

# Ontology-grounded Automatic Knowledge Graph Construction by LLM under Wikidata schema

Xiaohan Feng<sup>1</sup>, Xixin Wu<sup>1</sup> and Helen Meng<sup>1,\*</sup>

<sup>1</sup>Department of System Engineering and Engineering Management, Chinese University of Hong Kong

## Abstract

We propose an ontology-grounded approach to Knowledge Graph (KG) construction using Large Language Models (LLMs) on a knowledge base. An ontology is authored by generating Competency Questions (CQ) on knowledge base to discover knowledge scope, extracting relations from CQs, and attempt to replace equivalent relations by their counterpart in Wikidata. To ensure consistency and interpretability in the resulting KG, we ground generation of KG with the authored ontology based on extracted relations. Evaluation on benchmark datasets demonstrates competitive performance in knowledge graph construction task. Our work presents a promising direction for scalable KG construction pipeline with minimal human intervention, that yields high quality and human-interpretable KGs, which are interoperable with Wikidata semantics for potential knowledge base expansion.

## Keywords

Knowledge Graph, Relation Extraction, Large Language Model, Wikidata, Interpretable AI

## 1. Introduction

Knowledge Graphs (KGs) are structured representations of information that capture entities and their relationships in a graph format. By organizing knowledge in a machine-readable way, KGs enable a wide range of intelligent applications, such as semantic search, question answering, recommendation systems, and decision support [1]. The ability to construct high-quality, comprehensive KGs is thus critical for harnessing the power of these technologies across various domains.

Traditionally, the process of constructing KGs has relied heavily on manual effort by domain experts to define the relevant entities and relationships, populate the graph with valid facts, and ensure logical consistency [2]. However, this manual curation approach is time-consuming, expensive, and difficult to scale to large, evolving domains. There is a strong need for (semi-)automatic methods that can aid the KG construction process by extracting structured knowledge from unstructured data sources such as text.

Recent years have seen growing interest in leveraging Large Language Models (LLMs) for various knowledge capture and reasoning tasks [3]. Pre-trained on vast amounts of text data, LLMs can generate fluent natural language and have been shown to memorize and recall factual knowledge [4], [5]. However, directly applying LLMs to KG construction still faces several challenges. First, LLMs may generate inconsistent or redundant facts due to the lack

---

HI-AI@KDD, Human-Interpretable AI Workshop at the KDD 2024, 26<sup>th</sup> of August 2024, Barcelona, Spain

\*Corresponding author.

✉ xhfeng@se.cuhk.edu.hk (X. Feng); wuxx@se.cuhk.edu.hk (X. Wu); hmmeng@se.cuhk.edu.hk (H. Meng)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

of an explicit, unified schema [6]. Second, the generated KGs may be incomplete or biased towards the knowledge present in the LLM’s training data, which may not fully cover the target domain, especially for proprietary documents not included in pre-training set. Finally, it can be challenging to integrate LLM-generated KGs with existing knowledge bases due to misalignment with standard ontologies.

In this work, we propose a novel approach that harnesses the reasoning power of LLMs and the structured schema of Wikidata to construct high-quality KGs for proprietary knowledge domains. Our approach begins by discovering the scope of knowledge through the generation of Competency Questions (CQ) and answers from unstructured documents. We then summarize the relations and properties from these QA pairs into an ontology, matching candidate properties against those defined in Wikidata and extending the schema as needed. Finally, we use the resulting ontology to ground the transformation of CQ-answer pairs into a structured KG. By incorporating the Wikidata schema into our pipeline and grounding generation of KG on the same ontology, we aim to reduce redundancy, leverage the implicit knowledge captured during LLM pretraining while improving interpretability, and ensure interoperability with public knowledge bases. The generated KGs could be parsed with RDF parsers and used in downstream applications, or audited for correctness.

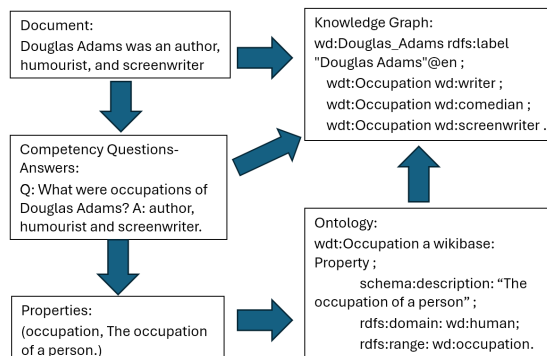


Figure 1: Flowchart of proposed approach.

The main contributions of this work are as follows:

1. We propose a novel ontology-grounded approach to LLM-based KG construction that leverages ontology based on Wikidata schema to guide the extraction and integration of knowledge from unstructured text.
2. We introduce a pipeline that combines competency question generation, ontology alignment, and KG grounding to systematically construct high-quality KGs that are consistent, complete, and interoperable with existing knowledge bases.
3. We demonstrate the effectiveness of our approach through experiments on benchmark datasets, showing improvements in KG quality compared to traditional methods alongside with interpretability and utility of generated KGs.

## 2. Literature Review

Knowledge graph construction has been an active area of research in recent years, with a wide range of approaches proposed for extracting structured knowledge from unstructured data sources [2]. Early methods relied heavily on rule-based systems and hand-crafted features to identify entities and relations in text [7]. With the advent of deep learning, neural network-based approaches have become increasingly popular, enabling more flexible and scalable KG construction [8].

One prominent line of work focuses on using distant supervision to automatically generate training data for relation extraction [9]. These methods assume that if two entities are mentioned together in a sentence and also appear in a knowledge base as subject and object of a relation, then that sentence is likely to express the relation. While distant supervision has been shown to be effective at scale, it often suffers from noise and incomplete coverage.

Another important direction is the development of unsupervised and semi-supervised methods for KG construction [10]. These approaches aim to reduce the reliance on large amounts of labeled data by leveraging techniques such as bootstrapping, graph-based inference, and representation learning. However, they often struggle with consistency and quality control issues. More recently, there has been growing interest in using large language models for KG construction [11], [12], [13]. These methods take advantage of the vast knowledge captured in pretrained Language Models (LM) to generate KG triples through prompt engineering and fine-tuning. While promising, these LM-based approaches only produce triplets without canonicalization, which makes portability and interoperability difficult. Additionally, some methods rely on vector-based similarity measures to deduce relationships between entities in KG, which yields good performance but falls short in interpretability [14].

As mentioned in Introduction, despite the significant progress in KG construction and LLM applications, performance, interpretability, coverage of proprietary documents, and interaction with other knowledge base remain issues. Our pipeline addresses these by grounding KG generation on ontology based on Wikidata schema, which ensures that output KG is human-readable and makes integrating with Wikidata or other KG easier; In the experiments below we show that these benefits can be also achieved on private documents with decent performance.

### 3. Method: Ontology-grounded KG Construction

Our proposed approach for ontology-grounded KG construction using LLMs consists of four main stages: 1) Competency Question Generation, 2) Relation Extraction and Ontology Matching, 3) Ontology Formatting, and 4) KG Construction. Figure 1 provides an overview of the pipeline.

#### 3.1. Competency Question (CQ)-Answer Generation

**The first step** in our pipeline is to generate a set of competency questions (CQs) and answers that capture the key information needs of the target domain. We employ an LLM to generate CQs based on the input documents. The LLM is provided with a set of instructions and examples to guide the generation process, encouraging the creation of well-formed, relevant questions that can be answered using the given documents. This step helps to scope the KG construction task within the knowledge domain, and ensure that the resulting KG aligns with the intended use cases. This also allows further ontology expansion by incorporating user-submitted domain-defining questions when interacting with the knowledge base, which serves as a user friendly interface of refining ontology by submitting new CQs and use our proposed pipeline to attach the incremental knowledge scope to existing ontology.

### 3.2. Relation Extraction and Ontology Matching

During our preliminary experiments of prompting LLMs to directly generate ontology on documents, we noted that the LLM spontaneously recalled Wikidata knowledge in response, consistent with previous works [15]. In our preliminary experiments, this behaviour also transfers to small 7B/14B models.

Following this direction, in **the second step** we extract relations from CQs and match them against Wikidata properties to better elicit model memories on Wikidata when constructing and using ontology. We first prompt LLMs to extract properties from CQ and write brief description on usage of extracted properties, including their domain and range, following editing guidelines of Wikidata. To match these properties against existing entries in Wikidata ontology, we pre-populate a candidate property list with all Wikidata properties after filtering out properties related to external database/knowledge base IDs. These extracted properties are then matched against the candidate list by a vector similarity search between description of properties. The representation for property is sentence embedding constructed from description of properties, and the top 1 closest candidate is retrieved for each extracted property. This matching result between each pair of extracted property and matched top 1 candidate is then vetted by LLM to see if they are really semantically similar as a final deduplication step. If a match is validated, the candidate property is added to the final property list; otherwise, the newly minted property is kept in the final list if we allow expansion from the candidate property list derived from Wikidata, and discarded when the final property list is required to be a subset of candidate property list. The first scenario is suitable for cases when no prior schema is known for the domain and some new properties outside of common ontology are expected, whereas the latter is for a known target list of possible properties.

### 3.3. Ontology Formatting

**In the third stage**, we use LLM to generate an OWL ontology based on the matched and newly created properties. We copy the description, domain and range field from all properties under Wikidata semantics. For new properties, LLM is prompted to infer and summarize classes for the domain and range of the relations to output a complete OWL ontology, following the format of copied Wikidata properties. This step ensures that the resulting KG is grounded in a formal, machine-readable ontology that captures relationships between entities, and close to the semantics of Wikidata for interoperability.

### 3.4. KG Construction

**In the final stage**, we use the LLM to construct a KG based on the CQs and related answers grounded by the generated ontology in the previous stage. For each (CQ, answer) pair, LLM extracts relevant entities and maps them to the ontology using the defined properties. The output is a set of RDF triples that constitutes the final KG.

## 4. Experiments and Discussion

### 4.1. Experiment settings

We evaluate our ontology-grounded approach to KG construction (KGC) on three datasets for KGC datasets: Wiki-NRE [16], SciERC [17], and WebNLG [18]. As Wiki-NRE and WebNLG are partially based on Wikidata and DBpedia (derived from Wikipedia contents), and in our proposed pipeline, Wikidata schema is utilized, we include SciERC for a more robust evaluation, since SciERC contains relation types that are not equivalent by nature to properties in Wikidata.

We used a subset Wiki-NRE’s test dataset containing 1,000 samples with 45 relation types following the split in [19], due to cost constraints. SciERC’s test set contain 974 samples under a schema with 7 relation types. For WebNLG, we used test set in Semantic Parsing (SP) task, with 1,165 samples and 159 relation types. For evaluation, we adopt partial F1 on KG triplets based on standards in [18]. All experiments are conducted for one-pass

We note in the previous reports that annotation in KGC reports may be incomplete in terms of both possible relation types and KG triplets [19], [20].

As our pipeline is designed to autonomously uncover knowledge structure with no prior assumption on knowledge schema, we report our result in two ways, corresponding to the two configurations of final de-duplication step in Section 3.2:

1. *Target schema constrained*: In this setting, we match all relation types in test sets to its closest equivalent in Wikidata and constrict ontology to the relation universe in test set.
2. *No schema constraint*: In this setting, we do not filter matched ontology, even if they are not in schema of test dataset. This setting is close to real-life applications when processing documents with unknown schema.

For property *conjunction, evaluate for, compare, feature of* in SciERC, we select the closest properties proposed by LLM based on our subjective opinion.

To highlight our system’s competency, rather than directly prompting triplets, we parse output KG with RDF parser and extract all valid RDF triples for KG related to each document in test set, and present triplets to evaluation script for assessment. This ensures that our evaluation is on the generated KG ready to be consumed in downstream application.

We test our pipeline on both Mistral-7B-instruct [21] and GPT-4o<sup>1</sup>. Due to cost constraints, we have only tested GPT-4o on *target schema constrained* setting. For embedding property usage comment, we select bge-small-en [22]. We use GenIE [23], PL-Marker [24], and ReGen [25] as fine-tuned baseline for Wiki-NRE, SciERC, and WebNLG dataset, respectively (collectively named **Non-LLM Baseline**). For LLM-based systems, we use results reported in [19] for Wiki-NRE and WebNLG on the same Mistral model, and GPT-4 results in [3] for SciERC. (collectively named **LLM Baseline**). We note that it is highly unlikely that Mistral-7B poses an advantage over an earlier version of GPT-4, when interpreting result of SciERC.

### 4.2. Result

---

<sup>1</sup><https://openai.com/index/hello-gpt-4o/>

Table 1 shows the performance of our method compared to state-of-the-art baselines on this subset. Our proposed approach exceeds all baseline under *target schema constrained* setting on Wiki-NRE and SciERC datasets, while displaying a small performance regression when without schema constraint. On WebNLG dataset, our pipeline maintained competitiveness against fine-tuned SOTA when constrained on target schema. These results validate the quality of KG generated by our pipeline, especially SciERC, whose semantics contains properties that are not native to Wikidata. We also note performance improvement when using GPT-4o.

Method	Wiki-NRE	SciERC	WebNLG
Non-LLM Baseline	0.484	0.532	<b>0.767</b>
LLM Baseline	0.647	0.07	0.728
Proposed (Mistral)	0.66/0.60	0.73/0.58	0.74/0.68
Proposed (GPT-4o)	<b>0.71</b> /N/A	<b>0.77</b> /N/A	0.76/N/A

Table 1: Partial F1 scores on test datasets. Best result is bolded. Results of proposed pipeline under two settings are presented as *Target schema constrained/no schema constraint*.

These results validate the quality of KG generated by our pipeline, especially SciERC, whose semantics contains properties that are not native to Wikidata. We also note performance improvement when using GPT-4o.

### 4.3. Discussion

#### 4.3.1. Performance discrepancy on different grounding ontology

It is worth noting that the relatively lower performance on *no schema constraint* setting across all datasets is due to the fact that the LLM discovers a richer ontology than the predefined target schema. While this expanded schema may capture additional relevant information, it can hinder extraction performance when evaluated solely against the limited target schema. This showcases the trade-off between schema completeness and strict adherence to a predefined ontology, and our pipeline performs best on a large set of documents with a limited scope of knowledge, requiring a concise schema.

Furthermore, the flipside of performance deficit in an absence of schema constraints, i.e. additional ontology entries outside of dataset-defined properties, cannot be evaluated against the dataset directly, as the ontology is not entirely covered by test set annotations. Hence, the virtue of *no schema constraint* setting is to demonstrate that our pipeline can indeed provide a coverage of the properties in test set, though somewhat limited compared to baselines, when also capturing ontology outside test set schema, which is potentially more useful when discovering ontology on a novel document set with no expert knowledge in its schema composition. This ability may be validated by manual evaluation on the full set of captured ontology in a future work down the line.

Nevertheless, the marginal performance deficit leaves room for improvement. Recent reports explored that long input context may pose challenge to LLMs even if such long context length is technically supported [26]. We conjecture that aside from trimming grounding ontology, which hinders the knowledge coverage of our pipeline, few-shot fine-tuning on the new ontology or general pretraining in KG construction task may be helpful. We leave these as possible future directions.

### 4.3.2. Utility of generated KG

It should be emphasized that, while the selected evaluation tasks evaluate the correctness of extracted triplets, the extracted knowledge graph can do more than that. With ongoing discussion related to grounding LLM knowledge on trusted knowledge sources to reduce hallucination [6], explicitly generating KG provides a path to audit knowledge elicited when interacting with LLM, and with evidence demonstrating that LLM has the potential to reason on graph and generate an explicit path to retrieve required knowledge [27], our pipeline may serve as a foundation for an interpretable QA system, where an LLM autonomously extracts ontology and deduces correct retrieval query based on the ontology when handling a set of unstructured document. The interpretability arise from the fact that KG and query could be understood and verified by users. Moreover, our usage of Wikidata schema offers potential interoperability with the whole Wikidata knowledge base, which safely expands the knowledge scope of QA system. We propose to continue research on this significant direction.

### 4.3.3. Computational resources

We note the growing concern of sustainability in LLM applications due to intensive requirement on computational resources. This pipeline consumes three separate LLM calls per document, plus one call per extracted relation. It is not straight forward to compare the carbon footprint of our approach compared to Non-LLM baselines, as our work at this stage does not require model fine-tuning, whereas all of the Non-LLM baselines employed various tuning techniques when producing the result. On the other hand, our smallest model adopted, Mistral-7B, is more than 10x larger in terms of parameter size compared to T5 models used in Non-LLM baselines. Larger models naturally require more powerful GPU clusters in terms of both GPU quantity and capability, but our zero-shot approach may provide an advantage in terms of resource cost compared to Non-LLM baselines when processing a small number of documents with no training requirement.

When comparing with LLM baselines, we note that the approach by [14], [3] consumes 1 and 2 LLM calls per document, respectively. However, we note that these baselines treat knowledge triplet as evaluation target, while we generate a formatted ontology at the end, which is more useful. Nevertheless, we recognize the performance burden and propose to explore techniques in fine-tuning and guided decoding to achieve better performance with smaller model and better reproducibility.

## 5. Conclusion

We have demonstrated the effectiveness of our ontology-grounded approach to KG construction using LLMs. By leveraging the structured knowledge in Wikidata, pretrained on LLM, and grounding KG construction with generated ontology, our pipeline is able to construct high-quality KGs across various domains while maintaining competitive performance with state-of-the-art baselines. Generated KGs that are conformant with Wikidata schema leaves possibly wide open, of building an interpretable QA system that has robust access to both common knowledge and proprietary knowledge base.

## Acknowledgments

This work is supported by Centre for Perceptual and Interactive Intelligence (CPII) Ltd, a CUHK-led InnoCentre under InnoHK scheme of Innovation and Technology Commission.

## References

- [1] A. Hogan, E. Blomqvist, M. Cochez, C. D’amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann, Knowledge graphs, *ACM Comput. Surv.* 54 (2021). URL: <https://doi.org/10.1145/3447772>. doi:10.1145/3447772.
- [2] S. Ji, S. Pan, E. Cambria, P. Marttinen, P. S. Yu, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE Transactions on Neural Networks and Learning Systems* 33 (2020) 494–514. URL: <https://api.semanticscholar.org/CorpusID:211010433>.
- [3] Y. Zhu, X. Wang, J. Chen, S. Qiao, Y. Ou, Y. Yao, S. Deng, H. Chen, N. Zhang, LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities, 2024. doi:10.48550/arXiv.2305.13168. arXiv:2305.13168.
- [4] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, S. Riedel, Language models as knowledge bases?, in: *Conference on Empirical Methods in Natural Language Processing*, 2019. URL: <https://api.semanticscholar.org/CorpusID:202539551>.
- [5] OpenAI, J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altschmidt, S. Altman, S. Anadkat, R. Avila, I. Babuschkin, S. Balaji, V. Balcom, P. Baltescu, H. Bao, M. Bavarian, J. Belgum, I. Bello, J. Berdine, G. Bernadett-Shapiro, C. Berner, L. Bogdonoff, O. Boiko, M. Boyd, A.-L. Brakman, G. Brockman, T. Brooks, M. Brundage, K. Button, T. Cai, R. Campbell, A. Cann, B. Carey, C. Carlson, R. Carmichael, B. Chan, C. Chang, F. Chantzis, D. Chen, S. Chen, R. Chen, J. Chen, M. Chen, B. Chess, C. Cho, C. Chu, H. W. Chung, D. Cummings, J. Currier, Y. Dai, C. Decareaux, T. Degry, N. Deutsch, D. Deville, A. Dhar, D. Dohan, S. Dowling, S. Dunning, A. Ecoffet, A. Eleti, T. Eloundou, D. Farhi, L. Fedus, N. Felix, S. P. Fishman, J. Forte, I. Fulford, L. Gao, E. Georges, C. Gibson, V. Goel, T. Gogineni, G. Goh, R. Gontijo-Lopes, J. Gordon, M. Grafstein, S. Gray, R. Greene, J. Gross, S. S. Gu, Y. Guo, C. Hallacy, J. Han, J. Harris, Y. He, M. Heaton, J. Heidecke, C. Hesse, A. Hickey, W. Hickey, P. Hoeschele, B. Houghton, K. Hsu, S. Hu, X. Hu, J. Huizinga, S. Jain, S. Jain, J. Jang, A. Jiang, R. Jiang, H. Jin, D. Jin, S. Jomoto, B. Jonn, H. Jun, T. Kaftan, Łukasz Kaiser, A. Kamali, I. Kanitscheider, N. S. Keskar, T. Khan, L. Kilpatrick, J. W. Kim, C. Kim, Y. Kim, J. H. Kirchner, J. Kiros, M. Knight, D. Kokotajlo, Łukasz Kondraciuk, A. Kondrich, A. Konstantinidis, K. Kosic, G. Krueger, V. Kuo, M. Lampe, I. Lan, T. Lee, J. Leike, J. Leung, D. Levy, C. M. Li, R. Lim, M. Lin, S. Lin, M. Litwin, T. Lopez, R. Lowe, P. Lue, A. Makanju, K. Malfacini, S. Manning, T. Markov, Y. Markovski, B. Martin, K. Mayer, A. Mayne, B. McGrew, S. M. McKinney, C. McLeavey, P. McMillan, J. McNeil, D. Medina, A. Mehta, J. Menick, L. Metz, A. Mishchenko, P. Mishkin, V. Monaco, E. Morikawa, D. Mossing, T. Mu, M. Murati, O. Murk, D. Mély, A. Nair, R. Nakano, R. Nayak, A. Neelakantan, R. Ngo, H. Noh, L. Ouyang, C. O’Keefe, J. Pachocki, A. Paino, J. Palermo, A. Pantuliano, G. Parascandolo, J. Parish, E. Parparita, A. Passos, M. Pavlov, A. Peng,



- A. Perelman, F. de Avila Belbute Peres, M. Petrov, H. P. de Oliveira Pinto, Michael, Pokorny, M. Pokrass, V. H. Pong, T. Powell, A. Power, B. Power, E. Proehl, R. Puri, A. Radford, J. Rae, A. Ramesh, C. Raymond, F. Real, K. Rimbach, C. Ross, B. Rotsted, H. Roussez, N. Ryder, M. Saltarelli, T. Sanders, S. Santurkar, G. Sastry, H. Schmidt, D. Schnurr, J. Schulman, D. Selsam, K. Sheppard, T. Sherbakov, J. Shieh, S. Shoker, P. Shyam, S. Sidor, E. Sigler, M. Simens, J. Sitkin, K. Slama, I. Sohl, B. Sokolowsky, Y. Song, N. Staudacher, F. P. Such, N. Summers, I. Sutskever, J. Tang, N. Tezak, M. B. Thompson, P. Tillet, A. Tootoonchian, E. Tseng, P. Tuggle, N. Turley, J. Tworek, J. F. C. Uribe, A. Vallone, A. Vijayvergiya, C. Voss, C. Wainwright, J. J. Wang, A. Wang, B. Wang, J. Ward, J. Wei, C. Weinmann, A. Welihinda, P. Welinder, J. Weng, L. Weng, M. Wiethoff, D. Willner, C. Winter, S. Wolrich, H. Wong, L. Workman, S. Wu, J. Wu, M. Wu, K. Xiao, T. Xu, S. Yoo, K. Yu, Q. Yuan, W. Zaremba, R. Zellers, C. Zhang, M. Zhang, S. Zhao, T. Zheng, J. Zhuang, W. Zhuk, B. Zoph, Gpt-4 technical report, 2024. [arXiv:2303.08774](https://arxiv.org/abs/2303.08774).
- [6] G. Agrawal, T. Kumarage, Z. Alghamdi, H. Liu, Can Knowledge Graphs Reduce Hallucinations in LLMs? : A Survey, 2024. [doi:10.48550/arXiv.2311.07914](https://doi.org/10.48550/arXiv.2311.07914). [arXiv:2311.07914](https://arxiv.org/abs/2311.07914).
- [7] E. Agichtein, L. Gravano, Snowball: extracting relations from large plain-text collections, in: Proceedings of the Fifth ACM Conference on Digital Libraries, DL '00, Association for Computing Machinery, New York, NY, USA, 2000, p. 85–94. URL: <https://doi.org/10.1145/336597.336644>. [doi:10.1145/336597.336644](https://doi.org/10.1145/336597.336644).
- [8] Y. Zhang, V. Zhong, D. Chen, G. Angeli, C. D. Manning, Position-aware attention and supervised data improve slot filling, in: M. Palmer, R. Hwa, S. Riedel (Eds.), Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Copenhagen, Denmark, 2017, pp. 35–45. URL: <https://aclanthology.org/D17-1004>. [doi:10.18653/v1/D17-1004](https://doi.org/10.18653/v1/D17-1004).
- [9] M. Mintz, S. Bills, R. Snow, D. Jurafsky, Distant supervision for relation extraction without labeled data, in: K.-Y. Su, J. Su, J. Wiebe, H. Li (Eds.), Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, Association for Computational Linguistics, Suntec, Singapore, 2009, pp. 1003–1011. URL: <https://aclanthology.org/P09-1113>.
- [10] X. L. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, W. Zhang, Knowledge vault: A web-scale approach to probabilistic knowledge fusion, in: The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014, 2014, pp. 601–610. URL: <http://www.cs.cmu.edu/~nlao/publication/2014.kdd.pdf>, evgeniy GabrilovichWilko HornNi LaoKevin MurphyThomas StrohmmanShaohua SunWei ZhangJeremy Heitz.
- [11] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Celikyilmaz, Y. Choi, COMET: Commonsense transformers for automatic knowledge graph construction, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 4762–4779. URL: <https://aclanthology.org/P19-1470>. [doi:10.18653/v1/P19-1470](https://doi.org/10.18653/v1/P19-1470).
- [12] Y. Chen, Y. Liu, L. Dong, S. Wang, C. Zhu, M. Zeng, Y. Zhang, AdaPrompt: Adaptive model training for prompt-based NLP, in: Y. Goldberg, Z. Kozareva, Y. Zhang (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2022, Associ-

- ation for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2022, pp. 6057–6068. URL: <https://aclanthology.org/2022.findings-emnlp.448>. doi:10.18653/v1/2022.findings-emnlp.448.
- [13] S. Yu, T. He, J. Glass, AutoKG: Constructing Virtual Knowledge Graphs from Unstructured Documents for Question Answering, 2021. doi:10.48550/arXiv.2008.08995. arXiv:2008.08995.
- [14] B. Chen, A. L. Bertozzi, AutoKG: Efficient Automated Knowledge Graph Generation for Language Models, 2023. doi:10.48550/arXiv.2311.14740. arXiv:2311.14740.
- [15] S. J. Semnani, V. Z. Yao, H. C. Zhang, M. S. Lam, WikiChat: Stopping the Hallucination of Large Language Model Chatbots by Few-Shot Grounding on Wikipedia, 2023. arXiv:2305.14292.
- [16] B. D. Trisedya, G. Weikum, J. Qi, R. Zhang, Neural relation extraction for knowledge base enrichment, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 229–240. URL: <https://aclanthology.org/P19-1023>. doi:10.18653/v1/P19-1023.
- [17] Y. Luan, L. He, M. Ostendorf, H. Hajishirzi, Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction, in: Proc. Conf. Empirical Methods Natural Language Process. (EMNLP), 2018.
- [18] T. Castro Ferreira, C. Gardent, N. Ilinykh, C. van der Lee, S. Mille, D. Moussallem, A. Shimorina, The 2020 bilingual, bi-directional WebNLG+ shared task: Overview and evaluation results (WebNLG+ 2020), in: T. Castro Ferreira, C. Gardent, N. Ilinykh, C. van der Lee, S. Mille, D. Moussallem, A. Shimorina (Eds.), Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+), Association for Computational Linguistics, Dublin, Ireland (Virtual), 2020, pp. 55–76. URL: <https://aclanthology.org/2020.webnlg-1.7>.
- [19] B. Zhang, H. Soh, Extract, Define, Canonicalize: An LLM-based Framework for Knowledge Graph Construction, 2024. arXiv:2404.03868.
- [20] R. Han, T. Peng, C. Yang, B. Wang, L. Liu, X. Wan, Is information extraction solved by chatgpt? an analysis of performance, evaluation criteria, robustness and errors, 2023. arXiv:2305.14450.
- [21] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mistral 7b, 2023. arXiv:2310.06825.
- [22] J. Chen, S. Xiao, P. Zhang, K. Luo, D. Lian, Z. Liu, Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation, 2023. arXiv:2309.07597.
- [23] M. Josifoski, N. De Cao, M. Peyrard, F. Petroni, R. West, GenIE: Generative information extraction, in: M. Carpuat, M.-C. de Marneffe, I. V. Meza Ruiz (Eds.), Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Seattle, United States, 2022, pp. 4626–4643. URL: <https://aclanthology.org/2022.naacl-main.342>. doi:10.18653/v1/2022.naacl-main.342.
- [24] D. Ye, Y. Lin, P. Li, M. Sun, Packed levitated marker for entity and relation extraction, in:

- S. Muresan, P. Nakov, A. Villavicencio (Eds.), Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 4904–4917. URL: <https://aclanthology.org/2022.acl-long.337>. doi:10.18653/v1/2022.acl-long.337.
- [25] P. Dognin, I. Padhi, I. Melnyk, P. Das, ReGen: Reinforcement learning for text and knowledge base generation using pretrained language models, in: M.-F. Moens, X. Huang, L. Specia, S. W.-t. Yih (Eds.), Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 1084–1099. URL: <https://aclanthology.org/2021.emnlp-main.83>. doi:10.18653/v1/2021.emnlp-main.83.
- [26] T. Li, G. Zhang, Q. D. Do, X. Yue, W. Chen, Long-context llms struggle with long in-context learning, 2024. [arXiv:2404.02060](https://arxiv.org/abs/2404.02060).
- [27] F. Brei, J. Frey, L.-P. Meyer, Leveraging small language models for text2sparql tasks to improve the resilience of ai assistance, 2024. [arXiv:2405.17076](https://arxiv.org/abs/2405.17076).

## A. Sample generated KG

This KG was generated under *no schema constraint* setting for this document: Mohammad Firouzi ( Born 1958 Tehran ) is a prolific Iranian musician , whose primary instrument is the barbat .

```
<Prefixes and definition of dependencies omitted>
wd: Mohammad_Firouzi a wd:human ;
    rdfs:label "Mohammad Firouzi"@en ;
    wdt:occupation wd:Musician ;
    wdt:CountryOfCitizenship wd:Iran ;
    wdt:PlaceOfBirth wd:Tehran ;
    wdt:DateOfBirth "1958"^^xsd:date .
```

Note that in official annotation, only triplets related to place of birth and nationality exist, hence the evaluation will be penalized with low precision.

## B. Preprocessing of Wikidata schema

To save space in LLM input context and mitigate performance drop on selected target schema when ontology is large, we only include commonly used properties by restricting data type on *item*, *quantity*, *string*, *monolingual text*, *point in time*.<sup>2</sup> To align with common pretraining objectives of LLM, we substitute entity identifiers (e.g. *P19*) with its literal label (*rdfs:label* in PascalCase (e.g. *PlaceOfBirth*)).

---

<sup>2</sup>[https://www.wikidata.org/wiki/Help:Data\\_type](https://www.wikidata.org/wiki/Help:Data_type)

## C. Prompts

All prompts are reused across all datasets.

### C.1. CQ generation

We prompt LLM to generate up to 3 CQs per document for efficiency considering nature of test datasets, but note that this may be adjusted.

Write competency questions based on the abstract level concepts in the document. Write questions that can be answered using the document only.

Write up to 3 questions per document.

Below are the examples and follow the same format when generating competency questions:

####

Document: Douglas Noel Adams (11 March 1952 – 11 May 2001) was an English author, humourist, and screenwriter, best known for The Hitchhiker's Guide to the Galaxy (HHGTTG). Originally a 1978 BBC radio comedy, The Hitchhiker's Guide to the Galaxy developed into a "trilogy" of five books that sold more than 15 million copies in his lifetime. It was further developed into a television series, several stage plays, comics, a video game, and a 2005 feature film. Adams's contribution to UK radio is commemorated in The Radio Academy's Hall of Fame.

####

Questions:

- CQ1. What is the date of birth of Douglas Noel Adams?
- CQ2. What is the date of death of Douglas Noel Adams?
- CQ3. What is the occupation of Douglas Noel Adams?
- CQ4. What is the country of citizenship of Douglas Noel Adams?
- CQ5. What is the most notable work of Douglas Noel Adams?
- CQ6. What is the original medium of The Hitchhiker's Guide to the Galaxy?
- CQ7. In what year was The Hitchhiker's Guide to the Galaxy originally broadcast?
- CQ8. How many books are in The Hitchhiker's Guide to the Galaxy "trilogy"?
- CQ9. What other media adaptations were created based on The Hitchhiker's Guide to the Galaxy?

####

Document:  
{document to be processed}

####

Questions:

## C.2. CQ answering

Use the provided document to answer user query. If you don't know the answer, just say that you don't know, don't try to make up an answer.

Passage: {doc}

Query: {query}

## C.3. Relation extraction

You are an assistant in building a knowledge graph. Analyze the following competency questions and identify all relationships and concepts mentioned in the question.

Extract relation first, then describe the usage of each relation based on your understanding given the context of competency questions.

Afterwards, extract all relation-related concepts.

You should only extract properties between entities and literals, not entities themselves, or classes of entities.

Therefore, not all CQs contain valid properties.

If you don't know the answer, just say that you don't know, don't try to make up an answer.

Merge all relations into one list and all concepts into one list.

Do not reply using a complete sentence, and only give the answer in the following format.

Below are the examples and follow the same format to extract the relations:

####

Document: Douglas Noel Adams (11 March 1952 - 11 May 2001) was an English author, humourist, and screenwriter, best known for The Hitchhiker's Guide to the Galaxy (HHGTTG).

Originally a 1978 BBC radio comedy, The Hitchhiker's Guide to the Galaxy developed into a "trilogy" of five books that

sold more than 15 million copies in his lifetime. It was further developed into a television series, several stage plays, comics, a video game, and a 2005 feature film. Adams' s contribution to UK radio is commemorated in The Radio Academy's Hall of Fame.

####

Questions:

- CQ1. What is the date of birth of Douglas Noel Adams?
- CQ2. What is the date of death of Douglas Noel Adams?
- CQ3. What is the occupation of Douglas Noel Adams?
- CQ4. What is the country of citizenship of Douglas Noel Adams?
- CQ5. What is the most notable work of Douglas Noel Adams?
- CQ6. What is the original medium of The Hitchhiker's Guide to the Galaxy?
- CQ7. In what year was The Hitchhiker's Guide to the Galaxy originally broadcast?
- CQ8. How many books are in The Hitchhiker's Guide to the Galaxy "trilogy"?
- CQ9. What other media adaptations were created based on The Hitchhiker's Guide to the Galaxy?

####

Relations:

- (date of birth, The date on which the subject was born.)
- (date of death, The date on which the subject died.)
- (occupation, The occupation of a person.)
- (country of citizenship, The country of which the subject is a citizen.)
- (notable work, The most notable work of a person.)
- (genre, The genre or type of work.)
- (publication date, The date or period when a work was first published or released.)
- (has part, Indicates that the subject has a certain part, component, or element.)
- (series, Indicates that the subject is part of a series, such as a book series, film series, or television series.)

####

Document:

{document to be processed}

####

Questions:

{CQs}

####

Relations:

#### C.4. Ontology matching

Decide if the two properties are semantically similar in an ontology.

You should say yes if you decide that these propties are similar, or if they are inverse properties.

Answer in "yes" or "no" only.

Property 1: {p1}

Property 2: {p2}

#### C.5. Ontology formatting

For properties under Wikidata schema, we retrieve *schema:description*, *rdfs:domain*, *rdfs:range* for each property and include it in resulting ontology. Otherwise LLM is prompted to author ontology as so:

Use the relations (properties) and their usage comments to build an ontology in RDF format.

If you don't know the answer, just say that you don't know, don't try to make up an answer.

Don't provide anything other than an ontology in RDF format.

Infer and summarize classes for domain and range of the relations across the concepts provided, and add these classes to relations only if required for clousre of relations.

For each relation, add relevant ontology entry for it.

Add *rdfs:comment* based on the usage comments.

Use *wdt:* namespace for all relations discovered. Use entities under these prefixes if necessary:

@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .

@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

@prefix owl: <http://www.w3.org/2002/07/owl#> .

@prefix wikibase: <http://wikiba.se/ontology#> .

@prefix schema: <http://schema.org/> .

@prefix wd: <http://www.wikidata.org/entity/> .

@prefix wdt: <http://www.wikidata.org/prop/direct/> .

Use turtle syntax.

Below is an example:

```
####
Relations:
(results , results: results of a competition such as sports or
elections)

####
Ontology:
wdt:Results a wikibase:Property ;
    schema:description "results of a competition such as sports
or elections" ;
    rdfs:label "results" ;
    rdfs:domain wd:referendum , wd:competition , wd:party
conference , wd:sporting event ;
    rdfs:range wd:electoral result , wd:voting result , wd:sport
result , wd:race result .

####
Relations:
{relation}

####
Ontology:
```

## C.6. KG generation

Your task is to construct a knowledge graph based on the provided ontology.  
Focus on understanding relationships from the question answer pair and document,  
and extract related entities , then mapping them to the ontology using the properties defined in the ontology.  
Do not include new properties other than those in ontology.  
Only use those properties in the ontology.  
Output in turtle format following the ontology provided.  
You should only include knowledge in question answer pairs and the document.  
Do not make up answers.

Use this ontology based on Wikidata as the starting point:  
{ont}



Below is an example:

####

Document:

Douglas Noel Adams (11 March 1952 - 11 May 2001) was an English author, humourist, and screenwriter, best known for The Hitchhiker's Guide to the Galaxy (HHGTTG). Originally a 1978 BBC radio comedy, The Hitchhiker's Guide to the Galaxy developed into a "trilogy" of five books that sold more than 15 million copies in his lifetime. It was further developed into a television series, several stage plays, comics, a video game, and a 2005 feature film. Adams's contribution to UK radio is commemorated in The Radio Academy's Hall of Fame.

####

Question answer pairs:

Q: What is Douglas Adams an instance of?

A: Douglas Adams is an instance of human.

Q: What is Douglas Adams' sex or gender?

A: Douglas Adams' sex or gender is male.

Q: Where was Douglas Adams born?

A: Douglas Adams was born in Cambridge.

Q: Where did Douglas Adams die?

A: Douglas Adams died in Santa Barbara, California.

Q: When was Douglas Adams born?

A: Douglas Adams was born on 1952-03-11.

Q: On what date did Douglas Adams die?

A: Douglas Adams died on 2001-05-11.

Q: What occupation did Douglas Adams have?

A: Douglas Adams was a writer, comedian, and dramatist.

Q: What languages did Douglas Adams speak, write, or sign?

A: Douglas Adams spoke, wrote, or signed English.

Q: Where was Douglas Adams educated?

A: Douglas Adams was educated at St John's College, Cambridge

and Brentwood School , Essex .

Q: What institution is Douglas Adams an alumni of?

A: Douglas Adams is an alumni of St John 's College .

Q: What are some notable works by Douglas Adams?

A: Some notable works by Douglas Adams include The Hitchhiker 's Guide to the Galaxy and Dirk Gently 's Holistic Detective Agency .

Q: Was Douglas Adams a member of any notable organizations?

A: Yes , Douglas Adams was a member of Monty Python and The Independent on Sunday .

Q: What award did Douglas Adams receive?

A: Douglas Adams received the Locus Award for Best Science Fiction Novel .

Q: What is the Commons Category for Douglas Adams?

A: The Commons Category for Douglas Adams is "Douglas Adams" .

####

Ontology :

```
@prefix rdf: http://www.w3.org/1999/02/22-rdf-syntax-ns# .
@prefix rdfs: http://www.w3.org/2000/01/rdf-schema# .
@prefix wdt: http://www.wikidata.org/prop/direct/ .
@prefix wd: http://www.wikidata.org/entity/ .
@prefix xsd: http://www.w3.org/2001/XMLSchema# .
wd:Douglas_Adams rdfs:label "Douglas Adams"@en ;
    wdt:InstanceOf wd:human ;
    wdt:SexOrGender wd:male ;
    wdt:PlaceOfBirth wd:Cambridge ;
    wdt:PlaceOfDeath wd:Santa_Barbara_California ;
    wdt:DateOfBirth "1952-03-11"^^xsd:date ;
    wdt:DateOfDeath "2001-05-11"^^xsd:date ;
    wdt:Occupation wd:writer ;
    wdt:Occupation wd:comedian ;
    wdt:Occupation wd:dramatist ;
    wdt:LanguagesSpokenWrittenOrSigned wd:English ;
    wdt:EducatedAt wd:St_Johns_College_Cambridge ;
    wdt:EducatedAt wd:Brentwood_School_Essex ;
    wdt:AlumniOf wd:St_Johns_College ;
    wdt:NotableWork wd:The_Hitchhikers_Guide_to_the_Galaxy ;
    wdt:NotableWork wd:Dirk_Gentlys_Holistic_Detective_Agency ;
```

wdt:MemberOf wd:Monty\_Python ;  
wdt:MemberOfOrganization wd:The\_Independent\_on\_Sunday ;  
wdt:Award wd:Locus\_Award\_for\_Best\_Science\_Fiction\_Novel ;  
wdt:CommonsCategory "Douglas Adams"@en .

####

Document:

{doc}

####

Questions and Answer pairs:

{qa}

####

Ontology: