

Describing Reasoning Results with RVO, the Reasoning Violations Ontology

Bojan Božić, Rob Brennan, Kevin C. Feeney, Gavin Mendel-Gleason

Knowledge and Data Engineering Group,
Trinity College Dublin
College Green
Dublin 2

{bojan.bozic, rob.brennan, kevin.feeney, mendelgg}@scss.tcd.ie

Abstract. This paper presents a new OWL RL ontology, the Reasoning Violations Ontology (RVO), which describes both ABox and TBox reasoning errors produced by DL reasoners. This is to facilitate the integration of reasoners into data engineering tool-chains. The ontology covers violations of OWL 2 direct semantics and syntax detected on both the schema and instance level over the full range of OWL 2 and RDFS language constructs. Thus it is useful for reporting results to other tools when a reasoner is applied to linked data, RDFS vocabularies or OWL ontologies, for example for quality evaluations such as consistency, completeness or integrity.

RVO supports supervised or semi-supervised error localisation and repair by defining properties that both identify the statement or triple where a violation is detected, and by providing context information on the violation which may help the repair process. In a case study we show how the ontology can be used by a reasoner and a supervised repair process to accelerate high quality ontology development and provide automated constraint checking feedback on instance data. RVO is also being used to enable integration of reasoning results into multi-vendor data quality tool chains within the ALIGNED H2020 project.

Keywords: Ontology Engineering, Data Integrity, Consistency Checking, Reasoning.

1 Introduction

A common task performed with Semantic Web reasoners is the detection and reporting of errors or inconsistencies found in an ontology. This task frequently occurs within the ontology authoring, interlinking, classification, quality analysis and evolution phases of the linked data lifecycle [1]. However, to manage the entire data lifecycle or even the full range of activities within a single lifecycle stage, typically requires many tools to be integrated into a tool-chain. Current research on standard mechanisms for linked data tool-chain integration is still immature. However, the use of ontologies to support the interchange of data for tool-chain integration has been explored [2]. The existing gap in related work, as we show in Section 2 is the absence of a complete ontology for modelling reasoning errors. This would enable the integration

of standardised ontology violation detection services into linked-data management tool-chains. Other relevant work in the field focuses on pitfalls, prevention of mistakes and general ontology design and structural issues.

The two goals that have driven this research are firstly, to produce a service which will assist in high-quality ontology development by identifying reasoning violations and producing semi-automated repair recipes. Secondly, to produce a service which can detect constraint violations on instance data according to a schema (ontology), and produce semi-automated repair recipes. In order to support these services, we need a rich, highly structured, general purpose way of expressing reasoning violations.

This paper presents the RVO (Reasoning Violations Ontology). It describes OWL and RDF(S) reasoning errors, in two categories: those involving classes, properties and axioms (schema / TBox) and those involving instances (ABox). This ontology is used by a custom reasoner implemented in SWI-Prolog, the Dacura Quality Service [3], to report the errors that it detects to other linked data lifecycle tools. Our first client application is the Dacura Schema Manager [4], which can consume RVO and presents the reasoner output with filtered, categorised, detailed error information, as well as information about the source of the error.

The principal contributions of this paper are: a description of the published RVO ontology, a discussion of the integration of RVO in our own data lifecycle web platform Dacura and documentation of the violation identification process for ontology developers.

The rest of the paper is structured as follows: Section 2 is dedicated to related work. Section 3 presents the RVO ontology and provides insights into its design. Section 4 validates the ontology in a toolchain integration case. Finally, Section 5 concludes the paper and provides an answer to the research question as well as an outlook on future work.

2 Related Work

To the best of our knowledge, there are no ontologies that have been developed for the specific purpose of describing reasoning violations. However, similar problems have been addressed from differing perspectives.

One of the most relevant contributions is OOPS! [5] which is a tool with a catalogue for validating ontologies by spotting common pitfalls. The catalogue contains 41 pitfalls which the tool checks for. The ontology can be inserted directly in a text-box or referenced by URI. Although, OOPS! identifies many common pitfalls, it detects design flaws rather than logical errors and does not use an ontology for error reporting.

Other research [6] has identified the types of flaws that can occur in the object property box and proposed corresponding compatibility services. However, this work is very specific and focuses on properties and their compatibility. Our approach addresses a far broader palette of violations, across the ABox and TBox, incorporating class and property violations.

In [7], a very similar approach to OOPS! was proposed, covering logical and non-logical anti-patterns, but it is quite limited as it covers only 4 logical and 3 non-logical anti-patterns as well as 4 guidelines.

The work presented in [8] will be combined with our Reasoning Violations Ontology in order to extend Linked Data Quality in an ALIGNED project use case. We have also published differences between SHACL and RVO in the deliverable [15].

The Shapes Constraint Language (SHACL) introduced in [9], is a language for describing and constraining the contents of RDF graphs. As part of its ongoing development, through the W3C's RDF Data Shapes Working Group, it is defining a standard error reporting format. Our ontology can be considered as an extension of SHACL's error reporting, as it can express a superset of the violations that can be expressed in SHACL, also while SHACL detects bad triples which caused an error, RVO is able to detect a whole subgraph which was involved in producing the violation. We plan to link RVO to SHACL errors through reuse of their predicates once the format achieves standardisation and stability.

Another W3C vocabulary is EARL¹ which can be used for validation results. Although it has been defined in the context of validating accessibility tools, it contains several terms for describing validation results in RDF.

There are also some publications about preventing errors in ontology development, such as [10]. They are extremely useful for defining best practices and fueling the discussion about ontology engineering style and error prevention, but they provide no insights into error reporting for existing ontologies. Closer to that is a publication about debugging OWL ontologies [11]. They have integrated a number of simple debugging cues generated from the description logic reasoner, Pellet², in the hyper-textual ontology development environment, Swoop³.

3 The Reasoning Violations Ontology (RVO)

The purpose of RVO is to enable a reasoner to describe reasoning errors detected in an input ontology in order to facilitate the integration of reasoners into semantic web toolchains. It is defined as a simple OWL 2 ontology that is amenable to RDFS-based interpretations or use as a linked data vocabulary without any dependence on reasoning. In future, an RDFS version of the ontology is planned, in order to support interpretation by RDFS reasoners. A permanent identifier for the ontology has been registered with the W3C permanent identifier community group. The full source of the ontology is published online⁴ and meta-data have been added to facilitate LODE-based [12] documentation generation⁵.

This ontology is used to describe RDF and OWL reasoning violation messages in the Dacura Quality Service [13]. These are generated by running an

¹ <https://www.w3.org/TR/EARL10-Schema/>

² <https://github.com/complexible/pellet>

³ <https://github.com/ronwalf/swoop>

⁴ <https://w3id.org/rvo>

⁵ <http://www.essepuntato.it/lode/closure/reasoner/https://w3id.org/rvo>

RDF/RDFS/OWL-DL reasoner over an RDF-based ontology model and allowing the Dacura quality service to report any integrity violations detected at schema or instance level. These violations report areas where the input model is logically inconsistent or breaks RDFS/OWL semantics or axioms. Violations may be reported as based on open world or closed world assumptions. The open world is the default OWL semantics and can typically only detect a limited number of problems due to incomplete knowledge. The closed world interpretation assumes that you have provided all relevant aspects of the model and is able to detect a much wider range of violations, e.g. missing or misspelled term definitions. This is often useful during ontology development or in a system that interprets OWL as a constraint language.

3.1 The Ontology

The ontology can be divided into two layers. The top layer consists of the base classes and their properties and the bottom layer is a vocabulary which defines the hierarchical structure of violations identified so far (see 3.2).

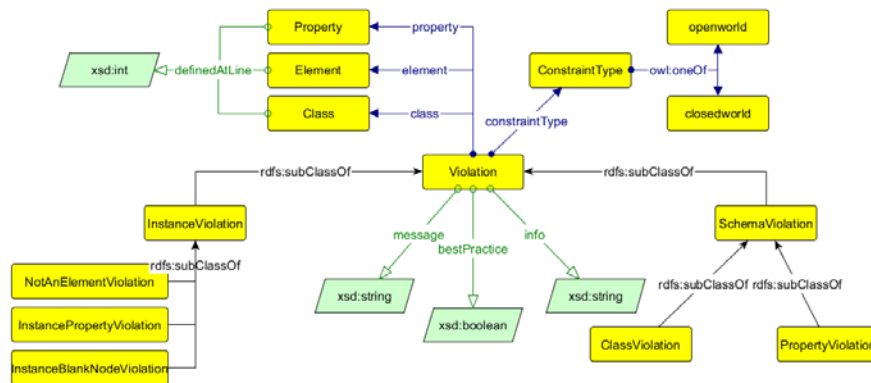


Figure 1 Base Classes and Properties in RVO

Figure 1 shows the top tier of the ontology which represents general metadata about a Violation as well as related properties, elements and classes.

Class violations are used for reporting issues regarding the TBox and instance violations ABox in general. Therefore, class violations are reported when e.g. property domains are missing, subsumption errors are detected, or class and property cycles are found. Instance violations show instances which are not elements of valid classes, cardinalities which are incorrect, property constraints that are violated, literals and objects which are confused, etc.

3.2 Error Classes

We have organised reasoning errors in violation classes and put them in a hierarchical structure. The top level differentiation of violations is instance or schema, de-

pending on whether the violation occurred in the ABox (instances) or TBox (schema). Here is an overview of all errors currently documented:

- **Instance:**
 - o InstanceBlankNode
 - o NotAnElement: NotRestrictionElement, ObjectInvalidAtClass, EdgeOrphanInstance, DataInvalidAtDatatype (NotBaseTypeElement)
 - o InstanceProperty: NoPropertyDomain, InvalidEdge, NotFunctionalProperty, LocalOrphanProperty, NotInverseFunctionalProperty
- **Schema:**
 - o ClassViolation: NotUniqueClassLabel, NotUniqueClassName, NotDomainClass, ClassCycle, NoImmediateClass (NotSuperClassOfClass), OrphanClass (NotIntersectionOfClass, NoSubclassOfClass, NotUnionOfClass)
 - o PropertyViolation: PropertyCycle, NotUniquePropertyName, SchemaBlankNode, PropertyTypeOverload, PropertyAnnotationOverload, OrphanProperty (NotSubpropertyOfProperty), PropertyDomain (InvalidDomain, DomainNotSubsumed, NoExplicitDomain), PropertyRange (InvalidRange, RangeNotSubsumed, NotExplicitRange)

3.3 Example of RVO in Use – Class Violation

In our example a reasoning error is asserted first in JSON as raw data and then converted to RDF triples using RVO in order to be consumed in Dacura. The example shows a `ClassViolation` which is a `SchemaViolation` and more specifically a `ClassCycleViolation`. Such specific violation detection results make it possible to provide exact suggestions to ontology developers or repair agents and trigger ontology improvements.

```
{ "rdf:type", "ClassCycleViolation",
  "bestPractice": { "type": "xsd:boolean", "data": "false" },
  "message": "Class UnitOfSocialOrganisation has a class cycle with
  path: [TemporalEntity,UnitOfSocialOrganisation]",
  "path": [
    "http://dacura.cs.tcd.ie/data/seshat#TemporalEntity",
    "http://dacura.cs.tcd.ie/data/UnitOfSocialOrganisation"
  ],
  "class": "http://dacura.cs.tcd.ie/data/seshat#UnitOfSocialOrganisation" }
```

The response RDF graph provides a much better way to interpret the results:

The instance is called *_example1* and is a *ClassCycleViolation*. *bestpractice* is false, so it is an error rather than a warning. The message provides a summary of the cause of the violation, but the important parts are the next two properties. The *path* property marks all classes which were involved in the cycle and the *class* property marks the class where the cycle has been detected. Another important feature is that RVO provides us direct links to the Seshat ontology (an ontology which models the Seshat Global History Databank⁶) and hence to the OWL classes from the external ontology which were involved in the violation process.

⁶ <https://evolution-institute.org/project/seshat/>

4 Validation through Integration into ALIGNED

The Dacura Toolchain Case Study covers toolchaining and reporting to users. At this point we want to have a closer look into supervised ontology repair and publishing.

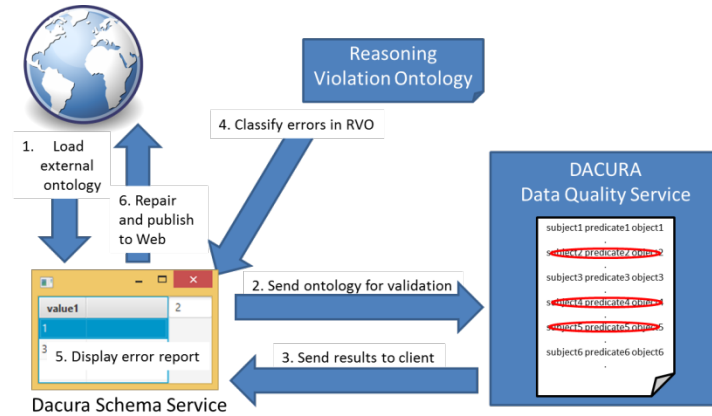


Fig. 2. Dacura Ontology Repair Use Case

In the use case scenario in Figure 2, the Dacura Shema Service gets an external ontology which is loaded by the Dacura Schema Manager and needs to be validated into the knowledge base (step 1). The ontology is then send to DQS and checked for errors by the DQS reasoner (step 2), whose rules comply with the violation classes of RVO. The detected violations are sent back to the client (step 3). The results are classified and assigned classes from RVO (step 4). Dacura then integrates the report in the UI and presents it to the user (step 5). Finally, the user can repair the ontology in an editor and republish it to the Web (step 6).

So how did RVO help in this specific scenario? After the Data Quality Service checked the external ontology for errors, it used the RVO structure to provide information about a specific violation by creating individuals. This provided us with a classification of DQS' results. RVO has been used by the client for classification of the errors and for providing relevant information about violations to the user. An additional benefit would be to archive the history of errors in a knowledge base and be able to query for certain occurrences of violations for an ontology.

Table 1. Ontology validation results

	DIO	DLO	SLO	SIP	EIPDM	SDO	DIOPP
# of triples	263	127	56	126	147	102	145
# of errors	0	1	25	10	37	6	4
# of warnings	4	11	5	5	5	7	11
h for validation	1	1	1	1	1	1	1

h for correction	1	2	4	2	4	2	2
------------------	---	---	---	---	---	---	---

Table 1 shows the validation of ALIGNED ontologies⁷ which have all been developed by using different approaches (Protégé, RDF2RDF, human checks, etc.). We have used the Data Quality Service and RVO to validate several project ontologies and report the errors and warnings found in a first run. Although, this is only a case study and especially the validation and correction efforts are estimated, the table signals the potential for improvement of existing ontologies.

5 Conclusion and Future Work

In this paper we have shown that a dedicated reasoning error ontology improves error reporting with structured data, and integration of the ontology in the Dacura toolchain case study.

The Reasoning Violations Ontology not only benefits the interpretation and further processing of reasoning errors in tools, platforms, and Web UIs which present results of the reasoning process or ontology validation, but can also be used as a common format to represent violations found in ontologies during the whole software toolchain process. We have shown an example of supervised ontology repair use case and explained the advantages of our approach. Furthermore, we have given some examples for reasoning violations and constructed RDF graphs to present the results.

Our future work will continue with the integration of the ontology in the ALIGNED toolchain and linking of the ontology to SHACL constraints as well as using it together with RDFUnit. Finally, we plan to evaluate the benefits in a case study with ontology engineers and investigate their work with the ontology and our tools in order to improve the quality of their ontology or repair it.

Acknowledgement

This research has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 644055, the ALIGNED project (www.aligned-project.eu) and from the ADAPT Centre for Digital Content Technology, funded under the SFI Research Centres Programme (Grant 13/RC/2106) and co-funded by the European Regional Development Fund.

References

- [1] S. Auer, L. Bühmann, C. Dirschl, O. Erling, M. Hausenblas, R. Isele, J. Lehmann, M. Martin, P. N. Mendes and B. Van Nuffelen, »Managing the life-cycle of linked data with the LOD2 stack,« in *International Semantic Web Conference*, Boston, 2012.

⁷ <http://aligned-project.eu/data-and-models/>

- [2] C. Keßler, M. d'Aquin and S. Dietze, »Linked Data for Science and Education,« in *Semantic Web*, pp. 1-15, 2013.
- [3] G. Mendel-Gleason, K. Feeney and R. Brennan, »Ontology Consistency and Instance Checking for Real World Linked Data,« in *Extended Semantic Web Conference*, Portorož, 2015.
- [4] K. Feeney, G. Mendel-Gleason and R. Brennan, »Linked data schemata: fixing unsound foundations (submitted),« in *Semantic Web Journal - Special Issue on Quality Management of Semantic Web Assets*, 2015.
- [5] M. Poveda Villalon, M. C. Suárez-Figueroa and A. Gómez-Pérez, »Validating ontologies with OOPS!,« in *18th International Conference on Knowledge Engineering*, Galway, 2012.
- [6] M. Keet, »Detecting and revising flaws in OWL object property expressions,« in *18th International Conference on Knowledge Engineering and Knowledge*, Galway, 2012.
- [7] C. Roussey, O. Corcho and L. M. Vilches-Blázquez, »A catalogue of OWL ontology antipatterns,« in *International Conference on Knowledge Capture*, 2009.
- [8] D. Kontokostas, P. Westphal, S. Auer, S. Hellmann, J. Lehmann, R. Cornelissen and A. Zaveri, »Test-driven evaluation of linked data quality,« in *Proceedings of the 23rd international conference on World Wide Web*, 2014.
- [9] A. Ryman, »Z Specification for the W3C Editor's Draft Core SHACL Semantics,« in arXiv preprint arXiv:1511.00384, 2015.
- [10] A. Rector, N. Drummond, M. Horridge, J. Rogers, H. Knublauch, R. Stevens, H. Wang and C. Wroe, "OWL pizzas: Practical experience of teaching OWL-DL: Common errors & common patterns," in *Engineering Knowledge in the Age of the Semantic Web*, Springer Berlin Heidelberg, 2004, pp. 63-81.
- [11] B. Parsia, E. Sirin and A. Kalyanpur, »Debugging OWL ontologies,« in *14th international conference on World Wide Web*, Chiba, 2005.
- [12] S. Peroni, D. Shotton and F. Vitali, »Tools for the automatic generation of ontology documentation: a task-based evaluation,« in *International Journal on Semantic Web and Information Systems (IJSWIS)*, pp. 21-44, 2013.
- [13] G. E. Mendel-Gleason, R. Brennan and K. C. Feeney, »Ontology Consistency and Instance Checking for Real World Linked Data,« in *2nd Workshop on Linked Data Quality*, Portorož, 2015.
- [14] A. W. Appel, »Verified Software Toolchain,« *Programming Languages and Systems*, Saarbrücken, Springer, 2011, pp. 1-17.
- [15] R. Brennan, B. Božić, M. Solanki, D. Kontokostas, A. Koller and C. Dirschl, »D2.7 Meta-model Phase 2,« in *Public ALIGNED project deliverable*.