

FIONA

Trajectory governance for AI agents.

FIONA is an external, model-independent layer that evaluates each proposed action against the accumulated state of the task before it is executed.

PRIVATE CORE

RESTRICTED HTTPS CLIENT


MODEL-AGNOSTIC

INTERNAL PROTOTYPE

THE PROBLEM IT ADDRESSES

In long tasks, every step can look reasonable, and yet the sequence can lose commitments, repeat actions, exhaust the budget or end without a usable result.

GOVERNANCE PATTERN



The agent proposes an action; FIONA evaluates it against the accumulated state and returns one of four decisions; the environment acts accordingly.

WHAT EXISTS TODAY

- Private SDK **0.7.0b1** with action and trajectory contracts.
- Restricted client **0.1.0b1** with typed contracts and HTTPS transport.
- **659 tests passed, 2 skipped**, plus type, build and packaging checks.

WHAT STILL HAS TO BE DEMONSTRATED

- Effectiveness in real workflows: completeness, quality, cost and non-interference.
- Production service: authentication, tenant isolation, trace export and performance.
- Independent validation and deployment readiness.

HOW TO READ THE EVIDENCE

Controlled results show **the mechanism's potential, not validation of the product**. An experimental pilot on an external benchmark was mixed and revealed a deadline-control defect. This guides the next validation; **it does not prove general superiority**.

Working objectives, not certifications.

The table below **does not claim compliance**. It shows where FIONA's mechanism *could* contribute to obligations under [Regulation \(EU\) 2024/1689](#) and to risks in the [MIT AI Risk Repository](#), and gives the actual status of each hypothesis. **No row has been validated by an independent third party.**

Timeline. The Regulation has applied generally since 2 August 2026. The rules for high-risk systems under Annex III will apply from 2 December 2027, and those for systems embedded in regulated products under Annex I from 2 August 2028, under [Regulation \(EU\) 2026/1744](#).

Regulatory-fit hypotheses and the current validation status of each.

OBLIGATION OR RISK	WHAT IT REQUIRES	FIONA FIT HYPOTHESIS	CURRENT STATUS
EU AI Act · Art. 9	A documented, continuous risk-management system.	Pre-action evaluation <i>could</i> provide a documented control point.	TO BE VALIDATED IN A PILOT
EU AI Act · Art. 12	Automatic, reconstructable event logging.	The decision trace <i>could</i> cover part of the requirement.	EXPORT FORMAT NOT FINALISED
EU AI Act · Art. 14	Effective human oversight of relevant decisions.	Justify mode <i>could</i> implement the human intervention point.	NOT TESTED WITH REAL OPERATORS
EU AI Act · Art. 15	Robustness and cybersecurity against errors and manipulation.	Separating the decision from the prompt <i>could</i> reduce some manipulation risks, but the net security effect and the new API and service surface have yet to be audited.	NO SECURITY AUDIT
MIT AIRI · 7.3	Lack of capability or robustness.	Deterministic validation of preconditions against the accumulated state.	CONTROLLED RESULTS, NOT IN A REAL WORKFLOW
MIT AIRI · 7.4	Lack of transparency or interpretability.	Rule and state citable for each decision; depends on the quality of the client's ontology.	DEPENDS ON THE ONTOLOGY

Actual regulatory conformity of a high-risk system requires a conformity assessment under the procedures laid down by the Regulation itself. FIONA can contribute technical evidence; it does not replace or guarantee it.

PROPOSED FIRST PILOT

One agent. One high-value workflow. A baseline fixed before we start. We measure **completeness, quality, non-interference, latency and cost** to reach a technical **go/no-go** decision on deployment. If the results do not beat the baseline, the correct outcome of the pilot is *not to deploy*.

LET'S EXPLORE YOUR USE CASE

pere@fiona-labs.com

fiona-labs.com

FIONA provides governance and technical evidence. It is not legal advice, a compliance certification or a guarantee of results. Version and test figures refer to the internal code as of September 2026 and have not been externally audited.