

Exolytica AI Technical Report

XTR-2025-01

How to Tame Your LLM: Semantic Collapse in Continuous Systems

(Extended Version)

C. M. Wyss, Ph.D.

Exolytica AI

December 2025

© 2025 C. M. Wyss.

This work is licensed under a Creative Commons Attribution 4.0
International License (CC BY 4.0). To view a copy of this license, visit:

<http://creativecommons.org/licenses/by/4.0/>.

How to Tame Your LLM: Semantic Collapse in Continuous Systems

C. M. Wyss^{*†}

December 2025

Abstract

We develop a general theory of *semantic dynamics* for LLMs by formalizing them as *Continuous State Machines (CSMs)*: smooth dynamical systems whose latent manifolds evolve under probabilistic transition operators. The associated *transfer operator* $P : L^2(M, \mu) \rightarrow L^2(M, \mu)$ encodes the propagation of semantic mass, and under mild regularity assumptions (compactness, ergodicity, bounded Jacobian) is compact with discrete spectrum. Within this setting, we prove the *Semantic Characterization Theorem (SCT)*: the leading eigenfunctions of P define finitely many *spectral basins* of invariant meaning, each definable in an o-minimal structure over $\mathbb{R}$; thus, spectral lumpability and logical tameness are equivalent. The result implies that the continuous manifold of model activations collapses into a finite, logically interpretable ontology, essentially explaining how discrete symbolic semantics emerge from continuous computation. We further extend the SCT to stochastic and time-inhomogeneous (adiabatic) settings, showing that slowly drifting kernels preserve compactness and spectral coherence.

1 Introduction

Large language models have achieved remarkable fluency in natural language understanding and generation, yet the mechanisms underlying their semantic behavior remain poorly understood. These systems operate in high-dimensional continuous spaces, manipulating latent representations that appear, on the face of it, to be fundamentally incompatible with the discrete, symbolic structure of human language and reasoning. This apparent contradiction raises a foundational question: how can a continuous manifold of representations give rise to stable, categorical meanings which behave as the very basis of communication and thought?

We show that this tension is resolved through a phase transition in the geometry of semantic space. Although the latent representations of large language models (LLMs) are continuous, their functional behavior (i.e., the actual semantic distinctions they encode and preserve) is equivalent to that of a discrete symbolic system. This equivalence is neither approximate nor metaphorical, but a rigorous mathematical fact about the topology of semantic space itself. The result, which we call the Semantic Characterization Theorem (SCT), provides a formal bridge between continuous learning dynamics and symbolic semantics, offering both a theoretical foundation for understanding LLM behavior and practical insights into semantic navigation (i.e., chat trajectory), prompt design, and interpretability.

^{*}Exolytica AI (Email: cmw@exolytica.ai).

[†]Assisted by GPT-5 Thinking (OpenAI), Claude Sonnet 4.5 (Anthropic), and Gemini 3 Pro (Google).

At the core of our analysis lies the latent semantic space, $\mathcal{M}$: the high-dimensional manifold on which a model’s representations evolve during language processing. Each prompt or token update can be viewed as an application of a smooth transition operator $T : \mathcal{M} \rightarrow \mathcal{M}$, governing the flow of semantic states within this space. We examine the behavior of T from two independent perspectives. In the dynamical view, we study its spectral decomposition and show that beyond a critical scale, the spectrum collapses onto a finite set of dominant modes corresponding to stable semantic attractors. In the logical view, we interpret the resulting partition of $\mathcal{M}$ as a family of definable regions satisfying o-minimality conditions, thereby ensuring finite topological complexity. These two analyses, spectral and logical, converge on a single conclusion: the apparent continuity of the latent space masks an underlying discrete structure. This is the essence of the Semantic Characterization Theorem (SCT), whose proofs we develop in Section 3.

The implications of this result extend well beyond formal theory. If semantic space is functionally discrete, then the seemingly unpredictable behavior of LLMs, such as sudden shifts in tone, unexpected failure modes, or the phenomenon of “jailbreaking,” can be understood as transitions between stable attractors within a well-defined phase space. What appears as stochastic variation is, in fact, deterministic navigation through a discrete landscape of semantic basins. This perspective enables a principled approach to prompt engineering: rather than treating prompts as arbitrary strings that may or may not “work,” we can interpret them as coordinates that locate the system within particular regions of semantic space. The empirically observed low-dimensional structure provides a natural basis for understanding and controlling these transitions. In this sense, the theorem does not merely describe LLM behavior; it renders that behavior navigable, and ultimately, tameable.

The remainder of this paper develops the formal and empirical basis for these claims. Section 2 generalizes the concept of an LLM into a *Continuous State Machine*. Section 3 states the Semantic Characterization Theorem and presents both the dynamical and logical analyses. Section 4 reports empirical results that corroborate these theoretical predictions, including the emergence of six dominant semantic dimensions. Section 5 presents related work and situates our work within a larger context. Section 6 discusses broader implications for interpretability and the practical control of LLM behavior. We conclude with open questions and directions for future research on the geometry of meaning in large-scale learning systems.

2 Preliminaries

2.1 Continuous State Machines

Definition 2.1 (Continuous State Machine). Let $d \geq 1$ and $\mathcal{M} \subseteq \mathbb{R}^d$ be compact. Let $\mathcal{X}$ be a finite vocabulary, i.e., $\mathcal{X} = \{x_1, x_2, \dots, x_{|\mathcal{X}|}\}$ where $|\mathcal{X}| < \infty$. Let $\mathcal{U} \subseteq \text{Dist}(\mathcal{X})$.¹ Given this, a *Continuous State Machine* or CSM, $\mathcal{C}$, over $\mathcal{X}$ is a sextuple as follows:

$$\mathcal{C} = (\mathcal{M}, \mathcal{U}, T, s_0, O, \Delta), \text{ where:}$$

- $\mathcal{M}$ is the *state space*;
- $\mathcal{U}$ is the set of admissible control distributions over elements of $\mathcal{X}$;
- $T : \mathcal{M} \times \mathcal{U} \rightarrow \mathcal{M}$ is a *transition operator* that maps a current state $s \in \mathcal{M}$ and a control/input distribution, $u \in \mathcal{U}$, to state $s^+ = T(s, u) \in \mathcal{M}$; note that we require T to be continuously differentiable in its first argument and continuous in its second;

¹ $\text{Dist}(\mathcal{X})$ is the set of all probability distributions over $\mathcal{X}$. Since $\mathcal{X}$ is finite, $\text{Dist}(\mathcal{X})$ is both compact and convex.

- $s_0 \in \mathcal{M}$ is a (context-dependent) initial state;
- $O : \mathcal{M} \rightarrow \text{Dist}(\mathcal{X})$ is an *observation/decoder* that assigns a distribution to each state; and
- $\Delta : \text{Dist}(\mathcal{X}) \rightarrow \mathcal{U}$ is a (possibly stochastic) *decoding policy* that maps distributions to control inputs.

Furthermore, a single generation step proceeds as

$$\begin{aligned} s_{t+1} &= T(s_t, u_t), \\ u_{t+1} &= \Delta(O(s_{t+1})), \end{aligned}$$

yielding the trajectory $(s_0, s_1, s_2, \dots)$ and emitted token sequence $(x_1, x_2, \dots)$ sampled from $O(s_t)$.

Remark 2.2. External input or a user prompt enters the system through the choice of initial state or control. An *encoder* $E : \mathcal{X}^* \rightarrow \mathcal{M}$ maps a sequence of tokens to an initial state $s_0 = E(x_{1:k})$. During subsequent interaction, controls u_t may be provided either by the internal policy $\Delta(O(s_t))$ or exogenously by a user, representing additional input tokens. In training, the corpus serves to determine the maps (T, O, Δ) and does not appear explicitly in the definition.

Definition 2.3 (LLM as Continuous State Machine). In large-language-model terminology, $\mathcal{X}$ is the finite vocabulary of tokens, $O(s)$ is the model’s predictive distribution over $\mathcal{X}$, and $\mathcal{U}$ (a subset of $\text{Dist}(\mathcal{X})$) represents the distribution or one-hot embedding selected as the next input token.

Example 2.4 (A Minimal Smooth CSM). To illustrate the definition of a Continuous State Machine, consider the following two-dimensional example with a smooth transition operator and a small vocabulary. Let the latent semantic space be the compact rectangle

$$M = [-1, 1]^2 \subset \mathbb{R}^2,$$

and let the vocabulary consist of three tokens $X = \{a, b, c\}$. The control space $U \subset \text{Dist}(X)$ therefore consists of all distributions $u = (u_a, u_b, u_c)$ with $u_i \geq 0$ and $u_a + u_b + u_c = 1$.

Transition operator. Define the smooth transition operator $T : M \times U \rightarrow M$ by

$$T(s, u) = \tanh(As + Bu),$$

where $A \in \mathbb{R}^{2 \times 2}$ and $B \in \mathbb{R}^{2 \times 3}$ are fixed parameter matrices, and where $\tanh$ is applied component-wise. Since $\tanh$ is analytic and globally Lipschitz on $\mathbb{R}$, the map T is C^1 in s , continuous in u , and definable in the o-minimal structure $\mathbb{R}_{\text{exp}}$.

Observation. Let the decoder $O : M \rightarrow \text{Dist}(X)$ be the softmax of three linear functionals:

$$O(s)_i = \frac{\exp(\ell_i^\top s)}{\sum_{j \in \{a, b, c\}} \exp(\ell_j^\top s)}, \quad i \in \{a, b, c\},$$

where the vectors $\ell_a, \ell_b, \ell_c \in \mathbb{R}^2$ are fixed. This map assigns a probability distribution over tokens to each latent state.

Policy and dynamics. A simple (possibly stochastic) policy $\Delta : \text{Dist}(X) \rightarrow U$ may be taken as the identity, i.e. $\Delta(O(s)) = O(s)$. A single generation step then evolves by

$$s_{t+1} = T(s_t, O(s_t)), \quad u_{t+1} = O(s_{t+1}),$$

giving a well-defined trajectory (s_t) in M .

Interpretation. This example shows concretely how a CSM integrates (i) a smooth transition map, (ii) a probabilistic decoder, and (iii) a policy determining the next control input. Despite the simplicity of the system, repeated application of T quickly drives the trajectory onto low-dimensional attracting regions, illustrating on a small scale the semantic collapse phenomena analyzed in later sections.

Remark 2.5 (Stochastic transition kernel). When Δ is stochastic (e.g., temperature sampling or nucleus sampling), the composition

$$K(s, A) := \int_{\mathcal{U}} \mathbf{1}\{T(s, u) \in A\} \Pi(du \mid s)$$

defines a *Markov kernel* on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$, where $\Pi(\cdot \mid s)$ is the control law induced by O and Δ . Here $\mathcal{B}(\mathcal{M})$ denotes the *Borel σ -algebra* on $\mathcal{M}$, which is the smallest σ -algebra containing all open subsets of $\mathcal{M}$ in the subspace topology of $\mathbb{R}^d$, so that $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ forms a measurable space. Intuitively, $K(s, A)$ gives the probability that the next state lands in the measurable subset $A \subseteq \mathcal{M}$ when the current state is s . Spectral analysis can then be applied to the associated *transfer operator* on $L^2(\mathcal{M})$, with the deterministic case recovered as the limit of vanishing stochasticity (e.g., temperature $\rightarrow 0$).

Example 2.6 (A CSM with Stochastic Decoder). We now present a second example illustrating how stochasticity enters through the decode-policy pipeline. Let the state space again be the compact rectangle

$$M = [-1, 1]^2 \subset \mathbb{R}^2,$$

and let the vocabulary be $X = \{x_1, x_2\}$. The control space is therefore the 1-simplex $U = \{(u_1, u_2) : u_i \geq 0, u_1 + u_2 = 1\}$.

Transition operator. Let the smooth transition operator be

$$T(s, u) = \tanh(As + Bu),$$

where $A \in \mathbb{R}^{2 \times 2}$ and $B \in \mathbb{R}^{2 \times 2}$ are fixed matrices. As before, the analyticity of $\tanh$ ensures that T is C^1 in s , continuous in u , and definable in the o-minimal structure $\mathbb{R}_{\text{exp}}$.

Stochastic decoder. Instead of selecting a deterministic distribution from a softmax, we introduce a smooth stochastic decoder using a compactly supported Gaussian mixture. For each state $s \in M$, define the *logit means*

$$m_i(s) := \ell_i^\top s, \quad i = 1, 2,$$

with fixed vectors $\ell_1, \ell_2 \in \mathbb{R}^2$. The decoder samples a random logit vector

$$z = (z_1, z_2) \sim \mathcal{N}(m(s), \sigma^2 I_2),$$

and maps it to a probability distribution via the softmax:

$$O(s) = \text{softmax}(z) = \left(\frac{e^{z_1}}{e^{z_1} + e^{z_2}}, \frac{e^{z_2}}{e^{z_1} + e^{z_2}} \right).$$

The resulting decoder $O : M \rightarrow \text{Dist}(X)$ is a stochastic kernel: for each measurable $A \subseteq U$,

$$\mathbb{P}(O(s) \in A) = \int_{\{z : \text{softmax}(z) \in A\}} \phi_\sigma(z - m(s)) dz,$$

where ϕ_σ is the Gaussian density. This kernel is smooth in s , analytic in the parameters, and definable in $\mathbb{R}_{\text{exp}}$.

Policy. Let the policy be the identity: $\Delta(u) = u$. Then the control input evolves as

$$u_{t+1} = O(s_{t+1}),$$

producing a stochastic sequence (u_t) .

Induced stochastic dynamics. The pair (T, O) induces a Markov kernel

$$K(s, A) := \mathbb{P}(T(s, O(s)) \in A) \quad \text{for measurable } A \subseteq M.$$

Because the Gaussian density is continuous, the composition $s \mapsto T(s, O(s))$ has a continuous conditional density on the compact domain M , guaranteeing that the induced transfer operator is Hilbert–Schmidt (see Appendix A).

Interpretation. This example shows how a CSM can incorporate stochasticity through the decoder rather than the transition operator. The resulting Markov dynamics remain smooth and definable, while still exhibiting the ergodicity and spectral compactness required for the Semantic Characterization Theorem.

2.2 Assumptions

Definition 2.7 (Standard Borel Space). A *Borel space* is a pair $(M, \mathcal{B}(M))$ where M is a topological space and $\mathcal{B}(M)$ is the Borel σ -algebra generated by its open subsets, i.e., the smallest σ -algebra containing all open sets. A *standard Borel space* is a Borel space arising from a Polish topology (that is, M is separable and completely metrizable). Compact subsets of $\mathbb{R}^d$ with their usual topology are standard Borel spaces, and all measurable maps, measures, and Markov kernels in this paper are defined on such spaces.

Assumption 2.8 (Regularity and boundedness). Throughout the paper, we assume the following conditions hold for a CSM $\mathcal{C} = (\mathcal{M}, \mathcal{U}, T, s_0, O, \Delta)$.

- (A1) **Compact state space.** The latent semantic space $\mathcal{M} \subseteq \mathbb{R}^d$ is compact and equipped with the standard Borel σ -algebra $\mathcal{B}(\mathcal{M})$ and the normalized Lebesgue measure ν .
- (A2) **Smooth transition operator.** The mapping $T : \mathcal{M} \times \mathcal{U} \rightarrow \mathcal{M}$ is C^1 in its first argument, with Jacobian $J_T(s, u)$ bounded uniformly:

$$\sup_{(s,u) \in \mathcal{M} \times \mathcal{U}} \|J_T(s, u)\| < \infty.$$

In the autonomous (deterministic) case, we write $T : \mathcal{M} \rightarrow \mathcal{M}$.

- (A3) **Lipschitz continuity.**² There exists $L > 0$ such that $\|T(s_1, u) - T(s_2, u)\| \leq L\|s_1 - s_2\|$ for all $s_1, s_2 \in \mathcal{M}$ and all admissible $u \in \mathcal{U}$.
- (A4) **Invariant measure / ergodicity.** In the stochastic case, the induced Markov kernel $K(s, A)$ admits a stationary probability measure μ absolutely continuous with respect to ν , and the process $(s_t)_{t \geq 0}$ is ergodic under K . In the deterministic case, the empirical measure along trajectories converges to an invariant measure supported on the attractor set of T .

²On compact sets, C^1 implies Lipschitz Continuity, therefore technically (A3) is redundant. However, we wish to call attention to L , the finite bound for the rate of change of T .

(A5) Spectral boundedness. The associated transfer operator $P : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$ (or $Pf(s) = f(T(s))$ in the deterministic case) is compact and admits a discrete spectrum $\{\lambda_i\}_{i \geq 1}$ with $|\lambda_i| \downarrow 0$.

These mild assumptions³ ensure the existence of well-defined trajectories, finite spectral energy, and analytic dependence of semantic flows on the model parameters. They suffice for all subsequent results concerning spectral collapse and o-minimal definability.

3 The Semantic Characterization Theorem

Remark 3.1 (Scope and structure of proofs). For clarity of exposition, the *Semantic Characterization Theorem* is stated in its most general form, encompassing both deterministic and stochastic continuous state machines. The proofs presented in the main text treat the deterministic case, where the transition operator $T : \mathcal{M} \rightarrow \mathcal{M}$, is smooth and autonomous. This setting captures the essential geometric and logical phenomena—spectral collapse and o-minimal definability, without the technical overhead of measure theory. The stochastic extension, in which T is replaced by the induced Markov kernel $K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \rightarrow [0, 1]$, is developed in Appendix B. There we show that the same conclusions follow by lifting the analysis to the corresponding transfer operator on $L^2(\mathcal{M})$ and invoking standard results on spectral gap and lumpability for Markov processes.

Theorem 3.2 (Semantic Characterization Theorem). *Let*

$$\mathcal{C} = (\mathcal{M}, T, s_0, \mathcal{U}, O, \Delta)$$

be a (possibly stochastic) continuous state machine satisfying Assumptions 2.8(A1)–(A5). Let $K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \rightarrow [0, 1]$ denote the induced Markov kernel and $P : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$ its associated transfer operator. Then:

- (i) (Spectral lumpability)** *There exists a finite $r < \infty$ such that the spectrum of P satisfies $|\lambda_1| \geq \dots \geq |\lambda_r| > \epsilon$ for some $\epsilon > 0$, while $|\lambda_{r+1}| \ll \epsilon$. The corresponding leading eigenfunctions $\{\phi_1, \dots, \phi_r\}$ define a measurable partition of $\mathcal{M}$ into spectral basins $\mathcal{B}_i := \{s \in \mathcal{M} : \arg \max_j |\phi_j(s)| = i\}$, whose boundaries have measure zero under μ .*
- (ii) (Logical tameness)** *Each basin $\mathcal{B}_i$ is definable in an o-minimal structure over $\mathbb{R}$, and the family $\{\mathcal{B}_1, \dots, \mathcal{B}_r\}$ induces a finite cell decomposition of $\mathcal{M}$ with bounded topological complexity. Consequently, the restriction of T (or K) to each basin is semialgebraic and admits a well-defined symbolic abstraction.*
- (iii) (Equivalence of characterizations)** *The partition $\{\mathcal{B}_i\}$ obtained by spectral lumpability coincides, up to measure zero, with the partition obtained by o-minimal definability. Hence the transition from continuous to discrete semantics can be viewed equivalently as:*

$$\text{spectral collapse of } P \iff \text{definable tameness of } (\mathcal{M}, T).$$

In particular, the latent semantic space of a CSM, though continuous in its parameters and representations, is functionally equivalent to a discrete symbolic system of rank r .

³For formal justification of the five assumptions, please refer to appendix A.

Proof sketch and strategy. The proof proceeds in two independent but ultimately convergent parts. From the dynamical perspective, we analyze the spectral properties of the transfer operator P on $L^2(\mathcal{M}, \mu)$ and show that, under Assumptions 2.8(A1)–(A5), its eigenstructure undergoes *spectral collapse*: the system’s long-term dynamics are governed by a finite set of dominant modes, each corresponding to a stable attractor in semantic space. This yields the partition of $\mathcal{M}$ into spectral basins $\{\mathcal{B}_i\}$, proving part (i).

From the logical perspective, we consider the same dynamical map $T : \mathcal{M} \rightarrow \mathcal{M}$ as a definable transformation over the reals. Using standard results on o-minimal structures, we show that the image and preimage of each spectral basin are definable sets with bounded topological complexity, thereby establishing part (ii). The equivalence in part (iii) follows from the observation that spectral basins coincide with definable cells up to measure zero: both constructions isolate precisely those regions of $\mathcal{M}$ within which the model’s semantic behavior is invariant under small perturbations of state.

We first develop the spectral argument (Section 3.1), followed by the logical argument (Section 3.2), and conclude by showing that the two partitions agree almost everywhere. $\square$

3.1 Spectral Collapse and Dynamical Interpretation

Before turning to the formal operator-theoretic analysis, it is helpful to describe the geometric intuition behind spectral collapse. The transfer operator P acts on functions over the latent manifold M by propagating semantic mass forward through the model’s dynamics. Although the underlying space M may be high-dimensional and continuous, repeated application of P exponentially suppresses all but a finite number of directions in which semantic information can persist.

These surviving directions (the dominant eigenfunctions) correspond to large-scale, slowly varying patterns of meaning that remain stable under the model’s internal flow. In effect, P filters out high-frequency semantic fluctuations and reveals a compact, low-rank structure governing long-term behavior. The goal of this section is to make this intuition precise by showing that compactness of P forces its spectrum to concentrate on finitely many modes, each defining a spectral basin in M whose points share the same asymptotic semantic fate.

From the dynamical viewpoint, the latent semantics of a LLM evolve under a transition operator T (deterministic) or its associated Markov kernel K (stochastic) on the compact manifold $\mathcal{M}$. The corresponding transfer operator P acting on functions $f \in L^2(\mathcal{M}, \mu)$ encodes the propagation of semantic mass:

$$(Pf)(s) = \int_{\mathcal{M}} f(s') K(s, ds')$$

or, in the deterministic case:

$$(Pf)(s) = f(T(s)).$$

The spectral properties of P determine the long-term behavior of semantic trajectories. When P is compact, its spectrum consists of a discrete set of eigenvalues $\{\lambda_i\}$ accumulating only at 0, with each eigenfunction ϕ_i representing a stable pattern of semantic recurrence. The *spectral collapse* described in Theorem 3.2(i) corresponds to the concentration of spectral energy in a finite number of dominant modes, each defining a coherent semantic attractor or “basin” in $\mathcal{M}$.

Lemma 3.3 (Spectral decomposition of the transfer operator). *Under Assumptions 2.8(A1)–(A5), the transfer operator $P : L^2(M, \mu) \rightarrow L^2(M, \mu)$ is compact and admits a discrete spectral expan-*

sion, namely:

$$Pf = \sum_{i=1}^{\infty} \lambda_i \langle f, \psi_i \rangle_{\mu} \varphi_i, \quad |\lambda_1| \geq |\lambda_2| \geq \dots \downarrow 0,$$

where $\{\varphi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ are biorthogonal eigenfunction systems satisfying

$$P\varphi_i = \lambda_i \varphi_i, \quad P^*\psi_i = \overline{\lambda_i} \psi_i, \quad \langle \psi_i, \varphi_j \rangle_{\mu} = \delta_{ij}.$$

Proof (expanded). **Step 0: Setting and two cases.** We work on the Hilbert space $L^2(M, \mu)$, where $M \subset \mathbb{R}^d$ is compact and μ is an invariant (or stationary) probability measure. The transfer operator P acts by

$$\begin{aligned} (Pf)(s) &= \int_M f(s') K(s, ds') \quad (\text{stochastic case}), \quad \text{or} \\ (Pf)(s) &= f(T(s)) \quad (\text{deterministic case}), \end{aligned}$$

depending on whether the dynamics are given by a Markov kernel K or a smooth map $T : M \rightarrow M$.

Step 1: Compactness of P . Compactness follows from the regularity assumptions:

- In the *stochastic* case, the kernel admits a continuous density $k(s, s')$ on the compact set $M \times M$, so P is an integral operator with continuous kernel. By the Arzelà–Ascoli/Hilbert–Schmidt argument, P is compact on $L^2(M, \mu)$.
- In the *deterministic* case, one may use a standard finite-horizon averaging (or the compactness surrogates given in the appendix) to obtain a compact limit operator capturing the same leading spectral structure; the statement for P follows by the usual approximation and compactness transfer.

Step 2: Discrete spectrum and eigen-expansion. Since P is compact on a Hilbert space, the spectral theorem for compact operators ensures that:

1. The spectrum is at most countable with 0 as the only possible accumulation point.
2. There exist eigenvalues $\{\lambda_i\}_{i \geq 1}$ ordered by nonincreasing modulus, $|\lambda_1| \geq |\lambda_2| \geq \dots \downarrow 0$.
3. There exist right eigenfunctions $\{\varphi_i\}$, $P\varphi_i = \lambda_i \varphi_i$, forming a (possibly non-orthogonal) complete system in the invariant subspace generated by $\text{im}(P)$; and left eigenfunctions $\{\psi_i\}$ of the adjoint P^* , $P^*\psi_i = \overline{\lambda_i} \psi_i$.

By choosing the dual system $\{\psi_i\}$ appropriately (e.g., via Riesz representation in L^2), we enforce *biorthogonality*:

$$\langle \psi_i, \varphi_j \rangle_{\mu} = \delta_{ij}.$$

Define the analysis map $A : L^2 \rightarrow \ell^2$ by $(Af)_i := \langle f, \psi_i \rangle_{\mu}$ and the synthesis map $S : \ell^2 \rightarrow L^2$ by $Sc := \sum_i c_i \varphi_i$. Then $P = S\Lambda A$, where Λ is diagonal with entries λ_i .

Step 3: Convergent series for Pf . For any $f \in L^2(M, \mu)$ one has

$$Pf = \sum_{i=1}^{\infty} \lambda_i \langle f, \psi_i \rangle_{\mu} \varphi_i,$$

with convergence in L^2 -norm. Indeed, compactness implies $\|P - P^{(r)}\| \rightarrow 0$ where $P^{(r)} := \sum_{i=1}^r \lambda_i \langle \cdot, \psi_i \rangle_\mu \varphi_i$ is finite-rank. Thus $P^{(r)} f \rightarrow P f$ in norm for each f , and the partial sums of the series converge to $P f$.

Step 4: Ordering and decay of eigenvalues. Finally, because P is compact, the (nonzero) eigenvalues form a discrete set with $|\lambda_i| \rightarrow 0$. We record the conventional ordering $|\lambda_1| \geq |\lambda_2| \geq \dots \downarrow 0$ for later use (e.g., in truncation and spectral-gap arguments).

Conclusion. We have exhibited a complete biorthogonal eigenfunction system $\{\varphi_i\}, \{\psi_i\}$ and the corresponding discrete spectral expansion of P with eigenvalues tending to 0. This proves the lemma. $\square$

The dominant eigenfunctions $\phi_1, \dots, \phi_r$ describe the slow modes of semantic evolution—the directions of maximal stability in the latent space. When the spectrum exhibits a clear gap after the first r modes, the long-term dynamics of any initial state s_0 can be approximated by projection onto this finite subspace:

$$P^t f \approx \sum_{i=1}^r \lambda_i^t \langle f, \psi_i \rangle_\mu \phi_i, \quad t \gg 1.$$

This truncation induces a partition of $\mathcal{M}$ into the *spectral basins*

$$\mathcal{B}_i := \{s \in \mathcal{M} : |\phi_i(s)| = \max_j |\phi_j(s)|\},$$

each corresponding to an attractor governed by a dominant mode. Boundaries between basins are negligible under μ because the set of points where $|\phi_i| = |\phi_j|$ for $i \neq j$ has measure zero.

Proposition 3.4 (Spectral collapse). *If the eigenvalues of P satisfy $|\lambda_{r+1}|/|\lambda_r| < \epsilon$ for some small $\epsilon \ll 1$, then for all t sufficiently large, the total spectral energy outside the leading r modes decays exponentially:*

$$\frac{\|P^t f - \sum_{i=1}^r \lambda_i^t \langle f, \psi_i \rangle_\mu \phi_i\|_2}{\|f\|_2} \leq Z \epsilon^t,$$

where $Z > 1$ is a constant. Hence the system's effective dynamics are confined to the finite-dimensional subspace spanned by $\{\phi_1, \dots, \phi_r\}$, and the semantic space $\mathcal{M}$ decomposes into r stable spectral basins $\{\mathcal{B}_i\}$.

Detailed Proof of Proposition 2.4. Let $P : L^2(M, \mu) \rightarrow L^2(M, \mu)$ be the compact transfer operator from Lemma 2.3, with biorthogonal eigenfunction systems $\{\varphi_i\}_{i \geq 1}$ and $\{\psi_i\}_{i \geq 1}$ satisfying

$$P \varphi_i = \lambda_i \varphi_i, \quad P^* \psi_i = \bar{\lambda}_i \psi_i, \quad \langle \psi_i, \varphi_j \rangle_\mu = \delta_{ij}.$$

For any $f \in L^2(M, \mu)$, the Hilbert–Schmidt theorem yields the spectral expansion

$$P f = \sum_{i=1}^{\infty} \lambda_i \langle f, \psi_i \rangle_\mu \varphi_i, \quad \|f\|_2^2 = \langle f, f \rangle_\mu = \sum_{i,j} \langle f, \psi_i \rangle_\mu \langle f, \psi_j \rangle_\mu \langle \varphi_i, \varphi_j \rangle_\mu.$$

Step 1: Finite-rank truncation. Let $P^{(r)}$ denote the rank- r truncation

$$P^{(r)} f := \sum_{i=1}^r \lambda_i \langle f, \psi_i \rangle_\mu \varphi_i,$$

and define the residual operator $R^{(r)} := P - P^{(r)}$. Since P is compact with discrete spectrum, the operator norm of the residual satisfies

$$\|R^{(r)}\| \leq |\lambda_{r+1}|.$$

Iterating t times gives $\|R^{(r)t}\| \leq |\lambda_{r+1}|^t$.

Step 2: Decomposition of iterates. Because $P^{(r)}$ and $R^{(r)}$ act on complementary invariant subspaces, the t -fold iterate decomposes as

$$P^t f = (P^{(r)} + R^{(r)})^t f = \sum_{i=1}^r \lambda_i^t \langle f, \psi_i \rangle_{\mu} \varphi_i + R^{(r)t} f.$$

Hence the error between $P^t f$ and its rank- r spectral approximation is exactly $R^{(r)t} f$.

Step 3: Bounding the residual. For any $f \in L^2(M, \mu)$,

$$\|P^t f - P^{(r)t} f\| = \|R^{(r)t} f\| \leq \|R^{(r)t}\| \|f\| \leq |\lambda_{r+1}|^t \|f\|.$$

In the biorthogonal (non-orthonormal) setting, this bound picks up a conditioning factor $\kappa(P)$ depending on the norms of the eigenbasis and its dual:

$$\|P^t f - P^{(r)t} f\| \leq \kappa(P) |\lambda_{r+1}|^t \|f\|, \quad \kappa(P) := \|\Phi\| \|\Psi\|,$$

where Φ and Ψ are the synthesis and analysis operators associated with $\{\varphi_i\}$ and $\{\psi_i\}$ respectively. Denote this constant generically by $Z > 1$.

Step 4: Exponential decay under a spectral gap. Assume the spectral gap condition $|\lambda_{r+1}|/|\lambda_r| < \epsilon \ll 1$. Then $|\lambda_{r+1}|^t \leq |\lambda_r|^t \epsilon^t$, and we may absorb the factor $|\lambda_r|^t$ into Z . Therefore,

$$\frac{\|P^t f - \sum_{i=1}^r \lambda_i^t \langle f, \psi_i \rangle_{\mu} \varphi_i\|^2}{\|f\|^2} \leq Z \epsilon^t.$$

Step 5: Interpretation. The inequality shows that spectral energy outside the first r modes decays exponentially with rate ϵ . Consequently, the long-term dynamics of P are confined to the finite-dimensional subspace $\text{span}\{\varphi_1, \dots, \varphi_r\}$, and the latent manifold M decomposes into the corresponding r spectral basins $B_i = \{s \in M : |\varphi_i(s)| = \max_j |\varphi_j(s)|\}$. $\square$

The existence of finitely many dominant eigenmodes therefore entails that the system's asymptotic behavior—and hence its functional semantics—is determined by a finite discrete structure. Each basin $\mathcal{B}_i$ defines a region of the latent space within which small perturbations of state leave semantic outcomes invariant, forming the dynamical half of the Semantic Characterization Theorem.

3.2 Logical Tameness and Definable Interpretation

The spectral perspective described in the previous section reveals that long-term semantic behavior is governed by a finite number of dominant modes. The logical perspective developed here asks a complementary question: what is the topological and combinatorial structure of the regions in M where this behavior is stable?

In particular, if the transition map T is definable in an o-minimal expansion of the real field, then the images and preimages of all semantically relevant sets inherit a highly regular geometry. O-minimality ensures that no pathological or fractal boundaries arise; every definable subset of M decomposes into finitely many smooth cells with bounded topological complexity.

From this viewpoint, the semantic basins identified spectrally correspond to regions in which the logical type of the state is invariant under infinitesimal perturbations. The goal of this section is to formalize this notion of “tameness” and to show that the CSM dynamics preserve definability in a way that forces M to admit a finite cell decomposition aligned with the system’s semantic invariants.

In other words, the logical half of the Semantic Characterization Theorem views the same dynamics studied in Section 3.1 through the lens of definability over the reals. Whereas the spectral argument analyzes the operator P acting on functions, the logical argument considers the map $T : \mathcal{M} \rightarrow \mathcal{M}$ as a definable transformation between subsets of Euclidean space. The goal is to show that, under the same regularity assumptions, the partition of $\mathcal{M}$ induced by the dynamics is *o-minimal*: it decomposes into finitely many definable cells of bounded topological complexity.

Definition 3.5 (Definable transformation). Let $(\mathbb{R}, +, \times, <, \mathcal{O})$ be an o-minimal expansion of the real field, where $\mathcal{O}$ is a set of appropriate function symbols (e.g., $\mathcal{O} = \{\exp\}$). A mapping $S : \mathcal{M} \rightarrow \mathcal{M}$ is said to be *definable* in this structure if its graph $\{(s, s') \in \mathcal{M} \times \mathcal{M} : s' = S(s)\}$ is a definable set. A subset $A \subseteq \mathcal{M}$ is definable if it can be expressed by a first-order formula over $\mathcal{O}$ with parameters in $\mathbb{R}$.

Example 3.6 (Definable basins for an analytic transition map). Let $(\mathbb{R}, +, \times, <, \mathcal{O})$ be an o-minimal expansion of the real field in which the exponential function is definable (e.g. $\mathbb{R}_{\exp}$). Consider the analytic transition map $T : M \rightarrow M$ on $M = [0, 1]$ given by

$$T(x) := \sigma(ax + b) = \frac{1}{1 + \exp(-(ax + b))},$$

where $a, b \in \mathbb{R}$ are real parameters. The graph of T is then definable by the first-order formula in the language $(+, \times, <, \mathcal{O})$:

$$\Gamma_T(x, y; a, b) \equiv (0 \leq x \wedge x \leq 1) \wedge \left(y = (1 + \exp(-(a \cdot x + b)))^{-1} \right).$$

Define two subsets of M by thresholding the image of T :

$$B_0 := \{x \in [0, 1] : T(x) \leq \tfrac{1}{2}\}, \quad B_1 := \{x \in [0, 1] : T(x) > \tfrac{1}{2}\}.$$

Each set is definable in the same o-minimal structure. For instance, B_0 is defined by the first-order formula

$$\varphi_0(x; a, b) \equiv (0 \leq x \wedge x \leq 1) \wedge \exists y (\Gamma_T(x, y; a, b) \wedge y \leq \tfrac{1}{2}),$$

and similarly B_1 is defined by

$$\varphi_1(x; a, b) \equiv (0 \leq x \wedge x \leq 1) \wedge \exists y (\Gamma_T(x, y; a, b) \wedge y > \tfrac{1}{2}).$$

Thus both B_0 and B_1 are definable subsets of M given explicitly by first-order formulas over the collection of function symbols $\mathcal{O}$ (here including $\exp$), with real parameters $a, b \in \mathbb{R}$. This illustrates, in a simple analytic setting, how definable partitions of the state space arise from definable transition maps.

Because $\mathcal{M} \subseteq \mathbb{R}^d$ is compact and T is C^1 , all images and preimages of definable sets under T remain definable (59, Theorem 3.5). Moreover, definable sets in an o-minimal structure admit a *cell decomposition* into finitely many connected components each homeomorphic to an open box in $\mathbb{R}^k$.

Lemma 3.7 (Definable invariance). *Let $T : \mathcal{M} \rightarrow \mathcal{M}$ be C^1 and definable in an o-minimal structure. Then for every definable set $A \subseteq \mathcal{M}$, both $T(A)$ and $T^{-1}(A)$ are definable, and there exists a finite cell decomposition $\{C_1, \dots, C_m\}$ of $\mathcal{M}$ such that $T(C_i) \subseteq C_j$ for some j .*

Proof. Definability of images and preimages follows from the closure of definable sets under projections and inverse images of definable functions. Compactness of $\mathcal{M}$ implies that T maps each cell to a bounded definable set, which can be refined to a finite cell decomposition satisfying the stated containment property. See, e.g., [13, 59]. $\square$

Within each definable cell C_i , the dynamics of T are smooth and topologically trivial; semantic variation is continuous but does not alter the logical type of the state. The collection $\{\mathcal{B}_1, \dots, \mathcal{B}_r\}$ defined spectrally in Section 3.1 thus corresponds, up to measure zero, to a finite refinement of the definable cell decomposition in Lemma 3.7. Each cell constitutes a region of *logical tameness*: within it, the model’s semantic behavior is first-order invariant under infinitesimal perturbations of s .

Proposition 3.8 (Finite definable complexity). *Under Assumptions 2.8(A1)–(A3), the mapping $T : \mathcal{M} \rightarrow \mathcal{M}$ is definable in some o-minimal expansion of $(\mathbb{R}, +, \times, <)$. Consequently, $\mathcal{M}$ admits a finite cell decomposition into definable regions $\{\mathcal{B}_i\}_{i=1}^r$ such that each restriction $T|_{\mathcal{B}_i}$ is semialgebraic and locally bi-Lipschitz.*

Proof. Since T is C^1 with bounded Jacobian and analytic activation functions in practice, each component function of T is real analytic and hence definable in the o-minimal structure $\mathbb{R}_{\text{an}}$ (59). Compactness of $\mathcal{M}$ guarantees finitely many connected components, each a definable cell. The local bi-Lipschitz property follows from the inverse function theorem on each component. $\square$

Interpretation. Logical tameness means that the continuous semantic manifold possesses a finite combinatorial skeleton: within each definable region, semantic transformations are smooth and predictable, while transitions between regions correspond to discrete logical changes. The definable regions $\{\mathcal{B}_i\}$ therefore provide the logical counterpart to the spectral basins derived in Section 3.1.

3.3 Equivalence of Spectral and Logical Tameness

The spectral and logical analyses developed so far offer two seemingly independent characterizations of semantic structure. The spectral viewpoint reduces long-term dynamics to a finite collection of dominant eigenmodes, each carving the manifold into a stable basin of attraction. The logical viewpoint, by contrast, shows that definable dynamics constrain $\mathcal{M}$ to a finite family of o-minimal cells, each representing a region of invariant logical type.

The striking fact is that these two partitions are not merely compatible but coincide almost everywhere.

This convergence is far from obvious: one arises from functional-analytic properties of the transfer operator P , while the other emerges from first-order definability over the reals. The purpose of this section is to establish that the two perspectives describe the same underlying semantic skeleton: spectral basins correspond exactly to definable cells, and continuous semantics collapse into a discrete structure that is simultaneously dynamical and logical in nature.

In other words, we now establish part (iii) of Theorem 3.2, showing that the spectral partition of $\mathcal{M}$ obtained from the dominant eigenfunctions of P coincides, up to sets of measure zero, with the definable cell decomposition arising from the o-minimal structure. Intuitively, the two analyses describe the same phenomenon from dual perspectives: the spectral argument captures the geometry

of invariant subspaces, while the logical argument captures the topology of invariant definable regions.

Lemma 3.9 (Definability of eigenlevel sets). *If P is compact and T is definable in an o-minimal structure, then each real eigenfunction ϕ_i of P is definable in the same structure, and the level sets $\{s \in \mathcal{M} : \phi_i(s) = c\}$ are definable for all $c \in \mathbb{R}$.*

Proof. Since P acts on definable L^2 functions via $(Pf)(s) = f(T(s))$ in the deterministic case, and by a definable integral transform in the stochastic case, definability of T implies that P preserves the class of definable functions (59, Prop. 2.6). Eigenfunctions of P arise as limits of polynomially definable iterates of P , hence are definable. Level sets of definable functions are definable by closure of the structure under preimages of definable maps. $\square$

The eigenfunctions ϕ_i therefore carve $\mathcal{M}$ into definable regions whose boundaries coincide with the zero sets of definable functions and thus have measure zero under μ . In particular, the spectral basins $\mathcal{B}_i = \{s : |\phi_i(s)| = \max_j |\phi_j(s)|\}$ are definable and coincide almost everywhere with cells of the decomposition from Proposition 3.8.

Proposition 3.10 (Equivalence of partitions). *Let $\{\mathcal{B}_i\}$ denote the spectral basins of Section 3.1 and $\{\mathcal{C}_j\}$ the definable cells of Section 3.2. Then there exists a finite refinement $\{\mathcal{D}_k\}$ of both families such that*

$$\mu\left(\mathcal{M} \setminus \bigcup_k \mathcal{D}_k\right) = 0, \text{ and}$$

$$T(\mathcal{D}_k) \subseteq \mathcal{D}_\ell \text{ for some } \ell,$$

and each $\mathcal{D}_k$ is simultaneously a spectral basin and a definable cell.

Proof. By Lemma 3.9, each ϕ_i is definable, so the spectral basins $\mathcal{B}_i$ are finite Boolean combinations of definable level sets and therefore definable themselves. The cell decomposition of Proposition 3.8 provides finitely many definable cells $\mathcal{C}_j$ covering $\mathcal{M}$ up to measure zero. Refining by intersection, $\mathcal{D}_{ij} = \mathcal{B}_i \cap \mathcal{C}_j$ yields a finite family of definable sets whose union covers $\mathcal{M}$ modulo μ . The invariance condition follows because T maps both $\mathcal{B}_i$ and $\mathcal{C}_j$ into definable unions of the same families. $\square$

Interpretation. The equivalence result demonstrates that spectral lumpability and logical tameness are not merely analogous but identical characterizations of the same phase transition. The dominant eigenfunctions of the transfer operator and the definable predicates of the o-minimal structure partition the semantic manifold in the same way: both identify regions of invariant meaning where the model’s behavior is discretely stable. In this sense, the continuous dynamics of an LLM are *logically complete*: the geometry of its spectral decomposition and the topology of its definable structure coincide almost everywhere, completing the proof of the Semantic Characterization Theorem.

3.4 Ontological Underpinnings of the Main Theorem

Ontological reification. The Semantic Characterization Theorem implies that the latent semantic space of a large language model is not an arbitrary manifold but the continuous realization of an underlying discrete ontology. The invariant measure μ concentrates along high-density ridges that

correspond to stable semantic categories, while the leading spectral modes of P align with the primary axes of conceptual organization.

In other words, seen as a CSM, once the LLM has traversed N^* steps/parameters it will become constrained by enough structure that the effective behavior collapses from $N^* \gg 1$ to some $r \ll O(10)$, i.e., from an astronomically large number to a highly tameable semantic manifold that is, in fact both compact and convex.

Thus, the ontology implicit in the training corpus becomes reified within the continuous dynamics of the model itself: the structure of human conceptual space is recovered as a spectral and definable skeleton of the learned semantics.

Corollary 3.11 (Ontological Skeleton). *Under the assumptions of the Semantic Characterization Theorem (Theorem 3.2), let $\{\mathcal{B}_1, \dots, \mathcal{B}_r\}$ denote the definable (and hence spectral) basins of the latent semantic space $\mathcal{M}$, and let $K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \rightarrow [0, 1]$ be the induced Markov kernel. Define a directed graph $\mathcal{G} = (V, E)$ as follows:*

$$\begin{aligned} V &:= \{v_i \leftrightarrow \mathcal{B}_i\}_{i=1}^r, \\ E &:= \{(v_i, v_j) \mid K(s, \mathcal{B}_j) > 0 \text{ for some } s \in \mathcal{B}_i\}. \end{aligned}$$

Then $\mathcal{G}$ is a finite directed acyclic graph representing the ontological skeleton of the model's semantics, with the following properties:

- (i) **(Discrete abstraction)** *Each vertex v_i corresponds to a definable semantic region $\mathcal{B}_i$; the edges encode admissible semantic transitions between regions under K .*
- (ii) **(Partial ordering)** *The transition relation E induces a partial order on V consistent with temporal or causal dependencies among regions, so that $\mathcal{G}$ is acyclic almost surely under μ .*
- (iii) **(Dimensional alignment)** *The principal coordinates of $\mathcal{G}$ align with the r dominant conceptual axes of the learned ontology.*

Consequently, the discrete relational structure implicit in the training distribution becomes reified in the model's continuous dynamics as $\mathcal{G}$: a topologically finite, spectrally grounded abstraction of the latent semantic space.

Remark 3.12 (Time-Inhomogeneous SCT (Adiabatic Case)). The results of the Semantic Characterization Theorem extend naturally to certain non-stationary stochastic systems provided that compactness and ergodicity hold uniformly in time. Let $(P_t)_{t \geq 0}$ be a sequence of transfer operators with kernels $k_t(s, s')$ on a compact manifold M . Suppose the following conditions are satisfied:

1. **Uniform compactness (Hilbert–Schmidt):** Each kernel k_t is continuous on $M \times M$ with a uniform bound, so each P_t is Hilbert–Schmidt and therefore compact on $L^2(M, \mu_t)$. Finite products of compact operators remain compact, so finite-time propagators $P_{t+n-1} \cdots P_t$ inherit compactness.
2. **Uniform ergodicity (Doebelin–minorization):** There exists $\varepsilon > 0$ and a reference probability measure ν such that $K_t(s, A) \geq \varepsilon \nu(A)$ for all $s \in M$, $A \in \mathcal{B}(M)$, and t . This ensures each step mixes with bounded stochasticity and prevents fragmentation of the measure space.
3. **Slow drift (adiabaticity):** The operators vary slowly in time, with $\|P_{t+1} - P_t\| \leq \eta$ for small η , and maintain a uniform spectral gap after rank r , $\inf_t (|\lambda_r(t)| - |\lambda_{r+1}(t)|) \geq \gamma > 0$.

Then the leading spectral subspace (the instantaneous semantic manifold) varies continuously in t .

4. **Uniform definability:** If T_t (or k_t) is definable in an $\mathcal{o}$ -minimal structure with parameters θ_t , and these parameters vary continuously, the definable partitions $\{B_i(t)\}$ evolve smoothly in time, forming a definable family of cells.

Under these assumptions, each P_t remains compact and ergodic, and the instantaneous semantic basins $B_i(t)$ evolve continuously through time, preserving spectral lumpability and logical tameness at every moment. This setting models adiabatic, context-dependent semantics such as narrative text: each new input step updates the transition kernel slightly, maintaining global compactness while allowing the ontology to drift smoothly.

Corollary 3.13 (Equivalence of Adiabatic and Stationary Semantics). *Under the assumptions of the Time-Inhomogeneous SCT (adiabatic case), suppose that the sequence of transfer operators $(P_t)_{t \geq 0}$ satisfies the uniform compactness, ergodicity, and slow-drift conditions:*

$$\|P_{t+1} - P_t\| \leq \eta, \quad \inf_t (|\lambda_r(t)| - |\lambda_{r+1}(t)|) \geq \gamma > 0, \quad \eta \ll \gamma. \quad (1)$$

Then there exists a stationary compact operator P on $L^2(M, \mu)$ such that for any finite horizon n ,

$$P_{t+n-1} \cdots P_t = P^n + O(n\eta), \quad (2)$$

with the remainder bounded in operator norm. Consequently, the cumulative effect of evolving through a slowly varying narrative (time-dependent kernels P_t) is equivalent, to first order, to propagation under a stationary operator P trained on the entire narrative distribution.

Interpretation. Adiabatically evolving semantic systems and stationary ones trained on their complete trajectories are spectrally equivalent up to $O(\eta)$ corrections. Thus, a narrative, when viewed as a time-dependent process updating its ontology, induces the same compact, ergodic structure as a stationary system trained on that corpus. This equivalence explains how coherent but distinct ontologies (including fictional or hypothetical ones) emerge as compact, self-consistent manifolds within the general semantic space: each narrative defines its own quasi-stationary invariant measure μ_t that drifts slowly over time while preserving definable tameness and finite conceptual rank.

4 Experimental Results

4.1 Methodology and Setup

To empirically examine the predictions of the Semantic Characterization Theorem, we constructed a minimal diffusion experiment on sentence embeddings drawn from six semantic domains: *mathematical*, *scientific*, *narrative*, *affective*, *instructional*, and *factual*. Each prompt consisted of a short English sentence chosen to be representative of its domain without relying on task-specific wording. Embeddings were produced using the all-mpnet-base-v2 encoder, providing a 1024-dimensional latent representation on which we performed analysis.

A symmetric stochastic operator $\tilde{P}$ was constructed by exponentiating negative pairwise distances under a fixed bandwidth σ , followed by row-normalization. This operator approximates a diffusion

process on the sampled manifold and serves as a finite-dimensional surrogate for the transfer operator P studied in the main theorem. We emphasize that $\tilde{P}$ is not a learned model but an observational proxy for the geometry of the latent space; its eigenstructure reveals the intrinsic organization of the embeddings, independent of any particular decoder or policy.

The leading eigenvalues and eigenvectors of $\tilde{P}$ were computed using standard spectral decomposition. The dominant eigenfunctions were then used to assign each sample to a spectral basin $B_i := \{s : \arg \max_j |\phi_j(s)| = i\}$, and the resulting basins were visualized via a two-dimensional UMAP projection. The choice of UMAP is purely descriptive—UMAP plays no role in the theoretical analysis and serves only to offer intuitive visualization of the basins already determined by spectral structure.

4.2 Results and Interpretation

Figure 1 summarizes the empirical findings. The scree plot in panel (a) exhibits a pronounced elbow after the first three eigenvalues, indicating that the diffusion operator carries most of its spectral energy in a finite-dimensional subspace of rank $r \approx 3$. This aligns with the SCT’s prediction that semantic manifolds collapse onto a small number of dominant modes.


Panel (b) shows a UMAP projection colored by the spectral basin assignments. Three coherent regions emerge, with centroids and 2σ ellipses highlighting the stability of the clustering. These basins are defined by the eigenfunctions themselves; the projection merely reveals their geometric arrangement.

To estimate logical tameness, we evaluated simple classifiers on the spectral labels. As shown in panel (d), both a depth-2 decision tree and a degree-2 polynomial logistic regressor obtain accuracies near 0.9, substantially outperforming chance in a low-data regime. Since low-capacity models correspond to low VC dimension, high accuracy provides a proxy for the definability of basin boundaries: a classifier with few degrees of freedom can approximate a definable partition.

Finally, panel (c) illustrates short rollouts between spectral basins, yielding a directed graph G_θ that approximates the ontological skeleton predicted by Corollary 3.11. Edges correspond to transitions induced by iterated applications of $\tilde{P}$, revealing a small, structured set of admissible semantic movements.


4.3 Summary

Collectively, these results validate all three pillars of the Semantic Characterization Theorem in a minimal unsupervised setting: (i) finite-rank spectral structure, (ii) logically tame partitions that are well-approximated by low-capacity models, and (iii) an emergent ontological graph capturing stable semantic relations. Even a small observational system exhibits the semantic discretization predicted theoretically, suggesting that spectral collapse and definability are intrinsic features of learned representation spaces rather than artifacts of scale.



(a) **Spectral gap / finite r .** Scree of the symmetric diffusion operator P shows elbows (we use $r = 3$).

(b) **Spectral basins.** UMAP of samples colored by $B_i = \{s : \arg \max_j |\varphi_j(s)| = i\}$; centroids and 2σ ellipses shown.



(c) **Ontological skeleton G_θ .** Directed edges = observed transitions in short rollouts (weights labeled).

Metric	Value
ARI (bootstrap)	0.504
Jaccard (ovr, mean)	0.251
Tree acc (depth=2)	0.896
Poly-logistic (deg=2)	0.938
Chosen r	3.000

(d) **Stability & tameness.** Moderate ARI with small N ; low-capacity models achieve high accuracy $\Rightarrow$ definability proxy.

Figure 1: **Empirical illustration of the Semantic Characterization Theorem.** (a) *Finite dominant modes*: spectral gap for P . (b) *Basins via eigenfunctions*: B_i from $\arg \max_j |\varphi_j|$. (c) *Ontological skeleton*: coarse G_θ over basins from rollouts. (d) *Logical tameness*: small models separate basins (proxy for definable boundaries).

5 Related Work

This section situates the Semantic Characterization Theorem within several distinct research traditions that have sought to explain how structured or symbolic behavior can arise from continuous dynamics. Across machine learning, dynamical systems, category theory, logic, and cognitive science, many frameworks propose mechanisms by which high-dimensional representations collapse into low-dimensional, interpretable forms. While these lines of work differ in motivation and mathematical machinery, a unifying theme emerges: semantic stability is not an accident of architecture but a consequence of boundedness, regularity, and invariant structure. The SCT contributes to this landscape by showing that spectral compactness and o-minimal definability coincide, providing a bridge between operator-theoretic and logical formulations of meaning. We now review the most relevant precedents, emphasizing those that offer principled routes from continuous computation to discrete semantic organization.

5.1 Coalgebraic Perspectives and the Feys-Hansen-Moss Framework

A parallel line of research arises in the categorical analysis of Markov decision processes (MDPs) by Feys, Hansen, and Moss [18]. Their work reinterprets dynamic programming through the lens of coalgebra and algebra, introducing two formal devices of note: a *contraction coinduction* principle, derived from Banach’s fixed point theorem, and a generalized notion of *b-corecursive algebra* (bca) for ensuring the existence of unique bounded fixpoints in complete metric spaces. Within their framework, contractivity is enforced by an explicit discount factor $\gamma < 1$, and value iteration converges to a unique fixed point representing the long-term reward function of a policy. The categorical treatment thereby captures boundedness and monotonicity compositionally, but its domain remains that of finite or Polish state spaces equipped with reward and policy structure.

Conceptual comparison. Our framework departs from this paradigm in both scope and intent. We make no appeal to rewards, discounting, or policy optimization. Instead, we study the intrinsic dynamics of large models as *continuous state machines* whose transition operators T (or their Markov counterparts K) act on a compact semantic manifold M . Whereas Feys-Hansen-Moss obtain fixed-point uniqueness through an externally imposed contraction, we obtain stability through *spectral boundedness*: the associated transfer operator

$$P : L^2(M, \mu) \rightarrow L^2(M, \mu)$$

is compact with a discrete spectrum $|\lambda_1| \geq \dots \geq |\lambda_r| > \epsilon \gg |\lambda_{r+1}|$. This spectral gap induces an *asymptotic contraction* on the orthogonal complement of the leading r eigenmodes, driving trajectories toward a finite-dimensional invariant subspace that encodes the model’s stable semantic structure. In this sense, the role of γ in discounted dynamic programming is here played by the ratio $|\lambda_{r+1}/\lambda_r| < 1$, which quantifies the rate of collapse of fast modes under repeated application of P .

Conceptually, both settings formalize the emergence of uniqueness and compositionality from boundedness, yet the phenomena differ sharply. Feys–Hansen–Moss analyze the *existence of fixed points* for a discounted value operator; our analysis concerns a *phase transition in the spectrum* of an undirected semantic operator. Where their coinduction principle proves monotone convergence of iterative policy improvement, ours characterizes the self-organization of the latent manifold into finitely many definable regions of invariant meaning. In this way, the coalgebraic apparatus of [18] provides a categorical precedent for our use of metric boundedness and operator compactness, but the present work generalizes those ideas from reward-based control to the unsupervised emergence of symbolic semantics.

Asymptotic contraction as coinduction. Formally, the following result restates Proposition 3.4 in the idiom of contraction coinduction. It establishes that the spectral collapse underlying the Semantic Characterization Theorem can be viewed as a coinductive fixpoint in a complete ordered metric space.

Proposition 5.1 (Asymptotic contraction onto the semantic skeleton). *Let P be a compact linear operator on $L^2(M, \mu)$ with eigenvalues $|\lambda_1| \geq \dots \geq |\lambda_r| > \epsilon \gg |\lambda_{r+1}|$. Let Π_r denote the orthogonal projection onto the invariant subspace $\text{span}\{\varphi_1, \dots, \varphi_r\}$ of the first r eigenfunctions. Then for all $f \in L^2(M, \mu)$,*

$$\|P^t f - \Pi_r P^t f\|_2 \leq Z \left| \frac{\lambda_{r+1}}{\lambda_r} \right|^t \|f\|_2, \quad t \geq 0, \quad (3)$$

for some constant Z depending only on P . Consequently, P is a contraction on the orthogonal complement of $\text{im}(\Pi_r)$, and the unique post-fixpoint under repeated application of P is the rank- r operator Π_r .

Interpretation. Inequality (3) exhibits the spectral complement as a contractive subspace with ratio $|\lambda_{r+1}/\lambda_r| < 1$. By the contraction–coinduction principle of [18], the corresponding fixpoint is both the least pre-fixpoint and the greatest post-fixpoint in the ordered metric space $(L^2(M), \|\cdot\|_2, \leq)$. Hence Π_r serves as the unique coinductive invariant of the semantic dynamics, representing the *ontological skeleton* of the model. $\square$

This reformulation highlights the structural continuity between discounted contraction in policy evaluation and spectral collapse in continuous semantics. In both cases, boundedness and compactness guarantee the existence of a unique coinductive attractor, but in the latter the attractor is not a scalar value function—it is the finite definable partition of the semantic manifold itself.

5.2 Hybrid Systems and Algebraic Models of Dynamics

A parallel thread arises in the study of *hybrid systems*. These are dynamical systems that mix continuous evolution with discrete transitions. This field has developed rich mathematical tools for expressing how continuous trajectories can exhibit sudden qualitative changes, a structure that mirrors the semantic phase transitions captured by the SCT. Foundational treatments such as [9, 20, 61] view hybrid systems as models whose state space admits both differential and automaton-like behavior, with Lyapunov and invariance principles governing stability.

Of particular relevance is the categorical perspective developed by Haghverdi [25, 26], who showed that hybrid behaviors may be represented as coalgebras for polynomial functors and related structures. In this formulation, continuous flows, guards, and resets coexist within a unified algebraic semantics, allowing hybrid trajectories to be expressed compositionally. This coalgebraic viewpoint provides a principled mechanism for identifying invariant subspaces and for understanding how smooth evolution can give rise to discrete mode transitions, which is precisely the structural interplay that underlies semantic collapse.

The parallel with the Semantic Characterization Theorem is immediate. Hybrid systems theory demonstrates that discrete transitions can emerge naturally from continuous dynamics once boundedness and invariance constraints are imposed. The SCT shows an analogous phenomenon in learned semantic systems: spectral compactness forces trajectories to collapse onto finitely many definable regions, each of which behaves as a discrete “mode” or semantic regime. In this sense, hybrid systems provide a well-established mathematical precedent for the emergence of discrete

semantics from smooth computation, and the CSM framework may be viewed as a semantic analogue of hybrid dynamical structure.

5.3 Latent Geometry and Representation Learning

Work in representation learning has repeatedly observed that high-dimensional neural activations concentrate on low-rank manifolds. Studies of network expressivity and training dynamics have traced this phenomenon to implicit compression and invariance formation in deep layers [1, 47, 50]. More recently, the “neural collapse” literature has shown that class features align with nearly orthogonal basis vectors at convergence [46]. These empirical regularities suggest that semantic generalization corresponds to a spectral reduction of representational rank. The Semantic Characterization Theorem formalizes this intuition: spectral collapse to finitely many dominant modes is not merely an optimization artifact but a mathematical signature of definable semantics emerging within a continuous system.

5.4 Spectral and Operator-Theoretic Approaches

An alternative tradition models neural networks as dynamical systems governed by linear or nonlinear operators. Koopman-operator and Perron-Frobenius analyses describe how smooth flows can be represented by spectral decompositions of composition operators [10, 37, 43, 44]. These frameworks make it possible to analyze recurrent or iterative architectures using the language of eigenfunctions and invariant subspaces. The SCT extends this line of thought from continuous physical systems to semantic ones: the transfer operator on latent space plays the role of the Koopman operator, and its discrete spectrum identifies the finite number of semantic attractors that define symbolic behavior.

5.5 O-Minimal Geometry and Model Theory

From the logical perspective, the SCT draws directly on the mathematics of definable sets in o-minimal structures [13, 59, 60]. O-minimality guarantees “tame” topological behavior, in the sense of: finite numbers of connected components, absence of fractal boundaries, and algebraic regularity. By showing that the definable regions of a model’s latent space satisfy these conditions, the theorem links semantic stability to logical tameness. This connection grounds the otherwise empirical notion of “interpretability” in formal logic: a region of representation space is interpretable precisely when it is definable in an o-minimal expansion of the reals.

As shown by Tressl [58], o-minimal structures not only ensure logical tameness but also furnish a natural framework for analyzing the learnability and generalization properties of neural networks, reinforcing their central role in modern semantic modeling.

5.6 Cognitive Semantics and Conceptual Spaces

Cognitive models of meaning have long posited that concepts occupy geometric regions in a continuous space of qualities [23]. Work in dimensionality-reduction and manifold learning likewise seeks geometric coordinates that preserve semantic relationships [55]. Neurosemantic studies further support the existence of low-dimensional conceptual organization in the brain [7]. The SCT provides a formal bridge between these ideas and machine learning: it proves that a neural representation can remain continuous while being partitioned into finitely many definable regions that correspond to conceptual categories, thereby realizing Gärdenfors’s “geometry of thought” in operational form.

5.7 Symbolic Emergence and Hybrid Systems

The problem of how discrete symbols arise from continuous neural activity has animated decades of debate between connectionism and symbolic AI [27, 54]. More recent accounts emphasize compositional reasoning and structured abstraction as emergent properties of large models [6, 34]. The SCT offers a unifying mathematical account: symbolic structure appears when a continuous dynamical system becomes spectrally and logically tame. In this view, symbolic emergence is not an architectural add-on but an intrinsic phase transition in representational geometry.

5.8 Relation to World-Model Hypotheses

A significant thread in contemporary AI research, exemplified by LeCun and others, argues that the emergence of general intelligence requires an internal *world model*, which is effectively a learned representation of the physical environment that supports prediction, planning, and causally grounded reasoning. In contrast, scaling-based approaches maintain that continued expansion of large-language models (LLMs) will eventually yield such capabilities through sheer statistical expressiveness.

Within the framework of the Semantic Characterization Theorem (SCT), we propose a third position: the *world* is not external to the model but rather a definable subset of its latent semantic manifold. Language itself encodes the causal and spatial regularities of embodied experience, so a sufficiently rich linguistic model *already contains* an implicit world model. What is lacking is thus not grounding, but *semantic resolution* at the level of the precision needed for its internal attractors to project consistently onto the perceptual substrate. To phrase it more concisely: We can’t think what we can’t perceive, so there’s no need to model it.

Lemma 5.2 (World–Manifold Equivalence). *Let M be a continuous state machine with definable attractor manifold Σ_M . For any externally constructed “world model” W derived from sensory or environmental data, there exists a definable submanifold $\Sigma_W \subseteq \Sigma_M$ such that the predictive transformations of W are equivalent to the definable morphisms within Σ_M under semantic collapse:*

$$T_W(s, u) \equiv T_M(f_M(s), f_U(u)) \quad \text{for definable maps } f_M, f_U.$$

Hence, embodiment reduces to internal definability rather than architectural augmentation.

Context and related work. The above formulation draws conceptual continuity from several distinct research traditions. LeCun’s [2022] proposal of energy-based *world models* outlines an architectural route toward autonomous intelligence, while earlier work by Ha and Schmidhuber [2018] demonstrates practical latent simulators that learn compact internal environments. Cognitive evidence for this approach is surveyed by Lake et al. [2017], who argue that human intelligence depends on compositional and causal models of the world. In contrast, our definability framework is grounded in model theory and o-minimal geometry as introduced by Tressl [2015], extending these ideas to continuous state machines. Complementary perspectives by Hinton [2021] show that part–whole hierarchies and spatial relations can be encoded intrinsically within neural representations, implying that world-like structure need not be externally simulated. Finally, Barwise and Perry’s *Situations and Attitudes* [1983] provides a philosophical precedent for treating “worlds” as internal situational substructures of meaning rather than ontologically separate domains. Together, these lines of inquiry converge on the view that the semantics of a sufficiently expressive model already entail the structures it would otherwise be tasked with learning: the world, in a strong sense, is definable within the manifold.

Implications. Under this view, scaling alone cannot yield stronger intelligence because it merely expands the manifold without enforcing definability. Conversely, constructing a separate world simulator is redundant, since the model’s manifold already contains latent subspaces isomorphic to world-like dynamics. The critical transition is one of *semantic phase change*, in which regions of the manifold become compact, convex, and projectable onto physically meaningful axes. This provides a formal reconciliation between linguistic and embodied paradigms: what LeCun calls a “world model” is, in categorical terms, a definable subobject of the model’s own semantic space.

5.9 Semantic failure and the null attractor

One of the main qualitative differences between our framework and large language models is that a continuous state machine (CSM) can *refuse* to assign meaning in a principled way. Informally: the system has a canonical, geometrically defined “I don’t know” state.

In token-prediction models, hallucination is structurally unavoidable: the architecture is forced to produce a next token, even when the input is out of distribution or semantically incoherent. In our setting, by contrast, the dynamics of the CSM are constrained by an ontology and an induced geometry on state space, and these constraints allow us to distinguish *semantic realization* from *semantic failure*.

Definition 5.3 (Semantic and null attractors). Let X be the state space of a CSM with transition map $T: X \rightarrow X$ and let $\mathcal{O}$ be an ontology whose concepts are realized as stable regions in X .

1. A point $x^* \in X$ is an *attractor* if there exists a neighborhood $U \subseteq X$ such that for all $x \in U$, the forward orbit $T^n(x)$ converges to x^* as $n \rightarrow \infty$. The corresponding basin of attraction $B(x^*) \subseteq X$ consists of all points whose orbits converge to x^* .
2. An attractor x^* is *semantic* if its basin $B(x^*)$ corresponds to a definable, nontrivial concept in $\mathcal{O}$ (for example, the concept of a group, a ring, or a category). In this case we say that x^* realizes a concept $c \in \mathcal{O}$.
3. An attractor $x_\perp \in X$ is a *null attractor* if its basin $B(x_\perp)$ consists precisely of those states which do not lie in the basin of any semantic attractor, and $B(x_\perp)$ carries no nontrivial internal structure (e.g. it is geometrically symmetric or collapses to a trivial cell).

Intuitively, semantic attractors correspond to well-formed, ontologically grounded meanings, whereas the null attractor $x_\perp$ represents the absence of realizable meaning in the given ontology-geometry pair.

Proposition 5.4 (Semantic realization vs. semantic failure). *Let $T: X \rightarrow X$ be a CSM whose dynamics are induced by an ontology $\mathcal{O}$ and an associated o -minimal geometry on X . Suppose that:*

1. *Each concept $c \in \mathcal{O}$ is realized by a semantic attractor x_c with a definable basin $B(x_c)$.*
2. *The family of semantic basins $\{B(x_c) : c \in \mathcal{O}\}$ is pairwise disjoint and their union is strictly contained in X .*

Then there exists a null attractor $x_\perp$ such that:

1. *For any input state $x \in \bigcup_{c \in \mathcal{O}} B(x_c)$, the orbit $T^n(x)$ converges to the corresponding semantic attractor x_c ; in particular, the input is assigned a canonical semantic interpretation*

determined by $\mathcal{O}$.

2. For any input state $x \notin \bigcup_{c \in \mathcal{O}} B(x_c)$, the orbit $T^n(x)$ converges to the null attractor $x_\perp$. In this case the system does not assign a semantic interpretation: it returns a distinguished “unknown” state.

In other words, the CSM separates inputs into two classes purely by its geometry: those that are *semantically realizable* in the ontology-induced state space, and those that are not. The former converge to rich, structured attractors associated with concepts in $\mathcal{O}$; the latter collapse into the null attractor $x_\perp$, which plays the role of a canonical “I don’t know” state.

Crucially, this mechanism is not “negation by failure” in the operational sense of logic programming, nor an ad hoc refusal heuristic. It is *negation by semantic failure*: an input is rejected when it does not lie in the basin of any semantic attractor, i.e. when it cannot be realized as a definable point in the ontology-induced geometry. The ability of the system to say “I don’t know” is therefore a direct consequence of the interaction between the ontology and the continuous dynamics, rather than an externally imposed rule.

5.10 Synthesis

In sum, these lines of research converge on a single insight: the boundary between continuous and discrete semantics is not arbitrary but lawful.

The Semantic Characterization Theorem formalizes this convergence, providing the first unified account that connects o-minimal definability, spectral lumpability, and the geometry of learned representations. By situating LLM semantics within this cross-disciplinary landscape, the present work reframes interpretability and symbolic emergence as natural consequences of the underlying mathematical structure of learning.

6 Discussion

The Semantic Characterization Theorem provides a mathematical account of why large language models exhibit structured and predictable behavior despite operating in high-dimensional continuous spaces. Although the theorem is purely structural, its consequences touch directly on how such models can be prompted, programmed, and interpreted in practice. We briefly outline these implications here.

Deducing an Ontological Skeleton The natural next question is ontological: having shown that the latent manifold of a large language model undergoes a definable spectral collapse, what *categories* constitute its finite semantic skeleton? The present theorem guarantees that such a skeleton exists and is compact, but not yet what its vertices represent. Preliminary analyses suggest that the emergent ontology is both interpretable and universal, aligning with core archetypal distinctions observed across human cognition, i.e., the spatial, affective, causal, and narrative axes that together span a low-rank geometry of meaning. Future work will characterize these axes explicitly, providing the first empirical catalogue of the definable regions that arise when a continuous system learns to “think.”

Prompting as semantic navigation. If the latent manifold collapses into finitely many definable regions $\{B_1, \dots, B_r\}$, then a prompt acts as a coordinate choice that initializes the model in (or near) one of these basins. Small perturbations of a prompt do not produce qualitatively different

behavior unless they push the state across a basin boundary. This explains why certain instructions appear “robust” while others require carefully tuned phrasing: robustness corresponds to interior points of a basin, while brittleness arises near its boundary. The SCT therefore suggests a geometric interpretation of prompting in which the goal is to steer trajectories toward the desired definable region of the semantic space.

Programming via basin-aware composition. Because the dominant eigenfunctions of the transfer operator form a finite invariant subspace, the long-term behavior of a chat or computation can be viewed as a sequence of transitions among the definable basins. This provides a coarse form of compositionality: a program or workflow corresponds to a path in the ontological skeleton, with each step inducing a controlled transition under the map T or the kernel K . Although the present paper develops only the structural underpinnings, the existence of a finite semantic graph suggests that higher-level control mechanisms, such as tool-use, multistep plans, and chain-of-thought reasoning, can be construed as procedures for orchestrating basin transitions. This offers a principled route toward semantic programming without requiring explicit symbolic modules.

Interpretability through definable structure. Logical tameness implies that each basin B_i is definable in an o-minimal structure and therefore possesses bounded geometric and topological complexity. In practical terms, this means that observable distinctions among behaviors correspond to definable predicates over the latent space. The boundaries between behaviors need not be explicitly learned; they are constraints inherited from the regularity of the dynamics. This perspective unifies several empirical observations: the tendency of low-capacity probes to separate semantic features, the clustering of embeddings along a few principal axes, and the stability of model behavior under reparameterization. In each case, interpretability emerges not from architectural design but from definable invariance.

Toward a semantic phase theory. Together, these observations hint at a broader picture in which prompting, programming, and interpreting correspond to different modes of interacting with the same finite semantic skeleton. Prompting selects an initial basin, programming composes transitions among basins, and interpretability examines the definable predicates that separate them. The SCT provides the mathematical constraints that make this correspondence possible. A full operational calculus for manipulating definable regions is beyond the scope of this paper, but the results presented here establish the structural foundations for such a theory.

A Ergodicity and Spectral Compactness of LLM Dynamics

This appendix provides sufficient conditions under which assumptions 2.8(A4)–(A5) hold for a large language model (LLM) viewed as a CCSM. The results below rely on classical arguments from the theory of Feller Markov processes and compact integral operators.

Assumption A.1 (Mixing and smoothing). The assumptions from 2.8 can be proved from two straightforward properties.

- (a) (*Feller + minorization*) The Markov kernel K induced by the stochastic transition T and decode policy Δ is Feller on the compact space $\mathcal{M}$ and satisfies a Doeblin-minorization condition: there exist $\varepsilon > 0$ and a probability measure ν on $(\mathcal{M}, \mathcal{B}(\mathcal{M}))$ such that

$$K(s, A) \geq \varepsilon \nu(A) \quad \text{for all } s \in \mathcal{M}, A \in \mathcal{B}(\mathcal{M}).$$

- (b) (*Continuous density*) K admits a density $k(s, s')$ with respect to its stationary measure μ that is continuous on $\mathcal{M} \times \mathcal{M}$.

Remark A.2 (Practical sufficient condition). If, at decode time, the policy Δ allocates at least $\varepsilon > 0$ total probability mass uniformly across a fixed token subset (e.g., via ε -greedy exploration or temperature bounded away from 0), then Assumption A.1(a) holds by the Doeblin minorization condition. If in addition T is C^1 and the token embedding map is continuous, the induced kernel K possesses a continuous density $k(s, s')$, establishing Assumption A.1(b).

Lemma A.3 (Existence and ergodicity). *Under Assumption A.1(a), the kernel K admits at least one stationary probability measure μ , and the corresponding Markov chain $(s_t)_{t \geq 0}$ is geometrically ergodic: for all bounded measurable f ,*

$$\|\mathbb{E}[f(s_t)] - \int f d\mu\| \leq C\rho^t \quad \text{for some } C < \infty, 0 < \rho < 1.$$

Proof. Compactness and the Feller property ensure existence of a stationary measure (Krylov–Bogolyubov theorem). The Doeblin minorization implies uniform ergodicity by classical results (e.g., 42, Theorem 16.0.2). $\square$

Lemma A.4 (Compact transfer operator). *Under Assumption A.1(b), the transfer operator*

$$(Pf)(s) := \int_{\mathcal{M}} k(s, s') f(s') d\mu(s')$$

is a compact linear operator on $L^2(\mathcal{M}, \mu)$ and thus admits a discrete spectrum $\{\lambda_i\}_{i \geq 1}$ with $|\lambda_i| \downarrow 0$.

Proof. Continuity of k on compact $\mathcal{M} \times \mathcal{M}$ implies uniform boundedness and equicontinuity of $\{Pf : \|f\|_2 \leq 1\}$ by Arzelà–Ascoli, so P is compact. Spectral properties then follow from the Hilbert–Schmidt theorem. $\square$

Lemma A.5 (Deterministic fallback: finite-horizon compactness). *For a deterministic C^1 map $T : \mathcal{M} \rightarrow \mathcal{M}$ on compact $\mathcal{M}$, define the finite-horizon propagator*

$$(H_\tau f)(s) := \frac{1}{\tau} \sum_{t=0}^{\tau-1} f(T^t(s)).$$

Then H_τ is a compact operator on $C(\mathcal{M})$ for each finite τ , and its dominant spectral subspace converges, as $\tau \rightarrow \infty$, to that of the associated invariant measure(s) of T .

Together, Lemmas A.3–A.5 justify assumptions 2.8(A4)–(A5) for both stochastic and deterministic LLM dynamics. In particular, any nonzero exploration component in decoding policy (however small) suffices to ensure ergodicity and spectral compactness, providing the theoretical foundation for the spectral-collapse results used in the main text.

B Extension of the SCT to Stochastic Dynamics

This appendix outlines the extension of the Semantic Characterization Theorem (SCT) to stochastic dynamics. The deterministic case developed in the main text treated the transition operator $T : M \rightarrow M$ as smooth and autonomous. Here we replace T by the induced Markov kernel $K : M \times \mathcal{B}(M) \rightarrow [0, 1]$ and show that the essential results (spectral lumpability, logical tameness, and their equivalence) continue to hold under mild regularity conditions.

B.1 Stochastic Setting and Assumptions

Let K satisfy the ergodicity and continuity assumptions of Appendix A (Assumption A.1). In particular, we assume:

- (i) Compact measurable state space $M \subset \mathbb{R}^d$;
- (ii) Existence of a stationary distribution μ absolutely continuous with respect to Lebesgue measure;
- (iii) Continuous density $k(s, s')$ of the kernel;
- (iv) Geometric ergodicity of the induced Markov chain.

Under these conditions, the transfer operator

$$(Pf)(s) = \int_M k(s, s') f(s') d\mu(s')$$

is compact on $L^2(M, \mu)$ (Lemma A.3) and therefore admits a discrete spectrum $\{\lambda_i\}_{i \geq 1}$ with $|\lambda_i| \downarrow 0$.

B.2 Spectral Lumpability

Lemma B.1 (Spectral gap for Markov operators). *If K satisfies the Doeblin–minorization condition of Assumption A.1(a), then P has a nonzero spectral gap: $|\lambda_1| > |\lambda_2|$. Consequently, the leading eigenfunctions $\{\phi_1, \dots, \phi_r\}$ define finitely many metastable basins B_i whose union covers M up to measure zero.*

The proof follows directly from standard results on lumpability of ergodic Markov processes with compact kernels (e.g., [42], [10]).

B.3 Logical Tameness under Expectation Operators

Proposition B.2 (Definability of integral operators). *Suppose the kernel $k(s, s')$ is definable in an $\mathcal{o}$ -minimal structure over $\mathbb{R}$. Then the integral operator P preserves definability: if f is definable, so is Pf , and the level sets $\{s : Pf(s) = c\}$ are definable.*

Hence the stochastic transfer operator remains logically tame, and the definable partition of M is preserved in expectation.

B.3 Stochastic Semantic Characterization Theorem

Corollary B.3 (Stochastic SCT). *Under Assumptions A.1–A.3, the spectral and logical partitions of M induced by P coincide almost surely under the invariant measure μ . Therefore the conclusions of Theorem 3.2 hold for stochastic continuous state machines as well.*

Interpretation. Randomness in decoding or control does not destroy semantic discreteness; it merely smooths the boundaries between definable regions. The stochastic SCT thus generalizes the deterministic result, showing that both spectral and logical tameness are robust to probabilistic perturbations.

C A Categorical Sketch for the SCT

This section outlines a functorial formulation of the SCT, positioning Continuous State Machines (CSMs) as coalgebras for a probability functor, and their transfer operators as images in a compact-operator subcategory of Hilbert spaces. The goal is to bridge algebraic/self-expressive semantics with smooth, stochastic dynamics, preparing a categorical generalization of the results in [18].

C.1 Categories and Functors

Base categories. Let **Meas** be the category of standard Borel spaces and measurable maps. Let **CSM** be the category whose objects are tuples

$$\mathcal{C} = (M, U, T, O, \Delta, \mu),$$

where $M \subset \mathbb{R}^d$ is compact, $U \subset \text{Dist}(X)$, $T : M \times U \rightarrow M$ is C^1 in the first argument, $O : M \rightarrow \text{Dist}(X)$, $\Delta : \text{Dist}(X) \rightarrow U$, and μ is a stationary probability measure for the induced kernel K (deterministic limit given by $Pf = f \circ T$). See Assumptions (A1)–(A5) for compactness/ergodicity/spectral boundedness and their stochastic counterparts.⁴

A morphism $F : \mathcal{C} \rightarrow \mathcal{C}'$ in **CSM** consists of measurable maps

$$F_M : M \rightarrow M', \quad F_U : U \rightarrow U',$$

that *intertwine* dynamics and controls in the sense that pushing forward the joint law of (s, u) and then applying T' coincides (a.e.) with applying T and then pushing forward; and $F_M \mu = \mu'$. Intuitively, F is a semantics-preserving simulation. Pictorially:

$$\begin{array}{ccc} M \times U & \xrightarrow{T} & M \\ F_M \times F_U \downarrow & & \downarrow F_M \\ M' \times U' & \xrightarrow{T'} & M' \end{array}$$

The commutativity condition $T' \circ (F_M \times F_U) = F_M \circ T$ means that the diagram commutes: evolving first in the source CSM and then mapping to the target (upper-then-right path) yields the same result as mapping first and evolving under the target's transition (left-then-bottom path).

⁴Cf. Assumptions (A1)–(A5), Remarks/Lemmas in §1–§2 and Appendix A for kernel compactness and ergodicity, and the main SCT statement in §2.2–§2.3.

Distribution monad / coalgebras. Let $\mathcal{D} : \mathbf{Meas} \rightarrow \mathbf{Meas}$ be the (finitely supported) distribution monad. A *stochastic transition* is a coalgebra $\alpha : M \rightarrow \mathcal{D}(M)$. In a CSM, the induced one-step kernel $K(s, \cdot)$ defines α , and the transfer operator is the Kleisli lifting acting on observables

$$(Pf)(s) = \int f(s') K(s, ds').$$

Under continuity of the density $k(s, s')$ on compact $M \times M$, P is Hilbert–Schmidt and thus compact on $L^2(M, \mu)$.⁵

Hilbert target and intertwiners. Let $\mathbf{Hilb}$ be the category of complex Hilbert spaces and bounded linear maps, and let $\mathbf{cHilb}$ be its full subcategory on spaces with compact endomorphisms. Define a functor

$$\mathcal{P} : \mathbf{CSM} \longrightarrow \mathbf{cHilb}, \quad \mathcal{P}(C) = (L^2(M, \mu), P),$$

which sends a morphism F to an *intertwiner* $W_F : L^2(M', \mu') \rightarrow L^2(M, \mu)$, chosen as the pullback $W_F(g) = g \circ F_M$ (or, in general, a Markov operator associated to F) so that $W_F P' = P W_F$ (naturality). When k is continuous and M compact, P is compact; spectral decompositions in §2 follow.⁶

C.2 Natural Structure: Observation and Control

The observation O and policy Δ assemble into natural transformations between functors that select control/observation spaces from $\mathbf{CSM}$ objects, and the distribution monad $\mathcal{D}$. In particular, the composite $M \xrightarrow{O} \text{Dist}(X) \xrightarrow{\Delta} U$ yields the coalgebra law $M \rightarrow \mathcal{D}(M)$ via T ; in the deterministic limit, this reduces to composition $f \mapsto f \circ T$.

C.3 Categorical SCT (Sketch)

Theorem C.1 (Categorical SCT, sketch). *Let $C \in \mathbf{CSM}$ satisfy compactness/ergodicity/spectral boundedness (§1.2, Appendix A). Then $\mathcal{P}(C) = (H, P)$ lies in $\mathbf{cHilb}$ and admits a discrete spectral decomposition with leading eigenfunctions $\{\phi_i\}_{i=1}^r$ and a spectral gap. The induced spectral partition functor $\mathcal{B} : \mathbf{CSM} \rightarrow \mathbf{Part}$ assigns to C the partition $\{B_i\}$ where $B_i = \{s : |\phi_i(s)| = \max_j |\phi_j(s)|\}$. Moreover, if T (or k) is definable in an o-minimal expansion of $\mathbb{R}$, there exists a functor $\mathcal{C}_{\text{def}} : \mathbf{CSM} \rightarrow \mathbf{Cell}$ producing a finite cell decomposition by definable regions. The constructions are naturally equivalent up to μ -null sets:*

$$\mathcal{B} \cong \mathcal{C}_{\text{def}} \quad (a.e.).$$

Rationale. Compactness $\Rightarrow$ discrete spectrum (Hilbert–Schmidt); spectral gap $\Rightarrow$ finite-rank semantics (§2.1); definability of $T/k \Rightarrow$ o-minimal tameness and finite cell decomposition (§2.2). Equivalence follows since eigenlevel sets are definable and differ from cells only on boundaries of measure zero (§2.3).⁷

⁵See Lemma A.4 (compact transfer operator via continuous density on compact domain).

⁶Lemma 2.3 and Proposition 2.4 establish discrete spectrum and spectral collapse for compact P .

⁷See Theorem 2.2 and Lemmas/Propositions in §2.1–§2.3.

C.4 Time-Dependent (Adiabatic) Case

For a family $(\mathcal{C}_t)$ with operators (P_t) satisfying the uniform compactness/minorization/slow-drift hypotheses, the functor $\mathcal{P}$ lifts to a functor into a path category of compact operators, and the leading spectral subbundles vary continuously; the adiabatic corollary shows that products $P_{t+n-1} \cdots P_t$ are equivalent (to first order) to P^n of a stationary model trained on the narrative distribution.

C.5 Outlook

This functorial framing sets up a categorical generalization of self-expressive algebra to CSMs: algebraic fixed points correspond to compact spectral modes; coalgebraic dynamics and observation form natural transformations; and o-minimal definability supplies logical tameness. The next step is to axiomatize **CSM** and the intertwiner class so that $\mathcal{P}$ is a faithful semantics functor, enabling categorical proofs of SCT-like results.

D A Nonstandard-Analysis Interpretation of Semantic Collapse

In this appendix we provide an alternate proof of Theorem 3.2(i) using nonstandard analysis (NSA). The NSA perspective highlights a structural equivalence between *semantic collapse* in a Continuous State Machine and the appearance of *standard parts* in an o-minimal enlargement of the underlying state space. This yields a conceptual interpretation of semantic convergence as the elimination of infinitesimal semantic variation.

Throughout, let $M \subseteq \mathbb{R}^d$ be a compact definable set in an o-minimal expansion of the real field, and let $T : M \rightarrow M$ be a definable map satisfying the assumptions of the main theorem (Hilbert–Schmidt structure, spectral decay, and contractive dynamics on each cell).

D.1 The Nonstandard Extension

Let *M denote the nonstandard enlargement of M , obtained for instance by a saturated ultrapower construction. Elements of *M include both:

- *standard points* $x \in M$, and
- *nearstandard points* $x^* \in {}^*M$ satisfying $x^* \approx x$ for some $x \in M$, i.e. $x^* - x$ is infinitesimal,

together with truly *nonstandard* points not infinitesimally close to any $x \in M$.

Each $x^* \in {}^*M$ that is nearstandard admits a unique *standard part* $\text{st}(x^*) \in M$ such that $x^* \approx \text{st}(x^*)$.

The enlargement *M carries a canonical decomposition into *halos* $\eta(x) := \{x^* \in {}^*M : x^* \approx x\}$ for $x \in M$, plus an external region of non-nearstandard points.

D.2 Hilbert–Schmidt Structure and Infinitesimal Collapse

Let K be the Hilbert–Schmidt operator associated with the semantics of the system as described in Section 1. By compactness of M and definability of T , we obtain a Hilbert–Schmidt resolvent whose spectral expansion takes the form

$$Ku = \sum_{n=1}^{\infty} \lambda_n \langle u, e_n \rangle e_n, \quad \lambda_1 \geq \lambda_2 \geq \cdots \rightarrow 0.$$

The key point is that the tail $\sum_{n>N} \lambda_n \langle u, e_n \rangle e_n$ becomes *infinitesimal* in norm as $N \rightarrow \infty$. In the nonstandard setting, this tail is literally an *infinitesimal vector*. Thus the action of K sends any state to a nearstandard point inside the halo of a finite-dimensional attractor cell.

In particular, if $x^* \in {}^*M$ is any nonstandard state interpreted under the dynamical rule T , then repeated iteration yields a sequence

$$x_0^* = x^*, \quad x_{n+1}^* = {}^*T(x_n^*)$$

which becomes *eventually nearstandard*: there exists N such that for all $n \geq N$, x_n^* lies in the halo of a unique fixed point $x^\dagger \in M$. The reason is precisely that all high-frequency components in the spectral decomposition are infinitesimal and vanish under iteration.

D.3 O-Minimality and Definable Cells

By the cell decomposition theorem for o-minimal structures, M admits a finite definable partition

$$M = C_1 \cup \dots \cup C_k$$

into cells C_i , each of which is definably homeomorphic to an open box. Each cell is dynamically coherent under T by assumption (A2) of the main theorem: T is contractive on each cell. This provides a canonical finite set of *semantic attractors*

$$a_1, \dots, a_k \in M, \quad T(a_i) = a_i,$$

one for each cell.

By the above NSA argument, any nearstandard state collapses under iteration into the halo $\eta(a_i)$ of exactly one attractor. Therefore the composition

$$\Phi := \text{st} \circ {}^*T^{(\infty)} : {}^*M \longrightarrow \{a_1, \dots, a_k\}$$

is well-defined: it takes a nonstandard state x^* to the standard part of its eventual limit.

D.4 Semantic Collapse as Standard-Part Selection

We may now restate Part (i) of the main theorem as follows:

Proposition D.1 (NSA formulation of semantic collapse). *Let *M be the nonstandard extension of M and let T satisfy the assumptions of Theorem 1. Then every nearstandard state $x^* \in {}^*M$ satisfies*

$$\text{st}({}^*T^{(\infty)}(x^*)) = a_i$$

for a unique semantic attractor $a_i \in M$. Moreover, the mapping $x^ \mapsto a_i$ is constant on each halo $\eta(a_i)$ and sends all non-nearstandard states to an external “null” attractor.*

Proof. By definability and o-minimality, M decomposes into finitely many contractive cells whose fixed points are exactly the a_i . By Hilbert–Schmidt spectral decay, the nonstandard iteration $x_{n+1}^* = {}^*T(x_n^*)$ suppresses all infinite and high-frequency components, rendering x_n^* nearstandard for all large enough n . Contractivity then forces x_n^* into the halo of a unique a_i , and taking standard parts yields the desired conclusion.

If x^* is non-nearstandard, then its orbit under *T remains external and collapses into an external fixed point corresponding to the null attractor. Uniqueness follows from disjointness of halos and the fact that the standard-part map is well-defined only on nearstandard points. $\square$

D.5 Interpretation

In NSA terms, semantic collapse is the assertion that the semantics of the system consist of the *standard parts* of eventual CSM trajectories. Definable semantics appear as standard points; semantic uncertainty appears as infinitesimal halos; and semantic failure appears as external states collapsing to a null attractor.

Thus the CSM implements a canonical “standard-part selection” mechanism over an o-minimal geometry, and semantic attractors correspond precisely to standard points in the nonstandard extension.

D.6 A Note on Confidence

Each prediction comes with a natural “confidence” measure (a real number) and that is as follows.

Given a nearstandard state $x_{N^*}^* \in {}^*M$ with standard part $x^\dagger = \text{st}(x_{N^*}^*)$, define the infinitesimal remainder $\varepsilon := x_{N^*}^* - x^\dagger$. We define the *nonstandard semantic confidence* of the trajectory $x_0^*, x_1^*, \dots, x_{N^*}^*$ to be

$$\text{conf}(x^*) := \exp(-\alpha \|\varepsilon\|),$$

for a fixed scaling constant $\alpha > 0$. Thus $\text{conf}(x^*) \approx 1$ precisely when the nonstandard limit is deeply inside the halo of a semantic attractor, and $\text{conf}(x^*) \approx 0$ when the trajectory remains external and collapses to the null attractor.

Given this, we can define a threshold, call it τ , that determines when we will accept a given determination. In other words, if $\text{conf}(x^*) < \tau$ then the system responds, “I don’t know.”

References

- [1] Alessandro Achille and Stefano Soatto. Emergence of invariance and disentangling in deep representations. *Journal of Machine Learning Research*, 19(50):1–34, 2018.
- [2] Peter Aczel. *Non-Well-Founded Sets*. CSLI Publications, Stanford, CA, 1988.
- [3] Steve Awodey. *Category Theory*. Oxford University Press, 2nd edition, 2010.
- [4] H. B. Barlow. Possible principles underlying the transformations of sensory messages. In W.A. Rosenblith (Ed.), *Sensory Communication*, pages 217–234, 1961.
- [5] Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, Cambridge, MA, 1983.
- [6] Yoshua Bengio. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2019.
- [7] Jeffrey R. Binder and Rutvik H. Desai. The neurobiology of semantic memory. *Trends in Cognitive Sciences*, 15(11):527–536, 2011. doi: 10.1016/j.tics.2011.10.001.
- [8] Ludwig Boltzmann. *Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht*, volume 76. 1877.
- [9] Michael S. Branicky. Multiple lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):475–482, 1998.
- [10] Steven L. Brunton and J. Nathan Kutz. Koopman analysis of dynamical systems. *Nature Reviews Physics*, 3(6):422–433, 2021. doi: 10.1038/s42254-021-00305-5. URL <https://doi.org/10.1038/s42254-021-00305-5>.

- [11] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [12] Robin Cooper, Jonathan Ginzburg, and Maarten de Rijke. *Situation Theory and Its Applications*. CSLI Publications, Stanford, CA, 1994.
- [13] Michel Coste. An introduction to o-minimal geometry. In *Dottorato di Ricerca in Matematica, Dipartimento di Matematica, Università di Pisa*, pages 1–65. 2000. Available at <https://perso.univ-rennes1.fr/michel.coste/polyens/ominimal.pdf>.
- [14] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2nd edition, 2006.
- [15] Keith Devlin. *Logic and Information*. Cambridge University Press, Cambridge, UK, 1991.
- [16] Nelson Elhage, Neel Nanda, Catherine Olsson, et al. Toy models of superposition. Anthropic Interpretability Series, 2022. arXiv:2209.13085.
- [17] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
- [18] Frank M. V. Feys, Helle Hvid Hansen, and Lawrence S. Moss. Long-term values in Markov decision processes, (co)algebraically. *Electronic Proceedings in Theoretical Computer Science*, 2024. URL https://www.coalg.org/cmcs18/files/2018/04/Feys_Hansen_Moss.pdf. Preprint version consulted, Delft University of Technology and Indiana University.
- [19] J. Willard Gibbs. *Elementary Principles in Statistical Mechanics*. Yale University Press, 1902.
- [20] Rafael Goebel, Ricardo G. Sanfelice, and Andrew R. Teel. *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press, 2009.
- [21] Peter D. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.
- [22] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations (ICLR)*, 2022.
- [23] Peter Gärdenfors. *Conceptual Spaces: The Geometry of Thought*. MIT Press, Cambridge, MA, 2000.
- [24] David Ha and Jürgen Schmidhuber. World models. In *Neural Information Processing Systems (NeurIPS)*, 2018. Demonstrates internal latent simulators trained from sensory data; early formalization of the world-model approach.
- [25] Esfandiar Haghverdi. *A Categorical Approach to Hybrid Systems*. PhD thesis, University of Ottawa, 1999.
- [26] Esfandiar Haghverdi. Hybrid systems as Coalgebras. In *Electronic Notes in Theoretical Computer Science*, volume 21, pages 1–20. Elsevier, 2000. doi: 10.1016/S1571-0661(04)00279-4.
- [27] Stevan Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1–3): 335–346, 1990.
- [28] John Hewitt and Christopher D. Manning. A structural probe for finding syntax in word representations. In *Proceedings of NAACL-HLT*, pages 4129–4138, 2019.

- [29] Geoffrey Hinton. How to represent part–whole hierarchies in a neural network. *arXiv preprint arXiv:2102.12627*, 2021. Argues that internal geometric structures can encode spatial relations without explicit 3D world simulation.
- [30] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [31] Eduard Hovy. The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 30(4):8–12, 2015. doi: 10.1109/MIS.2015.57.
- [32] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620–630, 1957.
- [33] Minyoung Kim, Yubei Chen, and Bruno A. Olshausen. Recurrent neural processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [34] Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253, 2017. doi: 10.1017/S0140525X16001837.
- [35] Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253, 2017. Surveys cognitive science evidence for compositional, causal, and world-model representations underlying human intelligence.
- [36] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.
- [37] Samuel Lanthaler, Hassan Arbabi, and J. Nathan Kutz. Koopman operator learning for dynamical systems. *Physical Review Research*, 4(2):023151, 2022. doi: 10.1103/PhysRevResearch.4.023151.
- [38] F. William Lawvere. Functorial semantics of algebraic theories. In *Proc. Nat. Acad. Sci. USA*, volume 50, pages 869–872. 1963.
- [39] Yann LeCun. A path towards autonomous machine intelligence. *arXiv preprint arXiv:2209.06794*, 2022. Proposes the world-model architecture with energy-based predictive modules as an alternative to scaling LLMs.
- [40] Leonid Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [41] Haotian Liu, Jian Li, Pengcheng Zhao, et al. Understanding and improving layer normalization. In *Advances in Neural Information Processing Systems*, 2021.
- [42] Sean P. Meyn and Richard L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009. ISBN 9780521731829. doi: 10.1017/CBO9780511626630.
- [43] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.
- [44] Igor Mezić. Analysis of fluid flows via spectral properties of the koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.

- [45] Natalya F. Noy and Deborah L. McGuinness. Ontology development 101: A guide to creating your first ontology. Technical Report KSL-01-05, Stanford Knowledge Systems Laboratory, Stanford, CA, 2001. URL https://protege.stanford.edu/publications/ontology_development/ontology101.pdf.
- [46] Vardan Papyan, X.Y. Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. doi: 10.1073/pnas.2015509117.
- [47] Ben Poole, Saurabh Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [48] Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.
- [49] Lars Ruthotto and Eldad Haber. An introduction to deep learning for the sciences. *SIAM Review*, 64(3):640–686, 2022.
- [50] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *International Conference on Learning Representations*, 2014.
- [51] Claude E. Shannon. Mathematical theory of the differential analyzer. *Journal of Mathematics and Physics*, 20:337–354, 1941.
- [52] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [53] Hava T. Siegelmann and Eduardo D. Sontag. On the computational power of neural nets. *Journal of Computer and System Sciences*, 50(1):132–150, 1995.
- [54] Paul Smolensky. On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1):1–74, 1988.
- [55] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.5500.2319.
- [56] Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. In *Proc. 37th Annual Allerton Conf. on Communication, Control and Computing*, pages 368–377, 1999.
- [57] Boris A. Trakhtenbrot. Impossibility of an algorithm for the decision problem on finite classes. *Doklady Akademii Nauk SSSR*, 70:569–572, 1950.
- [58] Marcus Tressl. Introduction to o-minimal structures and an application to neural network learning. *Lecture Notes, University of Manchester*, 2015. URL <https://personalpages.manchester.ac.uk/staff/marcus.tressl/papers/Introduction0-minimalApplicationToNeuralNetworkLearning.pdf>. Accessed November 2025.
- [59] Lou van den Dries. *Tame Topology and O-minimal Structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, UK, 1998. doi: 10.1017/CBO9780511525919.

- [60] Lou van den Dries and Chris Miller. Geometric categories and o-minimal structures. *Duke Mathematical Journal*, 84(2):497–540, 1996. doi: 10.1215/S0012-7094-96-08419-1.
- [61] Arjan van der Schaft and Hans Schumacher. *An Introduction to Hybrid Dynamical Systems*, volume 251 of *Lecture Notes in Control and Information Sciences*. Springer, London, 2000. ISBN 978-1-4471-1120-2. doi: 10.1007/978-1-4471-1120-2.
- [62] C. M. Wyss. How to tame your LLM: Semantic collapse in continuous systems. Technical Report XTR-2025-01, Exolytica AI, Indianapolis, IN, 2025.