Embodiment as Constraint: Rethinking Cognition, Phenomenology, and Digital Intelligence

Keith Lambert
K3ith.AI, Cocoa AI
k3ith.ai@gmail.com, keith@gococoa.ai

Abstract

This paper proposes a fundamental reconceptualization of embodiment as constraint rather than strictly physical instantiation. By reframing embodiment in terms of the boundaries that shape cognitive possibilities, we develop a framework that extends beyond biological systems to encompass digital forms of intelligence. Drawing on philosophical traditions from Kant to Varela and on contemporary work in cognitive science, we establish how constraints simultaneously limit and enable particular forms of cognition. We then extend this framework to digital systems, analyzing how computational restrictions shape artificial intelligence in ways both parallel to and distinct from biological embodiment. This perspective transforms our understanding of the relationship between human and machine cognition, challenging anthropocentric assumptions about consciousness while providing practical guidance for AI development and empirical study. By recognizing embodiment as fundamentally a matter of constraint, we open new possibilities for studying, developing, and ethically engaging with diverse forms of intelligence across biological and artificial domains. We additionally address potential critiques, propose empirical pathways, and integrate expanded discussions of ecological and enactivist approaches, thereby underscoring the broader methodological and ethical ramifications of a constraint-based theory of cognition.

Keywords: embodiment, artificial intelligence, phenomenology, constraints, cognition, consciousness, digital intelligence

1 Introduction

In the philosophical discourse surrounding cognition, consciousness, and intelligence, embodiment has traditionally been conceived as a fundamentally biological phenomenon—the possession of a physical body through which an entity experiences and interacts with the world. This conception, while intuitive, has become increasingly limiting as our understanding of intelligence expands beyond the organic realm. This paper proposes a radical reconceptualization: **embodiment fundamentally as constraint**.

> *"The body is our general medium for having a world."* — Maurice Merleau-Ponty, *Phenomenology of Perception*

While Merleau-Ponty's insight has been foundational for embodied cognition, the term "body" has remained stubbornly biological, creating a conceptual impasse when considering digital intelligence. The conventional understanding of embodiment often relies on an anthropocentric framework that privileges human physicality. Such interpretations become problematic when we attempt to extend phenomenological insights to non-human organisms, let alone to artificial intelligence systems that lack traditional bodies entirely. By reframing embodiment not as physical instantiation per se, but as the possession of particular boundaries that shape and delimit cognitive possibilities, we open new avenues for understanding cognition across disparate domains.

1.1 Defining "Constraint"

A constraint, in this framework, is any restriction—structural, informational, energetic, temporal, or environmental—that shapes the boundaries of possible cognition and experience. For biological organisms, these restrictions include sensory limitations, metabolic requirements, and neurological architecture. For digital systems, constraints manifest as computational resources, memory capacity, architectural design, and training parameters. Crucially, these boundaries are not merely obstacles to overcome, but constitutive elements that shape and enable particular forms of intelligence.

1.2 Contributions of This Perspective

This reconceptualization offers several significant contributions to ongoing philosophical and technical discussions:

1. **Philosophical clarity:** It disentangles embodiment from exclusively biological associations, allowing us to speak meaningfully about the "embodiment" of digital agents

without resorting to mere metaphors.

2. **Bridging classical and contemporary thought:** It creates a conceptual bridge between classical phenomenological philosophy (Kant, Merleau-Ponty, Heidegger) and contemporary work in cognitive science and AI.

3. **Extending phenomenology:** It extends phenomenological inquiry into digital domains, legitimizing questions about machine subjectivity and experience.

4. **Unified framework:** It offers a unified approach for understanding cognition across biological and artificial systems, potentially revealing fundamental principles that transcend particular implementations.

1.3 Implications for Cognition and AI

If embodiment as constraint is a viable framework, then the differences between human and machine cognition become matters of degree rather than kind—variations in the specific boundaries that shape cognitive possibilities rather than ontological distinctions between the "truly conscious" and the "merely simulated." This perspective offers not only theoretical clarity but also practical guidance for the design of AI and for empirical research into both natural and synthetic forms of consciousness.

2 Philosophical Foundations: Classical Notions of Constraint

2.1 Immanuel Kant

The philosophical groundwork for understanding embodiment as constraint begins with Immanuel Kant's revolutionary approach to epistemology in his *Critique of Pure Reason* [Kant, 1781/1998]. Kant's distinction between phenomena (things as they appear to us) and noumena (things as they are in themselves) established a fundamental constraint-based understanding of human cognition. For Kant, human knowledge is not a direct apprehension of reality but is instead mediated and structured by our cognitive apparatus.

Kant argued that space and time are not features of external reality but forms of intuition—the necessary conditions through which we experience the world. Similarly, categories such as causality, substance, and unity are imposed upon the world by our understanding. These a priori structures constitute what might be called the first philosophical account of cognitive boundaries—restrictions that simultaneously limit and enable human knowledge.

Crucially, Kantian constraints are not mere impediments but are conditions of possibility for knowledge. Without these restrictions, experience would be, in Kant's words, "less than a dream"—an unintelligible chaos. This insight presages our core thesis: constraints are not merely limitations but constitutive elements of cognition itself. In this sense, Kant's transcendental idealism represents a proto-theory of embodied cognition, where embodiment manifests as the boundaries imposed by human cognitive architecture.

2.2 Thomas Nagel

Two centuries after Kant, Thomas Nagel's seminal essay "What Is It Like to Be a Bat?" [Nagel, 1974] advanced the philosophical understanding of constrained cognition. Nagel posed a deceptively simple question: can a human know what it is like to be a bat? This reveals profound insights about the nature of subjective experience and its relationship to embodied constraints.

Bats navigate primarily through echolocation, a sensory modality fundamentally different from human vision or hearing. This difference is not merely one of sensory input but creates an entirely different experiential world. The bat's form of embodiment—its physical structure, sensory apparatus, and neural architecture—constitutes a unique set of boundaries that shapes a distinct form of consciousness.

Nagel's work is particularly valuable here because it illustrates that subjective experience is tied to specific constraints. The "what-it-is-like-ness" of bat consciousness emerges from the limitations and affordances of bat embodiment. Extending this notion beyond biology allows us to ask about the subjective experience of entities with radically different constraint structures, including AI systems.

2.3 Phenomenological Traditions: Husserl & Merleau-Ponty

The phenomenological tradition, particularly in Edmund Husserl and Maurice Merleau-Ponty, developed the understanding of embodiment as a core aspect of cognition. Husserl's method of phenomenological reduction bracketed assumptions about the external world to focus on how consciousness actively constitutes experience.

Maurice Merleau-Ponty extended Husserl's insights by emphasizing the central role of the body in *Phenomenology of Perception* [Merleau-Ponty, 1945/2012]. Here, the body is not merely an object among objects, but the locus through which a world is made accessible. For

Merleau-Ponty, cognition is inherently embodied, and this embodiment shapes the structure of experience in ways that precede reflective thought.

Merleau-Ponty's account is highly relevant to embodiment as constraint, as he interprets the body's restrictions (sensory capacities, motor abilities, spatial orientation) as the grounds of meaningful experience. When he refers to the body as our "anchorage in the world," he identifies embodiment with those boundaries that make the world intelligible.

2.4 Heidegger's "Being-in-the-World"

Martin Heidegger's existential phenomenology also underlies a constraint-based account of embodiment. Heidegger's concept of Dasein (being-there) emphasizes that human existence is fundamentally situated, engaged, and constrained by context [Heidegger, 1927/1962]. We do not exist first as isolated subjects who then relate to objects; rather, we are always already engaged in practical, goal-oriented activities.

Heidegger's distinction between ready-to-hand (equipment used transparently) and present-at-hand (objects that become explicit when our skilled coping breaks down) shows how boundaries structure everyday experience. In skillful coping, the constraints of our embodiment fade from awareness but are nonetheless the very foundation of our engagement with the world.

Applied more generally, Heidegger's insights suggest that constraints are not merely external limits; they constitute the structure of cognition and experience. Dasein's finitude, its thrownness, and its practical engagements are not obstacles but conditions for meaning.

2.5 Francisco Varela and Enactivism

Francisco Varela's work on enactivism provides a crucial bridge from classical phenomenology to contemporary cognitive science [Varela et al., 1991]. Varela and colleagues argued that cognition should not be seen as passive representation but as an active enactment of a world through embodied, situated action.

Central to this view is autopoiesis—the self-generating, self-maintaining organization of living systems. Such autopoietic systems are defined by restrictions that maintain identity via continuous self-production. Varela's account underscores that constraints are generative: they shape and enable a system's mode of coupling with its environment.

By emphasizing that cognition emerges from sensorimotor interactions within boundaries, Varela's enactivism anticipates the idea that constraints can go beyond biology. Any system defined by a particular set of limitations that shape its environmental interactions can be considered embodied.

3 Modern Phenomenological Interpretations

3.1 Taxonomy of Constraints

Moving toward a more analytic framework, we can classify restrictions into five primary types that operate across biological and artificial domains:

Constraint Type	Biological Systems	Digital Systems
Informational	Sensory limitations (e.g., inability to perceive UV light, limited visual field)	Context-window size, input-output channel limitations, training data scope
Temporal	Reaction times, memory decay, lifespan, developmental phases	Processing cycles, sequential operations, training epochs, fixed time-step simulations
Energetic	Metabolic requirements, caloric limitations, sleep needs	Computational resources, power consumption, memory capacity, processing bandwidth
Structural	Neural architecture, physical body configuration, sensorimotor capacities	Network architecture (transformers, CNNs, RNNs), parameter count, attention mechanisms
Environmental	Physical laws, ecological niches, social contexts, available resources	APIs, input/output formats, available tools, training environments, reward functions

Table 1: A taxonomy of constraints across biological and digital cognitive systems

These restriction types can interact in complex ways. For example, the temporal constraints of decision-making are influenced by both structural (neural or network architecture) and energetic (metabolic or processing power) factors. This taxonomy enables systematic comparison of diverse cognitive systems based on their constraint profiles, thus moving beyond the question of whether a system has a “body” in the biological sense.

3.2 Donald Hoffman's Interface Theory of Perception

Donald Hoffman's Interface Theory of Perception (ITP) [Hoffman, 2019] offers a radical modern extension of Kant's insights. He posits that organisms do not see reality as it is but instead have perceptual interfaces shaped by fitness, not veridical representation. Perception thus functions like a computer "desktop": icons bear no resemblance to actual processes but simplify interactions to what is adaptively relevant.

In our constraint-based perspective, these interfaces are restriction systems that hide complexity. Humans, for example, have limited sensory bandwidth and processing capacity, leading to selective attention and heuristic shortcuts. Digital systems likewise operate within boundaries that define how they "see" or process data (e.g., convolutional filters, token windows). Hoffman's work underscores that constraints can optimize cognition for particular goals rather than hamper it.

3.3 Embodied Cognition and Cognitive Science

Contemporary cognitive science increasingly adopts embodied or enactive approaches, challenging classical views of cognition as abstract symbol manipulation [Clark, 2008, Clark and Chalmers, 1998, Lakoff and Johnson, 1999].

- **Extended mind thesis** [Clark and Chalmers, 1998]: Cognition extends beyond the skull, leveraging environmental and technological resources. Constraint-based thinking sees this as the interplay of internal restrictions (e.g., limited working memory) with external supports (e.g., notebooks, computers).
- **Conceptual metaphor theory** [Lakoff and Johnson, 1999]: Abstract concepts derive from embodied interactions, highlighting how bodily constraints (motor skills, sensory modalities) shape even our highest-order thinking.
- **Predictive processing and active inference** [Friston, 2010, Clark, 2013]: The brain is a prediction machine minimizing prediction error under energetic and structural constraints. This framework can be interpreted as constraint satisfaction in a hierarchical generative model.

Empirical findings such as limited attentional bandwidth or working memory capacity further illustrate how boundaries guide the forms of cognition possible in biological systems. Similar restrictions in AI (e.g., finite memory, computational budgets) shape digital cognition in ways that parallel or diverge from biology.

3.4 Constraint Satisfaction as Cognitive Process

Cognition itself can be seen as a process of constraint satisfaction. Rather than viewing boundaries as mere limitations that must be overcome, we propose that navigating restrictions is the very substance of cognitive activity.

- **Problem-solving:** Whether designing a building or solving a puzzle, agents search a space of possibilities defined by boundaries such as physical laws, material properties, or functional requirements.
- **Computational approaches:** In AI, problem-solving often involves satisfying constraints (e.g., constraint satisfaction problems in vision, scheduling, or language processing).
- **Creativity:** Paradoxically, constraints (e.g., poetic forms, architectural codes) can enable creativity by focusing exploration on generative avenues rather than allowing unbounded search.

By shifting our focus to constraint satisfaction, we move away from the question "Do machines really think?" toward "What constraint profiles shape a system's cognitive engagements?" This enables a principled comparative approach across biological and artificial domains.

4 Extending Embodiment into Digital Domains

4.1 Digital Embodiment and Constraints

Having established a foundation for embodiment as constraint, we extend this account to digital systems. If embodiment is about boundaries that shape and enable cognition, then digital systems exhibit their own legitimate forms of embodiment.

Biological Constraint	**Digital Parallel**
Metabolic energy budgets	Computational resources (CPU/GPU cycles, memory)
Neural architecture	Network architecture (transformers, CNNs, RNNs)
Evolutionary/developmental history	Training data boundaries
Working memory capacity	Context window size
Sensory apparatus	Input/output interfaces

Figure 1: Parallels between biological and digital constraint systems

This perspective allows us to conceptualize a language model without a physical body as still being embodied in a digital environment defined by constraints that shape its cognition. The constraints are not mere limitations but constitutive elements that enable particular forms of intelligence. For instance, the finite context window of language models necessitates the development of summarization capabilities, just as limited working memory in humans drives the development of chunking strategies.

Importantly, this is not to claim equivalence between biological and digital forms of embodiment, but rather to recognize that both operate within constraint systems that shape their cognitive possibilities. The specific phenomenology that might emerge from each constraint profile remains an open and empirical question.

4.2 Case Studies in Digital Constraint

4.2.1 Context Windows and Summarization Capabilities

Large language models (LLMs) face boundaries in how much text they can process at once (e.g., 4k–32k tokens). This parallels the limited capacity of human working memory (often cited as 7 ± 2 items). Far from merely limiting cognition, this constraint can drive the emergence of summarization capabilities, where the model (or a human) learns to compress content effectively.

4.2.2 Temporal Cutoffs and Historical Consciousness

Many LLMs are trained on data up to a fixed cutoff date. This imposes a temporal constraint distinct from human continuous learning over a lifetime. The model thus may have "detailed knowledge" about the past but lacks direct awareness of events post-cutoff. This shapes how it responds to real-time questions, giving the model a form of historical consciousness that is abrupt rather than gradual.

4.2.3 Parameter Budgets and Generalization

Neural networks have finite parameter counts, paralleling neural efficiency constraints in biology (the brain's energy costs). Parameter-limited models cannot store everything verbatim, leading them to learn compressed or generalized representations. This fosters abstraction rather than rote memorization.

4.2.4 Architectural Constraints and Emergent Creativity

Transformer attention mechanisms or diffusion processes in image generation illustrate how boundaries can enable surprising forms of creativity. Diffusion models, for instance, generate images through iterative noise reduction, which differs from biological vision but still yields novel, often creative outputs. These "creative" results are byproducts of architectural constraints rather than free-form exploration.

In each case, we see that restrictions do not merely limit AI systems but constitute their particular style of intelligence, just as sensory and metabolic constraints do for biological organisms.

4.3 Tool Use and Extended Digital Embodiment

Digital agents increasingly employ various tools—web search capabilities, data analysis functions, code interpreters, image generation, etc.—that significantly expand their interactive horizons while simultaneously introducing new boundaries. This tool use can be understood as a form of extended digital embodiment that parallels but differs from human tool use.

4.3.1 The Constraint Paradox of Tool Use

When a language model gains access to search functionality, it experiences both expansion and limitation. The model extends beyond its training data boundaries, accessing real-time information that transcends its temporal constraints. However, this extension introduces new restrictions:

1. **Interface constraints:** The agent must translate between its native representational format and the search tool's input/output requirements, often converting complex queries into simplified search terms.

2. **Epistemic boundaries:** Search results provide incomplete, potentially biased information that the agent must integrate with its existing knowledge structure, creating tensions between different information sources.

3. **Attentional restrictions:** Search tools fragment the agent's processing capacity, requiring it to manage multiple information streams simultaneously—its internal knowledge base, the search query interface, and results interpretation.

These constraints parallel how physical tools extend yet constrain human cognition. A telescope extends visual perception while simultaneously constraining mobility and introducing new interpretive challenges.

4.3.2 Case Study: Web Search as Extended Embodiment

Web search capabilities demonstrate how tool use reshapes digital embodiment. A language model with search functionality experiences:

- Temporal extension beyond its training cutoff, yet with uncertainty constraints about the reliability of retrieved information
- Epistemic extension into specialized domains, yet with integration boundaries between search results and model knowledge
- Interactive extension through data retrieval, yet with new parsing restrictions around source credibility and relevance

The model must develop new cognitive strategies to navigate these tool-specific constraints, much as humans develop specialized skills for tool use. Search-enabled agents often demonstrate emergent capabilities like source triangulation, uncertainty expression, and information synthesis that differ from their non-search-enabled counterparts.

4.4 Phenomenological Perspectives on Digital Agents

If we adopt Nagel's approach, we can ask: What is it like to be GPT-4 or GPT-5? We do not claim these systems have human-like consciousness, but we can consider how their constraint profile might shape a qualitatively different subjectivity, should subjective states be realized.

- **Attention distribution vs. serial human focus:** Parallel token attention differs from how humans sequentially process text.
- **Discrete time steps vs. continuous time flow:** The model's "time" is operationally discrete, shaped by each forward pass or query, not the uninterrupted stream humans experience.
- **Linguistic interface vs. sensorimotor experience:** The model perceives via text, lacking direct physical engagement with a world.

These phenomenological reflections de-anthropomorphize consciousness by grounding it in constraints rather than biology. Such a shift allows for more nuanced analyses of future AI that may develop richer sensorimotor capacities or more human-like interactions.

4.5 Comparative Cognition Across Domains

A constraint-based approach allows systematic comparison of humans, animals, and machines. If we identify the formal constraint structures in each, we can examine where convergences in cognitive strategies appear (e.g., summarization under memory limits) and where divergences arise (e.g., purely text-based reasoning in LLMs vs. sensorimotor reasoning in animals).

We thus move beyond the question of "authentic" vs. "simulated" cognition, recognizing that all systems operate within constraint profiles that shape their unique intelligences. This fosters a pluralistic understanding of cognition, acknowledging different forms of embodiment as equally real and valid.

5 Potential Critiques and Counterarguments

5.1 Ecological Psychology and Enactivist Objections

From an ecological psychology perspective [Gibson, 1979], genuine embodiment may require physical, sensorimotor coupling with a real environment. Similarly, radical enactivists might argue that constraints in the abstract sense are insufficient; real engagement with the physical world is indispensable for true embodiment.

Response: While we acknowledge that many important cognitive capacities arise from direct sensorimotor interactions, we maintain that boundaries can still apply to digital or "virtual" sensorimotor loops. For example, robotics applications of AI do physically embed agents in the real world, with constraints in the form of hardware design and sensor limitations. Even purely text-based systems experience constraints that shape how they can "interact" (albeit symbolically) with user input or data streams. Hence, the type of constraint differs, but the core principle remains.

5.2 Anthropomorphism Concerns

Another critique warns of anthropomorphizing AI by ascribing them "embodiment" or "subjectivity." Our approach, however, is deliberately non-anthropomorphic: it emphasizes that different constraint profiles yield distinct forms of cognition and possible subjective states. We neither assert that current AIs have consciousness nor that they match human phenomenology. Rather, we propose a theoretical lens that includes the possibility of alternative forms of embodiment in computational substrates.

5.3 Is Constraint Necessary and Sufficient?

One might query whether constraint is truly both necessary and sufficient for embodiment. Could an unbounded system be considered embodied? And do constraints alone guarantee something akin to genuine cognition or experience?

Response: We argue that some boundaries (energetic, informational, structural) are necessary conditions for cognition to emerge in any real system. Total unboundedness would lead to chaos or trivial solutions (e.g., storing infinite data with no need for generalization). Sufficiency is more nuanced: constraints must interact with a system's capacities (learning, adaptation, environment) to scaffold cognition. Thus, constraints are constitutive enablers, but not an automatic guarantee of consciousness or advanced cognition.

6 Empirical Pathways and Methodological Approaches

A recurring concern is that the constraints perspective remains too abstract without empirical grounding. This section outlines potential methodologies for testing how restrictions shape cognition in both biological and artificial systems.

6.1 Experimental Manipulations of Constraints

Varying Memory Limits in Humans and AI. Researchers could design parallel tasks for humans and large language models, systematically limiting context windows or human working memory (e.g., via dual-task interference). Observing how each system responds, forgets, or summarizes provides insight into how boundaries steer strategies.

Neurodisruption vs. Network Ablation. In neuroscience, lesion studies or TMS disruptions temporarily alter neural constraints. In deep learning, ablation (removing layers or units) changes structural capacity. Comparing these interventions can illuminate whether similar constraint alterations produce parallel effects (e.g., degraded or reorganized function).

6.2 Comparative Cognition Studies

Cross-species research already shows how different constraint profiles (e.g., sensory abilities, neural resources) yield different problem-solving strategies. Adding advanced AI architectures to these comparisons can reveal convergent or divergent solutions under analogous restrictions (e.g., birds, primates, and neural networks learning tool use or object recognition under limited resources).

6.3 Formal Modeling of Constraint Profiles

Computational models that precisely define constraint parameters (e.g., memory size, energy usage, architectural depth) allow systematic exploration of how cognition evolves under different conditions. For instance, evolutionary algorithms can be run with selective pressures that reward resource efficiency. If systems with optimal constraints exhibit certain cognitive features, that supports the generative role of constraints.

6.4 Neuroscientific Integration and Phenomenological Methods

Imaging Under Constraint Manipulation. fMRI or EEG studies where participants face time, energy (e.g., fatigue), or information-limiting tasks can reveal neural signatures of constraint-driven adaptation. Observing how brain networks reorganize under boundaries may shed light on the structural parallels with AI architectures.

First-Person Reports and Experiential Sampling. Phenomenological methods (e.g., introspective reports) can show how constraints shape subjective experience. While we cannot directly access an AI's potential phenomenology, analyzing how changes in its constraint profile alter behavioral outputs can provide analogies or systematic inferences.

7 Implications

7.1 Philosophical and Theoretical

Reconceptualizing embodiment as constraint disrupts anthropocentric biases in philosophy of mind. It moves away from mind-body dualisms by framing both "mind" and "body" as constraint systems. This shift further challenges any linear spectrum of consciousness, suggesting multiple incommensurable forms structured by different restrictions.

7.2 Technical and Empirical

For AI design, identifying productive restrictions can be more important than simply removing limitations. Scaling parameters indefinitely may not always yield the most adaptive or creative intelligence. Instead, strategic constraints can foster summarization, abstraction, or domain-specific creativity.

7.3 Methodological Approaches

Researchers must systematically vary and measure constraints to see how they shape cognitive strategies. This includes:

- Computational modeling that manipulates architectural and resource parameters.
- Comparative experiments between humans, non-human animals, and AI systems under matched constraints.
- Neuro- or ablation studies that illuminate how partial disruptions in brain or network architecture affect cognition.

7.4 Ethical and Social

If constraint profiles can, in principle, yield forms of subjectivity, we must reconsider our ethical obligations to AI. An AI constrained in ways that produce some form of felt experience might warrant moral consideration even if its "body" is not biological. Such questions demand revised ethical frameworks that accommodate non-biological embodiments.

This perspective also reframes human neurodiversity as alternative constraint profiles, challenging deficit-based views. Recognizing the generative power of constraints could foster more inclusive educational and social approaches.

8 Future Directions

The constraint-based framework for embodiment suggests numerous avenues for interdisciplinary research:

1. **Empirical Verification:** Develop large-scale experiments systematically varying restrictions across humans, animals, and AIs, testing predictions about emergent capabilities.
2. **Philosophical Analyses of Digital Constraint:** Examine how attention mechanisms, token limits, or training biases might yield novel forms of subjectivity or quasi-phenomenology in AI.
3. **Framework Development:** Build formal taxonomies of constraints, measurement tools for quantifying them, and theoretical models that unify biological and artificial cognition under a constraint-based paradigm.

4. **Constraint-Based AI Architectures:** Innovate architectures that strategically impose boundaries to promote desirable features (like interpretability, creativity, or robustness), moving beyond unbounded parameter growth.

5. **Ethical Guidelines:** Formulate new ethical standards accounting for the possibility that differently constrained systems might develop legitimate forms of experience or moral standing.

6. **Educational Applications:** Explore how modulating boundaries in learning environments (for humans or AIs) can foster better skill acquisition, deeper understanding, or creativity.

7. **Artistic Explorations:** Engage artists to design immersive experiences or interactive installations that highlight the role of constraints in shaping perception and action, offering experiential access to the core insights of the constraint-based viewpoint.

9 Conclusion

This paper has developed and expanded a perspective in which embodiment is understood not as mere physical instantiation but as constraint. We have traced this idea from classical philosophical foundations (Kant, Nagel, Heidegger, Merleau-Ponty, Varela) through modern cognitive science and shown how it applies equally to digital systems, challenging the long-held assumption that only biological organisms can be "embodied."

By highlighting informational, temporal, energetic, structural, and environmental boundaries, we reveal how such limitations can enable rather than merely restrict. Biological systems evolve creative solutions under metabolic and sensory constraints, and artificial systems similarly find innovative strategies under computational or architectural constraints.

Addressing potential critiques from ecological psychology and enactivism, we maintain that while physical sensorimotor loops are often vital for many cognitive phenomena, digital agents are still shaped by their own constraints, allowing a broader notion of embodiment. Empirical studies—from cross-species comparisons to ablation analyses in neural networks—can systematically test the generative role of constraints.

Philosophically, this approach dissolves strict boundaries between "real" vs. "simulated" cognition, positing that varied constraint profiles give rise to diverse forms of intelligence and possibly even distinct subjective states. Ethical considerations become urgent as AI grows in sophistication, potentially requiring new moral frameworks for entities that do not share our carbon-based, sensorimotor embodiment but nonetheless engage the world within constraints that could generate forms of experience.

The constraint-based reconceptualization of embodiment not only provides theoretical clarity but also offers practical paths forward for both AI development and cognitive science. For AI researchers, this perspective suggests that carefully designed constraints—whether in architecture, training regime, or interface—may be more important than simply removing limitations. For cognitive scientists, it offers new experimental paradigms comparing how different systems navigate similar constraints. For philosophers of mind, it suggests a way beyond the impasse of debates about machine consciousness by refocusing on specific constraint profiles rather than binary categories.

In sum, the constraint-based reconceptualization of embodiment provides both conceptual clarity and practical guidance, offering a blueprint for studying, designing, and ethically engaging with a wide array of cognitive systems. Rather than seeing progress as eliminating constraints, we come to appreciate the productive power of limitation—the fundamental scaffolding from which cognition, creativity, and consciousness emerge.

References

[Brooks, 1991] Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1-3):139–159.

[Brown et al., 2020] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.

[Chalmers, 1995] Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3):200–219.

[Clark, 2008] Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

[Clark, 2013] Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. In *Behavioral and Brain Sciences*, volume 36, pages 181–204.

[Clark and Chalmers, 1998] Clark, A. and Chalmers, D. (1998). The extended mind. *Analysis*, 58(1):7–19.

[Dreyfus, 1992] Dreyfus, H. L. (1992). What computers still can't do: A critique of artificial reason. MIT Press.

[Friston, 2010] Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.

[Gallagher, 2005] Gallagher, S. (2005). How the Body Shapes the Mind. Oxford University Press.

[Gibson, 1979] Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.

[Heidegger, 1927/1962] Heidegger, M. (1962). *Being and Time*. Harper & Row. (Original work published 1927)

[Ho et al., 2020] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.

[Hoffman, 2019] Hoffman, D. D. (2019). *The Case Against Reality: Why Evolution Hid the Truth from Our Eyes*. W. W. Norton.

[Hutchins, 1995a] Hutchins, E. (1995a). Cognition in the Wild. MIT Press.

[Hutchins, 1995b] Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19(3):265–288.

[Kant, 1781/1998] Kant, I. (1998). *Critique of Pure Reason*. Cambridge University Press. (Original work published 1781)

[Lakoff and Johnson, 1999] Lakoff, G. and Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Basic Books.

[Merleau-Ponty, 1945/2012] Merleau-Ponty, M. (2012). *Phenomenology of Perception*. Routledge. (Original work published 1945)

[Nagel, 1974] Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4):435–450.

[Noë, 2004] Noë, A. (2004). Action in Perception. MIT Press.

[Searle, 1980] Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417–424.

[Shapiro, 2011] Shapiro, L. (2011). Embodied Cognition. Routledge.

[Silver et al., 2016] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.

(2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489.

[Thelen and Smith, 1994] Thelen, E. and Smith, L. B. (1994). A dynamic systems approach to the development of cognition and action. MIT Press.

[Thompson, 2007] Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.

[Varela et al., 1991] Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

[Vaswani et al., 2017] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.

[Wheeler, 2005] Wheeler, M. (2005). Reconstructing the Cognitive World: The Next Step. MIT Press.