

# Building Thoroughly Tested Semantic Rule Engines in the Age of GenAI

Towards realizing Interoperable, Legally Compliant and Trustworthy Data Flows

Patrick Hochstenbach<sup>1,2,\*</sup>, Jos De Roo<sup>1</sup>, Wout Slabbinck<sup>1</sup>, Beatriz Esteves<sup>1</sup> and Pieter Colpaert<sup>1</sup>

<sup>1</sup>Ghent University, IMEC-IDLab, Ghent, Belgium

<sup>2</sup>Ghent University, Ghent University Library, Ghent, Belgium

## Abstract

Data sovereignty starts before data leaves a user's device. As data processing moves toward the edge, future Web applications will need to evaluate rules, policies, and explanations locally, rather than delegating every decision to server-side infrastructure. At the same time, data will increasingly flow across organizations, applications, and regulatory contexts. We argue that such flows must be interoperable across ecosystems, legally compliant with applicable policies, and trustworthy to end-users. This requires a rule language that independent engines can implement consistently. In this paper, we introduce 1k+ Notation3 compliance tests for validating implementations against the Notation3 specification, developed in parallel with Eyeling: a JavaScript, browser-deployable reference implementation for interoperable in-browser trust engines. Applied to cwm, EYE, and jen3, the suite reveals conformance gaps with the published specification, while Eyeling reaches full conformance and shows that LLM-driven synthesis can bootstrap spec-compliant rule engines from a specification and conformance suite. These results suggest that in-browser semantic rule engines can provide a practical foundation for evaluating interoperability, legal compliance, and trust where data is accessed and used. While some may see Semantic Web tooling as deprecated in the era of genAI, our work can serve as an example of a symbiosis leveraging the best from both worlds.

## Keywords

Notation3, rule engine, Javascript, LLM code generation, Neurosymbolic AI

## 1. Introduction

The future Web will not be organized around a single global data model, a single integration platform, or a single trusted authority. Data will flow across organizations, applications, sectors, and regulatory contexts, while remaining shaped by the local needs of the actors that produce, manage, and consume it. In such an environment, data sovereignty starts before data leaves a user's device: users and applications must be able to understand, evaluate, and act on data locally, rather than delegating every decision to remote infrastructure. This requires building Web applications in which interoperability, legal compliance, and end-user trust are taken into account [1].

**Interoperability** cannot depend solely on prior agreement. Organizations will continue to use different vocabularies, schemas, policies, and operational assumptions. Rather than forcing all actors into a common model from the start, interoperability needs to be achieved progressively, when concrete integration needs arise. This idea of eventual interoperability shifts the problem from universal standardization to explicit interpretation: mappings, alignments, constraints, and transformations must be expressed in a form that other systems can inspect, reuse, and execute. Semantic rules are a natural fit for this role. They make translation assumptions explicit, provide a formal account of how one

---

*The International Workshop on Semantic Architectures for Governance and Explainability (SAGE26), co-located with SEMANTiCS 2026: 22nd International Conference on Semantic Systems, September 15, 2026, Ghent, Belgium*

\*Corresponding author.

✉ patrick.hochstenbach@ugent.be (P. Hochstenbach)

ORCID 0000-0001-8390-6171 (P. Hochstenbach); 0000-0001-8862-0666 (J. De Roo); 0000-0002-3287-7312 (W. Slabbinck); 0000-0003-0259-7560 (B. Esteves); 0000-0001-6917-2167 (P. Colpaert)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

representation is interpreted in terms of another, and allow heterogeneous data to be aligned without erasing local semantics.

The same need for explicit, executable interpretation appears in **legal compliance** [2]. Data flows are constrained not only by technical formats, but also by permissions, prohibitions, obligations, purposes, legal bases, and contextual conditions. Policy languages such as the Open Digital Rights Language (ODRL) [3] and vocabularies such as the Data Privacy Vocabulary (DPV) [4] make it possible to describe aspects of these constraints in machine-interpretable form. However, descriptions alone are not enough. Applications must be able to evaluate whether a given data use is allowed, whether additional obligations apply, or whether access should be denied under a specific policy. For such evaluations to be interoperable across implementations, the underlying rule semantics must be shared, testable, and independent of a single proprietary engine.

**Trust** then becomes a question of execution as much as representation. If interoperability mappings and compliance policies are evaluated only by remote services, users and applications must trust those services before they can trust the outcome. Client-side reasoning changes this relation. Rules, policies, input data, and derivations can be executed and inspected close to the user, for instance inside a browser, a data wallet, or an edge application. Yet this only helps if the local engine itself can be trusted. A client-side rule engine must therefore be validated against a shared specification and through implementation-independent compliance tests.

**Contributions.** This paper addresses that need by leveraging the decades of work done creating Notation3 (a first team submission was published in 2008<sup>1</sup>, that was later adopted by the Notation3 W3C Community Group). We present a Notation3 conformance test suite of more than thousand cases, each traceable to a specific normative clause of the N3 specification, publicly released for community use. Applied to the established reasoners cwm [5], EYE [6], and jen3 [7], the suite quantifies the conformance gap between current implementations and the published specification. We further present Eyeling, a Notation3 reasoner constructed via human-in-the-loop LLM synthesis directly from the specification text, supporting both forward and backward chaining with complete coverage of the N3 built-in functions and achieving 100% on the suite. Beyond the system itself, our work illustrates a broader methodological pattern: well-scoped Semantic Web specifications, paired with compliance tests, can serve as a viable basis for AI agents to implement core interoperable infrastructure.

## 2. Symbolic rules for trust and interoperability in the age of AI

**Assertions and alignments** Interoperability and trust are closely related. Every assertion is made relative to a vocabulary, a context, and an intended interpretation. When an organization states that something is true, it implicitly commits to the meanings of the terms used, the modelling choices behind them, and the conditions under which others may rely on the statement. Across ecosystems, different actors may use different identifiers for similar concepts, reuse the same term with different assumptions, or publish data at different levels of abstraction. By using a rule language and engine, one enables eventual interoperability<sup>2</sup>. Alignments, mappings, constraints, and rules express how one actor's assertions may be understood in another actor's context, and under which assumptions additional knowledge may be inferred. A rule can state that, for a specific purpose, a locally used class or property can be treated as corresponding to another one; that a value must be normalized before comparison; or that a conclusion only follows if a particular provenance, permission, or policy condition is present. For these resulted alignments and assertions to be trusted, one certainly also needs to trust the rule engine used, which thus needs heavy testing.

**Agent spaces** In an idea coined by Polleres et al. [8], data agents are a further evolution of data spaces, where not only data assets, metadata, policies, contracts, and credentials are first-class citizens, but also

---

<sup>1</sup><https://www.w3.org/TeamSubmission/2008/SUBM-n3-20080114/>

<sup>2</sup>Eventual interoperability was coined by Pieter Colpaert in a blog post published in January 2026 <https://pietercolpaert.be/interoperability/2026/01/08/eventual-interoperability>

knowledge graphs, services, tools, and agentic language models. Their central observation is that simply allowing LLM-based agents to operate over data spaces does not by itself preserve sovereignty or trust: agents may discover data, call tools, compose services, and act on behalf of humans or organizations, but their non-deterministic behavior also makes it harder to guarantee which data was used, for which purpose, under which policy, and with which justification. The proposed answer is a neuro-symbolic architecture in which agentic models are governed by **symbolic policy guardrails**. Such guardrails make policies, constraints, contracts, and process rules explicit and executable, so that an agent may propose an action, formulate a query, or mediate between heterogeneous sources, while a deterministic symbolic component checks whether the action is permitted, whether the relevant constraints hold, and which derivation justifies the result. In this view, symbolic rules are the mechanism through which agent spaces can become controllable, explainable, and policy-compliant: they define the boundary between what an agent can suggest and what the surrounding space can accept as valid, allowed, or trustworthy.

**Executable Legal Compliance** Deterministic rule engines in the form of programming-language-based formalisms have recently gained renewed interest as foundations for policy and rights languages such as ODRL [9, 10, 11], eFLINT [12], and LegalRuleML [13, 14, 15]. Their expressive power enables the modeling of complex conditions found in law, regulations, and data-sharing agreements: conditions that cannot be captured by OWL alone. While OWL offers a starting point for interoperability, no single-logic implementation has emerged that can express all legal requirements; existing approaches therefore resort to cumbersome combinations of OWL with SPARQL, SHACL, and ASP. What is needed instead is a unified logic that can express both the policies and their consequences within a single formalism [16].

**Notation3** (N3), as a Semantic Web language, can express rules over existing legal and policy vocabularies such as ODRL [17] and DPV [4]. Combined with semantic annotations of Terms of Service, N3 rules can infer the presence, absence, or inconsistency of legally relevant clauses, such as missing safeguards for cross-border transfers or missing transparency provisions for profiling and automated decision-making [18]. Its support for rule chaining, date-time arithmetic, string parsing, cryptography, and (scoped) negation-as-failure (SNAF) makes it possible to express policies and consequences within one deterministic execution environment, rather than relying on combinations of OWL, SPARQL, SHACL, and ASP. Recent attempts to apply Notation3 in legal compliance use-cases can be found in [19, 2, 16, 20, 21].

Thus, whether it's about scaling up trust in data sources expressed using heterogeneous data models, about providing safeguards for agents in agent spaces, or for providing legal compliance, the basis for trust anywhere is a well-tested rule engine.

### 3. Notation3 Conformance Test Suite

The Notation3 specification is accompanied by an official conformance test suite <sup>3</sup>, which provides a valuable foundation for any compliance evaluation of N3 implementations. The suite was, however, assembled during the drafting phase and has not been revised alongside subsequent editorial and normative changes. As a consequence, some tests reflect intermediate drafts of the specification, deprecated language features, or the behavior of early reference implementations such as cwm [5], rather than the normative content of the team submission published in July 2023 [22]. Conversely, several grammar and built-in requirements introduced, or refined in the later stages of standardization remain unexercised. In addition the original test suite contains tests for built-ins and language features that did not make it into the team submission. This motivates the revised conformance test suite presented in this article.

To test the adherence of emerging N3 reasoners against the published N3 specification, we developed a new handcrafted test suite, `notation3tests` [23] (two of us are part of the Notation3 Community

---

<sup>3</sup><https://w3c.github.io/N3/tests/>

Group). The suite provides extensive coverage of N3 grammar, semantics, and adherence to the official list of built-in functions [24] together with their argument modes and datatype casting. Each test case exercises a single, atomic grammar feature, semantic feature, built-in, argument mode or datatype cast, and cites the normative section of the specification it covers. Test outcomes are partitioned into four categories:

- *conformant* : the reasoner produces the expected output;
- *non-conformant* : the reasoner produces an output where none is expected;
- *incomplete* : the expected output is absent from the reasoner’s response;
- *crash* : the reasoner terminates during execution.

Within the crash category, the suite further distinguishes *expected* crashes, those arising from tests that deliberately introduce grammar errors or contradictions in the input theory, from *unexpected* crashes, which indicate implementation defects.

**Table 1**

Overview of compliance tests results for the cwm, EYE and jen3 N3 reasoners (as of April 29,2026).

| Test (N=972)   | <b>cwm</b> | <b>EYE</b> | <b>jen3</b> |
|----------------|------------|------------|-------------|
| conformant     | 681 (70%)  | 786 (81%)  | 735 (76%)   |
| non-conformant | 20 (2%)    | 22 (2%)    | 48 (5%)     |
| incomplete     | 220 (23%)  | 147 (15%)  | 146 (15%)   |
| crashed        | 51 (5%)    | 17 (2%)    | 43 (4%)     |

As of April 29, 2026, `notation3tests` comprises 972 test cases (today this is over 1000). The results of running the compliance suite against three principal existing N3 implementations (cwm commit 7a71cfd [5], EYE v11.23.10 [6], and jen3 commit 3e6c11d [7]) are presented in Table 1.


These three scores indicate a conformance gap between the existing implementations and the final published N3 specification, but the gap is not a measure of implementation quality. Each reasoner completed its main development cycle before the final Notation3 community group report was available, and each project validates against its own project-specific test suite or against the original W3C conformance suite mentioned at the start of this section.

In order to bridge the conformance gap and provide a fully compliant implementation, we set out to create a new N3 reasoner using LLM-driven synthesis, with our test suite made available to the LLM from the outset. The resulting Eyeling reasoner (described in detail in Section 4) reached a near-100% pass rate by version 1.6.13 using the live version of `notation3test` that existed at that time (the 73rd version in chronological order). At that point, 440 of the eventual 972 tests had been written, so this milestone reflects compliance against roughly 45% of the final suite.

This headline figure, however, overstates true compliance: the suite was incrementally extended by the testing team throughout the project, so each historical version of the Eyeling reasoner was scored only against the tests that existed at the time. A more faithful measure can be obtained *post factum*, by running the 972-test suite against every historical version of the Eyeling reasoner. The results of this retrospective evaluation of `notation3tests` are shown in Figure 1 (up to April 29, 2026). The single outlier visible in the figure is version 1.11.10 (the 167th version, chronologically), for which the LLM synthesis run did not yield an executable reasoner.<sup>4</sup>

Eyeling’s seemingly “100%” score is a consequence of its development being carried out in lockstep with the compliance suite. This co-development is an obvious limitation of our compliance claims. To provide confidence in our results, two authors not involved in the prompting sessions reviewed the parser, rule normalizer, and built-in dispatch module for case-specific branches and test-string matching (such as the names of test cases, names of predicates used in the test cases). No such patterns were found, which suggest a genuine reasoning behavior rather than overfitting to the test results. As an

<sup>4</sup>The LLM-driven synthesis run produced source code that could not be executed, because the resulting file lacked the required executable permission bits. The testing team chose not to correct the synthesized reasoner even for such simple cases.




**Figure 1:** Eyeling conformer over tagged versions. Top: share of failing tests broken down by failure mode (crashed, non-conformant, incomplete) on a logarithmic axis. Bottom: percentage of conforming tests over the same versions.

additional external check, on May 20, 2026 we ran the latest Eyeling version at that time, v1.24.27, against the official W3C N3 Turtle and syntax tests; it passed 308 of 339 (91%). Some of the remaining failures correspond to language features the team submission excluded, as discussed in Section 3.

We hope that publishing our compliance suite will encourage the development of further comparable implementations and provide a common benchmark against which their conformance can be measured.

## 4. Eyeling

Eyeling is an LLM-driven synthesized N3 reasoner produced by the project’s reasoning team using *GPT-5 Extended Thinking* (using versions 5.1, 5.2, 5.4 and 5.5) throughout the course of the project. The inputs used to develop the Eyeling reasoner comprised the three N3 specifications [22, 24, 25] and our notation3tests compliance test suite. Beyond compliance, the synthesis had to produce a mature implementation whose performance was on par with the EYE reasoner. This required thousands of prompts guided by human feedback. When the LLM was stuck, selected fragments of the EYE prolog source were provided as hints. Eyeling is therefore not a translation of EYE: the underlying V8 engines (Eyeling) and SWI-Prolog (EYE) differ in ways that materially affect implementation choices, clause indexing as the most consequential. As of April 29, 2026, this process had produced 294 chronological versions.<sup>5</sup>



**Figure 2:** High-level architecture of the Eyeling reasoner.

Figure 2 provides a high-level architectural overview of the Eyeling reasoner. Eyeling accepts one or more N3 documents as input, each parsed individually into an internal AST representation. A *Rule Normalization* step then applies blank node normalization to the AST, promoting blank nodes appearing in rule premises to universally quantified variables, as required by the N3 specification. This step also

<sup>5</sup>This corresponds to approximately five months of human-assisted development time.

handles RDF collections, ensuring that both `rdf:first/rdf:rest` chains and native N3 list notation are treated uniformly throughout the reasoning pipeline. The normalized rules and initial facts are then passed to the reasoning engine, which iteratively applies forward and backward chaining until a fixpoint is reached. Derived solutions are emitted as N3 output in a streaming fashion.

Both forward and backward chaining are supported within a single N3 input document, in accordance with the N3Logic semantics as specified in [25]. In addition, the synthesized reasoner implements complete coverage of the N3 Built-in Functions specification.

Despite its LLM-driven synthesis, Eyeling has no runtime dependencies on any language model: every execution consists entirely of deterministic symbolic reasoning steps. Once generated, Eyeling operates as a fully self-contained deterministic symbolic N3 reasoner. The implementation is written in pure JavaScript and has no external dependencies and no runtime requirements (e.g., no WebAssembly runtime). It can be executed as a command-line tool, consumed via an API, or embedded directly in a web browser.<sup>6</sup>

In addition to its score on the compliance suite (Section 3), Eyeling was also the fastest of the four N3 implementations (eye, eyeling, cwm and jen3) on the compliance-suite throughput. These timings should be read with caution: the compliance suite is correctness-oriented and does not exercise the math-heavy or large-scale workloads that engines like EYE are typically used for.

We compared the latest Eyeling version (v2.35.9 at the time of writing) against four benchmark families: Deep Taxonomy [26], LUBM(1) [27], OpenRuleBench (join1/tc/sg) [28], and eight Gabriel compute kernels [29]. On LUBM(1), Eyeling answers all fourteen queries correctly at its parse floor and is the fastest of the N3 reasoners tested (EYE, jen3, cwm), though still about 10× slower than the fastest Datalog engines. On join-heavy Datalog workloads it is 6–18× slower than the best Datalog engine.

## 5. Discussion: AI-Assisted Semantic Infrastructure

Eyeling is a symbolic component that can be embedded next to AI agents. In an agent space, an LLM may communicate with users and services, but a N3 reasoner such as Eyeling can act as a meaning boundary: it decides which conclusions follow from the formalized knowledge available to the system. The LLM can help author, refactor, or explain N3 rules, but the rule engine remains responsible for deterministic derivation and validation. This separation is particularly important for trust, because the outcome of a policy or interoperability decision should not depend solely on a probabilistic model invocation.

A browser-native rule engine further strengthens this trust relation. If a remote server is the only place where policies are evaluated, then the user must first trust that server before trusting the decision it returns. By contrast, an in-browser reasoner allows facts, policies, checks, and explanations to be evaluated close to the user and close to the data. For example, an agent may ask whether it may use a dataset for a given purpose, whether a planned action is compatible with a contract, or which facts justify a recommendation. Eyeling can represent the relevant data, rules, checks, and answers in one N3 artifact, enabling governed and inspectable decisions without requiring the user to expose all raw data to a remote policy engine.

This also suggests an alternative to unconstrained data sharing in agent spaces. Instead of giving agents broad access to raw data, a local rule engine can derive narrow, purpose-limited, and auditable insights. The Answer–Reason Why–Check pattern supported by Eyeling follows this idea: the system can compute an answer, provide the derivation that explains why the answer follows, and verify constraints that must not be violated. Such a pattern is well suited to agent spaces, where agents should operate over governed insights rather than unrestricted data dumps. It allows symbolic policies to determine not only whether an action is allowed, but also which explanation can be exposed to the user, another agent, or an auditor.

Finally, agent spaces require portability. Symbolic guardrails are only useful if they can be deployed where agents actually run: in Web applications, data-space connectors, policy editors, sandboxed

---

<sup>6</sup><https://eyereasoner.github.io/eyeling/playground>

runtimes, personal data stores, and client-side agent environments. Eyeling targets this requirement by providing a JavaScript implementation of N3 reasoning that can run in browsers and standard Node.js environments.

A natural question is whether this synthesis methodology generalizes beyond Eyeling’s JavaScript target. As a check, we used Claude Opus 4.7 to synthesize a C++ N3 reasoner from the W3C specifications, the `notation3tests` suite, and Eyeling’s source, with minimal human input [30]. After 29 commits and about six hours we had a spec-compliant reasoner: it matches Eyeling on *compliance*, but not on *maturity* or *performance*. One caveat: the N3 specification leaves room for interpretation, and the C++ implementation resolved the ambiguous cases the same way Eyeling does, so agreement on the test suite is not independent conformance. Still, the experiment suggests that once a reference implementation exists, porting to new platforms via LLM synthesis is feasible and cheap. Independent LLM-generated N3 reasoners have since appeared in Rust [31], Python [32] and Go [33], each with different features and performance.

## 6. Conclusion

We position the role of interoperable rule engines, like Eyeling, as a solution for symbolic guardrails in agent spaces, as a solution for automating legal compliance, and as a solution for scaling up assertions across data sources using heterogeneous data models. This requires a rule language with extensive tests, for which in this case we’ve chosen the RDF-native N3 language, and made sure running those rules becomes inspectable, portable, and testable across implementations. The results presented in this paper indicate that such rule engines can be built for the browser, validated against a shared conformance suite, and even synthesized with substantial assistance from LLMs when the underlying specification is precise enough. The blue-sky vision emerging from this research is that GenAI does not deprecate Semantic Web infrastructure, but become a means to accelerate the construction of the deterministic components needed for building trustworthy data ecosystems.

## Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-5.6 order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication’s content.

## References

- [1] B. De Meester, J. A. Rojas, F. Ongenaes, P. Colpaert, R. Verborgh, Tackling the write-to-read web of data with trustflows, in: Published in the Proceedings of the Industry, Doctoral Consortium, Posters and Demos at ISWC 2025, 2025. URL: <https://ceur-ws.org/Vol-4085/paper37.pdf>.
- [2] W. Slabbinck, J. Rojas Meléndez, B. Esteves, P. Colpaert, R. Verborgh, Interoperable Interpretation and Evaluation of ODRL Policies, in: E. Curry, M. Acosta, M. Poveda-Villalón, M. van Erp, A. Ojo, K. Hose, C. Shimizu, P. Lisena (Eds.), The Semantic Web, Springer Nature Switzerland, Cham, 2025, pp. 192–209. doi:10.1007/978-3-031-94578-6\_11.
- [3] R. Iannella, S. Villata, ODRL Information Model 2.2, 2018. URL: <https://www.w3.org/TR/odrl-model/>.
- [4] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data Privacy Vocabulary (DPV) – Version 2.0, in: G. Demartini, K. Hose, M. Acosta, M. Palmonari, G. Cheng, H. Skaf-Molli, N. Ferranti, D. Hernández, A. Hogan (Eds.), The Semantic Web – ISWC 2024, Springer Nature Switzerland, Cham, 2024, pp. 171–193. doi:10.1007/978-3-031-77847-6\_10.
- [5] T. Berners-Lee, cwm – a general purpose data processor for the semantic web, <https://github.com/linkedata/swamp>, 2022. Accessed: 2026-04-29.

- [6] J. De Roo, eyereasoner/eye: Euler yet another proof engine, <https://github.com/eyereasoner/eye>, 2026. Accessed: 2024-05-25.
- [7] W. Van Woensel, william-vw/jen3, <https://github.com/william-vw/jen3>, 2025. Java. Accessed: 2026-04-29.
- [8] A. Polleres, D. Rincon-Yanez, A. Anjomshoaa, D. Dobriy, E. Sallinger, From Data Spaces to Agent Spaces, Dubrovnik, Croatia, 2026. URL: [https://aic.ai.wu.ac.at/~polleres/publications/poll-etal-2026SDS\\_ESWC\\_Agent\\_Spaces.pdf](https://aic.ai.wu.ac.at/~polleres/publications/poll-etal-2026SDS_ESWC_Agent_Spaces.pdf).
- [9] N. Fornara, A. Chiappa, M. Colombetti, Using Semantic Web Technologies and Production Rules for Reasoning on Obligations and Permissions, in: M. Lujak (Ed.), *Agreement Technologies*, Springer International Publishing, Cham, 2019, pp. 49–63. doi:10.1007/978-3-030-17294-7\_4.
- [10] T. Pellegrini, V. Mireles, S. Steyskal, O. Panasiuk, A. Fensel, S. Kirrane, Automated rights clearance using semantic web technologies: The DALICC framework, in: T. Hoppe, B. Humm, A. Reibold (Eds.), *Semantic Applications: Methodology, Technology, Corporate Use*, Springer, Berlin, Heidelberg, 2018, pp. 203–218. doi:10.1007/978-3-662-55433-3\_14.
- [11] M. De Vos, S. Kirrane, J. Padget, K. Satoh, ODRL policy modelling and compliance checking, in: P. Fodor, M. Montali, D. Calvanese, D. Roman (Eds.), *Rules and Reasoning*, Springer International Publishing, Cham, 2019, pp. 36–51. doi:10.1007/978-3-030-31095-0\_3.
- [12] L. T. van Binsbergen, M. G. Kebede, J. Baugh, T. van Engers, D. G. van Vuurden, Dynamic generation of access control policies from social policies, *Procedia Computer Science* (2022). doi:10.1016/j.procs.2021.12.221.
- [13] F. Gandon, G. Governatori, S. Villata, Normative requirements as linked data, in: *Legal Knowledge and Information Systems*, IOS Press, 2017, pp. 1–10. doi:10.3233/978-1-61499-838-9-1.
- [14] E. Francesconi, G. Governatori, Patterns for legal compliance checking in a decidable framework of linked open data, *Artificial Intelligence and Law* 31 (2023) 445–464. doi:10.1007/s10506-022-09317-8.
- [15] L. Robaldo, F. Pacenza, J. Zangari, R. Calegari, F. Calimeri, G. Siragusa, Efficient compliance checking of RDF data, *Journal of Logic and Computation* 33 (2023) 1753–1776. doi:10.1093/logcom/exad034.
- [16] P. Hochstenbach, B. Esteves, R. Verborgh, Automated policy negotiation: A four-course meal, in: *Proceedings of the OPAL Workshop*, volume 3977 of *CEUR Workshop Proceedings*, Portorož, Slovenia, 2025. URL: <https://ceur-ws.org/Vol-3977/OPAL2025-2.pdf>.
- [17] B. Esteves, V. Rodríguez-Doncel, Analysis of ontologies and policy languages to represent information flows in GDPR, *Semantic Web* 15 (2024) 709–743. doi:10.3233/SW-223009.
- [18] E. Molino-Peña, J. De Roo, B. Esteves, J. M. García, A. Ruiz-Cortés, Detecting Contractual Risk in ODRL Policies Using DPV, in: *Accepted to be Published in the Proceedings of the 2nd ODRL And Beyond: Practical Applications And Challenges For Policy-Based Access And Usage Control (OPAL 2026) workshop, co-located with the Extended Semantic Web Conference (ESWC 2026)*, 2026.
- [19] Slabbinck, Wout and Rojas Melendez, Julian Andres and Gonçalves Crisóstomo Esteves, Beatriz and Verborgh, Ruben and Colpaert, Pieter, May the FORCE be with you? A framework for ODRL rule compliance through evaluation, in: Chaves-Fraga, D. and Heibi, I. and Garijo, D. and Collarana, D. and Salatino, A. and Vahdati, S. (Ed.), *SEMANTICS-PDWT 2025*, volume 4064, CEUR, 2025, p. 6. URL: <https://ceur-ws.org/Vol-4064/>.
- [20] R. Zhao, J. Zhao, Perennial Semantic Data Terms of Use for Decentralized Web, in: *Proceedings of the ACM Web Conference 2024, WWW '24*, Association for Computing Machinery, New York, NY, USA, 2024, pp. 2238–2249. URL: <https://dl.acm.org/doi/10.1145/3589334.3645631>. doi:10.1145/3589334.3645631.
- [21] E. Molino-Peña, W. Slabbinck, J. M. García, A. Ruiz-Cortés, B. Esteves, Automated Validation of ODRL Policies for Usage Control and Data Spaces | OpenReview, 2026. URL: <https://openreview.net/forum?id=IKYFFh0qab#discussion>.
- [22] W. Van Woensel, D. Arndt, P.-A. Champin, D. Tomaszuk, G. Kellogg, Notation3 language, <https://w3c.github.io/N3/spec/>, 2026. Accessed: 2026-09-10.

- [23] P. Hochstenbach, notation3tests, <https://codeberg.org/phochste/notation3tests>, 2026. Accessed: 2026-04-29.
- [24] W. Van Woensel, P. Hochstenbach, Notation3 builtin functions, <https://w3c-cg.github.io/n3Builtins/>, 2026. Accessed: 2026-09-10.
- [25] D. Arndt, P.-A. Champin, Notation3 semantics, <https://w3c.github.io/N3/reports/20230703/semantics.html>, 2024. Accessed: 2024-12-20.
- [26] T. M. Osmun, WellnessRules: N3 Implementation Through the Euler Engine, 2009. URL: <https://web.archive.org/web/20220119222608/http://responder.ruleml.org/WellnessRules/files/WellnessRulesN3-2009-11-10.pdf>.
- [27] Y. Guo, Z. Pan, J. Heflin, LUBM: A benchmark for OWL knowledge base systems, *Journal of Web Semantics* 3 (2005) 158–182. URL: <https://www.sciencedirect.com/science/article/pii/S1570826805000132>. doi:10.1016/j.websem.2005.06.005.
- [28] S. Liang, P. Fodor, H. Wan, M. Kifer, OpenRuleBench: an analysis of the performance of rule engines, in: *Proceedings of the 18th international conference on World wide web, WWW '09*, Association for Computing Machinery, New York, NY, USA, 2009, pp. 601–610. URL: <https://dl.acm.org/doi/10.1145/1526709.1526790>. doi:10.1145/1526709.1526790.
- [29] R. P. Gabriel, *Performance and Evaluation of LISP Systems*, The MIT Press, 1985. URL: <https://direct.mit.edu/books/book/3870/Performance-and-Evaluation-of-LISP-Systems>. doi:10.7551/mitpress/5298.001.0001.
- [30] P. Hochstenbach, Tango: a C++20 spec compliant Notation3 reasoner, 2026. URL: <https://doi.org/10.5281/zenodo.22140562>. doi:10.5281/zenodo.22140562.
- [31] J. De Roo, eyereasoner/eyeron, 2026. URL: <https://github.com/eyereasoner/eyeron>, original-date: 2026-06-30T21:51:57Z.
- [32] P. Colpaert, pyling-n3, 2026. URL: <https://pypi.org/project/pyling-n3/>.
- [33] P. Hochstenbach, romeo, 2026. URL: <https://patrickhochstenbach.net/git/phochste/romeo>.