

Homotopy and Identity

A Dynamical Foundation for Upper-Level Ontology

Flyxion

February 17, 2026

Abstract

Upper-level ontology plays a central role in stabilizing meaning across heterogeneous scientific and institutional systems. By imposing disciplined distinctions among continuants, occurrents, and dependence relations, realist frameworks have provided durable foundations for large-scale interoperability. Yet most such frameworks remain fundamentally entity-centric. Objects are treated as ontologically primary, while history, irreversibility, informational structure, and constraint propagation are rendered derivative.

This paper argues that persistent difficulties in ontology engineering irreversibility, mapping fragility, semantic drift, and the ontological status of information reflect a deeper structural limitation in entity-first metaphysics. In response, a history-first alternative is developed in which admissible trajectories, entropy bounds, and stabilization conditions are taken as primitive. Within this framework, entities are reconceived as low-entropy invariants emerging from constrained historical structure rather than as foundational atoms of description.

The argument proceeds in four stages. First, the necessity of upper-level ontology is clarified and a detailed exposition of entity-centric formal ontology is provided. Second, structural limitations arising from entity primacy are analyzed. Third, a field-theoretic reconstruction grounded in admissible histories and entropy constraints is introduced. Finally, it is shown how traditional upper ontologies may be embedded within this broader dynamical architecture, preserving realist discipline while extending ontological analysis to domains characterized by irreversibility and adaptive complexity.

Contents

1	The Necessity of Upper-Level Ontology	3
2	Core Structure of the Basic Formal Ontology	4
2.1	Continuants and Occurrents	5
2.2	Independent and Dependent Continuants	5
2.3	Relations and Participation	6
2.4	Strengths of the BFO Framework	6
3	Irreversibility and the Limits of Entity Primacy	7
4	A History-First Ontology of Entropic Fields	9
4.1	Scalar Density and Stabilization	10
4.2	Vector Fields and Directed Constraint Propagation	10
4.3	Entropy as Degeneracy of Futures	10
4.4	Derived Ontological Categories	11
4.5	Regime-Dependence and Ontological Scope	11
5	Embedding the Basic Formal Ontology as a Low-Entropy Subtheory	12
6	Implications for Ontology Engineering and Interoperability	13
7	Formal Comparison with Existing Upper-Level Ontologies	15
7.1	Comparison with the Basic Formal Ontology	15
7.2	Comparison with DOLCE and Process Ontologies	16
8	Logic, Statistics, and Ontological Commitments in Artificial Intelligence	16
9	Historical Cycles and Ontological Regimes	18

9.1	Comparison with Information-Centric Ontologies	19
9.2	Regime-Relativity and Ontological Scope	19
10	Conclusion	20
A	Admissible Histories and Entropy Bounds	22
A.1	Admissible History Space	22
A.2	Entropy as Degeneracy of Futures	22
A.3	Scalar Density and Stability	23
A.4	Vector Field and Directed Constraint Flow	23
A.5	Low-Entropy Subtheories	24
B	Regime Structure and Identity Conditions	24
B.1	History Space as a Directed Category	24
B.2	Entropy Bounds as Subcategory Constraints	24
B.3	Scalar Density as Stability Functional	25
B.4	Derived Identity Conditions	25
B.5	Process Structure	26
B.6	Regime Transition	26
C	Analysis of Logical and Stochastic AI	26
C.1	Logical AI as a Zero-Entropy Regime	26
C.2	Brittleness as Entropy Discontinuity	27
D	Categorical Embedding and Homotopy Structure	27
D.1	Embedding of Entity-Centric Ontology	27
D.2	Sheaf Structure and Ontology Alignment	28
D.3	Homotopy-Type of Regime Structure	28

D.4	Derived Stack Interpretation	29
D.5	Entropy Gradient and Morse Structure	29
D.6	Stochastic AI as High-Entropy Regime	30
D.7	Gradient Flow and Entropy Suppression	30
D.8	Regime Transition Theorem	31
D.9	Weak vs Strong AI Formalization	31

1 The Necessity of Upper-Level Ontology

Ontology engineering did not arise from abstract metaphysical curiosity alone, but from concrete failures of interoperability. As scientific domains expanded in scale and specialization, representational systems proliferated. Distinct communities developed their own terminologies, classificatory schemes, and modeling conventions. What initially appeared as harmless variation gradually accumulated into structural incompatibility. Data integration became fragile. Semantic drift increased with each version revision. Mappings required continual repair. Under such conditions, the absence of shared ontological discipline revealed itself not as a philosophical oversight but as an engineering liability.

Upper-level ontologies emerged as a response to this instability. Their central promise is not expressive richness but structural constraint. By supplying a small set of domain-neutral primitives categories intended to apply across all scientific domains they impose boundary conditions on admissible modeling moves. The function of an upper ontology is therefore regulative rather than descriptive. It does not aim to enumerate the furniture of the world in detail; instead, it seeks to constrain how domain ontologies may describe that furniture without contradiction or category error.

This constraint function is best understood against the background of semantic fragmentation. When two independently evolving ontologies attempt alignment, they often rely on entity-to-entity correspondences. Such correspondences assume that the mapped entities possess sufficiently stable identity conditions across contexts. Yet in practice, ontological structures evolve under distinct institutional pressures, modeling purposes, and historical contingencies. Mappings that initially appear coherent degrade under version drift. The resulting fragility is not accidental but structural: static correspondences are imposed upon systems that evolve under incompatible dynamics. Without a shared constraint layer governing admissible extensions, alignment becomes a perpetual repair process rather than a convergence mechanism.

Upper-level ontologies attempt to address this fragility by stabilizing the space of admissible historical developments. They do so by enforcing realist commitments. Reality, on this view, exists independently of conceptual schemes. Ontological categories are not artifacts of language but reflect structures that hold regardless of representation. This realist discipline is intended to prevent ontologies from collapsing into pragmatic taxonomies tailored to local use cases. By demanding that categories track mind-independent structure, upper ontologies introduce a form of ontological hygiene that curbs uncontrolled proliferation of types.

The Basic Formal Ontology represents one of the most rigorous articulations of this approach. It insists that ontological modeling must respect distinctions grounded in reality itself, rather than in linguistic convenience. The separation between continuants and occurrents, for example, is not presented as a mere heuristic but as a reflection of how entities persist or unfold in time. Such commitments aim to reduce ambiguity at the highest structural level, thereby decreasing entropy within the representational system. When identity conditions are clearly articulated and participation relations precisely defined, downstream modeling decisions inherit this stability.

Nevertheless, the very strength of upper-level ontologies their emphasis on entity classification reveals a deeper assumption. They presuppose that stable entities are ontologically primary, and that histories, processes, and informational transformations are derivative descriptions of what happens to those entities. This presupposition works effectively in domains characterized by low entropy and well-defined persistence conditions. Anatomical structures, chemical substances, and institutional artifacts often display sufficient stability that entity-centric classification yields durable models.

The necessity of upper-level ontology, then, is not in question. Large-scale knowledge systems require constraint layers to prevent semantic collapse. The question is whether entity primacy is the correct primitive upon which such constraint should be grounded. Persistent difficulties surrounding irreversibility, adaptive systems, and informational structure suggest that stability itself may require deeper explanation. If identity conditions are not primitive but emergent, then an ontology that begins with entities may capture only a special case of a more general historical structure.

The remainder of this essay proceeds under the hypothesis that upper-level ontology remains indispensable, but that its primitives require reconsideration. We therefore turn first to a detailed exposition of the Basic Formal Ontology in order to clarify both its achievements and its implicit commitments.

2 Core Structure of the Basic Formal Ontology

The Basic Formal Ontology is designed as a domain-neutral upper ontology intended to support interoperability across scientific and biomedical domains. Its guiding ambition is to provide a small, rigorously articulated set of top-level categories that reflect mind-independent structure. These categories are not intended to exhaustively describe reality, but to constrain how more specific domain ontologies may do so without contradiction.

2.1 Continuants and Occurrents

The central structural distinction within BFO is that between continuants and occurrents. This distinction functions as the primary ontological bifurcation upon which all further refinements depend.

Continuants are entities that persist through time while maintaining their identity. They are wholly present at every moment of their existence. An anatomical organ, a molecule, a person, or an institution are paradigmatic examples. A continuant may undergo qualitative change, but it does not unfold in time as a process; rather, it endures through time.

Occurrents, by contrast, are entities that unfold temporally. They are not wholly present at any single instant but extend across temporal intervals. Processes, events, and temporal regions fall under this category. A metabolic process, a conversation, or a geological transformation exemplifies an occurrent. Such entities are characterized by temporal parts rather than enduring presence.

The continuant/occurrent distinction is not merely classificatory but metaphysically significant. It enforces a discipline in which persistence and unfolding are categorically separated. In modeling practice, this separation prevents the conflation of entities that endure with processes that transpire. Participation relations mediate the interaction between these categories: continuants participate in occurrents, while occurrents depend on continuants for their realization.

This architecture reflects a commitment to ontological realism. The distinction is not treated as a modeling convenience but as a structural feature of reality itself. It aims to capture the intuition that a human being and a human life process are ontologically different kinds of entities, even if they are intimately related.

2.2 Independent and Dependent Continuants

Within the category of continuants, BFO introduces a further refinement between independent and dependent continuants. Independent continuants are those that can exist on their own, in the sense that their existence does not inhere in another entity. Material objects, organisms, and physical artifacts exemplify independent continuants.

Dependent continuants, in contrast, are entities that inhere in or depend upon independent continuants for their existence. Qualities such as color, mass, and temperature are specifically dependent continuants. They cannot exist without the entities of

which they are qualities. Similarly, roles, dispositions, and functions are treated as dependent continuants, inhering in their bearers while not being reducible to physical structure alone.

BFO further distinguishes between specifically dependent continuants, which depend on a particular bearer, and generically dependent continuants, which may be instantiated across multiple bearers. Information artifacts often fall into the latter category. A digital document, for example, may be realized in multiple physical media while preserving its informational identity.

These refinements aim to prevent ontological conflation. By clearly separating independent entities from the qualities, roles, and informational patterns that depend upon them, BFO enforces clarity in modeling relations between structure and attribute. The architecture ensures that ontological categories correspond to different modes of dependence rather than to superficial linguistic variation.

2.3 Relations and Participation

Relations in BFO are not primitive in the same sense as continuants and occurrents, but they are rigorously constrained. Participation, inherence, and dependence are formally articulated to preserve ontological discipline. A continuant participates in an occurrent; a dependent continuant inheres in an independent continuant; a generically dependent continuant is concretized in a physical bearer.

These relations function as structural connectors between ontological strata. They are designed to prevent illicit cross-category inferences. A process cannot inhere in an object, nor can a quality unfold in time as an occurrent. By sharply delineating the admissible relational structure, BFO reduces category errors and promotes interoperability across domain ontologies.

2.4 Strengths of the BFO Framework

The strength of BFO lies in its disciplined minimalism. By restricting the set of admissible primitives and clarifying their interrelations, it provides a stable foundation for large-scale ontology engineering. Domains characterized by relatively stable identity conditions benefit particularly from this approach. Anatomical structures, chemical entities, and institutional artifacts often exhibit sufficient persistence that entity-centric classification yields durable, interoperable models.

Moreover, BFOs realist stance guards against conceptual relativism. By insisting that ontological categories track mind-independent structure, it discourages ad hoc modeling driven solely by pragmatic convenience. In this respect, BFO performs an entropy-reducing function within representational systems. It limits the proliferation of incompatible ontological commitments and enforces consistency across domains.

Yet this same structure invites further scrutiny. The primacy accorded to continuants and their persistence conditions presupposes that identity is given rather than achieved. Processes are accommodated within the framework, but they are articulated as unfolding entities rather than as constitutive grounds of stability. The possibility that persistence itself may be a derivative phenomenon remains outside the formal architecture.

It is to this structural presuppositionentity primacythat we now turn.

3 Irreversibility and the Limits of Entity Primacy

The entity-centric architecture of the Basic Formal Ontology derives much of its strength from the clarity with which it articulates persistence. Continuants endure; occurrents unfold; dependent entities inhere; independent entities bear. Within domains characterized by stable identity conditions, this structure yields considerable explanatory and engineering power. Yet its underlying presuppositionthat entities are ontologically prior to the histories in which they participateintroduces a structural limitation when confronted with irreversibility.

Irreversibility is not merely the asymmetry of temporal ordering. It is the constitutive feature of physical law, biological development, learning, and computation whereby later states are not simply rearrangements of earlier ones but are constrained by paths taken. Thermodynamic processes, evolutionary trajectories, and epistemic revisions all exhibit path dependence. Once a system has traversed a particular historical trajectory, the space of admissible futures is altered. Identity conditions may persist, but they do so against a background of constrained and irrecoverable change.

Within BFO, time is represented primarily as a dimension along which continuants persist and occurrents extend. The framework accommodates processes, but it does so by treating them as entities that unfold. This formulation captures temporal extension but does not render irreversibility primitive. Processes are described, but the generative role of historical constraint in producing and stabilizing entities remains external to the ontological core. The distinction between continuants and occurrents

presupposes that persistence and unfolding are separable categories, rather than mutually constitutive aspects of constrained historical regimes.

The difficulty becomes more pronounced in domains where stability itself is contingent. In adaptive systems, identity is not a fixed given but an achievement maintained through continuous constraint management. Organisms maintain homeostasis; institutions sustain legitimacy; computational systems preserve invariants under transformation. In such contexts, persistence is not simply endurance through time but active stabilization against entropy. The entity persists because certain historical trajectories remain admissible while others are excluded.

An entity-centric ontology may describe these stabilization processes as occurrents in which continuants participate. However, this descriptive accommodation does not elevate irreversibility to ontological primacy. The fact that certain continuations become impossible after specific transitions is treated as a feature of process description rather than as a fundamental structuring principle of reality. Irreversibility thus remains derivative: it is something that happens to entities rather than something from which entities emerge.

The consequences of this orientation become visible in ontology interoperability. When two ontological systems evolve independently, their categories accumulate distinct historical commitments. Version drift alters definitions, refines scopes, and shifts boundaries. Mappings between such systems presuppose that entities correspond across historical divergence. Yet if identity conditions are themselves stabilized through historical constraint, then static correspondences attempt to align regimes whose admissible futures have diverged. The resulting fragility is not a failure of diligence but a manifestation of misaligned historical substrates.

Similarly, informational entities expose the limits of entity primacy. Information artifacts are often treated as generically dependent continuants. They may be realized in multiple physical bearers while preserving informational identity. Yet informational identity is not simply an intrinsic property; it depends on constraint regimes that maintain pattern stability across transformations. Compression, transmission, and reinterpretation alter the degeneracy of admissible futures associated with an informational structure. Without a principled account of how such stability is achieved and maintained, informational ontology oscillates between abstraction and embodiment without clear grounding.

These observations suggest that persistence may be better understood as a low-entropy condition achieved within constrained historical trajectories rather than as a primitive ontological category. Identity becomes intelligible not as a given but as the stabilization

of admissible continuations. In such a framework, irreversibility is constitutive: it defines the asymmetry through which stabilization occurs. Once a system traverses a particular historical path, certain futures become inaccessible, and others become reinforced. The entity is then recognized as a region of low entropy within this constrained space.

The critique advanced here does not deny the practical efficacy of entity-centric ontologies. Rather, it questions whether their primitives are sufficiently general to account for the full range of phenomena encountered in contemporary ontology engineering. If irreversibility and stabilization are ontologically basic, then an ontology that begins with entities may capture only the stabilized end state of a deeper dynamical process.

The next section introduces an alternative framework in which admissible histories, entropy constraints, and directional flows are treated as primitive. Within this history-first ontology, entities emerge as low-entropy invariants rather than as ontological starting points.

4 A History-First Ontology of Entropic Fields

If irreversibility is constitutive rather than derivative, then ontology must begin not with stabilized entities but with the space of admissible historical trajectories. The primitive question is no longer “What exists?” but “Which continuations remain possible under constraint?” Entities, processes, and relations must then be reconstructed as structured regions within this constrained historical manifold.

Let H denote the space of admissible histories. A history $h \in H$ is defined as an irreversible, temporally ordered trajectory subject to ontological constraints. These constraints are not merely logical consistency conditions but structural restrictions governing which extensions of a partial history remain admissible. The ontology is thus grounded in admissibility rather than in enumeration of objects.

Within this framework, three primitive fields characterize ontological structure: a scalar field Φ , a vector field \mathbf{v} , and an entropy field S . These are not introduced as physical hypotheses but as abstract ontological primitives capturing stability, directionality, and degeneracy within H .

4.1 Scalar Density and Stabilization

The scalar field Φ represents ontic density or stability. Regions of high Φ correspond to historical configurations that persist under perturbation. Such regions exhibit low sensitivity to minor variations in admissible continuation and therefore manifest as stable structures. What classical ontology calls an “object” is here interpreted as a region of sustained high Φ across a family of histories.

Persistence is thus reconstructed as stability across constrained trajectories. An entity does not endure because it is primitive; rather, it is primitive for representational purposes because it occupies a low-entropy, high-density region within H . Stability is an achievement of constrained history rather than a metaphysical starting point.

4.2 Vector Fields and Directed Constraint Propagation

The vector field \mathbf{v} encodes directed constraint propagation. It represents asymmetries in admissible continuation. Causal, inferential, functional, and normative flows are interpreted as manifestations of \mathbf{v} within particular regimes. Directionality is therefore not reduced to temporal ordering but grounded in structural constraint.

Processes are reinterpreted accordingly. Rather than being occurrents that unfold alongside entities, they are directional flows within constrained history space. A process corresponds to a trajectory guided by \mathbf{v} through regions of varying Φ and S . Participation relations, central to entity-centric ontologies, are subsumed under directional constraint coupling: stable flows linking regions of density.

4.3 Entropy as Degeneracy of Futures

The entropy field S measures the degeneracy of admissible futures. Formally, entropy at a historical configuration corresponds to the logarithm of the number of admissible continuations compatible with current constraints. Low entropy signifies restricted continuation and therefore stability; high entropy indicates branching possibility and potential instability.

Crucially, entropy is not interpreted epistemically but ontologically. It does not measure ignorance but structural openness of continuation. In this sense, entropy governs the fragility or robustness of identity conditions. An entity is stable when its continuation remains confined to a narrow band of admissible futures; it destabilizes when degeneracy increases.

4.4 Derived Ontological Categories

Within this field-theoretic ontology, classical categories emerge as stabilized configurations. An object corresponds to a persistent low-entropy region of high scalar density. A process corresponds to a directed flow along \mathbf{v} . A relation emerges as a stable coupling of vector flows between dense regions. Information is interpreted as an entropy-constrained, projectable pattern whose stability permits reliable propagation. Agency appears as a subsystem that actively maintains low entropy within its boundary by regulating constraint flow.

None of these categories are primitive. Each is a historically stabilized configuration within the coupled fields (Φ, \mathbf{v}, S) . Ontology engineering becomes the practice of articulating and managing these constraint regimes so that semantic collapse is prevented and interoperability remains viable.

4.5 Regime-Dependence and Ontological Scope

A history-first ontology does not abolish entity-centric frameworks; it relativizes them. In regimes characterized by low entropy and high scalar density, entities appear stable and primary. In regimes marked by rapid branching, adaptive restructuring, or informational flux, stability becomes contingent and must be explained rather than presupposed.

The advantage of the field-theoretic approach lies in its scalability. By treating admissible histories as primitive and entities as emergent invariants, it accommodates both low-entropy and high-entropy domains within a unified architecture. Interoperability problems, mapping fragility, and semantic drift are reinterpreted as entropy misalignments between regimes rather than as purely linguistic inconsistencies.

In this way, ontology shifts from static classification to dynamical constraint management. The question becomes not merely whether a category corresponds to reality, but whether it stabilizes admissible histories across scales and domains.

Having articulated this alternative framework, we now return to the Basic Formal Ontology in order to demonstrate how it may be embedded as a low-entropy subtheory within the broader entropic field ontology.

5 Embedding the Basic Formal Ontology as a Low-Entropy Subtheory

The critique advanced in the preceding sections does not entail the rejection of the Basic Formal Ontology. On the contrary, the structural discipline of BFO can be preserved within a broader history-first framework. The claim is not that entity-centric ontology is false, but that it is regime-specific. When entropy remains bounded and scalar density is sufficiently high, entity primacy is not only defensible but operationally optimal.

Within the scalarvectorentropy ontology, a low-entropy regime is characterized by restricted admissible continuation. The entropy field S remains uniformly bounded within a narrow band; the scalar field Φ exhibits stable attractors; directional constraint propagation \mathbf{v} is laminar rather than turbulent. In such regimes, identity conditions remain robust under perturbation. Entities exhibit persistence not as an unexplained given but as a consequence of structural stabilization.

Under these conditions, the BFO distinction between continuants and occurrents acquires a natural interpretation. Continuants correspond to regions of sustained high scalar density whose admissible futures remain tightly constrained. Occurrents correspond to directional flows within this stabilized region. Participation relations express stable couplings between dense regions and constrained flows. In this reinterpretation, BFOs primitives are not replaced but derived as low-entropy invariants of the coupled fields.

This embedding clarifies why BFO performs effectively in domains such as anatomy, chemistry, and regulated institutional systems. These domains exhibit strong stabilization mechanisms. Biological organisms maintain homeostasis; molecular structures possess well-defined bonding constraints; institutional artifacts are governed by explicit normative frameworks. The entropy of admissible futures within these systems is restricted by physical, biological, or legal constraints. As a result, entity-centric classification captures genuine structural invariants.

However, the embedding also clarifies the limits of BFO. In regimes where entropy increases where branching possibilities proliferate and identity conditions shift rapidly the assumption of stable continuants becomes fragile. Adaptive learning systems, evolving computational architectures, and socially negotiated institutional meanings may not maintain narrow entropy bands. Under such conditions, entities are transient attractors rather than enduring primitives. A field-theoretic description remains applicable, but the derived entity categories may fluctuate or dissolve.

The embedding relationship may be formalized as follows. Let (Φ, \mathbf{v}, S) denote a field configuration over history space H . Suppose there exists a region $R \subset H$ such that S remains bounded above by a small constant ϵ and Φ exhibits stable local maxima across R . Within R , identity conditions persist across admissible continuations. One may then define a derived ontology in which such stable maxima are treated as continuants and the associated flows as occurrents. The resulting entity-centric ontology is valid so long as the entropy bound holds.

Thus BFO can be interpreted as a low-entropy subtheory embedded within a more general entropic ontology. Its primitives correspond to stabilized configurations rather than to metaphysical atoms. The distinction between continuant and occurrent remains meaningful, but its scope is recognized as conditional upon regime stability.

This perspective transforms the perceived conflict between entity-centric and process-centric ontologies. The dispute is not over which category is ontologically fundamental in all circumstances, but over the regime within which modeling occurs. In low-entropy regimes, entity primacy yields clarity and interoperability. In high-entropy regimes, stabilization must be explained rather than presupposed, and a history-first ontology provides the necessary generality.

By embedding BFO rather than discarding it, the scalarvectorentropy framework preserves realist discipline while extending ontology to domains characterized by irreversibility, adaptive constraint, and semantic drift. Ontology engineering becomes the art of recognizing regime conditions and selecting primitives appropriate to the entropy profile of the system under study.

6 Implications for Ontology Engineering and Interoperability

Reconceiving ontology in terms of admissible histories and entropic regimes has direct implications for ontology engineering practice. If entities are stabilized invariants within constrained historical fields, then ontology design must account explicitly for regime conditions rather than presupposing universal stability. Engineering decisions become questions of constraint placement, entropy tolerance, and stabilization strategy.

In traditional upper-level ontology, interoperability is pursued through alignment of entity categories. Mappings are constructed between classes presumed to correspond across systems. These mappings assume that identity conditions remain sufficiently stable across independently evolving ontologies. Yet when the underlying historical

regimes diverge, such correspondences amplify entropy rather than reduce it. Version drift, scope refinement, and local adaptation introduce branching continuations that static mappings cannot absorb.

Within a history-first ontology, interoperability is reinterpreted as synchronization of constraint regimes rather than alignment of static entities. A successful mapping reduces entropy across systems by constraining admissible continuations toward convergence. Once convergence is achieved, the mapping becomes redundant; the systems share sufficient structural invariants that explicit correspondence operators are no longer required. Mapping fragility is thus explained not as technical failure but as entropy amplification within misaligned regimes.

This shift alters how ontology engineers evaluate success. Instead of measuring adequacy solely by consistency and coverage, one evaluates whether ontological commitments stabilize admissible histories across scale. An ontology that proliferates ad hoc categories may satisfy immediate modeling needs but increase entropy in future extensions. Conversely, an ontology that enforces overly rigid primitives may reduce entropy at the cost of adaptability, impeding legitimate evolution.

The scalarvectorentropy framework provides conceptual tools for articulating these tradeoffs. Scalar density corresponds to the inertia of ontological commitments. High-density regimes resist change and favor stability; low-density regimes permit rapid adaptation but risk fragmentation. Directed constraint flow captures the pathways through which revisions propagate across dependent systems. Entropy measures the degeneracy of future modeling trajectories permitted by present commitments.

Under this perspective, governance of top-level ontologies becomes an exercise in entropy management. Conservative regimes prioritize low entropy and high scalar density, favoring long-term stability. Agile regimes tolerate higher entropy to accommodate emerging domains and evolving practices. Neither orientation is universally correct; each reflects a distinct placement of constraints within history space.

This analysis also reframes debates concerning artificial intelligence, information ontology, and institutional modeling. Questions about the ontological status of informational entities reduce to questions about the stability of entropy-constrained patterns. Debates about machine intelligence concern whether artificial systems can maintain low-entropy regimes analogous to biological agents. Disputes over ontology evolution concern the acceptable rate of entropy increase relative to scalar stabilization.

The history-first ontology therefore extends the project of upper-level ontology rather than abandoning it. It preserves realist discipline while situating entity categories within a broader dynamical architecture. By treating irreversibility, constraint propa-

gation, and entropy as primitive, it provides a framework capable of accommodating both stable scientific domains and rapidly evolving computational or cognitive systems.

Ontology engineering, on this view, is not merely the classification of what exists. It is the deliberate management of admissible histories. Its task is to constrain representational trajectories so that identity, meaning, and interoperability remain viable across time. The question guiding ontology design thus becomes not only what entities populate the world, but which histories remain possible under the commitments we choose to encode.

7 Formal Comparison with Existing Upper-Level Ontologies

The preceding sections have advanced a history-first ontology grounded in scalar density, directed constraint flow, and entropy as degeneracy of admissible futures. In order to clarify its scope and its relation to existing work, it is necessary to compare this framework formally with established upper-level ontologies. The aim of this comparison is not polemical contrast but structural clarification.

7.1 Comparison with the Basic Formal Ontology

The Basic Formal Ontology enforces a bifurcation between continuants and occurrents and refines this distinction through dependence relations. Formally, one may represent BFO as a typed domain (E, P, D) where E denotes continuants, P denotes occurrents, and D encodes dependence and participation relations between these types. Identity conditions for E are taken as primitive, and temporal extension is expressed through relations linking E and P .

In the scalarvectorentropy framework, let H denote the space of admissible histories, and let (Φ, \mathbf{v}, S) represent the coupled fields governing stability, directionality, and degeneracy. An entity e in the BFO sense corresponds to a region $R_e \subset H$ such that the scalar field Φ attains a persistent local maximum across admissible continuations and the entropy field S remains bounded above by a regime-specific threshold ϵ . A process corresponds to a directed trajectory through H guided by \mathbf{v} within or between such regions.

Under this mapping, the BFO primitive distinction is recovered as a low-entropy invariant within the broader dynamical system. Continuants correspond to stabilized

regions; occurrents correspond to constrained flows. Participation relations correspond to stable couplings between dense regions and directional trajectories. The formal difference lies in ontological priority: BFO treats (E, P, D) as primitive, whereas the scalarvectorentropy ontology treats H and its constraint fields as primitive, deriving (E, P, D) under bounded-entropy conditions.

This difference yields distinct explanatory capacities. BFO provides clarity in domains where identity conditions are stable and entropy remains low. The field-theoretic ontology extends this clarity to regimes in which stability itself must be explained. The relationship is therefore one of embedding rather than opposition: BFO is recoverable as a special case under entropy constraints.

7.2 Comparison with DOLCE and Process Ontologies

Process-oriented upper ontologies such as DOLCE emphasize events and temporal unfolding more explicitly than entity-centric frameworks. They refine the ontology of occurrences, qualities, and participation structures in ways that foreground dynamical aspects of reality. Formally, such ontologies enrich the occurrent domain and articulate relations between events and endurants.

However, even process-centric frameworks typically retain identity conditions for their primitives as given. Events are categorized and related, but the stabilization of event types across histories is not treated as an ontological problem. The entropy of admissible continuation remains implicit. Without a principled measure of degeneracy or constraint, process ontologies describe unfolding but do not quantify regime stability.

In contrast, the scalarvectorentropy ontology integrates process and persistence through a shared dynamical substrate. Processes are not simply occurrents but directional flows within a constrained historical manifold. Stability and change are co-articulated through Φ and \mathbf{v} , while S regulates admissible branching. This integration permits explicit modeling of regime shifts, bifurcations, and entropy amplification phenomena that remain external to purely classificatory process ontologies.

8 Logic, Statistics, and Ontological Commitments in Artificial Intelligence

Debates within artificial intelligence provide a concrete illustration of the ontological tensions discussed in the preceding section. The contrast between Good Old-Fashioned

AI (GOFAI) and contemporary stochastic models is not merely methodological. It reflects divergent assumptions about knowledge, structure, and the role of ontology in reasoning systems.

GOFAI emerged from a commitment to explicit logical structure. Its foundational premise was that intelligent behavior could be achieved through formal manipulation of symbolic representations governed by logical axioms. Ontologies were treated as explicit rule sets encoding the structure of the world. Inference proceeded deterministically from these axioms, typically within first-order logic or closely related formal systems. The systems reasoning was transparent in principle: conclusions followed from premises via formally valid derivations.

This approach required that common sense be representable as a stable collection of entities, relations, and axioms. Researchers attempted to enumerate and encode these structures directly. Ontology, in this context, was not an auxiliary engineering tool but the central mechanism of intelligence. Knowledge representation was therefore primary. Learning, where present at all, played a subordinate role to explicit formalization.

Contemporary stochastic AI systems operate under a radically different premise. Rather than beginning with hand-crafted logical primitives, modern models rely on statistical regularities extracted from large corpora. Their internal representations are not explicit ontological axioms but distributed patterns in high-dimensional parameter spaces. Inference is probabilistic rather than deductive. Outputs are sampled from learned distributions conditioned on input prompts.

This statistical paradigm replaces explicit ontology with implicit structure. Instead of encoding common sense directly, the model approximates patterns of usage observed in data. The representations it learns may correlate with ontological distinctions, but such distinctions are emergent rather than prescribed. Reasoning becomes an act of probabilistic continuation rather than logical derivation.

The divergence between these paradigms reflects deeper ontological assumptions. GOFAI presupposed that the world could be captured through stable entity categories and deterministic rules. Its failures were often described as brittleness: systems performed well within narrow domains but collapsed under slight deviations from expected conditions. The rigidity of explicit axioms amplified modeling errors when confronted with the combinatorial vastness of human common sense.

Modern stochastic systems exhibit a complementary profile. They are flexible across domains, capable of generating plausible responses in varied contexts. Yet this flexibility arises from probabilistic continuation rather than logical constraint. The absence of hard ontological boundaries permits broad descriptive power but introduces

instability in the form of hallucinations. Because outputs are drawn from learned distributions rather than derived from axioms, the system may generate coherent but false continuations.

These contrasting architectures illustrate two distinct strategies for entropy management. GOFAI attempted to minimize entropy by sharply constraining admissible inferences through explicit axioms. The result was low degeneracy within defined domains but catastrophic instability outside them. Stochastic AI tolerates higher entropy, permitting a wide range of continuations, but sacrifices deterministic guarantees of correctness. Stability becomes statistical rather than logical.

9 Historical Cycles and Ontological Regimes

The historical trajectory of artificial intelligence reveals recurring cycles of expansion and contraction, often described as periods of hype followed by “AI winters.” The first major cycle, associated with early symbolic systems, was fueled by optimism that logical representation could scale to general intelligence. When these systems failed to meet ambitious expectations particularly in military and strategic planning contexts funding contracted and enthusiasm waned.

The present era represents a third major expansion, driven by large-scale stochastic models and unprecedented computational resources. These systems have achieved remarkable performance in language modeling and pattern recognition tasks. Yet they are also characterized by substantial training costs, infrastructural centralization, and persistent error modes. Concerns about financial sustainability, energy expenditure, and epistemic reliability suggest that another contraction may be possible.

From the perspective of entropic ontology, these cycles can be interpreted as oscillations between regimes of constraint. Symbolic systems imposed strong logical constraints, resulting in low entropy but limited adaptability. Stochastic systems permit high entropy in internal representation and output generation, trading determinism for flexibility. Each regime encounters scaling limits when its entropy profile becomes misaligned with the domains it attempts to model.

The lesson for upper ontology is not that one paradigm must supplant the other. Rather, it is that ontology must be capable of articulating both low-entropy, rule-governed regimes and high-entropy, probabilistic regimes within a unified framework. An entity-centric ontology aligns naturally with deterministic symbolic systems. A history-first, entropic ontology can account for the adaptive, statistical character of

contemporary models without abandoning realist discipline.

By situating artificial intelligence within the broader architecture of admissible histories and entropy management, we obtain a principled account of both its successes and its instabilities. GOFAI and stochastic AI are not simply technological alternatives; they instantiate distinct placements of constraint within history space. Understanding these placements clarifies the role that upper-level ontology must play in future intelligent systems.

9.1 Comparison with Information-Centric Ontologies

Information-centric ontologies attempt to elevate informational entities to primary status. They distinguish abstract informational content from its physical carriers and attempt to articulate relations between symbol structures and referents. Yet such ontologies often oscillate between realism and conceptualism. Informational identity is asserted, but the criteria for its stabilization across transformation remain under-specified.

Within the scalarvectorentropy framework, informational structure is interpreted as an entropy-constrained, projectable pattern. Let I denote a pattern over H . I qualifies as informationally stable if the degeneracy of admissible futures compatible with I remains bounded and if scalar density persists across realizations. Informational identity thus depends on the maintenance of low entropy under projection and transmission.

This formulation grounds information in constraint rather than in abstraction alone. It avoids conceptualism by treating informational stability as a structural feature of history space rather than as a mental construct. At the same time, it explains why informational entities may be multiply realizable: distinct physical realizations correspond to trajectories within the same low-entropy region of H .

9.2 Regime-Relativity and Ontological Scope

Across these comparisons, a unifying theme emerges. Existing upper-level ontologies articulate valuable structural distinctions, but they do so under implicit regime assumptions. Entity-centric frameworks presuppose low entropy and high scalar stability. Process-centric frameworks presuppose stable event types. Information-centric frameworks presuppose projectable patterns without formalizing degeneracy conditions.

The scalarvectorentropy ontology makes these regime conditions explicit. It treats entropy bounds, stabilization thresholds, and directional constraint propagation as primitive features of ontology rather than as background assumptions. As a result, it subsumes existing frameworks as specializations valid within appropriate regions of history space.

The formal comparison therefore supports a pluralistic but structured conclusion. Upper-level ontologies need not compete for exclusive metaphysical primacy. Instead, they may be understood as distinct projections of a deeper dynamical substrate. By grounding ontology in admissible histories and entropic constraint, the scalarvectorentropy framework provides a unifying architecture within which entity-centric, process-centric, and information-centric approaches can be situated without mutual exclusion.

10 Conclusion

Upper-level ontology arose from a genuine need: the stabilization of meaning across heterogeneous scientific and institutional systems. Frameworks such as the Basic Formal Ontology have demonstrated that realist discipline and carefully articulated primitives are indispensable for large-scale interoperability. By enforcing distinctions between continuants and occurrents and by clarifying dependence relations, BFO has provided an architecture capable of reducing semantic fragmentation within domains characterized by stable identity conditions.

The critique developed in this paper does not deny these achievements. Rather, it questions whether entity primacy is sufficiently general to serve as the ultimate ontological ground. Persistent difficulties surrounding irreversibility, adaptive systems, informational stability, and mapping fragility suggest that identity itself may require explanation. When stability is treated as primitive, the historical and entropic conditions under which it emerges remain unarticulated. As ontology engineering extends into domains of rapid evolution, distributed computation, and institutional flux, these conditions become increasingly central.

The history-first ontology advanced here inverts the order of explanation. Instead of beginning with entities and describing processes as what happens to them, it begins with admissible histories and treats entities as low-entropy invariants within constrained trajectories. Scalar density captures stabilization; vector flow captures directional constraint propagation; entropy measures the degeneracy of admissible futures. Within low-entropy regimes, classical entity-centric categories emerge naturally. In such

regimes, the BFO distinction between continuants and occurrents is not undermined but derived.

This embedding relationship clarifies the scope of existing upper ontologies. They remain valid within domains where entropy is bounded and stabilization mechanisms are robust. Their limitations arise not from internal inconsistency but from implicit regime assumptions. By making entropy and constraint primitive, the scalarvectorentropy framework generalizes upper ontology to domains in which stability must be achieved rather than assumed.

The broader implication is that ontology engineering is not merely classificatory but dynamical. It is the practice of constraining representational histories so that meaning, identity, and interoperability remain viable across time. The success of an ontology is measured not only by internal coherence but by its capacity to maintain low-entropy regimes under extension and revision.

Upper ontology therefore remains necessary. What changes is the understanding of its foundations. Entities are no longer treated as metaphysical atoms but as stabilized regions within a field of admissible histories. Ontological realism is preserved, but its grounding shifts from static enumeration to structural constraint. In this way, the discipline inaugurated by frameworks such as BFO is retained while its scope is extended to encompass irreversibility, adaptive complexity, and informational dynamics.

The task of future ontology engineering is thus twofold: to preserve the stabilizing discipline of upper-level categories and to articulate explicitly the entropic and historical conditions under which such categories remain valid. Only by integrating stability and change within a unified formal architecture can ontology keep pace with the dynamical systems it seeks to describe.

Appendices

A Admissible Histories and Entropy Bounds

A.1 Admissible History Space

Let H denote the space of admissible histories. A history $h \in H$ is defined as a temporally ordered sequence of states

$$h = (x_0, x_1, x_2, \dots)$$

such that each transition $x_t \rightarrow x_{t+1}$ satisfies a constraint operator

$$\mathcal{C}(x_t, x_{t+1}) = 1.$$

The constraint operator \mathcal{C} encodes structural admissibility. It is not restricted to logical consistency but may represent physical, biological, institutional, or informational constraints.

Irreversibility is captured by the asymmetry of \mathcal{C} :

$$\mathcal{C}(x_t, x_{t+1}) = 1 \Rightarrow \mathcal{C}(x_{t+1}, x_t) = 1.$$

The space H is therefore a directed graph or category whose morphisms represent admissible transitions.

A.2 Entropy as Degeneracy of Futures

For a given state x_t , define the admissible continuation set

$$\mathcal{A}(x_t) = \{x_{t+1} \mid \mathcal{C}(x_t, x_{t+1}) = 1\}.$$

The entropy field S at x_t is defined as

$$S(x_t) = \log |\mathcal{A}(x_t)|.$$

This entropy measures the degeneracy of admissible futures rather than epistemic uncertainty. A low value of S indicates tightly constrained continuation; a high value

indicates branching possibility.

A region $R \subset H$ is said to satisfy an entropy bound ϵ if

$$\sup_{x \in R} S(x) \leq \epsilon.$$

A.3 Scalar Density and Stability

Define a scalar stability field $\Phi : H \rightarrow \mathbb{R}_{\geq 0}$ such that $\Phi(x)$ measures persistence under perturbation. Formally, $\Phi(x)$ may be defined as a function of the local entropy gradient:

$$\Phi(x) = f\left(-\frac{dS}{dt}\Big|_x\right),$$

where f is monotone increasing.

Intuitively, high scalar density corresponds to regions where entropy remains bounded and variations in admissible continuation are damped. A region R is stable if $\Phi(x)$ attains a local maximum for all $x \in R$ under admissible perturbations.

An entity in the derived ontology corresponds to a connected component $R \subset H$ satisfying

$$\sup_{x \in R} S(x) \leq \epsilon \quad \text{and} \quad \Phi(x) \geq \delta > 0.$$

A.4 Vector Field and Directed Constraint Flow

Let \mathbf{v} denote a vector field on H assigning to each state x a preferred direction of admissible transition:

$$\mathbf{v}(x) \in T_x H.$$

A process corresponds to an integral curve $\gamma(t)$ satisfying

$$\frac{d\gamma}{dt} = \mathbf{v}(\gamma(t)),$$

subject to $\mathcal{C}(\gamma(t), \gamma(t + \Delta t)) = 1$.

Stable participation relations in the derived ontology correspond to persistent couplings between integral curves and low-entropy regions.

A.5 Low-Entropy Subtheories

Let $R \subset H$ be a region satisfying the entropy bound and stability condition above. Define a projection

$$\pi : R \rightarrow E$$

mapping stable connected components to equivalence classes interpreted as continuants. Directed trajectories within R are mapped to occurrents.

Under these conditions, an entity-centric ontology (E, P, D) may be constructed as a quotient structure over R . The validity of this quotient depends on maintenance of the entropy bound. If S exceeds the threshold ϵ , the projection π ceases to preserve identity conditions.

This formalization renders precise the embedding claim advanced in the main text: entity-centric upper ontologies are valid within regions of bounded entropy and scalar stabilization, but their primitives are derived from a more general historical field structure.

B Regime Structure and Identity Conditions

B.1 History Space as a Directed Category

Let H be a small category whose objects are admissible states and whose morphisms are irreversible transitions.

Definition 1. *A history is a functor*

$$h : \mathbb{N} \rightarrow H$$

such that for each t , the morphism $h(t) \rightarrow h(t + 1)$ exists in H .

Irreversibility is encoded by non-invertibility of morphisms.

Definition 2. *A regime $R \subset H$ is a full subcategory closed under admissible morphisms.*

B.2 Entropy Bounds as Subcategory Constraints

Let $\mathcal{A}(x)$ denote the set of outgoing morphisms from object x .

Definition 3. *The entropy at x is*

$$S(x) = \log |\mathcal{A}(x)|.$$

Definition 4. *A regime R satisfies entropy bound ϵ if*

$$\sup_{x \in R} S(x) \leq \epsilon.$$

Proposition 1. *If R satisfies $S(x) = 0$ for all $x \in R$, then R is a thin category.*

Proof. If $S(x) = 0$, then $|\mathcal{A}(x)| = 1$. Hence there is at most one outgoing morphism per object. Thus R is thin. \square

B.3 Scalar Density as Stability Functional

Define a functional

$$\Phi : \text{Ob}(H) \rightarrow \mathbb{R}_{\geq 0}.$$

Definition 5. *A region R is δ -stable if*

$$\inf_{x \in R} \Phi(x) \geq \delta.$$

Definition 6. *An invariant region is a connected component R satisfying both*

$$\sup_{x \in R} S(x) \leq \epsilon \quad \text{and} \quad \inf_{x \in R} \Phi(x) \geq \delta.$$

B.4 Derived Identity Conditions

Definition 7. *Two states $x, y \in H$ are identity-equivalent, written $x \sim y$, if they lie in the same invariant region.*

Proposition 2. *The relation \sim is an equivalence relation.*

Proof. Reflexivity and symmetry follow from set membership in invariant regions. Transitivity follows from connectedness of invariant regions. \square

Definition 8. *An entity is an equivalence class under \sim .*

B.5 Process Structure

Definition 9. *A process is a non-constant morphism chain*

$$x_0 \rightarrow x_1 \rightarrow \cdots \rightarrow x_n$$

within a regime R .

Proposition 3. *If R is invariant and entropy-bounded, then processes preserve identity class.*

Proof. If entropy remains bounded and scalar density exceeds δ , then all states along the chain lie within the same invariant region, hence same equivalence class. \square

B.6 Regime Transition

Definition 10. *A bifurcation point is an object x such that*

$$|\mathcal{A}(x)| > 1.$$

Definition 11. *A regime shift occurs when a trajectory exits an invariant region.*

Proposition 4. *Identity failure corresponds to regime shift.*

Proof. If a trajectory leaves an invariant region, either entropy exceeds ϵ or scalar density falls below δ . Hence equivalence class is no longer preserved. \square

C Analysis of Logical and Stochastic AI

C.1 Logical AI as a Zero-Entropy Regime

Let \mathcal{O} be a finite set of first-order axioms over signature Σ . Let $\mathcal{M}(\mathcal{O})$ denote the class of models satisfying \mathcal{O} .

Define the admissible continuation operator

$$\mathcal{A}_{\mathcal{O}}(x) = \{y \mid y \models \mathcal{O}\}.$$

Definition 12. *The entropy of a logical regime at state x is*

$$S_{\mathcal{O}}(x) = \log |\mathcal{A}_{\mathcal{O}}(x)|.$$

Proposition 5. *If \mathcal{O} is complete and consistent, then for all admissible x ,*

$$S_{\mathcal{O}}(x) = 0.$$

Proof. Completeness implies that all sentences are either derivable or refutable. Consistency ensures model non-emptiness. Hence admissible continuations are uniquely determined up to isomorphism. Therefore $|\mathcal{A}_{\mathcal{O}}(x)| = 1$. \square

Logical AI corresponds to a zero-entropy RSVP regime.

C.2 Brittleness as Entropy Discontinuity

Let \mathcal{C} be a constraint operator induced by \mathcal{O} .

Definition 13. *A brittleness point is a state x such that for perturbation δx ,*

$$\mathcal{C}(x, x + \delta x) = 0.$$

Proposition 6. *In zero-entropy regimes, any constraint violation produces total inadmissibility.*

Proof. If $S = 0$, admissible continuation set is singleton. Any violation removes the unique admissible continuation. Hence no admissible successor exists. \square

D Categorical Embedding and Homotopy Structure

D.1 Embedding of Entity-Centric Ontology

Let H be the directed category of admissible histories and let $\mathcal{R} \subset H$ be an invariant region.

Let **Ent** denote the category whose objects are entities and whose morphisms are participation or dependence relations.

Definition 14. Define a functor

$$F : \mathcal{R} \rightarrow \mathbf{Ent}$$

such that each invariant connected component of \mathcal{R} is mapped to a single object of \mathbf{Ent} .

Proposition 7. If \mathcal{R} satisfies entropy bound ϵ and scalar lower bound δ , then F is well-defined.

Proof. Invariant regions correspond to equivalence classes under \sim . Each class is mapped to a unique object. Morphisms internal to the region preserve identity; hence functoriality holds. \square

Definition 15. An entity-centric ontology is an image category $\text{Im}(F)$ for some invariant region \mathcal{R} .

D.2 Sheaf Structure and Ontology Alignment

Let $\{R_i\}$ be a cover of H by invariant regimes.

Definition 16. Define a presheaf \mathcal{F} over H such that for each regime R_i ,

$$\mathcal{F}(R_i) = \text{Im}(F_i),$$

where $F_i : R_i \rightarrow \mathbf{Ent}_i$ is the local embedding.

Restriction maps are given by inclusion of subregions.

Definition 17. \mathcal{F} is a sheaf if for any compatible family $\{s_i \in \mathcal{F}(R_i)\}$ agreeing on overlaps $R_i \cap R_j$, there exists a unique global section $s \in \mathcal{F}(\bigcup R_i)$.

Proposition 8. Ontology interoperability corresponds to existence of global sections of \mathcal{F} .

Proof. If local ontological embeddings agree on overlaps, a global consistent ontology exists. Failure of gluing indicates entropy amplification across regime boundaries. \square

D.3 Homotopy-Type of Regime Structure

Let $|H|$ denote the geometric realization of H .

Definition 18. A regime R has trivial homotopy type if $|R|$ is contractible.

Definition 19. A regime transition is homotopy-nontrivial if inclusion

$$R \hookrightarrow H$$

induces nontrivial change in homotopy groups.

Proposition 9. Bifurcation points correspond to branching in $|H|$ producing nontrivial π_1 or higher homotopy groups.

Proof. If $|\mathcal{A}(x)| > 1$, then geometric realization contains branching paths. Non-contractible loops may arise when paths reconverge. Hence homotopy nontriviality. \square

D.4 Derived Stack Interpretation

Let \mathcal{S} be a stack over H assigning to each regime R the groupoid of admissible entity embeddings.

Definition 20. A derived entity is an object in the homotopy limit

$$\text{holim}_{R \subset H} \mathcal{S}(R).$$

Proposition 10. Entity identity across regime shifts requires invariance under homotopy equivalence in \mathcal{S} .

Proof. If regime transition alters homotopy type, identity persists only if corresponding objects remain equivalent in homotopy limit. \square

D.5 Entropy Gradient and Morse Structure

Assume entropy function $S : H \rightarrow \mathbb{R}$ is smooth on $|H|$.

Definition 21. A stabilization point is a critical point of S satisfying

$$\nabla S = 0, \quad \text{Hessian}(S) \text{ positive definite.}$$

Proposition 11. Invariant regions correspond to neighborhoods of local minima of S .

Proof. Local minima imply bounded entropy and stability under perturbation. Hence scalar density maximal. \square

Theorem 1. *Regime transitions correspond to Morse bifurcations in entropy landscape.*

Proof. If Hessian signature changes, topology of sublevel sets changes, inducing homotopy-type transition in $|H|$. \square

D.6 Stochastic AI as High-Entropy Regime

Let $\theta \in \mathbb{R}^n$ denote model parameters. Let $p_\theta(y | x)$ be a conditional distribution.

Definition 22. *The local entropy of the stochastic regime at input x is*

$$S_\theta(x) = - \sum_y p_\theta(y | x) \log p_\theta(y | x).$$

Proposition 12. *If p_θ has full support over output space, then*

$$S_\theta(x) > 0.$$

Definition 23. *A hallucination event occurs when*

$$\exists y \in \text{supp}(p_\theta(\cdot | x)) \text{ such that } y \not\models \mathcal{W},$$

where \mathcal{W} denotes world-consistency constraints.

Proposition 13. *If entropy exceeds a threshold ϵ , hallucination probability is strictly positive.*

Proof. If $S_\theta(x) > \epsilon$, then p_θ assigns non-zero mass to multiple incompatible continuations. If \mathcal{W} excludes at least one such continuation, probability of violation is non-zero. \square

D.7 Gradient Flow and Entropy Suppression

Let parameter dynamics follow

$$d\theta_t = -\nabla_\theta \mathcal{L}(\theta_t) dt + \sigma dW_t.$$

Define hypothesis entropy

$$S_t = \log |\mathcal{H}_t|,$$

where \mathcal{H}_t is the hypothesis set compatible with θ_t .

Proposition 14. *If $\nabla_\theta \mathcal{L}$ eliminates incompatible hypotheses monotonically, then*

$$\frac{d}{dt} S_t \leq 0.$$

Definition 24. *Overfitting occurs when $S_t \rightarrow 0$ while true admissible hypothesis space \mathcal{H}^* satisfies $|\mathcal{H}^*| > 1$.*

D.8 Regime Transition Theorem

Let R_0 be a zero-entropy logical regime and R_s a stochastic regime with entropy function S_θ .

Theorem 2. *There exists no entropy-preserving homomorphism*

$$\phi : R_s \rightarrow R_0$$

unless $S_\theta(x) = 0$ for all x .

Proof. Entropy preservation requires $S_\theta(x) = S_{\mathcal{O}}(\phi(x))$. Since $S_{\mathcal{O}} = 0$, equality implies $S_\theta(x) = 0$. \square

D.9 Weak vs Strong AI Formalization

Let H denote the space of admissible cognitive histories.

Definition 25. *A system is weak if its admissible history set H_w is a strict subset of full cognitive space H .*

Definition 26. *A system is strong if $H_s = H$.*

Proposition 15. *Current LLM regimes satisfy $H_w \subsetneq H$.*

Proof. LLM continuation is constrained to statistical training distribution and lacks arbitrary domain transfer with preserved scalar density across all cognitive regimes. \square

References

- [1] R. Arp, B. Smith, and A. D. Spear, *Building Ontologies with Basic Formal Ontology*, MIT Press, 2015.
- [2] B. Smith and W. Ceusters, “Ontology as the Core Discipline of Biomedical Informatics,” *Studies in Health Technology and Informatics*, 2003.
- [3] J. McCarthy and P. J. Hayes, “Some Philosophical Problems from the Standpoint of Artificial Intelligence,” in *Machine Intelligence 4*, Edinburgh University Press, 1969.
- [4] J. McCarthy, “Circumscription — A Form of Non-Monotonic Reasoning,” *Artificial Intelligence*, 13(12), 1980.
- [5] N. Guarino and C. Welty, “An Overview of OntoClean,” in *Handbook on Ontologies*, Springer, 2009.
- [6] R. Brachman and H. Levesque, *Knowledge Representation and Reasoning*, Morgan Kaufmann, 2004.
- [7] C. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, 27(3), 1948.
- [8] A. Kolmogorov, “Three Approaches to the Quantitative Definition of Information,” *Problems of Information Transmission*, 1965.
- [9] D. Rumelhart, G. Hinton, and R. Williams, “Learning Representations by Back-Propagating Errors,” *Nature*, 323, 1986.
- [10] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.
- [11] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2018.
- [12] S. Mac Lane, *Categories for the Working Mathematician*, Springer, 1998.
- [13] F. Borceux, *Handbook of Categorical Algebra*, Cambridge University Press, 1994.
- [14] M. Kashiwara and P. Schapira, *Sheaves on Manifolds*, Springer, 1990.
- [15] J. Lurie, *Higher Topos Theory*, Princeton University Press, 2009.
- [16] M. Hirsch, *Differential Topology*, Springer, 1976.
- [17] J. Milnor, *Morse Theory*, Princeton University Press, 1963.

- [18] I. Prigogine, *From Being to Becoming*, W. H. Freeman, 1980.
- [19] E. T. Jaynes, *Probability Theory: The Logic of Science*, Cambridge University Press, 2003.