



PS 703301 – WS 2021/22
Current Topics in Computer Science

Final Report

Shaping Knowledge Graphs

Philipp Gritsch
Jamie Hochrainer
Kristina Magnussen
Danielle McKenney
Valerian Wintner

supervised by
M.Sc. Elwin Huaman

Contents

1. Introduction	3
2. Related Work	3
3. Approach	3
3.1. Technology Stack	3
3.2. Generating Constraints	3
3.2.1. Integrating RDF2Graph with our framework	4
3.3. Validating Constraints	5
3.4. Front-end	5
4. Results	6
4.1. Future work	6
5. Evaluation	7
6. Conclusion	7
References	7
A. Contribution Statements	7
B. Appendix	7

I added a comment colour for everyone. ;Comment colour Jamie; ;Comment colour Danielle;
;Comment colour Philipp; ;Comment colour Valerian; ;Comment colour Kristina;

remove this
part

1. Introduction

We used *CommonCrawl* datasets as the base for the *knowledge graph* which we wanted to assess. The data contained in those datasets is often inconsistent and might contain errors. In order to work with this data properly, it is necessary to shape the *knowledge graph* in which this data is contained. This shaping is done by inferring constraints over the data and validating it based on these constraints. Validating a graph against constraints gives important insight into the structure of the data. For instance, when all nodes of a type conform to constraints, then it may be useful to define these as required attributes for all future nodes to ensure uniformity in the data. Non conforming nodes may also deliver important insight into where information is missing. For example, if 99% of nodes of a given type conform to some constraints, it may be worthwhile to investigate the remaining 1% to see if they are missing necessary information or otherwise corrupt.

add introduction to KGs

2. Related Work

3. Approach

Our framework offers a way to evaluate a *knowledge graph* in an automated way. For this, we used *knowledge graphs* from the *CommonCrawl* datasets as a basis. The *knowledge graphs* are imported as a static file. After this, our framework infers constraints over this data set (see Section 3.2). These are validated automatically in the last step, see Section 3.3. The user can interact with this framework over the front-end, see Section 3.4. These different steps were implemented and tested separately. Once this was done, we consolidated them. The structure of our project can be seen in Fig. 1.

Add thesis
Werkmeister
RDF2Graph,
also add another work,
maybe from sources in thesis, done by Philipp

3.1. Technology Stack

The framework was implemented in *Java*. We used *Maven* as a project management tool. We also used *Jena*, which offers an *RDF* API as well as support for *SPARQL* queries and the *ShEx* language. The front-end was implemented using *Vue3*[1] as a front-end framework and *PrimeVue* as a library for the different components. For the deployment of our application we use single virtual machine. Access to the front-end is done via a single *Apache* server. The front-end accesses the back-end via an internal *REST-API*.

update figure

refer to the
readme here?
Or should this
happen somewhere else?

add reference
to our github
repo!

3.2. Generating Constraints

For the generation of constraints, we used the tool *RDF2Graph* [2] and adapted it for our purposes. As input, *RDF2Graph* takes a *knowledge graph* from *CommonCrawl*. The properties of the graph are read out with several *SPARQL* queries. These properties are saved in a new *RDF* graph. As output, we receive a graph containing constraints for the initial input data. We use *RDF2Graph* queries to extract the constraints in *ShEx* syntax.

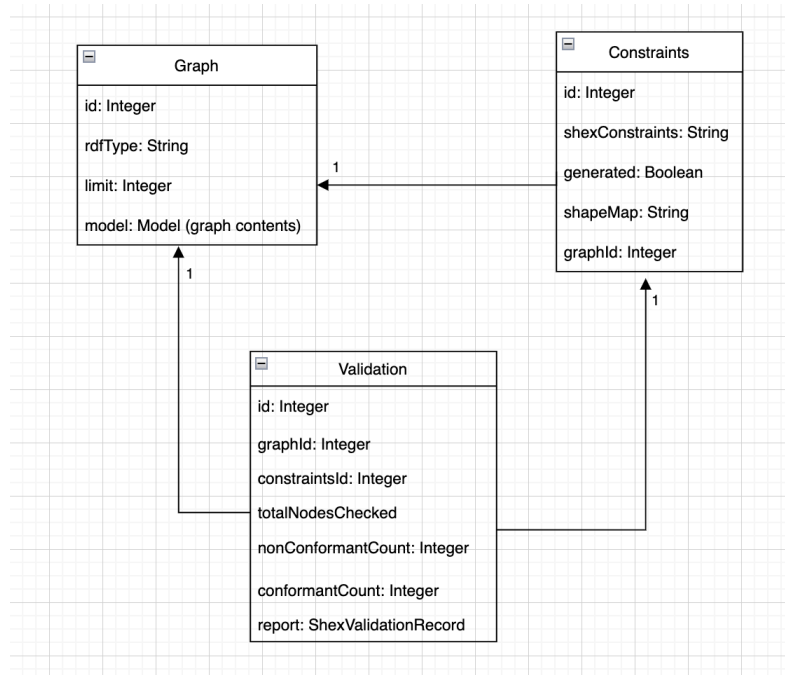


Figure 1: UML diagram of the framework structure



add query to graph (chosen by Philipp), e.g. multiplicity of argument etc.

3.2.1. Integrating RDF2Graph with our framework

We implemented the following steps in order to integrate *RDF2Graph* into our project. We added *RDF2graph* to our framework so that they could be compiled together. In addition, we changed some of the initial parameters of the *RDF2Graph*, since it originally was intended as a stand-alone application. As we are handling *Models* in our software, we changed the input to *RDF2Graph* to a *Model*. In our application, *RDF2Graph* does not use any other storage apart from the *Model* data structure. Previously, such a *Model* needed to be created by *RDF2Graph*, now it is provided by our framework. We did this so we could have full control over the files handled by *RDF2Graph*. *RDF2Graph* allows for multithreaded execution, which requires a thread pool. This thread pool was initially created by *RDF2Graph*. In our framework, it is provided by our application. In addition, resources which are used by *RDF2Graph* had to be provided in a different way so that they are still available when running from a server environment. We also changed some of the queries. *RDF2Graph* supports multiple output graphs, however, this did not work. As we only work on one *Model* at a time, we only use one output graph.

Add explanation for *Model*? May in glossary?

should we explain this in more detail?

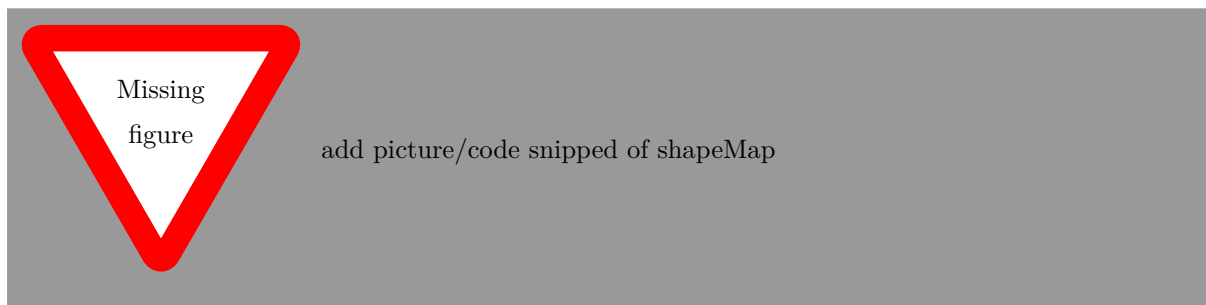
Add explanation of limit this section?

3.3. Validating Constraints

Given a *RDF* graph and a set of constraints, the validation consists of verifying that every node in the graph fulfils the requirements given in the constraints. A graph consists of several different types. Each of those types must conform to its definition outlined in the constraints. The results of the validation is be a boolean flag for every single node in the graph, indicating whether or not it conforms to its type's constraints. In case of nonconformity, a reason will be given.

In our code, this is implemented in the following way. As input, we receive a *RDF* subgraph as well as a set of constraints. We use this to generate a *shape map*, which contains all of the types which need to be validated. For the actual validation, the *ShExValidator* provided by the *Jena* library was used. The validator requires a set of constraints defined in valid *ShEx* syntax and a *shape map*. The *shape map* describes which types of nodes need to be validated against which *ShEx* constraint definitions.

add reference
to Jena libra
here?

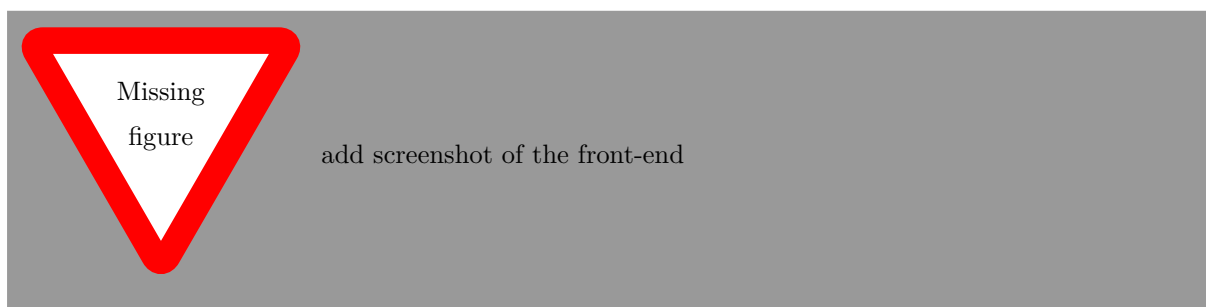


The class *ShexValidationRecord* stores the result of the validation for every single node of the graph. Not only is the individual result of every node checked against its relevant constraints, but we also calculate the percentage of nodes that conform to their constraints.

3.4. Front-end

We implemented a front-end where the user can choose a *knowledge graph* as well as a type of knowledge graph. In addition, the user can also set a limit. As output, *ShEx* constraints as well as a validation of those constraints are given. The constraints can be edited by the user and those edited constraints can be revalidated. If a node is deemed invalid, a reason is given, e.g. "Cardinality violation (min=1): 0" The user can download the subset of the graph which was validated. The interaction between user, front-end and server can also be seen in Fig. 2.

check whether
this is what
we are doing
in the finished
version



explain this
in more detail,
maybe also put
the explanation
in the query

explain how
different limits
influence the
output

update and
scale sequence
diagram and
refer to it

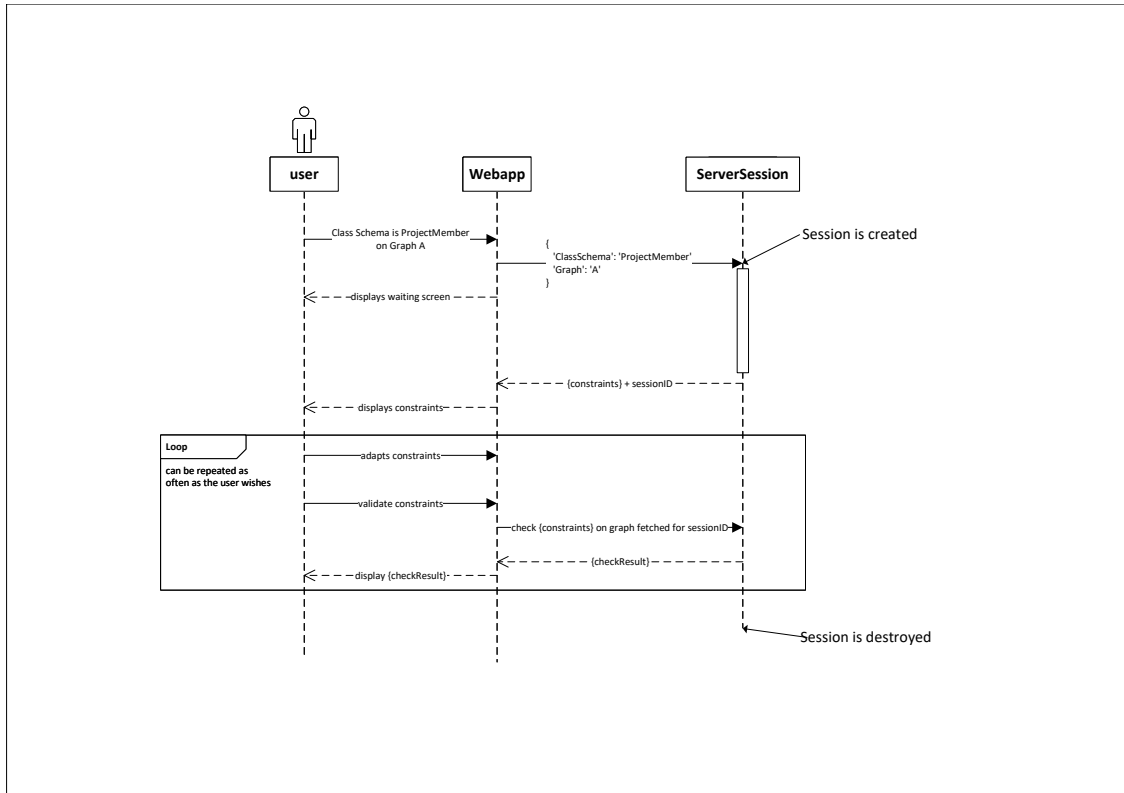



Figure 2: Sequence diagram showing the interaction between web application, user and server

4. Results

Our framework automatically infers constraints and validates the given data based on those constraints. This can be done on two different *CommonCrawl* datasets. The user can choose one of those datasets and a limit using the front-end. User can also edit constraints.



Missing figure

Maybe add small figure that shows workflow of project here? Something similar like we did in presentation but more professional?

explain this limit in more depth, maybe in front-end?

4.1. Future work

Our application currently only handles two different datasets. For future work, this could be expanded so that the framework could handle more and bigger datasets. Currently, the size of the datasets that can be handled is limited by the RAM on the virtual machine. One possible solution for this could be to only work on parts of the graph. One problem we encountered when handling datasets from *CommonCrawl* was the quality of these datasets. Many datasets include *non-unicode* characters, which are replaced by Jena with *unicode* characters. This takes a lot of computing time. In addition, many files include invalid

describe results of benchmark tests here

Possible future work could be: more datasets, more possibilities for user inputs

RDF syntax or are otherwise damaged. This means that in order to handle additional datasets, some way of processing these datasets would have to be implemented. Processing could include filtering for broken files and invalid syntax and fixing this before handling the dataset in the framework. In addition, more possibilities for user interaction could be added. For instance, a feature could be added where a user can upload their own dataset and have it validated.

5. Evaluation

6. Conclusion

References

- [1] Vue.js, documentation. <https://v3.vuejs.org/>.
- [2] J. van Dam. Rdf2graph. <https://github.com/jessevdam/RDF2Graph/>, 2022. Accessed: 2022-01-16.

A. Contribution Statements

Please write down a short contribution statement for each member of your group. You may evaluate the contribution along the three common categories: i) conception (i.e., problem framing, ideation, validation, and method selection), ii) operational work (e.g., setting up your tech stack, algorithm implementation, data analysis, and interpretation), and iii) writing & reporting (i.e., report drafting, literature review, revision of comments, presentation preparations, etc.).

B. Appendix

You may use appendices to include any auxiliary results you would like to share, however cannot insert in the main text due to the page limit.

Should we add proper SPARQL endpoints here? Might not be possible?

add benchmarks here

check what Ewin said concerning Evaluation on meeting 20.01.2022

Which challenges did we face during the implementation? (Maybe depth of SPARQL query, outdated RDF2Graph)

Did we achieve what we wanted to do? How well and reliably does the framework work?