

# NLQxform: A Language Model-based Question to SPARQL Transformer

Ruijie Wang<sup>1,2,\*</sup>, Zhiruo Zhang<sup>1</sup>, Luca Rossetto<sup>1</sup>, Florian Ruosch<sup>1</sup> and Abraham Bernstein<sup>1</sup>

<sup>1</sup>Department of Informatics, University of Zurich, Switzerland

<sup>2</sup>University Research Priority Program “Dynamics of Healthy Aging”, University of Zurich, Switzerland

## Abstract

In recent years, scholarly data has grown dramatically in terms of both scale and complexity. It becomes increasingly challenging to retrieve information from scholarly knowledge graphs that include large-scale heterogeneous relationships, such as authorship, affiliation, and citation, between various types of entities, e.g., scholars, papers, and organizations. As part of the Scholarly QALD Challenge, this paper presents a question-answering (QA) system called NLQxform, which provides an easy-to-use natural language interface to facilitate accessing scholarly knowledge graphs. NLQxform allows users to express their complex query intentions in natural language questions. A transformer-based language model, i.e., BART, is employed to translate questions into standard SPARQL queries, which can be evaluated to retrieve the required information. According to the public leaderboard of the Scholarly QALD Challenge at ISWC 2023 (Task 1: DBLP-QUAD – Knowledge Graph Question Answering over DBLP), NLQxform achieved an F1 score of 0.85 and ranked first on the QA task, demonstrating the competitiveness of the system.

## Keywords

Question answering, entity linking, language model, knowledge graph

## 1. Introduction

In the past decade, tremendous efforts [1, 2, 3, 4] have been devoted to organizing scholarly data in the form of knowledge graphs (KGs), which store heterogeneous relational information in an integrated and extensible manner, supporting querying with complex semantics. The DBLP scholarly KG<sup>1</sup> is one representative of such efforts. It stores the entire DBLP<sup>2</sup> data in the format of RDF [5] and provides an endpoint<sup>3</sup> for accessing the data with SPARQL [6] queries. Conventional tools for searching scholarly information, e.g., Google Scholar<sup>4</sup>, DBLP

---

*Scholarly QALD at ISWC 2023*

\*Corresponding author.

✉ ruijie@ifi.uzh.ch (R. Wang); zhiruo.zhang@uzh.ch (Z. Zhang); rossetto@ifi.uzh.ch (L. Rossetto);

ruosch@ifi.uzh.ch (F. Ruosch); bernstein@ifi.uzh.ch (A. Bernstein)

ORCID 0000-0002-0581-6709 (R. Wang); 0009-0008-7115-9429 (Z. Zhang); 0000-0002-5389-9465 (L. Rossetto);

0000-0002-0257-3318 (F. Ruosch); 0000-0002-0128-4602 (A. Bernstein)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

<sup>1</sup><https://blog.dblp.org/2022/03/02/dblp-in-rdf/>

<sup>2</sup><https://dblp.org/>

<sup>3</sup><https://dblp-kg.ltdemos.informatik.uni-hamburg.de/sparql>

<sup>4</sup><https://scholar.google.com/>

bibliography<sup>2</sup> and publisher-provided services<sup>5</sup>, only support text search with limited metadata-based filtering and sorting. In contrast, scholarly KGs are more powerful and versatile, as they support queries with complex constraints (triple patterns in SPARQL) and operations (functions and modifiers in SPARQL). Nevertheless, scholarly KGs are significantly less commonly used than conventional tools. A main hurdle is the complexity of the SPARQL language (cf. [7]).

To address the above issue, we developed a QA system that allows users to query the DBLP scholarly KG with natural language questions. The system is called **NLQxform**, as it answers Natural Language Questions (**NLQs**) using a **transformer**-based language model (LM). Specifically, NLQxform uses BART [8]—a sequence-to-sequence model, which has achieved promising performance in natural language translation. We first fine-tuned it with NLQ-SPARQL query pairs. Given a question, NLQxform employs BART to generate a logical form that has the basic structure of the target SPARQL query. Then, NLQxform links entities mentioned in the given question to their corresponding URLs in the underlying KG and corrects minor syntax and grammar errors in the logical form based on a SPARQL template base. Finally, NLQxform generates explicit candidate SPARQL queries and evaluates them via the DBLP endpoint to retrieve answers.

This work is part of the Scholarly QALD Challenge,<sup>6</sup> which provides valuable benchmarking for QA over scholarly KGs. We participated in both the Question Answering and Entity Linking sub-tasks in Task 1: DBLP-QUAD — Knowledge Graph Question Answering over DBLP.<sup>7</sup> The remainder of this paper is structured as follows: Related work on the two sub-tasks is introduced in Section 2. We elaborate on the NLQxform system in Section 3. The experimental setup and final results of NLQxform in the challenge are reported in Section 4. Finally, we conclude in Section 5.

## 2. Related Work

In this section, we first survey QA approaches with a focus on previous QALD challenges [9, 10, 11, 12, 13, 14, 15]. Then, we present an overview of QA-related entity linking.

**Question Answering.** Banerjee et al. [16] propose to employ pre-trained language models, such as T5 [17] and BART [8], as well as Pointer Generator Networks [18] to construct SPARQL queries using pre-identified entities and relations in given questions. The fine-tuned T5-base model achieved an F1-score of 0.87 on the DBLP-QuAD dataset [19], which is also used for training and validation in this iteration of the Scholarly QALD Challenge. WDAqua [20, 21] participated in several of the past QALD challenges. Diefenbach et al. [20] propose a new approach by not focusing on the syntax but on the semantics of tokens in given questions to generate SPARQL queries. First, recognized tokens are expanded to possible entities in the underlying KG. Next, all possible SPARQL queries are constructed and ranked based on a linear combination of features. Finally, top-ranked queries are selected and used to retrieve answers. The authors integrated their approach to Qanary [22], modular architecture for QA

---

<sup>5</sup>Examples include IEEE Xplore (<https://ieeexplore.ieee.org/Xplore/home.jsp>) and the ACM Digital Library (<https://dl.acm.org/>).

<sup>6</sup><https://kgqa.github.io/scholarly-QALD-challenge/2023/>

<sup>7</sup><https://codalab.lisn.upsaclay.fr/competitions/14264>

systems, and evaluated it on QALD iterations three through seven [23]. Zou et al. [24] recast the problem of NLQ to SPARQL translation with a graph-based solution. They form subgraphs by mapping phrases from given questions to entities in the KG. This relegates the problem of entity disambiguation to the query generation phase of their approach. Finally, the previously created subgraphs are matched to the KG with heuristic rules to retrieve top SPARQL queries and answers.

**Entity Linking.** DBLPLink [25] is a web application specifically developed for entity linking over the DBLP scholarly KG. It employs pre-trained language models to identify entity mentions in given questions and links them to KG entities using labels and entity embeddings. Steinmetz [26] presents an approach where input phrases are parsed to an abstract meaning representation graph based on a pre-trained BART model. From this representation graph, named entities are extracted and subsequently mapped to the KG using various information, including alternative labels, links, and node degrees. The system was evaluated on QALD-9 [27], outperforming competing systems. Diomedi and Hogan [28] propose a novel system called ElNeuKGQA, which combines entity linking with neural machine translation. To alleviate the issue of unseen words, the authors utilize a neural machine translation model to generate template queries wherein out-of-vocabulary words are replaced by placeholders. These are then filled by identified entities afterward in an “entity filling” phase, where sequence labeling is used to determine the role of placeholders. EARL [29] unifies entity linking and relation linking. To this endeavor, the authors present two different approaches. The first approach recasts the linking problem as the Generalised Traveling Salesman Problem, while the second uses an xgboost [30] classifier based on the connection density as determined by links in the KG. The two approaches also differ in that the former does not require training data, while the latter does. EARL was evaluated on the QALD-7 dataset [13] and achieved superior performance over its previous competitors.

### 3. The NLQxform System

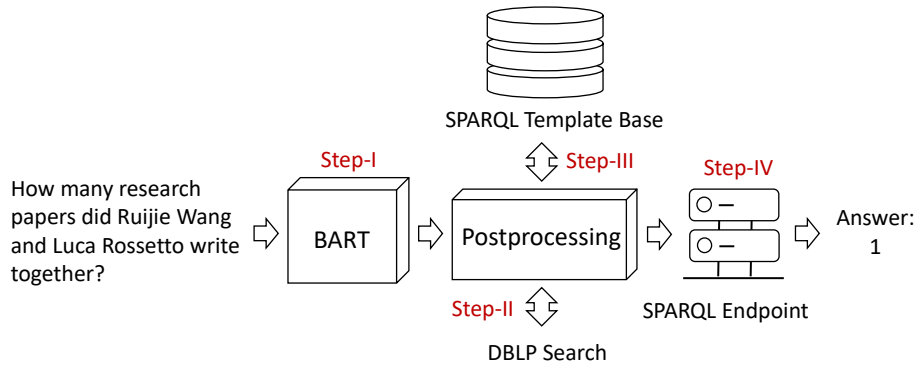
In this section, we elaborate on the NLQxform system. An overview of the system structure is presented in Fig. 1, where a natural language question is answered via four steps:

**Step-I: BART-based question to logical form translation.** In the first step, a BART [8] model (facebook/bart-base<sup>8</sup>) is used to translate the given question into a logical form that is very close to the final standard SPARQL query, except that entities in the logical form are still natural language mentions instead of URLs of entities. For example, given the question “*how many research papers did Ruijie Wang and Luca Rossetto write together,*” the BART model is expected to generate the logical form:

```
SELECT COUNT(DISTINCT ?answer) AS ?count
WHERE {
  ?answer <https://dblp.org/rdf/schema#authoredBy> <Ruijie Wang> .
  ?answer <https://dblp.org/rdf/schema#authoredBy> <Luca Rossetto> .
}
```

---

<sup>8</sup><https://huggingface.co/facebook/bart-base>



**Figure 1:** An overview of the NLQxform system and its four steps question-answering process.

Please note that the adopted BART model was not pre-trained on the question-to-SPARQL translation task. Therefore, it needs to be fine-tuned with training question-SPARQL query pairs.<sup>9</sup> Also, we make the following modifications to better suit the task: First, basic elements of SPARQL syntax, including clauses, functions, and modifiers (e.g., “SELECT”, “COUNT”, and “ORDER BY”), as well as some tokens frequently used in SPARQL queries (e.g., “{”, “}”, “(”, “)”, and “.”) are added as special tokens to the tokenizer of the language model (LM). Second, considering that the number of relations in the underlying knowledge graph is limited, we also add all relations as special tokens to the LM.<sup>10</sup> Therefore, the LM is able to recognize relations in given questions and directly generate logical forms with relation URLs.

**Step-II: DBLP search-based entity linking.** In the above step, entity mentions in the given question are recognized and used to construct the logical form. In this step, we search specific URLs of recognized entities using DBLP search APIs.<sup>11</sup> For example, regarding entities <Ruijie Wang> and <Luca Rossetto>, their correct URLs <<https://dblp.org/pid/57/5759-3>> and <<https://dblp.org/pid/156/1623>> can be obtained via requesting <https://dblp.org/search/author/api?q=Ruijie%20Wang> and <https://dblp.org/search/author/api?q=Luca%20Rossetto>.

**Step-III: SPAQRL template-based query correction.** This step is proposed to alleviate possible syntax and grammar errors in the generated logical forms. During training, we utilize ground-truth SPARQL queries to construct a SPARQL template base. For example, assuming the above question is a training question, the following template can be added:

```
SELECT COUNT(DISTINCT ?answer) AS ?count
WHERE {
    ?answer <https://dblp.org/rdf/schema#authoredBy> <entity_1> .
    ?answer <https://dblp.org/rdf/schema#authoredBy> <entity_2> .
}
```

<sup>9</sup>Please refer to Section 4 for information about training questions.

<sup>10</sup>Only 12 relations are used in our system for this challenge.

<sup>11</sup><https://dblp.org/faq/How+to+use+the+dblp+search+API.html>

It provides a gold example for future questions that require the count of publications authored by two given scholars. During inference, regarding each generated logical form, we find top-3 similar templates from the template base according to string similarity and pass them to the next step.

**Step-IV: SPARQL endpoint-based answer retrieval.** In this step, we fill in retrieved query templates with entity linking results to generate specific SPARQL queries and evaluate them via the SPARQL endpoint<sup>12</sup> of DBLP scholarly KG to retrieve answers. For example, the generated SPARQL query for the above question is expected to be

```
SELECT COUNT(DISTINCT ?answer) AS ?count
WHERE {
  ?answer <https://dblp.org/rdf/schema#authoredBy>
      <https://dblp.org/pid/57/5759-3> .
  ?answer <https://dblp.org/rdf/schema#authoredBy>
      <https://dblp.org/pid/156/1623> .
}
```

The final answer retrieved with this query is *1*.<sup>13</sup> It is worth mentioning that there could be several candidate URLs returned for one entity from Step-II. In this case, we follow the order of these URLs returned by the DBLP API and adopt the first generated SPARQL query that can return an answer in the endpoint.

## 4. Experiments

In this section, we report the experimental setup and evaluation results of NLQxform in the Scholarly QALD Challenge (Task 1: DBLP-QUAD – Knowledge Graph Question Answering over DBLP).<sup>14</sup>

**Dataset.** The employed QA dataset is DBLP-QuAD [19], which includes natural language questions posed over the DBLP scholarly KG. The training and validation sets of DBLP-QuAD, respectively, include 7,000 and 1,000 questions with annotations of correct answers, ground-truth SPARQL queries, and entity-linking results. They were used for the fine-tuning of the BART model. In the final phase of the challenge, 500 newly generated test questions<sup>15</sup> without answer annotations were provided for system evaluation.

**Evaluation Setup.** The entity linking results from Step-II and final answers of NLQxform were submitted to the Codalab platform,<sup>16</sup> where final results regarding entity linking and QA were automatically calculated. F1-score was used as the criterion for both tasks.

---

<sup>12</sup><https://dblp-kg.ltdemos.informatik.uni-hamburg.de/sparql>

<sup>13</sup>This result is with respect to the snapshot of the KG used by the endpoint.

<sup>14</sup><https://codalab.lisn.upsaclay.fr/competitions/14264>

<sup>15</sup><https://github.com/debayan/scholarly-QALD-challenge/blob/main/2023/datasets/codalab/finalphase/dblp-kgqa/dblp.heldout.500.questiononly.json>

<sup>16</sup><https://codalab.lisn.upsaclay.fr/competitions/14264#participate>

**Table 1**

The final evaluation results regarding entity linking and question answering. (best performance in **bold**, second best underlined)

Submission	F1 Score (Entity Linking)	F1 Score (Question Answering)
ID-544291	0.8283	0.0000
ID-544863	<u>0.8320</u>	0.0000
ID-557116	<b>0.8353</b>	0.0000
ID-545920	0.7100	0.0018
ID-556670	0.6235	0.2175
ID-547129	0.0000	<u>0.6619</u>
ID-557036 (NLQxform)	0.7961	<b>0.8488</b>

**Final Results.** We report the final results of public submissions<sup>17</sup> in Table 1. The following can be observed:

- On the entity linking task, NLQxform is ranked fourth regarding the F1 score. Compared to the best-performing system, the F1 score of our system is 4.7% lower. The main reason is that there is a lack of explicit supervision for the recognition of entities in Step-I, as the fine-tuning objective is the matching of whole queries without emphasis on entities. Also, due to an implementation issue, the linked entities that cannot result in valid SPARQL queries in Step-IV were not removed from our submission. Nevertheless, this shows that there is still room for improvement in our system. In addition, the precision and recall regarding entity linking are 0.81 and 0.79, respectively.
- On the QA task, NLQxform achieved the best performance, significantly improving over the second-best system (+28.2%). This demonstrates the overall effectiveness of our system and shows that the LM was successfully fine-tuned to learn the syntax and grammar of SPARQL. In addition to the above results, the precision and recall of our system on QA are 0.83 and 0.87, respectively.
- The QA task seems more challenging than the entity linking task in this challenge. On entity linking, most submissions achieved an F1 score above 0.6. However, on QA, only two systems are above 0.6, and the other two systems with a result are both below 0.3. This further demonstrates the competitiveness of NLQxform.

## 5. Conclusion and Outlook

In this paper, we present our QA system, NLQxform, which fine-tunes a transformer-based BART model and answers NLQs over the DBLP Scholarly KG via four steps: 1. BART-based question to logical form translation, 2. DBLP search-based entity linking, 3. SPARQL template-based query correction, and 4. SPARQL endpoint-based answer retrieval. NLQxform participated in the Scholarly QALD Challenge and ranked first on the QA task over DBLP-QUAD with a significant improvement over the second-best system (+28.2%), demonstrating the system’s effectiveness and competitiveness.

<sup>17</sup>[https://codalab.lisn.upsaclay.fr/competitions/public\\_submissions/14264](https://codalab.lisn.upsaclay.fr/competitions/public_submissions/14264)

A limitation of NLQxform is that the recognition of entities is not explicitly emphasized during the fine-tuning of the BART model, which led to an inferior entity linking performance. We will investigate this issue in future work. Also, some technical details of the system are not fully presented in this paper due to the limitation of space and readability considerations. They will be more comprehensively explained and analyzed in an extended version of this work.

## Acknowledgments

This work was partially funded by the Digital Society Initiative of the University of Zurich, the University Research Priority Program “Dynamics of Healthy Aging” at the University of Zurich, and the Swiss National Science Foundation through Projects “CrowdAlytics” (Grant Number 184994) and “MediaGraph” (Grant Number 202125).

## References

- [1] K. Wang, Z. Shen, C. Huang, C. Wu, Y. Dong, A. Kanakia, Microsoft academic graph: When experts are not enough, *Quant. Sci. Stud.* 1 (2020) 396–413. URL: [https://doi.org/10.1162/qss\\_a\\_00021](https://doi.org/10.1162/qss_a_00021). doi:10.1162/qss\_a\_00021.
- [2] J. Priem, H. A. Piwowar, R. Orr, Openalex: A fully-open index of scholarly works, authors, venues, institutions, and concepts, *CoRR abs/2205.01833* (2022). URL: <https://doi.org/10.48550/arXiv.2205.01833>. doi:10.48550/arXiv.2205.01833. arXiv:2205.01833.
- [3] S. Auer, V. Kovtun, M. Prinz, A. Kasprzik, M. Stocker, M. Vidal, Towards a knowledge graph for science, in: R. Akerkar, M. Ivanovic, S. Kim, Y. Manolopoulos, R. Rosati, M. Savic, C. Badica, M. Radovanovic (Eds.), *Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics, WIMS 2018, Novi Sad, Serbia, June 25-27, 2018, ACM, 2018*, pp. 1:1–1:6. URL: <https://doi.org/10.1145/3227609.3227689>. doi:10.1145/3227609.3227689.
- [4] D. Banerjee, S. Awale, R. Usbeck, C. Biemann, Dblp-quad: A question answering dataset over the DBLP scholarly knowledge graph, *CoRR abs/2303.13351* (2023). doi:10.48550/arXiv.2303.13351. arXiv:2303.13351.
- [5] J. Z. Pan, Resource description framework, in: S. Staab, R. Studer (Eds.), *Handbook on Ontologies, International Handbooks on Information Systems, Springer, 2009*, pp. 71–90. URL: [https://doi.org/10.1007/978-3-540-92673-3\\_3](https://doi.org/10.1007/978-3-540-92673-3_3). doi:10.1007/978-3-540-92673-3\_3.
- [6] W. W. W. Consortium, et al., Sparql 1.1 overview (2013). URL: <https://www.w3.org/TR/sparql11-overview/>.
- [7] E. Kaufmann, A. Bernstein, Evaluating the usability of natural language query languages and interfaces to semantic web knowledge bases, *Journal of Web Semantics* 8 (2010) 377–393. URL: <https://www.sciencedirect.com/science/article/pii/S1570826810000582>. doi:<https://doi.org/10.1016/j.websem.2010.06.001>, semantic Web Challenge 2009 User Interaction in Semantic Web research.
- [8] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, L. Zettlemoyer, BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension, in: *Proceedings of the 58th Annual Meeting of the As-*

- sociation for Computational Linguistics, ACL 2020, Online, July 5-10, 2020, Association for Computational Linguistics, 2020, pp. 7871–7880. doi:10.18653/v1/2020.acl-main.703.
- [9] E. Cabrio, P. Cimiano, V. López, A. N. Ngomo, C. Unger, S. Walter, QALD-3: multilingual question answering over linked data, in: P. Forner, R. Navigli, D. Tufis, N. Ferro (Eds.), Working Notes for CLEF 2013 Conference, Valencia, Spain, September 23-26, 2013, volume 1179 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2013. URL: <https://ceur-ws.org/Vol-1179/CLEF2013wn-QALD3-CabrioEt2013.pdf>.
- [10] C. Unger, C. Forascu, V. López, A. N. Ngomo, E. Cabrio, P. Cimiano, S. Walter, Question answering over linked data (QALD-4), in: L. Cappellato, N. Ferro, M. Halvey, W. Kraaij (Eds.), Working Notes for CLEF 2014 Conference, Sheffield, UK, September 15-18, 2014, volume 1180 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2014, pp. 1172–1180. URL: <https://ceur-ws.org/Vol-1180/CLEF2014wn-QA-UngerEt2014.pdf>.
- [11] C. Unger, C. Forascu, V. López, A. N. Ngomo, E. Cabrio, P. Cimiano, S. Walter, Question answering over linked data (QALD-5), in: L. Cappellato, N. Ferro, G. J. F. Jones, E. SanJuan (Eds.), Working Notes of CLEF 2015 - Conference and Labs of the Evaluation forum, Toulouse, France, September 8-11, 2015, volume 1391 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2015. URL: <https://ceur-ws.org/Vol-1391/173-CR.pdf>.
- [12] C. Unger, A. N. Ngomo, E. Cabrio, 6th open challenge on question answering over linked data (QALD-6), in: H. Sack, S. Dietze, A. Tordai, C. Lange (Eds.), Semantic Web Challenges - Third SemWebEval Challenge at ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Revised Selected Papers, volume 641 of *Communications in Computer and Information Science*, Springer, 2016, pp. 171–177. URL: [https://doi.org/10.1007/978-3-319-46565-4\\_13](https://doi.org/10.1007/978-3-319-46565-4_13). doi:10.1007/978-3-319-46565-4\_13.
- [13] R. Usbeck, A. N. Ngomo, B. Haarmann, A. Krithara, M. Röder, G. Napolitano, 7th open challenge on question answering over linked data (QALD-7), in: Semantic Web Challenges - 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28 - June 1, 2017, Revised Selected Papers, volume 769 of *Communications in Computer and Information Science*, Springer, 2017, pp. 59–69. doi:10.1007/978-3-319-69146-6\_6.
- [14] R. Usbeck, A. N. Ngomo, F. Conrads, M. Röder, G. Napolitano, 8th Challenge on Question Answering over Linked Data (QALD-8) (invited paper), in: Joint proceedings of the 4th Workshop on Semantic Deep Learning (SemDeep-4) and NLIWoD4: Natural Language Interfaces for the Web of Data (NLIWOD-4) and 9th Question Answering over Linked Data challenge (QALD-9) co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, California, United States of America, October 8th - 9th, 2018, volume 2241 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 51–57. URL: <https://ceur-ws.org/Vol-2241/paper-05.pdf>.
- [15] R. Usbeck, R. H. Gusmita, A. N. Ngomo, M. Saleem, 9th Challenge on Question Answering over Linked Data (QALD-9) (invited paper), in: Joint proceedings of the 4th Workshop on Semantic Deep Learning (SemDeep-4) and NLIWoD4: Natural Language Interfaces for the Web of Data (NLIWOD-4) and 9th Question Answering over Linked Data challenge (QALD-9) co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, California, United States of America, October 8th - 9th, 2018, volume 2241 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 58–64. URL: <https://ceur-ws.org/Vol-2241/paper-06.pdf>.



- [16] D. Banerjee, P. A. Nair, J. N. Kaur, R. Usbeck, C. Biemann, Modern baselines for SPARQL semantic parsing, in: SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022, ACM, 2022, pp. 2260–2265. doi:10.1145/3477495.3531841.
- [17] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *J. Mach. Learn. Res.* 21 (2020) 140:1–140:67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [18] A. See, P. J. Liu, C. D. Manning, Get to the point: Summarization with pointer-generator networks, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, Association for Computational Linguistics, 2017, pp. 1073–1083. doi:10.18653/v1/P17-1099.
- [19] D. Banerjee, S. Awale, R. Usbeck, C. Biemann, Dblp-quad: A question answering dataset over the DBLP scholarly knowledge graph, *CoRR abs/2303.13351* (2023). doi:10.48550/arXiv.2303.13351. arXiv:2303.13351.
- [20] D. Diefenbach, K. D. Singh, P. Maret, Wdaqua-core0: A question answering component for the research community, in: Semantic Web Challenges - 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28 - June 1, 2017, Revised Selected Papers, volume 769 of *Communications in Computer and Information Science*, Springer, 2017, pp. 84–89. doi:10.1007/978-3-319-69146-6\_8.
- [21] D. Diefenbach, K. Singh, P. Maret, Wdaqua-core1: a question answering service for rdf knowledge bases, in: Companion Proceedings of the The Web Conference 2018, 2018, pp. 1087–1091.
- [22] D. Diefenbach, K. D. Singh, P. Maret, Wdaqua-core1: A question answering service for RDF knowledge bases, in: Companion of the The Web Conference 2018 on The Web Conference 2018, WWW 2018, Lyon, France, April 23-27, 2018, ACM, 2018, pp. 1087–1091. doi:10.1145/3184558.3191541.
- [23] D. Diefenbach, K. D. Singh, P. Maret, Wdaqua-core1: A question answering service for RDF knowledge bases, in: P. Champin, F. Gandon, M. Lalmas, P. G. Ipeirotis (Eds.), Companion of the The Web Conference 2018 on The Web Conference 2018, WWW 2018, Lyon, France, April 23-27, 2018, ACM, 2018, pp. 1087–1091. URL: <https://doi.org/10.1145/3184558.3191541>. doi:10.1145/3184558.3191541.
- [24] L. Zou, R. Huang, H. Wang, J. X. Yu, W. He, D. Zhao, Natural language question answering over RDF: a graph data driven approach, in: International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014, ACM, 2014, pp. 313–324. doi:10.1145/2588555.2610525.
- [25] D. Banerjee, Arefa, R. Usbeck, C. Biemann, Dbplink: An entity linker for the DBLP scholarly knowledge graph, *CoRR abs/2309.07545* (2023). URL: <https://doi.org/10.48550/arXiv.2309.07545>. doi:10.48550/arXiv.2309.07545. arXiv:2309.07545.
- [26] N. Steinmetz, Entity linking for KGQA using AMR graphs, in: The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings, volume 13870 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 122–138. doi:10.1007/978-3-031-33455-9\_8.
- [27] R. Usbeck, R. H. Gusmita, A. N. Ngomo, M. Saleem, 9th challenge on question answering over linked data (QALD-9) (invited paper), in: K. Choi, L. E. Anke, T. Declerck, D. Gromann,

- J. Kim, A. N. Ngomo, M. Saleem, R. Usbeck (Eds.), Joint proceedings of the 4th Workshop on Semantic Deep Learning (SemDeep-4) and NLIWoD4: Natural Language Interfaces for the Web of Data (NLIWOD-4) and 9th Question Answering over Linked Data challenge (QALD-9) co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, California, United States of America, October 8th - 9th, 2018, volume 2241 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018, pp. 58–64. URL: <https://ceur-ws.org/Vol-2241/paper-06.pdf>.
- [28] D. Diomedi, A. Hogan, Entity linking and filling for question answering over knowledge graphs, in: Proceedings of the 7th Natural Language Interfaces for the Web of Data (NLIWoD) co-located with the 19th European Semantic Web Conference (ESWC 2022), Hersonissos, Greece, May 29th, 2022, volume 3196 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 9–24. URL: <https://ceur-ws.org/Vol-3196/paper2.pdf>.
- [29] M. Dubey, D. Banerjee, D. Chaudhuri, J. Lehmann, EARL: joint entity and relation linking for question answering over knowledge graphs, in: 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*, Springer, 2018, pp. 108–126. doi:10.1007/978-3-030-00671-6\_7.
- [30] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016, pp. 785–794. doi:10.1145/2939672.2939785.