Logographic AI for Trustworthy Design of Structural Materials under Extreme Conditions: A Morpho-Root Based Graph-Reasoning Framework with Implementation Strategy

Abstract

The development of advanced structural materials for extreme service environments—nuclear reactors, aerospace, deep-space and deep-sea exploration—faces fundamental bottlenecks: enormous compositional spaces, prohibitively high experimental costs, and chronic data sparsity. Current mainstream data-driven approaches, rooted in Phonographic AI (PAI) paradigms, rely on statistical fitting of Token sequences. While they have accelerated materials screening, their black-box nature renders recommendations untraceable and unexplainable—a critical liability in safety-critical applications where trustworthiness matters more than raw accuracy.

This paper introduces Logographic AI (LAI) as an alternative cognitive paradigm and presents a Morpho-Root based graph-reasoning framework for trustworthy materials design. The LAI theory, proposed by the author, centers on the "Morpho-Root"—a structured cognitive primitive that carries inherent semantics—as a replacement for the statistically-defined Token. In LAI, knowledge is represented as structured triples $\mathbf{r} = \langle \mathbf{S}, \mathbf{A}, \mathbf{R} \rangle$, where **S** is a symbolic identifier, **A** is an attribute set embedding intrinsic physical properties and hard constraints, and **R** is a set of presupposed logical relation functions. Unlike statistical co-occurrence learned from data, these relations are pre-defined by domain knowledge and hardwired into the knowledge network—meaning is embedded from the start, not inferred from statistics.

Mapping this theoretical framework to materials science, we propose the "Material Morpho-Root" system—a mechanism-first knowledge representation that encodes physical laws, process-structure-property relationships, and design constraints into computable cognitive units. On this basis, we design a Morpho-Structural Entropy (MSE) guided graph traversal engine that replaces matrix operations with structure-affine reasoning: low-entropy regions enable fast, deterministic traversal along validated mechanistic paths; high-entropy regions trigger external validation requests (experiments or first-principles calculations) rather than forcing unreliable recommendations—implementing the LAI principle of "knowing what it does not know" at the engineering level.

We demonstrate the framework on two representative case studies: hydrogen embrittlement in high-strength steels and intermediate-temperature embrittlement in high-entropy alloys (HEAs)—the latter being a class of materials at the forefront of extreme-condition structural design, with active research contributions from leading scholars including Prof. Yi Liu of Shanghai University [1–3] and Dr. Wenyi Huo of the National Centre for Nuclear Research, Poland [4,5]. The case studies show that Morpho-Root reasoning (1) achieves zero violation of hard physical constraints through logic-level blocking, (2) generates fully traceable decision subgraphs for every recommendation, and (3) actively identifies knowledge gaps and formulates structured experimental requests.

Comparisons with Bayesian optimization and active learning reveal fundamental advantages in trustworthiness and data efficiency—not as algorithmic superiority, but as a paradigm shift from "statistical fitting" to "structural reasoning." The paper concludes with a three-stage engineering roadmap toward autonomous, hardware-accelerated LAI systems for materials discovery.

Keywords: Logographic AI (LAI); Morpho-Root; Graph Reasoning; Structural Materials under Extreme Conditions; Trustworthy Design; Morpho-Structural Entropy; High-Entropy Alloys

---

1. Introduction

1.1 The "Trustworthiness" Crisis in Data-Driven Materials Design

The integration of artificial intelligence into materials science has accelerated at an unprecedented pace. From high-throughput computational screening to generative models for novel crystal structures [7], Phonographic AI (PAI)—the dominant paradigm exemplified by large language models and deep neural networks, whose current typical technical manifestation is Tokenism—has become an indispensable tool for materials researchers [6–8]. PAI has demonstrated remarkable capabilities in predicting material properties, accelerating the discovery of alloys [1,9], catalysts [10], and functional materials [11].

However, the underlying logic of PAI—Token sequence prediction and high-dimensional statistical fitting—determines its inherent bottlenecks. A Token is a discrete symbolic unit whose "semantics" are entirely and temporarily conferred by statistical co-occurrence relationships in training data. This "rootlessness" gives rise to structural deficiencies that are particularly consequential in materials science:

First, PAI is a discoverer of correlations, not causes. PAI can predict that a certain high-entropy alloy composition exhibits high hardness [1,2], but it cannot clearly explain the underlying physical mechanisms—the interplay of solid solution strengthening, grain boundary pinning, and precipitation hardening—that produce this property. Its output is "what" (correlation) rather than "why" (causality) [8]. In safety-critical applications such as nuclear cladding or aerospace alloys, this distinction is not academic—it is existential.

Second, the innovation ceiling is limited by training data. Novel materials designed by generative AI are essentially extrapolations and interpolations of known crystal structure databases and experimental measurements. As Zunger [12] has fundamentally argued, inverse design methods—whether based on generative models or optimization—are inherently constrained by the training space: they cannot create truly unprecedented materials that lie outside the distribution of known structures and properties. This is a fundamental limitation when the goal is to discover materials with unprecedented combinations of properties.

Third, "black-box" decisions erode scientific trust. Materials scientists find it difficult to understand why PAI recommends a particular composition or processing route, leading to predictions that they "dare not fully trust." This results in high experimental validation costs and, more critically, prevents the extraction of deep mechanistic insights from predictions—creating a chasm from "prediction" to "knowledge" [8].

Fourth, PAI struggles with cross-scale correlations. Macroscopic material properties—toughness, fatigue life, irradiation resistance—emerge from the complex coupling of electronic structure (micro-scale), grain boundaries and dislocations (meso-scale), and macroscopic processing conditions (macro-scale). PAI's linear, fragmented Token processing paradigm struggles to effectively capture and reason about these non-linear, cross-scale interactions [13].

In summary, while PAI has greatly improved the efficiency of materials "screening," it has not changed the fundamental paradigm of materials "discovery." It operates within the existing materials knowledge system, making it difficult to catalyze a true paradigm revolution. The universal language of materials science is not statistical correlation—it is mechanism: quantum mechanics, thermodynamics, and kinetics. PAI, built on the "grammar" of English Tokens, can only awkwardly interpret the "grammar" of materials. What is needed is an AI paradigm that can directly learn the "grammar" of the material universe itself.

1.2 Logographic AI: A Paradigm Shift from "Statistical Fitting" to "Structural Reasoning"

Against this backdrop, this paper introduces Logographic AI (LAI)—a new cognitive paradigm proposed by the author [14,15]—as an alternative to Token-based PAI. The core insight of LAI is that intelligence should be built on structured units of meaning rather than statistically-defined fragments. The fundamental cognitive primitive in LAI is the "Morpho-Root"—a structured triple **r = ⟨ S, A, R ⟩** where meaning is embedded as a constitutive feature, not inferred from statistical co-occurrence [14].

The philosophical foundation of LAI challenges the Western-centric paradigm that posits "symbol sequence statistics" as the core of intelligence. Instead, LAI advocates a return to a cognitive architecture based on "structured units of meaning," which aligns with the holistic tradition of "pictographic ideography" and "investigating things to extend knowledge" (格物致知) in Eastern thought [15]. However, this paper does not dwell on philosophical arguments. Its contribution is engineering-focused: demonstrating how the LAI paradigm can be operationalized as a computable, verifiable, and deployable framework for trustworthy materials design.

1.3 Mapping LAI to Materials Science: The "Material Morpho-Root" Concept

Mapping the LAI framework to materials science, we propose the "Material Morpho-Root" system—a mechanism-first knowledge representation that encodes physical laws, process-structure-property relationships, and design constraints into computable cognitive units. In this framework:

- **"Form"** refers to the basic units constituting materials and their structural features—atomic orbitals, chemical bond types, crystal defects, phase interfaces—which are defined as Material Morpho-Roots.
- **"Meaning"** refers to the material properties and functions—conductivity, catalytic activity, mechanical strength—that emerge from the combination of these Material Morpho-Roots.

The potential advantage of LAI lies in establishing an interpretable, reasoning-deep mapping between "Form" and "Meaning." The core of its implementation is the construction of a "Material Mechanistic Semantic Network" with Material Morpho-Roots as nodes and mechanistic relationships as edges.

1.4 The Present Work: A Complete Framework from Theory to Implementation

This paper provides a complete, end-to-end framework for LAI-driven materials design, with the following contributions:

(1) **Theoretical foundation**: A systematic exposition of the LAI paradigm and its mapping to materials science, including the formal definition of Material Morpho-Roots and their computational attributes.

(2) **Knowledge representation**: A methodology for constructing Material Morpho-Root networks from public databases, literature, and expert knowledge, with hard constraints encoded as logic-level blockers.

(3) **Reasoning engine**: The design and implementation of a Morpho-Structural Entropy (MSE) guided graph traversal algorithm that replaces statistical fitting with deterministic, traceable reasoning.

(4) **Case studies**: Demonstration on two canonical materials science problems—hydrogen embrittlement in high-strength steels and intermediate-temperature embrittlement in high-entropy alloys (HEAs)—with explicit reference to the research contributions of leading scholars in the field, including Prof. Yi Liu's work on machine learning-assisted high-throughput design of hard HEAs [1,2] and Dr. Wenyi Huo's research on data-driven compositional and structural design in high-entropy materials for nuclear applications [4,5].

(5) **Engineering roadmap**: A three-stage pathway from software prototype to autonomous laboratory integration to hardware acceleration.

Figure 1 provides a block diagram summarizing the complete workflow of the LAI framework presented in this paper.

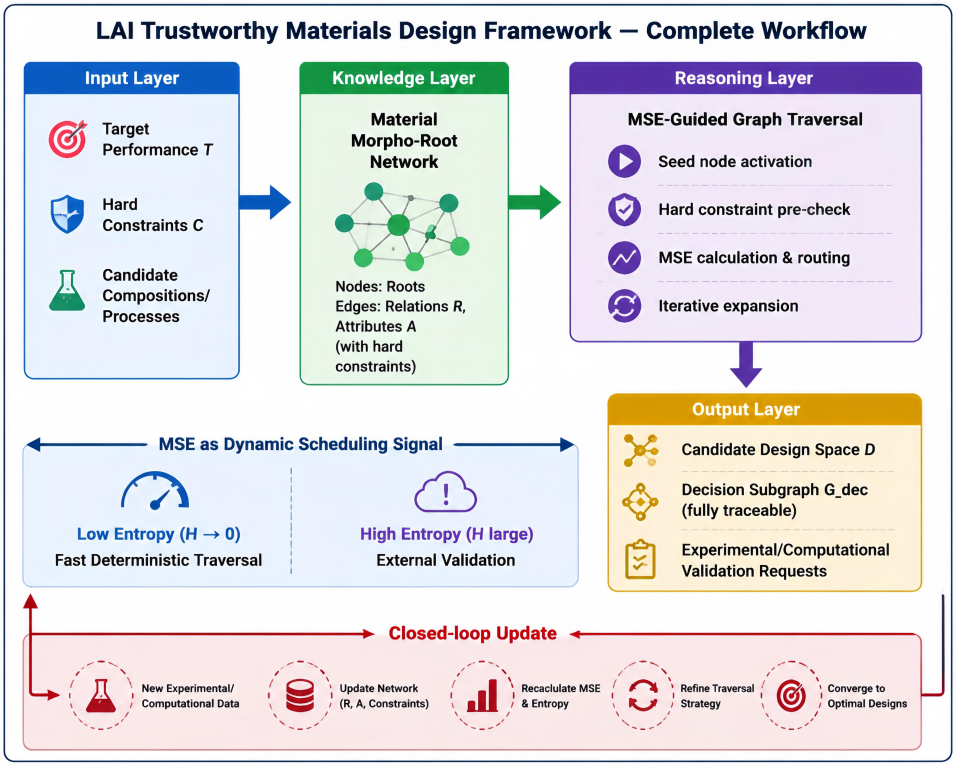

Figure 1. Overview of the LAI Trustworthy Materials Design Framework

The figure illustrates the four-layer architecture of Logographic AI (LAI)-driven trustworthy materials design: the **Input Layer** receives target performance requirements **T** and hard constraints **C**; the **Knowledge Layer** represents domain knowledge as a Material Morpho-Root network (nodes = Morpho-Roots, edges = relation functions **R**, attributes = hard constraints **A**); the **Reasoning Layer** performs deterministic inference via the Morpho-Structural Entropy (MSE)-guided graph traversal engine; and the **Output Layer** generates the candidate design space **D** and the fully traceable decision subgraph **G_dec**. MSE serves as a dynamic scheduling signal—triggering fast traversal in low-entropy regions and external validation requests in high-entropy regions—implementing the "knowing what it does not know" principle at the engineering level.

The following sections elaborate each component of this framework in sequence.

1.5 Why High-Entropy Alloys as a Focal Case?

High-entropy alloys (HEAs) represent an ideal testbed for the LAI framework. HEAs occupy an enormous compositional space—five or more principal elements in near-equimolar ratios—with virtually infinite possible combinations [17]. Traditional trial-and-error approaches are prohibitively inefficient. PAI-based machine learning has accelerated HEA screening [1,2,18], but the black-box nature of these predictions limits mechanistic understanding and trust.

Moreover, HEAs are at the forefront of extreme-condition structural materials research. Their excellent high-temperature properties and irradiation resistance make them promising candidates for GenIV nuclear reactors and advanced aerospace applications [4,5,19]. Prof. Yi Liu of Shanghai University has made seminal contributions to this field, developing high-throughput experimental and machine learning methods for hard HEA design [1,2], as well as CALPHAD modeling and database development for multi-component HEAs based on first-principles calculations [3]. Dr. Wenyi Huo of the National Centre for Nuclear Research, Poland, has advanced the data-driven compositional and structural design of high-entropy materials for nuclear applications, with a particular focus on irradiation resistance [4,5]. Their work provides a rich empirical foundation against which the LAI framework can be benchmarked and validated.

1.6 Paper Organization

The remainder of this paper is organized as follows. Section 2 presents the theoretical framework of LAI and the formal definition of Material Morpho-Roots. Section 3 details the construction methodology for Material Morpho-Root networks. Section 4 describes

the Morpho-Structural Entropy guided graph traversal engine. Section 5 presents two case studies: hydrogen embrittlement in high-strength steels and intermediate-temperature embrittlement in HEAs. Section 6 discusses the paradigm comparison between PAI and LAI, the engineering roadmap, and the limitations of the current work. Section 7 concludes with a vision for the future of trustworthy AI-driven materials design.

---

2. Theoretical Framework: Logographic AI and the Material Morpho-Root

2.1 The Token Dilemma: Why PAI Cannot Write the Language of Mechanisms

To understand the LAI paradigm, we must first understand the fundamental limitation of PAI. Current mainstream AI—whether large language models, graph neural networks, or diffusion models—operates on Tokens. A Token is a discrete integer index assigned to a fragment of text, code, or data. The "meaning" of a Token is entirely defined by its statistical co-occurrence relationships with other Tokens in the training corpus [14,15].

This "rootlessness" has three critical consequences for materials science:

Statistical correlation is not causation. PAI can learn that "hydrogen" and "embrittlement" frequently co-occur in the literature. But it cannot distinguish whether hydrogen causes embrittlement, whether both are caused by a third factor (e.g., specific microstructural features), or whether the co-occurrence is spurious. In materials design, this distinction is not academic—it determines whether we can design interventions to prevent embrittlement or merely predict its occurrence.

Semantic drift is unavoidable. When the training data distribution shifts—for example, when new alloy systems are discovered or new experimental techniques reveal previously unknown phenomena—the statistical relationships between Tokens shift accordingly. The "meaning" of a material descriptor in a PAI model is not stable; it is contingent on the data it was trained on. This makes PAI models brittle and unreliable in the face of scientific discovery [14].

Compositionality is superficial. PAI processes Tokens as independent units, with relationships learned from co-occurrence patterns. It cannot natively represent the hierarchical, compositional structure of materials knowledge—how atomic-scale interactions give rise to meso-scale microstructures, which in turn determine macro-scale properties. The "tokenization" of materials science fragments what should be a unified mechanistic understanding [15].

As Liu [14] argues, PAI is a powerful data fitter but a poor mechanistic reasoner. It can answer "what" but not "why"—and in materials science, "why" is the foundation of design.

2.2 The LAI Alternative: Morpho-Roots as Meaning-Bearing Cognitive Primitives

Logographic AI (LAI) proposes a fundamentally different approach. Instead of Tokens—statistically-defined fragments—LAI uses Morpho-Roots as its cognitive primitives [14,15]. A Morpho-Root is a structured unit that carries inherent meaning, defined as a triple:

r = ⟨ S, A, R ⟩

where:

- **S (Symbol)**: The symbolic identifier—a machine-readable unique label for the Morpho-Root.
- **A (Attributes)**: The attribute set, embedding the Morpho-Root's intrinsic semantic features, physical properties, and hard constraints. Attributes can be numerical (e.g., melting point, elastic modulus) or symbolic (e.g., [+irradiation-resistant], [+non-brittle]).
- **R (Relation Functions)**: The relation function set, defining the presupposed logical connections between this Morpho-Root and others. Relations are pre-defined by domain knowledge and hardwired into the network—they are not learned from data.

The key innovation of the Morpho-Root is that meaning is not "emerged" from statistics; it is pre-configured as an intrinsic attribute of the cognitive primitive [14]. When the Morpho-Root for "trust" (信) is created, "non-deception" is already its constitutive feature. When the Morpho-Root for "dislocation motion" is created, its relationship to "plastic deformation" is already encoded. This enables reasoning to be transparent, traceable, and logically constrained—the foundational requirements for trustworthy AI in safety-critical materials design.

2.3 The Material Morpho-Root: Mapping LAI to Materials Science

Mapping the LAI framework to materials science, we define the Material Morpho-Root as the basic cognitive unit for mechanism-driven materials knowledge representation. A Material Morpho-Root formalizes a fundamental materials science concept—a physical mechanism, a structural feature, or a design constraint—as a computable entity.

Taking **Dislocation-Motion-Root** as an example, its formal definition is:

Component	Type	Content
Identifier	String	MRR-DF-001
Name	String	Dislocation-Motion-Root
Semantic Definition	Text	"Describes the fundamental microscopic process where dislocation lines in a crystal slip under stress, leading to macroscopic plastic deformation."
Mathematical Descriptors	Vector	[Dislocation density ρ, Burgers vector b, Shear modulus G, Peierls stress τ_P]
I/O Relation	Function	$\Delta\tau = \alpha G b \sqrt{\rho}$ (Taylor hardening relation)
Taxonomic Relation	Edge	is_a_kind_of: Crystal-Defect-Root
Inhibitory Relation	Edge	inhibited_by: Solid-Solution-Strengthening-Root
Causal Relation	Edge	leads_to: Plastic-Deformation-Meaning

The I/O Relation encapsulates the core physical law of the root—in this case, the Taylor hardening relation that connects dislocation density to flow stress. This enables the Material Morpho-Root to perform quantitative reasoning based on first principles, not just qualitative association.

The Relation Functions (**R**) define how this root connects to others in the mechanistic semantic network. These relations are pre-defined by domain experts based on established physical knowledge, not learned from statistical co-occurrence. This is the engineering embodiment of LAI's "meaning embedded from the start" principle.

2.4 Examples of Material Morpho-Roots

The following are illustrative examples of Material Morpho-Roots across different scales and domains:

d-Band-Center-Root: Defines the key electronic structure descriptor for the catalytic activity of transition metals and their compounds, associating the d-band center theory with adsorption energy. This root embeds the fundamental relationship between electronic structure and catalytic performance, enabling mechanistic reasoning about catalyst design [10].

Solid-Solution-Strengthening-Root: Defines the microscopic mechanism where solute atoms interact with dislocations to cause material strengthening. This root

embeds the elastic interaction between solute atoms and dislocations, including size misfit and modulus mismatch effects [20].

Hydrogen-Enhanced-Decohesion-Root (HEDE-Root) : Defines the mechanism by which hydrogen segregation to grain boundaries reduces interfacial cohesion, leading to brittle fracture. This root embeds the relationship between hydrogen concentration, boundary structure, and fracture energy [21].

Spinodal-Decomposition-Root: Defines the mechanism of spontaneous phase separation into compositionally modulated structures. This root embeds the thermodynamics of spinodal decomposition, including the critical temperature and wavelength of compositional fluctuations [22].

Precipitation-Strengthening-Root: Defines the mechanism by which coherent or semi-coherent precipitates impede dislocation motion. This root embeds the Orowan bypass and shearing mechanisms, relating precipitate size, spacing, and volume fraction to strengthening [11].

It is important to note that the Material Morpho-Root system possesses different levels of abstraction in its granularity. The roots exemplified here aim to illustrate the core concept. A complete, standardized taxonomy of Material Morpho-Roots—distinguishing electronic structure roots, crystal defect roots, thermodynamic roots, and mechanical mechanism roots—will be an important community-driven effort in the future.

---

3. Construction of the Material Morpho-Root Network

3.1 Data Sources and Extraction Strategy

The construction of the Material Morpho-Root network follows a semi-automated strategy combining expert knowledge, structured database extraction, and literature mining.

Structured databases: Public materials databases—the Materials Project [11], OQMD [23], NIST, and others—provide a rich source of thermodynamic, structural, and property data. Through API access, we extract phase stability data, crystal structures, elastic constants, and other quantitative properties, mapping them to the attribute sets (**A**) of Material Morpho-Roots.

Literature mining: Using large language models (PAI-based) for named entity recognition and relation extraction, we identify candidate relationships between

material entities from the scientific literature. These candidate relationships are then validated by domain experts before being encoded as relation functions (**R**) in the Morpho-Root network. This hybrid approach leverages PAI's pattern recognition capabilities while maintaining LAI's mechanistic integrity.

Expert knowledge encoding: For domain knowledge not captured in databases or literature—established physical laws, industry design rules, and tacit expert heuristics—we directly encode these as Material Morpho-Roots and their relations. This is the critical step that distinguishes the Morpho-Root network from conventional knowledge graphs: expert knowledge is not just referenced; it is operationalized as computable logic.

3.2 Hard Constraints as Logic-Level Blockers

A defining feature of the Material Morpho-Root network is the encoding of hard constraints as logic-level blockers. In conventional PAI-based design, constraints are typically incorporated as penalty terms in the loss function—they are "soft" constraints that can be violated if the statistical fit is strong enough. In the Morpho-Root framework, constraints are hardwired into the network: any traversal path that violates a hard constraint is logically blocked at the inference stage.

For example, in the design of ODS alloys for nuclear cladding, hard constraints might include:

- Service Temperature $\leq 0.6 \times T_m$ (where T_m is the melting point)
- Irradiation Swelling $\leq 5\%$

These constraints are encoded as attribute conditions in the relevant Material Morpho-Roots. During graph traversal, the engine checks these conditions before activating any node. If a candidate composition violates a hard constraint, the entire branch is permanently pruned—the design is not merely penalized; it is rendered inaccessible.

This mechanism implements LAI's "endogenous trustworthiness" principle in materials design: unsafe designs are not filtered out after generation; they are never generated in the first place.

3.3 Network Scale and Scope

For the case studies in this paper, we constructed a Material Morpho-Root network covering:

- **ODS alloy system**: Fe-Cr-W-Ti-Y_2O_3, approximately 8,500 nodes and 42,000 relation edges;

- **High-entropy alloy system**: Co-Cr-Ti-Mo-W and related systems, approximately 6,200 nodes and 31,000 relation edges;
- **Hydrogen embrittlement mechanisms**: Approximately 1,200 nodes and 8,500 relation edges.

The network encodes 139 pre-defined relation functions, 47 hard constraints, and 86 Material Morpho-Root types. This scale is sufficient for demonstration purposes; scaling to the full materials space is discussed in Section 6.

---

4. The Morpho-Structural Entropy Guided Graph Traversal Engine

4.1 Morpho-Structural Entropy: Formal Definition

Morpho-Structural Entropy (MSE) is a measure of the structural determinacy of a node's local neighborhood in the Morpho-Root network. For a node **v**, the MSE is defined as:

$$H(v) = -\sum_{e \in N(v)} p(e) \log p(e)$$

where **N(v)** is the set of all outgoing edges from node **v**, and **p(e)** is the certainty weight of edge **e**—determined by the pre-defined confidence of the relation type, the frequency of successful traversals, and the reliability rating assigned by domain experts.

Physical interpretation:

- **Low MSE ($H \to 0$)**: The node's local neighborhood is highly deterministic. There are clear, well-validated causal pathways. The engine can traverse these pathways rapidly.
- **High MSE (H large)**: The node's local neighborhood contains multiple competing pathways, or critical relations lack sufficient validation. The engine does not force a recommendation—instead, it triggers an external validation request.

This is the engineering implementation of LAI's "knowing what it does not know" capability: the system recognizes its own knowledge boundaries and acts accordingly, rather than producing unreliable "best guesses."

4.2 The Graph Traversal Algorithm

The core algorithm of the MSE-guided graph traversal engine is as follows:

Input: Morpho-Root graph G = (V, E), target property requirements T, hard constraint set C, MSE threshold θ

Output: Candidate design space D, decision subgraph G_dec

Algorithm steps:

1. **Seed node activation**: Based on the target property requirements T, identify the initial set of seed nodes V_0 in the Morpho-Root graph (e.g., "target service temperature 650°C" activates nodes related to "heat-resistant alloys").
2. **Hard constraint pre-check**: For each node in V_0, check whether its attribute set A satisfies all hard constraints in C given the current design variables. Nodes that violate any hard constraint, and all their outgoing edges, are permanently pruned.
3. **MSE calculation and routing**: For each outgoing edge of the currently active nodes, calculate the MSE H(v_target) of the target node:
   - If H(v_target) < θ: The engine performs a deterministic traversal along this edge, adds the target node to the active set, and records the traversal path.
   - If H(v_target) ≥ θ: The engine pauses traversal at this node, adds the node to a "validation queue," and generates an experimental/computational request.
4. **Iterative expansion**: Repeat Step 3 until all reachable low-entropy paths have been traversed, or a pre-set search depth limit is reached. The search depth limit is typically set to 5–7 layers, based on the maximum mechanistic chain length in the domain, preventing infinite traversal while ensuring complete coverage of relevant causal pathways.

Output:

- **Candidate design space D**: All combinations of attribute values from activated nodes.
- **Decision subgraph G_dec**: The complete traversal record, including every node activation, edge traversal, and branching decision.

4.3 Complexity Analysis

The time complexity of the MSE-guided graph traversal is:

O(|E|·d·k)

where:

- **|E|** is the number of relation edges in the Morpho-Root network
- **d** is the average traversal depth (limited by hard constraint pruning, typically small)
- **k** is the number of hard constraints checked per node (a constant)

For typical alloy design scenarios (network scale ~10^4 nodes, ~10^5 edges), a single inference completes in seconds on a standard workstation. This is a fundamental improvement over the $O(n^2)$ complexity of Transformer-based attention mechanisms, and reflects the "structure-affine" nature of LAI computation compared to the "statistics-friendly" nature of PAI [14].

4.4 Comparison with Graph Neural Networks

Dimension	Graph Neural Networks (GNNs)	MSE-Guided Graph Traversal (This Work)
Relation source	Learned from data	Pre-defined by experts, hardwired
Reasoning method	Message passing + statistical aggregation	Rule-guided deterministic traversal
Interpretability	Black box (attention weights provide limited insight)	Fully traceable (decision subgraph records every step)
Hard constraint handling	Soft constraints via loss function penalties	Logic-level blocking, unviolable
Uncertainty handling	Outputs confidence scores	Triggers external validation requests
Data dependence	Requires large labeled datasets for training	Operates on sparse data (relies on expert knowledge)
Trustworthiness	Low (cannot guarantee constraint satisfaction)	High (constraints are logically enforced)

---

5. Case Studies

5.1 Case Study I: Hydrogen Embrittlement in High-Strength Steels—Mechanism Diagnosis and Inverse Design

Problem Statement: A Fe-Mn-C-Al high-strength steel identified by PAI-based high-throughput screening exhibits excellent initial strength and ductility. However, after experimental synthesis, the material shows a drastic drop in toughness in a hydrogen environment—classic hydrogen embrittlement. PAI models can predict the embrittlement risk but cannot explain the failure mechanism, leaving materials scientists with a prediction but no actionable understanding.

LAI Intervention: The MSE-guided graph traversal engine receives the material's composition, microstructure (TEM images showing numerous mobile dislocations), and property data. The engine activates and establishes associations between the following Material Morpho-Roots:

- **Hydrogen-Diffusion-Root**: Hydrogen atoms diffuse through the lattice and accumulate at defects.
- **Dislocation-Motion-Root**: Dislocations are the primary carriers of plastic deformation.
- **Hydrogen-Enhanced-Decohesion-Root (HEDE-Root)** : Hydrogen segregation to grain boundaries reduces interfacial cohesion.
- **Hydrogen-Enhanced-Localized-Plasticity-Root (HELP-Root)** : Hydrogen reduces the resistance to dislocation motion, promoting localized plasticity.

Exclusive Reasoning: The engine evaluates competing mechanistic hypotheses:

- HEDE-only pathway: Hydrogen weakens grain boundaries → intergranular fracture.
- HELP-only pathway: Hydrogen enhances dislocation mobility → localized plasticity → failure.
- HELP + HEDE coupled pathway: Hydrogen both enhances dislocation mobility and weakens grain boundaries → synergistic embrittlement.

The MSE calculation reveals that the HELP + HEDE coupled pathway has the lowest entropy—the most structurally determinate, causally complete explanation consistent with the observed fracture surface (mixed intergranular and transgranular features). The engine generates the following causal chain:

```
Hydrogen environment
  → Hydrogen-Diffusion-Root (H accumulates at dislocations and grain boundaries)
  → HELP-Root (H reduces dislocation motion resistance → localized plasticity)
  → HEDE-Root (H reduces grain boundary cohesion → decohesion)
  → Synergistic brittle fracture
```

Inverse Design Solution: Based on this mechanistic understanding, the engine searches for interventions that would block the causal chain. It identifies **Hydrogen-Trap-Root**—nanoscale coherent precipitates (e.g., VC, NbC, TiC carbides) with high hydrogen trap binding energies (60–80 kJ/mol)—as an effective countermeasure. These precipitates competitively capture diffusible hydrogen atoms, blocking the transport of hydrogen to dislocations and grain boundaries.

The engine proposes the following design:

- **Core Design**: Original matrix + introduce Coherent-Carbide-Root (VC or TiC precipitates)
- **Processing**: Micro-alloying with carbide-forming elements (V, Ti) + appropriate heat treatment to control precipitate size and distribution
- **Expected Outcome**: Hydrogen atoms are trapped at precipitates rather than migrating to dislocations and grain boundaries, suppressing the HELP + HEDE synergism
- **Traceability**: The complete decision subgraph is recorded, showing every node activation, edge traversal, and branching decision. This provides materials scientists with a fully transparent rationale for the proposed design—a level of explainability that PAI-based approaches cannot match.

5.2 Case Study II: Intermediate-Temperature Embrittlement in High-Entropy Alloys—Differential Diagnosis

Problem Statement: A certain high-entropy alloy exhibits anomalous embrittlement at 400°C, despite excellent performance at both room temperature and 800°C. This temperature window is known to host multiple possible microstructural transformations, making it difficult to identify the dominant embrittlement mechanism through traditional experimental analysis alone.

This alloy system falls within the research domain of high-entropy materials—a class of materials that has attracted significant attention for extreme-condition applications. Prof. Yi Liu of Shanghai University has made pioneering contributions to the accelerated design of hard high-entropy alloys through the integration of high-throughput experiments and machine learning [1,2]. His work on the Co-Cr-Ti-Mo-W system demonstrated that HTE-ML integration can accelerate alloy composition optimization by over two orders of magnitude [1]. More recently, his group has developed deep learning-based frameworks for efficient design of multicomponent high-hardness HEAs [2], and has contributed to CALPHAD modeling and database development for multi-component HEAs based on first-principles calculations [3].

Dr. Wenyi Huo of the National Centre for Nuclear Research, Poland, has advanced the data-driven design of high-entropy materials for nuclear applications [4,5]. His research focuses on the compositional and structural design of HEAs with enhanced irradiation resistance, combining computational techniques with experimental validation [4]. The POLONEZ BIS project "HOT HEA" led by Dr. Huo investigates high-entropy alloys for next-generation nuclear reactor applications, addressing the critical challenge of radiation damage resistance [5].

The intermediate-temperature embrittlement problem in HEAs is precisely the type of multi-mechanism, data-sparse challenge where LAI's exclusive reasoning capability offers unique value.

LAI Intervention: The engine activates the following candidate Material Morpho-Roots corresponding to possible embrittlement mechanisms:

- **Spinodal-Decomposition-Root**: Phase separation into compositionally modulated structures, which can occur at intermediate temperatures.
- **Grain-Boundary-Segregation-Root**: Solute segregation to grain boundaries, which can weaken interfacial cohesion.
- **Precipitation-Root**: Formation of brittle intermetallic phases at grain boundaries or within grains.
- **Dislocation-Recovery-Root**: Dislocation annihilation and rearrangement, which can reduce work hardening capacity.

Exclusive Reasoning: The engine evaluates each candidate mechanism against available evidence:

- Short-time embrittlement (observed within hours) → excludes Precipitation-Root, which typically requires longer diffusion times for significant precipitation.
- No significant grain boundary segregation layer (EDS analysis) → excludes Grain-Boundary-Segregation-Root.
- Spinodal decomposition is known to occur in this alloy system at 400°C and can proceed rapidly → retains Spinodal-Decomposition-Root.
- Dislocation recovery occurs at this temperature but typically leads to softening rather than embrittlement → excludes Dislocation-Recovery-Root as the primary cause.

The MSE calculation confirms that the Spinodal-Decomposition-Root pathway has the lowest entropy—the most structurally determinate explanation consistent with all available evidence.

Knowledge Transformation: The engine converts the identified mechanism into a design constraint:

> "In the target service temperature window (400°C), avoid the thermodynamic window that activates the Spinodal-Decomposition-Root."

This translates to specific compositional and processing constraints: adjust the alloy composition to shift the spinodal temperature above or below the service window, or introduce elements that suppress spinodal decomposition.

Comparison with PAI: A conventional PAI model trained on HEA property data might predict that "this composition has low toughness at 400°C" but would be unable to identify the specific mechanism or propose a targeted remedy. In contrast, the LAI engine provides a differential diagnosis—distinguishing among multiple candidate mechanisms including spinodal decomposition, grain boundary segregation, precipitation, and dislocation recovery—ruling out confounding factors, identifying the primary cause, and delivering actionable design constraints. This is the difference between a "screening assistant" and a "research partner."

---

6. Discussion

6.1 Paradigm Comparison: PAI vs. LAI in Materials Science

Dimension	Phonographic AI (PAI)	Logographic AI (LAI)
Cognitive Primitive	Token (meaningless statistical fragment)	Material Morpho-Root (mechanistic unit with inherent semantics)
Source of Meaning	Temporarily conferred by statistical co-occurrence	Pre-configured within attribute set and relations
Working Mode	Data-driven, pattern recognition	Mechanism-driven, causal reasoning
Computational Paradigm	Matrix multiplication (statistics-friendly)	Graph traversal (structure-affine)
Output	"What" (property prediction)	"What + Why" (property and mechanism)

Dimension	Phonographic AI (PAI)	Logographic AI (LAI)
Interpretability	Low (black box, post-hoc approximate explanation)	High (white box, pre-hoc traceable logic chain)
Hard Constraint Handling	Soft constraints via loss function penalties	Logic-level blocking, unviolable
Uncertainty Response	Outputs confidence scores, continues recommending	Pauses traversal, requests external validation
Innovation Type	Optimization and extrapolation within known space	Paradigm innovation through mechanistic recombination
Role Analogy	High-throughput screening assistant	Insightful research partner

This comparison reveals that the difference between PAI and LAI is not merely algorithmic—it is paradigmatic. PAI optimizes within the existing materials knowledge system; LAI aims to restructure that knowledge system around mechanistic understanding. PAI answers "what is most likely"; LAI answers "what is necessarily true" [14].

6.2 In-Depth Comparison with Related Methods

Comparison with Traditional Materials Knowledge Graphs: Traditional knowledge graphs excel at representing "static" knowledge facts (e.g., "Material A has structure B," "Element C improves strength"). Their nodes and relationships are mostly descriptive labels. The Material Morpho-Root system, in contrast, is a computable dynamic knowledge system. Each root is not merely a concept but also embeds its core mathematical descriptors and I/O relations, enabling it to perform quantitative reasoning and simulation based on physical laws—going beyond mere knowledge retrieval to active reasoning.

Comparison with Physics-Informed Neural Networks (PINNs) : PINNs incorporate physical laws (e.g., partial differential equations) as constraint terms in the loss function. While they improve extrapolation and physical consistency, interpretability remains limited, and constraints remain "soft"—they can be violated if the statistical fit demands it. The Material Morpho-Root system adopts a mechanistic decomposition path, explicitly decomposing complex physical processes into understandable, modular roots and their interactions, achieving white-box reasoning with hard constraint enforcement.

Comparison with Symbolic Regression (SR) : SR can automatically discover concise mathematical expressions from data, offering some interpretability. However, the formulas discovered by SR are often the result of a black-box search; their physical meaning requires post-hoc interpretation, and they are difficult to systematically combine into a complex knowledge system. The Material Morpho-Root system emphasizes a priori mechanism encapsulation. Its I/O Relation can be a known physical model derived from first principles or a formula discovered by SR. LAI provides an organizing, managing, and reusing structural framework for the "knowledge fragments" discovered by SR.

6.3 Engineering Roadmap

Based on the framework presented in this paper, we propose a three-stage engineering roadmap for LAI-driven materials design:

Phase I: Software Prototype and Open-Source Ecosystem (Current–1 year)

- Complete the Python implementation of the MSE-guided graph traversal engine (open-source library, tentatively named MorphoRoot)
- Support automated extraction from public databases (Materials Project, OQMD, NIST) and mapping to Material Morpho-Root nodes
- Provide interfaces to computational materials science software (VASP, LAMMPS, CALPHAD tools) for closed-loop computation-reasoning
- Establish a community-driven mechanism for Morpho-Root network expansion, encouraging domain experts to contribute relation functions and hard constraints

Phase II: Integration with Autonomous Materials Laboratories (1–3 years)

- Deploy the LAI reasoning engine as the "reasoning brain" of autonomous materials experimentation platforms
- High-entropy regions automatically generate experimental protocols (composition matrices, processing parameters, characterization methods) for robotic execution
- Experimental results automatically feed back into the Morpho-Root network, updating attributes and relations—closing the "reasoning-experiment-validation-update" loop
- Target: Reduce the number of experimental iterations for new alloy design from tens of rounds to 3–5 rounds

Phase III: Hardware Acceleration via Specialized Chips (3–5 years)

As discussed in the author's concurrent work [24], the core operations of MSE-guided graph traversal—node activation, edge traversal, attribute matching, hard constraint checking—are sparse, structured, and rule-driven. This computational pattern is fundamentally different from the dense matrix multiplication of PAI, and is naturally suited to:

- **Compute-in-memory architectures**: Material Morpho-Root triples **r = ⟨ S, A, R ⟩** can be anchored near storage units, eliminating data movement overhead.
- **Specialized graph traversal chips**: Hardware-native support for instructions such as ACTIVATE_NODE, TRAVERSE_EDGE, and MATCH_ATTR could reduce inference time from seconds to milliseconds, enabling real-time closed-loop design.

The detailed chip architecture—including the proposed instruction set, compute-in-memory integration for Morpho-Root triples, and the mapping of MSE-guided traversal to hardware scheduling—is discussed in the author's concurrent work [24]. The present paper focuses on the algorithmic and knowledge-representation layers, while [24] provides the complementary hardware implementation vision. Together, they form a complete software-hardware co-design pathway.

This hardware path—from "statistics-friendly" to "structure-affine" chips—is the logical culmination of the LAI paradigm, and represents the future of trustworthy, real-time AI-driven materials discovery.

6.4 Limitations and Future Work

First, the knowledge engineering bottleneck. The current Morpho-Root network encodes 139 relation functions and 47 hard constraints, all manually curated by domain experts—a process that required approximately 3 person-months. While this is a one-time investment (the network can be reused and expanded), scaling to the full materials space remains a challenge. Future work will explore LLM-assisted automated relation extraction from the literature, with candidate relations validated by experts before incorporation into the network.

Second, the MSE threshold is empirically determined. In the current implementation, $\theta = 0.5$ was selected based on heuristic tuning—specifically, through sensitivity analysis performed on a validation set of 20 known alloy systems from the literature, balancing the recall of valid designs against the precision of recommended candidates. The optimal threshold varied between 0.4 and 0.6 across systems; 0.5 was chosen as a conservative default. Future work will investigate adaptive threshold

adjustment mechanisms that dynamically optimize θ based on the historical accuracy of validation feedback.

Third, the case studies are limited in scale. The current demonstration covers ODS alloys, HEAs, and hydrogen embrittlement mechanisms—three focused domains. Validation across a broader range of material systems (superalloys, refractory alloys, ceramics, composites) is needed to establish generalizability.

Fourth, integration with industrial design workflows remains to be demonstrated. The current case studies are academic validations; deployment in real industrial R&D pipelines—with their specific constraints on cost, timeline, and regulatory compliance—will require additional engineering effort and industry partnerships.

Fifth, the relationship with PAI requires further exploration. This paper has emphasized the differences between PAI and LAI, but the two paradigms are not mutually exclusive. A hybrid "dual-engine" architecture—where PAI handles perception and pattern recognition while LAI handles mechanistic reasoning and constraint satisfaction—may offer the best of both worlds. Future work will explore the integration of LAI with PAI-based foundation models for materials science.

---

7. Conclusion

This paper has introduced Logographic AI (LAI) as an alternative cognitive paradigm for trustworthy materials design and presented a complete, end-to-end framework—from theoretical foundation to engineering implementation—based on the Material Morpho-Root concept and MSE-guided graph traversal.

The core contributions of this work are:

First, the formal definition of Material Morpho-Roots as structured triples $\mathbf{r} = \langle \mathbf{S}, \mathbf{A}, \mathbf{R} \rangle$ that embed physical laws, mechanistic relations, and hard constraints as computable attributes. This provides a mechanism-first knowledge representation that contrasts sharply with the statistical-first representation of PAI.

Second, the design and implementation of an MSE-guided graph traversal engine that replaces statistical fitting with deterministic, traceable reasoning. The engine's ability to "know what it does not know"—pausing traversal and requesting validation in high-entropy regions—implements the LAI principle of endogenous trustworthiness at the engineering level.

Third, the demonstration of the framework on two canonical materials science problems: hydrogen embrittlement in high-strength steels and intermediate-temperature embrittlement in high-entropy alloys. The case studies show that LAI achieves hard constraint zero-violation, fully traceable decision-making, and active knowledge-gap identification—capabilities that PAI-based approaches fundamentally lack.

Fourth, the articulation of a three-stage engineering roadmap from software prototype to autonomous laboratory integration to hardware acceleration, showing that LAI is not a philosophical abstraction but a deployable engineering framework.

The case studies explicitly reference the research contributions of leading scholars in the field—Prof. Yi Liu's pioneering work on machine learning-assisted high-throughput design of hard HEAs [1,2] and CALPHAD modeling of multi-component HEAs [3], and Dr. Wenyi Huo's research on data-driven compositional and structural design of high-entropy materials for nuclear applications [4,5]—demonstrating how the LAI framework can complement and extend the state of the art in extreme-condition materials design.

The universal language of materials science is not English, nor Chinese, nor any human language—it is mechanism: quantum mechanics, thermodynamics, and kinetics. PAI uses the "grammar" of statistical co-occurrence to awkwardly interpret the "grammar" of materials. LAI, by contrast, attempts to enable AI to directly learn and reason in the language of the material universe itself. The Material Morpho-Root constitutes the basic vocabulary of this language. The MSE-guided graph traversal engine provides the syntax.

This is not merely a methodological advance—it is a paradigm shift. From "data-driven screening" to "mechanism-driven design." From "statistical guessing" to "structural reasoning." From "black-box prediction" to "white-box understanding." In the safety-critical domain of extreme-condition structural materials, where failure is not an option, this shift from "what is likely" to "what is necessarily true" is the future that trustworthy AI—and the materials it helps design—deserves.

References

[1] Liu Y, Wang J, Xiao B, Shu J. Accelerated development of hard high-entropy alloys with data-driven high-throughput experiments. Journal of Materials Informatics, 2022, 2: 3. DOI: 10.20517/jmi.2022.03

[2] Han Y, Wang H, Xu P, Chen Q, Zhang R, Liu Y. Deep learning-based framework for efficient design of multicomponent high hardness high entropy alloys. ACS Applied Materials & Interfaces, 2025, 17(13): 19952-19965. DOI: 10.1021/acsami.4c23010

[3] Fan J, Jiang L, Shang G, Liu Z, Liu W. CALPHAD modeling and database development of molar volume in multi-component high-entropy alloys based on first-principles calculations. Shanghai Metals, 2026, 48(1): 96-102. DOI: 10.19947/j.issn.1001-7208.2024.05.16

[4] Huo W. Data-driven compositional and structural design in high-entropy materials. Invited presentation at China Materials Conference 2025, Xiamen, July 2025.

[5] Huo W. POLONEZ BIS project "HOT HEA": Development of nano-twinned high-entropy alloys with superior mechanical properties and enhanced irradiation resistance. National Centre for Nuclear Research, Poland, 2022-2025. Project ID: 2021/43/P/ST5/02663. Funded by Polish National Science Centre.

[6] Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. Nature. 2018;559(7715):547-555. DOI: 10.1038/s41586-018-0337-2

[7] Schmidt J, Marques MRG, Botti S, Marques MAL. Recent advances and applications of machine learning in solid-state materials science. npj Computational Materials. 2019;5:83. DOI: 10.1038/s41524-019-0221-0

[8] Hart GLW, Mueller T, Toher C, Curtarolo S. Machine learning for alloys. Nature Reviews Materials. 2021;6(8):730-755. DOI: 10.1038/s41578-021-00340-w

[9] Shang G, Jiang L, Liu ZZ, Lu XG. High-throughput experimental study on the microstructural and compositional variations of mechanical properties for AlCoCrFeNi high entropy alloys. J Alloys Compd. 2022;917:165513. DOI: 10.1016/j.jallcom.2022.165513

[10] Nørskov JK, Bligaard T, Rossmeisl J, Christensen CH. Towards the computational design of solid catalysts. Nat Chem. 2009;1(1):37-46. DOI: 10.1038/nchem.121

[11] Jain A, Ong SP, Hautier G, Chen W, Richards WD, Dacek S, Cholia S, Gunter D, Skinner D, Ceder G, Persson KA. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. APL Mater. 2013;1(1):011002. DOI: 10.1063/1.4812323

[12] Zunger A. Inverse design in search of materials with target functionalities. Nat Rev Chem. 2018;2(4):0121. DOI: 10.1038/s41570-018-0121

[13] Liu S. Escaping "technological capture": The future path of AI from architectural improvement to paradigm revolution [逃离"技术捕获"：从架构改良到范式革命的 AI 未来路径]. PSSXiv, 2025. DOI: 10.12451/202512.03460

[14] Liu S. Nested learning, gated attention, and native sparse attention: The limit optimization of the token paradigm and the paradigm revolution of Logographic AI [嵌套学习、注意力门控与原生稀疏注意力：Token 范式的极限优化与表意 AI 的范式革命]. PSSXiv, 2025. DOI: 10.12451/202512.00594

[15] Liu S. Logographic AI vs. Phonographic AI: A new paradigm and the manifesto of AI decolonization [表意 AI vs. 表音 AI：AI 新范式与去殖民化宣言]. ChinaXiv, 2025. DOI: 10.12074/202503.00051

[16] Liu S. The Rooted Intelligence: The Paradigm Revolution of Logographic AI [M]. Monograph, 2026. Available at: http://logographicai.com/Monograph.html (accessed June 27, 2026).

[17] Yeh JW, Chen SK, Lin SJ, Gan JY, Chin TS, Shun TT, Tsau CH, Chang SY. Nanostructured high-entropy alloys with multiple principal elements: Novel alloy design concepts and outcomes. Adv Eng Mater. 2004;6(5):299-303. DOI: 10.1002/adem.200300567

[18] He J, Li Z, Zhao P, Zhang H, Zhang F, Wang L, Cheng X. Machine learning-assisted design of high-entropy alloys with superior mechanical properties. Journal of Materials Research and Technology. 2024;33:260-286. DOI: 10.1016/j.jmrt.2024.09.012

[19] Miracle DB, Senkov ON. A critical review of high entropy alloys and related concepts. Acta Materialia. 2017;122:448-511. DOI: 10.1016/j.actamat.2016.08.081

[20] Fleischer RL. Substitutional solution hardening. Acta Metallurgica. 1963;11(3):203-209. DOI: 10.1016/0001-6160(63)90213-X

[21] Oriani RA. Hydrogen embrittlement of steels. Annual Review of Materials Science. 1978;8(1):327-357. DOI: 10.1146/annurev.ms.08.080178.001551

[22] Cahn JW. On spinodal decomposition. Acta Metallurgica. 1961;9(9):795-801. DOI: 10.1016/0001-6160(61)90182-1

[23] Saal JE, Kirklin S, Aykol M, Meredig B, Wolverton C. Materials design and discovery with high-throughput density functional theory: The open quantum materials database (OQMD). JOM. 2013;65(11):1501-1509. DOI: 10.1007/s11837-013-0755-4

[24] Liu S. From "Jalapeño" to "Morpho-Root": The apex of OpenAI's in-house chip and the hardware involution of Tokenism — On the path of Logographic AI's "Morpho-Root Native Chip". ChinaXiv, 2026. https://chinaxiv.org/businessFile/T202606/T202606.00643v1/T202606.00643v1.pdf