The Battlefront of Combating Misinformation and Coping with Media Bias

Yi R. Fung[1], **Kung-Hsiang Huang**[1], **Preslav Nakov**[2], **Heng Ji**[1]
[1]University of Illinois Urbana-Champaign
[2]Mohamed bin Zayed University of Artificial Intelligence
{yifung2, khhuang3, hengji}@illinois.edu
preslav.nakov@mbzuai.ac.ae

Abstract

Misinformation is a pressing issue in modern society. It arouses a mixture of anger, distrust, confusion, and anxiety that cause damage on our daily life judgments and public policy decisions. While recent studies have explored various fake news detection and media bias detection techniques in attempts to tackle the problem, there remain many ongoing challenges yet to be addressed, as can be witnessed from the plethora of untrue and harmful content present during the COVID-19 pandemic and the international crises of late. In this tutorial, we provide researchers and practitioners with a systematic overview of the frontier in fighting misinformation. Specifically, we dive into the important research questions of how to (*i*) develop a robust fake news detection system, which not only fact-check information pieces provable by background knowledge but also reason about the consistency and the reliability of subtle details for emerging events; (*ii*) uncover the bias and agenda of news sources to better characterize misinformation; as well as (*iii*) correct false information and mitigate news bias, while allowing diverse opinions to be expressed. Moreover, we discuss the remaining challenges, future research directions, and exciting opportunities to help make this world a better place, with safer and more harmonic information sharing.

1 Introduction

The growth of online platforms has greatly facilitated the way people communicate with each other and stay informed about trending events. However, it has also spawned unprecedented levels of inaccurate or misleading information, as traditional journalism gate-keeping fails to keep up with the pace of media dissemination. These undesirable phenomena have caused societies to be torn over irrational beliefs, money lost from impulsive stock market moves, and deaths occurred that could have been avoided during the COVID-19 pandemic, due to the infodemic that came forth with it, etc. (Allcott and Gentzkow, 2017; Rapoza; Solomon et al., 2020). Even people who do not believe the misinformation may still be plagued by the pollution of unhealthy content surrounding them, an unpleasant situation known as *information disorder* (Wardle et al., 2018). Thus, it is of pertinent interest for our community to better understand, and to develop effective mechanisms for remedying misinformation and biased reporting.

The emerging nature of news events, which also span diverse domains (e.g., economy, military, health, sports, etc.) and reporting style (e.g., long text vs. short text, realistic image vs. artistic image, etc.), makes misinformation detection and characterization challenging. Combating fake news and biased reports involve an interdisciplinary research area of reasoning on the semantics, style, cross-media contextualization, background knowledge, and propagation patterns, among others (Saquete et al., 2020; Pennycook and Rand, 2021; Collins et al., 2021). Moreover, the recent trends towards a more comprehensive understanding of the source stance, reporting intent, target audience, and propaganda technique behind a problematic piece of news (Zhang and Ghorbani, 2020) require greater socio-cultural norm and common sense awareness.

In this half-day tutorial, we aim to present a systematic overview of technological advancement in tackling interconnected tasks related to misinformation, media bias, malicious intent monitoring, and corrective actions. First, we will review prevailing paradigms and data resources for misinformation detection and characterization. Moreover, we will discuss the latest approaches to automatically explain why a news piece is inaccurate or misleading, and perform rectification of biased reporting. The participants will learn about trends and emerging challenges, representative deep neural network models, ready-to-use training resources, as well as how state-of-the-art language (and multimedia)

Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: Tutorial Abstracts, pages 28–34
November 20-21, 2022.

techniques can help build applications for the social good.

2 Outline of Tutorial Content

2.1 Background and Motivation [20min]

We begin motivating the tutorial topic with a selection of real-world examples of fake news and their harmful impacts to society, followed by a pedagogical exercise of how humans tend to approach the problem of fake news detection, characterization, and correction. We will point out conceptual distinctions amongst various types of fake news, including serious fabrication in news journalism about misattributed or nonexisting events, oversensationalized clickbaits, hoaxes which are false with the intention to be picked up by traditional news websites and satire which mimic genuine news but contain irony and absurdity (Rubin et al., 2015). For example, in general, news articles more likely involve serious fabrications, while social media posts involve more humour such as satire and hoaxes. We will also describe the cognitive, social and affective factors that lead people to form or endorse misinformed views (e.g., intuitive thinking, illusory truths, source cues, emotions, etc.), and the psychological barriers to knowledge revision after misinformation has been corrected, including correction not integrated, selective retrieval, and continued influence theories (Ecker et al., 2022).

2.2 Fake News Detection [60min]

Bearing these properties in mind, we introduce:

- ***stylistic*** approaches that focus on lexical features, readability, and syntactic clues (Pérez-Rosas et al., 2018; Rashkin et al., 2017; Choshen et al., 2019)

- ***fact-checking*** approaches that compare check-worthy content with background knowledge, such as external knowledge bases (FreeBase, WikiBase, etc) and previously fact-checked claims (Baly et al., 2018; Shaar et al., 2020; Hu et al., 2021; Liu et al., 2021; Guo et al., 2022)

- ***semantic-consistency*** approaches that extract features related to single-document discourse-level coherence (Karimi and Tang, 2019) and cross-document event-centric coherence (Wu et al., 2022) in text. Extending to cross-media domain, the common strategy is to check text–image consistency (Tan et al., 2020; Huang et al., 2022; Aneja et al., 2021) and text–video consistency (Wang et al., 2022).

- ***propagation patterns*** that capture confounding factors from the dynamics of how a news topic spreads and the social network interactions (Lu and Li, 2020; Shu et al., 2020; Cheng et al., 2021).

We will discuss the merits and the limitations of these different lines of fake news detection approaches. For example, fact-checking approaches may not fare well for early rumours or breaking news not yet groundable to an established background knowledge (Zhou et al., 2019; Guo et al., 2022), in which case, the credibility of the news source can offer complementary assistance (Cheng et al., 2021). Stylistic approaches may be simple but yet effective for detecting low-quality human-written fake news, but not so good for machine-generated misinformation, which is stylistically consistent regardless of the underlying motives (Schuster et al., 2020). We then cover recent approaches (Lee et al., 2021b; Fung et al., 2021) that leverage a combination of these elements for greater representation power and robustness. Importantly, we also cover works that explore the diachronic bias of fake news detection and portability across data in different time and language settings (Murayama et al., 2021; Gereme et al., 2021).

Special Note on Neural Fake News Generation & Detection:
Advancements in natural language generation spawn the rise of news generation models which represent a double-edged sword (Zellers et al., 2019). On one hand, malicious actors may irresponsibility take advantage of the technology to influence opinions and gain revenue. But, on the other hand, it can also be used as a source of machine-synthesized training data for detector models to overcome data scarcity since real-world fake news tends to be eventually removed by platforms, as well as a tool for threat modeling to develop proactive defenses against potential threats. We review how popular detectors perform on fake news created from large-scale language and vision generator model (Zellers et al., 2019; Güera and Delp, 2018; Agarwal et al., 2019). We also review progress in

generating fake news that better aligns with the key topic and facts (Mosallanezhad et al., 2021; Shu et al., 2021; Fung et al., 2021), and work towards applying topic/fact-constrained fake news generation to construct silver-standard data annotations for finer-grained fake news detection (Fung et al., 2021).

2.3 Fake News Characterization [30min]

To better understand and fight fake news, we next address some fundamental questions of characterizing fake news based on underlying source bias, reporting agenda, propaganda techniques, and target audience (Buchanan, 2020). First, we introduce modeling approaches for detecting political and socio-cultural biases in news articles (Kulkarni et al., 2018; Fan et al., 2019; Baly et al., 2020; Forbes et al., 2020). Next, we introduce the recent EMU benchmark that require models to answer open-ended questions capturing the intent and the implications of a media edit (Da et al., 2021). We cover methodologies for identifying the specific propaganda techniques used, *e.g.*, *smears*, *glittering generalities*, *association transfer*, etc. (Dimitrov et al., 2021). We also discuss the latest explorations in predicting the intended target of harmful media content, *e.g.*, the person, the organization, the community, or the society level (Pramanick et al., 2021).

2.4 Corrective Actions for Misinformation and Biased News Reporting [30min]

After misinformation has been detected and categorized based on its various characteristics, there is naturally follow-up interest in corrective explanations on why a piece of information is fake or misleading, and how to report less biased and more comprehensive news in general. Hence, we cover frameworks for explaining why a given piece of news is actually fake news through the leverage of reader comments, as well as appropriate strategies for placing the corrective explanations based on user studies (Shu et al., 2019; Brashier et al., 2021). We also cover research on mitigating media bias, such as through neutral article generation (Lee et al., 2021a).

Industry Initiatives: We further point out recent actions by tech companies with media-hosting platforms for fighting fake news. With urges from the government, they experiment with removing economic incentives for traffickers of misinformation, promoting media literacy, suspending improper posts and accounts, and adding colored labels, with corrections constructed from a community-based point system similar to Wikipedia, directly beneath misinformation posted by public figures[1].

2.5 Concluding Remarks & Future Directions [30min]

Finally, we summarize the major remaining challenges in this space, including the detection of subtle inconsistencies, enforcing schema or logical constraints in the detection, identifying semantically consistent but misattributed cross-media pairings, and greater precision in fine-grained explanations for the detected misinformation.

3 Specification of the Tutorial

The proposed tutorial is a cutting-edge tutorial that introduces new frontiers in research on battling misinformation and news bias. The presented topic has not been covered by previous ACL/NAACL/AACL tutorials in the past four years. While there has been an EMNLP'20 tutorial on "Fact-Checking, Fake News, Propaganda, and Media Bias: Truth Seeking in the Post-Truth Era" (Nakov and Da San Martino, 2020) and a COLING'20 tutorial on "Detection and Resolution of Rumors and Misinformation with NLP" (Derczynski and Zubiaga, 2020), fake news is a continuously evolving and extremely important societal problem. In our tutorial, we place particular emphasis on the latest lines of development, including an emphasis on multimedia contextualization, sociocultural awareness in characterization, and corrective actions. We estimate at least 75% of the work we reference has not been covered in the two previous aforementioned tutorials. We further estimate that at least 75% of the research covered in this tutorial is by researchers other than the instructors.

Audience and Prerequisites Based on the level of interest in this topic, we expect around 100 participants. While no specific background knowledge is assumed of the audience, it would be best for the attendees to know basic deep learning, pre-trained word embeddings (e.g., Word2Vec) and language models (e.g., BERT).

Reading List We recommend the literature cited in this paper, particularly: the rising threats of neural fake news (Zellers et al., 2019; Chawla, 2019),

[1] https://www.nbcnews.com/tech/tech-news/twitter-testing-new-ways-fight-misinformation-including-community-based-points-n1139931

knowledge-driven misinformation detection (Hu et al., 2021; Fung et al., 2021; Guo et al., 2022), intent characterization (Buchanan, 2020; Da et al., 2021), and study of fake news impact from a psychological point of view (Ecker et al., 2022).

Desired Venue The most desired venue for this tutorial would be AACL-IJCNLP'2022. The majority of our tutorial speakers have educational experience in Asia. At the same time, we also represent a global diversity in our research work.

Open Access We agree to allow the publication of the tutorial materials and presentation in the ACL Anthology. All the materials will be openly available at the UIUC Blender Lab website.

4 Tutorial Instructors

Below, we give the biographies of the speakers.

Yi R. Fung is a second-year Ph.D. student at the Computer Science Department of UIUC, with research interests in knowledge reasoning, misinformation detection, and computation for the social good. Her recent works include the INFOSURGEON fake news detection framework, and multiview news summarization. Yi is a recipient of the NAACL'21 Best Demo Paper, the UIUC Lauslen and Andrew fellowship, and the National Association of Asian American Professionals Future Leaders award. She has also been previously selected for invited talk (1 hour presentation) at the Harvard Medical School Bioinformatics Seminar. Additional information is available at `https://yrf1.github.io`.

Kung-Hsiang Huang is a first-year Ph.D. student at the Computer Science Department of UIUC. His research focuses on fact-checking and fake news detection. Prior to joining UIUC, he obtained his B.Eng. in Computer Science from the Hong Kong University of Science and Technology, and his M.S. in Computer Science is from USC. He is also a co-founder of an AI startup, Rosetta.ai. Additional information is available at `https://khuangaf.github.io`.

Preslav Nakov is a Principal Scientist at the Qatar Computing Research Institute (QCRI), HBKU, who received his PhD degree from the University of California at Berkeley (supported by a Fulbright grant). Dr. Nakov is President of ACL SIGLEX, Secretary of ACL SIGSLAV, a member of the EACL advisory board, as well as a member of the editorial board of Computational Linguistics, TACL, CS&L, IEEE TAC, NLE, AI Communications, and Frontiers in AI. His research on fake news was featured by over 100 news outlets, including Forbes, Boston Globe, Aljazeera, MIT Technology Review, Science Daily, Popular Science, The Register, WIRED, and Engadget, among others. He has driven relevant tutorials such as:

- WSDM'22: Fact-Checking, Fake News, Propaganda, Media Bias, and the COVID-19 Infodemic.
- CIKM'21: Fake News, Disinformation, Propaganda, and Media Bias.
- EMNLP'20: Fact-Checking, Fake News, Propaganda, and Media Bias: Truth Seeking in the Post-Truth Era.

Additional information is available at `https://en.wikipedia.org/wiki/Preslav_Nakov`.

Heng Ji is a Professor at the Computer Science Department of the University of Illinois Urbana-Champaign, and an Amazon Scholar. Her research interests focus on NLP, especially on Multimedia Multilingual Information Extraction, Knowledge Base Population and Knowledge-driven Generation. She was selected as "Young Scientist" and a member of the Global Future Council on the Future of Computing by the World Economic Forum. The awards she received include "AI's 10 to Watch" Award, NSF CAREER award, Google Research Award, IBM Watson Faculty Award, Bosch Research Award, Amazon AWS Award, ACL2020 Best Demo Paper Award, and NAACL2021 Best Demo Paper Award. She has given a large number of keynotes and 20 tutorials on Information Extraction, Natural Language Understanding, and Knowledge Base Construction in many conferences including but not limited to ACL, EMNLP, NAACL, NeurIPS, AAAI, SIGIR, WWW, IJCAI, COLING and KDD. A selected handful of her recent tutorials include:

- AAAI'22: Deep Learning on Graphs for Natural Language Processing. Language Processing.
- EMNLP'21: Knowledge-Enriched Natural Language Generation.
- ACL'21: Event-Centric Natural Language Processing.

Additional information is available at `https://blender.cs.illinois.edu/hengji.html`.

Ethical Considerations

Technological innovations often face the dual usage dilemma, in which the same advance may offer potential benefits and harms. For the news probing methodologies introduced in this tutorial, the distinction between beneficial use and harmful use depends mainly on the data and intention. Proper use of the technology requires that input corpora be legally and ethically obtained, with the target goal to fight misinformation and mal-intents. Besides, training and assessment data may be biased in ways that limit the system performance on less well-represented populations and in new domains – causing performance discrepancy for different ethnic, gender, and other sub-populations. Thus, questions concerning generalizability and fairness need to be carefully considered when applying news analysis techniques to specific settings. A general approach to ensure proper application of dual-use technology should incorporate ethical considerations as the first-order principles in every step of the system design, maintain transparency and interpretability of the data, algorithms, and models, and explore counter-measures to protect vulnerable groups.

References

Shruti Agarwal, Hany Farid, Yuming Gu, Mingming He, Koki Nagano, and Hao Li. 2019. Protecting world leaders against deep fakes. In *CVPR workshops*, volume 1.

Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–36.

Shivangi Aneja, Chris Bregler, and Matthias Nießner. 2021. Cosmos: Catching out-of-context misinformation with self-supervised learning. *arXiv preprint arXiv:2101.06278*.

Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. We can detect your bias: Predicting the political ideology of news articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, EMNLP '20, pages 4982–4991.

Ramy Baly, Georgi Karadzhov, Dimitar Alexandrov, James Glass, and Preslav Nakov. 2018. Predicting factuality of reporting and bias of news media sources. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3528–3539, Brussels, Belgium. Association for Computational Linguistics.

Nadia M Brashier, Gordon Pennycook, Adam J Berinsky, and David G Rand. 2021. Timing matters when correcting fake news. *Proceedings of the National Academy of Sciences*, 118(5).

Tom Buchanan. 2020. Why do people spread false information online? The effects of message and viewer characteristics on self-reported likelihood of sharing social media disinformation. *PLOS One*, 15(10):e0239666.

Ronit Chawla. 2019. Deepfakes: How a pervert shook the world. *International Journal of Advance Research and Development*, 4(6):4–8.

Lu Cheng, Ruocheng Guo, Kai Shu, and Huan Liu. 2021. Causal understanding of fake news dissemination on social media. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 148–157.

Leshem Choshen, Dan Eldad, Daniel Hershcovich, Elior Sulem, and Omri Abend. 2019. The language of legal and illegal activity on the Darknet. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, ACL '19, pages 4271–4279, Florence, Italy.

Botambu Collins, Dinh Tuyen Hoang, Ngoc Thanh Nguyen, and Dosam Hwang. 2021. Trends in combating fake news on social media–a survey. *Journal of Information and Telecommunication*, 5(2):247–266.

Jeff Da, Maxwell Forbes, Rowan Zellers, Anthony Zheng, Jena D Hwang, Antoine Bosselut, and Yejin Choi. 2021. Edited media understanding frames: Reasoning about the intent and implications of visual misinformation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, ACL-IJNCLP '21, pages 2026–2039.

Leon Derczynski and Arkaitz Zubiaga. 2020. Detection and resolution of rumors and misinformation with nlp. In *Proceedings of the 28th International Conference on Computational Linguistics: Tutorial Abstracts*, pages 22–26.

Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. Detecting propaganda techniques in memes. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, ACL-IJCNLP '21, pages 6603–6617.

Ullrich KH Ecker, Stephan Lewandowsky, John Cook, Philipp Schmid, Lisa K Fazio, Nadia Brashier, Panayiota Kendeou, Emily K Vraga, and Michelle A Amazeen. 2022. The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1):13–29.

Lisa Fan, Marshall White, Eva Sharma, Ruisi Su, Prafulla Kumar Choubey, Ruihong Huang, and Lu Wang. 2019. In plain sight: Media bias through the lens of factual reporting. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, EMNLP-IJCNLP '19, pages 6343–6349, Hong Kong, China.

Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. 2020. Social chemistry 101: Learning to reason about social and moral norms. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, EMNLP '20, pages 653–670.

Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, and Avirup Sil. 2021. InfoSurgeon: Cross-media fine-grained information consistency checking for fake news detection. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, ACL-IJCNLP '21, pages 1683–1698.

Fantahun Gereme, William Zhu, Tewodros Ayall, and Dagmawi Alemu. 2021. Combating fake news in "low-resource" languages: Amharic fake news detection accompanied by resource crafting. *Information*, 12(1):20.

David Güera and Edward J Delp. 2018. Deepfake video detection using recurrent neural networks. In *2018 15th IEEE international conference on advanced video and signal based surveillance (AVSS)*, pages 1–6. IEEE.

Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.

Linmei Hu, Tianchi Yang, Luhao Zhang, Wanjun Zhong, Duyu Tang, Chuan Shi, Nan Duan, and Ming Zhou. 2021. Compare to the knowledge: Graph neural fake news detection with external knowledge. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, ACL-IJCNLP '21, pages 754–763.

Mingzhen Huang, Shan Jia, and Siwei Lyu. 2022. Text-Image De-Contextualization Detection Using Vision-Language Models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Singapore.

Hamid Karimi and Jiliang Tang. 2019. Learning hierarchical discourse-level structure for fake news detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT '19, pages 3432–3442, Minneapolis, MN, USA.

Vivek Kulkarni, Junting Ye, Steve Skiena, and William Yang Wang. 2018. Multi-view models for political ideology detection of news articles. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3518–3527, Brussels, Belgium. Association for Computational Linguistics.

Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021a. Mitigating media bias through neutral article generation. *arXiv preprint arXiv:2104.00336*.

Nayeon Lee, Belinda Z. Li, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih, and Madian Khabsa. 2021b. On unifying misinformation detection. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT '21, pages 5479–5485.

Lihui Liu, Boxin Du, Yi Ren Fung, Heng Ji, Jiejun Xu, and Hanghang Tong. 2021. Kompare: A knowledge graph comparative reasoning system. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery Data Mining*, KDD '21, page 3308–3318, New York, NY, USA. Association for Computing Machinery.

Yi-Ju Lu and Cheng-Te Li. 2020. GCAN: Graph-aware co-attention networks for explainable fake news detection on social media. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online. Association for Computational Linguistics.

Ahmadreza Mosallanezhad, Kai Shu, and Huan Liu. 2021. Generating topic-preserving synthetic news. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 490–499.

Taichi Murayama, Shoko Wakamiya, and Eiji Aramaki. 2021. Mitigation of diachronic bias in fake news detection dataset. In *Proceedings of the Seventh Workshop on Noisy User-generated Text*, W-NUT '21, pages 182–188.

Preslav Nakov and Giovanni Da San Martino. 2020. Fact-checking, fake news, propaganda, and media bias: Truth seeking in the post-truth era. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 7–19.

Gordon Pennycook and David G Rand. 2021. The psychology of fake news. *Trends in cognitive sciences*, 25(5):388–402.

Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. Automatic detection of fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*, COLING '18, pages 3391–3401, Santa Fe, NM, USA.

Shraman Pramanick, Shivam Sharma, Dimitar Dimitrov, Md. Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2021. MOMENTA: A multimodal framework for detecting harmful memes and their targets. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4439–4455, Punta Cana, Dominican Republic.

Kenneth Rapoza. Can 'fake news' impact the stock market? *Forbes*.

Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2931–2937.

Victoria L Rubin, Yimin Chen, and Nadia K Conroy. 2015. Deception detection for news: three types of fakes. *Proceedings of the Association for Information Science and Technology*, 52(1):1–4.

Estela Saquete, David Tomás, Paloma Moreda, Patricio Martínez-Barco, and Manuel Palomar. 2020. Fighting post-truth using natural language processing: A review and open challenges. *Expert systems with applications*, 141:112943.

Tal Schuster, Roei Schuster, Darsh J Shah, and Regina Barzilay. 2020. The limitations of stylometry for detecting machine-generated fake news. *Computational Linguistics*, 46(2):499–510.

Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. 2020. That is a known lie: Detecting previously fact-checked claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, ACL '20, pages 3607–3618.

Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. dEFEND: Explainable fake news detection. In *Proceedings of the 25th International Conference on Knowledge Discovery & Data Mining*, KDD '19, pages 395–405, Anchorage, AK, USA.

Kai Shu, Yichuan Li, Kaize Ding, and Huan Liu. 2021. Fact-enhanced synthetic news generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13825–13833.

Kai Shu, Deepak Mahudeswaran, Suhang Wang, and Huan Liu. 2020. Hierarchical propagation networks for fake news detection: Investigation and exploitation. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 626–637.

Daniel H Solomon, Richard Bucala, Mariana J Kaplan, and Peter A Nigrovic. 2020. The "infodemic" of covid-19.

Reuben Tan, Bryan Plummer, and Kate Saenko. 2020. Detecting cross-modal inconsistency to defend against neural fake news. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, EMNLP '20, pages 2081–2106.

Kehan Wang, David Chan, Seth Z Zhao, John Canny, and Avideh Zakhor. 2022. Misinformation detection in social media video posts. *arXiv preprint arXiv:2202.07706*.

Claire Wardle, Hossein Derakhshan, et al. 2018. Thinking about 'information disorder': formats of misinformation, disinformation, and mal-information. *Ireton, Cherilyn; Posetti, Julie. Journalism,'fake news'& disinformation. Paris: Unesco*, pages 43–54.

Xueqing Wu, Kung-Hsiang Huang, Yi Fung, and Heng Ji. 2022. Cross-document misinformation detection based on event graph reasoning. In *NAACL 2022*, Seattle, WA, USA.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. *Advances in neural information processing systems*, 32.

Xichen Zhang and Ali A Ghorbani. 2020. An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 57(2):102025.

Kaimin Zhou, Chang Shu, Binyang Li, and Jey Han Lau. 2019. Early rumour detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT '19, pages 1614–1623, Minneapolis, MN, USA.