Governing Agentic AI Across the Industrial Lifecycle

Abstract

The increasing adoption of agent-based AI systems across engineering, manufacturing and operations introduces new challenges for safety, traceability, compliance and lifecycle control. While many current initiatives focus on individual AI components, platforms or use cases, regulated industrial domains require a system-level perspective that ensures coherent interaction between engineering, deployment and production environments.

This contribution addresses the need for a **transversal governance and lifecycle framework for agentic AI systems** that spans the full industrial lifecycle. The focus lies on defining open reference architectures, lifecycle models, human-in-the-loop concepts and traceability mechanisms that enable agent-based AI systems to be designed, validated, deployed and operated in a safe and auditable manner.

Rather than proposing new AI platforms or product implementations, the approach emphasizes **open, reusable and interoperable governance artefacts** that can be adopted across sectors and industrial contexts. By clearly separating upstream governance and orchestration concerns from downstream deployment and execution, the framework supports integration with heterogeneous AI infrastructures, platforms and production environments without creating functional overlap or vendor lock-in.

Embedded in the context of European industrial collaboration, the work highlights how shared reference architectures and open governance models can prevent fragmentation, enable cross-project interoperability and create a foundation for scalable and trustworthy agentic AI in safety-critical domains. The presentation illustrates key design principles and validation patterns using representative industrial scenarios and discusses their relevance for standardisation, SME adoption and long-term European digital sovereignty.