Introduction to the Special Issue on Combating Digital Misinformation and Disinformation

NAEEMUL HASSAN, University of Mississippi
CHENGKAI LI, University of Texas at Arlington
JUN YANG, Duke University
CONG YU, Google Research

1 INTRODUCTION

We are delighted to present this special issue of the *Journal of Data and Information Quality* (ACM *JDIQ*) on *Combating Digital Misinformation and Disinformation*. This issue presents an overview of innovative research primarily at the intersection of information credibility, machine learning, and data science, from theory to practice, with a focus on combating misinformation and disinformation.

Spread of misinformation and disinformation is one of the most serious challenges facing the news industry, and a threat to democracy worldwide. The problem has reached an unprecedented level via social media, where contents can be created and disseminated to a large audience with little to zero cost and revenues are driven by clicks. Researchers from multiple disciplines have proposed various strategies, built automated and semiautomated systems [1, 3], and recommended policy changes across the media ecosystem [2, 4]. Recently, researchers also explored how artificial intelligence techniques, particularly machine learning and natural language processing, can be leveraged to combat falsehoods online.

In this special issue of *JDIQ*, we provide a representative collection of insightful articles at the intersection of data quality and credibility, from theory to practice, with a focus on improvements in veracity and value. The articles went through a rigorous procedure of review involving at least three expert reviewers for each article. After two rounds of review, we selected five articles that made contributions to both research and practice.

Zannettou et al., in "The Web of False Information: Rumors, Fake News, Hoaxes, Clickbait, and Various Other Shenanigans," provide a typology of the false information content on the Web and surveys the latest research directions. It identifies several lines of works in the false information ecosystem. In particular, it surveys the research works from false information propagation, perception, and identification perspectives. Then, the authors specifically attend the false information spread in the political domain and investigate the velocity and consequence of the spread in communities. Finally, the authors delineate several future research directions that can help understand and mitigate this misinformation problem.

Xue et al., in "Content-Aware Trust Propagation Toward Online Review Spam Detection," provide algorithms to detect spam reviews that are prevalent online. The authors propose using the inconsistency in opinions expressed in one's review and the consensus opinions shared by a group of anonymous users to determine the trustworthiness of a review and its reviewer. Specifically, they model the opinion deviations from the majority in the form of a deviation-based penalty and integrate this penalty into a three-layer trust propagation framework to compute the trust scores for users, reviews, and review targets, respectively, in an iterative fashion. The authors experimented with three different datasets and demonstrated the efficacy of the proposed frameworks. In addition, they described how the framework can be extended to transfer users' trust scores from one domain to other domains.

The next three articles in this special issue focus on methodologies and algorithms for automated fact-checking.

Atanasova et al., in "Automatic Fact Checking Using Context and Discourse Information," present methodologies for detecting "check-worthy" claims in political debates and for fact checking answers to questions posed in online community forums. Their methods are based on supervised learning, using neural networks and kernel-based support vector machines. The key innovation lies in the use of additional contextual information—for political debates, they consider the both the full turn of a debate participant and the preceding and following turns by others; for checking answers in forums, they consider the whole question-answer thread and related posts in the entire forum. The authors have developed additional annotations for the 2016 U.S. presidential and vice-presidential debates, as well as community question answering data from the Qatar Living Forum; results on these datasets confirm the benefit of considering contextual features in these automated fact-checking tasks.

Lin et al., in "Discovering Patterns for Fact Checking in Knowledge Graphs," study how to characterize useful graph patterns in modeling facts based on large knowledge graphs and how to efficiently discover such graph patterns to support fact checking. The idea is based on the concept of graph fact checking rules (GFCs). A GFC defines graph patterns that distinguish between true and false statements. The article provides statistical measures to capture whether patterns are discriminative and diverse. It makes a good theoretical contribution in that it presents an algorithm to discover GFCs in large graphs with optimality guarantees. The article further discusses how to develop two classifiers for fact checking that use top-ranked GFCs as predictive rules and instance-level features of the pattern matches induced by GFCs. It discusses the results of comprehensive experiments using several real-world datasets, including general knowledge bases such as YAGO, DBpedia, and Wikidata and academic, social, and event networks.

Borges et al., in "Combining Similarity Features and Deep Representation Learning for Stance Detection in the Context of Checking Fake News," investigate a core research task in automated fact checking: stance detection and natural language inference. Stance detection is crucial for automated fact checking given a claim statement—the ability to correctly detect whether a particular article is generally supportive or against a given claim statement is one of the key steps in performing automated fact checking, especially without the context of a knowledge graph. The authors build on prior works using the Fake News Challenge[1] dataset, and they incorporate pretrained embeddings and encoder models from the deep learning literature to construct a model that is performing on par/slightly better than the previous state of art. This is a nontrivial achievement considering that deep learning models typically require a substantial amount of training data, which is not really available in the Fake News Challenge dataset.

---

[1]http://www.fakenewschallenge.org.

REFERENCES

[1] Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by ClaimBuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 1803–1812.

[2] David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, et al. 2018. The science of fake news. *Science* 359, 6380 (2018), 1094–1096.

[3] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. 2016. Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th International Conference Companion on World Wide Web.* 745–750.

[4] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter* 19, 1 (2017), 22–36.