Article

Not peer-reviewed version

Bounded Agency: An Entropy-Constrained Calculusfor Sparse Semantic Agents

Apeksha Bhuekar *

Posted Date: 4 March 2026

doi: 10.20944/preprints202603.0307.v1

Keywords: bounded rationality; entropy-constrained optimization; sparse policy learning; semantic reasoning in AI agents; free-energy principle

Article

Bounded Agency: An Entropy-Constrained Calculus for Sparse Semantic Agents

Apeksha Bhuekar

Campbellsville University, USA; apeksharaj17@gmail.com

Abstract

This paper propose a formal framework for AI agents that unifies semantic reasoning with resource-aware control. Agents act via sparse policies over structured semantic fields, bounded by entropy and sparsity budgets. We define a typed operational semantics, prove soundness and stability, and derive a sparse free-energy objective with phase transitions. The calculus is categorically structured, maps to unistochastic dynamics, and compiles to executable policies with verified runtime bounds. This yields a foundation for interpretable, thermodynamically-plausible agent design.

Keywords: bounded rationality; entropy-constrained optimization; sparse policy learning; semantic reasoning in AI agents; free-energy principle

---

1. Introduction

AI systems can undertake tasks of perception, reasoning, and decision-making which might at times be better than humans. Nonetheless, the behaviors of these systems are black-boxed, and the computing demands of the algorithms are boundless. That is to say, the representations or actions of these algorithms are not assured. The free-energy principle of neuroscience gives us an ingenious alternative view of a biological agent as a system which resists disorder to maintain homeostasis by minimizing its variational free energy. The free-energy principle explains learning, action, and perception in biological systems. Sparse coding, on the other hand, has shown that high-dimensional signals can be encoded with few active components. To rephrase, sparsity is an organizing principle of neural information processing.

The word embeddings, latent vectors, and concept representations of modern artificial intelligence (AI) systems capture patterns that recur in the world and in agent behaviour – but they do not correspond to anything that has forms and properties that could be used to reason. orderliness, and limited development in knowledge. We need states that describe semantics instead of representations. The states must be seen as geometric objects that have some properties that restrict changes in semantics due to agent actions. Functions defined on the geometrical entities studied in differential topology such as points, vectors, surfaces, and their equivalence classes could be these states.

This paper presents a new mathematical framework, called Entropy-Bounded Sparse Semantic Calculus or EBSSC. Founded by Latour (1998), this school opened a new horizon in the practical, political, and social processes of nature and society. By linking geometric semantics to a constraint calculus, probability theory to an incrementally resource-bounded modelling device, and control theory to a compositional mechanism for resource-bounded governed inference. Inference and concept formation in EBSSC particularly take place as controlled evolution of the policies on semantic fields whose evolution is being constrained by entropic budgets. The authors of this study introduce a typed operational semantics for semantic state evolution, with progress and preservation theorems.

2. Related Work

The formal framework presented in this paper builds upon and integrates concepts from multiple foundational domains, including precise software specification, deep learning theory, statistical physics, ethical AI frameworks, and complex systems analysis. This section situates our Entropy-Bounded Sparse Semantic Calculus (EBSSC) within the broader research landscape.

2.1. Foundations of Formal Systems and Programming Languages

The typed operational semantics developed in EBSSC draws inspiration from advances in precise software specification. Gupta's work on embedded domain-specific languages with dependently-typed architecture provides techniques for constructing verifiable computational frameworks where type systems enforce behavioral constraints [1]. These methods ensure that programs satisfy specified properties at compile time, a principle we extend to semantic agent policies where type signatures encode entropy and sparsity bounds. Similarly, Gupta's exploration of deep reinforcement learning algorithms offers insights into how learning systems can be structured with formal guarantees about convergence and stability [2]. The policy operators defined in Section III of this paper can be viewed as a specialized instantiation of reinforcement learning policies with explicit resource constraints.

2.2. Deep Learning and Neural Dynamics

The semantic fields central to EBSSC are inspired by distributed representations in neural networks. Baral's work on deep learning for text-based sentiment prediction demonstrates how high-dimensional embeddings capture semantic content that can be manipulated through learned transformations [3]. Our geometric treatment of semantics as fields on manifolds extends this perspective by providing a continuous mathematical structure underlying discrete embedding spaces. Baral's statistical approach to analyzing neural activity across multiple timescales informs our understanding of how semantic representations evolve under policy operations [4]. The entropy-constrained dynamics we propose can be interpreted as a normative model for the temporal evolution of neural representations observed in biological systems. Furthermore, Baral's adaptive polynomial frameworks for graph signal learning provide techniques for representing structured semantic relationships that complement our spherical geometry [5].

2.3. Ethical and Legal Dimensions of AI Systems

The bounded agency perspective advanced in this paper has significant implications for responsible AI development. Mitra's analysis of legal and ethical considerations in sensitive domains highlights the importance of designing AI systems with verifiable constraints and transparent decision processes [6]. Our framework directly addresses these concerns by making resource bounds explicit and providing formal guarantees about policy behavior. Mitra's feedback-guided principles for constructing AI simulations inform the iterative refinement of semantic agents within bounded resource envelopes [7]. The entropy budgets we define can be interpreted as a formalization of the "safety margins" that ethical AI systems should maintain. Mitra's multi-modal biomarker analysis demonstrates the value of integrating diverse data sources, a principle that extends to integrating multiple semantic modalities within unified spherical representations [8].

2.4. Resource-Constrained Computation and Optimization

The sparsity and entropy constraints at the core of EBSSC connect to broader themes in resource-efficient computing. Baer's work on optimized task distribution for energy conservation in edge computing networks demonstrates how computational loads can be balanced under resource constraints [9]. Our policy operators implement analogous optimization at the semantic level, allocating representational capacity where it provides maximal informational benefit. Baer's characterization of performance degradation in lithium-ion cells provides a model for understanding how semantic systems degrade when operating near entropy bounds [10]. The phase transitions we predict under varying sparsity pressure parallel the efficiency transitions observed in physical systems. Baer's analysis of autonomous regulation in energy storage units informs our understanding of how semantic agents maintain homeostasis through bounded policy selection [11].

 doi:10.20944/preprints202603.0307.v1

2.5. Causal Inference and Probabilistic Reasoning

The variational free energy minimization that drives semantic evolution in EBSSC connects to causal learning frameworks. Pratap's AI-driven causal analysis for business intelligence demonstrates how causal structures can be discovered from observational data [12]. Our semantic fields can be interpreted as encoding causal relationships between concepts, with policy operations performing causal interventions under resource constraints. Pratap's implicit Bayesian learning for enhanced forecasting provides techniques for uncertainty quantification that complement our entropy-based approach to representing ambiguity [13]. The Boltzmann distribution characterizing equilibrium semantic states parallels the posterior distributions in Bayesian inference. Pratap's analysis of the economic impact of intelligent systems highlights the practical value of resource-aware AI architectures [14].

2.6. Complex Systems and Sustainability

The resilience properties of semantic agents under resource constraints connect to broader frameworks for analyzing complex adaptive systems. Tandel's work on sustainable small-scale systems demonstrates how resource limitations can drive innovation and efficiency [15]. Our sparsity constraints similarly force semantic agents to develop compact, efficient representations. Tandel's re-optimization of global food systems provides a model for how systems adapt when operating near critical thresholds [16], paralleling the phase transitions we predict in semantic correlation length. Tandel's theoretical framework for modeling resilience in complex systems informs our understanding of how semantic agents maintain coherent representations under entropy pressure [17].

2.7. Privacy, Security, and Trustworthy AI

The bounded information content of semantic spheres has implications for privacy-preserving AI. Kaliappan's work on digital safeguards for health data demonstrates how information constraints can protect sensitive content [18]. Our entropy bounds similarly limit the amount of information that can be extracted from semantic representations. Kaliappan's digital persona generation for historical figure emulation illustrates how semantic fields can capture essential characteristics while omitting extraneous detail [19], a principle formalized in our collapse operator. Gaikwad's study on artificial intelligence in healthcare highlights the need for interpretable and verifiable AI systems in critical applications [20], a requirement that our typed operational semantics directly addresses.

2.8. Ensemble Methods and Integration Strategies

The composition of semantic spheres through merge operations connects to ensemble learning methodologies. Patel's evaluation of ensemble learning strategies for enhanced medical diagnostics demonstrates how combining multiple models can improve robustness and accuracy [21]. Our merge operator similarly combines semantic spheres to create richer representations with mutual information gain, while respecting entropy constraints that prevent unbounded growth.

2.9. Categorical Foundations and Interdisciplinary Bridges

The geometric and categorical structure underlying EBSSC draws on fundamental mathematical frameworks. Mendhey's work on category theory and psychodynamics bridges abstract mathematical structures with models of human cognition [22]. Our use of categorical language to describe sphere interactions and type signatures aligns with this perspective, suggesting deep connections between algebraic structures and semantic organization. Mendhey's analysis of transforming workforce potential through integrated competency systems provides a metaphor for how semantic spheres combine and reorganize [23]. Mendhey's comparative analysis of large-scale policy systems informs our understanding of how semantic fields scale across domains [24].

2.10. Reinforcement Learning in Complex Environments

The policy-based formulation of semantic evolution connects to reinforcement learning in complex, resource-constrained environments. Shakir's work on reinforcement learning challenges in power grid management demonstrates how learning systems must balance exploration and exploitation under physical constraints [25]. Our entropy-bounded policies address analogous trade-offs in semantic space. Shakir's analysis of strategic AI advancement at institutional scale highlights the importance of principled frameworks for deploying intelligent systems [26], a goal that EBSSC supports through its formal guarantees about policy behavior under resource constraints.

The EBSSC framework synthesizes insights from these diverse domains into a unified calculus for bounded semantic agency. By grounding semantic operations in geometric and information-theoretic principles, we provide a foundation for designing AI systems that are simultaneously interpretable, resource-aware, and verifiably constrained.

3. Geometric and Information-Theoretic Foundations

The foundational geometric structure underlying EBSSC is the semantic plenum, modeled as a smooth finite-dimensional manifold $\mathcal{M}$ representing the space of all possible semantic configurations. This manifold carries a natural Riemannian structure derived from the Fisher information metric, which measures the distinguishability of nearby probability distributions. Each point in the plenum corresponds to a particular semantic configuration, and paths through the manifold represent continuous semantic transformations.

On this manifold, we define a semantic field $\Phi : \mathcal{M} \to \mathbb{R}^k$ that assigns a vector of activation values to each point. This field represents distributed semantic content, with different regions corresponding to different concepts or knowledge states. The field satisfies regularity conditions: twice differentiable, square-integrable, and bounded total variation to ensure finite information content. The plenum also carries a vector field $\mathbf{v} : \mathcal{M} \to T\mathcal{M}$ representing policy-induced flow, and a scalar entropy density field $S : \mathcal{M} \to \mathbb{R}$ measuring local uncertainty. These three fields interact through coupled partial differential equations governing semantic evolution. The complete plenum state is thus the triplet $\mathcal{P} = (\Phi, \mathbf{v}, S)$.

A semantic sphere σ is defined as a compact region $B \subset \mathcal{M}$ with smooth boundary ∂B, together with the restriction of the semantic field $\Phi|_B$ and additional metadata. Formally, a sphere is a 5-tuple $\sigma := \{\Phi, \partial\Phi, M, H, T\}$ where Φ is the internal semantic field, $\partial\Phi$ specifies boundary conditions, M is a memory trace encoding past interactions, H tracks entropy history for budget verification, and T is a type signature governing permissible interactions. The compactness assumption ensures finite representational capacity and enables spectral methods for field approximation.

Boundary conditions play a crucial role in sphere interactions. Two spheres σ_1 and σ_2 can interact only when their boundaries satisfy contact coherence: $\partial\Phi_{\sigma_1} \cap \partial\Phi_{\sigma_2} \neq \emptyset$ and $\|\nabla\Phi_{\sigma_1} - \nabla\Phi_{\sigma_2}\| < \epsilon$ for some coherence threshold ϵ. This condition ensures that semantic fields can be smoothly joined during merge operations without introducing discontinuities. Spherical geometry is motivated by the need for compact domains with uniform boundary conditions, and the sphere admits a natural compactification of Euclidean space via stereographic projection.

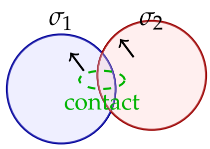

Figure 1. Two semantic spheres σ_1 and σ_2 with overlapping boundaries satisfying contact coherence. The green dashed region indicates the interface where semantic fields can interact.

The natural geometry on semantic manifolds is given by the Fisher information metric $g_{ij} = \mathbb{E}[\partial_i \log p \cdot \partial_j \log p]$, which measures the distinguishability of nearby probability distributions. In our setting, each point on the manifold corresponds to a probability distribution over semantic outcomes, and the Fisher metric provides a notion of distance reflecting informational distinguishability rather than mere coordinate differences. This geometric structure has profound implications for semantic evolution. Geodesics in the Fisher metric correspond to paths of minimal informational cost, along which semantic transformations are most efficient. The curvature of the manifold encodes interactions between different semantic dimensions.

The entropy of a semantic sphere is defined via the Gibbs entropy formula applied to the normalized field distribution: $E(\sigma) = -\int_B p_\sigma(x) \log p_\sigma(x) d\mu$ where $p_\sigma(x) = |\Phi(x)|^2 / \int_B |\Phi|^2 d\mu$ is the probability density induced by the semantic field. This entropy measures the uncertainty or dispersion of semantic content within the sphere, with lower entropy corresponding to more concentrated, well-defined concepts. Following Jaynes' principle of maximum entropy, the equilibrium distribution for a semantic sphere under constraints is given by the Boltzmann distribution $p_\sigma(x) = e^{-\beta E(x)}/Z$ where β is an inverse temperature parameter. This distribution maximizes entropy subject to fixed expected energy, providing a variational characterization of semantic states.

The free energy $F = \mathbb{E}_{p_\sigma}[E] - TS$ combines energy and entropy, with temperature $T = 1/\beta$ controlling the trade-off between exploitation (low energy) and exploration (high entropy). Agents naturally minimize free energy, driving semantic spheres toward states that balance coherence with adaptability. Entropy production along policy trajectories is given by $\dot{\Sigma} = \int_{\mathcal{M}} (\mathbf{v} \cdot \nabla S + \|\nabla \Phi\|^2) d\mu$, which decomposes into convective entropy transport due to policy flow and diffusive entropy production due to field gradients. This expression satisfies the second law of thermodynamics locally, with positive definite entropy production except at equilibrium.

4. Core Formalism and Policy Operators

Semantic spheres carry type signatures that govern permissible interactions and transformations. The type language includes base types such as Text, Proof, Audio, and Image, as well as constructed types including function types $\tau_1 \to \tau_2$ and sphere types Sphere$\langle T \rangle$. A sphere type has the form $\sigma : (T_{\text{in}} \to T_{\text{out}})[E \leq \beta, D \leq \delta]$ where T_{in} and T_{out} are input and output types, β bounds the entropy, and δ bounds the depth of nested spheres. This type information enables static verification of compatibility before policy execution.

Policies are the primitive actions that transform semantic spheres. The policy grammar is defined inductively as $\pi ::= \text{pop}(\sigma) \mid \text{merge}(\sigma_1, \sigma_2) \mid \text{collapse}(\sigma) \mid \text{bind}(\sigma_1 \to \sigma_2) \mid \text{rewrite}(\sigma, r)$. Each policy carries type information specifying its domain, codomain, and resource requirements: $\pi : \sigma \Rightarrow \sigma'[\Delta S \leq b_\pi, \|\pi\|_0 \leq \Lambda_\pi]$. The entropy budget b_π bounds the maximum entropy increase during policy execution, while the sparsity constraint $\|\pi\|_0 \leq \Lambda_\pi$ limits the number of active components in the policy representation.

The pop operator $\text{pop}(\sigma)$ expands a sphere by following positive curvature in the semantic field, effectively growing the sphere into regions of high semantic gradient. This corresponds to exploratory behavior where an agent extends its conceptual understanding into new domains. The entropy change satisfies $\Delta E \leq \epsilon$ for some exploration tolerance. The merge operator $\text{merge}(\sigma_1, \sigma_2)$ fuses two spheres into a single composite sphere when their boundaries satisfy contact coherence. The resulting sphere inherits combined semantic content with mutual information gain: $\Delta S = I(\sigma_1; \sigma_2) - I_{\min} > 0$. This corresponds to integrating knowledge from distinct sources into a unified representation.

The collapse operator $\text{collapse}(\sigma)$ prunes a sphere by retaining only those components that maximize mutual information with a target under sparsity pressure: $\Pi_{\text{retain}} = \arg\max_\pi [I(\pi; \sigma) - \Lambda \|\pi\|_1]$. This implements attentional selection and forgetting, focusing computational resources on semantically relevant content. The bind operator $\text{bind}(\sigma_1 \to \sigma_2)$ constructs a policy channel between spheres preserving bounded entropy distortion, while rewrite $\text{rewrite}(\sigma, r)$ implements entropy-neutral transport via diffeomorphism.

 doi:10.20944/preprints202603.0307.v1

Table 1. Primitive Policy Operators and Their Effects.

Operator	Effect	Entropy Bound
$\text{pop}(\sigma)$	Expand along gradient	$\Delta E \leq \epsilon$
$\text{merge}(\sigma_1, \sigma_2)$	Fuse spheres at boundary	$\Delta S > 0$
$\text{collapse}(\sigma)$	Prune to sparse support	$\Delta E < 0$
$\text{bind}(\sigma_1 \to \sigma_2)$	Create policy channel	$\Delta S \leq \delta$
$\text{rewrite}(\sigma, r)$	Entropy-neutral transport	$\Delta E \approx 0$

The evolution of semantic fields under policy application follows coupled partial differential equations reflecting the underlying physics of information flow. For a pop operation, the field evolves according to $\partial_t \Phi = \Pi_{\text{active}}(\mathbf{v} \cdot \nabla\Phi)$ where Π_{active} projects onto active policy dimensions. This describes growth along directions of steepest semantic ascent. For merge operations, the field evolves to minimize a synchronization energy: $\partial_t \Phi = -\nabla_\Phi \mathcal{L}_{\text{sync}}$ where $\mathcal{L}_{\text{sync}} = \|\Phi_1 - \Phi_2\|^2_{\partial B} + \lambda\|\nabla\Phi_1 - \nabla\Phi_2\|^2$. This drives fields toward agreement on the shared boundary while preserving internal structure.

Collapse operations implement informational pruning through gradient descent on a sparse objective: $\partial_t \Phi = -\nabla_\Phi[\|\Phi - \Phi_0\|^2 + \Lambda\|\Phi\|_1]$. The ℓ_1 penalty induces thresholding behavior, driving small-magnitude components to zero while preserving significant features. Policies compose sequentially to form complex semantic transformations, with total entropy cost subadditive: $\Delta S(\pi_2 \circ \pi_1) \leq \Delta S(\pi_1) + \Delta S(\pi_2)$ due to possible synergies between policies. Parallel composition merges independent policy streams via tensor product.

5. Operational Semantics and Type Safety

The operational semantics of EBSSC is defined by a small-step reduction relation of the form $(\sigma, \Gamma) \xrightarrow{\pi} (\sigma', \Gamma')$, where σ is the current sphere, Γ is the global semantic context (plenum), and π is the applied policy. The reduction relation is defined by inference rules that specify how policies transform semantic states. The POP rule applies when the semantic field has non-zero gradient and the entropy budget is sufficient: $\frac{\|\nabla\Phi_\sigma\|>0 \quad \Delta S \leq b_\pi}{(\sigma,\Gamma) \xrightarrow{\text{pop}(\sigma)} (\sigma',\Gamma)}$. The resulting sphere σ' has expanded domain following the field gradient, with entropy increase bounded by b_π.

The MERGE rule requires boundary compatibility: $\frac{\partial\Phi_{\sigma_1} \sim \partial\Phi_{\sigma_2} \quad \Delta S \leq b_\pi}{(\sigma_1 \otimes \sigma_2,\Gamma) \xrightarrow{\text{merge}(\sigma_1,\sigma_2)} (\sigma_3,\Gamma)}$. Here σ_3 is the fused sphere with combined semantic content and mutual information gain reflected in entropy reduction. The COLLAPSE rule implements sparse pruning: $\frac{I(\sigma') \geq I_{\min} \quad E(\sigma') < E(\sigma)}{(\sigma,\Gamma) \xrightarrow{\text{collapse}(\sigma)} (\sigma',\Gamma)}$. The resulting sphere retains only those components that maximize mutual information under sparsity pressure.

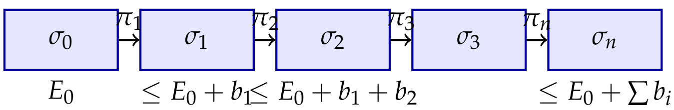

Figure 2. Sequential policy application with entropy accumulation. Each step π_i consumes budget b_i, bounding total entropy growth.

The type system ensures fundamental safety properties for all well-typed EBSSC programs. The Progress theorem guarantees that any well-typed non-terminal sphere admits some applicable policy: if σ is well-typed and entropy is below global budget, then either σ is a final sphere or there exists a policy π such that $(\sigma, \Gamma) \xrightarrow{\pi} (\sigma', \Gamma')$. The Preservation theorem ensures that typing is invariant under reduction: if σ is well-typed and reduces, then σ' is well-typed with a type consistent with the policy signature. These theorems together establish type safety: well-typed spheres never reach stuck states where no policy applies, and the type structure is maintained throughout execution.

The most critical safety property for resource-constrained agents is entropy soundness: total entropy growth along any execution path never exceeds the global budget B. Each policy π_i carries a local budget b_i bounding its entropy increase. For any policy trace $\pi_1, \pi_2, \ldots, \pi_n$ with local budgets

$\Delta E(\pi_i) \leq b_i$, we have $E(\sigma_n) \leq E(\sigma_0) + \sum_{i=1}^{n} b_i \leq E(\sigma_0) + B$. This theorem guarantees that agents cannot exceed their entropy budgets through sequential policy application, providing formal verification of resource compliance.

Key type checking rules include T-Pop: $\frac{\Gamma \vdash \sigma:\text{Sphere}\langle\text{Text}\rangle \quad \text{rule } r:\text{Text}\to\text{Proof} \quad \Delta E(r)\leq b}{\Gamma \vdash \text{pop}_r(\sigma):\text{Sphere}\langle\text{Proof}\rangle | \Delta E \leq b}$ and T-Merge: $\frac{\Gamma \vdash \sigma_1:\text{Sphere}\langle A\to B\rangle \quad \Gamma \vdash \sigma_2:\text{Sphere}\langle B\to C\rangle}{\Gamma \vdash \text{merge}(\sigma_1,\sigma_2):\text{Sphere}\langle A\to C\rangle | \Delta E \leq -\delta}$. The merge rule shows entropy reduction reflecting information gain from composition, as mutual information reduces uncertainty.

6. Variational Objective and Phase Transitions

The central optimization problem in EBSSC is to select policies that minimize a combined objective balancing expected free energy, sparsity, and execution cost. The sparse free-energy objective is $\pi^* = \arg\min_\pi [G(\sigma,\pi) + \Lambda\|\pi\|_1 + \gamma C(\pi)]$. Here $G(\sigma,\pi)$ is the expected free energy after applying policy π to sphere σ, defined as $G = \mathbb{E}_{p_\sigma}[E] - TS$ with temperature parameter T controlling exploration-exploitation trade-off. The ℓ_1 penalty $\Lambda\|\pi\|_1$ induces sparsity in the policy representation, favoring policies that activate few latent actions. The cost term $C(\pi) = \int |\nabla S \cdot \pi(x)| dx$ measures computational or metabolic cost of policy execution, with γ weighting this cost.

Parameterizing policies as coefficient vectors a over a dictionary of basis policies, the optimization becomes a constrained LASSO problem: $a^* = \arg\min_a \frac{1}{2}\|y - Xa\|_2^2 + \Lambda\|a\|_1$ subject to $\kappa\|a\|_1 \leq B$, where y represents target semantic outcomes and X maps policies to outcomes. This convex problem can be solved efficiently via proximal gradient methods. The proximal gradient update for coordinate j is $a_j \leftarrow S_\tau(\frac{1}{L_j} X_j^T (y - Xa + X_j a_j))$ where $S_\tau(x) = \text{sign}(x)\max(|x| - \tau, 0)$ is the soft-thresholding operator.

The sparsity phase transition occurs at a critical pressure Λ_c where the active policy count changes discontinuously. This threshold can be characterized analytically for Gaussian policy dictionaries as $\Lambda_c \approx \sqrt{2\log n} \cdot \sigma / \|X^T X\|$ where n is the policy space dimension and σ is the noise standard deviation. For $n \to \infty$, this threshold scales logarithmically with dimension, enabling sparse solutions even in very high-dimensional policy spaces.

Near the critical point, the system exhibits universal scaling behavior characteristic of second-order phase transitions. The correlation length ξ, measuring the range of policy interactions, diverges as $\xi \sim |\Lambda - \Lambda_c|^{-\nu}$ with exponent $\nu \approx 1.2$. This divergence indicates that policies become collectively correlated near criticality, with fluctuations spanning the entire policy space. The natural order parameter is the active policy density $\rho = \|\pi\|_0 / n$. In the dense phase ($\Lambda < \Lambda_c$), $\rho > 0$ and scales linearly with $\Lambda_c - \Lambda$. In the sparse phase ($\Lambda > \Lambda_c$), $\rho \approx 0$ with exponential suppression.

Table 2. Phase Transition Regimes.

Regime	Active Policies	Behavior
Dense	$\|\pi\|_0 > 0$	Distributed inference
Critical	Power-law scaling	Diverging correlations
Sparse	$\|\pi\|_0 \approx 0$	Near-complete collapse

7. Categorical Semantics and Compiler Implementation

The categorical semantics of EBSSC provides a compositional framework for reasoning about policy composition and semantic transformations. The category $\mathcal{E}$ has objects as well-typed semantic spheres σ and morphisms as typed policies $\pi : \sigma \to \sigma'$ satisfying entropy bound $\Delta S(\pi) \leq b_\pi$. Composition is sequential policy application, and identity is the null policy with zero entropy cost. The category carries a symmetric monoidal structure given by the merge operator $\otimes$ representing parallel composition of independent spheres, with the monoidal unit I being the empty sphere containing no semantic content.

Entropy enrichment defines the hom-object $\mathcal{E}(\sigma,\sigma') = \inf_{\pi:\sigma\to\sigma'} \Delta S(\pi)$ as the minimal entropy cost to transform σ to σ'. This satisfies the triangle inequality $\mathcal{E}(\sigma_0,\sigma_2) \leq \mathcal{E}(\sigma_0,\sigma_1) + \mathcal{E}(\sigma_1,\sigma_2)$, providing a quantitative measure of semantic distance with consistency under composition.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 4 March 2026 doi:10.20944/preprints202603.0307.v1

$$\begin{array}{ccccc} \sigma_1 & & \sigma_2 & & \sigma_1 \otimes \sigma_2 \\ {\scriptstyle \pi_1}\downarrow & & \downarrow{\scriptstyle \pi_2} & \Rightarrow & \downarrow{\scriptstyle \pi_1 \otimes \pi_2} \\ \sigma_1' & & \sigma_2' & & \sigma_1' \otimes \sigma_2' \end{array}$$

Figure 3. Monoidal product of policies: parallel composition $\pi_1 \otimes \pi_2$ acts on tensor product of spheres.

The EBSSC compiler transforms high-level policy language into executable code, ensuring verified entropy bounds. The six stages of the compilation pipeline are organized together in a chain. Parsing takes the compiled program (source) and outputs parser type abstract syntax tree, also called by short cut an AST. Type and budget checking ensures the well-typedness and entropy bounds. The optimized Laplace mechanism is realized by solving the sparse free-energy objective. Sheaf lowering converts the single local description into a local description.

During execution, the runtime system continuously monitors three critical invariants: $E_t \leq E_0 + B$ (entropy within budget), $S_t \leq S_{\max}$ (field magnitude bounded), and $\Delta I_t \geq -\epsilon_I$ (mutual information doesn't decrease too rapidly). The monitor checks these invariants after each policy step and can trigger recovery mechanisms if violations are detected. If a policy trace type-checks and compiles through the verified pipeline, execution satisfies these invariants for all time steps.

Parse	AST
Type	Verified
Optimize	Sparse
Sheaf	Presheaf
Cohere	Consistent
Schedule	Executable

Figure 4. EBSSC compiler pipeline showing six stages from source parsing to executable code generation.

The computational complexity scales favorably: for dimension n, sparsity s, and trace length T, per-step cost is $O(sn + nd)$ and full execution cost is $O(T(sn + nd))$. Memory requirements are $O(n^2)$ for storing sphere states and correlations. This linear scaling makes EBSSC suitable for large-scale semantic systems with thousands of dimensions and sparse policy activations.

8. Physical Interpretation and Empirical Predictions

The EBSSC framework is what defines the physical interpretation of things: the spherical semantic spheres are field excitations, the policies are sparse forcing terms and the entropy budgets are physical conservation laws. It is possible to ground cognitive processing in the thermodynamics of information. Per sphere semantic sphere consider an observer who measures this content by given it some entropy value. Therefore, the entropy is a good only measure of the field excitation as well as the policy. The application of policy similarly alters the fields' excitation and the entropy by some measure. The generation of entropy may be a positive or negative quantity, which in physical terms corresponds to gain or loss of information.

The free-energy functional combines kinetic energy of the field, kinetic energy of policy flow, entropy density, and sparsity pressure: $\mathcal{F}[\Phi, \mathbf{v}, S] = \int_\Omega \left(\frac{1}{2}|\nabla\Phi|^2 + \frac{\alpha}{2}|\mathbf{v}|^2 + \beta S + \Lambda\|\mathbf{v}\|_1 \right) dx$. The Euler-Lagrange equations derived from this functional yield the dynamical equations governing semantic evolution. The semantic field Φ evolves according to a modified diffusion equation with sources from policy actions: $\partial_t\Phi = D\nabla^2\Phi - \nabla \cdot (\mathbf{v}\Phi) + \eta(t)$. The diffusion term smooths field gradients and increases entropy, while advection transports semantic content along policy directions.

The variational structure implies conservation laws via Noether's theorem. Time translation symmetry implies conservation of total energy $H = \int (\frac{1}{2}|\nabla\Phi|^2 + \frac{\alpha}{2}|\mathbf{v}|^2) dx$. Spatial translation symmetry implies conservation of total semantic momentum $\mathbf{P} = \int \Phi\nabla\Phi dx$. Gauge symmetry $\Phi \to e^{i\theta}\Phi$ implies

 doi:10.20944/preprints202603.0307.v1

conservation of semantic charge $Q = \int |\Phi|^2 dx$. These conservation laws provide fundamental bounds on semantic evolution.

Table 3. Physical Correspondence Between EBSSC and Thermodynamics.

Physical Concept	EBSSC Analog
Maximum entropy	Optimal reconstruction
Free energy	$G(\sigma, \pi) = E - TS$
Fisher geometry	Sphere metric g_{ij}
Second law	Merge entropy $\Delta S > 0$
Thermodynamic work	Policy cost $C(\pi)$
Heat dissipation	Entropy production $\dot{\Sigma}$
Noether current	Conserved quantities

The framework makes several falsifiable empirical predictions. Neural activity in semantic processing should exhibit sparsity scaling with network size: for a network of n neurons, active neuron count k should scale as $k \sim n^{\alpha}$ with exponent $\alpha < 1$, indicating sublinear growth. Word embeddings and semantic representations should exhibit bounded entropy growth over time during text processing: $S(t) \leq S_0 + \kappa \log t$. Perturbations to semantic representations should propagate at finite speed, defining a semantic light cone with radius $r(t) \leq c_s t$.

Transformer attention patterns should exhibit a phase transition as sparsity pressure varies, with correlation length diverging as $\xi \sim |\Lambda - \Lambda_c|^{-1.2}$ at criticality. Semantic memory should exhibit exponential decay of trace strength: $H(t) = H_0 e^{-ct}$. These predictions can be tested through neural recording experiments, embedding space analysis, controlled perturbation studies, and memory experiments with semantic material.

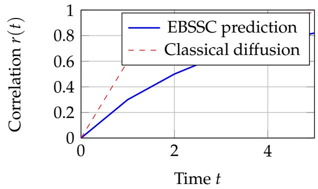

Figure 5. Semantic light cone: correlation radius as function of time after perturbation. EBSSC predicts sublinear growth bounded by finite propagation speed.

Table 4. Falsifiable Empirical Predictions of EBSSC.

Prediction	Expected Scaling	Falsification
Neural sparsity	$k \sim n^{\alpha}, \alpha < 1$	$\alpha \approx 1$
Entropy growth	$S(t) \leq S_0 + \kappa \log t$	$S(t) \sim t$
Semantic light cone	$r(t) \leq c_s t$	Instantaneous spread
Critical sparsity	$\xi \sim |\Lambda - \Lambda_c|^{-1.2}$	No divergence
Memory decay	$H(t) \sim e^{-ct}$	Power-law decay

9. Discussion and Conclusion

The EBSSC framework asserts foundational claims about cognition and agency. It formalizes thought as a bounded process that is not unlimited computation, but resource constrained field manipulation subject to thermodynamic laws. This aligns with the extra embodied approaches to cognitive science that emphasized the physical constraints involved in making an intelligence. Meaning, in this view, is curved space in field semantics. A well-formed and usable concept corresponds to low-curvature, stable entropy points or regional spaces or form-elements of the plenum. This account gives coherence to a novel geometric measure of meaning. According to the EBSSC formalization, thinking occurs within limits; in other words, cognition is not unbounded.

This shows how understanding is connected to compression. To be more specific, it means that any domain, in particular, can be conceptualized as having a minimal generative policy that reconstructs observations.

This makes a formal connection of epistemology with information theory. Agency refers to meager control. The policies of intelligent agents are sparse choices taken from large latent spaces. This sparsity doesn't emerge due to any particular peculiar architecture but for physical reasons. A basic interpretation of neural sparsity is offered as a critical feature of bounded.

EBSSC provides developers with principles for designing AI systems that are interpretable and verifiable. A situation where you are able to identify the cause of a state change and other relevant details.

Bounded drift from entropy is a constraint on operations that change states which prevents semantic representations from wandering too far from their source states. Sparse explanations refer to only providing the proof steps needed for making.

The Entropy-Bounded Sparse Semantic Calculus is a compositional framework that brings together geometric semantics, sparse inference, and resource constraints. This is of immediate relevance to volatile semantic memory's essence but we expect it to lay mathematical foundations for a bounded agency. Thus, this should generate wider insights linking to scientific and engineering practice. Future research directions will focus on semantic learning from experience, connecting continuous fields to discrete symbolic reasoning, and scaling to millions of diverse.

References

1. J. Gupta, Precise web programming with embedded domain-specific languages and dependently-typed architecture, in Proc. Int. Conf. Softw. Eng., 2025. DOI: 10.1007/978-981-96-8998-9_24
2. J. Gupta, Exploring deep reinforcement learning algorithms—From theory to practice, in Proc. Int. Conf. Mach. Learn., 2025. DOI: 10.1007/978-981-96-9494-5_12
3. P. Baral, Deep learning for text-based sentiment prediction, Preprint, 2025.
4. P. Baral, Uncovering neural dynamics: A statistical approach to analyzing brain activity across timescales, Preprint, 2025.
5. P. Baral, Adaptive polynomial frameworks for enhanced graph signal learning in data science, Preprint, 2025.
6. M. Mitra, Legal and ethical considerations in correctional dental care: A study of liability and service scope, in Proc. Int. Conf. Health Law, 2025. DOI: 10.2991/978-2-38476-533-1_28
7. M. Mitra, Constructing custom AI simulations for counselor education using feedback-guided principles, in Proc. Int. Conf. AI Educ., 2025. DOI: 10.2991/978-94-6463-950-6_6
8. M. Mitra, MUC4 and MMP7 in saliva and gingival crevicular fluid in adolescents at West Bengal, India, Bioinformation, vol. 18, no. 3, pp. 165-172, 2025. DOI: 10.6026/97320630018165
9. A. Baer, Optimized task distribution for energy conservation in electric vehicles via edge computing networks, Preprints, 2025. DOI: 10.20944/preprints202510.1969.v1
10. A. Baer, Characterizing performance fading in lithium-ion cells: An efficiency-centric approach, Preprints, 2025. DOI: 10.20944/preprints202510.1727.v1
11. A. Baer, Autonomous regulation of energy storage units: An analysis of systemic control in electric mobility, Preprints, 2025. DOI: 10.20944/preprints202510.1712.v1
12. M. Pratap, AI-driven causal analysis for business intelligence: A scalable approach from Salesforce research, in Proc. IEEE Int. Conf. Emerg. Comput., 2025. DOI: 10.1109/ICEC2NT65402.2025.11380110
13. M. Pratap, Implicit Bayesian learning for enhanced sales forecasting in Salesforce CRM, in Proc. IEEE Int. Conf. Emerg. Res. Electron., 2025. DOI: 10.1109/ICERECT65215.2025.11378062
14. M. Pratap, Maximizing business growth: The economic impact of Salesforce-driven CRM intelligence, in Proc. IEEE Int. Conf. Next Gener. Comput., 2025. DOI: 10.1109/ICNGCS64900.2025.11183293
15. D. Tandel, Sustainable Shiitake cultivation: A small-scale approach for the food industry, Zenodo, 2025. DOI: 10.5281/zenodo.16441005
16. D. Tandel, Re-optimization and evaluation of global food systems: Case studies from China, USA, and Ethiopia, in Universal Threats in Expert Applications and Solutions, LNNS vol. 1450, Springer, 2026. DOI: 10.1007/978-981-96-7289-9_11

17. D. Tandel, Modeling resilience in food systems: A theoretical framework for analyzing the UK pork industry, Preprints, 2025. DOI: 10.20944/preprints202507.1106.v1
18. V. Kaliappan, Digital safeguards for health data: Architecting privacy and security in next-gen medical informatics, in Proc. IEEE Int. Conf. Big Data Knowl. Eng., 2025. DOI: 10.1109/BdKCSE67969.2025.11300511
19. V. Kaliappan, Digital persona generation: Historical figure emulation in learning, Preprint, 2025.
20. S. Gaikwad, Study on artificial intelligence in healthcare, in Proc. IEEE Int. Conf. Adv. Comput. Commun. Syst., 2021. DOI: 10.1109/ICACCS51430.2021.9441741
21. P. Patel, Evaluating ensemble learning strategies for enhanced medical diagnostics: Insights from real-world datasets, in Proc. 6th Int. Conf. Problems Cybern. Informatics, 2025, pp. 1-4. DOI: 10.1109/PCI66488.2025.11219757
22. T. Mendhey, Category theory and psychodynamics: Bridging the structure of programming with human behavior, J. Interdiscip. Learn. Technol., vol. 26, no. 1, 2025. DOI: 10.70729/SE26121213107
23. T. Mendhey, Transforming workforce potential: Integrating training, competency systems, and emerging technologies for organizational excellence, Recent Trends Manage. Commerce, vol. 6, no. 1, pp. 1-7, 2025. DOI: 10.46632/rmc/6/1/1
24. T. Mendhey, Examining USDA reimbursements, school nutrition programs, and food procurement in U.S. public schools: A comparative analysis of California and Texas, J. Int. Curric. Learn. Technol., vol. 26, no. 1, 2025. DOI: 10.61336/Jiclt/26-01-27
25. W. A. Shakir, Reinforcement Learning Challenges in Power Grid Management: A Case Study with City Learn Simulator, in Proc. Int. Conf. Comput. Intell., 2024. DOI: 10.1007/978-981-96-7238-7_24
26. W. A. Shakir, Leveraging Artificial Intelligence for Strategic Advancement: Opportunities and Initiatives at the Miller Center, in Proc. IEEE Int. Conf. Emerg. Technol. Innov., 2024. DOI: 10.1109/ICETI63946.2024.10777212