

Health Misinformation in Search and Social Media

Amira Ghenai

School of Computer Science

University of Waterloo

Canada

aghenai@uwaterloo.ca

ABSTRACT

People regularly use web search and social media to investigate health related issues. This type of Internet data might contain “*misinformation*” i.e incorrect information which contradicts current established medical understanding. If people are influenced by the presented misinformation in these sources, they can make harmful decisions about their health. Our research goal is to investigate the affect of Internet data on people’s health. Our current findings suggest that people can be potentially harmed by search engine results. Furthermore, we successfully built a high precision approach to track misinformation in social media. In this paper, we briefly discuss our ongoing work results. Thereafter, we propose a research plan to understand possible mechanisms of misinformation’s effect on people and possible impacts of these misinformation on public health.

KEYWORDS

Health search, social computing for health, misinformation, rumor, user study

ACM Reference format:

Amira Ghenai. 2017. Health Misinformation in Search and Social Media. In *Proceedings of DH '17, London, United Kingdom, July 2–5, 2017*, 3 pages. <https://doi.org/http://dx.doi.org/10.1145/3079452.3079483>

1 INTRODUCTION

The massive amount of information and its accessibility on the Internet makes it easier for people to seek health related issues using online search. It has been reported that 44% of people had changed their decisions about how to treat an illness after consulting online search results [12]. However, online information is not always correct which suggests that there is a huge threat on people’s health when using online resources. Towards this direction, my PhD research thesis tackles the research question of how web search and social media interactions affect human beliefs and how to protect public health awareness against harmful content. In the following sections, I briefly present a description of my proposed research followed by the research methodology and future plans.

2 DESCRIPTION OF PROPOSED RESEARCH

During my doctoral dissertation, I plan to use mixed-methods approach in order to understand how online incorrect information affects people’s health. Similar to [3], we define *misinformation* as a peace of information spreading in the media confirmed to be false by reliable sources. The proposed research has the following two major research aspects:

- **Misinformation in web search:** Using controlled experiments to measure the positive and negative influence of search results on people’s health-care decisions (whether or not to take a treatment) and the main causes of this influence. Examples of factors that may affect people’s decision making, which we plan to explore, are: prior-knowledge, the trustworthiness of sources, visual display quality and the quality of information.
- **Misinformation in social media:** Using machine learning methods to detect misinformation messages. In this regard, we plan to explore how the volume of misinformation relates to real-world human behavior through social media. We will investigate how misinformation dynamics affect people’s health-care decisions (whether or not to visit hospitals or to immunize) or explore the affect of misinformation on anxiety level and what are the resulting implications (high hospitals visits, low birth rate). Further, we plan to explore online health misinformation susceptible cohorts.

3 CURRENT RESULTS

My proposed research will be built on prior work as well as my own recent conducted work which I summarize as follows:

Building on prior work, where authors [2, 10, 13, 14] designed user studies to understand how trust, source reliability and prior knowledge influence health decision making, my colleagues and I conducted a controlled laboratory study, with 60 participants (all students), where we biased search results towards correct or incorrect information for 10 different medical treatments aiming to determine the extent to which people can be influenced by search engine results. We found that search engine results can significantly influence people both positively and negatively. Results biased towards correct information can lead to increased accuracy up to 70%, while lowering harmful decisions from 20% to 6%. Conversely, results biased towards incorrect information significantly reduces accuracy from 43% to 23%. Furthermore, people tend to be more negatively influenced when exposed to incorrect web pages than when no search engine results exist. The findings of this project suggest that search engine researchers need to build more robust algorithms that take into account the correctness of results in addition to relevance when evaluating pages.

Moving to social media where literature work [6–9] aimed at detecting sickness and understanding rumor spread. In this context,

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

DH '17, July 2–5, 2017, London, United Kingdom

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5249-9/17/07.

<https://doi.org/http://dx.doi.org/10.1145/3079452.3079483>

I was recently working with my colleagues in a project where we aimed at detecting health related rumors and tracking rumors' spread in Twitter. This work is the first of its kind to relate health informatics, machine learning and social network analysis to detect health-related rumors in a social media site. We chose *Zika virus* as a case study in tracking misinformation around health concerns in Twitter. Using Twitter API, we collected more than 13 million tweets regarding the Zika outbreak. We used high-precision approaches such as Latent Dirichlet allocation topic model and keyword search to capture rumors, outlined by the World Health Organization. Later, we identified their geographic distribution and propagation over time. Finally, we used machine learning techniques to build an automatic classifier with a 98% precision to identify rumor-bearing tweets for a specific Zika related questionable topic. Results showed an extremely bursty behavior of rumor-related topics: Tweets of humor/jokes nature and the ones of pedestrian nature maintained a longer presence than those originated by advocacy websites. The findings of this work help public health authorities control the communication during health crisis and act upon rising rumors.

4 PROPOSED RESEARCH METHODOLOGY

Extending the work about the positive and negative influence of search engine results, we are planning to design controlled laboratory studies in order to understand the reasons behind people being influenced by online search results. People might be influenced because of a large number of factors, such as: prior-knowledge, the trustworthiness of sources, visual display aids. For space limitation reasons, we will only explain the experimental setup of one user study that looks into the affect of the prior-knowledge factor. During this specific study, we will ask about the efficacy of two types of medical treatments: treatments that are popular for a specific medical condition, assuming people already have a prior knowledge about them, and rare treatments for medical condition that few people know about. Later, we will ask participants to decide whether the treatment is helpful or not for the specific medical condition after being exposed to correct and incorrect search result pages. Based on the participants' answers, if there is a statistical significance between popular and non-popular treatments, then prior-knowledge has an impact on people's health-care decisions.

Second, building on our prior work that detects and tracks rumors in social media, we plan to understand the effect of health-related rumors in social media on people's behavior i.e. whether people talking about a rumor are more likely to have some specific event/outcome/personal experience. To achieve this, we will first build a medical rumor classifier that identifies who is spreading rumors and we plan to incorporate the following set of features that, to the best of our knowledge, have not yet been used for similar purposes. *Word embedding features*: we will compute these features as we expect that text rumors share certain semantic features. *Medical features*: we will use MetaMap tool in order to identify Metathesaurus concepts referred in a text [1]. *Reliability of resources*: we will identify the truthfulness of sources using *Health On the Net* tool¹ that lists over 7,300 certified trusted medical websites and *BotOrNot* tool that computes the likelihood that a Twitter account is

a bot [5]. *Memorability features*: based on Danescu-Niculescu-Mizil et al. [4]'s work, we plan to measure a text memorability score.

Next, we aim to understand how the detected rumors affect people's health-care decisions. Inspired by the work done by Olteanu et al. [11], we plan to use stratified propensity score matching analysis in order to understand possible personal experiences rumor cohorts might have. We define the treatment group as people who spread specific rumor about a health related topic and the control group as people who do not spread the health rumor but rather general information or clarification about the health topic. Possible outcomes we might be looking at are: word likelihood/distribution explaining personal experiences behavior [7, 8], anxiety level and emotions level changes, possible medical related outcomes: hospital visits, immunization decisions, social network behavior: are people spreading rumors more likely to have larger network than other? Do people spread rumors to be more famous? and the tweeting behavior: rates after talking about rumors: we can look if control people are more important than influential/treatment people in passing on a rumor.

5 CONCLUSION

Through this report, we present two approaches to address misinformation in search and in social media. After measuring the effect of misinformation in search results, we propose future directions to understand the reasons behind misinformation influence on people's health-care decisions. This work helps better understand people's health decision when searching for information online, and also helps researchers improve search engine results for the public health benefit. Next, we aim at understanding how rumors affect real health-care behavior. This work is important to improve health crisis management, where health authorities need to manage efficiently resources during massive waist caused by rumors.

REFERENCES

- [1] Alan R Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program.. In *Proceedings of the AMLA Symposium*. 17.
- [2] Pamela Briggs, Bryan Burford, Antonella De Angeli, and Paula Lynch. 2002. Trust in online advice. *Social science computer review* 20, 3 (2002), 321–332.
- [3] David Coady. 2006. Rumour has it. *International Journal of Applied Philosophy* 20, 1 (2006), 41–53.
- [4] Cristian Danescu-Niculescu-Mizil, Justin Cheng, Jon Kleinberg, and Lillian Lee. 2012. You had me at hello: How phrasing affects memorability. In *ACL '12*.
- [5] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A system to evaluate social bots. In *WWW*.
- [6] Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting Depression via Social Media.. In *ICWSM*. 2.
- [7] Munmun De Choudhury and Emre Kiciman. 2017. The Language of Social Support in Social Media and its Effect on Suicidal Ideation Risk. (2017).
- [8] Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *CHI*. 2098–2110.
- [9] Mark Dredze, David A Broniatowski, and Karen M Hilyard. 2016. Zika vaccine misconceptions: A social media analysis. *Vaccine* 34 (2016), 3441–3442.
- [10] Yvonne Kammerer, Ivar Bråten, Peter Gerjets, and Helge I Strømso. 2013. The role of Internet-specific epistemic beliefs in laypersons's source evaluations and decisions during Web search on a medical issue. *CHB* 29, 3 (2013), 1193–1203.
- [11] Alexandra Olteanu, Onur Varol, and Emre Kiciman. 2017. Distilling the Outcomes of Personal Experiences: A Propensity-scored Analysis of Social Media. *CSCW*.
- [12] Ronald E Rice. 2006. Influences, usage, and outcomes of Internet health information searching: multivariate results from the Pew surveys. *IJMI* 75, 1 (2006).
- [13] Elizabeth Silience, Pam Briggs, Lesley Fishwick, and Peter Harris. 2004. Trust and mistrust of online health sites. *SIGCHI*, 663–670.
- [14] Elizabeth Silience, Pam Briggs, Peter Richard Harris, and Lesley Fishwick. 2007. How do patients evaluate and make use of online health information? *Social science & medicine* 64, 9 (2007), 1853–1862.

¹<https://www.hon.ch/>