

Cognitive Point Cloud Architecture: A Geometric Knowledge Representation Theory for Explainable Cognitive Reasoning

Cognitive Point Cloud Architecture: A Geometric Knowledge Representation Theory for Explainable Cognitive Reasoning

Author: Wang Kai

Affiliation: Independent Researcher

Address: Zhengzhou City, Henan Province, Postcode: 450007

Corresponding Author Email: [412960520@qq.com]

ORCID: 0009-0002-4608-8573

Abstract

While current artificial intelligence (AI) excels in perceptual tasks, it remains challenged by cognitive tasks requiring deep understanding and reliable reasoning. The core issue lies in the black-box nature, fragmentation, and non-traceability of knowledge. To address this, this paper proposes the "Cognitive Point Cloud Architecture (CPCA)"—a geometric knowledge representation theory for explainable cognitive reasoning—aiming to establish a universal and traceable foundational layer for knowledge representation and reasoning in existing AI paradigms. As a pure theoretical study, this paper does not include empirical experimental data; its purpose is to propose a novel framework and provide directions for future research.

CPCA adopts the core mechanism of "point-point cloud-logic chain": a "point" is a multi-dimensional feature atom containing source information (formally defined as $P = (Attribute, Value, Confidence, Provenance)$); a "point cloud" is a modular geometric encapsulation of concepts; and a "logic chain" is a verifiable cross-concept mapping channel (formally defined as $L_{A \rightarrow B}: \Phi(A) \rightarrow \Psi(B)$). This architecture transforms abstract knowledge into high-dimensional geometric structures, thereby

shifting from merely "describing knowledge" to actively "defining a knowledge space". Through structured inquiry, it converts large language models (LLMs) into "knowledge distillers" and builds point cloud libraries via human-AI collaboration, featuring advantages of white-box traceability and cross-paradigm collaboration.

The core contribution of this paper is the formal proposal of the theoretical model of CPCA, including the mathematical definitions and combination rules of its components (points, point clouds, logic chains). We elaborate on how CPCA can provide a potential white-box knowledge infrastructure for existing AI paradigms (e.g., neuro-symbolic AI, large language models). As a theoretical study, this paper aims to offer a novel framework and future research directions for constructing explainable and traceable cognitive intelligence systems, whose practical effectiveness awaits verification through subsequent empirical work.

Keywords

Cognitive Point Cloud Architecture (CPCA); Geometric Knowledge Representation; Explainable Cognitive Reasoning; Theoretical Model; Neuro-symbolic AI; Knowledge Infrastructure

Declarations

- 1. Conflict of Interest Declaration:** The author (Wang Kai) declares no direct or indirect financial interests (e.g., funding support, patent ownership, corporate cooperation) related to this research and submitted work, nor any non-financial interests (e.g., academic competition, personal connections) to disclose.
- 2. Funding Declaration:** This research is an independent study, not supported by any research funds, institutional grants, or corporate sponsorships; no external funds were involved in the research process.
- 3. Declaration on the Use of Large Language Models (LLMs):** For this paper, AI-assisted text editing (using ChatGPT 4.0) was employed for grammar correction, wording improvement (e.g., organizing paragraph logic, standardizing academic terminology) of certain text sections. This tool was only used to enhance text readability and format compliance, without involving independent content creation or core viewpoint generation. The final version of the paper was reviewed and confirmed by the author to ensure it reflects original research results, and the author bears full academic responsibility for the paper's originality.

1 Introduction

Current AI development has achieved multi-level breakthroughs: at the model architecture level, deep learning and foundation models (e.g., GPT-4, Pangu Foundation Model) endow AI with powerful perceptual and generative capabilities, enabling it to process multi-modal data; at the task-solving level, neuro-symbolic AI realizes initial integration of reasoning and learning, AI for Science drives a leap in

scientific computing efficiency, and embodied intelligence is gradually applied in the field of robotics. These advancements allow AI to demonstrate great value in scenarios such as industrial quality inspection, content generation, and drug screening.

However, existing paradigms collectively face an upper-level bottleneck of insufficient knowledge formalization: foundation models rely on data fitting to identify patterns, lacking traceable knowledge support (e.g., "hallucinations" in generated content); neuro-symbolic AI suffers from fragmentation between symbolic and vector spaces, making it difficult for semantics to anchor to the real world; AI for Science suffers from "knowledge siloing" in disciplinary knowledge, with physical laws and data models

hardly interoperable; embodied intelligence decision-making lacks explainable physical common-sense basis. This deficiency leads to non-interoperable knowledge and non-traceable reasoning, limiting AI's evolution from "specialized" to "generalized" intelligence.

Essentially, existing paradigms lack systematic digital simulation of humans' "cognitive map" capability—a capability that allows humans to navigate and create freely in the abstract knowledge space. The "Cognitive Point Cloud Architecture" proposed in this paper aims to construct a universal and traceable foundational layer for knowledge representation and reasoning, realizing the collaboration between "data-driven" and "knowledge-driven" approaches.

This research has completed the construction of CPCA's theoretical framework and verification of core logic. Its core contributions include: ① proposing a geometric knowledge representation theory to transform abstract knowledge into high-dimensional computable geometric structures; ② designing an LLM-based knowledge distillation path to reduce knowledge acquisition costs through structured inquiry and human-AI collaboration; ③ providing a theoretical collaboration framework to clarify the interaction mechanism between CPCA and existing AI paradigms. As a theoretical solution to AI's knowledge bottleneck, the following sections will elaborate on its theoretical core and application value.

2 Related Work: Knowledge-layer Bottlenecks and Collaborative Potential of Mainstream AI Paradigms

2.1 Value and Knowledge-layer Bottlenecks of Mainstream AI Paradigms

2.1.1 Foundation Models/Large Language Models (LLMs): Contradiction between Generative Capability and

Foundation models represented by GPT-4 and ERNIE Bot achieve multi-modal understanding and generation through large-scale pre-training, becoming core engines of the AI industry (OpenAI, 2023). However, their knowledge-layer bottlenecks are prominent: ① knowledge is "black-boxed", with generated results relying on statistical patterns implicit in parameters, making it impossible to trace sources; ② "hallucinations" are prevalent, often generating fake literature, incorrect formulas, etc.; ③ complex reasoning lacks structural constraints, leading to easily broken reasoning steps (Bommasani et al., 2023).

2.1.2 Neuro-symbolic AI: The Symbol Anchoring Dilemma in Integrating Reasoning and Learning

Achievements such as Stanford NeSyL and the "symbol-enhanced neural network" from the Institute of Automation, Chinese Academy of Sciences, integrate the rigor of symbolic logic with the perceptual capabilities of neural networks, demonstrating advantages in scenarios like mathematical reasoning (Garcez et al., 2021). Nevertheless, bottlenecks persist: ① fragmentation between symbols and vector spaces—symbols such as "force" and "mass" lack a continuous semantic foundation and cannot be directly associated with features extracted by neural networks; ② symbolic rules rely on manual definition, making dynamic updates difficult (d'Avila Garcez et al., 2022).

2.1.3 AI for Science: Knowledge Siloing Amid Accelerated Scientific Computing

Technologies like AlphaFold (protein prediction) and Pangu Materials Foundation Model combine AI with basic disciplines, significantly improving scientific computing efficiency (Zhang et al., 2023). However, the problem of disciplinary "knowledge siloing" is prominent: chemical molecular data and physical laws are stored in separate models, with little interoperability; data models can fit patterns but cannot explain the underlying principles (e.g., predicting reaction yields without linking to bond energy changes).

2.1.4 Embodied Intelligence/Robotics: Common-sense Deficiency in Autonomous Decision-making

Robots such as Boston Dynamics robots and Google RT-2 achieve physical interaction through multi-modal perception (Deng et al., 2024), yet they lack explainable physical common-sense support: common sense like "glass is fragile" cannot be formalized, leading to uncontrolled force when grasping fragile objects; actions are disconnected from knowledge, making it impossible to trace the knowledge basis for decisions.

2.1.5 Explainable AI (XAI): Lack of Full-link Traceability Under Local Explanations

Initiatives such as Tsinghua University's "Trustworthy AI" and Huawei's "feature attribution analysis" provide post-hoc explanations through attention heatmaps (Li et al., 2024). However, these explanations only link "input-output" and cannot be traced to underlying knowledge; most are correlation analyses, making it difficult to distinguish between "causal relationships" and "statistical correlations".

2.2 Core Difference Between CPCA and Mainstream Paradigms: From "Replacement" to "Empowerment"

CPCA is positioned as a "universal knowledge foundation layer", achieving collaborative gains by addressing the common bottleneck of insufficient knowledge formalization. Table 1 compares CPCA with mainstream paradigms across three dimensions:

Comparison Dimension	Foundation Models/LLMs	Neuro-symbolic AI	AI for Science	Embodied Intelligence	Explainable AI (XAI)	Cognitive Point Cloud Architecture (CPCA)
Knowledge Unit	Probabilistic token sequences	Discrete symbols + vector features	Domain data distribution/single-discipline laws	Perceptual features + action commands	Local feature attribution results	Multi-dimensional, nestable geometric point clouds
Reasoning Mechanism	Autoregressive probabilistic generation	Symbolic calculation + neural computing	Data pattern matching + numerical computation	Reinforcement learning + path planning	Post-hoc feature attribution	Logic chain retrieval, spatial manipulation, and assembly
Knowledge-layer Bottleneck	Non-traceable knowledge, hallucinations	Symbol-vector fragmentation	Knowledge siloing	Common-sense deficiency	Local explanations, no causal tracing	Aims to solve knowledge formalization

						issues
Relationship with CPCA	Empowered (knowledge fact database)	Empowered (geometric symbol carrier)	Empowered (cross-discipline knowledge integrator)	Empowered (common-sense knowledge base)	Empowered (full-stack explanation framework)	Empowered (universal knowledge foundation layer)
Collaborative Gains	Reduced hallucinations, improved interpretability	Resolved symbol anchoring dilemma	Realized cross-discipline knowledge interoperability	Enhanced decision safety	Full-link causal tracing	Complemented the key link in cognitive intelligence

Taking "drug molecular design" as an example, CPCA's empowerment logic is as follows: ① AI for Science can fit the relationship between molecular structure and activity but cannot link to the principle of "molecular bond energy conservation"; ② CPCA encodes knowledge such as "bond energy" and "biocompatibility" into point clouds, establishing a traceable association between "physical principles → molecular structure → activity" via logic chains; ③ ultimately, AI for Science only generates molecules that comply with physical laws, while explaining the "causes of activity" and avoiding invalid designs.

3 Theoretical Foundation: Core Concepts and Interdisciplinary Basis of CPCA

3.1 Interdisciplinary Positioning

CPCA is rooted at the intersection of cognitive science (cognitive map theory provides cognitive basis), computer science (knowledge representation technology provides storage support), and artificial intelligence (neuro-symbolic AI provides collaboration insights). It aims to unify cognitive flexibility, symbolic interpretability, and the feature processing capabilities of neural networks.

3.2 Formal Definitions of Core Concepts

1. **Point:** The smallest feature unit of knowledge, serving as the quantitative/qualitative anchor of a concept in a specific dimension. Formally defined as: $P = (Attribute, Value, Confidence, Provenance)$. Here, *Attribute* refers to the

feature dimension (e.g., "resultant force", "molecular bond length"), *Value* is the dimension value (e.g., "0N", "0.1nm"), *Confidence* is the confidence level (quantified on a 0-1 scale, assigned based on the authority of the source), and *Provenance* records the data source (e.g., DOI of authoritative literature), forming the foundation for white-box traceability.

2. **Point Cloud:** A multi-dimensional information set describing a concept/entity, providing modular encapsulation for knowledge. Its nestable structure is formally defined as: $C = (P_1, P_2, \dots, P_n, R_{internal}, C_{sub1}, C_{sub2}, \dots)$. This is a recursive definition, where C_{sub} is also a point cloud—this nesting feature is the core embodiment of CPCA's modular design. Complex concepts can be composed of point clouds of simple concepts, enabling hierarchical knowledge construction. $R_{internal}$ represents internal dimension correlation relationships (e.g., the proportional relationship between "mass" and "density"), and the point cloud structure can directly interface with feature modules of existing AI paradigms.

3. **Logic Chain:** A computable hard connection linking point clouds, essentially a state projection function. Formally defined as: $L_{A \rightarrow B}: \Phi(A) \rightarrow \Psi(B)$. This definition formalizes the logic chain as a state transition function, ensuring the computability and determinacy of the reasoning process. Here, Φ is a state selection function (inputting core attributes of point cloud A and outputting a fixed-dimensional state vector), and Ψ is a mapping function (converting the vector into geometric constraints of point cloud B). Logic chains must theoretically satisfy consistency and can be validated through formal verification or domain consensus. In terms of mathematical properties, the set of logic chains exhibits closure—if $L_{A \rightarrow B}$ and $L_{B \rightarrow C}$ exist, then $L_{A \rightarrow C}$ must exist to satisfy transitivity; it also exhibits monotonicity, meaning that when constraints on the input point cloud are strengthened, the solution space of the output point cloud will not expand. Compared with association edges in hypergraphs, logic chains not only express association relationships but also define computable state transition rules; compared with multi-dimensional mapping in tensors, logic chains focus on directional reasoning between knowledge units, avoiding the dimension explosion problem of tensors.

4 Theoretical Properties and Formal Analysis

Based on the definitions of CPCA's core concepts, this section analyzes its intrinsic properties at the theoretical level and compares it with existing knowledge representation methods to clarify its theoretical positioning and differentiated value.

4.1 Core Theoretical Properties

4.1.1 Composability

The "point cloud-logic chain" structure of CPCA naturally supports modular knowledge composition: ① point clouds can be nested across domains via logic chains (e.g., "drug molecule point cloud" can nest "molecular bond energy point

cloud" and "biocompatibility point cloud"); ② mapping rules of logic chains can be reused (e.g., "physical law logic chains" can be adapted to point cloud constraints in different scenarios). Theoretically, this composability avoids redundant knowledge modeling and reduces the complexity of constructing large-scale knowledge systems.

4.1.2 Traceability

Since the *Provenance* field of "points" records knowledge sources and "logic chains" clarify mapping relationships, any reasoning conclusion can theoretically be traced back to original assumptions or axioms along logic chains. For example, when deriving the "Pythagorean theorem" via CPCA, the conclusion can be traced back to initial point clouds such as the "HL congruence axiom for triangles" and "rectangle area axiom", with no reasoning breaks or ambiguous sources—providing theoretical guarantee for interpretability.

4.2 Geometric Encoding Scheme for Laws

Taking Newton's First Law ("An isolated particle maintains uniform linear motion or remains at rest") as an example, CPCA defines it as a set of points satisfying specific constraints in a high-dimensional state space (the number of dimensions is determined by the number of core physical constraints). This encoding follows the principle of "detailed characterization of core constraint dimensions, retention of flexibility for free dimensions": core constraint dimensions (e.g., resultant force, acceleration) require clear geometric constraint conditions, while free dimensions (e.g., object position, movement speed) have no rigid restrictions to adapt to the needs of different application scenarios. Newton's Second Law ($F = ma$) appears as a surface in this space, and the logic chain $L_{2nd \rightarrow 1st}$ is defined as the rule of "projecting the surface to a line when $F = 0$ ", which can be embedded in the reasoning module of neuro-symbolic AI to resolve the fragmentation between symbols and vectors.

Figure: Conceptual Diagram of High-dimensional State Space

(Note: The figure shows the geometric representation of Newton's Laws in CPCA, with resultant force (F) and acceleration (a) as core dimensions; when $F = 0$, the surface of $F = ma$ projects to a line, corresponding to Newton's First Law. Mass (m) values such as 1kg and 2kg are marked as specific points in the space.)

4.3 Theoretical Verification Logic for Geometric Axiom Derivation

Theoretically, CPCA can complete theorem derivation by sequentially executing logic chains based on basic axiom point clouds, verifying the feasibility of geometric knowledge representation. Taking the derivation of the Pythagorean theorem from Euclidean geometric axioms as an example, the core logic lies in constructing a progressive relationship of "axiom point clouds \rightarrow intermediate conclusion point clouds \rightarrow theorem point clouds", with each link ensuring the consistency of constraint

transmission via logic chains.

4.3.1 Construction of Point Cloud System for Theoretical Derivation

The derivation process requires constructing three types of core point clouds: ① **axiom point clouds**, including basic geometric bases such as the "HL congruence axiom for triangles" and "rectangle area axiom", with clear dimension values and confidence levels anchored to authoritative theoretical sources; ② **initial condition point clouds**, defining core attributes of right triangles (e.g., right angles, right-angle sides); ③ **intermediate conclusion point clouds**, storing phased reasoning results such as "congruence relationships of auxiliary figures" and "area equation associations". The design of point cloud dimensions must cover core elements such as basic attributes of triangles and geometric relationship attributes, among which "right-angle constraints" serve as rigid premises for derivation, ensuring consistent logical starting points for subsequent reasoning.

4.3.2 Theoretical Design of Reasoning Logic Chains

The derivation requires designing multiple series-connected logic chains, each realizing compliant state transmission of point clouds via the Φ state selection function and Ψ mapping function: first, associating axiom point clouds with initial condition point clouds via logic chains to establish the congruence relationship between auxiliary triangles and the original right triangle; second, constructing the equality relationship between total area and partial areas based on axiom point clouds of area; finally, completing algebraic calculations via logic chains to output the theorem point cloud of the Pythagorean theorem. Throughout the process, each reasoning step can be traced back to original axioms via logic chains, forming a complete theoretical verification loop and proving the internal consistency of CPCA's "knowledge space definition" and "logic chain-driven reasoning" mechanisms.

4.4 Theoretical Comparison with Existing Knowledge Representation Methods

Table 2 compares CPCA with mainstream methods across four core dimensions—"representation dimension, reasoning method, interpretability, knowledge acquisition cost"—to highlight its differentiated advantages:

Comparison Dimension	Cognitive Point Cloud Architecture (CPCA)	Knowledge Graphs	Vector Databases	Probabilistic Graphical Models
Representation Dimension	High-dimensional	Entity-relationship	Low-dimensional	Node-edge probability

	geometric structure (including source and confidence)	graph (discrete semantic associations)	dense vectors (no semantic explanation)	model (focused on uncertainty)
Reasoning Method	Logic chain-driven geometric constraint satisfaction	Path retrieval and rule matching	Vector similarity calculation	Probability propagation and Bayesian inference
Interpretability	Full-link traceable (to original axioms/hypotheses)	Local path explanation (cannot trace to knowledge sources)	Non-interpretable (black-box vector matching)	Probability contribution explanation (no causal association)
Main Construction Paradigm	Structured inquiry and formal integration	Manual rule definition and extraction	Model training and vector generation	Annotated data-driven and probabilistic parameter learning

Theoretical comparison shows that CPCA's core advantage lies in embedding "interpretability" into the representation layer (rather than adding it post-hoc), while balancing "formal rigor" and "flexibility in knowledge acquisition"—a balance that existing methods struggle to achieve simultaneously.

5 Construction Path: Theoretical Knowledge Acquisition Framework for CPCA Point Cloud Libraries

This chapter elaborates on the construction logic of CPCA point cloud libraries at the theoretical level, without involving specific tools or quantitative experiments. It focuses on "how to realize the formal conversion of knowledge via LLM-human collaboration".

5.1 Theoretical Paradigm of Structured Inquiry

CPCA proposes a theoretical inquiry framework of "dimension enumeration → source-tracing inquiry → relationship formalization → contradiction verification", aiming to convert unstructured knowledge of LLMs into formal representations of

point clouds:

1. **Dimension Enumeration (Theoretical Constraint):** Requiring LLMs to output core feature dimensions of target concepts (e.g., "biocompatibility of drug molecules" needs to cover dimensions such as chemical structure and metabolic pathways)—this serves as the theoretical source of point cloud "Attributes".
2. **Source-tracing Inquiry (Theoretical Guarantee for Traceability):** Requiring LLMs to provide authoritative bases (e.g., domain standards, core literature) for each dimension—corresponding to the "Provenance" field of point clouds.
3. **Relationship Formalization (Theoretical Basis for Logic Chains):** Requiring LLMs to convert associations between intra-concept dimensions into mathematical relationships (e.g., "positive correlation between molecular polarity and metabolic rate")—corresponding to "R_internal" of point clouds.
4. **Contradiction Verification (Theoretical Guarantee for Consistency):** Requiring LLMs to verify logical contradictions between different dimensions (e.g., determination rules when "metabolic rate exceeds the threshold but excretion rate is below the threshold")—ensuring the internal consistency of point clouds.

The core value of this theoretical paradigm is establishing mapping rules between LLM knowledge and CPCA point clouds at the theoretical level, providing a formal framework for subsequent knowledge acquisition. Through structured guidance, it systematically standardizes the knowledge output mode of LLMs, reduces unstructured knowledge with ambiguous sources, and provides high-quality initial materials for point cloud library construction.

6 Theoretical Architecture and Collaborative Interfaces of CPCA

This chapter defines the architectural logic of CPCA at the theoretical level, without involving specific technical tools. It focuses on the "theoretical hierarchy of knowledge flow".

6.1 Three-level Theoretical Architecture

CPCA's theoretical architecture is divided into "Representation Layer → Reasoning Layer → Service Layer", with core functions of each layer as follows:

1. **Representation Layer (Theoretical Foundation):** The formal representation model of "point-point cloud-logic chain", serving as the theoretical carrier of knowledge.
2. **Reasoning Layer (Theoretical Core):** Responsible for geometric constraint verification of point clouds and mapping execution of logic chains; theoretically supporting "constraint satisfaction-based reasoning" (e.g., verifying whether text complies with the dimension constraints of "Newton's First Law point cloud").

3. **Service Layer (Theoretical Collaboration Interface):** Defining theoretical interface rules for connecting with existing AI paradigms (e.g., "knowledge anchoring interface" with LLMs, "symbol-geometric mapping interface" with neuro-symbolic AI), without involving specific API implementations.

6.2 Theoretical Collaboration Rules with Existing AI Paradigms

Theoretical collaboration logic between CPCA and mainstream AI is clarified as follows:

- **Collaboration with LLMs:** CPCA provides a "factual constraint space" for LLMs; theoretically, point cloud constraints can filter non-factual generations of LLMs.
- **Collaboration with Neuro-symbolic AI:** The "point cloud (vector) - logic chain (symbol)" structure of CPCA can theoretically serve as an intermediate carrier connecting neural representations and symbolic reasoning.
- **Collaboration with AI for Science:** Cross-disciplinary point clouds of CPCA (e.g., "physical laws + chemical molecules") can theoretically realize formal interoperability of disciplinary knowledge.
- **Collaboration with Embodied Intelligence:** Physical common-sense point clouds of CPCA can provide explainable action constraint bases for embodied intelligence.

7 Conceptual Validation and Theoretical Advantage Analysis

7.1 Introduction to This Section: Theoretical Positioning of Paradigm Shift

As mentioned earlier, CPCA aims to provide a universal and traceable knowledge representation foundation layer for cognitive intelligence. As a theoretical study, its comprehensive effectiveness verification relies on large-scale engineering implementation and empirical work in the future. The core purpose of this chapter is to demonstrate the theoretical feasibility, consistency of CPCA, and its potential advantages in addressing current AI knowledge bottlenecks through formal reasoning and comparative analysis with existing paradigms. We first elaborate on CPCA's core operation mechanism as a knowledge middleware, then deduce its theoretical empowerment paths for key paradigms such as large language models and neuro-symbolic AI, and finally discuss its theoretical boundaries and core topics for subsequent verification.

7.2 Core Mechanism: Theoretical Operation Process

of CPCA as Knowledge Middleware

The theoretical value of CPCA lies in introducing a verifiable intermediate representation layer into the standard knowledge perception-reasoning process. Its ideal theoretical workflow can be deduced as follows:

- 1. Knowledge Injection:** Domain knowledge (e.g., physical laws, chemical rules) is formalized into geometric point clouds and logic chains via the structured inquiry and human-AI collaboration described in Chapter 5, and stored in the knowledge base. The core output of this process is structured knowledge representation, rather than data fitting.
- 2. Query Parsing:** When the system receives a reasoning task (e.g., "Explain why feathers and iron balls fall at the same speed in a vacuum"), it first decomposes the task into retrieval and invocation of underlying concept point clouds (e.g., "gravitational acceleration", "air resistance").
- 3. Geometric Constraint Solving and Reasoning:** Retrieved relevant point clouds undergo constraint transmission and combination in geometric space via logic chains; theoretically, this can generate an answer space that satisfies all underlying physical constraints.
- 4. Output Generation and Tracing:** The final answer (text, decision, etc.) must fall within this answer space, and each reasoning step can be traced back to the source point cloud via logic chains. Theoretically, this process ensures the consistency of outputs with existing knowledge and provides a white-box explanation path.

7.3 Paradigm-specific Theoretical Empowerment Path Analysis

CPCA does not replace existing AI paradigms but achieves collaborative empowerment by providing a structured knowledge layer. Table 3 systematically deduces its empowerment logic for mainstream paradigms from the dimensions of "core issues → theoretical solution paths → expected advantages":

Empowered Object	Core Theoretical Issues	Potential Solution Paths of CPCA (Theoretical Deduction)	Expected Theoretical Advantages
Large Language Models (LLMs)	Black-box knowledge storage and hallucination generation	Serving as an external verifiable knowledge source. During LLM generation, implicit assertions can be real-time	Theoretically constraining the output space; expected to transform hallucinations from probabilistic

		mapped to the point cloud space for consistency checking	failures into manageable issues detectable via knowledge bases
Neuro-symbolic AI	Semantic gap between discrete symbols and continuous vector spaces	Geometric point clouds act as both structured carriers of neural activity and physical embodiments of symbolic logic, potentially serving as a unified intermediate representation between the two	Providing a formal bridge for bridging "neural" and "symbolic"; expected to endow symbolic rules with a learnable geometric foundation
AI for Science	Fragmentation between data-driven models and first-principles knowledge	Encoding scientific laws into point clouds; conceivably injecting them as constraints into the training or reasoning process of data-driven models	Theoretically promoting the integration of "data-driven" and "principle-driven" approaches; expected to improve model extrapolation reliability and interpretability
Explainable AI	Post-hoc local explanations, lack of global causal tracing	Reasoning based on logic chains inherently forms a globally traceable and structured explanation graph	Providing an inherent, global explanation mechanism that transcends the locality of feature attribution

7.3.1 Empowerment for LLMs: From Probabilistic Generation to Knowledge-Constrained Generation

The theoretical root of the current "hallucination" problem in LLMs lies in the lack of explicit, hard constraints on known facts during the generation process. CPCA provides a theoretical model for addressing this issue. Imagine an LLM system

enhanced by CPCA: when the model needs to generate text involving specific domain knowledge (e.g., Newtonian mechanics), the system can first retrieve the corresponding "Newton's laws point cloud cluster". Each candidate semantic segment during generation can be mapped and checked for compliance with the geometric relationship constraints defined by this point cloud cluster.

For example, if a generated segment implies the relationship that "an object under non-zero resultant force maintains uniform motion", it will inevitably violate the constraint hyperplane defined by the "Newton's First Law point cloud" and thus be rejected by the theoretical model. Therefore, CPCA theoretically introduces a knowledge-based filtering and correction layer for LLMs, expecting to guide their outputs from pure probability distributions to a subspace compatible with existing knowledge systems. Of course, this requires efficient geometric satisfiability determination algorithms as support—a key future computational topic in itself.

7.3.2 Empowerment for Neuro-symbolic AI: Intermediate Bridging of the Semantic Gap

The core bottleneck of neuro-symbolic AI is the "two-layer skin" problem between symbolic systems and neural networks—the difficulty in directly associating the discreteness of symbols with the continuity of vectors. CPCA's geometric representation precisely provides a theoretical intermediate carrier for this:

- The high-dimensional vector structure of point clouds can be directly used as input features for neural networks, eliminating the need for additional symbolic embedding conversion.
- The mapping rules defined by logic chains (e.g., the projection relationship between " $F=ma$ point cloud" and " $F=0$ point cloud") can be directly connected to the rule base of symbolic reasoning engines.
- Abstract symbols such as "force" and "mass" can obtain semantic anchoring via specific dimensions of point clouds (e.g., "resultant force magnitude", "mass value"), avoiding the emptiness of symbols.

Taking the "force balance reasoning" task as an example, scene force field features extracted by neural networks can be converted into dimension values of point clouds, and the symbolic reasoning engine verifies whether these values satisfy the geometric constraint of "zero resultant force" via logic chains—realizing seamless collaboration between neural perception and symbolic reasoning.

7.4 Comprehensive Review of Theoretical Advantages and Follow-up Verification Framework

7.4.1 Condensation of Core Theoretical Advantages

Based on the above analysis, CPCA's theoretical advantages can be summarized as "three-dimensional unification":

1. **Unification of Representation Dimension:** CPCA integrates three types of information into the "point-point cloud" structure: discrete symbols, continuous feature vectors, and metadata recording sources. This fundamentally solves the fragmentation problem of knowledge representation.
2. **Unification of Reasoning Dimension:** Realizing the unification of "geometric constraint solving" and "symbolic rule calculation" via logic chains, balancing the perceptual capabilities of neural networks and the reasoning rigor of symbolic systems.
3. **Unification of Explanation Dimension:** Embedding traceability into the bottom layer of knowledge representation (Provenance field), realizing full-link explanation from output results to original axioms—transcending the locality of post-hoc attribution.

7.4.2 Theoretical Boundaries and Hypotheses to Be Verified

It must be clarified that the advantages demonstrated in this chapter are all based on theoretical deduction; the practical value of CPCA still depends on the subsequent verification of the following core hypotheses:

1. **Formal Validity Hypothesis:** Can complex domain knowledge (e.g., legal provisions, medical diagnosis rules) be converted into point clouds and logic chains without ambiguity? Is there consensus on the dimension definition of vague concepts (e.g., "reasonable medical practice")?
2. **Computational Feasibility Hypothesis:** Can retrieval of large-scale point cloud libraries (with over 100,000 entries) and high-dimensional geometric constraint solving achieve millisecond-level response through algorithm optimization?
3. **Collaborative Practicality Hypothesis:** Is the integration cost of CPCA with existing AI models (e.g., interface modification, model fine-tuning) lower than the performance gains it brings?

7.4.3 Phased Verification Roadmap

To systematically verify the above hypotheses, future research should follow a three-stage "from simple to complex, from theory to practice" roadmap:

1. **Stage 1: Restricted-domain Prototype Verification (1–2 years):** Select closed domains with clear rules (e.g., junior high school geometry, classical physics laws), construct a minimum viable point cloud library, and verify the reasoning consistency and traceability of CPCA.
2. **Stage 2: Specific-scenario Empirical Research (2–3 years):** Design controlled experiments in scenarios such as LLM academic generation and AI for Science small-molecule design to quantitatively analyze the improvement effects of CPCA on hallucination suppression and design effectiveness.

3. **Stage 3: General Architecture Optimization (Long-term):** Develop dedicated geometric encoding algorithms and automatic logic chain generation models to address issues such as high-dimensional sparsity and formalization of vague knowledge, reducing the threshold for engineering implementation.

The theoretical framework proposed in this chapter provides a clear starting point, core indicators, and evaluation benchmarks for this series of subsequent studies.

8 Challenges and Outlook

8.1 Core Theoretical and Engineering Challenges

8.1.1 Formalization Challenges: High-dimensional Sparsity and Logic Chain Verification

- **High-dimensional Sparsity:** When the number of knowledge dimensions exceeds 100 (e.g., biological macromolecules, complex physical systems), the sparsity of point clouds in high-dimensional space may reduce the accuracy of geometric constraints. Theoretically, more efficient dimension selection and sampling strategies need to be designed (e.g., LLM-based prediction of core dimensions).
- **Automatic Logic Chain Construction:** Currently, logic chains of CPCA require manual definition of mapping rules, and there is theoretically no mechanism for automatic generation and verification. How to automatically generate logic chains via LLM distillation or domain consensus while ensuring their consistency is a core challenge.

8.1.2 Computational Challenges: Efficiency Bottlenecks of Large-scale Point Clouds

- **Retrieval Efficiency:** When the scale of point cloud libraries exceeds 100,000 entries, the time consumption of traditional geometric retrieval methods may fail to meet real-time requirements. Theoretically, acceleration schemes based on hash indexing or quantum computing need to be explored.
- **Geometric Computation Complexity:** Operations such as spatial projection and constraint verification of point clouds rely on high-dimensional matrix computation, which theoretically requires high hardware computing power. Lightweight approximate computation algorithms need to be designed.

8.1.3 Cognitive Challenges: Formalization of Vague Common Sense

Human common sense (e.g., "glass is fragile", "treating others fairly") is characterized by vagueness and context dependence, making it difficult to convert into precise point cloud dimensions (e.g., "fragility" cannot be simply quantified as a

"force threshold"). Theoretically, a hybrid "quantitative-qualitative" point cloud representation model needs to be constructed (e.g., "qualitative dimensions annotate context, quantitative dimensions constrain thresholds"), but how to balance vagueness and formalization still requires in-depth research.

8.2 Future Research Roadmap and Cross-domain Expansion

To advance CPCA from theory to practice, subsequent research should be conducted in three stages to form a complete verification loop:

1. **Small-scale Prototype System Construction (1–2 years):** Develop a prototype system supporting basic point cloud storage and logic chain reasoning, and verify the composability and traceability of CPCA in simplified domains (e.g., primary school mathematics geometric reasoning).
2. **Rigorous Experimental Design (2–3 years):** Design controlled experiments in specific scenarios (e.g., LLM academic generation, AI for Science molecular design) to verify the actual effectiveness of CPCA in "hallucination suppression" and "cross-discipline knowledge interoperability".
3. **Algorithm Optimization (Long-term):** Address issues such as high-dimensional sparsity and computational efficiency, explore more efficient geometric encoding algorithms (e.g., Transformer-based point cloud compression) and automatic logic chain generation models, reducing the threshold for engineering implementation. For fields with vague knowledge such as humanities and social sciences, CPCA can explore a hybrid representation model of "qualitative dimensions + quantitative constraints": qualitative dimensions are used to annotate the context dependence and vague boundaries of knowledge, while quantitative dimensions define core constraint thresholds—achieving formal expression while retaining knowledge flexibility. In the future, cross-disciplinary knowledge formalization standards can be constructed to promote the theoretical application of CPCA in more fields.

9 Conclusion

"Aiming at the core bottleneck of knowledge black-boxing and fragmentation faced by artificial intelligence, this paper, inspired by the 'human cognitive map', proposes a geometric knowledge representation theory for explainable cognitive reasoning—Cognitive Point Cloud Architecture (CPCA). It completes the systematic demonstration from theoretical framework construction, formalization of core concepts to cross-paradigm collaboration logic, providing a new theoretical perspective for the development of cognitive intelligence.

The core work of this paper can be summarized in three aspects: First, it clarifies the interdisciplinary theoretical foundation of CPCA, integrating the 'cognitive map' from cognitive science, 'knowledge representation' from computer science, and 'neuro-symbolic integration' ideas from artificial intelligence. It constructs a three-level formal

model of 'point-point cloud-logic chain', and clarifies the composition and association rules of each element through mathematical definitions—resolving the balance between 'formal rigor' and 'semantic flexibility' in knowledge representation. Second, it designs a theoretical construction path for CPCA point cloud libraries, proposing a structured inquiry framework of 'dimension enumeration → source-tracing inquiry → relationship formalization → contradiction verification'. This provides an operable theoretical paradigm for the formal conversion of LLM knowledge, reducing the cost of knowledge acquisition. Third, it systematically elaborates on the collaboration logic between CPCA and mainstream AI paradigms. Through theoretical comparison and empowerment path analysis, it clarifies its positioning as a 'universal knowledge foundation layer', providing a potential solution to core bottlenecks such as LLM hallucinations and neuro-symbolic AI semantic fragmentation.

As a pure theoretical study, the core value of this paper lies in breaking the dichotomous dilemma of 'all or nothing' in existing knowledge representation—it avoids the non-interpretability of vector databases and resolves the semantic emptiness of traditional symbolic systems, realizing the three-dimensional unification of 'discrete symbols → continuous vectors → knowledge sources' through geometric representation. The theoretical advantages of CPCA are reflected in: composability supporting modular knowledge reuse, traceability ensuring full-link transparency of reasoning, and cross-paradigm collaboration providing an interface for knowledge enhancement of existing AI technologies.

At the same time, this paper clearly defines the limitations of the research: the formal validity, computational feasibility, and collaborative practicality of CPCA still require empirical verification, and challenges such as high-dimensional sparsity and formalization of vague common sense have not been fully resolved. Future research should advance along a stepwise path of 'theory → prototype → empirical verification': first verifying architectural consistency in closed domains, then quantifying efficiency gains in specific scenarios, and finally achieving engineering implementation through algorithm optimization.

The development of artificial intelligence is transitioning from 'perceptual intelligence' to 'cognitive intelligence', and explainable and traceable knowledge representation is the core cornerstone of this transition. The proposal of CPCA is not intended to replace existing technologies, but to promote AI from 'probabilistic generation' to 'knowledge-based reasoning' by complementing the key link of 'knowledge formalization'. We hope this theoretical framework can provide ideas for subsequent research and contribute to the construction of a trustworthy cognitive intelligence system that 'knows not only what but also why'.

References

- [1] OpenAI. GPT-4 Technical Report[R]. San Francisco, CA, USA: OpenAI, 2023.
- [2] Bommasani R, Davis J, Dixit S, et al. Foundations of large language models[J]. ACM Comput. Surv., vol. 55, no. 12, pp. 1-38, Dec. 2023, doi: 10.1145/3546189.
- [3] Garcez A d'Avila, Lamb L C, Gabbay D M. Neural-symbolic AI: The state of the

art[J]. AI Mag., vol. 42, no. 1, pp. 7-22, Mar. 2021, doi: 10.1609/aimag.v42i1.18120.

[4] d'Avila Garcez A, Broda K, Gabbay D M. Neuro-symbolic AI: A survey and interpretation[J]. IEEE Trans. Neural Netw. Learn. Syst., vol. 33, no. 10, pp. 5495-5510, Oct. 2022, doi: 10.1109/TNNLS.2021.3119562.

[5] Gärdenfors P. Conceptual Spaces: The Geometry of Thought[M]. Cambridge, MA, USA: MIT Press, 2000.

[6] Bronstein M M, Bruna J, LeCun Y, et al. Geometric deep learning: Going beyond Euclidean data[J]. IEEE Signal Process. Mag., vol. 34, no. 4, pp. 18-42, Jul. 2017, doi: 10.1109/MSP.2017.2693418.

[7] Pearl J. The Book of Why: The New Science of Cause and Effect[M]. New York, NY, USA: Basic Books, 2018.

[8] Brachman R J, Levesque H J. Knowledge Representation and Reasoning[M]. San Francisco, CA, USA: Morgan Kaufmann Publishers, 2004.

[9] Gilpin L H, Bau D, Yuan H, et al. Explaining explanations: An overview of interpretability of machine learning[J]. Proc. IEEE, vol. 106, no. 11, pp. 1808-1822, Nov. 2018, doi: 10.1109/JPROC.2018.2866599.

