

Weak-TransWeave Cohesion: A Quantale-Hyperseed Reframing of Kemendo’s General Theory of Cohesion with Formal Theorems for Composition, Intelligence, and Self-Modification

A synthesis drawing on Kemendo’s GTC, Hyperseed, Quantale Weakness, and TransWeave

June 10, 2026

Abstract

Andrew Kemendo’s General Theory of Cohesion (GTC) proposes that systems emerge, persist, adapt, and collapse according to how well their components align action vectors, share energy, mediate boundaries, and use internal simulation to preserve viability. This paper preserves that core intuition while replacing GTC’s concrete scalar-vector formalism with a more general weak-transweave framework built from Hyperseed, quantale weakness, Pareto weakness tradeoff theory, geodesic control, and TransWeave. In the resulting picture, a cohesive system is a self-maintaining, resource-bounded, quantale-valued pattern-flow subobject whose self/other boundary remains viable because its internal dynamics, boundary translations, and self-models remain approximately transweavable at acceptable genenergy cost. Kemendo’s formulas then emerge naturally as a rank-one, Euclidean, scalarized, one-step approximation of this richer enriched-categorical viability calculus.

Beyond the conceptual translation, the paper develops three theorem families that make the framework mathematically operational. First, cohesion is shown to be compositional: sequential, parallel, and boundary-glued systems preserve cohesion up to explicitly bounded transweave and glue defects. Second, cohesion is related to Hyperseed-style general intelligence: cohesion is not identical to intelligence, but it bounds the leakage between ideal and actual cognition, controls transfer of competence across task graphs, and determines when cognitive synergy exceeds interface friction. Third, cohesion is linked to robust iterative goal-guided self-modification: if each self-modification step admits a low-defect transweave certificate and the defects are summably bounded, then identity drift and goal drift remain bounded, while capability can improve cumulatively whenever gains dominate drift losses. Together these theorems make precise a strong thesis: general intelligence is cohesive transfer of competence across task-space under resource, boundary, and self-continuity constraints.

Contents

1	Introduction	1
1.1	What this paper is trying to do	1
1.2	The main thesis	2
1.3	The plan of the paper	2
1.4	Methodological caveat	3

2	Kemendo’s General Theory of Cohesion	3
2.1	The basic intuition	3
2.2	Component attributes	3
2.3	Local adjacency field	4
2.4	Alignment cost and boundary formation	4
2.5	Projection and system presentation	5
2.6	Self-modeling and intelligence	5
2.7	The cohesion function	5
2.8	Where the formalism is overly restrictive	6
3	Key Ideas from Hyperseed, Quantale Weakness, Pareto Weakness, and TransWeave	6
3.1	Hyperseed: distinctions, p-bits, patterns, genenergy, and self/other cuts	6
3.1.1	P-bits	6
3.1.2	Distinctions and weakness	7
3.1.3	Genenergy	7
3.1.4	Self/other boundary	7
3.2	Quantale weakness	7
3.3	Quantale-valued process graphs	8
3.4	Pareto weakness	8
3.5	TransWeave	9
3.6	Geodesic or Wu Wei control	9
4	A More General Version: Weak-TransWeave Cohesion	9
4.1	From components to process nodes	10
4.2	From alignment to distinction management	10
4.3	From boundary components to a self/other cut	11
4.4	From inter-system cost Ψ to TransWeave defect	11
4.5	From self-modeling cost to self-transweaving	12
4.6	From scalar viability to a quantale-valued viability profile	12
4.7	Weak-TransWeave cohesion	13
4.8	General action principle	13
5	Kemendo GTC as a Special Case	13
6	Comparison Table	14
7	What the Generalization Buys	15
7.1	It removes the false universality of vector alignment	15
7.2	It treats boundaries as graded and conflict-tolerant	15
7.3	It makes intelligence a self-transweave condition	15
7.4	It makes interaction an operator-theoretic interface problem	15
7.5	It replaces scalar viability with Pareto viability	16
8	Formal Cohesion Certificates and Compositionality Theorems	16
8.1	Plain-language orientation	16
8.2	Normed update model	16
8.3	Cohesion certificates	16
8.4	Sequential composition	17

8.5	Parallel composition	18
8.6	Boundary gluing	18
8.7	Supported Pareto composition	19
8.8	A compact compositionality slogan	20
9	Cohesion and Hyperseed-Style General Intelligence	20
9.1	Plain-language orientation	20
9.2	Task ecologies and GTGI-style scores	20
9.3	No unconditional implication	21
9.4	Cohesion defect as intelligence leakage	21
9.5	Transfer-general intelligence theorem	22
9.6	Cognitive synergy threshold	23
9.7	Cohesion as an effective-resource multiplier	23
10	Cohesion and Robust Iterative Goal-Guided Self-Modification	24
10.1	Plain-language orientation	24
10.2	Transported self-modification trajectories	24
10.3	Bounded drift from summable defects	24
10.4	Goal drift bound	25
10.5	Capability growth under drift loss	25
10.6	Lyapunov-barrier formulation	26
10.7	Stochastic robustness	27
10.8	Design principle for self-modifying AGI	27
11	Synthesis: Cohesion, Intelligence, and Self-Modification	28
12	Conclusion	28

1 Introduction

1.1 What this paper is trying to do

Kemendo’s General Theory of Cohesion is interesting because it asks the right kind of question. Rather than asking only, “What does a system optimize?”, it asks the more basic question, “What lets something remain a system at all?” The answer GTC gives is that systems persist when their parts align, route energy, mediate exchanges through boundaries, and forecast which futures preserve cohesion.

This paper keeps that intuition while replacing the concrete mathematics with a more abstract and general framing. The central move is to replace

components with Euclidean action vectors and scalar energy costs

by

processes in a quantale-valued pattern-flow network, equipped with a paraconsistent self/other boundary and approximate TransWeave maps between internal, external, and simulated dynamics.

In plain terms: a system is not just a pile of parts pointing in roughly the same direction. A system is a pattern that keeps re-making itself while spending limited resources. To do this it must continually decide which distinctions to preserve, which boundaries to maintain, which interactions to accept, which futures to simulate, and which transformations between contexts can be trusted. The generalized framework introduced below gives each of these decisions a precise mathematical home.

1.2 The main thesis

The proposed generalization can be summarized as follows.

A cohesive system is a self-weaving, resource-bounded, quantale-valued pattern-flow subobject whose self/other boundary persists over time because its internal dynamics and boundary interactions remain approximately transweavable at acceptable genenergy cost.

Kemendo's GTC becomes a special case obtained by choosing a Euclidean state space, scalar energy channels, vector alignment as the main coherence proxy, and a scalarized viability function.

1.3 The plan of the paper

The paper proceeds in two movements. The first movement is a conceptual translation: GTC's vectors, energies, boundaries, self-models, and interaction terms are reinterpreted using quantale weakness, Hyperseed p-bit boundaries, Pareto viability profiles, and TransWeave defects. The second movement turns the translation into mathematics, developing formal theorem statements organized around three guiding questions.

First, does cohesion compose? The answer is yes, but only with explicit error budgets. If local interfaces have small TransWeave defect, then sequential and parallel composites have bounded defect. If two systems are glued along a boundary whose self/other cuts are compatible, then the glued system has bounded boundary defect. And if subsystem and glue policies are selected by a compatible Pareto normal, then the composite remains Pareto efficient within the typed feasible class.

Second, what is the relation between cohesion and general intelligence? The answer is not that cohesion equals intelligence. A rock may be cohesive without being intelligent. Rather, cohesion bounds the amount of intelligence lost to self-fragmentation, boundary drift, bad transfer, unnecessary distinctions, interface friction, and genenergy waste. Under task ecologies that reward transfer, modular synergy, and long-horizon self-maintenance, weak-transweave cohesion becomes a lower-bound certificate for effective general intelligence.

Third, what happens under iterative goal-guided self-modification? The answer is that robust self-modification is possible when each modification step comes with a low-defect transweave certificate. If per-step cohesion defects are summable and remain inside the radius of the identity/value basin, then the system does not drift arbitrarily far from itself. If, additionally, capability gains exceed drift losses, capability improves cumulatively. This yields a precise form of the slogan: open-ended self-improvement requires summably bounded self-loss.

1.4 Methodological caveat

The theorems below are deliberately conditional. They do not assert any mystical identity among cohesion, intelligence, life, and goodness. Each theorem says: once one has specified the relevant

state spaces, norms, quantales, task ecologies, boundary metrics, and defect tolerances, cohesion yields checkable composition, transfer, intelligence-leakage, and self-modification guarantees. The strength of the framework lies precisely in this conditionality: it replaces vague claims about “alignment” or “coherence” with explicit diagrams and error terms that can in principle be measured and audited.

2 Kemendo’s General Theory of Cohesion

2.1 The basic intuition

Kemendo’s core idea is that systems form when components line up their actions well enough to share energy through a boundary. Crucially, a system is not assumed as a fixed object. It is an emergent arrangement of components that can hold together under constraint. Boundary components play a special role because they are the points where internal organization meets external pressure.

To make this concrete, imagine a biological organism, a web service, a political coalition, or a multi-agent AI. Each has parts. Some parts sit deep inside the system and mostly support internal stability. Other parts face outward and mediate exchange with the environment. The system persists when the cost of keeping all these parts coordinated is less than the energy, reward, or capability gained by remaining organized.

2.2 Component attributes

Kemendo models each component by a small tuple of measurable attributes. The important conceptual point is that every component has an influence, a degree of alignment with the whole, a resource budget, a consumption rate, and a capacity to mediate interactions.

Formally, the component tuple is:

$$c_x = (\|A_x\|, \theta_x, E_s, E_c, M_x). \tag{1}$$

Here:

- $\|A_x\|$ is the magnitude of the component’s action vector.
- θ_x is the misalignment between the component and the system vector field.
- E_s is stored energy.
- E_c is energy consumption.
- M_x is mediative capacity, the ability to broker energy or interaction.

Put plainly: a component matters if it can do something, has enough energy to keep doing it, does not fight the rest of the system too much, and can help connect other components to one another.

2.3 Local adjacency field

A system does not arise from isolated components. It arises from mediated local influences: each component is pushed and pulled by its neighbors, weighted by how strongly they can act on it. Kemendo captures this with a local adjacency field.

For a component x_j with neighbor set $N(j)$:

$$L_{x_j}(t) = \sum_{i \in N(j)} f(i, j, t) \hat{A}_i(t), \quad (2)$$

where $\hat{A}_i(t)$ is a normalized action vector and $f(i, j, t)$ is a scalar coupling strength.

A typical GTC-style coupling has the shape:

$$f(i, j, t) = \frac{M_i(t) \|A_i(t)\| E_{\text{available},i}(t) \cos(\theta_{ij}(t))}{\log(1 + d_{ij}) E_{c,j}(t)}. \quad (3)$$

Reading the formula directly: influence is stronger when mediation, action magnitude, available energy, and alignment are high; it is weaker when distance and the recipient's consumption burden are high. Influence, in other words, is something a component must be able to afford to exert and the recipient must be able to afford to absorb.

2.4 Alignment cost and boundary formation

Kemendo's alignment cost measures the energetic price of keeping components pointed in the system's common direction.

$$C_{\text{align}}(t) = \sum_{i=1}^n \theta_i(t) \frac{E_{c,i}(t)}{E_{s,i}(t)}. \quad (4)$$

In nontechnical terms: a misaligned component is expensive, especially if it consumes a lot of energy and has little stored reserve. In the limiting case where $E_{s,i}(t) = 0$, the component cannot realign itself at all, so its alignment cost is effectively infinite.

This formula underlies Kemendo's picture of boundary formation. Interior components are relatively aligned and cheap to maintain. Boundary components are more exposed, more mediative, more variable, and more expensive. They are valuable because they enable exchange; they are fragile because they absorb constraint.

2.5 Projection and system presentation

Kemendo also distinguishes actual energetic influence from visible or perceived influence. A component can project a presentation toward another system. This distinction matters because interactions are mediated not only by what a component is, but by how it appears at a boundary.

A GTC-style projection map can be written:

$$\psi_{i \rightarrow j}(t) = I(M_i(t), E_{\text{available},i}(t), E_{s,i}(t), \text{Surface}_i(t), \text{Abstractions}_i(t)), \quad (5)$$

where I is an interpretation function.

The sign of $\psi_{i \rightarrow j}$ indicates the apparent posture of the interaction:

$$\psi_{i \rightarrow j}(0) \begin{cases} > 0, & \text{mutual or cohesion-gaining,} \\ = 0, & \text{neutral,} \\ < 0, & \text{ablative or cohesion-losing.} \end{cases} \quad (6)$$

This is one of the places where Kemendo's theory already reaches beyond physics-style energy accounting. The projection map is implicitly semiotic: systems interact through interpretable boundary presentations, not merely through energy exchange.

2.6 Self-modeling and intelligence

GTC becomes more interesting when the system does not merely react. A system can contain an internal model of itself and use that model to simulate possible futures. Kemendo calls the cost of maintaining this internal predictive model $\tau(t)$.

A simulated future receives predicted usable return:

$$U(\hat{S}_r(t+1)) = \hat{R}_{\text{expected}}(t+1) - \hat{C}_{\text{align}}(t+1) - \hat{\tau}(t+1). \quad (7)$$

In plain language: a future is not good merely because it contains reward. It is good only if the system can afford the alignment and modeling costs required to reach and maintain it.

This leads to a functional definition of intelligence within GTC. A system becomes intelligent when it uses an internal model to choose actions that preserve future boundary viability, rather than only reacting after damage has already occurred.

2.7 The cohesion function

Kemendo's main scalar viability function is:

$$\gamma(S_i, t) = \frac{R_{\text{actual}}(t-1) - C_{\text{align}}(t) - B_t(t') - \tau(t) + U(\hat{S}_r(t))}{\|\nabla V(r, t) + \nabla \Psi(S_i, S_j, t)\|}. \quad (8)$$

The numerator combines past realized gain, present structural cost, baseline maintenance, self-modeling cost, and predicted future usable reward. The denominator penalizes sensitivity to internal and external gradients, so that a system poised on a knife edge counts as less viable than one with the same net return and more slack.

A cumulative version over systems and time is:

$$\Gamma(\{S_1, \dots, S_n\}, [t_0, t_1]) = \int_{t_0}^{t_1} \sum_{i=1}^n \gamma(S_i, t) dt. \quad (9)$$

The intended interpretation is simple: γ is a momentary vital sign; Γ is the trajectory-level viability of a system or system ensemble.

2.8 Where the formalism is overly restrictive

Kemendo's formalism is valuable as an engineering sketch, but it builds in several strong simplifying assumptions that limit its generality:

1. It assumes Euclidean action vectors are the right universal primitive.
2. It treats alignment as an angle rather than as a general relation among constraints.
3. It treats energy, reward, cost, and viability as scalar quantities.
4. It treats the boundary as a component role rather than as a graded self/other cut.
5. It compresses inter-system interaction into a scalar function $\Psi(S_i, S_j)$.
6. It treats simulation cost $\tau(t)$ as an external scalar rather than as the cost of maintaining an approximate correspondence between real and simulated dynamics.

None of these assumptions undermines the underlying insight. The rest of this paper keeps the insight and generalizes the mathematics, so that each of these restrictions becomes an optional modeling choice rather than a built-in limitation.

3 Key Ideas from Hyperseed, Quantale Weakness, Pareto Weakness, and TransWeave

Before stating the generalized theory, we briefly introduce the four bodies of ideas it draws on. Each supplies a replacement for one of GTC’s restrictive primitives: Hyperseed supplies graded, conflict-tolerant boundaries and a multi-channel notion of resource; quantale weakness supplies a general calculus of graded costs and distinctions; Pareto weakness supplies multi-objective viability; and TransWeave supplies a precise theory of when structure can be transported between systems.

3.1 Hyperseed: distinctions, p-bits, patterns, genenergy, and self/other cuts

Hyperseed begins from the idea that cognition and reality-modeling are built from distinctions, patterns, processes, contexts, effort, and observer-relative interpretation. The important point for present purposes is that a system does not merely have facts. It has evidence states, limited resources, boundaries, and patterns that may be partly coherent and partly conflicting.

3.1.1 P-bits

A p-bit records positive and negative support separately:

$$p(\varphi) = (p^+(\varphi), p^-(\varphi)) \in [0, 1]^2. \quad (10)$$

The key feature of this representation is that it allows a system to hold genuine tension. Something can be partly self-like and partly other-like, partly supported and partly opposed, partly coherent and partly dissonant. Contradiction does not explode the system; it is simply recorded and managed.

3.1.2 Distinctions and weakness

Hyperseed treats information as tied to distinctions. To know something is, in part, to know which differences matter. But bounded systems cannot preserve all distinctions; they must continually decide which distinctions are worth the cost of maintaining.

This is where weakness enters. A weaker representation is one that makes fewer unnecessary distinctions. It remains broad, flexible, and cheap, while still preserving the distinctions needed for the task at hand. Weakness, in this sense, is not sloppiness but disciplined parsimony.

3.1.3 Genenergy

Where Kemendo speaks of energy, Hyperseed uses a broader notion: genenergy, or generalized capacity-for-action. Genenergy may include physical energy, computation, attention, memory, coordination bandwidth, social trust, and any other resource a system needs in order to act.

The natural move, then, is to generalize GTC’s energy terms from scalar energy to multi-channel genenergy.

3.1.4 Self/other boundary

In Hyperseed terms, a self is not just a spatial interior. It is a synergetic, self-maintaining grouping with coordinated intent, and its boundary is observer- and context-relative. Something may be partially inside and partially outside at the same time.

This suggests replacing Kemendo’s boundary component role with a p-bit-valued self/other cut:

$$\sigma_t : X_t \rightarrow [0, 1]^2, \quad \sigma_t(x) = (\sigma_t^+(x), \sigma_t^-(x)). \quad (11)$$

Here $\sigma_t^+(x)$ is evidence that x belongs to the system, while $\sigma_t^-(x)$ is evidence that x belongs to the environment or other.

3.2 Quantale weakness

A quantale is an algebraic structure for composing graded values. One can think of it as a flexible generalization of the nonnegative reals used for costs: it has an order for comparing values, an operation for combining sequential or interacting quantities, and joins for aggregating alternatives.

A commutative unital quantale is a tuple:

$$Q = (Q, \leq, \oplus, \otimes, 0, 1), \quad (12)$$

where \oplus aggregates alternatives and \otimes composes values.

In this setting, weakness theory typically measures how many important pairs are left undistinguished by a hypothesis, relation, partition, or model. A generic pattern is:

$$w(H) = \bigoplus_{(u,v) \in H} \mu(u) \otimes \mu(v), \quad (13)$$

or, in a relation-weighted version,

$$w(H) = \bigoplus_{(u,v) \in H} \mu(u, v). \quad (14)$$

The intuition is the one introduced above: the weaker model is the one that avoids needless distinctions while still doing the job. It is not sloppy. It is maximally permissive subject to adequacy.

3.3 Quantale-valued process graphs

With quantales in hand, Kemendo’s components can be generalized to a quantale-valued pattern-flow graph:

$$G_t : X_t \times X_t \rightarrow Q. \quad (15)$$

Here X_t is a set of current processes, components, agents, symbols, cells, services, institutions, or other nodes. The value $G_t(x, y)$ is not just a number. It may record resource flow, evidence, weakness, interface cost, modeling cost, value effect, and boundary pressure at once.

A useful product quantale is:

$$Q = Q_{\text{gen}} \times Q_{\text{evid}} \times Q_{\text{weak}} \times Q_{\text{bdry}} \times Q_{\text{model}} \times Q_{\text{value}}. \quad (16)$$

The point of the product structure is that every relation can carry multiple channels of meaning simultaneously. A component can be cheap but unreliable, evidentially strong but value-dangerous, or boundary-stabilizing but model-expensive. Scalar energy cannot express any of this without hiding important tradeoffs.

3.4 Pareto weakness

Once costs and benefits are multi-channel, optimization changes character. Real systems rarely optimize one scalar. They sit on a Pareto frontier: one cannot improve one important channel without worsening another.

Let $K(A) \in \mathbb{R}_+^m$ be a vector of costs and $I(A)$ be information, evidence, or task success. The attainable image is:

$$\text{Att} = \{(I(A), K(A)) : A \text{ feasible}\}. \quad (17)$$

A point is Pareto efficient when no other feasible point has more evidence and no greater cost, with at least one strict improvement.

In the product-quantale framing, weakness can be written as:

$$W(f) = \iota(f) \otimes \phi(\kappa(f)), \quad (18)$$

where $\iota(f)$ measures evidence or adequacy, $\kappa(f)$ measures cost, and ϕ maps cost into a penalty in the evidence quantale.

The guiding principle can be stated in one line:

Use the weakest model that still preserves the evidence, and spend explicit coupling cost only when the evidence forces you to.

At smooth Pareto points, optimal allocation equalizes marginal gain per unit cost. Abstractly, if budget is split between two routes A and B , then:

$$\lambda_I \frac{\partial I}{\partial k_A} = \lambda_K \cdot \frac{\partial K}{\partial k_A}, \quad \lambda_I \frac{\partial I}{\partial k_B} = \lambda_K \cdot \frac{\partial K}{\partial k_B}. \quad (19)$$

Where the frontier is nonsmooth, this should be read in normal-cone or subgradient form rather than as ordinary calculus.

3.5 TransWeave

TransWeave is a theory of transfer through approximate commutation. The basic question it answers is: when can a solution, policy, model, or value function from one system be transported to another without breaking the dynamics?

The intuition is as follows. Suppose you learn in one domain and translate to another. It is not enough that the states look similar. What matters is that updating first and translating later approximately agrees with translating first and updating later. When this holds, the translation respects the dynamics rather than merely the snapshots.

Let B_1 and B_2 be update operators for two systems, and let $T : X_1 \rightarrow X_2$ be a transport or translation map. A TransWeave condition is:

$$T \circ B_1 \approx B_2 \circ T. \quad (20)$$

The defect is:

$$\Delta_{\text{TW}}(T; B_1, B_2) = \|T \circ B_1 - B_2 \circ T\|. \quad (21)$$

In the Bellman-Darboux version, where B_1 and B_2 are Bellman operators, this becomes:

$$R(T; M_1 \rightarrow M_2) = T \circ B_1 - B_2 \circ T. \quad (22)$$

If the residual is small, transferred values, policies, or solutions degrade in a controlled way as they cross the interface. If it is large, transfer is unreliable no matter how similar the two systems may appear superficially.

3.6 Geodesic or Wu Wei control

The final ingredient is a control principle. The geodesic control idea is that good action is not maximum forcing but least unnecessary forcing. The system should move from what it has to what it needs along a low-effort path, preserving the relevant resources along the way.

A Schrodinger-bridge-style action functional is:

$$SB(\rho, v) = \int_{t_0}^{t_1} \int_X \rho(x, t) \left(\frac{1}{2} \|v(x, t)\|^2 + \epsilon \log \frac{\rho(x, t)}{p(x)} \right) dx dt, \tag{23}$$

subject to:

$$\partial_t \rho + \nabla \cdot (\rho v) = 0, \quad \rho(\cdot, t_0) = \rho_0, \quad \rho(\cdot, t_1) = \rho_1. \tag{24}$$

The optimal bridge has intermediate density proportional to a forward factor and backward factor:

$$\rho(x, t) \propto f(x, t)g(x, t). \tag{25}$$

In words: prefer states that are easy to reach from the past and useful for reaching the future, while keeping effort smooth and conserving the resource that matters.

4 A More General Version: Weak-TransWeave Cohesion

We now assemble the ingredients of the previous section into a generalized theory of cohesion. The strategy is to take each of GTC's primitives in turn and show what it becomes in the richer setting.

4.1 From components to process nodes

Kemendo starts with components carrying action vectors. The more general theory starts with processes in a quantale-valued network.

A process node may be a molecule, cell, neuron, agent, API service, institution, argument, proof state, habit, or symbolic pattern. What matters is not its physical type but the role it plays in a pattern-flow relation.

Definition 1 (Quantale-valued pattern-flow substrate). *A system substrate at time t is a pair (X_t, G_t) where X_t is a set of process nodes and*

$$G_t : X_t \times X_t \rightarrow Q \tag{26}$$

assigns a quantale-valued relation to each ordered pair.

Kemendo's vector field is recovered from this substrate only after choosing a Euclidean embedding and projecting G_t into vector coordinates. In other words, the vector picture is one representation among many, not the foundation.

4.2 From alignment to distinction management

Kemendo's alignment angle θ_i measures how far a component points away from the system vector. But in a general cognitive, biological, social, or symbolic system, coherence is not always geometric. The deeper question is:

Which distinctions must be preserved, and which distinctions can be collapsed without harming viability?

Answering this question costs resources in both directions, which gives a general replacement for alignment cost.

Let $D_{S,t}^{\text{needed}}$ be distinctions needed for viability, and let $D_{S,t}^{\text{preserved}}$ be distinctions currently preserved by the system. Let $D_{S,t}^{\text{unnecessary}}$ be distinctions that are not worth preserving, and $D_{S,t}^{\text{over}}$ be distinctions the system is currently over-modeling.

Define a distinction-management cost:

$$C_{\text{dist}}(S, t) = \text{Cost}_Q \left(D_{S,t}^{\text{needed}} \setminus D_{S,t}^{\text{preserved}} \right) \oplus \text{Cost}_Q \left(D_{S,t}^{\text{unnecessary}} \cap D_{S,t}^{\text{over}} \right). \quad (27)$$

The first term penalizes losing distinctions that matter; the second penalizes wasting resources on distinctions that do not. Together they capture the two ways a system can fail at distinction management: under-modeling and over-modeling.

Kemendo's C_{align} is then a geometric special case:

$$C_{\text{align}}(t) \approx s(C_{\text{dist}}(S, t)) \quad (28)$$

for a scalarization $s : Q \rightarrow \mathbb{R}$ and a Euclidean distinction policy in which misalignment angle proxies distinction loss.

4.3 From boundary components to a self/other cut

Kemendo treats boundary components as components with high mediation load and high exposure to external constraint. This is useful, but too concrete: it identifies the boundary with particular parts rather than with a graded property the system maintains. The more general object is a self/other cut.

Definition 2 (Paraconsistent self/other cut). *A system boundary at time t is a p -bit-valued function*

$$\sigma_t : X_t \rightarrow [0, 1]^2, \quad \sigma_t(x) = (\sigma_t^+(x), \sigma_t^-(x)), \quad (29)$$

where $\sigma_t^+(x)$ is evidence that x is self/internal and $\sigma_t^-(x)$ is evidence that x is other/external.

Under this definition, interior nodes have high positive and low negative support, exterior nodes have low positive and high negative support, and boundary nodes carry significant positive and negative support at once. The boundary is thus not a separate kind of component but a region of graded, simultaneous self-and-other evidence.

Kemendo's boundary components reappear when this p -bit cut is sharpened into three zones:

$$\text{inside}, \quad \partial S, \quad \text{outside}. \quad (30)$$

4.4 From inter-system cost Ψ to TransWeave defect

Kemendo's $\Psi(S_i, S_j)$ measures the way one system imposes alignment or mediation costs on another. The more general construction replaces this scalar with an approximate intertwining between the two systems' update dynamics.

Let:

$$B_i : X_i \rightarrow X_i, \quad B_j : X_j \rightarrow X_j \quad (31)$$

be update operators for two systems. Let

$$T_{ij} : X_i \rightarrow X_j \quad (32)$$

be a boundary translation, interface map, bridge, protocol, perception map, or semantic transport. The TransWeave defect is:

$$\Delta_{\text{TW}}(S_i, S_j) = \|T_{ij}B_i - B_jT_{ij}\|. \quad (33)$$

The interpretation is the same as before: interaction works well when “update then translate” is close to “translate then update.” If this fails, the interface is unstable, and the instability is now quantified rather than merely asserted.

Kemendo’s interaction term is recovered as a scalarization:

$$\Psi(S_i, S_j) \approx \lambda_{\text{TW}}\Delta_{\text{TW}}(S_i, S_j) + \lambda_{\text{res}}K_{\text{transport}}(S_i, S_j) + \lambda_{\text{bdry}}K_{\text{repair}}(S_i, S_j). \quad (34)$$

A pleasant side effect of the operator framing is that it yields cleaner classifications of inter-system relationships:

- mutualism: low bidirectional defect and net viability gain for both systems;
- parasitism: low defect and gain in one direction, high cost or loss in the reverse direction;
- competition: no available low-cost transweave preserves both systems’ viability;
- collapse: all feasible boundary transweaves exceed the system’s repair budget.

4.5 From self-modeling cost to self-transweaving

Kemendo’s $\tau(t)$ is the cost of maintaining an internal predictive model. The abstract version asks what that cost is actually paying for, and the answer is: keeping real dynamics and simulated dynamics in approximate correspondence.

Let B_S be the real update operator and \widehat{B}_S the internal simulator. Let:

$$R : X_S \rightarrow \widehat{X}_S \quad (35)$$

encode registration or modeling, and let:

$$E : \widehat{X}_S \rightarrow X_S \quad (36)$$

encode enactment or control.

The self-model is good when:

$$RB_S \approx \widehat{B}_SR, \quad B_SE \approx E\widehat{B}_S. \quad (37)$$

Define the self-transweave defect:

$$\Delta_{\text{self}}(S) = \left\| RB_S - \widehat{B}_SR \right\| \oplus \left\| B_SE - E\widehat{B}_S \right\|. \quad (38)$$

Then the generalized modeling cost is:

$$\tau_Q(t) = \inf \left\{ K_{\text{model}}(R, E, \widehat{B}_S) : \Delta_{\text{self}}(S) \leq \varepsilon \right\}. \quad (39)$$

Put plainly: intelligence costs resources because the system must keep its internal simulator woven to the real dynamics closely enough to guide action. Modeling cost is not an arbitrary tax; it is the price of a structural correspondence.

4.6 From scalar viability to a quantale-valued viability profile

Kemendo's γ compresses everything into one number. The general theory resists this compression and keeps the channels separate, for the same reason a physician tracks several vital signs rather than one composite health score.

Definition 3 (Quantale-valued viability profile). *For a system S , interval I , and policy or trajectory π , define*

$$\mathcal{V}(S, I, \pi) = (G(S, I, \pi), K(S, I, \pi), W(S, I, \pi), \Delta_{\text{TW}}(S, I, \pi), R_{\text{bdry}}(S, I, \pi), M(S, I, \pi)) \in Q. \quad (40)$$

Here G is goal or value support, K is genenergy cost, W is weakness or simplicity, Δ_{TW} is transweave defect, R_{bdry} is boundary repair burden, and M is modeling cost.

A system is viable when its profile is non-dominated:

$$S \text{ is viable over } I \iff \mathcal{V}(S, I, \pi) \text{ lies on the feasible Pareto frontier.} \quad (41)$$

Kemendo's scalar cohesion function is recovered by choosing a monotone scalarization:

$$s : Q \rightarrow \mathbb{R}, \quad \gamma(S, t) = s(\mathcal{V}(S, \{t\}, \pi)). \quad (42)$$

This is the central conceptual upgrade: viability is not fundamentally one number. The one-number version is a decision or visualization convenience, obtained only after a choice of priorities has been made.

4.7 Weak-TransWeave cohesion

All the pieces are now in place, and the general theory can be stated in a single definition.

Definition 4 (Weak-TransWeave cohesion). *A system S over interval I is weak-transweave cohesive when there exist:*

- a quantale-valued pattern-flow graph $G_t : X_t \times X_t \rightarrow Q$;
- a p -bit self/other cut $\sigma_t : X_t \rightarrow [0, 1]^2$;
- an internal update operator B_S ;
- boundary transweave maps T_{S_j} to relevant adjacent systems;
- a self-model update \widehat{B}_S with registration and enactment maps R, E ;

such that the resulting viability profile $\mathcal{V}(S, I, \pi)$ is Pareto non-dominated and the boundary and self-transweave defects remain below resource-dependent tolerances.

In plain language: a cohesive system is one that keeps a workable self/other boundary, preserves only the distinctions it needs, and maintains low-defect translations among its internal dynamics, external interfaces, and simulated futures.

4.8 General action principle

Finally, Kemendo’s rollout selection can be generalized into a least-unnecessary-effort control problem in the geodesic spirit of Section 3.

Let ρ_t be the system’s state or belief distribution, σ_t its self/other cut, T_t its boundary transweave, \widehat{B}_t its simulator, and P_t its passive or natural dynamics.

Define:

$$\mathcal{A}[\rho, T, \sigma] = \int_I \left[K_{\text{move}}(\dot{\rho}_t) \oplus K_{\text{resist}}(\rho_t \| P_t) \oplus K_{\text{bdry}}(\dot{\sigma}_t) \oplus K_{\text{model}}(\widehat{B}_t) \oplus \Delta_{\text{TW}}(T_t) \ominus G_t \right] dt. \quad (43)$$

Then cohesive action is:

$$(\rho^*, T^*, \sigma^*) \in \arg \min_{\rho, T, \sigma} \mathcal{A}[\rho, T, \sigma] \quad (44)$$

subject to viability, value, and boundary constraints.

Here $\ominus G_t$ should be read schematically: goal/value support reduces the net burden in whatever residuated or scalarized form is appropriate for the chosen quantale. If the quantale does not support subtraction, this is handled by moving G_t into a separate Pareto coordinate rather than literally subtracting.

5 Kemendo GTC as a Special Case

The general theory is not meant to discard Kemendo’s equations. On the contrary, it explains why they are plausible as a first approximation: they are what the general theory looks like after a specific stack of simplifying choices.

Theorem 1 (Kemendo GTC as Euclidean scalarization). *Let Q be a product quantale of resource, evidence, weakness, boundary, model-cost, and value channels. Let $G_t : X_t \times X_t \rightarrow Q$ be a quantale-valued pattern-flow graph with p -bit self/other cut σ_t , update operator B_t , boundary maps T_{ij} , and self-model update \widehat{B}_t . Choose:*

1. a Euclidean embedding $e : X_t \rightarrow \mathbb{R}^d$;
2. a scalarization $s : Q \rightarrow \mathbb{R}$;
3. a rank-one approximation $G_t(i, j) \mapsto f(i, j, t) \widehat{A}_i(t)$;
4. a scalar energy channel $E_s - E_c$;
5. a scalar mediation channel M_i ;
6. a scalar boundary-defect proxy $\Psi(S_i, S_j)$;
7. a one-step rollout approximation to geodesic control.

Then weak-transweave cohesion reduces approximately to Kemendo-style local adjacency, alignment cost, and scalar cohesion dynamics:

$$L_{x_j}(t) = \sum_{i \in N(j)} f(i, j, t) \widehat{A}_i(t), \quad (45)$$

$$C_{\text{align}}(t) = \sum_i \theta_i(t) \frac{E_{c,i}(t)}{E_{s,i}(t)}, \quad (46)$$

$$\gamma(S_i, t) = \frac{R_{\text{actual}}(t-1) - C_{\text{align}}(t) - B_t(t') - \tau(t) + U(\widehat{S}_r(t))}{\|\nabla V(r, t) + \nabla \Psi(S_i, S_j, t)\|}. \quad (47)$$

Proof sketch. The Euclidean embedding converts process nodes into vector coordinates. The rank-one approximation replaces the full quantale-valued relation with a scalar coupling multiplying a normalized action vector. Scalar energy and mediation channels recover Kemendo’s component tuple. The p-bit boundary cut is sharpened into interior, boundary, and exterior components. The TransWeave defect is compressed into Ψ . The Pareto viability profile is converted into one real number by s . Finally, the geodesic control problem is truncated to one-step rollout evaluation. These reductions yield Kemendo’s formulas as projected shadows of the general construction. \square

6 Comparison Table

The following table summarizes the translation between GTC’s primitives and their weak-transweave generalizations, item by item.

Kemendo GTC	Weak-TransWeave Cohesion
Euclidean action vectors	Quantale-valued process relations or enriched morphisms. A vector is only one representation of a process relation.
Component tuple ($\ A_x\ , \theta_x, E_s, E_c, M_x$)	Process node with multi-channel state: genenergy, evidence, weakness, boundary role, model cost, value effect, and context.
Alignment angle θ_i	Cost of preserving or relaxing distinctions. Alignment is a geometric proxy for distinction-management.
Scalar stored and consumed energy	Multi-channel genenergy, including physical energy, attention, compute, memory, trust, coordination bandwidth, and repair capacity.
Boundary components	Paraconsistent self/other cut $\sigma_t : X_t \rightarrow [0, 1]^2$, with boundary nodes carrying simultaneous self-like and other-like evidence.
Projection $\psi_{i \rightarrow j}$	Semiotic or interface map from internal state to externally interpretable presentation; can be modeled as a context-indexed morphism.
Inter-system interaction $\Psi(S_i, S_j)$	TransWeave commutator or boundary defect $\ T_{ij}B_i - B_jT_{ij}\ $, plus transport, repair, and resource costs.
Self-modeling cost $\tau(t)$	Cost of maintaining low self-transweave defect between real dynamics B_S and simulated dynamics \hat{B}_S .
Scalar cohesion $\gamma(S, t)$	Quantale-valued viability profile $\mathcal{V}(S, I, \pi)$, evaluated by Pareto non-domination or by a chosen scalarization.
Reward filtered by viability	Multi-channel value/evidence/genenergy tradeoff; reward is a derived scalar, not the primitive objective.
Rollout-based planning	Geodesic or Wu Wei control: choose least-unnecessary-effort paths preserving viability and boundary coherence.
Collapse from high alignment/modeling cost	Collapse from leaving the feasible Pareto frontier: boundary repair, modeling, or transweave defect exceeds available genenergy.
System persistence as aligned component motion	System persistence as stable self-weaving pattern flow under resource-bounded approximate morphisms.
CML as markup for components, variables, interactions, rollouts	A possible front-end for declaring a special scalar/vector approximation of the richer quantale-valued system.

7 What the Generalization Buys

It is fair to ask what is gained by all this added abstraction. The answer is that each generalization removes a specific limitation of the original formalism.

7.1 It removes the false universality of vector alignment

Vector alignment is useful in many physical and engineering models, but it is not a universal primitive. Social, cognitive, symbolic, and biological systems often cohere by preserving constraints, interfaces, conventions, codes, memories, or causal dependencies. Quantale-valued process relations can express all of these without forcing them into angles.

7.2 It treats boundaries as graded and conflict-tolerant

Real boundaries are often ambiguous. A microbiome, an outsourced team, a satellite service, a memory prosthesis, or a cultural institution may be partly self and partly other. P-bit self/other cuts model this directly.

7.3 It makes intelligence a self-transweave condition

Kemendo’s intuition that intelligence involves internal simulation is strong. The more exact formulation is that an intelligent system maintains low-defect maps between real dynamics, simulated dynamics, and enacted control. This recasts modeling cost as a structural correspondence cost rather than an arbitrary scalar.

7.4 It makes interaction an operator-theoretic interface problem

Kemendo’s Ψ says that other systems impose constraint. TransWeave says how to measure whether two systems can update through each other without breaking coherence. The operator formulation is both more precise and, as the next section shows, more compositional.

7.5 It replaces scalar viability with Pareto viability

A system can be good on one channel and bad on another. It can gain reward while losing boundary integrity, or simplify while losing necessary distinctions. A Pareto profile preserves this structure, whereas a scalar γ is meaningful only after a choice of priorities has been made.

8 Formal Cohesion Certificates and Compositionality Theorems

8.1 Plain-language orientation

The previous sections define weak-transweave cohesion conceptually. This section turns the concept into checkable certificates. The guiding idea is simple: if a system claims to remain itself across time, compose with another system, or transfer a policy across an interface, then there should be a map witnessing that the relevant update dynamics nearly commute. The failure of the diagram to commute is the defect.

On this view, cohesion is not a vague holistic property. It is a bundle of witnesses: self/other cuts, update operators, weakness-resource profiles, transweave maps, and explicit tolerances. Once cohesion is expressed this way, compositionality becomes a theorem rather than a metaphor.

8.2 Normed update model

We first state the operator-theoretic part of the theory in ordinary normed spaces. This is not the only possible setting, but it is the cleanest one for deriving explicit error bounds, and it covers most cases of practical interest.

Definition 5 (Update system). *An update system is a triple*

$$A = (X_A, B_A, \|\cdot\|_A), \quad (48)$$

where X_A is a normed vector space and $B_A : X_A \rightarrow X_A$ is an update operator. In applications, B_A may be a Bellman operator, a belief-update operator, a learning update, a policy improvement operator, a self-model update, or a coarse-grained physical dynamics.

Definition 6 (TransWeave defect). *Let $A = (X_A, B_A)$ and $B = (X_B, B_B)$ be update systems and let $T : X_A \rightarrow X_B$ be a bounded linear map. The TransWeave defect of T is*

$$\Delta(T; A, B) = \|TB_A - B_B T\|_{A \rightarrow B}. \quad (49)$$

If $\Delta(T; A, B) \leq \varepsilon$, then T is called an ε -transweave from A to B .

In plain language: a small defect means that translating after an update is almost the same as updating after translation, so the map T respects the dynamics rather than just relabeling states.

8.3 Cohesion certificates

Definition 7 (Weak-transweave cohesion certificate). *A weak-transweave cohesion certificate for a system S over an interval I consists of data*

$$\text{Cert}(S, I) = (X_t, G_t, \sigma_t, B_t, \widehat{B}_t, R_t, E_t, \{T_{tj}\}, \mathcal{V}_t, \Theta_t)_{t \in I}, \quad (50)$$

where:

- X_t is the state or process-node space;
- $G_t : X_t \times X_t \rightarrow Q$ is a quantale-valued pattern-flow graph;
- $\sigma_t : X_t \rightarrow [0, 1]^2$ is a p-bit self/other cut;
- $B_t : X_t \rightarrow X_t$ is the actual update operator;
- $\widehat{B}_t : \widehat{X}_t \rightarrow \widehat{X}_t$ is the internal simulator update;
- $R_t : X_t \rightarrow \widehat{X}_t$ and $E_t : \widehat{X}_t \rightarrow X_t$ are registration and enactment maps;
- T_{tj} are boundary or inter-system transweave maps;
- $\mathcal{V}_t \in Q$ is the quantale-valued viability profile;
- Θ_t is a collection of tolerances for boundary defect, self-model defect, transweave defect, weakness loss, and genenergy cost.

The certificate is valid at tolerance Θ_t when all named defects lie below their corresponding tolerance and \mathcal{V}_t is not dominated among feasible profiles of the same declared type.

The phrase “of the same declared type” is doing real work here. A web service, an organism, a proof search, and a society do not share the same feasible policy space. Cohesion claims must always be typed by the class of transformations and interfaces that are admissible.

8.4 Sequential composition

Theorem 2 (Sequential composition of TransWeave defects). *Let*

$$A_1 = (X_1, B_1), \quad A_2 = (X_2, B_2), \quad A_3 = (X_3, B_3) \quad (51)$$

be normed update systems. Let $T_1 : X_1 \rightarrow X_2$ and $T_2 : X_2 \rightarrow X_3$ be bounded linear maps with

$$\Delta(T_1; A_1, A_2) \leq \varepsilon_1, \quad \Delta(T_2; A_2, A_3) \leq \varepsilon_2. \quad (52)$$

Then the composite $T_2T_1 : X_1 \rightarrow X_3$ satisfies

$$\Delta(T_2T_1; A_1, A_3) \leq \|T_2\|\varepsilon_1 + \varepsilon_2\|T_1\|. \quad (53)$$

In particular, if T_1 and T_2 are nonexpansive, then

$$\Delta(T_2T_1; A_1, A_3) \leq \varepsilon_1 + \varepsilon_2. \quad (54)$$

Proof. Compute

$$T_2T_1B_1 - B_3T_2T_1 = T_2T_1B_1 - T_2B_2T_1 + T_2B_2T_1 - B_3T_2T_1 \quad (55)$$

$$= T_2(T_1B_1 - B_2T_1) + (T_2B_2 - B_3T_2)T_1. \quad (56)$$

Taking operator norms and using submultiplicativity gives

$$\|T_2T_1B_1 - B_3T_2T_1\| \leq \|T_2\|\|T_1B_1 - B_2T_1\| + \|T_2B_2 - B_3T_2\|\|T_1\|. \quad (57)$$

Substitute the two defect bounds. If $\|T_1\|, \|T_2\| \leq 1$, the simplified bound follows. \square

In words: cohesive transfer degrades additively along a chain, up to the expansion factors of the interfaces. Long chains of mediocre interfaces are therefore predictably worse than short chains of good ones, with the loss quantified at each step.

8.5 Parallel composition

Theorem 3 (Parallel composition of TransWeave defects). *Let $A_S = (X_S, B_S)$, $A_R = (X_R, B_R)$, $A'_S = (Y_S, B'_S)$, and $A'_R = (Y_R, B'_R)$ be update systems. Suppose bounded maps $T_S : X_S \rightarrow Y_S$ and $T_R : X_R \rightarrow Y_R$ satisfy*

$$\|T_S B_S - B'_S T_S\| \leq \varepsilon_S, \quad \|T_R B_R - B'_R T_R\| \leq \varepsilon_R. \quad (58)$$

Equip tensor products with a crossnorm satisfying $\|A \otimes C\| \leq \|A\|\|C\|$. Define

$$B_{S \otimes R} = B_S \otimes B_R, \quad B'_{S \otimes R} = B'_S \otimes B'_R, \quad T_{S \otimes R} = T_S \otimes T_R. \quad (59)$$

Then

$$\Delta(T_{S \otimes R}; A_S \otimes A_R, A'_S \otimes A'_R) \leq \varepsilon_S \|T_R B_R\| + \|B'_S T_S\| \varepsilon_R. \quad (60)$$

If $T_R B_R$ and $B'_S T_S$ are nonexpansive, then

$$\Delta(T_{S \otimes R}) \leq \varepsilon_S + \varepsilon_R. \quad (61)$$

Proof. Let $A = T_S B_S$, $A' = B'_S T_S$, $C = T_R B_R$, and $C' = B'_R T_R$. Then

$$(T_S \otimes T_R)(B_S \otimes B_R) - (B'_S \otimes B'_R)(T_S \otimes T_R) = A \otimes C - A' \otimes C' \quad (62)$$

$$= (A - A') \otimes C + A' \otimes (C - C'). \quad (63)$$

Take norms and use the crossnorm inequality:

$$\|A \otimes C - A' \otimes C'\| \leq \|A - A'\| \|C\| + \|A'\| \|C - C'\|. \quad (64)$$

Substitute the two one-factor defect bounds. □

This theorem justifies modularity: independently cohesive modules remain cohesive when run in parallel, provided the product operation does not introduce large new cross-dependencies.

8.6 Boundary gluing

The previous two theorems concern update operators. Boundaries are a different kind of object, and gluing two systems together requires a separate compatibility statement about their self/other cuts.

Definition 8 (Boundary compatibility defect). *Let (V, d_V) be a metric space of boundary values, for example $V = [0, 1]^2$ with an ℓ_1 or ℓ_∞ metric. Let $\sigma_S : X_S \rightarrow V$ and $\sigma_R : X_R \rightarrow V$ be two self/other cuts. If $I_S \subseteq X_S$, $I_R \subseteq X_R$, and $\iota : I_S \rightarrow I_R$ is a declared interface identification, define*

$$\delta_\partial(S, R; I) = \sup_{x \in I_S} d_V(\sigma_S(x), \sigma_R(\iota x)). \quad (65)$$

Theorem 4 (Boundary gluing bound). *Suppose systems S and R have internal boundary variation bounds $D_\partial(S)$ and $D_\partial(R)$, meaning that whenever two points are identified as adjacent inside the declared boundary region of the same system, their boundary values differ by at most the corresponding D_∂ . Let $S \cup_I R$ be the glued system obtained by identifying $x \in I_S$ with $\iota x \in I_R$. Then the glued boundary defect satisfies*

$$D_\partial(S \cup_I R) \leq \max\{D_\partial(S), D_\partial(R), \delta_\partial(S, R; I)\}. \quad (66)$$

In an additive cost convention the corresponding bound is

$$D_\partial(S \cup_I R) \leq D_\partial(S) + D_\partial(R) + \delta_\partial(S, R; I). \quad (67)$$

Proof. Every boundary comparison in the glued system is of one of three types: wholly inside S , wholly inside R , or across the glued interface. The first two are bounded by $D_\partial(S)$ and $D_\partial(R)$ by assumption. The cross comparison is bounded by $\delta_\partial(S, R; I)$ by definition. Taking the maximum gives the idempotent bound. Replacing maximum aggregation by additive aggregation gives the additive bound by the triangle inequality. □

The moral is that a social organism, multicellular organism, mindplex, web application, or federated AI is boundary-cohesive only when the glue between self/other cuts is itself low defect. Two internally tidy systems can still make an incoherent whole if they disagree about where the seam lies.

8.7 Supported Pareto composition

A tempting but false statement would be: “if every part is locally Pareto optimal, then the whole is globally Pareto optimal.” This fails in general because local tradeoff choices may use incompatible exchange rates between evidence, cost, weakness, and genenergy: one module may treat compute as cheap and trust as precious while a neighbor does the reverse. The correct theorem requires a shared supporting normal, or equivalently a compatible scalar selector.

Definition 9 (Profile order). *Let profiles have the form $q = (I, K) \in \mathbb{R} \times \mathbb{R}_+^m$, where larger I is better and smaller K is better. We write $q' \succ q$ if*

$$I(q') \geq I(q), \quad K(q') \leq K(q), \quad (68)$$

componentwise, with at least one strict inequality. A feasible profile is Pareto efficient if no feasible profile strictly dominates it.

Theorem 5 (Supported Pareto composition). *Let \mathcal{P}_i be feasible profile sets for typed subsystems and let \mathcal{P}_g be the feasible profile set for the glue. Assume composite profiles are generated additively by*

$$\Phi(q_1, \dots, q_n, q_g) = \left(\sum_{i=1}^n I_i + I_g, \sum_{i=1}^n K_i + K_g \right). \quad (69)$$

Let $\alpha > 0$ and $\beta \in \mathbb{R}_+^m$ have all entries strictly positive, and define the strictly monotone selector

$$L(q) = \alpha I(q) - \beta \cdot K(q). \quad (70)$$

If each q_i^ maximizes L over \mathcal{P}_i and q_g^* maximizes L over \mathcal{P}_g , then*

$$q^* = \Phi(q_1^*, \dots, q_n^*, q_g^*) \quad (71)$$

maximizes L over the product-generated composite feasible set. Consequently q^ is Pareto efficient in that typed composite feasible set.*

Proof. For any feasible tuple (q_1, \dots, q_n, q_g) ,

$$L(\Phi(q_1, \dots, q_n, q_g)) = \alpha \left(\sum_i I_i + I_g \right) - \beta \cdot \left(\sum_i K_i + K_g \right) \quad (72)$$

$$= \sum_i L(q_i) + L(q_g). \quad (73)$$

Each term is maximized by assumption, so the sum is maximized by the tuple of individual maximizers. If q^* were strictly dominated, the strict positivity of α and β would imply a strictly larger value of L , contradicting maximality. \square

In words: compositional cohesion requires not just locally good modules, but locally good modules selected at compatible tradeoff prices.

8.8 A compact compositionality slogan

In cost/error notation, the theorem family above can be summarized as

$$\text{Defect}(S \circ R) \leq \text{Defect}(S) \otimes \text{Defect}(R) \otimes \text{GlueDefect}(S, R), \quad (74)$$

with the exact meaning of \otimes determined by the chosen quantale or cost convention. In additive nonexpansive cases this reduces to the familiar bound

$$\text{Defect}(S \circ R) \leq \text{Defect}(S) + \text{Defect}(R) + \text{GlueDefect}(S, R). \quad (75)$$

In words:

Cohesion composes up to explicitly budgeted glue defect.

9 Cohesion and Hyperseed-Style General Intelligence

9.1 Plain-language orientation

Cohesion is not intelligence. A rock may be cohesive without being intelligent, and a brittle theorem prover may solve a narrow benchmark while having poor organismic cohesion. The useful claim is subtler:

Weak-transweave cohesion bounds the loss of effective general intelligence due to self-fragmentation, boundary drift, bad transfer, over-modeling, under-modeling, and interface friction.

General intelligence becomes visible when competence transfers across many kinds of tasks, representations, goals, time horizons, and collaborators. That transfer requires low-defect transweaves. Cohesion is therefore a substrate for robust general intelligence, not a synonym for it. This section makes the substrate relationship precise.

9.2 Task ecologies and GTGI-style scores

Definition 10 (Task ecology). *A task ecology is a probability space*

$$\mathcal{E} = (\Omega \times \mathcal{G} \times \mathcal{T}, \nu), \quad (76)$$

where Ω is a class of environments, \mathcal{G} a class of goals or value functionals, \mathcal{T} a class of horizons or resource regimes, and ν a probability measure over triples (ω, g, T) .

Definition 11 (GTGI-style score). *Let S be a cognitive system with induced policy or behavior π_S . Let*

$$V_{\omega, g, T}^S \in \mathbb{R} \quad (77)$$

be the value achieved by S on (ω, g, T) . The task-ecology score of S is

$$\Pi_{\mathcal{E}}(S) = \mathbb{E}_{(\omega, g, T) \sim \nu} [V_{\omega, g, T}^S]. \quad (78)$$

If cost is included, define the effective score

$$\Pi_{\mathcal{E}}^{\lambda}(S) = \mathbb{E}[V_{\omega, g, T}^S] - \lambda K(S), \quad (79)$$

where $K(S)$ is the relevant genenergy, computation, or attention cost.

These definitions deliberately leave open the exact GTGI measure. The theorems below apply to any measure expressible as an expected task value with a well-defined defect sensitivity.

9.3 No unconditional implication

Before stating positive results, it is worth recording the negative result that keeps the framework honest.

Proposition 1 (Cohesion does not imply general intelligence without a task ecology). *There is no nontrivial theorem of the form*

$$\text{Coh}(S) \geq c \quad \Rightarrow \quad \Pi_{\mathcal{E}}(S) \geq a \quad (80)$$

that holds uniformly for all task ecologies \mathcal{E} and all systems S , unless a is no larger than the trivial worst-case lower bound of the value scale.

Proof. Fix any cohesive system S with a restricted behavior repertoire. Choose a task ecology \mathcal{E} concentrated on tasks whose success condition requires an action or output that S never produces. Then $V_{\omega,g,T}^S$ equals the minimum value on the support of ν , so $\Pi_{\mathcal{E}}(S)$ is the trivial lower bound, regardless of $\text{Coh}(S)$. Therefore no positive lower bound follows from cohesion alone without assumptions on the task ecology. \square

This proposition is useful precisely because it blocks an overstatement. Cohesion becomes intelligence-relevant only when the task ecology rewards persistence, transfer, self-maintenance, adaptation, and synergy. The positive theorems below all carry such an ecological assumption.

9.4 Cohesion defect as intelligence leakage

Definition 12 (Weighted cohesion defect). *Let S^* be an ideal coherent version of a system S in the same typed architecture class. Define the normalized cohesion defect*

$$D_{\text{coh}}(S, S^*) = a_{\Delta} \Delta_{\text{TW}}(S, S^*) + a_{\partial} D_{\partial}(S, S^*) + a_W D_W(S, S^*) + a_K D_K(S, S^*) + a_M D_M(S, S^*), \quad (81)$$

where all weights are nonnegative. The terms measure transweave defect, boundary/self-other drift, weakness or distinction-management defect, excess genenergy cost, and self-modeling defect.

Assumption 1 (Task sensitivity to cohesion defect). *A task ecology \mathcal{E} is L -sensitive to cohesion defect if there is a measurable function $L(\omega, g, T)$ with finite expectation \bar{L} such that for all relevant systems S and idealizations S^* ,*

$$\left| V_{\omega,g,T}^S - V_{\omega,g,T}^{S^*} \right| \leq L(\omega, g, T) D_{\text{coh}}(S, S^*). \quad (82)$$

Theorem 6 (Cohesion bounds effective general-intelligence leakage). *If \mathcal{E} is L -sensitive to cohesion defect and $\bar{L} = \mathbb{E}[L(\omega, g, T)] < \infty$, then*

$$\Pi_{\mathcal{E}}(S) \geq \Pi_{\mathcal{E}}(S^*) - \bar{L} D_{\text{coh}}(S, S^*). \quad (83)$$

If cost is included and the cost gap satisfies $K(S) - K(S^) \leq D_K(S, S^*)$, then*

$$\Pi_{\mathcal{E}}^{\lambda}(S) \geq \Pi_{\mathcal{E}}^{\lambda}(S^*) - \bar{L} D_{\text{coh}}(S, S^*) - \lambda D_K(S, S^*). \quad (84)$$

Proof. From the sensitivity assumption,

$$V_{\omega,g,T}^S \geq V_{\omega,g,T}^{S^*} - L(\omega, g, T) D_{\text{coh}}(S, S^*). \quad (85)$$

Take expectations over ν . Since $D_{\text{coh}}(S, S^*)$ is fixed with respect to the task draw, it factors out of the expectation. This gives the first inequality. The effective-score inequality follows by subtracting $\lambda K(S)$ and using $K(S) \leq K(S^*) + D_K(S, S^*)$. \square

In plain terms: cohesion does not create intelligence from nothing, but it bounds how much intelligence leaks away when the actual system fails to realize its ideal coherent organization. The defect D_{coh} is exactly the leak rate.

9.5 Transfer-general intelligence theorem

General intelligence depends strongly on transfer. A system that learns one thing and cannot carry the result anywhere else is narrow, however impressive the single result may be. By contrast, a system whose task solutions live in a low-defect transweave atlas can generalize, because each mastered task radiates usable competence to its neighbors.

Definition 13 (Task transweave graph). *A task transweave graph is a directed graph $\mathcal{H} = (\mathcal{N}, \mathcal{E}_H)$ whose nodes are tasks or task families. An edge $e : i \rightarrow j$ carries a transfer map T_e and a defect $\varepsilon_e \geq 0$. A path $p = i_0 \rightarrow i_1 \rightarrow \dots \rightarrow i_k$ has transfer map $T_p = T_{i_{k-1}i_k} \dots T_{i_0i_1}$ and path defect*

$$d_\varepsilon(p) = \sum_{r=0}^{k-1} \varepsilon_{i_r i_{r+1}}. \quad (86)$$

Let $d_\varepsilon(i, j)$ be the shortest such path defect from i to j .

Assumption 2 (Edgewise transfer loss). *For each edge $e : i \rightarrow j$, transferred policies satisfy*

$$V_j(T_e \pi_i) \geq V_i(\pi_i) - L_e \varepsilon_e. \quad (87)$$

For simplicity assume $L_e \leq L$ uniformly.

Theorem 7 (Low-defect task atlases lower-bound transfer generality). *Let $M \subseteq \mathcal{N}$ be a set of mastered tasks, and for each task j let $i(j) \in M$ be a mastered task minimizing $d_\varepsilon(i, j)$. If the edgewise transfer-loss assumption holds with uniform L , then the policy obtained by transferring from the nearest mastered task satisfies*

$$V_j(T_{p_{i(j),j}} \pi_{i(j)}) \geq V_{i(j)}(\pi_{i(j)}) - L d_\varepsilon(i(j), j). \quad (88)$$

Consequently, for a task distribution γ on \mathcal{N} ,

$$\mathbb{E}_{j \sim \gamma} \left[V_j(T_{p_{i(j),j}} \pi_{i(j)}) \right] \geq \mathbb{E}_{j \sim \gamma} \left[V_{i(j)}(\pi_{i(j)}) - L d_\varepsilon(i(j), j) \right]. \quad (89)$$

Proof. Apply the edgewise transfer-loss bound along a path $p : i = i_0 \rightarrow \dots \rightarrow i_k = j$. Induction gives

$$V_{i_k}(T_p \pi_i) \geq V_{i_0}(\pi_i) - \sum_{r=0}^{k-1} L \varepsilon_{i_r i_{r+1}}. \quad (90)$$

Choosing a path of minimal defect gives the pointwise bound. Taking expectation over $j \sim \gamma$ gives the distributional bound. \square

This theorem is a precise version of a Hyperseed-flavored claim:

General intelligence increases when competence lives on a dense, low-defect transweave graph over task space.

9.6 Cognitive synergy threshold

Cohesion also matters when multiple cognitive modules are combined. A neural module, a symbolic module, a planner, a memory system, and a value model may each be individually useful, but their combination improves general intelligence only when the gains from cross-module interaction exceed the costs of maintaining the interfaces between them.

Definition 14 (Expected synergy and glue defect). *Let M_1, \dots, M_n be modules with individual policies π_i , and let $\pi_\otimes = \text{Glue}_\chi(\pi_1, \dots, \pi_n)$ be a composite policy. For a task instance $z = (\omega, g, T)$, define the pointwise synergy surplus*

$$\text{Syn}_z = V_z^{\pi_\otimes} - \max_i V_z^{\pi_i} + D_{\text{glue},z}(\chi). \quad (91)$$

Equivalently,

$$V_z^{\pi_\otimes} = \max_i V_z^{\pi_i} + \text{Syn}_z - D_{\text{glue},z}(\chi). \quad (92)$$

Theorem 8 (Synergy beats interface friction). *If*

$$\mathbb{E}_z[\text{Syn}_z] > \mathbb{E}_z[D_{\text{glue},z}(\chi)], \quad (93)$$

then

$$\Pi_{\mathcal{E}}(\pi_\otimes) > \mathbb{E}_z[\max_i V_z^{\pi_i}] \geq \max_i \Pi_{\mathcal{E}}(\pi_i). \quad (94)$$

Thus the composite is more generally intelligent, on this task ecology, than any single module in expectation.

Proof. Take expectations in

$$V_z^{\pi_\otimes} = \max_i V_z^{\pi_i} + \text{Syn}_z - D_{\text{glue},z}(\chi). \quad (95)$$

The strict inequality on expectations gives the first strict bound. The second inequality follows because the expectation of the pointwise maximum is at least the maximum of the expectations. \square

This is the mindplex theorem in miniature: a collective or modular cognitive system outperforms its parts when the value created by cross-module synergy is larger than the weak-transweave cost of maintaining the interfaces.

9.7 Cohesion as an effective-resource multiplier

Let $K_{\text{raw}}(S)$ be the cost of running the parts without coordination. Let $R_{\text{coh}}(S)$ be the resource saving from avoiding duplicated work, contradiction, double counting, boundary repair, and representational thrashing. Let $D_{\text{coh}}(S)$ be the residual cohesion defect. Then

$$K_{\text{actual}}(S) = K_{\text{raw}}(S) - R_{\text{coh}}(S) + D_{\text{coh}}(S). \quad (96)$$

For an effective score $\Pi^\lambda = \mathbb{E}[V] - \lambda K$, a cohesion improvement from S_0 to S_1 improves effective intelligence whenever

$$\mathbb{E}[V(S_1) - V(S_0)] + \lambda (R_{\text{coh}}(S_1) - R_{\text{coh}}(S_0)) > \lambda (D_{\text{coh}}(S_1) - D_{\text{coh}}(S_0) + K_{\text{raw}}(S_1) - K_{\text{raw}}(S_0)). \quad (97)$$

This is not a separate theorem so much as an accounting identity, but it clarifies why cohesion can raise intelligence without changing the nominal algorithm: it can free genenergy that was previously wasted on incoherence.

10 Cohesion and Robust Iterative Goal-Guided Self-Modification

10.1 Plain-language orientation

A self-modifying system changes the machinery that will make later changes. A goal-guided self-modifying system also changes, or may change, the very evaluators that judge later changes. This creates a drift problem: how can the system improve without losing the identity, values, boundary, or competence that made its improvement meaningful in the first place?

Weak-transweave cohesion gives a clean answer. Each self-modification step must come with a low-defect transweave certificate. The certificate says that the new self is not an arbitrary replacement of the old self, but a controlled deformation of the old dynamics, goals, and boundary. The theorems of this section show that when such certificates exist and their defects are summably bounded, drift stays bounded too.

10.2 Transported self-modification trajectories

Definition 15 (Self-modification trajectory). *A self-modification trajectory is a sequence*

$$S_0, S_1, S_2, \dots \quad (98)$$

where each

$$S_t = (X_t, B_t, G_t, \sigma_t, \pi_t) \quad (99)$$

contains a state space, cognition/update operator, goal structure, self/other cut, and policy or meta-policy. A transition $S_t \rightarrow S_{t+1}$ is equipped with transport maps

$$T_t^X : X_t \rightarrow X_{t+1}, \quad T_t^G : G_t \rightarrow G_{t+1}, \quad T_t^\sigma : \sigma_t \rightarrow \sigma_{t+1} \quad (100)$$

whenever those objects live in different representation spaces.

Definition 16 (Per-step self-modification defect). *Let $d^X, d^G, d^\sigma, d^W, d^K$ be the declared metrics or defect measures for state/update dynamics, goals, boundary cuts, weakness profiles, and genenergy profiles. Define*

$$D_t = a_X \|T_t^X B_t - B_{t+1} T_t^X\| + a_G d^G(T_t^G G_t, G_{t+1}) + a_\sigma d^\sigma(T_t^\sigma \sigma_t, \sigma_{t+1}) \quad (101)$$

$$+ a_W d^W(T_t^W W_t, W_{t+1}) + a_K d^K(T_t^K K_t, K_{t+1}), \quad (102)$$

with all weights nonnegative.

This is the formal version of a simple requirement: a good self-modification preserves what matters about cognition, goals, boundaries, weakness policy, and resource budget, up to an explicitly budgeted defect.

10.3 Bounded drift from summable defects

Assumption 3 (Nonexpansive transport atlas). *For each t , the transport from S_t to S_{t+1} is nonexpansive in the identity metric d_t in the following sense. If $T_{s:t}$ denotes the composed transport from time s to time t , then for all a, b in the same transported chart,*

$$d_{t+1}(T_t a, T_t b) \leq d_t(a, b). \quad (103)$$

Assume further that the per-step identity drift is bounded by D_t :

$$d_{t+1}(T_t S_t, S_{t+1}) \leq D_t. \quad (104)$$

Theorem 9 (Summable cohesion defects imply bounded identity drift). *Under the nonexpansive transport-atlas assumption,*

$$d_n(T_{0:n}S_0, S_n) \leq \sum_{t=0}^{n-1} D_t. \quad (105)$$

Consequently, if

$$\sum_{t=0}^{\infty} D_t < R, \quad (106)$$

and the robust identity basin around the transported initial self has radius R , then S_n remains inside that basin for every n .

Proof. The proof is by induction. For $n = 1$ the claim is the per-step drift bound. Assume it holds for n . Then, using the triangle inequality and nonexpansive transport,

$$d_{n+1}(T_{0:n+1}S_0, S_{n+1}) \leq d_{n+1}(T_n T_{0:n}S_0, T_n S_n) + d_{n+1}(T_n S_n, S_{n+1}) \quad (107)$$

$$\leq d_n(T_{0:n}S_0, S_n) + D_n \quad (108)$$

$$\leq \sum_{t=0}^n D_t. \quad (109)$$

The basin claim follows by taking the supremum over n and using the strict radius bound. \square

The lesson is that robustness is not achieved by forbidding change. It is achieved by making self-change a summably bounded sequence of low-defect deformations.

10.4 Goal drift bound

Theorem 10 (Bounded cohesion defect implies bounded goal drift). *Assume goal transports T_t^G are nonexpansive in a goal metric d_t^G , and assume for some $L_G \geq 0$ that*

$$d_{t+1}^G(T_t^G G_t, G_{t+1}) \leq L_G D_t. \quad (110)$$

Then

$$d_n^G(T_{0:n}^G G_0, G_n) \leq L_G \sum_{t=0}^{n-1} D_t. \quad (111)$$

In particular, if $\sum_t D_t < R$, then total goal drift is bounded by $L_G R$.

Proof. The proof is identical to the identity-drift proof, replacing S_t by G_t and D_t by $L_G D_t$. \square

This is a rigorous version of controlled value change. The theorem does not say goals never change. It says goal change remains a bounded deformation of the original goal structure.

10.5 Capability growth under drift loss

Assumption 4 (Gain-minus-drift improvement). *Let $\Pi(S_t)$ be a capability or GTGI-style score. Suppose each accepted self-modification step has proposed gain $g_t \geq 0$ and loses at most LD_t to cohesion drift:*

$$\Pi(S_{t+1}) - \Pi(S_t) \geq g_t - LD_t. \quad (112)$$

Theorem 11 (Robust cumulative self-improvement). *Under the gain-minus-drift assumption,*

$$\Pi(S_n) \geq \Pi(S_0) + \sum_{t=0}^{n-1} (g_t - LD_t). \quad (113)$$

If also $\sum_t D_t < R$, then the system remains inside the robust identity basin while achieving the above cumulative improvement. If

$$\sum_{t=0}^{\infty} g_t = \infty \quad \text{and} \quad \sum_{t=0}^{\infty} D_t < \infty, \quad (114)$$

then the lower bound diverges to $+\infty$. If Π is externally bounded above, then these assumptions cannot all hold indefinitely; the theorem should then be read as convergence toward the ceiling with finite total drift loss.

Proof. Sum the one-step inequalities from $t = 0$ to $n - 1$; the intermediate terms $\Pi(S_t)$ telescope. The identity-basin statement is the previous theorem. The divergence statement follows because a divergent nonnegative gain sum minus a finite drift-loss sum diverges. \square

This is the cleanest open-endedness result in the paper:

A system can improve indefinitely without losing itself if improvement gains accumulate while self-loss remains summably bounded.

10.6 Lyapunov-barrier formulation

Sometimes robustness is more naturally expressed not by a metric ball but by a barrier or Lyapunov function: a single potential whose level sets separate safe configurations from unsafe ones.

Theorem 12 (Cohesion barrier certificate). *Let $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ be a robustness potential, where unsafe systems satisfy*

$$\mathcal{R}(S) < \mathcal{R}_{\text{crit}}. \quad (115)$$

Suppose the self-modification trajectory satisfies

$$\mathcal{R}(S_{t+1}) \geq \mathcal{R}(S_t) - \rho_t \quad (116)$$

with $\rho_t \geq 0$. If

$$\mathcal{R}(S_0) - \sum_{t=0}^{\infty} \rho_t > \mathcal{R}_{\text{crit}}, \quad (117)$$

then S_t is never unsafe.

Proof. By induction,

$$\mathcal{R}(S_n) \geq \mathcal{R}(S_0) - \sum_{t=0}^{n-1} \rho_t. \quad (118)$$

The infinite-sum assumption implies

$$\mathcal{R}(S_n) > \mathcal{R}_{\text{crit}} \quad (119)$$

for every n , so S_n does not satisfy the unsafe condition. \square

A useful robustness potential has the schematic form

$$\mathcal{R}(S) = \alpha \text{GoalSim}(S, S_0) + \beta \text{SelfCont}(S, S_0) + \gamma \text{Weakness}(S) + \eta \text{GenEff}(S) - \lambda \text{TWDefect}(S, S_0). \quad (120)$$

The exact terms depend on the implementation. The theorem only requires that \mathcal{R} separates the unsafe region by a positive margin.

10.7 Stochastic robustness

Self-modifications proposed by search, program synthesis, evolution, or learning will usually come with uncertain defect estimates rather than exact values. The deterministic theorem therefore needs, and has, a high-probability analogue.

Theorem 13 (Finite-horizon stochastic defect bound). *Let D_0, \dots, D_{n-1} be nonnegative random per-step defects adapted to a filtration \mathcal{F}_t . Suppose*

$$0 \leq D_t \leq b_t, \quad \mathbb{E}[D_t \mid \mathcal{F}_t] \leq \delta_t. \quad (121)$$

Then for any $\alpha \in (0, 1)$, with probability at least $1 - \alpha$,

$$\sum_{t=0}^{n-1} D_t \leq \sum_{t=0}^{n-1} \delta_t + \sqrt{\frac{1}{2} \left(\sum_{t=0}^{n-1} b_t^2 \right) \log \frac{1}{\alpha}}. \quad (122)$$

Consequently, if the right-hand side is less than the robustness radius R , then the identity-drift bound holds with probability at least $1 - \alpha$.

Proof. Let

$$X_t = D_t - \mathbb{E}[D_t \mid \mathcal{F}_t]. \quad (123)$$

The sequence of partial sums of X_t is a martingale with increments bounded in an interval of length at most b_t . Hoeffding-Azuma gives

$$\Pr \left(\sum_t X_t \geq a \right) \leq \exp \left(-\frac{2a^2}{\sum_t b_t^2} \right). \quad (124)$$

Set the right side equal to α and solve for a . Since $\sum_t \mathbb{E}[D_t \mid \mathcal{F}_t] \leq \sum_t \delta_t$, the stated bound follows. The robustness claim follows from the deterministic bounded-drift theorem. \square

10.8 Design principle for self-modifying AGI

Taken together, the theorem family gives a practical proof obligation for each proposed self-modification:

Accept a self-modification only if it increases expected capability or value and comes with a weak-transweave certificate whose defect fits inside the remaining identity, goal, boundary, weakness, and genenergy budgets.

In symbols, a simple acceptance rule is

$$\Delta \Pi_t(S_{t+1}) - \lambda D_t - \mu C_t > 0, \quad D_t \leq \delta_t, \quad \sum_t \delta_t < R. \quad (125)$$

The first inequality says the modification is worth doing. The second says it is locally safe. The third says that local safety budgets compose into global robustness, which is precisely what the drift theorems guarantee.

11 Synthesis: Cohesion, Intelligence, and Self-Modification

The three theorem families can be read together as one larger picture.

First, compositionality says that systems can be built from parts if the glue between parts has controlled defect. This upgrades Kemendo’s notion of alignment: the important object is not just an angle between action vectors, but a typed transweave certificate showing that update dynamics, boundaries, and resource profiles remain compatible.

Second, the intelligence theorems say that cohesion is a substrate for general intelligence. General intelligence is not mere internal order. It is the ability to preserve, transport, and compound competence across a task ecology. Low cohesion defect bounds the gap between the system one designed or idealized and the system that actually operates under finite genenergy, boundary pressure, self-model error, and interface friction.

Third, the self-modification theorems say that open-ended improvement is not a matter of taking arbitrarily large leaps. It is a matter of accumulating gains while making self-loss summable. The system may change its own architecture, representations, and goals, but each accepted change must remain a controlled deformation of the previous self.

A compact formula for the whole theory is:

$$\text{robust open-ended intelligence} \approx \text{cumulative gain} - \text{cumulative cohesion defect}, \quad (126)$$

subject to a boundary condition:

$$\sum_t D_t < R_{\text{identity/value basin}}. \quad (127)$$

Thus the strong slogan is:

General intelligence is cohesive transfer of competence across task-space under resource, boundary, and self-continuity constraints.

12 Conclusion

Kemendo’s General Theory of Cohesion points in a fruitful direction: systems should be evaluated by what they can hold together under energy, boundary, alignment, and prediction constraints. Its limitation is not the intuition but the formal substrate. Euclidean vectors, scalar energy, scalar viability, and a single interaction function are too narrow to serve as a general theory of cohesion.

The more general framing proposed here is Weak-TransWeave Cohesion. It treats systems as quantale-valued pattern-flow networks with paraconsistent self/other cuts, weakness-guided distinction management, transweave defects at boundaries, self-transweave defects for internal simulation, and Pareto-valued viability profiles. In this setting, Kemendo’s formulas are not rejected. They are recognized as a crude but useful special case: a rank-one, Euclidean, scalarized, one-step approximation to a richer enriched-categorical theory of viable self-maintaining systems.

The theorem sections make the reframing mathematically useful rather than merely suggestive. The compositionality theorems show that cohesion survives sequential composition, parallel product, and boundary gluing when the appropriate transweave and glue defects are bounded. The intelligence theorems show that cohesion does not equal general intelligence, but does bound the loss of effective intelligence due to internal incoherence, poor transfer, boundary drift, excess distinction-making, and resource waste. The self-modification theorems show that iterative goal-guided self-modification remains robust when per-step defects are low and summably bounded, and that cumulative capability improvement follows whenever gains dominate drift losses.

The final picture is this:

Kemendo says: systems persist when components align under energy constraints. Weak-TransWeave Cohesion says: systems persist, compose, generalize, and self-improve when a self/other cut remains transweavably stable under resource-bounded, weakness-maximizing pattern-flow dynamics.

And the intelligence-oriented slogan is:

General intelligence is cohesive transfer of competence across task-space under resource, boundary, and self-continuity constraints.

This framing suggests a concrete research program. For any proposed biological, social, software, or AGI system: define its pattern-flow graph, self/other cut, viability quantale, update operators, and transweave maps; measure the local defects; and check whether they compose within the available genenergy budget. For self-modifying systems, require every accepted modification to carry a weak-transweave certificate. Carried out systematically, this program turns cohesion from a metaphor of “holding together” into an auditable mathematics of persistence, transfer, intelligence, and robust self-transformation.

References

- [1] Andrew Kemendo. *A General Theory of Cohesion*. August 9, 2025.
- [2] Ben Goertzel. *Hyperseed Ontology: Formal Presentation and Theoretical Development*. March 3, 2026.
- [3] Ben Goertzel. *Hyperseed: A Scrutability-Style Core Ontology and a Practical Inference-Control Scaffold for Probabilistic Logic Networks*. March 3, 2026.
- [4] Ben Goertzel. *Weakness is All You Need: Notes on Quantale Weakness as a Unifying Generalized-Occam Principle for Cognitive Science and AGI*. Rough version, 2026.
- [5] Ben Goertzel. *TransWeave: Transfer Learning and Cognitive Synergy via Geodesic Cognition Guidance and Compositional Solution Transfer using Bellman-Darboux Operator Intertwining*. 2025.
- [6] Ben Goertzel. *Practical Intelligence Lives on the Pareto Frontier Between Independence Enhancement and Dependency Modeling*. November 3, 2025.
- [7] Ben Goertzel. *Geodesic Inference Control*. Rough initial draft, September 3, 2025.