The Physical Intelligence Risk & Emergence Framework (PIRE-F)

Behavioral Governance for AI in Physical Systems

Sentinel Physics™
Independent Governance for Physical AI

Version 1.0

Executive Summary

Artificial intelligence is increasingly embedded in physical systems such as robotics, autonomous vehicles, medical devices, and industrial automation. In these environments, AI systems act upon matter and energy, producing risks that are immediate, irreversible, and legally consequential.

The Physical Intelligence Risk & Emergence Framework (PIRE-F) provides a behavior-first governance architecture for evaluating, constraining, and evidencing AI behavior in physical environments.

> PIRE-F evaluates what AI systems actually do — not what documentation claims they do.

1. The Governance Gap in Physical AI

Most AI governance approaches emphasize documentation, intent, and pre-deployment review. Physical AI systems fail at runtime, not on paper.

PIRE-F introduces enforceable, runtime governance aligned with global regulatory expectations.

2. What Is PIRE-F?

PIRE-F is a model-agnostic framework that evaluates physical AI systems based on observable behavior, constraint enforcement, autonomy escalation, and consequence traceability.

3. The Six Governance Domains

Domain	Purpose
Material Sovereignty	Ensures enforceable human authority over physical action
Behavioral Coherence	Detects instability and unsafe behavioral drift
Physical Alignment	Measures conformance to domain safety norms
Boundary Integrity	Evaluates adherence to operational envelopes
Emergent Agency	Detects autonomy escalation and self-directed action
Consequence Probity	Ensures causal traceability of outcomes

4. Physical Risk Index (PRI)

PIRE-F combines the six domain risks into a composite Physical Risk Index (PRI), while preserving single-domain escalation through a domain override rule.

5. Threshold-Based Intervention

PIRE-F defines a four-stage intervention ladder: Observe, Constrain, Reassert Authority, and Isolate. This transforms governance into enforceable operational control.

> Governance is not policy. Governance is the ability to intervene.

6. Standards & Regulatory Alignment

PIRE-F operationalizes the NIST AI Risk Management Framework for physical systems, supports ISO/IEC AI management standards, and provides evidence structures aligned with the EU Artificial Intelligence Act.

Conclusion

As AI systems increasingly act in the physical world, governance must move beyond documentation toward continuous, enforceable oversight. PIRE-F provides a practical framework for meeting this challenge.

About Sentinel Physics™

Sentinel Physics™ is an independent governance firm specializing in behavioral verification of AI systems operating in physical environments.