Theory and Practice of Science and Technology
2026, VOL. 7, NO. 2, 94-96
DOI: 10.47297/taposatWSP2633-456918.20260702

Rethinking Embodied AI: The Minimalist Model of Embodiment (LS-MME) for Decoupling Morphology and Computation and Its Socio-Technical Implications

Shen Lin

Guangdong Huayan Robot Co., Ltd., Shenzhen, Guangdong, 518000, China

ABSTRACT
The current obsession in the embodied AI industry with developing "all-in- one" humanoid robots is leading to unsustainable edge computing demands and prohibitive capital expenditures (CAPEX). This paper proposes the "Lin- Shen Minimalist Model of Embodiment" (LS-MME), advocating a fundamental paradigm shift that completely decouples a robot's physical morphology, envi- ronmental perception, and edge computing. The model establishes a systematic methodology for hardware simplification through three core principles: Minimum Viable Morphology (MVM), Locally Shared-Perception (LSP), and Minimum Compute Requirement (MCR). Research shows that this framework not only pro- vides a highly feasible engineering stepping stone to solve the PSPACE-complete complexity disaster in Heterogeneous Embodied Agent Training (HEAT), but also unlocks massive economic and practical value. By exponentially compressing the Total Cost of Ownership (TCO). Based on the macroeconomic and techno- logical deductions of this framework, we ofer a disruptive prediction: that AI will move away from isolated humanoids and evolve into a "distributed macro- organism network," a shift that will exponentially lower deployment costs and reshape our physical workspaces.

KEYWORDS
Embodied AI; Lin-shen minimalist model of embodiment; Decoupling architecture; Distributed macro-organism

1 Introduction: The Scalability Paradox and Economic Challenges of Humanoid Robots

The commercial deployment of embodied AI faces a critical structural bottleneck: academia and industry are largely trapped in a mindset that equates ″Artificial Gen- eral Intelligence (AGI)″ with ″humanoid robots.″ This architectural obsession forces engineers to cram massive arrays of multimodal sensors, heavy computing chips, and highly complex kinematic joints into a single bipedal shell.

From the perspectives of economics and systems engineering, this creates a severe scalability paradox. High physical complexity not only pushes the thermodynamic and physical limits of edge computing but also keeps unit manufacturing and operational costs unsustainably high. In real-world industrial, logistics, and domestic service scenarios, it is extremely difficult for complex humanoid robots to achieve a favorable Return on Investment (ROI) compared to traditional human labor. We must return to first principles: the fundamental purpose of a robot is to solve specific tasks in the physical world as efficiently and economically as possible, rather than mimicking the biological form of its creators. Based on this logic, this paper introduces the Lin-Shen Minimalist Model of Embodiment (LS-MME), aiming to redefine hardware design, computing power allocation, and the ecological boundaries between robots and human society.

2 Decoupling Principles and Formalization of the LS-MME

The core of the LS-MME framework is to shatter the myth of the ″all-in-one″ machine by thoroughly decoupling perception, computation, and physical movement at the sys- tem level. To illustrate this theory more intuitively, we define an embodied AI system as a ternary dynamic system consisting of structural complexity (Sc), local comput- ing power (Cl), and environmental intelligence (Ei). System Optimization Objective: $\min(\text{Costonboard}) \propto f(\text{Sc}, \text{Cl}, \text{Plocal})$

The LS-MME applies strict dimensionality reduction to this system through three core principles:

2.1 Minimum Viable Morphology (MVM): $\min(\text{Sc}) \rightarrow \text{Smin}$

The MVM principle emphasizes task-based reverse engineering: as long as the task success rate is met, the robot's physical structural complexity (Sc) must be reduced to the absolute lower bound permitted by physics and topology (Smin). Recent frontier research clearly reveals the ″Morphology-Control Trade-Of″ in robotics, proving that forcibly adding active joints triggers a non-linear explosion in con- trol dimensions. MVM dictates that simple tasks should be handled by the most minimalist end-efectors. Stripping away redundant degrees of freedom (DoF) that consume high energy and computing power is a physical prerequisite for large-scale commercialization.

2.2 Locally Shared-Perception (LSP): $\max(\text{Ei}) \Longrightarrow \text{Plocal} \rightarrow 0$

LSP argues that high-precision sensors (Plocal) should be stripped from the robot's body and replaced by environmental infrastructure (Ei). This ″perception offload- ing″ already has a solid empirical foundation in Cloud

* **Corresponding Author:** Shen Lin, shenlinbc@gmail.com

Robotics. For example, recent agricultural robotics research successfully deployed a two-phase cloud architecture that processes high-definition images via overhead cameras to guide lightweight end- efectors for harvesting, significantly lowering per-unit costs [1]. Furthermore, the latest research on 6G edge networks has successfully demonstrated seamless perception offloading for autonomous mobile agents [2]. LSP refines this into a universal paradigm: turning smart factories or buildings into the "eyes" themselves, constructing a global 3D environment, and allowing massive numbers of cheap, "blind" efectors to achieve precise operations by querying the central system.

2.3 Minimum Compute Requirement (MCR) and Graceful Degradation: min(Cl) → Csafe

Conventional wisdom suggests that to cope with network latency (e.g., 50 to 200 ms), robots must carry massive local computing power (Cl) for real-time planning. The MCR principle solves this pain point by introducing the mechanism of Graceful Degradation and Compliant Braking. Under MCR, high-level semantic and long- horizon planning computations are completely offloaded to the cloud. The robot's body only retains an ultra-low-power base control chip (Csafe). The sole core task of this chip is: upon encountering network jitter or unexpected obstacles, to execute microsecond-level energy dissipation and smooth deceleration based on classical New- tonian mechanics, entering a safe pause state to await new instructions from the cloud. This brief safety pause, mimicking biological "neural reflex latency," trades a tiny time compromise for the complete removal of expensive computing hardware.

3 Macro Techno-Economic Model and Its Theoretical Boundaries

Modern robotics relies on extremely complex dynamic modeling (such as finite element deformation spaces, Lie groups, etc.) to achieve local control.[1] However, to quantify the evolutionary trends driven by the LS-MME from a macro perspective, we combine its three principles to propose a heuristic scaling equation describing the collapse of physical redundancy in embodied AI.

3.1 The Techno-Economic Heuristic Equation for Embodied AI

LS-MME proposes the following macro equation to represent the dynamic inter- play among task requirements, robotic equipment complexity, and environmental intelligence:

$$S_c \times C_l = \frac{k}{E_i^{\lambda}}$$

Constraints (Subject to): Sc ⩾ Smin , Cl ⩾ Csafe

Logical Derivation: The numerator k represents the inherent complexity of the task. The denominator Ei represents environmental shared intelligence (maximized by the LSP principle, where λ is the network infrastructure scaling factor). The left side of the equation represents the cost drivers of a single robot: physical structural complexity Sc (constrained by MVM) and local computing load Cl (constrained by MCR).

This equation intuitively reveals a "Moore's Law" for embodied AI: as Ei (e.g., 5G/6G networks, edge servers, global sensors) grows exponentially, the right side of the equation collapses rapidly. To maintain economic equilibrium, the redundant complexity and expensive chip payloads of standalone robots exceeding the Smin and Csafe boundaries must inevitably drop at a drastic rate.

3.2 Dynamic Synergy in Application Boundaries and Unstructured Resistance

Any rigorous scientific framework must define its limitations. The economic premise of the LS-MME relies on the initial capital expenditure (CAPEX) of environmental intel- ligence (Ei). In newly built smart factories, modern hospitals, or warehouses, the Total Cost of Ownership (TCO) for retrofitting the environment is amortized across numer- ous minimalist robots, ofering an overwhelming commercial advantage. However, in highly open, unstructured wild environments (such as disaster rescue zones) or old buildings lacking retrofit budgets, the sunk cost of fixed environmental modifications will far exceed the benefits. Yet, this does not mean the LS-MME principles fail here. In these edge scenarios, the system can rapidly establish a dynamic network by deploying "Ad-hoc Perception Infrastructure"—such as utilizing satellite imagery, Unmanned Aerial Vehicles (UAVs) equipped with high-precision visual sensors, or other embodied devices acting as mobile perception nodes to provide global visual guidance for disaster recovery [3]. In scenarios with such dynamic scheduling and cooperative capabilities, the macro architecture of LS-MME can still operate efficiently. Therefore, highly independent and autonomous complex humanoid or quadruped robots are merely an "Optimal Option" in these environments, rather than the irreplaceable "Mandatory Requirement."

4 Advancing Frontier Challenges: Difusing the HEAT Complexity Disaster

The academic community is currently striving to conquer the "Heterogeneous Embod- ied Agent Training (HEAT)" challenge—using a unified foundation model to control robots of varying morphologies. Authoritative theory proves that cross-morphology policy training can be formalized as a structured Partially Observable Markov Decision Process (POMDP), with a computational complexity reaching the PSPACE-complete level [4].

Some might argue that dismantling a complex robot into a "cloud brain + mas- sive simple terminals" mathematically transforms the problem into a decentralized multi-agent POMDP (Dec-POMDP), which theoretically belongs to the even

harder NEXP-complete class. [4] However, this purely ignores the true physical engineering context: if we need to coordinately schedule thousands of humanoid robots, each with dozens of degrees of freedom in a factory, their joint multi-agent state space would result in a catastrophic dimensional explosion. In contrast, scheduling minimal- ist blind devices down-dimensioned by LS-MME involves a vastly smaller state space, and real-world tasks are often spatially decoupled. Therefore, the MVM principle of LS-MME essentially provides a highly commercially viable ″dimensionality reduction springboard″ for tackling macro-scale HEAT challenges: first using large models to solve the generalized control of simple tools, and then gradually breaking down high- dimensional coordination as the system evolves. Furthermore, combined with the latest software-layer dynamic resource allocation algorithms like ConfigBot [5], the system can maintain ultimate scheduling stability in edge-constrained environments.

5 Economic Application Value and the Dissolution of "Humanoid Ethics"

On an economic and application level, minimalist execution units adhering to LS-MME principles eliminate expensive dynamic bipedal balancing systems and onboard GPUs, easily compressing their manufacturing costs into the low-value durable goods range. Meanwhile, reduced mechanical complexity translates to a sharp drop in downtime. This extremely low CAPEX and high operational stability will drive unprece- dented application value and investment returns for embodied AI in manufacturing, warehousing logistics, and automated agriculture. Furthermore, as humanoid robots integrate into society on a large scale, they easily fall into the ″Humanoid Ethics Trap″: triggering the Uncanny Valley efect, exac- erbating fears of human labor replacement, and sparking endless debates over the moral status of machines [6]. The LS-MME provides a sociological antidote. Through mandatory ″De-personification,″ the physical representation of robots is degraded to ″execution sliders″ or ″simple chassis.″ This paradigm of ″the environment as intel- ligence, the machine as a tool″ completely strips away the anthropomorphic halo of robots, drastically lowering the psychological defenses of human society and clearing away non-technical friction for large-scale deployment.

6 Concluding Predictions: The Rise of the Distributed Macro-Organism

During the current market testing phase, humanoid robots are highly sought after for their ″plug-and-play″ features, primarily because existing physical infrastructure (like stairs and tools) is tailor-made for humans. However, from the perspective of first prin- ciples and long-term technological evolution, this is merely a temporary compromise based on historical path dependency. Based on the deductions of the LS-MME framework, this paper makes a disruptive prediction for the future: the ultimate commercial form of embodied AI will, in the vast majority of cases, no longer be isolated humanoid individuals that enclose perception, computation, and execution entirely within a single physical shell.

With the reconstruction of underlying logic, AI will evolve into an omnipresent ″distributed macro-organism.″ In this grand vision, the sensors of fixed buildings or mobile fleets serve as its optic nerves; the remote cloud data center acts as its unified superbrain; and the countless structurally minimalist, low-cost terminal devices shut- tling through factory assembly lines, smart homes, and even wild environments are merely the muscle fibers mobilized on demand by this network. Thoroughly decoupling cognition from morphology is not only an engineering necessity to break through individual computing bottlenecks but also the safest, most economical, and most ethically sound evolutionary endgame as humanity transitions into a highly automated society.

7 Declarations

Declaration of Generative AI and AI-assisted technologies in the writing process During the preparation of this work, the author used Google Gemini (version: Gemini 1.5 Pro) to assist in translating the original manuscript from Chinese into plain English and to improve the overall readability and language coherence. After using this tool/service, the author thoroughly reviewed, verified, and edited all the generated texts, and takes full responsibility for the ultimate content and academic integrity of this publication. The AI tool was strictly limited to language translation and refinement; it was not used to generate data, perform theoretical deductions, or create any of the underlying scientific models.

About the Author

Shen Lin, Han ethnicity, born on November 23,1992, from Henan Province, holds a master's degree in AI, robotics, and embodied, intelligence.

References

[1] Baig, H. L. Y. D.. D. M. M. A. A cloud robotics architecture for greenhouse straw- berry harvesting with single-pass data collection. Industrial Robot: the interna- tional journal of robotics research and application (2026).

[2] Mostardinha, G. J. V. P. R. P. R. D. S. S. M. C. M. M.. P. F. C. T. Edge computing and 5g network integration for mobility-aware service deployments. Pervasive and Mobile Computing (2026).

[3] C., L., Ay.n. How morphological computation shapes integrated information in embodied agents. PMC (2021).

[4] M¨uller.V.C., Hofmann.M. What is morphological computation? on how the body contributes to cognition and control. Artificial Life (2017).

[5] Dasanayake, B.V.D.G.B.L. P. S. A review of mathematical models in robotics. Applied Sciences (2025).

[6] Dwivedula, M. S. A. A. B. J. K. D. R. C. J. R. Configbot: Adaptive resource allocation for robot applications in dynamic environments. IEEE/RSJ International Confer- ence on Intelligent Robots and Systems (IROS) (2025).