ODRL Open Issues*

Víctor Rodríguez-Doncel
Universidad Politécnica de Madrid
Madrid, Spain

Nicoletta Fornara
Università della Svizzera Italiana
Lugano, Switzerland

Andrea Cimmino
Universidad Politécnica de Madrid
Madrid, Spain

Abstract

The Open Digital Rights Language (ODRL) version 2.2 was published in 2018, and since then a substantial body of academic and practitioner work has built upon, extended, and critiqued it. This short position paper consolidates that effort into a structured set of requirements and issues for ODRL 3.0. Using artificial intelligence to assist in extracting and synthesising evidence, we analysed three complementary sources: 22 scientific publications (2018–2026); 155 issues in the W3C ODRL GitHub repository; and 713 messages on the public ODRL mailing list (January 2019–June 2026)—the work of about 100 people. From this corpus we identified over 30 requirements in five thematic areas (formal semantics and evaluation, expressivity, operationalisation and profiles, interoperability, and vocabulary quality) and offer them as input to the ODRL 3.0 standardisation process. We argue that a normative formal semantics is the crucial requirement for ODRL 3.0. The full paper, to be made available by the time of the Workshop, will additionally describe the business value of some of the key issues and requirements.

1 Introduction

ODRL is a W3C standard, published as two Recommendations—an information model [7] and a vocabulary and expression specification [6]—that together provide a policy expression language for permissions, prohibitions, and obligations over digital content and services. Since version 2.2 (2018), ODRL has attracted sustained attention at venues such as the Web Conference, ISWC, ESWC, SEMANTiCS, and dedicated ODRL workshops, spanning formal semantics, rule-based reasoning, domain-specific profiles, and machine-executable enforcement. Industrial needs have reinforced this: the rise of data spaces (federated infrastructures for sovereign data sharing, championed by Gaia-X and the IDSA) has placed ODRL at the heart of real-world policy exchange, with requirements beyond the standard's original scope.

The W3C ODRL Community Group is actively preparing the next version, and a broad, community-sourced picture of requirements is exactly what is needed. We survey the scientific literature published since 2018, the GitHub issues, and the public mailing list to extract requirements, feature requests, and critiques. We make no claim that every requirement should be incorporated into ODRL 3.0—some will conflict with the language's design principles or with one another. Productive requirements elicitation begins with a broad, inclusive collection of ideas before selection and prioritisation; our contribution is to provide that initial picture.

*The authors thank Renato Ianella for his valuable input to these ideas.

2 Methodology

We followed a lightweight systematic literature review protocol inspired by Kitchenham and Charters [8], drawing on three sources.

Scientific literature. Starting from the AI-assisted tool Elicit and complemented by queries to Google Scholar, Semantic Scholar, the ACM Digital Library, and IEEE Xplore (search string `ODRL OR "Open Digital Rights Language"`, restricted to 2018 onwards), we screened candidates by title and abstract to retain technical work on policy modelling, evaluation, or extension. For each paper we recorded the use case, limitations explicitly mentioned, proposed extensions or workarounds, and feature requests. Large language models surfaced candidate items, which the authors then reviewed and cross-checked to discard hallucinated or unsupported claims.

GitHub issues. We analysed the 155 open and closed issues of the W3C ODRL repository (June 2026), covering vocabulary terms, semantic clarifications, profile guidance, evaluation semantics, and explicit ODRL 3.0 feature requests, together with the wiki *Roadmap brainstorm* page [10].

Mailing list. We reviewed all threads on the public list between January 2019 and June 2026—90 monthly archive pages, 713 messages—screening by subject line and reading substantive threads in full, again with AI assistance and author verification. Key topics include formal evaluation semantics and the Open- vs. Closed-World debate, temporal modelling, the Constraint vs. Refinement distinction, profiles (data sovereignty, security, temporal, verifiable credentials, AI), enforcement and validation, agent delegation, and alignment with IDSA/DCAT.

Extracted items from all three sources were grouped thematically, de-duplicated, assigned unique identifiers, and classified by theme.

3 Requirements

The requirements are grouped into five thematic areas. For brevity, only the requirement statements are given here; the full paper documents the supporting literature for each.

3.1 Formal Semantics and Evaluation

A recurring theme is the absence of a normative, machine-interpretable semantics for ODRL 2.2 [1–5,9].

R-S1 Normative Evaluation Semantics. A normative formal evaluation model (inference rules or a reference algorithm) that unambiguously defines whether a request satisfies a policy.

R-S2 State-of-the-World Representation. A standard representation of executed actions, environmental facts, and temporal state as input to the evaluator.

R-S3 Evaluation Request/Response Format. A standard request and response format, aligned with emerging standards such as AuthZen/RFC 9396.

R-S4 Closed-World Assumption. Explicit adoption of the Closed-World Assumption (prohibited-by-default) for policy evaluation, with an explicit gap-resolution strategy.

R-S5 Operand/Operator Specification. Normative implementation guidance for `LeftOperand`, `RightOperand`, and `Operator` terms (data types, value sources, comparison semantics), including how knowledge bases are linked for KB-dependent operands.

R-S6 Conflict Resolution. Formal semantics for resolving overlapping or contradictory rules, supporting legal priority principles (*lex specialis*, *lex posterior*, *lex superior*).

R-S7 Policy Comparison Semantics. Formal semantics for comparing two policies (equivalent, more permissive, or more restrictive), essential for data-space matching and license composition.

R-S8 Policy Requirements Querying. Declarative querying of which rights are granted, which obligations apply, and which constraints must hold—without a full state-of-the-world evaluation.

R-S9 Semantics of Permissions and Duties. Clarification of what it means for a duty to be *fulfilled* and how this affects the associated permission, including a dedicated `Obligation` class with activation conditions, deadlines, and triggers.

3.2 Expressivity

R-E1 Composite Logical Constraints. Nested combinations of `isAllOf`/`isAnyOf`/`isNoneOf`.

R-E2 Multi-Dimensional Constraints. Per-axis semantics for multi-dimensional operands (e.g. `absoluteSize`), as in the proposed Axis-Aligned Profile.

R-E3 Dynamic External Data. Incorporating dynamic values (geolocation, trusted datetime) at evaluation time, with continuous re-evaluation.

R-E4 Spatial and Jurisdictional Constraints. Native geographic and legal-jurisdiction constraints, bindable independently to action, assignee, and target.

R-E5 Temporal Constraints. Recurring windows, event-relative deadlines, and duration-based constraints.

R-E6 Delegation and Trust Chains. Delegated permissions and trust chains of arbitrary length, plus rights revocation.

R-E7 Policy Lifecycle and States. Formal modelling of `Set` → `Offer` → `Agreement` transitions and validity conditions.

R-E8 Policy Templates. Abstract/template policies with variables resolved at instantiation time.

R-E9 Multiple Targets/Parties per Rule. Sequences of assets, parties, and actions as rule parameters.

R-E10 Publisher/Creator Actions. A top-level *creation/publishing* action complementing `use` and `transfer`.

R-E11 AI Training Action. A standardised term for permitting/prohibiting use of assets in AI model training.

R-E12 Payment by Percentage. A `payPercentage` left operand complementing `payAmount`.

R-E13 Hierarchical Class Operators. `subClassOf`/`superClassOf` operators with inclusive/exclusive variants.

R-E14 Agent Delegation. Ring-fenced, time- and count-limited delegation from a human principal to an AI agent.

R-E15 Post-Processing and Derived Assets. Governing assets *derived from* the licensed asset (e.g. RAG/AI pipeline outputs).

R-E16 Constrained Assignee Types. Restricting the *type* of eligible assignee via a referenced taxonomy.

R-E17 Temporal Profile. A standardised temporal profile maintained alongside the core specification.

R-E18 Security Profile. A profile for authentication, authorisation tokens, and secure policy exchange.

3.3 Operationalisation and Profiles

R-O1 Conformance Testing. A standard evaluator report format and test suite, with conformance levels for distinct processor types (validator, normaliser, converter, profiler, authoriser, monitor, conflict resolutor).

R-O2 Profile Definition Mechanism. A machine-readable mechanism for defining profiles (constraint/rule subclasses, namespaces).

R-O3 Deprecation of `andSequence`. Deprecate `odrl:andSequence` in favour of event-based sequencing, given its underspecified semantics.

R-O8 Decoupled Reasoning and Materialisation. Separate the core evaluation model from reasoning over ontological dependencies, via a standardised normalisation process.

R-O4 Audit Trails and Provenance. Attaching provenance and audit information to evaluations and executed actions.

R-O5 Enforcement and Validation Tooling. Endorsed best practices and reference tools, including SHACL-based validators.

R-O6 ODRL Shapes. SHACL shapes expressing the structural constraints of valid ODRL documents.

R-O7 AI Vocabulary Profile. A community-maintained profile for AI actions and constraints (training, inference, output use).

3.4 Interoperability

R-I1 Alignment with Related Standards. Maintained mappings to DCAT 3, DPV, DUO, and the IDSA Information Model.

R-I2 Data Spaces. Explicit support for data-space architectures (Gaia-X, IDSA): negotiated, exchanged, and enforced policies; usage control beyond access control.

R-I3 Policy Discovery and Management. A standard way to link resources to governing policies and a reference architecture for policy management, discovery, and enforcement.

R-I4 Verifiable Credentials Integration. An ODRL profile for use within Verifiable Credentials.

R-I5 Market Data and Domain Profiles. Domain profiles for regulated sectors (financial market data, archival records, healthcare consent).

R-I6 Cross-Language Interoperability. Guidance on translating, aligning, or federating policies across languages (ODRL, Solid WAC, XACML, RBAC).

R-I7 Machine-Readable Copyright in the AI Era. Positioning ODRL as the vehicle for machine-readable copyright/licensing metadata for AI training, aligned with emerging EU and international frameworks.

3.5 Specification Quality

R-V1 Refinement vs. Constraint. Clarify the distinction, with worked examples; it is load-bearing for automated license composition.

R-V2 `hasPart/isPartOf`. Improve definitions to avoid ambiguity with set-membership operators.

R-V3 State Diagrams. Non-normative state and sequence diagrams for the policy lifecycle and deontic states.

R-V4 Errata and Corrections. Fix typographical and syntactic errors (e.g. `neq`/`neg`, invalid JSON-LD examples).

4 Conclusions

We have consolidated requirements for ODRL 3.0 from three complementary sources—scientific publications since 2018, the 155 GitHub issues, and roughly seven years of mailing-list deliberation (2019–2026)—into over 30 requirements across five thematic areas. The most consistently recurring needs are: a normative evaluation model under the Closed-World Assumption; a defined state-of-the-world representation; richer constraint expressivity (composite logical constraints, multi-dimensional operands, spatial and temporal conditions, delegation chains); agent delegation for AI use cases; explicit support for data-space deployments; and a family of community-maintained profiles (AI vocabulary, temporal, security, verifiable credentials, and domain-specific). Not all requirements are mutually compatible, and not all will reach ODRL 3.0; the purpose is to provide a broad, inclusive starting point for the prioritisation and design work ahead.

The full version of this paper will extend the present analysis in three ways. First, it will include a table mapping the identified issues to the use cases that motivate them. Second, it will broaden the source corpus by also analysing the Community Group meeting notes. Third, it will give greater weight to formal semantics, which we regard as the distinctive feature of ODRL 3.0.

Use of Artificial Intelligence. Artificial intelligence (Anthropic Claude) was used substantially to analyse the surveyed papers and to assist in drafting. All AI-generated content was reviewed, verified, and edited by the authors, who take full responsibility for the final text.[1]

[1]This work has been supported by HARNESS (EU Horizon Europe, grant agreement no. 101169409) and

References

[1] Piero Andrea Bonatti, Nicoletta Fornara, and Andreas Harth. Towards a formal semantics of the open digital rights language (ODRL 2.2). In *ODRL and Beyond: Practical Applications and Challenges for Policy-based Access and Usage Control (OPAL 2025), co-located with ESWC 2025*, 2025. CEUR Workshop Proceedings.

[2] Andrea Cimmino, Juan Cano-Benito, and Raúl García-Castro. Practical challenges of ODRL and potential courses of action. In *Companion Proceedings of the ACM Web Conference 2023 (WWW'23 Companion)*, pages 1428–1431. ACM, 2023.

[3] Andrea Cimmino and Nicoletta Fornara. Improving ODRL 2.2: Current limitations and theoretical solutions. In *ODRL and Beyond: Practical Applications and Challenges for Policy-based Access and Usage Control (OPAL 2025), co-located with ESWC 2025*, 2025. CEUR Workshop Proceedings.

[4] Marina De Vos, Sabrina Kirrane, Julian Padget, and Ken Satoh. ODRL policy modelling and compliance checking. In *Rules and Reasoning (RuleML+RR 2019)*, volume 11784 of *Lecture Notes in Computer Science*, pages 36–51. Springer, 2019.

[5] Patrick Hochstenbach, Beatriz Esteves, and Ruben Verborgh. Automated policy negotiation: a four-course meal. In *ODRL and Beyond: Practical Applications and Challenges for Policy-based Access and Usage Control (OPAL 2025), co-located with ESWC 2025*, 2025. CEUR Workshop Proceedings.

[6] Renato Iannella, Michael Steidl, Stuart Myles, and Víctor Rodríguez-Doncel. ODRL vocabulary & expression 2.2. W3C recommendation, W3C, February 2018. `https://www.w3.org/TR/odrl-vocab/`.

[7] Renato Iannella and Serena Villata. ODRL information model 2.2. W3C recommendation, W3C, February 2018. `https://www.w3.org/TR/odrl-model/`.

[8] Barbara Kitchenham and Stuart Charters. Guidelines for performing systematic literature reviews in software engineering. EBSE Technical Report EBSE-2007-01, Keele University and Durham University, 2007.

[9] Jaime Osvaldo Salas, Paolo Pareti, Semih Yumusak, Soulmaz Gheisari, Luis-Daniel Ibáñez, and George Konstantinidis. Evaluation and comparison semantics for ODRL. In *arXiv preprint*, 2025. arXiv:2509.03915.

[10] W3C ODRL Community Group. ODRL roadmap brainstorm, 2024. `https://github.com/w3c/odrl/wiki/Roadmap-brainstorm`.

by MALTA PID2024-159504OB-I00 funded by MICIU/AEI/10.13039/501100011033.