

An Axiomatic Approach to General Intelligence: $\text{SANC}(E_3)$

Self-organizing Active Network of Concepts with Energy E_3

Daesuk Kwon
Won-gi Paeng
Hyntel, Inc.
Seoul, Republic of Korea

HYNTEL@HYNTEL.NET
 WGPANG@HYNTEL.NET

Editor: Alexey Potapov

Abstract

General intelligence must reorganize experience into internal structures that enable prediction and action under finite resources. Existing systems implicitly presuppose fixed primitive units—tokens, subwords, pixels, or predefined sensor channels—thereby bypassing the question of how representational units themselves emerge and stabilize.

This paper proposes Self-organizing Active Network of Concepts with Energy E_3 – $\text{SANC}(E_3)$, an axiomatic framework in which representational units are not given *a priori* but instead arise as stable outcomes of competitive selection, reconstruction, and compression under finite activation capacity, governed by the explicit minimization of an energy functional $E_3 = \lambda_1 L_{\text{rec}} + \lambda_2 C_{\text{struct}} + \lambda_3 C_{\text{update}}$.

$\text{SANC}(E_3)$ draws a principled distinction between **system tokens**—structural anchors such as {here, now, I} and sensory sources—and tokens that emerge through self-organization during co-occurring events. Five core axioms formalize finite capacity, association from co-occurrence, similarity-based competition, confidence-based stabilization, and the reconstruction–compression–update trade-off. Residual capacity $C_{\text{remain}}(t)$ is introduced as an explicit control variable that dynamically modulates representation creation, stabilization, and deletion thresholds.

A key feature of the framework is a **pseudo-memory-mapped I/O mechanism**, through which internally replayed Gestalts are processed via the same axiomatic pathway as external sensory input. At the theoretical level, perception, imagination, prediction, planning, and action are unified within a single representational and energetic process. This unification is intended as a structural account: it does not claim that a complete implemented AGI system is provided here, but specifies how these processes can be described within one axiomatic pathway. The same account naturally extends the framework to embodied and physical agents, where interaction with the environment and motor behavior are treated as continuations of Gestalt completion rather than as separate control modules.

From the axioms, twelve propositions are derived, suggesting how category formation, automatic threshold tuning, hierarchical organization, unsupervised learning, and high-level cognitive activities—dialogue, authoring, perception, causal inference, and action—can be understood as instances of **Gestalt completion under E_3 minimization**.

Keywords: axiomatization of intelligence; competitive selection; system tokens; reconstruction–compression trade-off; category formation; self-similar hierarchy; Gestalt completion

1. Introduction

For millennia, the question “What is intelligence?” has been posed in myriad forms. Contemporary artificial intelligence has inherited this question yet still lacks a rigorous, operational definition of

general intelligence. Despite impressive recent advances—most notably attention-based Transformers (Vaswani et al., 2017)—humanity has arguably not achieved AGI. We contend that this shortfall stems from attempts to implement isolated facets of intelligence on the basis of accumulated empirical success and local mechanistic insight, without first defining what intelligence is, or worse, by adopting an inadequate definition. In the absence of a definition, one cannot deliberately engineer general intelligence; at best, one may stumble toward something that resembles it.

Rather than attempting a direct definition of ‘general intelligence,’ we take human intelligence as the paradigmatic case: whatever humans do when we say they are ‘being intelligent’—that is what we mean by general intelligence. Our approach is to ask: what structural and dynamical properties must a system possess to exhibit such intelligence? The axioms we propose are intended as candidate necessary structural conditions motivated by features that human intelligence appears to satisfy. Any system satisfying all these axioms is termed *SANC(E_3)-intelligent* within the terminology of this paper. To our knowledge, human intelligence remains the only known system that appears to fully satisfy these axioms; whether artificial systems can do so is precisely the question this framework aims to clarify.

This paper proposes **SANC(E_3)** (*Self-organizing Active Network of Concepts with Energy E_3*), a minimal axiomatic framework for characterizing the emergence of intelligence.

SANC(E_3) distinguishes *system tokens* (given by the architecture) from *tokens* that emerge through self-organization. It defines an intelligent system as one that, under finite resources and an explicit energy-minimization objective, autonomously generates and organizes representational units whose experiential completion yields prediction and can trigger action. Such representational units are termed Gestalts.

In SANC(E_3), association—not binding—is taken as the primitive of intelligence: intelligence is defined as the capacity to map co-occurring events to associated internal representations under finite resources, while compositional operators such as G are treated solely as formal realizations of this associative relation.

We do not claim a complete axiomatization in the strict sense of mathematical logic; rather, we present a principled axiomatic framework specifying candidate structural conditions for general intelligence—understood here as the capacity to autonomously form, organize, and deploy internal representations that enable prediction and action across open-ended domains under finite resources.

Scope and limits. The present paper is primarily a theoretical and axiomatic proposal. The axioms are not claimed to constitute a complete mathematical sufficiency theorem for all possible intelligent systems; sufficiency in the empirical sense also depends on non-degenerate parameter regimes, scalable implementation, and adequate interaction with an environment. Claims about reasoning, dialogue, authoring, social interaction, and embodied action should therefore be read as structural claims within the SANC(E_3) formalism, not as claims that a complete implemented AGI system has already demonstrated all of these capacities.

A crucial feature of the framework is that characteristics such as morpheme-level tokens and unsupervised learning are not externally imposed design goals but natural consequences of the axiomatic system. Likewise, causal inference and social interaction are explained naturally by the same Gestalt-completion principle, without dedicated specialized mechanisms.

1.1 Related Work

Formal definitions of intelligence. Efforts to define intelligence mathematically or axiomatically have been central to AGI research. Hutter’s AIXI model (Hutter, 2005) defines an optimal agent by combining Solomonoff induction with sequential decision theory, maximizing expected reward while favoring simpler environment models. AIXI is, however, incomputable and presupposes infinite resources; moreover, it does not address how representational units (tokens) form. Legg and Hutter (Legg and Hutter, 2007) extend this idea to “universal intelligence,” defined as a weighted sum of goal achievement across environments, but the same limitations persist.

Wang’s NARS (Wang, 2019) defines intelligence as adaptation under insufficient knowledge and resources (AIKR). SANC’s A1 (finite capacity) aligns with AIKR’s resource constraints, yet SANC goes further: it explains how representational units themselves emerge (T3) and how hierarchical structure forms (T7, T8) under an explicit energy decomposition.

Chollet (Chollet, 2019) characterizes intelligence as skill-acquisition efficiency and proposes an algorithmic-information-theoretic formalization; the ARC benchmark targets measurement. SANC is not primarily a measurement framework but a generative one: it specifies, via axioms, what an intelligent system must satisfy and how its core properties arise.

Friston’s Free Energy Principle (Friston, 2010) posits that agents minimize variational free energy, emphasizing prediction-error minimization. SANC likewise adopts energy-based dynamics but decomposes E_3 explicitly into three terms (L_{rec} , C_{struct} , C_{update}) to formalize the trade-off among reconstruction accuracy, structural simplicity, and update stability.

Bostick (Bostick, 2024) proposes an axiomatic definition of AGI as “Lawful Autonomous Continuity,” emphasizing identity continuity and autonomy. As a purely definitional approach, it does not specify mechanisms for representation formation or learning; SANC renders these mechanisms explicit through axioms and derived propositions.

Distinctive features of SANC(E_3). Compared to prior work, SANC(E_3) is distinguished by: (i) tokens are emergent outcomes of E_3 dynamics, not primitives (T3); (ii) system tokens Σ_{sys} formalize Kantian *a priori* anchors to resolve bootstrapping; (iii) the trade-off between token explosion and sequence explosion is made explicit (T4); (iv) unsupervised learning is a natural consequence of the axiomatic system (T10); and (v) dialogue, authoring, causal inference, perception, and self-improvement (extended definition D15) are given a common structural interpretation as Gestalt completion.

Connections to recent AI research. Forward–Forward learning (Hinton, 2022) demonstrates local learning without backpropagation; SANC’s A2 (similarity-based competition) and A3 (confidence-based stabilization) provide an axiomatic account of local competitive dynamics. JEPA and I-JEPA (LeCun, 2022; Assran et al., 2023) propose predictive representation learning in embedding space; SANC’s T6 generalizes this insight: *Gestalt completion (replay) is prediction*. Nested Learning (Behrouz and others, 2025) emphasizes multiple time scales; SANC captures this via A3 and self-similar levels (T8). Large Concept Models (Meta AI, 2024) aim to operate at concept level; SANC is more fundamental, treating both vocabulary-level tokens and concept-level tokens as outcomes generated identically by the axioms (T3). Finally, fixed tokenization schemes such as BPE (Sennrich, Haddow, and Birch, 2016) assume a static vocabulary; SANC requires the token inventory itself to evolve with accumulated experience (T3, T4).

What we contribute. (i) Core axioms A1–A5. (ii) Mechanism axioms A6–A10. (iii) Derived propositions T1–T12. (iv) Extended definitions D10–D15. (v) Systematic comparison with existing

theories. (vi) Appendix C provides the original postulates, first formulated in the mid-1990s, in informal language for intuitive accessibility.

2. Notation

Terminological convention (levels). The core theoretical entities in SANC are the *associated event* and its corresponding *associated representation*; an associated event is said to form a *Gestalt* when it constitutes a coherent, unified whole that is distinguished and mapped to an internal associated representation. The term *token* is used only as an implementation convenience (e.g., Transformer symbols used to instantiate representations). Wherever possible, we prefer the theoretical terms *associated event / associated representation / Gestalt*.

We introduce a formal vocabulary. The notation is implementation-agnostic: it constrains what must exist without prescribing whether the realization is neural, symbolic, or hybrid.

- \mathcal{E} : set of events (external stimuli or internal activations).
- \mathcal{E}^* : set of finite event sequences (Kleene star); $e_{1:t} \in \mathcal{E}^*$ denotes the sequence from time 1 to t .
- \mathcal{U} : universe of representational units (discrete symbols, vectors, subspaces, distributions, etc.).
- $A_t \subseteq \mathcal{U}$: active state at time t .
- κ, C : capacity measure and finite upper bound; $\kappa(A_t) \leq C < \infty$.
- C_{tot} : total available activation capacity.
- $C_{\text{remain}}(t)$: residual capacity, defined as $C_{\text{tot}} - \kappa(A_t)$.
- L_{rec} : reconstruction loss.
- C_{struct} : structural complexity cost.
- C_{update} : update cost.
- $\mathcal{P}(X)$: power set of X .
- $\text{co}_t(x, y)$: predicate indicating that x and y are co-detected or co-active at time t .
- $\text{co}^*(x, y)$: accumulated co-occurrence statistic over time; the specific increment/decrement scheme is implementation-dependent, but to maintain finite capacity, items not reactivated for an extended period may be weakened or removed.
- Π_t : selection operator (the concrete rule may vary but must satisfy A2).
- $S_t \subseteq \mathcal{U}$: candidate set at time t (active units, association candidates, retrieval candidates, etc.).
- $\Pi_t(S_t) \subseteq S_t$: units actually selected (or reinforced) from S_t .
- $\text{CD}_t(u)$: confidence degree of unit u at time t .
- G : association (compositional) operator, $G : \mathcal{U} \times \mathcal{U} \rightarrow \mathcal{U}$.

- sim_t : similarity predicate on elements of \mathcal{U} .
- $[u]_{\text{sim}}$: similarity neighborhood of u (not assumed to be an equivalence class): $[u]_{\text{sim}} := \{v \in S_t : \text{sim}_t(u, v)\}$.
- E_3 : energy function, weighted sum of L_{rec} , C_{struct} , C_{update} (Axiom A4).
- θ_{stab} : stabilization threshold (token-formation threshold).
- θ_{err} : reconstruction-error threshold (reconstruction succeeds iff $L_{\text{rec}} \leq \theta_{\text{err}}$).
- $P_t(\cdot)$: predictive distribution over the next event at time t (e.g., a normalized score distribution).
- H_t : predictive uncertainty (e.g., entropy $H_t = -\sum P_t(e) \log P_t(e)$, sum over $e \in \mathcal{E}$).
- θ_{unc} : uncertainty threshold (context segmentation and “completed Gestalt candidate” detection).
- B_t : the current associated event (Gestalt) formed up to time t (A8).
- $\text{terminate}(B_t)$: operator that finalizes B_t as a “completed Gestalt candidate” and ends the current association (A8).
- $\text{reset}_{\text{assoc}}(t+1)$: operator that resets the association state so a new Gestalt association begins at time $t+1$ (A8).
- S_i : the i -th SANC(E_3)-intelligent subsystem within a larger system (A10).
- $\text{Comm}(i, j)$: a communication relation/channel enabling Gestalt exchange between subsystems S_i and S_j (A10).
- $\text{send}_{i \rightarrow j}(\mathbf{g})$: transfer operator that sends a Gestalt (or a Gestalt sequence/set) \mathbf{g} from S_i to S_j (A10).
- $\text{recv}_{j \leftarrow i}(\mathbf{g})$: reception operator by which S_j receives \mathbf{g} from S_i ; received Gestalts are elements of the shared universe \mathcal{U} and are included in the candidate set S_t , participating in A5–A8 (A10).
- Σ_t : set of stabilized tokens at time t , $\Sigma_t := \{u \in \mathcal{U} : \text{CD}_t(u) \geq \theta_{\text{stab}}\}$.
- τ : a token at time t , $\tau \in \Sigma_t$.
- Σ_{sys} : system tokens: $\Sigma_{\text{anchor}} \cup \Sigma_{\text{sense}} \cup \Sigma_{\text{apriori}}$ (see D1).
- Σ_{anchor} : anchor tokens (e.g., {here, now, I}; see D1).
- Σ_{sense} : sensory tokens.
- Σ_{apriori} : (*a priori*) intrinsic tokens.

Auxiliary notation for selection bias:

- $\text{score}_t(u)$: scalar selection score for candidate u (implementation-defined).

- $\Delta_t^{\text{will}}(u)$: will/intent bias term added to the score.
- $\text{score}_t^{\text{will}}(u) := \text{score}_t(u) + \Delta_t^{\text{will}}(u)$.
- $\Pi_t^{\text{will}}(S_t)$: will-biased selection (see D11).
- δ_I : I-proximity distance threshold (see D12).

Remark (distributed vs. discrete). Elements of \mathcal{U} need not be discrete symbols. A “token” is a stabilized unit in Σ —an outcome of stabilization, not an *a priori* primitive. (**Only system tokens** are tokens in Σ_t that are given *a priori*.)

Remark (homogeneity across levels). In this paper, ‘homogeneity across levels’ means that the axiomatic rules governing event–representation correspondence, association/generation, selection/retention (and deletion of non-*a priori* representations), and finite capacity constraints operate identically regardless of level.

3. Core Definitions

Definition (D0: SANC(E_3)-intelligence). A system is SANC(E_3)-intelligent if there exist a representational universe \mathcal{U} , an activation dynamic A_t over \mathcal{U} , a capacity measure κ with finite upper bound C , a candidate set S_t , a selection operator Π_t , a confidence degree CD_t , system tokens Σ_{sys} , and an energy $E_3 = \lambda_1 L_{\text{rec}} + \lambda_2 C_{\text{struct}} + \lambda_3 C_{\text{update}}$ such that Axioms A1–A10 hold and the system’s dynamic update mechanism operates to minimize E_3 (at least approximately) under the constraints of A1–A3—specifically, the update dynamics exhibit a descent property that decreases E_3 over time (in expectation).

Definition (D1: representational unit, identification, segmentation, system tokens). A representational unit is any $u \in \mathcal{U}$. An identification operator $\text{id} : \mathcal{E} \rightarrow \mathcal{U}$ assigns a representational unit to an event. A segmentation operator $\text{seg} : \mathcal{U} \rightarrow \mathcal{P}(\mathcal{U})$ decomposes a unit into parts. System tokens $\Sigma_{\text{sys}} \subseteq \mathcal{U}$ are representational units given by the system’s structure, rather than emerging from self-organization through event association.

- **Anchor tokens** Σ_{anchor} : structurally given tokens such as {here, now, I} that co-occur with, and are embedded with, all representations.
- **Sensory-source tokens**: outputs of sensory receptors—level-0 materials.

Anchor tokens participate in the same network as other tokens, serving as relational reference points, but their existence is presupposed. Their interpretation (what “here” or “now” refers to) may evolve through experience.

Relation to Kant. System tokens can be viewed as formalizing Kantian *a priori* forms that structure experience. Unlike Kant’s forms (Kant, 1781), however, system tokens are themselves tokens: they participate as nodes in the relational network.

Definition (D2: distinction and consistency). The system can distinguish identical, similar, and different events, and can assign consistent internal representations to them. Specifically: identical

events map to the same representation; similar events decompose into shared and differing parts (see D3).

Definition (D3: similarity). $\text{sim}_t(u, v) \iff \exists a, b, c : u = G(a, b), v = G(a, c), b \neq c$. Here a is the shared subcomponent; b and c are the differing parts. Thus similarity requires (i) at least one shared component, (ii) at least one differing component, and (iii) recomposability into u and v .

Example. Suppose $ab = G(a, b)$ corresponds to “with Mom” and $ac = G(a, c)$ to “with Dad.” Then $\text{sim}_t(ab, ac)$ holds: $a =$ “with” is the shared part; $b =$ “Mom” and $c =$ “Dad” are the differing parts.

Admissible instantiations. In a symbolic implementation, similarity may be defined by shared prefix, suffix, substring, dependency context, or feature subset. In a vector-space implementation, it may be defined by cosine similarity, projection onto a shared subspace, or another task-dependent metric. The axioms require that similarity support competition (A2) and splitting (A7); they do not require a unique metric.

Definition (D4: association and associated representation). Association is the fundamental relation by which co-occurring events are related. Formally, association may be realized by a compositional operator $G : \mathcal{U} \times \mathcal{U} \rightarrow \mathcal{U}$, which constructs an associated representation from existing representations. The statement “ x is associated with y ” can be understood as: $\text{co}_t(x, y)$ or $\text{co}^*(x, y)$ causes $G(x, y)$ to become a candidate for selection. An association (and its realized associated representation) can be strengthened or weakened by selection Π_t and by the frequency of observation (experience), and, unlike *a priori* representations, is not permanent. $\text{co}^*(x, y)$ denotes the temporal accumulation of co-occurrence, increasing or decreasing with co-occurrence frequency; the specific increment/decrement scheme is implementation-dependent.

Definition (D5: Gestalt). A Gestalt is an associated structure that (i) relates co-occurring components, (ii) can become a selectable unit under Π_t , (iii) is internally replayable, and (iv) can enter Σ_t upon stabilization. It is a representational unit in formation or already formed.

This formalizes the Gestalt-psychology dictum (Wertheimer (Wertheimer, 1923)) that “the whole is more than the sum of its parts.” Once formed, a Gestalt can be treated as an event at a higher level and can be associated to others again (see A6).

Definition (D6: stabilized set and tokens). $\Sigma_t := \{u \in \mathcal{U} : \text{CD}_t(u) \geq \theta_{\text{stab}}\}$. A token at time t is any $\tau \in \Sigma_t$. (System tokens are presupposed separately.)

Remark. Tokens other than system tokens are understood to form when Gestalt candidates activated by experience undergo association, selection, and reinforcement and stabilize into Σ_t . Moreover, thresholds such as θ_{stab} are not assumed to be fixed constants but **are adaptively set to avoid increasing E_3** so as not to worsen the internal objective E_3 (see Proposition T12).

Definition (D7: category). A *category* is a token that captures shared parts or context across multiple representational units and thereby lowers E_3 .

Example. Having repeatedly experienced $ab =$ “with Mom” and $ac =$ “with Dad,” followed by $db =$ “for Mom” and $dc =$ “for Dad,” two new categories {with, for} and {Mom, Dad} may emerge—corresponding, for instance, to preposition and noun categories.

Definition (D8: retrieval, replay/regeneration, and processing equivalence).

- (a) **Partial-match retrieval.** Partial-match retrieval is an operator $\text{retr}_t : \mathcal{E}^* \rightarrow \mathcal{P}(S_t)$ that, given an observed event history $e_{1:t}$, returns a set of candidate Gestalts in S_t sharing matching parts with the current evidence. That is, $\text{retr}_t(e_{1:t}) \subseteq S_t$, and this set may contain stabilized tokens (Σ_t) as well as pre-stabilized Gestalt candidates (not yet in Σ_t). (Analogous to spreading activation in semantic networks (Quillian (Quillian, 1968)) and to Gestalt completion.)
- (b) **Replay/regeneration.** $\text{replay}(u)$ denotes an internal operation that reactivates or regenerates a Gestalt u so as to reduce L_{rec} under A4, typically via retr . Replay is a mechanism for lowering L_{rec} and, more broadly, part of the system’s E_3 -minimizing dynamics.

A minimal procedural form is: given current evidence $e_{1:t}$ and token set Σ_t , retrieval returns a candidate set $R \subseteq S_t$; candidates are evaluated by their partial reconstruction of $e_{1:t}$; A2 selects compatible winners; replay regenerates missing components of the selected Gestalts; the regenerated components contribute to the predictive distribution $P_t(\cdot)$ used by A8. This procedure specifies the role of replay without prescribing a unique architecture.

- (c) **Pseudo-memory-mapped I/O.** Internal replay of a Gestalt is subject to the same axiomatic decision rules as external sensory input—namely similarity-based competition (A2), selection and stabilization (A3), evaluation of reconstruction loss L_{rec} , and global minimization of E_3 (A4).

As a result, internally replayed Gestalts drive downstream systems—including reasoning, inference, dialogue, and action generation—in a manner indistinguishable from externally perceived events. The equivalence concerns not the origin of the input but the identity of the axiomatic processing path and its causal effects.

This organization is functionally analogous to memory-mapped I/O, but is not based on physical address mapping; rather, it relies on isomorphism of representational processing rules. We therefore refer to this mechanism as pseudo-memory-mapped I/O.

Definition (D9: embedding / relational distance vector and learning).

- (a) For a representation $u \in \mathcal{U}$, define an **embedding** as:

$$\text{emb}(u) = (\{d(u, v) : v \in \Sigma_t, \text{co}^*(u, v) > 0\}, \{d(u, \tau) : \tau \in \Sigma_{\text{anchor}}\}),$$

where Σ_t is the stabilized token set at time t , $\Sigma_{\text{anchor}} = \{\text{here, now, } I\} \subseteq \Sigma_{\text{sys}}$ are *a priori* coordinate-like tokens, and $d(u, v)$ is a relational (latent) distance. In this paper, the ‘**meaning**’ of a token u is the **relational profile** given by its relationships (similarity/distance/association) with other tokens at the same time.

Because anchor tokens are assumed to co-occur with all representations, every embedding contains distances to Σ_{anchor} . Tokens such as *past, future, or world* are not *a priori* tokens but can be interpreted as tokens emerging from relational structure (such as distance to Σ_{anchor}).

Remark 1: $\text{emb}(u)$ denotes an (**unordered**) **profile** of the above distances; in implementation, this may be arranged into a fixed-order/dimension vector or summarized as statistics.

Remark 2 (Saussure and relational meaning). Defining meaning as a relational profile is consistent with Saussure’s insight (de Saussure, 1916) that the meaning of a sign is not a

fixed entity but is determined by its relations (context) with other signs (arbitrariness of the sign / system of differences). Definition D9 formalizes this as distance-based embedding.

Remark 3 (admissible distance metrics). The distance $d(u, v)$ may be Euclidean distance, cosine distance, graph distance, information distance, or another implementation-dependent metric, provided it is used consistently and becomes smaller, in expectation, for representations whose accumulated association $\text{co}^*(u, v)$ is stronger.

- (b) **Learning.** In this paper, ‘learning’ refers to the process by which Gestalts (tokens) in Σ_t are created, deleted, or stabilized as experience accumulates, and their relationships (similarity/association/confidence degree/inter-token relational profiles) are updated so as to better satisfy the internal objective E_3 .

4. Core Axioms

The first five axioms establish foundational constraints; the remaining five (Section 5) specify compositional and structural mechanisms. Together, they force the emergence of tokens, categories, and hierarchy without hand-designed vocabulary or rules.

Axiom (A1: finite active capacity). $\kappa(A_t) \leq C < \infty$. The active state has finite capacity at every time t . Finiteness necessitates competition.

Necessity and limit. Without a finite capacity bound, the selection operator could admit all candidates and the compression pressure in C_{struct} would become vacuous. A1 distinguishes SANC(E_3) from idealized models that assume unbounded computational or representational resources. A1 does not prescribe the numerical value of C ; an under-capacity regime may fail empirically, and the viable range of C is implementation- and environment-dependent.

Axiom (A2: similarity-based competition and single selection). $|\Pi_t(S_t) \cap [u]_{\text{sim}}| \leq 1$, and $\kappa(\Pi_t(S_t)) \leq C$. Similar candidates compete, and at most one is selected from each similarity neighborhood.

Necessity and limit. Without A2, similar candidates could stabilize redundantly, causing structural growth without representational gain. A2 imposes only an upper bound within a similarity neighborhood; it does not require that neighborhoods form a partition, nor does it prescribe whether selection is implemented by argmax, softmax, sampling, or another rule.

Axiom (A3: stabilization and deletion by identification and selection).

- (a) **Stabilization by confidence degree.**

- If $u \in \Pi_t(S_t)$: $\text{CD}_{t+1}(u) > \text{CD}_t(u)$ [selection reinforces]
- If $u \notin \Pi_t(S_t)$ repeatedly: $\text{CD}_t(u) \rightarrow 0$ [non-selection extinguishes]

The specific increment/decrement scheme is implementation-dependent, subject to A4.

- (b) **Forgetting and deletion.** Non-*a priori* representations—including those derived from experienced events or from *a priori* representations—may have their co^* values decrease if the corresponding event (or equivalent reactivation evidence) is not observed for a certain period, eventually leading to deletion from Σ_t . (The decay scheme is left to implementation.)

- (c) **A priori representation exception.** *A priori* representations are exempt from the forgetting/deletion rules above.

Necessity and limit. A3 separates transient activation from stabilized representation. Without reinforcement and decay, tokens could neither become stable nor be removed under finite capacity. A3 specifies monotonic tendencies—selection reinforces, persistent non-selection weakens—but leaves the concrete update law, decay schedule, and reinforcement magnitude to implementation.

Axiom (A4: energy E_3 and reconstruction–update dynamics).

- (a) **Definition.** $E_3 = \lambda_1 L_{\text{rec}} + \lambda_2 C_{\text{struct}} + \lambda_3 C_{\text{update}}$, where:

- L_{rec} (reconstruction loss): error in reconstructing events using the current token set Σ_t .
- C_{struct} (structural cost): complexity of the token inventory (e.g., number of tokens, total description length).
- C_{update} (update cost): instability incurred by changing tokens (e.g., invalidating existing links).

A minimal concrete form of $C_{\text{update}}(\Sigma_t \rightarrow \Sigma_{t+1})$ is any function that is monotonically non-decreasing in: (i) token-set change, such as $|\Sigma_t \Delta \Sigma_{t+1}|$; (ii) association invalidation, such as the number or weight of broken $\text{co}^*(u, v) > 0$ relations when u or v is removed or replaced; and (iii) re-stabilization burden, such as the number of newly created units whose confidence degree must reach θ_{stab} . For example, one may instantiate

$$C_{\text{update}} = \alpha |\Sigma_t \Delta \Sigma_{t+1}| + \beta B_t + \gamma N_t,$$

where B_t counts or weights invalidated associations and N_t counts newly created non-stabilized units. This example is not unique; it specifies the structural role of update cost.

- (b) **Intelligence condition.** A $\text{SANC}(E_3)$ -intelligent system acts to minimize E_3 .
- (c) **Trade-off.** The three terms constituting E_3 trade off against each other as follows, and E_3 minimization finds a balance:

- Minimizing only $L_{\text{rec}} \rightarrow$ token explosion ($C_{\text{struct}} \uparrow$)
- Minimizing only $C_{\text{struct}} \rightarrow$ reconstruction failure ($L_{\text{rec}} \uparrow$)
- Minimizing only $C_{\text{update}} \rightarrow$ inability to adapt ($L_{\text{rec}} \uparrow$ on new experience)

- (d) **Reconstruction–update dynamics.** For each event e :

- The system attempts reconstruction using current Σ_t .
- If $L_{\text{rec}}(e, \Sigma_t) \leq \theta_{\text{err}}$ (success): the CD of the Gestalts used is reinforced.
- If $L_{\text{rec}}(e, \Sigma_t) > \theta_{\text{err}}$ (failure): association (A5), splitting (A7), and competition (A2) are triggered, updating $\Sigma_t \rightarrow \Sigma_{t+1}$.

Admissible optimization families. The minimization of E_3 need not be exact or global. It may be realized by local competitive dynamics, gradient-based updates in a differentiable instantiation, sampling-based search, or hybrids of these. What is required by the axioms is a descent tendency under finite capacity, not a unique optimization algorithm.

Relation to backpropagation. The error-driven update in (d) is compatible with error backpropagation (EBP) as an implementation mechanism. However, an EBP-only system that optimizes only reconstruction loss L_{rec} ignores C_{struct} and C_{update} , and therefore need not minimize E_3 as a whole. SANC does not preclude EBP; it requires that learning operate under A1–A3 and the full E_3 trade-off.

- (e) **Capacity–Energy Coupling.** In a SANC(E_3)-intelligent system, the dynamics of E_3 minimization are modulated by the residual capacity $C_{\text{remain}}(t)$.

The thresholds associated with representation creation, stabilization, and deletion (θ_{create} , θ_{stab} , θ_{delete}) are not fixed constants but are dynamically adjusted as functions of $C_{\text{remain}}(t)$.

As $C_{\text{remain}}(t)$ decreases, the system suppresses the creation of new Gestalts and favors reuse and compression of existing ones, preventing uncontrolled growth of C_{struct} while minimizing degradation of L_{rec} .

Necessity and limit. A4 is required because minimizing any single pressure alone yields a degenerate regime: perfect reconstruction can overfit by proliferating units, pure compression can destroy reconstruction, and pure stability can prevent adaptation. The weights λ_i and the concrete forms of the three terms are not fixed by the axiom; they define a family of SANC(E_3) instantiations.

Axiom (A5: co-occurrence generates association candidates). $\text{co}_t(x, y) \vee \text{co}^*(x, y) \Rightarrow \exists t' \geq t : G(x, y) \in S_{t'}$. Co-detected events become candidates for associated representation (Gestalt). This axiom formalizes Hebb’s principle (Hebb, 1949) that “neurons that fire together wire together.” (Association candidates thus generated may be weakened or deleted by A3 if not sufficiently reinforced/reactivated.)

Necessity and limit. Without A5, associated representations would need to be supplied externally, contradicting the claim that tokens are outcomes rather than primitives. A5 does not state that every co-occurrence stabilizes; most candidates may disappear through A2–A4 dynamics.

5. Mechanism Axioms

Axiom (A6: closure under association). For all $x, y \in \mathcal{U}$, $G(x, y) \in \mathcal{U}$. The universe \mathcal{U} is closed under association. Consequently, recursive composition and hierarchy become possible.

Axiom (A7: similarity-based splitting). If $\text{sim}_t(G(a, b), G(a, c))$ holds, a splitting occurs that identifies the shared part a and the differing parts b and c .

Axiom (A8: uncertainty-gated segmentation and Gestalt completion). If the uncertainty H_t of the predictive distribution $P(e_{t+1} \mid e_{1:t})$ for the next event exceeds threshold θ_{unc} , the current associated event (Gestalt) B_t is deemed complete and association restarts at $t + 1$.

$$H_t \geq \theta_{\text{unc}} \Rightarrow \text{terminate}(B_t) \wedge \text{reset}_{\text{assoc}}(t + 1).$$

Remark. The system may compute a predictive distribution $P_t(\cdot)$ for the next event given the observed history $e_{1:t}$, and its uncertainty H_t (e.g., entropy). As H_t decreases (predictions become more certain), the current context tends to cohere into a single Gestalt. Conversely, when a context (Gestalt) is completed and a new context begins, H_t may spike; if $H_t \geq \theta_{\text{unc}}$, the system treats the current history $e_{1:t}$ as a ‘‘completed Gestalt candidate,’’ includes it among association/stabilization candidates, and segments subsequent events as a ‘new Gestalt.’

Necessity and limit. Without an uncertainty-gated boundary, an event stream may fail to form finite Gestalt candidates, preventing reusable higher-level structure. The uncertainty measure may be entropy, predictive variance, surprise, or another implementation-dependent signal, and θ_{unc} may itself be tuned by the system pressure described in T12.

Axiom (A9: cross-level input–output isomorphism, homogeneity of levels).

- (a) **Input–output isomorphism / format-preserving compatibility.** Each level ℓ operates over the same representational space \mathcal{U} (i.e., each level can be understood as an operation mapping \mathcal{U} to \mathcal{U}), and all levels share the same representational schema (association G , similarity sim , confidence degree CD , segmentation/completion predicates). Moreover, because associated representations also belong to \mathcal{U} (A6), lower-level outputs can become higher-level inputs without conversion, and higher-level Gestalts can be fed back to lower levels without conversion.
- (b) **Homogeneity of levels.** The axiomatic rules governing event–representation correspondence, association/generation, selection/retention (and deletion of non-*a priori* representations), and finite capacity constraints operate identically regardless of level. Thus the same axiomatic dynamics can be repeated at each level, providing the formal basis for fractal-like self-similar organization.

Axiom (A10: interconnection of subsystems via Gestalt exchange). A $\text{SANC}(E_3)$ -intelligent system may be composed of multiple $\text{SANC}(E_3)$ -intelligent subsystems $\{S_i\}$. For any pair (i, j) , there exists a communication relation $\text{Comm}(i, j)$ and transfer/reception operators $\text{send}_{i \rightarrow j}$, $\text{recv}_{j \leftarrow i}$ that exchange Gestalts (associated events/representations) as units/sets/sequences. Received Gestalts are admitted into the recipient’s candidate set S_i and can participate in association (A5–A6), splitting (A7), stabilization and deletion (A3), and segmentation (A8), thereby mutually shaping each subsystem’s dynamics. Because a subsystem can itself be recursively organized as a network of subsystems, this axiom supports fractal-like self-similar organization and formalizes Hawkins’ ‘‘thousand brains’’ hypothesis (Hawkins, 2021), extending naturally to multi-agent interaction (Minsky (Minsky, 1986)).

Necessity and limit. A10 makes multi-agent and multi-modal Gestalt exchange internal to the same formalism rather than an external add-on. It does not prescribe channel bandwidth, fidelity, latency, encoding, or synchronization; these are implementation constraints.

6. Derived Propositions

Each proposition states a consequence generally expected in systems satisfying the axioms. Except T1, T2, T10, the ‘proofs’ for each item are sketches that reveal dependencies on axioms and definitions rather than formal proofs.

Proposition (T1: no privileged invariant tokens, except system tokens). Any token other than Σ_{sys} is not exempt from competition, revision, or extinction.

Dependencies: A1, A2, A3.

Proof. Let $u \in \mathcal{U} \setminus \Sigma_{\text{sys}}$ be arbitrary. By Axiom A2, u is subject to similarity-based competition: whenever there exists $v \in [u]_{\text{sim}} \cap S_t$ with $v \neq u$, at most one of $\{u, v\}$ can be selected. By Axiom A3(a), non-selection causes CD to decrease. By A3(b)–(c), only *a priori* representations are exempt from forgetting and deletion. Since $u \notin \Sigma_{\text{sys}}$, u receives no exemption: it is subject to competition under A2 and to potential deletion under A3(b). Therefore, u is not exempt from competition, revision, or extinction. ■

Proposition (T2: inevitable turnover under finite capacity). For any SANC(E_3)-intelligent system continuously engaged with an environment, turnover (forgetting) of some Gestalt candidates is unavoidable.

Dependencies: A1, A2, A3, A5.

Proof. A system continuously engaged with an environment receives events over unbounded time. By A5, co-occurring events generate association candidates, so infinitely many distinct candidates enter S_t over time. By A1, the active capacity is finite at every t . By A2, only a limited number can be selected per time step, so some candidates are not selected infinitely often. For any such candidate u , A3(a) yields $\text{CD}_t(u) \rightarrow 0$. By D6, once $\text{CD}_t(u) < \theta_{\text{stab}}$, $u \notin \Sigma_t$, hence u eventually leaves the stabilized set. Therefore turnover is unavoidable. ■

Proposition (T3: all tokens except system tokens are generated). Tokens (excluding system tokens) are generated as stabilized structures preferred under E_3 minimization. Only system tokens are structural premises, not products of competition.

Dependencies: A3, A4, D1, D6.

Proof sketch. A4 drives E_3 minimization. Reducing L_{rec} requires representational structures that successfully reconstruct events. A3 reinforces structures that succeed, raising their CD above θ_{stab} and thereby conferring token status (D6). Tokens are thus interpretable as emergent outcomes of E_3 dynamics. ■

Proposition (T4: token–sequence balance). Minimizing only L_{rec} leads to token explosion; minimizing only C_{struct} leads to sequence explosion; E_3 induces a balance between the two explosions.

Dependencies: A4.

Proof sketch. Direct consequence of A4(c). ■

Proposition (T5: category formation as compression). Shared parts or contexts across multiple Gestalts stabilize as new representations (categories), because such categories enable reuse and lower C_{struct} while also reducing L_{rec} , thereby lowering E_3 doubly.

Dependencies: A4, A5, A7, D7.

Proof sketch. A7 identifies shared parts across similar Gestalts. Reifying the shared part as a token enables reuse in reconstructing multiple events, lowering C_{struct} . A4 therefore favors category formation. ■

Proposition (T6: prediction is completion). Given representations that only partially match a Gestalt, E_3 minimization drives completion of the missing parts of the Gestalt; this completion constitutes prediction.

Dependencies: A4, D5, D8(a), D8(b).

Proof sketch. Partial-match retrieval (D8(a)) retrieves candidate Gestalts sharing parts with current evidence. Replay (D8(b)) internally regenerates them, activating as-yet-unobserved components. This activation is prediction. The derivation formalizes Wertheimer’s (Wertheimer, 1923) Gestalt-completion principle. ■

Remark (causal inference is a consequence). “Causal inference” is not a separate faculty but a corollary of T6. When events form a temporal Gestalt (e.g., [striking match, flame]), presenting the earlier part activates completion of the unobserved later part, completing the Gestalt—this is experienced as causal inference.

Proposition (T7: hierarchy emerges from closure). If a single level cannot sufficiently minimize E_3 , closure under association (A6) permits recursive composition; E_3 preference (A4) drives formation of multi-level hierarchical structure; subsystem interconnection (A10) allows different subsystems to specialize at different hierarchical levels while sharing representations.

Dependencies: A4, A6, A10.

Proof sketch. A6 entails that associated structures re-enter \mathcal{U} , enabling recursive processing. Reconstructing complex events at a single level incurs high C_{struct} or high L_{rec} . Forming higher-level tokens reduces both, so A4 favors hierarchical multi-level structure. A6 and A10 enable horizontal specialization and vertical serialization across subsystems. ■

Proposition (T8: self-similar, homogeneous levels). Systems satisfying A1–A10 can form recursively connected multi-level structures, where each level implements dynamics satisfying A1–A10. Thus the overall organization has a self-similar (fractal-like) structure with homogeneity across levels.

Dependencies: A6, A9, A10.

Proof sketch. By A6, higher-level inputs can be composed from lower-level outputs; by A10, interconnection across levels is possible. Moreover, by A9, the axiomatic role structure repeats in the same form regardless of level, so each level can implement the same axiomatic dynamics (A1–A10). Recursive repetition yields a self-similar structure. ■

Proposition (T9: system tokens are re-interpretable). Anchor tokens such as {here, now, I} are presupposed as a kind of ‘relational anchor’ that mediates relationships among all tokens. However, their meaning is determined relationally, and as the surrounding tokens and their distances change, the meaning of anchor tokens also evolves.

Dependencies: D1, D9, T3.

Proof sketch. D1 presupposes the existence of system tokens. D9 defines meaning relationally (via distances to other tokens). By T3, surrounding tokens change and evolve. Hence the relational meaning of system tokens also changes and evolves. ■

Proposition (T10: unsupervised learning is a consequence). Internal E_3 minimization under A1–A4 causes learning to occur without explicit external supervision.

Dependencies: A1, A3, A4, D9.

Proof. By A4(b), the system minimizes $E_3 = \lambda_1 L_{\text{rec}} + \lambda_2 C_{\text{struct}} + \lambda_3 C_{\text{update}}$. By A4(d), reconstruction success strengthens the participating gestalt(s) (increasing CD), while failure triggers updates of Σ_t via segmentation/recombination mechanisms. By A3(a), reinforced representations become more stable, while non-reinforced ones decay. By D9(b), learning is precisely the creation,

stabilization, modification, or removal of representations and their relations. Since these state changes are driven solely by internal E_3 -minimization and reconstruction-error thresholds (and do not require labels/rewards), learning occurs without external supervision. ■

Proposition (T11: concepts generated from events form a network: SANC). The set of stabilized tokens Σ_t forms a **graph structure** connected by association (e.g., co^* and association operator G); in this paper, we call this a ‘concept network.’ Such a network ‘self-organizes’ in the sense that its structure changes and reorganizes without external supervision as tokens are created, deleted, and reinforced according to experience. Anchor tokens serve as globally referenceable reference nodes throughout the system, providing global coordinates (relational anchors) for the network.

Dependencies: A5, A6, D1, D4, D6.

Proof sketch. A5 forms associations between co-occurring representations. D4 extends association relationships to closure/structure via operator G . By D6, the stabilized token set Σ_t is formed, and the connections among tokens and association relationships constitute a graph (network). Anchor tokens from D1 serve as globally referenceable reference points, providing relational coordinates. ■

Proposition (T12: threshold tuning via system pressure). E_3 optimization biases the choice of thresholds such as θ_{stab} . (Too low yields token explosion ($C_{\text{struct}} \uparrow$); too high prevents token formation ($L_{\text{rec}} \uparrow$.) Thresholds are therefore systemically determined near the E_3 optimum.

Dependencies: A4, T4, D6.

Proof sketch. T4 establishes that E_3 balances token count against sequence complexity. The same trade-off applies to θ_{stab} : a suboptimal threshold shifts the balance unfavorably. ■

Corollary (dialogue, reasoning, and related processes). Under Definition D8(c) (pseudo-memory-mapped I/O) and Axiom A4(d), external sensory input, internal replay (imagination), and action-related signals are processed through the same logical update pathway. Consequently, dialogue, reasoning, and embodied interaction can be modeled, at this level of abstraction, as different operational modes of a single representational completion mechanism.

Justification—formal argument, not a full proof. By Definition D8(c), externally perceived events and internally replayed representations enter the system through an identical logical data path, differing only in origin, not in processing structure. By Axiom A4(d), each such event—regardless of origin—triggers reconstruction evaluation, coherence reinforcement, or structural update of Σ_t according to reconstruction success or failure. Hence, no formal distinction is made within the update dynamics between perception-driven input and internally generated input. Action-related signals, insofar as they are evaluated by their reconstructive consequences on subsequent sensory input, are integrated into the same update loop. Therefore, dialogue (linguistic interaction), reasoning (internal replay and completion), and physical interaction (action–perception cycles) are unified as instances of representational completion under a common update mechanism. ■

Remark (embodiment and physical AI). The corollary extends naturally to embodied agents. Motor action can be understood as the external realization of an internally completed Gestalt: the system replays a Gestalt whose completion includes motor-effector components, and this replay causes physical changes in the environment. The resulting environmental change produces new sensory input, which re-enters the same axiomatic pathway (D8(c)). Thus the perception–action loop can be described as a continuous chain of Gestalt completions under E_3 minimization.

Physical AI, robotics, and embodied cognition may therefore be expressible within the same formal vocabulary, although implementing such systems requires additional architectural and empirical work.

7. Extended Definitions: Higher Cognition as Gestalt Completion

Definition (D10: desire/drive). Desire is a bias toward internally completing and confirming a selected Gestalt.

Definition (D11: will/intent; selection bias). Will is a bias that disposes the system to select a particular Gestalt among multiple candidates—a bias applied to the selection process itself. Assume that at time t each candidate $u \in S_t$ has a scalar selection score $\text{score}_t(u) \in \mathbb{R}$, and that will introduces an additive bias $\Delta_t^{\text{will}}(u) \in \mathbb{R}$. The selection score is defined as:

$$\text{score}_t^{\text{will}}(u) := \text{score}_t(u) + \Delta_t^{\text{will}}(u).$$

Define a score-based selection operator (e.g., argmax or top- k), and define will-biased selection as:

$$\Pi_t^{\text{will}}(S_t) := \arg \max_{u \in S_t} \text{score}_t^{\text{will}}(u) \quad (\text{or top-}k \text{ for } k > 1).$$

Remark. This may be connected to Schopenhauer’s notion of “will” (Schopenhauer, 1819), but this is an interpretive gloss, not part of the axioms.

Definition (D12: self). Let the system token “I” $\in \Sigma_{\text{anchor}}$ function as an anchor for first-person self-reference. At time t , define the set of tokens proximate to “I” as:

$$X_{I,t} := \{\tau \in \Sigma_t : d(\tau, I) \leq \delta_I\}.$$

For any set $X \subseteq \mathcal{U}$, define the association-closure under G $\text{cl}_G(X) \subseteq \mathcal{U}$ as “the smallest set containing X and closed under association operator G (A6).” The self-representation u_{self} is a representative Gestalt selected (and reinforced) within this closure:

$$u_{\text{self}} \in \Pi_t(\text{cl}_G(X_{I,t})).$$

Thus the self is a Gestalt composed from tokens proximate to the anchor “I” and strengthened through selection; the specific association order or implementation is not prescribed.

Remark. Because this construction forms a closure containing the anchor token “I,” it suggests the possibility of interpreting phenomena related to Hofstadter’s “strange loop” (Hofstadter, 1979) and self-referential representations in logic and set theory. This remark is interpretive, not part of the axioms.

Definition (D13: thought, emotion, sleep).

- (a) **Thought:** The process of retrieving higher-level Gestalts consistent with lower-level Gestalts, replaying subsequent Gestalts, and having the self treat them again as events to be represented.
- (b) **Emotion:** A Gestalt containing evaluative information (whether beneficial or detrimental) relevant to maintaining the existence of the self, together with associated sensory-source tokens.

- (c) **Sleep:** The process of testing whether (temporarily) formed Gestalt candidates can replay to reconstruct experience while reducing or blocking external input. Replay failure triggers reconstruction or new formation to reduce E_3 and improve restoration performance.

Definition (D14: dialogue, authoring, perception, social interaction).

- (a) **Dialogue:** In the SANC(E_3) formalism, dialogue can be modeled as follows: given a question Gestalt Q , retrieving stabilized question–answer (Q, A) Gestalts and activating the answer Gestalt A to complete the QA Gestalt, or generating a new (Q, A) Gestalt.
- (b) **Authoring:** Authoring can be described as a Gestalt-completion process: under partial cues, desire and will (D11) select a compatible higher-level Gestalt; subsequent Gestalts are internally replayed and cause corresponding external events (realizing the replayed representations as scores, sentences, or drawings). Authoring is complete when the system verifies that the selected Gestalt completes consistently with the intended bias (D11).
- (c) **Perception (vision/speech):** Partial sensory streams activate candidate Gestalts, which compete to complete latent structures that explain the input. The percept is the operation of internally finding a stabilized Gestalt that best satisfies finite capacity and E_3 , and mapping the given events to the Gestalt.
- (d) **Social interaction as a structural consequence:** At this abstract level, multi-agent interaction can be represented without adding a new primitive axiom. Other agents are external event sources to an individual intelligence (agent); D14 formalizes turn-taking as question–answer completion; A10 models theory of mind via interconnected subsystems—a formalization of Minsky’s “society of mind” (Minsky, 1986).

Definition (D15: self-improvement). ‘Self-improvement’ refers to the process in which a ‘self (D12)’ representation is generated within the system, emotion (D13) induces will (D11) in a direction favorable to maintaining the self’s existence, biasing/guiding the progression of thought (D13), and as a result, a new self is authored (D14).

8. Discussion – Relationship to Existing Theories

SANC(E_3) shares motivations with several established theories of cognition and intelligence, while differing in the structural commitments it makes. Table 1 summarizes these relations: for each prior framework, we identify the insight that SANC inherits or echoes, and the specific point at which SANC’s formulation diverges. The table is not a critique of these theories but a positioning statement, intended to make explicit which of SANC’s axioms or propositions corresponds to (or extends) ideas already present in the literature.

9. Scope and Limitations

SANC(E_3) is a theoretical framework rather than a complete implemented architecture. The axioms are proposed as candidate structural conditions for general intelligence, not as a complete mathematical sufficiency theorem for all possible intelligent systems. Empirical sufficiency also

Theory	Shared Insight	SANC Distinction
Global Workspace (Baars, 1988)	Competition for access	Similarity-based, not broadcast-based
IIT (Tononi, 2004)	Integration under constraints	Binding-based, not Φ -maximizing
Predictive processing (Clark, 2013)	Prediction at the core	Prediction = completion (T6)
Free Energy Principle (Friston, 2010)	Energy-based dynamics	E_3 explicitly decomposed
Transformers/LLMs (Vaswani et al., 2017)	Token-based prediction	Tokens are outcomes (T3)
Forward–Forward (Hinton, 2022)	Local learning	Formalized via A2, A3
Nested Learning (Behrouz and others, 2025)	Multiple time scales	Formalized via A3, T8
Hofstadter (GEB) (Hofstadter, 1979)	Strange loops	Derived from A6, A9, D12
Minsky (Society of Mind) (Minsky, 1986)	Interacting agents	Formalized via A10
Kantian <i>a priori</i> (Kant, 1781)	Structural premises	System tokens as network nodes
AIXI (Hutter, 2005)	Formal optimal agent	Computable, representation-emergent
NARS (Wang, 2019)	Finite resources, adaptation	Adds token emergence, hierarchy, energy decomposition
Chollet (Chollet, 2019)	Generalization efficiency	Generative constraints, not measurement

Table 1: Comparison of SANC(E_3) with existing theories.

depends on viable parameter regimes, scalable algorithms, suitable embodiment or input channels, and adequate environmental interaction.

Several components remain intentionally abstract. The similarity predicate, the selection operator, confidence-degree updates, the concrete form of C_{update} , and the optimization procedure for E_3 may be instantiated in different ways. This abstraction is useful for identifying common structural requirements across symbolic, neural, hybrid, and embodied systems, but it also means that the paper does not by itself provide a directly reproducible AGI implementation.

The unification claimed in this paper is structural rather than empirical: the framework shows how perception, prediction, reasoning, dialogue, authoring, social interaction, and action can be described within a common axiomatic pathway, but it does not claim that a current implementation has already demonstrated all of these capacities. Establishing such claims requires scalable implementations, ablation studies, and comparison with existing cognitive architectures and machine-learning systems.

10. Conclusion

SANC(E_3) formalizes general intelligence as an axiomatic framework of an event–representation system oriented toward E_3 minimization under finite capacity. This paper has proposed an axiomatic characterization of general intelligence. Ten axioms—five foundational (A1–A5) and

five structural (A6–A10)—together with the energy functional E_3 yield twelve derived propositions. These propositions suggest that token formation, forgetting, category emergence, hierarchy, and self-similar organization are not independent design choices but mutually entailed structural pressures. Three results merit emphasis:

Tokens as outcomes. Proposition T3 establishes that all tokens except system tokens are products of E_3 dynamics. This inverts the standard assumption: vocabulary is not an input to intelligence but an output of it.

Cognition as completion. Proposition T6 and the Corollary give prediction, dialogue, authoring, perception, causal inference, and embodied action a common structural interpretation as Gestalt completion. At the formal level considered here, no separate primitive mechanisms need to be postulated; practical implementations may nevertheless use distinct architectural modules.

Inevitability of forgetting. Proposition T2 proves that any system satisfying the axioms must forget. Memory is not merely storage but active curation under finite capacity—a feature, not a limitation.

10.1 Implications

If these axioms capture important structural conditions for general intelligence, then scaling alone may be insufficient to bridge the gap between current systems and AGI. What is missing may be not only more parameters but the right structural constraints: autonomous token formation (T3), similarity-based competition (A2), and the reconstruction–compression–stability trade-off (A4). Conversely, within the $SANC(E_3)$ formalism, systems satisfying A1–A10 should exhibit the derived properties under non-degenerate instantiations. The question “What is intelligence?” admits no single answer. $SANC(E_3)$ offers one: intelligence is E_3 -minimizing Gestalt completion under finite capacity, anchored by system tokens, generating its own representational vocabulary. Whether this characterization is correct is ultimately an empirical question—one we intend to pursue.

10.2 Future Work

This axiomatic system specifies what should exist in a system called intelligent under the $SANC(E_3)$ characterization, but does not prescribe a unique algorithm that satisfies it. Many ‘near-intelligent’ systems exist, from attention-based Transformers (Vaswani et al., 2017) to GPT (Radford et al., 2018; Brown and others, 2020), but human intelligence remains the only known system that appears to satisfy all these axioms. Implementing algorithms that satisfy this axiomatic system is a necessary step toward testing whether $SANC(E_3)$ can serve as a path toward artificial general intelligence (AGI).

We are developing a prototype that induces tokens from raw character streams across three hierarchical levels. Demonstrating that such a system satisfies A1–A10 and manifests T1–T12 would provide empirical grounding. Extending the framework to embodied agents—coupling E_3 to irreversible motor consequences—remains the principal theoretical task.

Acknowledgments

This work stands on a long intellectual lineage of humanity. From Confucius’s insight (Confucius, ca 500 BCE) that “to know what you know and to know what you do not know—that is knowledge,”

to Plato’s epistemology (Plato, ca 369 BCE) that “knowledge is justified true belief”; from the debates of Yi Hwang (Yi Hwang, 1568) and Yi I (Yi I, 1575) on the essence and priority of the Four Beginnings and Seven Emotions, to Kant (Kant, 1781) who viewed space and time as *a priori* forms of experience, and Schopenhauer (Schopenhauer, 1819) who grasped the world as will and representation.

From Wertheimer’s (Wertheimer, 1923) Gestalt psychology—“the whole is more than the sum of its parts”—to Saussure’s (de Saussure, 1916) structural linguistics in which linguistic units are defined relationally; from Hebb’s (Hebb, 1949) principle that “neurons that fire together wire together,” to Quillian (Quillian, 1968) who proposed associative networks of concepts.

From Hofstadter’s (Hofstadter, 1979) “strange loops” and “self-reference,” to Minsky’s (Minsky, 1986) “society of mind,” to Penrose (Penrose, 1989) who raised the non-algorithmic nature of consciousness—for over 2,000 years, all have asked the same question: What is knowledge? **What is intelligence?**

To this question—perhaps the oldest since humanity’s emergence—SANC(E_3) attempts to offer one answer. By stating in axiomatic form what any generally intelligent system must satisfy, and by integrating the insights of our predecessors into a minimal formal framework. We hope this attempt proves meaningful to all who have not stopped asking this question throughout humanity’s long existence, and to all who will continue to ask it.

References

- Assran, M.; Duval, Q.; Misra, I.; Bojanowski, P.; Vincent, P.; Rabbat, M.; LeCun, Y.; and Ballas, N. 2023. Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Behrouz, A., et al. 2025. Nested Learning: The Illusion of Deep Learning Architectures. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Bostick, D. 2024. What Is AGI? An Axiomatic Definition of General Intelligence as Lawful Autonomous Continuity. PhilArchive.
- Brown, T. B., et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Chollet, F. 2019. On the Measure of Intelligence. *arXiv preprint arXiv:1911.01547*.
- Clark, A. 2013. Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences* 36(3):181–204.
- Confucius. ca. 500 BCE. *The Analects*. (original work; multiple translations available).
- de Saussure, F. 1916. *Cours de linguistique générale (Course in General Linguistics)*. Payot, Paris.
- Friston, K. 2010. The Free-Energy Principle: A Unified Brain Theory? *Nature Reviews Neuroscience* 11:127–138.
- Hawkins, J. 2021. *A Thousand Brains: A New Theory of Intelligence*. Basic Books.

- Hebb, D. O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. Wiley.
- Hinton, G. 2022. The Forward-Forward Algorithm: Some Preliminary Investigations. *arXiv preprint arXiv:2212.13345*.
- Hofstadter, D. R. 1979. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books.
- Hutter, M. 2005. *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer.
- Kant, I. 1781. *Kritik der reinen Vernunft (Critique of Pure Reason)*. (original work; multiple modern editions available).
- LeCun, Y. 2022. A Path Towards Autonomous Machine Intelligence. OpenReview.
- Legg, S., and Hutter, M. 2007. Universal Intelligence: A Definition of Machine Intelligence. *Minds and Machines* 17(4):391–444.
- Meta AI. 2024. Large Concept Models: Language Modeling in a Sentence Representation Space. *arXiv preprint arXiv:2412.08821*.
- Minsky, M. 1986. *The Society of Mind*. Simon & Schuster.
- Penrose, R. 1989. *The Emperor's New Mind*. Oxford University Press.
- Plato. ca. 369 BCE. *Theaetetus*. (original work; multiple translations available).
- Quillian, M. R. 1968. Semantic Memory. In Minsky, M., ed., *Semantic Information Processing*. MIT Press. 227–270.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving Language Understanding by Generative Pre-Training. Technical report, OpenAI.
- Schopenhauer, A. 1819. *The World as Will and Representation*. (original work; multiple modern editions available).
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Tononi, G. 2004. An Information Integration Theory of Consciousness. *BMC Neuroscience* 5(42).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Wang, P. 2019. On Defining Artificial Intelligence. *Journal of Artificial General Intelligence* 10(2):1–37.
- Wertheimer, M. 1923. Laws of Organization in Perceptual Forms. *Psychologische Forschung* 4:301–350.

Yi Hwang. 1568. *Ten Diagrams on Sage Learning* (Seonghak Sipdo). (Korean classic; multiple modern editions available).

Yi I. 1575. *Essentials of the Learning of the Sages* (Seonghak Jibyo). (Korean classic; multiple modern editions available).

Appendix A. Summary of Token Types, Axioms, and Propositions

Token types

- System tokens (Σ_{sys}): Structurally given, including anchor tokens {here, now, I} and sensory-source tokens.
- Tokens ($\Sigma_t \setminus \Sigma_{\text{sys}}$): Representational units that emerge via self-organization and E_3 dynamics.

Axioms

Axiom	Description
A1	Finite capacity
A2	Similarity-based competition
A3	CD stabilization
A4	Energy E_3
A5	Co-occurrence \rightarrow association candidates
A6	Closure
A7	Splitting
A8	Uncertainty-gated segmentation & Gestalt completion
A9	Isomorphism across levels
A10	Subsystem interconnection

Propositions

Thm	Description	Dependencies
T1	No privileged tokens (except Σ_{sys})	A1, A2, A3
T2	Memory requires forgetting	A1, A2, A3, A5
T3	Tokens are outcomes	A3, A4, D1, D6
T4	Token–sequence balance	A4
T5	Categories as compression	A4, A5, A7, D7
T6	Completion = prediction	A4, D5, D8(a), D8(b)
T7	Hierarchy from closure	A4, A6, A10
T8	Self-similar levels	A6, A9, A10
T9	System tokens re-interpretable	D1, D9, T3
T10	Unsupervised learning	A1, A3, A4, D9
T11	Concept network	A5, A6, D1, D4, D6
T12	Threshold tuning	A4, T4, D6

Appendix B. Mapping from Original SANC Postulates

The SANC(E_3) formalization is based on the authors' original SANC postulates (C.1–C.31). The table below maps the original postulates to definitions, axioms, and propositions in this paper.

Original Postulate	This Paper
C.1 (events \rightarrow representations)	id in D1
C.2 (distinction and consistency)	D2
C.3 (co-occurrence and association/Gestalt)	A5, D4, D5
C.4 (similarity and splitting)	D3, A7
C.5 (category generation)	D7, T5
C.6 (partial-match retrieval)	D8(a)
C.7 (internal replay/regeneration)	D8(b)
C.8 (forgetting)	A3(b)
C.9 (restoration and active update)	A4(d)
C.10 (completion candidacy via uncertainty)	A8
C.11 (reinforcement/exclusion)	A3, D11
C.12 (competition and survival)	A2
C.13 (capacity constraint: token vs. sequence explosion)	A1, T4
C.14 (automatic threshold determination)	T12
C.15 (multi-level identical structure)	T8
C.16 (I/O isomorphism across levels)	A9
C.17 (concept network: SANC)	T11
C.18 (interconnected subsystems)	A10
C.19 (embedding by co-occurrence)	D9
C.20 (<i>a priori</i> representations)	D1, Σ_{sys}
C.21 (self as embedding)	D12, T8
C.22 (desire and will)	D10, D11
C.23 (thought)	D13
C.24 (emotion)	D13
C.25 (sleep)	D13
C.26 (dialogue)	D14
C.27 (creativity/writing)	D14
C.28 (vision/speech processing)	D14
C.29 (unsupervised learning)	T10
C.30 (self-improvement; environment improvement)	T10, Section 7
C.31 (unintentional intelligence)	Section 7

Appendix C. Original SANC Postulates (Extended)

This appendix presents the original SANC postulates, first formulated in the mid-1990s, in informal language. Many readers may find these more intuitive and accessible than the formal axioms in the main text.

C.1 From events to representations. A detected event generates a corresponding internal representation. An “event” may be sensory, motor, interoceptive, linguistic, abstract, or any detectable form.

C.2 Distinction and consistency. Intelligence can distinguish identical, similar, and different events, and can assign consistent internal representations to them.

C.3 Co-occurrence and association (Gestalt). Events detected together become associated. This associated event is a “Gestalt.” *Corollary:* A Gestalt can be treated as an event at a higher level and can be recombined.

C.4 Similarity and splitting. Similar events split into shared and differing parts (e.g., $ab =$ ‘with mom’ and $ac =$ ‘with dad’ share $a =$ with but differ in $b =$ mom vs. $c =$ dad).

C.5 Category (company) generation. If the differing parts co-occur with other events, categories are formed both for the shared parts with those other events and for the differing parts themselves. *Example:* If b and c also appear with d (i.e., $db =$ ‘for mom’, $dc =$ ‘for dad’ are experienced), then two categories $[a, d]$ and $[b, c]$ are formed (e.g., [with, for] and [mom, dad]).

C.6 Partial-match retrieval. Using stored Gestalts, the system can retrieve Gestalts that partially match a given event sequence.

C.7 Internal regeneration and reactivation. The system retrieves and internally regenerates/reactivates Gestalts that match the event sequence so far. *Corollary:* Internal regeneration fills “gaps” in the given sequence.

C.8 Forgetting. Non-*a priori* representations (including those generated from experienced events and *a priori* representations) are deleted if the corresponding event is not re-experienced within a certain time.

C.9 Restoration and active updating. The system restores all events so far using Gestalts (and actively updates Gestalts until restoration becomes complete).

C.10 Completion candidacy via uncertainty. As an experienced sequence grows, predicting the next event becomes easier and more certain; however, when a context is completed and a new context begins, next-event prediction suddenly becomes difficult. In this case, the sequence up to that point is treated as a completed-Gestalt candidate.

C.11 Reinforcement and exclusion. The more often a Gestalt is selected, the stronger it becomes; associated events that frequently lose in competition are excluded (via CD).

C.12 Competition and survival. Candidate Gestalts compete to explain the currently given events, and only one is selected.

C.13 Capacity constraint (token explosion vs. sequence explosion). Associated-event creation should be minimized. Creating higher-level Gestalts at every level causes a combinatorial explosion of representational units (token explosion). Failing to create higher-level Gestalts forces storage of too many event representations to restore experience (sequence explosion). With finite resources, intelligence balances restorability of experienced sequences against creation of new higher-level tokens.

C.14 Automatic threshold determination. Thresholds that suppress new higher-level token creation (balancing token vs. sequence explosion) and thresholds on CD that make a token persistent are not fixed *a priori*; they are determined dynamically by competition (C.12).

C.15 Multi-level identical structure. Intelligence consists of multiple levels with the same structure (fractal-like).

C.16 Input–output isomorphism across levels. Every level takes inputs of the same form and produces outputs of the same form.

C.17 Network of concepts (SANC). Each level produces higher-level events as associations of input events; such higher-level events form a network (graph) of co-occurring or co-occurable lower-level events. Viewing events as concepts, intelligence is the system that automatically and actively generates/organizes concepts and their networks in response to events.

C.18 Interconnected subsystems. Intelligence is a system of multiple such multi-level subsystems that are interconnected. Gestalt sets are exchanged among these vertical and horizontal interconnections, enabling formation and processing of different levels of concepts (representations) and different modalities.

C.19 Embedding (based on co-occurrence). Every event is embedded by the events it co-occurs with or can co-occur with, and by distances within that relation network. That is, embedding is the information contained in the ‘co-occurrence/co-occurrence’ relation network.

C.20 *A priori* representations exist. Some representations exist intrinsically rather than being generated from others (e.g., raw sensory signals; character codes in computers). Some *a priori* representations are embedded in all events (e.g., “now,” “here,” “I”).

C.21 Self. The self is a representation (embedding) of the intelligence itself, and a set of representations lying within a certain distance from the *a priori* representation “I.”

C.22 Desire and will (selection bias). Desire is the tendency to internally complete (realize) a Gestalt from given events. Will is a bias that makes certain Gestalts among many selectable candidates be selected/regenerated more often.

C.23 Thought. A process in which a top-level Gestalt retrieves and internally represents/constructs the next Gestalt at the next-lower level.

C.24 Emotion. A Gestalt that encodes information about being favorable or unfavorable for maintaining the existence of the “self.”

C.25 Sleep. A process that tests whether (temporarily) formed Gestalts can reconstruct experienced reality while reducing/blocking external input. If reconstruction fails, the system reconstructs/creates Gestalts and repeats internal replay to improve restoration performance.

C.26 Dialogue. Treating a question–answer (QA) pair as a unit Gestalt, a dialogue is a sequence of QA Gestalts. When only Q is given from outside, completion pressure selects/regenerates an appropriate A to complete the QA Gestalt; this is answer generation.

C.27 Creativity/writing. Completing a partially specified Gestalt (topic/constraints/structure) via replay/regeneration in a direction that reduces E_3 , producing a novel structured artifact.

C.28 Visual/speech processing. When partial sensory streams activate candidate Gestalts, the candidates compete to complete a latent structure that explains the input. Recognition is the stabilized completion that best fits under finite capacity and E_3 .

C.29 Unsupervised learning. A system with the above properties can learn from event sequences without explicit external supervision.

C.30 Self-improvement and environment improvement. Such a system can select or create environments advantageous for survival and can improve itself.

C.31 Unintentional intelligence. Functions often called the “essence” of intelligence—prediction, memory, learning, classification, creativity, reasoning, etc.—arise necessarily and as by-products in systems that satisfy the SANC postulates.