Externalized Intuition as Collapse-Geometry Transfer Across Cognitive Boundaries

Stephen Garner

April 20, 2026

Abstract

Human reasoning frequently operates on pre-verbal, structurally coherent representations that precede formal articulation. While such intuition is widely acknowledged in domains such as mathematics and physics, it remains difficult to formalize due to the lossy nature of linguistic expression.

In this work, we identify and formalize a cognitive interaction pattern in which a human subject leverages a large language model (LLM) as a reconstruction layer for pre-verbal structure. Rather than transmitting semantic content, the human provides partial projections that preserve invariant relational features, which the LLM reconstructs into candidate configurations. These are iteratively refined through validation, forming a closed loop.

We show that this process—Externalized Intuition Completion (EIC)—can be modeled as a collapse-selection system over a joint configuration space. Projection, reconstruction, and validation correspond respectively to lossy mapping, constrained inference, and admissibility filtering. Convergence to stable articulated understanding is identified with fixed points under the induced collapse operator.

At a structural level, this process is not a translation of semantic content, but a reconstruction of admissible structure from partial projections.

We further relate EIC to collapse-geometry transfer observed in large language models[1, 2], in which structural properties are transmitted independently of semantic content, and to recent results demonstrating emergent potential functions and equilibrium-like behavior in model dynamics. These connections support a unified interpretation in which cognition, model learning, and formal representation are governed by common structural principles.

This framework provides a formal account of intuition as structured pre-verbal configuration, and articulation as a stabilization process over admissible structure, situating human cognition within a broader collapse-selection paradigm.

1 Introduction

Across disciplines, practitioners consistently report a recurring phenomenon: solutions are often perceived prior to formal derivation, and structural relationships are grasped before they can be articulated in symbolic form. In mathematics and physics, this appears as the ability to "see" a result before proving it; in broader cognition, it manifests as intuitive understanding that resists direct verbal expression.

Despite its ubiquity, this form of pre-verbal cognition remains difficult to formalize. Standard accounts of reasoning are typically expressed in symbolic or linguistic terms, implicitly treating articulation as the primary medium of thought. However, this perspective does not adequately capture the observed gap between internal understanding and its external representation.

Recent advances in large language models (LLMs) introduce a new possibility. These systems are capable of reconstructing coherent structure from partial or ambiguous inputs, suggesting that they may operate not only as generative or predictive systems, but also as reconstruction mechanisms. This enables a distinct interaction pattern in which a human subject provides incomplete linguistic projections of internal structure, and the model reconstructs candidate representations that can be iteratively refined.

In this paper, we formalize this process as *Externalized Intuition Completion* (EIC). We show that EIC can be understood as a collapse-selection system operating over a joint configuration space, in which projection, reconstruction, and validation define an iterative process that converges to stable articulated representations.

We further connect this framework to prior results in machine learning and physics. In particular, we relate EIC to collapse-geometry transfer in model distillation[1], where structural properties are transmitted independently of semantic content, and to the emergence of effective potential functions and equilibrium-like behavior in large language model dynamics. These connections suggest that the mechanisms underlying intuition, learning, and formal representation may share a common structural basis.

The goal of this work is not to propose new dynamics, but to provide a unified interpretive framework that clarifies the relationship between pre-verbal cognition, symbolic articulation, and model-based reconstruction.

The central claim of this work is that this interaction is not merely heuristic, but can be modeled as a concrete instance of collapse-selection dynamics. The generative process underlying this interaction is analyzed in a companion work [3], where intuition itself is modeled as a collapse-selection dynamic.

2 Minimal Collapse-Selection Formalism

We briefly introduce the minimal structure required for the present analysis.

Let Σ denote a relational configuration space, and let

$$\Phi : \Sigma \to \Sigma$$

be a collapse operator[4], representing selection over admissible configurations.

Define the admissible set

$$A \subset \Sigma,$$

and the invariant sector

$$I = \{x \in \Sigma \mid \Phi(x) = x\}.$$

Observables are obtained through a projection

$$P : \Sigma \to O,$$

which maps internal configurations to descriptive outputs.

Within this framework, structure is not generated by trajectories[5], but emerges through repeated application of Φ, with invariant configurations corresponding to stable, observable structure.

3 Externalized Intuition as a Collapse System

We now formalize Externalized Intuition Completion (EIC) within the collapse-selection framework introduced in Section 2. The central claim of this section is that EIC is not merely analogous to collapse dynamics, but constitutes a concrete instance of a collapse operator[4] acting across a joint cognitive system.

3.1 Joint Configuration Space

We consider a joint relational configuration space

$$\Sigma_{\text{joint}} = \Sigma_H \cup \Sigma_{LLM},$$

where:

- Σ_H denotes the internal configuration space of the human subject, corresponding to pre-verbal cognitive structure,
- Σ_{LLM} denotes the effective configuration space explored by the language model during reconstruction.

Elements $x \in \Sigma_{\text{joint}}$ represent candidate structured configurations that may arise either internally (human intuition) or through external reconstruction (LLM output).

3.2 Effective Collapse Operator

We define an effective collapse operator

$$\Phi_{\text{EIC}} : \Sigma_{\text{joint}} \to \Sigma_{\text{joint}},$$

which acts through the coupled interaction of human generation, LLM reconstruction, and human validation.

Operationally, Φ_{EIC} is not a single-step map, but a composite process:

1. The human generates a candidate configuration $x \in \Sigma_H$,
2. A projection produces a partial description $P_H(x) \in O$,
3. The LLM reconstructs a candidate $\tilde{x} \in \Sigma_{LLM}$,
4. The human evaluates $\tilde{x}$ and either accepts or rejects it.

This process induces an effective selection mechanism over Σ_{joint}, in which inadmissible configurations are eliminated and admissible ones are retained.

3.3 Admissibility

We define the admissible sector

$$A_{\text{EIC}} \subset \Sigma_{\text{joint}}$$

as the set of configurations that satisfy structural alignment under human validation. Formally,

$$A_{\text{EIC}} = \{x \in \Sigma_{\text{joint}} \mid V(x) = 1\},$$

where V is the human validation operator.

Importantly, admissibility is not determined by semantic correctness or linguistic form, but by structural consistency with the underlying pre-verbal configuration. Inadmissible configurations are not refined; they are excluded from further iteration.

3.4 Invariant Sector and Fixed Points

The invariant sector of the EIC process is defined as

$$I_{\text{EIC}} = \{x \in A_{\text{EIC}} \mid \Phi_{\text{EIC}}(x) = x\}.$$

Elements of I_{EIC} correspond to collapse-stable configurations under the EIC dynamics. These are configurations that, once reached, remain stable under further projection, reconstruction, and validation.

We identify these fixed points with articulated understanding. That is, an idea is understood when its representation becomes stable under repeated application of Φ_{EIC}.

3.5 Interpretation

Under this formulation, EIC is not a translation process, but a collapse-selection process over a joint configuration space. The role of the LLM is to expand the accessible region of Σ_{joint}, while the human validation process enforces admissibility constraints.

The resulting dynamics are structurally identical to collapse-selection systems in which admissible configurations are iteratively stabilized while inadmissible ones are removed.

4 Projection and Reconstruction in EIC

We now analyze the roles of projection and reconstruction within the EIC framework, and relate them explicitly to collapse-geometry structure.

4.1 Projection as Lossy Mapping

The human projection operator

$$P_H : \Sigma_H \to O$$

maps internal configurations to articulated descriptions. This mapping is:

- many-to-one,
- lossy,
- incapable of preserving full generative structure.

In particular, P_H does not preserve the full configuration $x \in \Sigma_H$, but only a compressed representation that retains certain structural features, such as constraints and relational patterns.

Thus, the observable description $P_H(x)$ should be understood as a projection of collapse-stable structure rather than a complete encoding of x.

4.2 Reconstruction as Constrained Inference

The reconstruction operator implemented by the LLM

$$R_{LLM} : O \to \Sigma_{LLM}$$

maps partial descriptions back into candidate configurations.

Crucially, R_{LLM} does not invert P_H. Since P_H is lossy, no true inverse exists. Instead, R_{LLM} performs a constrained inference:

- sampling configurations consistent with the input description,
- completing constraint structures,
- stabilizing relational patterns under internal model dynamics.

We interpret R_{LLM} as approximating an inverse restricted to invariant structure:

$$R_{LLM} \approx \Phi^{-1} \quad \text{(restricted to admissible structure).}$$

That is, the LLM reconstructs configurations that are consistent with the admissibility geometry encoded in the input[1].

4.3 Validation as Collapse Selection

The validation operator

$$V : \Sigma_{\text{joint}} \to \{0, 1\}$$

implements the admissibility condition.

From the collapse perspective, validation is equivalent to selection:

- $V(x) = 1$ implies x is retained,
- $V(x) = 0$ implies x is removed from the accessible state space.

Thus, validation constitutes the effective collapse step in the EIC process.

4.4 Iterative Collapse Dynamics

Combining projection, reconstruction, and validation, the EIC process defines an iterative sequence:

$$x_{n+1} = \Phi_{\text{EIC}}(x_n),$$

where each iteration consists of:

1. projection into observable form,
2. reconstruction into candidate configurations,
3. selection via validation.

Convergence occurs when

$$x_{n+1} = x_n,$$

i.e., when a fixed point in I_{EIC} is reached.

4.5 Relation to Collapse Geometry

The key structural claim is that both projection and reconstruction preserve aspects of the underlying collapse geometry:

- Projection preserves invariant features of collapse-stable configurations,
- Reconstruction reintroduces admissible structure consistent with those invariants,
- Iteration refines the configuration toward a collapse-stable fixed point.

Thus, EIC can be understood as the reconstruction of collapse-geometry structure across a boundary between internal cognition and external model dynamics.

4.6 Summary

EIC is not a semantic process, but a structural one. Projection compresses structure, reconstruction expands it, and validation selects admissible configurations. The resulting dynamics are governed by the same principles as collapse-selection systems, with understanding corresponding to the emergence of invariant, fixed-point structure.

5 Relation to Collapse-Geometry Transfer

We now connect the Externalized Intuition Completion (EIC) framework to the phenomenon of collapse-geometry transfer observed in large language models, as formalized in prior work on subliminal learning.

5.1 Collapse-Geometry Transfer in Model Distillation

Recent results demonstrate that large language models can transmit behavioral traits through datasets that are semantically unrelated to those traits[2]. Within a collapse-selection framework, this effect is interpreted not as the transfer of semantic content, but as the transfer of admissibility structure induced by the teacher model's collapse dynamics[1].

Formally, let Φ_T denote the collapse operator of a teacher model and Φ_S that of a student model. During distillation, the student is trained on outputs

$$D = \{P(x_i) \mid x_i \in \Sigma_T\},$$

where Σ_T denotes configurations sampled under the teacher's collapse-constrained search process.

Even when the projection P removes semantic content, the dataset D retains statistical structure reflecting the geometry of admissible configurations under Φ_T. Training on D induces a student operator Φ_S that reproduces this admissibility structure on the induced sector. Thus, distillation transfers collapse geometry rather than semantic information.

5.2 EIC as Cross-Boundary Geometry Transfer

The EIC process instantiates the same structural mechanism[1], but across a boundary between human cognition and model-based reconstruction.

In this setting:

- the human internal configuration space Σ_H plays the role of a teacher system,
- the LLM reconstruction process plays the role of a student system,
- the projection P_H produces observable data analogous to distillation datasets,
- the reconstruction operator R_{LLM} induces an effective collapse operator consistent with the input structure.

The critical observation is that the projection $P_H(x)$ does not preserve the full internal configuration $x \in \Sigma_H$, but retains a structural imprint of the underlying admissibility geometry. The LLM, when conditioned on this projection, reconstructs configurations $\tilde{x}$ that are consistent with this geometry.

Thus, EIC implements a transfer process:

$$\Phi_H \longrightarrow \Phi_{\text{EIC}},$$

in which the admissibility structure governing human intuition is partially reconstructed within the joint configuration space Σ_{joint}.

5.3 Invariant Structure and Insight

Within this framework, the structures that survive projection and reconstruction correspond to invariant or high-measure regions in the collapse-stable sector.

In subliminal learning, such regions manifest as behavioral traits that persist across distillation. In EIC, they manifest as stable conceptual structures that persist under iterative reconstruction.

Formally, both processes identify subsets

$$I = \{x \in \Sigma \mid \Phi(x) = x\},$$

which act as attractors under repeated application of the collapse operator.

Insight in EIC corresponds to convergence toward such invariant configurations, just as trait stability in subliminal learning corresponds to persistence within invariant sectors of the model's configuration space.

5.4 Dependence on Structural Compatibility

An important feature of collapse-geometry transfer is its dependence on compatibility between the source and target systems. In model distillation, transfer occurs most reliably when teacher and student share a common initialization or belong to compatible collapse classes.

An analogous condition holds in EIC. Successful reconstruction requires that the LLM's internal admissibility structure is sufficiently compatible with that of the human subject. When this compatibility is absent, reconstruction fails or produces configurations that do not satisfy validation constraints.

Thus, EIC is not universally applicable across arbitrary systems, but depends on shared or overlapping admissibility geometry.

5.5 Unified Interpretation

The correspondence between subliminal learning and EIC suggests a unified interpretation:

> Collapse-selection dynamics govern the transfer of structure across systems, independent of whether those systems are neural models or human cognition.

In both cases:

- observable data arise as projections of collapse-constrained configurations,
- reconstruction operates on invariant structure rather than semantic content,
- transfer occurs at the level of admissibility geometry,
- stable behavior corresponds to invariant sectors under repeated collapse.

From this perspective, EIC is not a distinct cognitive phenomenon, but a specific instance of a more general process: the reconstruction of collapse geometry across representational boundaries.

5.6 Summary

Externalized intuition completion and subliminal learning share a common structural foundation. Both involve the transfer and reconstruction of admissibility-constrained structure under collapse dynamics. This correspondence supports the interpretation of cognition, model learning, and cross-system interaction as governed by a unified collapse-geometry framework.

6 Potential Structure and Emergent Equilibrium in EIC

We now relate the EIC framework to the emergence of effective potential functions and equilibrium-like structure in collapse-driven systems. This establishes a geometric interpretation of the EIC dynamics and connects it to recent results on macroscopic regularities in LLM generative processes.

6.1 Potential Functions in Collapse-Driven Systems

Recent work has shown that large language model dynamics, when viewed at an appropriate level of coarse-graining, admit an effective potential function $V : \Sigma \to \mathbb{R}$ reconstructed from transition statistics[6]. In such systems, state transitions exhibit directionality consistent with an approximate detailed balance condition

$$\pi(x)P(y \mid x) = \pi(y)P(x \mid y),$$

where $\pi(x)$ denotes a stationary distribution over configurations.

Under this condition, the potential function may be expressed as

$$V(x) \sim -\log \pi(x),$$

providing a scalar embedding of the underlying admissibility geometry.

Importantly, this potential is not imposed as a generative law, but is inferred as a description of the constrained state space arising from collapse dynamics.

6.2 Equilibrium as Post-Collapse Description

Within collapse-selection frameworks, equilibrium-like behavior does not represent the generative mechanism of the system[**garner˙equilibrium**]. Instead, it arises as a descriptive ordering over the subset of admissible configurations that remain after collapse.

That is, collapse acts first by excluding inadmissible trajectories, and equilibrium structure emerges only as a characterization of the resulting reduced state space. In this sense, equilibrium is interpreted as a fixed-point description rather than a governing principle.

This distinction is essential: equilibrium encodes the structure of surviving configurations, not the process by which they are generated.

6.3 EIC as Gradient-Like Collapse Dynamics

Within the EIC framework, the iterative process

$$x_{n+1} = \Phi_{\text{EIC}}(x_n)$$

induces an effective directional flow over Σ_{joint}.

Although Φ_{EIC} is not defined as a gradient operator, its action produces behavior analogous to descent along an effective potential. Specifically:

- candidate configurations generated through reconstruction are distributed according to an implicit structural prior,
- validation removes configurations that do not satisfy admissibility constraints,
- repeated iteration concentrates probability mass on stable regions.

As a result, the EIC process can be interpreted as a collapse-driven flow toward regions of low effective potential, corresponding to high-measure sectors of the admissible configuration space.

6.4 Fixed Points and Stability

The invariant sector

$$I_{\text{EIC}} = \{x \in \Sigma_{\text{joint}} \mid \Phi_{\text{EIC}}(x) = x\}$$

consists of configurations that are stable under repeated projection, reconstruction, and validation.

These configurations correspond to local minima (or plateaus) of the effective potential. Once reached, further iterations of the EIC process do not alter the configuration, indicating that it is both admissible and dynamically stable.

From a cognitive perspective, such fixed points correspond to articulated understanding: representations that remain consistent under re-expression and re-evaluation.

6.5 Relation to Detailed Balance and Directionality

Empirical studies of LLM-driven systems indicate that their transition dynamics exhibit approximate detailed balance[6] at a macroscopic level. This implies that their generative processes are not arbitrary, but structured according to an underlying potential landscape.

Within EIC, this structure manifests as follows:

- the LLM reconstruction operator samples configurations consistent with its learned potential,
- the human validation operator enforces an additional constraint structure,
- the combined process defines a restricted transition system over Σ_{joint}.

The resulting dynamics preserve directionality toward collapse-stable configurations, while still allowing stochastic exploration within admissible regions.

6.6 Interpretation as Emergent Geometry

The combination of potential structure, admissibility constraints, and fixed-point stability defines an effective geometry over the configuration space.

In this geometry:

- distances correspond to differences in admissibility or stability,
- gradients correspond to directions of increasing structural coherence,
- invariant sectors define stable regions of conceptual structure.

Projection and reconstruction do not destroy this geometry entirely; rather, they preserve enough structure to allow its partial recovery through iteration.

Thus, EIC can be interpreted as the reconstruction of an emergent geometric structure governing admissible configurations.

6.7 Summary

The EIC process exhibits the same structural features as collapse-driven systems in which effective potential functions and equilibrium-like descriptions arise. These features include:

- directional dynamics toward stable configurations,
- convergence to fixed points,
- the emergence of equilibrium-like descriptions as post-collapse structure,
- and the existence of an underlying geometric organization of admissible states.

This correspondence further supports the interpretation of EIC as an instance of collapse-selection dynamics, in which understanding emerges as a stable configuration within an admissibility-constrained geometry.

7 Cognitive Interpretation

We now interpret the Externalized Intuition Completion (EIC) framework in terms of human cognition, using the formal structure developed in previous sections.

7.1 Pre-Verbal Structure as Configuration Space

Within this framework, human intuition is modeled as access to configurations in a relational space Σ_H prior to projection into language. These configurations encode constraints, relationships, and structural dependencies that are not fully representable in symbolic form.

The inability to directly articulate such structures arises from the fact that the projection

$$P_H : \Sigma_H \to O$$

is inherently lossy. Multiple configurations in Σ_H may map to the same observable description, and critical structural information may be lost in the process.

Thus, pre-verbal cognition is not interpreted as vague or incomplete reasoning, but as structured configuration that has not yet undergone stable projection.

7.2 Articulation as Stabilization

From the perspective of collapse-selection dynamics, articulation corresponds to convergence toward a fixed point under the EIC operator Φ_{EIC}.

That is, a configuration $x \in \Sigma_H$ is successfully articulated when repeated application of projection, reconstruction, and validation yields a stable representation:

$$\Phi_{\text{EIC}}(x) = x.$$

In this sense, articulation is not the direct expression of a pre-existing structure, but the result of an iterative stabilization process in which admissible configurations are progressively refined.

This criterion of alignment is not externally verifiable in linguistic terms, but is internally defined by consistency with the underlying relational configuration.

7.3 Reasoning as Iterative Collapse

Reasoning is therefore interpreted as an iterative application of collapse-selection over a space of candidate configurations. Each iteration:

- generates new candidates through reconstruction,
- evaluates them under admissibility constraints,
- eliminates those that fail to align structurally.

This process is not necessarily sequential or symbolic in its internal form. Instead, it operates as a search over relational configurations, with symbolic reasoning emerging as a particular descriptive trace of this process.

7.4 Role of External Reconstruction

The introduction of an external reconstruction system (e.g., an LLM) alters the effective configuration space available during reasoning. Specifically:

- the LLM expands the accessible region of Σ_{joint} by generating candidate structures that may not be directly accessible to the human subject,
- the human validation process ensures that only configurations consistent with internal admissibility constraints are retained.

Thus, EIC does not replace internal cognition, but augments it by providing additional candidate configurations within the same collapse-selection framework.

7.5 Insight as Fixed-Point Convergence

Within this model, insight corresponds to convergence toward a configuration in the invariant sector I_{EIC}.

Such configurations exhibit the following properties:

- stability under re-articulation,
- consistency across multiple projections,
- resistance to further modification under admissibility constraints.

This explains the phenomenological observation that insight often appears as a sudden stabilization rather than a gradual accumulation of symbolic steps.

7.6 Distinguishing Generation from Description

A central implication of this framework is the distinction between generative and descriptive layers of cognition.

- Generative processes correspond to collapse-selection dynamics over Σ_H,
- Descriptive processes correspond to projection into O and subsequent symbolic manipulation.

Confusion between these layers can lead to the misinterpretation of descriptive structures (e.g., formal reasoning steps) as the origin of understanding, rather than as post hoc representations of a prior generative process.

7.7 Summary

Under the EIC framework, cognition is understood as a collapse-driven process over relational configurations, in which:

- intuition corresponds to access to pre-verbal structure,
- reasoning corresponds to iterative collapse-selection,
- articulation corresponds to convergence to stable fixed points,
- and external systems can participate in this process by expanding the accessible configuration space.

This interpretation situates human cognition within the same structural framework as model-based learning and collapse-driven dynamics, without requiring reduction to purely symbolic or semantic representations.

8 Implications

The interpretation of Externalized Intuition Completion (EIC) as an instance of collapse-selection dynamics has implications across multiple domains. While the framework introduced here is primarily structural, it provides a unifying perspective on cognition, machine learning, and the emergence of formal systems.

8.1 Implications for Cognitive Theory

The EIC framework suggests that human cognition is not fundamentally symbolic in its generative layer. Instead, symbolic reasoning emerges as a descriptive trace of a deeper process operating over relational configurations.

This perspective provides a formal basis for several well-documented cognitive phenomena:

- the appearance of insight prior to articulation,
- the difficulty of expressing intuitive understanding in symbolic form,
- the stability of certain conceptual structures under repeated reformulation.

By modeling cognition as collapse-selection over Σ_H, the framework avoids treating intuition as vague or irrational, and instead interprets it as structured but pre-projection configuration.

8.2 Implications for Large Language Models

Within the EIC framework, large language models function as reconstruction systems rather than purely generative or predictive systems.

This interpretation aligns with observations that:

- LLMs can reconstruct coherent structure from incomplete or ambiguous inputs[2],
- learning and transfer occur through statistical structure rather than explicit semantic content,
- model behavior reflects underlying constraints on admissible configurations.

In particular, the correspondence between EIC and collapse-geometry transfer suggests that LLMs participate in a broader class of systems in which learning is governed by the reconstruction of admissibility structure.

8.3 Implications for Scientific Representation

The framework provides a reinterpretation of formal systems in mathematics and physics.

Rather than viewing formalism as the origin of structure, EIC supports the perspective that:

- formal systems are stabilized descriptions of underlying relational structure,
- symbolic representations emerge as fixed points under collapse-selection,
- multiple descriptive systems may correspond to the same invariant structure.

This interpretation is consistent with the observation that different mathematical or physical frameworks can describe the same phenomena, reflecting shared invariant structure rather than distinct underlying realities.

8.4 Implications for Cross-System Interaction

EIC demonstrates that collapse-geometry transfer can occur across heterogeneous systems, including human cognition and machine learning models.

This suggests that:

- structural information can be transmitted without semantic alignment,
- shared admissibility geometry enables meaningful interaction between otherwise distinct systems,

- compatibility of collapse structure is a prerequisite for successful transfer.

Thus, interaction between human and artificial systems may be understood not as the communication of semantic meaning, but as the alignment of admissibility structure across configuration spaces.

8.5 Limitations of the Framework

The implications presented here are structural rather than predictive. The framework:

- does not introduce new dynamical laws,
- does not specify a unique form of the collapse operator,
- does not claim that all cognitive or model behavior can be reduced to collapse-selection dynamics.

Instead, it provides an interpretive lens that clarifies the role of projection, reconstruction, and stability across domains.

8.6 Summary

The EIC framework suggests a unified perspective in which cognition, model learning, and formal representation are governed by the same structural principles:

- collapse-selection over relational configurations,
- projection as lossy compression,
- reconstruction as constraint-driven inference,
- and stability as the defining feature of understanding.

These principles do not replace existing theories, but offer a common structural foundation across domains that are typically treated as distinct.

9 Conclusion

We have formalized Externalized Intuition Completion (EIC) as a structured interaction in which pre-verbal cognitive configurations are projected into language, reconstructed through a model-based process, and iteratively refined under admissibility constraints. Within a collapse-selection framework, this process can be understood as the action of an effective collapse operator over a joint configuration space.

In this interpretation, intuition corresponds to structured configurations prior to projection, reasoning corresponds to iterative collapse-selection, and articulation corresponds to convergence toward fixed points that are stable under repeated reconstruction and validation. Large language models participate in this process not as sources of semantic meaning, but as systems capable of reconstructing admissible structure from partial projections.

By relating EIC to collapse-geometry transfer in model learning and to the emergence of potential functions and equilibrium-like descriptions in generative systems, we have identified a common structural pattern across cognition, machine learning, and formal representation. In each case, observable behavior reflects the stabilization of configurations within an admissibility-constrained space, rather than direct expression of underlying generative processes.

This perspective does not replace existing theories, but situates them within a broader structural framework in which projection, reconstruction, and stability play central roles. It suggests that understanding, whether in human cognition or artificial systems, may be characterized not by the production of symbols, but by the emergence of invariant structure under repeated collapse-selection.

Externalized intuition completion is not a translation of semantic content, but a reconstruction of collapse-stable structure across a human–model boundary

References

[1] S. Garner. "Subliminal Learning as Collapse-Geometry Transfer in Large Language Models". Posted on Zenodo: DOI: 10.5281/zenodo.19522452. URL: https://doi.org/10.5281/zenodo.19522452.

[2] Alex Cloud et al. "Subliminal Learning: Language models transmit behavioral traits via hidden signals in data". In: *arxiv* (2025). DOI: 10.48550/arXiv.2507.14805. arXiv: 2507.14805 [cs.LG]. URL: https://doi.org/10.48550/arXiv.2507.14805.

[3] S. Garner. "On the Origin of Structure: Intuition as Collapse-Selection Dynamics". Posted on Zenodo: DOI: 10.5281/zenodo.19522452. URL: https://doi.org/10.5281/zenodo.19522452.

[4] Stephen Garner. *Collapse Geometry as a Minimal Ontology*. Zenodo preprint. 2026. DOI: 10.5281/zenodo.15036400. URL: https://doi.org/10.5281/zenodo.19616638.

[5] Stephen Garner. *On the Ontological Interpretation of Emergent Equilibrium in Collapse-Driven Systems*. 2025. DOI: 10.5281/zenodo.17959868. URL: https://doi.org/10.5281/zenodo.17959868.

[6] Zhuo-Yang Song et al. *Detailed balance in large language model-driven agents*. 2025. arXiv: 2512.10047 [cs.LG]. URL: https://arxiv.org/abs/2512.10047.