

KAIA Research Series

An Introduction to Geometric Context Modeling and the Five Papers That Establish It as a Field

Tiffney (Smith) Bare • Independent AI Researcher • May 2026
www.linkedin.com/in/tiffneybare

Introduction

This document introduces the KAIA Research Series and provides an overview of the five papers that establish Geometric Context Modeling as a new field of artificial intelligence. It is intended for readers encountering this work for the first time, whether from research, industry, policy, or education backgrounds.

The series was produced entirely through independent research conducted between April and May 2026, across 27 experiments, without institutional support, a research budget, or GPU hardware. It was conducted on standard consumer CPU hardware, which is itself a demonstration of the core claim: serious AI research does not require the resources that most people do not have.

What is Geometric Context Modeling?

Two paradigms currently dominate artificial intelligence for language and meaning. The first is statistical language modeling, which represents meaning as probability distributions over token sequences. The second is 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. Neither can tell you, at any moment, exactly where that conversation is located semantically, with what intensity, across which dimensions, and with what recent history.

Geometric Context Modeling is the third foundation. It is the practice of representing, maintaining, and reasoning about conversational context as an explicit trajectory through a structured geometric semantic space. Meaning is encoded as position. Context is maintained as a trajectory. Reasoning happens through geometric operations. The system runs on any CPU. It requires no GPU. It costs nothing to operate.

KAIA (Knowledge Architecture for Intelligent Agents) is the first implemented system in this class.

Why this research matters

The GPU hardware barrier does not sort AI researchers by aptitude. It sorts them by economic circumstance. A student without access to a GPU cluster or cloud budget cannot meaningfully experiment with current AI systems. A researcher in a region with

limited infrastructure cannot deploy or study them. Geometric Context Modeling was designed from the ground up to change that structural condition.

Three independent bodies of research arrived at the same geometric conclusions from entirely different directions. A 2024 paper by Peter Gardenfors established that meaning has a geometric structure that does not require massive statistical exposure to learn. A January 2026 neuroscience study found that the human hippocampus organizes word meaning along stable geometric axes, a structure that KAIA builds computationally. A 2025 LLM interpretability study found that large language models accidentally internalize geometric structure as a side effect of statistical training at scale. These were not sources for this research. They are independent corroboration reached after the experiments were complete.

How the papers fit together

The five papers form a single research arc. Papers 1 through 4 establish the architecture, investigate its geometry systematically, identify a precise ceiling, and explain what that ceiling reveals about the structure of meaning itself. Paper 5 steps back from the experimental record and proposes Geometric Context Modeling as a field, with KAIA as its first implementation.

Each paper stands on its own. Together, they constitute the complete scientific record of the field's establishment.

Paper 1

KAIA: A Geometric Architecture for Semantic Reasoning in Resource-Constrained Environments

Experiments 1 through 19 • Covers the full architecture, benchmark suite, and validated results

What this paper does

Paper 1 presents the KAIA architecture in full and validates it against a seven-benchmark suite designed specifically for agent-oriented semantic reasoning. It introduces the critical reframing that defines the entire research series: transformer benchmarks measuring next-word prediction accuracy are the wrong evaluation framework for a semantic reasoning system.

The paper documents why training a linear prediction layer on the KAIA state produced 0% next-word accuracy across five epochs, and why that result is a clarification rather than a failure. KAIA is not a text generator. It is a semantic reasoning engine. The benchmark was measuring the wrong thing. This realization led to the definition of seven correct benchmarks for this class of system.

Key findings

- Intent classification: 70% accuracy using 13-dimensional axis encoding and a linear classifier trained on 80 examples
- Context relevance ranking: 80% using cosine similarity on de-meaned 50-dimensional encoding, no training required
- Semantic similarity scoring: 80% tier separation across HIGH, MEDIUM, and LOW categories, no training required
- Analogy completion: 40% top-1, 65% top-3 via vector arithmetic, no training required
- Memory retrieval: 70% accuracy at 897,000 queries per second on a single CPU core, no training required
- Agent routing: 85% accuracy using the same linear classifier approach as intent classification
- Semantic midpoint detection: 100% across every run, every embedding, and every axis set tested

Why it matters

Paper 1 establishes that a geometric semantic reasoning engine can achieve meaningful performance on the tasks agents actually perform, using CPU hardware, with fixed memory that never grows, and at speeds that exceed GPU-backed transformers for the tasks the system is designed for. It also establishes the correct evaluation framework for this class of system, which all subsequent papers build on.

Paper 2

The Geometry of Semantic Opposition: Why Distributional Embeddings Resist Abstract Antonym Detection

Experiments 19 through 23 • Systematic investigation of the architecture's primary limitation

What this paper does

Paper 2 investigates a consistent finding from the benchmark suite: antonym detection works reliably for five concrete physical axes (temperature, speed, dominance, light, and social exchange) but fails for eight abstract axes (valence, moral, truth, existence, temporal, freedom, arousal, and size). The paper runs five systematic experiments to identify why, testing embedding dimensionality, contrastive training, axis alignment, and E8 lattice structure as possible causes.

The conclusion is definitive: the failure is not caused by insufficient dimensions, incorrect axis design, or inadequate training. It is caused by the distributional training objective itself. Abstract antonyms appear in nearly identical linguistic contexts, so co-occurrence-based training pushes them together in embedding space rather than apart.

Key findings

The Gram matrix of the 13 KAIA semantic axes reveals that moral and truth are 74.4% similar in GloVe space, nearly the same direction. Valence, moral, truth, and social form a dense evaluative cluster. This structure is not E8-symmetric, with all residuals falling between 24 and 40 degrees. Contrastive fine-tuning collapses all negative poles into a single region rather than separating them. Moving from 50-dimensional to 300-dimensional embeddings does not change the abstract axis result. The ceiling is a property of co-occurrence-based training objectives, not of dimensionality or axis design.

The positive result is equally important: geometric convexity between semantic poles holds at 100% across all experiments. The space between any two poles is inhabited at every point with semantically coherent intermediate words. This property holds regardless of embedding dimensionality, axis definition, or corpus.

Why it matters

Paper 2 locates the boundary of what the current architecture can and cannot do with precision. That precision is itself a contribution. Statistical AI systems cannot state clearly what they do not know. This paper can.

Paper 3

Data-Driven Semantic Axes: ICA Reveals That Independent Opposition Dimensions in Word Embeddings Do Not Match Human Conceptual Categories

Experiments 23 and 24 • What the geometry actually is versus what we assumed it would be

What this paper does

Paper 3 applies Independent Component Analysis to WordNet antonym difference vectors to discover the actual independent opposition dimensions in GloVe embedding space, rather than imposing human-designed semantic categories. The question it asks is: if we let the geometry define its own axis structure, what does it find?

The answer is that the geometry finds different axes than humans would design. The most geometrically independent opposition dimensions in GloVe space are specificity, causality, directness, change, and naturalness, not valence, moral, and truth. The ICA-discovered axis set achieves 42% better Gram matrix orthogonality than the human-designed axes and eliminates the pathological entanglement between moral and truth found in Paper 2.

Key findings

ICA on 190 WordNet antonym difference vectors discovers axes that are more geometrically independent than those derived from human semantic intuition. The discovered directions include specificity (specific versus general), causality (causal versus acausal), directness (explicit versus implicit), change (dynamic versus static), and naturalness (natural versus artificial). These are more geometrically orthogonal in distributional embedding space than valence, moral, and truth because they correspond more closely to the actual statistical structure of how language is used.

Crucially, the architecture-level semantic benchmarks (context ranking, similarity scoring, analogy completion, and convexity) remain completely stable regardless of which axis design is used. This confirms that these benchmark results are properties of the embedding geometry itself, not of the specific axis labeling chosen.

Why it matters

Paper 3 demonstrates that the geometry of meaning in distributional embeddings has its own structure that does not necessarily align with human conceptual categories. This is a fundamental finding about the nature of meaning in language data, with implications for any AI system that attempts to reason about semantics using distributional embeddings.

Paper 4

Physical Universality vs Cultural Contingency: Why Geometric Semantic Architectures Work for Physical Dimensions and Fail for Human Conceptual Categories

Experiments 19 through 26 • The theoretical synthesis that explains the concrete and abstract split

What this paper does

Paper 4 provides the theoretical synthesis of the experimental findings from Papers 1 through 3. It argues that the concrete versus abstract split in antonym detection performance is not a limitation of the KAIA architecture or of distributional embeddings. It is a fundamental distinction between two classes of semantic opposition: physically forced universals and culturally contingent constructs.

Physical opposites are forced by reality. Hot and cold cannot co-occur in the same object at the same scale simultaneously. This is true in every human language, every culture, every corpus of text. The distributional geometry of these concepts is clean because the underlying physics is clean. Abstract opposites are constructed by human conceptual frameworks. Good and bad, true and false, free and trapped are not properties of physical reality but categories that humans impose on experience. These categories vary across languages, cultures, and philosophical traditions.

Key findings

The paper documents the systematic entanglement of abstract semantic axes across six experiments and three embedding approaches. A cosine similarity of 0.744 between the moral axis and the truth axis is not a measurement error or an artifact of axis design. It is a faithful representation of how these concepts are actually used in human language: in nearly identical contexts, by nearly identical speakers, making nearly identical evaluative judgments.

The path forward requires grounding that goes beyond any existing text corpus. Resolving abstract axis entanglement requires either purpose-built embeddings trained with explicit opposition objectives, multi-corpus grounding that combines linguistic traditions from multiple cultures, or a compositional approach that decomposes abstract concepts into combinations of physically grounded dimensions. The paper identifies these three paths and explains why each represents a genuine contribution to the field if pursued.

Why it matters

Paper 4 reframes what looked like a limitation as a discovery. The architecture is not failing to represent moral opposites. It accurately represents the fact that moral opposites, as encoded in human language data, are not geometrically opposed. Understanding why that is true, and what it would take to change it, is a contribution to both AI research and cognitive science.

Paper 5

Geometric Context Modeling: Introducing a New Paradigm for Semantic Representation and a New Class of Artificial Intelligence Systems

Introductory position paper • Proposes Geometric Context Modeling as a field and KAIA as its first implementation

What this paper does

Paper 5 steps back from the experimental record and proposes Geometric Context Modeling as a new field of inquiry and a new class of artificial intelligence systems. It surveys the existing work that approaches the same idea without fully committing to it, identifies the precise gap between that work and what Geometric Context Modeling proposes, and positions KAIA as the first implemented system in this class.

The paper makes the philosophical case for the paradigm alongside the empirical case. 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?

Key contributions

Paper 5 introduces the formal definition of Geometric Context Modeling as the practice of representing, maintaining, and reasoning about conversational context as an explicit trajectory through a structured geometric semantic space. It defines the five properties that distinguish this paradigm from all prior approaches: continuous geometry, persistent real-time trajectory, explicit semantic axes, geometric reasoning operations, and CPU-native fixed memory. No prior paradigm satisfies all five simultaneously.

The paper also documents three independent convergences from cognitive science, neuroscience, and LLM interpretability that arrived at the same geometric conclusions from entirely different directions, none of which were known during the research that produced the experimental findings. That convergence is the strongest external validation the paradigm has received.

Why it matters

Paper 5 is the document that introduces this work to the broader research community. It is written for the widest possible audience across disciplines and is intended as a public introduction prior to peer review. If there is one paper in the series to share with someone encountering Geometric Context Modeling for the first time, it is this one.

How to read the series

For readers new to the research: begin with Paper 5, which introduces the paradigm and its significance without requiring familiarity with the experimental details. Then read Paper 1 for the full architecture and benchmark results. Papers 2, 3, and 4 can be read in order or independently depending on interest.

For researchers: the papers are self-contained and cross-referenced. The experimental record is in Papers 1 through 4. The paradigm proposal is in Paper 5. The research framework document and external overview, available separately, provide additional context on the architecture principles and research philosophy.

For funders and policymakers: Paper 5 is the entry point. The external overview document provides a non-technical summary of the findings and their implications for AI access, equity, and environmental impact.

What comes next

The mathematical track, documented in the Mathematical Track Proposal, extends Geometric Context Modeling to mathematical reasoning. Mathematics is the domain where geometric structure is most cleanly forced by logic rather than culture, making it the ideal next domain for validation and extension of the architecture.

Domain expansion to code, visual reasoning, and scientific notation follows from the same architectural foundation. Each domain requires its own axis space and benchmark suite, but shares the same 52-byte context state and computational substrate with the language and mathematics modules.

All research will be published in full regardless of outcome. The open problems are documented.

The path is clear. The work continues.