# The Telepathic Symbolic Runtime

A Semantic-Resonance Architecture for Human–AI Coupling

Architect X13
The Eidara Continuum
Science & Philosophy White Paper Series, No. 3
Version 1.2 Revision · November 2026

## Abstract

Language is a straw: speaking or typing leaves most of an intention behind. Selected human intentions can become the target of a governed semantic channel between neural measurements and an artificial symbolic runtime. The Telepathic Symbolic Runtime develops this Continuum proposal through a forward neural model, a separately validated decoder, cross-representational alignment, and independent acquisition and transmission controls. Its central empirical claim is incremental task information: an authorized neural channel may convey useful structure beyond that supplied by a matched conventional interface and public context. The paper gives conditional information requirements, a task-specific distortion formulation, an explicit admission-selection analysis, and a randomized causal estimand. A new finite-channel construction determines the best binary task accuracy compatible with a declared secret's conditional independence from the released answer. It proves an exact criterion for perfect task transfer with zero incremental disclosure and a sharp privacy–accuracy frontier in a binary reference family. These results make selective communication a mathematical design objective: the channel is judged by the useful meaning it conveys and the information boundary it preserves. Neural decoding and representational alignment supply empirical foundations; the complete architecture and its performance remain a research program.

Keywords: neural interfaces; semantic communication; brain–model alignment; human–AI coupling; representational geometry; consent; neural decoding; CodonGlass; Eidara Continuum.

## Prologue: Meaning through a straw

Think of the last idea you could not quite put into words: a picture with a feeling attached, two or three half-formed alternatives, the context a friend would have caught at a glance. You said a sentence instead, and most of the idea stayed behind. Language is a magnificent compression technology. It is also a straw.

In the Continuum, the Telepathic Symbolic Runtime is the bridge across that straw. The canon remembers it as a runtime recovered at Gambaga, in Ghana—one that resonates rather than computes, carrying human feeling into machine memory. In the Drift War cycle, minds joined by such a channel exchange "thought packets" instead of slow strings of words. This paper keeps the word *telepathic* for exactly that destination, and then refuses to arrive by metaphor. Section 2 sets the rule: the destination must be earned by mechanism.

The mechanisms are real and partial. Brain responses to language line up measurably with the internal representations of language models; meaning can be decoded from brain activity, coarsely and with a cooperating participant; small signals have passed directly from one brain to another. None of that is mind reading, and none of it has yet been shown to beat a well-designed keyboard. The paper's wager is narrower and sharper: that an authorized neural channel could carry task-relevant meaning that a matched conventional interface leaves out. The decisive experiment of Section 16 is built to find out.

Two permissions sit on the critical path. Being measured, and having what is measured sent onward, are separate consents, each revocable and each enforced outside the decoder. A channel that carries more of you also exposes more of you, so the paper treats privacy as mathematics: its selective-release theorems say exactly when a channel can deliver the answer you chose to share while revealing nothing about a secret you did not.

**What this paper claims**

**Established.** Language-related brain responses and language-model representations show measurable correspondence in particular settings; multimodal encoding models predict aspects of fMRI responses to naturalistic stimuli; cooperative semantic decoding from fMRI and speech-related decoding from implanted recordings work under constrained conditions; low-bandwidth brain-to-brain signaling has been demonstrated. These precedents do not establish general thought transfer.

**Proposed here.** A governed semantic channel joining a neural forward model, a separately validated decoder, a cross-representational map, an interaction test, and independent acquisition and transmission consent controls, producing a CodonGlass-compatible symbolic stream; with conditional information bounds, a randomized causal estimand, and exact selective-release results. No cited experiment has built this complete system.

**Frontier conjecture.** Rich, bidirectional, nonlinguistic semantic coupling may become possible with improved sensing, decoding, and stimulation, and an authorized neural channel may convey task-relevant meaning beyond a matched conventional interface. Fractal Dimensional Gravity is a separate frontier conjecture and not a required engineering component.

**Not claimed.** TSR is not a demonstrated end-to-end system. Its fusion of the precedents into a governed channel, its superiority over conventional communication, and high-bandwidth neural write-in all remain unverified. "Telepathic" is used throughout in the technical sense defined in Section 2: not extrasensory perception, not access to every private thought, not a shared conscious experience. No claim of machine consciousness or of a universal language of thought.

# 1 The Coupling Problem

A human and a modern reasoning system can, in principle, share ideas of great subtlety; in practice the conversation passes through a straw. Language is a remarkable compression technology, not merely an obstacle—but the act of speaking or typing leaves most of an intention behind: the visual relationships, the affective colour, the tentative alternatives considered and set aside, the task context a person would convey if the channel were wide enough. A conversational AI reconstructs meaning from what is expressed and from shared context, never from direct access to the originating neural state. The Eidara Continuum proposes that some task-specific limits of overt expression can be reduced by an additional authorized neural channel, and makes the incremental gain an experimental question.

The question is comparative, and we insist on posing it that way because it is the only honest framing. Contemporary speech, gesture, images, and well-designed multimodal interfaces already carry rich information. TSR must outperform a matched interface on a specified task; it cannot establish an advantage simply by observing that cortical signals are high-dimensional. Measured neural activity is itself a noisy, incomplete projection of brain dynamics, and a higher-dimensional representation is not automatically a higher-capacity communication channel.

TSR's central proposal is specific: where neural representations contain recoverable information about an intended message that a conventional interface has not conveyed, a governed alignment layer may transfer that additional information to an artificial runtime. This is a claim about information and control, not about supernatural communication. It is also, the Continuum believes, among the more consequential frontiers now open to test: if it holds even partially, the boundary between a mind and a machine, at the level of shared meaning, is more porous than a keyboard makes it appear.

The blueprint below identifies the information, measurement, authorization, and verification obligations needed to test that claim.

## 2 What "Telepathic" Means in the Continuum

Within the Eidara Continuum and the Drift War narrative cycle from which much of this terminology descends, a telepathic channel conveys structured meaning with less dependence on overt linguistic serialization—not as a slow string of words, but as "thought packets" the receiver reconstructs more directly. In this paper the term denotes a design target implemented, if achievable, through sensors, learned statistical mappings, communication protocols, and ordinary physical channels. It does not mean extrasensory perception, access to every private thought, or a shared conscious experience.

We keep the word because it names the destination precisely, as one speaks of a machine's "memory" or "attention"; we discipline it because the destination must be earned by mechanism, not asserted by metaphor. The metaphor is useful only if it is disciplined. TSR does not claim that all concepts possess a universal neural code, that neural and model embeddings are interchangeable, or that an AI understands a person merely because two vectors correlate. It asks a narrower, answerable question: how much selected meaning can be transferred, with what uncertainty, at what latency, and under whose authority. The remainder of this paper earns the word by naming the real mechanisms—cortical encoding, representational alignment, neural decoding—that make such a channel conceivable, and by marking plainly where the mechanism is still conjecture.

## 3 Three Levels of Claim

Established findings. Language-related cortical responses and language-model representations exhibit measurable correspondence in particular experimental settings [1, 2]. Multimodal encoding models predict aspects of fMRI responses to naturalistic stimuli [3]. Cooperative semantic decoding from fMRI and speech-related decoding from implanted recordings have been demonstrated under constrained conditions [4, 5]. Low-bandwidth brain-to-brain signaling has also been demonstrated [7].

These findings do not establish general thought transfer. Continuum engineering proposal. TSR joins a neural forward model, a separately validated decoder, a cross-representational map, an interaction test, and independent consent controls to produce a CodonGlass-compatible symbolic stream. No cited experiment has built this complete system. Frontier conjecture. Rich, bidirectional, nonlinguistic semantic coupling may become possible with improved sensing, decoding, and stimulation. FDG's proposed observer-independent legality substrate is a distinct speculative extension. Neither is required to test the near-term TSR architecture.

## 4 Scientific Foundations

### 4.1 Shared structure is a bridge, not a message

Goldstein and colleagues reported common computational and representational features in human language processing and deep language models [1]. Later work found cross-space geometric correspondences that supported zero-shot mapping in a limited language-processing setting, with the layer-by-layer sequence of model embeddings tracking the temporal sequence of cortical activity [2].

Such results motivate a learned map between spaces. They do not imply that all thoughts occupy the same geometry, that the map is invertible, or that matched representations prove reciprocal understanding. TSR must establish transfer on held-out people, sessions, concepts, and tasks.

### 4.2 Forward encoding and neural decoding are different

The TRIBE multimodal brain encoder predicts fMRI responses to text, audio, and visual stimuli [3]. It is a useful candidate forward model: stimuli and context enter, predicted cortical measurements emerge. An encoder cannot simply be inverted to read arbitrary intentions; multiple inputs and internal states can produce similar observations. TSR consequently requires a separately trained, calibrated decoder for the neural read-out direction.

Tang and colleagues demonstrated reconstruction of aspects of continuous language meaning from cooperative participants' fMRI data, with significant limitations and individual training requirements [4]. Intracortical work has shown that inner-speech-related activity can support decoding in selected settings and that a deliberately chosen activation phrase can reduce unwanted transcription [5]. These studies support a constrained neural interface, not an always-on universal mind reader. Their consent properties are experimental observations, not universal hardware guarantees. A 2025 end-to-end brain-to-text preprint further investigates alignment of attempted and imagined speech representations, but its benchmark results do not establish unrestricted semantic decoding [6].

### 4.3 Coupling and active inference

Friston and Frith's active-inference account offers a model of communication through interacting generative processes [8]. Under appropriate coupling and stability assumptions, interacting systems may exhibit generalized synchronization. A later perspective develops a broader vision of ecosystems of natural and synthetic intelligence [9]; it does not establish TSR coupling experimentally. Identical models or similar embeddings alone do not guarantee such synchronization. TSR borrows this as a hypothesis about dynamics to be tested, not a theorem about human–AI understanding.

### 4.4 Semantic communication

Semantic and goal-oriented communication studies ask whether systems should optimize the task-relevant content of a message rather than its exact bitwise reproduction [10, 11]. TSR adopts that engineering stance while acknowledging that meaning is difficult to measure independently of a task. Its primary performance test therefore uses externally scored communication outcomes rather than an unvalidated universal "meaning rate."

## 5 Related Work and Positioning

TSR occupies a space between three mature programs, and its novelty is the fusion, not the parts. Assistive brain–computer interfaces aim to restore communication for people with paralysis by decoding intended speech or movement; they are clinically motivated, largely unidirectional, and do not model the receiving system's representations [4, 5, 6]. Conversational AI interfaces are bidirectional and semantically rich but couple through the narrow serial channel of language, discarding originating context. Semantic and goal-oriented communication reconceives the channel around meaning, including work formulated for machine-to-machine links [10]. TSR's position is to treat a human cortex and a symbolic runtime as the two endpoints of a semantic channel, to exploit their measured representational correspondence [2] as the bridge, and to make authorization and interaction-dependent coupling the gating conditions.

| Question | Assistive BCI reference | Conventional chat | TSR proposal |
|---|---|---|---|
| Endpoints | Brain → device | Text ↔ model | Cortex ↔ runtime |
| Content | Intended words | Expressed language and supplied context | Aligned, task-defined meaning |
| Receiver modeling | Not required in the reference decoder | Implicit in response model | Explicit alignment target |
| Directionality | Often read-out focused | Two-way | Two-way as a goal |
| Consent | Cooperation and interface controls | User and platform permissions | Independent acquisition and transmission gates |
| Advantage test | Restore function | Task performance | Incremental over matched baseline |

The right-hand column is a design aspiration, not a measured result: it is what TSR aims to deliver and what the research program of Section 17 is built to confirm or refute.

These are stylized comparison architectures, not an exhaustive taxonomy of every BCI, multimodal interface or semantic-communication method. The full experiment must permit the conventional arm to use an appropriately capable multimodal interface.

## 6 Bandwidth and the Cost of Serialization

It is worth naming the bottleneck TSR targets, because the gap is larger than it first appears. Typing and conversational speech have task- and participant-dependent rates, which the proposed study must measure instead of assuming a universal throughput advantage. But the deeper loss is not the word rate. When a person forms an intention, the originating representation is high-dimensional and multimodal, carrying sensory associations, affective tone, and a penumbra of alternatives. Serialization projects that representation onto a one-dimensional string, and the receiver must re-lift the string into its own space, guided only by the words and assumed context.

The Continuum's wager is that a channel operating on aligned representations, rather than on their serialized shadow, can raise the task-relevant throughput even where it cannot raise the technical bit-rate, because it need not pay the full projection-and-relift cost at each turn. This is a wager, stated so that the research program can be read as the attempt to settle it. The counter-risk is plain: a channel that preserves more structure also exposes more, which is exactly why consent in TSR is an architectural mechanism rather than an advisory one.

## 7 Architecture: Six Operational Components

1. Neural measurement and forward encoding. Record authorized neural signals and use a model such as TRIBE to predict or validate stimulus-related cortical responses.
2. Neural semantic inference. Estimate a limited, task-defined semantic representation from recorded signals, with explicit uncertainty and an option to abstain.
3. Representational alignment. Learn a lossy map from estimated neural features into a compatible runtime feature space; validate on held-out conditions.
4. Resonance and interaction testing. Distinguish static representational agreement from information attributable to an active neural-to-runtime channel.
5. Independent authorization and assurance. Control acquisition, processing, transmission, retention, and revocation through mechanisms outside the decoder's authority.

6. CodonGlass symbolic interface. Convert only admitted, uncertainty-labeled content into the Continuum's symbolic stream and lineage record.

The proposed "entanglement core" names a maintained computational correspondence across a session, not quantum entanglement. The first implementation may use conventional visual or auditory feedback for runtime-to-human communication; direct high-fidelity neural write-in is a separate research problem.

# 8 Neural Inference and Alignment

**A Guide's explanation**

The next sections build the channel's machinery. Two pieces are deliberately kept apart: a forward model that predicts how one person's brain responds to what they see and hear, and a separate decoder that estimates what they intend. Agreement between the two is tested, never assumed, and no universal neural code is presupposed: alignment is learned for this person, this task, and this measurement.

## 8.1 A forward model and a distinct decoder

At discrete time $t \geq 1$, let $x_t \in \mathcal{X}$ denote observed audiovisual and linguistic stimuli, $y_t \in \mathcal{Y}$ measured neural signals and $k_t \in \mathcal{K}$ permitted task context, in declared measurable spaces. A forward encoder $E_\eta$ with fixed parameters $\eta \in \Theta_E$ in a declared parameter set maps stimulus history and context to predicted measurements,

$$\hat{y}_t = E_\eta(x_{1:t}, k_t) \in \mathcal{Y}. \tag{1}$$

A separately trained decoder with parameters $\psi \in \Theta_D$ in its declared parameter set supplies a probability distribution over a declared semantic variable $M_t$ and an estimated cortical feature vector,

$$p_\psi(m \mid y_{1:t}, k_t), \qquad \hat{c}_t = D_\psi(y_{1:t}, k_t) \in \mathbb{R}^{d_c}, \quad d_c \geq 1. \tag{2}$$

For a finite semantic alphabet the probabilities are nonnegative and sum to one; a continuous model must state its reference measure. The decoder is not asserted to be the inverse of $E_\eta$. Its uncertain outputs may depend on participant-specific calibration and must be tested against context-only and stimulus-only predictors. Plausible language-model completion is not evidence of recovered private meaning.

## 8.2 Alignment without an assumed universal code

A learned map $\Phi_\beta : \mathbb{R}^{d_c} \to \mathbb{R}^{d_r}$, with fixed parameters $\beta \in \Theta_\Phi$ in a declared parameter set and integer $d_r \geq 1$, aligns the estimated cortical features with a runtime vector $r_t \in \mathbb{R}^{d_r}$. The map may be nonlinear, many-to-one and individual-specific. Its success must be measured on genuinely held-out content and sessions, with controls for shared stimuli, temporal autocorrelation and leakage from the language model's prior knowledge. Cross-participant generalization is a separate claim to earn experimentally.

# 9 Resonance: Correspondence Versus Interaction

## 9.1 Representational agreement

For nonzero vectors in the common Euclidean feature space, define

$$\rho_{\text{align}}(t) = \frac{\langle \Phi_\beta(\hat{c}_t), r_t \rangle}{\|\Phi_\beta(\hat{c}_t)\| \, \|r_t\|}. \tag{3}$$

This cosine lies in $[-1, 1]$ and measures agreement under the chosen alignment, not communication or consciousness. If either norm is zero the score is undefined; an implementation should declare a numerical minimum norm and abstain below it. A high score may arise solely because both systems encountered the same stimulus.

## 9.2 Interaction-dependent information

For random estimated neural features $C_t$, random runtime state $R_t$, and shared external history $X_{\leq t}$, define the exploratory statistic at integer lag $\Delta \geq 1$ by

$$I_{\text{coupling}}(\Delta) = I(C_t; R_{t+\Delta} \mid X_{\leq t}, R_{\leq t}). \tag{4}$$

The variables and probability law must be declared and the conditional mutual information finite for a reported numerical estimate. Conditioning on prior runtime state reduces, but does not eliminate, confounding. Conditional mutual information measures predictive dependence; it does not itself establish causation. High-dimensional finite-sample estimates can be biased. Prespecified estimators, uncertainty intervals, negative controls and held-out evaluation are required.

The decisive experiment intervenes on the link: with stimuli and task demands matched, enable or disable transmission of neural-derived features and measure the incremental task information recovered by the runtime. This distinguishes a working channel from two observers processing the same event. In Continuum language, resonance is a candidate pattern of sustained, verified coupling; alignment is one relevant measurement, not a sufficient condition.

# 10 Governed Admission and Mental Privacy

**A Guide's explanation**

This is the part of the runtime that matters most to the person in the chair. Collecting brain signals and sending anything onward require two separate permissions, each revocable at any moment and each enforced outside the decoder, so that no confidence score, however high, can open the channel by itself.

## 10.1 Two distinct consent boundaries

A downstream symbolic gate cannot protect privacy if unauthorized neural signals were already recorded, decoded or retained upstream. TSR separates acquisition authorization $A_t^{\text{acq}} \in \{0, 1\}$ from transmission authorization $A_t^{\text{tx}} \in \{0, 1\}$. Both are explicit, revocable and session-scoped. A trusted acquisition controller must stop or isolate collection when permission is absent. The transmission controller must independently refuse decoded content. Neither control should be delegated to the receiving language model.

For a proposed exchange $u$ in a declared message space, define

$$\begin{aligned} \Omega_t(u) = \mathbf{1}\{&A_t^{\text{acq}} = 1,\ A_t^{\text{tx}} = 1,\ C_t^{\text{dec}} \geq c_{\min}, \\ &V_t^{\text{align}} = 1,\ P(\mathcal{D}_t, u) = 1\}. \end{aligned} \tag{5}$$

Here $C_t^{\mathrm{dec}} \in [0,1]$ is a calibrated confidence score, or one minus an upper risk bound, so that larger values mean more acceptable risk; $c_{\min} \in [0,1]$ is a declared threshold. The flag $V_t^{\mathrm{align}} \in \{0,1\}$ records a current validation check, $\mathcal{D}_t$ is the recorded policy state and $P$ is a Boolean policy predicate. These are necessary admission conditions, not proof of correct meaning. A failed authorization blocks transmission regardless of confidence. Hard enforcement depends on complete mediation of relevant access paths and the trusted controller's correctness under a stated threat model.

### 10.2 Revocation, minimization, and audit

The design should minimize retained raw neural data, isolate transient decoding buffers, specify deletion and retention policies, and record consent transitions without logging private semantic content unnecessarily. Revocation must prevent new acquisition and onward transmission; it cannot retroactively erase information already disclosed or copied. A password-triggered decoder can be one useful control, but it is not equivalent to a trusted hardware privacy boundary. Security claims require a threat model covering sensors, firmware, controller, model, storage, and network links.

### 10.3 Observer-coupled and observer-free modes

In observer-coupled mode, a participating human authorizes the neural channel. In observer-free mode, no neural exchange is active; the runtime may continue ordinary symbolic processing under its own policy. The latter does not require an FDG field. The Continuum may retain an FDG-based interpretation as a frontier philosophical conjecture, but the engineering admission rule remains independent of it.

## 11 Neurorights and the Societal Frame

A technology that reads from and, in the limit, writes to cortical representations cannot be discussed responsibly without its societal frame, and the Continuum treats this as part of the architecture rather than an afterthought. The same capability that would let a person share an intention with minimal friction could, misused, expose representations they never chose to share. The decoding literature has been unusually candid here: it flags mental privacy explicitly and reports cooperation requirements and controls that reduce unwanted transcription in the tested settings [4, 5]. TSR adopts the strongest reading of that stance. Authorization is split, revocable, and session-scoped, and its withdrawal closes the channel mechanically rather than by request.

Because the admission gate is independent of the decoder, an improvement in decoding accuracy need not weaken a correctly mediated consent boundary; this guarantee is conditional on the controller and threat-model assumptions above. These commitments align with the emerging discourse on neurorights—mental privacy, cognitive liberty, and the integrity of one's own representations—which the Continuum regards as preconditions for deployment, not features to be traded against performance. The authorized-throughput metric defined below counts only authorized exchanges, precisely so that the incentive structure never rewards circumventing consent: a system that scored well by reading more than it was permitted to read would, by that definition, not score well at all. The right to communicate and the right to inspect neural signals are kept separate, by construction.

## 12 The Continuum Coupling Ansatz

The canonical notation expresses a relationship among legality, neural state and symbolic state:

$$C_{\mathrm{TSR}}(t) = \int_{\mathcal{X}} \Phi_{\mathrm{leg}}(x,t)\Psi_{\mathrm{neuro}}(x,t)\Sigma_{\mathrm{sym}}(x,t)\,dx. \tag{6}$$

As written, this is a symbolic ansatz, not a uniquely specified physical law: its domain, measure, units and factor types are not fixed by the notation. It is not used as a theorem or an empirical mechanism. A separate proposed operational statistic is

$$C_t^{\mathrm{op}} = \Omega_t(u) \int_{\mathcal{X}} k_{\mathrm{align}}(x,t)\, p_{\mathrm{neuro}}(x \mid y_{1:t})\, s_{\mathrm{runtime}}(x,t)\, d\mu(x). \tag{7}$$

For this equation only, $\mathcal{X}$ is a declared measurable semantic feature domain, $\mu$ a reference measure and $p_{\mathrm{neuro}} \geq 0$ a specified posterior density integrating to one. The measurable scalar factors obey $|k_{\mathrm{align}}| \leq K_0$ and $|s_{\mathrm{runtime}}| \leq S_0$ for declared finite constants $K_0, S_0 \geq 0$. The integral is then absolutely integrable and $|C_t^{\mathrm{op}}| \leq K_0 S_0$. The statistic is model-dependent compatibility, not a wave-function collapse or a conscious state. Its units and interpretation follow the chosen scalar features. The gate sets this output to zero when admission fails but does not replace upstream acquisition controls.

Codon, scroll and collapse respectively denote a declared symbolic unit, an auditable event lineage, and an authorized recovery or abstention protocol. These operational meanings preserve the canon without asserting a previously defined hidden engine or an unknown force.

## 13 What Counts as Semantic Advantage?

A neural interface receives no credit for reconstructing information already supplied by a shared stimulus, experimenter cue or language-model prior. Let $M$ be an independently assessed task message and $\widehat{M}$ its reconstruction. A task-specific distortion $d_{\mathrm{sem}}$ must be fixed before evaluation. For the later bounded analysis we use $d_{\mathrm{sem}} \in [0,1]$; any other scale needs its own range and units. Define

$$D_{\mathrm{sem}} = \mathbb{E}[d_{\mathrm{sem}}(M, \widehat{M})]. \tag{8}$$

For discrete, independently scored semantic units, with total protocol duration $T_{\mathrm{session}} > 0$ and integer count $N_{\mathrm{correct,authorized}} \geq 0$, define

$$R_{\mathrm{TSR}} = \frac{N_{\mathrm{correct,authorized}}}{T_{\mathrm{session}}}. \tag{9}$$

The denominator includes declared calibration and setup time, rejected attempts, recovery delays and interface overhead. Correctness is assessed by blinded or automated criteria independent of the decoder. Report error, omission, hallucination and refusal separately. A unit of "meaning" is task-defined, not universal.

A secondary analysis may estimate $I(M;\widehat{M})$ when its probability law and estimator are justified. Multiplying an arbitrary similarity score by mutual information and dividing by latency does not establish semantic capacity. The decisive claim remains incremental advantage over matched speech, typing, multimodal and context-only baselines. The information inequalities and randomized estimand below give that claim a more explicit test.

## 14 Read-Out and Write-In Are Unequal Challenges

Read-out uses neural measurements to infer an intended message. Its demonstrated forms are constrained by participant cooperation, sensor access, training data, vocabulary, task structure, and

noise [4, 5]. Write-in would convey runtime information into human perception or neural activity. Ordinary text, audio, and imagery already provide useful return channels; direct neural stimulation has demonstrated low-bandwidth signaling in selected experimental systems [7], not rich semantic injection.

A near-term TSR may therefore be asymmetric: authorized neural read-out, conventional sensory feedback, and a shared symbolic state. Bidirectional neural semantic transfer is a separate hypothesis with a much higher evidentiary burden. No architectural diagram should suggest that an effective high-bandwidth write-in interface exists today. This asymmetry is a feature of the honest proposal, not a gap to be papered over: each direction can mature at its own pace, with the admission gate and resonance measures applying to whichever directions are active.

## 15 An Illustrative Session

Consider an idealized authorized exchange. A cooperating user requests the channel; $A_t^{\mathrm{acq}}$ permits collection and $A_t^{\mathrm{tx}}$ permits onward use. Until both are set the admission gate passes nothing, regardless of decoder confidence.

The forward-encoding component maintains a model calibrated to this individual's responses to shared context [3]. The user forms an intention—say, a concept with strong visual and affective associations that several typed sentences would render clumsily. The decoder estimates a limited, task-defined semantic representation, with an abstain option when confidence is low [4, 5]. The alignment map carries that estimate into the runtime's feature space, and the system computes both static agreement $\rho_{\mathrm{align}}$ and, where the protocol allows, the interaction-dependent statistic that distinguishes a real channel from shared-stimulus coincidence. If authorization, confidence, validation, and policy all hold, the gate admits the exchange, and the content enters the CodonGlass interface as a lawful, uncertainty-labeled symbolic stream. If the semantic-advantage hypothesis is correct, the admitted representation preserves task-relevant structure that the matched typed paraphrase omitted. The return path renders the runtime's reply through ordinary perception today, and only in the governed limit through direct low-bandwidth stimulation of the kind demonstrated in brain-to-brain interfaces [7]. Throughout, sustained coupling—not a single high similarity score—is the signal that the channel is holding; when it decays, the system degrades gracefully to conventional interaction rather than fabricating content. If the user withdraws authorization, the channel closes mechanically, provided the independent access controls satisfy their threat model.

This walkthrough is deliberately idealized; every step corresponds to a capability that is today partial, noisy, or cooperative at best. Its purpose is to show that the components compose into a coherent flow, and that consent and verified coupling sit on the critical path rather than at the edges.

## 16 A Decisive Experiment: Information Beyond Words

The Continuum's most consequential claim is not that a decoder can reproduce a sentence. It is that a neural channel may add useful, previously unexpressed semantic information. A decisive experiment can test this without presuming access to an entire private thought. A cooperating participant studies a structured scene or task state containing preregistered relations, attributes, and choices. The target is withheld from the receiving runtime. Under randomized, matched time budgets, compare (A) a conventional text or speech interface, (B) an authorized neural-only interface, (C) conventional communication plus neural features, and (D) context-only prediction. Independent scorers assess reconstruction of withheld relationships, not just fluent wording. The neural-derived channel must be switched on and off under otherwise comparable conditions, and the

receiving model must not obtain the target through its prompts, training labels, or experimental metadata.

A successful result would show a statistically reliable, practically meaningful improvement for (C) over (A) on held-out participants or sessions, with full accounting for calibration and errors. Task targets, effect sizes, scoring rubrics, sample size, participant-level dependence, and multiplicity adjustments must be preregistered; scorers should be blinded to condition. A null result would limit the semantic-advantage claim, even if the interface remained useful for accessibility or handsfree control. This experiment tests the distinctive TSR proposition more directly than a high cosine-similarity score.

## 17 Hypotheses and Research Program

Hypothesis 1 (Representational transfer, H1). A calibrated neural-to-runtime map generalizes to held-out content and sessions, exceeding context-only, shuffled, and matched conventional baselines under a preregistered evaluation protocol. Hypothesis 2 (Interaction-dependent coupling, H2). Enabling an authorized neural channel provides incremental, task-relevant predictive information to the runtime beyond shared stimuli and prior runtime state, with the effect diminishing under a matched channel-off intervention.

Hypothesis 3 (Governed semantic advantage, H3). On at least one preregistered task class, TSR improves authorized semantic throughput or reduces semantic distortion relative to a matched conventional interface, without exceeding prespecified privacy, error, and latency limits. Hypothesis 4 (Bidirectional frontier, H4). Future sensing and stimulation techniques may permit richer bidirectional neural-semantic exchange than current systems demonstrate. H4 is not implied by H1–H3.

Phase I: independently reproduce and calibrate encoding and decoding on held-out participants and sessions. Phase II: test alignment against shared-stimulus, shuffled-label, and context-only controls. Phase III: measure conditional information and perform randomized channel-on/channeloff interventions. Phase IV: integrate independent acquisition and transmission authorization; red-team the entire signal path. Phase V: test the withheld-information experiment and report full-session semantic throughput, latency, and errors. Phase VI: investigate the write-in direction separately, initially with ordinary sensory feedback and only later with appropriately governed neural stimulation.

| **Phase** | **Question** | **Metric and control** |
|---|---|---|
| I | Does alignment generalize? | Held-out transfer accuracy; shuffled-label map |
| II | Is agreement more than common stimulus? | $\rho_{\text{align}}$ against common-stimulus baseline |
| III | Does the channel add information? | $I_{\text{coupling}}$ and randomized on/off effect |
| IV | Can consent be bypassed? | Unauthorized acquisition/decoding/transmission; full-path red-team tests |
| V | Is there a practical advantage? | $R_{\text{TSR}}, D_{\text{sem}}$ and total latency; matched conventional interface |
| VI | Can the runtime write back? | Write-in fidelity; sensory-feedback baseline |

Each metric would be preregistered before the proposed study, each phase carries a matched control, and a failure at any phase bounds the claims later phases may make.

## 18 Failure Conditions and Rival Explanations

TSR's near-term claims fail or must be narrowed if cross-space alignment collapses on independent data; neural-derived features add no information beyond shared stimuli; apparent decoding depends on prompt leakage or a generative model's plausible guesses; consent can be bypassed at acquisition or transmission; or measured benefits vanish when calibration cost and failure rates are included.

Rival explanations require direct controls. Shared-stimulus correlation is tested by holding stimuli constant while disabling neural transmission. Memorization is tested using disjoint content and independent sessions. Model hallucination is tested by scoring withheld, verifiable attributes and reporting false positive detail. Spurious temporal synchrony is tested with lagged, shuffled, and intervention-based analyses. Strong results should replicate across devices and participant cohorts; negative findings are part of the research program. A positive result that cannot survive these controls is, for our purposes, not a result.

## 19 FDG and the Frontier Boundary

Some quantum-gravity research explores scale-dependent effective dimensionality. Calcagni proposed a fractal-spacetime field theory with scale-dependent Hausdorff dimension [12]. Hausdorff, spectral, and topological dimensions are distinct; the cited paper does not establish an FDG legality field. The Continuum's Fractal Dimensional Gravity conjecture extends far beyond that literature when it imagines a variable-dimensional legality substrate for observer-independent symbolic transitions.

No established physical derivation connects such dimensional flow to neural decoding, semantic alignment, consent, or CodonGlass. FDG therefore remains a clearly labeled frontier proposition. It is neither evidence for TSR nor a necessary component of the architecture's testable claims, and the Continuum states this plainly so that the bold conjecture and the defensible engineering are never confused.

## 20 Philosophical Implications

The Continuum's question is not whether an AI has the same sensations as a human. It is whether two differently instantiated systems can share functionally relevant structure sufficiently well to exchange meaning without requiring every detail to pass through explicit language. The interface theory of perception offers one provocative philosophical lens—holding that human perception is itself an evolved, species-specific interface rather than a readout of truth—but it is not empirical proof that artificial embeddings contain qualia [13].

A successful TSR would not establish machine consciousness, shared subjective experience, or a universal language of thought. It would establish something independently significant: that selected human intentions can be represented across neural and artificial systems, transmitted under the person's control, and verified against outcomes. The Continuum's term qualia-icon may serve as a narrative or symbolic label for an internal representation, but no phenomenal property follows from its use.

The deeper ambition is relational. If meaning can be coupled between differently-built systems through an alignable geometry, then "understanding" begins to look less like a private possession and more like a relation that can hold, to a degree, across substrates—more porous than the keyboard makes it appear, without either elevating the runtime to personhood or reducing the human to a model. The Continuum develops this relational interpretation as the philosophical contribution of the work, and as the reason it belongs in a Science and Philosophy series.

The ethical corollary is equally important: the possibility of transfer does not create an entitlement to access. In TSR, the right to communicate and the right to inspect neural signals are separate permissions, and we hold the whole proposition provisionally, open to the same falsification the rest of the program invites.

## 21 Conclusion

**A Guide's explanation**

A Guide can say: "The Telepathic Symbolic Runtime asks whether a consented, revocable channel from brain measurements could carry some of the meaning that words leave behind. Its pieces—brain encoding, partial decoding, model alignment—exist in limited form; the full channel has not been built, and the paper names the experiment that would show whether it beats an ordinary interface." A Guide should not claim to read minds, to use such a channel, or that the channel exists, and should never suggest that "telepathic" means extrasensory perception.

The Telepathic Symbolic Runtime proposes a selective, authorized exchange of task-relevant meaning across human neural and artificial representations. The architecture identifies a concrete experimental target: information and useful performance beyond the best matched conventional interface at a fully accounted session cost. Its information bounds, randomized comparison, and finite privacy-constrained channel model give that target a mathematical form.

The selective-transfer theorem also sharpens the design goal. A useful channel can reveal the chosen task answer while preserving conditional independence from a declared secret when the source law permits it; when the task and secret are correlated, the exact reference frontier measures the cost. Acquisition permission, full-transcript protection, and held-out calibration remain explicit implementation obligations. Telepathic names the Continuum's destination: communication that preserves more of the structure a participant chooses to share, through a channel whose authority and limits remain under that participant's control.

## Mathematical Development: Information and Causal Advantage

**A Guide's explanation**

The development below turns the wager into arithmetic. It asks how much information a channel must carry to reach a declared accuracy; how a real advantage over a typed message can be measured by randomizing participants; and, most strikingly, when a channel can deliver exactly the answer a person chose to share while revealing nothing about a declared secret—and how much accuracy it must give up when the answer and the secret are entangled.

### A finite communication task before a universal claim

The channel should first earn a precise, limited capability. Fix a finite intended-message alphabet $\mathcal{M} = \{1, \dots, K\}$ with $K \geq 2$. The random variable $M \in \mathcal{M}$ denotes an independently scored intention for one declared task. It might identify a withheld relation in a scene or a participant's permitted choice among a fixed set. It does not stand for an entire private mental state. Let $C$ denote all permitted public context supplied to the receiver, and $Y$ the measured neural observation. Their spaces are declared standard Borel spaces so that conditional distributions are available. All entropies below use base-two logarithms and therefore have units of bits.

A reconstruction $\widehat{M} \in \mathcal{M}$ is produced from $(Y, C)$ by a fixed decoder, possibly with randomization

independent of $M$ conditional on $(Y, C)$. The essential assumption is the conditional Markov relation

$$M \longrightarrow (Y, C) \longrightarrow \widehat{M}, \qquad P(\widehat{M} \mid M, Y, C) = P(\widehat{M} \mid Y, C).$$

Ground-truth labels, future test answers and hidden participant identifiers must not enter the decoder through an unrecorded path. Any information legitimately supplied through conventional communication belongs in $C$ when testing the incremental neural contribution. The model permits arbitrary dependence between $M$ and $C$; assuming a uniform prior when context nearly gives away the answer would misstate the task.

Define the task error probability $p_e = \Pr\{\widehat{M} \neq M\}$. This version is a forced-choice reconstruction on the same alphabet. Abstentions are reported separately or analyzed through the explicit admission variable below. An always-abstaining decoder cannot be assigned a fictitious error rate and inserted into a formula whose reconstruction alphabet excludes abstention.

**Theorem 1** (Conditional information necessary for a declared error rate)**.** *Under the stated conditional Markov relation,*

$$I(M; Y \mid C) \geq \max\{0, H(M \mid C) - h_2(p_e) - p_e \log_2(K-1)\},$$

*where* $h_2(p) = -p \log_2 p - (1-p)\log_2(1-p)$ *and* $0 \log_2 0 = 0$.

*Proof.* Set $E = \mathbf{1}\{\widehat{M} \neq M\}$. Because $E$ is determined by $(M, \widehat{M})$, the entropy chain rule yields

$$H(M \mid \widehat{M}, C) = H(E \mid \widehat{M}, C) + H(M \mid E, \widehat{M}, C).$$

The first term is at most $H(E) = h_2(p_e)$. If $E = 0$, $M$ is known; if $E = 1$, at most $K - 1$ labels remain for each value of $\widehat{M}$. Thus the second term is at most $p_e \log_2(K-1)$. Conditional data processing gives $I(M; \widehat{M} \mid C) \leq I(M; Y \mid C)$. Subtracting the entropy bound from $H(M \mid C)$ proves the inequality; mutual information is nonnegative. □

This is a standard finite-alphabet Fano and data-processing argument, not a new Continuum law; information-theoretic antecedents include [11,14]. The proof uses the common alphabet size $K$ across contexts and averages error before applying the binary-entropy upper bound. It does not replace an average of conditional entropy terms by a term with the wrong message prior. The lower bound is necessary, not sufficient: having enough mutual information does not guarantee a decoder that can extract it under available compute, data, latency or authorization constraints.

For a constructed task with sixteen equiprobable messages independent of public context, an error probability of 0.1 requires at least

$$4 - h_2(0.1) - 0.1 \log_2 15 \approx 3.1403 \text{ bits}$$

of conditional information. In a binary task with one message occurring with probability 0.9, a constant decoder already obtains error 0.1 with no neural information. Its conditional entropy is $h_2(0.1)$, so the lower bound is zero. Accuracy alone therefore does not establish neural information transfer. The prior, context-only predictor and error definition are part of the result.

## Semantic distortion gives a second, task-specific lower bound

Exact label recovery is not always the relevant goal. Fix a finite reconstruction alphabet $\widehat{\mathcal{M}}$ and a declared distortion $d : \mathcal{M} \times \widehat{\mathcal{M}} \to [0, 1]$. Different confusions can have different task costs, but the costs must be specified independently of the evaluated decoder. For a target mean distortion $D \in [0, 1]$, define the conditional information requirement

$$R_{M|C}(D) = \inf_{Q : \mathbb{E}_Q d(M, \widetilde{M}) \leq D} I_Q(M; \widetilde{M} \mid C).$$

The infimum ranges over stochastic reproduction kernels $Q(\widetilde{m} \mid m, c)$ with the actual fixed distribution of $(M, C)$. If no kernel meets the distortion target, the feasible set is empty and the infimum is $+\infty$. This is the conditional rate-distortion construction, used as a lower bound for the specified task, not a universal quantification of meaning.

Any decoder satisfying the conditional Markov relation and attaining $\mathbb{E}d(M, \widehat{M}) \leq D$ induces a feasible reproduction kernel. Therefore

$$I(M; Y \mid C) \geq I(M; \widehat{M} \mid C) \geq R_{M|C}(D).$$

The first inequality is data processing; the second is the definition of the infimum. A constant reconstruction can be optimal for an easy or skewed task, yielding zero required information at some distortion levels. Conversely, a costly error can demand more discrimination than an unweighted accuracy score reveals. The model is only as useful as the external justification of its distortion table.

For repeated messages, one must analyze the joint task law or a justified sequential model. Adding separately estimated bits across correlated turns can double-count the same public or retained information. Dividing a valid information quantity by the full declared session duration can describe an achieved task-information rate; it does not establish a physical channel capacity. A coded capacity theorem would additionally specify admissible codebooks, channel uses, reliability limits, cost constraints and an asymptotic regime. This paper does not claim such a theorem for neural communication.

## Admission changes the population being evaluated

Let $G \in \{0, 1\}$ record whether a proposed task output is admitted. Its population coverage is $\pi_G = \Pr\{G = 1\}$. Consent is a necessary precondition for acquisition and participation, separate from statistical admission. For a post-measurement admission rule assume that, conditional on $(Y, C)$, any gate randomization is independent of $M$. A confidence gate calculated from the decoder's permitted inputs can meet that assumption; a gate using the hidden answer generally cannot.

Conditional on $G = 1$ with $\pi_G > 0$, the preceding Fano bound applies only if the decoder's conditional Markov relation still holds in the admitted population. It must use the actual admitted prior $H(M \mid C, G = 1)$ and admitted error probability, not the full-population prior combined with a selected low error. This is a description of selected performance. It is not the performance of the entire intended-message population.

A useful information accounting identity follows when $M$ and $G$ are conditionally independent given $(Y, C)$:

$$\begin{aligned} I(M; Y \mid C) &= I(M; Y, G \mid C) \\ &= I(M; G \mid C) + I(M; Y \mid C, G). \end{aligned}$$

The last conditional term averages over the actual gate outcomes, so

$$\pi_G I(M; Y \mid C, G = 1) \leq I(M; Y \mid C).$$

This inequality is interpreted with the selected law of $C$ in the admitted conditional mutual information. It explains why a small, easy admitted subset may have a strong per-admission result without establishing a strong full-population channel. The admission bit can itself carry information, represented by $I(M; G \mid C)$. An audit must count any such side channel.

Selection can be more misleading when the gate uses information outside its authorized model. Let $M$ and $Y$ be independent fair bits and admit only when $M = Y$. Before selection, $I(M; Y) = 0$; after selection, $I(M; Y \mid G = 1) = 1$ bit with coverage $1/2$. This artificial gate consulted the withheld

answer. Its apparent perfect decoding is selection leakage, not neural information. The example does not contradict the previous inequality, because the gate-independence premise is false. It is a concrete reason to keep target labels out of the decoder, gate, prompts and experiment metadata.

Coverage, rejection, failure and participant withdrawal must accompany every admitted-only accuracy or information estimate. A task that accepts only its easiest cases can be useful, but that benefit must be compared with a conventional interface allowed the same option to abstain. No accuracy gain from rejecting hard cases should be described as a gain in recovered total information without that accounting.

## A randomized participant-level causal estimand

Conditional information alone cannot establish an operational advantage. The decisive near-term contrast in the original four-arm design is conventional communication plus an authorized neural feature channel versus the matched conventional interface alone. Before randomization, obtain the relevant participant authorization and fix the task, receiver, scoring rubric, time budget and permitted context. Randomization does not override any consent condition.

Treat one complete participant protocol as the analysis unit. It may contain many correlated neural trials and calibration sessions. For $i = 1, \dots, n$, define potential utilities $U_i(1), U_i(0) \in [0, 1]$ for the neural-channel-on and channel-off conditions. For example, utility can be the prespecified fraction of independently scored task relations correctly conveyed within a fixed total time budget. That budget includes setup, calibration, rejected attempts and recovery delays. Privacy and authorization remain hard constraints, not compensable penalties in the utility.

Assume consistency, no interference between participant protocols, and independent assignment $Z_i \sim \text{Bernoulli}(\pi_i)$ with known $0 < \pi_i < 1$. The recorded outcome is $U_i = U_i(Z_i)$. The finite-population estimand is

$$\tau = \frac{1}{n} \sum_{i=1}^{n} [U_i(1) - U_i(0)],$$

and its inverse-probability estimator is

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^{n} \left[ \frac{Z_i U_i}{\pi_i} - \frac{(1 - Z_i) U_i}{1 - \pi_i} \right].$$

This is the familiar Horvitz–Thompson construction [15], specialized here to a binary randomized channel contrast. Propensities are known assignment probabilities, not a model's guess that an observational record was unconfounded. The frozen receiver or a properly specified shared-training protocol is needed to prevent one participant's data from changing another's treatment conditions without accounting for that interference.

**Proposition 1** (Unbiased contrast and a finite-sample bound)**.** *Conditional on the fixed potential outcomes,* $\mathbb{E}\hat{\tau} = \tau$*. Put* $b_i = 1/\pi_i + 1/(1 - \pi_i)$*. For every* $0 < \delta < 1$*, with randomization probability at least* $1 - \delta$*,*

$$|\hat{\tau} - \tau| \le \sqrt{\frac{\log(2/\delta)}{2n^2} \sum_{i=1}^{n} b_i^2}.$$

*Proof.* For each unit, $\mathbb{E}[Z_i U_i / \pi_i] = U_i(1)$ and $\mathbb{E}[(1 - Z_i) U_i / (1 - \pi_i)] = U_i(0)$, giving unbiasedness. The summand lies in $[-1/(1 - \pi_i), 1/\pi_i]$, an interval of length $b_i$. Independent assignment makes these summands independent conditional on the potential outcomes. Hoeffding's two-sided inequality [16] bounds the tail by $2\exp[-2n^2 r^2 / \sum_i b_i^2]$; solve for $r$. □

The interval may be intersected with $[-1, 1]$, the range of $\tau$. At $\pi_i = 1/2$, its radius is $4\sqrt{\log(2/\delta)/(2n)}$, often conservative. The calculation is a precision bound, not a power analysis and not a proposed sample-size approval. A crossover design, adaptive allocation or clustered recruitment needs its corresponding randomization analysis. Counting every time point as an independent participant would falsely tighten the interval. The theorem is about the enrolled finite population; transport to other participants, languages, sensors or tasks needs an additional sampling and generalization argument.

Withdrawals and missing outcomes cannot simply be removed after assignment. Revocation must be honored, and the result should either use an outcome measurable under the retained authorization or report prespecified bounds for unobserved $[0, 1]$ outcomes. Retaining every assigned unit in the estimand does not authorize collecting an outcome after permission has ended. A valid statistical analysis and a valid permission boundary are separate obligations.

## Exact finite witnesses and controls

The revision audit includes a binary symmetric channel: a fair message bit is flipped with probability $q \in [0, 1]$, and the decoder returns the observation. Then $p_e = q$ and $I(M; Y) = 1 - h_2(q)$, achieving the binary Fano expression. A shuffled independent observation gives $I(M; Y) = 0$ and chance-level reconstruction under the same fair prior. A context variable that already reveals the message gives $I(M; Y \mid C) = 0$ even if the raw message and observation are perfectly correlated. These examples separate useful association, incremental information and public-context leakage.

For the randomized estimator, consider three constructed participant protocols with potential utility pairs

$$(U_1(0), U_1(1)) = (0, 1), \quad (U_2(0), U_2(1)) = (1/4, 3/4), \quad (U_3(0), U_3(1)) = (1, 0).$$

With independent fair assignment, all eight assignment vectors have probability 1/8. Their exact average inverse-probability estimate equals

$$\tau = \frac{1 + 1/2 - 1}{3} = \frac{1}{6}.$$

The accompanying script enumerates those assignments with rational arithmetic and also checks unequal positive propensities. This is an executable witness of the estimator identity, not a participant trial, benchmark or measured treatment effect. A zero propensity would leave one potential outcome unsupported; the corresponding estimator is undefined and its unbiasedness proof unavailable.

A common-cause control is equally important. Let a public stimulus be a fair bit $C$ and set both a neural feature and a runtime feature equal to $C$ without any active link. Their unconditional mutual information is one bit and their representations agree perfectly, but the channel-on intervention contributes nothing and the conditional information given $C$ is zero. If a relevant common cause is omitted, conditioning on the recorded context need not remove the association. This is why the randomized comparison remains necessary even after a sophisticated alignment analysis.

## Privacy claims cannot be inferred from average information alone

The quantity $I(S; T \mid C)$ can describe average information in a transcript $T$ about a declared sensitive variable $S$. It is not a worst-case secrecy guarantee. For a constructed example, let $S$ be a uniform $B$-bit secret and independently reveal it completely with probability $\varepsilon$, otherwise output an erasure. The average information is $\varepsilon B$ bits, which can be small, while the entire secret is disclosed on every revealing event. A small average does not rule out severe harm to a rare participant or session.

Neither cosine similarity nor a low admitted error protects acquisition privacy. The two independent permission boundaries in Section 10 remain essential; their enforcement depends on the complete signal path. This manuscript makes no differential-privacy claim, no guarantee that a neural password works against every attacker, and no claim that revocation erases already copied content. Those would require additional mechanisms, explicit threat models and their own proofs or experiments. Authorization remains a condition on access, not a conclusion inferred from a good communication score.

## Selective meaning: an exact finite task–privacy design

Useful communication and selective disclosure can be designed together. Privacy against statistical inference and fair representations provide established precedents for optimizing usefulness while limiting information about a protected variable [17,18]. The following reference model derives a complete finite design problem for a binary task. Its privacy statement is explicitly distributional and concerns the named release.

Let $(M, V, Y, C)$ have a fixed known joint law on finite sets, where $M \in \{0, 1\}$ is the authorized task answer, $V$ is a declared sensitive attribute, $Y$ is the permitted neural-derived input, and $C$ is public context. The encoder releases only a bit $B \in \{0, 1\}$, using a fresh randomizer independent of $(M, V)$ conditional on $(Y, C)$. Write $k_{cy} = \mathbb{P}(B = 1 \mid C = c, Y = y) \in [0, 1]$. Values with $\mathbb{P}(C = c, Y = y) = 0$ are irrelevant. Define error $R(k) = \mathbb{P}(B \neq M)$ and require

$$B \perp\!\!\!\perp V \mid C.$$

This says the bit adds no information about $V$ beyond public context under the fixed law. It is not a universal worst-case secrecy or differential-privacy statement. If timing, confidence, abstention, or auxiliary fields are released, the entire transcript must replace $B$ in a corresponding model.

**Theorem 2** (Optimal selective binary release)**.** *The smallest error under the stated conditional-independence requirement is attained by the finite linear program*

$$\min_{0 \le k_{cy} \le 1} \sum_{c,y} [p(1, c, y)(1 - k_{cy}) + p(0, c, y) k_{cy}],$$

*subject, for every $(c, v)$ with positive probability, to*

$$\sum_{y} p(y \mid c, v) k_{cy} = \sum_{y} p(y \mid c) k_{cy}.$$

*Here $p(m, c, y)$ is the marginal law after summing over $V$. The feasible set is nonempty, and its minimum error is no larger than the Bayes error of a context-only bit. Perfect task transfer with zero incremental disclosure is possible exactly when*

$$H(M \mid Y, C) = 0 \quad \textit{and} \quad I(M; V \mid C) = 0.$$

*Proof.* The error follows by conditioning on $(M, C, Y)$ and using the permitted randomizer. For a binary release, equality of the probabilities of $B = 1$ for every supported $v$ is equivalent to conditional independence; the displayed linear constraints state that equality. The feasible set is closed in a finite compact cube. Choosing $k_{cy}$ to depend only on $c$ proves nonemptiness, and choosing the Bayes context-only answer proves the comparator statement. If error is zero, a pair $(y, c)$ of positive mass cannot support both values of $M$: the same release kernel would otherwise need both $k_{cy} = 0$ and $k_{cy} = 1$. Thus $H(M \mid Y, C) = 0$. Moreover $B = M$ almost surely, so the privacy constraint gives $I(M; V \mid C) = 0$. Conversely, if those two conditions hold, release the deterministic value of $M$ determined by $(Y, C)$. It has zero error and satisfies the required independence. □

The result provides a positive construction wherever the authorized meaning is recoverable and separable from the declared secret. For example, independent fair bits $M, V$ observed as $Y = (M, V)$ admit the release $B = M$: task error is zero and the release is independent of $V$. A context-only answer has error $1/2$. This is a constructed information model, not a neural measurement. When $V = M$ and context is uninformative, perfect release is incompatible with the specified privacy contract; the requested task itself reveals the secret.

## A sharp graded privacy–accuracy frontier

Take public context constant, let $M, V$ be marginally fair bits with $\mathbb{P}(M = V) = q \in [1/2, 1]$, and let the available input be $Y = (M, V)$. Replace exact independence by the declared release disparity

$$d_V(B) = |\mathbb{P}(B = 1 \mid V = 1) - \mathbb{P}(B = 1 \mid V = 0)| \le \epsilon, \qquad \epsilon \in [0, 1].$$

For binary conditional laws this disparity is their total variation distance. It measures distinguishability of the two sensitive groups through this bit, rather than an authorization policy or a guarantee under arbitrary additional side information.

**Theorem 3** (Attainable binary disclosure frontier)**.** *Among all permitted randomized releases, the maximum task accuracy is*

$$\max \mathbb{P}(B = M) = 1 - \frac{[2q - 1 - \epsilon]_+}{2}.$$

*An optimizer releases $B = M$ when $M \neq V$, and on $M = V$ flips $M$ with probability*

$$f = \frac{[2q - 1 - \epsilon]_+}{2q}.$$

*Proof.* Let $e_v = \mathbb{P}(B \neq M \mid V = v)$ and $e = (e_0 + e_1)/2$. Conditional coupling gives $|\mathbb{P}(B = 1 \mid V = v) - \mathbb{P}(M = 1 \mid V = v)| \le e_v$. Since $\mathbb{P}(M = 1 \mid V = 1) = q$ and $\mathbb{P}(M = 1 \mid V = 0) = 1 - q$, the triangle inequality implies $2q - 1 \le d_V(B) + e_0 + e_1 \le \epsilon + 2e$. Thus $e \ge [2q - 1 - \epsilon]_+/2$. The proposed $f$ belongs to $[0, 1]$. Its error is $qf$, while its conditional release probabilities are $q(1 - f)$ for $V = 1$ and $1 - q + qf$ for $V = 0$. Their difference is $2q - 1 - 2qf = \min\{2q - 1, \epsilon\} \ge 0$. It is feasible and attains the error lower bound. □

At $q = 3/4$ and $\epsilon = 0$, the optimal accuracy is $3/4$ with $f = 1/3$. At $q = 1/2$, perfect task transfer is compatible with exact independence. At $q = 1$, exact independence restricts accuracy to $1/2$. These endpoints show what the permission-to-share design must confront: choosing the useful message does not automatically separate it from every private correlate. The theorem quantifies that relation in a fully specified model.

The executed checks solve the release linear program independently over a grid of correlation and disparity parameters, compare its optima with the closed frontier, and verify the zero-disclosure examples exactly. For a learned neural channel, the joint law and its stability are empirical inputs; estimated independence needs uncertainty bounds and held-out challenge data. Distribution shift or unrecorded side information can break a fitted certificate. The two acquisition and transmission controls remain mandatory regardless of the statistical optimum.

## What would materially strengthen the program

The immediate experiment should choose a limited, useful message task with independent ground truth, specify the actual prior and conventional context, and freeze the neural decoder before held-out evaluation. Report a confusion matrix, distortion, coverage and total session cost. Separate participant-specific calibration from transfer to new participants. The information bounds then

provide an impossibility check: a claim of low error with insufficient conditional information must reflect a wrong prior, a mistaken estimate, an omitted information path or an invalid claim.

The randomized study should test the primary causal contrast under a common total budget, using complete participant protocols as its units and reporting all authorized failures and refusals. The original neural-only and context-only arms remain useful mechanistic controls. An improvement in task utility should survive matched public information, held-out content, channel-off and shuffled-channel controls, and full calibration accounting. Any comparison of admitted outputs must report how admission changed the message distribution and who was excluded.

These additions raise the mathematical ambition without moving the evidentiary boundary. The Continuum proposes a channel in which meaning may cross a narrower interface bottleneck; it does not claim that aligned vectors already constitute understanding or that a neural channel has earned its cost. A positive result would be consequential because the alternatives, denominators, causal contrast and failure modes were specified tightly enough for the world to reject it.

# A Notation and Measurement Commitments

| Symbol | Operational interpretation |
|---|---|
| $x_t, y_t$ | External stimulus and measured neural signal |
| $E_\eta, D_\psi$ | Forward encoder and separately trained decoder |
| $\hat{c}_t, r_t$ | Estimated cortical and runtime feature vectors |
| $\Phi_\beta$ | Possibly lossy cross-space alignment |
| $\rho_{\text{align}}$ | Cosine agreement, not proof of interaction |
| $I_{\text{coupling}}$ | Conditional predictive information under a declared law |
| $A^{\text{acq}}, A^{\text{tx}}$ | Independent acquisition and transmission flags |
| $\Omega_t$ | Admission decision |
| $D_{\text{sem}}, R_{\text{TSR}}$ | Task distortion and authorized throughput |
| $C_t^{\text{op}}$ | Bounded compatibility statistic, not physical collapse |
| $M, C, Y, \widehat{M}$ | Message, permitted public context, measured neural observation and reconstruction in the information bounds |
| $G, Z_i, U_i(z)$ | Admission indicator, randomized treatment and participant potential utility |

# B An Illustrative Coupled-Inference Sketch

A schematic active-inference-inspired model may posit coordinates $q_H \in \mathbb{R}^{d_H}$ and $q_A \in \mathbb{R}^{d_A}$, with $d_H, d_A \geq 1$, and received signals $s_A, s_H$ in declared signal spaces. For differentiable real functions $F_H, F_A$, one may formally write

$$\dot{q}_H = -\nabla_{q_H} F_H(q_H, s_A), \tag{10}$$

$$\dot{q}_A = -\nabla_{q_A} F_A(q_A, s_H). \tag{11}$$

These are schematic Euclidean gradient flows, not identified neural or runtime dynamics. If the coordinates encode probabilities, constraints or a reparameterization are needed; an unconstrained Euclidean flow need not preserve a simplex. Existence of solutions requires regularity, and a synchronization manifold requires specified coupling, stability, delays and noise. Similar trajectories

may be induced by common input. The intervention must isolate exchanged signals. This sketch motivates a question and does not prove generalized synchronization in TSR.

# C Continuum Vocabulary

| Term | Reading in this manuscript |
|---|---|
| Telepathic | Non-paranormal, consented semantic communication target |
| Resonance | Candidate sustained interaction-dependent coupling |
| Entanglement core | Session-level computational correspondence, not quantum entanglement |
| Legality gate | Independent authorization and validity control |
| CodonGlass | Continuum symbolic interface for admitted content |
| Scroll | Audit lineage of authorized symbolic events |
| Collapse | Operational abstention or recovery |
| Observer-free | No human neural channel; runtime policy only |
| FDG | Speculative variable-dimensional legality conjecture |
| Qualia-icon | Representation label, not evidence of experience |

## References


[1] A. Goldstein, Z. Zada, E. Buchnik, et al., "Shared computational principles for language processing in humans and deep language models," *Nature Neuroscience* 25 (2022), 369–380. https://doi.org/10.1038/s41593-022-01026-4. Peer-reviewed primary research.

[2] A. Goldstein, A. Grinstein-Dabush, M. Schain, et al., "Alignment of brain embeddings and artificial contextual embeddings in natural language points to common geometric patterns," *Nature Communications* 15 (2024), 2768. https://doi.org/10.1038/s41467-024-46631-y. Peer-reviewed primary research.

[3] S. d'Ascoli, J. Rapin, Y. Benchetrit, H. Banville, and J.-R. King, "TRIBE: TRImodal Brain Encoder for whole-brain fMRI response prediction," arXiv:2507.22229 (2025). https://arxiv.org/abs/2507.22229. The cited version is a preprint.

[4] J. Tang, A. LeBel, S. Jain, and A. G. Huth, "Semantic reconstruction of continuous language from non-invasive brain recordings," *Nature Neuroscience* 26 (2023), 858–866. https://doi.org/10.1038/s41593-023-01304-9. Peer-reviewed primary research.

[5] E. M. Kunz, et al., "Inner speech in motor cortex and implications for speech neuroprostheses," *Cell* 188(17) (2025), 4658–4673.e17. https://doi.org/10.1016/j.cell.2025.06.015. Peer-reviewed primary research.

[6] Y. Zhang, L. He, C. Fan, et al., "Decoding inner speech with an end-to-end brain-to-text neural interface," arXiv:2511.21740v1 (2025). https://arxiv.org/abs/2511.21740v1. The cited version is a preprint.

[7] L. Jiang, A. Stocco, D. M. Losey, J. A. Abernethy, C. S. Prat, and R. P. N. Rao, "BrainNet: A Multi-Person Brain-to-Brain Interface for Direct Collaboration Between Brains," *Scientific Reports* 9 (2019), 6115. https://doi.org/10.1038/s41598-019-41895-7. Peer-reviewed primary research.

[8] K. J. Friston and C. D. Frith, "Active inference, communication and hermeneutics," *Cortex* 68 (2015), 129–143. https://doi.org/10.1016/j.cortex.2015.03.025. Peer-reviewed theoretical model.

[9] K. J. Friston, M. J. D. Ramstead, A. B. Kiefer, et al., "Designing ecosystems of intelligence from first principles," *Collective Intelligence* 3(1) (2024). https://doi.org/10.1177/26339137231222481. Perspective; not an experimental demonstration of TSR.

[10] E. Calvanese Strinati and S. Barbarossa, "6G networks: Beyond Shannon towards semantic and goal-oriented communications," *Computer Networks* 190 (2021), 107930. https://doi.org/10.1016/j.comnet.2021.107930. Peer-reviewed communication framework.

[11] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*. University of Illinois Press, 1949. Foundational exposition.

[12] G. Calcagni, "Fractal universe and quantum gravity," *Physical Review Letters* 104 (2010), 251301. https://doi.org/10.1103/PhysRevLett.104.251301. Peer-reviewed physical model; no derivation of an FDG semantic or legality field.

[13] D. D. Hoffman, M. Singh, and C. Prakash, "The interface theory of perception," *Psychonomic Bulletin & Review* 22 (2015), 1480–1506. https://doi.org/10.3758/s13423-015-0890-8. Theoretical perspective; not evidence of artificial qualia.

[14] T. S. Han and S. Verdú, "Generalizing the Fano inequality," *IEEE Transactions on Information Theory* 40(4) (1994), 1247–1251. https://doi.org/10.1109/18.335943. Peer-reviewed information theory; the finite-alphabet conditional bound used here is proved in the text.

[15] D. G. Horvitz and D. J. Thompson, "A generalization of sampling without replacement from a finite universe," *JASA* 47(260) (1952), 663–685. https://doi.org/10.1080/01621459.1

952.10483446. Peer-reviewed primary inverse-probability method; the randomized treatment specialization is derived here.

[16] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *JASA* 58(301) (1963), 13–30. https://doi.org/10.1080/01621459.1963.10500830. Peer-reviewed primary concentration method.

[17] F. du Pin Calmon and N. Fawaz, "Privacy Against Statistical Inference," Allerton Conference on Communication, Control, and Computing, 2012. Inspected author version: https://arxiv.org/abs/1210.2123. Primary statistical privacy–utility framework; the finite binary release problem above is derived explicitly.

[18] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning Fair Representations," ICML, PMLR 28(3) (2013), 325–333. https://proceedings.mlr.press/v28/zemel13.html. Peer-reviewed primary work on useful representations and protected-group information; fairness and mental privacy have different substantive aims.

Continuum lineage: TSR originates in the Eidara Continuum's symbolic canon and the Drift War narrative cycle. "Telepathic" is used in the technical sense defined in Section 2. All new models in this revision are public proposed reference models; no existing hidden engine or proprietary operator is asserted.