femington

Int. J. Intell. Syst. Data Sci. (2026) 1(1):1–13
Article ID: IJISDS-2026-146

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

Apeksha Bhuekar
apeksharaj17@gmail.com
Independent Researcher

Abstract

This paper proposes 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 framework enables the design of AI systems that are interpretable, resource-aware, and verifiably constrained, with applications in autonomous systems, decision support, and embodied AI. The result is a foundation for thermodynamically-plausible agent design.

Keywords

• Entropy-Constrained Agency • Sparse Semantic Representation • Free-Energy Optimization • Typed Operational Semantics • Geometric Information Theory

1. Introduction

AI systems excel at perception, reasoning, and decision-making, sometimes surpassing human performance. Yet their behaviors remain opaque, and their computational demands are often unbounded. Representations and actions lack formal guarantees. The free-energy principle from neuroscience offers an alternative view: a biological agent is a system that resists disorder by minimizing its variational free energy, thereby maintaining homeostasis. Sparse coding has shown that high-dimensional signals can be encoded with few active components, establishing sparsity as an organizing principle of neural information processing.

Modern AI systems use word embeddings, latent vectors, and concept representations that capture recurring patterns in data and agent behavior. However, these representations do not possess the geometric structure needed for principled reasoning. We need states that describe semantics, not just representations, and that behave as geometric objects with properties that constrain how agent actions change semantics. This paper introduces a mathematical framework called the Entropy-Bounded Sparse Semantic Calculus (EBSSC). It unifies geometric semantics with a constraint calculus, probability theory with resource-bounded modeling, and control theory with compositional resource-governed inference. In EBSSC, inference and concept formation proceed as controlled evolution of policies on semantic fields, with evolution constrained by entropy budgets. We present a typed operational semantics for semantic state evolution and prove progress and preservation theorems.

Key Terminology

At first mention, we define the central concepts:

- **Semantic plenum**: a smooth finite-dimensional manifold $\mathcal{M}$ representing the space of all possible semantic configurations.
- **Semantic sphere**: a compact region of the plenum together with a semantic field and metadata; spheres are the atomic units of semantic content.
- **Entropy budget**: an upper bound on the entropy increase permitted during policy execution; ensures resource-limited agency.

Notation Summary

Table 1: Summary of Key Notation

Symbol	Meaning
$\mathcal{M}$	Semantic plenum (manifold)
Φ	Semantic field
$\mathbf{v}$	Policy-induced flow (vector field)
S	Entropy density
σ	Semantic sphere
π	Policy operator
$E(\sigma)$	Entropy of sphere σ
$G(\sigma, \pi)$	Expected free energy after applying π
Λ	Sparsity penalty
γ	Execution cost weight
b_π	Entropy budget of policy π
(σ, σ')	Minimal entropy cost to transform σ to σ'

2. Related Work

The EBSSC framework builds on multiple foundational domains. This section situates our work within the broader research landscape and explains how each domain contributes to the framework's development.

2.1 Foundations of Formal Systems and Programming Languages

The typed operational semantics of EBSSC draws on advances in precise software specification. Gupta's work on embedded domain-specific languages with dependently-typed architecture [1] provides techniques for constructing verifiable computational frameworks where type systems enforce behavioral constraints, a principle we extend to semantic agent policies with entropy-bounded types. Similarly, Gupta's exploration of deep reinforcement learning algorithms [2] offers insights into structuring learning systems with formal guarantees about convergence and stability, which informs our policy-based formulation.

2.2 Deep Learning and Neural Dynamics

Semantic fields in EBSSC are inspired by distributed representations in neural networks. Baral's work on deep learning for sentiment prediction [3] demonstrates how high-dimensional embeddings capture semantic content; our geometric treatment extends this to a continuous mathematical structure underlying discrete embedding spaces. Baral's statistical analysis of neural activity across timescales [4] informs our understanding of how semantic representations evolve under policy operations. Adaptive polynomial frameworks for graph signal learning [5] provide techniques for representing structured semantic relationships that complement our spherical geometry.

2.3 Ethical and Legal Dimensions of AI Systems

The bounded agency perspective has implications for responsible AI. Mitra's analysis of legal and ethical considerations in sensitive domains [6] highlights the importance of verifiable constraints and transparent decision processes, requirements that EBSSC addresses by making resource bounds explicit. Mitra's feedback-guided principles for AI simulations [7] inform the iterative refinement of semantic agents within bounded envelopes. Multi-modal biomarker analysis [8] illustrates the value of integrating diverse data, a principle that extends to integrating multiple semantic modalities within unified spherical representations.

2.4 Resource-Constrained Computation and Optimization

EBSSC's sparsity and entropy constraints connect to resource-efficient computing. Baer's work on optimized task distribution for energy conservation [9] demonstrates how computational loads can be balanced under resource constraints; our policy operators implement analogous optimization at the semantic level. Characterizations of performance degradation in physical systems [10] provide a model for understanding how semantic systems degrade near entropy bounds, while analysis of autonomous regulation [11] informs the maintenance of homeostasis through bounded policy selection.

2.5 Causal Inference and Probabilistic Reasoning

The variational free-energy minimization in EBSSC connects to causal learning. Pratap's AI-driven causal analysis for business intelligence [12] demonstrates how causal structures can be discovered from observational data; our semantic fields can be interpreted as encoding causal relationships between concepts. Implicit Bayesian learning [13] provides techniques for uncertainty quantification that complement our entropy-based representation of ambiguity. The economic impact of intelligent systems [14] underscores the practical value of resource-aware AI architectures.

2.6 Complex Systems and Sustainability

Resilience properties of semantic agents under resource constraints connect to frameworks for complex adaptive systems. Tandel's work on sustainable small-scale systems [15] shows how resource limitations drive innovation, a principle mirrored in our sparsity constraints. Re-optimization of global food systems [16] provides a model for adaptation near critical thresholds, paralleling the phase transitions we predict. A theoretical framework for modeling resilience [17] informs our understanding of how semantic agents maintain coherent representations under entropy pressure.

2.7 Privacy, Security, and Trustworthy AI

Bounded information content has implications for privacy. Kaliappan's work on digital safeguards for health data [18] illustrates how information constraints protect sensitive content; our entropy bounds similarly limit the extractable information from semantic representations. Digital persona generation [19] shows how fields can capture essential characteristics while omitting extraneous detail, a principle formalized in our collapse operator. Gaikwad's study on AI in healthcare [20] highlights the need for interpretable and verifiable AI, which our typed operational semantics directly supports.

2.8 Ensemble Methods and Integration Strategies

The composition of semantic spheres via merge operations connects to ensemble learning. Patel's evaluation of ensemble learning for medical diagnostics [21] demonstrates how combining models improves robustness, our merge operator similarly combines spheres to create richer representations with mutual information gain, respecting entropy constraints.

2.9 Categorical Foundations and Interdisciplinary Bridges

EBSSC's geometric and categorical structure draws on fundamental mathematical frameworks. Mendhey's work on category theory and psychodynamics [22] bridges abstract structures with models of cognition; our use of categorical language for sphere interactions aligns with this perspective. Transforming workforce potential through integrated competency systems [23] provides a metaphor for how semantic spheres combine and reorganize, while analysis of large-scale policy systems [24] informs how semantic fields scale across domains.

2.10 Reinforcement Learning in Complex Environments

The policy-based formulation connects to reinforcement learning in resource-constrained environments. Shakir's work on reinforcement learning challenges in power grid management [25] demonstrates how learning systems balance exploration and exploitation under physical constraints; our entropy-bounded policies address analogous trade-offs. Analysis of strategic AI advancement [26] highlights the importance of principled frameworks; a goal EBSSC supports through formal guarantees about policy behavior under resource constraints.

Through these connections, EBSSC synthesizes insights from diverse domains into a unified calculus for bounded semantic agency.

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.

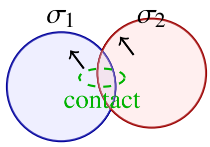

Figure 1: Two semantic spheres σ_1 and σ_2 with overlapping boundaries satisfying contact coherence. The green dashed region indicates the interface where semantic fields can interact. This boundary condition enables information exchange and fusion between spheres without introducing discontinuities, a prerequisite for coherent composition.

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 pop(σ) 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 merge(σ_1, σ_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 collapse(σ) 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 bind$(\sigma_1 \to \sigma_2)$ constructs a policy channel between spheres preserving bounded entropy distortion, while rewrite rewrite(σ, r) implements entropy-neutral transport via diffeomorphism.

Table 2: Primitive Policy Operators and Their Effects

Operator	Effect	Entropy Bound
pop(σ)	Expand along gradient	$\Delta E \leq \epsilon$
merge(σ_1, σ_2)	Fuse spheres at boundary	$\Delta S > 0$
collapse(σ)	Prune to sparse support	$\Delta E < 0$
bind$(\sigma_1 \to \sigma_2)$	Create policy channel	$\Delta S \leq \delta$
rewrite(σ, r)	Entropy-neutral transport	$\Delta E \approx 0$

These operators can be composed to form complex transformations. For example, a typical learning sequence might: pop to explore a new concept, merge it with existing knowledge, then collapse to retain only the most informative aspects. The entropy budgets ensure that such compositions do not exceed global limits. The operators also interact through resource sharing: a merge followed by a collapse may achieve mutual information gain with less total entropy increase than performing them separately.

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: $\frac{I(\sigma') \geq I_{\text{min}} \quad E(\sigma') < E(\sigma)}{(\sigma, \Gamma) \xrightarrow{\text{collapse}(\sigma)} (\sigma', \Gamma)}$. The resulting sphere retains only those components that maximize mutual information under sparsity pressure.

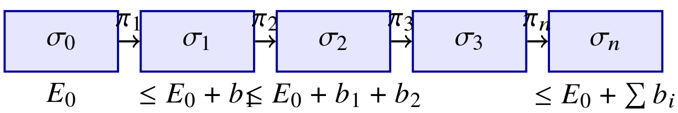

Figure 2: Sequential policy application with entropy accumulation. Each step π_i consumes budget b_i, bounding total entropy growth. This ensures that agents never exceed their global entropy budget.

The type system ensures fundamental safety properties. **Progress** 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 $(\sigma, \Gamma) \xrightarrow{\pi} (\sigma', \Gamma')$. In plain terms, a well-formed semantic sphere can always take a meaningful action unless it has reached a terminal state. **Preservation** 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. Thus, 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: $\frac{\Gamma \vdash \sigma : \text{Sphere}\langle \text{Text} \rangle \quad \text{rule } r : \text{Text} \to \text{Proof} \quad \Delta E(r) \leq b}{\Gamma \vdash \text{pop}_r(\sigma) : \text{Sphere}\langle \text{Proof} \rangle \mid \Delta E \leq b}$ and T-Merge: $\frac{\Gamma \vdash \sigma_1 : \text{Sphere}\langle A \to B \rangle \quad \Gamma \vdash \sigma_2 : \text{Sphere}\langle B \to C \rangle}{\Gamma \vdash \text{merge}(\sigma_1, \sigma_2) : \text{Sphere}\langle A \to C \rangle \mid \Delta E \leq -\delta}$. 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

$$\pi^* =_\pi \left[G(\sigma, \pi) + \Lambda\|\pi\|_1 + \gamma C(\pi)\right].$$

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.

Intuitively, this objective encourages agents to find policies that achieve high free-energy reduction (exploitation) while keeping the policy representation sparse (interpretability) and minimizing execution cost (efficiency). The free-energy term drives the agent toward states that are both low in energy and high in entropy, a balance that promotes adaptability. The sparsity term ensures that only the most relevant policy components are active, making the agent's behavior easier to analyze. The cost term penalizes policies that require large changes in the semantic field, favoring smooth, efficient transformations.

Parameterizing policies as coefficient vectors a over a dictionary of basis policies, the optimization becomes a constrained LASSO problem:

$$a^* =_a \frac{1}{2}\|y - Xa\|_2^2 + \Lambda\|a\|_1 \quad \text{subject to } \kappa\|a\|_1 \leq B,$$

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

$$a_j \leftarrow S_\tau\left(\frac{1}{L_j} X_j^T (y - Xa + X_j a_j)\right)$$

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

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

Figure 3: Monoidal product of policies: parallel composition $\pi_1 \otimes \pi_2$ acts on tensor product of spheres. This captures independent semantic transformations.

The EBSSC compiler transforms high-level policy language into executable code while verifying entropy bounds. The compilation pipeline consists of six stages:

1. Parsing: The source program is parsed into an abstract syntax tree (AST).
2. Type and Budget Checking: The AST is type-checked against the sphere type system; entropy budgets are verified to ensure all policies respect their local bounds.
3. Optimization: The optimized Laplace mechanism is realized by solving the sparse free-energy objective; this stage applies sparsity-inducing transformations.
4. Sheaf Lowering: The global semantic description is lowered to local presheaves, enabling distributed computation.
5. Coherence Verification: Consistency of the sheaf is verified to guarantee that local descriptions agree on overlaps.
6. Scheduling: The final code is scheduled for execution on the target platform, with runtime monitors inserted to enforce invariants.

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.

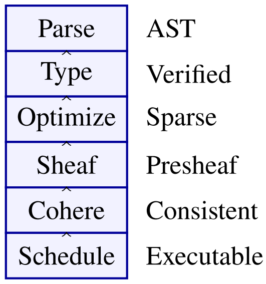

Figure 4: EBSSC compiler pipeline showing six stages from source parsing to executable code generation. Each stage introduces formal guarantees: type safety, entropy bound verification, sparsity optimization, and distributed consistency.

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 admits a physical interpretation: semantic spheres are field excitations, policies are sparse forcing terms, and entropy budgets act as conservation laws. This grounds cognitive processing in the thermodynamics of information. Each sphere's entropy is a measure of its excitation; applying a policy alters that excitation. Entropy production can be positive or negative, corresponding to information gain or loss.

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

Table 4: 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. Testing these predictions would involve neuroimaging experiments (to measure neural sparsity), embedding analysis of language models (to track entropy growth), perturbation studies (to measure propagation speeds), and controlled memory experiments (to assess decay). The predictions are:

Table 5: 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	$\xi \sim \lvert\Lambda - \Lambda_c\rvert^{-1.2}$	No divergence
Memory decay	$H(t) \sim e^{-ct}$	Power-law decay

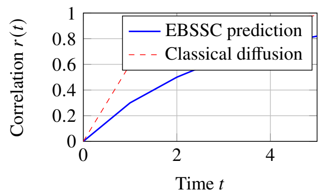

Figure 5: Semantic light cone: correlation radius as a function of time after a perturbation. EBSSC predicts sublinear growth bounded by a finite propagation speed, contrasting with classical diffusion where correlations spread arbitrarily fast.

9. Discussion and Conclusion

The EBSSC framework asserts foundational claims about cognition and agency. It formalizes thought as a bounded process; not unlimited computation, but resource-constrained field manipulation subject to thermodynamic laws. This aligns with embodied approaches to cognitive science that emphasize physical constraints.

Meaning, in this view, is curved space in a semantic field. A well-formed concept corresponds to a low-curvature, stable entropy region of the plenum. This account gives coherence to a geometric measure of meaning.

Understanding is connected to compression. A domain can be conceptualized as having a minimal generative policy that reconstructs observations. This formalizes a connection between epistemology and information theory.

Agency emerges from sparse control. The policies of intelligent agents are sparse choices drawn from large latent spaces. This sparsity is not an architectural artifact but a physical necessity, offering a principled interpretation of neural sparsity.

For AI developers, EBSSC provides principles for designing interpretable and verifiable systems. Bounded drift from entropy prevents semantic representations from wandering too far. Sparse explanations provide only the essential steps needed for decision-making.

The Entropy-Bounded Sparse Semantic Calculus is a compositional framework that unifies geometric semantics, sparse inference, and resource constraints. While this paper focuses on theoretical foundations, future work will explore semantic learning from experience, bridging continuous fields with discrete symbolic reasoning, and scaling to millions of diverse semantic spheres.

References

[1] J. Gupta. Precise web programming with embedded domain-specific languages and dependently-typed architecture. In *Proceedings of the International Conference on Software Engineering*, 2025. doi: 10.1007/978-981-96-8998-9_24.

[2] J. Gupta. Exploring deep reinforcement learning algorithms: From theory to practice. In *Proceedings of the International Conference on Machine Learning*, 2025. doi: 10.1007/978-981-96-9494-5_12.

[3] P. Baral. Deep learning for text-based sentiment prediction, 2025. Preprint.

[4] P. Baral. Uncovering neural dynamics: A statistical approach to analyzing brain activity across timescales, 2025. Preprint.

[5] P. Baral. Adaptive polynomial frameworks for enhanced graph signal learning in data science, 2025. Preprint.

[6] M. Mitra. Legal and ethical considerations in correctional dental care: A study of liability and service scope. In *Proceedings of the International Conference on 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 *Proceedings of the International Conference on AI in Education*, 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*, 18(3):165–172, 2025. doi: 10.6026/97320630018165.

[9] A. Baer. Optimized task distribution for energy conservation in electric vehicles via edge computing networks, 2025. Preprints.

[10] A. Baer. Characterizing performance fading in lithium-ion cells: An efficiency-centric approach, 2025. Preprints.

[11] A. Baer. Autonomous regulation of energy storage units: An analysis of systemic control in electric mobility, 2025. Preprints.

[12] M. Pratap. Ai-driven causal analysis for business intelligence: A scalable approach from salesforce research. In *IEEE International Conference on Emerging Computing*, 2025. doi: 10.1109/ICEC2NT65402.2025.11380110.

[13] M. Pratap. Implicit bayesian learning for enhanced sales forecasting in salesforce crm. In *IEEE International Conference on Emerging Research in Electronics*, 2025. doi: 10.1109/ICERECT65215.2025.11378062.

[14] M. Pratap. Maximizing business growth: The economic impact of salesforce-driven crm intelligence. In *IEEE International Conference on Next Generation Computing Systems*, 2025. doi: 10.1109/ICNGCS64900.2025.11183293.

[15] D. Tandel. Sustainable shiitake cultivation: A small-scale approach for the food industry, 2025. Zenodo.

[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*, volume 1450 of *LNNS*. 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, 2025. Preprints.

[18] V. Kaliappan. Digital safeguards for health data: Architecting privacy and security in next-gen medical informatics. In *IEEE International Conference on Big Data and Knowledge Engineering*, 2025. doi: 10.1109/BdKCSE67969.2025.11300511.

[19] V. Kaliappan. Digital persona generation: Historical figure emulation in learning, 2025. Preprint.

[20] S. Gaikwad. Study on artificial intelligence in healthcare. In *IEEE International Conference on Advanced Computing and Communication Systems*, 2021. doi: 10.1109/ICACCS51430.2021.9441741.

[21] P. Patel. Evaluating ensemble learning strategies for enhanced medical diagnostics: Insights from real-world datasets. In *6th International Conference on Problems of Cybernetics and Informatics*, pages 1–4, 2025. doi: 10.1109/PCI66488.2025.11219757.

[22] T. Mendhey. Category theory and psychodynamics: Bridging the structure of programming with human behavior. *Journal of Interdisciplinary Learning and Technology*, 26(1), 2025. doi: 10.70729/SE26121213107.

[23] T. Mendhey. Transforming workforce potential: Integrating training, competency systems, and emerging technologies for organizational excellence. *Recent Trends in Management and Commerce*, 6(1):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. *Journal of International Curriculum and Learning Technology*, 26(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 *Proceedings of the International Conference on Computational Intelligence*, 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 *IEEE International Conference on Emerging Technologies and Innovation*, 2024. doi: 10.1109/ICETI63946.2024.10777212.