Semantic Web 0 (0) 1
IOS Press

RDF Graph Validation Using Rule-Based Reasoning

Ben De Meester [a,*], Pieter Heyvaert [a], Dörthe Arndt [a], Anastasia Dimou [a], and Ruben Verborgh [a]
[a] *Ghent University – imec – IDLab, Department of Electronics and Information Systems,*
Technologiepark-Zwijnaarde 122, 9052 Ghent, Belgium
E-mails: ben.demeester@ugent.be, pheyvaer.heyvaert@ugent.be, doerthe.arndt@ugent.be,
anastasia.dimou@ugent.be, ruben.verborgh@ugent.be

Editor: Axel Polleres, Vienna University of Economics and Business, Vienna, Austria
Solicited reviews: Jose Emilio Labra Gayo, University of Oviedo, Oviedo, Spain; Ognjen Savković, Free University of Bozen-Bolzano, Bolzano, Italy; Simon Steyskal, Vienna University of Economics and Business, Siemens AG Österreich, Vienna, Austria; Two anonymous reviewers

Abstract. The correct functioning of Semantic Web applications requires that given RDF graphs adhere to an expected shape. This shape depends on the RDF graph and the application's supported entailments of that graph. During validation, RDF graphs are assessed against sets of constraints, and found violations help refining the RDF graphs. However, existing validation approaches cannot always explain the root causes of violations (inhibiting refinement), and cannot fully match the entailments supported during validation with those supported by the application. These approaches cannot accurately validate RDF graphs, or combine multiple systems, deteriorating the validator's performance. In this paper, we present an alternative validation approach using rule-based reasoning, capable of fully customizing the used inferencing steps. We compare to existing approaches, and present a formal ground and practical implementation "Validatrr", based on N3Logic and the EYE reasoner. Our approach – supporting an equivalent number of constraint types compared to the state of the art – better explains the root cause of the violations due to the reasoner's generated logical proof, and returns an accurate number of violations due to the customizable inferencing rule set. Performance evaluation shows that Validatrr is performant for smaller datasets, and scales linearly w.r.t. the RDF graph size. The detailed root cause explanations can guide future validation report description specifications, and the fine-grained level of configuration can be employed to support different constraint languages. This foundation allows further research into handling recursion, validating RDF graphs based on their generation description, and providing automatic refinement suggestions.

Keywords: Constraints, Rule-based Reasoning, Validation

1. Introduction

Semantic Web data is represented using the Resource Description Framework (RDF), forming an *RDF graph* [36]. The *quality* of an RDF graph – its "fitness for use" [135] – heavily influences the results of a Semantic Web *application* [84]. An RDF graph's fitness for use depends on its *shape*, i.e., the RDF graph itself and the application's *supported entailments* of that RDF graph. For example, some applications support inferring `rdfs:subClassOf` entailments [29], whereas other applications require the RDF graph to explicitly contain all classifying triples (i.e., `rdfs:subClassOf` entailment is not supported).

RDF graphs are *validated* by assessing their adherence to a set of *constraints* [82], and different applications (i.e., different *use cases*) specify different sets of constraints. Via validation, we discover (portions of) RDF graphs that do not conform to these constraints, i.e., the *violations* that occur. These violations guide the user to the resources and relationships related to the constraints. *Refining* these resources and relationships results in an RDF graph of higher quality [47], thus,

---

[*] Corresponding author. E-mail: ben.demeester@ugent.be.

RDF graph validation is an important element for the correct functioning of Semantic Web applications.

1.1. Validation problems

Let us consider the following example: an RDF graph containing people and their birthdates is validated. The use case dictates the set of constraints and the supported entailments. Specifically, we validate formula (1)[1], with a relevant ontology represented in formula (2).

```
:Bob :firstname "Bob" ;                              (1)
     :birthdate "1970-01-01"^^xsd:date .
:birthdate rdfs:domain :Person .                     (2)
:Bob a :Person .                                     (3)
```

Problem 1 (P1): finding the root causes of violations. For example, a use case dictates that every resource should have either a firstname and lastname, or a nickname. This constraint, $c_{compound}$, is thus a *compound* of several constraints. When the RDF graph contains formula (1) and formula (3), `:Bob` should be marked as a resource violating $c_{compound}$. However, the RDF graph cannot be refined solely by knowing which resources violate the constraint. The *root cause* of the violation is needed: does the resource lack firstname, lastname, or nickname?

For constraint types such as compound constraints, existing validation approaches typically return the resource that violates the constraint. More detailed descriptions are typically not provided, and manual inspection is needed to discover the root cause of a violation, i.e., *why* a resource violates a constraint. Without the root cause, it is hard to (automatically) refine the RDF graph and improve its quality.

Problem 2 (P2): the number of found violations depends on the supported entailments. A mismatch between which entailments are supported during validation and which entailments are supported by the use case influences, e.g., whether formula (3) is inferred or not. Thus, either too many or too few violations can be returned [26]. This difference in number of found violations gives a biased idea of the real quality of the validated RDF graph.

Too many violations: formula (2) specifies the domain of `:birthdate`. Let us validate that “every resource in the RDF graph that has a birthdate, is a person” given formula (1). When the entailments of formula (2) are not supported, this would result in a violation: formula (3) is missing in the RDF graph. However, when the entailments of formula (2) are supported, we can infer formula (3), and no violation is returned.

Too few violations: Let us validate that “every person in the RDF graph adheres to constraint $c_{compound}$” given formula (1). Formula (3) is not explicitly stated and the entailments of formula (2) are not supported. No violations are found: `:Bob` is not explicitly classified as a `:Person`, thus `:Bob` is not targeted by $c_{compound}$. However, supporting those entailments can create new statements to be validated, and lead to new violations. For example, by inferring formula (3) using formula (2), `:Bob` is targeted by – and violates – $c_{compound}$. Such violations are not found in the original RDF graph, but discovered due to the supported entailments.

Customizing the *set of inferencing steps* during validation (e.g., whether `rdfs:domain` entailments are supported or not) allows to match the entailments supported by the use case with those of the validation approach. However, support for customizable inferencing steps is limited. When a fixed set (or no set) of inferencing steps is supported, a separate reasoning process is needed to infer unsupported entailments, and edge cases handling this fixed set cannot be validated accurately. For example, let us look at the W3C recommended Shapes Constraint Language (SHACL): a language for validating RDF graphs against a set of constraints [78]. SHACL specifies a fixed set of inferencing steps during validation, namely, `rdfs:subClassOf` entailment when targeting resources of a certain class. Thus, one cannot validate, e.g., whether an RDF graph explicitly contains all triples that link resources to all their classes given a set of `rdfs:subClassOf` axioms, as `rdfs:subClassOf` triples are inferred by a conform SHACL validator[2]. RDF graphs that do not contain all classifying triples will be valid according to SHACL validators, however, they are handled poorly by applications that do not support `rdfs:subClassOf` entailment.

Problem 3 (P3): Combining validation with a reasoning preprocessing step decreases performance Entailments can be inferred by performing reasoning as a preprocessing step prior to validation [26], thus com-

---

[1] For the remainder of the paper, empty prefixes denote the fictional schema http://example.com/, other prefixes are conform with the results of https://prefix.cc.

[2] For a detailed example, please see https://idlabresearch.github.io/validatrr/blog/2019/09/shacl-subclassof.html.

bining multiple systems. The resulting RDF graph then explicitly contains all supported entailments, given that the reasoner can be configured to only infer the entailments that are supported by the use case. The number of found violations is then accurate with respect to the use case (solving P2). However, this requires a sequence of independent systems. Thus, the preprocessing step possibly produces entailments not relevant for validation [26]. This independent generation of unnecessary entailments can decrease the performance compared to a single validation system. More, due to this sequence of independent systems, finding the root causes involves investigating the results of both systems: the validator who detects violations, and the reasoner who infers entailments.

1.2. Hypotheses

To solve aforementioned observed validation problems, we pose following hypotheses.

Hypothesis 1 Root causes can be explained more accurately compared to existing validation approaches when using a logical framework that can be configured declaratively.

Hypothesis 2 A more accurate number of violations are found compared to existing validation approaches when supporting a custom set of inferencing steps.

Hypothesis 3 A validation approach supporting more accurate root cause explanations and a custom set of inferencing steps can support an equivalent number of constraint types compared to existing approaches.

Hypothesis 4 A validation approach supporting a custom set of inferencing steps is faster than an approach including the same inferencing as a preprocessing step.

1.3. Contributions

In this paper, we propose an approach for RDF graph validation that uses a rule-based reasoner as its underlying technology. Rule-based reasoners can generate a proof stating which rules were triggered for which returned violation. Thus, the root causes of violations can be accurately explained (*solving P1*).

A validation approach using rule-based reasoning natively support the inclusion of a custom set of inferencing steps by adding custom rules. The supported entailments during validation can thus be matched to the entailments supported by the use case, and the validation returns an accurate number of found violations (*solving P2*).

Moreover, rule-based reasoners only need a single language to declare both the constraints and the set of inferencing rules, and only a single system to execute the validation. Compared to a combination of a reasoner and a validation system, this approach does not lead to the generation of entailments unnecessary to the validation step, making it potentially faster than including an inferencing preprocessing step (*solving P3*).

Our contributions are as follows.

i An analysis of existing validation approaches and comparison to a rule-based reasoning approach.
ii A formal ground for using rule-based reasoning for validation.
iii An application of that formal ground by providing an implementation using N3Logic [20] to define the inferencing and validation rules, executed using the EYE reasoner [131], supporting general constraint types as described by Hartmann et al. [62, 64].
iv An evaluation of our approach, positioning it within the state of the art by functionally validating the hypotheses and comparing the validation speed.

We validated that (a) the formal logical proof explains the root cause of a violation more detailed than the state of the art; (b) an accurate number of violations is returned by using a custom set of inferencing rules up to at least OWL-RL complexity and expressiveness; (c) the number of supported constraint types is equivalent to existing validation approaches; and (d) our implementation is faster than a combined system, and faster than an existing validation approach when RDF graphs are smaller than one hundred thousand triples.

The remainder of the paper is organized as follows. We start by giving an overview of the state of the art (Section 2), after which we position and compare rule-based reasoning as validation approach (Section 3). Then, we discuss the logical requirements (Section 4) and apply them to achieve a practical implementation (Section 5). Finally, we evaluate our proposed approach (Section 6) and summarize our conclusions (Section 7).

2. State of the art

In this work, we propose an alternative validation approach using rule-based reasoning. We first provide a background on validation and reasoning in Section 2.1. Then, we give an overview of existing validation approaches in Section 2.2, and of related vocabularies and ontologies in Section 2.3. We conclude with an overview of general constraint types in Section 2.4, which allows us to functionally compare validation approaches. Our categorization is derived from the general quality surveys of Zaveri et al. [135], Ellefi et al. [49], and Tomaszuk [127], and from the "Validating RDF Data" book [84]. The related works are extended with recent works published in, among others, the major Semantic Web conferences (ESWC and ISWC), and the major Semantic Web journals (Journal of Web Semantics and Semantic Web Journal).

2.1. Background

Validation Data quality can be assessed by employing a set of *data quality assessment metrics* [22]. Quality assessment for the Semantic Web – and more specifically, for Linked Data – spans multiple dimensions, further categorized in *accessibility*, *intrinsic*, *trust*, *dataset dynamicity*, *contextual*, and *representational* dimensions [135]. Validating an RDF graph directly relates to intrinsic quality dimensions, as defined by Zaveri et al. [135]: (i) independent of the user's context, and (ii) checking if information correctly and compactly represents the real world data and is logically consistent in itself, i.e., the graph's adherence to a certain schema or shape. In this paper, we specifically focus on RDF graph validation, i.e., the intrinsic dimensions.

Validation of an RDF graph can be automated by using a set of *test cases*, each assessing a specific *constraint* [82]. *Violations* of those constraints are then indicated when a validation returns negative results. Validation is typically achieved following Closed World Assumption (CWA): what is not known to be true must be false. For example, a validation assesses for a specific RDF graph if all objects linked via the predicate `schema:birthdate` are a valid `xsd:date`, or if all subjects and objects linked via the predicate `foaf:knows` are explicitly listed to be of type `:Human`. Negative results are returned, indicating violations.

Reasoning Ontologies are prevalent in the Semantic Web community to represent the knowledge of a domain. *Ontology languages* are used to annotate asserted facts (*axioms*). Examples include RDF Schema (RDFS) [29] and the Web Ontology Language (OWL) [68]. *Reasoning* on top of these axioms is achieved, as the calculus of the used *logic* specifies a set of *inferencing steps*, inferring logical consequences (*entailments*) from these axioms [45]. Logics for the Semantic Web – given the open nature of the Web – typically follow the Open World Assumption (OWA): what is not known to be true is simply unknown.

Semantic Web reasoners are typically *description logic-based reasoners* supporting OWL-DL or subprofiles such as OWL-QL [94], or *rule-based reasoners* [100]. Description logic-based reasoners are typically optimized for specific description logics, such as KAON2[3] for $\mathcal{SHIQ}$ and FaCT++[4] for $\mathcal{SROIQ}$. Rule-based reasoners typically follow two types of inferencing algorithms: *forward chaining* and *backward chaining* [100]. Whereas forward chaining tries to infer as much new information as possible, backward chaining is goal-driven: the reasoner starts with a list of goals and tries to verify if there are statements and rules available that support any of these goals [100]. The employed rules define the logic followed by rule-based reasoners such as EYE [131] or cwm [17]. Whereas description logic-based reasoners have (optimized) inferencing steps for, e.g., `rdfs:subClassOf` and other RDFS or OWL constructs embedded, rule-based reasoners commonly rely on the general "implies" construct. Each rule specifies "A implies B", where both the *antecedent* "A" and the *consequence* "B" can consist of statements [100]. Certain constructs such as `rdfs:subClassOf` can be translated into one or more rules[5].

There is a clear distinction between ontologies and the constraint set for RDF graph validation: ontologies focus on the representation of a domain, whereas RDF graph validation checks whether the resources of that graph conform to a desired schema [84]. It is not required that the representation of a domain aligns with the schema for validation. However, they can complement each other. The usage of ontologies prescribes a set of inferencing steps, for example, the FOAF ontology declares the `rdfs:range` of the `foaf:knows` predicate as `foaf:Person` [30]. Whether these inferencing steps are taken into account during validation or not, influences the number of found violations [26].

2.2. Validation Approaches

In this section, we discuss RDF graph validation approaches. Tools and surveys that cover quality dimensions other than the intrinsic dimensions such as accessibility or representational dimensions are out of scope. We discuss the approaches roughly in chronological

---

[3] http://kaon2.semanticweb.org/

[4] http://owl.cs.manchester.ac.uk/tools/fact/

[5] http://eulersharp.sourceforge.net/#theories

order: *hard-coded*, using *integrity constraints*, *query-based*, and using a *high-level language*. Except from *hard-coded* systems, these validation approaches propose or use some kind of declarative means to describe RDF graph constraints.

2.2.1. *Hard-coded*

Hard-coded systems are a black box where the business logic lies within the code base: the implementation embeds both description and validation of constraints. Hogan et al. analyzed common quality problems both for publishing and intrinsic quality dimensions [69], providing an initial set of best practices [70]. Efforts focus on a limited set of configurable settings (i.e., turning constraint rules on or off) [90].

2.2.2. *Integrity Constraints*

For these validation approaches (so-called "logic-based approaches"), the axioms of vocabularies and ontologies used by the validated RDF graph are *interpreted as integrity constraints* [93, 101, 124]. For example, disjointness forces a description logic-based reasoner to throw an error, which is interpreted as a violation. To combine CWA typically assumed for validation with OWA assumed in ontology languages, alternative semantics for these ontology languages are proposed. The underlying technology used is a description logic-based reasoner or a SPARQL endpoint.

Description logic-based reasoner Motik et al. [93] propose semantic redefinitions, where a certain subset of axioms are designated as constraints. To know which alternative semantics for OWL apply, constraints have to be marked as such. They propose to integrate their implementation with KAON2. Furthermore, custom integrity constraints for Wordnet have been verified using Protégé [95] with FaCT++ [33].

SPARQL endpoint Tao et al. [124] propose using OWL expressions with Closed World assumption and a weak variant of Unique Name assumption to express integrity constraints. OWL semantics are redefined, without being explicitly stated as such during validation. They use SPARQL [2] for axioms described in RDF, RDFS, and OWL [124], e.g., using SPARQL property paths to simulate `rdfs:subClassOf` entailment. Tao et al. work in a general OWL setting, where their approach is sound but not complete. In an RDF setting the approach is both sound and complete, as there is only a single model that needs to be considered [101]. This implementation is incorporated into Stardog ICV [107]. Patel-Schneider separates validation into integrity constraints and Closed World recognition [101], showing that RDF and RDFS entailment can be implemented for both by translation to SPARQL queries.

2.2.3. *Query-based*

In query-based approaches, constraints are described and interpreted similar to SPARQL queries [64, 102]: only RDF graphs whose structure is compatible with the defined structure are returned. These approaches use an embedded or external SPARQL endpoint as underlying technology.

CLAMS [51] is a system to discover and resolve Linked Data inconsistencies. They define a violation as a minimal set of triples that cannot coexist. The system identifies all violations by executing a SPARQL query set. Knublauch et al. propose the SPARQL Inference Notation (SPIN) [80]: a SPARQL-based rule and constraint language. The SPARQL query is described using RDF statements instead of using the original SPARQL syntax. Kontokostas et al. [82] propose Data Quality Test Patterns (DQTP): tuples of typed pattern variables and a SPARQL query template to declare test case patterns. The validation framework that validates these DQTPs is called RDFUnit. The DQTPs are transformed into SPARQL queries, where every SPARQL query is a test case. RDFUnit additionally allows automatically generated test cases, depending on the used schema.

RDFUnit is also used to validate Linked Data generation rules in the RDF Mapping Language (RML) [46], by manually defining different DQTPs to target the generation description instead of the generated RDF graph [47]. This means the RDF graph can be validated before any data is generated, as the generation description reflects how the RDF graph will be formed.

2.2.4. *High-level language*

These approaches use a terse high-level language specifically designed to describe constraints for validation [84]. These languages are independent of underlying technologies, and alternative implementation strategies can be devised. We first discuss initial high-level languages, after which we discuss high-level languages with wide adoption from the community: ShEx and SHACL.

Description Set Profiles (DSP) [97] define a set of constraints using Description Templates, targeted specifically to Dublin Core Application Profiles, and implemented using SPIN [25]. Other high-level languages to describe constraints include OSLC Resource Shapes [117] – part of IBM Resource Shapes – and RDF Data Descriptions [53]. Luzzu [42] uses a custom declarative constraint language (Luzzu Quality Metric Language, LQML). Any metric that can be expressed in a SPARQL query can be defined using LQML. Moreover, quality dimensions other than the intrinsic dimensions are also expressible using LQML. Luzzu supports basic metrics and custom JAVA code allowing users to implement custom metrics.

ShEx Shape Expressions (ShEx) [111, 112] is a structural schema language which can be used for RDF graph validation. The grammar of ShEx is inspired by Turtle and RelaxNG, its semantics are well-founded, and its complexity and expressiveness are formalized [23, 123]. ShEx provides an extension point to handle advanced constraints via Semantic Actions, which allows to evaluate a part of the validated RDF graph using a custom function.

SHACL The Shapes Constraint Language (SHACL) is the W3C Recommendation for validating RDF graphs against a set of constraints [78]. The core of SHACL is independent of SPARQL, which promotes the development of new algorithms and approaches to validate RDF graphs [85]. The original specification does not include a denotational semantics such as ShEx, however, the recent work of Corman et al. propose a concise formal semantics for SHACL's core constraint components, and a way of handling recursion in combination with negation [34]. Advanced features of SHACL include SHACL Rules (to derive inferred triples from the validated RDF graph) and SHACL Functions (to evaluate a part of the validated RDF graph using a custom function) [81].

2.3. *Validation reports*

Validation reports handle identification of which data quality dimensions are assessed in general, and the representation of violations in particular.

To identify data quality dimensions, Radulvic et al. extended the Dataset Quality Ontology (daQ) [43] to include all data quality dimensions as identified by Zaveri et al. [135], leading to the Data Quality Vocabulary [113]. This allows the comparison of data quality dimension coverage of different frameworks.

The violations report itself allows to distribute and compare the violations found in an RDF graph, and can refer to the dimension specifications using aforementioned general vocabularies. For example, the Quality Problem Report Ontology assembles detailed quality reports for all data quality dimensions [42]. The Reasoning Violations Ontology (RVO) is used to represent integrity constraint violations [28], and Kontokostas et al [82] use the RDF Logging Ontology[6] (RLOG) to describe RDFUnit's violation results. Both ShEx and SHACL provide violation report descriptions, with means to specify the violating resources, using a ShapeMap [112] and a Focus node [78], respectively.

2.4. *Constraint types*

Hartmann *né* Bosch et. al identify eighty-one general *constraint types* [27, 64]. These constraint types are an abstraction of specific constraints, independent of the constraint language used to describe them. A constraint type can be defined in different ways. For example, the *property domain* constraint type specifies that resources that use a specific property should be classified via a specific class, e.g., all resources using the `:birthdate` property that are not classified as a `:Person` are violating resources. Using RDFS [29], the property domain constraint type can be assessed by interpreting `rdfs:domain` as an integrity constraint. Using SHACL, this can be achieved by defining a `sh:property` with `sh:class` for a `sh:targetSubjectsOf` shape [78].

Moreover, Hartmann et al. provide a logical underpinning stating the requirements for a validation approach to support all constraint types [26]. For thirty-five out of eighty-one constraints types (43.2%), reasoning (up to OWL-DL expressiveness) can improve the validation: without reasoning, either too many or too few violations can be returned.

3. Comparative analysis

Different types of validation approaches are proposed in the state of the art. The most prominent approaches are *hard-coded*, based on *integrity constraints*, *query-based*, and using *high-level languages*. In this section, we compare them with our proposed *rule-based reasoning* approach. Our analysis is summarized in Table 1.

We adapt the framework presented by Pauwels et al. [105], which introduces comparative factors of key implementation strategies for compliance checking applications. We adjust these factors with respect to the validation problems identified in Section 1.1. We generalize the factors *time*, *customization*, and *inferencing steps*, and introduce *explanation* and *reasoning preprocessing* as validation-specific factors.

---

[6] http://persistence.uni-leipzig.org/nlp2rdf/ontologies/rlog#

Table 1

Comparing the prominent validation approaches with rule-based reasoning, using factors *explanation*, *time*, *customization*, *inferencing steps*, and *reasoning preprocessing*. The *time* row indicates which approaches' execution time is influenced due to the reasoning preprocessing using an asterisk. The asterisk in the *inferencing steps* row indicates that approaches based on integrity constraints cannot combine with a custom set of inferencing steps that overlaps with the integrity constraints, as their semantics are redefined.

Approach	**Hard-coded**	**Integrity constraints**	**Query-based**	**High-level language**	***Rule-based reasoning***
Explanation	No	Limited / Yes	Limited	Limited	Yes
Time	Short*	Long	Long*	Short*	Long
Customization	Limited	Limited	Open	Open	Open
Inferencing steps	No / Limited	Yes*	Limited / Yes	Limited / Yes	Yes
Reasoning preprocessing	Yes	Limited	Yes	Yes	N/A

Explanation The explanation as to why a certain violation occurs (i.e., the root cause). The more specific a validator can explain, the easier it is to (automatically) refine the RDF graph and improve its quality. Existing approaches typically have the means to explain violations up to the level of which resource violates which constraint. Explanations of *hard-coded* approaches either need to be explicitly implemented, or are provided by inspecting the code base. When using *integrity constraints*, approaches exist for resolving inconsistencies. These approaches perform some sort of root cause analysis, but are usually targeted at refining the axioms of the ontologies themselves [60]. It is not a standard feature to produce proofs of the results of description logic-based reasoners [99]. In a *query-based* approach, the used SPARQL endpoint returns bindings [2]. In the case of validation, it returns the violating resources, without additional explanation. *High-level languages* can have mechanisms to additionally include the violating resources in the validation report. For example, ShEx and SHACL provide ShapeMaps [112] and Focus nodes [78], respectively. SHACL's Focus nodes can further specify which predicate and object cause the violation, except for, e.g., compound constraints. Using *rule-based reasoning* allows the generation of a logical proof, as rule-based reasoning relies on a general "implies" construct to describe rules, and rule-based reasoners typically do not contain description logic optimizations. Such a logical proof declares which rules were triggered to arrive at a certain conclusion, giving a precise explanation for the root causes of constraint violations. Where existing approaches typically have the means to explain violations up to the level of which resource violates which shape, a logical proof can provide a more detailed explanation.

Time The time needed to execute the validation: short versus long. Typically, specialized approaches allow for optimizations, making them faster than general approaches. *Hard-coded* is usually the fastest and needs the shortest processing time, followed by systems that use *high-level languages*: both can be optimized for validation tasks. The other approaches (using *integrity constraints*, *query-based*, and *rule-based reasoning*) are typically built using an underlying existing technology (description logic-based reasoners, SPARQL endpoints, and rule-based reasoners, respectively). They are not built (or optimized) for validation tasks. This makes them independent of the constraint language, but can also slow down the validation. The total execution time of validation approaches depends on whether a reasoning preprocessing step to include additional inferencing steps is required or not. Using rule-based reasoning is thus potentially slower than existing approaches, however, it does not require inclusion of reasoning preprocessing.

Customization The extent of customization each type of approach enables. Typically, ease of customization is improved by using a declarative language. Customization of a *hard-coded* system requires development effort, as the business logic is embedded within the code. Other approaches rely on declarations to customize the validation. Declarations are decoupled, i.e., independent of the tool's implementation. Thus, they can be shared and easier customized to a certain use case. Description logic-based reasoners used to identify *integrity constraints* are typically optimized for description logics such as OWL-QL and OWL-DL. Customization is limited to the description logic that the reasoner is optimized for. *Query-based* approaches allow customization by defining additional SPARQL queries and registering custom functions [61]. Systems using *high-level languages* are customized using the declarations as specified by the used language. The adoption of ShEx and SHACL shows that these languages provide sufficient customization. The extension mechanisms of these languages such as Semantic Actions [112] and SHACL Advanced Features [81], respectively, allow to customize the validation even further. Using *rule-based reasoning* allows customization by adding and removing rules. As opposed to existing approaches, users can customize both the constraint types and the set of inferencing steps within the same declarative language.

Inferencing steps Whether the validation approach supports a (custom) set of inferencing steps. *Hard-coded* systems can support a fixed set of inferencing steps, but this set cannot be inspected or altered without investigating the code base. Approaches that use *integrity constraints* for validation propose alternative semantics of commonly agreed upon ontology languages to include, among others, some form of CWA [93, 124]. This leads to ambiguity in the Semantic Web as an existing, globally agreed upon logic, is changed [7]. It is not possible to combine such validation with a (custom) set of inferencing steps within a description logic: the same inferencing step has different semantics whether it is used for validation or for inferring new statements. SPARQL endpoints used for *query-based* approaches can support up to OWL-RL reasoning [77], or support up to RDF and RDFS entailment via translation of the SPARQL queries using property paths [124]. *High-level languages* such as SHACL allow specifying the entailment regime used [58]: SHACL validators may operate on RDF graphs that include entailments using the `sh:entailment` property [78]. Furthermore, SHACL Rules [81] can be used to a certain extent to generate inferred statements during validation. By design, *rule-based reasoning* allows inclusion of a set of additional (custom)) inferencing rules [100]. Whereas existing approaches mostly allow configuration to support, e.g., a specific entailment regime, the customization of the set of inferencing steps is more fine-grained for rule-based reasoners. This can increase complexity, but also allows catering the validation to use cases that depend on a specific set of inferencing steps. The importance of such use cases is evidenced by the fact that SHACL Rules is proposed as an advanced feature to the SHACL specification [81].

Reasoning preprocessing Existing approaches have no support for including a custom set of inferencing steps, propose alternative semantics, or allow a specific entailment regime. By including a reasoning step as preprocessing step to these approaches (see Fig. 1.1), the entailments valid during validation can be matched with the entailments valid for the use case, even when that use cases requires a custom set of inferencing steps [26]. First, a reasoner – optionally, hence the dashed line – infers all valid entailments of the original RDF graph (Fig. 1.1, *Reasoner*), taking into account the axioms of the relevant ontologies and vocabularies (*Axioms*). Then, the newly generated RDF graph (*RDF graph**) is validated with respect to the specified constraints (Fig. 1.1, *Validator*).

By using a preprocessed inferred RDF graph, multiple systems (i.e., the reasoner and the validator) need to be combined, configured, and maintained. This separates concerns, however, this also means that different languages may need to be learned and combined for specifying the inferencing steps and constraints. As these multiple systems are not aligned, the reasoner could infer a large number of new triples that are irrelevant to the defined constraints, which could lead to bad scaling (Fig. 1.1, *RDF graph**). Also, explaining the violation is hindered. Even when the reasoner can differentiate between the original triples and the inferred triples, finding the root causes involves investigating the output of both systems: the validator detecting the violations, and the reasoner inferring the supported entailments.

Reasoning preprocessing is not required when using *rule-based reasoning*. The set of inferencing steps and the set of constraints can be defined using the same declaration (Fig. 1.2, *Inferencing rules* and *Constraints**), and executed simultaneously on the RDF graph and the axioms. Which statements need to be inferred can be optimized guided by the set of constraints, and only the output of a single system needs to be investigated to explain the found violations.

4. Logical Requirements

In this section, we discuss the logical requirements needed for RDF graph validation, and argue for using a rule-based logic.

Constraint languages need to cope with different constraint types depending on users' needs. Each constraint type implies logical requirements. The constraint types and the requirements they entail are investigated by Hartmann et al., claiming that Closed World Assumption (CWA) and Unique Name Assumption (UNA) are crucial for validation [26]. These requirements typically do not apply to logics for the Semantic Web, as data on the Web is decentralized, information is spread ("anyone can say anything about anything" [36]), and single resources can have multiple URIs. Instead, relevant logics such as OWL-DL assume OWA and in

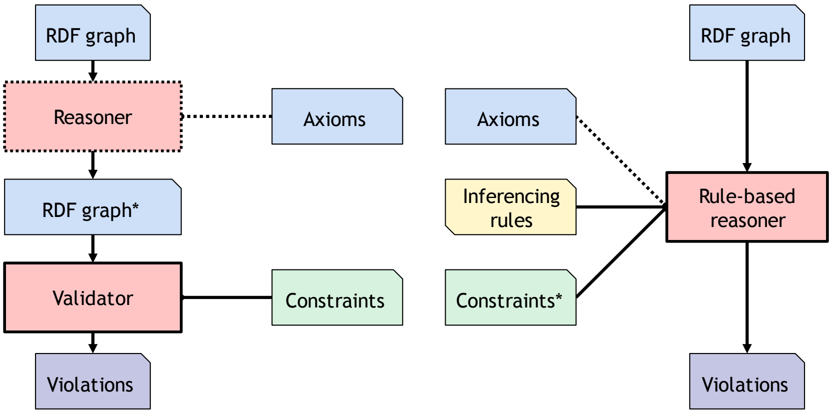

(1) Preprocessing approach (2) Rule-based reasoning approach

Figure 1. The preprocessing approach: first (optionally, hence the dashed line), a *reasoner* is used to generate intermediate data (*RDF graph**). That intermediate data is then the input data for the *Validator*. Using a rule-based reasoner only needs a single system and language to combine reasoning and validation.

general non-Unique Name Assumption [94]. Hartmann et al. emphasize the difference between reasoning and validation, and favor query-based approaches for validation. When needed, query-based approaches can be combined with reasoning (e.g., OWL-DL or OWL-QL) as a preprocessing step.

However, in this section, we show how rule-based reasoning can be used for validation in a Semantic Web context, even though this reasoning typically does not follow CWA and UNA. Specifically, we state that the requirements for using rule-based reasoning are (i) supporting Scoped Negation as Failure (SNAF) [37, 76, 108] instead of CWA (Section 4.1), (ii) containing predicates to compare URIs and literals instead of supporting UNA (Section 4.2), and (iii) supporting expressive built-ins, as validation often deals with, e.g., string comparison and mathematical calculations (Section 4.3).

4.1. Scoped Negation as Failure

Existing works claim that CWA is needed to perform validation [26, 101, 124]. Given that most Web logics assume OWA, this would require semantic redefinitions to include inferencing during validation [93], which leads to ambiguity. However, as validation copes with the *local knowledge base*, and not the entire Web, we claim Scoped Negation as Failure (SNAF) is sufficient. This is an interpretation of logical negation: instead of stating that ρ does not hold (i.e., $\neg\rho$), it is stated that reasoning fails to infer ρ within a specific *scope* [37, 76, 108]. This scope needs to be explicitly stated. As such, SNAF keeps monotonicity.

To understand the idea behind Scoped Negation as Failure, let us validate following RDF graph:

$$\texttt{:Kurt a :Researcher;} \tag{4}$$

$$\texttt{:name "Kurt01".} \tag{5}$$

We validate the constraint "every individual which is declared as a researcher is also declared as a person". This thus means a violation is returned when an individual is found during validation which is a researcher, but not a person:

$$\begin{aligned} &\forall \texttt{x} : ((\texttt{ x a :Researcher })\wedge \\ &\quad \neg(\texttt{ x a :Person })) \\ &\quad \rightarrow (\texttt{ :constraint :isViolated "true". }) \end{aligned} \tag{6}$$

As stated, this constraint cannot be tested with OWA: the knowledge base contains the triple of formula (4), but not of:

$$\texttt{:Kurt a :Person.} \tag{7}$$

The rule is more general: given its open nature, we cannot guarantee that there is no document in the entire Web which declares the triple of formula (7).

This changes if we take into account SNAF. Suppose that $\mathcal{K}$ is the set of triples we can derive (either with or without reasoning) from our knowledge base of formulas (4) and (5). Having $\mathcal{K}$ at our disposal, we can test:

$$
\begin{aligned}
&\forall \texttt{x} : (((\ \texttt{x a :Researcher}\ ) \in \mathcal{K}) \land \\
&\quad \neg((\ \texttt{x a :Person}\ ) \in \mathcal{K})) \\
&\quad \rightarrow (\ \texttt{:constraint :is :violated.}\ )
\end{aligned}
\tag{8}
$$

The second conjunct is not a simple negation, it is a negation with a certain scope, in this case $\mathcal{K}$. If we add new data to our knowledge base, e.g., the triple of formula (7), we would have a different knowledge base $\mathcal{K}'$ for which other statements hold. The truth value of formula (8) would not change since this formula explicitly mentions $\mathcal{K}$. SNAF is what we actually need for validation: we do not validate the Web in general, we validate a specific RDF graph.

4.2. *Predicates for Name Comparison*

UNA is deemed required for validation [26], i.e., every resource taken into account can only have one single name (a single URI in our case) [75]. UNA is in general difficult to obtain for the Semantic Web and Web logics due to its distributed nature: different RDF graphs can – and actually do – use different names for the same individual or concept. For instance, the URI `dbpedia:London` refers to the same place in Britain as, e.g., `dbpedia-nl:London`. That fact is even stated in the corresponding datasets using the predicate `owl:sameAs`. The usage of `owl:sameAs` conflicts with UNA and influences validation [26].

Let us look into the following example. We assume `dbo:capital` is an `owl:InverseFunctionalProperty`. Our knowledge base contains:

$$\texttt{:Britain dbo:capital :London.} \tag{9}$$

$$\texttt{:England dbo:capital :London.} \tag{10}$$

Since both `:Britain` and `:England` have `:London` as their capital and `dbo:capital` is an inverse functional property, an description logic-based reasoner would derive that

$$\texttt{:Britain owl:sameAs :England.} \tag{11}$$

This thus influences the validation result. Such a derivation cannot be made if UNA holds, since UNA explicitly excludes this possibility.

The related constraint – defined as `INVFUNC` by Kontokostas et al. [82] – specifies that each resource should contain exactly one relationship via `dbo:capital`, i.e., the capital is different for every resource. The constraint `INVFUNC` is related to `owl:InverseFunctionalProperty`, but it is slightly different: while OWL's inverse functional property refers to *the resources* that are in the domain of `dbo:capital`, the validation constraint `INVFUNC` refers to *the representation of those resources*. The RDF graph of formulas (9) and (10) thus violates the `INVFUNC` constraint. Even if our logic does not follow UNA, this violation can be detected if the logic offers predicates to compare the (string) representation of resources.

4.3. *Expressive Built-ins*

Validation often deals with, e.g., string comparison and mathematic calculations. These functionalities are widely spread in rule-based logics using *built-in functions*. While it typically depends on the designers of a logic which features are supported, there are also common standards. One of them is the Rule Interchange Format (RIF), whose aim is to provide a formalism to exchange rules in the Web [74]. Being the result of a W3C working group consisting of developers and users of different rule based languages, RIF can also be understood as a reference for the most common features rule based logics might have.

Let us take a closer look to the comparison of URIs from the previous section. `func:compare` can be used to compare two strings. This function takes two string values as input, and returns `-1` if the first string is smaller than the second one regarding a string order, 0 if the two strings are the same, and 1 if the second is smaller than the first. The example above gives:

$$
\begin{aligned}
&\texttt{("http://example.com/Britain"} \\
&\ \texttt{"http://example.com/England")} \\
&\qquad\qquad\qquad \texttt{func:compare -1.}
\end{aligned}
\tag{12}
$$

To refer to a URI value, RIF provides the predicate `pred:iri-string` which converts a URI to a string and vice versa. To enable a rule to detect whether the two URI names are equal or not, an additional function is needed: the reasoner has to detect whether the comparison's result is different from zero. That can be checked using the predicate `pred:numeric-not-equal`, which is the RIF version of ≠ for numerical values. In the example, the comparison would be `true` since $0 \neq -1$. Using these RIF built-ins, a reasoner can check the name equality between `:Britain` and `:England`, and return a violation. Whether a rule based Web logic is suited for validation highly depends on its built-ins. If it supports all RIF predicates, this can be seen as a strong indication that it is expressive enough.

5. Application

In this section, we present our approach that uses rule-based reasoning for validation. We discuss the different components and the workflow in Section 5.1, the underlying technologies in Section 5.2, and implementation in Section 5.3. We end with an example using rules in Section 5.4.

5.1. Customizable validation

Our validator consists of multiple components that can be configured by adjusting the different rule sets (Fig. 2). The execution is primarily handled using the rule-based reasoner as underlying technology.

The set of *Inferencing rules* specifies the supported entailments during validation. This set can either be a predefined set to support, e.g., RDFS entailment [29], or can be fully customized. Optionally, the relevant axioms are provided during validation. As such, the entailments supported by the use case can be matched during validation.

The set of rules forming the *Constraint translation* allows our validator to infer the general constraint types – common across existing constraint languages [27, 64] – from specific constraint descriptions. It can thus infer these types from the constraints described in a specific language such as SHACL [78]. The general constraint types are described using RDF-CV, which generalizes the constraint types into a coherent structure [27]. The purpose of RDF-CV is not to invent a new constraint language: it is a concise ontology which is deemed universal enough to describe constraints expressible by other constraint languages such as SHACL[7]. Our rule-based validator is thus constraint language-independent.

The set of rules forming the *Validation* allows our validator to infer violations on the RDF graph with all supported entailments, based on the general constraint types. This set of rules specifies how to detect each constraint type.

The set of rules forming the *Report* allows our validator to infer the resulting violations in the required format. This set can be adapted to, e.g., the SHACL report format [78].

As a result, this declarative approach is decoupled from ontology language, constraint language, and report format. When no additional rule sets are included (i.e., only the *Validation* rule set is used), this validator does not infer any entailments, only validates constraints described using RDF-CV, and returns a report in a format based on RDF-CV.

All rule sets and input data are taken into account during a single reasoner execution. As opposed to using a reasoning preprocessing step, the inferred entailments can be geared towards the specified constraints (when making use of a backward chaining reasoner), and no unnecessary entailments are produced. For example, when an axiom specifies the range of a certain path, but no constraints are related to that path, this range might not need to be inferred. Moreover, as you only have a single system, finding the root cause does not require investigation of multiple systems: the logical proof contains the complete overview of which rules were used to generate which entailments and which violations.

5.2. Used Technologies

The most important technological considerations are the rule-based web logic and reasoner in accordance with that logic.

Rule-based web logic Rule-based web logics include the Semantic Web Rule Language (SWRL) [71], the Datalog+/- framework [32] and N3Logic [20][8]. We use N3Logic as it fulfills all requirements: SWRL does not support the logical requirement SNAF[9], and the Datalog+/- framework does not support production of logical proofs. N3Logic is being actively supported and used, as evidenced by recent papers and patents [40, 106, 128, 134], and by the recently founded W3C Notation 3 (N3) Community Group fostering development, implementation, and standardization[10].

---

[7] For a detailed description of RDF-CV, we refer to the original papers [24, 27], or the source: https://github.com/boschthomas/RDF-Constraints-Vocabulary

[8] For a more thorough discussion of relevant rule languages, we refer to Section 3.2 of [132].

[9] https://github.com/protegeproject/swrlapi/wiki/SWRLLanguageFAQ#Does_SWRL_support_Negation_As_Failure

[10] https://www.w3.org/community/n3-dev/

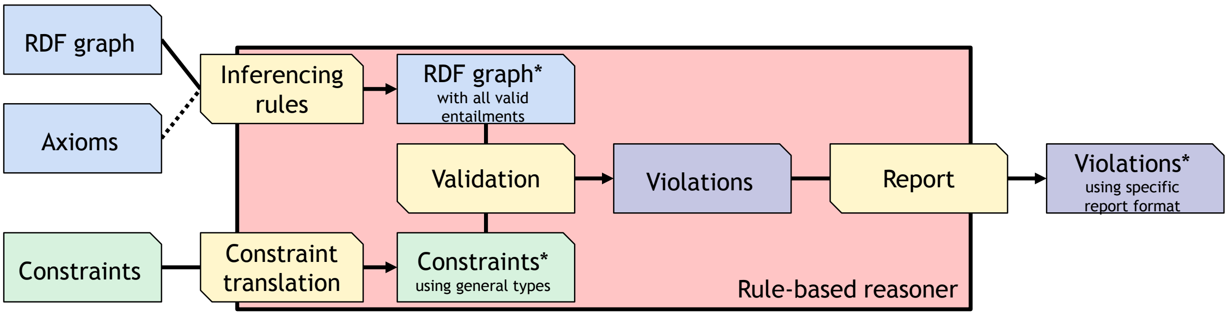

Figure 2. Components view of our approach. All double-snipped rectangles are rule sets, the single-snipped rectangles are RDF graphs or constraint declarations. The large overlapping rectangle is the rule-based reasoner. By taking all rule sets into account, the rule-based validator is formed. Four parts can be identified within the validation execution: (i) possibly guided by provided *Axioms*, all supported entailments of the given *RDF graph* can be generated using the *Inferencing rules*, resulting in *RDF graph**; (ii) the general *Constraints** are inferred from the given *Constraints* using a set of rules for *Constraint translation*; (iii) the rules for *Validation* generate *Violations*; and (iv) the returned *Violations** are structured given a set of rules that specify the *Report* format.

We follow the formalized semantics of N3Logic [132] as implemented in the EYE reasoner [9]: a clear formal definition of Notation3's semantics was missing from its initial proposal [20]. Verborgh et al. formalised the basics of the model theory of a logic with similar properties to N3Logic, excluding the constructs which lead to different interpretations (mainly nested implicit quantification) [132]. This work also proves the correctness of the calculus N3 reasoners use. Thus: the results of the reasoners are correct if the defined model theory is followed. Arndt et al. expanded on this work, specifically investigating the excluded constructs, and defined two different mappings from N3 syntax to core logic syntax covering two possible interpretations of N3Logic [9]. Even though this work defines two possible semantics, the difference between these two semantics does not influence the use of N3 in our paper since the semantic differences are only relevant for deeply nested formulas, our formulas are not of that nature (see Listing 3).

More, N3Logic supports at least OWL-RL inferencing [6, 8], which can be included during validation: the rules for OWL-RL are specified[11] and are supported by every rule language that is at least as expressive as Datalog. This includes N3Logic: the concrete realisation of these rules in N3 can be found online[12].

N3Logic, among others, covers existential rules, thus typically rendering the logic undecidable. This brings three trade-offs. First, we note that decidability does not imply that reasoning times are acceptable: even decidable logics can result in reasoner time-outs. Second, we expect the validation rules to be used in a distributed context (the Web). Thus, even though relevant research investigates the maximal subset of existential rules which are still decidable [13, 31, 125], we have no control over all potential rules used together with our validation rules, and cannot use these well-studied mechanisms ensuring a set of existential rules to be decidable. These mechanisms need to consider all rules together. Third, for example, the logic framework Prolog is a widely used Turing complete programming language. Even though this is a desirable property for a programming language, making it very expressive, checking properties over a Turing complete language is undecidable. Prolog remains a popular choice: we can conclude that using this undecidable logic allows for expressiveness, without necessarily introducing a performance bottleneck.

The *rule language* introduced together with N3Logic is N3 [18, 20]. Everything covered by RDF 1.1 Semantics [65] is covered in N3. Syntactically, it is a superset of Turtle [15]. N3 allows declaring inferencing rules, axioms, and constraints in the same language. As in RDF, blank nodes are understood as existentially quantified variables and the co-occurrence of two triples as in the RDF graph of formulas (9) and (10) is understood as their conjunction. More, N3 supports universally quantified variables, indicated by a leading question mark `?`.

```
?x :likes :IceCream.                    (13)
```

---

[11] https://www.w3.org/TR/owl2-profiles/#Reasoning_in_OWL_2_RL_and_RDF_Graphs_using_Rules

[12] http://eulersharp.sourceforge.net/2003/03swap/eye-owl2.html

stands for *"Everyone likes ice cream."*, or in first order logic

$$\forall x : \text{likes}(x, \text{ice-cream}) \tag{14}$$

Rules are written using curly brackets `{ }` and the implication symbol =>. An `rdfs:subClassOf` relation such as `:Person rdfs:subClassOf :Researcher` can be expressed as:

```
{?x a :Researcher} => {?x a :Person}.
```
(15)

The general `rdfs:subClassOf` relation can be expressed as:

```
{?C rdfs:subClassOf ?D. ?X a ?C}
          => {?X a ?D}.
```
(16)

Reasoner Reasoners that support N3Logic include FuXi, cwm, and EYE. FuXi[13] is a forward chaining production system for N3 whose reasoning is based on the Rete algorithm [55]. The forward chaining cwm [17] reasoner is a general-purpose data processing tool which can be used for querying, checking, transforming and altering information. EYE[14] [131] is a high-performance reasoner written in Prolog, enhanced with Euler path detection, allowing the creator of the rules to decide when to do forward reasoning and when backwards. EYE has generous support for built-in functions[15], among which, the RIF functions.

We choose the EYE reasoner as it fulfills the requirements as presented in Section 4. Furthermore, its ability to combine forward and backward chaining proves especially useful since constraint types are mostly localized to single relationships [26]. This means backward chaining has a potentially large impact on the performance: reasoning during validation can be very targeted, and in most cases, only facts that are relevant to the defined constraints are inferred.

5.3. Implementation

Our implementation is dubbed "Validatrr": a validator using rule-based reasoning. A Node.js JavaScript framework was created to discover and retrieve the vocabularies and ontologies as required by the use case, manage the commandline arguments, etc. The implementation is available at https://github.com/IDLabResearch/validatrr, and the set of validation rules (Fig. 2, center) is available at https://github.com/IDLabResearch/data-validation.

5.4. Execution example

As example, we validate an RDF graph with a custom set of inferencing steps using SHACL constraints. We take into account the example of the introduction (formula (1)), but the case where `:Bob` has two birthdates defined. The implications of `rdfs:domain` (formula (2)) should be taken into account as defined in RDFS [29] during validation, and the SHACL constraint states that each person should have exactly one birthdate (Listing 1). The result should be in the SHACL validation report format. Using this example, we can detail every step as show in Fig. 2: the RDF graph with all supported entailments (*RDF graph**) and general constraint types (*Constraints**) are inferred using a (custom) set of inferencing rules (*Inferencing rules*) and constraint translation rules (*Constraint translation*), after which the validation occurs (*Validation*), and the resulting violations are translated via rules (*Report*) in a specific report format (*Violations**).

```
:PersonShape a sh:NodeShape ;
  sh:targetClass :Person ;
  sh:property [
    sh:path :birthdate ;
    sh:minCount 1 ; sh:maxCount 1 ;
    sh:datatype xsd:date ] .
```

Listing 1: Person Shape in SHACL

To make sure `rdfs:domain` is correctly interpreted during validation, we include additional inferencing rules[16] (*Inferencing rules*), described in N3 as

```
{?P rdfs:domain ?C. ?X ?P ?Y}
          => {?X a ?C} .
```
(17)

Given formula (17), it is inferred that `:Bob` is a person (*RDF graph**).

To make sure SHACL constraints are correctly interpreted, SHACL translation rules need to be included during validation (*Constraint translation*). The gen-

---

[13] http://code.google.com/p/fuxi/

[14] https://github.com/josd/eye

[15] http://eulersharp.sourceforge.net/2003/03swap/eye-builtins.html

[16] http://eulersharp.sourceforge.net/2003/03swap/rdfs-domain.html

eral “Exact Qualified Cardinality Restrictions” RDF-CV constraint is inferred from the SHACL constraint of Listing 1, using the rules of Listing 2 (*Constraints**).

```
{
 ?sh a sh:NodeShape ;
  sh:targetClass ?Class ;
  sh:property [
   sh:path ?p ;
   sh:minCount ?v ; sh:maxCount ?v1 ;
  sh:datatype ?C ] .
 ?v pred:numeric-equal ?v1
} => {
 ?constraint a rdfcv:SimpleConstraint ;
  :originalShape ?sh ;
  :constraintType :ExQualCardRestr ;
  rdfcv:constrainingElement
   :exact-cardinality ;
  rdfcv:contextClass ?Class ;
  rdfcv:leftProperties ?p ;
  rdfcv:classes ?C ;
  rdfcv:constrainingValue ?v
} .
```

Listing 2: Translate the SHACL shape to a general constraint type

Validation makes use of general rules, i.e., Listing 3 (*Validation*). Lines 11–14 define how to find a violation, relying on built-ins: gather a set of resources in a list (`e:findall`), calculate the length of that list (`e:length`), and mathematically compare numbers (`math:notEqualTo`). For all objects of a certain class or datatype related using predicate `?p` (in this case `:birthdate`) where the number of objects is different from the constraint value `?v` (in this case `1`), a violation is returned (lines 16–21).

```
 1 {
 2  ?constraint a rdfcv:SimpleConstraint ;
 3   :constraintType :ExQualCardRestr ;
 4   rdfcv:constrainingElement
 5    :exact-cardinality ;
 6   rdfcv:contextClass ?Class ;
 7   rdfcv:leftProperties ?p ;
 8   rdfcv:classes ?C ;
 9   rdfcv:constrainingValue ?v .
10  ?x a ?Class.
11  _:x e:findall
12   ( ?C {?x ?p ?o. ?o a ?C} ?list) .
13  ?list e:length ?l .
14  ?l math:notEqualTo ?v
15 } => {
16  _:v a :constraintViolation ;
17   :violatedConstraint ?constraint ;
18   :class ?Class ;
19   :instance ?x ;
20   :objectClass ?C ;
21   :property ?p
22 } .
```

Listing 3: Validate using general constraint types

The general violations are translated into a report format (Fig. 2, *Violations**), e.g., using the SHACL Validation Report [78] (see Listing 4). The result is a set of triples using the exact same input and output as a SHACL processor. However, the RDF graph’s supported entailments can be matched to the use case, and the process is a single reasoning execution with transparent rule sets.

```
{
 _:v a :constraintViolation ;
  :violatedConstraint [
   :originalShape ?sh ;
   :constraintType :exact-cardinality
  ] ;
  :class ?Class ;
  :instance ?x ;
  :objectClass ?C ;
  :property ?p
} => {
 _:y a sh:ValidationReport ;
  sh:conforms false ;
  sh:result [
   a sh:ValidationResult ;
   sh:resultSeverity sh:Violation ;
   sh:focusNode ?x ;
   sh:resultPath ?p ;
   sh:resultMessage "No exact match" ;
   sh:sourceShape ?sh ]
} .
```

Listing 4: Translate the general violations to the SHACL validation report

Moreover, different constraint descriptions are easily supported via the general constraint types. Given the OWL restriction of Listing 5: using a different set of rules, we can translate this restriction into the same constraint type (Listing 6). The validation process continues exactly the same.

```
:Person rdfs:subClassOf _:x .
_:x a owl:Restriction ;
 owl:onProperty :birthdate ;
 owl:qualifiedCardinality
```

```
      "1"^^xsd:nonNegativeInteger ;
  owl:onDataRange xsd:date .
```

Listing 5: An OWL restriction

```
{
  ?Class rdfs:subClassOf ?c .
  ?c a owl:Restriction ;
    owl:onProperty ?x ;
    owl:qualifiedCardinality ?v ;
    owl:onDataRange ?C
} => {
  ?constraint a rdfcv:SimpleConstraint ;
     :originalShape _:x ;
     :constraintType :ExQualCardRestr ;
    rdfcv:constrainingElement
      :exact-cardinality ;
    rdfcv:contextClass ?Class ;
    rdfcv:leftProperties ?p ;
    rdfcv:classes ?C ;
    rdfcv:constrainingValue ?v
} .
```

Listing 6: Translate the OWL restriction to the general constraint type

6. Hypothesis validation

To validate the hypotheses of Section 1.2, we compare Validatrr to different validation approaches. We show that Validatrr (i) accurately explains the root cause of why a violation occurs in more cases than specified in SHACL, given the SHACL core constraint components (accepting Hypothesis 1, see Section 6.1); (ii) returns an accurate number of validation results with respect to the used set of inferencing steps, compared to an integrity constraints validator with a fixed set of inferencing steps using RDFUnit (accepting Hypothesis 2, see Section 6.2); and (iii) supports an equivalent number of constraint types than existing approaches (accepting Hypothesis 3, see Section 6.3). The performance evaluation shows that our implementation is faster than the state of the art when combining inferencing and validation for commonly published datasets (accepting Hypothesis 4, see Section 6.4).

6.1. Root cause explanation of constraint violations

Using the logical proof, we increase the explanation's accuracy compared to what is currently expected of a validation approach. SHACL is a W3C Recommendation standardizing the description of constraints and violation reports for RDF graph validation. We show that the logical proof produced by the rule-based reasoning execution provides more detailed root cause explanations of constraint violations, compared to SHACL's violation report description.

The SHACL recommendation provides a set of test cases, enabling implementations to prove compliance[17]. The validation report denotes the violating resources via `sh:focusNode`, and in some cases can further specify the violating path via `sh:resultPath` and the violating value via `sh:value` [78]. However, it is not always possible to retrieve such additional information about the root cause. We revisit the previous example constraint that given a resource r, this resource has $(r_{firstname} \wedge r_{lastname}) \vee (r_{nickname})$[18]. Validation of formula (1) using a conforming SHACL implementation results in a validation report similar to Listing 7. The validation report does not provide any further details to explain why `:Bob` is invalid[19].

```
[ rdf:type sh:ValidationReport ;
  sh:conforms "false"^^xsd:boolean ;
  sh:result [
    rdf:type sh:ValidationResult ;
    sh:focusNode :Bob ;
    sh:resultSeverity sh:Violation ;
    sh:sourceConstraintComponent
      sh:OrConstraintComponent ;
    sh:sourceShape :PersonNameShape ;
    sh:value :Bob ; ] ; ]
```

Listing 7: Validation report of an OR constraint

The rule-based reasoning execution of Validatrr can generate a proof, showing the rules used to reach a conclusion. This logical proof allows to determine, for each violation, which part of the RDF graph is the root cause of the violation, and which axiom of the used ontology triggered an inference causing the violation. Listing 8 shows the part of the proof which contains the rules deriving the violation. For `:firstname`, `:lastname`, and `:nickname`, we query objects that are linked using the respective predicate (Listing 8, lines 12–15, 18–21, and

---

[17] https://github.com/w3c/data-shapes/tree/gh-pages/data-shapes-test-suite/tests

[18] This example is similar to the following SHACL test case: https://github.com/w3c/data-shapes/blob/gh-pages/data-shapes-test-suite/tests/core/node/or-001.ttl

[19] https://www.w3.org/TR/shacl/#validator-OrConstraintComponent

24–27). $\mathcal{K}$ is the scope of our knowledge base, in which we look for violations. We count the number of objects found and compare them with the needed number. For `:firstname`, one linked object is found (Listing 8, lines 16–17), however, no linked object is found for `:lastname` nor `:nickname` (Listing 8, lines 22–23 and 28–29): a violation is returned.

```
 1  <#lemma20> a r:Inference;
 2  r:gives {
 3  _:b1 a :constraintViolation.
 4  _:b1 :violatedConstraint _:b2.
 5  _:b1 :class :Man.
 6  _:b1 :instance :Bob.
 7  _:b1 :property :lastname.
 8  _:b1 :property :nickname. };
 9  r:evidence (
10  ...
11  <#lemma37>
12  [ a r:Fact; r:gives { (𝒦 1) e:findall
13  (1
14  {:Bob :firstname _:b3}
15  (1))}]
16  [ a r:Fact; r:gives {(1) e:length 1}]
17  [ a r:Fact; r:gives {1 math:greaterThan 0}]
18  [ a r:Fact; r:gives {(𝒦 1) e:findall
19  (1
20  {:Bob :lastname _:b3}
21  ())}]
22  [ a r:Fact; r:gives {() e:length 0}]
23  [ a r:Fact; r:gives {0 math:lessThan 1}]
24  [ a r:Fact; r:gives {(𝒦 1) e:findall
25  (1
26  {:Bob :nickname _:b3}
27  ())}]
28  [ a r:Fact; r:gives {() e:length 0}]
29  [ a r:Fact; r:gives {0 math:lessThan 1}]).
```

Listing 8: Validation proof of an OR constraint

Due to this proof, Validatrr can provide detailed explanations for the root causes of violations for all SHACL core constraint components, compared to 46%–75% of SHACL-conforming implementations. Analysis of the SHACL specification shows that, out of the 28 core constraint components, 13 (46%) provide a full explanation of the root cause (summarized in Table 2). For eight of the remaining components (an additional 29%), the validation report returns which resource violates which constraint, but does not return a detailed explanation. For example, a `sh:class` violation occurs when the targeted node is a literal, *or* when the targeted node is not classified accordingly, but this disjunction is not reflected in the validation report. For the remaining seven components, the validation report does not provide an explanation at all. For example, violations of nested shapes are not reflected in the validation report, only violations of top-level shapes.

Table 2

Analysis of root cause explanation of violations for SHACL core constraint components. Validatrr can provide more detailed explanations for up to 56% of the components compared to SHACL-conforming implementations.

SHACL Name	Root Cause Explanation	Comment
`sh:class`	~	disjunction
`sh:datatype`	~	disjunction
`sh:nodeKind`	~	disjunction
`sh:minCount`	✗	no explanation
`sh:maxCount`	✗	no explanation
`sh:minExclusive`	✓	
`sh:minInclusive`	✓	
`sh:maxExclusive`	✓	
`sh:maxInclusive`	✓	
`sh:minLength`	~	disjunction
`sh:maxLength`	~	disjunction
`sh:pattern`	✓	
`sh:languageIn`	~	disjunction
`sh:uniqueLang`	✗	no explanation
`sh:equals`	✓	
`sh:disjoint`	✓	
`sh:lessThan`	✓	
`sh:lessThanOrEquals`	✓	
`sh:not`	✓	
`sh:and`	~	conjunction
`sh:or`	~	disjunction
`sh:xone`	✓	
`sh:node`	✗	nesting
`sh:property`	✗	nesting
`sh:qualifiedValueShape`, `sh:qualifiedMinCount`, `sh:qualifiedMaxCount`	✗	nesting
`sh:close`, `sh:ignoredProperties`	✓	
`sh:hasValue`	✓	
`sh:in`	✗	nesting

Compared to SHACL-conforming implementations, Validatrr supports, among others, explanation of disjunction and nested shapes. Our approach provides detailed explanations for all core components of W3C's recommended high-level language to describe constraints. We thus accept Hypothesis 1.

6.2. Accurate number of found violations

Validatrr finds a more accurate number of violations compared to the state of the art. To prove this, we first compare Validatrr with the state of the art functionally, and then include a set of inferencing steps to clarify the difference.

Specifically, we compare with RDFUnit [82]. Hartmann et. al explicitly proposed using query-based approaches for validation [24], and RDFUnit is such a query-based approach, relying on a SPARQL endpoint, and describing the constrains using SPARQL templates named Data Quality Test Patterns (DQTP). As such, RDFUnit is highly configurable and one of the implementations that supports SHACL[20].

Functional comparison We compare with the original pattern library of RDFUnit [82]. This pattern library is the closest to the constraint types as introduced by Hartmann et al. [27, 64]: the mapping between those two is presented in previous work [7]. We test all unit tests defined by RDFUnit[21] after retrieving them as-is from the RDFUnit repository. As Validatrr validates general constraint types, a custom profile was created that translates the RDFUnit patterns to general constraint types. For a detailed explanation of the different test cases, we refer to the original RDFUnit paper [82].

The validation results depend on the used set of inferencing steps. RDFUnit implicitly takes "every resource is an `rdfs:Resource`" and the `rdfs:subClassOf` construct into account, forming the custom set of inferencing steps υ. We compare RDFUnit with Validatrr using three sets of inferencing steps, taking into account (i) no entailment at all ($\emptyset$), (ii) the custom set of inferencing steps (υ), and (iii) full RDFS entailment (ρ).

Table 3 summarizes the results. For each constraint, we mention the test case's name, the number of violations that RDFUnit detects, and the number of violations that Validatrr detects using the different sets of inferencing steps. The table shows the impact of using different sets of inferencing steps: depending on the set, Validatrr finds a different number of violations. More, Validatrr detects more violations using the same set of inferencing steps: there is a higher number of found violations for Validatrr under υ compared to RDFUnit.

Table 3

Comparing RDFUnit to Validatrr using different sets of inferencing steps ($\emptyset$, υ, and ρ). Validatrr finds more violations given the same set of inferencing steps, and the set of inferencing steps used impacts the result. Test cases where Validatrr outperforms RDFUnit are starred. Rows where Validatrr and RDFUnit differ are marked gray.

Test Case	**# found violations**			
	RDFUnit	Validatrr		
	υ	$\emptyset$	υ	ρ
INVFUNC_correct	0	0	0	0
INVFUNC_wrong	2	0	2	2
OWLCARDT_correct	0	0	0	0
OWLCARDT_wrong_exact	6	6	6	6
OWLCARDT_wrong_max	2	2	2	2
OWLCARDT_wrong_min	2	2	2	2
OWLDISJC_correct	0	0	0	2
OWLDISJC_wrong	6	2	6	6
OWLQCARDT_correct	0	0	0	0
OWLQCARDT_wrong_exact	6	6	6	6
OWLQCARDT_wrong_max	2	2	2	2
OWLQCARDT_wrong_min	2	2	2	2
RDFLANGSTRING_correct	0	0	0	0
RDFLANGSTRING_wrong	2	2	2	0
*RDFSRANGE-MISS_wrong**	1	3	3	0
RDFSRANGED_correct	0	0	0	0
*RDFSRANGED_wrong**	2	3	3	0
*RDFSRANGE_correct**	0	5	4	0
*RDFSRANGE_wrong**	1	3	3	3
RDFSRANG_LIT_correct	0	0	0	0
RDFSRANG_LIT_wrong	3	3	3	1

Validatrr finds more violations and supports more constraint types than RDFUnit, denoted as starred test cases RDFSRANGE-MISS_wrong, RDFSRANGED_wrong, RDFSRANGE_correct, and RDFSRANGE_wrong. RDFUnit does not yet support the constraint type *multiple ranges*: when a certain predicate is used, each resource linked as an object to that predicate should be classified into multiple classes. In all other cases, both solutions identify the same number of violations when using the same set of inferencing steps. Validatrr thus functionally outperforms the pattern library (i.e., the corresponding constraint types) of RDFUnit.

Impact of including sets of inferencing steps during validation Running Validatrr using different sets of inferencing steps impacts the number of found violations. Validatrr is designed to easily configure this set using inferencing rules (Fig. 2, top-left). The results are found in Table 3, comparing the different Validatrr columns. On the one hand, certain violations are not found without entailment ($\emptyset$), as is the case for INVFUNC_wrong and OWLDISJC_wrong. On the other hand, violations are

---

[20] https://w3c.github.io/data-shapes/data-shapes-test-suite/

[21] https://github.com/AKSW/RDFUnit/tree/master/rdfunit-core/src/test/resources/org/aksw/rdfunit/validate/data

resolved early-on when including RDFS entailment (ρ), as is the case for RDFLANGSTRING_wrong.

Compared to existing validation approaches, our approach allows including custom sets of inferencing steps during validation. The inferencing provenance is retained in the proof, as all inferencing occurs during a single reasoning execution. The logical proof can thus distinguish between violations that are caused due to constraint violations in the original RDF graph, or due to entailment during validation. We thus accept Hypothesis 2.

6.3. Equivalent number of constraint types

Validatrr can support an equivalent number of constraint types compared to existing validation approaches such as RDFUnit and SHACL. In the previous section, we showed we functionally outperform the original pattern library of RDFUnit whilst including a custom set of inferencing steps during validation. In this section, we compare our number of supported constraint types to that of SHACL [78].

We test Validatrr against general constraint types [63, 64], to show that the number of supported constraint types is equivalent to SHACL. We do not test specifically against SHACL's test cases, as Validatrr is independent of the constraint language. We provide a set of test cases, used to test these different constraint types[22].

Hartmann et al. investigated the constraint type support of SHACL, and stated that its coverage is 52% [63]. We updated the coverage report as presented by Hartmann et al. to take the latest SHACL specification and advanced features into account [78, 81]. The relevant data is available at Appendix A, and online[23]. This updated report shows that SHACL's constraint type coverage is 84%.

Validatrr can cover up to 94% of all constraint types – given the current expressive support for built-ins – and has been tested to cover a similar number of constraint types as SHACL[24]. After including the rules for the remaining constraint types, we support an equivalent number of constraint types compared to SHACL. We thus accept Hypothesis 3.

Achieving 100% coverage (i.e., the remaining five constraint types) requires additional development on the reasoner to support specific built-ins. "Whitespace Handling" and "HTML Handling" require parsing built-ins, and "Valid Identifiers" requires a built-in to test URIs' dereferencability. The remaining two types ("Structure" and "Data Model Consistency") are general constraint types, defined by Hartmann et al., requiring SPARQL support. Supporting these constraint types requires a translation from SPARQL queries to N3 rules, for which we refer to related work [122].

6.4. Speed

A validation approach that supports a custom set of inferencing steps is faster than a validation system that includes a reasoning preprocessing step. We first compare the performance of Validatrr to that of RDFUnit, both without and with a custom set of inferencing steps.

For these performance evaluations, we used 300 data sets with sizes ranging from ten to one million triples, and an executing machine consisting of 24 cores (Intel Xeon CPU E5-2620 v3 @ 2.40GHz) and 128GB RAM. All evaluations were performed using untampered docker images for both approaches to maintain reproducibility, the different tests were orchestrated using custom scripts. All timings include the docker images' initialization time. The data is available online[25].

Performance comparison We compare the execution time of Validatrr to RDFUnit, following RDFUnit's original evaluation method. We use a default set of constraints for a fixed set of schemas, as defined by Kontokostas et al. [82]. We consider six commonly used schemas: FOAF, GeoSPARQL, OWL, DC terms, SKOS, and Prov-O. For each schema, we use RDF graphs of varying size. The validated RDF graphs' size range from ten triples to one million triples, in logarithmic steps of base ten. At most ten different RDF graphs – per schema, per RDF graph size – were downloaded, by querying LODLaundromat's SPARQL endpoint [16].

We validate the different RDF graphs against their respective schema using the default set of constraints and set of inferencing steps (υ) of RDFUnit, and measure total execution time of Validatrr and RDFUnit. The median execution time across all schemas is plotted against RDF graph size per approach in a log-log scale (see Fig. 3). To make sure we can combine execution times across schemas, we tested the null hypothesis that no significant difference in execution time was found

---

[22] https://github.com/IDLabResearch/data-validation

[23] https://github.com/IDLabResearch/constraint-types-coverage

[24] The test report is available at https://github.com/IDLabResearch/validatrr/blob/v0.2.0/reports/validatrr-rdfcv-earl.ttl

[25] https://github.com/IDLabResearch/validation-benchmark/tree/master/data/validation-journal

between schemas, by performing an ANOVA statistical test with single factor "used schema" for measurement variable "execution time per triple", executed pairwise for all used schemas. The null hypothesis with $\alpha = 0.05$ was accepted for every pair. The number of found violations are not plotted, as statistical analysis shows no large correlation between execution time and number of found violations, neither for Validatrr or RDFUnit (-0.0203 and 0.0458, respectively).

Validatrr's execution time is highly correlated with the number of triples of the validated RDF graph. Regression analysis shows an R square value of 0.9998, the null hypothesis with $\alpha = 0.05$ is accepted: Validatrr's execution time grows linearly with respect to the size of the validated RDF graph. Meanwhile, the execution time of RDFUnit remains constant at around 30s. This could largely be due to the set-up time required by RDFUnit, however, the timings attained via RDFUnit's docker image does not allow us to draw further conclusions. The set-up time of RDFUnit thus possibly dominates the total execution time.

Without customizing the set of inferencing steps and docker images, Validatrr is faster for small RDF graphs. Validatrr is about an order of magnitude faster until 10,000 triples, namely, 1-2s per RDF graph compared to 30s per RDF graph for RDFUnit. After 100,000 triples, Validatrr is slower than RDFUnit, as Validatrr's linearly growing execution time surpasses RDFUnit's execution time.

Custom inferencing steps' performance impact We compare the execution time of Validatrr to RDFUnit when using a custom set of inferencing steps. We use RDFS entailment (ρ): it is commonly used, and the evaluation of Section 6.2 showed it affects the number of violations found. For Validatrr, we include the RDFS rules during validation. For RDFUnit, we include an RDFS entailment preprocessing step, as RDFUnit's docker image does not allow configuration to use a SPARQL engine that has inferencing capabilities. However, even if it would be possible to use a different SPARQL engine, a reasoning preprocessing step would still be needed for use cases that require support for a specific set of inferencing steps, not covered by typical entailment regimes [2].

To keep the measures comparable, we use the EYE reasoner as used in Validatrr with the same RDFS entailment rule set to execute the reasoning preprocessing step. This also precludes the need to compare with other sets of inferencing steps than RDFS entailment: the conclusions will be similar due to the usage of the same reasoner. Fig. 4 depicts the timings of RDFUnit and Validatrr. For RDFUnit, it depicts the combined timings of RDFS entailment as preprocessing step and validation on the newly inferred RDF graph(RDFUnit (ρ)), and it depicts solely the validation timings on the newly inferred graph (RDFUnit). For Validatrr, it depicts the timings of the validation with the two sets of inferencing rules (Validatrr (ρ) and Validatrr (υ), respectively).

Validatrr's performance is not affected by using a different set of inferencing steps, whereas the preprocessing step deteriorates RDFUnit's performance. This effect is noticable starting from RDF graphs of 10,000 triples. For RDF graphs of one million triples, compared to the previous evaluation, median execution time rises from 27s to 210s for RDFUnit, largely due to the reasoning preprocessing step.

The number of found violations inversely affects the validation execution speed. Most original violations handle missing domain and range classes, which is inferred in RDFS entailment. Statistical analysis does not allow us to accept the null hypothesis that the number of violations found is inversely correlated to the execution time. However, we notice increased performance for both approaches when less violations need to be handled. Compared to previous evaluation, for one million triples, execution time (without reasoning preprocessing) drops from 27s to 21s for RDFUnit, and from 116s to 80s for Validatrr.

The performance evaluations show that the execution time of Validatrr outperforms RDFUnit for small RDF graphs up to 100,000 triples, and its linear scaling behavior is not affected by including RDFS entailment during validation. Validatrr outperforms RDFUnit when reasoning preprocessing is needed, i.e., when the used SPARQL endpoint does not support inferencing up to the needed expressiveness, or cannot be sufficiently customized to the use case. Where RDFUnit first needs to infer all implicit data before validation, Validatrr can infer this data during validation, and thus performs better. We thus accept Hypothesis 4.

7. Conclusion and future work

In this section, we discuss our proposed rule-based reasoning validation approach and introduced implementation. We provide concluding remarks and guide towards future work with respect to (i) the detailed root cause explanations, (ii) the fine-grained level of configuration, (iii) the number of constraint types sup-

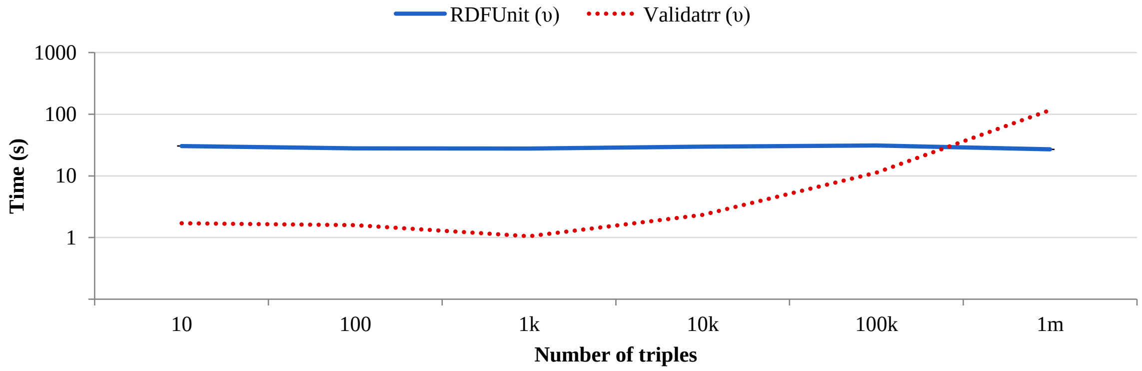

Figure 3. Validatrr's execution speed (dotted line) is up to an order of magnitude faster than RDFUnits's (solid line) when the number of triples per RDF graph is below 100,000 triples

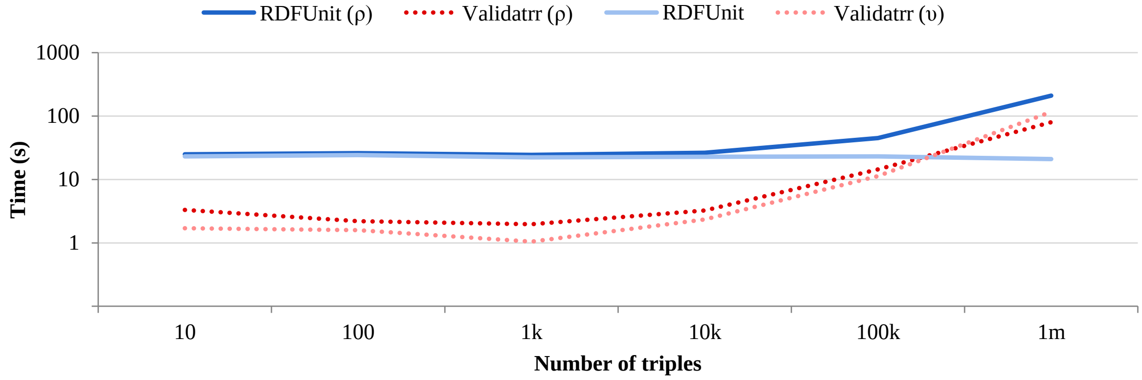

Figure 4. Validatrr's performance is not affected when including the RDFS inferencing rules (dotted line, compared to the lighter dotted line), whereas the reasoning preprocessing time deteriorated RDFUnit's performance (solid line, compared to the lighter solid line).

ported by our approach, and (iv) the scaling behavior of Validatrr's performance. We close by providing some further research perspectives.

The logical proof of a validation execution, generated by the rule-based reasoner, provides a more detailed root cause explanation of why a violation occurs than the state of the art. Our evaluation does not imply that existing approaches and implementations are not capable of providing a similar level of detail. However, it does show the feasibility of more detailed explanations, and the capability of our approach to generate them. To improve the level of detail of explanations provided in the validation report, our work can guide future iterations of, e.g., SHACL's validation report descriptions, and the algorithms that generate them.

Our approach is fully configurable by adjusting different rule sets: only a single declaration and single implementation is needed to support different constraint languages, sets of inferencing steps, and validation report descriptions. This level of control considerably increases expressiveness and complexity of the validator, and a small change in a rule set could have large effects on the validation results. However, such fine-grained configuration is not needed for every use case. Future work requires investigation into configuration defaults for, among others, ShEx and SHACL: to what extend can Validatrr be configured to function as a compliant ShEx or SHACL validator, and how will the combination of inferencing rule sets look like? A short-term goal is showing that Validatrr with the right configuration passes the core SHACL tests and is included as a compliant SHACL validator in the respective W3C documentation[26]. As such, we can provide a compliant SHACL validator where `sh:entailment` is accurately supported: the user can choose exactly which inferencing rule set is supported during validation, and can choose not to rely on the predefined custom set of inferencing steps (i.e., support for `rdfs:subClassOf`,

---

[26] https://w3c.github.io/data-shapes/data-shapes-test-suite/

but no other RDFS entailments) as currently specified in SHACL [78].

Our approach supports an equivalent number of constraint types compared to the state of the art, with description logic-expressiveness up to at least OWL-RL. An important point of interest is handling recursion, one of the main differences between ShEx and SHACL. The semantics of ShEx are defined, also for recursion [23], and – as it is currently undefined in the SHACL specification [78] – current works are investigating recursion in combination with negation for SHACL [34]. Future work for our approach is investigating recursion, taking into account the conclusions and mentioned complexity issues of aforementioned works. Accepting that the general problem is NP-Hard, using rule-based reasoning gives us a strong tool to handle recursion. A rule-based reasoner such as the EYE reasoner has path detection: different validations calling each other can be handled, as path detection prevents the reasoner from applying the same rule to the same data twice. In this regard, we can further investigate whether the strategies of Answer Set Programming [48] help to solve related problems, taking into account their two kinds of negation (Negation as Failure and strong negation). After investigating which rules are needed to handle recursion, the user can choose whether or not recursion should be supported during validation, as these extra rules can be added or not.

The performance of Validatrr is up to an order of magnitude faster than RDFUnit for RDF graphs up to 100,000 triples, and scales linearly w.r.t. the number of triples in the RDF graph. However, it scales less than RDFUnit, making Validatrr less suitable for large RDF graphs. As such, a trade-off must be made: our approach, which performs better for smaller RDF graphs, allows fine-grained configuration and detailed explanation, whereas other approaches scale better but do not provide the same level of detail. For future work, further investigation into related works that aim to improve the performance of rule-based reasoners, such as the work of Arndt et al. [6], can be used to improve the current scaling behavior of Validatrr.

Further research perspectives include validation of RDF graph generation descriptions, and automatic graph refinement based on violation explanations. The combination reduces the effort required to provide high-quality RDF graph generation descriptions, and is being further investigated by Heyvaert et al. [67].

On the one hand, a declarative description for generating an RDF graph – e.g., using the RDF Mapping Language (RML) [46] – can be validated, to show whether that description produces a valid RDF graph [47]. Certain constraints that apply to the description can be inferred based on the constraints that apply to the RDF graph. By including a custom inferencing rule set that reflects such inferencing in Validatrr, the generation description can be validated based on the set of constraints that apply to the RDF graph. As such, only a single set of constraints needs to be maintained and understood. The requirements of this custom inferencing rule set, and which constraint types can be applied to generation descriptions, is future work.

On the other hand, rules that handle the accurate explanations of why a violation is returned, can provide suggestions to (automatically) resolve the violation. For example, the constraint specifying “every book should have either an ISSN or an ISBN number” is violated by a resource that has both numbers. Suggestions include removing the ISSN number and removing the ISBN number. Which types of suggestions can be provided, and in which order these should be applied, is future work.

Acknowledgements

The described research activities were funded by Ghent University, imec, Flanders Innovation & Entrepreneurship (VLAIO), and the European Union. Ruben Verborgh is a postdoctoral fellow of the Research Foundation – Flanders (FWO).

Appendix A. Updated Constraint Types Coverage

Table 4

Coverage of validation approaches w.r.t. constraint types. Taken from Hartmann et al. [63], updated take the recent advancements of SHACL into account [78, 81]. Changes w.r.t. the original tabe of Hartmann et al. are marked grey.

Type	Name	SHACL	Validatrr
A01	*Functional Properties	✓	✓
A02	*Inverse-Functional Properties	✗	✓
A03	*Primary Key Properties	✗	✓
A04	*Subsumption	✓	✓
A05	*Sub-Properties	✗	✓
A06	*Object Property Paths	✗	✓
A07	Allowed Values	✓	✓
A08	Not Allowed Values	~	✓
A09	*Class Equivalence	✗	✓
A10	*Equivalent Properties	✓	✓
A11	Literal Value Comparison	✓	✓
A12	Value is Valid for Datatype	✓	✓
A13	*Property Domains	✓	✓
A14	*Property Ranges	✓	✓
A15	*Class-Specific Property Range	✓	✓
A16	Data Property Facets	✓	✓
A17	Literal Ranges	✓	✓
A18	Negative Literal Ranges	~	✓
A19	IRI Pattern Matching	✓	✓
A20	Literal Pattern Matching	✓	✓
A21	Negative Literal Pattern Matching	~	✓
A22	*Existential Quantifications	✓	✓
A23	*Universal Quantifications	✓	✓
A24	*Value Restrictions	✓	✓
A25	Use Sub-Super Relations in Validation	✗	✓
A26	Negative Property Constraints	~	✓
A27	Language Tag Matching	✓	✓
A28	Language Tag Card.	~	✓
A29	Whitespace Handling	✗	✗
A30	HTML Handling	✗	✗
A31	Structure	✓	✗
A32	*Minimum Unqualified Card.	✓	✓
A33	*Minimum Qualified Card.	✓	✓
A34	*Maximum Unqualified Card.	✓	✓
A35	*Maximum Qualified Card.	✓	✓
A36	*Exact Unqualified Card.	✓	✓
A37	*Exact Qualified Card.	✓	✓
A38	*Cardinality Shortcuts	~	~
A39	Vocabulary	✓	✓
A40	Provenance	✓	~
A41	Required Properties	~	✓
A42	Optional Properties	~	✓

Table 5

Coverage of validation approaches w.r.t. constraint types. Taken from Hartmann et al. [63], updated take the recent advancements of SHACL into account [78, 81]. Changes w.r.t. the original tabe of Hartmann et al. are marked grey (2).

Type	Name	SHACL	Validatrr
A43	Repeatable Properties	~	✓
A44	Conditional Properties	✓	✓
A45	Recommended Properties	✓	✓
A46	Severity Levels	✓	~
A47	Labeling and Documentation	✓	~
A48	Context-Sp. Property Groups	✓	✓
A49	Context-Sp. Exclusive OR of P.	✓	~
A50	Context-Sp. Exclusive OR of P. Groups	✓	~
A51	Context-Sp. Inclusive OR of P.	✓	✓
A52	Context-Sp. Inclusive OR of P. Groups	✓	✓
A53	Mathematical Operations	~	~
A54	Ordering	✓	~
A55	*Inverse Object Properties	✓	~
A56	*Symmetric Object Properties	✓	~
A57	*Asymmetric Object Properties	✓	✓
A58	*Transitive Object Properties	✗	~
A59	*Self Restrictions	✓	~
A60	Valid Identifiers	✗	✗
A61	Recursive Queries	✓	~
A62	*Reflexive Object Properties	✓	~
A63	*Class-Sp. Reflexive Object P.	✓	~
A64	*Irreflexive Object Properties	✓	✓
A65	*Class-Specific Irreflexive Object Properties	✓	~
A66	Data Model Consistency	✓	✗
A67	Handle RDF Collections	✓	~
A68	Membership in Controlled Vocabularies	✓	~
A69	Disjoint Properties	✗	✓
A70	Disjoint Classes	~	✓
A71	String Operations	~	~
A72	Aggregations	✓	~
A73	*Individual Equality	✓	~
A74	Individual Inequality	✓	~
A75	Context-Specific Valid Classes	✗	~
A76	Context-Specific Valid Properties	✓	~ ~
A77	Property Assertions	~	~
A78	*Intersection	✓	✓
A79	*Disjunction	✓	~
A80	*Negation	✓	✓
A81	*Default Values	✓	✓

References

[1] Ziawasch Abedjan and Felix Naumann. Improving RDF Data Through Association Rule Mining. *Datenbank-Spektrum*, 13 (2):111–120, May 2013. DOI: 10.1007/s13222-013-0126-x.

[2] Carlos Buil Aranda, Olivier Corby, Souripriya Das, Lee Feigenbaum, Paula Gearon, Birte Glimm, Steve Harris, Sandro Hawke, Ivan Herman, Nicholas Humfrey, Nico Michaelis, Chimezie Ogbuji, Matthew Perry, Alexandre Passant, Axel Polleres, Eric Prud'hommeaux, Andy Seaborne, and Gregory Todd Williams. SPARQL 1.1 Overview. Recommendation, World Wide Web Consortium (W3C), March 2013. URL http://www.w3.org/TR/sparql11-overview/.

[3] Marcelo Arenas, Sebastián Conca, and Jorge Pérez. Counting beyond a Yottabyte, or how SPARQL 1.1 property paths will prevent adoption of the standard. In Alain Mille, Fabien Gandon, and Jacques Misselis, editors, *Proceedings of the 21st international conference on World Wide Web*, pages 629–638, New York, NY, United States, April 2012. Association for Computing Machinery. DOI: 10.1145/2187836.2187922.

[4] Marcelo Arenas, Oscar Corcho, Elena Simperl, Markus Strohmaier, Mathieu d'Aquin, Kavitha Srinivas, Paul Groth, Michel Dumontier, Jeff Heflin, Krishnaprasad Thirunarayan, and Steffen Staab, editors. *The Semantic Web – ISWC 2015: 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, Cham, October 2015. Springer. DOI: 10.1007/978-3-319-25010-6.

[5] Dörthe Arndt, Ruben Verborgh, Jos De Roo, Hong Sun, Erik Mannens, and Rik Van de Walle. Semantics of Notation3 Logic: A solution for implicit quantification. In Nick Bassiliades, Georg Gottlob, Fariba Sadri, Adrian Paschke, and Dumitru Roman, editors, *Rule Technologies: Foundations, Tools, and Applications*, volume 9202 of *Lecture Notes in Computer Science*, pages 127–143, Cham, July 2015. Springer. DOI: 10.1007/978-3-319-21542-6_9.

[6] Dörthe Arndt, Ben De Meester, Pieter Bonte, Jeroen Schaballie, Jabran Bhatti, Wim Dereuddre, Ruben Verborgh, Femke Ongenae, Filip De Turck, Rik Van de Walle, and Erik Mannens. Improving OWL RL reasoning in N3 by using specialized rules. In Valentina Tamma, Mauro Dragoni, Rafael Gonçalves, and Agnieszka Ławrynowicz, editors, *Ontology Engineering: 12th International Experiences and Directions Workshop on OWL*, volume 9557 of *Lecture Notes in Computer Science*, pages 93–104, Cham, April 2016. Springer. DOI: 10.1007/978-3-319-33245-1_10.

[7] Dörthe Arndt, Ben De Meester, Anastasia Dimou, Ruben Verborgh, and Erik Mannens. Using Rule-Based Reasoning for RDF Validation. In Stefania Constantini, Enrico Franconi, William Van Woensel, Roman Kontchakov, Fariba Sadri, and Dumitru Roman, editors, *Rules and Reasoning: International Joint Conference, RuleML+RR 2017, London, UK, July 12–15, 2017*, volume 10364 of *Lecture Notes in Computer Science*, pages 22–36, Cham, July 2017. Springer. DOI: 10.1007/978-3-319-61252-2_3.

[8] Dörthe Arndt, Pieter Bonte, Alexander Dejonghe, Ruben Verborgh, Filip De Turck, and Femke Ongenae. SENSdesc: Connect Sensor Queries and Context. In Reyer Zwiggelaar, Hugo Bamboa, Ana Fred, and Sergi Bermúdez i Badia, editors, *Proceedings of the 11th International Joint Conference on Biomedical Engineering Systems and Technologies*, volume 5, pages 671–679, Setúbal, Portugal, 2018. SCITEPRESS - Science and Technology Publications. DOI: 10.5220/0006733106710679.

[9] Dörthe Arndt, Tom Schrijvers, Jos De Roo, and Ruben Verborgh. Implicit quantification made explicit: How to interpret blank nodes and universal variables in Notation3 Logic. *Journal of Web Semantics*, 58(100501), October 2019. DOI: 10.1016/j.websem.2019.04.001.

[10] Ahmad Assaf, Raphaël Troncy, and Aline Senart. Roomba: An Extensible Framework to Validate and Build Dataset Profiles. In Fabien Gandon, Christophe Guéret, Serena Villata, John Breslin, Catherine Faron-Zucker, and Antoine Zimmerman, editors, *The Semantic Web: ESWC 2015 Satellite Events*, volume 9341 of *Lecture Notes in Computer Science*, pages 325–339, Cham, 2015. Springer. DOI: 10.1007/978-3-319-25639-9_46.

[11] Sören Auer, Jens Lehmann, and Axel-Cyrille Ngonga Ngomo. Introduction to Linked Data and Its Lifecycle on the Web. In Axel Polleres, Claudia d'Amato, Marcelo Arenas, Siegfried Handschuh, Paula Kroner, Sascha Ossowski, and Peter F. Patel-Schneider, editors, *Reasoning Web. Semantic Technologies for the Web of Data*, volume 6848 of *Lecture Notes in Computer Science*, pages 1–75, Berlin, Heidelberg, 2011. Springer. DOI: 10.1007/978-3-642-23032-5_1.

[12] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2nd edition, 2007.

[13] Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. On rules with existential variables: Walking the decidability line. *Artificial Intelligence*, 175(9-10):1620–1654, June 2011. DOI: 10.1016/j.artint.2011.03.002.

[14] Jorge Baier, Dietrich Daroch, Juan L. Reutter, and Domagoj Vrgoc. Evaluating navigational RDF queries over the Web. In Peter Dolog and Peter Vojtas, editors, *Proceedings of the 28th ACM Conference on Hypertext and Social Media*, New York, NY, United States, 2017. Association for Computing Machinery. DOI: 10.1145/3078714.3078731.

[15] David Beckett, Tim Berners-Lee, Eric Prud'hommeaux, and Gavin Carothers. RDF 1.1 Turtle – Terse RDF Triple Language. Recommendation, World Wide Web Consortium (W3C), February 2014. URL http://www.w3.org/TR/turtle/.

[16] Wouter Beek, Laurens Rietveld, Hamid R. Bazoobandi, Jan Wielemaker, and Stefan Schlobach. LOD Laundromat: A Uniform Way of Publishing Other People's Dirty Data. In Mika et al. [91], pages 213–228. DOI: 10.1007/978-3-319-11964-9_14.

[17] Tim Berners-Lee. Cwm, October 2000. URL http://www.w3.org/2000/10/swap/doc/cwm.html.

[18] Tim Berners-Lee. Notation 3 Logic, August 2005. URL http://www.w3.org/DesignIssues/N3Logic.

[19] Tim Berners-Lee and Dan Connolly. Notation3 (N3): A readable RDF syntax. Team submission, World Wide Web Consortium (W3C), March 2011. URL http://www.w3.org/TeamSubmission/n3/.

[20] Tim Berners-Lee, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler. N3Logic: A logical framework for the World Wide Web. *Theory and Practice of Logic Programming*, 8(3): 249–269, May 2008. DOI: 10.1017/S1471068407003213.

[21] Abraham Bernstein and Natasha Noy. Is this really science? the semantic webber's guide to evaluating research contributions. Technical report, University of Zurich – Department of Informatics (IFI), 2014. URL https://www.merlin.uzh.ch/contributionDocument/download/6915.

[22] Christian Bizer and Richard Cyganiak. Quality-driven information filtering using the WIQA policy framework. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(1):1–10, January 2009. DOI: 10.1016/j.websem.2008.02.005.

[23] Iovka Boneva, Jose Emilio Labra Gayo, and Eric Prud'hommeaux. Semantics and Validation of Shapes Schemas for RDF. In d'Amato et al. [38], pages 104–120. DOI: 10.1007/978-3-319-68288-4_7.

[24] Thomas Bosch and Kai Eckert. Requirements on RDF constraint formulation and validation. In William Moen and Amy Rushing, editors, *Proceedings of the International Conference on Dublin Core and Metadata Applications*, pages 95–108. Dublin Core Metadata Initiative, October 2014.

[25] Thomas Bosch and Kai Eckert. Towards Description Set Profiles for RDF using SPARQL as Intermediate Language. In William Moen and Amy Rushing, editors, *Proceedings of the International Conference on Dublin Core and Metadata Applications*, pages 129–137. Dublin Core Metadata Initiative, October 2014.

[26] Thomas Bosch, Erman Acar, Andreas Nolle, and Kai Eckert. The role of reasoning for RDF validation. In Sebastian Hellmann, Josiane Xavier Parreira, and Axel Polleres, editors, *Proceedings of the 11th International Conference on Semantic Systems*, pages 33–40, New York, NY, United States, 2015. Association for Computing Machinery. DOI: 10.1145/2814864.2814867.

[27] Thomas Bosch, Andreas Nolle, Erman Acar, and Kai Eckert. RDF Validation Requirements – Evaluation and Logical Underpinning. arXiv preprint, July 2015. URL http://arxiv.org/abs/1501.03933.

[28] Bojan Bozic, Rob Brennan, Kevin Feeney, and Gavin Mendel-Gleason. Describing Reasoning Results with RVO, the Reasoning Violations Ontology. In Jeremy Debattista, Jürgen Umbrich, and Javier D. Fernándex, editors, *Joint Proceedings of the 2nd Workshop on Managing the Evolution and Preservation of the Data Web (MEPDaW 2016) and the 3rd Workshop on Linked Data Quality (LDQ 2016) co-located with 13th European Semantic Web Conference (ESWC 2016)*, volume 1585 of *CEUR Workshop Proceedings*, pages 62–69. CEUR-WS.org, May 2016.

[29] Dan Brickley and R. V. Guha. RDF Schema 1.1. Recommendation, World Wide Web Consortium (W3C), February 2014. URL http://www.w3.org/TR/rdf-schema/.

[30] Dan Brickley and Libby Miller. FOAF Vocabulary Specification 0.99. Namespace document, January 2014. URL http://xmlns.com/foaf/spec/.

[31] A. Calì, G. Gottlob, and M. Kifer. Taming the Infinite Chase: Query Answering under Expressive Relational Constraints. *Journal of Artificial Intelligence Research*, 48:115–174, October 2013. DOI: 10.1613/jair.3873.

[32] Andrea Calì, Georg Gottlob, Thomas Lukasiewicz, and Andreas Pieris. Datalog+/-: A Family of Languages for Ontology Querying. In Oege de Moor, Georg Gottlob, Tim Furche, and Andrew Sellers, editors, *Datalog Reloaded*, volume 6702 of *Lecture Notes in Computer Science*, pages 351–368, Berlin, Heidelberg, March 2011. Springer. DOI: 10.1007/978-3-642-24206-9_20.

[33] Fabricio Chalub and Alexandre Rademaker. Verifying Integrity Constraints of a RDF-based WordNet. In *Global WordNet Conference*, page 309, 2016.

[34] Julien Corman, Juan L. Reutter, and Ognjen Savković. Semantics and Validation of Recursive SHACL. In Vrandečić et al. [133], pages 318–336. DOI: 10.1007/978-3-030-00671-6_19.

[35] Umutcan Şimşek, Elias Kärle, Omar Holzknecht, and Dieter Fensel. Domain Specific Semantic Validation of Schema.org Annotations. In A. Petrenko and A. Voronkov, editors, *Perspectives of System Informatics. PSI 2017*, volume 10742 of *Lecture Notes in Computer Science*, pages 417–429, Cham, 2018. Springer. DOI: 10.1007/978-3-319-74313-4_31.

[36] Richard Cyganiak, David Wood, and Markus Lanthaler. RDF 1.1 Concepts and Abstract Syntax. Recommendation, World Wide Web Consortium (W3C), February 2014. URL http://www.w3.org/TR/rdf11-concepts/.

[37] Carlos Viegas Damásio, Anastasia Analyti, Grigoris Antoniou, and Gerd Wagner. Supporting Open and Closed World Reasoning on the Web. In Jóse Júlio Alferes, James Bailey, Wolfgang May, and Uta Schwertel, editors, *Principles and Practice of Semantic Web Reasoning*, volume 4187 of *Lecture Notes in Computer Science*, pages 149–163, Berlin, Heidelberg, 2006. Springer. DOI: 10.1007/11853107_11.

[38] Claudia d'Amato, Miriam Fernandez, Valentina Tamma, Freddy Lecue, Philippe Cudré-Mauroux, Juan Sequeda, Christoph Lange, and Jeff Heflin, editors. *The Semantic Web – ISWC 2017: 16th International Semantic Web Conference, Vienna, Austria, October 21–25, 2017, Proceedings, Part I*, volume 10587 of *Lecture Notes in Computer Science*, Cham, October 2017. Springer. DOI: 10.1007/978-3-319-68288-4.

[39] Ben De Meester, Pieter Heyvaert, Anastasia Dimou, and Ruben Verborgh. Towards a Uniform User Interface for Editing Data Shapes. In Valentina Ivanova, Patrick Lambrix, Steffen Lohmann, and Catia Pesquita, editors, *Proceedings of the 4th International Workshop on Visualization and Interaction for Ontologies and Linked Data*, volume 2187 of *CEUR Workshop Proceedings*, pages 13–24. CEUR-WS.org, October 2018.

[40] Jos De Roo, Giovanni Mels, Hong Sun, and Dirk Colaert. Specialisation Mechanism for Terminology Reasoning, 2017.

[41] Jeremy Debattista, Christoph Lange, and Sören Auer. daQ, an Ontology for Dataset Quality Information. In Christian Bizer, Tom Heath, Sören Auer, and Tim Berners-Lee, editors, *Proceedings of the Workshop on Linked Data on the Web co-located with the 23rd International World Wide Web Conference (WWW 2014)*, volume 1184 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014.

[42] Jeremy Debattista, Sören Auer, and Christoph Lange. Luzzu – A Methodology and Framework for Linked Data Quality Assessment. *J. Data and Information Quality*, 8(1):4:1–4:32, October 2016. DOI: 10.1145/2992786.

[43] Jeremy Debattista, Makx Dekkers, Christophe Guéret, Deirdre Lee, Nandana Mihindukulasooriya, and Amrapali Zaveri. Data on the web best practices: Data quality vocabulary. Working group note, World Wide Web Consortium, December 2016. URL https://www.w3.org/TR/vocab-dqv/.

[44] Jeremy Debattista, Christoph Lange, Sören Auer, and Dominic Cortis. Evaluating the Quality of the LOD Cloud: An Em-

pirical Investigation. *Semantic Web Journal*, 9(6):859–901, September 2018. DOI: 10.3233/sw-180306.

[45] Kathrin Dentler, Ronald Cornet, Annette ten Teije, and Nicolette de Keizer. Comparison of Reasoners for Large Ontologies in the OWL 2 EL Profile. *Semantic Web Journal*, 2(2): 71–87, April 2011. DOI: 10.3233/SW-2011-0034.

[46] Anastasia Dimou, Miel Vander Sande, Pieter Colpaert, Ruben Verborgh, Erik Mannens, and Rik Van de Walle. RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data. In Christian Bizer, Tom Heath, Sören Auer, and Tim Berners-Lee, editors, *Proceedings of the 7th Workshop on Linked Data on the Web*, volume 1184 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014.

[47] Anastasia Dimou, Dimitris Kontokostas, Markus Freudenberg, Ruben Verborgh, Jens Lehmann, Erik Mannens, Sebastian Hellmann, and Rik Van de Walle. Assessing and Refining Mappings to RDF to Improve Dataset Quality. In Arenas et al. [4], pages 133–149. DOI: 10.1007/978-3-319-25010-6_8.

[48] Thomas Eiter, Giovambattista Ianni, and Thomas Krennwallner. Answer Set Programming: A Primer. In Sergio Tessaris, Enrico Franconi, Thomas Eiter, Claudio Gutierrez, Siegfried Handschuh, Marie-Christine Rousset, and Renate A. Schmidt, editors, *Reasoning Web. Semantic Technologies for Information Systems*, volume 5689 of *Lecture Notes in Computer Scinece*, pages 40–110, Berlin, Heidelberg, 2009. Springer. DOI: 10.1007/978-3-642-03754-2_2.

[49] Mohammed Ben Ellefi, Zohra Bellahsene, John Breslin, Elena Demidova, Stefan Dietze, Julian Szymanski, and Konstantin Todorov. RDF Dataset Profiling – a Survey of Features, Methods, Vocabularies and Applications. *Semantic Web Journal*, 9(5):677–705, August 2018. DOI: 10.3233/SW-180294.

[50] Ivan Ermilov, Jens Lehmann, Michael Martin, and Sören Auer. LODStats: The data web census dataset. In Paul Groth, Elena Simperl, Alasdair Gray, Marta Sabou, Markus Krötzsch, Freddy Lecue, Fabian Flöck, and Yolanda Gil, editors, *The Semantic Web – ISWC 2016*, volume 9982 of *Lecture Notes in Computer Science*, pages 38–46, Cham, 2016. Springer. DOI: 10.1007/978-3-319-46547-0_5.

[51] Mina Farid, Alexandra Roatis, Ihab F. Ilyas, Hella-Franziska Hoffmann, and Xu Chu. CLAMS: Bringing Quality to Data Lakes. In Fatma Özcan and Georgia Koutrika, editors, *Proceedings of the 2016 International Conference on Management of Data (SIGMOD)*, pages 2089–2092, New York, NY, United States, 2016. Association for Computing Machinery. DOI: 10.1145/2882903.2899391.

[52] Kevin Chekov Feeney, Gavin Mendel Gleason, and Rob Brennan. Linked data schemata: Fixing unsound foundations . *Semantic Web Journal*, 9(1):53–75, February 2017. DOI: 10.3233/sw-170271.

[53] Peter M. Fischer, Georg Lausen, Alexander Schätzle, and Michael Schmidt. RDF Constraint Checking. In Peter M. Fischer, Gustavo Alonso, Marcelo Arenas, and Floris Geerts, editors, *Proceedings of the Workshops of the EDBT/ICDT 2015 Joint Conference (EDBT/ICDT 2015)*, volume 1330 of *CEUR Workshop Proceedings*, pages 2015–2012. CEUR-WS.org, March 2015.

[54] Daniel Fleischhacker, Heiko Paulheim, Volha Bryl, Johanna Völker, and Christian Bizer. Detecting Errors in Numerical Linked Data Using Cross-Checked Outlier Detection. In Mika et al. [91], pages 357–372. DOI: 10.1007/978-3-319-11964-9_23.

[55] Charles L. Forgy. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19(1):17–37, September 1982. DOI: 10.1016/0004-3702(82)90020-0.

[56] Christian Fürber and Martin Hepp. Using SPARQL and SPIN for data quality management on the Semantic Web. In Witold Abramowicz and Robert Tolksdorf, editors, *Business information systems*, volume 47 of *Lecture Notes in Business Information Processing*, pages 35–46, Berlin, Heidelberg, 2010. Springer. DOI: 10.1007/978-3-642-12814-1_4.

[57] Christian Fürber and Martin Hepp. Swiqa - a semantic web information quality assessment framework. In Virpi Kristiina Tuunainen, Matti Rossi, and Joe Nandhakumar, editors, *European Conference on Information Systems (ECIS)*, page 76. Association for Information Systems, 2011.

[58] Birte Glimm and Chimezie Ogbuji. SPARQL 1.1 Entailment Regimes. Recommendation, World Wide Web Consortium (W3C), March 2013. URL https://www.w3.org/TR/sparql11-entailment/.

[59] Christophe Guéret, Paul Groth, Claus Stadler, and Jens Lehmann. Assessing Linked Data Mappings Using Network Measures. In Elena Simperl, Philipp Cimiano, Axel Polleres, Oscar Corcho, and Valentina Presutti, editors, *The Semantic Web: Research and Applications. Proceedings of the 9th Extended Semantic Web Conference (ESWC)*, volume 7295 of *Lecture Notes in Computer Science*, pages 87–102, Berlin, Heidelberg, 2012. Springer. DOI: 10.1007/978-3-642-30284-8_13.

[60] Peter Haase and Guilin Qi. An Analysis of Approaches to Resolving Inconsistencies in DL-based Ontologies. In Giorgos Flouris and Mathieu d'Aquin, editors, *Proceedings of the International Workshop on Ontology Dynamics (IWOD-07)*, pages 97–109, 2007.

[61] Steve Harris and Andy Seaborne. SPARQL 1.1 Query Language. Recommendation, World Wide Web Consortium (W3C), March 2013. URL https://www.w3.org/TR/sparql11-query/.

[62] Thomas Hartmann. *Validation Framework for RDF-based Constraint Languages*. PhD thesis, Karlsruher Institut für Technologie (KIT), 2016.

[63] Thomas Hartmann. Validation framework for rdf-based constraint languages - phd thesis appendix. Technical report, Karlsruher Institut für Technologie (KIT), 2016. URL http://digbib.ubka.uni-karlsruhe.de/volltexte/1000054062.

[64] Thomas Hartmann, Benjamin Zapilko, Joachim Wackerow, and Kai Eckert. Validating RDF Data Quality using Constraints to Direct the Development of Constraint Languages. In *IEEE Tenth International Conference on Semantic Computing (ICSC)*, pages 116–123. IEEE, 2016. DOI: 10.1109/icsc.2016.43.

[65] Patrik J. Hayes and Peter F. Patel-Schneider. RDF 1.1 Semantics. Recommendation, World Wide Web Consortium (W3C), February 2014. URL http://www.w3.org/TR/rdf11-mt/.

[66] Jelle Hellings, Marc Gyssens, Jan Paredaens, and Yuqing Wu. Implication and Axiomatization of Functional Constraints on Patterns with an Application to the RDF Data Model. In Christoph Beierle and Carlo Meghini, editors, *Foundations of Information and Knowledge Systems*, volume 8367 of *Lecture Notes in Computer Science*, pages 250–269, Cham, 2014. Springer. DOI: 10.1007/978-3-319-04939-7_12.

[67] Pieter Heyvaert, Anastasia Dimou, Ben De Meester, and Ruben Verborgh. Rule-driven inconsistency resolution for knowledge graph generation rules. *Semantic Web Journal*, 10 (6):1071–1086, October 2019. DOI: 10.3233/SW-190358.

[68] Pascal Hitzler, Markus Krötzsch, Bijan Parsia, Peter F. Patel-Schneider, and Sebastian Rudolph. OWL 2 Web Ontology Language – Primer (Second Edition). Recommendation, World Wide Web Consortium (W3C), December 2012. URL http://www.w3.org/TR/owl2-primer/.

[69] Aidan Hogan, Andreas Harth, Alexandre Passant, Stefan Decker, and Axel Polleres. Weaving the Pedantic Web. In Christian Bizer, Tom Heath, Tim Berners-Lee, and Michael Hausenblas, editors, *3rd International Workshop on Linked Data on the Web*, volume 628 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2010.

[70] Aidan Hogan, Jürgen Umbrich, Andreas Harth, Richard Cyganiak, Axel Polleres, and Stefan Decker. An Empirical Survey of Linked Data Conformance. *Journal of Web Semantics*, 14: 14–44, July 2012. DOI: 10.1016/j.websem.2012.02.001.

[71] Ian Horrocks, Peter F Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosof, and Mike Dean. SWRL: A semantic web rule language combining OWL and RuleML. Member submission, World Wide Web Consortium (W3C), May 2004. URL https://www.w3.org/Submission/SWRL/.

[72] J. M. Juran. *Juran's Quality Control Handbook*. Mcgraw-Hill, Texas, USA, 4th edition, August 1988.

[73] Michael Kifer. Nonmonotonic Reasoning in FLORA-2. In Chitta Baral, Gianluigi Greco, Nicola Leona, and Giorgio Terracina, editors, *Logic Programming and Nonmonotonic Reasoning*, volume 3662 of *Lecture Notes in Computer Science*, pages 1–12, Berlin, Heidelberg, 2005. Springer. DOI: 10.1007/11546207_1.

[74] Michael Kifer. Rule Interchange Format: The Framework. In Diego Calvanese and Georg Lausen, editors, *RR 2008: Web Reasoning and Rule Systems*, volume 5341 of *Lecture Notes in Computer Science*, pages 1–11, Berlin, Heidelberg, 2008. Springer. DOI: 10.1007/978-3-540-88737-9_1.

[75] Michael Kifer, Georg Lausen, and James Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42(4):741–843, July 1995. DOI: 10.1145/210332.210335.

[76] Michael Kifer, Jos de Bruijn, Harold Boley, and Dieter Fensel. A Realistic Architecture for the Semantic Web. In Asaf Adi, Suzette Stoutenburg, and Said Tabet, editors, *Rules and Rule Markup Languages for the Semantic Web*, volume 3791 of *Lecture Notes in Computer Science*, pages 17–29, Berlin, Heidelberg, 2005. Springer. DOI: 10.1007/11580072_3.

[77] Holger Knublauch. OWL 2 RL in SPARQL. Documentation, TopBraid. URL http://topbraid.org/spin/owlrl-all.html.

[78] Holger Knublauch and Dimitris Kontokostas. Shapes Constraint Language (SHACL). Recommendation, World Wide Web Consortium (W3C), July 2017. URL https://www.w3.org/TR/shacl/.

[79] Holger Knublauch and Pano Maria. SHACL JavaScript Extensions. Working group note, World Wide Web Consortium (W3C), June 2017. URL https://www.w3.org/TR/shacl-js/.

[80] Holger Knublauch, James A. Hendler, and Kingsley Idehen. SPIN – Overview and Motivation. Member submission, World Wide Web Consortium (W3C), February 2011. URL https://www.w3.org/Submission/spin-overview/.

[81] Holger Knublauch, Dean Allemang, and Simon Steyskal. SHACL Advanced Features. Working group note, World Wide Web Consortium (W3C), June 2017. URL https://www.w3.org/TR/shacl-af/.

[82] Dimitris Kontokostas, Patrick Westphal, Sören Auer, Sebastian Hellmann, Jens Lehmann, Roland Cornelissen, and Amrapali Zaveri. Test-driven evaluation of linked data quality. In Chin-Wan Chung, editor, *Proceedings of the 23rd international conference on World Wide Web*, pages 747–757, New York, NY, United States, April 2014. Association for Computing Machinery. DOI: 10.1145/2566486.2568002.

[83] Jose Emilio Labra Gayo, Eric Prud'hommeaux, Harold Solbrig, and Jose María Álvarez Rodríguez. Validating and Describing Linked Data Portals using RDF Shape Expressions. In Magnus Knuth, Dimitris Kontokostas, and Harald Sack, editors, *Proceedings of the 1st Workshop on Linked Data Quality co-located with 10th International Conference on Semantic Systems (SEMANTiCS 2014)*, volume 1215 of *CEUR Workshop Proceedings*. CEUR-WS.org, September 2014.

[84] Jose Emilio Labra Gayo, Eric Prud'hommeaux, Iovka Boneva, and Dimitris Kontokostas. *Validating RDF Data*, volume 7. Morgan & Claypool Publishers LLC, September 2017. DOI: 10.2200/s00786ed1v01y201707wbe016.

[85] Jose Emilio Labra Gayo, Eric Prud'hommeaux, Harold Solbrig, and Iovka Boneva. Validating and Describing Linked Data Portals using Shapes. arXiv preprint, January 2017. URL https://arxiv.org/abs/1701.08924.

[86] André Langer and Martin Gaedke. DaQAR – An Ontology for the Uniform Exchange of Comparable Linked Data Quality Assessment Requirements. In Tommi Mikkonen, Ralf Klamma, and Juan Hernández, editors, *Web Engineering*, volume 10845 of *Lecture Notes in Computer Science*, pages 234–242, Cham, 2018. Springer. DOI: 10.1007/978-3-319-91662-0_18.

[87] Huiying Li, Yuanyuan Li, Feifei Xu, and Xinyu Zhong. Probabilistic Error Detecting in Numerical Linked Data. In Qiming Chen, Abdelkader Hameurlain, Farouk Toumani, Roland Wagner, and Hendrik Decker, editors, *Database and Expert Systems Applications*, volume 9261 of *Lecture Notes in Computer Science*, pages 61–75, Cham, 2015. Springer. DOI: 10.1007/978-3-319-22849-5_5.

[88] André Melo and Heiko Paulheim. Detection of Relation Assertion Errors in Knowledge Graphs. In *Proceedings of the Knowledge Capture Conference*, pages 1–8, New York, NY, United States, 2017. Association for Computing Machinery. DOI: 10.1145/3148011.3148033.

[89] André Melo and Heiko Paulheim. Learning SHACL Constraints for Validation of Relation Assertions in Knowledge Graphs. 2018. ESWC2018 Submission.

[90] Pablo N. Mendes, Hannes Mühleisen, and Christian Bizer. Sieve: linked data quality assessment and fusion. In Divesh Srivastava and Ismail Ari, editors, *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, pages 116–123, New York, NY, United States, 2012. Association for Computing Machinery. DOI: 10.1145/2320765.2320803.

[91] Peter Mika, Tania Tudorache, Abraham Bernstein, Chris Welty, Craig Knoblock, Denny Vrandečić, Paul Groth, Natasha Noy, Krzysztof Janowicz, and Carole Goble, editors. *The Semantic Web – ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part I*, volume 8796 of *Lecture Notes*

in Computer Science, Cham, October 2014. Springer. DOI: 10.1007/978-3-319-11964-9.

[92] Boris Motik, Ian Horrocks, and Ulrike Sattler. Adding Integrity Constraints to OWL. In Christine Golbreich, Aditya Kalyanpur, and Bijan Parsia, editors, *OWL: Experiences and Directions*, volume 258 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2007.

[93] Boris Motik, Ian Horrocks, and Ulrike Sattler. Bridging the gap between OWL and relational databases. *Journal of Web Semantics*, 7(2):74–89, April 2009. DOI: 10.1016/j.websem.2009.02.001.

[94] Boris Motik, Bernardo Cuenca Grau, Ian Horrocks, Zhe Wu, Achille Fokoue, and Carsten Lutz. OWL 2 Web Ontology Language Profiles (Second Edition). Recommendation, World Wide Web Consortium (W3C), December 2012. URL https://www.w3.org/TR/owl2-profiles/.

[95] Mark A. Musen. The Protégé Project: A Look Back and a Look Forward. *AI Matters*, 1(4):4–12, June 2015. DOI: 10.1145/2757001.2757003.

[96] Sebastian Neumaier, Jürgen Umbrich, and Axel Polleres. Automated Quality Assessment of Metadata across Open Data Portals. *Journal of Data and Information Quality (JDIQ)*, 8 (2):1–29, November 2016. DOI: 10.1145/2964909.

[97] Mikael Nilsson. Description set profiles: A constraint language for dublin core application profiles. Working draft, Dublin Core Metadata Initiative (DCMI), 2008. URL http://dublincore.org/documents/2008/03/31/dc-dsp/.

[98] Chimezie Ogbuji et al. FuXi 1.4: A Python-based, bidirectional logical reasoning system for the semantic web. URL http://code.google.com/p/fuxi/.

[99] Bijan Parsia, Nicolas Matentzoglu, Rafael Gonçalves, Birte Glimm, and Andreas Steigmiller. The OWL Reasoner Evaluation (ORE) 2015 Competition Report. *Journal of Automated Reasoning*, 59:455–482, February 2017. DOI: 10.1007/s10817-017-9406-8.

[100] Adrian Paschke. Rules and Logic Programming for the Web. In Axel Polleres, Claudia d'Amato, Marcelo Arenas, Siegfried Handschuh, Paula Kroner, Sascha Ossowski, and Peter F. Patel-Schneider, editors, *Reasoning Web. Semantic Technologies for the Web of Data*, volume 6848 of *Lecture Notes in Computer Science*, pages 326–381, Berlin, Heidelberg, 2011. Springer. DOI: 10.1007/978-3-642-23032-5_6.

[101] Peter F. Patel-Schneider. Using Description Logics for RDF Constraint Checking and Closed-World Recognition. In Blai Bonet and Sven Koenig, editors, *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 247–253. AAAI Press, 2015.

[102] Peter F. Patel-Schneider. Diverging Views of SHACL, October 2016. URL https://research.nuance.com/diverging-views-of-shacl/.

[103] Heiko Paulheim. Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods. *Semantic Web Journal*, 8(3):489–508, December 2016. DOI: 10.3233/SW-160218.

[104] Heiko Paulheim and Christian Bizer. Improving the Quality of Linked Data Using Statistical Distributions. *International Journal on Semantic Web and Information Systems*, 10(2): 63–86, April 2014. DOI: 10.4018/ijswis.2014040104.

[105] Pieter Pauwels and Sijie Zhang. Semantic rule-checking for regulation compliance checking: an overview of strategies and approaches. In Jakob Beetz, Léon van Berlo, Timo Hartmann, and Robert Amor, editors, *Proceedings of the 32rd international CIB W78 conference*, 2015.

[106] Pieter Pauwels, Tarcisio Mendes de Farias, Chi Zhang, Ana Roxin, Jakob Beetz, Jos De Roo, and Christophe Nicolle. A performance benchmark over semantic rule checking approaches in construction industry. *Advanced Engineering Informatics*, 33:68–88, August 2017. DOI: 10.1016/j.aei.2017.05.001.

[107] Héctor Pérez-Urbina, Evren Sirin, and Kendall Clark. Validating RDF with OWL integrity constraints. Technical report, LLC, 2012. URL https://www.stardog.com/docs/4.1.3/icv/icv-specification.

[108] Axel Polleres, Cristina Feier, and Andreas Harth. Rules with Contextually Scoped Negation. In York Sure and John Domingue, editors, *The Semantic Web: Research and Applications: 3rd European Semantic Web Conference, ESWC 2006 Budva, Montenegro, June 11-14, 2006 Proceedings*, volume 4011 of *Lecture Notes in Computer Science*, pages 332–347, Berlin, Heidelberg, 2006. Springer. DOI: 10.1007/11762256_26.

[109] Axel Polleres, Harold Boley, and Michael Kifer. RIF Datatypes and Built-Ins 1.0 (Second Edition). Recommendation, World Wide Web Consortium (W3C), February 2013. URL https://www.w3.org/TR/rif-dtb/.

[110] María Poveda-Villalón, Asunción Gómez-Pérez, and Mari Carmen Suárez-Figueroa. OOPS! (OntOlogy Pitfall Scanner!): An on-line tool for ontology evaluation. *International Journal on Semantic Web and Information Systems (IJSWIS)*, 10(2):7–34, 2014. DOI: 10.4018/ijswis.2014040102.

[111] Eric Prud'hommeaux, Jose Emilio Labra Gayo, and Harold Solbrig. Shape expressions: an RDF validation and transformation language. In Harald Sack, Agata Filipowska, Jens Lehmann, and Sebastian Hellmann, editors, *Proceedings of the 10th International Conference on Semantic Systems*, pages 32–40, New York, NY, United States, 2014. ACM, Association for Computing Machinery. DOI: 10.1145/2660517.2660523.

[112] Eric Prud'hommeaux, Iovka Boneva, Jose Emilio Labra Gayo, and Gregg Kellogg. Shape Expressions Language 2.1. Draft community group report, World Wide Web Consortium (W3C), November 2018. URL http://shex.io/shex-semantics/.

[113] Filip Radulovic, Nandana Mihindukulasooriya, Raúl García-Castro, and Asunción Gómez-Pérez. A comprehensive quality model for Linked Data. *Semantic Web Journal*, 9(1):3–24, November 2017. DOI: 10.3233/sw-170267.

[114] Michael Röder, Ricardo Usbeck, and Axel-Cyrille Ngonga Ngomo. GERBIL – Benchmarking Named Entity Recognition and Linking Consistently. *Semantic Web Journal*, 9(5):605–625, August 2017. DOI: 10.3233/SW-170286.

[115] Edna Ruckhaus, Oriana Baldizán, and María-Esther Vidal. Analyzing Linked Data Quality with LiQuate. In Yan Tang Demey and Hervé Panetto, editors, *On the Move to Meaningful Internet Systems: OTM 2013 Workshops*, volume 8186 of *Lecture Notes in Computer Science*, pages 629–638, Berlin, Heidelberg, 2013. Springer. DOI: 10.1007/978-3-642-41033-8_80.

[116] Arthur Ryman. Resource Shape 2.0. Member submission, World Wide Web Consortium (W3C), February 2014. URL https://www.w3.org/Submission/shapes/.

[117] Arthur G. Ryman, Arnaud J. Le Hors, and Steve Speicher. OSLC Resource Shape: A language for defining constraints on Linked Data. In Christian Bizer, Tom Heath, Tim Berners-Lee, Michael Hausenblas, and Sören Auer, editors, *Proceedings of the WWW2013 Workshop on Linked Data on the Web*, volume 996 of *CEUR Workshop Proceedings*. CEUR-WS.org, May 2013.

[118] Schematron. Information technology – Document Schema Definition Languages (DSDL) – Part 3: Rule-based validation, Schematron. International standard, International Organization for Standardization, Geneva, Switzerland, 2016. URL http://schematron.com/.

[119] K. Bernhard Schiefer and Gary Valentin. DB2 universal database performance tuning. *IEEE Data(base) Engineering Bulletin*, 22(2):12–19, 1999.

[120] Shawn Simister and Dan Brickley. Simple Application-Specific Constraints for RDF Models. In *RDF Validation Workshop. Practical Assurances for Quality RDF Data, Cambridge, Ma, Boston*, 2013.

[121] Evren Sirin and Jiao Tao. Towards Integrity Constraints in OWL. In Rinke Hoekstra and Peter F. Patel-Schneider, editors, *Proceedings of the 6th International Conference on OWL: Experiences and Directions – OWLED'09*, volume 529 of *CEUR Workshop Proceedings*, pages 79–88. CEUR-WS.org, 2009. DOI: 10.5555/2890046.2890055.

[122] José Hiram Soltren. Query-based database policy assurance using semantic web technologies. mathesis, Massachusetts Institute of Technology, 2009.

[123] Slawek Staworko, Iovka Boneva, Jose Emilio Labra Gayo, Samuel Hym, Eric Prud'hommeaux, and Harold Solbrig. Complexity and Expressiveness of ShEx for RDF. In Marcelo Arenas and Martín Ugarte, editors, *LIPIcs-Leibniz International Proceedings in Informatics*, volume 31 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 195–211, Dagstuhl, Germany, 2015. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik. DOI: 10.4230/LIPIcs.ICDT.2015.195.

[124] Jiao Tao, Evren Sirin, Jie Bao, and Deborah L. McGuinness. Integrity Constraints in OWL. In Maria Fox and David Poole, editors, *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 1443–1448, Menlo Park, California, July 2010. AAAI Press.

[125] Michaël Thomazo. Compact Rewritings for Existential Rules. In Francesca Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1125–1131, Menlo Park, California, August 2013. AAAI Press.

[126] Dominik Tomaszuk. Inference rules for RDF(S) and OWL in N3Logic. arXiv preprint, 2016. URL http://arxiv.org/abs/1601.02650.

[127] Dominik Tomaszuk. RDF Validation: A Brief Survey. In Stanislaw Kozielski, Dariusz Mrozek, Paweł Kasprowski, and Bożena Małysiak-Mrozek, editors, *Beyond Databases, Architectures and Structures. Towards Efficient Solutions for Data Analysis and Knowledge Representation*, volume 716 of *Communications in Computer and Information Science*, pages 344–355, Cham, 2017. Springer. DOI: 10.1007/978-3-319-58274-0_28.

[128] Dominik Tomaszuk. Inference rules for OWL-P in N3Logic. In M. Ganza, L. Maciaszek, and M. Paprzycki, editors, *Communication Papers of the 2018 Federated Conference on Computer Science and Information Systems*, volume 17 of *ACSIS*, pages 27–33. Polskie Towarzystwo Informatyczne, September 2018. DOI: 10.15439/2018f102.

[129] Peter Van-Roy and Seif Haridi. *Concepts, Techniques, and Models of Computer Programming*. MIT Press, 2004.

[130] Pierre-Yves Vandenbussche, Ghislain A. Atemezing, María Poveda-Villalón, and Bernard Vatant. Linked Open Vocabularies (LOV): a gateway to reusable semantic vocabularies on the Web. *Semantic Web Journal*, 8(3):437–452, December 2016. DOI: 10.3233/SW-160213.

[131] Ruben Verborgh and Jos De Roo. Drawing Conclusions from Linked Data on the Web: The EYE Reasoner. *IEEE Software*, 32(5):23–27, May 2015. DOI: 10.1109/MS.2015.63.

[132] Ruben Verborgh, Dörthe Arndt, Sofie Van Hoecke, Jos De Roo, Giovanni Mels, Thomas Steiner, and Joaquim Gabarró. The Pragmatic Proof: Hypermedia API Composition and Execution. *Theory and Practice of Logic Programming*, 17(1):1–48, January 2017. DOI: 10.1017/S1471068416000016.

[133] Denny Vrandečić, Kalina Bontcheva, Mari Carmen Suárez-Figueroa, Valentina Presutti, Irene Celino, Marta Sabou, Lucie-Aimée Kaffee, and Elena Simperl, editors. *The Semantic Web – ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, October 8–12, 2018, Proceedings, Part I*, volume 11136 of *Lecture Notes in Computer Science*, Cham, 2018. Springer. DOI: 10.1007/978-3-030-00671-6.

[134] Mustafa Yuksel, Suat Gonul, Gokce Banu Laleci Erturkmen, Ali Anil Sinaci, Paolo Invernizzi, Sara Facchinetti, Andrea Migliavacca, Tomas Bergvall, Kristof Depraetere, and Jos De Roo. An Interoperability Platform Enabling Reuse of Electronic Health Records for Signal Verification Studies. *BioMed Research International*, 2016:1–18, March 2016. DOI: 10.1155/2016/6741418.

[135] Amrapali Zaveri, Anisa Rula, Andrea Maurino, Ricardo Pietrobon, Jens Lehmann, and Sören Auer. Quality assessment for linked data: A survey. *Semantic Web Journal*, 7(1): 63–93, March 2015. DOI: 10.3233/SW-150175.

[136] Hong Zhu, Patrick A. V. Hall, and John H. R. May. Software unit test coverage and adequacy. *ACM Computing Surveys*, 29 (4):366–427, December 1997. DOI: 10.1145/267580.267590.