Knowledge Graphs and their Applications in Civil Security

Simon Ott, Daria Liakhovets, Mina Schütz, Medina Andresel,
Mihai Bartha, Sven Schlarb, Alexander Schindler
Austrian Institute of Technology GmbH
Giefinggasse 4, 1210 Vienna, Austria
`{firstname.lastname}@ait.ac.at`

Abstract

In civil security, knowledge graphs play a crucial role by representing complex relationships among key entities, enhancing operational efficiency and ultimately ensuring the safety and well-being of communities. In this paper we introduce the technical tasks and methodology used for the applications of knowledge graphs in the context of civil security and describe the design of a software system for creating and leveraging knowledge graphs in use-cases such as hybrid threats identification, supply-chain resilience and situational awareness.

1 Introduction

Civil security refers to measures and efforts implemented by governments and organizations to protect civilians and maintain public order, encompassing emergency preparedness, disaster response, law enforcement, and the safeguarding of critical infrastructure to ensure the safety and well-being of the population. In the realm of civil security, the effective management and utilization of vast and diverse information sources play a pivotal role in ensuring the safety and well-being of communities. Knowledge graphs, with their ability to represent complex relationships and interconnections within data, offer a potent solution to enhance the efficiency and efficacy of civil security operations. Knowledge graphs serve as a unified framework for organizing and integrating heterogeneous data from various sources, such as news data, social media, and government databases. By structuring data into interconnected nodes and edges, knowledge graphs facilitate the discovery of valuable insights and data patterns useful for decision making.

In the following, we will give an overview over different tasks that deal with the creation and utilization of knowledge graphs. These task include the extraction of knowledge graphs from unstructured text data and the fusion of multi-model data, the linkage and comparison of extracted facts to an existing curated knowledge base, the aggregation of features of entities, and explainability and interpretability of machine learning models infusing knowledge based on graphs. Finally, we will describe the design of a software system for storing and querying knowledge graphs, text, embeddings and other metadata and how the different components interact with each other.

2 Methodology

In many civil security applications, heterogeneous data coming from governmental sources must be combined with news articles or social media information. The KG-related tasks required for such use-cases are as follows: (a) knowledge graph construction, (b) knowledge graph completion and linkage to general purpose KGs, and (c) the use of KGs for further downstream tasks. We describe below the approaches we use for each task.

The First Austrian Symposium on AI, Robotics, and Vision (AIROV24).

Extraction of KG Triples and Information Fusion. In the context of decision-making applications there is a need to fuse information that comes from multi-modal data such as text, images, video and audio [33, 7]. KGs represent a suitable solution to represent and interlink relevant information coming from various sources, given that they offer formal guarantees and information reliability. The problem of (semi-)automatically constructing a knowledge graph involves the extraction, transformation and linking of raw information into an RDF graph. The methods used for this tasks depend on the type of the underlying data. To build a knowledge graph from raw text involves the identification of relevant entities and relations between them. Recently, zero-shot approaches for this task relying on large language models (LLMs) have received a lot of attention [34, 31, 5] since they do not require annotated data or expensive re-training. Retrieval augmented generation (RAG) approaches [15] firstly identify relevant existing knowledge for the input text which is then fed to the LLM model for improving the extraction results. Domain ontologies (also in combination with RAG) can be useful to guide the LLM-based extraction process to identify entities and relations that are only covered by the ontology [5]. For other data modalities, such as images and audio, we rely upon models that output a textual description which is then used for the KG construction. The output of the information extractors is then fused with existing entities and relations in the KG. For this purpose, representation learning techniques such as KG embedding models are used to rank similar existing entities in the graph, which are then being evaluated by a human expert before being merged into the KG. In this way, the validated KG represents a reliable source of information which can then be analysed and further used for other downstream tasks. In addition, logical rules can also be useful in integrating background and expert knowledge.

Linking to General-Purpose KGs. Existing knowledge graphs of specific domains can also be enriched by linking entities to entities of large-scale general-purpose knowledge graphs, such as DBPedia [2] and Wikidata [30]. For example one database highly suitable for analysing news and events in civil security is the Global Database of Events, Language, and Tone (GDELT) [14], which is a large-scale, open-source collection of online news and event metadata which gets updated nearly in real time. GDELT provides a Global Knowledge Graph (GKG), a knowledge graph which links nodes of type *News article* to nodes of type *Organisation*, *Person*, *Location* and *Theme*. However, since there are only relations between nodes of type *News article* and the extracted entities, the ability of reasoning on this knowledge graph is rather limited. To address this issue, we enrich the knowledge graph by linking the extracted entities to entities of general purpose knowledge bases such as DBpedia. For this, language models can be used to recognize, disambiguate and link named entities. An example for such a pipeline is DBPedia Spotlight [20]. The final enriched graph is created, by sampling subgraphs using the extracted nodes as roots (e.g. k-hop neighborhood sampling).

Fact Analysis. One use case of knowledge graphs is to compare claimed facts from e.g. news articles with a collection of curated facts. To achieve this, link prediction approaches can be used. The task of link prediction is to predict a tail (head) entity *t* given a head (tail) entity and relation $r(h, ?)$ and $r(?, t)$, e.g. $bornIn(Obama, ?)$. Most of these approaches yield a score for the likelihood of a triple being true. This score can also be interpreted as a score that measures how much the link prediction model - that was trained on the curated facts - agrees with the claimed fact and thus can be used for assessing an estimation about the support of a claimed fact. For example the claimed fact $bornIn(Obama, Kenya)$ should receive a very low score if there is little evidence about the fact being true, while $locatedIn(London, UK)$ should receive a high score even if these facts are unknown and were not in the underlying training data. Several approaches have been developed for this task, latent approaches that embed entities and relations in a vector space (knowledge graph embedding such as TransE [3], ComplEx [3] or graph neural network based such as RGCN [24]), symbolic approaches that try to mine rules from the graph such as AMIE [10] or AnyBURL [19], and hybrid neuro-symbolic approaches such as [23] that try to combine both the latent and symbolic space. One limitation of current link prediction approaches is, that they cannot handle entities not seen during training. For the previously introduced application, this poses a problem, as news data is highly dynamic and training the models on the full set of knowledge contained in current general purpose knowledge bases is computationally not feasible due to their large-scale nature. However approaches for inductive link prediction is an active research field [11]. These approaches can handle new unseen entities to a certain extent, however this comes with a price of lower link prediction performance. Symbolic approaches that learn rules perform on-par or better than latent knowledge graph embedding approaches when choosing the right aggregation function [22]. One advantage they have for the application of fact checking is that 1) they are explainable by default, meaning that you can explain predictions in terms of which rule generated the prediction and 2) they learn some types of rules which are inductive by default.

Feature Aggregation using GNNs. Embeddings of nodes in a knowledge graph are useful for downstream tasks such as clustering, visualization or classification. Often nodes of a knowledge graph have other features, such as as text embeddings or outputs of text classification models. In order to fuse features of the graph and node features, we propose to use a graph neural network as an encoder for the embeddings. These embeddings can be trained using a knowledge graph embedding scoring function such as [32]. In the context of civil security, we created high-level node embeddings from the previously introducted GDELT knowledge graph. For example, the node features of type *News article* consist of the embedding of the article text and outputs of 14 text classification models which classify texts based on a variety of criteria in the civil security domain [25], such as disinformation [27], toxicity [26], hate speech [6], and sexism [4].

Explainability & Knowledge Infusion. Deep learning (DL) models have become progressively challenging to understand and are therefore also referred to as black-boxes [17]. Since those models are often used in text analysis in the civil security area - and often have issues with specialized domain knowledge [16] - the models should be more comprehensible. Through methods of explainable artificial intelligence (XAI) the models themselves and their outputs can be understood better [12]. This can be achieved through the incorporation of knowledge graphs [16] into the ML models (i.e. attention mechanism or information loss[29]). Many approaches have been proposed to achieve this infusion: 1) concatenating KG embeddings[13] or entities[18] with content representations and metadata [13]; 2) incorporating knowledge with an additional visible matrix where each sentence is enhanced by information extracted from a KB [16]; 3) enhancing BERT with KG embeddings about author information [21]; 4) add reasoning datasets in the domain of few-shot hate speech detection [1]; 5) combining WordNet/ConceptNet with BERT model to deal with ambiguous entities [8]. Especially in domains, such as toxicity and hate speech, external knowledge can improve the ability to derive explanations and from DL models, which also helps the comparison with claims and factual information in news articles [29]. Using KGs in DL models promotes interpretability by presenting the results in information visualizations [9]. We use KGs not only for infusing their embeddings to the classification models during and before training but also use the visual representation of information from the GDELT news articles to make them more understandable to end-users [25, 28].

3 System Design

In this section we outline different software components that can be used in various applications to utilize knowledge graphs. These components are shown in Figure 1: A central *Data Storage*, *Data Providers* and *Data Enrichment* modules. Applications can consume data from the *Data Storage*, interact with the *Data Enrichment* pipelines or act as building blocks for more complex applications. Below, we briefly describe some of the software components we have used to implement this architecture.

Data Providers include applications which can fetch data from APIs or extract data from web pages. For example, to usage of the texts of the news articles provided by the GDELT dataset, requires the extraction of the text from the web page as GDELT only provides links. Another example would be SPARQL[1] queries on DBPedia via its SPARQL REST API.

Data Storage includes (1) a document-oriented database for storing raw documents and metadata, e.g. crawling results; (2) a vector-oriented database for storing features and vector representations, e.g. embeddings and classification results; (3) a graph-oriented database for storing general-purpose and extracted KGs. In our system we use (1) MongoDB[2], (2) Milvus[3] and (3) Neo4j[4].

Data Enrichment modules are designed to extract/infer additional features. These include vector representations (embeddings) from language models, classification labels from a custom set of ML-models and explanations for their predictions using XAI. As described in section 2, further modules

[1]https://www.w3.org/TR/sparql11-query/
[2]https://www.mongodb.com/
[3]https://milvus.io/
[4]https://neo4j.com/

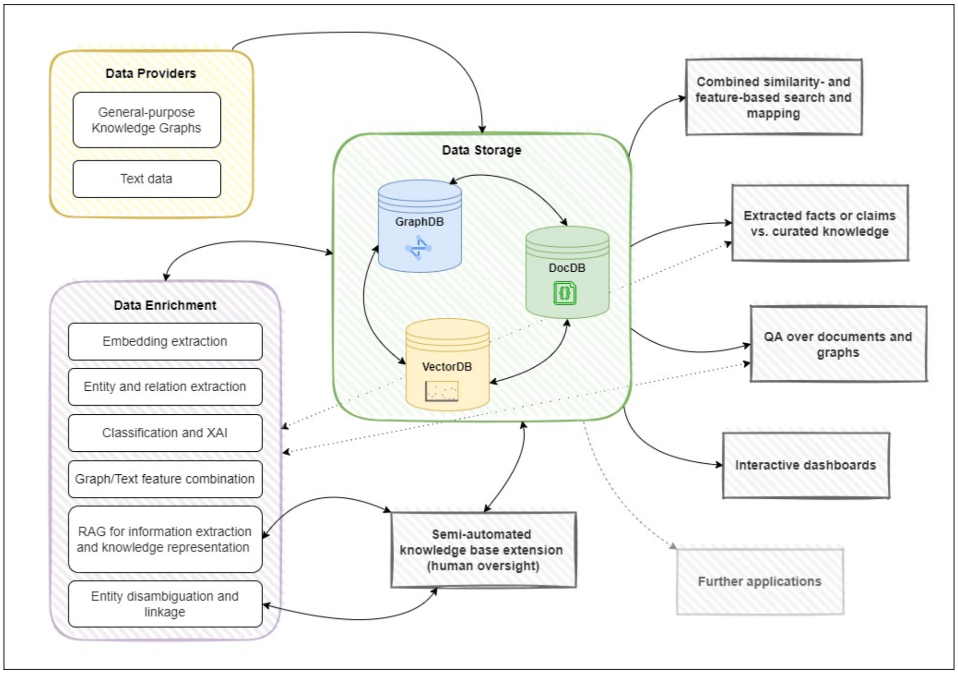

Figure 1: Architecture scheme.

extract structured knowledge from text, which includes entity recognition, disambiguation and linkage to existing knowledge bases, as well as extracting structured knowledge in accordance with case-specific ontologies via RAG-based information extraction. Another module is used to enable high-level concept modeling by aggregating text, graph and other features in a single embedding using GNNs. These modules are mainly built with Python, utilizing the PyTorch ML-framework and well-known frameworks like HuggingFace and LangChain for feature extraction. For LLM serving we use Ollama. We use PyTorch Geometric, a popular machine learning library for learning GNNs. The *Application layer* is built upon the described components and utilizes hybrid search capabilities of the Data Storage to query data based on similarity measures, metadata and further features like labels and entities. Some applications may trigger Data Enrichment pipelines, as for semi-automated knowledge base extension, where the extracted information has to be requested and approved by human, or for extraction and comparison of claims with curated knowledge. Other applications may only fetch and represent data from the Data Storage, e.g. interactive dashboards. More complex applications combine and connect simpler modules to enable advanced workflows. In the application layer, we build the front-end with SvelteKit[5] and use OpenAPI[6] to interact with the back-end processing, which is mainly Python-based.

4 Conclusion

In this paper, we introduced tasks on the creation and utilization of knowledge graphs in the context of civil security and described the design of a software system for knowledge graphs, textural data and other features such as embeddings and metadata.

Acknowledgments and Disclosure of Funding

This work is funded by the Austrian Research Promotion Agency (FFG) projects "HYBRIS" (grant no. 44155055), "DAGMAR" (grant no. 52224305) in the Austrian security research program KIRAS of the Federal Ministry of Finance (BMF)), and partially funded by the Austrian Defense Research Programme FORTE of the Federal Ministry of Finance (BMF) through project BOOST.

[5]https://kit.svelte.dev/
[6]https://www.openapis.org/

References

[1] Badr AlKhamissi, Faisal Ladhak, Srinivasan Iyer, Veselin Stoyanov, Zornitsa Kozareva, Xian Li, Pascale Fung, Lambert Mathias, Asli Celikyilmaz, and Mona Diab. ToKen: Task decomposition and knowledge infusion for few-shot hate speech detection. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2109–2120, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.

[2] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Golbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber, and Philippe Cudré-Mauroux, editors, *The Semantic Web*, pages 722–735, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[3] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795, 2013.

[4] Jaqueline Böck, Mina Schütz, Daria Liakhovets, Nathanya Queby Satriani, Andreas Babic, Djordje Slijepcevic, Matthias Zeppelzauer, and Alexander Schindler. Ait_fhstp at exist 2023 benchmark: Sexism detection by transfer learning, sentiment and toxicity embeddings and hand-crafted features. In Mohammad Aliannejadi, Guglielmo Faggioli, Nicola Ferro, and Michalis Vlachos, editors, *Proceedings of the Working Notes of CLEF 2023 - Conference and Labs of the Evaluation Forum*, pages 878–890, Thessaloniki, Greece, 2023.

[5] J. Harry Caufield, Harshad Hegde, Vincent Emonet, Nomi L. Harris, Marcin P. Joachimiak, Nicolas Matentzoglu, HyeongSik Kim, Sierra A. T. Moxon, Justin T. Reese, Melissa A. Haendel, Peter N. Robinson, and Christopher J. Mungall. Structured prompt interrogation and recursive extraction of semantics (SPIRES): A method for populating knowledge bases using zero-shot learning. *CoRR*, abs/2304.02711, 2023.

[6] Christoph Demus, Jonas Pitz, Mina Schütz, Nadine Probol, Melanie Siegel, and Dirk Labudde. Detox: A comprehensive dataset for German offensive language and conversation analysis. In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 143–153, Seattle, Washington (Hybrid), July 2022. Association for Computational Linguistics.

[7] Refiz Duro, Medina Andresel, Christoph Singewald, Veronika Siska, Axel Weißenfeld, and Dražen Ignjatović. Boosting the automated information processing for reconnaissance. In *2023 IEEE International Workshop on Technologies for Defense and Security (TechDefense)*, pages 214–219, 2023.

[8] Keyur Faldu, Amit P. Sheth, Prashant Kikani, and Hemang Akabari. KI-BERT: infusing knowledge context for better language and domain understanding. *CoRR*, abs/2104.08145, 2021.

[9] Guiseppe Futia and Antonio Vetrò. On the integration of knowledge graphs into deep learning models for a more comprehensible ai—three challenges for future research. *Information*, 11(2), 2020.

[10] Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian Suchanek. Amie: association rule mining under incomplete evidence in ontological knowledge bases. In *Proceedings of the 22nd international conference on World Wide Web*, pages 413–422. ACM, 2013.

[11] Mikhail Galkin, Etienne Denis, Jiapeng Wu, and William L. Hamilton. Nodepiece: Compositional and parameter-efficient representations of large knowledge graphs. In *International Conference on Learning Representations*, 2022.

[12] Manas Gaur, Keyur Faldu, and Amit Sheth. Semantics of the black-box: Can knowledge graphs help make deep learning systems more interpretable and explainable? *IEEE Internet Computing*, 25(1):51–59, 2021.

[13] Boshko Koloski, Timen Stepišnik Perdih, Marko Robnik-Šikonja, Senja Pollak, and Blaž Škrlj. Knowledge graph informed fake news classification via heterogeneous representation ensembles. *Neurocomputing*, 496:208–226, 2022.

[14] Kalev Leetaru and Philip A Schrodt. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, volume 2(4), pages 1–49. Citeseer, 2013.

[15] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc., 2020.

[16] Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. K-BERT: enabling language representation with knowledge graph. *Proceedings of AAAI 2020*, 2020.

[17] Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. Hatexplain: A benchmark dataset for explainable hate speech detection. *CoRR*, abs/2012.10289, 2020.

[18] M. Mayank, S. Sharma, and R. Sharma. Deap-faked: Knowledge graph based approach for fake news detection. In *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 47–51, Los Alamitos, CA, USA, nov 2022. IEEE Computer Society.

[19] Christian Meilicke, Melisachew Wudage Chekol, Patrick Betz, Manuel Fink, and Heiner Stuckenschmidt. Anytime bottom-up rule learning for large scale knowledge graph completion. *Under Submission in: The VLDB Journal—The International Journal on Very Large Data Bases*, 2023.

[20] Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. Dbpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th International Conference on Semantic Systems*, I-Semantics '11, page 1–8, New York, NY, USA, 2011. Association for Computing Machinery.

[21] Malte Ostendorff, Peter Bourgonje, Maria Berger, Julián Moreno Schneider, Georg Rehm, and Bela Gipp. Enriching BERT with knowledge graph embeddings for document classification. *CoRR*, abs/1909.08402, 2019.

[22] Simon Ott, Patrick Betz, Daria Stepanova, Mohamed H. Gad-Elrab, Christian Meilicke, and Heiner Stuckenschmidt. Rule-based knowledge graph completion with canonical models. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, page 1971–1981, New York, NY, USA, 2023. Association for Computing Machinery.

[23] Ali Sadeghian, Mohammadreza Armandpour, Patrick Ding, and Daisy Zhe Wang. Drum: End-to-end differentiable rule mining on knowledge graphs. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[24] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam, editors, *The Semantic Web*, pages 593–607, Cham, 2018. Springer International Publishing.

[25] Mina Schütz. Disinformation detection: Knowledge infusion with transfer learning and visualizations. In Jaap Kamps, Lorraine Goeuriot, Fabio Crestani, Maria Maistro, Hideo Joho, Brian Davis, Cathal Gurrin, Udo Kruschwitz, and Annalina Caputo, editors, *Advances in Information Retrieval. ECIR 2023. Lecture Notes in Computer Science Vol. 13982*, pages 468–475, Cham, 2023. Springer Nature Switzerland.

[26] Mina Schütz, Christoph Demus, Jonas Pitz, Nadine Probol, Melanie Siegel, and Dirk Labudde. DeTox at GermEval 2021: Toxic comment classification. In *Proceedings of the GermEval 2021 Shared Task on the Identification of Toxic, Engaging, and Fact-Claiming Comments*, pages 54–61, Duesseldorf, Germany, September 2021. Association for Computational Linguistics.

[27] Mina Schütz, Daniela Pisoiu, Daria Liakhovets, Alexander Schindler, and Melanie Siegel. Gerdisdetect: A german multilabel dataset for disinformation detection. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Turin, Italy, 2024.

[28] Mina Schütz, Alexander Schindler, and Melanie Siegel. Disinformation detection: An explainable transfer learning approach. In *CODE2021*, Munich, Germany, 2021.

[29] Amit Sheth, Valerie L. Shalin, and Ugur Kursuncu. Defining and detecting toxicity on social media: context and knowledge are key. *Neurocomputing*, 490:312–318, 2022.

[30] Denny Vrandečić and Markus Krötzsch. Wikidata: a free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85, sep 2014.

[31] Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, Yong Jiang, and Wenjuan Han. Zero-shot information extraction via chatting with chatgpt. *CoRR*, abs/2302.10205, 2023.

[32] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.

[33] Xiaojuan Zhao, Yan Jia, Aiping Li, Rong Jiang, Haocheng Xie, Yichen Song, and Weihong Han. Multi-source knowledge fusion: A survey. In *2019 IEEE Fourth International Conference on Data Science in Cyberspace (DSC)*, pages 119–127, 2019.

[34] Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. Llms for knowledge graph construction and reasoning: Recent capabilities and future opportunities. *CoRR*, abs/2305.13168, 2023.