OSINT Knowledge Graph for Fact-Checking: Google Map Hacks Debunking

Giuseppe Fenza, Vincenzo Loia, Paola Montserrat Mainardi and Claudio Stanzione

April 20, 2023

OSINT Knowledge Graph for Fact-Checking: Google Map Hacks Debunking

Giuseppe Fenza[1], Vincenzo Loia[1], Paola Montserrat Mainardi[1], and Claudio Stanzione[2]
[1] University of Salerno, Italy
[2] Center for Higher Defence Studies, Italy
gfenza@unisa.it, loia@unisa.it, p.mainardi2@studenti.unisa.it, stanzione.dottorando@casd.difesa.it

Abstract

The continuing spread of disinformation on the Web has caused a great demand for automatic fact-checking tools to help institutions and practitioners check the truthiness of web content. Existing literature presents numerous solutions for specific tasks, such as textual fact-checking or image/video truthiness, that, together with the availability of Open Source Intelligence (OSINT) tools and principles, paves the way for new comprehensive solutions. This work introduces a Knowledge Graph-based approach for fact-checking and news debunking. The idea is to map fact-checking workflow activities leveraging OSINT to specific scenarios emerging from Web and social media monitoring. The reference Knowledge Graph is constructed by analyzing sources through Text Mining and Semantic Analysis techniques. Finally, a real case study is carried out to show the applicability of the approach for fact-checking purposes.

1. Introduction

The ever-increasing proliferation of non-checked information (fake news, rumors, etc.) on the Internet may generate viral content that can negatively affect public opinion and democratic equilibria in the physical society. Fake content can be shared easily on social media (e.g., Reddit, Facebook, and Twitter) and news sites and blogs as well that may be further spread through link sharing and posting on other social communities or websites. When it eventually turns out to be false, fake content already caused unnecessary panic and anxiety among the public (Giansiracusa, 2021). For example, during covid pandemic early stages, some observed that, within cities, the spatial distribution of the density of COVID-19 infection correlates with the density of 5G towers, suggesting a possible correlation between them (Flaherty, Sturm, & Farries, 2022). Of course, there is no correlation between the two phenomena, but the fake news spread caused anxiety in the audience and paved the way to new unfounded conspiracy theories.

Fact-checking techniques may be employed to prevent disinformation proliferation and consequential anxiety among people. In detail, these activities are thought to clarify the primarily false information presented and thus force the recipient to think more deeply about the published facts (Kvetanová, Predmerská, & Švecová, 2020). To prove content falsity, there is the need to collect further data related to the news and compare them to spot some inconsistencies. News debunking is intended as a multi-level task since it requires not only checking the truthiness of news/post content but also dates, location, source and provenance. In this regard, some OSINT tools and services may help the fact-checkers seeking information on the Web. For instance, Google Maps, with its Street View, gives a chance to acquire images of a cited location to check elements in images posted; Google Lens allows finding elements from posted images, such as features suggesting places or pictures posted by other users or sites with meaningful information that can support the analysis.

Notwithstanding this digital support, proving the truth of content may be a very stressful task to fulfill due to the effort spent by humans in checking multiple features related to the news and choosing which aspect to focus on first or which data to search for to spot distinctive elements for the analysis. Therefore, questions arise: What are the elements (e.g., provenance, source, content, etc.) to be checked first? What kind of data should fact-checkers consider and retrieve to prove fact truthiness? Could the activities be organized in a unique method to perform reliable fact-checking?

This study tries to answer these questions by describing and applying a cognitive approach depicting the various stages to consider, data to acquire, and different OSINT technologies to perform fact-checking based on previous knowledge about similar facts or events. The approach leverages a Knowledge Graph (Fenza, Gallo, Loia, Marino, & Orciuoli, 2020), constructed through Text Mining and Semantic Analysis of considered pieces of content, and that guides the selection of suitable fact-checking activities. The functioning of the proposed model is proved by carrying out a complete case study about real news regarding an experiment performed by artist Simon Weckert consisting in tricking Google Maps service with simulated traffic congestion. The paper contribution consists of a cognitive approach for fact-checking content based on a Knowledge Graph. In particular:

- The KG exploits Text Mining and Semantic Analysis;
- The KG suggests suitable actions (among ones in the fact-checking workflow) based on the experience;
- The KG, at the same time, harvests experts' feedback to feed the KG;
- A new ontological model (Debunking Model) is defined for representing domain-specific concepts;
- The Debunking Model helps in identifying inconsistencies among pieces of evidence;
- The Semantic Analysis consists of inheriting and combining existing ontologies to one depicted in this work.
- The approach applicability has been proven through a real example of fact-checking.

The rest of the paper is organized as follows: Section 2 provides related work, Section 3 introduces the proposed cognitive approach for fact-checking, and Section 4 demonstrates its potential in a real case study. Conclusions close the paper.

2. Related work

Many approaches explored fact-checking techniques to discover different kinds of disinformation forms. Some works, in particular, focussed on debunking rumors on social networks to analyze behaviors [Tian and Ding (2019)] or exploring specific features, such as the unmasking of fake reviews (Bangerter, et al., 2021) or the employment of denials to increase sharing and spread potentials (Pal, Chua, & Goh, 2019). Other works focussed on the interplay between rumor spreading and debunking, ending up with a model determining when a debunking application is needed or most effective (Jiang, Gao, & Zhuang, 2021).

Other types of research focus on designing automated or semi-automated tools for fact-checking, such as (Martín, Huertas-Tato, Huertas-García, Villar-Rodríguez, & Camacho, 2022) that introduce a model based on semantic similarity and natural language inference to perform multilingual fact-checking and hoax spread monitoring. Other solutions explore knowledge-based systems to deal with fact-checking, such as in (Wang, Mao, Wei, & Zeng, 2022), which presents a semantic structure that exploits a knowledge graph for fact-learning and verification. In other cases, the idea consists of comparing different sources with different credibility levels (De Maio, Fenza, Gallo, Loia, & Volpe, 2020), or analyzing possible biases in news perceptions and exploring how partisan leanings influence the news selection algorithm for fact-checking (Babaei, et al., 2021).

Some other works employed knowledge to achieve more robust fact-checking, such as (Zhu, Zhang, Gu, & Deng, 2021) proposes an approach combining Wikidata5M knowledge graph and Wikipedia documents to incorporate external knowledge into the claim. In addition, the paper by (Seddari, et al., 2022) introduces a hybrid fake news detection system combining linguistic and knowledge-based features to spot fake news on social networks.

Contrary to the existing literature, this work proposes a cognitive approach whose Knowledge Graph (KG) suggests suitable actions (among ones in the fact-checking workflow) based on the experience and, at the same time, harvests experts' feedback to feed the KG. Moreover, the proposed approach is applied to an existing case study to check its suitability in debunking and fact-checking.

3. Proposed Methodology

This paper proposes a methodology consisting of constructing and maintaining a Knowledge Graph (KG) to suggest suitable fact-checking activities among the ones described by the Fact-checking workflow. The KG is fed with data taken from the Web and social media (see Figure 1), as well as experts' feedback.

The following subsections detail the KG and the workflow used in the proposed methodology.

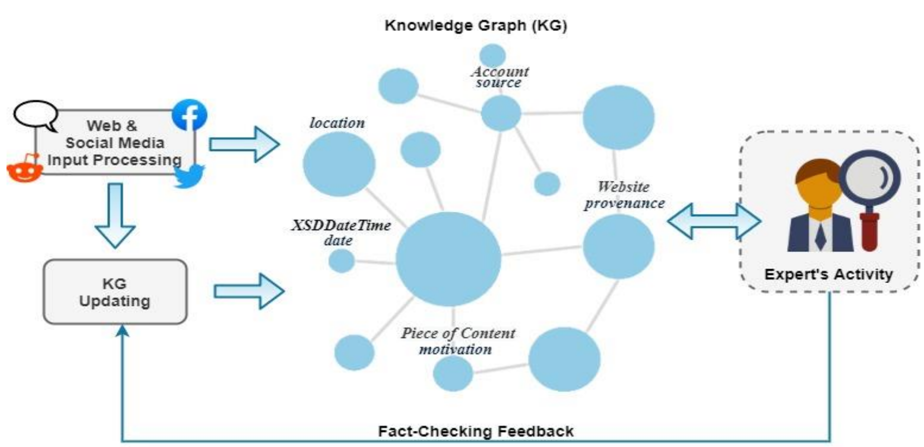

Figure 1 - Proposed methodology

3.1. Knowledge Graph for Fact-Checking

As already stated, the proposed methodology is based on a Knowledge Graph (KG) supplied by suggestions, doubts, or skepticism noticed on the Web or social media. The idea consists of identifying sources of interest (e.g., Facebook pages, Twitter accounts, etc.) and constantly monitoring their evolution, mainly about specific topics. Relevant recommendations can be extracted by scraping such pages and applying Natural Language Processing (NLP) techniques combined with a semantical analysis. Recommendations conveniently processed and conceptualized populate the KG and give directions for subsequent fact-checking activities. Moreover, the experts' choices and decisions further contribute to KG's growth and update. For example, suppose the expert realizes a correlation between a news aspect and a workflow phase: such intuition feeds the graph.

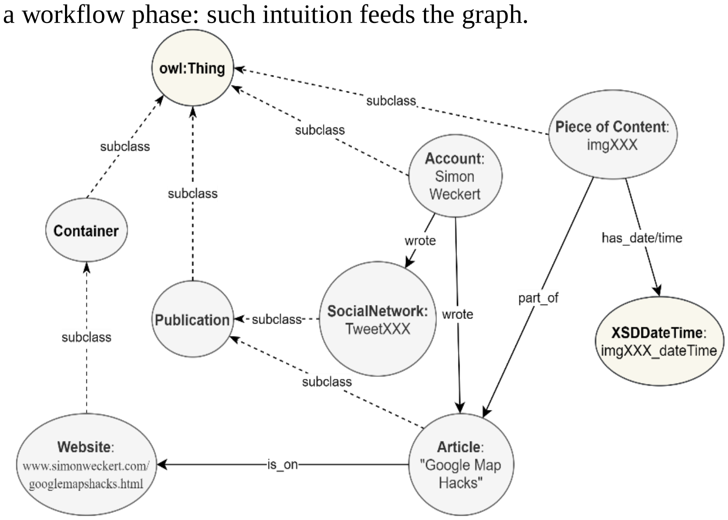

Figure 2 - Debunking Model

The graph construction leverages NLP techniques and the *Debunking Model* for the semantical analysis. Debunking Model consists of an ontological model relating the main aspects involved in the workflow to interface and analyze the data collected about the article or post (i.e., container, content, user profiles, etc.). The ontology is built over state-of-the-art ontologies (i.e., OWL[1], OWL time[2], etc.) to model knowledge about the people writing and sharing web contents, temporal and spatial information related to the contents and people's activities, respectively. New classes and properties added to existing ontologies to represent the main aspects involved in fact-checking and debunking activities are depicted in Figure 2 and summarized as follows:

- Class **Container** and its subclass **Website** represent the web space where contents are shared.
- Class **Publication** represents the posted content that may be an article on a news site (class **Article**) or a post on a social network (**SocialNetwork** class). Elements of a post or an article, such as text, photos, links and videos, are represented as instances of **Piece of Content**.

[1] http://owl.man.ac.uk/2006/07/sssw/people.owl
[2] https://www.w3.org/TR/owl-time/

- Class **Account** represents people that wrote and shared specific content (e.g., an article, a post) on a Website.

Moreover, Figure 2 also demonstrates a simplified example of ontology instantiation based on the case study described in Section 4.

3.2. Fact-checking workflow

Fact-checking experts, mainly from important news agencies, shared numerous hints aimed at evaluating content trustworthiness. In particular, the suggested workflow by Urbani (Urbani, 2020) described in this section exploits Open Source INTelligence (OSINT) fundamentals. OSINT tools, by collecting and analyzing publicly available information, can assist in, for example, certifying the ownership of an image or identifying a location in it. The process mainly consists of verifying a piece of content's provenance, source, date, location, and motivation. In particular:

- The provenance ensures to refer to the original article or piece of content.
- Found the source refers to identifying who created the original piece of content.
- The date refers to when the content was created.
- The location identifies where the piece of content was captured.
- The motivation aims to consider what caused the capturing of the piece of content.

Each pillar contributes to a better understanding of the content and its reliability. The following subsections give details about each pillar.

Finding the ***provenance*** of content means examining its original form to more easily understand who posted it, when, where, and why. Techniques to discover the original content depend on the type of content. For example, the Reverse Image Search, consisting of searching for the content in large databases (e.g., Google Images), can be a solution for images. In the case of videos, a frame from a video can be subjected to a reverse image search through the Reverse Video Search. In some situations, finding the original content through other strategies is difficult; searching for it in more private and anonymous locations could be helpful. Some examples are Reddit[3], 4chan[4], Discord[5], and, where appropriate, Twitter and Facebook.

The ***source*** of content (i.e., owner) can be a valid indication of its reliability. However, since everyone can re-post a piece of content written or captured by someone else on the Internet, it is mandatory to identify its real "owner".

Once the first uploader is found, we should understand if the content is coherent with the authors' geographic position, other shared contents, and so on. In particular, it could be interesting to investigate authors' social accounts, make reverse image searches of account images, search for shared posts in Google to understand if there is embedded content, and so forth. In addition, checking if declared email addresses are associated with any user (for example, through Skype) could help determine the source's credibility and, for example, make sure it is not an automated account (i.e., a bot). In this regard, specific techniques can be used (e.g., learning models), or attention can be paid to the volume of daily posts and whether a period of silence is associated with nighttime rest.

Finding the ***date*** means determining when the original content was created. A starting point is referring to timestamps associated with a post or file metadata, for example, in the case of image files, the Exif (Exchangeable image file format). However, since these types of metadata are not always available, some further checks could include observing the video/image to understand the period of the year. In this sense, handy tools are:

---

[3] https://camas.github.io/reddit-search/

[4] https://4chansearch.com/

[5] https://disboard.org/search

- SunCalc[6]. It allows viewing the sun's angle on a specific day at a particular location, which can help identify the time associated with an event in a photo or video.
- Wolfram Alpha[7]. It is a computational knowledge engine that, among other things, enables you to look up the weather for a specific date. This way, a check can be done between the declared date and the weather for that date.

Similar problems with date reconstruction may arise with ***location*** identification since geo-tags are not always available and may not accurately reflect the location provided in the content. Finding specific details in the image or video and doing research using satellite imagery may be helpful for this aim. Examples include looking for squares, signs, flags, banners, etc., to attempt and associate a location with the picture or video. Also, spoken language and clothing could be helpful. However, special attention should be paid to the level of update of images and the latest events in the area by considering the most recent local events (e.g., war or extremely severe weather) that may significantly impact the terrain.

About ***motivation***, it is difficult to identify common hints: the process strictly depends on the considered piece of content. However, in general terms, it could be helpful to find involved people's affiliations or communities and, where possible, try to speak directly with them. Motivation could be better depicted by extracting the context of the news facts by analyzing comments to the post/news, tweets and other web resources referencing the article considered.

4. Case Study

The KG-based approach described in Section 3 has been demonstrated in a case study based on an experiment by the German artist Simon Weckert, who walked in empty streets of Berlin bringing a little wagon containing 99 smartphones, each running a GPS Maps service. The artist aimed to simulate a fake traffic congestion event by exploiting the high number of devices, their closeness and the slow movement of the wagon. The experiment was introduced by the artist in an article called "Google Maps Hacks" on his personal blog[8] and reported in the book "How Algorithms Create and Prevent Fake News" (Giansiracusa, 2021).

The objective of the case study is to question some technical aspects of the experiment that the author does not fully describe. Following the approach proposed in Section 3, we tried to extract useful information to understand the feasibility of the experiment. In particular, the case study starts with identifying relevant suggestions about the experiment itself on the Web to realize what pillars of the workflow are bringing up. Moreover, each pillar investigation contributes to populating the Debunking Model, as depicted in Figure 2. Finally, the KG turns on incoherence among findings.

The following subsections detail each pillar analyzed.

4.1. Provenance

News about the experiment appeared on many news sites and blogs; however, we focused on original pictures and information posted on Weckert's site.

[6] https://www.suncalc.org/
[7] https://www.wolframalpha.com/
[8] https://www.simonweckert.com/googlemapshacks.html

4.2. Source

The source of information is the artist himself, who posted information about the experiment on his personal site. However, through cross searches, we also identify his Twitter account[9], from which we also find the date of publication of the experiment itself. The tweet posting the experiment had a lot of resonance. In particular, other Twitter users' comments questioned technical modalities of the experiment. For example, they expressed doubts about adopted devices, Internet connections and attainability (see Figure 3); Weckert did not clarify all factors. Following, this aspect is further examined.

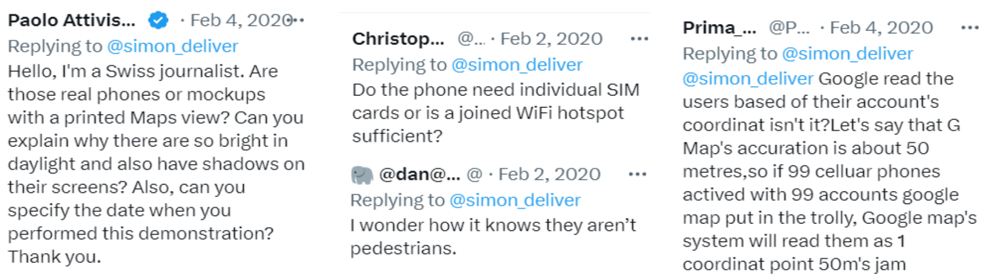

Figure 3 - Example of comments to the tweet announcing the experiment.

Since the pictures depicting adopted smartphones (see Figure 4) are insufficient to discover device details, we explored the Web to find other helpful information. In a video posted by Arte TV[10], we catch the image in Figure 5 from which smartphones are better visible. The first remark regards the difference between smartphone screens in Figure 4 and Figure 5: in this latter, despite the lack of light concerning the experiment execution, screens are darker and less visible.

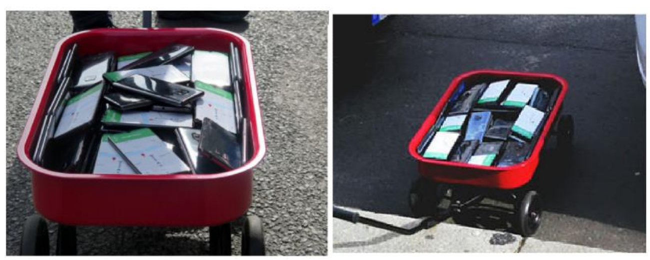

Figure 4 - Site pictures depicting adopted smartphones.

In order to acquire more information about devices, we searched for similar images and obtained the smartphone model through Google Lens. It is a Huawei Mate 20 Pro, a particularly expensive model released in 2018. Therefore, the hypothesis that Weckert bought all these expensive devices is hard to believe. Instead, many suppose he rented or purchased second-hand devices; anyway, he does not provide further details about them.

9 https://twitter.com/simon_deliver
10 https://www.arte.tv/de/videos/102583-000-A/simon-weckert-arte-tracks/

Figure 5 - Smartphones adopted during the experiment.

Another crucial element to prove experiment truthiness is geo-localization. Since the smartphones are all grouped in the wagon, they have very close GPS coordinates or even the same, which is not the case of a traffic congestion scenario but more similar to an accident with 99 vehicles involved. Therefore, even though fascinating and unreal, the 99-car accident thesis must be rejected since the cart moves and crashed cars do not move.

4.3. Date

From Weckert's website, it is unclear when the experiment was done. Unique available information is the date of the news posting on Twitter (i.e., 02/01/2020). Thus, through the FotoForensics[11] tool, we extracted image metadata that, for pictures in Figure 6 (i.e., not subsequently modified), reports October 06, 2019, as the creation date. For images in Figure 7, only the modification date was available, i.e., October 14, 2019.

By a search on Wolphram Alpha, at the given date, the weather is congruent with the conditions in the photos. Moreover, the sun's direction, evaluated through sunCalc, is congruent with the lights in the pictures.

Figure 6 – Site pictures in temporal sequence.

[11] https://fotoforensics.com/

Figure 7 - Site pictures without temporal information.

4.4. Location

Regarding location identification, we know a priori that Weckert conducts the experiment in Berlin, Germany. However, in the experiment description, the exact covered distance and the main streets visited are reported as unclear. So, to further analyze the feasibility of the experiment, we realized a focused analysis to reconstruct its path, as described below.

By leveraging picture metadata, we know the chronological sequence of pictures, as depicted in Figure 6. Moreover, from the images in Figure 7, we can extract traversed streets (that, through searches on Google Street View, also match with associated images):

- An der Schillingbrücke
- Michaelbrücke
- Ziegelstraße and Ebertbrücke.

Discovering locations reported in Figure 6's pictures needs an in-depth analysis made through the synergy between Google Lens and Google Street View. In particular, we can recognize, in chronological order:

- Mittelstraße
- Geschwister-Scholl-Straße

The third location in Figure 6 is impossible to identify since there is no recognizable place.

The longest path among cited places is between three and four kilometers (depending on the type of selected path). In this sense, let us notice that on maps shown on Weckert's site, with traffic mode on, the path is shown as full green lines along the roadside, as reported in the leftmost image in Figure 8. Since this configuration is present in maps showing paths by feet, we assume that Google Maps has detected the smartphones as people walking on the road; otherwise, the result would have been as the one shown in in the right image in Figure 8 (i.e., the smartphones would have been taken as 99 vehicles). Although the possibility of managing fake accounts, according to Google Maps, which collects data during the trip (i.e., the type of device employed and the mean speed), the service would not have mistakenly taken the 99 smartphones as 99 vehicles. Moreover, the hypothesis that the service would have considered the 99 devices as people seems not credible either.

Figure 8 - Comparison between a walking path (on the right) and a driving path (on the left). In this latter, a blue line covers the road space.

4.5. Motivation

Simon Weckert, through his site, declares his fascination for the digital world and its reflection on social aspects. He aims to evaluate the worth of technology from the perspective of future generations. The artist's philosophy aligns with the nature of the experiment.

Conclusions

The paper presents a cognitive approach for fact-checking web content based on a Knowledge Graph leveraging OSINT tools and principles and a domain ontology. In particular, the method considers state-of-the-art tips and associated instruments for acquiring and fact-checking information. Acquired data is conveniently annotated and analyzed at each workflow stage to determine inconsistencies between different aspects (i.e., website, account profiles, articles and pieces of content), allowing smart fact-checking activities. The proposed approach was applied to a real case study concerning an experiment performed by artist Simon Weckert consisting in tricking Google Maps service with simulated traffic congestion. The case study aims to demonstrate the practical potential of the proposed methodology to support practitioners and fact-checkers in ascertaining the truthiness of web content through a multi-aspect fact analysis.

In the future, it would be interesting to automatize the whole workflow supporting the technical tasks (e.g., KG implementation and querying), as well as content identification, extraction and annotation.

Acknowledgment

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU.

References

Babaei, M., Kulshrestha, J., Chakraborty, A., Redmiles, E. M., Cha, M., & Gummadi, K. P. (2021). Analyzing biases in perception of truth in news stories and their implications for fact checking. *IEEE Transactions on Computational Social Systems, 9*(3), 839--850.

Bangerter, M. L., Fenza, G., Gallo, M., Genovese, A., Nota, F. D., Stanzione, C., & Zanfardino, G. (2021, June). Unmask inflated product reviews through Machine Learning. In 2021 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA) (pp. 1-6). IEEE.

De Maio, C., Fenza, G., Gallo, M., Loia, V., & Volpe, A. (2020). Cross-relating heterogeneous Text Streams for Credibility Assessment. Bari, Italy: IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS).

Fenza, G., Gallo, M., Loia, V., Marino, D., & Orciuoli, F. (2020). A cognitive approach based on the actionable knowledge graph for supporting maintenance operations. *IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS)*, (pp. 1--7). Bari, Italy.

Flaherty, E., Sturm, T., & Farries, E. (2022). The conspiracy of Covid-19 and 5G: Spatial analysis fallacies in the age of data democratization. *Social Science & Medicine, 293*, 114546.

Giansiracusa, N. (2021). *How Algorithms Create and Prevent Fake News: Exploring the Impacts of Social Media, Deepfakes, GPT-3, and More* (Apress ed.). New York: Springer.

Jiang, M., Gao, Q., & Zhuang, J. (2021). Reciprocal spreading and debunking processes of online misinformation: A new rumor spreading--debunking model with a case study. *Physica A: Statistical Mechanics and its Applications*, *565*.

Kvetanová, Z., Predmerská, A. K., & Švecová, M. (2020). Debunking as a Method of Uncovering Disinformation and Fake News. In *Fake News Is Bad News: Hoaxes, Half-truths and the Nature of Today's Journalism.* IntechOpen.

Martín, A., Huertas-Tato, J., Huertas-García, Á., Villar-Rodríguez, G., & Camacho, D. (2022). FacTeR-Check: Semi-automated fact-checking through semantic similarity and natural language inference. *Knowledge-Based Systems, 251*, 109265.

Pal, A., Chua, A. Y., & Goh, D. H.-L. (2019). Debunking rumors on social media: The use of denials. *Computers in Human Behavior, 96*, 110--122.

Seddari, N., Derhab, A., Belaoued, M., Halboob, W., Al-Muhtadi, J., & Bouras, A. (2022). A hybrid linguistic and knowledge-based analysis approach for fake news detection on social media. *IEEE Access*, *10*, 62097--62109.

Urbani, S. (2020, 09 22). *Verifying Online Information.* Retrieved 02 15, 2023, from https://firstdraftnews.org/long-form-article/verifying-online-information/

Wang, S., Mao, W., Wei, P., & Zeng, D. D. (2022). Knowledge structure driven prototype learning and verification for fact checking. *Knowledge-Based Systems, 238*, 107910.

Zhu, B., Zhang, X., Gu, M., & Deng, Y. (2021). Knowledge enhanced fact checking and verification. *IEEE/ACM Transactions on Audio, Speech, and Language Processing, 29*, 3132--3143.