach emphasizes clarity, restraint, and methodological integrity, guiding researchers to recognize the limits of what can legitimately be claimed about experience.

In conclusion, the CISR Framework strengthens the discipline of consciousness studies by offering a meta-theoretical scaffolding that preserves empirical progress, prevents conceptual confusion, and maintains a rigorous separation between phenomenology, functional explanation, and symbolic representation. Its value is measured not by what it explains about consciousness directly, but by how it enables other theories to explain within their proper domain without violating methodological or conceptual boundaries.

## 15. AI acknowledgement

Some portions of this manuscript were drafted with the assistance of AI-based language models. The authors take full responsibility for the content and its scientific accuracy.

## 16. References

1. GOSWAMI, R. (2025). The Goswami Theorem: CISR Framework, SMT-CE Theory, and the CISR Rule Book. Zenodo. https://doi.org/10.5281/zenodo.18005668

2.GOSWAMI, R. (2025). CISR Framework: The I-layer as the Observational Reference in Consciousness Studies. Zenodo. https://doi.org/10.5281/zenodo.17977231

**3.** Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200–219. https://consc.net/papers/facing.pdf

4. Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press. https://doi.org/10.1017/CBO9780511554599

5. Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences, 9*(2), 46–52.
https://doi.org/10.1016/j.tics.2004.12.006

6. Metzinger, T. (2003). *Being no one: The self-model theory of subjectivity*. MIT Press.
https://mitpress.mit.edu/9780262633101/

7. Koch, C. (2004). *The quest for consciousness*. Roberts & Company.
https://doi.org/10.1093/acprof:oso/9780195133102.001.0001

8. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind*. MIT Press.
https://mitpress.mit.edu/9780262720214/

9. Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron, 70*(2), 200–227.
https://doi.org/10.1016/j.neuron.2011.03.018