

Research Article

Towards Electromagnetic-Field Physical AI for 6G: Concepts, Architecture, and Standardization

Yajun Zhao^{1,2}, Mengnan Jian^{1,2}, Yongpeng Wu³

1. ZTE Corporation, Shenzhen, China; 2. State Key Laboratory of Mobile Network and Mobile Multimedia Technology, China; 3. School of Integrated Circuits (School of Information Science and Electronic Engineering), Shanghai Jiao Tong University, Shanghai, China

Physical AI is conventionally associated with robots, vehicles, and other mechanically embodied systems. This article extends Physical AI to electromagnetic-field embodiment, where RF-programmable hardware acts on real electromagnetic fields and propagation environments via phased arrays, reconfigurable intelligent surfaces, and programmable wireless environments. We define electromagnetic-field Physical AI through a perception–modeling–decision–execution–feedback loop, introduce a four-level hierarchy from passive perception to closed-loop field-state regulation, and propose five identification criteria to distinguish field-level intelligence from conventional parameter optimization. We further position this perspective relative to emerging world-model-based wireless intelligence, arguing that the two are complementary: world models emphasize predictive internal cognition, whereas electromagnetic-field Physical AI emphasizes physical-layer embodiment through which intelligent policies act on propagation environments. We examine the control-theoretic structure of distributed-parameter field dynamics, discuss reduced-order modeling and structured physical priors, and analyze implications for 6G architecture and standardization. A concept-validation example suggests that explicit field-level optimization can outperform parameter-centric baselines under simplified but explicit controlled assumptions. This work also proposes a layered standardization framework with field-level intent descriptors and capability abstractions, offering a starting point for future programmable wireless environment standards.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

I. Introduction

Physical AI is emerging as an important concept for next-generation intelligent systems ^{[1][2][3]}. While it is often discussed in the context of mechanical embodiment, its underlying principle is broader: intelligence can interact with the physical world through any executable physical substrate. Unlike conventional AI that mainly processes information in digital space, Physical AI operates through closed loops connecting sensing, modeling, decision-making, execution, and feedback in the real world. Most existing discussions focus on robots, autonomous vehicles, humanoids, and related embodied systems, where embodiment is naturally understood through bodies, joints, motors, contact forces, and motion planning. This robotics-centered perspective has produced remarkable advances in manipulation, locomotion, and autonomous navigation.

However, this view, while essential, is incomplete. In wireless and electromagnetic systems, intelligent agents can also act on the physical world through continuous fields rather than discrete mechanical motion. Phased arrays, reconfigurable intelligent surfaces (RIS), Network-Controlled Repeaters (NCR), and distributed radio nodes increasingly enable AI-driven regulation of propagation paths, spatial energy concentration, multipath structure, and regional coverage patterns ^[4]. When such systems sense the electromagnetic environment, infer control actions, execute them through RF hardware, and update based on feedback, they generally exhibit the defining features of Physical AI without moving a single mechanical joint. In hybrid scenarios (e.g., movable antennas or UAV-mounted RIS), field control may be jointly executed with mechanical actions.

We argue these systems merit recognition as a distinct complementary branch of Physical AI: electromagnetic-field Physical AI. We do not relabel every AI-enabled wireless method as Physical AI, nor replace mechanical embodiment with an electromagnetic one. Rather, we clarify when an electromagnetic system transitions from conventional parameter optimization to field-level physical intelligence. This distinction affects research paradigms, architecture, evaluation, interoperability, and standardization, determining whether a wireless system is merely a communication optimizer or an intelligent agent that shapes its propagation environment. Embodiment is not confined to mechanics; in electromagnetic systems, intelligence can directly act on real propagation environments via RF hardware.

Three concurrent trends make this discussion timely. First, infrastructure programmability (RIS), NCR ^[5], large arrays, and distributed nodes are turning the radio environment into a partially controllable substrate, shifting wireless systems from passive adaptation toward active environmental control.

Second, physics-aware AI, including physics-informed neural networks, neural operators, and model-based learning, is making it easier to embed electromagnetic structure into inference and control [\[6\]\[7\]\[8\]\[9\]](#). Third, 6G convergence integrating communications, sensing, localization, and environmental awareness makes purely link-centric objectives insufficient in some scenarios. Networks may need to illuminate one region for sensing, maintain coverage in another for communication, and suppress leakage elsewhere for privacy or interference control [\[10\]](#).

Recent work on world-model-enabled closed-loop intelligence in wireless cyber-physical systems [\[11\]](#) reinforces the practical stakes of this discussion. These systems learn internal models to forecast environmental changes, communication conditions, and downstream effects before taking physical actions. Yet predictive modeling alone does not resolve the embodiment question: by what physical means does intelligence actually act on the environment? World-model-based reasoning and electromagnetic-field Physical AI thus complement, rather than oppose, each other. The former excels at cognitive prediction and latent policy search; the latter specializes in field-level actuation and direct manipulation of real propagation channels. A future 6G system could use world models to screen candidate policies, while relying on field-embodied RF hardware to carry them out in space.

In summary, this article provides:

1. We extend Physical AI to the electromagnetic field domain, defining electromagnetic-field Physical AI and proposing a four-level hierarchy together with five identification criteria. We explain how electromagnetic-field Physical AI complements the emerging predictive world model paradigm in future wireless intelligent systems.
2. We develop an engineering interpretation grounded in approximate target-region field regulation, structured physical priors, and enabling hardware paths.
3. We examine 6G architectural implications and propose a candidate standardization perspective with layered abstractions and building blocks for smart radio environments.

The remainder of this article is organized as follows. Section II presents the framework and criteria. Section III discusses engineering realities and physical priors. Section IV analyzes 6G relevance. Section V provides concept validation and open challenges. Section VI explores standardization implications. Section VII concludes.

II. Concept and Framework

This section aims to establish a systematic conceptual foundation for electromagnetic-field Physical AI. Given that intelligent behavior in RF systems is often equated with conventional parameter optimization, we propose a four-level progression (L1–L4) to clearly distinguish systems that merely process or adapt to electromagnetic phenomena from those that explicitly regulate real propagation fields as controllable physical objects. On this basis, we provide an engineering-oriented working definition, explain the unique structure of field embodiment from a control-theoretic perspective as a closed-loop system linking sensing, state estimation, decision-making, and actuation through physically executable interactions with the propagation field, and then distill five practical identification criteria to offer operational guidelines for system classification and evaluation.

A. A Four-Level Progression

The evolution of AI is often framed as four stages: perceptual AI, generative AI, agentic AI, and Physical AI. To distinguish conventional electromagnetic intelligence from electromagnetic-field Physical AI, this article introduces an L1–L4 hierarchy corresponding to the above-mentioned four stages of conventional AI evolution ^[12], which traces EM-related intelligent systems from reading the electromagnetic world, to generating EM objects, to optimizing EM parameters, and finally to closed-loop field regulation (Fig.1).

- **L1: Passive Electromagnetic Perception**

At L1, electromagnetic waves are used to observe the world, while the electromagnetic field itself is not treated as a control target. AI supports radar recognition, channel estimation, spectrum sensing, and electromagnetic imaging. L1 systems interpret the environment via electromagnetic signals, without regulating propagation structure.

- **L2: Digital Electromagnetic Generation**

At L2, AI generates electromagnetic objects in digital space, such as channel samples, propagation maps, synthetic echoes, radiation patterns, or simulation environments. These outputs are useful for design and training, but they remain representational rather than physically enacted. Typical use cases include AI-generated channel datasets, digital-twin propagation scenes, and synthetic radar-echo libraries. This supports design, training, and evaluation, but the generated object remains digital rather than physically enacted.

▪ L3: Parameter-Level Adaptive Optimization

At L3, AI optimizes frequencies, power, beam weights, schedules, or other link-level variables. Cognitive radio, adaptive beam management, and many RIS-assisted communication schemes fit here [4]. The system reacts intelligently to channel conditions, but its objective remains link metrics such as SINR, throughput, latency, or BER. L3 systems are adaptive users of the channel, not explicit shapers of the field as a spatial control object. Their objective remains link-centric, whereas L4 elevates the field itself to the primary control target.

▪ L4: Closed-Loop Electromagnetic Environment Reconstruction

At L4, the system uses RF hardware to regulate regional field-state distributions, propagation topology, or spatial energy patterns based on feedback. In this sense, L4 represents electromagnetic-field embodiment rather than merely region-based optimization. L4 is not merely a region-based extension of link optimization; rather, it treats the field distribution itself as an explicit control object within a task-relevant spatial domain. The defining shift is that the optimization target moves from link parameters to field behavior over a task-relevant region [4]. Typical scenarios include uniform coverage shaping in an indoor area, interference nulling over a protected zone, and joint communication-sensing energy allocation across a spatial region. Here, the field distribution itself serves as an explicit control target, rather than merely its impact on individual links.

The division between layers is not clear-cut but gradual. From a practical perspective, the distinction becomes operationally meaningful when field state, physical constraints, and feedback-based actuation dominate the control loop. While control parameters affect electromagnetic fields mathematically, the key differentiation hinges on whether field distribution serves as the dominant explicit control entity, embedded within optimization objectives, state characterization, and feedback loops. In short, L3 optimizes link metrics, treats the environment as static and the field as a hidden intermediary; field variations are merely secondary outcomes of parameter tuning. L4 directly programs spatial field distribution, regards the environment as a dynamic, reconfigurable substrate, and takes the field as the primary state to govern. The contrast lies not in region size but in control ontology: whether the field is a byproduct of parameter tuning (L3) or the primary state to be regulated (L4).

A system moves closer to L4 when several features appear simultaneously: 1) the objective is explicitly region-based rather than only terminal-based; 2) the state or feedback includes field-relevant spatial descriptors, mode coefficients, energy maps, or regional statistics; 3) physics is embedded into the

control structure through field-consistent models or constraints ^{[6][7][13][9]}; and 4) the system acts on the real field through executable hardware with repeatable feedback ^[4]. Not every L4-oriented system must satisfy all of these features in the same form, but the more of them that are present, the stronger the case for field-level embodiment.

For example, a beam-tracking system that maximizes received signal strength at one terminal remains in L3. By contrast, a system that aims to maintain near-uniform coverage across a factory floor while suppressing leakage of electromagnetic interference signals outside the facility is better described as L4, provided that regional field behavior enters the control loop explicitly.

To make this progression more operational, Table I contrasts conventional link-centric optimization, L3 parameter-level adaptation, and L4 field-level regulation. The table is not intended as a rigid taxonomy for every wireless system; rather, it provides a compact diagnostic view of where the dominant control abstraction resides.




Figure 1. Electromagnetic-field Physical AI: Four-Level Progression

Aspect	Conventional Link-Centric Optimization	L3 Parameter-Level Adaptation	L4 Field-Level Regulation
Primary objective	Fixed or predesigned link operation	Adaptive optimization of link/service metrics	Electromagnetic environment reconstruction via multi-region field programming
State representation	Static parameters and basic link measurements	Channel state and adaptation context	Field-relevant spatial descriptors, modes, or regional statistics
Role of physics	Mostly background modeling assumption	Partially embedded in optimization or learning	Explicit structural constraint in control and inference
Feedback type	Limited link measurement or preset configuration	Terminal measurements plus control response	Regional or field-aware observations with closed-loop update
Execution target	Preset transmitter/resource settings	Adaptive beam/power/RIS parameter updates	Mapping of environment-level control intents to physical field states through RF-programmable hardware, enabling global reprogramming of spatial propagation conditions
View of the field	Implicit effect behind channel metrics	Secondary consequence of parameter tuning	Explicit control object; the environment is a reconfigurable resource to be programmed, not a fixed medium to be adapted to

Table I. Comparison of Conventional Link-Centric Optimization, L3 Parameter-Level Adaptation, and L4 Field-Level Regulation

Aspect	Existing Link-Centric / L3 Practice	L4 Field-Level Regulation Need
Objective expression	UE-centric QoS metrics such as SINR, throughput, and latency	Region-level field intent such as coverage uniformity, null-zone protection, and joint ISAC illumination
State abstraction	CSI, CQI, RSSI, and adaptation context	Reduced-order field state, such as spatial modes, regional energy maps, or task-relevant mode coefficients
Feedback type	Terminal measurements and control responses	Field-aware observations with closed-loop spatial update
Execution target	Beam weights, power levels, and RIS coefficient tuning	Environment-level intent mapping: translating multi-region field objectives into executable hardware commands that jointly reprogram spatial propagation conditions
Role of physics	Background modeling assumption or partial constraint	Explicit structural constraint embedded in control and inference
Standard implication	Largely supported by existing link-centric abstractions	May require additional field-aware abstractions in selected scenarios

Table II. From Link-Centric Control to Field-Level Regulation: Implications for Standard Abstractions

B. Working Definition and Control-Theoretic Structure

We define electromagnetic-field Physical AI as an intelligent system in which RF-programmable hardware forms the agent boundary; the real electromagnetic field is the primary agent-environment interaction medium; field evolution is governed by Maxwell’s equations and shaped through controllable excitation sources and programmable boundary/interface conditions, subject to hardware limits; and a perception-modeling-decision-execution-feedback loop approximately reshapes the field over task-relevant spatial regions [6][7][4][13][9]. The intended scope is engineering-level approximate field regulation under finite resources.

This definition avoids two common mistakes. One is to treat any tunable RF system as field intelligence. The other is to imply unrealistic global control over all space. We deliberately use the term approximate regulation rather than exact control to acknowledge the limits imposed by finite aperture, finite sensing, finite bandwidth, and finite control resolution.

From a control-theoretic perspective, the contrast between mechanical and field embodiment is instructive. Mechanical Physical AI typically operates on finite-dimensional state spaces of joint angles, velocities, and forces, where the agent boundary is a rigid or soft body with localized contact [1][2]. Field embodiment, by contrast, operates on distributed-parameter systems governed by partial differential equations [13]. The state is a field over a continuous spatial domain rather than a finite vector, while the control input acts through excitation sources, array weights, or programmable boundary conditions. The observation space is also indirect: the full field is not directly available, but only sparsely sampled through terminals, antenna ports, or sensing nodes.

This structural difference has important implications for control design. In mechanical systems, controllability and observability are often framed through finite-dimensional models. In field systems, exact controllability over the full state space is generally unrealistic with finite actuators and sparse observations. The more appropriate notion here is engineering-oriented approximate reachability over task-relevant regions: given finite hardware, can the field in a designated region be driven sufficiently close to a desired pattern for the task? This reframing shifts the question from “can we control everything?” to “can we control what matters?” and makes field-level Physical AI practically meaningful.

C. Five Identification Criteria

A practical way to judge whether a system qualifies as electromagnetic-field Physical AI is to check five criteria. Nevertheless, the distinction is operationally meaningful when field state, physical constraints, and feedback-based actuation become dominant elements of the control loop.

- **Physical reality:** The action target must be a real electromagnetic field or propagation structure, not merely a digital estimate or offline artifact.
- **Physical constraint:** The solution must remain feasible under the physical laws of wireless electromagnetic waves (i.e., Maxwell's equations), hardware capabilities, power budgets, and regulatory requirements [6][7][4][13].

- **Executability:** Control decisions must map to real hardware actions, such as array weights, RIS states, or impedance settings ^[4].
- **Closed-loop operation:** The system must sense, decide, execute, and update based on feedback in a repeatable loop.
- **Field-level objective capability:** The primary objective should involve regional field control, spatial coverage shaping, or multipath restructuring, rather than only pointwise link metrics. In practice, this means controlling resolvable spatial modes within the target region, subject to limits imposed by aperture, wavelength, sensing resolution, and available control degrees of freedom ^{[4][13]}.

These criteria provide a practical filter for distinguishing electromagnetic-field Physical AI from conventional communication optimization. They also suggest why systems satisfying all five criteria may eventually require new interfaces and abstractions beyond current link-centric standards ^[4].

Systems satisfying these criteria are further distinguished from conventional wireless optimization by three concise features. First, the electromagnetic field itself becomes an explicit control state rather than a hidden intermediary behind channel abstractions. Second, physical consistency is enforced as a structural ingredient of inference and control rather than treated merely as a post hoc validation check. Third, feedback closes through regional or field-aware observations, not only through terminal-centric link metrics.

D. A Formulated View

At an engineering level, the core task is to adjust controllable hardware settings, such as array amplitudes and phases, RIS reflection states, and tunable impedances, to improve task performance, whether the goal is stronger regional focusing, deeper suppression in a protected zone, or better energy efficiency. The defining feature is not merely that parameters are tuned, but that the adjustments serve the closed-loop shaping of real spatial field distributions under explicit physical constraints ^{[6][7][4]}.

To make the L3/L4 distinction more concrete, an L3 system typically solves a problem of the form

$$\min_{u \in \mathcal{U}} J_{\text{link}}(\text{SINR}(u), R(u), \text{BER}(u), \dots) \quad (1)$$

where $u \in \mathcal{U}$ is the control parameter vector (e.g., beam weights, power levels, RIS coefficients), $\mathcal{U}$ is the feasible set constrained by hardware and regulations, and J_{link} is a scalarized link-centric objective that aggregates point-wise quality metrics (note that metrics such as achievable rate $R(u)$ are typically

derived from SINR but are kept as separate arguments here for notational generality). In this formulation, the electromagnetic field affects the objective only implicitly through the channel response.

An L4 system, by contrast, moves closer to a formulation such as

$$\begin{aligned} \min_{u \in \mathcal{U}} J_{\text{field}} \left(\mathbf{E}|_{\mathcal{D}_{\text{target}}}, \mathbf{H}|_{\mathcal{D}_{\text{target}}} \right) \\ \text{s.t. } M(\mathbf{E}, \mathbf{H}, \mathbf{u}) = 0 \end{aligned} \quad (2)$$

where $M(\cdot) = 0$ compactly denotes Maxwell's equations with programmable boundary and source terms determined by the control parameters $\mathbf{u}$. Here, the field distribution over the target region enters the objective explicitly, and propagation physics appears as a central constraint rather than a background assumption [6][7][13][9]. The exact formulation will vary by application, but the conceptual distinction remains the same: implicit field effects versus explicit field objectives. The critical distinction is that L3 optimization aggregates per-link metrics over terminals, whereas L4 optimization directly programs the spatial field over regions, enabling trade-offs that link-centric formulations cannot capture. The objective is not merely to minimize link-level error, but to realize a task-relevant field state, such as region-specific coverage enhancement, interference suppression, leakage confinement, or joint sensing-communication illumination.

Aspect	World-Model-Based Wireless Intelligence	Electromagnetic-Field Physical AI	Relationship
Core role	Predictive internal cognition and latent-state estimation	Physical actuation and embodiment through RF-programmable hardware	Complementary
Primary object	Imagined environmental dynamics and policy rollouts	Real field distribution in physical space	Sequentially coupled: predict, then act
Key strength	Screening candidate policies before physical execution	Executable field regulation with physical-consistency enforcement	Co-design opportunity in closed-loop systems
Hardware dependence	Indirect, typically through cloud or edge computation	Direct, through phased arrays, RIS, and programmable boundaries	Distinct but integrable
Typical operational tendency	Often slower-timescale model inference and policy evaluation	Often faster-timescale hardware execution and feedback-driven adaptation	Naturally supports hierarchical control splits

Table III. Positioning World-Model-Based Wireless Intelligence and Electromagnetic-Field Physical AI

III. Engineering Realities and Physical Priors

A. From Infinite-Dimensional Fields to Practical Control

Practical field control is limited by finite aperture, element count, phase resolution, update rate, power budget, and hardware losses ^[4]. Exact full-field controllability is infeasible; only task-relevant, reduced-dimensional representations can be steered ^[13]. The core issue is approximate reachability over a target region: given fixed resources, can the system shape the field in that region close enough to the desired pattern? Overall performance is degraded by reduction errors, sparse observations, noise, coupling, environmental mismatch, and implementation imperfections. Consequently, field-level Physical AI is viable only in bounded, partially structured environments, e.g., indoor ISAC or RIS-assisted coverage. The fundamental limit is that finite sensing prohibits exact global regulation; what remains is approximate control over relevant regions, constrained by aperture size, wavelength, and sensing density.

B. Reduced-Order and Surrogate Modeling Paths

All practical controllers rely on compressed field representations. Two mainstream modeling paradigms are adopted.

- **Modal or Galerkin-type reduction:** This method characterizes electromagnetic fields via finite basis functions (e.g., dominant eigenmodes and POD modes) and regulates corresponding modal coefficients. It delivers clear physical interpretability yet is sensitive to geometric variations, material properties, and environmental dynamics [13].
- **Operator-learning or surrogate parameterization:** This paradigm learns the implicit mapping between control inputs and field responses without explicit modal decomposition, with typical models including FNO and DeepONet. It is well-suited for scenarios with complex geometries or time-varying boundaries. Different from classical reduced-order models, such operator-level surrogates implement task-oriented engineering approximation for field control. [8][9].

For both schemes, the effective controllable degrees of freedom are determined by the numerical rank of the control-to-field and observation mappings within the target region. Compact low-dimensional representations function well under favorable propagation conditions, while rich multipath components and sharp field discontinuities will increase the effective state dimension.

C. Structured Physical Priors

Embedding inherent electromagnetic physical structures outperforms pure data-driven statistical modeling. Representative physical priors are summarized as follows.

- **Spatiotemporal continuity:** Electromagnetic fields vary smoothly in both space and time, which improves sampling efficiency and spatial interpolation accuracy. [6][7].
- **Topology-aware smoothness:** Adjacent array elements exhibit correlated responses due to shared scattering effects and electromagnetic coupling, and appropriate regularization can significantly enhance system robustness. [4][5][14].
- **Hard engineering constraints:** Practical constraints including power limits, EIRP masks, phase quantization, dynamic range, interference thresholds, and spectral regulations should be embedded into the solution space during modeling, rather than being imposed as post-hoc corrections. [6][7][4].

These physical priors improve the physical plausibility of control outputs, reduce computational search complexity, and effectively narrow the sim-to-real gap.

D. Enabling Hardware Paths

Three promising hardware implementations support field-level Physical AI.

- **Active phased arrays:** Capable of precise amplitude and phase control. This technology is mature but constrained by high power consumption and hardware costs.
- **RIS and programmable boundaries:** Enable partial propagation environment control. They redistribute electromagnetic energy with ultra-low power consumption yet the mainly reshape rather than actively radiate energy, while active RIS can mitigate this limitation to a certain extent. [4][6][7][8].
- **Programmable electromagnetic environments:** Integrate and coordinate base stations, RIS, sensors, and edge controllers. This architecture aligns well with the 6G vision but faces inherent challenges in system complexity, real-time control latency, and deployment costs. [4].

Practical deployment does not require online solving of full field equations at each time step. A feasible architecture offline learns and calibrates physics-consistent surrogate models, executes low-dimensional control generation and inference at the edge side, and implements high-speed RF control at the hardware layer [6]–[8][9].

IV. Potential Impact on 6G System Architecture

The shift from link-centric optimisation to field-aware control has several architectural implications for 6G.

A. From Link-Centric to Environment-Centric Control

Traditional wireless design is link-centric (rate, delay, interference, beam alignment). In 6G, the objective may broaden to environment-centric control, making propagation itself partially programmable. EM-field Physical AI provides a useful language for this transition.

B. Natural Support for ISAC

ISAC inherently requires region-level objectives—illuminating one area, preserving coverage elsewhere, and suppressing leakage in a third. These spatial multi-region tasks are not well captured by single-link

metrics, yet they are identified as key 6G capabilities in standardization roadmaps^[10].

C. Split Intelligence Across Cloud, Edge, and RF Hardware

A hierarchical deployment model is likely: world modelling and policy refinement run at cloud/edge timescales, while RF hardware (beam weights, RIS states) executes compact control signals at much faster timescales.

This split clarifies the relationship between predictive intelligence and embodiment. A world model can estimate latent states, forecast propagation, or evaluate candidate policies via rollouts—but these predictions still need a physically meaningful actuation substrate. EM-field Physical AI provides that substrate by linking high-level policy outputs to executable field actions via arrays, RIS panels, and programmable radio nodes.

In a co-designed pipeline, the world model generates candidate policies (slower timescale) and screens them; the field-embodied layer executes the selected policy (faster timescale) and returns real observations to refine the model. The key interface is a field-state abstraction layer: the world model outputs spatial intents/target descriptors, and the field-embodied layer maps them to hardware actions under real-time physical constraints. In essence, world models decide which policy to try; field-embodied Physical AI determines how to physically realise it.

D. Implications for AI-Native Air Interfaces

Future AI-native air interfaces may extend optimisation beyond coding/modulation/beam selection to include region-specific field objectives, exposure-aware shaping, and joint communication-sensing field design^[4]. This does not replace link-level abstractions but suggests that field-aware control primitives may become useful in selected scenarios.

V. Concept Validation and Open Challenges

A. A Concept-Validation Example

This concept validation considers a simplified indoor multipath scenario (28 GHz, 2D single-frequency TE approximation) and compares three schemes: static (no adaptation), L3 Sparse (parameter-level link optimization), and L4 Dense (explicit field-level regulation), along with regularized L4 variants. Key findings (shown in Fig.2 and Fig.3):




Figure 2. Gains Comparison




Figure 3. Performance metrics of Balance Comparison

- **Focus gain (dB):** Static 12.33, L3 11.77 (−0.56), L4 13.07 (+0.73). The sign reversal shows that link-centric optimization misdirects energy away from the target region, while field-level regulation recovers and enhances regional focusing.
- **Gain improvement over static (dB):** L3 degrades regional gain (−0.56 dB) vs. static, while L4 improves it (+0.73 dB)—the sign reversal proves that field-level regulation consistently outperforms link-centric adaptation.
- **Null suppression (dB):** Static (−1.64) and L3 (−1.53) are negative (ineffective); L4 achieves +1.74, a qualitative shift to active spatial energy control.
- **Background leakage (dB):** L3 marginally reduces it (15.33 vs. 16.94); L4 maintains comparable leakage (15.50) while improving target-region performance.

- **Target contrast (dB):** Improves progressively from -4.60 dB (static) to -3.56 dB (L3) and -2.43 dB (L4), reflecting enhanced isolation of the target region from background interference.

In summary, field-level objectives consistently outperform link-centric optimization in focusing, nulling, contrast, and uniformity. The L3/L4 distinction is qualitative—L3 even degrades relative to static, while L4 consistently improves. This is an algorithm-in-the-loop validation under controlled assumptions; hardware impairments, wideband effects, sparse observability, and protocol implementation remain open challenges.

B. Open Challenges

Several challenges remain before electromagnetic-field Physical AI can mature into a deployable 6G capability.

- **Reduced-order modeling:** Compact field representations must remain physically meaningful and online-updatable, especially in dynamic wideband scenarios. ^{[8][13][9]}
- **Sparse observation and physical inconsistency:** Sparse sampling may cause controllers to overfit measured data and yield “physical hallucinations”—plausible control behaviors that violate physical consistency or induce unintended field effects in unobserved regions. Mitigation requires rigorous physical constraints, uncertainty-aware inference, and rapid consistency validation. ^{[6][7]}
- **Sim-to-real transfer:** Quantization errors, electromagnetic coupling, system nonlinearity, parameter drift, fabrication imperfections, and calibration mismatches can severely degrade real-world control performance compared with simulation results. ^[4]
- **Scalable coordination:** Large-scale programmable environments demand hierarchical or distributed control paradigms, as centralized optimization becomes intractable with an increasing number of controllable elements ^[4].
- **Safety and compliance:** Field-level regulation may impact unintended receivers and open-space propagation, raising concerns regarding electromagnetic exposure, interference management, spectrum compliance, and technical accountability.
- **Fusion with mechanical embodiment:** Future hybrid systems integrating robots, drones, mobile platforms, and programmable wireless environments require unified control frameworks, which remain largely underexplored. ^{[1][2]}

These challenges demand coordinated research efforts spanning wireless theory, control engineering, hardware design, and standardization development.

VI. Standardization Implications

A. Why Standardization May Become Relevant

Large-scale and multi-vendor deployment of programmable wireless environments will require unified interoperability specifications ^[4]. Current wireless standards focus on device-level and link-level functions, including scheduling, CSI reporting, beam management, and resource allocation. However, they cannot natively support region-level field optimization and environment-oriented control, which has been recognized in recent 6G standardization studies ^[15].

Three key technical gaps remain:

1. The absence of standardized interfaces for field-level control intent, such as regional coverage shaping and protected-area nulling.
2. Limited feedback mechanisms that rely only on conventional CSI and pointwise link measurements.
3. Heterogeneous and proprietary control interfaces for programmable boundary hardware ^{[4][15]}.

Instead of rebuilding the entire standard stack, future 6G systems can introduce additional field-oriented abstractions on top of existing link-layer frameworks, consistent with global 6G roadmaps ^[10].

B. A Candidate Layered Architecture and Building Blocks

This section presents a four-layer conceptual architecture that decouples universal interoperability functions from vendor-specific optimization algorithms.

- **Sensing layer:** Collects spatial measurements from terminals, antenna arrays, and sensing nodes while maintaining backward compatibility with existing standards. It extends current CSI-RS, SRS, and beam reporting frameworks. For RIS-enabled scenarios, UE-transparent sensing and control align with ongoing NCR standardization progress ^[5].
- **Field-state abstraction layer:** Converts raw measurements into hardware-independent spatial representations, including beamspace coefficients, dominant propagation modes, spatial covariance

features, and uncertainty statistics. Different from conventional NWDAF and RAN controllers that process link-level KPIs, this layer focuses on spatial field characteristics.

- **Control-intent layer:** Defines hardware-agnostic field-level optimization targets, such as regional focusing, null-zone suppression, exposure-aware shaping, and customized coverage masks. It extends existing intent-based management from service-level requirements to spatial field control and distinguishes field-oriented optimization from traditional link optimization.
- **Execution layer:** Translates high-level spatial intents into device-specific configurations, including array weight vectors, RIS phase states, and control commands, while reserving internal optimization algorithms for vendor implementation.

C. Governance and Phased Evolution

Electromagnetic field regulation affects open-space propagation beyond individual devices or vendors, so technical interfaces alone are insufficient. Governance mechanisms—auditability, policy enforcement, access control, anomaly detection, and explainability—are essential for reliable deployment. Field-level actuation must specifically ensure EMF exposure compliance and coexistence interference management, preventing regulatory violations or unintended disruptions to other services.

A phased standardization roadmap is feasible. In the near term, enhanced capability reporting and environment sensing will be prioritized. In the mid term, 6G standardization will incorporate spatial intent abstraction and programmable environment interfaces. In the long term, cross-domain coordination, safety certification, and hybrid mechanical-electromagnetic control will be investigated. This staged evolution is illustrative and conforms to the official 6G deployment roadmap ^[10].

VII. Conclusion

We argue that field-domain level control is a distinct branch of Physical AI, complementary to mechanical embodiment, and central to 6G, ISAC, and programmable environments—marking a shift from environment adaptation to environment reconstruction. We propose a four-level hierarchy (L1–L4) and five criteria to distinguish conventional EM intelligence from field-embodied Physical AI, emphasising approximate regional regulation under physical constraints rather than global exact control. This perspective complements emerging world-model-based architectures: predictive internal cognition and field-domain level embodiment address related but separate problems, and future systems

may require both. We also outline structured physical priors, hardware paths, and implications for 6G standardization, offering a layered view and building blocks for future discussion.

The core takeaway is that when AI acts on real EM fields through executable hardware, with explicit constraints and feedback, it transcends ordinary parameter tuning and enters field-embodied Physical AI. Recognising this distinction provides the communications community with a clearer conceptual framework and a solid basis for discussing the evolution of programmable wireless environments in 6G and beyond.

Statements and Declarations

Funding

This work was supported in part by the China National Key R&D Program under Grant 2026ZD1307100, Research on Key Technologies of AI Inference Services for 6G Networks.

Use of Generative AI

During the preparation of this manuscript, the authors used Deepseek for language refinement and proofreading purposes. The authors have reviewed and verified all AI-assisted edits and take full responsibility for the content of the manuscript.

References

1. ^{a, b, c}"From Embodied Intelligence to Physical AI." (2026). *Nat Mach Intell.* 8(4):491–492. doi:[10.1038/s42256-026-01239-3](https://doi.org/10.1038/s42256-026-01239-3).
2. ^{a, b, c}Wu EH, Liu YQ, Xu TC, Ren LX, Qin YM, Wei MY, He XW, Yuan DY, Hou WC, Ma ZW, Sheng B (2026). "Physical AI: Evolution, Progress, Challenges, and Prospects." *J Comput Sci Technol.* 41(1):271–288. doi:[10.1007/s11390-026-6258-x](https://doi.org/10.1007/s11390-026-6258-x).
3. ^ΔNvidia (2025). "What Is Physical AI?" Nvidia. <https://www.nvidia.com/en-us/glossary/generative-physical-ai/>.
4. ^{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t}Di Renzo M, Zappone A, Debbah M, et al. (2020). "Smart Radio Environments Empowered by Reconfigurable Intelligent Surfaces: How It Works, State of the Art, and the Road Ahead." *IEEE J Sel Areas Commun.* 38(11):2450–2525.

5. ^{a, b, c}Zhao Y, Dai L, Yuan Y, et al. (2024). "Evolution From NCR for 5G-Advanced to RIS for 6G: A Unified System Framework and Standardization Roadmap." *IEEE Commun Stand Mag.* Accepted.
6. ^{a, b, c, d, e, f, g, h, i, k}Raissi M, Perdikaris P, Karniadakis GE (2019). "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations." *J Comput Phys.* **378**:686–707.
7. ^{a, b, c, d, e, f, g, h, i}Karniadakis E, Kevrekidis IG, Lu L, et al. (2021). "Physics-Informed Machine Learning." *Nat Rev Phys.* **3**(6):422–440.
8. ^{a, b, c, d, e}Li Z, Kovachki N, Azizzadenesheli K, et al. (2021). "Fourier Neural Operator for Parametric Partial Differential Equations." in *Proc. ICLR*.
9. ^{a, b, c, d, e, f, g}Lu L, Jin P, Pang G, et al. (2021). "Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators." *Nat Mach Intell.* **3**(3):218–229.
10. ^{a, b, c, d}ITU-R (2023). "Framework and Overall Objectives of the Future Development of IMT for 2030 and Beyond." *Recommendation ITU-R M.2160-0*.
11. ^aHan X, Wu Y, Shen M, et al. (2026). "Towards World Model-Empowered Integrated Sensing, Communication, and Decision for Complex Unmanned Systems." *IEEE Communications Magazine.* Accepted. <https://arxiv.org/abs/2606.30568>.
12. ^aNVIDIA (2025). *CES 2025 Highlights Deck*. Santa Clara: NVIDIA Corporation.
13. ^{a, b, c, d, e, f, g, h, i}Lions JL (1971). *Optimal Control of Systems Governed by Partial Differential Equations*. Berlin, Germany: Springer.
14. ^aSatorras VG, Hoogeboom E, Welling M (2021). "E (n) Equivariant Graph Neural Networks." in *Proc. ICML*, p. 9323–9332.
15. ^{a, b}ETSI (2023). "Reconfigurable Intelligent Surfaces (RIS); Technological Challenges, Architecture and Impact on Standardization." *GR RIS 002 V1.1.1*.

Declarations

Funding: This work was supported in part by the China National Key R&D Program under Grant 2026ZD1307100, Research on Key Technologies of AI Inference Services for 6G Networks.

Potential competing interests: No potential competing interests to declare.