

Leveraging small language models for Text2SPARQL tasks to improve the resilience of AI assistance

Felix Brei^{1,*}, Johannes Frey^{2,3} and Lars-Peter Meyer^{1,3}

¹ETi Competence Center @ Institute for Applied Informatics, Germany, <https://cc-eti.org>

²KMI Competence Center @ Institute for Applied Informatics, Germany, <https://kmi-leipzig.de>

³Institute of Computer Science, Leipzig University, Germany, <https://cs.uni-leipzig.de>

Abstract

In this work we will show that language models with less than one billion parameters can be used to translate natural language to SPARQL queries after fine-tuning. Using three different datasets ranging from academic to real world, we identify prerequisites that the training data must fulfill in order for the training to be successful. The goal is to empower users of semantic web technology to use AI assistance with affordable commodity hardware, making them more resilient against external factors.

Keywords

Language models, SPARQL generation, Question Answering

1. Introduction

The usage of Large Language Models (LLMs) has increased exponentially since the advent of ChatGPT. According to Similarweb, the website of OpenAI alone was visited more than 1.6 billion times by February 2024¹. In addition to that, Microsoft has launched several AI assistants called 'Copilots' which are based on LLMs^{2,3}, as well as Google releasing their AI called Bard which is now known as Gemini^{4,5}. This suggests that the big tech companies believe in the potential of LLMs to become part of our daily lives, just like smartphones or computers in general. But do they hold up to the expectations?

Several test suites were derived to assess the generative capabilities of LLMs, for example TruthfulQA[1], HellaSwag [2] or the Abstraction and Reasoning Corpus (ARC) [3]. These test suites, among others, are run regularly on the latest entries to the LLM circus and the results for open LLMs are presented publicly on the Huggingface OpenLLM leaderboard [4]. We can see that the performance increases drastically over time, with Bloom [5] scoring an average of


Third International Workshop on Linked Data-driven Resilience Research (D2R2'24) co-located with ESWC 2024, May 27th, 2024, Hersonissos, Greece

*Corresponding author.

✉ brei@infai.org (F. Brei); frey@informatik.uni-leipzig.de (J. Frey); lpmeier@infai.org (L. Meyer)

🆔 0009-0008-5245-6655 (F. Brei); 0000-0003-3127-0815 (J. Frey); 0000-0001-5260-5181 (L. Meyer)

© 2024 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)

¹<https://www.similarweb.com/website/openai.com/#overview>

²<https://blogs.microsoft.com/blog/2023/09/21/announcing-microsoft-copilot-your-everyday-ai-companion/>

³<https://github.blog/2023-11-08-universe-2023-copilot-transforms-github-into-the-ai-powered-developer-platform/>

⁴<https://blog.google/technology/ai/bard-google-ai-search-updates/>

⁵<https://blog.google/technology/ai/google-gemini-ai/>

46.06 in August 2023 and Smaug-72B⁶ holding the record in February 2024 with a score of 80.48, only half a year later.

These test suites however cover mostly natural language domains, like the task to continue a sentence, answer questions or extract information from a paragraph of text. Based on the experience from early experiments [6], a test suite [7] was developed that evaluates capabilities of LLMs to interface knowledge graphs and assist in knowledge engineering tasks. While the smaller open-source GPT4All models severely struggled, the state-of-the-art commercial LLMs GPT4 and Claude showed promising results [8] and a trend of performance improvements over the course of 2023 [9] in dealing with KGs in Turtle format.

Alas, these results come with several caveats:

1. The commercial LLMs that were tested are all hosted externally. This can be problematic regarding data protection, because a user has to send a information to a third party.
2. Because of their sheer size (GPT4 has one trillion parameters⁷), running these models locally is prohibitively costly and therefore not an option for a lot of research institutes and other parties. On top of that, training a model of such size is also extremely expensive⁸.
3. Even these commercial models were at the time of writing still significantly challenged by SPARQL query generation or RML mapping generation [8, 10] indicating a need for specific training or fine-tuning of all models w.r.t. handling those tasks in a reliable and efficient way.
4. Since all these larger are hosted on third party platforms, users are at the mercy of the vendors to keep the services running and affordable. However, vendors suddenly changing their licensing and cost model has already happened in the past⁹, as well as deep sea cables being damaged¹⁰, separating certain areas of the world from the internet and leaving local companies only with the computational resources they have on site.

So we ask ourselves the following question: Given a single task that we want so solve using LLMs, is it possible to achieve a similar performance of these large models with a much smaller one? This would enable small businesses to use AI assistance with affordable hardware they can host on site, increasing their resilience against outages, vendors changing their pricing models, disruption due to trade embargoes and other external factors.

As a first step into this direction, we study the task of translating a natural language question into a SPARQL query because we think that this task enables people who are not familiar with SPARQL to extract knowledge and insights from a knowledge graph which would otherwise not be possible for them. The paper is organized as follows: First, we look at related research in this field and explain where we fit into the big picture. Then we explain the setup of our experiments, namely which model families were chosen and why and which datasets we trained them on. After that, we present and explain the results of our work and finally, we draw conclusions and give an outlook on the directions that our research will head next.

⁶<https://huggingface.co/abacusai/Smaug-72B-v0.1>

⁷<https://www.semafor.com/article/03/24/2023/the-secret-history-of-elon-musk-sam-altman-and-openai>

⁸<https://www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over/>

⁹<https://www.theverge.com/2023/9/12/23870547/unit-price-change-game-development>

¹⁰<https://edition.cnn.com/2024/03/04/business/red-sea-cables-cut-internet/index.html>

2. Related Work

Current approaches focus on fine-tuning large language models. For example the authors of [11] propose a methodology for fine-tuning OpenLLaMA to generate SPARQL queries over life science knowledge graphs using data augmentation techniques, such as providing meaningful variable names and inline comments, improving the performance of the model in generating accurate SPARQL queries. The authors of [12] use Llama as their basis for fine-tuning to generate SPARQL queries over Wikidata.

These two papers have shown that translating natural language to SPARQL queries is possible, but they use models with at least three (OpenLLaMA) resp. seven (LLaMA) billion parameters. The hardware required to train these models can be expensive, which is why we want to explore models that are even smaller.

Smaller, fine-tuned models for one specific task are also able to beat the performance of LLMs, e.g. SQLCoder-7B¹¹ performs better on SQL than state of the art GPT4. Our research is comparable to that, but with much less parameters and SPARQL instead of SQL.

[13] manages to fine-tune T5 on SPARQL queries for Wikidata, but to achieve these results, the data had to be preprocessed in a way that is specific to T5. Furthermore, while this paper explores other ways to tackle this task in general, it only looks at T5 instead of other model families as we do.

[14] gives a comprehensive overview and performs a comparison of pre-trained LMs (PLMs), non-pre-trained LMs (NPLMs), and LLMs, testing various fine-tuning methods using LLMs. [15] fine-tunes a lightweight model for SPARQL generation using synthetic training data generated by the FlexKBQA framework on a target knowledge graph (sampling structured query templates that are converted into SPARQL query instances and translated into natural language questions using LLMs). The light-weight model can perform further self-guided training on real queries to address a distribution shift between synthetic and real queries. [16] uses a GPT model to investigate what parts of the Text2SPARQL task are the hardest for the model to solve so appropriate countermeasures can be taken.

[17] proposes a whole new architecture specific for SPARQL generation based on GPT. This direction we assume promising for the future, but here we are focusing on more foundational research first to understand which model families work best on a given dataset and why.

3. Experimental Setup

3.1. Model families

As was mentioned in the introduction, the focus of our work is to fine-tune language models that can be considered small by modern standards. We chose one billion parameters as an arbitrary limit on the number of parameters, but as a general guideline we consulted the Steam Hard- and Software Survey¹² and found that 57.22% of their users use a GPU with 8GB of VRAM or more (January 2024). A model with less than one billion parameters should fit into this amount of VRAM comfortably, showing again that these LLMs can be trained and run locally.

¹¹<https://huggingface.co/defog/sqlcoder-7b-2>

¹²<https://store.steampowered.com/hwsurvey/>

Another consideration is the public availability of the models. We believe that research should be available to anyone who is interested and this should be reflected in the choice of models. Therefore, we only select models that are openly available on Huggingface.

Following these criteria, we observe quickly that there are only three large model families that fit the bill, which we introduce here briefly. A full list of models evaluated is given in table 1

3.1.1. T5 and Flan-T5

In June 2020 Google released an LLM called Text-To-Text Transfer Transformer, or T5 in short [18]. The base version consists of roughly 220 million parameters, with smaller and larger versions available. With T5, Google wanted to provide a single LLM that can solve any NLP task like text classification, sentiment analysis and so on. A user must provide a prefix like *'Translate the following sentence to french:'* and the LLM then infers how to process the rest of the prompt. In 2022, researchers at Google released new versions of T5 called FLAN-T5 [19] (FLAN stands for *fine-tuning language models* [20]) which, according to the authors, should outperform T5 on any given task.

3.1.2. BART

BART was developed by Facebook and released in October 2019 [21]. It consists of 139 million parameters and is a combination of a BERT-like encoder [22] with a GPT-like autoregressive decoder [23]. In August 2020, a multilingual version called mBART was released [24]. The authors put special emphasis on the fact that BART is just a pretrained model and needs to be fine-tuned for a given specific task. We also included mREBEL models as a specialized version of BART for multilingual relation extraction [25] since it was finetuned with knowledge graphs in mind.

3.1.3. M2M100 and NLLB-200

The M2M100 model was introduced in 2020 [26] as a many-to-many translation tool for 100 languages. The original version consists of 1.3 billion parameters which exceeds the upper bound we imposed. But there is a distilled version available directly from the Facebook research team at Huggingface called M2M100-418M¹³ which we use in our experiments.

Its successor, the NLLB-200 model, was introduced in 2022 [27] and stands for *'no language left behind'*. Again we use the distilled version NLLB-200-Distilled-600M¹⁴ instead of the 3.3 billion full version of the model. As the authors state, the model is *'primarily intended for research in machine translation'* which fits our bill perfectly.

This leaves us with a selection of models to be assessed in our experiment that can be seen in table 1.

¹³https://huggingface.co/facebook/m2m100_418M

¹⁴<https://huggingface.co/facebook/nllb-200-distilled-600M>

Table 1

Model names and their number of parameters, as used in our experiments.

Name	No. parameters	Name	No. parameters
T5-Small	60.5M	BART-Base	139M
T5-Base	223M	BART-Large	406M
T5-Large	738M	mBART-LARGE-50	611M
FLAN-T5-Small	77M	mREBEL-Base	484M
FLAN-T5-Base	248M	mREBEL-Large	611M
FLAN-T5-Large	783M	M2M100-418M	418M
		NLLB200-Distilled-600M	600M

3.2. Datasets used for Fine Tuning and Evaluation

In order to study how well the models can be fine-tuned towards a target KG, we use three evaluation datasets from different domains and with varying complexity. These datasets are comprised of a number of natural language questions, which are mapped to a SPARQL query w.r.t. the target KG. For the first two datasets (organizational graph and CoyPu graph) we generate questions and queries by sending the graph via the OpenAI API to GPT4 and prompting it to generate tuples of natural language question, matching SPARQL query, and the expected result of the query. These tuples are filtered by checking if the results that the SPARQL query returns match with the expected results. Both datasets are then augmented by sending each remaining question again to GPT and asking it to paraphrase the question, giving us a total of two natural language questions per SPARQL query.

3.2.1. Organizational Graph

Introduced in [6], this small knowledge graph uses established vocabularies to describe an organization with departments and employees. There is a clear schema that maps person and department names to their corresponding RDF resource, for example *"Anne Miller"* maps to :anne while *"Bob Tanner"* maps to :bob. In this dataset and the next we also let the language model omit the prefix definitions for the queries and assume they are already present in the preamble of the executed SPARQL query. Using GPT4 we generated a dataset consisting of 69 datapoints, which were split into 53 tuples for training and 16 for testing.

3.2.2. A subset of the CoyPu graph

The CoyPu project¹⁵ aims to improve supply chain resiliency for corporations by combining different data sources about public infrastructure, trades and trade agreements, events like disasters and conflicts and many more into a large knowledge graph. Querying this knowledge graph has the potential to help businesses identify risks like single points of failures and mitigate them. This usefulness combined with the fact that the other two datasets have more of an academic background made us decide that we use a subset of the CoyPu knowledge graph as another dataset for training. Creating a viable subset lead to its own difficulties and hurdles

¹⁵<https://coypu.org/>

however, which we consider as future work. This dataset contains 131 tuples in total, which were split into 105 for training and 26 for testing.

3.2.3. QALD10

The *Question Answering over Linked Data* (QALD) dataset is a standard benchmark¹⁶ with QALD10 being the latest incarnation [28]. It consists of SPARQL queries along with matching questions in different natural languages, w.r.t. Wikidata. In this work, we focus on English and filter the dataset accordingly. This dataset is especially difficult for a language model to handle because there is no clear indication how to link entities from a given question like "Barack Obama" to their Wikidata entity ID (:Q76), giving rise to a whole field of research called Entity Linking [29].

3.3. Fine-tuning

For every evaluation dataset individually, we perform fine-tuning of the selected models using PyTorch (100 epochs). Since a single run of fine-tuning does not hold much statistical significance and involves random parameters, we performed isolated runs of the training for a total of ten times. For each run we shuffle the training data with a predetermined random seed to make the results reproducible. Specifically, each run is given an ID from *R01* to *R10* and the seeds are generated by calculating the SHA512 sum of the ID and taking the first eight digits, so *R01* results in the seed 99975818, *R02* in 56899599 and so on.

4. Results

In the following subsections we only include those language models in the plots that generated at least one correct query. The T5 family consistently generated not a single correct query on the organizational graph which is why it is absent in the result tables and figures. In fact, all T5 models did not produce a single correct result across all runs.

To generate the datapoints for each plot, we interrupted the training every five epochs and made the language models translate the questions from the evaluation split into SPARQL queries. We then executed the queries and compared the result sets to determine whether the answers is correct.

4.1. Organizational Graph

Figure 1 shows that all models from the BART and M2M100 families manage to learn the structure of the knowledge graph at least to a certain degree. When taking the best results for each model, aside from NLLB-200, all models turn at least eleven of the sixteen questions into correct SPARQL queries. The performance however fluctuates extremely during the course of the training which is indicative of overfitting.

Repeating the experiment we can see that performance varies a bit depending on the order the training data is ingested into the network. The statistics are shown in table 3 and the raw

¹⁶<https://www.nliwod.org/challenge>

Table 2

An example of the errors that were made in the context of the organizational graph: no binding variable, wrong entity ID :charles, wrong property foaf:firstName, wrong literal BobTanner

Question	What is the surname of Bob Tanner?
Gold answer	SELECT ?surname WHERE { :bob foaf:surname ?surname . }
Generated query	SELECT ?surname WHERE { :charles foaf:firstName 'BobTanner' }




Figure 1: Number of correct SPARQL translations across epochs for the organizational graph for one fine-tuning run per model. 16 questions were presented.

data is plotted on the left side of figure 3. We can see that for this dataset, BART-L performs best (as well as the other sizes of BART), with M2M100 being close behind. Another thing we see from the left plot in figure 3 is that except for one outlier coming from mREBEL-L, the success of fine-tuning is reliable and reproducible.

Looking at common errors made during the translation, we found that the best models rarely generated SPARQL that could not be parsed, but they rather mixed up terms and injected parts of the training data into the queries. An example is shown in table 2.

4.2. CoyPu

In figure 2 we can see that the performance during the first run of the experiment varies less drastically than for the Organizational Graph. The standard deviations seen in table 5 are similar though so we think this is just a coincidence. Again, the (FLAN-)T5 models never generate even a single correct query, so they are excluded from consecutive runs.

We can also see that for this dataset, M2M100 outperformed the other models and BART-L is in fact one of the worst, which is a complete shift from the results before. This again shows

Table 3

Average number of correctly translated questions in the context of the organizational graph. Standard deviation is also shown as a percentage of the average to measure the reliability of the fine-tuning.

Model name	Average	Standard deviation	Std. dev percent
BART	12.90	1.14	8.80
BART-L	13.30	0.64	4.81
mBART-50	12.80	0.75	5.85
mREBEL-L	11.10	1.92	17.31
M2M100	12.50	1.02	8.20
NLLB-200	8.20	0.75	9.13

Table 4

An example of the errors that were made in the context of the CoyPu graph: latitude and longitude got mixed.

The ellipsis on the IRI was inserted by ourselves to keep the line shorter. In fact, the language model generated the correct IRI to use in the query.

Question	What is the latitude of the port with the ID 'AUDKB'?
Gold answer	SELECT ?latitude WHERE { <https://data.coypu.org/... ns2:hasLatitude ?latitude }
Generated query	SELECT ?latitude WHERE { <https://data.coypu.org/... ns2:hasLatitude ?longitude }

Table 5

Average number of correctly translated questions in the context of the CoyPu graph. Standard deviation is also shown as a percentage of the average to measure the reliability of the fine-tuning.

Model name	Average	Standard deviation	Std. dev percent
BART	11.90	0.83	6.98
BART-L	14.00	1.18	8.45
mBART-50	18.50	1.20	6.51
mREBEL-L	17.50	1.02	5.86
M2M100	19.30	0.90	4.66
NLLB-200	17.80	0.87	4.90

that one should always evaluate more than a single model since the performance seems to be tied to the structure of the underlying data for fine-tuning.

Again BART-L did not generate a lot of parsing errors, but instead mixes up terms from the supply chain domain, see for example table 4.

On the other hand, the main errors that M2M100 made were generating SPARQL queries that could not be parsed (i.e. grammatically incorrect).

4.3. QALD10

The language models have a really hard time with the QALD10 dataset. While the structure of the generated query comes close to the correct ones, the models cannot handle the translation from entity names like *Kobe Bryant* to their corresponding Wikidata IDs like Q25369. We




Figure 2: Number of correct SPARQL translations across epochs for the CoyPu dataset for one fine-tuning run per model. 26 questions were presented.

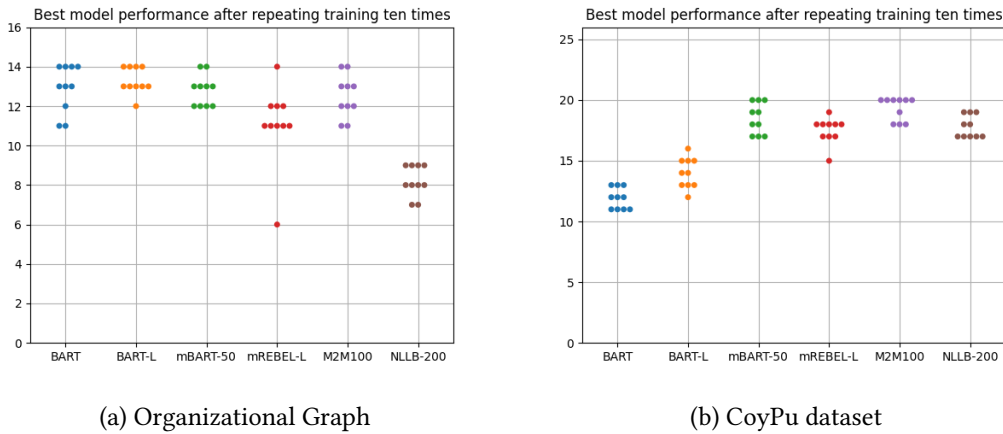


Figure 3: Repeating the experiment ten times for both datasets then selecting the best model from each run

expected this to happen since the whole field of entity linking is ongoing research and far from trivial.

Another problem here is that the QALD10 dataset requires the inclusion of all necessary prefix definitions as part of the query, which was not a requirement both in the organizational as well as the CoyPu graph datasets.

To provide some numbers for clarification: The best performing model was M2M100-418M.

The validation dataset contains 394 questions, and only 104 were turned into SPARQL queries that could be parsed. Out of these 104 parsable questions, 51 returned an empty result set. The remaining queries except for three used COUNT and returned 0 because the result set of the underlying query was empty. The three final ones did return wrong bindings.

BART-L only managed to translate one single question into a valid SPARQL query, but the result set was not correct. Interestingly, mBART-L generated 101 parsable SPARQL queries, which makes it a close second to M2M100-418M. The error distribution is about the same as for M2M100-418M though, so no question was correctly answered.

5. Conclusion and Future Work

In this paper we have shown that fine-tuning language models for the translation of natural language to SPARQL queries is indeed possible, although there are some limitations like the requirement of a clear and concise mapping from entities in a question to entities in the knowledge graph, like *Anne Miller* to `:anne` instead of `:person1234`. If this requirement is met, both the BART family as well as the M2M100 family are able to fulfill this task.

There is a large amount of avenues that can be explored from here on. First, we should find a better way to define the limit on the number of parameters that a model is allowed to have. Here we have focused on maximum one billion, as this is a limit for most consumer GPUs, but probably there is a connection between model size and sparql generation capabilities.

Secondly, we want to explore how these results can be used to deploy a fine-tuned language model next to a RAG agent to improve its question answering capabilities. So far, LLMs used by RAG agents often lack the ability to correctly apply aggregate functions, which could be remedied by offering the RAG agent a SPARQL query as another source of information.

Since all these models are open source, we can also modify them by manipulating existing layers as well as removing some or inserting new ones. This might be a way to reduce inference time and improve the performance even further. One could also derive completely new models from scratch, since most pre-training datasets are openly available and pre-training is fast due to the small size of the models.

And on top of that, we still have the problem that both the organizational graph dataset as well as the CoyPu dataset were generated using GPT which defeats the purpose of being independent from third parties. We will also investigate in the future how the training data can be generated with open source LLMs like Falcon, Bloom and others so even this step of the pipeline can be executed locally. Here it does not matter if we have enough GPU memory available, since the creation of the training and testing datasets is only done once, so it is not an issue if this step takes a bit longer.

The goal of this paper was to do a small survey of the out-of-the-box capabilities of readily available language models. What we have seen so far looks promising and there is a lot of intriguing research to be done in the near future.

Acknowledgments

This work was partially supported by grants from the German Federal Ministry for Economic Affairs and Climate Action (BMWK) to the CoyPu project (01MK21007A) and KISS project (01MK22001A) as well as from the German Federal Ministry of Education and Research (BMBF) to the project StahlDigital (13XP5116B).

References

- [1] S. Lin, J. Hilton, O. Evans, Truthfulqa: Measuring how models mimic human falsehoods, 2022. [arXiv:2109.07958](https://arxiv.org/abs/2109.07958).
- [2] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, Y. Choi, Hellaswag: Can a machine really finish your sentence?, 2019. [arXiv:1905.07830](https://arxiv.org/abs/1905.07830).
- [3] F. Chollet, On the measure of intelligence, 2019. [arXiv:1911.01547](https://arxiv.org/abs/1911.01547).
- [4] E. Beeching, C. Fourrier, N. Habib, S. Han, N. Lambert, N. Rajani, O. Sanseviero, L. Tunstall, T. Wolf, Open llm leaderboard, https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard, 2023.
- [5] B. Workshop, Bloom: A 176b-parameter open-access multilingual language model, 2023. [arXiv:2211.05100](https://arxiv.org/abs/2211.05100).
- [6] L.-P. Meyer, C. Stadler, J. Frey, N. Radtke, K. Junghanns, R. Meissner, G. Dziwis, K. Bulert, M. Martin, Llm-assisted knowledge graph engineering: Experiments with chatgpt, in: C. Zinke-Wehlmann, J. Friedrich (Eds.), First Working Conference on Artificial Intelligence Development for a Resilient and Sustainable Tomorrow (AITomorrow) 2023, Informatik aktuell, 2023. doi:10.1007/978-3-658-43705-3_8. [arXiv:2307.06917](https://arxiv.org/abs/2307.06917).
- [7] L.-P. Meyer, J. Frey, K. Junghanns, F. Brei, K. Bulert, S. Gründer-Fahrer, M. Martin, Developing a scalable benchmark for assessing large language models in knowledge graph engineering, in: N. Keshan, S. Neumaier, A. L. Gentile, S. Vahdati (Eds.), Proceedings of the Posters and Demo Track of the 19th International Conference on Semantic Systems (SEMANTICS 2023), 2023. URL: <https://ceur-ws.org/Vol-3526/paper-04.pdf>. [arXiv:2308.16622](https://arxiv.org/abs/2308.16622).
- [8] J. Frey, L. Meyer, N. Arndt, F. Brei, K. Bulert, Benchmarking the abilities of large language models for RDF knowledge graph creation and comprehension: How well do llms speak turtle?, in: M. Alam, M. Cochez (Eds.), Proceedings of the Workshop on Deep Learning for Knowledge Graphs (DL4KG 2023) co-located with the 21th International Semantic Web Conference (ISWC 2023), Athens, November 6-10, 2023, volume 3559 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: <https://ceur-ws.org/Vol-3559/paper-3.pdf>. [arXiv:2309.17122](https://arxiv.org/abs/2309.17122).
- [9] J. Frey, L.-P. Meyer, F. Brei, S. Gründer-Fahrer, M. Martin, Assessing the evolution of llm capabilities for knowledge graph engineering in 2023, in: A. M. Peñuela, O. Corcho, P. Groth, E. Simperl, V. Tamma, A. Nuzzolese, M. Poveda-Villalón, M. Sabou, V. Presutti, I. Celino, A. Revenko, J. Raad, B. Sartini, P. Lisena (Eds.), ESWC 2024 Satellite Events, Hersonissos, Crete, Greece, May 26 - 30, 2024, Proceedings., 2024. URL: https://www.researchgate.net/publication/378804553_Assessing_the_Evolution_of_LLM_capabilities_for_Knowledge_Graph_Engineering_in_2023.

- [10] M. Hofer, J. Frey, E. Rahm, Towards self-configuring knowledge graph construction pipelines using llms - a case study with rml, in: Fifth International Workshop on Knowledge Graph Construction @ ESWC2024, 2024.
- [11] J. C. Rangel, T. M. de Farias, A. C. Sima, N. Kobayashi, Sparql generation: an analysis on fine-tuning openllama for question answering over a life science knowledge graph, 2024. [arXiv:2402.04627](https://arxiv.org/abs/2402.04627).
- [12] S. Xu, S. Liu, T. Culhane, E. Pertseva, M.-H. Wu, S. Semnani, M. Lam, Fine-tuned LLMs know more, hallucinate less with few-shot sequence-to-sequence semantic parsing over Wikidata, in: H. Bouamor, J. Pino, K. Bali (Eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Singapore, 2023, pp. 5778–5791. URL: <https://aclanthology.org/2023.emnlp-main.353>. doi:10.18653/v1/2023.emnlp-main.353.
- [13] D. Banerjee, P. A. Nair, J. N. Kaur, R. Usbeck, C. Biemann, Modern baselines for sparql semantic parsing, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22, ACM, 2022. doi:10.1145/3477495.3531841.
- [14] P. A. K. K. Diallo, S. Reyd, A. Zouaq, A comprehensive evaluation of neural sparql query generation from natural language questions, 2024. [arXiv:2304.07772](https://arxiv.org/abs/2304.07772).
- [15] Z. Li, S. Fan, Y. Gu, X. Li, Z. Duan, B. Dong, N. Liu, J. Wang, Flexkbqa: A flexible llm-powered framework for few-shot knowledge base question answering, ArXiv abs/2308.12060 (2023). URL: <https://api.semanticscholar.org/CorpusID:261076103>.
- [16] D. Bustamante, H. Takeda, Sparql generation with entity pre-trained gpt for kg question answering, ArXiv abs/2402.00969 (2024). URL: <https://api.semanticscholar.org/CorpusID:267406567>.
- [17] M. R. A. H. Rony, U. Kumar, R. Teucher, L. Kovriguina, J. Lehmann, Sgpt: A generative approach for sparql query generation from natural language questions, IEEE Access 10 (2022) 70712–70723. doi:10.1109/ACCESS.2022.3188714.
- [18] 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, Journal of Machine Learning Research 21 (2020) 1–67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [19] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, S. Narang, G. Mishra, A. Yu, V. Zhao, Y. Huang, A. Dai, H. Yu, S. Petrov, E. H. Chi, J. Dean, J. Devlin, A. Roberts, D. Zhou, Q. V. Le, J. Wei, Scaling instruction-finetuned language models, 2022. doi:10.48550/ARXIV.2210.11416.
- [20] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, Q. V. Le, Finetuned language models are zero-shot learners, in: International Conference on Learning Representations, 2022. URL: <https://openreview.net/forum?id=gEZrGCozdqR>.
- [21] 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, CoRR abs/1910.13461 (2019). [arXiv:1910.13461](https://arxiv.org/abs/1910.13461).
- [22] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [23] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, Language models are

- unsupervised multitask learners (2019).
- [24] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, A. Fan, Multilingual translation with extensible multilingual pretraining and finetuning (2020). [arXiv:2008.00401](https://arxiv.org/abs/2008.00401).
 - [25] P.-L. Huguet Cabot, S. Tedeschi, A.-C. Ngonga Ngomo, R. Navigli, Red^{fm}: a filtered and multilingual relation extraction dataset, in: Proc. of the 61st Annual Meeting of the Association for Computational Linguistics: ACL 2023, Association for Computational Linguistics, Toronto, Canada, 2023. URL: <https://arxiv.org/abs/2306.09802>.
 - [26] A. Fan, S. Bhosale, H. Schwenk, Z. Ma, A. El-Kishky, S. Goyal, M. Baines, O. Celebi, G. Wenzek, V. Chaudhary, N. Goyal, T. Birch, V. Liptchinsky, S. Edunov, E. Grave, M. Auli, A. Joulin, Beyond english-centric multilingual machine translation, 2020. [arXiv:2010.11125](https://arxiv.org/abs/2010.11125).
 - [27] N. Team, M. R. Costa-jussà, J. Cross, O. Çelebi, M. Elbayad, K. Heafield, K. Heffernan, E. Kalbassi, J. Lam, D. Licht, J. Maillard, A. Sun, S. Wang, G. Wenzek, A. Youngblood, B. Akula, L. Barrault, G. M. Gonzalez, P. Hansanti, J. Hoffman, S. Jarrett, K. R. Sadagopan, D. Rowe, S. Spruit, C. Tran, P. Andrews, N. F. Ayan, S. Bhosale, S. Edunov, A. Fan, C. Gao, V. Goswami, F. Guzmán, P. Koehn, A. Mourachko, C. Ropers, S. Saleem, H. Schwenk, J. Wang, No language left behind: Scaling human-centered machine translation, 2022. [arXiv:2207.04672](https://arxiv.org/abs/2207.04672).
 - [28] A. Perevalov, D. Diefenbach, R. Usbeck, A. Both, Qald-9-plus: A multilingual dataset for question answering over dbpedia and wikidata translated by native speakers, in: 2022 IEEE 16th International Conference on Semantic Computing (ICSC), IEEE Computer Society, Los Alamitos, CA, USA, 2022, pp. 229–234. doi:10.1109/ICSC52841.2022.00045.
 - [29] D. Diomedi, A. Hogan, Entity linking and filling for question answering over knowledge graphs, in: Proceedings of NLIWoD2022, 2022.

A. Online Resources

Source code for the training, organizational graph dataset and CoyPu dataset can be found at <https://github.com/AKSW/LMs4Text2SPARQL> and at DOI:10.5281/zenodo.10996425.