Can ODRL Compliance Be Determined from Logs?

A Causal Observability Framework for Runtime Policy Verification

Position Paper for the W3C ODRL Future Directions Workshop (2026)

Yuqiang Wang, Beijing Cognitive Emergence Technology Co., Ltd.

Email: yuqiang@humanjudgment.org | LinkedIn: https://www.linkedin.com/in/yuqiang-wang-14570213a/

Abstract

ODRL has matured as a language for expressing permissions, prohibitions, duties, and constraints. However, future deployments in data spaces, AI agents, financial systems, and automated policy enforcement face a further question: when is a policy-compliance fact actually determinable from available runtime evidence? This position paper introduces target determinability under partial causal observation as a formal lens for ODRL runtime compliance verification. It argues that future ODRL profiles should define not only policy vocabulary and constraints, but also evidence requirements, counterexample criteria, and compliance-receipt structures needed to determine key compliance targets. Judgment-event records and verifiable runtime traces are discussed as one possible evidence substrate for ODRL enforcement, auditing, and interoperability across agentic systems.

1. Motivation: From Policy Expression to Compliance Determinability

ODRL provides a shared model for expressing rights-related policies over assets, parties, actions, permissions, prohibitions, duties, and constraints. This has made it useful across digital rights, data spaces, industry consortia, creative media, and regulated information flows. The next stage of adoption will increasingly involve automated runtime environments: policy decision points, policy enforcement points, data-space connectors, AI agents, and compliance automation tools.

In such settings, the central question is no longer only whether a policy can be expressed. It is whether a policy-relevant target fact can be determined from the evidence available after execution. For example: Was a duty fulfilled before a permission was exercised? Was a prohibition violated? Was a constraint satisfied at the time of access? Did an agent delegate a policy decision to another component, and was that delegation authorized?

This paper proposes that the ODRL community treat compliance determinability as a first-class future direction. A policy language can specify what should hold; a runtime evidence framework must show whether the relevant target facts are actually settled by logs, traces, credentials, receipts, and causal event records.

2. The Causal Observability Gap in ODRL Runtime Enforcement

A recurring problem in runtime compliance is the causal observability gap: global facts are produced by causal event histories, while observers receive only partial projections of those histories. A system may keep logs without preserving the causal predecessors relevant to a compliance target. It may verify signatures or message formats while still failing to determine whether the target action was authorized under the applicable ODRL policy.

This distinction is important for ODRL. Having a policy is not the same as having execution evidence. Having a log is not the same as having causal coverage. Having causal coverage is not the same as determining a compliance target. Finally, determining a target does not require reconstructing the full internal history; it requires retaining enough target-relevant evidence to distinguish compliant from non-compliant histories.

A useful future question for ODRL profiles is therefore: for each important compliance target, what observations are sufficient, insufficient, or minimally necessary to determine the target with zero error in the intended deployment model?

3. Formal Lens: Histories, Observations, and Compliance Targets

The proposed framework models an external target-determination problem as P = (H, O, Q), where H is the set of possible real execution histories, O is the observation available to the verifier, and Q is the target fact to be determined. In an ODRL setting, H may contain possible policy-execution histories, O may include logs, credentials, policy decisions, enforcement events, and receipts, and Q may state whether an ODRL permission, prohibition, duty, or constraint was satisfied.

A faithful reduction maps the external problem to a formal triple (F, Omega, D), consisting of a finite family of causal event configurations, an observation function, and a formal target. The core criterion is simple: D is determinable from Omega exactly when D is constant on every observational equivalence class. Equivalently, there exists a decision function g such that D = g composed with Omega.

If two possible configurations have the same observation but different target values, they form a certificate of non-determinability. In ODRL terms, if a compliant execution and a non-compliant execution produce the same available runtime evidence, no verifier using only that evidence can correctly determine compliance for both cases.

4. Example: ODRL Policy Execution by an AI Agent

Consider an AI agent that requests access to a licensed dataset. An ODRL policy permits research use, prohibits redistribution, requires attribution, and imposes a duty to delete derived temporary files after processing. A runtime verifier later asks whether the agent's use was compliant.

If the observation contains only the final output, the compliance target is likely not determinable. If it contains only the access decision, it may still fail to show whether duties were fulfilled. If it contains a policy-decision flag without causal evidence, it may not distinguish a valid decision from a forged or stale one. Additional evidence may be needed: policy version, asset identifier, party identifier, purpose binding, decision timestamp, duty-fulfillment event, deletion proof, delegation chain, and tamper-resistant hashes of verification events.

The point is not that every deployment needs every evidence item. The point is that each ODRL profile should be able to specify which compliance targets it cares about and which evidence families are sufficient to determine those targets under the profile's threat model and operating assumptions.

5. Evidence Requirements for ODRL Profiles

Future ODRL profiles could extend beyond vocabulary and constraint definitions to include evidence requirements. The following mapping illustrates the idea:

ODRL target	Possible evidence requirement	Determinability risk if missing
Permission exercised validly	Policy version, party, asset, action, purpose, time, authorization record	Access may be logged but authorization validity remains unsettled
Prohibition not violated	Negative evidence over relevant action scope; distribution/use trace; audit boundary	Absence of a violation record may not prove no violation
Duty fulfilled	Causal predecessor event showing attribution/payment/deletion/reporting completion	The system may show output without showing required duty completion
Constraint satisfied	Runtime measurement or credential proving time, geography, role, purpose, or count condition	A policy decision flag may not prove the constraint held at execution time
Delegated decision valid	Delegation event, authority chain, policy scope, verifier identity	A downstream agent may act without recoverable authorization path

The constrained-evidence view also gives profile authors a concrete tool-development target. Given a finite model, a checker can group possible configurations by observation value and test whether each group has a unique compliance target value. If yes, it can output a decision table. If not, it can output an indistinguishable counterexample pair and the evidence gaps that must be closed.

This suggests a practical ODRL tooling roadmap: policy editors, validators, and conformance tests could evolve toward determinability-aware tooling that checks not only whether a policy is syntactically valid, but whether the associated evidence model can determine the intended compliance targets.

6. Toward Compliance Receipts and Judgment-Event Evidence

ODRL runtime environments could produce compliance receipts that bind policy identifiers, asset identifiers, decision events, enforcement events, duty events, and verification evidence. Such receipts need not expose full internal histories. They should expose enough target-relevant causal evidence to make the selected compliance facts determinable.

Judgment-event records are one possible substrate for this evidence layer. They can record when an automated system interpreted a policy, made a policy decision, delegated a decision, verified a duty or constraint, or terminated an action. These records can then be linked with verifiable credentials, decentralized identifiers, signed receipts, or data-space governance logs.

This approach complements ODRL rather than replacing it. ODRL remains the policy-expression layer. Determinability analysis provides a formal method for asking whether runtime observations settle policy-relevant facts. Judgment-event evidence and compliance receipts provide an implementation direction for automated and auditable enforcement.

7. Proposed Discussion Questions for the ODRL Community

- Should future ODRL profiles define required evidence for core compliance targets, in addition to policy vocabulary and constraints?
- Can ODRL conformance testing include determinability tests that output counterexample pairs when runtime evidence is insufficient?

- How should ODRL integrate with verifiable credentials, DIDs, signed event records, and data-space receipts without over-specifying implementation details?
- Which compliance targets are common enough across domains to justify shared evidence patterns, and which should remain profile-specific?
- Can the ODRL community define a minimal compliance-receipt model for runtime enforcement and audit interoperability?

8. Conclusion

ODRL's future importance will grow as policies are executed by automated systems rather than interpreted only by humans. In that environment, policy interoperability must be joined by compliance determinability: the ability to decide whether policy-relevant target facts are settled by available runtime evidence.

This position paper proposes target determinability under partial causal observation as a formal lens for that transition. It invites the ODRL community to consider evidence requirements, counterexample-based testing, compliance receipts, and judgment-event records as part of the next stage of ODRL runtime verification and profile governance.