From "Magnetic Cage" to "Intelligent Web": A Paradigm Interpretation of the Breakthrough in Controlled Nuclear Fusion Superconducting Magnets from the Perspective of Logographic AI

Abstract

In June 2026, the Hefei Institutes of Physical Science, Chinese Academy of Sciences, announced that China's self-developed toroidal field superconducting magnet and high-temperature superconducting central solenoid coil for fusion reactors had successively completed acceptance tests, achieving 100% domestic manufacturing of core technologies and ranking among the world's best in overall performance. This breakthrough is a systematic triumph of materials science, superconducting technology, and extreme-condition engineering design. It also marks a critical phase in which fusion engineering is transitioning from "experience-driven" to "intelligent design."

However, current AI applications in the fusion domain remain dominated by the data-driven paradigm—whether in high-throughput material screening, plasma control, or knowledge graph construction, the cognitive primitive is the statistically associated Token, which is fundamentally incapable of meeting the essential demands for interpretability, causal reasoning, and cross-scale knowledge integration in fusion engineering.

Based on the theory of Logographic AI (LAI), this paper argues, from three dimensions—knowledge engineering, complex system reasoning, and materials design—that replacing the Token with the "Morpho-Root" as the cognitive primitive provides a paradigm pathway for fusion superconducting magnet R&D, moving from "statistical fitting" to "structural reasoning." Furthermore, it proposes a "Neuro-Symbolic Dual-System Architecture" and a "Morpho-Root Network" construction framework for fusion engineering.

Keywords

controlled nuclear fusion; superconducting magnet; Logographic AI; Morpho-Root; structural reasoning; materials design; neuro-symbolic system

I. Introduction: The "Question of Design" within the Magnet Breakthrough

On June 27, 2026, the Institute of Plasma Physics, Hefei Institutes of Physical Science, Chinese Academy of Sciences, announced two major breakthroughs: the toroidal field superconducting magnet had completed its development and acceptance, and the high-temperature superconducting central solenoid coil had passed full-operation parameter tests.

The toroidal field magnet measures 21 meters in length, 12 meters in width, and 3.3 meters in height, weighing 582 tons. It is the largest fusion reactor superconducting magnet in the world, with a stored energy three times that of comparable ITER magnets [2]. The high-temperature superconducting central solenoid coil stably carries a current of 60 kA with a stored energy of 6.03 MJ, placing its key indicators at the international leading level [2][3]. More critically, both magnet systems achieved full-chain 100% domestic manufacturing, from superconducting strands and high-strength low-temperature stainless steel to specialized insulating materials, completely breaking the foreign technological monopoly [2].

Behind these breakthroughs lie a series of arduous design decisions—decisions concerned not merely with the optimization of individual indicators but with fundamental conflicts among multiple physical constraints. This conflict can be abstracted as a "multi-objective physical conflict" problem:

- Pursuing high confinement fields demands high engineering current density and a high fill factor in superconducting strands, but this exacerbates AC losses in the conductor and induces severe thermo-mechanical coupled stress during pulsed operation.
- To achieve a 60-year design life, materials must possess extremely high resistance to irradiation damage and fatigue, which often necessitates the introduction of substantial stabilizing matrix and structural redundancy, directly reducing the engineering current density.
- To control manufacturing tolerances and costs for such a giant magnet, the structure must be simplified, yet physical effects such as stress concentration, quench propagation, and low-temperature thermal contraction compensation call for more complex geometric designs.

The project team innovatively adopted a stress-dispersing support structure and a high-low temperature hybrid magnet design [2], overcoming over ten key technical challenges. However, this "experience-driven" breakthrough model is facing the challenge of diminishing marginal returns. A deeper bottleneck lies in the "computational barrier of cross-scale simulation": traditional multi-physics coupling methods struggle to unify first-principles-scale material irradiation damage (nanometer), superconducting vortex dynamics (micrometer), and global magnet stress distribution (meter) within a single computational framework. Parameter transfer between scales relies on manual experience, causing enormous information loss and preventing the formation of a closed-loop end-to-end design optimization.

As fusion engineering moves from "principle verification" to "engineering realization," how can we make AI a true "design partner"—one that can not only screen candidate solutions but also understand "why a certain design works," trace the "causal chain of design decisions," integrate "cross-scale physical knowledge," and find unexplored Pareto-optimal solutions within the forest of physical conflicts? This is precisely the core proposition that Logographic AI theory seeks to address.

II. Achievements and Limitations of Current Fusion AI

In recent years, significant progress has been made in applying machine learning to the fusion domain. From plasma state prediction and control, to superconducting magnet design and optimization, and on to high-throughput screening of fusion materials—by accelerating the solution of Maxwell's equations through AI surrogate models, the simulation time for the electromagnetic behavior of superconducting magnets has been dramatically shortened [7]. At the plasma physics level, deep reinforcement learning has been successfully applied to real-time magnetic control of tokamaks, achieving high-precision maintenance of plasma configuration [8]. AI technology is being rapidly integrated into the entire chain of magnet design, optimization, operation and maintenance, and intelligent manufacturing, demonstrating indispensable capabilities in design optimization, performance evaluation, predictive analysis, and real-time decision-making [6]. At the knowledge management level, researchers have begun constructing fusion domain knowledge graphs, utilizing large language models for named entity recognition and relation extraction, and answering complex multi-hop scientific queries through Retrieval-Augmented Generation (RAG). The research team led by Researcher Liu Pan from the Tianfu Innovation Energy Research Institute is also exploring "AI collaborating with big data and high-performance computing to drive fusion physics design," establishing fusion knowledge graphs and material parameter databases in an attempt to bridge data silos [4].

However, these advances remain confined to the "statistical fitting" paradigm, and their fundamental limitations are manifested in three dimensions:

First, the "descriptive" limitation of knowledge graphs. Current fusion knowledge graphs store knowledge as entity-relation-entity triples (e.g., "superconducting magnet — generates — magnetic field"), with nodes and relations mostly being descriptive labels. This representation can answer "what," but struggles with "why"—it cannot reveal the physical causal chains behind "why the stress-dispersing structure can reduce thermal stress" or "how the high-low temperature hybrid design works synergistically." A more fatal flaw is that the triple structure cannot support "counterfactual reasoning." Engineers naturally ask during the design process: "If we do not use the stress-dispersing structure but increase the radial support plate thickness, how will the thermal stress distribution change?" or "If tantalum were used instead of titanium, would the low-temperature plasticity improve similarly?" These questions require the AI to perform interventions on causal graphs (do-calculus)[9], a task for which statically correlated knowledge graphs are entirely powerless.

Second, the "correlation" trap in material screening. Current machine learning tools find associations between elemental combinations and properties through statistical learning, but their reasoning process is a black box. When an AI recommends a specific alloy composition, it cannot trace why the addition of a particular element improved the material's low-temperature plasticity—because it does not understand the physical causality between dislocation motion, grain boundary sliding, and alloying elements. Even more worryingly, statistical models are prone to the "spurious correlation" trap: an alloy composition may co-occur with excellent performance in training data due to research trends and publication bias from a specific historical period, rather than a genuine causal effect. In fusion materials—a domain characterized by "low data, high cost, and long cycle time"—scarcity of samples makes this problem particularly severe.

Third, the "fragmentation" dilemma in multi-disciplinary knowledge integration. Fusion magnet design involves multiple disciplines, including superconductivity physics, structural mechanics, thermal engineering, materials science, and neutronics. The Token-based processing method of current AI systems fragments knowledge into statistical units, making it impossible to natively represent cross-scale causal chains such as "how atomic-scale lattice defects affect macroscopic stress distribution." More critically, the physical effects in fusion engineering inherently cross disciplinary boundaries—for example, "quench propagation" simultaneously involves the destruction of the superconducting state (physics), Joule heat generation (electrical science), heat conduction and convection (thermal science), mechanical deformation due to thermal stress (mechanics), and insulation breakdown (materials science). The Tokenization approach rigidly dissects these effects into independent features, destroying their intrinsic coupling.

In short, the essence of current fusion AI is a "screening assistant" in statistical space—it can efficiently find "likely optimal" candidates from known solutions but cannot "reason" about designs in the sense of causal understanding, much less execute the core cognitive cycle of engineering design: "design → counterfactual testing → causal attribution → solution iteration."

III. The Response of Logographic AI: From "Statistical Fitting" to "Structural Reasoning"

Logographic AI theory [1] offers an alternative path by replacing the Token with the "Morpho-Root" as the cognitive primitive.

In The Rooted Intelligence: The Paradigm Revolution of Logographic AI [1], Liu Shen systematically expounds the core diagnosis of this theory: the current mainstream AI paradigm—"Phonographic AI" (PAI), or "Tokenism"—is essentially an "ungrounded symbolism" system. The meaning of symbols is entirely defined by statistical co-occurrence, rather than being embedded in the structure of the cognitive primitive itself. This diagnosis has profound implications for fusion AI: when we view "stress-dispersing support structure" merely as a Token that frequently co-occurs with "thermal stress" and "mechanical load," the AI can never understand the physical causal mechanism among these three. It "recognizes" the symbol but does not "comprehend" the meaning.

As an alternative, the Morpho-Root is formalized as a structured triple:

$r = \langle S, A, R \rangle$

where S is the symbolic identifier, A is the attribute set embedding inherent physical properties and causal constraints, and R is the preset set of logical relation functions. The essential difference from a knowledge graph triple (entity-relation-entity) is that the Morpho-Root's R is not an externally attached descriptive tag, but a pre-encoded function of physical laws—it defines what physical causality can occur between this Morpho-Root and others, what constraints apply, and how these relationships are expressed mathematically. It should be emphasized that what is pre-encoded in the R set is the functional form of physical laws (e.g., "Joule heat is a quadratic function of current"), not the specific parameter values under actual operating conditions. The latter belong to a learnable instantiation layer within the Morpho-Root network and are outside the definitional scope of the cognitive primitive.

Mapping this framework onto fusion superconducting magnet design gives rise to the concept of "Fusion Morpho-Roots." Taking the "stress-dispersing support structure" as an example, its Morpho-Root can be defined as:

```
r_stress_dispersion = ⟨
  "Stress Dispersing Support Structure",
  {
    [+Structural Mechanics], [+Thermal Stress Management], [+Mechanical Strength],
    σ_yield (Yield Strength), E (Elastic Modulus), α (Thermal Expansion Coefficient)
  },
  {
    disperses(Stress Dispersing Support Structure, Thermal Stress),
    transfers(Stress Dispersing Support Structure, Mechanical Load),
    compatible_with(Stress Dispersing Support Structure, High-Low Temperature Hybrid Design),
    constrained_by(Material Yield Strength, Design Safety Factor)
  }
⟩
```

For material Morpho-Roots with deeper physical content, the R set should directly encode physical functional relationships. Taking "Niobium-tin superconductor" as an example:

```
r_Nb3Sn = ⟨
  "Niobium-tin Superconductor",
  {
    [+Superconducting], [+Brittle Ceramic], [+Strain Sensitive],
    T_c(B, ε): Critical temperature as a function of magnetic field B and mechanical strain ε,
    J_c(B, T, ε): Critical current density function,
    σ_f: Fracture Strength,
    E: Young's Modulus (~150 GPa, ceramic-like behavior)
  },
  {
    determines(T_c, J_c),
    is_a_function_of(J_c, [B, T, ε]),
    causes(Mechanical Strain ε, Reduces J_c),
    reversed_by(Reducing ε, Stress Dispersing Support Structure),    // Causal edge directly connecting to the engineering Morpho-Root
    constrains(σ_f, Safety Design),
    degraded_by(Neutron Irradiation, T_c)          // Causal edge for irradiation damage
  }
⟩
```

In this design, physical laws are no longer externally invoked tools or post-hoc verification rules, but constitutive features of the cognitive primitive. When an agent reasons "why adopt the stress-dispersing structure," it traverses the preset causal path: thermal stress concentration → stress-dispersing structure redistributes stress through geometric design → material yield strength constraint ensures safety margin → high-low temperature hybrid design synergistically optimizes the thermo-mechanical coupled response. When reasoning "how irradiation affects magnet stability," it walks along another causal edge: neutron irradiation → degradation of superconductor T_c →

narrowing of operational temperature margin → increased quench risk → protection system trigger threshold must be reset. Each step is traceable and interpretable.

This design fundamentally alters the reasoning approach of fusion AI. The table below presents a paradigm comparison:

Dimension	Current Paradigm (Statistical Fitting)	LAI Paradigm (Structural Reasoning)
Knowledge Representation	Token / Descriptive Triples	Morpho-Root (Structural unit embedding physical causality)
Reasoning Method	Statistical Association / Black-box Prediction	Causal Traversal / Walking along preset physical relation edges
Interpretability	Post-hoc Attribution ("The model says...")	Ex-ante Transparency ("Because A, and A leads to B, therefore C")
Design Capability	Screening "looks like" feasible solutions	Reasoning about "why it works" for designs
Counterfactual Reasoning	Not Supported (requires retraining)	Natively Supported (do-intervention along causal edges)
Cross-scale Integration	Token fragmentation, broken cross-scale links	Preset relation edges connecting Morpho-Roots at different scales, forming physical causality bridges

Taking cross-scale integration as an example, a typical cross-scale chain in a Morpho-Root network might be: "Irradiation Defect" Morpho-Root → leads to → "Pinning Force Enhancement" Morpho-Root → bidirectional influence (pinning enhancement increases J_c, but irradiation embrittlement decreases σ_f) → "Magnet Stability" Morpho-Root. This "bridge-like" causal connectivity is precisely what neither traditional simulation methods nor current AI can provide.

IV. A "Morpho-Root Network" Concept for Fusion Engineering

Based on the above framework, we can envision a "Morpho-Root Network" for the design of fusion superconducting magnets, embedded within a broader "Neuro-Symbolic Dual-System Architecture."

4.1 Construction of a Fusion Morpho-Root Lexicon

The Fusion Morpho-Root lexicon covers three levels, constituting a complete causal chain from microphysics to the macroscopic system:

Material-level Morpho-Roots: such as "Niobium-tin Superconductor Morpho-Root," "High-strength Stainless Steel Morpho-Root," "High-temperature Superconducting Tape Morpho-Root," "Copper Stabilizer Morpho-Root," "Low-temperature Insulation Material Morpho-Root." Each Morpho-Root embeds its key physical attributes (critical temperature, critical current density, yield strength, irradiation damage threshold, thermal conductivity, electrical resistivity) and their logical relationships with process parameters. The peculiarity of material-level Morpho-Roots is that their attributes are often functions of multiple independent variables—$J_c(B, T, \varepsilon, \text{neutron fluence})$—making the ability to handle multivariate functional causality a necessity for the R set.

Structure-level Morpho-Roots: such as "D-shaped Coil Morpho-Root," "Stress-Dispersing Support Structure Morpho-Root," "High-Low Temperature Hybrid Magnet Morpho-Root," "Pancake Winding Morpho-Root," "Box Support Structure Morpho-Root." Each embeds its mechanical behavior model (force-displacement relationship, stress-strain distribution) and constraint relationships with material-level Morpho-Roots. The key mission of structure-level Morpho-Roots is to transform abstract mechanical principles (e.g., "arch-shaped load bearing," "flexible compensation") into causal edges directly coupled with material properties.

System-level Morpho-Roots: such as "Plasma Confinement Morpho-Root," "Magnet Stability Morpho-Root," "Quench Protection Morpho-Root," "Cryogenic System Morpho-Root," "Neutron Shielding Morpho-Root." Each embeds system-level physical laws (Ampere's Law, heat conduction equation, circuit equation) and causal chains with structure-level Morpho-Roots. System-level Morpho-Roots embody the emergent effect that "the whole is greater than the sum of its parts"—for example, "Magnet Stability" is not a simple sum of the stability of individual coils, but involves complex system behavior such as multi-coil electromagnetic coupling, cooling channel distribution, and quench propagation dynamics.

4.2 Relation Function System

Five preset semantic relations constitute the grammatical basis for fusion mechanism reasoning:

Relation Type	Function Name	Physical Semantics	Example
Causality	leads_to	A directly causes a change in B	Current increase → Magnetic field strengthening → Improved plasma confinement
Inhibition	inhibited_by	A is weakened/offset by the effect of B	Thermal stress concentration → Inhibited by the stress-dispersing structure suppressing local yielding
Composition	composed_of	A is physically composed of B	D-shaped coil → Superconducting winding + Support structure + Insulation layer + Cooling pipes
Constraint	constrained_by	The value range of A is restricted by B	Operating current density ≤ min(Critical current density, Engineering current density upper limit)
Synergy	synergizes_with	The combined effect of A and B exceeds	High-low temperature hybrid design →

Relation Type	Function Name	Physical Semantics	Example
		the sum of their independent effects	Low-temperature stage provides high field + High-temperature stage provides flexible compensation

Here, the "Constraint" relation has been extended from the original "material yield strength → design safety factor" into an inequality form, more precisely expressing the boundary conditions in engineering design. The "Synergy" relation specifically captures the nonlinear coupling effects in fusion magnet design—the very aspects most difficult to master and most prone to "surprises" in empirical design.

4.3 Reasoning Mechanism: An Algorithmic Description

When a fusion engineer proposes a design objective (e.g., "the central solenoid coil must operate stably at 60 kA for 100,000 pulses"), the Morpho-Root network executes the following reasoning process. This process can be formalized as a constraint satisfaction and causal traversal problem on a heterogeneous causal graph:

(1) **Seed Activation**: The query parser identifies core concepts in the design objective, activating the "Central Solenoid Coil Morpho-Root," "High-Temperature Superconducting Morpho-Root," "High-Current Conductor Morpho-Root," and "Pulsed Operation Morpho-Root." Activation means not just "lighting up" these nodes, but loading their A sets (attributes) and R sets (relational functions) into the reasoning working memory.

(2) **Constraint Propagation**: Following preset constraint edges, the design requirements are decomposed into sub-goals and propagated to downstream Morpho-Roots. This step essentially performs a form of variational message passing on the causal graph—60 kA current → conductor cross-section ÷ critical current density → material selection → fabrication process feasibility; 100,000 pulses → fatigue life constraint → stress amplitude upper limit → structural design parameters. Each constraint edge carries the physical equation, ensuring physical consistency throughout the propagation process.

(3) **Causal Traversal**: Starting from the core variable "current," a breadth-first traversal is performed along causal edges—current → Joule heat → temperature gradient → thermal stress → stress-dispersing structure → material yield strength check → safety margin evaluation. Simultaneously, the quench causal chain is activated—current → disturbance energy → temperature disturbance → risk of superconducting state destruction → quench protection response time constraint.

(4) **Conflict Detection and Relaxation**: Check for incommensurable constraint conflicts between the attributes of activated Morpho-Roots. This is a typical Partial MAX-SAT problem—for example, the conductor cross-section required for the "target

60 kA current" may conflict with the intrinsic constraints of the "structural compactness" Morpho-Root; high current requires a large amount of copper stabilizer for quench safety, but this increases coil mass, conflicting with "support structure lightweighting." Rather than simply reporting "no solution," the Morpho-Root network seeks a set of feasible solutions on the Pareto frontier by relaxing low-priority constraints.

(5) **Solution Generation and Explanation**: Candidate design solutions are output, each accompanied by its complete decision subgraph—a full path from design objective to material selection, from physical causality to constraint satisfaction. Engineers can click on any node in the subgraph to view the complete causal chain of "why this parameter was chosen." The traceability of each design iteration ensures the possibility of knowledge accumulation and error auditing.

V. Conclusion: From "Magnetic Cage" to "Intelligent Web"—Proposing the Neuro-Symbolic Dual-System Architecture

The breakthrough in China's controlled nuclear fusion superconducting magnets is not only a triumph of engineering capability but also raises new questions about the role of AI in the fusion domain.

The "statistical fitting" paradigm of current fusion AI can efficiently "screen" solutions in massive data but struggles to "understand" the physical causality behind design decisions. The "structural reasoning" paradigm of Logographic AI offers a potential path for fusion engineering to evolve AI from a "screening assistant" into a "design partner."

We hereby explicitly propose a complementary architecture—the "Neuro-Symbolic Dual-System Architecture" for fusion engineering:

- **System 1 (Fast, Intuitive System)**: Data-driven Phonographic AI (PAI), responsible for high-throughput material screening, real-time plasma control, anomaly detection, and pattern recognition. It is the "peripheral vision" of fusion AI—fast, broad, but uninterpretable.
- **System 2 (Slow, Reasoning System)**: Morpho-Root network based on Logographic AI (LAI), responsible for causal reasoning, design solution verification, counterfactual analysis, and cross-scale knowledge integration. It is the "fovea" of fusion AI—slow, deep, but traceable.

These two systems are not in competition but in a synergistic relationship. PAI extracts candidate relations from vast literature and experimental data, while LAI solidifies these relations into a computable Morpho-Root network; PAI rapidly generates a pool of candidate solutions, while LAI performs causal verification, conflict detection, and Pareto optimization on these solutions. The fusion knowledge graph under construction by Researcher Liu Pan's team [4] can precisely serve as the initial data source for building the Morpho-Root network—the knowledge graph provides "relational hypotheses," which the Morpho-Root network upgrades into "causal imperatives."

When an AI system no longer sees "stress-dispersing structure," "high-low temperature hybrid design," and "niobium-tin superconductor" as meaningless Token fragments but as "Morpho-Roots" embedding physical causality, the AI can "walk" on a causal map—every step of reasoning is transparent and traceable, every design decision is evidence-based, and every engineering breakthrough can be deconstructed, learned, and reused.

From "Magnetic Cage" to "Intelligent Web"—what fusion engineering requires is not only stronger magnetic fields but also smarter intelligence. And the meaning of "smarter" is not "calculating faster," but "understanding more deeply." As Logographic AI theory declares: Intelligence must have roots, and civilization must have a soul.

Finally, it should be noted that the Morpho-Root network proposed in this paper is currently at the conceptual framework stage. The completeness of its formal definition, the computational efficiency of its reasoning mechanism, and its effectiveness in real fusion engineering design tasks require subsequent engineering verification. As the first step towards engineering deployment, the basic data structure and relational reasoning prototype of the Morpho-Root have been preliminarily implemented in the GitHub open-source repository (WindSeaAI/OpenMorpho) [10]. However, its integration with real fusion design data and full-chain testing remain to be advanced. Progressing this framework from theoretical conception to a deployable engineering design tool is a key task for the next stage of research.

References

[1] Liu, S. (2026). The Rooted Intelligence: The Paradigm Revolution of Logographic AI. ChinaXiv. T202606.00352. https://chinaxiv.org/businessFile/T202606/T202606.00352v1/T202606.00352v1.pdf

[2] Dai, W. (2026-06-27). China achieves new breakthrough in R&D of fusion reactor superconducting magnets. Xinhua News Agency. https://baijiahao.baidu.com/s?id=1869116568792069049&wfr=spider&for=pc

[3] CNR News. (2026-06-28). Domestic "magnetic cage" solidifies the foundation of "artificial sun": How far are we from unlimited clean energy? https://tech.cnr.cn/gstj/20260628/t20260628_527679854.shtml

[4] School of Nuclear Science and Technology, University of South China. (2023-05-30). Researcher Liu Pan from Tianfu Innovation Energy Research Institute was invited to give an academic report for our faculty and students [EB/OL]. https://hjxy.usc.edu.cn/info/1410/8100.htm.

[5] Zhang, X., Duan, S., Gao, B., et al. (2025). Application and prospect of big data and AI technology in magnetic confinement fusion. Experimental Technology and Management, 42(4), 1-13. DOI:10.16791/j.cnki.sjg.2025.04.001.

[6] Yazdani-Asrami, M., et al. (2023). Advancements in artificial intelligence for superconducting magnet design and operation. Superconductor Science and Technology, 36(7), 073001. DOI:10.1088/1361-6668/acd9e1

[7] Lu, L., et al. (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. Nature Machine Intelligence, 3(3), 218–229. https://doi.org/10.1038/s42256-021-00302-5

[8] Degrave, J., Felici, F., Buchli, J., et al. (2022). Magnetic control of tokamak plasmas through deep reinforcement learning. Nature, 602(7897), 414-419. https://doi.org/10.1038/s41586-021-04301-9

[9] Pearl, J. (2009). Causality: Models, Reasoning, and Inference (2nd ed.). Cambridge, UK: Cambridge University Press.

[10] WindSeaAI. OpenMorpho: Logographic AI Morpho-Root data structure and relational reasoning prototype [CP/OL]. GitHub. https://github.com/WindSeaAI/OpenMorpho/tree/main/OpenMorpho.