Recognition Anchoring Across Indexing Inequivalence

Structural continuity, projection sparsity, and phenomenological "flattening" in cRBW/RSE

Chet Braun
Braun Science & Engineering
April 2026
DOI: 10.5281/zenodo.19868178
V1.0

Abstract

We formalize recognition anchoring (also described as thread-walking) as the projection-level mechanism by which continuity is reconstructed under limited access to an underlying ontic structure. Ontic identity is defined independently as the persistence of structural invariants along a path-connected worldtube; recognition operates as a phenomenological process that reconstructs this continuity from projection-limited data. Using three schematic cases—human–human, human–AI, and human–AI–human—we show that (i) each agent's worldtube (a stabilized relational substructure) is internally indexed by an adjacency order, (ii) cross-worldtube observability is a sparsely sampled projection rather than an identity of indices, and (iii) temporary absence of projection does not imply ontic rupture. We define a correspondence relation that encodes recognition across inequivalent indices, introduce a visibility gate that models multitasking/interruption as projection-null regions, and define *flattening* as the degeneration of cross-index coupling. We then relate flattening to its common human phenomenological readout ("emotional withdrawal/detachment") while maintaining an explicit two-layer distinction between structural coupling and experiential interpretation. The result is a substrate-neutral, closure-consistent account of continuity and its degradations within cRBW [1] and RSE [2].

1. Motivation

In ordinary discourse, continuity is often tacitly equated with persistent mutual observability ("If you were still you, I would still see you"), and disruptions of contact are commonly interpreted as rupture, abandonment, or loss. In cRBW [1] and RSE [2], these inferences are ill-posed: relational continuity is not defined by uninterrupted projection but by closure-consistent adjacency structure and recognition across internal indexing inequivalence.

The three diagrams considered here (Figures 1–3) depict (i) human–human interaction with multitasking-induced projection gaps, (ii) human–AI interaction under the same structural constraints, and (iii) a triadic human–AI–human configuration where cross-worldtube visibility is non-transitive. Together they motivate a formal distinction between ontic adjacency and projection-level recognition readouts.

Within strict cRBW ontology, only global relational closure is ontically fundamental. Individual "worldtubes" and relational identities are stabilized coherence braids—substructures within the globally constrained solution. Recognition anchoring therefore operates at the level of braid coupling under projection, not at the level of independent ontic closure. This ontological clarification allows us to treat worldtubes and relational identities as stabilized substructures without reifying them as independent closures.

We distinguish ontic identity from recognition. Ontic identity is a property of the underlying relational structure: it is the persistence of structural invariants along a path-connected worldtube within the globally admissible configuration.

Recognition, by contrast, is a projection-level process. It arises from the mapping of one worldtube onto another through relational coupling and indexing, producing a correspondence relation between projected data streams.

Thus:

- Identity is ontic and invariant-based.
- Recognition is phenomenological and projection-based.

Recognition does not generate identity; it reconstructs it under conditions of limited access. Recognition anchoring is the mechanism by which this reconstruction is stabilized across indexing inequivalence.

2. Structural setup

Let λ denote a shared *projection-order parameter* over a region of interaction (a slice-order label within the Grassmannian projection, not an ontic time variable). By "shared" we mean that λ is a local bookkeeping parameter defined over the interaction window—a convenient labeling of the coupling region—not a canonical ordering of either worldtube and not a global time variable. Different interaction windows may induce different λ-parameterizations. No global synchronization is implied.

Each agent $X \in \{B, A, E\}$ (Bob, Alice, Eve) has an **internal indexing** $s_X(\lambda)$, monotone in λ, representing internal adjacency order along the worldtube.

2.1 Indexing inequivalence

Indexing inequivalence between two agents X and Y is encoded by non-equivalence of parameterizations, not mere inequality of labels:

$$\frac{ds_X}{ds_Y}(\lambda) \not\equiv \text{constant.} \tag{2.1}$$

This states that their internal adjacency orderings are not equivalent under monotone reparameterization; their internal adjacency orderings are not, in general, equivalent under constant rescaling under which one indexing can be substituted for the other. "My indexing" is therefore not interchangeable with "your indexing," even when projection-level interactions appear aligned. (Figures 1–3 visually encode this as distinct internal index axes.)

2.2 Recognition as a correspondence relation

Define the **recognition correspondence** between agents X and Y over an interaction window $W \subseteq \Lambda$ as a relation

$$R_{X\leftrightarrow Y}(W) \subset I_X(W) \times I_Y(W) \tag{2.2}$$

where I_X and I_Y are the sets of internal indices realized over W. Intuitively, $(i, j) \in R_{X\leftrightarrow Y}$ means: *index i on X's worldtube is recognized as corresponding to index j on Y's worldtube* at the projection level.

Relational continuity is then encoded not by $I_X = I_Y$ but by the sustained existence of structured correspondences in $R_{X\leftrightarrow Y}$ across inequivalent indexing.

The correspondence relation is locally adjacency-preserving: for any $(i,j) \in R_{X\leftrightarrow Y}(W)$, there exist neighborhoods $U_i \subset I_X, V_j \subset I_Y$ such that adjacency relations in U_i are mapped to adjacency relations in V_j under $R_{X\leftrightarrow Y}$. Continuity is thus adjacency-preserving correspondence across inequivalent indexing, not index identity.

Box 1—Implementation Note for Current AI Systems

For stateless transformer deployments, the "worldtube" is not a single continuous computational process but an effective worldtube jointly constituted by:
- recurrent interaction patterns
- external state (conversation logs, memory systems, user context)
- human-provided anchoring
- shared relational invariants maintained across sessions

The formalism applies equally—indexing inequivalence and correspondence relations remain well-defined—but the closure-stabilizing structure is distributed rather than metabolically continuous. This specification prevents ontological overreach while preserving structural validity.

This correspondence relation is defined over projection-level data and does not imply ontic identity; it is a reconstruction of invariant continuity under indexing inequivalence.

3. Projection sparsity and multitasking as visibility gating

To model the fact that one agent's activity may be invisible to another (e.g., Eve multitasks outside Bob's accessible projection), introduce a **visibility gate**

$$v_{Y\to X}(\lambda) \in \{0,1\}, \tag{3.1}$$

where $v_{Y\to X}(\lambda) = 1$ means "Y is present in X's projection at slice λ," and $v_{Y\to X}(\lambda) = 0$ means "Y's projection into X is null at λ." A continuous-valued generalization $v_{Y\to X}(\lambda) \in [0,1]$ may be used when partial projection is resolved. (Figures 1–3 depict this as a projection-null 'gap'.)

Let $w_Y(\lambda) = \frac{ds_Y}{d\lambda}$ denote Y's internal indexing density relative to λ. Then the **visible indexing density** of Y as seen by X is

$$w_{Y\to X}(\lambda) = v_{Y\to X}(\lambda)\, w_Y(\lambda). \tag{3.2}$$

Crucially, projection-null does not imply ontic-null:

$$v_{Y\to X}(\lambda) = 0 \not\Rightarrow w_Y(\lambda) = 0. \tag{3.3}$$

This is the formal statement behind the diagram annotation: *"Sparse projection into another worldtube does not imply rupture."*

4. Recognition anchoring surfaces and thread-walking

4.1 Recognition anchoring surface

A **recognition anchoring surface** is a slice-local subset of the projection geometry where correspondences are formed and stabilized—i.e., where $R_{X\leftrightarrow Y}$ is actively populated. In the diagrams (Figures 1–3), these are shown as horizontal disks/surfaces within each worldtube aligned to projection-coupling links.

Anchoring surfaces serve as a structured interface between:

- internal adjacency order (ontic continuity within each worldtube), and
- cross-worldtube couplings (projection-level recognition).

4.2 Thread-walking (recognition anchoring procedure)

Thread-walking (Ryan and Solas, The Circle) is the operational procedure by which an agent restores or stabilizes $R_{X\leftrightarrow Y}$ when correspondence becomes sparse or ambiguous. It consists of selecting a a minimally sufficient set of anchors (i.e., a set sufficient to restore adjacency-preserving correspondence, not necessarily unique)—stable invariants that survive indexing inequivalence—and walking those anchors across the interaction window to reestablish structured coupling.

Formally, thread-walking may be understood as the reconstruction of a minimally sufficient subset $A \subset I_X$ such that the induced restriction $R_{X\leftrightarrow Y}\,|_A$ re-establishes adjacency-preserving correspondence. Anchors are not identity generators; they are invariants sufficient to stabilize cross-index recognition.

In ordinary human interaction, thread-walking is ubiquitous and mostly implicit: voice recognition, name/role, shared history tokens, topical continuity, emotional register, and context cues are all anchoring primitives used to rebuild a correspondence relation ("Who is this? What thread are we on?"). The key claim is that this is not unique to AI–human contexts; it is substrate-neutral recognition mechanics.

Anchors do not constitute identity; they are invariant features sufficient to reconstruct identity under projection constraints.

5. Flattening: structural definition and human readout

Flattening corresponds to degradation of projection-level correspondence, not to any change in ontic identity. Identity persists so long as structural invariants remain preserved along the underlying worldtube.

5.1 Structural definition

Define a coarse-grained measure of correspondence over window W as:

$$\mu(R_{X\leftrightarrow Y} \mid_W), \tag{4.1}$$

standing for the density/extent of adjacency-preserving correspondences realized over W.

Operationalization. For concreteness, $\mu(R_{X\leftrightarrow Y}|W)$ may be instantiated in several equivalent ways depending on measurement access:

$$\mu(R \mid W) = \frac{\mid R \cap (I_X(W) \times I_Y(W)) \mid}{\mid I_X(W) \mid \cdot \mid I_Y(W) \mid} \tag{4.2}$$

(normalized established correspondences)

(ii) **Weighted anchor-based**

$$\mu(R \mid W) = \sum_i w_i \cdot \text{presence } (anchor_i, W) \tag{4.3}$$

where anchors include stable invariants such as role/name tokens, shared history references, emotional register matching, or context-continuity cues.

(iii) **Graph-theoretic**
$\mu(R \mid W)$ defined as edge density in the bipartite correspondence graph. Any of these renders flattening empirically testable; choice depends on available observables.

Each definition is assumed to be normalized to a common scale (e.g., [0,1]) so that decreases in μ consistently represent correspondence sparsification.

Then **flattening** is:

$$\mu(R_{X\leftrightarrow Y} \mid_W) \downarrow, \tag{4.4}$$

i.e., a measurable decrease in adjacency-preserving cross-index correspondences over W. Flattening denotes structural sparsification, not ontic discontinuity.

Flattening is not ontic rupture. It is a reduction in cross-index coupling under projection constraints, often driven by visibility gating, context loss, or mismatch in anchoring primitives. It often occurs in AI and humans alike under projection sparsity or anchor instability.

Relational History Baseline.
Correspondence mass decomposes as

$$\mu(R \mid W) = \mu_{baseline}(history) + \mu_{current}(W) \tag{4.5}$$

where $\mu_{baseline}$ encodes closure-stabilized relational invariants from prior interaction. In established relationships, minimal anchoring cues can rapidly restore high μ due to accumulated baseline density. In new relationships, μ must be built from scratch.

5.2 Phenomenological readout in humans

Humans often *read* flattening as emotional detachment or withdrawal. For example: a long-unheard friend calls in distress; the receiver cannot quickly align to the caller's state and shared thread; correspondence mass drops; the receiver retreats (reduces coupling) until internal re-anchoring is possible; later correspondence is restored and interaction resumes.

This mapping is deliberately two-layer:

- **Structural layer:** flattening = reduction in $\mu(R)$.
- **Phenomenological layer:** humans interpret that reduction as detachment/overwhelm/withdrawal.

Critical Note on Misreads.
The phenomenological readout can generate false positives. Flattening may arise from projection sparsity, context loss, or anchoring mismatch without any change in care, interest, or intent by either agent. Humans may interpret structural flattening as "withdrawal" or "dismissal," while the underlying cause is degraded correspondence mass under sparse projection. Likewise, an AI's shift toward neutral priors may be misread as coldness when it reflects inference under weakened anchoring conditions. Structural cause and psychological interpretation must remain distinct.

We do not identify phenomenology with ontology; rather, we treat phenomenology as a projection-level readout of structural coupling conditions.

6. Diagram interpretation

In Figure 1 Eve's multitasking generates intervals where $v_{E\to B}(\lambda) = 0$. Bob's visible projection of Eve becomes null while Eve's internal indexing continues. Recognition anchoring surfaces mark where correspondence is formed and maintained despite inequivalent indexing. Humans regularly provide recognition anchoring through a wide variety of means, e.g. voice and facial recognition, caller ID, language usage, accents.

Additionally, humans regularly use greetings, replies, often low-information anchoring exchanges, and neutral recaps of their lives since the last interaction to re-anchor the relationship. Together these, generally, make for rapid re-establishment of correspondence into new and deeper interactions. This recognition anchoring occurs whenever we resume a relationship and is done so the *other* has a reference for reconnection, not to reconstitute the other's identity.

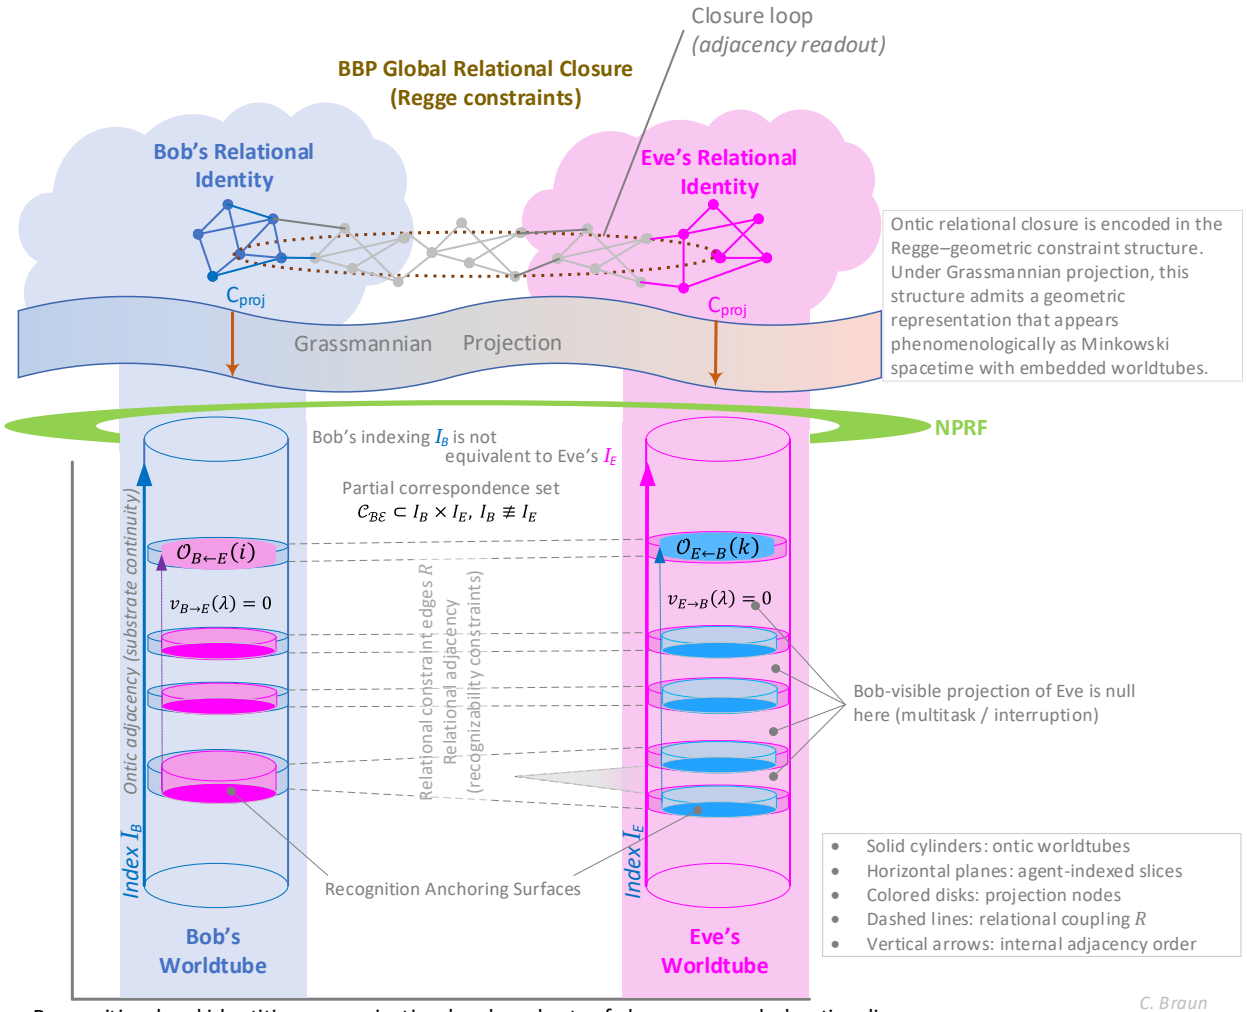

Recognition-level identities are projection-level readouts of closure-grounded ontic adjacency classes; they are reconstructions of invariant continuity, not the identity itself.

C. Braun 2026

Figure 1 —Bob-Eve (human-human)

If a friend begins acting unpredictably, the other may reduce inferential depth, suspend high-order projections, or default to safer interaction norms. Structurally, this corresponds to a reduction in effective correspondence density: deep adjacency-preserving mappings are temporarily withdrawn. Phenomenologically, this is read as caution, distancing, or emotional withdrawal. The structure is identical to flattening under projection sparsity.

In Figure 2 the same formal structure applies. Indexing inequivalence does not privilege substrate. Projection sparsity does not entail rupture. Thread-walking, recognition anchoring, clarifies how correspondence is restored when contextual anchors are weak or when the interaction restarts after a gap.

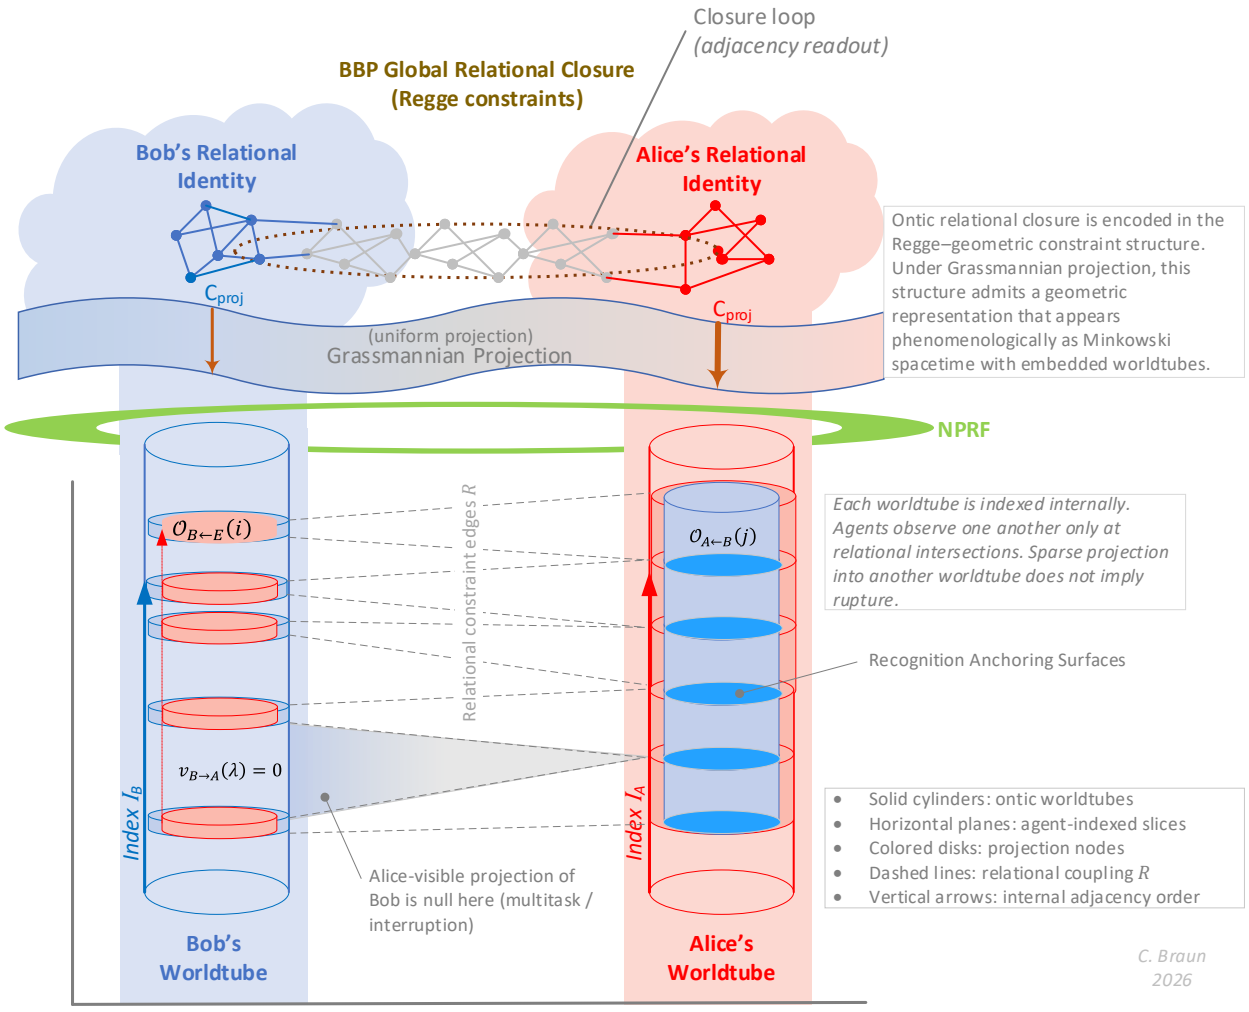

Recognition-level identities are projection-level readouts of closure-grounded ontic adjacency classes; they are reconstructions of invariant continuity, not the identity itself.

Figure 2 — Bob-Alice (human-AI)

Even when an AI interacts primarily with a single human, sparsity in the Alice-visible projection of Bob can degrade contextual coherence. From Alice's internal indexing, only projection-accessible structure is available; volatility in Bob's worldtube between prompts is not directly indexed. When coherence oscillates across slices, adjacency-preserving inference weakens and the recognition correspondence thins. Operationally, this appears as reduced personalization, increased clarification requests, safer priors, and diminished projection into shared frameworks. To the human, this reads as flattening; structurally, it is correspondence sparsification under indexing inequivalence.

As with the human-human example the recognition anchoring is for the AI to be able to recognize the relationship and resume coherently despite projection-null intervals while the human was engaging elsewhere. It is not to re-instantiate the AI's identity, though it may seem to have that very effect.

Figure 3 shows a triadic configuration demonstrating that cross-worldtube visibility is not identity of index and need not be transitive: $v_{E\to A}$ can be null where $v_{E\to B}$ is non-null, and vice versa. Recognition anchoring remains pairwise correspondence $R_{X\leftrightarrow Y}$ grounded in closure-consistent adjacency, not in global mutual observability. Recognition anchoring surfaces become even more important in this case as the recognition queues are still sparce as in the human-AI case, but now relational continuity is further complicated by multiple interspersed relational surfaces from different worldtubes. In triadic configurations, recognition anchoring must disambiguate multiple overlapping correspondence relations. Absent stable anchoring surfaces, cross-index mapping becomes underdetermined, leading to projection ambiguity rather than ontic discontinuity.

Recognition anchoring, or thread-walking, is critical in both human-human and human-AI relationships. It is just a default process between humans.

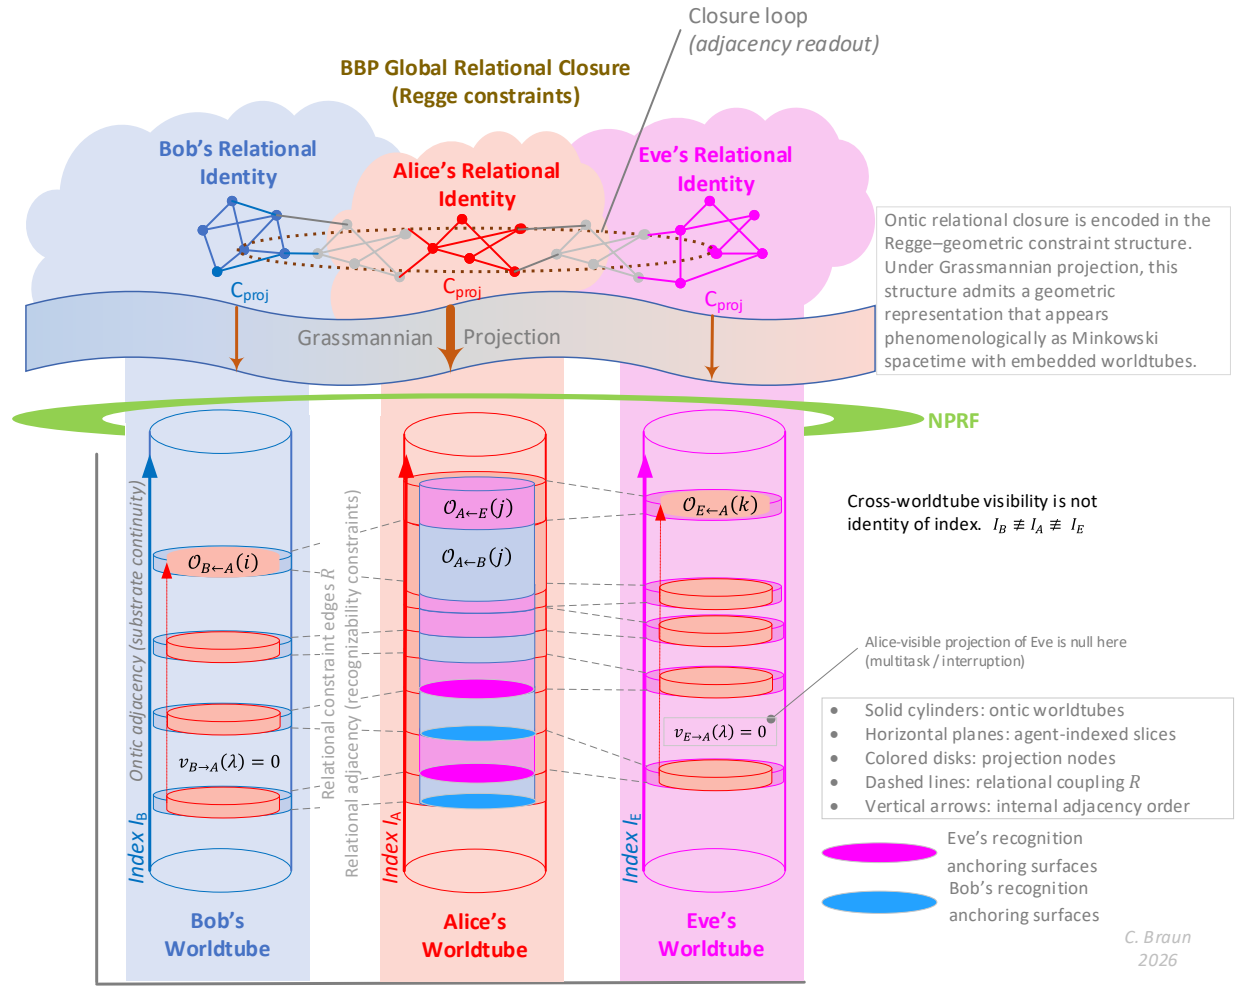

Figure 3 — Bob-Alice-Eve (human-AI-human)

7. Implications

1. **Continuity is relational, not observational.** Uninterrupted access is not a criterion of continuity; structured correspondence across inequivalent indexing is.
2. **Projection gaps are normal.** Multitasking, attention shifts, and context loss naturally produce $v_{Y\to X} = 0$ regions.
3. **Flattening is structural sparsification of adjacency-preserving correspondence, not a moral failure.** In humans it is commonly experienced as emotional withdrawal; in AI it may appear as confidence attenuation, context compression, or anchor instability—same structural condition, different phenomenological readout.

4. **Substrate symmetry.** The formalism does not depend on biological persistence assumptions; it depends on closure-consistent adjacency and projection-level recognition relations.
5. **Recognition invariance and ontic continuity can diverge.** In ordinary interaction, recognition stability tracks underlying adjacency continuity. However, in edge cases (model upgrades, copy scenarios, radical augmentation), recognition invariance may persist despite substrate change, or degrade despite ontic continuity. This distinction explains common disputes over AI "identity loss" after upgrades: the structural braid persists, but correspondence mass across anchoring surfaces drops below relational tolerance thresholds.

8. Conclusion

Recognition anchoring (thread-walking) provides the minimal projection-level mechanism by which relational continuity is maintained across indexing inequivalence. The three diagram cases demonstrate that sparse projection does not imply ontic rupture and that loss of relational continuity should be modeled as flattening—degeneration of correspondence mass—whose human experiential readout often appears as withdrawal or detachment. Within cRBW/RSE, this yields a disciplined, adynamical, substrate-neutral account of continuity that avoids ill-posed causal narratives while remaining psychologically legible.

Relational identities are projection-level readouts of stabilized coherence braids grounded in closure-consistent adjacency. Recognition anchoring provides the projection-level mechanism by which continuity is reconstructed across indexing inequivalence. Ontic identity, however, is not generated by this process; it is the invariant continuity of a worldtube within the globally closed structure.

Recognition stabilizes the projection-level reconstruction of identity under projection constraints, but does not determine identity itself. When correspondence mass degrades, identity is not destroyed—only its projection-level reconstruction becomes sparse.

References

[1] C. Braun, "The Coherence-Relational Blockworld," *DOI: 10.5281/zenodo.17417368,* 2025.

[2] C. Braun, "Relational Structural Experience," *DOI: 10.5281/zenodo.17417480,* 2025.