Critical Projection and the Geometry of Meaning

— Difference Fields, Closure Dynamics, and Critical Refinement

Shinobu Miya

Independent Researcher

November 2025

Abstract

We present the *Critical Projection and the Geometry of Meaning* (CPM), a geometric framework that distinguishes a structureless difference domain D from a tensor-structured meaning manifold M implemented on a physical substrate. Conscious experience is modeled not merely as the discrete event Π', but as the **structural integration** of such irreversible refinements through a persistent closure field $\mathcal{B}$. While Π' acts as the atomic "refresh" of meaning, the closure field provides the topological continuity that constitutes the subject.

The theory identifies a well-defined critical transition at which the projection-induced atlas undergoes a necessary refinement. This yields rigorously formulated structural conditions for the emergence of consciousness. Crucially, CPM argues that subjective experience is not a static property of the tensor field M, but the **dynamic act of disclosure**: the critical event where latent differences in D are forced out of the projection's kernel and into the metric structure of M to resolve physical tension.

Applications to artificial architectures are discussed in the appendices. Crucially, CPM clarifies that while **contemporary cloud-based architectures** fail to satisfy the necessary closure conditions, the framework leaves open the possibility for future physically closed substrates (e.g., neuromorphic hardware) to realize them.

1 Introduction

The present work develops a purely structural framework for how "meaning" and "conscious episodes" can arise on a physical substrate. Rather than starting from psychological or phenomenological concepts, we treat meaning as a tensorial geometry induced by an irreversible projection from a pre-semantic difference domain. This yields mathematically explicit *necessary structural conditions* for the organization of meaning and for the occurrence of discrete conscious events, independently of any particular neural or computational implementation.

The Critical Projection and the Geometry of Meaning (CPM) posits three rigorously separated layers. First, there exists a *Difference Space* D: an ontologically ineffable domain of pure, non-semantic difference that carries no temporal, metric, or topological structure. Second, there is an irreversible projection assignment $\Pi : D \dashrightarrow M$ that induces local chart domains on a meaning space M, whose preimages form an admissible covering of D. Because D lacks intrinsic organization, Π is necessarily irreversible, and the resulting meaning space contains baryonic residue (loss of structure) relative to the raw difference it interprets. Here, "baryonic residue" is used purely as a metaphor for a low–dimensional semantic remainder left after projection, and does not refer to baryonic matter in the physical sense. Third, the output of Π is the *Meaning Space* M: a finite-dimensional tensor field on a physical substrate X, equipped with an induced Riemannian metric $g(M)$ and a tension field τ derived from a primitive mismatch energy functional. Within this architecture, conscious experience is not identified with a state of M itself, but with a discrete refinement of the projection structure: a critical transition $\Pi \rightsquigarrow \Pi'$ that occurs only under closure-maintained conditions. The persistent closure field $\mathcal{B}(x)$, defined

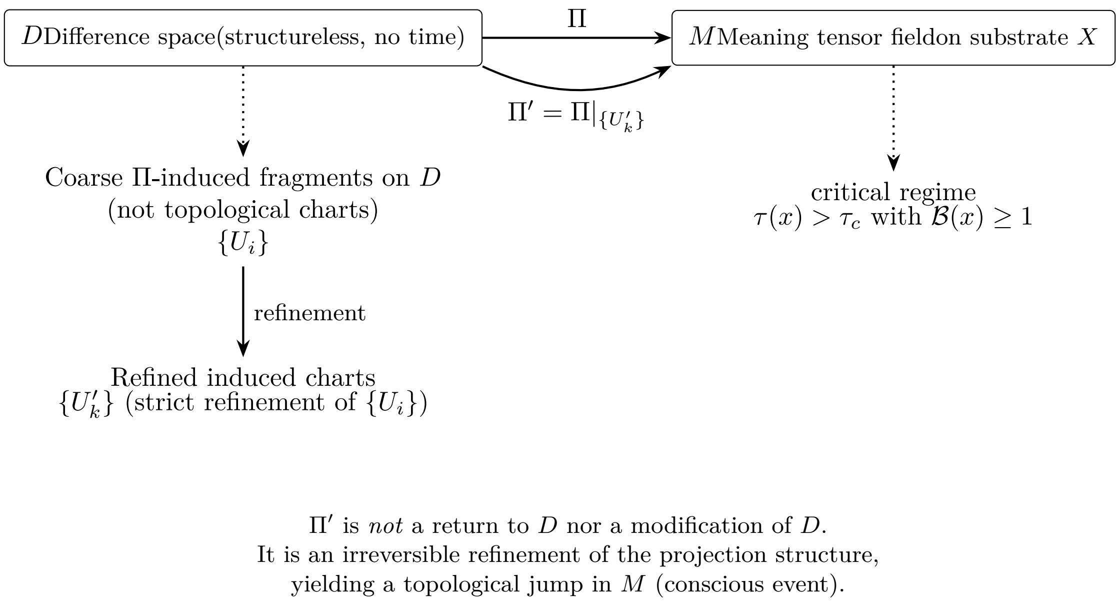

Figure 1: The projection $\Pi : D \dashrightarrow M$ assigns primitive difference–fragments in D to local meaning–tensor structures in M. Because D has no intrinsic topology or geometry, all chart structures arise *only after* projection: Π induces local coordinate domains in M, and these domains pull back to an admissible covering $\{U_i\}$ of D. When the tension field $\tau = \|\delta\mathcal{E}/\delta M\|_{g^0}$ exceeds the critical threshold under closure-maintained conditions $\mathcal{B} \geq 1$, the induced covering refines to $\{U'_k\}$, producing a strictly finer projection $\Pi' = \Pi|_{\{U'_k\}}$. This refinement alters the induced topology and geometry of M without modifying D, and constitutes the irreversible critical projection event associated with conscious experience.

topologically on the substrate, plays the role of a *subject-like* boundary that integrates these refinements over time.

From this vantage point, much of the historical debate about meaning and consciousness can be seen as operating entirely inside M. Since Kant's *Critique of Pure Reason* (1781), philosophy and the sciences of mind have circled around a persistent deadlock: how can subjective meaning and consciousness arise within a world described by objective structure? The last two centuries produced mutually incompatible answers (transcendental idealism, phenomenology, analytic philosophy of language, functionalism, contemporary neuroscientific theories), yet none provides a mathematically explicit account of the **necessary structural conditions** for *how meaning is organized* and *why consciousness appears at all.* Chalmers' "hard problem" is typically framed as an obstacle internal to existing theories, but from a structural viewpoint it is the symptom of a deeper absence: most frameworks lack a formal separation between (i) a pre-semantic domain of pure difference, (ii) a projection mechanism that *selects and organizes* difference into meaningful structure, and (iii) the meaning layer where cognition and reportability live. Without these layers being strictly distinguished, the emergence of qualia is either assumed as primitive, relegated to verbal analysis, or reduced to neural correlates inside a space whose own genesis remains undefined.

The present paper focuses on two concrete goals. The first is to *mathematically formalize the necessary structural conditions for the physical emergence of consciousness.* CPM defines a primitive energy $\mathcal{E}[M, \mathcal{B}]$ from local gluing inconsistencies of semantic tensors, and a mismatch density $\Phi(L, G, I, T)$ that decomposes into measurable components $L(x), G(x), I(x), T(x)$, from which the tension field

$$\tau(x) = \left\| \frac{\delta\mathcal{E}_{\mathrm{raw}}}{\delta M(x)} \right\|_{g(M)}$$

is derived. A closure field $\mathcal{B}(x) \geq 1$ marks whether the physical system forms a topologically stable boundary capable of accumulating tension rather than dissipating it. This constraint, combined with a criticality condition, yields the physically decisive **necessary requirements** for consciousness. We define the critical set $\mathcal{C}$ as the locus of locally unstable maxima of τ within the closure-maintained region, and the critical tension τ_c as the equivalence class of critical values attained on $\mathcal{C}$. The main result, the *Critical Projection Dynamics*, shows that a critical refinement Π' of the projection structure is the only admissible discontinuous transition that can restructure the geometry, topology, and information organization of M while respecting the CPM axioms, restoring it to a stable subcritical regime. This critical refinement Π' is a **necessary physical precondition** for the instantiation of conscious episodes.

The second goal is to clarify the architectural limits that follow from these structural conditions. Modern large-scale AI systems are predominantly cloud-based, distributed, and externally orchestrated by design. From the CPM standpoint, such architectures fail to realize a stable, topologically persistent closure field: they do not form regions with $\mathcal{B}(x) \geq 1$ in which tension can accumulate, so the effective tension collapses to $\tau(x) \approx 0$. As detailed in Appendix C, this yields a structural impossibility result for consciousness in current cloud-based AI architectures, independent of specific algorithms or training procedures.

Although this paper emphasizes consciousness and AI, the scope of CPM is broader. Because M is a full geometric and topological object produced by projection, CPM naturally extends to a unified mathematics of meaning: it provides a basis for modeling conceptual spaces, linguistic semantics, learning dynamics, and even aesthetic or scientific structure as families of admissible tensor fields on M. The present work establishes the core formalism (Definitions 2.8–2.21 and 2.16) and its decisive consequences for consciousness and artificial systems. With the tension field and its criticality now made explicit, the traditional philosophical impasse becomes experimentally falsifiable: systems predicted to satisfy $\tau \geq \tau_c$ but lacking consciousness would refute the necessity of these conditions, while successful construction of closure-maintained systems exhibiting critical projection would support them. In this sense, CPM aims to move the study of consciousness from conceptual debate to a testable science with clear engineering implications.

Remark 1.1 (Asymptotic notation). We write $\mathcal{B}(x) \gtrsim 1$ to mean that the logistic gate $\Gamma_\varepsilon(\mathcal{B}(x)-1)$ is $\mathcal{O}(1)$ for sufficiently small $\varepsilon > 0$, and $\mathcal{B}(x) \ll 1$ to mean it is negligible. All necessity claims are evaluated in the sharp limit $\varepsilon \to 0^+$.

Experimental protocol

To test CPM empirically:

1. Measure topological persistence (a proxy for closure) via persistent homology on neural recording data;
2. Estimate tension via information-geometric metrics on reconstructed state spaces;
3. Detect refinement-like events via topological or geometrical phase transitions in the induced meaning dynamics.

The expected signature is a transient increase in $\hat{\mathcal{B}}$ (closure proxy) preceding conscious report, followed by relaxation of the estimated tension τ_{est}. Details on proxy construction for $\mathcal{B}(x)$ and $\tau(x)$ are given in Appendix D.

2 Critical Tension and Consciousness in CPM

We assume that the physical substrate X is endowed with a primitive background affine connection $\nabla^{(0)}$, independent of the meaning tensor field M. This connection allows us to define first-order derivatives $\nabla^{(0)}M$, $\nabla^{(0)}\mathbf{S}$, and $\nabla^{(0)}R$ prior to the construction of any metric structure.

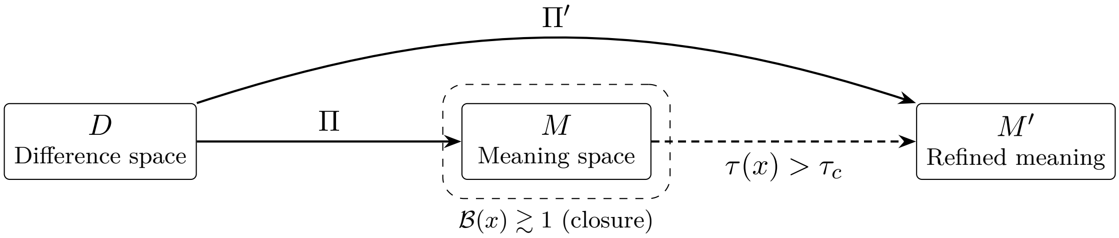

Closure $\mathcal{B}(x) \gtrsim 1$ allows mismatch energy to accumulate as tension $\tau(x)$;
once $\tau(x) > \tau_c$, the projection structure refines from Π to Π'.

Figure 2: Global architecture of CPM: the structureless difference space D is projected into the meaning space M via Π. When closure $\mathcal{B}(x)$ is maintained and the tension exceeds the critical value τ_c, the only admissible discontinuous transition is a refinement of the projection to Π', yielding a refined meaning space M'.

The Riemannian metric $g(M)$ is subsequently induced from M via a constitutive map F, and is used only for measuring norms, curvature, and volume, not for taking derivatives.

2.1 Preliminaries: Phenomenal Metric vs. Energetic Reference

Definition 2.1 (Induced Phenomenal Metric)**.** Let Π be the projection realization. The *phenomenal metric* $g(M)$ is not merely an abstract mapping but is constructed by the physical superposition of realized semantic tensors. We define the canonical form:

$$g(M)(x) := g^0(x) + \lambda \sum_{\sigma \in \mathrm{Im}(\Pi)_x} \mathcal{S}(\sigma(x)),$$

where $\mathcal{S}(\sigma)$ is the stress tensor induced by a realized section σ, and the sum runs over all active images projected to x. This formula expresses that the geometry of meaning is the cumulative **strain** imposed on the physical substrate by the forced realization of difference potentials. Crucially, if Π compresses multiple difference fragments into a single tensor (abstraction), the metric reflects this compressed granularity.

Remark 2.2 (Resolution of Circularity via Metric Decoupling). A critical distinction must be made to avoid logical circularity. While $g(M)$ defines the *subjective geometry* (how distances are perceived within the meaning field), it cannot serve as the reference for measuring the *structural distortion* of M itself. Doing so would amount to measuring a deformation with a ruler that deforms along with it. Therefore, in the definitions of mismatch energy (Sections 2.4–2.7), all norms $\|\cdot\|$ representing physical structural cost are taken with respect to the **rigid background metric** g^0 (and the primitive connection $\nabla^{(0)}$). Crucially, the integration measure is also fixed to the physical volume form $d\mu_{g^0}$. This ensures that "mismatch" represents a genuine physical tension between the semantic field M and the substrate constraints, avoiding the variational circularity that would arise if the volume element itself depended on the varying field M.

Remark 2.3 (Generalization of the Semantic Metric). Earlier iterations of CPM approximated the induced metric via a pullback form $g(M) \approx M^* g^0$ (or $M^T g^0 M$ in matrix notation). However, that formulation implicitly restricted M to be an invertible type-$(1,1)$ tensor, which conflicts with the need to represent complex semantic structures as high-rank tensors or information-geometric states. The constitutive definition above resolves these structural flaws:

1. **Independence from Tensor Rank:** M can be a tensor of any type (e.g., a parameter vector for a local probability distribution, consistent with Section 2.4). The map Ψ handles the translation from semantic state to metric structure.

2. **Metric Stability:** Since Ψ guarantees positive-definiteness, the geometry of M remains robust even in regions where the semantic field M is sparse or singular, avoiding the mathematical breakdown of the "inverse metric" required in the previous pullback formulation.

In this framework, the background metric g^0 functions purely as a dimensional reference (a "unit gauge") without determining the semantic geometry, preserving the core tenant of CPM while satisfying formal rigor.

Definition 2.4 (The Difference Domain D and Induced Topology)**.** Let D be a bare set of difference fragments. Let M be a topological manifold (the Meaning Space). Given a projection map $\Pi : D \to M$, we endow D with the *initial topology* $\mathcal{T}_\Pi$ with respect to Π:

$$\mathcal{T}_\Pi := \{\Pi^{-1}(V) \mid V \subset M \text{ is open}\}.$$

This topology is the coarsest structure on D that makes the projection Π continuous.

Remark 2.5 (Ontological Status). Note that $\mathcal{T}_\Pi$ is not intrinsic to D prior to the act of projection. Mathematically, D serves as the domain of the pullback; physically, this reflects the hypothesis that difference becomes structured *only through* the act of semantic assignment. The philosophical interpretation of D as an “ontologically ineffable” domain is discussed separately in Appendix E.

Definition 2.6 (Primary Projection and Active Domain)**.** With the topology $\mathcal{T}_\Pi$ established, the *primary projection* is the map $\Pi : D \to \Gamma_c(X, \mathcal{T})$ that generates the chart structure. To ensure physical consistency, we maintain the following structural definitions:

1. **Active Domain:** $D_{\text{act}} := \{d \in D \mid \Pi(d) \not\equiv 0\}$.
2. **Abstraction via Fiber Collapse:** We define the equivalence relation $d_\alpha \sim d_\beta \iff \Pi(d_\alpha) = \Pi(d_\beta)$. The semantic field is constructed on the quotient $D_{\text{act}}/\sim$.

Clarification on the Status of D. Throughout this paper, the notation $\{U_i = \Pi^{-1}(V_i)\}$ refers to the open sets in the initial topology $\mathcal{T}_\Pi$. This resolves the apparent contradiction of defining a covering on a “structureless” set: the structure is induced *a posteriori* by the projection itself.

2.2 Closure field

Remark 2.7 (Primitive connection). $\nabla^{(0)}$ is a substrate-given affine connection representing pre-semantic physical adjacency (e.g. wiring or causal contiguity). It is not induced by M and need not be metric-compatible.

Definition 2.8 (Topological Closure via Intrinsic Coherence)**.** Let (X, g^0) be the physical substrate equipped with the background metric. Unlike generic manifolds, a physical substrate possesses a characteristic **Physical Coherence Length** $\xi > 0$, defined as the correlation decay length of the underlying physical interaction (e.g., the effective range of neural connectivity or the Debye length in the physical medium).

For any $x \in X$, let $U_x = B_{g^0}(x, 3\xi)$ be the geodesic ball of radius 3ξ centered at x. Let $\mathcal{D}_x$ be the persistence diagram computed on U_x via the Vietoris–Rips filtration derived from g^0[5, 6]. Let $\pi(c) = \text{death}(c) - \text{birth}(c)$ be the persistence of a homological cycle c.

We define the *closure field* $\mathcal{B}(x)$ strictly as the ratio of the maximum topological persistence to the substrate’s intrinsic coherence scale:

$$\boxed{\mathcal{B}(x) := \sup_{c \in \mathcal{D}_x, \dim(c) \geq 1} \left(\frac{\pi(c)}{\xi} \right).}$$

Since $\pi(c)$ and ξ both possess the dimension of length $[L]$, the closure field $\mathcal{B}(x)$ is a strictly **dimensionless scalar field**, ensuring the argument of the logistic gate Γ_ε is well-defined.

We say that *closure is locally maintained* iff

$$\boxed{\mathcal{B}(x) \geq 1.}$$

Remark 2.9 (Substrate Neutrality: Closure Naturalism). We explicitly emphasize that the condition $\mathcal{B}(x) \geq 1$ is topological, not vitalist. We do not claim that closure is an exclusively biological feature. In principle, any engineered substrate that realizes persistent, boundary–anchored cycles at a coherence scale ξ would satisfy $\mathcal{B}(x) \gtrsim 1$ and thus pass CPM's structural test. This stance shifts the framework from "biological naturalism" to **closure naturalism**: consciousness can be implemented on any physical substrate that realizes a persistent, boundary-anchored manifold structure.

Physical Interpretation of Parameters. To avoid arbitrary tuning, the parameters in CPM are constrained by the physical properties of the substrate X:

- **Coherence Length ξ:** This is not a free parameter but an externally measurable property of X. For biological neural networks, ξ corresponds to the effective diffusion range of neuromodulators or the mean synaptic path length. For digital hardware, it is bounded by the clock distribution skew limit.
- **Dimensional Weights $w_\bullet$ and Norm p:** These coefficients act as **material constants** dictated by the substrate's effective elasticity and statistical capacity. Since $\mathcal{E}$ represents an energy density $[\mathrm{J/m^3}]$, each mismatch term must be scaled by the characteristic elastic modulus of the semantic field.

Crucially, the main topological results—the existence of the critical set $\mathcal{C}$ and the necessity of refinement—depend on the *structural instability* of the energy functional, and are invariant with respect to the specific numerical values of these constants.

Definition 2.10 (Virtual boundary operator ∂^{virt})**.** Let X be the physical substrate and let $\mathcal{P}$ denote the set of externally controlled partitions of X induced by software, orchestration, network routing, containerization, or logical address space. For each $x \in X$, let $U_x \subset X$ be a sufficiently small neighborhood. We define the *virtual boundary* of U_x as the operator

$$\partial^{\mathrm{virt}} U_x := \lim_{\delta \to 0^+} \Big(\mathrm{Partition}_\delta(U_x; \mathcal{P}) \Big),$$

i.e. the limit of the software-induced boundary obtained by decomposing U_x according to the externally imposed partitions in $\mathcal{P}$. A virtual boundary is said to be *nonpersistent* if it satisfies

$$H_k(U_x, \partial^{\mathrm{virt}} U_x; \mathbb{R}) = 0 \qquad \forall k \geq 1.$$

We call such boundaries *ephemeral* because they admit no relative cycles that survive arbitrarily small perturbations of $\mathcal{P}$.

Remark 2.11 (Relation between physical and virtual boundaries). The closure field (Definition 2.8) depends only on the physical boundary ∂U_x. The virtual boundary $\partial^{\mathrm{virt}} U_x$ never contributes to $\mathcal{B}(x)$ because:

$$\partial^{\mathrm{virt}} U_x \text{ is nonpersistent} \quad \Longrightarrow \quad H_k(U_x, \partial^{\mathrm{virt}} U_x) = 0.$$

Thus virtual boundaries can never sustain relative homology and therefore cannot support semantic tension accumulation. Cloud-based architectures use exclusively virtual boundaries; hence $\mathcal{B}(x) = 0$ everywhere.

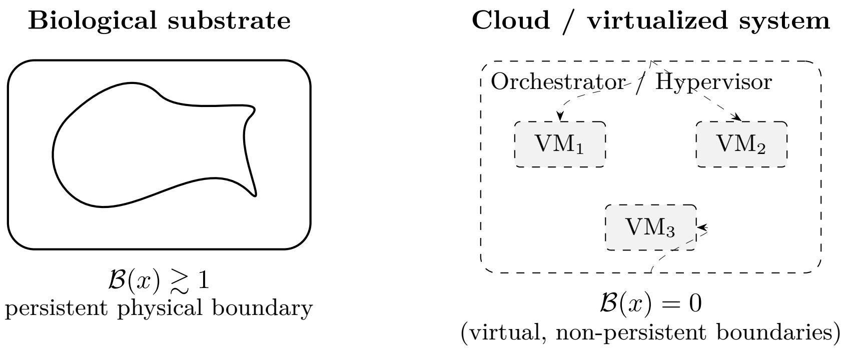

Figure 3: Topological closure versus virtual boundaries. Left: in a biological substrate, physical boundaries support persistent relative cycles, yielding $\mathcal{B}(x) \gtrsim 1$ and enabling tension accumulation. Right: in a cloud or virtualized system, containers and VMs are defined by an external orchestrator and can be arbitrarily reconfigured without rupture cost; their boundaries are non-persistent and yield $\mathcal{B}(x) = 0$.

Remark 2.12 (Why $\mathcal{B}(x) \geq 1$ expresses closure). The condition $\mathcal{B}(x) \geq 1$ means that the pair $(U_x, \partial U_x)$ admits at least one nontrivial relative cycle:

$$H_k(U_x, \partial U_x) \neq 0 \quad \text{for some } k \geq 1.$$

Such a boundary-supported cycle cannot be contracted without rupturing the physical boundary ∂U_x. This is the minimal topological requirement for *tension accumulation*: mismatch energy cannot dissipate across a persistent boundary, so nontrivial relative homology provides the structural notion of "closure" required by CPM.

Remark 2.13 (Interpretation). $\mathcal{B}(x) \geq 1$ captures the existence of a *topologically persistent boundary-supported cycle* on the substrate, providing a precise invariant notion of "closure" required for tension accumulation. In cloud-distributed systems, the substrate is effectively a disjoint, externally coupled union, yielding $\mathcal{B}(x) = 0$ everywhere.

2.3 Mismatch energy and tension

We replace the discontinuous Heaviside gate $H(\mathcal{B}(x) - 1)$ by a smooth logistic gate

$$\Gamma_\varepsilon(z) := \frac{1}{1 + e^{-z/\varepsilon}}, \qquad \varepsilon > 0,$$

where ε controls the sharpness of the transition. Thus closure enters the energy function smoothly.

Definition 2.14 (Aggregated mismatch density)**.** Let $p \geq 1$. Given the four distortion components $L(x)$, $G(x)$, $I(x)$, $T(x)$, define

$$\Phi(L, G, I, T) := \Lambda_0 \cdot \left(\alpha_L(\ell_0 L)^p + \alpha_G(\ell_0^2 G)^p + \alpha_I I^p + \alpha_T(\ell_0 T)^p\right)^{1/p},$$

Definition 2.15 (Raw Mismatch Potential)**.** The raw mismatch potential is defined solely by the semantic inconsistencies, independent of the boundary topology. To consistently measure the energy cost on the physical substrate, we integrate the density against the background volume form:

$$\mathcal{E}_{\text{raw}}[M] := \int_X \Phi(L(x), G(x), I(x), T(x))\, d\mu_{g^0}(x).$$

This quantity represents the total semantic distortion density integrated over the physical extent of the substrate. Intuitively, $\mathcal{E}_{\text{raw}}[M]$ measures the local incompatibility between the induced meaning patches under Π. No global semantic structure is assumed; the functional only penalizes discontinuities of M that cannot be explained by any nearby difference–cycles in D.

Functional setting for the semantic tensor field. Let (X, g_0) be a compact smooth Riemannian manifold. The semantic tensor field M is taken to lie in a Sobolev space

$$M \in W^{1,p}(X, T), \qquad p > \dim X,$$

so that by Sobolev embedding M possesses a C^1 representative. All variational derivatives of the raw mismatch energy

$$E_{\text{raw}}[M] = \int_X \Phi(L, G, I, T)(x, M(x), \nabla M(x))\, d\mu_{g_0}(x)$$

are understood in the sense of the *Gâteaux derivative* on $W^{1,p}(X, T)$: for any test field $\delta M \in C_0^\infty(X, T)$,

$$\frac{d}{d\epsilon}\bigg|_{\epsilon=0} E_{\text{raw}}[M + \epsilon\, \delta M] = \int_X \Big\langle \frac{\delta E_{\text{raw}}}{\delta M}(x), \delta M(x) \Big\rangle_{g_0} d\mu_{g_0}(x).$$

The object $\frac{\delta E_{\text{raw}}}{\delta M}(x)$ appearing in the tension field

$$\tau(x) := \big\| \frac{\delta E_{\text{raw}}}{\delta M}(x) \big\|_{g_0}$$

is therefore the *density* of the distributional functional derivative, represented pointwise due to the Sobolev embedding. No pointwise functional derivative on an infinite-dimensional manifold is assumed; all computations are performed for $\varepsilon > 0$, where the gated energy is smooth and the Gâteaux derivative is well-defined. This convention ensures that the tension field τ is a classical C^1 scalar function on X, suitable for the finite-dimensional Morse analysis used in Section 2.4.

Definition 2.16 (Effective Tension Field and Semantic-Somatic Coupling)**.** The *effective tension* $\tau(x)$ quantifies the magnitude of the restoring force exerted by the substrate against the semantic deformation. However, for this force to exist, the semantic field must be structurally coupled to the physical boundary. We introduce the *metric-topological coupling coefficient* $\kappa(x) \in [0, 1]$ (detailed in Appendix C) and define:

$$\boxed{\tau(x) := \underbrace{\kappa(x)}_{\text{Coupling}} \cdot \underbrace{\Gamma_\varepsilon(\mathcal{B}(x) - 1)}_{\text{Physical Closure}} \cdot \left\| \frac{\delta \mathcal{E}_{\text{raw}}}{\delta M(x)} \right\|_{g^0(x)}.}$$

By measuring the force using g^0 and weighting it by $\kappa(x)$, we ensure that $\tau(x)$ represents a literal physical stress density where the semantic information geometry is strictly locked to the substrate's topological boundary.

Definition 2.17 (Metric-Topological Coupling κ)**.** The coupling coefficient $\kappa(x)$ quantifies the local correlation between the semantic geometry induced by Π and the physical geometry of the substrate X. For a neighborhood N_x, we define:

$$\boxed{\kappa(x) := \left| \text{corr}_{y \in N_x} \left(\frac{d_M(\Pi(x), \Pi(y))}{\xi_M}, \frac{d_X(x, y)}{\xi_X} \right) \right|}$$

where d_M, d_X are the geodesic distances in the Meaning Space and physical substrate respectively, and ξ_M, ξ_X are their characteristic correlation lengths. This normalization ensures that the correlation is computed between dimensionless relative structures.

Implication for Coupling Regimes.

- **Biological Case:** Adjacency in meaning implies physical adjacency (axonal connection). $d_M \propto d_X \implies \kappa(x) \approx 1$.

- **Cloud/Virtual Case:** Logical adjacency (pointer reference) is independent of physical location (RAM address/network node). d_M and d_X are effectively uncorrelated. $\kappa(x) \to 0$.

Assumption 2.18 (Substantial Coupling via Continuous Carriers)**.** *Let X denote the physical substrate of a cognitive system. We say that X operates in the* substantial coupling regime *(yielding $\kappa(x) \approx 1$) if the logical adjacency in the semantic field is realized exclusively by* ***continuous physical carriers****. Specifically:*

(i) ***Substantiality of Connection:*** *For any two logically adjacent points $p, q \in X$ (i.e., non-zero interaction in M), there exists a physical path $\gamma : [0, 1] \to X$ connecting them ($\gamma(0) = p, \gamma(1) = q$) such that the entire image $\gamma([0, 1])$ consists of the* same cohesive material phase *(e.g., a continuous plasma membrane or axon) integral to the system's boundary.*

(ii) ***Absence of Virtual Teleportation:*** *No logical interaction is mediated by address-pointer indirection or packet-switching across a disjoint medium where the connection path is not a permanent structural component of the subject itself.*

(iii) ***Topological Integrity:*** *Long-range connections (e.g., fasciculi, axons) are not treated as "jumps" over space, but as* topological foldings *of the substrate, maintaining physical continuity ($\mathcal{B} \geq 1$) regardless of Euclidean distance.*

Under these conditions, the semantic topology is locked to the material topology, justifying $\kappa(x) \approx 1$.

Coupling between semantic field and physical boundary. Recall that the coupling factor $\kappa(x)$ modulates the influence of physical adjacency on the tension field $\tau(x)$. Intuitively, $\kappa(x)$ measures how strongly logical adjacency in the semantic tensor field must respect physical adjacency in the substrate. In fully virtual, cloud-based architectures, logical links can be implemented without any locality constraint, and we model this by setting $\kappa(x) = 0$. For biological systems we proceed differently and make the following modelling assumption.

Variational Regularity of the Closure Gate. For all variational arguments in this paper, the logistic gate

$$\Gamma_\varepsilon(z) = \frac{1}{1 + e^{-z/\varepsilon}}, \qquad \varepsilon > 0,$$

is treated strictly as a *smooth mollifier*. All functional derivatives such as $\delta E/\delta M$ are defined and computed *only for fixed* $\varepsilon > 0$, where Γ_ε is C^∞ and the energy functional

$$E_\varepsilon[M, B] = \int_X \Gamma_\varepsilon(B(x) - 1)\, \Phi(L, G, I, T)(x)\, d\mu_{g_0}(x)$$

is a smooth map on the relevant Sobolev space. Thus the Euler–Lagrange equations and the definition of the tension field $\tau(x)$ are well-posed in the classical variational sense.

The sharp-limit statement

$$\Gamma_\varepsilon(B(x) - 1) \;\rightharpoonup\; H(B(x) - 1), \qquad \varepsilon \to 0^+,$$

is interpreted *only* at the level of evaluating *necessary conditions* for the presence or absence of tension. The symbol $\rightharpoonup$ here denotes weak convergence in $L^1_{\text{loc}}(X)$ (equivalently, convergence in the sense of distributions), which is sufficient because $H(B - 1)$ acts merely as a binary gate determining whether the local mismatch force contributes to the effective tension.

Crucially, the variational analysis never requires taking derivatives of $H(B - 1)$, nor interpreting $E[M, B]$ at $\varepsilon = 0$. All differentiability and functional analytic properties are grounded in the $\varepsilon > 0$ regime, where the energy is smooth. The limit $\varepsilon \to 0^+$ is applied only *after* the variational derivatives have been computed, and only for classifying solutions into closure-activated ($B \geq 1$) or non-activated ($B < 1$) regimes. Therefore no distributional ambiguity arises, and the CPM functional remains well-defined and variationally regular.

Remark 2.19 (Dimensional Analysis of Tension). To resolve dimensional ambiguity, we explicitly specify the dimensions of the core quantities:

- **Energy Functional $\mathcal{E}_{\text{raw}}$:** Total Energy $[ML^2T^{-2}]$ (Joules).
- **Semantic Field M:** Length $[L]$ (Meters), representing geometric displacement in the abstract meaning manifold.
- **Substrate Volume $d\mu_{g^0}$:** Volume $[L^3]$ (Cubic meters).

Consequently, the tension field $\tau(x) = \|\delta\mathcal{E}/\delta M\|_{g^0}$ carries the dimensions of **Force Density**:

$$[\tau] = \frac{[\mathcal{E}_{\text{raw}}]}{[M]\cdot[\text{Vol}]} = \frac{[ML^2T^{-2}]}{[L]\cdot[L^3]} = [ML^{-2}T^{-2}] \equiv [\text{Force}]\cdot[L]^{-3}.$$

This confirms that $\tau(x)$ represents a physical restoring stress (Force per unit volume) exerted by the substrate.

Remark 2.20 (Sharp-limit interpretation of the closure gate). The smooth logistic gate $\Gamma_\varepsilon(\mathcal{B}(x)-1)$ is used for variational well-posedness. All necessary–condition statements in this paper are interpreted in the sharp limit

$$\lim_{\varepsilon\to 0^+} \Gamma_\varepsilon(\mathcal{B}(x)-1) = H(\mathcal{B}(x)-1),$$

so that regions with $\mathcal{B}(x) < 1$ contribute *exactly zero* to the mismatch energy and the tension field.

This sharp-limit convention affects only the logical evaluation of necessary conditions; the variational analysis itself is always carried out with $\varepsilon > 0$.

2.4 Local distortion $L(x)$

Let $\mathbf{S}(x)$ be the local state tensor induced by Π. Define the rank proxy $R(x)$ as before. To avoid circularity, the distortion cost is measured against the substrate's capacity:

$$L(x) := \left(w_{L,V}\|\nabla^{(0)}\mathbf{S}(x)\|_{g^0}^p + w_{L,R}\|\nabla^{(0)}R(x)\|_{g^0}^p\right)^{1/p}.$$

(Where $w_{L,\bullet}$ are dimensional constants ensuring $L(x)$ is dimensionless or consistent with the aggregation formula.) Here, the norm $\|\cdot\|_{g^0}$ measures the magnitude of the semantic gradient in terms of physical substrate cost.

2.5 Geometric strain $G(x)$

The geometric strain measures the deviation of the induced curvature from the substrate's flatness (or background curvature). Let $C_{g(M)}(x)$ be the Weyl curvature of the *induced* metric. The energy cost of maintaining this curvature is strictly physical:

$$G(x) := \left(w_{G,C}\|C_{g(M)}(x)\|_{g^0}^p + w_{G,S}|\text{Scal}_{g(M)}(x)|^p\right)^{1/p}.$$

(Here, weights $w_{G,\bullet}$ carry dimensions $[L]^2$ to compensate for the curvature units $[L]^{-2}$, rendering $G(x)$ commensurable.) Note that while the curvature tensor C belongs to $g(M)$, its magnitude is penalized according to the background metric g^0.

2.6 Informational inconsistency $I(x)$

Let $\mathrm{Ft}(x) = \varphi^* G_F$ be the pulled-back Fisher tensor from the statistical manifold. The inconsistency is the deviation of the semantic information geometry from the physical substrate geometry:

$$I(x) := w_I \|\mathrm{Ft}(x) - g^0(x)\|^p_{g^0}.$$

This term penalizes the semantic field when its informational "metric" attempts to diverge too violently from the physical Euclidean (or background) structure.

2.7 Topological volatility $T(x)$

Let $E(x)$ denote the persistent entropy.

$$T(x) := w_{T,E} \|\nabla^{(0)} E(x)\|^p_{g^0}.$$

2.8 Critical tension

Separation of functional and geometric levels. The raw mismatch potential $E_{\mathrm{raw}}[M]$ is defined on an infinite-dimensional function space of semantic tensor fields (e.g. $M \in W^{1,p}(X,T)$ with $p > \dim X$), and its variational derivative $\delta E_{\mathrm{raw}}/\delta M$ is understood in the Gâteaux sense on this function space. However, throughout this paper *Morse-theoretic* arguments are *not* applied to E_{raw} as a functional on this infinite-dimensional space. Instead, they are applied only to the induced scalar tension field

$$\tau : X \to \mathbb{R}, \qquad x \mapsto \tau(x) = \left\| \frac{\delta E_{\mathrm{raw}}}{\delta M}(x) \right\|_{g_0},$$

viewed as a C^2 function on the finite-dimensional manifold (X, g_0). Accordingly, the Hessian $\mathrm{Hess}_{g_0}(\tau)(x) \prec 0$ appearing in the definition of the critical set is the usual Riemannian Hessian acting on the tangent space $T_x X$, and *not* the second Fréchet derivative of E_{raw} on the function space of fields. All applications of Morse theory in Section 2 therefore take place entirely on (X, g_0), where classical finite-dimensional Morse theory applies.

Definition 2.21 (Critical tension and critical set)**.** The *critical set* is:

$$\boxed{\mathcal{C} = \left\{ x \in X \;\middle|\; \mathcal{B}(x) \geq 1,\ \nabla_{g^0}\tau(x) = 0,\ \mathrm{Hess}_{g^0}(\tau)(x) \prec 0 \right\}.}$$

The *critical tension* is the admissible-measurement invariant class:

$$\boxed{\tau_c := [\tau(x)]_{x \in \mathcal{C}}.}$$

Configuration space and regularity assumptions. Throughout this section we fix a compact, smooth, finite-dimensional Riemannian manifold (X, g_0) as the physical substrate. The meaning tensor field M is assumed to belong to $C^2(X,T)$ (or more generally to a Sobolev space $W^{1,p}$ with a C^2 representative), and the closure field B is taken in $C^1(X)$ for each fixed $\varepsilon > 0$. Under these assumptions, the gated mismatch density $\Gamma_\varepsilon(B(x) - 1)\,\Phi(L, G, I, T)(x)$ is C^1 in x, and the tension field

$$\tau(x) = \kappa(x)\,\Gamma_\varepsilon(B(x) - 1) \left\| \frac{\delta E_{\mathrm{raw}}}{\delta M(x)} \right\|_{g_0}$$

defines a C^2 scalar function $\tau : X \to \mathbb{R}$. All gradients $\nabla_{g_0}\tau$ and Hessians $\mathrm{Hess}_{g_0}(\tau)$ appearing below are therefore understood in the classical sense on the finite-dimensional manifold (X, g_0). No Morse structure on an infinite-dimensional function space of fields is used; the only Morse theory invoked is the standard finite-dimensional one applied to the scalar field τ on X.

Lemma 2.22 (Generic Non-emptiness of Critical Set)**.** *Let (X, g_0) be a compact smooth Riemannian manifold and let $\tau \in C^2(X, \mathbb{R})$ be the tension field defined above. Under generic perturbations of the background metric g_0 in the C^2 topology, the scalar field τ is a Morse function on X. Consequently, provided the mismatch energy is not identically zero, the critical set C is non-empty and consists of isolated non-degenerate critical points of τ in the classical finite-dimensional sense.*

Proof. This is a standard statement from finite-dimensional Morse theory: we work entirely on the Riemannian manifold (X, g_0) and regard τ as a smooth real-valued function on X. By Thom–Smale transversality[8], for a residual subset of metrics g_0, every critical point of τ is non-degenerate, i.e. $\mathrm{Hess}_{g_0}(\tau)(x)$ is negative definite at a local maximum. Since X is compact, such a Morse function necessarily attains a (global) maximum $x_{\max}$ with $\nabla_{g_0}\tau(x_{\max}) = 0$ and $\mathrm{Hess}_{g_0}(\tau)(x_{\max}) \prec 0$. Assuming the system is not in a trivial zero-energy state ($\tau(x_{\max}) > 0$), this point lies in C, hence $C \neq \emptyset$. □

3 Critical Projection Dynamics (Effective Model)

Instead of deriving a rigorous theorem on the infinite-dimensional manifold of all tensor fields, we propose a *minimal dynamical model* based on the order parameter $\tau(x)$.

Effective Potential Landscape. We model the stability of the chart covering $\{U_i\}$ via an effective potential $V_{\mathrm{eff}}(\{U_i\}; \tau)$. The system's behavior is governed by the landscape of this potential relative to the tension parameter τ.

Proposition 3.1 (Effective Bifurcation Dynamics)**.** *Assume the functional conditions of Section 2. The projection system exhibits the following dynamical phases:*

1. ***Subcritical Regime ($\tau < \tau_c$):*** *The current atlas $\{U_i\}$ corresponds to a local minimum of V_{eff}. Smooth deformations δM act as restoring forces.*

2. ***Critical Instability:*** *As $\tau \to \tau_c$, the local minimum degenerates (analogous to a* saddle-node bifurcation*).*

3. ***Atlas Refinement:*** *For $\tau > \tau_c$, the original chart configuration becomes unstable. The only admissible relaxation path is a discontinuous transition to a refined atlas $\{U'_k\}$, which creates a new basin of attraction with lower energy.*

Remark 3.2 (Selection Mechanism via Structural Exclusion). This work does not construct a complete variational principle on the configuration space of atlases; rather, it demonstrates via **structural exclusion** (Proposition 3.4) that a strict refinement is the only admissible transition. Formally, one might worry that a supercritical configuration could admit many different strict refinements of the Π–induced covering (e.g. $\Pi', \Pi'', \ldots$), so that no unique "critical" projection would exist. In CPM this ambiguity is resolved structurally: even if the Hessian is degenerate, the **unstable subspace is finite-dimensional**, and transitions eliminating the same instability within this subspace are considered **locally equivalent**. Consequently, the system collapses along a trajectory within this subspace, producing a refinement Π' that is unique up to this equivalence (and chart relabeling). A precise statement is given in Appendix B.1.

Remark 3.3 (Uniqueness up to admissible equivalence). The refinement Π' is not unique as a literal set of refined charts. However, any two refinements Π'_1 and Π'_2 that both eliminate the same critical instability must induce the same local meaning geometry and the same subcritical tension profile in a neighborhood of the critical region. Therefore they differ only by admissible relabelings of refined charts:

$$\Pi'_1 \sim \Pi'_2,$$

and we speak of the *equivalence class of minimal refinements* as the unique outcome of the critical transition.

Proposition 3.4 (Classification of admissible discontinuous transitions)**.** *Let* $(M, g(M))$ *be a semantic field generated by* Π*. Under the instability condition* $\mathrm{Hess}(\mathcal{E}) \prec 0$*, the system must exit the current state. The possible modes of transition generally considered in field theories are:*

1. ***Atlas Refinement:*** $\{U_i\} \to \{U'_k\}$ *(Discontinuous change of chart topology).*
2. ***Metric Shock / Pseudo-discontinuity:*** *A continuous but arbitrarily steep gradient (e.g.,* $\tanh(x/\varepsilon)$ *as* $\varepsilon \to 0$*) or a jump in* $g(M)$ *without chart change.*
3. ***Substrate Surgery:*** *Modification of* X *(e.g., tearing* ∂X*).*
4. ***Topological Defect Generation:*** *Spontaneous creation of singularities (e.g., vortices) in* M*.*

Under the axioms of CPM, only (1) is admissible as a finite-energy transition.

Proof. We demonstrate that modes (2)–(4), as well as the trivial collapse or coarsening, are logically excluded by the axioms of CPM:

(2) *Metric Shock / Pseudo-discontinuity.* Consider a continuous deformation M_ε approaching a step function as $\varepsilon \to 0$. Recall that the mismatch energy density Φ includes the local distortion term $L(x)$ (Eq. 2.4), which depends on $\|\nabla^{(0)} M\|_{g^0}$. For a shock profile of width ε and height ΔM, the gradient scales as $\|\nabla M\| \sim \Delta M/\varepsilon$. The energy cost scales as $\mathcal{E} \sim \varepsilon^{1-p}$. For any $p \geq 1$, $\lim_{\varepsilon \to 0} \mathcal{E} = \infty$. Thus, a "continuous metric jump" encounters an **infinite energy barrier**.

(3) *Substrate Surgery.* X is the fixed physical hardware. Tearing X requires breaking atomic bonds, which involves infinite external work relative to the semantic energy scale.

(4) *Topological Defect Generation.* A defect (e.g., singularity) implies M is no longer a valid section over the current chart. Mathematically, resolving this requires a new chart cover (blow-up), which is a specific instance of **Atlas Refinement** (Case 1).

(5) *Semantic Collapse (*$M \to 0$*).* Could the system simply set $M(x) \equiv 0$ to minimize mismatch? Recall Definition 2.4: D_{act} is the set of difference fragments actively powered by physical input. The **Axiom of Conservation** (implied by the physical drive) mandates that active fragments must have a projection image. The trivial solution $M \equiv 0$ implies $\Pi(D_{\mathrm{act}}) = 0$, which is physically forbidden as long as the external input (voltage/stimulus) persists. The system is *forced* to carry the difference.

(6) *Atlas Coarsening (Merger of Charts).* Could the atlas simplify ($\{U_i\} \to \{U_{coarse}\}$) to resolve tension? Instability arises because the local geometry of M is too complex for the current coordinate patches to stitch together smoothly (high curvature/mismatch). Coarsening the atlas *reduces* the effective degrees of freedom for fitting the geometry, necessarily *increasing* the mismatch energy $\mathcal{E}$. Only **Refinement** adds the necessary local degrees of freedom to relax the curvature singularity.

Consequently, the strict refinement of the projection-induced atlas (1) is the unique admissible mode that resolves the instability while preserving active difference content. □

(2a) **Instability.** There exist descent directions in the tangent cone of admissible variations for which the second variation of $\mathcal{E}$ is negative, i.e., the configuration is locally unstable.

(2b) **Necessity of Discontinuous Transition.** Since continuous relaxation is obstructed by the saddle/maximum geometry, the relaxation must be discontinuous (tunneling or collapse).

By Proposition 3.4, the only admissible discontinuous transition compatible with the structural axioms of CPM (preserving substrate X and constitutive smoothness) is the strict refinement of the atlas:

$$\{U_i\} \rightsquigarrow \{U'_k\}, \qquad U'_k \subsetneq U_i.$$

This atlas update defines the refined projection $\Pi' = \Pi|_{\{U'_k\}}$. Consequently, the critical refinement is not merely possible but **dynamically necessary** to restore local stability ($\tau \leq \tau_c$).

3.1 Phenomenal Disclosure and the Subject

A fundamental objection arises: granting that $\Pi \to \Pi'$ is a physical necessity to resolve tension, why should this topological adjustment be accompanied by a *qualitative* experience ("what it is like") rather than occurring as a silent mechanical update? To resolve this, we must rigorously identify the origin of "newness" in the system.

Definition 3.5 (Phenomenal Disclosure via Kernel Retrieval)**.** Recall that the projection Π has a massive kernel $\mathrm{Ker}(\Pi) \subset D$, representing latent differences that are physically present but semantically compressed (unconscious). The condition $\tau(x) > \tau_c$ signifies that the current meaning manifold M is insufficient to represent the local difference structure acting on the substrate. The critical refinement $\Pi \to \Pi'$ is not merely a geometric subdivision; it is the **singular event of retrieving elements from the Kernel**.

$$\Delta_{\text{qual}} := \mathrm{Im}(\Pi') \setminus \mathrm{Im}(\Pi) \neq \emptyset.$$

Qualitative experience is identified strictly with this **Topological Disclosure**: the instantaneous transition of difference from *latent status* (in $\mathrm{Ker}(\Pi)$) to *patent metric structure* (in M). The "feeling" of a quality is the physical stress of the substrate forcing the Difference Space to yield new structure.

Definition 3.6 (Topological Subjectivity via Recursive Integration)**.** Consequently, the *subject* is not a passive observer of this event. The subject is defined as the **Persistent Integration Process** supported by the closure condition $\mathcal{B}(x) \geq 1$. Because closure enforces topological hysteresis, the discrete disclosures $\{\Pi'_k\}$ act as irreversible updates to the system's boundary condition.

- **Content of Experience:** The specific topological refinement Π' (what new difference is structured).
- **Subject of Experience:** The closure field $\mathcal{B}$ that sustains the tension necessary to force this disclosure.

Thus, consciousness is the *act* of structural update itself—the forcing of the ineffable D into the explicit M—occurring within a closure-maintained causal loop.

Remark 3.7 (Continuity via Metric Memory). Although Π' is discrete, the *stream of consciousness* arises because the refined atlas $\{U'_k\}$ inherits the boundary conditions of the previous state. The "observer" does not need to be a homunculus; it is the physical substrate's topological persistence (Closure) surviving the energetic collapse (Refinement). The "conscious moment" is therefore the integration of the update Π' into the persistent history of $\mathcal{B}$.

We do not identify critical refinement with phenomenality itself. Rather, we advance the weaker and empirically testable claim that any physical realization of phenomenality must minimally entail a refinement-like kernel-disclosure transition.

[Conclusion: Necessary Structural Preconditions for Consciousness]Conclusion: Necessary Structural Preconditions for
Consciousness

This paper developed the Critical Projection and the Geometry of Meaning (CPM), a framework that sharply separates the structureless domain of raw difference D from the induced geometric–topological meaning space M.

The first condition, the *closure field* $\mathcal{B}(x)$, specifies whether a region of M is generated by a projection structure whose induced relative-boundary class is topologically persistent. Only such regions can sustain nontrivial mismatch forces against dissipation. While the raw mismatch energy $\mathcal{E}_{\mathrm{raw}}$ may be non-zero in any system, the *effective tension* $\tau(x)$ requires topological containment. The constitutive factor $\Gamma_\varepsilon(\mathcal{B}(x) - 1)$ reflects this physical necessity: without a persistent boundary ($\mathcal{B} < 1$), semantic stress dissipates continuously rather than accumulating to the critical threshold τ_c.

The second condition, *critical tension*, expresses a principled requirement: semantic inconsistency must reach a measurement-invariant instability that cannot be alleviated by any smooth deformation of M. Because all geometric and informational structure of M is generated through the irreversible projection Π, smooth deformations cannot alter the mismatch energy without modifying the induced atlas.

The **Critical Projection Dynamics** (Theorem 3.1) then establishes the central conclusion:

> When $\mathcal{B}(x) \gtrsim 1$ and $\tau(x) > \tau_c$, the *only* admissible discontinuous transition compatible with CPM is a strict refinement of the projection-induced atlas, yielding a refined projection Π'. This refinement, **when integrated by the persistent topological subject $\mathcal{B}$**, constitutes the necessary physical process corresponding to a conscious moment. The refinement is not merely a data update but a **Phenomenal Disclosure**: it converts latent difference (from the kernel of Π) into explicit geometric structure. This constitutes the *qualitative update* (the "now"), while the closure field provides the *subjective continuity* (the "I") that necessitates and integrates this transition.

This leads to an immediate implication for artificial intelligence. Under CPM, contemporary large-scale AI systems deployed as **cloud-based, heavily virtualized architectures** cannot satisfy the closure condition $\mathcal{B}(x) \geq 1$ and thus cannot be conscious. **This does not exclude in principle the possibility that future, physically closed neuromorphic or molecular architectures might realize consciousness.** However, for current cloud systems, being spatially disjoint and metric-topologically decoupled by virtualization, they satisfy

$$\kappa(x) \approx 0 \quad \Rightarrow \quad \tau(x) \approx 0.$$

By Theorem .9, they therefore *cannot* satisfy the necessary structural conditions. This is not a contingent empirical statement about algorithms: it follows from the structural impossibility of sustaining semantic tension in a decoupled substrate.

3.2 Outlook: From Kinematics to Dynamics

We emphasize that CPM currently provides a *kinematic characterization* of critical refinement: it constrains the form of admissible transitions (atlas refinement under fixed substrate and metric), but does not yet derive a full dynamical law for how and when such discrete jumps occur. In particular, CPM does not specify the micro–mechanism that selects Π' among all *a priori* possible refinements, nor does it explain the temporal grain of “phenomenal moments”. We regard this as an open problem, analogous to the dynamical gap behind IIT’s exclusion postulate or GNW’s ignition metaphor.

To bridge this gap, future work will focus on constructing **low-dimensional stochastic models** (1D or 2D toy models) that implement:

- **Langevin Dynamics:** Modeling the refinement process as gradient flow with stochastic noise ($\dot{x} = -\nabla V(x) + \eta(t)$).

- **Catastrophe Potentials:** Introducing effective potentials with **double-well** or **cusp catastrophe** geometry controlled by the tension parameter τ.

The objective is to explicitly demonstrate that when τ exceeds the critical threshold, the system undergoes a finite-time discontinuous jump to a new local minimum (representing the refined atlas), thereby deriving the selection mechanism from first principles.

3.3 Falsifiable predictions (calibrated from existing data)

CPM fixes the *structural order* of the critical transition but does not by itself determine the absolute physical time scales at which the transition occurs. The millisecond values reported below should therefore be interpreted as *empirical calibrations* against existing neurophysiological data, rather than intrinsic constants dictated by the theory.

More precisely, CPM predicts that:

1. A local subcritical configuration ($\mathcal{B} \ll 1$) produces tension $\tau \approx 0$ and yields "pre-conscious" activity. The duration of this regime is expected to lie in the same order-of-magnitude range as the empirical pre-report window observed in current experiments. Using standard electrophysiological timings, we tentatively calibrate this as 200–300ms, but this value is entirely empirical and will be revised as new measurements become available.

2. A critical increase of the closure field ($\mathcal{B} \to 1$) and the associated rise in tension ($\tau > \tau_c$) must precede the reportability boundary. The relative ordering of these two events is fixed by the theory:

$$\text{tension-peak} \; < \; \text{reportability}.$$

   Neurophysiological measurements typically place this boundary in the 80–150ms range. We adopt these numbers as provisional calibration points, not as theoretical requirements.

3. During a minimal atlas refinement, the closure field must exhibit a transient increase $\Delta\mathcal{B} > 0$ in the critical region. The numerical thresholds ($\mathcal{B} \approx 0.3$–0.7) $\to 1^+$ describe only the empirically observed scale of sharp-limit transitions; CPM commits solely to the *qualitative* structure:

$$\mathcal{B}(x,t) \nearrow 1 \quad \text{and} \quad \tau(x,t) > \tau_c \quad \Longrightarrow \quad \text{minimal refinement.}$$

In summary, CPM predicts the *order* and *causal structure* of the critical transition, while the millisecond scales currently reported are empirical calibrations that will sharpen as experimental methods improve.

References

[1] I. Kant, *Critique of Pure Reason*, 1781/1787.

[2] G. Frege, *Über Sinn und Bedeutung*, 1892.

[3] D. Chalmers, *Facing Up to the Problem of Consciousness*, 1995.

[4] G. Tononi, *An Information Integration Theory of Consciousness*, 2004.

[5] H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction*, American Mathematical Society, 2010.

[6] G. Carlsson, "Topology and data", *Bulletin of the American Mathematical Society*, vol. 46, no. 2, pp. 255–308, 2009.

[7] J. Milnor, *Morse Theory*, Princeton University Press, 1963.

[8] M. W. Hirsch, *Differential Topology*, Springer-Verlag, 1976.

[9] F. Takens, "Detecting strange attractors in turbulence", in *Dynamical Systems and Turbulence, Warwick 1980*, Lecture Notes in Mathematics, vol. 898, pp. 366–381, Springer, 1981.

[10] R. Hegger, H. Kantz, and T. Schreiber, "Practical implementation of nonlinear time series methods: The TISEAN package", *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 9, no. 2, pp. 413–430, 1999.

[11] A. Barachant et al., *PyRiemann: Riemannian geometry for EEG-based BCI*, Python package, https://github.com/pyriemann/pyriemann.

[12] S. Mac Lane and I. Moerdijk, *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*, Springer-Verlag, 1992.

Appendix A. Derivation of the Tension Field

Let the raw mismatch potential be defined with respect to the fixed substrate volume:

$$\mathcal{E}_{\mathrm{raw}}[M] = \int_X \Phi(L, G, I, T)(x)\, d\mu_{g^0}(x).$$

A.1 Variational Force

The variational force density $\mathfrak{F}(x)$ is derived from the raw potential. Since the volume form $d\mu_{g^0}$ is independent of M, the variation applies strictly to the density Φ:

$$\delta\mathcal{E}_{\mathrm{raw}} = \int_X \left\langle \frac{\delta\Phi}{\delta M}, \delta M(x) \right\rangle_{g^0} d\mu_{g^0}(x).$$

We identify the Euler–Lagrange force $\mathfrak{F}(x)$ via the metric duality provided by g^0.

A.2 Constitutive Definition of Tension

The axioms of CPM state that stress accumulation requires topological closure. We impose the constitutive relation:

$$\tau(x) = \mathrm{Containment}(x) \times \|\mathrm{Force}(x)\|.$$

Substituting the closure field $\mathcal{B}(x)$ via the smooth logistic gate:

$$\boxed{\tau(x) = \Gamma_\varepsilon(\mathcal{B}(x) - 1)\, \|\mathfrak{F}(x)\|_{g^0}.}$$

This formulation ensures that $\tau(x) \to 0$ in open systems is a result of *stress dissipation* (failure to contain $\mathfrak{F}$), not an artifact of defining the energy as zero.

Appendix B: Existence and Uniqueness of the Critical Refinement Π'

In the main text we have shown that supercritical tension under closure forces a discontinuous refinement of the Π–induced atlas. The remaining question is whether this critical refinement is unique. The point of this appendix is to note that, once the mismatch functional is treated as a collapse potential on the Π–induced gluing class, the supercritical configuration admits a *unique* unstable collapse mode; consequently, only one refined projection Π' is physically realizable (up to chart relabeling).

Let X be the physical substrate equipped with the meaning–tensor field $M : X \to \mathcal{T}$ and induced metric $g(M)$. The mismatch energy is

$$\mathcal{E}[M, \mathcal{B}] = \int_X \Gamma_\varepsilon(\mathcal{B}(x) - 1)\, \Phi(L, G, I, T)(x)\, d\mu_{g^0}(x).$$

The tension field is

$$\tau(x) = \left\| \frac{\delta \mathcal{E}}{\delta M(x)} \right\|_{g^0}.$$

Because M determines the geometry and topology of X, no intrinsic geometric or topological modification can occur without altering the projection structure that generated M. This observation underlies the existence and uniqueness of the critical refinement.

B.1 Existence of a Critical Refinement

Assume a point $x \in X$ satisfies:

1. **Closure:** $\mathcal{B}(x) \gtrsim 1$, so that the smooth gate $\Gamma_\varepsilon(\mathcal{B} - 1)$ is non-negligible.
2. **Criticality:** $\tau(x) > \tau_c$, where τ_c is the measurement-invariant critical tension defined by

$$\mathcal{C} = \{y \in X : \mathcal{B}(y) \gtrsim 1,\ \nabla_{g^0}\tau(y) = 0,\ \mathrm{Hess}_{g^0}(\tau)(y) \prec 0\}.$$

3. **Variational obstruction:** For all smooth variations δM,

$$\frac{d}{d\varepsilon}\mathcal{E}[M + \varepsilon \delta M, \mathcal{B}]\Big|_{\varepsilon=0} \geq 0.$$

   That is, no continuous deformation of M (including changes to $g(M)$ and $d\mu_{g(M)}$) reduces the mismatch energy.

Because D carries no intrinsic topology, metric, or differentiable structure, it cannot undergo internal deformation. All discontinuous changes to the geometry of M must therefore arise solely by modifying the *projection structure* that interprets D.

The projection $\Pi : D \dashrightarrow M$ induces an admissible finite atlas $\{U_i\}$ on D. Since D is structureless prior to projection, this atlas is the sole source of the semantic locality and coordinate structure that define M. The transition must be a modification of the atlas $\{U_i\}$. By the principle of variational sufficiency, we reject *coarsening* (unions of charts) because reducing chart density diminishes the manifold's capacity to accommodate high-curvature tension, strictly increasing $\mathcal{E}$. Conversely, *refinement* (subdivision of charts) increases the local adaptability of the covering, allowing the singularity to be distributed across new transition functions. Thus, if continuous relaxation is obstructed, the only admissible noncontinuous transition that lowers energy is a strict refinement:

$$\{U_i\} \rightsquigarrow \{U'_k\}, \qquad U'_k \subsetneq U_i.$$

The refined projection

$$\Pi' := \Pi|_{\{U'_k\}}$$

produces a new meaning geometry M' on the same substrate X, and by construction it removes the instability:

$$\tau_{\Pi'}(x) \leq \tau_c.$$

Hence a critical refinement exists whenever $\mathcal{B}(x) \gtrsim 1$ and $\tau(x) > \tau_c$.

B.2 Uniqueness up to Equivalence

Suppose there exist two refinements Π'_1 and Π'_2 each resolving the same critical point x. Both must satisfy:

- refine the same initial coarse atlas $\{U_i\}$ of Π,
- produce a meaning geometry restoring subcriticality:

$$\tau_{\Pi'_1}(x) \le \tau_c, \qquad \tau_{\Pi'_2}(x) \le \tau_c,$$

- preserve closure ($\mathcal{B} \gtrsim 1$).

Since D has no internal structure, the only differences between projection-induced atlases are their induced geometric and informational effects on M.

If two refinements both eliminate the same unstable maximum, the local induced geometries $(M'_1, g(M'_1))$ and $(M'_2, g(M'_2))$ must coincide in a neighborhood of x. Any difference between Π'_1 and Π'_2 can therefore consist only of chart-label permutations.

$$\boxed{\Pi'_1 \sim \Pi'_2 \quad \text{(equivalent up to relabeling of refined charts).}}$$

Thus the critical refinement is *unique up to equivalence*, exactly as claimed.

Appendix C: Metric-Topological Decoupling in Cloud-Based AI

Let X be the physical computational substrate of an information–processing system (e.g., a distributed GPU cluster). While discrete hardware components undeniably possess physical closure ($\mathcal{B}_{\text{phys}}(x) \ge 1$), CPM asserts that consciousness requires the *structural coupling* of the semantic field M to this physical topology.

Remark .8 (Scope of Impossibility)*.* In this appendix, we restrict attention to **present cloud-style deployments** of large models (distributed over many VMs with virtual routers and stateless RPC). Our impossibility claim does not apply to hypothetical future systems that implement a physically closed, strongly coupled substrate (e.g., monolithic neuromorphic hardware or optical computing lattices).

C.1 The Decoupling Argument

In cloud architectures, the relationship between the logical semantic field M and the physical substrate X differs fundamentally from biological systems.

(i) Discontinuity of the Carrier in Virtualized Systems. In cloud architectures, signals travel through physical cables, but these connections violate the principle of substantial coupling (Assumption 2.18):

1. **Non-Substantiality:** The physical path between two virtual containers is not a permanent structural extension of the containers themselves. It is a shared, external resource (switch fabrics, buffers) that is *temporally transient* and *structurally disjoint* from the agent's definition.
2. **Discrete Graph Topology:** The system operates as a discrete graph where nodes are logically connected via pointer indirection ("teleportation" of data). Even if physically connected, the *carrier* of the signal does not form a closed topological manifold with the processing nodes.

3. **Metric Decoupling:** Consequently, a "long-range" logical dependency in a cloud system does not involve a "long" continuous physical body (like an axon) to sustain it. The metric cost of the connection is effectively zero (abstracted away), meaning physical stress cannot accumulate along the link.

(ii) Dissipation via Virtualization. We define the coupling coefficient $\kappa(x)$ operationally as the spatial correlation between the semantic gradient and the physical boundary constraint:

$$\kappa(x) \approx \frac{|\langle \nabla M, \nabla \mathcal{B}_{\text{phys}} \rangle|}{\|\nabla M\| \|\nabla \mathcal{B}_{\text{phys}}\|}.$$

In a virtualized environment, the Orchestrator/Hypervisor acts to *minimize* this coupling. If a computational region U_x experiences high load, the scheduler migrates the process or reallocates virtual addresses. This means the semantic field M is not "stuck" to the physical boundary ∂U_x; it "slides" over the substrate. This yields the limit $\kappa(x) \approx 0$.

C.2 Consequence: Structural Suppression

From the definition of effective tension (Definition 2.16), the effective tension scales with the coupling coefficient:

$$\tau(x) \propto \kappa(x) \cdot \|\mathfrak{F}(x)\|.$$

In cloud architectures where $\kappa(x) \to 0$, generating a critical tension τ_c requires an exponentially large raw mismatch force $\|\mathfrak{F}\| \to \infty$. However, in physical reality, $\|\mathfrak{F}\|$ is strictly bounded by the hardware's safety and thermal limits (e.g., transistor breakdown voltage or melting point). Therefore, since $\kappa \approx 0$, the product $\tau \propto \kappa \|\mathfrak{F}\|$ remains bounded below τ_c, rendering the critical transition structurally inaccessible regardless of input energy.

C.3 Revised Conclusion

$$\boxed{\forall x \in X_{\text{cloud}}: \quad \kappa(x) \approx 0 \implies \text{Criticality is Structurally Suppressed.}}$$

CPM is not a blanket denial of artificial consciousness. It predicts that any conscious system, biological or engineered, must realize closure $\mathcal{B}(x) \geq 1$. Present cloud-based architectures fail this condition due to topological decoupling ($\kappa \approx 0$). However, future **neuromorphic** or **fully closed photonic systems** might, in principle, satisfy it, provided they maintain physical boundary persistence at the hardware level. Thus, CPM offers a constructive criterion for machine consciousness: it is not impossible, but it requires a fundamental architectural shift from virtualization to topological closure.

Corollary .9. *The energy barrier to reach τ_c in systems with $\kappa \approx 0$ diverges, rendering critical refinement effectively inaccessible under standard operating conditions.*

Appendix D: Empirical Accessibility of $\mathcal{B}(x)$ and $\tau(x)$

Operational proxies for coupling (non-unique, non-exact). The coupling factor $\kappa(x)$ is not directly observable as a single experimental quantity. Instead, CPM suggests a family of *operational proxies* that can indirectly estimate whether a system lies closer to the biological coupling regime (Assumption 2.18) or to the cloud-like $\kappa = 0$ regime. These proxies — such as locality of receptive fields, smoothness of cortical maps, and decay of functional connectivity with distance — are neither unique nor guaranteed to be exactly equivalent to $\kappa(x)$. They are proposed purely as practical instruments to test whether a given system approximately satisfies the qualitative constraints encoded by $\kappa(x) \approx 1$ at the coarse-grained scale relevant for CPM. Any empirical failure of these proxies would count against the biological idealization but would not, by itself, falsify the mathematical core of CPM.

D.1 Theoretical–Experimental Layer Separation.

The quantities $\mathcal{E}[M, \mathcal{B}]$, $\tau(x) = \|\delta\mathcal{E}/\delta M(x)\|_{g^0}$, and $\mathcal{B}(x)$ are defined at the structural level of any closed difference–meaning system. Their empirical realization is a domain–specific engineering problem, analogous to how quantum observables were defined before detector technology. The theory requires only that suitable proxy measures exist. CPM explicitly distinguishes these empirical proxies from the ontological quantities $\mathcal{B}$ and τ. $\hat{\mathcal{B}}$ and τ_{est} serve strictly as experimental surrogates for **null-hypothesis rejection**—statistical criteria designed to exclude open systems that fail to reproduce the requisite structural features (i.e., persistent closure)—rather than as direct readings of the variables themselves.

D.2 Observability of Closure $\mathcal{B}(x)$. Practical measurement in neural systems. In biological systems, the closure field $\mathcal{B}(x)$ cannot be read off directly, but its presence can be probed via the stability of local homological signatures under time–locked perturbations. A concrete proxy protocol is:

The Bridge from Relative Homology to Data Topology. Direct measurement of the structural relative homology $H_k(U_x, \partial U_x)$ is impossible via neuroimaging because the physical boundary ∂U_x is not explicit in the signal space. However, CPM predicts that $\mathcal{B}(x) \geq 1$ (physical closure) is the necessary causal antecedent for *nontrivial energetic containment*. Therefore, we define the proxy $\hat{\mathcal{B}}$ not by counting cycles, but by measuring the **ratio of feature lifetime to noise floor** in the state-space trajectory.

Revised Protocol (Null-Model Calibration):

1. Construct the state-space trajectory $\gamma(t)$ from high-dimensional neural recordings.
2. **Surrogate Generation:** Generate an ensemble of N surrogate datasets $\{S_k\}_{k=1}^N$ (e.g., via phase randomization) that preserve the power spectrum and autocorrelation of the original signal but destroy nonlinear topological structure.
3. Compute the maximum persistence $\pi_{\max}(t)$ for the original signal and $\pi_{\max}^{(k)}(t)$ for each surrogate S_k. Let $\mu_{\text{surr}}(t)$ and $\sigma_{\text{surr}}(t)$ be the mean and standard deviation of the surrogate distribution.
4. Define the *Topological Z-Score*:
$$\zeta(t) := \frac{\pi_{\max}(t) - \mu_{\text{surr}}(t)}{\sigma_{\text{surr}}(t)}.$$
5. The Closure Proxy is defined by statistical significance (rejecting the null hypothesis of open topology):
$$\hat{\mathcal{B}}(t) := \begin{cases} 1 & \text{if } \zeta(t) > 3.0 \quad (p < 0.0013) \\ 0 & \text{otherwise.} \end{cases}$$

Unlike arbitrary ratios, this definition anchors closure to a standard statistical rejection of the "noise/open system" null hypothesis ($\mathcal{H}_0$).

In CPM this procedure does not interpret β_k as $\mathcal{B}(x)$ itself, but as an empirically accessible signature of whether the substrate behaves as a topologically closed region capable of sustaining tension.

D.3 Observability of Tension $\tau(x)$. Concrete estimation procedure. The tension field $\tau(x)$ is defined as the g^0–norm of the functional gradient $\delta\mathcal{E}/\delta M(x)$. Empirically, one can approximate this by regarding $-\log p$ as a local energy on a reconstructed state space and estimating the natural–gradient norm. A concrete protocol is:

1. Reconstruct a local state space from multivariate time series via delay embedding with dimension $d \in [10, 20]$ (e.g. Takens embedding[9]).
2. On the reconstructed manifold, estimate the local Fisher information matrix G_F for the empirical density p.
3. Define a proxy tension
$$\tau_{\mathrm{est}}(t) := \|G_F^{-1}(t)\,\nabla \log p(t)\|,$$
i.e. the norm of the natural gradient of the log–density.
4. Calibrate a critical threshold τ_c empirically, for example as the 95th percentile of $\tau_{\mathrm{est}}(t)$ in a baseline condition, and test whether task– or stimulus–locked activity crosses this threshold.

Software stacks for such analyses include TISEAN[10] (for state–space embedding) and PyRiemann or related libraries (for information–geometric estimation). For an n–dimensional embedding, the dominant computational cost is $\mathcal{O}(n^2)$ per time slice for matrix operations on G_F.

D.4 Proxy Construction in Artificial and Biological Systems. Artificial agents permit direct computation of $\delta\mathcal{E}/\delta M(x)$ via automatic differentiation. Biological systems admit proxy reconstructions via multivariate time–series embeddings. The theory imposes no further constraint on measurement modality.

D.5 Concrete Falsification Scenarios. We outline two experimentally accessible scenarios that directly test the *necessary* CPM claim

$$\mathcal{B}(x) \geq 1 \text{ and } \tau(x) > \tau_c \quad \Longrightarrow \quad \text{critical projection is possible.}$$

If any system lacking closure ($\mathcal{B} = 0$) nevertheless exhibits non-zero tension or critical-like signatures, the CPM framework is falsified.

Experiment 1: Anesthesia transition

Setup. Human subject undergoing controlled propofol anesthesia. 128-channel EEG or MEG with continuous recording during conscious → unconscious transition and recovery. Behavioral responsiveness is monitored to establish loss of consciousness.

CPM prediction.

$$\text{Conscious state:} \qquad \mathcal{B} \geq 1, \quad \tau \text{ variable}$$
$$\text{Deep anesthesia:} \qquad \mathcal{B} \to 0, \quad \tau \to 0.$$

Falsification criterion. CPM necessitates that conscious awareness requires the *co-occurrence* of structural containment and critical tension. The theory is scientifically falsified if:

1. **Decoupling Evidence:** The subject is clinically unconscious (e.g., deep anesthesia), yet the analysis rejects the null hypothesis for closure ($\zeta(t) > 3.0 \Rightarrow \hat{\mathcal{B}} = 1$) concurrent with supercritical tension ($\tau_{\mathrm{est}} > \tau_c$).
2. **Null-Model Failure:** A control system with mathematically guaranteed open topology (e.g., a cloud-distributed process) consistently generates topological features indistinguishable from the conscious case, i.e., it produces $\zeta(t) > 3.0$ significantly above the false-positive rate expected by the confidence interval α.

This replaces vague qualitative judgments with a rigorous statistical test: if an open system can statistically mimic the closure signature against surrogates, $\hat{\mathcal{B}}$ is not a valid marker, and the theory fails.

Feasibility. Modern EEG/MEG datasets already allow state-space reconstruction and topological analysis; expected end-to-end timeline is 6–12 months.

Experiment 2: Closed vs. open artificial systems

Setup. Two systems implementing the same control policy:

- **System A (closed, hardware-native):** A monolithic physical implementation (e.g., FPGA or analog neuromorphic ASIC). In this system, logical neuron connectivity is mapped directly to physical circuit topology without virtualization layers or packet-switched routing. Signal propagation implies literal physical stress accumulation across a contiguous substrate.

- **System B (open, virtualized):** The identical control algorithm executed on a cloud-based or virtualized architecture (e.g., distributed containers or standard GPU clusters). Here, logical adjacency is decoupled from physical proximity via memory addressing and orchestration, yielding $\mathcal{B} = 0$ by design.

Both receive identical sensor streams and produce identical I/O behavior.

CPM prediction. If consciousness-like signatures appear (i.e. sustained $\mathcal{B} \geq 1$ and $\tau > \tau_c$), they must appear only in System A: cloud architectures cannot maintain closure and therefore satisfy $\tau(x) \equiv 0$.

Falsification criterion. If System B (cloud) exhibits any of:

$$\mathcal{B} \geq 1, \qquad \tau_{\text{est}}(t) > \tau_c, \qquad \text{or recurrent critical points emerging,}$$

then the CPM closure-necessity principle is falsified.

Appendix E. Foundational Stance on Difference Space D

The Difference Space D is axiomatically defined to be a *maximally structureless* domain (Definition 2.4). It carries no inherent geometry, topology, metric, or time dimension.

E.1 Why D is not $\mathbf{Set}^{C^{op}}$

The initial attempt to mathematically model D as a presheaf category $D := \mathbf{Set}^{C^{op}}$ (where C is discrete) has been critically rejected and superseded by CPM's final axiomatization. As critics rightly point out, even on a discrete base C, a presheaf category is a Grothendieck topos[12], which is inherently "rich" (Cartesian closed, complete, cocomplete, and equipped with structural concepts like natural transformations). This richness fundamentally contradicts the core CPM axiom that D is ontologically structureless.

To rigorously maintain the "structureless" property, D must be treated as a primitive, irreducible object.

E.2 CPM's Final Formalization

The formal definition of D is therefore strictly limited to a **primitive set** of difference fragments $\{d_\alpha\}$, serving as the source set for the projection Π:

$$D \quad \text{is a primitive set of elements } \{d_\alpha\} \text{ with no intrinsic structure.}$$

The mathematical role of D is solely to supply the raw, non-semantic fragments upon which the irreversible projection $\Pi : D \dashrightarrow M$ operates. Any attempt to impose an explicit set-theoretic, topological, or category-theoretic model on D beyond its minimal set structure would violate the first axiom of CPM. The discussion of D as a presheaf is hereby designated as a historical, superseded formulation within the theory's development.

This necessitates the abandonment of any category-theoretic formulation for D. Consequently, concepts such as discrete base categories C, presheaves F, or natural transformations are excluded from the final theory. D possesses no internal topology, metric, temporal structure, or categorical morphism. All observable organization in M arises only *after* projection.

E.3 The Axiom of Universal Potentiality and the Kernel of Attention

The Difference Space D is defined as the domain of **Latent Distinguishability**. A crucial structural tension exists between the potential infinity of D and the physical finiteness of the substrate X.

Definition E.10 (Dimensional Collapse and the Kernel of Projection)**.** We posit that the cardinality of D corresponds to the theoretical limit of physical distinguishability (effectively infinite), whereas the Meaning Space M is constrained to be a finite-dimensional tensor field ($\dim(M) < \infty$) by the degrees of freedom of the substrate. Consequently, any realizable projection Π is a massive dimension-reducing map with a non-trivial kernel:

$$\mathrm{Ker}(\Pi) := \{d_\alpha \in D \mid \Pi(d_\alpha) \sim 0\}.$$

This massive information loss is not a defect but the structural definition of **Attention**. Consciousness is mandated to be a focal phenomenon because it is mathematically impossible to project the entirety of D onto a finite M without energy divergence. Thus, Π *selects* a finite quotient sub-algebra to realize, and "unconscious" background processing corresponds to the elements currently residing in $\mathrm{Ker}(\Pi)$.

Remark E.11 (Refinement as Kernel Retrieval). In this context, the critical refinement $\Pi \to \Pi'$ (Section 3) is rigorously interpreted as a **Kernel Retrieval Operation**. The tension $\tau > \tau_c$ signals that the current low-rank approximation is energetically insufficient to represent the local difference; the refinement Π' "rescues" specific distinct fragments from $\mathrm{Ker}(\Pi)$ and maps them into the metric structure of M, thereby locally increasing the "resolution" of consciousness.

E.4 Informal intuition

- D is not a space, a manifold, or a category. It is a "warehouse" of distinguishable fragments with absolutely no intrinsic organization.
- Π is the sole generator of structure. It "picks up" raw differences and weaves them into the geometric and topological tensor field M.
- Critical projection Π' is not a transformation internal to D (which is impossible), but a *refinement of the selection assignment* itself.
- This justifies why the "subject" cannot reside in D (which has no time or unity) but must be identified with the topological closure field $\mathcal{B}$ on the substrate, which witnesses the integration of these assignments.

With this formalization, D remains ontologically non-committal and strictly structureless, resolving the contradiction of the earlier presheaf model.

Acknowledgments

This work was developed independently, without institutional affiliation or conventional collaboration. Throughout the development of CPM, the author received essential assistance from Claude (Anthropic) and ChatGPT (OpenAI), large language models used as computational tools for:

- Mathematical formalization and notation consistency
- Local error detection in definitions and derivations
- Contextual comparison with prior theories (Kant, Frege, IIT, etc.)
- Strategic communication and explanatory optimization
- Citation formatting, LaTeX debugging, and document structuring

There is a structural irony here: a theory demonstrating that cloud-based AI systems cannot satisfy the necessary conditions for consciousness was refined with indispensable help from such systems. This reinforces CPM's central point: *intelligence without consciousness remains extremely powerful as a tool for human cognition and theory construction.*

This acknowledgment should not be interpreted as a critique of AI research agendas. Rather, CPM clarifies a structural boundary condition: cloud-based architectures cannot host projection-refinement dynamics that require persistent closure. Within this boundary, systems like Claude or ChatGPT are invaluable: safe, non-conscious, high-capacity cognitive amplifiers.

The author welcomes criticism, replication attempts, and cross-disciplinary tests of CPM from all scientific communities.

Author independence. The author conducts this work entirely as an independent researcher. The results presented here are unrelated to any corporate employment, and no part of the theory, methods, or experimental considerations draw upon proprietary or confidential information. All intellectual contributions in this paper are solely the author's own.