

# Animals Beat ChatGPT at Physics: How Ecological Physics informs Embodied AGI

Leonardo Fabbri

February 2026

## Abstract

Artificial Intelligence research is currently bifurcated between the manipulation of linguistic symbols and the mastery of physical states. While Large Language Models (LLMs) excel at statistical token prediction, achieving Artificial General Intelligence (AGI) requires robust internal simulations of the material world. This paper builds upon our previous work regarding infant physical intuition, extending the inquiry into comparative animal cognition. We analyze the “Predictive Engine” of aquatic predators and the dragonfly’s interceptive steering as benchmarks for Physical AI. Finally, we discuss the transition toward Joint Embedding Predictive Architectures (JEPA) and the economic mandate of the embodied AI frontier.

## 1 Introduction

The fundamental limitation of current generative architectures is their detachment from physical causality. Predicting a linguistic token is a task fundamentally distinct from predicting the “next state” of a dynamic environment. As global capital pivots toward Physical AI, the biological world provides a rigorous benchmark. Evolution has prioritized the ability to model environment dynamics—gravity, inertia, and causality—as a prerequisite for survival, long before the emergence of symbolic language.

## 2 The Biological Mandate: Internal Simulation

Intelligence in biological entities emerged not from language, but from the necessity to navigate a three-dimensional medium. Internal simulation is a prerequisite for existence, allowing agents to predict the consequences of their actions and the trajectories of external agents.

This paper builds and expands upon our previous research, *Babies Beat ChatGPT at Physics* (Fabbri, 2025), which established that human infants possess an innate “intuitive physics” engine that outperforms the statistical pattern-matching of current LLMs. We here extend this thesis to the domain of comparative cognition, arguing

that the “World Model” is a universal biological solution to environmental uncertainty across species.

## 3 Ecological Physics: Case Studies in Interception

The verification of world models in biological systems is most evident in high-stakes predatory behavior, where agents must nullify sensory latency through predictive projection. We examine two landmark studies that demonstrate how primitive nervous systems implement sophisticated “Ecological Physics.”

### 3.1 The Dragonfly: Forward and Inverse Models

The research by Mischiati et al. (2014) on the dragonfly (*Plathemis lydia*) provides definitive evidence of internal models in invertebrates. The study utilized high-speed motion capture with reflective markers on the insect’s head, body, and wings to reconstruct 3D trajectories during prey capture.

The dragonfly faces a critical 50ms neural delay between visual stimulus and motor response. To resolve this, it does not merely “react” to prey position. Instead, the brain employs a **forward model** to predict the sensory consequences of its own maneuvers and an **inverse model** to calculate the motor commands required to align its body with the prey’s future flight path.

Crucially, the dragonfly’s head movements are predictive; the head “locks onto” the target independently of the body. The resulting head-body angle then serves as an error signal that guides the systematic rotation of the body. This hierarchical control allows the agent to maintain the prey image on the high-resolution fovea even during maneuvers exceeding  $40.28 \text{ m/s}^2$ .

### 3.2 The Zebrafish: Stochastic 3D Prediction Engines

While the dragonfly relies on continuous flight models, the larval zebrafish (*Danio rerio*) demonstrates physical intelligence in a discrete, bout-based environment. Bolton

et al. (2019) revealed what they termed a “stochastic 3D prediction engine.”

The study found that the zebrafish does not just orient toward the current position of the prey; it combines position and velocity perception to estimate a **preferred future positional estimate**. The zebrafish employs a recursive strategy: in each swim bout, it calculates the future location and systematically halves the angle of attack.

A significant discovery in this study was the role of “graded variance.” The zebrafish’s movements are not deterministic but stochastic, with the variance around the mean result scaling according to the distance from the target. This probabilistic approach allows the agent to optimize its hunting trajectory in a complex 3D medium, confirming that even the most “primitive” vertebrates utilize an implicit, adaptive world model.

## 4 Cognitive Foundations: Latent Learning

The shift toward world models in AI revisits Edward Tolman’s mid-20th-century work on “latent learning.” Tolman demonstrated that rats create “cognitive maps”—internal representations of their environment—even in the absence of immediate reinforcement. For modern AGI, this implies that “pre-training” on the structure of reality is more efficient than training on specific task rewards.

## 5 The LeCun Paradigm: JEPA vs. LLMs

Yann LeCun characterizes autoregressive models as a “dead end” for superintelligence. He proposes the Joint Embedding Predictive Architecture (JEPA), which focuses on predicting latent representations of the world rather than raw pixels or tokens. By predicting in a shared latent space, JEPA models can ignore irrelevant noise and focus on high-level causal structures, mirroring the efficient “Inattentive Blindness” seen in human and animal cognition.

## 6 The Rise of Physical AI: Economic Mandate

The transition from Digital to Physical AI represents a massive shift in capital allocation. The Global Embodied AI Market is projected to grow from \$2.5 billion in 2024 to \$10.75 billion by 2034. As shown in Table 1, the period between 2020 and 2026 has seen a radical transformation in the unit economics of humanoid platforms.

The value of Physical AI is increasingly tied to “Software-as-the-Soul.” Humanoid robotics require a uni-

Table 1: Comparative Humanoid Metrics (2020 vs. 2026)

Metric	2020 (Prototype)	2026 (Production)
Battery Runtime	1.5 – 2 Hours	8 – 12 Hours (Full Shift)
Inference Latency	100+ ms	18 – 20 ms
Unit Cost (Est.)	\$250k – \$1M	\$15k – \$30k
Hand Dexterity	5 – 11 DoF	22 DoF (Gen 3)
Torque Density	Moderate	High (Solid-State Gears)

versal foundation model that understands the rigid body dynamics of the real world. This “Physical Soul”—the internal model of physics—is becoming the most valuable asset in the global economy.

## 7 Conclusion

The analysis of biological intelligence and current AI trajectories confirms that the “Chatbot Era” is a precursor to an era of “Physical Intelligence.” The ability to simulate and navigate the material world is the true frontier of AGI. Those who prioritize the mastery of environment laws over linguistic syntax will be the architects of the next generation of technological wealth.

## References

- [1] Fabbri, L. (2025). *Babies Beat ChatGPT at Physics: Evaluating World Models in Early Developmental Stages*.
- [2] Mischiati, M., et al. (2014). Internal models direct dragonfly interception steering. *Nature*.
- [3] Bolton, A. D., et al. (2019). Elements of a stochastic 3D prediction engine in larval zebrafish. *Neuron*.
- [4] Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*.
- [5] LeCun, Y. (2024). *A Path Towards Autonomous Machine Intelligence*.