

Coherentie-Interlock Layer (CIL)

A structural safety architecture for autonomous systems.

"No execution without coherence."

Versie 1.0 · 13 maart 2026 · English · Reading time: 5 min. · © Niels van den Hoek

Table of contents

1. Why CIL exists
2. Core architectural pattern
3. Demonstration architecture
4. Positioning
5. Contribution to AI for Good
6. Contribution to the Sustainable Development Goals
7. Conceptual origin

1. Why CIL exists

Autonomous systems increasingly operate across critical domains such as urban infrastructure, embodied robotics, and automated decision pipelines. While monitoring and auditing mechanisms are widely used, most operate retrospectively — identifying issues only after execution has already occurred.

CIL introduces a structural interlock between decision formation and execution. It verifies whether a proposed action remains coherent with the currently active normative state before execution is allowed.

CIL does not determine what is “good.” It ensures that what is declared as normative remains explicit, traceable, versioned, and enforceable.

Rather than preventing change, CIL prevents unreviewed normative drift.

2. Core architectural pattern

CIL operates through a minimal, composable structure:

- **Normative State Registry (NSR)** — explicit, versioned normative declaration
- **Propose Decision** — proposed operational weights or parameters
- **Coherence Verification** — comparison against the active baseline and tolerance
- **Execution Controller** — allow, review, or block

This core mechanism can be extended with declared operational stances and time-bounded context overrides, without altering the underlying interlock logic.

3. Demonstration architecture

CIL is demonstrated across three structural variants:

- **Version 1 – Minimal Interlock**
- **Version 2 – Operational Stances**
- **Version 3 – Context-Aware (TTL Override)**

Each structural variant is accompanied by a scenario verification suite that validates the normative control flow under multiple operational conditions (declaration, review, stance switching, contextual override and registry reset).

Each variant can also be viewed through different domain lenses, such as:

- Urban traffic coordination
- Embodied robotic systems
- Autonomous decision pipelines

The architectural pattern remains identical; only the semantic interpretation of normative variables changes.

4. Positioning

CIL is not a governance document. It is not a monitoring dashboard.

It is an architectural safety layer designed to make unreviewed normative drift structurally impossible.

CIL does not replace regulation or ethical deliberation. It provides the structural enforcement layer that makes normative commitments technically binding within autonomous systems.

The architecture operates independently of ideology, jurisdiction, or sector.

5. Contribution to AI for Good

CIL contributes to AI for Good not by prescribing ethical outcomes, but by structurally reducing the risk of invisible normative drift in autonomous systems.

In complex environments — such as urban infrastructure, robotics, and automated decision pipelines — priorities can shift gradually due to optimization pressure, parameter tuning, or system updates. Without structural safeguards, these shifts may remain undetected until after real-world impact occurs.

CIL addresses this risk at the architectural level. By enforcing explicit, versioned normative baselines and placing an interlock between decision formation and execution, it ensures that significant deviations cannot occur without explicit review and traceable recalibration.

- **Risk Reduction:** Prevents silent priority shifts before execution.
- **Transparency:** Makes normative states explicit and version-controlled.
- **Accountability:** Requires documented review before baseline changes become active.
- **Domain Independence:** Applicable across sectors without relying on specific regulatory regimes.

6. Contribution to the Sustainable Development Goals

As autonomous systems become embedded in infrastructure, cities and public services, maintaining alignment between declared priorities and operational behaviour becomes increasingly important.

CIL contributes to the Sustainable Development Goals by strengthening the governance architecture of autonomous systems. Rather than targeting a single sector, it provides a structural mechanism that helps ensure that automated systems remain aligned with declared institutional objectives.

In particular, the architecture supports goals related to resilient infrastructure, sustainable urban systems, and accountable institutions.

- **SDG 9 – Industry, Innovation and Infrastructure:** Supporting safe deployment of AI within critical technological infrastructures.
- **SDG 11 – Sustainable Cities and Communities:** Enabling reliable governance of autonomous systems operating in urban environments.
- **SDG 16 – Peace, Justice and Strong Institutions:** Improving transparency, traceability, and accountability in automated decision systems.

By ensuring that declared normative priorities remain explicit and enforceable within autonomous architectures, CIL helps align technological autonomy with long-term societal goals.

7. Conceptual origin

CIL emerged alongside the broader AIAS (Artificial Intelligence Alternative Social) exploration, which investigates the societal, organizational and human impact of artificial intelligence.

