

Combating Fake News: A Data Management and Mining Perspective

Laks V.S. Lakshmanan[‡], Michael Simpson[‡], Saravanan Thirumuruganathan[†]

[‡]University of British Columbia; [†]QCRI, HBKU

[‡]{laks,mesimp}@cs.ubc.ca, [†]sthirumuruganathan@hbku.edu.qa

ABSTRACT

Fake news is a major threat to global democracy resulting in diminished trust in government, journalism and civil society. The public popularity of social media and social networks has caused a contagion of fake news where conspiracy theories, disinformation and extreme views flourish. Detection and mitigation of fake news is one of the fundamental problems of our times and has attracted widespread attention. While fact checking websites such as snopes, politifact and major companies such as Google, Facebook, and Twitter have taken preliminary steps towards addressing fake news, much more remains to be done. As an interdisciplinary topic, various facets of fake news have been studied by communities as diverse as machine learning, databases, journalism, political science and many more.

The objective of this tutorial is two-fold. First, we wish to familiarize the database community with the efforts by other communities on combating fake news. We provide a panoramic view of the state-of-the-art of research on various aspects including detection, propagation, mitigation, and intervention of fake news. Next, we provide a concise and intuitive summary of prior research by the database community and discuss how it could be used to counteract fake news. The tutorial covers research from areas such as data integration, truth discovery and fusion, probabilistic databases, knowledge graphs and crowdsourcing from the lens of fake news. Effective tools for addressing fake news could only be built by leveraging the synergistic relationship between database and other research communities. We hope that our tutorial provides an impetus towards such synthesis of ideas and the creation of new ones.

PVLDB Reference Format:

Laks V.S. Lakshmanan, Michael Simpson, Saravanan Thirumuruganathan. Combating Fake News: A Data Management and Mining Perspective. *PVLDB*, 12(12): 1990-1993, 2019. DOI: <https://doi.org/10.14778/3352063.3352117>

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 12, No. 12

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3352063.3352117>

1. MOTIVATION

Fake news is one of the major problems facing governments, society, academia, and industry. The dissemination of false and misleading information has a corrosive effect on the trust of public toward various institutions [1]. While fake news has always existed in history, the advent of new technology has exacerbated its reach and potential for damage. Social media and online social networks have allowed weaponization of fake news at an unprecedented scale [15, 30]. Not surprisingly, the academic community has mobilized to counteract this phenomenon. Many research communities including databases, machine learning, data mining, journalism and political science have focused on different aspects of this problem.

Our tutorial provides a unifying framework for categorizing prior research focusing on four facets of fake news: *detection*, *propagation*, *mitigation* and *intervention*. The work on detection seeks to identify which items are fake through diverse tools such as machine learning, content analysis, propagation analysis, computational fact checking and so on. The study of propagation seeks to understand and model how fake news spreads in various media such as social media. Once the fake news has spread, there are many techniques for addressing it. The mitigation based approaches seek to minimize the severity by “inoculating” users from falling for fake news, reduce the impact of filter bubbles, showing diverse viewpoints and so on. Finally, the intervention based approaches take an active role in reducing the spread of fake news. They could involve mild interventions such as amplifying real news alongside fake news to severe ones such as removing content or accounts.

Goal of Tutorial. The goal of this tutorial is to provide an intuitive summary of research from other communities and discuss how prior work in the database community could be used to enrich them. Instead of delineating database and non-database research, we offer a synthesis based approach that interleaves non-database and database research that emphasizes the strong connection between them. We review various technical challenges, recent solutions and highlight a number of intriguing open problems at the intersection of various communities. We hope that this will empower database researchers to make impactful contributions in the fight against fake news.

Scope of Tutorial. Given the huge amount of prior research on fake news and the limited amount of time, this tutorial provides a carefully selected subset of topics that we believe are most relevant to database researchers. For

example, while we touch upon machine learning based approaches for fake news detection, we focus extensively on how prior database research on data integration and fusion could also be used for detecting fake news items. Similarly, while we briefly introduce propagation models for fake news, we weave them with influence maximization works that have been carried out in the database community.

Target Audience and Prerequisites. The target audience consists of database researchers who wish to familiarize themselves with recent developments in fake news research and contribute to novel algorithms and systems for combating it. This tutorial is also appropriate for practitioners who want an intuitive overview of the state of the art. The tutorial assumes familiarity with databases, data mining, and fundamentals of machine learning.

Connection with Related tutorials. Due to the timeliness and relevance of fake news, there have been a number of tutorials in data mining and machine learning communities on this topic. Related tutorials include: (a) Computational Fact Checking: A Content Management Perspective by Sylvie Cazalens, Julien Leblay, Philippe Lamarre, Ioana Manolescu and Xavier Tannier in VLDB 2018 and WWW 2018; (b) Fake News: Fundamental Theories, Detection Strategies and Challenges by Xinyi Zhou, Reza Zafarani, Kai Shu and Huan Liu in WSDM 2019; (c) Computational Solutions against Fake News: AI vs. DB Approaches by Naeemul Hassan and Dongwon Lee in AAAI 2018; (d) Mining Misinformation in Social Media: Understanding Its Rampant Spread, Harm, and Intervention by Liang Wu, Giovanni Luca Ciampaglia and Huan Liu in ICDM 2017; and (e) Fact Checking: Theory and Practice by Xin Luna Dong, Christos Faloutsos, Xian Li, Subhabrata Mukherjee and Prashant Shiralkar in KDD 2018.

The proposed tutorial differs from existing ones in several aspects: (1) we provide a novel unifying framework for research on fake news that seamlessly combines prior work from database and other communities. (2) we highlight how various innovations made by the database community could be used to address fake news. (3) none of the previous tutorials above offer the whole spectrum of detection, propagation, mitigation and intervention unlike us.

2. TUTORIAL ORGANIZATION

In this section, we describe a high level overview of the major topics covered in the tutorial.

2.1 Primer on Fake News (15%)

In this introductory part of the tutorial, we describe various attempts at defining fake news [28]. Fake news is often used as an umbrella term for various specific concepts such as rumors, conspiracy theories, hoaxes, clickbaits, media hype, satire, misreporting, and sensationalism. While fake news has always existed, a number of recent phenomena such as the advent of social media, increasing polarization and filter bubbles, perverse incentives and virality have exacerbated their reach. We briefly describe the ecosystem of fake news and why it has become a serious threat [18, 2]. Finally, we describe the impact of fake news that requires a sustained focus from researchers from various communities including our own.

2.2 Propagation of Fake News (15%)

In the second part of the tutorial, we provide a concise summary of previous attempts to model the propagation of fake news such as [30]. Fake news often becomes more viral than real news and more often. A thorough understanding of the propagation of fake news could be used for: (1) detection: fake news often propagates in a distinct way compared to real news that could be used for early detection [32]; (2) mitigation: by understanding the key contributors for propagation, we could mitigate it by attenuating the factors for virality [26]. We introduce how influence propagation models studied by researchers in the database and other communities could be used for understanding the spread of fake news.

2.3 Detection of Fake News (40%)

The objective of this part of the tutorial is two-fold. First, we wish to familiarize the database community with prior attempts by the ML and AI communities for detecting fake news. Second, we describe how various research techniques pioneered by the database community could be used for fake news detection. We hope that the synthesis of these ideas would result in a more sophisticated next generation mechanism for combating fake news.

2.3.1 ML based Approaches

Prior research [27, 33, 26] could be categorized along two dimensions. The first dimension controls the specific ML technique used: supervised or unsupervised. The second dimension is based on the specific features used for the chosen ML technique. The unsupervised approaches often operate under the assumption that real and fake news have distinct features that could be learned automatically and used to distinguish them. If most of the news in a social network is real, then fake news could be considered as an anomaly allowing us to use prior work on outlier/anomaly detection. In contrast, supervised approaches begin with a labeled training dataset that is used to train a machine learning or deep learning model. This model is then used to categorize news as real and fake. Of course, the model is periodically re-trained to prevent adversarial users from gaming it. Both these models work by extracting features for categorization. Popular features include content of the news, credibility of the creator, propagation trace of the article etc.

2.3.2 Database Approaches

There have been some preliminary efforts by our community for fake news detection and computational fact checking [7, 12]. At its core, the problem boils down to the following: given a claim and a database of facts, how can one determine if the claim is real or fake news? Most of the approaches rely on closed world assumption whereby all true facts are assumed to be those stored in an appropriate repository such as relational database or knowledge graph. By considering different ways by which claims and the database are instantiated, it is possible to use a wide variety of prior work from databases. Some promising avenues include:

Data Integration is an important body of work [10] that allows one to combine a number of heterogeneous data sources (that could be structured or unstructured) into a single unified data repository. We first describe a scenario where all the data sources used for fact checking are reliable and discuss the challenges inherent in integrating them. We

discuss common data integration architectures and how the integrated database could be used for fact checking.

Truth Discovery and Fusion. Even usually reliable sources could contain facts that are incorrect or inconsistent. Typically, most of the sources in the web are user generated that contain conflicting or erroneous information. This results in a scenario where there are many sources that could possibly disagree with each other. Given this setting, it is important to identify the veracity of the claims so that eventually fake claims can be identified. There has been extensive research [8, 9, 5, 16], on algorithms for truth discovery using techniques for evaluating data source reputation/authoritativeness, claim veracity etc. However, extending them to fake news [13] is often challenging and requires additional research. Convincing a user of the falsehood of a claim calls for additional capabilities such as explanations to concisely describe how the decision was made [21].

Knowledge Graphs such as Freebase & DBpedia have been successfully used to store facts that are collected from the web. It is possible to construct a knowledge graph based on trusted entities and repositories such as fact checkers. The knowledge graph can then be used to answer simple queries – and in our context – whether a claim is supported [8, 6, 17, 23, 24, 25, 11]. Typically, checking a claim is related to finding a path between the subject and object of an RDF triple. We provide a brief overview of fake news detection approaches such as ClaimBuster and DeFacto that use knowledge graphs.

Crowdsourcing has been extensively studied by our community for solving various data management problems. Fact checking could be considered as a crowdsourcing problem that involves either experts (such as professional fact checkers) or non-experts (such as users of a social media). Techniques for estimating the skill of a worker and identifying the consensus from the responses of workers could be used and extended to automatically establish the veracity of a claim [27]. Given that the current manual fact checking approach is not scalable, crowdsourcing could offer a viable alternative. Example uses include identifying the quality of online news sources [20], fast detection of fake news in social media websites such as Twitter [29], and even in combating and reducing the spread of fake news [14].

2.4 Mitigation and Intervention of Fake News (20%)

In this part, we summarize prior techniques to mitigate the spread of fake news. Specifically, we focus on:

1. **Epidemiological approaches** that focus on identifying appropriate nodes for inoculation such that the spread of the disease is minimized [19]. This could be done before or after a specific fake news has started spreading.
2. **Influence Maximization models** that consider the fake news as an active viral campaign and seek to identify a competing campaign [3, 4] to minimize the reach or exposure of the fake news.

We consider a wide variety of intervention tactics ranging from removal of fake news items to termination of accounts that consistently spread fake news [22] and describe effective algorithms for each of those tactics and describe their

relative efficacy. We also consider some “softer” touch approaches recently used by major companies such as Google and Facebook that gently nudge the users away from fake news.

2.5 Future Opportunities (20%)

In our final part of the tutorial, we make a call to arms to the database community to investigate the problem of fake news in its various aspects, by identifying major research problems. We highlight a number of promising avenues for collaboration including building of effective systems and increasing the scalability of fact checking systems. We believe that our community has pioneered a number of tools that could be effective weapons against fake news. As just one example, there are intriguing connections between query processing on uncertain databases [31] and identifying veracity of news.

3. BIOGRAPHICAL SKETCHES

Laks V.S. Lakshmanan is a professor in the department of Computer Science at the University of British Columbia. He is a Research Fellow of the BC Advanced Systems Institute and was named ACM Distinguished Scientist in November 2016. His research interests span a wide spectrum of topics in Database Systems and related areas, including: relational and object-oriented databases, advanced data models for novel applications, OLAP and data warehousing, database mining, data integration, semi-structured data and XML, directory-enabled networks, querying the WWW, information and social networks and social media, recommender systems, and personalization.

Michael Simpson is a Postdoctoral Researcher in the Department of Computer Science at the University of British Columbia. He earned his PhD from the University of Victoria. His research interests include data mining, social network analysis, and the design of scalable algorithms for graph problems.

Saravanan (Sara) Thirumuruganathan is a scientist in the Data Analytics group of QCRI, HBKU. He earned his PhD from University of Texas at Arlington. He is broadly interested in data integration/cleaning and machine learning for data management. Saravanan’s work has been selected among best papers of VLDB 2018 and 2012.

4. REFERENCES

- [1] H. Allcott and M. Gentzkow. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–36, 2017.
- [2] V. Bakir and A. McStay. Fake news and the economy of emotions: Problems, causes, solutions. *Digital Journalism*, 6(2):154–175, 2018.
- [3] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *International workshop on web and internet economics*, pages 306–311. Springer, 2007.
- [4] A. Borodin, Y. Filmus, and J. Oren. Threshold models for competitive influence in social networks. In *International workshop on internet and network economics*, pages 539–550. Springer, 2010.

- [5] C. Castillo, M. Mendoza, and B. Poblete. Information credibility on twitter. In *WWW*, pages 675–684. ACM, 2011.
- [6] G. L. Ciampaglia, P. Shiralkar, L. M. Rocha, J. Bollen, F. Menczer, and A. Flammini. Computational fact checking from knowledge networks. *PLoS one*, 10(6):e0128193, 2015.
- [7] S. Cohen, C. Li, J. Yang, and C. Yu. Computational journalism: A call to arms to database researchers. 2011.
- [8] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *KDD*, pages 601–610. ACM, 2014.
- [9] X. L. Dong, L. Berti-Equille, and D. Srivastava. Integrating conflicting data: the role of source dependence. *PVLDB*, 2(1):550–561, 2009.
- [10] A. Halevy, A. Rajaraman, and J. Ordille. Data integration: the teenage years. In *PVLDB*, pages 9–16. VLDB Endowment, 2006.
- [11] W. Hamilton, P. Bajaj, M. Zitnik, D. Jurafsky, and J. Leskovec. Embedding logical queries on knowledge graphs. In *Advances in Neural Information Processing Systems*, pages 2026–2037, 2018.
- [12] N. Hassan, G. Zhang, F. Arslan, J. Caraballo, D. Jimenez, S. Gawsane, S. Hasan, M. Joseph, A. Kulkarni, A. K. Nayak, et al. Claimbuster: The first-ever end-to-end fact-checking system. *PVLDB*, 10(12):1945–1948, 2017.
- [13] Z. Jin, J. Cao, Y. Zhang, and J. Luo. News verification by exploiting conflicting social viewpoints in microblogs. In *AAAI*, 2016.
- [14] J. Kim, B. Tabibian, A. Oh, B. Schölkopf, and M. Gomez-Rodriguez. Leveraging the crowd to detect and reduce the spread of fake news and misinformation. In *WSDM*, pages 324–332. ACM, 2018.
- [15] D. M. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018.
- [16] Y. Li, J. Gao, C. Meng, Q. Li, L. Su, B. Zhao, W. Fan, and J. Han. A survey on truth discovery. *ACM Sigkdd Explorations Newsletter*, 17(2):1–16, 2016.
- [17] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- [18] N. Mele, D. Lazer, M. Baum, N. Grinberg, L. Friedland, K. Joseph, W. Hobbs, and C. Mattsson. Combating fake news: An agenda for research and action, 2017.
- [19] R. Pastor-Satorras and A. Vespignani. Immunization of complex networks. *Physical review E*, 65(3):036104, 2002.
- [20] G. Pennycook and D. Rand. Crowdsourcing judgments of news source quality. *SSRN. com*, 2018.
- [21] K. Papat, S. Mukherjee, J. Strötgen, and G. Weikum. Where the truth lies: Explaining the credibility of emerging claims on the web and social media. In *WWW*, pages 1003–1012, 2017.
- [22] C. Shao, G. L. Ciampaglia, O. Varol, A. Flammini, and F. Menczer. The spread of fake news by social bots. *arXiv preprint arXiv:1707.07592*, pages 96–104, 2017.
- [23] B. Shi and T. Weninger. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowledge-based systems*, 104:123–133, 2016.
- [24] B. Shi and T. Weninger. Proje: Embedding projection for knowledge graph completion. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [25] P. Shiralkar, A. Flammini, F. Menczer, and G. L. Ciampaglia. Finding streams in knowledge graphs to support fact checking. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 859–864. IEEE, 2017.
- [26] K. Shu, H. R. Bernard, and H. Liu. Studying fake news via network analysis: detection and mitigation. In *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*, pages 43–65. Springer, 2019.
- [27] K. Shu, A. Sliva, S. Wang, J. Tang, and H. Liu. Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, 19(1):22–36, 2017.
- [28] E. C. Tandoc Jr, Z. W. Lim, and R. Ling. Defining “fake news” a typology of scholarly definitions. *Digital Journalism*, 6(2):137–153, 2018.
- [29] S. Tschitschek, A. Singla, M. Gomez Rodriguez, A. Merchant, and A. Krause. Fake news detection in social networks via crowd signals. In *WWW*, pages 517–524. WWW, 2018.
- [30] S. Vosoughi, D. Roy, and S. Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018.
- [31] J. Widom. Trio: A system for integrated management of data, accuracy, and lineage. Technical report, Stanford InfoLab, 2004.
- [32] L. Wu and H. Liu. Tracing fake-news footprints: Characterizing social media messages by how they propagate. In *WSDM*, pages 637–645. ACM, 2018.
- [33] X. Zhou and R. Zafarani. Fake news: A survey of research, detection methods, and opportunities. *arXiv preprint arXiv:1812.00315*, 2018.