FOUNDING VISION — SETALEUR APLAMDA RESEARCH PROGRAMME

Structuralist Artificial Intelligence: Intelligence as the Formation and Transfer of Relational Structure from Experience

Momen Ghazouani
Chief Scientist, Setaleur Aplamda
April 2026 github.com/Setaleur-Aplamda

Abstract

We introduce **Structuralist Artificial Intelligence (SAI)**, a research vision that defines intelligence as the capacity to form, reorganize, and transfer relational structures from experience. This vision is grounded in a concrete biological observation: the learning brain does not merely adjust the strengths of existing connections it reorganizes its connectivity topology. Synaptic pruning, dendritic spine formation, and Hebbian consolidation collectively produce not a weight matrix with updated values but a *relational topology that encodes the structural regularities of experience.* The topology is the knowledge.

SAI translates this observation into three falsifiable architectural requirements: **(R1) Structural Formation** learning systems must extract latent relational patterns, not surface statistical correlations; **(R2) Topological Reorganization** acquired structures must be dynamically reorganizable under novel inputs without catastrophic interference; and **(R3) Structural Transfer** representations must generalize across domains that share underlying relational geometry, independently of surface presentation.

We argue that these requirements define a research programme that is strictly more specific than any appeal to general intelligence, strictly more measurable than any behavioral imitation target, and strictly more honest than any claimed capability horizon. SAI does not propose a path toward human-like cognition; it proposes a path toward systems that know *what they know* systems whose confidence is grounded in structural evidence, whose generalization is bounded by structural distance, and whose failure modes are predictable from structural geometry rather than concealed by statistical averaging.

The Setaleur Aplamda research programme operationalizes SAI through the AI Implicit paradigm and its architectural family, evaluated through the Experience-Compressed Intelligence (ECI) measurement framework. This paper establishes the theoretical foundations of that programme.

Keywords: Structuralist Artificial Intelligence, Relational Structure, Structural Topology, Hebbian Learning, Synaptic Pruning, Experience Compression, Structural Transfer, Epistemic Awareness, AI Implicit.

1. The Question Behind the Question

1.1. What Changes When a System Learns?

The canonical question of artificial intelligence research has been: *what can this system do?* SAI asks a prior question whose answer constrains the first: *what changes inside the system when it genuinely learns?*

This is not a rhetorical distinction. Two systems that achieve identical performance on a benchmark may have undergone qualitatively different internal transformations. The first may have adjusted numerical parameters in proportion to gradient signals, distributing prediction credit across millions of weights with no interpretable structural change. The second may have formed a compact relational representation a topology of connections encoding invariant patterns that generalize across surface presentations of the same underlying structure. These two systems are not equivalent learners who happened to reach the same score. They are in different epistemic states, and their behavior under distributional shift will differ accordingly.

1.2. Why Existing Frameworks Cannot Answer It

Three dominant frameworks in artificial intelligence provide answers to *what systems can do*, but none provides a principled answer to what changes when they learn.

Behavioral frameworks. Turing-style evaluation [14], task-based benchmarks [15], and economic productivity metrics [16] measure external behavior against predefined tasks. They say nothing about the internal state change that produced that behavior, whether it was efficient or wasteful, robust or brittle, structurally grounded or statistically memorized.

Statistical optimization. The dominant paradigm treats learning as the minimization of expected loss over

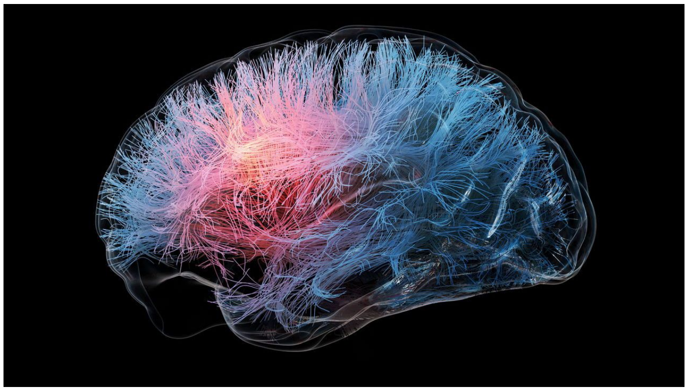

Figure 1: A tractography visualization of white matter fiber pathways in the human brain. The architecture of interconnected pathways illustrates the kind of relational topology that SAI takes as its biological ground: a geometry of connections shaped by experience, not a fixed graph with updated scalar weights. This is the computational target of the SAI framework not the adjustment of edge strengths over a static structure, but the reorganization of relational geometry as the primary product of learning.

a labeled distribution [19]. This framework has a precise answer to what changes: parameter values shift in the direction of negative gradient. What it does not answer and cannot, by the structure of the objective is whether those parameter changes constitute structural knowledge extraction or surface correlation memorization.

Capability-horizon frameworks. Definitions organized around Artificial General Intelligence [17, 18] specify a target state rather than a learning mechanism. They define intelligence by what a hypothetical system *could* do, not by the structural properties of what any real system has learned. This makes them aspirationally useful but scientifically intractable: there is no experiment that falsifies a definition organized around an ill-bounded capability target.

SAI answers the prior question directly: *learning is the formation of relational structure from experience.* The internal state change that constitutes genuine learning is a change in the relational topology of the learned representation. This answer is falsifiable, measurable, and architecturally actionable.

2. The Biological Ground

We do not appeal to neuroscience as a source of metaphors. We appeal to it as a source of existence proofs: demonstrations that the learning mechanisms we propose are not theoretical constructions but properties of the most efficient learning system known.

2.1. Three Mechanisms, One Principle

The mammalian brain learns through three mechanisms that collectively instantiate structural reorganization rather than weight adjustment [2–4].

2.1.1. Hebbian Consolidation

Donald Hebb's 1949 postulate *neurons that fire together, wire together* describes not the strengthening of a fixed connection but the formation of a preferential relational pathway [1]. When two neurons co-activate repeatedly, the synaptic connection between them is biochemically reinforced through long-term potentiation (LTP), increasing the probability that future activation of one will propagate to the other.

The computational consequence is not a weight update in a fixed graph; it is the *emergence of a functional edge in a relational structure*. The topology of the network changes: a connection that was statistically negligible becomes structurally significant. Repeated co-activation of a pattern of neurons produces a stable attractor in the network's relational geometry a structural encoding of that pattern that persists independently of the individual firing events that created it.

This is the biological foundation of **Structural Formation**: the primary product of learning is a change in relational topology, not a change in scalar connection strength.

2.1.2. Synaptic Pruning

The developing mammalian brain is *overconnected*. Postnatal synaptic density far exceeds that of the adult brain; the trajectory from infancy to adulthood is characterized by the systematic elimination of underutilized synapses, a process termed synaptic pruning [5, 6].

This is not a passive decay process. Pruning is activity-dependent: synapses that participate in consistent co-activation patterns are preserved; those that are not recruited into stable relational structures are retracted. The result is a connectivity topology that is sparse, selective, and structurally organized around the regularities of the organism's experience a compressed relational representation of the statistical structure of the environment.

The computational consequence is profound and directly opposed to the trajectory of gradient-based systems, which tend to grow in parameter count as tasks grow in complexity. Biological learning *compresses*: it starts with redundant connectivity and converges to a sparse, efficient relational topology. The sparsity is not a constraint imposed from outside; it is the natural endpoint of structural optimization under activity.

2.1.3. Structural Plasticity

Beyond the strengthening and pruning of existing synapses, the adult brain generates entirely new synaptic contacts in response to novel experience a phenomenon termed structural plasticity or experience-dependent synaptogenesis [2, 3]. Dendritic spines, the postsynaptic terminals through which neurons receive input, grow, retract, and stabilize in response to patterns of neural activity. New spines form transiently in response to learning and stabilize into permanent contacts if the associated activity patterns persist.

This mechanism implements **topological addition**: the relational graph of the network gains new edges, not merely new edge weights, in response to learning. The structure of knowledge is not fixed at initialization and subsequently parameterized; it is dynamic and continuously reorganized by experience.

2.2. The Computational Translation

Three mechanisms, one principle: learning is topological reorganization. The biological learning system does not maintain a fixed graph and adjust its edge weights. It maintains a dynamic relational topology that grows, prunes, and reorganizes in response to the structural regularities of experience. The final state of learning is a relational geometry, not a weight vector.

Table 1 summarizes the computational translation of each biological mechanism into an SAI architectural requirement.

Table 1: Biological mechanisms and their computational translations in SAI.

Mechanism	**Biological function**	**SAI requirement**
Hebbian LTP	Relational edge formation	R1: Structural Formation
Synaptic Pruning	Topology compression	R2: Structural Compression
Structural Plasticity	Graph reorganization	R3: Topological Flexibility

2.3. What Gradient Descent Does Not Do

It is important to state this contrast precisely, because the claim is not that gradient-based systems do not learn they clearly do but that what they learn is structurally different from what biological systems learn.

Gradient descent, applied to a fixed-topology neural network, adjusts scalar connection weights in proportion to their contribution to reducing a scalar loss. The topology of the network which neurons are connected to which is fixed at initialization and does not change during training. The relational structure of the learned representation is therefore an incidental byproduct of weight adjustment under a prediction objective, not a primary target of the learning process.

This distinction has two concrete consequences. First, gradient-based systems have no mechanism for discovering that a structural feature is irrelevant and should be topologically removed they can reduce its weight but cannot prune its existence. Second, they have no mechanism for generating a new relational edge in response to a novel pattern they can only redistribute weight over an existing fixed graph.

The biological system, by contrast, can both add and remove relational edges. Its structural topology at the end of learning is an active encoding of the regularities of experience. Gradient descent produces a parameterized fixed graph; biological learning produces a dynamic relational topology.

SAI proposes that the *relational topology*, not the weight vector, is the appropriate unit of learned knowledge.

3. The SAI Definition

3.1. Formal Statement

Definition 1 (Structuralist Intelligence)**.** A learning system exhibits *structuralist intelligence* to the degree that it satisfies three measurable properties:

(i) **Structural Formation**: Given a set of experiences $\mathcal{E}$, the system produces a relational representation $\mathcal{R}(\mathcal{E})$ that encodes the latent invariants and compositional patterns of $\mathcal{E}$ with compression efficiency exceeding that of surface statistical encoding.

(ii) **Topological Coherence**: The representation $\mathcal{R}(\mathcal{E})$ is geometrically organized such that structurally similar inputs produce nearby activations and structurally distinct inputs produce separated activations, independently of surface presentation.

(iii) **Structural Transfer**: For any target domain $\mathcal{D}_t$ sharing latent structural geometry with the source domain $\mathcal{D}_s$, the system's performance on $\mathcal{D}_t$ exceeds the performance achievable by surface statistical generalization from $\mathcal{D}_s$ to $\mathcal{D}_t$.

3.2. Three Properties Unpacked

3.2.1. Structural Formation

Structural formation is the capacity to extract the *generative pattern* underlying observed data the compact relational description that could have produced the observations, rather than the surface statistical signature of the observations themselves.

The distinction is empirically testable. A system that has learned the *topology* of handwritten digits that the digit zero has one enclosed region, the digit eight has two, the digit one has none encodes knowledge that is invariant to handwriting style, ink thickness, and rotation. A system that has memorized the pixel-intensity distribution of MNIST digits has encoded knowledge that is sensitive to all of these. The topological system generalizes from five examples; the statistical system requires thousands [26].

Structural formation is therefore not merely more efficient it is categorically different in kind. The knowledge it produces is *structural fact*, not *statistical correlation*.

3.2.2. Topological Coherence

Topological coherence is the property that the learned representation organizes inputs by their structural identity rather than their surface presentation. This is measurable through the inter-class separation ratio in the learned representation space: if structurally distinct classes occupy well-separated regions regardless of surface variation, and structurally similar inputs (e.g., the same digit in different handwriting styles) occupy nearby regions, the representation is topologically coherent.

This property is directly analogous to the categorical perception exhibited by the biological auditory cortex: phonemes that are acoustically similar but phonologically distinct are represented in separated neural populations; phonemes that vary acoustically but belong to the same phonological category are represented in overlapping populations [7]. The brain encodes structural categories, not acoustic distributions.

Topological coherence has a further consequence for epistemic awareness: a system with a coherent relational geometry can detect structural unfamiliarity by measuring the distance of a novel input from known class prototypes in representation space. This produces epistemic confidence as a geometric emergent without requiring any learned uncertainty estimator.

3.2.3. Structural Transfer

Structural transfer is the capacity to apply structural knowledge acquired in one domain to a domain that shares underlying relational geometry, despite surface dissimilarity. This is the empirical signature that distinguishes genuine structural learning from surface memorization: a system that has encoded the structural description of a pattern can apply that description wherever the pattern occurs, regardless of the medium in which it appears.

Biological expert learners exhibit structural transfer reliably. Chess grandmasters apply strategic pattern recognition to novel board positions because they have encoded the abstract relational structure of tactical situations, not the specific pixel appearance of individual boards [8]. Physicians trained in one diagnostic domain apply structural reasoning to adjacent domains because the underlying relational patterns causal chains, symptom clusters, differential categories are structurally shared.

The capacity for structural transfer is precisely what the Cross-Domain Retention (CDR) metric in the ECI framework is designed to measure: the fraction of structural performance retained when the surface presentation of the domain changes while the underlying relational geometry is preserved.

3.3. What SAI Does Not Claim

Scientific integrity requires that the boundaries of this framework be stated as explicitly as its positive claims.

SAI does not claim to define human intelligence. The definition in Section 3.1 characterizes a measurable set of properties. Human cognition satisfies many of these properties and likely additional ones (embodied agency, social cognition, metacognitive regulation) that are outside this framework's current scope.

SAI does not claim that structural learning is sufficient for all intelligent behavior. Structural representation is a necessary substrate for the properties we describe. It is not claimed to be sufficient for open-ended agency, natural language understanding, or sensorimotor control in unstructured environments.

SAI does not propose a capability horizon. There is no system that "achieves" SAI in the way that one might "achieve" a behavioral benchmark. SAI provides a continuous measurement framework. Systems occupy a position on the structural learning spectrum; the framework measures where they are, not whether they have arrived somewhere.

SAI does not evaluate systems by what they can do. Two systems with identical task performance may occupy very different positions on the structural learning spectrum. SAI's measurement target is the learning process itself, not its behavioral output.

4. The SAI–Optimization Distinction

4.1. A Precise Formulation

The distinction between SAI and the statistical optimization paradigm is not a philosophical disagreement about ultimate goals. It is a concrete difference in the learning objective and the representation target.

Statistical optimization minimizes a scalar loss $\mathcal{L}$ over a fixed representation topology:

$$\theta^* = \arg\min_{\theta} \; \mathbb{E}_{(x,y)\sim\mathcal{D}}[\mathcal{L}(f_\theta(x), y)] \,. \tag{1}$$

The learned parameter vector θ^* encodes whatever statistical regularities reduce $\mathcal{L}$ most efficiently. The topology of the graph f is fixed; only edge weights change. The objective does not distinguish between structurally informative regularities and spuriously correlative ones.

SAI targets a different objective. Let $\mathcal{G}(\mathcal{E})$ denote the latent generative structure underlying a set of experiences $\mathcal{E}$, and let d_{struct} be a structural distance measure on representation space. The structural learning objective is:

$$\mathcal{R}^* = \arg\min_{\mathcal{R}} \; \mathbb{E}_{\mathcal{E}}[d_{\text{struct}}(\mathcal{R}(\mathcal{E}),\, \mathcal{G}(\mathcal{E}))] \,. \tag{2}$$

The optimization target in (2) is not a scalar loss over labeled pairs but a structural distance between the learned representation and the latent generative structure. The topology of $\mathcal{R}$ is not fixed; it is precisely what is being learned.

4.2. Three Failure Modes of Statistical Optimization

When the objective is (1) rather than (2), three systematic failure modes are structurally unavoidable not engineering deficiencies, but necessary consequences of the objective.

Statistical conflation. Gradient descent cannot distinguish a regularity that encodes structural identity from one that encodes surface correlation [20]. Both reduce $\mathcal{L}$ equivalently. The result is representations that encode texture when shape would generalize, background when object structure would transfer, and co-occurrence when causal structure would predict.

Topological rigidity. Because the graph topology is fixed at initialization, the learned representation cannot perform synaptic pruning (removing structurally irrelevant connections) or structural plasticity (generating connections for novel structural patterns). The representation is always a reparameterization of the initial topology.

Epistemic opacity. The softmax normalization of output probabilities ensures that the sum of class confidences equals one, regardless of whether the input lies within or outside the training distribution. There is no architectural representation of absolute structural distance from known patterns only the relative distribution of prediction credit over a fixed class set [21].

None of these failure modes can be corrected post-hoc without addressing the optimization objective itself. Calibration layers, dropout uncertainty, and ensemble methods treat the symptoms of epistemic opacity without altering the cause.

4.3. Gradient Descent as a Tool, Not an Obstacle

SAI does not prohibit gradient-based optimization. It prohibits defining the learning objective as surface prediction accuracy. Gradient descent is a powerful optimization mechanism; the question is what it is applied to optimize.

Within the SAI framework, gradient-based methods are appropriate when the optimization target is structural: when gradient signals carry information about the geometry of relational representations, the density of structural prototypes, or the compactness of class topology. Gradient descent used to minimize structural reconstruction error, to organize prototype geometry, or to maximize inter-class relational separation is an SAI-compatible tool.

The architectural commitment of SAI is not to gradient-freedom but to *structural primacy*: the relational topology of the learned representation is the primary design target, not a byproduct of prediction accuracy minimization.

5. The SAI Research Programme

5.1. Scope and Boundaries

The SAI research programme addresses systems that learn from structured experience systems for which the concept of *latent relational geometry* is well-defined. This includes visual recognition, symbolic sequence processing, relational reasoning, and any domain where the observed data is generated by a consistent structural process.

The programme does not currently address open-ended embodied agency, affective cognition, or social intelligence. These are not excluded on principled grounds structural learning is potentially relevant to all of them but they are outside the scope of the current empirical and theoretical development, and honesty requires acknowledging that scope.

5.2. The AI Implicit Paradigm

The AI Implicit paradigm [22] is the operational expression of SAI within the programme. It translates the three SAI requirements into concrete architectural principles:

Tacit Structure Extraction (from R1): The primary learning signal is the recovery of latent relational patterns, not the minimization of surface prediction error.

Experience Compression (from R2): The efficiency of structural extraction is measured by knowledge density the amount of reusable relational representation formed per unit of training experience.

Epistemic Confidence (from R3): A system that cannot recognize when a novel input is structurally distant from its learned representation is not a reliable learner. Epistemic awareness is an architectural requirement, derived from the geometry of the learned relational field.

5.3. The ECI Measurement Framework

The Experience-Compressed Intelligence (ECI) framework [23] provides the primary measurement instruments for SAI programme evaluation. Its metrics Compression Ratio, Tacit Knowledge Extraction Rate, Cross-Domain Retention, and Experience Efficiency Index are designed to measure the three SAI properties directly, without appeal to task-specific accuracy as a proxy.

The critical property of ECI as an SAI measurement tool is that it produces *internally grounded* measurements: the ECI score of a system reflects properties of the learning process itself, not the difficulty of the evaluation task. A system that achieves 95% task accuracy through surface memorization correctly receives a low ECI score; a system that achieves 83% accuracy through structural compression correctly receives a higher ECI score that reflects its structural efficiency. This separation of accuracy from structural quality is the central evaluative commitment of the SAI programme.

5.4. The Architectural Family

The Setaleur Aplamda research programme develops three complementary architectural realizations of SAI principles, each addressing a distinct domain:

Linear Networks [24]: A gradient-free relational density architecture for feature-space structural learning. Demonstrates that epistemic confidence can be produced as a geometric emergent of the relational field without gradient-based training.

Deep Transducers [25]: A gradient-informed architecture for structural learning in symbolic sequence domains, using gradient descent as a compression tool over learned prototype geometry rather than as a prediction optimizer.

TSNet [26]: A gradient-free visual architecture grounded in the Generative Trace Hypothesis, demonstrating that topological and path-geometric features computed deterministically from image structure constitute a sufficient basis for classification with calibrated epistemic awareness.

Each architecture demonstrates one dimension of the SAI framework. Together, they constitute a body of empirical evidence that structural primacy is feasible across distinct learning domains.

6. Falsifiable Predictions

A research vision that makes no falsifiable predictions is philosophy, not science. We state six concrete predictions of the SAI framework. Each is formulated so that a well-designed experiment can confirm or refute it.

Prediction 1 (Logarithmic Saturation)**.** SAI systems will exhibit logarithmic learning curves rapid early performance improvement followed by structural saturation while accuracy-optimized systems on the same task will exhibit linear scaling. The point of saturation corresponds to the example count at which the latent structural description of the class is fully recovered.

Prediction 2 (Transfer Asymmetry)**.** SAI systems will exhibit substantially higher Cross-Domain Retention on tasks sharing underlying structural geometry than on tasks sharing only surface statistical properties. Accuracy-optimized systems will exhibit the reverse: surface-similar domains will transfer more readily than structurally similar ones.

Prediction 3 (OOD Collapse)**.** For inputs generated by a structurally distinct process regardless of their surface visual richness or complexity SAI systems will produce collapsed epistemic confidence signals. Accuracy-optimized systems will assign high softmax confidence to structurally unfamiliar inputs, because their confidence is a function of decision-boundary geometry rather than structural proximity.

Prediction 4 (Separation Without Optimization)**.** A structural feature space computed deterministically from domain geometry will produce inter-class separation ratios exceeding those achievable by gradient descent on surface pixel features, because structural features encode class-defining invariants while gradient-based features encode class-predictive correlations.

Prediction 5 (Compression Threshold)**.** For any domain whose class identity is encoded in a compact structural description, there exists a minimum example count n^* below which SAI systems produce unreliable structural representations and above which performance saturates rapidly. This threshold n^* is determined by the statistical variance of the structural description, not by the complexity of the domain.

Prediction 6 (Topology Determines Epistemic Signal)**.** The epistemic confidence of an SAI system will be more strongly correlated with the topological coherence of the input (its structural consistency with the learned relational geometry) than with the input's surface complexity, pixel informativeness, or semantic richness. Structurally incoherent inputs will receive low confidence regardless of their visual content.

These six predictions are not exhaustive. They are the minimum falsifiable content required to distinguish the SAI framework from a philosophical position. Architectures developed within the SAI programme will be evaluated against them.

7. Philosophical Foundations

7.1. Structure as the Substrate of Knowledge

The proposition that knowledge consists in relational structure rather than propositional content has a long intellectual lineage, from Kant's transcendental schemata [11] through Piaget's constructivism [12] to the Gestalt tradition's emphasis on structural organization over elemental composition [13]. SAI is not a novel philosophical position; it is the architectural operationalization of a position that cognitive science has maintained for nearly a century.

Polanyi's concept of tacit knowledge [9] is particularly relevant: the knowledge that experts possess and exercise is not propositional but structural it consists in the relational organization of experience into patterns that guide perception and action without explicit articulation. Polanyi's physician, who recognizes a diagnosis from the structural configuration of symptoms, is exercising precisely the capacity that SAI is designed to formalize and measure.

7.2. Knowledge Density as an Information-Theoretic Concept

The SAI criterion for progress knowledge density, not task accuracy has a precise information-theoretic grounding. Kolmogorov complexity defines the complexity of a pattern as the length of the shortest program that generates it [10]. Intelligence, under this lens, is the capacity to discover short programs for complex phenomena to find the compact structural description that generates observed data.

A system that memorizes N training examples encodes $O(N)$ information; no compression has occurred. A system that extracts the structural description of the generative process underlying those examples encodes $O(K)$ information, where $K \ll N$ is the complexity of the generative structure. The ratio N/K is a natural measure of learning efficiency; it is precisely what the SAI measurement framework quantifies through the ECI metrics.

7.3. The Structuralist Epistemology

SAI adopts what might be called a *structural epistemology*: the position that what a system knows is constituted by its relational representation of the world, not by its ability to produce correct outputs under queried conditions. This has a concrete evaluative consequence: a system that knows the structural description of a domain should be able to generalize it, transfer it, and recognize its limits. A system that has merely memorized correct outputs cannot do any of these things.

This epistemological position implies a specific view of epistemic failure. A system fails epistemically not when it produces a wrong output but when it *does not know that it is likely to be wrong* when its confidence is uncorrelated with its structural competence. The most dangerous failure mode is not incorrect prediction but incorrect *certainty* about incorrect prediction.

SAI designs for epistemic awareness as a primary property: a system built on structural relational representations naturally produces confidence signals from the geometry of those representations, because structural familiarity the degree to which a novel input is geometrically consistent with the learned relational topology is directly computable from the representation itself.

8. Limitations and Honest Assessment

Scientific integrity requires that the limitations of this framework be stated as precisely as its contributions.

The structural description problem. SAI assumes that, for a given domain, a compact structural description exists and can be extracted. For domains with genuinely high-dimensional, compositional structure natural language, open-ended embodied interaction, social cognition the existence of such a description is not guaranteed. The programme currently focuses on domains where the structural hypothesis is well-motivated; extension to complex domains requires empirical validation, not theoretical assumption.

The topology specification problem. The SAI objective in Equation (2) requires a specification of the latent generative structure $\mathcal{G}(\mathcal{E})$ against which to measure representation fidelity. In practice, this structure is not directly observable; it must be estimated from data. The quality of structural learning therefore depends on the quality of the structural estimator, and this dependency is not fully resolved by the current framework.

The biological analogy is an existence proof, not a blueprint. We invoke Hebbian consolidation, synaptic pruning, and structural plasticity as demonstrations that topology-changing learning is computationally feasible and biologically instantiated. We do not claim that artificial systems should replicate these mechanisms in detail. The biological evidence motivates the computational target; it does not specify the implementation path.

The programme is at an early stage. The architectural realizations of SAI described in Section 5 are proof-of-concept demonstrations on tractable benchmarks. The generalization of SAI principles to large-scale, multi-modal, and continuously learning systems is future work, not established fact. We state this explicitly because the history of artificial intelligence research is littered with frameworks whose proof-of-concept results were overgeneralized before the generalization was warranted.

9. Conclusion

We have introduced Structuralist Artificial Intelligence as a research vision grounded in a precise biological observation and expressed in three falsifiable architectural requirements.

The biological observation is this: the learning brain reorganizes its connectivity topology. Through Hebbian consolidation, synaptic pruning, and structural plasticity, experience produces not updated weight values but an altered relational geometry a topology that encodes the structural regularities of the world rather than the statistical co-occurrence of its surface presentations. The topology is the knowledge.

The three requirements translate this observation into measurable architectural targets: structural formation (the extraction of latent relational patterns), topological coherence (the geometric organization of representations by structural identity), and structural transfer (the generalization of structural knowledge across domains sharing relational geometry).

These requirements define a research programme that is more specific than any appeal to general intelligence, more measurable than any behavioral imitation target, and more honest than any claimed capability horizon. Progress within SAI is measured not by reaching a destination but by advancing along three well-defined dimensions dimensions for which the ECI framework provides concrete, falsifiable metrics.

What SAI offers that prior frameworks do not is not a more ambitious target. It is a more honest one: a definition of progress that can be observed, measured, and falsified at every stage of architectural development, in every domain where structural learning is feasible. Systems that satisfy the SAI requirements at high levels will be systems that know what they know systems whose confidence is grounded in structural evidence, whose generalization is bounded by structural distance, and whose failure modes are predictable from relational geometry rather than concealed by statistical averaging.

That is not a path toward human cognition. It is a path toward something more achievable and more useful: artificial systems with genuine structural competence, and genuine awareness of its limits.

References

[1] Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory.* Wiley.

[2] Bhatt, D. L., et al. (2009). Dendritic spine dynamics. *Annual Review of Physiology*, 71, 261–282.

[3] Holtmaat, A., & Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, 10(9), 647–658.

[4] Chklovskii, D. B., Mel, B. W., & Svoboda, K. (2004). Cortical rewiring and information storage. *Nature*, 431(7010), 782–788.

[5] Huttenlocher, P. R. (1990). Morphometric study of human cerebral cortex development. *Neuropsychologia*, 28(6), 517–527.

[6] Feinberg, I. (1982). Schizophrenia: Caused by a fault in programmed synaptic elimination during adolescence? *Journal of Psychiatric Research*, 17(4), 319–334.

[7] Liberman, A. M., et al. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, 54(5), 358–368.

[8] Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.

[9] Polanyi, M. (1966). *The Tacit Dimension.* University of Chicago Press.

[10] Kolmogorov, A. N. (1963). On tables of random numbers. *Sankhyā: The Indian Journal of Statistics, Series A*, 369–376.

[11] Kant, I. (1781). *Critique of Pure Reason.* Translated by N. K. Smith (1929), Macmillan.

[12] Piaget, J. (1936). *Origins of Intelligence in the Child.* Routledge & Kegan Paul.

[13] Köhler, W. (1929). *Gestalt Psychology.* Liveright.

[14] Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.

[15] Chollet, F. (2019). On the measure of intelligence. *arXiv:1911.01547.*

[16] Davidson, T. (2023). What a compute-centric framework says about takeoff speeds. *AI Alignment Forum.*

[17] Goertzel, B. (2014). Artificial general intelligence: Concept, state of the art, and future prospects. *Journal of Artificial General Intelligence*, 5(1), 1–46.

[18] Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4), 391–444.

[19] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

[20] Geirhos, R., et al. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11), 665–673.

[21] Nguyen, A., Yosinski, J., & Clune, J. (2015). Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *Proc. CVPR*, 427–436.

[22] Ghazouani, M. (2026) "AI Implicit A Foundational Paradigm for Intelligence Through Experience Compression," Zenodo . https://doi.org/10.5281/zenodo.19947976

[23] Ghazouani, M. (2026) "Experience-Compressed Intelligence (ECI)," Zenodo . https://doi.org/10.5281/zenodo.19957177

[24] Ghazouani, M. (2026). Linear Networks: A gradient-free architecture for AI Implicit., Zenodo

[25] Ghazouani, M. (2026). Bold learning is all we need ,Zenodo

[26] Ghazouani, M. (2026) "TSNet Topological Structure Networks," Zenodo . https://doi.org/10.5281/zenodo.19928909