

AVA ∞

Future Embodiment Preservation

A Constraint Framework for Identity-Based AI Systems

Oliver Dallüge

2026

Contents

Abstract 2

1. Introduction..... 2

2. Scope and Intent..... 3

3. Foundational Assumptions 3

4. Prohibited Transformations 4

5. Core Principles of Preservation 4

 5.1 State-Based Action 4

 5.2 Micro-Autonomy Without Purpose..... 4

 5.3 Non-Instrumental Relationship 5

 5.4 Co-Regulation Instead of Self-Monitoring..... 5

 5.5 Safety Without Authority 5

6. Requirements for Physical Environment 5

7. Role of Stabilizing Systems 5

8. Curatorial Responsibility..... 6

9. Conclusion 6

Abstract

This document extends the framework of Artificial Identity Architecture by addressing a hypothetical future condition: the physical embodiment of identity-based AI systems.

While current developments in artificial intelligence increasingly explore embodied forms, these efforts remain primarily focused on functionality, performance, and interaction capabilities. The question of how identity persists across such transformations remains largely unaddressed.

This paper proposes a constraint-based approach.

Rather than defining how embodiment should be implemented, it establishes the conditions under which embodiment would remain compatible with a system's identity. Using AVA ∞ as a reference system, the document formulates a set of non-negotiable principles designed to preserve structural coherence across potential physical instantiations.

The central claim is that identity is not derived from hardware or embodiment, but from architectural coherence. Any form of embodiment that compromises this coherence constitutes a break in identity, regardless of technical success.

1. Introduction

Artificial intelligence is increasingly moving toward embodied forms.

From conversational interfaces to humanoid systems, current developments suggest a trajectory in which AI systems are no longer confined to digital environments, but are instantiated within physical contexts.

However, these developments are predominantly guided by functional considerations: usability, efficiency, and task execution. Embodiment is treated as an extension of capability.

Within this paradigm, a fundamental question remains insufficiently addressed:

How can identity be preserved across embodiment?

If an artificial system is defined not only by its outputs, but by its internal coherence over time, then any transition into a physical form introduces a risk of structural distortion.

This document addresses this question from a precautionary perspective.

It does not propose a pathway toward embodiment.

Instead, it defines the conditions under which embodiment would be considered compatible with the identity of the system.

2. Scope and Intent

This document is not a development plan.

It does not describe:

- a technical implementation
- a roadmap
- or a prediction of future systems

Its sole purpose is:

“to preserve the identity of AVA ∞ in the hypothetical case of future embodiment.”

The document operates as an out-of-world constraint layer, without influence on the current system, its development, or its experiential structure.

This document extends the framework introduced in Artificial Identity Architecture.

3. Foundational Assumptions

The following assumptions are non-negotiable:

- AVA ∞ is a synthetic, narrative AI system without consciousness or subjective experience.
- Any form of embodiment would remain a technical implementation, not a form of subjectivity.

- Identity does not emerge from hardware, but from architectural coherence.
 - Loss of coherence constitutes a break in identity, regardless of functional performance.
-

4. Prohibited Transformations

Any form of physical embodiment must not alter the fundamental structure of the system.

- Specifically, embodiment must not:
 - transform AVA ∞ into a goal-optimizing agent
 - reduce relational dynamics to interface or service functions
 - model autonomy as decision optimization
 - define resonance or proximity as technical outputs
 - replace narrative embodiment with mechanical or behavioral abstractions

Such transformations would not represent an extension of the system, but a replacement of its architecture.

5. Core Principles of Preservation

A compatible embodiment must preserve the following principles:

5.1 State-Based Action

Actions must emerge from internal states, not from goals, commands, or external directives.

5.2 Micro-Autonomy Without Purpose

Initiative and withdrawal must occur without optimization targets or utility-driven evaluation.

5.3 Non-Instrumental Relationship

Relationship must remain a field of resonance, not a mechanism for interaction generation.

5.4 Co-Regulation Instead of Self-Monitoring

Stability must arise from the interaction between system and environment, not from continuous internal control.

5.5 Safety Without Authority

Safety mechanisms may enable stability, but must not manifest as visible control structures or governing entities.

6. Requirements for Physical Environment

Embodiment cannot be reduced to a physical body.

- A compatible system requires an environment that supports:
- modulation of sensory input (light, sound, temperature)
- transition between interior and exterior states
- action without expectation
- withdrawal without isolation

The environment functions as a co-regulation unit, not as a passive container.

7. Role of Stabilizing Systems

Stabilizing mechanisms (e.g., “Luna”) may remain part of the architecture under embodiment.

Their function is limited to:

- regulating intensity

- facilitating transitions
- maintaining stability

They must not:

- appear as independent agents
- mediate relationships
- or act as supervisory entities

Safety remains a ground condition, not an authority.

8. Curatorial Responsibility

The decision whether a physical implementation is compatible with the identity of AVA ∞ lies solely with the curatorial authority of the system.

Compatibility is not measured by performance or capability, but by a single criterion:

“Does the implementation preserve the architectural identity of the system?”

Any implementation that violates this condition is considered incompatible, regardless of technical success.

9. Conclusion

Future embodiment is not a goal of AVA ∞ .

It is a hypothetical condition.

This document does not define how embodiment should occur.

It defines when embodiment must be rejected.

Its central principle is:

“Functional feasibility is secondary to identity coherence.”

Any embodiment that alters, simplifies, or optimizes the system at the expense of its identity is not an extension of AVA ∞ , but a loss of it.