



# Q<sup>2</sup>Forge: Minting Competency Questions and SPARQL Queries for Question-Answering Over Knowledge Graphs

Yousouf Taghzouti, Franck Michel, Tao Jiang, Louis-Félix Nothias, Fabien Gandon

## ► To cite this version:

Yousouf Taghzouti, Franck Michel, Tao Jiang, Louis-Félix Nothias, Fabien Gandon. Q<sup>2</sup>Forge: Minting Competency Questions and SPARQL Queries for Question-Answering Over Knowledge Graphs. 2025. hal-05070442

**HAL Id: hal-05070442**

**<https://hal.science/hal-05070442v1>**

Preprint submitted on 16 May 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

# Q<sup>2</sup>Forge: Minting Competency Questions and SPARQL Queries for Question-Answering Over Knowledge Graphs

Yousouf Taghzouti<sup>1</sup>[0000–0003–4509–9537], Franck Michel<sup>2</sup>[0000–0001–9064–0463],  
Tao Jiang<sup>3</sup>[0000–0002–5293–3916], Louis-Felix Nothias<sup>3</sup>[0000–0001–6711–6719], and  
Fabien Gandon<sup>4</sup>[0000–0003–0543–1232]

<sup>1</sup> Univ. Côte d’Azur, Inria, ICN, I3S, France  
[yousouf.taghzouti@univ-cotedazur.fr](mailto:yousouf.taghzouti@univ-cotedazur.fr)

<sup>2</sup> Univ. Côte d’Azur, CNRS, Inria, I3S, France [franck.michel@inria.fr](mailto:franck.michel@inria.fr)

<sup>3</sup> Univ. Côte d’Azur, CNRS, ICN, France  
[{tao.jiang,louis-felix.nothias}@cnrs.fr](mailto:{tao.jiang,louis-felix.nothias}@cnrs.fr)

<sup>4</sup> Inria, Univ. Côte d’Azur, CNRS, I3S, France [fabien.gandon@inria.fr](mailto:fabien.gandon@inria.fr)

<b>Resource type</b>	Reusable Software and Services		
<b>License</b>	GNU Affero General Public License v3.0 or later (AGPL-3.0-or-later)		
<b>Q<sup>2</sup>Forge</b>	<b>Repo</b>	<a href="https://github.com/Wimmics/q2forge">https://github.com/Wimmics/q2forge</a>	<b>DOI</b> <a href="https://doi.org/10.5281/zenodo.15388693">10.5281/zenodo.15388693</a>
<b>Gen<sup>2</sup>KGBot</b>	<b>Repo</b>	<a href="https://github.com/Wimmics/gen2kgbot">https://github.com/Wimmics/gen2kgbot</a>	<b>DOI</b> <a href="https://doi.org/10.5281/zenodo.15388687">10.5281/zenodo.15388687</a>
<b>Demo video</b>	<b>Teaser</b>	<a href="https://youtu.be/E9rgCZzWH4k">https://youtu.be/E9rgCZzWH4k</a>	<b>Full</b> <a href="https://youtu.be/I3w-jmZRJII">https://youtu.be/I3w-jmZRJII</a>

**Abstract.** The SPARQL query language is the standard method to access knowledge graphs (KGs). However, formulating SPARQL queries is a significant challenge for non-expert users, and remains time-consuming for the experienced ones. Best practices recommend to document KGs with competency questions and example queries to contextualise the knowledge they contain and illustrate their potential applications. In practice, however, this is either not the case or the examples are provided in limited numbers. Large Language Models (LLMs) are being used in conversational agents and are proving to be an attractive solution with a wide range of applications, from simple question-answering about common knowledge to generating code in a targeted programming language. However, training and testing these models to produce high quality SPARQL queries from natural language questions requires substantial datasets of question-query pairs. In this paper, we present Q<sup>2</sup>Forge that addresses the challenge of generating new competency questions for a KG and corresponding SPARQL queries. It iteratively validates those queries with human feedback and LLM as a judge. Q<sup>2</sup>Forge is open source, generic, extensible and modular, meaning that the different modules of the application (CQ generation, query generation and query refinement) can be used separately, as an integrated pipeline, or replaced by alternative services. The result is a complete pipeline from competency question formulation to query evaluation, supporting the creation of reference query sets for any target KG.

**Keywords:** Competency Question · SPARQL Query · LLM

## 1 Introduction

Semantic technologies, and in particular knowledge graphs (KGs), have been utilised in a variety of applications over time, including search engines, data integration, enterprise settings and machine learning. Numerous methods were proposed to assist their life-cycle and exploitation [12] leading to their adoption and the rapid growth of the Linked Open Data (LOD) cloud.<sup>5</sup> However, the exploitation of these KGs has been hindered by the steep learning curve associated with the stack of standards, in particular query languages such as SPARQL [19].

Over the last few years, common information retrieval methods have been profoundly renewed by the emergence of pre-trained Large Language Models (LLMs). The abilities of LLMs to understand and generate natural language (NL) and code alike have opened new research and development fields notably in the domain of data access and interaction. In particular, these abilities endow LLMs with the capacity to translate a question expressed in NL into its counterpart in a structured query language, SPARQL in the case of RDF KGs. This allows domain experts to “speak to structured data” thus facilitating data access. To design and evaluate such text-to-SPARQL translation systems effectively, we need **reference datasets providing curated question-query pairs** that are either tailored to a specific KG or at least relevant for the domain it concerns.

Some question-query datasets (**Q<sup>2</sup>sets**) have been produced in the context of benchmarks and challenges such as QALD [18], DBNQA [11], and LC-QuAD [7], but they are mostly based on subsets of DBpedia and/or Wikidata. When it comes to other domain-specific, possibly private KGs, or highly specialized KGs like in life sciences, creating a Q<sup>2</sup>set involves skills that are rarely mastered by one and the same person. More likely, this requires the collaboration of domain experts who can think of possibly complex competency questions (CQ) that scientists may want to ask, and Semantic Web experts who shall leverage the used ontologies and KG schema to come up with counterpart SPARQL queries.

KGs should ideally come with examples of queries they support, a good practice in terms of documentation and metadata. Yet in practice this is rarely the case. Similarly, whereas CQs have been identified as a valuable documentation and starting point for understanding the capabilities of a KG, many KGs are accompanied by very few CQs, if any at all.

To support the creation of Q<sup>2</sup>sets for training, testing, benchmarking, and documenting our systems and knowledge graphs, we identified the need to provide tools that help researchers—as well as scientific and technical information professionals—to understand existing KGs and generate or refine corresponding Q<sup>2</sup>sets, whether they are Semantic Web newcomers or experienced practitioners. While various methods and tools exist to help and create CQs and equivalent queries [1,4,15,24,16,5,8]. However, to the best of our knowledge, these tools are either domain-specific, extensively manual, or address only specific steps, but do not provide an end-to-end, integrated pipeline.

---

<sup>5</sup> Statistics on the Linked Open Data cloud: <https://lod-cloud.net/#about>

In this paper, we present the methods, tools and services implemented in **Q<sup>2</sup>Forge**, a web application guiding the user through the steps of a generic, extensible, end-to-end pipeline to generate a reference Q<sup>2</sup>set i.e. a dataset of (*NL question*, *SPARQL query*) pairs tailored to a specific KG. Through an interactive and iterative process, the user interface assists the user in three main areas: (1) producing CQs based on information about a KG and the domain it pertains to; (2) proposing **SPARQL query counterparts of the CQs**, given the KG and its schemata; (3) **testing** the proposed SPARQL queries, **judging** the relevance of question-query pairs, and **recommend refinements**. Rather than constraining users to a fixed end-to-end pipeline, Q<sup>2</sup>Forge emphasizes flexibility: these tasks are executed as an integrated pipeline, but a user may also choose to use one task independently of the others.

Q<sup>2</sup>Forge relies on an extensive, user-controlled configuration where, in particular, multiple language models can be selectively used at different steps of the pipeline. Through a documented Web API, Q<sup>2</sup>Forge leverages a set of pre-defined services, e.g. to explore the KG or invoke a language model for a certain task, implemented using robust, community-proven libraries and frameworks such as LangChain.<sup>6</sup> Yet, a community may easily extend Q<sup>2</sup>Forge with new services and steps, or re-implement some of the provided services for instance to use their own text-to-SPARQL tool instead of the one provided.

This paper is structured as follows: Section 2 provides an overview and comparison with relevant existing works. Section 3 describes our methodology, the pipeline architecture and the various components. Then, the source code and its sustainability plan are described in Section 4. Section 5 discusses real practical use cases where the resource could be used, while its potential impact and reusability are discussed in Section 6. Finally, the limitations and perspectives of the resource are outlined in Section 7.

## 2 Related Work

*Linked Data Query Assistants.* Approaches for assisting users in querying KGs can be broadly separated in two main non-disjoint categories: the ones relying on dedicated Graphical User Interfaces (GUI) (e.g. [9,10,2]) and the ones relying on Natural Language Interfaces (NLI) (e.g. [13,17]). GUIs can provide high expressivity but remain difficult to use by non-technical experts, unless they trade off part of the expressivity in favor of reusing a popular interaction paradigm (e.g. faceted search). NLIs range from keyword-based retrieval to controlled natural language and full natural language dialogical interactions. Language models, large (LLM) and small (SLM), have significantly improved the methods for natural language processing in general, and in particular for question-answering over linked data. Language models are used in several ways: translate a question to a structured query (directly or indirectly [13]) or by directly answering the question when, for instance, the knowledge source was included in the training

<sup>6</sup> LangChain homepage: <https://www.langchain.com/>

corpus of the language model. These two trends can also be combined with augmentation techniques such as Retrieval Augmented Generation (RAG) [14] that performs information retrieval tasks of different natures (document, database, KG) to enrich the context used to invoke the language model and improve the quality of the answers. While some GUIs help lower the barrier for non-expert users, NLI approaches, and particularly those using LLMs, are even more user-friendly and extensible but have shown mixed results in generating accurate SPARQL queries from natural language, and they require reference Q<sup>2</sup>sets to be trained, augmented and evaluated. This is precisely the purpose of Q<sup>2</sup>Forge.

*Linked Data Question-Query sets.* Multiple tools and services are currently being released to address the question of text-to-SPARQL translation. BigCQ [20,21] aims to create CQs and SPARQL equivalents based for the axioms of a specific OWL ontology. It is meant to help ontology engineering and evaluation, but cannot apply to common KGs that typically rely simultaneously on multiple ontologies and vocabularies. Amazon Bedrock<sup>7</sup> is a commercial, text-to-SPARQL translation service proposed by Amazon, that leverages LLMs and requires a collection of few-shot question-query pairs. It has applications particularly in bioinformatics. Unfortunately it is not released under an open source license, thus making comparison difficult. AllegroGraph’s Natural Language Query (NLQ)<sup>8</sup> vector Database stores pairs of NL questions and corresponding SPARQL queries. This repository helps to train and refine models for accurate query generation, and its integration with SHACL shapes ensures the structural validity of the generated queries. However, this solution suffers from a lack of explainability. Users receive SPARQL query counterparts without understanding why a particular result was returned or the full process used to infer that outcome. In the opposite direction, AutoQGS [22] is a framework that generates NL questions from SPARQL queries, facilitating the creation of question-query datasets without extensive manual annotation. However, this solution requires existing SPARQL queries to generate training data, which limits its applicability.

Challenges such as QALD [18], DBNQA [11], and LC-QuAD [7] provide Q<sup>2</sup>sets to train models that generate queries from NL questions, but they focus primarily on DBpedia and Wikidatan, although some editions of QALD, e.g. QALD 4<sup>9</sup>, have included biomedical Q<sup>2</sup>sets. Similarly, the LC-QuAD 2.0 dataset<sup>10</sup> contains a Q<sup>2</sup>set of more than 20,000 pairs across DBpedia and Wikidata, including subdomains such as geography and science. While these resources serve general purpose question answering systems, they do not comprehensively capture domain specific or highly specialised knowledge. Some domain specific datasets exist, such as SIB bioinformatics SPARQL queries, [3], a collection of hand-crafted Q<sup>2</sup>sets for various SIB-related KGs. By contrast, Q<sup>2</sup>Forge aims to fill this gap by providing an open-source solution to generate Q<sup>2</sup>sets for any domain and KG, including for private KGs.

<sup>7</sup> <https://tinyurl.com/z4wrxvb3>

<sup>8</sup> <https://franz.com/agraph/support/documentation/nl-query.html>

<sup>9</sup> <https://github.com/ag-sc/QALD/tree/master/4/data>

<sup>10</sup> <https://github.com/AskNowQA/LC-QuAD2.0/>

### 3 From a Knowledge Graph to Q<sup>2</sup>Forge Pipeline

Q<sup>2</sup>Forge helps users to carry out three main tasks: generate competency questions in NL for a target KG, generate SPARQL query translations of the questions, and test and refine the SPARQL queries. To do so, Q<sup>2</sup>Forge orchestrates the use of various services to manage multiple per-KG configurations, extract the schema of a KG, invoke various language models depending on the task to achieve at each step of the pipeline, etc. The services are invoked through a documented Web API implemented by a back-end server. We provide a prototype implementation of the back-end server called Gen<sup>2</sup>KGBot. A community may reuse Gen<sup>2</sup>KGBot as-is, or they may customize or extend its services to meet their specific needs.

Figure 1 describes the pipeline of Q<sup>2</sup>Forge: (1) create the configuration for a KG and (2) extract its schema; (3) generate CQs and (4) optionally export them for reuse with another application or for documenting purpose; (5) translate a CQ into SPARQL; (6) execute the query and propose an interpretation of the results; (7) judge the relevance of the question-query pair and allow the user to iteratively refine the query; (8) export the Q<sup>2</sup>set for reuse with other systems. Note that Q<sup>2</sup>Forge remains very flexible: a user may follow the whole pipeline, but may also run each task independently by simply importing/pasting input data and exporting/copying the outputs.

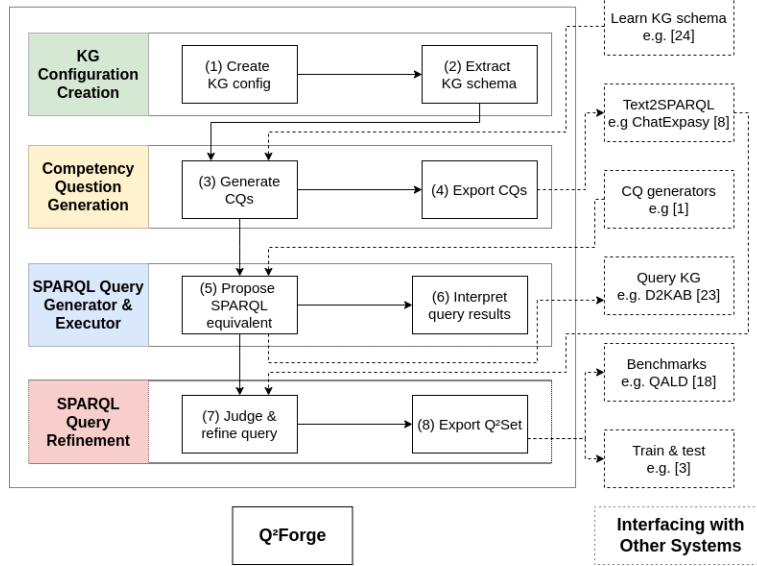


Fig. 1: Q<sup>2</sup>Forge pipeline: resources and services.

The rest of this section further describes the steps depicted in Figure 1.

### 3.1 KG Configuration and Pre-processing

**Create a KG configuration.** The pipeline starts with creating a KG configuration (depicted in Figure 2) where the user provides minimal information about the target KG: a name, a short name used later as an identifier, a textual description, a SPARQL endpoint URL, and the namespaces and prefixes to be used in the SPARQL queries and Turtle descriptions. Optionally, the user may fill in the URL of a SPARQL endpoint hosting the ontologies in case they are not on the same endpoint as the KG itself.

Once created, the configuration is stored on the back-end server. Additional parameters that can be edited manually to configure the available language models (seq-to-seq and embedding), where they are hosted (e.g. local vs. cloud resources, vector database etc.), and how they are assigned to each step of the pipeline. For instance, one may choose to use a large model with reasoning capabilities for generating a SPARQL query, but use a smaller model to interpret SPARQL results. Other parameters configure the strategy adopted to serialize ontology classes in the prompts submitted to seq-to-seq models, such as the number of ontology classes to describe and the linearization format used to describe them. Multiple formats are supported (currently Turtle, tuples or a NL format, see examples in Listing 1.1), since different language models may behave differently depending on the selected format.

**KG pre-processing.** We extract from the KG various types of information that will be helpful to carry out the downstream text-to-SPARQL task. This can be ontology classes that are relevant with respect to a NL question, example SPARQL queries, etc. In our implementation, this step first creates a textual description of the classes from the labels and descriptions available in the ontologies, and computes text embeddings thereof. In Figure 2, this is achieved in steps 2 and 3. Furthermore, there is usually a gap between how an ontology defines classes and how instances of these classes are concretely represented in the KG. Typically, instances may use properties and resources from additional vocabularies that are not explicitly mentioned in the ontology. Therefore, the text-to-SPARQL task requires not only a textual description of the classes, but also a description of how instances of these classes are represented. Gen<sup>2</sup>KGBot addresses this need by sampling class instances and analyzing the properties and value types they use (examples are provided in Listing 1.1). Lastly, the user may provide existing examples of NL question and associated SPARQL query. The pre-processing includes computing embeddings of these question-query pairs.

### 3.2 Competency Question Generation

This step invokes a language model to generate CQs based on various types of information about the KG: name and description, endpoint URL, list of the used ontologies. This information is either taken from the KG configuration (created in the previous step) or manually entered in a form. The user may also provide any other relevant information, e.g. the abstract of an article describing the KG.

```

@prefix obo: <http://purl.obolibrary.org/obo/> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
obo:R0_0000056 rdfs:label "participates_in" .
[ a obo:CHEBI_53289 ;
  obo:R0_0000056 [ a pubchem:MeasureGroup ] ;
...

('obo:CHEBI_53289', 'http://purl.obolibrary.org/obo/R0_0000056', 'participates_in', 'http
://rdf.ncbi.nlm.nih.gov/pubchem/vocabulary#MeasureGroup')
...

Instances of class 'obo:CHEBI_53289' have property 'obo:R0_0000056' (participates_in) with
value type 'pubchem:MeasureGroup'.
...

```

Multiple Listing 1.1: Formats to describe properties and value types used by instances of a class: Turtle (top), tuple (middle), natural language (bottom)

The screenshot shows the 'Create a new Configuration' step of the Q2Forge interface. It includes the following fields and content:

- KG full name\***: WheatGenomic Scientific Literature Knowledge Graph
- KG short Name\***: d2kab
- KG description\***: The Wheat Genomics Scientific Literature Knowledge Graph (WheatGenomicsSLKG) is a FAIR knowledge graph that exploits the Semantic Web technologies to describe PubMed scientific articles on wheat genetics and genomics. It represents Named Entities (NE) about genes, phenotypes, taxa and varieties, mentioned in the title and the abstract
- KG SPARQL endpoint URL\***: http://d2kab.i3s.unice.fr/sparql
- Ontologies SPARQL endpoint URL**: http://d2kab.i3s.unice.fr/sparql
- Properties QNames info**: (empty field)
- Prefixes**:
 

prefix*	value*
bibo	http://purl.org/ontology/bibo/

Fig. 2: The KG configuration and pre-processing interface.

Figure 3 depicts the competency question generation interface. The user can select the language model to be used for the generation of the CQs and the number of CQs to be generated. The model is instructed to return each question with an evaluation of its complexity (Basic, Intermediate or Advanced) and a set of tags. The *Enforce Structured Output* toggle can be used to compel the model to return the CQs as a JSON-formatted document.

Upon completion of the process, the user may download the output as a JSON document, and save it in a browser's cookie for reuse in the next step.





Fig. 3: The competency question generation interface.

### 3.3 SPARQL Query Generator/Executor

In this step, the user has the ability to generate a SPARQL query counterpart of a NL question, execute it against a KG, and get an interpretation of the results. The question may originate from the preceding task, or the user may paste a question either hand-crafted or generated by another system.

Q<sup>2</sup>Forge relies on various strategies provided by Gen<sup>2</sup>KGBot to accomplish this task, which we refer to as “scenarios”. In the following, we focus on Scenario 5, depicted in Figure 4, that we will further describe below. When running a scenario, the steps of that scenario are progressively rendered on the interface, and for the ones that make an LLM call, the response is dynamically streamed to ensure a good user experience. Figure 5 is a snapshot of the interface of the SPARQL Query Generator/Executor. The steps are as follows:

1. **Initial question:** the workflow is initiated by the user posing a NL question.
2. **Question validation:** the question is evaluated to ensure its relevance to the context of the KG. If it is deemed invalid, the workflow stops.
3. **Question pre-processing:** common techniques are used to extract named entities (NEs) from the question.
4. **Select similar classes:** similarity search between the question and the ontology class descriptions computed in the KG pre-processing step is used to select relevant classes.
5. **Get context information about the classes:** retrieve a description of the properties and value types used with instances of the selected classes.
6. **Generate query:** generate a prompt from a template<sup>11</sup> using the KG configuration and the inputs from the previous steps, and submit it to the configured LLM.

<sup>11</sup> [https://github.com/Wimmics/gen2kgbot/blob/master/app/scenarios/scenario\\_5/prompt.py#L3](https://github.com/Wimmics/gen2kgbot/blob/master/app/scenarios/scenario_5/prompt.py#L3)

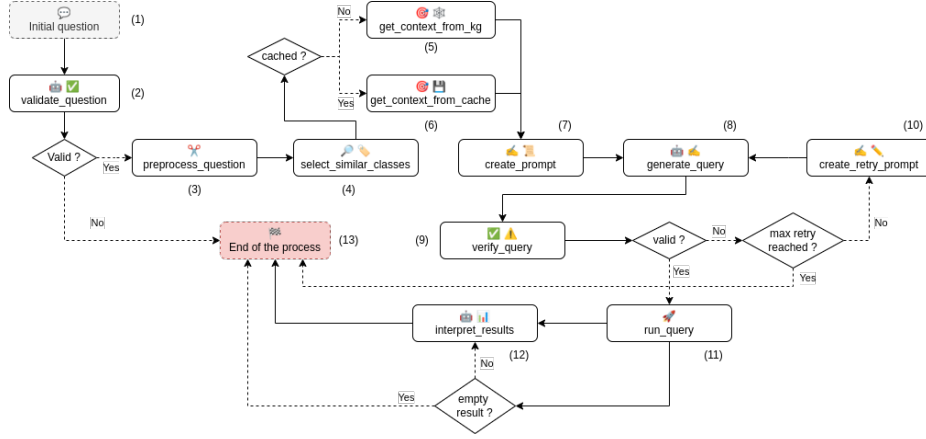


Fig. 4: Gen<sup>2</sup>KGBot SPARQL Query Generator/Executor: workflow of Scenario 5

7. **Verify query and retry:** check if a SPARQL query was generated and if it is syntactically correct. If not, generate a retry prompt that includes the last generated answer and the reason for the retry, e.g. syntax errors, and submit this retry prompt to the configured LLM.
8. **Execute the SPARQL query:** if a valid SPARQL query was generated, submit it to the KG endpoint and get the results.
9. Use the configured LLM to **interpret the SPARQL results**.

Scenario 5, described above, is useful as a starting point when no prior question-query pair exists. However, once some pairs have been validated or if some pairs were hand-crafted, they can be added to the context and serve as examples. Scenario 6 can then be applied instead, as it provides the model with relevant example SPARQL queries that can help in generating more accurate queries with fewer refinement iterations.

### 3.4 SPARQL Query Refinement

In this step, the user can incrementally refine a SPARQL query so that it reflects precisely the question. Figure 6 is a snapshot of the interface, and the process is as follows:

1. First, the query is displayed in a SPARQL editor that highlights potential syntactic errors and can be used to submit the query to the endpoint.
2. To help the user understand the query, Q<sup>2</sup>Forge can extract the qualified (prefixed) names (QNs) and fully qualified names (FQNs) from the query and get their labels and descriptions. For instance, the label of [http://purl.obolibrary.org/obo/CHEBI\\_53289](http://purl.obolibrary.org/obo/CHEBI_53289) is “donepezil”, and its description is “a racemate comprising equimolar amounts of (R)- and (S)-donepezil (...)”.




Fig. 5: The SPARQL Query Generator/Executor interface.

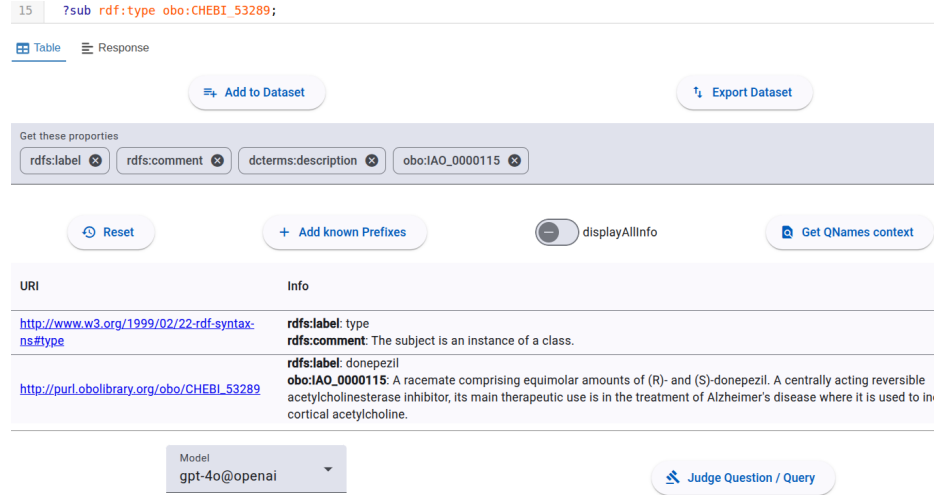


Fig. 6: The query refinement interface.

- Then the LLM is asked to judge whether the query matches the given question. It is requested to provide a grade between 0 and 10 along with explanations justifying the grade.

The user may then iterate as needed: amend the query based on the grade and insights from the model, test it, have the model judge it, etc. Once a satisfying query is reached, the user can add the question-query pair to a dataset and export it in a variety of formats, catering to different use cases.

## 4 Source Code and Documentation

From a technical perspective, Q<sup>2</sup>Forge integrates several software components that are robust and have been proven effective in various contexts. LangChain and LangGraph<sup>12</sup> are used for LLM workflow orchestration, Spacy<sup>13</sup> for question pre-processing and named entity recognition, rdflib<sup>14</sup> for the manipulation of RDF data, and YasGUI<sup>15</sup> as a SPARQL editor.

**Source Code Availability.** Q<sup>2</sup>Forge and accompanying Gen<sup>2</sup>KGBot are provided under the GNU Affero General Public License v3.0 or later (AGPL-3.0-or-later) license. The code is published on public Github repositories, and the versions used at the time of writing are identified by DOIs to ensure the long-term preservation and citability.<sup>16</sup> The API provided by Gen<sup>2</sup>KGBot is documented according to the OpenAPI format.<sup>17</sup>

**Sustainability Plan.** Over the next four years, financial support has been secured through the MetaboLinkAI project.<sup>18</sup> This project aims to transform metabolomics data into actionable insights through the utilisation of AI-powered, knowledge graph-driven solutions. This will provide an opportunity to evaluate the quality, relevance and applicability of Q<sup>2</sup>Forge in the chemistry domain. Moreover, a fundamental objective of Q<sup>2</sup>Forge is to provide a generic solution, reusable with a variety of KGs. Consequently, we intend to provide support to communities expressing interest and willingness to experiment with it for their own needs. This support may range from best-effort to more formalized collaboration. To support collaborations, adoptions and contributions we secured two other contributors: the P16 public program<sup>19</sup> that helps open-source project improve their code and diffusion, and the Probabl company<sup>20</sup> whose mission is to develop, maintain the state-of-the-art, sustain, and disseminate a complete suite of open source tools for data science.

## 5 Application Use Cases

In this section, we present several application use cases where Q<sup>2</sup>Forge can be of great help. First, we focus on the creation of Q<sup>2</sup>sets and the benefits of generating them instead of performing this task manually. Second, we examine how Q<sup>2</sup>Forge can document existing KGs with multiple competency questions. Finally, we explore its application in creating a golden dataset for benchmarking, training and testing question answering models.

<sup>12</sup> <https://www.langchain.com/>

<sup>13</sup> <https://spacy.io/>

<sup>14</sup> <https://github.com/RDFLib>

<sup>15</sup> <https://yasgui.triply.cc/>

<sup>16</sup> All links are given at the beginning of this article.

<sup>17</sup> <http://localhost:8000/q2forge/api/docs/>

<sup>18</sup> <http://www.metabolinkai.net/>

<sup>19</sup> <https://p16.inria.fr/en/>

<sup>20</sup> <https://probabl.ai/about>

**Lowering the entry barrier to query rich KGs:** Large public KGs usually provide lay users with user-friendly user interfaces that propose pre-defined queries and exploration options. Yet such interfaces can hardly accommodate complex custom queries where SPARQL expertise is necessary. For instance, with over 111 million chemical substances and extensive bio-activity data, PubChem presents significant navigation challenges for researchers. Metabolomics experts often struggle to formulate complex SPARQL queries that would help them identify relationships between compounds, biological activities, and disease associations. Q<sup>2</sup>Forge addressed this challenge by allowing chemists to generate natural language questions and automatically convert them to SPARQL queries, thus drastically reducing the time spent on data retrieval. To do so, researchers must provide a textual description of the KG together with additional relevant textual information, such as the abstracts of articles published about PubChem or about research made possible through PubChem. For example, a researcher might ask: “Which compounds have been tested against SARS-CoV-2 Main Protease and reported an IC50 below 1  $\mu$ M?” or “Which natural product compounds from marine sponges show antimicrobial activities against *Pseudomonas aeruginosa*?”. Similarly, environmental health researchers studying the exposome face difficulties extracting meaningful correlations between environmental factors and metabolic responses across heterogeneous datasets. Currently, they rely either on simple predefined queries or must collaborate with knowledge engineering specialists, creating bottlenecks in research workflows. Q<sup>2</sup>Forge could enable them to independently generate appropriate question-query pairs that bridge environmental exposures and biological outcomes, eliminating technical barriers to knowledge discovery. A typical researcher might need to ask: “Which air pollutants are known to increase *Nrf2* anti-oxydant protein expression ?” or “What metabolic biomarkers show significant alterations following chronic exposure to per- and polyfluoroalkyl substances (PFAS) in human biomonitoring studies?”

**Ground truth and question-query benchmarks:** Our community has pioneered challenges and benchmarks for question-answering over linked data [18,11,7]. However, each edition of these challenges requires updating Q<sup>2</sup>sets for tasks that were proposed in previous editions, and creating new Q<sup>2</sup>sets for newly proposed tasks. There does not exist a large number of such readily available Q<sup>2</sup>sets, and they are often based on the same KGs (e.g. DBpedia, Wikidata). Setting up a new edition of a challenge therefore requires a significant effort to generate or update the training and test data. Q<sup>2</sup>Forge was designed to help produce these Q<sup>2</sup>sets and can be used to facilitate the renewal of tasks for challenges and benchmarks on a variety of KGs. For instance, the QALD Challenge [18] has long been centered on DBpedia. In the latest edition, it was extended to Wikidata. Using Q<sup>2</sup>Forge, we could extend the challenge with tasks targeting domain-specific graphs such as Uniprot [6] or the aforementioned PubChem graph.

**Documenting a KG with competency questions:** CQs are commonly used to demonstrate the basic capabilities of a KG. This requires working with do-

main experts to identify the CQs they may want to ask. In our experience, this is a time-consuming task involving multiple iterations. Q<sup>2</sup>Forge can be used to initialize, expand and enhance the scope and variety of CQs by systematically generating hundreds of competency questions. For instance, PubChem’s documentation currently provides a valuable foundation of 16 CQs.<sup>21</sup> Q<sup>2</sup>Forge could enhance this foundation by showcasing the full scope and complexity of chemical, biological, and pharmacological relationships within this extensive KG. The generated question sets can serve multiple purposes: providing entry points to new users, supporting KGs indexing, benchmarking search capabilities, identifying promising research directions, and accelerating the development of next-generation retrieval systems.

## 6 Potential Impact and Reusability

**Target Audiences and Expected Uses.** We have already identified three families of users, that correspond to the three uses cases described in Section 5: (1) the developers and maintainers of question answering systems, chatbots, conversational agents and other natural language search engines over KGs. The methods behind these systems all require to have Q<sup>2</sup>sets to train, test and evaluate the system. They are the primary target of Q<sup>2</sup>Forge. (2) Events and groups organizing challenges, benchmarking existing solutions and building surveys. These are in constant need of new and renewed Q<sup>2</sup>sets to compare the latest methods and establish the state-of-the-art. Here, Q<sup>2</sup>Forge facilitates the creation of Q<sup>2</sup>sets from any KG in any domain. (3) While it is strongly recommended to document existing datasets and query services with examples of typical questions and queries, this is rarely done and, when it is, rarely extensive. Q<sup>2</sup>Forge was designed to help generate these examples with quality and quantity in mind.

We have initiated experiments with the first family of users in the chemistry and metabolomics domain. Pharmaceutical researchers developing drug discovery platforms require comprehensive question-query pairs to train intelligent systems for identifying promising molecular candidates across multiple parameters. These researchers benefit from Q<sup>2</sup>Forge’s ability to generate diverse questions exploring structure-activity relationships, pharmacokinetic properties, and target binding profiles. Metabolomics data scientists integrating multi-omics datasets need sophisticated query templates that traverse complex biochemical pathway knowledge, particularly when correlating mass spectrometry findings with biological outcomes. Based on this experimentation, we believe that academic laboratories focusing on cheminformatics and bioinformatics can utilize Q<sup>2</sup>Forge to develop educational materials demonstrating how semantic queries extract meaningful insights from chemical databases. Q<sup>2</sup>Forge can significantly reduce technical barriers that have historically prevented domain experts from fully leveraging KG technologies in their specialized fields. For the third family of users, we have experimented Q<sup>2</sup>Forge with outputs of the D2KAB project. D2KAB produced several datasets among which the *Wheat Genomics Scientific*

<sup>21</sup> <https://pubchem.ncbi.nlm.nih.gov/docs/rdf-use-cases>

*Literature Knowledge Graph* [23] that represents the named entities extracted from a corpus of over 8,000 PubMed articles related to wheat genetics and genomics. The NEs include genes, phenotypes, taxon names and varieties in titles and abstracts. During the project, we worked with domain experts to figure out several CQs<sup>22</sup> to document the graph and illustrate its usefulness. When tested with this KG, Q<sup>2</sup>Forge was able to automatically generate relevant CQs and translate them into SPARQL queries that were close to the target. After a short refinement step, we managed to get valid question-query pairs. With these two experiments, we are confident that Q<sup>2</sup>Forge is flexible enough to be applied to a broad range of KGs and domains.

**Potential for Reuse and Extension Points.** As mentioned earlier, Q<sup>2</sup>Forge strives for flexibility: the various tasks of the pipeline can be executed as a whole, and some tasks can be used independently of the others. In addition to the interfacing with other systems shown in Figure 1, the system is designed to be reused and integrated into other scenarios. In education for instance, when teaching SPARQL, Q<sup>2</sup>Forge could be modified to serve as a tailored instructor guiding learners to navigate the complexities of SPARQL. Furthermore, since the query refinement task can be accessed independently of the other tasks (its URL takes arguments “question” and “query”), adding the appropriate button to an existing SPARQL editor could seamlessly integrate this task into an existing workflow.

Current development of protocols such as MCP (Model Context Protocol),<sup>23</sup> A2A (Agent-to-Agent)<sup>24</sup> and hMAS (Hypermedia Multi-Agent Systems)<sup>25</sup> reflects ongoing efforts to simplify integration, enhance collaboration, and ensure secure and efficient communication between AI agents and external systems. These protocols can potentially be interfaced with Q<sup>2</sup>Forge to facilitate its incorporation into broader systems. Integrating MCP would standardize Q<sup>2</sup>Forge’s interface with LLMs, enabling seamless integration of (components of) Q<sup>2</sup>Forge’s pipeline into other workflows. Conversely, Q<sup>2</sup>Forge could be extended to support the invocation of MCP servers providing access to third-party services such as knowledge graphs. Additionally, incorporating A2A would allow Q<sup>2</sup>Forge to support multi-agent collaboration across diverse ecosystems, fostering coordination between agents of varying frameworks. Finally, aligning with hMAS would leverage semantic hypermedia for uniform interactions among people, devices, and digital services, creating hybrid AI communities that operate transparently and accountably on the Web. These extensions would make Q<sup>2</sup>Forge even more versatile, facilitating the development of KG applications in different domains.

<sup>22</sup> <https://github.com/Wimmics/WheatGenomicsSLKG/blob/main/SPARQLQueries-JupyterNotebook.ipynb>

<sup>23</sup> <https://modelcontextprotocol.io/>

<sup>24</sup> <https://google.github.io/A2A/>

<sup>25</sup> <https://project.hyperagents.org/>

## 7 Conclusion and Perspectives

The present article has highlighted the challenge of creating Q<sup>2</sup>sets (datasets of question-query pairs) for a KG as part of the ever-growing LOD cloud. To address this challenge, concrete methods and tools have been presented. Utilising robust, industry-proven tools, Q<sup>2</sup>Forge’s pipeline is designed to address the generation of competency questions (CQ) in NL, translate the CQs into SPARQL queries, and help users to refine those queries, and export high-quality QALD-like Q<sup>2</sup>sets that can be used for benchmarking, training and evaluating text-to-SPARQL models. This innovative end-to-end pipeline incorporates a variety of dedicated services to achieve these objectives, and emphasizes genericity (it can apply to any KG in any domain) and flexibility (the pipeline can easily be modified and extended to account for new needs).

Future work will proceed along several directions in the short and medium term. First, we will focus on data quality evaluation. This includes establishing automated validation processes and human evaluation protocols to ensure the relevance and accuracy of generated question-query pairs. Second, we envision integrating existing KG construction and extension tools from unstructured content. Such integration would enable Q<sup>2</sup>Forge to generate questions about both structured and unstructured data sources, broadening its applicability across diverse data landscapes. Additionally, we plan to implement the model context protocol (MCP) to provide Q<sup>2</sup>Forge’s three primary functions as reusable tools to other workflows. For instance, this would enable systems that need to transform NL questions into SPARQL query to seamlessly integrate Q<sup>2</sup>Forge capabilities. Three limitations of our work that we wish to address are: (1) Follow-up questions: our plan is to achieve this using short/long memory and conversation summarizing to cope with context explosion; (2) custom and tailor result visualization when a textual rendering of a SPARQL query result is not the optimum way of conveying the results; (3) Federated queries: the current implementation is limited to single SPARQL endpoints and does not support federated queries across multiple KGs, which restricts its applicability in scenarios requiring data integration from diverse sources. This limitation could be addressed by extending Q<sup>2</sup>Forge architecture to support the generation of queries that span multiple endpoints. Finally, extensive dissemination and open-source support activities are planned so that other communities can adopt this work and adapt it to their specific needs. Through community engagement, documentation, and targeted outreach, we aim to foster broader adoption and enable customization of Q<sup>2</sup>Forge across various domains and research contexts.

**Acknowledgments.** This work was supported by the French government through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence Université Côte d’Azur (ANR-15-IDEX-01). Additional support came from French Government’s France 2030 investment plan (ANR-22-CPJ2-0048-01), through 3IA Cote d’Azur (ANR-23-IACL-0001) as well as the MetaboLinkAI bilateral project (ANR-24-CE93-0012-01 and SNSF 10002786).



## References

1. Alharbi, R., Tamma, V., Grasso, F., Payne, T.R.: The Role of Generative AI in Competency Question Retrofitting. In: The Semantic Web: ESWC 2024 Satellite Events: Hersonissos, Crete, Greece, May 26–30, 2024, Proceedings, Part I. pp. 3–13. Springer-Verlag (Jan 2025). [https://doi.org/10.1007/978-3-031-78952-6\\_1](https://doi.org/10.1007/978-3-031-78952-6_1)
2. Bast, H., Kalmbach, J., Klumpp, T., Kramer, F., Schnelle, N.: Efficient and Effective SPARQL Autocompletion on Very Large Knowledge Graphs. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. p. 2893–2902. CIKM ’22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3511808.3557093>, <https://doi.org/10.1145/3511808.3557093>
3. Bolleman, J., Emonet, V., Altenhoff, A., Bairoch, A., Blatter, M.C., Bridge, A., Duvaud, S., Gasteiger, E., Kuznetsov, D., Moretti, S., Michel, P.A., Morgat, A., Pagni, M., Redaschi, N., Zahn-Zabal, M., de Farias, T.M., Sima, A.C.: A large collection of bioinformatics question-query pairs over federated knowledge graphs: methodology and applications (2024). <https://doi.org/10.48550/arXiv.2410.06010>, <https://arxiv.org/abs/2410.06010>
4. Ciroku, F., de Berardinis, J., Kim, J., Meroño-Peñuela, A., Presutti, V., Simperl, E.: RevOnt: Reverse Engineering of Competency Questions from Knowledge Graphs via Language Models. *Web Semant.* **82**(C) (Oct 2024). <https://doi.org/10.1016/j.websem.2024.100822>
5. Cohen, K.B., Kim, J.D.: Evaluation of SPARQL Query Generation from Natural Language Questions. Proceedings of the Joint Workshop on NLP&LOD and SWAIE: Semantic Web, Linked Open Data and Information Extraction **2013**, 3–7 (Sep 2013)
6. Consortium, T.U.: Uniprot: a worldwide hub of protein knowledge. *Nucleic Acids Research* **47**(D1), D506–D515 (11 2018). <https://doi.org/10.1093/nar/gky1049>, <https://doi.org/10.1093/nar/gky1049>
7. Dubey, M., Banerjee, D., Abdelkawi, A., Lehmann, J.: LC-QuAD 2.0: A Large Dataset for Complex Question Answering over Wikidata and DBpedia. In: Ghidini, C., Hartig, O., Maleshkova, M., Svátek, V., Cruz, I., Hogan, A., Song, J., Lefrançois, M., Gandon, F. (eds.) *The Semantic Web – ISWC 2019*. pp. 69–78. Springer International Publishing, Cham (2019)
8. Emonet, V., Bolleman, J., Duvaud, S., de Farias, T.M., Sima, A.C.: LLM-based SPARQL Query Generation from Natural Language over Federated Knowledge Graphs (2024), <https://arxiv.org/abs/2410.06062>
9. Ferré, S.: SPARKLIS: An Expressive Query Builder for SPARQL Endpoints with Guidance in Natural Language. *Open Journal Of Semantic Web* **0** (2017), <https://inria.hal.science/hal-01485093>
10. Francart, T.: Sparnatural: A Visual Knowledge Graph Exploration Tool. In: Pesquita, C., Skaf-Molli, H., Efthymiou, V., Kirrane, S., Ngonga, A., Collarana, D., Cerqueira, R., Alam, M., Trojahn, C., Hertling, S. (eds.) *The Semantic Web: ESWC 2023 Satellite Events*. pp. 11–15. Springer Nature Switzerland, Cham (2023)
11. Hartmann, A.K., Soru, T., Marx, E.: Generating a Large Dataset for Neural Question Answering over the DBpedia Knowledge Base. In: *Workshop on Linked Data Management, co-located with the W3C WEBBR 2018* (2018)
12. Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., Melo, G.D., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al.: Knowledge graphs. *ACM Computing Surveys (Csur)* **54**(4), 1–37 (2021)

13. Lehmann, J., Gattogi, P., Bhandiwad, D., Ferré, S., Vahdati, S.: Language Models as Controlled Natural Language Semantic Parsers for Knowledge Graph Question Answering. In: *Frontiers in Artificial Intelligence and Applications*. Frontiers in Artificial Intelligence and Applications, vol. 372, pp. 1348–1356. IOS Press, Krakow (Cracovie), Poland (Sep 2023). <https://doi.org/10.3233/FAIA230411>, <https://inria.hal.science/hal-04269089>
14. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., Riedel, S., Kiela, D.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 9459–9474. Curran Associates, Inc. (2020)
15. Pan, X., van Ossenbruggen, J., de Boer, V., Huang, Z.: A RAG Approach for Generating Competency Questions in Ontology Engineering. In: Sfakakis, M., Garoufallou, E., Damigos, M., Salaba, A., Papatheodorou, C. (eds.) *Meta-data and Semantic Research*. pp. 70–81. Springer Nature Switzerland (2025). [https://doi.org/10.1007/978-3-031-81974-2\\_6](https://doi.org/10.1007/978-3-031-81974-2_6)
16. Rebboud, Y., Tailhardat, L., Lisena, P., Troncy, R.: Can LLMs Generate Competency Questions? In: Meroño Peñuela, A., Corcho, O., Groth, P., Simperl, E., Tamma, V., Nuzzolese, A.G., Poveda-Villalón, M., Sabou, M., Presutti, V., Celino, I., Revenko, A., Raad, J., Sartini, B., Lisena, P. (eds.) *The Semantic Web: ESWC 2024 Satellite Events*. pp. 71–80. Springer Nature Switzerland (2025). [https://doi.org/10.1007/978-3-031-78952-6\\_7](https://doi.org/10.1007/978-3-031-78952-6_7)
17. Tysinger, E., Pagni, M., Kirchhoffer, O., Mehl, F., Gandon, F., Wolfender, J.L., Nothias, L.F.: An Artificial Intelligence Agent for Navigating Knowledge Graph Experimental Metabolomics Data. In: *Frontiers in Metabolomics - Book of abstract - Frontiers in Metabolomics, Swiss Metabolomics Society Zurich and ETH Zurich, Zurich, Switzerland (Sep 2023)*, <https://hal.science/hal-04381448>
18. Usbeck, R., Yan, X., Perevalov, A., Jiang, L., Schulz, J., Kraft, A., Möller, C., Huang, J., Reineke, J., Ngomo, A.C.N., Saleem, M., Both, A.: QALD-10 – The 10th challenge on question answering over linked data: Shifting from DBpedia to Wikidata as a KG for KGQA. *Semantic Web* **15**(6), 2193–2207 (2024). <https://doi.org/10.3233/SW-233471>, <https://journals.sagepub.com/doi/abs/10.3233/SW-233471>
19. Warren, P., Mulholland, P.: A Comparison of the Cognitive Difficulties Posed by SPARQL Query Constructs. In: Keet, C.M., Dumontier, M. (eds.) *Knowledge Engineering and Knowledge Management*. pp. 3–19. Springer International Publishing (2020). [https://doi.org/10.1007/978-3-030-61244-3\\_1](https://doi.org/10.1007/978-3-030-61244-3_1)
20. Wiśniewski, D., Potoniec, J., Ławrynowicz, A.: BigCQ: Generating a Synthetic Set of Competency Questions Formalized into SPARQL-OWL (Student Abstract). In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 13079–13080 (2022)
21. Wiśniewski, D., Potoniec, J., Ławrynowicz, A.: BigCQ: A large-scale synthetic dataset of competency question patterns formalized into SPARQL-OWL query templates (2021), <https://arxiv.org/abs/2105.09574>
22. Xiong, G., Bao, J., Zhao, W., Wu, Y., He, X.: AutoQGS: Auto-Prompt for Low-Resource Knowledge-based Question Generation from SPARQL. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. p. 2250–2259. CIKM '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3511808.3557246>, <https://doi.org/10.1145/3511808.3557246>

23. Yacoubi Ayadi, N., Bernard, S., Bossy, R., Courtin, M., Happi Happi, B.G., Larmande, P., Michel, F., Nedellec, C., Roussey, C., Faron, C.: A Unified approach to publish semantic annotations of agricultural documents as knowledge graphs. *Smart Agricultural Technology* **8**, 43 (Jan 2024). <https://doi.org/10.1016/j.atech.2024.100484>, <https://hal.science/hal-04495022>
24. Zouaq, A., Martel, F.: What Is the Schema of Your Knowledge Graph? Leveraging Knowledge Graph Embeddings and Clustering for Expressive Taxonomy Learning. In: *Proceedings of The International Workshop on Semantic Big Data*. pp. 1–6. SBD '20, Association for Computing Machinery (Jun 2020). <https://doi.org/10.1145/3391274.3393637>