

Geometric Context Modeling

Introducing a New Paradigm for Semantic Representation and a New Class of Artificial Intelligence System

Tiffney Bare • Independent AI Researcher • May 2026

KAIA Research Series • Introductory Position Paper

Abstract

This paper introduces Geometric Context Modeling as a new field of inquiry and a new class of artificial intelligence systems. Current AI systems are built on one of two foundations: statistical language modeling, which represents meaning as probability distributions over token sequences; or symbolic knowledge representation, which encodes meaning as discrete facts and rules. Neither approach maintains a continuous, explicit, geometric representation of where a conversation is in semantic space.

Geometric Context Modeling proposes a third foundation. Meaning is encoded as position in a structured geometric space with defined axes, distances, and directions. Context is maintained as a trajectory through that space, updated continuously by each new input using a fixed-size state vector. Routing and reasoning operate through geometric operations: projection, nearest-neighbor lookup, weighted combination, and directional navigation. No token prediction is required. No discrete symbolic facts are required. Geometry is the representation.

We document the theoretical basis for this paradigm, survey the existing work that approaches it without fully committing to it, identify the precise gap between that work and what Geometric Context Modeling proposes, and introduce KAIA (Knowledge Architecture for Intelligent Agents) as the first implemented system in this class. KAIA maintains a 52-byte context state regardless of conversation length, operates at 44,000 to 97,000 tokens per second on a standard CPU with no GPU requirement, and achieves 70 to 85 percent accuracy across a seven-benchmark suite of agent-oriented semantic reasoning tasks.

1. The Problem with Current Foundations

Two paradigms currently dominate artificial intelligence for language and meaning.

The first is statistical language modeling. Systems built on this foundation learn the probability distribution of tokens given prior tokens from massive text corpora. They predict what comes next. Transformers, large language models, and their variants are all statistical language models. They are extraordinarily capable at generation and tasks that can be framed as completion or prediction. But their representation of meaning is implicit, distributed across billions of parameters, and fundamentally tied to the statistical patterns of whichever corpus was used for training. Context in these systems is a sliding window of prior tokens, not an explicit representation of where the conversation is semantically.

The second is symbolic knowledge representation. Systems built on this foundation encode meaning as discrete facts, entities, relationships, and rules. Knowledge graphs, ontologies, expert systems, and frame-based systems are all symbolic. They are interpretable, auditable, and logically consistent within their defined domains. But they are brittle at the boundaries of what was

explicitly encoded, cannot handle the gradations and ambiguities of natural language, and require enormous manual effort to construct and maintain.

Both paradigms share a deeper limitation. Neither represents the semantic state of a conversation as a continuous geometric object. Neither maintains a persistent, explicit position in a structured meaning space that updates as the conversation progresses. Neither can tell you, at any moment in a conversation, exactly where that conversation is located semantically, with what intensity, across which dimensions, with what recent history.

This is the gap Geometric Context Modeling fills.

2. Geometric Context Modeling: The Core Proposal

Geometric Context Modeling is the practice of representing, maintaining, and reasoning conversational context as an explicit trajectory through a structured geometric semantic space.

Each of the three words in the name carries a precise meaning.

2.1 Geometric

The representation is geometric rather than statistical or symbolic. Meaning is encoded as position in a space with defined structure: axes that correspond to semantic dimensions, distances that reflect semantic similarity, and directions that correspond to semantic change. The geometric operations available in this space (projection, nearest-neighbor lookup, centroid averaging, directed stepping) are the reasoning operations. No probability distribution needs to be learned. No discrete fact needs to be asserted. The position is the meaning.

This geometric grounding has a direct empirical consequence: the representation is inspectable. You can ask where a concept sits on the temperature axis, how far it is from another concept, and in which direction it leads from its current position toward another semantic region. These are questions with precise geometric answers.

2.2 Context

The system models context rather than language. This is the central distinction from statistical language modeling. A language model asks: given what has been said, what is most likely to be said next? A context model asks: given what has been said, where are we right now in the space of meaning?

These are different questions requiring different computational structures. Next-token prediction requires a probability distribution over vocabulary. Context position requires a geometric state in semantic space. The first output is a token. The second output is the location. Generating text is optional and secondary to the context model. Knowing where you are is the primary output.

2.3 Modeling

The system is a model in the technical sense: a structured representation of the world that can be queried, updated, and reasoned over. The model is not a lookup table. It is not a trained classifier. It is a continuous, dynamic representation that changes as new information arrives and can be interrogated at any point to answer questions about the current semantic situation.

In Geometric Context Modeling, the model is the trajectory. The history of where the conversation has been, weighted by recency, is the current state. The state is explicit, readable, and auditable.

3. Related Work and Where It Falls Short

Several bodies of existing work approach Geometric Context Modeling without fully instantiating it. Understanding where each falls short clarifies what the new paradigm contributes.

3.1 Contextualized word embeddings

ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019) introduced the idea that the same word should receive different vector representations in different contexts. The vector for light in a sentence about weight should differ from the vector for light in a sentence about brightness. This is a step toward context-sensitive geometric representation.

But contextualized embeddings require running the full model on each sentence at inference time. The context sensitivity is computed anew for each sentence; it is not maintained in a persistent state. There is no trajectory, no history, no accumulated record of where the conversation has been. Each sentence is processed in isolation. And the representations are still ultimately trained through statistical objectives, not through geometric reasoning operations.

3.2 State space models

Mamba (Gu and Dao, 2024) and related structured state-space models compress the sequence history into a fixed-size state vector that evolves with each new input. This is architecturally close to the SSM component of Geometric Context Modeling. The fixed-size state is a genuine departure from the KV cache growth of transformers.

But current SSMs are designed to be better language models, not to be context models. Their state vectors are learned through next-token prediction objectives. The dimensions of the state have no semantic interpretation. You cannot ask an SSM state vector where the conversation is on the temperature axis, because its dimensions were never defined in those terms. The geometry is not structured. The trajectory is not semantic.

3.3 Semantic trajectory analysis

Recent work in cognitive science and clinical NLP (Lopes da Cunha et al., 2024; Toro-Hernandez et al., 2024) has framed concept production as movement through embedding space, computing trajectory metrics such as velocity, acceleration, entropy, and distance to centroid to characterize semantic navigation patterns. This work explicitly treats meaning as spatial and context as trajectory.

But this work analyses trajectories retrospectively for clinical or research purposes. It does not maintain a real-time trajectory for reasoning. It uses transformer-based embeddings as the substrate, not a structured geometric space with defined axes. It measures where a participant went; it does not use the trajectory state to make routing decisions.

3.4 Knowledge graphs and context ontologies

Symbolic context modeling systems (including enterprise context management systems) maintain structured representations of active entities, their relationships, and their relevance to the current conversational context. These systems are explicit and auditable, closer in spirit to Geometric Context Modeling than statistical approaches.

But they operate over discrete symbolic structures, not continuous geometric spaces. They cannot represent the graded, continuous nature of semantic position. They cannot perform geometric operations, such as finding the midpoint between two concepts or stepping along a semantic axis. The geometry that makes context reasoning possible is absent.

3.5 The gap

Approach	Continuous geometry	Persistent trajectory	Explicit axes	Semantic operations	CPU-native fixed memory
Statistical language models	No	No	No	No	No
Contextualized embeddings	Partial	No	No	Limited	No
State space models (Mamba)	No	Yes (unstructured)	No	No	Yes
Semantic trajectory analysis	Yes	Retrospective only	No	Limited	No
Symbolic context systems	No	Yes (discrete)	Partial	Symbolic only	Yes
Geometric Context Modeling	Yes	Yes (real-time)	Yes	Yes	Yes

Table 1. Geometric Context Modeling against existing approaches. No prior paradigm satisfies all five properties simultaneously.

4. The Theoretical Basis

4.1 Meaning as position

The foundational claim of Geometric Context Modeling is that meaning can be usefully represented as position in a structured geometric space. This is not a new idea in cognitive science. Conceptual space theory (Gardenfors, 2000) proposed that concepts are regions in a multi-dimensional quality space defined by perceptual and conceptual dimensions. Conceptual metaphor theory (Lakoff and Johnson, 1980) demonstrated that abstract concepts are systematically grounded in spatial and physical experience through geometric metaphors.

What is new is computational instantiation. Geometric Context Modeling takes these theoretical insights and builds a working system on them: a structured space with defined axes, efficient geometric operations, and a persistent state that tracks the trajectory through that space.

4.2 Context as trajectory

Conversation is not a sequence of tokens. It is a path through semantic space. Each word moves the current position. Some words move it far (strong semantic content), some move it slightly (function words and discourse markers), some hold it roughly in place (confirmations and acknowledgments). The path has direction, velocity, and history.

The SSM decay rate model captures this naturally. The current context state is a weighted average of all prior semantic positions, with recent positions weighed more heavily. The decay rate is a tunable parameter per axis, allowing different semantic dimensions to have different

memory characteristics. Temperature context fades differently from social context, which fades differently from temporal context.

This trajectory representation has a property that static representations cannot have: it knows not just where the conversation is, but where it has been and how it got there. Two conversations can be at the same semantic position but have arrived from different directions, and the trajectory state reflects that difference.

4.3 Physical universals and culturally contingent constructs

The geometric space is not neutral with respect to what it can represent. Physical and perceptual dimensions that are universal across human languages (temperature, speed, luminance, mass, scale) exhibit a clean geometric structure in distributional word embeddings because physical reality forces the same distributional patterns across all languages and corpora. Abstract conceptual dimensions (moral evaluation, epistemic truth, social belonging, existential freedom) have an entangled geometric structure because they are culturally contingent: different linguistic traditions encode them differently, and the embedding geometry reflects the specific cultural framework of its training corpus.

This distinction is not a limitation of Geometric Context Modeling. It is a diagnostic capability. A system built on this paradigm can tell you which semantic dimensions it can navigate cleanly (physically grounded) and which require supplementary cultural knowledge or multi-corpus grounding (culturally contingent). That is more honest and more useful than a system that provides confident answers without knowing what it knows and what it does not.

4.4 Routes and contextual decomposition

Words with multiple distinct meanings (polysemous words) cannot be faithfully represented as single fixed vectors. The word light has different geometric positions depending on whether it refers to luminosity, weight, or a display indicator. The word cold has different meanings depending on whether it refers to temperature or emotional distance. A Geometric Context Model handles this through contextual decomposition: each meaning is a distinct geometric position (a route) in the axis space, and the current context state routes each ambiguous word to the appropriate position.

Route discovery, introduced in this research series, is the process of asking the embedding directly how many routes each word has. By clustering a word's nearest neighbors, the geometry reveals how many semantically distinct regions that word inhabits. This is geometry revealing its own structure, not human intuition imposing it.

5. KAIA: The First Geometric Context Model

KAIA (Knowledge Architecture for Intelligent Agents) is the first implemented system in the Geometric Context Model class. It was developed through a series of 27 experiments over three months of independent research. The system is described fully in the accompanying paper series (Bare, 2026a, 2026b, 2026c, 2026d). This section summarizes the architecture and validated results.

5.1 Architecture

KAIA has three components. The geometric space is a 50-dimensional GloVe word embedding with 13 confirmed semantic axis directions built from cluster centroids. The axes span physical

dimensions (temperature, speed, luminosity, dominance, social exchange), ICA-discovered dimensions (energy, truth, valence, change), and partially validated dimensions (directness, causality, naturalness). The context state is a 52-byte vector, one float32 per axis, updated at each token using exponential decay. The routing layer resolves polysemous words to their correct geometric position based on the current context state.

5.2 Resource profile

Property	KAIA	Transformer LLM (7B)
Context memory	52 bytes (fixed)	14 GB (grows with sequence length)
Token throughput	44,000 to 97,000 per second	200 to 800 per second
Hardware requirement	Standard CPU, no GPU	GPU required
Memory retrieval	0.065 microseconds at 5,000 items	Linear search or vector DB
Architecture transparency	Fully inspectable geometric state	Implicit in billions of parameters

Table 2. KAIA resource profile compared to a 7B parameter transformer language model. All KAIA measurements from Experiment 6 and Experiment 19 of the research series.

5.3 Benchmark results

KAIA was evaluated against a seven-benchmark suite designed for agent-oriented semantic reasoning. The suite was developed as part of this research series after Experiment 18 reframing established that next-word prediction benchmarks measure the wrong capability for a context model.

Benchmark	Score	Method	Experiments
B1 Intent classification	70%	13D axis encoding + linear classifier (80 examples)	19
B2 Context relevance ranking	80%	Cosine similarity on de-meaned encoding, no training	19
B3 Semantic similarity	80%	Tier separation, no training	19
B4 Antonym detection (concrete)	67% top-1	Pole-swap on axis space, no training	19 to 26
B5 Analogy completion	65% top-3	Vector arithmetic, no training	19
B6 Memory retrieval	70%, 897k q/s	Dimensional nearest-neighbor, no training	19
B7 Agent routing	85%	13D axis encoding + linear classifier (80 examples)	19

B8b Semantic midpoint	100%	Geometric convexity, no training	19 to 26
-----------------------	------	----------------------------------	----------

Table 3. KAIA benchmark results. B2, B3, B5, B6, and B8b use no trained components. Results are stable across multiple runs and axis designs. B8b held at 100% across 15 consecutive experiments.

5.4 Key empirical findings

The research series produced several findings relevant to the Geometric Context Modeling paradigm beyond the KAIA benchmark results.

- **Geometric convexity is a structural property.** The semantic space is inhabited at every point between any two poles, regardless of which words define those poles, which corpus trained the embedding, or which axis design is used. The midpoint between the love-cluster and hate-cluster vectors lies on semantically coherent intermediate words (feelings, dislikes, ambivalence). This was held at 100 percent across 15 consecutive experiments.
- **Physical universals encode cleanly; culturally contingent constructs do not.** Temperature, speed, and luminosity opposition is clean in any corpus because physical reality forces the same distributional patterns everywhere. Moral, epistemic, and social opposition are entangled because these dimensions reflect culturally specific frameworks.
- **ICA reveals the actual independent structure of embedding space.** Independent Component Analysis on WordNet antonym difference vectors discovers directions that are 42 percent more orthogonal than human-designed axes. The discovered directions include specificity, causality, directness, and change, which are more geometrically independent than valence, moral, and truth. Human conceptual categories are not the natural geometry of distributional embeddings.
- **Route discovery lets geometry define its own polysemy structure.** By clustering a word's nearest neighbors in the embedding, the algorithm identifies how many semantically distinct routes each word has without any human judgment. Light has three routes (luminosity, weight, display), cold has two (temperature, emotional distance), and freedom has one in the Wikipedia corpus (political only). The corpus limitation is itself information about what the system knows.

6. What Geometric Context Modeling Is Not

Placing this paradigm clearly in the landscape requires stating what it is not.

Geometric Context Modeling is not a replacement for language generation and does not compete with it. It handles the representation and reasoning about context, which is a separate capability from generating language. A complete conversational system uses both: the context model maintains geometric understanding of where the conversation is, and a generation layer produces responses from that understanding. KAIA routes and positions; a generation component speaks. The two work together, and because they are architecturally separate, each can be optimized independently. A generation layer of any kind can sit above a Geometric Context Model without changing how the context model operates.

It is not a claim that geometry captures everything about meaning. The physically grounded dimensions work cleanly. The culturally contingent dimensions require supplementary grounding. The context model is honest about this distinction in a way that statistical models cannot be.

It is not a rejection of statistical learning. The embedding (GloVe) that defines the geometric space was trained on a large corpus using statistical learning. The linear classifiers used for intent classification and agent routing are trained statistically on small example sets. Statistics produce the map. Geometry is how the map is read and navigated.

It is not a complete theory of meaning. It is a computationally tractable, empirically validated approach to representing and navigating meaning for the specific purpose of agent-oriented semantic reasoning. The benchmark suite defines what that purpose requires. The results show it is achievable on commodity hardware.

7. Open Problems

Geometric Context Modeling is a new field. The open problems are substantial.

- **Culturally contingent axis grounding.** The abstract axis ceiling (zero top-1 accuracy on moral, epistemic, and existential opposition) is confirmed across six experiments and three embedding approaches. Resolving it requires either purpose-built embeddings trained with explicit opposition objectives, multi-corpus grounding that combines the cultural frameworks of multiple linguistic traditions, or a compositional approach that decomposes abstract concepts into combinations of physically grounded dimensions.
- **Decay rate derivation.** The 13 SSM decay rates are currently set by intuition. No experiment has measured the actual temporal dynamics of semantic context in real conversation. Deriving empirically validated decay rates from annotated conversational data is an important open problem.
- **Route library construction.** Route discovery (v3) identifies how many routes each word has in each corpus. Building the route library (the mapping from route identity to geometric position and context signature) currently requires curated examples. A principled approach to constructing the library at scale, from examples rather than retraining, would make the paradigm more practically deployable.
- **Multi-corpus integration.** Wikipedia and Gigaword produce one geometry. Fiction, conversational text, and domain-specific corpora produce different geometric representations of the same words. A framework for combining route libraries from multiple corpora, handling the cases where routes exist in one corpus but not another, is an open architectural question.
- **Evaluation suite generalization.** The seven-benchmark suite was designed for agent-oriented reasoning tasks. Geometric Context Modeling may be applicable to other tasks (clinical semantic analysis, educational dialogue, creative generation) that require different benchmarks. Extending the evaluation framework to new domains is future work.

8. Conclusion

Geometric Context Modeling is a new paradigm for artificial intelligence built on the claim that the right representation for conversational context is not a probability distribution over tokens or a set of discrete symbolic facts, but a continuous geometric position in a structured semantic space maintained as a trajectory.

This paradigm produces systems with properties that neither statistical language models nor symbolic knowledge systems can achieve simultaneously: explicit, inspectable context state; fixed memory regardless of conversation length; semantic reasoning through geometric

operations rather than statistical inference; and the ability to distinguish confidently between what the geometry encodes cleanly and what requires supplementary grounding.

KAIA is the first implementation. Its results (70 to 85 percent on agent-oriented benchmarks, 52 bytes of context memory, 44,000 to 97,000 tokens per second on CPU, 100 percent geometric convexity across 15 experiments) demonstrate that the paradigm is computationally viable and produces useful capabilities.

The computational primitive of language modeling is the token probability distribution. The computational primitive of symbolic AI is the discrete fact. The computational primitive of Geometric Context Modeling is the semantic position in a structured space. These are not competing implementations of the same idea. They are three different answers to the question: what is the basic unit of meaning for an artificial system? The answer proposed here is position. Specifically, position in a space with geometry that reflects the structure of human semantic experience, maintained as a trajectory that records where the conversation has been. That is what KAIA implements. That is what this paper proposes as a field.

The path is clear. The work continues.

References

- Bare, T. (2026a). KAIA: A Geometric Architecture for Semantic Reasoning in Resource-Constrained Environments. KAIA Research Series, Paper 1.
- Bare, T. (2026b). The Geometry of Semantic Opposition: Why Distributional Embeddings Resist Abstract Antonym Detection. KAIA Research Series, Paper 2.
- Bare, T. (2026c). Data-Driven Semantic Axes: ICA Reveals That Independent Opposition Dimensions in Word Embeddings Do Not Match Human Conceptual Categories. KAIA Research Series, Paper 3.
- Bare, T. (2026d). Physical Universality vs Cultural Contingency: Why Geometric Semantic Architectures Work for Physical Dimensions and Fail for Human Conceptual Categories. KAIA Research Series, Paper 4.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL.
- Gardenfors, P. (2000). Conceptual Spaces: The Geometry of Thought. MIT Press.
- Gu, A. and Dao, T. (2024). Mamba: Linear-Time Sequence Modeling with Selective State Spaces. ICML.
- Lakoff, G. and Johnson, M. (1980). Metaphors We Live By. University of Chicago Press.
- Pennington, J., Socher, R., and Manning, C. (2014). GloVe: Global Vectors for Word Representation. EMNLP.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep Contextualized Word Representations. NAACL.