



King's Research Portal

Document Version

Version created as part of publication process; publisher's layout; not normally made publicly available

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Simperl, E. (Accepted/In press). A Decade of Knowledge Graphs in Natural Language Processing: A Survey. In *2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

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

Phillip Schneider¹, Tim Schopf¹, Juraj Vladika¹, Mikhail Galkin²,
Elena Simperl³ and Florian Matthes¹

¹Technical University of Munich, Department of Informatics, Germany

²Mila Quebec AI Institute & McGill University, School of Computer Science, Canada

³King’s College London, Department of Informatics, United Kingdom

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

Abstract

In pace with developments in the research field of artificial intelligence, 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), experiencing a rapid spread and wide adoption within recent years. Given the increasing amount of research work in this area, several KG-related approaches have been surveyed in the NLP research community. However, a comprehensive study that categorizes established topics and reviews the maturity of individual research streams remains absent to this day. Contributing to closing this gap, we systematically analyzed 507 papers from the literature on KGs in NLP. Our survey encompasses a multifaceted review of tasks, research types, and contributions. As a result, we present a structured overview of the research landscape, provide a taxonomy of tasks, summarize our findings, and highlight directions for future work.

1 Introduction

Knowledge acquisition and application are inherent to natural language. Humans use language as a means of communicating facts, arguing about decisions, or questioning beliefs. Therefore, it is not surprising that computational linguists started already in the 1950s and 60s to work out ideas on how to represent knowledge as relations between concepts in semantic networks (Richens, 1956; Quillian, 1963; Collins and Quillian, 1969).

More recently, knowledge graphs (KGs) have emerged as an approach for semantically representing knowledge about real-world entities in a machine-readable format. They originated from research on semantic networks, domain-specific ontologies, as well as linked data, and are thus not an entirely new concept (Hitzler, 2021). Despite

their growing popularity, there is still no general understanding of what exactly a KG is or for what tasks it is applicable. Although prior work has already attempted to define KGs (Pujara et al., 2013; Ehrlinger and Wöß, 2016; Paulheim, 2017; Färber et al., 2018), the term is not yet used uniformly by researchers. Most studies implicitly adopt a broad definition of KGs, where they are understood as *“a graph of data intended to accumulate and convey knowledge of the real world, whose nodes represent entities of interest and whose edges represent relations between these entities”* (Hogan et al., 2022).

KGs have attracted a lot of research attention in both academia and industry since the introduction of Google’s KG in 2012 (Singhal, 2012). Particularly in natural language processing (NLP) research, the adoption of KGs has become increasingly popular over the past 5 years, and this trend seems to be accelerating. The underlying paradigm is that the combination of structured and unstructured knowledge can benefit all kinds of NLP tasks. For instance, structured knowledge from KGs can be injected into that of the contextual knowledge found in language models, which improves the performance in downstream tasks (Colon-Hernandez et al., 2021). Furthermore, with the growing importance of KGs, there are also expanding efforts to construct new KGs from unstructured texts.

Ten years after Google coined the term knowledge graph in 2012, a plethora of novel approaches has been proposed by scholars. Therefore, it is important to assemble insights, consolidate existing results, and provide a structured overview. However, to our knowledge, there are no studies that offer an overview of the whole research landscape of KGs in the NLP field. Contributing to closing this gap, we performed a comprehensive survey to analyze all research performed in this area by classifying established topics, identifying trends, and outlining areas for future research. Our three main contributions are as follows:

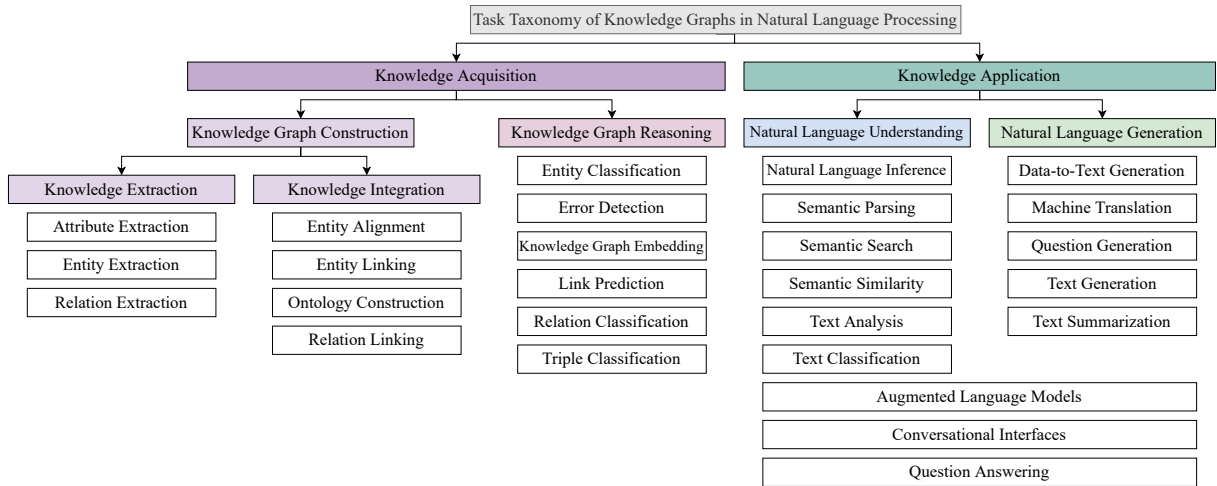


Figure 1: Taxonomy of tasks in the literature on KGs in NLP.

1. We systematically extract information from 507 included papers and report insights about tasks, research types, and contributions.
2. We provide a taxonomy of tasks in the literature on KGs in NLP shown in Figure 1.
3. We assess the maturity of individual research streams, identify trends, and highlight directions for future work.

Our survey sheds light on the evolution and current research progress regarding KGs in NLP. Although we cannot achieve complete coverage of all relevant papers on this topic, we aim at providing a representative overview that can help both NLP scholars and practitioners by offering a starting point in the literature. Moreover, our multifaceted analysis can guide the research community in closing existing gaps and finding novel ways how to combine KGs with NLP.

2 Related Work

Related literature that includes both KGs and NLP seems to be relatively scarce. Most survey papers focus either only on KGs or only on NLP. In their broad introduction to KGs, Hogan et al. (2022) point out that existing surveys on KGs tend to revolve around specific aspects of KGs, most commonly their construction and embedding.

Such surveys with a KG focus usually bring up NLP only in the context of employed NLP methods, like information extraction, being used to populate and refine graphs (Nickel et al., 2016). Other surveys on KGs mention some downstream applications of KGs for NLP tasks, such as for con-

structing augmented language models, question answering over knowledge bases (KBQA), or recommender systems (Ji et al., 2021).

As noted previously, related work that includes both KGs and NLP strictly focus on a specific application or task. For example, Safavi and Koutra (2021) provide an overview on applying relational world knowledge from KGs to augment large contextual language models. Other surveys on specific applications include KG reasoning (Chen et al., 2019), biomedical KGs (Nicholson and Greene, 2020), and the task of KBQA (Fu et al., 2020).

The survey on graphs in NLP by Nastase et al. (2015) covers only smaller graphs such as dependency graphs and dialogue trees. Even though it does not include KGs, the survey concludes that graphs are a powerful representation formalism and how NLP tasks can benefit from harnessing the potential of data presented in graph structures.

To the best of our knowledge, this is the first survey covering a wide spectrum of techniques, methods as well as applications of KGs within the NLP research field.

3 Method

To achieve our objective of providing a thorough overview of the research landscape, we conducted a systematic mapping study following the process defined by Petersen et al. (2008). Its three main steps are explained in the next subsections.

3.1 Research Questions

The goal of our study is a multifaceted analysis of KGs in the field of NLP, such as identifying and quantifying research topics, domains, and out-

comes. These objectives are reflected in the research questions (RQs) stated below.

RQ1: What are the characteristics and trends of the research literature on KGs in NLP?

RQ2: What are the different tasks mentioned in the existing research studies?

RQ3: What are the research types and main contributions of the studies?

3.2 Search and Screening Procedure

After specifying the RQs, we defined a set of related keywords for KGs and NLP to be used for the database search of relevant studies. From initial test searches, we observed that including terms associated with KGs (e.g., “semantic network” or “ontology”) yielded too many irrelevant results. To restrict the research scope to the concept of KGs, we decided to use the following search string:

("knowledge graph") AND ("NLP" OR "natural language processing" OR "computational linguistics"). The search string was applied to title, abstract, and keywords. If a given paper had no keywords, we used index keywords from the database if they were available.

For our search of relevant publications, we queried six academic databases, as listed in Table 1. The ACL Anthology is a digital archive of prestigious conferences and journals in NLP. ACM and IEEE provide access to publications of additional reputable venues in the broader computer science field. The remaining databases are commonly chosen in other related surveys to further increase the coverage of the respective field of interest.

In the first week of 2022, we applied our search string to the databases and restricted the time window to ten years from 2012 until 2021. Then, the exported files were merged, ensuring that each publication record was either a conference or a journal paper. We automatically identified and removed duplicate records as well. Through this, we obtained a dataset of 746 unique papers. Given this initial dataset, we further filtered down the truly relevant studies by screening for the following inclusion criteria: (1) peer-reviewed studies from conferences or journals, (2) studies with a clear focus on KGs in NLP, (3) studies are written in English and full texts are electronically accessible. In reverse, this implies the publications that did not satisfy all three inclusion criteria were excluded from the dataset.

As part of the screening procedure, two of the authors read title, abstract, and keywords to deter-

Academic Database	No. of Papers
ACL Anthology	164
ACM Digital Library	26
IEEE Xplore	76
ScienceDirect	34
Scopus	200
Web of Science	7
Total	507

Table 1: Overview of academic databases and number of included papers.

mine if a paper matched the inclusion criteria. In ambiguous cases, the full text of the paper was examined. The two authors screened all papers and decided together on keeping or dropping records from the dataset. The final dataset included a total of 507 papers, as listed in Table 1. We make our annotated dataset available through a public GitHub repository.¹

3.3 Classification Scheme and Data Extraction

According to our RQs, the included papers had to be categorized with respect to three facets: task, research type, and contribution. Established classification schemes from Wieringa et al. (2006) and Shaw (2003) were adapted for the research and contribution type as presented in Appendix A. For classifying tasks, we constructed a task taxonomy, following the iterative procedure suggested by Petersen et al. (2008), in which an initial classification scheme derived from keywords continuously evolves through adding, merging, or splitting categories during the classification process. Our task taxonomy is based on existing schemes from Paulheim (2017), Liu et al. (2020a), and Ji et al. (2021). Once the initial schemes were set up, all papers were sorted into the classes as part of the data extraction process. The 507 included studies were divided between two of the authors. In regular sessions, they discussed changes to the classification schemes or clarified uncertain labels. While each paper got assigned one label for the research type assigned, multiple labels were possible with regard to tasks and contributions. To assess the reliability of the inter-annotator agreement, the two authors independently classified a random sample of 50 papers. We calculated Cohen’s Kappa coefficient of these annotations for each facet (Cohen, 1960).

¹<https://github.com/sebischair/KG-in-NLP-survey>

The annotations of the task, research, and contribution facets had coefficients of 0.73, 0.87, and 0.76, respectively. Cohen suggested interpreting Kappa values from 0.61 to 0.80 as substantial and from 0.81 to 1.00 as almost perfect agreement.

4 Results

In this chapter, we report the results of the data extraction process. It is arranged into subsections according to the formulated RQs.

4.1 Characteristics of the Research Landscape (RQ1)

In regard to the literature on KGs in NLP, we started our analysis by looking at the number of studies as an indicator of research interest. The distribution of publications over the ten-year observation period is illustrated in Figure 2. While the first publications appear in 2013, the annual publications grew slowly between 2013 and 2016. From 2017 onwards, the number of publications doubled almost every year. Because of the significant rise in research interest within these years, more than 90% of all included publications originate from these five years. Even though the growth trend seems to stop in 2021, this is likely due to the data export which happened in the first week of 2022, leaving out many studies from 2021 that were enlisted in the databases later in 2022. Nonetheless, the trend in Figure 2 clearly indicates that KGs are receiving increasing attention from the NLP research community. Considering the 507 included papers, the number of conference papers (402) was nearly four times as high as that of journal papers (105).




Figure 2: Distribution of number of papers from 2012 to 2021 (database export was performed in the first week of the year 2022).

We also investigated institutional affiliations by country to determine what countries are most active in the field of KGs in NLP. In total, we identified 44 countries contributing to the research literature. As part of the Appendix, we provide a world map with all countries in Figure 7 and a list of the top 20 countries by the number of affiliated papers in Table 7. China ranks first and holds a major proportion with 199 papers, accounting for 39% of all publications. The United States and India come in second and third with 119 and 49 papers, respectively. Germany, the United Kingdom, and Italy follow in the ranking. All European countries had a combined total of 141 affiliated publications.

Another finding of the data extraction process concerns the diverse application areas of KGs in NLP. We observed that the number of domains explored in the research literature grew rapidly in parallel with the annual count of papers. To reveal the great variety of areas, we list all 20 discovered domains and their subdomains in Table 6 in the Appendix. In Figure 3, the ten most frequent domains are displayed. It is striking that health is by far the most prominent domain. The latter appears more than twice as often as the scholarly domain, which ranks second. Other popular areas are engineering, business, social media, or law. In view of the domain diversity, it becomes evident that KGs are naturally applicable to many different contexts, as has been stated in prior work (Abu-Salih, 2021; Ji et al., 2021; Zou, 2020).

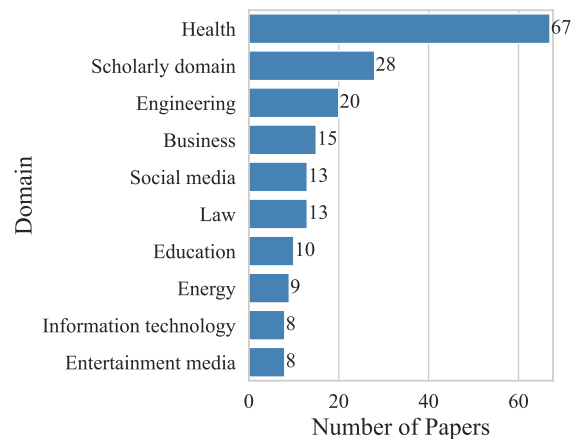


Figure 3: Number of papers by most popular application domains.

4.2 Tasks in the Research Literature (RQ2)

Based on the tasks identified in the literature on KGs in NLP, we developed the empirical taxon-

Task	No. of Papers	Representative Papers
Relation extraction	144	Peng et al. (2017), Wang et al. (2018b), Zhang et al. (2019a)
Entity extraction	143	Rospocher et al. (2016), Luan et al. (2018), Wang et al. (2018a)
Question answering	103	Bao et al. (2016), Zhang et al. (2018), Feng et al. (2020)
Semantic search	91	Speer et al. (2017), Wang et al. (2020), Gaur et al. (2021)
Augmented language models	84	Zhang et al. (2019b), Bosselut et al. (2019), Liu et al. (2020b)
Knowledge graph embedding	61	Shi and Weninger (2018), Ali et al. (2021), Wang et al. (2021b)
Entity linking	38	Kartsaklis et al. (2018), Moon et al. (2018), Chen et al. (2018)
Ontology construction	32	Gangemi et al. (2016), Haussmann et al. (2019), Li et al. (2020)
Conversational interfaces	29	Zhou et al. (2018) Moon et al. (2019), Wu et al. (2019)
Link prediction	26	Lv et al. (2019), Sun et al. (2020), Wang et al. (2021a)

Table 2: Overview of most popular tasks in the literature on KGs in NLP.

omy shown in Figure 1. The two top-level categories consist of knowledge acquisition and knowledge application. Knowledge acquisition contains NLP tasks to construct KGs from unstructured text (knowledge graph construction) or to conduct reasoning over already constructed KGs (knowledge graph reasoning). KG construction tasks are further split into two subcategories: knowledge extraction, which is used to populate KGs with entities, relations, or attributes, and knowledge integration, which is used to update KGs. Knowledge application, being the second top-level concept, encompasses common NLP tasks, which are enhanced through structured knowledge from KGs.

As might be expected, the frequency of occurrence in the literature for the tasks from our taxonomy varies greatly. While Table 2 gives an overview of the most popular tasks, Figure 5 compares their popularity over time. Figure 4 displays the number of detected domains for the most prominent tasks. It shows that certain tasks are adopted to more domain-specific contexts than others.

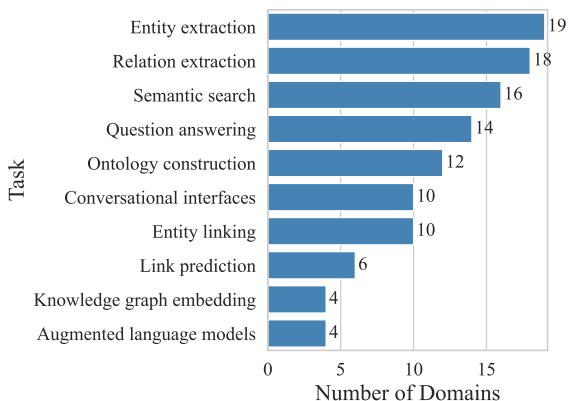


Figure 4: Overview of most popular tasks by number of application domains.

4.2.1 Knowledge Graph Construction

The task of **entity extraction** is a starting point in constructing KGs and is used to derive real-world entities from unstructured text (Al-Moslmi et al., 2020). Once the relevant entities are singled out, relationships and interactions between them are found with the task of **relation extraction** (Zhang et al., 2019a). A lot of papers use both entity extraction and relation extraction to construct new KGs, e.g., for news events (Rospocher et al., 2016) or scholarly research (Luan et al., 2018).

Entity linking is a task of linking entities recognized in some text to already existing entities in KGs (Moon et al., 2018; Wu et al., 2020). Since synonymous or similar entities often exist in different KGs or in different languages, **entity alignment** can be performed to reduce redundancy and repetition in future tasks (Gangemi et al., 2016; Chen et al., 2018). Coming up with the rules and schemes of KGs, i.e., their structure and format of knowledge presented in it, is done with the task of **ontology construction** (Haussmann et al., 2019).

4.2.2 Knowledge Graph Reasoning

Once constructed, KGs contain structured world knowledge and can be used to infer new knowledge by reasoning over them. Thereby, the task of classifying entities is called **entity classification**, while **link prediction** is the task of inferring missing links between entities in existing KGs often performed via ranking entities as possible answers to queries (Shi and Weninger, 2018; Bosselut et al., 2019; Wang et al., 2019; Ali et al., 2021).

Knowledge graph embedding techniques are used to create dense vector representations of a graph so that they can then be used for downstream machine learning tasks. While this problem can be

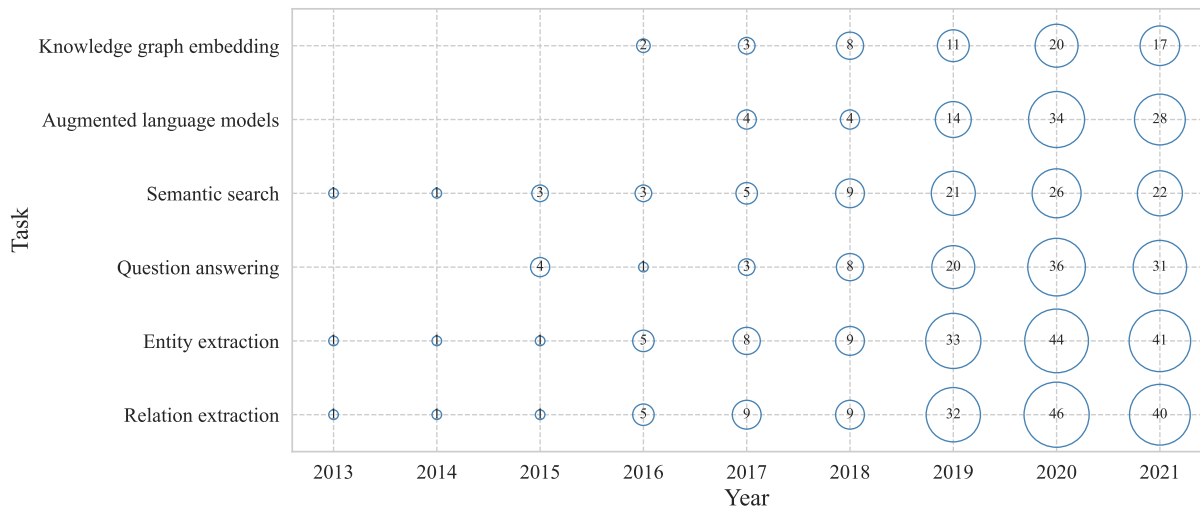


Figure 5: Distribution of number of papers by most popular tasks from 2013 to 2021.

related solely to KGs, in our survey this label refers to approaches that jointly learn text and graph embeddings (Chen et al., 2018; Wang et al., 2021b).

4.2.3 Knowledge Application

Existing KGs can be used in a multitude of popular NLP tasks. Here we outline the most popular ones.

Question answering (QA) was found to be the most common NLP task using KGs. This task is typically divided into textual QA and question answering over knowledge bases (KBQA). Textual QA derives answers from unstructured documents while KBQA does so from predefined knowledge bases (Fu et al., 2020). KBQA is naturally tied to KGs while textual QA can also be approached by using KGs as a source of common-sense knowledge when answering questions. As Zhu et al. (2021) conclude, this approach is desired not only because it is helpful for generating answers, but also because it makes answers more interpretable.

Semantic search refers to "search with meaning", where the goal is not just to search for literal matches, but to understand the search intent and query context as well (Bast et al., 2016). This label denoted studies that use KGs for search, recommendations, and analytics. Examples are a big semantic network of everyday concepts called ConceptNet (Speer et al., 2017) and a KG of scholarly communications and the relationships, among them the Microsoft Academic Graph (Wang et al., 2020).

Conversational interfaces constitute another NLP field that can benefit from world knowledge contained in KGs. Zhou et al. (2018) utilize the knowledge from KGs to generate responses of con-

versational agents that are more informative and appropriate in a given context. Knowledge-aware dialogue generation was also explored by Moon et al. (2019), Wu et al. (2019), Liu et al. (2019).

Natural language generation (NLG) is a subfield of NLP and computational linguistics that is concerned with models which generate natural language output from scratch. KGs are used in this subfield for producing natural language text from KGs (Koncel-Kedziorski et al., 2019), generating question-answer pairs (Reddy et al., 2017), the multi-modal task of image captioning (Lu et al., 2018), or data augmentation in low-resource settings (Sharifirad et al., 2018).

Text analysis combines various analytical NLP techniques and methods that are applied to process and understand textual data. Exemplary tasks are sentiment detection (Kumar et al., 2018), topic modeling (Li et al., 2019), or word sense disambiguation (Kumar et al., 2019).

Augmented language models are a combination of large pretrained language models (PLMs) such as BERT (Devlin et al., 2019) and GPT (Radford et al., 2018) with knowledge contained in KGs. Since PLMs derive their knowledge from huge amounts of unstructured training data, a rising research trend is in combining them with structured knowledge. Knowledge from KGs can be infused into language models in their input, architecture, output, or some combination thereof (Colon-Hernandez et al., 2021). Some notable examples we outline are ERNIE (Zhang et al., 2019b), COMET (Bosselut et al., 2019), K-BERT (Liu et al., 2020b), and KEPLER (Wang et al., 2021b).

4.3 Research Types and Contributions (RQ3)

Table 3 shows the distribution of papers according to the different research and contribution types as defined in Table 4 and 5 in the Appendix. It shows that most papers conduct validation research, investigating new techniques or methods that have not yet been implemented in practice. A considerable number of papers, although significantly less, focus on solution proposals of approaches by demonstrating their advantages and applicability by a small example or argumentation. However, these papers usually lack a profound empirical evaluation. Secondary research accounts for only a small number of papers and is severely underrepresented in the research field of KGs in NLP. As already mentioned in Section 1 and Section 2, there is a notable lack of studies that summarize, compile, or synthesize existing research regarding KGs in NLP. Moreover, evaluation research papers that implement and evaluate approaches in an industry context are equally scarce. Opinion papers are almost non-existent.

In terms of contribution types, techniques, methods, and tools are predominant. Resources and guidelines, as opposed to this, are rather underrepresented. This is in accordance with the distribution of research types, which indicates that mainly new methods and techniques are researched, but hardly any secondary research is conducted. Additionally, the research area of KGs in NLP is lacking new resources such as text corpora, benchmarks, or constructed graphs.

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 3: Number of papers by research type and contribution type.

Figure 6 depicts the different tasks of the analyzed studies and their relative share of contribution types. We can notice that entity extraction and relation extraction, which encompass the most works

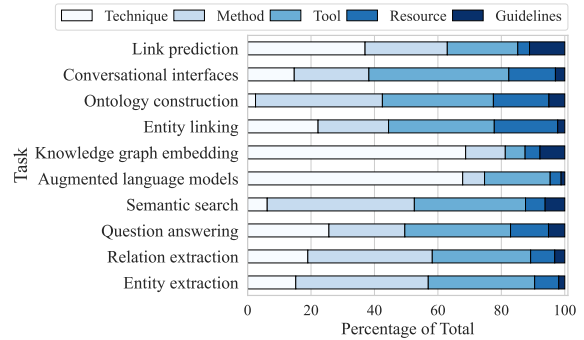


Figure 6: Percentage of contribution type by tasks.

in line with Table 2, have a very balanced distribution of contribution types. These tasks, which build the foundation for KG construction, have been researched for a long time and the number of studies in these areas is continually increasing, as can be seen in Figure 5. Furthermore, a comparison of Figure 5 with Figure 6 shows that tasks, such as relation extraction or semantic search, which have existed for some time and continue to grow steadily have a rather balanced ratio of contribution types, too. This is an indication that these tasks are already reasonably mature, as some extensive preliminary work is required, for example, to use multiple techniques in a new method.

Additionally, mature research areas already focus on industrialization, investigating how to use techniques in different domains and developing tools. Figure 4 strengthens the impression that tasks such as relation extraction or semantic search are already reasonably mature, as they are used in many different domains. In contrast, immature research areas still primarily focus on investigating new techniques and are used in a few domains only. For instance, the augmented language models and knowledge graph embedding tasks have mainly techniques as the contribution type and are not used in many different domains. Therefore, they can still be considered relatively immature. This may be a result of the fact that these tasks are still relatively young and less investigated. Figure 5 shows that the two tasks have only seen a sharp increase in studies from 2018 onwards and attracted a lot of interest since then.

5 Discussion

The observations of our comprehensive survey reveal several insights. It is important to situate these findings with respect to related work and industry reports in the artificial intelligence (AI) field.

Since the first publications in 2013, researchers worldwide have paid increasing attention to study KGs from a NLP perspective, especially in the past five years. This observed growth in research interest is in line with the KG survey of [Chen et al. \(2021\)](#). We identified China and the United States as the most active countries shaping the research landscape, which is to be expected considering both countries regularly claim the top ranks in the popular "AI Index Report" from Stanford University ([Zhang et al., 2021](#)). The report further highlights a soaring AI investment in the health domain. The latter was also the most dominant domain in our results (see Figure 3). However, research in the health domain has to be considered critically, since these works compare poorly to other domains regarding reproducibility metrics, such as dataset and code accessibility ([McDermott et al., 2021](#)).

Table 3 shows evidently that the research field of KGs in NLP is lacking new resources such as text corpora, benchmarks, or KGs. This leads to the assumption that most works train and evaluate using the same limited available datasets and benchmarks. As a result, novel approaches are often optimized only for certain available benchmarks which may not hold up in practice. Furthermore, the lack of secondary research visible in Table 3 reveals the need for more works that present an overview of the research field.

The frequency of tasks in our survey greatly varies, as reflected in Table 2. Studies concerning KG construction account for the majority of all papers. Applied NLP tasks such as QA and semantic search also have a strong research community. The most emergent topics in recent years have been augmented language models, QA, and KG embedding. Some of the outlined tasks are still confined to the research community, while others have found practical application in many real-life contexts. From Figure 4 it is evident that the KG construction tasks and semantic search over KGs are the most widely applied ones. Of the NLP tasks, QA and conversational interfaces have been adopted to many real-life domains, usually in the form of digital assistants. Tasks like KG embedding and augmented language models are still only being researched and lack a widespread practical adoption in real-world scenarios. We anticipate that as the research areas of augmented language models and KG embedding mature, more methods and tools will be investigated for these tasks.

6 Limitations

Although we employed a rigorous study design and paid careful attention to executing each search and analysis step, our study is subject to limitations.

Given the restriction to one search string and six databases, there should be some relevant publications that we did not retrieve. This is the case for studies that did not mention our search terms in title, abstract, or keywords. To mitigate the risk of incompleteness, we chose common databases with a large number of publications in the examined research area. Further, we performed a preliminary search to optimize the completeness of results. Whenever possible, we replaced missing keywords with index keywords from the source database.

Moreover, the screening for relevant studies depends on the personal assessment of the researchers, which can bias the study selection. As a countermeasure, we defined selection criteria for the inclusion and exclusion of studies. During the study selection, two researchers assessed of selection criteria in parallel and discussed contradicting decisions until they reached a consensus to mitigate subjective bias.

The accuracy of the classification results constitutes another threat to the validity of our study. Data extraction bias may negatively affect the accuracy of the classification results. To mitigate this risk, the authors regularly discussed the used classification schemes and assigned labels to establish a common understanding of each class. In addition, we calculated Cohen's Kappa coefficient to quantify the reliability of the inter-annotator agreement.

7 Conclusion

Recent years have witnessed a rising prominence of KGs in NLP research. Despite the rapidly growing body of literature, until now, no study has been published that summarizes the progress so far. To provide an overview of this maturing research area, we performed a multifaceted survey of tasks, research types, and contributions.

Our findings show that a large number of tasks concerning KGs in NLP have been studied across various domains, including emerging topics like knowledge graph embedding or augmented language models. However, we observed a lack of secondary research and evaluations in practice, both of which are crucial to reflect the major scientific progress of the field as a whole. Our study lays the grounds for further research in this direction.

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

- Bilal Abu-Salih. 2021. [Domain-specific knowledge graphs: A survey](#). *Journal of Network and Computer Applications*, 185:103076.
- Tareq Al-Moslmi, Marc Gallofré Ocaña, Andreas L. Opdahl, and Csaba Veres. 2020. [Named entity extraction for knowledge graphs: A literature overview](#). *IEEE Access*, 8:32862–32881.
- Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Mikhail Galkin, Sahand Sharifzadeh, Asja Fischer, Volker Tresp, and Jens Lehmann. 2021. [Bringing light into the dark: A large-scale evaluation of knowledge graph embedding models under a unified framework](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1.
- Junwei Bao, Nan Duan, Zhao Yan, Ming Zhou, and Tiejun Zhao. 2016. [Constraint-based question answering with knowledge graph](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2503–2514, Osaka, Japan. The COLING 2016 Organizing Committee.
- Hannah Bast, Björn Buchhold, Elmar Haussmann, et al. 2016. [Semantic search on text and knowledge bases](#). *Foundations and Trends® in Information Retrieval*, 10(2-3):119–271.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. [COMET: Commonsense transformers for automatic knowledge graph construction](#). In *Proceedings of the 57th Annual Meeting of the ACL*, pages 4762–4779. Association for Computational Linguistics.
- Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018. [Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment](#). In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 3998–4004. International Joint Conferences on Artificial Intelligence Organization.
- XiaoJun Chen, Shengbin Jia, and Yang Xiang. 2019. [A review: Knowledge reasoning over knowledge graph](#). *Expert Systems with Applications*, 141:112948.
- Xieling Chen, Haoran Xie, Zongxi Li, and Gary Cheng. 2021. [Topic analysis and development in knowledge graph research: A bibliometric review on three decades](#). *Neurocomputing*, 461:497–515.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and psychological measurement*, 20(1):37–46.
- Allan M. Collins and M. Ross Quillian. 1969. [Retrieval time from semantic memory](#). *Journal of Verbal Learning and Verbal Behavior*, 8(2):240–247.
- Pedro Colon-Hernandez, Catherine Havasi, Jason B. Alonso, Matthew Huggins, and Cynthia Breazeal. 2021. [Combining pre-trained language models and structured knowledge](#). *CoRR*, abs/2101.12294.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186. Association for Computational Linguistics.
- Lisa Ehrlinger and Wolfram Wöß. 2016. [Towards a definition of knowledge graphs](#). In *SEMANTiCS*.
- Michael Färber, Frederic Bartscherer, Carsten Menne, Achim Rettinger, Amrapali Zaveri, Dimitris Kontokostas, Sebastian Hellmann, and Jürgen Umbrich. 2018. [Linked data quality of dbpedia, freebase, open-cyc, wikidata, and yago](#). *Semant. Web*, 9(1):77–129.
- Yanlin Feng, Xinyue Chen, Bill Yuchen Lin, Peifeng Wang, Jun Yan, and Xiang Ren. 2020. [Scalable multi-hop relational reasoning for knowledge-aware question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1295–1309, Online. Association for Computational Linguistics.
- Bin Fu, Yunqi Qiu, Chengguang Tang, Yang Li, Haiyang Yu, and Jian Sun. 2020. [A survey on complex question answering over knowledge base: Recent advances and challenges](#). *CoRR*, abs/2007.13069.
- Aldo Gangemi, Mehwish Alam, Luigi Asprino, Valentina Presutti, and Diego Reforgiato Recupero. 2016. [Framester: A wide coverage linguistic linked data hub](#). In *European knowledge acquisition workshop*, pages 239–254. Springer.
- Manas Gaur, Keyur Faldu, and Amit Sheth. 2021. [Semantics of the black-box: Can knowledge graphs help make deep learning systems more interpretable and explainable?](#) *IEEE Internet Computing*, 25(1):51–59.
- Steven Haussmann, Oshani Seneviratne, Yu Chen, Yarden Ne’eman, James Codella, Ching-Hua Chen, Deborah L. McGuinness, and Mohammed J. Zaki. 2019. [Foodkg: A semantics-driven knowledge graph for food recommendation](#). In *The Semantic Web – ISWC 2019*, pages 146–162, Cham. Springer International Publishing.

- Pascal Hitzler. 2021. [A review of the semantic web field](#). *Commun. ACM*, 64(2):76–83.
- Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D’amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. 2022. [Knowledge graphs](#). *ACM Computing Surveys*, 54(4):1–37.
- Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and S Yu Philip. 2021. [A survey on knowledge graphs: Representation, acquisition, and applications](#). *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514.
- Dimitri Kartsaklis, Mohammad Taher Pilehvar, and Nigel Collier. 2018. [Mapping text to knowledge graph entities using multi-sense LSTMs](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1959–1970, Brussels, Belgium. Association for Computational Linguistics.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. [Text Generation from Knowledge Graphs with Graph Transformers](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2284–2293, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abhishek Kumar, Daisuke Kawahara, and Sadao Kurohashi. 2018. [Knowledge-enriched two-layered attention network for sentiment analysis](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 253–258, New Orleans, Louisiana. Association for Computational Linguistics.
- Sawan Kumar, Sharmistha Jat, Karan Saxena, and Partha Talukdar. 2019. [Zero-shot word sense disambiguation using sense definition embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5670–5681, Florence, Italy. Association for Computational Linguistics.
- Dingcheng Li, Siamak Zamani, Jingyuan Zhang, and Ping Li. 2019. [Integration of knowledge graph embedding into topic modeling with hierarchical dirichlet process](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 940–950.
- Xinyu Li, Chun-Hsien Chen, Pai Zheng, Zuoxu Wang, Zuhua Jiang, and Zhixing Jiang. 2020. [A Knowledge Graph-Aided Concept–Knowledge Approach for Evolutionary Smart Product–Service System Development](#). *Journal of Mechanical Design*, 142(10).
- Shuang Liu, Hui Yang, Jiayi Li, and Simon Kolmanič. 2020a. [Preliminary study on the knowledge graph construction of chinese ancient history and culture](#). *Information*, 11(4):186.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020b. [K-bert: Enabling language representation with knowledge graph](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2901–2908.
- Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. [Knowledge aware conversation generation with explainable reasoning over augmented graphs](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Di Lu, Spencer Whitehead, Lifu Huang, Heng Ji, and Shih-Fu Chang. 2018. [Entity-aware image caption generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4013–4023, Brussels, Belgium. Association for Computational Linguistics.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Xin Lv, Yuxian Gu, Xu Han, Lei Hou, Juanzi Li, and Zhiyuan Liu. 2019. [Adapting meta knowledge graph information for multi-hop reasoning over few-shot relations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 3376–3381, Hong Kong, China. Association for Computational Linguistics.
- Matthew B. A. McDermott, Shirley Wang, Nikki Marinsek, Rajesh Ranganath, Luca Foschini, and Marzyeh Ghassemi. 2021. [Reproducibility in machine learning for health research: Still a ways to go](#). *Science Translational Medicine*, 13(586):eabb1655.
- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. 2018. [Multimodal named entity disambiguation for noisy social media posts](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 2000–2008, Melbourne, Australia. Association for Computational Linguistics.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. [OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy. Association for Computational Linguistics.

- Vivi Nastase, Rada Mihalcea, and Dragomir R Radev. 2015. [A survey of graphs in natural language processing](#). *Natural Language Engineering*, 21(5):665–698.
- David N. Nicholson and Casey S. Greene. 2020. [Constructing knowledge graphs and their biomedical applications](#). *Computational and Structural Biotechnology Journal*, 18:1414–1428.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2016. [A review of relational machine learning for knowledge graphs](#). *Proceedings of the IEEE*, 104(1):11–33.
- Heiko Paulheim. 2017. [Knowledge graph refinement: A survey of approaches and evaluation methods](#). *Semantic Web*, 8(3):489–508.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. [Cross-sentence n-ary relation extraction with graph LSTMs](#). *Transactions of the Association for Computational Linguistics*, 5:101–115.
- Kai Petersen, Robert Feldt, Shahid Mujtaba, and Michael Mattsson. 2008. [Systematic mapping studies in software engineering](#). In *12th International Conference on Evaluation and Assessment in Software Engineering (EASE) 12*, pages 1–10.
- Jay Pujara, Hui Miao, Lise Getoor, and William Cohen. 2013. [Knowledge graph identification](#). In *The Semantic Web – ISWC 2013*, pages 542–557, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Ross Quillian. 1963. *A Notation for Representing Conceptual Information: an Application to Semantics and Mechanical English Paraphrasing*. Systems Development Corporation.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#). *OpenAI Blog*.
- Sathish Reddy, Dinesh Raghu, Mitesh M. Khapra, and Sachindra Joshi. 2017. [Generating natural language question-answer pairs from a knowledge graph using a RNN based question generation model](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 376–385, Valencia, Spain. Association for Computational Linguistics.
- R. H. Richens. 1956. [Preprogramming for mechanical translation](#). *Mech. Transl. Comput. Linguistics*, 3.
- Marco Rospocher, Marieke van Erp, Piek Vossen, Antske Fokkens, Itziar Aldabe, German Rigau, Aitor Soroa, Thomas Ploeger, and Tessel Bogaard. 2016. [Building event-centric knowledge graphs from news](#). *Web Semant.*, 37(C):132–151.
- Tara Safavi and Danai Koutra. 2021. [Relational World Knowledge Representation in Contextual Language Models: A Review](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1053–1067, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sima Sharifirad, Borna Jafarpour, and Stan Matwin. 2018. [Boosting text classification performance on sexist tweets by text augmentation and text generation using a combination of knowledge graphs](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 107–114, Brussels, Belgium. Association for Computational Linguistics.
- Mary Shaw. 2003. [Writing good software engineering research papers](#). In *25th International Conference on Software Engineering, 2003. Proceedings.*, pages 726–736. IEEE.
- Baoxu Shi and Tim Wener. 2018. [Open-world knowledge graph completion](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Amit Singhal. 2012. [Introducing the knowledge graph: Things, not strings](#). *Google Blog*.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*, page 4444–4451. AAAI Press.
- Zhiqing Sun, Shikhar Vashishth, Soumya Sanyal, Partha Talukdar, and Yiming Yang. 2020. [A re-evaluation of knowledge graph completion methods](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5516–5522, Online. Association for Computational Linguistics.
- Bo Wang, Tao Shen, Guodong Long, Tianyi Zhou, Ying Wang, and Yi Chang. 2021a. [Structure-augmented text representation learning for efficient knowledge graph completion](#). In *Proceedings of the Web Conference 2021*. ACM.
- Chengbin Wang, Xiaogang Ma, Jianguo Chen, and Jingwen Chen. 2018a. [Information extraction and knowledge graph construction from geoscience literature](#). *Comput. Geosci.*, 112:112–120.
- Guanying Wang, Wen Zhang, Ruoxu Wang, Yalin Zhou, Xi Chen, Wei Zhang, Hai Zhu, and Huajun Chen. 2018b. [Label-free distant supervision for relation extraction via knowledge graph embedding](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2246–2255, Brussels, Belgium. Association for Computational Linguistics.
- Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul Kanakia. 2020. [Microsoft Academic Graph: When experts are not enough](#). *Quantitative Science Studies*, 1(1):396–413.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021b.

- KEPLER: A unified model for knowledge embedding and pre-trained language representation.** *Transactions of the Association for Computational Linguistics*, 9:176–194.
- Zihao Wang, Kwunping Lai, Piji Li, Lidong Bing, and Wai Lam. 2019. **Tackling long-tailed relations and uncommon entities in knowledge graph completion.** In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 250–260, Hong Kong, China. Association for Computational Linguistics.
- Roel Wieringa, Neil Maiden, Nancy Mead, and Colette Rolland. 2006. **Requirements engineering paper classification and evaluation criteria: a proposal and a discussion.** *Requirements engineering*, 11(1).
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. **Scalable zero-shot entity linking with dense entity retrieval.** In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6397–6407, Online. Association for Computational Linguistics.
- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. **Proactive human-machine conversation with explicit conversation goal.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804, Florence, Italy. Association for Computational Linguistics.
- Daniel Zhang, Saurabh Mishra, Erik Brynjolfsson, John Etchemendy, Deep Ganguli, Barbara J. Grosz, Terah Lyons, James Manyika, Juan Carlos Niebles, Michael Sellitto, Yoav Shoham, Jack Clark, and C. Raymond Perrault. 2021. **The AI index 2021 annual report.** *CoRR*, abs/2103.06312.
- Ningyu Zhang, Shumin Deng, Zhanlin Sun, Guanying Wang, Xi Chen, Wei Zhang, and Huajun Chen. 2019a. **Long-tail relation extraction via knowledge graph embeddings and graph convolution networks.** In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3016–3025, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J. Smola, and Le Song. 2018. **Variational reasoning for question answering with knowledge graph.** In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence and 30th Innovative Applications of Artificial Intelligence Conference and 8th AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’18/IAAI’18/EAAI’18. AAAI Press.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019b. **ERNIE: Enhanced language representation with informative entities.** In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, Florence, Italy. Association for Computational Linguistics.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. **Commonsense knowledge aware conversation generation with graph attention.** In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI’18*, page 4623–4629. AAAI Press.
- Fengbin Zhu, Wenqiang Lei, Chao Wang, Jianming Zheng, Soujanya Poria, and Tat-Seng Chua. 2021. **Retrieving and reading: A comprehensive survey on open-domain question answering.** *arXiv preprint arXiv:2101.00774*.
- Xiaohan Zou. 2020. **A survey on application of knowledge graph.** In *Journal of Physics: Conference Series*, volume 1487. IOP Publishing.

A Supplementary Material

Table 4 shows the classification scheme for research types. Wieringa et al. (2006) introduced this scheme in order to categorize different types of research papers with differing approaches to what is being studied. Although the categories of evaluation research and validation research seem to be similar, there is a key difference. A paper is considered to be evaluation research only if the investigated problem is implemented and evaluated in practice. Papers labeled as validation research investigate properties of proposed solutions that have not been implemented in practice, while solution proposal papers introduce new solutions without a rigorous empirical validation.

Table 5 shows the classification scheme of contribution types employed in this study. It is based on the classification scheme of Shaw (2003) and adapted to the field of KGs in NLP. Here, special attention needs to be paid to the distinction between method and technique. While a technique concentrates on solving a single specific task, a method involves a set of different techniques as well as procedures that must be executed in a systematic way to achieve a concrete objective.

Table 6 contains an overview of the 20 domains we discovered in the literature on KGs in NLP. For each domain, we identified a set of subdomains, which is listed as well.

Table 7 and the world map in Figure 7 give information about the number of papers by affiliated countries. While the table only shows the top 20 most active countries, the world map presents a global overview of all 44 countries contributing to the research literature.

Research Type	Description
Evaluation research	The implementation of an existing technique or method is evaluated in practice within an industry context.
Opinion paper	Report of the personal opinion of somebody on the suitability of a certain technique or method without relying on related work and research methods.
Secondary research	Analysis and synthesis of findings from multiple studies to systematically review a research field or gather evidence on a topic.
Solution proposal	Proposal of novel solution or extension for a technique or method by demonstrating their advantages and applicability by a small example or argumentation.
Validation research	Empirical investigation of characteristics from proposed techniques or methods that have not been implemented in practice yet.

Table 4: Classification scheme for research types adapted from [Wieringa et al. \(2006\)](#).

Contribution Type	Description
Guidelines	List of advices or recommendations derived from the obtained research results.
Method	A method contains a set of techniques and procedures that need to be systematically executed to achieve a concrete goal.
Resource	A resource is a published data set that supports techniques, methods, or tools, e.g., text corpora, benchmarks, or knowledge graphs.
Technique	A technique is the manner in which a concrete task within our task taxonomy is performed, often in the form of an algorithm or mathematical model.
Tool	A tool is a documented implementation of a technique or method in the form of a software library, prototype, or full application system.

Table 5: Classification scheme for contribution types adapted from [Shaw \(2003\)](#).

Domain	Identified Subdomains
Agriculture	Agricultural production, agricultural plant species
Business	E-commerce, finance, human resources, product design, real estate
Culture	Cultural heritage, ethnic minorities, film culture, museums, poetry
Education	Curriculum design, digital library, e-learning, moral education
Energy	Oil and gas industry, power grid fault disposal, smart grid
Engineering	Mechanical engineering, software engineering, electrical engineering
Entertainment media	Computer games, media recommendation, movies, music, television
Food	Dietary choices, recipe search
Health	Biomedicine, traditional Chinese medicine, pharmacology, mental health
History	Genealogy, historical events, retrieval of historical documents
Information technology	App ecosystems, Internet of Things, technical support, cybersecurity
Law	Law enforcement, patents, privacy policies, identity fraud detection
Natural science	Mineralogy, oceanography, petroleum geology
Scholarly domain	Bibliometrics, grant datasets, research collaborations, scientific corpora
News	Fake news detection, journalism, news exploration
Public sector	Government, military, poverty reduction, public safety organizations
Social media	Insight extraction from posts, misinformation detection, opinion mining
Social science	Open-source social science, social network analysis
Sports	Basketball, football
Tourism	Tourism question answering system, travel guide

Table 6: Overview of identified application domains and subdomains.

Rank	Country	No. of Affiliated Papers
1	China	199
2	United States	119
3	India	49
4	Germany	47
5	United Kingdom	34
6	Italy	21
7	Canada	19
8	Spain	16
9	France	15
10	Singapore	14
11	Australia	13
12	Hong Kong	10
13	Ireland	9
14	Netherlands	8
15	Japan	8
16	South Korea	6
17	Switzerland	6
18	Greece	5
19	Brazil	5
20	Portugal	4

Table 7: Overview of top 20 countries by number of affiliated papers.

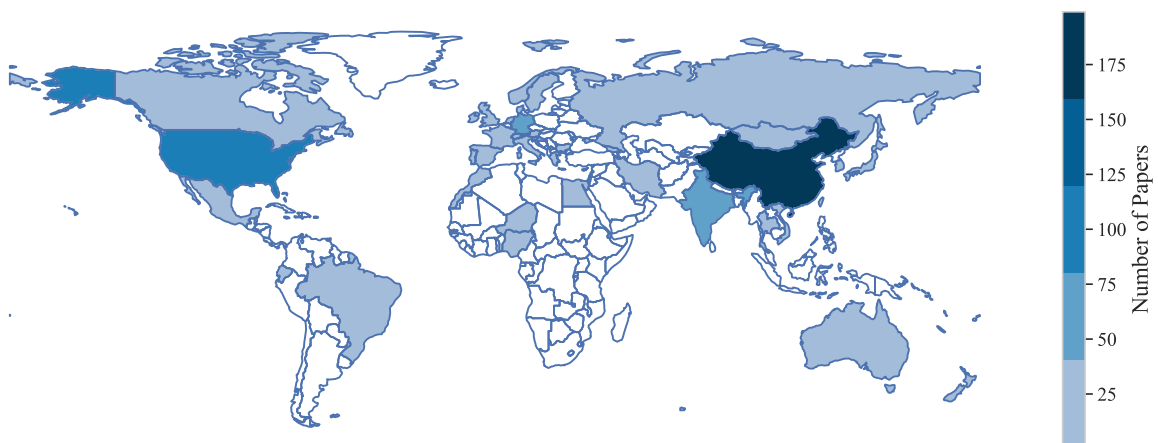


Figure 7: Global overview of number of papers by affiliated country.