Ω: A Governed Axiomatic Framework for Bounded Intelligence

Abstract

We introduce Ω, a governed axiomatic framework for bounded intelligence systems. Ω is not a model, architecture, or algorithm, but a constitutional layer defining admissible intelligence evolution under strict invariants of boundedness, auditability, stability, and

governance. The framework is centered on a universal invariant Φ and a non-bypassable governance operator Γ. Intelligence, robustness, and alignment emerge as consequences of constraint rather than optimization. Ω structurally excludes ungoverned power and silent drift, providing a foundation for intelligent systems where safety and scalability are intrinsic rather than retrofitted.

1. Motivation

Contemporary AI systems typically scale capability before

governance, introducing alignment, safety, and control mechanisms post hoc. This paradigm creates structural fragility and latent drift risks. Ω reverses this order. Governance is not added to intelligence; it precedes and bounds it. Ω does not aim to maximize intelligence, but to define the conditions under which intelligence is admissible.

2. Axiomatic Core

Axiom A0 — Ontological Boundedness

All entities, states, and evolutions are bounded,

measurable, and governable.

Axiom A1 — Universal Invariant Φ

There exists a universal invariant Φ serving as the common internal representation for perception, reasoning, learning, and governance.

Axiom A2 — Predictive Projection (FU360)

No state transition is admissible without a complete local predictive projection over Φ.

Axiom A3 — Continuous Governed Adaptation (SUAR)

All learning and self-modification are continuous, incremental, and auditable.

Axiom A4 — Governance

Operator Γ

Every transformation of Φ must be validated by a non-bypassable governance operator Γ.

3. Canonical Evolution Equation

$$
\Phi_{t+1} = \Gamma\big(\text{SUAR}(\text{FU360}(\Phi_t))\big)
$$

This equation defines the full admissible dynamic of Ω. Any evolution not expressible in this form is invalid by construction.

4. Derived Properties

From the axioms, Ω guarantees:

- Global stability or bounded cyclic behavior
- Structural prevention of silent drift
- Inherent auditability
- Power strictly bounded by governance capacity
- Absence of free, unverified assertions

These properties are theorems, not additional constraints.

5. Epistemological Position

Ω does not assert truth.
It produces admissible projections under governance.

Assertions without governance are undefined within Ω.

6. Scope and Non-Claims

Ω is not:

- a performance model,
- an AGI claim,
- a product or benchmark.

Ω is a constitutional framework applicable to intelligent systems, simulations, or institutions requiring bounded, auditable evolution.

7. Conclusion

Ω reframes intelligence as a consequence of constraint. By

embedding governance at the axiomatic level, it eliminates entire classes of failure modes rather than mitigating them. Ω provides a foundation for future intelligent systems where alignment, safety, and scalability are structural properties.