

Empowering COVID-19 Fact-Checking with Extended Knowledge Graphs

Paolo Mengoni[1(✉)] and Jinyu Yang[2]

[1] School of Communication and Film, Hong Kong Baptist University, Hong Kong, China
pmengoni@hkbu.edu.hk
[2] Hong Kong Baptist University, Hong Kong, China
18252591@life.hkbu.edu.hk

Abstract. During the COVID-19 outbreak, fake news regarding the disease have spread at an increasing rate. Let's think, for instance, to face masks wearing related news or various home-made treatments to cure the disease. To contrast this phenomenon, the fact-checking community has intensified its efforts by producing a large number of fact-checking reports. In this work, we focus on empowering knowledge-based approaches for misinformation identification with previous knowledge gathered from existing fact-checking reports. Very few works in literature have exploited the information regarding claims that have been already fact-checked. The main idea that we explore in this work is to exploit the detailed information in the COVID-19 fact check reports in order to create an extended Knowledge Graph. By analysing the graph information about the already checked claims, we can verify newly coming content more effectively. Another gap that we aim to fill is the temporal representation of the facts stored in the knowledge graph. At the best of our knowledge, this is the first attempt to associate the temporal validity to the KG relations. This additional information can be used to further enhance the validation of claims.

Keywords: Fact checking · Knowledge graphs · Natural Language Processing

1 Introduction

During the current COVID-19 outbreak, we are witnessing a new wave of misinformation spreading on different media. Social media, forums and traditional news outlets are evenly affected by the diffusion of fake news, disinformation and misinformation. On the opposite side, the efforts from the scientific community have brought to a wide range of initiatives in fact-checking and veracity assessment.

The common approaches to misinformation identification involve professional fact-checkers or experts to produce reports that assess the veracity (i.e. veracity

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
O. Gervasi et al. (Eds.): ICCSA 2022 Workshops, LNCS 13377, pp. 138–150, 2022.
https://doi.org/10.1007/978-3-031-10536-4_10

level) of the rumours. Those fact-checkers rely on their expertise, that can vary from journalism to the medical fields, to annotate the misleading content. On the other hand, computational approaches for misinformation are introduced to tackle the increased volume and speed of fake news creation. These approaches can be categorized in four different categories: (i) knowledge-based methods that classify the veracity of content on the basis of pre-existing knowledge bases; (ii) style-based approaches that analyse the information semantic and syntactic features; (iii) credibility-based techniques that take into account the credibility of the content spreading subjects; (iv) propagation-based approaches that exploit social network analyses to understand the content propagation [10,12]. The first two categories focus on analysing the content of the news, while the latter two rely on the analysis of the medium where the content spreads.

The result of this process - manual or automatic - is usually a fact-check report that contains a classification of claims given a specific assessment label. The veracity assessment can vary from the simple assignment of {True; False} labels, to a more detailed classification using {True, Partially True, False, Unsubstantiated}, to fine-grain classification like in *Snopes* that use an extended set of labels including {True; Mostly true; Mixture; Mostly false; False; Unproven; Outdated; Miscaptioned; Correct attribution; Misattributed; Scam; Legend, Labeled Satire, Lost Legend}. The fact-check reports also contain details about the assessment process and include a wide range of information pertaining to the assessed claim and related facts.

In this work, we focus on a novel knowledge-based approach that will take in account multiple sources to assess the veracity level of a new claim. The knowledge base is then transformed to a Knowledge Graph (KG): a graph structure that represents the facts in a knowledge base as nodes and their relationships as edges. Usually the content of the KG is assumed to be true, while it is not possible to make any assumption on the information that is absent. In our approach we will associate a veracity label to the facts, gathering this information from the existing fact-checking reports. This kind of approach has been rarely explored in literature and will be exploited to create a specialized, thematic KG to support fact-checking.

The research questions that we will try to answer can be summarised as follows: (i) how to transform a fact-checking report into a Knowledge Graph, (ii) how to integrate temporal information into the Knowledge Graph, and (iii) how the Knowledge Graph can be used to support fact-checking.

The rest of the paper is organized as follows. In Sect. 2, we discuss related works using Knowledge Graphs to support fact-checking. In Sect. 3, we introduce in detail the methodology used for this work. Section 4 reports the results of our approach and in Sect. 5 we analyse and discuss them. In Sect. 6, we draw conclusions and introduce future works.

2 Related Works

The knowledge-based approaches to fact-checking, like the one we propose in this work, usually rely on pre-existing knowledge bases that contain general

information. The knowledge bases may contain information stored in structured way (e.g. like in DBPedia[1]), semi-structured (e.g. as in Wikipedia[2]), or unstructured like in other open access repositories and ontologies. Some techniques will directly rely on the knowledge base content of one (single-source approaches) or more (open-source approaches) to assess the truthfulness of a claim.

The knowledge base can be transformed into a Knowledge Graph (KG), a graph structure that represents the facts in a knowledge base as nodes and their relationships as edges. Facts in the knowledge base are usually stored as sets of (Subject, Predicate, Object) (SPO) triples. The Subject and Object of the SPO triples will become the nodes of the knowledge graph, while the Predicate will be the label associated with the edges between the nodes. The Predicate will define the relationship between the nodes [1]. Metadata will be associated with the relationship between the nodes and usually will define the relationship type and other additional information deemed needed for the fact-checking task. The metadata annotations can be crowdsourced [6], manually annotated by experts [5], or automatically generated by using Natural Language Processing (NLP) techniques [11].

Some assumptions are made on the content of the KG. In general, the information present in the graph is always considered to be True, while absent SPO triples may take any possible veracity value [12]. Given these assumptions, to support fact-checking there has been an emerging need to add additional information to the knowledge graphs. Recent works associated to the KG metadata the truth values of the statements as assessed by fact-checking organisations (such as *Snopes*[3], AFP Fact Check[4], etc.) [11].

The KG information is then used to support various fact-checking strategies, such as the Entity Location, Relation Verification, and Knowledge Inference [12]. Given a claim, the fact-checking goal is to evaluate the probability (or possibility) that an edge exists or may exist in the KG [8,9].

The ClaimsKG [11] is a similar work that is focusing more on Fact Retrieval, while our approach is focusing more on structuring the KG to empower the Relation Verification and, in future works, Knowledge Inference tasks. The novel approach we here introduce, very less explored in literature until recently [7], is motivated from making good use of the previously fact-checked claims to assess the veracity of new ones.

3 Methodology

In this work we will explore the Relation Verification task. This aspect of the knowledge-based fact checking is using information in the KG to verify newly coming content.

[1] https://wiki.dbpedia.org/.
[2] https://www.wikipedia.org/.
[3] https://www.snopes.com/.
[4] https://factcheck.afp.com/.

In the approach introduced in this work, the fact-checking reports are transformed to an extended KG *FactsKG* that can be used to assess claim veracity using various knowledge-based techniques. The approach we are introducing here is general but given the availability of data and the importance of the topic, we focus on COVID-19 related claims and the associated fact-checking reports.

By analysing the fact-checking reports, we will extract facts that have been already assessed by fact-checking organisations. These facts will be stored in the KG in the form of SPO triples. Unlike many other KG, the graph that we are going to generate will include facts with various veracity levels as assessed in the fact-checking reports. This is one of the most innovative aspect of our approach, that will extend the classical KG definition by adding additional information about the facts in the metadata of the SPO triples. The edge metadata will store the relation veracity that has been assessed by the fact checking reports. Additionally, *FactsKG* will also include temporal information as the relationship between Subject and Object may vary during time. Fact-checkers will be able to immediately and effectively assess newly coming claims using directly the *FactsKG* or through automated Relation Verification techniques.

The methodology used in this work includes the fact-checking report dataset creation, SPO generation from the dataset, and finally the transformation of SPO triples into *FactsKG*.

The first step of the process is the creation of a dataset from previously produced fact-checking reports. This task involves the retrieval and collection of public fact-checking reports from fact-checking websites. We implemented a scraper for obtaining automatically the information from the fact-checking websites, using Python and the Beautifulsoup library. After analysing the data scraped from the various fact-checking websites, we implemented the dataset schema including information about the report *retrieved date*, *url*, *title*, *publisher name*, *publisher website*, *textual veracity rating*, *language code*, *report text*, *claimant*, *claim date and time*. During analysis and early testing of the methodology we have found that we should save a permanent copy of the reports. This was needed to keep our dataset consistent as some reports were updated and deleted by the fact-checking organisations. The permanent copy of the reports found by our scraper has been saved online using Internet Archive[5] and PermaCC[6] services and stored in the dataset field *claim permanent url*. An additional offline copy has been saved offline as html file and stored in the dataset field *report permanent html filename*.

Once the fact-checking reports have been collected, we perform the second step of our methodology. By using Natural Language Processing (NLP) techniques we automatically transform the fact-checking reports in (Subject, Predicate, Object) (SPO) triples. The dataset is analysed through the NLP pipeline included in the *spaCy* Python library [2]. The process includes stop word removal, tokenization, token transformation to vector using convolutional neural networks (*tok2vec*), part of speech tagging, sentence segmentation, dependency

[5] https://archive.org/.
[6] https://perma.cc/.

parsing, sentence entity recognition, named entity recognition, attribute ruler, and lemmatization. All the NLP steps allow to extract informative SPO triples from the fact checking reports, thus transforming the unstructured textual data in a structured form. Contextually extended metadata will be associated with the SPO triples. This includes the triples' veracity level, temporal information and other features deemed to be useful for future tasks.

The last step of our methodology is the construction of the Knowledge Graph (KG). First we analysed and designed various possible solutions for the KG structure and its storage. For the data storage we selected the Neo4J[7] database management system that supports native graph data storage. In the analysis phase we explored the transformation of the SPOs in the dataset with a 2-nodes and a 3-nodes graph model. The 2-nodes model involves the storage of Subject and Object as nodes of the graph, while the Predicate is the label of the edge between the pair of nodes. In the 3-nodes model each element is stored as a node with Subject/Object nodes and Predicate nodes. The Subject node have outbound connections to Predicate nodes and the Object nodes have inbound connection from the Predicate nodes. The Predicate edges and nodes attributes, for the 2-nodes and 3-nodes respectively, will contain the veracity level, temporal information, and other features.

4 Results

Working with real-life data is always difficult. During the phases of data collection, cleaning, transformation, and knowledge graph creation we faced various challenges.

4.1 Dataset Creation

The structure of the fact-checking reports is different for every website, making difficult to access the information from automated systems. As shown in Fig. 1, AFP Fact check report[8] is a textual report with a simple structure as shown in Fig. 1a, that makes the collection task easy. However, there is no clear textual indication of the claim veracity assessment. AFP is using different colours (i.e., red, orange, etc.) for the cross drawn over the assessed claim to display the veracity of the claim, as shown in Fig. 1b. In this work, we extract the veracity information from the text, but in this case it is more difficult to give a clear classification of AFP reports.

On the other hand, as shown in Fig. 2a, the content of the fact-check report from *Snopes*[9] is more structured, requiring a different scraping strategy. How-

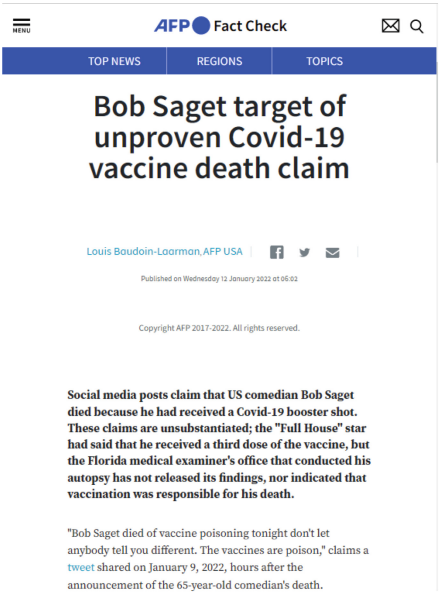
[7] https://neo4j.com.

[8] Available at the address: https://factcheck.afp.com/http%253A%252F%252Fdoc.afp.com%252F9VY4M3-1.

[9] Permanent copy available at the address: https://web.archive.org/web/20220315225545/https://www.snopes.com/fact-check/pfizer-covid-19-vaccine-fetal-cells/.

ever, the veracity of the claim is clearly displayed in a specific section of the report as shown in Fig. 2b.

During the scraping process we faced also the problem that some reports structure is not consistent, also internally to each source. Moreover, the reports structure is changing during time requiring manual intervention for adjusting the scraping parameters. Moreover, some of the reports are missing some parts requiring specific error management strategies.



(a) Title and content of the report



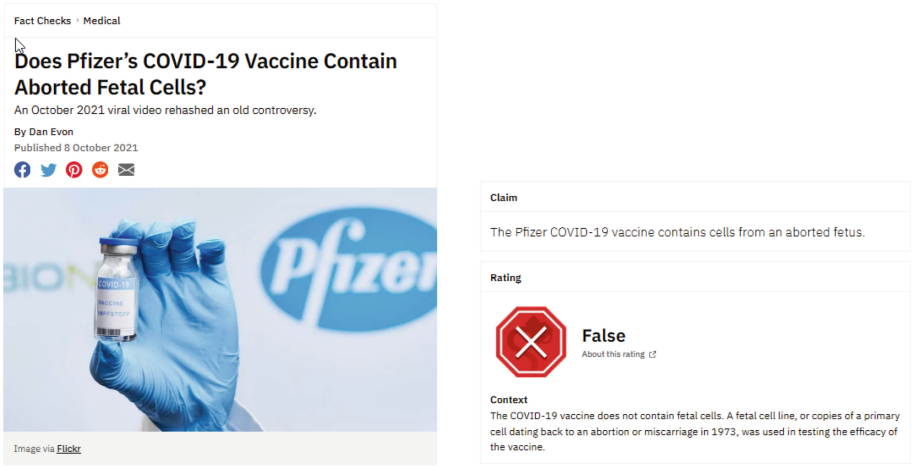
(b) Claim assessment

Fig. 1. AFP fact-checking report

4.2 SPO Extraction

After scraping the content, we applied NLP techniques to extract the SPO triples from the reports content. As introduced in Sect. 3, we used Python and *spaCy* library to analysis the content at syntactic and semantic level.

The first challenge we faced came from the library itself. With *spaCy* version 2.x we were able to use *neuralcoref* [4], a very powerful co-reference library that uses artificial neural networks to correctly replace the pronouns with the corresponding subjects. However, the predicate's tense extraction is only available in *spaCy* version 3.x and *neuralcoref* still doesn't support it. As an alternative, we started to use the *coreferee* [3] coreference library. This library is available for the latest version of *spaCy* and after evaluation it achieves comparable results to *neuralcoref* but with a more convoluted approach.



(a) Title and content of the report (b) Claim assessment

Fig. 2. Snopes fact-checking report

Additional challenges for the SPO triples extraction task are related to the fact-checking reports authors' writing styles. When the phrasing become more complex, the NLP libraries tend to struggle in identifying the correct tokens and complete sentences.

4.3 Knowledge Graph Creation

Starting from the SPO triples, a knowledge graph has been build. After analysing and testing the features of the 2-nodes and 3-nodes KG models, introduced in Sect. 3, the latter model was chosen. After careful evaluation and testing, the 3-nodes model has several advantages over the 2-nodes KG model. Given the characteristics of the Neo4J's Cypher query language, the 3-nodes model has higher computational speed, it allows to store more information in the node's attributes, and makes available the relation information when executing queries.

These characteristics also make the model suitable for the Relation Verification task. In Fig. 3 we show the partial view of the *FactsKG* representing the *AFP* fact-checking report shown in Fig. 1. The red colour nodes are Predicates nodes, in blue are shown the Subject/Object nodes. The links are also labelled with the report id stored in our dataset, as shown in Fig. 4.

5 Discussion

Processing the fact-checking reports with our methodology allows the creation of a knowledge graph. In *FactsKG* we were able to incorporate valuable metadata for the fact-checker that includes the assessed veracity value from multiple sources, the temporal information, original report and claim link, and so on.

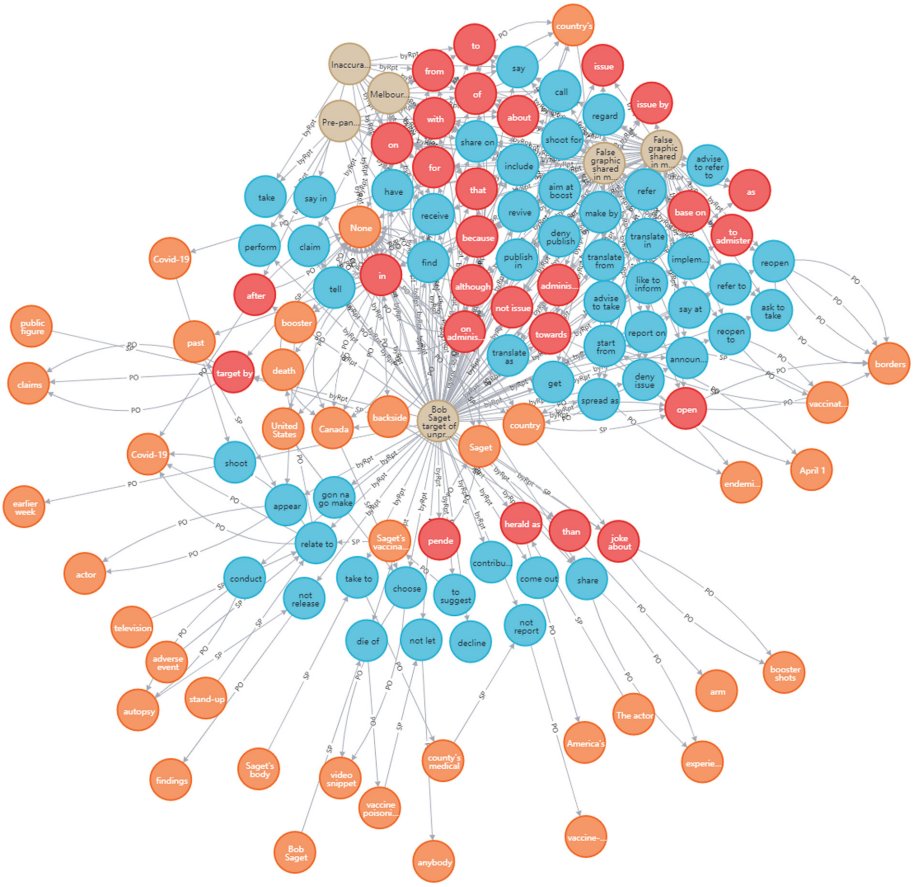
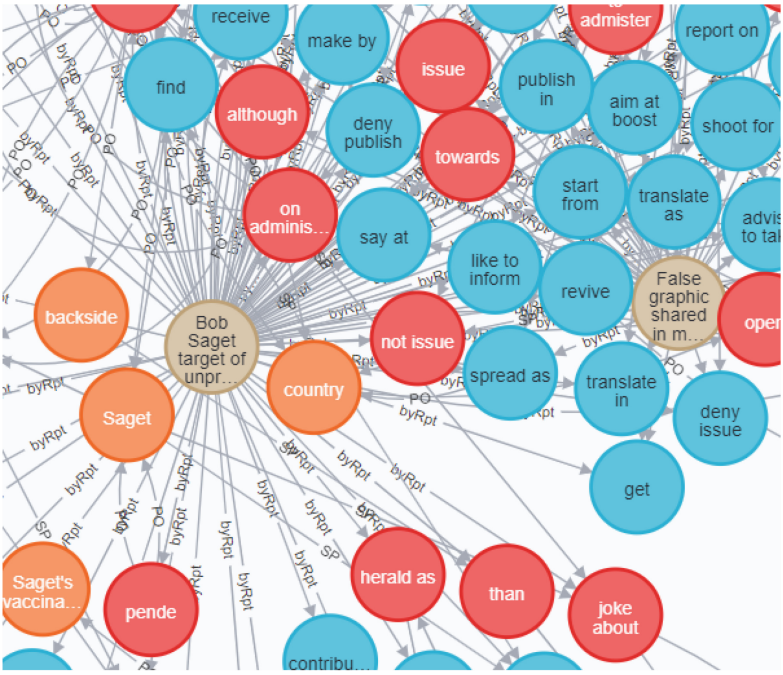
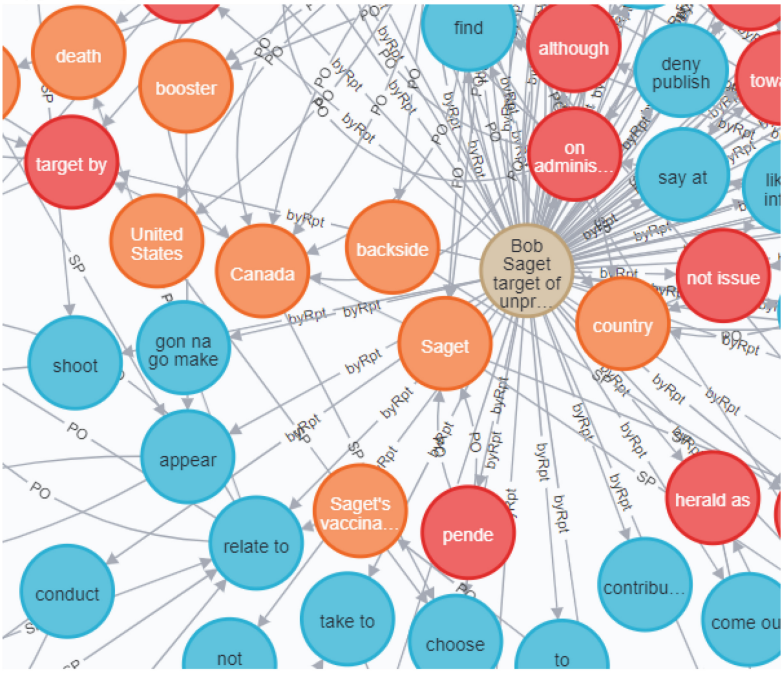


Fig. 3. Knowledge graph representing fact-checking report from AFP. (Color figure online)



(a) Two reports are connected by relationships on the KG through various nodes.



(b) Multiple connections from various predicates connect to the same subject/object.

Fig. 4. Zoom on a section of the KG (Color figure online)

This information has been made accessible for the fact-checkers through a web interface. The process of claim-checking starts with the input of a claim. The claim text will be processed with the same SPO extraction procedure introduced in Sect. 4.2. The SPOs extracted from the claim will be matched with the SPOs in *FactsKG*. Each SPO can be extracted by using a combination of complete or partial matching and with a direct or indirect search. The matching depends on whether we may consider the full SPO triple (complete matching) for the KG query or a part of it such as considering only subject or object as starting nodes (partial strategy). Moreover, the search strategy can be either direct or indirect, depending on whether we may consider only direct edges on the graphs or allow a longer path search. For the scope of this work we limited the number of steps for the partial indirect matching to 3, but the parameter can be freely adjusted.

Claim:
Bob Saget died of vaccine poisoning tonight.

Extracted SPOs:

Subject	Predicate	Object	Tense
Bob Saget	die of	vaccine poisoning tonight	Past

Direct complete matchup:

Subject	Predicate	Object	Textual Rating	Report id	Claim Date	Action
Bob Saget	die of	vaccine poisoning tonight	False	1	2022/01/9	Explore this triple in KG

Direct Partial matchup: Explore in KG

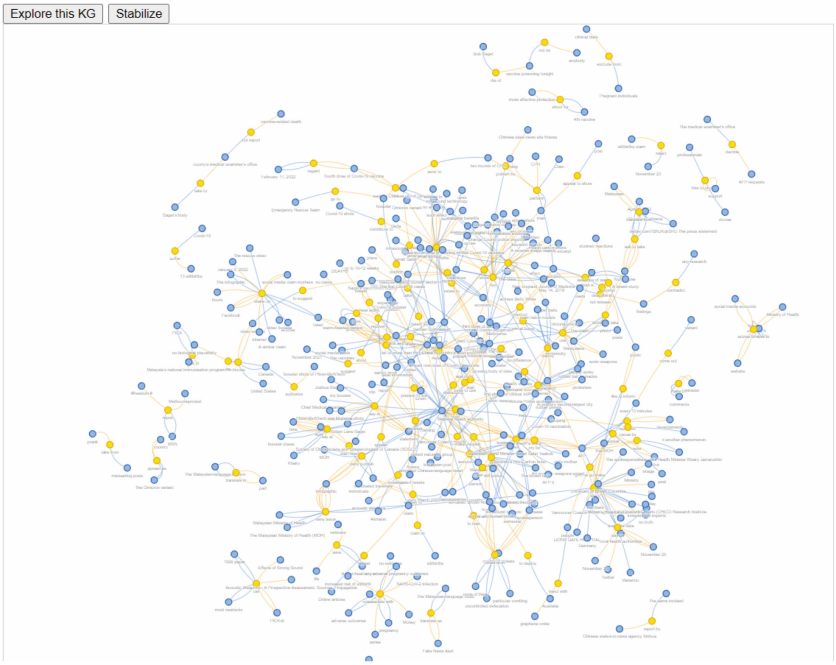
Subject	Predicate	Object	Textual Rating	Report id	Claim Date	Action
any adverse pregnancy outcomes	associate with	vaccine	False	3	2021/11/14	Explore this triple in KG
Saget	receive	Florida medical examiner's office	False	1	2022/01/9	Explore this triple in KG
Saget	receive	third dose of vaccine	False	1	2022/01/9	Explore this triple in KG
Bob Saget	die of	vaccine poisoning tonight	False	1	2022/01/9	Explore this triple in KG
vaccine poisoning tonight	not let	anybody	False	1	2022/01/9	Explore this triple in KG
poison	claim	tweet	False	1	2022/01/9	Explore this triple in KG
Saget	tell	person	False	1	2022/01/9	Explore this triple in KG
Saget	have	Omicron variant	False	1	2022/01/9	Explore this triple in KG
Saget	have	Delta	False	1	2022/01/9	Explore this triple in KG
Saget	have	Covid-19 booster	False	1	2022/01/9	Explore this triple in KG
Orange County police deputies	find	Saget	False	1	2022/01/9	Explore this triple in KG
Orange County police deputies	find	Saget	False	1	2022/01/9	Explore this triple in KG
Saget	say in	serious tone	False	1	2022/01/9	Explore this triple in KG
Saget	share	experience	False	1	2022/01/9	Explore this triple in KG
Saget	have	Omicron variant	False	1	2022/03/12	Explore this triple in KG
Saget	have	Delta	False	1	2022/03/12	Explore this triple in KG
Saget	have	Covid-19 booster	False	1	2022/03/12	Explore this triple in KG
Orange County police deputies	find	Saget	False	1	2022/03/12	Explore this triple in KG
Saget	receive	Florida medical examiner's office	False	1	2022/03/12	Explore this triple in KG
Saget	receive	third dose of vaccine	False	1	2022/03/12	Explore this triple in KG

Indirect complete matchup:

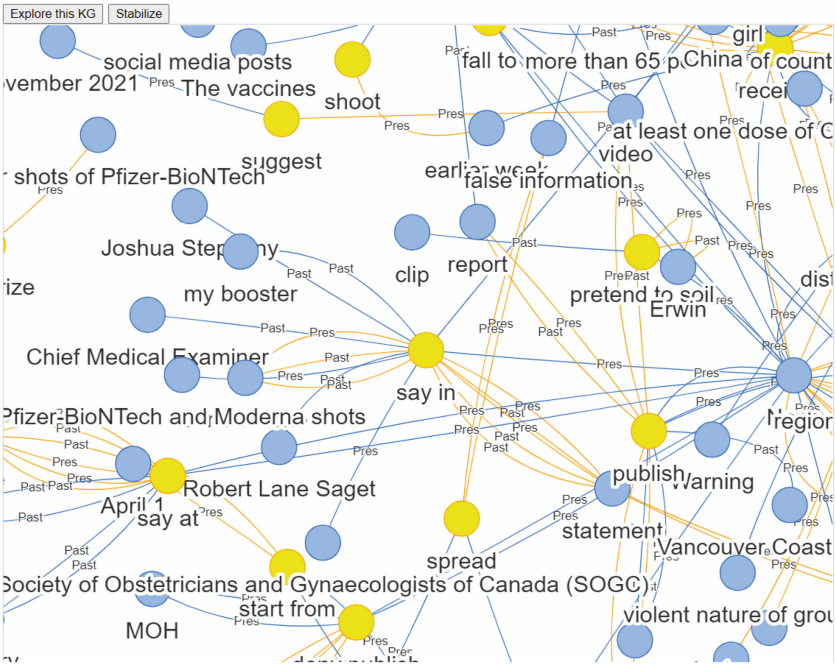
Subject	Predicate	Object	Textual Rating	Report id	Claim Date	Action
any adverse pregnancy outcomes	associate with	vaccine	False	3	2021/11/14	Explore this triple in KG

Fig. 5. Screenshot of the tabulated results of the claim matching.

The results of the matching are visualised using tabulated lists and graphical visualisation. In Fig. 5 is shown the tabulated results view, that reports the results for each matching and search combination. The textual list includes also



(a) Visualisation of the partial direct matching.



(b) Zoom on a part of the partial direct matching KG.

Fig. 6. Visualisation of the partial direct matching results. (Color figure online)

the option to visualise the results in graphical way. For each section there is also the additional button to show all the results within a single visualisation, as shown in Fig. 6. The graphical visualisation include the nodes and the relations involved in the matching. The blue colour nodes are subjects or objects, the yellow colour nodes are predicates, the blue links connect a subject to a predicates, and the orange links connect the predicates to the objects as stored in the KG. By using this interface it makes easier for fact-checkers to check the path and the reasoning that lead to the veracity assessment on multiple sources combined together in the *FactsKG*.

6 Conclusions and Future Works

In this work, we introduced a novel methodology for extracting a Knowledge Graph from fact-checking reports. The novel knowledge-based approach takes in account fact-checking reports from multiple sources to help the veracity level assessment by fact-checkers. The knowledge base storing the fact-checking reports, is first transformed in Subject-Predicate-Object triples and then integrated into the Knowledge Graph (KG). The *FactsKG* is an extended KG as integrates additional metadata about the facts. The metadata includes the veracity label associated to the facts, temporal information about the facts, and additional information useful for the task. The visual interface developed for matching new claims on *FactsKG*, when integrated in the fact-checking process, may facilitate the work of the fact-checkers.

In future works, we will introduce complex network analysis for the automatic Relation Verification task. The assessment of new claims will be performed on the KG using automated complex network analysis techniques.

Acknowledgements. This work is partly supported by the "AI-Info Communication Study (AIS) Scheme 2021/22 (Ref. AIS 21–22/05)" and "Teaching Development Grant" - School of Communication and Film, Hong Kong Baptist University, Hong Kong, China.

References

1. Hogan, A., et. al.: Knowledge Graphs. Synthesis Lectures on Data, Semantics, and Knowledge, vol. 12, no. 2, pp. 1–257 (2021). https://doi.org/10.2200/S01125ED1V01Y202109DSK022
2. Honnibal, M., Montani, I., Van Landeghem, S., Boyd, A.: SpaCy: Industrial-Strength Natural Language Processing in Python (2020). https://doi.org/10.5281/zenodo.1212303
3. Hudson, R.P., msg systems ag: coreferee: coreference resolution for multiple languages. https://github.com/msg-systems/coreferee
4. Huggingface: neuralcoref: fast coreference resolution in spaCy with neural networks. https://github.com/huggingface/neuralcoref
5. Mihaylova, T., et. al.: Fact checking in community forums. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)

6. Pérez-Rosas, V., Mihalcea, R.: Experiments in open domain deception detection. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 1120–1125 (2015)
7. Shaar, S., Martino, G.D.S., Babulkov, N., Nakov, P.: That is a known lie: detecting previously fact-checked claims. arXiv preprint arXiv:2005.06058 (2020)
8. Shi, B., Weninger, T.: Discriminative predicate path mining for fact checking in knowledge graphs. Knowl. Based Syst. **104**, 123–133 (2016)
9. Shiralkar, P., Flammini, A., Menczer, F., Ciampaglia, G.L.: Finding streams in knowledge graphs to support fact checking. In: 2017 IEEE International Conference on Data Mining (ICDM), pp. 859–864. IEEE (2017)
10. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: a data mining perspective. ACM SIGKDD Explor. Newsl. **19**(1), 22–36 (2017)
11. Tchechmedjiev, A., et al.: ClaimsKG: a knowledge graph of fact-checked claims. In: Ghidini, C., et al. (eds.) ISWC 2019. LNCS, vol. 11779, pp. 309–324. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30796-7_20
12. Zhou, X., Zafarani, R.: A survey of fake news: fundamental theories, detection methods, and opportunities. ACM Comput. Surv. (CSUR) **53**(5), 1–40 (2020)