Chair of Software Engineering for Business Information Systems
School of Computation, Information and Technology
Department of Computer Science
Technical University of Munich

A Decade of Knowledge Graphs in Natural Language Processing: A Survey

Phillip Schneider, Tim Schopf, Juraj Vladika, Mikhail Galkin, Elena Simperl and Florian Matthes

{phillip.schneider, tim.schopf, juraj.vladika, matthes}@tum.de, mikhail.galkin@mila.quebec, elena.simperl@kcl.ac.uk

Motivation

- In recent years, knowledge graphs (KGs) have attracted a surge of interest from both academia and industry.
- As a representation of semantic relations between entities, KGs have proven to be particularly relevant for natural language processing (NLP).
- Despite the increasing amount of research work in this area, a comprehensive study that categorizes established topics and reviews emerging research streams remains absent to this day.
- To provide an overview of the research landscape, we analyzed 507 papers in a multifaceted survey of tasks, research types, and contributions.

Method

- As research method, we chose to conduct a systematic mapping study according to the procedure defined by Petersen et al. [1].
- For our search of relevant publications, we queried six academic databases using the following search string: *("knowledge graph") AND ("NLP" OR "natural language processing" OR "computational linguistics").* The search time frame spanned 10 years from 2012 to 2021.
- After screening the initial dataset of 746 papers, we included 507 relevant papers and assigned labels for task, research type, and contribution type according to specified classification schemes [2, 3].

Results

- There was a major rise in research interest from 2017 onwards, as over 90% of all included papers were published in these 5 years.
- The diversity of studied application domains grew rapidly in parallel to the annual count of papers with 20 identified domains (e.g., health, scholarly research, engineering, business, or social media).
- Figure 1 shows the empirically developed taxonomy of tasks and Table 1 lists the 5 most popular tasks in the literature.

Task	No. of Papers	No. of Domains
Relation extraction	144	18
Entity extraction	143	19
Question answering	103	14
Semantic search	91	16
Augmented language models	84	4

Table 1: Overview of most popular tasks in the literature on KGs in NLP.

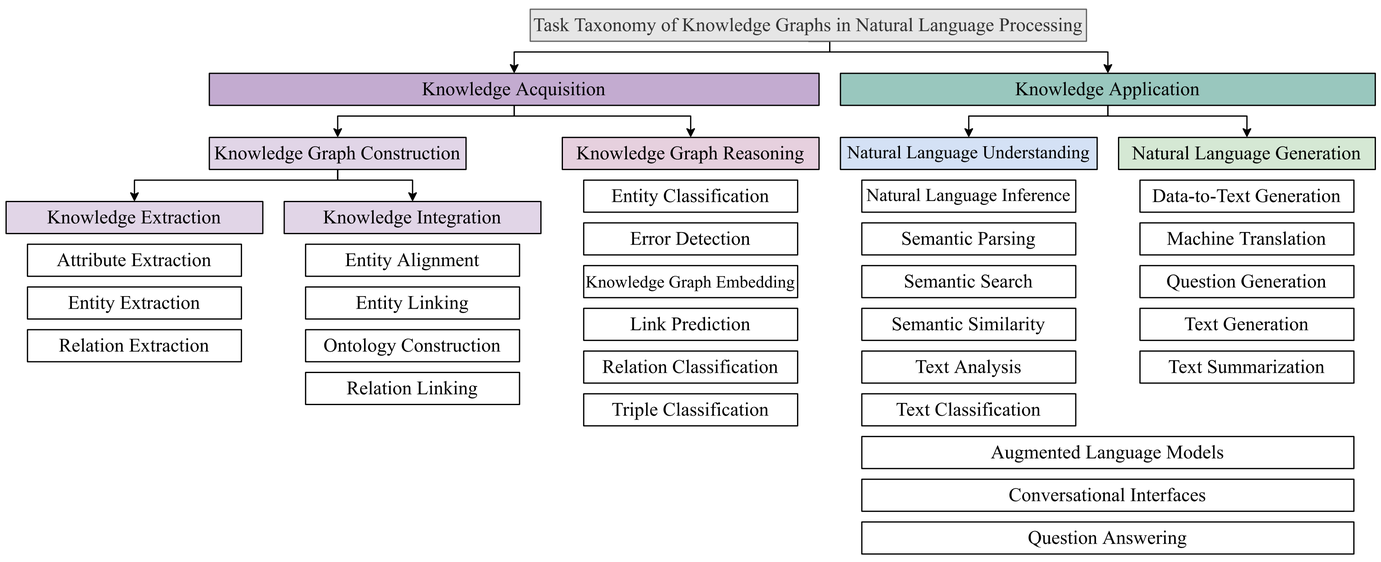

Figure 1: Taxonomy of tasks in the literature on KGs in NLP.

- According to Table 2, there is a notable lack of secondary studies that consolidate existing research. Moreover, evaluation research papers that implement approaches in an industry context are equally scarce.
- Figure 2 illustrates that long-established tasks, such as relation extraction or semantic search, have a balanced ratio of diverse contribution types, indicating that they are reasonably mature, whereas novel topics like augmented language models or KG embedding have mostly techniques as contributions.

Research Type	No. of Papers
Validation research	338
Solution proposal	149
Secondary research	10
Evaluation research	7
Opinion paper	3

Contribution Type	No. of Papers
Technique	186
Method	154
Tool	139
Resource	50
Guidelines	24

Table 2: Number of papers by research type [3] and contribution type [2].

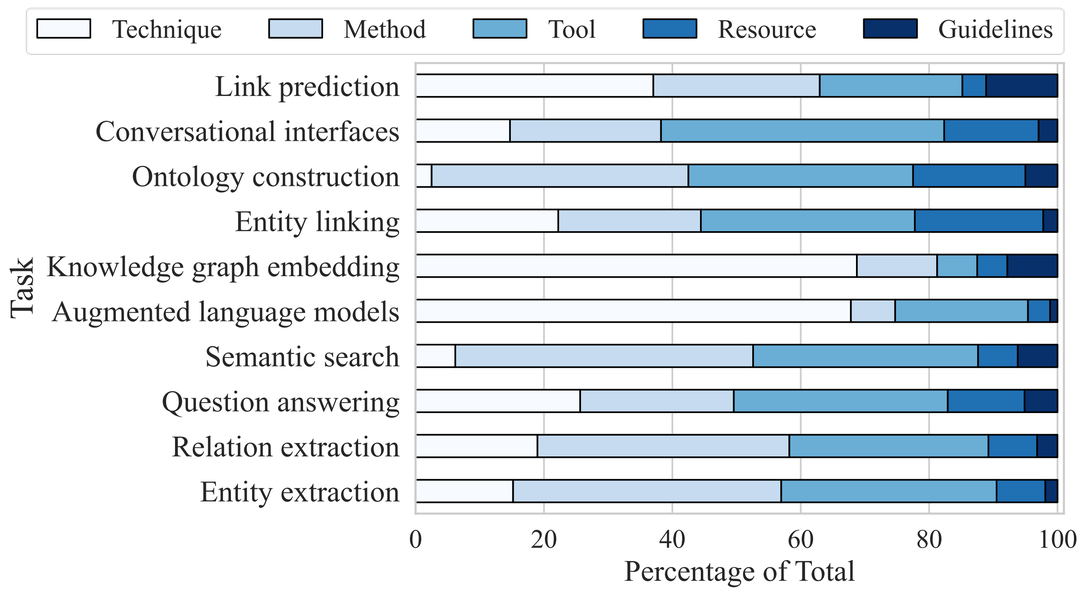

Figure 2: Percentage of contribution type by tasks.

Conclusion

- Our findings show a rising prominence of KGs in NLP research. A large number of tasks in the literature have been studied across various application domains.
- Aside from established topics like ontology construction and semantic search, there are emerging topics like augmented language models, KG embedding, or conversational interfaces.
- We also observed a lack of secondary research and evaluations in practice, both of which are crucial to reflect the scientific progress of the field as a whole. Our study lays the grounds for further research in this direction.

Data repository with annotated collection of 507 included papers.

github.com/sebischair/KG-in-NLP-survey

Technical University of Munich
School of Computation, Information and Technology
Department of Computer Science
Chair of Software Engineering for Business Information Systems (sebis)

Acknowledgements
This work has been supported by the German Federal Ministry of Education and Research (BMBF) Software Campus grant 01IS17049. The fourth author is partially supported by the Canada CIFAR AI Chair Program and Samsung AI grant (held at Mila).

References
[1] Petersen, Kai, et al. "Systematic mapping studies in software engineering." *12th International Conference on Evaluation and Assessment in Software Engineering (EASE)*. 2008.
[2] Shaw, Mary. "Writing good software engineering research papers." *25th International Conference on Software Engineering*. 2003.
[3] Wieringa, Roel, et al. "Requirements engineering paper classification and evaluation criteria: a proposal and a discussion." *Requirements engineering*. 2006.