

Comprehensive Event Representations using Event Knowledge Graphs and Natural Language Processing

Tin Kuculo

Supervised by Prof. Dr. Wolfgang Nejdl, Prof. Dr. Elena Demidova and Dr. Simon Gottschalk
L3S Research Center, Leibniz Universität Hannover
Hannover, Germany
kuculo@L3S.de

ABSTRACT

Recent work has utilised knowledge-aware approaches to natural language understanding, question answering, recommendation systems, and other tasks. These approaches rely on well-constructed and large-scale knowledge graphs that can be useful for many downstream applications and empower knowledge-aware models with commonsense reasoning. Such knowledge graphs are constructed through knowledge acquisition tasks such as relation extraction and knowledge graph completion. This work seeks to utilise and build on the growing body of work that uses findings from the field of natural language processing (NLP) to extract knowledge from text and build knowledge graphs. The focus of this research project is on how we can use transformer-based approaches to extract and contextualise event information, matching it to existing ontologies, to build comprehensive knowledge graph-based event representations. Specifically, sub-event extraction is used as a way of creating sub-event-aware event representations. These event representations are then further enriched through fine-grained location extraction and contextualised through the alignment of historically relevant quotes.

CCS CONCEPTS

• **Computing methodologies** → *Multi-task learning*; • **Information systems** → *Ontologies*; Document representation; Question answering; • **Theory of computation** → *Structured prediction*.

KEYWORDS

knowledge graph, event extraction, event representation, quotes, event geotagging

ACM Reference Format:

Tin Kuculo. 2022. Comprehensive Event Representations using Event Knowledge Graphs and Natural Language Processing. In *Companion Proceedings of the Web Conference 2022 (WWW '22 Companion)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3487553.3524199>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WWW '22 Companion, April 25–29, 2022, Virtual Event, Lyon, France

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9130-6/22/04...\$15.00
<https://doi.org/10.1145/3487553.3524199>

1 INTRODUCTION

Knowledge graphs, structured representations of facts comprising entities, relationships, and semantic descriptions, have recently gained both academic and industrial attention [13]. By creating comprehensive machine-readable knowledge graphs about the world's entities, their semantic properties, and their relationships, remarkable possibilities open up. Potential applications include the creation of machine-readable encyclopedias that can be queried with high precision, similar to semantic databases; quick and accurate mapping of textual phrases to entities in the knowledge graph; an ability to answer natural language questions about entities such as who, where, when. Finally, world knowledge is a crucial component in providing disambiguating context for machine translation and can be a catalyst for the acquisition of further knowledge, as well as for the growth and maintenance of the knowledge graph [35].

Knowledge graphs mainly cover static information such as a person's date of birth and birthplace, illustrating the state of the world rather than changes over time. However, considering the dynamic nature of events can be very important for the reconstruction of someone's past or the extensive past of local industries, regions, or organisations. It is, therefore, necessary to build a different type of structured database based on events rather than entities and entity-based real facts. Capturing this dynamic knowledge requires considering events as the unit for storing knowledge [31]. This need has resulted in the study and creation of event knowledge graphs.

There are two different approaches to understanding and processing events, one more from the Semantic Web and one more from the NLP perspective. In knowledge graphs such as DBpedia and Wikidata, typically, events are clearly distinguished and named events such as Brexit, FIFA World Cup 2014. This is often seen in knowledge graphs. The second view is that of fine-grained events found in text, often as sub-events in a bigger context. For example, a text about Brexit might cover the visit of the UK's prime minister Theresa May in Berlin: this is the approach often used in NLP. In my work, I want to bring these two approaches together. For example, consider a properly modelled *Visit* event, which is linked to the *Brexit* event in a knowledge graph. To extract and model this kind of information, we use Wikipedia as our source of unstructured information and Wikidata to structure the outputs of our NLP components.

In this work, by utilising event extraction, we find new events and form their representations. Event representations can still benefit from additional contextualisation. We provide this additional context in two ways:

(i) Quotes: while static information such as the time and location of a given event is understood across languages, the impact of a given event on specific populations varies across linguistic and cultural contexts of those populations. In our work, we argue that this impact can be captured and analysed across time through the lens of quotes. Specifically, the quotes said by persons of public interest are often relevant from the historical perspective. Examples include *"Tear down this wall"* by the United States President Ronald Reagan in the context of the fall of the Berlin Wall.

(ii) Fine-grained locations: another perspective is the provision of locations with higher precision. While knowledge graphs often provide locations of events, they are often not fine-grained. For example, for an event such as the 2022 Winter Olympics, knowledge graphs might only contain the country, China, or the city, i.e., Beijing, but not the actual competition sites. In this example, we want to do better and find the exact venues, such as the Beijing Wukesong Sports Center.

2 PROBLEM

Given a set of named events, we want to extract all relevant knowledge about them and contextually interlink them to persons and their quotes. This is accomplished through three separate components.

- Sub-event Extraction: A processing pipeline is created such that for each event, all of its sub-events are extracted and linked with its entity arguments to corresponding Wikidata objects and classes.
- Fine-Grained Location Extraction: For additional contextualization of events, OpenStreetMap is used to link geographic entities present in the text of an event's Wikipedia page. Through a representation learning approach, correct geographic entities are then selected as the most precise available fine-grained location information about the event.
- Quote Extraction and Cross-lingual Alignment: The third processing pipeline utilizes Wikiquote – an online collection of quotes¹. Using Wikiquote as a data source, we focus on Wikiquote pages about persons. The quotes by these persons are then extracted and interlinked with their contexts. Finally, encoding the quotes using a cross-lingual language model, we cluster them together across languages.

3 PROPOSED APPROACH

In this section, we briefly present our approach to the specific tasks we focus on. These tasks are event extraction, and event enrichment and contextualization. The general approach can be seen in Figure 1.

For the event extraction, a set of events of interest is defined. For the selected events, we take their corresponding Wikipedia pages and run our sub-event extraction pipeline. The pipeline consists of three components: the event type extraction, the argument extraction component, and the knowledge modelling component that matches our extracted information to the Wikidata ontology.

To enrich and contextualize our events, we take on two tasks. In the first, we address the extraction of fine-grained locations. While some events' locations can be accurately described by the names of countries or cities they have occurred in, others would benefit from

more fine-grained location information, e.g., the specific building in which the event occurred instead of the city. To increase the precision of event location information, we take a mixed approach of relation extraction, entity linking, and link prediction. Lastly, for the contextualization of events, we extract quotes of public figures and interlink them with event and entity information.

4 METHODOLOGY

In this section, the main components of the research project will be described—namely, the event extraction and the contextualisation and enrichment of events.

4.1 Event Extraction

In the first step, a set of events is selected. For each of these events, we take their corresponding Wikipedia page and run an event extraction model [22]. The event extraction component runs on each paragraph of the Wikipedia page and outputs a set of sub-events found in the text. For example, consider the following excerpt in Figure 2:

As of 20 January 2022, at least 1,488 protesters and bystanders, have been shot and killed by police forces and at least 8,702 people detained.

Figure 2: Taken from the English Wikipedia page on the Myanmar Protests and modified for simplicity.²

The event extraction component identifies two sub-events: a *Die* sub-event and an *Arrest* sub-event. In the following step, these are mapped to Wikidata classes *death* (Q4)³ and *legal action* (Q2135540). We extract all their subclasses to be candidate Wikidata classes representing our subevents. Among the subclasses of *death* (Q4) there exists a Wikidata class *extra-judicial killing* (Q2717573), similarly among the subclasses of *legal action* (Q2135540) there is a *detention* (Q1071447). For each of the candidate subclasses, we generate a list of questions to be asked by the question answering model [16]. According to the returned answers, we select the correct Wikidata class to be assigned as the sub-event type.

To evaluate the performance of this processing pipeline, we generate a test dataset based on existing sub-events in Wikipedia texts. Specifically, here we are referring to sub-events already defined in Wikidata.

For example, in the sentence shown in Figure 3, the 2016 Republican National Convention is an existing event in Wikidata, with properties such as location, country, point in time. In addition to this automatically generated test dataset based on Wikipedia links, we also plan to manually annotate a selected subset of Wikipedia event articles. To evaluate, we compare our model's extracted properties with the ones present in Wikidata. This evaluation procedure is done in tandem on different baselines. Firstly, to validate the need for the question answering component, we ablate the question answering part of the pipeline and match ACE ontology events directly to Wikidata. This is followed by an evaluation of a SOTA triple extraction model on existing sub-events. Finally, we compare

¹https://en.wikiquote.org/wiki/Main_Page

²https://en.wikipedia.org/wiki/2021-2022_Myanmar_protests

³<https://www.wikidata.org/wiki/Q4>

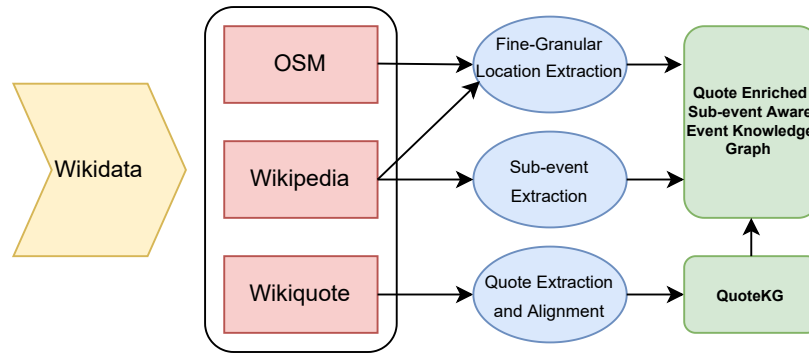


Figure 1: General structure of the research project.

On July 19, the second night of the *2016 Republican National Convention*, Pence won the Republican vice presidential nomination by acclamation.

Figure 3: Example sentence in the Wikipedia article on the 2016 Republican Party vice-presidential candidate selection. Links are italicised.

our full approach to an approach that extracts entities and does triple extraction on text.

4.2 Event Contextualization and Enrichment

To further enhance our event representations, we look to two directions: increasing the precision of our location information and contextualising events through related quotes of persons of public interest.

4.2.1 Fine-Grained Location Extraction. To further enrich event representations, a processing pipeline is used to fill empty location nodes and retrieve more precise location information for existing ones. Specifically, we use a combined approach of link prediction and location extraction. A multitask framework [32] is used to train a model to jointly predict locations based on existing nodes in EventKG [10] and extract fine-grained location information from Wikipedia texts of corresponding events. The fine-grained location information is determined via entity-linking to OpenStreetMap (OSM) using WorldKG [7], and a relation extraction model trained to extract candidate locations and rank them to determine the correct location.

4.2.2 Quote Extraction and Cross-lingual Alignment. To support the research of events through the lens of impactful quotes by public figures, we align quotes across languages and enrich them with context, creating the first multilingual knowledge graph of quotes, *QuoteKG* [15]. *QuoteKG* is based on Wikiquote, another wiki-based project. Like Wikipedia, Wikiquote has independent versions for different languages. It consists of pages, each of which covers a particular topic and is divided into subsections and sections. For *QuoteKG*, we focus on Wikiquote pages about persons that contain quotes attributed to them.

QuoteKG's creation pipeline, shown in Figure 4, begins by processing Wikiquote in all available languages and selecting the pages about specific persons.

After we identify and enrich the quotes, we need to detect which of them represent mentions of the same quote said by a person of public interest. This task of cross-lingual alignment of quote mentions is treated as a clustering task at the end of which each cluster represents a quote with a set of mentions.

The clustering task is performed for each individual person separately. Given a person's quote mention in a set of languages, the aim is to create clusters of highly similar mentions. To derive a similarity between two mentions, potentially from different languages, we compute the cosine similarity of sentence embeddings derived from the mentions' texts. As an embedding model, a language-agnostic transformer model pre-trained on millions of multilingual paraphrase examples is used, namely, XLM-RoBERTa [5]. Given these embeddings and the cosine similarity function, clustering is performed by detecting communities of quotes using a nearest-neighbour search. To do so, we utilise UKPLab's Fast Clustering algorithm⁴ as it is optimised for efficient similarity computations of our embeddings.

5 PRELIMINARY RESULTS

So far, we have created the *QuoteKG* schema based on schema.org as well as a pipeline that extracts quotes from the Wikiquote corpus and aligns them across languages. The *QuoteKG*'s schema was made to model quotes, their relationships with persons and other entities, as well as their different mentions, e.g., translations, typically in different contexts. To this end, *QuoteKG* is based on an extension of the *schema.org* vocabulary that provides a `so:Quotation`⁵ class which is re-used. *QuoteKG* has been made publicly available⁶ and includes nearly one million quotes in 52 languages, said by nearly 69,000 people of public interest. Table 1 provides detailed statistics for selected languages.

For the sub-event extraction framework, the question answering pipeline has been completed. We have established the mapping between ACE event types and Wikidata classes. In addition, the

⁴<https://github.com/UKPLab/sentence-transformers/blob/master/examples/applications/clustering/>

⁵Prefix "so:" stands for <https://schema.org/>

⁶<https://quotekg.l3s.uni-hannover.de/>




Figure 4: Pipeline to create *QuoteKG* from Wikiquote.

Table 1: Statistics of selected languages in *QuoteKG*.

Language	Persons	Quotes	Mentions	Mentions with Contexts
English	19,076	267,994	271,879	193,848
Italian	18,795	145,258	146,128	48,107
German	3,463	16,017	16,454	4,330
Croatian	2,720	11,019	12,944	2,045
Welsh	241	457	503	247
All Languages	69,539	882,969	964,009	412,801

entity linking and the question answering part of the pipeline has been built as well. So far, we have identified many cases of potential new sub-events, for example, consider the following sentence taken from the Wikipedia article on *Battle of Latakia*:

“the Israeli ships engaged a 560-ton Syrian T43-class minesweeper”

While a sub-event extraction model might detect an *Attack* sub-event, our pipeline detects *naval battle* (*Q1261499*).

Lastly, although the fine-grained location extraction is still in an early stage of development, we find the research well-motivated. After our data collection, we have found 700,000 events, of which 500,000 have locations, but only 50,000 have coordinates.

6 STATE OF THE ART

In this section, the latest research in event extraction is addressed, prominent event knowledge graphs, and work done on geolocation. This is followed by an overview of other knowledge graphs containing quotes.

6.1 Event Knowledge Representation

Event knowledge graphs have come about as a way to address the dynamic nature of the real world and as a way to capture the temporal dimension of relations. This is shown in initial work on temporal knowledge graphs such as the Integrated Conflict Early Warning System (ICEWS) dataset and the Global Database of Events, Language and Tone (GDEL). Recently, more work has been done in the financial domain and work motivated by the news domain [1, 3]. Work that steps out of domain-specific focus [10] focuses on the big-picture event properties such as the time, location, and participants of a particular event, not delving into the specific sub-events of an event. One notable exception is the NewsReader project [31] where the authors focus on news articles and use a synset detection based approach, whereas, in my work, the focus is creating enriched representations of important events that contain sub-events as they relate to them.

6.2 Event Extraction

Event extraction, a challenging research task in information extraction, is at the heart of information extraction research. Event

extraction effectively converts unstructured plain texts into structured event information that describes “who, when, where, what, why” and “how” of real-world events. Traditional event extraction methods mainly relied on complex feature engineering and faced error propagation issues. Some of these issues are ameliorated using deep learning approaches. The early deep learning approaches [4, 28] used Convolutional Neural Networks (CNN) to automatically extract lexical and sentence-level features. Other research explores RNN-based architectures. JRNN [27] is proposed with a bidirectional RNN for event extraction. While previous approaches relied heavily on language-specific knowledge and existing NLP tools, [9] show that learning automatically from data can result in accurate multilingual event detection.

Sentence-level sequential modelling suffers from the low efficiency in capturing very long-range dependencies. The modelling of structural information in the attention mechanism has gradually attracted the attention of researchers. The JMEE [21] uses a graph convolution network based on attention to jointly model the graph information to extract multiple event triggers and arguments. [26] investigate a CNN based on dependency trees to perform event detection and are the first to integrate syntax into neural event detection. The network explicitly models the information from entity mentions to improve performance for event detection.

Since the arrival of transformer models, research has shown its efficacy on event extraction [36, 37]. [34] is a BERT-based framework that models text spans and captures within-sentence and cross-sentence context. Similar to our own work, [8, 19, 20] frame event extraction as a question answering problem.

6.3 Event Geolocation

The task of event geolocation has most recently been studied as a part of event detection in social media analysis [23, 24]. Other research focuses on Geolocating entities in knowledge graphs [30] or geotagging text content in general [2, 14, 17]. On the other hand, work addressing geotagging of events [6, 12] tends to focus on a specific domain. While there is work focusing on fine-grained location estimation [18], to my best knowledge, this has not been applied specifically to events and OpenStreetMap geographic objects.

6.4 Quotes in Knowledge Graphs

DBQuote [29] is a system for collection and semantic annotation of quotes extracted from Twitter and Wikiquote; however, it covers only two languages (English and Korean) and has not been made available. Other work [25] focuses on English news articles, quotes, and topics to create a database intended for journalistic research. Event-centric knowledge graphs, such as EventKG [11], offer insights into human history and events of importance, but they lack quotations that provide context to public figures and their actions.

As a consequence, quotes have only had insufficient coverage in the Semantic Web: for example, Wikidata [33] has fewer than 400 instances of the class “Phrase”⁷ that are attributed to an author or creator – most of them only consisting of few words (e.g., “cogito ergo sum” and “covfe”). To the best of my knowledge, our knowledge graph of quotes, *QuoteKG*, is the first publicly available knowledge graph of quotes.

7 CONCLUSIONS

In this paper, I have described my research project and its focus on the creation of quote enriched sub-event aware event representations with very fine-grained location information. I have defined the framework to be used for this purpose and have shown some preliminary results. These include *QuoteKG*, a multilingual knowledge graph of quotes, the state of event extraction framework, and the statistics motivating my work on fine-grained event extraction.

8 ACKNOWLEDGEMENTS

I thank my supervisors, Prof. Dr. tech. Wolfgang Nejdl, Prof. Dr. Elena Demidova, and Dr. Simon Gottschalk for their feedback and support. This work has been funded by European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 812997 (Cleopatra).

REFERENCES

- [1] Tareq Al-Moslimi and Marc Gallofré Ocaña. 2020. Lifting news into a journalistic knowledge platform. (2020).
- [2] Marco Avvenuti, Stefano Cresci, Leonardo Nizzoli, and Maurizio Tesconi. 2018. GSP (Geo-Semantic-Parsing): geoparsing and geotagging with machine learning on top of linked data. In *European Semantic Web Conference*. Springer, 17–32.
- [3] Rudnik Charlotte, Ehrhart Thibault, Ferret Olivier, Teyssou Denis, Troncy Raphael, and Tannier Xavier. 2019. Searching news articles using an event knowledge graph leveraged by Wikidata. In *International World Wide Web Conference Committee*.
- [4] Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. Event extraction via dynamic multi-pooling convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. 167–176.
- [5] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised Cross-lingual Representation Learning at Scale. *arXiv preprint arXiv:1911.02116* (2019).
- [6] Agung Dewandaru, Dwi Hendratno Widyantoro, and Saiful Akbar. 2020. Event Geoparser with Pseudo-Location Entity Identification and Numerical Argument Extraction Implementation and Evaluation in Indonesian News Domain. *ISPRS International Journal of Geo-Information* 9, 12 (2020), 712.
- [7] Alishiba Dsouza, Nicolas Tempelmeier, Ran Yu, Simon Gottschalk, and Elena Demidova. 2021. WorldKG: A World-Scale Geographic Knowledge Graph. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4475–4484.
- [8] Xinya Du and Claire Cardie. 2020. Event extraction by answering (almost) natural questions. *arXiv preprint arXiv:2004.13625* (2020).
- [9] Xiaocheng Feng, Bing Qin, and Ting Liu. 2018. A language-independent neural network for event detection. *Science China Information Sciences* 61, 9 (2018), 1–12.
- [10] Simon Gottschalk and Elena Demidova. 2018. EventKG: a multilingual event-centric temporal knowledge graph. In *European Semantic Web Conference*. Springer, 272–287.
- [11] Simon Gottschalk and Elena Demidova. 2019. EventKG—the Hub of Event Knowledge on the Web—and Biographical Timeline Generation. *Semantic Web* 10, 6 (2019), 1039–1070.
- [12] Andrew Halterman. 2019. Geolocating political events in text. *arXiv preprint arXiv:1905.12713* (2019).
- [13] 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* (2021).
- [14] Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Ioannis Kompatsiaris. 2017. Geotagging text content with language models and feature mining. *Proc. IEEE* 105, 10 (2017), 1971–1986.
- [15] Tin Kuculo, Simon Gottschalk, and Elena Demidova. 2022. QuoteKG: A Multilingual Knowledge Graph of Quotes. In *European Semantic Web Conference*. Springer.
- [16] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942* (2019).
- [17] Egoitz Laparra and Steven Bethard. 2020. A Dataset and Evaluation Framework for Complex Geographical Description Parsing. In *Proceedings of the 28th International Conference on Computational Linguistics*. 936–948.
- [18] Jennifer Lautenschlager, James Starz, and Ian Warfield. 2017. A Statistical Approach to the Subnational Geolocation of Event Data. In *Advances in Cross-Cultural Decision Making*. Springer, 333–343.
- [19] Fayuan Li, Weihua Peng, Yuguang Chen, Quan Wang, Lu Pan, Yajuan Lyu, and Yong Zhu. 2020. Event extraction as multi-turn question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 829–838.
- [20] Jian Liu, Yubo Chen, Kang Liu, Wei Bi, and Xiaojiang Liu. 2020. Event extraction as machine reading comprehension. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1641–1651.
- [21] Xiao Liu, Zhunchen Luo, and Heyan Huang. 2018. Jointly multiple events extraction via attention-based graph information aggregation. *arXiv preprint arXiv:1809.09078* (2018).
- [22] Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. Text2event: Controllable sequence-to-structure generation for end-to-end event extraction. *arXiv preprint arXiv:2106.09232* (2021).
- [23] Stuart E Middleton, Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Yiannis Kompatsiaris. 2018. Location extraction from social media: Geoparsing, location disambiguation, and geotagging. *ACM Transactions on Information Systems (TOIS)* 36, 4 (2018), 1–27.
- [24] Paul Mousset, Yoann Pitarch, and Lynda Tamine. 2020. End-to-End Neural Matching for Semantic Location Prediction of Tweets. *ACM Transactions on Information Systems (TOIS)* 39, 1 (2020), 1–35.
- [25] Chris Newell, Tim Cowlshaw, and David Man. 2018. Quote extraction and analysis for news. In *Proceedings of the Workshop on Data Science, Journalism and Media, KDD*. 1–6.
- [26] Thien Nguyen and Ralph Grishman. 2018. Graph convolutional networks with argument-aware pooling for event detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [27] Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 300–309.
- [28] Thien Huu Nguyen and Ralph Grishman. 2015. Event detection and domain adaptation with convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*. 365–371.
- [29] Guangyuan Piao and John G Breslin. 2015. DBQuote: A Social Web based System for Collecting and Sharing Wisdom Quotes. In *Proceedings of the 5th Joint International Semantic Technology Conference, Poster and Demonstrations*.
- [30] Peiyuan Qiu, Jialiang Gao, Li Yu, and Feng Lu. 2019. Knowledge embedding with geospatial distance restriction for geographic knowledge graph completion. *ISPRS International Journal of Geo-Information* 8, 6 (2019), 254.
- [31] 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. *Journal of Web Semantics* 37 (2016), 132–151.
- [32] George Stoica, Emmanouil Antonios Platanios, and Barnabás Póczos. 2020. Improving Relation Extraction by Leveraging Knowledge Graph Link Prediction. *arXiv preprint arXiv:2012.04812* (2020).
- [33] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a Free Collaborative Knowledge Base. *Commun. ACM* 57, 10 (2014), 78–85.
- [34] David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. *arXiv preprint arXiv:1909.03546* (2019).
- [35] Gerhard Weikum and Martin Theobald. 2010. From information to knowledge: harvesting entities and relationships from web sources. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 65–76.
- [36] Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5284–5294.
- [37] Tongtao Zhang, Heng Ji, and Avirup Sil. 2019. Joint entity and event extraction with generative adversarial imitation learning. *Data Intelligence* 1, 2 (2019), 99–120.

⁷<https://www.wikidata.org/wiki/Q187931>