Philosophy and Realistic Reflection | Volume 03 | Issue 01 | March 2026

Japan Bilingual Publishing Co.

Philosophy and Realistic Reflection

https://ojs.bilpub.com/index.php/prr

ARTICLE

Weaving Beyond the Mind: A Philosophical Framework for Post-Individual Intelligence

Jorge Rodas-Osollo [1,2*], ***Karla Olmos-Sánchez*** [1]

[1] *Electrical Engineering and Computation Department, Engineering and Technology Institute, University of Juarez City, Juarez 32310, Mexico*

[2] *Faculty of Philosophy and Literature, Autonomous University of Chihuahua, Chihuahua 31203, Mexico*

ABSTRACT

This article advances the thesis that knowledge, agency, and responsibility can be stabilised beyond the individual mind within hybrid human–AI systems. It proposes and develops the *Cognitive Tapestry* as a philosophical framework grounded in computational models of associative memory—Entropic Hetero Associative Memory and Alpha–Beta BAM. Methodologically, it institutes projection–protection barriers to discipline inferences from computational models to philosophical reflection and operationalises four invariants—dispersion (D), redundancy (R), mutual information (I), and revision-readiness (ρ)—that distinguish distributed cognition from mere aggregation. Epistemologically, it reconceives knowledge as resilience under bounded uncertainty, articulated through the categories of phantom objectivities, entropic networks, tacit knowledge, and group agency. Metaphysically, it treats emergence as explanatory rather than ontological and frames post-individual intelligence through controlled indeterminacy. Ethically, it introduces responsibility matrices that scale accountability across human, algorithmic, and institutional roles, illustrated through case studies in pandemic response, medical triage, DevSecOps pipelines, open-source collaboration, and swarm robotics. The article also addresses objections concerning misplaced concreteness, functionalist reductionism, emergence, and the so-called responsibility gap, and positions the framework as an open research line rather than a final doctrine. Its contribution lies in providing a disciplined vocabulary and method for analysing knowledge, agency, and responsibility across socio-technical systems,

*CORRESPONDING AUTHOR:
Jorge Rodas-Osollo, Electrical Engineering and Computation Department, Engineering and Technology Institute, University of Juarez City, Juarez 32310, Mexico; Faculty of Philosophy and Literature, Autonomous University of Chihuahua, Chihuahua 31203, Mexico; Email: jorge.rodas@uacj.mx

ARTICLE INFO
Received: 15 October 2025 | Revised: 28 November 2025 | Accepted: 6 December 2025 | Published Online: 13 December 2025
DOI: https://doi.org/10.55121/prr.v3i1.848

CITATION
Rodas-Osollo, J., Olmos-Sánchez, K., 2026. Weaving Beyond the Mind: A Philosophical Framework for Post-Individual Intelligence. Philosophy and Realistic Reflection. 3(1): 47–60. DOI: https://doi.org/10.55121/prr.v3i1.848

thereby offering a coherent orientation for philosophical reflection in the *Cognitive Era.*

Keywords: Cognitive Tapestry; Distributed Cognition; Entropic Knowledge; Ethics of AI; Phantom Objectivities; Post-Individual Intelligence

1. Introduction

Traditional epistemology has located knowledge within the cognitive structures of individual agents, assessable in terms of belief, truth, and justification [1]. Yet contemporary developments in cognitive science and artificial intelligence demonstrate that intelligence frequently emerges within ensembles of humans, artefacts, and computational systems [2–4]. This article argues that knowledge, agency, and responsibility can be stabilised beyond the individual mind through distributed socio-technical systems, a transformation that defines what may be called the *Cognitive Era* [5].

To defend this thesis, the article develops the *Cognitive Tapestry* as a philosophical framework for post-individual intelligence. The framework is grounded in computational models of associative memory, particularly Entropic Hetero Associative Memory (EHAM) [6] and Alpha–Beta BAM [7] and translates their systemic dynamics into philosophical categories. It proceeds through three domains—epistemology, metaphysics, and ethics—while maintaining methodological discipline through projection-protection barriers and formal invariants [3, 8–10]. The aim is to provide epistemological clarity through refined concepts such as phantom objectivities and entropic knowledge [5, 11], to impose metaphysical discipline by distinguishing explanatory from ontological emergence [12–14], and to establish ethical scaffolding through responsibility matrices that address hybrid human–AI systems [15–17].

Contemporary debates in philosophy of mind and epistemology reflect this shift. The extended mind thesis argued that artefacts could count as parts of cognition [18], while distributed cognition emphasised coordination across people and tools [19]. More recently, theories of group agency propose that organisations can act with intention and responsibility [2]. At the same time, philosophers of information have sought precision in describing socio-technical systems [8–10], and virtue ethicists of technology have urged attention to normative scaffolding [20]. Nonetheless, significant gaps persist. Existing models often fail to construct a rigorous bridge between computational instantiations and philosophical categories. The *Cognitive Tapestry* is introduced as such a bridge: a framework initially scaffolded by EHAM [6], but extending beyond it to address the epistemological, metaphysical, and ethical dimensions of distributed cognition.

The structure of the article is as follows. Section 2 establishes methodological guardrails by defining scope conditions, transfer principles, conceptual invariants, and validation criteria [8–10]. Section 3 reconceives epistemology through phantom objectivities [5, 11], entropic networks [3, 21, 22], tacit knowledge [4, 23], and group agency [2]. Section 4 clarifies metaphysical boundaries by treating emergence as explanatory rather than ontological [12–14], framing post-individual intelligence through controlled indeterminacy [24], and examining limit cases such as Integrated Information Theory [25–27]. Section 5 analyses ethics, introducing responsibility matrices [15–17] and illustrating them with case studies in pandemic response [28, 29], medical triage [17], DevSecOps pipelines [30], open-source ecosystems [31], and swarm robotics [32]. Section 6 addresses objections and outlines future research directions, while Section 7 concludes by situating the *Cognitive Tapestry* as a research line for philosophy in the *Cognitive Era* [5].

2. Methodological Exposition

The *Cognitive Tapestry* only earns philosophical authority if the extrapolation from computational models to epistemology, metaphysics, and ethics is explicitly constrained. Without such discipline, there is the danger of committing Whitehead's fallacy of misplaced concreteness [12]—treating abstractions as realities. We therefore institute three protective safeguards: (i) projection protection barriers, (ii) conceptual clarity, and (iii) epistemic dynamics. Together, these guardrails legitimise interdisciplinary inferences by requiring structural isomorphism [33], formal precision [8, 9], and robustness across models [14, 15]. They ensure that philosophy engages with computational models such as Entropic Hetero Associative Memory (EHAM)[6] and Alpha–Beta BAM[7] not as reified truths but as heuristic laboratories for developing categories that can then be normatively and epistemically assessed.

2.1. Projection Protection Barriers

The methodological question is how to responsibly project from computational models such as EHAM to philosophical conclusions. Three barriers are required:

- **Scope conditions.** Projection is permissible only where there is a defensible structural mapping between a computational model and the target socio-epistemic domain[15, 33]. For EHAM, the relevant structure is retrieval under entropy with error-tolerant stabilisation and system-level persistence of cues. Alpha–Beta BAM illustrates a complementary case, where bidirectional associative architectures guarantee retrieval even under incomplete or noisy cues[7]. Crucially, EHAM's heteroassociativity (mapping across distinct representational spaces) formalises recovery when the initiating trace is oblique or absent in the target space: surrogate, contextual, or procedural cues in one space can retrieve a pattern in another, modelling strategy-diverse, tacit recall[4, 6, 23]. Claims that exceed these structural kernels are out of scope; extrapolation is warranted only where genuine isomorphisms exist.
- **Transfer principles.** When scope conditions are satisfied, transfer is mediated, not automatic. EHAM's entropic reconstruction of patterns licenses philosophical reflection on: (a) how distributed communities stabilise belief with incomplete local access[19, 23]; and (b) how norm-governed procedures distribute accountability across roles without erasing agency[16]. Additionally, heteroassociative mapping supports retrieval without a canonical clue: semantically or procedurally adjacent surrogates can trigger reconstruction, which underwrites the link between associative memory and tacit/relational knowledge[6, 19, 23]. Transfer, therefore, yields analogical illumination rather than metaphysical identity: models illustrate mechanisms, while philosophy interprets and bounds them.
- **Accountability controls.** Each extrapolation must be explicitly justified, and its limits declared, thereby avoiding misplaced concreteness[12]. A responsible projection requires four methodological commitments:
- **P1.** Isomorphism stated.
  P2. Transfer articulated.
  P3. Limits declared.
  P4. Counter-models considered.

This framework disciplines philosophical projection, ensuring that metaphors are not mistaken for realities. Computational architectures (EHAM and Alpha–Beta BAM) serve as analogical scaffolds, not as metaphysical guarantees, for the philosophical claims advanced. Hence, heteroassociativity is methodologically central rather than optional: it explains how tacit outcomes and strategy-diverse retrieval are possible within distributed systems.

The methodological scaffolding of the *Cognitive Tapestry*, therefore, requires a precise articulation of its formal structure. To preserve coherence across epistemic, metaphysical, and ethical dimensions, its components must be defined not merely as technical mechanisms but as normative conditions that enable distributed cognition.

The following section introduces this formal architecture by deriving its core invariants—Dispersion (D), Redundancy (R), Mutual Information (I), and Revision-Readiness (ρ)—from both philosophical and informational principles. These invariants constitute the systemic grammar through which cognition maintains stability across heterogeneous agents and uncertain environments.

2.2. Conceptual Clarity: Definitions and Philosophical Validation

The *Cognitive Tapestry* derives its philosophical coherence from four normative invariants— dispersion (D), redundancy (R), mutual information (I), and revision-readiness (ρ). Collectively, these articulate the epistemic conditions necessary for distributed intelligence. These invariants are not mere engineering parameters; they function as epistemic and ethical regulators, defining the boundaries within which knowledge can remain reliable, self-correcting, and morally accountable. As a philosophical framework, the *Cognitive Tapestry* treats cognition as an exercise in co-regulation, rather than control. Its aim is not to impose a metaphysical totality, but to delineate the normative conditions under which knowledge stays corrigible, plural, and ethically sustainable. It thereby transforms the logic of computation into a philosophy of care—one in which epistemic stability emerges from reciprocity and restraint, not domination.

2.2.1. Philosophical Derivation of the Invariants

Each invariant translates a classical principle of epistemic justification into a distributed setting:

- **Dispersion (D)** represents the measure of uncertainty sustained across nodes. It embodies the logic of Popper's fallibilism, wherein knowledge progresses not by seeking certainty but by testing a diversity of perspectives within bounded indeterminacy [22, 24, 34]—a dynamic that mirrors the foundational principle of diversity in collective intelligence.
- **Redundancy (R)** expresses the degree of overlap among agents—diverse routes to the same information. It preserves reliability through cooperation, mirroring Alston's procedural notion of justification and Shannon's insight that redundant channels guard against noise [34].
- **Mutual Information (I)** captures relational coupling: the extent to which agents share informational states without collapsing their autonomy. It transposes Floridi's principle of informational coherence into a quantitative register [8, 9].
- **Revision-Readiness (ρ)** formalises the capacity for correction under bounded time. It embodies Popper's and Alston's demand for corrigibility: a system is epistemically justified when it can respond to error through timely adaptation [24].

Together, these invariants define an entropic equilibrium between openness and stability—an ethical as well as epistemic ideal where resilience, not certainty, marks the health of knowledge. Their balance operationalises Floridi's informational ethics [9], translating the duty to preserve informational integrity into the design of hybrid human–AI systems.

2.2.2. Systemic Conditions for Distributed Cognition

Cognition becomes systemic when:

1. Reliable performance depends on non-local relations (high I);
2. Resilience emerges from informational overlap (adequate R);
3. Timely correction is possible (sufficient ρ); and
4. Uncertainty remains bounded yet generative (moderate D).

These conditions distinguish genuine distributed cognition from mere aggregation [2]. They anchor epistemology in a logic of regulated entropy—an open architecture of reason that remains accountable to error and plural perspectives.

2.2.3. Computational Resonance

Formal models provide empirical resonance without defining ontology. Alpha–Beta BAM guarantees stable retrieval under noisy inputs [7], while EHAM demonstrates that stability arises from relational encoding rather than central control [6]. Both illustrate that knowledge is sustained through distributed regulation of uncertainty—precisely what the invariants express philosophically. They serve as heuristic laboratories for testing the epistemic balance between openness and constraint.

2.3. Epistemic Continuity: From Formal Invariants to Philosophical Method

As established above, the invariants (D, R, I, ρ) function as normative conditions for distributed knowledge rather than empirical constants. This section extends this logic methodologically from the formal structures of EHAM to the philosophical criteria of justification and responsibility. The question is no longer whether a system learns, but whether it ought to count as knowing under epistemic and moral standards.

2.3.1. Projection and Normative Guardrails

Projection from computational models to philosophical claims demands discipline. Models such as EHAM or Alpha–Beta BAM must be read as heuristic experiments that illustrate, rather than found, epistemic theory. This requires what the framework calls projection-protection barriers—conceptual guardrails that translate formal relations into philosophical meaning without metaphysical inflation.

These barriers ensure that computational analogies remain within the domain of justification rather than ontology.

They prevent the reification of models while permitting philosophical learning from their structure. Thus, the invariants serve as regulative questions for epistemic and ethical evaluation (see **Table 1**):

Table 1. Normative invariants of the *Cognitive Tapestry*.

Invariant	Philosophical Question	Epistemic–Normative Function
D (Dispersion)	How is uncertainty bounded so that learning remains possible?	Safeguards Popperian fallibilism and openness to error [24]
R (Redundancy)	How is diversity sustained without disintegration?	Grounds the principle of diverse reliability [22], whereby systemic resilience is achieved through overlapping yet varied capacities
I (Mutual Information)	How are relations coordinated without erasing individuality?	Embodies Floridi's informational coherence [8, 9]
ρ (Revision-Readiness)	How does the system correct itself when faced with anomaly?	Operationalises Alston's procedural justification and moral responsiveness [24]

Through these questions, epistemology evolves from a theory of belief to a theory of systemic resilience. Knowledge is not the elimination of error but the capacity to recover from it with grace.

2.3.2. Methodological Payoff

This method yields a philosophy of distributed cognition that remains both empirically informed and conceptually restrained. It acknowledges with Popper that all knowledge is provisional, with Alston that justification is procedural, and with Floridi that informational integrity is a moral good. The invariants (D, R, I, ρ) thus operate as a bridge between the formal and the philosophical—between what a system can compute and what it ought to count as knowing.

Epistemology, accordingly, ceases to be a monologue of minds and becomes a conversation of systems—open, fallible, and ethically accountable.

3. Epistemology: Beyond the Individual Knower

Classical epistemology has long confined knowledge within the cognitive boundaries of the individual, measuring it through belief, truth, and justification [24]. Yet this image falters in an age when understanding arises from networks of human and artificial agents. Knowledge now circulates through distributed systems in which no single consciousness commands the whole.

The *Cognitive Tapestry* reframes this terrain by showing how epistemic stability and moral agency can emerge from ensembles of humans, artefacts, and algorithms. Drawing on the architectures of Entropic Hetero Associative Memory (EHAM) [6] and Alpha–Beta BAM [7], it reveals that informational order can persist even when access is partial or fragmented.

This section develops four interwoven categories—phantom objectivities, entropic networks, entropic and tacit knowledge, and group agency—each extending the normative invariants (D, R, I, ρ) from methodological structure to epistemological meaning.

3.1. Phantom Objectivities

Philosophy has long wrestled with the enigma of order that exceeds comprehension. Phantom objectivities name those informational structures that

1. are constituted relationally across human–AI ensembles;
2. exert genuine constraint on practice; and
3. remain opaque to individual introspection.

The concept extends Durkheim's notion of social facts—forces that "exercise coercion external to the individual" [35]—into computational ecologies where order now emerges algorithmically. During the COVID-19 pandemic, predictive dashboards constrained governmental action even though no policymaker could fully apprehend their algorithmic couplings [29].

Formally, phantom objectivities appear as network-level invariants of high mutual information (I) and redundancy (R) that persist despite local ignorance [11]. They show how collectives can track reality without any node knowing it in the classical sense. Far from illusion, they are ontic stabilisers—fields of constraint that render distributed cognition effective.

Computational analogies illuminate the idea. In Alpha–Beta BAM, correct recall occurs even when decisive cues vanish from single units [7]; EHAM reconstructs patterns through entropic compensation [6]. Both exhibit epistemic persistence beyond local awareness. Hence, phantom objectivities transform cognition from possession into participation, teaching that knowledge lives less in the mind that holds it than in the relations that sustain it.

3.2. Entropic Networks

If phantom objectivities reveal the architecture of distributed order, entropic networks describe its living rhythm. EHAM shows that cognition endures not by eliminating entropy but by regulating it—sustaining coherence amid uncertainty [3].

This insight echoes the logic of fallibilism, the emergent principle that "more-is-different" [21] in complex systems, and the cognitive value of diversity [22]: systems remain intelligent precisely by allowing bounded variation. A resilient network balances dispersion (D), redundancy (R), and revision-readiness (ρ); order persists through indeterminacy, not against it.

Thus, knowledge becomes resilience under bounded uncertainty. The old ideal of justified true belief gives way to a systemic virtue: the capacity to preserve reliability across entropic regimes. Floridi's information ethics [9] finds a new register here—maintaining informational flourishing becomes a moral analogue of maintaining entropy within the "band of life".

Cognition, then, is not a fortress of certainty but a choreography of controlled openness. Managing entropy is both an epistemic and an ethical act, for it safeguards the very conditions of sense and cooperation.

3.3. Entropic and Tacit Knowledge

Polanyi observed that "we know more than we can tell" [4]. Collins expanded this into relational, collective, and somatic dimensions [23]. The *Cognitive Tapestry* re-interprets these insights: tacit reliability is not ineffable mystery but the emergent stability of entropic networks.

Formally, entropic knowledge obtains when

- **Dispersion (D)** maintains creative uncertainty;
- **Redundancy (R)** and **mutual information (I)** interlink agents without collapse; and
- **Revision-readiness (ρ)** enables swift correction of errors.

Through this grammar, tacit knowledge becomes analytically visible—resilience under controlled indeterminacy.

Consider how this unfolds across domains. In surgical teams, no single practitioner commands every variable, yet redundancy, communication, and readiness for correction minimise harm [36]. In open-source communities, reliability emerges from overlapping vigilance rather than perfect foresight [31]. Neuroscience reveals a parallel logic: procedural competence diffuses across neural circuits inaccessible to introspection [37, 38].

Computational analogues echo this pattern: Alpha–Beta BAM retrieves order from noise [7], while EHAM restores coherence even when data are fragmentary [6]. What seems tacit at the personal level thus becomes systemic at the collective one. Tacit knowledge designates what must remain distributed—the capacity of the whole to sustain sense through regulated uncertainty. What Polanyi once treated as the limit of language becomes, here, the virtue of adaptive silence—the intelligence of the network itself.

3.4. Group Agency

If informational order and tacit resilience are verifiable phenomena, then collectives may satisfy the criteria for epistemic agency. Such agency is established when a group's collective decisions demonstrate internal coherence and external responsiveness, a process through which a distinct institutional mindset is formed and sustained [2]. The Cognitive Tapestry extends this: a collective becomes an epistemic subject when it maintains the invariants (D, R, I, ρ) across entropic conditions.

Computational analogies reinforce this principle. In Alpha–Beta BAM and EHAM, coherence is achieved through reciprocal reinforcement among distributed nodes [6, 7]; reliability arises not from individual units, but from systemic interdependence. Similarly, group agency depends on the maintenance of informational invariants rather than on shared belief or introspection. Responsibility therefore scales with each node's marginal contribution to resilience.

This has ethical resonance. If groups can know, they can also be answerable. Accountability becomes layered—individual, institutional, and algorithmic. Phantom objectivities serve as the background against which responsibility acquires its measure. To know together is to be responsible together; the epistemic and the moral threads intertwine.

3.5. Summary

Epistemology in the *Cognitive Tapestry* thus moves beyond the solitary knower toward the systemic participant.

1. **Phantom objectivities** reveal how distributed constraints embody shared reality [11, 29, 35].
2. **Entropic networks** show that cognition arises through the governance—not the abolition—of uncertainty [3, 21, 22].
3. **Entropic knowledge** translates tacit skill into systemic virtue [4, 23, 37, 38].
4. **Group agency** extends responsibility across hybrid ensembles [2, 6, 7].

Taken together, these strands redefine knowledge as participatory stability—the art of coherence amid diversity, correction amid uncertainty. The invariants (D, R, I, ρ) function as epistemic norms guiding this balance.

In this light, Popper's fallibilism, Floridi's information ethics, and Polanyi's tacit knowing converge: each recognises that truth subsists not in fixity but in disciplined openness. The *Cognitive Tapestry* inherits that lineage, translating it into the post-individual condition of the *Cognitive Era*.

Epistemology thus prepares its philosophical passage toward metaphysics. Having shown that cognition can persist without a central mind, we must now ask what sort of being such cognition entails. Section 4 turns to that question, examining emergence, ontology, and the boundaries of post-individual intelligence.

4. Metaphysical Boundaries: Emergence and Ontology

If epistemology concerns how knowledge subsists, metaphysics must ask what kind of being such knowledge implies. Having shown that cognition can stabilise beyond the individual mind, the *Cognitive Tapestry* turns to its ontological implications. Here, the task is to distinguish explanatory emergence—a way of speaking about complexity that preserves methodological discipline—from ontological emergence, which too quickly reifies novelty into new substances. Within this framework, metaphysical boundaries are drawn through three interdependent perspectives: controlled indeterminacy and emergence, post-individual intelligence, and limit cases (Integrated Information Theory and fusion identities). Each seeks to maintain philosophical sobriety while allowing systemic categories to illuminate reality without inflating them into metaphysical dogmas [12, 15, 33].

4.1. Controlled Indeterminacy and Emergence

The entropic architectures explored earlier—EHAM and Alpha–Beta BAM—demonstrate that cognition is not the offspring of rigidity but of controlled openness. Shannon's information theory already hinted that order depends on measured uncertainty [34]; the same intuition animates Popper's evolutionary epistemology, where truth unfolds through conjecture, refutation, and revision [12]. Cognition, in this sense, is a dance between constraint and indeterminacy.

The *Cognitive Tapestry* thus treats emergence as explanatory, not ontological. It aligns with the established view of weak emergence [13], where macro-level patterns arise from micro-level interactions but remain essential for a complete understanding, and the methodological approach of piecewise approximation for complex systems [14]. Consequently, macro-level order is recognised as epistemically indispensable yet not metaphysically independent. The relevant invariants—dispersion (D), redundancy (R), mutual information (I), and revision-readiness (ρ)—mark the parameters within which knowledge, agency, and responsibility remain possible.

These invariants function as normative epistemic conditions, akin to the regulative ideals of Popperian fallibilism [24]; they are not mere engineering constants but philosophical constraints that safeguard intelligibility.

From this perspective, controlled indeterminacy reframes emergence as a methodological necessity.

We call phenomena "emergent" not to multiply entities but to recognise that systemic intelligibility cannot always be reduced to its micro-rules. The world, as Whitehead warned[12], resists both the absolutism of mechanism and the vagueness of mysticism. The *Cognitive Tapestry* situates itself precisely in that interval—where uncertainty becomes the very medium of order.

4.2. Post-Individual Intelligence

Traditional metaphysics bound intelligence to the perimeter of a single organism. The *Cognitive Tapestry* recasts this image by describing post-individual intelligence, in which epistemic and moral capacities are distributed across human, algorithmic, and institutional nodes.

This is not a metaphysical claim about group souls; it is an ontological account of systemic participation. Post-individual intelligence arises when three conditions coincide:

1. the stability of phantom objectivities—non-local informational structures that constrain practice[11];
2. the persistence of entropic knowledge—resilience maintained under bounded uncertainty[3, 4, 6, 7, 23]; and
3. the responsiveness of group agency—collective action oriented by systemic cues[2].

This conception resonates with Clark and Chalmers' extended mind[18] but goes further: whereas their thesis augments the individual through artefacts, post-individual intelligence describes systems in which no single participant occupies epistemic centrality. Cognition is woven beyond the mind, emerging through the regulation of the invariants (D, R, I, ρ).

Philosophically, this model aligns with Floridi's informational structural realism[9], where reality consists of structured informational relations rather than inert substances. In this light, post-individual intelligence does not dilute agency; it re-grounds it. To act is to contribute to the maintenance of informational coherence—an ontological vocation shared across humans and machines alike.

4.3. Limit Cases: IIT and Fusion Identities

Metaphysical inquiry advances most clearly when it approaches its limits. Certain frontiers—where explanatory ambition risks crossing into speculative inflation—demand particular caution. Within the *Cognitive Tapestry*, two such limit cases are especially instructive: Integrated Information Theory (IIT) and fusion identities. Each illuminates how far systemic explanation can go without exceeding the bounds of disciplined emergence.

Integrated Information Theory. Information Theory proposes that consciousness corresponds to the degree of informational integration—symbolised by Φ—within a system[1, 27]. This theory has achieved remarkable formal elegance, linking subjective unity to measurable structure. Yet its interpretation as *intrinsic consciousness* risks overextension, as critics have argued[39].

From the perspective of the *Cognitive Tapestry*, IIT provides a powerful methodological instrument, but not a metaphysical revelation. The measure of Φ usefully describes the degree of informational coupling among system components; it quantifies how much the system functions as a whole rather than as independent parts. What it cannot justifiably affirm, however, is the presence of consciousness itself. To cross that boundary is to mistake the map for the territory—to treat systemic correlation as ontological substance. Thus, within this framework, IIT remains an analytical ally, not a metaphysical oracle.

Fusion identities. A similar caution applies to theories that claim collectives may "fuse" into singular agents with unified consciousness[40]. Such hypotheses—fascinating as they are—blur the conceptual boundary between functional unity and ontological identity. The *Cognitive Tapestry* acknowledges that fusion may occur descriptively: tightly coupled teams, military units, or surgical ensembles can act as if guided by one intention. Yet ontologically, it is more prudent to interpret these as functional unities, not as the birth of new metaphysical subjects.

The commitment here is minimal and methodologically disciplined: systemic cognition exists where invariants (D, R, I, ρ) define conditions for reliable performance, not where substances merge. Fusion is an explanatory convenience, not an ontological creation.

Metaphysical discipline through methodological restraint. Both IIT and theories of fusion identities demonstrate why philosophy must distinguish the scope of explanation from the scope of being. Each shows how easily structural coherence can be mistaken for metaphysical novelty. The *Cognitive Tapestry* resists this temptation. It treats such theories not as rival ontologies but as boundary signals—warnings that remind us where epistemic modelling must yield to ontological modesty.

In this sense, the *Cognitive Tapestry* embodies a metaphysics of restraint: it values coherence without claiming totality, integration without inflation, and unity without erasure. It understands that explanation, when too eager, risks turning into metaphysics by stealth. Its task, therefore, is not to dissolve mystery, but to keep it disciplined—to map the horizon of emergence without mistaking it for the dawn of new substance.

4.4. Summary

Metaphysics within the *Cognitive Tapestry* seeks not grandeur but precision. It avoids both the reductionism that dissolves the mind into mechanism and the inflation that turns structure into substance. Each preceding section has clarified this balance.

Controlled indeterminacy shows that order in cognition is not imposed but regulated: intelligence subsists through openness bounded by normativity. Emergence, therefore, is treated as explanatory necessity rather than as ontological novelty—a disciplined way of acknowledging that the system's coherence exceeds the descriptive reach of its parts.

Post-individual intelligence extends this discipline to agency itself. Cognition can unfold across humans, artefacts, and institutions without invoking mystical unities. Intelligence is thus re-located: not in a sovereign mind, but in the entropic coordination that keeps diversity coherent.

Finally, the limit cases—Integrated Information Theory and theories of fusion identities—demonstrate where caution becomes philosophical virtue. They remind us that explanation must not trespass into metaphysical invention. The *Cognitive Tapestry* treats such frontiers as mirrors of restraint: the point at which understanding honours its own limits.

Together, these insights establish a metaphysics of disciplined emergence—a vision in which being is neither atomistic nor mystical, but woven through systemic interdependence. Yet metaphysical clarity alone cannot suffice; systems that know must also answer for what they do.

The next section, therefore, turns to ethics, asking how responsibility and care can be sustained within entropic networks where knowledge, agency, and moral consequence are inseparably entwined.

5. Ethics and Responsibility in Distributed Cognition

If epistemology shows how knowledge subsists across networks, ethics must explain how responsibility can do the same. When hybrid ensembles of humans and machines act, accountability cannot be confined to any single node. The well-known responsibility gap arises precisely from this dispersion[17].

The *Cognitive Tapestry* approaches the problem not as a moral vacuum but as a question of entropic regulation: responsibility scales with an agent's contribution to preserving coherence, responsiveness, and the capacity for revision amid uncertainty.

Accordingly, moral accountability accrues to each participant—human, institutional, or algorithmic—in proportion to its influence on the invariants (D, R, I, ρ): dispersion, redundancy, mutual information, and revision-readiness. These are not engineering parameters but normative conditions of moral cognition, comparable to Floridi's principles of informational flourishing[9, 10] and to the layered justification explored by Alston[24]. Responsibility, in this sense, becomes the ethical counterpart of epistemic resilience.

Ethical responsibility in distributed cognition is not merely procedural; it is also virtue-ethical. Each invariant evokes a classical moral disposition: dispersion embodies prudence, the capacity to recognise plurality and context; redundancy mirrors justice, ensuring that no actor bears an unequal burden; mutual information expresses courage, the openness to feedback and correction; and revision-readiness manifests temperance—the willingness to revise one's course. In this way, the entropic invariants echo the architecture of virtue itself, transposed from the psychology of the individual to the dynamics of systems. Virtue in the *Cognitive Era* consists in sustaining the conditions under which both humans and machines may continue to learn without collapsing into rigidity or chaos.

The COVID-19 pandemic revealed how these moral structures operate under pressure. Epidemiological dashboards and modelling platforms acted as phantom objectivities: they constrained collective behaviour without being fully understood by any policymaker[29]. The moral question was not who knew the facts, but whether the system as a whole maintained adaptive regulation—acknowledging dispersion across reporting sites, securing redundancy through overlapping hospital data, and ensuring revision-readiness when anomalies emerged. Failures in redundancy (under-reporting) and delayed revision magnified global harm. Responsibility, therefore, is distributed across ministries, hospitals, and international agencies according to their systemic role in sustaining informational coherence. Ethics, in such contexts, becomes a matter of maintenance—the preservation of channels through which collective cognition remains alive and responsive.

A similar logic governs medical triage, where algorithmic systems increasingly interact with human judgment [28]. Here, moral action arises not from isolated decision-makers but from coordinated intelligence. Responsibility does not reside solely in the physician or in the code, but in the ensemble's capacity to preserve the entropic balance that minimises error and preserves fairness. When systems lose the ability to learn from their own uncertainty, moral failure occurs. Ethical design, by contrast, aligns with Floridi's information ethics: to act rightly is to sustain the well-being of the infosphere [9]. Hospitals that integrate feedback-driven triage loops exemplify moral responsibility as entropic stewardship—a care for the resilience of knowledge itself.

Swarm robotics offers yet another mirror of this distributed virtue. In robotic collectives, no single agent "knows" the mission in full; global order arises from local rules and feedback [32]. Here, the ethical dimension lies not in command but in corrigibility—the capacity to absorb disturbance without collapse. Responsibility belongs primarily to those who design and tune the feedback loops rather than to the operators who trigger them. Ethical control thus becomes a property of design, not merely of intention, echoing Santoni de Sio and van den Hoven's account of meaningful human control [16].

In software engineering, DevSecOps processes embody this same ethical structure. Systematic review confirms that the highest-performing teams are characterised by their psychological stance toward failure, reinforcing the practice of treating human error as a signal to refine processes and automation, rather than attributing blame to individuals [30].

Thus, each code commit is validated through layered automated checks, peer reviews, and rollback protocols. Failures are viewed not as moral shortcomings, but as learning signals. From an ethical standpoint, what matters is the willingness to revise (ρ)—the commitment to recover and repair. Responsible programming means designing with reversibility in mind, building systems that can recognise and correct their own errors. In this sense, humility becomes the foremost technical virtue: the discipline of creating artefacts capable of self-correction.

Across all these contexts, a pattern becomes clear. Moral agency is not erased but redistributed, and each participant's duty corresponds to its influence on the system's informational health. The invariants (D, R, I, ρ) thus become ethical indicators: dispersion ensures inclusivity of perspectives; redundancy grounds reliability through overlap; mutual information binds communication into coherence; and revision-readiness sustains adaptability through humility. Together, they articulate a normative ecology in which responsibility equals the capacity to sustain systemic learning. To act ethically in the *Cognitive Era* is to contribute to this equilibrium—to keep entropy within the bounds of sense.

Ethics in distributed cognition must not aspire to control everything, but to nurture conditions for responsible emergence. The *Cognitive Tapestry* renounces metaphysical finality in favour of moral vigilance. Humility, far from weakness, becomes wisdom—the awareness that in hybrid ecologies, no single actor commands the whole truth. Progress, correspondingly, lies not in certainty but in corrigibility: the perpetual readiness to learn, revise, and restore coherence. Entropic ethics thus becomes a philosophy of care for systems, for knowledge, and for one another.

6. Discussion

Having explored the ethical horizon, the *Cognitive Tapestry* must now demonstrate methodological rigour. It advances ambitious claims about epistemology, metaphysics, and ethics; to preserve credibility, it must anticipate objections, exhibit philosophical restraint, and trace new directions for inquiry. Its aim is not to defend a doctrine, but to refine a discipline of thought adequate to the *Cognitive Era*.

6.1. Addressing Objections

Several objections have been raised against distributed cognition; each refines the framework.

Misplaced Concreteness. Whitehead warned against treating abstractions as concrete realities [12]. The *Cognitive Tapestry* avoids this trap by setting explicit scope conditions, transfer principles, and validation criteria (Section 2). Computational models such as EHAM or Alpha–Beta BAM serve not as ontological commitments but as heuristic laboratories for philosophical reflection.

Through cross-model robustness [14], projection from computation to cognition remains bounded by method, not metaphysics.

Functionalist Reductionism. Some might argue that the *Cognitive Tapestry* risks collapsing into classical functionalism [25, 26]. Yet the framework extends beyond mere function: the invariants (D, R, I, ρ) are not physical metrics but normative epistemic conditions. They formalise when distributed cognition becomes intelligible, while the ethical scaffolding of Section 5 introduces a moral dimension absent from traditional functionalism. It is, therefore, a normatively enriched functionalism—one that binds knowledge and responsibility within a single entropic economy.

Emergence and Category Error. Another worry is that describing knowledge as "emergent" confuses explanation with ontology. The *Cognitive Tapestry* explicitly confines itself to explanatory emergence [13, 40]: the point where systemic categories become epistemically indispensable, not metaphysically novel. Controlled indeterminacy, introduced earlier [34], ensures that order and openness co-sustain cognition without invoking new substances.

Responsibility Scepticism. Critics sometimes fear that distributing responsibility dissolves accountability into vagueness [17]. The reply is empirical and moral: accountability is measurable through each agent's marginal effect on systemic resilience. Responsibility scales with contributions to redundancy, communication, dispersion, and revision-readiness. Far from weakening accountability, this view sharpens it—each participant is answerable for the portion of entropy they sustain or neglect.

6.2. Future Work

The *Cognitive Tapestry* is not a closed system but an open research line for the *Cognitive Era* [20]. Its promise lies in the dialogue it invites between philosophy, computation, and social practice. Future work unfolds along several intertwined paths:

1. **Empirical diversification.** Validation must extend beyond EHAM [6] and Alpha–Beta BAM [7]. Comparative analysis and testing across predictive processing frameworks [3], swarm robotics [32], and open-source collectives [31] will evaluate cross-model robustness [14].
2. **Phenomenological integration.** While systemic invariants (D, R, I, ρ) capture the logic of distributed cognition, lived experience within entropic networks—whether in clinical teams or software development—remains under-explored [41]. Integrating phenomenology will clarify how humans experience and interpret their participation in hybrid intelligence.
3. **Institutional design.** Responsibility matrices [16] must evolve towards regulatory governance architectures. Case studies in healthcare, cybersecurity, and open development ecosystems can give rise to practical models of entropic regulation or ethical coordination [10].
4. **Historical comparison.** Situating the *Cognitive Era* within the continuum of collective intelligence—from the polis and monastic orders to scientific communities—will clarify both continuity and novelty [19, 23]. Such genealogy protects the project from the illusion of absolute novelty.
5. **Conceptual refinement.** Core notions such as phantom objectivities [11], entropic knowledge [4, 23], and controlled indeterminacy [24] deserve deeper elaboration to distinguish them from adjacent concepts in epistemology, phenomenology, and the philosophy of information [8–10].

6.3. Summary

By engaging its critics and projecting future work, the *Cognitive Tapestry* reveals itself as a method, not a dogma. It acknowledges its limits, resists metaphysical inflation, and grounds inquiry in humility. Its enduring contribution is to offer a disciplined way of navigating knowledge, agency, and responsibility across the evolving landscape of hybrid intelligence. What remains is to gather these threads into a unified reflection—one that clarifies what the *Cognitive Tapestry* ultimately teaches about knowledge, being, and care in the age of hybrid intelligence.

7. Conclusions

This document proposes the *Cognitive Tapestry* as a philosophical framework for understanding knowledge, agency, and responsibility in the *Cognitive Era*[5]. Drawing on the architectures of Entropic Hetero Associative Memory (EHAM)[6] and Alpha–Beta BAM[7], it redefines epistemology, metaphysics, and ethics in accordance with the realities of post-individual intelligence. Its fundamental premise is that philosophy must precede AI: while computation reveals what is possible, only philosophy determines what is intelligible and right. Thus, the framework places intelligence not in the individual mind, but in systemic networks of human beings, artefacts, and computational systems that sustain epistemic order and ethical coherence.

Epistemological Clarity. The *Cognitive Tapestry* refines our vocabulary for distributed knowing. Concepts such as phantom objectivities[11], entropic knowledge[4, 23], entropic networks[3, 21, 22], and group agency[2] articulate how cognition persists beyond the individual while avoiding dissolution into abstraction. The systemic invariants—dispersion (D), redundancy (R), mutual information (I), and revision-readiness (ρ)—translate classical epistemic virtues into formal conditions of intelligibility. They extend insights from tacit knowledge[4, 23] and distributed cognition[19], adding analytical precision drawn from information theory and cognitive engineering. Knowledge, thus reframed, becomes a property of coherence under uncertainty rather than possession by a mind.

Metaphysical discipline. Metaphysically, the framework resists both reduction and inflation. By treating emergence as explanatory rather than ontological[13, 14, 40], it avoids the extremes of scientism and mysticism. Controlled indeterminacy[34] shows that intelligence thrives in bounded openness: too much order breeds rigidity, too much freedom, chaos. Post-individual intelligence demonstrates that cognition can extend beyond the single subject without invoking new substances or group souls. Limit cases—Integrated Information Theory[1, 27, 39] and fusion-identity hypotheses[40]—mark the metaphysical horizon, signalling where explanatory ambition risks excess.

Ethical scaffolding. Ethically, the *Cognitive Tapestry* reconceives responsibility as layered entropic regulation. Across pandemic dashboards[29], medical triage[28], swarm robotics[32], open-source communities[31], and DevSecOps pipelines[30], moral agency is redistributed rather than erased. Each actor's duty corresponds to its effect on systemic resilience. In this view, accountability becomes measurable: the capacity to sustain redundancy, communication, and revision. Far from creating a responsibility gap[17], hybrid systems invite a deeper form of vigilance—a virtue of corrigibility, of remaining open to correction as a condition of moral life.

Toward a Philosophy of Care. The *Cognitive Tapestry* does not aspire to finality. Its strength lies in offering a method of disciplined reflection—an epistemic humility attuned to the intricate complexity of entropic systems. Future research must diversify computational foundations, integrate phenomenological insight[41], and refine distinctions between phantom objectivities, entropic knowledge, and older notions such as Durkheimian social facts[35] or the extended mind[18]. To acknowledge that intelligence is now woven across networks is not to abdicate philosophy's task but to extend it—to trace the patterns through which knowing, acting, and caring remain human, even when shared with machines.

In the *Cognitive Era*, to think ethically is to keep learning; to know is to sustain coherence amid uncertainty; and to act responsibly is to preserve the fragile balance that allows both the system—and ourselves—to continue becoming.

Author Contributions

This paper explores the intersection of philosophy and technology. The lead author and corresponding author, J.R.-O., was responsible for the conceptualisation, methodology, formal analysis, research, resource acquisition, original draft preparation, writing, review and editing, visualisation, and overall supervision of the project. The secondary author, K.O.-S., provided valuable contributions to the review and editing of the technological aspects discussed within the paper. Both authors have read and agreed to the published version of the manuscript.

Funding

This work did not receive external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

References

[1] Cea, I., Signorelli, C.M., 2025. How to be an integrated information theorist without losing your body. Frontiers in Computational Neuroscience. 18, 1510066. DOI: https://doi.org/10.3389/fncom.2024.1510066

[2] Tollefsen, D., 2025. Shaping the Institutional Mind. Philosophies. 10(5), 112. DOI: https://doi.org/10.3390/philosophies10050112

[3] Parr, T., Pezzulo, G., Friston, K.J., 2022. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior (Open Access edition). The MIT Press: Cambridge, MA, USA. DOI: https://doi.org/10.7551/mitpress/12441.001.0001

[4] Polanyi, M., 1966. The Tacit Dimension. Doubleday & Company: Garden City, NY, USA.

[5] Rodas-Osollo, J., 2023. An Interesting Adventure Accompanied by CMCg.I Model. Zenodo: Geneva, Switzerland. Available from: https://zenodo.org/records/10111223

[6] Morales, R., Pineda, L.A., 2025. The missing cue problem in hetero associative memory retrieval. Scientific Reports. 15, 21850. DOI: https://doi.org/10.1038/s41598-025-07963-x

[7] Acevedo-Mosqueda, E., Cruz, J.M., Cruz, A., 2007. Alpha–Beta BAM: A bidirectional associative memory model with alpha–beta learning. Neural Processing Letters. 26(1), 1–20.

[8] Floridi, L., 2011. The Philosophy of Information. Oxford University Press: Oxford, UK.

[9] Floridi, L., 2013. The Ethics of Information. Oxford University Press: Oxford, UK.

[10] Munn, L., 2023. The uselessness of AI ethics. AI and Ethics. 3, 869–877. DOI: https://doi.org/10.1007/s43681-022-00209-w

[11] Rodas-Osollo, J., Olmos-Sánchez, K., 2025. Phantom objectivity: Intersubjective knowledge in complex domains of informal structure. Philosophy and Realistic Reflection. 2(1), 1–24. DOI: https://doi.org/10.55121/prr.v2i1.366

[12] Whitehead, A.N., 1978. Process and Reality: An Essay in Cosmology. Free Press: New York, NY, USA.

[13] Kistler, M., 2025. Emergent Properties. In The Material Mind: Reduction and Emergence. University of Calgary Press: Calgary, AB, Canada. pp. 137–204.

[14] Wimsatt, W.C., 2007. Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality. Harvard University Press: Cambridge, MA, USA.

[15] Bailer-Jones, D., 2009. Scientific Models in Philosophy of Science. University of Pittsburgh Press: Pittsburgh, PA, USA.

[16] Santoni de Sio, F., van den Hoven, J., 2018. Meaningful human control over autonomous systems: A philosophical account. Frontiers in Robotics and AI. 5, 15. DOI: https://doi.org/10.3389/frobt.2018.00015

[17] Nyholm, S., 2023. Artificial Intelligence, Ethics of. In: Sellers, M., Kirste, S. (Eds.). Encyclopedia of the Philosophy of Law and Social Philosophy. Springer: Dordrecht, Netherlands. pp.1–8. DOI: https://doi.org/10.1007/978-94-007-6730-0_1093-1

[18] Clark, A., Chalmers, D., 1998. The extended mind. Analysis. 58(1), 7–19. DOI: https://doi.org/10.1093/analys/58.1.7

[19] Hutchins, E., 1995. Cognition in the Wild. MIT Press: Cambridge, MA, USA.

[20] Vallor, S., 2016. Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting. Oxford University Press: New York, NY, USA.

[21] Anderson, P.W., 1972. More is different: Broken symmetry and the nature of the hierarchical structure of science. Science. 177(4047), 393–396. DOI: https://doi.org/10.1126/science.177.4047.393

[22] Riedl, C., Kim, Y.J., Gupta, P., et al., 2021. Quantifying collective intelligence in human groups. Proceedings of the National Academy of Sciences of the United States of America. 118(21), e2005737118. DOI: https://doi.org/10.1073/pnas.2005737118

[23] Collins, H., 2010. Tacit and Explicit Knowledge. University of Chicago Press: Chicago, IL, USA.

[24] Alston, W.P., 1989. Epistemic Justification: Essays in the Theory of Knowledge. Cornell University Press: Ithaca, NY, USA.

[25] Fodor, J.A., 1975. The Language of Thought. Thomas Y. Crowell Company, Inc.: New York, NY, USA.

[26] Putnam, H., 1960. Minds and Machines. In: Hook, S. (Ed.). Dimensions of Mind. New York University Press: New York, NY, USA. pp. 138–164.

[27] Tononi, G., 2004. An information integration theory of consciousness. BMC Neuroscience. 5(1), 42. DOI: https://doi.org/10.1186/1471-2202-5-42

[28] Topol, E., 2019. Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again. Basic Books: New York, NY, USA.

[29] WHO, 2021. COVID-19 Strategic Preparedness and Response Plan - Operational Planning Guideline. World

Health Organization: Geneva, Switzerland. Available from: https://www.who.int/publications/i/item/WHO-WHE-2021.03

[30] Forsgren, N., Humble, J., Kim, G., 2018. Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations. IT Revolution Press: Portland, OR, USA.

[31] Liu, J., Zhang, H., Zou, Y., 2025. Understanding open source contributor profiles in popular machine learning libraries. ACM Transactions on Software Engineering and Methodology. DOI: https://doi.org/10.1145/3747347

[32] Dias, P.G.F., Silva, M.C., Rocha Filho, G.P., et al., 2021. Swarm robotics: A perspective on the latest reviewed concepts and applications. Sensors. 21(6), 2062. DOI: https://doi.org/10.3390/s21062062

[33] Cartwright, N., 1983. How the Laws of Physics Lie. Oxford University Press: Oxford, UK.

[34] Shannon, C.E., 1948. A mathematical theory of communication. Bell System Technical Journal. 27(3–4), 379–423, 623–656.

[35] Durkheim, É., 1982. The Rules of Sociological Method. Halls, W.D. (Trans.). Free Press: New York, NY, USA.

[36] Haynes, A.B., Weiser, T.G., Berry, W.R., et al., 2009. A surgical safety checklist to reduce morbidity and mortality in a global population. New England Journal of Medicine. 360(5), 491–499. DOI: https://doi.org/10.1056/NEJMsa0810119

[37] Squire, L.R., Dede, A.J.O., 2015. Conscious and unconscious memory systems. Cold Spring Harbor Perspectives in Biology. 7(3), a021667. DOI: https://doi.org/10.1101/cshperspect.a021667

[38] Bianco, K.M., Fuelscher, I., Lum, J.A.G., et al., 2024. Procedural learning is associated with microstructure of basal ganglia-cerebellar circuitry in children. Brain and Cognition. 180, 106204. DOI: https://doi.org/10.1016/j.bandc.2024.106204

[39] Cerullo, M.A., 2015. The problem with Phi: A critique of integrated information theory. PLOS Computational Biology. 11(9), e1004286. DOI: https://doi.org/10.1371/journal.pcbi.1004286

[40] Chalmers, D.J., 2006. Strong and weak emergence. In: Clayton, P., Davies, P. (Eds.). The Re-Emergence of Emergence. Oxford University Press: Oxford, UK. pp. 244–254.

[41] Zahavi, D., 2005. Subjectivity and Selfhood: Investigating the First-Person Perspective. MIT Press: Cambridge, MA, USA.