

# TrueHealth: Workshop on Combating Health Misinformation for Social Wellbeing @ ICWSM 2023

Gabriella Pasi , Rishabh Upadhyay , and Marco Viviani 

University of Milano-Bicocca, Department of Informatics, Systems, and Communication  
Information and Knowledge Representation, Retrieval, and Reasoning (IKR3) Lab  
Edificio U14 (ABACUS), Viale Sarca, 336 – 20126 Milan, Italy  
{gabriella.pasi, rishabh.upadhyay, marco.viviani}@unimib.it  
<https://ikr3.disco.unimib.it/>

## Abstract

In the last decade, people are increasingly relying on online content as a source of health information, facing the potential harm of encountering false information, conspiracy theories, and unreliable remedies, both as individuals and as a community. Research in this area becomes hence fundamental, given the many challenges and open issues related to the sensitive nature of health information and its use by users. In this context, the TrueHealth 2023 Workshop aims to encourage interdisciplinary research involving computer scientists, physicians, lawyers, and communication experts. The goal is to develop technological solutions that help users access truthful and reliable health information for their social well-being. The workshop seeks to combine expertise from different perspectives to address the problem of health misinformation comprehensively.

## Workshop Overview

In recent years, people have increasingly referred to the Web and social media as sources of information about health-related problems and solutions, as confirmed by the U.S. Pew Research Center (Fox and Duggan 2013), and other European and international studies (Eurostat 2021; Jia, Pang, and Liu 2021). Although, on the one hand, these platforms favor easier and more direct access to information sources by users without the intermediation of experts (Eysenbach 2007), on the other hand, it is precisely such democratization of health information that constitutes a potential danger for people, especially those with a low health literacy (Cline and Haynes 2001). As we have seen especially in the last period, linked to the pandemic, the proliferation of false information, conspiracy theories, and unreliable remedies risk compromising the health not only of individuals, but that of the community as a whole (Southwell et al. 2019; Swire-Thompson, Lazer et al. 2020; Tran, Valecha, and Rao 2022).

## Scope

In the context described above, several literature solutions have been presented in recent years to address the problem of the spread of online health misinformation in different ways. Some treated the problem as a binary classification

task (Cui et al. 2020; Mayank, Sharma, and Sharma 2022; Upadhyay, Pasi, and Viviani 2021, 2023), to distinguish reliable information from health misinformation. Other solutions are approaching it as an Information Retrieval task, to guarantee access to estimated truthful health information without filtering out that identified by the algorithms as misinformation, following the latest recommendations on trustworthy AI (Fröbe et al. 2022; Pradeep et al. 2021; Upadhyay, Pasi, and Viviani 2022; Zhang et al. 2022). With respect to the explainability of results obtained from AI approaches, regardless of which task is used to tackle the problem, some recent solutions are moving in this direction (Kou et al. 2022; Markus, Kors, and Rijnbeek 2021; Upadhyay et al. 2023). However, research in this area remains in its infancy and rather limited, while many problems remain open, given the particularly sensitive nature of health information and how it is used by users.

It is with this in mind that the TrueHealth 2023 Workshop aims to promote the study and development of technological solutions that help users connect with truthful and reliable information for their social well-being. To this end, it is essential to promote research of an interdisciplinary nature, involving computer scientists, physicians, lawyers, and communication experts who can address the problem of health misinformation from different points of view by combining their expertise.

## Topics

The topics of interest of the TrueHealth 2023 Workshop at ICWSM include, but are not limited to:

- Assessing the credibility/truthfulness of Online Health Information (OHI);
- Consumer Health Search (CHS) and information credibility/truthfulness;
- Debunking health misinformation;
- Fake news/rumors and healthcare;
- Measures, evaluation methods, and datasets for health misinformation detection;
- Health misinformation detection;
- Health literacy and information credibility/truthfulness;
- Fact-checking in Online Health Information (OHI);
- Misinformation and public opinion on health;

\*Corresponding author.

- Relationship between access to misinformation and danger to public health;
- Relationship between psychological characteristics and perceptions of health misinformation;
- Techniques for accessing and retrieving credible/truthful Online Health Information (OHI).

## Submissions

The TrueHealth 2023 Workshop received ten submissions, in the form of 2-page Abstracts, Short, and Full Papers. Five submissions were accepted, i.e., one 2-page Abstract, two Short Papers, and two Full Papers, so with an acceptance rate of 50%. Papers have been submitted from five different countries, i.e., Canada, China, Portugal, UK, and USA.

The accepted papers, collected in these Proceedings, have mainly considered COVID-19, touching on aspects such as fact-checking, explainability, conspiracy theories, and toxicity of health-related online content.

## Abstracts

The only 2-page Abstract presented at the Workshop, entitled: “Disinformation and Health: A Systematic Review Study on Health Automatic Fact-checking During the Pandemic”, provides an overview of fact-checking research areas and focuses on a systematic literature review of computational approaches. It also discusses essential elements of the fact-checking pipeline, such as check-worthiness estimation, detecting fact-checked claims, stance detection, and source reliability estimation. The paper addresses the uncertainty surrounding the use of automation in disseminating results, especially in debunking health-related hoaxes. The author proposes a more robust interdisciplinary collaboration between journalists and artificial intelligence developers in specialized domains as a solution to this challenge.

## Short Papers

The first Short Paper presented at TrueHealth 2023, entitled: “Towards Developing a Measure to Assess Contagiousness of Toxic Tweets”, examines the prevalence of toxicity on social media platforms using COVID-19-related datasets. The researchers focus on discussions about face masks, vaccines, 5G technology, and Bill Gates conspiracy theories. They investigate the contagiousness of tweets containing hateful or toxic content related to these topics. The study highlights that while most people believe in the effectiveness of face masks and vaccines against COVID-19, a minority hold opposing views. Additionally, conspiracy theories linking 5G technology and Bill Gates to the virus were found to contribute to worries during the pandemic. The researchers measure the intensity and type of hate speech in each dataset and propose a toxicity contagiousness score to study how toxicity spreads within these datasets. The findings reveal that in the 5G dataset, low toxic content is positively correlated with the contagiousness score. Similarly, highly toxic content in the Bill Gates and pro-mask datasets also shows a positive correlation. However, in datasets related to anti-community sentiments such as anti-mask and anti-vaccine,

highly toxic content has a negative correlation with the contagiousness score. These results provide insights into how different types of content and contexts influence the propagation of toxicity within online communities.

The second Short Paper presented, entitled: “Ecologies of Conspiracy Influencers: A Creator-Centered Lens on Understanding Multiplatform Strategies”, focuses on the role of conspiracy influencers in spreading health misinformation, particularly during the COVID-19 pandemic. The researchers highlight that a small number of these influencers are responsible for a significant amount of misinformation related to public health. While previous research has examined the spread of misinformation, there has been limited investigation into the specific strategies employed by conspiracy influencers across various platforms. To gain a better understanding of these strategies, the researchers conducted a cross-sectional study of 55 influencers, analyzing factors such as their platform usage, audience size, account creation date, and content originality. The results indicate that these influencers utilize multiple platforms to evade algorithmic discrimination and de-platforming. They also tailor their content to monetization channels. Interestingly, despite the popularity of unmoderated platforms, the study reveals that there is still a reliance on moderated platforms to build an audience. The findings of this research can be used to inform strategies aimed at combating the spread of health misinformation within the online ecosystem.

## Full Papers

The first Full Paper presented, entitled: “A Large-Scale Comparative Study of Accurate COVID-19 Information versus Misinformation”, presents a comprehensive comparison study between non-misinformation and misinformation about COVID-19. The authors use a machine learning classifier to identify and extract misinformation. To ensure accuracy, they employ a data enrichment process and post-processing steps. The study reveals clear differences between COVID-19 misinformation and non-misinformation tweets. Non-misinformation tweets primarily focus on prominent actors and community spread, while approximately one-third of misinformation source tweets are related to conspiracy theories. Linguistic analysis supports these findings. The top 10 Bag-Of-Words features associated with misinformation are related to conspiracy theories. Additionally, the LIWC (Linguistic Inquiry and Word Count) analysis indicates that misinformation is frequently associated with negative emotions, anger, and death, whereas non-misinformation is associated with positive emotions, authenticity, and social categories. The paper highlights that addressing misinformation related to conspiracy theories is a primary concern for social media platforms. Twitter, for instance, removes over 40% of such misinformation tweets. In contrast, tweets related to the virus’s origin receive the least attention, with nearly 70% of those tweets remaining accessible. Misinformation tweets have a spreading power that is 158% higher than non-misinformation tweets. Notably, tweets related to conspiracy theories have a longer spreading period w.r.t. other topics, with more than two rounds of spreading even after 32 hours.

The second Full Paper presented, entitled: “A BERT-based Explainable System for COVID-19 Misinformation Identification”, addresses the problem of COVID-19 misinformation and its potential negative impact on public health efforts. In particular, the contribution of this paper is the presentation of an explainable COVID-19 misinformation detection system. The system utilizes an existing BERT-based classification model to classify claims related to COVID-19 as true, false, or partially true. It then provides rationales for its predictions using the LIME XAI (Local Interpretable Model-Agnostic Explanations) tool. This allows machine learning practitioners to understand how the model works, enabling them to debug, fine-tune, and optimize it. Moreover, the system offers explanations for its predictions by extracting relevant sentences from news articles, which serve as the ground truth for the classification process. The sources of the news articles are provided, along with a credibility ranking for each publisher (e.g., high or medium). This information helps end users understand how the classification decision is made, which data sources were utilized, and whether those sources are trustworthy. By providing such information and explanations, the system aims to build trust among its users. The authors claim that their system is the first explainable fact-checking system designed specifically to combat COVID-19 misinformation. They provide examples of the system’s explainability output to demonstrate its effectiveness. Furthermore, they suggest that the proposed explainability framework can be applied to misinformation identification models in other domains beyond COVID-19, such as politics, finance, sports, and entertainment.

### Keynote Speeches

As part of the Workshop, two Keynote Speeches were given. The first, entitled: “Anti-vaccination Echo Chambers”, was given by Yelena Mejova. The second, entitled: “Studying the COVID-19 Infodemic at Scale”, was given by Anatoliy Gruzd. Further details are in the following.

#### Anti-vaccination Echo Chambers

**Abstract:** Despite vaccinations saving uncounted lives, vaccination hesitancy continues to limit the impact of this highly effective intervention. Anti-vaccination views pervade online social media, fueling distrust in scientific expertise and increasing vaccine-hesitant attitudes. Unfortunately, social media discourse has been shown to form “echo chambers” wherein individuals receive the information they are likely to agree with, and surround themselves with like-minded others. In this talk, I will present our efforts to measure and track social media echo chambers around vaccination at a national and international level. We use network analysis to show the clusters of like-minded users, with those sharing anti-vaccination content communicating rarely with those promoting the vaccines, and this separation persists throughout the pandemic. Then, we turn to 28 countries, and show that many display such polarization in their vaccine debates. Moreover, we show that the anti-vaccination clusters tend to be connected across countries, making a kind of international Twitter anti-vaccination echo chamber.




**Yelena Mejova:** she is a Senior Research Scientist at the ISI Foundation in Turin, Italy, working in the area of Data Science for Social Impact and Sustainability. Her research concerns the use of social media in health informatics, especially in lifestyle diseases and mental health, as well as for tracking political speech and other cultural phenomena. Since 2023, she is a co-Editor-in-Chief of EPJ Data Science. Previously as a scientist at the Qatar Computing Research Institute, Yelena was a part of the Social Computing Group working on computational social science, especially as applied to tracking real-life health signals.

URL: <https://www.isi.it/en/people/yelena-mejova>

#### Studying the COVID-19 Infodemic at Scale

**Abstract:** False narratives about COVID-19 have gone global and have been spreading almost as fast as the virus itself. Since January 2020, there have been over 10,000 false and unproven COVID-19-related claims shared via social media and other channels. The talk will discuss how researchers at Toronto Metropolitan University’s Social Media Lab, in partnership with the World Health Organization, spearheaded an international effort to help stem the rise and counter COVID-19 misinformation via the COVID-19 Misinformation Portal.



**Anatoliy Gruzd:** he is a Professor of Information Technology Management and holds the Canada Research Chair in Privacy-Preserving Digital Technologies at the Ted Rogers School of Management, Toronto Metropolitan University. He is also the Director of Research at the Social Media Lab

at the university and a Member of the Royal Society of Canada’s College of New Scholars, Artists, and Scientists. The broad aim of Dr. Gruzd’s various research initiatives is to understand how social media data can be used ethically to tackle a wide variety of societal problems from combating dis/misinformation to helping educators navigate social media for teaching and learning. URL: <https://anatoliygruzd.ca/>

### Workshop Organization

Three Co-organizers and twelve Program Committee members participated in the organization and operation of the Workshop.

#### Co-organizers



**Gabriella Pasi:** she is a Full Professor at the University of Milano-Bicocca / Department of Informatics, Systems, and Communication (DISCo), where she leads the Information and Knowledge Representation, Retrieval, and Reasoning (IKR3) Research Lab. She currently holds the role of Head of Department and she is also Pro-rector for International Rela-

tion

tions. Her main research interests are related to Natural Language Processing (in particular in connection with the tasks of Information Retrieval and Information Filtering), Knowledge Representation and Reasoning, and User Modeling. She is an Associate Editor of several international journals and has participated in the organization of several international events, in both roles of general and program chair. She is a Fellow of the IFSA and she is co-director of the EL-LIS Unit of Milano. She has been both coordinator and PI of several international research projects. She has published more than 250 papers in international journals, books, and in the Proceedings of international conferences.

URL: <https://ikr3.disco.unimib.it/people/gabriella-pasi/>



**Rishabh Upadhyay:** he is a Research Fellow at the University of Milano-Bicocca / Department of Informatics, Systems, and Communication (DISCo). His research interests are related to Machine and Deep Learning, Information Retrieval, and Social Computing. He is currently

working within the EU Horizon 2020 ITN/ETN DoSSIER project on Domain-Specific Systems for Information Extraction and Retrieval, particularly on the project “Assessing Credibility, Value, and Relevance”.

URL: <https://en.unimib.it/rishabh-gyanendra-upadhyay>



**Marco Viviani** is an Associate Professor at the Università degli Studi di Milano-Bicocca / Department of Informatics, Systems, and Communication (DISCo), on topics connected to Social Computing, Trust and Reputation Management, Information Retrieval, Text Mining, and Natural Language Processing.

He is involved in the organization of several research initiatives. He was General Co-chair of MDAI 2019 and organized several Workshops and Special Tracks at International Conferences, including ROM-CIR 2021, 2022, and 2023: Workshops on Reducing Online Misinformation through Credible Information Retrieval. He is Associate Editor of “Social Network Analysis and Mining”, and Guest Editor of several Special Issues in International Journals, among which the “Special Issue on Health Misinformation on Social Media”, in the International Journal of Environmental Research and Public Health. He has published more than 80 research works in international journals, at international conferences, as monographs and book chapters.

URL: <https://ikr3.disco.unimib.it/people/marco-viviani/>

### Program Committee Members

- **Sanda Harabagiu**, The University of Texas at Dallas, USA
- **Liadh Kelly**, Maynooth University, Ireland
- **Dongwon Lee**, Penn State University, USA
- **Yelena Mejova**, ISI Foundation, Italy
- **Marinella Petrocchi**, IIT-CNR, Italy

- **Michael Sirivianos**, Cyprus University of Technology, Cyprus
- **Xingyi Song**, The University of Sheffield, UK
- **Francesca Spezzano**, Boise State University, USA
- **Angelo Spognardi**, Sapienza University of Rome, Italy
- **Hanna Suominen**, The Australian National University, Australia
- **Bei Yu**, Syracuse University, USA
- **Arkaitz Zubiaga**, Queen Mary University of London, UK

### Acknowledgments

We would like to thank the authors of the submitted articles for their interest in the considered problem, the Keynote Speakers for the interest aroused in new research directions, and the members of the Program Committee for their valuable contribution to the success of the TrueHealth 2023 Workshop.

The TrueHealth 2023 Workshop was supported by the EU Horizon 2020 ITN/ETN on Domain Specific Systems for Information Extraction and Retrieval (DoSSIER), H2020-EU.1.3.1., ID: 860721, <https://dossier-project.eu/>

### References

- Cline, R. J.; and Haynes, K. M. 2001. Consumer health information seeking on the Internet: the state of the art. *Health education research*, 16(6): 671–692.
- Cui, L.; Seo, H.; Tabar, M.; Ma, F.; Wang, S.; and Lee, D. 2020. Deterrent: Knowledge guided graph attention network for detecting healthcare misinformation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 492–502.
- Eurostat. 2021. One in two EU citizens look for health information online. <https://ec.europa.eu/eurostat/web/products-eurostat-news/-/edn-20210406-1>. Accessed on June 15, 2023.
- Eysenbach, G. 2007. From intermediation to disintermediation and apomediation: new models for consumers to access and assess the credibility of health information in the age of Web2. 0. *Studies in health technology and informatics*, 129(1): 162.
- Fox, S.; and Duggan, M. 2013. Health Online 2013. <https://www.pewresearch.org/internet/2013/01/15/health-online-2013/>. Accessed on June 15, 2023.
- Fröbe, M.; Günther, S.; Bondarenko, A.; Huck, J.; and Hagen, M. 2022. Using Keyqueries to Reduce Misinformation in Health-Related Search Results. In *Proceedings of the 2nd Workshop Reducing Online Misinformation through Credible Information Retrieval*, 1–10.
- Jia, X.; Pang, Y.; and Liu, L. S. 2021. Online health information seeking behavior: a systematic review. In *Healthcare*, volume 9, 1740. MDPI.
- Kou, Z.; Shang, L.; Zhang, Y.; and Wang, D. 2022. Hc-covid: A hierarchical crowdsourced knowledge graph approach to explainable covid-19 misinformation detection.

*Proceedings of the ACM on Human-Computer Interaction*, 6(GROUP): 1–25.

Markus, A. F.; Kors, J. A.; and Rijnbeek, P. R. 2021. The role of explainability in creating trustworthy artificial intelligence for health care: a comprehensive survey of the terminology, design choices, and evaluation strategies. *Journal of Biomedical Informatics*, 113: 103655.

Mayank, M.; Sharma, S.; and Sharma, R. 2022. DEAP-FAKED: Knowledge graph based approach for fake news detection. In *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 47–51. IEEE.

Pradeep, R.; Ma, X.; Nogueira, R.; and Lin, J. 2021. Vera: Prediction techniques for reducing harmful misinformation in consumer health search. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2066–2070.

Southwell, B. G.; Niederdeppe, J.; Cappella, J. N.; Gaysynsky, A.; Kelley, D. E.; Oh, A.; Peterson, E. B.; and Chou, W.-Y. S. 2019. Misinformation as a misunderstood challenge to public health. *American journal of preventive medicine*, 57(2): 282–285.

Swire-Thompson, B.; Lazer, D.; et al. 2020. Public health and online misinformation: challenges and recommendations. *Annu Rev Public Health*, 41(1): 433–451.

Tran, T.; Valecha, R.; and Rao, H. R. 2022. Health-related Misinformation Harm during the COVID-19 Pandemic: An Investigation of Non-comparative and Comparative Harm Perceptions. *AIS Transactions on Human-Computer Interaction*, 14(2): 185–206.

Upadhyay, R.; Knoth, P.; Pasi, G.; and Viviani, M. 2023. Explainable Online Health Information Truthfulness in Consumer Health Search. *Frontiers in Artificial Intelligence*, 6.

Upadhyay, R.; Pasi, G.; and Viviani, M. 2021. Health Misinformation Detection in Web Content: A Structural-, Content-Based, and Context-Aware Approach Based on Web2Vec. In *Proceedings of the Conference on Information Technology for Social Good, GoodIT '21*, 19–24. New York, NY, USA: Association for Computing Machinery. ISBN 9781450384780.

Upadhyay, R.; Pasi, G.; and Viviani, M. 2022. An Unsupervised Approach to Genuine Health Information Retrieval Based on Scientific Evidence. In *Web Information Systems Engineering–WISE 2022: 23rd International Conference, Biarritz, France, November 1–3, 2022, Proceedings*, 119–135. Springer.

Upadhyay, R.; Pasi, G.; and Viviani, M. 2023. Vec4Cred: a model for health misinformation detection in web pages. *Multimedia Tools and Applications*, 82(4): 5271–5290.

Zhang, D.; Vakili Tahami, A.; Abualsaud, M.; and Smucker, M. D. 2022. Learning trustworthy web sources to derive correct answers and reduce health misinformation in search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2099–2104.